NAKIVO Backup & Replication v11.1

# User Guide for VMware

# Table of Contents

**NAKIVO Backup & Replication Overview** ..... **5**

Data Protection ..... 6

Data Recovery ..... 27

Disaster Recovery ..... 36

Reliability ..... 46

Performance ..... 67

Administration ..... 82

Automation ..... 88

Integration ..... 92

BaaS ..... 99

NAKIVO Licensing Policy ..... 106

**Deployment** ..... **119**

Architecture ..... 120

System Requirements ..... 130

Deployment Scenarios ..... 215

Installing NAKIVO Backup & Replication ..... 231

Updating NAKIVO Backup & Replication ..... 296

Uninstalling NAKIVO Backup & Replication ..... 329

**Getting Started** ..... **333**

Logging in to NAKIVO Backup & Replication ..... 334

First Steps with NAKIVO Backup & Replication ..... 341

Web Interface Components ..... 344

Managing Jobs and Activities ..... 354

**Settings** ..... **405**

General ..... 406

Inventory ..... 510

Nodes ..... 564

Backup Repositories ..... 610

Federated Repositories ..... 696

Tape ..... 723

Expert Mode ..... 762

Maintenance Mode ..... 786

Virtual Appliance Configuration ..... 789

Multi-Tenant Mode Configuration ..... 814

Support Bundles ..... 829

Built-in Support Chat ..... 831

**Backup** ..... **839**

Creating VMware Backup Jobs ..... 839

Creating Backup Copy Jobs ..... 872

Backing Up to Tape ..... 901

Staging (Seeding) Initial Backup ..... 923

Deleting Backups ..... 924

**Replication** ..... **933**

Creating VMware Replication Jobs ..... 934

Creating VMware Real-Time Replication Jobs ..... 968

Staging (Seeding) VM Replication ..... 989

**Recovery** ..... **990**

Granular Recovery ..... 991
Full Recovery ..... 1060
Planning Disaster Recovery ..... 1106

**Performing Flash Boot Recovery ..... 1185**

Creating VMware Flash Boot Jobs ..... 1186
Migrating Recovered VMs Using Flash Boot ..... 1204

**Monitoring ..... 1206**

Metrics ..... 1206
Alerts ..... 1211
Reports ..... 1220

**Integration and Automation ..... 1233**

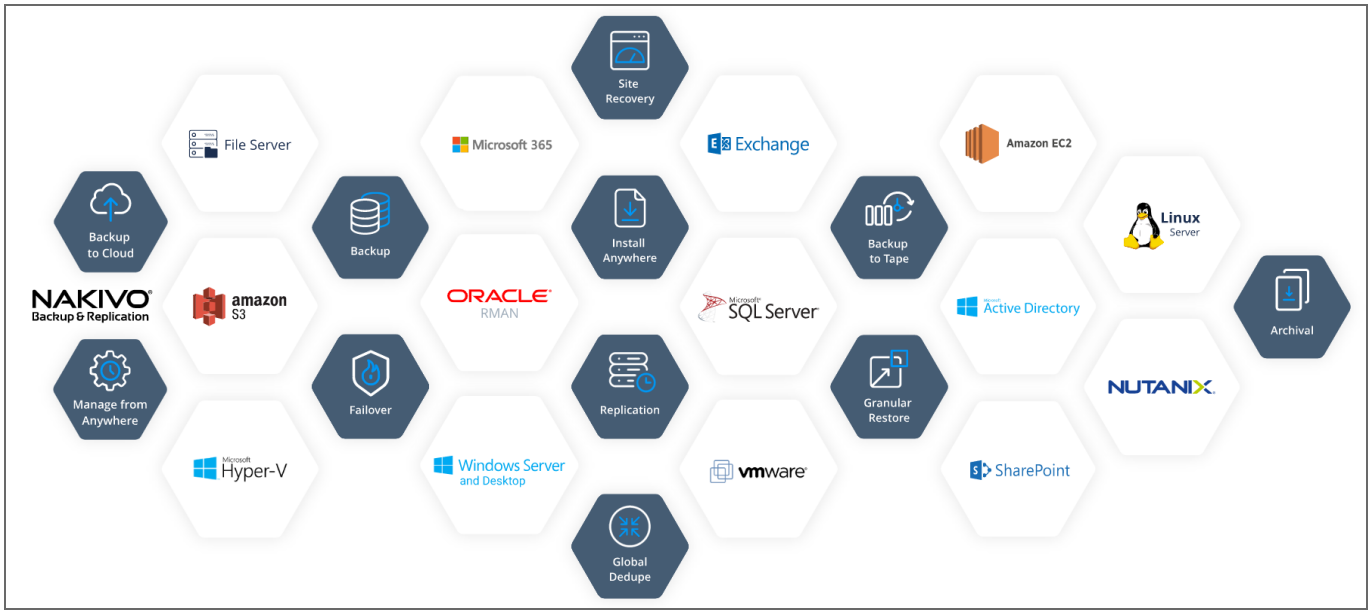
Command Line Interface ..... 1234
Automation with HTTP API ..... 1248
Aptare IT Analytics Integration ..... 1249

**Multi-Tenant Mode ..... 1254**

Creating a Local Tenant ..... 1254
Creating a Remote Tenant ..... 1261
Using the MSP Console ..... 1265
Using the MSP Dashboard ..... 1278
Granting Self-Service Access ..... 1289

# NAKIVO Backup & Replication Overview

NAKIVO Backup & Replication offers backup, replication, failover, backup to cloud, backup to tape, backup copy, backup data reduction, instant verification, granular restore and disaster recovery orchestration for virtual, physical, cloud and SaaS environments - all in one convenient web interface.




The product provides image-based, application-aware, incremental backup and replication. You can easily schedule jobs using the calendar in the product's web interface and save up to 1,000 recovery points for each backup, rotating them on a GFS basis. You can also protect your VMs and instances more efficiently by taking advantage of Changed Block Tracking (for VMware), Resilient Change Tracking (for Hyper-V), or Changed Regions Tracking (for Nutanix), LAN-Free Data Transfer, Network Acceleration, and other product features.

The solution includes an advanced disaster recovery (DR) functionality. It allows you to automate and orchestrate DR activities across multiple sites. Build advanced site recovery workflows to failover an entire site in just a few clicks, perform non-disruptive recoverability testing, and make sure you have a workable DR plan in place to help minimize downtime and prevent loss of revenue or data.

NAKIVO Backup & Replication allows you to simplify data protection management through the automation of core tasks such as backup, replication, and backup copy. Instead of tracking every change in your environment and manually adding VMs or physical machines to jobs, you can set up policies based on a VM/physical machine name, tag, size, location, power state, configuration, or other parameters. NAKIVO Backup & Replication can regularly scan your infrastructure and automatically protect VMs, physical machines, and Amazon EC2 instances that match policy rules.

With NAKIVO Backup & Replication, you can also ensure the safety and integrity of your Microsoft Office 365 data. The product allows you to reliably protect Microsoft Exchange mailboxes, OneDrives for Business, and SharePoint Online sites.

# Data Protection

Data protection is the process of safeguarding business-critical information from loss, corruption or compromise. NAKIVO Backup & Replication offers a complete suite of backup features to protect physical, virtual, and cloud environments. By providing you with great flexibility and multiple automation options, the product can save you time and resources. For more information about the data protection offered by NAKIVO Backup & Replication, refer to the following topics:


- "Amazon EC2 Backup" on the next page
- "Amazon EC2 Concepts" on the next page
- "Backup Copy" on page 10
- "Backup to Cloud" on page 15
- "Backup to Tape" on page 17
- "Backup Encryption" on page 19
- "Container Protection" on page 23
- "Virtual Machine Backup" on page 26


## Agent-Based VM Backup and Recovery

You can use the agent-based backup approach in NAKIVO Backup & Replication to discover and back up virtual machines. For the list of supported hypervisors, refer to "Feature Requirements" on page 169.

The following recovery scenarios are supported for VMs protected using the agent-based backup approach:

- File-level recovery
- Object-level recovery
- Universal object recovery

**Notes**

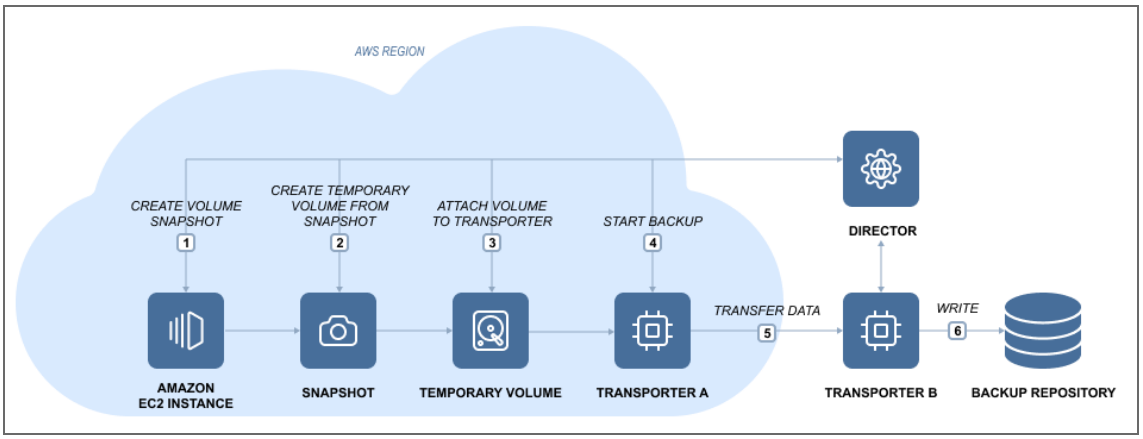
- Immutability is supported.
- To resolve the possible controller incompatibility when recovering to the guest hypervisor (for example, you cannot boot to ISO of bare-metal recovery), consider changing the controller to VMware SCSI.

For system requirements, refer to the "System Requirements" on page 130 page.

For a full list of supported backup/backup copy jobs destinations, see the corresponding section in "Feature Requirements" on page 169.

# Amazon EC2 Backup

NAKIVO Backup & Replication allows you to create native backups of Amazon EC2 Instances. An Amazon EC2 instance backup is a point in time copy of an entire instance that is stored in a special folder called a Backup Repository.




Here is how NAKIVO Backup & Replication performs Amazon EC2 instance backup:

1. Takes snapshots of the EBS volumes attached to the Amazon EC2 instance.
2. Converts snapshots to temporary volumes and attaches them to the Transporter instance.
3. Reads data from the temporary volumes and sends it to the backup repository.
4. Detaches and removes the volumes.

To back up VMware VMs to Amazon EC2 you need to do the following:

1. Add an Amazon EC2 Account to the product's Inventory.
2. Deploy a Transporter to the Amazon EC2 Region where you wish to create a Backup Repository.
3. Create a Backup Repository in the Amazon EC2 Region.

# Amazon EC2 Concepts

- Instance
- EBS Volume
- Region
- Availability Zone
- VPC
- Subnet
- Security Group
- Key Pair
- Elastic Network Adapter

## Instance

An *Amazon EC2 Instance* is a virtual server in Amazon's Elastic Compute Cloud (EC2). Amazon EC2 provides different Instance types so you can choose the CPU, memory, storage, and networking capacity you need.

## EBS Volume

An *Amazon EBS Volume* is a virtual disk that can be attached to any Amazon EC2 Instance that is in the same Availability Zone. Amazon EBS volumes persist independently from the life of the instance, i.e. deleting an Amazon EC2 Instance does not delete EBS Volumes that were connected to it.

## Region

An *Amazon EC2 Region* is a geographic area where an Amazon EC2 Instance is hosted. Amazon EC2 provides multiple Regions so you can create and run your Amazon EC2 Instances in locations that meet your requirements. Each Region is completely independent and isolated from others.

## Availability Zone

An *Amazon EC2 Availability Zone* is a location within an Amazon EC2 Region. Each Availability Zone is isolated from failures in other Availability Zones, yet all Availability Zones within the same region are connected with low-latency network connectivity to others in the same Region.

## VPC

A *virtual private cloud* (VPC) is a virtual network in Amazon EC2. A VPC is dedicated to your AWS Account and is logically isolated from other virtual networks in the AWS cloud. Similar to regular networks, you can configure your VPCs: select IP address ranges, create subnets, configure route tables, network gateways, and security settings. After you have created and configured a VPC, you can connect your Amazon EC2 Instances to the VPC.

## Subnet

A *subnet* is a range of IP addresses in a VPC. You can connect Amazon EC2 Instances to a subnet that you select: public subnets provide access to the Internet, while private subnets don't.

## Security Group

A *security group* is a virtual firewall that controls the traffic for one or more instances. When you create an Amazon EC2 Instance, you associate one or more security groups with the Instance. You add rules to each security group that allows traffic to or from its associated instances. You can modify the rules for a security group at any time; the new rules are automatically applied to all instances that are associated with the security group. When we decide whether to allow traffic to reach an instance, we evaluate all the rules from all the security groups that are associated with the instance.

## Key Pair

Amazon EC2 uses *key pairs* to encrypt and decrypt login information. A key pair consists of a Public Key that is used to encrypt passwords, and a Private Key is used to decrypt them. When creating a new Amazon EC2 Instance, you need to either create a new Key Pair for it or assign an existing key pair for the Instance. To log in to your Amazon EC2 Instance, you must provide the private key for it. Note that Linux instances have no password, and you use a key pair to log in using SSH. With Windows instances, you use a key pair to obtain the administrator password and then log in using RDP.

## Elastic Network Adapter

*Elastic Network Adapter* (ENA) is a custom network interface with accompanying drivers providing Enhanced Networking on EC2 instances. ENA is optimized to deliver high throughput and packet per second performance and consistently low latencies on EC2 instances. Depending on the type of EC2 instance, you can utilize up to 20 Gbit/s of network bandwidth with ENA. For more information, refer to the corresponding article on the AWS website.

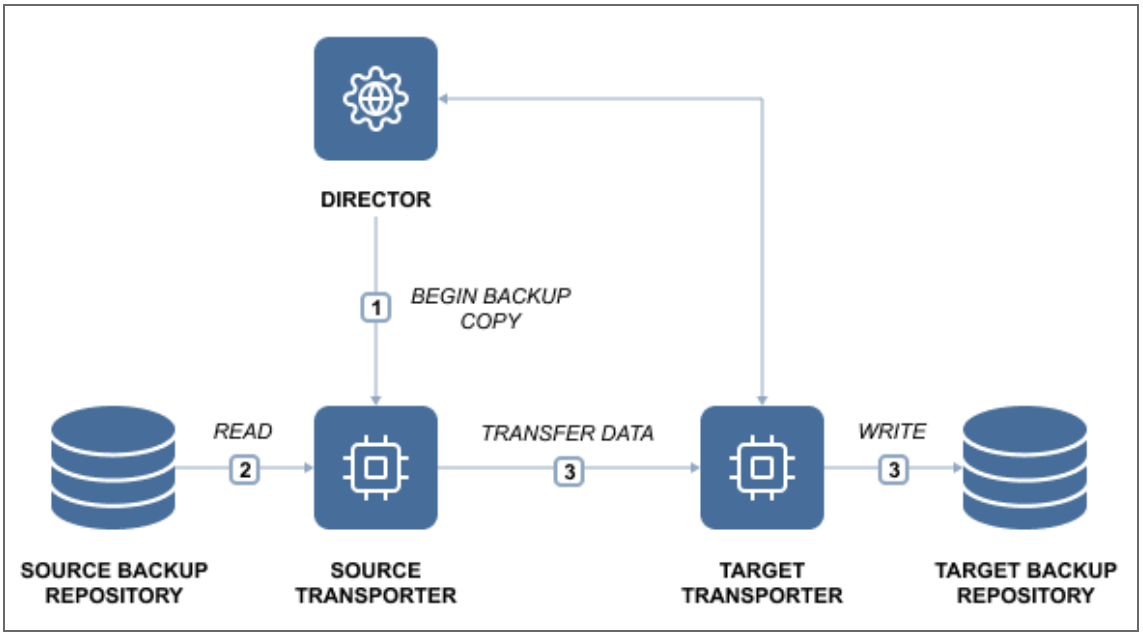
# Backup Copy

Backups can be lost on account of a number of reasons, so having more than one copy of your business-critical backups is vital for ensuring that your data can be recovered in case of disaster. Backup Copy jobs provide a simple yet powerful way to create and maintain copies of your backups. Backup copy jobs copy backups from one Backup Repository to another without affecting the source ESXi hosts, VMs, or Amazon EC2 instances. This way, your source VMs or Amazon EC2 instances are read-only once while backups can be copied to one or multiple locations.

- Create Mirrored Copy of your Backup Repository
- Copy Most Important Backups
- Copy Backups Created by Particular Backup Jobs
- Resource Efficiency and Variable Data Compression
- Copy Backups Offsite
- Copy Backups to Amazon Cloud
- Copy Recovery Points that You Need
- Schedule Backup Copy to Suit Your Needs

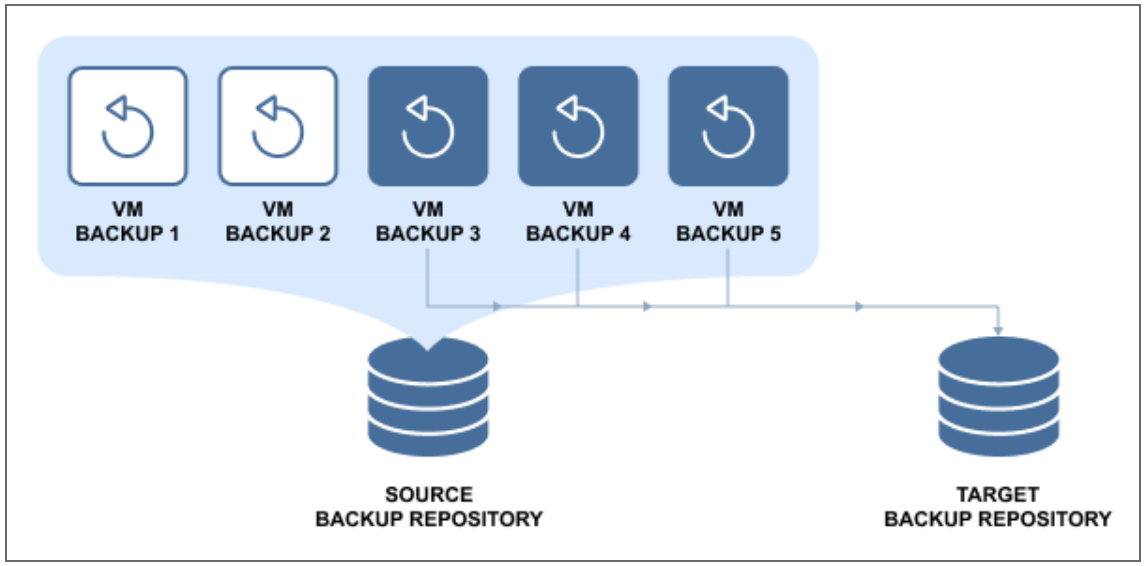
## Create Mirrored Copy of your Backup Repository

With a Backup Copy job, you can create and maintain a mirrored copy of your primary Backup Repository, which is the simplest and the most reliable way to protect all your backups. Think of it as a Backup Repository replication: all backups and recovery points that appear in the Backup Repository A will be automatically sent to Backup Repository B:



## Copy Most Important Backups

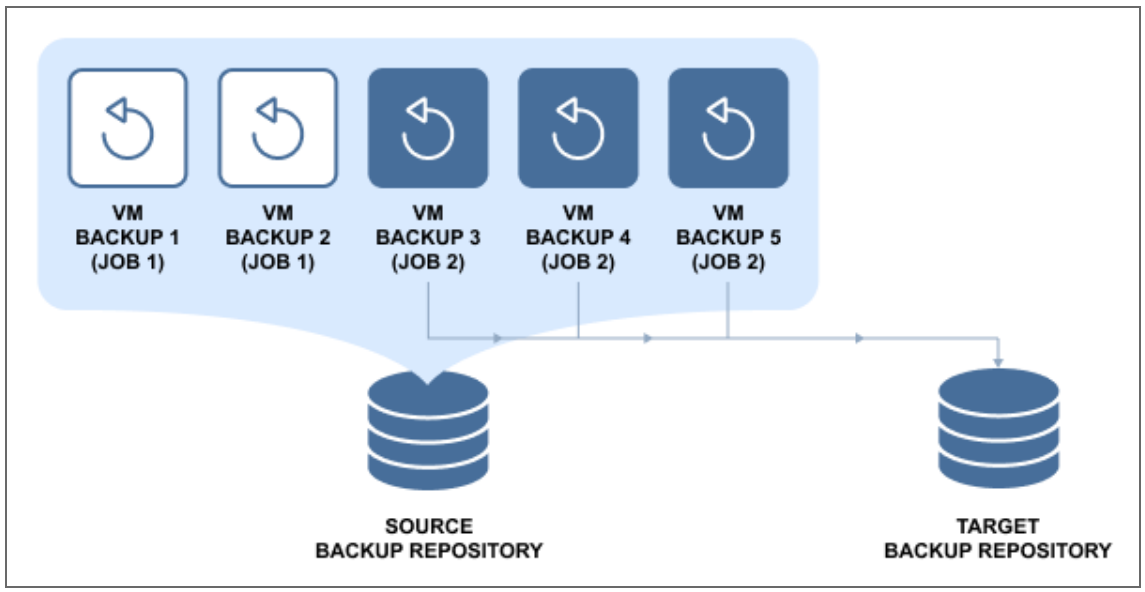
To save storage space on your secondary Backup Repository and to speed up data transfer, you can choose to create a Backup Copy job for only the most important backups:




This way, only the selected backups (and their recovery points) will be transferred to the secondary Backup Repository.

## Copy Backups Created by Particular Backup Jobs

NAKIVO Backup & Replication enables you to create and maintain copies of backups created by particular Backup jobs:




This way, you can ensure that all backups created by important Backup jobs are copied to a secondary Backup Repository.

## Resource Efficiency and Variable Data Compression

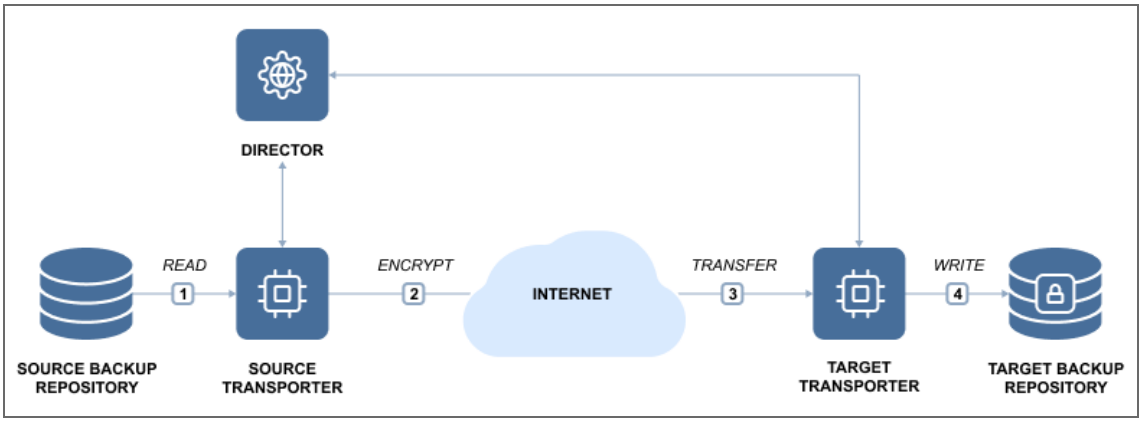
In addition to global data deduplication, NAKIVO Backup & Replication automatically compresses backed up data to reduce the amount of space that backups occupy in storage. By default, the compression level in the new Backup Repositories is set to “Fast,” so that your Backup jobs will run faster. When creating a secondary Backup Repository, you can set the compression level to “Best,” which uses more CPU, but delivers better compression levels. This way, the strongest compression algorithm will be used to compress backup data, resulting in smaller backups in your secondary Backup Repository.

> **Note**
> Compression is always performed before encryption. This ensures that backup data is efficiently compressed, even when encryption is enabled.

Similarly, if source and target Repositories already share the same type and compression, NAKIVO Backup & Replication automatically skips data pack and unpack stages during Backup Copy jobs to cut down on time and resource usage.

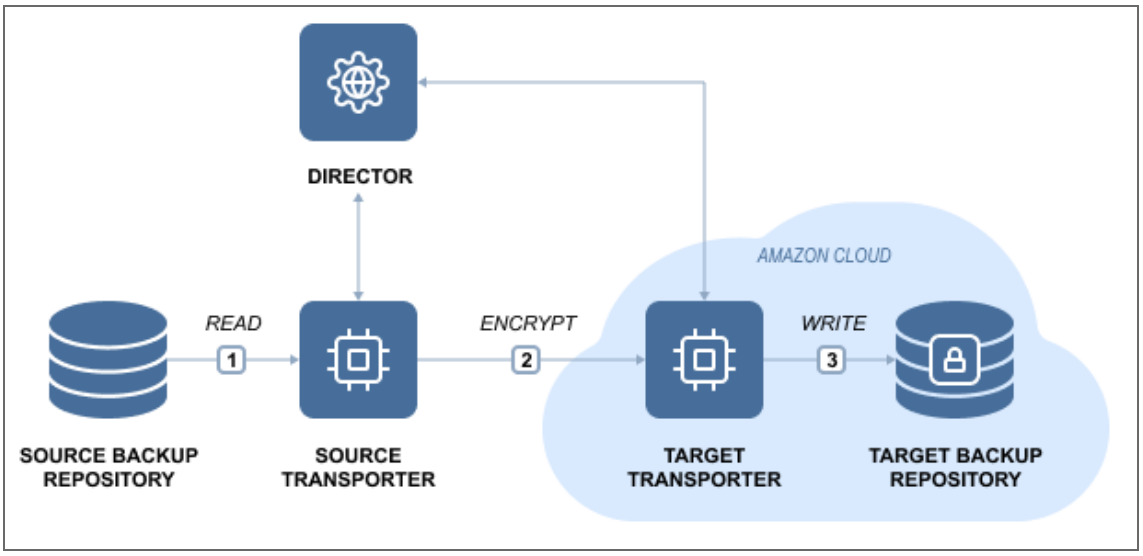
## Copy Backups Offsite

While you can keep copies of your backups locally, having at least one copy of your most critical backups offsite can save you a lot of trouble in case a local disaster should wipe your primary backups. The secondary Backup Repository can be placed in any location that has a connection to the Internet, since backup data can be transferred via AES 256 encrypted link, and your secondary backup repository can be encrypted as well.



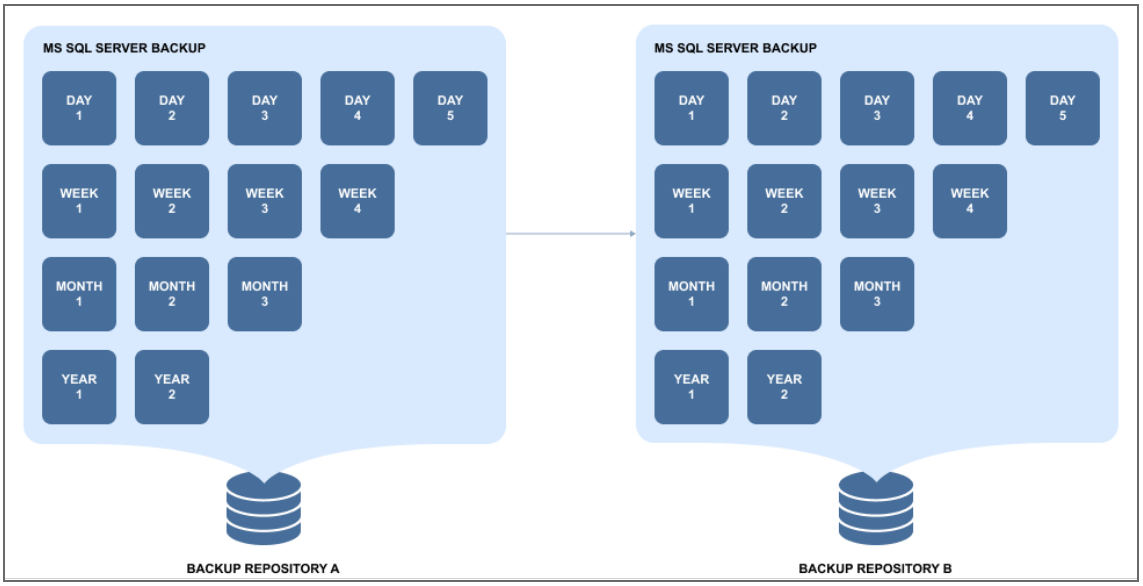

## Copy Backups to Amazon Cloud

Amazon provides one of the most reliable and affordable cloud services in the industry. With NAKIVO Backup & Replication, you can use Amazon's fast, reliable, and affordable cloud to store copies of your backups.

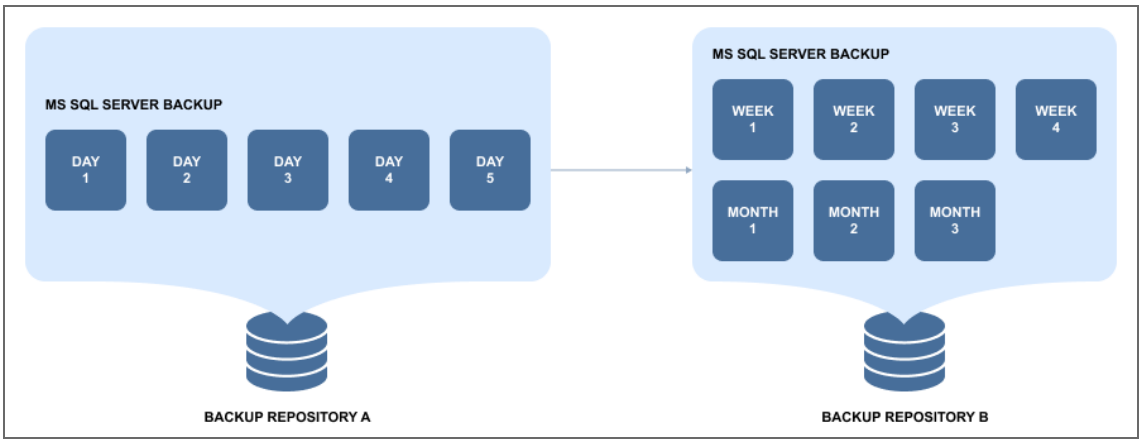



## Copy Recovery Points that You Need

Each backup can contain up to 10,000 recovery points, which are saved based on recovery point retention policy, i.e. how many recovery points you want to have and for how long you want to keep them. With Backup Copy jobs, you can choose to create a mirrored copy of each backup: all recovery points that are available in Backup Repository A will be copied to Backup Repository B.




However, Backup and Backup Copy are different jobs, so you can set different retention policies for your primary backups and their copies in a different Backup Repository. This way, for example, you can store several daily backups onsite, and keep (archive) weekly, monthly, and yearly copies of backups in a secondary Backup Repository for long-term storage.




Also, you can use fast storage for a subset of backups and use slower, but more reliable storage for long-term archiving.

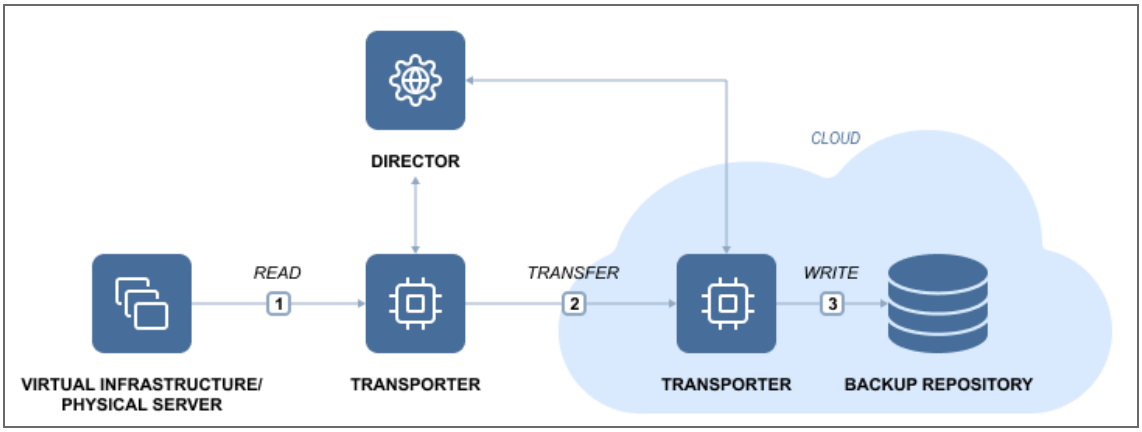
## Schedule Backup Copy to Suit Your Needs

Backup Copy jobs have their own schedule, so you can set them up to run whenever it suits your needs. For example, you can set up a Backup Copy job to run every night on workdays, or set it up to run on weekends to send all backups made during the week to a secondary Backup Repository.

To learn how to create and run backup copy jobs with NAKIVO Backup & Replication, refer to “Creating Backup Copy Jobs” on page 872.

# Backup to Cloud

NAKIVO Backup & Replication allows you to send backups and backup copies to Amazon EC2, Amazon S3, generic S3-compatible storage, Microsoft Azure, Backblaze B2, and Wasabi Hot Cloud Storage..




Keeping backups in the cloud provides a number of benefits, including:

- Safe backup storage. Storing backups in the cloud keeps them safe even if the local infrastructure becomes unavailable.
- Backup immutability. Backups stored in the cloud can also be made immutable to protect them from new ransomware attacks or accidental deletions.
- Flexible backup storage. Cloud storage can be easily expanded as required, eliminating the need to choose, order, install, and configure new servers or hard drives for your growing environment.
- Easy and quick data recovery. Backups can be accessed at any time and from anywhere.
- Affordable backup storage. Instead of buying, configuring and maintaining an offsite backup infrastructure, you can simply use your existing hardware.
- Simple backup management. The set it and forget it approach in NAKIVO Backup & Replication allows you to schedule regular backup jobs to the cloud.

While cloud providers offer cloud storage at an affordable price, NAKIVO Backup & Replication helps further reduce offsite backup costs with additional features like incremental backup, exclusion of swap files and partitions, and backup compression, among others. With NAKIVO Backup & Replication, you can use Amazon EC2, Amazon S3, generic S3-compatible storage, Microsoft Azure, Backblaze B2, or Wasabi as your primary or secondary backup storage destination.

# How Backup to Cloud Works

A backup represents a point-in-time copy of a VM or physical machine that is stored in the Backup Repository. A Backup Repository is the destination for storing backup data. NAKIVO Backup & Replication allows you to create backup repositories in public clouds such as Amazon EC2, Amazon S3, generic S3-compatible storage, Microsoft Azure, Backblaze B2, or Wasabi. In NAKIVO Backup & Replication, a backup job is performed as follows:

1. The product automatically creates temporary snapshots of the source VMs/physical machines.
2. The data blocks that were changed since the last backup are identified and sent to the Backup Repository.
3. The temporary snapshots created in the process are removed.

However, backups can also get lost or damaged as a result of unexpected events. With NAKIVO Backup & Replication, you can also run backup copy jobs, which allow you to create copies of VMware vSphere VM, Microsoft Hyper-V, Amazon EC2, or physical machine backups. Creating copies of critical backups provides an additional level of data protection to avoid a single point of failure.

NAKIVO Backup & Replication enables you to copy backups from one Backup Repository to another without using the source hosts/VMs or physical machines. For more information, refer to “Backup Copy” on page 10.

NAKIVO Backup & Replication also includes an automated backup verification feature, which reads backups at the block level, compares the data written to the Backup Repository with the data from the source machine, then checks whether the data on both sites is identical and can be recovered in case of disaster.

For more details on backing up to cloud, refer to the following topics:

- “Backup Repository in Amazon EC2” on page 630
- “Backup Repository in Amazon S3” on page 636
- “Backup Repository in Generic S3-Compatible Object Storage ” on page 640
- “Backup Repository in Microsoft Azure Blob Storage” on page 644
- “Backup Repository in Backblaze B2 Cloud Storage” on page 648
- “Backup Repository in Wasabi Hot Cloud Storage” on page 653

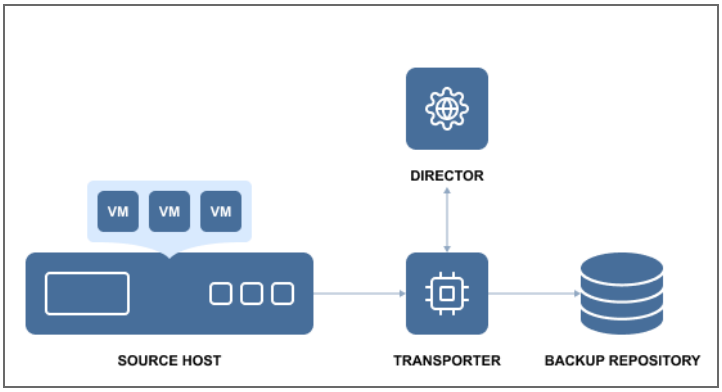
# Backup to Tape

NAKIVO Backup & Replication provides native tape support for automated tape libraries, including virtual tape libraries (VTL), as well as standalone tape drives.

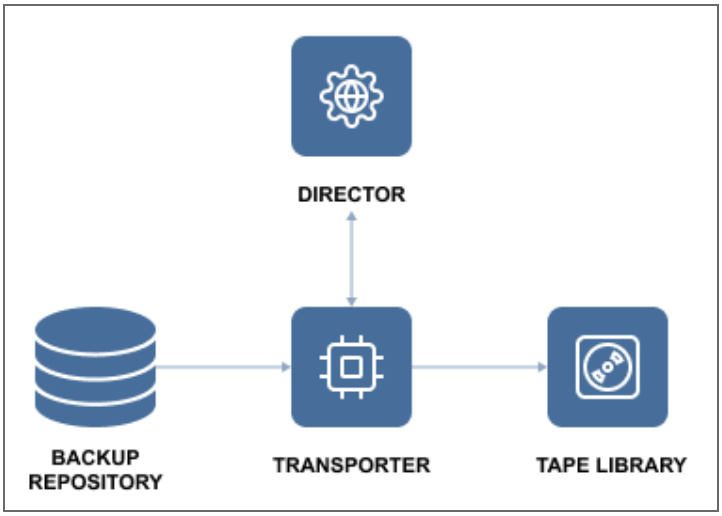
Backup to Tape is the process of backing up critical data to a tape cartridge. In essence, backing up to tape means creating a backup, storing it in the repository and then moving it to a tape cartridge for safekeeping. NAKIVO Backup & Replication supports backups of the following platforms: VMware, Hyper-V, Nutanix AHV, Amazon Amazon EC2, and physical machines. The backups can be sent to physical tape libraries or VTL for storing. NAKIVO Backup & Replication allows for realizing the Disk Staging (D2D2T) backup strategy, where disks are used as an additional, temporary stage of the backup process before finally storing backup to tape.

In NAKIVO Backup & Replication, the process of storing backups to tape consists of two stages:

**Stage 1** – creating backups and storing them in the Backup Repository:




**Stage 2** – copying backups from the repository to the tape library:




Recovering from tape is the reverse of backing up: the backups stored on the tape cartridges are first recovered to the Backup Repository and then recovered using NAKIVO Backup & Replication's tools.

Before you back up/recover to/from tape (physical or VTL), you need to configure NAKIVO Backup & Replication by adding tape libraries, discovering cartridges, etc.

The Native Tape support is fully integrated into NAKIVO Backup & Replication solution and allows you to administer all backup and restore operations on tapes directly from the application's user interface. Saving data on tapes presents you with the same data managing options as disk repositories: you can store full and incremental backups, apply user-defined retention settings to the archived data, select restore points and so on.

NAKIVO Backup & Replication supports Linear Tape-Open tape libraries and standalone tape drives starting from generation 3 (LTO3) or later as well as VTL. Using the solution, you can discover not only tape libraries and standalone devices, but also the tape cartridges in those devices.

> **Note**
> All the tape cartridges discovered within a Robotic Tape Library should have barcodes for the best performance of the product. For standalone tape devices, this is not essential.

Also, any changes to the tape infrastructure (moving or removing cartridges, changing their order, etc.) made by any other means (i.e. manually or via command line) rather than with NAKIVO Backup & Replication is the user's responsibility, since the system is unaware of such changes.

NAKIVO Backup & Replication supports writing/reading backups to/from discovered tape cartridges, as well as other operations, like moving cartridges between slots, erasing, scanning, etc.

The table below provides a description of some of the tape-related terms:

| Term | Description |
|---|---|
| **Tape Library** | A storage device that includes one or more tape drives, a number of slots and a media changer (robot). |
| **Tape Drive** | A device component (or a standalone device) used to read and write the tape cartridge. |
| **Slot** | A place in the tape library designed to hold a single cartridge. |
| **Mail Slot** | A slot in the tape library that allows you to physically add or remove a tape cartridge without disturbing the operation of the tape library. |
| **Media Changer** | A device component used to move a single tape cartridge between slots and load/unload the cartridge to/from the tape drive. |
| **Tape Cartridge (Tape)** | A unit of sequential magnetic medium and an optional barcode used for identification. |
| **Media Pool** | A logical container that contains tape cartridges. |

| Term | Description |
|---|---|
| **Backup (Tape)** | A logical entity containing one or more recovery points on one or more tape cartridge(s) that belong to a single source object. |
| **Recovery Point (Tape)** | A complete or incomplete data set required to rebuild a VM or instance as of a particular moment in time. |

# Backup Encryption

With NAKIVO Backup & Replication you can configure encryption to protect backup data against breaches, perform recovery from encrypted backups, and manage passwords.

## How Backup Encryption Works

To encrypt backup data, do the following:

1. Enable **Backup Encryption** in the **Options** step of the corresponding wizard. In case of a backup copy job, you can select the encryption on source and target.
2. Set up and confirm a password.
3. Optionally, in the **Settings > General > System Settings > Encryption** tab, enable the **(AWS) Key Management Service**. For more information, refer to Enabling KMS.
4. Run the job.

The product automatically generates the password hash based on the user password.

The cryptographic salt used for hash creation is saved in the recovery point metadata.

The password hash is used to generate a single-use encryption key to encrypt the backup and File System Indexing (FSI) data. For this, the **FSI** option must be enabled in the job.

> **Note**
> Data compression is performed before encryption to maintain its efficiency. Encryption does not interfere with or degrade the results of compression.

For more information about browsing through encrypted FSI data when the password hash is available in the product database, refer to Searching Indexed Files.

You can find more details on how to enable backup encryption for the supported jobs in the following articles:

- “Backup Job Wizard for VMware: Options” on page 857
- “Backup Copy Job Wizard: Options” on page 891

For more details on how to enable the **Backup Encryption** feature, refer to Enabling Backup Encryption.

## Trust Zones

Trust Zones refer to secure areas or system components where encrypted data are processed, transmitted, or stored. These zones ensure the integrity and confidentiality of data by facilitating encryption and decryption processes. Trust zones create a secure framework for handling sensitive data during backup, recovery, or the backup copying operations.

For **Backup jobs**, a trust zone includes the source transporter where data is read from the backup source and prepared for secure transmission.

For **Recovery Jobs**, a trust zone includes the target transporter responsible for securely writing the restored data to its target location.

For **Backup Copy Jobs**, you can define the trust zones as the source or target transporter, depending on the configuration.

- Select **Options** > **Backup Encryption** > **Enable on source**, to encrypt the backup data with the source transporter. In this case, the source transporter works in the trusted zone and the encrypted backup can be sent to the target transporter in the untrusted zone.
- Select **Options** > **Backup Encryption** > **Enable on target**, to encrypt the backup data with the target transporter. In this case, the source transporter works in the untrusted zone and the backup encryption can be done with the target transporter in the trusted zone.

> **Note**
> If you selected encryption on source or target, the **Network acceleration** option can not be used.

**Reverse Connection**

Instead of the source transporter directly initiating the connection, the target transporter in the untrusted zone establishes a connection back to the source using a proxy host. This approach mitigates security risks by allowing the system to operate within strict firewall policies, ensuring that only validated and encrypted requests are processed.

## Restoring from Encrypted Backups

NAKIVO Backup & Replication keeps all encryption/decryption details safe so that you do not need to enter your password every time you need to restore data from encrypted backups. With the **Backup Encryption** feature, you get your backed-up data decrypted and restored even if you forget your passwords.

To recover encrypted backup data, do the following:

1. In the **Backups** step of the corresponding wizard, select a backup object and a recovery point. Proceed as described below:

- If the password hash is available, the product uses it to decrypt the backup and perform the recovery

- If the password hash is not available, but the AWS KMS was enabled when performing encryption on the data:
  - The product discovers an existing AWS account to get access to AWS KMS and the cryptographic keys stored therein – no need to configure the AWS KMS service in the **Encryption** tab.
  - The product verifies that the AWS account is correct and AWS KMS has a corresponding cryptographic key for decrypting the password hash.
  - The product restores the password hash.
  - Proceed to recovery from an existing encrypted backup. A password hash is used to decrypt the backup.
- If a password hash is not available and AWS KMS was not enabled when performing encryption on the data but salt is available:
  - In the **Backups** step of the corresponding wizard, select a backup object and a recovery point.
  - Enter the password manually.
  - The hash is generated based on the available salt and the provided password. The product uses the password hash to decrypt the backup and perform the recovery.

2. Proceed to the next step of the wizard.

**Important**

If the salt is not available in the recovery point metadata, recovery cannot be performed and the corresponding encrypted recovery point is considered corrupted.

You can find more details on how to perform restore from encrypted backups in the following articles:

- "Recovery Job Wizard for VMware: Backups" on page 1064
- "Recovery Job Wizard for VMware Cloud Director: Backups" on page 1082
- "VMware Flash Boot Job Wizard: Backups" on page 1188
- "Backup Export Wizard: Backups" on page 1098

## How Encryption/Decryption of System Configuration Works

To safely encrypt your system configuration bundle, initiate exporting system configuration in the **Settings > General > System Settings > Configuration** tab, set a password, and proceed to exporting.

To recover from the system configuration bundle, initiate importing system configuration in the **Settings > General > System Settings > Configuration** tab, provide the password to decrypt and import the configuration, and proceed to importing.

Refer to "System Migration" on page 465 for more information.

## Password Management

With NAKIVO Backup & Replication, you can create and manage your passwords for encrypting backups, system configuration bundles stored as self-backup, and FSI data.

**Notes**

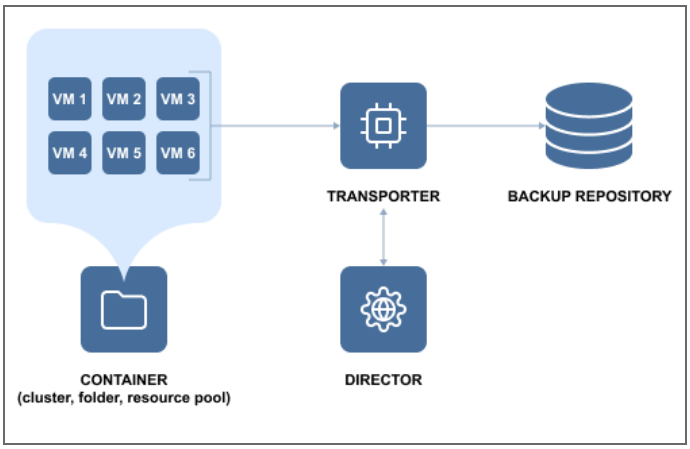
- It's recommended that you enable the **(AWS) Key Management Service** in the **Settings > General > System Settings > Encryption** tab. If **AWS** is enabled, all backup encryption passwords are encrypted with the Key Management Service cryptographic key to be available for recovery in case of product re-installation. For more information, refer to Enabling KMS.
- AWS Key Management Service is not applied to self-backup and system configuration encryption.

Refer to "Managing Passwords" on page 432 for more details.

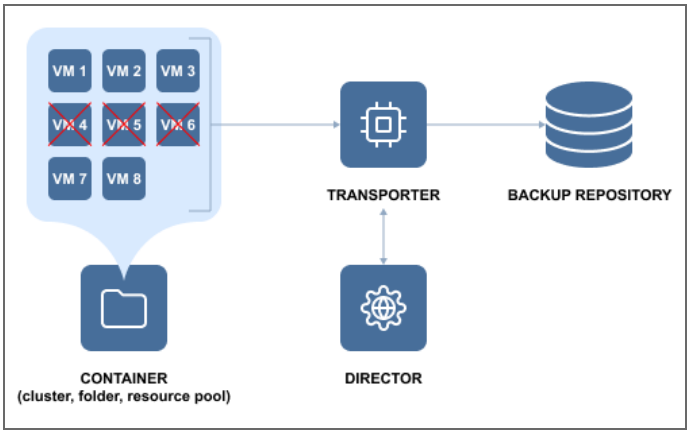
# Container Protection

VMs can be organized into containers, such as resource pools, clusters, and folders. This form of organization allows you to easily add resources upon request and unload them when they are no longer necessary. NAKIVO Backup & Replication allows you to add an entire container to a backup or replication job. All changes in the container (i.e. adding to or removing from) are automatically reflected in a backup or replication job. Thus, all important VMs are continuously protected. If certain VMs inside a container are not required to be backed up or replicated frequently, you can exclude them from a backup or replication job. The container will still be protected but will not include less important VMs. This will save space in the Backup Repository and increase the speed of backup or replication jobs. For example, you set up a backup job for a cluster to run daily, but this cluster contains a couple of rather massive VMs that do not require frequent backups; you can edit the job by excluding those VMs. NAKIVO Backup & Replication will ask you whether to keep or remove backups made on previous job runs.

- Day 1




- Day 2



# File Level Backup

With NAKIVO Backup & Replication, you can select and back up volumes and folders located on Windows and Linux physical machines, and create backup copies of these volumes and folders. When a volume or a folder is selected, all its sub-folders and files are automatically backed up. The product also offers you multiple recovery options such as recovering selected files and folders to a custom share, downloading them to the browser, or forwarding them via email as attachments. You can also send file-level backup copies to tape for long-term data archiving and then recover these backup copies to a target destination. For more details, refer to Backup Copy to Tape Recovery.

To protect your volumes and folders stored on Windows and Linux physical machines, proceed with the following steps:

1. Add a physical machine to the NAKIVO Backup & Replication inventory.
2. Create a file-level backup job.
3. Recover individual volumes and folders using the File Level Recovery Wizard when needed.

For more details on the requirements for this feature, refer to "Feature Requirements" on page 169.

# Proxmox VE Backup

Proxmox VE (Virtual Environment) is an open-source virtualization platform for enterprises. Proxmox VE is Debian GNU/Linux based and allows deployment and management of Virtual machines (QEMU/KVM hypervisor) and Containers (LXC).

This feature allows you to backup Proxmox VE infrastructures, allowing you to:

- Add supported Proxmox VE clusters and standalone hosts (nodes) to the Inventory.
- Back up Proxmox VE VMs and templates.
- Copy Proxmox VE VM backups.
- Create Proxmox VE VM replicas.
- Instantly recover full Proxmox VE VMs from backups using Flash VM Boot.
- Recover full Proxmox VE VMs from backups.
- Perform File Recovery from Proxmox VE VM backups to source or a custom VM.
- Perform Object Recovery from Proxmox VE VM backups to a custom destination.
- Perform Universal Object recovery from Proxmox VE VM backups for a VM disk.

The feature supports **Full** and **Incremental** backups and offers two methods for **Change Block Tracking** (CBT): the native Proxmox VE change tracking and the proprietary method developed by NAKIVO.

The feature supports **Backup encryption** to ensure data is securely encrypted and transferred, preventing unauthorized access and maintaining data confidentiality.

Immutability is supported for backups stored in supported repositories.

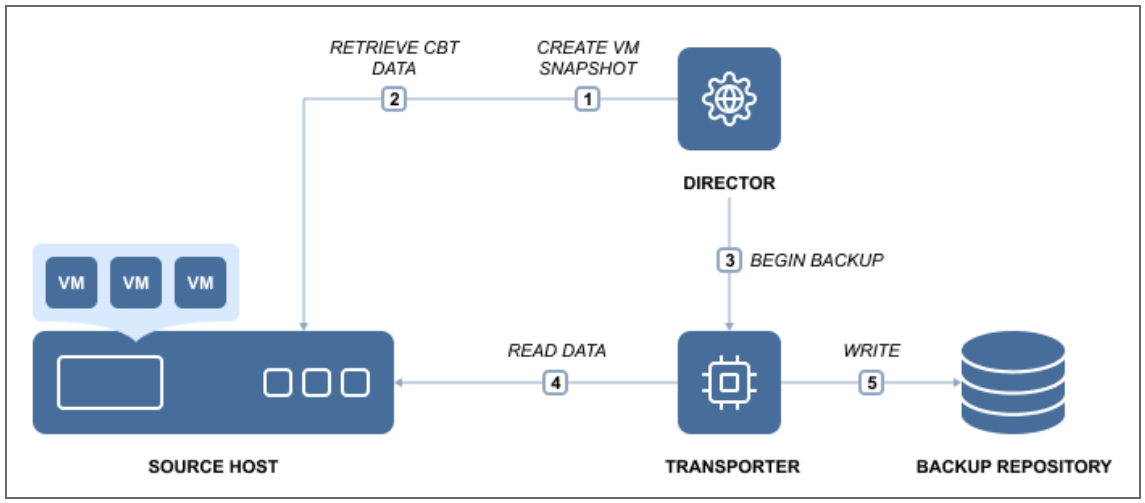
A transporter injected inside a Proxmox VE host or a Proxmox VE node can be edited when locked within a job run.
For a full list of supported functionality, see the corresponding section in Feature Requirements.

# Virtual Machine Backup

NAKIVO Backup & Replication works in a virtual environment and uses an image-based approach to VM backup. It is an agentless application that does not require you to install any additional software inside the VM guest OS to retrieve VM data. It exploits virtualization platforms' snapshot capabilities to back up VMs. When you initiate a VM backup, NAKIVO Backup & Replication requests a virtualization platform to create a VM snapshot which is basically a point-in-time copy of a VM including its configuration, OS, applications, associated data, system state, and so on. The snapshot is used as a source of data for backup. Copying of the data from the source datastore is performed at a block level. NAKIVO Backup & Replication fetches the VM data, performs compression and deduplication, and finally stores the backup files in the repository.

In NAKIVO Backup & Replication, backing up is performed via a job that must be created and configured prior to the backup itself. A backup job is a configuration unit of the backup activity that defines when, what, how and where is to be backed up. One or several VMs can be processed by a single backup job. A job can be started manually or scheduled for execution. The initial job's run always produces a complete backup of the VM image. The following sessions can create full or incremental backups. During incremental backups, NAKIVO Backup & Replication copies only blocks of data that have changed since the last backup job session. Tracking of changed data blocks is performed using the virtualization platforms' capabilities (CBT/RCT/CRT) or with NAKIVO Backup & Replication's proprietary method.

Technically, the VM backup process is performed according to the following flow depicted below:




To learn how to create VM backups with NAKIVO Backup & Replication, refer to “Backup” on page 839.

# Data Recovery

One of the key elements of an effective protection strategy is ensuring that data can be restored quickly after any corruption or loss. NAKIVO Backup & Replication provides several recovery options for maintaining the operational backup of data and business continuity/disaster recovery:

Refer to the following topics for more information about data recovery:


- “Cross-Platform Recovery” on page 28
- “Instant File Recovery to Source” on page 30
- “Instant Object Recovery” on page 31
- “Instant VM Recovery - Flash Boot” on page 32
- “Recovery From Tape” on page 33
- “Backup Copy To Tape Recovery” on page 33
- “Universal Object Recovery” on page 35

# Cross-Platform Recovery

With Cross-Platform Recovery, you can seamlessly protect VM/physical machine data across multiple platforms and virtualized environments. You can also benefit from the following other advantages:

- **Data Migration** – whether a disaster renders one of your hypervisors/physical servers unavailable, or you simply make the decision to switch to a single-platform virtualized environment, Cross-Platform Recovery can be of help. Export your VM or physical machine backup data in the desired format, and know that you can recover on a different platform without encountering any incompatibility issues.
- **Long-Term Data Archiving** – the specifics of your line of business or legislative requirements may require you to store backups for years. With NAKIVO Backup & Replication, you can easily export and store data offsite for as long as you need. Moreover, if your choice of virtualization software changes over time, you shouldn't have any problems recovering from your old backups in the new environment.
- **Recoverability Testing** – the fact that you have a backup does not automatically mean you can recover from that backup. Cross-Platform Recovery gives you the freedom to test different scenarios of recoverability in multiple environments, thus helping ensure business continuity. With Cross-Platform Recovery, no disaster can catch you off guard.

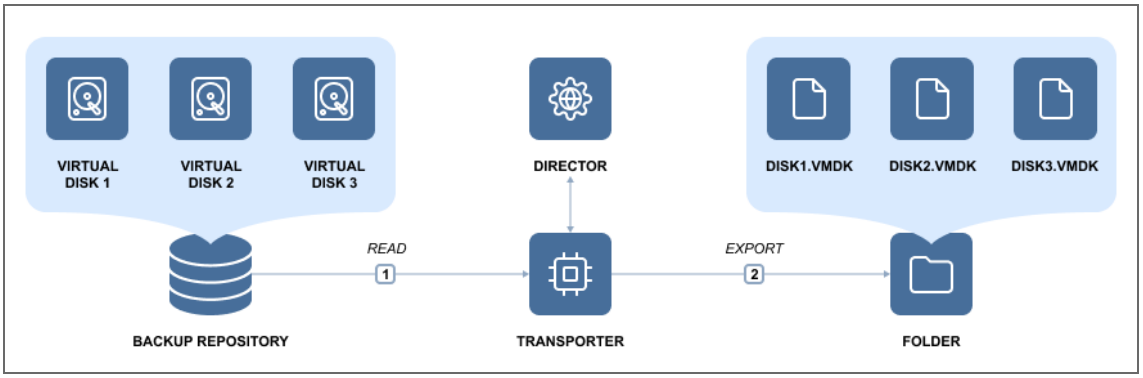
You can export VM/physical machine data from any backup into the format of your choice in four simple steps:

1. Select a backup (VMware, Hyper-V, Nutanix AHV, Proxmox VE, or physical server).

   **Note**

   Backup Export is not supported for *.raw* and *.qcow2* formats of Proxmox.

2. Choose one or multiple virtual disks that you would like to export.
3. Specify the target location and export format (VMDK, VHD, or VHDX).
4. Click a button and have the data of each selected disk exported into a separate file.

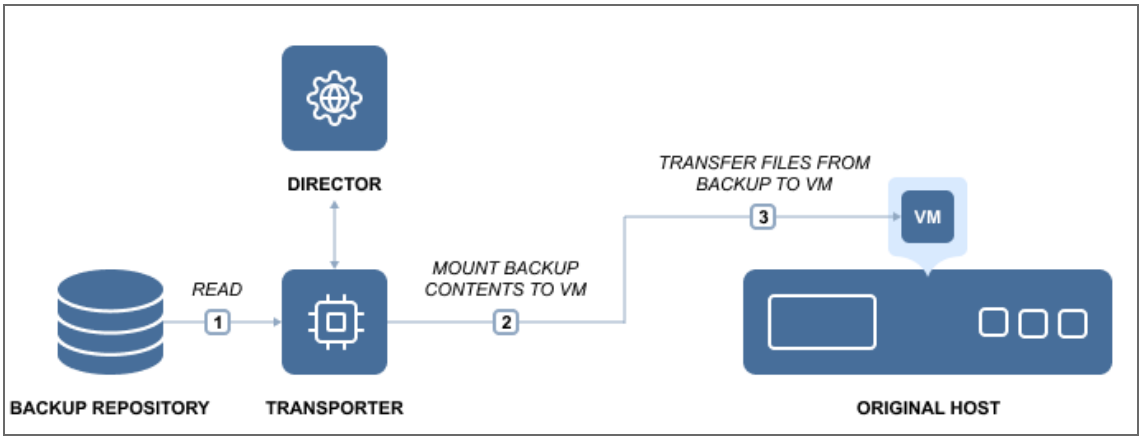


Once exported, the files can be used for recovery or long-term storage. Cross-Platform Recovery allows for unrestricted data protection across different hypervisors, physical machines and cloud platforms. Whether one of your hypervisors or physical machines is down or you need to migrate data from one platform to another, Cross-Platform Recovery gives you the necessary tools for seamless cross-platform data protection and recovery.

# Instant File Recovery to Source

The Instant File Recovery to Source feature allows you to recover files and folders to their original location (or any custom location) in a single click. NAKIVO Backup & Replication can instantly recover files right from compressed and deduplicated backups. Files can be recovered from both Windows and Linux-based machines. With the push of a button, the selected files can be reinstated in their original location or in a new custom location on any VM/physical machine, downloaded to the local machine, or sent via email. When restoring files back to the original location, the file permissions are all restored as well. The Instant File Recovery feature works both via LAN and WAN. Thus, even if local backups are unavailable, you can recover from a backup copy located, for example, in an Amazon EC2 cloud a thousand miles away. Note that recovery to the source is executed via a system account.

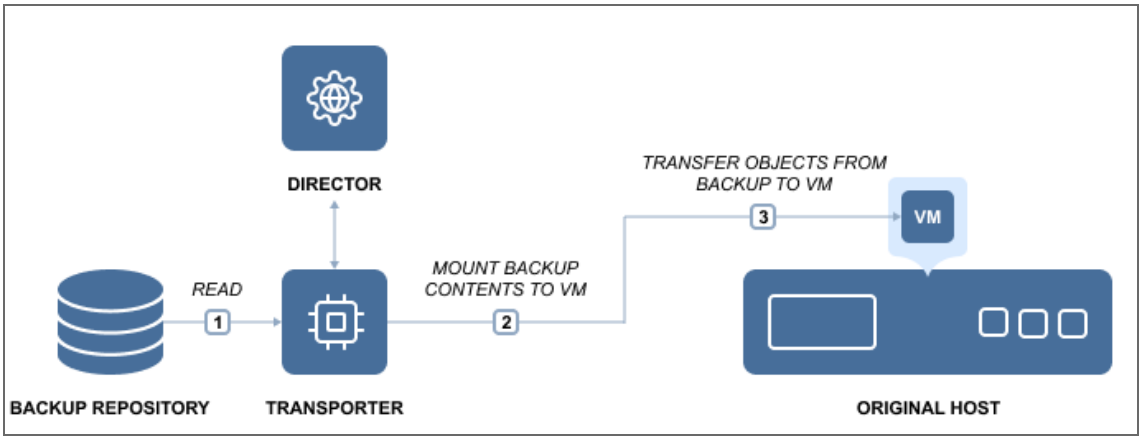
The file recovery process is simple and straightforward. First, select a backup and recovery point from which you wish to recover files. The files and folders available for recovery are displayed right in the NAKIVO Backup & Replication web interface. Browse or search for files, select the files you wish to recover, specify where you want them, click the button, and behold! The files are instantly recovered.




To learn how to recover files with NAKIVO Backup & Replication, refer to "File Recovery" on page 992.

# Instant Object Recovery

NAKIVO Backup & Replication provides you with the ability to instantly browse, search, and recover Microsoft Active Directory, Microsoft Exchange, and Microsoft SQL Server directly from compressed and deduplicated backups. The objects can be restored to the source server, to a different server, or exported to a custom location. The feature streamlines, automates, and speeds up the process of restoring your data, and is available out-of-the-box in NAKIVO Backup & Replication. For more information, refer to "Granular Recovery" on page 991.



# Instant VM Recovery - Flash Boot

The Flash boot feature allows you to boot a VM directly from compressed and deduplicated backups for fast recovery during an outage. When a business-critical machine goes down, every minute of downtime has costly and damaging consequences. With NAKIVO Backup & Replication, you can recover entire machines from their backups in minutes. The Flash Boot feature allows you to boot machines directly from compressed and deduplicated backups without recovering entire machines first. This feature works right out of the box without any special setup. Just choose a backup, a recovery point, and a recovery location (a host, a resource pool, or a cluster where you want to run the recovered machine). Then press the button and your machine is booted in no time.

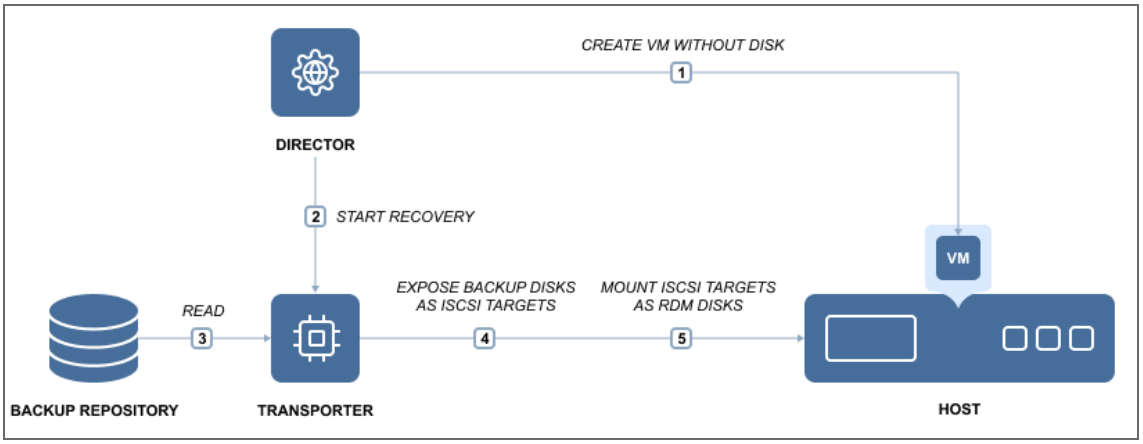
Once the machine is running, you can migrate it to production for permanent recovery. Note that the backup from which the VM is booted is not affected. Changes you make to the running VM will not modify or remove the data in your VM backup. In addition to the VM recovery capabilities, the Flash boot feature offers other useful functions. For example, it allows you:

- Access the files, folders, and application objects of any application on any OS.
- Test system updates and application patches before applying them to your production machine.
- Verify the backup to ensure that the OS and applications run properly.
- Copy a VMDK or VHDX file, and then delete the virtual machine.

This is how the Flash boot feature works:

NAKIVO Backup & Replication consists of two main components: the Director, which is the management component, and the Transporter, which performs actual data protection and recovery tasks. By default, both components are automatically installed to enable all features out of the box.

When you run a Flash boot job, the Director creates a new VM without any disks on the target server, then commands the Transporter to expose the machine disks from the Backup Repository as iSCSI targets. Finally, the Director mounts the exposed disks to the newly created VM.



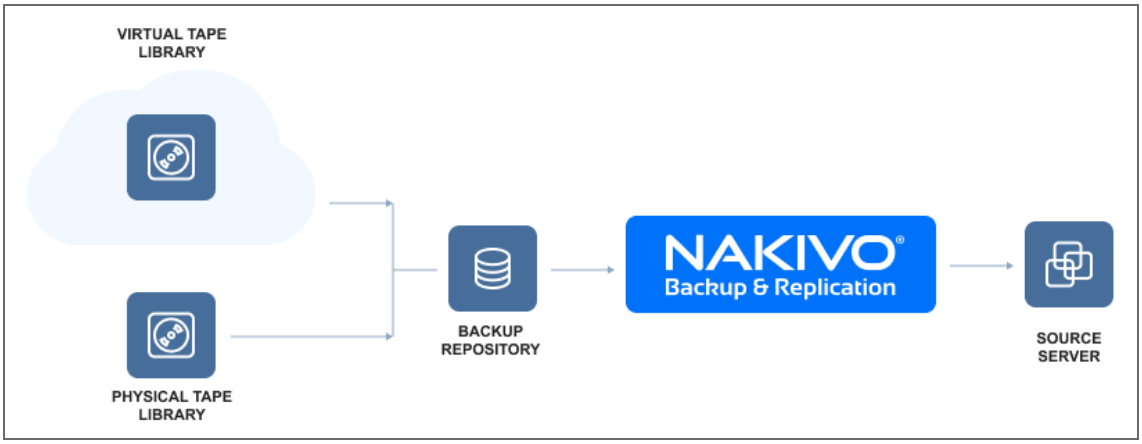
This process is fully automated and takes mere seconds to complete, after which the machine OS boot is started. Once booted, the machine can be migrated to the production environment using the hypervisor's native live migration feature.

With NAKIVO Backup & Replication you can also perform Flash boot to run VMware VMs directly from physical machine backups. If a business-critical physical machine goes down, you can use Cross-Platform Flash Boot for instant recovery without having to manually install a new OS and applications on the new machine. The machine recovered this way can be used as a testing environment and can later be migrated for permanent use.

To learn how to create recovery jobs using the Flash boot feature, refer to "Performing Flash Boot Recovery" on page 1185

# Recovery From Tape

NAKIVO Backup & Replication allows you to recover VMs and EC2 instances directly from tape backups using the standard recovery tools. You can also move backed up data from a tape cartridge to an existing **Backup Repository** if needed.




Refer to the following topics for more information:

- "Starting Recovery from Tape" on page 1093
- "Starting VMware VM Recovery" on page 1062
- "Starting VMware Cloud Director Recovery" on page 1080

# Backup Copy To Tape Recovery

With NAKIVO Backup & Replication, you can send the backups to tape for long-term data archiving and then recover these backup copies to a target backup repository.

The following backup types are supported:

- File Share Backup
- Microsoft 365 Backup

**Note**

File share and Microsoft 365 backup options are disabled in the following recovery job wizards when recovering backups from tape:

- New Recovery Job Wizard for VMware vSphere
- New Recovery Job Wizard for Amazon EC2
- New Recovery Job Wizard for Microsoft Hyper-V
- New Recovery Job Wizard for Nutanix AHV
- New Recovery Job Wizard for Physical Machine

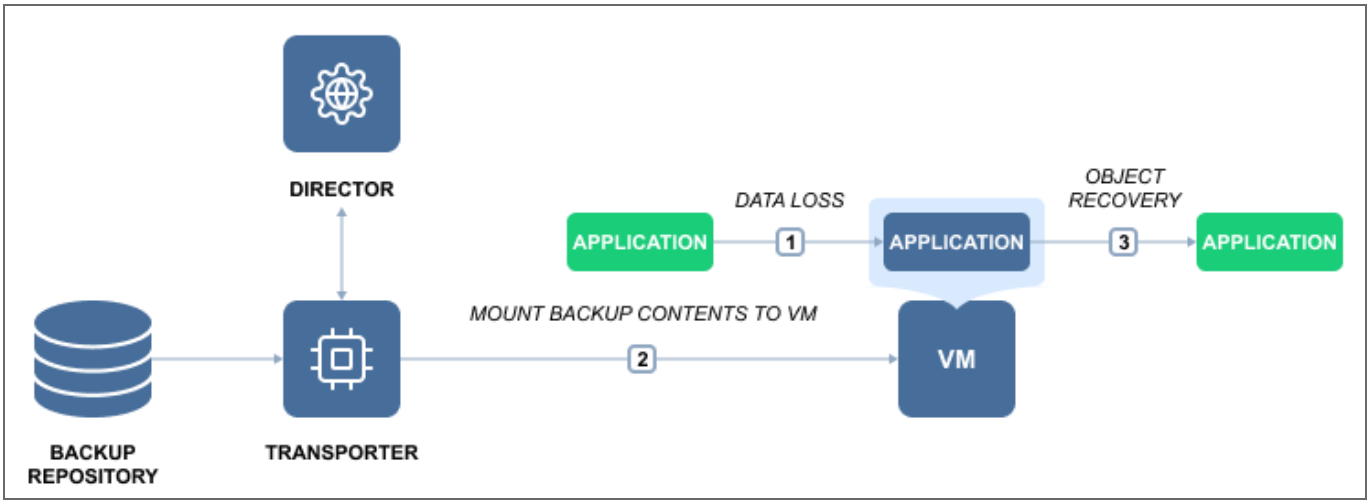
# Universal Object Recovery

The Universal Object Recovery feature allows you to recover any object in the infrastructure – whatever the application or file system – in a matter of minutes by mounting the appropriate backup to a VM or physical machine and then recovering the necessary data using the native application tools.

Universal Object Recovery provides multiple recovery options, increases the flexibility of the recovery process, and saves a significant amount of time.

- Versatility – with Universal Object Recovery, you are not limited to certain applications or file systems: you can recover any object at any time (provided you have a recent backup). Moreover, the feature allows you to recover individual objects back to the source, to another VM or instance, or even to a physical machine.
- Lower Overhead – Universal Object Recovery lets you restore individual objects without having to recover the entire VM or physical machine. Thus, the feature eliminates the complexity of full machine recovery, saving you time that can be better used for other important tasks.
- Faster Recovery – recovering an entire machine from a deduplicated and compressed backup takes time, affecting your ability to meet your RTOs. With Universal Object Recovery, you can instantly mount disks from a backup, decreasing recovery time and ensuring minimal interruptions in your business operations.

You can use NAKIVO Backup & Replication to recover application objects in a few simple steps: just open the Universal Object Recovery Wizard and select the appropriate recovery point. Once you choose the disks you wish to be mounted, NAKIVO Backup & Replication attaches said disks to the specified VM or physical machine. All you need to do after the mount is log into the corresponding VM or physical machine and use native application tools to recover the data.




To learn how to create object recovery jobs with NAKIVO Backup & Replication, refer to the corresponding topics of the "Granular Recovery" on page 991 section.

# Disaster Recovery

Disaster Recovery (DR) is a practice intended to support an organization's ability to remain fully operational after an emergency event. DR serves to limit risks by getting an organization's infrastructure to run as close to normal as possible after an abrupt intermission. NAKIVO Backup & Replication allows you to address all major DR planning points by creating automated DR workflows for VMware, Microsoft Hyper-V, Nutanix AHV and Amazon EC2 environments. The solution allows you to protect VMs running within a cluster, replicate VMs, including creation of real-time replication jobs, perform failover to replicas and replica failback operations.

When utilizing Site Recovery (SR), you can include up to 200 actions to a single job, including failover, failback, start/stop VMs and instances, run/stop jobs, run script, attach or detach repository, send email, wait, and check condition. By arranging actions and conditions into one automated algorithm, you can create site recovery jobs of any complexity level. You can check the “Site Recovery” on page 42 page to learn more.

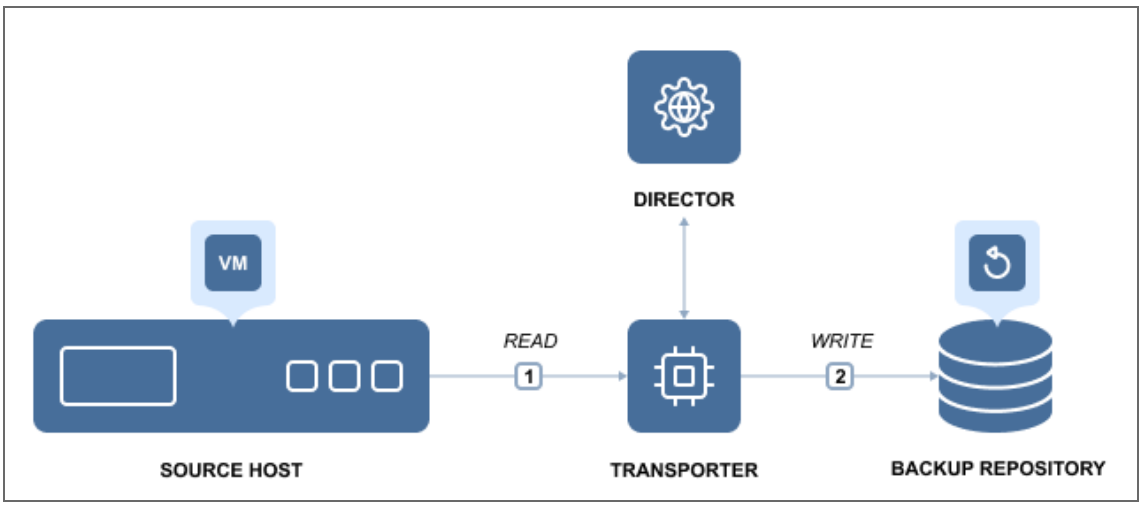
This section contains the following topics:


- “Replication From Backup” on page 37
- “Real-Time Replication for VMware” on page 38
- “Replication Types” on page 40
- “Site Recovery” on page 42

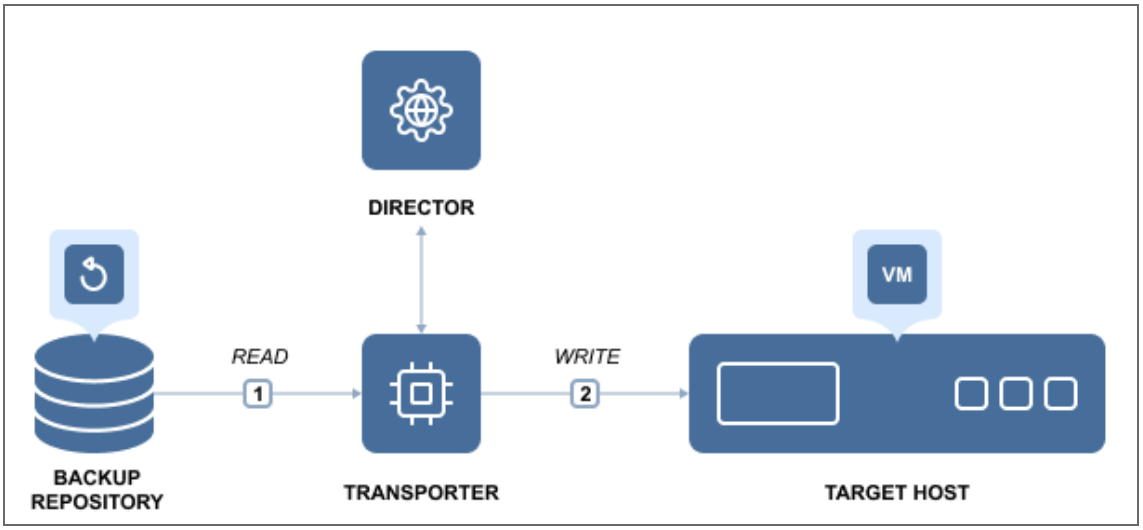
# Replication From Backup

The Replication From Backup feature allows for offloading the production environment by replicating VMs directly from backups.

**Step 1** - Create a backup




**Step 2** - Replicate VM from backup




Setting up a replication from backup job for VMware and Hyper-V environments is no more time-consuming than setting up a traditional replication job. Once you launch a new replication job wizard and select VM backup as the source, NAKIVO Backup & Replication proceeds to read the data from the repository and injects it into the replica.

To learn how to create replication jobs with NAKIVO Backup & Replication, refer to “Replication” on page 933.

# Real-Time Replication for VMware

NAKIVO Backup & Replication allows you to perform real-time replication (RTR) for VMware vSphere VMs. Unlike regular replication jobs, RTR jobs continuously copy data from the source host as it is being generated. In the event of a disaster, this feature allows you to perform a failover to a state mere minutes or even seconds before the disaster occurred.

Real-time replication jobs for VMware vSphere work as follows:

1. A target Transporter is deployed as a VMware vSphere appliance.
2. The **I/O Filter** is installed on source hosts, and **Journal Service** is installed on target transporters located on target hosts.
3. A Recovery Point Objective (**RPO**) is set in the **Real-Time Replication Job Wizard for VMware** .
4. The VMware RTR job is run with the configured **RPO** and **I/O Journal** settings.
5. Data changes made on the source VM are continuously transferred to the installed **I/O Journal** as quickly as possible. The Journal mode defines how data changes are managed on replicas, offering two options:
   - Rollback journal: New changes are written to the replica directly while old data from the replica is saved to the journal.
   - Roll Forward journal: New changes are stored in the journal and later are merged into the replica based on the journal settings.
6. If a disaster occurs, a VMware real-time replication failover job is created and run to restore the VM to a selected recovery point (reference point that is a snapshot of data changes made to a source VM) of a replica VM and can be used for failover.

   A Site Recovery job can also be configured to streamline disaster recovery operations by managing actions such as Run, Stop, and Real-Time Replication Failover in both production and test modes. This job allows automated recovery workflows, including seamless integration of replication failover based on predefined conditions.

**Important**

Before starting a real-time replication job, make sure the time on the source ESXi host(s) and the target **Transporter** VA is synchronized.

Otherwise, the **RPO** value and/or recovery point date-time data may be incorrect. Also, the **Real-Time Replication Failover Job Wizard for VMware** process may yield the incorrect recovery points.

For more information on using real-time replication and the above terms, refer to the following pages:

- “Installing I/O Filter” on page 583
- “Installing Journal Service” on page 590

- “Creating VMware Real-Time Replication Jobs” on page 968
- “Real-Time Replication Failover for VMware” on page 1136

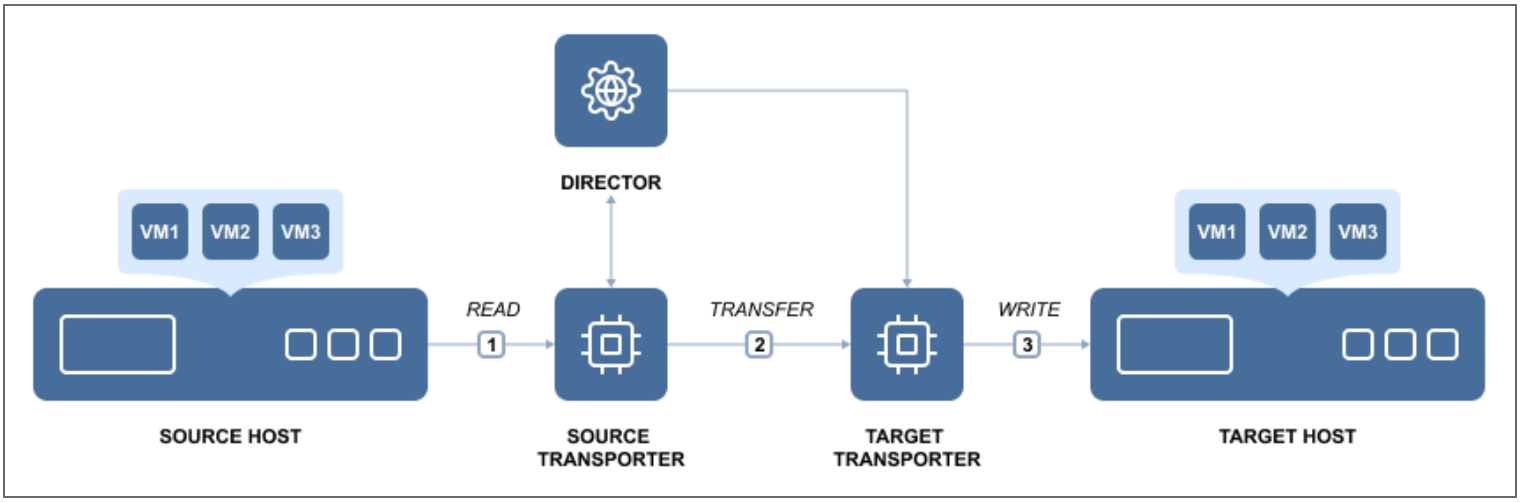
# Replication Types

NAKIVO Backup & Replication allows you to replicate virtual machines and Amazon EC2 instances.

- VM Replication
- Amazon EC2 Replication

## VM Replication

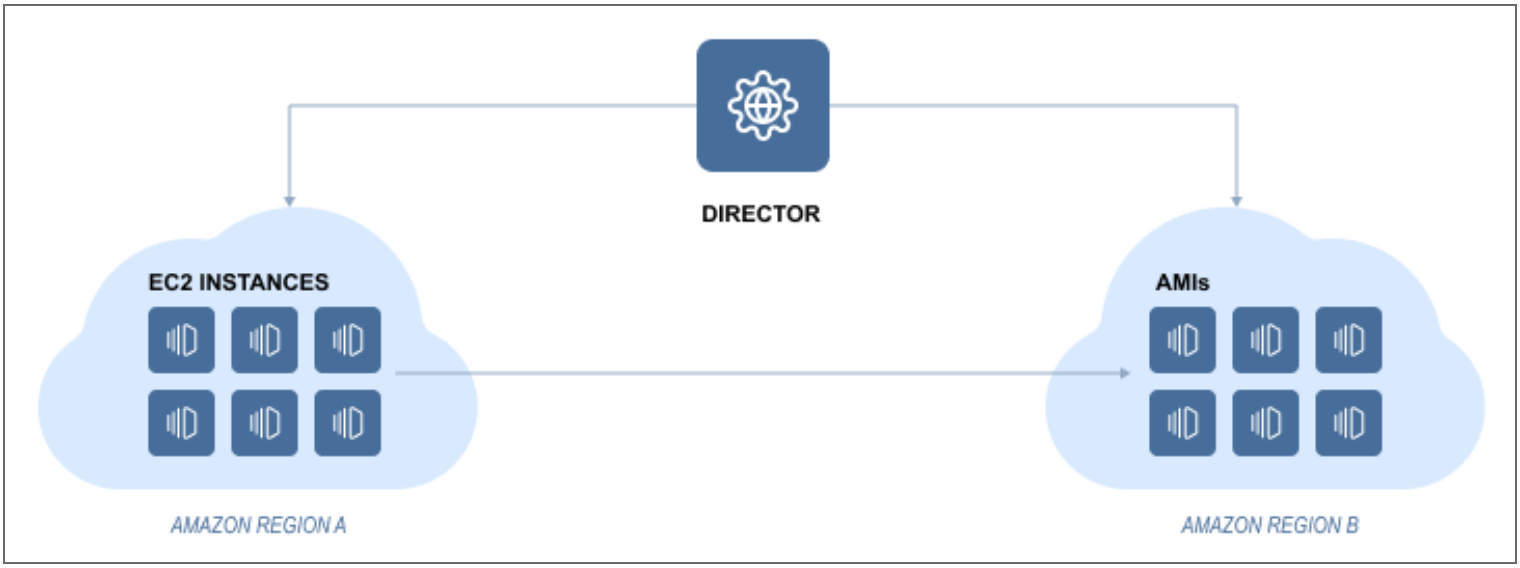
A VM replica is an exact copy of an entire VMware or Hyper-V VM created on a target host. VM replication ensures business continuity as it lets you immediately power on the replica of any failed primary VM at any time.




Replicas are stored on the target hosts in a powered-off state, and so do not consume any resources. If the source VM has been damaged, you simply need to power the replica without NAKIVO Backup & Replication. NAKIVO Backup & Replication provides you with the ability to add both individual VMs and all VMs within selected VMware containers (such as resource pools, folders, hosts, clusters, etc.) to a replication job. That is to say, all your new VMs that are created in or moved to a protected container are automatically added to your replication job. You can store up to 30 recovery points for each replica. Even if the source VM was replicated after an error, you can always revert to the last working copy.

## Amazon EC2 Replication

Amazon EC2 replication creates identical copies (i.e. replicas) of your Amazon EC2 instances as AMIs, ensuring business continuity in case primary instances or the whole region become unavailable. NAKIVO Backup & Replication allows for replicating instances either inside the same region or to another one.

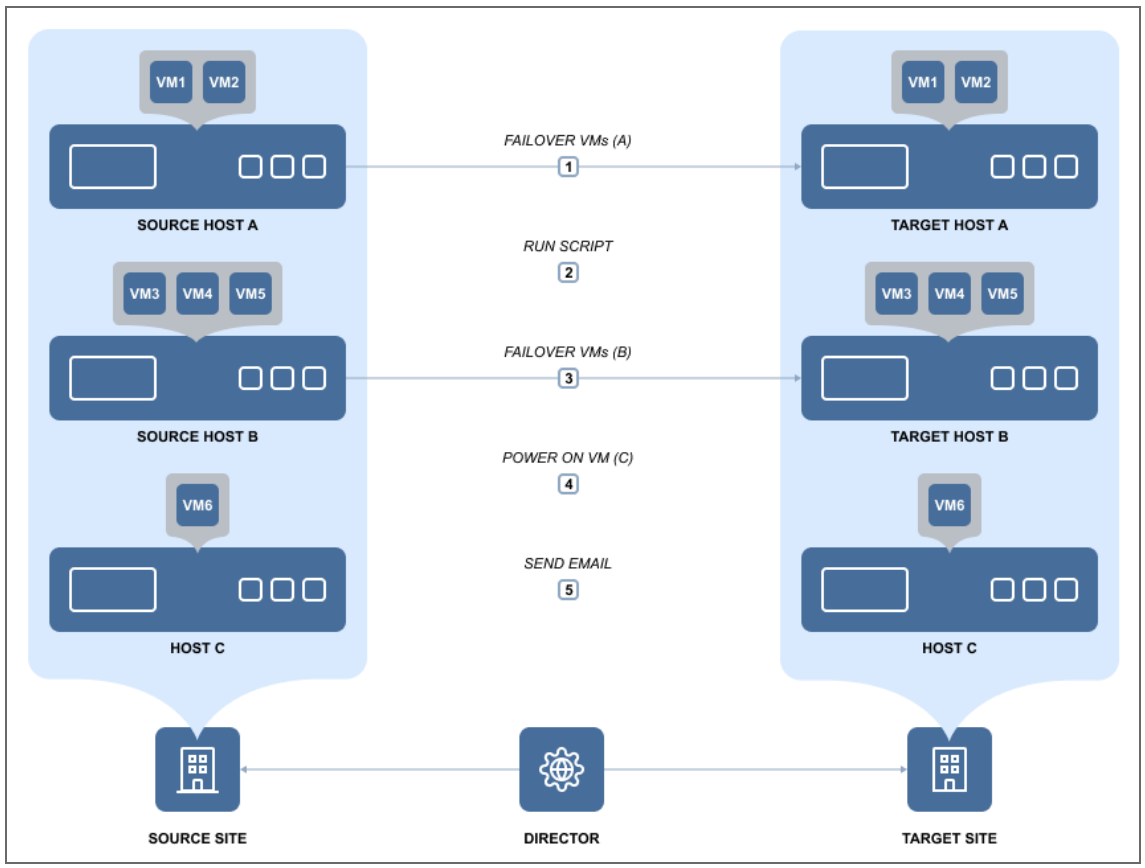



In the case of same region replication, NAKIVO Backup & Replication first initiates the creation of snapshots of selected source volumes and then creates AMIs from the source instance configuration and created snapshots. In case of cross-region replication, NAKIVO Backup & Replication initiates the creation of snapshots of selected source volumes in the source region. Then the product copies those snapshots to the target region and creates AMI from source instance configuration and copied snapshots. Finally, NAKIVO Backup & Replication deletes the snapshots in the source region.

To learn how to create VMware, Hyper-V, and Amazon EC replication jobs, refer to the corresponding topics in “Replication” on page 933.

# Site Recovery

Using Site Recovery Jobs, NAKIVO Backup & Replication can automate the execution of one or more actions. An action is a single task that can be included in a Site Recovery Job. The list of available actions can be found under "Site Recovery Job Wizard: Actions" on page 1149.




Special Actions used in recovering your IT environment with a Site Recovery Job are Failover and Failback:

- **Failover** switches workloads from the primary location to a secondary recovery location. With Failover action, you can temporarily suspend workloads on the primary location, and start them from the recovery location.
- **Failback** is the process of synchronizing data that has changed since Failover finished, back to the primary location. With Failback action, you can stop workloads on the secondary location and switch them back to the primary location.

Failover and Failback actions are applicable to replicas, and switch replica states from *Failover* to *Normal* correspondingly. The Site Recovery Job can be executed in one of the following modes:

- **Test mode** is designed to verify the Site Recovery Job workflow and results. You can execute a Site Recovery Job in the test mode on demand or on schedule. Refer to "Running Site Recovery Job in Test Mode" on page 1182 for details.

- **Production mode** is designed to recover the environment from a disaster. You can execute a Site Recovery Job in the production mode on demand only. Refer to "Running Site Recovery Job in Production Mode" on page 1183 for details.

When the Site Recovery Job is run in the production mode, Failover may be either of the following types:

- **Planned failover** is designed to achieve zero data loss when disaster happens. The application will sync replica data with the source VM before switching workloads to the replica.
- **Emergency failover** is designed to minimize downtime. The application will switch workloads from the source VM to the replica immediately.

The topic includes the following sections:


- Workflow of Site Recovery Job
- Cleanup of Site Recovery Job Testing


## Workflow of Site Recovery Job

If your Site Recovery Job contains a Failover action, the action will be executed as follows:

- Site Recovery Job is executed in **production** mode as **Emergency Failover** is being carried out:
  1. Replication from the source VM to the replica will be disabled.
  2. The replica will be rolled back to a specified recovery point (optional, as the latest recovery point is used by default).
  3. The replica will be connected to a new network (optional).
  4. The static IP address of the replica will be modified (optional).
  5. The source VM will be powered off (optional).
  6. The replica will be powered on.
  7. The replica will be switched to the Failover state.
- Site Recovery Job is executed in **production** mode as **Planned Failover** is being carried out:
  1. Replication from the source VM to the replica will be disabled.
  2. An incremental replication from the source VM to the replica will be run once.
  3. The source VM will be powered off.
  4. An incremental replication from the source VM to the replica will be run once more.
  5. The replica will be connected to a new network (optional).
  6. The static IP address of the replica will be modified (optional).
  7. The replica will be powered on.
  8. The replica will be switched to the Failover state.
- Site Recovery Job is executed in **test** mode:
  1. Replication from the source VM to the replica will be disabled.
  2. An incremental replication from the source VM to the replica will be run once.
  3. The replica will be connected to an isolated network (optional).

4. The static IP address of the replica will be modified (optional).
5. The replica will be powered on.
6. The replica will be switched to the Failover state.

If your Site Recovery Job contains a Failback action, the action will be executed as follows:

- Site Recovery Job is executed in the **production** mode:
  1. The source VM will be powered off (if it exists and is powered on).
  2. A protective snapshot of the source VM will be created.
  3. An incremental or full replication from the replica to the source VM will be run once.
  4. The replica will be powered off (optional).
  5. An incremental replication from replica to the source VM will be run once more.
  6. The source VM will be connected to a new network (optional).
  7. The static IP address of the source VM will be modified (optional).
  8. The source VM will be powered on.
- Site Recovery Job is executed in **test** mode:
  1. The source VM will be powered off (if it exists and is powered on).
  2. A protective snapshot of the source VM will be created.
  3. An incremental or full replication from replica to the source VM will be run once.
  4. The source VM will be connected to an isolated network (optional).
  5. The static IP address of the source VM will be modified (optional).
  6. The source VM will be powered on.

# Cleanup of Site Recovery Job Testing

After executing a Site Recovery Job in test mode, the cleanup will be carried out as follows:

1. VMs that have been powered on during the Site Recovery Job testing will be powered off, and vice versa.
2. Repositories that have been attached during the Site Recovery Job testing will be detached, and vice versa.
3. Jobs that have been enabled during the Site Recovery Job testing will be disabled, and vice versa.
4. If the **Failover** action was part of the Site Recovery Job testing:
   a. The replica will be powered off.
   b. The replica will be reverted to the pre-Failover state via the snapshot.
   c. The replica will be switched to the Normal state.
   d. Replication from the source VM to the replica will be enabled.
5. If the **Failback** action was part of the Site Recovery Job testing:
   a. The source VM will be removed (if it did not exist before the Site Recovery Job testing), or else:
   b. The source VM will be reverted to the protective snapshot.

c. The source VM will be powered on (if it exists and was powered off).
d. The protective snapshot will be removed from the source VM.

# Reliability

NAKIVO Backup & Replication employs various techniques to ensure that data is stored, transferred and recovered correctly and consistently.

This section contains the following topics :


- “Application and Database Support” on page 47
- “Backup Immutability” on page 48
- “Backup Size Reduction” on page 50
- “Direct Connect” on page 54
- “Encryption in Flight and at Rest” on page 56
- “External Product Database Support” on page 57
- “Log Truncation” on page 58
- “Recovery Point Retention” on page 59
- “Self-Backup Feature” on page 63
- “Two-Factor Authentication” on page 65
- “VM Verification” on page 66

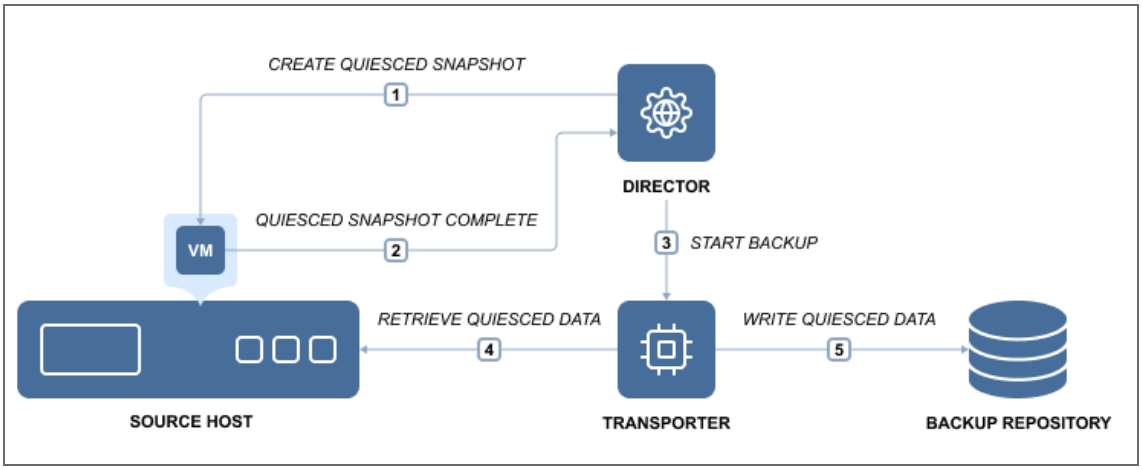
# Application and Database Support

When you back up a VM that runs Active Directory, Microsoft SQL Server, Microsoft Exchange, or any other application or database, it is crucial to ensure that all data inside of those applications remain consistent in the backup. This is important because portions of data and some transactions kept in memory may be incomplete when the VM backup is made. If you take no actions to flush memory and I/O operations, the backups will be crash-consistent. It is similar to pulling the plug on a physical server and then powering it back on. Therefore, most modern applications and databases offer ways to recover from this state. However, in most cases you'll still need to spend some time on manual restore operations and run the risk of losing important data.

To ensure that all data is consistent in the backups, NAKIVO Backup & Replication allows you to use the application awareness feature which is called app-aware mode. To perform consistent backups and replicas of Windows-based environments, the product relies on the Microsoft Volume Shadow Copy (VSS) service running inside VMs. If your application is not VSS-aware or runs on Linux, it provides you with the ability to run custom pre-freeze and post-thaw scripts to enable application-consistent VM backup and replication. A pre-freeze script is executed before a snapshot of a VM is taken, and post-thaw script is executed after the snapshot has been taken.

With the app-aware mode turned on, your backups and replicas will contain consistent application and database data, so you won't need to take any extra configuration steps. As a result, you will be able to instantly recover not only full VMs, but also Microsoft Exchange and Active Directory objects, such as emails or users, directly from a compressed and deduplicated backup. If app-aware mode is disabled, NAKIVO Backup & Replication will create normal (standard) snapshots of source volumes instead of quiesced ones. In case of failure, the product will copy data directly from source volumes without displaying an error.

The app-aware mode can be enabled/disabled on the page of the backup and replication job wizard of all supported platforms.



# Backup Immutability

When creating a backup job and selecting the **Amazon EC2**, **Amazon S3**, **generic S3-compatible storage**, **Wasabi**, **Azure Blob Storage**, **Backblaze B2 Cloud Storage**, **Dell EMC DataDomain**, **NEC HYDRAstor**, or **Local Folder** type of Backup Repository, NAKIVO Backup & Replication allows you to make the recovery points in these repositories immutable. With immutability enabled, the recovery points are stored using the write-once-read-many (WORM) model. Immutability adds another layer of security to backups by protecting recovery points from encryption by ransomware or accidental deletions/modifications.

For the **Amazon EC2**, **Amazon S3**, **generic S3-compatible storage**, **Wasabi**, **Azure Blob Storage**, **Dell EMC DataDomain**, **NEC HYDRAstor**, or **Backblaze B2 Cloud Storage** type of Backup Repository, Object Lock or version-level immutability support should be enabled for the bucket or blob container used to store backups. This type of immutability cannot be shortened or lifted, not even by the root user.

For **Google Cloud Storage** type of Backup Repository, the Backup Immutability is not supported.

With the **Local Folder** type of Backup Repository, immutable recovery points cannot be overwritten, deleted, or changed by anyone except the root user before the specified period expires.

When the **Local Folder** type of Backup Repository is deployed as part of a VMware vSphere, Nutanix AHV virtual appliance, or a pre-configured AMI in Amazon EC2, NAKIVO Backup & Replication provides an even higher level of ransomware protection. You can make recovery points stored in this type of repository immutable, and no user, not even the root user, can lift the immutability.

You can find more details on how to enable immutability in the following articles:

- “Deploying VMware Virtual Appliance” on page 232
- “Deploying Nutanix AHV Virtual Appliance” on page 239
- “Deploying Amazon Machine Image in Amazon EC2” on page 247
- “Backup Job Wizard for VMware: Retention” on page 856
- “Backup Copy Job Wizard: Retention” on page 889

For more details on the requirements for this feature, refer to this page in the User Guide.

# Backup Malware Scan

With this feature, NAKIVO Backup & Replication can scan the selected backups for malware during the recovery process and perform specified actions if malware is detected.

When the feature is enabled, the product uses a supported antivirus software installed on a Repository Transporter or on a designated **Scan Server** added to the inventory to detect if malware is present in a backup.

You can find more details on how to enable recovery with malware detection in the following articles:

- Adding Scan Servers
- Recovery Job Wizard for VMware: Options

- Recovery Job Wizard for VMware Cloud Director: Options
- VMware vSphere Flash Boot Job Wizard: Options
- Universal Object Recovery Wizard: Options

For more details on the requirements for this feature, refer to the Feature Requirements.

# Backup Size Reduction

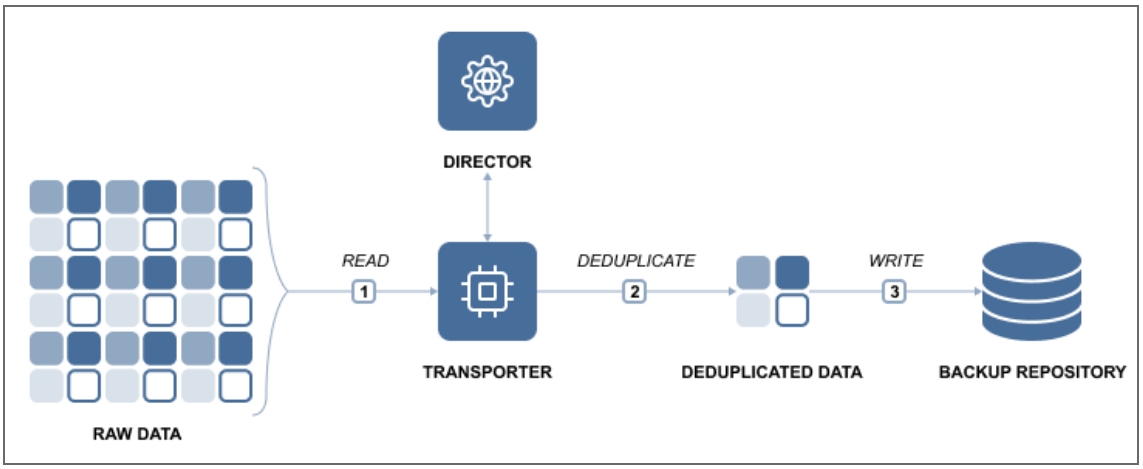
NAKIVO Backup & Replication utilizes multiple methods, such as deduplication and compression, to optimize the size of stored backups. The main purpose of these methods is to reach the correct balance between the amount of data read and transferred during backup.

This section contains the following topics:


- "Global Data Deduplication" on page 51
- "Excluding Swap Files and Partitions" on page 52
- "Excluding Unused Blocks" on page 53

# Global Data Deduplication

Backup deduplication is a method for reducing backup size by excluding duplicate data blocks from the backup. In any given organization, VMs contain duplicates of data, such as VMs deployed from the same template, VMs with the same OS, and VMs that have some (semi) identical files, such as database entries. Block-level data deduplication enables you to reduce backup size by saving only unique data blocks to the Backup Repository while replacing duplicated blocks with references to existing ones.




NAKIVO Backup & Replication automatically deduplicates all backups in a given forever-incremental Backup Repository if this feature is enabled. This means that all data blocks are taken into account by backup deduplication, even if you back up your VMware VMs, Hyper-V VMs, and Amazon EC2 instances to the same Backup Repository. Global deduplication can be enabled during Backup Repository creation process. You can also use hardware-based data deduplication device such as an EMC Data Domain instead of enabling it for the repository.

**Note**

The backup deduplication method can be enabled/disabled during the Backup Repository creation process. For details, refer to one of the Creating Backup Repositories sections.

VM backup deduplication can provide a 10X to 30X reduction in storage capacity requirements. For example, you have 10 VMs running Windows Server 2016, which occupies 10 GB each. While the total amount of data is 100 GB, only one copy of OS data (10 GB) will be written to a backup repository with data deduplication, which provides 10 to 1 storage space savings.

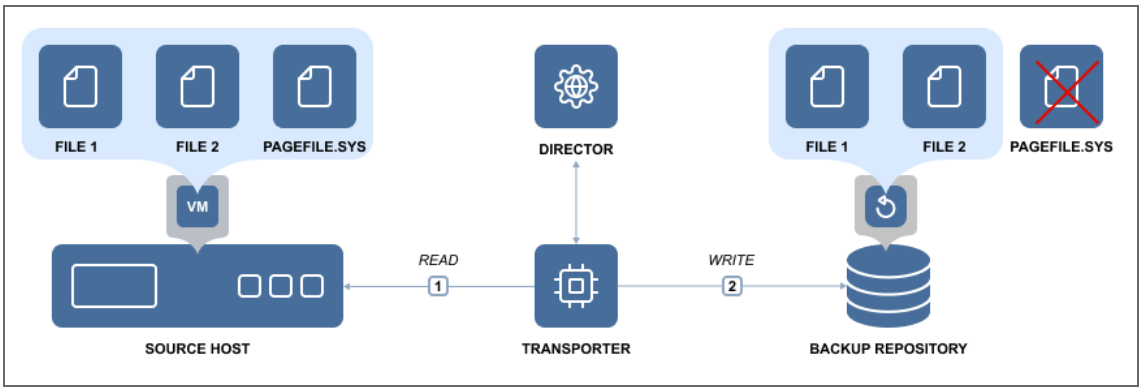
More efficient disk space utilization allows for storing more recovery points per VM backup. In addition, lower storage space requirements save money on direct storage costs (as fewer disks are needed to store the same amount of information) and on related costs (such as cooling, electricity, and maintenance).

# Excluding Swap Files and Partitions

Swap files on Windows OS and swap partitions on Linux OS serve as “virtual memory” and store temporary runtime data that is not in use by RAM. Swap files and partitions improve OS performance: Once the physical memory is full, the OS can send less frequently used data to a swap file/partition and use the freed-up physical memory to perform high priority tasks. While this approach is great for OS and application performance, it has a negative effect on VM backup and replication.

The contents of the swap file change constantly, so each time you run a VM backup or VM replication, the swap file/partition is included in the backup/replica. Since the swap file can automatically grow up to 3x the size of RAM, gigabytes of unnecessary data are processed, transferred and stored each time you back up a VM. The impact of swap files and partitions on backup and replication is significant even in small environments. For example, if you run a backup for 10 VMs and each VM has just 2 GB of swap data, you will transfer and store: 10 VMs x 2 GB x 22 working days = 440 GB of useless data in one month alone.

NAKIVO Backup & Replication automatically excludes swap files and partitions in VMware VMs, Hyper-V VMs, and Amazon EC2 instances, which results in faster and smaller backups and replicas. Note that the application-aware mode instructs applications and databases running inside VMs to flush their data from memory to disk, which means that all important data will be included in your VM backups and replicas. This option can be enabled on a per-job basis.



## Excluding Unused Blocks

In addition to excluding swap files and partitions, NAKIVO Backup & Replication allows you to exclude unused disk blocks during the backup or replication process. This includes the following fragments within the file system:

- Never used volume area.
- File area used by deleted files (without hard reference).

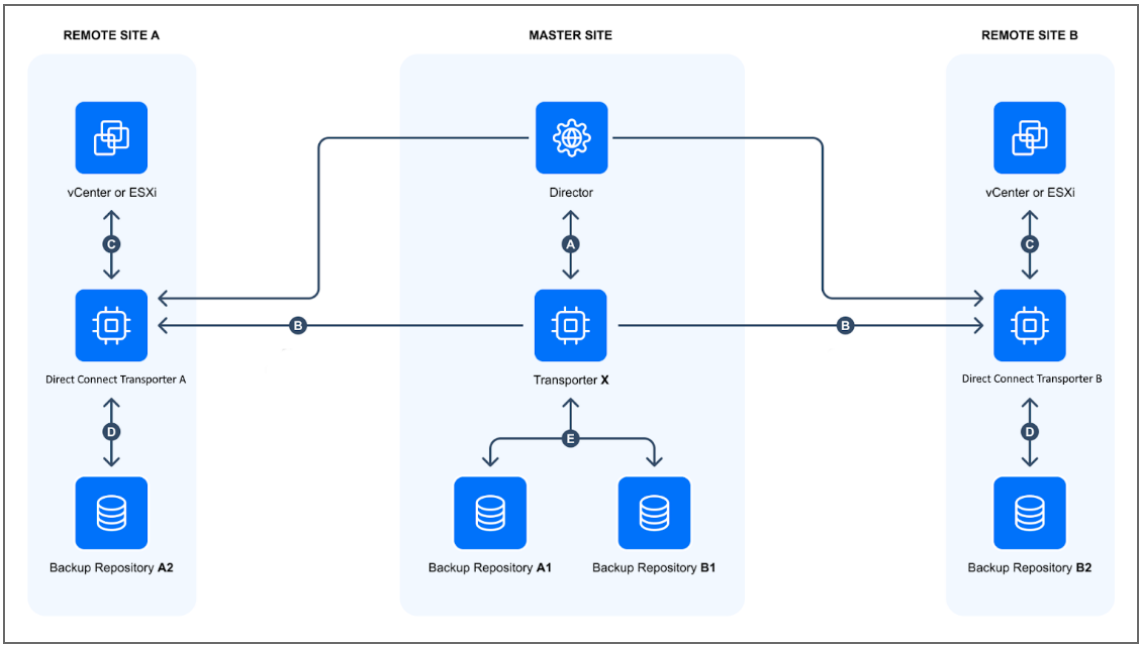
Enabling this option reduces the size of backups and replicas, ensuring that only relevant data is copied. Excluding these blocks of data also means that less processing power and time are required for the workflow to finish.

This option can be configured on a per-job basis on the **Options** page of backup and replication jobs and is enabled by default. The feature supports processing source objects running on Windows OS. It is available for the NTFS file system.

# Direct Connect

With NAKIVO Backup & Replication, you can access remote environments for backup and recovery operations securely and efficiently with the Direct Connect feature. Direct Connect is particularly beneficial for MSPs managing client infrastructures remotely, allowing them to provide data protection services without relying on persistent VPN connections.

To use Direct Connect, a Transporter is installed on the client's site with Direct Connect enabled for it and then connected to the MSP Director. You do not need to install the Director or any other component besides a Transporter on the client's site. After you add the Direct Connect Transporter to the product, you can use it for to discover/refresh the client's Inventory, create/manage backup repositories, export backups, and create/run backup, replication, recovery, and Site Recovery jobs.




A **Direct Connect Transporter** is a Transporter deployed at a remote environment with the Direct Connect feature enabled.

To enable Direct Connect for a Transporter, follow the instructions in this article. Make sure to check the Feature requirements.

There are two Direct Connect approaches available, depending on the solution license: Direct Connect and MSP Direct Connect.

# Direct Connect

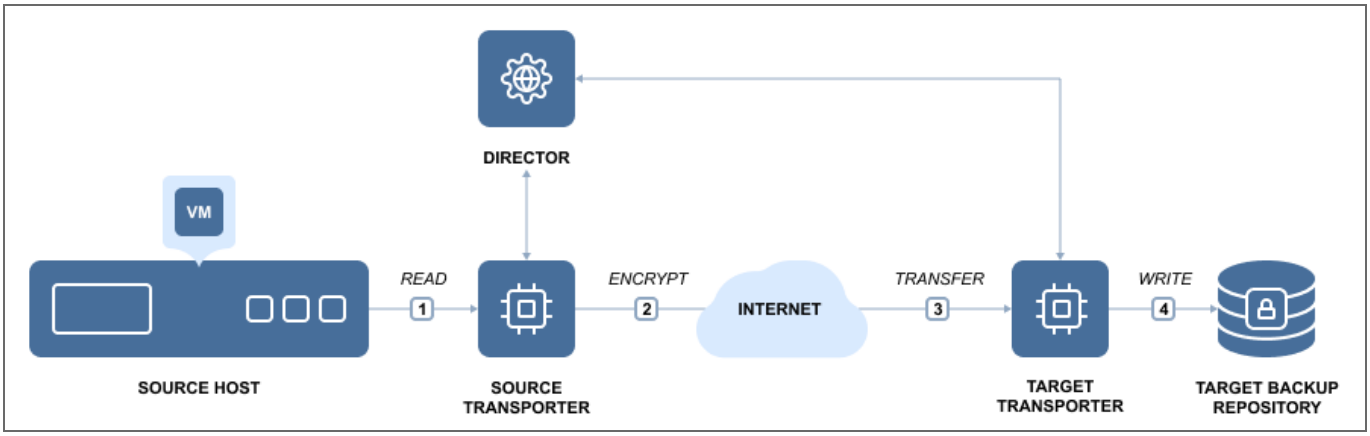
The Direct Connect feature requires exposing the Transporter port on the remote site's local machine to make it externally accessible via the internet, allowing the Transporter to communicate with the multi-tenant Director for backup and recovery operations.

# MSP Direct Connect

The MSP Direct Connect feature does not require exposing any ports at the remote site, providing enhanced security and simplified configuration. MSP Direct Connect requires an MSP License.

# Encryption in Flight and at Rest

VM backup encryption uses a mathematical algorithm that transforms source information into a non-readable cipher text. The goal of VM backup encryption is to make your data unintelligible to unauthorized readers and impossible to decipher when attacked. VM backups that are sent over the Internet should be encrypted before the first bit leaves your organization and travels over the WAN (backup encryption in flight). If the destination is not secure, your data should remain encrypted as well (backup encryption at rest).




NAKIVO Backup & Replication uses AES 256 encryption to protect VM backups, which is the de facto worldwide encryption standard that secures online information and transactions by financial institutions, banks, and e-commerce sites.

- VM Backup Encryption in Flight
- VM Backup Encryption at Rest

## VM Backup Encryption in Flight

VM backup encryption in flight is performed by a pair of Transporters. The Transporter is a component of NAKIVO Backup & Replication that performs all data protection and recovery tasks: data read, compression, deduplication, encryption, transfer, write, verification, granular and full VM recovery, and so on.

The source Transporter for the offsite backup encrypts and sends the encrypted data. The target Transporter receives and decrypts data. For example, when you back up VMs over the WAN to an offsite location, the Transporter installed in the source site compresses and encrypts VM data before transferring it over WAN. Then, the Transporter installed in the Target site receives and unencrypts the data prior to writing it to the Backup Repository.

> **Note**
> Data compression is performed before encryption to maintain its efficiency. Encryption does not interfere with or degrade the results of compression.

## VM Backup Encryption at Rest

It is equally important for the data at rest to be secured by encryption. NAKIVO Backup and Replication provides you with the ability to encrypt Backup Repositories so that backup data at rest, housed in the repository itself, is secure. You can set up encryption on the Options page of the repository creation wizard. For details, refer to the following topics:

- “Local Backup Repository” on page 615
- “Backup Repository on CIFS Share” on page 620
- “Backup Repository on NFS Share” on page 625
- “Backup Repository in Amazon EC2” on page 630
- “Backup Repository on Deduplication Appliance” on page 657

# External Product Database Support

With NAKIVO Backup & Replication, you can use an external database for the Director instead of the built-in database. This feature can help you avoid corruption of the built-in database, which can sometimes occur in large environments. You can migrate the existing database to a supported external database at any time. The feature is available for both the single-tenant and the multi-tenant modes of the product. For more information, refer to the following articles:

- Database Options
- Troubleshooting External Database Connection Issues

# Log Truncation

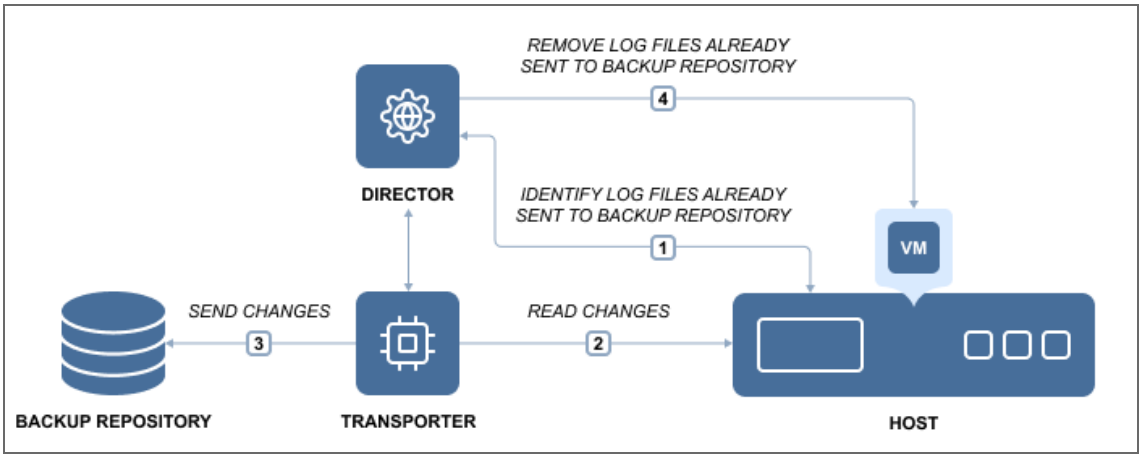
With NAKIVO Backup & Replication, you can remove (truncate) transaction log files of Microsoft Exchange and Microsoft SQL servers which will allow you to reduce the size of backups and, as a result, to optimize the use of storage space. Log truncation can be enabled on the **Options** page of backup and replication jobs.

- Microsoft Exchange Server Log Truncation
- Microsoft SQL Server Log Truncation

## Microsoft Exchange Log Truncation

Microsoft Exchange is the industry's leading platform for email, calendaring, and messaging services. To protect data from undesired deletion or modification, each change that is made to a Microsoft Exchange server database is recorded in transaction logs. These logs can be replayed to recover data that was removed or changed in the database. While this approach improves data protection, it has a downside. Since the Microsoft Exchange database is constantly changing (as data is written and removed in the database), transaction logs grow over time. If not periodically removed, they will eventually fill up the disk and may crash the entire server.

NAKIVO Backup & Replication can create consistent backups of VMware and Hyper-V VMs as well as remove transaction log files of Microsoft Exchange 2013, 2016, and 2019 servers. After creating a successful backup, NAKIVO Backup & Replication connects to your Microsoft Exchange server, identifies which transaction log files have already been written to the database and removes or truncates those log files.




As a result, NAKIVO Backup & Replication creates regular, application-consistent backups of your Microsoft Exchange server and also removes the transaction log files so they don't consume all free disk space on the server.

## Microsoft SQL Server Log Truncation

Any Microsoft SQL server tracks all database transactions (modifications) completed by the server and records them to the transaction logs. Transaction log files (identified with the .ldf extension) are very important, as they are used to ensure database integrity and allow restoring data by replaying the changes. However, these files grow over time and can eventually fill all the free space. This may result in the Microsoft SQL Server crash, or loss of valuable data. That is where Transaction Log Truncation might help.

On one hand, you need to keep the transaction logs, so you can recover Microsoft SQL Server data in case any data deletion, undesired modification, or corruption occurs. On the other hand, you need to remove transaction logs to save space, but without any transaction records you will be unable to successfully recover, should any unpredictable situation occur.

The best practice is to first back up the whole VMware or Hyper-V VM running Microsoft SQL Server and all log files stored therein, and then delete or truncate those files on the source VM freeing up the storage space.

NAKIVO Backup & Replication supports transaction log truncation for Microsoft SQL Server 2008 and later. The product follows the best practice of performing the log truncation process while ensuring ease of use and simplicity. NAKIVO Backup & Replication can automatically truncate transaction log files after successful VM backup and replication. All you need to do is just set it and forget it.

To free up the VM storage space, NAKIVO Backup & Replication performs the following operations:

- Backs up/replicates the entire VMware or Hyper-V VM running Microsoft SQL Server.
- After completing a successful backup/replication, identifies Microsoft SQL Server transaction log files, which were already committed to the database.
- Truncates (deletes) the committed transaction log files on the source VM, thus freeing up storage space.

Consequently, you get a VM backup/replica with all transaction log files. Even though the backed up log files can be pretty large, NAKIVO Backup & Replication easily reduces the size of the VM backup by using backup deduplication and compression features. In its turn, the original VM is left logs-free and can be recovered at a certain recovery point using the aforementioned VM backup/replica, should something go wrong.

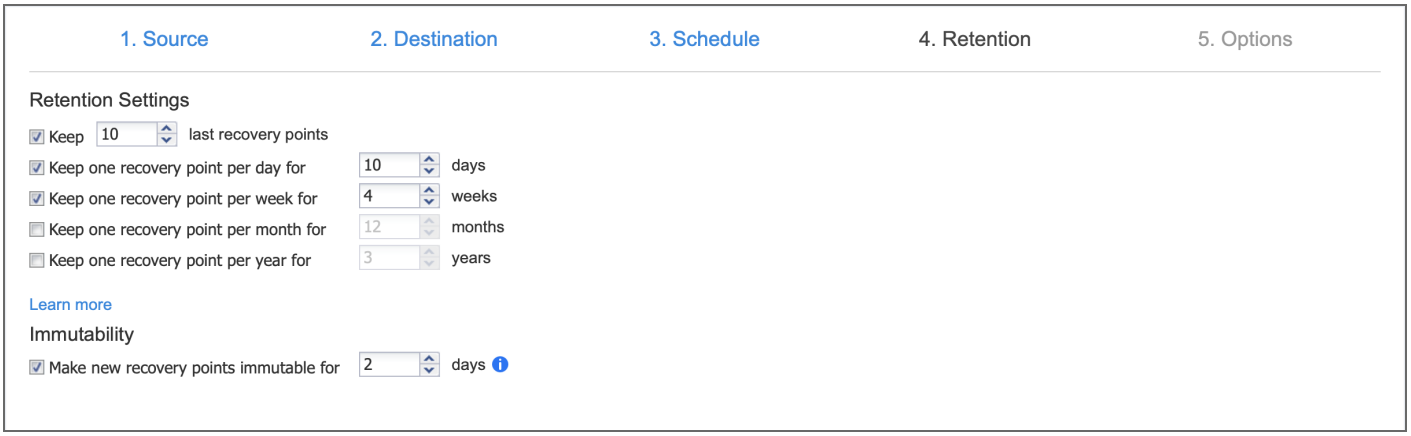
# Recovery Point Retention

After each job run, NAKIVO Backup & Replication creates a recovery point for each VM, object, instance, machine, or account in a **Backup Repository**. A recovery point represents the backup of the respective source as of a particular moment in time and allows you to recover individual files, application objects, or the entire VM/object/instance/machine/account from the **Backup Repository**.

## Legacy Retention Approach

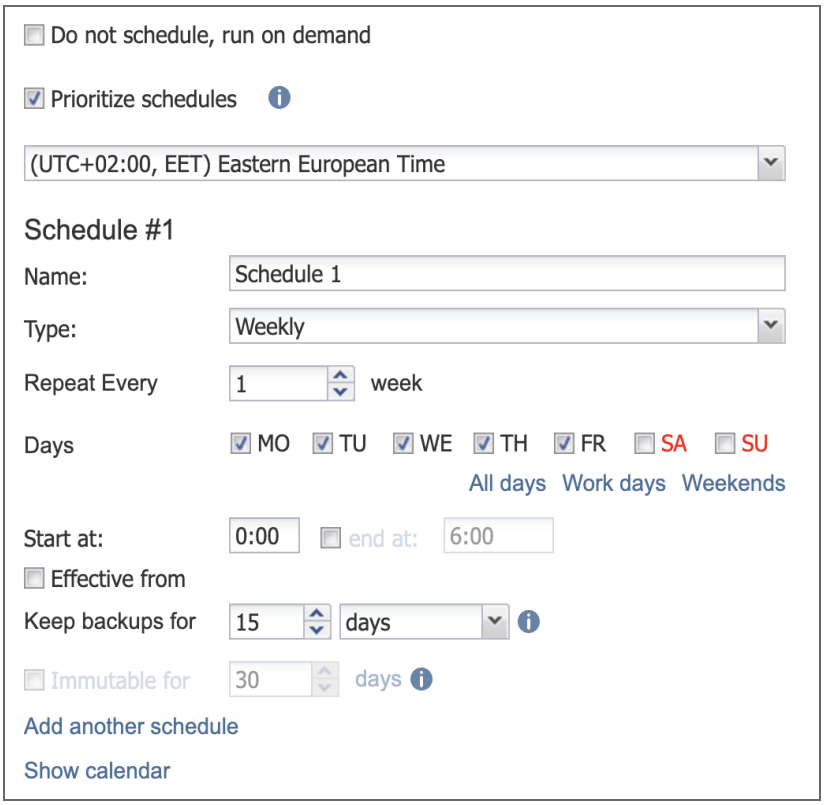
With the legacy retention method, NAKIVO Backup & Replication offers Grand-Father-Son (GFS) retention. This method allows you to save storage space while retaining the recovery points for any period that you need with the following options:

- Retain a specified number of last recovery points: after the specified number of recovery points in the backup repository is exceeded, the oldest recovery point is deleted.
- Retain one recovery point for a specified period of time: one recovery point is stored for the specified period of time, after which this recovery point is deleted.
- Make new recovery points immutable: this option sets an immutability flag on new recovery points, preventing their deletion or modification for a specified period of time.




## Schedule Retention Approach

With the schedule retention method, NAKIVO Backup & Replication allows you to set retention settings directly in the scheduling step of the job creation/editing process. This method allows you to set multiple schedules at chosen intervals for one job. These schedules can each be configured with their own retention settings with the "**Keep backups for**" option. This method is available for all backup and backup copy jobs with the exception of Oracle database backups.




With schedule retention settings, you can set up a clear recovery point retention policy for each job schedule and time interval. For example, if you set one job schedule to "**Keep backups for: 3 days**" and the job runs every weekday at noon, then a recovery point created with this schedule on a Monday will expire at noon on that Thursday. To ensure timely removal, NAKIVO Backup & Replication performs hourly status checks of all recovery points and deletes those that have expired.

**Notes**

Recovery points created with or migrated to the scheduler retention scheme are given expiration dates.

- To view the expiration dates and other details of recovery points created with this approach, refer to "Viewing Backup Repository Details" on page 687.
- To learn more about how expiration dates are assigned to recovery points, refer to this article in the Knowledge Base.

If NAKIVO Backup & Replication tries to remove the recovery point and fails, the following occurs:

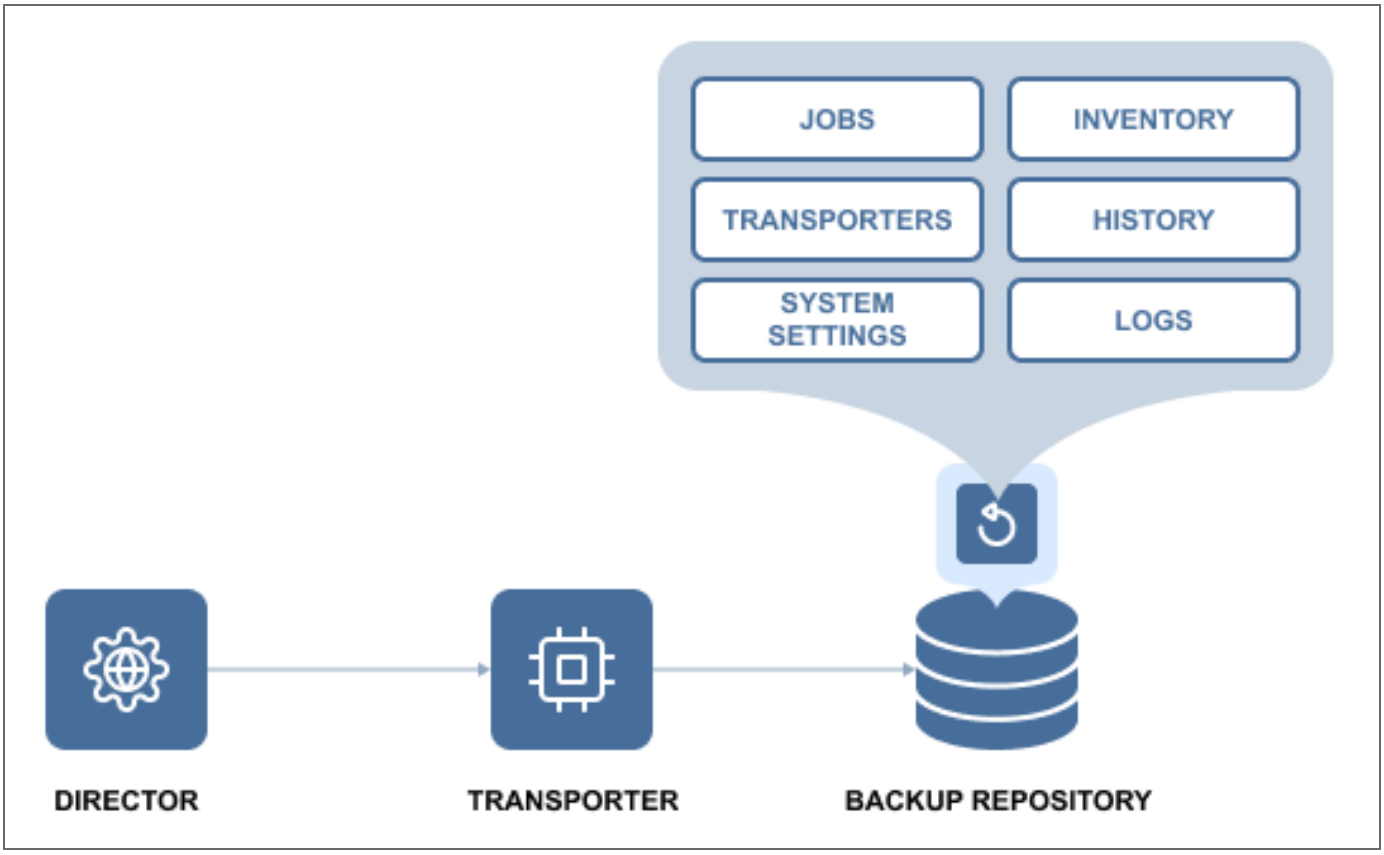
- If you made the removal attempt, the process can now be retried.
- If the removal attempt was made due to retention settings, the removal is retried depending on the retention approach.
- If you physically remove the recovery point from the folder, the database still keeps the record of the recovery point and marks it as missing in the UI. These records are removed later, depending on the retention approach.

# Self-Backup Feature

The Self-Backup feature provides automated protection of everything you have configured in NAKIVO Backup & Replication.

A truly complete data protection solution needs to back up not only your VMs, but also itself. There are good reasons for that. For example, the VM running the product may become corrupted, struck by a virus attack, or accidentally deleted. Regardless of the cause, you will need to restore the disrupted product as quickly as possible. Fortunately, a new instance of NAKIVO Backup & Replication can be installed in less than one minute. However, you will still need to restore the product configuration (such as jobs). Also, you do not want to lose the backup history. To save you time, NAKIVO Backup & Replication automatically backs up the entire configuration, including all jobs, inventory, information about connected **Transporters**, **Backup Repositories** and other.

The Self-Backup feature is enabled by default, and NAKIVO Backup & Replication sends daily self-backups to the first five backup repositories available in the product. Each self-backup is kept for five days, by default. Should you like to, you can fine-tune the backup targets, schedule, and retention policy.




If you accidentally make some undesired changes in the product, you can easily roll back to a previous system state from the backup. Migrating the system configuration to a new product instance is simple: just install a new copy of NAKIVO Backup & Replication, import a **Backup Repository** that contains a self-backup, and select a recovery point. The previous product configuration is restored along with all settings.

The Self-Backup feature saves you time and brings you peace of mind, ensuring reliable protection of everything you configure in NAKIVO Backup & Replication.

For information on the Self-Backup configuration, refer to "Self-Backup" on page 447.

# Two-Factor Authentication

NAKIVO Backup & Replication allows you to add an additional layer of security with two-factor authentication (2FA). By enabling 2FA, you add another step to the user login process to prevent malicious access to the solution and the organization's backup data. User authentication requires entering a code generated in one of the following ways:

- A code generated by the Google Authenticator mobile app
- A code sent to the specified email address
- One of the single-use backup codes

You can find more information in the following articles:


- "Logging in to NAKIVO Backup & Replication" on page 334
- "Configuring Two-Factor Authentication" on page 502

# VM Verification

VM verification is a process of checking the integrity of a backup or replica by booting a VM from a backup or starting a replica and interacting with it. With the VM verification feature, you have proof that your VM backups or replicas are usable, and can rest assured that your VMs can be recovered in case of disaster. VM backups and replicas can be corrupted or not bootable, even if the data protection software performed properly. The worst time to find out that your backup is bad is when your VM is down. If you don't have backup copies or VM replicas at an offsite location, you are left without any viable means of quickly restoring business processes.

VM verification involves the following entities:

- **Source Object**: Backup recovery point or replica recovery point which is used as a source of data for VM verification.
- **Target Object**: An entity that is subject to VM verification. It can be a replica or a temporary VM created via Flash boot.
- **Guest OS Agent**: An entity in the target object which allows remote interaction with the guest OS of this object (VMware Tools for VMware vSphere; Hyper-V integration services for Microsoft Hyper-V). Guest OS agent is required to be installed on the target object in order to perform VM verification.

There are two VM verification methods:

- **Boot Verification**: Verifying the target VM via starting target VM and checking whether hypervisor tools are running.
- **Screenshot Verification**: Verifying the target VM via starting the target VM and taking a screenshot of the VM screen.

To verify VMware and Hyper-V backups, NAKIVO Backup & Replication relies on the Flash boot feature. After a VM backup job has completed the data transfer, the product performs the following actions:

1. Instantly runs the VM from the newly created backup (with networking turned off).
2. Waits until the OS has booted.
3. Checks if guest OS agents are run successfully (if Boot Verification is selected).
4. Makes a screenshot of a running VM (If Screenshot Verification is selected).
5. Discards the test-recovered VM.

You can view the results of the verification procedure in the Dashboard or choose to receive an email report. VM verification, being an option for the jobs listed below, can be run on demand or scheduled to run automatically, saving you time and effort. VM verification option is available for the following jobs:

- VMware VM Backup Job
- VMware VM Replication Job
- VMware Flash boot Job
- Backup Copy Job

# Performance

A backup process can handle a huge amount of data, thus it is imperative to ensure that the data flow is efficient, and every resource used in the backup process is optimized. NAKIVO Backup & Replication provides the following techniques to increase performance:


- "Advanced Bandwidth Throttling" on page 68
- "Backup from Storage Snapshots" on page 69
- "Deduplication Appliance Support" on page 71
- "Full Synthetic Data Storage" on page 73
- "Incremental Jobs" on page 75
- "Jobs and Concurrent Tasks" on page 76
- "LAN-Free Data Transfer" on page 77
- "Network Acceleration" on page 80
- "Transporter Load Balancing" on page 81

# Advanced Bandwidth Throttling

NAKIVO Backup & Replication was designed to transfer data at the maximum available speeds for the purposes of completing VM backup, replication, and recovery jobs as quickly as possible. However, if you run data protection jobs during business hours, your LAN or WAN networks risk being overloaded. This can affect the performance of applications and degrade user experience (think of email messages taking too long to be sent, excessive load times for websites, etc.). NAKIVO Backup & Replication addresses this issue with the flexible Advanced Bandwidth Throttling feature. With Advanced Bandwidth Throttling, you can set limits for your data protection jobs and make sure they don't take more bandwidth than you can afford to allocate.

Advanced Bandwidth Throttling allows you to set global rules that limit the data transfer speeds of your backup processes. Such rules can apply to different jobs and on different schedules. For instance, you can create a global rule preventing your backup jobs from consuming more than 50 MByte/s during business hours, but leave the bandwidth unrestricted for Sunday backups. You can also create bandwidth throttling rules on a per-job basis, if you want to have more granular control over the whole process. Individual limits override global rules, sparing you the need to adjust the global rule for every job.

The Advanced Bandwidth Throttling feature of NAKIVO Backup & Replication is an effective means of optimizing backup operations and controlling your network traffic. With global and individual limits on data transfer speeds, the feature can help you ensure the performance of your business applications is never affected by backup workloads – even if you have little bandwidth to spare. With bandwidth rules, usage of LAN/WAN bandwidth by NAKIVO Backup & Replication jobs may be restricted to a specific amount. For more information, refer to the following sections:

- About Bandwidth Rules

## About Bandwidth Rules

A bandwidth rule specifies the bandwidth amount that can be used by one job, by multiple jobs, or by all applicable job. When a job containing multiple VMs starts running with a bandwidth rule active, the rule divides bandwidth between tasks.

Bandwidth rules are applicable to the following types of NAKIVO Backup & Replication jobs:

- Backup Job
- Backup Copy Job
- Replication Job (except for Amazon EC2)
- Recovery Job
- Replica Failback (except for Amazon EC2)

Bandwidth rules may be always active, active on schedule, or disabled. Refer to “Bandwidth Throttling” on page 407 for more details.

A bandwidth rule can be:

- **Global Rule** – a bandwidth rule applied to all applicable Jobs.
- **Per Job Rule** – a bandwidth rule only applied to specific Jobs.

**Per Job** rules have higher priority than **Global Rules**. A per job rule will be applied to the job when both the per job rule and a global rule are active for the same job.

Bandwidth rules (up to 100 rules can be created and enforced at the level of a Single-tenant product / Tenant of Multi-tenant product) are applied at the Transporter level, stored at the Director, and enforced while starting processing of a specific job object that falls into the limits of the current rules. Bandwidth rules can be enabled/disabled individually.

When enabled, the rule can limit bandwidth of JODTs that are covered by this rule.

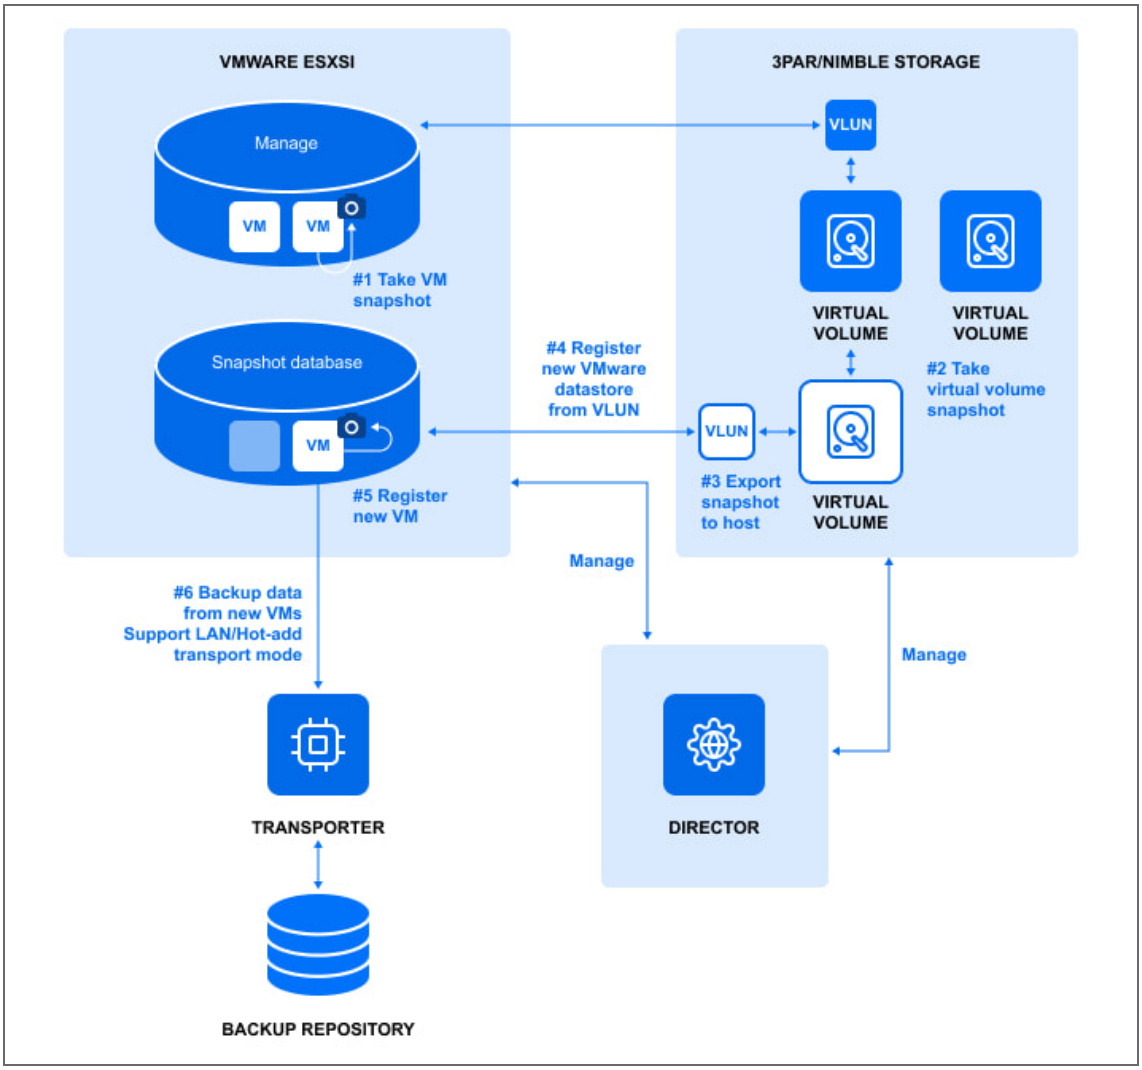
> **Note**
> Keep in mind that if the job starts when the Bandwidth throttling rule is active, it will stay at that speed even if the rule time is over.

# Backup from Storage Snapshots

NAKIVO Backup & Replication allows you to use HPE 3PAR, HPE Nimble, HPE Primera, HPE Alletra, or NetApp storage devices to back up from storage snapshots. This backup approach offers advantages when your VMs are processing large amounts of data while being backed up. Backup from storage snapshots includes the following steps:

1. A VM snapshot is created.
2. Soon after, a storage snapshot is created.
3. The VM snapshot is removed.

The VM snapshot only exists for a short period of time because the storage snapshot takes a small amount of time to be created, and it contains all required data for performing VM backup (delta and CBT data). A storage snapshot can be created within seconds, and it consumes a small amount of space which, in turn, reduces the impact of backup activities on the production environment and improves RPOs.




To back up from HPE 3PAR/Nimble/Primera/Alletra storage snapshots, do the following:

1. Add the HPE 3PAR, HPE Nimble, HPE Primera, HPE Alletra or NetApp storage device to the Inventory.
2. During VMware Backup Job creation, add VMs that have their disks residing on the supported HPE 3PAR/Nimble/Primera/HPE Alletra/NetApp storage devices.
3. Enable **Backup from storage snapshot** on the **Options** page of the backup job wizard.
4. Complete the wizard and run the job.

# Deduplication Appliance Support

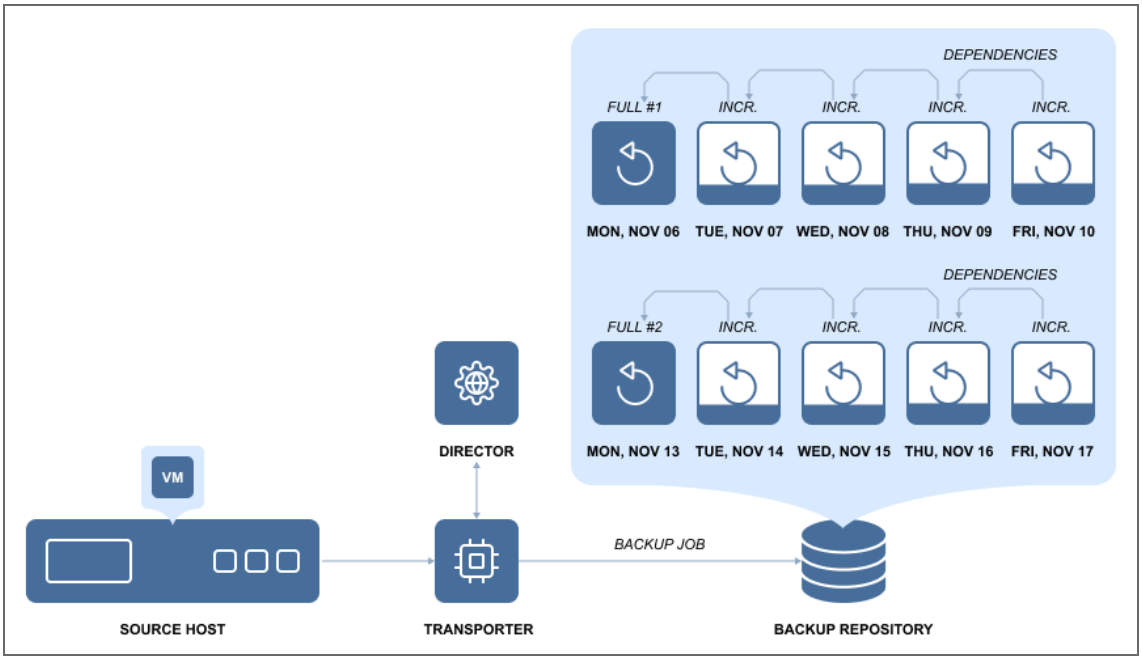
Deduplication appliances are solutions that implement specialized data reduction techniques to eliminate duplicate copies of repeated data. Deduplication appliances are leveraged across a range of data protection solutions, regardless of whether network-attached storage, disk, and/or tape is used. The biggest advantage of deduplication appliances is their ability to reduce datastore space used – sometimes by ratios of 20:1 or more.

NAKIVO Backup & Replication supports integration with deduplication appliances. For details, refer to the following sections:

- NAKIVO Optimization for Deduplication Appliances
- Deduplication Appliance Configuration Details

## NAKIVO Optimization for Deduplication Appliances

NAKIVO Backup & Replication provides a special type of Backup Repository (stream repository) optimized for high performance with deduplication appliances. With this type of Backup Repository, NAKIVO Backup & Replication supports virtually any type of deduplication appliance as a primary or a secondary backup destination. The architecture of such Backup Repository is based on sequential block write operations through a restricted number of data streams and storing backup blocks in dedicated data files. Data blocks are stored in incremental backup files and full backup files. This means that the repository stores VM backup chains consisting of periodic full backups and several increments between these full backups.




In terms of integration with deduplication appliances, a stream repository:

- Creates fewer data streams in read/write operations during VM backup and recovery.
- Does not leverage the global data deduplication feature of NAKIVO Backup & Replication.

## Deduplication Appliance Configuration Details

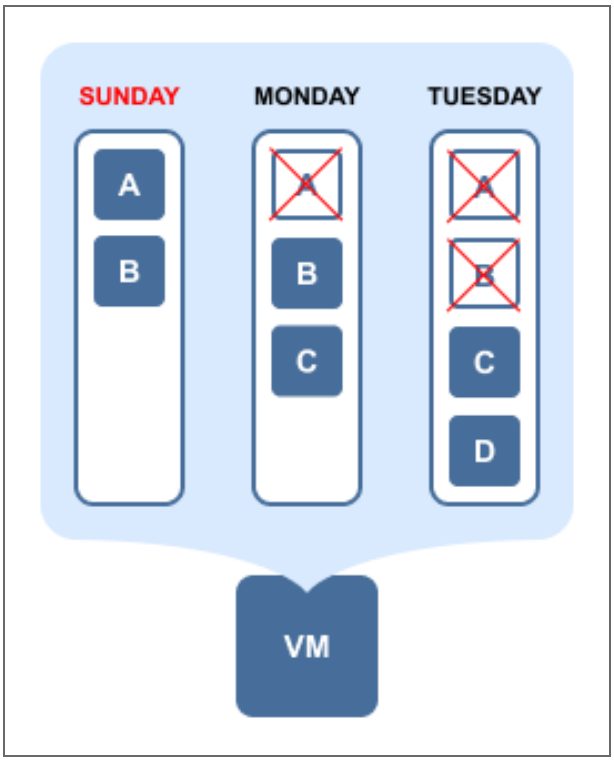
When a Backup Repository is created on a deduplication appliance, NAKIVO's built-in data deduplication functionality is disabled. Additionally, the incremental-with-full-backups option is enabled by default. This configuration ensures that no extra resources are spent for double deduplication and reclaiming repository space is not required.

# Full Synthetic Data Storage

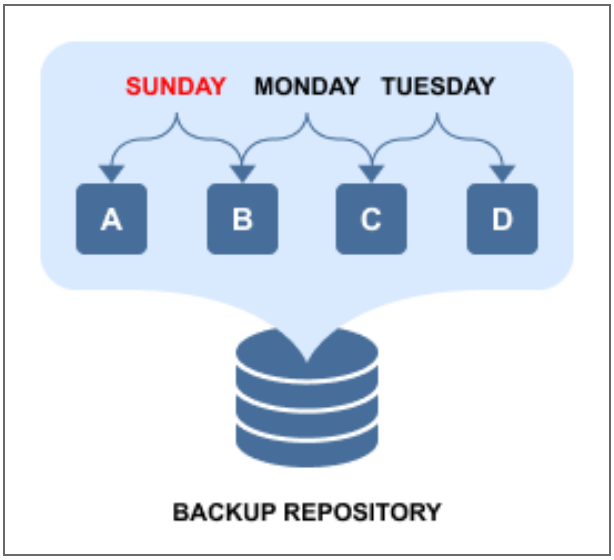
With **forever incremental** (**Store backups in separate files** option is not selected) Backup Repositories, NAKIVO Backup & Replication uses the full synthetic mode to store backups: all unique data blocks are stored in a single pool, while recovery points serve as references to the data blocks that are required to reconstruct a machine at a particular moment in time.

**Example**

You run the first backup of a VM on Sunday. For the sake of simplicity, let's say that the VM consists only of 2 data blocks: A and B. Then on Monday, you run an incremental backup, which finds that the block A has been deleted, but a new block C has been added. Then on Tuesday, the incremental backup finds that the block B has been deleted and a new block D has been added. Here's how the VM would look like during the three days:




And here's how the data will be stored in the **forever incremental** (**Store backups in separate files** option is not selected) **Backup Repository** if the job is set to keep 3 or more recovery points:




As you can see from above, each unique data block is stored only once to save space, while recovery points are just references to data blocks that are required to reconstruct the VM as of a particular moment in time. If, for example, you delete Monday's recovery point, then no actual data removal will occur, as its data blocks (B and C) are required for recovery points of Sunday and Tuesday. If, on the other hand, you change the recovery point retention policy to keep only the last two recovery points (Mon and Tues in our case), then only block A will be deleted, as it's not being used anywhere else.

The full synthetic data storage approach provides a number of benefits:

- **Smaller backups**: Unique data blocks are stored only once and can be referenced by multiple recovery points, as opposed to storing the same data again in different increments.
- **Faster backups**: There is no need to run full backups periodically or transform legacy increments into virtual full backups, as each recovery point already "knows" which data blocks should be used to reconstruct an entire machine.
- **Safer backups**: With a legacy incremental backup approach, losing one increment in a chain means losing the entire chain of recovery points after that increment. With NAKIVO Backup & Replication losing a data block or an increment (such as A or B in the example above) can still leave you with recoverable increments.
- **Faster recovery**: A legacy incremental backup consists of a chain of increments that you must apply one by one in order to get to a particular machine state. With NAKIVO Backup & Replication, each recovery point already "knows" which data blocks should be used to reconstruct an entire machine.

# Incremental Jobs

NAKIVO Backup & Replication allows you to create in incremental backup and replication jobs. For more information refer to:

- Backup Jobs
- Replication Jobs

## Backup Jobs

- When a **forever incremental** (**Store backups in separate files** option is not selected) Backup Repository is utilized as a destination, the full backup will be performed only on the first backup job run. All consequent job runs will send only changed data (increments) to the Backup Repository. This approach reduces backup time and network load. For example, if NAKIVO Backup & Replication determines that the amount of data that has been changed on a 100 GB VM is just 1 MB, only 1 MB of data will be transferred to the Backup Repository, but the created recovery point will reference all data blocks (from previous job runs) which are required to restore the entire 100 GB VM. With this approach, each recovery point "knows" all data blocks that are needed for recovery, so there is no need to apply increments one by one to get to a particular point or periodically transform backed up data blocks.
- When an **incremental with full backups** (**Store backups in separate files** option is selected) Backup Repository is utilized as a destination, NAKIVO Backup & Replication performs a full backup on the first backup job run. Consequently, NAKIVO Backup & Replication runs incremental backups and periodically creates full backups according to the specified settings. Every VM backed up to said Backup Repository will produce full backup files and incremental backup files.

## Replication Jobs

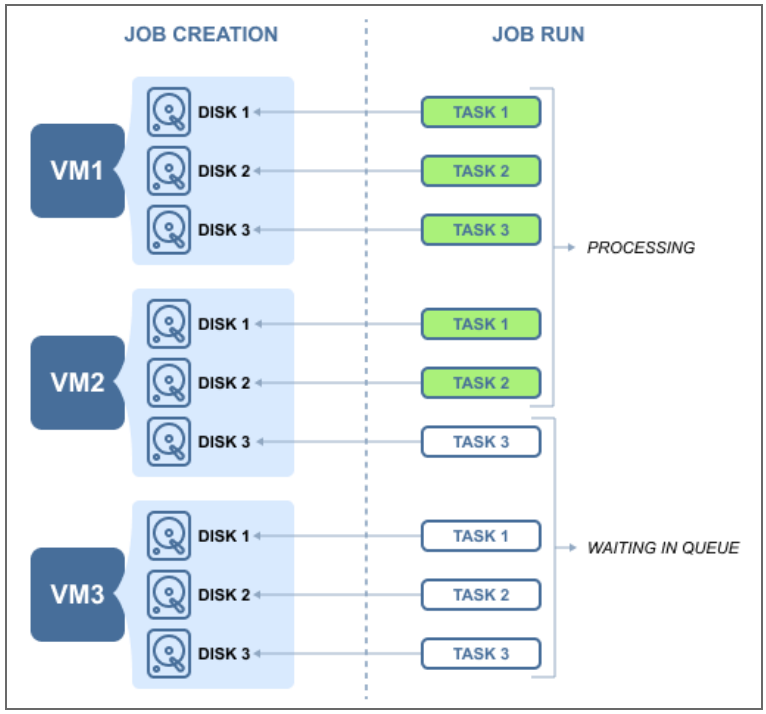
Replication jobs in NAKIVO Backup & Replication are forever incremental. This means that after the initial full replication, all subsequent job runs will send only changed data (increments) to the replica. This approach reduces the replication time and network load. For example, if NAKIVO Backup & Replication determines that the amount of changed data on a 100 GB VM constitutes just 1 MB, only 1 MB of data will be transferred to the replica VM.

# Jobs and Concurrent Tasks

Job is a data protection activity that is performed by NAKIVO Backup & Replication in accordance with a distinct configuration. These are the main types of NAKIVO Backup & Replication jobs:

- Backup jobs
- Replication jobs
- Recovery Jobs

In NAKIVO Backup & Replication, a job can have one or more job objects to process. Depending on your preferences, job objects may be reordered for processing within a job. See the example below.




Each job object may consist of one or more machine disks, Oracle databases, Exchange Online mailboxes, OneDrive for Business instances or SharePoint Online sites that have to be processed within a job run. Data processing that is related to a specific VM disk or service constitutes a single task, in the scope of the corresponding job. Such tasks are processed by a Transporter. For the sake of managing the load over the infrastructure, any Transporter is configured to process a limited number of concurrent tasks. When a task is processed, the Transporter starts processing another task if available. A task can be one disk, file or recovery session, Oracle database, Exchange Online mailbox, OneDrive instance, or a SharePoint Online site. By default, NAKIVO Backup & Replication is set to process 6 concurrent tasks per one Transporter. Refer to “Editing Nodes” on page 593 to learn how to change the Transporter maximum load.

# LAN-Free Data Transfer

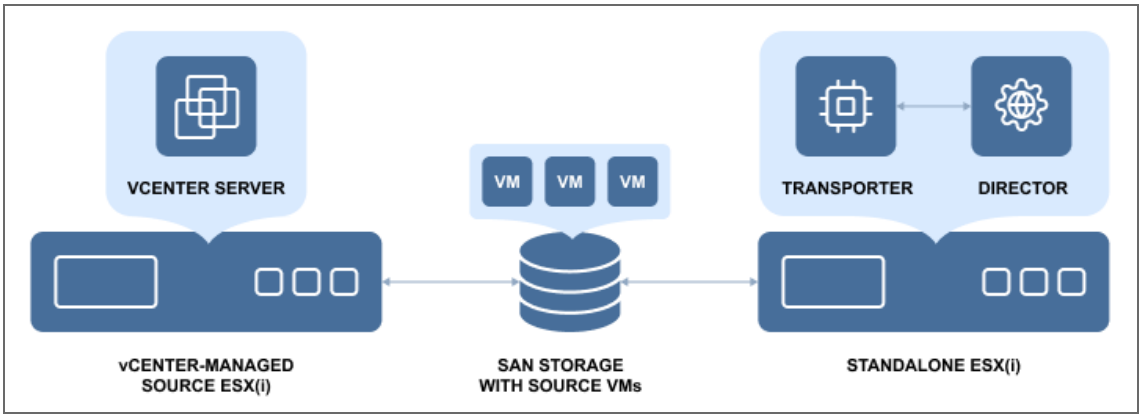
NAKIVO Backup & Replication automatically uses a LAN-free data transfer mode for VMware backup, Hyper-V backup, Nutanix AHV backup, VMware replication, and Hyper-V replication. The LAN-free data transfer mode boosts VM backup and replication speed in addition to reducing the load on your network.

On the VMware platform, this is achieved with the Hot Add and Direct SAN Access features.


- "Direct SAN Access for VMware" on page 78
- "Hot Add for VMware" on page 79

# Direct SAN Access for VMware

If your VMs are located on a Fiber Channel or iSCSI Storage Area Network (SAN) device, NAKIVO Backup & Replication can use direct SAN access for data retrieval. Using this storage access mode can significantly increase the speed of backup and replication while decreasing the load on your production network.

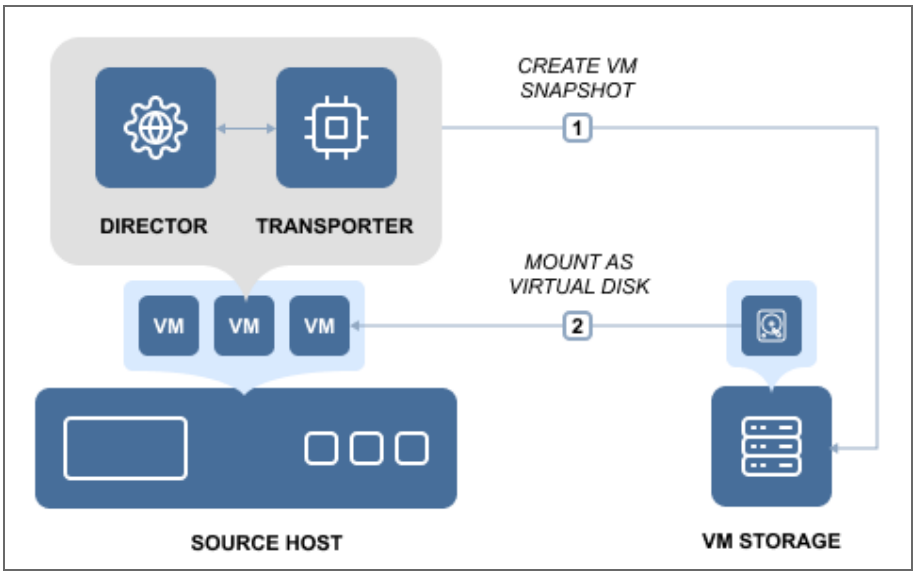



NAKIVO Backup & Replication relies on the VMware VDDK library that provides SAN support and does not make any write/modification operations on the SAN LUNs. That is to say, the product will not compromise data integrity or interfere with the data of running VMs on your SAN datastores.

For information about Transporter deployment requirements as well as recommendations on setting up the SAN access, refer to Transporter Deployment for SAN Access.

# Hot Add for VMware

The Hot Add Data Transfer mode significantly improves VM backup and replication speed and reduces the load on the network. NAKIVO Backup & Replication can read data directly from VM datastores, bypassing the host's TCP/IP stack that would otherwise impact every VM on the host, and slow down the data transfer. NAKIVO Backup & Replication can mount (Hot Add) VM snapshots, and read VM data directly from VM datastores through the host's storage I/O stack.




By default, NAKIVO Backup & Replication will automatically attempt to use the Hot Add mode for VM backup and replication jobs. Please check the appropriate feature requirements section for prerequisites and limitations.

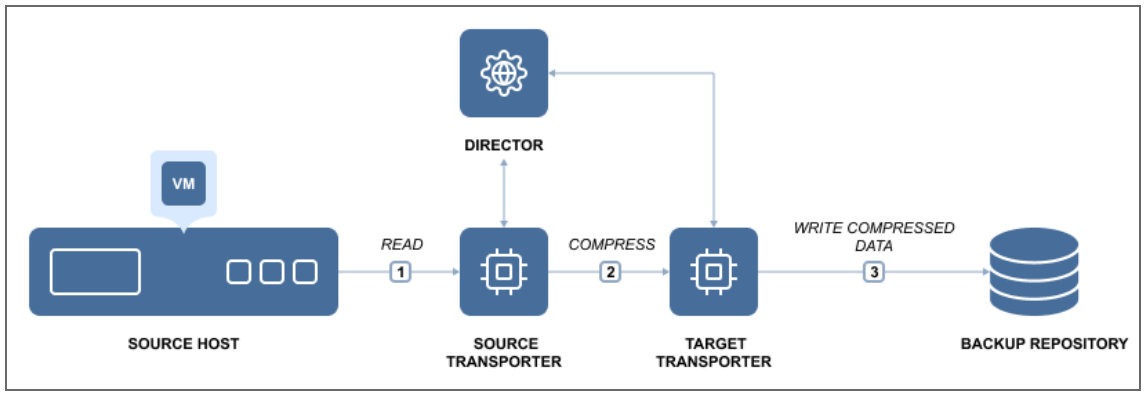
# Network Acceleration

Whether you run VM backup and replication jobs during business hours or send VM backups and replicas offsite over the Internet, saving network bandwidth is of the essence. NAKIVO Backup & Replication provides the Network Acceleration feature to speed up VM backup and replication jobs, shorten backup windows, and reduce network load at the same time. With network acceleration enabled, you can increase VM backup, replication, and recovery speed by 2X in WAN and busy LAN networks.

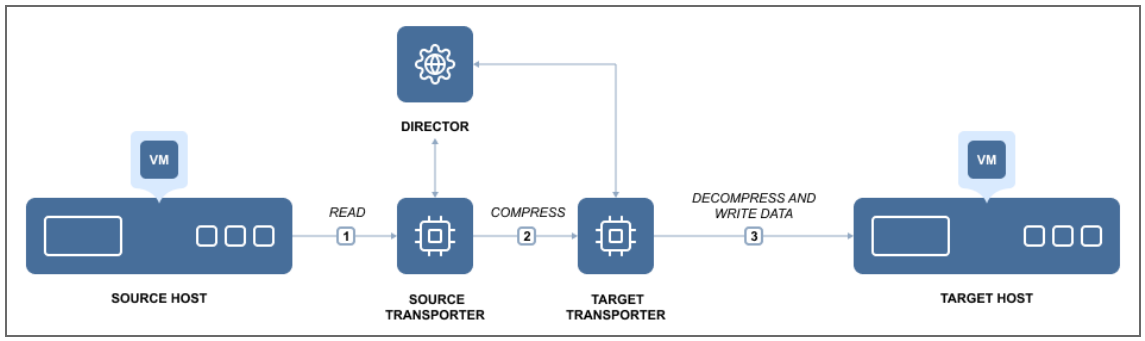
Network acceleration is achieved by the use of two instances of Transporter. Transporter is the product component that performs all data protection and recovery tasks, such as backup, replication, recovery, encryption, and so on. To simplify deployment and configuration, one instance of Transporter is automatically installed with NAKIVO Backup & Replication.

To enable Network Acceleration, you just need to install another Transporter instance locally or offsite and then enable Network acceleration in your job. When the job is executed, the source Transporter will read the data, compress and optimize it, and then send the data to the target Transporter. By using Network Acceleration, you can reduce the amount of data that is transferred over the network, which also means that your jobs will complete faster.

**Network Acceleration for Backup**




**Network Acceleration for Replication**



# Transporter Load Balancing

The **Transporter load balancing** feature allows you to distribute the tasks within a single job among available transporters to maximize data processing speed. The product achieves this by:

- Determining the priority of tasks to be processed by a transporter
- Determining possible data paths for processing each task, that is:
    - Determining the transporters available for a job
    - Measuring transporter proximity
    - Verifying cross-transporter connectivity
    - Determining available transporter load while accounting for a transporter's existing load, which can include maintenance tasks like repository verification
- Determining a candidate data path for processing the task
- Prioritizing the candidate data path based on transporter round-trip time (RTT) and available load
- Distributing tasks across the candidate data path to evenly load transporters based on task and candidate data path priorities

# Administration

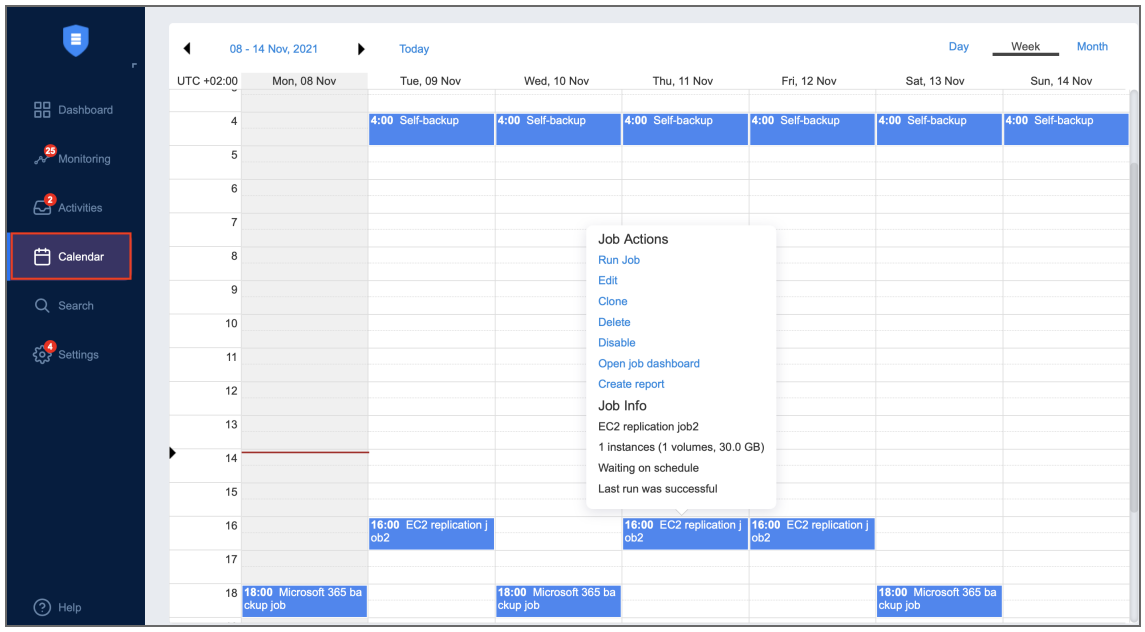
NAKIVO Backup & Replication strives to make the user experience as intuitive and easy-to-use as possible, and provides users with the following features:


- “Calendar” on page 83
- “File System Indexing” on page 83
- “Global Search” on page 85
- “Monitoring” on page 85
- “Policy-Based Data Protection” on page 87

# Calendar

Backing up VMs is a resource-intensive process, which places extra load on your infrastructure, be it VMware, Hyper-V, or AWS. This is particularly noticeable when it comes to large environments with thousands of VMs. Too many backup jobs running concurrently on the same host or on the same network may affect the performance of your virtual environment and slow down your VMs. To reduce the load on your resources, you need to carefully schedule and structure your backup jobs, to ensure the shortest backup windows possible.

Scheduling data protection jobs may be tricky in large virtual environments, where you need to fit multiple jobs into a backup window and avoid possible overlaps. To resolve this issue, NAKIVO Backup & Replication features the **Calendar** dashboard, which is aimed at greatly improving job scheduling. The **Calendar** dashboard displays all your jobs in the calendar view, the time it took different jobs to run in the past, and the predicted job duration in the future. Here, you can get a bird's eye view of all your jobs, and you can easily find open time slots for new jobs, which you can create right in the dashboard. You can also visit past jobs to view the status and details of the jobs that have been completed and drill down to their details. The **Calendar** dashboard has an intuitive interface and navigation, similar to those of the most popular calendar applications.




# File System Indexing

With **File System Indexing** enabled, each backup job run saves an index of files and folders from the source machine and maps it to the corresponding recovery point. This allows you to search for specific files and folders inside a backup using the **Global Search** dashboard.

The index of supported Windows/Linux operating systems only includes files, links, and directories; files and folders in NFS/CIFS shared folders, as well as removable disks and product component installation folders are excluded from indexing.

For more information about indexing file system files and folders, refer to the following topics:

- File System Indexing for VMware vSphere VM backup jobs
- “Searching Indexed Files” on page 402

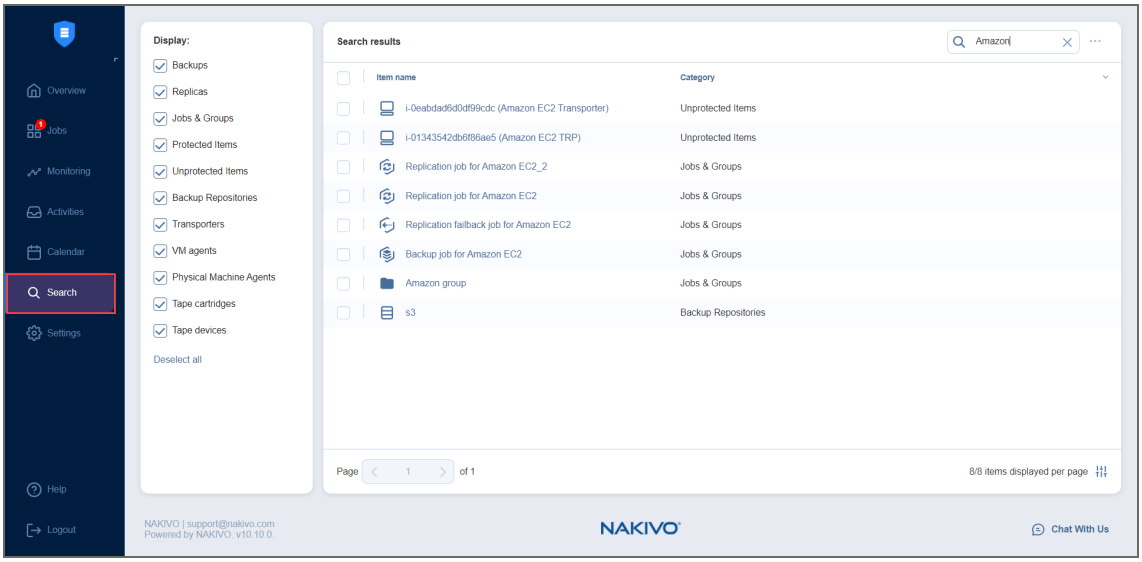
# Global Search

NAKIVO Backup & Replication includes the powerful global search feature that allows you to find any item quickly by entering the name of the item (or part of the name) into the search box. You can refine the search results by using filters (for example, choose to view only VM backups). In addition, you can select items in the search results and instantly perform mass actions on them, such as creating a new job for unprotected VMs or adding items to an existing job. The ability to perform such actions simplifies the management of your backup infrastructure.

With the global search feature, you can:

- **Search**: Instantly search for VMs, backups, replicas, jobs, groups, Backup Repositories, Transporters, tape cartridges, and tape devices.
- **Filter**: Choose to view a subset of results – for example, unprotected VMs only.
- **Get information**: View item details, such as size, host, datastores, networks, and protection status.
- **Act**: After finding what you were looking for, you can take an action – add multiple unprotected VMs to a job, start a recovery, run a job, etc.

The global search feature in NAKIVO Backup & Replication is an easy-to-use tool that helps you manage large backup infrastructures and saves you time.




# Monitoring

With NAKIVO Backup & Replication, you can monitor your VM infrastructure and its performance, create and configure reports and alert templates, and view triggered alerts.

Refer to the following topics to learn how to use **Monitoring**:

- “Metrics” on page 1206
- “Alerts” on page 1211

- “Reports” on page 1220

# Policy-Based Data Protection

Policy-Based Data Protection relieves you of the need to chase new VMs or changes in your infrastructure. Once a policy is created, all the matching VMs are protected automatically. Whenever a VM's status changes, the policy recognizes this change and excludes or adds the VM to jobs accordingly. The feature is designed to reduce complexity and add more flexibility to data protection processes such as backup, replication, or backup copy. You can set rules based on the VM name, tag, size, location, VM configuration, power state, or any combination of these parameters. A newly-created VM or instance is automatically added to data protection jobs if they match your policy rules; you don't have to keep track of all the changes in your infrastructure or manually manage data protection for new VMs. You can add as many new VMs and instances as you need because NAKIVO Backup & Replication can automatically protect all of them for you, as long as you have policy-based jobs in place.

This functionality can be a great time-saver if your virtualized infrastructure is actively expanding, includes numerous VMs and instances, or has a complex multilayer architecture. The Policy-Based Data Protection feature contributes greatly to the overall usability of NAKIVO Backup & Replication, making it an even more efficient data protection tool. Policies can be created for VM backup, replication, and backup copy jobs in just a few steps. Simply select the criteria (e.g., a VM's name, size, tag, etc.), enter the necessary search parameters, and have all the matching items included in the job automatically. For instance, you can choose to back up all VMware VMs tagged "Accounting" which exceed 100 GB in size and have more than 2 GB of allocated RAM. Once the policy has been created, NAKIVO Backup & Replication recognizes newly added VMs or instances with the same characteristics and automatically includes them into the existing job.

Refer to the following topics to know how to use the feature:


- “Managing Job Policies” on page 382
- “Managing Policy Rules” on page 385

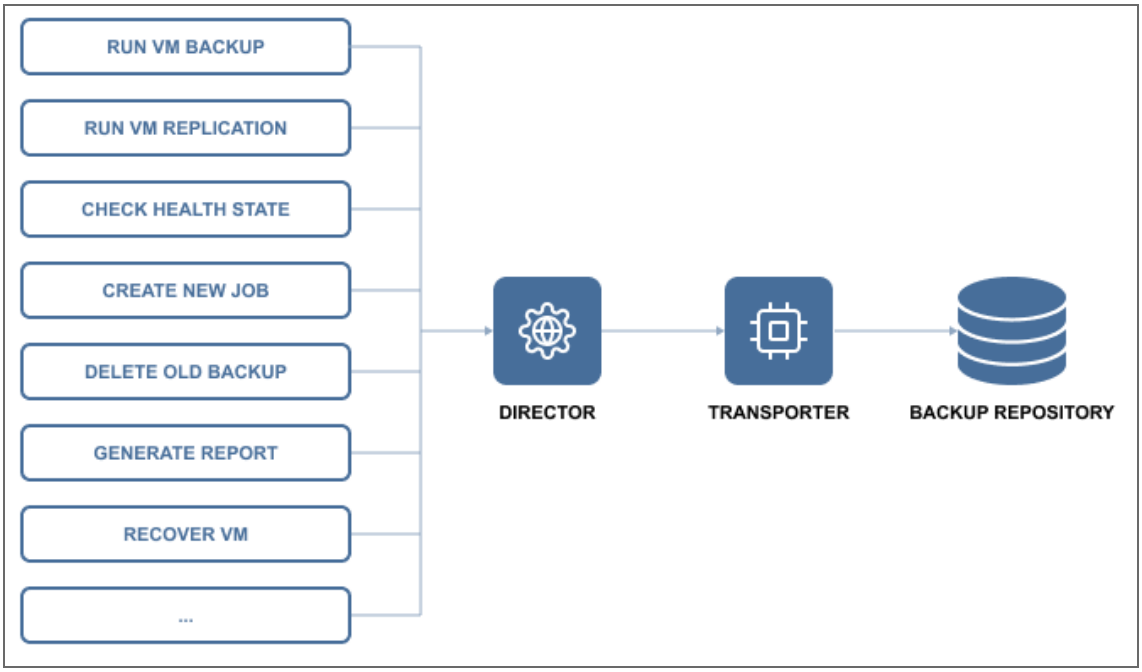
# Automation

The following features help users eliminate repetitive routine work and automate their activities:


- "HTTP APIs" on page 89
- "Job Chaining" on page 90
- "Pre and Post Job Scripts" on page 91

# HTTP APIs

NAKIVO Backup & Replication provides a simple HTTP API that lets you automate and orchestrate VM backup, replication, and recovery tasks. The API provides complete coverage of the product features, that is, you can use the API to perform all tasks that are available in the product's Web interface.



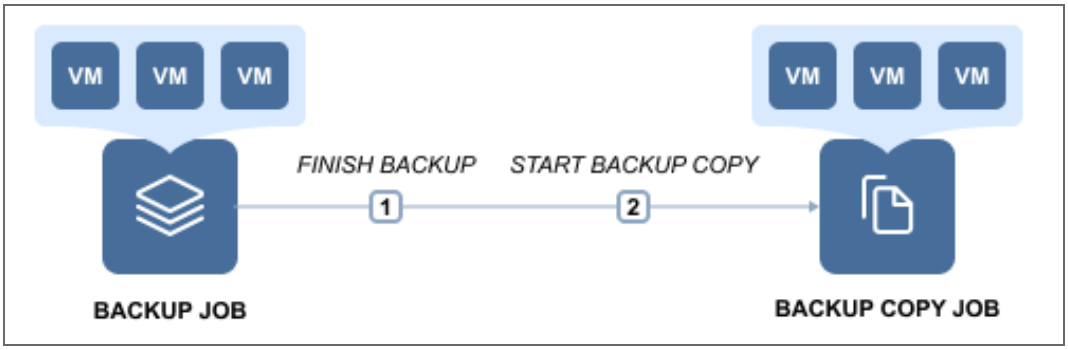

The API allows you to easily integrate NAKIVO Backup & Replication with monitoring, automation, and orchestration solutions to reduce time spent on backup management and reduce data protection costs. To speed up integration time, the API comes as part of an Integration Kit, which includes API documentation and code examples.

By using the API, you can:

- Save time on backup administration by automating the data protection process from VM provisioning to VM decommissioning.
- Ensure an uninterrupted backup process by monitoring the health status of the product components.
- Prevent failed jobs and out of space errors by monitoring backup repositories.
- Reduce storage space by automating backup decommissioning.
- Improve compliance by automating data protection reporting.
- Align data protection with your business processes by triggering VM backup and replication jobs with your orchestration and automation tools.
- Increase recovery speed by automating recovery.

# Job Chaining

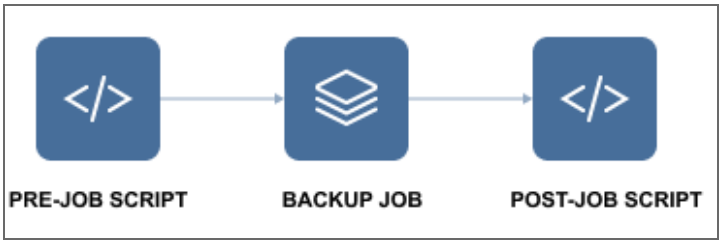
Job Chaining allows you to link jobs so that they run one directly after another. For example, you can set up a VM backup job that saves backups locally and then starts a Backup Copy job that copies the newly created backups to Amazon cloud.




You can link any type of jobs together – backup, backup copy, replication and recovery – and add any number of jobs to the chain. For instance, you can set up a series of backup jobs that trigger one another in the order of priority, or set up a series of Backup Copy jobs, which first send weekly backups to a DR repository and then send monthly backups to Amazon cloud for archiving.

# Pre and Post Job Scripts

NAKIVO Backup & Replication provides you with the ability to run a script before a job begins (a pre-job script) and after the job has been completed (a post-job script).




By running your pre- and post- job scripts, you can do just about anything: start custom pre-freeze and post-thaw scripts on Linux systems to create application-aware backups and replicas, wake servers, establish connections, mount volumes, start and stop services, send commands to 3rd-party reporting, monitoring and automation tools, and etc.

# Integration

NAKIVO Backup & Replication provides support for enterprise-grade deduplication appliances, such as EMC Data Domain and NEC HYDRAstor. Deduplication appliances are servers designed to reduce data size, and can be used as backup targets. Deduplication appliances operate best with sequential large block I/O from backup software. Therefore, when backing up your VMs to a deduplication appliance, it is important to make sure that the architecture of your Backup Repository is optimized for these devices and your VM backups have a large block I/O. Only by doing this, you will be able to maximize your VM backup speeds.

NAKIVO Backup & Replication offers you two different types of backup repositories to choose from:

- The regular Backup Repository, which is optimized for generic storage systems and performs **forever-incremental** (when the **Store backups in separate files** option is not enabled) VM backups along with global data deduplication and compression.
- The special Backup Repository with an architecture optimized for efficient operation on deduplication appliances. This is known as **Incremental-with-full-backups** (when the **Store backups in separate files** option is enabled during the Backup Repository creation process). The repository performs incremental-with-full VM backups, and proprietary VM backup deduplication and compression by NAKIVO Backup & Replication are turned off. The file structure is also improved, and each backup, along with its recovery points, is stored in a separate folder for easier manageability.

When tested in a customer environment on a high-end NEC HYDRAstor deduplication appliance, the product's special Backup Repository demonstrated a 53X boost in backup speed over the regular Backup Repository. NAKIVO Backup & Replication backed up the customer's VMs at an incredible 3.2 GByte/s. NAKIVO Backup & Replication ensures that you can use existing storage hardware while achieving top VM backup performance.

Integration with the following solutions allow NAKIVO Backup & Replication to further increase backup speed and save storage space:


- “Active Directory” on page 93
- “EMC DD Boost” on page 94
- “HPE 3PAR Storage” on page 95
- “HPE Nimble Storage” on page 95
- “HPE Primera Storage” on page 95
- “HPE Alletra Storage” on page 96
- “HPE StoreOnce Catalyst” on page 97
- “NEC HYDRAstor” on page 98
- “NetApp” on page 98

# Active Directory

Microsoft Active Directory is a leading directory service, which provides you with the ability to authenticate and authorize users and computers in a Windows domain type network. To simplify user management, NAKIVO Backup & Replication provides integration with Microsoft Active Directory. You can easily map Active Directory groups to NAKIVO Backup & Replication user roles, which will allow domain users to log in to NAKIVO Backup & Replication with their domain credentials. With this feature, you can align NAKIVO Backup & Replication with your company's security policy and seamlessly provide Admin and Guest access to NAKIVO Backup & Replication.

For more information, refer to the following topics:


- “Configuring Active Directory Integration” on page 480
- “Managing Active Directory Users” on page 475

# EMC DD Boost

The Dell/EMC Data Domain Boost technology allows for the reduction of storage consumption by up to 17X, greatly accelerating the VM backup process. The aggregate quantity of business data produced has drastically increased in recent years, which results in two major problems for modern companies. The first is the amount of storage space that backups occupy, and the second is the significant load on the production network created by backup operations, especially if they are run during business hours.

NAKIVO Backup & Replication and Dell/EMC Data Domain Boost offer a combined solution for both of these challenges. By using NAKIVO Backup & Replication along with source-side deduplication of Dell/EMC Data Domain Boost, you can perform VM backups 50% faster while reducing the size of your backups by up to 94%. This means that you can offload your network and save storage space at the same time.

For more information about the integration of NAKIVO Backup & Replication with EMC DD Boost, refer to the following articles:


- "Storage Integration Requirements" on page 135
- Integrating with EMC DD Boost
- "Backup Repository on Deduplication Appliance" on page 657

# HPE 3PAR Storage

HPE 3PAR StoreServ storage is built to meet the requirements of consolidated cloud service providers. HPE 3PAR can easily handle large workloads and guarantees 99.9999% data availability. HPE 3PAR StoreServ storage uses massively scalable and flash-optimized Tier-1 architecture, which allows for agile and efficient responses. HPE 3PAR enables you to ensure service levels with QoS optimization tools and consistent, sub-millisecond latency.

When using HPE 3PAR StoreServ Storage with NAKIVO Backup & Replication, you can use the backup from storage snapshots approach and significantly reduce the load on your production environment.

For details about integration of HPE 3PAR devices with NAKIVO Backup & Replication, refer to "Backup from Storage Snapshots" on page 69.

# HPE Nimble Storage

HPE Nimble Storage is designed to meet the demand for cloud storage and combat data disruption. Nimble Storage can easily handle large workloads with a guaranteed 99.9999% data availability and lower operational storage expenses. The Nimble Storage model can analyze millions of sensors per second to effectively prevent and predict data problems. This enables you to ensure service levels with QoS optimization tools and data mobility across clouds.

When using HPE Nimble Storage with NAKIVO Backup & Replication, you can use the backup from storage snapshots approach. This approach allows you to significantly reduce the load of data protection activities on your production environment.

For details about the integration of HPE Nimble Storage devices with NAKIVO Backup & Replication, refer to "Backup from Storage Snapshots" on page 69.

# HPE Primera Storage

HPE Primera Storage combines the agility of cloud computing with the formidable intelligence of HPE InfoSight, providing instant data accessibility, rapid setup within minutes, transparent upgrades, and service-based delivery. What's more, it comes with a robust 100% availability guarantee, assuring continuous operations and ensuring that all your mission-critical applications run smoothly.

When using HPE Primera Storage with NAKIVO Backup & Replication, you can use the backup from storage snapshots approach. This approach allows you to significantly reduce the load of data protection activities on your production environment.

For details about the integration of HPE Primera Storage devices with NAKIVO Backup & Replication, refer to "Backup from Storage Snapshots" on page 69.

# HPE Alletra Storage

HPE Alletra Storage offers a robust and adaptable storage solution tailored to elevate your enterprise's storage infrastructure to unprecedented levels of efficiency, agility, and performance.

Powered by HPE Alletra AI, this storage solution delivers exceptional agility, ensuring seamless adaptation to the changing demands of dynamic workloads without compromising on performance. Its flexible architecture allows for easy scalability and customization, enabling your storage infrastructure to grow in tandem with your business needs.

When using HPE Alletra Storage with NAKIVO Backup & Replication, you are free to use the backup from storage snapshots approach significantly reducing the load on your production environment.

For details about integration of HPE Alletra storage devices with NAKIVO Backup & Replication, refer to “Backup from Storage Snapshots” on page 69.

# HPE StoreOnce Catalyst

HPE StoreOnce Systems from Hewlett Packard Enterprise provide a disk-based data protection platform. This platform addresses data growth by applying HPE StoreOnce deduplication software for efficient and long-term backup data retention. HPE StoreOnce Catalyst, a data protection protocol optimized for disk-based data protection, is the most efficient way to transfer data to a StoreOnce System. When using HPE StoreOnce Catalyst for your Backup Repository, you get the following advantages:

- Reduction in network bandwidth as only unique chunks of data are transferred
- Lower physical storage space requirements with data deduplication
- Better backup copy job performance between HPE StoreOnce storage devices.

Starting from NAKIVO Backup & Replication version 10.1, you can create a Backup Repository on a StoreOnce appliance with HPE StoreOnce Catalyst support. Refer to the following topics for details:


- “Deduplication Appliance Support” on page 71
- “Storage Integration Requirements” on page 135
- “Backup Repository on Deduplication Appliance” on page 657

# NEC HYDRAstor

HYDRAstor is an award-winning product developed by the NEC Corporation. It is a disk-based grid storage platform offering long-term data retention through its maximized capacity of legacy storage solutions and scalability of performance. A HYDRAstor storage system can be composed of multiple nodes – from one to over 100. Each node consists of standard hardware including disk drives, memory, CPU, and network interfaces. The system is integrated with the HYDRAstor software, thus creating a single storage pool. The software incorporates multiple features of distributed storage systems. The features include content-addressable storage, variable block size, inline global data deduplication, erasure codes, data encryption, Rabin fingerprinting, and load balancing.

HYDRAstor can be scaled from one node to 165 in a multi-rack grid appliance. Its bandwidth and capacity can be scaled separately by using different types of nodes:

- **Hybrid nodes**: add both performance and capacity.
- **Storage nodes**: add capacity.

HYDRAstor supports online expansion with automatic data migration and zero downtime. With a standard configuration, the product provides resiliency up to 3 concurrent disk/node failures. Failures are detected automatically, and data reconstruction is also performed automatically. This means that if the time between failures is sufficient for reconstructing data, the system will withstand any number of them. For more information about NEC HYDRAstor, refer to the NEC official website.

To know more about the integration of NAKIVO Backup & Replication with NEC HYDRAstor, refer to the following articles:

- Storage Integration Requirements
- Integrating with NEC HYDRAstor
- Backup Repository on Deduplication Appliance

# NetApp

NetApp storage systems are enterprise-class data storage and retrieval systems. These powerful data storage systems can boost the performance of your enterprise apps, provide unified storage to save money and free up space, simplify ongoing management, and streamline operations.

When using NetApp with NAKIVO Backup & Replication, you can use the backup from storage snapshots approach. This approach allows you to significantly reduce the load of data protection activities on your production environment.

For details about the integration of NetApp storage devices with NAKIVO Backup & Replication, refer to “Backup from Storage Snapshots” on page 69.

# BaaS

NAKIVO Backup & Replication allows for creating and managing multiple isolated tenants within one product instance.

This section contains the following topics:


- "Branding" on page 100
- "License Delegation" on page 101
- "MSP Console" on page 101
- "Multi-Tenancy" on page 104
- "Self-Service" on page 105

# Branding

Whether you plan to use NAKIVO Backup & Replication internally or provide backup/DR-as-a-Service to external customers, you may find it beneficial to align the product's look and feel with your company's brand. NAKIVO Backup & Replication provides a simple way to customize your product's interface so that it looks like an integral part of your organization. You can customize:

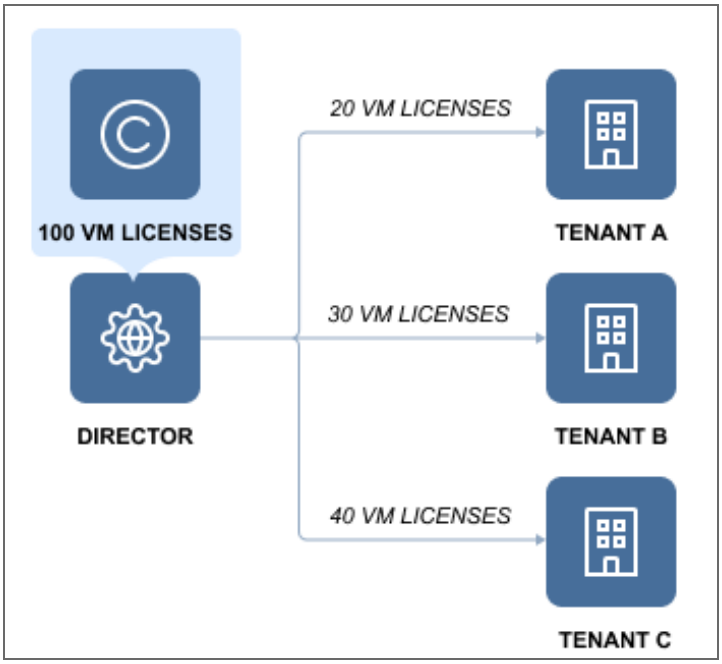
- **Product**: Product title and product logo.
- **Company information**: Company name and website URL.
- **Contact information**: Email, support email, and contact phone.

For information on branding configuration, refer to “Branding” on page 411.

# License Delegation

In Multi-tenant mode, NAKIVO Backup & Replication enables you to create multiple isolated tenants in a single copy of the product. The tenants can represent branch offices/departments in enterprise environments or clients in Cloud Provider environments.

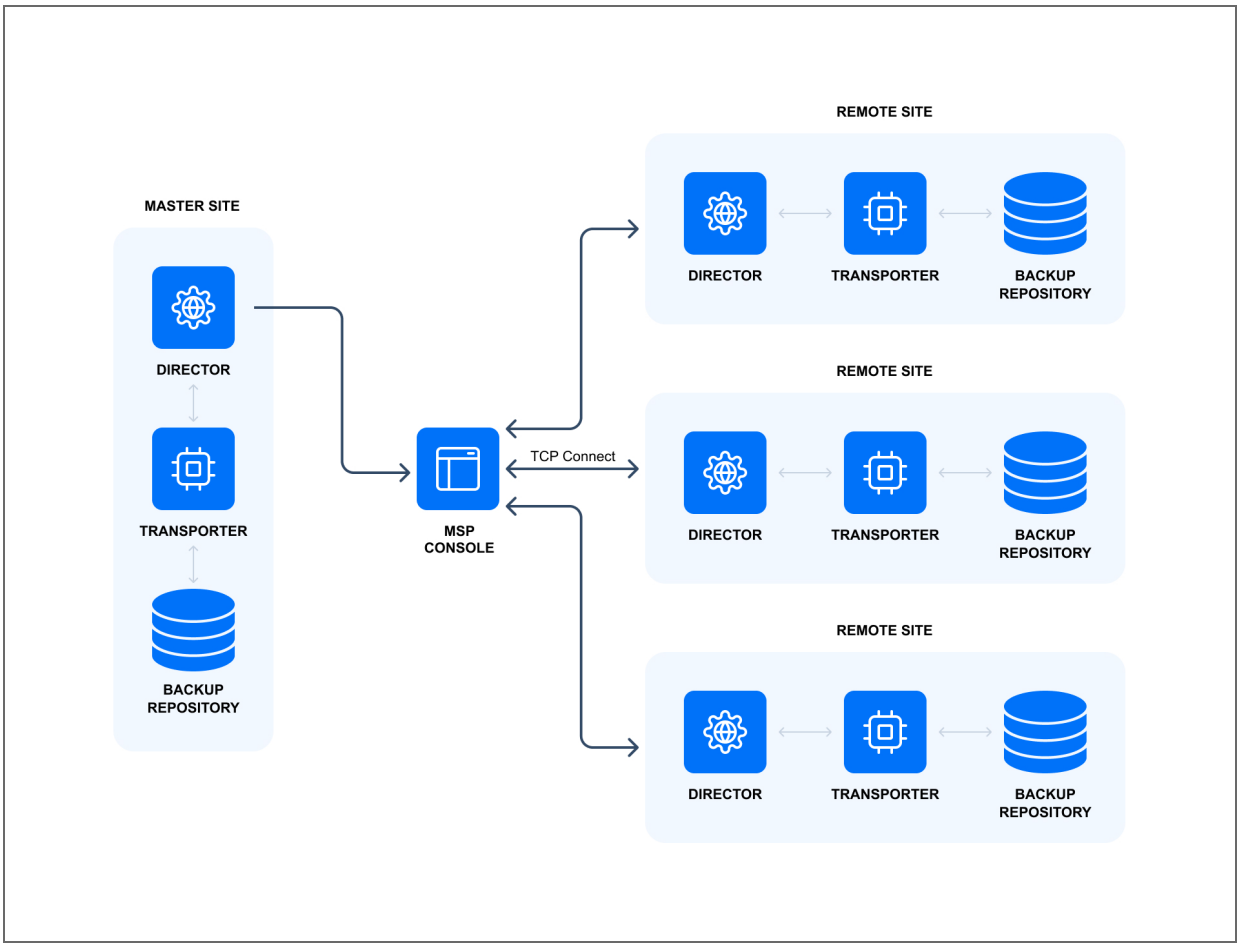
Since tenants are isolated and need to have a limit as to how many licenses each of them can use, NAKIVO Backup & Replication has provided the License Delegation feature. In Multi-tenant mode, a Master Admin (tenant manager) can install one multi-socket license in the product and then assign or delegate a specific number of licenses to each tenant. For example, the Master Admin can install a 20-socket license in the Multi-tenant mode of NAKIVO Backup & Replication, and assign 3 licenses to Tenant A, 2 licenses to Tenant B, and 4 licenses to Tenant C, and let 11 licenses remain unused.




At any moment, the Master Admin can redistribute licenses: revoke any number of licenses from any tenant, which will return them to the Master License Pool, and add licenses to another tenant. The License Delegation feature makes license management simple and manageable in large and distributed environments.

# MSP Console

Managed service providers (MSPs) leverage NAKIVO Backup & Replication to simplify configuring cloud backup infrastructures, empowering tenants to effortlessly and securely send and store their backup data in the cloud. This robust framework allows MSPs to deliver cloud backup repository services and disaster recovery capabilities seamlessly to their clients (tenants). This approach guarantees a reliable strategy for data protection and facilitates efficient disaster recovery in the cloud.




MSPs can also use NAKIVO Backup & Replication to manage clients with standalone instances of NAKIVO Backup & Replication by establishing bidirectional connectivity between MSP and client sites.

**MSP Console** allows managed service providers to manage the remote environments of clients centrally. MSPs can view local tenants as well as remote tenants with standalone instances of NAKIVO Backup & Replication.

Unlike local tenants in the traditional Multi-Tenancy workflow, remote tenants added to the **MSP Console** retain the ability to manage their resources in their data protection infrastructure.

See the topics below for more information:

- “MSP Architecture” on page 439
- “Using the MSP Console” on page 1265

For more information on using the **MSP Console** as a managed service provider (MSP), refer to the following topics:

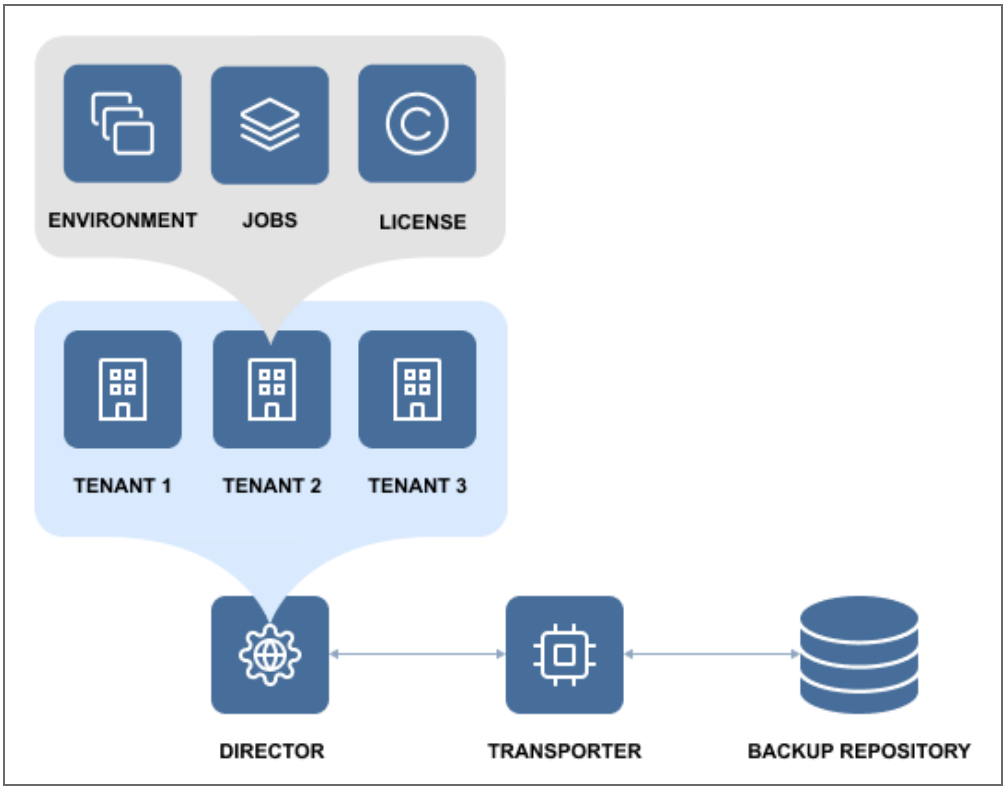
- “Creating New Tenants” on page 1267
- “Opening MSP Console” on page 1266

For more information on connecting a standalone instance to an MSP as a remote tenant, refer to the following topics:

- “MSP Tab in Single-Tenant Mode” on page 436
- “Adding an MSP” on page 437
- “Managing an MSP Connection” on page 438

# Multi-Tenancy

Multi-tenancy enables you to create and manage up to 100 isolated tenants within a single copy of the product. Tenants can represent business units, branch offices, departments, customers, and any other entities.




In Multi-tenant mode, each tenant can access their own environment through a self-service portal, and perform all data protection and recovery tasks. At the same time, tenants are isolated from each other and cannot access the environment and jobs of other tenants.

With Multi-tenancy, you can:

- Deliver Backup-as-a-Service, Replication-as-a-Service, and Disaster-Recovery-as-a-Service, for VMware, Hyper-V and AWS EC2 environments more efficiently and cost-effectively.
- Reduce complexity by managing multiple tenants in a single pane of glass.
- Offload data protection and recovery tasks to tenants.
- Reduce footprint by managing tenants in a single instance of the product.

# Self-Service

In the multi-tenant mode, you can provide tenants with access to their dashboards. By default, a tenant admin account is automatically created when you create a new tenant. If you assign the **Self-service administrator** role to the tenant admin, the tenant admin has full control over all product features inside the tenant dashboard. This includes editing and updating tenant **Inventory**, **Transporters**, and **Backup Repositories**, creating and managing jobs and groups, as well as managing local users and user roles. For each tenant, one guest account can be created. The tenant guest usually has limited permissions inside the tenant.

To provide a tenant with access to the self-service interface, send them the following information:

- Link to NAKIVO Backup & Replication Director
- Tenant login
- Tenant password

# NAKIVO Licensing Policy

This page offers an overview of the NAKIVO Backup & Replication licensing policy. The policy includes the licensing models for different platforms and the type of technical support provided with each model.


- Licensing for NAKIVO Backup & Replication
  - Perpetual Licenses
  - Per-Workload Subscription Licenses
  - License Units
  - NAS File Share Backup Rules
  - License Expiration
- Licensing for Backup for Microsoft 365
- Real-Time Replication for VMware Licensing
- IT Monitoring Licensing
- Frequently Asked Questions
- Additional Resources


## Licensing for NAKIVO Backup & Replication

NAKIVO Backup & Replication is available in 5 editions with a Perpetual License or a Subscription License depending on the platform to be protected and an organization's data protection requirements.

### Perpetual Licenses

Perpetual Licenses are available for virtual machines, physical machines, NAS, and Oracle Database on the following terms:

- For VMware vSphere, Microsoft Hyper-V, Nutanix AHV virtual machines (VMs), and Proxmox VEs the solution is licensed per CPU socket. That is, a license is required for each CPU socket on hosts with VMs to be backed up or replicated. Licensed sockets can be used for any of the four platforms and may be reassigned at any time.
- For physical machines, the solution is licensed per server and per workstation. Perpetual Licenses for physical machine backup (servers or workstations) are sold in bundles of 5 servers/workstations.

  **Note**
  A per-server Perpetual License cannot be used for physical workstations, and a per-workstation Perpetual License cannot be used for servers.

- It is possible to purchase a license for a single bundle of 5 physical servers in case you also purchase a license for a bundle of 10 physical workstations along with it.

- For NAS backup, the solution is licensed per terabyte (see NAS File Share Backup Rules for more details).
- For Oracle Database, the solution is licensed per database (available with the Enterprise Plus edition only).

Perpetual Licenses come with one year of Standard Support. Additional years of support can be purchased upfront. Upgrades to 24/7 Support are also available.

**Notes**

- Valid support is required to receive product updates.
- License calculation is based on the physical processors that are installed in the existing sockets. Therefore, empty sockets are not licensed.

See a breakdown of the different editions below. For a detailed comparison of each edition's features, refer to the Editions Comparison section on the Pricing and Editions page.

<table>
<tr><th>Edition</th><th>Platform</th><th>License unit limitations</th><th>Overview</th></tr>
<tr><td rowspan="6">Pro Essentials</td><td>VMware vSphere</td><td rowspan="4">Min. 2 sockets<br>Max. 6 sockets</td><td rowspan="6">All the features of the Pro edition but with a limit on the number of license units (see License Units)</td></tr>
<tr><td>Nutanix AHV</td></tr>
<tr><td>Microsoft Hyper-V</td></tr>
<tr><td>Proxmox VE</td></tr>
<tr><td rowspan="2">Windows/Linux Physical Machines</td><td>Min. 10 servers<br>Max. 50 servers</td></tr>
<tr><td>Min. 10 workstations<br>Max. 150 workstations</td></tr>
<tr><td>NAS</td><td>Min. 1 TB<br>Max. 50 TB</td></tr>
</table>

| | | | |
|---|---|---|---|
| **Enterprise Essentials** | VMware vSphere | Min. 2 sockets<br>Max. 6 sockets | All the features of the Enterprise edition but with a limit on the number of licensed units (see License Units) |
| | Nutanix AHV | | |
| | Microsoft Hyper-V | | |
| | Proxmox VE | | |
| | Windows/Linux Physical Machines | Min. 10 servers<br>Max. 50 servers | |
| | | Min. 10 workstations<br>Max. 150 workstations | |
| | NAS | Min. 1 TB<br>Max. 50 TB | |
| **Pro** | VMware vSphere | No limits | Includes most product features with limitations on backup to the cloud, administrative tools, and BaaS |
| | Nutanix AHV | | |
| | Microsoft Hyper-V | | |
| | Proxmox VE | | |
| | Windows/Linux Physical Machines | | |
| | NAS | | |
| **Enterprise** | VMware vSphere | No limits | Includes all product features except Oracle Database backup and a few administration features (see the Pricing and Editions page for the full list) |
| | Nutanix AHV | | |
| | Microsoft Hyper-V | | |
| | Proxmox VE | | |
| | Windows/Linux Physical Machines | | |
| | NAS | | |

| | | | |
|---|---|---|---|
| **Enterprise Plus** | VMware vSphere | No limits | The most complete edition of NAKIVO Backup & Replication |
| | Nutanix AHV | | |
| | Microsoft Hyper-V | | |
| | Proxmox VE | | |
| | Windows/Linux Physical Machines | | |
| | NAS | | |
| | Oracle Database | | |

## Per-Workload Subscription Licenses

The Per-Workload Subscription Licenses are available for virtual machines, physical machines, NAS, and Oracle Database on the following terms:

- For VMware vSphere, Microsoft Hyper-V, Nutanix AHV, Proxmox VE, and Amazon EC2, the solution is licensed per VM/instance.
- For physical machines, the solution is licensed per 1 server or 3 workstations.
- For NAS, the solution is licensed per 0.5 terabytes (see NAS File Share Backup Rules for more details).
- For Oracle Database, the solution is licensed per database (available with the Enterprise Plus edition only).

Subscription Licenses are annual subscriptions (1 to 5 years) that are billed upfront and include 24/7 Support for the licensed period.

See a breakdown of the different editions below. For a detailed comparison of each edition's features, refer to the Editions Comparison section on the Pricing and Editions page.

| Edition | Platform | License unit limitations | Overview |
|---|---|---|---|
| **Pro Essentials** | VMware vSphere | Min. 5 workloads<br>Max. 50 workloads | All the features of the Pro edition but with a limit on the number of license units (see License Units) |
| | Nutanix AHV | | |
| | Microsoft Hyper-V | | |
| | Proxmox VE | | |
| | Amazon EC2 | | |
| | Windows/Linux Physical Machines | | |
| | NAS | | |

| | | | |
|---|---|---|---|
| **Enterprise Essentials** | VMware vSphere | Min. 5 workloads<br>Max. 50 workloads | All the features of the Enterprise edition but with a limit on the number of licensed units (see License Units) |
| | Nutanix AHV | | |
| | Microsoft Hyper-V | | |
| | Proxmox VE | | |
| | Amazon EC2 | | |
| | Windows/Linux Physical Machines | | |
| | NAS | | |
| **Pro** | VMware vSphere | No limits | Includes most product features with limitations on backup to the cloud, administrative tools, and BaaS |
| | Nutanix AHV | | |
| | Microsoft Hyper-V | | |
| | Proxmox VE | | |
| | Amazon EC2 | | |
| | Windows/Linux Physical Machines | | |
| | NAS | | |
| **Enterprise** | VMware vSphere | No limits | Includes all product features except Oracle Database backup and a few administration features (see the Pricing and Editions page for the full list) |
| | Nutanix AHV | | |
| | Microsoft Hyper-V | | |
| | Proxmox VE | | |
| | Amazon EC2 | | |
| | Windows/Linux Physical Machines | | |
| | NAS | | |

| | | | |
|---|---|---|---|
| **Enterprise Plus** | VMware vSphere | No limits | The most complete edition of NAKIVO Backup & Replication |
| | Nutanix AHV | | |
| | Microsoft Hyper-V | | |
| | Proxmox VE | | |
| | Amazon EC2 | | |
| | Windows/Linux Physical Machines | | |
| | NAS | | |
| | Oracle Database | | |

# License Units

License units are defined differently for Perpetual Licenses and Per-Workload Subscription Licenses as shown below. In addition, there are limitations on the number of license units with the Pro Essentials and Enterprise Essentials editions.

## Units for Perpetual Licenses

| Platform | License Unit | Pro Essentials/Enterprise Essentials Editions Limits* |
|---|---|---|
| VMware vSphere | 1 CPU Socket | 2-6 Units (Sockets) |
| Microsoft Hyper-V | | |
| Nutanix AHV | | |
| Proxmox VE | | |
| Windows/Linux Physical Server | 5 Servers | 2-10 Units (10-50 Servers) |
| Windows/Linux Workstation | 5 Workstations | 2-30 Units (10-150 Workstations) |
| NAS File Share | 1 Terabyte | 1-50 Units (1-50 TB) |
| Oracle Database | 1 Database | N/A |

*A Perpetual License for Pro Essentials/Enterprise Essentials can cover up to 30 units of virtual machines, physical machines, and/or NAS combined.

Below is an example of a valid order for a single Pro Essentials/Enterprise Essentials Perpetual License that combines virtual, physical, and file share protection for a total of 30 units:

- 6 Sockets (6 units)
- 40 Physical Servers (8 units in bundles of 5)
- 40 Physical Workstations (8 units in bundles of 5)
- 8 TB of file share space (8 units)

## Workloads in Per-Workload Subscription License

<table>
<tr><th>Platform</th><th>License Unit (Workload)</th><th>Pro Essentials/Enterprise Essentials Editions Limits</th></tr>
<tr><td>VMware vSphere</td><td rowspan="4">1 VM</td><td rowspan="9">Minimum of 5 workloads<br>Maximum of 50 workloads</td></tr>
<tr><td>Microsoft Hyper-V</td></tr>
<tr><td>Nutanix AHV</td></tr>
<tr><td>Proxmox VE</td></tr>
<tr><td>Amazon EC2</td><td>1 Instance</td></tr>
<tr><td>Windows/Linux Physical Server</td><td>1 Server</td></tr>
<tr><td>Windows/Linux Workstation</td><td>3 Workstations</td></tr>
<tr><td>NAS File Share</td><td>0.5 Terabyte</td></tr>
<tr><td>Oracle Database</td><td>1 Database</td></tr>
</table>

# NAS File Share Backup Rules

File share backup has a few additional rules and details regarding licensing. Licenses are consumed based on the following rules:

- License consumption is calculated based on backed up source file share data, determined during each file share backup job run.
    - NAKIVO Backup & Replication sums up the last-known amount of protected source data across all file share backup jobs.

    - If the same file share and/or its contents are protected by multiple jobs, the source data is still summed.
    - If a job run reaches or exceeds the licensed data size, the job will become disabled. It will not be possible to create new file share backup jobs, and the current job cannot be re-enabled until it is edited to exclude a sufficient amount of backup data.
- Adding shares to inventory does **not** consume licenses.
- File share backup metadata does **not** contribute to licensed file share size.

In addition, there are specific rules regarding the calculation of licenses for protected source data:

- If the total source data size is greater than zero and less than or equal to 0.5 TB, one license unit is consumed.

> **Note**
> In this section, **license unit** refers only to 0.5 TB. While file share backup is licensed per-Terabyte in a Perpetual License, license **consumption** is counted in 0.5 TB increments for both Perpetual and Per-Workload Subscription license types.

- If the total source data size exceeds 0.5 TB, the number of consumed licenses is determined as follows:
    1. The total source data size is rounded down to the nearest multiple of 0.5 TB
    2. The rounded amount is divided by 0.5 TB
    3. The resulting value is the number of licensed units consumed

> **Example**: Total backed up source data of 1850 GB (1.85 TB) is rounded down to 1.5 TB and divided by 0.5 TB to get **3 license units consumed**.

## License Expiration

For both single-tenant and multi-tenant modes, NAKIVO Backup & Replication enters grace period mode for 15 days after license expiration. During this period, the product continues operating normally. After the grace period mode has ended or if the product does not have a license, NAKIVO Backup & Replication enters limited mode. In this mode, the following occurs:

- All backup, backup copy, replication, and site recovery jobs are disabled.
- Jobs that are currently in progress are disabled after completion.
- Real-time replication jobs are stopped and disabled.
- Recovery jobs can still be run in limited mode as normal.

Limited mode is disabled upon changing to the active license.

# Licensing for Backup for Microsoft 365

Backup for Microsoft 365 is licensed per user on an annual basis (1 to 3 years). A user is defined as a unique Microsoft 365 account that has access to Exchange Online, OneDrive for Business, SharePoint Online, and/or Teams. Each user is equivalent to one license unit.

Organizations may purchase a Subscription License for Backup for Microsoft 365 as a standalone offering or combine it with any existing NAKIVO Backup & Replication edition and license type (Perpetual or Per-Workload Subscription). Subscription Licenses come with 24/7 Support covering the licensed period.

See the overview below of possible pairings for a Microsoft 365 Subscription License with any edition of NAKIVO Backup & Replication.

| Purchased with | Support level | Coverage |
|---|---|---|
| Perpetual License (any edition) | 24/7 Support for Microsoft 365 License; Standard Support for Perpetual License | Minimum 10 license units (users) per order |
| | 24/7 Support across the board (requires Support Upgrade for Perpetual License) | |
| Subscription License (any edition) | 24/7 Support across the board | Minimum 10 license units (users) per order |

When combining a Subscription License for Backup for Microsoft 365 with a NAKIVO Backup & Replication Perpetual License of any type, the following technical support conditions apply:

- The end date for support coverage must be the same for both licenses.
- You may upgrade Perpetual License Standard Support to 24/7 Support, or keep it at the default Standard Support.

**Note**

Users with a Microsoft 365 Student SKU do not require a license. Users with a Microsoft 365 Faculty SKU or another non-student SKU still require a license.

The following Student SKUs are supported:

| Plan Name | SkuPartNumber | SkuID |
|---|---|---|
| Office 365 Education E3 for Students | ENTERPRISEPACK_STUDENT | 8fc2205d-4e51-4401-97f0-5c89ef1aafb |
| Office 365 A1 for Students | STANDARDWOFFPACK_STUDENT | 314c4481-f395-4525-be8b-2ec4bb1e9d91 |
| Office 365 A5 for Students | ENTERPRISEPREMIUM_STUDENT | ee656612-49fa-43e5-b67e-cb1fdf7699df |
| Office 365 A5 without PSTN Conferencing for Students | ENTERPRISEPREMIUM_NOPSTNCONF_STUDENT | 1164451b-e2e5-4c9e-8fa6-e5122d90dbdc |
| Office 365 Education E1 for Students | STANDARDPACK_STUDENT | d37ba356-38c5-4c82-90da-3d714f72a382 |
| Microsoft 365 A3 for Students | M365EDU_A3_STUDENT | 7cfd9a2b-e110-4c39-bf20-c6a3f36a3121 |
| Microsoft 365 A3 for Students use benefits | M365EDU_A3_STUUSEBNFT | 18250162-5d87-4436-a834-d795c15c80f3 |
| Microsoft 365 A5 for Students | M365EDU_A5_STUDENT | 46c119d4-0379-4a9d-85e4-97c66d3f909e |
| Microsoft 365 A5 Student use benefits | M365EDU_A5_STUUSEBNFT | 31d57bc7-3a05-4867-ab53-97a17835a411 |
| Microsoft 365 A5 without Audio Conferencing for Students | M365EDU_A5_NOPSTNCONF_STUDENT | a25c01ce-bab1-47e9-a6d0-ebe939b99ff9 |
| Microsoft 365 A5 without Audio Conferencing for Students use benefit | M365EDU_A5_NOPSTNCONF_STUUSEBNFT | 81441ae1-0b31-4185-a6c0-32b6b84d419f |
| Office 365 A3 for Students | ENTERPRISEPACKPLUS_STUDENT | 98b6e773-24d4-4c0d-a968-6e787a1f8204 |
| Office 365 A3 Student use benefit | ENTERPRISEPACKPLUS_STUUSEBNFT | 476aad1e-7a7f-473c-9d20-35665a5cbd4f |

| | | |
|---|---|---|
| Office 365 A5 Student use benefit | ENTERPRISEPREMIUM_STUUSEBNFT | f6e603f1-1a6d-4d32-a730-34b809cb9731 |
| Office 365 A5 without Audio Conferencing for Students use benefit | ENTERPRISEPREMIUM_NOPSTNCONF_STUUSEBNFT | bc86c9cd-3058-43ba-9972-141678675ac1 |
| Office 365 Education for Homeschool for Students | STANDARDWOFFPACK_HOMESCHOOL_STU | afbb89a7-db5f-45fb-8af0-1bc5c5015709 |
| Office 365 A1 for Students (for Device) | STANDARDWOFFPACK_STUDENT_DEVICE | 160d609e-ab08-4fce-bc1c-ea13321942ac |
| Windows 10/11 Enterprise A3 for students | WIN10_ENT_A3_STU | d4ef921e-840b-4b48-9a90-ab6698bc7b31 |
| Minecraft Education Student | MEE_STUDENT | 533b8f26-f74b-4e9c-9c59-50fc4b393b63 |

## SharePoint Online Backup Licensing Rules

In addition to regular Microsoft 365 Subscription Licensing rules, there are conditions specific to SharePoint Online backup licensing.

- License units for SharePoint Online backup:
  - A user (including a user in groups) that has "Edit" or "Full Control" permissions for a site consumes 1 license unit.
  - For personal sites, only the owner of the personal site consumes a license unit. Other users with access to this personal site do not consume any license units.
  - A personal site owner with access to a regular site requires only one license unit.
- License units are matched to a given email account, meaning the following rules apply for mailbox backup:
  - Users with access to a SharePoint Online site who also have a mailbox under the same email account require only one license unit.
  - If a mailbox does not correspond to a licensed email account, a separate license unit is required to back it up.
- License units are not given per SharePoint site or affected by the size of a site.

# Real-Time Replication for VMware vSphere Licensing

Real-Time replication for VMware vSphere is available with a Perpetual License (per socket) or Per-Workload Subscription License (per VM), and can be purchased separately from NAKIVO Backup & Replication. A Real-Time Replication license can also be combined with an existing license of the same type and edition of NAKIVO Backup & Replication.

## Perpetual Licenses for VMware Real-Time Replication

- Licensed per CPU socket: A license is required for each CPU socket on hosts with VMs selected for real-time replication.
- **Enterprise Essentials** edition requires a minimum of 2 sockets and allows up to 6 sockets.
- Perpetual Licenses come with one year of Standard Support. Additional years of support can be purchased upfront. Upgrades to 24/7 Support are also available.

> **Note**
> Valid support is required to receive product updates.

## Subscription Licenses for VMware Real-Time Replication

- Licensed per VM
- Annual subscriptions (1 to 5 years) that are billed upfront
- Include 24/7 Support for the licensed period

If purchased together with NAKIVO Backup & Replication, the licenses for both products must have the same:

- Edition
- Support end date
- Type of support (Standard or 24/7)

# IT Monitoring Licensing

IT Monitoring for VMware vSphere virtual machines is available with a Perpetual License (per socket) or Per-Workload Subscription License (per VM), and can be purchased separately from NAKIVO Backup & Replication. An IT Monitoring license can also be combined with an existing license of the same type and edition of NAKIVO Backup & Replication.

## Perpetual Licenses for IT Monitoring

- Licensed per CPU socket: A license is required for each CPU socket on hosts with VMs to be monitored.
- Perpetual Licenses come with one year of Standard Support. Additional years of support can be purchased upfront. Upgrades to 24/7 Support are also available.

**Note**

Valid support is required to receive product updates.

## Subscription Licenses for IT Monitoring

- Licensed per VM
- Annual subscriptions (1 to 5 years) that are billed upfront
- Include 24/7 Support for the licensed period
- License unit limits for the Pro Essentials/Enterprise Essentials editions: 5-50 workloads

If purchased together with NAKIVO Backup & Replication, the licenses for both products must have the same:

- Edition
- Support end date
- Type of support (Standard or 24/7)

# Frequently Asked Questions

**Q: What is a socket?**

A: A socket refers to the socket on the motherboard onto which a CPU is inserted. For a Perpetual license, only the number of sockets is counted; the number of CPU cores per socket is not taken into account.

**Q: Does adding another Transporter require an additional license?**

A: NAKIVO Backup & Replication is not licensed per Transporter. You can install additional Transporters regardless of the licensing model (Perpetual or Subscription).

**Q: Do I need to license both source and target hosts in a disaster recovery scenario?**

A: Only the source side of replication requires a license. For a scenario wherein you replicate a VM from Site A, recover it in Site B, then failback to Site A, only hosts on Site A side need to be licensed.

**Q: Will updating the product to its new version reset my current license?**

A: Updating to a major (e.g. 10.8 to 11.0), minor (e.g. 10.0 to 10.5), or sub-minor version (e.g. 10.0.0 to 10.0.1) will not reset the current license of the product.

# Additional Resources

NAKIVO Pricing & Editions

NAKIVO Customer Support Policy

NAKIVO Customer Support Agreement

End-User License Agreement

# Deployment

This section contains the following topics :


- “Architecture” on page 120
- “System Requirements” on page 130
- “Deployment Scenarios” on page 215
- “Installing NAKIVO Backup & Replication” on page 231
- “Updating NAKIVO Backup & Replication” on page 296
- “Uninstalling NAKIVO Backup & Replication” on page 329

# Architecture


- What is NAKIVO Backup & Replication?
- Solution Components


## What is NAKIVO Backup & Replication?

NAKIVO Backup & Replication is an all-in-one solution designed to back up, replicate, and recover virtual machines and cloud instances. The product can also back up and recover physical machines.

## Solution Components

NAKIVO Backup & Replication is a server application that can be installed on a virtual or physical machine. The application is designed to achieve top speeds for CPU and RAM to achieve the top speed of VM backup, replication, and recovery. Thus, NAKIVO Backup & Replication components should be installed on a machine designated for backup and replication so it does not interfere with the performance of other applications.

NAKIVO Backup and Replication consists of the following components:

- “Director” on page 121
- “Transporter” on page 123
- “Backup Repository” on page 127

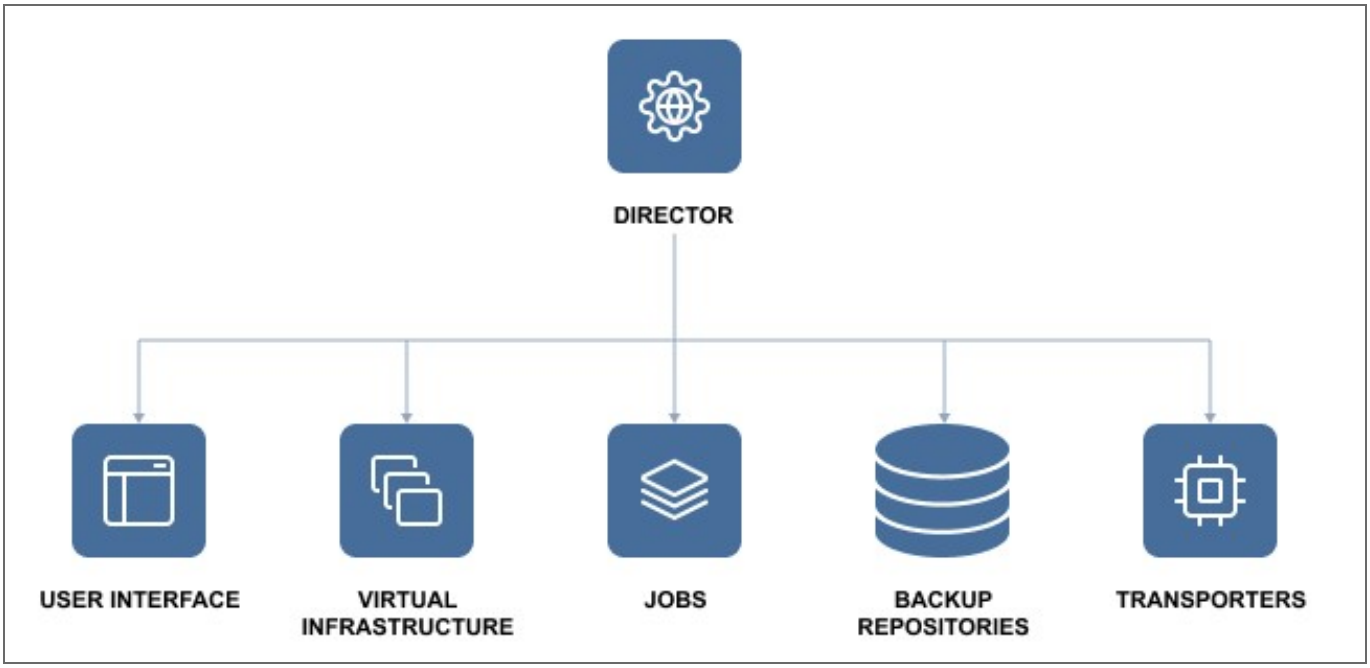
All components can be installed on a single machine or can be distributed across multiple machines and geographical locations.

# Director


- What is Director?
- How Many Directors Should be Deployed?

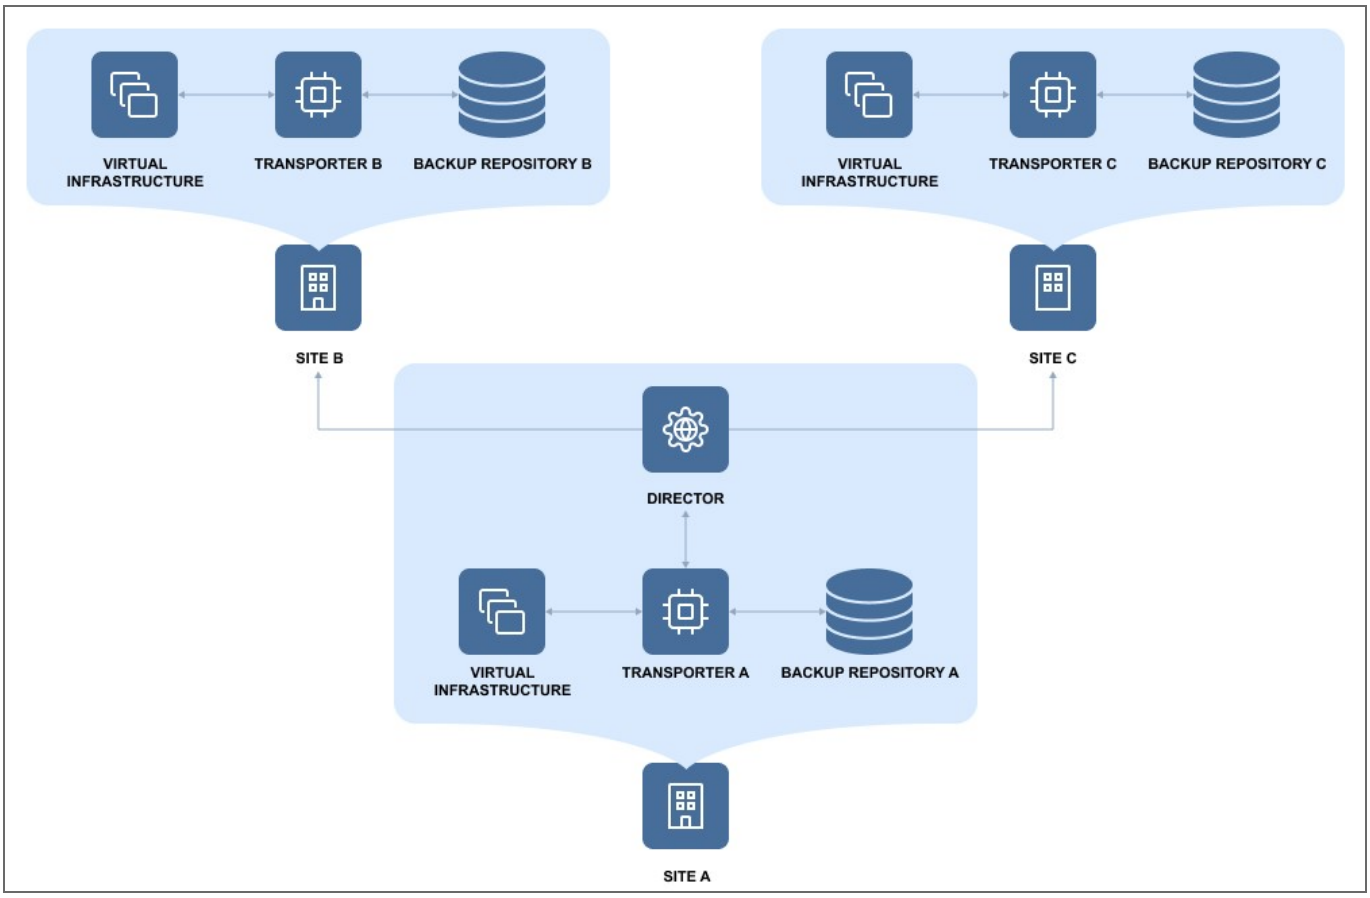

## What is Director?

**Director** is the central management instance of the product. It provides Web interface, locates and maintains the inventory, provides users with the ability to create and run jobs, manages **Backup Repositories**, **Transporters**, and other product elements.




## How Many Directors Should be Deployed

Only one instance of the **Director** should be installed per customer. As a central management point for data protection, one instance of the **Director** can manage multiple geographically distributed virtual and cloud environments, **Backup Repositories**, and **Transporters**. See the example below.

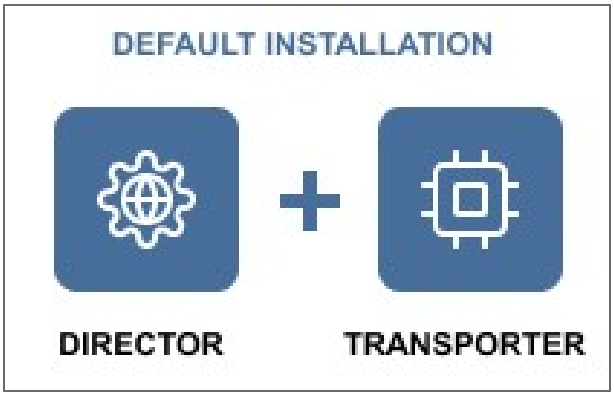

VIRTUAL INFRASTRUCTURE
TRANSPORTER B
BACKUP REPOSITORY B
VIRTUAL INFRASTRUCTURE
TRANSPORTER C
BACKUP REPOSITORY C
SITE B
SITE C
DIRECTOR
VIRTUAL INFRASTRUCTURE
TRANSPORTER A
BACKUP REPOSITORY A
SITE A

# Transporter

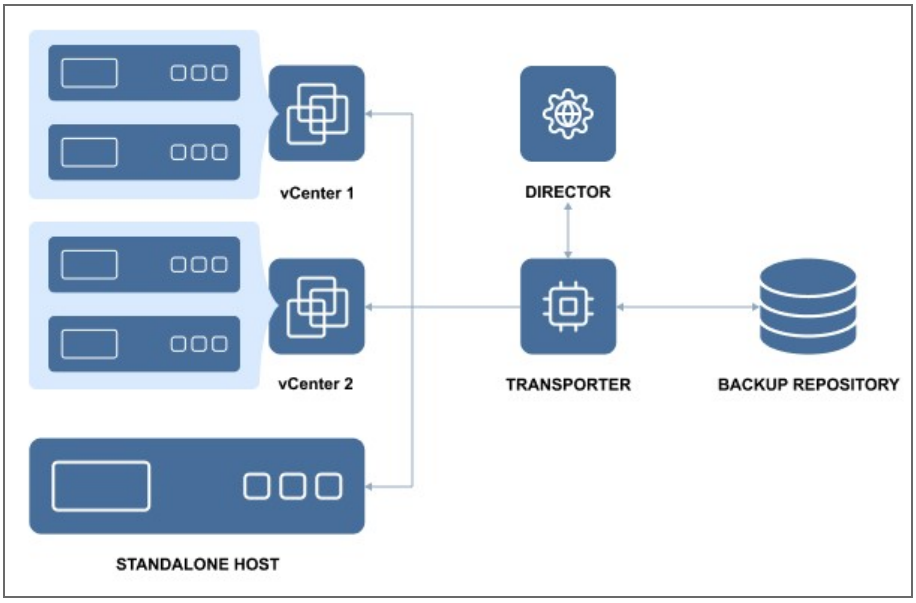

- What is a Transporter?
- How many Transporters Should be Deployed?
- How Transporters are Selected for Jobs
- Transporter Security
- Transporter as a VMware Appliance


## What is a Transporter?

The **Transporter** is the component of the product that does all of the heavy lifting. It performs backup, replication, and recovery, as well as data compression, deduplication, and encryption. An instance of the **Transporter** is automatically installed along with the **Director** to enable backup, replication, and recovery out of the box. The default **Transporter** is called "Onboard Transporter", and it must not be removed or added to the product by another **Director**.



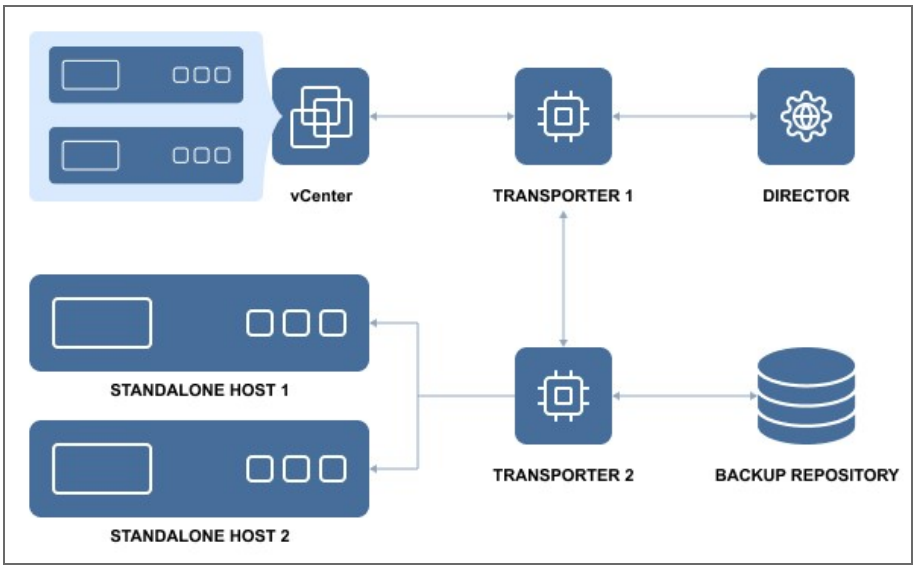

A single **Transporter** can back up, replicate, and recover multiple VMs and cloud instances.



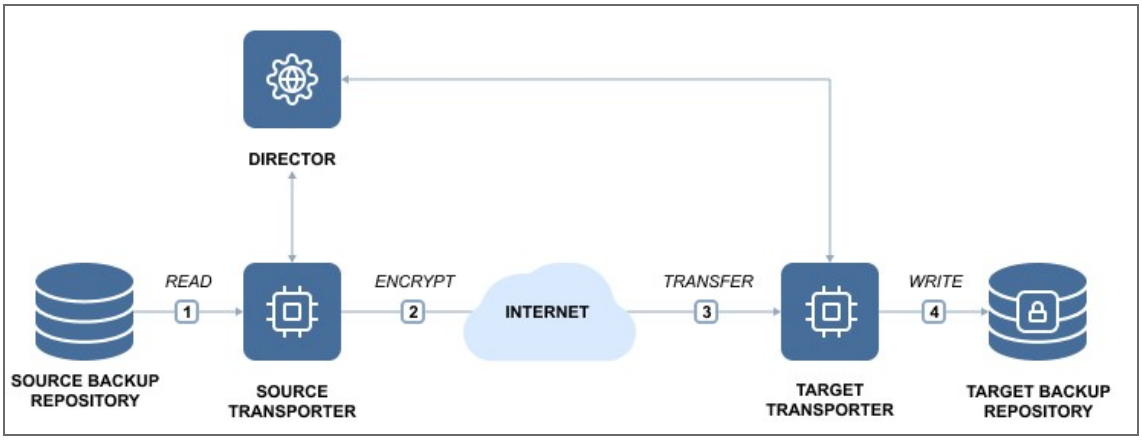
One Transporter can simultaneously process multiple source disks (6 by default) during backup, replication, and recovery. If jobs contain more disks than the **Transporter** is set to process simultaneously, the disks will be put in a queue and will be processed once the **Transporter** frees up.

## How Many Transporters Should be Deployed?

In most cases, it is sufficient to deploy only one **Transporter** per site. In large environments, where multiple source items need to be processed simultaneously, multiple **Transporters** can be deployed to distribute the workload.




Deploying multiple **Transporters** also enables network acceleration and AES 256 encryption of traffic between a pair of **Transporters**. For example, if VMs are replicated over WAN between two sites, the **Transporter** installed in the source site can compress and encrypt data before transferring it over WAN, and the **Transporter** installed in the target site can unencrypt and decompress the data prior to writing it to the target server.

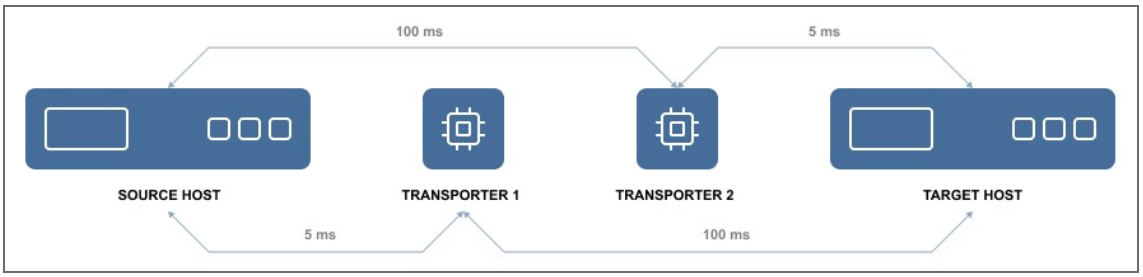


If you plan to transfer data over WAN without a VPN connection from your source site to the target site, make sure the source and target **Transporters** are added to the product using external IP addresses or DNS names that can be properly resolved in WAN, so that the two **Transporters** can connect to each other.

## How Transporters are Selected for Jobs

In large and geographically distributed environments multiple Transporters can be deployed to distribute the data protection workload, optimize network traffic, and improve data transfer speeds. Thus, if more than one **Transporter** is deployed for NAKIVO Backup & Replication, it is important to determine which one should be used to read data from a particular source and which one should be used to write data to a target.

By default, the product automatically determines which **Transporter** should be used based on the proximity of a **Transporter** to the source or target server. The proximity is measured by using the ping round trip time.




In the example above, **Transporter** 1 will be selected to read data from the Source ESXi, and **Transporter** 2 will be selected to write data to the Target ESXi.

The Transporter selection can also be configured manually during job creation.

## Transporter Security

It is possible to set a Master Password for the Transporter and use a CA certificate to make NAKIVO Backup & Replication more secure. The certificate can be set for the Onboard **Transporter** during the full installation of the product or for individual **Transporters** during **Transporter**-only installation, or by using the Windows Updater on Windows operating systems. The master password can be set only during the **Transporter**-only installation.

This option is available for the following supported target platforms:

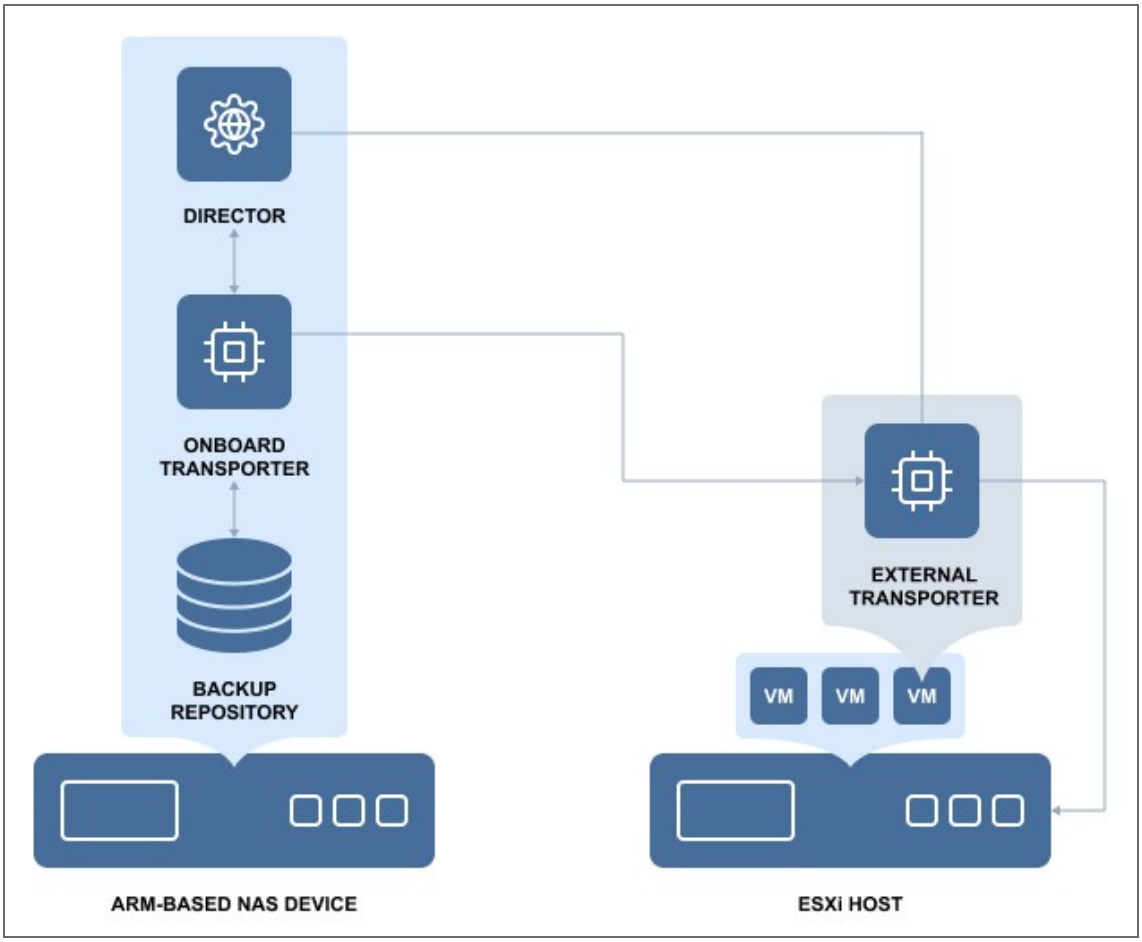
- VMware vSphere
- Microsoft Hyper-V
- Amazon EC2
- Nutanix AHV
- Supported NAS models
- Virtual Appliances
- Physical machines

To use CA certificates, make sure that they adhere to the necessary requirements. Refer to Custom CA-Signed Certificate Compatibility.

## Transporter as a VMware Appliance

Since VMware does not provide a VDDK library for ARM-based processors, the NAKIVO Onboard **Transporter** may not support VMware as some functionality necessary for working with VMware is missing for such ARM-based NAS devices.

In this case, you will need to deploy an additional **Transporter** as a VMware appliance to allow NAKIVO Backup & Replication to work with VMware vCenters and ESXi hosts, and protect your virtual infrastructure.




Once deployed, the additional **Transporter** allows the application to retrieve necessary data via Transporter-to-Transporter communication.

# Backup Repository


- What is a Backup Repository?
- How Much Data Can Be Stored in a Backup Repository?
- How is a Backup Repository Managed?


## What is a Backup Repository?

A **Backup Repository** is a folder used by NAKIVO Backup & Replication to store backups. When you add a **Backup Repository** to the product, NAKIVO Backup & Replication creates a folder named “NakivoBackup” in the specified location and keeps all backed up data and **Backup Repository** metadata in that folder.

**Important**

- Do not modify or delete any files inside the “NakivoBackup” folder. Modifying or deleting any file inside the “NakivoBackup” folder may irreversibly damage an entire **Backup Repository**.
- To avoid disrupting NAKIVO Backup & Replication processes and data corruption, add the application to the whitelist/exclusions list of the antivirus software running on the machine on which the NAKIVO Backup Repository is set up.

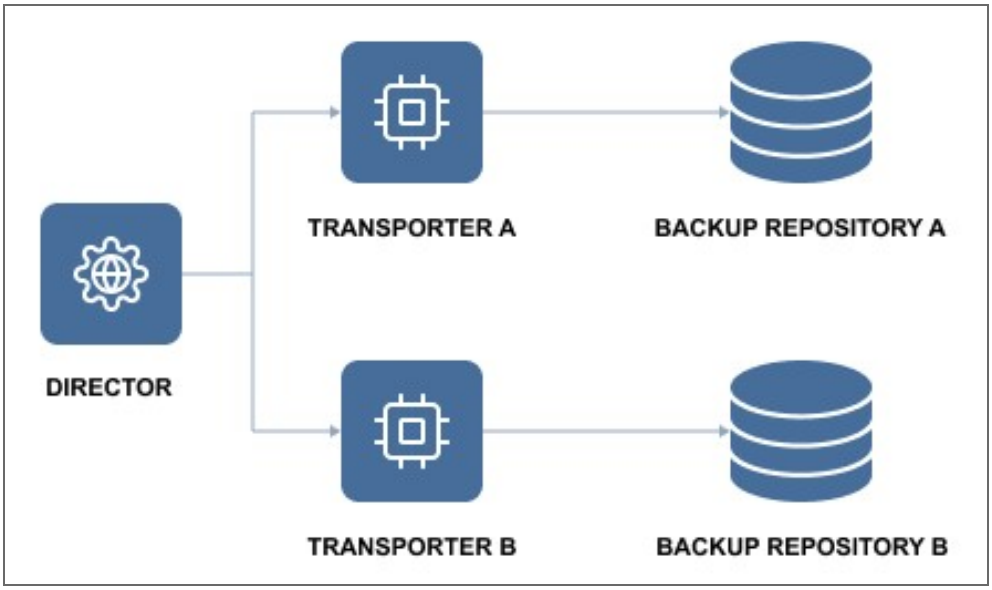
By default, a **Backup Repository** is created when the full solution (both **Director** and **Transporter**) is installed. The default **Backup Repository** is named “Onboard repository”.

## How Much Data Can Be Stored in a Backup Repository?

The maximum recommended size of a Backup Repository used with NAKIVO Backup & Replication is 128 TB of data after compression and deduplication. For repositories larger than 128 TB, it is recommended to use an **Incremental with full backups** type of **Backup Repository**. The number of **Backup Repositorie**s per installation is unlimited. Additionally, **Backup Repositories** can be configured to compress and deduplicate backups at the block level to save storage space.

## How is a Backup Repository Managed?

Each **Backup Repository** is managed by a single **Transporter** called an Assigned **Transporter**. In other words, only one **Transporter** can read data from and write data to a particular **Backup Repository**.




The Assigned **Transporter** is responsible for all interaction with its **Backup Repository**. A single **Transporter** can be assigned to and manage multiple **Backup Repositories**.

# Federated Repository

What Is a Federated Repository?
How Does a Federated Repository Work?
How Many Members Can Be Added to a Federated Repository?

## What Is a Federated Repository?

NAKIVO Backup & Replication supports creating a **Federated repository** to be used as a destination for new and existing backup/backup copy jobs, backup export, or backup migration. Also, the federated repository can be used as a source for recovery or backup copy jobs. With this feature, existing backup repositories with identical values for their data storage type attribute can be combined into a federated repository. These federated repositories are scaled horizontally to consume multiple backup repositories (federated repository members). Thus, backup/backup copy jobs can continue to run even if one of the members is unavailable or has insufficient storage space. You can use a local folder, NFS share, or CIFS share type of backup repository to create a federated repository.

**Important**

The **Federated repository** feature supports backing up data to **Incremental with full backups** repositories only.

## How Does a Federated Repository Work?

When you create a federated repository, the **Director** creates this federated repository consisting of one or more federated repository members. When you assign this federated repository as a target for new or existing backup/backup copy jobs, backup export, or backup migration, during the job run, the **Director** determines and assigns an available federated repository member with sufficient storage space to store the backup data.

A federated repository can also be used as a source for recovery and/or replication jobs from backup.

The **Federated repository** feature supports the following backup location policy:

- **Integrity**: All dependent recovery points are stored within a single federated repository member.

To use federated repository members (on NFS and CIFS shares or local folders) in your data protection workflows, follow these steps:

1. Create and set up a federated backup repository.
2. Select the federated repository as a destination for new or existing backup/backup copy jobs.
3. Select the federated repository as a source for recovery or backup copy jobs, as a source for replication from backup jobs.

## How Many Members Can Be Added to a Federated Repository?

The maximum recommended number of members assigned to one federated repository is 128. Members of the same federated repository can use different Transporters, and they must be of the **Incremental with full backups** repository type. The maximum total number of members in all federated repositories is 500 (1 member in each federated repository).

If needed, federated repository members can be managed/refreshed individually or added/removed to/from the federated backup repository.

**Note**

You cannot simultaneously select:

- The members that support immutability and contain at least one immutable object or are associated with job(s) configured to create an immutable recovery point(s) and
- The members that do not support immutability.

Selecting one type of member disables the member(s) of the other type.

# System Requirements

Before you start using NAKIVO Backup & Replication, make sure that the servers or machines that you plan to use as backup infrastructure components meet the requirements listed in the following topics:


- "Supported Platforms" on page 131
- "Storage Integration Requirements" on page 135
- "Deployment Requirements" on page 138
- "Feature Requirements" on page 169

# Supported Platforms

NAKIVO Backup & Replication provides data protection for the following platforms:

- VMware vSphere v5.5 - 8.0d
- VMware vSphere 9.0.0.0100 (Compatibility Support)
- Nutanix Acropolis v6.7.1.5 (STS) - v6.7.1.7 (STS)
- Nutanix Acropolis v6.8 (eSTS) - v6.8.1.8 (eSTS)

**Notes**

- To learn about the limitations of NAKIVO Backup & Replication related to supported platforms, refer to the Platform Limitations section of the latest Release Notes.
- To add a supported platform to NAKIVO Backup & Replication, make sure that your system has been updated with the latest patch and all the necessary requirements are met.
- The support for sub-versions that are not stated in the user guide can be clarified with the support team.
- NAKIVO Backup & Replication does not support backup and recovery of ACL data via NFSv3.

Find the necessary requirements below:

- Hypervisor Requirements
  - VMware vSphere
- Public Cloud Requirements
- Cloud Region Requirements

# Hypervisor Requirements

To provide data protection for your virtual environments, make sure the following requirements are met:

## VMware vSphere

- Full administrative permissions (recommended), or limited permissions are required.
- VPXA service must be running.
- NAKIVO Backup & Replication uses QEMU Guest Agent to quiesce the guest OS. It may not detect if the VSS service is not working. This may occur for several reasons, such as due to incorrect permissions.

# Public Cloud Requirements

NAKIVO Backup & Replication does not support the following:

- Backblaze deltas that are larger than 4 TiB
- AWS/Wasabi/S3-compatible deltas that are larger than 20 TiB
- Azure deltas that are larger than 800 TiB

## Amazon S3 and Backblaze

Refer to Required AWS IAM Permissions for Amazon S3 and Backblaze for details.

**Note**

There is also an option for granting full IAM permissions for NAKIVO Backup & Replication.

## Google Cloud Storage

The Google Cloud Storage can be added as a Generic S3-Compatible Object Storage. To add a Google Cloud Storage, the Google Cloud account must be preconfigured. Refer to Required Google Cloud Storage Permissions for details.

## Microsoft Azure

NAKIVO Backup & Replication supports the following storage account types:

| Storage account type | Supported services | Supported access tiers |
|---|---|---|
| Standard general-purpose v2 | Blob storage (including Data Lake Storage), Queue Storage, Table storage, and Azure Files | Hot, Cool |
| Standard general-purpose v1 | Blob storage (including Data Lake Storage), Queue Storage, Table storage, and Azure Files | N/A |

Refer to “Configuring a Microsoft Azure Storage Account” on page 526 for more information.

# Cloud Region Requirements

For **Amazon S3** or **Amazon EC2**, NAKIVO Backup & Replication supports the following regions:

- AWS GovCloud (US-West)
- AWS GovCloud (US-East)
- US East (Ohio)
- US East (N. Virginia)
- US West (N. California)

- US West (Oregon)
- Africa (Cape Town)
- Asia Pacific (Hong Kong)
- Asia Pacific (Hyderabad)
- Asia Pacific (Jakarta)
- Asia Pacific (Malaysia)
- Asia Pacific (Melbourne)
- Asia Pacific (Mumbai)
- Asia Pacific (Osaka)
- Asia Pacific (Seoul)
- Asia Pacific (Singapore)
- Asia Pacific (Sydney)
- Asia Pacific (Taipei)
- Asia Pacific (Thailand)
- Asia Pacific (Tokyo)
- Canada (Calgary)
- Canada (Central)
- EU (Frankfurt)
- EU (Zurich)
- EU (Ireland)
- EU (London)
- EU (Milan)
- EU (Spain)
- EU (Paris)
- EU (Stockholm)
- Israel (Tel Aviv)
- Mexico (Central)
- Middle East (Bahrain)
- Middle East (UAE)
- South America (Sao Paulo)

For **Wasabi**, NAKIVO Backup & Replication supports the following regions:

- Wasabi US East 1 (N. Virginia)
- Wasabi US East 2 (N. Virginia)
- Wasabi US Central 1 (Texas)
- Wasabi US West 1 (Oregon)
- Wasabi CA Central 1 (Toronto)

- Wasabi EU West 1 (London)
- Wasabi EU West 2 (Paris)
- Wasabi EU West 3 (London)
- Wasabi EU Central 1 (Amsterdam)
- Wasabi EU Central 2 (Frankfurt)
- Wasabi AP Northeast 1 (Tokyo)
- Wasabi AP Northeast 2 (Osaka)
- Wasabi AP Southeast 1 (Singapore)
- Wasabi AP Southeast 2 (Sydney, Australia)

# Storage Integration Requirements

NAKIVO Backup & Replication can be integrated with deduplication appliances including Dell EMC Data Domain, NEC HYDRAstor, and HP StoreOnce (Catalyst) appliances by using an **Incremental-with-full-backups** repository. Deduplication appliances are servers designed to reduce data size and can be used as backup targets. They operate best with sequential large block I/O from backup software. Therefore, when backing up to a deduplication appliance, it is important to make sure that the architecture of your Backup Repository is optimized for these devices and your backups have a large block I/O. Only by doing this will you be able to maximize your backup speed. NAKIVO Backup & Replication provides advanced integration with the following storage solutions:

- Dell-EMC Data Domain
- NEC HYDRAstor
- HPE StoreOnce with Catalyst Support
- HPE 3PAR Storage
- HPE Nimble Storage
- HPE Primera Storage
- HPE Alletra Storage
- NetApp

## Dell-EMC Data Domain

**Supported versions**:

- Dell-EMC Data Domain 6.1 - 8.1.0.0

## NEC HYDRAstor

**Supported systems**:

- NEC HYDRAstor v5.5.1 - 5.7.1
- NEC Storage HS Universal Express I/O Module Version v1.8.0 - 1.8.8

## HPE StoreOnce with Catalyst Support

**Supported versions**:

- HPE StoreOnce 3.18.18
- HPE StoreOnce 4.2.3
- HPE StoreOnce 4.3.2
- HPE StoreOnce 4.3.6
- HPE StoreOnce 4.3.7

**Integration requirements and limitations**:

- NAKIVO Backup & Replication installed on Windows (x64) and Linux (x64) machines must have HPE StoreOnce Catalyst API Library.
- HPE StoreOnce Catalyst integration is not supported on devices with ARM7 and ARM64 (AArch64) processors.

**Supported Maximums**

| StoreOnce Model | Maximum Sessions | Maximum Transporter Load | Maximum Recovery Points |
|---|---|---|---|
| **VSA** | | | |
| VSA Gen 4 (128+sessions) | 128-256 | 6 | 7 |
| **HPE ProLiant Gen 10 (StoreOnce 4.2.3)** | | | |
| 3620 | 128 | 6 | 7 |
| 3640 | 192 | 6 | 14 |
| 5200 | 512 | 10 | 21 |
| 5250 | 512 | 10 | 21 |
| 5650 | 1024 | 16 | 30 |
| **HPE ProLiant Gen 9 (StoreOnce 3.18.18)** | | | |
| 3500 | 192 | 6 | 14 |
| 5100 | 320 | 10 | 14 |
| 5500 | 1000 | 16 | 30 |
| 6600 | 1024 | 16 | 30 |
| **HPE ProLiant Gen 8 (StoreOnce 3.18.18)** | | | |
| 4500 | 128 | 6 | 7 |
| 4700 | 192 | 6 | 14 |
| 4900 | 500 | 10 | 21 |
| 6500 | 512 | 10 | 21 |

## HPE 3PAR Storage

An HPE 3PAR storage added to the NAKIVO Backup & Replication inventory allows you to back up VMware VMs from HPE 3PAR storage snapshots. For details, refer to HPE 3PAR Storage documentation.

**Supported versions**:

- HPE 3PAR OS 3.1.2 and above

## HPE Nimble Storage

An HPE Nimble storage device added to the NAKIVO Backup & Replication inventory allows you to back up VMware VMs from HPE Nimble storage snapshots. For details, refer to https://www.hpe.com/us/en/storage/nimble.html.

**Supported versions**:

- HPE Nimble OS 5.0.2 and above

## HPE Primera Storage

An HPE Primera storage device added to the NAKIVO Backup & Replication inventory allows you to back up VMware VMs from HPE Primera storage snapshots. For details, refer to https://www.hpe.com/us/en/storage/hpe-primera.html.

**Supported versions**:

- HPE Primera OS 4.5 and above

## HPE Alletra Storage

An HPE Alletra storage device added to the NAKIVO Backup & Replication inventory allows you to back up VMware VMs from HPE Alletra storage snapshots. For details, refer to https://www.hpe.com/us/en/hpe-alletra.html.

**Supported versions:**

- HPE Alletra OS 9.4

**Supported devices:**

- HPE Alletra 5000 series
- HPE Alletra 6000 series
- HPE Alletra 9000 series
- HPE Alletra MP B10000

## NetApp

The feature supports the following NetApp Operating Systems:

- ONTAP v9.6 and higher

The following NetApp data storage arrays are supported:

- AFF
- FAS

The NetApp user must have the *admin* role.

# Deployment Requirements

NAKIVO Backup & Replication can be deployed as a virtual appliance (VA) or installed directly onto a supported machine or network-attached storage (NAS). Below is the list of deployment requirements and performance-related recommendations.


- Hardware
  - VM or Physical Machine
  - Network Attached Storage
  - Scalability and UI Performance
- Operating Systems
- Networking Requirements
  - Required TCP Ports
  - “External Resources” on page 149
  - Network Conditions
- Web Browsers


## Hardware

### VM or Physical Machine

NAKIVO Backup & Replication can be installed on a machine with the following minimum hardware characteristics:

Director and Onboard Transporter:

- **CPU**: x86-64, 2 cores
- **RAM**: 4 GB + 250 MB for each concurrent task
  - For SaaS Backup Repository-related activities:
    - additional 2 GB
    - additional 150 MB for each concurrent Java Transporter task
- **Free space**: 10 GB

Transporter only:

- **CPU**: x86-64, 2 cores
- **RAM**: 2 GB + 250 MB for each concurrent task
  - For SaaS Backup Repository-related activities:
    - additional 2 GB
    - additional 150 MB for each concurrent Java Transporter task
- **Free space**: 5 GB

## Network Attached Storage

NAKIVO Backup & Replication can be installed on supported NAS with the following minimum hardware characteristics:

Director and Onboard Transporter:

- **CPU**: x86-64, 2 cores
- **RAM**: 1 GB + 250 MB for each concurrent task
  - For SaaS Backup Repository-related activities:
    - minimum total RAM: 4 GB
    - additional 150 MB for each concurrent Java Transporter task
- **Free space**: 10 GB

Transporter only:

- **CPU**: x86-64, 2 cores
- **RAM**: 512 MB
  - For SaaS Backup Repository-related activities:
    - minimum total RAM: 4 GB
    - additional 150 MB for each concurrent Java Transporter task
- **Free space**: 5 GB

**Note**

- To ensure smoother operation, at least 8 GB of RAM is recommended, particularly if you expect to run concurrent backup or recovery tasks.
- Onboard Transporters installed on NAS devices with ARM CPU do not support VMware infrastructures. Refer to Transporter Does Not Support VMware vSphere for a solution.

Supported NAS Devices

- **Synology**: For a full list of supported models, refer to "Supported Synology NAS Devices" on page 163
- **QNAP**: For a full list of supported models, refer to "Supported QNAP NAS Devices" on page 155
- **ASUSTOR**: For a full list of supported models, refer to "Supported ASUSTOR NAS Devices" on page 150
- **NETGEAR**: For a full list of supported. For a full list of supporter models, refer to "Supported NETGEAR NAS Devices" on page 152.
- **Western Digital**: For a full list of supported models, refer to "Supported Western Digital NAS Devices" on page 168.

## Generic ARM-based NAS devices

The device for installing NAKIVO Backup & Replication should meet the following requirements:

- Single-board computer with ARMv7/ARMv8 CPU

  > **Note**
  > ARM-based NetGear NAS devices are not supported.

  32/64-bit Linux-based OS supported by NAKIVO Backup & Replication
- Minimum 16 GB of onboard memory or microSD card for OS & software installation
- RAM: minimum 512 MB for Transporter-only installation; minimum 1 GB for full installation
- Separate microSD/HDD/SSD card for Repository storage
- Open ports for Director and Transporter (see Required TCP Ports)
- Enabled SSH protocol
- Active network connection

## Scalability and UI Performance

For optimal user interface performance, it's important to allocate an appropriate amount of resources to your NAKIVO Backup & Replication instance. The following are the guidelines for allocating RAM to your instance based on the number of jobs created by this instance:

1. Up to 29 jobs: 2 GB of allocated RAM
2. 30–49 jobs: 4 GB of allocated RAM
3. 50–99 jobs: 8 GB of allocated RAM
4. 100–199 jobs: 16 GB of allocated RAM
5. 200+ jobs: 20 GB of allocated RAM

> **Note**
> The above guidelines refer to both active and disabled jobs.

If your instance has less than the recommended amount of allocated RAM for the respective number of jobs, consider adding more resources to the machine hosting the instance.

The machines used to open product web UI should meet the following requirements:

- **Processor**: 1.5 GHz or higher
- **RAM**: 1 GB or more
- **Display resolution**: 1366x768 pixels or higher
- **Web browser**: Mozilla Firefox or Google Chrome
- Cookies, Javascript and images must be enabled in the web browser.

# Operating Systems

NAKIVO Backup & Replication can be installed on the following operating systems:

**Note**

- SELinux module must be disabled to install NAKIVO Backup & Replication on Linux.
- Installation on Windows Core is currently not supported.

**Windows**

- Microsoft Windows Server 2025 (x64)
- Microsoft Windows Server 2022 (21H2) (x64)
- Microsoft Windows Server 20H2 (x64)
- Microsoft Windows Server 2019 Standard (x64)
- Microsoft Windows Server 2016 Standard (x64)
- Microsoft Windows 11 (x64)
  - 21H2
  - 22H2
  - 23H2
- Microsoft Windows 10 Enterprise (x64)
- Microsoft Windows 10 Home (x64)
- Microsoft Windows 10 Professional (x64)

**Linux**

- Rocky Linux 9.5 (64-bit)
- Rocky Linux 9.4 (64-bit)
- Rocky Linux 9.3 (64-bit)
- Rocky Linux 9.2 (64-bit)
- Rocky Linux 9.1 (64-bit)
- Rocky Linux 9.0 (64-bit)
- Rocky Linux 8.9 (64-bit)
- Rocky Linux 8.8 (64-bit)
- Rocky Linux 8.7 (64-bit)
- Rocky Linux 8.6 (64-bit)
- Rocky Linux 8.5 (64-bit)
- Rocky Linux 8.4 (64-bit)
- Rocky Linux 8.3 (64-bit)
- Debian 12.10 (64-bit)
- Debian 12.9 (64-bit)
- Debian 12.8 (64-bit)
- Debian 12.7 (64-bit)
- Debian 12.6 (64-bit)

- Debian 12.5 (64-bit)
- Debian 12.4 (64-bit)
- Debian 12.3 (64-bit)
- Debian 12.2 (64-bit)
- Debian 12.1 (64-bit)
- Debian 11.7 (64-bit)
- Debian 11.6 (64-bit)
- Debian 11.5 (64-bit)
- Debian 11.4 (64-bit)
- Debian 11.3 (64-bit)
- Debian 11.2 (64-bit)
- Debian 11.1 (64-bit)
- Debian 11.0 (64-bit)
- Debian 10.13 (64-bit)
- Debian 10.12 (64-bit)
- Debian 10.11 (64-bit)
- Debian 10.10 (64-bit)
- Debian 10.9 (64-bit)
- Debian 10.8 (64-bit)
- Debian 10.7 (64-bit)
- Debian 10.6 (64-bit)
- Debian 10.5 (64-bit)
- Debian 10.4 (64-bit)
- Debian 10.3 (64-bit)
- Debian 10.2 (64-bit)
- Debian 10.1 (64-bit)
- Oracle Linux 9.5 (64-bit)
- Oracle Linux 9.4 (64-bit)
- Oracle Linux 9.3 (64-bit)
- Oracle Linux 9.2 (64-bit)
- Oracle Linux 9.1 (64-bit)
- Oracle Linux 9.0 (64-bit)
- Oracle Linux 8.9 (64-bit)
- Oracle Linux 8.8 (64-bit)
- Oracle Linux 8.7 (64-bit)
- Oracle Linux 8.6 (64-bit)

- Oracle Linux 8.5 (64-bit)
- Oracle Linux 8.4 (64-bit)
- Oracle Linux 8.3 (64-bit)
- Oracle Linux 8.2 (64-bit)
- Oracle Linux 8.1 (64-bit)
- Oracle Linux 8.0 (64-bit)
- Oracle Linux 7.9 (64-bit)
- Oracle Linux 7.8 (64-bit)
- Oracle Linux 7.7 (64-bit)
- Oracle Linux 7.6 (64-bit)
- Oracle Linux 7.5 (64-bit)
- Oracle Linux 7.4 (64-bit)
- Ubuntu 24.10 STS (x64)
- Ubuntu 24.04 Server LTS (x64)
- Ubuntu 22.04 Server LTS (x64)
- Ubuntu 20.04 Server (x64)
- Ubuntu 18.04 Server (x64)
- Ubuntu 16.04 Server (x64)
- SUSE Linux Enterprise Server 15 SP6 (x64)
- SUSE Linux Enterprise Server 15 SP5 (x64)
- SUSE Linux Enterprise Server 15 SP4 (x64)
- SUSE Linux Enterprise Server 15 SP3 (x64)
- SUSE Linux Enterprise Server 15 SP2 (x64)
- SUSE Linux Enterprise Server 15 SP1 (x64)
- SUSE Linux Enterprise Server 12 SP5 (x64)
- SUSE Linux Enterprise Server 12 SP4 (x64)
- SUSE Linux Enterprise Server 12 SP3 (x64)
- Red Hat Enterprise Linux 9.5 (x64)
- Red Hat Enterprise Linux 9.4 (x64)
- Red Hat Enterprise Linux 9.3 (x64)
- Red Hat Enterprise Linux 9.2 (x64)
- Red Hat Enterprise Linux 9.1 (x64)
- Red Hat Enterprise Linux 9.0 (x64)
- Red Hat Enterprise Linux 8.10 (x64)
- Red Hat Enterprise Linux 8.6 (x64)
- Red Hat Enterprise Linux 8.5 (x64)

- Red Hat Enterprise Linux 8.4 (x64)
- Red Hat Enterprise Linux 8.3 (x64)
- Red Hat Enterprise Linux 8.2 (x64)
- Red Hat Enterprise Linux 8.1 (x64)
- Red Hat Enterprise Linux 8.0 (x64)
- Red Hat Enterprise Linux 7.9 (x64)
- Red Hat Enterprise Linux 7.8 (x64)
- Red Hat Enterprise Linux 7.7 (x64)
- Red Hat Enterprise Linux 7.6 (x64)
- Red Hat Enterprise Linux 7.5 (x64)
- Red Hat Enterprise Linux 7.4 (x64)
- CentOS Stream 10 (x64)
- CentOS Stream 9 (x64)
- CentOS Stream 8 (x64)
- CentOS Linux 8.5 (x64)
- CentOS Linux 8.4 (x64)
- CentOS Linux 8.3 (x64)
- CentOS Linux 8.2 (x64)
- CentOS Linux 8.1 (x64)
- CentOS Linux 8.0 (x64)
- CentOS Linux 7.9 (x64)
- CentOS Linux 7.8 (x64)
- CentOS Linux 7.7 (x64)
- CentOS Linux 7.6 (x64)
- CentOS Linux 7.5 (x64)
- CentOS Linux 7.4 (x64)
- CentOS Linux 7.3 (x64)
- CentOS Linux 7.2 (x64)
- CentOS Linux 7.1 (x64)
- CentOS Linux 7.0 (x64)
- AlmaLinux 9.5 (64-bit)
- AlmaLinux 9.4 (64-bit)
- AlmaLinux 9.3 (64-bit)
- AlmaLinux 9.2 (64-bit)
- AlmaLinux 9.1 (64-bit)
- AlmaLinux 9.0 (64-bit)

- AlmaLinux 8.7 (64-bit)

**NAS**

- ASUSTOR ADM v3.5 - v5.0.0.RA82
- Netgear ReadyNAS OS v6.9 - v6.10.9
- Synology DSM v6.0 - v7.2.2
- QNAP QTS v4.3 - v5.2.4.3092
- QNAP QuTS Hero h4.5.3 - h5.2.0
- QNAP QuTScloud v4.5.1 - c5.1.0
- WD MyCloud v5
- TrueNAS CORE v13.0-U6.2 - v13.3-U1

**Supported Operating System Localizations**

NAKIVO Backup & Replication can be installed on a supported OS with the following OS localization:

- English
- Italian
- German
- French
- Spanish

# Networking Requirements

## Required TCP Ports

NAKIVO Backup & Replication requires the following TCP ports to be open for a successful operation:

| TCP Port (Default) | Where | Description |
|---|---|---|
| **NAKIVO Backup & Replication** | | |
| 80<br>8080 | Director | Used to redirect to the secure TCP port (HTTPS) 4443. |
| 4443 | Director | Used to access the Director web UI. Must be opened on the Director machine. |

| | | |
|---|---|---|
| 9446 | Transporter | Used by Director and Transporters to communicate with the Transporter. Must be opened on the Transporter machine. |
| 9448 - 10000 | Transporter | Used by Director and Transporters for cross-Transporter data transfer. Must be opened on the Transporter machine. |
| **VMware** | | |
| 443 | vCenter Server, ESXi host | Used by Director and Transporters to access VMware infrastructure. Must be opened on vCenter Servers and ESXi hosts. |
| 902 | ESXi host | Used by Transporters to access VMware infrastructure. Must be opened on ESXi hosts. |
| 3260 | Transporter, ESXi host | Used by Proxy Transporters to access VMware infrastructure during a Flash boot. Must be opened on the Transporter machine and the ESXi host used as the target for a Flash boot. |
| **Hyper-V** | | |
| 137 - 139 | Hyper-V hosts | Used by Director to upload files and install configuration service. Must be opened on Hyper-V servers. |
| 445 | Hyper-V hosts | Used by Director to upload files and install configuration service. |
| 5986 (opens automatically) | Hyper-V hosts | Used by Transporter to add a host to inventory and establish a connection with it. |

| | | |
|---|---|---|
| 9445 (opens automatically) | Hyper-V hosts | Used by Director to upload files and install configuration service. Must be opened on Hyper-V host if NAKIVO Backup & Replication is installed on a host and this host is added to inventory simultaneously. |
| 9446 (opens automatically) | Hyper-V hosts | Used by Director and Transporters to communicate with the Transporter. Used by Transporters for cross-Transporter data transfer. Must be opened on the Transporter machine. |
| 9448 -10000 (opens automatically) | Hyper-V hosts | Used by Transporters for cross-Transporter data transfer. Must be opened on the Transporter machine. |
| **Physical machine (Windows)** | | |
| 445 | Windows machine | Used by Director to upload files and install configuration service via SMB. |
| 9446 (opens automatically) | Windows machine | Used to create the Transporter installed by default. |
| 9448 -10000 (opens automatically) | Windows machine | Used by Transporters for cross-Transporter data transfer. Must be opened on the Transporter machine. |
| **Physical machine (Linux)** | | |
| 22 | Linux machine | Used by Director to access a Linux physical machine via SSH. |
| 9446 | Linux machine | Used to create the Transporter installed by default. |
| 9448 -10000 | Linux machine | Used by Transporters for cross-Transporter data transfer. Must be opened on the Transporter machine. |

| | | |
|---|---|---|
| **Real-time replication** | | |
| 33088 | ESXi host | Used by the I/O Filter (source daemon installed on ESXi host) for receiving requests (or commands) from the Director. |
| 33089 | Transporter, ESXi host | Used by the Journal Service (installed on the target Transporter) for receiving requests (or commands) from the Director and the I/O Filter. (Outgoing ports on the I/O Filter). |
| 33090 | Transporter, ESXi host | Used by the Journal Service (installed on the target Transporter) for receiving changed data from the I/O Filter. (Outgoing ports on the I/O Filter). |
| 33091 | Transporter, ESXi host | Used by the Journal Service (installed on the target Transporter) for receiving heartbeat requests from the I/O Filter. (Outgoing ports on the I/O Filter). |
| **Nutanix AHV** | | |
| 9440 | Nutanix REST APIs | Used to manage traffic to and from Prism Web Console and SSH console for nCLI. Prism is the central management interface for Nutanix clusters that offers an end-to-end management solution for virtualized datacenter environments that streamlines and automates common workflows. |
| **Proxmox VE** | | |
| 8006 | Proxmox VE hosts | Proxmox VE API |
| 22 | Proxmox VE hosts | Used by Director to manage (install/uninstall/upgrade) Transporter and check temp folder's existence via SSH. |

| | | |
|---|---|---|
| 9446 | Proxmox VE hosts | Used by Director and Transporters to communicate with the Transporter. |
| 9448-10000 | Proxmox VE hosts | Used by Transporters for cross-Transporter data transfer. |

Make sure that the specified port is open in your firewall.

## External Resources

NAKIVO Backup & Replication Director should have access to the following external resources:

| | |
|---|---|
| https://www.nakivo.com/nbr_update_check-/ga/<br>https://52.8.1.104:443/<br>https://upload.nakivo.com/ | Used for updating, checking, and supporting bundle upload. |
| https://nbr-va.s3.us-west-2.amazonaws.com/ | Used for downloading VA. |
| https://d96i82q710b04.cloudfront.net/ | Used for downloading installers. |

## Network Conditions

NAKIVO Backup & Replication has been tested to work in the following minimal network conditions:

- **Latency (RTT)**: Up to 250 ms
- **Packet loss**: Up to 1 %
- **Bandwidth**: 1 Mb/s or higher
- **ICMP ping traffic:** It should be allowed on all hosts on which NAKIVO Backup & Replication components are installed as well as on all source and target hosts.

# Web Browsers

NAKIVO Backup & Replication user interface can be accessed through the following web browsers:

- Google Chrome: Version 80 or higher
- Mozilla Firefox: Version 74 or higher

# Supported ASUSTOR NAS Devices

NAKIVO Backup & Replication supports the following ASUSTOR NAS devices :

## Director and Onboard Transporter

- AS3102T
- AS3102T v2
- AS3104T
- AS3202T
- AS3204T
- AS3204T v2
- AS4002T
- AS4004T
- AS5202T
- AS5304T
- AS5002T
- AS5004T
- AS5008T
- AS5010T
- AS6102T
- AS6104T
- AS6302T
- AS5102T
- AS5104T
- AS5108T
- AS5110T
- AS6202T
- AS6204T
- AS6208T
- AS6210T
- AS6404T
- AS6204RS / AS6204RD
- AS-609RS / AS-609RD
- AS7004T
- AS7008T
- AS7010T
- AS6212RD
- AS7009RD / AS7009RDX
- AS7012RD / AS7012RDX
- AS-602T
- AS-604RS / AS-604RD
- AS-604T

- AS-606T
- AS-608T
- AS6508T
- AS6510T
- AS7110T
- AS6602T
- AS6604T
- AS7116RDX
- AS7112RDX
- AS1102T
- AS1104T
- AS3302T
- AS3304T
- AS6504RD
- AS6504RS
- AS6512RD
- AS5402T
- AS5404T
- FS6706T
- FS6712X
- AS1102T Lite
- AS1104T Gen 2
- AS3304T Gen 2
- AS6702T
- AS6704T
- AS3302T v2
- AS3304T v2
- AS5004U
- AS6706T
- AS7212RDX
- AS7216RDX

## Transporter Only

- AS1002T
- AS1002T v2
- AS1004T
- AS1004T v2

For minimum hardware requirements, refer to “Network Attached Storage” on page 139.

# Supported NETGEAR NAS Devices

NAKIVO Backup & Replication supports the following NETGEAR NAS devices:

## Director and Onboard Transporter

- RN51600
- RN51661D
- RN51661E
- RN51662D
- RN51662E
- RN51663D
- RN51663E
- RN51664E
- ReadyNAS 524X
- ReadyNAS 526X
- ReadyNAS 528X
- ReadyNAS 626X
- ReadyNAS 628X
- RN716X
- RN628X
- RN626X
- RN528X
- RN526X
- RN524X
- RN31600
- RN31661D
- RN31661E
- RN31662D
- RN31662E
- RN31663D
- RN31663E
- RN31664E
- ReadyNAS 422
- ReadyNAS 424
- ReadyNAS 426
- ReadyNAS 428
- RN516
- RN426
- RN424
- RN422
- RN31400
- RN31421D

- RN31441D
- RN31441E
- RN31442D
- RN31442E
- RN31443D
- RN31443E
- RN316
- RN31200
- RN31211D
- RN31212D
- RN31221D
- RN31221E
- RN31222D
- RN31222E
- RN31223D
- RN314
- RN312
- RN322121E
- RN322122E
- RN322123E
- RN322124E
- RN32261E
- RN32262E
- RN32263E
- RN4220S
- RN4220X
- RN422X122
- RN422X123
- RN422X124
- RN422X62E
- RN422X63E
- RN422X64E
- RR2304
- RN21241D
- RN21241E
- RN21243D
- RN21243E
- RN3130
- RN31342E
- RN3138
- RN3220
- RR2312
- RR3312

- RN4220
- RR4312X
- RR4312S
- RR4360X
- RR4360S
- RN202
- RN204
- RN212
- RN214
- RN2120

## Transporter Only

- RN102
- RN10200
- RN10211D
- RN10221D
- RN10222D
- RN10223D
- RN104
- RN10400
- RN10421D
- RN10441D
- RN10442D
- RN10443D

For minimum hardware requirements, refer to “Network Attached Storage” on page 139

# Supported QNAP NAS Devices

NAKIVO Backup & Replication supports the following QNAP NAS Devices:

## Director and Onboard Transporter

- HS-251+
- HS-453DX
- TS-251
- TS-251+
- TS-251A
- TS-251B
- TS-253Be
- TS-328
- TS-332X
- TS-351
- TS-431P
- TS-431P2
- TS-431X
- TS-431X2
- TS-431XeU
- TS-432XU
- TS-432XU-RP
- TS-451
- TS-451+
- TS-451A
- IS-400 Pro
- IS-453S
- TBS-453A
- TBS-453DX
- TS-128A
- TS-131P
- TS-231P
- TS-231P2
- TS-253 Pro
- TS-253A
- TS-253B
- TS-228A
- TS-451U
- TS-453 mini
- TS-453 Pro
- TS-453A
- TS-453B

- TS-453Be
- TS-453Bmini
- TS-453BT3
- TS-453BU
- TS-453BU-RP
- TS-453U
- TS-453U-RP
- TS-463U
- TS-463U-RP
- TS-463XU
- TS-463XU-RP
- TS-473
- TS-473A
- TS-563
- TS-653 Pro
- TS-653A
- TS-653B
- TS-653B
- TS-673
- TS-673A
- TS-677
- TS-832X
- TS-832XU
- TS-832XU-RP
- TS-853 Pro
- TS-853A
- TS-853BU
- TS-853BU-RP
- TS-853U
- TS-853U-RP
- TS-863U
- TS-863U-RP
- TS-863XU
- TS-863XU-RP
- TS-873
- TS-873U
- TS-873U-RP
- TS-877
- TS-877XU
- TS-877XU-RP
- TS-883XU
- TS-883XU-RP
- TS-932X

- TS-963X
- TS-977XU
- TS-977XU-RP
- TS-983XU
- TS-983XU-RP
- TS-1232XU
- TS-1232XU-RP
- TS-1253BU
- TS-1253BU-RP
- TS-1253U
- TS-1253U-RP
- TS-1263U-RP
- TS-1263U
- TS-1263XU
- TS-1263XU-RP
- TS-1273U
- TS-1273U-RP
- TS-1277
- TS-1277XU-RP
- TS-1283XU-RP
- TS-1635AX
- TS-1673U
- TS-1673U-RP
- TS-1677X
- TS-1677XU-RP
- TS-1683XU-RP
- TS-1685
- TS-2477XU-RP
- TS-2483XU-RP
- TVS-463
- TVS-471
- TVS-472XT
- TVS-473e
- TVS-473
- TVS-663
- TVS-671
- TVS-672XT
- TVS-673
- TVS-673e
- TVS-682
- TVS-682T
- TVS-863
- TVS-863+

- TVS-871
- TVS-871T
- TVS-871U-RP
- TVS-872XT
- TVS-872XU
- TVS-872XU-RP
- TVS-873e
- TVS-873
- TVS-882
- TVS-882T
- TVS-882ST2
- TVS-882BR
- TVS-882BRT3
- TVS-882ST3
- TVS-951X
- TVS-972XU
- TVS-972XU-RP
- TVS-1271U-RP
- TVS-1272XU-RP
- TVS-1282
- TVS-1282T
- TVS-1282T3
- TVS-1582TU
- TVS-1672XU-RP
- TVS-2472XU-RP
- SS-EC1279U-SAS-RP
- SS-EC1879U-SAS-RP
- SS-EC2479U-SAS-RP
- TDS-16489U
- TES-3085U
- TES-1885U
- TS-EC880U
- TS-EC880U R2
- TS-EC1280U
- TS-EC1280U R2
- TS-EC1680U
- TS-EC1680U R2
- TS-EC2480U
- TS-EC2480U R2
- TVS-EC880
- TVS-EC1080
- TVS-EC1080+
- TVS-EC1280U-SAS-RP

- TVS-EC1580MU-SAS-RP
- TVS-EC1680U-SAS-RP
- TVS-EC1680U-SAS-RP R2
- TVS-EC2480U-SAS-RP
- TVS-EC2480U-SAS-RP R2
- TVS-EC2480U-SAS-RP R2
- TVS-EC1580MU-SAS-RP R2
- TVS-EC1280U-SAS-RP R2
- TDS-16489U-SE1-R2
- TDS-16489U-SE2-R2
- TDS-16489U-SF2-R2
- TDS-16489U-SF3-R2
- TS-2888X-W2195-512G
- TS-2888X-W2195-256G
- TS-2888X-W2195-128G
- TS-2888X-W2175-512G
- TS-2888X-W2175-256G
- TS-2888X-W2175-128G
- TS-2888X-W2145-512G
- TS-2888X-W2145-256G
- TS-2888X-W2145-128G
- TS-2888X-W2133-64G
- TS-2888X-W2123-32G
- ES2486dc
- TS-1886XU-RP
- TS-230
- TS-251C
- TS-251D
- TS-253D
- TS-451DeU
- TS-453D
- TS-653B
- TS-653D
- TS-h1277XU-RP
- TS-h1283XU-RP
- TS-h977XU-RP
- TVS-472XT-PT
- TVS-672N
- TVS-872N
- TVS-EC2480U-SAS-RP-R2
- TS-431P3
- TS-231P3
- TS-431X3

- TS-h686-D1602
- TS-h886-D1622
- TS-873AU
- TS-873AU-RP
- TS-1273AU-RP
- TS-1673AU-RP
- TS-932PX
- GM-1001
- TS-432PXU
- TS-432PXU-RP
- TS-832PXU
- TS-832PXU-RP
- TS-1232PXU-RP
- TS-451D2
- TS-h2490FU-7232P-64G
- TS-h2490FU-7302P-128G
- TS-h1886XU-RP
- TS-h1683XU-RP
- TS-h2483XU-RP
- TVS-h1288X
- TVS-h1688X
- TS-h973AX-8G
- TS-h973AX-32G
- TS-832PX
- TS-h3088XU-RP-W1270-64G
- TS-h3088XU-RP-W1250-32G
- TS-453DU-4G
- TS-873A
- TS-EC879U-RP
- TS-831X-4G
- TS-831X-8G
- TS-831X-16G
- TS-EC879U-RP
- TS-h987XU-RP
- TS-h3077AFU
- TS-h1277AXU-RP
- TS-h1677AXU-RP
- TVS-h674T
- TVS-h874T
- TBS-574TX
- TS-AI642
- TS-855X
- TS-879U-RP

- TS-262
- TS-462
- TS-h1887XU-RP
- TS-1655
- TS-855eU
- TS-473A
- TBS-h574TX
- TDS-h2489FU-4309Y
- TDS-h2489FU-4314
- TS-130
- TS-233
- TS-253E
- TS-264
- TS-364
- TS-410E
- TS-433
- TS-453E
- TS-464
- TS-464eU
- TS-464U-4G
- TS-464U-RP-4G
- TS-473A-8G
- TS-664
- TS-864eU
- TS-864eU-RP
- TS-h1090FU
- TS-h1290FX
- TS-h1677XU-RP
- TS-h2477XU-RP
- TS-hx87XU-RP
- TS-i410X
- TVS-672X-i3-8G
- TVS-675
- TVS-872X-i3-8G
- TVS-h474
- TVS-h674
- TVS-h874
- TS-216G
- TS-432X
- TS-632X
- TS-433eU
- TS-h2287XU-RP
- TS-h765eU

- TS-h1277AFX
- TS-1264U-RP
- TDS-h2489FU
- TS-873AeU-RP-4G
- TS-435XeU-4G
- ES2486dc
- ES1686dc R2
- TS-h2287XU-RP
- TS-h2490FU
- TS-433-4G

For minimum hardware requirements, refer to “Network Attached Storage” on page 139.

# Supported Synology NAS Devices

NAKIVO Backup & Replication supports the following Synology NAS devices:

## Director and Onboard Transporter

- FS3017
- FS2017
- FS1018
- RS18017xs+
- RS18016xs+
- RS10613xs+
- RS4017xs+
- RS3618xs
- RS3617xs+
- RS3617RPxs
- RS3617xs
- RS3614xs+
- RS3614RPxs
- RS3614xs
- RS3413xs+
- RS3412RPxs
- RS3412xs
- RS3411RPxs
- RS3411xs
- RS2818RP+
- RS2418RP+
- RS2418+
- RS2416RP+
- RS2416+
- RS2414RP+
- RS2414+
- RS2212RP+
- RS2212+
- RS2211RP+
- RS2211+
- RS1619xs+
- RS1219+
- RS818RP+
- RS818+
- RS816
- RS815RP+
- RS815+

- RS815
- RS814RP+
- RS814+
- RS814
- RS812RP+
- RS812+
- RS810RP+
- RS810+
- RC18015xs+
- DS3617xs
- DS3615xs
- DS3612xs
- DS3611xs
- DS3018xs
- DS2415+
- DS2413+
- DS2411+
- DS2015xs
- DS1819+
- DS1817+
- DS1817
- DS1815+
- DS1813+
- DS1812+
- DS1618+
- DS1517+
- DS1517
- DS1515+
- DS1515
- DS1513+
- DS1512+
- DS1511+
- DS918+
- DS916+
- DS718+
- DS716+II
- DS716+
- DS715
- DS713+
- DS712+
- DS710+
- DS418
- DS418play

- DS418j
- DS416
- DS416play
- DS415+
- DS414
- DS412+
- DS411+II
- DS411+
- DS218+
- DS218
- DS218play
- DS216+II
- DS216+
- DS216play
- DS215+
- DS214+
- DS118
- DS116
- DS1019+
- DS2419+
- DS420+
- DS420j
- DS620slim
- DS720+
- DS920+
- DS925+
- FS3400
- FS3600
- FS6400
- RS819
- RS820+
- RS820RP+
- SA3200D
- SA3400
- SA3600
- DS1520+
- DS1621+
- DS1621xs+
- DS1821+
- RS1221+
- RS1221RP+
- RS2421+
- RS2421RP+

- RS4021xs+
- RS2423+
- RS2423RP+
- DS124
- DS224+
- DS223j
- DS723+
- DS273+
- DS223
- DS1823xs+
- DS1821Plus
- DS423+
- DS423
- SA3400D
- SA3610
- SA3410
- DS120j
- DS1522
- DS1522+
- DS220j
- DS2422+
- DS3621xs+
- DS3622xs+
- DS923+
- FS2500
- FS3410
- RS2821RP+
- RS3621RPxs
- RS3621xs+
- RS422+
- RS822+
- RS822RP+
- DS1525+
- DS1825+
- DS725+
- DS425+
- RS2825RP+
- DS225+

## Transporter Only

- RS217
- RS214

- DS416slim
- DS416j
- DS414slim
- DS414j
- DS218j
- DS216
- DS216j
- DS215j
- DS214
- DS213j
- DS115
- DS114
- DS220j
- DS419slim

**Important**

For minimum hardware requirements, refer to “Network Attached Storage” on page 139.

# Supported Western Digital NAS Devices

NAKIVO Backup & Replication supports the following Western Digital NAS devices for Director and Onboard installation:

- MyCloud DL2100
- MyCloud DL4100
- MyCloud PR2100
- MyCloud PR4100

For minimum hardware requirements, refer to “Network Attached Storage” on page 139.

# Feature Requirements

Some NAKIVO Backup & Replication features require certain conditions to work properly. To learn about the limitations of NAKIVO Backup & Replication, refer to the **Feature Limitations** section of the latest Release Notes. The requirements for product features are listed below.

- App-Aware Mode
- Auto-Update
- Backup Encryption
- Backup Immutability
- Backup Malware Scan
- Cross-Platform Recovery
- Direct Connect
- Encrypted Backup Repository
- External Database
- Federated Repository
- File Recovery
- File Share Backup
- File System Indexing
- Generic S3-Compatible Object Storage
- Granular User Notifications
- Hot Add
- Merge Jobs
- Monitoring
- MSP Console
- Native Tape Support
- Object Recovery and Log Truncation for Microsoft Exchange
- Object Recovery and Log Truncation for Microsoft SQL Server
- Object Recovery for Microsoft Active Directory
- Permanent VM Agent
- Proxmox VE Support
- Real-Time Replication for VMware
- Site Recovery
- Transporter Load Balancing
- VM Limitation for Multi-Tenancy

## App-aware Mode

To enable application awareness for source objects, make sure the following requirements are met:

### VMware

- VMware Tools should be installed, running, and up to date on all source VMs.
- VMware Snapshot Provider service must be running on source VMs.

## Auto-Update

Auto-update is available for instances of NAKIVO Backup & Replication installed on the following operating systems:

- Linux
- Windows

> **Note**
> Auto-update is not supported for NAS systems.

Auto-update is available for the following types of Transporters:

- Auto-injected Transporters on Linux (includes physical Transporters)
- Manually installed Transporters on Linux
- Auto-injected Transporters on Windows (includes Hyper-V and physical Transporters)
- Manually installed Transporters on Windows
- Auto-injected Transporters in AWS (Linux)
- Auto-injected Transporter in VMware (Linux)

> **Note**
> Manually installed Transporters on Linux and Windows must be v10.8 or newer to support auto-update.

Before initiating an auto-update, make sure that the following conditions are met:

- If updating a manually installed Transporter on Linux or Windows, make sure that you have configured a **Master Password** for the Transporter in the Managing Credentials menu.
- At least 1.5 GB of free space is available on the machine on which the full solution is installed.
- If you have a perpetual license, your Maintenance & Support period is active. You can verify this on the product Licensing page.

We encourage you to update your version of NAKIVO Backup & Replication to the latest build available in order to maintain the most stable and secure experience.

**Important**

It is highly recommended that you update your product from any version within its support period. Check NAKIVO Backup & Replication Support Lifecycle Policy for more details.

**Note**

To maintain maximum security on your Linux OS, always keep it up-to-date; however, if necessary, you may disable auto-updates by editing the `/etc/apt/apt.conf.d/20auto-upgrades` file.

For more details, refer to https://linuxconfig.org/disable-automatic-updates-on-ubuntu-22-04-jammy-jellyfish-linux.

## Backup Encryption

With the **Backup encryption** feature, you can customize your protection preferences by encrypting backup data created by the product, managing passwords, and performing recovery from encrypted backups.

To use the **Backup encryption** feature, make sure you meet the following requirements:

**Supported Data Types**:

- The following backups stored on the Director can be encrypted:
  - Workloads
  - System configurations
  - FSI data

**Supported Job Types**:

- Backup jobs
  - Exceptions:
    - Microsoft 365 backup jobs with SaaS repositories selected as the only target repository
    - Oracle RMAN backup jobs
  - If a job is edited and the encryption details are changed, the next job run is performed in active full mode.
  - Data encryption must occur after data compression.
  - You cannot apply different passwords to the recovery points within the same chain.
- File share backup jobs
- Backup copy jobs
  - Backup copies can be encrypted.
  - Encrypted backups can be copied without decrypting them (tape to repository backup copy jobs).
  - It is possible to have different passwords within the same chain.

**Supported Destinations**:

The following repository types are supported for storing encrypted backups:

- Repositories of the **Incremental with full** data storage type:
  - Backup repository types:
    - Local folders
    - CIFS shares
    - NFS shares
    - Amazon EC2
    - Amazon S3
    - Generic S3 compatible object storages
    - Wasabi
    - Deduplication appliances
    - Azure Blob (if available)
    - Backblase B2 (if available)

      **Note**

      SaaS cannot be used for storing encrypted backups.
- Federated repositories
- Tape
  - Supported for the data copied to tape media
  - Not supported if hardware encryption is enabled on a target tape device

**Supported Recovery Jobs**:

- Full recovery jobs and Flash VM boot jobs are supported.
  - Decryption takes place before decompression.

**Supported Recovery Sessions**:

The following recovery sessions are supported to recover encrypted backups:

- Object level recovery
- Universal object recovery
- Microsoft 365 recovery
  - Explicit user agreement is required for sending decryption details to the target endpoint.

**Supported Key Management Service**:

Supported KMS: AWS Key Management Service (AWS KMS) is used for encrypting a password hash.

- The AWS account must be discovered for KMS integration.
- For encrypting a password hash:
  - The existing KMS cryptographic key must be selected or a new key must be generated.
  - Global KMS usage must be enabled.
  - The discovered AWS account must have sufficient permissions to use KMS.

**Note**

The list of required AWS permissions for creating the (AWS) KMS Keys:

- *kms:ListKeys*
- *kms:ListAliases*
- *kms:CreateKey*
- *kms:CreateAlias*
- *kms:PutKeyPolicy*
- *kms:TagResource*
- *tag:GetResources*
- *iam:CreateServiceLinkedRole*
- *iam:ListUsers*
- *iam:ListRoles*
- *iam:GetUser*

Refer to Creating keys - AWS Key Management Service for more details.

**General Requirements**:

- The backup encryption feature is not supported for protecting backups in repositories of the **Forever Incremental** storage type.
- Backup data is protected with *AES 256* block cipher encryption with a 256-bit key length.
- Passwords are saved per tenant and cannot be propagated between tenants.
- User passwords for data encryption must be set per job.
- The **Network acceleration** option is not available if the **Backup encryption** option is enabled.

**Supported Platforms**:

Encryption across all platforms is supported.

**Immutability**:

Immutability can be applied to encrypted backups.

# Backup Immutability

To make backups immutable, do the following with Backup Repositories located in:

- **Amazon S3**, **generic S3-compatible storage**, or **Wasabi**, the following options must be enabled for the buckets where the repository is located:
  - Object Lock
  - Versioning
- **Backblaze B2 Cloud Storage**, **File Lock** (also known as **Object Lock**) must be enabled.
- **Azure Blob Storage**, the following options must be selected for the Azure storage account or container:

- Enable version-level immutability support
- Enable versioning for blobs

**Notes**

- Disable Object Lock retention mode and retention period for the Amazon S3 or Wasabi bucket where the repository is located, as retention settings are set in NAKIVO Backup & Replication during job creation.
- Backing up to Wasabi with Object Lock enabled may take longer compared to when Object Lock is disabled.
- Backup Immutability is not supported for encrypted Backup Repositories.
- Backup Immutability is not supported for MinIO.
- Backup Immutability is not supported for Google Cloud Storage.

- **Local Folder** type of **Backup Repository**, the following conditions must be met:
  - Target file system must support extended attributes modified by **chattr** and **setfattr** commands.
  - The **Backup Repository** type must be **Local Folder**.
  - The **Backup Repository** must have **Store backups in separate files** selected.

    **Note**

    Only Linux OS and NAS OS specified in operating systems are supported.

- **Deduplication Appliance**, the following conditions must be met:
  - The backup repository must be created in the supported duplication appliance. For the purpose of Backup Immutability feature, only the following deduplication appliances are supported:
    - **NEC HYDRAstor**
    - **Dell EMC Data Domain**
  - For **NEC HYDRAstor** do the following:
    - Enable the WORM functionality when creating the deduplication appliance Repository.
    - Select the **Enable backup immutability** checkbox when creating the Repository.
  - For **Dell EMC Data Domain** do the following:
    - Enable the **Retention lock** functionality on the Dell EMC Data Domain system when creating the deduplication appliance Repository.
    - Select the **Enable backup immutability** checkbox when creating the Repository.

**Note**

The root user cannot modify, delete, or override the immutability settings on backups during the defined immutability period.

# Backup Malware Scan

The following recovery jobs with image-based backups are supported:

- VMware vSphere VM recovery
- Microsoft Hyper-V VM recovery
- Nutanix AHV VM recovery
- VMware Cloud Director recovery
- Physical Machine recovery
- Flash boot recovery
- Universal object recovery

The scan server must adhere to the following requirements:

- A supported version of antivirus software must be installed on the scan server.
- iSCSI must be available on the scan server.
- SMB must be enabled on the scan servers using Windows OS.
- SSH must be enabled on the scan servers using Linux OS.
- TCP port *9445* must not be blocked.
- SElinux service must be disabled on the scan servers using Linux OS.
- The NTFS-3G package must be enabled on the scan servers using Linux OS.
- Special permissions must be configured for NAKIVO Backup & Replication recovery service.

The following antivirus software is supported:

- **Windows OS**:
  - Microsoft Windows Defender
    - Antimalware Client Version: 4.10.14393 or higher
    - Engine Version: 1.1.12805 or higher
    - Must support the command line: *Scan -ScanType %type% -File %path% -DisableRemediation -BootSectorScan*
    - See more details here
  - ESET NOD32 Antivirus:
    - Version 14.2.24 or higher
    - Must support the command line: *%path% /clean-mode=None /no-symlink*
    - See more details here
  - Kaspersky Internet Security:
    - Version 2015 (15.0) or higher
    - Must support the command line: *scan %path% -i0*
    - See more details here
  - Sophos Intercept X:

        - Version 2.0.20 or higher
        - Must support the command line: *scan --noui --expand_archives %path%*
        - See more details here
- **Linux OS**:
    - Clam AV:
        - Version clamav-0.99.0 or higher
        - Must support the command line: *clamscan --infected --recursive %path%*
        - See more details here
    - Sophos Protection:
        - Version 1.1.8 or higher
        - Must support the command line: *avscanner %path% --scan-archives*
        - See more details here

**Notes**

- Sophos Protection can be used only for Linux server recovery.
- Sophos AV software is limited to the specific Linux operating systems.
- Antivirus software on Windows OS cannot scan Linux OS backups.
- Scanning process may not detect malware if the antivirus software has the run-time (real-time) protection enabled.
- NAKIVO Backup & Replication delegates scanning and detection to the installed antivirus engine.
- Any filename or path reporting limitations, such as missing Unicode characters in alerts, originate from the antivirus engine, not NAKIVO Backup & Replication.

## Cross-Platform Recovery

The following scenarios are supported if a VM is exported from backup and imported into a different hypervisor:

| | **Target Platforms** | | |
|---|---|---|---|
| **Source Platforms** | **VMware vSphere 8** | **Microsoft Hyper-V 2016/2019/20H1/2022/2025** | **Nutanix AHV 5.10/5.15** |

| | | | |
|---|---|---|---|
| **VMware vSphere 8** | • Windows Server 2016<br>• Windows Server 2019<br>• Windows Server 20H1<br>• Windows Server 2022<br>• Ubuntu Server 18.04<br>• RHEL 7.4<br>• CentOS 7.0 | • Windows Server 2016<br>• Windows Server 2019<br>• Windows Server 20H1<br>• Ubuntu Server 18.04<br>• RHEL 7.4<br>• CentOS 7.0 | • Windows Server 2016<br>• Windows Server 2019<br>• Windows Server 20H1<br>• Windows Server 2022<br>• Ubuntu Server 18.04<br>• RHEL 7.4<br>• CentOS 7.0 |
| **Microsoft Hyper-V 2016/2019/20H1/2022/2025** | • Windows Server 2016<br>• Windows Server 2019<br>• Windows Server 20H1<br>• Windows Server 2022<br>• Ubuntu Server 18.04<br>• RHEL 7.4<br>• CentOS 7.0 | • Windows Server 2016<br>• Windows Server 2019<br>• Windows Server 20H1<br>• Windows Server 2022<br>• Ubuntu Server 18.04<br>• RHEL 7.4<br>• CentOS 7.0 | • Windows Server 2016<br>• Windows Server 2019<br>• Windows Server 20H1<br>• Windows Server 2022<br>• Ubuntu Server 18.04<br>• RHEL 7.4<br>• CentOS 7.0 |

| | | | |
|---|---|---|---|
| **Proxmox VE** | • Windows Server 2016<br>• Windows Server 2019<br>• Windows Server 20H1<br>• Windows Server 2022<br>• Ubuntu Server 18.04<br>• RHEL 7.4<br>• CentOS 7.0 | • Windows Server 2016<br>• Windows Server 2019<br>• Windows Server 20H1<br>• Windows Server 2022<br>• Ubuntu Server 18.04<br>• RHEL 7.4<br>• CentOS 7.0 | • Windows Server 2016<br>• Windows Server 2019<br>• Windows Server 20H1<br>• Windows Server 2022<br>• Ubuntu Server 18.04<br>• RHEL 7.4 |
| **Physical Machines** | • Windows Server 2016<br>• Windows Server 2019<br>• Windows Server 20H1<br>• Windows Server 2022<br>• Ubuntu Server 22.04<br>• RHEL 9.2<br>• CentOS 8.5 | • Windows Server 2016<br>• Windows Server 2019<br>• Windows Server 20H1<br>• Windows Server 2022<br>• Ubuntu Server 18.04<br>• RHEL 7.4<br>• CentOS 7.0 | • Windows Server 2016<br>• Windows Server 2019<br>• Windows Server 20H1<br>• Windows Server 2022<br>• Ubuntu Server 18.04<br>• RHEL 7.4<br>• CentOS 7.0 |

| | | | |
|---|---|---|---|
| **Nutanix AHV 5.10 , 5.15** | • Windows Server 2016<br>• Windows Server 2019<br>• Windows Server 20H1<br>• Windows Server 2022<br>• Ubuntu Server 18.04<br>• RHEL 7.4<br>• CentOS 7.0 | • Windows Server 2016<br>• Windows Server 2019<br>• Windows Server 20H1<br>• Windows Server 2022<br>• Ubuntu Server 18.04<br>• RHEL 7.4<br>• CentOS 7.0 | • Windows Server 2016<br>• Windows Server 2019<br>• Windows Server 20H1<br>• Windows Server 2022<br>• Ubuntu Server 18.04<br>• RHEL 7.4<br>• CentOS 7.0 |

* To run a VM with RHEL or CentOS on Microsoft Hyper-V 2016/2019/20H1/2022/2025, the following option must be configured in grub boot parameters:
*ata_piix.prefer_ms_hyperv=0*
As an alternative, the source machine can be pre-configured with the command below:
*mkinitrd -f -v --with=hid-hyperv --with=hv_utils --with=hv_vmbus --with=hv_storvsc --with=hv_netvsc /boot/initramfs-$(uname -r).img $(uname -r)*

# Direct Connect

## Deployment Scenarios

Direct Connect

To use Direct Connect, the **Transporter** must be installed on one of the following operating systems at the remote site:

- Windows
- Linux
- NAS

Direct Connect supports the following Transporters:

- **Onboard Transporter**
- Installed service

- Transporter as a VMware Appliance

> **Note**
> Direct Connect is not supported for **Onboard Transporters** located on NAS devices.

The following deployment scenarios are supported:

- **Director** and **Transporter**(s) installed at the MSP site, and more than one **Direct ConnectTransporter** installed at each remote site.
- Primary repository at the remote site (managed by one of the **Direct ConnectTransporters**) and a secondary repository at the MSP site.

> **Note**
> **Direct Connect Transporters** support a maximum of 10 assigned repositories.

Additionally, the following requirements must be met:

- A static external IP address is required at the remote environment.
- A single TCP port must be exposed to the Internet at the remote environment.
- Port forwarding must be configured at the remote environment to forward requests from this external port to the deployed/installed **Direct Connect Transporter.**

MSP Direct Connect

To use MSP Direct Connect, the **Direct Connect Transporter** must be installed on one of the following operating systems at the client site:

- Windows
- Linux

The following deployment scenarios are supported:

- **Director** and **Transporter**(s) installed at the MSP site, and more than one **Direct Connect Transporter** installed at each client site.
- Primary repository at the client site (managed by one of the **Direct Connect Transporters**) and a secondary repository at the MSP site.

> **Note**
> **Direct Connect Transporters** support a maximum of 10 assigned repositories.

Additionally, to enable communication between the **MSP Director** and client-side **Direct Connect Transporters**, the following requirements must be met at the MSP site:

- The MSP site must have a publicly accessible IP address for communication between the **MSP Director** and **Direct Connect Transporters** at the client sites.

- Two public TCP ports must be opened on the MSP's firewall/router for the **MSP Director**:
  - TCP Port 4443 (for primary communication)
  - TCP Port 4442 (for additional communication)
- A dedicated public TCP port must be opened for each **MSP Transporter** so the client's **Direct Connect Transporters** can communicate directly with them.

  For example:
  - Public port 10055 for TCP Port 9446 of **MSP Onboard Transporter**
  - Public port 10059 for TCP Port 9446 of **MSP Transporter A**
  - Public port 10060 for TCP Port 9446 of **MSP Transporter B**

Each MSP Transporter that needs to communicate with a client's **Direct Connect Transporter** must have Direct Connect enabled and be assigned a Public IP and a Public Node Port. To do this:

1. Select the **Enable Direct Connect for this Transporter** checkbox in **"Editing Nodes" on page 593**.
2. Specify **Public IP** and **Public Node Port** of the **MSP Transporter**.

> **Note**
> You can enable **Direct Connect** through the **Local Tenant Creation Wizard** by configuring the following parameters:
> - Public IP address
> - Public node port of the MSP Transporter

## License

Direct Connect

Direct Connect is supported under the following licenses:

- Promo
- Trial
- Enterprise Essentials
- Enterprise
- Enterprise Plus
- MSP Pro
- MSP Enterprise
- MSP Enterprise Plus

MSP Direct Connect

MSP Direct Connect is supported under the following licenses:

- Promo
- Trial
- MSP Pro

- MSP Enterprise
- MSP Enterprise Plus

## Supported Platforms

Both Direct Connect approaches support the following platforms:

- VMware vSphere (including Free ESXi)
- Microsoft Hyper-V
- Physical machines (Windows, Linux)
- Proxmox VE

## Supported Features

Both Direct Connect approaches support the following product features:

- Inventory discovery/refresh
- Repository creation/management
- VM backup
- VM replication
- Backup copy
- Full VM recovery
- Flash boot
- File recovery to browser
- Failover/Failback
- Auto update
- File Recovery (to source)
- File Level Backup
- Object Recovery
- Backup Export
- VM verification
- Site Recovery

## Supported Scenarios

Direct Connect

- Performing **Backup** jobs (including File level backup):
    - From the client's local environment to:
        - The allocated MSP repository
        - The client’s local repository
    - From the allocated MSP environment to:

        - The client's local repository
        - The allocated MSP repository
- Performing **Replication** jobs (Real-Time Replication is not supported):
    - From the client's local environment to:
        - The allocated client's local environment
        - The allocated MSP environment
    - From the allocated MSP environment to:
        - The client's local environment
        - The allocated MSP environment
- Performing **Backup Copy** jobs:
    - From the client's local repository to:
        - The client's local repository
        - The allocated MSP repository
    - From the allocated MSP repository to:
        - The client's local repository
        - The allocated MSP repository
- Performing **Full VM recovery** jobs:
    - From the client's local repository to:
        - The client's local environment
        - The allocated MSP environment
    - From the allocated MSP repository to:
        - The client's local environment
        - The allocated MSP environment
- Performing **Flash Boot** jobs:
    - From the client's local repository to:
        - The client's local environment
    - From the allocated MSP repository to:
        - The allocated MSP environment
- Running **File Recovery** (to browser/via email):
    - From the client's local repository to the browser
    - From the allocated MSP's repository to the browser
- Running **File Recovery** (to source/share):
    - From the client's local repository to:
        - The client's local environment

    - From the allocated MSP repository to:
        - The allocated MSP environment
- **Object Recovery** (to source):
    - From the client's local repository to:
        - The client's local environment
    - From MSP allocated repository to:
        - The MSP environment
- Performing **Failover/Failback** jobs:
    - From the client's local environment to:
        - The client's local environment
        - The MSP allocated environment
    - From MSP allocated environment to:
        - The client's local environment
        - The MSP allocated environment
- Performing **Backup Export** jobs:
    - From the client's local repository to:
        - The client's local environment
    - From MSP repository to:
        - The MSP allocated environment
- Performing **VM verification**:
    - From client's local repository to:
        - The client's local environment
    - From the allocated MSP repository to:
        - The allocated MSP environment
- Performing **Site Recovery** jobs:
    - From the client's local environment to:
        - The client's local environment
        - The MSP allocated environment
    - From MSP allocated environment to:
        - The client's local environment
        - The MSP allocated environment

**Notes**

- Connection from client's local Environment to client's local Environment / Repository works with 1 Transporter. Actions between multiple Transporters are not supported.
- For File Recovery (to source) and Object Recovery (to source) (from client's repository to client's environment) you can recover your backup with **Recovery server** dropdown only. Manually entering Server hostname or IP is not supported. For more details refer to the corresponding articles below:
  - File Recovery Wizard: Recovery Method
  - Object Recovery Wizard for Microsoft Exchange: Recovery Method
  - Object Recovery Wizard for Microsoft AD Server: Recovery Server
  - Object Recovery Wizard for Microsoft SQL Server: Recovery Server
- To backup Physical Machines, Hyper-V, or Proxmox VMs from the client's local environment to the allocated MSP repository, a Direct Connect Transporter must be installed on these VMs and used for their discovery. The requirement applies only to this scenario and does not affect VMware environments.
- For Replication, Full VM Recovery, and Failover/Failback from MSP to client, Hyper-V, Proxmox, and Physical Machines on the client site must also be discovered via Direct Connect using a Universal Transporter. However, unlike the client side, the MSP doesn't need a separate Direct Connect Transporter for each machine, one Direct Connect Transporter can be used to discover multiple VMs.
- The Direct Connect approach enables backing up client workloads even when they reside in a LAN-only environment without internet access:
  - The **Direct Connect Transporter** must be installed on a machine within the same LAN as the protected workload, such as a Hyper-V server, Physical Machine, VMware host, or Proxmox cluster.
  - The workload does not require internet access.
  - The **Direct Connect Transporter** must have local network (LAN) access to the protected workload and internet connectivity to establish a connection with the MSP site.

  The MSP Direct Connect approach supports this scenario for VMware environments only. Hyper-V, Proxmox, and Physical Machines are not supported.

MSP Direct Connect

- Performing **Backup** jobs (including file backup):
    - From the client's local environment to:
        - The allocated MSP repository
        - The client's local repository
    - From the allocated MSP environment to:
        - The client's local repository
        - The allocated MSP repository
- Performing **Replication** jobs (Real-Time Replication is not supported):
    - From the allocated client's local environment to:
        - The allocated client's environment
        - The allocated MSP environment
    - From the allocated MSP environment to:
        - The allocated client's environment
        - The allocated MSP environment
- Performing **Backup Copy** jobs:
    - From the client's local repository to:
        - The client's local repository
        - The allocated MSP repository
    - From allocated MSP repository to:
        - The client's local repository
        - The allocated MSP repository
- Performing **Full VM recovery** jobs:
    - From the client's local repository to:
        - The client's local environment
        - The allocated MSP environment
    - From the allocated MSP repository to:
        - The client's local environment
        - The allocated MSP environment
- Performing **Flash Boot** jobs:
    - From the client's local repository to:
        - The client's local environment
    - From the allocated MSP repository to:
        - The allocated MSP environment
- Running **File Recovery** (to browser/via email):

        - From the client's local repository to the browser
        - From the allocated MSP's repository to the browser
- Running **File Recovery** (to source/share):
    - From the client's local repository to:
        - The client's local environment
    - From the allocated MSP repository to:
        - The allocated MSP environment
- Object **Recovery** (to source):
    - From the client's local repository to:
        - The client's local environment
    - From MSP allocated repository to:
        - The MSP environment
- Performing **Failover/Failback** jobs:
    - From the client's local environment to:
        - The client's local environment
        - The MSP allocated environment
    - From MSP allocated environment to:
        - The client's local environment
        - The MSP allocated environment
- Performing **Backup Export** jobs:
    - From the client's local environment to:
        - The client's local environment
    - From MSP allocated environment to:
        - The MSP allocated environment
- Performing **VM verification**:
    - From client's local repository to:
        - The client's local environment
    - From the allocated MSP repository to:
        - The allocated MSP environment
- Performing **Site Recovery** jobs:
    - From the client's local environment to:
        - The client's local environment
        - The MSP allocated environment
    - From MSP allocated environment to:

- The client’s local environment
- The MSP allocated environment

**Notes**

- Connection from client's local environment to client's local Environment / Repository works when using multiple Transporters on the same network. Actions between Transporters on different networks are not supported.
- For File Recovery (to source) and Object Recovery (to source) (from client's repository to client's environment) you can recover your backup with **Recovery server** dropdown only. Manually entering Server hostname or IP is not supported. For more details refer to the corresponding articles below:
  - File Recovery Wizard: Recovery Method
  - Object Recovery Wizard for Microsoft Exchange: Recovery Method
  - Object Recovery Wizard for Microsoft AD Server: Recovery Server
  - Object Recovery Wizard for Microsoft SQL Server: Recovery Server
- To backup Physical Machines, Hyper-V, or Proxmox VMs from the client's local environment to the allocated MSP repository, a Direct Connect Transporter must be installed on these VMs and used for their discovery. The requirement applies only to this scenario and does not affect VMware environments.
- For Replication, Full VM Recovery, and Failover/Failback from MSP to client, Hyper-V, Proxmox, and Physical Machines on the client site must also be discovered via Direct Connect using a Universal Transporter. However, unlike the client side, the MSP doesn't need a separate Direct Connect Transporter for each machine, one Direct Connect Transporter can be used to discover multiple VMs.
- The Direct Connect approach enables backing up client workloads even when they reside in a LAN-only environment without internet access:
  - The **Direct Connect Transporter** must be installed on a machine within the same LAN as the protected workload, such as a Hyper-V server, Physical Machine, VMware host, or Proxmox cluster.
  - The workload does not require internet access.
  - The **Direct Connect Transporter** must have local network (LAN) access to the protected workload and internet connectivity to establish a connection with the MSP site.

  The MSP Direct Connect approach supports this scenario for VMware environments only. Hyper-V, Proxmox, and Physical Machines are not supported.

## Encrypted Backup Repository

To enable encryption, the following requirements should be met:

- The Transporter assigned to the encrypted backup repository must be installed on Ubuntu, SLES or RHEL. Currently, **Transporters** installed on other Linux versions, Windows, and NAS do not support this feature.
- For certain SLES and RHEL versions, only full device/partition encryption is available. Learn more.

## External Database

The following external databases are supported:

- PostgreSQL v10-16x

The following system requirements apply to the machine housing the external database:

- **CPU**:
  - x86-64, 4+ cores (Single-Tenant mode)
  - x86-64, 8+ cores (Multi-Tenant mode)
- **RAM**:
  - 8+ GB (Single-Tenant mode)
  - 36+ GB (Multi-Tenant mode)
- **Free Space**: 50 GB
- **OS**: Windows and Linux operating systems.

**Notes**

- Using SSD is highly recommended.
- The external database can be created on a physical machine or VM or stored in a container.
- Database migration is supported for both the single-tenant and multi-tenant modes of the solution. However, only the Master Admin can perform database migration in multi-tenant mode.
- All tenants share the same database server after the migration, but each tenant has a separate database.
- All tenant databases must be the same type as the database of the Master Admin.
- Some NAS devices may already contain the PostgreSQL as inbox package.

## Federated Repository

For the **Federated repository** feature to work for new and existing backup/backup copy jobs, backup export, or backup migration, the following general requirements apply:

- The **Federated repository** feature supports backing up data to **Incremental with full backups** repositories only.
- You cannot create a federated repository if no supported backup repositories have been added to the product. Refer to Creating Federated Backup Repositories for more details.
- After a federated repository is saved, the jobs using the repositories now selected to be part of the federated repository are automatically reconfigured to use this federated repository as a target.
- After a backup repository is added as a member to a federated repository, it cannot be selected as a target for a new backup/backup copy job.
- A federated repository can be used as a destination for backup/backup copy jobs.
- You can select a federated repository as a source for recovery, backup copy, and replication jobs.
- You cannot select a federated repository member as a source for recovery and replication jobs.
- A federated repository cannot be used for self-backup, and federated repository members are not displayed in the list of repositories available for self-backup.
- When you remove a federated repository, federated repository members and their contents are not removed and are treated as standalone backup repositories. However, note that backup chains may be broken.
- You cannot perform recovery from recovery points located in an inaccessible federated repository member or those that depend on inaccessible recovery points.
- You cannot select a federated repository as a destination for recovery from tape media.
- If critical recovery points have been removed due to the removal of a member, a backup object that remains in the federated repository is marked accordingly. A full recovery point for this backup object is created during the next backup job run.
- Recovery of a workload is not possible if a federated repository member containing the required full recovery point is unavailable.
- You can modify the settings of an existing federated repository, detach it, or remove it.

**Notes**

- A federated repository cannot be detached while a backup job using this federated repository as a target is running.
- A federated repository cannot be removed while a backup job using this federated repository as a target is running. To remove such federated repository, delete (or edit) the corresponding job to ensure that no items are being backed up to this federated repository.
- A federated repository cannot be removed if it contains at least one backup object referenced by existing job(s).

- You can add/remove federated repository members with corresponding attributes to/from a federated backup repository and manage/refresh federated repository members individually.
- The federated repository member can only be removed at the federated repository level.

**Notes**

- You cannot remove the only remaining member storing full only recovery points.
- You cannot remove the only remaining member storing incremental only recovery points.

**Supported Backup Location Policy**

- **Integrity**: All dependent recovery points are stored on the same federated repository member.

**Supported Functionality**

- Backup/backup copy
- Full VM/instance recovery
- Object level recovery
- Backup export
- Universal object recovery
- Flash VM boot
- Backup migration

**Supported Federated Repository Member Types**

- Local Folder
- CIFS Share
- NFS Share

**Supported Repository Attributes**

- Compression level (fast/medium/best/disabled)
- Encryption (enabled/disabled)

# File Recovery

Recovered files can be downloaded or sent via email. They can also be recovered to a server or file share. Before using the feature, make sure the following packages and services are installed/running either on the (proxy) transporter or target server depending on the selected recovery method:

Microsoft Windows

- *Net Security* package should be installed
- *Microsoft iSCSI Initiator* service should be installed and running
- *net.exe* utility should be installed
- SMB (1.x, 2.x, 3.x)/CIFS File Sharing Support feature should be turned on

- PowerShell should be version 2.0 or above
- PowerShell ISE should be available

Ubuntu Linux

- *openssh-server* package should be installed
- *sshd* service should be running
- *parted* utility should be installed
- *fdisk* utility should be installed
- *open-iscsi* package should be installed
- *iscsiadm* utility should be installed
- *iscsid* utility should be installed
- *iscsid* service should be running (for v16.04 and v18.04)
- *iscsi_tcp* module should be installed
- *SElinux* service should be disabled

SUSE Linux Enterprise Server (SLES)

- *openssh-server* package should be installed
- *sshd* service should be running
- *parted* utility should be installed
- *fdisk* utility should be installed
- *open-iscsi* package should be installed
- *iscsiadm* utility should be installed
- *iscsid* utility should be installed
- *iscsid* service should be running (for v12)
- *iscsi_tcp* module should be installed
- *SElinux* service should be disabled

Red Hat Enterprise Linux (RHEL)

- *openssh-server* package should be installed
- *sshd* service should be running
- *parted* utility should be installed
- *fdisk* utility should be installed
- *iscsi-initiator* package should be installed
- *iscsiadm* utility should be installed
- *iscsid* utility should be installed
- *iscsid* service should be running
- *iscsi_tcp* module should be installed
- *SElinux* service should be disabled

> **Note**
> File recovery is not supported for 4K sector size drives and datastores. Refer to How to Check Drive and Datastore Sector Size for more information.

Below are the requirements which must be met for each recovery method.

## Downloading Files to Browser or Sending Files via Email

The following file systems are supported:

- If the Transporter assigned to the backup repository is installed on Windows:
  - NTFS
  - FAT32
  - ReFS

  > **Note**
  > If the OS build, patch level, or ReFS version on the machine hosting Transporter, that mounts backup data in FLR/OLR/UOR scenarios, differs from the source machine OS build, patch level, or ReFS version, recovery may fail or be incomplete.

- If the Transporter assigned to the backup repository is installed on Linux:
  - NTFS
  - FAT32
  - EXT2
  - EXT3
  - EXT4
  - XFS
  - SwapLinux
- For the ReiserFS file system, it is necessary to install the *linux-image-extra-virtual* package for Ubuntu.:

  ```
  apt-get -y install linux-image-extra-virtual
  ```

- Linux VMs where Transporter is deployed should have the *lvm2* package installed to allow mounting LVM volumes.
- The *ntfs-3g* package should be installed along with Transporter on Linux to allow recognizing NTFS partitions.

## Recovering Files to Server

To recover files to a server, make sure you meet the following requirements:

**Supported OS**

- **Windows**
    - Windows Server 2025 (24H2) (x64)
    - Windows Server 2022 (21H2) (x64)
    - Windows Server 20H2 (20H2) (x64)
    - Microsoft Windows Server 2022 (x64)
    - Microsoft Windows Server 2019 Standard (x64)
    - Microsoft Windows Server 2016 Standard (x64)
    - Microsoft Windows 11 (21H2/22H2) (x64)
    - Microsoft Windows 11 (x64)
    - Windows 10 Enterprise (20H2 / 21H1 / 21H2) (x64)
    - Microsoft Windows 10 Home (x64)
    - Microsoft Windows 10 Professional (x64)
- **Linux**
    - Debian 13 (64-bit)
    - Debian 12.10 (64-bit)
    - Debian 12.9 (64-bit)
    - Debian 12.8 (64-bit)
    - Debian 12.7 (64-bit)
    - Debian 12.6 (64-bit)
    - Debian 12.5 (64-bit)
    - Debian 12.4 (64-bit)
    - Debian 12.2 (64-bit)
    - Debian 12.1 (64-bit)
    - Debian 12.1 (64-bit)
    - Debian 11.7 (64-bit)
    - Debian 11.6 (64-bit)
    - Debian 11.5 (64-bit)
    - Debian 11.4 (64-bit)
    - Debian 11.3 (64-bit)
    - Debian 11.2 (64-bit)
    - Debian 11.1 (64-bit)
    - Debian 11.0 (64-bit)
    - Debian 10.13 (64-bit)
    - Debian 10.12 (64-bit)
    - Debian 10.11 (64-bit)
    - Debian 10.10 (64-bit)

- Debian 10.9 (64-bit)
- Debian 10.8 (64-bit)
- Debian 10.7 (64-bit)
- Debian 10.6 (64-bit)
- Debian 10.5 (64-bit)
- Debian 10.4 (64-bit)
- Debian 10.3 (64-bit)
- Debian 10.2 (64-bit)
- Debian 10.1 (64-bit)
- Ubuntu 25.04 Server (x64)
- Ubuntu 24.10 Server STS (x64)
- Ubuntu 24.04 Server LTS (x64)
- Ubuntu 22.04.4 Server LTS (x64)
- Ubuntu 22.04 Server LTS (x64)
- Ubuntu 20.04 Server (x64)
- Ubuntu 18.04 Server (x64)
- Ubuntu 16.04 Server (x64)
- SUSE Linux Enterprise Server 15 SP6 (x64)
- SUSE Linux Enterprise Server 15 SP5 (x64)
- SUSE Linux Enterprise Server 15 SP4 (x64)
- SUSE Linux Enterprise Server 15 SP3 (x64)
- SUSE Linux Enterprise Server 15 SP2 (x64)
- SUSE Linux Enterprise Server 15 SP1 (x64)
- SUSE Linux Enterprise Server 12 SP5 (x64)
- SUSE Linux Enterprise Server 12 SP4 (x64)
- SUSE Linux Enterprise Server 12 SP3 (x64)
- SUSE Linux Enterprise Server 12 SP2 (x64)
- SUSE Linux Enterprise Server 12 SP1 (x64)
- Red Hat Enterprise Linux 10.0 (x64)
- Red Hat Enterprise Linux 9.6 (x64)
- Red Hat Enterprise Linux 9.5 (x64)
- Red Hat Enterprise Linux 9.4 (x64)
- Red Hat Enterprise Linux 9.3 (x64)
- Red Hat Enterprise Linux 9.2 (x64)
- Red Hat Enterprise Linux 9.1 (x64)
- Red Hat Enterprise Linux 9.0 (x64)

- Red Hat Enterprise Linux 8.12 (x64)
- Red Hat Enterprise Linux 8.11 (x64)
- Red Hat Enterprise Linux 8.10 (x64)
- Red Hat Enterprise Linux 8.9 (x64)
- Red Hat Enterprise Linux 8.8 (x64)
- Red Hat Enterprise Linux 8.7 (x64)
- Red Hat Enterprise Linux 8.6 (x64)
- Red Hat Enterprise Linux 8.5 (x64)
- Red Hat Enterprise Linux 8.4 (x64)
- Red Hat Enterprise Linux 8.3 (x64)
- Red Hat Enterprise Linux 8.2 (x64)
- Red Hat Enterprise Linux 8.1 (x64)
- Red Hat Enterprise Linux 8.0 (x64)
- Red Hat Enterprise Linux 7.9 (x64)
- Red Hat Enterprise Linux 7.8 (x64)
- Red Hat Enterprise Linux 7.7 (x64)
- Red Hat Enterprise Linux 7.6 (x64)
- Red Hat Enterprise Linux 7.5 (x64)
- Red Hat Enterprise Linux 7.4 (x64)
- Red Hat Enterprise Linux 7.3 (x64)
- Red Hat Enterprise Linux 7.2 (x64)
- Red Hat Enterprise Linux 7.1 (x64)
- Red Hat Enterprise Linux 7.0 (x64)
- CentOS Stream 10 (x64)
- CentOS Stream 9 (x64)
- CentOS Stream 8 (x64)
- CentOS Linux 8.5 (x64)
- CentOS Linux 8.4 (x64)
- CentOS Linux 8.3 (x64)
- CentOS Linux 8.2 (x64)
- CentOS Linux 8.1 (x64)
- CentOS Linux 8.0 (x64)
- CentOS Linux 7.9 (x64)
- CentOS Linux 7.8 (x64)
- CentOS Linux 7.7 (x64)
- CentOS Linux 7.6 (x64)

- CentOS Linux 7.5 (x64)
- CentOS Linux 7.4 (x64)
- CentOS Linux 7.3 (x64)
- CentOS Linux 7.2 (x64)
- CentOS Linux 7.1 (x64)
- CentOS Linux 7.0 (x64)
- AlmaLinux 8.7 - 10 (x64)

**TCP Ports**

Connection to the following TCP ports should be allowed by the firewall of the target system:

- *22* – Used by SSH for secure logins, file transfers (scp, sftp) and port forwarding.
- *9445* – Used by NAKIVO Backup & Replication to communicate with the VM.
- *10000* – Used by NAKIVO Backup & Replication for iSCSI communication.

> **Note**
> ICMP Ping traffic should be allowed by the firewall of the target system.

**Permissions**

The following permissions for Microsoft Windows VMs should be granted:

- Users should be members of a local Administrators group.
- Users should have access to default administrative shares.
- Users should have permissions to access the corresponding folder\file.
- Users should have executive permissions for running some utilities, for example, `net.exe` utility.
- User Account Control (UAC) remote restrictions should be disabled for some Microsoft Windows versions.
- Users should have permissions to "Log on as a batch job".

The following permissions and settings should be set up for Linux VMs:

- Users should belong to the `sudo` group to complete recovering files to server successfully.
- Users should have executive permissions for running some utilities, for example, /sbin/parted, /sbin/fdisk, /sbin /iscsiadm, /sbin/iscsid.
- PasswordAuthentication should be set to "yes".
- Provide special permissions to NAKIVO recovery service. For more details, refer to Required Permissions for Linux Recovery Server.

## Recovering Files to File Share

The following file share types are supported:

- NFS
- CIFS

The target share must have one of the following protocols installed:

- NFS 3.x
- SMB 1.x
- SMB 2.x
- SMB 3.x

To recover file/folder attributes or encrypted files, the machine on which the Transporter is installed, recovery share, and backed up machine must have the same OS capabilities and file system type. To avoid potential issues while conducting file recovery, ensure the following:

- If the target is an NFS file share, check that nfs-utils is installed.
- If the target is a CIFS share, add a localhost domain to the user credentials (e.g. localhost\Administrator).

## File Share Backup

The following requirements must be met to use the feature:

- File shares must be located on a NAS, Windows Server, Linux Server, or Nutanix File Server that supports the NFS or CIFS/SMB protocol.
- Read permission is required to back up the file shares.
- Read/write permission is required to restore to a file share.
- To use NFS file shares with Windows, NFS client feature must be enabled.
- To add a Windows-configured NFS share to the Inventory, the `nfs` file share location must be reachable by the Transporter. For more details, refer to NFS File Share Location Cannot Be Reached by the Transporter.
- File shares with the following protocols are supported:
  - NFS 3.x
  - SMB 1.x
  - SMB 2.x
  - SMB 3.x
- The following operating systems do **not** support any NFS Server versions:
  - Windows 7
  - Windows 8.1
  - Windows 10

## File System Indexing

**File System Indexing** is supported on all supported browsers for NAKIVO Backup & Replication and runs on all supported operating systems.

**File System Indexing** skips files and folders on disks attached via FC and FCOE.

**File System Indexing** skips `tmpfs` mountpoints on Linux operating systems and the following system folders:

- Windows OSs:
  - %ProgramFiles%
  - %ProgramFiles(x86)%
  - %ProgramW6432%
  - %windir%
  - %TEMP%
- Linux OSs:
  - /boot
  - /dev
  - /cdrom
  - /media
  - /proc
  - /run
  - /selinux
  - /tmp
  - /sys

Before you can use the **File System Indexing** feature, make sure the source VM must meet the following requirements:

- VMware Tools or Hyper-V Integration Services must be running on VMware or Hyper-V VMs correspondingly.
- *iSCSI* must be available.
- *SMB* must be enabled on Windows OS.
- *SSH* must be enabled on Linux OS.
- TCP port `9445` must not be blocked.
- *SElinux* service must be disabled on Linux OS.
- Special permissions must be configured for Linux OS.
- Special permissions must be configured for Windows OS.
- To run **File System Indexing** on RHEL 7.9 and RHEL 7.4, the **Netstat** (network statistics) command line tool must be installed on the source VM.

## Generic S3-Compatible Object Storage

The following vendors for generic S3-compatible object storage are currently supported:

- MinIO
- Ceph
- Cloudian

- C2 Object Storage
- SeaGate LyveCloud
- IDrive e2
- OVH Cloud
- Cubbiit (cubbit.io)
- Pilvio (pilw.io)
- Google Cloud
- VAST Data
- DigitalOcean
- Google Cloud Storage

**Notes**

- This list only contains vendors that were specifically tested and will be updated as NAKIVO continues to test more vendors.
- The retention setting for bucket is not required, it can be disabled or set as none
- Immutability is supported only if object lock and S3 versioning are enabled on the vendor side and are supported by vendor APIs.
- Immutability is not supported for MinIO, SeaGate LyveCloud, and Google Cloud.
- Some S3-compatible vendors not listed above may be supported if they use the applicable APIs. The list of APIs used by NAKIVO Backup & Replication for generic S3-compatible object storage functionality can be found here.

## Granular User Notifications

The feature supports the following jobs:

- Backup jobs:
  - VMware vSphere
  - Amazon EC2
  - Microsoft Hyper-V
  - Physical machine
  - Nutanix AHV
  - Oracle Database
  - Microsoft 365
- Replication jobs:
  - VMware vSphere
  - Amazon EC2
  - Microsoft Hyper-V

## Hot Add

In order for the Hot Add feature to work for VMware VM backup, replication, and recovery, the following requirements must be met:

- The Transporter that will be reading or writing data from/to the VM disks should run on a VM.
- The Transporter VM should:
  - Be available in the product Inventory,
  - Run on a host that has access to the datastore(s) with the VM disks, Run in the same datacenter as the VM that is to be processed.

A single SCSI controller on the VM hosting NAKIVO Backup & Replication can support up to 15 disks including the system disk of the VM with NAKIVO Backup & Replication and mounted disks of the Backup Repository. To process VMs with a total number of disks that is larger than that limit, it is necessary to install one or more additional SCSI controllers.

## Merge Jobs

The feature supports the following types of jobs:

- Backup
- Backup copy
- Replication

Job merging can be performed in the following cases:

- Both source and target jobs are of the same type and platform.
- The source job is in an idle state.

Job merging cannot be performed in the following cases:

- One of the selected jobs is a backup copy job with the destination set to tape.
- The target job uses the **Policies** view.
- The Transporter selection settings of the target job cannot be applied to the source job objects.
- Both source and target jobs contain or reference the same workload.

## Monitoring

For the feature to function correctly, the following conditions must be met:

- For NAKIVO Backup & Replication to display disk usage, VMware Tools must be running in a VM.
- VM must be powered on for the monitoring to be available.
- The *fontconfig* library must be installed for reports to work correctly.

The feature supports the following platforms:

- Paid versions of VMware vCenter Server
- Paid versions of VMware ESXi host

The following types of items are supported:

- VMware vSphere host
- VMware vSphere datastore
- VMware vSphere VM

The full list of supported versions of vCenter and ESXi can be found here.

If monitoring over 1000 workloads, make sure your machine meets the following requirements:

- **CPU**: x86-64, 4+ cores
- **RAM**: 8 GB
- **Disk Free Space**: 5 GB

## MSP Console

To use the **MSP Console** to manage remote tenants, a managed service provider (MSP) needs to configure the following TCP ports:

- **MSP Director port**: This is the TCP port used by the Director for the MSP's instance of NAKIVO Backup & Replication. By default, this is TCP port *4443*. The MSP must provide a remote tenant with their Director port number during configuration. The remote tenant needs to enter this port number when adding the MSP to their standalone instance of NAKIVO Backup & Replication.
- **Listening port**: Additionally, the MSP must have a port open for listening to the remote tenant. By default, TCP port *6702* is used. The MSP may change the listening port used by changing the **system.msp.console.listening.port** parameter in Expert settings.

To use the **MSP Console** to create and manage local tenants, a managed service provider (MSP) can enable Direct Connect to establish a connection with client remote resources.

**Notes**

- Only users with an *MSP license*, *Beta instance*, *Promo license*, or *Trial license* can access the MSP Console and the Licensing > Tenants tab. This allows them to efficiently oversee all independent instances of NAKIVO Backup & Replication associated with a managed service provider (MSP) as well as local tenants from a unified interface, eliminating the need to navigate through individual tenants.
- Your version of NAKIVO Backup & Replication must be the same as the MSP's version. Otherwise, you cannot connect to MSP's instance of NAKIVO Backup & Replication. This note is only relevant for multi-tenancy product with remote tenants.
- The MSP uses a separate listening port for communication with a remote tenant's instance (port *6702* is used by default). If the MSP changes the listening port used, the connection may be interrupted.

## Native Tape Support

NAKIVO Backup & Replication supports tape environments with the following configurations:

- Robotic tape libraries and standalone devices of LTO3 and later generations.
- AWS Storage Gateway service with a Tape Gateway that functions as a Virtual Tape Library (VTL).
- FalconStor StorSafe VTL (for VTLs connected to Linux and Windows).
- The gateway VM deployed on-premises needs to have the following minimum resources:
  - **CPU**: x64, 4+ cores
  - **RAM**: 16+ GB
  - **Free Space**: 80 GB
- According to the requirements for Amazon EC2 instances, when deploying the gateway on Amazon EC2, the instance size should be at least `2xlarge` for the compute-optimized instance family.
- The instance type should be c4 or c5 instance types. The 2xlarge instance size or higher can be chosen to meet the required RAM requirements.
- All physical tape cartridges must have barcodes.
- Installation is supported on all Windows OS and Linux OS, as listed on the Supported Platforms page.
- Installation on NAS OS is not supported.
- The "mtx" and "lsscsi" utilities need to be installed on the Linux transporter server in order to detect the tape library changer.

## Object Recovery and Log Truncation for Microsoft Exchange

To successfully perform object recovery and log truncation for Microsoft Exchange, make sure you meet the following requirements:

**Supported Microsoft Exchange versions**

NAKIVO Backup & Replication supports the following versions of Microsoft Exchange for object recovery and log truncation:

- Microsoft Exchange 2019
- Microsoft Exchange 2016
- Microsoft Exchange 2013

**Permissions**

The following requirements should be met for log truncation:

- Selected users should have permissions to "Log on as a batch job".
- Active Directory Module For Windows PowerShell must be installed.
- The VM must be accessible over network.
- The following user permissions should be provided:
  - If NAKIVO Backup & Replication uses the administrator user account, it should belong to the following groups:
    - Administrators
    - Domain Users

        - Organization Management
    - If NAKIVO Backup & Replication uses accounts other than the `administrator` user account:
        - The user should belong to the following groups:
            - Administrators
            - Domain Users
            - Organizational Management
        - The user should have the Full control permission granted for the folder in which the Exchange database is located.

**Services and Settings**

NAKIVO Backup & Replication requires PowerShell v2 or later to be available on the Microsoft Exchange machine.

- VMware VM must be running on VMware ESXi 5.0 and later.
- VMware Tools or Hyper-V Integration Services must be running on VMware or Hyper-V VMs correspondingly.

# Object Recovery and Log Truncation for Microsoft SQL Server

To successfully perform object recovery and log truncation for a Microsoft SQL Server, you must meet general requirements as well as requirements for object recovery and log truncation.

**General Requirements**

To successfully perform object recovery and log truncation for a Microsoft SQL Server, make sure you meet the following general requirements:

**Supported Versions of Microsoft SQL Server**

NAKIVO Backup & Replication supports the following versions of Microsoft SQL Server for object recovery and log truncation:

- Microsoft SQL Server 2022
- Microsoft SQL Server 2019
- Microsoft SQL Server 2017
- Microsoft SQL Server 2016
- Microsoft SQL Server 2014
- Microsoft SQL Server 2012

**Supported Editions**

NAKIVO Backup & Replication supports the following editions of Microsoft SQL Server:

- Enterprise
- Standard
- Web

- Developer
- Express edition

**Permissions**

- A user logging in to Microsoft SQL Server must have a `sysadmin` role.
- The user running Microsoft SQL Service should have permissions to "Log on as a batch job".

**Services and Settings**

- NAKIVO Backup & Replication requires PowerShell v2 or later.
- VMware Tools or Hyper-V Integration Services must be running on VMware or Hyper-V VMs correspondingly.
- `sqlcmd` utility must be installed on the machine running Microsoft SQL server.
- Ports *137-139* must be opened for `cifs`.
- The SMB v1 or higher version of SMB protocol should be enabled.

## Requirements for Microsoft SQL Server Object Recovery

- The user running Microsoft SQL service must have executive permissions to the `Data` folder and all other folders in which the databases are located.
- If "Rename recovered item if such item exists" option is selected during the recovery, NAKIVO Backup & Replication skips keys, constraints, indexes, and statistical properties when recovering a table to an original location.
- If "Overwrite the original item if such item exists" option is chosen, all the above properties are preserved. Tables that contain a foreign key cannot be recovered with this option.
- Full administrative permissions are required.
- Default administrative shares must be enabled.
- The "File server" role must be enabled.
- Ports *445* and 9*445* must be opened on the instance.

## Requirements for Microsoft SQL Server Log Truncation

- VMware VM must be running on VMware ESXi 5.0 and later.
- System databases are skipped during the log truncation.
- Databases with the "Simple" recovery model are skipped during the log truncation.
- A database must be in the "online" state.
- The SMB v1 or higher version of SMB protocol should be enabled.

# Object Recovery for Microsoft Active Directory

## Supported Versions

NAKIVO Backup & Replication supports the following versions of Microsoft Active Directory for objects recovery:

- Windows Server 2022 (21H2)
- Windows Server 20H2 (20H2)
- Windows Server 2019
- Windows Server 2016

## Requirements for Object Recovery for Microsoft Active Directory

- The ISCI Initiator service must be running on the recovery server.
- The vc_redist.x86.exe (v.2015) file must be installed on the recovery server.
- Active Directory Web Services must be running.
- Port *5000* must not be blocked by other services and must be opened in the firewall of AD.
- Active Directory Module For Windows PowerShell must be installed.

# Permanent VM Agent

VM agents are available for the following platforms:

- VMware vSphere
- Microsoft Hyper-V
- Amazon EC2
- Nutanix AHV

VM agents support the following features:

- OS quiescing
- File recovery to source
- File System Indexing

VM agents can be used in the Windows and Linux operating systems listed in this section of Deployment Requirements.

# Proxmox VE Support

Proxmox Virtual Environments (VE) are supported on all supported browsers for NAKIVO Backup & Replication and all supported operating systems.

To successfully discover Proxmox VE items (clusters and standalone hosts (nodes)), add them to the **Inventory**, back up Proxmox VE VMs, create backup copies, and perform full and file level recovery, make sure you meet the following requirements:

- **Proxmox VE** are supported for the following Proxmox VE versions:
  - **Proxmox VE 8.x**. For more details, refer to Proxmox VE versions table.

> **Notes**
>
> - The Proxmox VE (host/cluster) of the version lower than the oldest supported version cannot be discovered.
> - The Proxmox VE (host/cluster) of the version higher than the newest supported version can be discovered, but confirmation of the operation is required to proceed.
> - The product can discover all Proxmox VE storage types (See Proxmox VE Storage Types), except the following:
>   - PBS
>   - ESXi

- A transporter must be available inside the Proxmox VE host for running backup and recovery jobs.

> **Notes**
>
> - For the Proxmox VE cluster, the transporter is injected into each cluster host.
> - Transporters injected into Proxmox VE hosts are supported for auto-update.

**Supported Functionality**

- Backup of Proxmox VE VMs to Incremental with Full and Forever Incremental repositories
  - Immutability
  - Change Block Tracking (CBT)
    - Proprietary method
    - Native change tracking
  - Application awareness

**Notes**

App-aware mode on Proxmox VE VM requires the following to function:

- QEMU Guest Agent must be enabled and running on the VM. You can test the communication qm agent by running the following command. If the QEMU Guest Agent runs correctly in the VM, it will return without an error message.

```
qm agent <vmid> ping
```

- Virtio-serial driver must be installed on the Windows based VM. For Linux based VMs, you only need to run the install command.
  - For Debian/Ubuntu: *apt-get install qemu-guest-agent*
- For FreeBSD-based VMs, create the script files in the source VM usr/sbin/ folder:
  - "pre-freeze-script"
  - "post-thaw-script"

  Make sure they have sufficient permissions.

    - Backup encryption
    - VM verification
    - Log truncation for Proxmox VE backups
- Backup copy of Proxmox VE VM can be saved to the following types of storage:
    - Repositories/stream repositories with creating immutable copies for Proxmox VE VM backups
    - Data storage types:
        - Forever Incremental
        - Incremental with periodic full
    - Repository types:
        - Local Folder
        - CIFS Share
        - NFS Share
        - Cloud & S3-Compatible Storages
            - Amazon S3
            - Wasabi
            - Backblaze B2 Cloud Storage
            - Microsoft Azure Blob Storage
            - Amazon EC2
            - Generic S3-compatible Storage

    - Deduplication Appliances
        - Dell EMC Data Domain Boost
        - HPE StoreOnce Catalyst
        - NEC HYDRAstor
- Backup copy to tape
- Replication of Proxmox VE VMs
- Flash VM Boot from Proxmox VM backups
- Full recovery of Proxmox VE VM backups
- File recovery from Proxmox VE VM backups
- Direct recovery of Proxmox VE VM backups from tape
- Object recovery from Proxmox VE VM backups
- Universal object recovery for Proxmox VE VM backups
- Backup export for Proxmox VE VM backups
- Granular notifications for Proxmox VE backup jobs

**Feature Limitations**

Native Change Tracking

- If the VM image is stored in QEMU format, the dirty bit map (CBT) is persistent, and VM restart does not affect CBT (we can continue conducting incremental backups).
- If it is stored in RAW format, a non-persistent dirty bit map is used, and VM restart leads to a dirty bitmap reset and the next backup should be full.

Backups can be performed on disks mapped to ISCSI storage LUNs but they cannot be used as target storage for recovery due to the following reasons:

- When an ISCSI LUN is mapped to multiple disks or VMs, the recovery process writes data to that LUN, which affects all other VMs on that LUN.
- There is a requirement for LUNs to be specified for every recovery disk, which is not supported by the current UI.
- NAKIVO Backup & Replication uses iSCSI storage for Flash VM Boot jobs. For this reason, VM disks cannot be mapped to storages and can only be either skipped or recovered.

**Proxmox VE Limitations**

- Proxmox VE VM network devices cannot be connected directly to a physical network; they must be connected to a bridge.
- Freeze and thaw hooks in the QUEMU guest agent are only supported on Linux VMs. Custom quiescing scripts for Proxmox VE are not supported on Windows VMs.
- Proxmox VE storages must contain **Disk Image** in their **Content** field to store VM disks.
- To successfully perform backups, the number of free nbd devices on a Proxmox VE host must be more than or equal to the number of VM disks to back up.

- For successful Proxmox VE VM replication, the target storage must support snapshot, and source disk format will be converted to “qcow2”.
- The Proxmox's log truncation function uses "QEMU Agent" to execute the script with "local system" account (NT AUTHORITY\SYSTEM):
  - For SQL Server, it requires using an account having "sysadmin" role.
  - If the account does not have "sysadmin" role by default, you cannot execute some SQL commands (Eg: BACKUP DATABASE).
  - It is also not possible to switch to an account having "sysadmin" role when executing a command.
  - For the log truncation to work, do the following:
    - Re-config "QEMU Agent" service to run using an account having SQL's "sysadmin" role, and restart the service.
    - Re-config SQL's "NT AUTHORITY\SYSTEM" account: Assign "sysadmin" role.
- Do not manually unlock or reboot a Proxmox VM from the Proxmox UI while a backup job is running. Doing so may cause temporary snapshot files to remain after the successful job run.

## Real-Time Replication for VMware

Real-time replication is supported for the following ESXi versions:

- Source hosts: ESXi 7.0, v8.0, v8.0U2, or v8.0U3a cluster managed by vCenter
- Target hosts: see “Supported Platforms” on page 131

NAKIVO Backup & Replication requires the following listening TCP ports to be open for a successful operation:

- *33088* – Used by the **I/O Filter** (source daemon installed on ESXi host) for receiving requests (or commands) from the **Director**.
- *33089* – Used by the **Journal Service** (installed on the target **Transporter**) for receiving requests (or commands) from the **Director** and the **I/O Filter**. (Outgoing ports on the **I/O Filter**).
- *33090* – Used by the **Journal Service** (installed on the target **Transporter**) for receiving changed data from the **I/O Filter**. (Outgoing ports on the **I/O Filter**).
- *33091* – Used by the **Journal Service** (installed on the target **Transporter**) for receiving heartbeat requests from the **I/O Filter**. (Outgoing ports on the **I/O Filter**).
- *33092* – Used by the **Journal Service** (installed on the target **Transporter**) for receiving encrypted data from the **I/O Filter**. (Outgoing ports on the **I/O Filter**).

The following requirements must be met to use the feature:

- Source hosts must have the **Host Image Profile Acceptance Level** setting set to **Community level**. Therefore, only clusters managed by vCenter are supported as a source (no clusters using a single image are supported). To configure this, select the host and navigate to **Configure** > **System** > **Security Profile**.
- Source hosts must have secure boot disabled.
- For each target ESXi host or cluster, a virtual appliance must be deployed with a transporter discovered on at least one host. The Journal Service is automatically installed on the target Transporter after Real-Time Replication job creation. For more information see "Installing Journal Service" on page 590.
- The target Transporter must use Hot Add mode to access the replica and uses 4 SCSI controllers. If the replica has fewer than 4 SCSI controllers, the product shall add more SCSI controllers.

**Notes**

- If a replica and a Transporter/Journal Service VM are located on different hosts in the cluster, and a single datastore is mounted to both cluster hosts, the real-time replication job will run given that this is a Hot Add requirement (**Transporter** should have access to replica disks).
- If a replica and a Transporter/Journal Service VM are located on different hosts in the cluster, and different (not local) datastores are mounted to each cluster host, the real-time replication job will run as this is a Hot Add requirement (**Transporter** should have access to replica disks).
- If a replica and a Transporter/Journal Service VM are located on different hosts, and different datastores (for example, local datastores) are visible per particular host, the real-time replication job will fail because the Transporter VM should have access to replica disks (based on Hot Add requirements).

- The time on the source ESXi host and the target **Transporter** appliance must be synchronized.
- The target datastore must have at least 5 GB of free space.
- The source ESXi host must have at least 16 GB of RAM.
- For one real-time replication job/real-time replication failover job, the **Transporter**'s maximum capacity is limited to 11 source VMs in case there is one disk per VM or up to 50 disks per source VM. For more information, see this article.
- The supported object maximums for real-time replication are:

| Type | Limit |
|---|---|
| VMs per host | 100 |
| Disks per VM | 50 |

| Type | Limit |
|---|---|
| Disks per host | 500 |
| Source VM size | 30 TB |
| Source VMs per job | 100 |
| Source VMs per product/tenant | 500 |
| Source VM disks per product/tenant | 1000 |
| Jobs per product/tenant | 100 |
| Journal size per VM | 20 TB |

**Permissions**

The vCenter/cluster users that are added to NAKIVO Backup & Replication must be created from **Global Permissions** and have the following privileges:

- **Host** > **Configuration** > **Query patch**
- For vSphere 7.0: **Profile-driven storage update** and **Profile-driven storage view**
- For vSphere 8.0: **Update VM storage policies** and **View VM storage policies**

## Site Recovery

To successfully perform a site recovery, make sure you meet the following requirements:

**Run Script Action**

The list of supported operating systems where the Run Script action may be run can be found in the Recovering Files to Server subsection above.

**TCP ports**

Connection to the following TCP ports should be allowed by the firewall of the target system:

- *22* – Used by SSH for secure logins, file transfers (scp, sftp) and port forwarding.
- *9445* – Used by NAKIVO Backup & Replication to communicate with the VM.

> **Note**
> ICMP ping traffic should be allowed by the firewall of the target system.

**Required permissions for Microsoft Windows VMs**:

- Users should be members of a local Administrators group.
- Users should have access to default administrative shares.
- Users should have permissions to access the corresponding folder/file.
- User Account Control (UAC) remote restrictions should be disabled for some Microsoft Windows versions.
- Users should have permissions to "Log on as a batch job".

**Services and Settings**

- For Windows source VMs, the SMB (1.x, 2.x, 3.x) / CIFS File Sharing Support feature should be turned on.
- For Linux VMs, users should belong to the `sudo`group.
- VMware Tools or Hyper-V Integration Services must be running on VMware or Hyper-V VMs, respectively.

## Transporter Load Balancing

The **Transporter load balancing** feature ensures the even distribution of tasks among available transporters while processing backup/replication/recovery jobs.

The following platforms are supported:

- VMware vSphere
- Nutanix AHV
- Amazon EC2

**Requirements**

To use the **Transporter load balancing** feature, ensure that the product is set to determine the transporter automatically for at least one VM/instance. To verify, go to the backup job **Options** tab and check that the **Data Transfer** > **Transporters** option is set to **Automatic selection**.

**Feature Limitations**

- Direct Connect (proxy transporter)
    - For items that were added using the Direct Connect feature, the product uses the transporter initially used to add an item.

## VM Limitation for Multi-Tenancy

The feature is only available if a license with Socket limit mode is installed.

The following hypervisors are supported:

- VMware vSphere
- VMware Cloud Director
- Microsoft Hyper-V
- Nutanix AHV

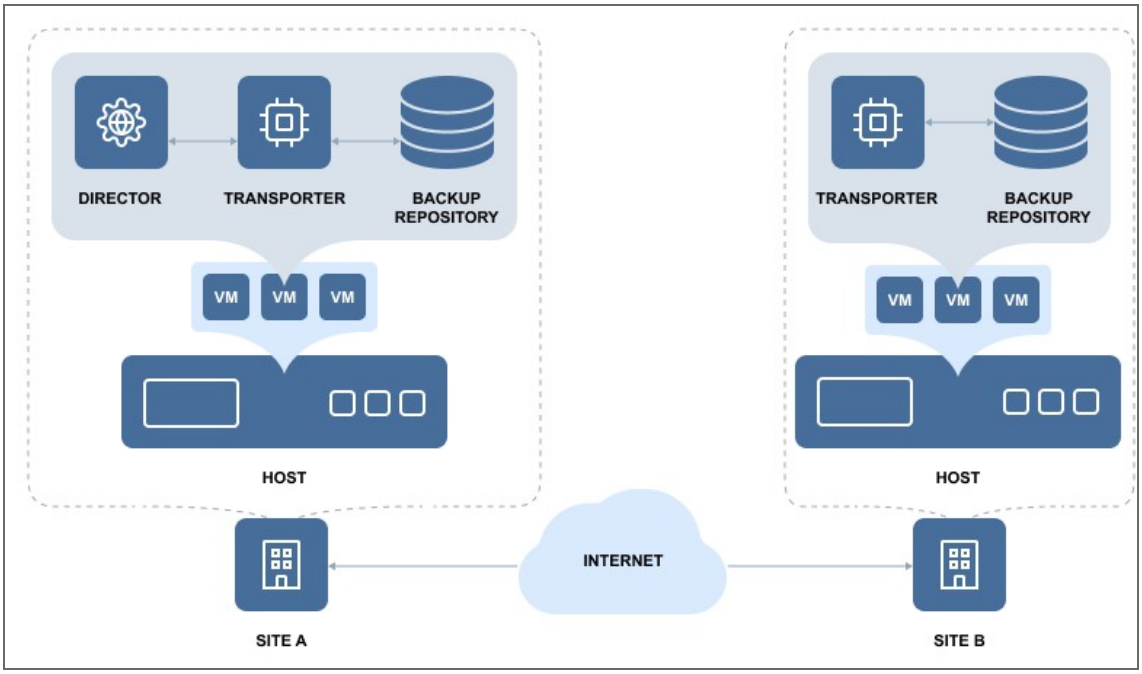
# Deployment Scenarios

NAKIVO Backup & Replication is a modular solution that can be fully installed on a single machine to protect small and mid-sized environments, as well as scale out horizontally and support large distributed environments. Refer to the sections below to learn more about the product deployment scenarios.


- “Distributed Deployment” on page 216
- “Multi-Tenant Deployment” on page 217
- “Single Site Deployment” on page 230

# Distributed Deployment

If you have multiple sites and need to back up and/or replicate over WAN, install the Director and Transporter on one site, and at least one Transporter on all other sites.




**Note**

Make sure the required ports are open on the appropriate endpoints. The full list of required ports can be found in Deployment Requirements.

# Multi-Tenant Deployment

Installation of a multi-tenant solution of NAKIVO Backup & Replication allows you to create multiple isolated tenants within a single product deployment and manage them from a single pane of glass. In the Multi-Tenant mode, tenants can access the self-service portal to offload backup, replication, and recovery tasks from the service provider.

For more information, refer to these topics:


- “Backup from Remote Site to Master Site” on page 218
- “Replication from Remote Site to Master Site” on page 220
- “Local Backup at Remote Site” on page 222
- “Local Replication at Remote Site” on page 224
- “Backup at Master Site” on page 226
- “Replication at Master Site” on page 228
- “Multi-Tenant Mode” on page 1254

# Backup from Remote Site to Master Site


- Deployment Scenario
- Deployment Diagram
- Deployment Steps
- Connections

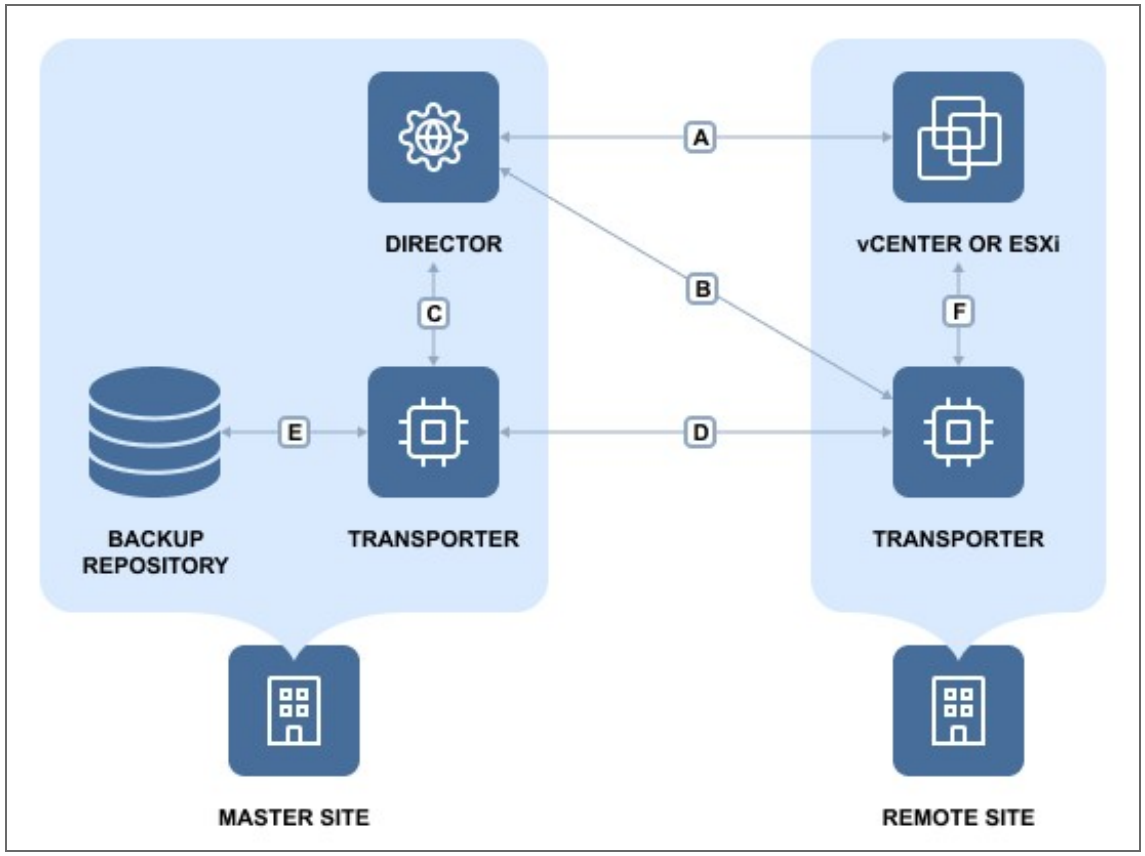

## Deployment Scenario

In this scenario, tenant VMs are running at remote sites and are backed up to a single master site.

> **Example**
> A service provider needs to back up customers' VMs to the service provider’s datacenter so that the customers don’t see each other’s backups and can recover their own files and emails through a self-service interface.

## Deployment Diagram

The deployment diagram for the above scenario is as follows:




## Deployment Steps

To deploy the above scenario, perform the following steps:

1. Install the Director in multi-tenant mode at the master site.
2. Install at least one Transporter at the master site.
3. Install at least one Transporter at each remote site.
4. For each tenant, prepare a separate folder at the master site for creating separate Backup Repositories.

## Connections

The implementation of the above scenario requires that the following connections be available:

| Connection | Description |
|---|---|
| **A** | Connection from the machine on which the Director is installed to vCenter servers and ESXi hosts that run source VMs at remote sites. The port used for communication with vCenter servers and ESXi hosts (443 by default) is open in firewalls. |
| **B** | Connection from the machine on which the Director is installed to machines at remote sites on which Transporters are installed. The port used for communication with the Transporters (9446 by default) is open in firewalls. |
| **C** | Connection from the machine on which the Director is installed to vCenter servers and ESXi hosts at the master site where VM replicas will be created. The port used for communication with vCenter servers and ESXi hosts (443 by default) is open in firewalls. |
| **D** | Connection from the machine on which the Director is installed to the machine at the master site where the Transporter is installed. The port used for communication with the Transporters (9446 by default) is open in firewalls. |
| **E** | Connection from the machine at the Master site where the Transporter is installed to ESXi hosts at the master site where VM replicas will be created. |
| **F** | Connection from the machine at the Master site where the Transporter is installed to machines at remote sites where Transporters are installed. The ports used for data transfer between a pair of Transporters are open in firewalls. |
| **G** | At remote sites, connections from machines on which Transporters are installed to vCenter servers and ESXi hosts running source VMs. |

**Note**

For security purposes, a VPN connection should be established between the master site and remote sites.

# Replication from Remote Site to Master Site


- Deployment Scenario
- Deployment Diagram
- Deployment Steps
- Connections

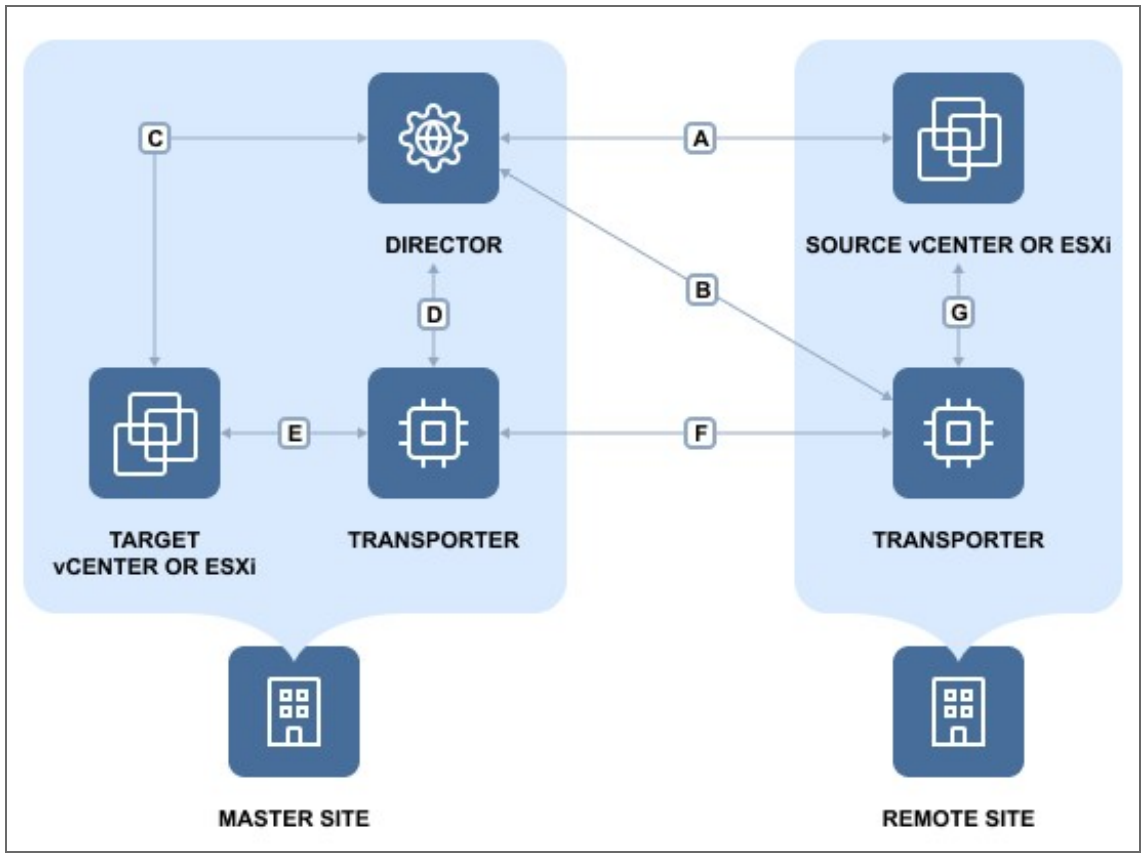

## Deployment Scenario

In this scenario, tenant VMs are running at remote sites and are replicated to a single master site.

> **Example**
> A service provider wants to introduce Replication-as-a-Service to customers and replicate their VMs to the service provider's datacenter.

## Deployment Diagram

The deployment diagram for the above scenario is as follows:




## Deployment Steps

To deploy the above scenario, perform the following steps:

1. Install the Director in multi-tenant mode at the master site.
2. Install at least one Transporter at the master site.
3. Install at least one Transporter at each remote site.
4. For each tenant, prepare a separate ESXi host that will serve as a replication target.

## Connections

The implementation of the above scenario requires that the following connections be available:

| Connection | Description |
|---|---|
| **A** | Connection from the machine on which the Director is installed to vCenter servers and ESXi hosts that run source VMs at remote sites. The port used for communication with vCenter servers and ESXi hosts (443 by default) is open in firewalls. |
| **B** | Connection from the machine on which the Director is installed to machines at remote sites on which Transporters are installed. The port used for communication with the Transporters (9446 by default) is open in firewalls. |
| **C** | Connection from the machine on which the Director is installed to vCenter servers and ESXi hosts at the master site where VM replicas will be created. The port used for communication with vCenter servers and ESXi hosts (443 by default) is open in firewalls. |
| **D** | Connection from the machine on which the Director is installed to the machine at the master site where the Transporter is installed. The port used for communication with the Transporters (9446 by default) is open in firewalls. |
| **E** | Connection from the machine at the master site where the Transporter is installed to ESXi hosts at the master site where VM replicas will be created. |
| **F** | Connection from the machine at the master site where the Transporter is installed to machines at remote sites where Transporters are installed. The ports used for data transfer between a pair of Transporters are open in firewalls. |
| **G** | At remote sites, connections from machines on which Transporters are installed to vCenter servers and ESXi hosts running source VMs. |

**Note**

For security purposes, a VPN connection should be established between the master site and remote sites.

# Local Backup at Remote Site


- Deployment Scenario
- Deployment Diagram
- Deployment Steps
- Connections

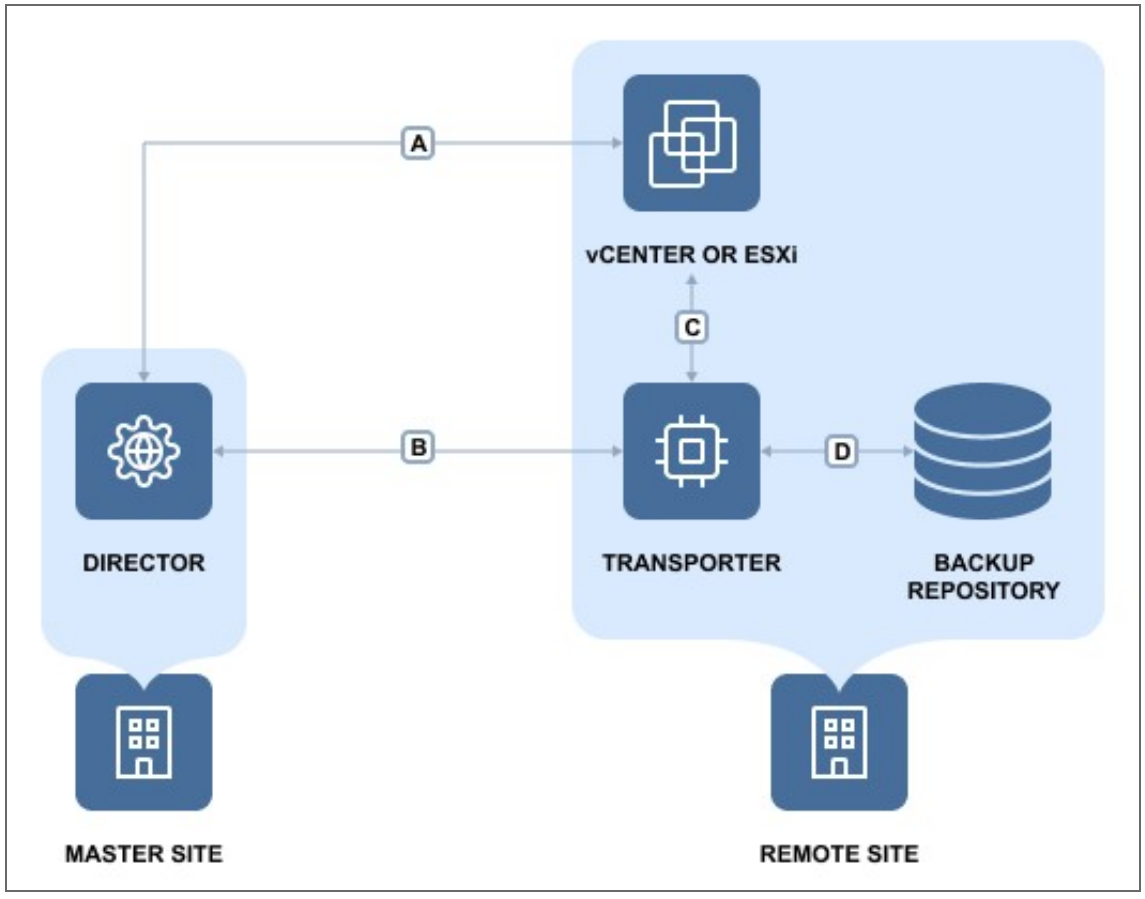

## Deployment Scenario

In this scenario, tenant VMs are running and backed up locally at the remote sites.

> **Example**
> An Enterprise has two branch offices running VMware virtual infrastructure. The IT manager located at the headquarters is responsible for the Enterprise data protection and needs to back up VMs locally at their branch offices to ensure fast operational recovery. Employees of the branch offices should have access to their VM backups and be able to recover their files and emails.

## Deployment Diagram

The deployment diagram for the above scenario is as follows:



## Deployment Steps

To deploy the above scenario, perform the following steps:

1. Install the Director in multi-tenant mode at the master site.
2. Install at least one Transporter at each remote site.
3. For each tenant, prepare a separate folder at a remote site for creating a Backup Repository.

## Connections

The implementation of the above scenario requires that the following connections be available:

| Connection | Description |
|---|---|
| **A** | Connection from the machine on which the Director is installed to vCenter servers and ESXi hosts that run source VMs at remote sites. The port used for communication with vCenter servers and ESXi hosts (443 by default) is open in firewalls. |
| **B** | Connection from the machine on which the Director is installed to machines at remote sites where the Transporters are installed. The port used for communication with the Transporters (9446 by default) is open in firewalls. |
| **C** | Connection from the machines on which the Transporters are installed at remote sites to vCenter servers and ESXi hosts running source VMs. |

**Note**

For security purposes, a VPN connection should be established between the master site and remote sites.

# Local Replication at Remote Site


- Deployment Scenario
- Deployment Diagram
- Deployment Steps
- Connections

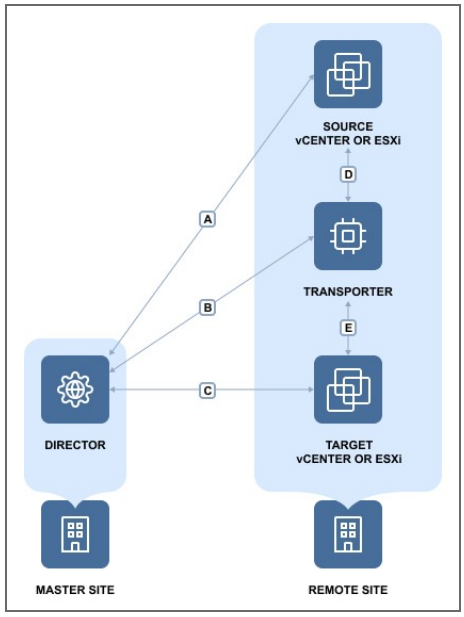

## Deployment Scenario

In this scenario, tenant VMs are running and replicated locally at the remote sites.

> **Example**
> An Enterprise has two branch offices running VMware virtual infrastructure. The IT manager located at the headquarters is responsible for the Enterprise data protection and needs to replicate business critical VMs locally at the branch offices for high availability.

## Deployment Diagram

The deployment diagram for the above scenario is as follows:




## Deployment Steps

To deploy the above scenario, perform the following steps:

1. Install the Director in multi-tenant mode at the master site.
2. Install at least one Transporter at each remote site.
3. For each tenant, prepare a separate folder at the remote site for creating a Backup Repository.

## Connections

The implementation of the above scenario requires that the following connections be available:

| Connection | Description |
|---|---|
| A | Connection from the machine on which the Director is installed to vCenter servers and ESXi hosts that run source VMs at remote sites. The port used for communication with vCenter servers and ESXi hosts (443 by default) is open in firewalls. |
| B | Connection from the machine on which the Director is installed to machines at remote sites where Transporters are installed. The port used for communication with the Transporters (9446 by default) is open in firewalls. |
| C | Connection from the machine on which the Director is installed to vCenter servers and ESXi hosts where VM replicas will be created at remote sites. The port used for communication with vCenter servers and ESXi hosts (443 by default) is open in firewalls. |
| D | At remote sites, connections from machines where Transporters are installed to vCenter servers and ESXi hosts running source VMs. |
| E | At remote sites, connections from machines where Transporters are installed to vCenter servers and ESXi hosts where VM replicas will be created. |

**Note**

For security purposes, a VPN connection should be established between the master site and remote sites.

# Backup at Master Site


- Deployment Scenario
- Deployment Diagram
- Deployment Steps
- Connections

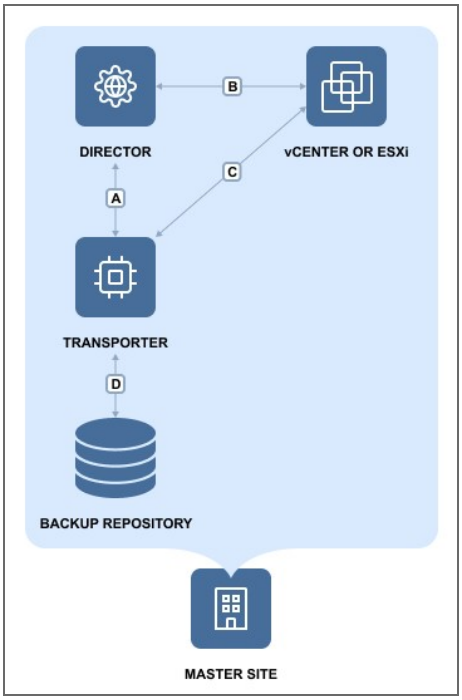

## Deployment Scenario

In this scenario, tenant VMs are running at the master site and the backing up of tenant VMs is also performed at the master site.

> **Example**
> A service provider runs VMs of customer A and customer B in the service provider's datacenter. The Service Provider seeks to offer Backup-as-a-Service to both customers. The customers should be able to recover their files and emails through a self-service interface without being able to see each other's backups.

## Deployment Diagram

The deployment diagram for the above scenario is as follows:




## Deployment Steps

To deploy the above scenario, perform the following steps:

1. Install the Director in multi-tenant mode at the master site.
2. Install at least one Transporter at the master site.

3. For each tenant, prepare a separate folder at the master site for creating a Backup Repository.

## Connections

The implementation of the above scenario requires that the following connections be available:

| Connection | Description |
|---|---|
| **A** | Connection from the machine on which the Director is installed to the machine on which the Transporter is installed. The port used for communication with the Transporters (9446 by default) is open in firewalls. |
| **B** | Connection from the machine on which the Director is installed to vCenter servers and ESXi hosts that run source VMs. The port used for communication with vCenter servers and ESXi hosts (443 by default) is open in firewalls. |
| **C** | Connection from the machine on which the Transporter is installed to vCenter servers and ESXi hosts running source VMs. |
| **D** | Connection from the machine on which the Transporter is installed to the folders where tenant Backup Repositories will be created. |

**Note**

For security purposes, a VPN connection should be established between the master site and remote sites.

# Replication at Master Site


- Deployment Scenario
- Deployment Diagram
- Deployment Steps
- Connections

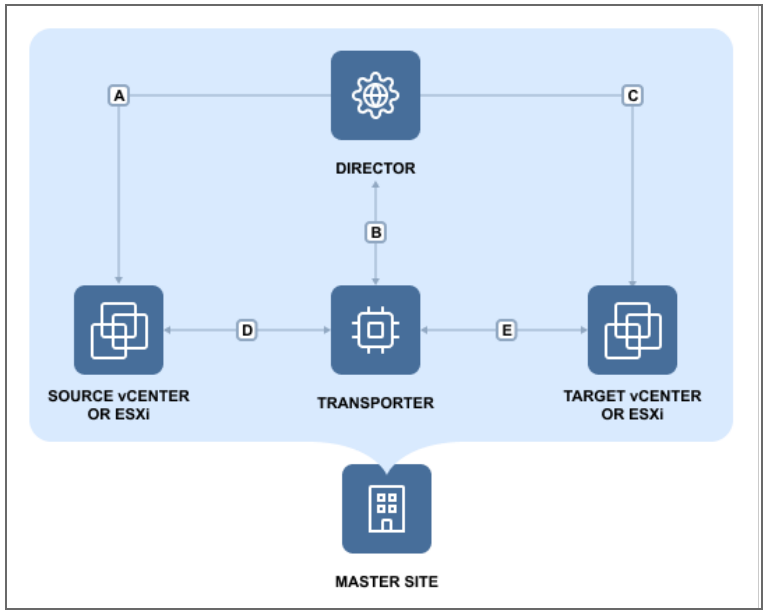

## Deployment Scenario

In this scenario, tenant VMs are running at the Master site and the replication of tenant VMs is also performed at the Master site.

> **Example**
> A service provider runs customers' VMs in the service provider's datacenter. To ensure high availability of tenant VMs, the service provider seeks to replicate customer VMs to a different server.

## Deployment Diagram

The deployment diagram for the above scenario is as follows:




## Deployment Steps

To deploy the above scenario, perform the following steps:

1. Install the Director in multi-tenant mode at the master site.
2. Install at least one Transporter at the master site.
3. For each tenant, prepare a separate ESXi host that will serve as a replication target.

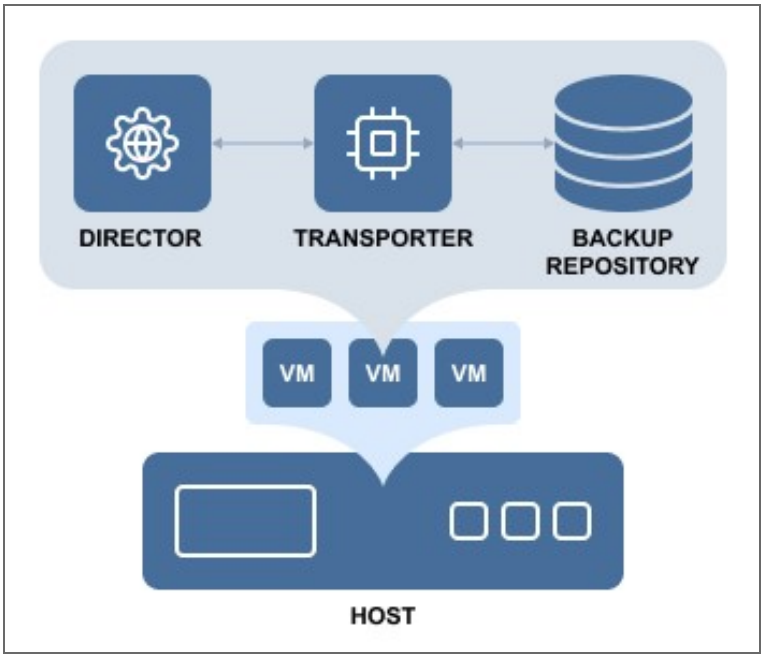
## Connections

The implementation of the above scenario requires that the following connections be available:

| Connection | Description |
|---|---|
| **A** | Connection from the machine on which the Director is installed to vCenter servers and ESXi hosts that run source VMs. The port used for communication with vCenter servers and ESXi hosts (443 by default) is open in firewalls. |
| **B** | Connection from the machine on which the Director is installed to the machine on which the Transporter is installed. The port used for communication with the Transporters (9446 by default) is open in firewalls. |
| **C** | Connection from the machine on which the Director is installed to vCenter servers and ESXi hosts where VM replicas will be created. The port used for communication with vCenter servers and ESXi hosts (443 by default) is open in firewalls. |
| **D** | Connection from the machine on which the Transporter is installed to vCenter servers and ESXi hosts running source VMs. |
| **E** | Connection from the machine on which the Transporter is installed to vCenter servers and ESXi hosts where VM replicas will be created. |

# Single Site Deployment

For a single site deployment, it is often sufficient to install both the Director and Transporter on a single VM/physical machine within your infrastructure.




This deployment provides you with the ability to back up, replicate, and recover multiple VMs from multiple source hosts.

# Installing NAKIVO Backup & Replication

Refer to the sections below to learn how to install NAKIVO Backup & Replication:


- "Deploying VMware Virtual Appliance" on page 232
- "Deploying Nutanix AHV Virtual Appliance" on page 239
- "Deploying Amazon Machine Image in Amazon EC2" on page 247
- "Installing on Windows" on page 250
- "Installing on Linux" on page 261
- "Installing on Synology NAS" on page 270
- "Installing on QNAP NAS" on page 277
- "Installing on Western Digital NAS" on page 283
- "Installing on ASUSTOR NAS" on page 285
- "Installing on NETGEAR ReadyNAS" on page 290
- "Installing on Generic ARM-Based Devices" on page 293
- "Installing on TrueNAS" on page 294


**Note**

Before installing your NAKIVO Backup & Replication solution, make sure that your system has been updated with the latest patch and all the necessary requirements are met.

# Deploying VMware Virtual Appliance

- Deploying Virtual Appliance with vSphere Web Client
- Virtual Appliance OS, Credentials, and Security
- Web Interface Login

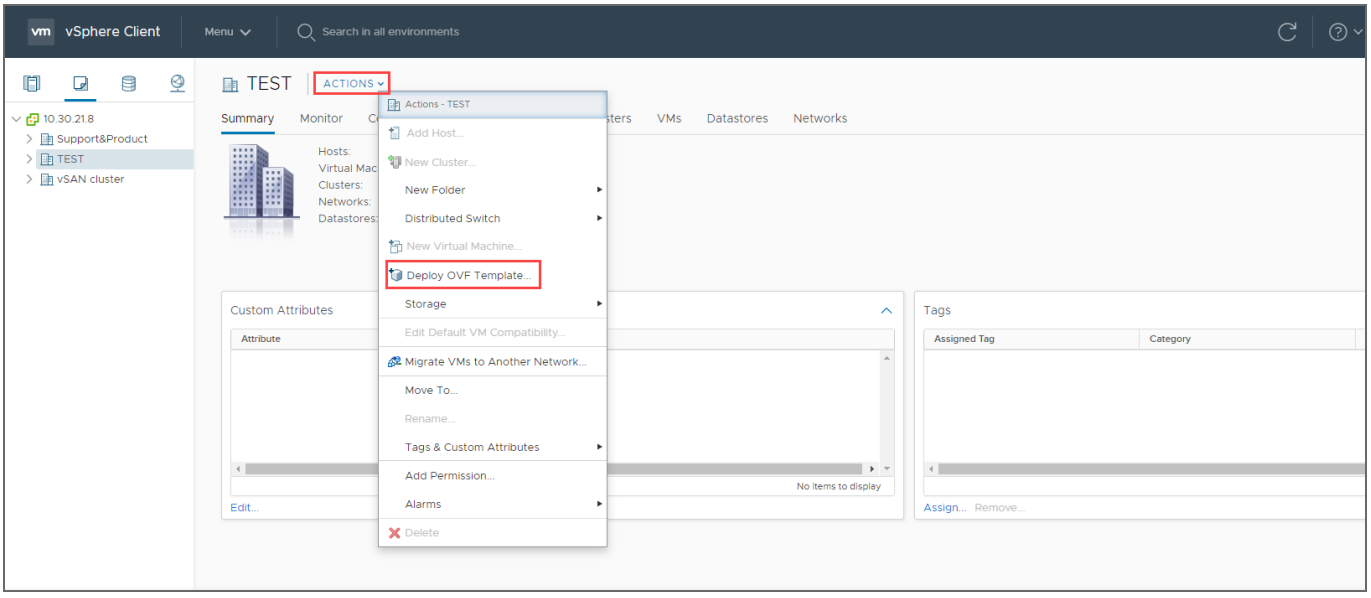
NAKIVO Backup & Replication offers the following VA deployment options:

- Full Solution
- Full Solution without **Backup Repository**
- **Transporter**-only
- **Transporter** with **Backup Repository**
- Multi-tenant **Director**

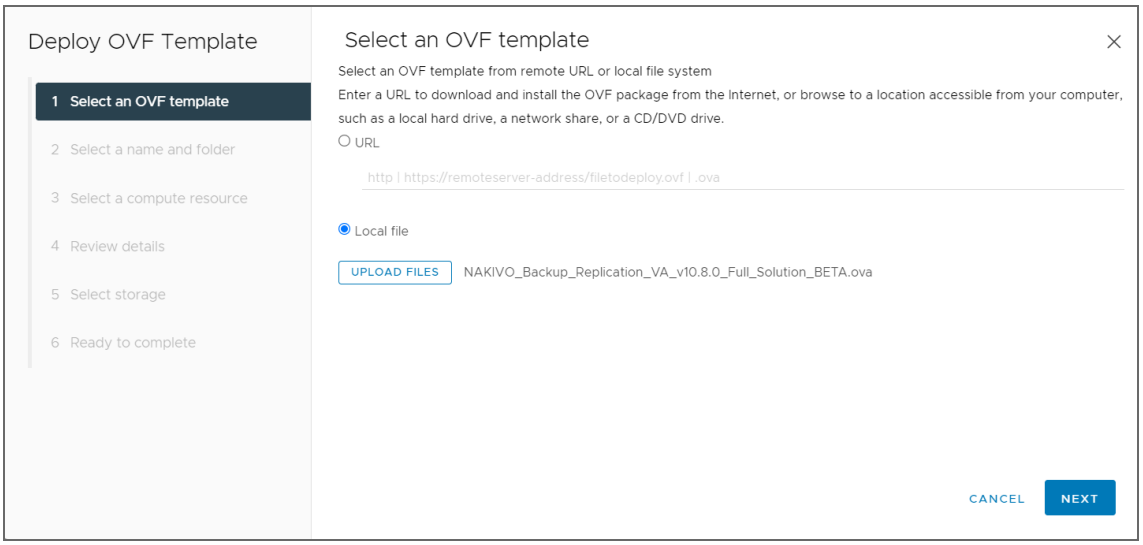
The Virtual Appliance (VA) has two disks: the first (30 GB) contains a Linux OS with NAKIVO Backup & Replication, and the second (500 GB) is used as a **Backup Repositor**y. If you deploy the Virtual Appliance disks using the **Thin Provision** option, then the disks will not reserve space on your datastore and will only consume space when actual data (such as your backups) is written to disks.

## Deploying Virtual Appliance with vSphere Web Client

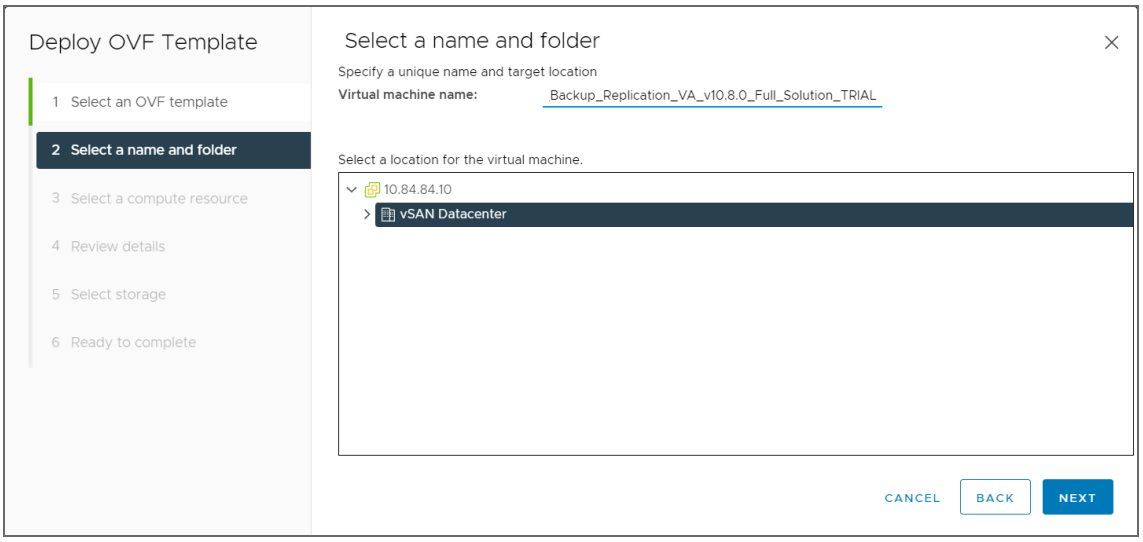
1. Download NAKIVO Backup & Replication VA.
2. Log in to your vSphere vCenter with the vSphere Web Client.
3. Select **Deploy OVF Template** from the **Actions** menu. Note that the Client Integration Plug-in must be installed to enable OVF functionality.



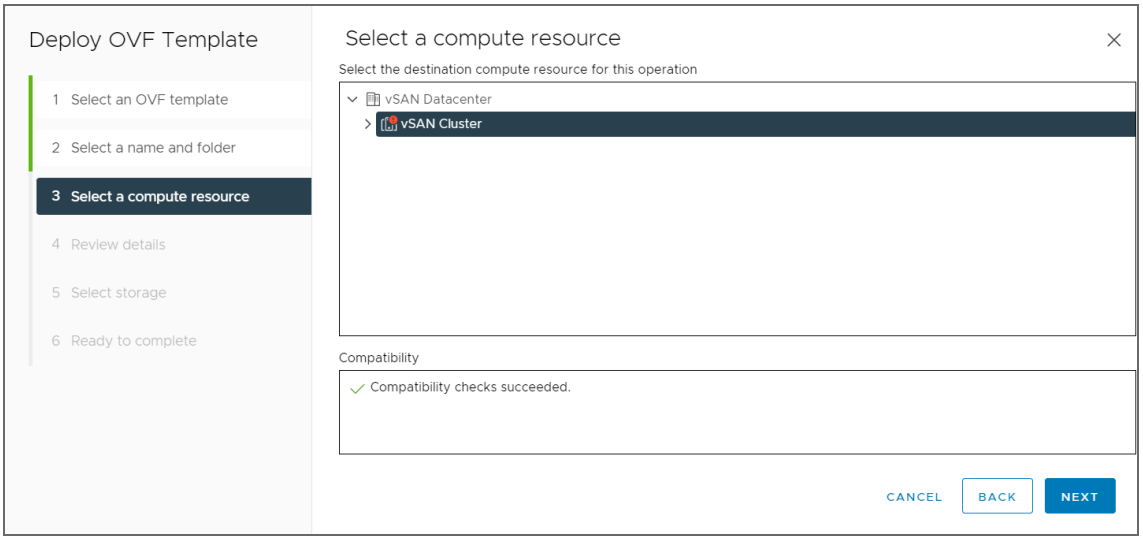

4. On the **Select an OVF template** page of the **Deploy OVF Template** wizard, select **Local file** and upload the VA file (.ova) you've downloaded. Click **Next**.



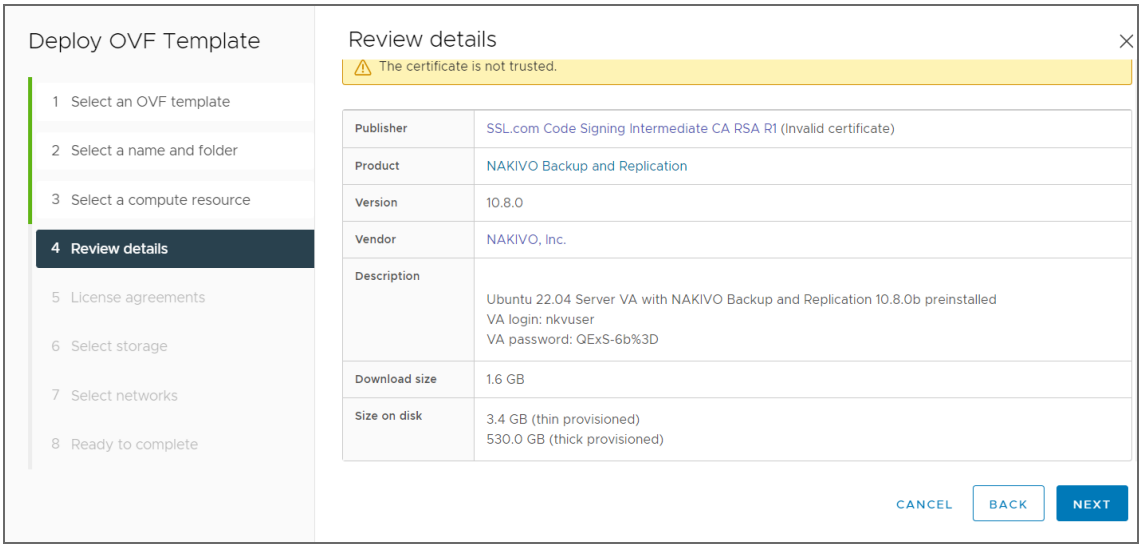

5. On the **Select a name and folder** page, specify a unique name and target location for the Virtual Appliance. Click **Next**.



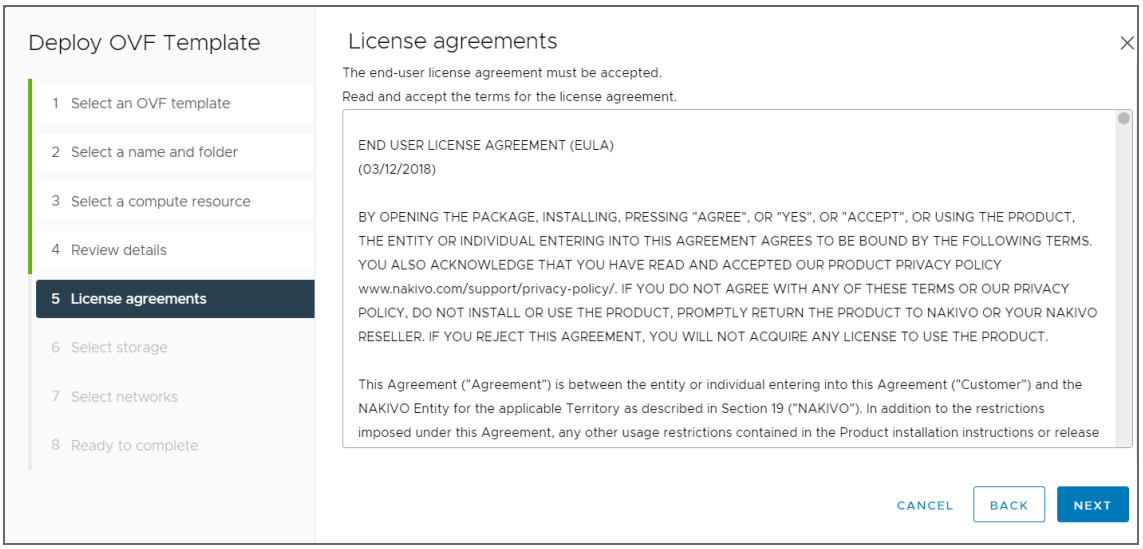

6. On the **Select a computer resource** page, select the resource pool within which you would like to deploy the Virtual Appliance and click **Next**.



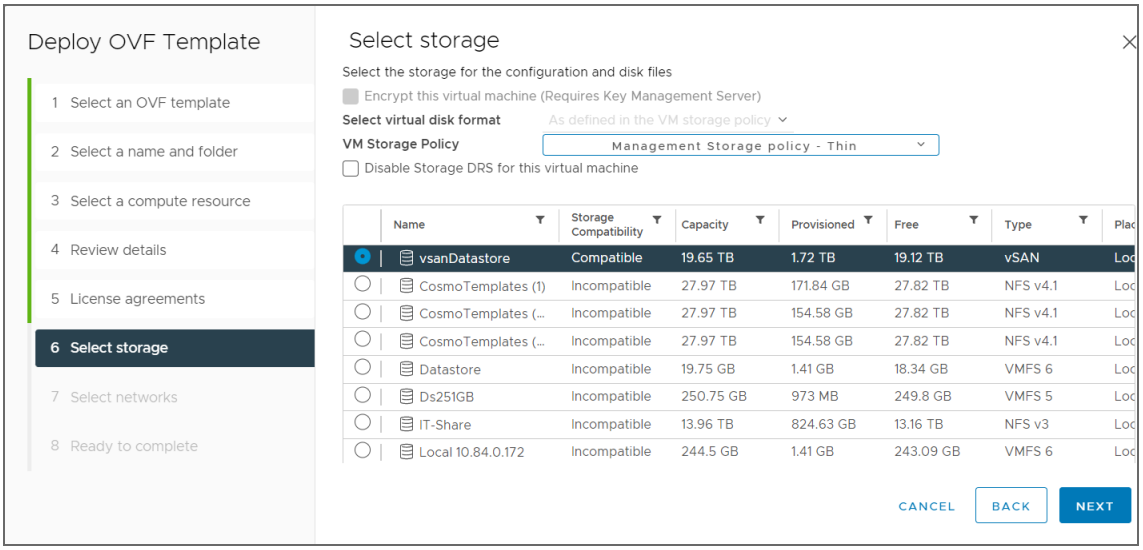

7. On the **Review details** page, review the template details and click **Next**.




8. On the **License agreements** page, read the end-user license agreement (EULA). If you agree to its terms, select **I accept all license agreements** and then click **Next**.



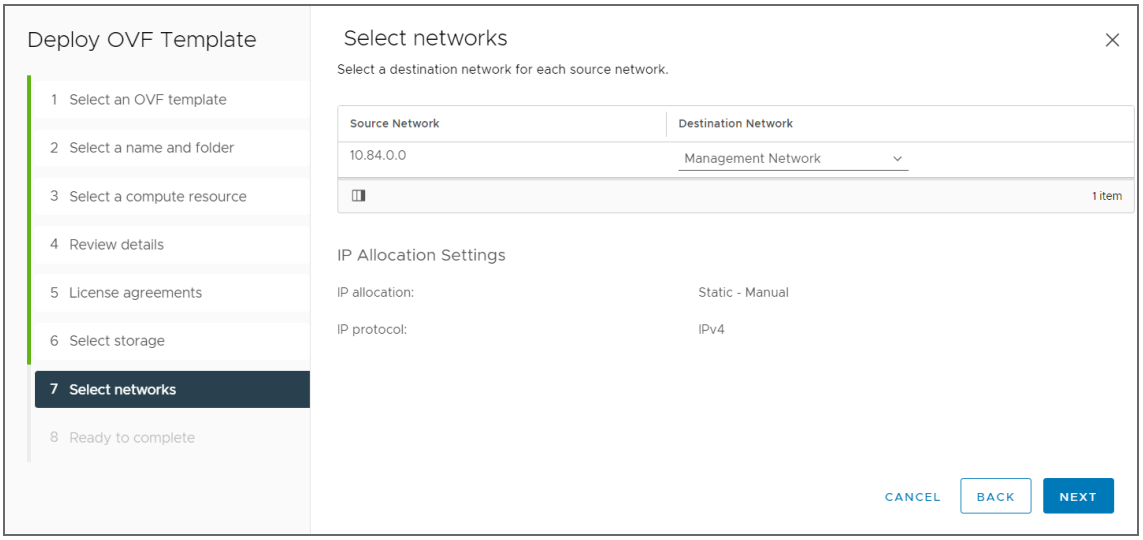

9. On the **Select storage** page, select a datastore in which you would like to keep the Virtual Appliance disk, virtual disk format (*Thin Provisioning* is recommended), VM storage policy and click **Next**.



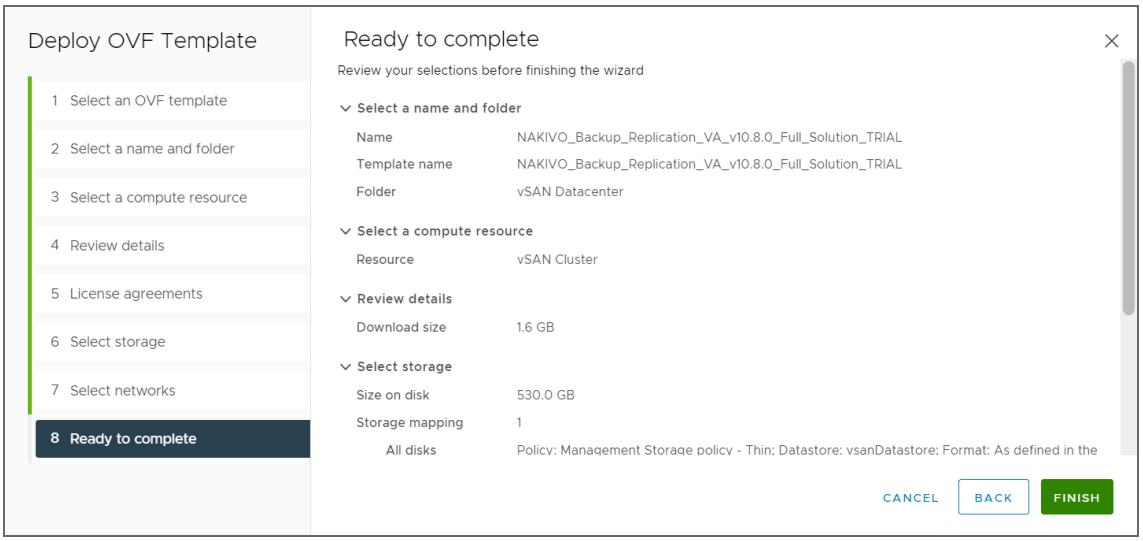

**Important**

If you use thick provisioning instead of thin provisioning, keep in mind that NAKIVO Backup & Replication can take up to 0,5 TB of data. Check to see if it is 0,5 TB by default for all cases.

10. On the **Select networks** page, select a network to which the Virtual Appliance will be connected. Opting for a network with DHCP and Internet access is recommended. Click **Next**.




11. On the **Ready to complete** page, review the summary of the setups you have configured and click **Finish** to complete deployment.




After the Virtual Appliance is deployed, you may need to configure it.

**Important**

If you plan to expose the Virtual Appliance to the Internet, change the default credentials and set up a login and password for the Web interface.

# Virtual Appliance OS, Credentials, and Security

The appliance runs Ubuntu 24.04, 64-bit. Use the following credentials to log in to the appliance:

- **Username**: nkvuser
- **Password**: QExS-6b%3D

For the versions of NAKIVO Backup & Replication older than 7.2, the password is `root`.

**Important**

- If you plan to expose the Virtual Appliance to the Internet, change the default VA credentials and set up a login and password for the Web interface.
- It is recommended to run an update on all packages in your Virtual Appliance at least once a month.

To enable Backup Immutability for **Amazon S3**, **Wasabi**, **Azure Blob Storage**, **Backblaze B2 Cloud Storage**, or **Local Folder** types of Backup Repository deployed as part of virtual appliance, NAKIVO Backup & Replication does the following:

- Creates a new user for all administrative needs and adds it to the *sudo* group
- Disables root user
- Changes default SSH port to 2221
- Configure the following kernel parameters via **sysctl.conf**:
  - Limits network-transmitted configuration for IPv4/IPv6
  - Prevents the common 'syn flood attack'
  - Turns on source IP address verification
  - Prevents a cracker from using a spoofing attack against the IP address of the server
  - Logs several types of suspicious packets, such as spoofed packets, source-routed packets, and redirects
  - Configures swap. Sets **vm.swappiness** to 15
  - Sets **kernel.unprivileged_bpf_disabled** to 1
  - Sets **kernel.core_pattern** to **/tmp/%e.%p.core**
  - Sets **kernel.core_uses_pid** to 1
  - Sets **kernel.dmesg_restrict** to 1
  - Sets **kernel.kptr_restrict** to 2
  - Sets **kernel.sysrq** to 0
- Secures */tmp* and */var/tmp*
- Secures Shared Memory
- Installs and configures **fail2ban**

**Notes**

- After **fail2ban** is installed on the hardened VA, the user IP may be banned for 10 minutes if mistakes have been made during the login procedure.
- Any additional packages installed manually on the system may cause a security breach.

## Web Interface Login

Open your browser and go to the VA's web interface. Refer to the Getting Started section to better understand how to continue working with NAKIVO Backup & Replication.

# Deploying Nutanix AHV Virtual Appliance


- Deploying Nutanix AHV Virtual Appliance
- Virtual Appliance OS, Credentials, and Security
- Web Interface Login


## Deploying Nutanix AHV Virtual Appliance

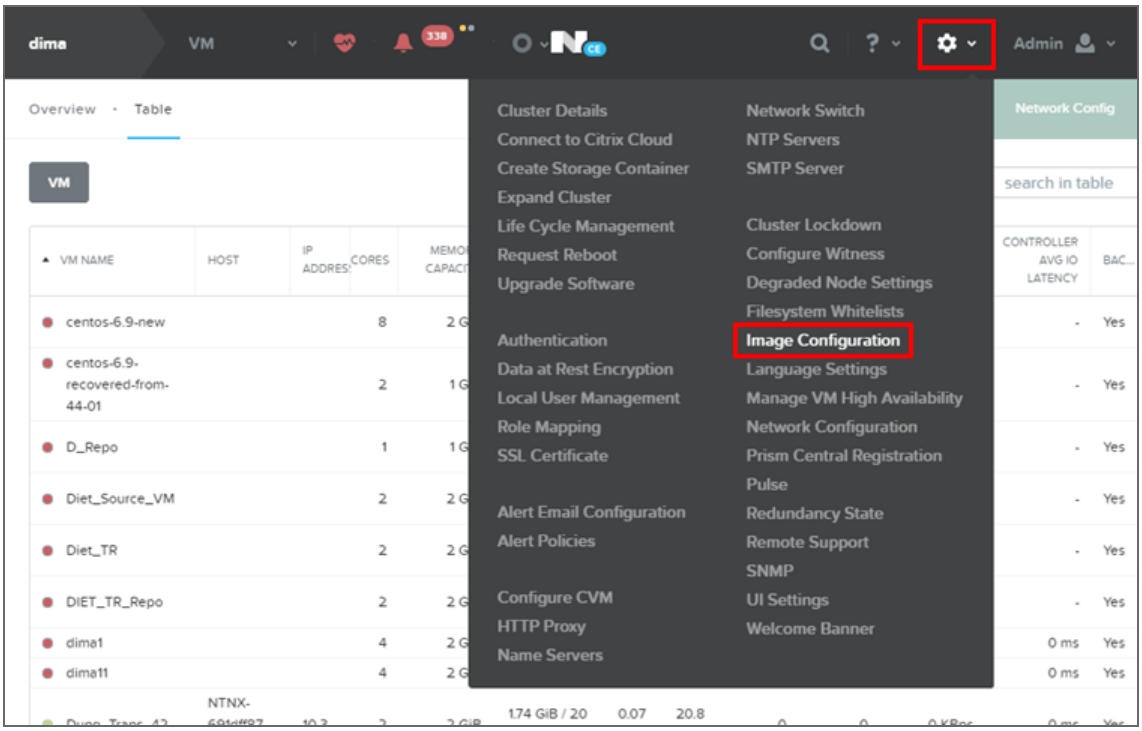
The NAKIVO Backup & Replication instance must be deployed in a Nutanix AHV cluster in order to enable backup and recovery functions.

NAKIVO Backup & Replication offers the following solutions:

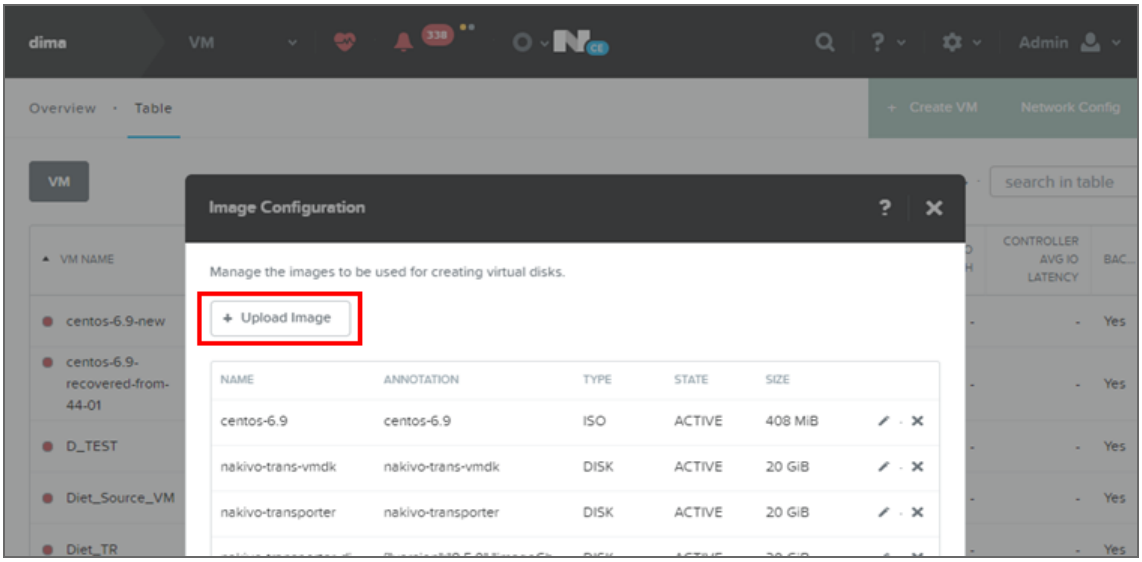
- Full Solution (Single Tenant) - requires a 100 GB thin provisioned disk
- **Transporter**-only - requires a 20 GB thin provisioned disk

To deploy a virtual appliance via the Nutanix Prism application, follow the steps below:

1. Download the `.VMDK` file with a full or transporter-only image from the Nakivo website and store it locally.
2. Log in to the Prism console.
3. From the **Configurations** menu, select **Image Configuration**.

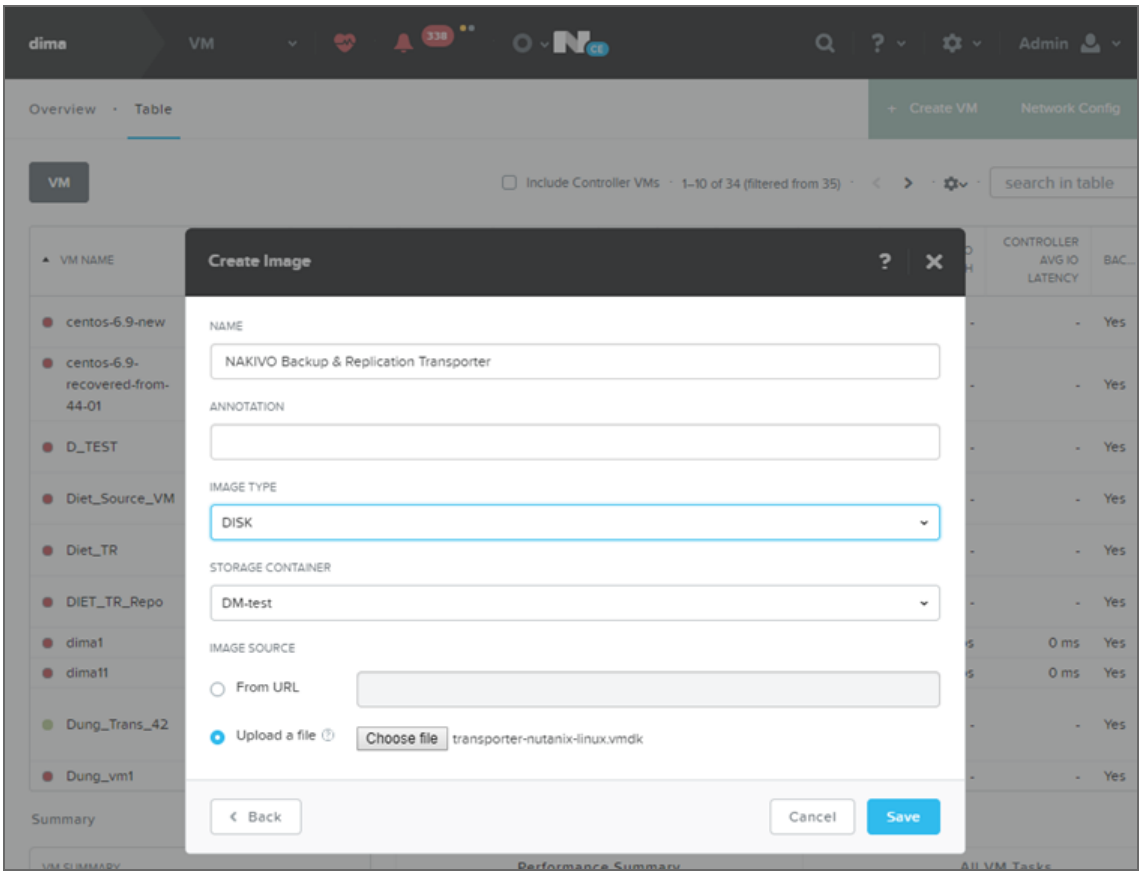



4. In the **Image Configuration** dialog, click **Upload Image**.

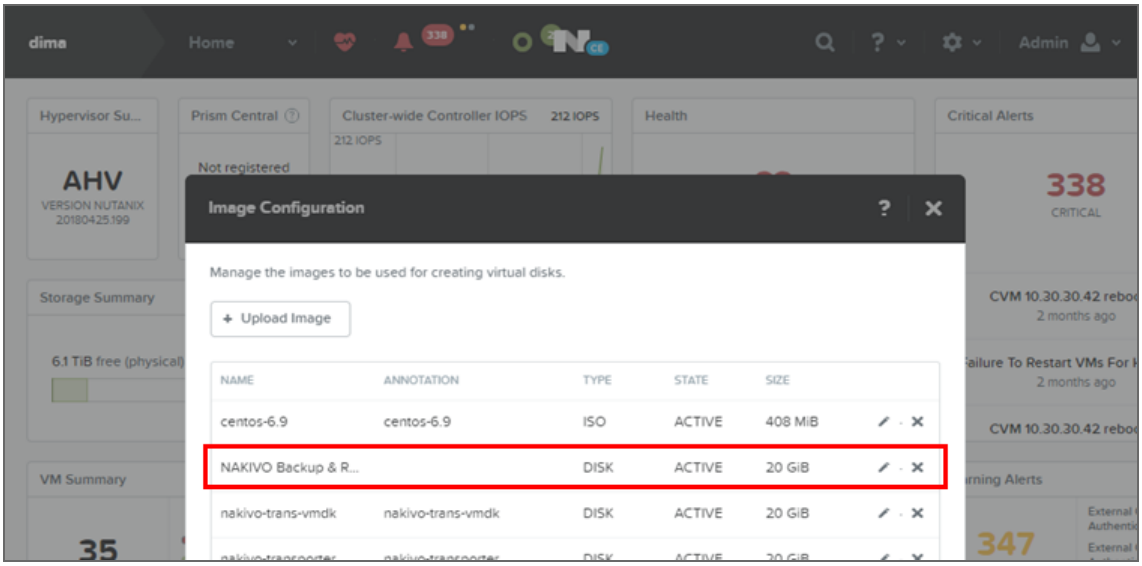



5. In the **Create Image** dialog, fill in the following options:
    - **Name**: Enter a name for the new image.
    - **Image Type**: From the drop-down list, select **DISK**.
    - **Storage Container**: Select the storage container you wish to use from the drop-down list. The list includes all storage containers created for this cluster. If there are no storage containers currently available, a **Create Storage Container** link is displayed.
    - **Image Source**: Click the **Upload a file** radio button to upload a file from your workstation. Click

the **Choose File** button and then select the file to upload from the file search window.



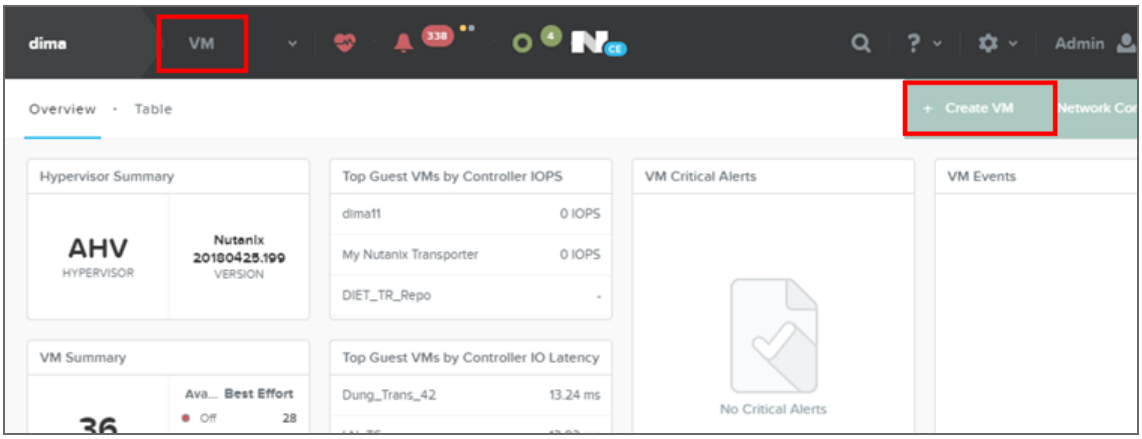

6. When all fields are correct, click the **Save** button.
   After the file uploading completes, the **Create Image** window closes and the **Image Configuration** window reappears with the new image present in the list.




**Note**

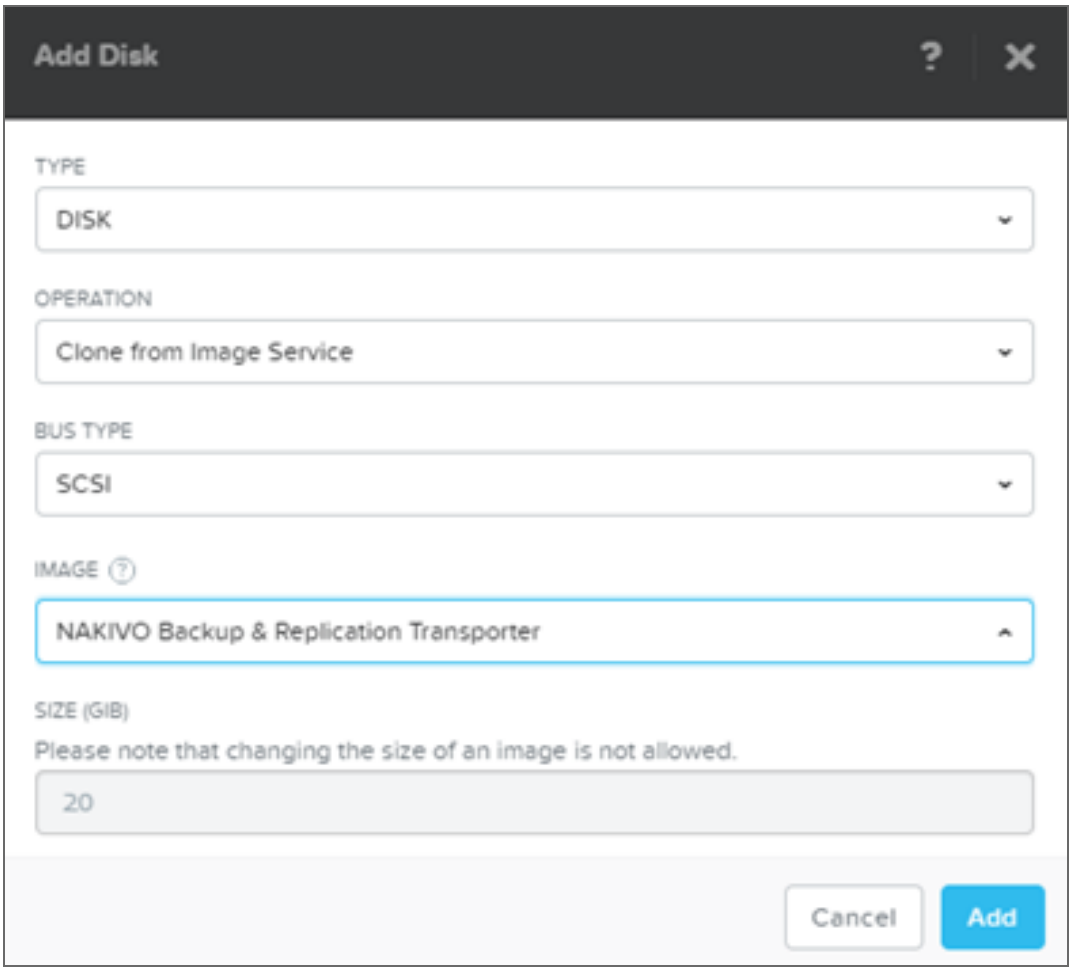
Make sure the status of the disk is **Active** before proceeding to the next step.

7. Close the **Image Configuration** window, go to the **VM** view and click **Create VM**.



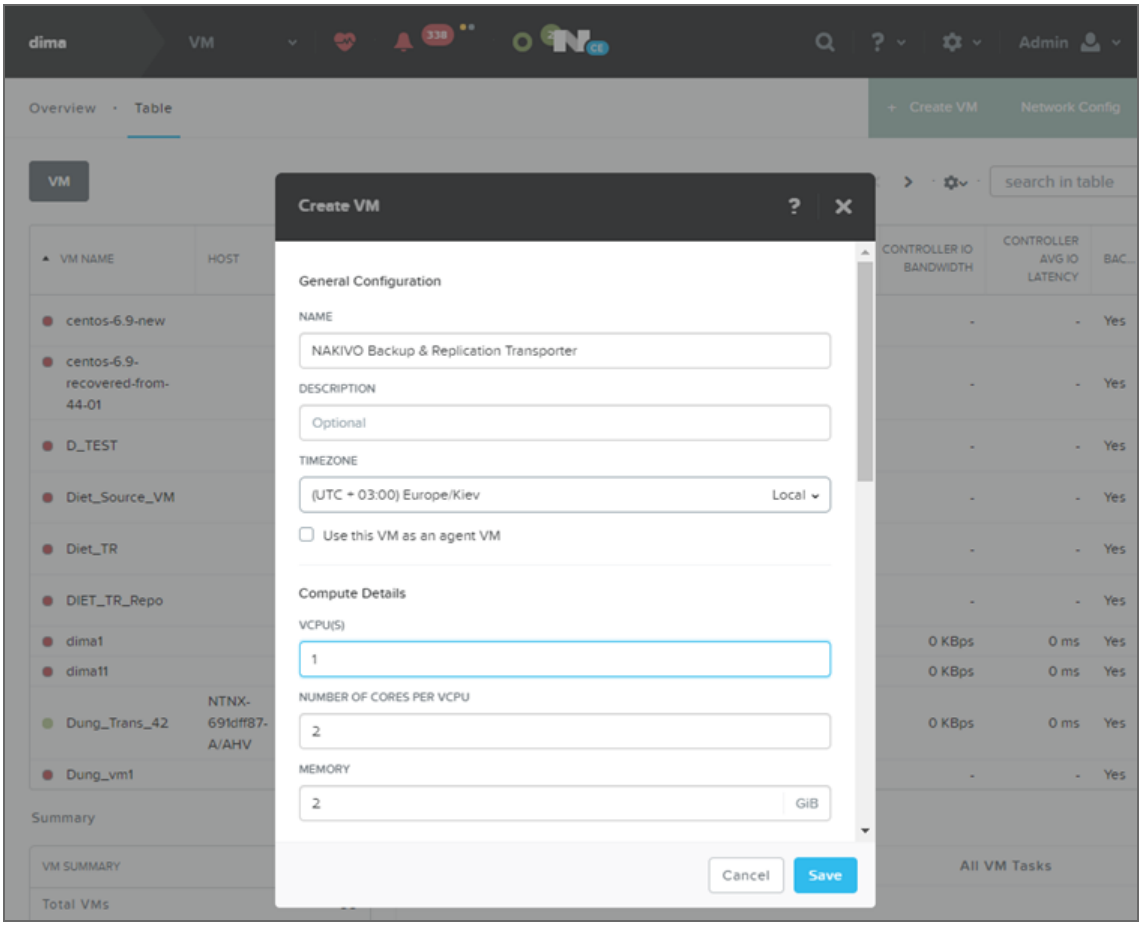

8. In the **Create VM** dialog, fill in the following options:
   - **Name**: Enter a name for the VM.
   - **vCPU(s)**: Enter the number of virtual CPUs to allocate to this VM (minimum 1).
   - **Number of Cores per vCPU**: Enter the number of cores assigned to each virtual CPU (minimum 2).
   - **Memory**: Enter the amount of memory (in GBs) to allocate to this VM (minimum 4 GB + 250 MB for each concurrent job for full solution/minimum 2 GB + 250 MB for each concurrent job **Transporter**-only solution).
   - In the **Disk** section, click **Add New Disk**, and specify the following settings in the **Add Disk** dialog:
     a. **Type**: Select **Disk**.
     b. **Operation**: Select **Clone from Image Service**.
     c. **Bus Type**: Select **SCSI**.

d. **Image**: Select your uploaded image from the list.

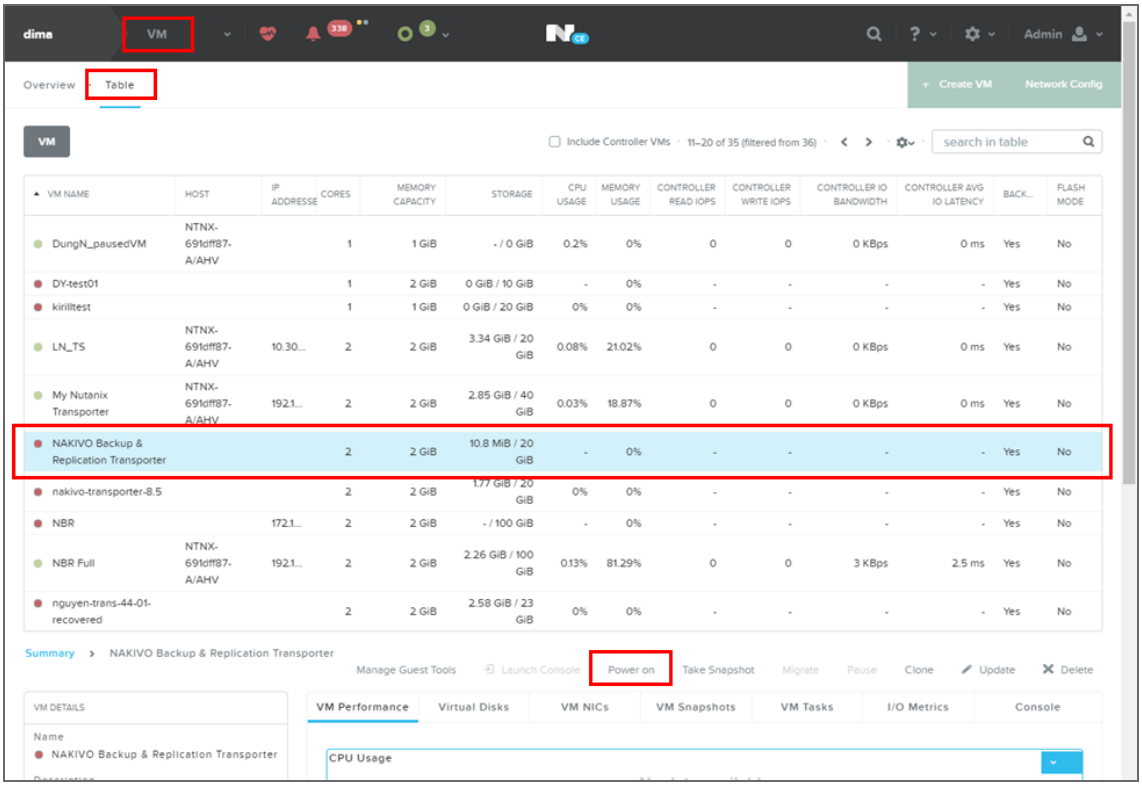



- In the **Network Adapters (NIC)** section, click **Add New NIC** and select an available VLAN from the list.

9. Click **Save**.




10. Wait until the process of VM creation is complete and locate your newly-created VM on the list.
11. Select your VM and click **Power On**.




12. After the Virtual Appliance is deployed and powered on, you may need to configure it.

## Virtual Appliance OS, Credentials, and Security

The appliance runs Ubuntu 24.04, 64-bit. Use the following credentials to log in to the appliance:

- **Username**: nkvuser
- **Password**: QExS-6b%3D

For the versions of NAKIVO Backup & Replication older than 7.2, the password is `root`.

**Important**

- If you plan to expose the Virtual Appliance to the Internet, change the default VA credentials and set up a login and password for the Web interface.
- It is recommended to run an update on all packages in your Virtual Appliance at least once a month.

To enable Backup Immutability for **Amazon S3**, **Wasabi**, **Azure Blob Storage**, **Backblaze B2 Cloud Storage**, or **Local Folder** types of Backup Repository deployed as part of virtual appliance, NAKIVO Backup & Replication does the following:

- Creates a new user for all administrative needs and adds it to the *sudo* group
- Disables root user
- Changes default SSH port to *2221*
- Configure the following kernel parameters via **sysctl.conf**:

    - Limits network-transmitted configuration for IPv4/IPv6
    - Prevents the common 'syn flood attack'
    - Turns on source IP address verification
    - Prevents a cracker from using a spoofing attack against the IP address of the server
    - Logs several types of suspicious packets, such as spoofed packets, source-routed packets, and redirects
    - Configures swap. Sets **vm.swappiness** to 15
    - Sets **kernel.unprivileged_bpf_disabled** to 1
    - Sets **kernel.core_pattern** to **/tmp/%e.%p.core**
    - Sets **kernel.core_uses_pid** to 1
    - Sets **kernel.dmesg_restrict** to 1
    - Sets **kernel.kptr_restrict** to 2
    - Sets **kernel.sysrq** to 0
- Secures */tmp* and */var/tmp*
- Secures Shared Memory
- Installs and configures **fail2ban**

**Notes**

- After **fail2ban** is installed on the hardened VA, the user IP may be banned for 10 minutes if mistakes have been made during the login procedure.
- Any additional packages installed manually on the system may cause a security breach.

## Web Interface Login

Open your browser and go to the VA web interface. Refer to "Getting Started" on page 333 to better understand how to continue working with NAKIVO Backup & Replication.

# Deploying Amazon Machine Image in Amazon EC2

You can deploy NAKIVO Backup & Replication as a pre-configured Amazon Machine Image (AMI) in Amazon EC2. After you complete the download form, you get a link to the AWS marketplace page where you can download the AMI.

## Configuring AMI Parameters

Configure the following AMI parameters:

1. **Instance Type**: More powerful instances can process tasks faster and run more tasks simultaneously. The minimum requirement for NAKIVO Backup & Replication is the t2.micro instance type; the t2 medium instance type is recommended.
2. **Instance Details**: Assign a public IP to the instance if you wish to access the instance over the Internet.
   - **Security Group**: Use the "All Traffic" rule or create a set of rules listed below:

| Type | Port Range | Source | Description |
|---|---|---|---|
| SSH | 2221 | 0.0.0.0/0 | Enables remote SSH access to the instance |
| Custom TCP | 80 | 0.0.0.0/0 | Enables access to the web interface |
| Custom TCP | 443 | 0.0.0.0/0 | Required for local **Transporter** import |
| Custom TCP | 902 | 0.0.0.0/0 | Required for local **Transporter** import |
| Custom TCP | 4443 | 0.0.0.0/0 | Enables access to the web interface |
| Custom TCP | 9446 | 0.0.0.0/0 | Enables access to a remote **Transporter** |
| Custom TCP | 9448-10000 | 0.0.0.0/0 | Enables access to a remote **Transporter** |
| All ICMP | 0-65535 | 0.0.0.0/0 | Enables access to a remote **Transporter** |

> **Note**
> Older AMIs may still use SSH Port *22* instead of *2221*.

4. **Key pair**: Select an existing key pair or create a new key pair for your instance. If you select an existing key pair, make sure you have access to the private key file.

> **Note**
> The AMI deliverable uses Ubuntu 22.04 OS and a standalone EC2 instance with a **Director** and **Transporter**. Instead of the default system user **ubuntu**, the AMI uses the username **nkvuser**.

Refer to "Getting Started" on page 333 to better understand how to continue working with NAKIVO Backup & Replication.

## Security

The security of your backups can be significantly improved with "Backup Immutability" on page 48. For this feature to be available, the backups must be stored in the **Amazon S3**, **Wasabi**, **Azure Blob Storage**, **Backblaze B2 Cloud Storage**,or **Local Folder** types of **Backup Repository** deployed via AWS AMI on your EC2 instance.

To enable Backup Immutability for a **Local folder** type of **"Backup Repository" on page 127** deployed via an AMI, NAKIVO Backup & Replication does the following:

- Creates a new user for all administrative needs and adds it to the *sudo* group
- Disables root user
- Changes default SSH port to *2221*
- Configures the following kernel parameters via **sysctl.conf**:
  - Limits network-transmitted configuration for IPv4/IPv6
  - Prevents the common 'syn flood attack'
  - Turns on source IP address verification
  - Prevents a cracker from using a spoofing attack against the IP address of the server
  - Logs several types of suspicious packets, such as spoofed packets, source-routed packets, and redirects
  - Configures swap. Sets **vm.swappiness** to 15
  - Sets **kernel.unprivileged_bpf_disabled** to 1
  - Sets **kernel.core_pattern** to **/tmp/%e.%p.core**
  - Sets **kernel.core_uses_pid** to 1
  - Sets **kernel.dmesg_restrict** to 1
  - Sets **kernel.kptr_restrict** to 2
  - Sets **kernel.sysrq** to 0
- Secures */tmp* and */var/tmp*
- Secures Shared Memory
- Installs and configures **fail2ban**
- Uninstalls **multipath**
- Disables **snapd**
- Installs the following packets:
  - **nfs-common**
  - **ecryptfs-utils**
  - **cryptsetup**

**Notes**

- After **fail2ban** is installed on the hardened AMI, the user IP may be banned for 10 minutes if mistakes have been made during the login procedure.
- Any additional packages installed manually on the system may cause a security breach.
- It is possible to ping a hardened AMI.

# Installing on Windows

NAKIVO Backup & Replication offers the following installation options for Windows machines:

- Full Solution
- **Transporter**-Only Solution
- Multi-Tenant Solution

After successful product installation, refer to the Getting Started section to learn how to continue working with NAKIVO Backup & Replication.

- Installing Full Solution on Windows
- Installing Transporter-Only on Windows
- Installing Full Solution in Multi-Tenant Mode on Windows
- Silent Installation

## Installing Full Solution on Windows

To install NAKIVO Backup & Replication with default options, simply run the NAKIVO Backup & Replication installer for Windows and click **Install**. This will install all product components (**Director**, **Transporter**, and **Backup Repository**) and you will be able to use all product features after installation.

1. Set the installation options as follows:
   - **Installation type**: Leave the **Full solution** option selected to install the key product components (Director and Transporter)
   - **Create repository**: Leave the checkbox selected to create a **Backup Repository** on the machine on which NAKIVO Backup & Replication is installed.
   - Optionally, click **Browse** and select a folder to change the default location of the **Backup Repository**.




2. Click **MORE OPTIONS** to set up more installation options:

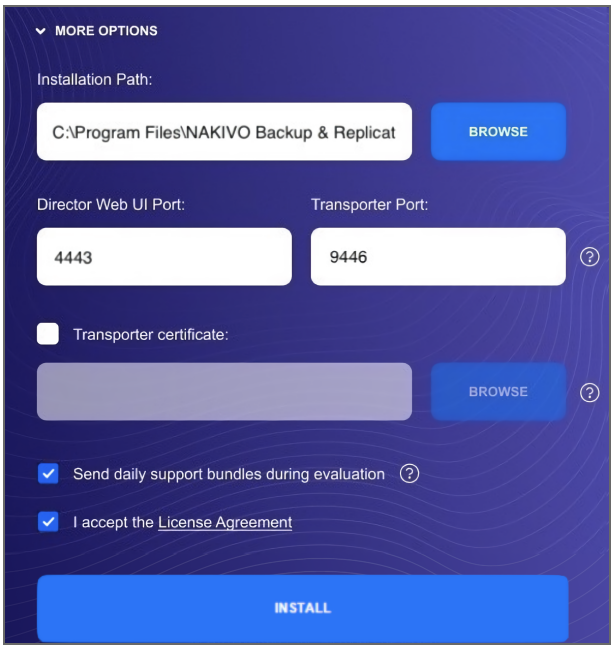
- **Installation path**: The location where NAKIVO Backup & Replication will be installed. If you want to change the default path to NAKIVO Backup & Replication, click **Browse** and select a new location.
- **Director Web UI port**: The default port that will be used to connect to the Web UI of NAKIVO Backup & Replication. Make sure that the port you specify is open in your firewall.
- **Transporter port**: The default port that will be used by the **Director** to communicate with the Onboard **Transporter**. Make sure that the port you specify is open in your firewall.
- **Transporter certificate**: This allows you to use a CA Certificate. Enter the path to the folder containing the certificate file in the field.

**Note**

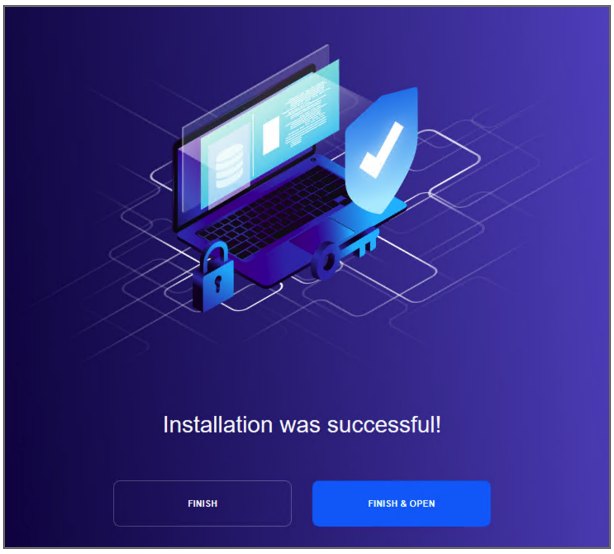
- When the checkbox is not selected, NAKIVO Backup & Replication automatically installs a self-signed certificate.
- It is possible to install a CA-signed certificate for the **Transporter** by conducting silent installation using the command-line arguments passed to the installer:
  - Use the following command for Windows OS:
    `installer.exe --cert C:\certificate.pem --eula-accept`
    The short option for the Windows OS command is the following:
    `installer.exe -ct C:\certificate.pem -ea`
  - Use the following command for Linux OS:
    `installer.sh --cert /tmp/certificate.pem --eula-accept`

- **Send daily support bundles during evaluation**: When this option is selected, NAKIVO Backup & Replication automatically creates, encrypts, and uploads support bundles once a day to a NAKIVO support server during the evaluation period. The NAKIVO support team may use this information to improve the product experience and may be able to identify and resolve product issues faster.

3. **I accept the License Agreement**: Select this option to confirm that you have read and agreed to the License Agreement.
4. Click **Install**.




5. Click **Finish** to complete the installation process or **Finish & Open** to complete installation and start using NAKIVO Backup & Replication.




6. To prevent unauthorized access to the product, create your user account. Fore more details, refer to “Logging in to NAKIVO Backup & Replication” on page 334 .

# Installing Transporter-Only on Windows

If you have already installed the full solution (both **Director** and **Transporter**) and wish to deploy an additional **Transporter**, run the NAKIVO Backup & Replication installer for Windows and follow the steps below:

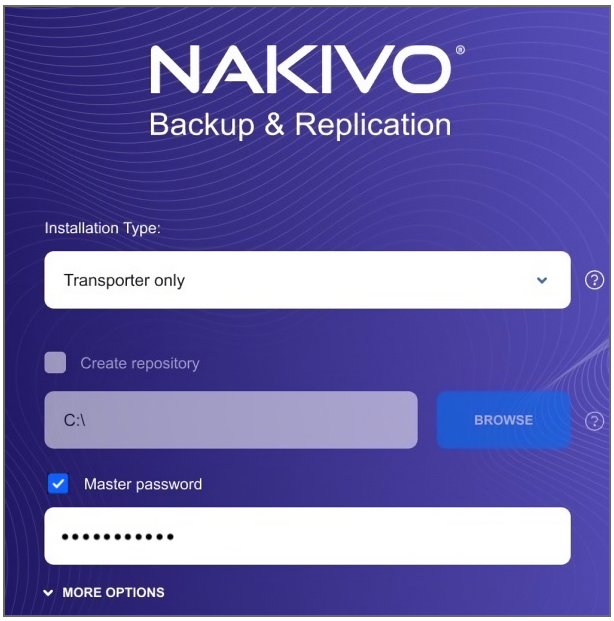
## Transporter Installation Prerequisites

Prior to installing the **Transporter**, make sure the following prerequisites are met:

- Make sure the machine on which you plan to install the **Transporter** has a connection to the relevant items below:
  - The machine on which the **Director** is installed.
  - VMware/Hyper-V/Nutanix AHV servers on which you plan to back up or replicate VMs (provided that you plan to retrieve VM data using the Transporter you are about to install)
  - Machines on which you have installed other **Transporters** (provided that you plan to set up data transfer between an existing Transporter and the one you are about to install)
  - **Backup Repositor**y (provided that you plan to assign the **Transporter** you are about to install to a **Backup Repository**)
  - VMware/Hyper-V/Nutanix AHV servers which you plan to use as a destination for replicated VMs (provided that you plan to write data to the target servers and datastores using the Transporter you are about to install)
- For VMware/Hyper-V/Nutanix AHV servers discovered with DNS names, make sure those DNS names can be resolved on the machine on which to install the Transporter.

## Transporter Installation Process

1. Run the NAKIVO Backup & Replication installer.
2. Choose **Transporter only** from the **Installation type** drop-down list.



3. Optionally, you can select the **Master password** checkbox and enter the password that will be used to generate a pre-shared key and secure the **Transporter**.

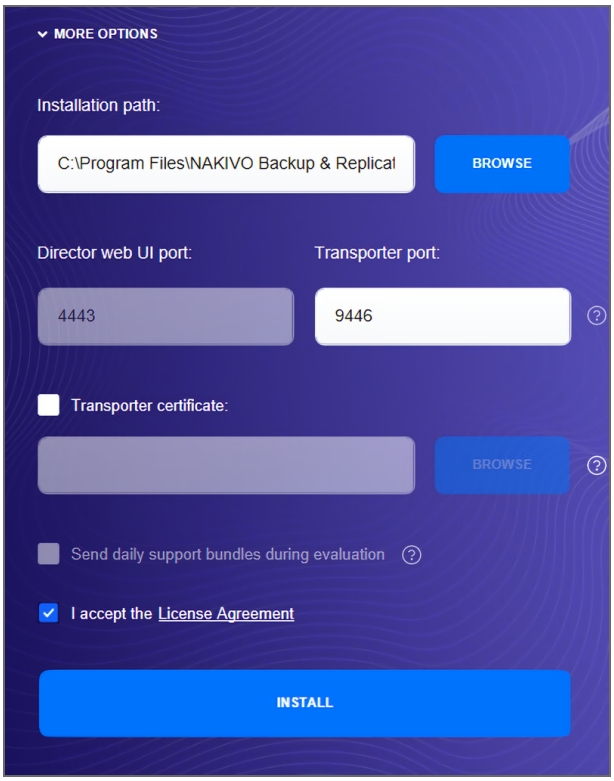
> **Note**
>
> - The master password must adhere to the following requirements:
>   - Minimal length - 5 characters.
>   - Maximum length - 50 characters.
> - The master password can be set and re-set manually by running the command on the machine housing the **Transporter**. Follow these steps:
>   - Enter the following command `bhsvc -b P@ssword123`
>   - Restart the **Transporter** service.

4. Click **MORE OPTIONS** and set up the following:
   - **Installation path**: The location where the **Transporter** will be installed. If you want to change the default path to the **Transporter** installation folder, click **Browse** and select a new location.
   - **Transporter port**: The default port that will be used by the **Director** to communicate with the Onboard **Transporter**. Make sure that the port you specify is open in your firewall.
   - **Transporter certificate**: This allows you to use a CA Certificate. Enter the path to the folder containing the certificate file in the field.

> **Note**
>
> - When the checkbox is not selected, NAKIVO Backup & Replication automatically installs a self-signed certificate.
> - It is possible to set up a master password and CA-signed certificate for the **Transporter** by conducting silent installation using the command-line arguments passed to the installer:
>   - Use the following command for Windows OS:
>     ```
>     installer.exe --cert C:\certificate.pem --master-pass P@ssword123 --eula-accept
>     ```
>     The short option for the Windows OS command is the following:
>     ```
>     installer.exe -ct C:\certificate.pem -b P@ssword123 -ea
>     ```
>   - Use the following command for Linux OS:
>     ```
>     installer.sh --cert /tmp/certificate.pem -b P@ssword123 --eula-accept
>     ```

- **Send daily support bundles during evaluation**: If this option is selected, NAKIVO Backup & Replication will automatically create, encrypt, and upload support bundles once a day to a NAKIVO support server during the evaluation period. NAKIVO Support team may use this information to improve the product experience and may be able to identify and resolve product issues faster.

5. **I accept the License Agreement**: Select this option to confirm that you have read and agreed to the License Agreement.
6. Click **Install**.




7. When the installation is complete the **Transporter installation was successful** notification appears.




8. Add the Transporter to NAKIVO Backup & Replication.

## Installing Full Solution in Multi-Tenant Mode on Windows

To install the full solution in multi-tenant mode on a Windows OS, run the NAKIVO Backup & Replication installer for Windows and follow the steps below:

1. Set the installation options as follows:
    - **Installation type**: Select the **Multi tenant solution** option from the **Installation type** drop-down list.
    - **Create repository**: Leave the checkbox selected to create a **Backup Repository** on the machine on which NAKIVO Backup & Replication is installed.

- Optionally, click **Browse** and select a folder to change the default location of the **Backup Repository**.




2. Click **MORE OPTIONS** to set up more installation options:
   - **Installation path**: The location where NAKIVO Backup & Replication will be installed. If you want to change the default path to the product, click **Browse** and select a new location.
   - **Director Web UI port**: The default port that will be used to connect to the Web UI of NAKIVO Backup & Replication. Make sure that the port you specify is open in your firewall.
   - **Transporter port**: The default port that will be used by the **Director** to communicate with the Onboard **Transporter**. Make sure that the port you specify is open in your firewall.

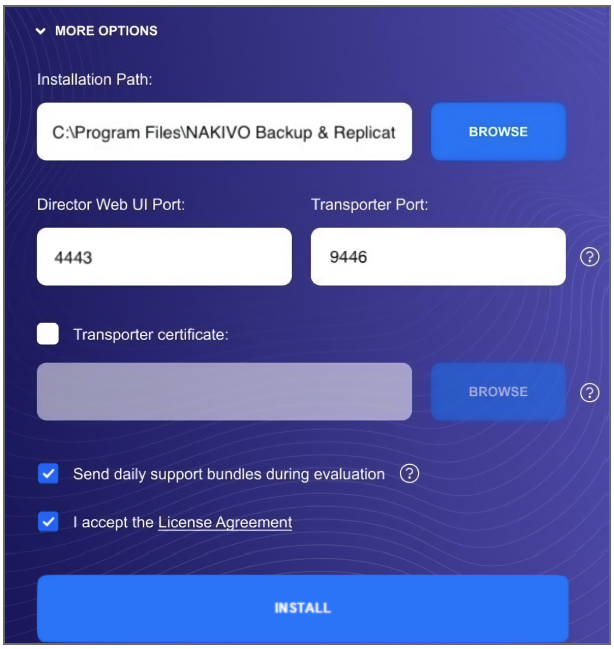
- **Transporter certificate**: This allows you to use a CA Certificate. Enter the path to the folder containing the certificate file in the field.

  **Note**

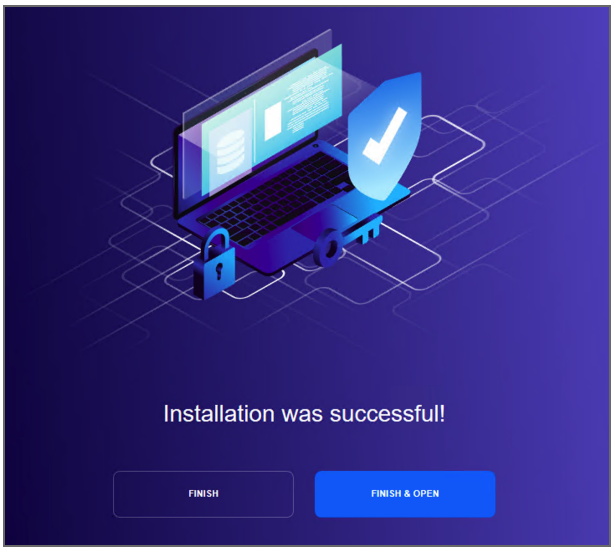
  - When the checkbox is not selected, NAKIVO Backup & Replication automatically installs a self-signed certificate.
  - It is possible to install a CA-signed certificate for the **Transporter** by conducting silent installation using the command-line arguments passed to the installer:
    - Use the following command for Windows OS:
      `installer.exe --cert C:\certificate.pem --eula-accept`
      The short option for the Windows OS command is the following:
      `installer.exe -ct C:\certificate.pem -ea`
    - Use the following command for Linux OS:
      `installer.sh --cert /tmp/certificate.pem --eula-accept`

- **Send daily support bundles during evaluation**: When this option is selected, NAKIVO Backup & Replication automatically creates, encrypts, and uploads support bundles once a day to a NAKIVO support server during the evaluation period. The NAKIVO support team may use this information to improve the product experience and may be able to identify and resolve product issues faster.

3. **I accept the License Agreement**: Select this option to confirm that you have read and agreed to the License Agreement.
4. Click **Install**.




5. Click **Finish** to complete the installation process or **Finish & Open** to complete installation and start using NAKIVO Backup & Replication.




> **Note**
> The onboard backup repository for the Master Tenant is automatically created after the installation.

6. Create an account by completing the form. For details, refer to “Logging in to NAKIVO Backup & Replication” on page 334.

Credentials are not required to log in as Master Admin after installation. However, the default credentials are required to log into the product after the first tenant is created. To log in as Master Admin, specify “admin” as the username and leave the password field empty. You can change credentials in the product configuration.

## Silent Installation on Windows

You can install NAKIVO Backup & Replication in silent mode via a command line by running the following command: **installer.exe -f --eula-accept**. This installs all product components (**Director**, **Transporter**, and **Backup Repository**), and you will be able to use all product features after installation.

The following arguments are available:

| Argument | Description |
|---|---|
| **-h** | Display the list of available arguments without starting the installation. |
| **--eula-accept, -ea** | Indicates that you have read and agree to the End User License Agreement. |
| **-f** | Performs the silent installation of the full solution (**Director** and **Transporter**). |
| **-t** | Performs the silent installation of **Transporter** only. |
| **-m** | Performs the silent installation of the full solution in multi-tenant mode. |
| **-u** | Performs the silent update of the installed product components. |
| **--release-notes, -n** | Indicates the user has read the release notes for the new release during an update. |
| **-sii** | Performs the silent install or update ignoring the single installer instance check. |
| **--ignore-pre-install-action-failures, -ipiaf** | All pre-install action failures are ignored. |
| **--cert** | Allows to set up a custom **Transporter** certificate. |
| **--master-pass** (short version: **-b**) | Allows to set up a custom master password for the Transporter. |

# Installing on Linux


- Linux Installation Prerequisites
- Silent Installation on Linux
- Installing Full Solution on Linux
- Installing Transporter on Linux
    - Transporter Installation Prerequisites
    - Transporter Installation
- Installing Full Solution in Multi-Tenant Mode on Linux


## Linux Installation Prerequisites

In order to install and use NAKIVO Backup & Replication on a Linux OS, make sure the following packages are installed:

- cifs-utils
- iscsi-initiator-utils
- ntfs-3g
- tar

## Silent Installation on Linux

You can install NAKIVO Backup & Replication in silent mode via a command line. To install the full solution, simply run the following command: `installer.sh -f --eula-accept` This will install all product components (Director, Transporter, Backup Repository) and you will be able to use all product features after installation.

The following arguments are available:

| Argument | Description |
|---|---|
| **-h, -help, help** | Display the list of available arguments without starting the installation. |
| **--eula-accept, -ea** | Indicates that you have read and agree to the End User License Agreement. |
| **-f** | Shall perform the silent installation of the full solution (**Director** and **Transporter**). |
| **-t** | Shall perform the silent installation of **Transporter** only. |
| **-m** | Shall perform the silent installation of the full solution in multi-tenant mode. |
| **-u** | Shall perform the silent update of the installed product components. |

| Argument | Description |
|---|---|
| **-e** | Shall install **Transporter** on Amazon EC2, or update **Transporter** installed on Amazon EC2. Refer to Updating on Amazon EC2 for details. |
| **-a** | Shall enable uploading support bundles to support team server (Call Home). Refer to System Settings for details. |
| **-y** | Shall accept limitations silently. |
| **-i <install_path>** | Shall install to the specified installation path. |
| **-d <director_port>** | Shall provide a custom **Director** port. |
| **-p <transporter_port>** | Shall provide a custom **Transporter** port. |
| **-r <port1>-<port2>** | Shall provide a custom transporter data ports range. |
| **-C** | Shall suppress creating the repository. |
| **-c <repo_path>** | Shall create the repository. The **<repo_path>** parameter is optional. |
| **--rt <repo_type>** | Shall create a repository of the specified type. The **<repo_type>** parameter may accept the following values: **1** – "Forever incremental with deduplication"; **2** – "Forever incremental without deduplication"; **3** – "Incremental with full backups (deduplication devices)". |
| **--rc <compress_level>** | Shall specify the repository compression level. The parameter may accept the following values: Disabled; Fast; Medium; Best. Refer to Creating Backup Repositories for details. |
| **--pnp-cleanup** | Shall clean up the database of the device manager for the Linux kernel. |
| **--cert** | Allows to set up a custom **Transporter** certificate. |
| --**master-pass** (short version: **-b**) | Allows to set up a custom master password for the **Transporter**. |

## Installing Full Solution on Linux

Follow the steps below to install all components of NAKIVO Backup & Replication (both **Director** and **Transporter**) on a Linux OS:

1. Upload the installer file to the machine on which you want to install NAKIVO Backup & Replication in the *binary transfer mode*. For example:

- Upload the installer from a Windows-based machine.
- Upload the product from a Linux-based machine: run the following command: `wget 'server_ip/shared/NAKIVO_Backup_&_Replication_TRIAL.sh'`

2. Log in to the Linux machine and allow the execution of the installer file.
   For example: `chmod +x NAKIVO_Backup_&_Replication_TRIAL.sh`
3. Execute the installer file with root privileges.
   For example: `sudo ./NAKIVO_Backup_&_Replication_TRIAL.sh`
4. Review the license agreement (press **Space** to go to the next page of the agreement). If you agree to the terms of the license agreement, press "Y" and then press **Enter**.

   **Note**
   On some Linux consoles you need to press "Q".

5. Type "S" to install the full solution and press **Enter**.
6. Optionally, you can install CA **Transporter** certificate. Enter the path to the folder containing the certificate file and press **Enter**.

   **Notes**
   - If no path to the CA certificate was provided, NAKIVO Backup & Replication automatically installs a self-signed certificate.
   - It is possible to set up CA-signed certificate for the **Transporter** by conducting silent installation using the command-line arguments passed to the installer. Use the following command:
     `installer.sh --cert /tmp/certificate.pem --eula-accept`

7. Specify the installation path for the product: Press **Enter** to accept the default installation path "/opt/nakivo" or enter a custom path and press **Enter**.
8. Specify the **Director** HTTPS port (which will be used to access the Web UI of NAKIVO Backup & Replication): Press **Enter** to accept the default port *4443* or enter a custom port number and press **Enter**. Make sure the port you specify is open in your firewall.
9. Specify whether to allow the product to automatically send support bundles to a NAKIVO server during the evaluation period (Call Home). If this option is enabled, NAKIVO Backup & Replication will automatically create, encrypt, and upload support bundles once a day to a NAKIVO support server during the evaluation period. NAKIVO Support team may use this information to improve the product experience and will be able to identify and resolve product issues faster.
10. Specify the Transporter port (which will be used to connect to the **Transporter** that is installed by default with the **Director**): Press **Enter** to accept the default port "9446" or enter a custom port number (1 to 65535) and press **Enter**. Make sure the port you specify is open in your firewall.

11. Specify a range of port numbers (from *1* to *65535*) that will be used to transfer data by the Onboard **Transporter** (default are *9448-10000*). The range you specify should contain at least 100 ports. Make sure that the ports you specify are open in your firewall.
12. Specify a path to the default Backup Repository: Press **Enter** to accept the default path "/opt/nakivo/repository" or enter a custom path and press **Enter** to begin the installation process.

After the installation is complete, you can log in to NAKIVO Backup & Replication by opening the product's web interface in your browser (use the machine's IP address or DNS name with the appropriate HTTPS port). By default, login name and password are not required to access NAKIVO Backup & Replication. To prevent unauthorized access to the product, you can set up credentials in **Configuration**.

# Installing Transporter on Linux

If you have already installed the full solution (both **Director** and **Transporter**) and want to deploy an additional **Transporter**, run the NAKIVO Backup & Replication installer for Windows and follow the steps below:

## Transporter Installation Prerequisites

Prior to installing a **Transporter**, make sure the following prerequisites are met:

1. Make sure the machine on which you plan to install the **Transporter** has a connection to the relevant items below:
   - The machine on which the **Director** is installed
   - VMware/Hyper-V servers on which you plan to back up or replicate VMs (if you plan to retrieve VM data using the **Transporter** you are about to install)
   - Machines on which you have installed other **Transporters** (if you plan to set up data transfer between an existing **Transporter** and the one you are about to install)
   - Backup repository (if you plan to assign the **Transporter** you are about to install to a **Backup Repository**)
   - VMware/Hyper-V servers which you plan to use as a destination for replicated VMs (if you plan to write data to the target servers and datastores using the **Transporter** you are about to install)
2. If you have discovered VMware/Hyper-V servers using DNS names, make sure those DNS names can be resolved on the machine on which you plan to install the **Transporter**.

## Transporter Installation

1. Upload the installer file to the machine on which you wish to install NAKIVO Backup & Replication in the *binary transfer mode*. For example:
   - Upload the installer from a Windows-based machine.
   - Upload the product from a Linux-based machine: run the following command: `wget 'server_ip/shared/NAKIVO_Backup_&_Replication_TRIAL.sh'`

2. Allow the execution of the installer file. For example: `chmod +x NAKIVO_Backup_&_Replication_TRIAL.sh`
3. Execute the installer file with root privileges. For example:`sudo ./NAKIVO_Backup_&_Replication_TRIAL.sh`
4. Review the license agreement (press **Space** to go to the next page of the agreement). If you agree to the terms of the license agreement, press "Y" and then press **Enter**.
5. Type "T" to install only the **Transporter** and press **Enter**.

> **Note**
> Alternatively, you can use the **-t** argument to install the **Transporter** silently:
> `sudo ./NAKIVO_Backup_&_Replication_TRIAL.sh -t`

6. Optionally, enter the master password that will be used to generate a pre-shared key and secure the **Transporter** and then press **Enter**.

> **Notes**
> - The master password must adhere to the following requirements:
>   - Minimal length - 5 characters.
>   - Maximum length - 50 characters.
> - The master password can be set and re-set manually by following these steps:
>   1. Switch to root using the following command:
>      `sudo -i`
>   2. Stop the **Transporter** service.
>   3. Go to the **Transporter** folder with the following command:
>      `cd /opt/nakivo/transporter`
>   4. Run the following command to set the master password:
>      `./bhsvc -b P@ssword123`
>   5. Restart the Transporter service.

7. Specify the installation path for the product: Press **Enter** to accept the default installation path "/opt/nakivo" or enter a custom path and press **Enter**.
8. Optionally, you can install CA Transporter certificate. Enter the path to the folder containing the certificate file and press **Enter**.

**Notes**

- If no path to the CA certificate was provided, NAKIVO Backup & Replication automatically installs a self-signed certificate.
- It is possible to set up a master password and a CA-signed certificate for the **Transporter** by conducting silent installation using the command-line arguments passed to the installer. Use the following command:
  `installer.sh --cert /tmp/certificate.pem -b P@ssword123 --eula-accept`

9. Specify the **Transporter** port (used to connect to the **Transporter**): Press **Enter** to accept the default port *9446* or enter a custom port number and press **Enter** to begin the installation process. Make sure the port you specify is open in your firewall.

After the installation is complete, add the **Transporter** to NAKIVO Backup & Replication.

## Installing Full Solution in Multi-Tenant Mode on Linux

Follow the steps below to install the full solution in multi-tenant mode on a Linux OS:

1. Upload the installer file to the machine on which you want to install NAKIVO Backup & Replication in the *binary transfer mode*. For example:
   - Upload the installer from a Windows-based machine.
   - Upload the product from a Linux-based machine: run the following command: `wget 'server_ip/shared/NAKIVO_Backup_&_Replication_TRIAL.sh'`
2. Log in to the Linux machine and allow the execution of the installer file.
   For example: `chmod +x NAKIVO_Backup_&_Replication_TRIAL.sh`
3. Execute the installer file with root privileges.
   For example: `sudo ./NAKIVO_Backup_&_Replication_TRIAL.sh`
4. Review the license agreement (press **Space** to go to the next page of the agreement). If you agree to the terms of the license agreement, press "Y" and then press **Enter**.
5. Type “M” to install the **Director** in Multi-tenant mode and press **Enter**.

**Note**

Alternatively, you can use the **-m** argument to install the solution in multi-tenant mode silently:
`sudo ./NAKIVO_Backup_&_Replication_TRIAL.sh -m`

6. Optionally, you can install CA **Transporter** certificate. Enter the path to the folder containing the certificate file and press **Enter**.

**Notes**

- If no path to the CA certificate was provided, NAKIVO Backup & Replication automatically installs a self-signed certificate.
- It is possible to set up CA-signed certificate for the **Transporter** by conducting silent installation using the command-line arguments passed to the installer. Use the following command:
  `installer.sh --cert /tmp/certificate.pem --eula-accept`

7. Specify the installation path for the product: Press **Enter** to accept the default installation path "/opt/nakivo" or enter a custom path and press **Enter**.
8. Specify the Director HTTPS port (which will be used to access the Web UI of NAKIVO Backup & Replication): Press **Enter** to accept the default port *4443* or enter a custom port number and press **Enter**. Make sure the port you specify is open in your firewall.
9. Specify whether to allow the product to automatically send support bundles to a NAKIVO server during the evaluation period. If this option is enabled, NAKIVO Backup & Replication will automatically create, encrypt, and upload support bundles once a day to a NAKIVO support server during the evaluation period. NAKIVO Support team may use this information to improve the product experience and will be able to identify and resolve product issues faster.
10. Specify the **Transporter** port (which will be used to connect to the **Transporter** that is installed by default with the **Director**): Press **Enter** to accept the default port "9446" or enter a custom port number (*1* to *65535*) and press **Enter**. Make sure the port you specify is open in your firewall.
11. Specify a range of port numbers (from 1 to 65535) that will be used to transfer data by the Onboard **Transporter** (default are *9448-10000*). The range you specify should contain at least 100 ports. Make sure that the ports you specify are open in your firewall.
12. The onboard backup repository for the Master Tenant is automatically created after the installation.
13. Specify a path to the default backup repository: Press **Enter** to accept the default path `/opt/nakivo/repository` or enter a custom path and press **Enter** to begin the installation process.

**Note**

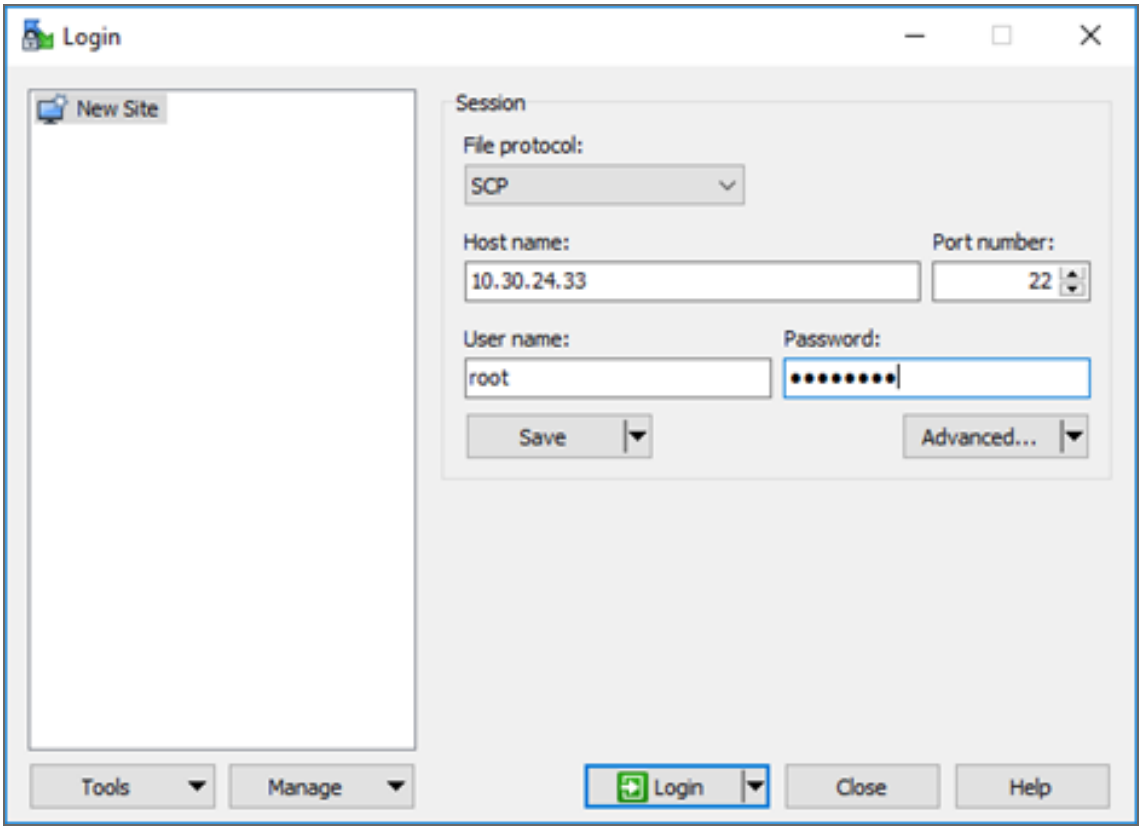
The onboard backup repository for the Master Tenant is automatically created after the installation.

After the installation is complete, you can log in to NAKIVO Backup & Replication by opening its web interface in your browser. Use the machine's IP address or DNS name together with the Director HTTPS port. Refer to "Getting Started" on page 333 to know how to continue working with NAKIVO Backup & Replication.

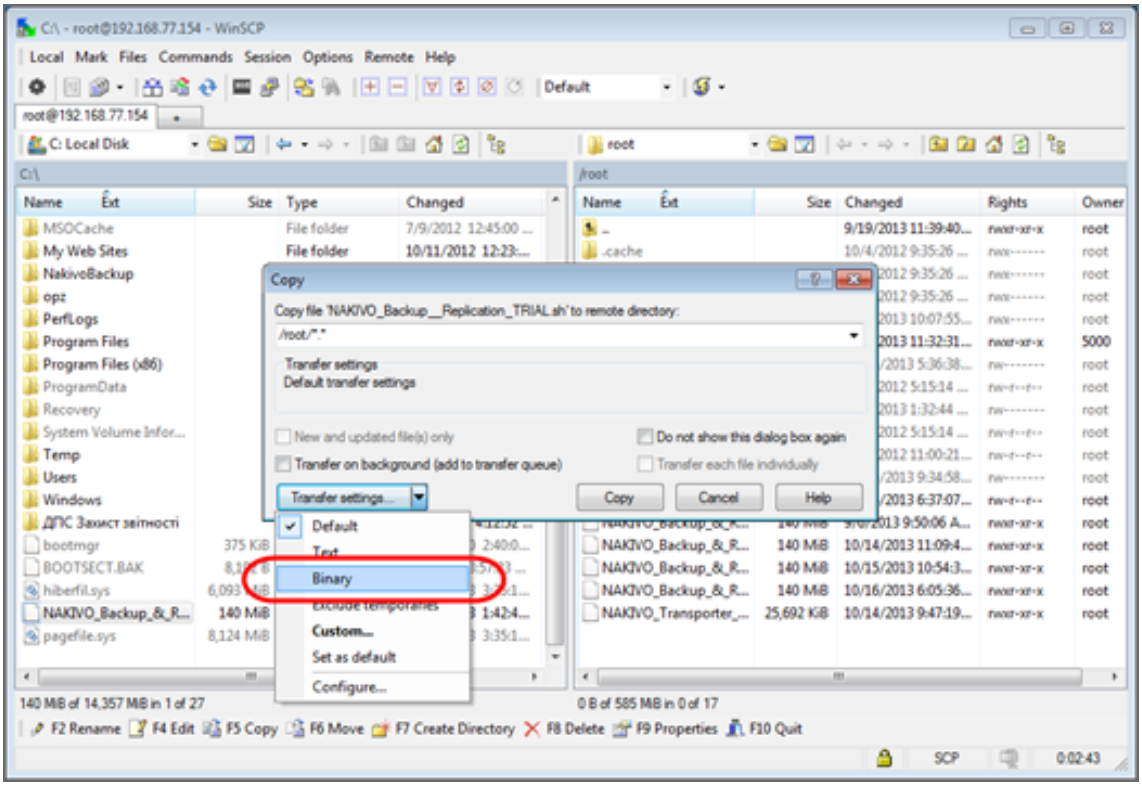
# Uploading Installer from Windows Machine to Linux Machine

To upload the installer from a Windows-based machine, follow the steps below:

1. Download the free WinSCP client from http://winscp.net, install, and run it.
2. Choose **SCP** from the **File protocol** list.
3. Specify the IP address or the hostname of the Linux machine on which you would like to install the product in the **Host name** field.
4. Specify the username and password to the Linux machine in the appropriate boxes.
5. Leave other options as is and click **Login**.




   Click **Yes** in the dialog box that opens.
6. In the left pane, find the folder that contains the Linux installer, in the right pane, go up to the root folder.
7. Drag and drop the installer from left to the right pane.
8. Choose **Binary** from the **Transfer settings** drop-down list in the **Copy** dialog box that opens.


C:\ - root@192.168.77.154 - WinSCP
Local Mark Files Commands Session Options Remote Help
Default
root@192.168.77.154
C: Local Disk
root
C:\
/root
Name Ext Size Type Changed
Name Ext Size Changed Rights Owner
MSOCache
My Web Sites
NakivoBackup
opt
PerfLogs
Program Files
Program Files (x86)
ProgramData
Recovery
System Volume Infor...
Temp
Users
Windows
ДПС Захист звітності
bootmgr
BOOTSECT.BAK
hiberfil.sys
NAKIVO_Backup_&_R...
pagefile.sys
Copy
Copy file 'NAKIVO_Backup__Replication_TRIAL.sh' to remote directory:
/root/*.*
Transfer settings
Default transfer settings
New and updated file(s) only
Do not show this dialog box again
Transfer on background (add to transfer queue)
Transfer each file individually
Transfer settings...
Copy
Cancel
Help
Default
Text
Binary
Custom...
Set as default
Configure...
NAKIVO_Backup_&_R... 140 MiB 10/14/2013 11:09:4...
NAKIVO_Backup_&_R... 140 MiB 10/15/2013 10:54:3...
NAKIVO_Backup_&_R... 140 MiB 10/16/2013 6:05:36...
NAKIVO_Transporter_... 25,692 KiB 10/14/2013 9:47:19...
140 MiB of 14,357 MiB in 1 of 27
0 B of 585 MiB in 0 of 17
F2 Rename F4 Edit F5 Copy F6 Move F7 Create Directory F8 Delete F9 Properties F10 Quit
SCP
0:02:43

9. Click **Copy**.

# Installing on Synology NAS

NAKIVO Backup & Replication can be installed directly on a supported Synology NAS to create your own, high-performance backup appliance. With the appliance, all VM data protection components are unified in a single system that is fast to deploy and easy to manage, while also not consuming your environment's valuable resources. Moreover, you are getting an all-in-one backup hardware, backup software, backup storage, and data deduplication in a single box. This results in a zero VMware footprint, less power and cooling, less required maintenance, time, money, and – most of all – higher VM backup performance.

You can install a Synology package with either all NAKIVO Backup & Replication components (**Director**, **Transporter**, **Backup Repository**) or a **Transporter** only. The product can be installed via Package Center or manually. For more details, refer to the corresponding topics below:

- "Installing on Synology NAS via Package Center" on page 271
- "Installing on Synology NAS Manually" on page 273

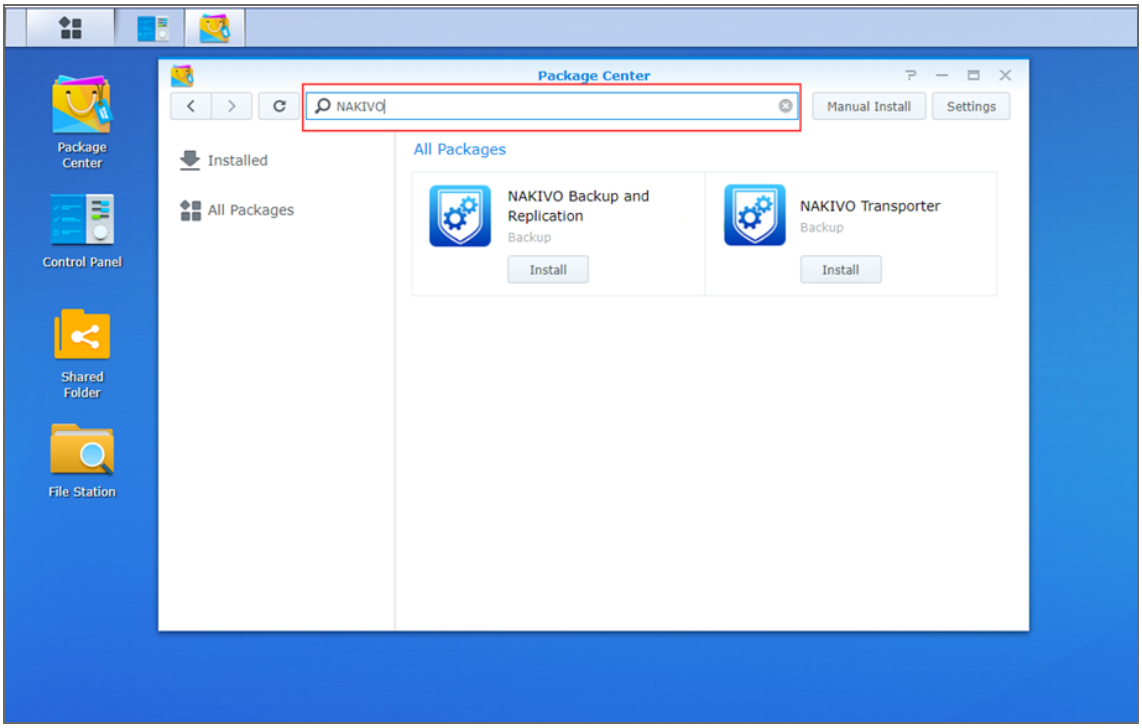
**Note**

A pre-shared key is not created during **Transporter**-only installation. When adding this Transporter to NAKIVO Backup & Replication, filling out the master password field is not required. The master password can be manually set and reset later. For details, refer to "Installed Service" on page 565.

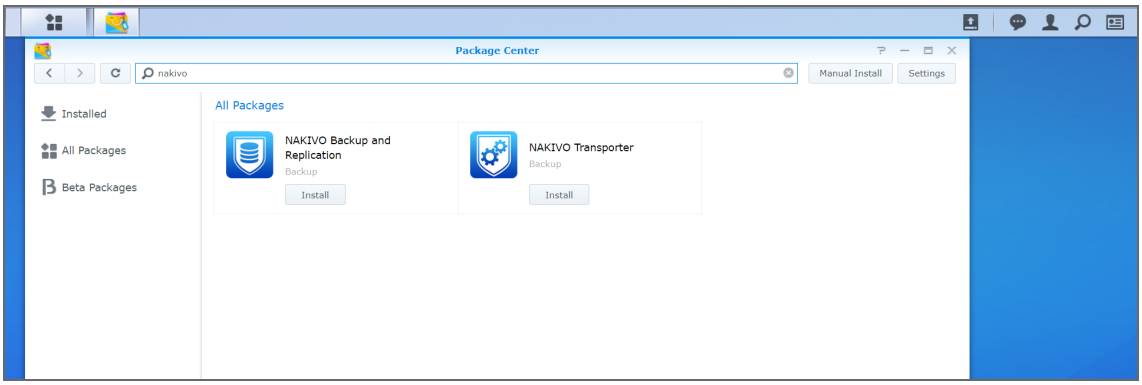
# Installing on Synology NAS via Package Center

To automatically install a NAKIVO Backup & Replication application on a Synology NAS, do the following:

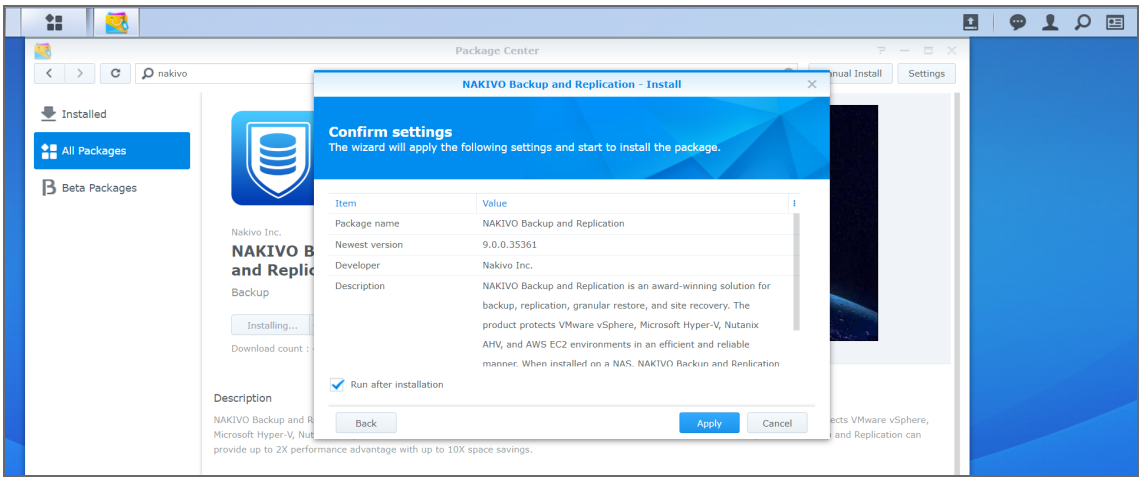
1. Log in to your Synology account and open **Package Center** in the management interface.
2. Use the search box to find NAKIVO Backup & Replication packages.




3. Click **Install** on one of the following:
   - **NAKIVO Backup and Replication** to install all product components.
   - **NAKIVO Transporter** to install a **Transporter** only.




4. Select the **I accept the terms of the license agreement** checkbox and click **Next**.
5. In the **Confirm settings** dialog box, click **Apply**.




**Note**

If you installed NAKIVO Backup & Replication on an ARM-based NAS, an additional **Transporter** needs to be added to your application outside of the ARM-based NAS to allow it to work with VMware vCenters and ESXi hosts. Please refer to the “Adding VMware vCenters and ESXi hosts” on page 511 topic for details.

Refer to “Getting Started” on page 333 to better understand how to continue working with NAKIVO Backup & Replication.

# Installing on Synology NAS Manually

If for any reason installation of NAKIVO Backup & Replication via Package Center is not available for your Synology NAS, you can install it manually.

The following packages are available for manual installation:

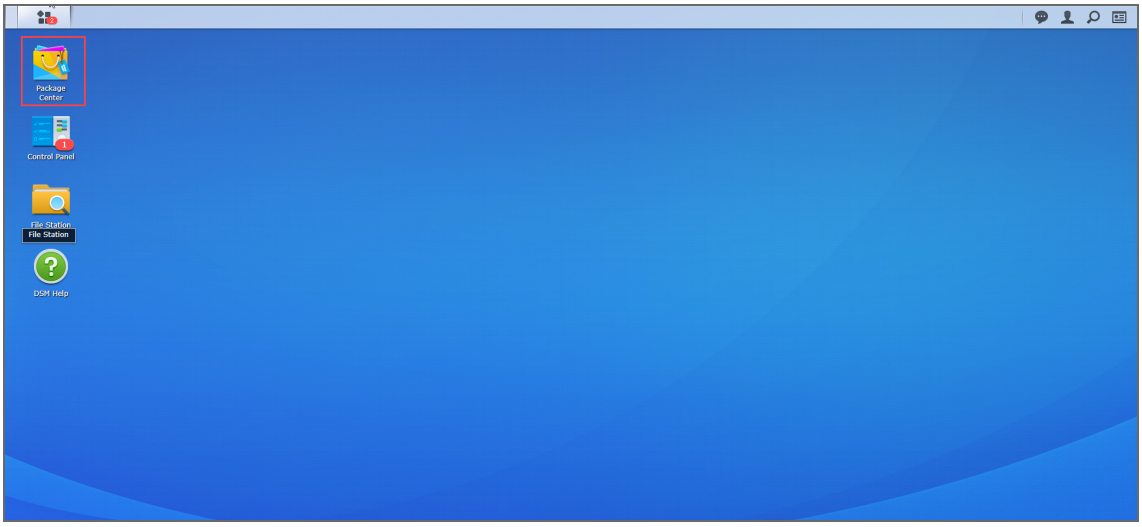
- Synology package
- Synology **Transporter** package
- Synology ARM package
- Synology ARM **Transporter** package

To manually install NAKIVO Backup & Replication on a Synology NAS, do the following:

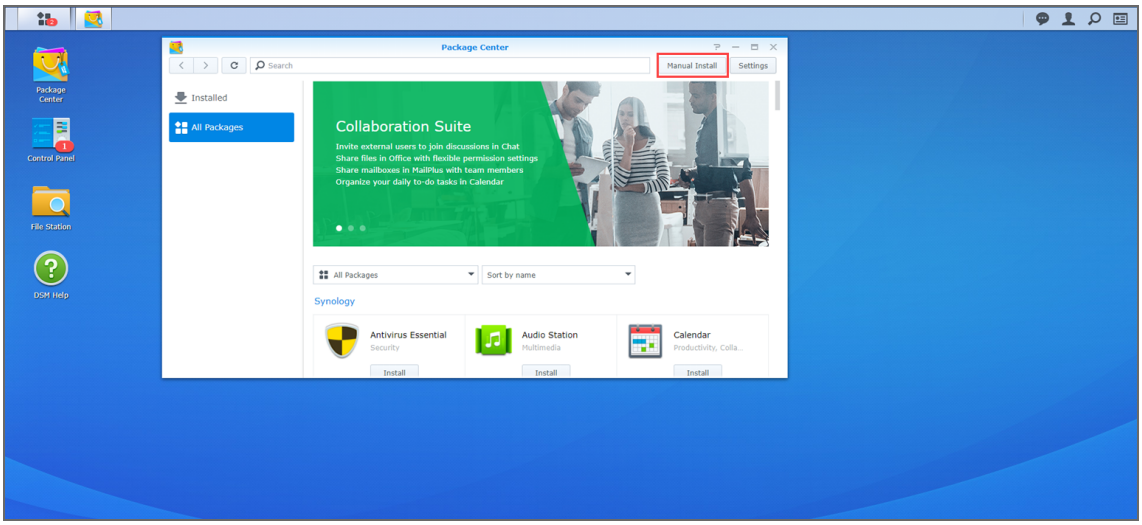
1. Download a Synology NAS package.

> **Note**
>
> Installing the NAKIVO Backup & Replication instance on your NAS-based backup appliance requires accurate knowledge of the hardware system information.
>
> Refer to the following page to learn how to identify your Synology NAS device CPU model: How to Identify the CPU Architecture of a Synology NAS Processor.

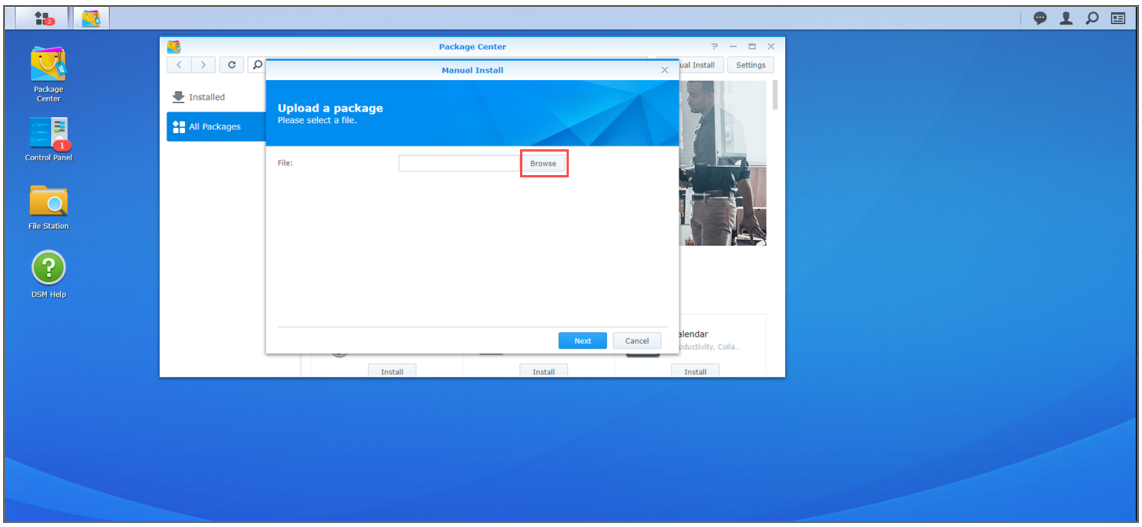
2. Log in to your Synology account and open the **Package Center** in the management interface.



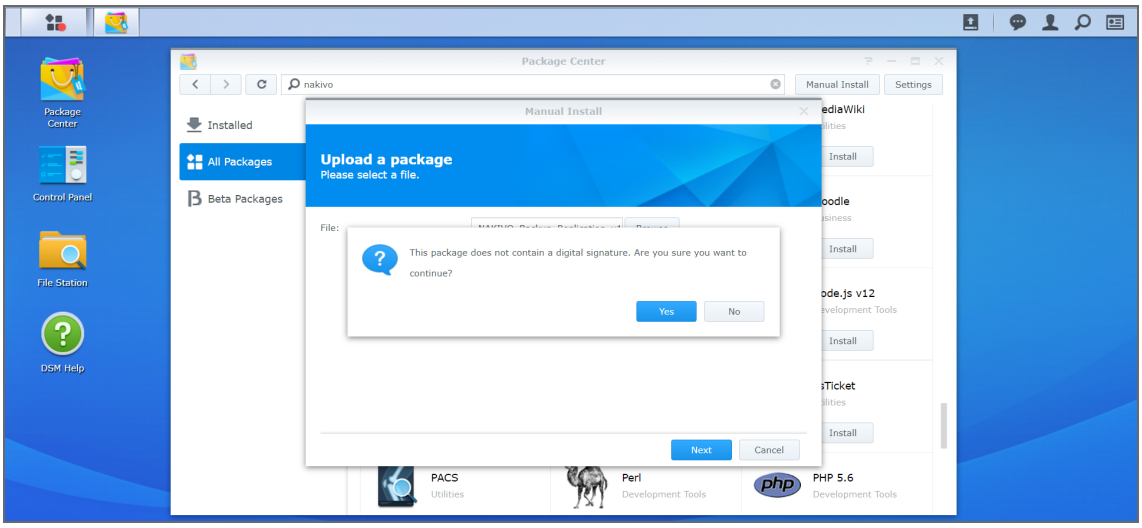

3. Click **Manual Install.**



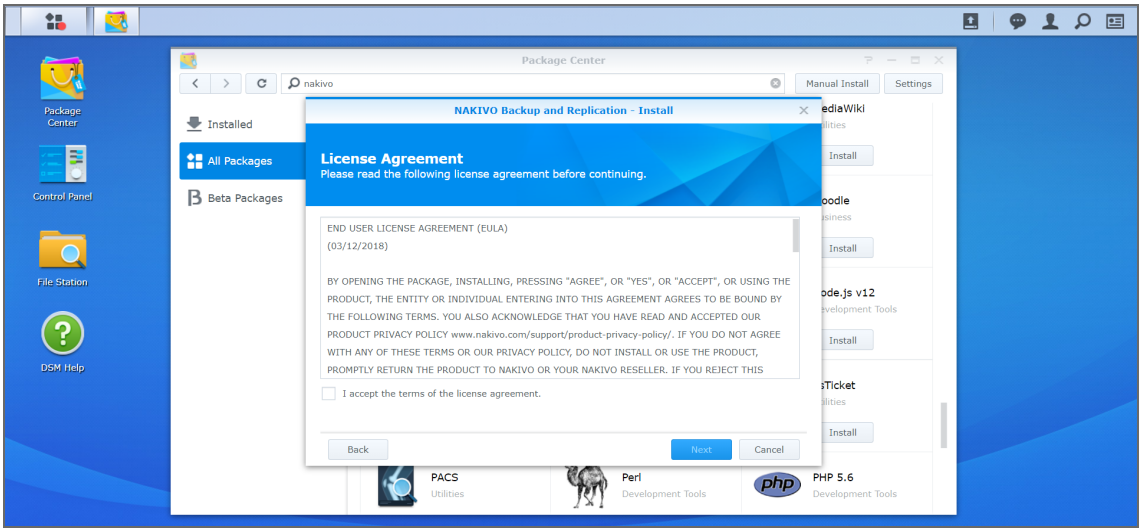

4. Click **Browse**, navigate to the Synology NAS package that you have downloaded, select it, and click **Open**.



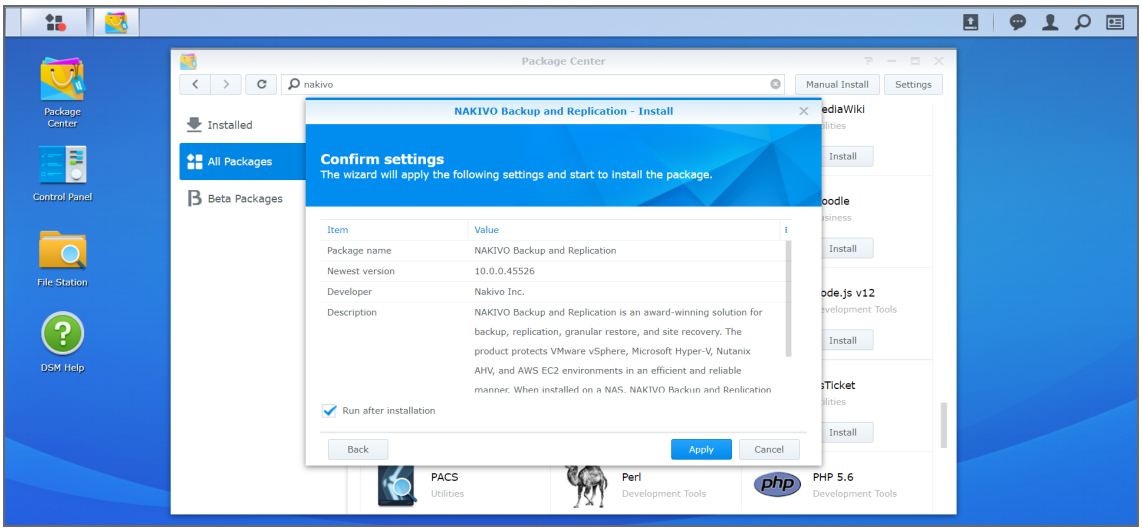

5. Click **Yes** to proceed.



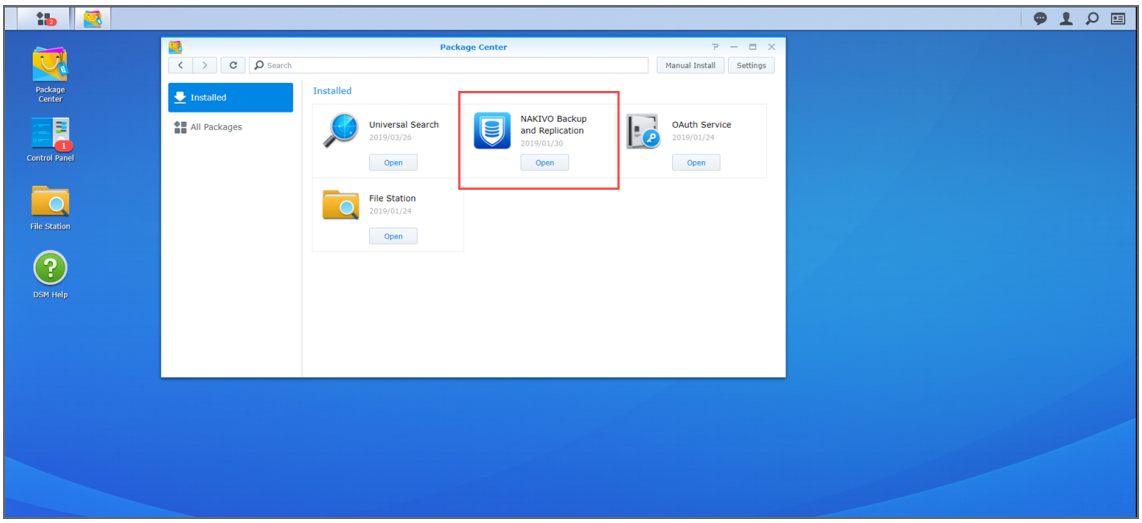

6. After reading through the License Agreement, check **I accept the terms of the license agreement** and click **Next**.




7. Optionally check **Run after installation** to start NAKIVO Backup & Replication immediately after the install process is finished. Click **Apply.**




8. Now NAKIVO Backup & Replication is installed on your NAS. To open the NAKIVO Backup & Replication Web interface, go to the following address in your web browser: `https://NAS_IP_address:4443`, or click the NAKIVO Backup & Replication icon in the main menu of the NAS.




> **Note**
>
> If you installed NAKIVO Backup & Replication on an ARM-based NAS, an additional **Transporter** needs to be added to your application outside of the ARM-based NAS to allow it to work with VMware vCenters and ESXi hosts. Please refer to the "Adding VMware vCenters and ESXi hosts" on page 511 topic for details.

Refer to "Getting Started" on page 333 to better understand how to continue working with NAKIVO Backup & Replication.

# Installing on QNAP NAS

You can install a QNAP package with either all NAKIVO Backup & Replication components (**Director**, **Transporter**, **Backup Repository**) or a **Transporter** only.

NAKIVO Backup & Replication can be installed directly on a supported QNAP NAS to create your own, high-performance backup appliance. With this appliance, all VM data protection components are unified in a single system that is quick to deploy and easy to manage, while also not consuming your environment's valuable resources. Moreover, you are getting an all-in-one backup hardware, backup software, backup storage, and data deduplication in a single box. This results in a zero VMware footprint, less power and cooling, less required maintenance, time, money, and – most of all – higher VM backup performance.

You can install NAKIVO Backup & Replication either via QNAP store or manually.


- "Installing on QNAP NAS via QNAP Store" on page 278
- "Installing on QNAP NAS Manually" on page 280


**Note**

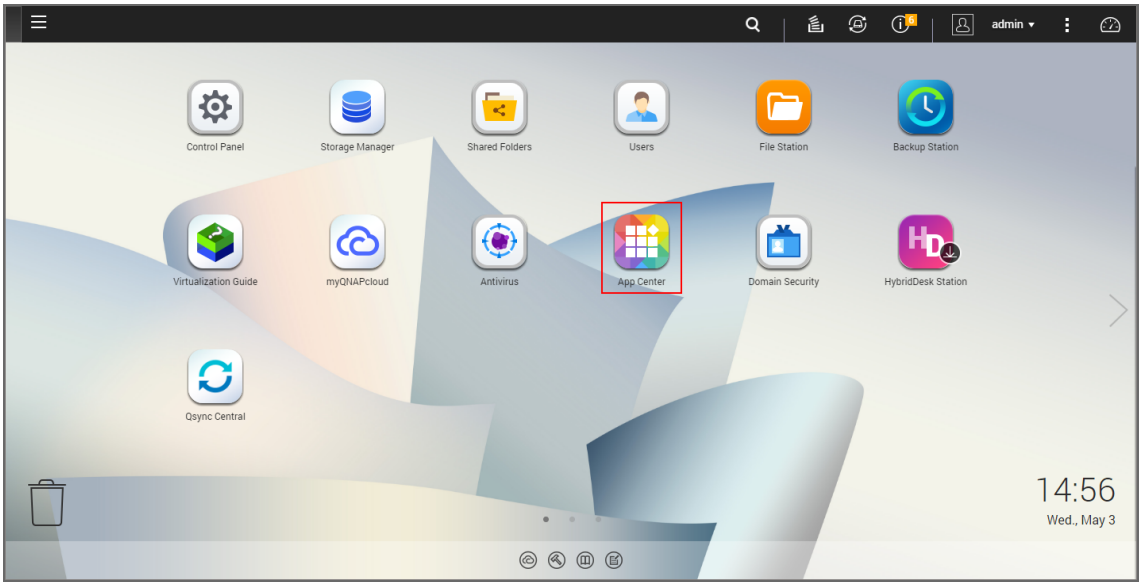
A pre-shared key is not created during **Transporter**-only installation. When adding this **Transporter** to NAKIVO Backup & Replication, filling out the master password field is not required. The master password can be manually set and reset later. For details, refer to "Installed Service" on page 565.

# Installing on QNAP NAS via QNAP Store

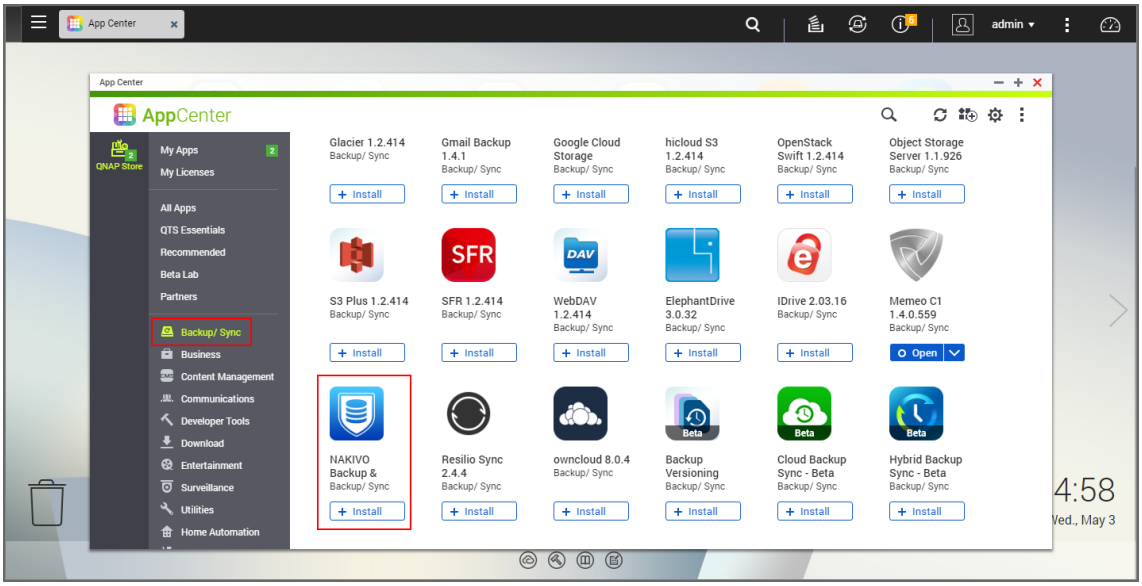
Check to see if your NAS model is supported before you begin installing NAKIVO Backup & Replication on a QNAP NAS.

To install NAKIVO Backup & Replication take the following steps:

1. Open the QNAP Desktop in your browser by entering the IP address of your QNAP NAS.




2. Go to **App Center**.
3. Select the **Backup/Sync** category and locate NAKIVO Backup & Replication. Alternatively, you can use the search bar at the top of the App Center window. Click on the magnifying glass icon and enter 'Nakivo'.




4. Click **Install**.
5. Wait till the installation is completed.

By default, NAKIVO Backup & Replication interface is available by the IP address of your QNAP NAS on the port *4443*: `https://<IP_address_of_QNAP_NAS>:4443`.

**Note**

If you installed NAKIVO Backup & Replication on an ARM-based NAS, an additional **Transporter** needs to be added to your application outside of the ARM-based NAS to allow it to work with VMware vCenters and ESXi hosts. Please refer to the "Adding VMware vCenters and ESXi hosts" on page 511 topic for details.

Refer to "Getting Started" on page 333 to know how to continue working with NAKIVO Backup & Replication.

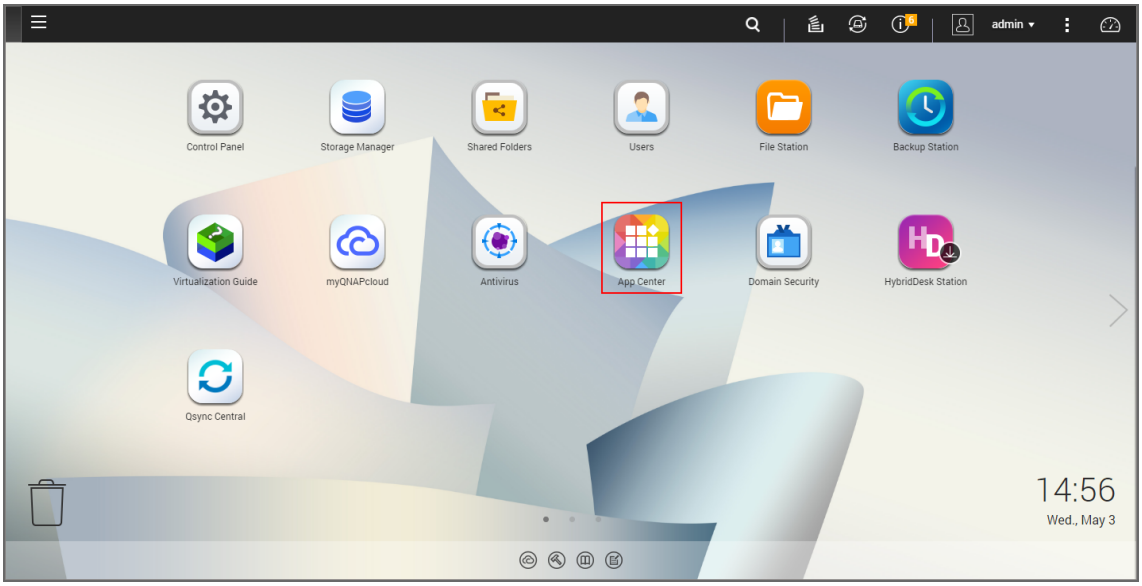
# Installing on QNAP NAS Manually

Before you begin installing NAKIVO Backup & Replication on a NAS, make sure your NAS model is supported and you have downloaded the installer (`.qpkg` file) for QNAP NAS.

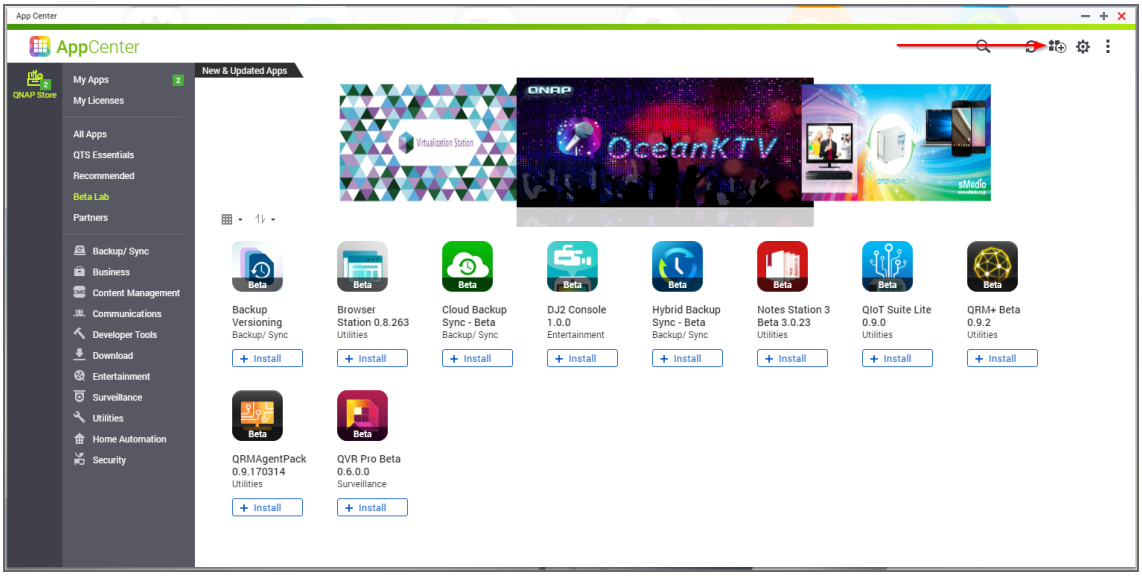
> **Note**
>
> Installing updates of NAKIVO Backup & Replication on your NAS-based backup appliance requires accurate knowledge of the hardware system information.
>
> Refer to the following page to learn how to identify your QNAP device CPU model: How to Identify the CPU Architecture of a QNAP NAS Processor.

To install NAKIVO Backup & Replication on a NAS:

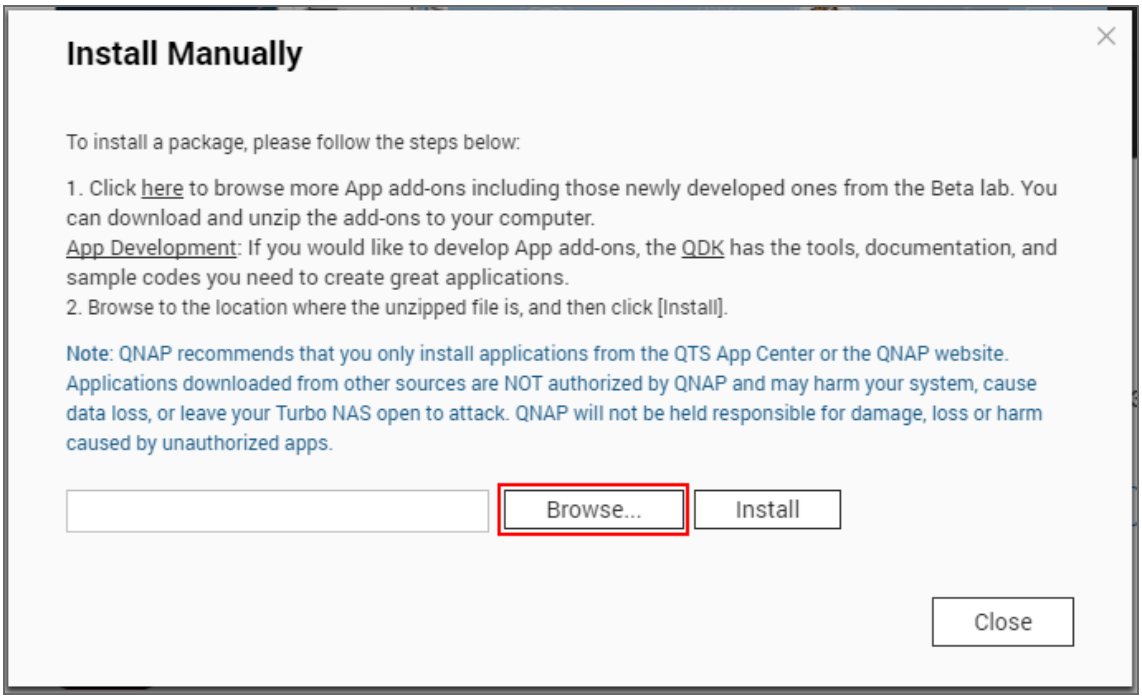
1. Open the QNAP Desktop in your browser by entering the IP address of your QNAP NAS.




2. Go to **App Center**.
3. Click the **Install Manually** icon.




4. Click **Browse** in the window that appears and locate the installer (`.qpkg` file) on your computer.




5. Click **Install**.
6. Wait until the installation is complete.

By default, NAKIVO Backup & Replication interface is available at the IP address of your QNAP NAS on the port *4443*: `https://<IP_address_of_QNAP_NAS>:4443`.

**Note**

If you installed NAKIVO Backup & Replication on an ARM-based NAS, an additional **Transporter** needs to be added to your application outside of the ARM-based NAS to allow it to work with VMware vCenters and ESXi hosts. Please refer to the Adding VMware vCenters and ESXi Hosts topic for details.

Refer to "Getting Started" on page 333 to better understand how to continue working with NAKIVO Backup & Replication.

# Installing on Western Digital NAS

You can install a Western Digital MyCloud package with either all NAKIVO Backup & Replication components (Director, Transporter, Backup Repository) or a **Transporter** only. The following packages are available:

- Western Digital MyCloud DL2100 package
- Western Digital MyCloud DL2100 **Transporter** package
- Western Digital MyCloud DL4100 package
- Western Digital MyCloud DL4100 **Transporter** package
- Western Digital MyCloud PR2100 package
- Western Digital MyCloud PR 2100 **Transporter** package
- Western Digital MyCloud PR 4100 package
- Western Digital MyCloud PR 4100 **Transporter** package

NAKIVO Backup & Replication can be installed directly on a Western Digital MyCloud NAS to create your own, high-performance backup appliance. With this appliance, all VM data protection components are unified in a single system that is quick to deploy and easy to manage, while also not consuming your environment's valuable resources. Moreover, you are getting an all-in-one backup hardware, backup software, backup storage, and data deduplication in a single box. This results in a zero VMware footprint, less power and cooling, less required maintenance, time, money, and – most of all – higher VM backup performance. NAKIVO Backup & Replication is installed on a NAS hard drive (not on the NAS Flash memory), so if you remove the hard drive from the NAS you will also remove the product from it.

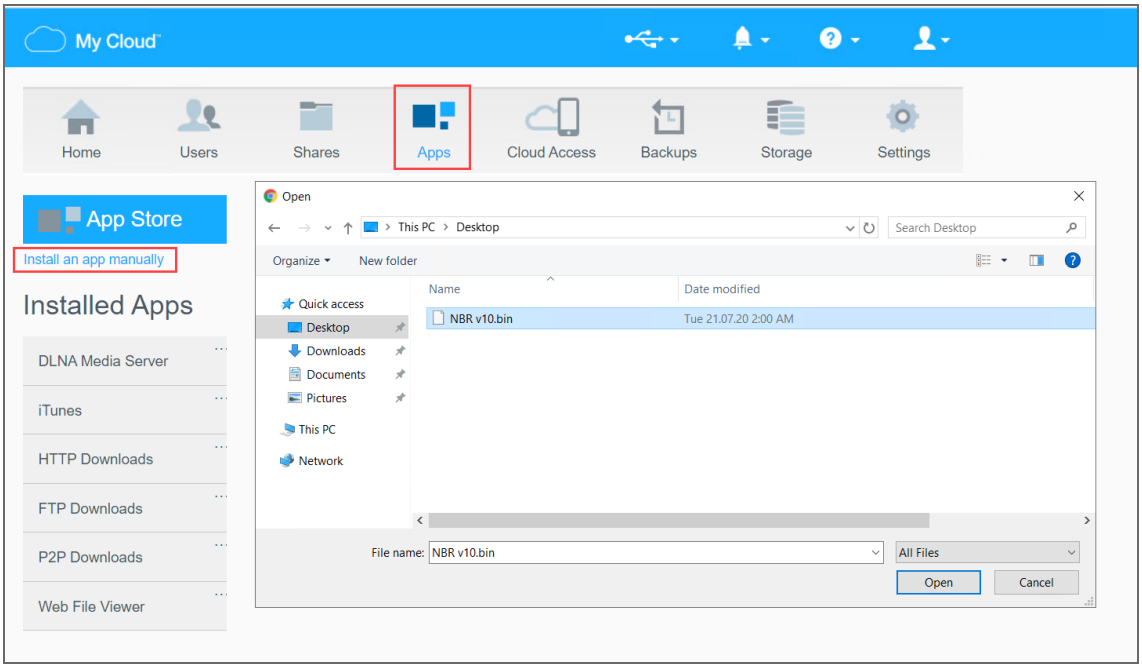
> **Note**
> A pre-shared key is not created during **Transporter**-only installation. When adding this **Transporter** to NAKIVO Backup & Replication, filling out the master password field is not required. The master password can be manually set and reset later. For details, refer to "Installed Service" on page 565.

Prior to installing NAKIVO Backup & Replication onto a Western Digital MyCloud NAS device, make sure the following requirements have been met:

1. Your Western Digital MyCloud NAS model is supported by NAKIVO Backup & Replication.
2. You have access to the NAS **My Cloud** Dashboard.
3. You have NAKIVO Backup & Replication installer for Western Digital NAS available on your computer.

Follow the steps below to install NAKIVO Backup & Replication on a Western Digital MyCloud NAS device:

1. On the **My Cloud** dashboard, click **Apps**. The list of installed NAS applications opens on the left side of the page.
2. Above the list of NAS installed applications, click **Install an app manually**. The **File Upload** dialog opens.




3. In the **File Upload** dialog, navigate to your copy of NAKIVO Backup & Replication installer and click **Open**. The installation progress bar opens.
4. When the installation finishes successfully, a dialog box opens with a message informing you about it. Click **OK** to close the dialog box.

After the installation is complete, NAKIVO Backup & Replication will appear in the list of installed NAS applications. To access the product, do either of the following:

- Open the `https://<NAS_IP>:4443` address in your browser.
- In the list of installed NAS applications, click **NAKIVO Backup & Replication** and then click **Configure**.

Refer to "Getting Started" on page 333 to better understand how to continue working with NAKIVO Backup & Replication.

# Installing on ASUSTOR NAS

You can install an ASUSTOR package with either all NAKIVO Backup & Replication components (Director, Transporter, Backup Repository) or a Transporter only.

NAKIVO Backup & Replication can be installed directly on a supported ASUSTOR NAS to create your own, high-performance backup appliance. With the appliance, all VM data protection components are unified in a single system that is fast to deploy and easy to manage, while also not consuming your environment's valuable resources. Moreover, you are getting an all-in-one backup hardware, backup software, backup storage, and data deduplication in a single box.


- "Installing on ASUSTOR NAS via App Central" on page 286
- "Installing on ASUSTOR NAS Manually" on page 288


**Note**

A pre-shared key is not created during Transporter-only installation. When adding this Transporter to NAKIVO Backup & Replication, filling out the master password field is not required. The master password can be manually set and reset later. For details, refer to "Installed Service" on page 565.

# Installing on ASUSTOR NAS via App Central

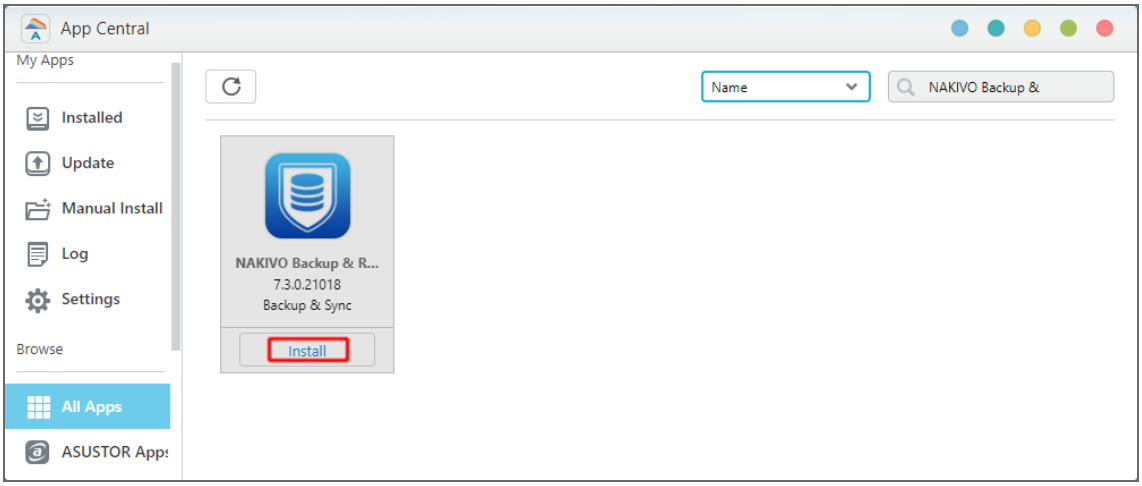
Before you begin installing NAKIVO Backup & Replication on a NAS make sure your NAS model is supported.

To install NAKIVO Backup & Replication on ASUSTOR NAS via App Central:

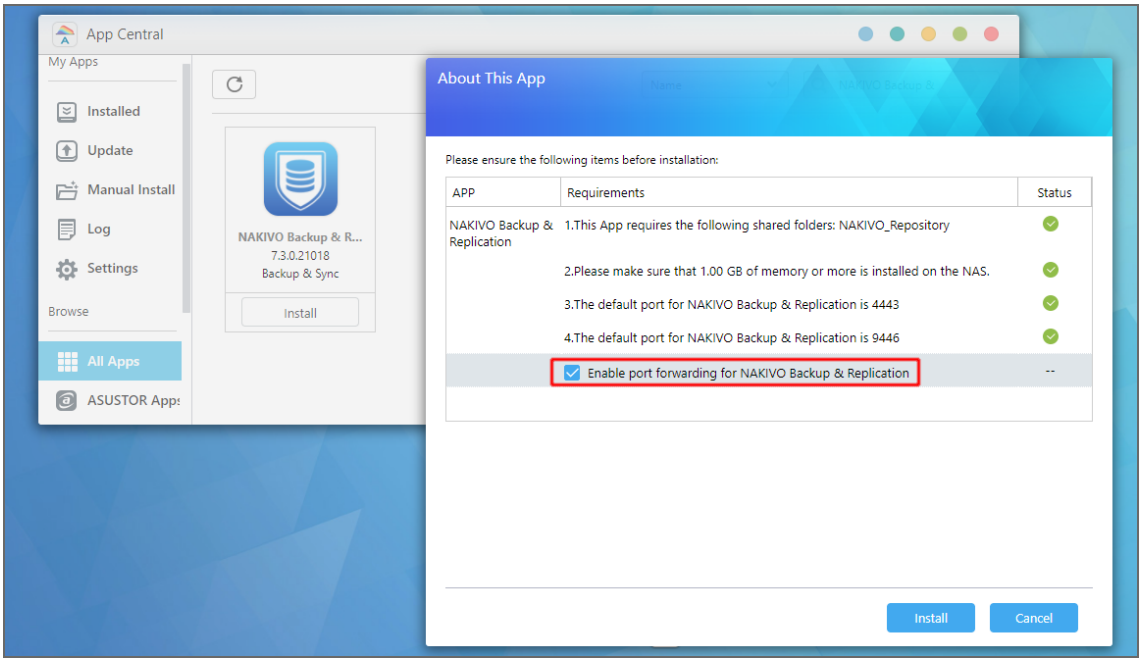
1. Open the ASUSTOR Desktop in your browser by entering the IP address of your ASUSTOR NAS.




2. Go to **App Central**.
3. Go to **Browse** > **All Apps**.
4. Find **NAKIVO Backup & Replication** in the store. Alternatively, enter **Nakivo** in the search box.
5. Click **Install**.




6. In the **About This App** dialog box that opens, select **Enable port forwarding for NAKIVO Backup & Replication** and then click **Install**.




7. Wait until the installation is complete.

By default, the NAKIVO Backup & Replication interface is available at the IP address of your ASUSTOR NAS on the port 4443: `https://<IP_address_of_ASUSTOR_NAS>:4443`.

> **Note**
>
> If you have installed NAKIVO Backup & Replication on an ARM-based NAS, an additional Transporter needs to be added to the application outside of the ARM-based NAS to allow working with VMware vCenters and ESXi hosts. Please refer to the "Adding VMware vCenters and ESXi hosts" on page 511 topic for details.

Refer to "Getting Started" on page 333 to understand better how to continue working with NAKIVO Backup & Replication.

# Installing on ASUSTOR NAS Manually

Before you begin installing NAKIVO Backup & Replication on a NAS, make sure your NAS model is supported and you have downloaded an installer (.apk file) for ASUSTOR NAS.

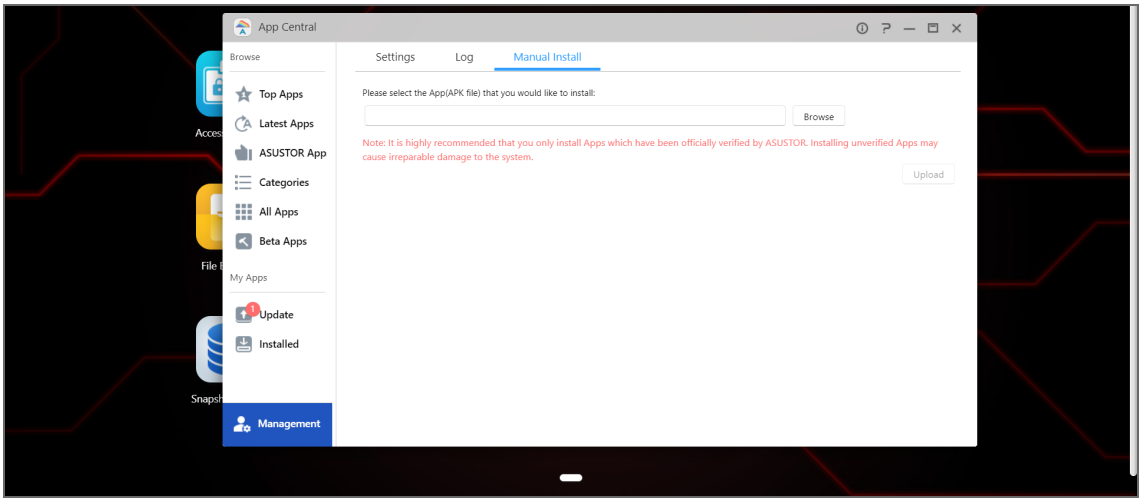
> **Note**
>
> Installing the NAKIVO Backup & Replication instance on your NAS-based backup appliance requires accurate knowledge of the hardware system information.
>
> Refer to the following page to learn how to identify your ASUSTOR device CPU model: How to Identify the CPU Architecture of an ASUSTOR NAS Processor.

To manually install NAKIVO Backup & Replication on ASUSTOR NAS:

1. Open the ASUSTOR Desktop in your browser by entering the IP address of your ASUSTOR NAS.
2. Go to **App Central**.

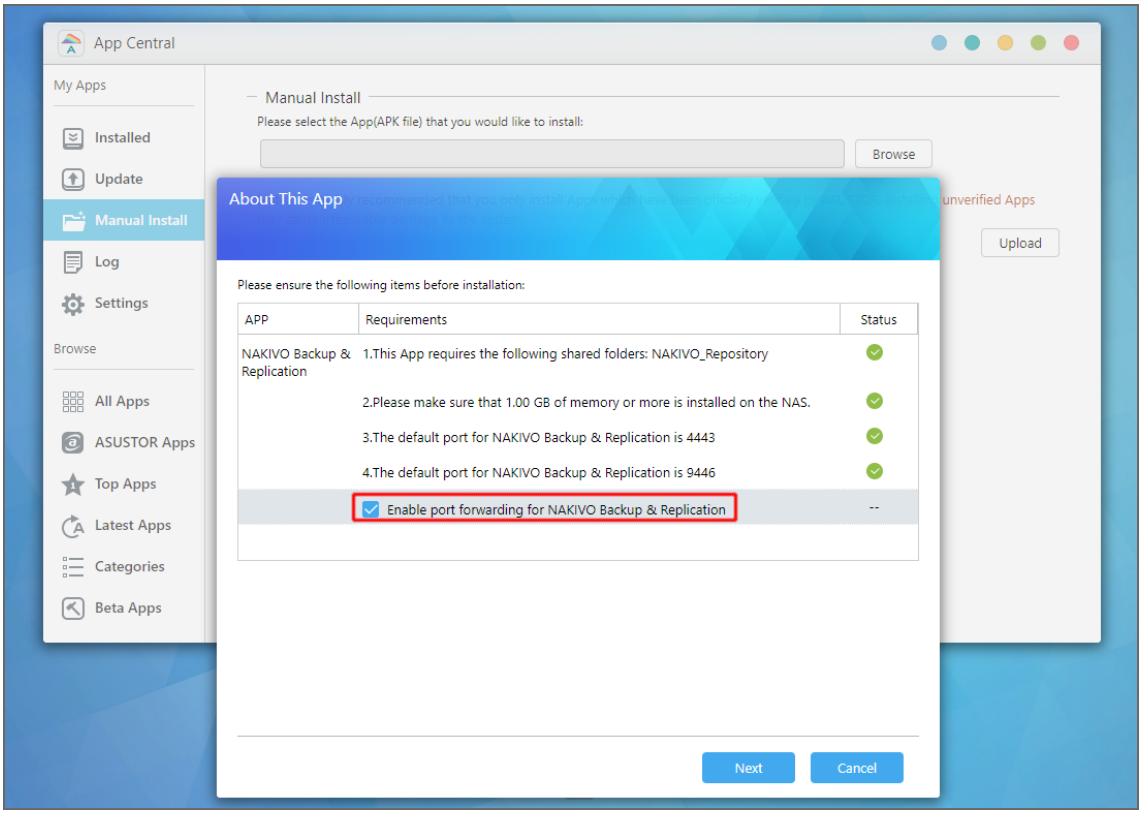



3. Click **Management** > **Manual Install**.




4. Click **Browse**. In the dialog box that opens, locate the installer (.apk file) on your computer.

5. Click **Upload**.
6. In the **About This App** dialog box that opens, check **Enable port forwarding for NAKIVO Backup & Replication**.




7. Click **Next**.
8. In the warning dialog box that opens, select **I understand the risks associated with installing unverified apps**.
9. Click **Install**.
10. Wait until the installation is complete.

By default, the NAKIVO Backup & Replication interface is available at the IP address of your ASUSTOR NAS on the port 4443: `https://<IP_address_of_ASUSTOR_NAS>:4443`.

> **Note**
>
> If you have installed NAKIVO Backup & Replication on an ARM-based NAS, an additional Transporter needs to be added to the application outside of the ARM-based NAS to allow working with VMware vCenters and ESXi hosts. Please refer to the "Adding VMware vCenters and ESXi hosts" on page 511 topic for details.

Refer to "Getting Started" on page 333 to understand better how to continue working with NAKIVO Backup & Replication.

# Installing on NETGEAR ReadyNAS

You can install the NETGEAR package that includes all NAKIVO Backup & Replication components (Director, Transporter, Backup Repository) or the NETGEAR Transporter package.

NAKIVO Backup & Replication can be installed directly on a supported NETGEAR ReadyNAS to create your own high-performance backup appliance. With the appliance, all VM data protection components are unified in a single system that is fast to deploy and easy to manage, while also not consuming your environment's valuable resources. Moreover, you are getting an all-in-one backup hardware, backup software, backup storage, and data deduplication in a single box. For installation instructions, refer to the following topics:

- “Installing on NETGEAR ReadyNAS via Available Apps” on page 291
- “Installing on NETGEAR ReadyNAS Manually” on page 292

**Note**

A pre-shared key is not created during **Transporter**-only installation. When adding this **Transporter** to NAKIVO Backup & Replication, filling out the master password field is not required. The master password can be manually set and reset later. For details, refer to “Installed Service” on page 565.

# Installing on NETGEAR ReadyNAS via Available Apps

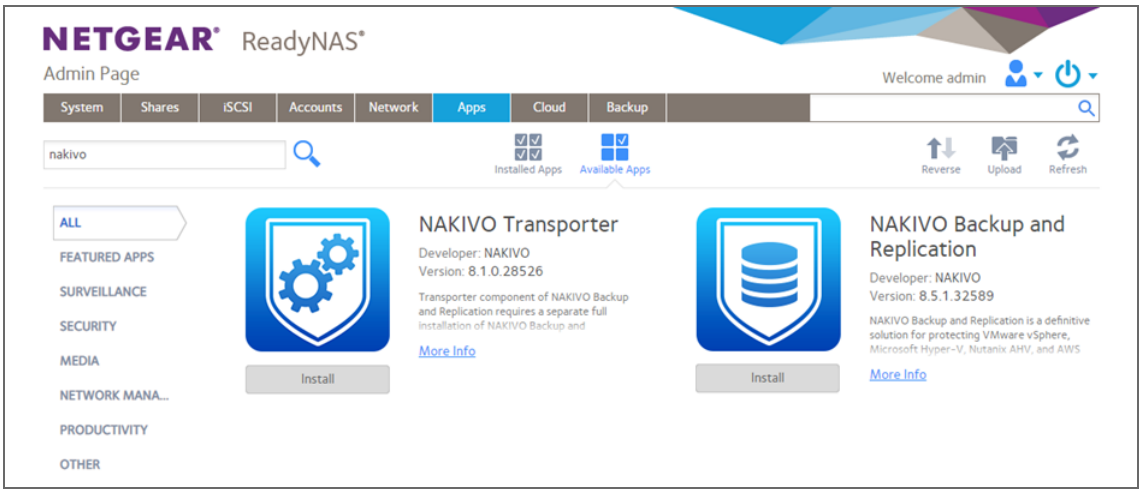
**Important**

This method is not possible for ReadyNAS OS 6.10.10. To use this method for installing NAKIVO Backup & Replication, you need to use or downgrade to ReadyNAS OS 6.10.9.

Before you begin installing NAKIVO Backup & Replication or NAKIVO Transporter on a NETGEAR ReadyNAS device, please check if your NETGEAR ReadyNAS model is supported.

To install NAKIVO Backup & Replication or NAKIVO **Transporter**, take the following steps:

1. Open the NETGEAR ReadyNAS **Admin Page** in your browser by entering the IP address of your NAS.
2. Go to **Apps** -> **Available Apps**.
3. Find **NAKIVO Backup & Replication** or **NAKIVO Transporter** in the list of available applications. Alternatively, you can enter `NAKIVO` to the filtering box in the upper left corner of the **Admin Page**.
4. Click the **Install** button below the corresponding item.




**Note**

Make sure that only one instance of the NAKIVO solution - either Full Product or **Transporter**-only - is installed on the device concurrently. Having both products installed at once may lead to incorrect operation.

5. Wait until the installation is completed.

By default, the **NAKIVO Backup & Replication** interface is available at the IP address of your NETGEAR ReadyNAS on the port 4443: `https://<IP_address_of_ NETGEAR_ReadyNAS>:4443`.

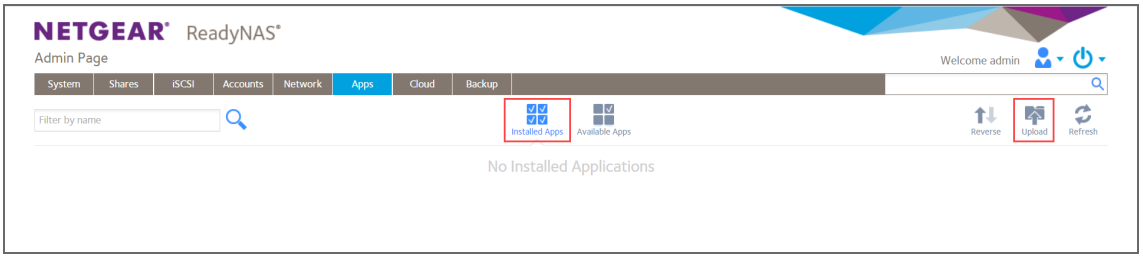
Refer to "Getting Started" on page 333 to know how to continue working with NAKIVO Backup & Replication.

# Installing on NETGEAR ReadyNAS Manually

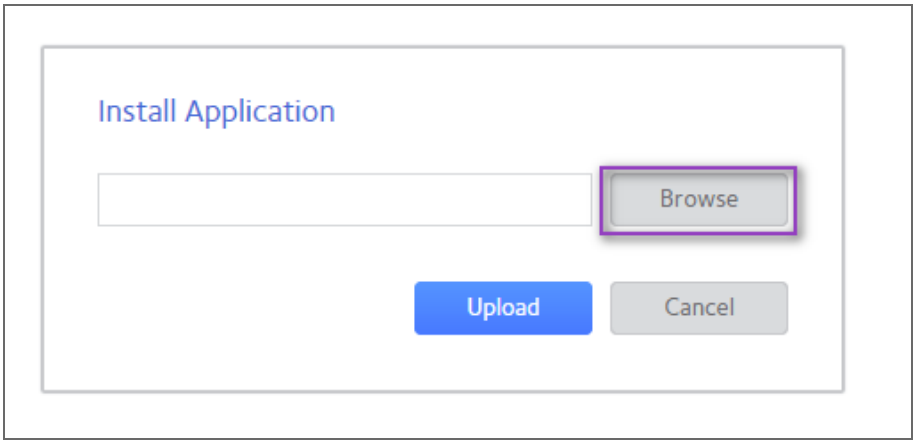
Before you begin installing NAKIVO Backup & Replication or NAKIVO Transporter on a NETGEAR ReadyNAS device, make sure your NAS model is supported and you have downloaded a relevant installer (`.deb` file) for NETGEAR ReadyNAS.

To install NAKIVO Backup & Replication or NAKIVO **Transporter**, take the following actions:

1. Open the NETGEAR ReadyNAS **Admin Page** in your browser by entering the IP address of your NAS.
2. Go to **Apps** and click **Upload**.




3. The **Install Application** dialog box opens. Click **Browse**.




4. In the dialog box that opens, locate the downloaded installer (`.deb` file) and then click **Upload**.
5. Wait until the installation has been completed.

> **Note**
>
> Make sure that only one instance of the NAKIVO solution - either Full Product or **Transporter**-only - is installed on the device concurrently. Having both products installed may lead to incorrect operations.

By default, NAKIVO Backup & Replication interface is available at the IP address of your NETGEAR ReadyNAS on the port 4443: `https://<IP_address_of_NETGEAR_ReadyNAS>:4443`.

Refer to "Getting Started" on page 333 to understand better how to continue working with NAKIVO Backup & Replication.

# Installing on Generic ARM-Based Devices

NAKIVO Backup & Replication can be deployed on ARMv7/ARMv8 computers by downloading and running an appropriate script within a Linux-based OS supported by NAKIVO Backup & Replication.

1. Download a package suitable for your setup from the downloads page.

   **Note**
   Installing the NAKIVO Backup & Replication instance on your ARM-based backup appliance requires accurate knowledge of the hardware system information.
   Refer to the following page to learn how to identify your ARM-based device CPU model: How to Identify the CPU Architecture of a Generic ARM-Based Device.

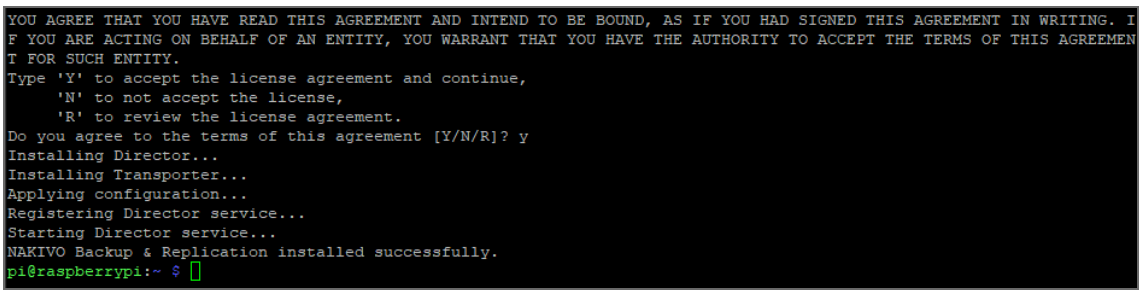
2. Upload the installer file to the machine on which you wish to install NAKIVO Backup & Replication in the binary transfer mode. For example:
   1. Upload the installer from a Windows-based machine
   2. Upload the product from a Linux-based machine: run the following command: wget 'server_ip/shared/NAKIVO Backup_&_Replication v8.5.0.30224 Installer-NAS-ARM-TRIAL.sh'
3. Log in to the Linux machine and allow for the execution of the installer file.

   **Example**
   chmod +x ./NAKIVO Backup_&_Replication v8.5.0.30224 Installer-NAS-ARM-TRIAL.sh

4. Execute the installer file with root privileges.

   **Example**
   sudo ./NAKIVO Backup_&_Replication v8.5.0.30224 Installer-NAS-ARM-TRIAL.sh

5. Accept the License Agreement by typing [Y] and hit **Enter**. You can review the license agreement by typing [R]. Rejecting [N] the license agreement will terminate the installation process and the product will not be installed.
6. The system will notify you when the installation is successfully completed.




Refer to "Getting Started" on page 333 to better understand how to continue working with NAKIVO Backup & Replication.

# Installing on TrueNAS

Make sure the following prerequisites are met:

- You have access to the TrueNAS system.
- Your TrueNAS system meets system requirements for installing NAKIVO Backup & Replication.
- The `iocage` jail/container manager is installed on your TrueNAS system. Refer to the `iocage` README page for a description.
- A storage pool is created on your TrueNAS system. Make sure the pool has enough storage for all NAKIVO Backup & Replication functionality. Refer to TrueNAS User Guide for more details on creating storage pools.
    - To create local Repository on TrueNAS local storage outside the jail, make sure you use the following mount points:
      Source: */mnt/share/*
      Destination: */mnt/test/iocage/jails/nbr/root/usr/repo*

      **Notes**
      - The *repo* is an empty folder created by the user.
      - The path for adding a repository on NAKIVO Backup & Replication is */usr/repo*.
- Make sure to check limitations for TrueNAS here.

Follow the steps below to install NAKIVO Backup & Replication on a TrueNAS system:

1. Log in to the TrueNAS system via SSH.
2. Go to the `tmp` folder: `cd /tmp`
3. Download the necessary `json` file:
    - for the full NAKIVO Backup & Replication installation on a TrueNAS v13:
      `wget` https://github.com/NAKIVO/iocage-plugin-nbr/raw/master/nbr.json
    - for the NAKIVO Backup & Replication Transporter installation on a TrueNAS v13:
      `wget` https://github.com/NAKIVO/iocage-plugin-nbr-transporter/raw/master/nbr-transporter.json

      **Notes**
      If a utility for downloading files like `wget` or `curl` is missing on your TrueNAS system, you can first download the necessary file to your local machine and then upload it to TrueNAS with a third-party tool like WinSCP or FileZilla.
4. Install NAKIVO Backup & Replication with the `iocage` jail/container manager:

**Note**

Make sure that the jail IP address is not the IP address of your TrueNAS system.

- For the full NAKIVO Backup & Replication installation on a TrueNAS:
  `iocage fetch -P nbr.json vnet="off" ip4="inherit" ip4_addr="em0|x.x.x.x/24"`
- For the NAKIVO Backup & Replication **Transporter** installation on a TrueNAS:
  `iocage fetch -P nbr-transporter.json vnet="off" ip4="inherit" ip4_addr="em0|x.x.x.x/24"`

**Note**

Run the following command to select the desired release if it was not found:
`iocage fetch --server ftp-archive.freebsd.org --root-dir /pub/FreeBSD-Archive/old-releases/amd64/`

5. For the NAKIVO Backup & Replication **Transporter** installation, add the **Transporter** to the **Director**. Refer to "Adding Existing Nodes" on page 565 for details.

# Updating NAKIVO Backup & Replication

NAKIVO Backup & Replication automatically checks for updates once each day. If an update is available, a notification is displayed in the product web interface. Click the notification link to view information about the update.

Starting from v8.5, a full solution of the NAKIVO Backup & Replication installed on Windows or Linux can be updated automatically. Should you find that product auto updating is not supported or there are some network issues, you can update the product manually. For more details, refer to the corresponding articles below.

**Notes**

- It is highly recommended that you update your product from any version within its support period. Check NAKIVO Backup & Replication Support Lifecycle Policy for more details.
- No new job runs can be started during the update.
- Updating to a major (e.g. 10.8 to 11.0), minor (e.g. 10.0 to 10.5), or sub-minor version (e.g. 10.0.0 to 10.0.1) will not reset the current license. Check NAKIVO Licensing Policy for more details about product licensing.
- Updating the product running the NAKIVO Backup & Replication version more than N-1 requires an interim update. For example, to update the product from v9.3 (N-2) to v11.0 (N), it first must be updated to v10.x, where N is the current major version.

To manually update any copy of NAKIVO Backup & Replication, go to the download page with updates.

To update your copy of the product to a newer version, you need to download an appropriate updater and run it on:

- Each machine on which you have additionally installed the **Transporter**.
- The machine on which the **Director** is installed.

Refer to the following topics for more information:


- "NAKIVO Backup & Replication Support Lifecycle Policy" on the next page
- "Updating Virtual Appliance" on page 301
- "Updating on Windows" on page 306
- "Updating on Linux" on page 308
- "Updating on Synology NAS" on page 310
- "Updating on Western Digital NAS" on page 313
- "Updating on Amazon EC2" on page 314
- "Updating on QNAP NAS" on page 320

- "Updating on ASUSTOR NAS" on page 323
- "Updating on NETGEAR ReadyNAS" on page 325
- "Updating on TrueNAS" on page 327
- "Updating on Generic ARM-Based Devices" on page 327
- "Software Update" on page 452


# NAKIVO Backup & Replication Support Lifecycle Policy

The level of support for various versions of NAKIVO Backup & Replication varies based on their lifecycle phase. The most recent product versions receive full support, which includes updates, hotfixes, and patches. Conversely, older versions may be covered with limited support.

The periods of the NAKIVO Backup & Replication Support Lifecycle are as follows:

- **Hotfix period:** 1 year after release. Within this period, NAKIVO can release hotfixes/patches to fix different product issues.
- **Support period**: 3 years after release. No support will be offered when the support period expires.

**Important**

It is highly recommended that you update your product from any version within its support period.

Please refer to the table below for the list of NAKIVO Backup & Replication product versions and their lifecycle periods.

| Version | Release date | End of Fix | Security Patches | End of Support |
|---|---|---|---|---|
| **11.1** | 6 October 2025 | 6 October 2026 | 6 October 2027 | 6 October 2028 |
| **11.0.4** | 1 September 2025 | 1 September 2026 | 1 September 2027 | 1 September 2028 |
| **11.0.3** | 22 April 2025 | 22 April 2026 | 22 April 2027 | 22 April 2028 |
| **11.0.2** | 17 March 2025 | 17 March 2026 | 17 March 2027 | 17 March 2028 |
| **11.0.1** | 13 January 2025 | 13 January 2026 | 13 January 2027 | 13 January 2028 |
| **11.0** | 4 November 2024 | 4 November 2025 | 4 November 2026 | 4 November 2027 |
| **10.11.3** | 17 September 2024 | 17 September 2025 | 17 September 2026 | 17 September 2027 |
| **10.11.2** | 17 June 2024 | 17 June 2025 | 17 June 2026 | 17 June 2027 |

| | | | | |
|---|---|---|---|---|
| **10.11.1** | 17 April 2024 | 17 April 2025 | 17 April 2026 | 17 April 2027 |
| **10.11** | 29 January 2024 | 29 January 2025 | 29 January 2026 | 29 January 2027 |
| **10.10.1** | 6 November 2023 | 6 November 2024 | 6 November 2025 | 6 November 2026 |
| **10.10** | 9 October 2023 | 9 October 2024 | 9 October 2025 | 9 October 2026 |
| **10.9** | 4 July 2023 | 4 July 2024 | 4 July 2025 | 4 July 2026 |
| **10.8** | 23 January 2023 | 23 January 2024 | 23 January 2025 | 23 January 2026 |
| **10.7.2** | 7 November 2022 | 7 November 2023 | 7 November 2024 | 7 November 2025 |
| **10.7.1** | 12 October 2022 | 12 October 2023 | 12 October 2024 | 12 October 2025 |
| **10.7** | 31 August 2022 | 31 August 2023 | 31 August 2024 | 31 August 2025 |
| **10.6.1** | 9 June 2022 | 9 June 2023 | 9 June 2024 | 9 June 2025 |
| **10.6** | 3 May 2022 | 3 May 2023 | 3 May 2024 | 3 May 2025 |
| **10.5.1** | 22 December 2021 | 22 December 2022 | 22 December 2023 | 22 December 2024 |
| **10.5** | 16 November 2021 | 16 November 2022 | 16 November 2023 | 16 November 2024 |
| **10.4.1** | 26 October 2021 | 26 October 2022 | 26 October 2023 | 26 October 2024 |
| **10.4** | 2 August 2021 | 2 August 2022 | 2 August 2023 | 2 August 2024 |
| **10.3** | 4 May 2021 | 4 May 2022 | 4 May 2023 | 4 May 2024 |
| **10.2** | 25 January 2021 | 25 January 2022 | 25 January 2023 | 25 January 2024 |
| **10.1.1** | 30 October 2020 | 30 October 2021 | 30 October 2022 | 30 October 2023 |
| **10.1** | 21 September 2020 | 21 September 2021 | 21 September 2022 | 21 September 2023 |

| | | | | |
|---|---|---|---|---|
| **10.0.1** | 26 August 2020 | 26 August 2021 | 26 August 2022 | 26 August 2023 |
| **10.0** | 23 July 2020 | 23 July 2021 | 23 July 2022 | 23 July 2023 |
| **9.4** | 3 June 2020 | 3 June 2021 | 3 June 2022 | 3 June 2023 |

**Important**

- All earlier versions of the product not mentioned here are not supported.
- Updating the product (Director and Transporter) to a new version is prohibited in case the current version is two or more major versions older than the new version. For example: you cannot upgrade your product from 9.x to 11.x.
- Updating any version of the product to any Beta version of the product is prohibited.
- Updating any Beta version to any version of the product is prohibited

# Product Lifecycle Terminology and Milestones

## Versioning Scheme

NAKIVO classifies releases as major or minor:

- A major release contains significant new features and/or structural changes. Changes may also include deprecating existing features and breaking backward compatibility.
- A minor release may include some new functionalities and changes that enhance performance and usability. These releases preserve backward compatibility within the related major release.

For information about changes in different releases, refer to the Release Notes.

## Release Lifecycle

The lifecycle of major and minor NAKIVO Backup & Replication releases consists of 3 phases:

- **Phase 1 - Early Access / Beta**: Releases in this phase are meant for testing and feedback purposes and should not be used in production environments.
- **Phase 2 - General Availability**: Releases in this phase are fully supported and can be used in production environments.

**Notes**

- These releases may receive minor fixes within the first year from the date of general availability. During first year after GA, fixes are provided to the current version. After 1 year passed, for the next 2 years fixes are provided to the version which is on 1st year.
- Support in this case means that NAKIVO continues to deliver technical support and updates for the duration of that period.
- For information on technical support services available to you if you have a valid NAKIVO Backup & Replication subscription and/or perpetual license, refer to the NAKIVO Technical Support Policy.

- **Phase 3 - End of Support**: End of support for a minor or major release means that NAKIVO will no longer provide technical support or issue updates, patches and fixes for a specific version.

A version reaches end of support after 3 years from the date of general availability. End of support for a version applies to all solution editions and license types, including NFR licenses.

If your version is nearing the end of support, we recommend that you plan an update to a newer NAKIVO Backup & Replication version to avoid support gaps and maintain optimal data protection. You can update your instance in the product UI or download an updater from the NAKIVO website.

**Notes**

- During the general availability period, if you have a valid perpetual and/or subscription license, you can update your NAKIVO Backup & Replication version to any newer supported version.
- NAKIVO does not provide patches or fixes for third-party technology and services that reach end of life within the 3-year support period.

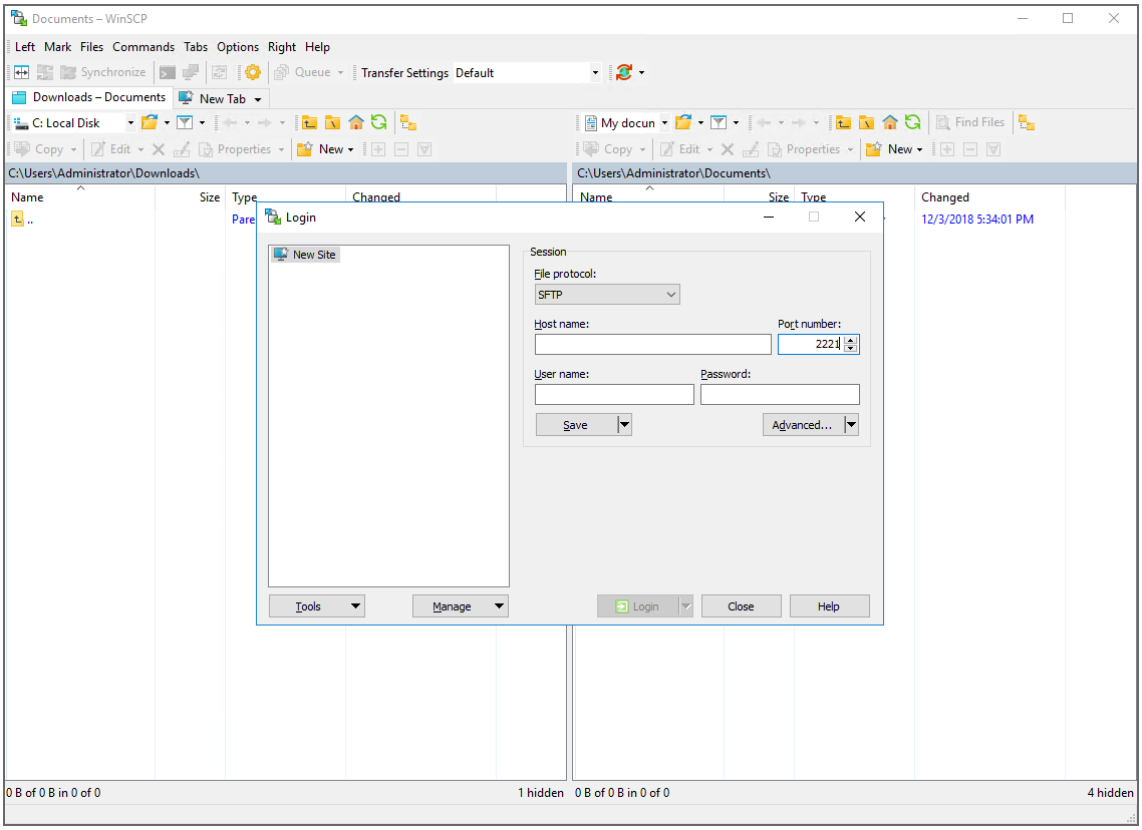
# Updating Virtual Appliance

Prior to updating your virtual appliance (VA):

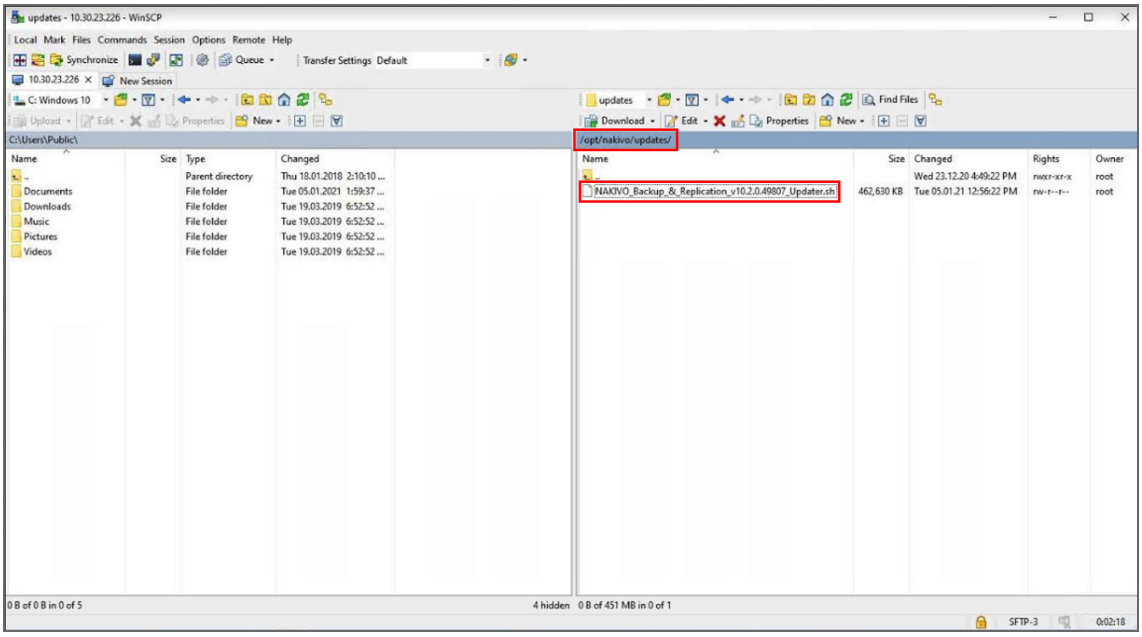
1. Make sure that no jobs or repository maintenance tasks are running in the product.
2. Create a snapshot of the VA to revert to the previous version in case any failure occurs.

Follow the steps below to update your VA:

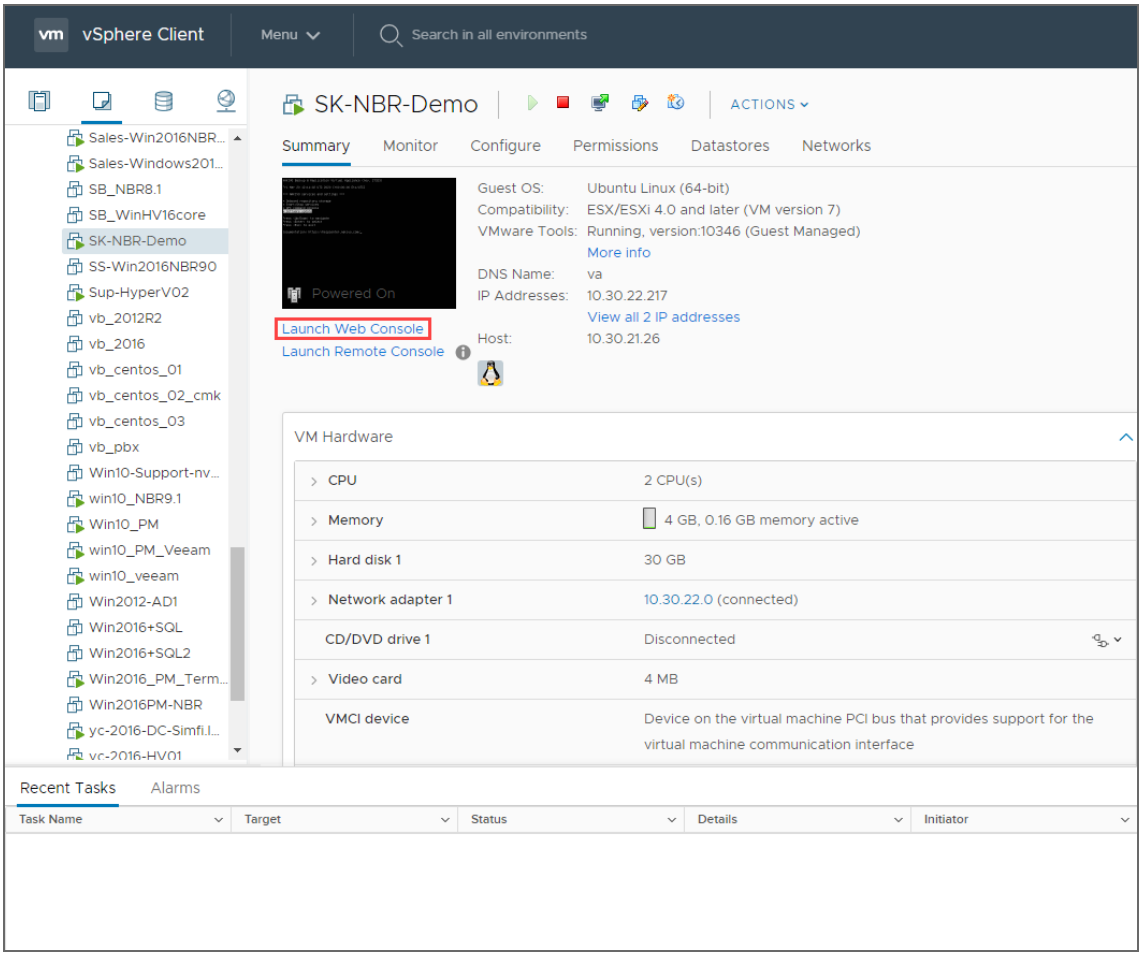
1. Using SSH client, log in to the VA that needs to be updated.




2. Download the latest *VA and Linux updater* from here.
3. Change the directory to `/opt/nakivo/updates` and locate the updater.

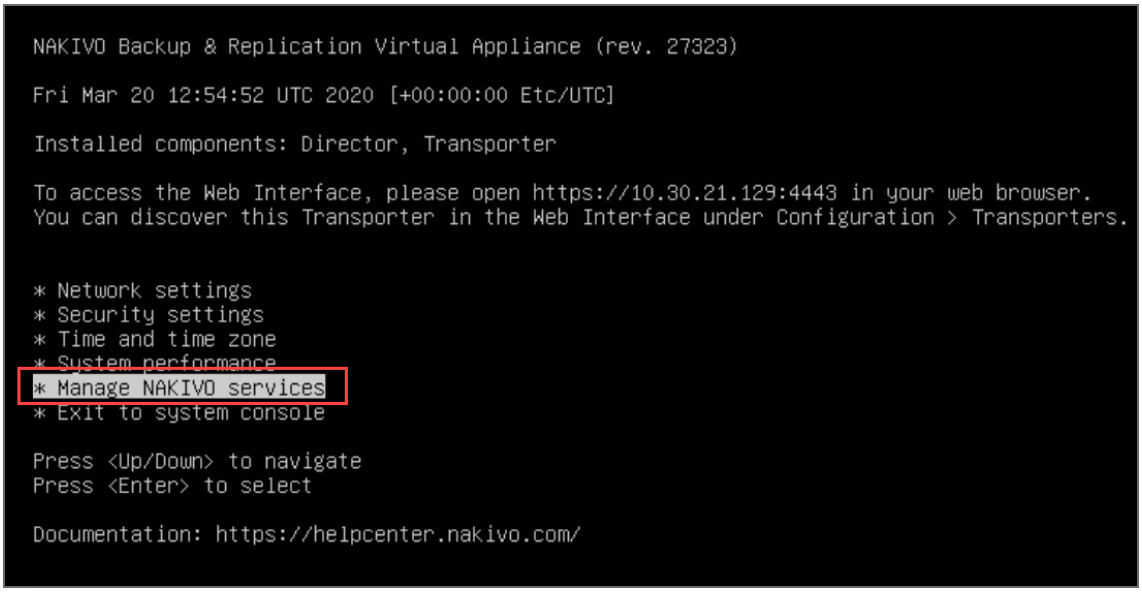



4. Log out from the SSH client.
5. Log in to your vSphere client, navigate to your VA and click **Launch Web Console**.

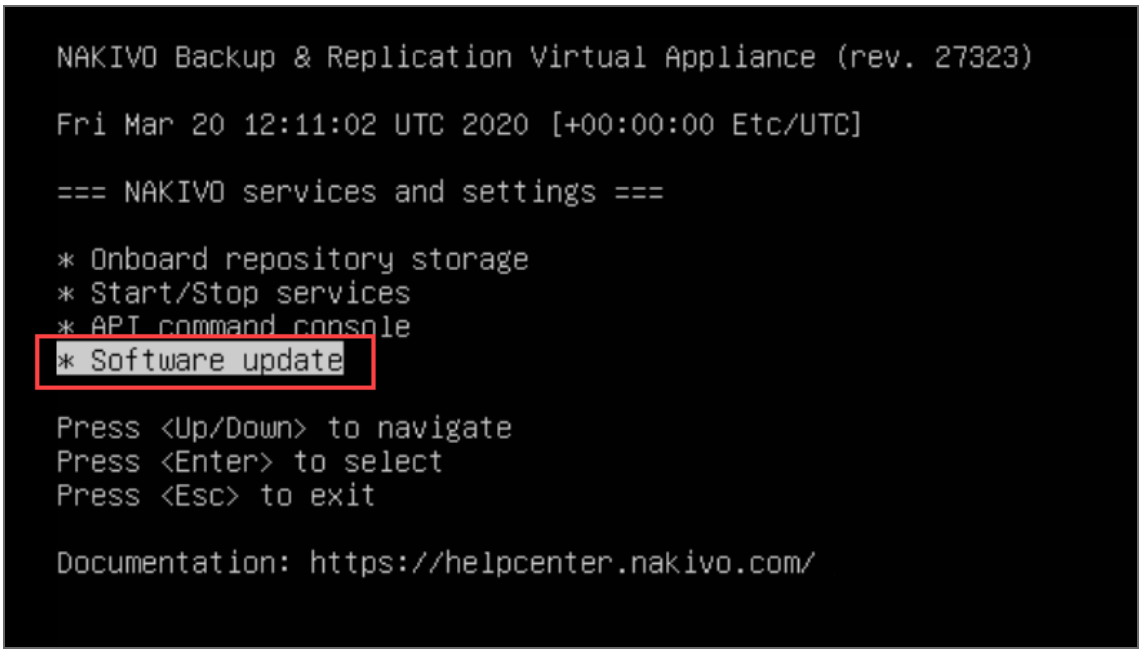



6. Do one of the following depending on the NAKIVO Backup & Replication version you use:

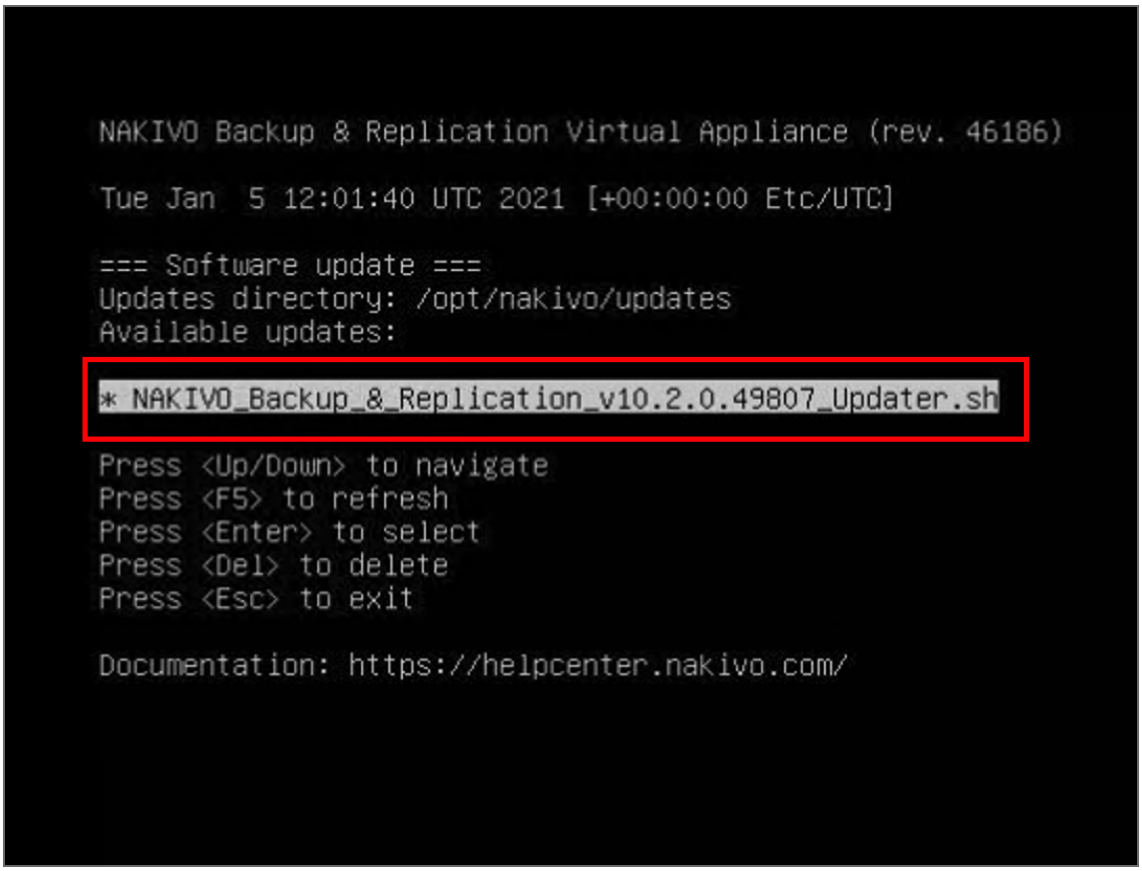
- For the product Version 8.1 and above:
  1. In the VA menu, select **Manage NAKIVO services** and press **Enter**.



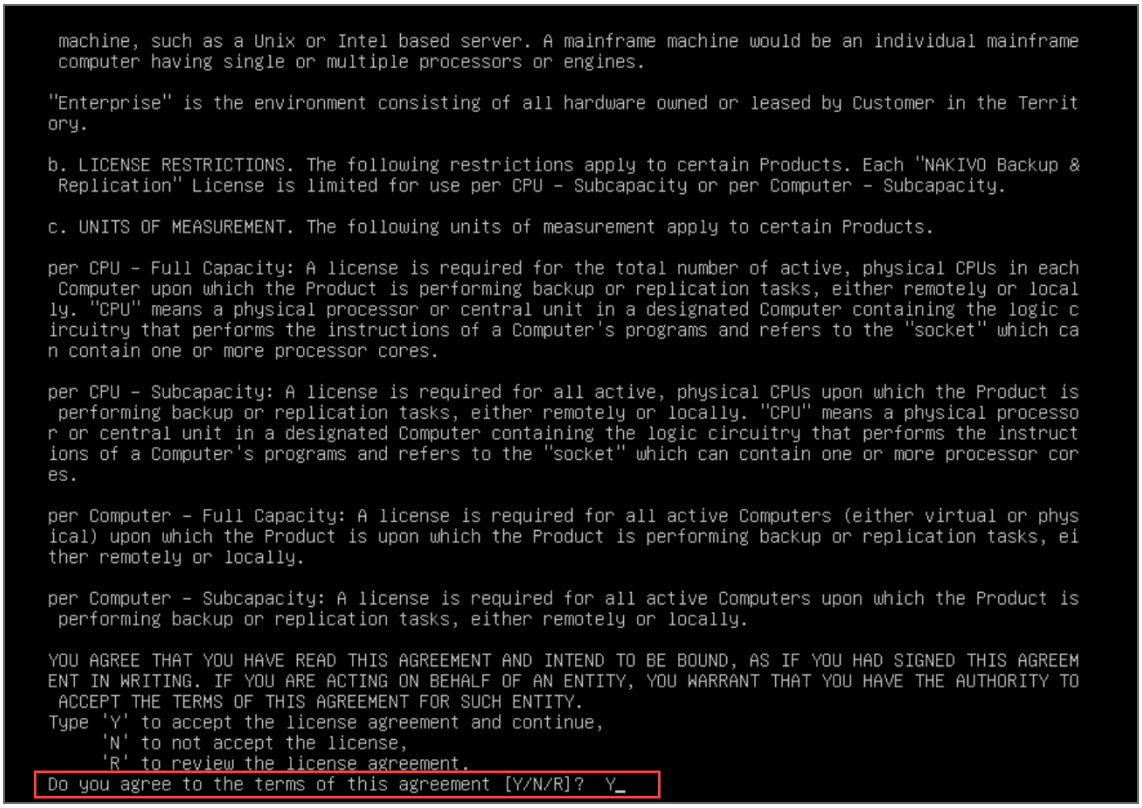

  2. In the menu that opens, select **Software update** and press **Enter**.




  3. Select the updater that you have downloaded and press **Enter**.




4. Review the End User License Agreement. Press **Space** to go to the bottom of it. If you agree to the terms of the agreement, type **Y** and then press **Enter** to begin the update process.




- For earlier product versions:
    1. In the VA menu, select **Software update** and press **Enter.**
    2. Select the updater that you have downloaded and press **Enter.**
    3. Review the End User License Agreement. Press **Space** to go to the bottom of it. If you agree to the terms of the agreement, type **Y** and then press **Enter** to begin the update process.

7. When the update process is complete, a message will appear to inform you about it. Exit the VA console.
8. Update all machines on which you have deployed an additional **Transporter**.

**Note**

Updating your VA with versions prior to the previous major version (for example, updating VA version 6.1 to version 9.0) is prohibited. Please update your VA to the next major version first.

# Updating on Windows

If auto-update within the NAKIVO Backup & Replication interface is not supported, follow the steps below to update the product manually:

1. Download the latest Windows updater from here.
2. Make sure that no jobs or repository maintenance tasks are running in the product.
   If NAKIVO Backup & Replication is installed on a VM, create a snapshot of the VM before updating the product.
3. Run the updater on the machine on which the Director is installed, and also on all machines on which you have additionally deployed a **Transporter**.
4. Optionally, you can select the **Master password** checkbox and enter the password that will be used to generate a pre-shared key and secure the **Transporter**. This option is available only for the **Transporter**-only update.

   **Notes**
   - The master password must adhere to the following requirements:
     - Minimal length - 5 characters.
     - Maximum length - 50 characters.
   - The master password can be set and re-set manually by running the command on the machine housing the **Transporter**. Follow these steps:
     - Enter the following command `bhsvc -b P@ssword123`
     - Restart the **Transporter** service.

5. Optionally, you can select the **Transporter certificate** checkbox. This allows you to use a CA Certificate. Enter the path to the folder containing the certificate file in the field.

   **Notes**
   - When the checkbox is not selected, NAKIVO Backup & Replication automatically installs a self-signed certificate.
   - If the **Transporter Certificate** checkbox is not selected, a warning window appears prompting you to install it. Click **Continue** to proceed.

6. Click **Update**.
7. When the update is complete, click **Finish**.
8. If you have entered the new master password on step 4, do the following:
   a. Go to **Settings > Transporters** and click on the **Transporter** you have changed the master password for.
   b. Select **Edit**.

c. Enter the new master password and click **Connect**.
d. The **Certificate Acceptance** dialog box appears. Verify the certificate details, and click **Accept**.
e. Click **Apply** to save the changes.
f. Click on the same **Transporter** once again and select **Refresh** to refresh the **Transporter**.

# Updating on Linux

If updating on a Linux OS within the NAKIVO Backup & Replication interface is not supported, follow the steps below to update the product manually:

1. Download the latest Linux/VA updater from here.
2. Upload the updater to the machine on which the **Director** is installed.

   **Important**
   Make sure you are using the *binary transfer mode* when uploading the updater to the machine with a Linux OS. For example:

   - Upload the installer from a Windows-based machine
   - Upload the product from a Linux-based machine: run the following command: `wget 'server_ip/shared/NAKIVO_Backup_Replication_vX.X.X_Updater.sh'`
3. Log in to the Linux machine and allow the execution of the updater file. For example: `chmod +x NAKIVO_Backup_Replication_vX.X.X_Updater.sh`
4. Make sure that no jobs or repository maintenance tasks are running in the product.
   If NAKIVO Backup & Replication is installed on a VM, create a snapshot of the VM prior to updating the product.
5. Run the updater file with root privileges. For example: `sudo ./NAKIVO_Backup_Replication_vX.X.X_Updater.sh`
6. Review the license agreement (press **Space** to go to the next page of the agreement). If you agree to the terms of the license agreement, press “Y” and then press **Enter**.
7. Enter the “Y” key and then press **Enter** to confirm that you wish to stop the services and begin the update process.
8. Update all machines on which you have additionally deployed a “Transporter” on page 123.

**Note**

In some rare cases, CIFS/NFS shares may remain mounted on the OS level even after detaching a repository in the NAKIVO web UI.

To avoid risks of damaging CIFS/NFS-based repositories during the update, follow the steps below:

1. Make sure there is no activity in NAKIVO.
2. Stop NAKIVO transporter services assigned to the attached backup repositories.
3. Check if any active mounts are left.
4. Run the `umount` command.
5. Run the NAKIVO update.

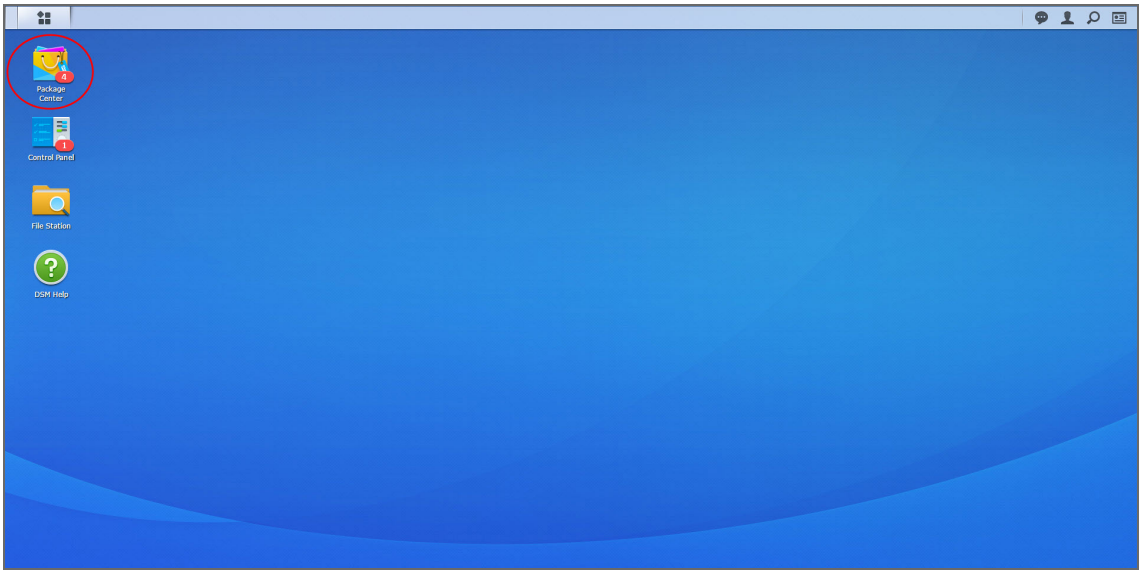
The same flow is recommended for M365 repositories.

# Updating on Synology NAS

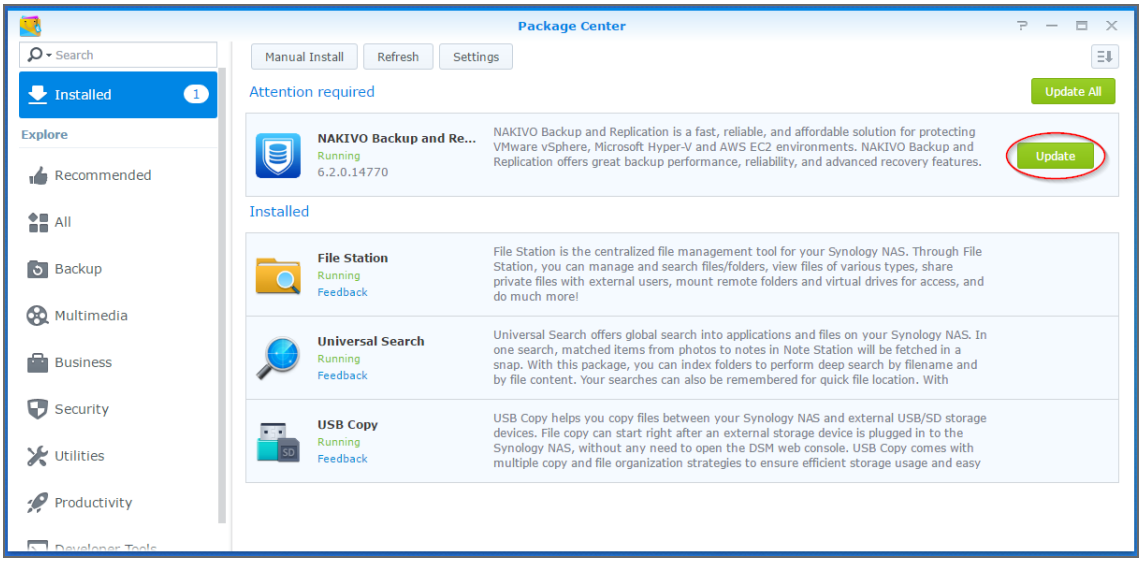
- Updating via Synology Package Center
- Updating Manually

## Updating via Synology Package Center

1. Make sure that no jobs or repository maintenance tasks are running in the product.
2. In the Synology NAS management interface, open the **Package Center**.




3. Go to the **Installed** section.
4. If there is a new version of NAKIVO Backup & Replication available, you will see an **Update** button.




5. Click **Update**.

6. Wait until the update is complete.
7. Repeat these steps on all Synology NAS where you have also installed a **Transporter**.

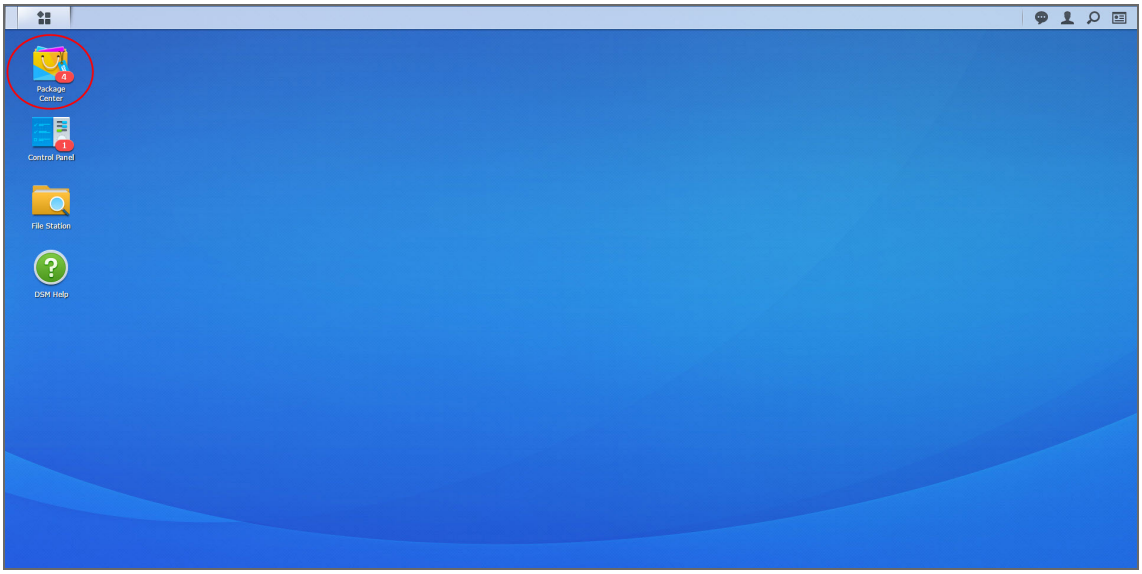
> **Note**
> If the latest version of NAKIVO Backup & Replication is not available in the **Synology Package Center**, you may update manually by following the instructions in this Knowledge Base article.

## Updating Manually

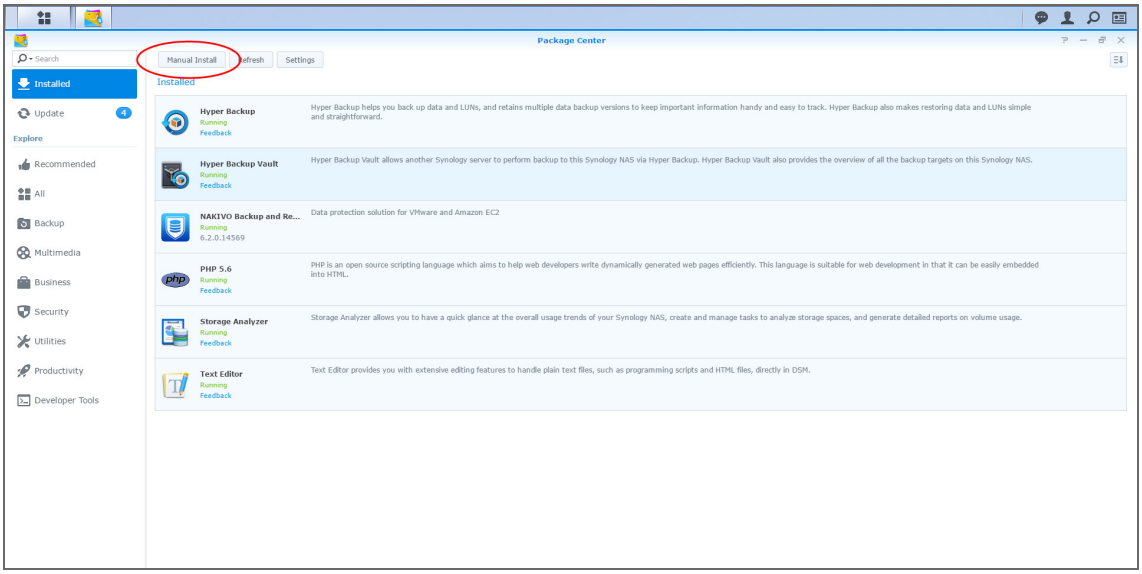
1. Download the latest Synology NAS updater from here.

> **Note**
> Installing updates of NAKIVO Backup & Replication on your NAS-based backup appliance requires accurate knowledge of the hardware system information.
> Refer to the following page to learn how to identify your Synology device CPU model: How to Identify the CPU Architecture of a Synology NAS Processor.

2. Make sure that no jobs or repository maintenance tasks are running in the product.
3. In the Synology NAS management interface, open the **Package Center**.




4. Click **Manual Install**.




5. Click **Browse**, navigate to the Synology NAS package that you have downloaded, select it, and click **Open**.
6. Click **Next**. the package is uploaded to your NAS.
7. Click **Apply**.
8. Run an appropriate updater on all machines on which you have also installed a **Transporter**.

Now, NAKIVO Backup & Replication has been updated.

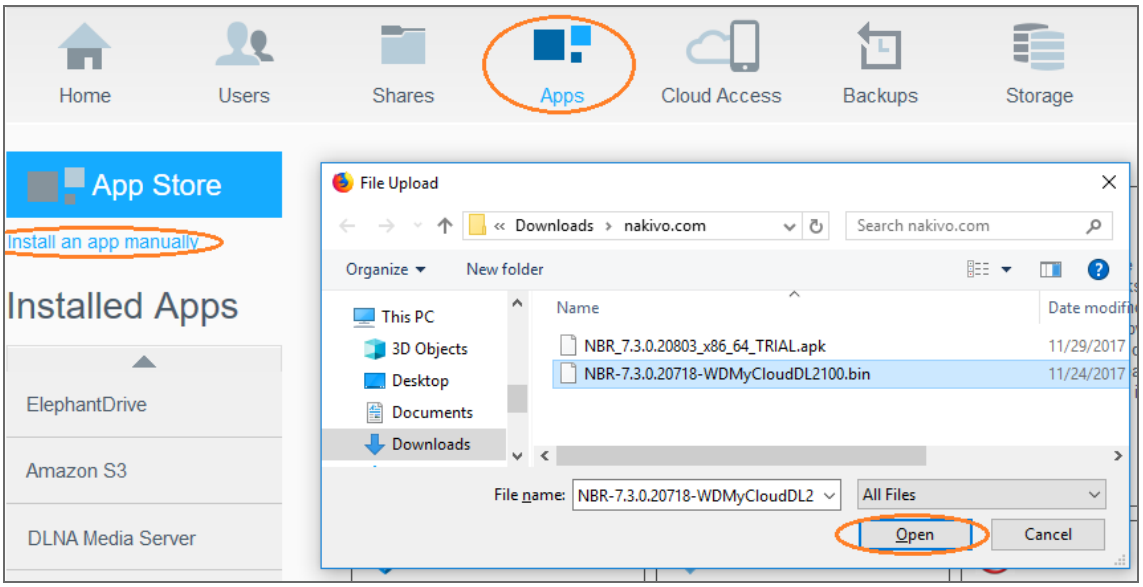
# Updating on Western Digital NAS

Prior to updating NAKIVO Backup & Replication on Western Digital MyCloud NAS, make sure the following requirements have been met:

- You have access to the Western Digital NAS MyCloud Dashboard.
- NAKIVO Backup & Replication installer is available for your Western Digital NAS.

Please follow the steps below to update NAKIVO Backup & Replication on a Western Digital MyCloud NAS device:

1. Make sure that no jobs or repository maintenance tasks are running in the product.
2. In the **My Cloud** Dashboard, click **Apps**. The list of installed NAS applications opens on the left side of the page.
3. Above the list of NAS installed applications, click **Install an app manually**. The **File Upload** dialog opens.
4. In the **File Upload** dialog, navigate to your copy of the NAKIVO Backup & Replication installer for Western Digital NAS and click **Open**. The update progress bar opens.




5. Once the update has successfully finished, a dialog box opens with a message including said information. Click **OK** to close the dialog box.

# Updating on Amazon EC2

The main installation of NAKIVO Backup & Replication (**Director** and **Transporter**) must be updated the way it is done on Linux.

**Notes**

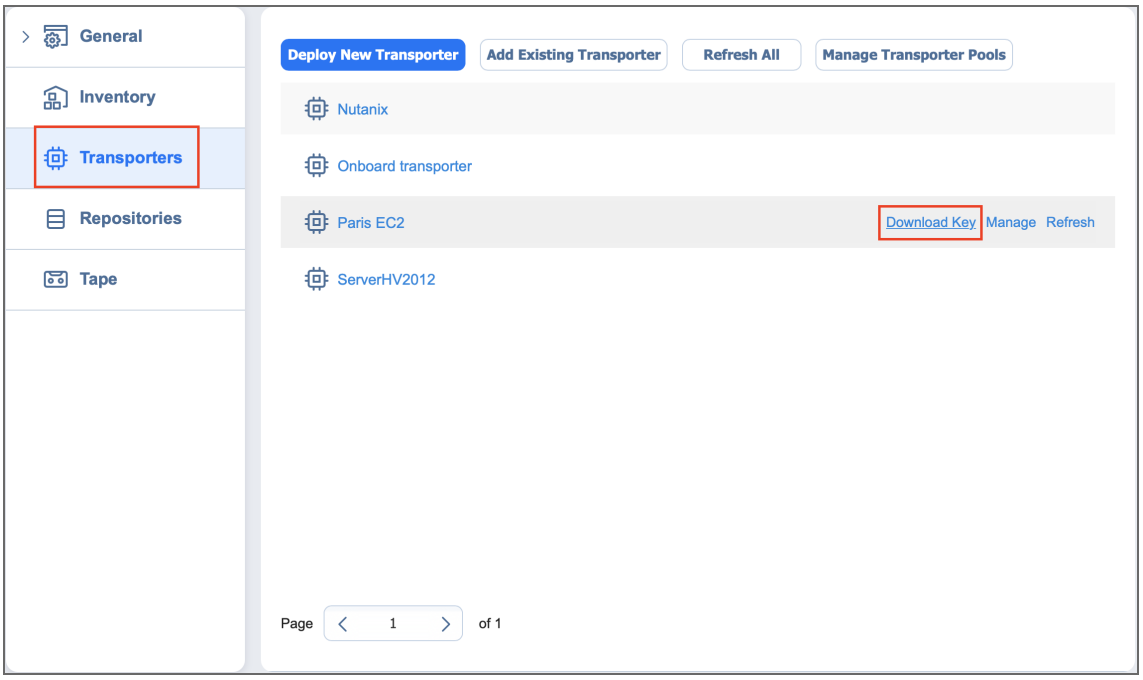
- You have to apply the **-e** argument for executing the installer, in order to avoid changing the Amazon EC2 **Transporter** with the regular Linux **Transporter**. Refer to "Installing on Linux" on page 261 for a description of the available arguments.
- Only the main installation of NAKIVO Backup & Replication needs to be updated manually. **Transporters** installed on Amazon EC2 instances are updated automatically.

# Connecting to Amazon EC2 from Windows

You can use the following free tools to connect to your Amazon EC2 instance:

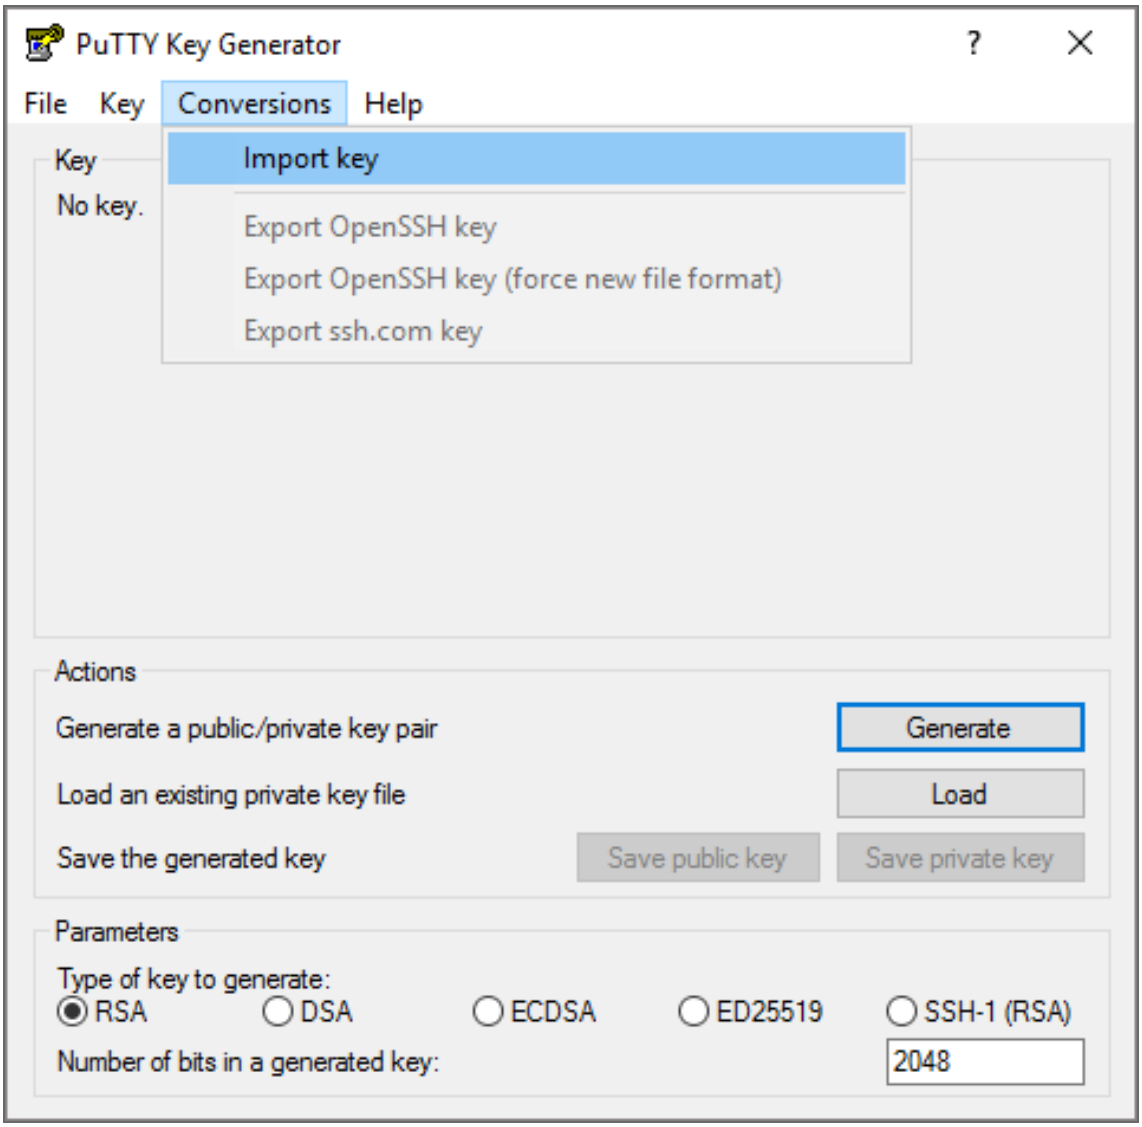
- WinSCP to upload the installer file.
- PuTTYgen tool to convert the private key.
- PuTTY tool to connect to an Amazon instance securely.

1. Log in to NAKIVO Backup & Replication.
2. Go to **Settings** > **Transporters**.
3. Download the keys of your Amazon instance.

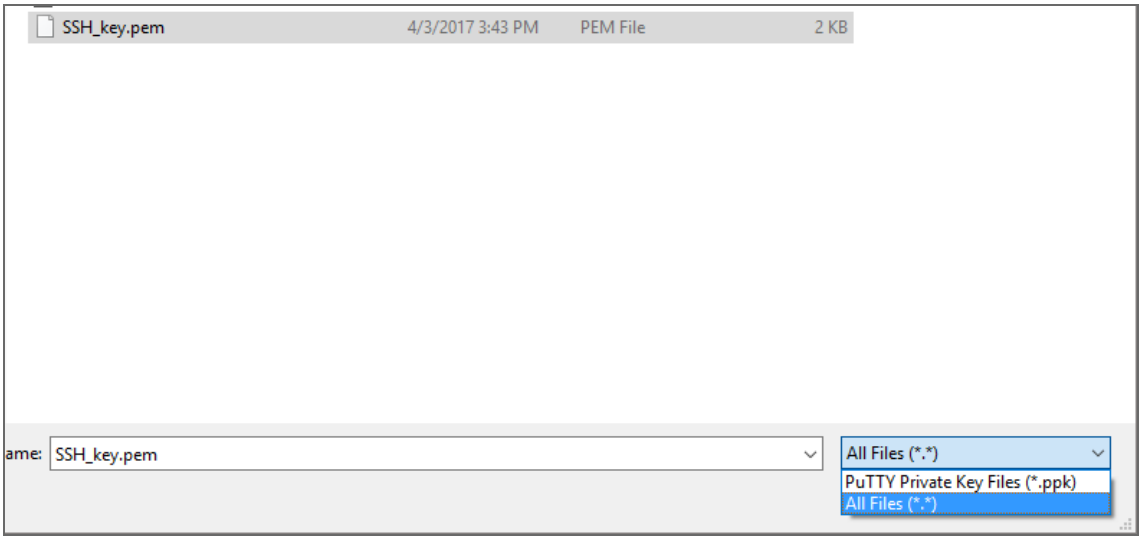



4. Click on the **Transporter** to view its details. Copy or remember the IP-address/hostname of the Amazon instance.
5. Unzip the folder with the key.
6. Convert the key using PuTTYgen:

1. In PuTTYgen menu, go to *Conversions > Import.*

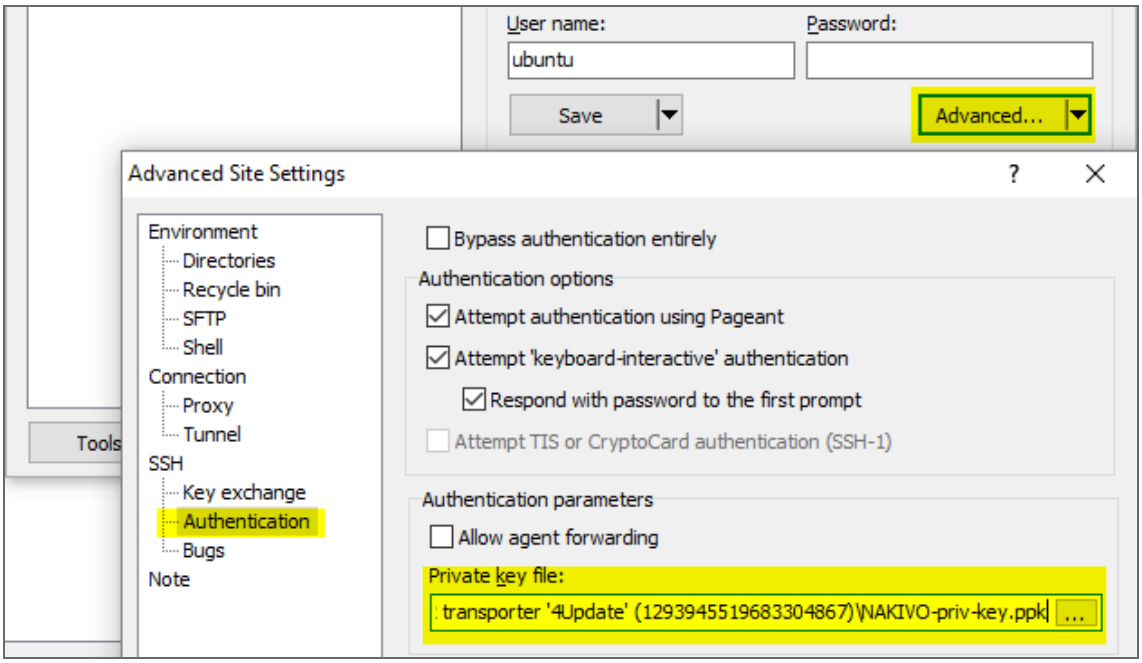



2. Locate the `SSH_key.pem` you just downloaded and unzipped. If you don't see it in the **Open**... dialogue box, change the file type to **All files**.




3. Click on **Save private key**. If PuTTYgen asks you to save the key without a passphrase, click **Yes**.

7. Open WinSCP.
8. Create a new session:
    a. Add the hostname or IP address of your Amazon instance you received on step 4 into the **Host Name** box.
    b. In the **Username** box, enter `nkvuser`.
    c. Leave the **Password** box empty.
    d. Add the private key to WinSCP:
        1. Click the **Advanced...** button.
        2. The **Advanced Site Settings** dialog box opens. Go to *SSH > Authentication > Private key file:* and select the key file you generated on step 6.




3. Click **OK**.

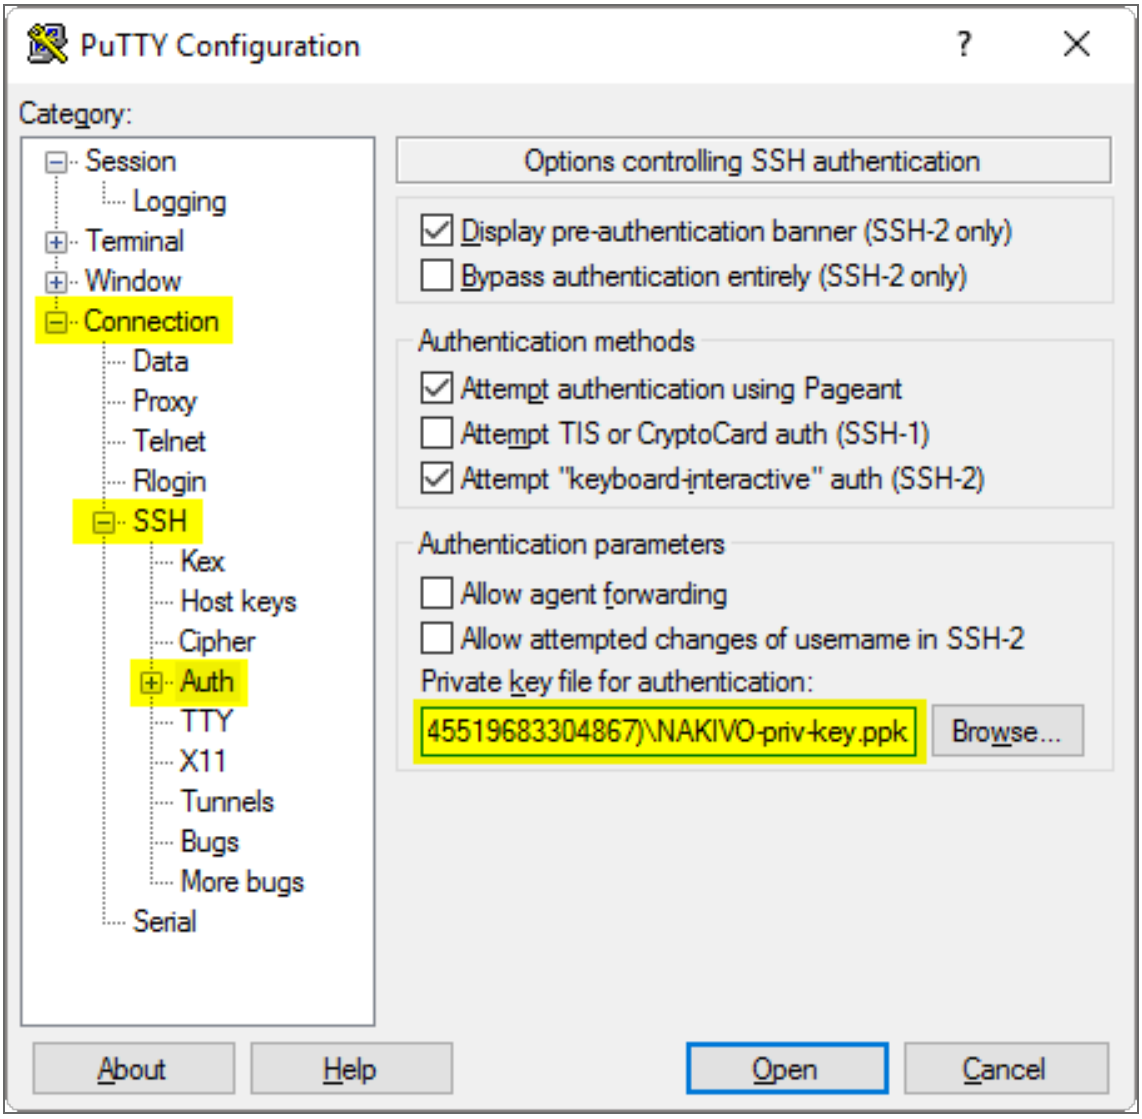
e. Click **Login**.

f. Upload the updater file.

g. Open PuTTY.

h. Enter the IP-address or hostname of the Amazon EC2 instance.

i. Go to *Connection > SSH > Auth* and add the private key in *Private key file for authentication:* box.




j. Click **Open**.

k. In the command line prompt that opens: log in to the Amazon EC2 instance:

   1. For **login**, enter `nkvuser`.
   2. For **password**, leave a blank line.

9. Update NAKIVO Backup & Replication following the instructions.

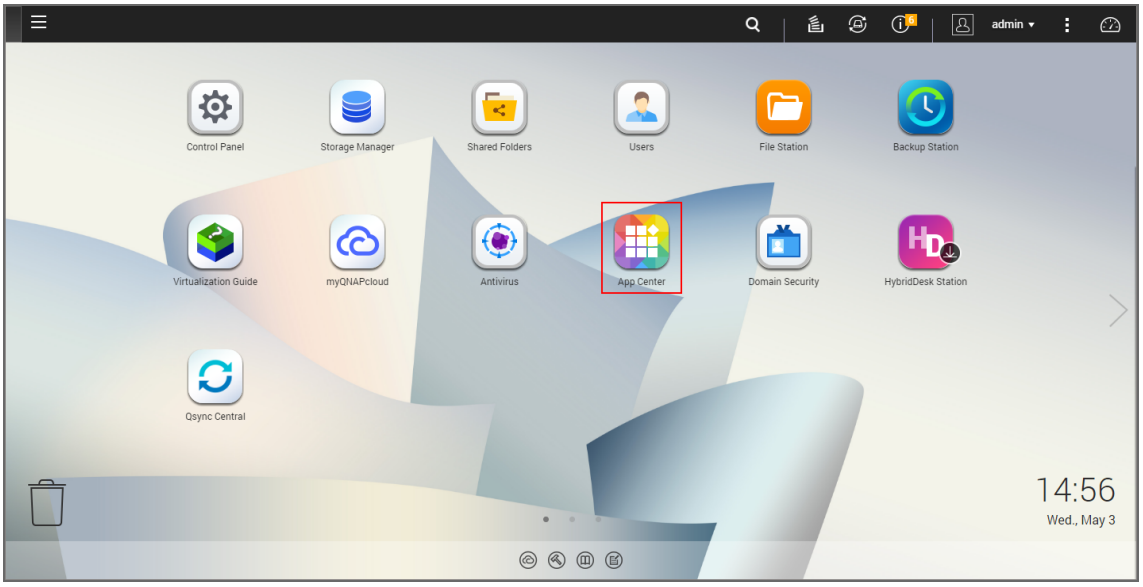
# Updating on QNAP NAS

You can update NAKIVO Backup & Replication via QNAP AppCenter or manually. Refer to the following subtopics for details:

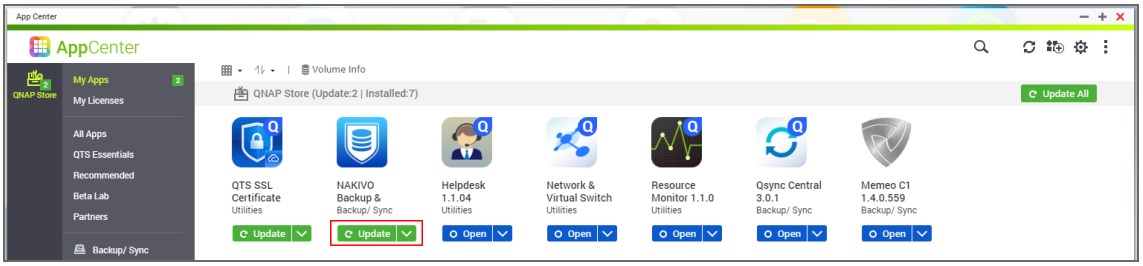
- Updating via QNAP AppCenter
- Updating Manually

## Updating via QNAP AppCenter

1. Open the QNAP Desktop in your browser by entering the IP address of your QNAP NAS.
2. Make sure that no jobs or repository maintenance tasks are running in the product.
3. Go to **App Center**.




4. Select the *Backup/Sync* category and find NAKIVO Backup & Replication. Alternatively, use the search box at the top of the App Center window: click on the magnifier icon and enter "Nakivo".
5. If the new version of NAKIVO Backup & Replication is available in the QNAP App Center, you will see a green **Update** button.



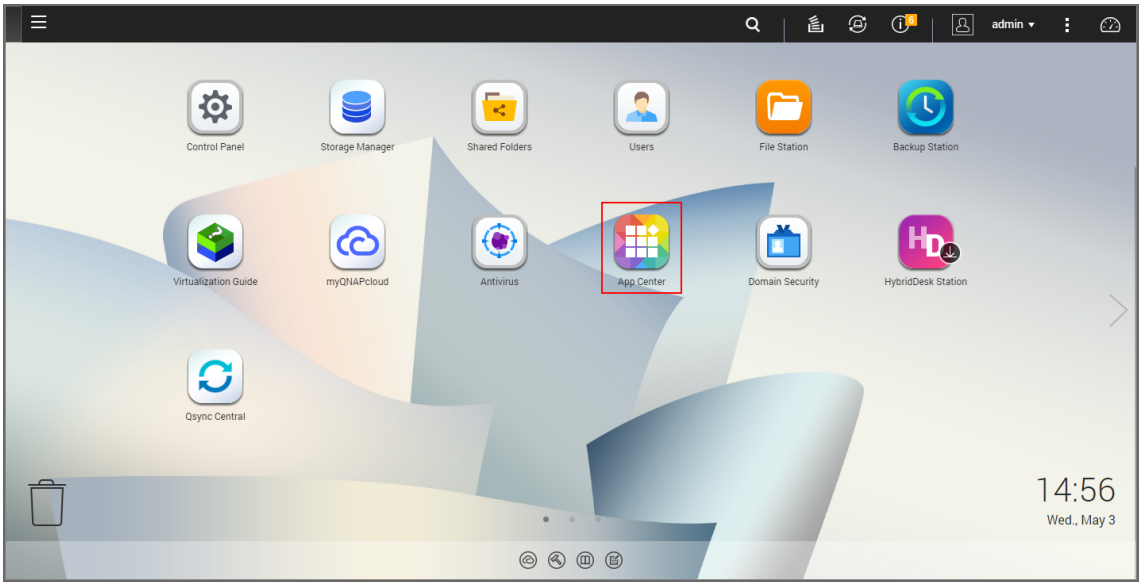

6. Click the **Update button** and wait till update finishes.

# Updating Manually

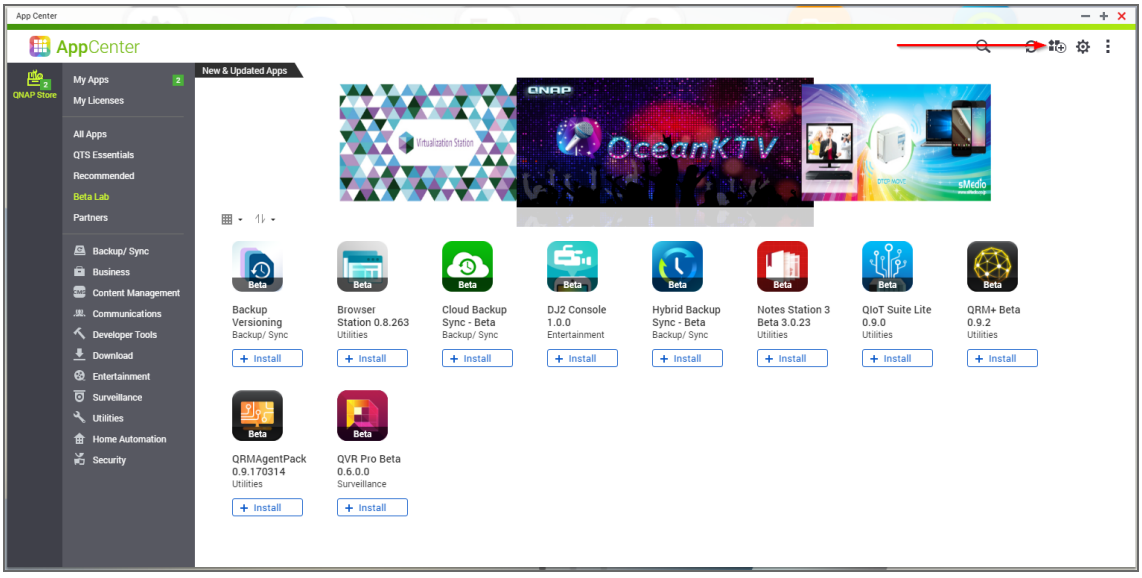
1. Download the update package from here.

> **Note**
>
> Installing updates of NAKIVO Backup & Replication on your NAS-based backup appliance requires accurate knowledge of the hardware system information.
>
> Refer to the following page to learn how to identify your QNAP device CPU model: How to Identify the CPU Architecture of a QNAP NAS Processor.

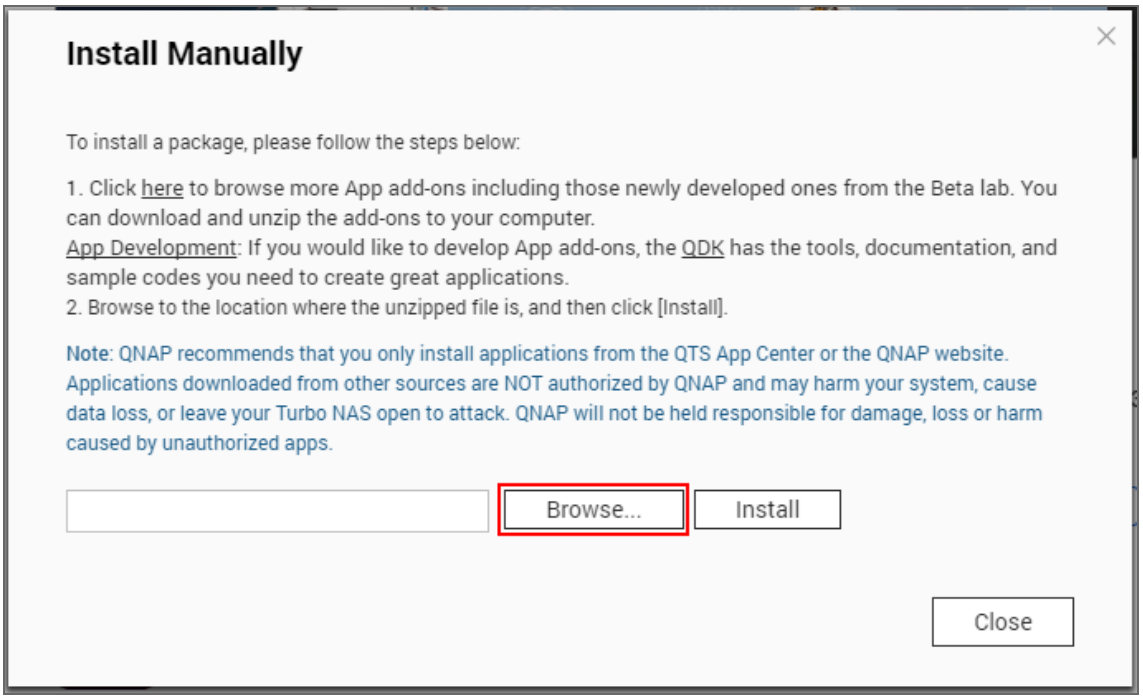
2. Open the QNAP Desktop in your browser by entering the IP address of your QNAP NAS.
3. Go to **App Center**.




4. Click the **Install Manually** icon.




5. Click **Browse**. In the window appears, locate the installer (`.qpkg` file) on your computer.




6. Click **Install**.
7. Wait until the update process is finished.

# Updating on ASUSTOR NAS


- Updating on ASUSTOR NAS Manually
- Updating on ASUSTOR NAS via App Central


## Updating on ASUSTOR NAS Manually

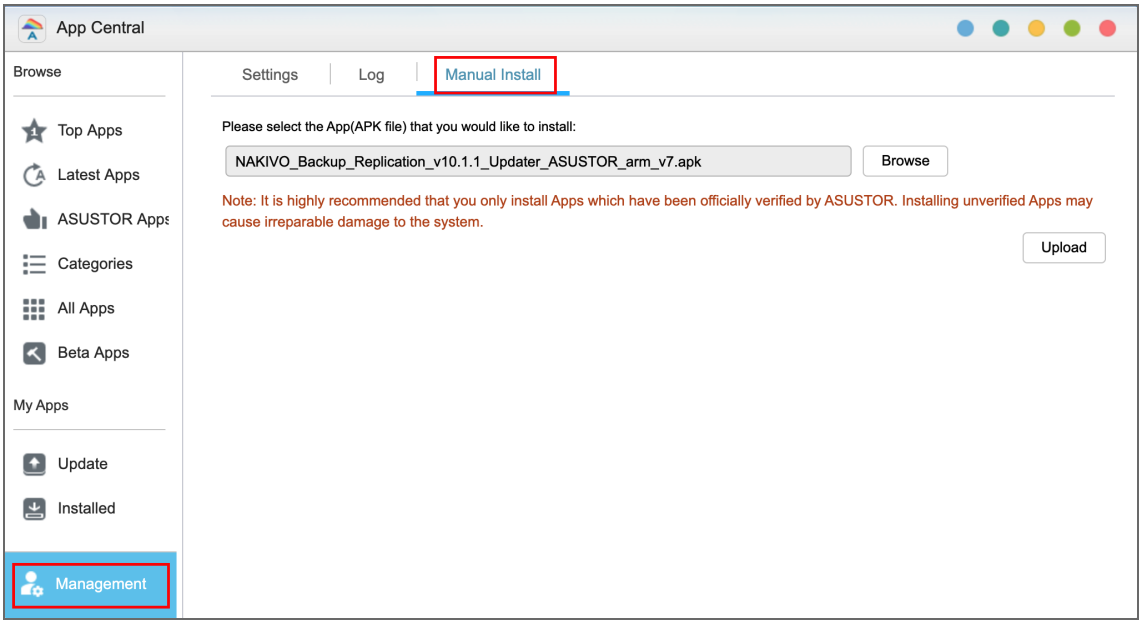
Prior to updating NAKIVO Backup & Replication on ASUSTOR NAS manually, make sure the following requirements are met:

- You have access to the ASUSTOR NAS.
- NAKIVO Backup & Replication installer is available for your ASUSTOR NAS.

> **Note**
>
> Installing updates of NAKIVO Backup & Replication on your NAS-based backup appliance requires accurate knowledge of the hardware system information.
>
> Refer to the following page to learn how to identify your ASUSTOR device CPU model: How to Identify the CPU Architecture of an ASUSTOR NAS Processor.

Follow the steps below to update NAKIVO Backup & Replication on ASUSTOR NAS manually:

1. Make sure that no jobs or repository maintenance tasks are running in the product.
2. Open the **App Central** from the ASUSTOR NAS **Desktop**.
3. Click **Management** in the bottom left corner and click **Manual Install**.
4. The **Manual Install** pane opens to the right of the **App Central**. Click **Browse**.
5. The **Open** dialog box opens. Locate your copy of NAKIVO Backup & Replication installer for ASUSTOR NAS and click the **Open** button.
6. The **Open** dialog closes, and the **Upload** button becomes enabled. Click the **Upload** button.
7. When the upload finishes, the **About This App** dialog opens. If you are sure the requirements are met, click the **Next** button.
8. The **About This App** dialog opens a message asking you to review the summary of the NAKIVO Backup & Replication update. Select the checkbox **I understand the risks associated with installing unverified Apps** and click **Install**.
9. The **About This App** dialog closes, and the **Installed** pane of the **App Central** opens.
10. Wait until the update of NAKIVO Backup & Replication is complete.




# Updating on ASUSTOR NAS via App Central

Follow the steps below to update NAKIVO Backup & Replication on ASUSTOR NAS via App Central:

1. Open the **App Central** from the ASUSTOR NAS **Desktop**.
2. In the **Browse** menu to the left, click **All Apps**. The list of applications available in **the App Central** opens in the right pane.
3. In the search box in the upper right corner of the pane, enter "Nakivo". Installations of the NAKIVO Backup & Replication application that are available at App Central are now displayed.
4. Click the **Update** button below the required NAKIVO Backup & Replication application to start uploading the update.
5. When the update is uploaded successfully, the **About This App** dialog opens. Click the **Update** button if you are sure that all the requirements are met.
6. The **About This App** dialog closes, and the **Installed** pane of the **App Central** opens. Wait until the update of the NAKIVO Backup & Replication is completed.

# Updating on NETGEAR ReadyNAS


- Updating on NETGEAR ReadyNAS Manually
- Updating on NETGEAR ReadyNAS via Available Apps

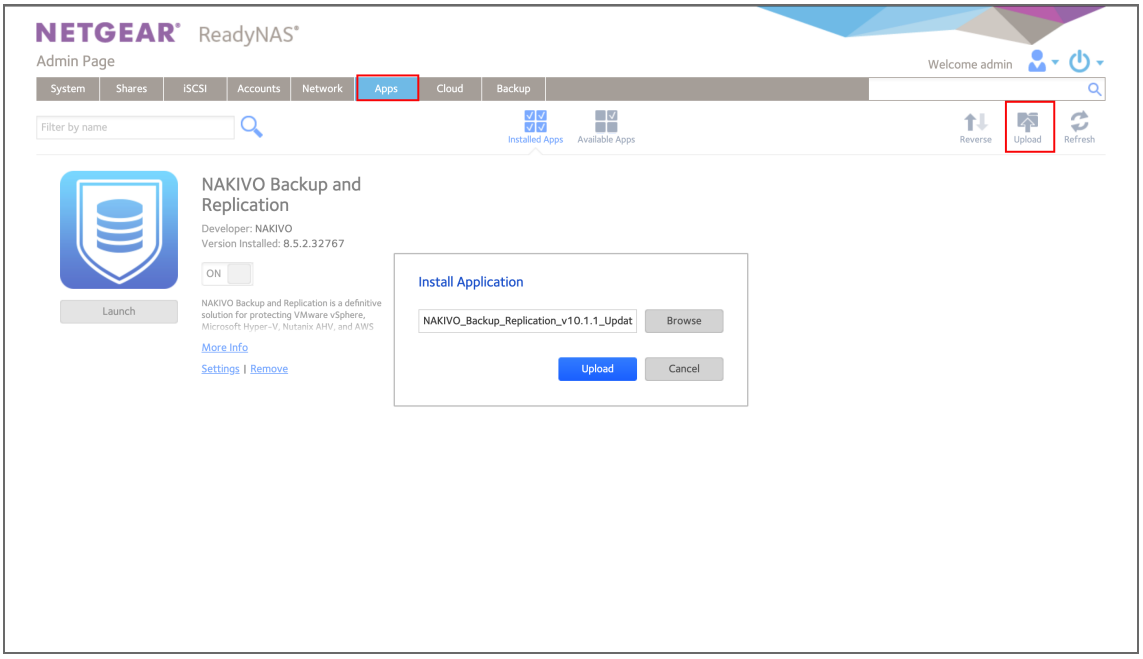

## Updating on NETGEAR ReadyNAS Manually

Prior to updating NAKIVO Backup & Replication on NETGEAR ReadyNAS manually, make sure the following requirements have been met:

- You have access to the NETGEAR ReadyNAS.
- NAKIVO Backup & Replication update is available for your NETGEAR ReadyNAS.

Follow the steps below to update NAKIVO Backup & Replication on NETGEAR ReadyNAS manually:

1. Make sure that no jobs or repository maintenance tasks are running in the product.
2. Open the NETGEAR ReadyNAS **Admin Page** in your browser by entering the IP address of your NAS.
3. Go to **Apps** and click **Upload**.
4. The **Install Application** dialog box opens. Click **Browse**.
5. In the dialog box that opens, locate the downloaded installer (.deb file) and then click **Upload**.
6. Wait until the update is completed.



# Updating on NETGEAR ReadyNAS via Available Apps

**Important**

This method is not possible for ReadyNAS OS 6.10.10. To use this method of updating NAKIVO Backup & Replication, you need to use or downgrade to ReadyNAS OS 6.10.9.

Prior to updating NAKIVO Backup & Replication on NETGEAR ReadyNAS via Available Apps, make sure that you have access to NETGEAR ReadyNAS.

Follow the steps below to update NAKIVO Backup & Replication on NETGEAR ReadyNAS via Available Apps:

1. Open the NETGEAR ReadyNAS **Admin Page** in your browser by entering the IP address of your NAS.
2. Go to **Apps** > **Available Apps**.
3. Find **NAKIVO Backup & Replication** in the list of available applications. Alternatively, enter NAKIVO to the filtering box in the upper left corner of the **Admin Page**.
4. If a new version of NAKIVO Backup & Replication is available in the NETGEAR **Available Apps**, the **Update** button will be available below the application item. Click the **Update** button.
5. Wait until the update is complete.

# Updating on Generic ARM-Based Devices

If auto updating of NAKIVO Backup & Replication is not supported, follow the steps below to update the product on a Generic ARM-based device manually:

1. Download the latest Generic ARM-based NAS updater from here.

   **Note**

   Installing updates of NAKIVO Backup & Replication on your ARM-based backup appliance requires accurate knowledge of the hardware system information.

   Refer to the following page to learn how to identify your ARM-based device CPU model: How to Identify the CPU Architecture of a Generic ARM-Based Device.

2. Upload the updater to the machine on which the **Director** is installed.

   **Important**

   Make sure you are using the *binary transfer mode* when uploading the updater to the machine with a Linux OS. For example:

   - Upload the installer from a Windows-based machine
   - Upload the product from a Linux-based machine: run the following command: `wget 'server_ip/shared/NAKIVO_Backup_Replication_vX.X.X_Updater.sh'`

3. Log in to the Generic ARM-based NAS machine and allow the execution of the updater file. For example: `chmod +x NAKIVO_Backup_Replication_vX.X.X_Updater.sh`
4. Make sure that no jobs or repository maintenance tasks are running in the product.
   If NAKIVO Backup & Replication is installed on a VM, create a snapshot of the VM prior to updating the product.
5. Execute the updater file with root privileges. For example: `sudo ./NAKIVO_Backup_Replication_vX.X.X_Updater.sh`
6. Review the license agreement (press **Space** to go to the next page of the agreement). If you agree to the terms of the license agreement, press “Y” and then press **Enter**.
7. Press the “Y” key and then press **Enter** to confirm that you wish to stop the services and begin the update process.
8. Update all machines on which you have additionally deployed a **Transporter**.

# Updating on TrueNAS

Prerequisites:

- You are logged in to the TrueNAS system with the TrueNAS GUI.
- The **Shell** button is enabled in the interface.

Follow the steps below to update NAKIVO Backup & Replication on your TrueNAS system:

1. Make sure that no jobs or repository maintenance tasks are running in the product.
2. Navigate to the **Jails** page of the TrueNAS GUI and click the jail of the NAKIVO Backup & Replication plugin to select it.
3. Click the **Shell** button to open a web shell.
4. In the web shell prompt, download the latest Virtual Appliance and Linux updater from the NAKIVO Backup & Replication update page with the `curl` command. For example:
   `curl -O https://<download_server>/<path>/NAKIVO_Backup_Replication_vX.X.X.XXXXX_Updater.sh`
5. Change the updater file permission with the `chmod` command:
   `chmod +x NAKIVO_Backup_&_Replication_vX.X.X.XXXXX_Updater.sh`
6. Run the updater in silent mode:
   `./NAKIVO_Backup_&_Replication_vX.X.X.XXXXX_Updater.sh -y -u --eula-accept`

# Uninstalling NAKIVO Backup & Replication


- Uninstalling on Windows
- Uninstalling on Linux or Generic ARM-based NAS
  - Uninstalling Director and Onboard transporter on Linux or Generic ARM-Based NAS
  - Uninstalling Transporter on Linux or Generic ARM-Based NAS
- Uninstalling on Synology NAS
- Uninstalling on Western Digital NAS
- Uninstalling on QNAP NAS
- Uninstalling on ASUSTOR NAS
- Uninstalling NETGEAR ReadyNAS
- Terminating on Amazon EC2
- Uninstalling on TrueNAS


## Uninstalling on Windows

To uninstall NAKIVO Backup & Replication, run the uninstaller:

1. Go to **Start** -> **Control Panel** and run **Programs and Features**.
2. Select **NAKIVO Backup & Replication** and click **Uninstall**.
3. In the **NAKIVO Backup & Replication Uninstallation** wizard, click **Uninstall**.
4. Click **Close** when the uninstallation process is completed.

## Uninstalling on Linux or Generic ARM-based NAS

Refer to the sections below to learn how to uninstall NAKIVO Backup & Replication on a Linux OS or a generic ARM-based NAS.

### Uninstalling Director and Onboard Transporter on Linux or Generic ARM-based NAS

To uninstall the **Director** and Onboard **Transporter**, which is installed with the **Director** by default, follow the steps below:

1. Run the "uninstall" script which is located in the **Director** folder inside the product installation folder. If the product is installed in the default location, run: /opt/nakivo/director/uninstall
2. Enter "U" and then press **Enter** to confirm uninstalling the application.

### Uninstalling Transporter on Linux or Generic ARM-based NAS

To uninstall the Transporter, follow the steps below:

1. Run the "uninstall" script which is located in the transporter folder inside the product installation folder. If the product is installed in the default location, run: /opt/nakivo/transporter/uninstall
2. Enter "U" and then press **Enter** to confirm uninstalling the application.

# Uninstalling on Synology NAS

Follow the steps below to uninstall NAKIVO Backup & Replication on a Synology NAS:

1. In the Synology NAS management interface, open the **Package Center**.
2. Click NAKIVO Backup & Replication.
3. Choose **Uninstall** from the **Actions** list.
4. Click **OK** in the message box that opens to confirm that you wish to uninstall the application.

When the uninstallation process is completed, NAKIVO Backup & Replication will be removed from the list of installed applications.

# Uninstalling on Western Digital NAS

Follow the steps below to uninstall NAKIVO Backup & Replication on a Western Digital NAS:

1. Open the NAS My Cloud Dashboard and click **Apps**.
2. In the **Installed Apps** list, select NAKIVO Backup & Replication.
3. The NAKIVO Backup & Replication item opens to the right of the installed applications list. Click the **Uninstall** button.
4. The **Uninstall NAKIVO Backup and Replication** dialog opens. Click **OK** to confirm that you wish to uninstall the application and delete all application data and settings.
5. The **Updating** progress bar opens. Wait until the uninstallation completes.

When the uninstallation process is completed, NAKIVO Backup & Replication will be removed from the list of installed applications.

# Uninstalling on QNAP NAS

Follow the steps below to uninstall NAKIVO Backup & Replication on a QNAP NAS:

1. Open the QNAP NAS Desktop and click **App Center**.
2. The **App Center** dialog opens. In the **My Apps** list, locate the NAKIVO Backup & Replication application and open the list of applicable actions by clicking the drop-down button.
3. In the list of applicable actions, click **Remove**.
4. In the dialog that opens, click **OK** to confirm removing the application and application-relevant user data.
5. Wait until the uninstallation is complete.

When the uninstallation process is completed, NAKIVO Backup & Replication will be removed from the list of installed applications.

# Uninstalling on ASUSTOR NAS

Follow the steps below to uninstall NAKIVO Backup & Replication on a QNAP NAS:

1. Open the ASUSTOR NAS Desktop and click **App Central**.
2. In the list of installed applications, locate NAKIVO Backup & Replication, select it and then click the **Remove** button.
3. In the dialog that opens, click **OK** to confirm that you wish to remove the application.
4. The **Removing** progress bar opens. Wait until the uninstallation is completed.

When the uninstallation process is completed, NAKIVO Backup & Replication will be removed from the list of installed applications.

# Uninstalling on NETGEAR ReadyNAS

Follow the steps below to uninstall NAKIVO Backup & Replication on NETGEAR ReadyNAS:

1. Open the NETGEAR ReadyNAS **Admin Page** and go to **Apps** > **Installed Apps**.
2. Locate **NAKIVO Backup & Replication** in the list of available applications. Alternatively, enter NAKIVO to the filtering box in the upper left corner of the **Admin Page**.
3. Click the **Remove** button below the application item.
4. The **Confirm Deletion** dialog box opens. Click **Yes** to confirm that you wish to uninstall NAKIVO Backup & Replication on NETGEAR ReadyNAS.
5. Wait until the uninstallation is completed.

When the uninstallation process is completed, NAKIVO Backup & Replication will be removed from the list of installed applications.

# Terminating on Amazon EC2

Follow the steps below to terminate NAKIVO Backup & Replication that is launched as an Amazon EC2 instance:

1. Open AWS Management Console and go to **EC2 Dashboard**.
2. In the **Instances** menu, click **Instances**.
3. In the list of instances, locate the necessary NAKIVO Backup & Replication instance and select it.
4. In the **Actions** menu, go to **Instance State** and click **Terminate**.
5. In the **Terminate Instances** dialog, click **Yes, Terminate** to confirm that you wish to terminate your instance of NAKIVO Backup & Replication.
6. Wait until the instance is terminated.

In about 60 minutes, the terminated NAKIVO Backup & Replication instance will be removed from the list of Amazon EC2 instances.

# Uninstalling on TrueNAS

Uninstalling a plugin deletes the associated TrueNAS jail because it is no longer required. Before uninstalling NAKIVO Backup & Replication, make sure that there is no data or configuration in the jail that needs to be saved.

Follow the steps below to uninstall NAKIVO Backup & Replication on a TrueNAS:

1. Log in to the TrueNAS system using the TrueNAS GUI.
2. In the left pane of the TrueNAS GUI, click **Plugins** -> **Installed**.
3. A list of installed plugins opens. For the desired NAKIVO plugin, click the **Options** button and then **Delete**.
4. The **Delete** dialog opens asking to confirm the operation. Click **Delete**.

When the uninstall process is completed, NAKIVO Backup & Replication will be removed from the list of installed plugins.

# Getting Started

When deployed, NAKIVO Backup & Replication is ready for use. The topics below will provide you with information on how to start working with the application.


- “Logging in to NAKIVO Backup & Replication” on page 334
- “First Steps with NAKIVO Backup & Replication” on page 341
- “Web Interface Components” on page 344
- “Managing Jobs and Activities” on page 354

# Logging in to NAKIVO Backup & Replication


- Getting to the Login Page
- Creating a User Account
- Resetting Password
- Default Password in Amazon EC2
- Passing Verification


## Getting to the Login Page

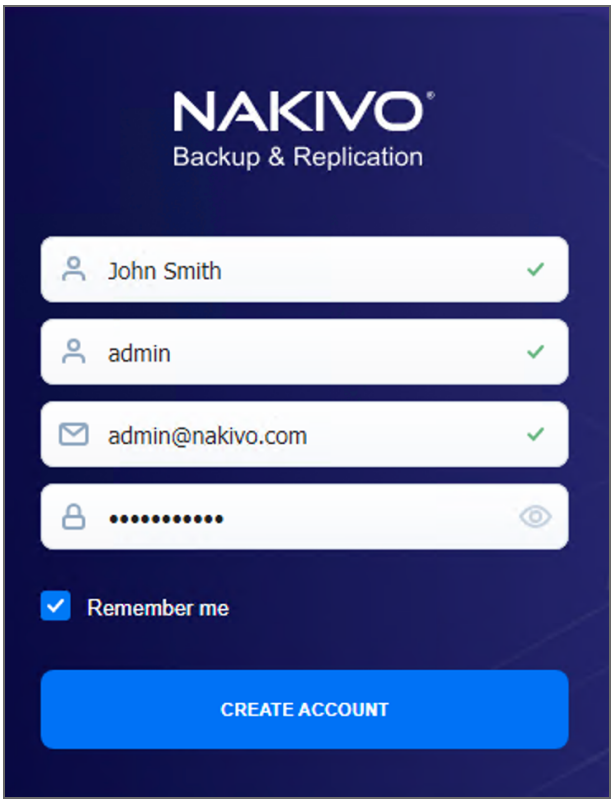
To open the NAKIVO Backup & Replication login page, enter the machine's IP address or DNS name in your web browser, followed by the HTTPS port (default port is 4443).

> **Note**
> If you selected a custom HTTPS port during installation, replace *4443* with the custom value.

## Creating a User Account

When you open the NAKIVO Backup & Replication login page for the first time, you are prompted to create a new user account. This user account is the admin account to be used to access your instance of NAKIVO Backup & Replication. Fill out the fields in the form:

1. **Name**: Provide your real name.
2. **Username**: Enter an admin username to log in to NAKIVO Backup & Replication.
3. **Email**: Provide an email.
4. **Password**: Enter a password.
5. Optionally, you can select **Remember me** to save your credentials.
6. Click **CREATE ACCOUNT**.




**Note**
If NAKIVO Backup & Replication is deployed in an Amazon EC2 instance, you will first be prompted to enter the Amazon EC2 instance ID.

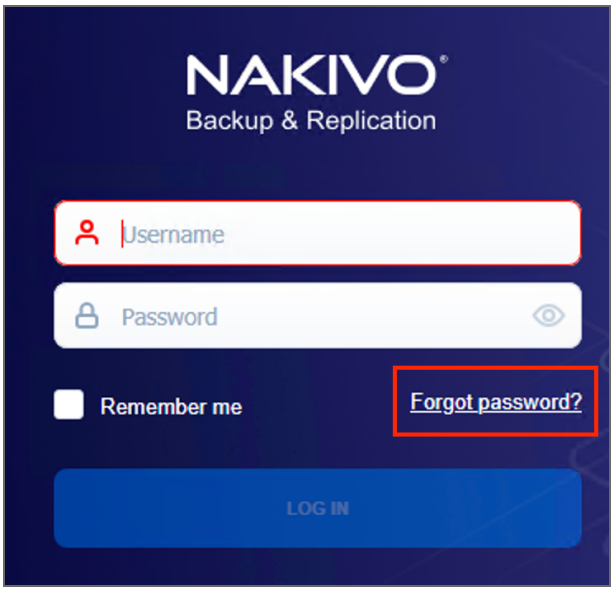
NAKIVO Backup & Replication opens in your browser displaying the configuration wizard. Refer to First Steps with NAKIVO Backup & Replication to learn how to start using NAKIVO Backup & Replication.

To log out, click **Logout** in the bottom left corner.

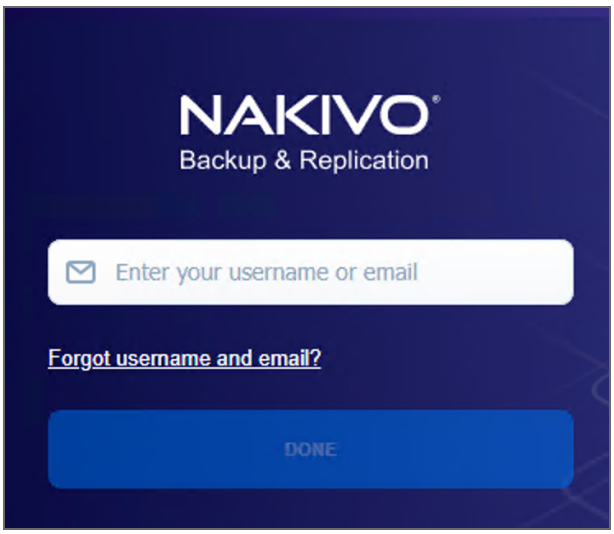
# Resetting Password

If you forget the password used to log in to NAKIVO Backup & Replication, you can restore it by following the steps below:

1. Go to NAKIVO Backup & Replication login page.
2. Click the **Forgot password** link.




3. Do one of the following:
   - If you have set up email settings in NAKIVO Backup & Replication, enter your email address and click **Done**.

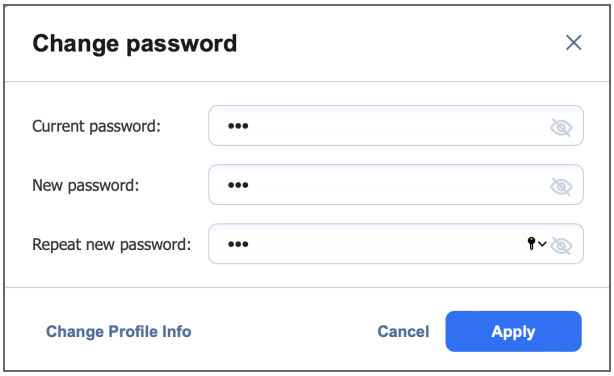



A temporary password, which is a security string, is sent to your inbox. Enter this password the next time you log in to your NAKIVO Backup & Replication instance. Once you are logged in, it's recommended that you change the temporary password for your user account. To change the temporary password:

a. Click **Logout** in the bottom left corner.
b. Select **Profile**.
c. Click **Change password**.
d. In the dialog box that opens, fill out the following fields:
   - **Current password**: Enter the temporary password that you received to your inbox.
   - **New password**: Enter a new password.

- **Repeat new password**: Enter the new password again.

e. Click **Change**.




You can also change your temporary password in **Settings**>**General**>**Users and Roles**

- If you have not set up email settings in NAKIVO Backup & Replication:
  a. Enter your username and click **Done**.
  b. Go to the product installation folder and locate the “forgot_password.txt” file.

> **Important**
> For security reasons, only a root user (Linux) or a member of the Administrators group (Windows) is allowed to access the installation folder and the "forgot_password.txt" file.

  c. Paste the security string from the file in the appropriate field.
  d. Click **Done**.

**Notes**

- If you are using a Virtual Appliance (VA), go to the VA console, then go to the command line and enter: `cat /opt/nakivo/director/forgot_password.txt` The security string will be displayed on the screen. You can copy and paste it into the web interface.
- For Windows OS, you can find forgot_password.txt file in the following default installation folder: *C:\Program Files\NAKIVO Backup & Replication*. Please note that if you changed the installation path, the file will be located in the folder you specified.
- If you are using a NAS, open an SSH connection to your device and read the forgot_password.txt file in the following folders:
  - For ASUSTOR NAS: `/usr/local/AppCentral/NBR`
  - For NETGEAR NAS: `/apps/nbr`
  - For QNAP NAS: `/share/CACHEDEV1_DATA/.qpkg/NBR`
  - For Synology NAS: `/volume1/@appstore/NBR`
  - For Western Digital NAS: `/mnt/HD/HD_a2/Nas_Prog/NBR`
- To learn how to open an `SSH` connection to your NAS device and read text files, refer to the NAS vendor documentation.

# Default Password in Amazon EC2

If you have deployed NAKIVO Backup & Replication as an Amazon machine image in Amazon EC2, use the following default credentials to log in:

- **Username**: admin
- **Password**: The password is the ID of the NAKIVO Backup & Replication instance in Amazon EC2.

# Passing Verification

If two-factor authentication was configured, verification needs to be passed after entering the credentials to access your NAKIVO Backup & Replication instance. This can be done in one of the following ways:

- Google Authenticator code from the mobile app
- A code sent to the specified email address
- One of the single-use backup codes

If Two-factor authentication was enabled but never configured, it must be configured now. Do the following:

1. Click **Continue**.
2. Optionally, click on the **change your email** link to enter the new email address for the user. Select **Continue** to proceed.

3. Enter the verification code that was sent to the specified email and click **Continue**. Optionally, click **Resend email** in case you did not receive it.
4. Optionally, enter the alternative email address that can be used in case the primary one becomes unavailable, and select **Continue**. Alternatively, select **Skip this step**.
5. If you have entered the alternative email address for the previous step, enter the verification code that was sent to the specified email, and click **Continue**. Optionally, click **Resend email** in case you did not receive it.
6. Follow instructions on screen to download and install Google Authenticator, and click **Continue**.
7. Add your NAKIVO Backup & Replication user account to Google Authenticator. Use one of the following methods:
    - Select **Scan QR Code** option and scan the QR code in the popup window.
    - Select **Enter a Code** option and follow the instructions to enter the shown code into the Google Authenticator app.
8. Enter the 6-digit verification code from Google Authenticator into the field. Note that the verification code is time-based. Click **Continue** to proceed.
9. A pairing key is displayed which can be used to add multiple devices to your account.

    **Important**
    It is highly recommended that you save the pairing key or write it down.

    You have the following options:
    - Optionally, click on the **Copy the key** link to copy your key and save it for future use.
    - Optionally, click on the **Download pairing information** link to download and save instructions on how to use the pairing key.
    - Click **Continue** when you're done.
10. Four backup codes are displayed on the next page. These one-time codes can be used to log in when you are unable to provide a verification code. Click on the **Save as PDF** link to download and save these codes in PDF format or write them down. Click **Continue**.
11. Enter one of the backup codes in the next popup window to confirm that you have saved them, and click **Finish**.

## Google Authenticator Verification

If you have selected the **Google Authenticator** verification method on the **Managing Two-Factor Authentication** page, do one of the following:

- Enter the verification code from Google Authenticator into the field, and click **Proceed**.
- Enter one of the one-time backup codes.
- Click **More verification options** to use email verification.

## Email Verification

If you have selected the **Email** verification method on the **Managing Two-Factor Authentication** page, do one of the following:

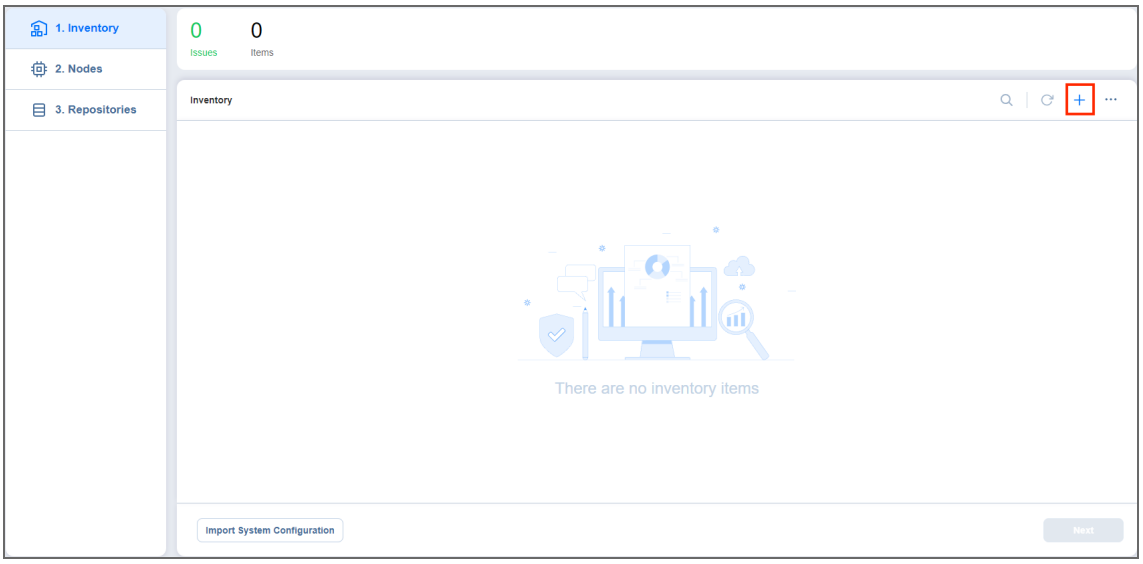
- Select one of the email addresses verified previously, and click **SEND VERIFICATION CODE**. Then click **OK**.
- Enter one of the one-time backup codes.
- Alternatively, click **More verification options** to choose a different email for verification.

# First Steps with NAKIVO Backup & Replication

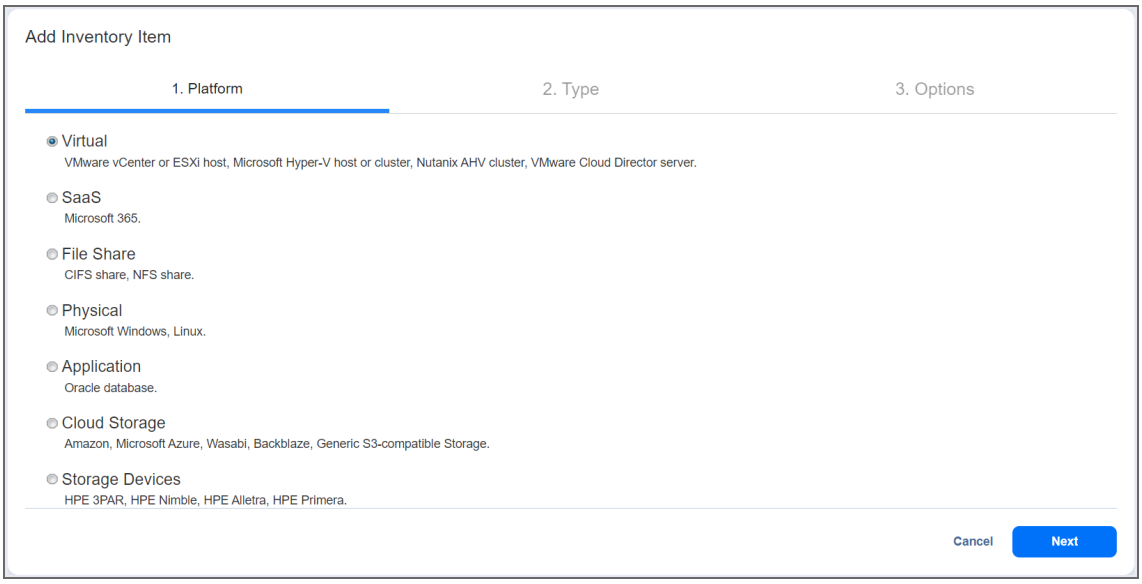
When you log in to NAKIVO Backup & Replication for the first time, the initial configuration wizard opens.

Proceed as follows:

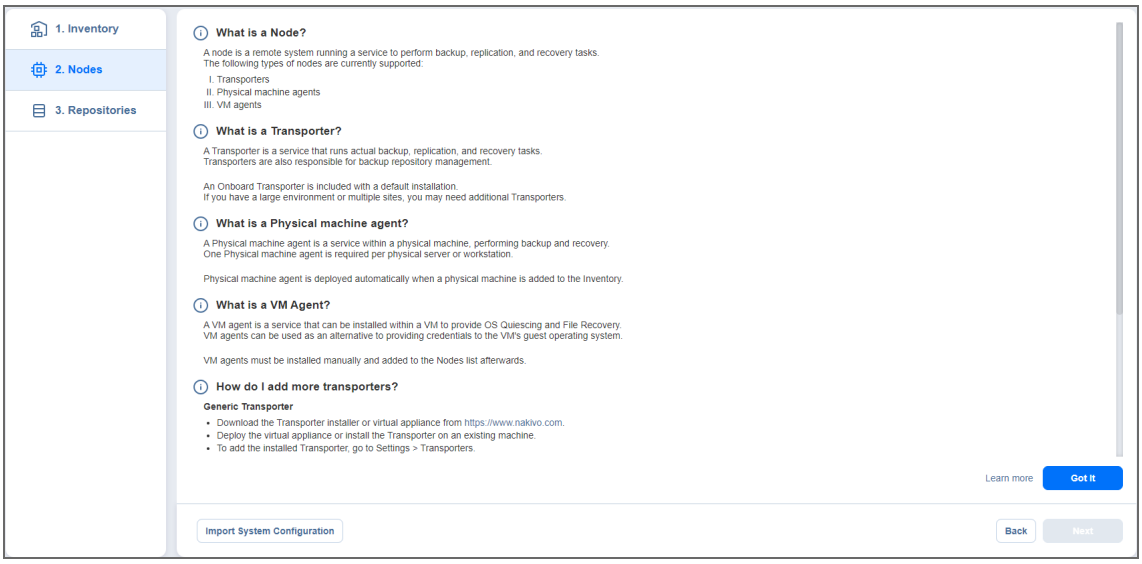
1. On the **Inventory** page of the wizard, click **Add New**.



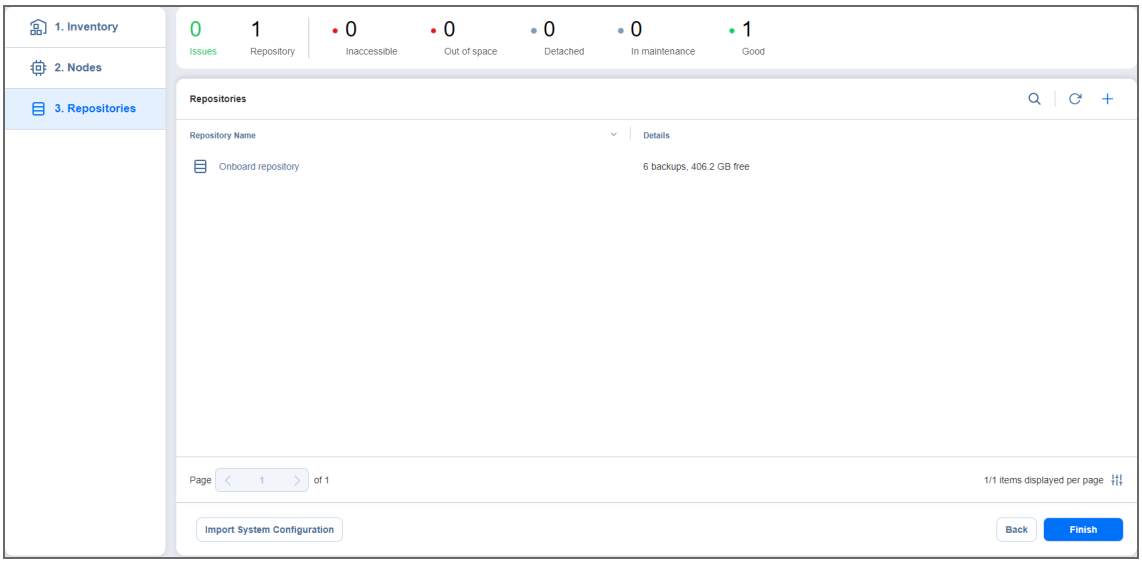

2. Select one of the given options:
    - Virtual
    - SaaS
    - File Share
    - Physical
    - Application
    - Cloud Storage
    - Storage Devices



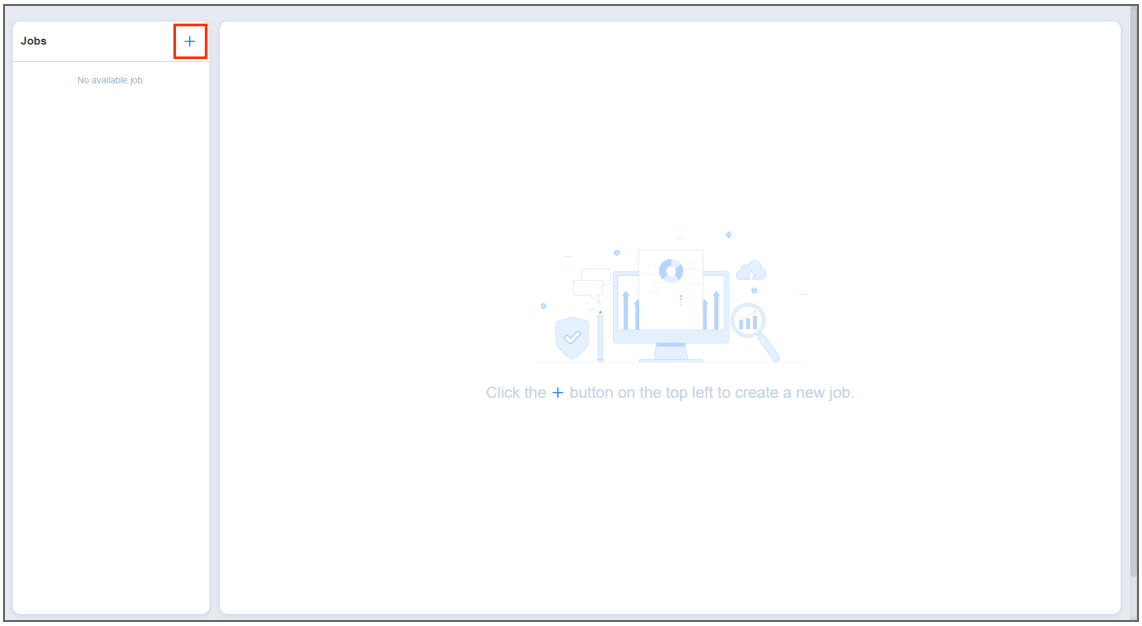

3. Proceed with adding items as described in the Inventory article.
4. On the **Nodes** page of the wizard, you will find information about the **Transporter** component of the NAKIVO Backup & Replication.




5. To deploy a new **Transporter** or add an existing one, click **Got it** and proceed as described in the **Transporters** article.
6. To move to the next page of the wizard, click **Next**.
7. On the **Repositories** page of the wizard, you can add a local or a remote **Backup Repository** to your application by clicking **Add Backup Repository.**




8. Click **Finish**.
9. The **Jobs** menu of the application opens. Proceed by creating backup and replication jobs.




If your current license type is **Free** and the **Trial** license has not yet been applied to the current deployment of NAKIVO Backup & Replication, a dialog box appears. Using this dialog box, you can contact the sales team to change your license type or try the full functionality of the solution for 15 days. If you do not want to upgrade your license type right away, you can do it at any time in the Help menu.

**Note**

If you switch the license type to **Trial**, the product will automatically go back to using your **Free** license after expiration.

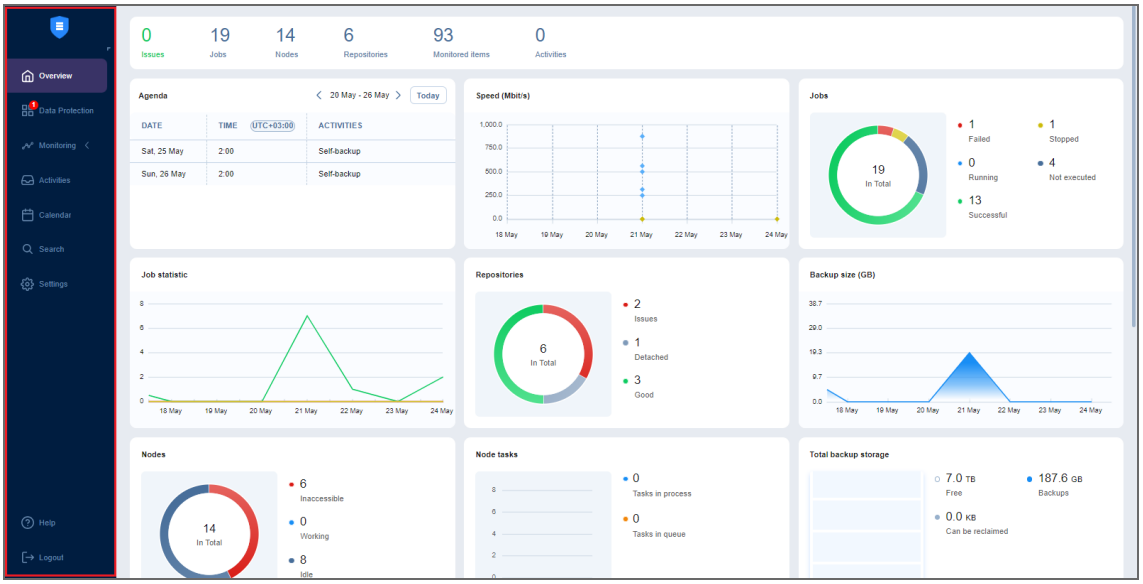
# Web Interface Components

The interface of NAKIVO Backup & Replication consists of the following components:

- Main Menu
- Overview
- Data Protection
- Monitoring
- Activities
- Calendar
- Search
- Settings
- Help Menu
- Online Chat Dialog
- Language Selection
- Special Offers Toolbar
- Tenants Dashboard

## Main Menu

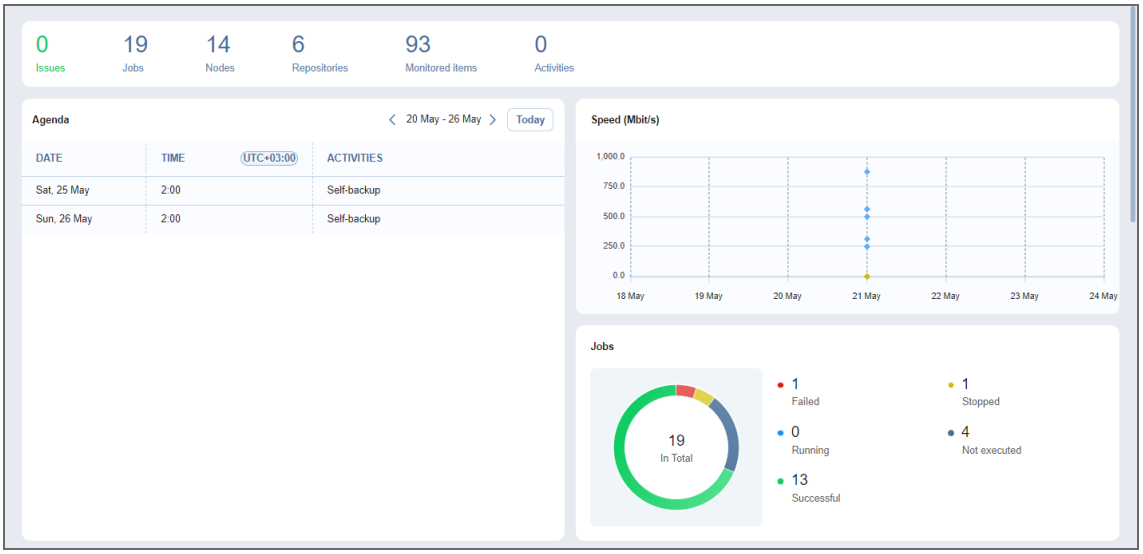
The main menu of NAKIVO Backup & Replication is located on the left side of the product interface. It provides access to the overview dashboard, jobs, activities, calendar, global search, and product settings. It also contains the **Help** menu and **Log Out** button.



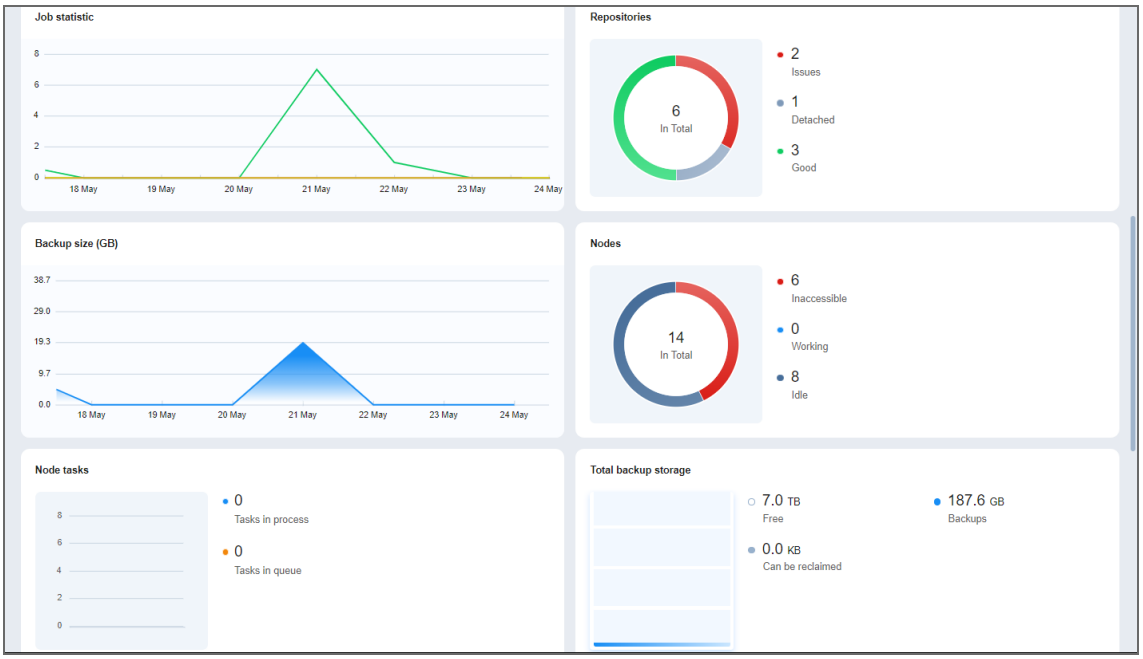
# Overview

The **Overview** page displays the key statistics for your instance of NAKIVO Backup & Replication. The information is displayed in the following widgets:

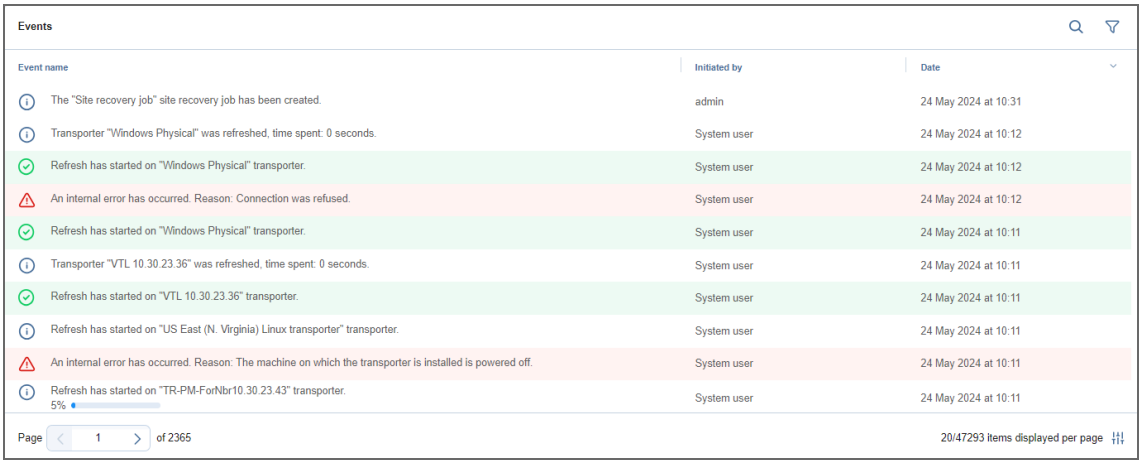
- **Summary bar**: Lists the total number of issues (errors and alarms), jobs, transporters, repositories, monitored items, and running activities.
- **Agenda**: Lists running and scheduled activities for a given week. By default, this widget displays the current week.
- **Speed**: Displays the speed at which raw data has been transferred during successful job runs in the previous seven days.
- **Jobs**: Displays the total number of available jobs and their respective last run statuses.




- **Job statistic**: Shows a graph of the number of successful, stopped, and failed jobs for each day in the previous seven days.
- **Repositories**: Displays the total number of available repositories and their statuses.
- **Backup size**: Displays the total size of backups created for each day in the previous seven days. Note that backups in forever incremental (**Store backups in separate files option** is not selected) Backup Repositories are considered 0KB, and thus are not reflected in the **Backup size** graph.
- **Transporters**: Displays the total number of available transporters that have been added or deployed successfully and their statuses.
- **Transporter tasks**: Displays the total number of tasks being processed or waiting to be processed by all transporters.
- **Total backup storage**: Displays the total amount of storage of all available repositories.




- **Events**: Lists all events, including errors, warnings, and general status information, sorted by date by default. Includes search and filter functions that simplify finding specific events by name, type, or date range.

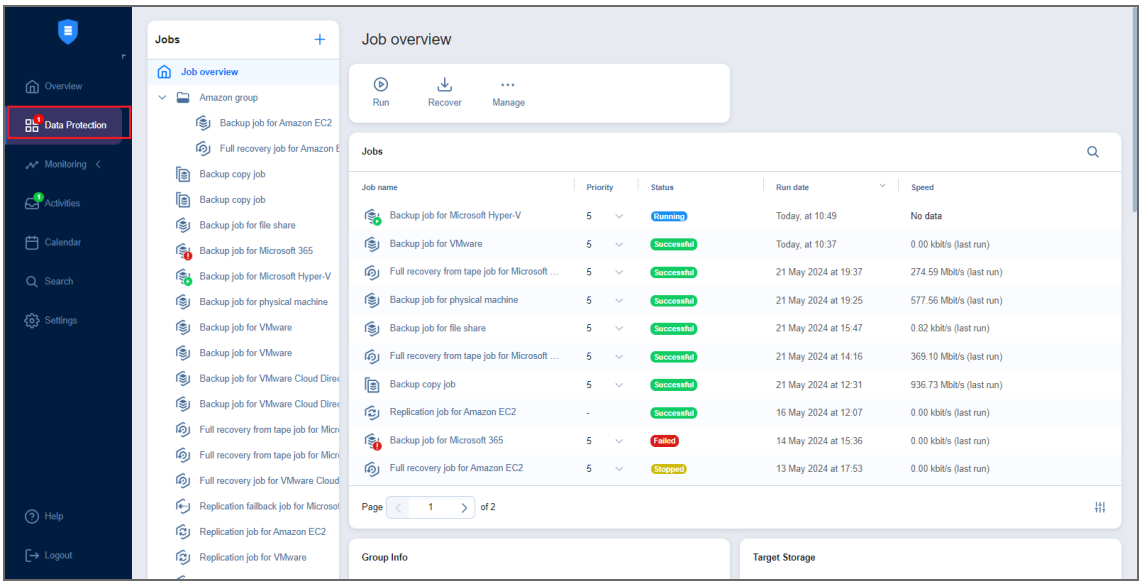



# Data Protection

Using the **Data Protection** page, you can:

- View, run, and stop jobs on demand or on schedule
- Recover files, objects, VMs, and sites
- Manage jobs
- Create backup, backup copy, replication, recovery, and flash boot jobs
- Create job reports

- Create and manage job groups

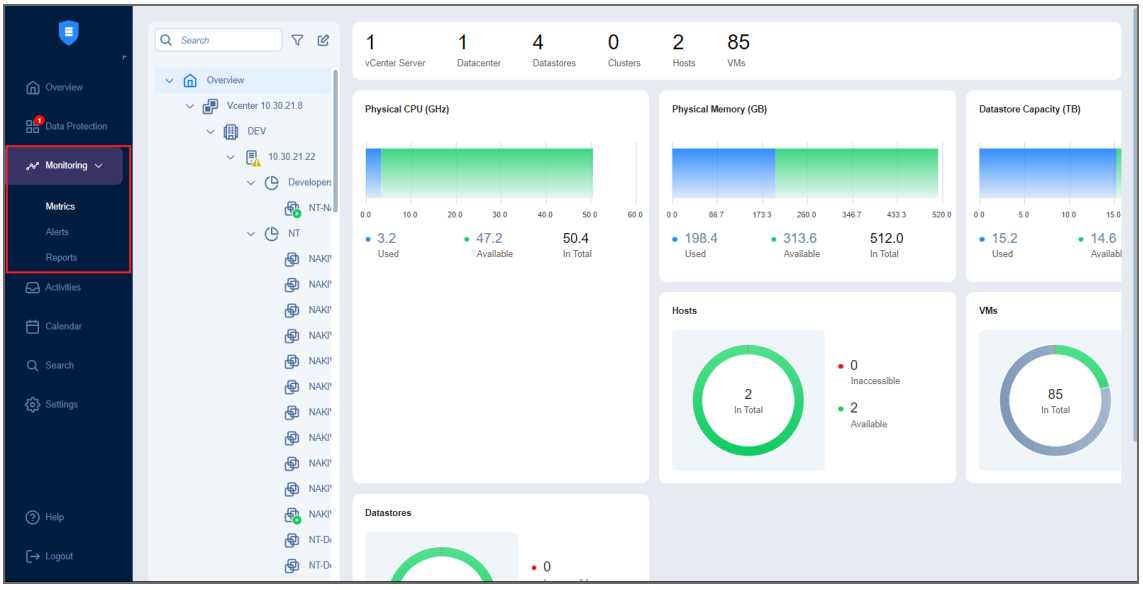



# Monitoring

On the **Monitoring** page, you can check the following metrics (current and historical):

- For VMware VMs: CPU usage, memory usage, and disk usage
- For VMware hosts: CPU usage and memory usage
- For VMware datastores: Disk usage

For additional information, refer to the Monitoring topic.

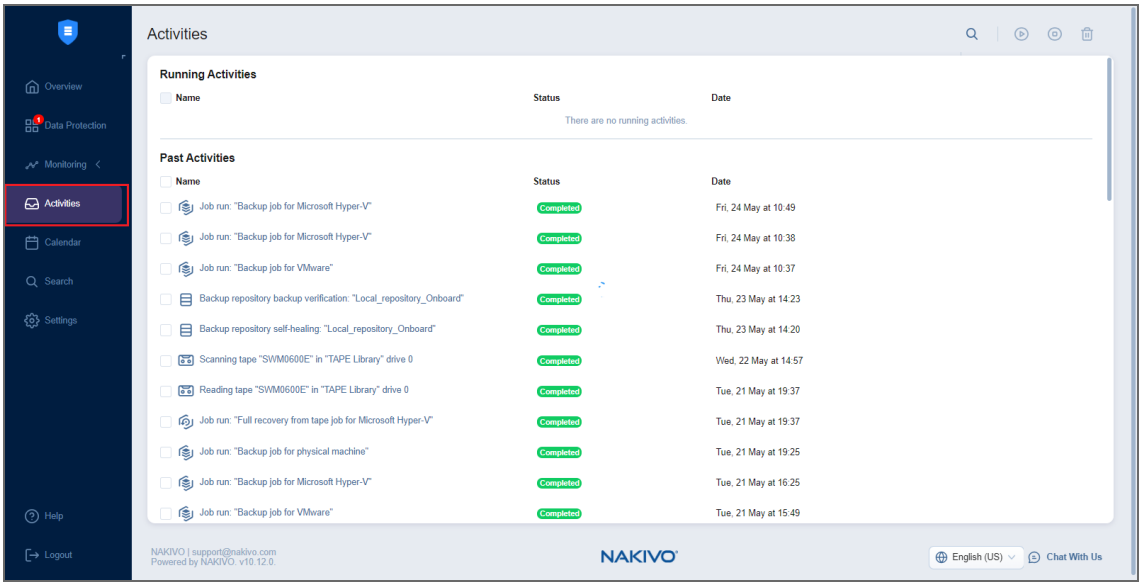



# Activities

The **Activities** page displays a list of all running and past activities, such as:

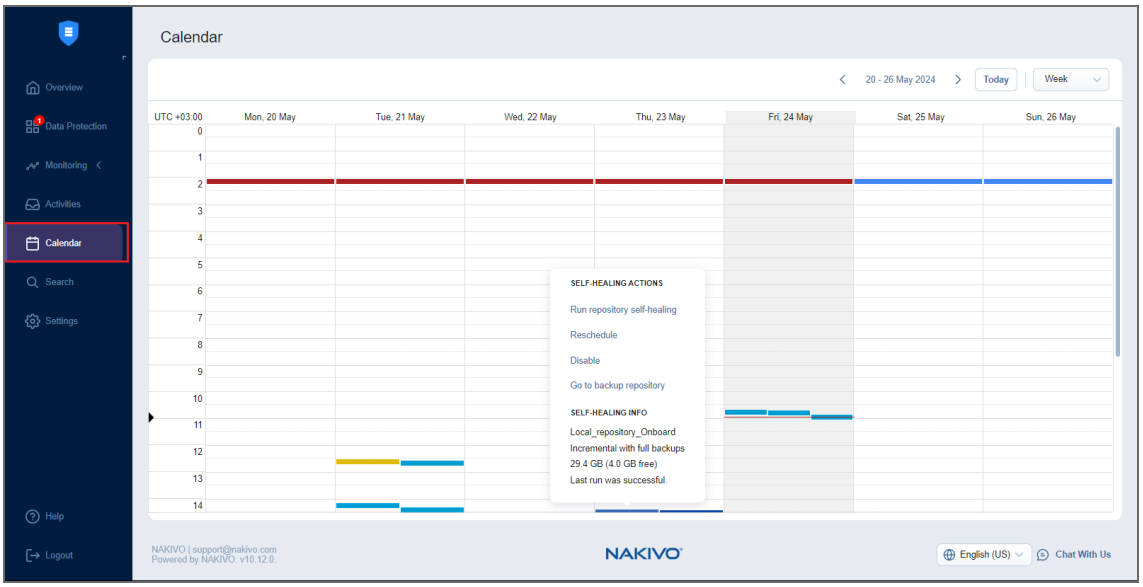
- Job run
- Repository Self-Backup
- File download
- Application object download
- Universal object recovery
- Repository space reclaim
- Repository self-healing
- Repository backup verification
- Backup export
- Tape-specific activities, namely: scanning, erasing, and reading Tape
- Other

For further details and information, refer to “Managing Activities” on page 392.



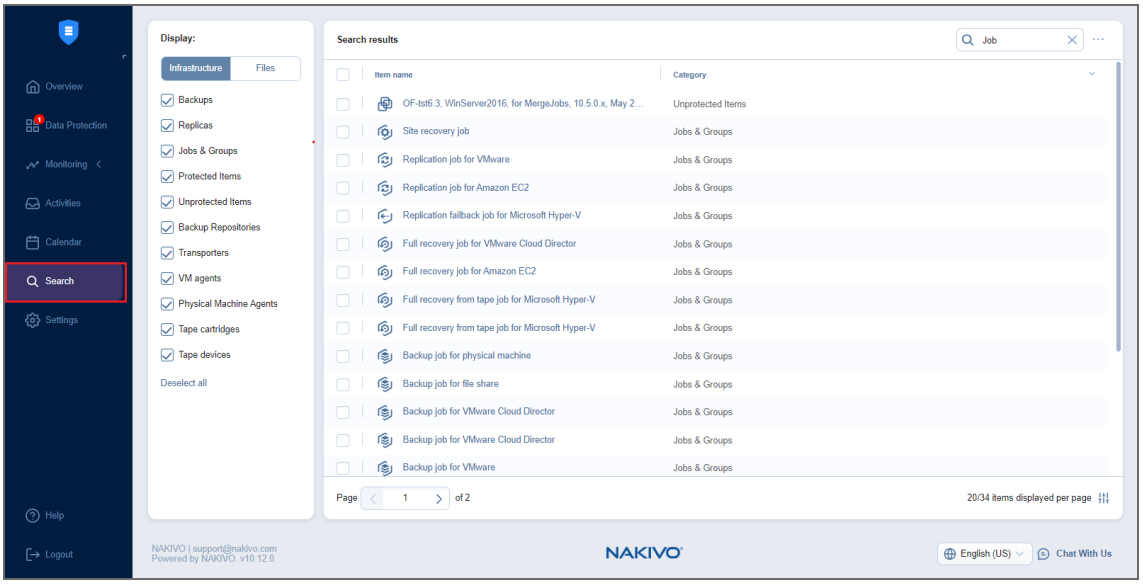

# Calendar

The **Calendar** allows you to schedule jobs and view the history of all job runs in one organized space. For more information, refer to “Using Calendar” on page 396.



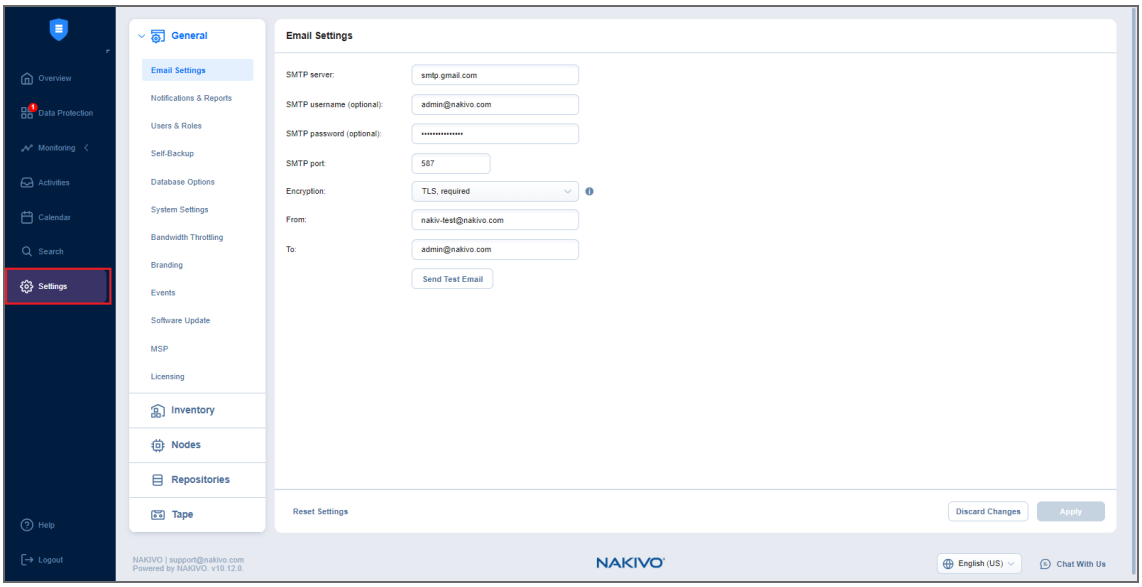

# Search

The **Search** page allows you to search for items within the entire NAKIVO Backup & Replication instance–the Inventory, Transporters, Repositories, tape devices, jobs, backups, replicas, and more. For more details, refer to “Using Global Search” on page 398.



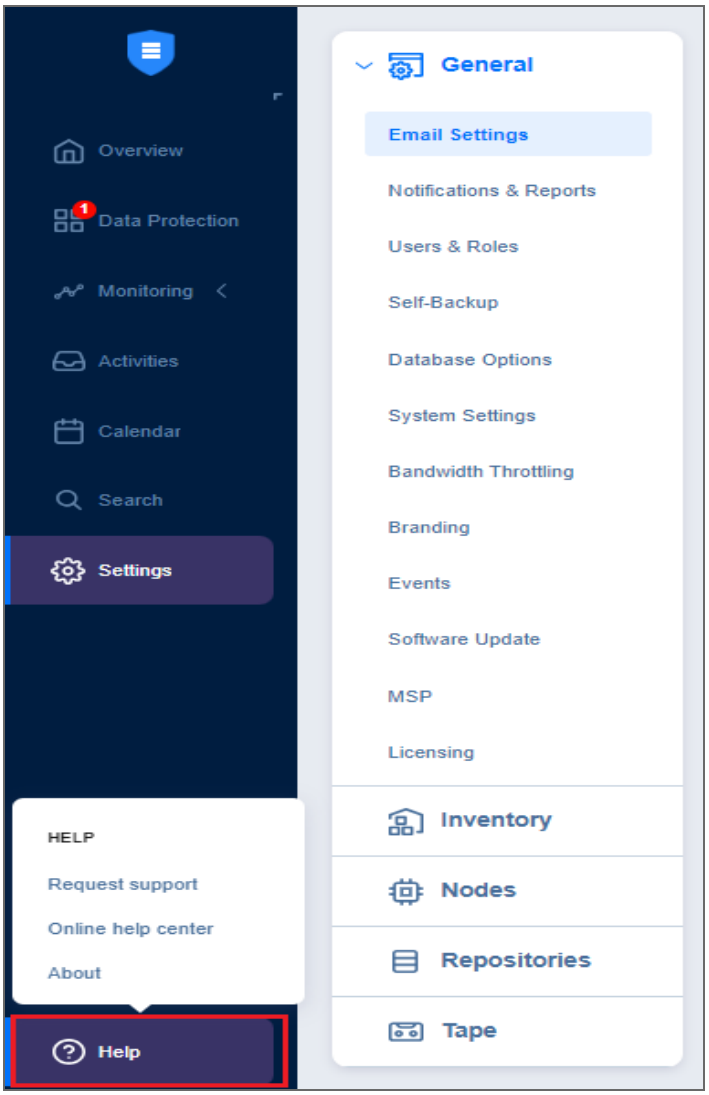

# Settings

On the **Settings** page, you can configure NAKIVO Backup & Replication General, Inventory, Transporters, Repositories, and Tape settings. Refer to “Settings” on page 405 for more detailed information.



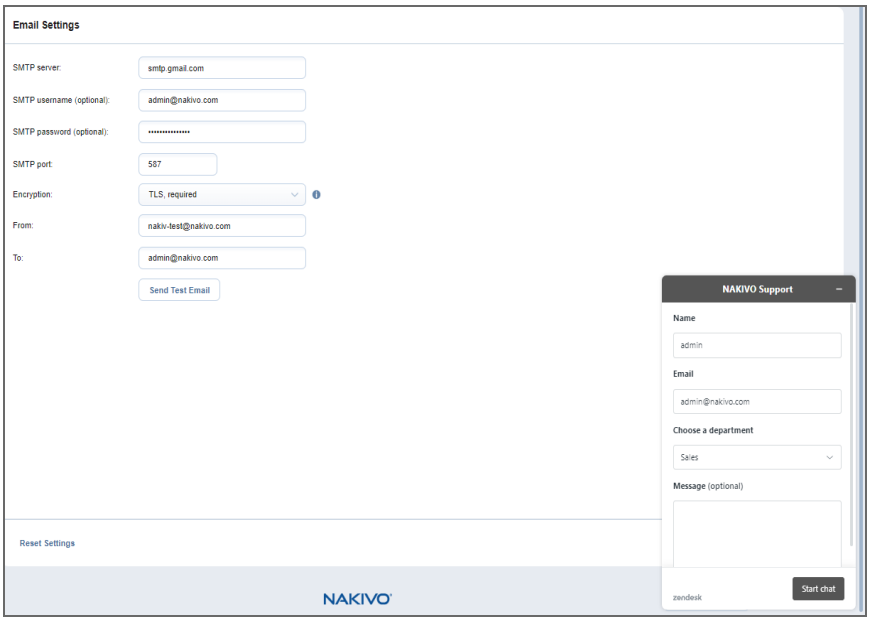

# Help Menu

Use the **Help** menu to request technical support and access the NAKIVO online help center or see information about your product. If you are evaluating NAKIVO Backup & Replication, you may also check the the **How to Buy** section of the **Help** menu to view pricing, request a live demo or quote, find a reseller, or contact Sales. If you are using a Free license, you may also upgrade to a Trial license for 15 days with the **Try full functionality** option.




# Online Chat Dialog

The **NAKIVO Support** online chat is located in the right bottom corner of the application. It enables you to quickly request help from a sales or technical support representative.



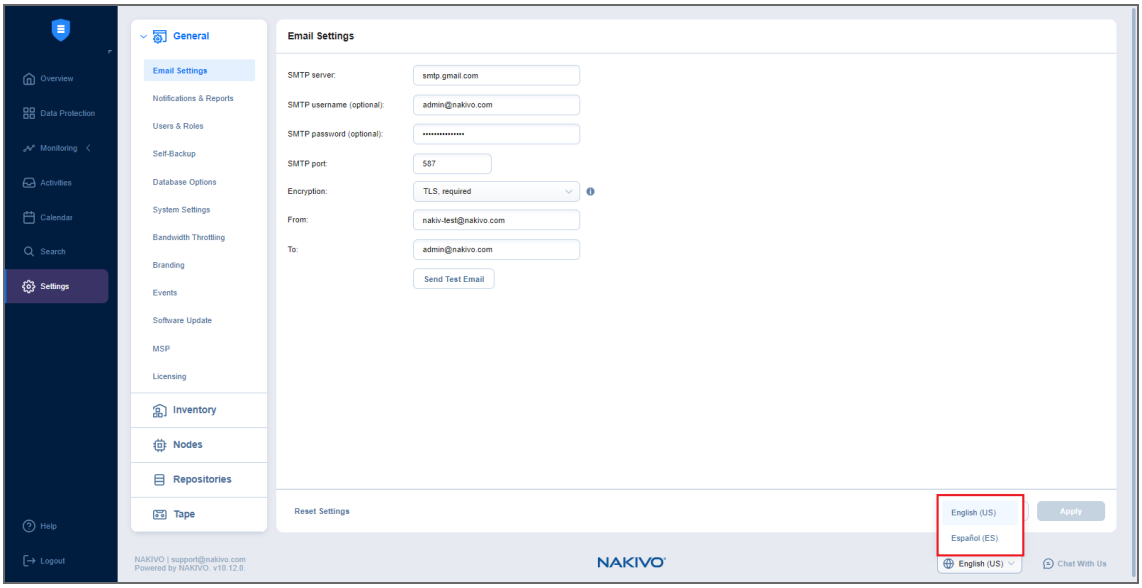

# Language Selection

NAKIVO Backup & Replication supports the following languages:

- Spanish
- French
- German
- Italian
- Polish
- Chinese (simplified)

To select the language for your NAKIVO Backup & Replication instance UI, do the following:

1. Log in to your NAKIVO Backup & Replication instance.
2. In the lower right corner of the NAKIVO Backup & Replication interface, select the needed language from the dropdown menu.




3. The current page reloads automatically with the UI and content switched to the selected language.

**Notes**

- The language switching menu is available to users with the **Administrator** role only; the language change is applied globally (across all users).
- In Multi-tenant deployment, it is available at the Master Admin level only; the language change is applied globally (across all local tenants/users).
- Only users with the **Administrator** role can change the language in a remote tenant.
- You can also change language in the product installer UI, in silent installation interface, and using the command line interface

# Special Offers Toolbar

This element of the interface is located to the left of the NAKIVO Backup & Replication dashboard. The toolbar contains special offers. If you click the button, a dialog opens displaying information about a specific offer. If needed, the **Special Offers** toolbar can be disabled. Refer to "System Settings" on page 455 for details.

# Tenants Dashboard

If you use NAKIVO Backup & Replication in a multi-tenant mode, the **Tenants** dashboard allows you to create, manage, and configure tenants.

# Managing Jobs and Activities

Using NAKIVO Backup & Replication interface, you can manage jobs and tasks. This section covers the following topics:


- "Jobs Dashboard" below
- "Running Jobs on Demand" on page 364
- "Managing Jobs" on page 371
- "Managing Job Policies" on page 382
- "Managing Policy Rules" on page 385
- "Job Alarms and Notifications" on page 390
- "Managing Activities" on page 392
- "Using Calendar" on page 396
- "Using Global Search" on page 398


## Jobs Dashboard

The **Data Protection** (**Jobs**) dashboard is a detailed interface where you can create and manage jobs, as well as get an overview of job details. For a detailed explanation of each component in the Jobs dashboard, see the sections below.


- Group/Job Overview Dashboard
  - Action Bar
  - Summary Bar
  - Jobs Table
  - Group Info
  - Overview Panes
- Job Dashboard
  - Action Bar
  - Summary Bar
  - Job Info
  - Job Settings
  - Job Objects
  - Overview Panes


## Group/Job Overview Dashboard

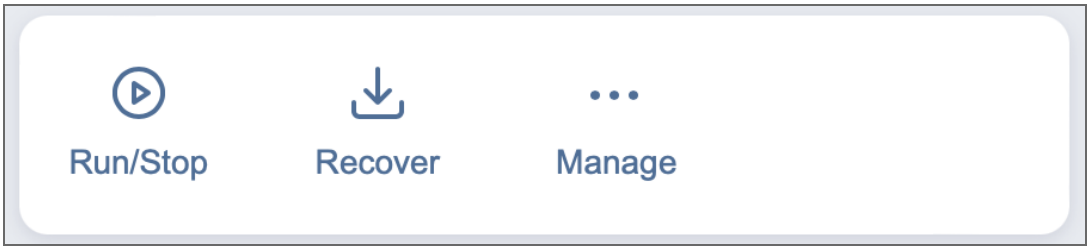
The **Job overview** and job group views offer an overview of multiple jobs. See the sections below for more information.

## Action Bar

The group and **Job overview** action bars contain the following three job actions:

- **Run/Stop**: Opens the Run/Stop Jobs dialog box
- **Recover**: Brings up a list of recovery options for the selected group of jobs
- **Manage**: Brings up the options to **Rename**, **Delete**, or **Disable** a job group, as well as change the destination for all backup jobs in the group

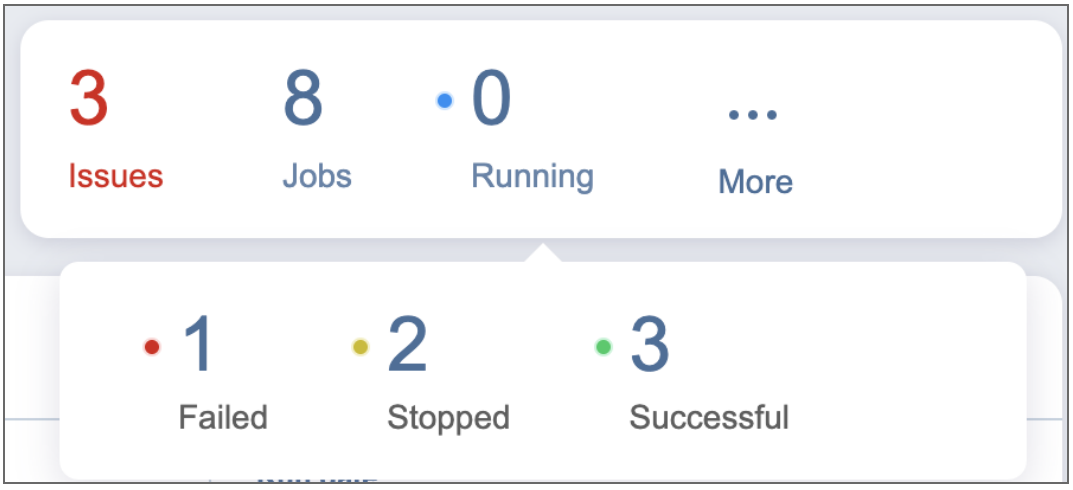
For more information on using the action bar, see "Running Jobs on Demand" on page 364.




## Summary Bar

The summary bar displays information about the jobs in a given group. The data displayed is as follows:

- **Issues**: Total number of issues/alarms for the group of jobs. When clicked, this displays the Alarms & Notifications dialog box.
- **Jobs**: Total number of jobs in the group.
- **Running**: The number of running jobs in the group.
- **Failed**: The number of failed jobs in the group.
- **Stopped**: The number of stopped jobs in the group.
- **Successful**: The number of successful jobs in the group.



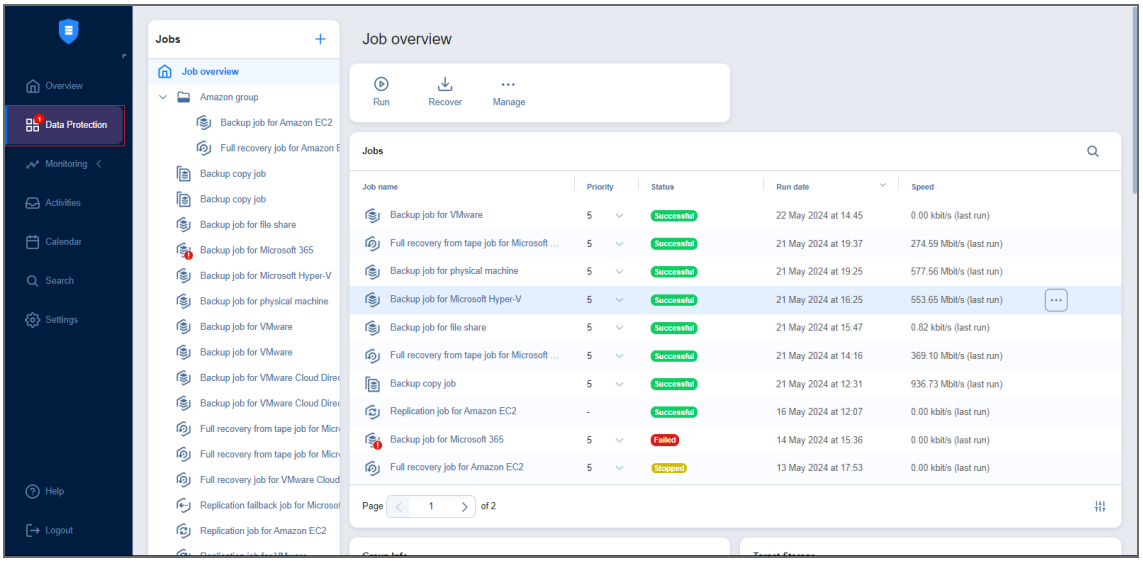

## Jobs Table

The **Jobs** table shows a list of jobs and the information about each job in the following columns:

- **Job name**: The name of a given job in the group.
- **Priority**: The priority level of a given job in the group. Click the arrow button to the right of this parameter to change the priority level of a job.

> **Note**
> This column is only available in the Enterprise, Enterprise Essentials, Enterprise Plus, MSP Enterprise, and MSP Enterprise Plus editions.

- **Status**: The status of a given job:
  - **Successful**: The last job run was successfully completed.
  - **Failed**: The last job run failed.
  - **Running**: The job is currently running.
  - **Stopped**: The last job run was stopped.
  - **Not executed yet**: The job has not been executed yet.
- **Run date**: The date of a given job's last run.
- **Speed**: If the job is currently running, displays the current job run speed. If the job is not currently running, displays the speed of the last job run.




To customize the sorting of the **Jobs** table, click the head of the column you wish to sort by. To change the order of the columns, drag and drop a column to the needed position. You may also search for a job by clicking the **Search** button at the top of the table. To manage a job in the table, hover over a job and click the ellipsis **Manage** button on the right side.

## Group Info

This pane displays current information about the jobs in the selected group. The information includes:

- Currently running jobs, displayed as a ratio to the total number of jobs
- The status of recent jobs; **Completed**, **Failed**, or **Stopped**

- The number of source objects and their respective total size (if applicable)




## Overview Panes

There are several other panes that give an overview for the chosen job group. These panes are as follows:

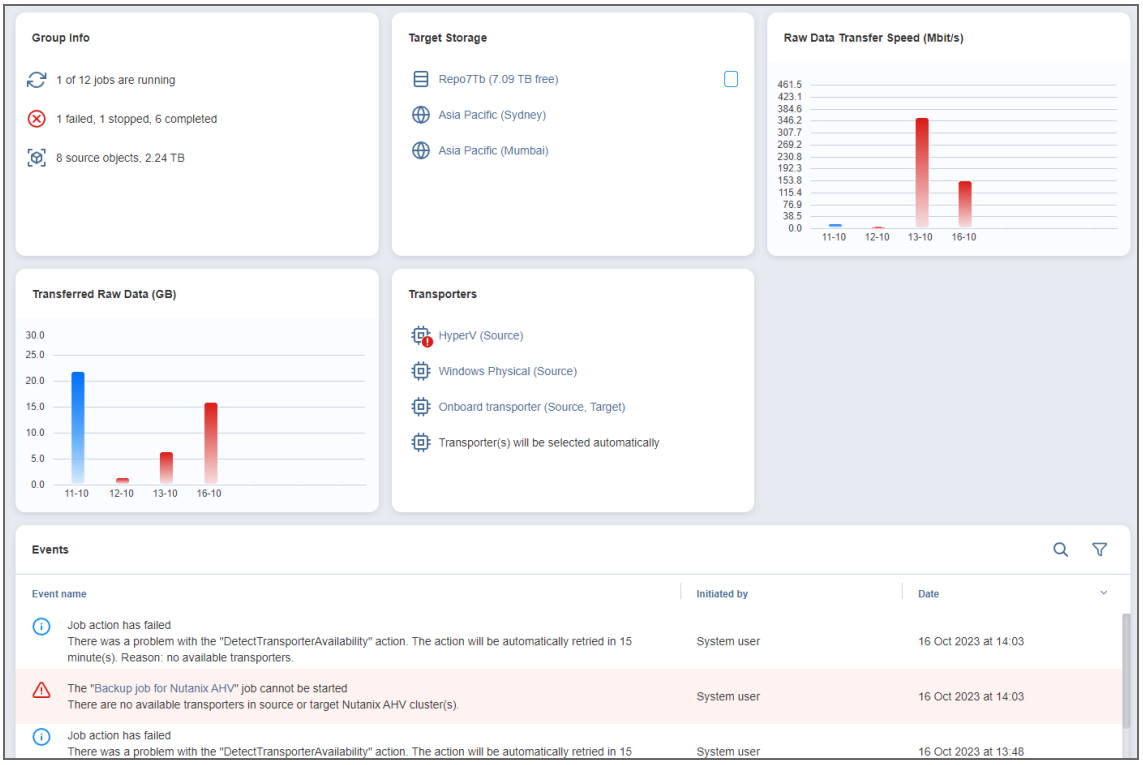
- **Target Storage**: The target storage(s) of the jobs in the chosen group. To open the popup with additional information about it (**Data storage type**, **Free space**, **Used space**, **Transporter**, **Path**, **Compression**, **Deduplication**, **Space savings** fields), click this target storage.

  > **Note**
  > If a federated repository is selected as a destination, only **Data storage type** (*Incremental with full backups*), **Free space**, **Used space** fields are displayed.

- **Raw Data Transfer Speed (Mbit/s)**: The raw data transfer speed for previous job runs (if no job in the group is currently running) or current job run (if a job is currently running). If a job run includes multiple backup objects, the aggregated data transfer speed of all backup objects is displayed.
- **Transferred Raw Data (GB)**: The amount of raw transferred data before compression/deduplication for the current job run or past job run(s).

  > **Note**
  > The **Raw Data Transfer Speed** and **Transferred Raw Data** widgets show:
  > - Raw (uncompressed) data size for disk-based backups.
  > - Actual transferred (compressed) data size for file-based backups (File-level backup and File share backup jobs).

- **Transporters**: Table of the Transporters used by the group of jobs
- **Events**: Table of alarms/notifications for the given group of jobs




# Job Dashboard

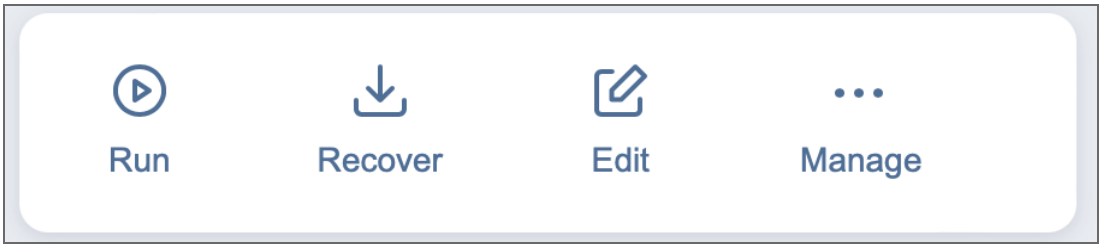
When selecting a specific job in the **Jobs** menu, the following information is displayed.

## Action Bar

The job action bar contains the following four actions:

- **Run**: Opens the job run dialog box
- **Recover**: Brings up a list of recovery options for the given job (backup and replication jobs only)
- **Edit**: Opens the job edit wizard
- **Manage**: Brings up the options to **Clone**, **Merge**, **Rename**, **Create report** for, **Enable**/**Disable**, or **Delete** the job.

For more information on using the action bar, see “Running Jobs on Demand” on page 364.



## Summary Bar

The summary bar displays information about a job. The data displayed is as follows:

- **Issues**: Total number of issues/alarms for the job. When clicked, the Alarms & Notifications dialog box is displayed.
- **Objects**: Total number of objects covered by the job
- **Source size**: Total size of the objects covered by the job

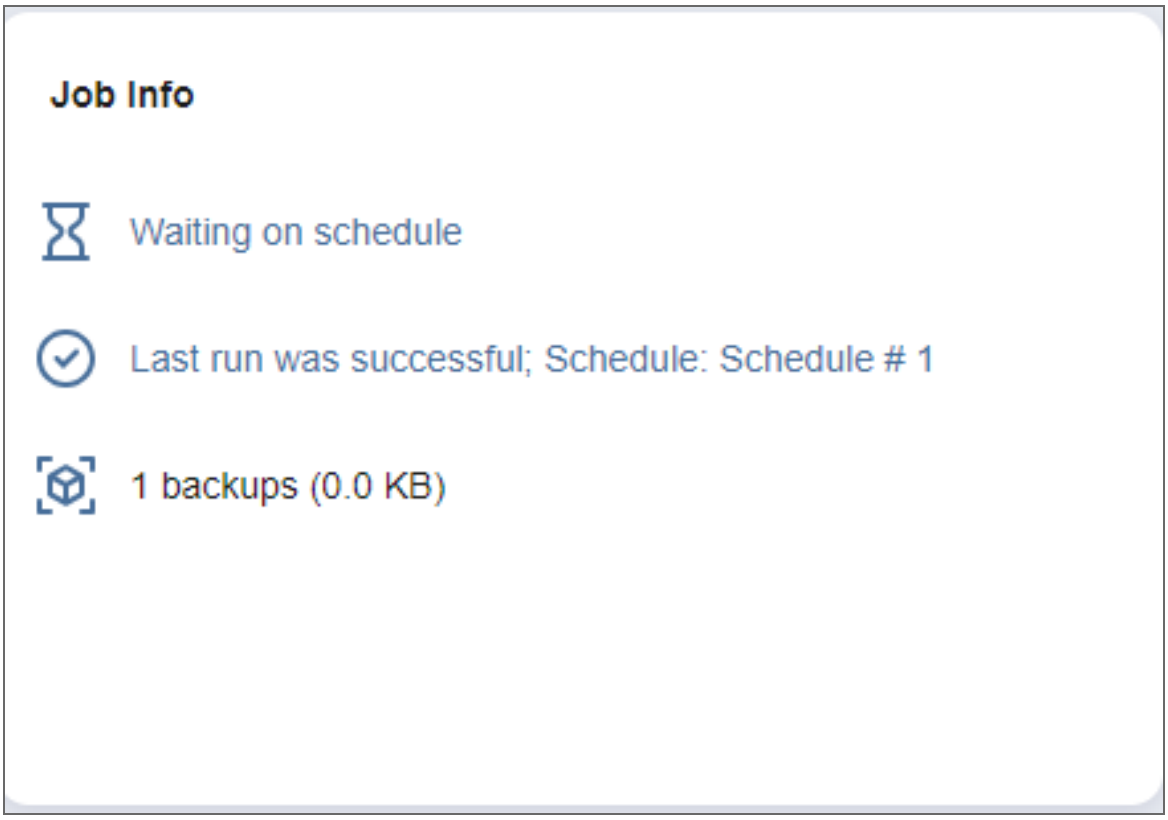



## Job Info

This pane displays current information about the job. The information includes:

- The running schedule for this job. NAKIVO Backup & Replication displays the name of the schedule associated with the last job run with the following exceptions:
  - If the schedule name is not set, the product displays a number of the schedule.
  - If the job was run manually and in the manual run no specific schedule settings were used, the “Manual run” is displayed instead of the schedule name.
- The status of this job; **Successful**, **Failed**, **Running**, **Stopped**, or **Not executed yet**

- The number of source objects and their respective total size (if applicable)



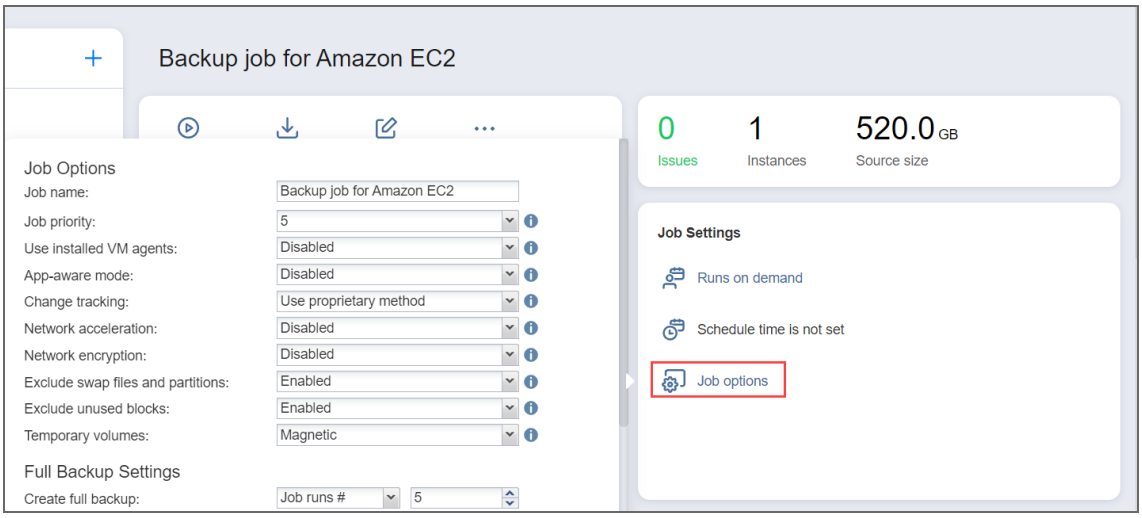

## Job Settings

This pane allows you to view and edit certain options for a job. The settings displayed are as follows:

- The running schedule for this job (if applicable)
- Scheduled run time(s) for this job (if applicable)
- **Recovery points retention**: Clicking this opens a dialog box with **Retention Settings** for this job (if applicable)

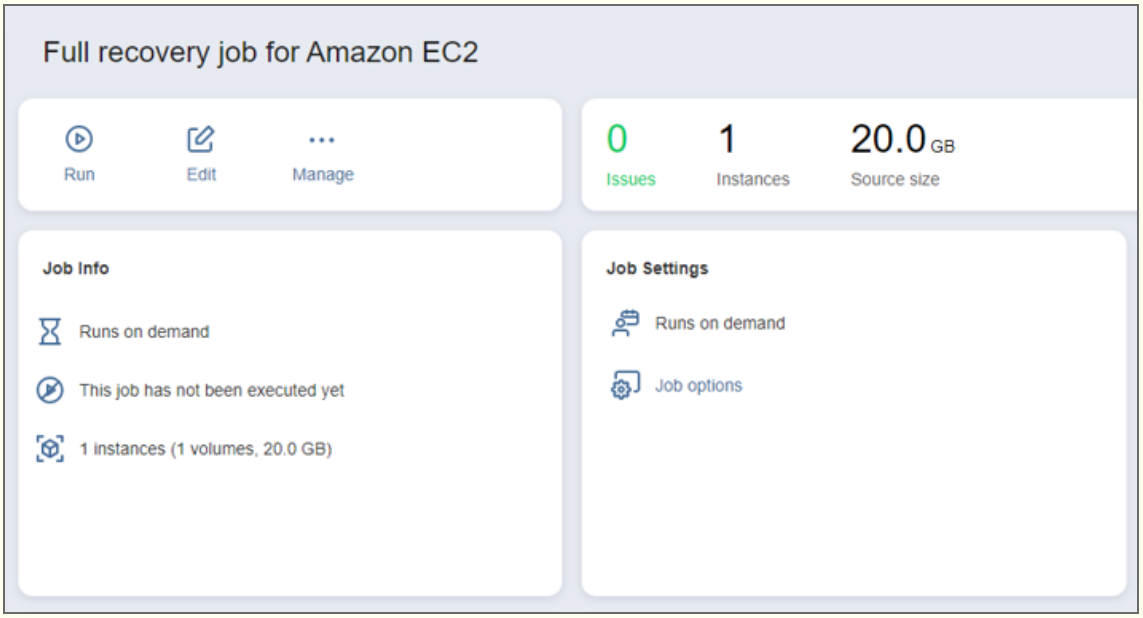
> **Note**
>
> This option is only available in the following cases:
>
> - Your version of NAKIVO Backup & Replication is older than v10.8.
> - You have updated NAKIVO Backup & Replication from a version older than v10.8 to v10.8 or newer and have not enabled the new scheduler for the respective job.
> - You enabled legacy retention in the Expert tab.

- **Job options**: Clicking this opens a dialog box with **Job Options** for this job.




**Note**

The **Job Settings** pane for NAKIVO Backup & Replication recovery jobs only displays the settings relevant to this particular job:

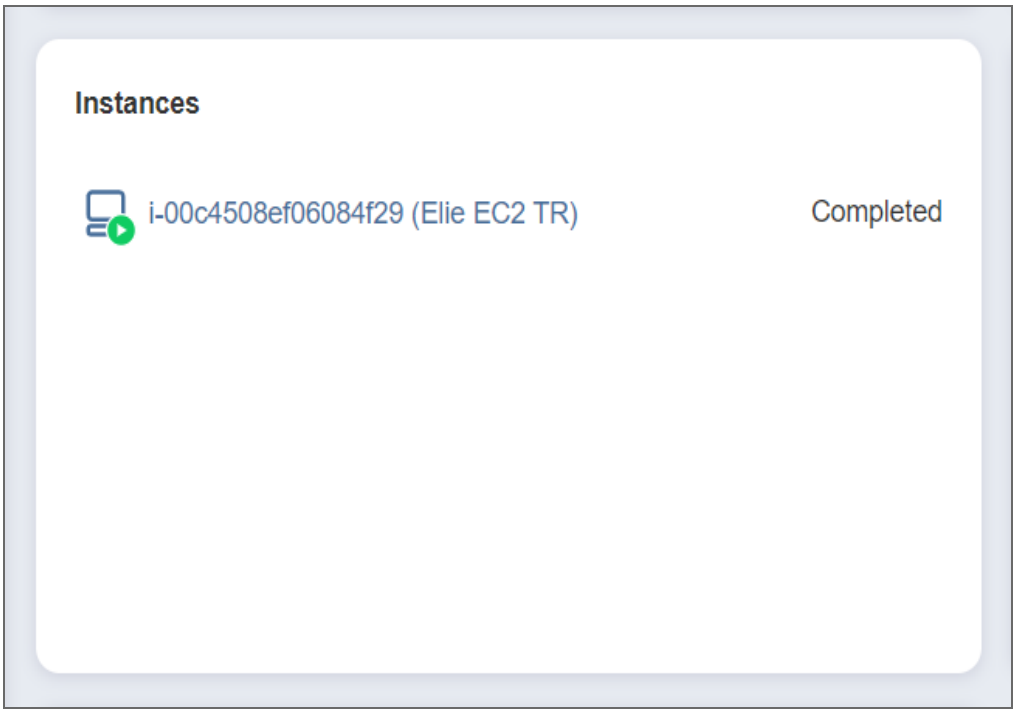



## Job Objects

This pane displays a list of backup/replication/recovery objects based on the respective object type. The objects can be one of the following:

- Virtual Machines
- Instances
- Backups
- Physical Machines

- Databases
- Microsoft 365 Items
- File Share Items



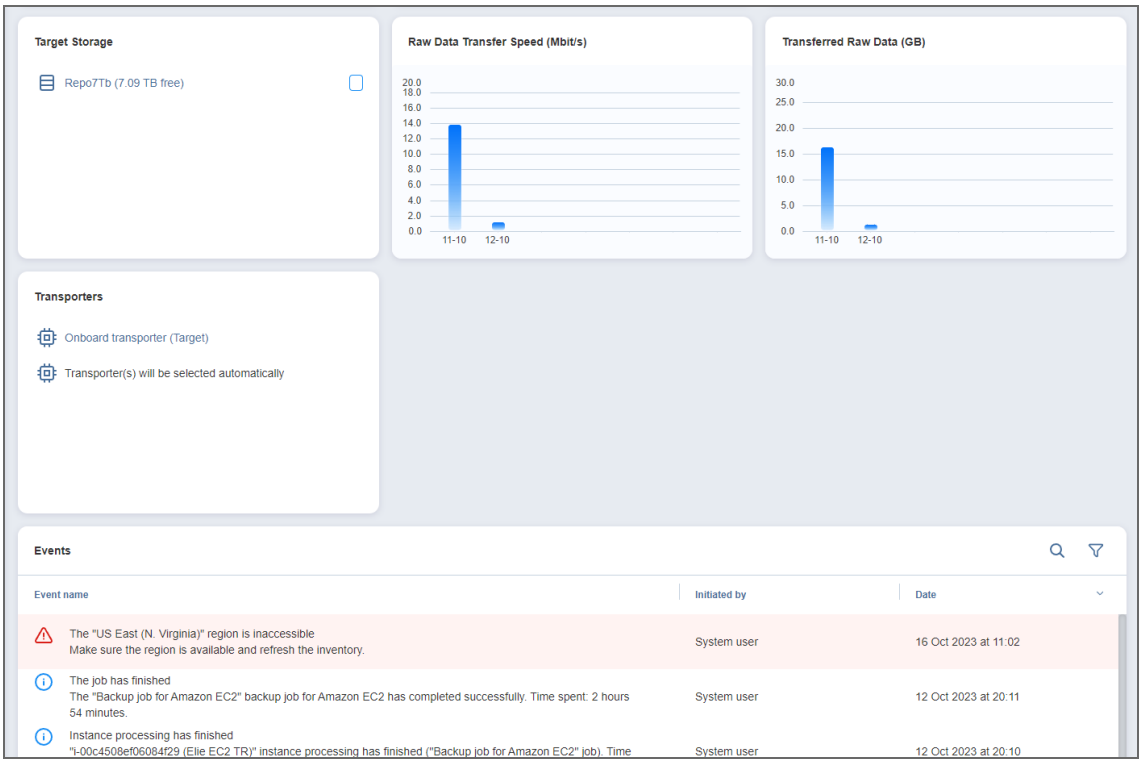

## Overview Panes

Several other panes give an overview of the chosen job. These panes are as follows:

- **Target Storage**: The target storage of the chosen job
- **Raw Data Transfer Speed (Mbit/s)**: The raw data transfer speed for the current job run or previous job runs if the job is not currently running. If a job run includes multiple backup objects, the aggregated data transfer speed of all backup objects is displayed.
- **Transferred Raw Data (GB)**: The amount of raw transferred data before compression/deduplication for a current job run or past job run(s).

> **Note**
> The **Raw Data Transfer Speed** and **Transferred Raw Data** widgets show:
> - Raw (uncompressed) data size for disk-based backups.
> - Actual transferred (compressed) data size for file-based backups (File-level backup and File share backup jobs).

- **Transporters**: Table of the Transporters used by the job
- **Events**: Table of alarms/notifications for the given job


Target Storage
Repo7Tb (7.09 TB free)
Raw Data Transfer Speed (Mbit/s)
20.0
18.0
16.0
14.0
12.0
10.0
8.0
6.0
4.0
2.0
0.0
11-10
12-10
Transferred Raw Data (GB)
30.0
25.0
20.0
15.0
10.0
5.0
0.0
11-10
12-10
Transporters
Onboard transporter (Target)
Transporter(s) will be selected automatically
Events
Event name
Initiated by
Date
The "US East (N. Virginia)" region is inaccessible
Make sure the region is available and refresh the inventory.
System user
16 Oct 2023 at 11:02
The job has finished
The "Backup job for Amazon EC2" backup job for Amazon EC2 has completed successfully. Time spent: 2 hours 54 minutes.
System user
12 Oct 2023 at 20:11
Instance processing has finished
"i-00c4508ef06084f29 (Elie EC2 TR)" instance processing has finished ("Backup job for Amazon EC2" job). Time
System user
12 Oct 2023 at 20:10

# Running Jobs on Demand

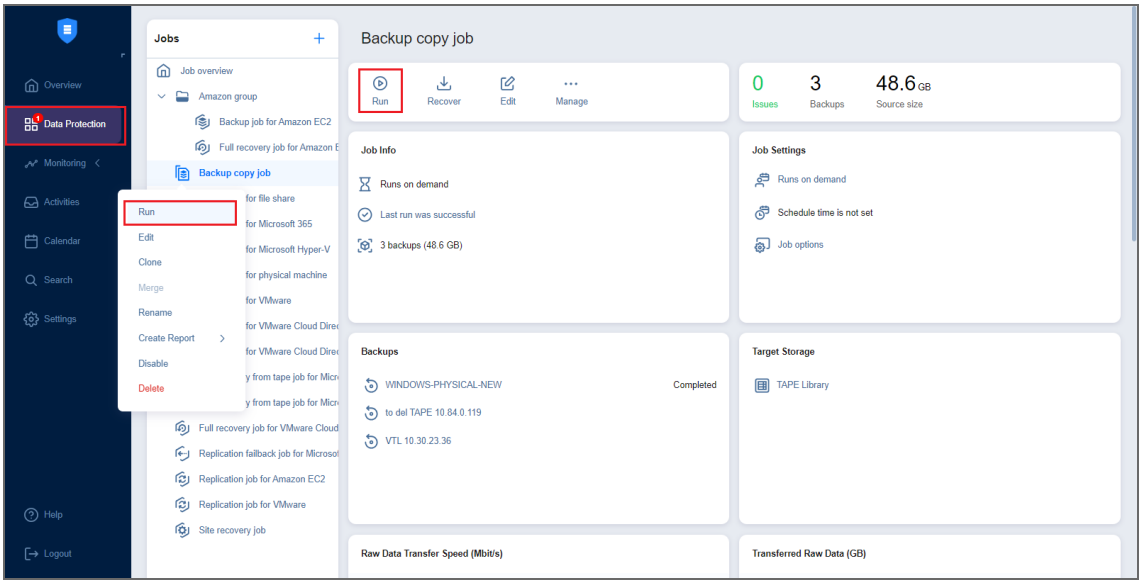
Use the **Data Protection** menu to start and stop jobs on demand.

- Starting Jobs
- Stopping Jobs
- Managing Grouped Jobs

## Starting Jobs

To start a job, follow the steps below:

1. Go to the **Data Protection** menu, select the job from the list of jobs, and click **Run**. Alternatively, right-click a job to bring up the action menu and click **Run**.

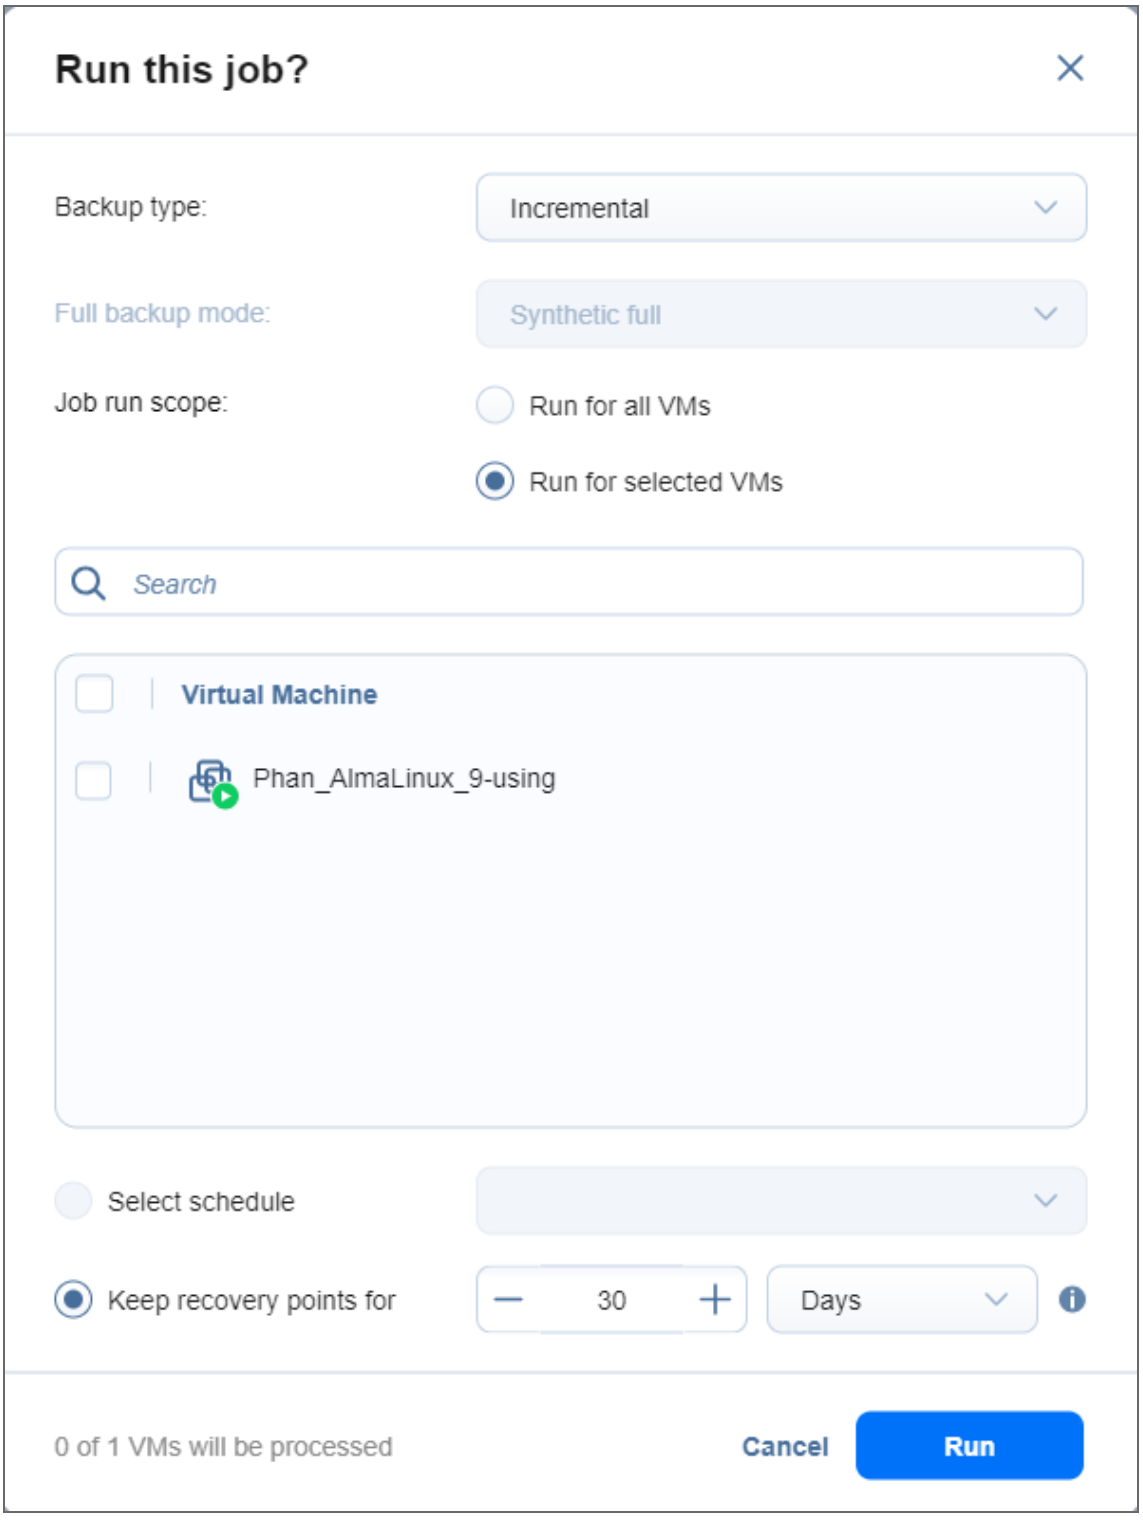



2. Choose one of the following options:
   - **Run for all VMs/VM templates/backups/physical machines/instances/databases/sites/accounts/items**: The job runs for all job objects.
   - **Run for selected VMs/VM templates/backups/physical machines/instances/databases/sites/accounts/items**: The job runs for the job objects that you select.
   - **Run for failed VMs/VM templates/backups/physical machines/instances/databases/sites/accounts/items**: If applicable, the job runs for previously failed job objects only.
   - **Run for stopped VMs/VM templates/backups/physical machines/instances/databases/sites/accounts/items**: If applicable, the job runs for previously job stopped objects only.

3. If backups in the Backup Repository selected for a job are stored in separate files, you have to choose between the following backup types:
    - **Incremental**: The job creates an incremental backup.
    - **Full**: The job creates a full backup. When you choose this option, choose one of the full backup modes:
        - **Synthetic full**: The application first creates an incremental backup—that is, transfers only the data that changed since the last backup—and then creates a new full backup using the last full backup and the chain of subsequent incremental recovery points.
        - **Active full**: The application reads all source machine data and transfers it to the backup repository to create a full backup.
4. For backup and backup copy jobs, you can use preconfigured retention settings by selecting **Use job retention** (legacy retention approach) or **Select schedule** (schedule retention approach), or specify custom retention settings for a manual job run by selecting **Keep recovery points for**.
    - **Use job retention**: Select this option to use the preconfigured legacy retention settings for a job run. If a previous run for this job was stopped or failed, the settings used for that run are selected by default.
    - **Select schedule**: Select this option to choose a preconfigured schedule and its retention settings for this job. If a previous run for this job was stopped or failed, the settings used for that run are selected by default. Recovery points created by a manual job run using this option are automatically assigned expiration dates.
    - **Keep recovery points for**: The recovery points created by this job run are kept for the specified period of time and then expire. The expired recovery points are removed during the following job run.

    > **Note**
    > If a job does not support retention or has the **Do not schedule, run on demand** option selected, only the **Keep recovery points for** option will be available.

5. Click the **Run** button to confirm your selection.

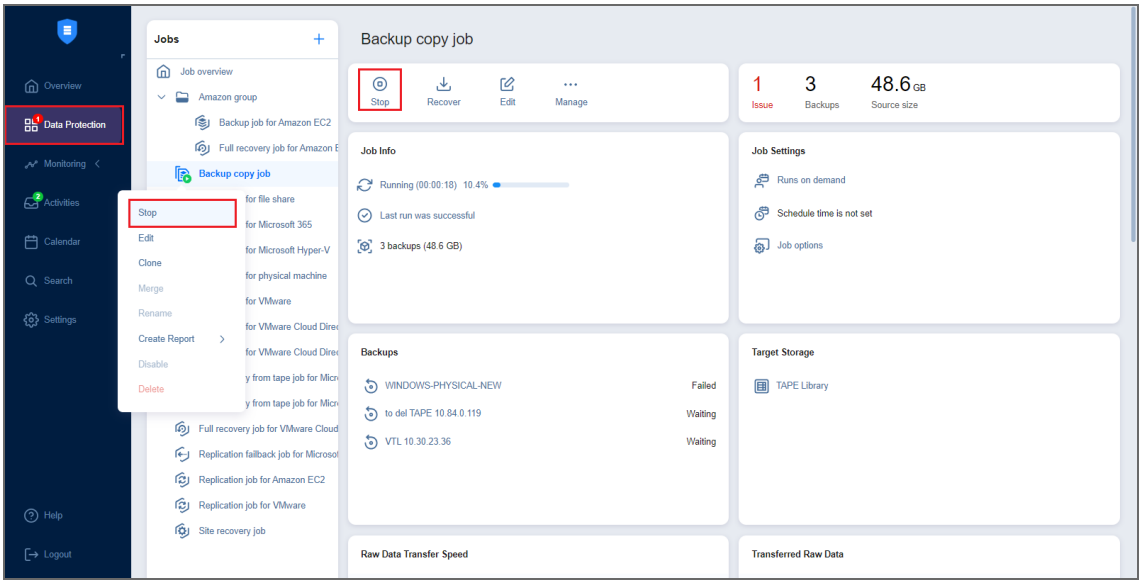



The product will close the dialog box and start running your job.

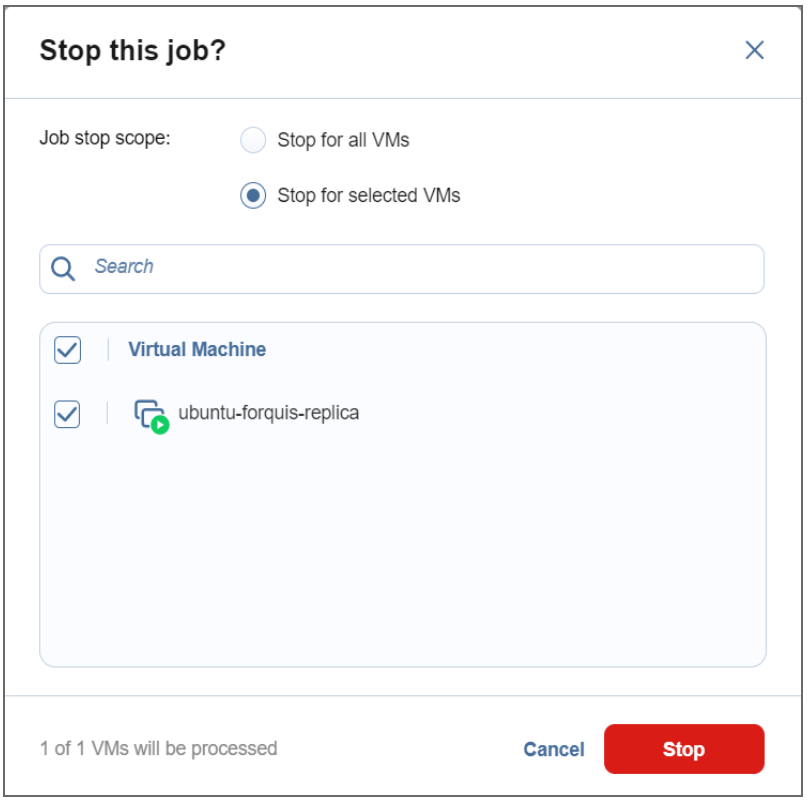
## Stopping Jobs

To stop a job that is running, follow the steps below:

1. Go to the **Data Protection** menu, select the job from the list of jobs, and click **Stop** . Alternatively, right-click a job to bring up the action menu and click **Stop**.




2. In the dialog box that opens, choose one of the following:
   - **Stop for all VMs/VM templates/backups/physical machines/instances/databases/sites/accounts/items**: The job stops for all job objects.
   - **Stop for selected VMs/VM templates/backups/physical machines/instances/databases/sites/accounts/items**: The job stops for the job objects that you select.
3. Click the **Stop** button in the dialog box to confirm your selection.

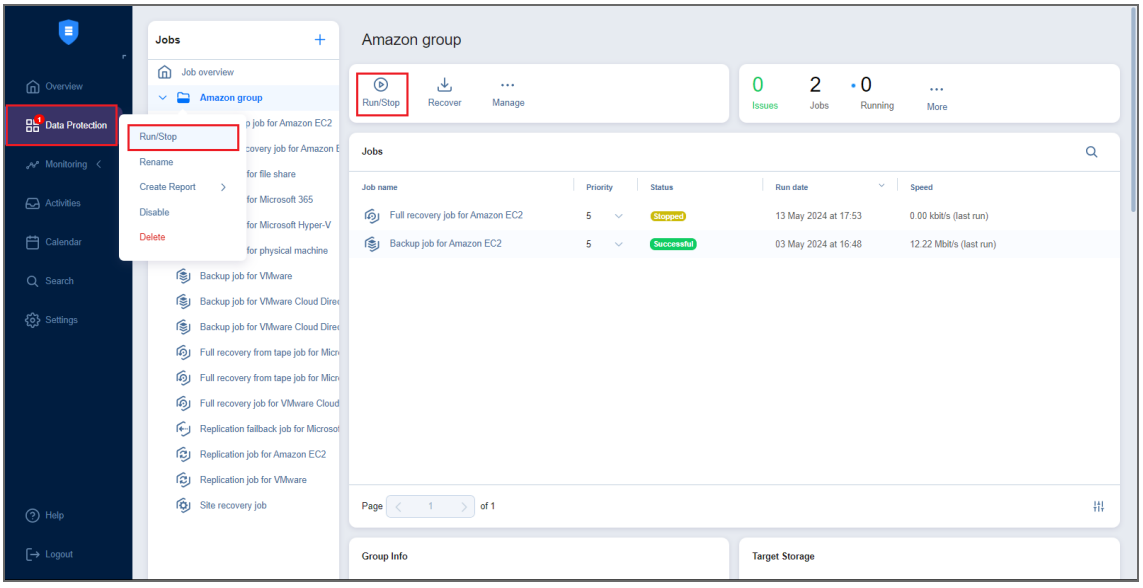



The application closes the dialog box and stops your job.

# Managing Grouped Jobs

To efficiently start or stop jobs in bulk (run all failed jobs, for example), follow these steps:

1. From the **Data Protection** menu, select the needed job group and click **Run/Stop**. To manage all jobs and groups at once, select **Overview** and click **Run/Stop**. Alternatively, right-click on the needed job group to bring up the action menu and click **Run/Stop**.



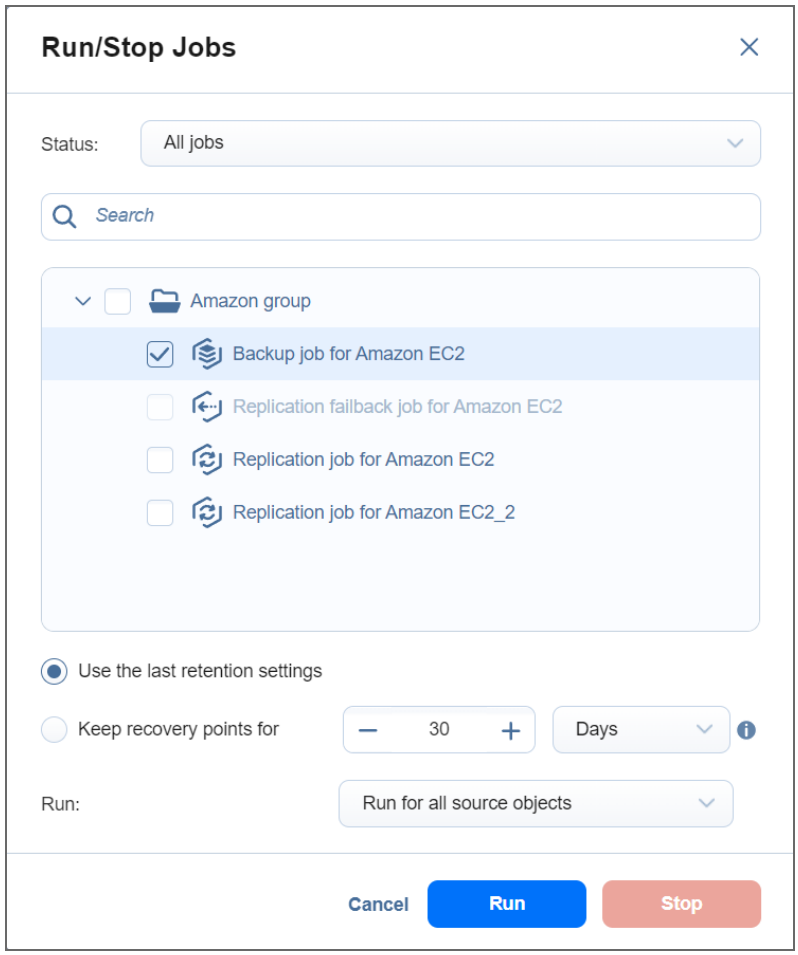

2. In the drop-down **Status** menu, select one of the following:

- **All jobs:** Displays all jobs in the group
- **Failed jobs:** Displays all failed jobs in the group
- **Stopped jobs:** Displays all stopped jobs in the group

3. Select the jobs you want to run/stop.
   a. When running backup or backup copy jobs, specify the retention settings with one of the following options:
      - **Use the last retention settings**: Select this option to use the retention settings from the last job run for the manual job run.
      - **Keep recovery points for**: The recovery points created by this job run are kept for the specified period of time and then expire. The expired recovery points are removed during the following job run.

      > **Note**
      > If the group of jobs contains at least one job that isn't a backup/backup copy job, does not support retention, or has had its retention settings changed since the previous run, only the **Keep recovery points for** option will be available.

   b. In the lowest drop-down menu, specify (if applicable) whether you want the operation to run for failed source objects, stopped source objects, or all source objects.
4. Click the **Run** or **Stop** button to confirm your selection.


Run/Stop Jobs
Status:
All jobs
Search
Amazon group
Backup job for Amazon EC2
Replication failback job for Amazon EC2
Replication job for Amazon EC2
Replication job for Amazon EC2_2
Use the last retention settings
Keep recovery points for
30
Days
Run:
Run for all source objects
Cancel
Run
Stop

# Managing Jobs

Using the **Data Protection** menu, you can easily manage jobs. Use the **Manage** menu to rename, edit, merge, delete, and enable/disable jobs.

- Viewing Job Run Details
- Renaming Jobs
- Editing Jobs
- Cloning Jobs
- Merging Jobs
- Deleting Jobs
- Disabling and Enabling Jobs
- Grouping Jobs
- Creating Groups
- Creating Job Reports

## Viewing Job Run Details

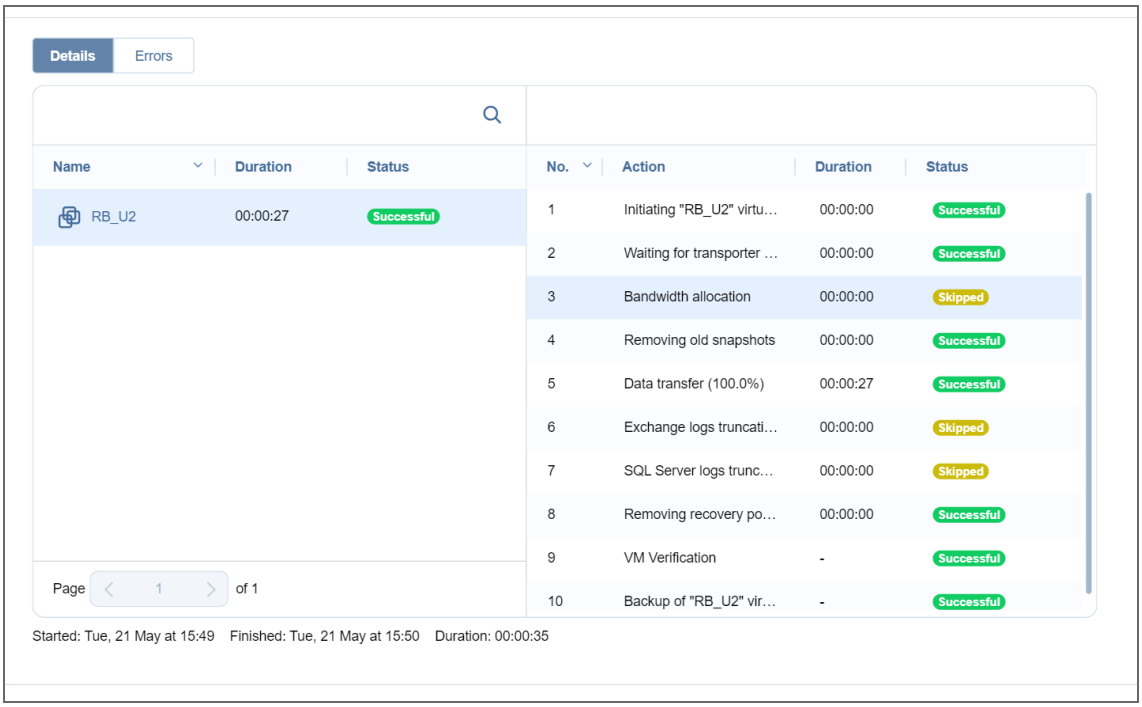
You can view job run details in the **Data Protection** or **Activities** dashboard.

- **Jobs dashboard**: In the **Jobs** pane on the right, click the job status to view details.
- **Activities dashboard**: Click on the job name and then click on the **details** link.

The **Job Run Details** popup window opens displaying the following details:

- On the **Details** tab:
  - **Name**: The column displays the name of the job object.
  - **Duration**: The column displays the real-time duration for job processing. For completed jobs, the column shows the total duration for job processing.

- **Status**: The column displays the last run status of the job.




**Notes**

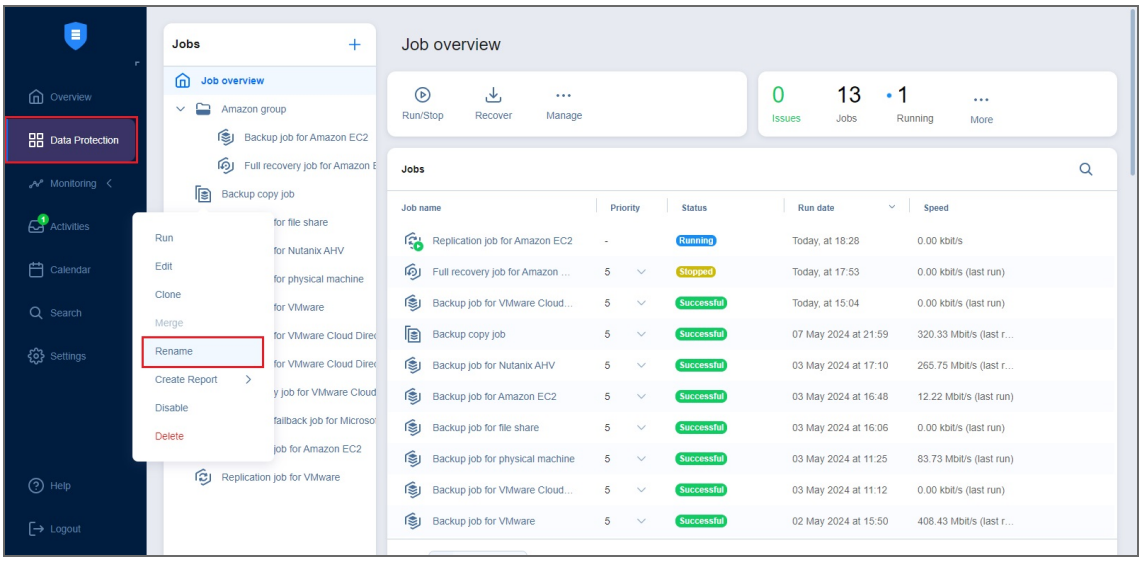
- You can click on any of the columns in the window to sort the list alphabetically or numerically. Additionally, you can use the search field for the specific job object.
- On the right side, you can find the list of the job runs for this job and their durations. Clicking each specific run displays the detailed list of actions that were performed during the job run, their status, and the amount of time each of these actions took.
- If the job was never retried, the right pane shows the actions that NAKIVO Backup & Replication took when running this job.

- On the **Errors** tab, you can view all the alarms associated with this job.

# Renaming Jobs

1. From the list of jobs, right-click on the job you want to rename.
2. Click **Rename**.

3. In the dialog box that opens, specify the new name for the job and click **Rename**.

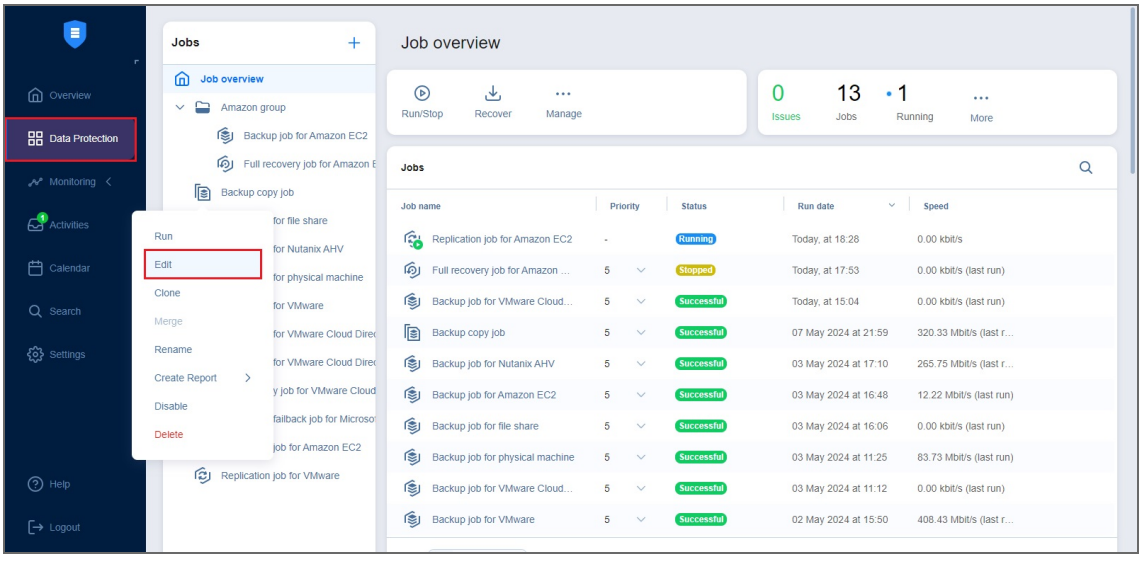



> **Note**
>
> You can also rename jobs by selecting a job and clicking **Manage** > **Rename**.

# Editing Jobs

To edit a job, follow the steps below:

1. Right-click on the job you want to edit from the list of jobs.
2. Click **Edit**.
3. In the **Edit** wizard, click the needed page to open it for editing.
4. Make the required changes and click **Save** or **Save & Run**.



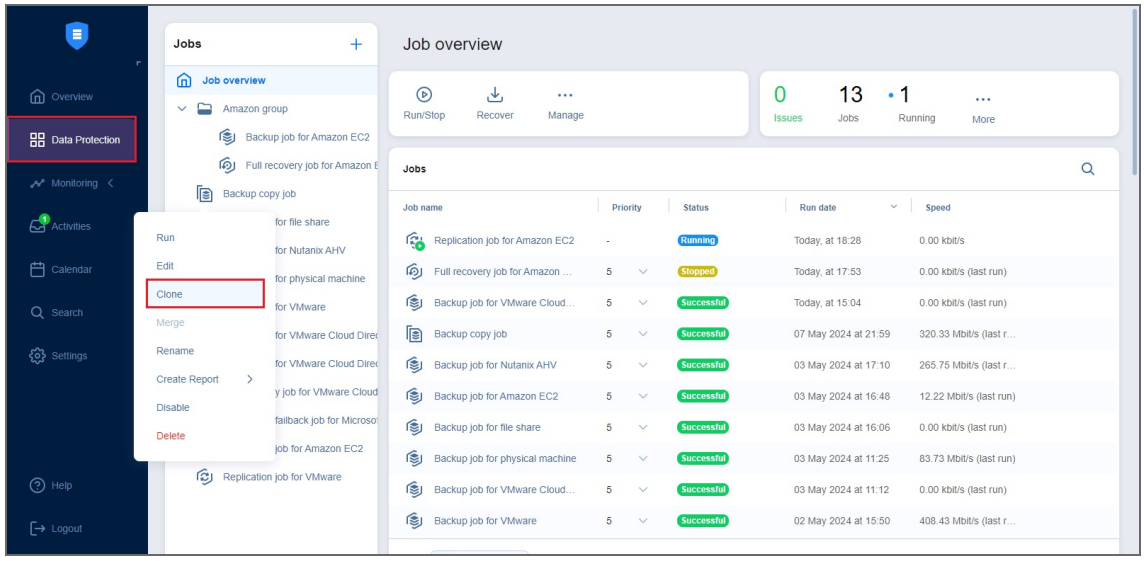
**Notes**

- You can edit the job while it is running, but the changes will be applied only when the job run has completed.
- You can also edit jobs by selecting a job and clicking **Manage** > **Edit**.

# Cloning Jobs

To clone a job, follow the steps below:

1. Right-click on the job you want to clone from the list of jobs.
2. Click **Clone**.



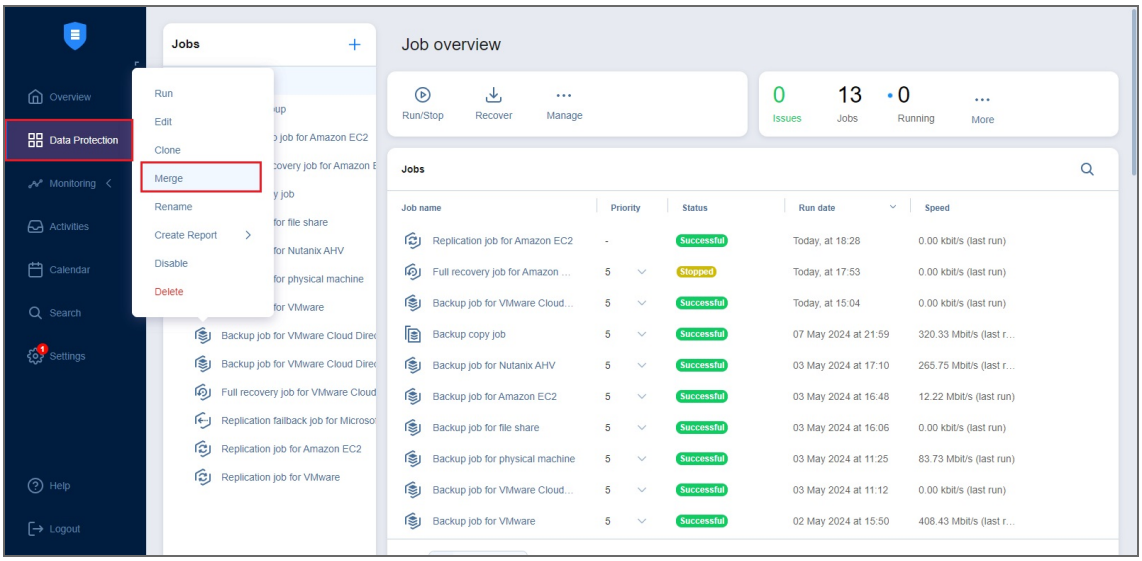

**Note**

You can also clone jobs by selecting a job and clicking **Manage** > **Clone**.

# Merging Jobs

NAKIVO Backup & Replication allows you to merge jobs of the same type. Before doing this, make sure to check feature requirements and how the feature works. To merge the jobs, do the following:

1. From the list of jobs, right-click on the source job you want to merge.
2. Click **Merge**.
3. Choose the target job for the merge and click **Apply**.
4. After the merge is finished, click **Close** to close the popup.




**Notes**

The **Merge** button may be unavailable in the following cases:

- The selected job does not meet the feature requirements.
- The selected job is currently running.
- There are no target jobs available to merge the selected job with.

# Deleting Jobs

To delete a job follow the steps below:

1. Right-click on the job you want to delete from the list of jobs.
2. Click **Delete**.
3. In the dialog box that opens, select one of the following:
   - **Delete job and keep backups**
   - **Delete job and backups**

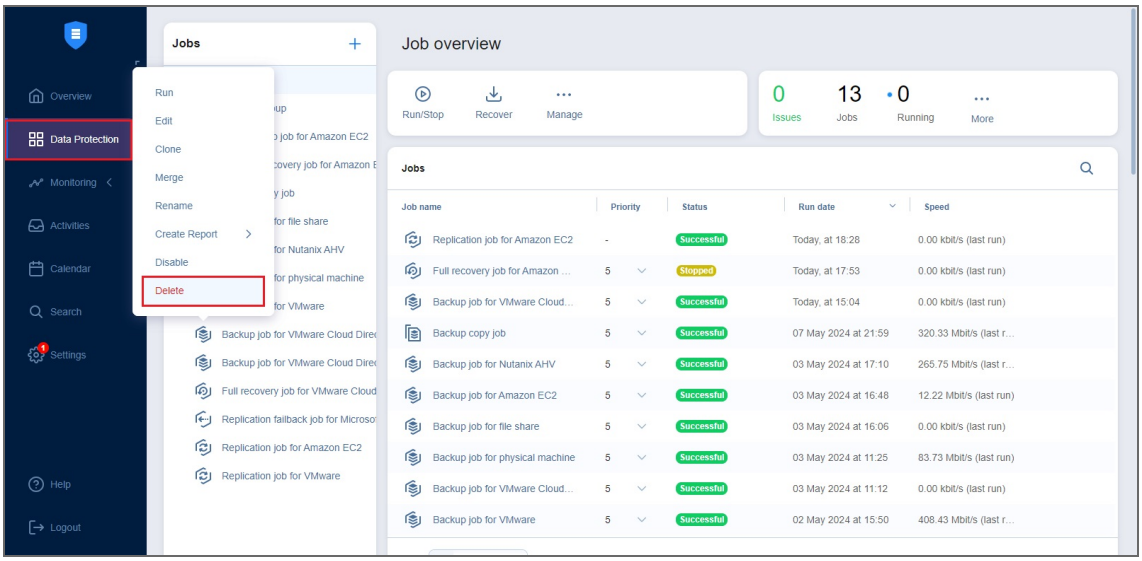
**Note**

If the job has target objects, the options in the dialogue box allow you to keep or delete the job and either of the following:

- backups
- replicas
- recovered VMs
- recovered instances
- replicas and their journals

Note that replicas, recovered VMs, and recovered instances are deleted in the powered off state.

4. Click **Delete**



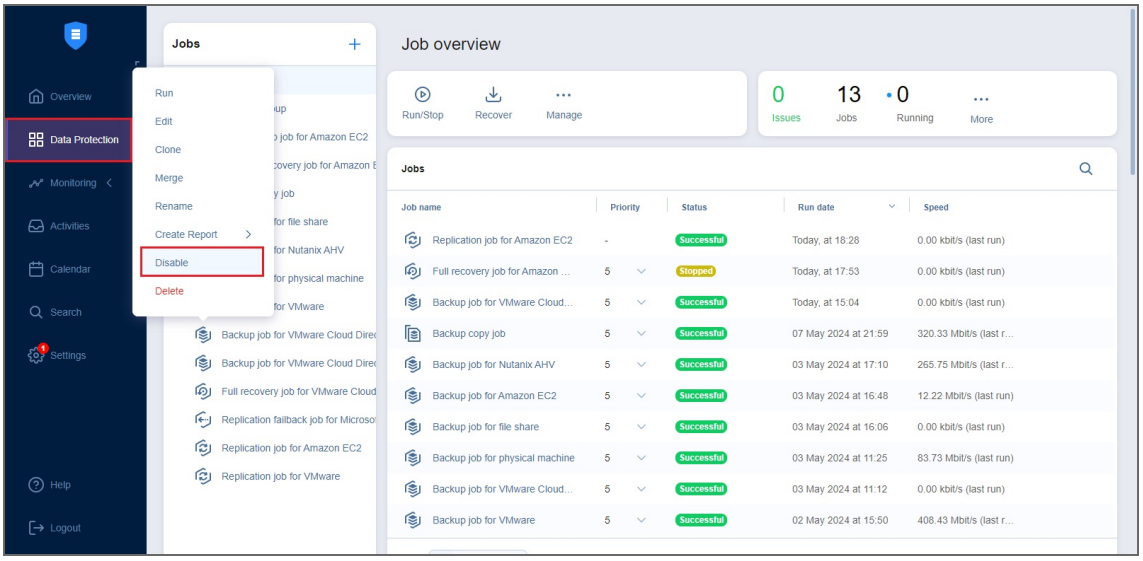

**Notes**

- You can also delete jobs by selecting a job and clicking **Manage** > **Delete**.
- Backups can also be deleted from Backup Repositories.

# Disabling and Enabling Jobs

NAKIVO Backup & Replication provides you with the ability to disable jobs. A disabled job does not run based on the schedule and cannnot be run on demand.

1. From the list of jobs, right-click on the job you want to disable.
2. Click **Disable**.




To enable a job, select **Enable** from the **Manage** menu.

**Note**

You can also manage jobs by selecting a job and selecting the desired action from the **Manage** menu.

# Grouping Jobs

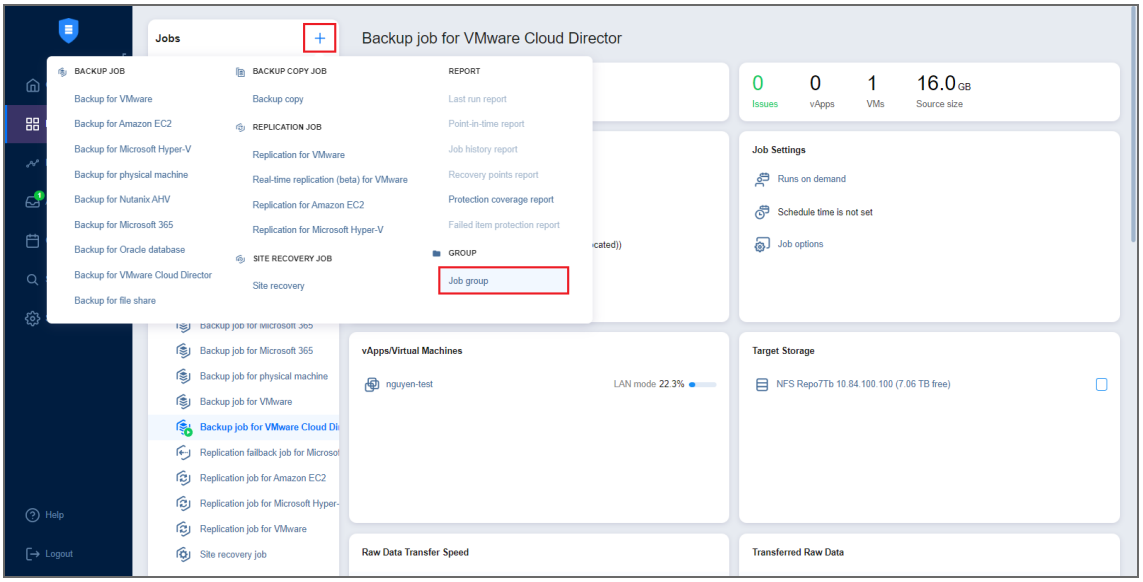
Groups are folders which allow you to:

- Logically arrange jobs (to represent organizations, locations, services, etc.).
- Perform bulk actions with all or selected jobs in a group.

## Creating Groups

To create a group, follow the steps below:

1. In the **Data Protection** menu, click **Create** and then click **Job group**.

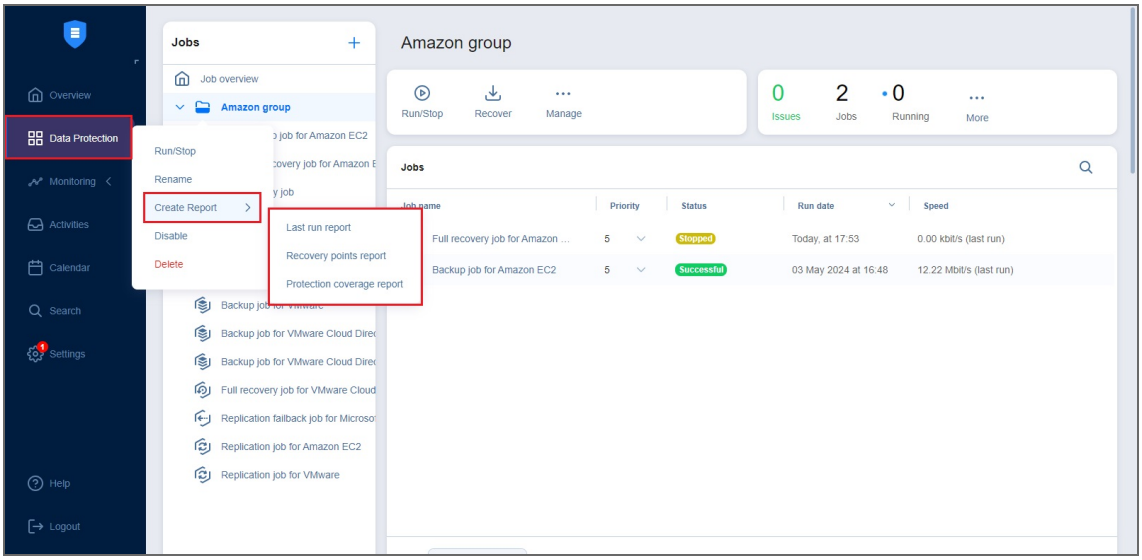



2. Type in the group name in the dialog box that opens and click **Create**.

The following actions are available to manage groups:

- To add a job to a group, simply drag the job into the group.
- To remove a job from the group, drag the job outside the group.
- To delete a group, right-click the group and choose **Delete** from the shortcut menu that opens. Confirm the group deletion when prompted to do so. Note that when deleting a group, the jobs in the group are not deleted. The jobs are moved to the parent group (or to *Overview*).
- To rename a group, double-click the group and enter a new name.
- To enable or disable all jobs inside a group, click the **Enable/Disable** switch.
- To run jobs available in a group, click **Run/Stop** and then click **Run** Jobs. In the dialog box that opens, select the jobs you want to run and click **Run Jobs**.

- To stop running the jobs available in a group, click **Run/Stop** and then click **Stop Jobs**. In the dialog box that opens, select the jobs you want to stop and click **Stop Jobs**.
- To create a report, click **Create Report**. In the dialog box that opens, select one of the following reports from the **Create Report** menu:
  - **Last run report**: Provides data on the last run of the job. Choose either the PDF or CSV format and click **Create**.
  - **Recovery points report**: Contains information regarding the sizes of recovery points for backups/replicas for the chosen job or jobs.
  - **Protection coverage report**: Contains information about all VMs and instances protected by backup/replication jobs, as well as about all unprotected VMs and instances.



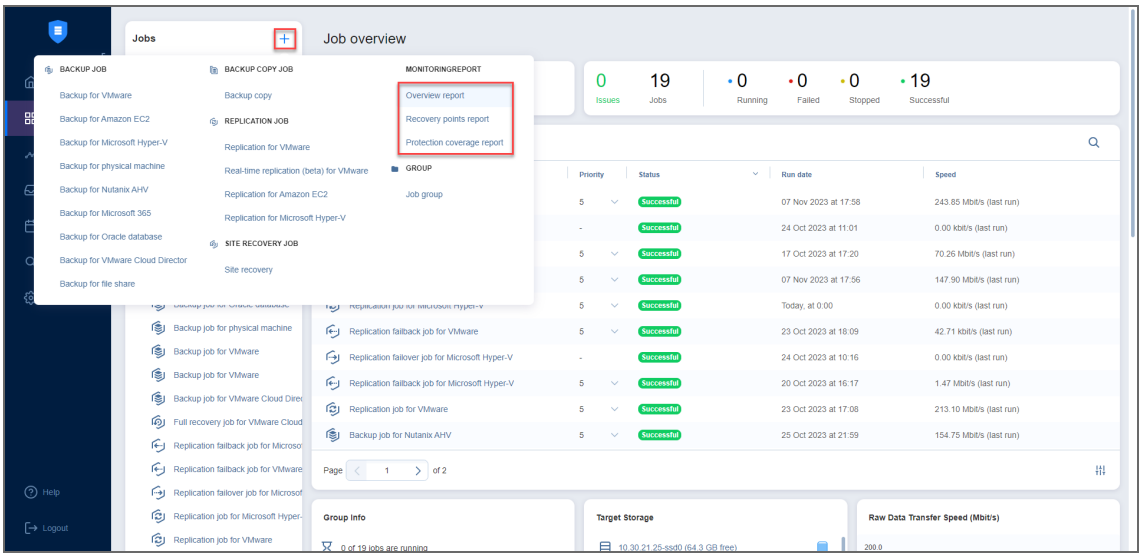

- Choose either the PDF or CSV format and click **Create**.

# Creating Job Reports

To create a general report for all your jobs:

1. Select **Overview** in the **Data Protection** menu.
2. Click **Create Report**.
3. Choose one of the following reports:
   - **Overview report**: Contains information about the status and errors of all jobs.
   - **Recovery points report**: Contains information about the recovery points of backups/replicas for the chosen job or jobs created within a specific period.
   - **Protection coverage report**: Contains information about all VMs and instances protected by backup/replication jobs, as well as about all unprotected VMs and instances. Choose either the PDF or CSV format for the **Protection coverage report** and click **Create**.

4. Choose a location to save the report and click **Save**.




To generate reports from for an individual job, do the following:

1. Go to the list of jobs.
2. Select the job that you need to generate a report for and right-click it or click **Create**.
3. Select one of the following reports from the **Create Report** menu:
   - **Last run report**: Provides data on the last run of the job.
   - **Job history report**: Provides data on job runs that occurred during a specified time period. To generate a report, pick a start date on the left and finish date on the right side of the popup and click **Create**.

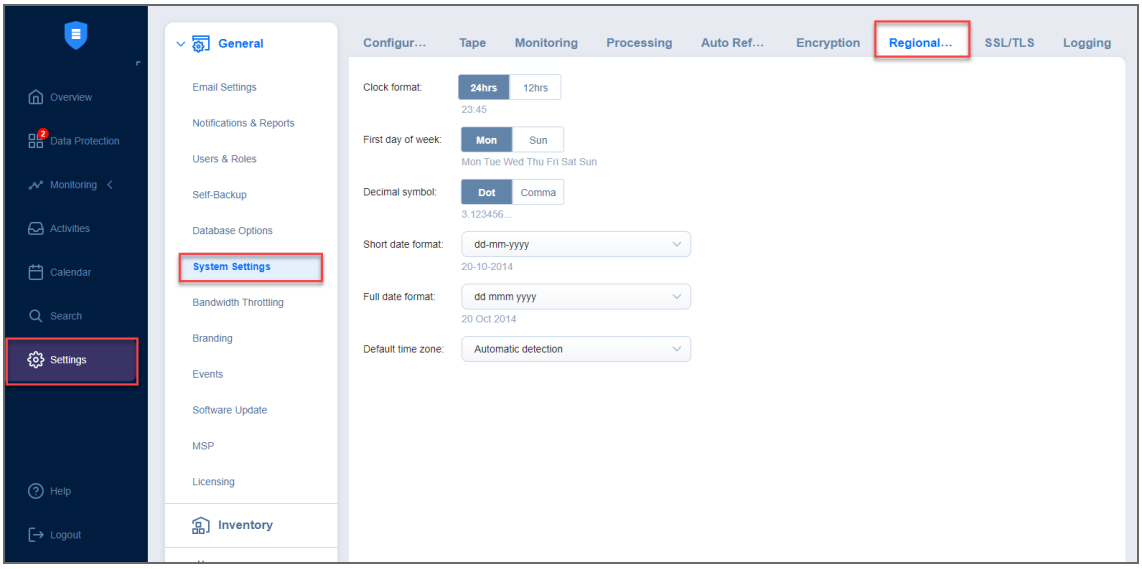
     Job history reports feature the **Calendar View** hyperlink. Clicking on it opens the high-level calendar view page with the summary of the backup status of every backup item within a single job.

     The report data is organized in separate blocks of weeks each consisting of two columns:
     - **Item Name** - displays the job object(s) name as ordered in the corresponding job
     - **Day of the week (Date)** - displays the status for processing each job object

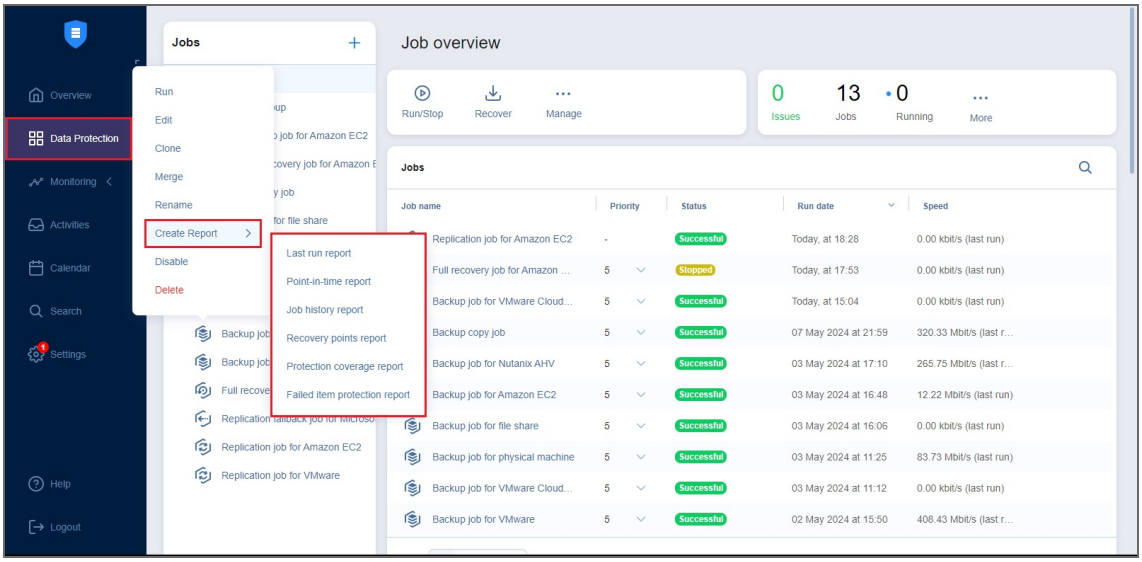
       Possible statuses:
       - **Successful** – displayed for successful object processing
       - **Failed** – displayed for failed object processing
       - **Skipped** – displayed for skipped object processing
       - **Stopped** – displayed for stopped object processing
     - The status field is empty if no object processing was scheduled on a particular day.
     - The columns representing weekdays that are not included in the reported range are empty.
     - Weekdays are displayed when there is at least one run on the week.

**Notes**

- The enhanced **Calendar View** report data is displayed from old to new for all reports where the **Job Runs** block is shown except for the **Monitoring** area and the **Point-in-time report**.
- To customize the first day of the week, go to the **Settings > General tab > System Settings > Regional Format** tab and select the needed day in the **First day of week** selector.

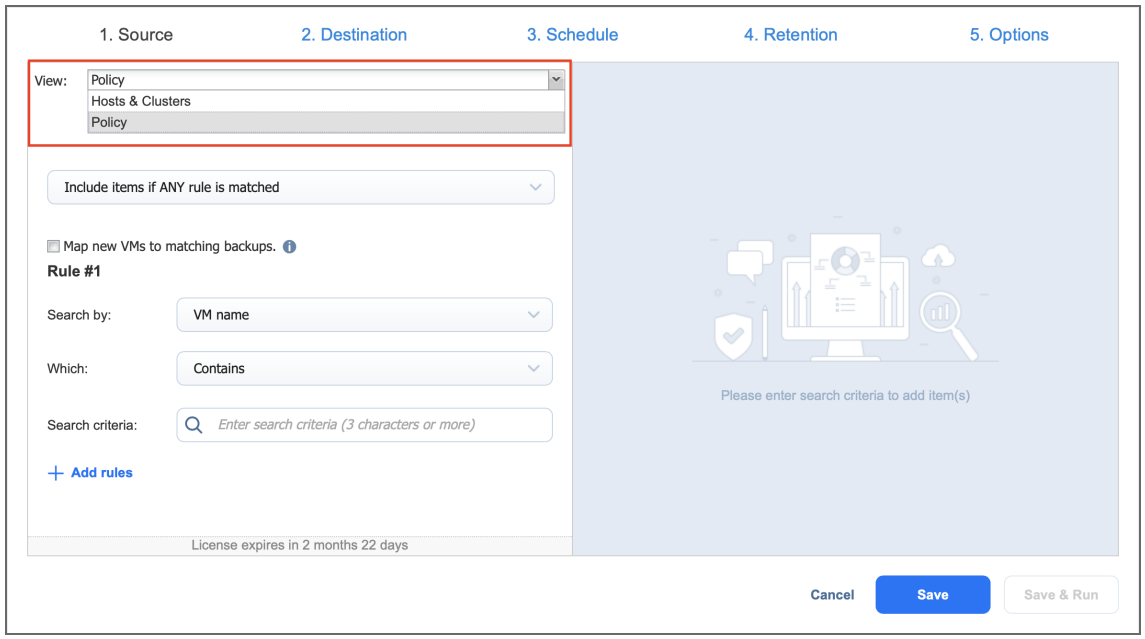



- **Point-in-time report**: Provides data on a particular job run. To generate a report, choose a date in the popup and click **Create.**
- **Recovery points report**: Contains data of recovery points created within specific period.
- **Protection coverage report**: Contains information about all VMs and instances protected by backup/replication jobs, as well as about all unprotected VMs and instances.
- **Failed item protection report**: Contains information about job objects for which processing failed during the last job run. Only backup and replication jobs are included.
- **Site recovery job report**: Contains a summary of the site recovery job, including the result of passing the **Recovery time objective** value, information about all actions performed, and all registered alarms and notifications.


Overview
Data Protection
Monitoring
Activities
Calendar
Search
Settings
Help
Logout
Jobs
Run
Edit
Clone
Merge
Rename
Create Report
Disable
Delete
Last run report
Point-in-time report
Job history report
Recovery points report
Protection coverage report
Failed item protection report
Replication job for Amazon EC2
Replication job for VMware
Job overview
Run/Stop
Recover
Manage
0
Issues
13
Jobs
0
Running
More
Jobs
Job name
Priority
Status
Run date
Speed
Replication job for Amazon EC2 - Successful Today, at 18:28 0.00 kbit/s (last run)
Full recovery job for Amazon ... 5 Stopped Today, at 17:53 0.00 kbit/s (last run)
Backup job for VMware Cloud... 5 Successful Today, at 15:04 0.00 kbit/s (last run)
Backup copy job 5 Successful 07 May 2024 at 21:59 320.33 Mbit/s (last r...
Backup job for Nutanix AHV 5 Successful 03 May 2024 at 17:10 265.75 Mbit/s (last r...
Backup job for Amazon EC2 5 Successful 03 May 2024 at 16:48 12.22 Mbit/s (last run)
Backup job for file share 5 Successful 03 May 2024 at 16:06 0.00 kbit/s (last run)
Backup job for physical machine 5 Successful 03 May 2024 at 11:25 83.73 Mbit/s (last run)
Backup job for VMware Cloud... 5 Successful 03 May 2024 at 11:12 0.00 kbit/s (last run)
Backup job for VMware 5 Successful 02 May 2024 at 15:50 408.43 Mbit/s (last r...

# Managing Job Policies

With policies, you can create rules that easily add matching items to NAKIVO Backup & Replication jobs. For example, you can create a backup job that meets the following criteria: (a) size of VM is more than 4 GB, (b) number of VM CPU sockets is more than 2, and (c) VM name contains "Ubuntu". Any policy is applied to a single job. In the NAKIVO Backup & Replication job wizard, job policy is accessible from the **Policy** view of the Source page.




Every job policy contains at least one rule. Refer to “Managing Policy Rules” on page 385 for details.

Job policies are available for the following job types:

- Backup jobs
- Replication jobs
- Backup copy jobs
- Failover jobs
- Failback jobs
- Several actions of the Site Recovery job

Learn how to save, edit, and remove job policies in these sections:

- Saving Job Policy
- Editing Job Policy
- Removing Job Policy

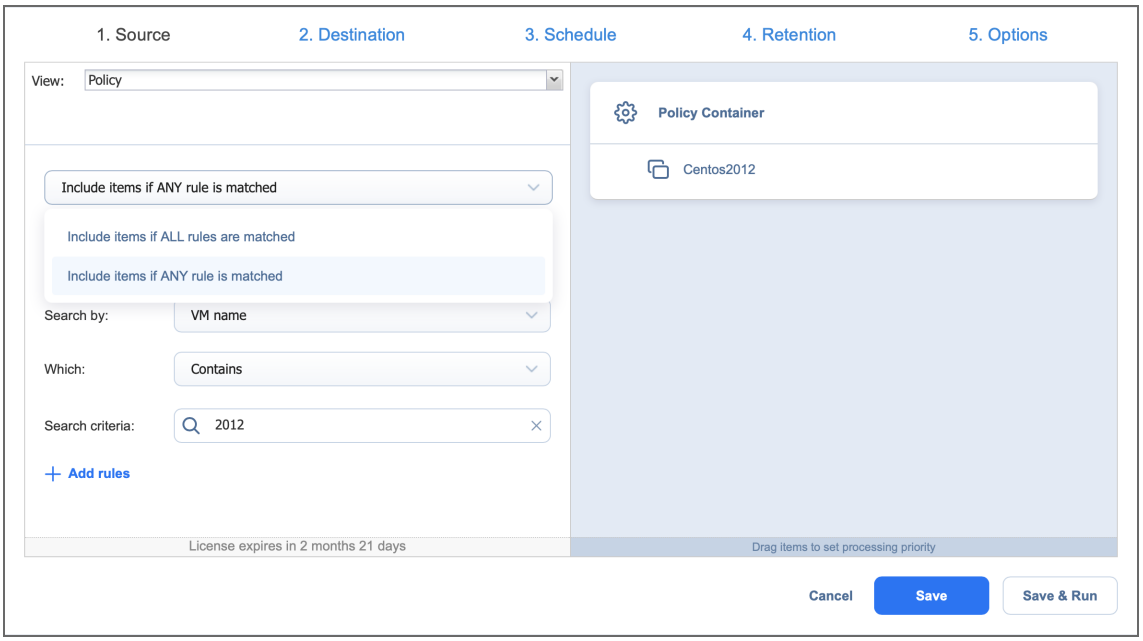
## Saving Job Policy

Follow the steps below to save a policy rule:

1. Make sure your job is opened in the **Policy** view.
2. Choose either of the following **Condition** for your job policy:

- **Include items if ALL rules are matched**: If selected, the logical AND will be applied to the set of policy rules.
- **Include items if ANY rule is matched**: If selected, the logical OR will be applied to the set of policy rules.

3. **Map new VMs/instances/machines to matching backups**: If the checkbox is selected, NAKIVO Backup & Replication maps new workloads, added to the job as compliant to the configured policy rules, to matching backups within the specified destination. This option is only available for VMware/Hyper-V/Amazon EC2/Physical machine backup jobs.
4. Provide the necessary policy rules. Refer to "Managing Policy Rules" on page 385 for details. Make sure that at least one item matches the available set of policy rules.
5. Save your job.




# Editing Job Policy

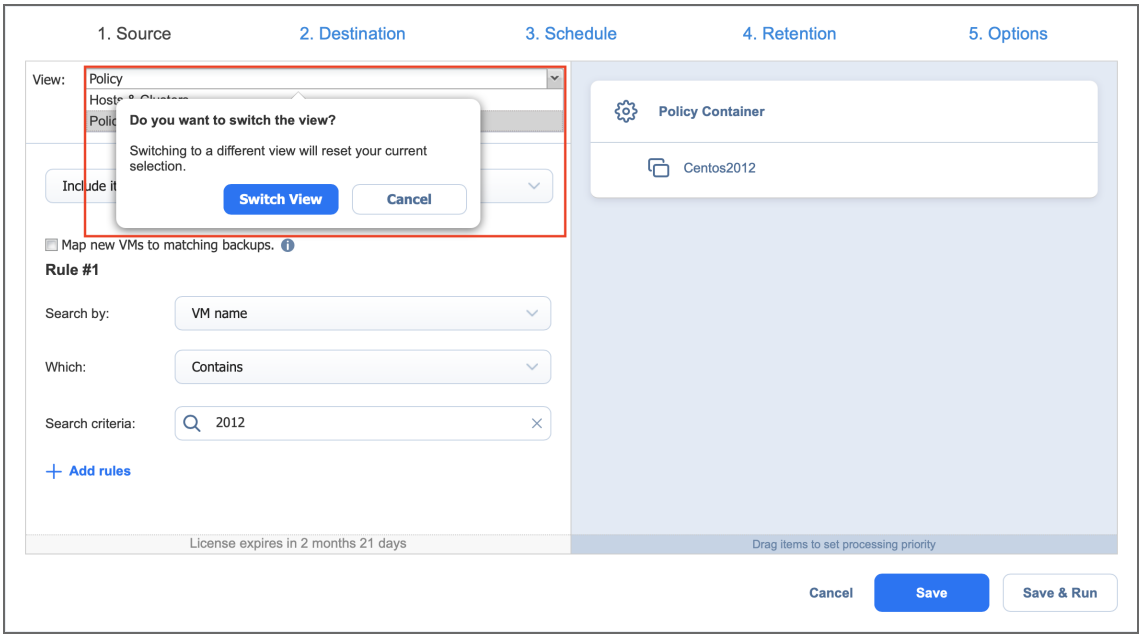
Follow the steps below to edit a job policy:

1. Make sure your job is opened in the **Policy** view.
2. Change the necessary parameters of your job policy:
    1. **Condition**.
    2. Add, edit or delete policy rules. Refer to "Managing Policy Rules" on page 385 for details.
3. Save your job.

# Removing Job Policy

Follow the steps below to remove an entire job policy:

1. Make sure your job is opened in the **Policy** view.
2. Switch to any other inventory view available on the list.
3. A dialog opens warning you that switching to a different view will reset your selection for the current job. Click **Switch View** to confirm your operation.
4. Save your job.



# Managing Policy Rules

Policy rules are an integral part of job policies. Refer to the following sections for details:

- About Policy Rules
- Editing Policy Rules
- Adding Policy Rule
- Removing Policy Rule

## About Policy Rules

In the **Policy** view of the inventory tree, policy rules are numbered by NAKIVO Backup & Replication for your convenience.

Every policy rule contains the following options:

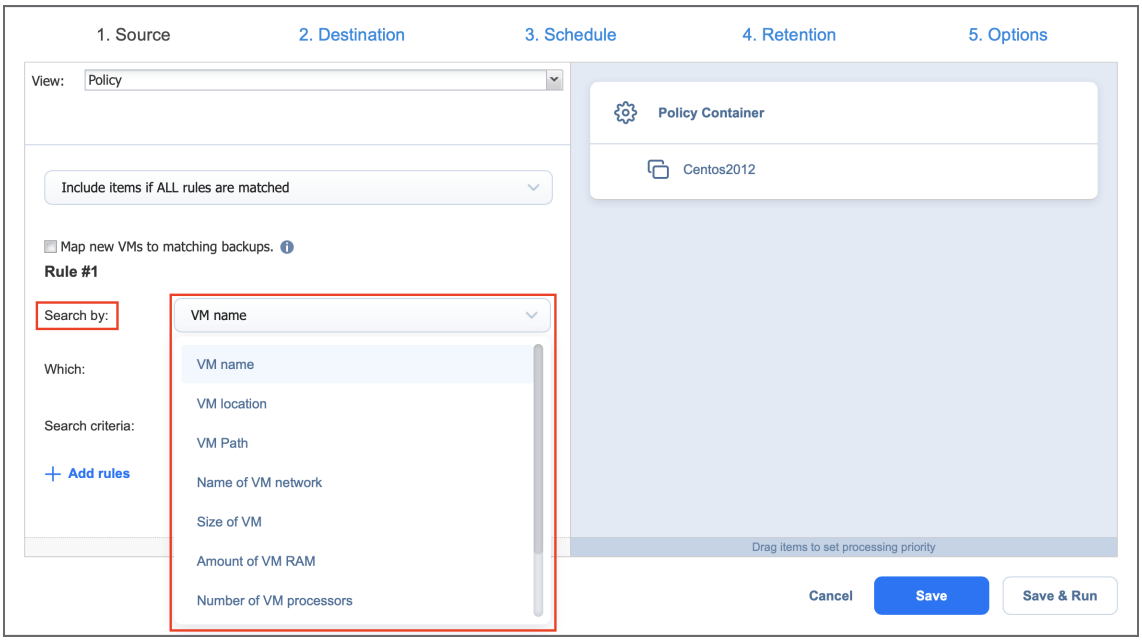
1. **Search by**: A drop-down list with the following search criteria:
    - **VM / VM Template / Instance / Backup / Replica / Machine name**: The rule is to be applied based on the name of the object.
    - **VM / Instance tag**: The rule is to be applied based on the tag of the object.
    - **VM / VM Template / Instance / Replica / Backup location**: The rule is to be applied based on the location of the object.
    - **Name of VM datastore / VM Template datastore / Replica datastore / VM Path / Replica Path / IP address**: The rule is to be applied based on the name of the datastore, path, or IP address.
    - **Name of VM network / VM Template network / Subnet / replica network**: The rule is to be applied based on the name of the network.
    - **Size of VM / VM Template / instance / replica / physical machine**: The rule is to be applied based on the size of the object.
    - **Amount of VM / instance / physical / replica RAM**: The rule is to be applied based on the amount of RAM for a given object.
    - **Number of VM CPU sockets / replica CPU sockets / VM processors / replica processors / Instance virtual CPUs / physical CPUs**: The rule is to be applied based on the number of CPU sockets, processors, or CPUs, depending on the object.

**Note**

The objects corresponding to the above criteria are as follows:

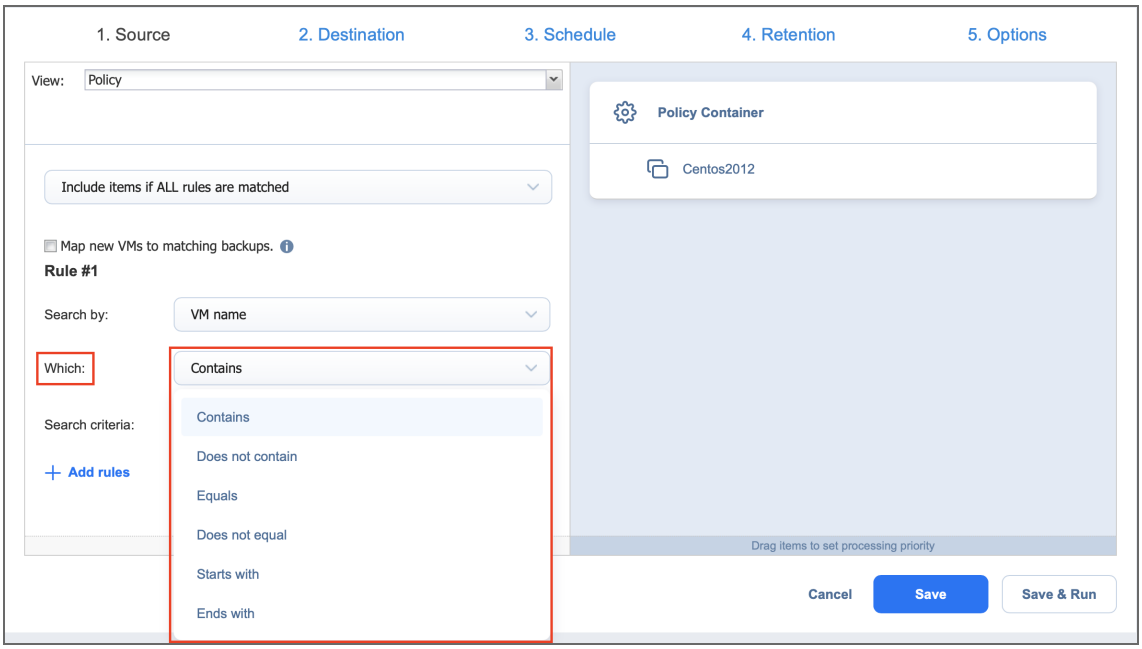
- **VM CPU sockets**: VMware VMs
- **Replica CPU sockets**: VMware VM replicas
- **VM processors**: Hyper-V VMs
- **Replica processors**: Hyper-V VM replicas
- **Instance virtual CPUs**: Amazon EC2 instances
- **Physical CPUs**: Physical machines

- **VM power state / Instance power state**: The rule is to be applied based on the power state of the object.
- **IP Address**: The rule is to be applied based on the IP address of the object.




2. **Search parameter**: You can choose either of the following:
   - For **VM / VM Template / Instance / Backup / Replica / Machine name**, **Name of VM network / VM Template network / Replica network / Subnet / VM datastore / VM Template datastore / Replica datastore**, **VM / VM Template / Replica Path**, **VM / Instance tag**, and **IP Address**:
     - **Contains**
     - **Does not Contain**
     - **Equals** (always applied to the VM tag)
     - **Does not equal**
     - **Starts with**
     - **Ends with**

- For **Amount of VM / Instance / Physical / Replica RAM**, **Number of VM CPU sockets / replica CPU sockets / VM processors / replica processors / instance virtual CPUs / physical CPUs**, and **Size of VM / VM Template / Instance / Replica / physical machine**, you can choose any of the following search parameters:
  - **Is more than**
  - **Is less than**
  - **Equals**
  - **Does not equal**
- For **VM / Instance power state** and **VM / VM Template / Instance / Replica / Backup location**:
  - **Is**
  - **Is not**




3. **Search criteria**: A text string or a numeric value to be used by the policy rule.

When you enter or edit parameters, the changes are immediately reflected in the list of selected items.

## Editing Policy Rule

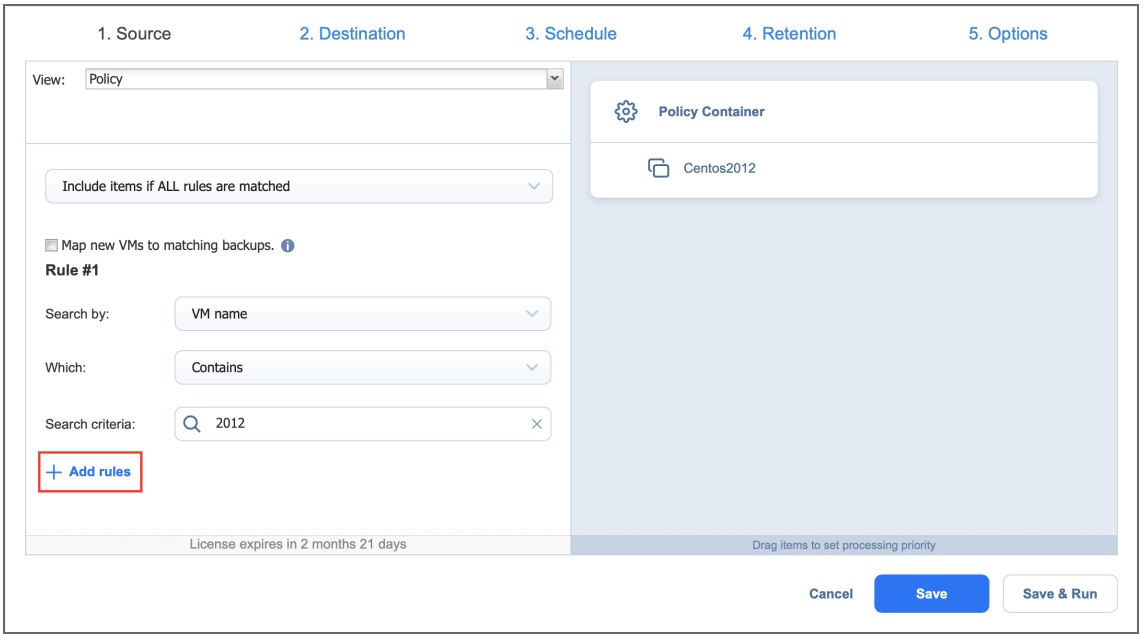
Follow the steps below to edit a policy rule:

1. Make sure your job is opened in the **Policy** view.
2. Locate your policy rule in the left pane of the view. If necessary, use the scroll bar.
3. Change the necessary parameters of your policy rule. Make sure that at least one item matches an available set of policy rules.
4. Click **Next**.

# Adding Policy Rule

Follow the steps below to add a policy rule:

1. Make sure your job is opened in the **Policy** view.
2. In the left pane of the wizard, click **Add rules**.
3. The wizard displays a new policy rule, *Rule #N*. Provide the necessary parameters of your new policy rule. Make sure that at least one item matches the available set of policy rules.
4. Click **Next** when all parameters are set.

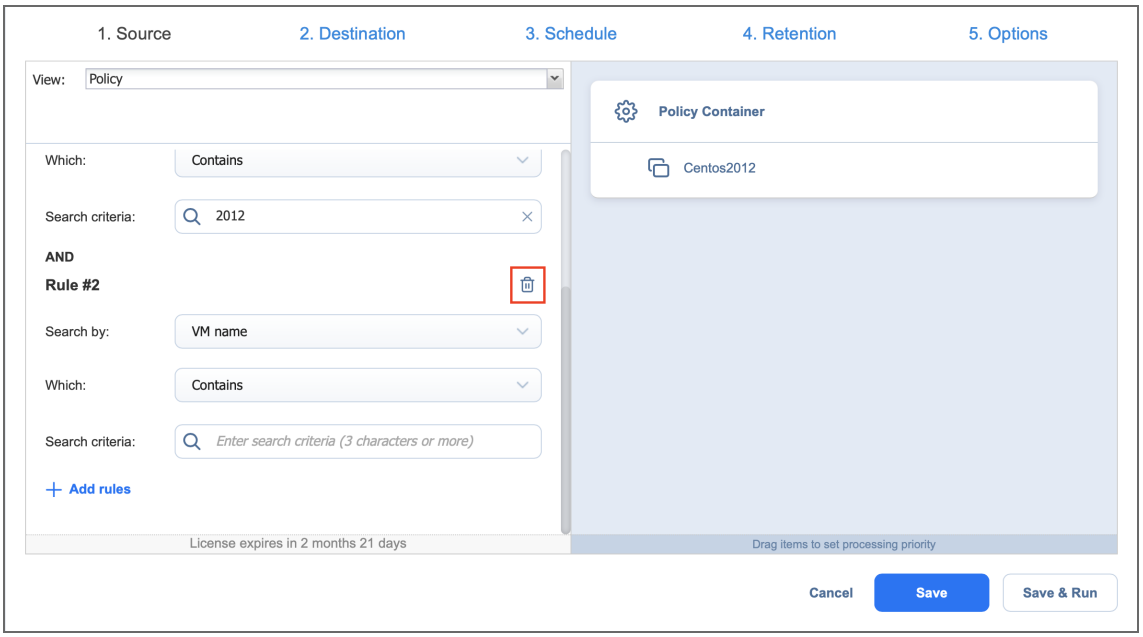



# Removing Policy Rule

Follow the steps below to remove a policy rule:

1. Make sure your job is opened in the **Policy** view.
2. Locate your policy rule in the left pane of the view. If necessary, use the scroll bar.
3. Hover over the rule you would like to remove to reveal the **Remove** icon to its right.

4. Click **Next** when all parameters are set.




**Note**

You cannot remove all policy rules. A job policy must have at least one rule.

# Job Alarms and Notifications

NAKIVO Backup & Replication displays:

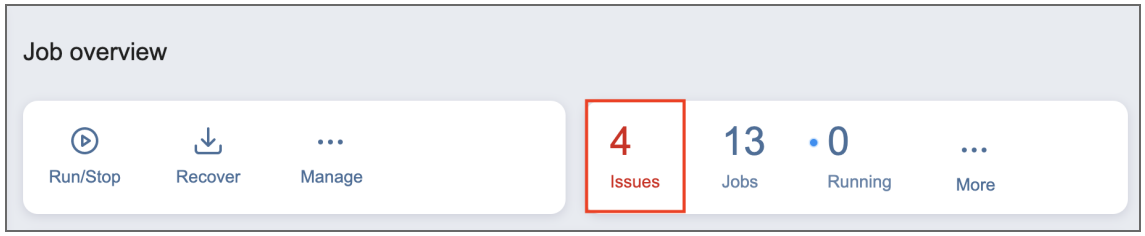
- **Alarms:** Job failures
- **Notifications**: Infrastructure changes and minor errors that do not lead to processing failure

For details, refer to the following sections:


- Viewing Alarms and Notifications
- Dismissing Alarms and Notifications

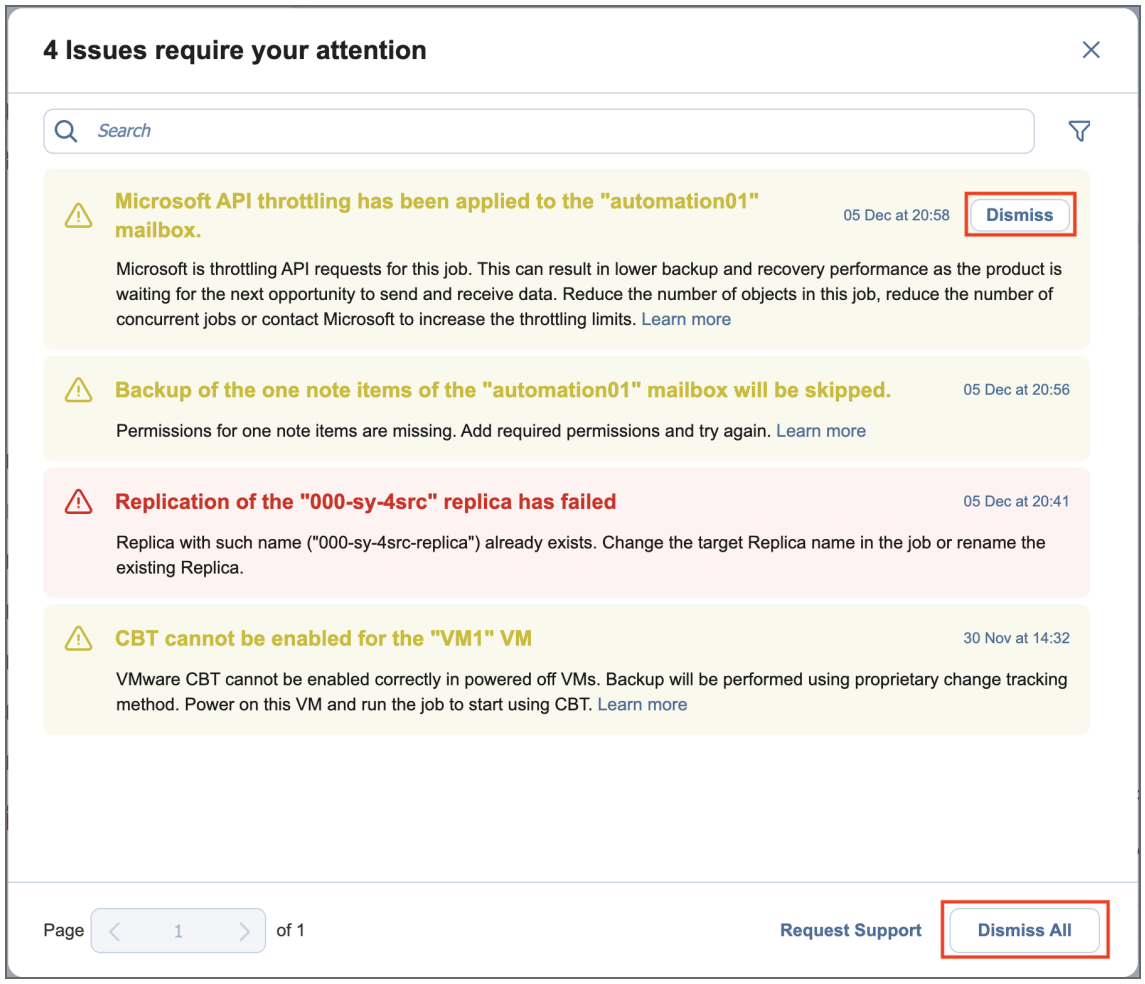

## Viewing Alarms and Notifications

To view alarms and notifications, click the red **Issues** number in the **Summary bar**.




## Dismissing Alarms and Notifications

To dismiss all alarms and notifications in a job or selected group, click **Dismiss All**. To dismiss an individual alarm or notification, hover the mouse pointer over the alarm or notification and click **Dismiss**.


4 Issues require your attention
Search
Microsoft API throttling has been applied to the "automation01" mailbox.
05 Dec at 20:58
Dismiss
Microsoft is throttling API requests for this job. This can result in lower backup and recovery performance as the product is waiting for the next opportunity to send and receive data. Reduce the number of objects in this job, reduce the number of concurrent jobs or contact Microsoft to increase the throttling limits. Learn more
Backup of the one note items of the "automation01" mailbox will be skipped.
05 Dec at 20:56
Permissions for one note items are missing. Add required permissions and try again. Learn more
Replication of the "000-sy-4src" replica has failed
05 Dec at 20:41
Replica with such name ("000-sy-4src-replica") already exists. Change the target Replica name in the job or rename the existing Replica.
CBT cannot be enabled for the "VM1" VM
30 Nov at 14:32
VMware CBT cannot be enabled correctly in powered off VMs. Backup will be performed using proprietary change tracking method. Power on this VM and run the job to start using CBT. Learn more
Page
1
of 1
Request Support
Dismiss All

# Managing Activities

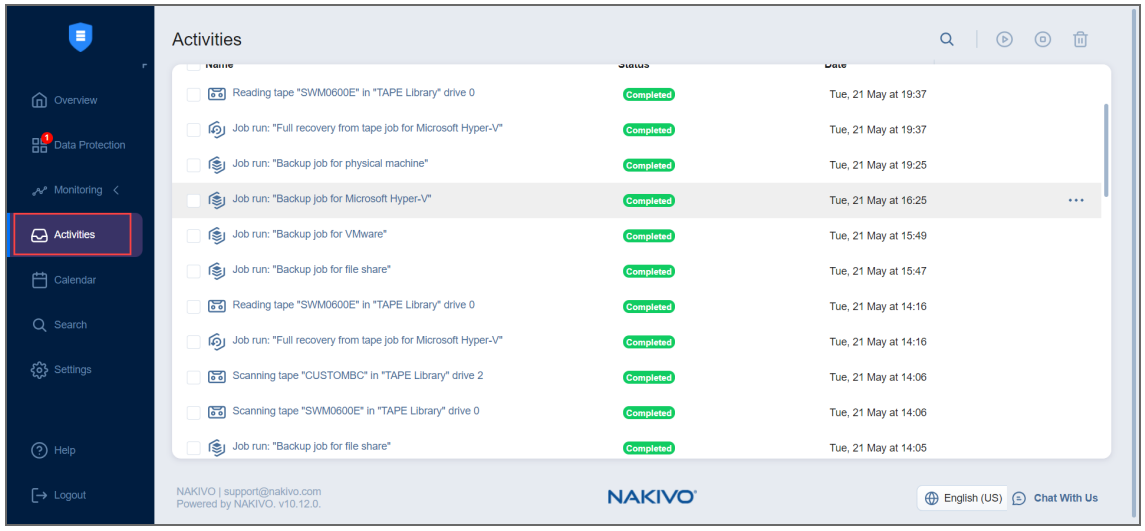
The **Activities** page displays current and past tasks performed by NAKIVO Backup & Replication. From this dashboard, the following actions can be done:

- Viewing Activities
- Searching for Activities
- Viewing Activity Details
- Stopping Running Activities
- Running Activities Again
- Removing Activities

Past activities are stored for the number of days specified in the **Store job history for the last X days** setting in the General tab.

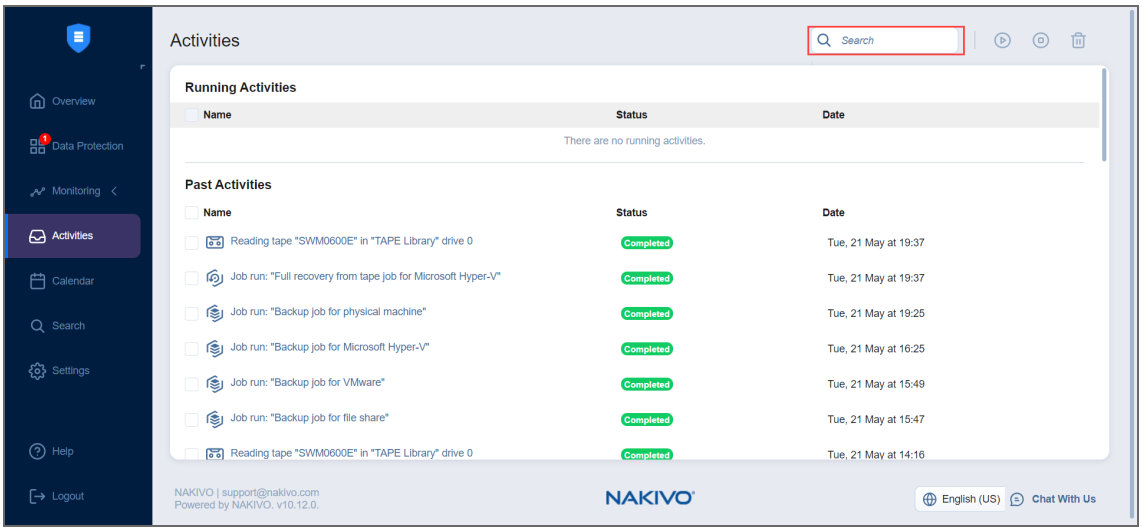
## Viewing Activities

The **Activities** dashboard allows viewing all your current and past activities in the application.



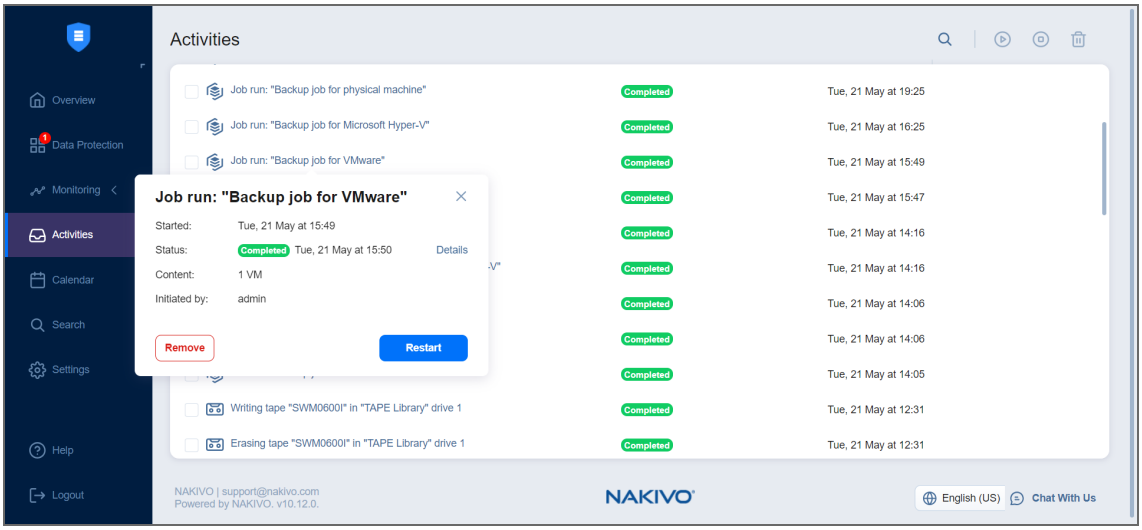

## Searching for Activities

Find activity by typing in part of its name in the **Search** field.



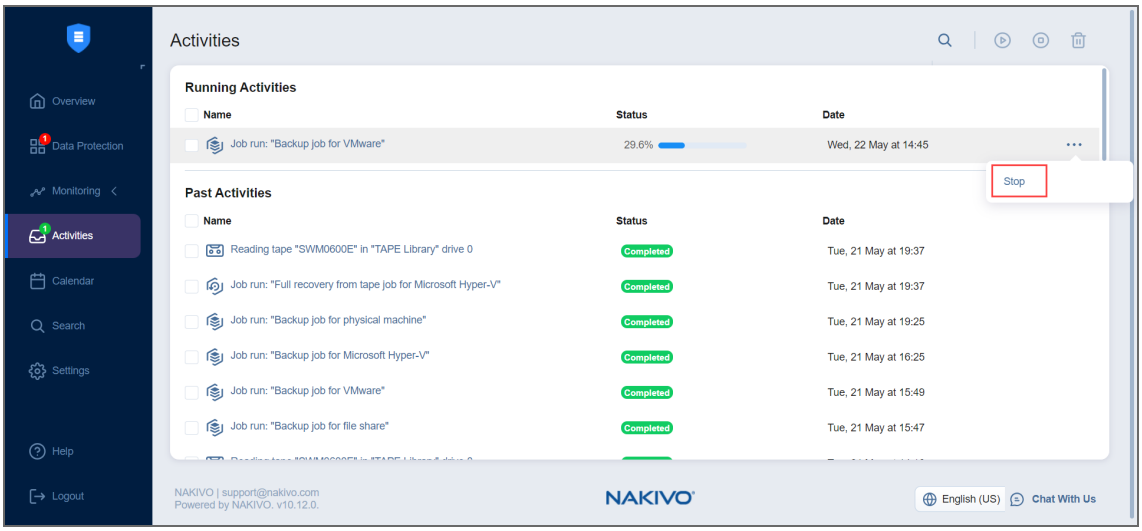

# Viewing Activity Details

View the details of an activity by selecting an activity name. Note that you can view job run details by clicking the **details** link.



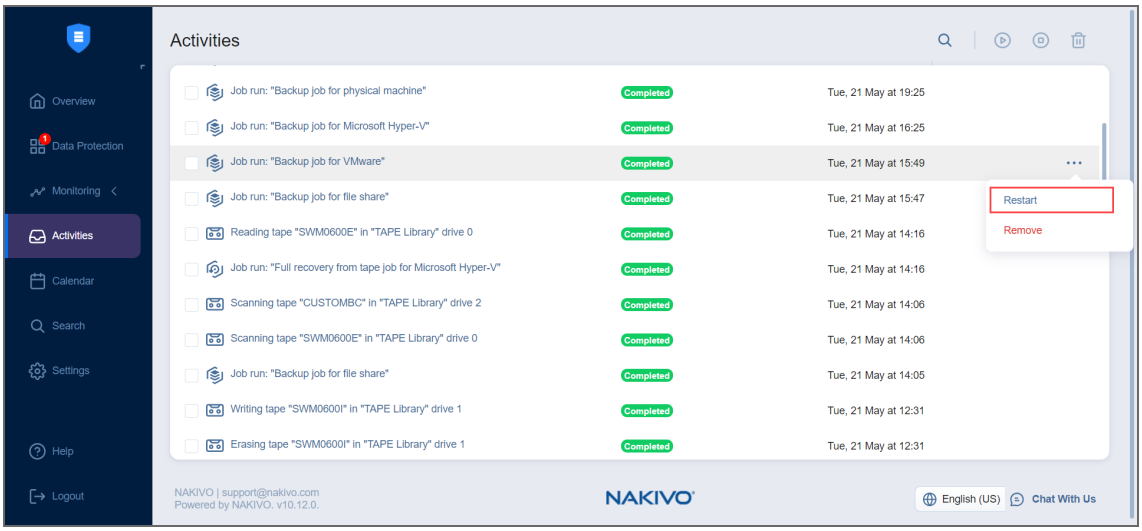

# Stopping Running Activities

To stop running activities, tick the checkbox next to each desired activity and click **Stop** in the toolbar above. To stop all running activities, tick the **Select/Deselect all** checkbox at the top and click **Stop**. You can also stop a single activity by clicking the **Stop** icon that appears when you hover over a specific running activity.



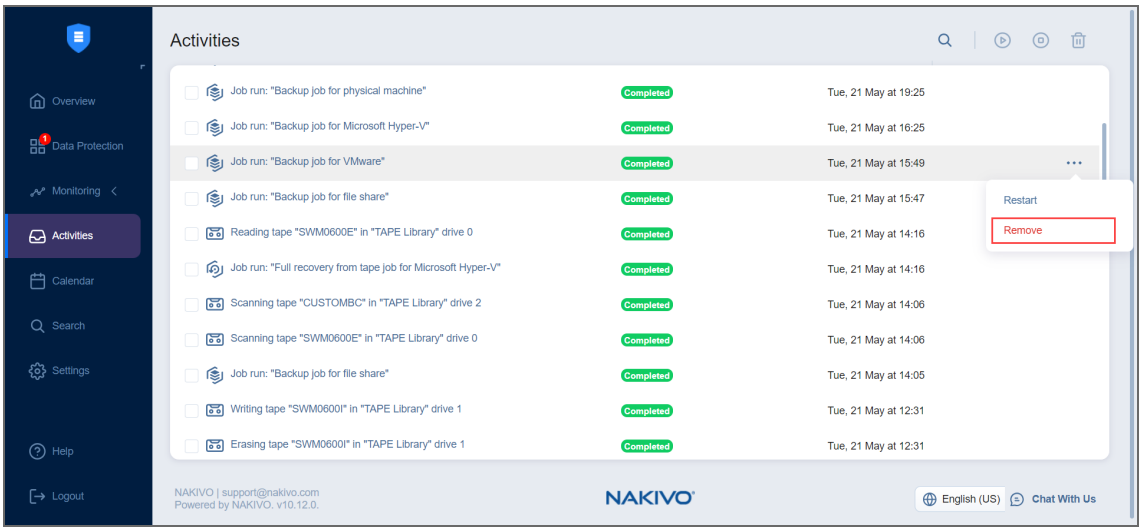

# Running Activities Again

To run activities again (if possible), tick the checkbox next to each desired activity and click **Start** in the toolbar above. To run all activities again at once, tick the **Select/Deselect all** checkbox at the top and click **Start**. You can also run a single activity by clicking the **Start** icon that appears when you hover over a specific activity.




# Removing Activities

To remove activities from the list, select the checkbox next to an activity, then click "**...**" and select **Remove**. To remove all activities from the list at once, tick the checkbox next to the **Name** and click **Remove** icon..


Activities
Overview
Data Protection
Monitoring
Activities
Calendar
Search
Settings
Help
Logout
Job run: "Backup job for physical machine" Completed Tue, 21 May at 19:25
Job run: "Backup job for Microsoft Hyper-V" Completed Tue, 21 May at 16:25
Job run: "Backup job for VMware" Completed Tue, 21 May at 15:49
Job run: "Backup job for file share" Completed Tue, 21 May at 15:47
Reading tape "SWM0600E" in "TAPE Library" drive 0 Completed Tue, 21 May at 14:16
Job run: "Full recovery from tape job for Microsoft Hyper-V" Completed Tue, 21 May at 14:16
Scanning tape "CUSTOMBC" in "TAPE Library" drive 2 Completed Tue, 21 May at 14:06
Scanning tape "SWM0600E" in "TAPE Library" drive 0 Completed Tue, 21 May at 14:06
Job run: "Backup job for file share" Completed Tue, 21 May at 14:05
Writing tape "SWM0600I" in "TAPE Library" drive 1 Completed Tue, 21 May at 12:31
Erasing tape "SWM0600I" in "TAPE Library" drive 1 Completed Tue, 21 May at 12:31
Restart
Remove
NAKIVO | support@nakivo.com
Powered by NAKIVO. v10.12.0.
NAKIVO
English (US)
Chat With Us

# Using Calendar

The **Calendar** allows you to schedule and view the history of past job runs.

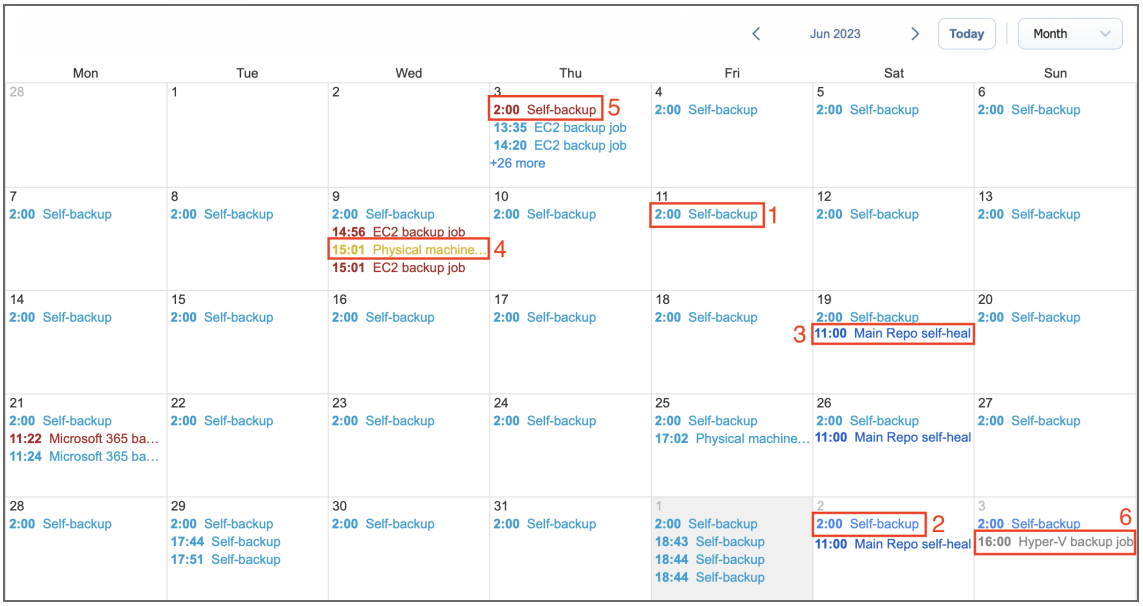
- Understanding Calendar Formatting
- Creating Jobs with Calendar
- Editing Jobs with Calendar

## Understanding Calendar Formatting

Jobs in the **Calendar** view are formatted by start/end time and color coded by status. The color coding format is as follows:

1. Successful job runs are marked in teal.
2. Future scheduled job runs and currently running jobs are marked in sky blue.
3. Repository maintenance jobs (such as scheduled self-healing) are marked in navy blue.
4. Stopped job runs are marked in yellow.
5. Failed job runs are marked in red.
6. Job runs belonging to disabled jobs are marked in gray.

> **Note**
> Job runs that complete later than their start date are marked in the Calendar for the appropriate number of days. In **Month** view, such jobs are also marked with background fill. For example, a job that started on a Monday and finished on a Wednesday will be marked in one continuous solid light blue line across three days.



# Creating Jobs with Calendar

To create a job:

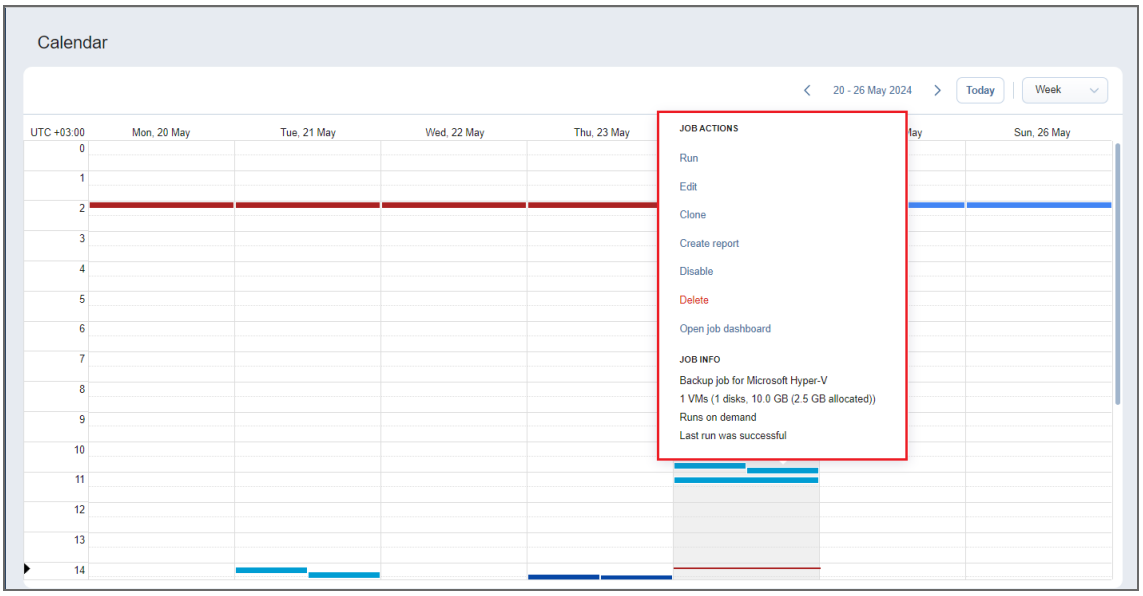
1. Click on the date and time when you'd like to run the job
2. Select the type of job you need.
3. On the **Schedule** page of the wizard, the time you've selected in the **Calendar** will be selected.

# Editing Jobs with Calendar

If you click on the job title on the **Calendar** dashboard, the **Job Actions** and **Job info** menus appear.

Using them, you can:

- Run a job on demand.
- Edit a job.
- Clone a job.
- Delete a job. If the job repeats on schedule, this action will affect all job runs.
- Disable/Enable a job. If the job repeats on schedule, this action will affect all job runs.
- Open the **Job Dashboard**.
- Create a report.
- View the job info.



# Using Global Search

Using the **Global Search** dashboard, search for items within the entire inventory of NAKIVO Backup & Replication, Transporters, Backup Repositories, jobs, backups, and replicas.

- Opening Global Search
- Running Global Search
- Filtering Search Results
- Applying Bulk Action
- “Searching Indexed Files” on page 402
- Viewing Object Info
- Viewing Group Info

**Note**

When the multi-tenant mode is enabled, **Global Search** operates within a specific tenant. For more information about multi-tenancy in NAKIVO Backup & Replication, please consult with the following resources:

- “Multi-Tenant Deployment” on page 217
- “Multi-Tenancy” on page 104
- “Multi-Tenant Mode” on page 1254

## Opening Global Search

To open **Global Search**, click the **Search** icon in the main toolbar of the application. The **Global Search** dashboard opens with the **Infrastructure** tab selected by default.

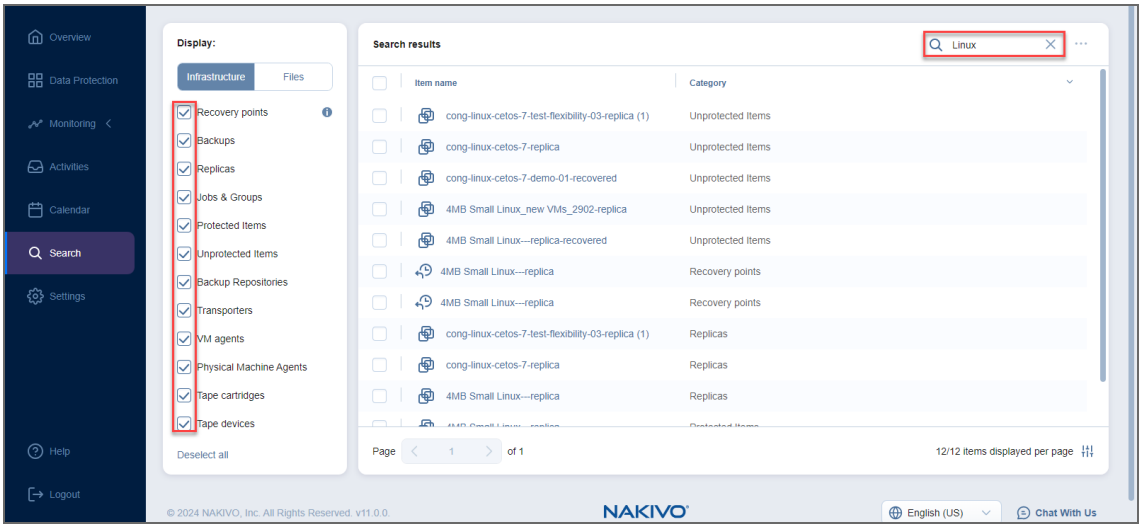


# Running Global Search

Click the **Search** button and enter your search string into the search box. The search results are immediately returned in the form of a list.

By default, your search results are non-filtered. This means that the search is applied to all categories of NAKIVO Backup & Replication objects that are enabled by default.

> **Note**
>
> To fine-tune search results, you can use the following wildcards:
>
> - "**?**" representing a single character.
>
>  "*****" representing zero or more characters.




Please note the following:

- Search is case insensitive.
- Search results are grouped by categories and the items within the same category are sorted alphabetically.
- Sorting by the Item name column is applied if multiple items of the same category are returned.

# Configuring Table Columns

You can show/hide columns or modify the number of items per page in the Search Results table. In the lower right corner, click the controls icon. In the dialog window that opens, select/deselect checkboxes and click **Apply**. The following options are available:

- Item name
- Category
- Path
- Recovery Point Date

**Note**

You cannot modify the **Item name** and **Category** columns in this dialog window.

Click **Cancel** to exit or **Reset Settings** to reset to default settings.

## Filtering Search Results

To narrow your search results, you can deselect some of the categories listed in the Infrastructure tab:

- Recovery points
- Backups
- Replicas
- Jobs & Groups
- Protected Items
- Unprotected Items
- Backup Repositories
- Transporters
- VM agents
- Physical Machine Agents
- Tape cartridges
- Tape devices

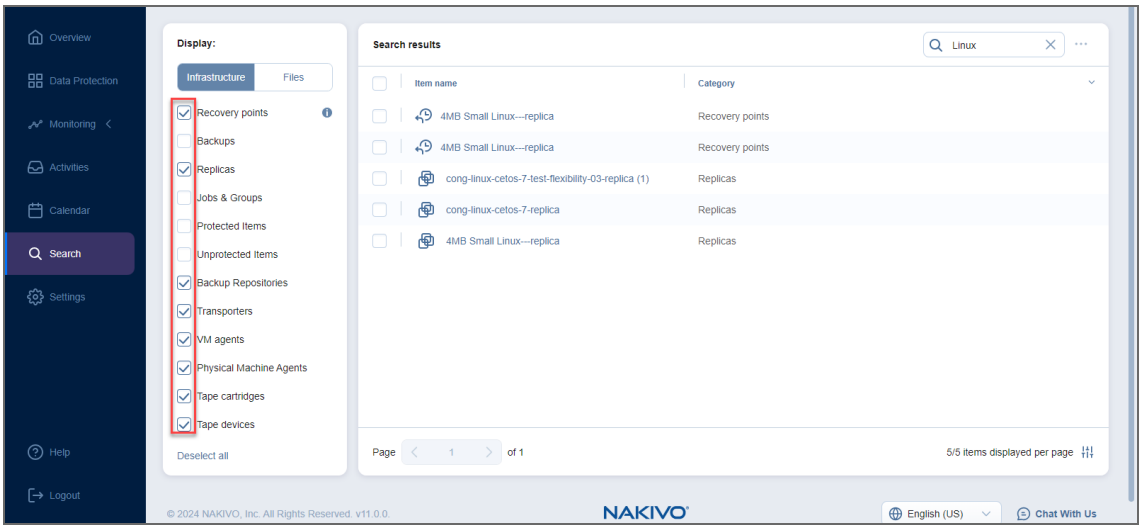
**Notes**

- When the search matches the recovery point name, all chains of recovery points of the source backup object to which this recovery point belongs are returned.

  **Exception**

  Self-backup recovery points are not returned in the search result.
- If you have added a federated repository to the product, **Global Search** displays the federated repository name in the search result instead of its members.
- The following fields are not displayed in the federated repository popup:
  - Backup repository path
  - Backup repository compression
  - Transporter
  - Deduplication
  - Space savings

The filtered search results are displayed immediately as the search results list.



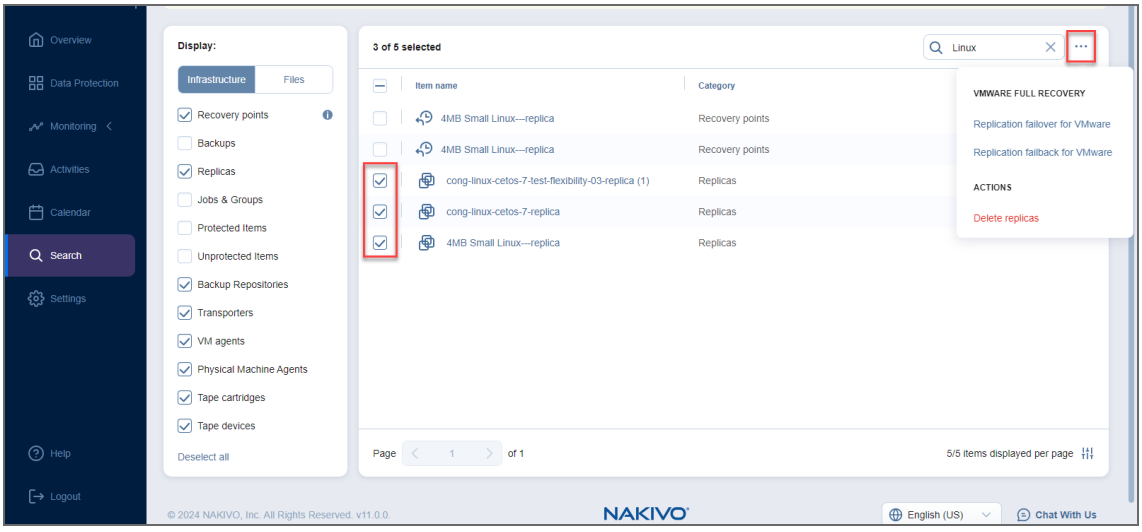

To get back to the default filtering settings, first click **Deselect all** and then **Select all** below the categories list.

# Applying Bulk Action

With NAKIVO Backup & Replication Global Search, you can apply a bulk action to objects belonging to the same category and of the same type.

Proceed as follows to apply a bulk action:

1. In the search result list, select similar objects.
2. The **Bulk Action** button becomes active in the upper right corner of the dialog. Click **Bulk Action**.



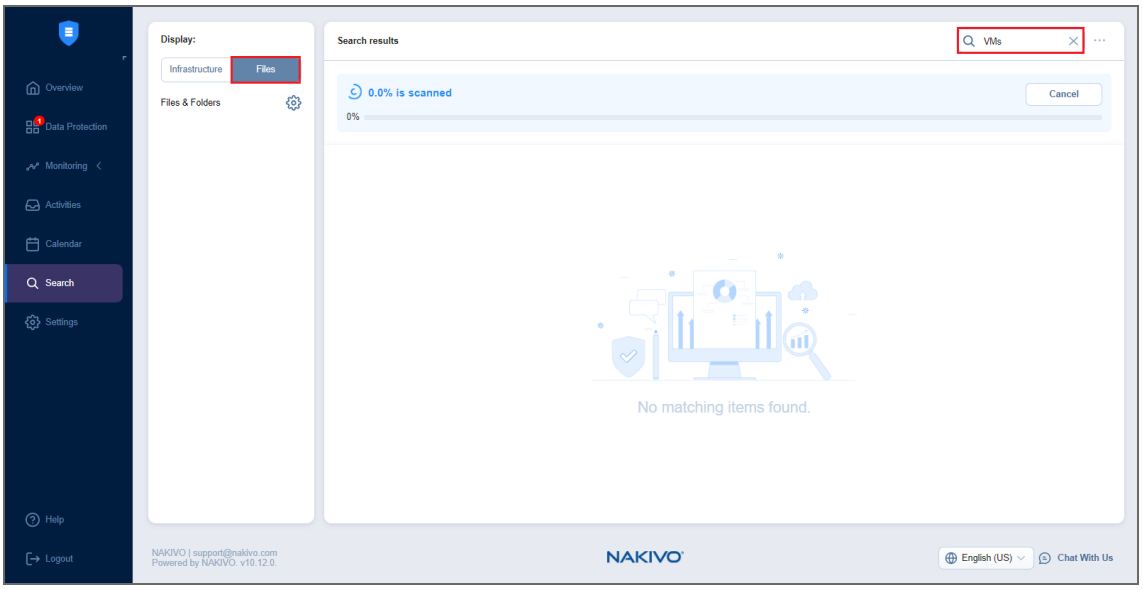

3. A dialog opens with the list of actions applicable to the selected items. To proceed with the necessary action, click the corresponding item in the list of actions.

**Notes**

- Bulk actions do not apply to NAKIVO Backup & Replication dissimilar objects.
- Bulk actions do not apply to items with mixed destinations (disk and tape).

## Searching Indexed Files

To search for a file or folder, switch to the **Files** tab. Click the **Search** button and enter your search string into the search box.

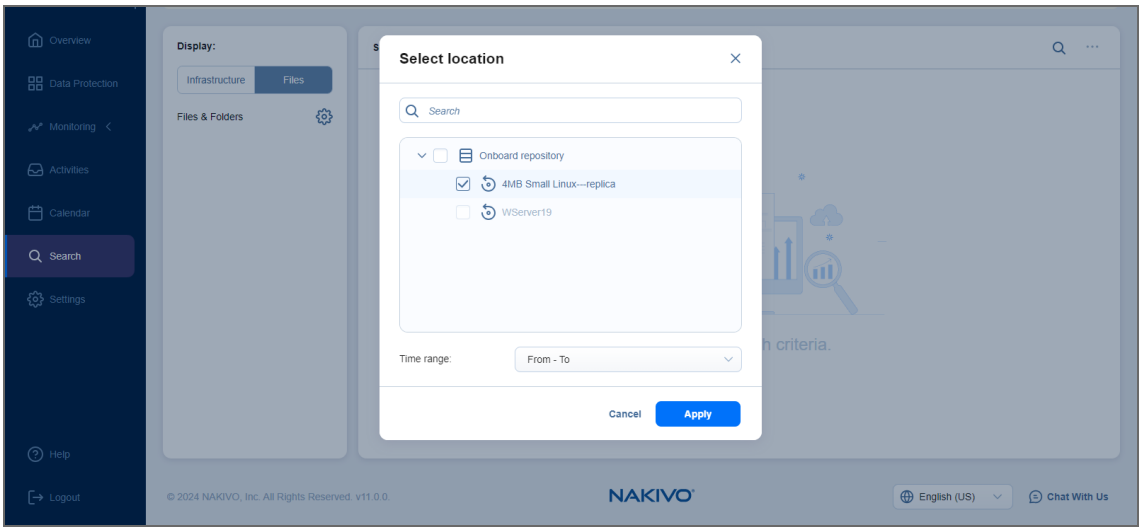



The **Search results** pane displays a progress bar (displayed when scanning is enabled only) and the results of the search.

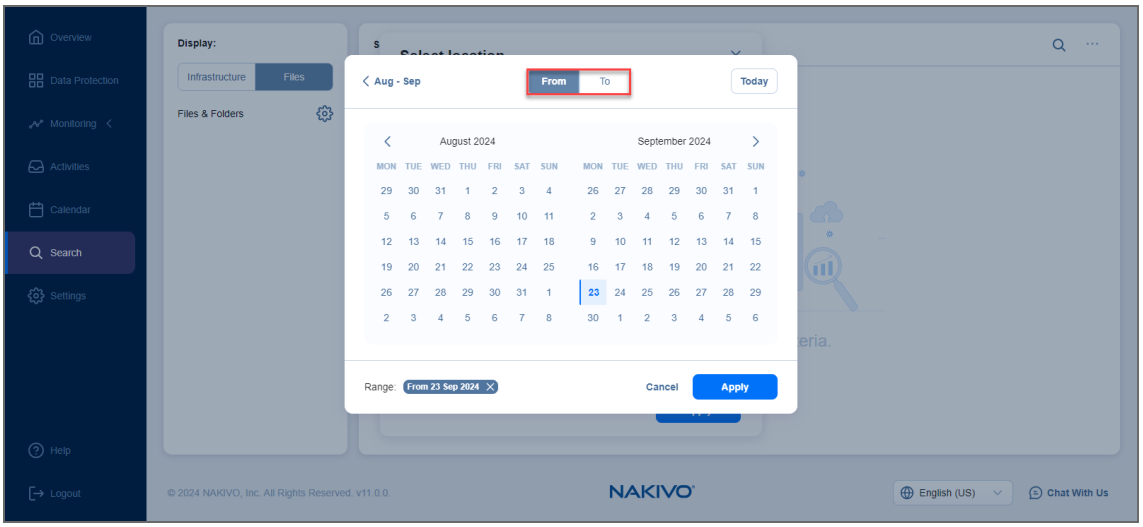
**Notes**

- You must enable **File System Indexing** for VMware vSphere backup jobs to allow the search for specific files and folders inside a VM backup.
- The **File System Indexing** search results for encrypted backups are not displayed if the password hash is not available for these backups.
- If the backups subject to scanning contain encrypted recovery points that cannot be decrypted automatically, these recovery points are skipped and the warning dialog box is displayed before scanning.
  - Click **Cancel** to dismiss the dialog and clean up the search criteria in the search field.
  - Click **Proceed** to proceed to scanning.
- Refer to Managing Backup Encryption for more details.

You can also narrow your search results in the **Files** tab. To do that, click the gear icon next to **Files & Folders** to open the **Select location** popup.




Open the date picker dropdown and set the needed search range by navigating between the **From** and **To** panes.



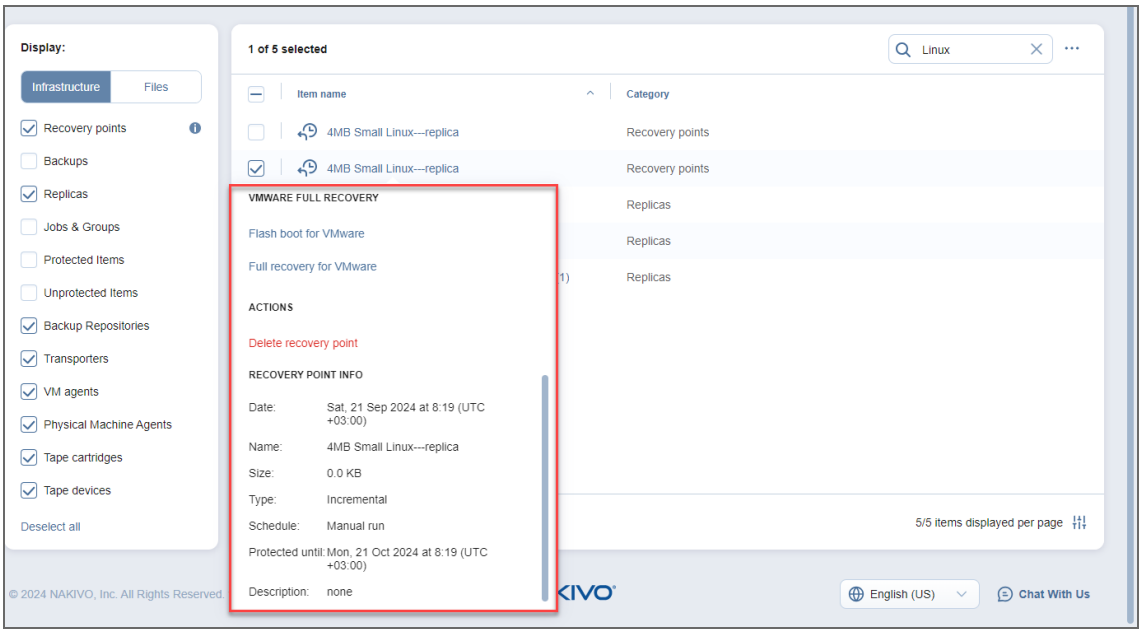

Click **Apply** to filter the search results.

The index of the following items included in the backup job is compressed and recorded into the **Director** file storage to be returned as **Global Search** results:

- File name
- File extension
- File path
- Linked recovery point
- Disk label on OS level
    - Windows OSs: disk name is a drive letter of volume disk (e.g. C, D, E, F, etc.)
    - Linux OSs: disk name is a block device name of disk (e.g: /dev/sda1, /dev/sdb2, etc.).

# Viewing Object Info

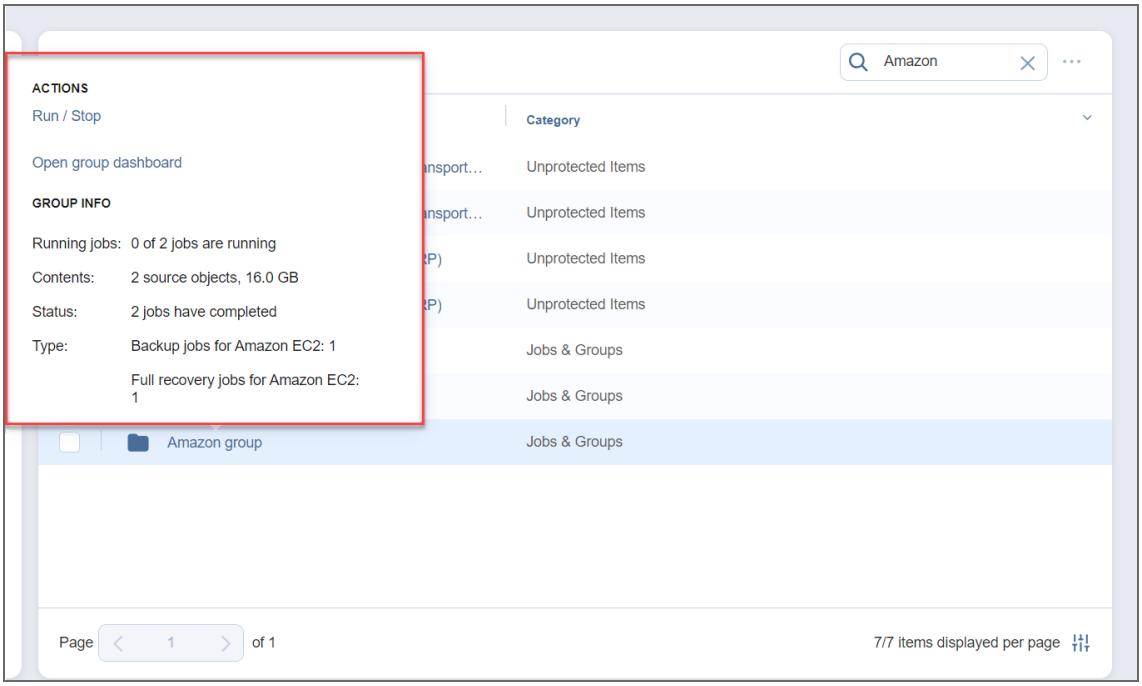
To view info on a specific object available in the search result, select the object.




A dialog opens displaying object info, along with the list of typical actions applicable to the object.

# Viewing Group Info

To view information about a specific group available in the search result, select the group.




A dialog opens displaying group info, along with the list of typical actions applicable to it.

# Settings

This section covers the following topics:


- "General" on page 406
- "Inventory" on page 510
- "Nodes" on page 564
- "Backup Repositories" on page 610
- "Federated Repositories" on page 696
- "Expert Mode" on page 762
- "Maintenance Mode" on page 786
- "Virtual Appliance Configuration" on page 789
- "Multi-Tenant Mode Configuration" on page 814
- "Support Bundles" on page 829
- "Built-in Support Chat" on page 831

# General

This section contains the following topics:


- "Bandwidth Throttling" on page 407
- "Branding" on page 411
- "Configuring Events" on page 413
- "Database Options" on page 414
- "Email Settings" on page 419
- "Licensing" on page 423
- "Managing Backup Encryption" on page 426
- "MSP Tab in Single-Tenant Mode" on page 436
- "Notifications & Reports" on page 444
- "Self-Backup" on page 447
- "Software Update" on page 452
- "System Settings" on page 455
- "Users and Roles" on page 468

# Bandwidth Throttling

With bandwidth throttling settings, you can control the throughput of the data processing by setting specific limits for all or for separate jobs. Bandwidth throttling is managed with bandwidth rules. When a bandwidth rule is applied to your job, the speed of data transfer from source to target will not exceed the specified limit.

To use bandwidth throttling, make sure the following requirements are met:

- At least one bandwidth rule is created.
- At least one source and target endpoint exist.
  - Hypervisor host, **Transporter**, or **Backup Repository** act as a source or target endpoint.
  - The current job object data transfer falls into the bandwidth rule.
- At least one **JODT** exists.

> **Note**
>
> **Source Endpoint** - an endpoint that sends data.
>
> **Target Endpoint** - an endpoint that receives data.
>
> **Job Object Data Transfer** (**JODT**) is a step of a single job object processing which transfers data of the job object from the source endpoint to the target endpoint.

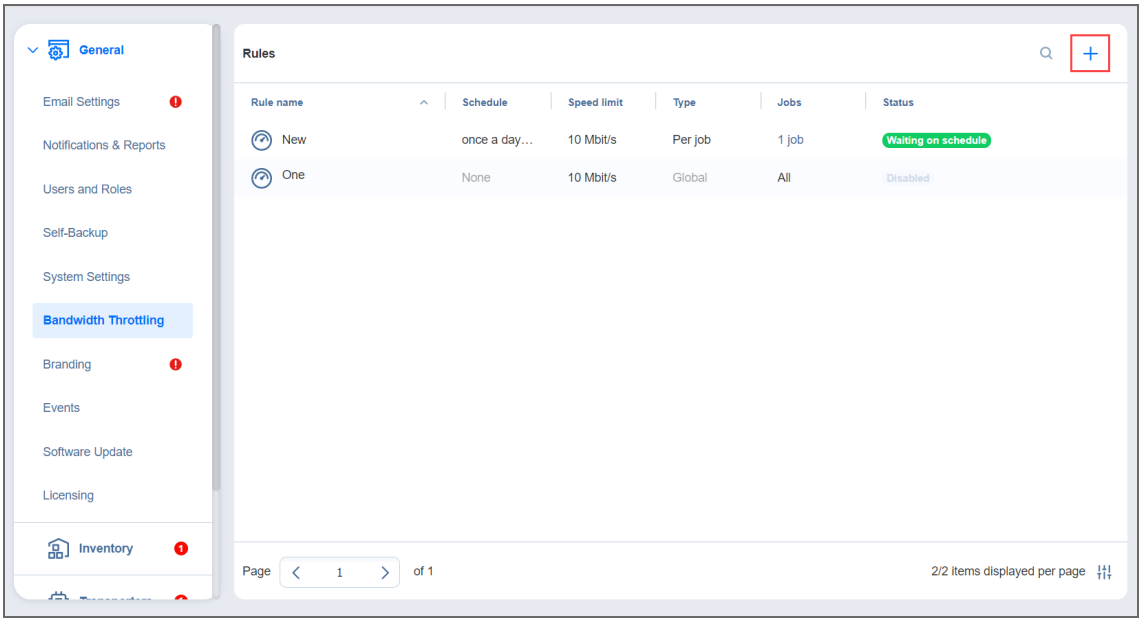
This topic contains the following instructions:

- Accessing Bandwidth Throttling Settings
- Creating Bandwidth Rules
- Managing Bandwidth Rules

## Accessing Bandwidth Throttling Settings

To access bandwidth throttling settings, follow the steps below:

1. Click **Settings** in the left pane of the application to open the **Settings** dashboard.
2. In the **General** tab of the **Settings** dashboard, click **Bandwidth Throttling**. The *Bandwidth Throttling* section opens.




# Creating Bandwidth Rules

Please follow the steps below to create a bandwidth rule:

1. In the *Bandwidth Throttling* section of the **General** tab of **Settings**, click the "**+**" icon.
2. The **New Bandwidth Rule** wizard opens. Proceed as follows:
    a. Choose a type for your bandwidth rule and click **Next**:
        - **GLOBAL**: The rule will be applied to all applicable jobs.
        - **PER-JOB**: The rule will be applied to the selected jobs.

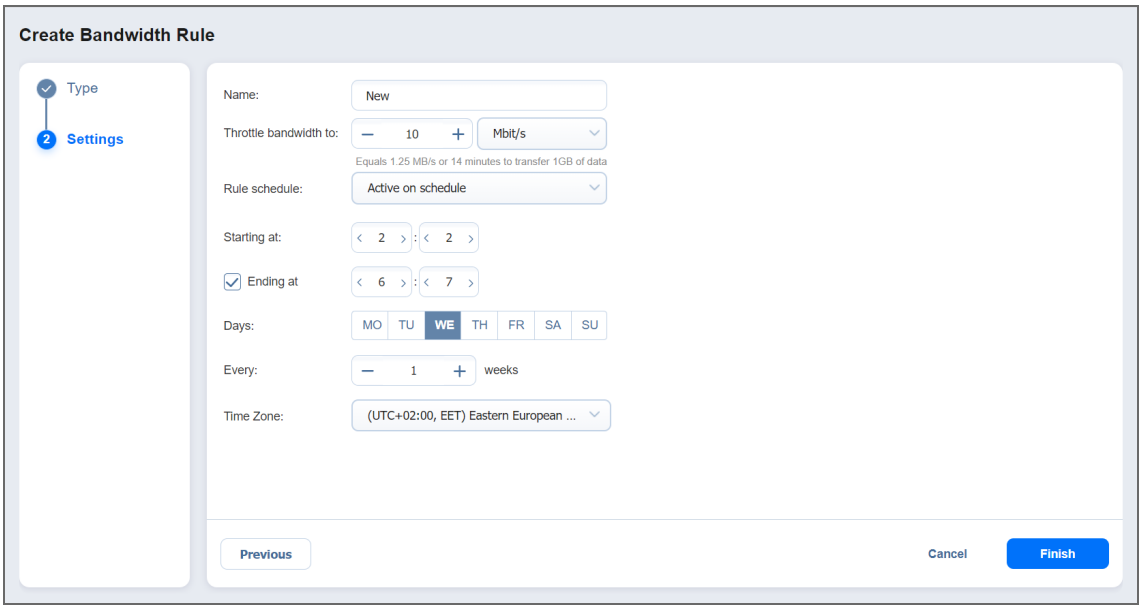
> **Note**
> When applied to specific jobs, **Per-Job** bandwidth rules have higher priority over **Global** bandwidth rules.

    b. **Job**: Choose a job to apply to the **Per-Job** bandwidth rule. Click **Next**.
    c. **Settings**: Configure the following settings:
        a. **Name**: Enter a name for your bandwidth rule.
        b. **Throttle bandwidth to**: Enter the value of the bandwidth limit; and choose the measurement unit: Mbit/s or Gbit/s.

> **Notes**
> - For your convenience, a description is available below the value you've entered, explaining what the value means.
> - In some cases, the actual data transfer speed may exceed the limit you set by up to 0.3 MByte/s or 2.4 Mbit/s.

c. **Rule schedule**: Choose either of the following:
   - **Always active**: The rule will always be active.
   - **Active on schedule**: The rule will be active on schedule. When chosen, the following options are available:
     a. **Starting at** and **ending at**: Enter the time, in hours and minutes, when the rule will be active.
     b. **Days**: Select weekdays for which the rule will be active.
     c. **Time Zone**: Choose a time zone of your rule.
   - **Disabled**: The rule will be disabled.

3. Click **Finish**.




Refer to "Advanced Bandwidth Throttling" on page 68 for more details about bandwidth rules.

> **Note**
> When set, the bandwidth rule will only limit the bandwidth itself and will not affect the number of VMs processed at the same time.

## Managing Bandwidth Rules

You can search for the specific rule by clicking the **magnifying glass** icon in the upper-right part of the screen and entering the name in the search box.

Click on the **ellipsis** to the right of the rule's name to manage bandwidth rules with the following commands:

- **Edit**: The **Edit Bandwidth Rule** dialog opens where you can modify your rule.
- **Disable/Enable**: When applied, the command will disable/enable the rule.
- **Delete**: When applied, a dialog will open asking you to confirm the operation. Click **Delete** to confirm that you wish to delete your rule.

**Note**

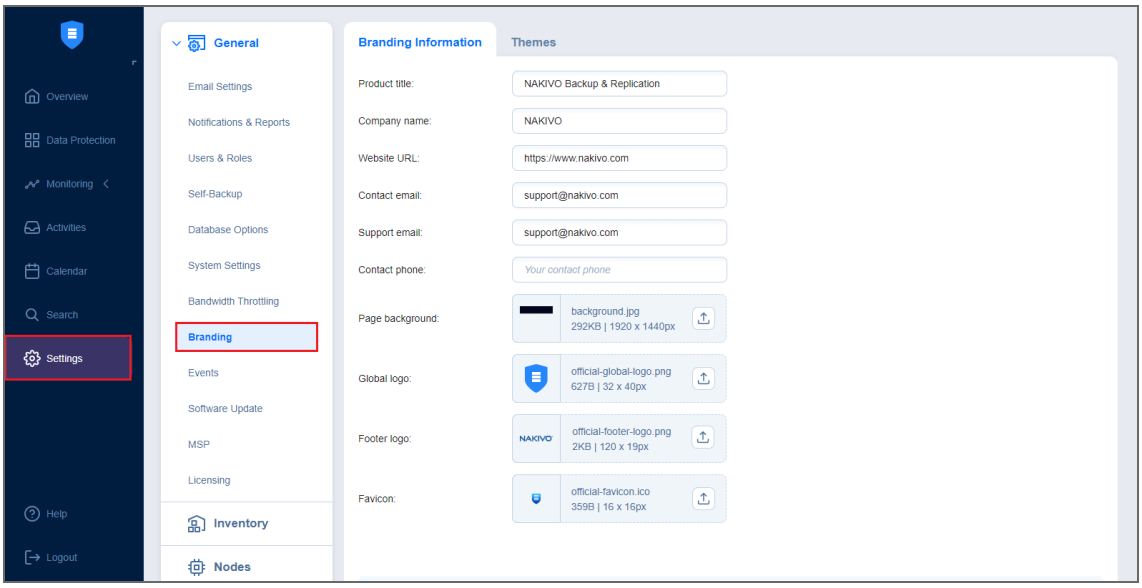
**Per-Job** bandwidth rules can also be created/managed on the **Options** page of the wizard during creating/editing the corresponding jobs. Please refer to the topics:

- “Creating VMware Backup Jobs” on page 839
- “Creating Backup Copy Jobs” on page 872
- “Creating VMware Replication Jobs” on page 934
- “VMware VM Recovery” on page 1061

# Branding

You can change the product branding settings such as product name, logo, background, and so on. To configure these product settings, follow the steps below:

1. Click **Settings** in the left pane of the product.
2. Go to the **General** tab and click **Branding**.




3. Change the following, as appropriate:
   - **Product title**
   - **Company name**
   - **Website URL**
   - **Contact email**
   - **Support email**
   - **Contact phone**
   - **Page background**
   - **Global logo**
   - **Footer logo**
   - **Favicon**
4. On the **Themes** tab, you can configure the colors of your NAKIVO Backup & Replication instance.
5. After making the necessary changes, click **Apply**. Alternatively, click **Discard Changes** to discard any changes you have made.
6. Optionally, click **Reset Settings** to return all the settings to their default values.

During upload, the logo and bookmark icon images are internally resized while preserving the aspect ratio. The background image is used as it is. To get the best image quality, follow the recommendations below.

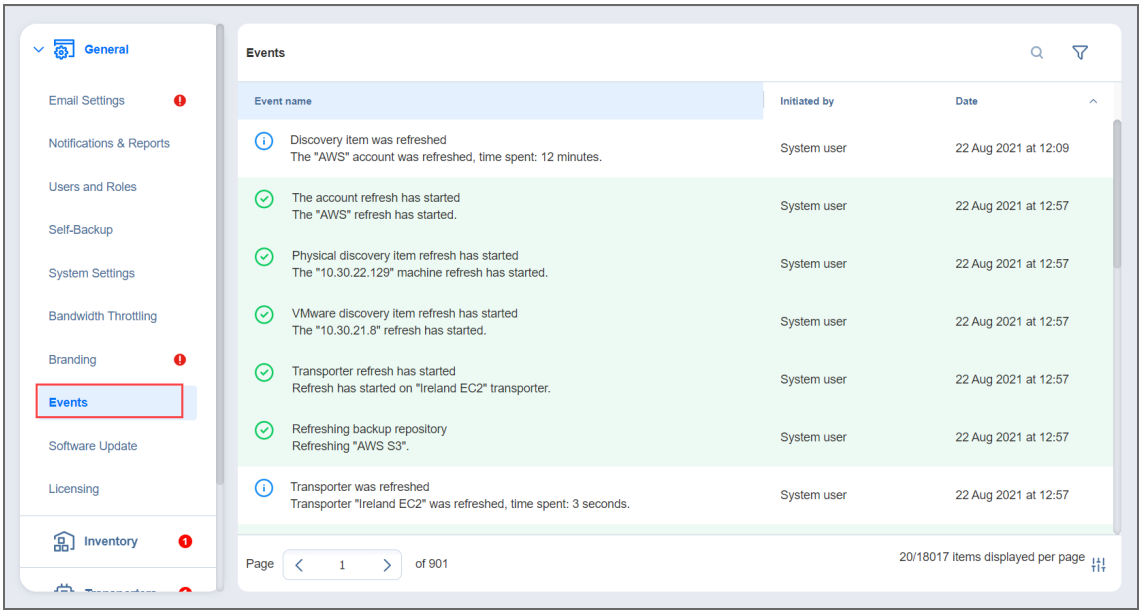
| Image | Best format | Best resolution |
|---|---|---|
| Global logo | .png | 32x40 |
| Footer logo | .png | 32x40 |
| Favicon | .png | 16x16 |

# Configuring Events

NAKIVO Backup & Replication can store and display system events. By default, events are stored for 60 days; you can change the time period in **Settings**.

To view events, follow the steps below:

1. Click **Settings** in the left pane of the product.
2. Open the **General** tab and click **Events**. The **Events** page opens, displaying the NAKIVO Backup & Replication system events.




3. Optionally, you can enter a search string to the **Search** box. This allows you to see events related only to NAKIVO Backup & Replication items – **Transporters**, repositories, jobs, backups, and replicas,– contained in your search string.
4. Optionally, you can select filter the events by the following parameters:
    - **Initiated by**: Select one of the users of the product in the dropbox
    - **Event type**: Choose among the following event types:
        - Info
        - Warning
        - Error
        - Debug
    - **Date**: After selecting this parameter, choose the start and end dates. This allows you to limit the events list within a specific time period.

# Database Options


- System Requirements for PostgreSQL
- Installing PostgreSQL
- Migrating Database


NAKIVO Backup & Replication allows you to migrate the internal H2 database used by the NAKIVO Backup & Replication Director to an external database. To do that, take the following steps:

**Important**

- If you migrate the internal H2 database to an external database, you will not be able to switch back to the internal database or an external database of the same type later.
- For multi-tenant mode, only the Master Admin can perform database migration. The functionality is not available for the local tenant. The remote tenant can still perform database migration as described below if they log in as a single tenant into NAKIVO Backup & Replication.
- The migration occurs for all local tenants at the same time. If the migration fails for one of the tenants, the product reverts to the previous database type automatically.
- Insufficient resources may cause PostgreSQL to crash during operations.
- PostgreSQL can be installed either on the same server as the Director or on a separate server. If you choose to install it on the Director's server, ensure sufficient resources are available to run both services.

## System Requirements for PostgreSQL

1. Make sure that the PostgreSQL server has at least:
   - 4 CPUs and 8GB RAM for Single-Tenant mode
   - 8 CPUs and 36GB RAM for Multi-Tenant mode
2. Set the following parameters in the **postgresql.conf** file:

- for Single-Tenant mode:

```
max_connections = 500
shared_buffers = 256MB
work_mem = 10MB
```

- for Multi-Tenant mode for around 100 tenants including 50 tenants with monitoring data:

```
max_connections = 3000
shared_buffers = 256MB
work_mem = 10MB
```

- for Multi-Tenant mode for around 100 tenants with monitoring data:

```
max_connections = 5000
shared_buffers = 256MB
work_mem = 6MB
```

See also "Feature Requirements" on page 169.

# Installing PostgreSQL

## Installing PostgreSQL on Windows

1. Download the PostgreSQL installer from the official website.
2. Run the installer and follow the instructions.
3. Edit the **postgresql.conf** to set the required system resources configuration.
4. Restart the PostgreSQL service to apply changes.

## Installing PostgreSQL on Linux

PostgreSQL is available integrated with the package management on most Linux platforms.

1. Install the PostgrSQL

   ```
   apt install postgresql
   ```

2. Configure PostgreSQL by editing the **postgresql.conf** to set the required system resources configuration.
3. Restart the PostgreSQL service to apply changes.

   ```
   systemctl restart postgresql
   ```

**Note**

When the PostgreSQL is installed, ensure that the database user has the following privileges:

- Create Database
- Create user with password
- Grant privileges on database to the created user
- Alter the created user

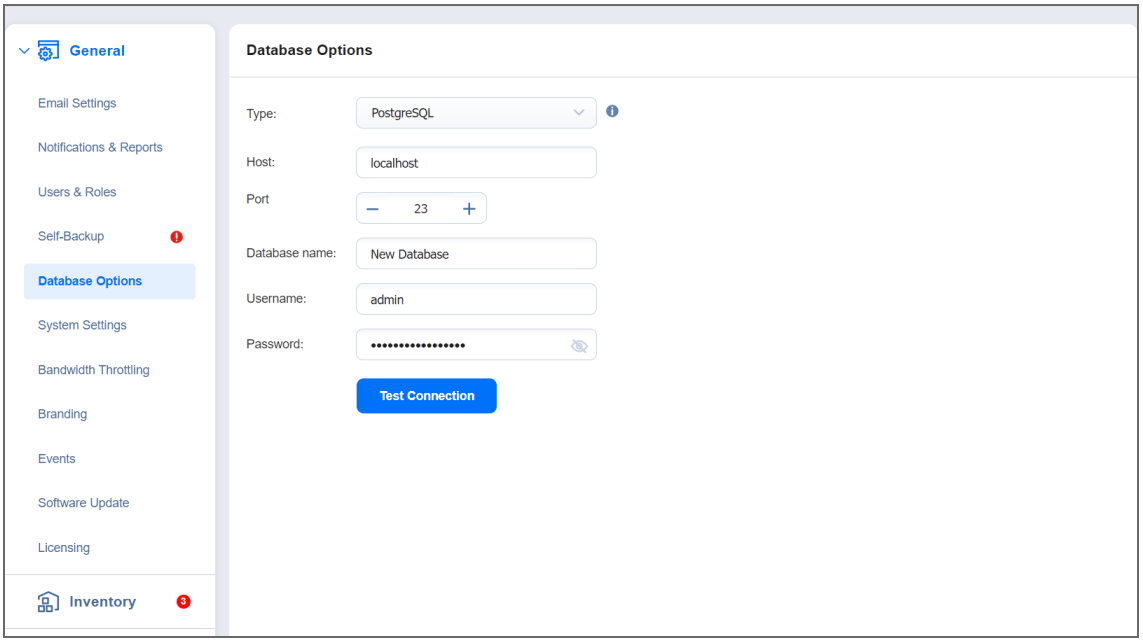
# Migrating Database

1. Go to **Settings > General > Database Options**.
2. Select the external database from the list of supported platforms in the **Type** drop-down list. Note that the internal database is selected by default.
3. In the **Host** field, enter the hostname or IP of the server housing the database.

4. In the **Port** field, enter the relevant port number.
5. Enter the name for your database in the **Database name** field.
6. Enter **Username** and **Password** in the corresponding fields.
7. Click **Test Connection**.
8. If the test is successful, click **Apply Settings**:
    - If the database does not exist, a dialog box appears asking if you would like to create one and proceed with the migration. Click **Migrate**.
    - If the database belongs to the current NAKIVO Backup & Replication installation, a dialog box appears asking if you would like to update the settings of the existing database. Click **Update** to proceed.
    - If the database already exists and is compatible with the current NAKIVO Backup & Replication installation, a dialog box appears asking if you would like to use it, cleanup all its records and replace the contents of the database with the new data. Click **Proceed**.

**Important**

Check ALL your repositories after migration and make sure they are still attached.



**Notes**

- If you have the Self-Backup feature enabled, the self-backup process starts before the database switch and runs again after the switch is completed.
- Self-backup of an external database is possible only with a single-tenant instance of the solution.
- If the external database is installed on another VM or is using an IP address instead of *localhost*, take the following steps before migration:
  1. Open the *pg_hba.conf* file located in the external database installation folder.
  2. Change **IPv4 local connections** settings *from 127.0.0.1/32* to *0.0.0.0/0*.
  3. Save changes.
  4. Restart external database service.
- If the connection between PostgreSQL and NAKIVO Backup & Replication cannot be established, add the following string to the *pg_hba.conf* file:

  *host DATABASE USER ADDRESS METHOD [OPTIONS]*

  *host all all 0.0.0.0/0 md5*

  or

  *host DATABASE USER ADDRESS METHOD [OPTIONS]*

  *host all all 0.0.0.0/0 scram-sha-256*

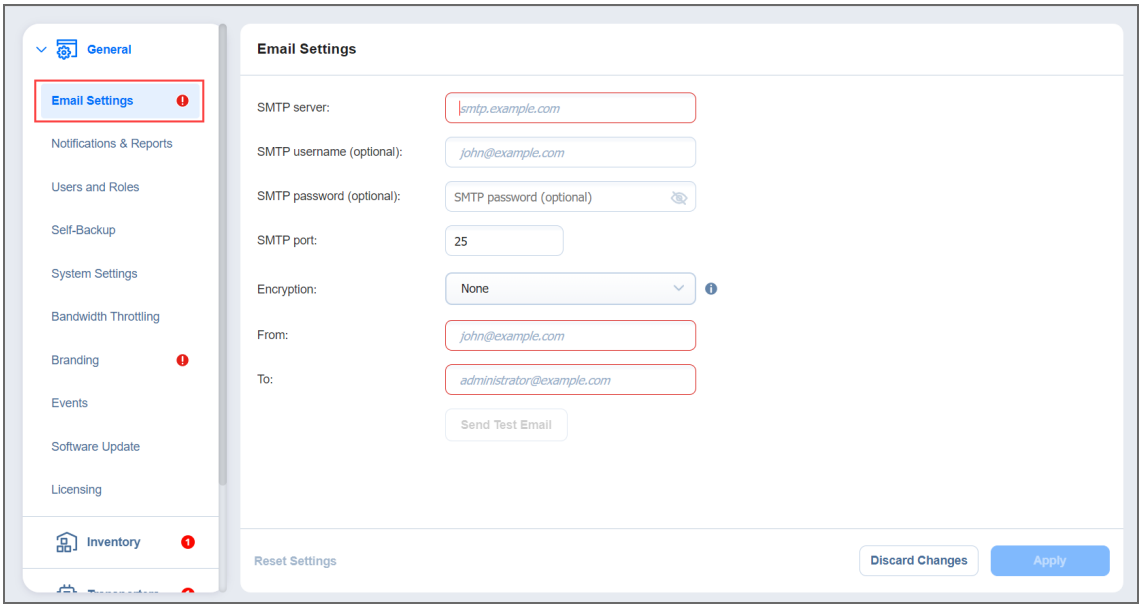
  Note that method (md5/scram-sha-256) may be different for some versions of PostgreSQL. Check the respective method for your version of PostgreSQL before applying the changes.
- If Master Tenant connects to existing database that already houses the data from previous migrations, such database is automatically mapped to the tenants during the new migration using the database UUID.
- It is not possible to recover from a self-backup and system migration in the following cases:
  - The NAKIVO Backup & Replication installation uses the H2 database while the self-backup contains data from an external database.
  - The NAKIVO Backup & Replication installation uses an external database while the self-backup contains data from the H2 database.
- It is not possible to edit the external database **Host**, **Port**, and **Database name** after a successful migration.
- If the internal database is used, the product checks the performance capability adequacy of this database to the current product workload:

- This check is performed every 10 days by default.
- If the total number of protected workloads for single tenant or per tenant for Multi-Tenant mode exceeds the limit of 100 VMs/instances/physical machines/oracle databases, the product displays the notification with recommendation to switch to the external database.

# Email Settings

On this page, you can configure your email settings. Do this by following the steps below:

- Log in to NAKIVO Backup & Replication.
- Click **Settings** in the left pane of the product.
- Go to the **General** tab.
- Click **Email Settings** to configure email settings on the page that opens.




**Important**

If you use an email with two-factor authentication, grant access permissions to NAKIVO Backup & Replication via your account security settings and generate a unique password. When configuring email setting of the product, enter this password in the **SMTP password** box.

As an example, use instructions for Google accounts provided in the" Using Email with Two-Factor Authentication" on the next page section below.

1. To set email settings, fill out the fields in the Email settings section:
    - **SMTP server**: The address of the server responsible for sending emails.
    - **SMTP username**: The username on the server (usually the same as the email username).
    - **SMTP password**: Usually the same as the password to your email.
    - **SMTP port**: Depends on encryption type.
    - **Encryption**: Select the type of encryption:
        - **None**: Always use a plaintext connection. Not recommended.
        - **TLS, if possible**: Start with plaintext, then use `STARTTLS` to switch to secure connection if supported by the server.

- **TLS, required**: Start with plaintext, then use `STARTTLS` to switch to secure connection; drop the connection if not supported by the server.
- **SSL, required**: Use the SSL-encrypted connection.

- **From**: Specify the sender email address
- **To**: Specify the receiver email address

Click **Send Test Email** to verify that the settings are correct.

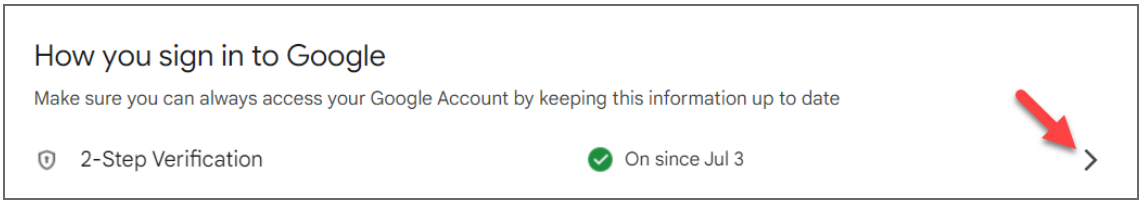
2. Click **Apply** to save the settings.
3. Alternatively, click **Discard Changes** to discard any changes you have made to the email configuration.
4. Optionally, click **Reset Settings** to return all the settings to their default values.

## Using Email with Two-Factor Authentication

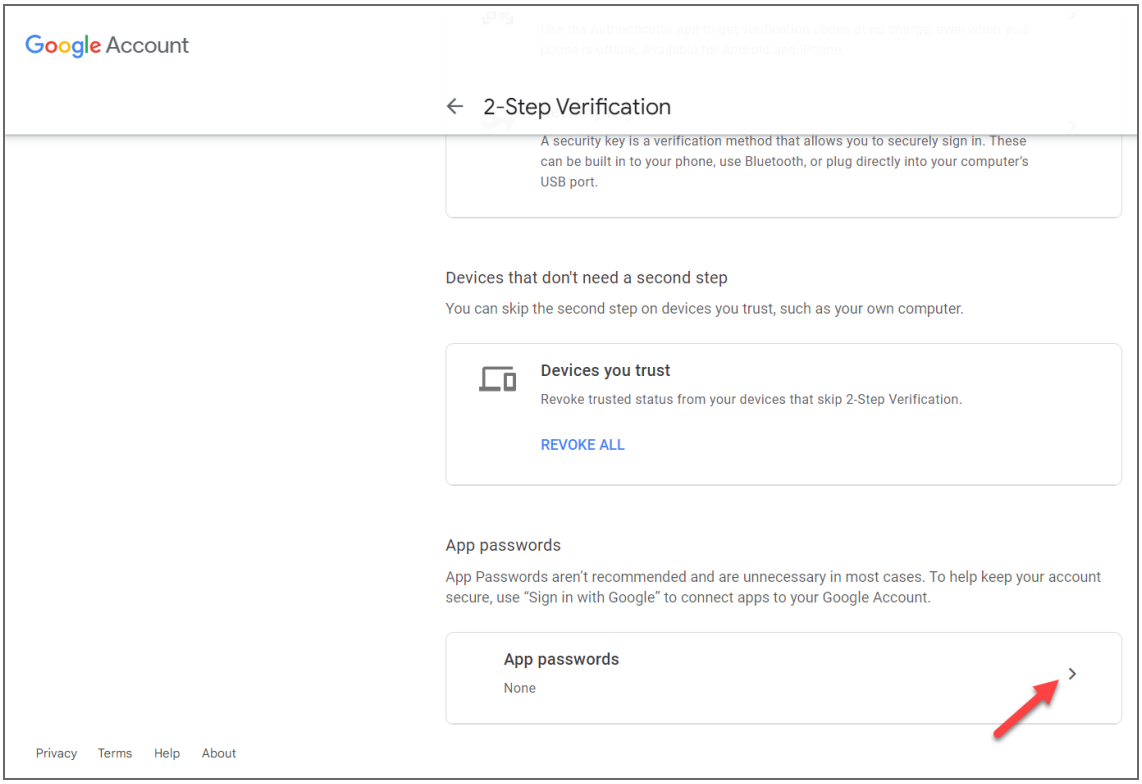
Many email providers require that the third-party apps or devices use two-factor authentication to sign into their accounts. Therefore, for correct email notification setup, you must grant access permissions to NAKIVO Backup & Replication via your account security settings and generate a unique **SMTP password**.

As an example, you can follow the instructions for setting up a Gmail account notifications below:

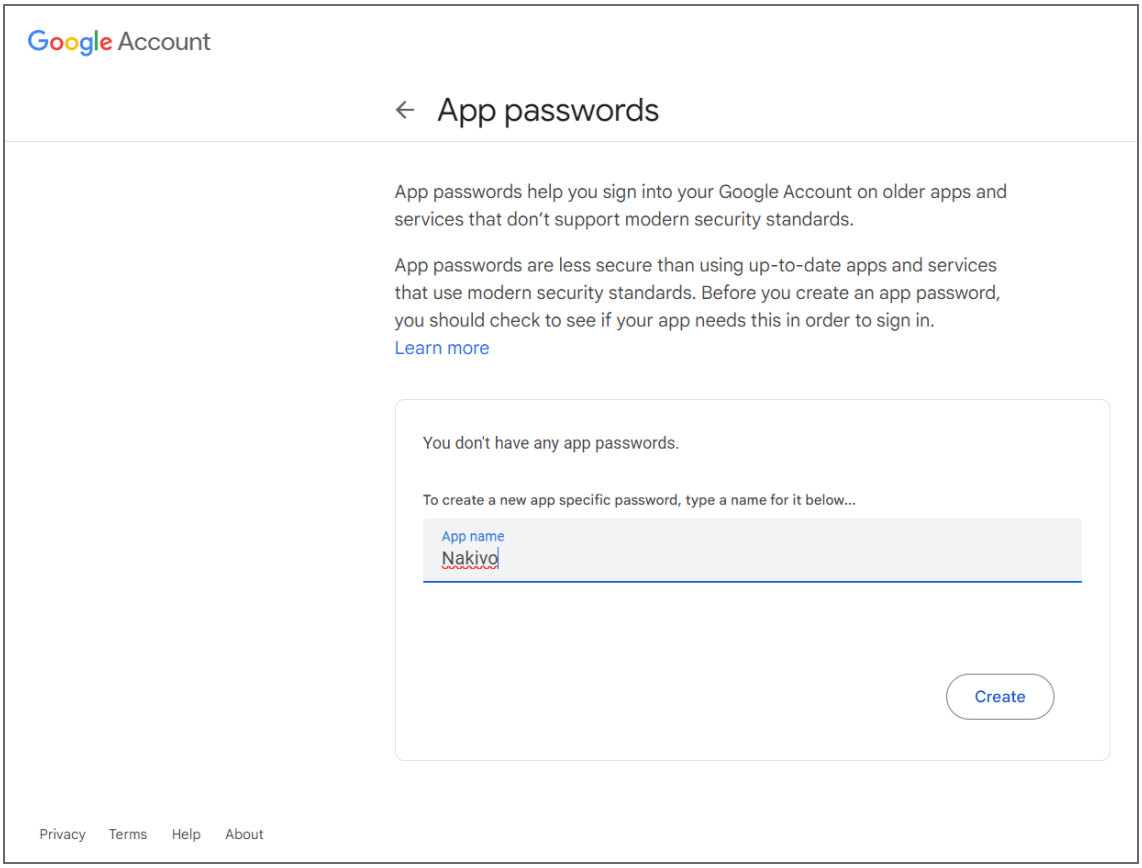
1. Log into your Google account and allow 2-step verification (you will be asked to complete a second step to verify it's you when you sign in):



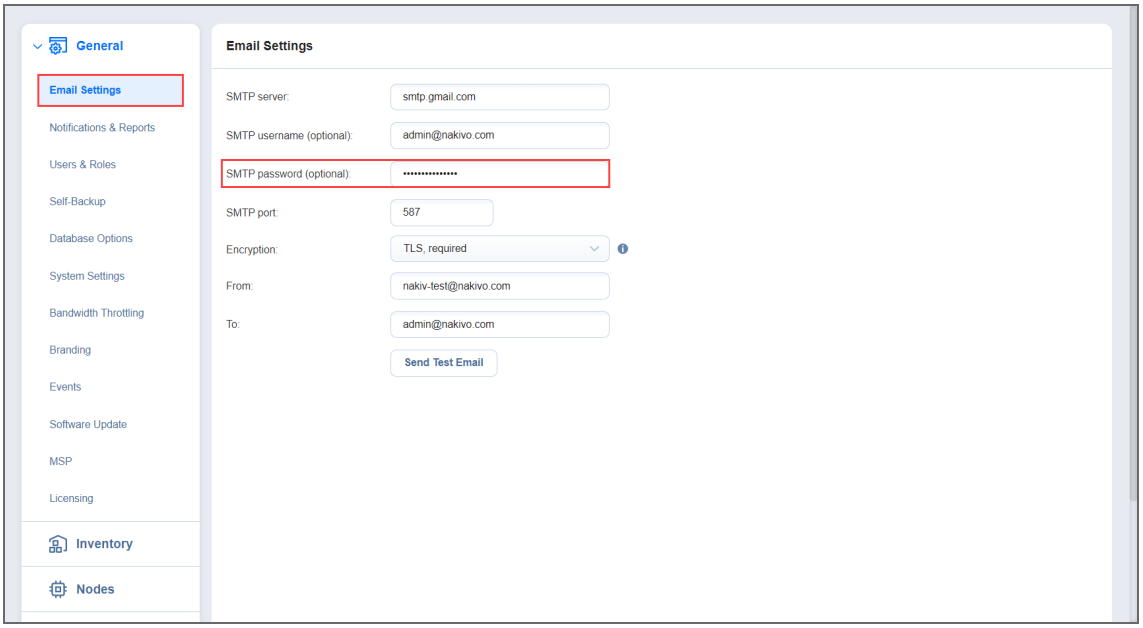

2. Scroll down to the *App passwords* section and click the arrow button to proceed with the password.




3. Type your app name and click **Create** to generate the password.



4. Now, copy the generated password and paste it into the **SMTP password** field in the **Settings > General > Email Settings** tab.

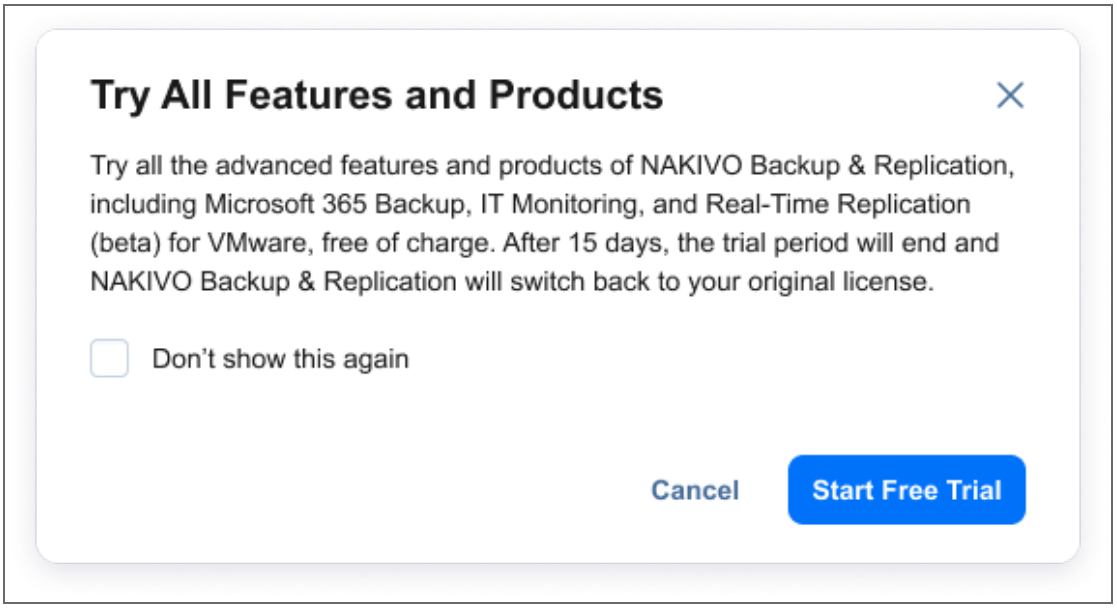


# Licensing

To check your license details, follow these steps:

1. Go to the main menu of NAKIVO Backup & Replication and click **Settings** > **General**.
2. Go to the **Licensing** tab to see license details.

Optionally, you can click the **Try All Features and Products** button to enable all Enterprise Plus license features for 15 days. After this time period ends, NAKIVO Backup & Replication automatically switches back to your original license.




**Notes**

The button is not displayed in the following cases:

- You are using NAKIVO Backup & Replication as a tenant in Multi-tenant mode.
- You are using one of the following license editions:
  - Free
  - Trial
  - Beta
  - Promo
  - Enterprise Plus
  - MSP Enterprise Plus

In the **License Information** section, you can find detailed license information, including:

- **Type**: Type of the license
- **Edition**: Edition of the license

- **Serial number**: Serial number of the license
- **License expiration date**: Date when the the license expires

In the **Perpetual licensing section**, you can see the following information:

- Number of licensed and used CPU sockets
- Number of licensed and used VMs
- Number of licensed and used physical servers
- Number of licensed and used physical workstations
- Number of licensed and used EC2 instances
- Number of licensed and used Oracle databases

In the **Per-workload subscription licensing** section, you can see the following information:

- Number of licensed and used workloads
- Subscription end date

In the **Microsoft 365 subscription licensing** section, you can see the following information:

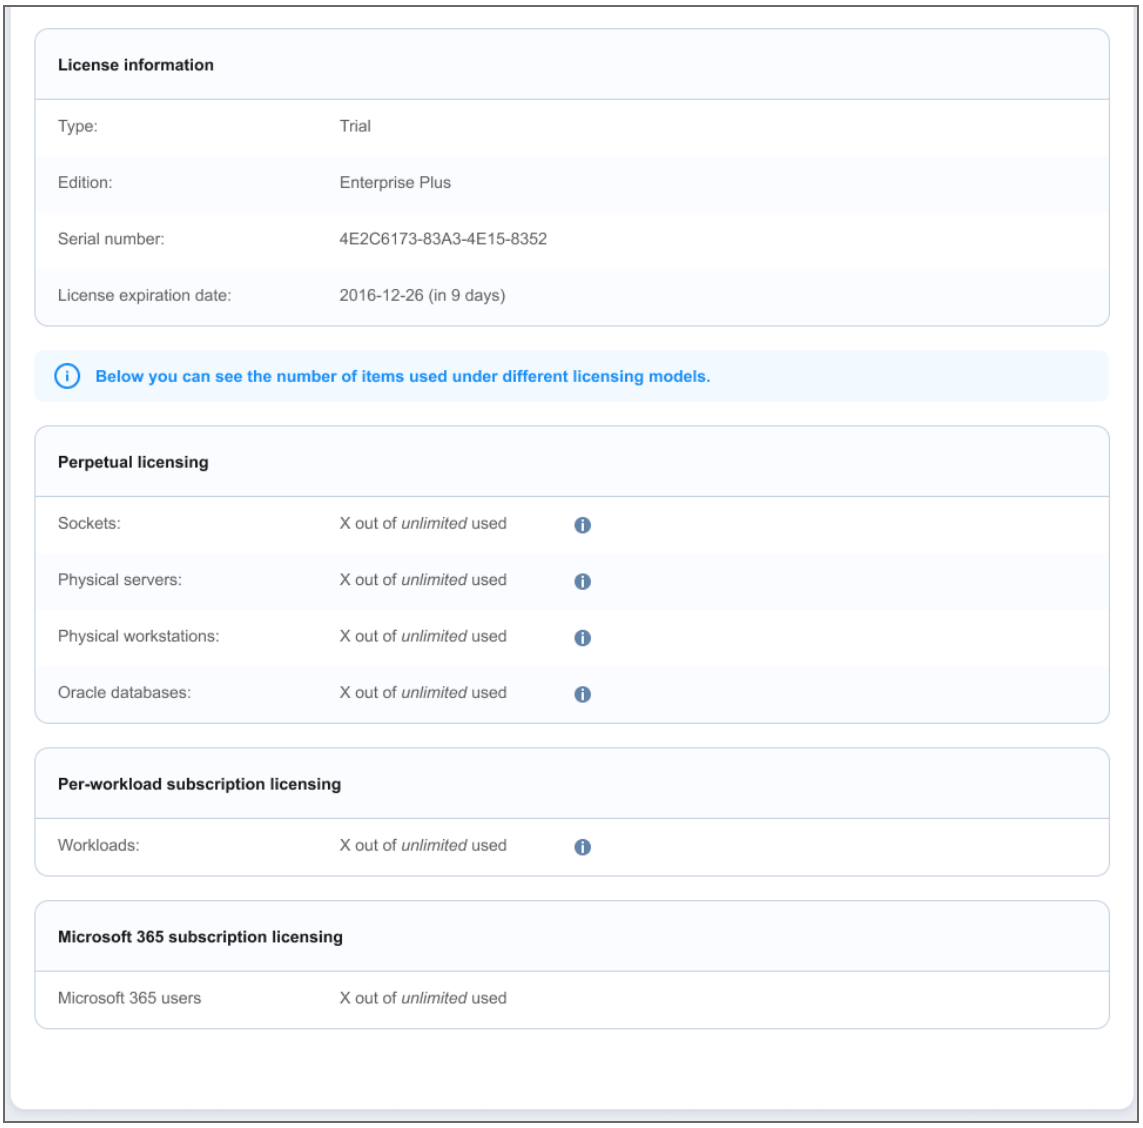
- Number of licensed and used Microsoft 365 users
- Subscription end date

If you are logged in as a tenant in multi-tenant mode, the following information is displayed in the **Obtain more licenses** section:

- **Email address** of the master tenant
- **Contact phone** of the master tenant
- **Company website** of the master tenant

To change your license, follow the steps below:

1. Go to the main menu of NAKIVO Backup & Replication and click **Settings**.
2. Go to the **Licensing** tab and click **Change License**.
3. Locate and open the license file in the window that appears.




# Upgrading from Free License

If your license type is **Free** and the **Trial** license has not yet been applied to you deployment of NAKIVO Backup & Replication, you can try the full functionality of the solution for 15 days. To do that:

1. Open the **Help Menu**.
2. Select the **Try full functionality** option. A new popup window appears.
3. Click **Start Free Trial.**

> **Note**
> Once the **Trial** license expires, the product automatically switches back to the **Free** license.

# Managing Backup Encryption

With NAKIVO Backup & Replication, you can personalize your protection settings through encryption of backup data, password management, and restoration from encrypted backups.

By using the **Backup Encryption** feature, you can encrypt the following:

1. Product workloads' backups
2. System configuration stored as a self-backup and used in the scope of system migration
3. File System Indexing (FSI) data stored on the Director

> **Note**
> Encryption does not alter the quality or effectiveness of the compression process, as compression reduces the data size before encryption is applied.

The product offers an efficient password management dialog box, so that you do not need to enter your password every time you restore encrypted data.

AWS KMS is highly recommended as an additional layer of protection and to restore the password hash in case you forget it.

Refer to the following sections to learn more about managing backup encryption:

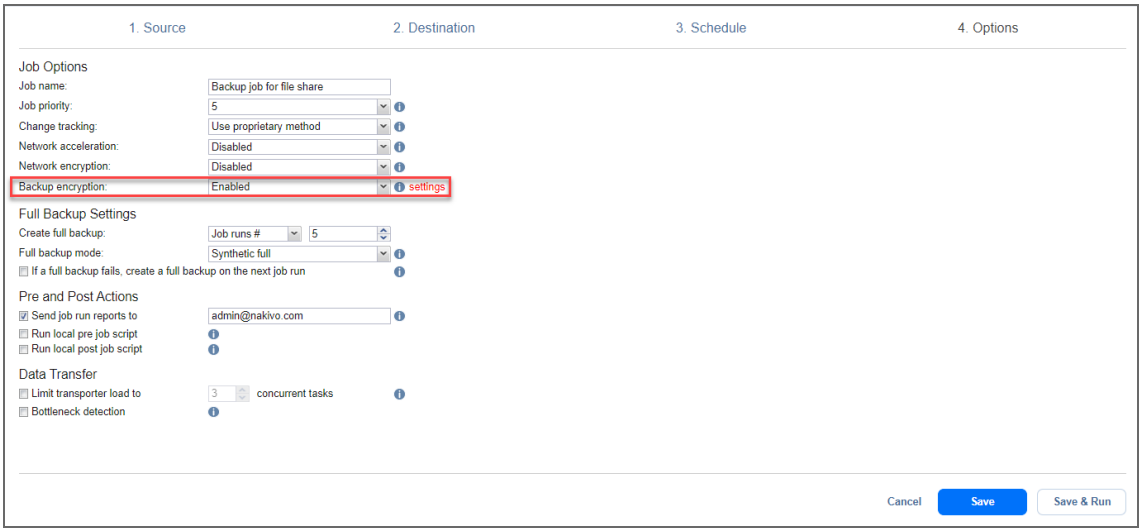
- "Enabling Backup Encryption" below
- "Managing Passwords" on page 432
- "Enabling KMS" on page 429

## Enabling Backup Encryption

To set up backup encryption, do the following:

1. In the **Options** step of the corresponding backup/backup copy job wizard, select **Enabled** from the **Backup encryption** drop-down list.
2. After the **Backup encryption** mode is enabled, the **settings** link becomes available.

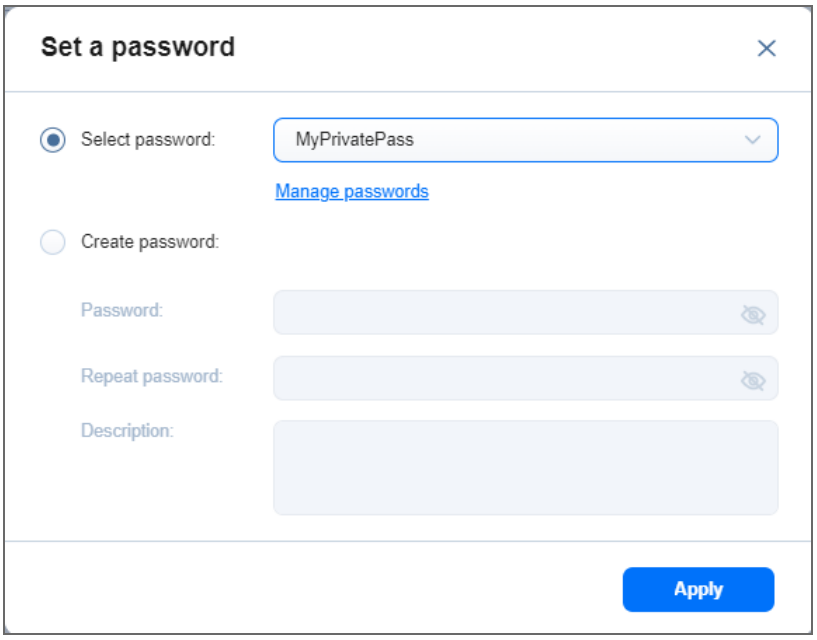
> **Notes**
> - If not configured, the settings link is highlighted in red.
> - For backup copy jobs, you can select from two options: **Enabled on source** and **Enabled on target** to enable **Backup encryption**. For more details, refer to Backup Encryption.



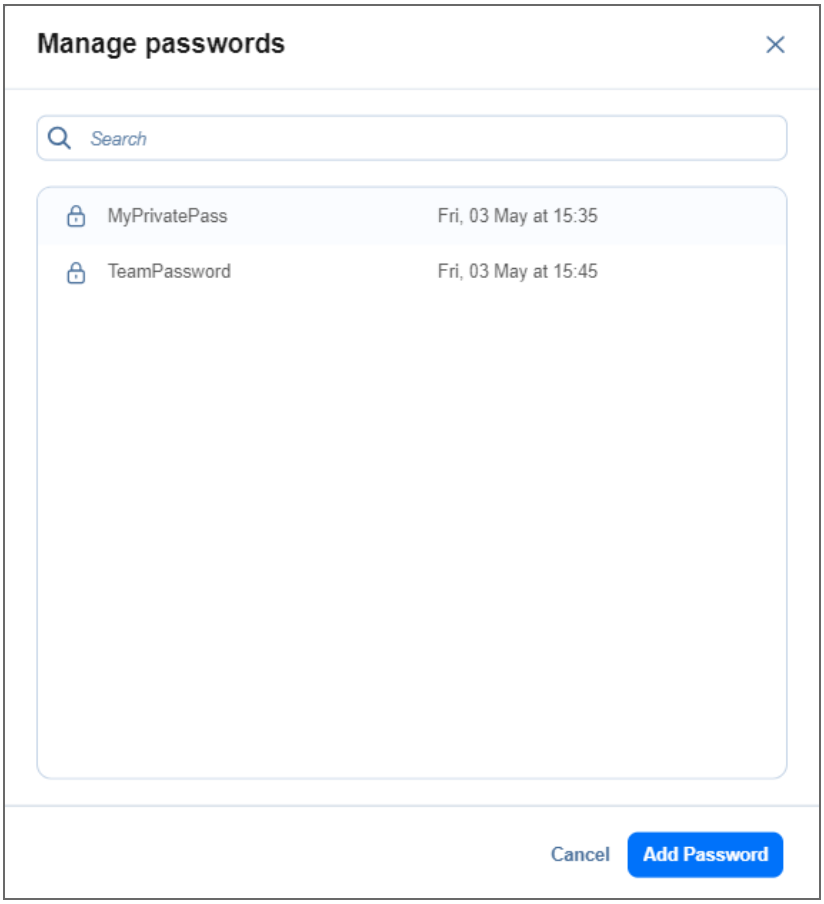

3. Click the **settings** link to open the **Set a Password** dialog box.

> **Notes**
>
> - If the **Key Management Service (KMS)** is not enabled, a warning message is displayed. To enable KMS, click the **Encryption Tab** link to go to the **Settings > General > System Settings > Encryption** tab. After enabling KMS, you can proceed with setting a password.
> - It's recommended that you enable the **(AWS) Key Management Service**. If **AWS** is enabled, all backup encryption passwords encrypted with the Key Management Service cryptographic key are available for recovery in case of product re-installation. For more information, refer to "Enabling KMS" on page 429.

4. In the dialog box that appears, select the needed password or create a new one. Refer to "Setting Password" on page 432 for more details.




5. Optionally, you can click the **Manage passwords** link to manage the existing or add a new password to the list of passwords.




6. Click **Apply** to proceed.

   The product automatically generates the password hash based on the user password

The cryptographic salt used for hash creation is saved in the recovery point metadata.

The password hash is used to generate a single-use encryption key to encrypt the backup and FSI data (if the **FSI** option is enabled in the job).

- If **KMS** is enabled, the password hash is encrypted with the AWS KMS key and saved in the recovery point metadata.
- If **KMS** is not enabled, a dialog box opens, warning you that if you lose the password, it will be impossible to decrypt your data, and this data will be lost forever.

7. Click the **Proceed** button to go to the next step of the wizard.
   - To enable KMS, click the **Encryption Tab** link to go to the **Settings > General > System Settings > Encryption** tab.
   - Click the **Cancel** or **X** button to close the dialog box without applying any changes.

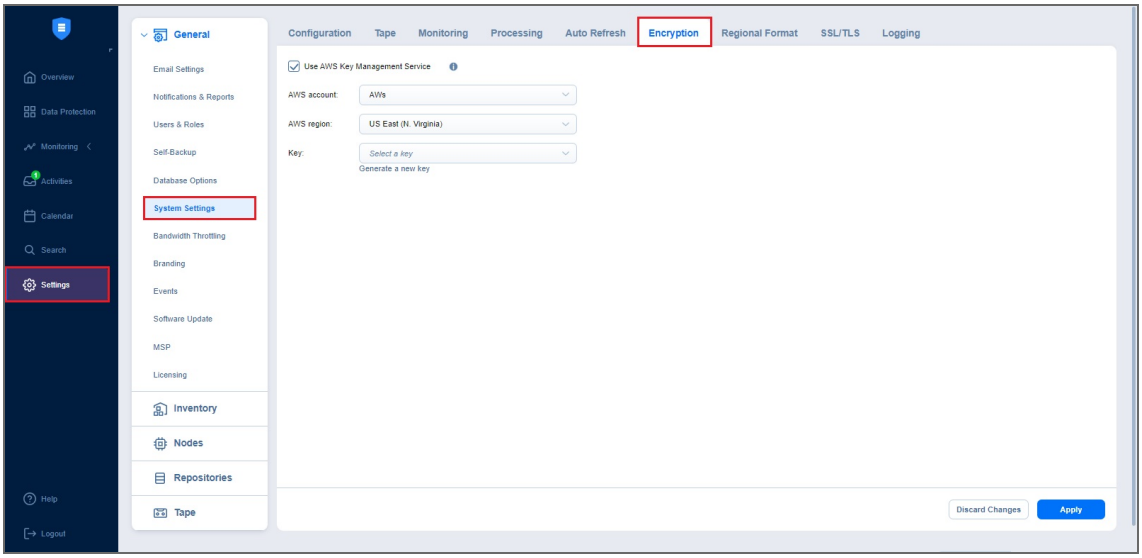
> **Notes**
>
> - Single chain of incremental recovery points has to have consistent encryption settings including encryption password.
> - Changing encryption settings in a job (including changing encryption password) results in creating an active full backup and starting a new chain or recovery points.

## Enabling KMS

For encrypting the password hash, the AWS Key Management Service is used.

To enable the **(AWS) Key Management Service** proceed as follows:

1. Open the **Settings > General > System Settings > Encryption** tab.
2. Check the **Use AWS Key Management Service** checkbox (disabled by default).
3. Specify the AWS account (the AWS account should be discovered first) by selecting from the **AWS account** dropdown (the option is disabled if no AWS accounts have been discovered or if the discovered AWS account has insufficient permissions to use KMS ). For the list of required AWS permissions for creating the (AWS) KMS Keys refer to the Feature Requirements section.
4. Specify the AWS region by selecting from the **AWS region** dropdown (the option is disabled if no AWS account has been discovered).



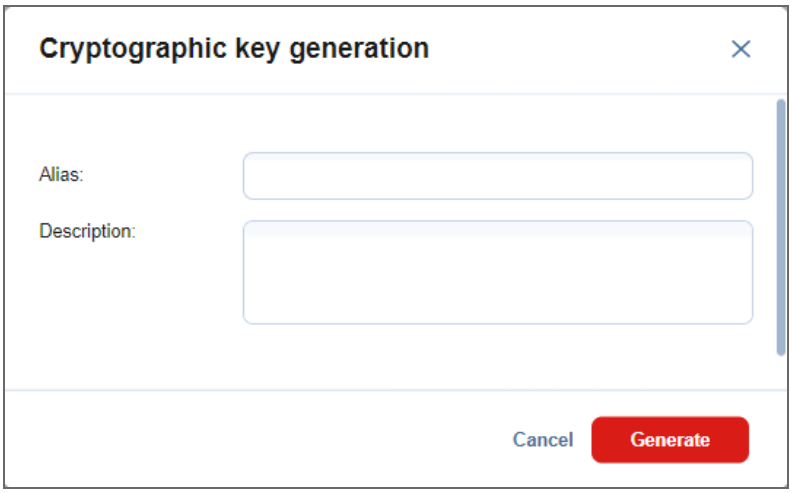

5. From the **Key** dropdown, select a key from the list of existing symmetric cryptographic keys available to the specified AWS account (the option is disabled if no keys are available).
6. Click **Apply** to apply the changes.
7. If encryption is set for a job, the password hash is generated based on the provided password. The password hash gets encrypted with the KMS cryptographic key with *base64* and is saved in the recovery point metadata.

**Notes**

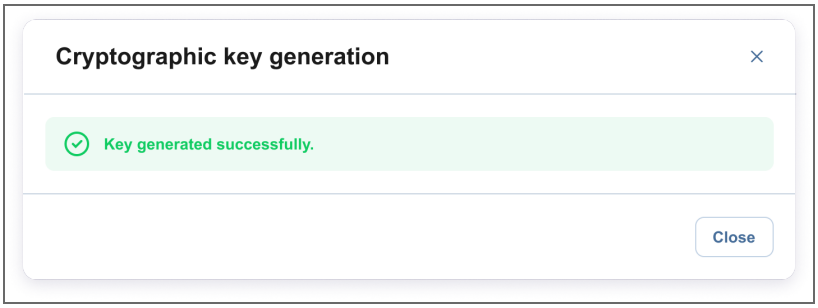
- If the AWS Key Management Service is enabled, the password hashes are automatically be restored in case of product reinstallation. Otherwise, passwords need to be provided manually in case the password hashes are not available.
- The AWS Key Management Service is not applied to self-backup and system configuration encryption.

## Generating New KMS Cryptographic Key

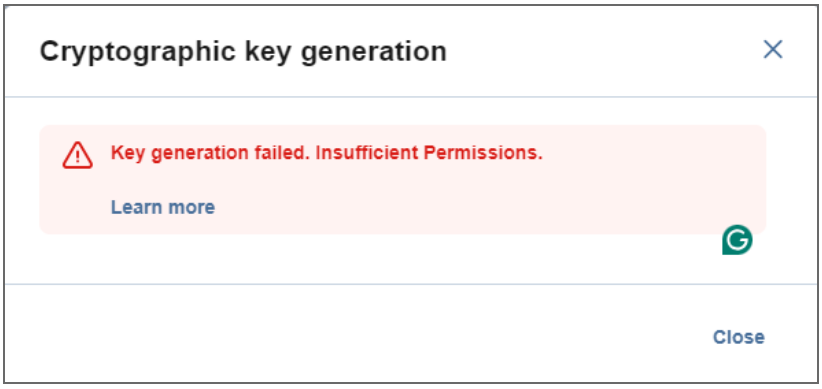
1. Optionally, click **Generate a new key** to open the **Cryptographic key generation** dialog box.




2. In the dialog box that appears, enter the alias and its description (optional) and click **Generate** to initiate generation of a new symmetric cryptographic key in the specified account. For more details, refer to the AWS::KMS::Alias page. Click the **Cancel** or **X** button to close the dialog box without applying any changes.
3. If the key was generated successfully, the following dialog appears:




4. If the key generation fails, a dialog box with the reason for the failure appears:




5. Click the **Close** or **X** button to proceed with key generation.

# Managing Passwords

With NAKIVO Backup & Replication, you can create and then manage your passwords for encrypting/decrypting backups and FSI data stored on the Director or exporting/importing the system configuration stored as self-backups.

## Setting Password

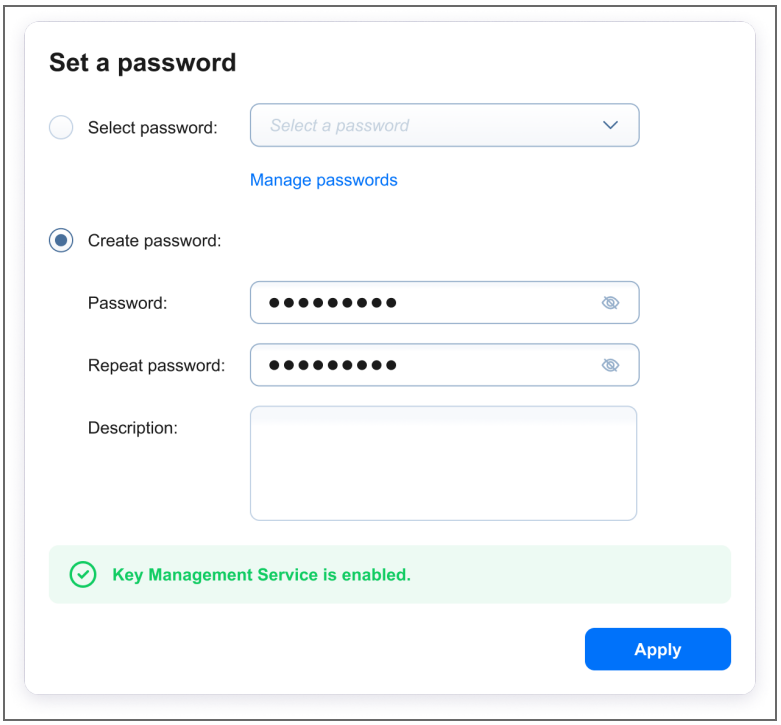
To set a password, proceed as follows:

1. In the **Set a Password** dialog box, choose the **Select password** option and select the needed password from the list of passwords to apply it to the job. Close the dialog box.

   **Note**

   The **Select Password** option is disabled if no passwords are available.

2. Optionally, you can click the **Manage passwords** link to manage the existing or add a new password to the list of passwords.
3. To create a new password, choose the **Create password** option. In the fields below, enter the password, repeat it, and add a description (required).



> **Note**
>
> By default, for the source-side backup encryption, the encryption key is applied to data in the trusted zone from the target side. If the Allow sending encryption key to source side. Not Recommended checkbox is selected, the encryption key will be sent to the non-trusted zone to arrange communication from the trusted to a non-trusted zone.

4. Click **Apply** to close the dialog box and apply the set password to the job.
   Click the **Cancel** or **X** button to close the dialog box without applying any changes.
   - If the **Key Management Service (KMS)** is enabled, the password hash is encrypted with the AWS KMS key and saved in the recovery point metadata.
   - If the **KMS** is not enabled, a warning message is displayed.
   - Click the **Proceed** button to dismiss the warning and apply the set password to the job and close the dialog box.
   - To enable KMS, click the **Encryption Tab** link to go to the **Settings > General > System Settings > Encryption** tab.

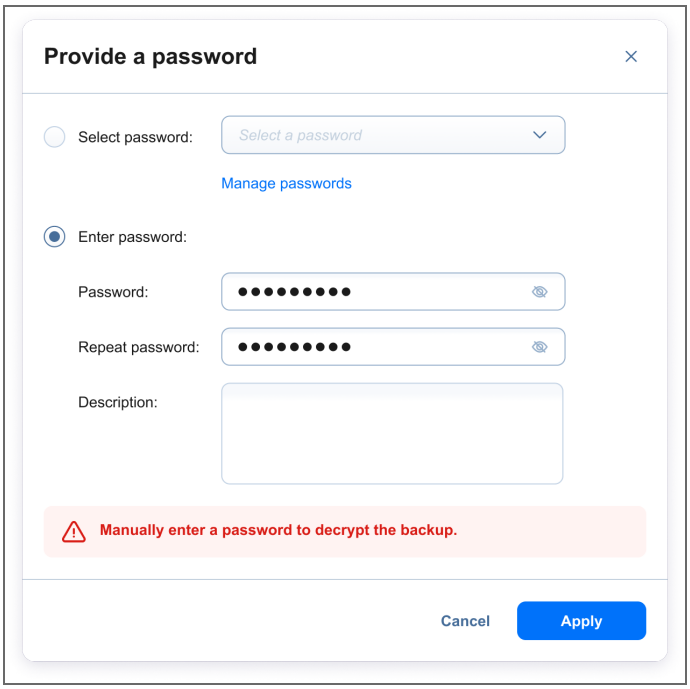
     > **Note**
     >
     > It's recommended that you enable the **(AWS) Key Management Service**. If **KMS** is enabled, all backup encryption passwords encrypted with the Key Management Service cryptographic key are available for recovery in case of product re-installation. For more information, refer to "Enabling KMS" on page 429.

## Providing Password

To provide the password, do the following:

1. In the **Provide a Password** dialog box, select the needed password from the list of passwords or enter it in the provided fields.
2. Optionally, you can click the **Manage passwords** link to manage the existing or add a new password to the list of passwords.
3. To enter a password, choose the **Enter password** option. In the fields below, enter the password, repeat it, and add a description (required).
4. Click **Apply** to proceed.

Click the **Cancel** or **X** button to close the dialog box without applying any changes.



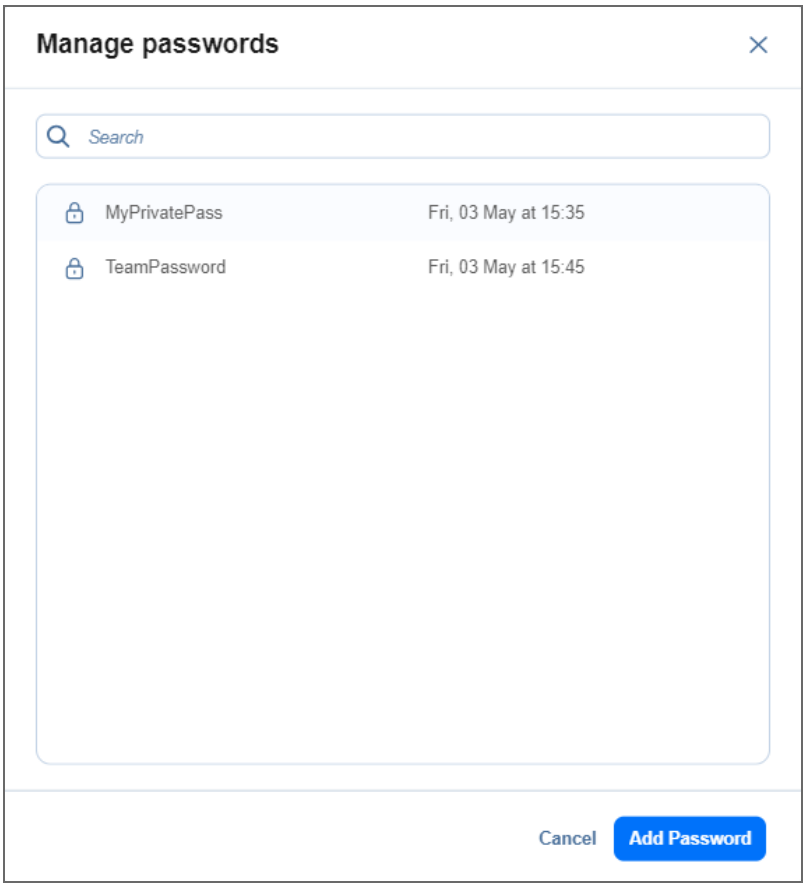

## Managing Passwords

In the **Manage passwords** dialog box, you can do the following:

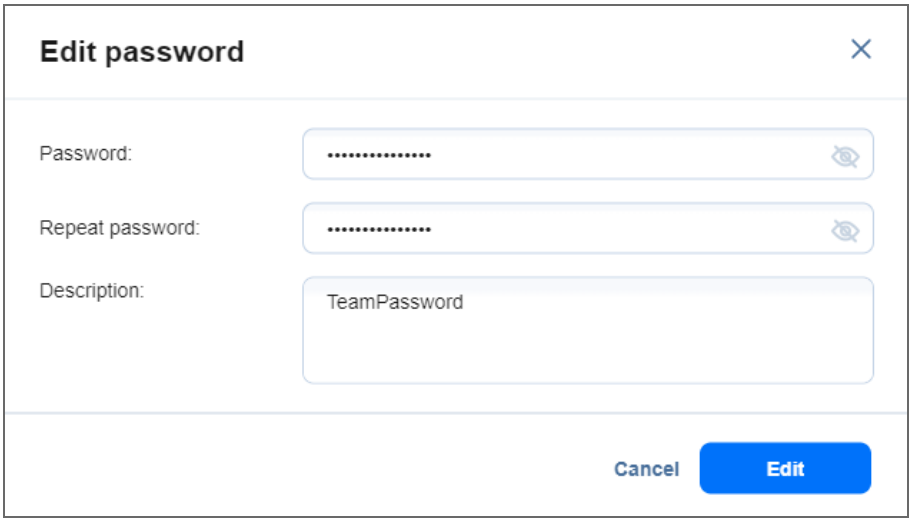
- View the list of the saved passwords in the current tenant with their descriptions and the timestamps of their last updates.

> **Note**
> If the description or timestamp does not fit it is truncated. To display full password information, hover over it.

- Search for passwords by entering the description (fully or partially) into the **Search** bar.
- Manage the existing passwords by hovering over the needed item and selecting from the list of typical actions applicable to it:




- Select **Edit** to open the **Edit password** dialog box, where you can edit the existing password.
  - Click the **Edit** button to confirm changes or the **Cancel** or **X** button to close the dialog box without applying any changes.

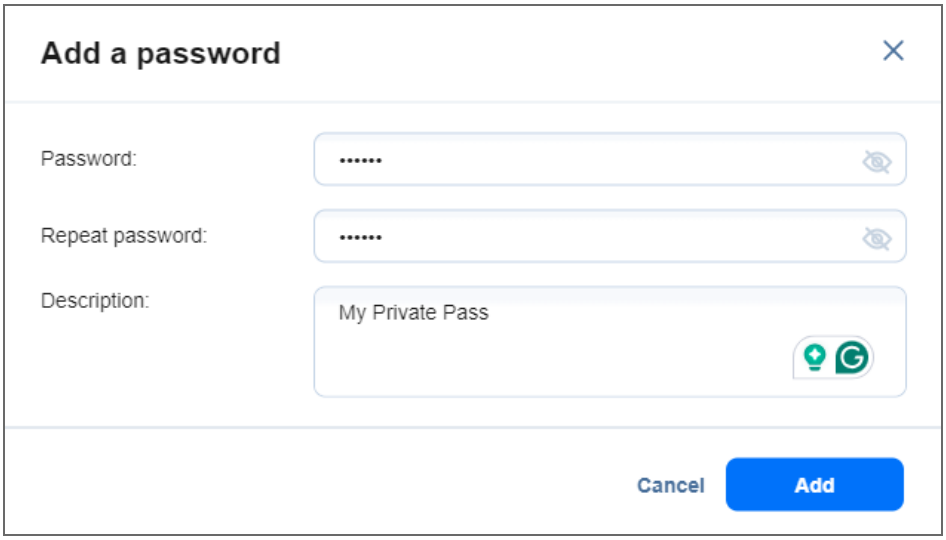



- Select **Delete** to display a confirmation dialog box, where you can confirm deleting the existing password. Click **Cancel** to dismiss the dialog box without applying any changes.

**Note**

The confirmation dialog box is displayed if the rule is not used by any jobs.

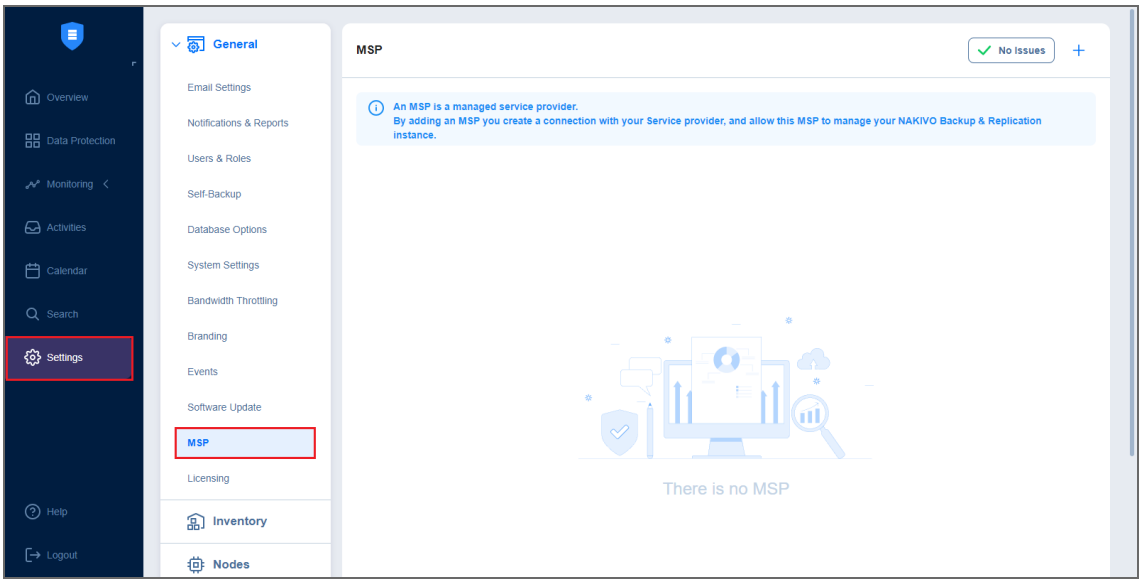
- Add a new password by clicking the **Add Password** button. In the dialog box that appears, enter the password, repeat it, and provide its description (optional), then click **Add**.




- Click the **Cancel** or **X** button to close the dialog box without applying any changes.

# MSP Tab in Single-Tenant Mode

To allow an MSP to manage your standalone instance of NAKIVO Backup & Replication, you need to link your instance to the MSP's instance of the solution. To do so, you should first add and configure the MSP details on the **MSP** page in **Settings**. These details can be provided by the MSP after Remote Tenant Configuration. Establishing a connection to an MSP as a remote tenant allows the MSP to monitor and manage your instance of NAKIVO Backup & Replication.



As a remote tenant, you can manage an established connection to an MSP on the **MSP** page.

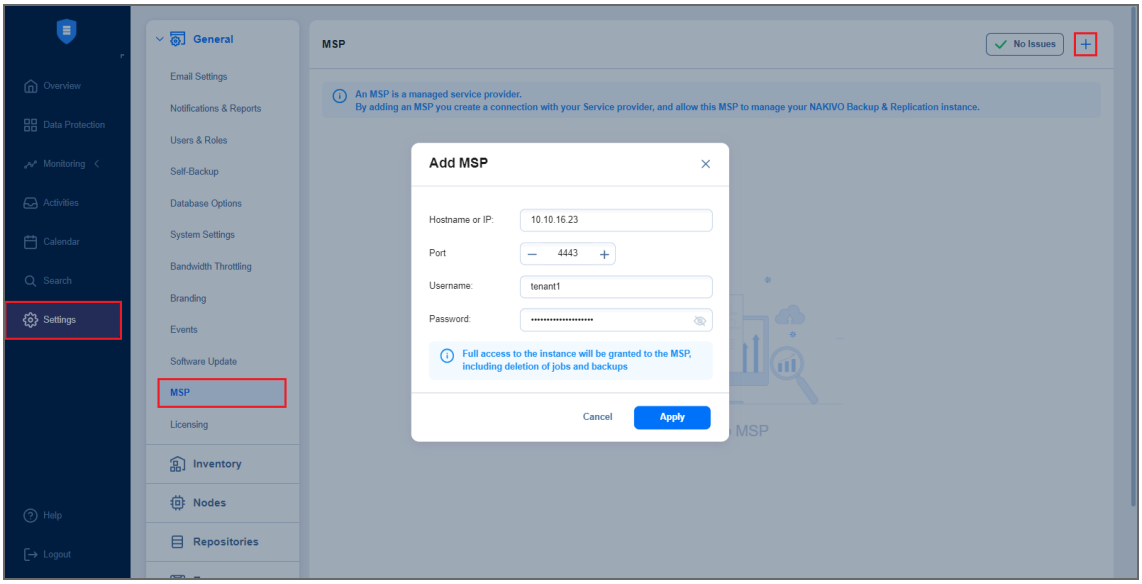
See the topics below for more information:

- “MSP Architecture” on page 439
- “Adding an MSP” below
- “Managing an MSP Connection” on the next page

## Adding an MSP

To add an MSP to which you would like to link your NAKIVO Backup & Replication instance, you need the Master tenant’s hostname or IP address, port number, and the remote tenant credentials they have generated for you. Once you obtain the above information, follow the steps below:

1. Go to **Settings** and click the **+** button in the top right corner.
2. Fill in the **Hostname or IP**, **Port**, **Username**, and **Password** fields based on the information provided by your Master tenant (MSP).

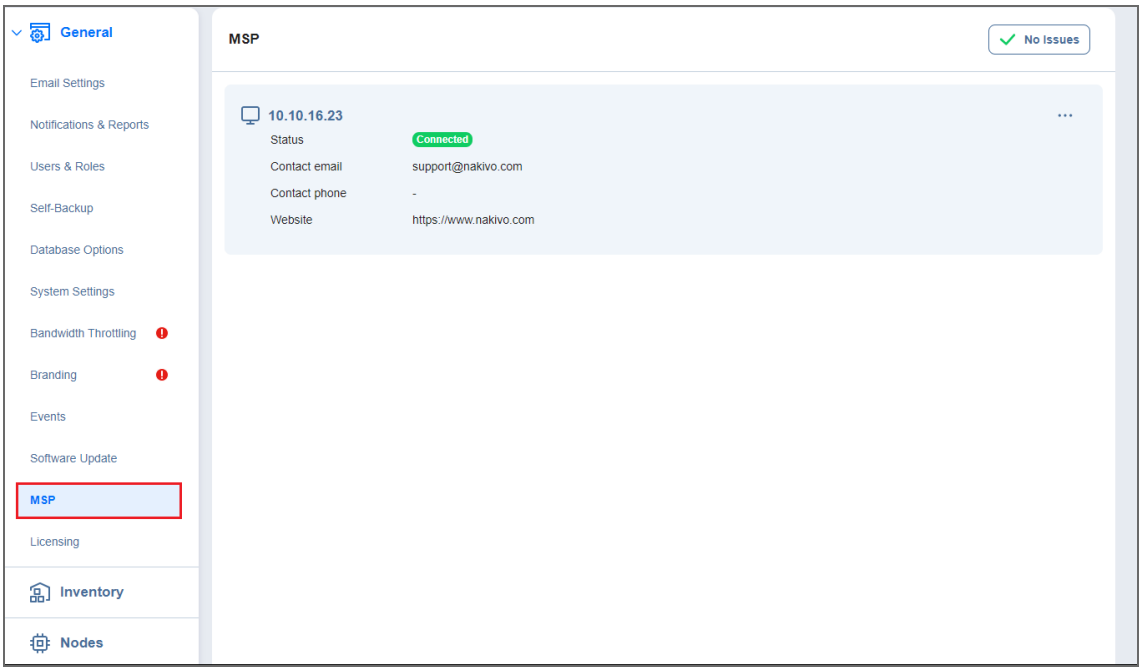



3. Click **Apply**. The following screen will display the MSP’s certificate details.

> **Notes**
>
> - Your version of NAKIVO Backup & Replication must be the same as the MSP’s version. Otherwise, you cannot connect to MSP’s instance of NAKIVO Backup & Replication.
> - The MSP uses a separate listening port for communication with a remote tenant's instance (port *6702* is used by default). If the MSP changes the listening port used, the connection may be interrupted. For more information on required TCP ports, see the **MSP Console** section in “Feature Requirements” on page 169.

4. Read through the MSP’s certificate details and click **Accept**.

5. The added MSP should now appear in your **MSP** menu.




Once the MSP is added, you have successfully established a connection with the Master tenant as a remote tenant. This allows the Master tenant to access your instance of NAKIVO Backup & Replication as long as the connection is active.

# Managing an MSP Connection

Several options are available for managing an established connection to an MSP. See the sections below for more information on managing yyour connection as a Remote tenant.

- Viewing MSP Details
- Disconnecting/Reconnecting an MSP
- Deleting an MSP

## Viewing MSP Details

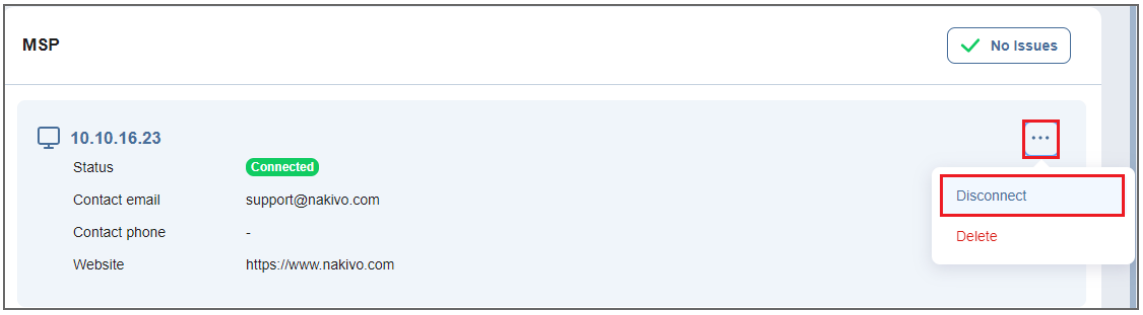
The MSP block in the **MSP** menu contains the following information:

- **Hostname or IP**: The hostname or IP address of the Master tenant.
- **Connection status**: Current status of your connection to the MSP.
  - **Connected**: Your NAKIVO Backup & Replication is connected to the MSP
  - **Disconnected**: Your NAKIVO Backup & Replication is disconnected from the MSP
  - **Connecting**: Your NAKIVO Backup & Replication is actively trying to establish a connection to the MSP
- **Contact information**: The Master tenant's email address, phone number, and website, updated automatically.




## Disconnecting/Reconnecting an MSP

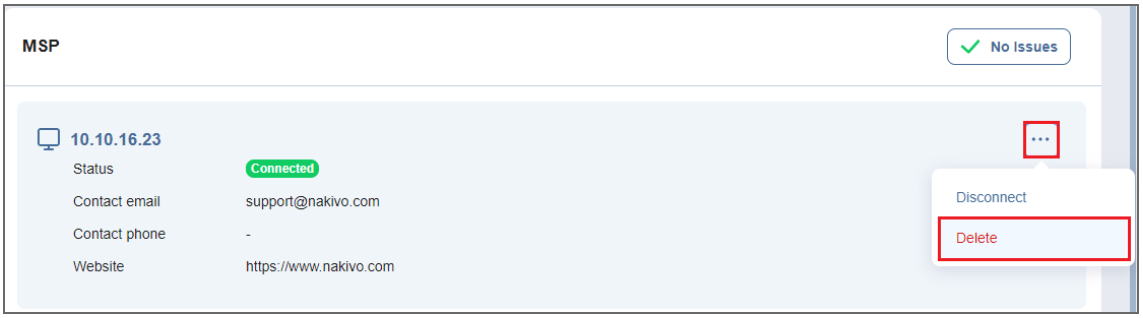
To disconnect your instance of NAKIVO Backup & Replication from an MSP, click the ellipsis **Manage** button in the top right corner of the MSP block. In the popup, click **Disconnect**. This will suspend the connection to the MSP until resumed.




To reconnect to the MSP, simply click the ellipsis **Manage** button in the top right corner of the MSP block and click **Connect** in the popup. You will not be asked to provide the same details you did when first connecting to the MSP unless the MSP has changed their certificate or your remote tenant credentials.

## Deleting an MSP

To delete an MSP connection, click the ellipsis **Manage** button in the top right corner of the MSP block. In the popup, click **Delete** and confirm the action. This action will erase all tenant data from the MSP's side and vice versa.




# MSP Architecture

The MSP infrastructure in NAKIVO Backup & Replication consists of the following components:

- **MSP Console**: The centralized control hub that enables service providers to efficiently manage the entire cloud backup and replication processes. It is available to users with an *MSP licens*e, *Beta instance*, *Promo license*, or *Trial license*.
- **MSP Dashboard**: The Master Tenant Dashboard that is available for a multi-tenancy product with a non-MSP license.
- **Director**: This core solution component orchestrates and coordinates the other components of the NAKIVO Backup & Replication infrastructure for MSPs, ensuring a cohesive and streamlined operation.
- **Transporter**: Transporters are responsible for efficiently moving data between different components and optimizing the backup and replication processes.
- **Backup Repository**: These are storage locations within the NAKIVO Backup & Replication cloud infrastructure. Backup repositories are designed to store the backups of local tenant workloads. As primary and secondary backup storage, backup repositories can be used to implement the recommended 3-2-1 backup strategy.
- **Master Site and Remote Site(s**): These terms encompass the main NAKIVO Backup & Replication infrastructure components (that is, **Director**, **Transporter(s)**, and **Backup Repositories**) in an MSP setting. The Master site (on the MSP side) serves as the central point for managing multiple Remote sites (client sites).
- **Remote Tenant**: This is a tenant with a standalone NAKIVO Backup & Replication instance installed in single-tenant mode in their infrastructure and added to the multi-tenant installation of the MSP for monitoring purposes by the Master tenant (the MSP).
- **Local Tenant**: This is a tenant created and managed by the Master tenant (MSP) to provide backup and disaster recovery as a service to clients. To establish a connection with client remote resources, the MSP should enable Direct Connect.

**Notes**

- Only users with an *MSP license*, *Beta instance*, *Promo license*, or *Trial license* can additionally access the MSP Console and the Licensing > Tenants tab. This allows them to efficiently oversee all independent instances of NAKIVO Backup & Replication associated with a managed service provider (MSP) as well as local tenants from a unified interface, eliminating the need to navigate through individual tenants. Users with multi-tenant instances of NAKIVO Backup & Replication without these license types only have access to the MSP Dashboard.
- Multiple remote tenants, each tied to a separate client environment, can connect to a single instance of the MSP Director.
- Only local tenants consume MSP workloads. Since remote tenants are licensed independently, they have their own workloads.
- The Master tenant (MSP) operates each remote tenant independently from other tenants, managing and facilitating a specific tenant's backup, replication, and recovery operations.
- The **Direct Connect** feature allows the NAKIVO Backup & Replication instance installed at the MSP site to access the resources at the remote site of a client via a single port connection (the port should be exposed on the client's infrastructure) without the need to establish a VPN connection.

The MSP architecture in NAKIVO Backup & Replication empowers the MSP Director to oversee and synchronize with each tenant (client). It ensures the establishment of secure and isolated communication channels, facilitating both data exchange and administration within the distinct data protection environment of each tenant.

To establish these isolated communication channels, the MSP initiates the creation of a remote tenant account for standalone users. After a remote tenant at the MSP Director is created, the client with a standalone instance of NAKIVO Backup & Replication can use the provided credentials to connect to the MSP's NAKIVO Backup & Replication instance. Refer to "Creating a Remote Tenant" on page 1261 for more details.

Once the setup is done on both sides, the MSP can manage the Remote tenant like a regular user instance on the remote tenant side. This way, the MSP (Master tenant) manages the remote tenants (client sites) with a single tenant from the NAKIVO Backup & Replication instance installed on another Director. For more information, refer to "Managing an MSP Connection" on page 438.

A managed service provider (MSP) can also use the MSP Console to create and manage local tenants. To do this, the MSP can enable Direct Connect to establish a connection with client remote resources.

## Running an MSP Connection

The general workflow for running a connection between the remote tenant instance and the MSP's instance of NAKIVO Backup & Replication is as follows:

- At startup, each remote tenant connects to the MSP Director to establish and maintain a single TCP connection (using port *4443*) for various types of traffic, including:
  - MSP Dashboard traffic (between MSP and remote tenant): the established connection on port *6702 TCP* for bidirectional communication between the MSP Dashboard and the storage transporter on the client side.
  - Remote tenant traffic: All remote tenant-related activities, including drill-through into data, remote tenant actions, and other communication between the MSP and the storage transporter.
- In the course of initial connection establishment, the MSP informs the remote tenant about the listening port used.
- MSP opens the port in the firewall.
- Each remote tenant then maintains a dedicated and separate TCP connection with the MSP Director (using port *6702 TCP*) for bidirectional communication.
- The MSP Director, in turn, uses a separate listening port (*6702 TC*P by default) for communication with a remote tenant's instance (while *4443 TCP* is used for initial MSP - remote tenant connection setup/checks only).
- These connections serve as the conduits for all data transfers, administrative commands, monitoring, and management actions related to the specific remote tenant.
- Both ports, *6702 TCP* and *4443 TCP*, use TLS for securing the communication channels between the MSP Director and each remote tenant. TLS ensures encryption, authentication, and data integrity for the traffic exchanged between these components, maintaining data security for each tenant's connection.
- The ports used by the MSP Director (*6702 TCP*) and the initial connection setup port (*4443 TCP*) can be configured or customized by the MSP using Expert Settings to suit specific deployment requirements or security protocols.
- If the MSP changes the listening port used, the connection may be interrupted. For more information on the required TCP ports, see the **MSP Console** section in Feature Requirements.
- During a product update (on the Master tenant or Remote tenant side) the connection is suspended. After the update is completed, the remote tenant continuously tries to reestablish the connection with the Master tenant. In turn, the Master tenant side starts listening to the connection attempts from the Remote tenant side until the connection is reestablished.

- In case of a network connectivity issue, the connection between the MSP and the remote tenant can be lost. Once the issue is resolved, the Remote tenant side initiates infinite attempts to reestablish the connection with the Master tenant side. In turn, the Master tenant side starts listening to the connection attempts from the remote tenant side until the connection is reestablished. Alternatively, the connection can be reestablished manually.

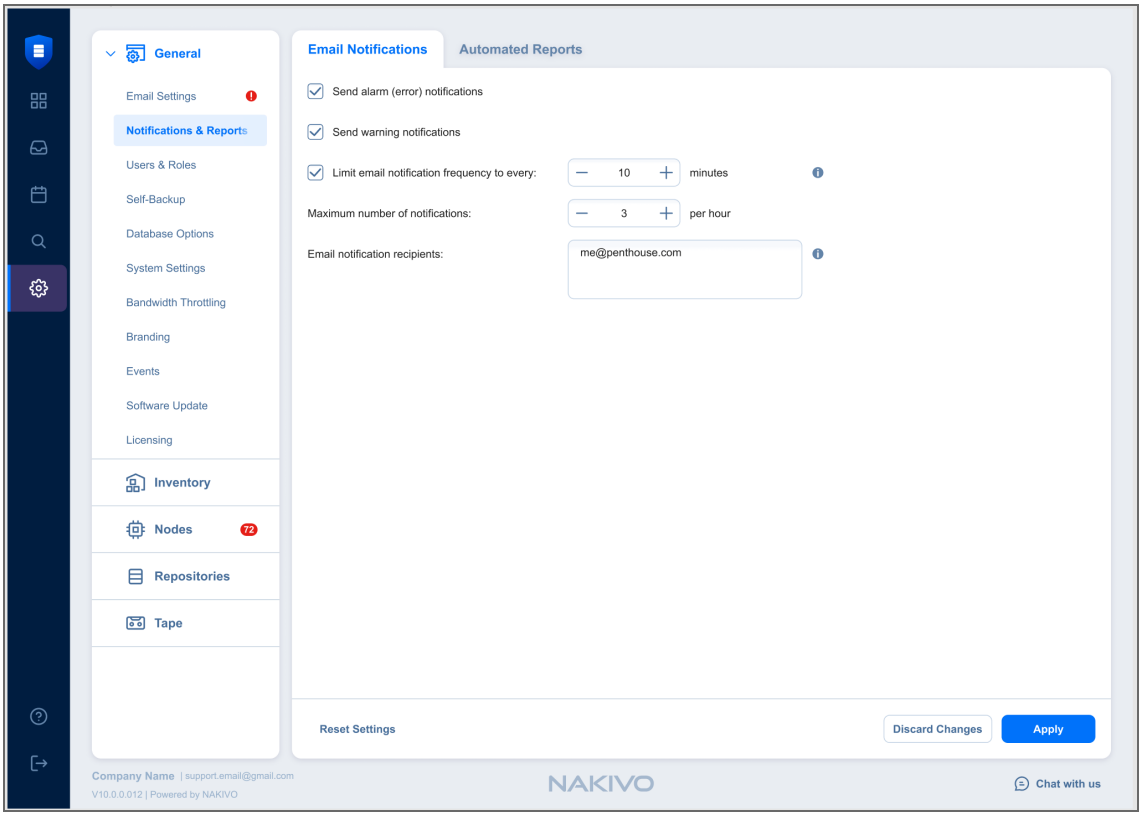
# Notifications & Reports

NAKIVO Backup & Replication can send notifications and reports over email.

- Email Notifications
- Automatic Reports

To receive automatic notifications, configure email settings by following the steps below:

1. Log in to NAKIVO Backup & Replication.
2. Click **Settings** in the left pane of the product.
3. Go to the **General** tab.
4. Click **Notifications & Reports** to configure notifications and automatic reports section on the page that opens.
5. Click **Apply** to save the settings after you're done.
6. Alternatively, click **Discard Changes** to discard any changes you have made to the email configuration.
7. Optionally, click **Reset Settings** to return all the settings to their default values.




**Note**

To configure email notifications and automatic reports, you must first configure email settings.

## Email Notifications

To set Email notifications, fill out the fields in the *Email notifications* section:

- **Send alarm (error) notifications**: If this option is selected, NAKIVO Backup & Replication will send email notifications to the specified recipients in case an error (for example, a job failure) occurs in the product. For users in Multi-Tenant Mode, these notifications also identify the relevant tenant and the instance where the error occurred.
- **Send warning notifications**: If this option is selected, NAKIVO Backup & Replication will send email notifications to the specified recipients in case the product generates a warning message (for example, lost connection to a host or **Backup Repository**). For users in Multi-Tenant Mode, these notifications also identify the relevant tenant and the instance that generated the warning.
- **Limit email notification frequency**: This option allows you to set up an email notification frequency in minutes. If deselected, notification emails will be sent every 5 minutes with no hourly limit.
- **Maximum number of notifications**: Use this option to change the limit of email notifications receivable per hour. If this limit is reached, any additional notifications will be delivered the following hour.
- **Email notification recipients**: Specify the recipients who will be receiving alarm and warning notifications (if enabled).

## Automatic Reports

To set automatic reports, fill out the fields in the Automatic Reports section:

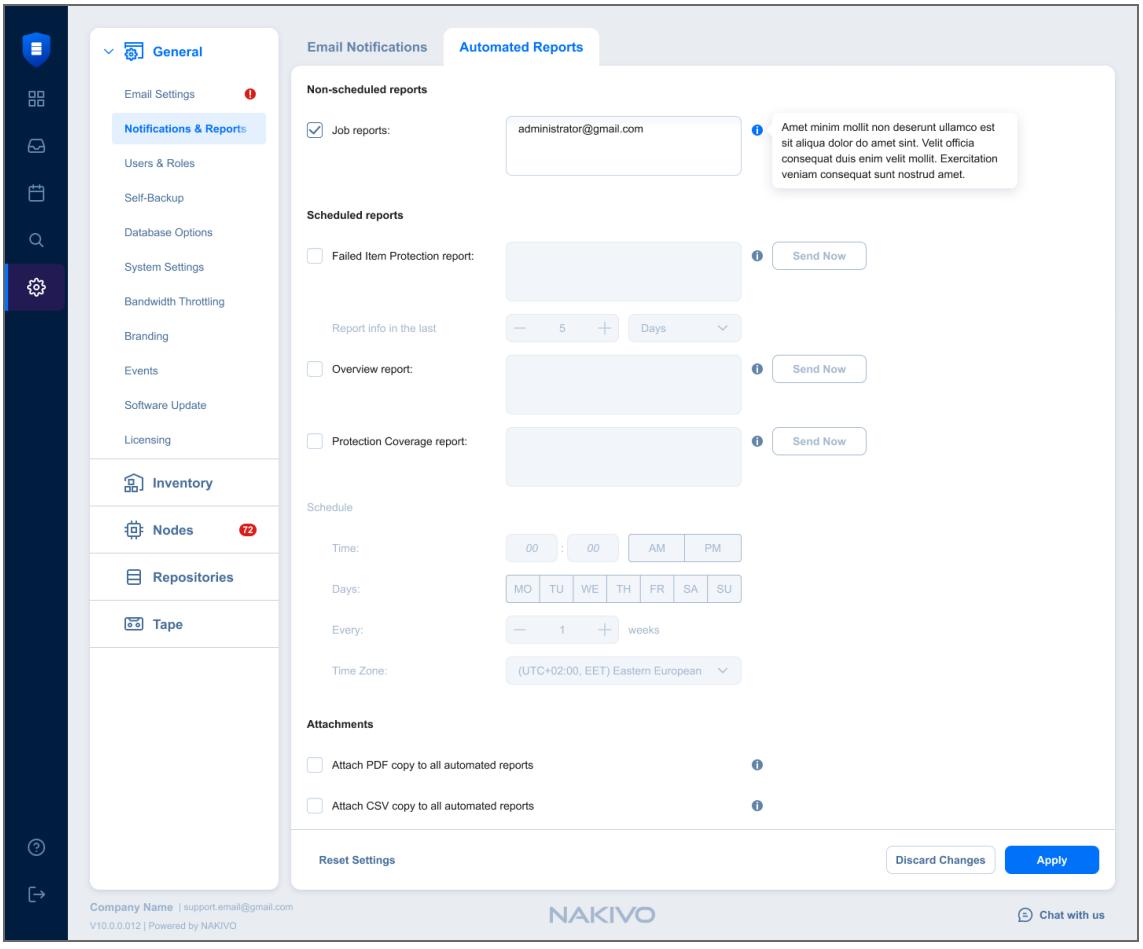
- **Job reports**: If this option is selected, NAKIVO Backup & Replication will send an HTML report after the completion of every job (regardless of the job success or failure) to email addresses specified in the text field. Use a semi-colon to separate multiple email addresses.
- **Failed Item Protection report**: Contains information about all items which had failed to be protected by backup and/or replication jobs, and the error message. Additionally, configure **Report info in the last** option by entering the number of days you want to get the report for.
- **Overview report**: If this option is selected, NAKIVO Backup & Replication will generate the Overview report (which includes information about all jobs and groups in the product) on the date and time specified in the scheduler and will send the report to the recipients specified in the text field. Use a semi-colon to separate multiple email addresses.
- **Protection Coverage**: If this option is selected, NAKIVO Backup & Replication will generate the Protection Coverage report. This includes information about all VMs & instances protected by backup and/or replication jobs as well as the information about all unprotected VMs & instances. The report will be sent to the recipients specified in the text field on the date and time specified in the scheduler. Use a semi-colon to separate multiple email addresses.
- **Schedule**: Configure the schedule at which you want to get the reports.

- **Attach PDF copy to all automated reports**: Select this option to get the additional attached copy of the report in the PDF format.
- **Attach CSV copy to all automated reports**: Select this option to get the additional attached copy of the report in the CSV format.

**Note**

NAKIVO Backup & Replication supports the following special characters in reports:

- US special characters
- Characters in the following languages:
  - Vietnamese
  - Japanese
  - Korean
  - Chinese
  - Arabic



# Self-Backup

The self-backup feature allows you to automatically protect configuration settings of your NAKIVO Backup & Replication instance. For more information, refer to “Self-Backup Feature” on page 63.

**Notes**

- Self-backup is not supported for the multi-tenant configuration.
- A federated repository cannot be used for self-backup.

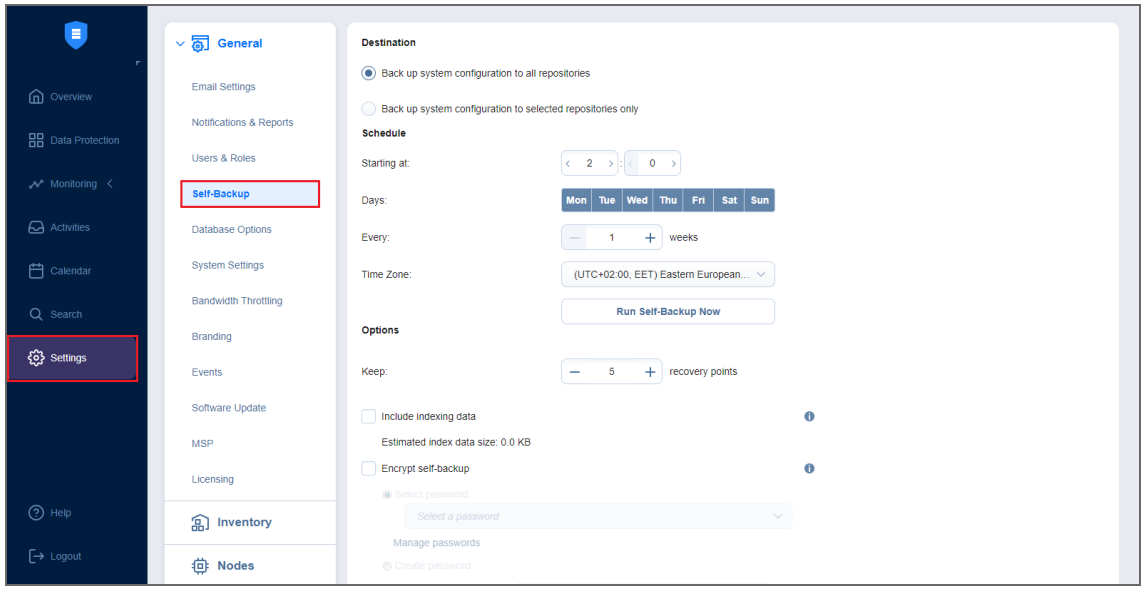
To configure self-backup options, proceed as described in the following sections:


- Accessing Self-Backup Options
- Setting Up Self-Backup Destination
- Self-Backup Schedule
- Self-Backup Options
- Self-Backup Encryption
- Recovering from Self-Backup


## Accessing Self-Backup Options

To access self-backup options, follow the steps below:

1. Click **Settings** in the left pane of NAKIVO Backup & Replication.
2. Go to the **General** tab and click **Self-Backup**.
3. After making the necessary changes, click **Apply**. Alternatively, click **Discard Changes** to discard any changes you have made.



## Setting Up Self-Backup Destination

To configure a self-backup destination, follow the steps below:

1. Select **Back up system configuration to all repositories** to enable all repositories in the list of repositories where system configuration is backed up. If deselected, you can remove specific repositories from the list.

   **Important**

   - Backing up your NAKIVO Backup & Replication system configuration to a DD Boost storage unit Backup Repository causes the DD Boost storage unit to be unmounted. Therefore, to avoid re-adding the DD Boost storage as an existing **Backup Repository** manually, exclude DD Boost storage unit repositories from the list of repositories for self-backup.
   - Self-backup cannot be performed to **SaaS** type of **Backup Repository**.

2. Alternatively, select **Back up system configuration to selected repositories only** and select specific repositories you wish to use for self-backup.
3. If necessary, add a **Backup Repository** to the list:
   - Click the "**+**" icon to add repositories to the list of repositories for system backing up.
   - In the **Add Backup Repositories** dialog that opens, select the necessary repositories and close the dialog.

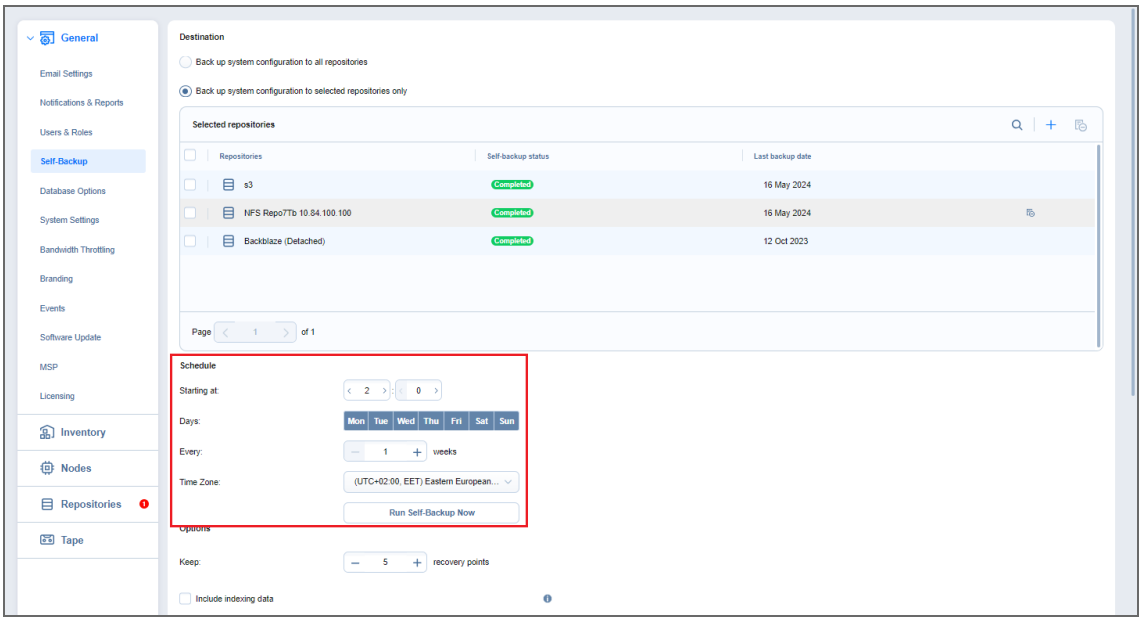
**Notes**

- Federated repositories are disabled on the list of repositories used as a target for self-backup.
- Federated repository members are not displayed on the list of repositories available for self-backup.

## Self-Backup Schedule

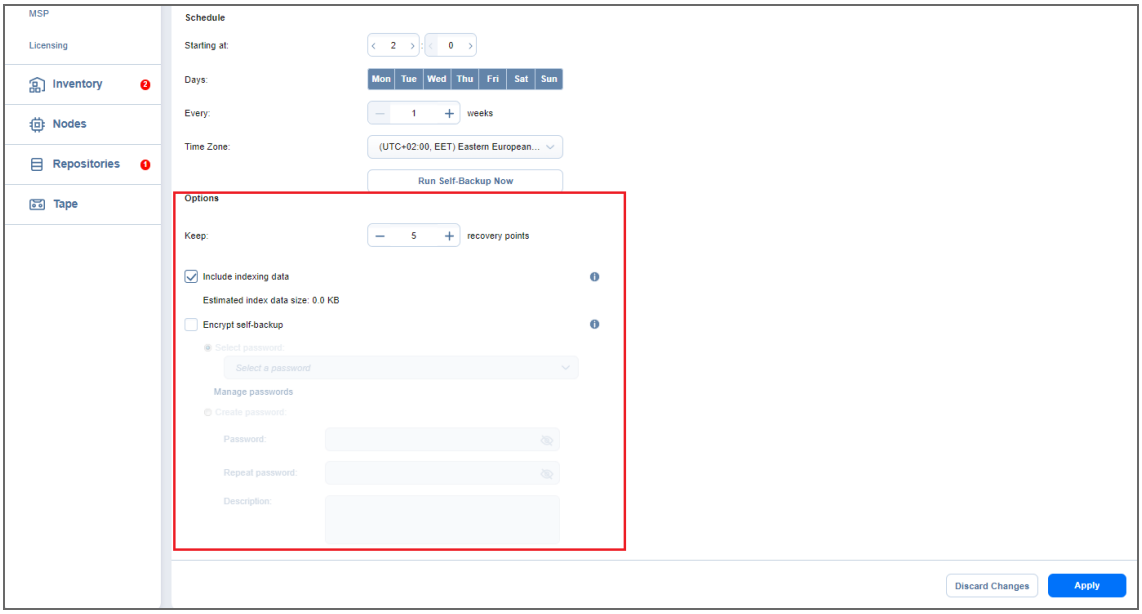
To configure the self-backup schedule, follow the steps below:

1. In the **Schedule section**, enter time to trigger starting the self-backup. You can choose a specific time zone from the list, enter the hours and minutes of the day, and select the necessary days of the week.
2. If you need to start the self-backup immediately, click **Run Self-backup Now**.
3. When ready with configuring the self-backup schedule, click **Apply**.




# Self-Backup Options

In the **Options** section of the self-backup settings, you can enter a number of recovery points to be kept for the self-backup. To apply your settings, click the **Apply** button.




# Self-Backup Encryption

A system configuration stored as a self-backup can be encrypted to safeguard data from loss, corruption, or unauthorized access to data.

Select **Encrypt self-backup** to encrypt your backup for additional security. Afterwards, select the password or create a new one in the corresponding fields. See Managing Passwords for more details.

**Note**

AWS Key Management Service is not applied to self-backup encryption.

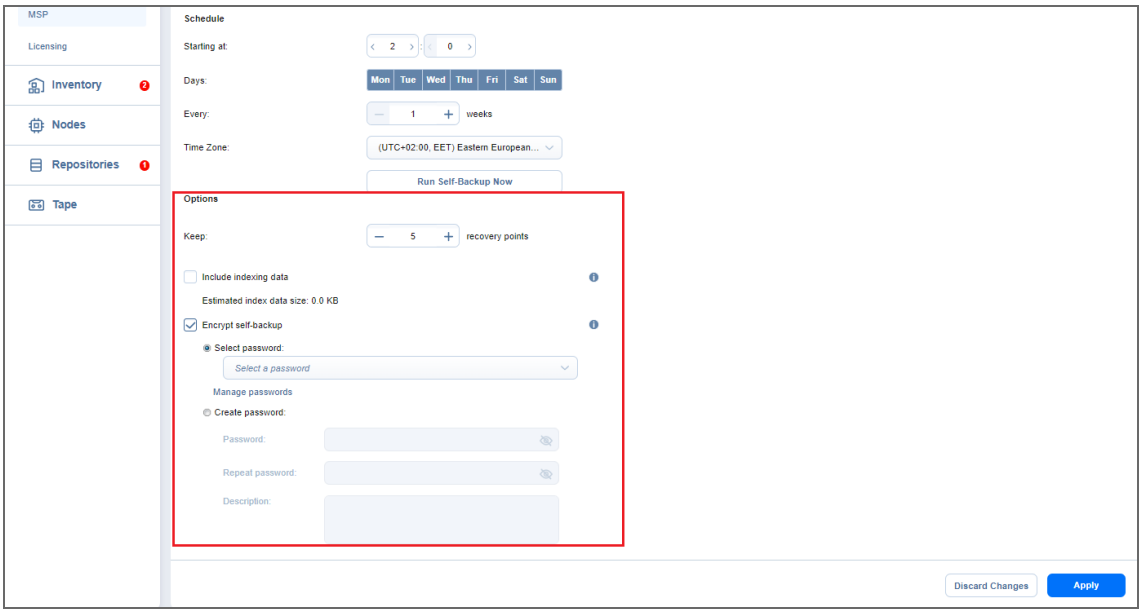
Set up self-backup encryption as follows:

1. Select the **Encrypt self-backup** checkbox (disabled by default) to store self-backups in an encrypted form.
2. Select the **Select password** option and click on the field to display a list of saved passwords with descriptions.
3. Select the needed password.

**Note**

The **Select Password** option is disabled if no passwords are available.

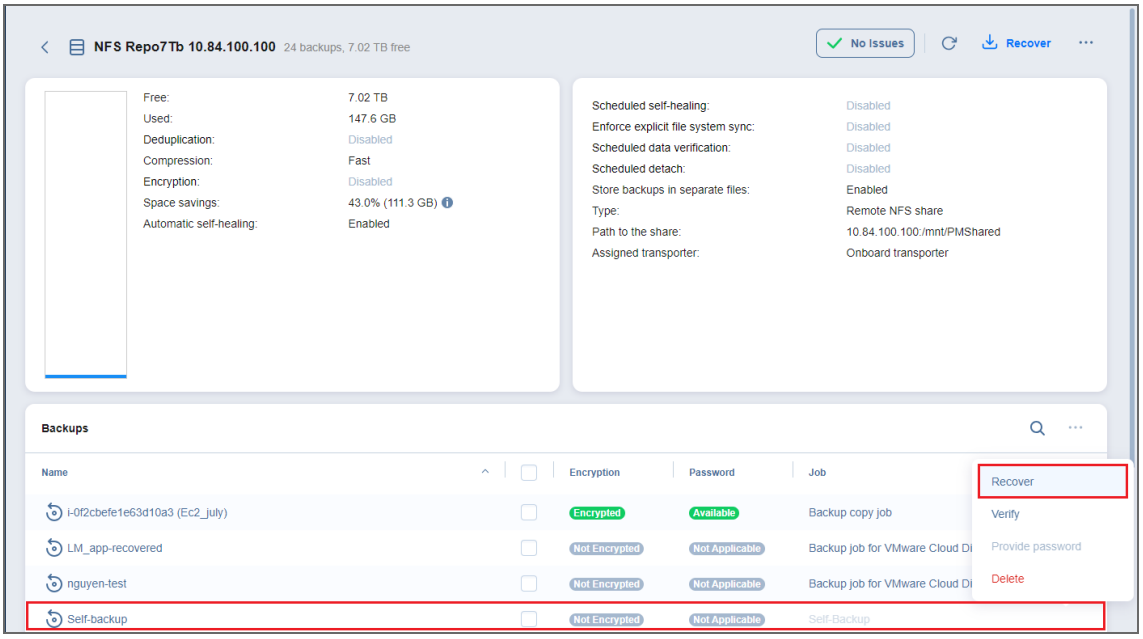
4. To create a new password, select the **Create password** option.
5. Enter the new password and its description and repeat it.
6. If a newly created password is applied, the **Select password** option becomes selected and the new password is pre-selected in the **Select password** dropdown and is added to the list of saved passwords.
7. Click the **Cancel** or **X** button to close the dialog box without applying any changes.
8. Optionally, you can click the **Manage passwords** link to manage the existing or add a new password to the list of passwords. For more information, refer to Managing Backup Encryption.




# Recovering from Self-Backup

To recover the configuration of NAKIVO Backup & Replication from a self-backup stored in a **Backup Repository**, do the following:

1. Go to **Settings** > **Repositories**.
2. Select one of the repositories that contain a self-backup.
3. Select the self-backup from the **Backups** list and click **Recover**.




4. Click **Restore.**
5. Wait for the system configuration to be restored. When the self-backup recovery process is completed, a message announcing success appears.

> **Note**
>
> After a system configuration backup, the encrypted recovery points passwords that were created and saved in the password manager become unavailable and providing them manually is required.

# Software Update

- Download & Update Option
- Download Option

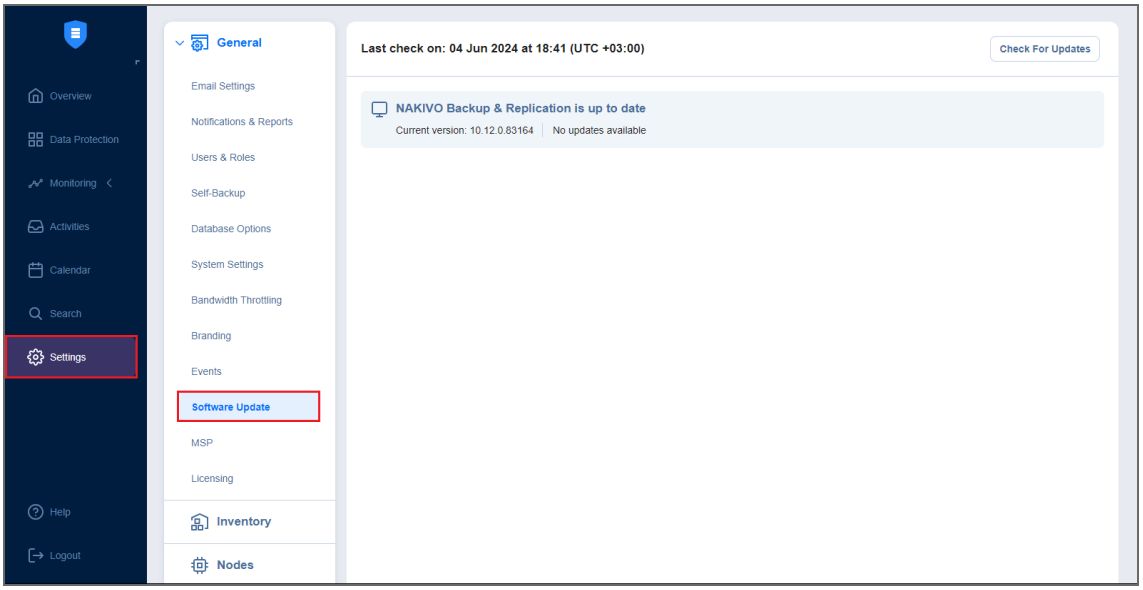
When the full solution of NAKIVO Backup & Replication (that is, the Director and the Transporter) is installed on a Windows or Linux machine, you can download product updates from the **Software Update** tab in the web interface. This feature automatically updates your NAKIVO Backup & Replication instance, the Onboard Transporter, and any other nodes that support auto-update.

For a list of supported nodes and requirements for the auto-update feature, see the **Auto-Update** section in "Feature Requirements" on page 169.

> **Note**
>
> In case **I/O Filter** was installed on the Cluster before the software update and Auto Refresh Inventory is disabled, manually refresh the **Inventory** after an software update to check that **I/O Filter Daemons** are displayed in the **Nodes** tab, and ensure Real-Time Replication jobs run successfully.

To check if an update is available, do the following:

1. Click **Settings** in the left pane of the product.
2. Go to the **General** tab.
3. Go to the **Software Update** page.
4. Click **Check for Updates** if needed.



**Note**

If you are using a multi-tenant solution, only master-tenant users with the appropriate permissions can see and manage software updates.

## Download & Update Option

To download and install the update, do the following:

1. Optionally, click **Release Notes** to see features and improvements implemented in the new product version.
2. Select the **I have read the Release Notes** checkbox.
3. Click **Download & update**.
4. Click **Update Now**.

Before downloading the update, the product performs a self-backup and stops all current activities including running jobs, recovery jobs, repository maintenance, etc. When the download is complete, the product updating process begins. The product downloads the update to the **Director** first. When the **Director** is updated, the update is automatically uploaded to the **Transporters** which are then updated simultaneously. If some **Transporters** are not updated, you can update them manually. Refer to the corresponding articles for details.

**Notes**

- For a list of supported **Transporters**, see the **Auto-Update** section in "Feature Requirements" on page 169.
- Only 20 **Transporters** can be updated simultaneously. All other **Transporters** will be sent to a queue and updated once the previous update is completed.

## Download Option

If you wish to postpone an update or schedule it, take the following steps to download the update without installing it:

1. Optionally, click **Release Notes** to see features and improvements implemented in the new product version.
2. Select the **I have read the Release Notes** checkbox.
3. Click **Download**.
4. After the download is completed, do one of the following:
    - Click **Update Now** if you want to start the updating process. Updating the product will stop all current activities, including running jobs, recovery jobs, repository maintenance, etc.
    - Click **Schedule Update** to update the solution at a specific time:

1. In the dialog box that opens, pick a day and time for updating. Click **Apply**.
2. On a working day before the scheduled update, you will see a notification in the product menu with the **Update Reminder** dialog box. By hovering over this notification, you can:
    a. Click **Reschedule** if you want to reschedule the update and pick a different time.
    b. Click **Cancel update** to cancel updating the full solution.

> **Note**
> A notification about the update will also be sent to your email if email settings are configured.

# System Settings

To configure the system settings, follow the steps below:

1. Click **Settings** in the main menu on the left.
2. Go to the **General** tab and click **System settings.**
3. Set the following options:
    - In the **Configuration** tab:
        - **Store system events for the last x days**: Events older than the specified number of days (can be from 5 to 365) are deleted.
        - **Store job history for the last x days**: The history of the jobs older than the specified number of days (can be from 5 to 90) is deleted.

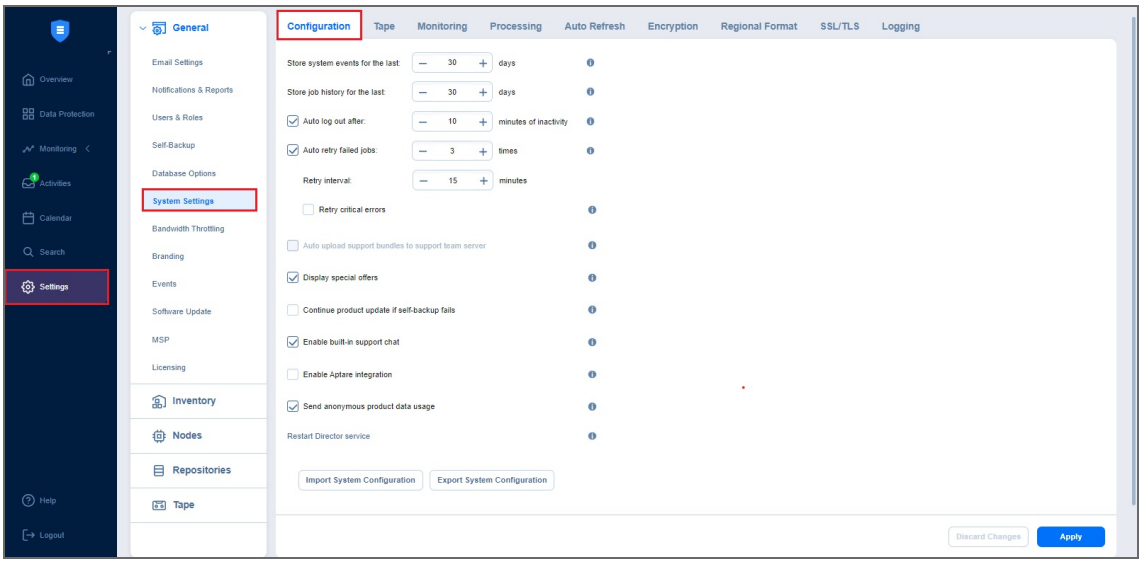
            **Note**
            This option is not displayed for the Master tenant in Multi-tenancy mode.

        - **Auto log out after x minutes of inactivity**: When this option is selected, the current user is automatically logged out of NAKIVO Backup & Replication after the specified period of inactivity.
        - **Auto retry failed jobs x times with y minutes interval**: When this option is selected, failed jobs are automatically retried the specified number of times (from 2 to 10) and with the specified time interval (from 1 to 60). Jobs with failed backup, replication, and recovery remain in the “running” state until all retries have either succeeded or failed.
        - **Retry critical errors**: When this option is selected, NAKIVO Backup & Replication tries to automatically rerun jobs with critical and non-critical errors a specified number of times.

            **Notes**
            - The term **critical error** refers to persistent errors that are unlikely to change without any additional intervention, that is, hardware failure.
            - The term **non-critical error** refers to non-persistent errors that are likely to change without any additional intervention, that is, unstable network connection.

        - **Auto upload support bundles to support team server**: When this option is selected, NAKIVO Backup & Replication automatically creates, encrypts, and uploads support bundles once a day to a NAKIVO support server. The NAKIVO Support Team may use this information to improve the product experience and to identify and resolve product issues faster.

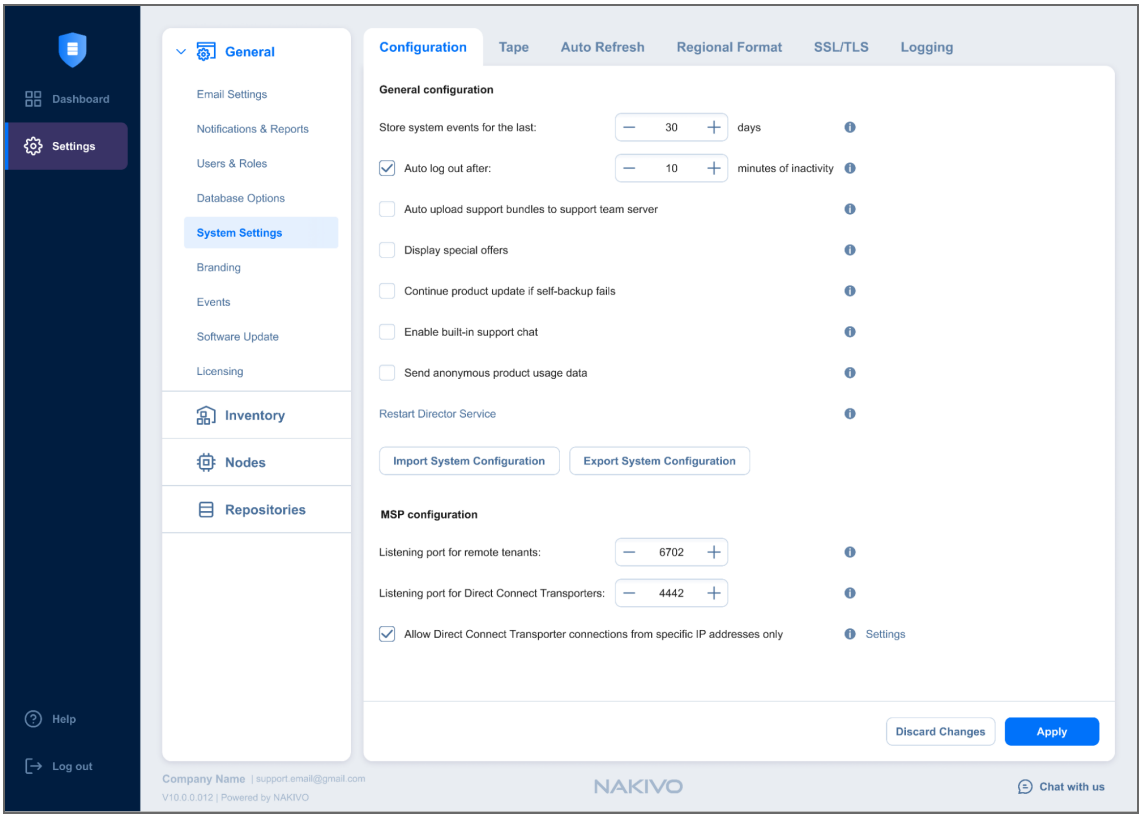
- **Display special offers**: When this option is enabled, the NAKIVO special offers toolbar appears in the NAKIVO Backup & Replication interface.
- **Continue product update if self-backup fails**: When this option is selected, updates proceed even if self-backup cannot be performed.
- **Enable built-in support chat**: When this option is selected, you can contact a NAKIVO representative via chat in the NAKIVO Backup & Replication interface. When selected in the multi-tenant mode, the built-in support chat is available to all tenants of the NAKIVO Backup & Replication instance.
- **Enable Aptare Integration**: Select this option to integrate the APTARE storage resource management platform with NAKIVO Backup & Replication. For integration details, refer to Aptare IT Analytics Integration.
- **Send anonymous product usage data**: Enable this option to send anonymous product usage data to NAKIVO for efficient product development and enhancement. Note that no personal data is collected.
- Click the **Use New Scheduler** link to enable the use of a new scheduler that merges the retention and schedule steps. The scheduler allows you to set backup retention settings per schedule and get expiration dates for recovery points.
- You can click **Restart Director service** to stop all current activities and restart the **Director**. After clicking the link, a confirmation window appears. Click **Reboot** to confirm the restart.
- **Import System Configuration**: Find more information on the topic here.
- **Export System Configuration**: Find more information on the topic here.



- You can also configure the following options in **Multi-tenant mode**:
  - **Listening port for remote tenants**:
    1. Specify the listening port for remote tenant connections in a multi-tenant environment (6702 is the default port). You can change this option only in the master tenant.
    2. Restart the **Director** to apply the changes.
  - **Listening port for Direct Connect Transporters**:
    1. Specify the listening port for a **Direct Connect Transporter**. The allowed range is 1 – 65,535 (4442 is the default port).

       > **Note**
       > The port must be different from the **MSP Director** port configured in the installer. You can change this option only in the master tenant.

    2. Restart the **Director** to apply the changes. For more information, see MSP Direct Connect.
  - **Allow Direct Connect Transporter connections from specific IP addresses only**: Select the checkbox to allow only whitelisted Direct Connect Transporters to connect to the MSP Director.
  - **Settings**: Click the link to open the Direct Connect Transporter IP whitelist dialog box.



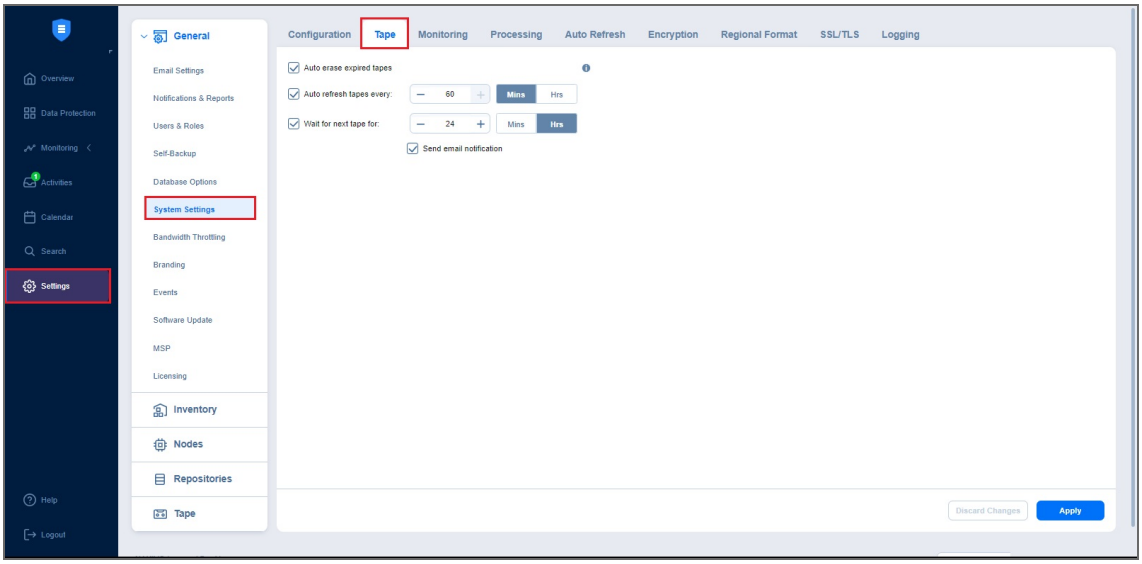
- In the **Tape** tab:
  - **Auto erase expired tapes**: When this option is selected, expired tapes are erased automatically.

    **Important**

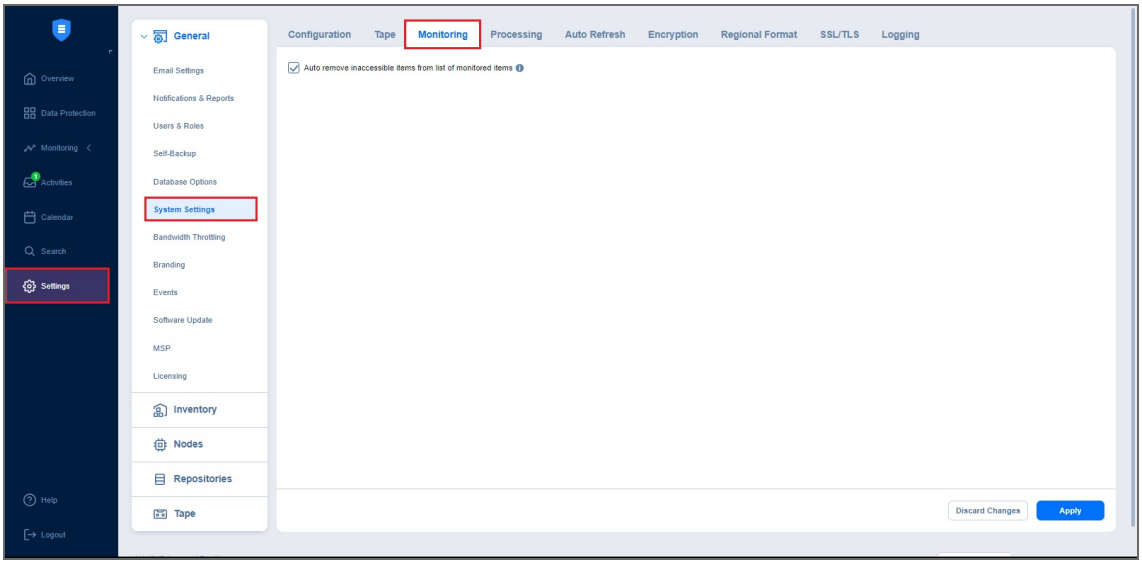
    When this option is selected, the following prerequisites must be met for a cartridge to be erased:

    - All recovery points within the tape cartridge have expired.
    - There are no dependent recovery points on other tape cartridges.
    - The product keeps at least one full chain of recovery points.

  - **Auto refresh tapes every**: Select how often the contents of the tapes are refreshed in minutes or hours. Deselect if refreshing is not required.
  - **Wait for next tape for**: Specify how long the system should wait for the next tape if there is no appropriate amount. Select the **Send email notification** checkbox to receive email notifications.




- In the **Monitoring** tab:
  - **Auto remove inaccessible items from list of monitored items**: When this option is selected, all inaccessible items are removed automatically from the list of monitored items.

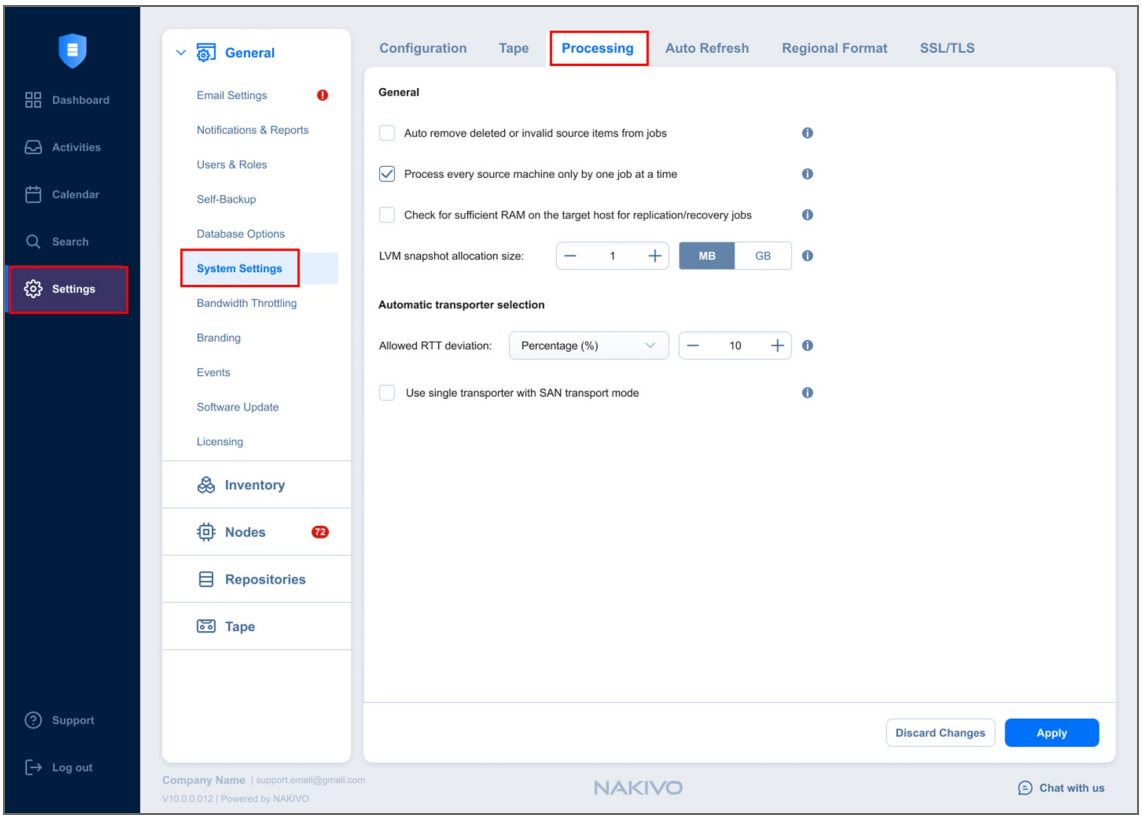



- In the **Processing** tab:
  - **General**:
    - **Auto remove deleted or invalid source items from jobs**: This option applies to a protected container (such as a VMware cluster or EC2 region). When this option is selected, if during an inventory refresh, NAKIVO Backup & Replication discovers that at least one VM and/or EC2 instance is no longer available in the protected container, these VMs and EC2 instances are automatically removed from all jobs.
    - **Process every source item only by one job at a time**: When this option is selected, all machines in backup and replication jobs are processed by one job at a time only. Running jobs and respective source objects are not affected after changing this setting. This option is always enabled for physical servers and real-time replication jobs.
    - **Check for sufficient RAM on the target host for replication/recovery jobs**: When this option is deselected, NAKIVO Backup & Replication does not check whether the amount of RAM on the target host is sufficient for replication and recovery jobs.
    - **LVM snapshot allocation size**: This option allows you to set an LVM allocation snapshot size for a Linux physical server backup. The default size is 1 GB. The maximum size is 1000 GB.
  - **Automatic transporter selection**: This section allows you to set the target transporter as a source transporter when finding the best data transfer path for a job object. To use a single transporter as a target and as a source, check the Use single transporter with SAN transport mode checkbox.

**Notes**

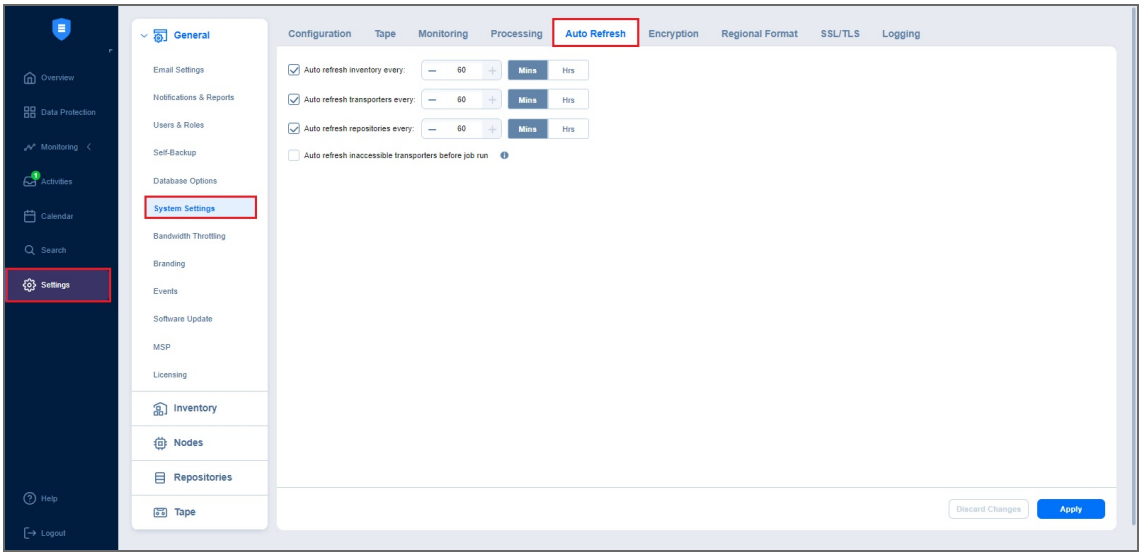
- This setting can be applied to the jobs with the Transporters option set to Automatic selection (Options step, DataTransfer section).
- This setting can be applied to jobs using the SAN transport mode.

- **Allowed RTT deviation**: This option allows you to define the acceptable range for selecting data paths during job execution. You can specify the deviation as a percentage (%) or a threshold (ms). This setting helps you optimize job performance and evenly distribute the workload across transporters. When a transporter is selected automatically, the data path with the best RTT is used during the job. If load balancing is needed during data transfer, data paths with RTT values within the specified range are chosen to ensure smooth operation.
- **Use a single transporter with SAN transport mode**: Select this checkbox to limit data transfer to a single transporter that is working as the target and source transporter when using SAN transport mode with automatic transporter selection enabled.




- In the **Auto Refresh** tab:
  - **Auto refresh inventory every X minutes**: Specify how often you want your inventories to be refreshed.

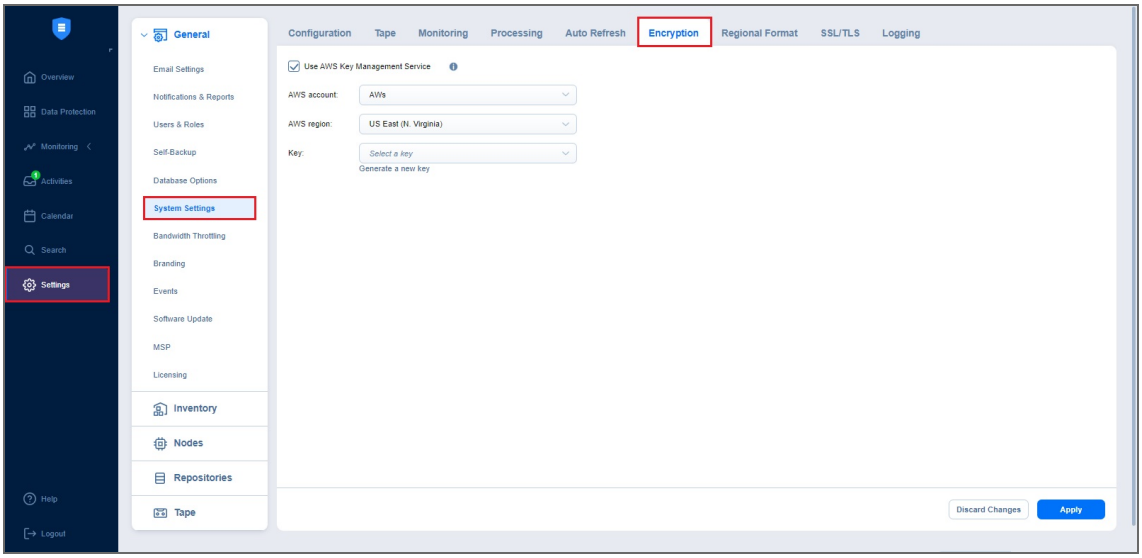
- **Auto refresh transporters every X minutes**: Specify how often you want your transporters to be refreshed.
- **Auto refresh repositories every X minutes**: Specify how often you want your repositories to be refreshed.
- **Auto refresh inaccessible transporters before job run**: If this option is enabled, the product refreshes all inaccessible transporters before a job is processed.




- In the **Encryption** tab:
  - **Use (AWS) Key Management Service**: If enabled, all backup encryption passwords encrypted with the Key Management Service cryptographic key are available for recovery in case of product re-installation.

    You can select an existing symmetric cryptographic key or create a new one.

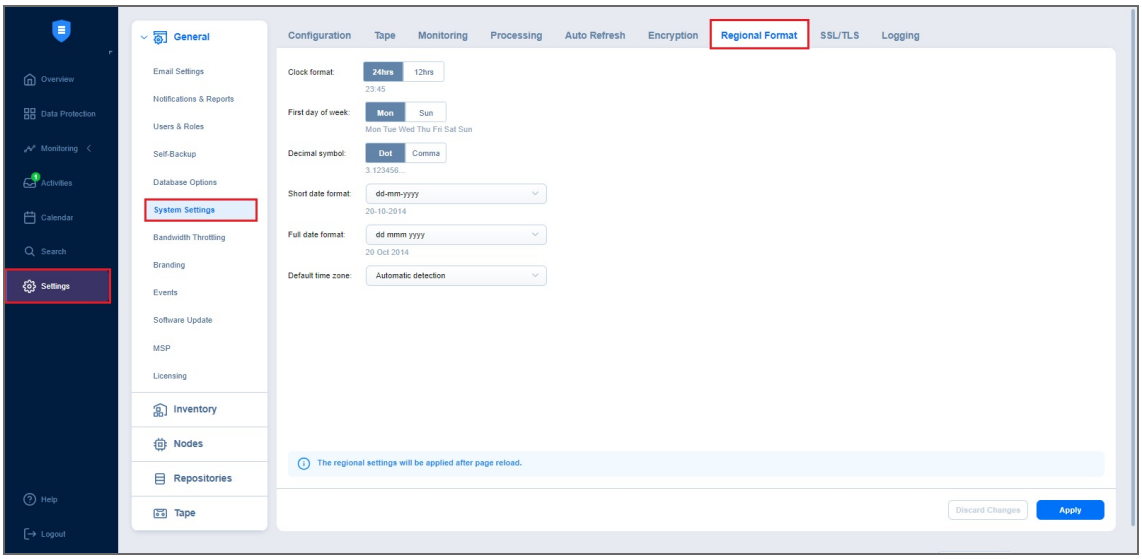
    For more details, refer to Managing Backup Encryption.



**Notes**

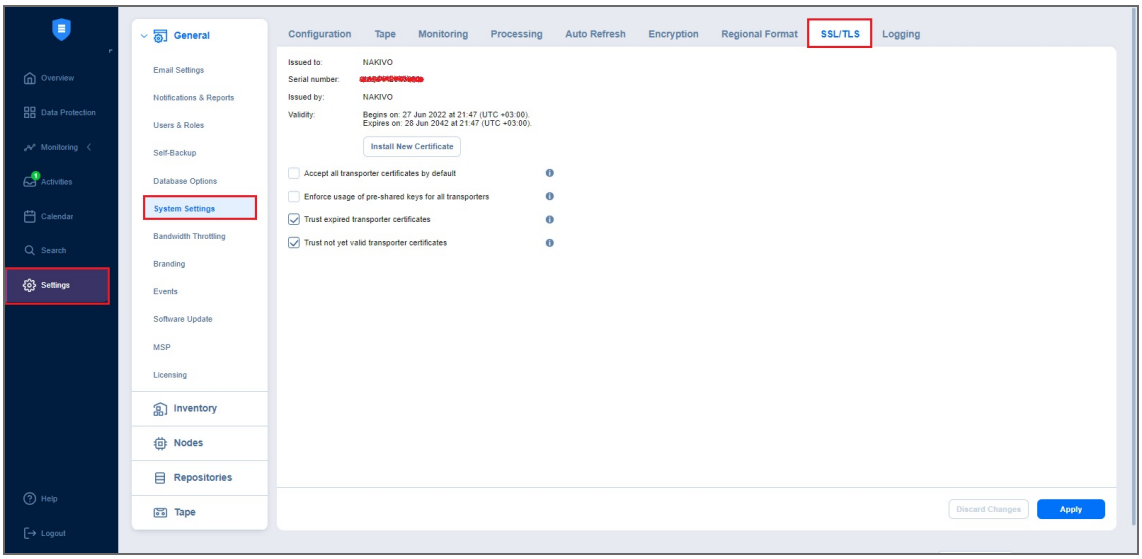
- The **Encryption** tab is displayed at the Tenant level and in single tenant mode and is hidden at the Master level.
- The passwords are not propagated from one tenant to another.
- In case encryption is set for a job, the password hash is generated based on the configured password.
- The password hash gets encrypted with the KMS cryptographic key and is saved in the recovery point metadata.
- The AWS Key Management Service is not applied to self-backup and system configuration encryption. For more details, refer to Self-Backup Encryption and System Migration.

- In the **Regional Format** tab, set:
  - **Clock format**
  - **First day of week**
  - **Decimal symbol**
  - **Short date format**
  - **Full date format**
  - **Default time zone**




- In the **SSL/TLS** tab, you can either:
  - **Install new certificate**: A dialog opens allowing you to install a new TLS/SSL certificate for the NAKIVO Backup & Replication web interface. Certificates are generated either internally or through certification authorities. Proceed as follows to install a new certificate:

- Click **Browse** and navigate to the location of either of the following certificate file types:
  - **Private key**: A file in the *.key format.
  - **Private key password (optional)**: A password for your private key.
  - **Certificate file**: A file in the *.pem, *.crt, *.cer, *.p7b, or *.p7s format.
  - **Intermediate certificate (optional)**: A file in one of the following formats: *.pem, *.crt, *.cer, *.p7b, *.p7s.
- **Accept all transporter certificates by default**: Select this option to automatically accept all transporter certificates. After selecting the option, click **Continue** in the warning popup window that appears to confirm the selection.
- **Enforce usage of pre-shared keys for all transporters**: Selecting this option makes sure that transport function only when pre-shared key is installed.
- **Trust expired self-signed transporter certificates**: Selecting this option makes the solution trust the expired self-signed transporter certificates.



**Notes**

- Repositories assigned to the inaccessible transporters are refreshed as well during the auto refresh.
- The **Auto refresh inaccessible transporters before job run** option is supported only for the following platforms:
  - VMware vSphere
  - AWS
  - Cloud Director
  - Nutanix AHV
- If any time zone other than **(UTC+00:00, UTC) Coordinated Universal Time** is chosen, daylight savings times are honored.

4. After making the necessary changes, click **Apply**. Alternatively, click **Discard Changes** to discard any changes you have made.

**Notes**

- NAKIVO Backup & Replication supports Certificates with the RSA algorithm only.
- In the *Web Interface TLS/SSL Certificate* section, you can see a notification about imminent TLS/SSL Certificate expiration in 30 days and onwards. If your certificate has expired, you will be asked to install a valid certificate.

# System Migration

NAKIVO Backup & Replication provides you with the ability to migrate all your settings (including inventory, jobs, credentials, transporter settings, and so on) to a new instance (copy) of the product.

**Important**

System configuration export and import are designed for migration purposes only, and not to serve as a system configuration backup. After you have exported system configuration from an old instance of the product, do not run jobs in that old instance. Doing so results in failed jobs in the new instance after the migration. All jobs have to be recreated, and full initial job run is required.

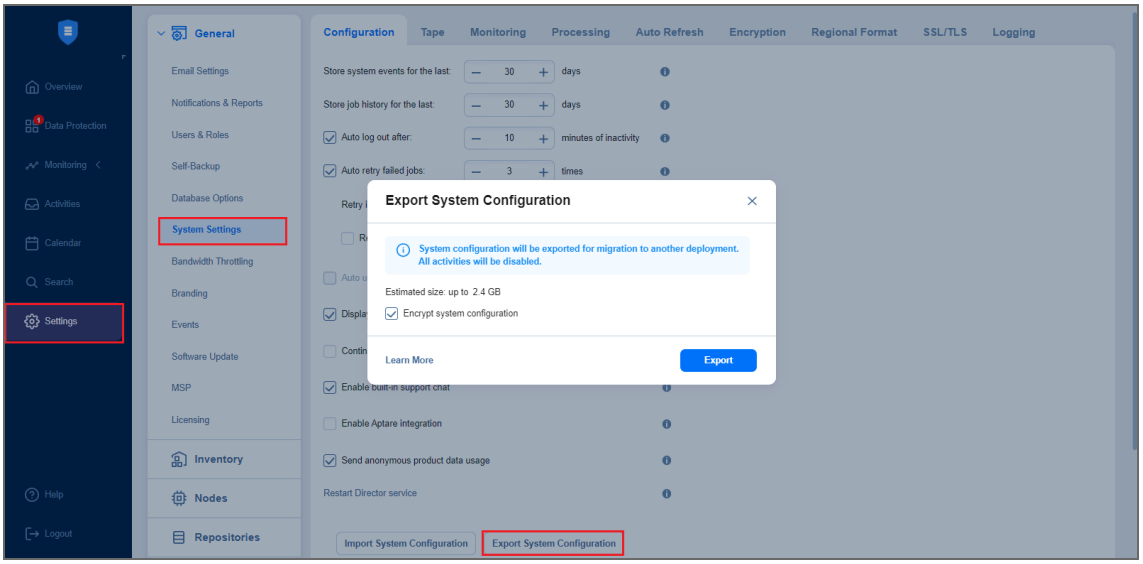
See the topics below for more information:

- Exporting System Configuration
- Importing System Configuration

## Exporting System Configuration

To export system configuration from the old deployment, follow the steps below:

1. Click **Settings** in the left pane of the product.
2. Select **System Settings** tab in the **General** section.
3. On the **Configuration** tab, click **Export System Configuration**.
4. In the dialog window that appears, click **Export**.



**Notes**

- Select the **Include indexing data** checkbox to include the indexing data into the package.
- Selecting the **Include indexing data** checkbox displays the **Estimated index data size** label and includes the size of index data in the displayed data size.
- Optionally, select the **Encrypt system configuration** checkbox.

5. Click **Export** to confirm the operation.
6. If the **Encrypt system configuration** checkbox is selected, clicking the **Export** button opens the **Set a password** dialog box. For more details, refer to Setting Password.

**Notes**

- All activities in the old instance (such as jobs and recovery sessions) are automatically stopped and all jobs are disabled.
- AWS Key Management Service is not applied to system configuration encryption.

7. Wait until the export is completed and download the export bundle.
8. Do not run jobs in the old instance.

## Importing System Configuration

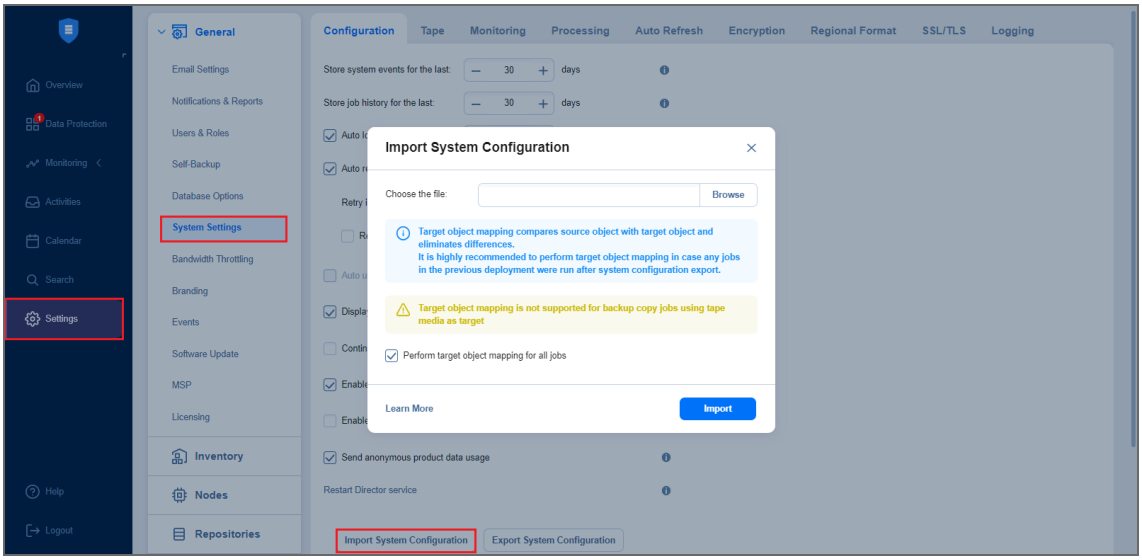
To import system configuration into a new instance of the product, follow the steps below:

1. Click **Settings** in the left pane of the product.
2. Select **System Settings** tab in the **General** section.
3. On the **Configuration** tab, click **Import System Configuration**.
4. In the dialog box that appears, locate the system configuration bundle using the **Browse** button.
5. Click **Import**.

**Note**

If the configuration bundle is encrypted, clicking on the **Import** button opens the **Provide a password** dialog box. In the dialog box that appears, select the needed password or create a new one. For more information, refer to Managing Passwords.

6. Click the **Cancel** or **X** button to close the dialog box without applying any changes.




7. Click **Proceed** to confirm the operation.

**Important**

- If there is any existing data in the new instance, it will be overwritten with the import operation.
- Target object mapping is skipped for backup copy jobs using tape media as target.
- If a physical configuration of your source deployment differs from a target deployment, a **Backup Repository** may become inaccessible after the bundle import is completed.

8. Wait until the import is completed, and close the dialog window.

**Notes**

- **Backup Repositories** are not migrated by the system configuration export and import. If you have a local **Backup Repository** on the old instance of the product, you may want to move it to the new location. After moving the **Backup Repository**, you may need to edit **Backup Repository** settings in the new instance so that the new settings refer to the actual **Backup Repository** location.
- In case a custom TLS/SSL certificate of the Web server was used in the old instance, a manual service restart is required in the new deployment.

# Users and Roles

Accessing NAKIVO Backup & Replication is possible either with a user account created in the product or with an account added to the product from Active Directory. Each user in the product is assigned a role, which is a set of specific permissions.


- Managing Users and Roles
- Navigating Users View
- Navigating Roles View
- "Navigating AD Domains View" on page 472
- Navigating AD Groups View


## Managing Users and Roles

Managing users and roles can be done by following these steps:

1. Log in to NAKIVO Backup & Replication.
2. Click **Settings** (cog icon) in the left pane of the product.
3. Go to the **General** tab and click **Users & Roles**.

## Navigating Users View

To see the list of all users, select the **Users** view in the upper pane. On this page of the solution you can do the following:

- See the list of all users added to NAKIVO Backup & Replication.
- Sort the list by **User Name**, **Role**, **Tenant**, **Domain**, **Access level**, **Group** or **2 Factor-authentication** by clicking on the respective name of the column.
- Search for users by entering the name of the user (fully or partially) into the **Search** bar.

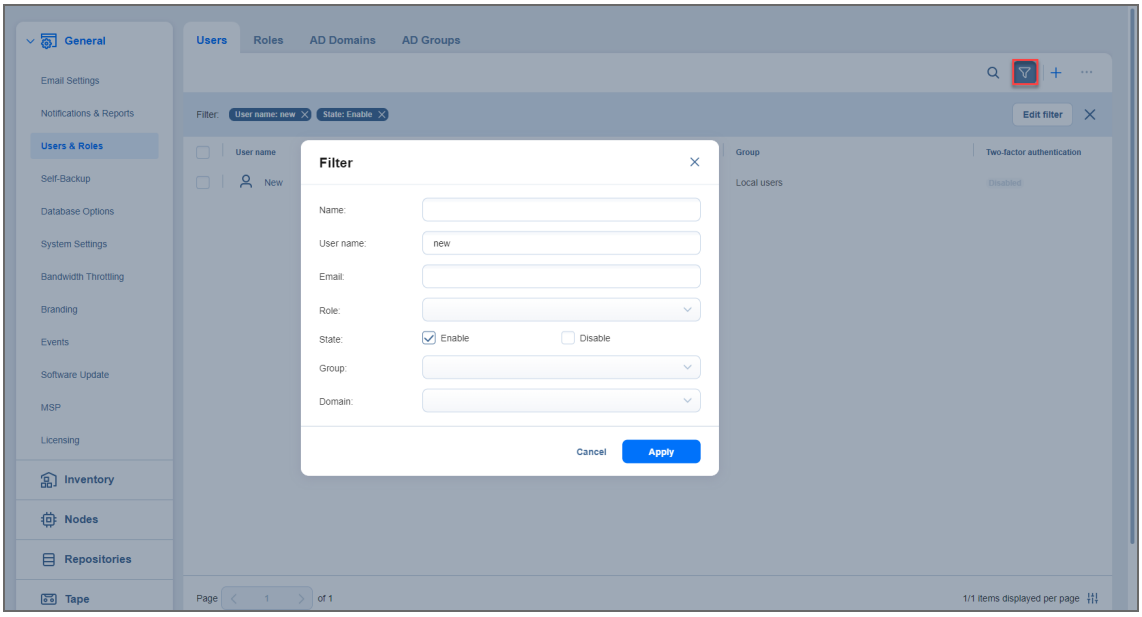
  **Notes**

  - The search phrase must be contained in the user registration name, username, or email address.
  - The **Search** option allows you to perform search on the following columns:
    - **User name**
    - **Role**
    - **2 Factor-authentication**
    - **Access level**
    - **Group**

- Filter the list of users by selecting the **Filter** option.

**Note**

Clicking **Filter** opens a new window that allows you to filter the list of local users according to **User name**, **Role**, **State**, **Group**, and **Domain**.




- Add a new local user by clicking the **+** icon.
- Delete, disable, enable two-factor authentication, and assign a new role to the local user individually. These actions, except **Edit**, can also be done in bulk by checking the box in the upper left pane to select all users and clicking the **...** (ellipsis) button.

**Notes**

- When selecting all local users to apply a bulk action, NAKIVO Backup & Replication selects only those users that are displayed on the screen.
- When a user is deleted, the name of the user becomes *Deleted* in places where this user is referenced to.
- You can create another user with the same username as the deleted user.

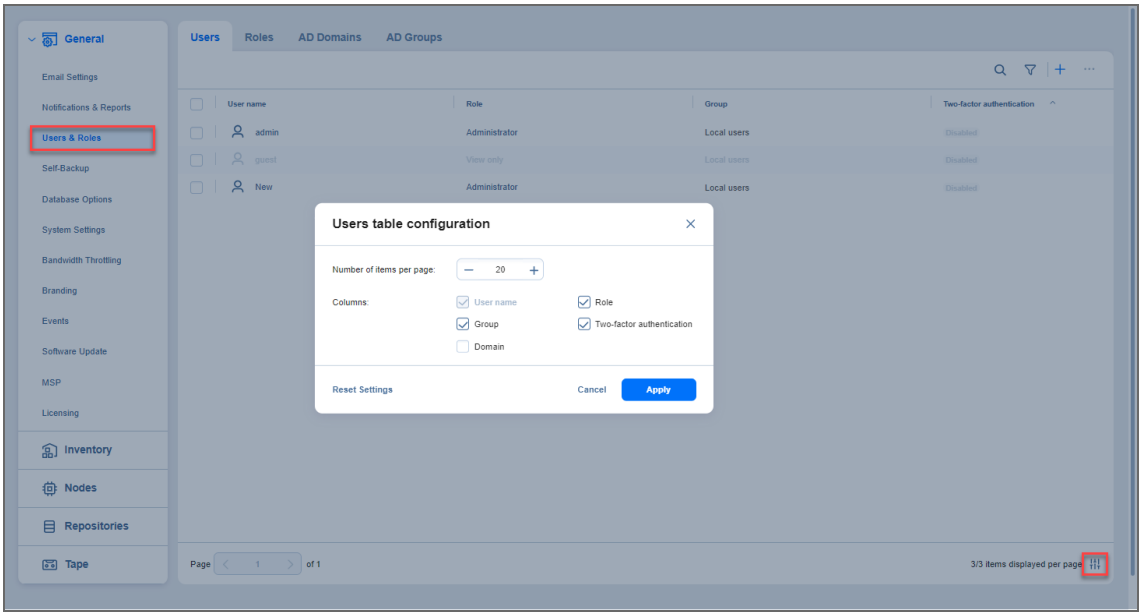
- Edit the role assigned to the local use by clicking on the name of the role in the respective column.

## Configuring Users Table

You can show/hide columns or modify the number of items per page in the **Users** table. In the lower right corner, click the controls icon. In the dialog window that opens, select/deselect checkboxes and click **Apply**. The following options are available:

- User name
- Group
- Domain

- Role
- Two-factor authentication




**Note**

You cannot modify the **User name** column in this dialog window. Refer to "Editing Local User" on page 490 to learn more about editing users.

## Navigating Roles View

To see the list of all local users, select the **Roles** view in the upper pane. On this page of the solution you can do the following:

- See the list of all user roles added to NAKIVO Backup & Replication.
- Sort the list by **Role name**, **Access level**, or **Number of users** by clicking on the respective name of the column.

  **Note**

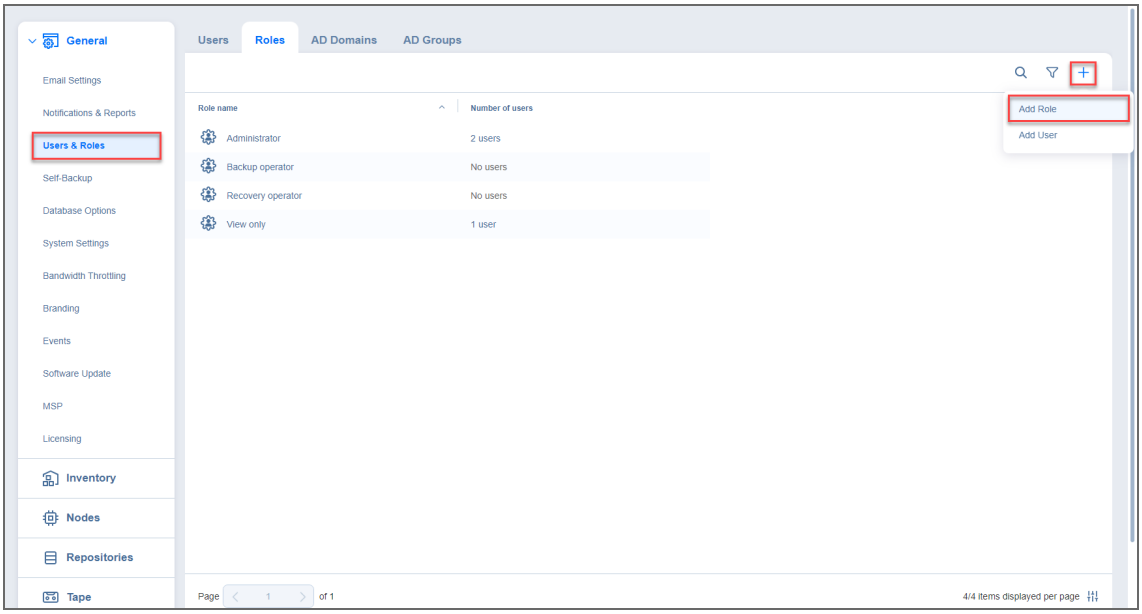
  The **Access level** column is displayed only for the **Master tenant** in Multi-tenant mode. It displays the access level that the role has.

- Filter the list of users by entering the name of the user fully or partially into the **Search** bar or by selecting the **Filter** option.

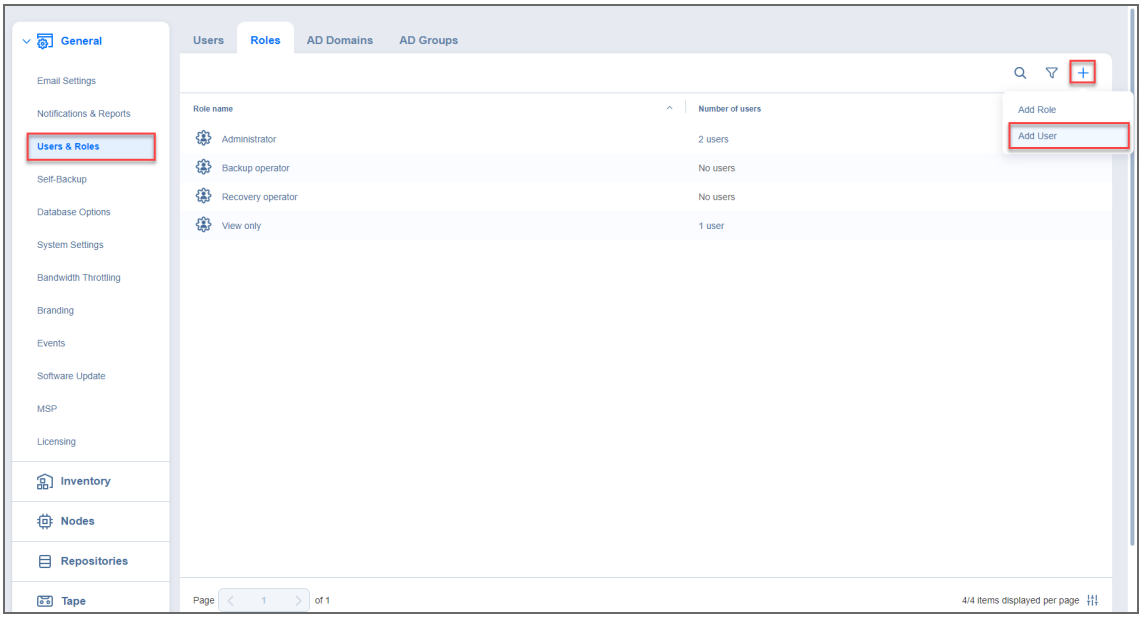
  **Note**

  Clicking **Filter** opens a new window that allows you to filter the list of local users according to **Role name** or **Number of users**.

- Add a new role by clicking the **+** icon and selecting **Add Role**.




- Add a new local user by clicking the **+** icon and selecting **Add User**.

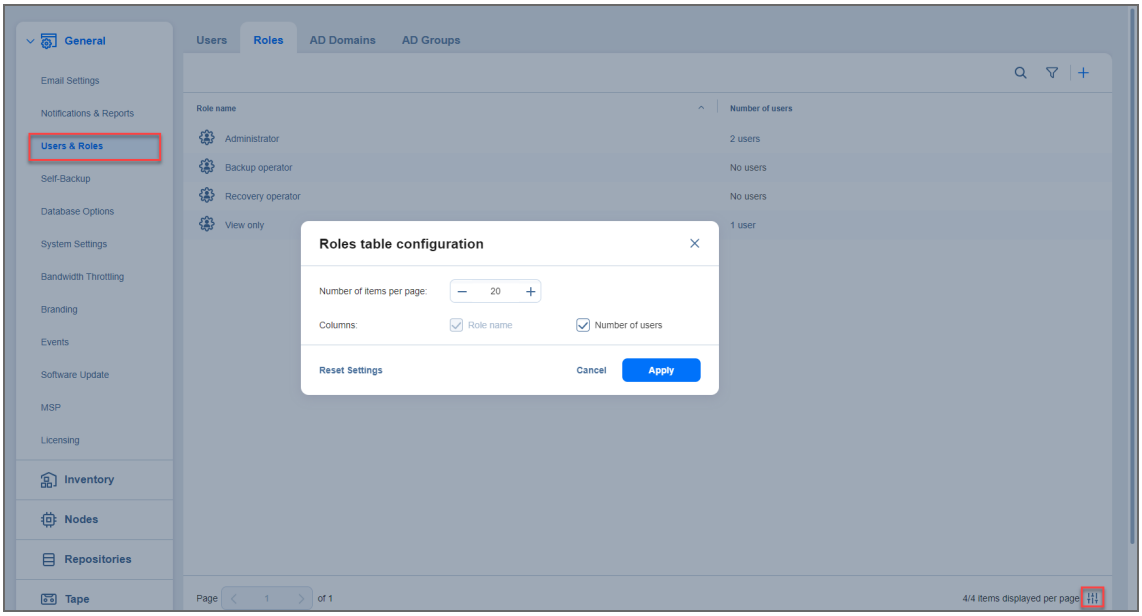



- Edit, delete, or clone the user roles individually. These actions, except **Edit**, can also be done in bulk by checking the box in the upper left pane to select all users and clicking the **...** (ellipsis) button. When selecting all local users to apply a bulk action, NAKIVO Backup & Replication selects only those users that are displayed on the screen.
- **Edit**, **Delete**, **Clone** the role by clicking the ellipses to the right of the role's name.

## Configuring Roles Table

You can show/hide columns or modify the number of items per page in the **Roles** table. In the lower right corner, click the controls icon. In the dialog window that opens, select/deselect checkboxes and click **Apply**. The following options are available:

- Role name
- Number of users



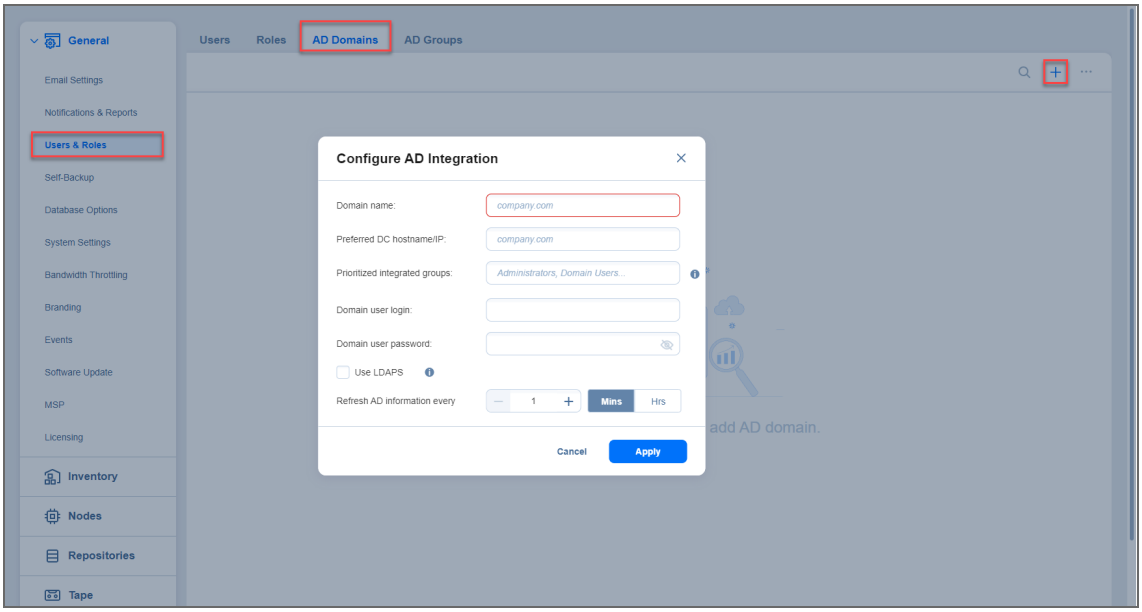

**Note**

You cannot modify the **Role name** column in this dialog window. Refer to "Editing User Role" on page 497 to learn more about editing alert user roles.

## Navigating AD Domains View

To see the list of all Active Directory domains, select the **AD  Domains** view in the upper pane. On this page of the solution you can do the following:

- See the list of all AD domains added to NAKIVO Backup & Replication.
- Sort the list by **Domain name**, **DC/Hostname IP**, **Prioritized integrated groups**, **Domain user**, **Refresh AD** or **Status** by clicking on the respective name of the column.
- Search the list of AD domains by entering the name of the user fully or partially into the **Search** bar.
- Add a new AD domain by clicking the **+** icon.




> **Note**
>
> You can add and then manage two or more AD domains and add AD groups from multiple domains. Refer to Configuring Active Directory Integration for more details.

- Edit, delete, disable/enable AD domains. These actions, except **Edit**, can also be done in bulk by checking the box in the upper left pane and clicking the **...** (ellipsis) button.

> **Note**
>
> The user associated with the disabled AD domain cannot log in to the solution.

## Navigating AD Groups View

To see the list of all Active Directory groups, select the **AD Groups** view in the upper pane.

On this page of the solution you can do the following:

- See the list of all AD groups added to NAKIVO Backup & Replication.
- Sort the list by **Group name**, **Logged in users**, **Access level**, or **Role** by clicking on the respective name of the column.

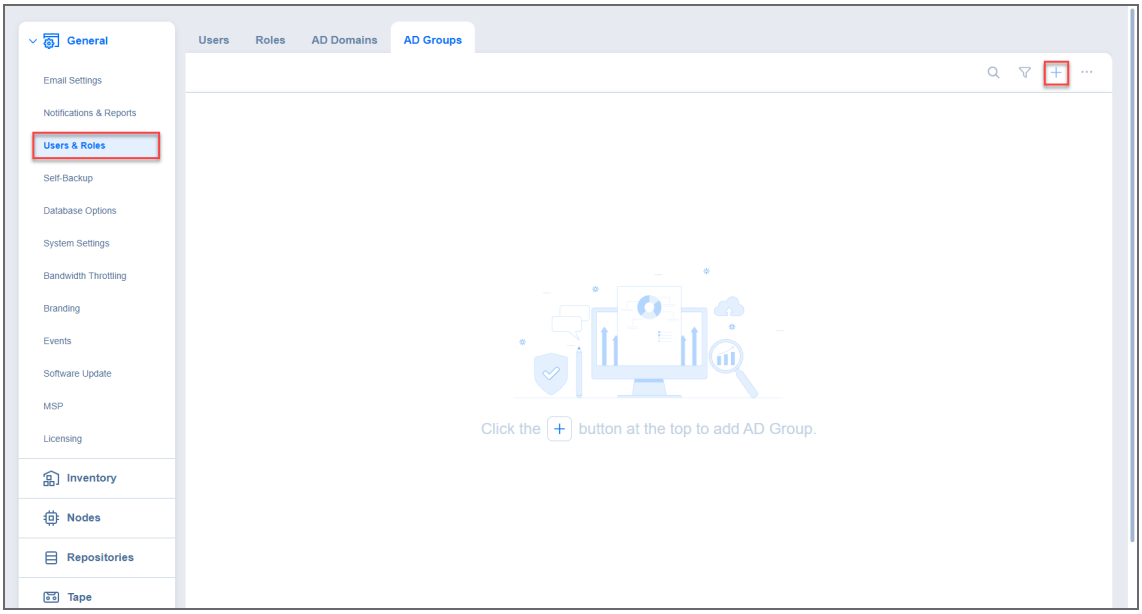
  > **Note**
  >
  > The **Access level** column is displayed only for **Master tenant** in Multi-tenant mode. It displays the access level assigned to the AD group.

- Filter the list of users by entering the name of the user (fully or partially) into the **Search** bar or by selecting the **Filter** option.

**Note**

Clicking **Filter** opens a new window that allows you to filter the list of local users according to **Group name**, **Role**, **Number of users**, and **Status**.

- Add a new AD group by clicking the **+** icon.




**Notes**

- You can add AD groups from multiple domains.
- The AD groups list is disabled if the domain is disabled.

- Edit, delete, disable, enable two-factor authentication, and assign a new role to the local user individually. These actions, except **Edit**, can also be done in bulk by checking the box in the upper left pane to select all users and clicking the **...** (ellipsis) button.

**Note**

When selecting all AD groups to apply a bulk action, NAKIVO Backup & Replication selects only those groups that are displayed on the screen.

- Edit the role assigned to the user by clicking on the name of the role in the respective column.

For details, refer to the following sections:

- “Managing Active Directory Users” on page 475
- “Managing Local Users” on page 483
- “Managing User Roles” on page 492
- “Configuring Two-Factor Authentication” on page 502
- “Managing Application Passwords” on page 500

# Managing Active Directory Users

With NAKIVO Backup & Replication, you can configure Active Directory integration at any time. You can also freely add, edit, disable AD users, or assign a role to them. For details, refer to the topics below:


- "Adding Active Directory User" on page 476
- "Assigning Role to Active Directory User" on page 479
- "Configuring Active Directory Integration" on page 480
- "Disabling Active Directory User" on page 482
- "Editing Active Directory User" on page 483

## Adding Active Directory User

After configuring AD integration in the **Active Directory Configuration** wizard, you can proceed with adding AD user(s). Alternatively, switch to **AD Groups** tab and then click on the "**+**" icon.

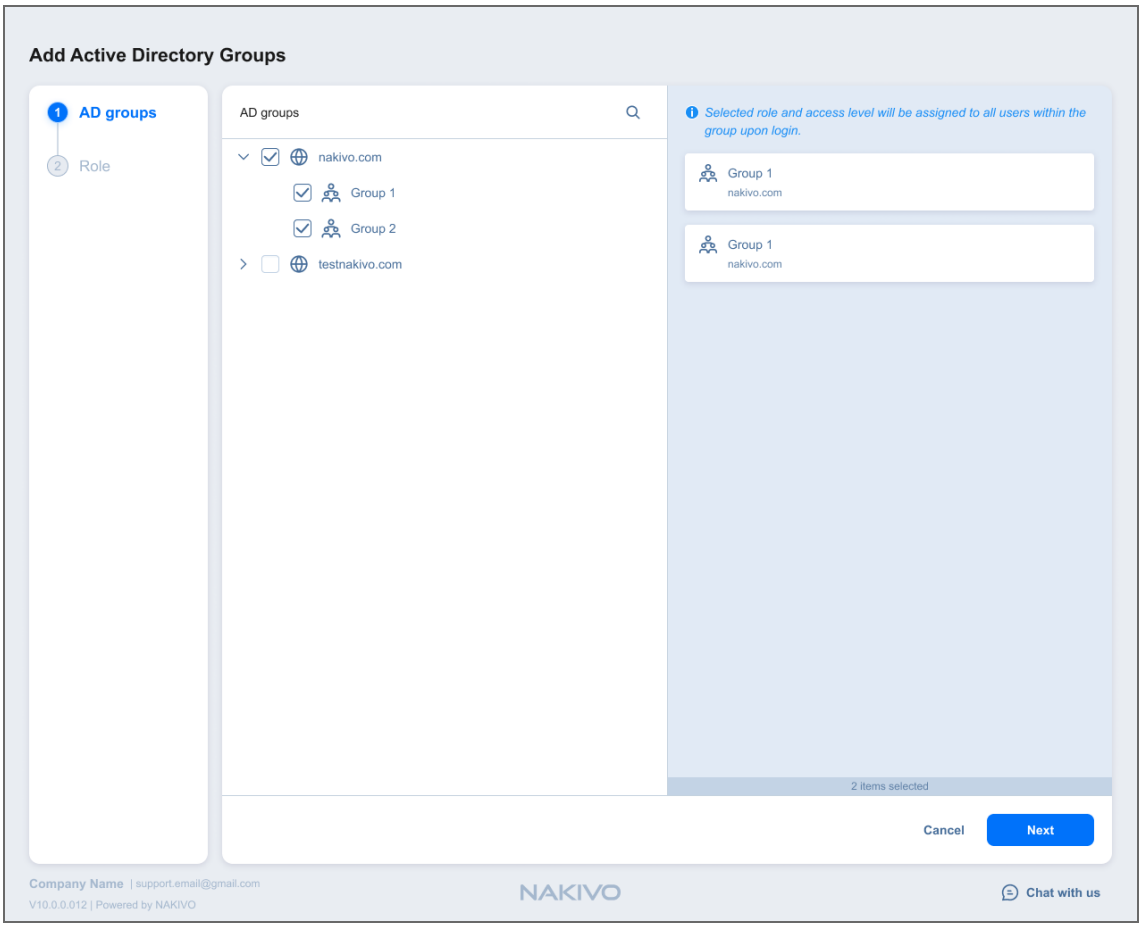
Proceed as follows:

1. Optionally, you can filter the tree of Active Directory users by entering a string to the **Search** box. You can enter a section or the whole name of the item.
2. Select **Active Directory users and groups** by placing a checkmark to their left.
3. The selected items appear in the right pane of the page. If necessary, reorder the selected items by dragging them to a new position. By doing so, you can specify to add the most important users and groups first.

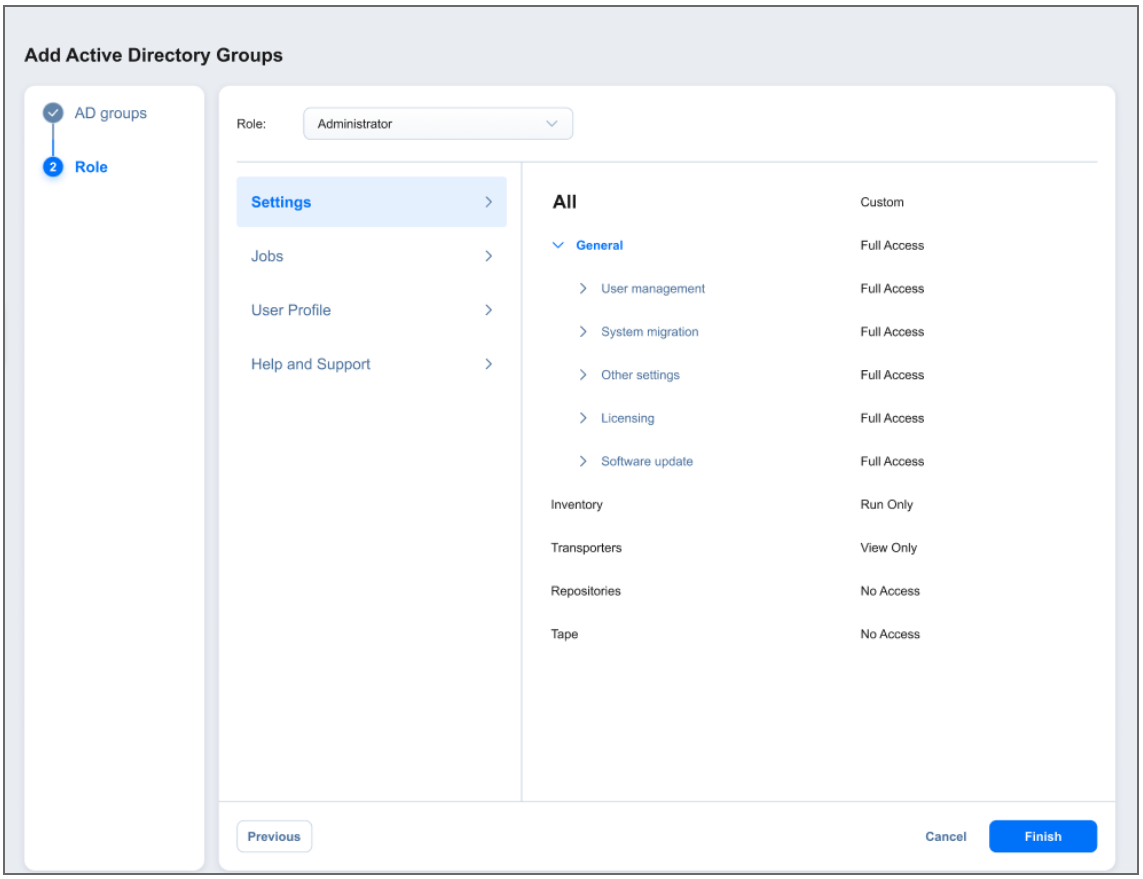
   > **Note**
   > Only logged in users that belong to the group can be added.

4. Review the list of selected items. If necessary, remove a selected user or group from the list in either of the following ways:
   - Deselect the item in the left pane. This will remove the item from the right pane.
   - In the right pane, hover the pointer over the item you wish to remove and click the **Remove**

button. This will deselect the item in the left pane.




5. Click **Next** to proceed to the **Role** tab.
6. On the **Role** tab, choose a user role to be assigned to the users.




7. In the lower right corner of the page, click **Finish**. The selected AD groups with selected access level and role appear in the **AD Groups** tab.

## Assigning Role to Active Directory User

Follow the steps below to assign a role to an Active Directory user:

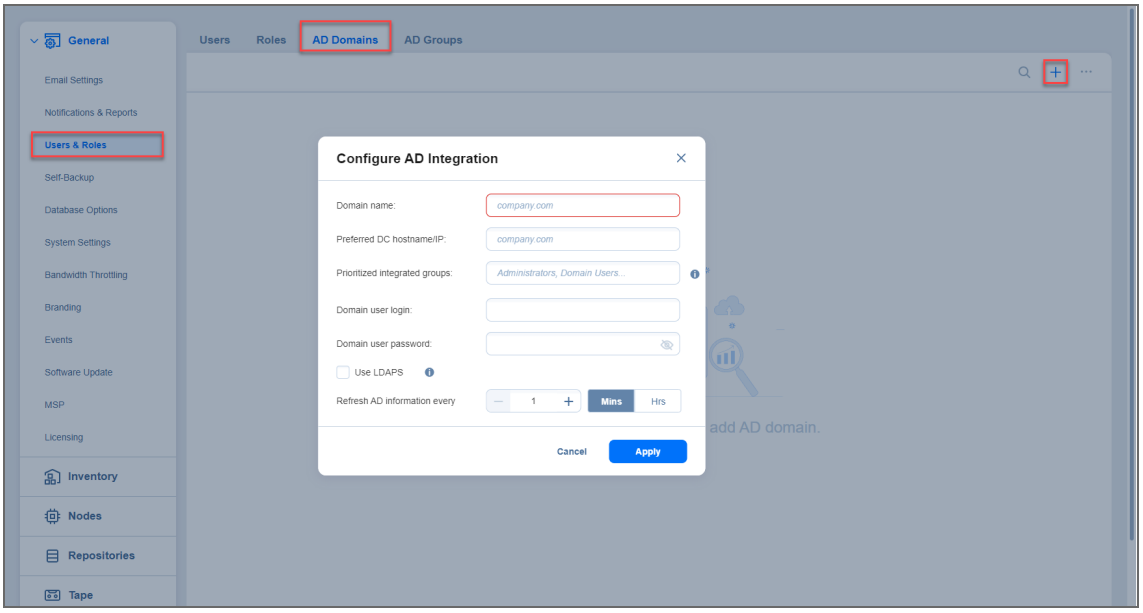
1. Go to **Settings > General > Users & Roles**.
2. The **Users & Roles** page opens in the **Users** view. Hover over the Active Directory user, and then click the **…** (ellipsis) button in the rightmost column of the row.
3. In the menu that opens, click **Assign role**.
4. In the dialog box that opens, select a new user role from the **Role** list and then click **Save**.

The Active Directory user appears in the list of users with the assigned role.

## Configuring Active Directory Integration

To configure Active Directory integration, follow these steps:

1. Go to **Settings > General > Users & Roles**.
2. Select the **AD Domains** tab and then click the "**+**" icon.




3. The **Active Directory Configuration Wizard** opens. Proceed as follows:
   a. In the **Domain name** box, enter the domain name.
   b. In the **Preferred DC hostname/IP** box, enter the name of the preferred domain controller or its IP address.
   c. Optionally, you can enter the name of the preferred Active Directory groups in the **Prioritized integrated groups** box.

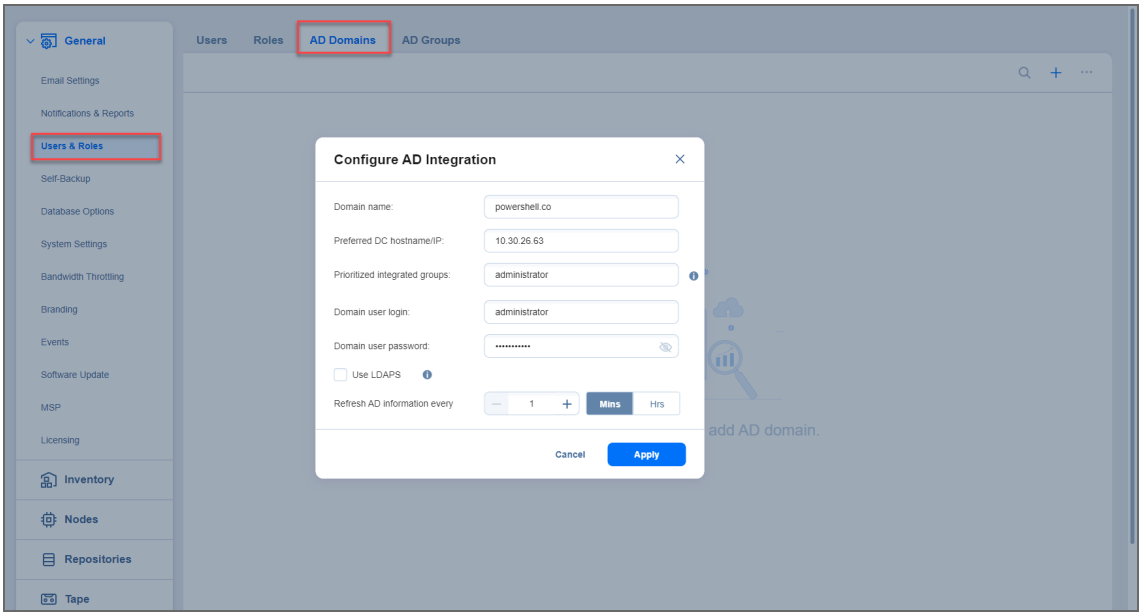
> **Note**
> If a user is a member of two or more Active Directory groups, enter the prioritized group's name in this field.

   d. In the **Domain user login** box, enter the username that will be applied when integrating Active Directory.
   e. In the **Domain user password** box, enter the user password that will be applied when integrating Active Directory.
   f. Optionally, enable **Use LDAPS** option. If checked, port *636* is used for LDAP (Lightweight Directory Access Protocol) over SSL.
   g. **Refresh AD information every**: Specify a periodicity of refreshing Active Directory information.

h. In case Active Directory integration was successfully completed before, you can optionally click **Remove AD Integration** to cancel the AD integration.

> **Note**
>
> The **Remove AD Integration** option is disabled if AD integration is not configured.

i. Click **Apply** after you're done.




j. On the **Users** page of the wizard, proceed with adding an Active Directory user.

**Notes**

- You can add up to 10 AD domains and up to 5.000 AD groups to your solution. Refer to "Navigating AD Domains View" on page 472 and "Navigating AD Groups View" on page 473 for more details about how to work with AD domains and groups.
- A new domain must be assigned a unique name.
- A new AD domain will be added to the list of Active Directory domains and displayed for the corresponding user on the **Users** tab.
- After a new AD domain is added, new AD groups that belong to this domain are created.
- After the existing AD domain is edited/changed, the AD groups that belong to this domain are edited/changed accordingly, for example:
  - If the AD groups are no longer associated with any AD domain, they are deleted.
  - After the items belonging to the AD groups marked as deleted are removed from the product (for example, by another user), these AD groups are deleted.

When the wizard closes, the **Users & Roles** page opens, displaying the newly-added Active Directory users in the list of users.

## Disabling Active Directory User

Follow the steps below to disable an Active Directory user:

1. Go to **Settings > General > Users & Roles**.
2. The **Users & Roles** page opens in the **Users** view. Hover over the Active Directory user you want to disable, and then click the **…** (ellipsis) button in the rightmost column of the row.
3. In the menu that opens, click **Disable**.
4. In the dialog box that opens, click **Disable** to confirm that you want to disable the Active Directory user.

The Active Directory user appears dimmed in the list of users.

> **Note**
> You cannot enable the AD user if its parent AD group is disabled.

## Editing Active Directory User

Follow the steps below to edit an Active Directory user:

1. Go to **Settings > General > Users & Roles**.

   The **Users & Roles** page opens in the **Users** view. In the list of users, do either of the following:

   a. Locate the Active Directory user and click its name.

   b. Hover over the Active Directory user, click the **...** (ellipsis) button in the rightmost column of the row.

   c. Click **Edit**.

2. The **Edit Active Directory User** page opens. Edit the Active Directory user properties if necessary:

   a. In the **General** tab, edit the user.

   b. In the **Role** tab, edit the user role.

   c. Click **Save** to save your modifications to the Active Directory user.

# Managing Local Users

With NAKIVO Backup & Replication, you can freely add, edit, disable, delete local users, or assign a role to them. For details, refer to the topics below:


- "Adding Local User" on page 484
- "Assigning Role to Local User" on page 487
- "Deleting Local User" on page 488
- "Disabling Local User" on page 489
- "Editing Local User" on page 490

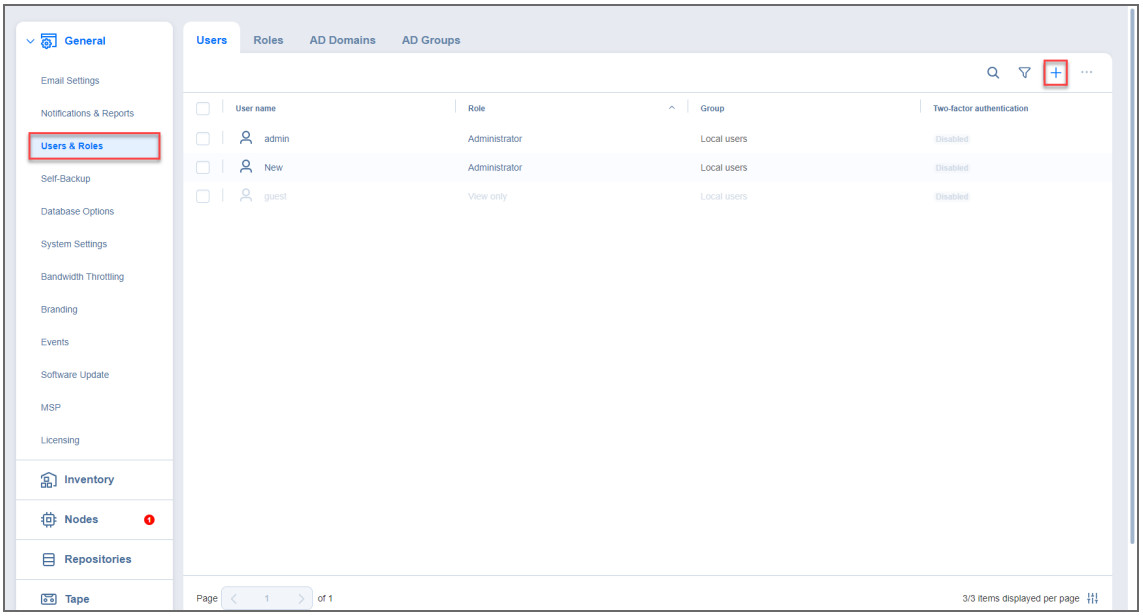

The application has the following built-in local users:

- **admin**: This user has the **Administrator** role assigned. You cannot delete it, disable it, or assign another role.
- **guest**: This user has the **View only** role assigned, with configurable file and object recovery permissions. By default, the account is disabled.

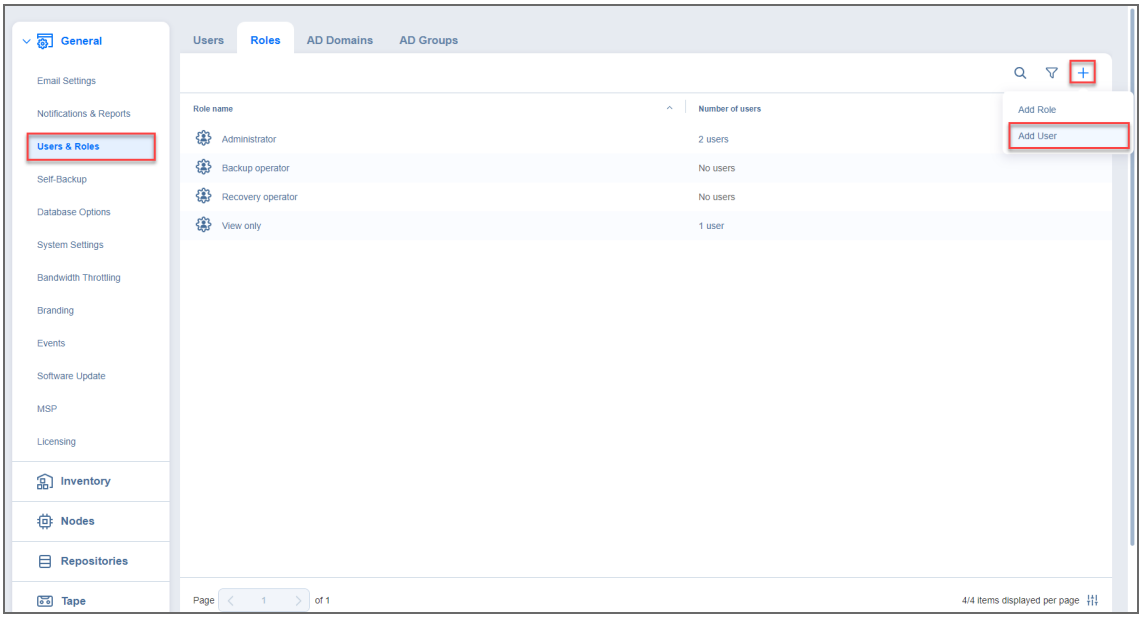
# Adding Local User

Follow the steps below to add a local user:

1. Go to **Settings > General > Users & Roles**
2. The **Users & Roles** page opens on the **Users** tab.
3. Click the **+** icon.

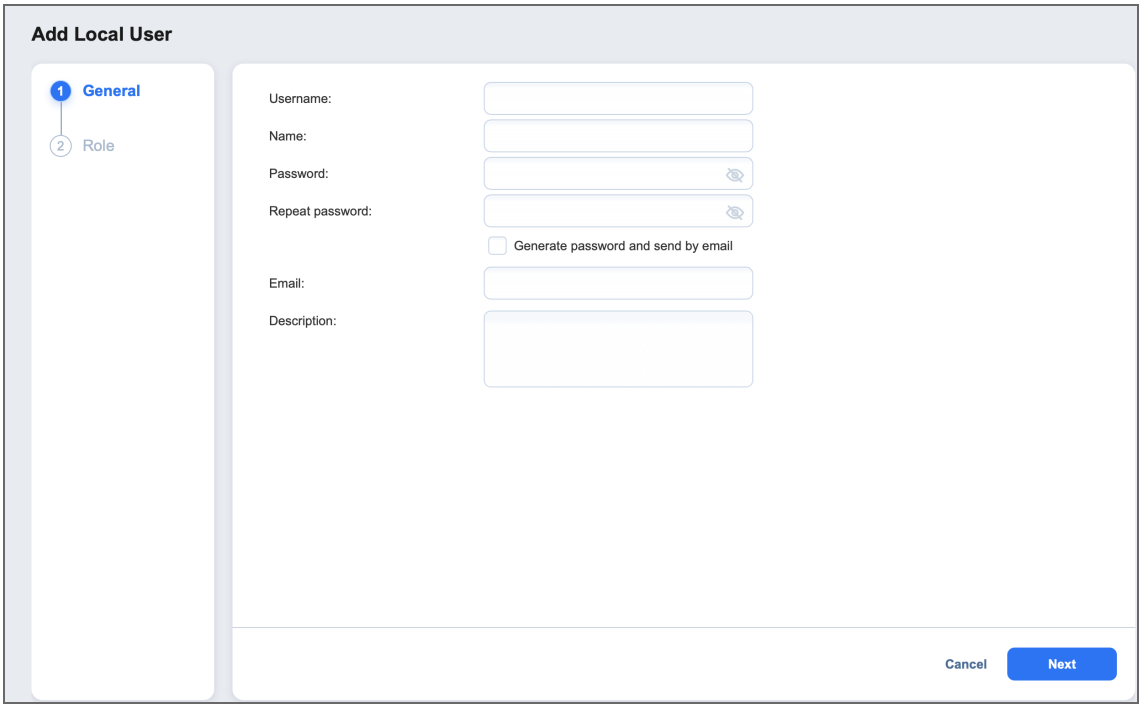



Alternatively, you can switch to the **Roles** tab, click on the "+" icon and select **Add User**.



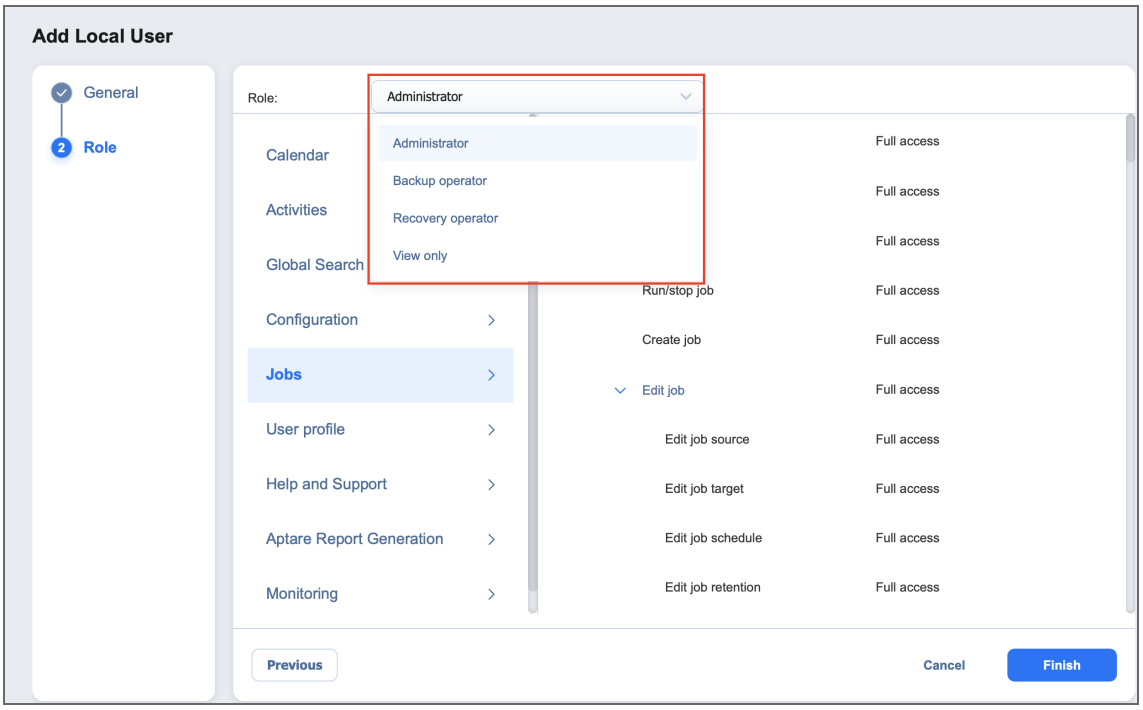

4. The **Add Local User** page opens. Proceed as follows:
   a. In the **Username** box, enter the user name.
   b. In the **Name** box, enter the user's real name.

c. In the **Password** box, enter the user password. To generate a password automatically and send it to the user, select **Generate password and send by email**.
d. In the **Repeat password** box, re-enter the user password.
e. In the **Email** box, enter the user's email address.
f. In the **Description** box, optionally enter a user description.




g. Click **Next** to proceed to the **Role** Tab.
h. In the **Access level** dropdown list, select an access level for the new user (for multi-tenant solutions only).
i. In the **Role** dropdown list, select a user role. Refer to "Managing User Roles" on page 492 for more details about user roles.
j. In the lower right corner of the page, click **Finish**. The local user will appear in the list of users.

5.


Add Local User
General
2 Role
Role:
Administrator
Administrator
Backup operator
Recovery operator
View only
Calendar
Activities
Global Search
Configuration
Jobs
User profile
Help and Support
Aptare Report Generation
Monitoring
Full access
Full access
Full access
Run/stop job
Full access
Create job
Full access
Edit job
Full access
Edit job source
Full access
Edit job target
Full access
Edit job schedule
Full access
Edit job retention
Full access
Previous
Cancel
Finish

## Assigning Role to Local User

Follow the steps below to assign a role to a local user:

1. Go to **Settings > General > Users & Roles**. The **Users & Roles** page opens in the **Users** view.
2. Hover over the local user, and then click the **…** (ellipsis) button in the rightmost cell of the row.

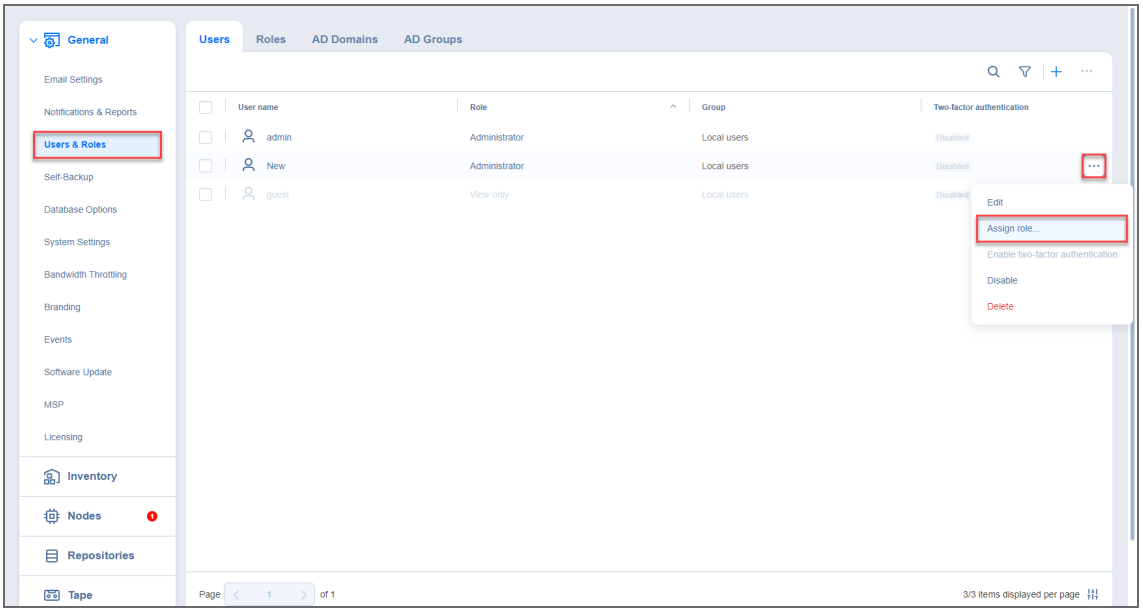
> **Note**
> To assign a role to multiple users, select them using the checkboxes, and then click the ellipsis button for bulk edit.

3. In the resulting menu, click **Assign role**.

> **Important**
> - The **Assign role** button is disabled if:
> - At least one Direct Connect user is selected
> - There are disabled users in the selected users
> - The default administrator is among the selected users

   There are local tenant users and remote tenant users among the selected ones

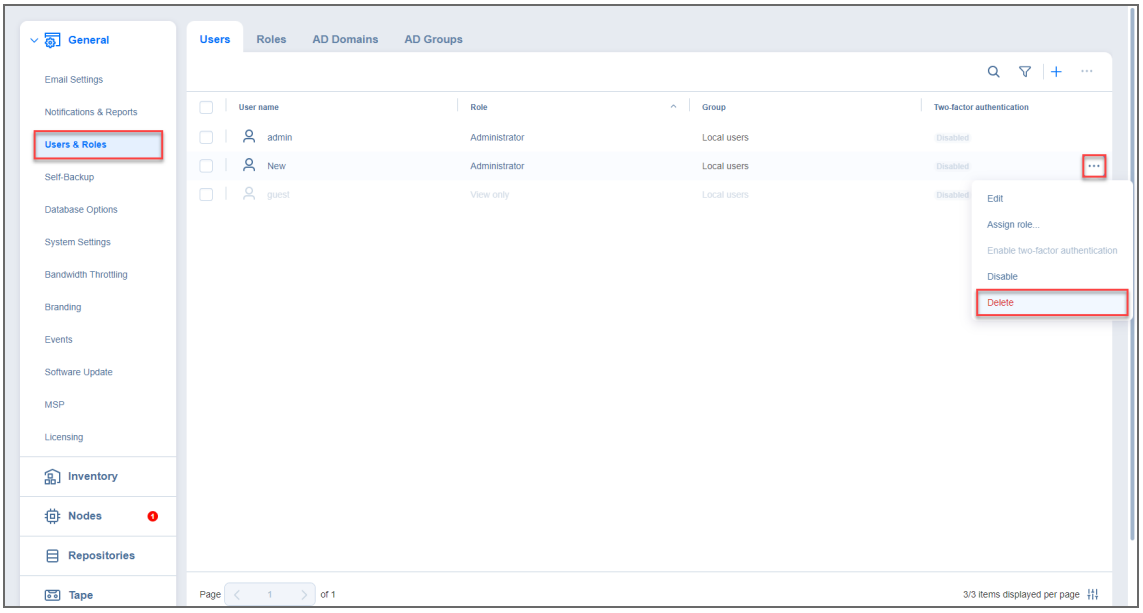
4. In the dialog box that opens, select a new user role from the **Role** drop-down list and then click **Save**.
   The local user will appear in the list of users with the assigned role.



## Deleting Local User

Follow the steps below to delete a local user:

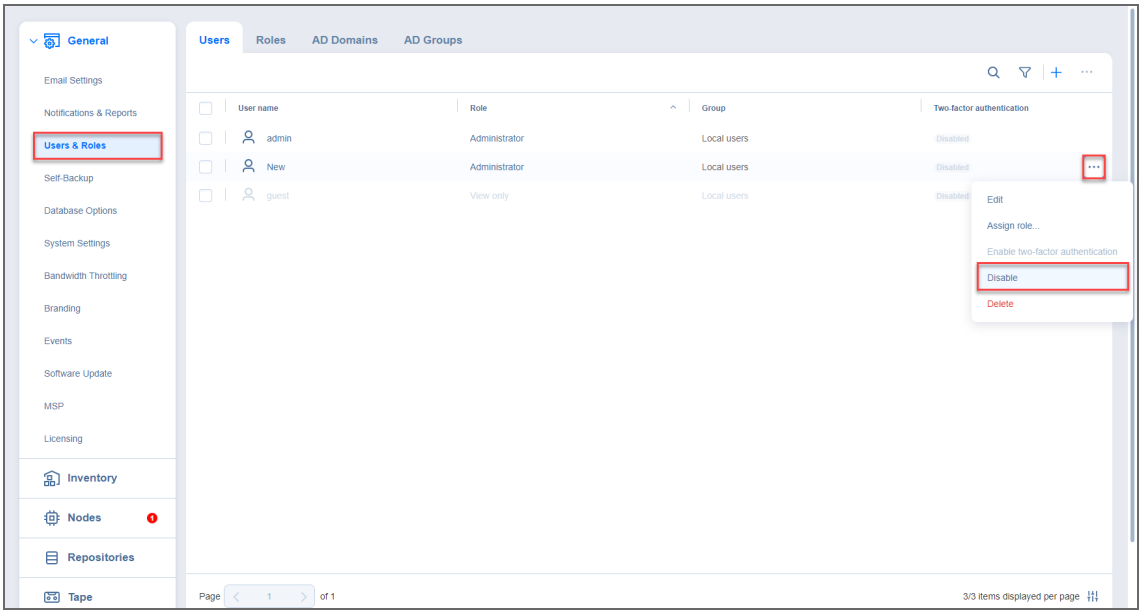
1. Go to **Settings > General > Users & Roles**.
2. The **Users & Roles** page opens in the **Users** view. Hover over the local user you wish to be deleted, and then click the **...** (ellipsis) icon in the rightmost cell of the row.
3. In the resulting menu, click **Delete**.
4. In the dialog box that opens, click **Delete** again to delete the local user. The user will disappear from the list of users.



## Disabling Local User

Follow the steps below to disable a local user:

1. Go to **Settings > General > Users & Roles**.
2. The **Users & Roles** page opens in the **Users** view. Hover over the local user you wish to be disabled, and then click the ... (ellipsis) button in the rightmost cell of the row.
3. In the resulting menu, click **Disable**.
4. In the dialog box that opens, click **Disable** again to disable the local user. The user will appear dimmed in the list of local users.

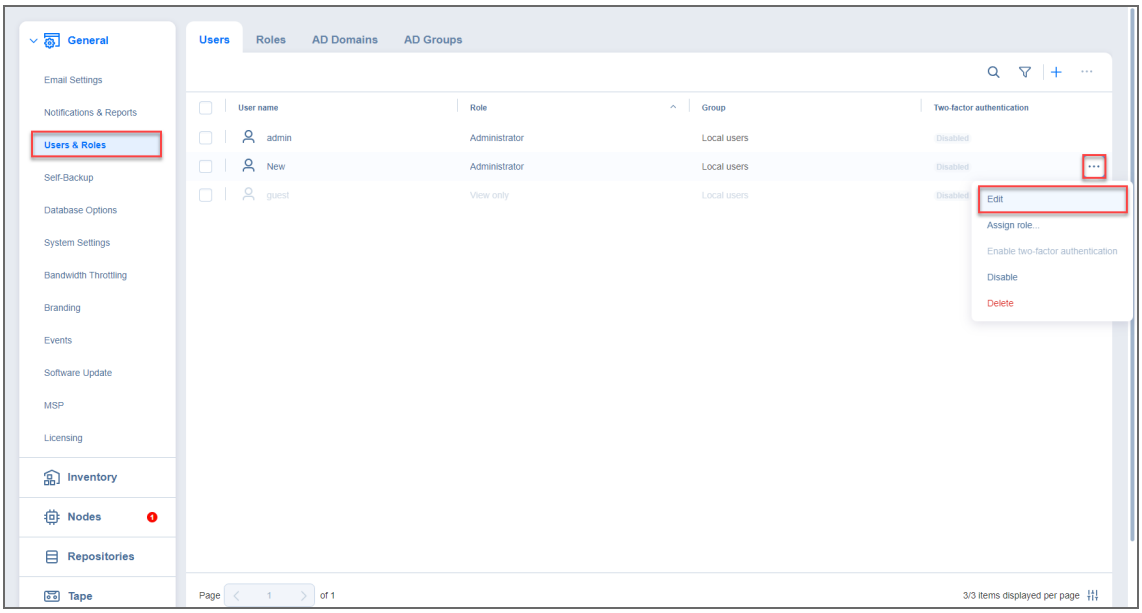



The disabled user cannot log in to the product.

## Editing Local User

Please follow the steps below to edit a local user:

1. Go to **Settings > General > Users & Roles**.
2. The **Users & Roles** page opens in the **Users** view. In the list of users, do either of the following:
    a. Locate the local user that you want to edit and click on the user name.
    b. Hover over the local user and click the **...** (ellipsis) button in the rightmost cell of the row. In the resulting menu, click **Edit**.




3. The **Edit User** page opens. Edit the local user properties if needed:
    a. In the **Username** box, edit the user name.
    b. In the **Name** box, edit the user's real name.
    c. In the **Password** box, edit the user password.
    d. If you edited the user password, re-enter the user password in the **Repeat password** box.
    e. In the **Email** box, edit the user's email address.
    f. Optionally, enable **Two-factor authentication**.

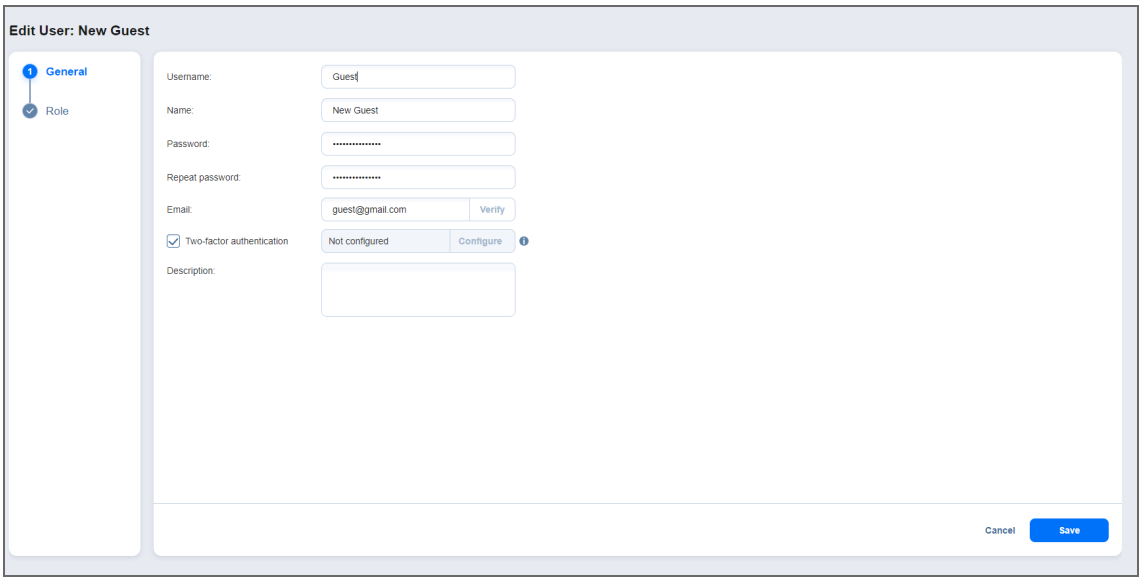
**Note**

This feature is disabled when no email address has been provided for the user.

g. In the **Description** box, edit the user description.

h. In the **Role** tab, edit the user's role.

i. Click **Save** to save your modifications to the local user.



# Managing User Roles

A user role with full access to the **User management** permission is assigned to your user profile to manage user roles. You cannot edit or delete the user role that is assigned to your user profile. The following topics describe how to manage roles of NAKIVO Backup &Replication users in detail:


- “Overview of User Roles” on page 493
- “Adding User Role” on page 495
- “Editing User Role” on page 497
- “Cloning User Role” on page 499
- “Deleting User Role” on page 500

# Overview of User Roles

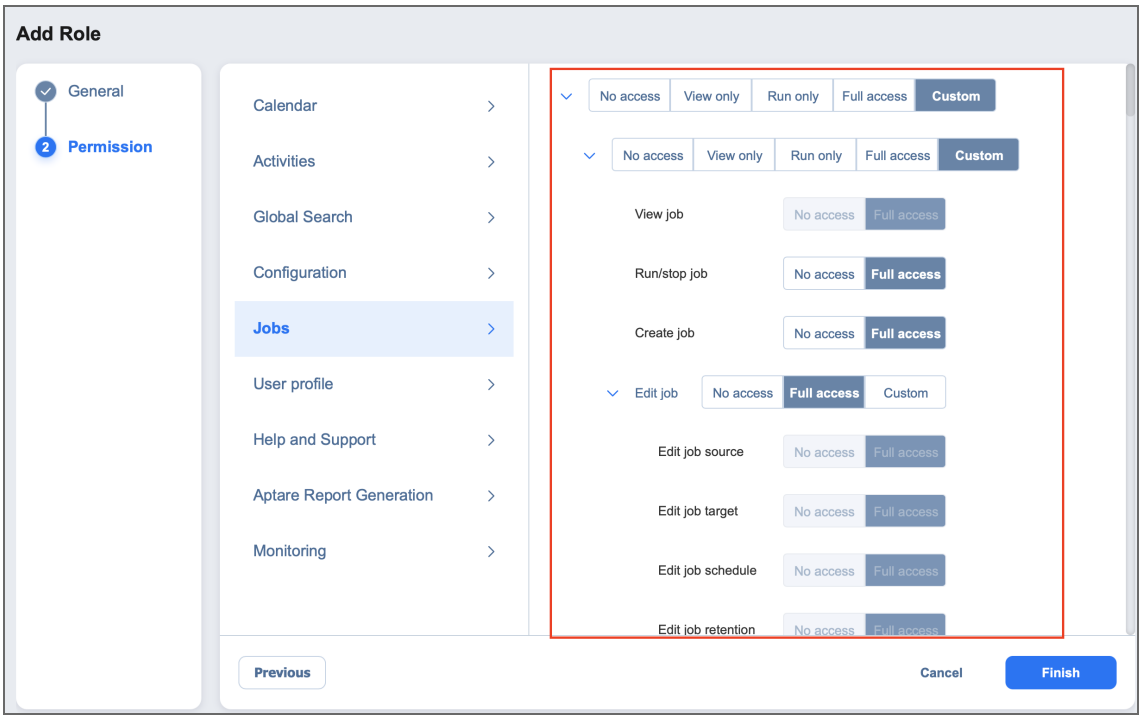
NAKIVO Backup & Replication allows you to assign roles and grant specific permissions to users of the product.

- User Roles
- Access Levels
- Built-in User Roles

User Roles

A user role consists of a set of permissions that can be granted to a NAKIVO Backup & Replication user. Available permissions are grouped by the following product objects:

- **Calendar**: Contains permissions for accessing the **Calendar** dashboard.
- **Activities**: Contains permissions for accessing the **Activities** dashboard.
- **Global Search**: Contains permissions for accessing **Global Search**.
- **Configuration**: Contains a series of permissions for accessing configuration of NAKIVO Backup & Replication.
- **Jobs**: Contains a series of permissions for managing jobs.
- **User profile**: Contains a series of permissions for managing user profile.
- **Help and Support**: Contains a series of permissions for accessing email support, online help center, chat support, and system information.
- **Aptare Report Generation**: Contains permissions for managing Aptare report generation.
- **Monitoring**: Contains permissions for managing the **Monitoring** feature.



Access Levels

There are the following access levels that can be set up for particular permission:

- **No access**: The user cannot view, edit, and run the commands, neither from the graphical interface nor from the command line.
- **View only**: The user can view the commands in the graphical interface but cannot edit or run them; using the command line, the user can only run the commands that do not change NAKIVO Backup & Replication objects.
- **Run only**: The user can only view and run commands, both from the graphical interface and the command line.
- **Full access**: The user can view, edit, and run the commands, both from the graphical interface and the command line.
- **Custom**: A custom set of permissions is configured for a product object.

Built-In User Roles

The product offers you a number of built-in user roles:

- Backup operator
- Recovery operator
- Self-service administrator
- Self-service user
- View only

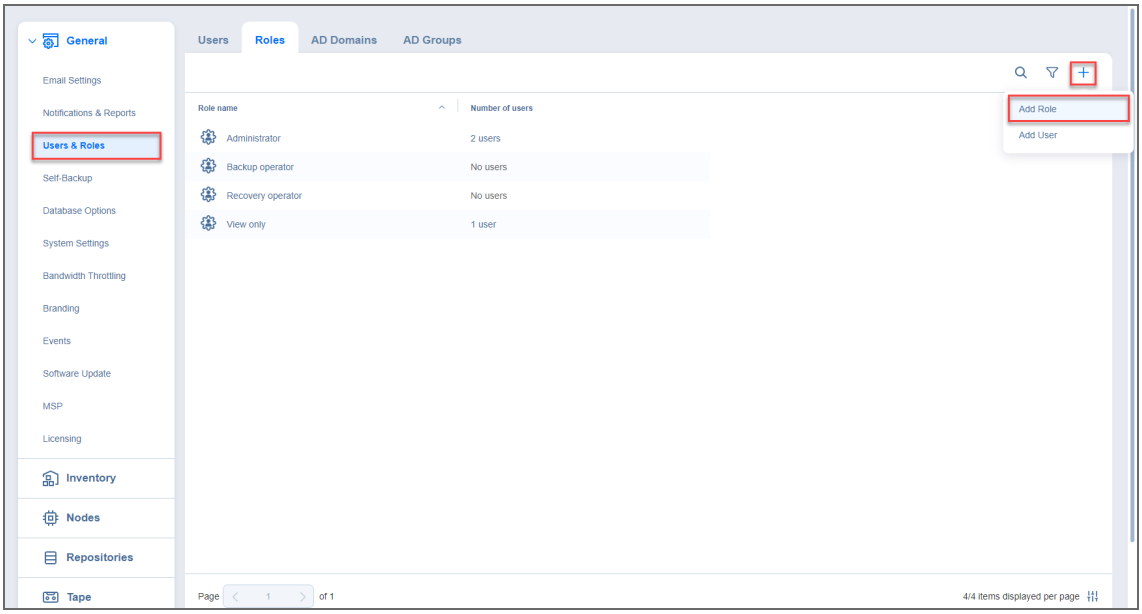
> **Note**
> For editions other than Enterprise Plus, the roles **Backup operator**, **Recovery operator**, and custom roles are not available.

Built-in user roles can be used for performing typical user management tasks. If you need an extra level of security, you can add a new user role or take a built-in user role as a starting point by cloning it.
The user profile can only have a single role assigned.

## Adding User Role

Follow the steps below to add a user role:

1. Go to **Settings > General > Users & Roles**.
2. On the **Users & Roles** page, switch to the **Roles** tab.
3. Click the **+** icon and then select **Add Role**.

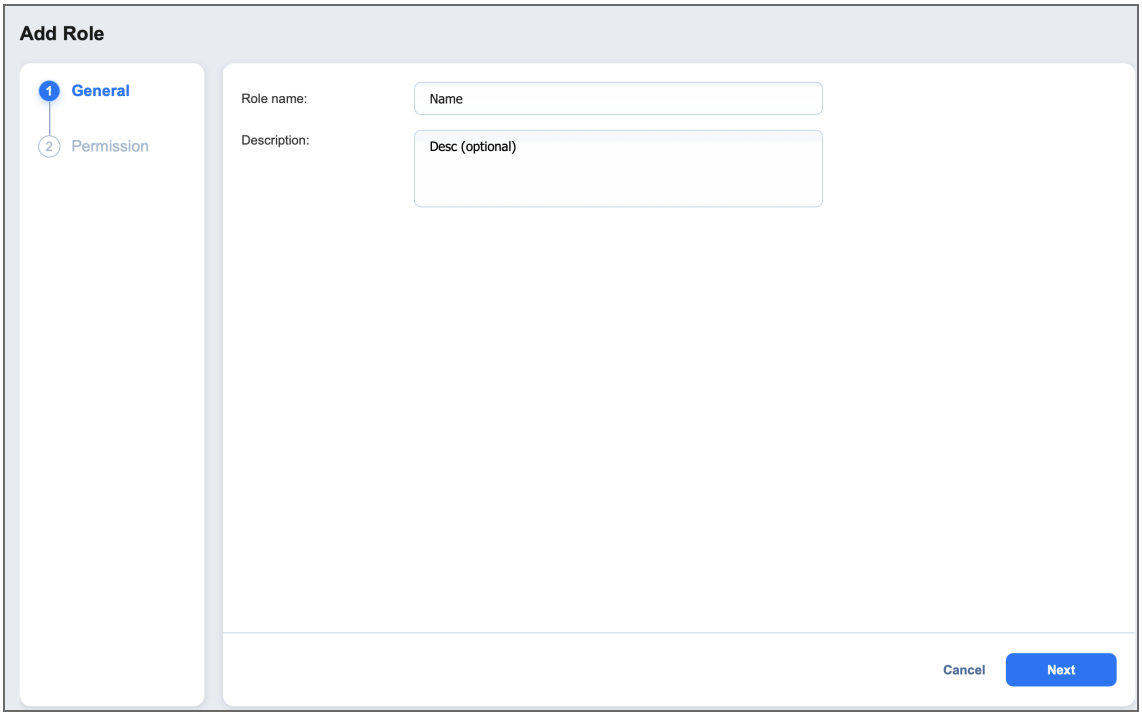



4. The **Add Role** page opens. Proceed as follows:
   a. In the **Role name** box, enter the role name.

   > **Note**
   > In the Multi-tenant mode, every local tenant must have a unique role name.

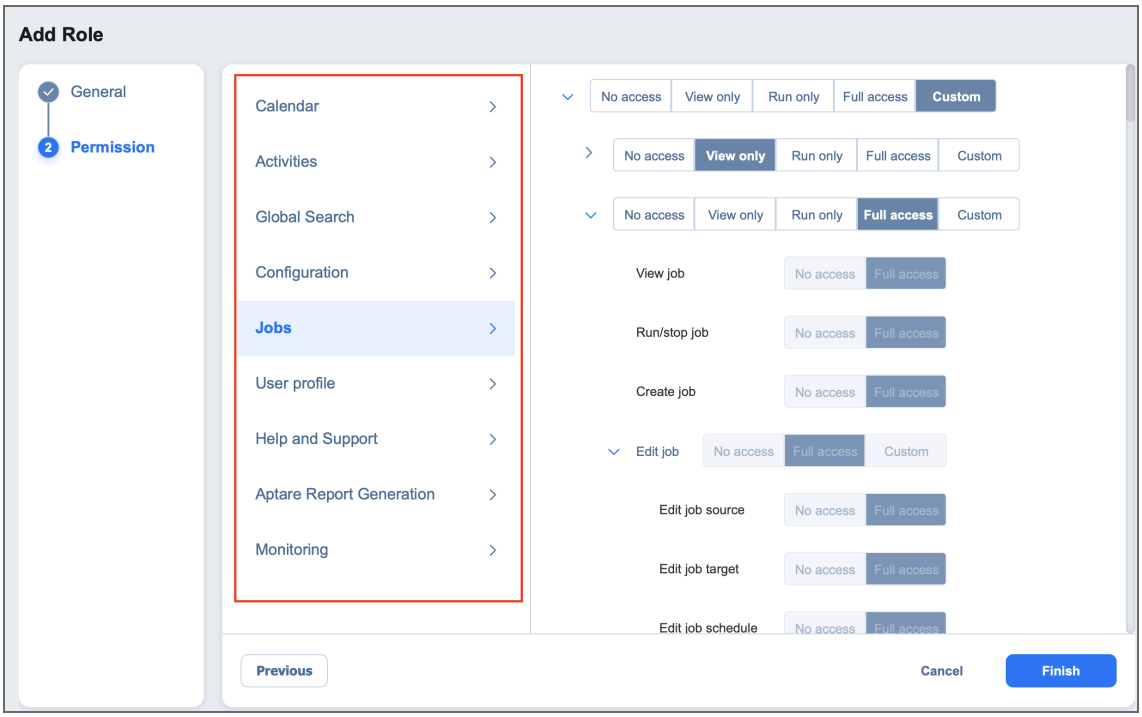
   b. If you are working with a multi-tenant environment, choose either a tenant, master tenant, or all tenants, from the **Access level** list.
   c. In the **Description** box, optionally enter a user description.

   > **Note**
   > For editions other than Enterprise Plus, the roles **Backup operator**, **Recovery operator**, and custom roles are not available.




d. Click **Next** to proceed to the **Permission** tab.

e. A list of permissions opens. Specify necessary permissions for the user role.



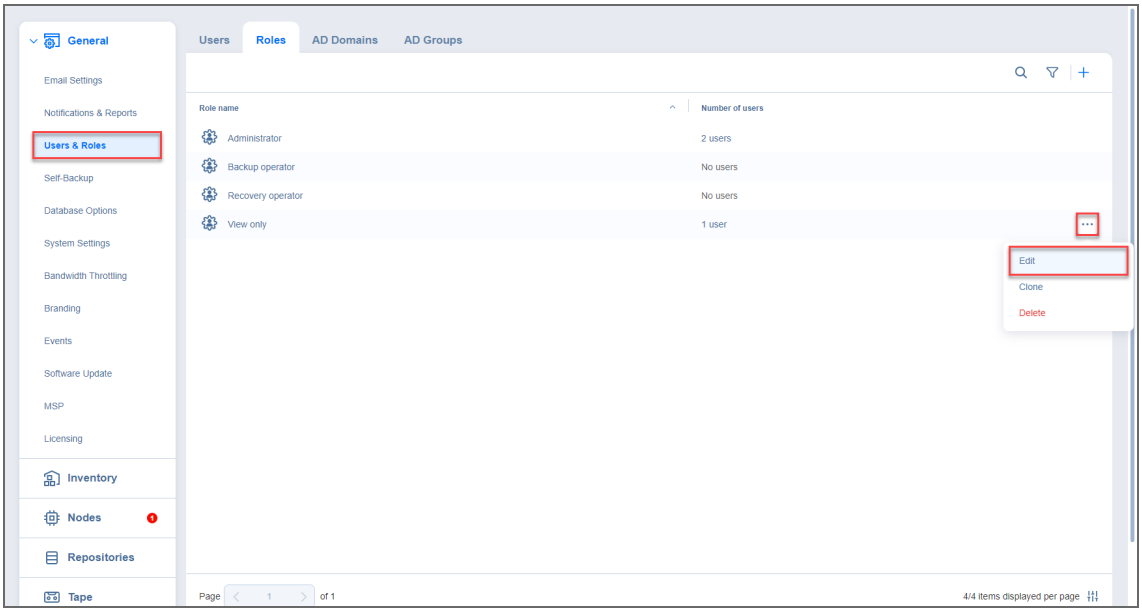

f. Click **Finish** in the lower right corner of the page.

The user role appears in the list of roles.

## Editing User Role

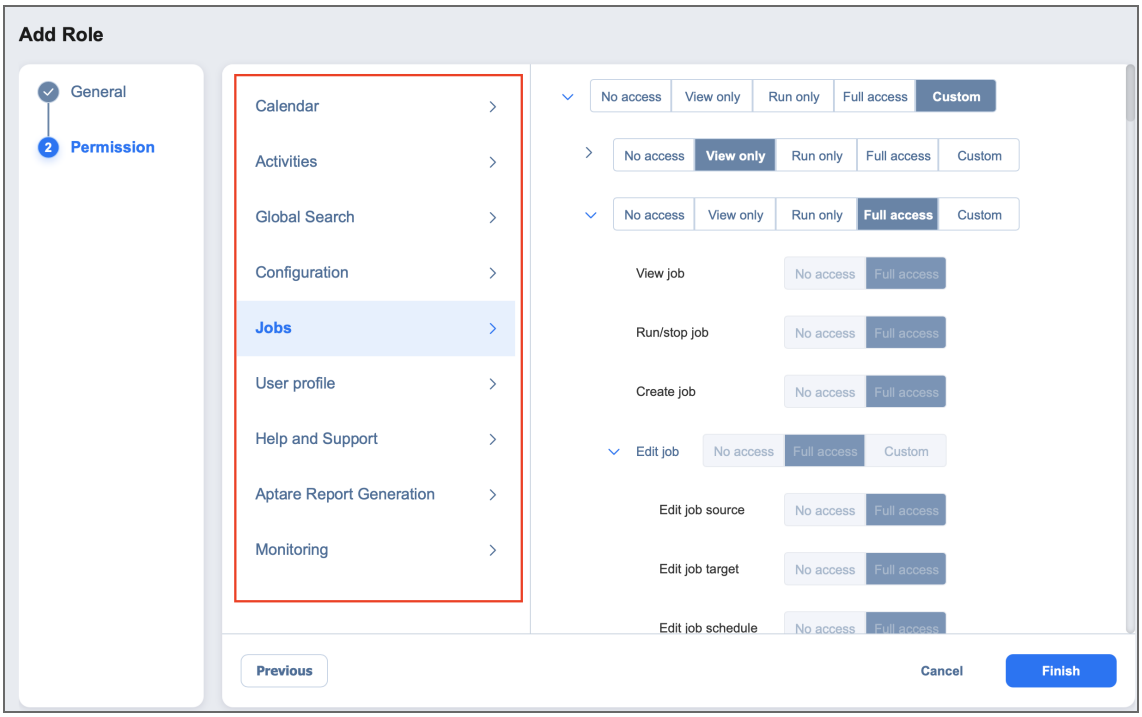
Follow the steps below to edit a user role:

1. Go to **Settings > General > Users & Roles**.
2. On the **Users & Roles** page, switch to the **Roles** tab.
3. In the list of roles, do either of the following:
   a. Locate the user role and click on it.
   b. Hover over the user role, click the **...** (ellipsis) button in the rightmost column of the row, and click **Edit**.




4. The **Edit Role** page opens. Edit the user role properties if needed:
   a. In the **Role name** box, edit the user role name.
   b. If you are working with a multi-tenant environment, you can change the access level for this role by choosing another tenant, master tenant, or all tenants in the **Access level** list.
   c. In the **Description** box, edit the user description.
   d. You can view the **Number of users** with this role, as well as view a full list by clicking the ***x* users** button.
   e. In the **Permissions** tab, you can edit all necessary permissions for the user role.

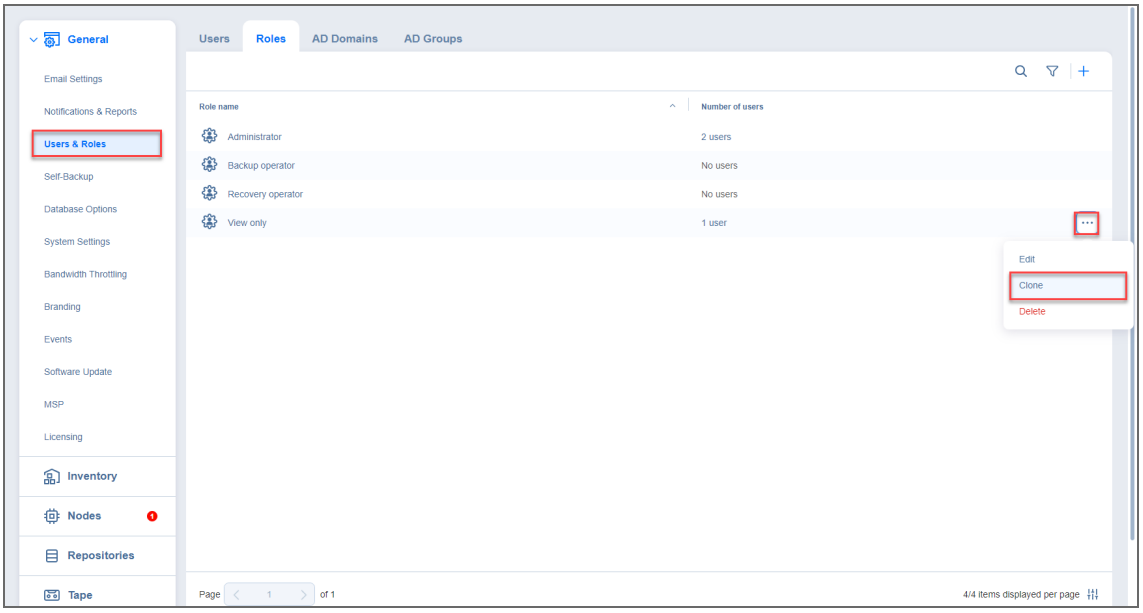
f. When finished, click **Save** in the lower right corner of the page.



# Cloning User Role

Follow the steps below to clone a user role:

1. Go to **Settings > General > Users & Roles**.
2. On the **Users & Roles** page, switch to the **Roles** tab.
3. Hover over the user role, click the **...** (ellipsis) button in the rightmost column of the row and then click **Clone**.
4. A dialog opens asking you to enter the name of the new user role. Enter the name of the new user role and click **Save**.

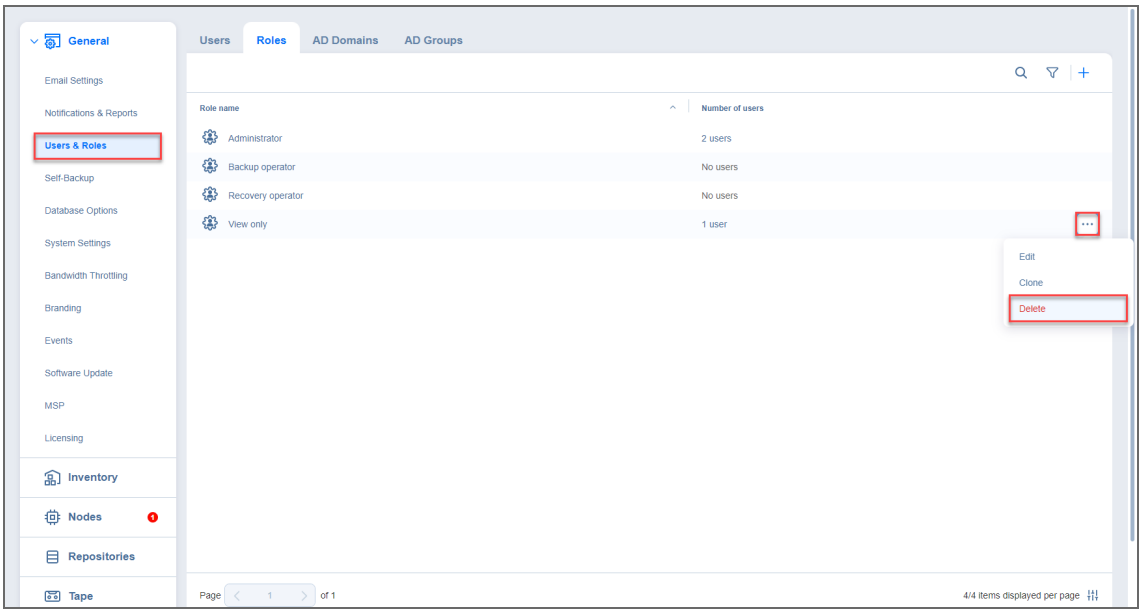



The new user role appears in the list of roles.

## Deleting User Role

Follow the steps below to delete a user role:

1. Go to **Settings > General > Users & Roles**.
2. On the **Users & Roles** page, switch to the **Roles** tab.
3. Hover over the user role, click the **...** (ellipsis) button in the rightmost column of the row and then click **Delete**.
4. In the dialog box that opens, click **Delete** to confirm deletion of the local user.




The user role disappears from the list of roles.

# Managing Application Passwords

An application password is a unique, 16-character credential that can be generated for a user with the Direct Connect role.

## Generate App Password

To generate an application password for the Direct Connect user:

1. Select the Direct Connect user in the list.
2. Click the … (ellipsis) and select **Generate app password**. The application password is generated for the user.

**Notes**

- If a password has already been generated, **Generate app password** is disabled.
- The user can have only one application password, which remains unaffected by changes to their regular login credentials.
- You can download the existing user application password as a text file.

## Download App Password

To download the application password, proceed as follows:

1. Select the Direct Connect user from the list.
2. Click the … (ellipsis) and select **Download app password**. The application password is downloaded as a text file in the browser.

**Note**

You can generate and download application passwords only for users with the **Direct Connect** role.

## Delete App Password

To delete the application password that was assigned to a user, perform the following steps:

1. Select a user with an application password from the list.
2. Click the … (ellipsis) and select **Delete app password**. NAKIVO Backup & Replication deletes the application password for the user.

# Configuring Two-Factor Authentication

NAKIVO Backup & Replication allows you to add an additional layer of security with two-factor authentication (2FA). For details, refer to the topics below:

- Enabling Two-Factor Authentication
- Managing Two-Factor Authentication
- Setting Up Google Authenticator

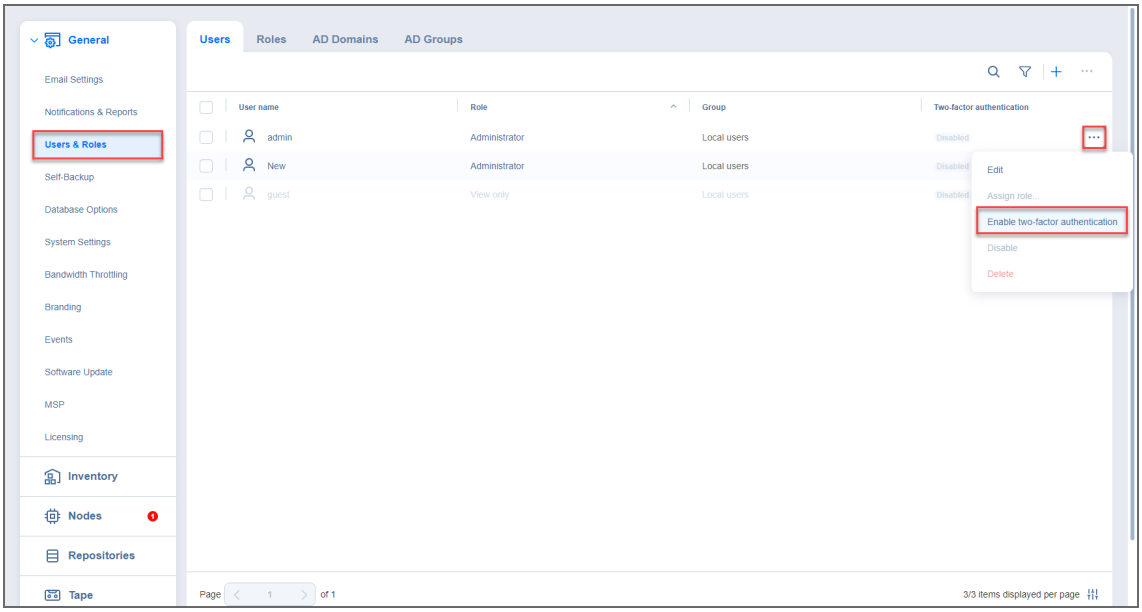
## Enabling Two-Factor Authentication

Two-factor authentication can be enabled in either of the following pages:

- On the **Editing local user** page, select the **Two-factor authentication** checkbox.

> **Notes**
>
> - The two-factor authentication is disabled for:
>   - Direct Connect users
>   - Users without **User management** permission
> - Users without **User management** permission cannot enable Two-factor authentication.
> - Users without **Administrator** role or **Configuration** permission can only configure Two-factor authentication on the login screen of NAKIVO Backup & Replication.
> - It is possible to enable Two-factor authentication only after configuring Email Notifications.

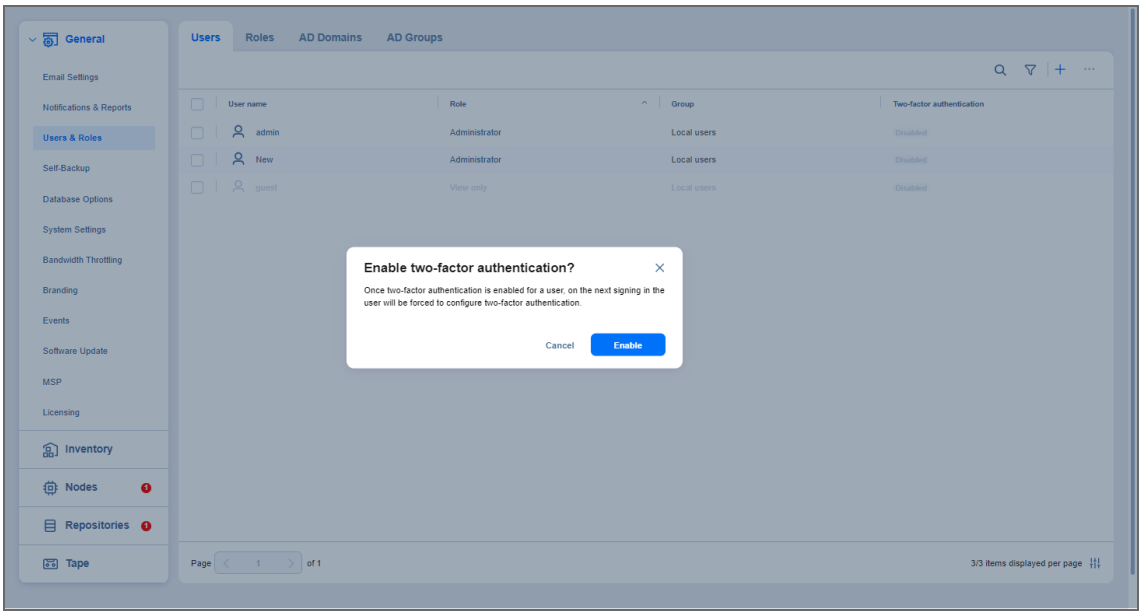
- On the **Users** view, hover over user's name and select **Manage > Enable two-factor authentication**.




> **Note**
>
> - Once two-factor authentication is enabled but not configured for a user, on the next signing in the user will be forced to configure it.

Proceed with configuring two-factor authentication:

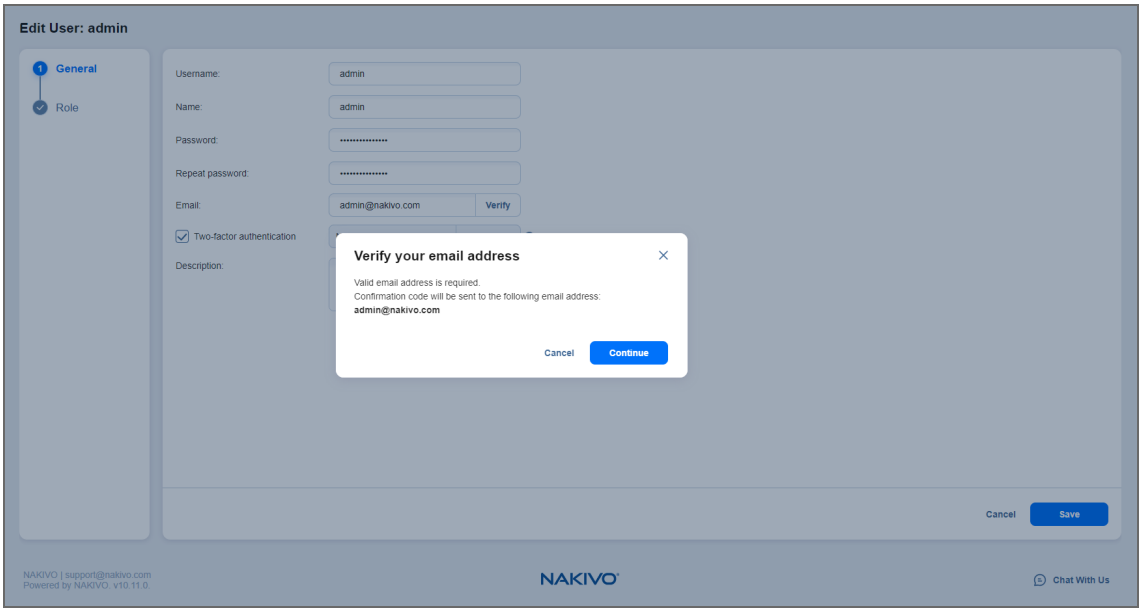
1. Click **Enable** in the dialogue window that appears.



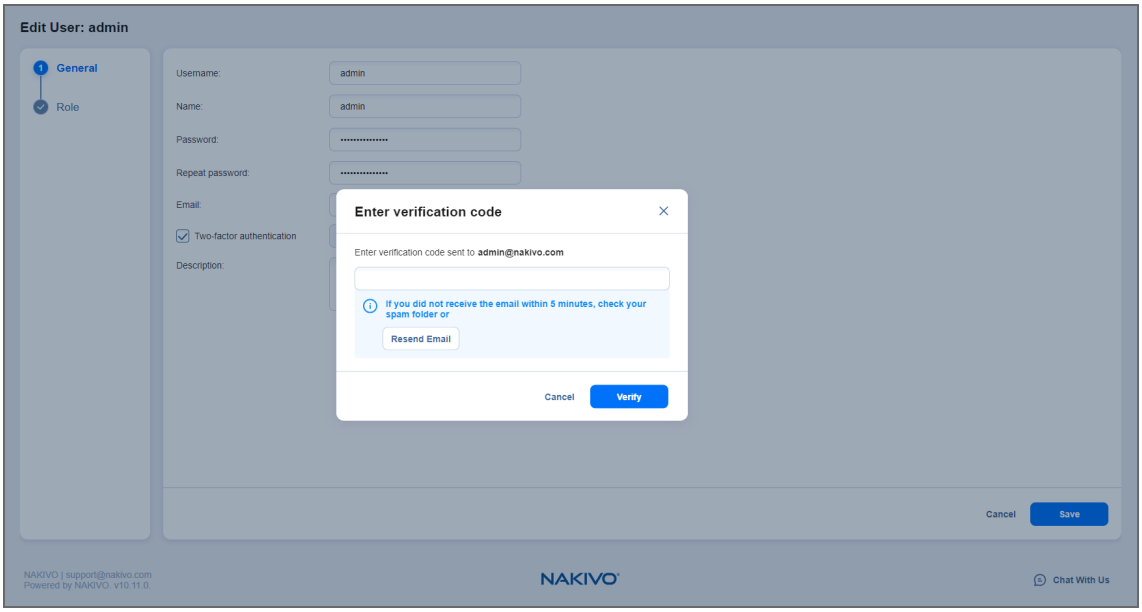

> **Note**
>
> - The two-factor authentication feature is disabled for MSP users drilled down into an ST instance (remote tenant).

2. Click **Continue** in the **Verify your email address** popup that appears.

> **Note**
>
> - Optionally, you can enter the new email address for the user.
> - Also, you can select **Configure** to proceed with 2FA configuration.




3. Enter the verification code that was sent to the specified email address, and click **Verify**.



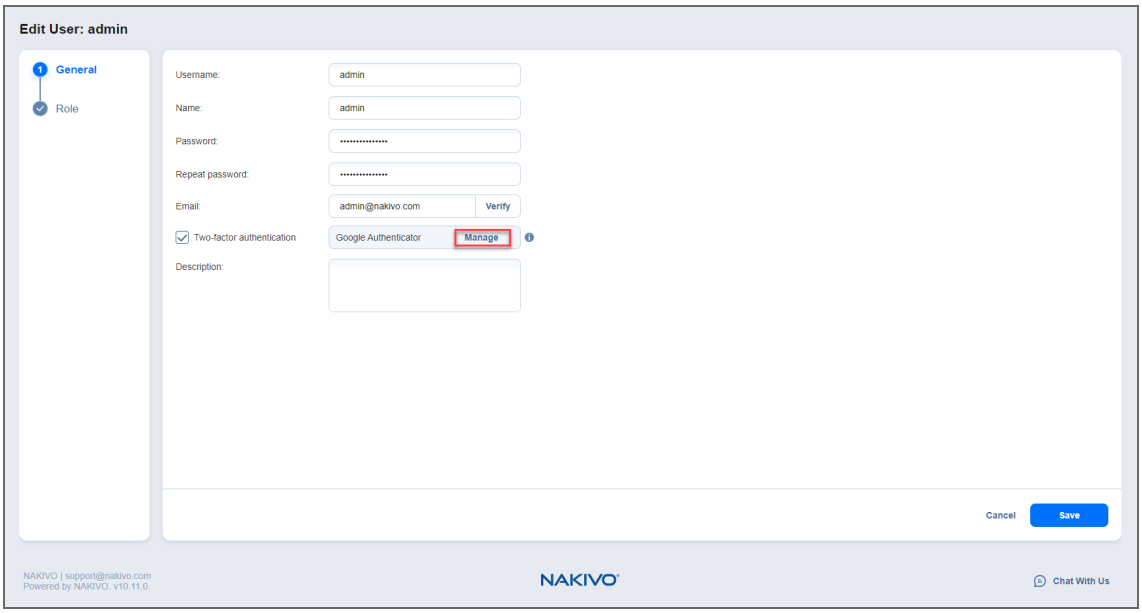

4. Optionally, enter the alternative email address that can be used in case the primary one becomes unavailable, and select **Continue**. Alternatively, select **Skip** to skip this step.
5. If you have entered the alternative email address during the previous step, enter the verification code that was sent to the specified email, and click **Continue** to proceed with Google Authenticator configuration. Alternatively, when configuring 2FA on the **Editing local user** page, select **Cancel** on the **Get Google Authenticator** popup to set up Google Authenticator later.

> **Note**
> When configuring 2FA on the login screen, clicking **Cancel** returns you to the main login screen.

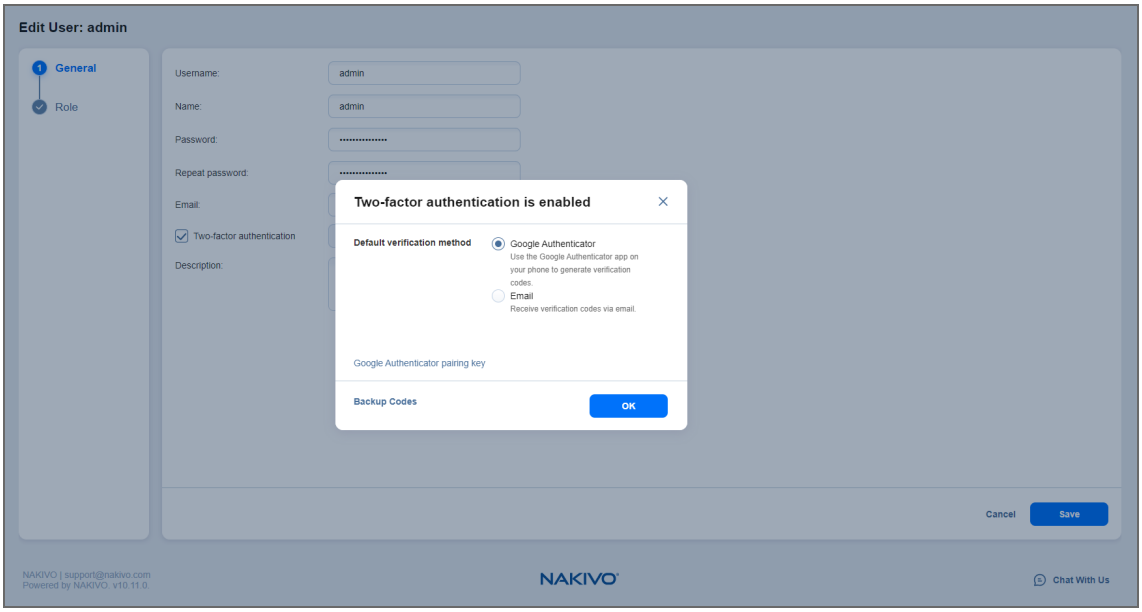
# Managing Two-Factor Authentication

You can manage two-factor authentication in the following way:

1. Click the **Manage** link to the right of **Two-factor authentication** checkbox.




2. Choose one of the following verification methods:
    - **Google Authenticator**: Choose this option to use the Google Authenticator app to generate verification codes. Optionally, click on the **Google Authenticator pairing key** link to see your pairing key or on the **Backup codes** link to view your backup codes.
    - **Email**: Choose this option to receive verification codes via email. Optionally, you can view and change your primary email and add an alternative email by clicking the **add** link. Here you can also view your backup codes by clicking the **Backup codes** link.

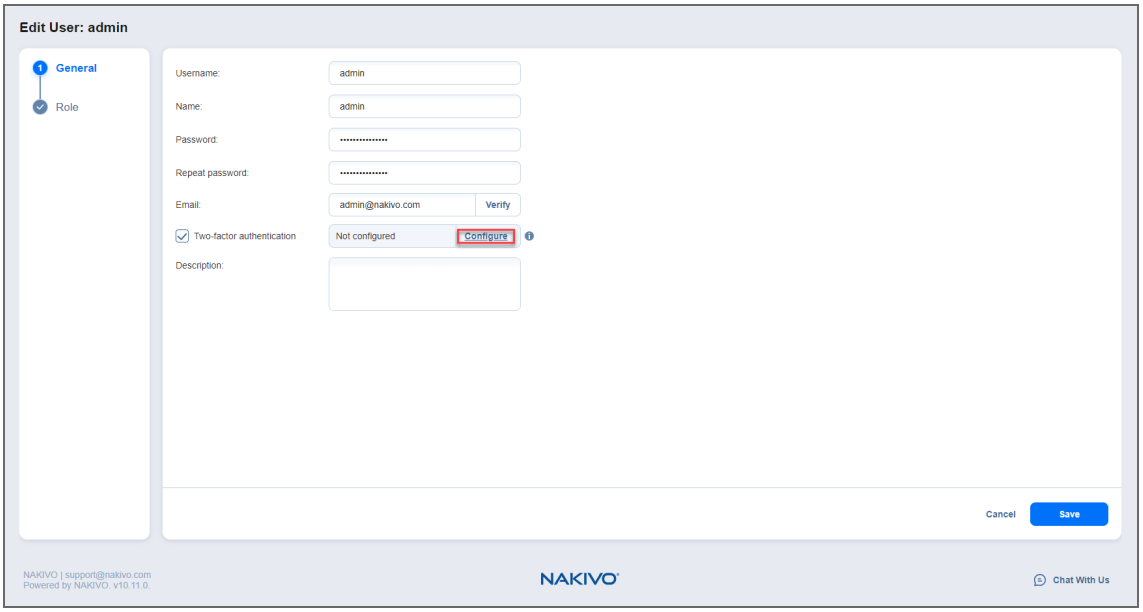



3. Click **OK** when you're done.

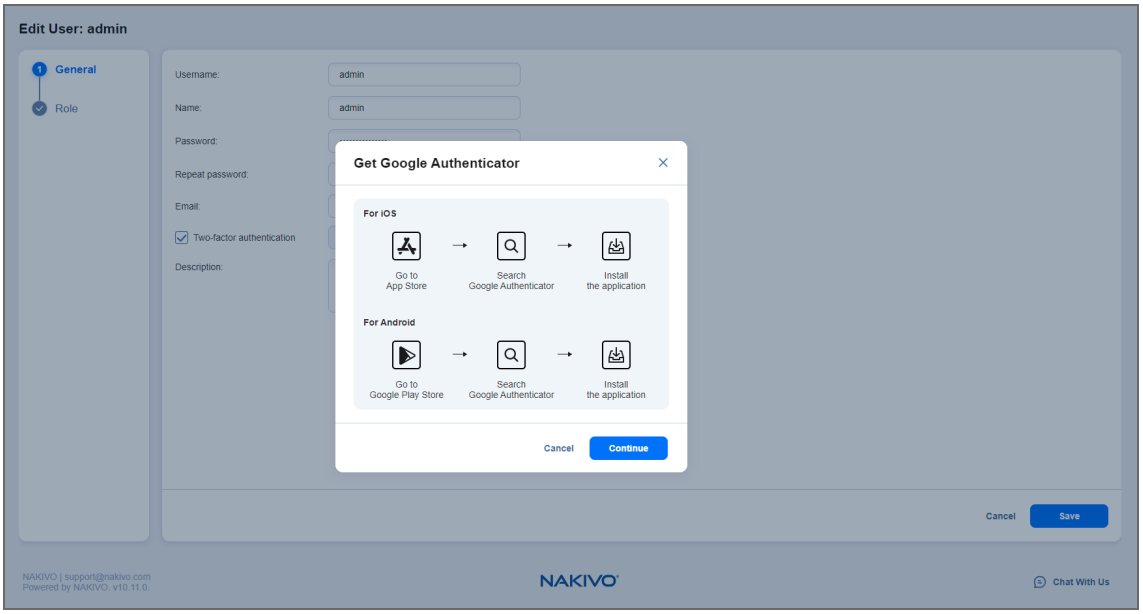
# Setting Up Google Authenticator

NAKIVO Backup & Replication uses Google Authenticator for two-factor authentication. To set up Google Authenticator, do the following:

1. Optionally, if you selected **Cancel** on the **Get Google Authenticator** popup, click the **Configure** link to the right of the **Two-factor authentication** checkbox if you are configuring.




2. Follow the instructions in the popup window to download and install Google Authenticator.



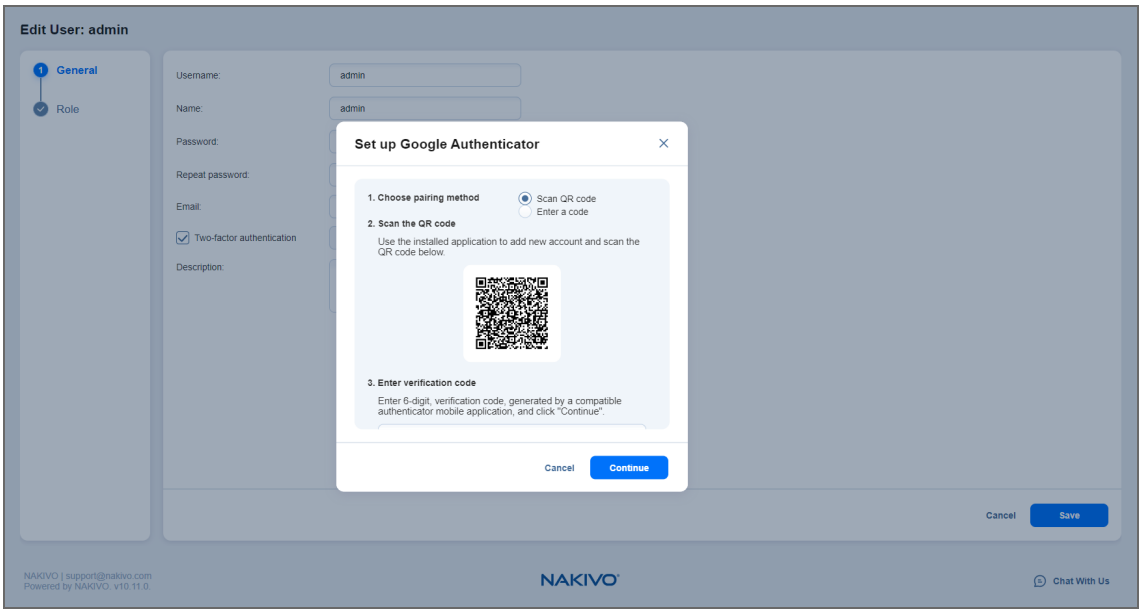

3. Add your NAKIVO Backup & Replication user account to Google Authenticator. Use one of the following methods:

- Select **Scan QR Code** option, and scan the QR code in the popup window.
- Select **Enter a Code** option, and follow the instructions in the popup window to enter the shown code into the Google Authenticator app.
- A popup window appears containing the pairing key, which can be used for adding multiple devices to your account.

> **Important**
> It is highly recommended that you save the pairing key or write it down.

You have the following options:

- Optionally, click on the **copy the Key** link to copy your key and save it for future use.
- Optionally, click on the **download pairing information** link to download and save instructions on how to use the pairing key.
- Click **Continue** when you're done.



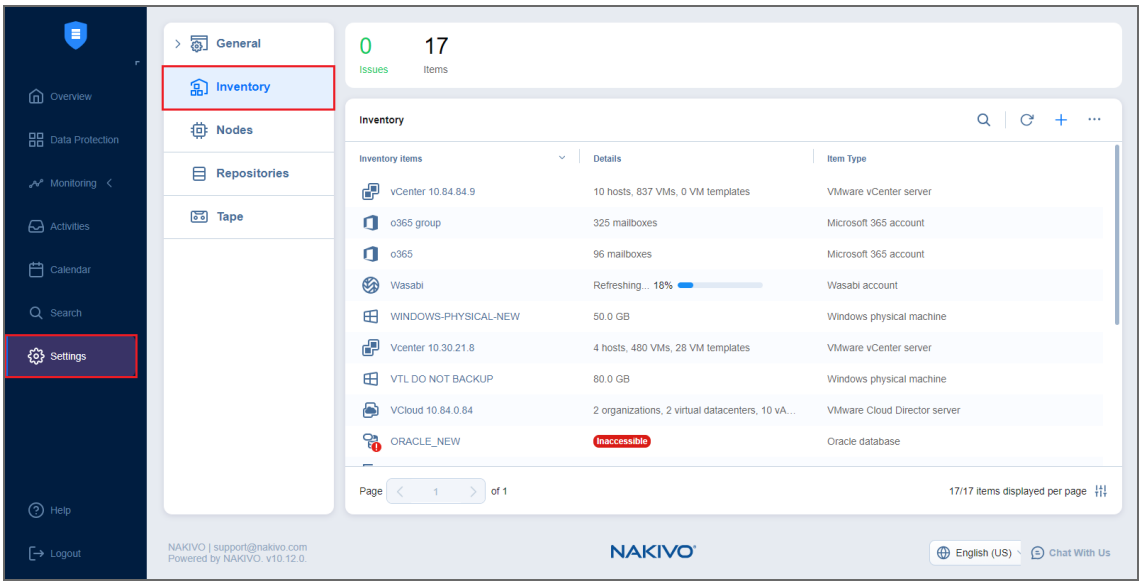

4. The **Backup codes** popup window with four backup codes appears. These one-time codes can be used to log in when you are unable to provide a verification code. Click on the **download as PDF** link to download and save these codes in PDF format or write them down. Additionally, you can click the **generate backup codes** link to generate new codes. Click **Continue**.
5. Enter one of the backup codes in the next popup window to confirm that you have saved them, and click **Finish**.

> **Notes**
> - The backup code used in this step remains valid for one more use.
> - The **Manage** link replaces the **Configure** link after this step has been completed.

# Inventory

Prior to creating backup, replication, or recovery jobs, you need to add your virtual/cloud/physical infrastructure, Microsoft 365 account, Oracle database, or supported storage device to the product's **Inventory**. The discovered item is added to the internal product database, which is refreshed every 1 hour by default. The **Inventory** tab contains a **Summary** bar, which offers an overview of all **Inventory** items. The data displayed is as follows:

- **Issues**: Total number of issues/alarms related to **Inventory** items
- **Items**: Total number of items in the **Inventory**




Refer to the following sections to learn more about adding and managing **Inventory** items:

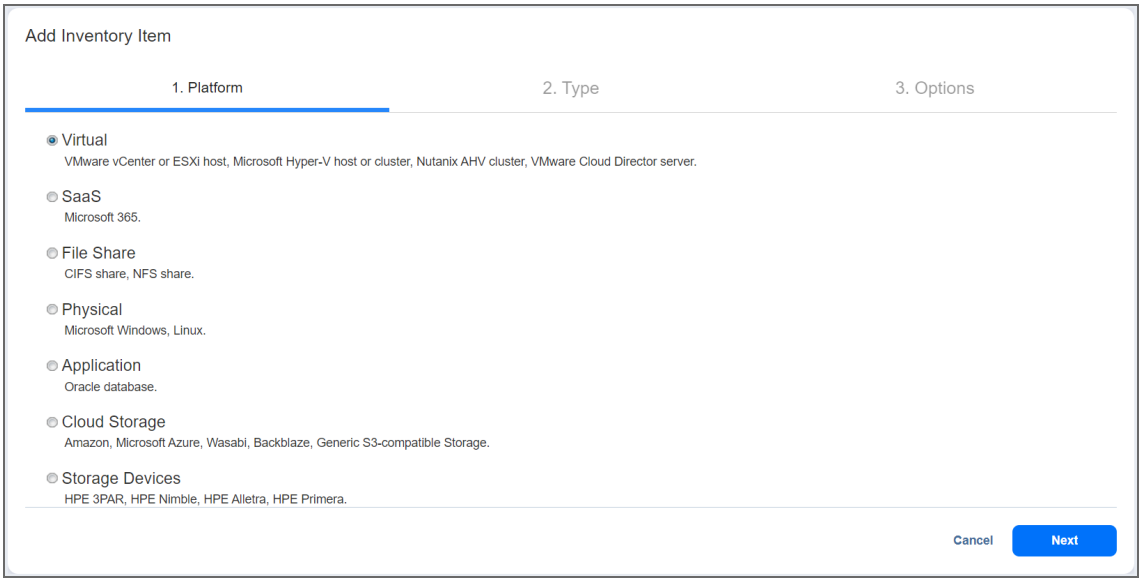

- "Adding VMware vCenters and ESXi hosts" on page 511
- "Adding Amazon EC2 Accounts" on page 513
- "Adding Generic S3-Compatible Object Storage " on page 515
- "Adding Wasabi Accounts" on page 518
- "Adding Backblaze Accounts" on page 521
- "Adding Microsoft Azure Storage Accounts" on page 526
- "Adding HPE 3PAR Storage Devices" on page 542
- "Adding HPE Nimble Storage Devices" on page 543
- "Adding HPE Primera Storage Devices" on page 545
- "Adding HPE Alletra Storage Devices" on page 546
- "Managing Inventory" on page 550
- "Managing Credentials" on page 559

# Adding VMware vCenters and ESXi hosts

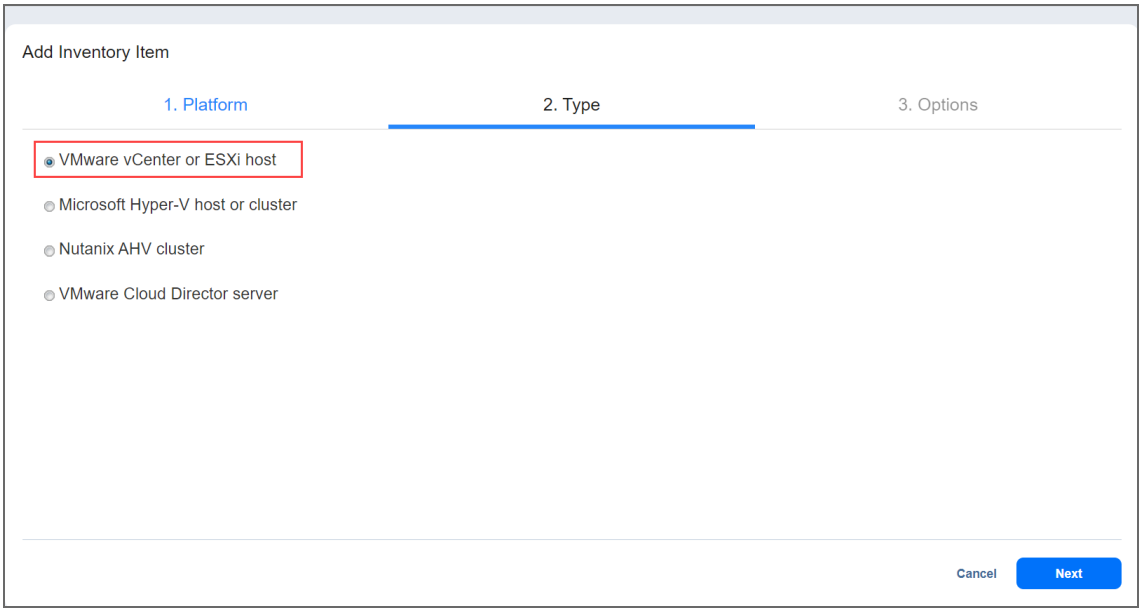
Before adding VMware vCenters and ESXi hosts, ensure your environment meets the platform and network requirements. Check the supported platforms and required TCP ports.

To add VMware vCenter servers and standalone ESXi hosts to the product, follow the steps below:

1. Navigate to **Settings**.
2. Go to the **Inventory** page and click **+**.
3. On the **Platform** page, select **Virtual** and then click **Next** to proceed.




4. On the **Type** page, select **VMware vCenter or ESXi** host and then click **Next** to proceed.




5. The **Options** page opens. Proceed as follows:

a. In the **Hostname or IP** field, specify the hostname or IP address of the vCenter server or standalone ESXi host that you wish to add to the **Inventory**.

> **Notes**
>
> - vCenter-managed ESXi hosts should not be discovered directly by their IP addresses or hostnames. Instead, you should add the vCenter that manages such ESXi hosts.
> - The Transporter installed on NAS with ARM architecture cannot be used for discovery of ESXI host with enabled **Direct Connect** option.

b. In the **Username** and **Password** fields, specify credentials of the vCenter server or standalone ESXi host that you want to add to the inventory.

> **Note**
>
> The credentials you specify should have full administrative privileges to the vCenter server or standalone ESXi host.

c. Specify the Web services port in the appropriate field.

d. Optionally, you can enable the **Use Direct Connect** option for this item to discover a vCenter server or standalone ESXi host located in a remote environment using the "Direct Connect" on page 54 functionality.. To do this:

   i. Select **Use Direct Connect**.

   ii. In the **Assigned transporter** drop-down list, choose the **Transporter** that was installed in the remote environment with **Direct Connect** enabled.

> **Notes**
>
> - This option is available only if there is a **Transporter** with **Direct Connect** enabled on the remote site. For details, refer to"Enabling Direct Connect" on page 596.
> - vSphere tags support is not available via **Direct Connect**.

6. Click **Finish**. After the process is completed successfully, you can exit **Settings** and create jobs with the newly discovered VMs.

7. When NAKIVO Backup & Replication is deployed on an ARM-based NAS, a dialog box may open stating that the current **Transporter** does not support VMware vSphere and asking you to deploy an additional **Transporter**. Click **Got It** to close the dialog box. Refer to "Deploying Transporter as VMware Appliance" on page 575 for more details about deploying additional **Transporters** that support VMware vSphere.

# Adding Amazon EC2 Accounts

Add an Amazon EC2 account to NAKIVO Backup & Replication as described in the sections below.

- "Creating AWS Access Key ID and Secret Access Key" below
- "Adding an Amazon EC2 Account to Inventory" below

## Creating AWS Access Key ID and Secret Access Key

Prior to adding your AWS account to the inventory, you need to create and retrieve an AWS Access Key ID and Secret Access Key. The Access Key ID and Secret Access Key are used by NAKIVO Backup & Replication to sign the programmatic requests sent to AWS, such as retrieving the list of instances, creating snapshots, and so on.

To create an AWS Access Key ID and a Secret Access Key, follow the steps below:

1. If you don't have an AWS account, create a new one at https://aws.amazon.com.
2. Open the IAM console.
3. In the left pane, click **Users**.
4. Click your IAM username (not the checkbox).
5. Go to the **Security Credentials** tab and then click **Create Access Key**.
6. Click **Download Credentials** and store the keys in a secure location.

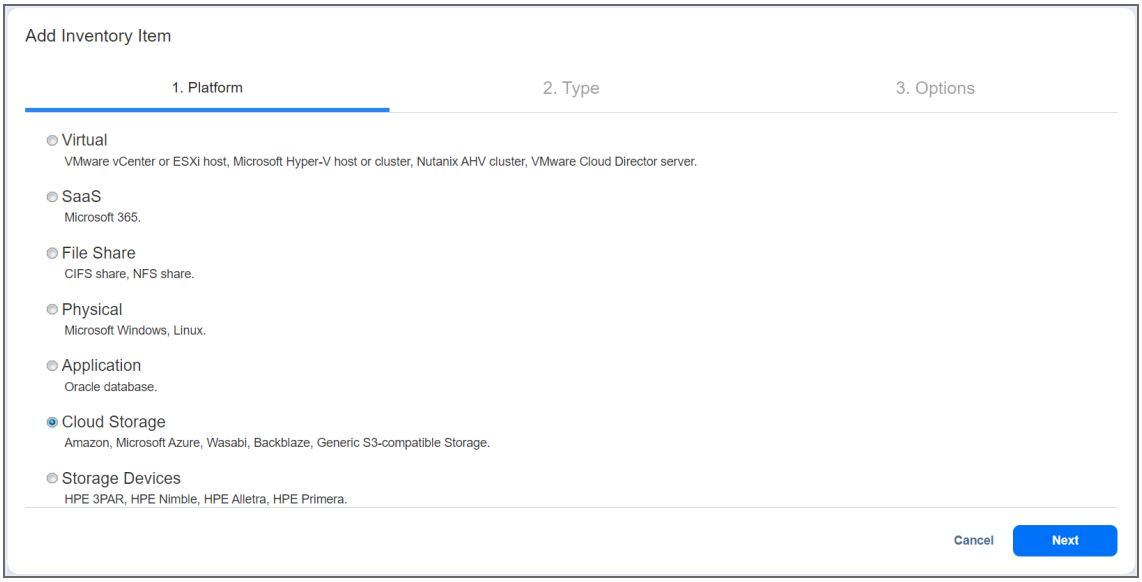
**Important**

Your Secret Access Key will no longer be available in the AWS Management Console; you will have one copy only. Store it in a secure location and do not share it in order to protect your account from unauthorized access.

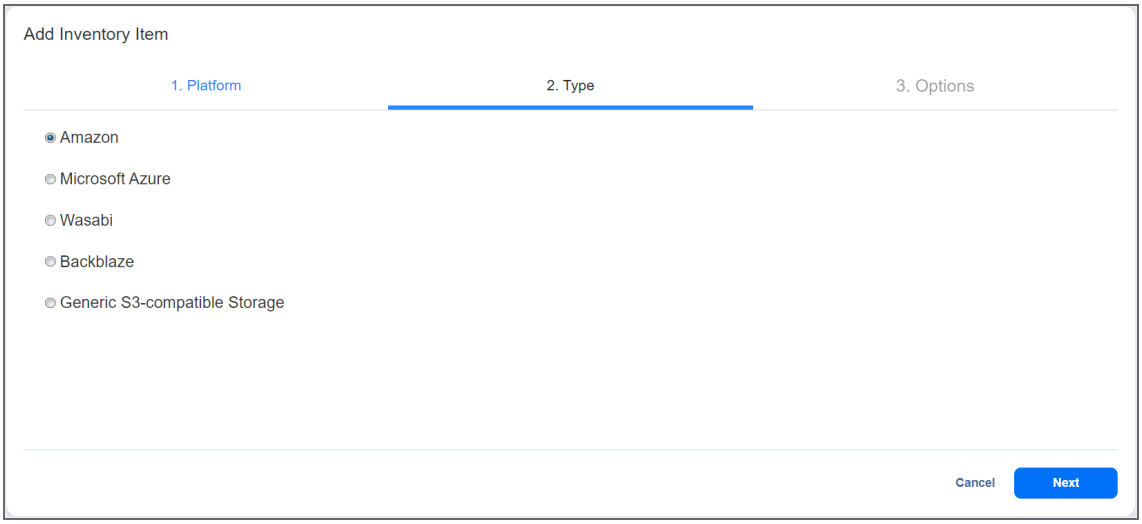
## Adding an Amazon EC2 Account to Inventory

To add an Amazon EC2 account to NAKIVO Backup & Replication, follow the steps below:

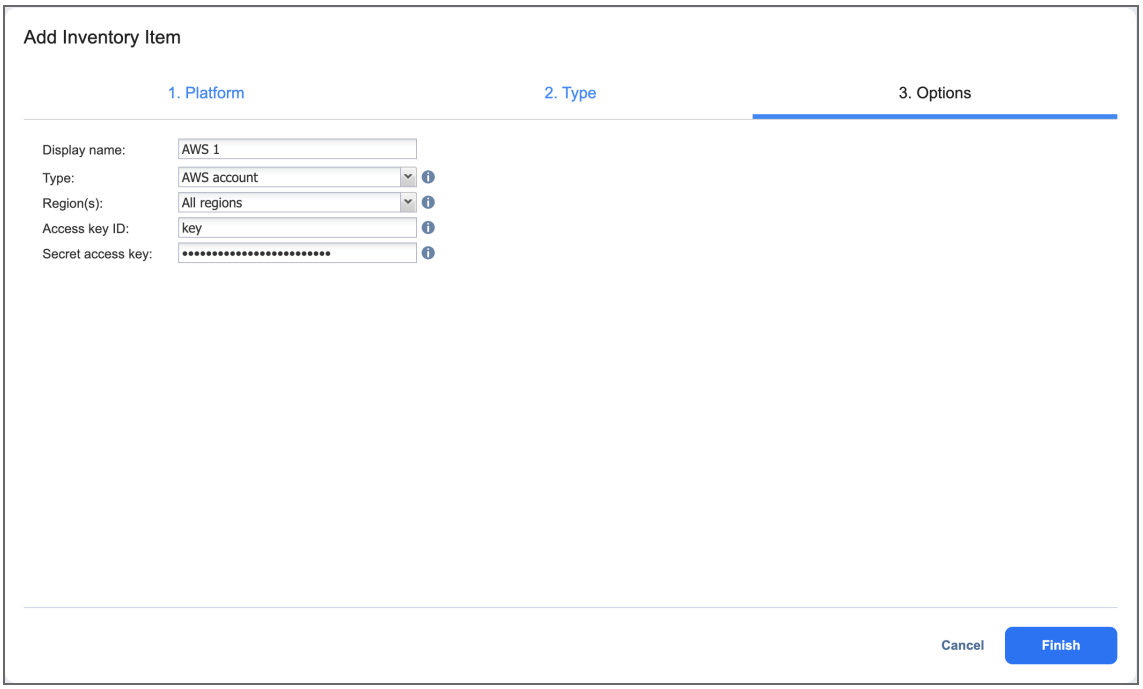
1. Click **Settings**.
2. Go to the **Inventory** page and click **+**.
3. On the **Platform** page of the wizard, select **Cloud Storage** and click **Next** to proceed.




4. On the **Type** page of the wizard, select **Amazon** and click **Next** to proceed.




5. On the **Options** page of the wizard, fill in the following fields:
   a. Enter the name in the **Display name** box.
   b. Select the AWS account from the **Type** drop-down list.
   c. Select the AWS region from the **Region(s)** drop-down list.
   d. Enter the **Access Key ID** and **Secret Access Key** of a root user or a sub-user in the corresponding fields.
6.

- 


# Adding Generic S3-Compatible Object Storage

> **Note**
>
> Only specific S3-compatible vendors are supported. Please see Feature Requirements for more information.

Before adding a generic S3-compatible object storage to the **Inventory**, obtain the following credentials for your storage service:

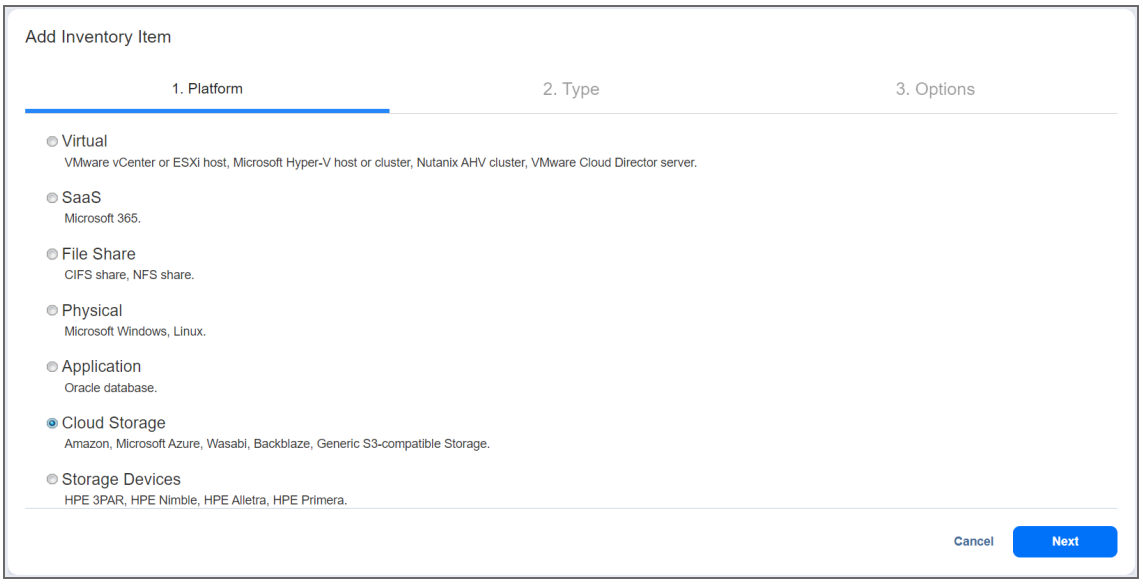
- Access Key
- Secret Key

If you don't have your storage credentials, check the service's web console for a corresponding section (or ask your storage service administrator for assistance). If needed, generate a new **Access Key** and **Secret Key** in the web console or via SSH.

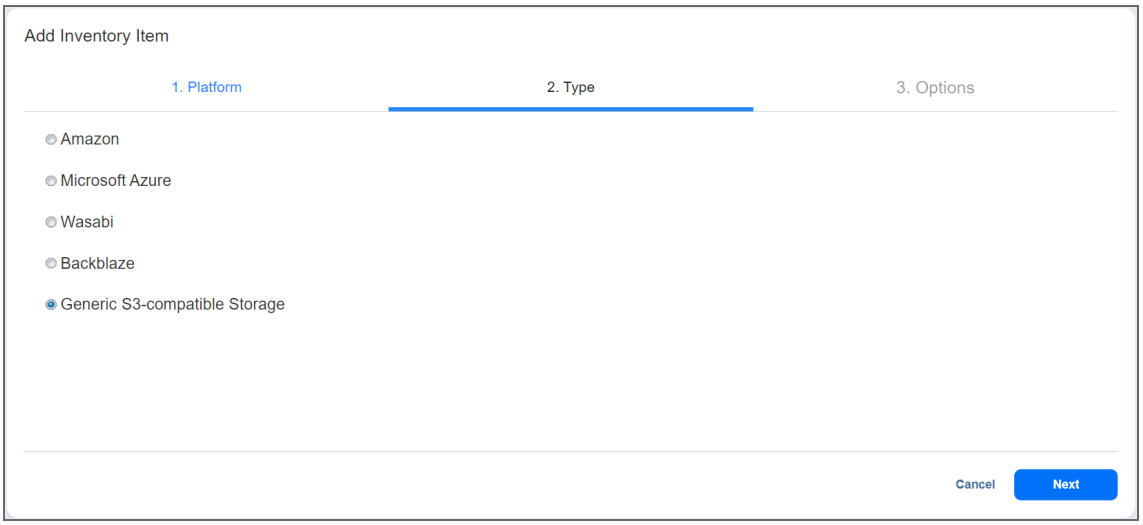
> **Note**
>
> For Google Cloud Storage, ensure the Google Cloud is properly configured:
>
> - Default project for interoperable access is set in Google Cloud Console.
> - HMAC keys are created.
> - Cloud Storage request endpoint configured.
>
> For more information, see How to add Google Cloud Storage.

To add S3-compatible object storage to Inventory in NAKIVO Backup & Replication, follow the steps below:

1. Verify that the S3-compatible object storage meets Feature Requirements.
2. Click **Settings** in the main menu of NAKIVO Backup & Replication.
3. Go to the **Inventory** page and click "**+**".
4. On the **Platform** page of the wizard, select **Cloud Storage** and click **Next** to proceed.




5. On the **Type** page of the wizard, select **Generic S3-compatible Storage** and click **Next** to proceed.

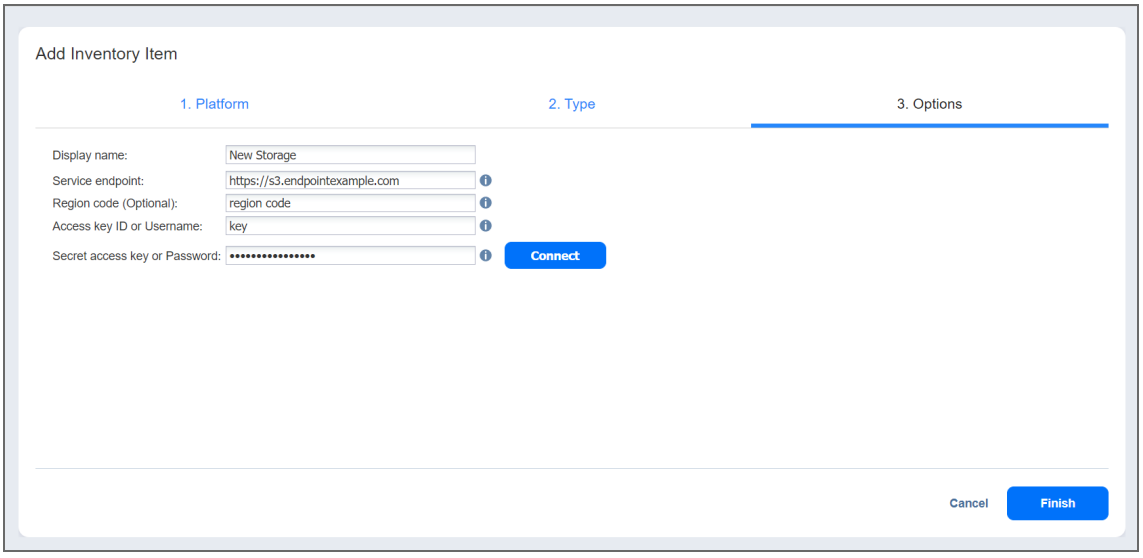


6. On the **Options** page of the wizard, provide the following information:
    - **Display name**: Specify a name for the S3-compatible object storage device. This name will be displayed in the **Inventory**.
    - **Service endpoint**: Enter a full HTTP/HTTPS URL that is used to access the storage.
    - **Region code**: Optionally, enter the technical region code where the data is stored. To enter several region codes, separate them using the semicolon ";" symbol. It is highly recommended to leave this field blank.
    - **Access key ID or Username**: Enter the storage access key ID or username that was created during account setup or on the App Keys page in your storage account.
    - **Secret access key or Password**: Enter the storage secret access key or password that was created during account setup or on the App Keys page in your storage account.
    - Click **Connect** to bring up the **Certificate Details** popup.

    **Note**

    The **Connect** button may be disabled if no HTTP/HTTPS certificates are detected or required for the storage to be added to the Inventory.

7. Optionally, you can select **Auto accept new certificate if this certificate is expired or changed**.
8. Click **Accept** to confirm the certificate.




9. Finally, click **Finish** when you are done.

    **Note**

    When configuring NGINX Proxy for MinIO Server used as the vendor for generic S3-compatible object storage, make sure you either turn caching off (*proxy_cache off;*) or set *proxy_cache_convert_head off;*.

    Find more information on the topic here.

# Adding Wasabi Accounts

Add a Wasabi account to NAKIVO Backup & Replication as described in the sections below.

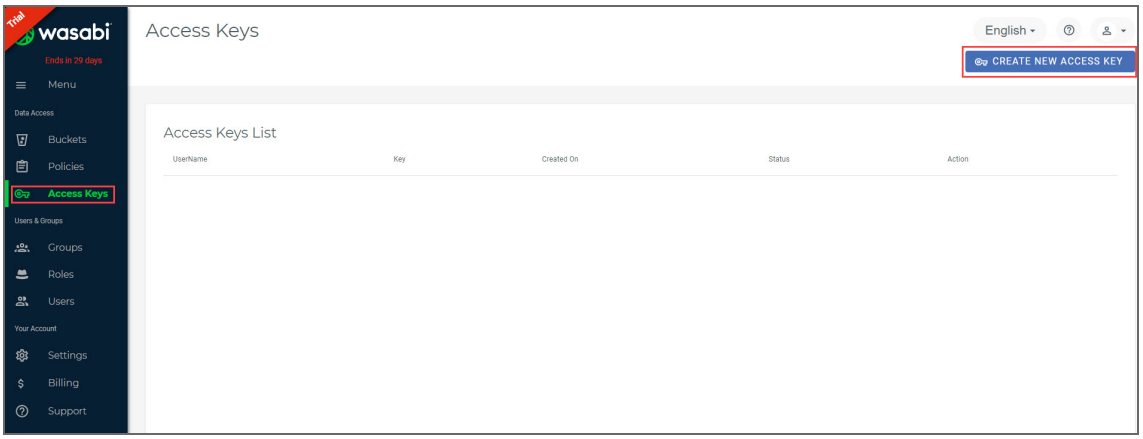
- Creating Wasabi Access Key ID and Secret Access Key
- Adding a Wasabi Account to Inventory

## Creating Wasabi Access Key ID and Secret Access Key

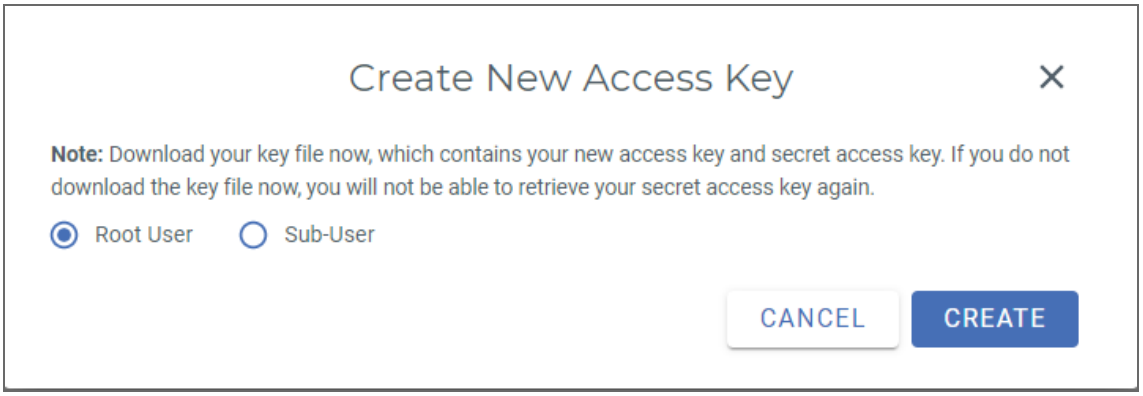
Prior to adding your Wasabi account to the inventory, you need to create and retrieve a Wasabi Access Key ID and Secret Access Key. They are used by NAKIVO Backup & Replication to sign the programmatic requests sent to Wasabi, such as retrieving the list of instances, creating snapshots, and etc.

To create a Wasabi Access Key ID and a Secret Access Key, follow the steps below:

1. If you don't have a Wasabi account, create a new one at wasabi.com/sign-up/.
2. Log in to your Wasabi account.
3. Navigate to the main menu and click **Access Keys**.
4. Click **Create New Access Key**.




5. In the dialog box that opens, select one of the following:
   - **Root user**: Select this option and click **Create**.

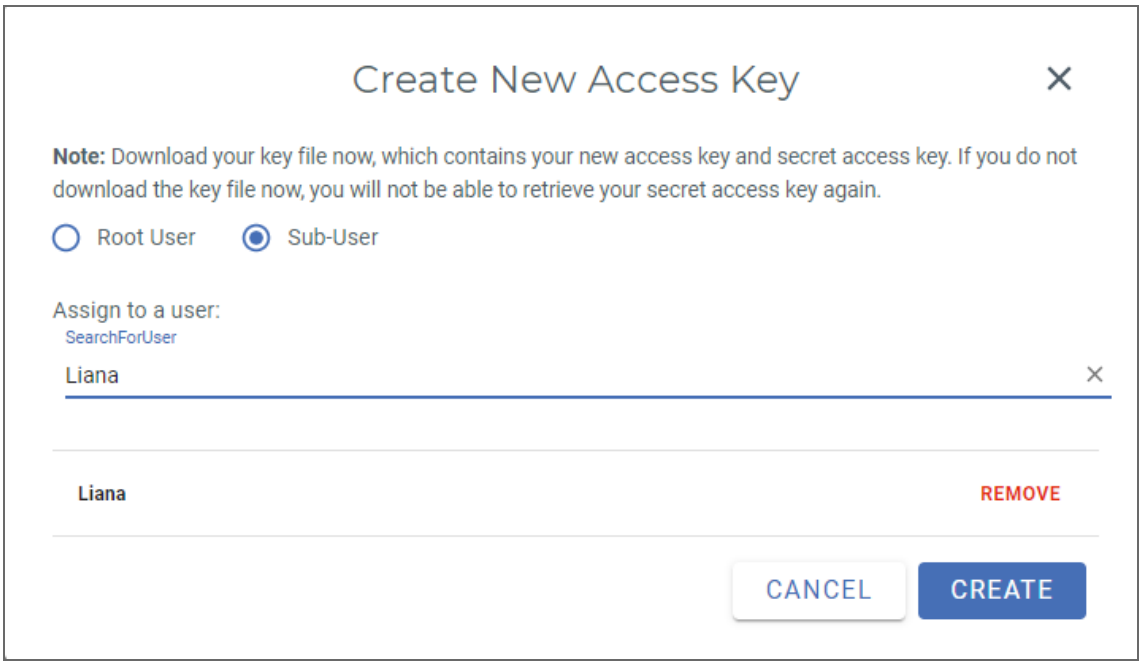


- **Sub-User**: Select a sub-user from the **Assign to a user** drop-down list and click **Create**. Note that the original user access key of the selected user will be changed.

**Notes**

- To use the **Sub-User** option, you need to have at least one user created in your Wasabi account. For details, refer to Creating a User and How do I set up Wasabi for user access separation?
- For the sub-user to use the immutable backup, the following set of IAM rules must be used:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "BucketPolicy",
            "Effect": "Allow",
            "Action": "s3:*",
            "Resource": [
                "arn:aws:s3:::<YOUR BUCKET NAME>/*",
                "arn:aws:s3:::<YOUR BUCKET NAME>"
            ]
        },
        {
            "Sid": "RepoRefresh",
            "Effect": "Allow",
            "Action": "s3:ListAllMyBuckets",
            "Resource": "*"
        }
    ]
}
```



6. Click **Download CSV** and save the file with the generated keys in a secure location. Keep the Access key confidential to protect your account.

## Adding a Wasabi Account to Inventory

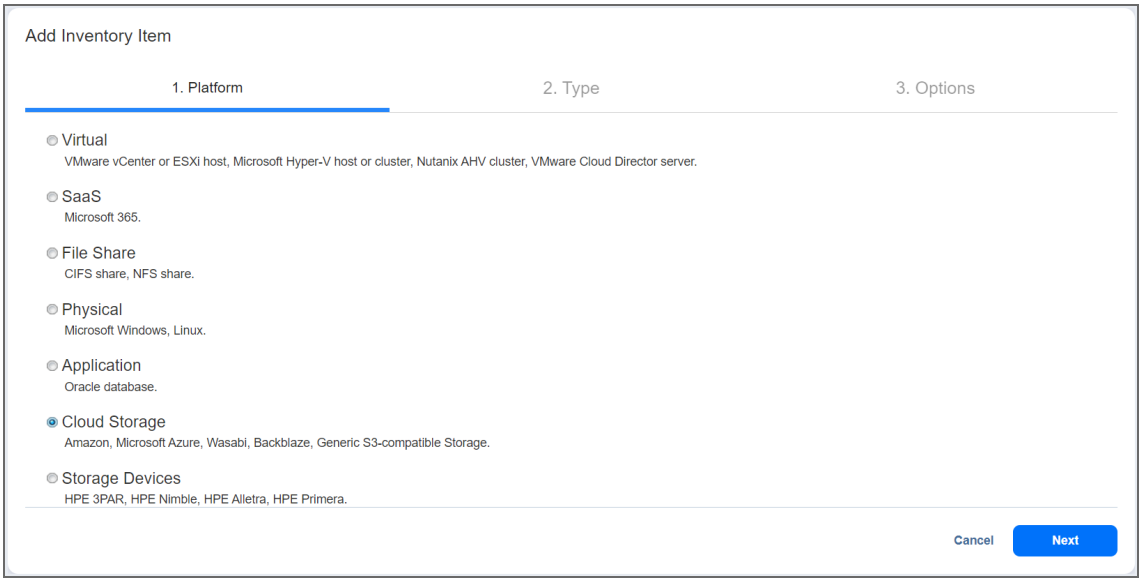
Before adding a Wasabi Account to **Inventory**, make sure that your system can reach the storage provider:

- Check if the hostname of the Wasabi Storage can be resolved correctly by the Transporter (that you assign to the Repository).
- Confirm that port 443 is open.
- Check SSL/TLS connectivity to the storage provider.

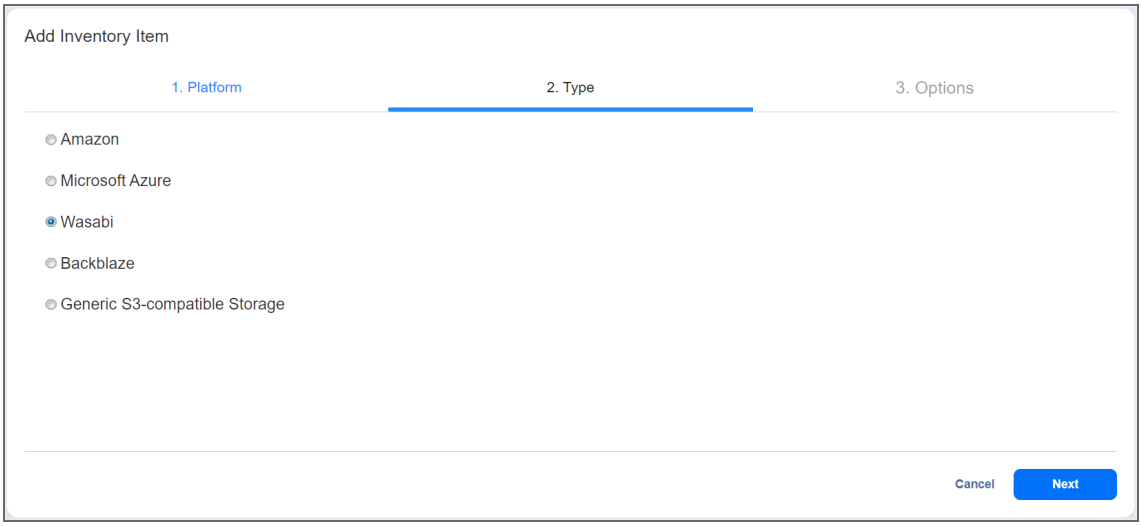
For more information, refer to Troubleshooting S3 Compatible Repositories.

To add a Wasabi account to NAKIVO Backup & Replication, follow the steps below:

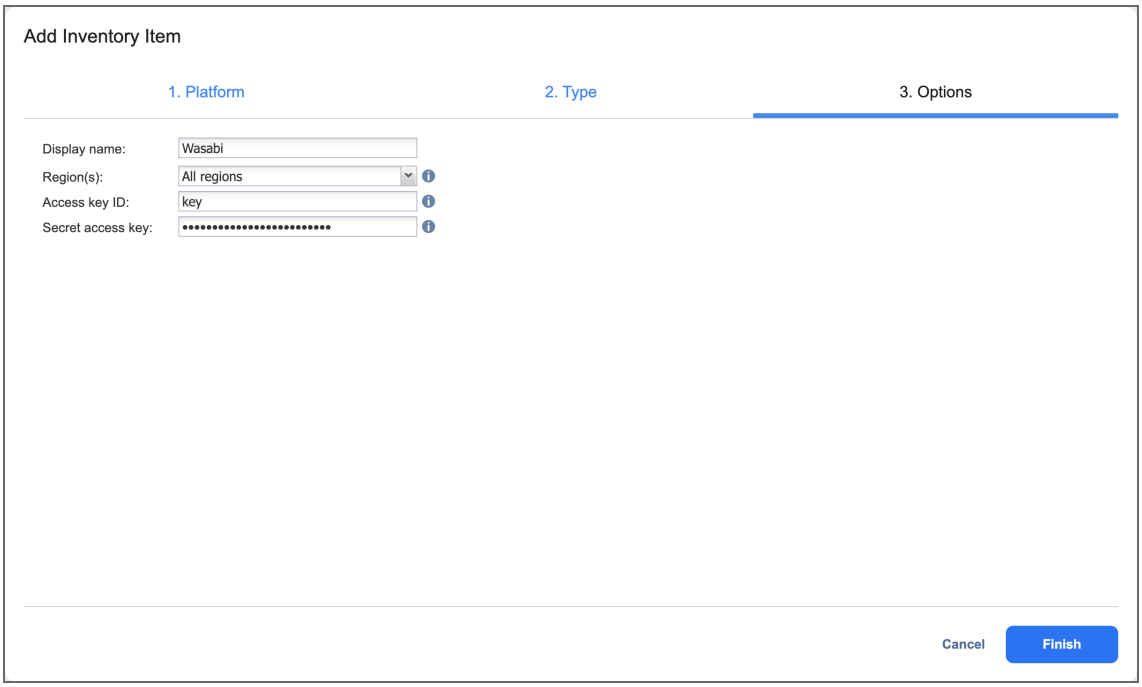
1. Click **Settings**.
2. Go to the **Inventory** page and click **+**.
3. On the **Platform** page of the wizard, select **Cloud Storage** and click **Next** to proceed.




4. On the **Type** page of the wizard, select **Wasabi** and click **Next** to proceed.




5. On the **Options** page of the wizard, fill in the following fields:
   a. Enter the name in the **Display name** box.
   b. Select the Wasabi region from the **Region(s)** drop-down list.
   c. Enter the **Access Key ID** and **Secret Access Key** of a root user or a sub-user in the corresponding fields.
6. Click **Finish**.




# Adding Backblaze Accounts

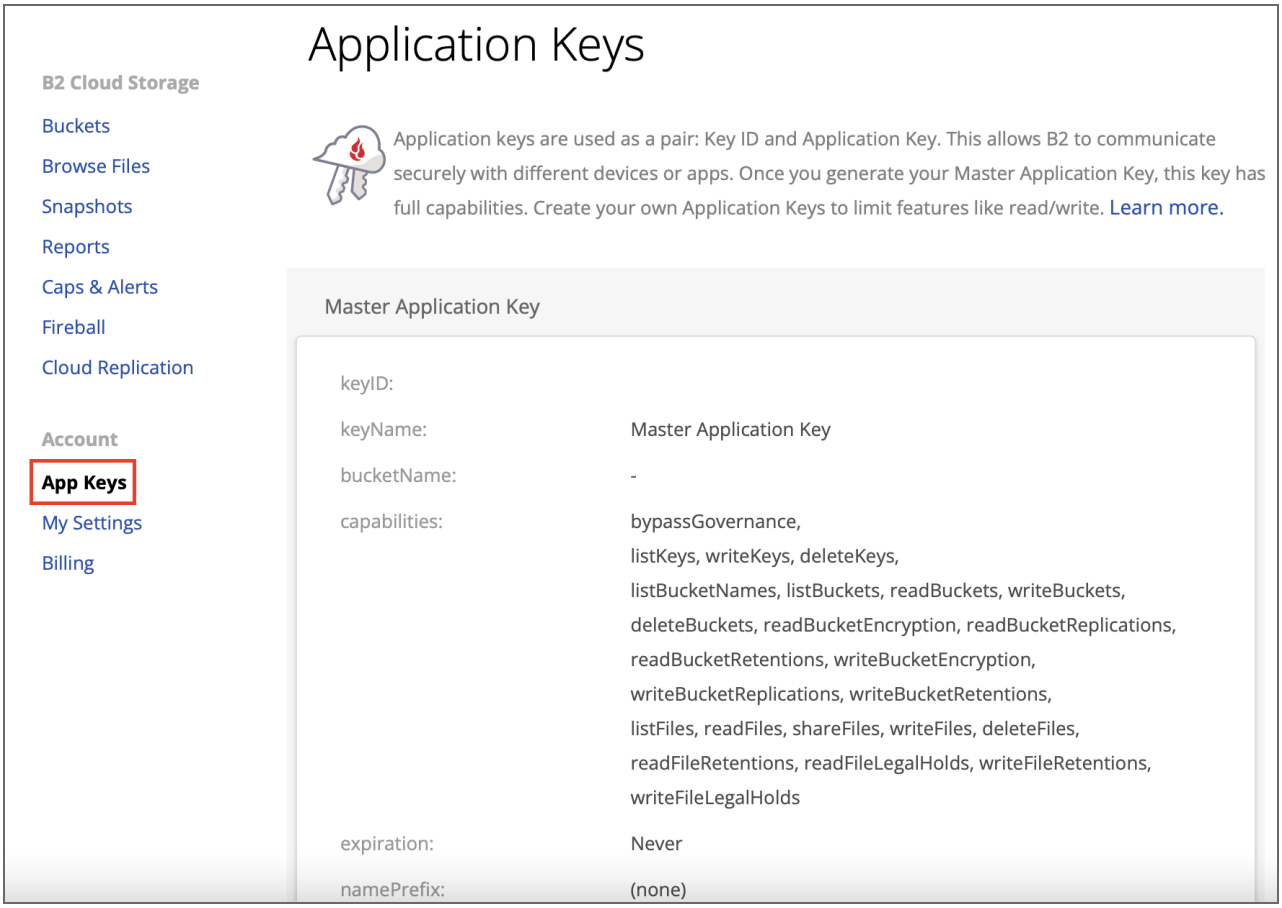
Add a Backblaze account to NAKIVO Backup & Replication as described in the sections below.

- Obtaining Backblaze Credentials
- Adding a Backblaze Account to Inventory

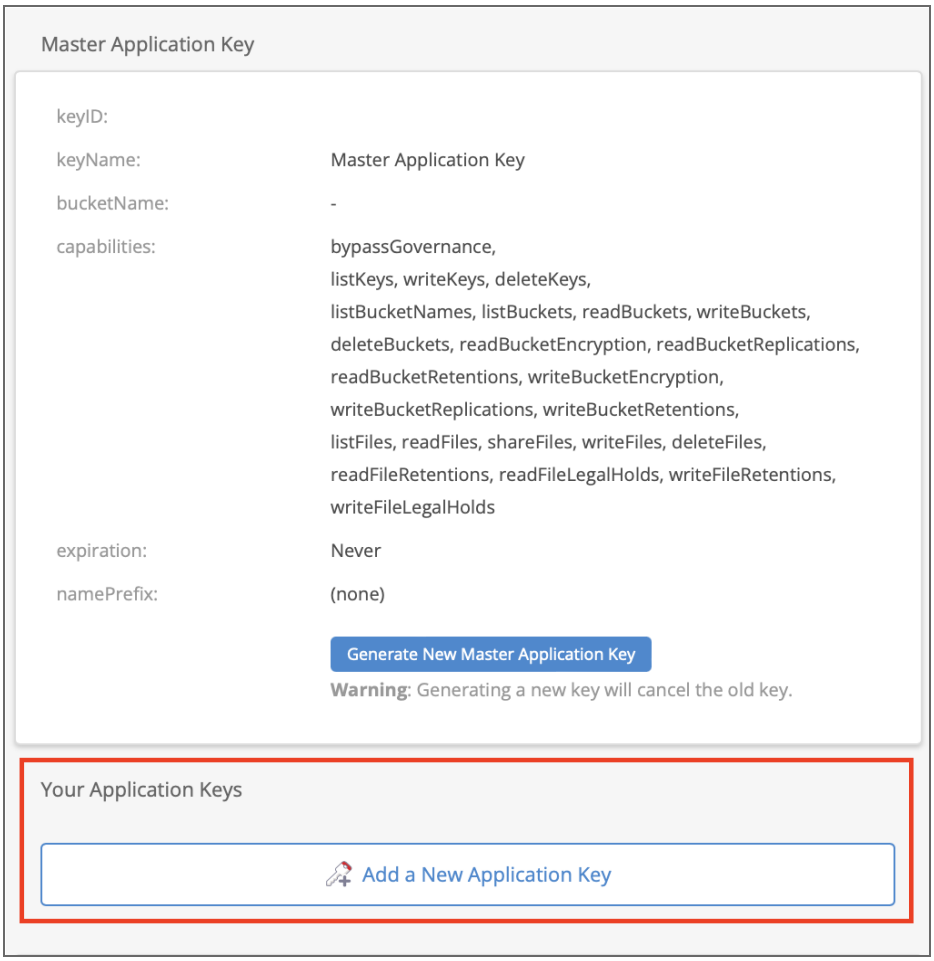
# Obtaining Backblaze Credentials

If you have not already generated and saved application key information in your Backblaze account, you will need to do so. To obtain the credentials required to add a Backblaze account to the NAKIVO Backup & Replication **Inventory**, follow the steps below:

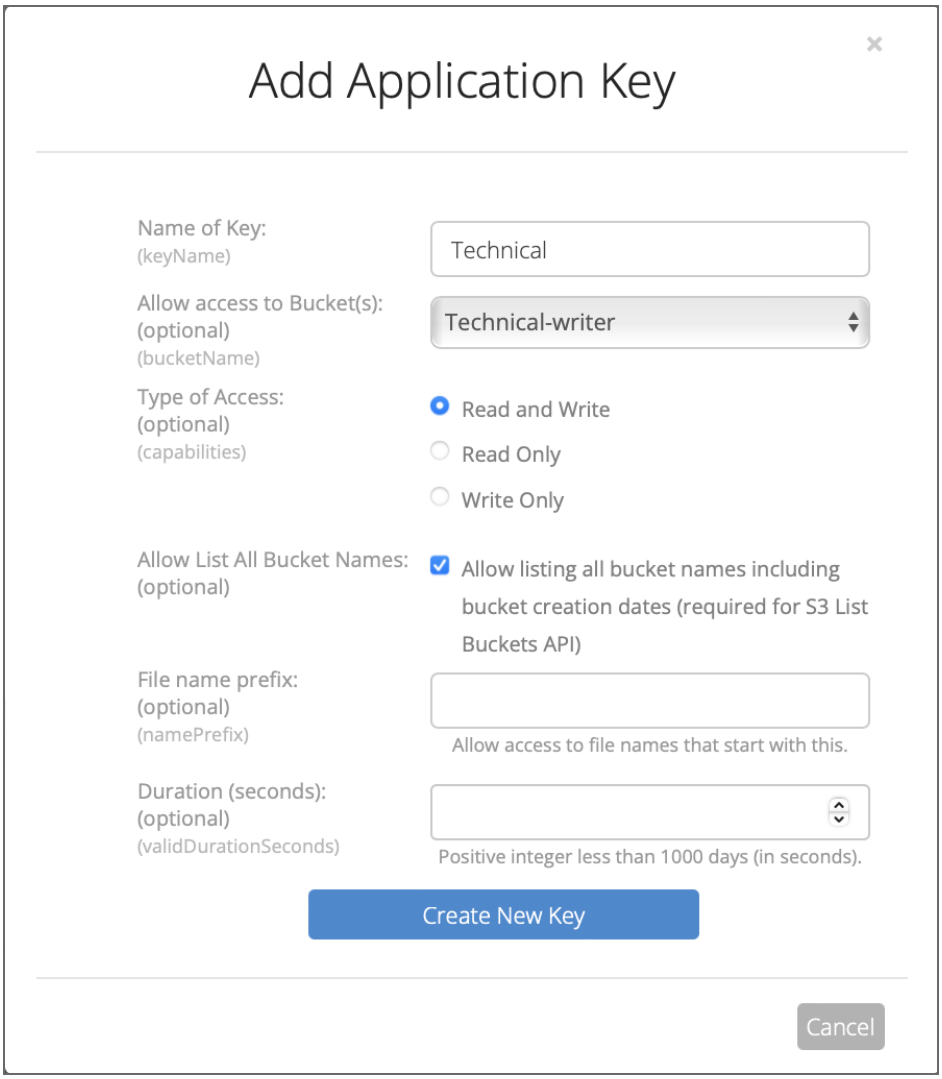
1. Log in to your Backblaze account.
2. Locate the **Account** tab on the left side and click **App Keys**.



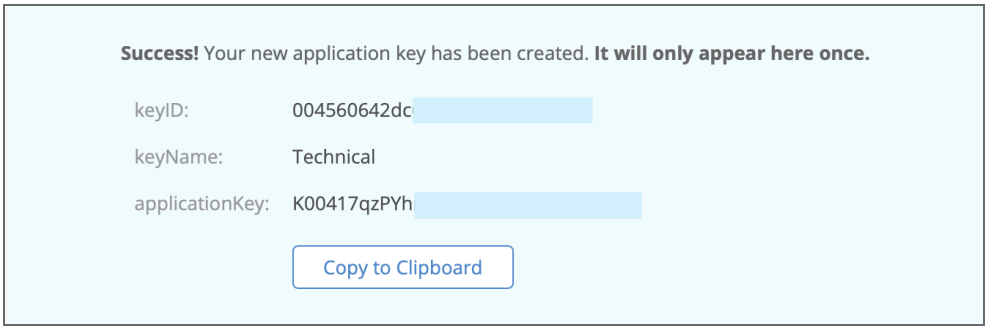

3. Find the **Add a New Application Key** button and click it.




4. Fill in the following information:
    - **Name of Key**: Enter a key name of your choice.
    - **Allow access to Bucket(s)**: Select a specific bucket or all buckets.
    - **Type of access**: Choose the level of access given to external applications.
    - **Allow List All Bucket Names**: Check to allow listing of all bucket names for S3 API purposes (required).
    - **File name prefix**: Limits access to files that begin with the specified text.
    - **Duration (seconds)**: Validity duration of the key in seconds (leave blank to keep it indefinite).
5. Click **Create New Key**.



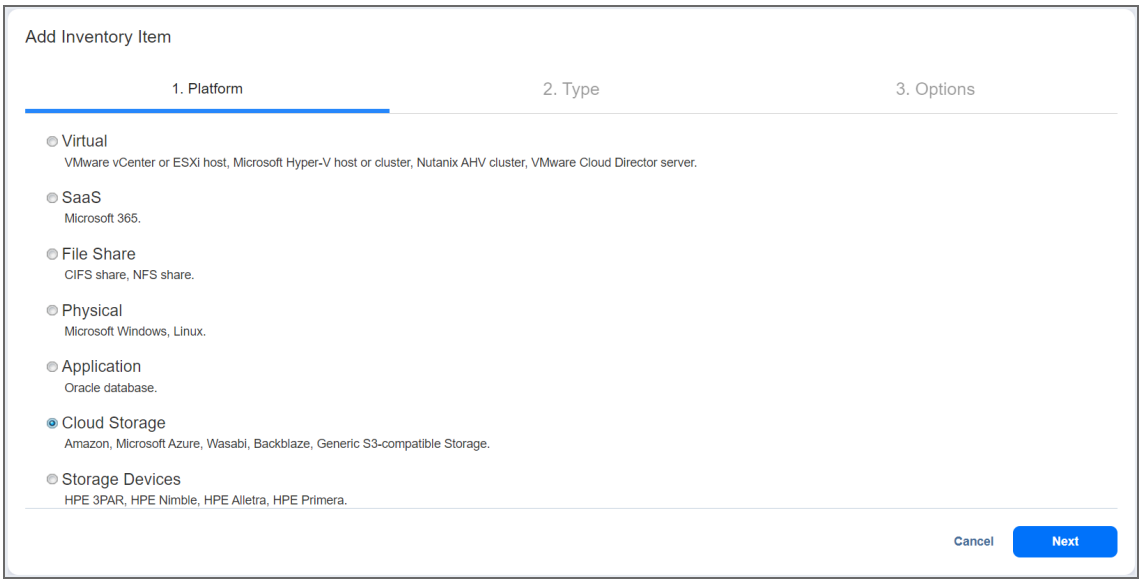

6. Locate the generated key and note the following information:
    - **keyID**
    - **Application Key**



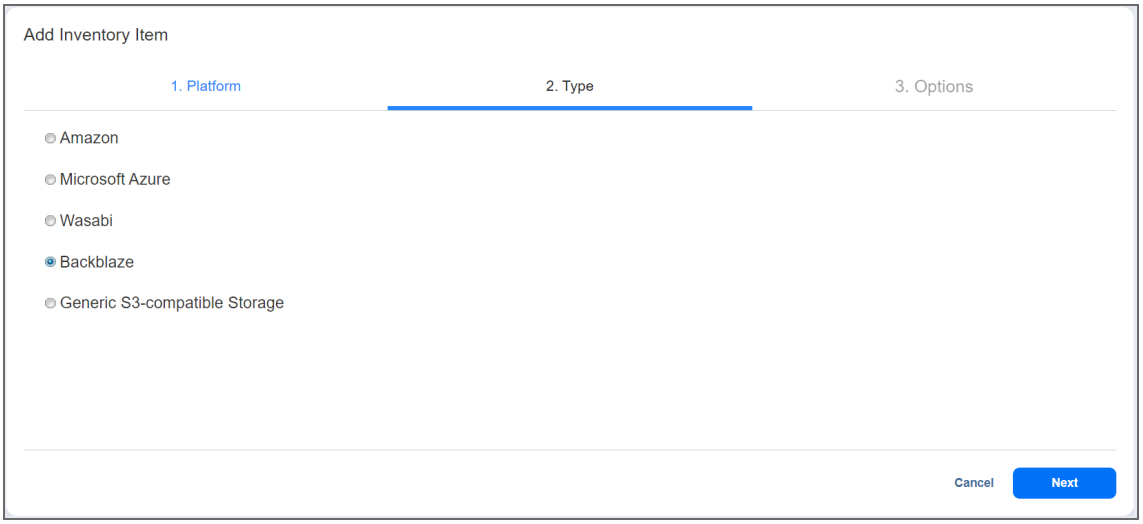

# Adding a Backblaze Account to Inventory

To add a Backblaze storage account to the **Inventory**, do the following:

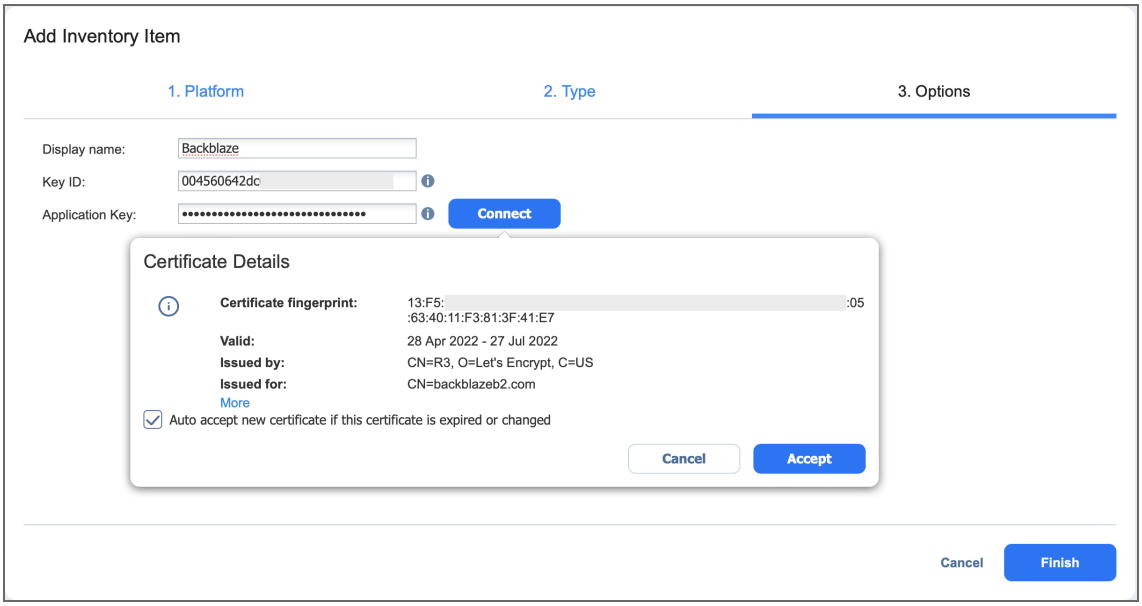
1. Click **Settings** in the left pane.
2. Go to the **Inventory** tab and click **+**.
3. On the **Platform** page of the wizard, select **Cloud Storage**. Click **Next** to proceed.




4. On the **Type** page, choose **Backblaze**. Click **Next** to proceed.




5. On the **Options** page, configure the following:
   - **Display name**: Enter a display name for the Backblaze storage account.
   - **Key ID**: Enter the **keyID** generated on the **App Keys** page in your Backblaze account.
   - **Application Key**: Enter the **Application Key** generated on the **App Keys** page in your Backblaze account.
6. Click **Connect**. This should bring up the **Certificate Details** pop-up window.




7. Optionally, you can select **Auto accept new certificate if this certificate is expired or changed**. Click **Accept** to confirm the certificate.
8. Click **Finish** to add the account to **Inventory**.

# Adding Microsoft Azure Storage Accounts

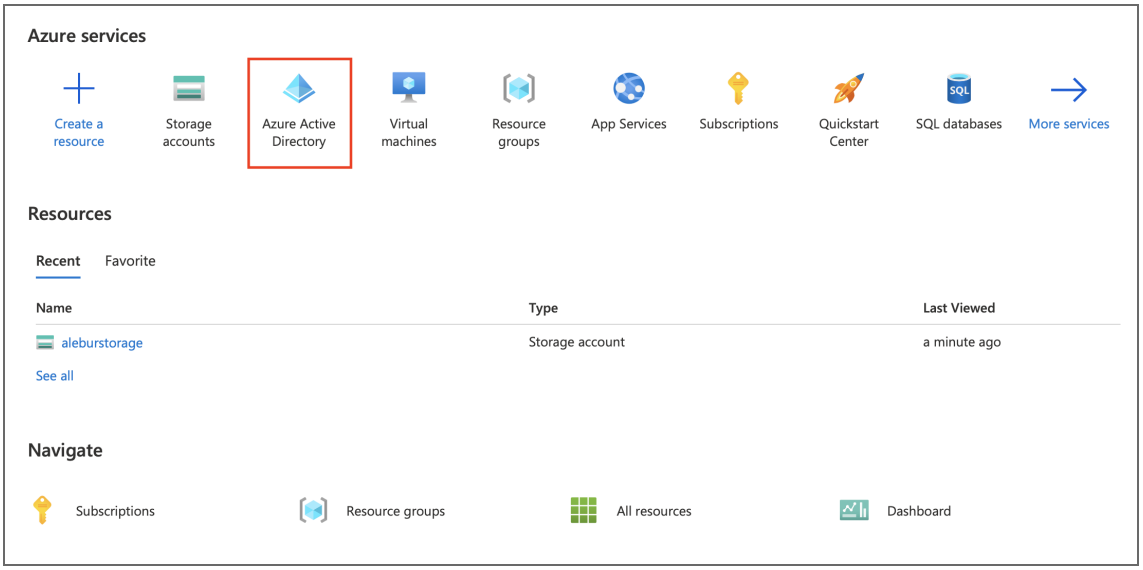
Configure and add a Microsoft Azure Storage account to NAKIVO Backup & Replication as described in the sections below.

- Configuring a Microsoft Azure Storage Account
- Obtaining Microsoft Azure Credentials
- Adding Microsoft Azure Storage Account to Inventory

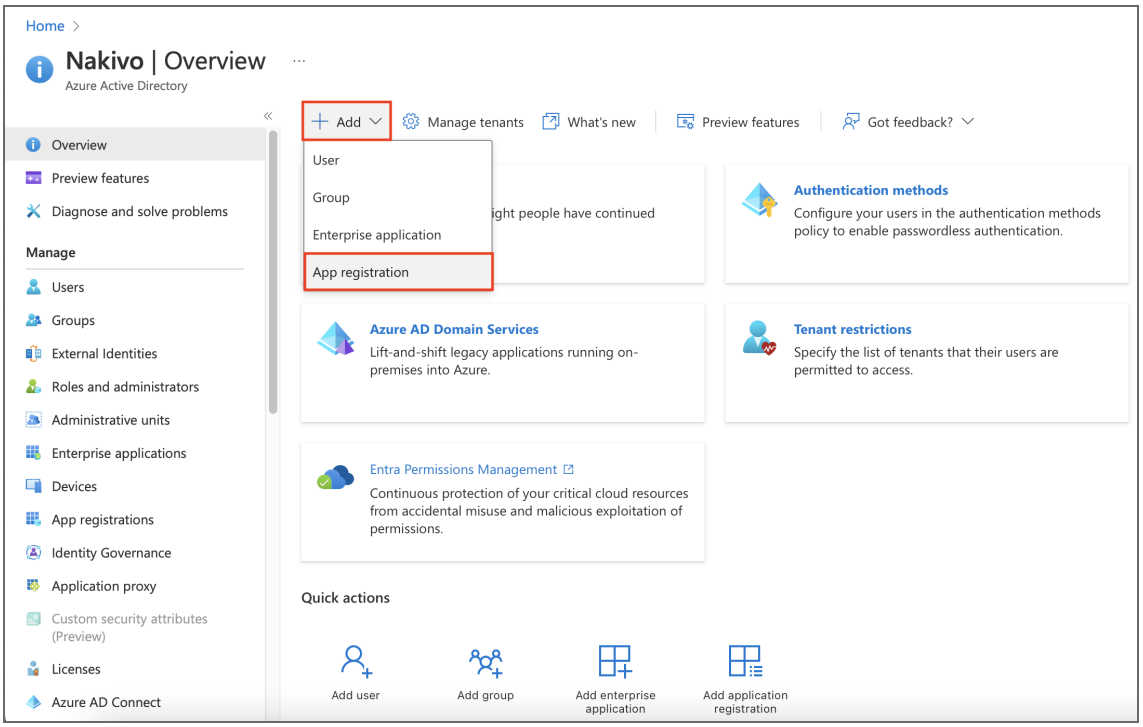
## Configuring a Microsoft Azure Storage Account

To configure a Microsoft Azure Storage account to work with NAKIVO Backup & Replication, follow the steps below.

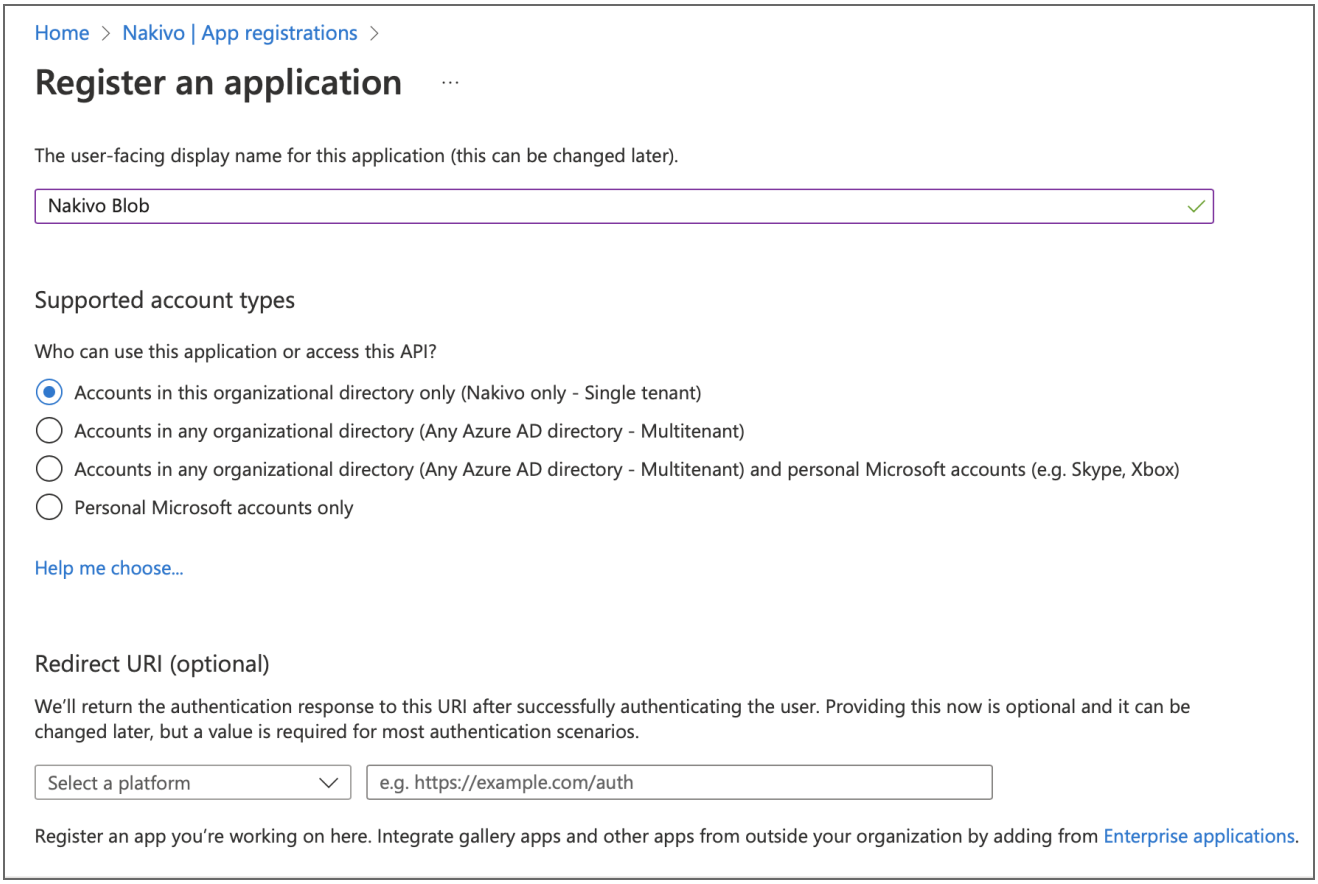
1. Open the Azure Portal by going to portal.azure.com.
2. Sign in to Microsoft Azure with your Microsoft account credentials.
3. Open **Azure Active Directory** from the services dashboard.



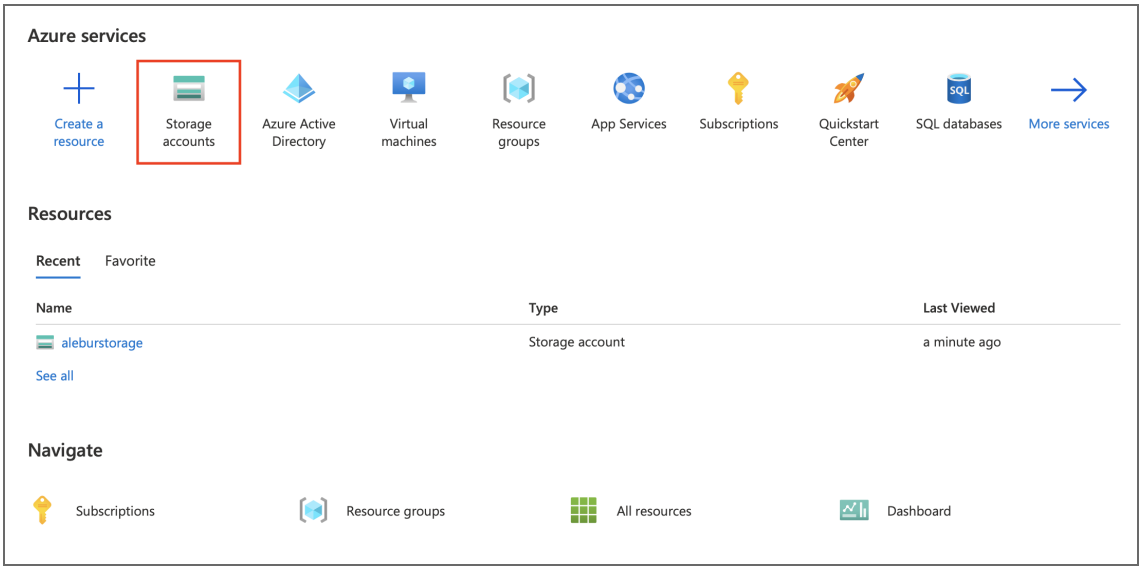

4. Register a new application by clicking **Add** > **App registration** from the **Overview** or **App registrations** menu. If you already have an application for use with NAKIVO Backup & Replication, skip to step 6.



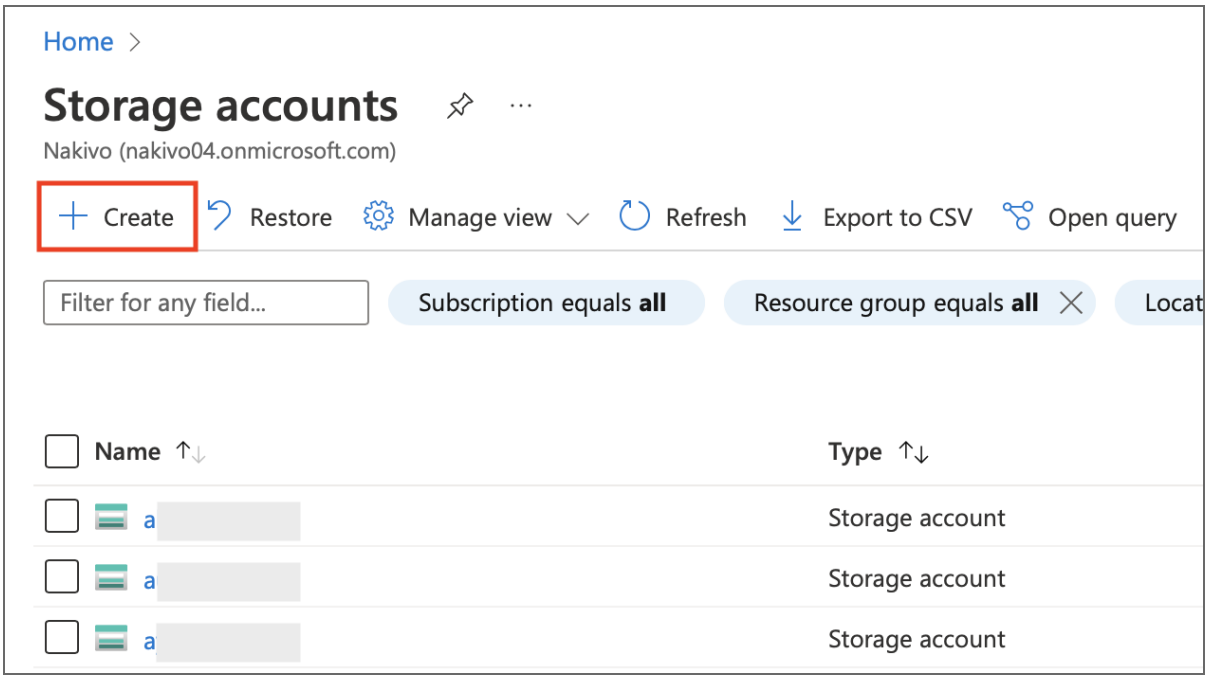

5. Enter a name for your application and set the access level per your requirements. When you're done, click **Register**.



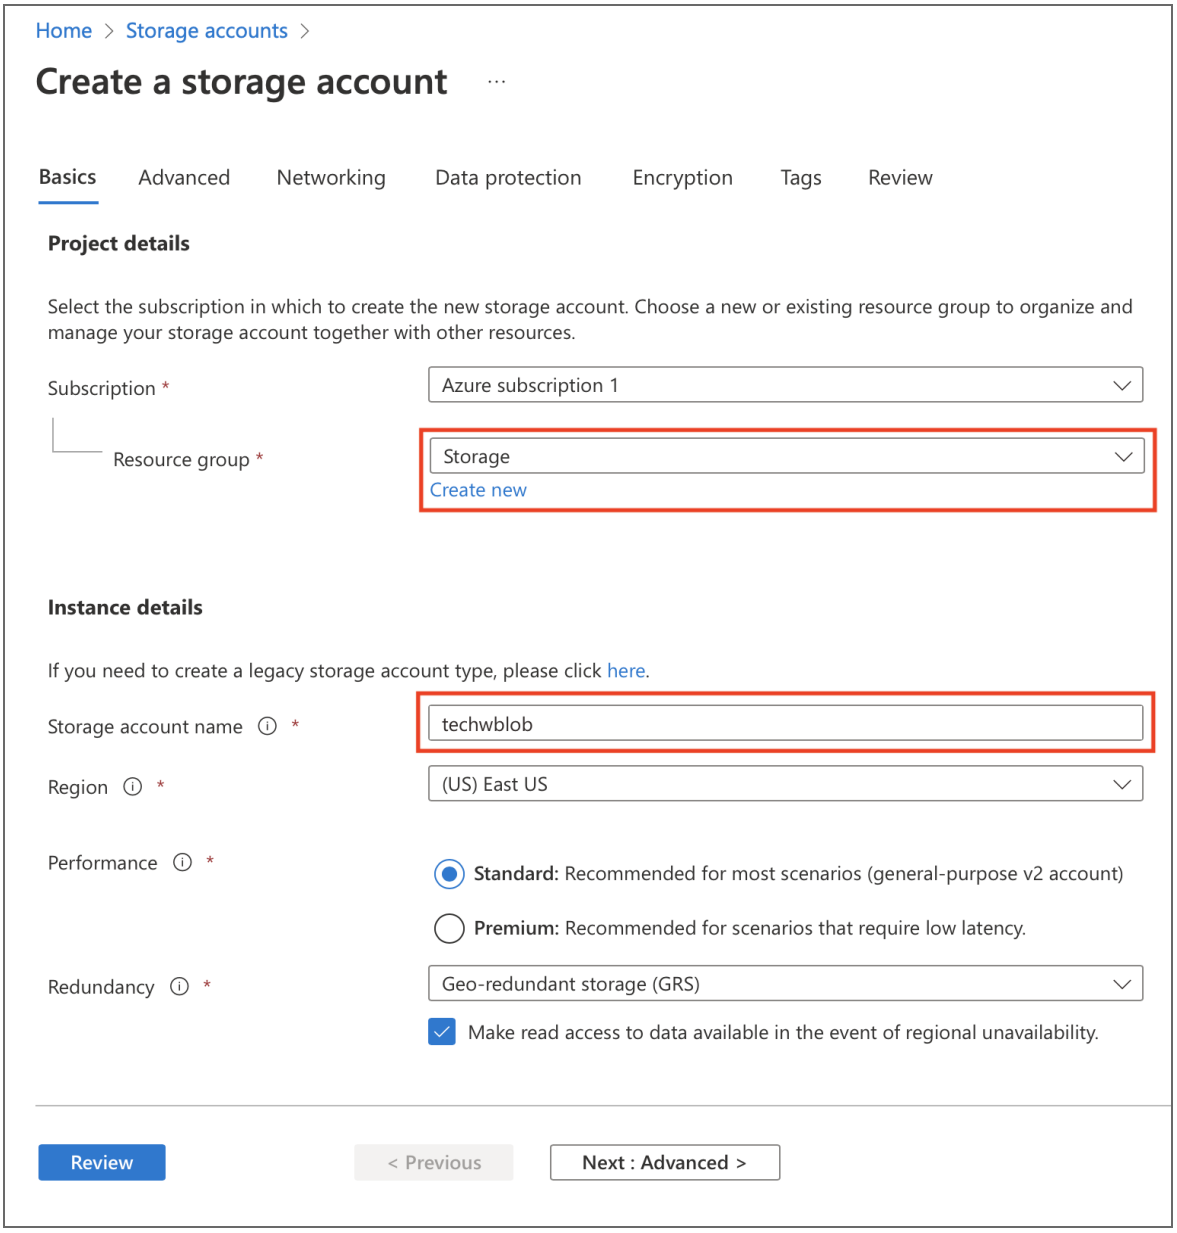

6. Next, return to the Azure homepage an open **Storage accounts** from the services dashboard.



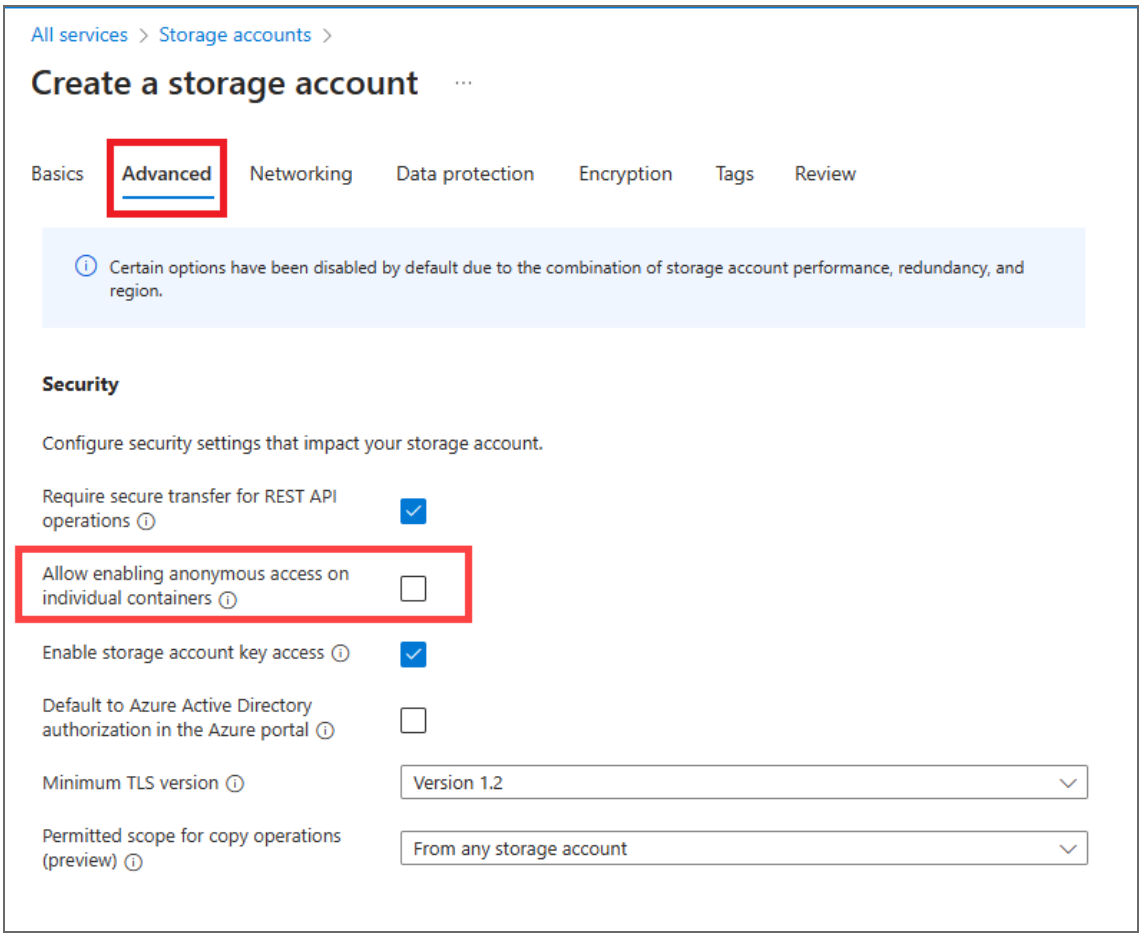

7. Click **Create** to create an Azure storage account. If you already have a storage account, skip to step 9.



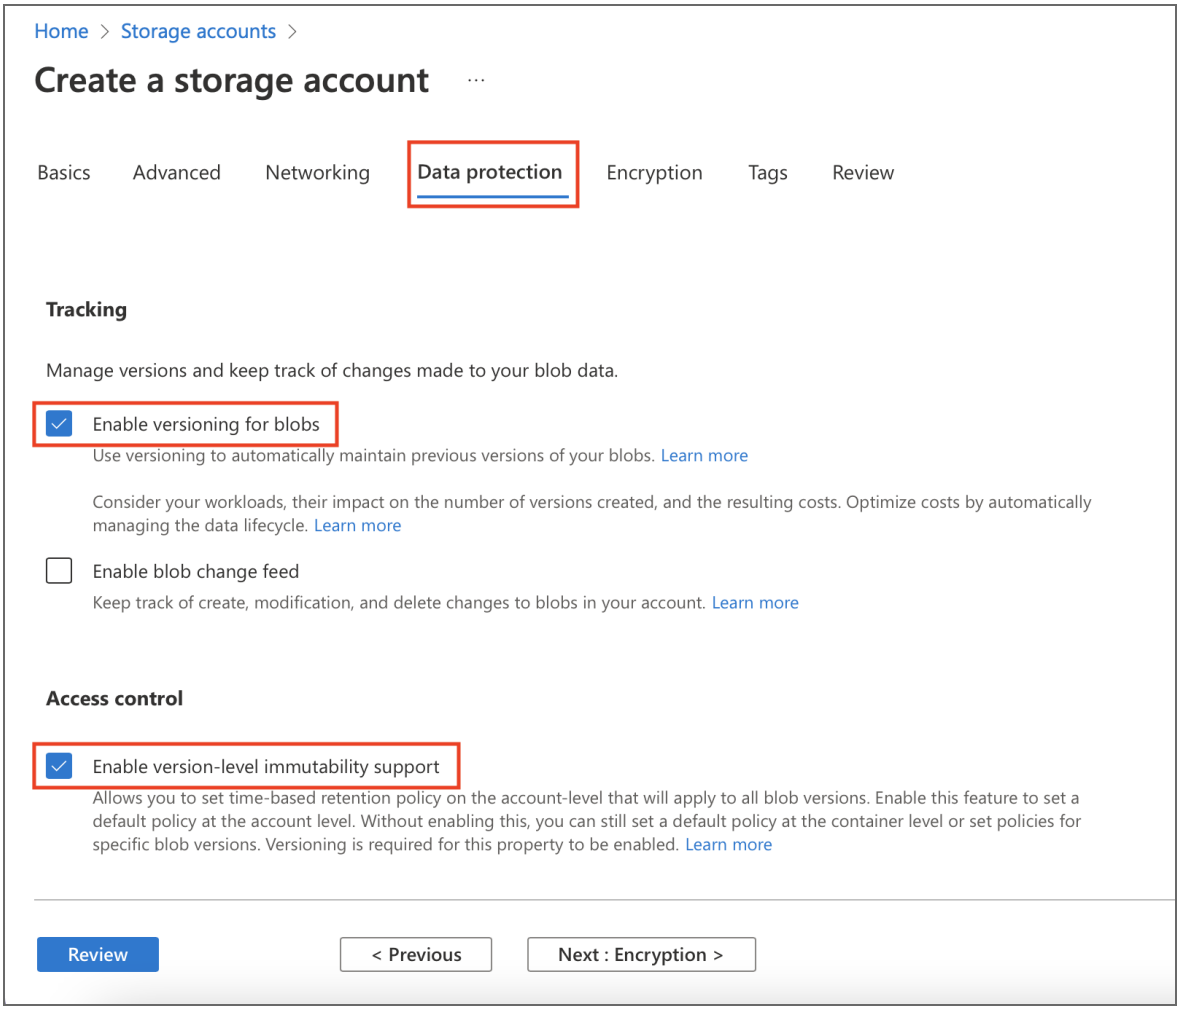

8. Select the appropriate **Subscription** and **Resource group** from the respective drop-down menus. You may also create a new resource group by clicking the **Create new** button under the **Resource group** drop-down menu. Name your storage account and configure the **Region**, **Performance**, and **Redundancy** settings based on your preference.




9. On the **Advanced** tab, it is recommended to disable **Allow enabling anonymous access on individual containers** option.



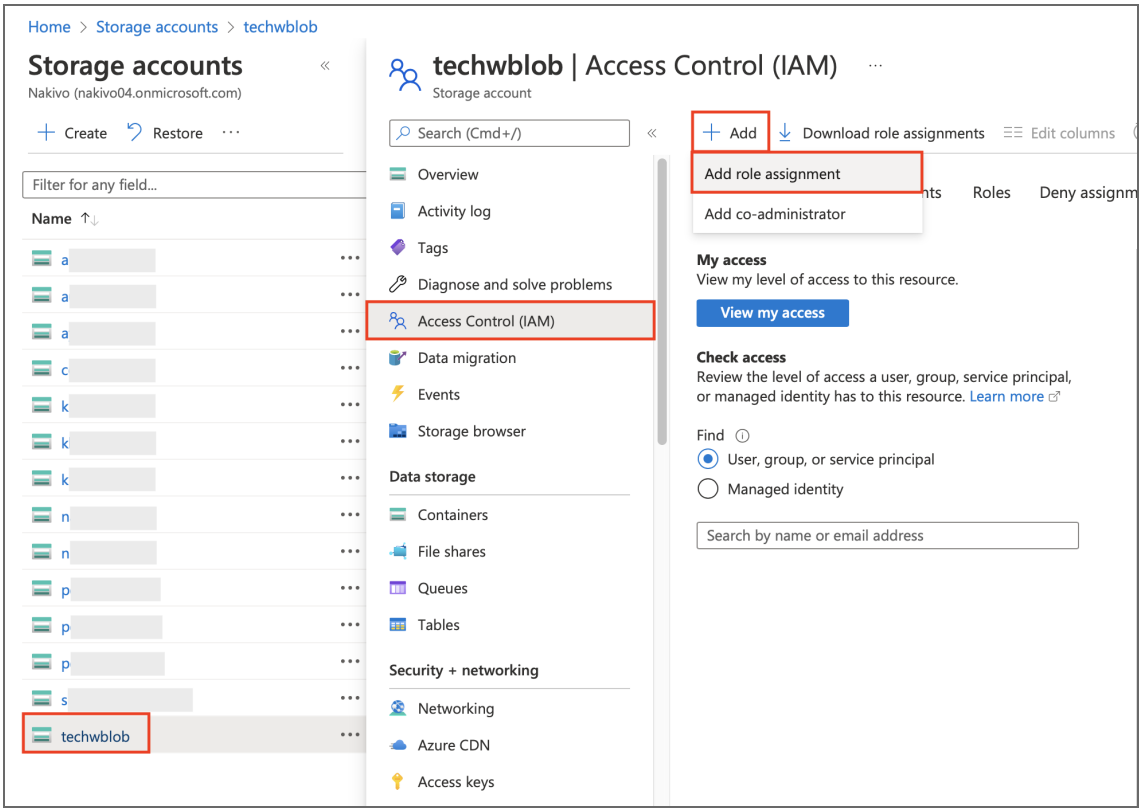

10. If you wish to enable Backup Immutability for this storage account, go to the **Data protection** tab. Under **Tracking**, find and enable the **Enable versioning for blobs** setting. Under **Access control**, find and enable the **Enable version-level immutability support** setting.



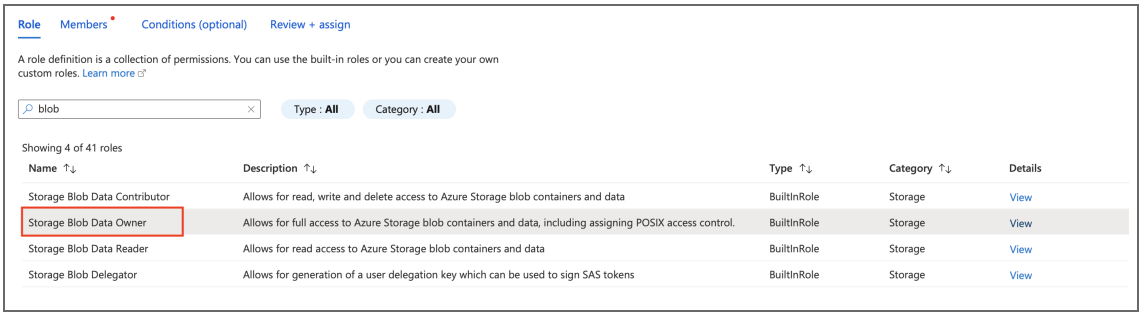

11. Optionally, configure advanced settings within the other tabs. When you're done, click **Review**. Review the account configuration and click **Create** if everything is in order.

> **Note**
> After clicking **Create**, the storage account will undergo a short deployment stage before it appears in the **Storage accounts** menu.

12. Locate your storage account in the **Storage accounts** menu and click on it to open the account settings. Go to the Access Control (IAM) tab and click **Add** > **Add role assignment**.




13. Find the **Storage Blob Data Owner** role and select it. Click **Next**.

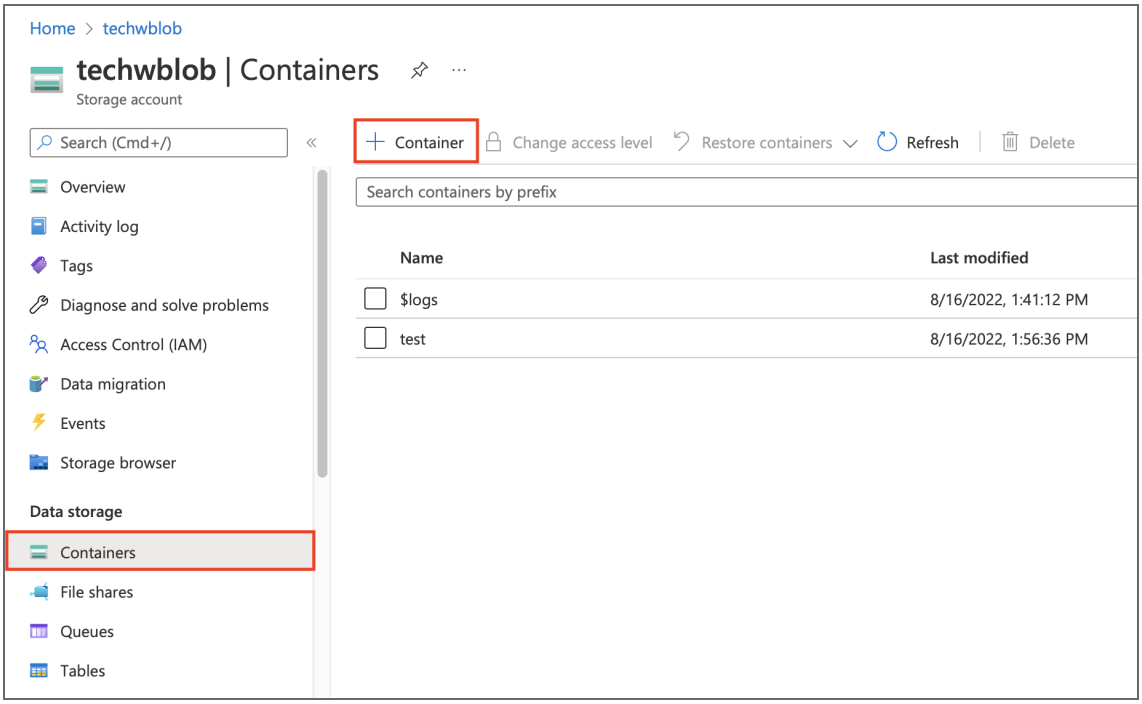



14. Click **Select members** and find the application registered in the previous steps using the search bar. Click on the application name and click **Select** below to confirm. Click **Review + assign** to add the role.

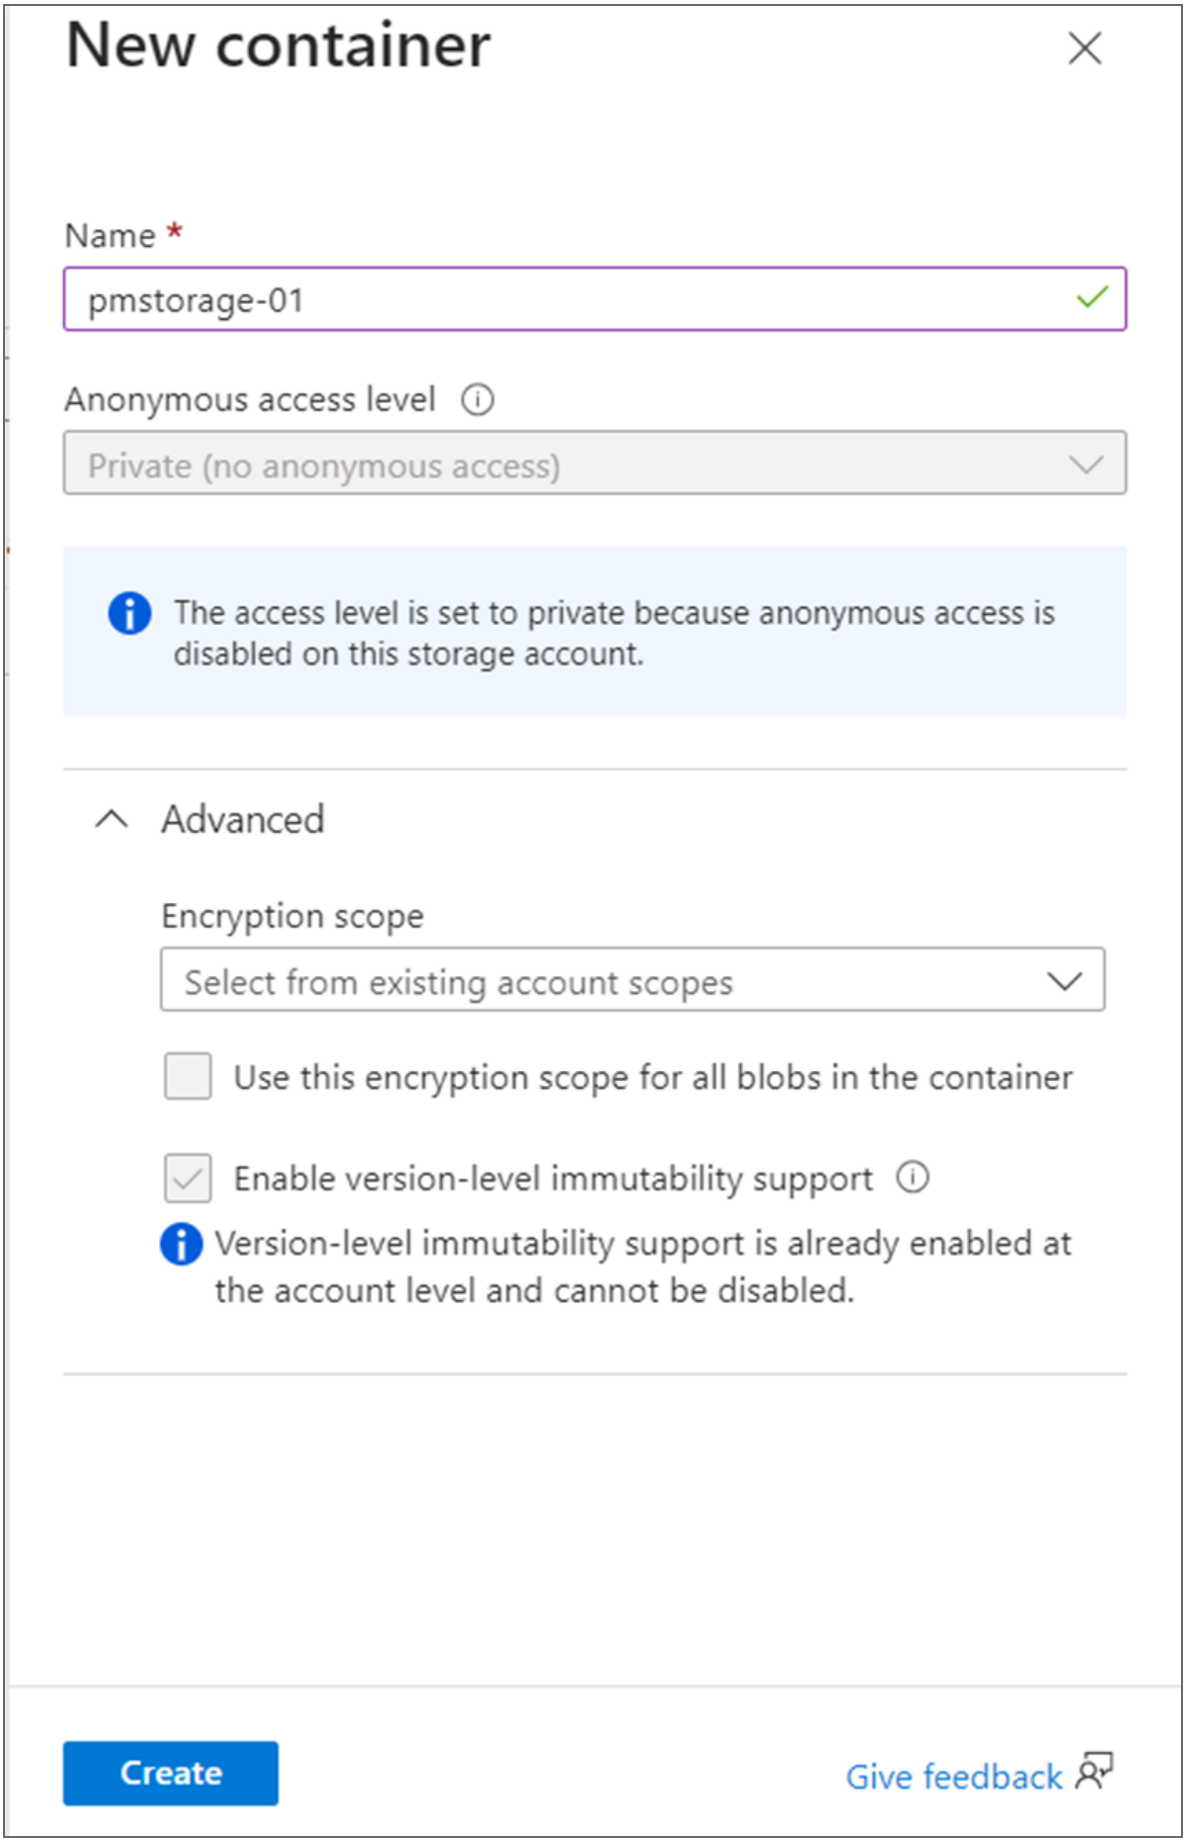


1. To add storage containers to the storage account and configure immutability, go to the **Containers** tab and click **+ Container**.

> **Note**
> NAKIVO Backup & Replication automatically detects containers created after adding a Microsoft Azure Storage account to inventory.

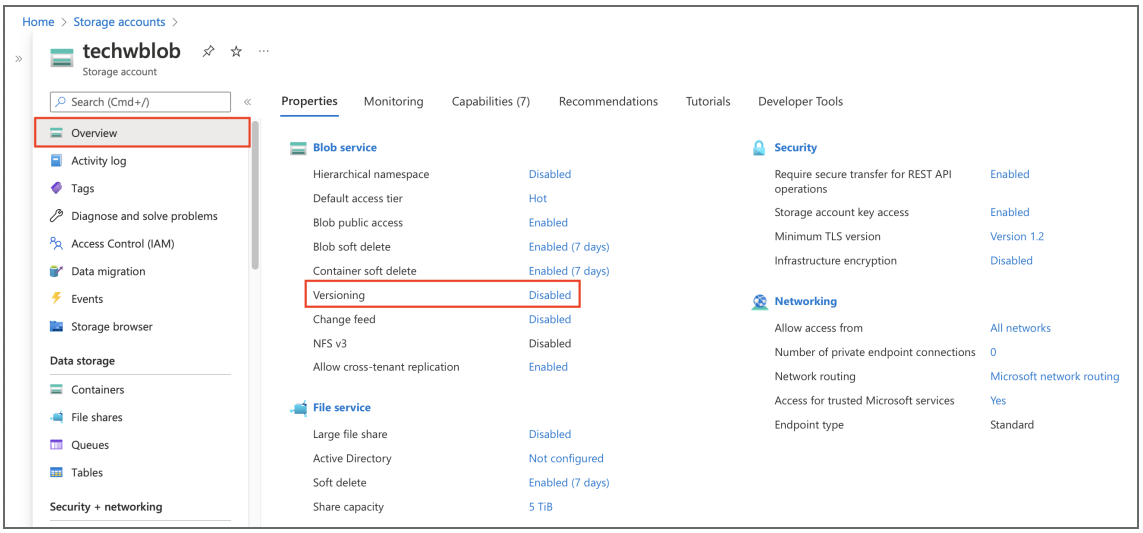



2. Name the container and configure its access level as needed. Select **Enable version-level immutability support** under **Advanced** settings if you wish to enable Backup Immutability for this container.

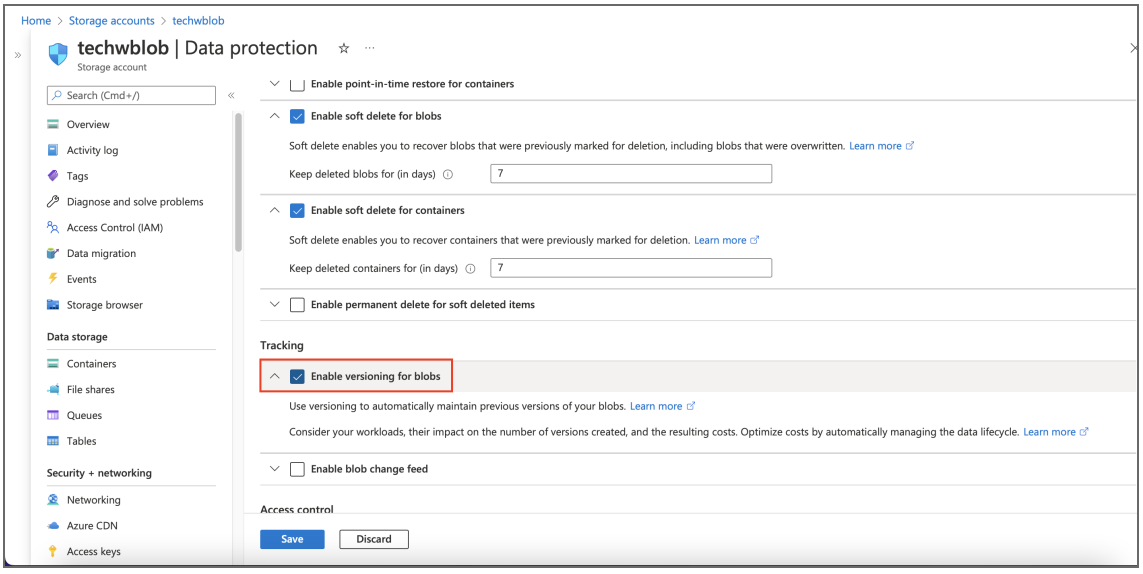

New container
Name *
pmstorage-01
Anonymous access level
Private (no anonymous access)
The access level is set to private because anonymous access is disabled on this storage account.
Advanced
Encryption scope
Select from existing account scopes
Use this encryption scope for all blobs in the container
Enable version-level immutability support
Version-level immutability support is already enabled at the account level and cannot be disabled.
Create
Give feedback

> **Note**
>
> If your storage account does not have version-level immutability support enabled (as described in step 9), you will need to enable this option per container. Existing containers without the **Enable version-level immutability support** option enabled will not be able to make use of Backup Immutability.

15. If you enabled version-level immutability support in any of the previous steps, also make sure that versioning for blobs is enabled. Return to the storage account's Overview menu and scroll down to find **Versioning** in the **Properties** tab. If your versioning is **Disabled**, click **Disabled**.



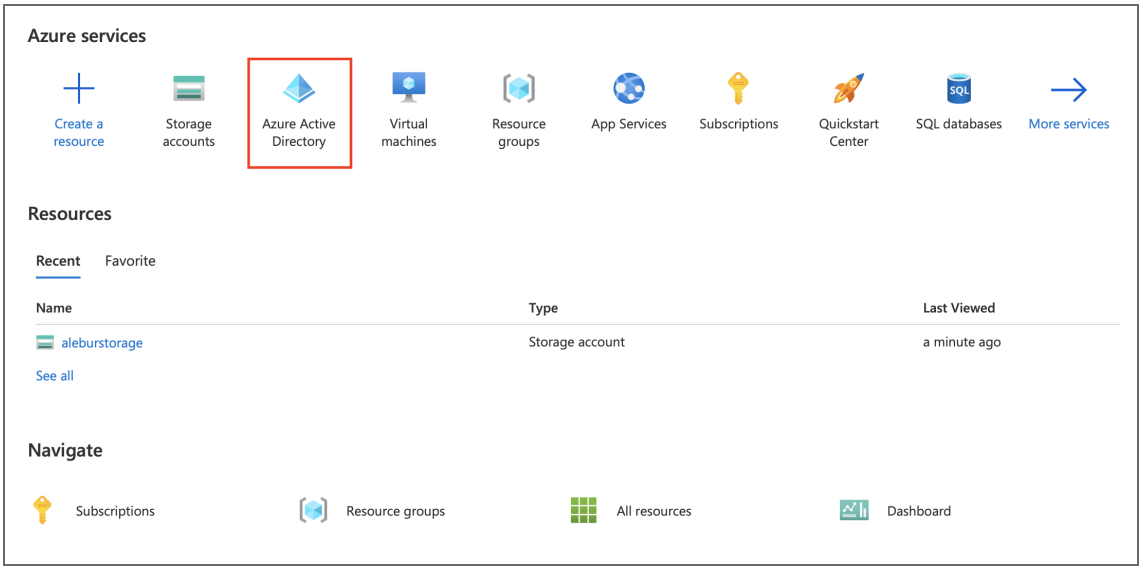

16. Scroll down to find **Enable versioning for blobs** under **Tracking**. Enable this feature and click **Save** below.



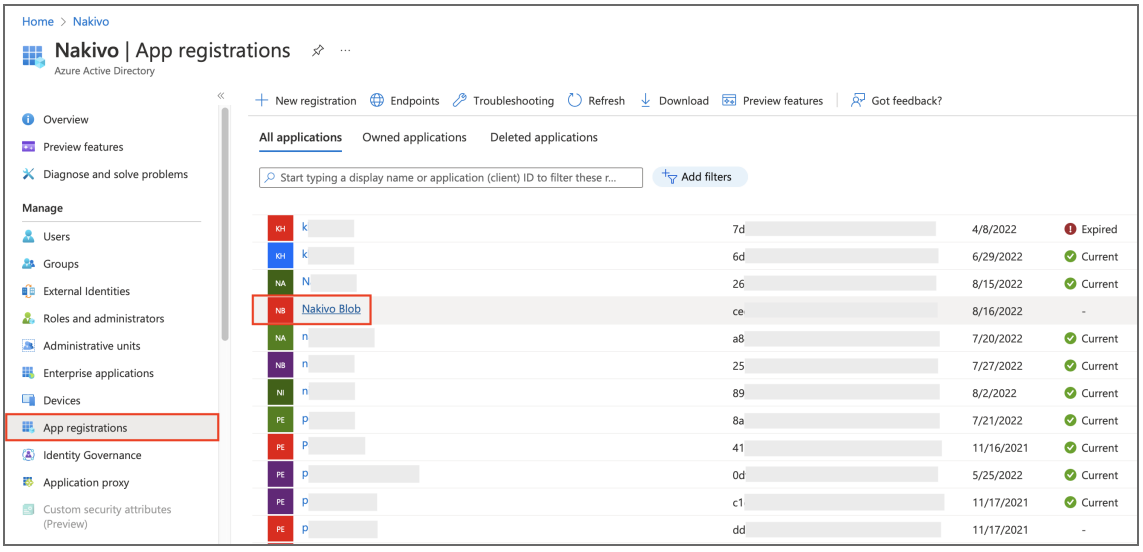
# Obtaining Microsoft Azure Credentials

To obtain the credentials required to add a Microsoft Azure Storage account to the NAKIVO Backup & Replication **Inventory**, follow the steps below.

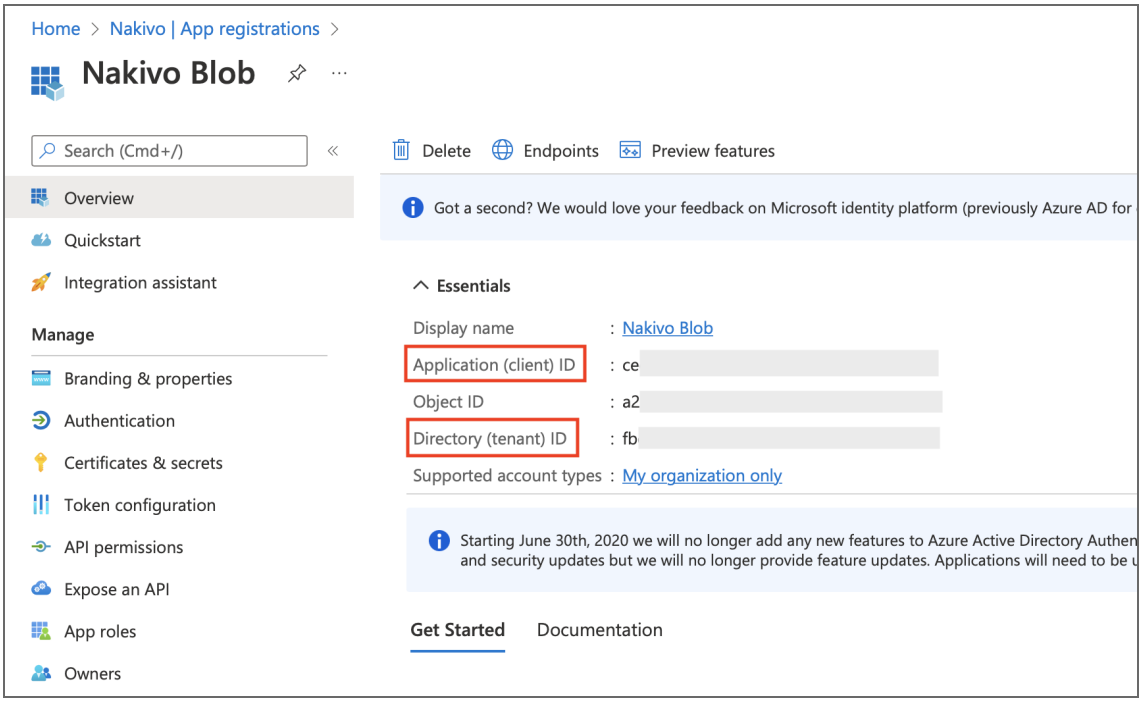
1. Open the Azure Portal by going to portal.azure.com
2. Sign in to Microsoft Azure with your Microsoft account credentials.
3. Select **Azure Active Directory** from the **Dashboard** or from the Portal Menu.



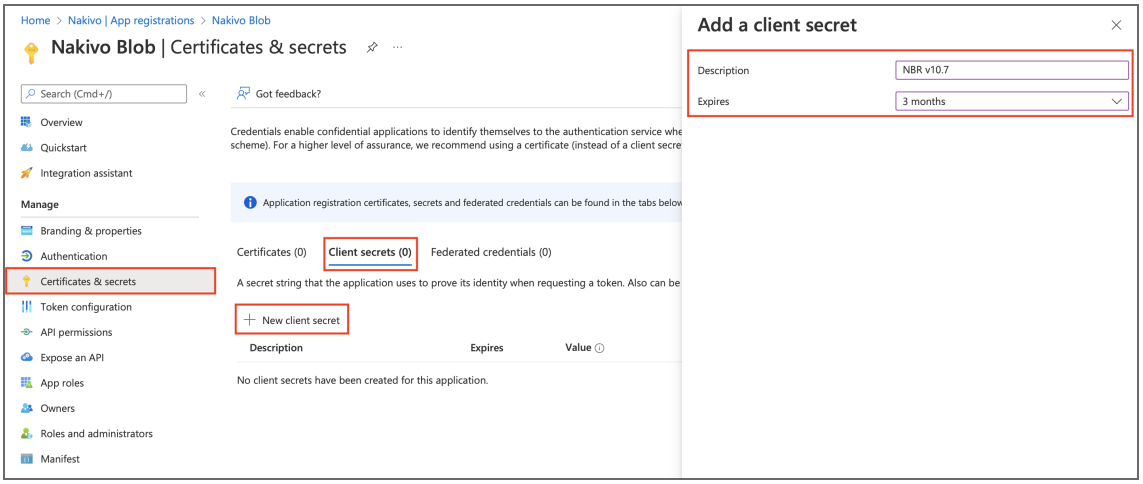

4. In the left menu, click **App registrations** and locate the application registered for use with NAKIVO Backup & Replication. Click on its name to open the application's settings.



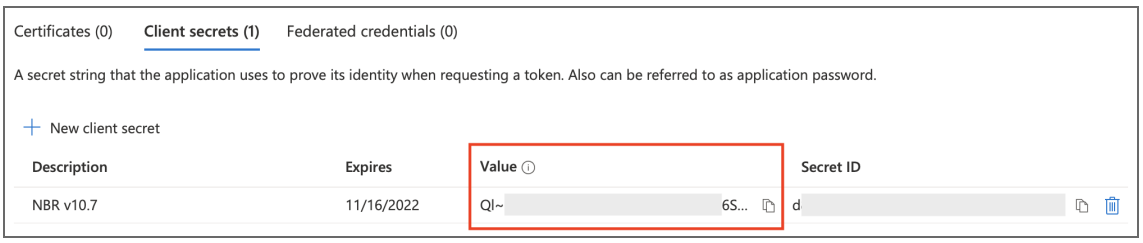

5. Locate and make a note of the Client ID and Tenant ID near the top of the **Overview** menu.




6. Go to the **Certificates & secrets** tab. If you already have a client secret for this application, skip this portion. Otherwise, generate a new client secret for the application by clicking **New client secret** in the **Client secrets** tab. Set a description and expiration period for your client secret and click **Add** below.



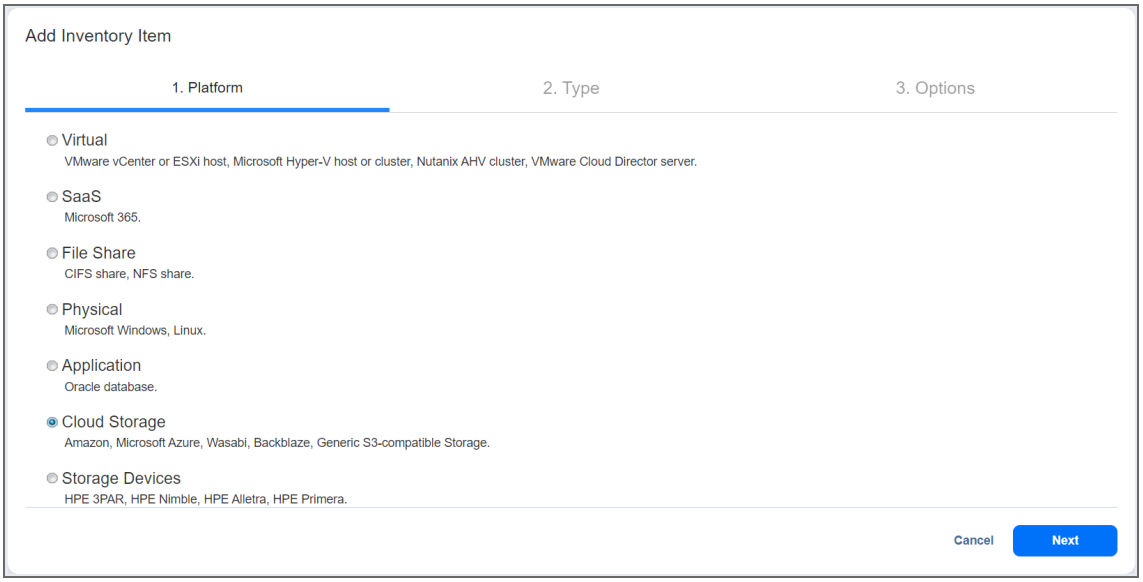

7. Find your newly generated client secret in the **Client secrets** tab in the **Value** column. Store the client secret in a reliable location, as you will have to generate a new one if you lose it.



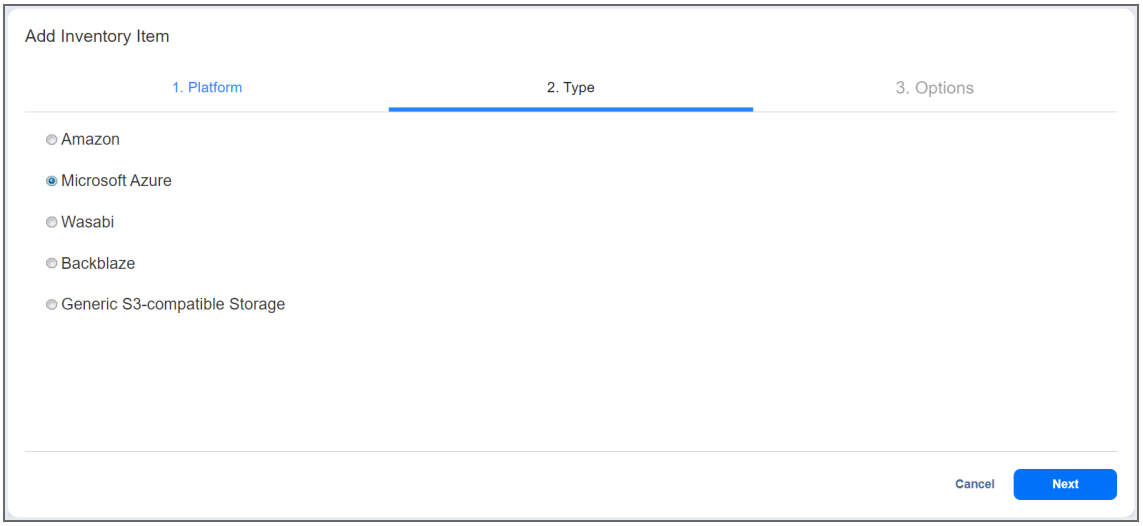
# Adding Microsoft Azure Storage Account to Inventory

To add a Microsoft Azure Storage account to the NAKIVO Backup & Replication **Inventory**, follow the steps below.

1. Click **Settings** in the left pane.
2. Go to the **Inventory** tab and click **+**.
3. On the **Platform** page of the wizard, select **Cloud Storage**. Click **Next** to proceed.




4. On the **Type** page, choose **Microsoft Azure**. Click **Next** to proceed.




5. On the **Options** page, configure the following:
   - **Display name**: Enter a desired Inventory display name for the Microsoft Azure Storage account.
   - **Storage account**: Enter the name of the storage account created in the Azure portal.

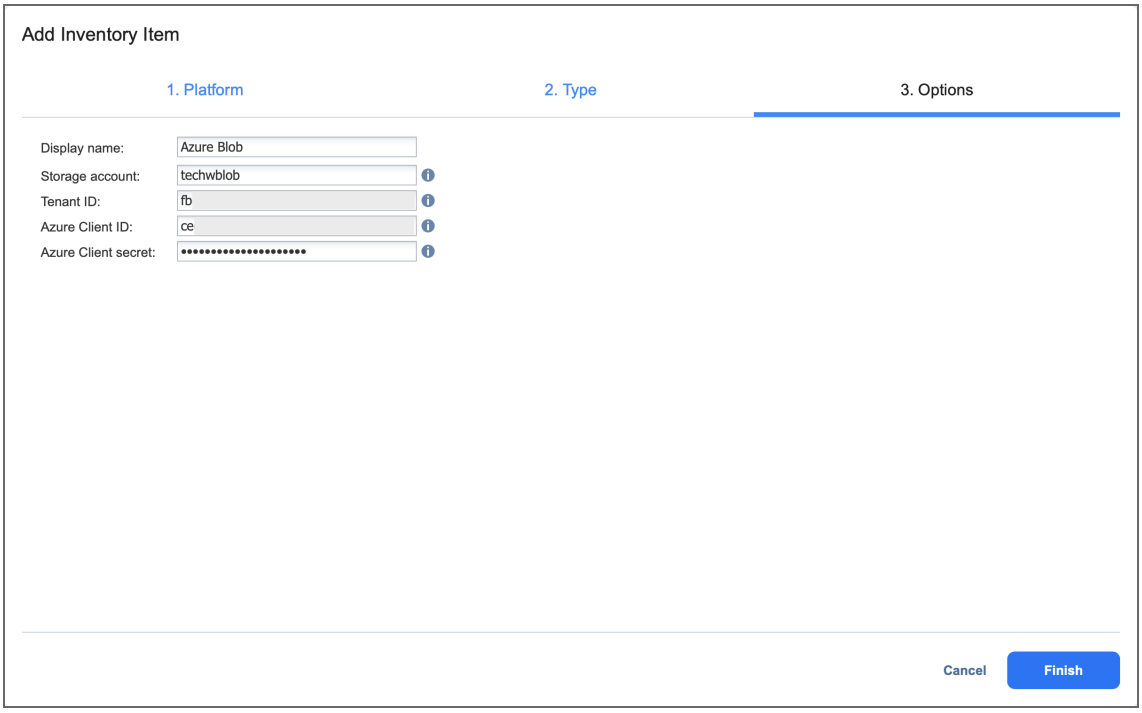
**Notes**

- The following Microsoft Azure storage account types are supported:
  - **General-purpose V2**: Blob storage (block blob, page blob)
  - **Premium Block blobs**: Blob storage (block blob only)
- Newly added General Purpose V2 accounts utilize block blob type instead of page blob type.
- Existing Azure storage accounts added in previous releases may still be used.

- **Tenant ID**: Enter the Azure Tenant ID created when registering your Microsoft Azure account in the Azure Portal.
- **Azure Client ID**: Enter the Azure Client ID created when registering your Microsoft Azure account in the Azure Portal
- **Azure Client Secret**: Enter the Azure Client Secret obtained in the Azure Portal. For more information on obtaining Azure credentials, refer to the Obtaining Microsoft Azure Credentials section above.

**Note**

In order to add a Microsoft Azure Storage account to NAKIVO Backup & Replication, the account must be registered in Azure Active Directory. In addition, NAKIVO Backup & Replication must be assigned an appropriate role within Azure's access control. See the Configuring a Microsoft Azure Storage Account section above for more details.

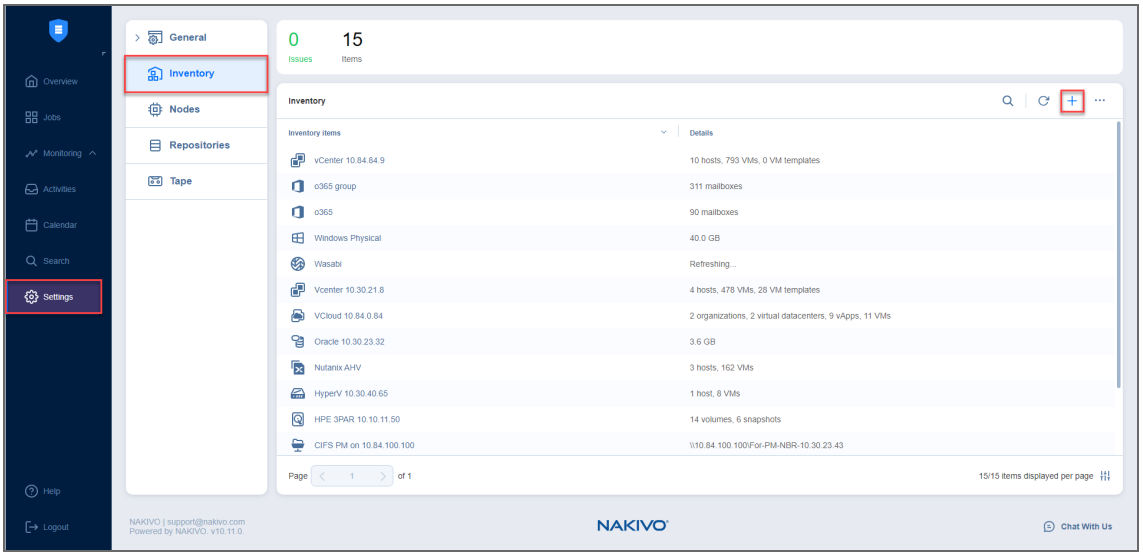


6. Click **Finish** to add the account to the **Inventory**.

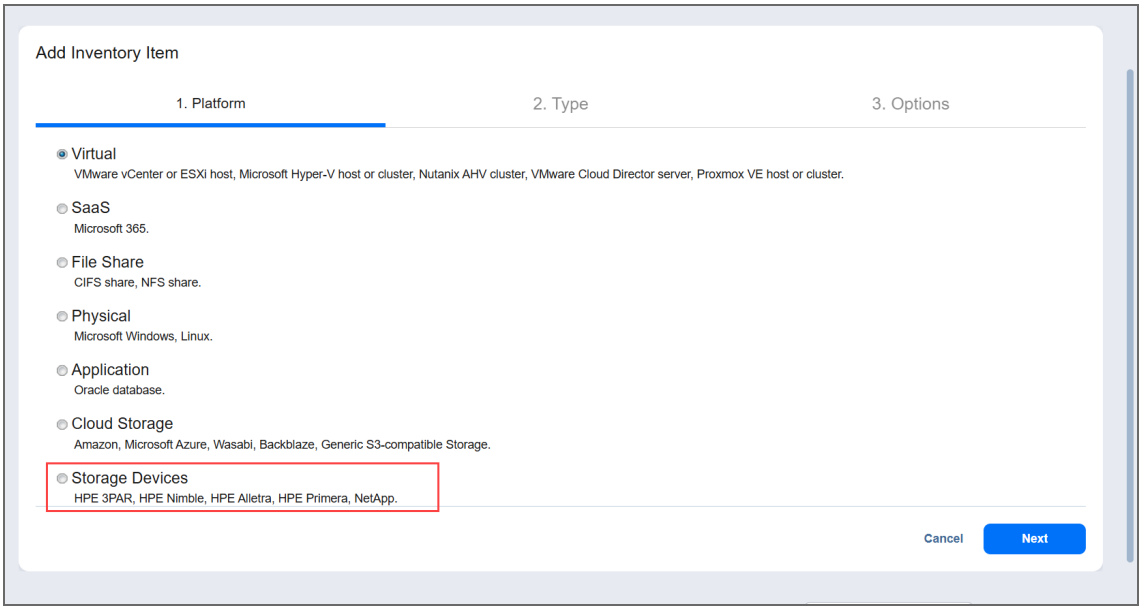
# Adding HPE 3PAR Storage Devices

To add an HPE 3PAR storage device to the **Inventory** in NAKIVO Backup & Replication, follow the steps below:

1. Verify that the HPE 3PAR device meets the integration requirements.
2. Click **Settings** in the main menu of NAKIVO Backup & Replication.
3. Go to the **Inventory** page and click **+**.

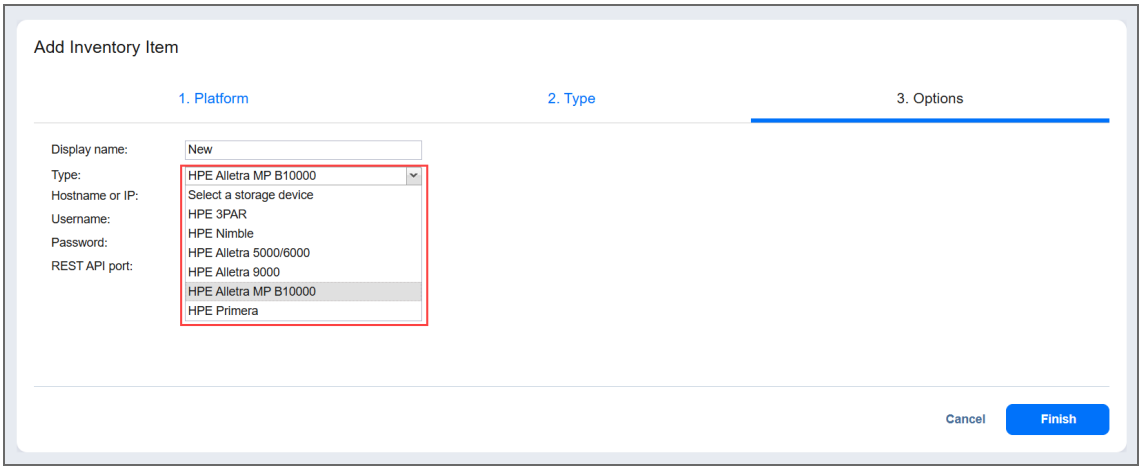



4. On the **Platform** page of the wizard, select **Storage Devices** and click **Next** to proceed.




5. On the **Type** page of the wizard, keep **Hewlett Packard Enterprise (HPE)** selected and click **Next**.
6. On the **Options** page of the wizard, provide the following information:

- **Display name**: Specify a name for the HPE 3PAR storage device. This name will be displayed in the **Inventory**.
- **Type**: Select the storage device type from the dropdown list: *HPE 3PAR*.
- **Hostname or IP**: Specify the hostname or the IP address of the HPE 3PAR device that you want to add to the **Inventory**.
- **Username**: Provide a username for the HPE 3PAR storage device.
- **Password**: Provide a password to the HPE 3PAR storage device.
- **Web services port**: Specify the web services port. By default, *8080* is selected.

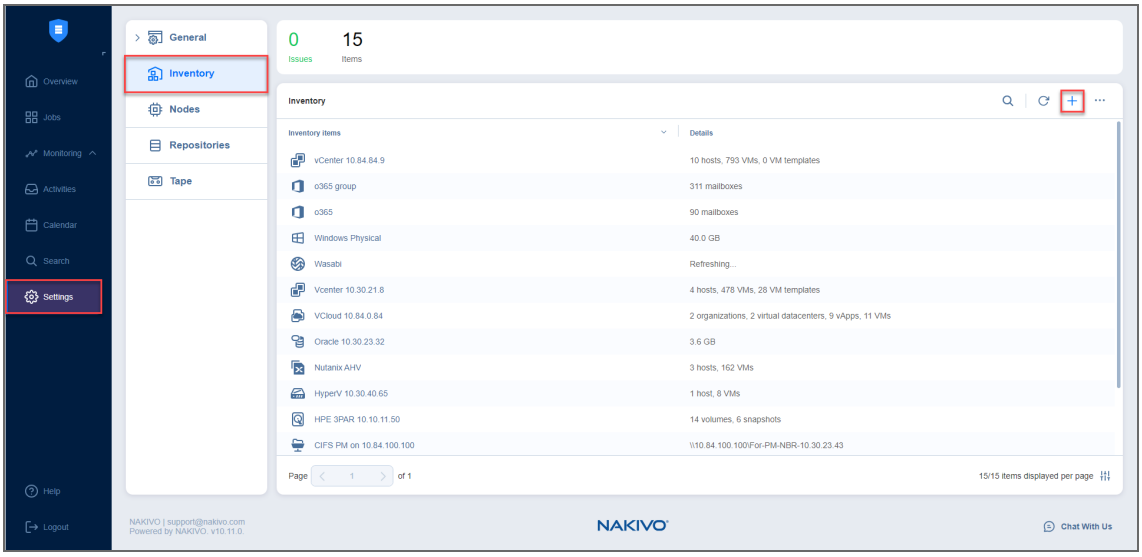



7. Click **Finish** when you are done.

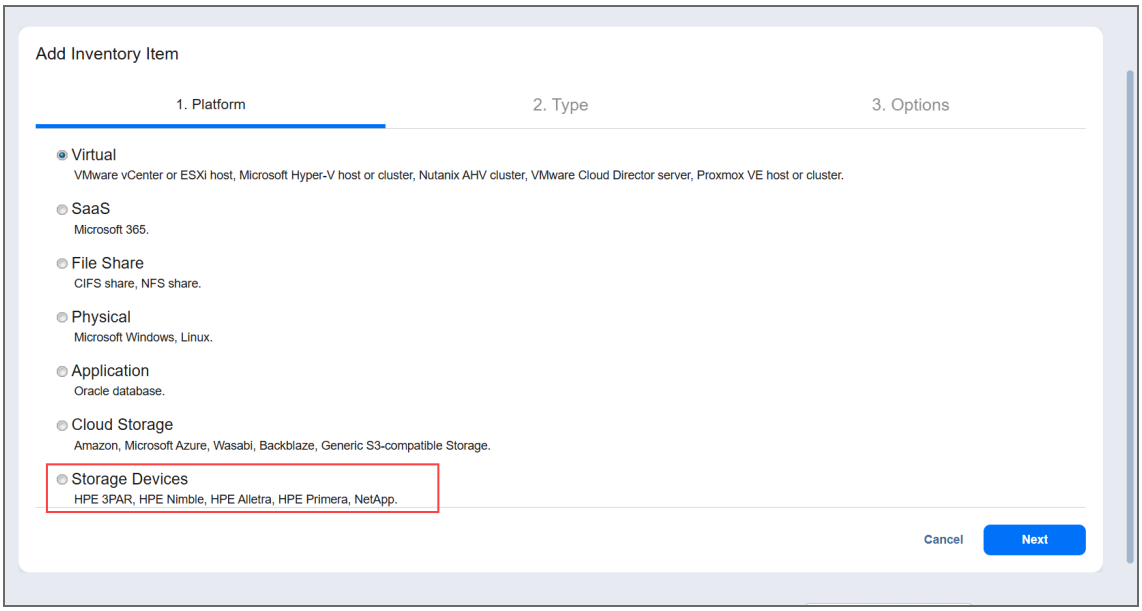
# Adding HPE Nimble Storage Devices

To add an HPE Nimble storage device to the NAKIVO Backup & Replication **Inventory**, follow the steps below:

1. Verify that the HPE Nimble device meets the integration requirements.
2. Click **Settings** in the main menu of NAKIVO Backup & Replication.
3. Go to the **Inventory** tab and click **+**.

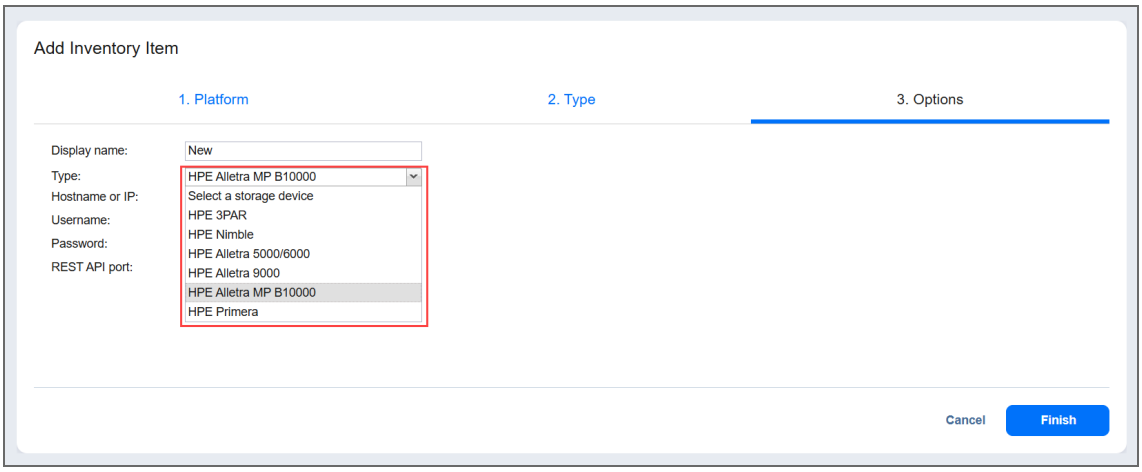



4. On the **Platform** page of the wizard, select **Storage Devices** and click **Next** to proceed.




5. On the **Type** page of the wizard, keep **Hewlett Packard Enterprise (HPE)** selected and click **Next**.
6. On the **Options** page of the wizard, provide the following information:
    - **Display name**: Specify a name for the HPE Nimble storage device. This name will be displayed in the **Inventory**.
    - **Type**: Select the storage device type from the dropdown list: *HPE Nimble*.
    - **Hostname or IP**: Specify the hostname or the IP address of the HPE Nimble device that you want to add to the **Inventory**.
    - **Username**: Provide a username for the HPE Nimble storage device.
    - **Password**: Provide a password to the HPE Nimble storage device.

- **REST API port**: Specify the REST API port.

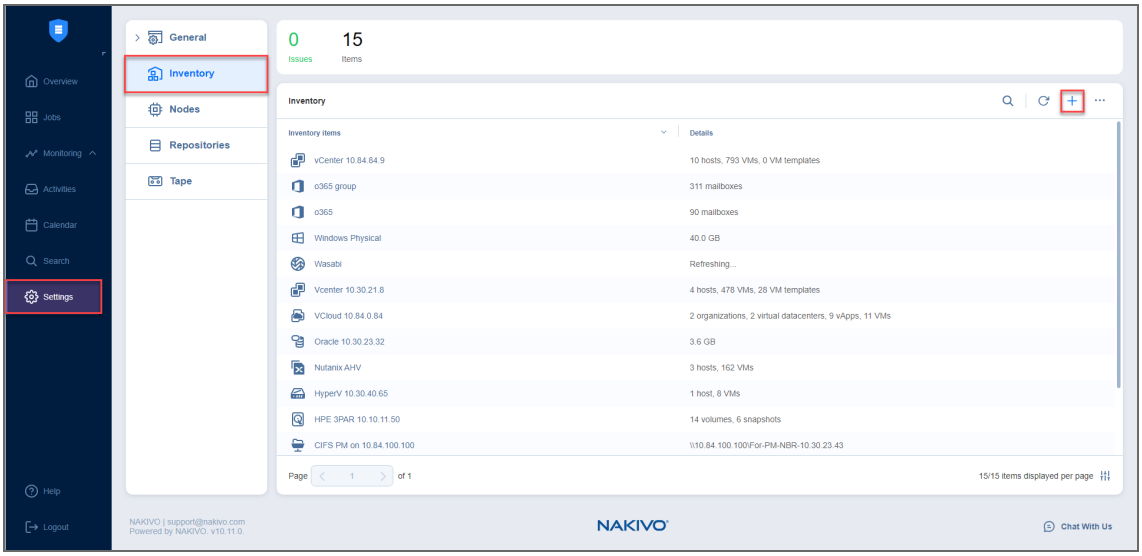



7. Click **Finish** when you are done.

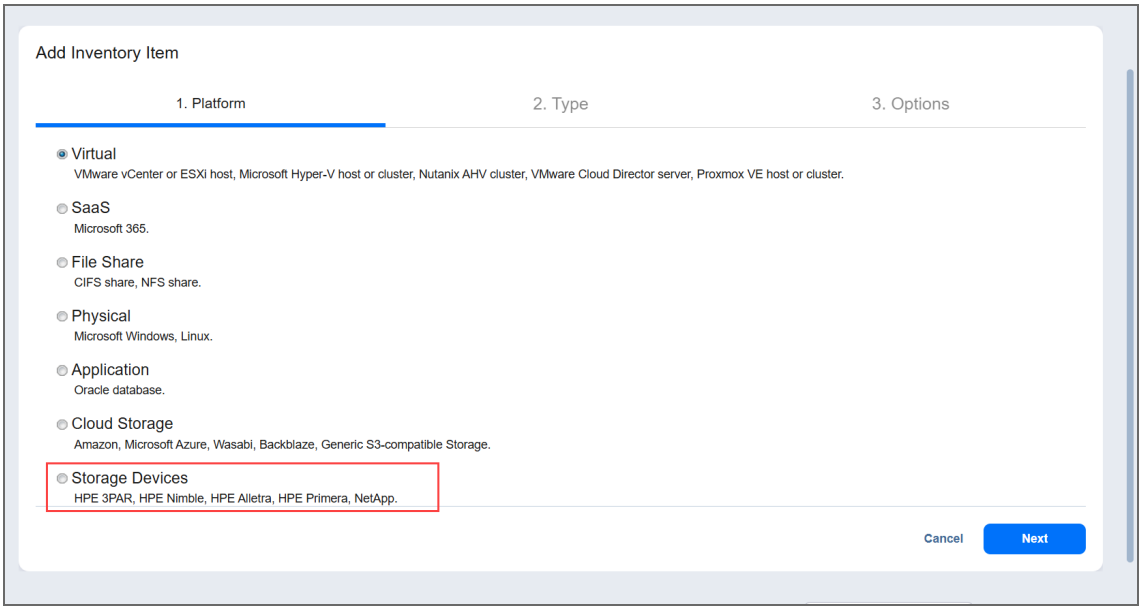
# Adding HPE Primera Storage Devices

To add an HPE Primera storage device to the NAKIVO Backup & Replication **Inventory**, follow the steps below:

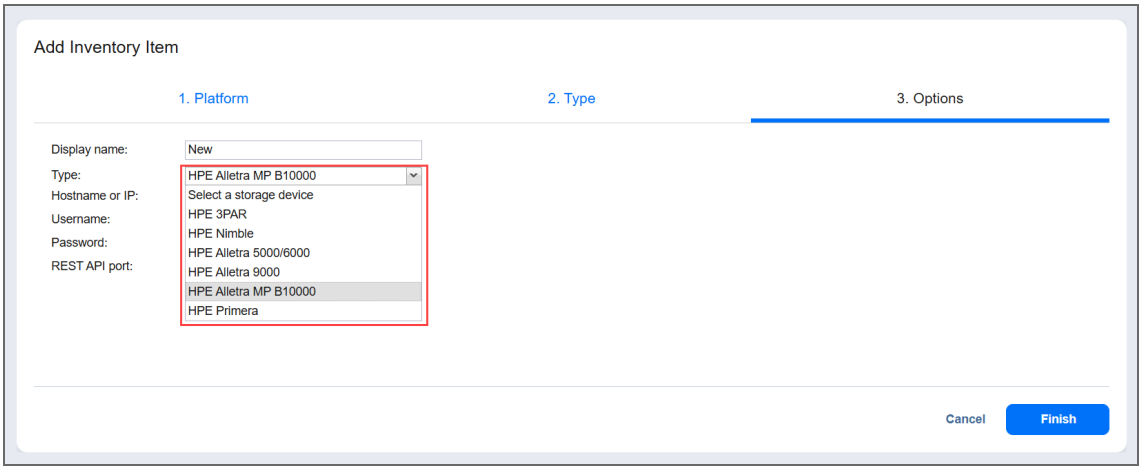
1. Verify that the HPE Primera device meets the integration requirements.
2. Click **Settings** in the main menu of NAKIVO Backup & Replication.
3. Go to the **Inventory** tab and click **+**.




4. On the **Platform** page of the wizard, select **Storage Devices** and click **Next** to proceed.




5. On the **Type** page of the wizard, keep **Hewlett Packard Enterprise (HPE)** selected and click **Next**.
6. On the **Options** page of the wizard, provide the following information:
   - **Display name**: Specify a name for the HPE Primera storage device. This name will be displayed in the **Inventory**.
   - **Type**: Select the storage device type from the dropdown list: *HPE Primera*.
   - **Hostname or IP**: Specify the hostname or the IP address of the HPE Primera device that you want to add to the **Inventory**.
   - **Username**: Provide a username for the HPE Primera storage device.
   - **Password**: Provide a password to the HPE Primera storage device.
   - **Web services port**: Specify the web services port. By default, *443* is selected.

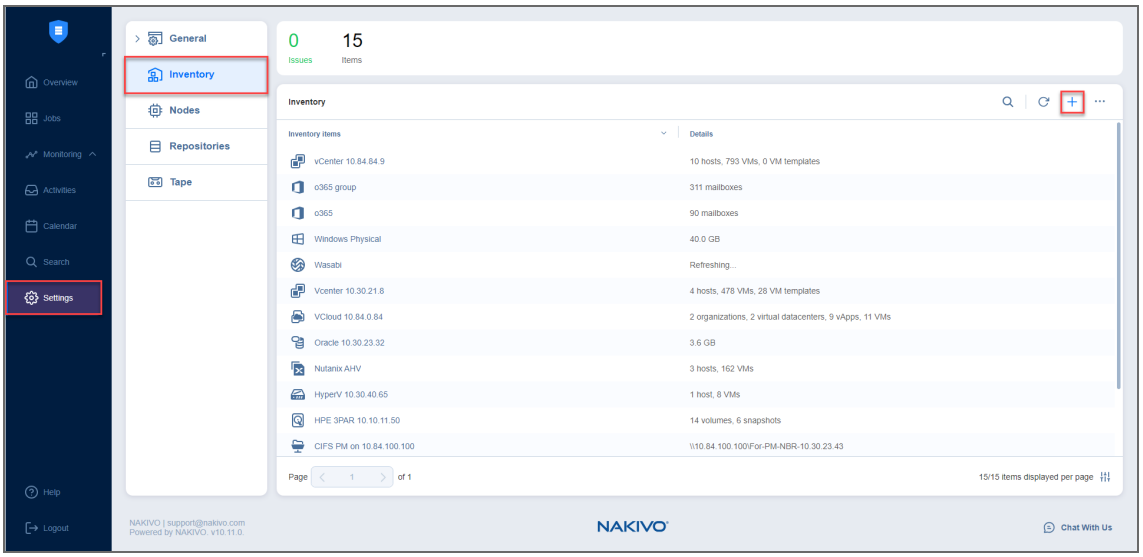



7. Click **Finish** when you are done.

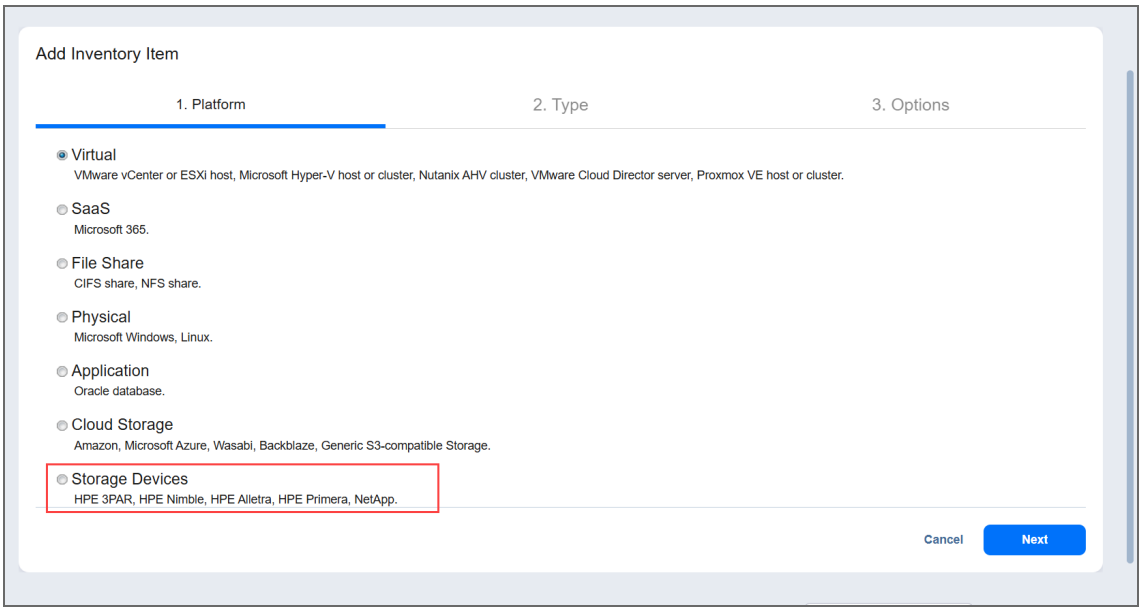
# Adding HPE Alletra Storage Devices

To add an HPE Alletra storage device to the NAKIVO Backup & Replication **Inventory**, follow the steps below:

1. Verify that the HPE Alletra device meets the integration requirements.
2. Click **Settings** in the main menu of NAKIVO Backup & Replication.
3. Go to the **Inventory** tab and click **+**.

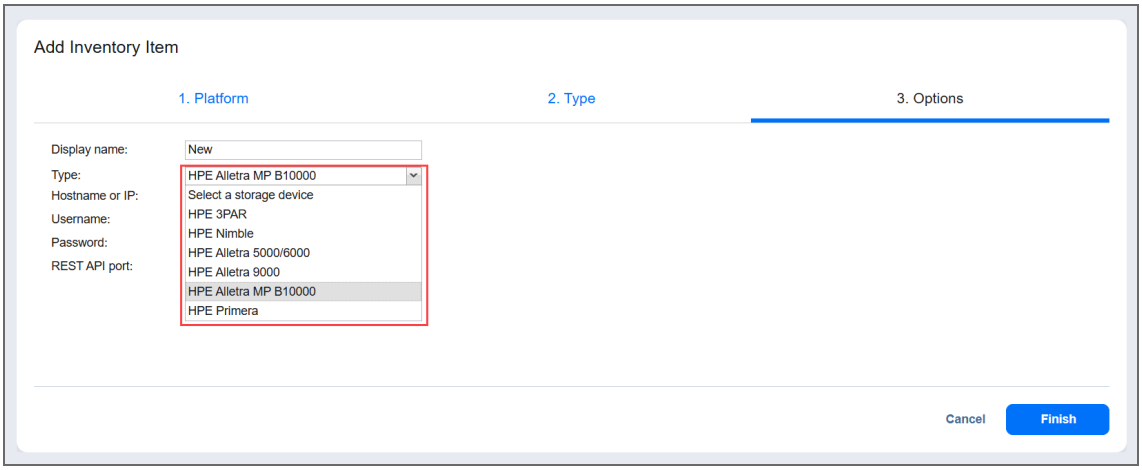



4. On the **Platform** page of the wizard, select **Storage Devices** and click **Next** to proceed.




5. On the **Type** page of the wizard, keep **Hewlett Packard Enterprise (HPE)** selected and click **Next**.
6. On the **Options** page of the wizard, provide the following information:
   - **Display name**: Specify a name for the HPE Alletra storage device. This name will be displayed in the **Inventory**.
   - **Type**: Select the storage device type from the dropdown list.
   - **Hostname or IP**: Specify the hostname or the IP address of HPE Alletra device that you want to add to the **Inventory**.
   - **Username**: Provide a username for the HPE Alletra storage device.

- **Password**: Provide a password to the HPE Alletra storage device.
- **REST API port**: For the *HPE Alletra 5000/6000/MP B10000* storage device type, specify the REST API port.
- **Web services port**: For the *HPE Alletra 9000* storage device type, specify the web services port. By default, *443* is selected.



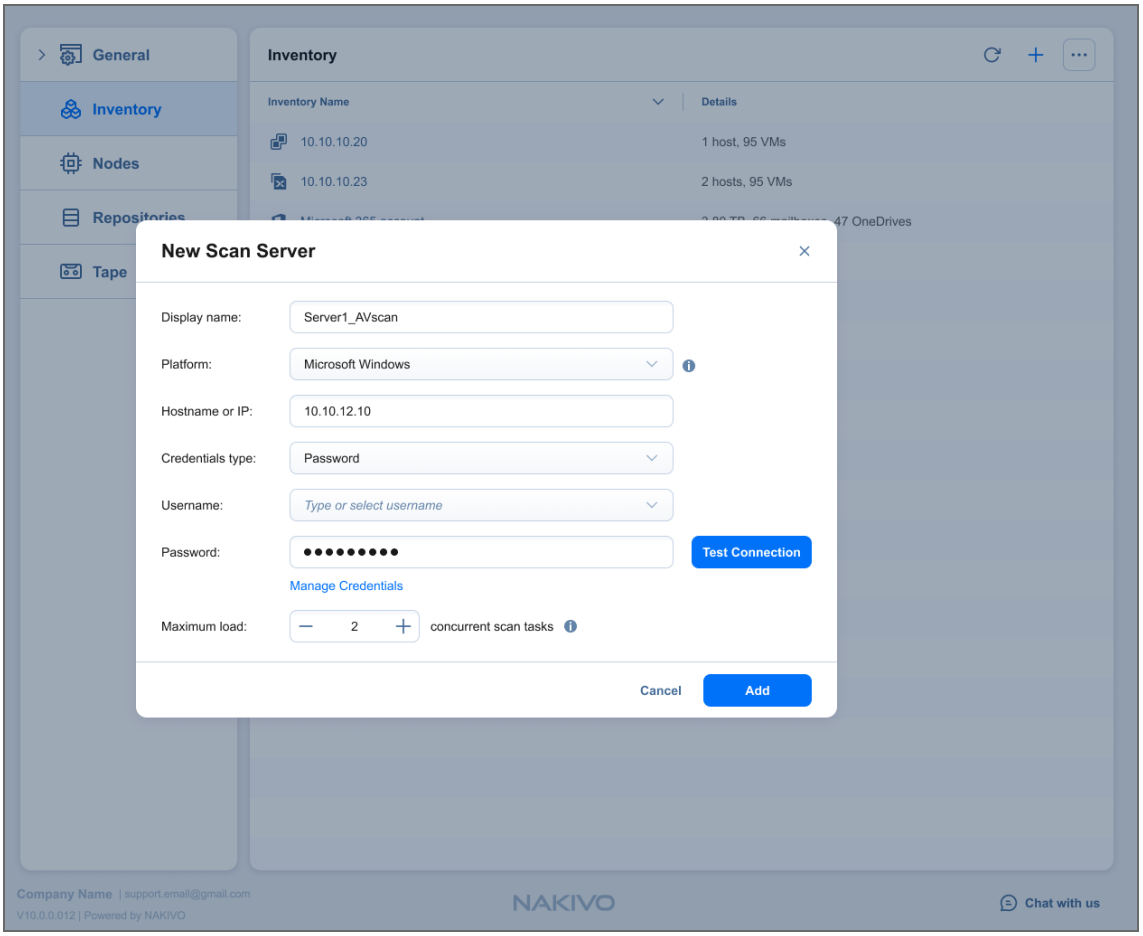

7. Click **Finish** when you are done.

# Adding Scan Servers

To add Scan Servers to the **Inventory**, do the following:

1. Go to **Settings** > **Inventory**.
2. Click the "**...**" button and select **Scan servers**.
3. The **Scan Servers** popup displays a list of added scan servers. Optionally, you can click the "**...**" button to the right of an added scan server and select **Edit** or **Delete** to either edit a scan server's settings or delete it from the list.
4. Click the "**+**" button.
5. In the **New Scan Server** popup, provide the following information:
   - **Display name**: Specify a name for the scan server.
   - **Platform**: Select either **Microsoft Windows** or **Linux**.
   - **Hostname or IP**: Specify the hostname or IP address of the scan server that you want to add to the **Inventory**.
   - **Credentials type**: Choose your preferred option and enter your respective credentials:
     - **Password**: Enter a **Username** with administrative privileges for the scan server entered above and your **Password**.
     - **Private key**: Select your private key from the drop-down list.
6. Optionally, you can add, manage, or delete your credentials using the **Manage credentials** functionality. Refer to Managing Credentials for more information.

7. Click **Test Connection** to make sure that NAKIVO Backup & Replication can successfully connect to your scan server.
8. Configure the **Maximum load** for the scan server, which is the maximum number of concurrent scan tasks the scan server can process.
9. After you're done, click **Add**.



# Managing Inventory

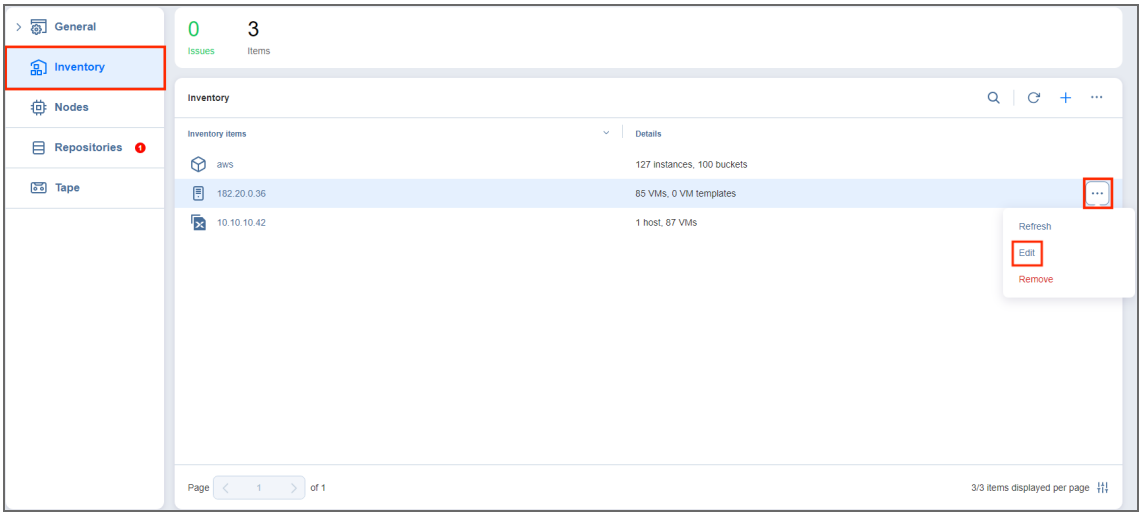
Refer to the following topics:


- “Editing Inventory Items” on page 551
- “Refreshing Inventory” on page 552
- “Removing Items from Inventory” on page 554

# Editing Inventory Items

If the credentials of an inventory item are no longer correct, the connection to the **Inventory** item will be lost. To re-establish a connection, update the required fields in the product by following the steps below:

1. Click **Settings** in the left pane of the product.
2. Go to the **Inventory** tab.
3. Hover over the item you would like to edit.
4. Click ••• on the right side and then click **Edit**.




5. Update the appropriate fields and click **Save**.

# Refreshing Inventory

NAKIVO Backup & Replication keeps the information about the discovered infrastructure in its internal database, which is refreshed every 1 hour by default. During the **Inventory** refresh, the product collects all required information about your virtual infrastructure, such as a list of hosts and VMs, their power state, and so on.

Only one item can be refreshed at a time. If you have added multiple items to the inventory, they will remain in the queue until they are able to be refreshed. Refer to the sections below to learn how to refresh the discovered infrastructure.

- Changing Inventory Refresh Frequency
- Manually Refreshing All Inventory
- Manually Refreshing a Discovered Item

## Changing Inventory Refresh Frequency

1. Click **Settings** in the left pane of the product.
2. Go to the **System setting > Auto refresh** tab.
3. Do either of the following:

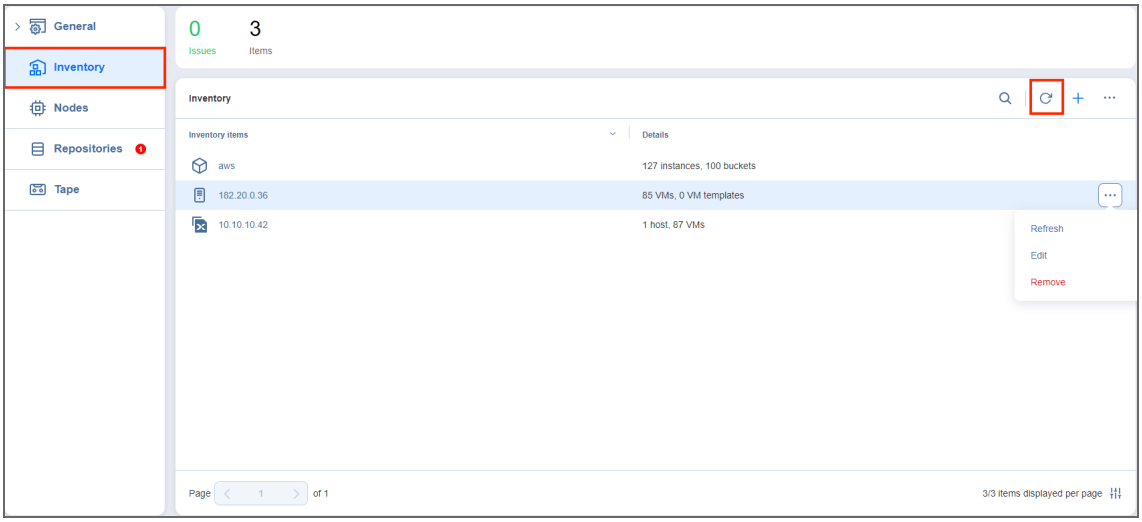
- To prevent the product from automatically refreshing the **Inventory**, deselect the **Refresh inventory every X [time period]** checkbox.
- To change the inventory refresh frequency, enter a new value in the **Refresh inventory every X [time period]** field (from 1 to 60 minutes or from 1 to 24 hours).

> **Note**
> New settings are applied instantly and do not need to be saved.

## Manually Refreshing the Entire Inventory

To refresh all **Inventory** items, follow the steps below:

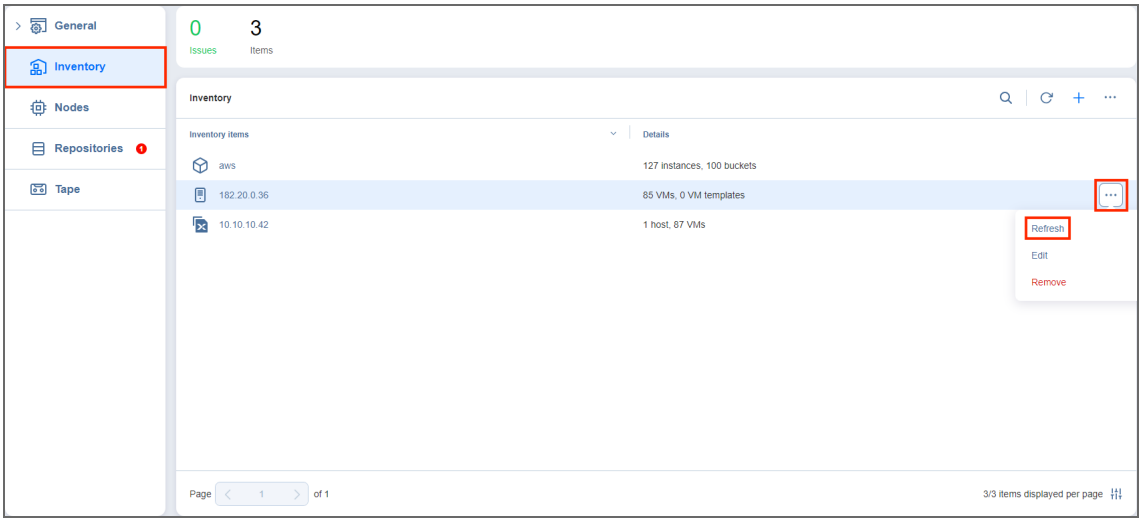
1. Click **Settings** in the left pane of the product and go to the **Inventory** tab.
2. Click the **Refresh All** button.




## Manually Refreshing an Inventory Item

To refresh a single **Inventory** item, follow the steps below:

1. Click **Settings** in the left pane of the product and go to the **Inventory** tab.
2. Click the ••• button next to the item that you would like to refresh.
3. Click **Refresh**.

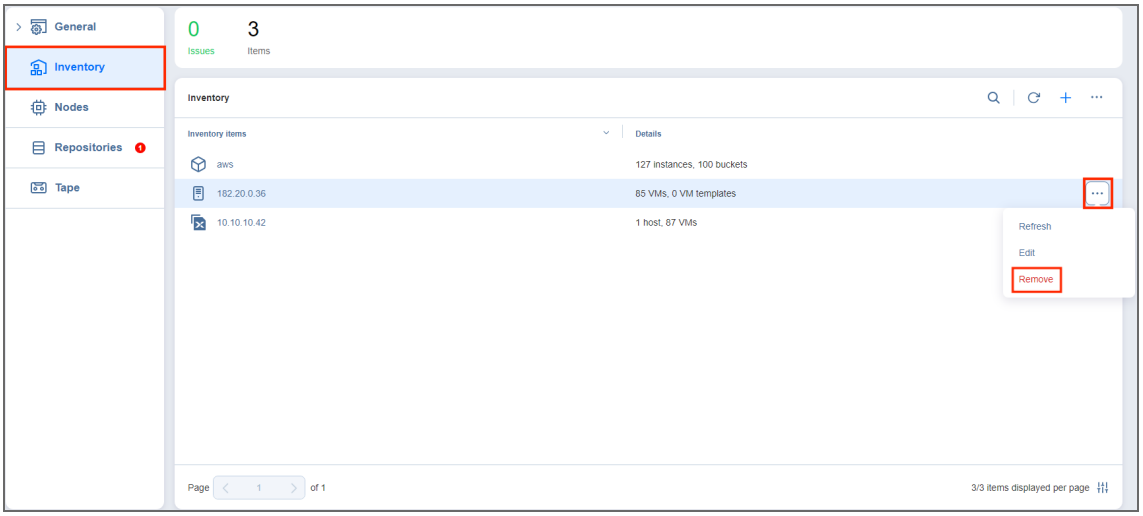


## Removing Items from Inventory

You cannot remove an inventory item if there is at least one backup or replication job that uses the item or its children. In order to remove such items from the **Inventory**, you first need to delete (or edit) the corresponding jobs so no VMs/Instances are backed up or replicated on the host/server/account being removed.

To remove an item from the **Inventory**, follow the steps below:

1. Click **Settings** in the left pane of the product and go to the **Inventory** tab.
2. Hover over the item that you would like to remove from the **Inventory**.
3. Click **Manage** on the right side and click **Remove**.




# Managing Certificates

NAKIVO Backup & Replication provides you with the ability to create and store certificates for Microsoft 365 authentication. Refer to the following topics:

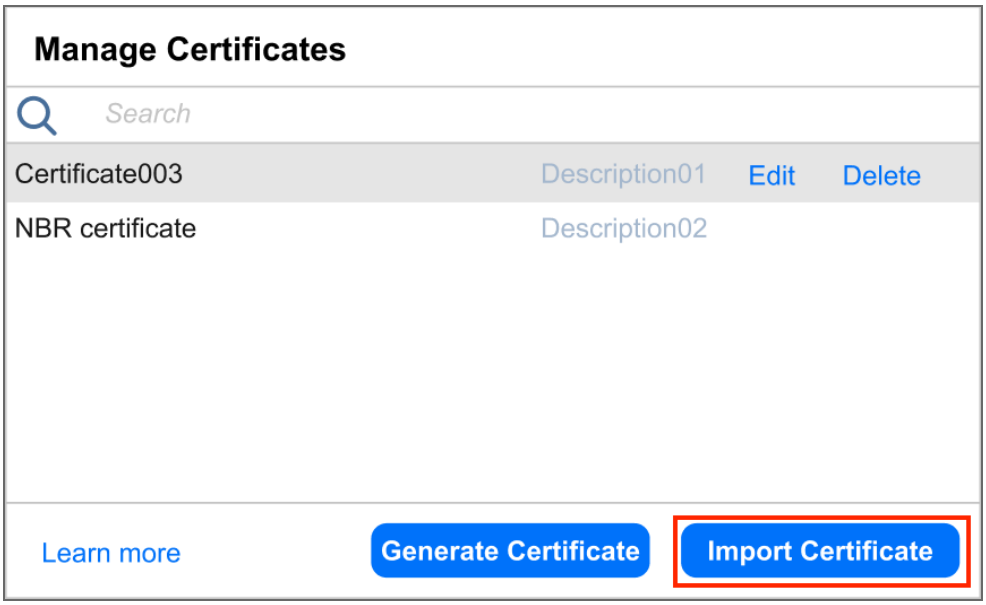
- Importing a Certificate
- Generating a New Certificate
- Editing Certificates
- Deleting Certificates

## Importing a Certificate

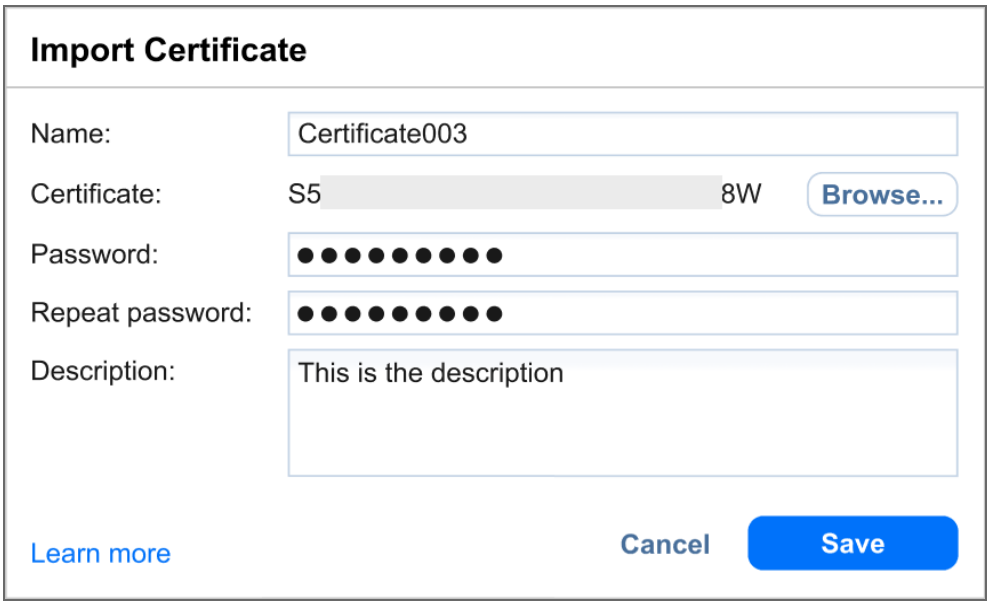
To import an existing certificate, do the following:

1. Click **Settings** in the left pane of the product.
2. Go to the **Inventory** tab.
3. Click **Manage**.
4. In the dialog that opens, click **Certificates**.

5. In the **Manage Certificates** dialog that opens, click **Import Certificate**.




6. In the resulting **Import Certificate** pop-up, enter a display name for the certificate in the **Name** field.
7. In the **Certificate** row, click **Browse…** and upload a certificate in the required `.pfx` format.
8. Enter the certificate password if needed in the **Password** and **Repeat password** fields.
9. Optionally, add a description for the certificate in the **Description** field.
10. Click **Save**.

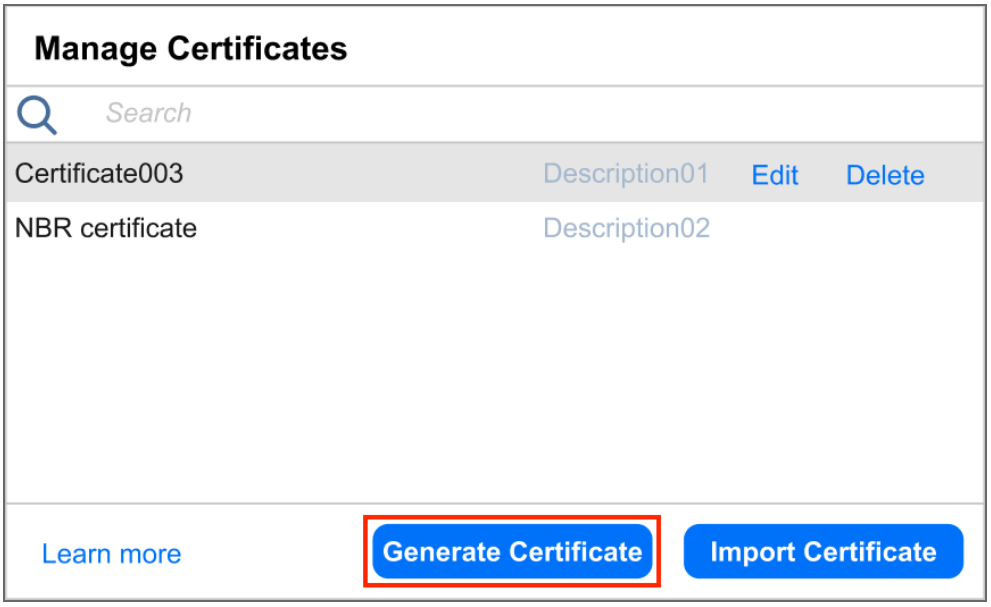



NAKIVO Backup & Replication will check if the certificate is valid. If it is, the certificate will become available for Microsoft 365 authentication.

## Generating a New Certificate

To generate a new self-signed certificate, do the following:

1. Click **Settings** in the left pane of the product.
2. Go to the **Inventory** tab.
3. Click **Manage**.
4. In the dialog that opens, click **Certificates**.
5. In the **Manage Certificates** dialog that opens, click **Generate Certificate**.




6. In the resulting **Generate New Self-Signed Certificate** pop-up, enter a display name for the certificate in the **Name** field.
7. Optionally, add a description for the certificate in the **Description** field.
8. Click **Save** to generate a new self-signed certificate and add it to the list of saved certificates. NAKIVO Backup & Replication will automatically download the new certificate to the browser.

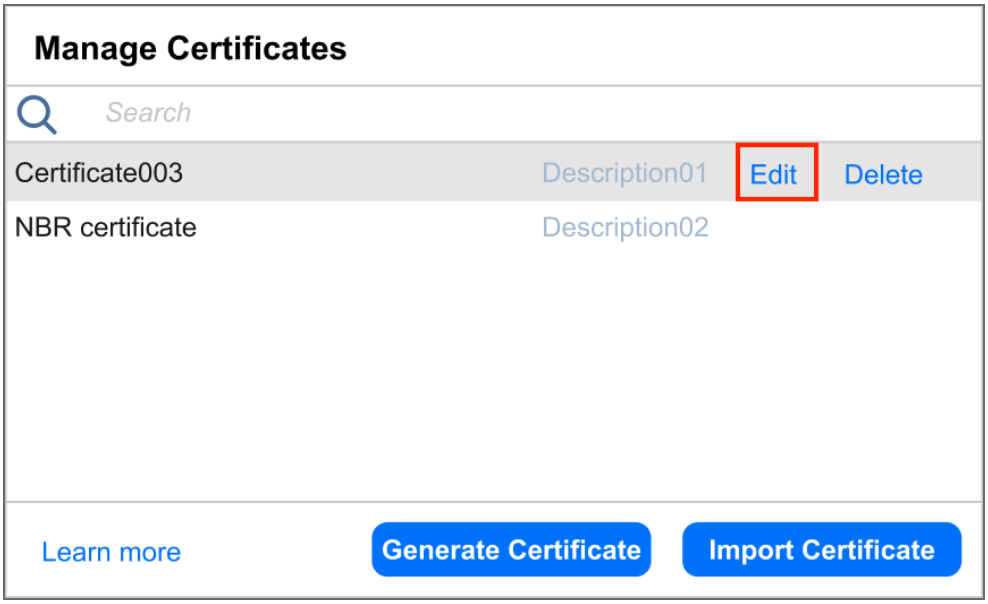



You can now use this certificate for Microsoft 365 authentication.

## Editing Certificates

To edit a certificate, do the following:

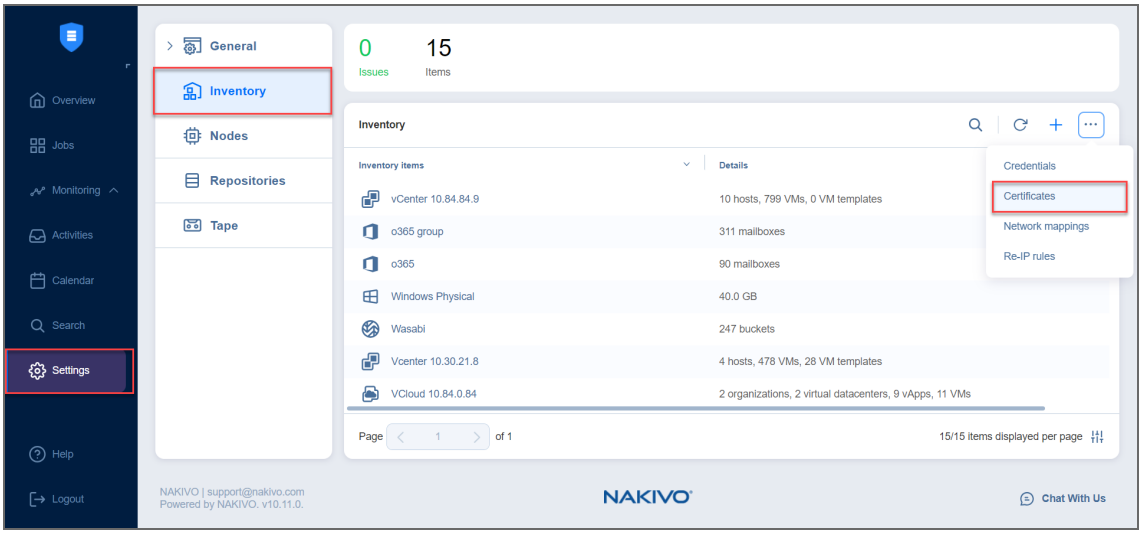
1. Click **Settings** in the left pane of the product.
2. Go to the **Inventory** tab.
3. Click **Manage**.
4. In the dialog that opens, click **Certificates**.
5. In the **Manage Certificates** dialog box that opens, hover over the certificate you want to edit and click **Edit**.
6. Edit any necessary fields and click **Save** when done.



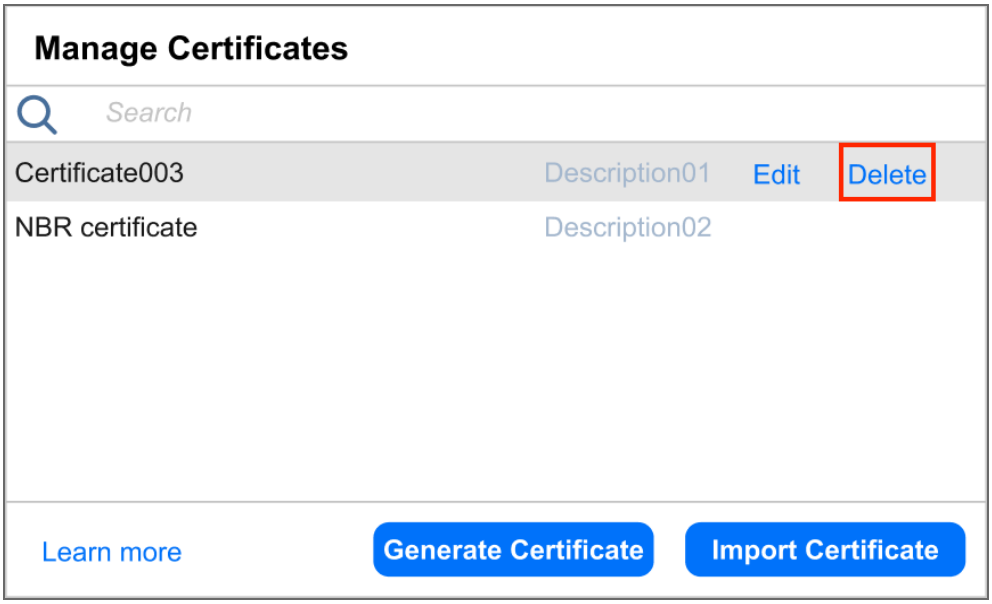

## Deleting Certificates

To delete a certificate, do the following:

1. Click **Settings** in the left pane of the product.
2. Go to the **Inventory** tab.
3. Click **Manage**.
4. In the dialog that opens, click **Certificates**.




5. In the **Manage Certificates** dialog box that opens, hover over the certificate you want to delete and click **Delete**.
6. In the confirmation pop-up that opens, click **Delete** to confirm deletion.



# Managing Credentials

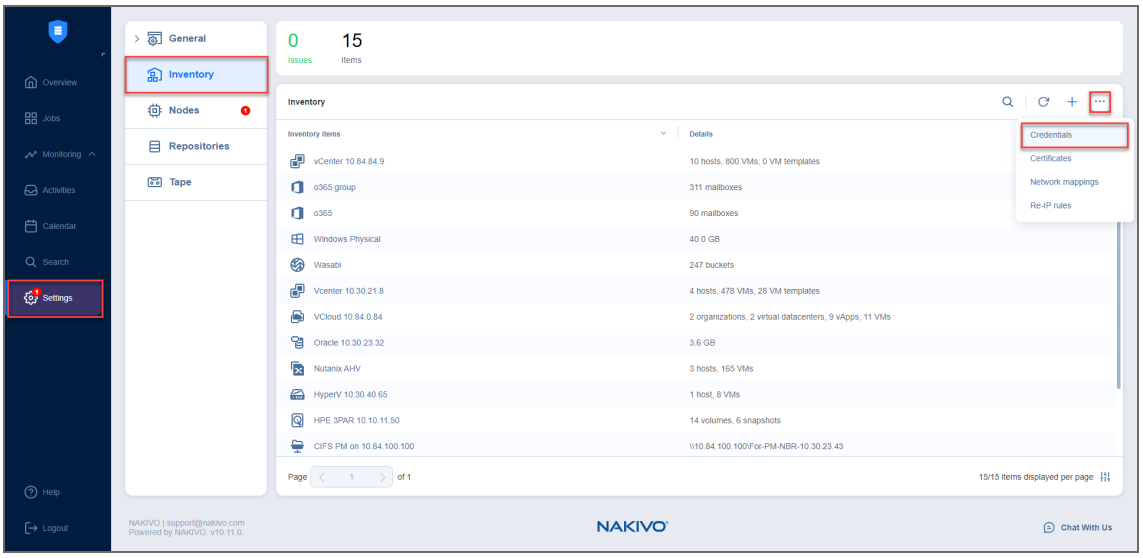
NAKIVO Backup & Replication provides you with the ability to store your OS login and password, Amazon EC2 instance private keys, and SSH keys to your Linux machines. Refer to the following topics:

- Adding Credentials
- Editing Credentials
- Deleting Credentials

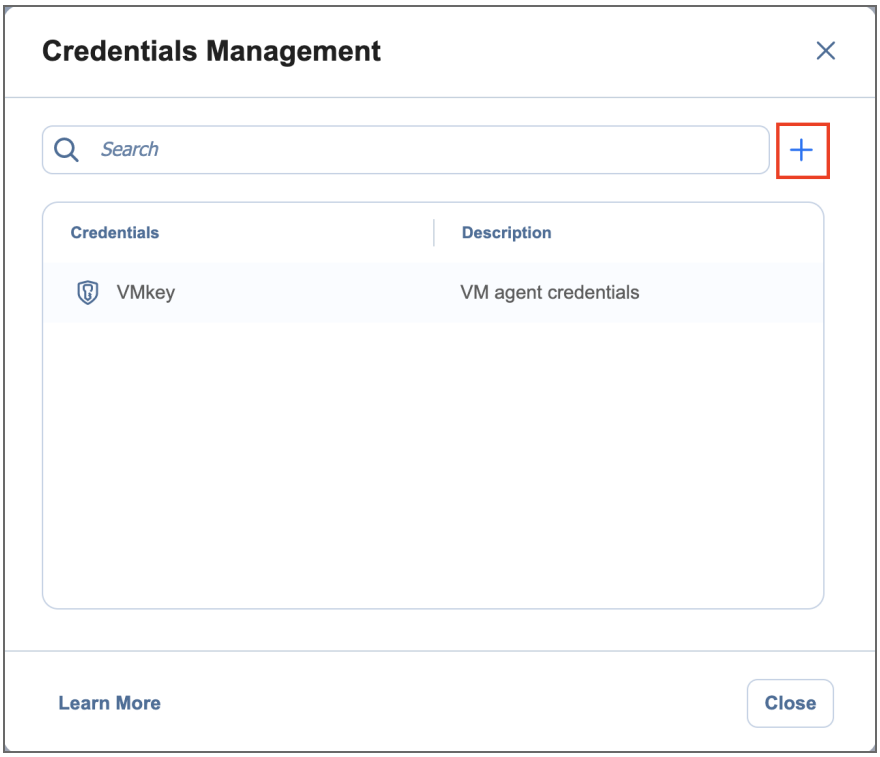
## Adding Credentials

To add new credentials, do the following:

1. Click **Settings** in the left pane of the product.
2. Go to the **Inventory** tab.
3. Click **Manage**.
4. In the dialog box that opens, click **Credentials**.




5. In the **Manage Credentials** dialog box that opens, click **Add Credentials**.




6. Then, do the following:
   - **Type**: Select the type of credentials:
     - To set up a basic username and password, fill out the **Username**, **Password**, and (optionally) **Description** fields and click **Save**.
     - To set up a master password, select **Master password** from the drop-down list and fill out the **Name**, **Password**, and (optionally) **Description** fields and click Save.
     - To add a private key to an Amazon EC2 instance or a Linux physical machine, do the following:
       a. **Private key**: Select **Private Key** from the **Type** menu.
       b. **Username**: Enter a username for the private key.
       c. **Password**: Create a password for the private key.
       d. **Repeat password**: Repeat password.

          > **Note**
          > If you generated your key with a passphrase, you have to enter this passphrase into the **Password** and **Repeat password** boxes.

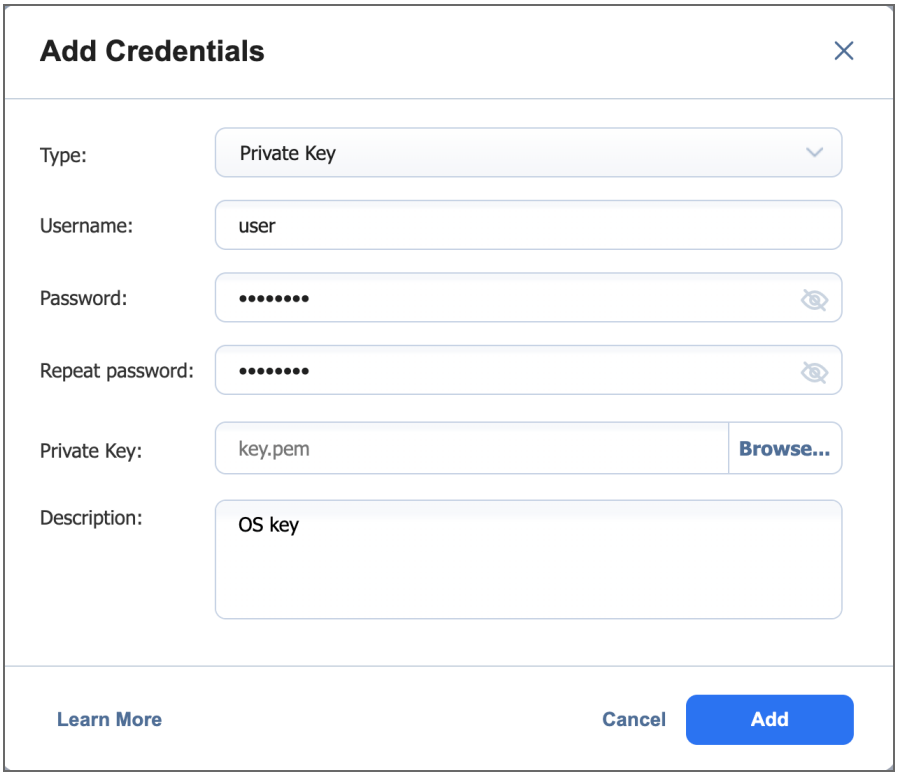
       e. Locate and select the private key.

**Notes**

- Supported key formats: RSA, *DSA*
- By default, newer versions of *ssh-keygen* generate keys with the unsupported -----*BEGIN OPENSSH PRIVATE KEY*----- format. To generate a key with the -----*BEGIN RSA PRIVATE KEY*----- format, include *-m PEM* in your *ssh-keygen* command.
- Supported file extensions: no extension, `.pem`, `.key`, `.cer`, `.der`, `.txt`

f. Fill out the **Description** box.

g. Click **Save**.

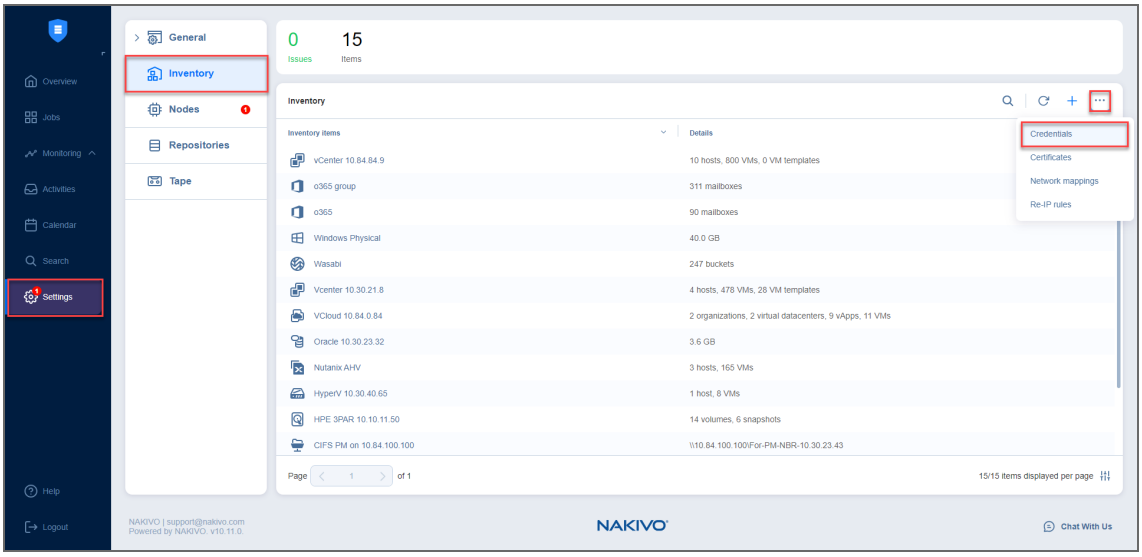



You can now assign the credentials while creating jobs or setting up VM agents.

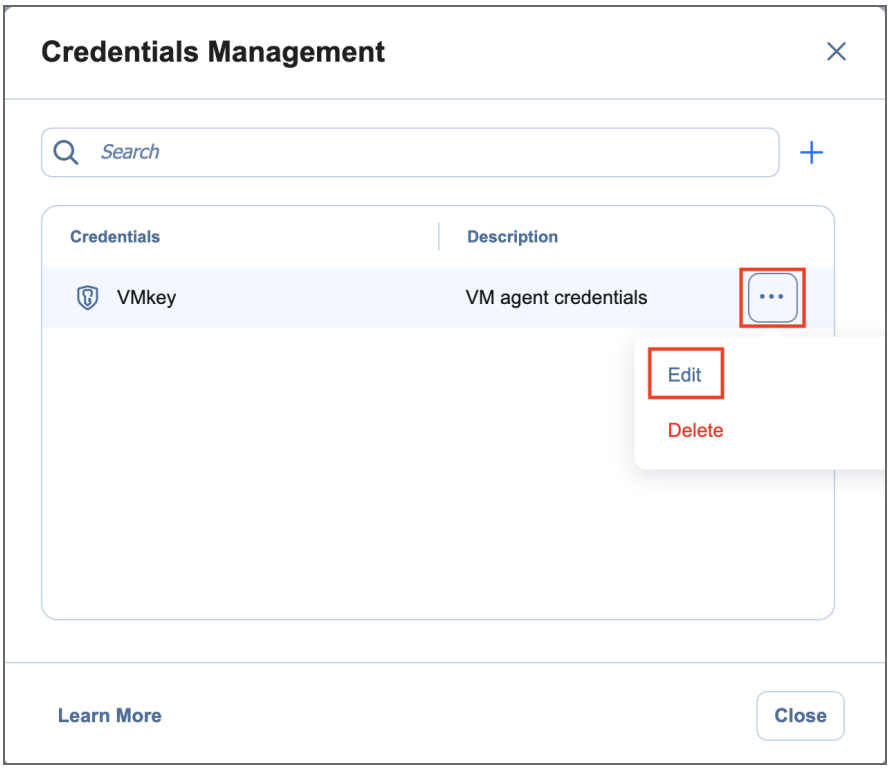
## Editing Credentials

To edit credentials, do the following:

1. Click **Settings** in the left pane of the product.
2. Go to the **Inventory** tab.
3. Click **Manage**.
4. In the dialog that opens, click **Credentials**.




5. Hover over the record that you would like to edit and click **Manage > Edit**.

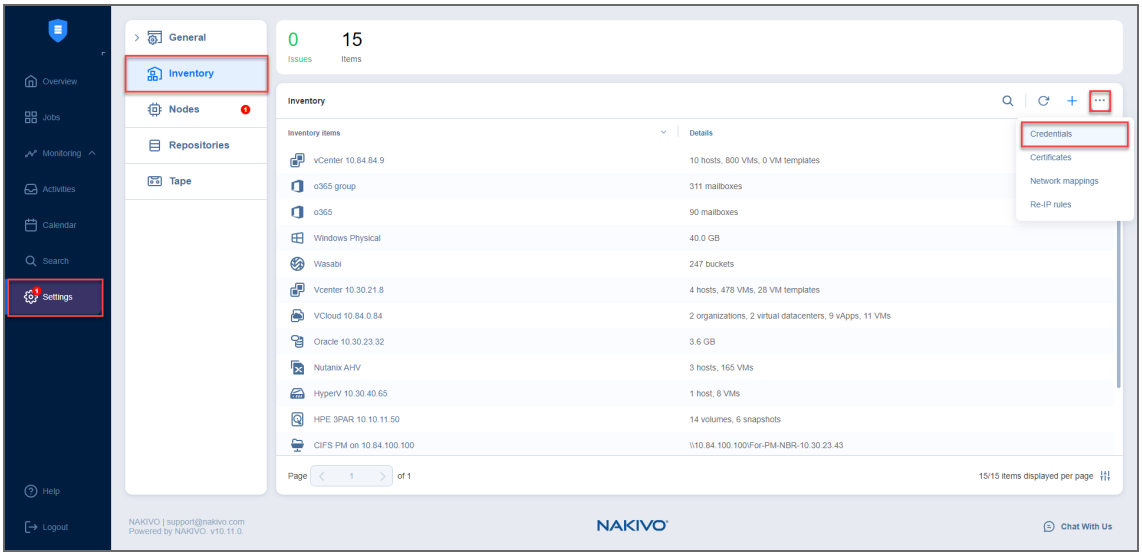



6. Make any required changes, and then click **Save**.

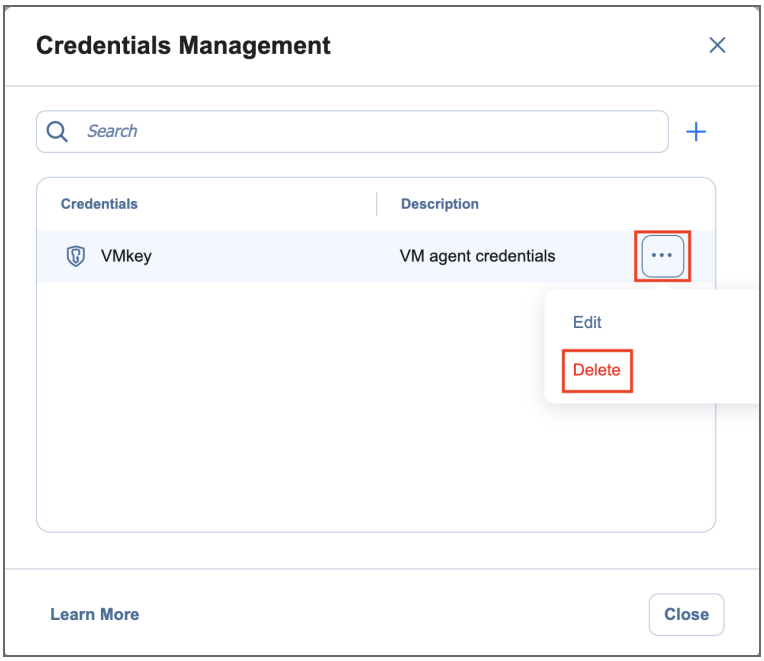
# Deleting Credentials

To delete credentials, do the following:

1. Click **Settings** in the left pane of the product.
2. Go to the **Inventory** tab.
3. Click **Manage**.
4. In the dialog that opens, click **Credentials**.




5. Hover over the record that you would like to delete and click **Manage > Delete**.

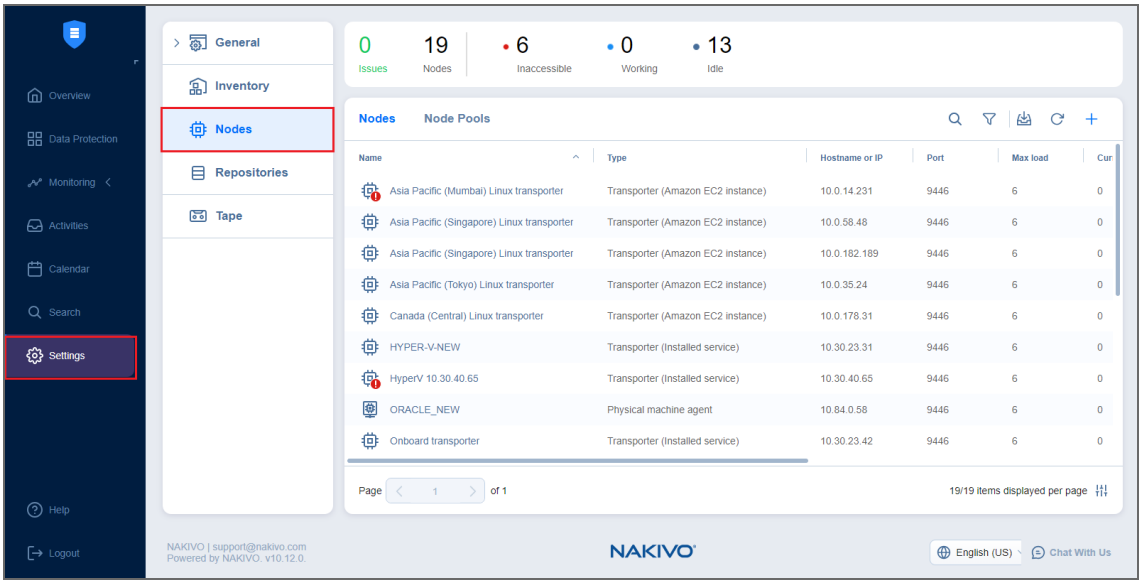



6. Click **Delete** in the confirmation dialog box that opens.

# Nodes

Nodes are an essential component of NAKIVO Backup & Replication. They include **Transporters**, VM Agents, and Physical Machine Agents. The **Transporter**, for example, performs backup, replication, and recovery, as well as data compression, deduplication, and encryption. The **Nodes** tab contains a **Summary** bar, which offers an overview of all nodes. The data displayed is as follows:

- **Issues**: Total number of issues/alarms related to nodes
- **Nodes**: Total number of nodes
- **Inaccessible**: Total number of inaccessible nodes
- **Working**: Total number of working nodes
- **Idle**: Total number of idle nodes




To learn how to add nodes and manage them, refer to the topics below:

- “Configuring Nodes” below
- “Managing Nodes” on page 593

## Configuring Nodes

Refer to the following topics:

- “Adding Existing Nodes” on the next page
- “Deploying Transporter as Nutanix AHV Appliance” on page 574
- “Deploying Transporter as VMware Appliance” on page 575
- “Deploying Transporters in Amazon EC2” on page 578
- “Installing a VM Agent” on page 582

- “Installing I/O Filter” on page 583
- “Installing Journal Service” on page 590

# Adding Existing Nodes

After you have installed a **Transporter** or Agent, you need to add it to NAKIVO Backup & Replication so that the **Transporter** or Agent can be used for backup, replication, and recovery tasks.

**Important**

Before adding the existing **Transporter** to your NAKIVO Backup & Replication, make sure that this **Transporter** is not used by any other **Director** as it may lead to unforeseen errors.

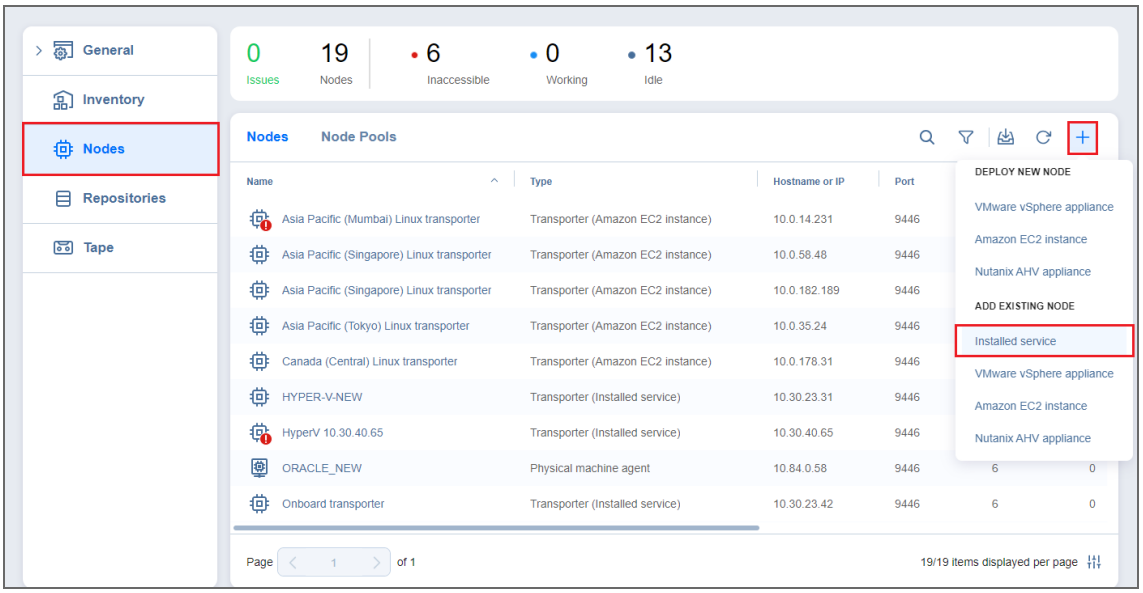
Refer to the following topics:

- Installed Service
- VMware Appliance
- Amazon EC2 Instance
- Nutanix AHV Appliance

## Installed Service

Follow the steps below to add a node that is installed as a service:

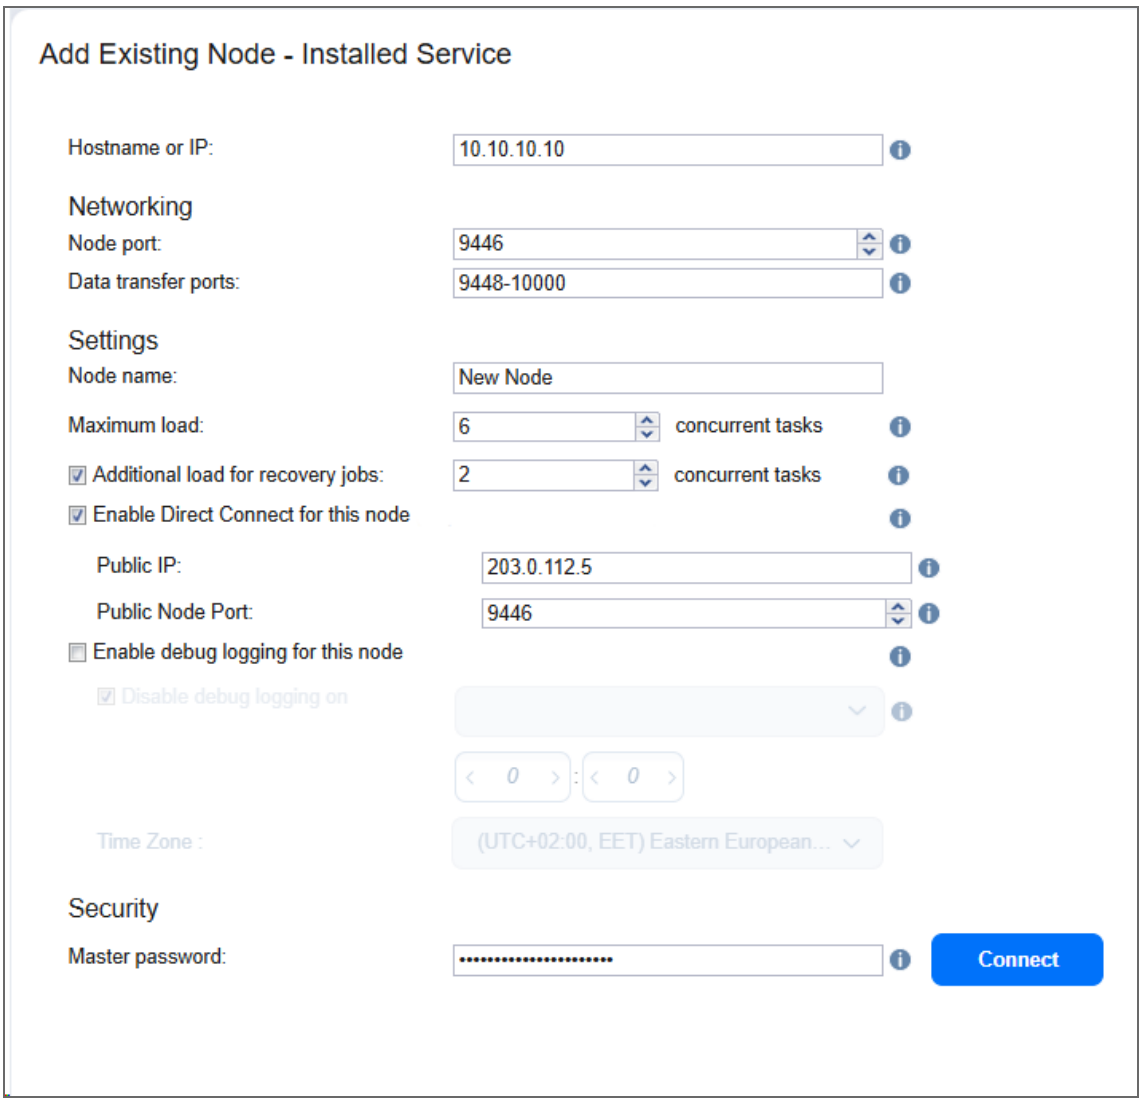
1. Click **Settings** in the left pane of the product, go to the **Nodes** tab, and click **+**.
2. In the **Add Existing Node** section, click **Installed service**.




- The **Add Existing Node - Installed Service** menu opens. In the **Hostname or IP** box, enter the IP address or hostname of the machine on which the node is installed.

**Note**

If you are adding the node by a DNS name, make sure this DNS name can be resolved on the machines on which the **Director** and any other nodes (which you plan to use in conjunction with the current one) are installed.




3. Click **More options...** to reveal and edit the following fields:
    - In the *Networking* section:
        - **Node port**: Specify the port number that will be used to connect to the node.
        - **Data transfer ports**: Specify a range of port numbers (from 1 to 65535) that will be used to transfer data. The range you specify should contain at least 100 ports. Make sure that the ports you specify are open in your firewall.

- In the *Settings* section:
  - **Node name**: Specify a display name for the node.
  - **Maximum load**: Specify the maximum number of tasks that the node should process simultaneously. A task, for example, is a backup or replication of a single VM disk, or one granular recovery session.
  - **Additional load for recovery jobs**: If selected, the specified amount of tasks will be added to set maximum node load to be used for recovery jobs exclusively. This allows running the specified amount of concurrent recovery jobs along with other types of jobs without the need to wait for their completion.
  - **Enable Direct Connect for this node**: When this option is enabled, you can access remote resources via a single port connection without establishing a VPN connection. For more information refer to Enabling Direct Connect.
    - **Public IP**: The IP address of external **MSP Transporter** that is reachable from the internet and used for communication with the **Direct Connect Transporter**.
    - **Public Node Port**: The port of external **MSP Transporter** that is reachable from the internet and used for communication with the **Direct Connect Transporter**. Ensure that the entered combination of the **Public Node Port** and **Public IP** address is not used by other transporters.

    The following conditions must be met at the remote infrastructure to enable this feature:
    - A NAKIVO **Transporter** or Agent must be installed.
    - A master password must be set for security reasons. A pre-shared key is generated based on the entered master password.
    - The node port on the local machine must be exposed to external availability via the Internet.
    - The **Public IP** and **Public Node Port** fields are available only with the license that supports MSP Direct Connect.
  - **Enable debug logging for this node**: If needed, enable debug level logging for the current node. It is not recommended to use this option on a permanent basis.
- In the *Security* section:
  - **Master Password**: Optionally, you can set a password to secure the connection. The set password must match the one configured on the **Transporter** or Agent. Note that setting a master password is required when the **Enable Direct Connect for this node** option is enabled. Proceed as follows:
    a. After entering the password, click **Connect**.
    b. The **Certificate Details** dialog box appears. Verify the certificate details, and click **Accept**.

**Notes**

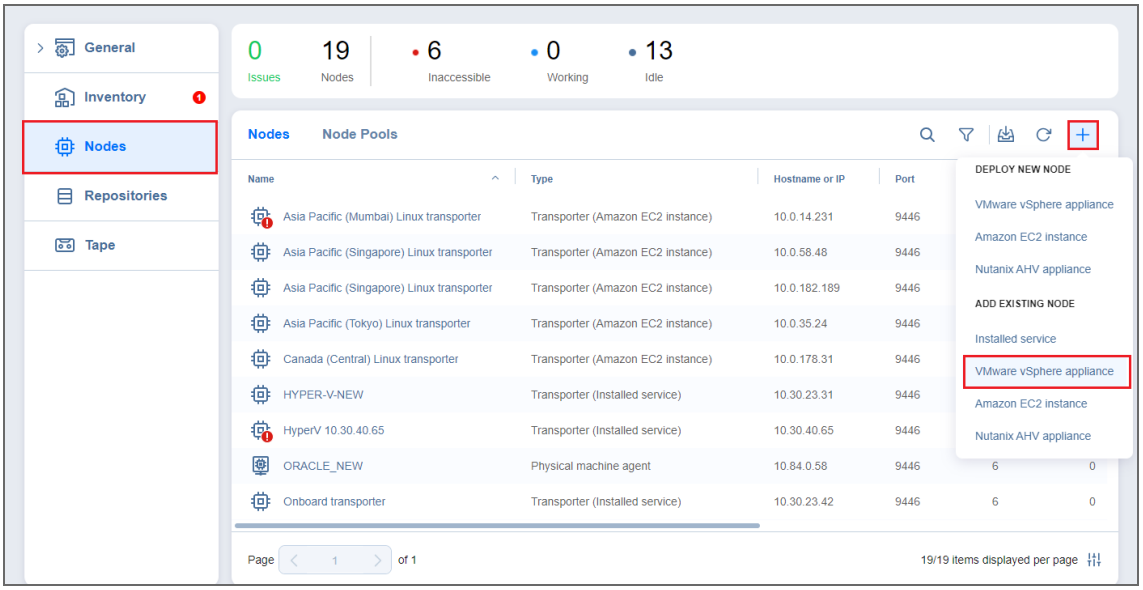
- The master password must adhere to the following requirements:
  - Minimal length - 5 characters.
  - Maximum length - 50 characters.
- The master password can be set and re-set manually by running the command on the machine housing the **Transporter** or Agent. Follow these steps:
  - Enter the following command `bhsvc -b password`, replacing "password" with your master password.
  - Restart the **Transporter** or Agent.

4. Click **Add**. The node is added to the product and can be used for backup, replication, and recovery jobs.

## VMware Appliance

Follow the steps below to add a **Transporter** that is deployed as a VMware appliance:

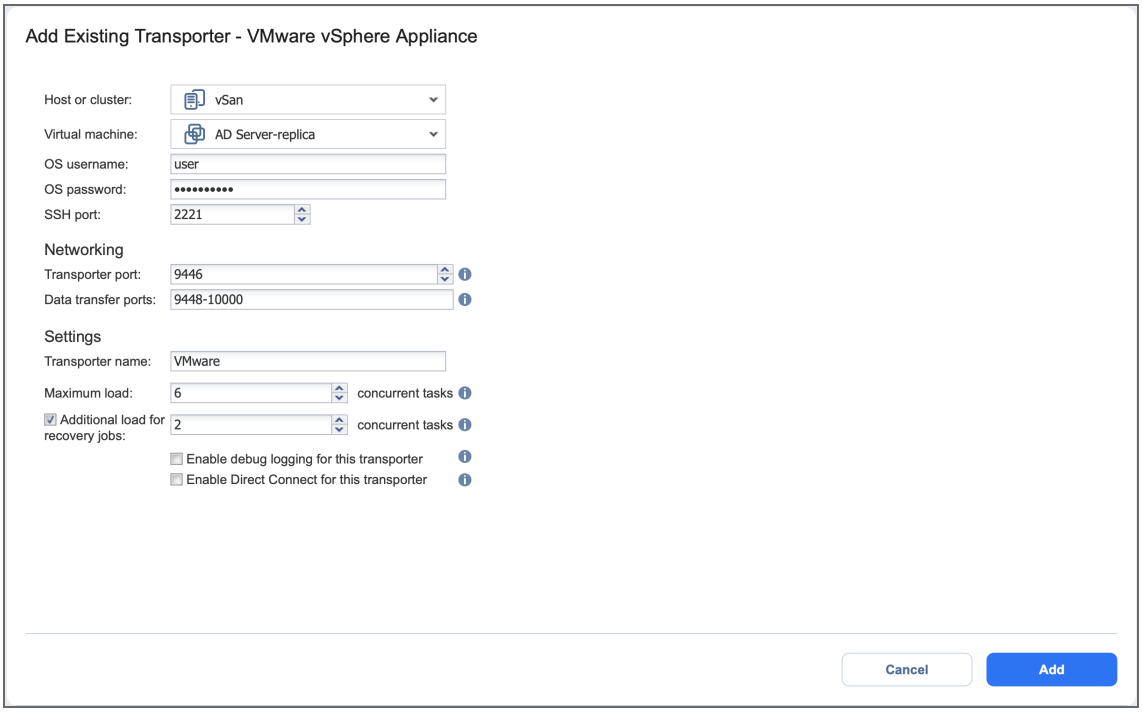
1. Click **Settings** in the left pane of the product, go to the **Nodes** tab, and click **+**.
2. In the **Add Existing Node** section, click **VMware vSphere appliance**.




3. The **Add Existing Transporter - VMware vSphere Appliance** dialog opens. Fill out the fields as described below:
   - In the **Host or cluster** box, specify the location of the host or cluster where the corresponding virtual machine is deployed.
   - In the **Virtual machine** box, specify the virtual machine on which the **Transporter** is installed.

- In the **OS Username** and **OS Password** fields, specify credentials for accessing the virtual machine.
- In the **SSH port** box, enter the SSH port if needed.
- Click **More options** to reveal and edit the following fields:
    - In the *Networking* section:
        - **Transporter port**: Specify the port number that will be used to connect to the Transporter.
        - **Data transfer ports**: Specify a range of port numbers (from 1 to 65535) that will be used to transfer data. The range you specify should contain at least 100 ports. Make sure that the ports you specify are open in your firewall.
    - In the *Settings* section:
        - **Transporter name**: Specify a display name for the **Transporter**.
        - **Maximum load**: Specify the maximum number of tasks that the **Transporter** should process simultaneously. A task, for example, is a backup or replication of a single VM disk, or one granular recovery session.
        - **Additional load for recovery jobs**: Selecting this option reserves the **Transporter**'s resources exclusively for recovery jobs. This allows you to run recovery jobs concurrently with other types of jobs without the need to wait for their completion. The **Transporter** resources will be reserved according to the specified number.
        - **Enable debug logging for this transporter**: If needed, enable debug level logging for the current transporter. It is not recommended to use this option on a permanent basis.
        - **Enable Direct Connect for this transporter**: When this option is enabled, you can access remote resources via a single port connection without establishing a VPN connection. The following conditions must be met at the remote infrastructure to enable this feature:
            - A NAKIVO Transporter must be installed.
            - A master password must be set for security reasons. A pre-shared key is generated based on the entered master password.
            - The **Transporter** port on the local machine must be exposed to external availability via the Internet.

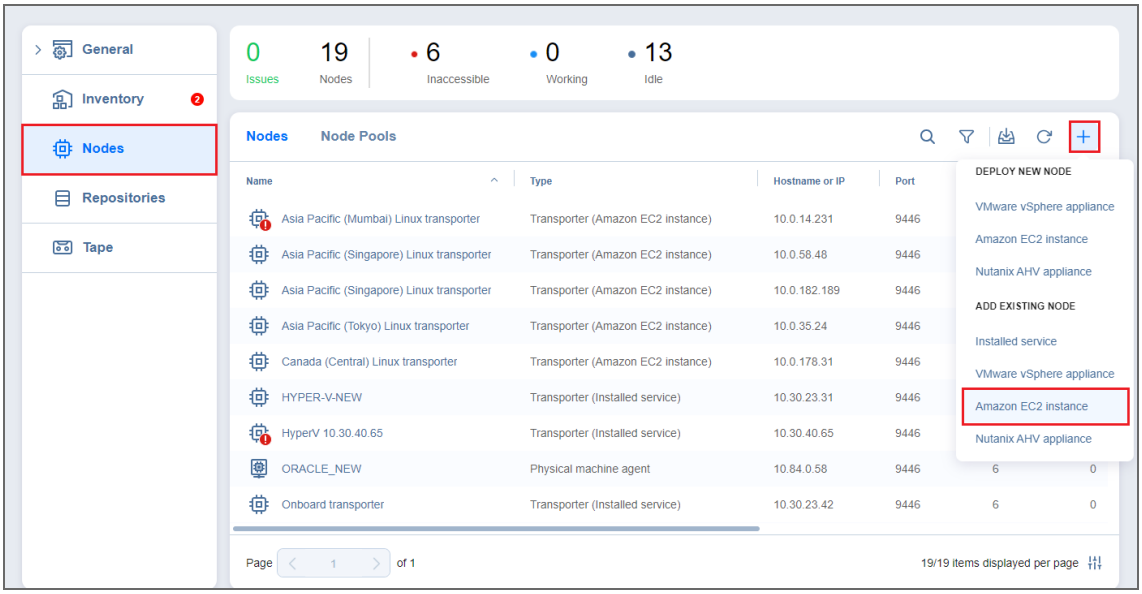
4. Click **Add**. The **Transporter** is added to the product and can be used for backup, replication, and recovery jobs.




## Amazon EC2 Instance

If you have already deployed a **Transporter** in Amazon EC2 and now wish to re-import the **Transporter** in a new instance of NAKIVO Backup & Replication, do the following:

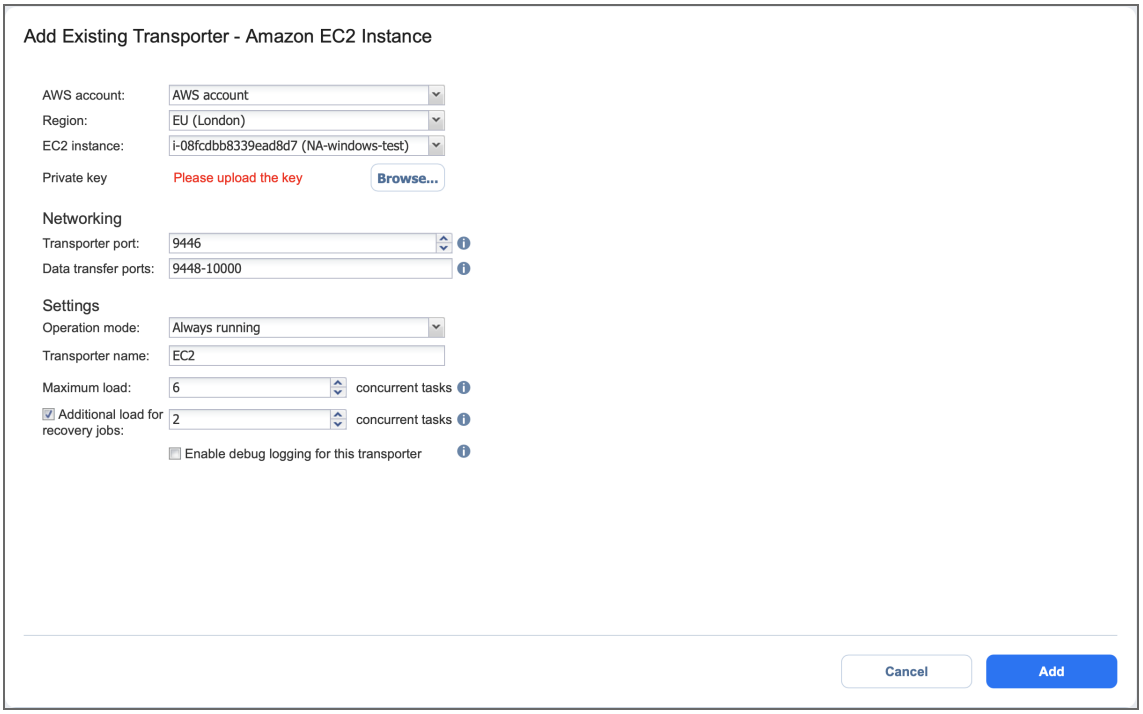
1. Click **Settings** in the left pane of the product, go to the **Nodes** tab, and click **+**.
2. In the **Add Existing Node** section, click **Amazon EC2 instance**.




3. The **Add Existing Transporter - Amazon EC2 Instance** dialog opens. Fill out the fields as described below:

- **AWS account**: Choose an appropriate Amazon AWS Account from the list of Amazon AWS Accounts added to the Inventory.
- **Region**: Choose a region in which an AWS EC2 instance with the **Transporter** is deployed.
- **EC2 Instance**: Select the Amazon EC2 Instance with the **Transporter** that you wish to add to the product.
- **Private key**: Click the **Browse** button to locate and upload the Private key for the **Transporter** Instance that was created when you deployed the **Transporter** in the cloud.
- Click **More options** to reveal and edit the following fields:
    - In the *Networking* section:
        - **Transporter port**: Specify the port number that will be used to connect to the **Transporter**.
        - **Data transfer ports**: Specify a range of port numbers (from 1 to 65535) that will be used to transfer data. The range you specify should contain at least 100 ports. Make sure that the ports you specify are open in your firewall.
    - In the *Settings* section:
        - **Operation mode**: Choose one of the following **Transporter** operation modes:
            - **Always running**
            - **Running while required**
        - **Transporter name**: Specify a display name for the **Transporter**.
        - **Maximum load**: Specify the maximum number of tasks that the **Transporter** should process simultaneously. A task, for example, is a backup or replication of a single VM disk, or one granular recovery session.
        - **Additional load for recovery jobs**: If selected, the specified amount of tasks will be added to set maximum transporter load to be used for recovery jobs exclusively.
        - **Enable debug logging for this Transporter**: If needed, enable debug level logging for the current transporter. It is not recommended that you use this option on a permanent basis.

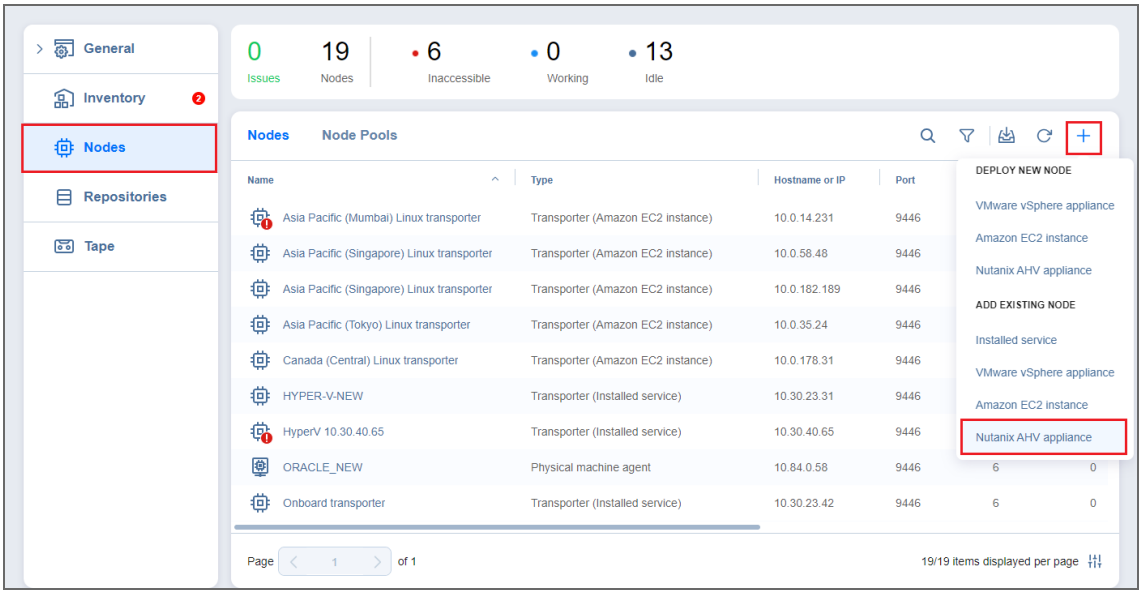
4. Click **Add**. The **Transporter** is added to the product and can be used for backup, replication, and recovery jobs.




## Nutanix AHV Appliance

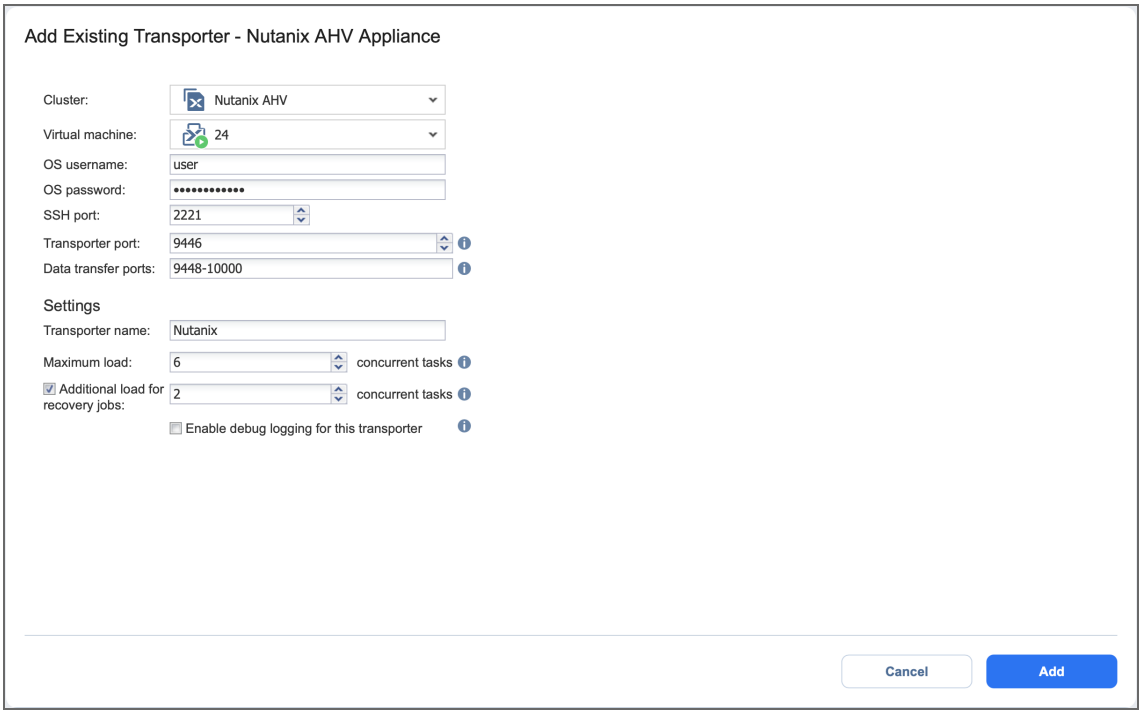
Follow the steps below to add a **Transporter** that is deployed as a Nutanix AHV appliance:

1. Click **Settings** in the left pane of the product, go to the **Nodes** tab, and click **+**.
2. In the **Add Existing Node** section, select **Nutanix AHV appliance**.




3. In the **Add Existing Transporter - Nutanix AHV Appliance** dialog, enter the following options:
   - In the **Cluster** box, select the cluster where the corresponding virtual machine is deployed.
   - In the **Virtual machine** box, specify the virtual machine on which the **Transporter** is installed.

- In the **OS Username** and **OS Password** fields, specify credentials for accessing the virtual machine.
- In the **SSH port** box, enter the SSH port if needed.
- Click **More options** to reveal and edit the following fields:
  - In the *Networking* section:
    - **Transporter port**: Specify the port number that will be used to connect to the **Transporter**.
    - **Data transfer ports**: Specify a range of port numbers (from 1 to 65535) that will be used to transfer data. The range you specify should contain at least 100 ports. Make sure that the ports you specify are open in your firewall.
  - In the *Settings* section:
    - **Transporter name**: Specify a display name for the **Transporter**.
    - **Maximum load**: Specify the maximum number of tasks that the **Transporter** should process simultaneously. A task, for example, is a backup or replication of a single VM disk, or one granular recovery session.
    - **Additional load for recovery jobs**: If selected, the specified amount of tasks will be added to set maximum transporter load to be used for recovery jobs exclusively.
    - **Enable debug logging for this Transporter**: If needed, enable debug level logging for the current transporter. It is not recommended that you use this option on a permanent basis.



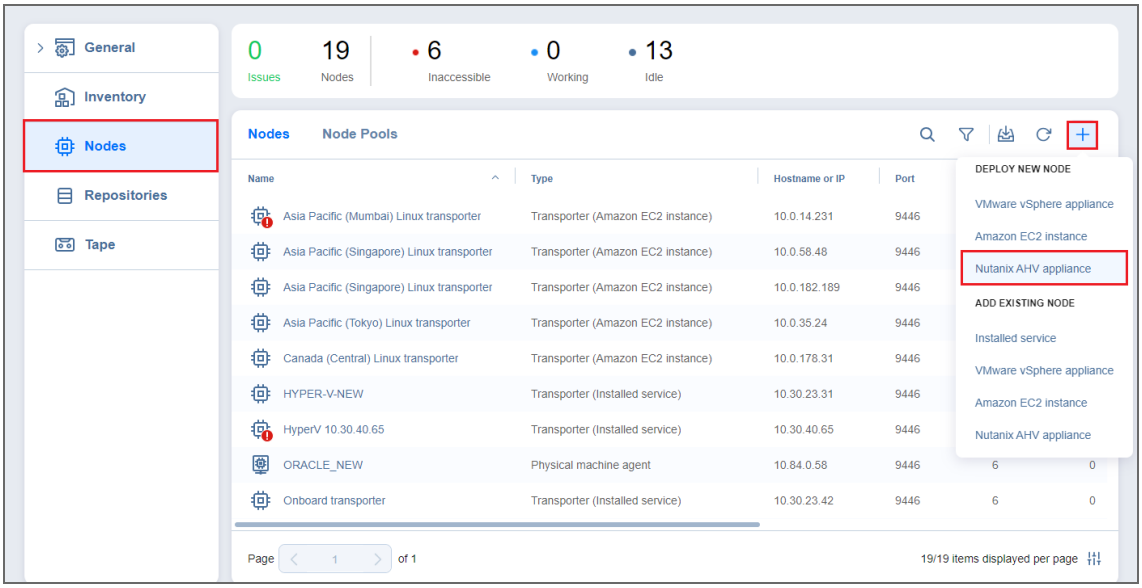
4. Click **Add**. The **Transporter** is added to the product and can be used for backup, replication, and recovery jobs.

## Deploying Transporter as Nutanix AHV Appliance

To enable NAKIVO Backup & Replication to create and run jobs within a Nutanix AHV cluster, a dedicated **Transporter** must be deployed as a Nutanix appliance in that cluster.

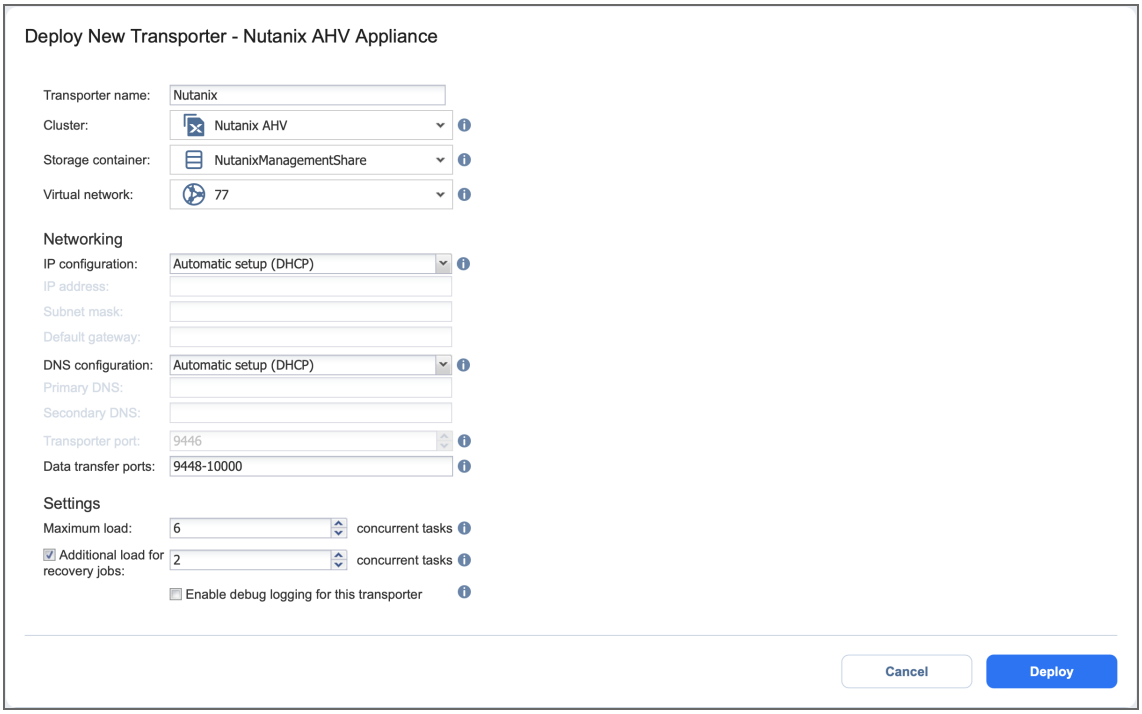
Please follow the steps below to add a transporter as a Nutanix appliance:

1. Click **Settings** in the left pane of the product, go to the **Nodes** tab, and click **+** .
2. In the **Deploy New Node** section, click **Nutanix AHV appliance**.




3. In the **Deploy New Transporter - Nutanix AHV Appliance** dialog, specify the following options:
   - **Transporter name**: Enter a name for the new **Transporter**.
   - **Cluster**: Select a cluster where the transporter VM will run.
   - **Storage container**: Select a storage container where the transporter VM will be located.
   - **Virtual network**: Select a virtual network where the transporter VM will be connected.
4. Click **Deploy** to proceed with the automatically selected networking options and default **Transporter** load configuration.
5. Alternatively, click **More options** if you wish to manually set the following options:
   - **IP configuration**: Can be either **Automatic setup (DHCP)** or **Manual setup**. With manual setup selected, specify an **IP address**, **Subnet mask** and **Default gateway**.
   - **DNS configuration**: Can be either **Automatic setup (DHCP)** or **Manual setup**. With manual setup selected, specify **Primary** and **Secondary DNS**.
   - **Transporter port**: Enter a communication port for your **Transporter**.
   - **Data transfer ports**: Enter a port range that will be used by your **Transporter** for actual data transfer.

- **Maximum load**: Specify the maximum number of tasks that the **Transporter** should process simultaneously. A task, for example, is a backup or replication of a single VM disk, or one granular recovery session.
- **Additional load for recovery jobs**: If selected, the specified quantity of tasks will be added to set the maximum **Transporter** load to be used for recovery jobs exclusively. This allows for running the specified amount of concurrent recovery jobs along with other types of jobs without the need to wait for their completion.
- **Enable debug logging for this transporter**: If needed, enable debug level logging for the current transporter. Using this option on a permanent basis is not recommended.



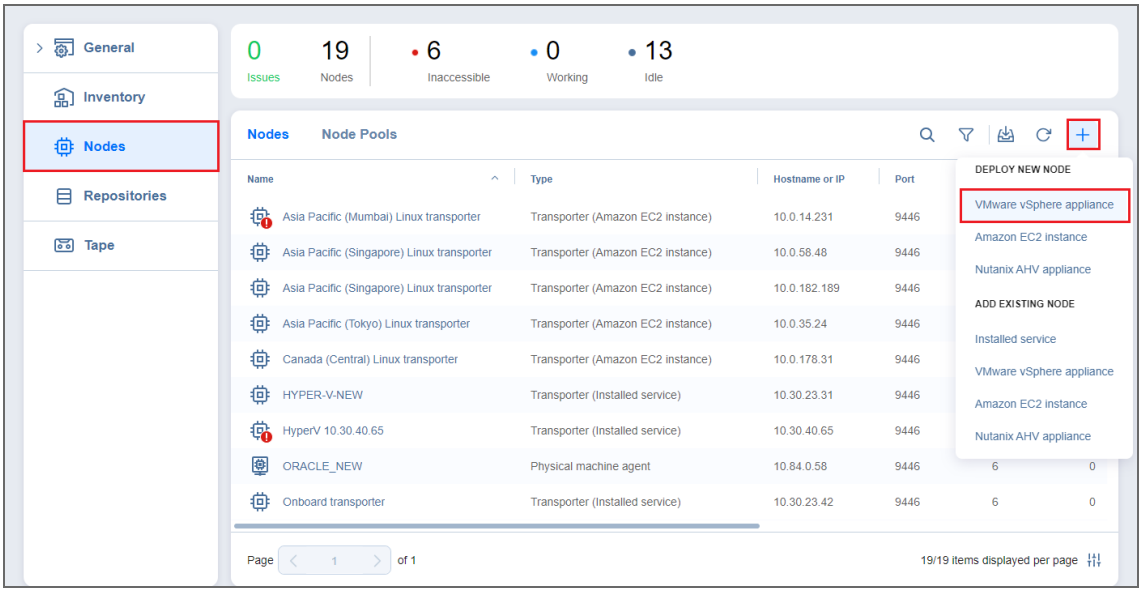

6. Click **Deploy** to begin the deployment process. Successfully deployed transporters are displayed in the **Transporters** tab.

# Deploying Transporter as VMware Appliance

> **Note**
> If your instance of NAKIVO Backup & Replication is installed on ARM-based NAS, an external **Transporter** needs to be deployed to work with VMware vCenters and ESXi hosts. This is because certain features are not supported by ARM-based NASes.

Please follow the steps below to deploy a Transporter that supports VMware vCenter:

1. Click **Settings** in the left pane of the product, go to the **Nodes** tab, and click **+** .
2. In the **Deploy New Node** section, click **VMware vSphere appliance**.




3. In the **Deploy New Transporter - VMware vSphere Appliance** dialog that opens, proceed as follows:
    - **Transporter name**: Enter a name for your **Transporter**.
    - **Host or cluster**: Select a target host or cluster.
    - **Datastore**: Select a target datastore.
    - **Virtual network**: Select a target virtual network.

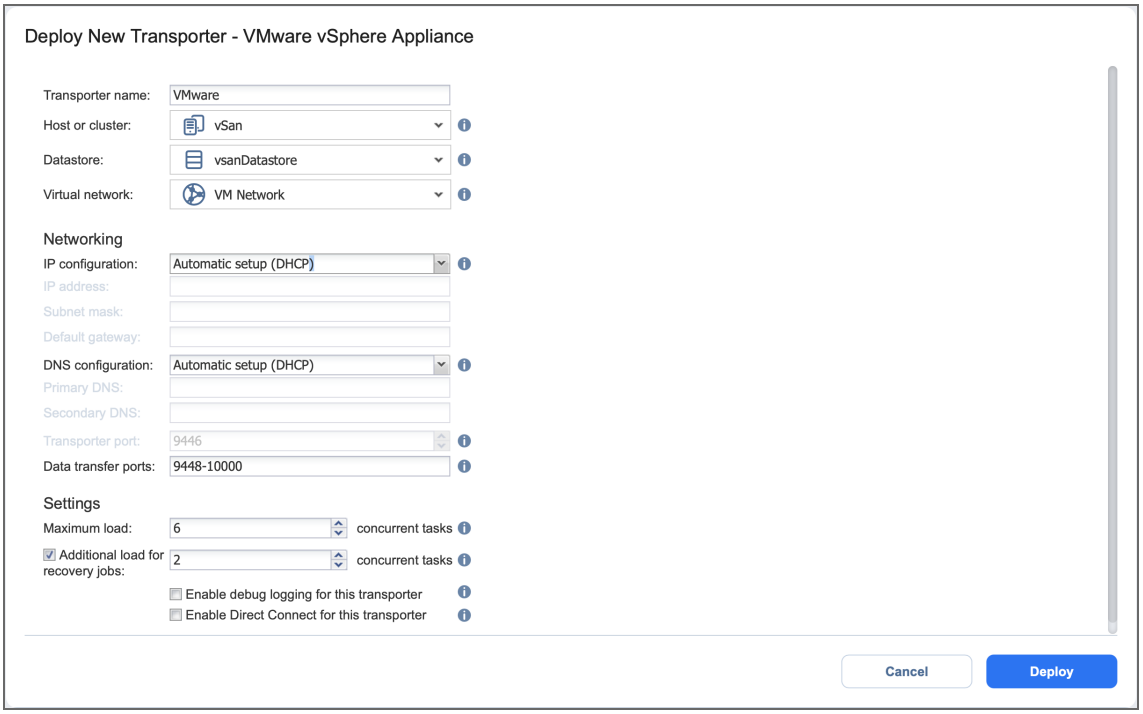
      **Note**
      An internet connection is required to deploy a new Transporter as a VMware appliance on the target host or cluster.

    - If necessary, access the advanced options for your **Transporter** by clicking **More options** and then entering data for the following parameters:
        - In the *Networking* section:
            - **IP configuration**: It can be either **Automatic setup (DHCP)**, or **Manual setup**.
            - **IP address**: If you have chosen **Manual setup** for the IP configuration, enter a **Transporter** IP address.
            - **Subnet mask**: If you have chosen **Manual setup** for the **IP configuration**, enter a subnet mask.
            - **Default gateway**: If you have chosen **Manual setup** for the **IP configuration**, enter a default gateway.
            - **DNS configuration**: It can be either **Automatic setup (DHCP)**, or **Manual setup**.
            - **Primary DNS**: If you have chosen **Manual setup** for the **DNS configuration**, enter a primary DNS server IP address.

- **Secondary DNS**: If you have chosen **Manual setup** for the **DNS configuration**, enter a secondary DNS server IP address.
- **Transporter port**: Enter a communication port for your transporter.
- **Data transfer ports**: Enter a port range that will be used by your transporter for actual data transfer.

- In the *Settings* section:
  - **Maximum load**: A number of tasks concurrently processed by the **Transporter**.
  - **Additional load for recovery jobs**: If selected, the specified amount of tasks will be added to set maximum transporter load to be used for recovery jobs exclusively. This allows for running the specified amount of concurrent recovery jobs along with other types of jobs without the need to wait for their completion.
  - **Enable debug logging for this transporter**: When selected, it enables debug level logging for the Transporter. It is not recommended to have this option selected on a permanent basis.
  - **Enable Direct Connect for this transporter**: When this option is enabled, you can access remote resources via a single port connection without establishing a VPN connection. The following conditions must be met at the remote infrastructure to enable this feature:
    - A NAKIVO **Transporter** must be installed.
    - A master password must be set for security reasons. A pre-shared key is generated based on the entered master password.
    - The **Transporter** port on the local machine must be exposed to external availability via the Internet.

4. Click **Deploy** to confirm deploying the **Transporter**.



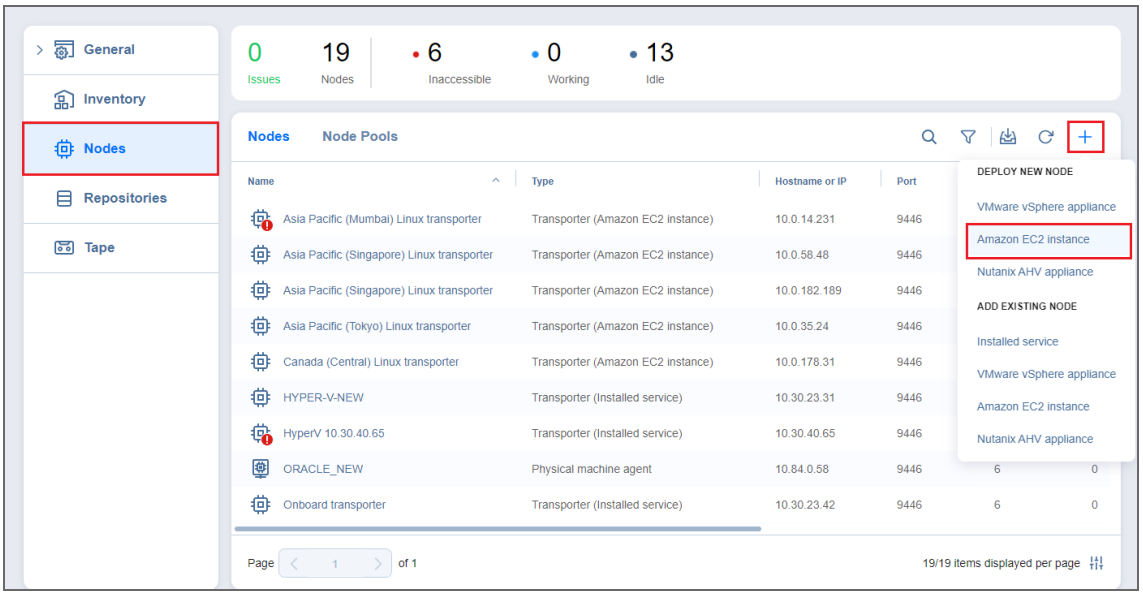

# Deploying Transporters in Amazon EC2

You need to deploy a **Transporter** in Amazon EC2 to enable the following features:

- Backing up VMware VMs and/or Amazon EC2 Instances to a backup repository located in Amazon EC2.
- Backing up Amazon EC2 Instances in a particular Amazon EC2 Region.

NAKIVO Backup & Replication automates deploying a **Transporter** in Amazon EC2. To deploy a **Transporter** in Amazon EC2 within the product interface, follow the steps below:

1. Click **Settings** in the left pane of the product, go to the **Nodes** tab, and click **+** .
2. In the **Deploy New Node** section, click **Amazon EC2 instance**.



3. The **Deploy New Transporter - Amazon EC2 Instance** dialog opens. Fill out the fields as described below:
    - **Transporter name**: Enter a name for the **Transporter**.
    - **Region**: Select an Amazon EC2 region where you wish to deploy the **Transporter**. This will enable you to create a backup repository in the region as well as back up Amazon EC2 Instances available in the region.
    - **Instance type**: Choose a type of Amazon EC2 Instance (for example, "t2.medium") that will be used to deploy the **Transporter**. Note that more powerful instances may be able to process data faster, but will cost more to run on Amazon EC2.

      > **Note**
      > ARM-based instances cannot be selected if you have chosen **Windows** for the **Platform** option.

    - Click **More options** to reveal and edit the following options:
        - In the *Networking* section:
            - **Automatically configure VPC for this transporter**: If selected, a new VPC with a single public subnet will be created and used to deploy this transporter. If you want to deploy the **Transporter** into a different VPC and subnet, deselect this option.
                - **Network**: Select a network to which the Amazon EC2 instance with the **Transporter** will be connected.
                - **Subnet**: Select a subnet for the Amazon EC2 Instance with the **Transporter**.
            - **Allowed traffic from**: Enter the IP addresses of the machines that can connect to the Amazon EC2 instance with the **Transporter**. Access from other IP addresses will be restricted.

              > **Important**
              > By default, the Amazon EC2 security group is not restricted; that is, the **Transporter** can be accessed by and receive tasks from any machine. For security purposes, restrict traffic to trusted IP addresses.

            - **Transporter Port**: Specify the port number that will be used to connect to the **Transporter**.

- **Data transfer ports**: Specify a range of port numbers (from 1 to 65535) that will be used to transfer data. The range you specify should contain at least 100 ports. Make sure that the ports you specify are open in your firewall.

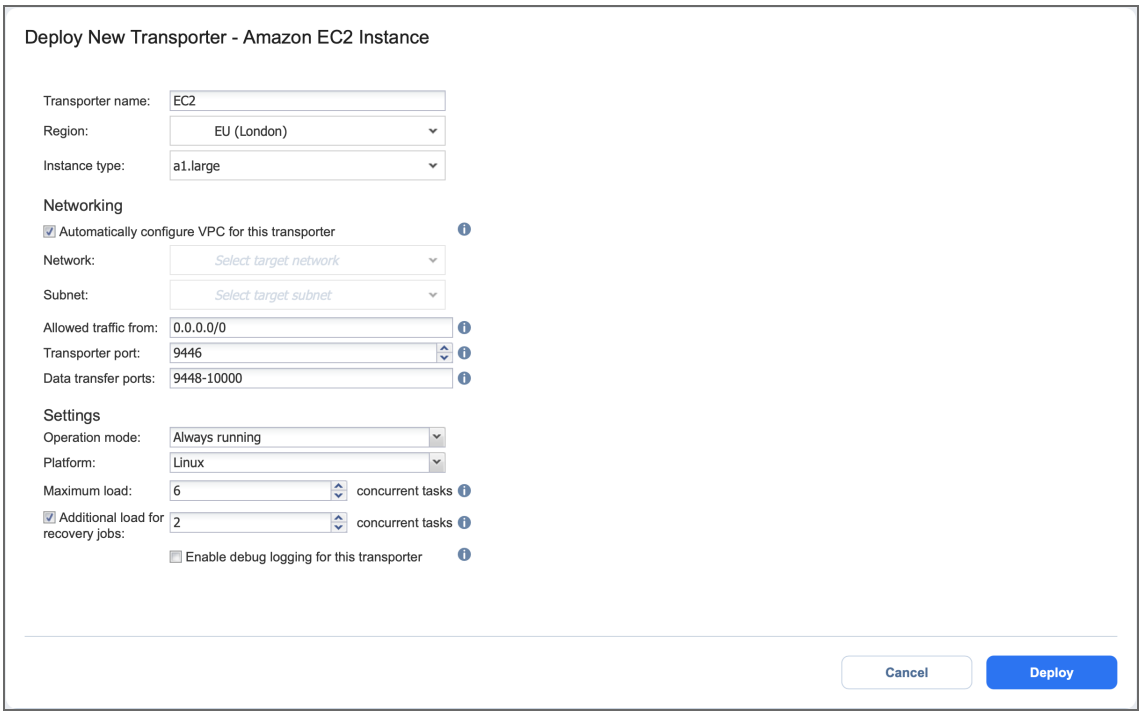
- In the *Settings* section:
  - **Operation mode**: If you select the **Running while required** option, the Amazon EC2 Instance with the **Transporter** will be powered on only when the **Transporter** is required to run a backup, replication, and recovery tasks.
  - **Platform**: Choose an OS for the instance where the **Transporter** will be deployed.

    **Note**
    Windows OS is not supported for ARM-based instances.

  - **Maximum load**: Specify the maximum number of tasks that the **Transporter** should process simultaneously. An example of a task is processing a single VM disk or a single file recovery session.
  - **Additional load for recovery jobs**: If selected, the specified amount of tasks will be added to set the maximum **Transporter** load to be used for recovery jobs exclusively. This allows for running the specified quantity of concurrent recovery jobs along with other types of jobs without the need to wait for their completion.
  - **Enable debug logging for this Transporter**: Enables debug level logging for the current **Transporter**. Since this feature slows down **Transporter** performance, it is recommended that you enable debug logging only for the investigation of support issues.

**Note**
Refer to "Amazon EC2 Concepts" on page 7 for the definitions of Amazon EC2-related terms.

4. Click **Deploy**.




**Note**

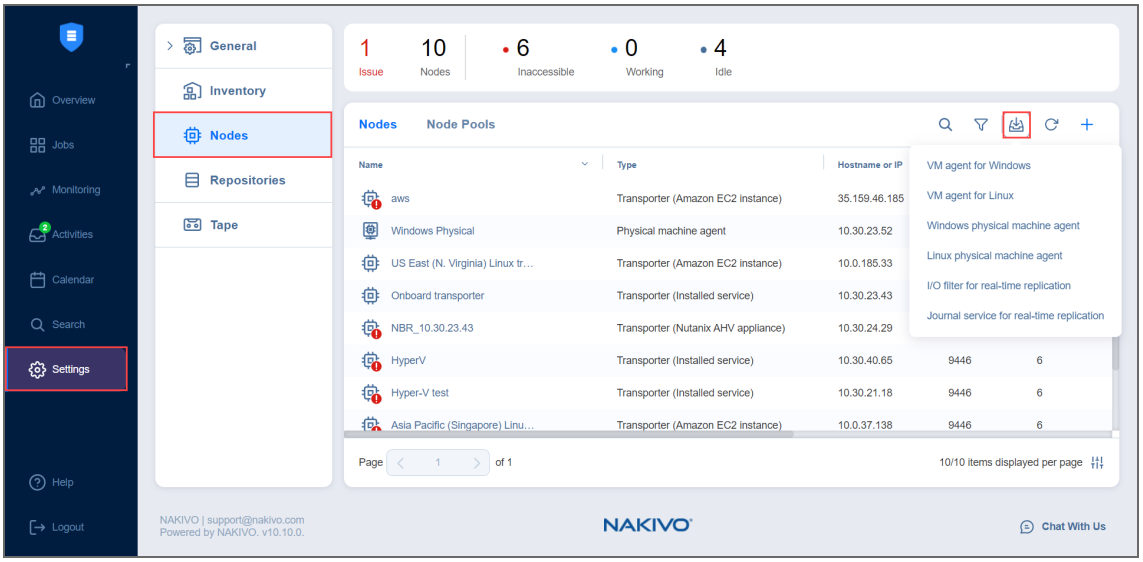
- After deploying a **Transporter** in Amazon EC2, you need to download the Transporter Key. A Transporter Key is used by NAKIVO Backup & Replication to access and manage the **Transporter** in Amazon EC2. If you lose the current instance of NAKIVO Backup & Replication and install a new copy of the product, you will need to provide the Transporter Key to access the **Transporter**.
- You may be additionally charged for using a 3rd-party resource. Please refer to the 3rd-party resource provider documentation for details.

# Installing a VM Agent

Each VM agent (VMA) has its own ID, certificate, and pre-shared key. The VMA ID must be unique, meaning that duplicate agents (that is, agents with the same ID) are not permitted. It is recommended that you use unique certificates and master passwords with each VMA. See below for installation details.

## Installation

You can download the VM agent (VMA) installer files from the application UI. To do this, go to **Settings** > **Nodes**, click the **Download** button, and select the needed OS.




Installation is done via the command line and requires you to set up a master password with a minimum length of 5 characters. Note that you cannot install a VMA on a machine on which a **Transporter** is installed, and vice versa. See command examples below:

- Windows: *installer.exe --cert C:\certificate.pem -p ExamplePassword --eula-accept*
  - To confirm installation success, check **Control Panel** > **Programs & Features**
  - Installation result is logged in *C:\install.log*
- Linux: *installer.sh -s 9445 --cert /tmp/certificate.pem -p ExamplePassword --eula-accept*
  - To confirm installation success, run the *systemctl status nkv-bhsvc* command.
  - Installation result is logged in */tmp/nkv-install.log*

After installing the VM agent, proceed as described in “Using a VM Agent” on page 604.

## Removal

Removing a VM agent in the product UI does not remove it from the VM. Removing VM agents is only possible by uninstalling them from the VM. See uninstallation details below:

- Windows: Run *Uninstall NAKIVO Backup & Replication Agent*, located in the *NAKIVO* folder within *Programs*.
- Linux: Run the */opt/nakivo/agent/uninstall* command.

### Certificate

Installed VM agents (VMAs) may use a CA-signed certificate or a self-signed certificate. After successful installation of the VMA, local copies of the provided certificate for VMA installation are automatically removed.

Multiple CA certificate chains are not supported. CA certificates must be placed in the **Director** installation folder. These certificates are trusted automatically by the **Director**.

If a CA-signed certificate is not provided, the VMA automatically generates a self-signed certificate.

## Installing I/O Filter

**I/O Filter** is a software module that installs into a cluster and allows intercepting and manipulating VM's I/O. In order to create and run a real-time replication job for VMware, the source cluster(s) must have the **I/O Filter** tool installed. See the sections below for details.

- Installation
- Update
- Removal
- Security

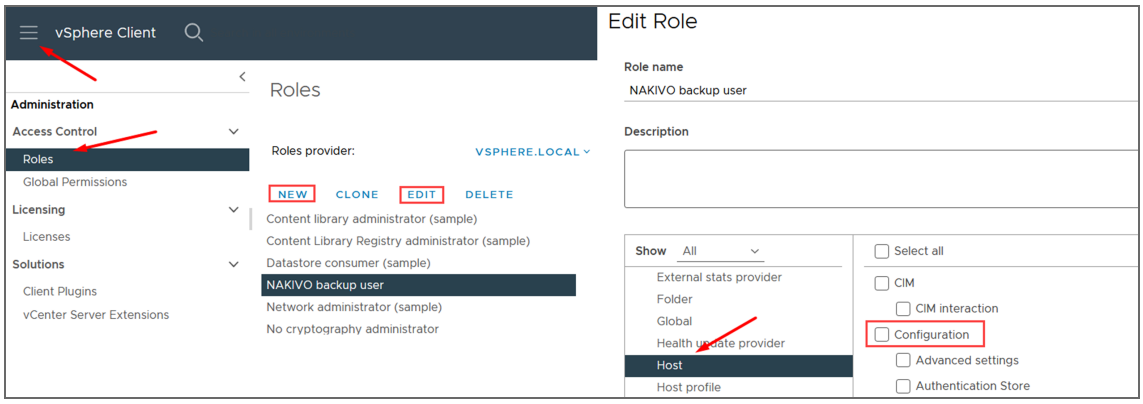
### Installation

**I/O Filter** is a software module that installs into a cluster and allows intercepting and manipulating VM's I/O. **I/O Filter** should be installed into a Cluster. There is no need to install the **I/O Filter** on the target side. The **Setup Real-Time Replication Service** wizard is used to install I/O Filter on Cluster and assign VM Storage policy to the source VMs.

Installing the **I/O Filter** on a cluster will create an instance of **I/O Filter Daemon** on every ESXi host of the cluster.

**Notes**

- Hosts may need to be manually put into maintenance mode during the **I/O Filter** installation.
- The ESXi host's acceptance level must be set to **Community level**. Therefore, only clusters using baselines are supported (no clusters using a single image are supported). To check a host's acceptance level, select the host and navigate to **Configure** > **System** > **Security Profile**. The acceptance level is displayed under **Host Image Profile Acceptance Level**.
- The *Query patch* vCenter permission must be granted to NAKIVO Backup & Replication to manage the I/O Filter (**Host** > **Configuration** > **Query patch**). Note that the *Query patch* vCenter permission must be granted for the user who used to add to the NAKIVO Inventory. Refer to the following screenshots for details:




**Note**

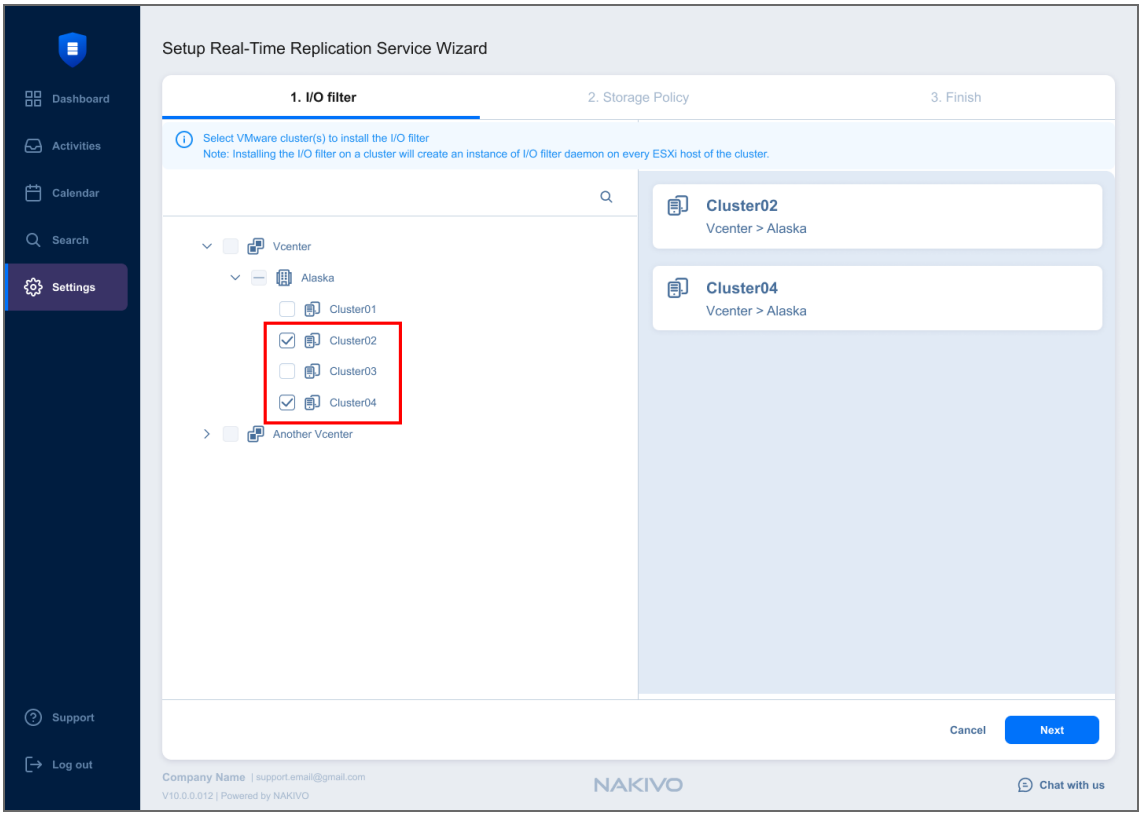
By default, the **I/O Filter** bundle file is located in the NAKIVO Backup & Replication *iofilter* folder:

- On Windows:
  - C:\Program Files\NAKIVO Backup &Replication\packages\iofilter-vaio7
  - C:\Program Files\NAKIVO Backup & Replication\packages\iofilter-vaio8
- On Linux:
  - /opt/nakivo/director/packages/iofilter-vaio7
  - /opt/nakivo/director/packages/iofilter-vaio8

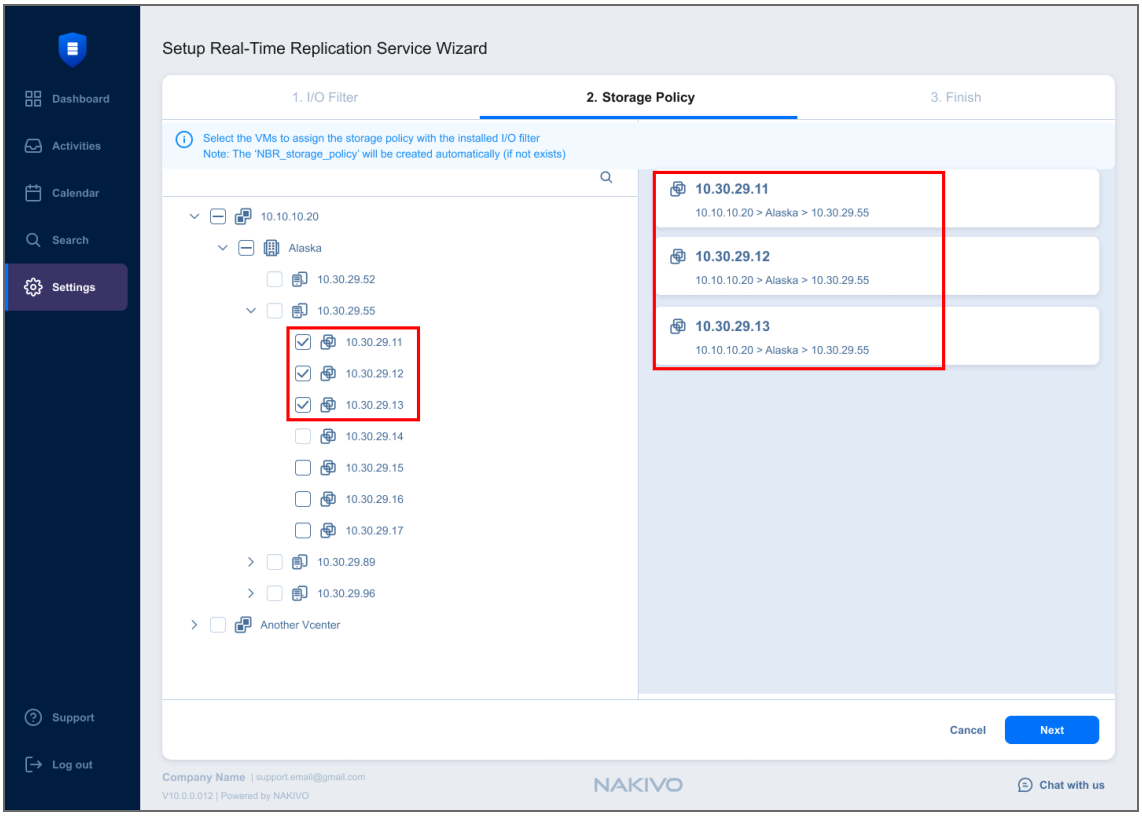
The VAIO SDK 7 can be used for vSphere versions 7.0 – 8.0.3 and is automatically installed on vSphere version 7.0. The VAIO SDK 8 can be used for vSphere versions 8.0 – 8.0.3 and is automatically installed according to the vSphere version.

Proceed as follows:

1. Navigate to the main UI and select the **Settings** > **Nodes** tab.
2. Click “**+**” (plus) button and select **Real-Time Replication Service** under the **Deploy New Node** section. This opens the **Setup Real-Time Replication Service Wizard**, showing a list of discovered VMware hosts and clusters.
3. Choose the VMware clusters where the **I/O Filter** package will be deployed. If a cluster has an **I/O Filter** installed, the VMware cluster is selected by default.




4. Click **Next** to proceed to the **Storage Policy** configuration step.
5. Select VMware VMs from the list on the left to assign a storage policy that uses the installed **I/O Filter**. The default storage policy will automatically be assigned to the selected VM.




6. If the selected VM already has the default storage policy assigned, a dialog window will appear, prompting you to remove the policy from the VM. Click **Continue** to proceed, and the VM will be deselected. Otherwise, click **Cancel** to dismiss the window and keep the policy assigned.
7. Click **Next** to proceed to the **Finish** step.
8. Click **Close** to exit the wizard. The **I/O Filter** will be automatically installed on the selected cluster(s).

**Notes:**

- The list of installed I/O filters is displayed in the **Nodes** UI. Click on the desired **I/O Filter Daemon** to view its details.
- I/O filters can be installed on up to 10 clusters simultaneously.
- If the installation fails, it will be retried 3 times with 1-minute intervals between attempts.
- If a cluster has an older version of the **I/O Filter** installed, the **I/O Filter** will be updated to the latest version.
- Clusters with the latest version will skip the installation process.
- In multi-tenant mode, if vCenter was added to the Inventory from the tenant, the tenant can install/uninstall I/O Filter and assign/unassign VM Storage Policy. In case the Master tenant has assigned resources to the tenant, the local tenant user can not install/uninstall them, only refresh action is available.

## Update

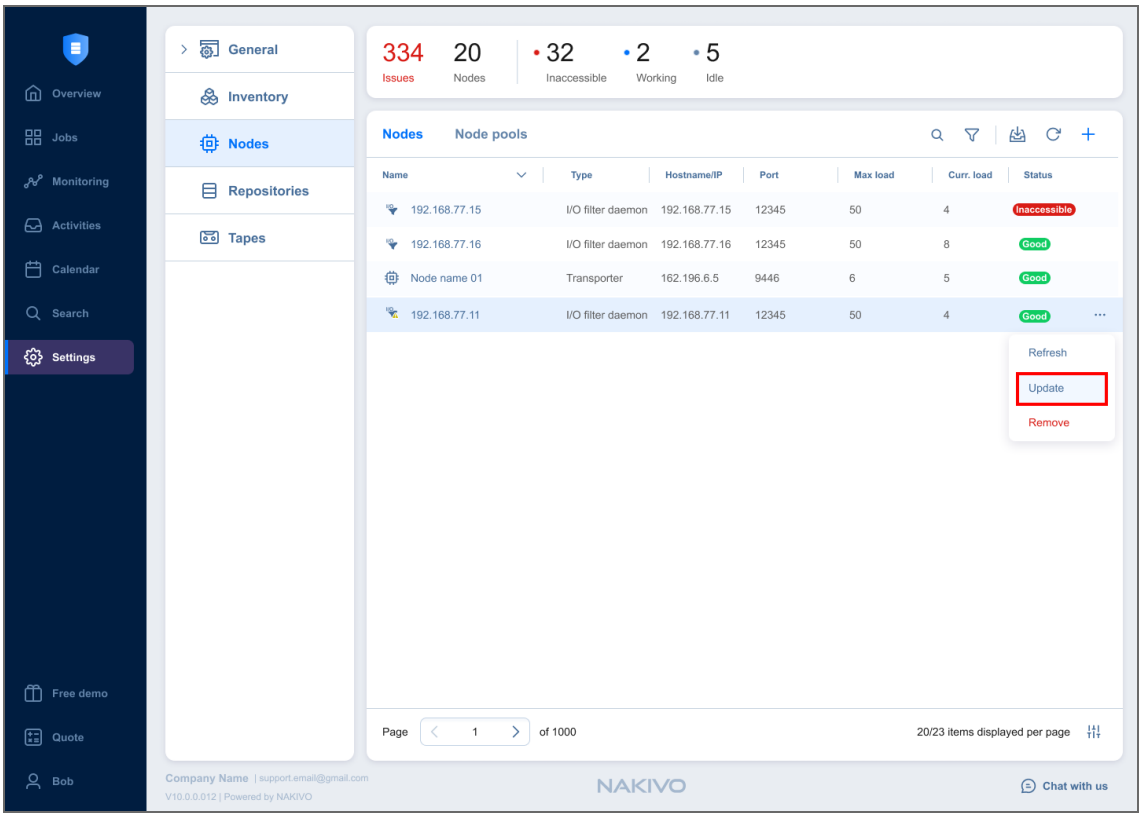
**Notes**

- Hosts may need to be manually put into maintenance mode during the **I/O Filter** update.
- The ESXi host's acceptance level must be set to **Community level**. Therefore, only clusters using baselines are supported (no clusters using a single image are supported). To check a host's acceptance level, select the host and navigate to **Configure** > **System** > **Security Profile**. The acceptance level is displayed under **Host Image Profile Acceptance Level**.

If an **I/O Filter** is not the latest version, a warning sign is displayed. To update an existing **I/O Filter**, do the following:

1. Navigate to the main menu and select the **Settings** > **Nodes** tab.
2. Hover the **I/O Filter Daemon** that requires updating from the list, click ellipses and select the **Update** option. **Setup Real-Time Replication Service** wizard opens.
3. On the **I/O Filter** step of the **Setup Real-Time Replication Service** wizard leave the Cluster selected and click **Next**.
4. On the **Storage Policy** step you may select/unselect the VMs if needed or leave them as is. Click **Next** to proceed.
5. On the **Finish** step you may click the **Activities** link to see the update result or click **Close** to exit the

Wizard and stay on the **Nodes** tab. The **I/O Filter** update will be initiated.




**Note**

During the **I/O Filter** Update the old **I/O Filter** is removed and a new **I/O Filter** is installed.

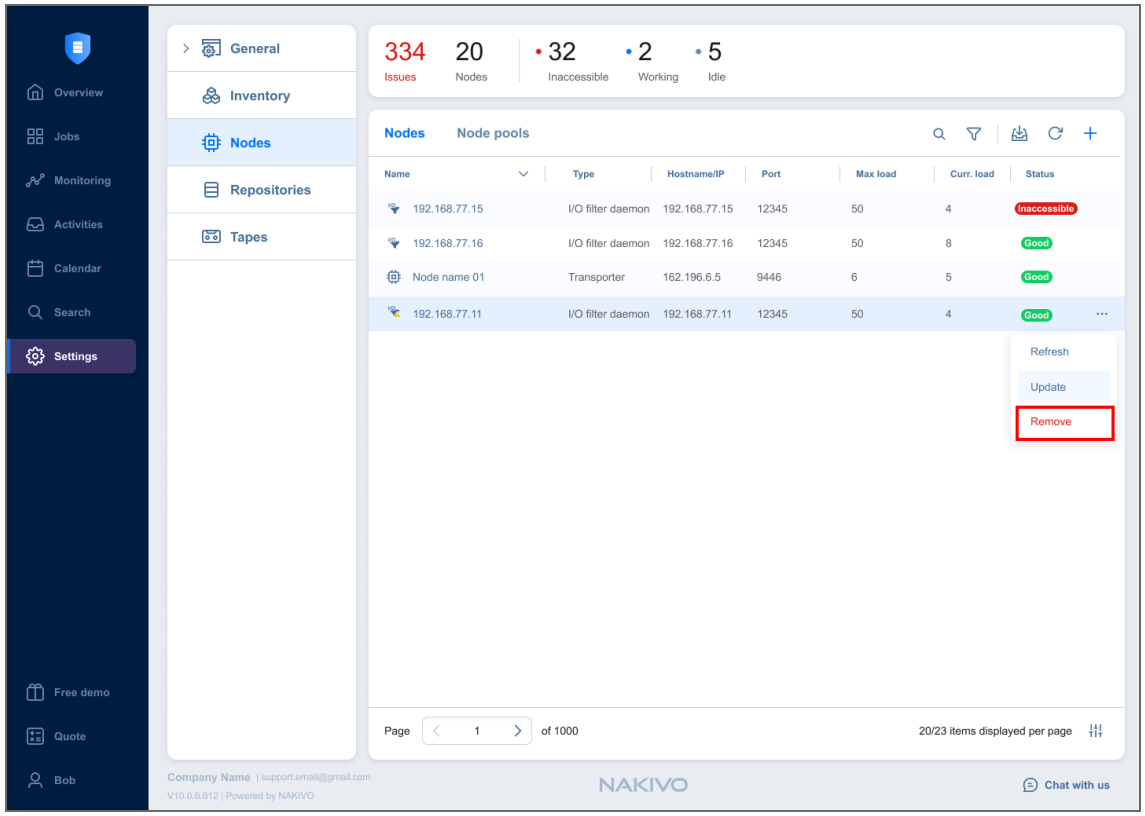
## Removal

To remove the **I/O Filter**:

1. Navigate to the main menu and select the **Settings** > **Nodes** tab.
2. Select the **I/O Filter Daemon** from the list, click ellipses and select **Remove**. The **Setup Real-Time Replication Service** wizard opens.
3. On the **I/O Filter** step of the **Setup Real-Time Replication Service** wizard, deselect the cluster and click **Next**.
4. On the **Storage Policy** step click **Next**.
5. On the **Finish** step, click the **Activities** link to see the result or click **Close** to exit the wizard and stay on the **Nodes** tab. The **I/O Filter** will be removed automatically.

**Notes**

- If Real-Time Replication jobs or Real-Time Replication Failover jobs are using a specified **I/O filter**, setting up the real-time replication service will fail. You can view the list of jobs using the specified **I/O Filter** by navigating to the **Activities** tab and clicking the **Details** link.
- VM storage policies will be unassigned if the cluster is deselected during the **I/O Filter** step of the **Setup Real-Time Replication Service** wizard.
- When removing an **I/O filter**, NAKIVO Backup & Replication verifies whether any VMs or jobs across the entire vCenter are using it. The uninstallation will be blocked if the **I/O filter** is still in use.
- If VMware vSphere Distributed Resource Schedule (DRS) is disabled, manually place cluster hosts in maintenance mode before proceeding. If DRS is enabled, hosts are placed in maintenance mode automatically.
- If the I/O filter is removed from the cluster (vSphere Client: **Cluster** > **Configure** > **I/O Filters**) but remains on the individual hosts (vSphere Client: **Host** > **Configure** > **I/O Filters**), run the following command on the host to remove it manually: *esxcli software vib remove -n nakfilter --maintenance-mode*



## Security

Communication between the Director and the **I/O Filter** is encrypted. Each **I/O Filter** instance has a certificate of the parent ESXi host and pre-shared key. The certificate and pre-shared key ensure trusted communication between the **Director** and **I/O Filter**.

To view the certificate details (including the certificate fingerprint and validity period), go to **Settings** > **Nodes** and select the desired **I/O Filter Daemon**.

> **Note**
> When encryption is enabled in the Real-Time Replication job settings, the **I/O Filter Daemon** sends data to **Journal Service** via port 33092. Otherwise, port 33090 is used.

# Installing Journal Service

In order to create and run a real-time replication job for VMware, target ESXi hosts must contain a VMware Virtual Appliance with the **Journal Service** installed. See the sections below for more details.


- "Installation" below
- "Update" on the next page
- "Removal" on the next page
- "Security" on the next page


## Journal Service

The **Journal Service** is a software module that installs into the target transporter and manages the **I/O Journal** created per VM disk and which is a set of records of I/O operations that keeps reference recovery points (snapshots of data changes made to a source VM created at the moment the replication job is stopped) of a replica VM and can be used for failover. The size and history limits of the **I/O Journal** can be configured in the **Retention** step of the Real-Time Replication Job Wizard for VMware.

**Journal Extent** is an internal logical part of the **I/O Journal** where reference recovery points of a replica VM are kept. One **Journal Extent** is created for every RTR replication job run and contains one or more **Frames**, which are logical components containing up to 65,536 data blocks, each block containing 4096 bytes of data.

## Installation

The **Journal Service** is automatically installed on **Transporter** Virtual Appliances (VA) where it is not already present when you complete the **Real-Time Replication Job Wizard**.

**Notes**:

- Ensure that the specified target **Transporter** is not used by other jobs.
- In environments with multiple **Transporter** VAs, the **Journal Services** can be installed simultaneously.
- If a target host includes multiple **Transporter** VAs, the first VA in the list is selected for installation by default.

## Update

The **Journal Service** is updated as part of the NAKIVO Backup & Replication update process. The **Journal Service** is updated only if its newer version is available in the updated version of NAKIVO Backup & Replication.

**Notes**:

- **Transporter** VAs are automatically updated after the inventory refresh when the new version of **I/O Filter** is found.
- Updates are applied regardless of the **Journal Service** current status (for example, stopped, failed, or running), including **Journal Services** installed in previous versions of NAKIVO Backup & Replication.

## Removal

To remove an installed **Journal Service** from a virtual appliance, run the following commands:

```
bash /opt/nakivo/journalservice/uninstall
OR
./opt/nakivo/journalservice/uninstall
OR
cd /opt/nakivo/journalservice
./uninstall
```

## Security

The Journal Service operates in a secure mode, ensuring both the integrity and confidentiality of its operations. Communication between the Director and the Journal Service is encrypted.

Certificate and Pre-Shared Key

Each Journal Service instance is assigned a unique certificate and a pre-shared key:

- **Certificate**: Establishes and maintains trust from the Director to the Journal Service.
- **Pre-Shared Key**: Provides trust in the opposite direction, from the Journal Service to the Director.

The Journal Service uses the certificate and pre-shared key associated with the Transporter where it is installed. During startup, the Journal Service retrieves these credentials directly from the Transporter.

### Installation and Update Lifecycle

The certificate and pre-shared key remain intact during updates to the Journal Service, whether the updates are made manually or automatically.

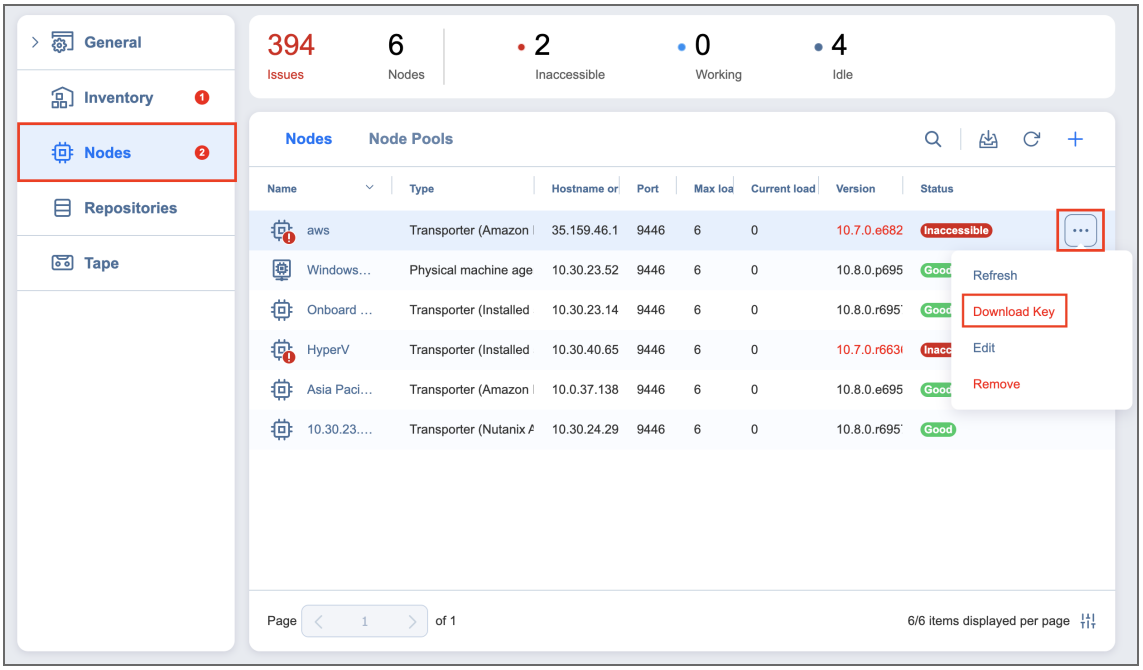
# Managing Nodes

Refer to the following topics:


- “Editing Nodes” below
- “Downloading Transporter Credentials” below
- “Enabling Direct Connect” on page 596
- “Managing Node Pools” on page 601
- “Refreshing Node Details” on page 602
- “Removing (Deleting) Nodes” on page 603
- “Using a VM Agent” on page 604


## Downloading Transporter Credentials

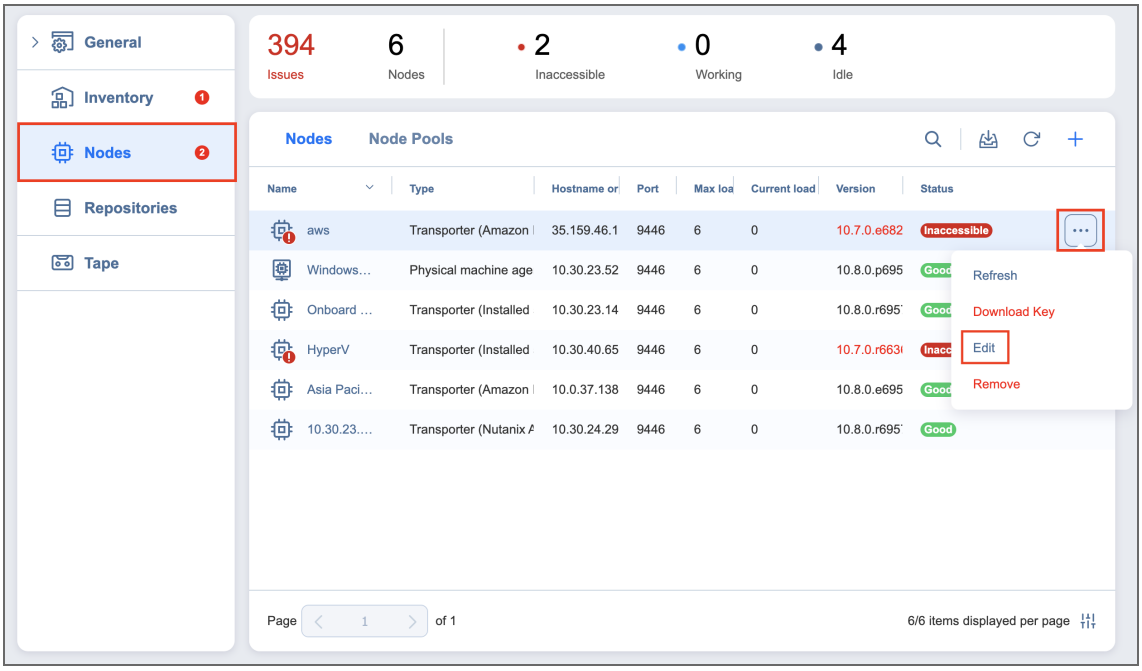
If you would like to import an Amazon EC2, Nutanix AHV, or VMware **Transporter** into another installation of NAKIVO Backup & Replication, you need to download the **Transporter**'s credentials. To obtain the credentials, hover over the desired **Transporter** and click the ellipsis **Manage** button on the right side. In the dialog box, click **Download Key**. This begins the download of a ZIP file containing the **Transporter**'s credentials.




## Editing Nodes

To modify the settings of an existing node, follow the steps below:

1. Click **Settings** in the left pane of the product.
2. Go to the **Nodes** tab and hover over the node you would like to edit.
3. On the right side, click the ellipsis **Manage** button and then click **Edit**.



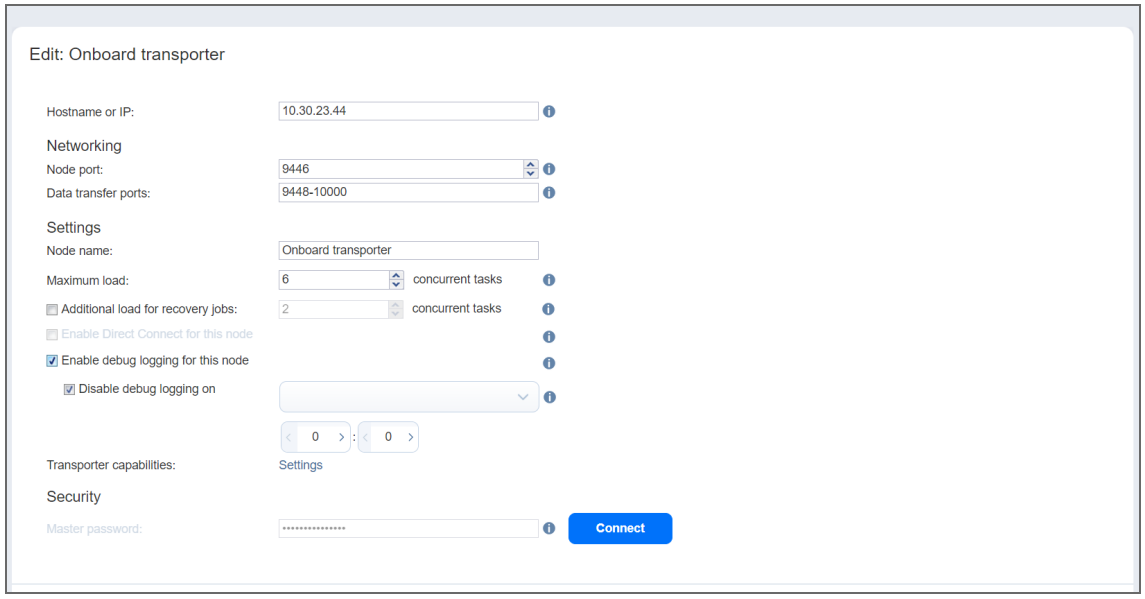

4. A dialog opens for editing the node's settings. Edit the settings as required:
    - **Hostname or IP**: Here you can edit the IP address or hostname of the machine on which the node is installed. Not applicable to Nutanix AHV Appliances, VMware vSphere appliances, or Amazon EC2 instances.
        - In the *Networking* section:
            - If editing a Nutanix AHV Appliance, VMware vSphere appliance, or Amazon EC2 instance, you can edit the following options:
                - **OS username**: Enter the username used to access the virtual machine.
                - **OS password**: Enter the password for the username entered previously (not applicable to EC2 instances).
                - **SSH port**: Enter the SSH port if needed.
            - If editing other node types:
                - **Node port**: Enter a communication port for your node.
                - **Data transfer ports**: Enter a port range that will be used by your node for actual data transfer.
        - In the *Settings* section:
            - **Node name**: Edit the name of your node.
            - **Maximum load**: Edit the number of tasks concurrently processed by the node.
            - **Additional load for recovery jobs**: If selected, the specified amount of tasks will be added to the set maximum node load to be used for recovery jobs exclusively. This allows for running the specified amount of concurrent recovery jobs along with other types of jobs without the need to wait for their completion.

- **Enable Direct Connect for this node**: When this option is enabled, you can access remote resources via a single port connection without establishing a VPN connection. The following conditions must be met at the remote infrastructure to enable this feature:
  - A NAKIVO **Transporter** or Agent must be installed.
  - A master password must be set for security reasons. A pre-shared key is generated based on the entered master password.
  - The node port on the local machine must be exposed to external availability via the Internet.
- **Enable debug logging for this node**: Enable/disable debug level logging for the node. Having this option enabled on a permanent basis is not recommended.
- **Transporter capabilities**: Click the **settings** button next to this option to modify the additional capabilities of the selected **Transporter**. Click **Apply** to save your changes.

**Notes**

- **Backup repository management** cannot be disabled if the **Transporter** is assigned to a repository.
- **VMware vSphere support** cannot be disabled if the **Transporter** is selected in a VMware backup, replication, or recovery job.
- **Transporter** capabilities is not supported for physical agent.

5. Click **Apply** to save your changes.



# Enabling Direct Connect

To start using this feature, you can choose different approaches. For feature requirements, refer to Direct Connect Requirements.

## Direct Connect

To start using this feature, take the following steps:

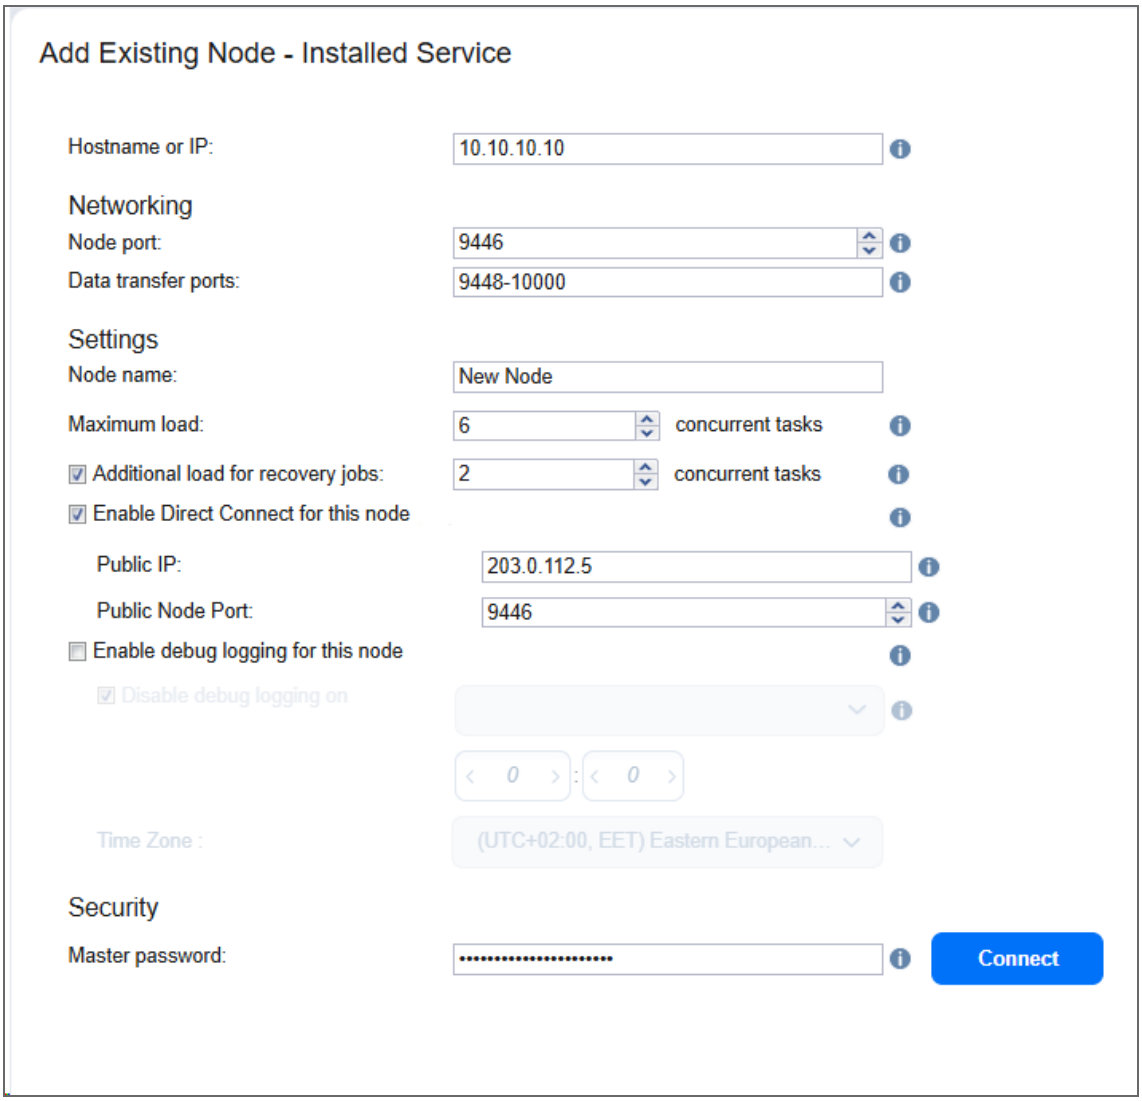
1. Download NAKIVO Backup & Replication installer or a VA with **Transporter**.
    1. Install Transporter at the remote environment and set the master password during the installation. Alternatively, deploy Transporter VA in the remote environment. The master password for the VA can be set after the deployment.

    **Notes**

    - For Windows, the installer UI and silent installations are supported.
    - For Linux, interactive and silent installations are supported.
    - For NAS, see the Supported NAS devices.
    - Master Password is required for Direct Connect Transporter installation.

2. If you didn't set up the master password during the installer setup or if you have deployed Transporter VA, you need to generate a pre-shared key. This can be done at any time after the installation. Execute the following command in **Transporter** installation folder:
   *bhsvc -b "password"*
3. Restart the Transporter service after generating the password to apply the change.
4. Expose the Transporter port on the local machine to be externally available via the Internet.
5. Add the **Transporter** to NAKIVO Backup & Replication by entering the **Transporter** details.
6. Select **Enable Direct Connect** for this transporter.
7. Enter **Public IP** and **Public Node Port** (IP and Port of the external **MSP Transporter** that is reachable from the internet and used for communication with the **Direct Connect Transporter**). Ensure that the entered combination of the **Public Node Port** and **Public IP** address is not used by other transporters.
8. Enter master password.
9. Click **Connect**.
10. View Transporter certificate information.
11. Accept received certificate.
12. Confirm adding the **Transporter**.

The client's **Direct Connect Transporter** is ready to be used for data protection activities.




## MSP Direct Connect

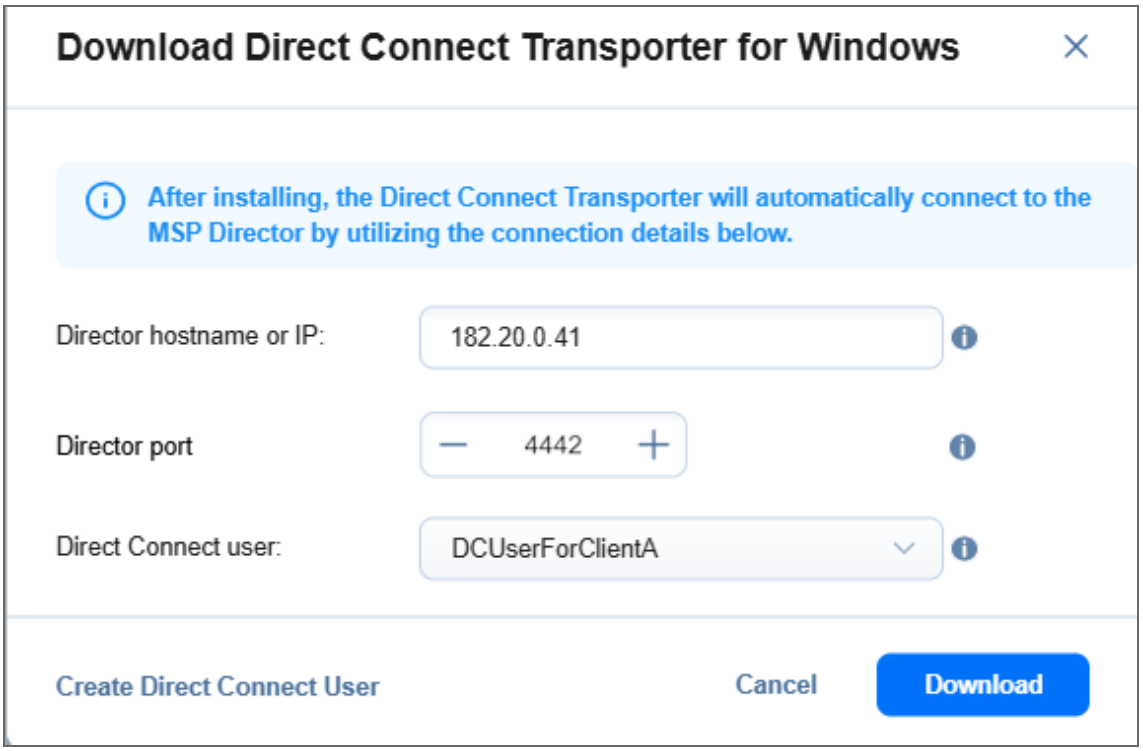
To enable MSP Direct Connect, the **Direct Connect Transporter** must be installed on the client's Windows/Linux operating system.

Installation and setting up the **Direct Connect Transporter** is performed according to the following flow:

1. Go to **Settings** > **Nodes**, click **Download**, and select **Direct Connect Transporter for Windows** or **Direct Connect Transporter for Linux** from the drop-down.

> **Note**
> The Direct Connect Transporter for Windows / Direct Connect Transporter for Linux options are disabled if your current license does not support the MSP Direct Connect feature.

2. In the dialog box, set the installation options as follows:




- **Director hostname or IP** field: Enter the IP address or hostname of the machine on which the **MSP Director** is installed. This should be an external address that is reachable from the Internet.
- **Director port** field: Enter the **MSP Director port** used for communication with the **Direct Connect Transporter**. By default, port 4442 is filled in automatically, so make sure the port is open and not being used by other applications. This setting is global and can only be changed inside the master tenant **Settings** > **System Settings** > **Configuration** tab.
  You can specify an external port that is reachable from the Internet instead of default value, for example:
  - Public port 10005 for TCP Port 4442 of MSP communication.
- **Direct Connect user** dropdown: Click to select the user with the **Direct Connect** role. By default, the newest created user with the **Direct Connect** role is assigned.
  **Notes**
  - This user is used to generate an application password required for authentication when the **Direct Connect Transporter** from the client site connects to the **MSP Director**. This option is disabled if no users with the **Direct Connect** role exist in the tenant.
  - To create a new user with the **Direct Connect** role, click the **Create Direct Connect user** link, and the **Add Local User** wizard that opens proceeds as described in Adding Local Users.
  - When the wizard is opened this way, the default value of the **Role** dropdown is set to **Direct Connect** and can't be changed.

3. Click **Download** to start downloading the installer.

> **Notes**
>
> - The **Download** button is disabled if the required dropdown/fields are not filled in or contain invalid values.
> - Optionally, click the **Cancel** or **X** button to close the dialog box without applying any changes.

4. After the Direct Connect Transporter installer is downloaded, run the installer at the client's site and follow the steps below:

> **Notes**
>
> - For Windows, the installer UI and silent installations are supported.
> - For Linux, interactive and silent installations are supported.
> - Master Password is required for Direct Connect Transporter installation.

5. After the **Direct Connect Transporter** is installed successfully at the client's site, the **MSP Director** automatically adds the **Direct Connect Transporter** to the **Nodes** tab of the local tenant from which it was downloaded.

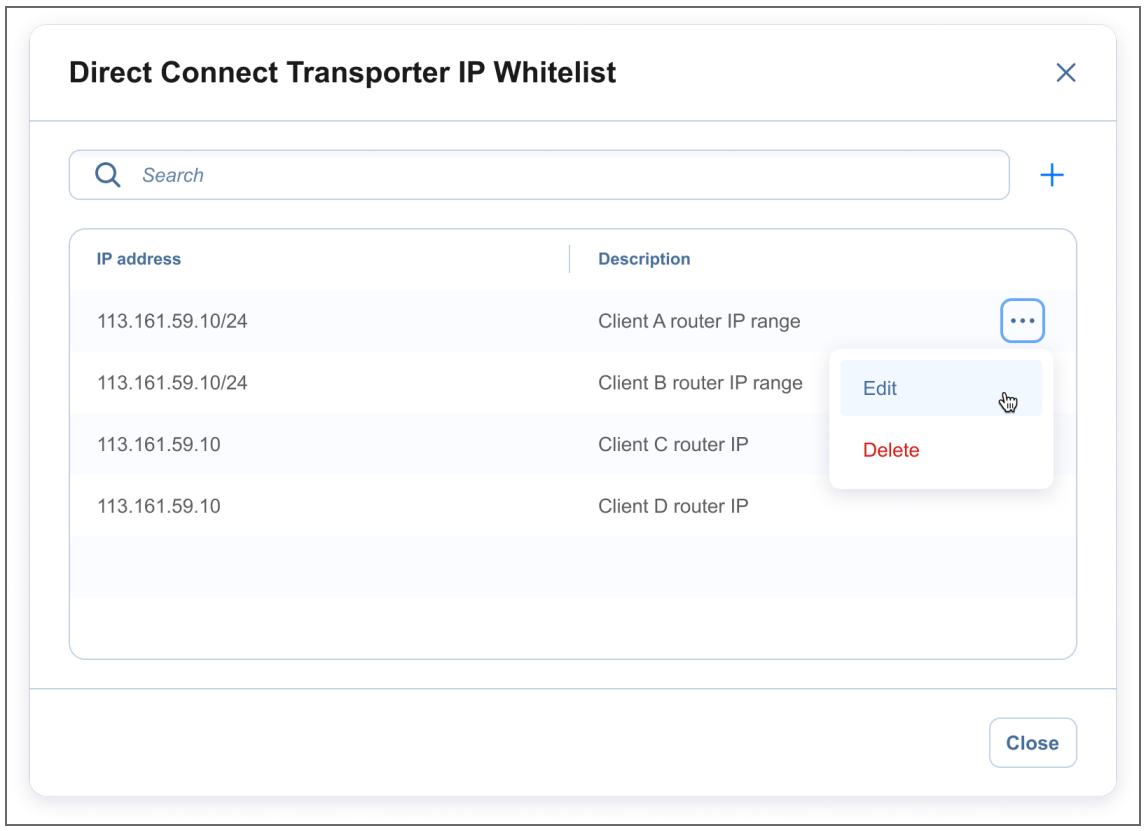
> **Note**
>
> If the installer is downloaded from the **Nodes** tab from the master tenant UI, the **Direct Connect Transporter** will be added there.

6. The added **Direct Connect Transporter** is displayed with the **Pending** status in the **Nodes** dashboard of the local tenant instance or master tenant, depending on where the installer was downloaded from.
7. Hover the Pending **Direct Connect Transporter**, click **Accept**, and enter the configured Master password to accept the **Direct Connect Transporter**. The client's Direct Connect Transporter is ready to be used for data protection activities.

**Notes**

- To back up items from the client's site to the MSP Repository, you must ensure at the MSP site, the **MSP Transporter** handling the backup must have Direct Connect enabled and be assigned to the MSP repository.
- You cannot add the **Direct Connect Transporter** manually.
- After the **Direct Connect Transporter** is installed, a config file is generated in the **Direct Connect Transporter** installation folder (**activedc.conf**). The **activedc.conf** file contains the following details about the connection to the **MSP Director**:
  - MSP Director hostname or IP
  - MSP Director port
  - Direct Connect user application password
  - MSP Director fingerprint
- The following permissions are required to edit the **activedc.conf** file on Linux:
  - Owner: bhsvc
  - Group: bhsvc
  - Permissions: 644
- If the **Direct Connect Transporter** does not appear in **MSP Director** after installation, check or modify this configuration file.

IP Whitelist Management



Only the **Direct Connect Transporters** from the specified IP addresses can connect to the **MSP Director** when enabled. If no whitelist is configured, any **Direct Connect Transporter** can connect to the **MSP Director** as long as its config file contains matching values (MSP Director IP, port, application password).
To do it, follow the steps below:

1. Go to the **General** > **System Settings** > **Configuration** tab.
2. Select the **Allow Direct Connect Transporter connections from specific IP addresses only** checkbox. The **Settings** link appears.
3. Click the **Settings** link.
4. In the **Direct Connect Transporter IP Whitelist** dialog box, click the **(+)** button to add the IP address to the whitelist.
5. Enter the required data and click **Add**. The IP address is added to the whitelist.

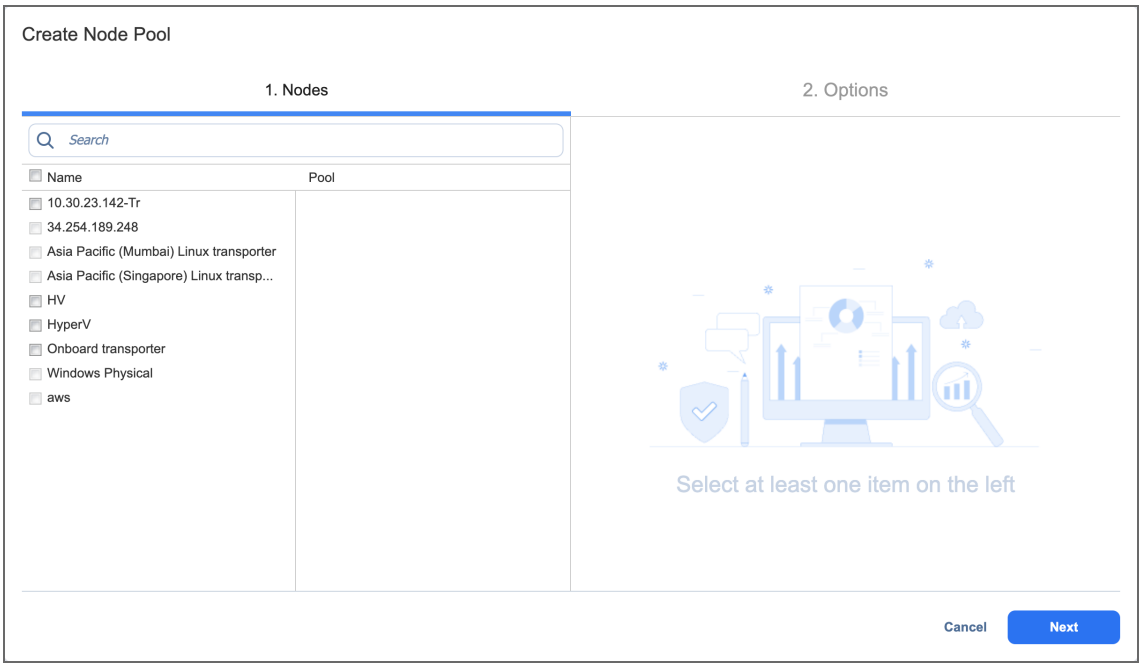
**Notes**

- If the IP address was on the whitelist and you removed it, the Direct Connect Transporter from that IP address remains connected to the MSP Director until you manually remove it from the Nodes tab.
- The whitelist can be configured at either the master tenant or the local tenant level, depending on where the Direct Connect Transporter will be added.

# Managing Node Pools

NAKIVO Backup & Replication allows you to group nodes into pools to optimize backup, replication, and recovery jobs. To create a node pool, do the following :

1. Navigate to **Settings**.
2. Click the **Nodes** tab.
3. Open the **Node Pools** tab, then click the plus **Add** button.
4. Complete the **Create Node Pool** wizard and click **Finish**.




A node pool can be selected in the *Data Transfer* section on the **Options** page of backup, replication, and recovery jobs. A node can be included in only one pool. To move a node from one pool to another, you need to remove it from the original pool first.

# Refreshing Node Details

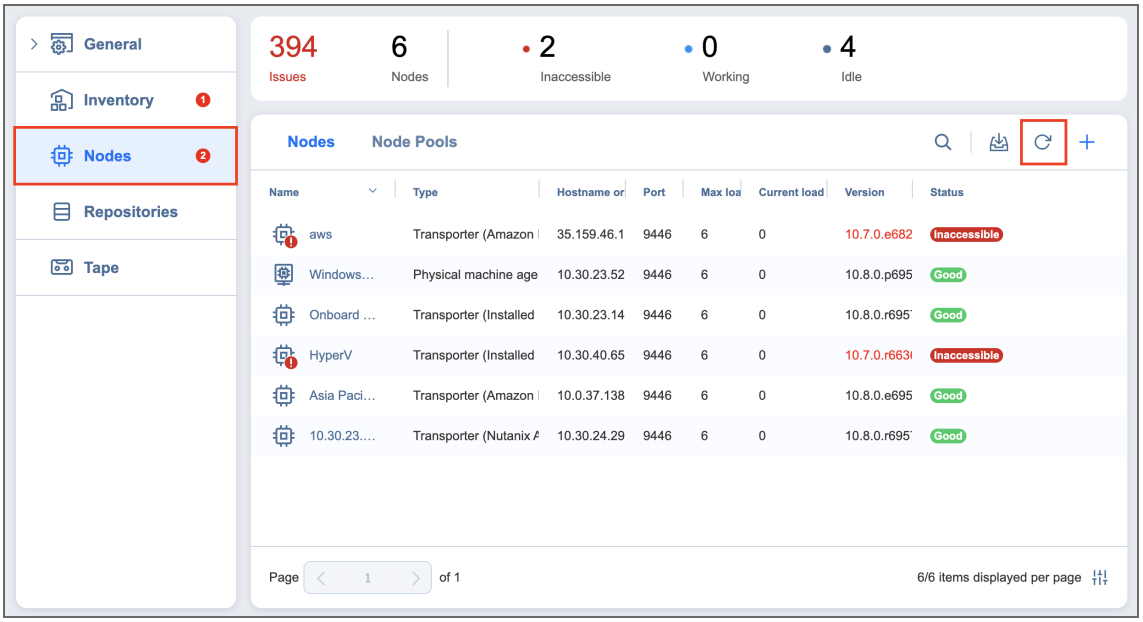
By default, NAKIVO Backup & Replication refreshes the information about **Transporters** every hour. During the refreshing process, the product collects all the required information about all transporters. Only one **Transporter** can be refreshed at a time. If you have more than one **Transporter**, all others will remain in the queue until they are able to be refreshed.

- Manually Refreshing All Nodes
- Manually Refreshing a Single Node

## Manually Refreshing All Nodes

To refresh all nodes, follow the steps below:

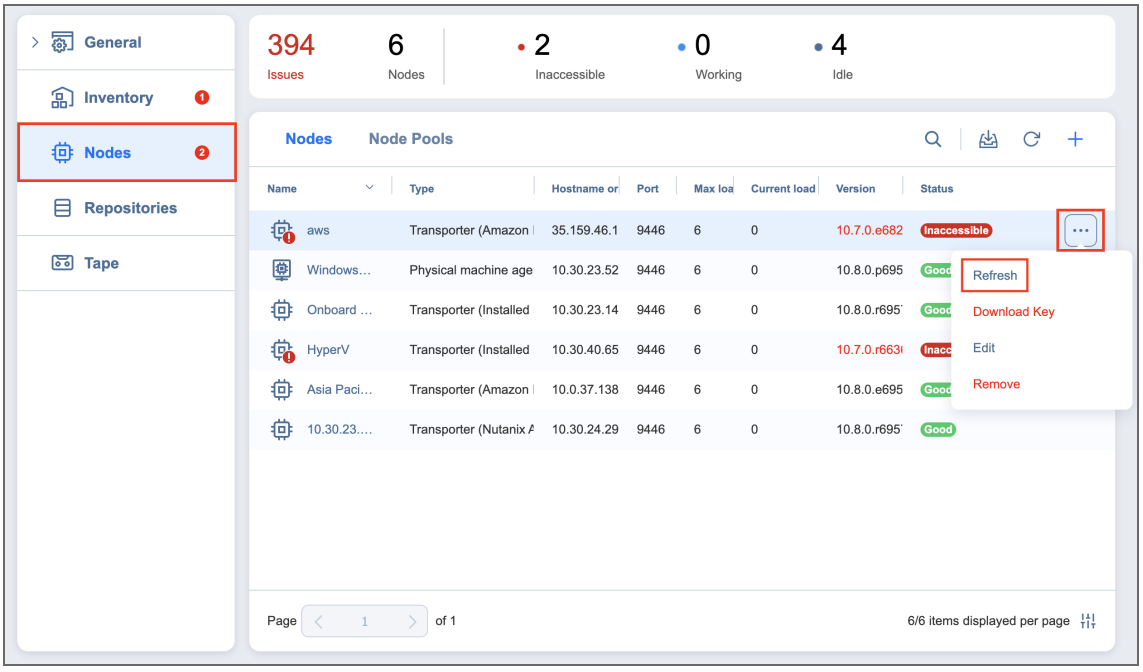
1. Click **Settings** in the left pane of the product and go to the **Nodes** tab.
2. Click the **Refresh** button above the **Nodes** table.




## Manually Refreshing a Single Node

To refresh a single node, follow the steps below:

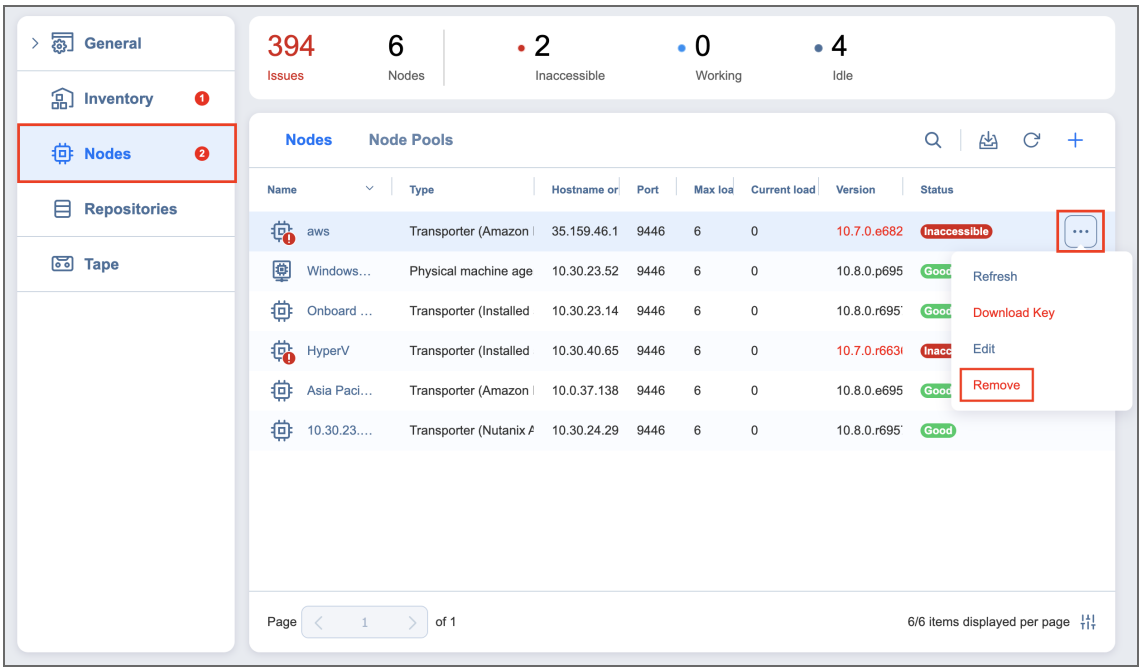
1. Click **Settings** in the left pane of the product.
2. Go to the **Nodes** tab.
3. Hover over the node you would like to refresh and click the ellipsis **Manage** button.
4. Click **Refresh**.




# Removing (Deleting) Nodes

To remove a Transporter from NAKIVO Backup & Replication, follow the steps below:

1. Click **Settings** in the left pane of the product.
2. Go to the **Nodes** tab.
3. Hover over the node you would like to remove.
4. On the right side, click the ellipsis **Manage** button and then click **Remove**.




**Note**

The following nodes cannot be removed:

- The Onboard **Transporter** (which is installed with the **"Director" on page 121** by default)
- Nodes manually assigned to a job
- **Transporters** assigned to backup repositories

# Using a VM Agent

With NAKIVO Backup & Replication, you can install a permanent virtual machine agent (VMA) rather than injecting a temporary agent for every job run. This agent simplifies OS quiescing and file-level recovery to the source virtual machine by eliminating the need to provide credentials for the VM's guest OS.

- The VM agent is used for OS quiescing if the relevant job option is enabled and only for Hyper-V. For other hypervisors, VMA is used only for file recovery.
- With the **File Restore through Permanent Agent** feature, you can manually select the proxy transporter to mount the backups and communicate with the auto-detected permanent virtual machine agent at the target destination for **File Level Recovery to Source** jobs. For more information, see Mounting a backup to a transporter.

For information on VMA system requirements, see “Feature Requirements” on page 169. For information on installing a VMA, see “Installing a VM Agent” on page 582.

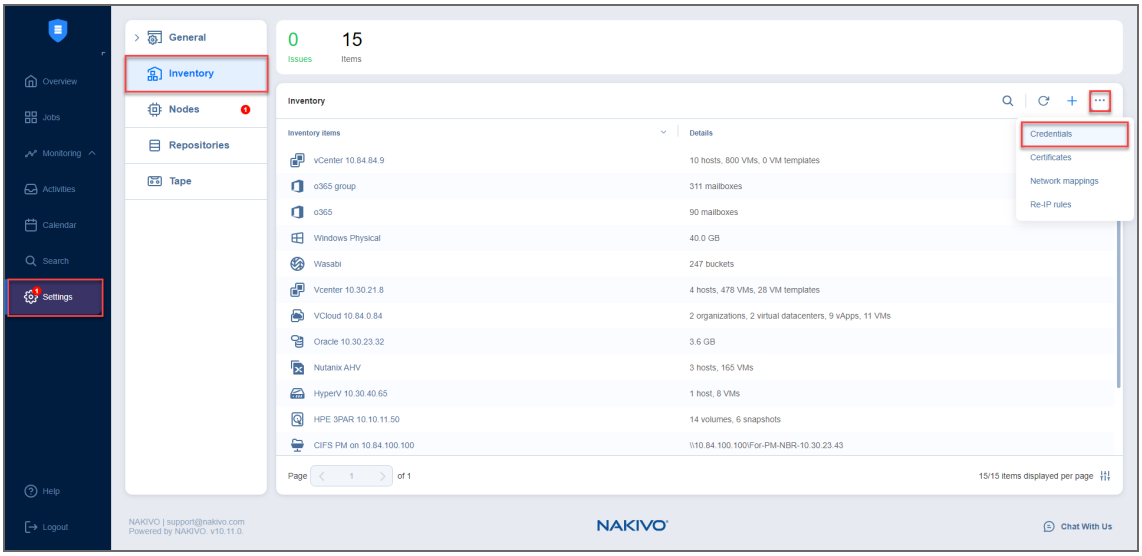
After installing a VM agent inside a VM and adding it to the **Nodes** tab, you may proceed with configuring it for use in jobs. See the sections below:

- Setting Default Credentials
- Enabling VM Agents
- “Updating VM Agents Manually” on page 607
- How a VM Agent Works

## Setting Default Credentials

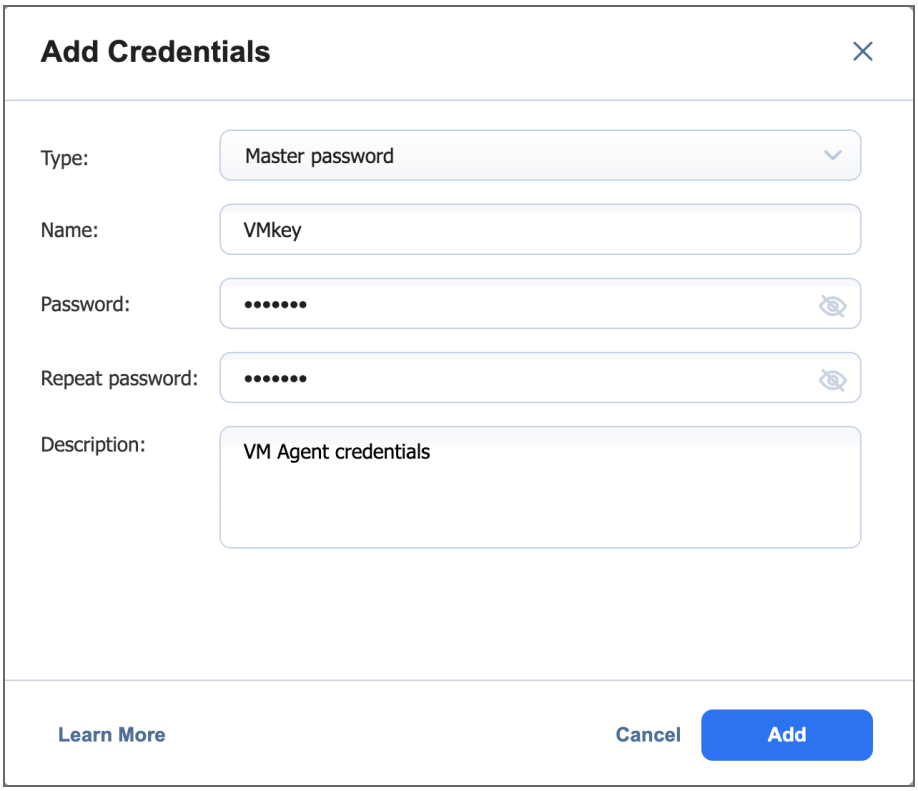
To configure a default master password for VM agents, do the following:

1. Go to **Settings** > **Inventory** > **Manage** > **Credentials**.




2. Click **Add Credentials**.
3. In the **Type** drop-down list, select **Master password**.
4. Enter a **Name** and **Password**.

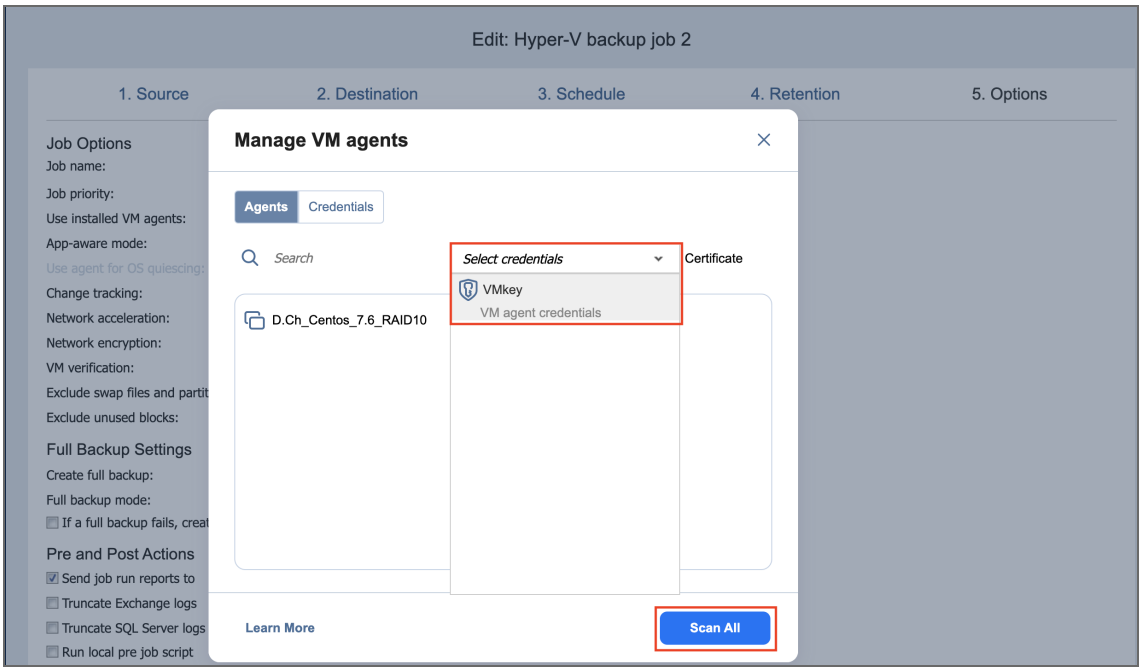
5. Optionally, add a **Description**.




## Enabling VM Agents

To enable your installed VM agents (VMAs) to be used in jobs, proceed as follows:

1. Click on a job containing the VM in which you have installed a permanent agent.
2. Go to **Manage** > **Edit** > **Options**.
3. Make sure that **Use installed VM** agents is enabled. Click **settings**.
4. Click **Scan All** to scan every VM in the job for a VM agent.
5. Once the scan is complete, select the master password you wish to use for the discovered VM agents from the **Select credentials** drop-down list.




6. In the **Certificate** column, click **Verify** to verify the validity of a VMA's certificate.
7. **Save** the updated job options.

Repeat this process for all jobs for which you want to enable the use of VMAs.

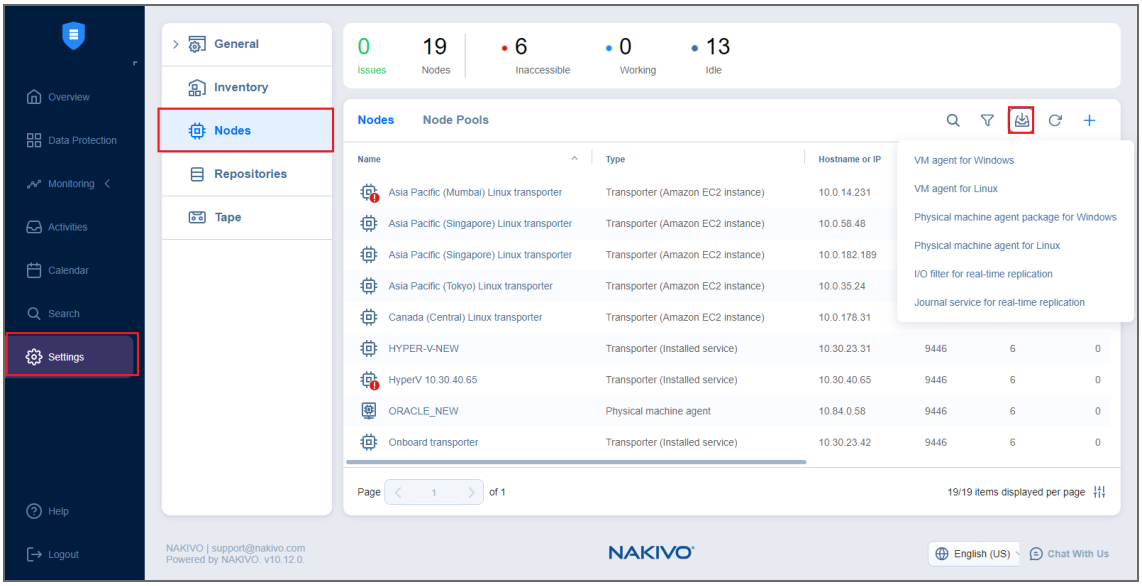
**Note**

- If a VM agent is installed on the recovery server when creating or editing a File Recovery job, the installed agent will automatically be chosen to perform recovery operations.
- If NAKIVO Backup & Replication cannot find installed VM agents or is prevented from doing so (such as by a firewall), the application will scan the VMs until a 2-minutes timeout. VMs are scanned in parallel.
- If there is no VM agent at the target VM, you may proceed by providing OS credentials to the target machine.

## Updating VM Agents Manually

To manually update a VM agent on a specific VM, proceed as follows:

1. Go to **Settings** > **Nodes**, click the **Download** button, and select to download the new VM agent build for a Window or Linux machine.




2. For silent update, start the agent installer on the machine where the VM agent resides:
   - For Windows OS, use `installer.exe -u --eula-accept`

> **Note**
>
> The following arguments can be used instead of `--eula-accept`
>
> - `--e`
> - `--eu`
> - `--eul`
> - `--eula`
> - `--eula-`
> - `--eula-a`
> - `--eula-ac`
> - `--eula-acc`
> - `--eula-acce`
> - `--eula-accep`
> - `--eula--accept`

   - For Linux OS, use `installer.sh -u --eula-accept`

The update success/failure is displayed in the appropriate log file.

## How a VM Agent Works

When supported VM actions are prompted, the **Director** checks for VM agent availability on the respective VM as follows:

1. If an installed VMA is found in the VM, the **Director** uses this VMA to perform the specified action(s).
2. If an installed VMA is not found in the VM, the **Director** injects a temporary agent or uses native tools (for example, VMware Tools) to perform the specified action(s). You must then enter VM credentials to proceed.
3. If the usage of an installed VMA is disabled in job options, the **Director** uses the injection approach or native tools to proceed.

# Backup Repositories

A **Backup Repository** is one of the key components of NAKIVO Backup & Replication and is a regular folder where the product stores backups and backup metadata. For more detailed information, refer to “Backup Repository” on page 127.

This section covers repository-related topics such as creation, management, etc. of **Backup Repositories** and contains the following articles:


- “Adding Existing Backup Repositories” on page 611
- “Creating Backup Repositories” on page 614
- “Managing Backup Repositories” on page 665
- “Viewing Backup Repository Details” on page 687

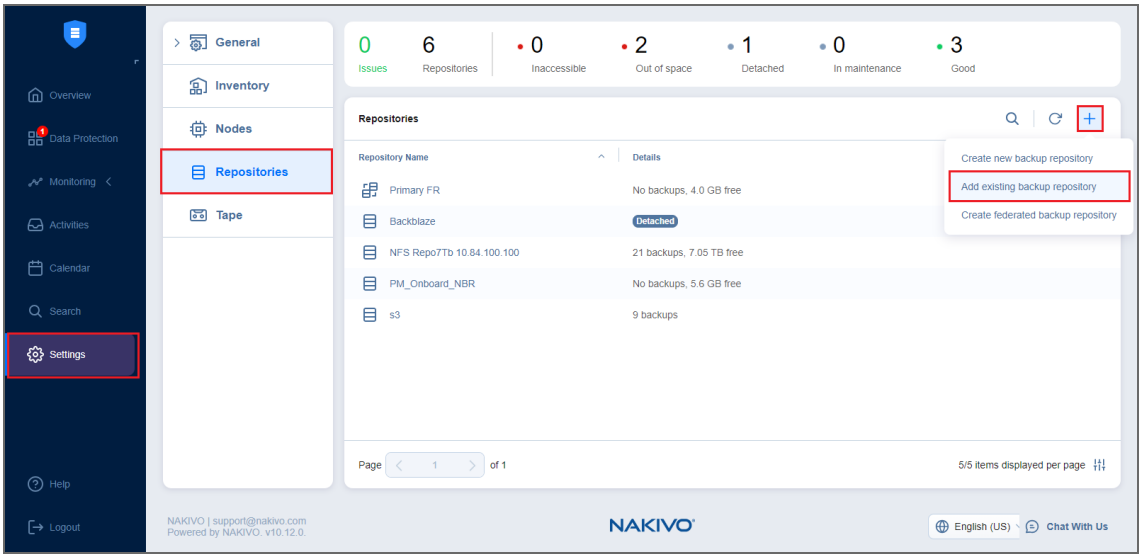
# Adding Existing Backup Repositories

NAKIVO Backup & Replication allows you to add an existing **Backup Repository** to a new copy of the product.

> **Note**
> During the import process, NAKIVO Backup & Replication searches for the *NakivoBackup* folder in the specified location. If your **Backup Repository** is located in *E:\backup\NakivoBackup*, you should specify the following path: *E:\backup*

To import an existing **Backup Repository**, do the following:

1. In the main menu, click **Settings**.
2. Go to the **Repositories** tab and click **+**.
3. Click **Add existing backup repository** in the dialog box that opens.




4. The **Add Existing Backup Repository** wizard opens. On the **Type** page of the wizard, select one of the following **Backup Repository** types:
   - **Local Folder**
   - **CIFS Share**
   - **NFS Share**
   - **Cloud & S3-Compatible Storage**
   - **Deduplication Appliance**
5. When you select **Cloud**, the **Vendor** page opens. Select the cloud storage vendor from the following options:
   - Amazon S3
   - Microsoft Azure
   - Wasabi

- Backblaze
- Amazon EC2
- Generic S3-compatible storage

6. When you select **Deduplication Appliance**, the **Device** page opens. Select the device from the following options:
   - Dell EMC Data Domain Boost
   - HPE StoreOnce Catalyst
   - NEC HYDRAstor
7. On the **Name & Location** page of the wizard, fill out all the necessary fields as described in the article for the corresponding **Backup Repository** type.
8. On the **Options** page of the wizard, depending on the repository type, the following options can be available for configuration:
   - **Encryption password**: If the **Backup Repository** is encrypted, type in the encryption password.
   - **Enable automatic repository self-healing**: Leave this option selected to automatically trigger repository self-healing in case the product detects symptoms of problems in the backup infrastructure (such as incorrect timestamps on metadata and data files). You can deselect this option and run self-healing manually.
   - **Run full data verification on schedule**: When this option is selected, the product runs full verification of all data available in the **Backup Repository** based on the specified schedule. The product reads each block of data to ensure that it is identical to the data block that was read on the source VM during the backup process. This way the product verifies each recovery point in the **Backup Repository**.

     > Backup verification is a time-consuming process and consumes CPU of the **Transporter** assigned to the **Backup Repository**. It is recommended to schedule backup verification during non-working hours.

   - **Run repository self-healing on schedule**: You can select this checkbox to additionally run repository self-healing based on a schedule. You can configure the schedule by clicking the **schedule** link when the option is selected. The default schedule is set to run every day at 11 AM.
   - **Reclaim unused space on schedule**: You can select this option to run the **Backup Repository** space reclaim process based on a schedule. You can configure the schedule by clicking the **schedule** link when the option is selected. The default schedule is set to run every Saturday at 12 PM.
   - **Enforce explicit file system sync**: When this option is selected, explicit sync with the file system is enforced during all backup operations to this repository. This setting is considered more reliable but may lead to lower performance on some storage devices.

- **Detach this repository on schedule**: Select this option if you want to detach and then reattach the **Backup Repository** based on a schedule. Detaching a **Backup Repository** saves its data and metadata in a consistent state and stops the product interaction with the **Backup Repository** (so that it can be copied or moved). You can use this feature, for example, for the disk-to-disk-to-tape (D2D2T) data protection approach: backups are stored on a disk for fast operational recovery and copied to a tape (while the repository is detached) for archiving and long-term storage.
- **Delete and re-create the repository on attach**: When this option is selected, all data in the **Backup Repository** is erased prior to attaching it to the product. As a result, jobs that write to this **Backup Repository** create full VM backups. You can use this option, for example, to create full daily, weekly, or monthly VM backups and write them to tape or removable media.

9. Click **Finish**. The **Backup Repository** is imported to the list.

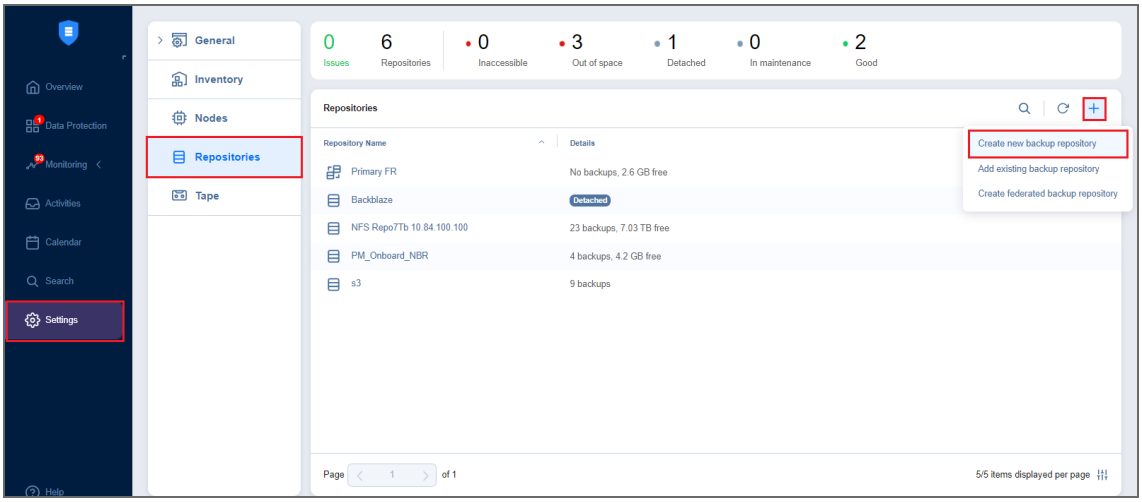
# Creating Backup Repositories

NAKIVO Backup & Replication allows you to create additional **Backup Repositorie**s for storing backups. You can use a local folder, NFS share, CIFS share, public cloud, or deduplication appliance as a **Backup Repository** location. To create a new **Backup Repository**, follow the steps below.

**Important**

Do not create **Backup Repositorie**s inside NAKIVO Backup & Replication installation folders. The data inside **Director** and **Transporter** folders may be lost after a solution update.

1. In NAKIVO Backup & Replication, navigate to **Settings**.
2. Go to the **Repositories** tab and click **+**.
3. Click **Create new backup repository**.




Choose one of the locations for storing your backups by completing the **Create Backup Repository** wizard as described in the sections below:


- "Local Backup Repository" on page 615
- "Backup Repository on CIFS Share" on page 620
- "Backup Repository on NFS Share" on page 625
- "Backup Repository in Amazon EC2" on page 630
- "Backup Repository in Amazon S3" on page 636
- "Backup Repository in Generic S3-Compatible Object Storage " on page 640
- "Backup Repository in Microsoft Azure Blob Storage" on page 644
- "Backup Repository in Backblaze B2 Cloud Storage" on page 648
- "Backup Repository in Wasabi Hot Cloud Storage" on page 653
- "Backup Repository on Deduplication Appliance" on page 657

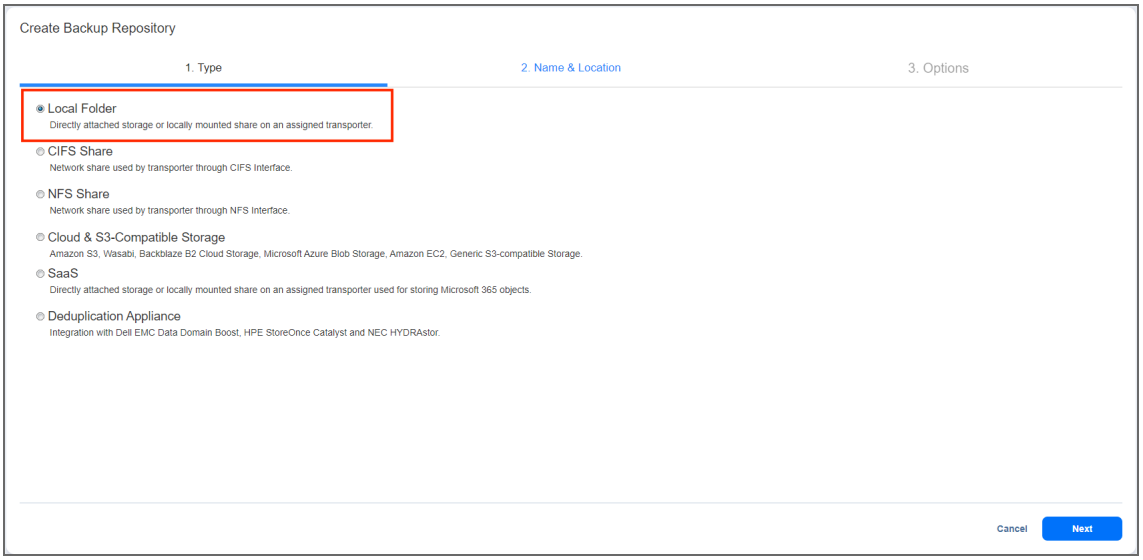
# Local Backup Repository

To create a **Backup Repository** locally on the machine on which the assigned **Transporter** is installed, choose a local folder. Proceed as described in the following sections:

- Create Backup Repository: Type
- Create Backup Repository: Name and Location
- Create Backup Repository: Options

## Create Backup Repository: Type

On the **Type** page of the **Create Backup Repository** wizard, select **Local Folder** and click **Next** to move to the next page of the wizard.




## Create Backup Repository: Name and Location

On the **Name & Location** page of the wizard, do the following:

1. Enter the name of the **Backup Repository** in the **Name** box.
2. Select the **Transporter** from the **Assigned transporter** drop-down list.
3. Enter the path to the local **Backup Repository** folder on the machine on which the assigned **Transporter** is installed.

   Below are the default paths, you can select another location/storage,make sure the read/write permissions are available.

   - For QNAP:

     ```
     /share/CACHEDEV1_DATA/.qpkg/NBR/
     ```

- For Synology:

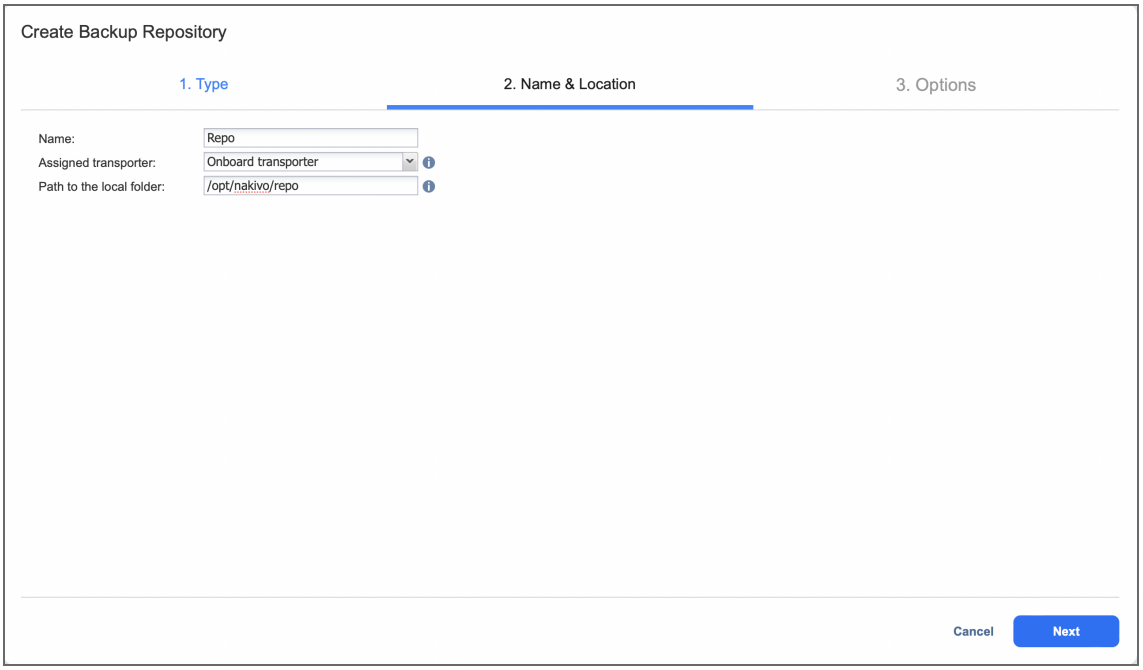
```
/"volumenumber"/@appstore/NBR
```

- For Linux:

```
/opt/nakivo/repository
```

4. Click **Next** to go to the next page of the wizard.

**Important**
Before choosing this location, make sure that you have read and write permissions for the folder that will be used as a repository. If needed, check and set the correct permissions using the steps outlined in the Backup Repository: Insufficient Permissions article.




## Create Backup Repository: Options

On the **Options** page, do the following:

1. Set up *Storage Savings & Encryption* options:
   - **Data size reduction**: If this option is enabled, NAKIVO Backup & Replication enables the use of data size reduction for this repository to save disk space. Note that this may put additional load on the CPU. Disabling data size reduction is required if the target is a deduplication storage appliance. Click settings to configure the settings. A popup window appears. Set the following:

- **Compression**: Select a compression level that will be used to reduce the data size in the **Backup Repository**. Note that higher compression levels consume considerably more CPU and may slow down VM backup speed. The following options are available:
  - **Disabled:** The data in the **Backup Repository** will not be compressed.
  - **Fast:** Lowest compression level.
  - **Medium:** Medium compression level.
  - **Best:** Maximum compression level.

  > **Note**
  > This option cannot be configured after creating the **Backup Repository**.

- **Store backups in separate files**: Select this option to enable this backup repository to store data of every machine in separate backup files. Enabling this option is highly recommended to ensure higher reliability and performance. Leave this option unchecked if you wish to enable deduplication on a given backup repository.
- **Deduplication**: Select this option to enable the backup deduplication method to reduce the backup size by excluding duplicate data blocks from the backup.

  > **Note**
  > This option is not available if the Store backups in separate files checkbox has been selected.

- **Encryption:** This option is available only if the **Backup Repository** is created locally on the machine on which the Assigned **Transporter** is installed, and the machine is running a Linux OS. Select **Enabled** from the drop-down list and specify an encryption password. (The password will be required for importing the **Backup Repository** into a new instance of the product.) The product will encrypt the repository destination (using `ecryptfs` for folders and `cryptsetup (crypt-md)` in LUKS mode for devices/partitions) prior to creating the **Backup Repository**.

  > **Notes**
  > - To avoid `ecryptfs` errors, make sure that there are no other folders and files except the NakivoBackup folder in the repository location.
  > - **Backup Repository** encryption can significantly influence backup speed.

2. Set up *Reliability & Maintenance* options:
   - **Enable automatic repository self-healing:** Leave this option selected to automatically trigger repository self-healing in case the product detects symptoms of problems in the backup infrastructure such as incorrect timestamps on metadata and data files. You can deselect this option and run self-healing manually.

- **Run repository self-healing on schedule:** If required, select this checkbox to run repository self-healing on schedule. You can configure the schedule by clicking the **schedule** link when the option is selected. The default schedule is set to run every day at 11 AM.
  If **Stop backup and recovery to run self-healing** is selected, any jobs or recoveries which use this repository will be stopped to run scheduled self-healing. Otherwise, scheduled self-healing will be skipped in case there are running jobs or recoveries on this repository.
- **Run full data verification on schedule:** If selected, NAKIVO Backup & Replication will run full verification of all data available in the Backup Repository on the specified schedule. The product will read each block of data and ensure that it is identical to the data block that was read on the source VM during the backup. This way, the product will verify each recovery points in the **Backup Repository**.
  If **Stop backup and recovery to run backup verification** is selected, any running jobs which use this **Backup Repository** will be stopped to run scheduled data verification. Otherwise, scheduled data verification will be skipped in case there are running jobs on this **Backup Repository**.

  > **Note**
  > Backup verification is a time-consuming process and consumes CPU of the **Transporter** assigned to the **Backup Repository**. It is recommended that you schedule backup verification during non-working hours.

- **Reclaim unused space on schedule**: If required, select this option to run the **Backup Repository** space reclaim process on schedule. Space reclaim will compact the data. Unused space will be reclaimed. Keep in mind that this process can be time-consuming.

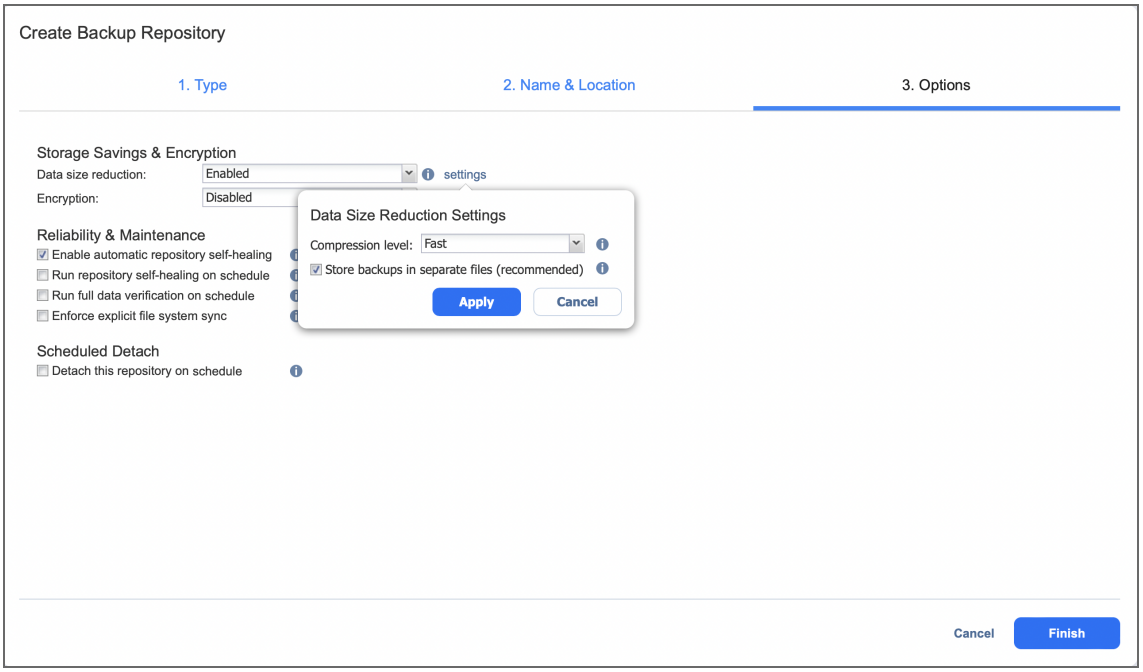
  > **Note**
  > This option is available only if **Store backups in separate files** is not enabled.

  If **Stop backup and recovery to run space reclaim** is selected, any running jobs which use this **Backup Repository** will be stopped to run scheduled space reclaiming. Otherwise, scheduled space reclaiming will be skipped in case there are running jobs on this **Backup Repository**.

  > **Important**
  > Do not reboot/disconnect the "null" **Transporter** and storage device while space reclaim is in progress to avoid **Backup Repository** corruption.

- **Enforce explicit file system sync**: When selected, explicit sync with the file system is enforced during all backup operations to this repository. This setting is considered more reliable but may lead to lower performance on certain storage devices. By default, the option is disabled.

4. Schedule detaching of the **Backup Repository**:
   - **Detach this repository on schedule:** Select this option if you want to detach and then attach the **Backup Repository** on a schedule. Detaching a **Backup Repository** saves its data and metadata in a consistent state and then stops the product's interaction with the **Backup Repository** (so that it can be copied or moved). You can use this feature, for example, for the disk-to-disk-to-tape (D2D2T) data protection approach, in which backups are stored on a disk for fast operational recovery, and copied to a tape (while the repository is detached) for archiving and long-term storage.
   - **Delete and re-create the repository on attach:** If this option is selected, all data in the **Backup Repository** will be erased prior to attaching it to the product. As a result, jobs that write to this **Backup Repository** will create full VM backups. You can use this option, for example, to create full daily, weekly, or monthly VM backups and write them to tape or removable media.
5. Click **Finish** to finish creating the **Backup Repository**.



# Backup Repository on CIFS Share

Choose this option if you want to create a **Backup Repository** on a Windows CIFS share. Before creating a **Backup Repository** on a CIFS share, make sure that all the necessary prerequisites are met:

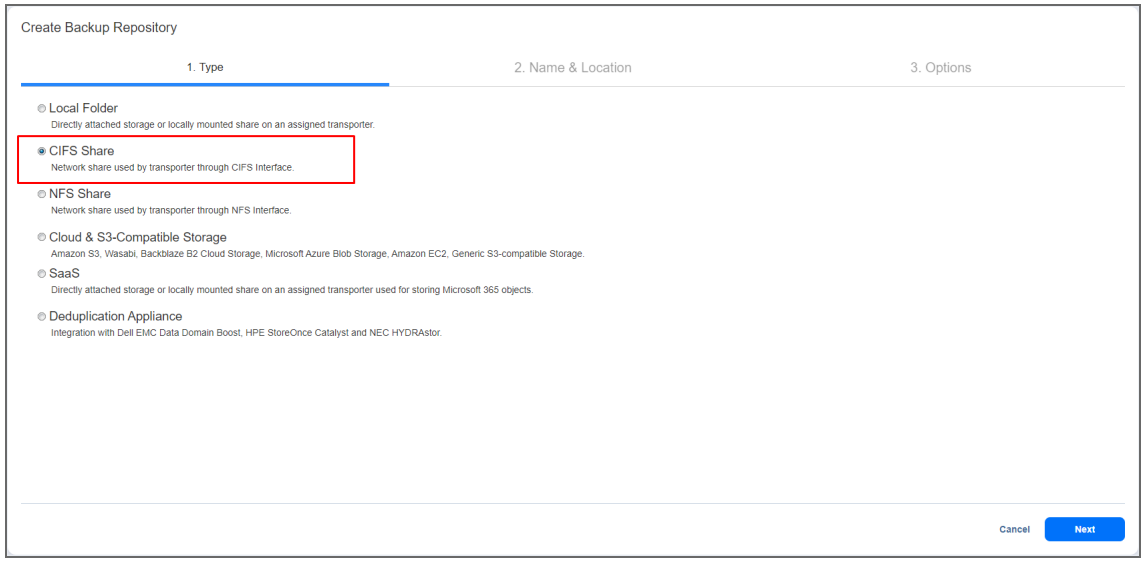
- The folder where you would like to create the **Backup Repository** exists on the share.
- The share can be accessed from the machine on which the Assigned **Transporter** is installed.
- You are using credentials with read and write permissions to the share.
- The share is compatible with Version 2 or later of the SMB protocol.

To create a **Backup Repository** on a Windows CIFS share, proceed as described in the following sections:

- Create Backup Repository: Type
- Create Backup Repository: Name and Location
- Create Backup Repository: Options

## Create Backup Repository: Type

On the **Type** page of the **Create Backup Repository** wizard, select **CIFS Share** and click **Next** to move to the next page of the wizard.




## Create Backup Repository: Name and Location

On the **Name & Location** page of the wizard, do the following:

1. Enter the name of the **Backup Repository** in the **Name** box.
2. Select the **Transporter** from the Assigned transporter drop-down list.
3. Enter the path to the CIFS share.

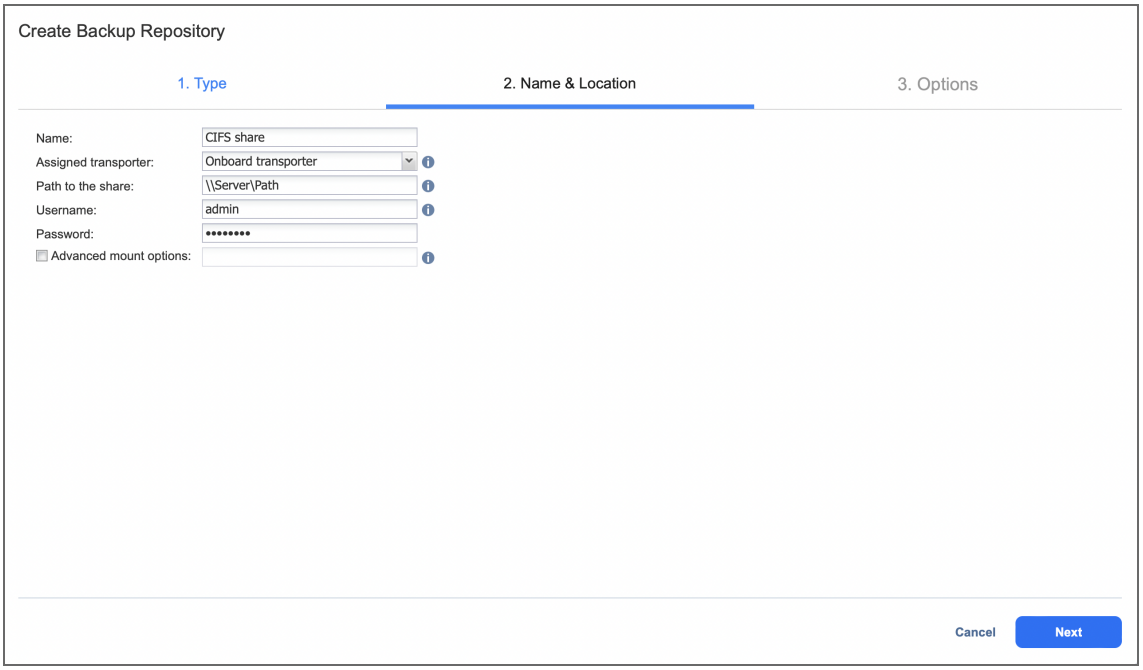
   **Example**

   Synology share path: `\\10.30.30.61\ayunt_cifs1`
4. Provide username and password in the appropriate boxes.

**Note**

If you're using domain credentials to access the share, enter your domain username via the following format: `domain\username`.

5. Select **Advanced mount options** if needed. Refer to the mount `man` pages for a detailed description of CIFS share mount options.
6. Click **Next** to go to the next page of the wizard.




## Create Backup Repository: Options

On the **Options** page, do the following:

1. Set up *Storage Savings & Encryption* options:
    - **Data size reduction**: If this option is enabled, NAKIVO Backup & Replication enables the use of data size reduction for this repository to save disk space. Note that this may put additional load on the CPU. Disabling data size reduction is required if the target is a deduplication storage appliance. Click settings to configure the settings. A popup window appears. Set the following:
        - **Compression**: Select a compression level that will be used to reduce the data size in the **Backup Repository**. Note that higher compression levels consume considerably more CPU and may slow down VM backup speed. The following options are available:
            - **Disabled:** The data in the **Backup Repository** will not be compressed.
            - **Fast:** Lowest compression level.
            - **Medium:** Medium compression level.
            - **Best:** Maximum compression level.

> **Note**
> This option cannot be configured after creating the **Backup Repository**.

- **Store backups in separate files**: Select this option to enable this backup repository to store data of every machine in separate backup files. Enabling this option is highly recommended to ensure higher reliability and performance. Leave this option unchecked if you wish to enable deduplication on a given backup repository.
- **Deduplication**: Select this option to enable the backup deduplication method to reduce the backup size by excluding duplicate data blocks from the backup.

> **Note**
> This option is not available if the Store backups in separate files checkbox has been selected.

- **Encryption:** This option is available only if the **Backup Repository** is created locally on the machine on which the Assigned **Transporter** is installed, and the machine is running a Linux OS. Select **Enabled** from the drop-down list and specify an encryption password. (The password will be required for importing the **Backup Repository** into a new instance of the product.) The product will encrypt the repository destination (using `ecryptfs` for folders and `cryptsetup (crypt-md)` in LUKS mode for devices/partitions) prior to creating the **Backup Repository**.

> **Notes**
> - To avoid `ecryptfs` errors, make sure that there are no other folders and files except the NakivoBackup folder in the repository location.
> - **Backup Repository** encryption can significantly influence backup speed.

2. Set up *Reliability & Maintenance* options:
   - **Enable automatic repository self-healing:** Leave this option selected to automatically trigger repository self-healing in case the product detects symptoms of problems in the backup infrastructure such as incorrect timestamps on metadata and data files. You can deselect this option and run self-healing manually.
   - **Run repository self-healing on schedule:** If required, select this checkbox to run repository self-healing on schedule. You can configure the schedule by clicking the **schedule** link when the option is selected. The default schedule is set to run every day at 11 AM.
     If **Stop backup and recovery to run self-healing** is selected, any jobs or recoveries which use this repository will be stopped to run scheduled self-healing. Otherwise, scheduled self-healing will be skipped in case there are running jobs or recoveries on this repository.

- **Run full data verification on schedule:** If selected, NAKIVO Backup & Replication will run full verification of all data available in the **Backup Repository** on the specified schedule. The product will read each block of data and ensure that it is identical to the data block that was read on the source VM during the backup. This way, the product will verify each recovery points in the Backup Repository.
  If **Stop backup and recovery to run backup verification** is selected, any running jobs which use this **Backup Repository** will be stopped to run scheduled data verification. Otherwise, scheduled data verification will be skipped in case there are running jobs on this Backup Repository.

  > **Note**
  > Backup verification is a time-consuming process and consumes CPU of the **Transporter** assigned to the **Backup Repository**. It is recommended that you schedule backup verification during non-working hours.

- **Reclaim unused space on schedule**: If required, select this option to run the **Backup Repository** space reclaim process on schedule. Space reclaim will compact the data. Unused space will be reclaimed. Keep in mind that this process can be time-consuming.

  > **Note**
  > This option is available only if **Store backups in separate files** is not enabled.

  If **Stop backup and recovery to run space reclaim** is selected, any running jobs which use this **Backup Repository** will be stopped to run scheduled space reclaiming. Otherwise, scheduled space reclaiming will be skipped in case there are running jobs on this **Backup Repository**.

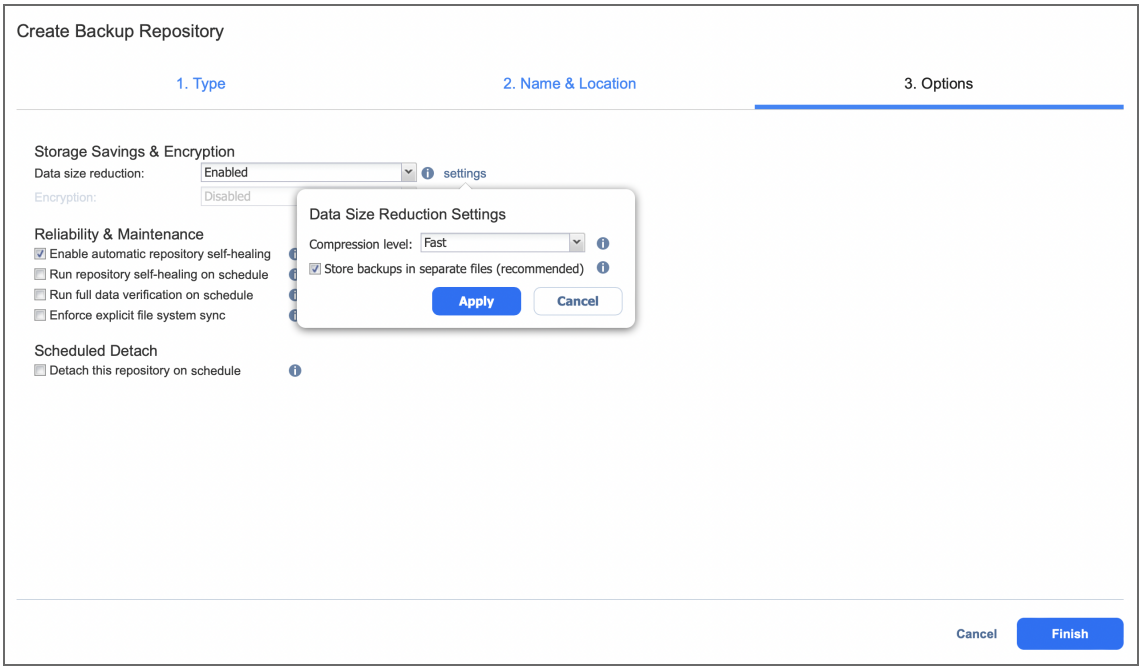
  > **Important**
  > Do not reboot/disconnect the "null" **Transporter** and storage device while space reclaim is in progress to avoid **Backup Repository** corruption.

- **Enforce explicit file system sync**: When selected, explicit sync with the file system is enforced during all backup operations to this repository. This setting is considered more reliable but may lead to lower performance on certain storage devices. By default, the option is disabled.

4. Schedule detaching of the **Backup Repository**:
   - **Detach this repository on schedule:** Select this option if you want to detach and then attach the **Backup Repository** on a schedule. Detaching a **Backup Repository** saves its data and metadata in a consistent state and then stops the product's interaction with the **Backup Repository** (so that it can be copied or moved). You can use this feature, for example, for the disk-to-disk-to-tape (D2D2T) data protection approach, in which backups are stored on a disk for fast operational recovery, and copied to a tape (while the repository is detached) for archiving and long-term storage.

- **Delete and re-create the repository on attach:** If this option is selected, all data in the **Backup Repository** will be erased prior to attaching it to the product. As a result, jobs that write to this **Backup Repository** will create full VM backups. You can use this option, for example, to create full daily, weekly, or monthly VM backups and write them to tape or removable media.

5. Click **Finish** to finish creating the **Backup Repository**.



# Backup Repository on NFS Share

Choose this option if you wish to create a **Backup Repository** on an NFS share. Before creating a **Backup Repository** on an NFS share, make sure that all the necessary prerequisites are met:

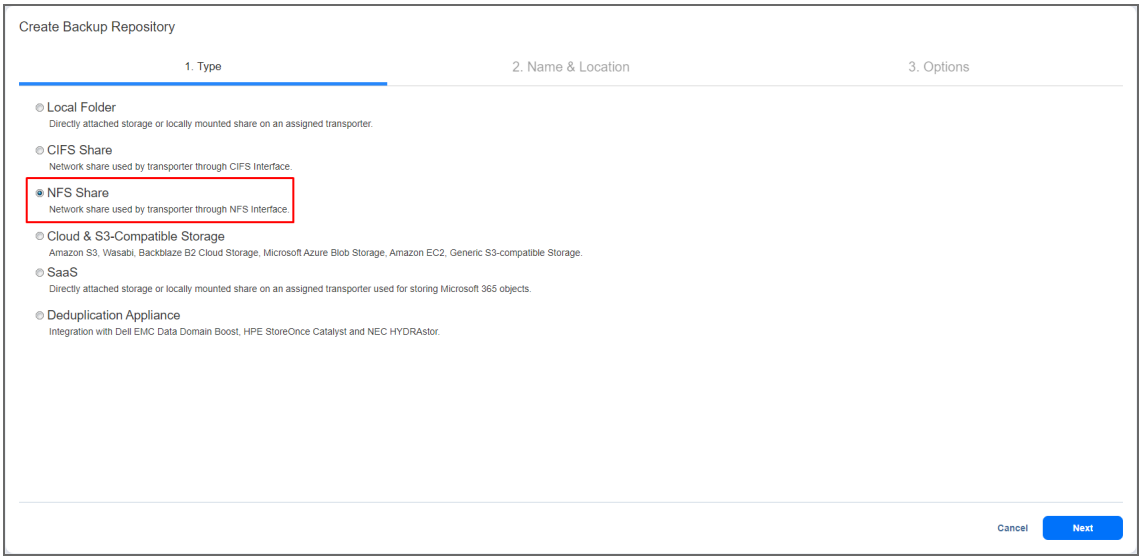
- The folder where you would like to create the **Backup Repository** exists on the share.
- The share can be accessed from the machine on which the Assigned **Transporter** is installed.
- You are using credentials with read and write permissions to the share.

To create a repository on an NFS share, proceed as described in the following sections:

- Create Backup Repository: Type
- Create Backup Repository: Name and Location
- Create Backup Repository: Options

## Create Backup Repository: Type

On the **Type** page of the **Create Backup Repository** wizard, select **NFS Share** and click **Next** to move to the next page of the wizard.




## Create Backup Repository: Name and Location

On the **Name & Location** page of the wizard, do the following:

1. Enter the name of the **Backup Repository** in the **Name** box.
2. Select the **Transporter** from the **Assigned transporter** drop-down list.
3. Enter the path to the NFS share.

   **Examples**

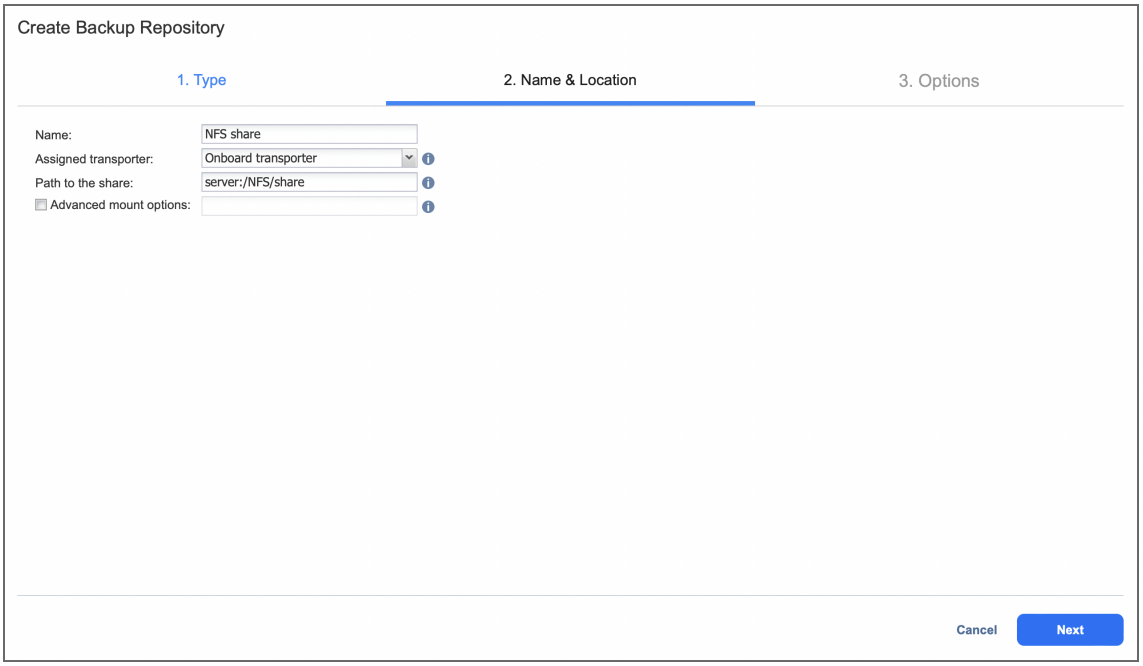
   QNAP share path: `10.30.30.109:/ayunt_nfs`

> **Note**
> If the Assigned **Transporter** is installed on a Windows OS, you need to enable the "Client for NFS" feature on the machine on which the **Transporter** is installed.

4. Select **Advanced mount options** if needed. Refer to the `mount` man pages for a detailed description of mount options.

> **Note**
> To create a **Backup Repository** on a NEC HydraStor deduplication appliance, refer to Integrating with NEC HydraStor.

5. Click **Next** to go to the next page of the wizard.




## Create Backup Repository: Options

On the **Options** page, do the following:

1. Set up *Storage Savings & Encryption* options:
   - **Data size reduction**: If this option is enabled, NAKIVO Backup & Replication enables the use of data size reduction for this repository to save disk space. Note that this may put additional load on the CPU. Disabling data size reduction is required if the target is a deduplication storage appliance. Click settings to configure the settings. A popup window appears. Set the following:
     - **Compression**: Select a compression level that will be used to reduce the data size in the **Backup Repository**. Note that higher compression levels consume considerably more CPU and may slow down VM backup speed. The following options are available:

- **Disabled:** The data in the **Backup Repository** will not be compressed.
- **Fast:** Lowest compression level.
- **Medium:** Medium compression level.
- **Best:** Maximum compression level.

> **Note**
> This option cannot be configured after creating the **Backup Repository**.

- **Store backups in separate files**: Select this option to enable this backup repository to store data of every machine in separate backup files. Enabling this option is highly recommended to ensure higher reliability and performance. Leave this option unchecked if you wish to enable deduplication on a given backup repository.
- **Deduplication**: Select this option to enable the backup deduplication method to reduce the backup size by excluding duplicate data blocks from the backup.

> **Note**
> This option is not available if the Store backups in separate files checkbox has been selected.

- **Encryption:** This option is available only if the **Backup Repository** is created locally on the machine on which the Assigned **Transporter** is installed, and the machine is running a Linux OS. Select **Enabled** from the drop-down list and specify an encryption password. (The password will be required for importing the **Backup Repository** into a new instance of the product.) The product will encrypt the repository destination (using `ecryptfs` for folders and `cryptsetup (crypt-md)` in LUKS mode for devices/partitions) prior to creating the **Backup Repository**.

> **Notes**
> - To avoid `ecryptfs` errors, make sure that there are no other folders and files except the NakivoBackup folder in the repository location.
> - **Backup Repository** encryption can significantly influence backup speed.

2. Set up *Reliability & Maintenance* options:
   - **Enable automatic repository self-healing:** Leave this option selected to automatically trigger repository self-healing in case the product detects symptoms of problems in the backup infrastructure such as incorrect timestamps on metadata and data files. You can deselect this option and run self-healing manually.

- **Run repository self-healing on schedule:** If required, select this checkbox to run repository self-healing on schedule. You can configure the schedule by clicking the **schedule** link when the option is selected. The default schedule is set to run every day at 11 AM.
  If **Stop backup and recovery to run self-healing** is selected, any jobs or recoveries which use this repository will be stopped to run scheduled self-healing. Otherwise, scheduled self-healing will be skipped in case there are running jobs or recoveries on this repository.
- **Run full data verification on schedule:** If selected, NAKIVO Backup & Replication will run full verification of all data available in the **Backup Repository** on the specified schedule. The product will read each block of data and ensure that it is identical to the data block that was read on the source VM during the backup. This way, the product will verify each recovery points in the **Backup Repository**.
  If **Stop backup and recovery to run backup verification** is selected, any running jobs which use this Backup Repository will be stopped to run scheduled data verification. Otherwise, scheduled data verification will be skipped in case there are running jobs on this **Backup Repository**.

  **Note**
  Backup verification is a time-consuming process and consumes CPU of the **Transporter** assigned to the **Backup Repositor**y. It is recommended that you schedule backup verification during non-working hours.

- **Reclaim unused space on schedule**: If required, select this option to run the **Backup Repository** space reclaim process on schedule. Space reclaim will compact the data. Unused space will be reclaimed. Keep in mind that this process can be time-consuming.

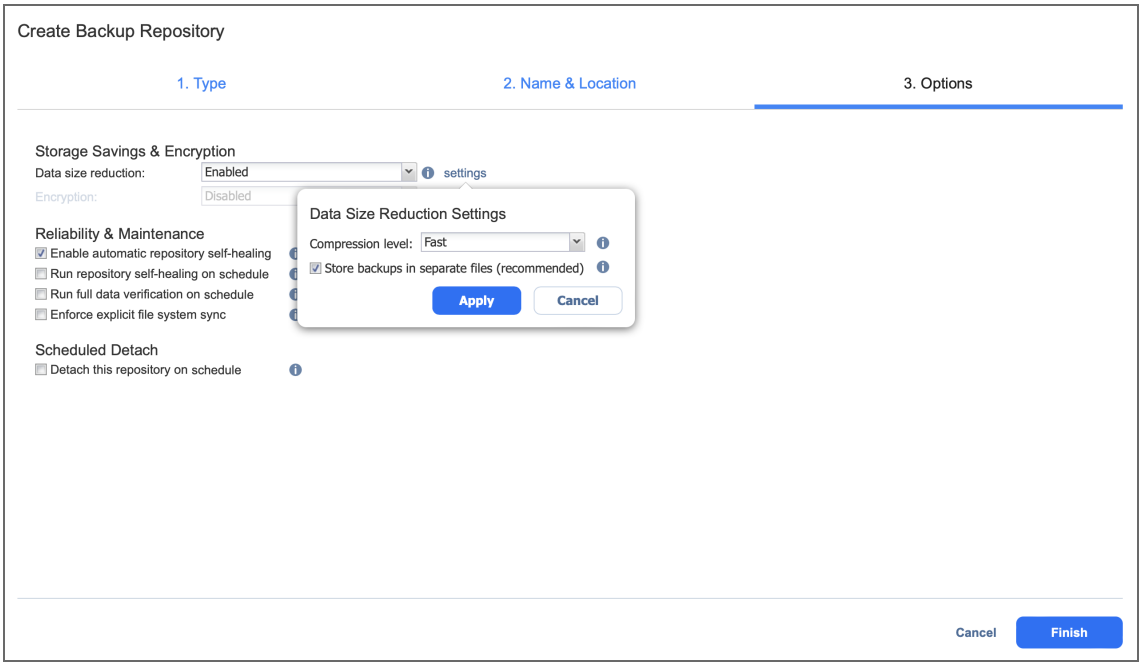
  **Note**
  This option is available only if **Store backups in separate files** is not enabled.

  If **Stop backup and recovery to run space reclaim** is selected, any running jobs which use this Backup Repository will be stopped to run scheduled space reclaiming. Otherwise, scheduled space reclaiming will be skipped in case there are running jobs on this **Backup Repository**.

  **Important**
  Do not reboot/disconnect the "null" **Transporter** and storage device while space reclaim is in progress to avoid **Backup Repository** corruption.

- **Enforce explicit file system sync**: When selected, explicit sync with the file system is enforced during all backup operations to this repository. This setting is considered more reliable but may lead to lower performance on certain storage devices. By default, the option is disabled.

4. Schedule detaching of the Backup Repository:
   - **Detach this repository on schedule:** Select this option if you want to detach and then attach the **Backup Repository** on a schedule. Detaching a **Backup Repository** saves its data and metadata in a consistent state and then stops the product's interaction with the **Backup Repository** (so that it can be copied or moved). You can use this feature, for example, for the disk-to-disk-to-tape (D2D2T) data protection approach, in which backups are stored on a disk for fast operational recovery, and copied to a tape (while the repository is detached) for archiving and long-term storage.
   - **Delete and re-create the repository on attach:** If this option is selected, all data in the **Backup Repository** will be erased prior to attaching it to the product. As a result, jobs that write to this **Backup Repository** will create full VM backups. You can use this option, for example, to create full daily, weekly, or monthly VM backups and write them to tape or removable media.
5. Click **Finish** to finish creating the **Backup Repository**.



# Backup Repository in Amazon EC2

Choose this option if you want to create a **Backup Repository** in Amazon EC2. The **Backup Repository** will be created in the same region where the assigned Transporter is located.

**Important**

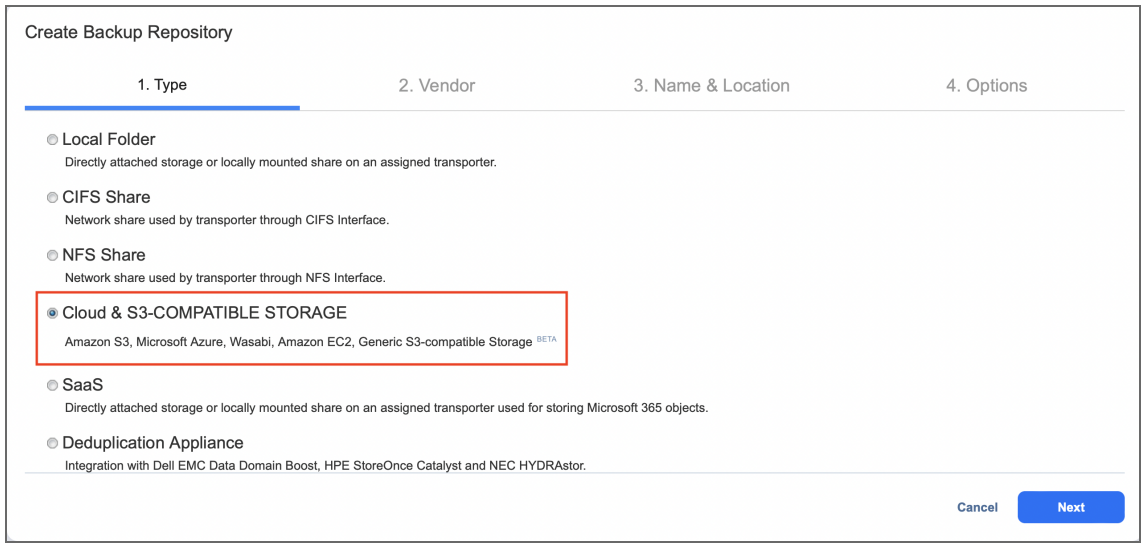
- To avoid disrupting NAKIVO Backup & Replication processes and data corruption, add NAKIVO Backup & Replication to the white/exclusions list of antivirus software running on the machine on which the NAKIVO **Backup Repository** is set up.
- You may be additionally charged for using a third-party resource. Refer to the third-party resource provider documentation for details.

To create a repository in Amazon EC2, proceed as described in the following sections:

- Create Backup Repository: Type
- Create Backup Repository: Vendor
- Create Backup Repository: Name & Location
- Create Backup Repository: Options

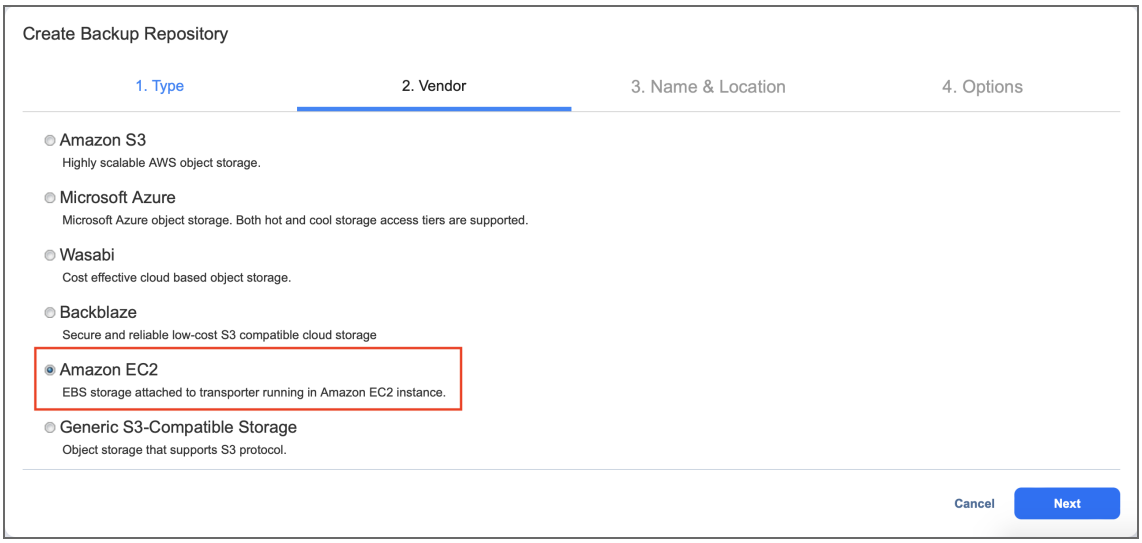
## Create Backup Repository: Type

On the **Type** page of the **Create Backup Repository** wizard, select **Cloud & S3-compatible Storage** and click **Next** to move to the next page of the wizard.




## Create Backup Repository: Vendor

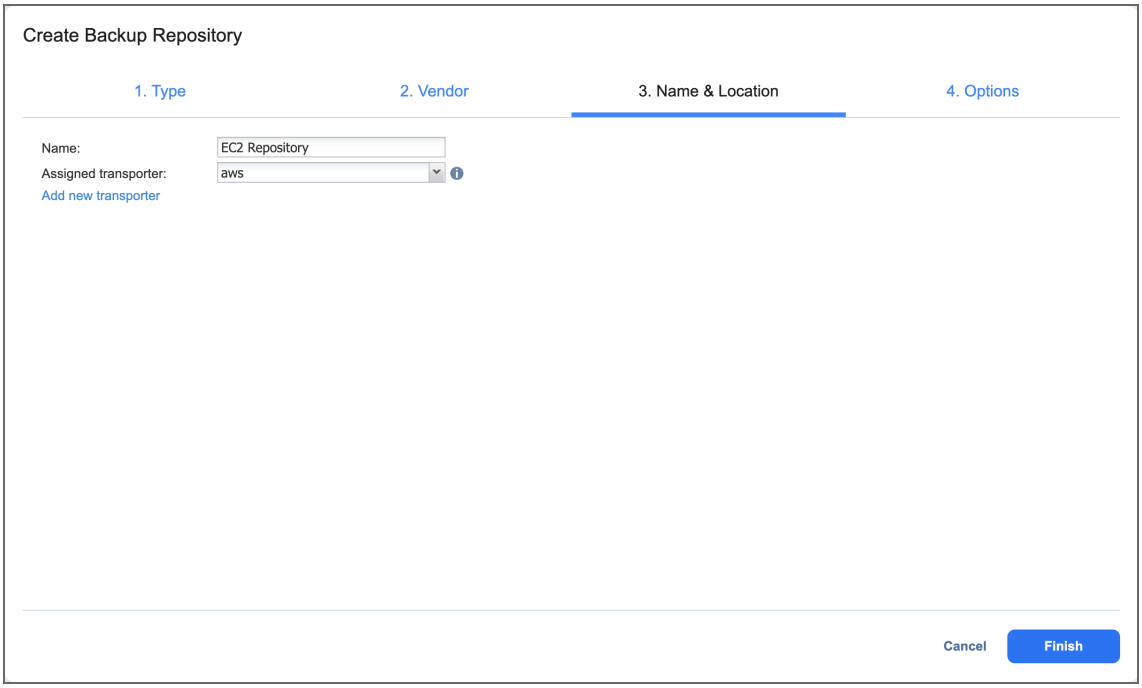
On the **Vendor** page of the wizard, select **Amazon EC2**. Click **Next** to proceed to the next step.




## Create Backup Repository: Name & Location

On the **Name & Location** page of the wizard, do the following:

1. Enter the name of the **Backup Repository** in the **Name** box.
2. Select the **Transporter** from the **Assigned transporter** drop-down list. To add a new **Transporter**, click **Add new transporter** and configure a new Transporter in the **Transporters** tab. Once the new **Transporter** is successfully added, it appears in the **Assigned transporter** drop-down list.
3. Click **Next** to go to the next page of the wizard.




## Create Backup Repository: Options

On the **Options** page, do the following:

Configure data storage options:

- **Volume type**: Choose one of the following EBS volumes that will be used for creating the **Backup Repository**:
  - Cold HDD (sc1)
  - Throughput Optimized HDD (st1)
  - General Purpose SDD (gp2)
  - General Purpose SDD (gp3)
  - Magnetic Standard
- **Storage**: Specify a size for the Backup Repository that will be allocated in Amazon EC2 using EBS Volumes. The volumes will be attached to the selected Amazon EC2 **Transporter**.
- **Storage chunk (GB)**: A **Backup Repository** in Amazon EC2 is created by using multiple EBS Volumes (chunks). The maximum size of the **Backup Repository** is limited to 50 EBS Volumes (chunks) or 16 TB (whichever occurs first). The size of a storage chunk defines the size of each individual EBS volume. Also, the storage will be resized (either manually or automatically) with the minimal step of the storage chunk specified here. To scale up to 16,000 GB, it is recommended that you have 400 GB storage chunk or bigger. Storage chunk cannot be changed later.
- **Automatically resize storage**: When this option is selected, the cloud storage is automatically increased and reduced as required.

Set up **Storage Savings & Encryption** options:

- **Data size reduction**: When this option is enabled, NAKIVO Backup & Replication enables the use of data size reduction for this repository to save disk space. Note that this may put additional load on the CPU. Disabling data size reduction is required if the target is a deduplication storage appliance. Click **Settings** to configure the settings. A popup window appears. Set the following:
  - **Compression**: Select a compression level to be used to reduce the data size in the **Backup Repository**. Note that higher compression levels consume considerably more CPU and may slow down VM backup speed. The following options are available:
    - **Disabled:** No compression.
    - **Fast:** Lowest compression level.
    - **Medium:** Medium compression level.
    - **Best:** Maximum compression level.

> **Note**
> This option cannot be configured after creating the **Backup Repository**.

- **Store backups in separate files**: Select this option to enable this **Backup Repository** to store data of every machine in separate backup files. Enabling this option is highly recommended to ensure better reliability and performance. Leave this option unselected if you wish to enable deduplication on a **Backup Repository**.
- **Deduplication**: Select this option to enable the backup deduplication method to reduce the backup size by excluding duplicate data blocks from the backup.

> **Note**
> This option is not available if the Store backups in separate files checkbox has been selected.

- **Encryption:** This option is available only if the **Backup Repository** is created locally on the machine on which the Assigned **Transporter** is installed, and the machine is running a Linux OS. Select **Enabled** from the drop-down list and specify an encryption password (the password is required for importing the **Backup Repository** to a new instance of the product). The product will encrypt the repository destination using `ecryptfs` for folders and `cryptsetup (crypt-md)` in LUKS mode for devices/partitions before creating the **Backup Repository**.

> **Notes**
> - To avoid `ecryptfs` errors, make sure that there are no other folders and files except the *NakivoBackup* folder in the repository location.
> - **Backup Repository** encryption can significantly affect backup speed.

Set up **Reliability & Maintenance** options:

- **Enable automatic repository self-healing:** Leave this option selected to automatically trigger repository self-healing in case the product detects symptoms of problems in the backup infrastructure (such as incorrect timestamps on metadata and data files). You can deselect this option and run self-healing manually.
- **Run repository self-healing on schedule:** You can select this checkbox to run repository self-healing based on a schedule. You can configure the schedule by clicking the **schedule** link when the option is selected. The default schedule is set to run every day at 11 AM.
  When **Stop backup and recovery to run self-healing** is selected, any jobs or recoveries that use this repository are stopped to run scheduled self-healing. Otherwise, scheduled self-healing is skipped if there are running jobs or recoveries on this repository.

- **Run full data verification on schedule:** When this option is selected, NAKIVO Backup & Replication runs full verification of all data available in the **Backup Repository** based on the specified schedule. The product reads each block of data to ensure that it is identical to the data block that was read on the source VM during the backup. This way, the product verifies each recovery point in the **Backup Repository**.
  When **Stop backup and recovery to run backup verification** is selected, any running jobs which use this **Backup Repository** are stopped to run scheduled data verification. Otherwise, scheduled data verification is skipped if there are running jobs on this Backup Repository.

  > **Note**
  > Backup verification is a time-consuming process and consumes CPU of the **Transporter** assigned to the Backup Repository. It is recommended that you schedule backup verification during non-working hours.

- **Reclaim unused space on schedule**: You can select this option to run the **Backup Repository** space reclaim process based on a schedule. Space reclaim compacts the data. Unused space is reclaimed. Keep in mind that this process can be time-consuming.

  > **Note**
  > This option is available only when **Store backups in separate files** is not enabled.

  When **Stop backup and recovery to run space reclaim** is selected, any running jobs that use this **Backup Repository** are stopped to run scheduled space reclaiming. Otherwise, scheduled space reclaiming is skipped if there are running jobs on this **Backup Repository**.

  > **Important**
  > Do not reboot/disconnect the "null" **Transporter** and storage device while space reclaim is in progress to avoid **Backup Repositor**y corruption.

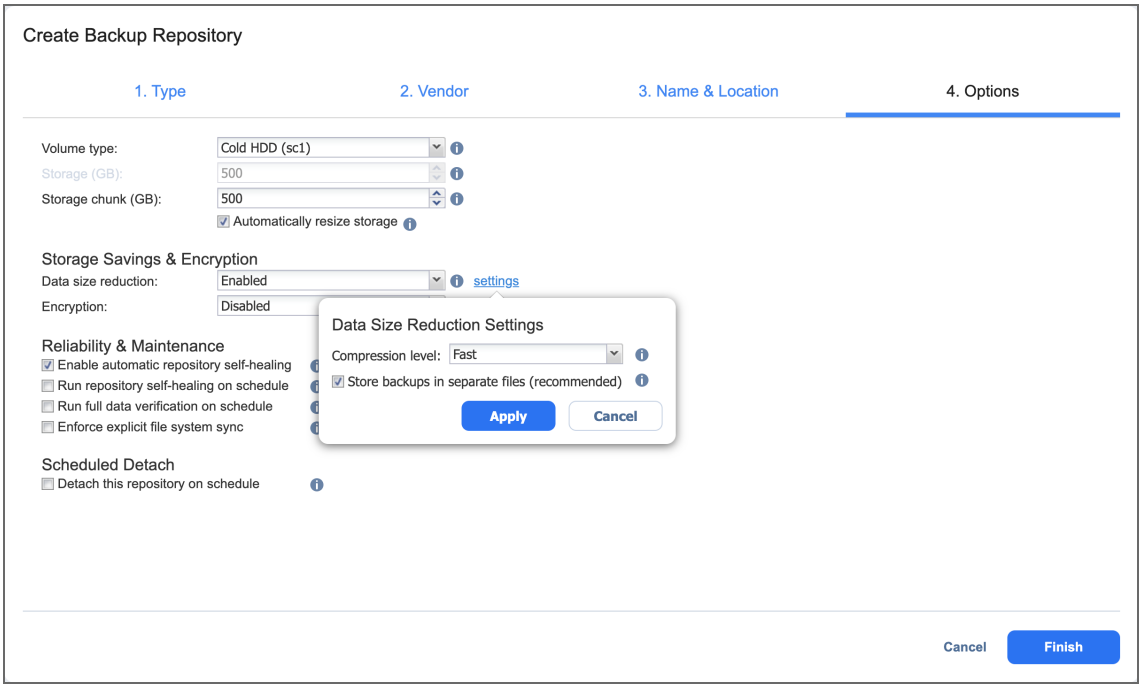
- **Enforce explicit file system sync**: When this option is selected, explicit sync with the file system is enforced during all backup operations to this repository. This setting is considered more reliable but may lead to lower performance on certain storage devices. By default, the option is disabled.

Schedule detaching of the Backup Repository:

- **Detach this repository on schedule:** Select this option if you want to detach and then attach the **Backup Repository** based on a schedule. Detaching a **Backup Repository** saves its data and metadata in a consistent state and then stops the product's interaction with the **Backup Repository** (so that it can be copied or moved). You can use this feature, for example, for the disk-to-disk-to-tape (D2D2T) data protection approach: backups are stored on a disk for fast operational recovery and copied to a tape (while the repository is detached) for archiving and long-term storage.

- **Delete and re-create the repository on attach:** When this option is selected, all data in the **Backup Repository** is erased prior to attaching it to the product. As a result, jobs that write to this **Backup Repository** create full VM backups. You can use this option, for example, to create full daily, weekly, or monthly VM backups and write them to tape or removable media.

Click **Finish** to finish creating the **Backup Repository**.



# Backup Repository in Amazon S3

Select the **Amazon S3** option if you want to create a **Backup Repository** in Amazon S3. Before creating a repository, grant the required S3 access permissions to NAKIVO Backup & Replication. For details, refer to Required AWS IAM Permissions for Amazon S3 and Backblaze and Permissions for the Amazon S3 Bucket. In addition, make sure to enable the following options for the relevant Amazon S3 bucket:

- Object Lock
- Versioning

Since retention settings are set by NAKIVO Backup & Replication during job creation, disable the Object Lock retention mode and retention period on the S3 bucket as well.

**Important**

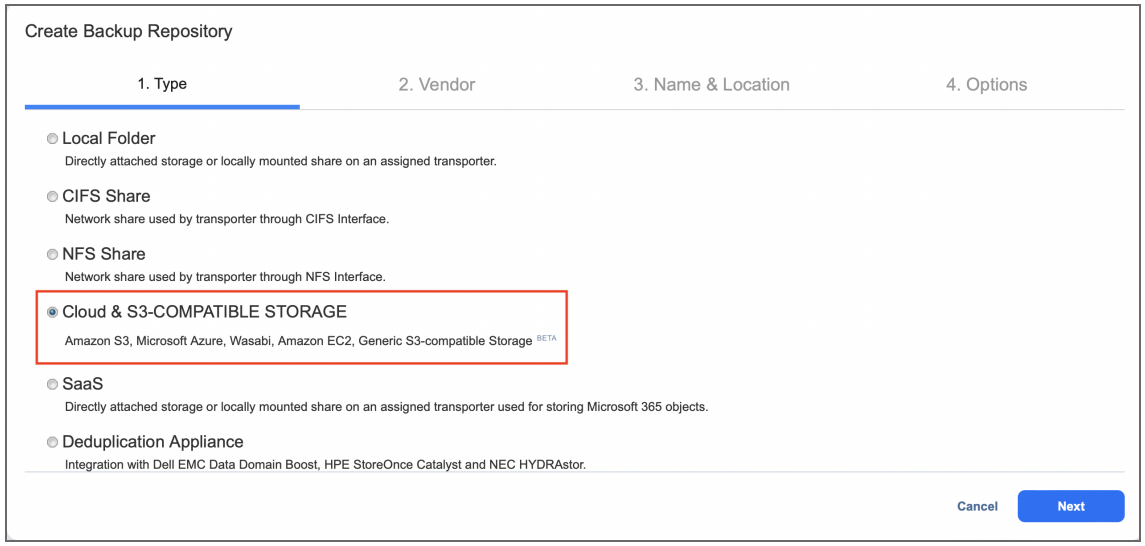
- You will be charged for Amazon S3 storage/traffic according to AWS tariffs.
- Forever incremental backups are not supported by this location.
- Only Amazon S3 Standard storage class is supported.

To create a **Backup Repository** in an Amazon S3 bucket, proceed as described in the following sections:

- Create Backup Repository: Type
- Create Backup Repository: Vendor
- Create Backup Repository: Name & Location
- Create Backup Repository: Options

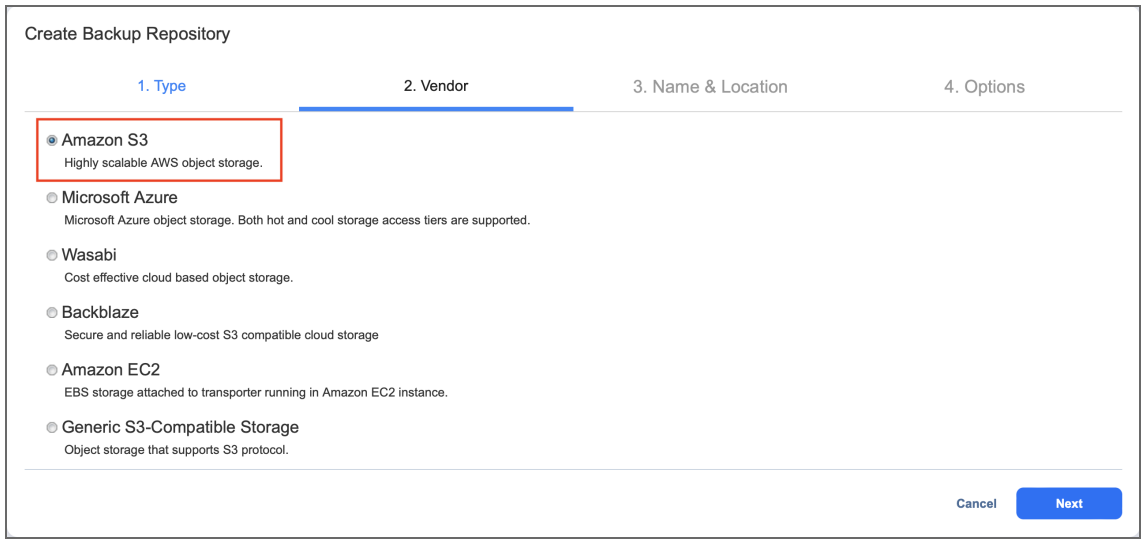
## Create Backup Repository: Type

On the **Type** page of the **Create Backup Repository** wizard, select **Cloud & S3-compatible Storage** and click **Next** to go to the next page of the wizard.



## Create Backup Repository: Vendor

On the **Vendor** page of the wizard, select **Amazon S3**. Click **Next** to proceed to the next step.

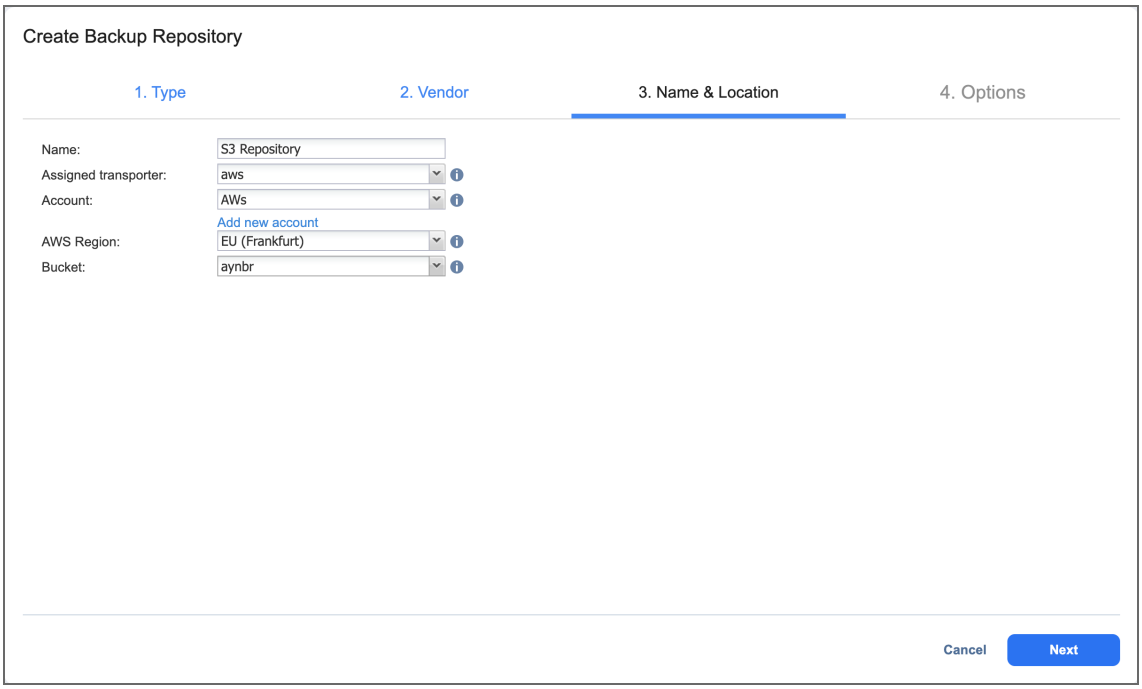



## Create Backup Repository: Name & Location

On the **Name & Location** page of the wizard, do the following:

1. Enter the name of the **Backup Repository** in the **Name** box.
2. Select the **Transporter** from the **Assigned transporter** drop-down list.
3. Select an AWS account from the **Account** drop-down list.
4. Select the **AWS region** connected to the bucket where you want to store your backups.
5. Select the bucket where you want to store your backups from the **Bucket** drop-down list.

6. Click **Next** to go to the next page of the wizard.




## Create Backup Repository: Options

On the **Options** page, do the following:

1. In the **Storage Savings** section, select a compression level for reducing the data size in the **Backup Repository**. Note that higher compression levels consume considerably more CPU and may slow down the backup speed. The following options are available:
   - **Disabled**: No compression.
   - **Fast**: Lowest compression level.
   - **Medium**: Medium compression level.
   - **Best**: Maximum compression level.

   > **Note**
   > This option cannot be configured after you create the **Backup Repository**.

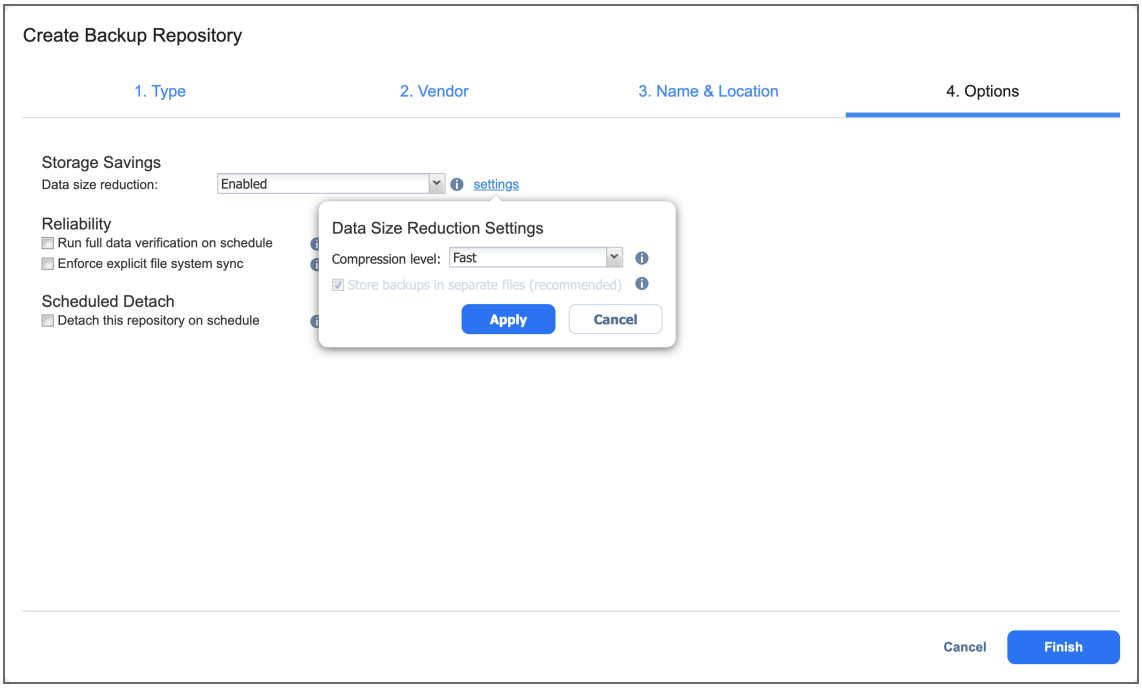
2. Set up **Reliability & Maintenance** options:
   - **Run full data verification on schedule**: When selected, the product runs full verification of all data available in the **Backup Repository** according to the specified schedule. The product reads each block of data and ensures that it is identical to the data block that was read on the source machine during the backup. This way, the product verifies each recovery point in the **Backup Repository**.

When **Stop backup and recovery to run full data verification** is selected, any running jobs that use this **Backup Repository** are stopped to run scheduled data verification. When this option is not selected, scheduled data verification is skipped if there are running jobs on this **Backup Repository**.

> **Note**
> Backup verification is a time-consuming process and utilizes the CPU resources of the **Transporter** assigned to the **Backup Repository**. It is recommended that you schedule backup verification during non-working hours.

- **Enforce explicit file system sync**: When selected, explicit sync with the file system is enforced during all backup operations to this repository. This setting is considered more reliable but may lead to lower performance on certain storage devices. By default, the option is disabled.

3. Schedule detaching of the **Backup Repository**:
   - **Detach this repository on schedule**: Select this option if you want to detach and then attach the **Backup Repository** based on a schedule. Detaching a **Backup Repository** saves its data and metadata in a consistent state and then stops the interaction of the product with the **Backup Repository** (so that it can be copied or moved). You can use this feature, for example, for the disk-to-disk-to-tape (D2D2T) data protection approach, in which backups are stored on a disk for fast operational recovery and copied to a tape (while the repository is detached) for archiving and long-term storage.
   - **Delete and re-create the repository on attach**: When this option is selected, all the data in the **Backup Repository** is erased prior to attaching it to the product. As a result, jobs that write to this **Backup Repository** create full backups. You can use this option, for example, to create full daily, weekly, or monthly VM backups and write them to tape or removable media.
4. Click **Finish** to complete **Backup Repository** creation.




# Backup Repository in Generic S3-Compatible Object Storage

> **Note**
>
> Only specific S3-comaptible vendors are supported. Please see Feature Requirements for more information.

Before creating a repository, enable the following options for the generic S3-compatible storageused:

- Object Lock
- Versioning

To create a **Backup Repository** in a generic S3-compatible object bucket, proceed as described in the following sections:

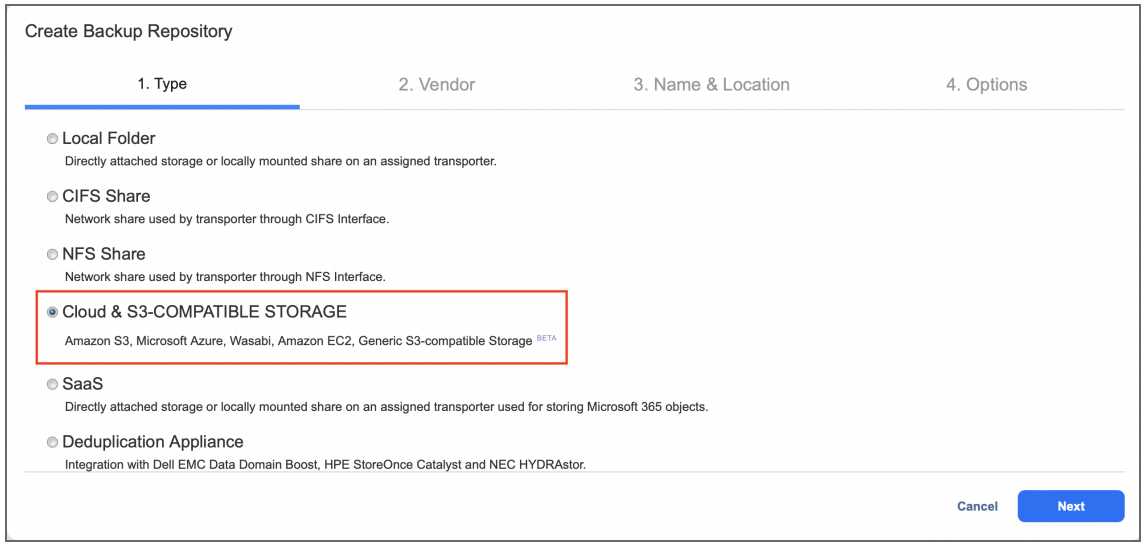
- Create Backup Repository: Type
- Create Backup Repository: Vendor
- Create Backup Repository: Name & Location
- Create Backup Repository: Options

**Important**

- Forever incremental backups are not supported by this location.
- The S3-compatible bucket is designated exclusively for use as a NAKIVO repository. Please refrain from placing any third-party data into it.
- Buckets with the following naming patterns will not be discovered:
  - Names shorter than 3 characters or longer than 63 characters
  - Names that do not start and end with a letter or number
  - Names containing unsupported special characters (allowed characters are: lowercase letters, numbers, dot . and hyphen -)
  - Names starting with the prefix xn–
  - Names ending with the suffix -s3alias
  - Names formatted as IP addresses

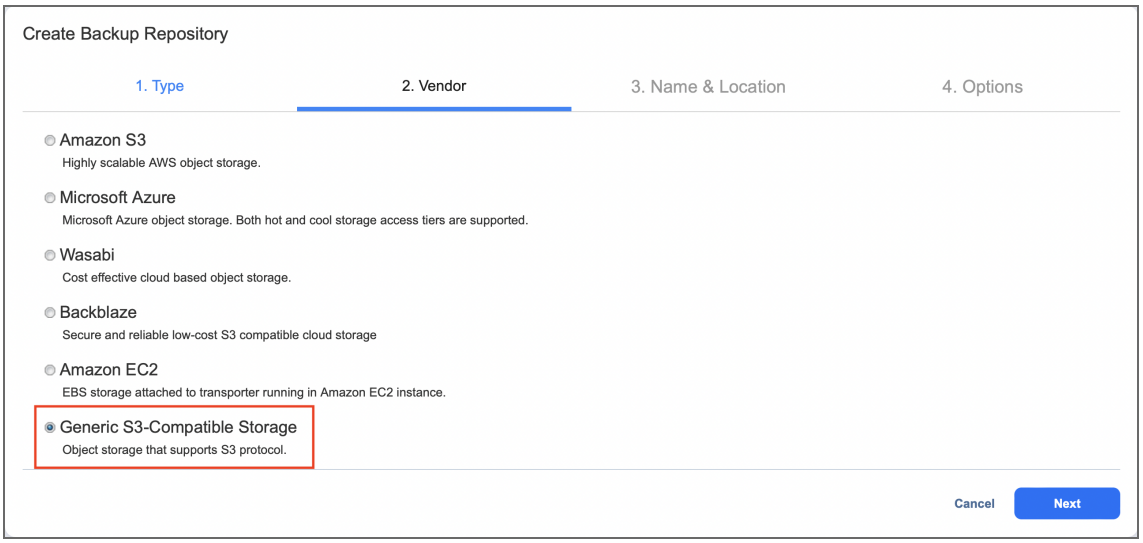
## Create Backup Repository: Type

On the **Type** page of the **Create Backup Repository** wizard, select **Cloud & S3-compatible Storage** and click **Next** to go to the next page of the wizard.




## Create Backup Repository: Vendor

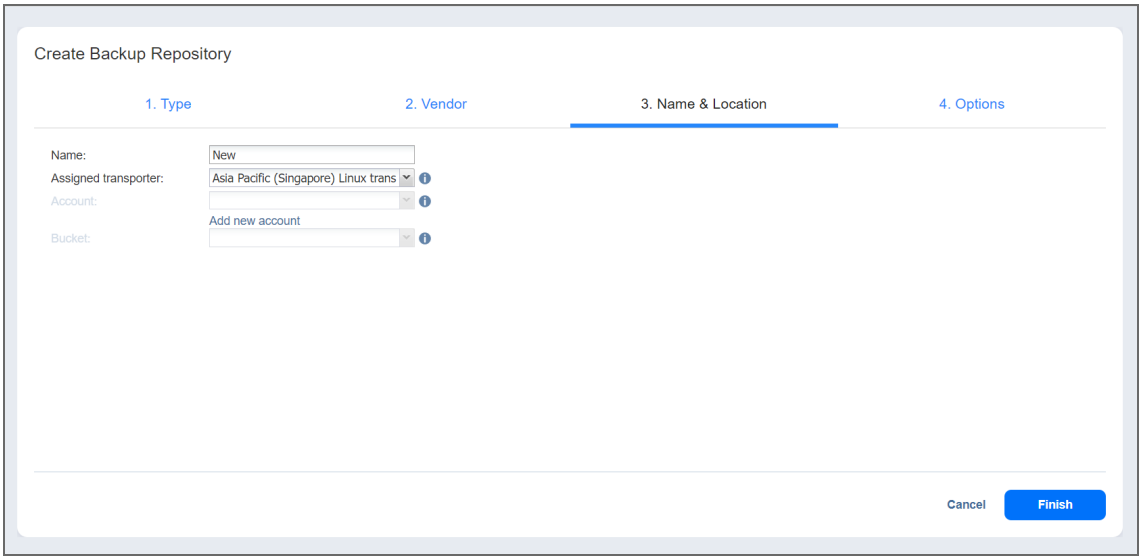
On the **Vendor** page of the wizard, select **Generic S3-Compatible Storage**. Click **Next** to proceed to the next step.




## Create Backup Repository: Name & Location

On the **Name & Location** page of the wizard, do the following:

1. Enter the name of the **Backup Repository** in the **Name** box.
2. Select the **Transporter** from the **Assigned transporter** drop-down list.
3. Select the existing generic S3-compatible object storage **Account** where the backup repository will be located.
4. Optionally, click the **Add new account** link if you have not yet added a generic S3-compatible object storage account to the **Inventory**.
5. Select the bucket where you want to store your backups from the **Bucket** drop-down list.
6. Click **Next** to go to the next page of the wizard.




## Create Backup Repository: Options

On the **Options** page, do the following:

1. In the **Storage Savings** section, select a compression level for reducing the data size in the **Backup Repository**. Note that higher compression levels consume considerably more CPU and may slow down the backup speed. The following options are available:
    - **Disabled**: No compression
    - **Fast**: Lowest compression level
    - **Medium**: Medium compression level
    - **Best**: Maximum compression level

    > **Note**
    > This option cannot be configured after you create the **Backup Repository**.

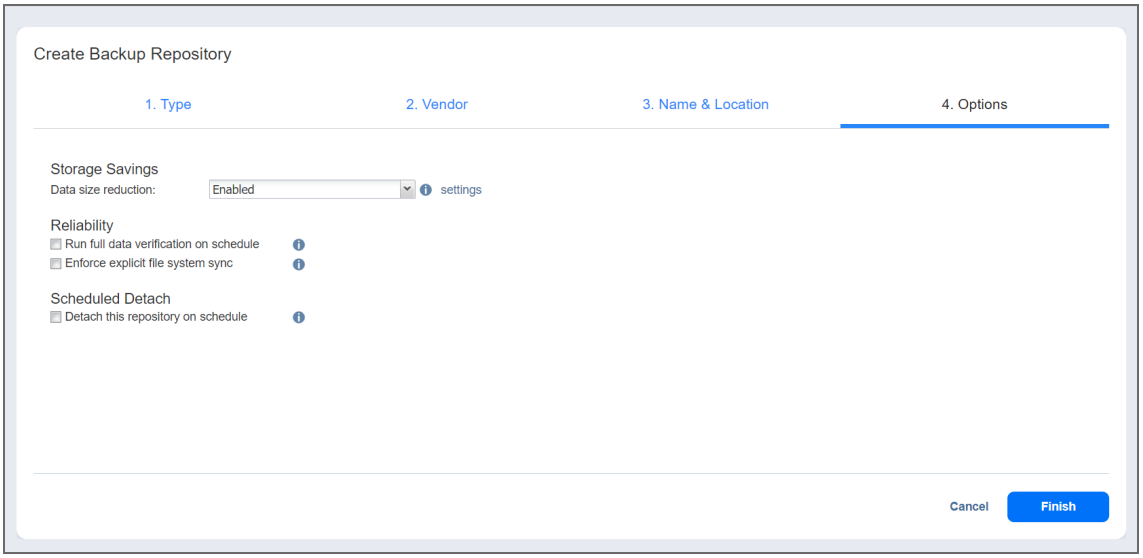
2. Set up **Reliability & Maintenance** options:
    - **Run full data verification on schedule**: When this option is selected, the product runs full verification of all data available in the **Backup Repository** according to the specified schedule. The product reads each block of data and ensures that it is identical to the data block that was read on the source machine during the backup. This way, the product verifies each recovery point in the **Backup Repository**.
        - When **Stop backup and recovery to run full data verification** is selected, any running jobs that use this Backup Repository are stopped to run scheduled data verification. When this option is not selected, scheduled data verification is skipped if there are running jobs on this **Backup Repository**.

        > **Note**
        > Backup verification is a time-consuming process and utilizes the CPU resources of the **Transporter** assigned to the **Backup Repository**. It is recommended that you schedule backup verification during non-working hours.

    - **Enforce explicit file system sync**: When this option is selected, explicit sync with the file system is enforced during all backup operations to this repository. This setting is considered more reliable but may lead to lower performance on certain storage devices. By default, the option is disabled.
3. Schedule detaching of the Backup Repository:

- **Detach this repository on schedule**: Select this option if you want to detach and then attach the **Backup Repository** on a schedule. Detaching a **Backup Repository** saves its data and metadata in a consistent state and then stops the product's interaction with the **Backup Repository** (so that it can be copied or moved). You can use this feature, for example, for the disk-to-disk-to-tape (D2D2T) data protection approach, in which backups are stored on a disk for fast operational recovery, and copied to a tape (while the repository is detached) for archiving and long-term storage.
- **Delete and re-create the repository on attach**: When this option is selected, all the data in the **Backup Repository** is erased prior to attaching the repository to the product. As a result, jobs that write to this **Backup Repository** create full backups. You can use this option, for example, to create full daily, weekly, or monthly backups and write them to tape or removable media.

4. Click **Finish** to complete **Backup Repository** creation.




## Backup Repository in Microsoft Azure Blob Storage

Before creating a Microsoft Azure Blob storage repository, you need to configure your Azure storage account to work with NAKIVO Backup & Replication. For details, refer to "Configuring a Microsoft Azure Storage Account" on page 526.

To create a **Backup Repository** in Azure Blob storage, proceed as described in the following sections:

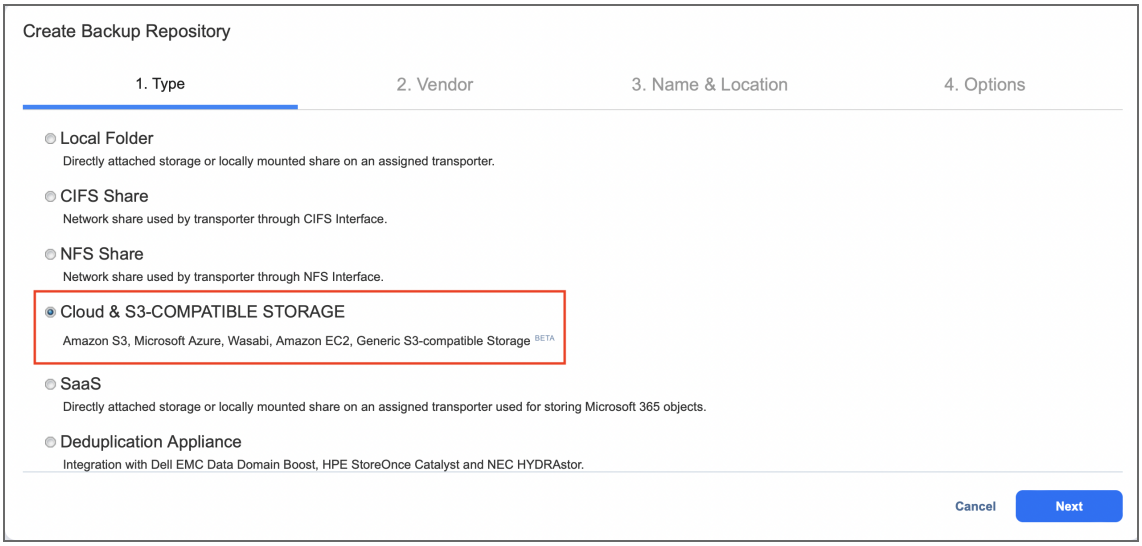
- Create Backup Repository: Type
- Create Backup Repository: Vendor
- Create Backup Repository: Name & Location
- Create Backup Repository: Options

**Important**

- Forever incremental backups are not supported by this location.

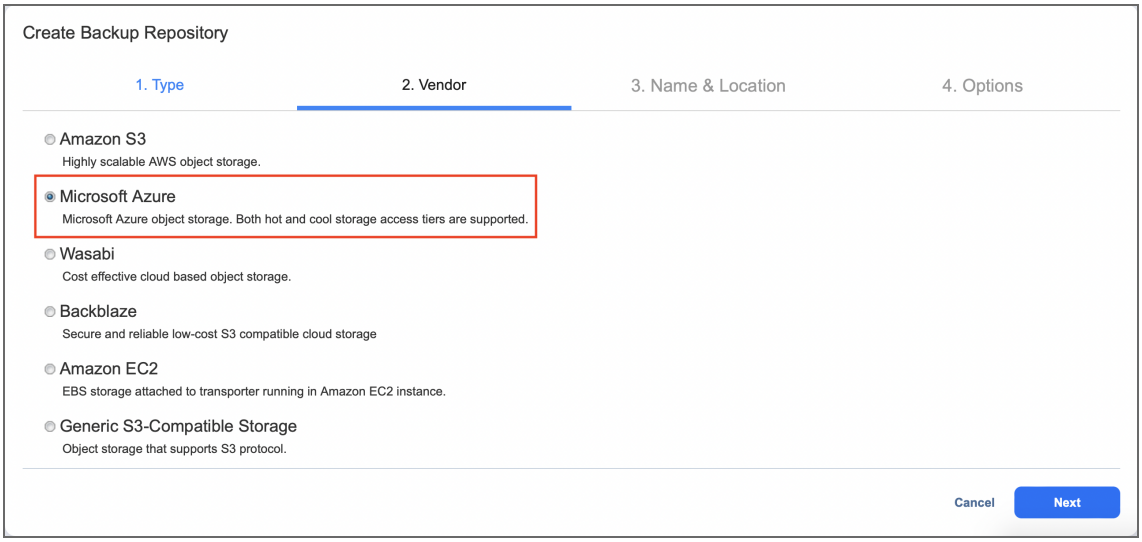
## Create Backup Repository: Type

On the **Type** page of the **Create Backup Repository** wizard, select **Cloud & S3-compatible Storage** and click **Next** to go to the next page of the wizard.




## Create Backup Repository: Vendor

On the **Vendor** page of the wizard, select **Microsoft Azure**. Click **Next** to proceed to the next step.



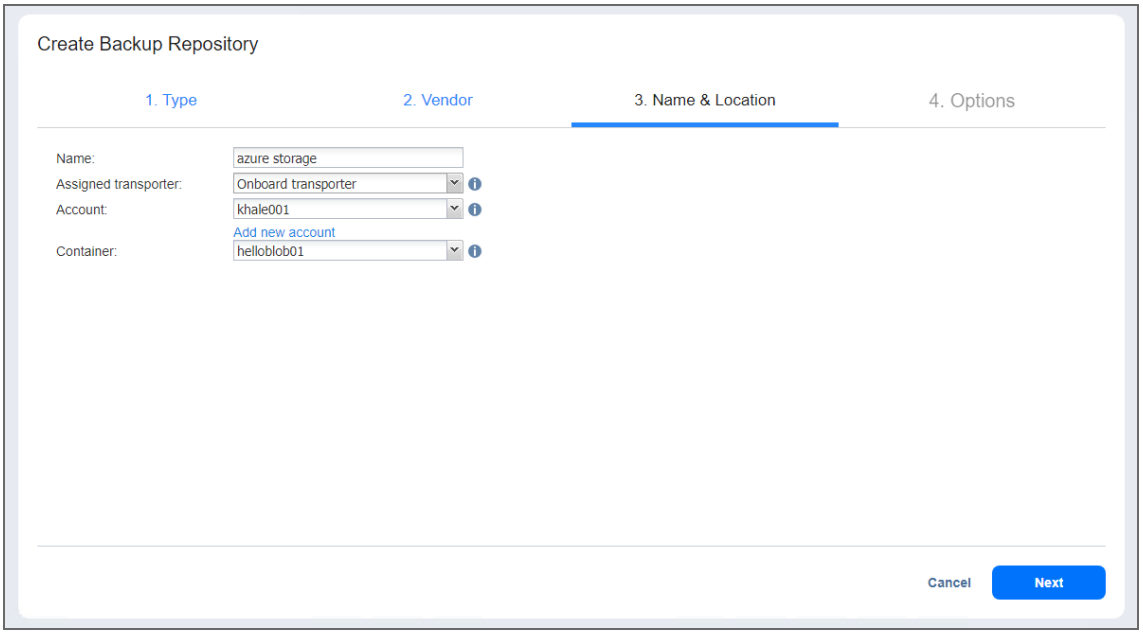
**Notes**

- The following Microsoft Azure storage account types are supported:
  - **General-purpose V2**: Blob storage (block blob, page blob)
  - **Premium Block blobs**: Blob storage (block blob only)
- New backup repository is created in the container utilizing block blob type only.
- Existing repositories created in a container utilizing page blob type may still be used. However, saving data in a container utilizing page blob type is not recommended.

## Create Backup Repository: Name & Location

On the **Name & Location** page of the wizard, do the following:

1. Enter the name of the **Backup Repository** in the **Name** box.
2. Select the **Transporter** from the **Assigned transporter** drop-down list.
3. Select a Microsoft Azure storage account from the **Account** drop-down list.
4. In the **Container** drop-down list, select the container within the chosen storage account where you want to store backups.
5. Click **Next** to go to the next page of the wizard.




## Create Backup Repository: Options

On the **Options** page, do the following:

1. In the **Storage Savings** section, select a compression level for reducing the data size in the **Backup Repository**. Note that higher compression levels consume considerably more CPU and may slow down the backup process. The following options are available:

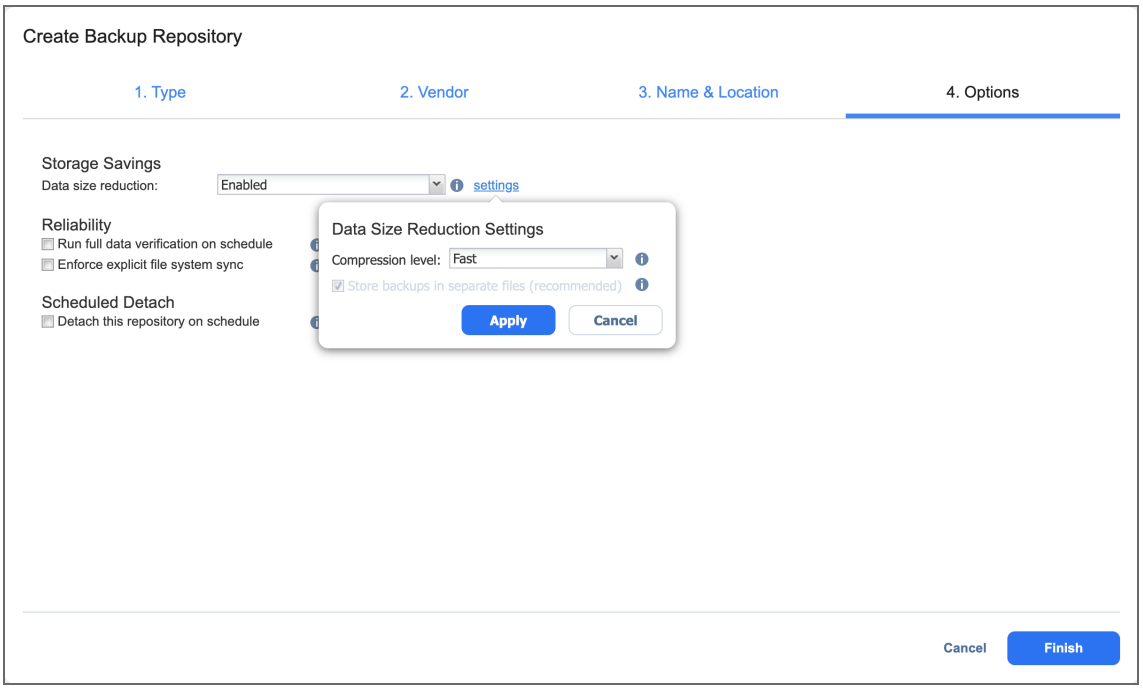
- **Disabled**: No compression.
- **Fast**: Lowest compression level.
- **Medium**: Medium compression level.
- **Best**: Maximum compression level.

> **Note**
> This option cannot be configured after you create the Backup Repository.

2. Set up **Reliability & Maintenance** options:
3. **Run full data verification on schedule**: When this option is selected, the product runs full verification of all data available in the Backup Repository based on the specified schedule. The product reads each block of data and ensures that it is identical to the data block that was read on the source machine during the backup. This way, the product verifies each recovery point in the **Backup Repository**. When **Stop backup and recovery to run full data verification** is selected, any running jobs that use this Backup Repository are stopped to run scheduled data verification. When this option is not selected, scheduled data verification is skipped if there are running jobs on this **Backup Repository**.

> **Note**
> Backup verification is a time-consuming process and utilizes the CPU resources of the **Transporter** assigned to the Backup Repository. It is recommended that you schedule backup verification during non-working hours.

   - **Enforce explicit file system sync**: When this option is selected, explicit sync with the file system is enforced during all backup operations to this repository. This setting is considered more reliable but may lead to lower performance on certain storage devices. This option is disabled by default.
4. Schedule detaching of the Backup Repository:
   - **Detach this repository on schedule**: Select this option if you want to detach and then attach the **Backup Repository** based on a schedule. Detaching a **Backup Repository** saves its data and metadata in a consistent state and then stops the interaction of the product with the **Backup Repository** (so that it can be copied or moved).
   - **Delete and re-create the repository on attach**: When this option is selected, all the data in the **Backup Repository** is erased prior to attaching it to the product. As a result, jobs that write to this **Backup Repository** create full backups. You can use this option, for example, to create full daily, weekly, or monthly VM backups and write them to tape or removable media.
5. Click **Finish** to complete Backup **Repository** creation.




# Backup Repository in Backblaze B2 Cloud Storage

To create a **Backup Repository** in Backblaze B2 storage, proceed as described in the following sections:

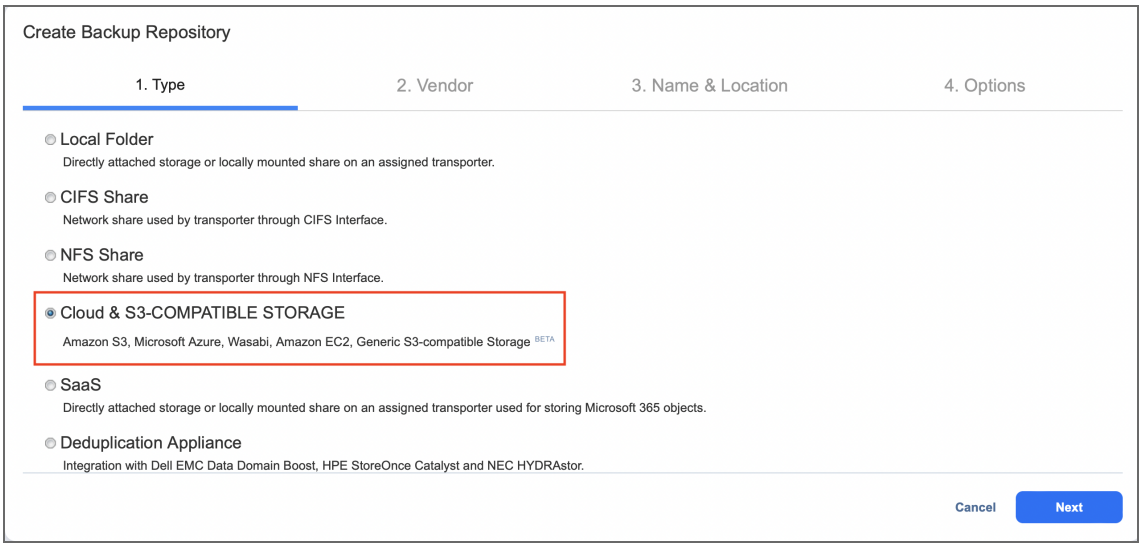
- Create Backup Repository: Type
- Create Backup Repository: Vendor
- Create Backup Repository: Name & Location
- Create Backup Repository: Options

**Important**

- Forever incremental backups are not supported by this location.

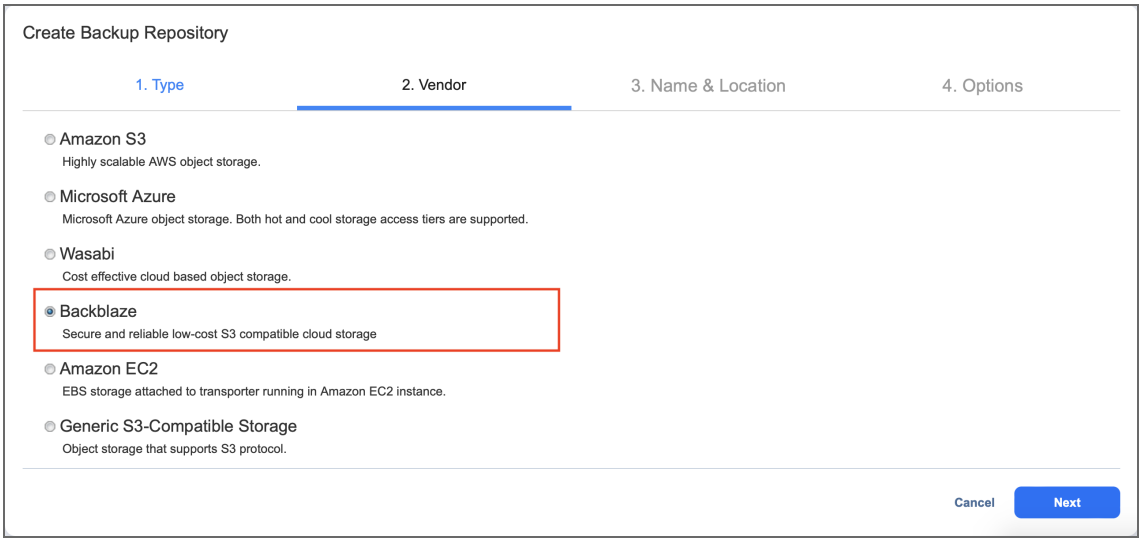
## Create Backup Repository: Type

On the **Type** page of the **Create Backup Repository** wizard, select **Cloud** and click **Next** to go to the next page of the wizard.




## Create Backup Repository: Vendor

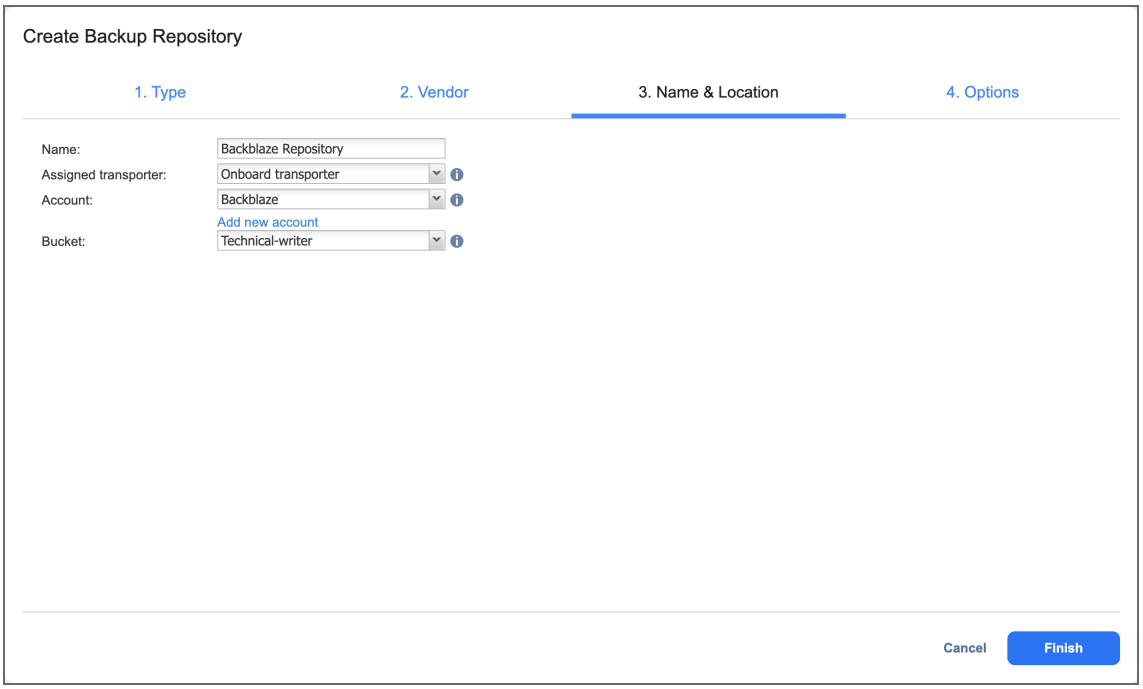
On the **Vendor** page of the wizard, select **Backblaze**. Click **Next** to proceed to the next step.




## Create Backup Repository: Name & Location

On the **Name & Location** page of the wizard, do the following:

1. Enter the name of the y**Backup Repositor** in the **Name** box.
2. Select the **Transporter** from the **Assigned transporter** drop-down list.
3. Select a Backblaze account from the **Account** drop-down list.
4. In the **Bucket** drop-down list, select the bucket within the chosen storage account where you want to store backups.
5. Click **Next** to go to the next page of the wizard.




## Create Backup Repository: Options

On the **Options** page, do the following:

1. In the **Storage Savings** section, select a compression level for reducing the data size in the **Backup Repository**. Note that higher compression levels consume considerably more CPU and may slow down the backup process. The following options are available:
   - **Disabled**: No compression.
   - **Fast**: Lowest compression level.
   - **Medium**: Medium compression level.
   - **Best**: Maximum compression level.

   > **Note**
   > This option cannot be configured after you create the **Backup Repository**.

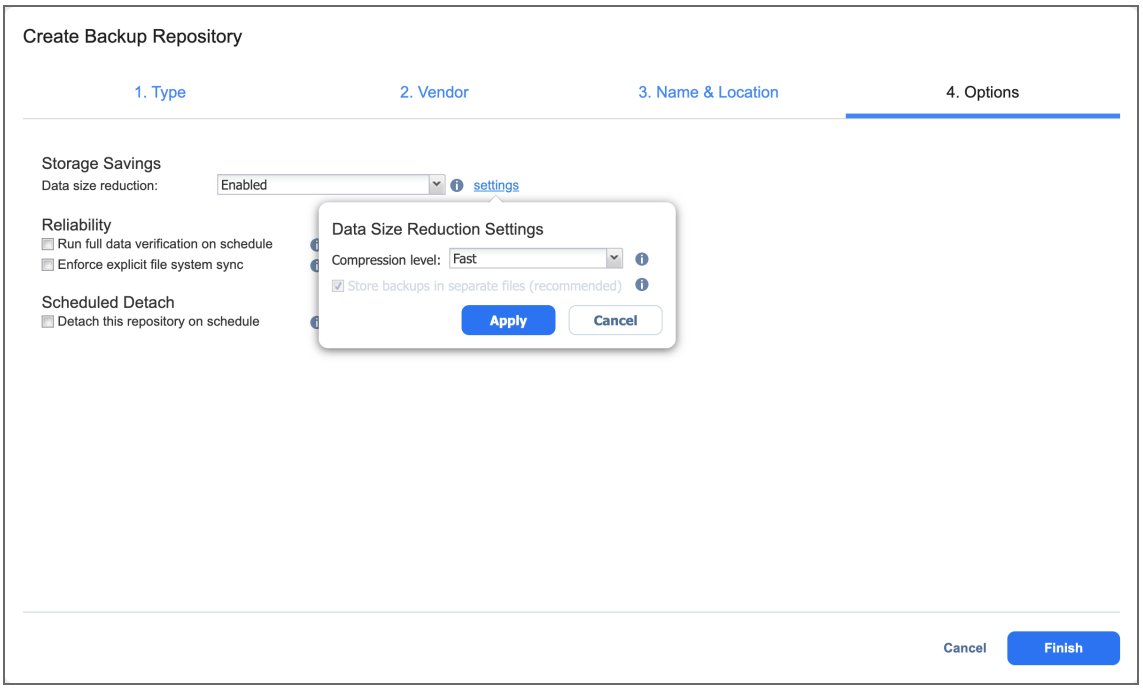
2. Set up **Reliability & Maintenance** options:
   - **Run full data verification on schedule**: When this option is selected, the product runs full verification of all data available in the **Backup Repository** based on the specified schedule. The product reads each block of data and ensures that it is identical to the data block that was read on the source machine during the backup. This way, the product verifies each recovery point in the Backup Repository.

When **Stop backup and recovery to run full data verification** is selected, any running jobs that use this **Backup Repository** are stopped to run scheduled data verification. When this option is not selected, scheduled data verification is skipped if there are running jobs on this **Backup Repository**.

> **Note**
> Backup verification is a time-consuming process and utilizes the CPU resources of the **Transporter** assigned to the **Backup Repository**. It is recommended that you schedule backup verification during non-working hours.

- **Enforce explicit file system sync**: When this option is selected, explicit sync with the file system is enforced during all backup operations to this repository. This setting is considered more reliable but may lead to lower performance on certain storage devices. This option is disabled by default.

3. Schedule detaching of the Backup Repository:
   - **Detach this repository on schedule**: Select this option if you want to detach and then attach the **Backup Repository** based on a schedule. Detaching a **Backup Repository** saves the **Backup Repository** data and metadata in a consistent state and then stops the interaction of the product with the **Backup Repository** (so that the Backup Repository can be copied or moved).
   - **Delete and re-create the repository on attach**: When this option is selected, all the data in the **Backup Repository** is erased prior to attaching it to the product. As a result, jobs that write to this **Backup Repository** create full backups. You can use this option, for example, to create full daily, weekly, or monthly VM backups and write them to tape or removable media.
4. Click **Finish** to complete **Backup Repository** creation.


Create Backup Repository
1. Type
2. Vendor
3. Name & Location
4. Options
Storage Savings
Data size reduction:
Enabled
settings
Reliability
Run full data verification on schedule
Enforce explicit file system sync
Scheduled Detach
Detach this repository on schedule
Data Size Reduction Settings
Compression level:
Fast
Store backups in separate files (recommended)
Apply
Cancel
Cancel
Finish

# Backup Repository in Wasabi Hot Cloud Storage

Select the **Wasabi** option if you want to create a **Backup Repository** in Wasabi. Before creating a repository, make sure to grant full access permissions to NAKIVO Backup & Replication in Wasabi.

**Important**

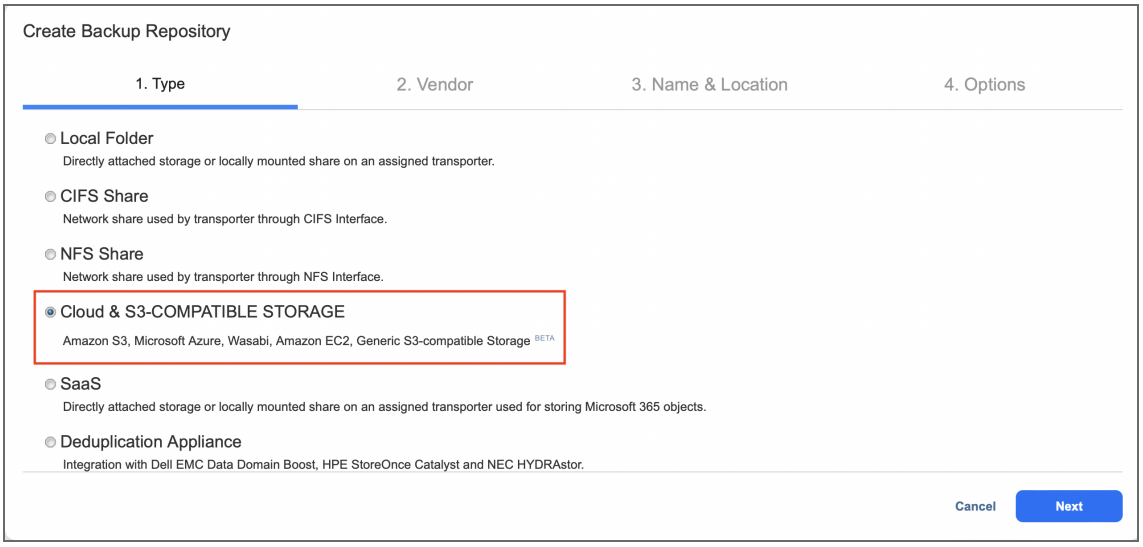
- You may be charged for Wasabi storage/traffic. Refer to Cloud Storage Pricing for details.
- Forever incremental backups are not supported by this location.
- The Wasabi bucket is designated exclusively for use as a NAKIVO repository. Please refrain from placing any third-party data into it.
- Make sure you are using full Wasabi account as the limited account does not have sufficient permissions.

To create a **Backup Repository** in a Wasabi bucket, proceed as described in the following sections:

- Create Backup Repository: Type
- Create Backup Repository: Vendor
- Create Backup Repository: Name & Location
- Create Backup Repository: Options

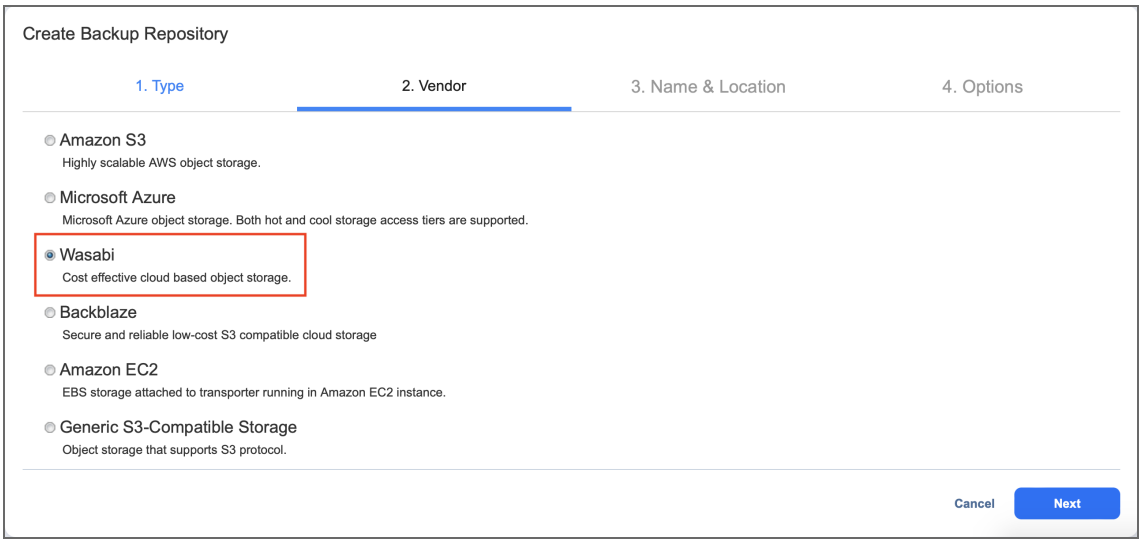
## Create Backup Repository: Type

On the **Type** page of the **Create Backup Repository** wizard, select **Cloud & S3-compatible Storage** and click **Next** to go to the next page of the wizard.




## Create Backup Repository: Vendor

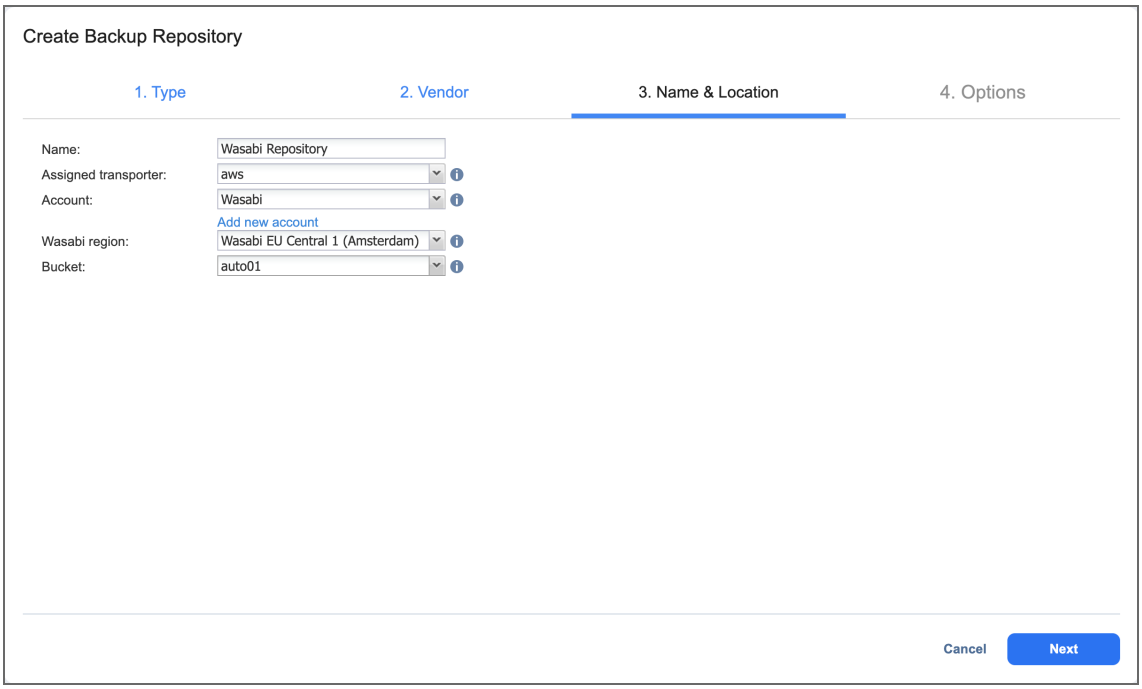
On the **Vendor** page of the wizard, select **Wasabi**. Click **Next** to proceed to the next step.




# Create Backup Repository: Name & Location

On the **Name & Location** page of the wizard, do the following:

1. Enter the name of the **Backup Repository** in the **Name** box.
2. Select the **Transporter** from the **Assigned transporter** drop-down list.
3. Select a Wasabi account from the **Account** drop-down list.
4. Select the **Wasabi region** connected to the bucket where you want to store your backups.
5. Select the bucket where you want to store your backups from the **Bucket** drop-down list.
6. Click **Next** to go to the next page of the wizard.



## Create Backup Repository: Options

On the **Options** page, do the following:

1. In the **Storage Savings** section, select a compression level for reducing the data size in the **Backup Repository**. Note that higher compression levels consume considerably more CPU and may slow down the backup speed. The following options are available:
   - **Disabled**: No compression.
   - **Fast**: Lowest compression level.
   - **Medium**: Medium compression level.
   - **Best**: Maximum compression level.

   > **Note**
   > This option cannot be configured after you create the **Backup Repository**.

2. Set up **Reliability & Maintenance** options:
   - **Run full data verification on schedule**: When selected, the product runs full verification of all data available in the **Backup Repositor**y according to the specified schedule. The product reads each block of data and ensures that it is identical to the data block that was read on the source machine during the backup. This way, the product verifies each recovery point in the **Backup Repository**.

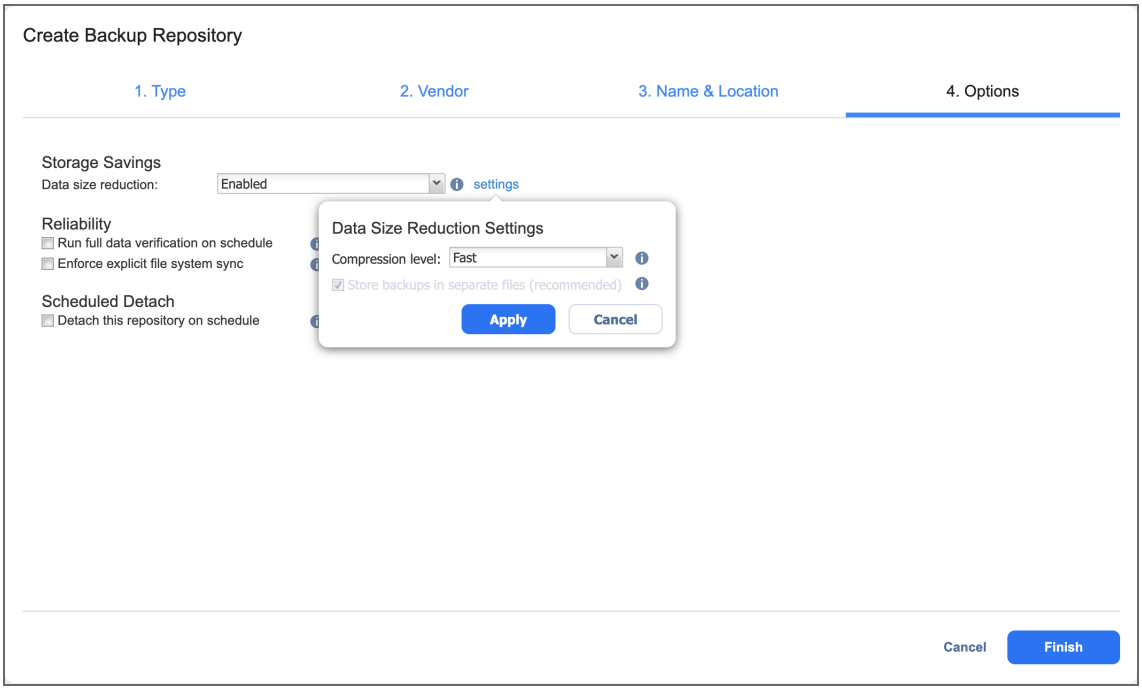
     When **Stop backup and recovery to run full data verification** is selected, any running jobs that use this **Backup Repository** are stopped to run scheduled data verification. When this option is not selected, scheduled data verification is skipped if there are running jobs on this **Backup Repository**.

     > **Note**
     > Backup verification is a time-consuming process and utilizes the CPU resources of the **Transporter** assigned to the **Backup Repository**. It is recommended that you schedule backup verification during non-working hours.

   - **Enforce explicit file system sync**: When selected, explicit sync with the file system is enforced during all backup operations to this repository. This setting is considered more reliable but may lead to lower performance on certain storage devices. By default, the option is disabled.
3. Schedule detaching of the **Backup Repository:**
   - **Detach this repository on schedule**: Select this option if you want to detach and then attach the **Backup Repository** based on a schedule. Detaching a **Backup Repository** saves its data and metadata in a consistent state and then stops the interaction of the product with the **Backup Repository** (so that it can be copied or moved).

- **Delete and re-create the repository on attach**: When this option is selected, all the data in the **Backup Repository** is erased prior to attaching it to the product. As a result, jobs that write to this **Backup Repository** create full backups. You can use this option, for example, to create full daily, weekly, or monthly VM backups and write them to tape or removable media.

4. Click **Finish** to complete **Backup Repository** creation.



# Backup Repository on Deduplication Appliance

NAKIVO Backup & Replication allows you to use advanced deduplication appliances for data protection.

**Notes**

- Before creating a **Backup Repository** on a Dell EMC DD, you need to install BoostFS Plugin and create a storage unit on the data domain backup appliance. Refer to Integrating with EMC DD Boost for details.
- Before creating a **Backup Repository** on an NEC HYDRAstor, you need to configure the NEC HYDRAstor and the machine on which NAKIVO Transporter is installed. Refer to Integrating with NEC HYDRAstor for details.
- To create a **Backup Repository** on other deduplication appliances, refer to "Backup Repository on NFS Share" on page 625.

To create a repository on a deduplication appliance, proceed as described in the following sections:

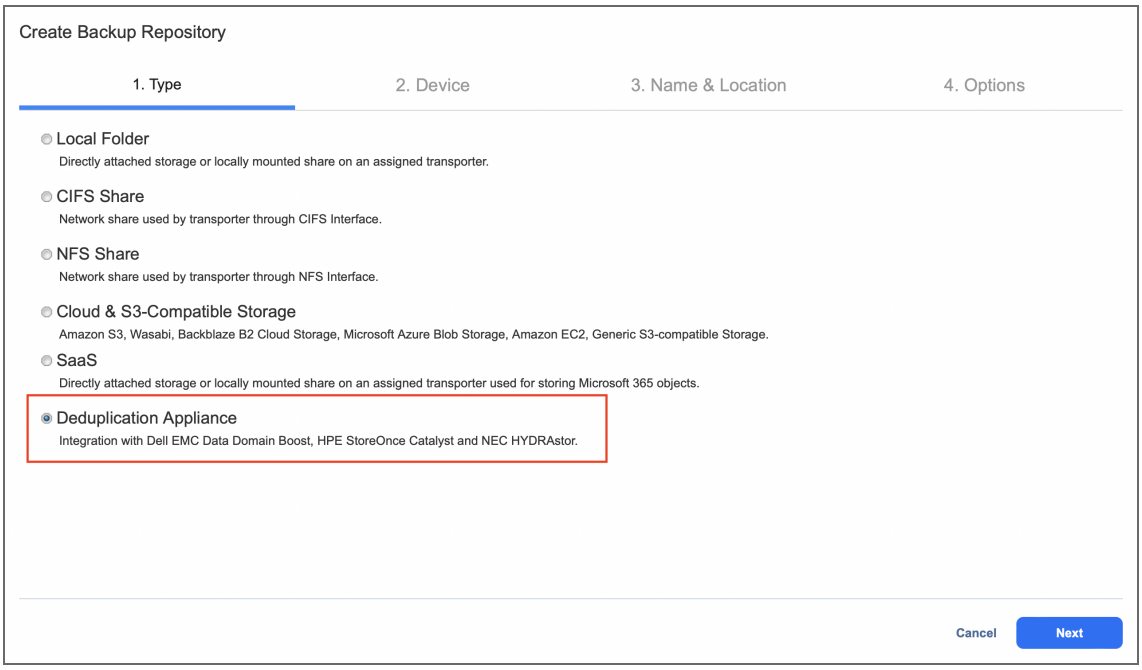
- Create Backup Repository: Type
- Create Backup Repository: Device
- Create Backup Repository: Name and Location
- Create Backup Repository: Options

## Create Backup Repository: Type

On the **Type** page of the **Create Backup Repository** wizard, select **Deduplication Appliance** and click **Next** to go to the next page of the wizard.

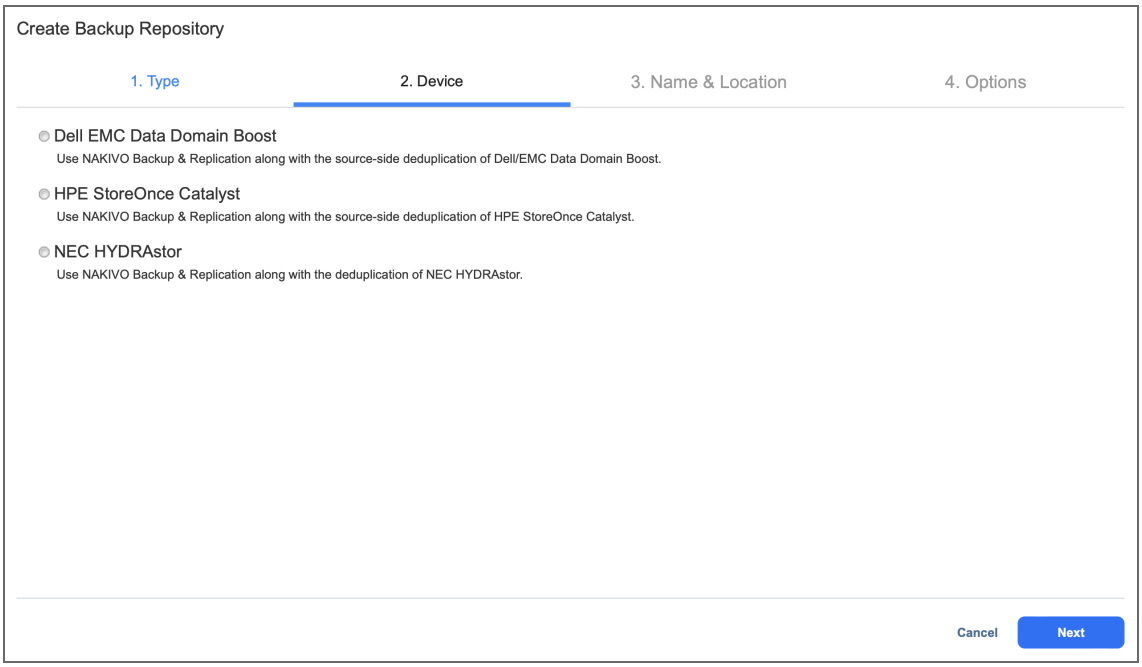
**Note**

Refer to "Storage Integration Requirements" on page 135 to see the list of supported advanced deduplication appliances.




# Create Backup Repository: Device

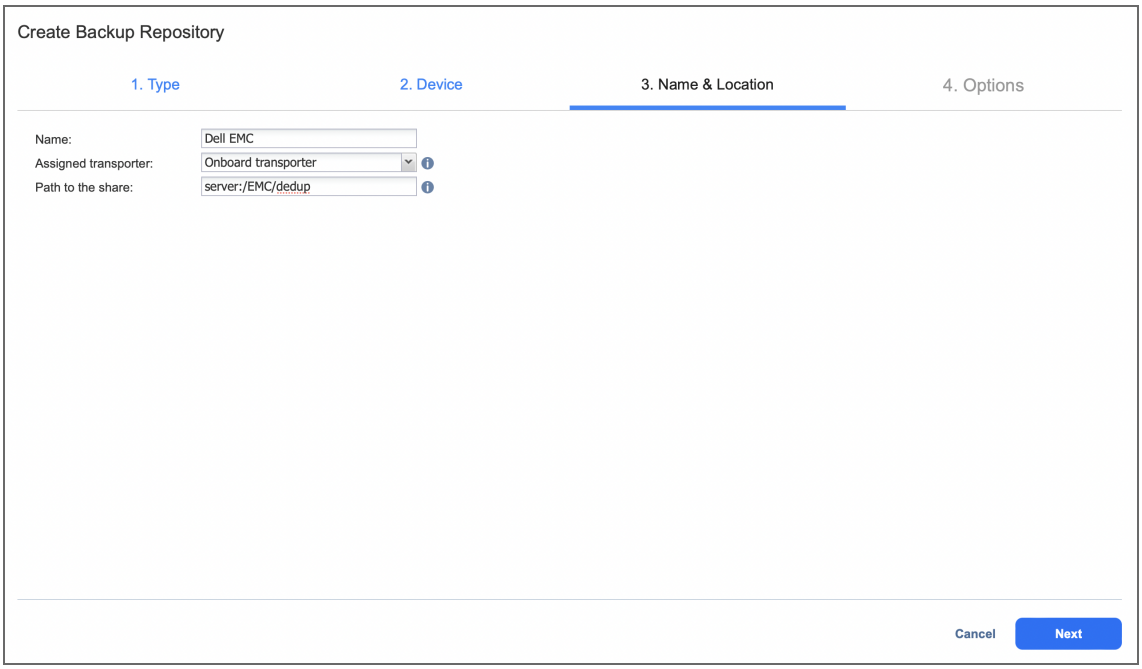
1. On the **Device** page, select one of the devices:
   - **Dell EMC Data Domain Boost**
   - **HPE StoreOnce Catalyst**
   - **NEC HYDRAstor**
2. Click **Next** to go to the next page of the wizard.



## Create Backup Repository: Name and Location

On the **Name & Location** page, specify the following:

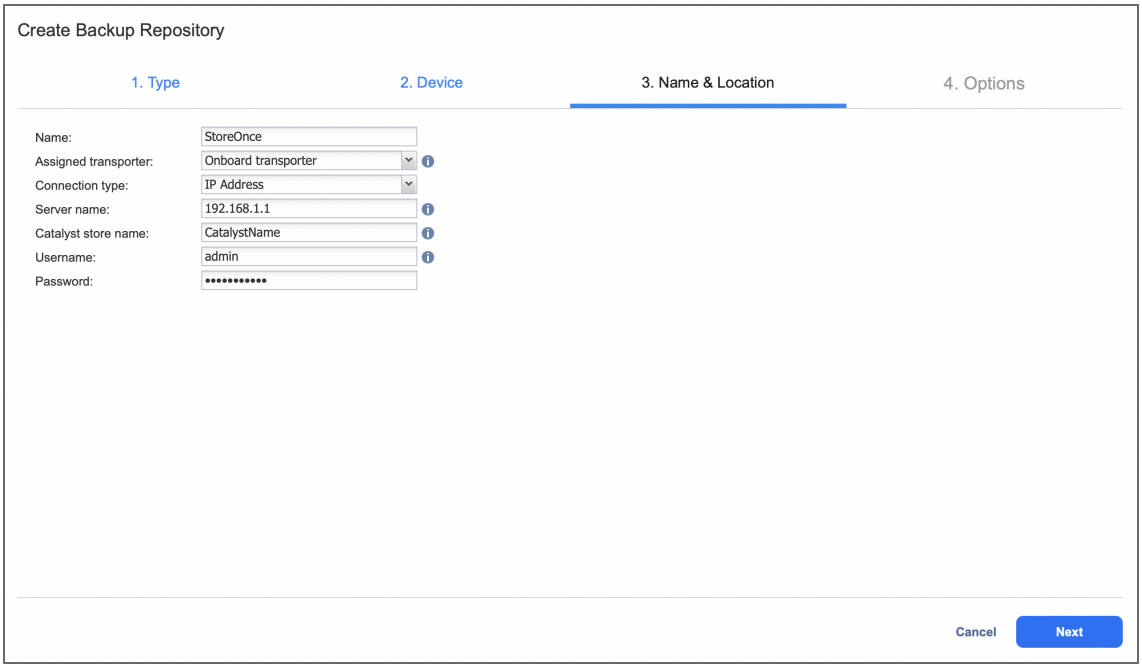
1. **Name**: Enter a name for the **Backup Repository**.
2. **Assigned transporter**: Choose a **Transporter** that will manage (that is, write data to and read data from) this **Backup Repository**.
3. Depending on the deduplication appliance, provide the following information:
    - **Dell EMC Data Domain Boost**
        1. **Name**: Enter the name of your **Backup Repository**.
        2. **Assigned transporter**: Select the assigned **Transporter**.
        3. **Path to the share**: Enter the path to the share folder in the following format: `<backup_appliance>:/<storage_unit>`. Refer to Creating a NAKIVO Backup & Replication Backup Repository on EMC Data Domain Backup Appliance for details.




- **HPE StoreOnce Catalyst**
    1. **Name**: Enter the name of your **Backup Repository**.
    2. **Assigned transporter**: Select the assigned **Transporter**.
    3. **Connection type**: Select one of the connection types to be used to access the **Backup Repository**:
        - **IP address**
        - **Fibre Channel**
    4. Depending on the connection type, do the following:

- **Server name** (if IP address connection type is selected): Enter the server name or IP address of the HPE StoreOnce Catalyst.
- **COFC identifier** (if Fibre Channel connection type is selected): Enter the COFC identifier. You can find your COFC identifier by going to **Catalyst Settings>Fibre Channel** in the **StoreOnce Management Console**.

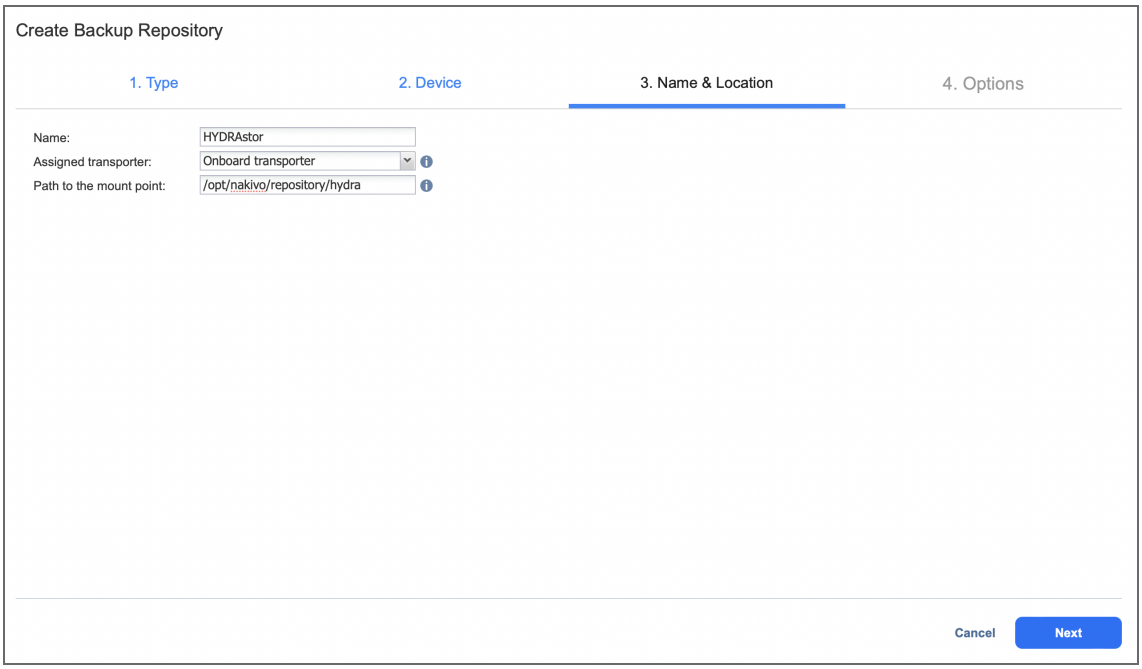
5. **Catalyst store name**: Enter the Catalyst store name.
6. **Username**: Provide the username to the Catalyst store.
7. **Password**: Provide the password to the Catalyst store.




- **NEC HYDRAstor**

1. **Name**: Enter the name of your **Backup Repository**.
2. **Assigned transporter**: Select the assigned **Transporter**.
3. **Path to the mount point**: Enter the path to the mount point in the following

format: `/opt/nakivo/repository/hsva`.




4. Click **Next** to go to the next page of the wizard.

## Create Backup Repository: Options

On the **Options** page, do the following:

1. Set up **Storage Savings & Encryption** options:
    - **Data size reduction**: If this option is enabled, NAKIVO Backup & Replication enables the use of data size reduction for this repository to save disk space. Note that this may put additional load on the CPU. Disabling data size reduction is required if the target is a deduplication storage appliance. Click **settings** to configure the settings. A popup window appears. Set the following:
        - **Compression**: Select a compression level that will be used to reduce the data size in the **Backup Repository**. Note that higher compression levels consume considerably more CPU and may slow down VM backup speed. The following options are available:
            - **Disabled:** The data in the **Backup Repository** will not be compressed.
            - **Fast:** Lowest compression level.
            - **Medium:** Medium compression level.
            - **Best:** Maximum compression level.

> **Note**
>
> This option cannot be configured after creating the **Backup Repository**.

- **Store backups in separate files**: Select this option to enable this backup repository to store data of every machine in separate backup files. Enabling this option is highly recommended to ensure higher reliability and performance. Leave this option unchecked if you wish to enable deduplication on a given backup repository.
- **Deduplication**: Select this option to enable the backup deduplication method to reduce the backup size by excluding duplicate data blocks from the backup.

> **Note**
> This option is not available if the Store backups in separate files checkbox has been selected.

- **Encryption:** This option is available only if the **Backup Repository** is created locally on the machine on which the Assigned **Transporter** is installed, and the machine is running a Linux OS. Select **Enabled** from the drop-down list and specify an encryption password. (The password will be required for importing the **Backup Repositor**y into a new instance of the product.) The product will encrypt the repository destination (using `ecryptfs` for folders and `cryptsetup (crypt-md)` in LUKS mode for devices/partitions) prior to creating the **Backup Repository**.

> **Note**
> - *Storage Savings and Encryption* settings are locked to the recommended settings.
> - To avoid `ecryptfs` errors, make sure that there are no other folders and files except the NakivoBackup folder in the repository location.
> - **Backup Repository** encryption can significantly influence backup speed.

2. Set up **Backup Immutability Support** options:
   - Enable backup immutability: If selected, you can create immutable backups on the Repository. The immutability period can be configured per Backup job, allowing granular control over retention policies. Depending on the storage system in use, additional configuration is required. Ensure that the corresponding feature is properly configured and supported on your storage device before enabling this option:
     - NEC HYDRAstor: enable the WORM feature on the NEC HYDRAstor system.
     - Dell EMC Data Domain: enable the Retention Lock feature on the Dell EMC Data Domain system.
3. Set up **Reliability & Maintenance** options:

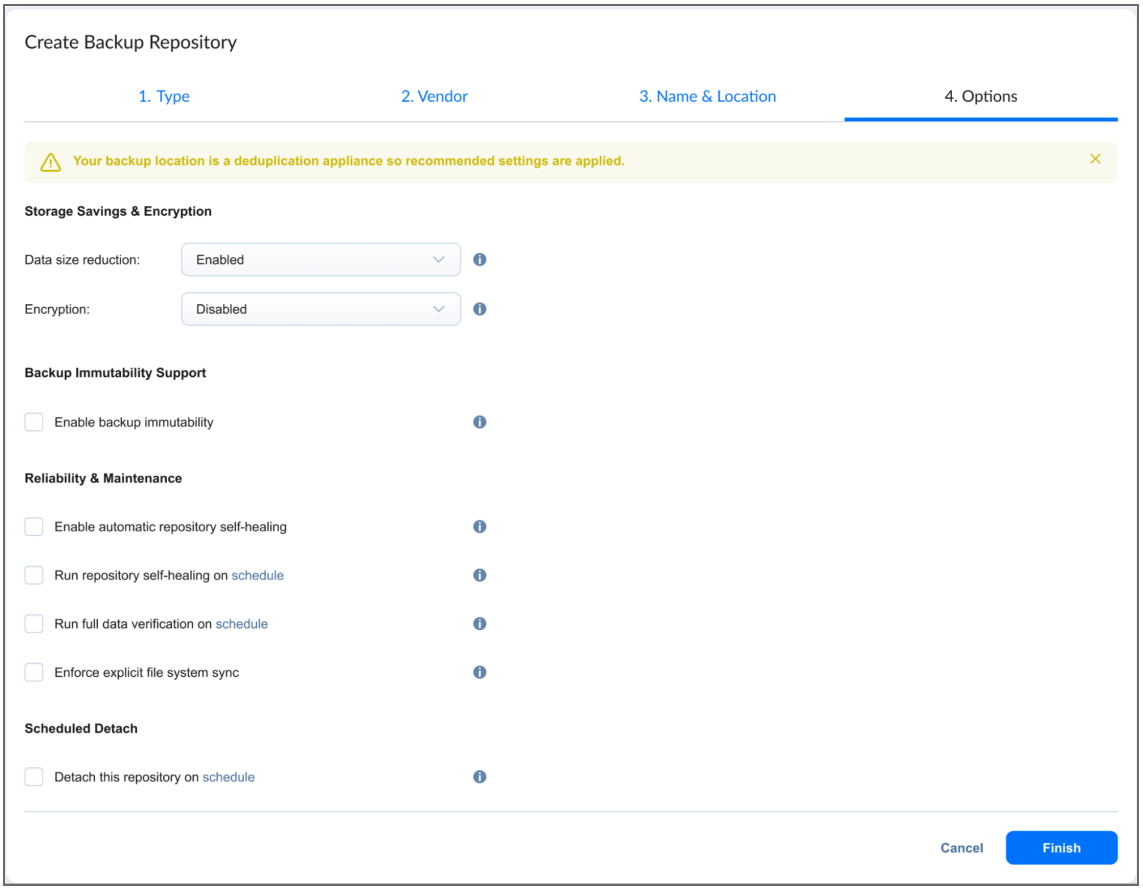
- **Enable automatic repository self-healing:** Leave this option selected to automatically trigger repository self-healing in case the product detects symptoms of problems in the backup infrastructure such as incorrect timestamps on metadata and data files. You can deselect this option and run self-healing manually.
- **Run repository self-healing on schedule:** If required, select this checkbox to run repository self-healing on schedule. You can configure the schedule by clicking the **schedule** link when the option is selected. The default schedule is set to run every day at 11 AM.
  If **Stop backup and recovery to run self-healing** is selected, any jobs or recoveries which use this repository will be stopped to run scheduled self-healing. Otherwise, scheduled self-healing will be skipped in case there are running jobs or recoveries on this repository.
  - **Run full data verification on schedule:** If selected, NAKIVO Backup & Replication will run full verification of all data available in the **Backup Repository** on the specified schedule. The product will read each block of data and ensure that it is identical to the data block that was read on the source VM during the backup. This way, the product will verify each recovery point in the **Backup Repository**.
    If **Stop backup and recovery to run backup verification** is selected, any running jobs which use this **Backup Repository** will be stopped to run scheduled data verification. Otherwise, scheduled data verification will be skipped in case there are running jobs on this **Backup Repository**.

    **Note**
    Backup verification is a time-consuming process and consumes CPU of the **Transporter** assigned to the **Backup Repository**. It is recommended that you schedule backup verification during non-working hours.
- **Enforce explicit file system sync**: When selected, explicit sync with the file system is enforced during all backup operations to this repository. This setting is considered more reliable but may lead to lower performance on certain storage devices. By default, the option is disabled.

4. Schedule detaching of the **Backup Repository**:
   - **Detach this repository on schedule:** Select this option if you want to detach and then attach the **Backup Repository** on a schedule. Detaching a **Backup Repository** saves its data and metadata in a consistent state and then stops the product's interaction with the **Backup Repository** (so that it can be copied or moved). You can use this feature, for example, for the disk-to-disk-to-tape (D2D2T) data protection approach, in which backups are stored on a disk for fast operational recovery, and copied to a tape (while the repository is detached) for archiving and long-term storage.

- **Delete and re-create the repository on attach:** If this option is selected, all data in the **Backup Repository** will be erased prior to attaching it to the product. As a result, jobs that write to this **Backup Repository** will create full VM backups. You can use this option, for example, to create full daily, weekly, or monthly VM backups and write them to tape or removable media.

5. Click **Finish** to finish creating the **Backup Repository**.



# Managing Backup Repositories

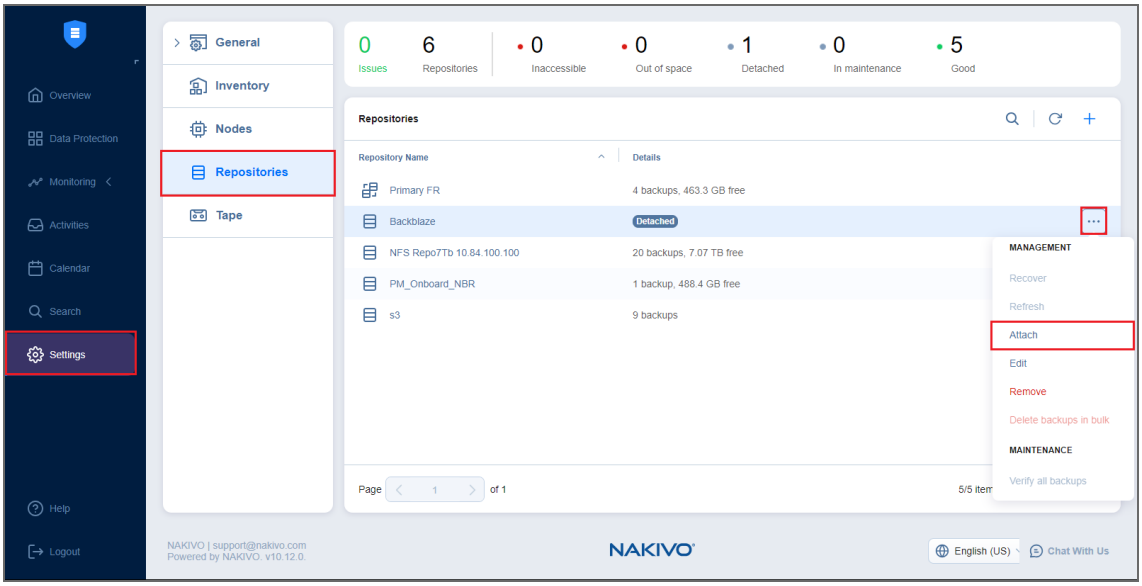
Refer to the following topics:


- “Attaching Backup Repositories” on page 666
- “Detaching Backup Repositories” on page 667
- “Editing Backup Repositories” on page 668
- “Reclaiming Backup Repository Space” on page 670
- “Refreshing Backup Repositories” on page 673
- “Removing and Deleting Backup Repositories” on page 675
- “Repairing Backup Repository” on page 677
- “Running Backup Repository Self-Healing” on page 680
- “Running Block-Level Backup Verification” on page 682
- “Managing Backup Encryption” on page 426

# Attaching Backup Repositories

If you have detached a **Backup Repository**, you can reattach it to the product by following the steps below:

1. From the main menu of NAKIVO Backup & Replication, click **Settings**.
2. Go to the **Repositories** tab and hover over a **Backup Repository**.
3. On the right side, click ••• and then click **Attach.**




The **Backup Repository** is reattached to NAKIVO Backup & Replication. You can now back up to the attached **Backup Repository**.

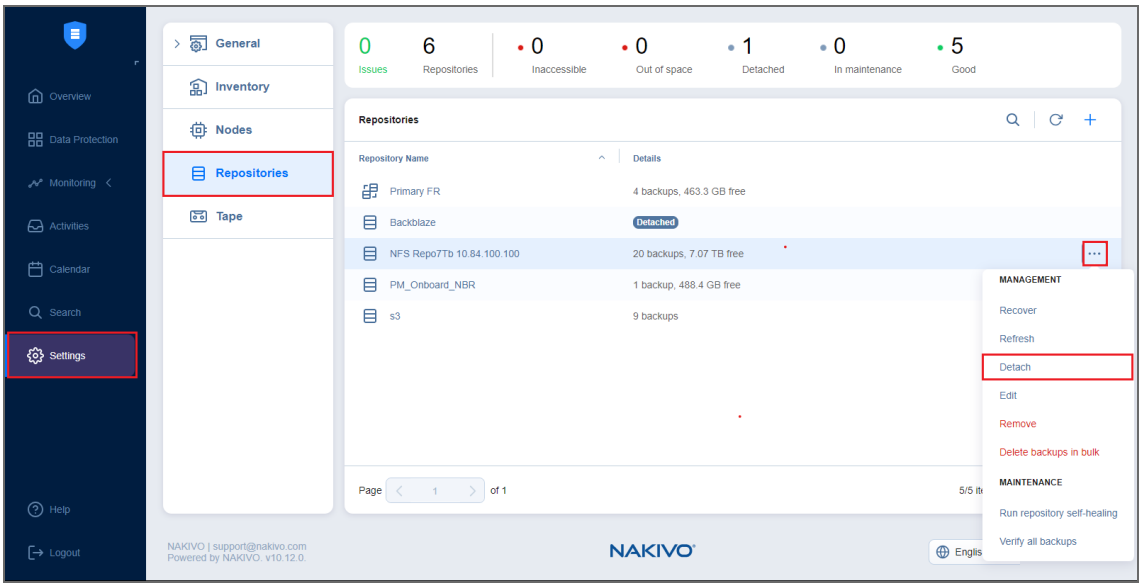
# Detaching Backup Repositories

Detaching a **Backup Repository** saves its data and metadata in a consistent state and stops the product's interaction with the repository (e.g. reading and writing of data or metadata). You may want to detach a **Backup Repository** in order to move it to a different location or to put the associated storage in maintenance.

> **Note**
> As the product does not interact with detached repositories, jobs with detached **Backup Repositories** as target storage will fail.

To detach a **Backup Repository**, follow the steps below:

1. From the main menu, click **Settings**.
2. Go to the **Repositories** tab and hover over a **Backup Repository**.
3. On the right side, click ••• and then click **Detach**.



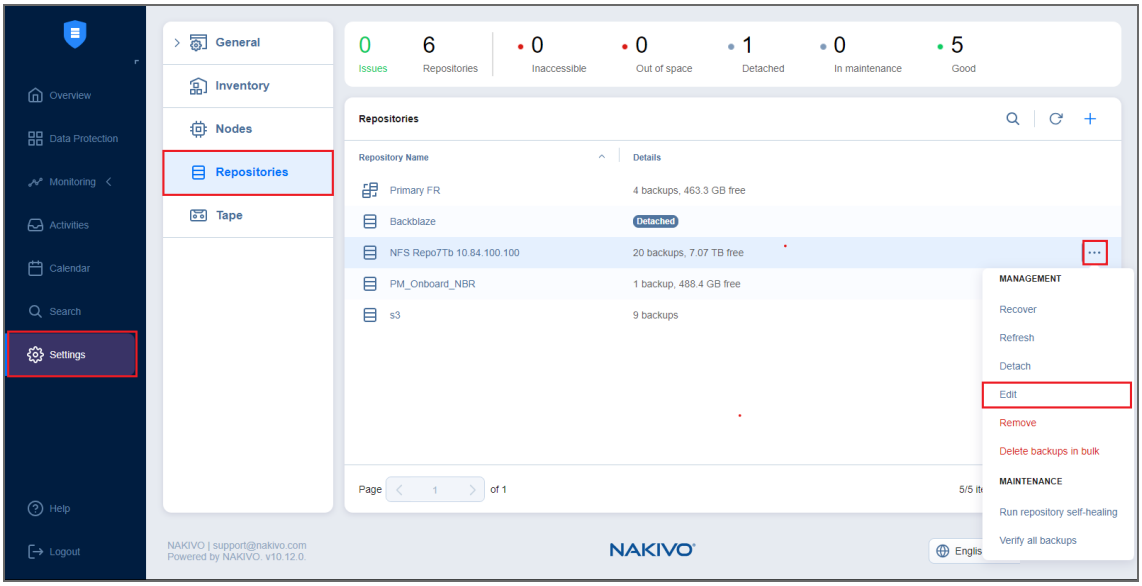

> **Note**
> A **Backup Repository** cannot be detached if a job that backs up to this **Backup Repository** is running.

The **Backup Repository** is detached from the product. You can reattach the **Backup Repository** to NAKIVO Backup & Replication when needed.

# Editing Backup Repositories

To modify the settings of an existing **Backup Repository**, follow the steps below:

1. From the main menu of NAKIVO Backup & Replication, click **Settings**.
2. Go to the **Repositories** tab and hover over a **Backup Repository**.
3. On the right side, click ••• and then click **Edit**.




> **Note**
>
> A **Backup Repository** cannot be edited while a job that backs up to this **Backup Repository** is running.

4. Update the fields as necessary.
5. Click **Apply**. Changes you have made are applied and the **Backup Repository** update starts.

# How to Copy Backup Repository to Tape

With NAKIVO Backup & Replication, you are able to use a disk-to-disk-to-tape (D2D2T) data protection approach. This approach allows to store backups on a disk for fast operational recovery and copy them to a tape for archival and long-term storage. To achieve this, you need to take these steps:

1. Create a Backup Repository on a disk or use the Onboard **Backup Repository** created with the product installation.

   > **Note**
   > By default, the Onboard **Backup Repository** stores backups in incremental and full backup files (**Store backups in separate files** option is enabled). If you want to store only incremental backups, you should create a new backup repository and configure it as forever incremental. This can be done by deselecting the **Store backups in separate files** option on the **Options** page of the **Create Backup Repository** wizard.

2. Create and run VM backup jobs to the **Backup Repository**.
3. After all backup jobs are complete, do either of the following:
   - Manually detach the **Backup Repository** to ensure its data is consistent.
   - Enable scheduled repository detach/attach in repository settings.
4. Copy the entire folder with the **Backup Repository** to a tape.

   > **Note**
   > To automate the folder copy process, you can use post-job scripts or 3rd-party utilities.

# Reclaiming Backup Repository Space

When a backup or recovery point is deleted in a **Backup Repository**, the space occupied by that backup or recovery point is marked as “free” and can be reused by new data blocks on the next job runs. However, the actual size of the **Backup Repository** may not change. The size of a **Backup Repository** can be reduced by rearranging the data blocks so there are no “free” ones occupying storage space. The amount of space that can be freed up is displayed in parentheses after the amount of used space. This is applicable if the repository type is **Forever-incremental**. Otherwise, if the repository type is **Incremental with full backups**, space reclaiming is not required. It is enough to delete the backups or recovery points to free up space and continue backing up to the repository.

For the incremental with full backup **Backup Repository** type, it is technically impossible to remove recovery points if there is no full backup after them. Make a full backup before deleting older recovery points.

Reclaiming free space can take the same amount of time as copying the entire **Backup Repository** to the storage where it is located (that is, if your repository size is 500 GB, reclaiming free space can take the same amount of time as copying 500GB of data to the storage where the **Backup Repository** is located).

Refer to the following topic to learn how to start and stop the reclaiming process:

- Starting the Space Reclaiming Process
- Stopping the Space Reclaiming Process

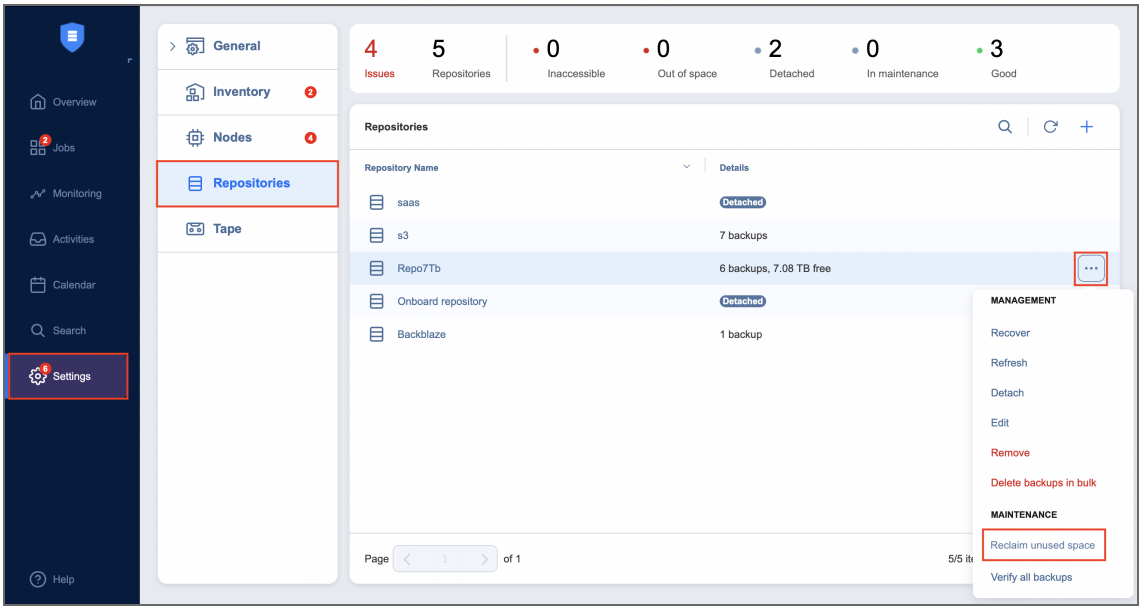
## Starting the Space Reclaiming Process

**Important**

Space reclaim requires at least 500 MB of free space on the repository storage in order to start.

To reclaim free space, follow the steps below:

1. Go to the main menu of NAKIVO Backup & Replication and click **Settings**.
2. Go to the **Repositories** tab and choose a **Backup Repository**.
3. In the title of the **Backup Repository**, click ••• and then click **Reclaim unused space.**




The space reclaiming process cannot be started if a job that backs up to this **Backup Repository** is concurrently running.

4. In the dialog box that opens, leave the **Interrupt space reclaim task if backup or recovery is started** option selected to pause the space reclaiming process when a backup or recovery is started. The space reclaiming process will be resumed once the backup or recovery job is completed. If you deselect the option, backup jobs will fail and recovery jobs will not start until the space reclaim process is completed.
5. Click **Start**. The process of rearranging data blocks is started, and progress is displayed in the title of the **Backup Repository**.

## Stopping the Space Reclaiming Process

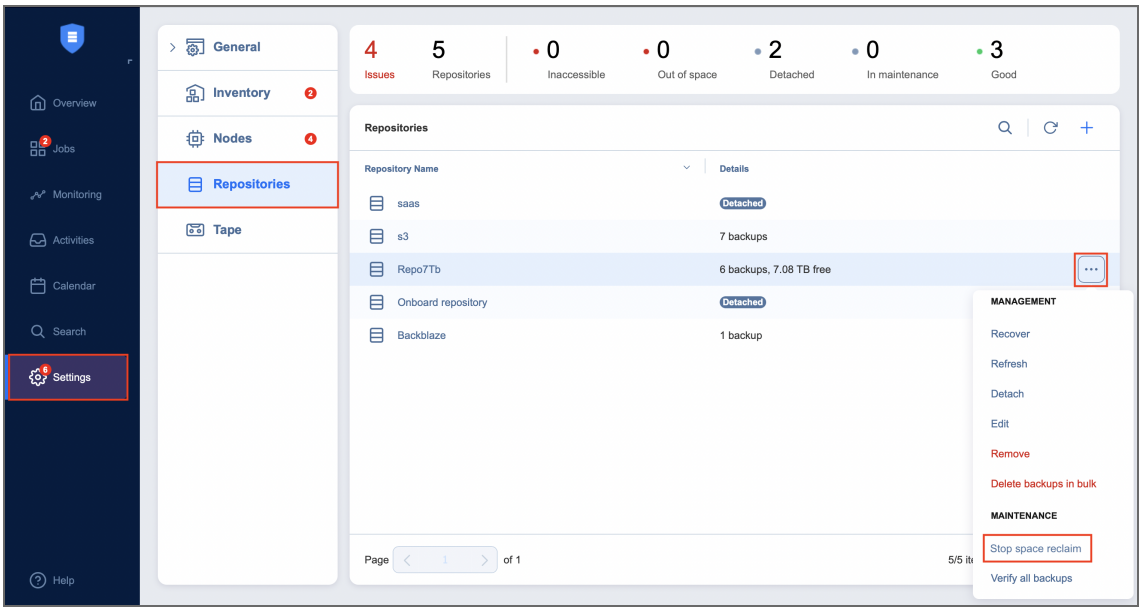
You can stop the space reclaim process at any time (for example to run a recovery job, move your **Backup Repository** to a new location, or put your backup storage on maintenance).

Before the space reclaiming process begins, the **Backup Repository** is detached from the product to keep data in a consistent state. The space reclaiming process stops if job that backs up VMs to such a **Backup Repository** is started and resumes after it is finished.

To stop the space reclaim process, follow the steps below:

1. Go to the main menu of NAKIVO Backup & Replication and click **Settings**.
2. Go to the **Repositories** tab and choose a **Backup Repository**.

3. In the title of the **Backup Repository**, click ••• and then click **Stop space reclaim**.



# Refreshing Backup Repositories

By default, NAKIVO Backup & Replication refreshes **Backup Repository** information hourly. During the refreshing process, the product collects all required information about **Backup Repositories**, such as the amount of free space, number of backups, and number of recovery points.

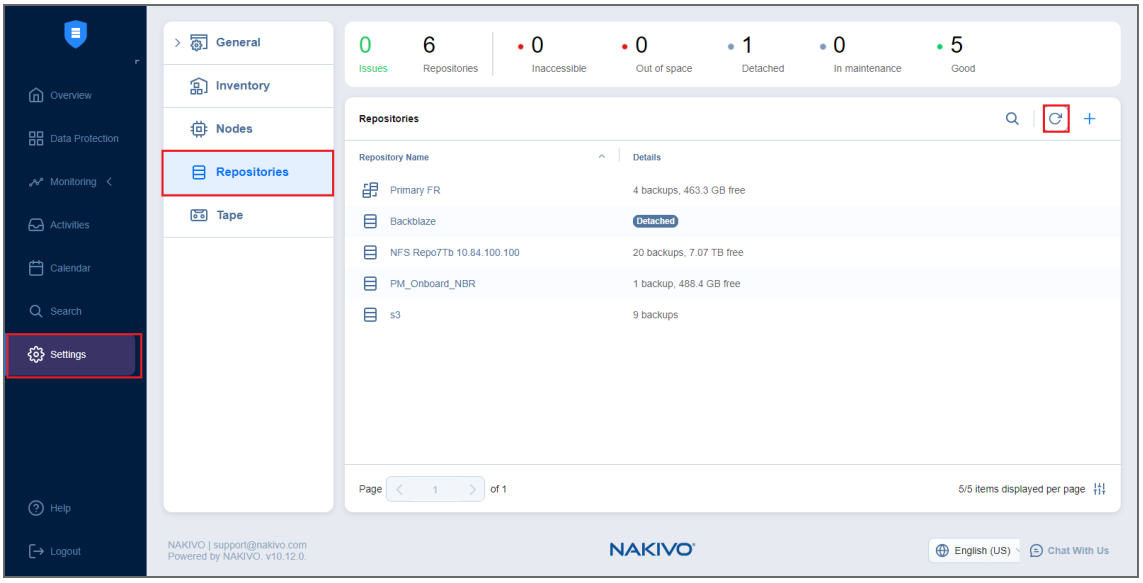
Only one **Backup Repository** is refreshed at a time. Therefore, if you attempt to refresh multiple **Backup Repositories**, all but one will be added to a queue.

- Refreshing All Backup Repositories
- Refreshing a Single Backup Repository

## Refreshing All Backup Repositories

To refresh all backup repositories, follow the steps below:

1. From the main menu of NAKIVO Backup & Replication, click **Settings**.
2. Go to the **Repositories** tab.
3. Click the **Refresh All** button.



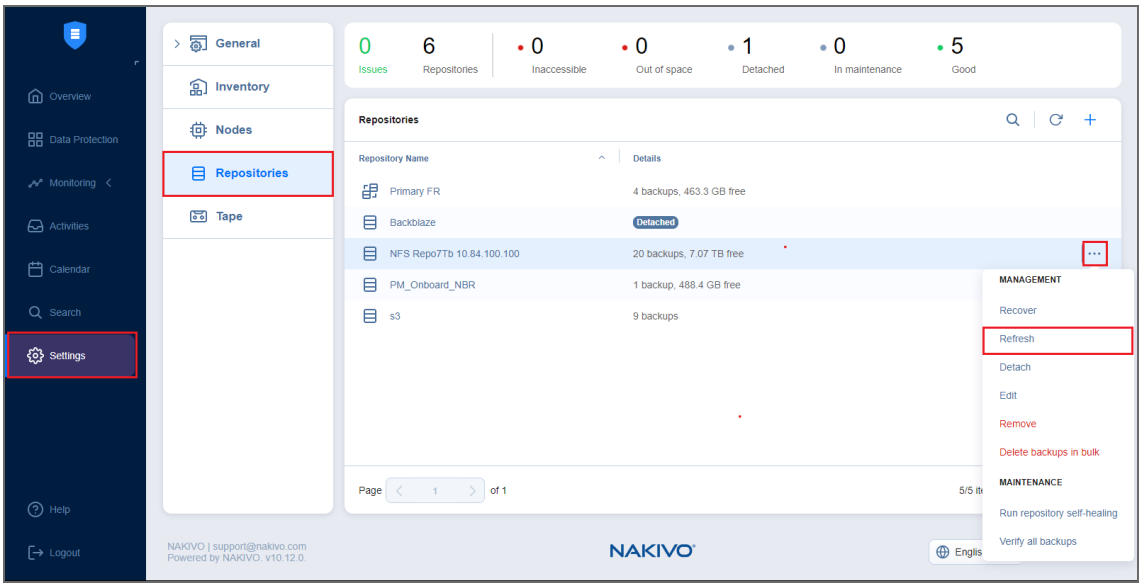

The **Backup Repository** refresh process begins.

## Refreshing a Single Backup Repository

To refresh a single **Backup Repository**, follow the steps below:

1. From the main menu of NAKIVO Backup & Replication, click **Settings**.
2. Go to the **Repositories** tab.
3. Hover over the **Backup Repository** that you wish to refresh and click •••.

4. Click **Refresh**.




The **Backup Repository** refresh begins.

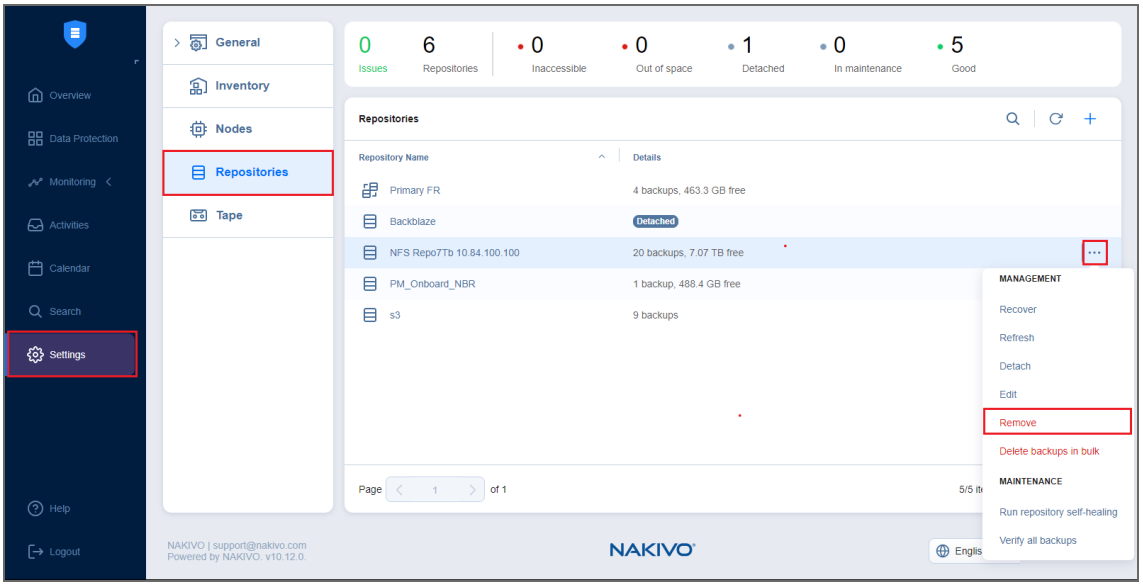
# Removing and Deleting Backup Repositories

In NAKIVO Backup & Replication, you can either permanently delete a **Backup Repository** and all of its data or remove only the **Backup Repository** from the product while maintaining all of its data. After removing a **Backup Repository** you will be able to import it into the same or a new instance of the product.

> **Note**
> You will not be able to remove a **Backup Repository** if there is a job that backs up to this **Backup Repository**. To remove such a **Backup Repository**, delete (or edit) the corresponding jobs so no items are backed up to the aforementioned repository.

To permanently delete or remove a **Backup Repository** from the product, follow the steps below:

1. From the main menu of NAKIVO Backup & Replication, click **Settings**.
2. Go to the **Repositories** tab.
3. Hover over a **Backup Repository.**
4. On the right side, click **Manage** and then click **Remove.**




5. Do the following when the confirmation message appears:
   - To remove the **Backup Repository** from NAKIVO Backup & Replication and keep the **Backup Repository** on a disk, select **Remove repository and keep backups**.

     > **Note**
     > You can import the removed **Backup Repository** back to the same instance or to a new installation.

   - To permanently delete the **Backup Repository** and all its data, select **Remove repository and delete backups**.

**Note**

This operation will permanently delete the **Backup Repository** and all its backups.

# Repairing Backup Repository

In case an immutable backup or the **Backup Repository** itself is corrupted, it is possible to initiate a repair process. During this process, NAKIVO Backup & Replication attempts to revert the **Backup Repository** or a specific backup to its uncorrupted state.

Refer to the following topics:

- Running the Repair Process for a Backup Repository
- Running the Repair Process for a Specific Backup Object

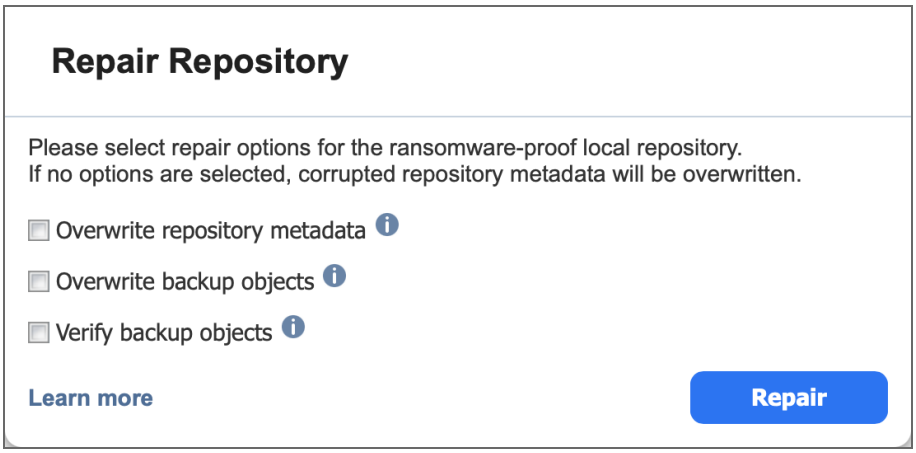
## Running the Repair Process for a Backup Repository

To run repair for a **Backup Repository**, do the following:

1. Go to **Settings > Repositories** and hover over the name of the **Backup Repository**.
2. Click ••• and select **Repair**. Alternatively, you can click on the name of the **Backup Repository** and then go to **Manage** > **Repair** to start the repair process.

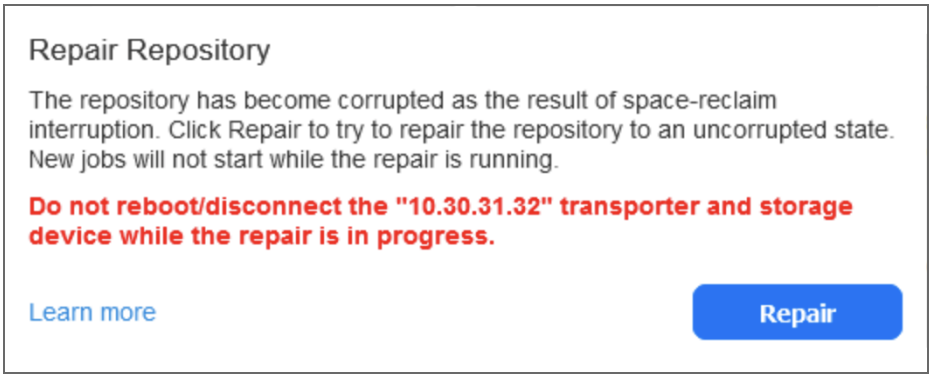
> **Note**
>
> The **Repair** option is only available in the following cases:
>
> - Non-immutable recovery points have been removed from the **Backup Repository** of the **Local folder** or **Amazon S3** type.
> - The local **Backup Repository** is inaccessible and meets the conditions specified in the feature requirements section.
> - A **forever incremental** repository becomes corrupted due to space reclaim interruption. This may occur as a result of rebooting the transporter assigned to the repository or disconnecting storage while space reclaim is in progress.

3. Select the desired options from the following:
   - **Overwrite repository metadata:** When this option is selected, the metadata file is overwritten even if it is present and valid. If the metadata file is not present, the new file is then created regardless of whether this option is selected or not.
   - **Overwrite backup objects:** When this option is selected, the locked backup objects are overwritten with the immutable data during the repair process.
   - **Verify backup objects:** When this option is selected, NAKIVO Backup & Replication runs verification of the backup object after the repair process is completed. When this option is not selected, NAKIVO Backup & Replication runs automatic self-healing after the repair process is completed.




**Note**

When initiating a repair for a **Forever Incremental** repository that has become corrupted as a result of space reclaim interruption, the following dialog will appear instead.

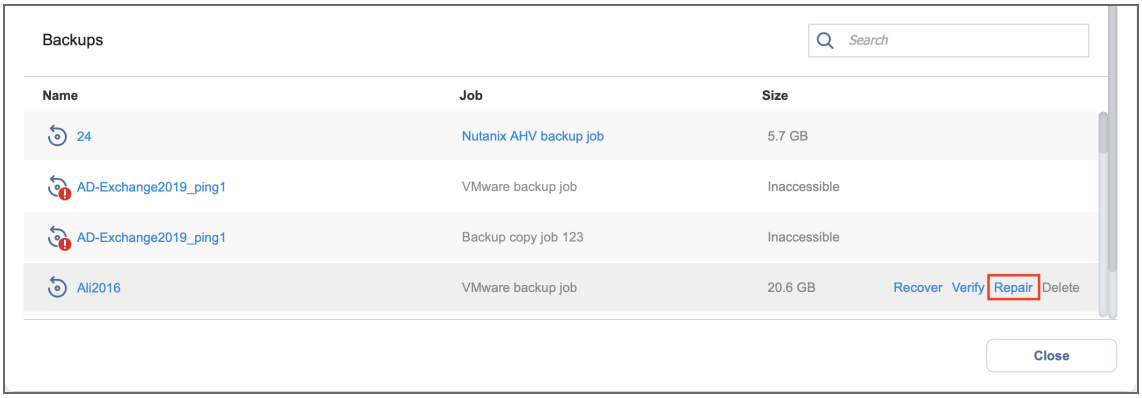



4. Click **Repair** to begin the repair process.

## Running the Repair Process for a Specific Backup Object

To run a repair for a specific backup object located in a **Local Folder** or **Amazon S3** type of **Backup Repository**, do the following:

1. Go to **Settings > Repositories** and can click on the name of the **Backup Repository**. Hover over the name of the backup and click **Repair** to start the repair process. Alternatively, you can click on the name of the backup and then click **Repair**.

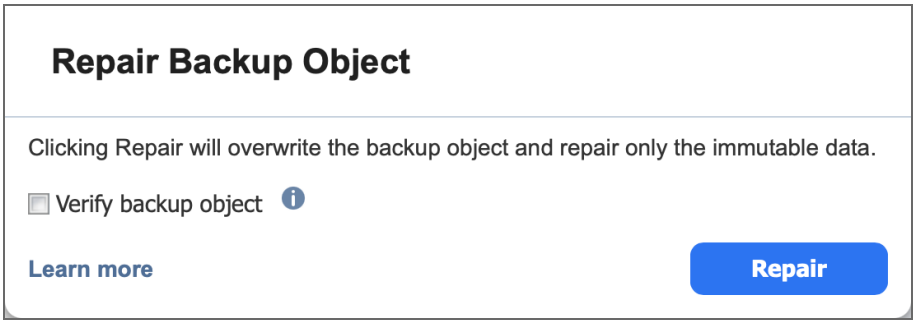



**Note**

You can also perform the **Repair** process for a backup object when all files except immutable files were manually deleted from the **Backup Repository**. The **Repair** option is only available in the following cases:

- The **Backup Repository** is inaccessible, was created in Amazon S3, and has **Object Lock** enabled.
- The local **Backup Repository** is inaccessible and meets the conditions specified in the feature requirements section.

2. Optionally, select the **Verify backup object** option. When you select this option, NAKIVO Backup & Replication runs verification on the backup object after the repair process has completed. In case **Verify backup object** is not selected, NAKIVO Backup & Replication runs automatic self-healing after the repair process is finished.




3. Click **Repair** to begin the repair process.

# Running Backup Repository Self-Healing

The self-healing process verifies **Backup Repository** integrity and automatically repairs errors wherever possible. Namely, the process performs the following tasks:

- Verifies that the data blocks of each recovery point are present in the **Backup Repository**.
- Cleans up “in progress” blocks of data from failed/crashed backup job runs that did not have a proper cleanup.
- Verifies and repairs Backup Repository metadata so that it correctly describes available data.
- Restores the consistent state of the **Backup Repository** to enable subsequent backup jobs.

Before the self-healing process begins, the **Backup Repository** is detached from the product to keep data in a consistent state. Jobs that back up VMs to such **Backup Repository** will fail while the self-healing process is in progress.

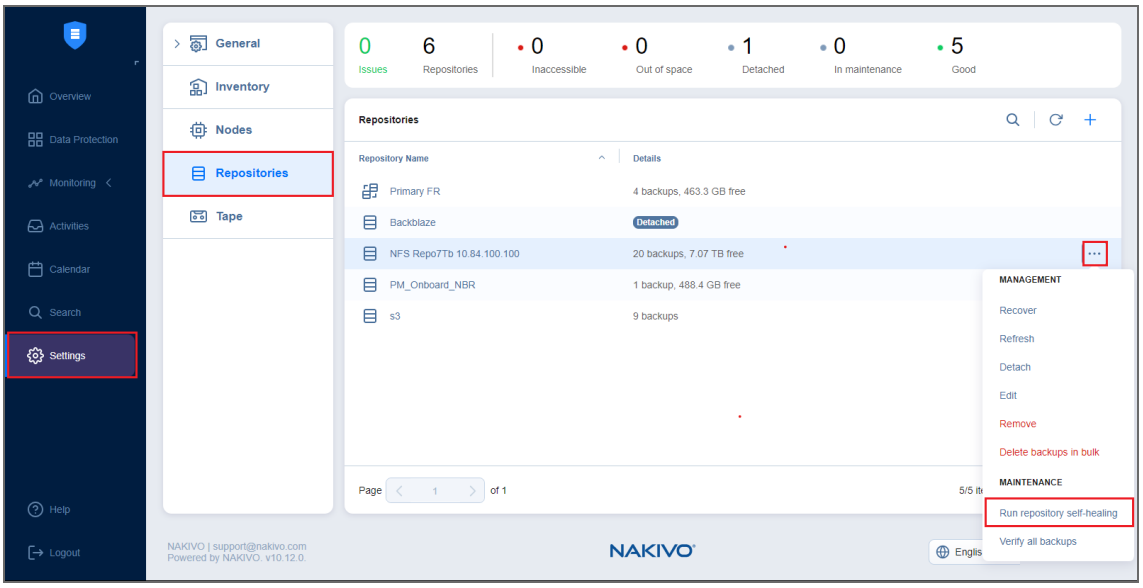
Refer to the following topics to learn more:

- “Starting the Self-Healing Process” below
- “Stopping the Self-Healing Process” on the next page

## Starting the Self-Healing Process

To run the **Backup Repository** self-healing, follow the steps below:

1. From the main menu of NAKIVO Backup & Replication, click **Settings**.
2. Go to the **Repositories** tab and hover over a **Backup Repository**.
3. On the right side, click ●●● and then click **Run repository self-healing.**



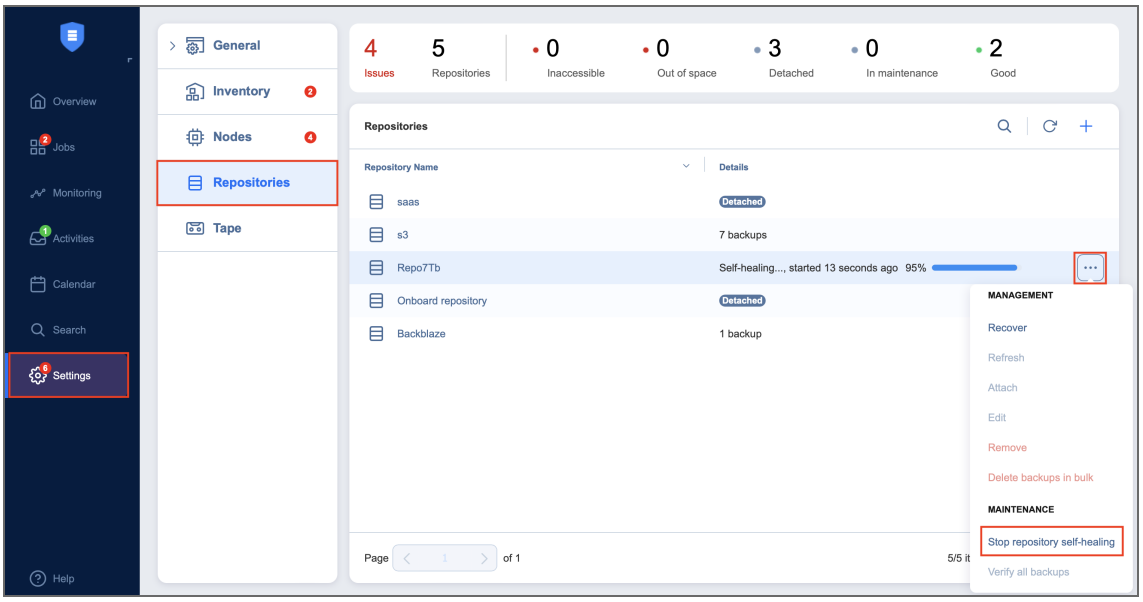

4. In the dialog box that appears, click **Start**. The self-healing process begins.

## Stopping the Self-Healing Process

You can stop the self-healing process at any time (for example, to run a recovery job, move your **Backup Repository** to a new location, or put your backup storage on maintenance).

To stop the self-healing process, follow the steps below:

1. From the main menu of NAKIVO Backup & Replication, click **Settings**.
2. Go to the **Repositories** tab and hover over a **Backup Repository**.
3. On the right side, click **Manage** and then click **Stop repository self-healing.** The self-healing process stops.



# Running Block-Level Backup Verification

Block-level backup verification reads each block of data in a **Backup Repository**, makes a hash of each data block, and then compares the newly created hashes to the originals that were created during the backup process. If the hashes match, this means that the data blocks in the **Backup Repository** are identical to the data blocks that were read on the source machines. This way, NAKIVO Backup & Replication verifies that backups are good and recoverable.

Refer to the following topics to learn more:


- "Verifying Backups" on the next page
  - "Verifying All VM Backups" on the next page
  - "Verifying a Single Backup" on page 684
- "Stopping the Backup Verification Process" on page 685
  - "Stopping Backup Verification for a Backup Repository" on page 685
  - "Stopping Backup Verification for a Single Backup" on page 685

# Verifying Backups

**Notes**

- Before backup verification begins, the **Backup Repository** is detached from the product to keep data in a consistent state. Backup jobs that write data to such a **Backup Repository** fail while the backup verification process is in progress.
- Self-healing is run prior to backup verification.
- **Backup verification** verifies all recovery points within the repository as follows:
  - Whether the metadata of each recovery point is in place and up to date
  - Whether the data of each recovery point is in place and corresponds to the checksums
  - Whether each recovery point is encrypted or not and, if encrypted, whether the hash required for its decryption is available

  The verification results display the state of all recovery points within the repository, which can be one of the following:
  - **OK**: Metadata is in place and up to date, all data is in place and correct.
  - **Not encrypted**: Recovery point is not encrypted.
  - **Encrypted**: Recovery point and the corresponding backup object are encrypted. The corresponding password hash required for its decryption is available or is not available.
  - **Corrupted**: Any of the metadata or data is not present or not correct.

  The results of recovery point full verification are marked as follows:
  - **OK recovery points**:
    - **Full verification**  icon
    - **Quick verification**  icon
    - Hover the mouse pointer over the icon to display the tooltip: *OK (verified on [maintenance end date, time, and timezone])*.
  - **Corrupted recovery points**  icon. Hover the mouse pointer over the icon to display the tooltip: *Corrupted (checked on [maintenance end date, time and timezone])*.

## Verifying All VM Backups

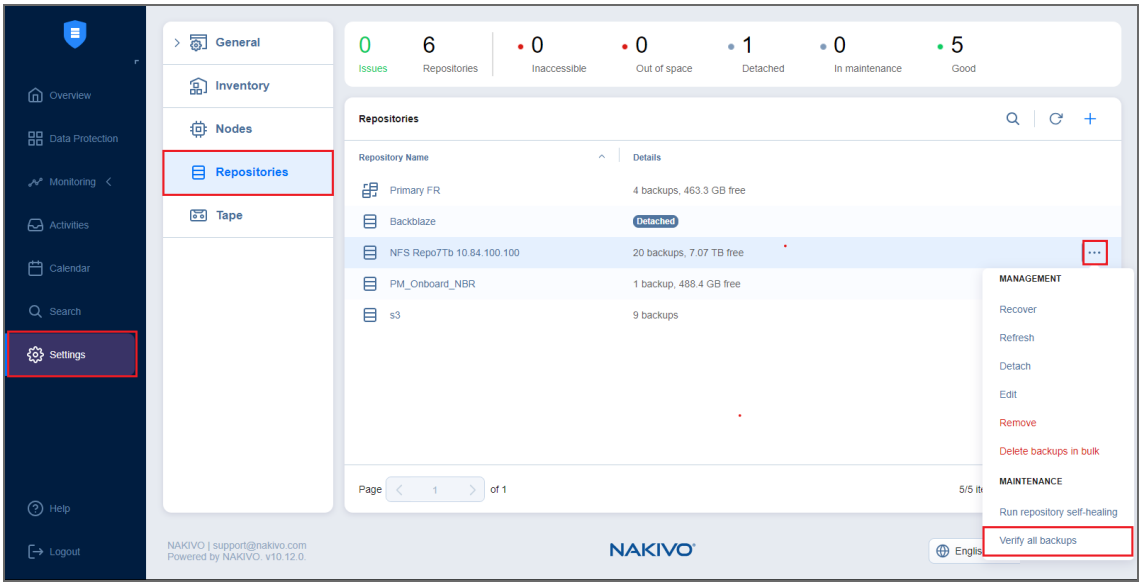
To verify all VM backups in a repository, follow the steps below:

1. From the main menu of NAKIVO Backup & Replication, click **Settings**.
2. Go to the **Repositories** tab and hover over a **Backup Repository**.

3. **On the right side, click Manage and then click Verify all backups.**

> **Note**
>
> The backup verification process cannot be started if a job that backs up to this **Backup Repository** is running.

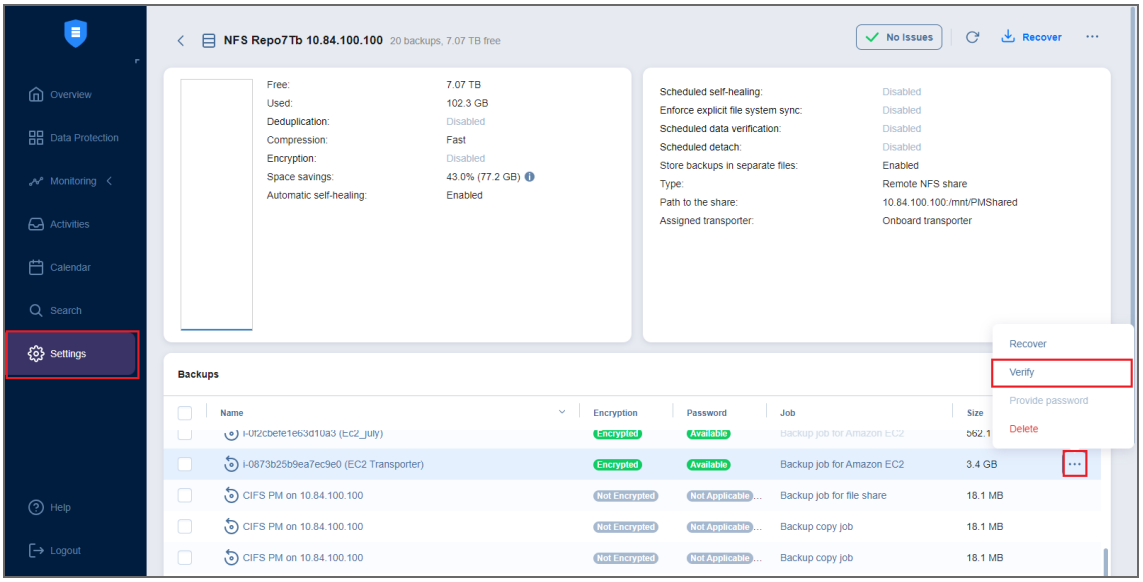
In the dialog box that opens, click **Start**. The backup verification process is started.




## Verifying a Single Backup

To verify a single backup in a repository, follow the steps below:

1. From the main menu of NAKIVO Backup & Replication, click **Settings**.
2. Go to the **Repositories** tab and click a **Backup Repository** to expand it.
3. Hover over the desired backup, click on the "**...**" button, and then select **Verify**.



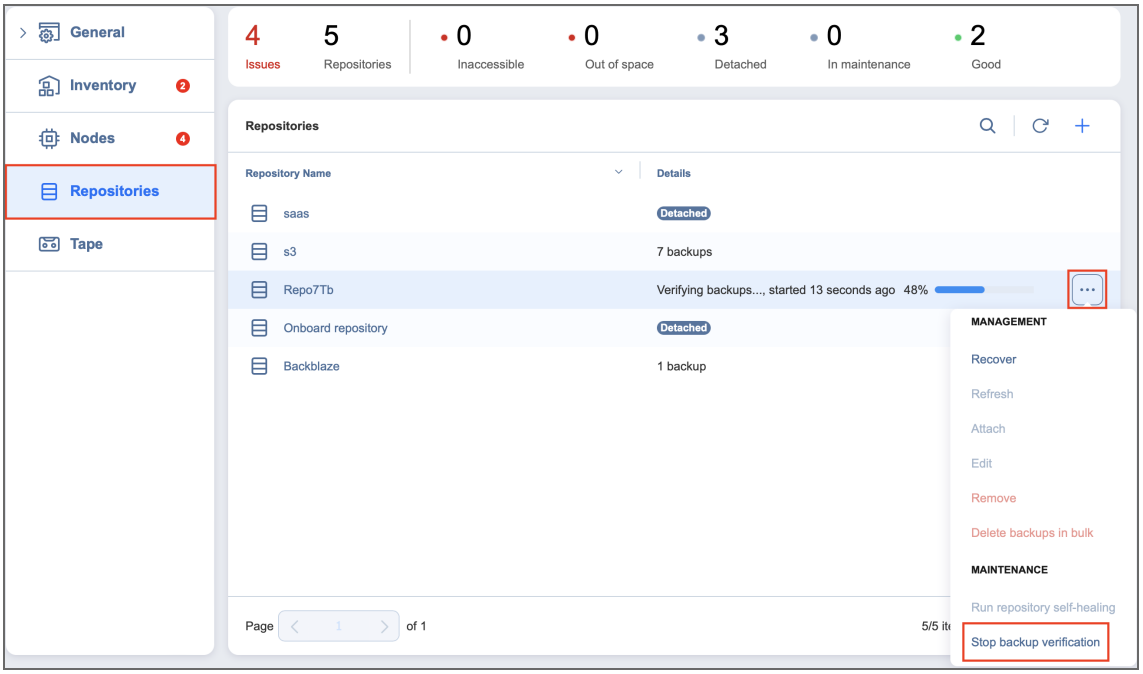
# Stopping the Backup Verification Process

You can stop the backup verification process at any time (for example, to run a recovery job, move your **Backup Repository** to a new location, or put your backup storage on maintenance).

## Stopping Backup Verification for a Backup Repository

To stop the backup verification process for a **Backup Repository**, follow the steps below:

1. From the main menu of NAKIVO Backup & Replication, click **Settings**.
2. Go to the **Repositories** tab and hover over a **Backup Repository**.
3. On the right side, click **Manage** and then click **Stop backup verification**. The backup verification process is stopped.

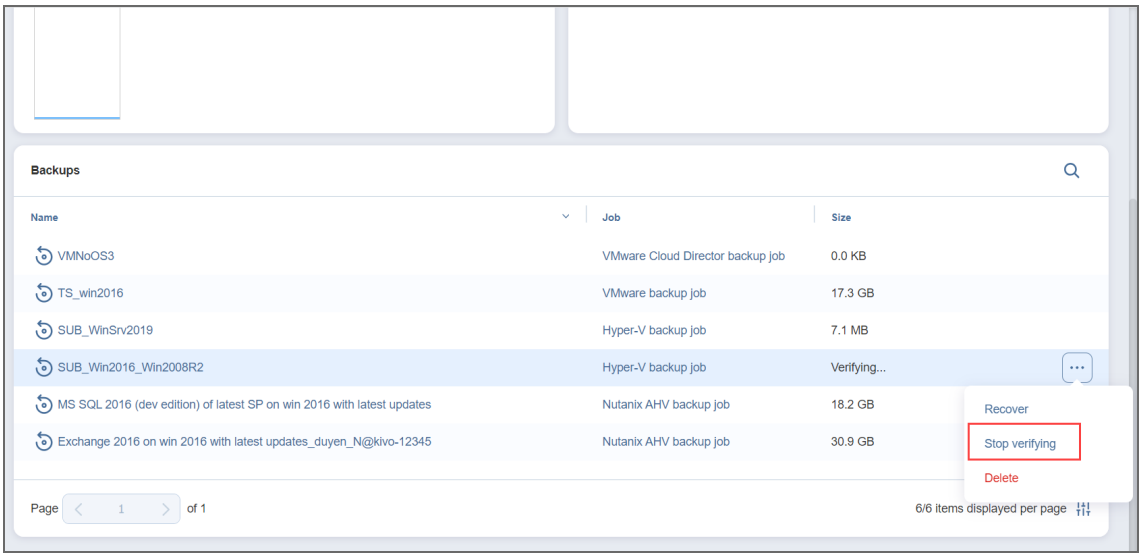



## Stopping Backup Verification for a Single Backup

To stop the backup verification process for a backup, follow the steps below:

1. From the main menu of NAKIVO Backup & Replication, click **Settings**.
2. Go to the **Repositories** tab and click a **Backup Repository** to expand it.

3. Hover over the desired backup, click on the "**...**" button, and then select **Stop verifying**.




For near-instant backup verification, refer to the “VM Verification” on page 66 feature.

# Viewing Backup Repository Details

To view the instance **Backup Repositories** details, follow the steps below:

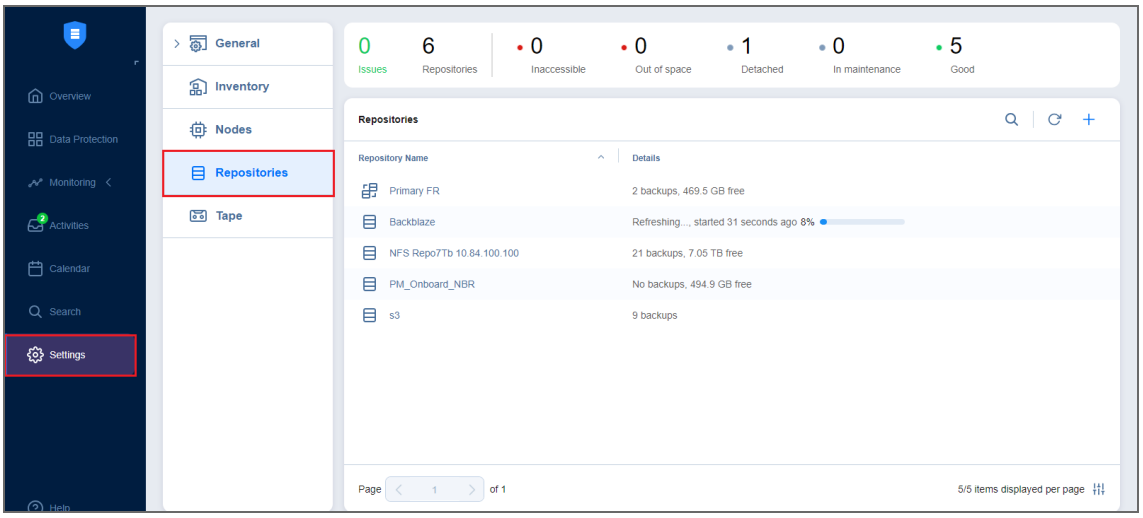
1. Go to the main menu of NAKIVO Backup & Replication and click **Settings**.
2. Go to the **Repositories** tab.

The **Repositories** tab contains a **Summary** bar, which offers an overview of all backup repositories both federated and standalone (not used as members of federated repositories). The data displayed is as follows:

- **Issues**: Total number of issues/alarms related to repositories
- **Repositories**: Total number of repositories
- **Inaccessible**: Total number of inaccessible repositories
- **Out of Space**: Total number of repositories that are out of storage space
- **Detached**: Total number of detached repositories
- **In Maintenance**: Total number of repositories in maintenance
- **Good**: Total number of usable repositories

To see information about specific standalone repositories, backups, and recovery points, see the sections below.

To see information about federated repositories and their members as well as the backups and recovery points assigned to them, refer to Viewing Federated Repository.




In this view, you can also search for a specific repository by entering its name (or part of its name) into the **Search** box.

In the **Repositories** table below, you can see information about specific repositories, their backups and recovery points.

The table data is displayed as follows:

- **Repository Name**: The name of the repository
- **Details**: The information about the repository:
  - Number of backups
  - Amount of free space in the backup repository
  - Repository status, which can be detached or inaccessible
  - The progress bar displayed when any of in progress statuses is in action (for example, self-healing or verifying backups)

To manage the **Backup Repository**, hover over it, and on the right side, click the ellipsis **Manage** button, and then select the necessary option in the drop-down list. For more details, refer to "Managing Backup Repositories" on page 665.

Refer to the following sections to learn more about viewing Backup Repository details:

- "Viewing Standalone Backup Repository Details" below
- "Viewing Backup Details" on page 690
- "Viewing Recovery Point Details" on page 693

## Viewing Standalone Backup Repository Details

To view the details for a standalone **Backup Repository**, follow the steps below:

1. Go to the main menu of NAKIVO Backup & Replication and click **Settings**.
2. Go to the **Repositories** tab.

> **Note**
> All repositories added to the product are listed in the **Repositories** table in the following order:
> - Federated repositories (for more details, refer to Federated Repositories)
> - Standalone repositories (not used as members of federated repositories)

3. Click a standalone **Backup Repository**  to open it in a new window.

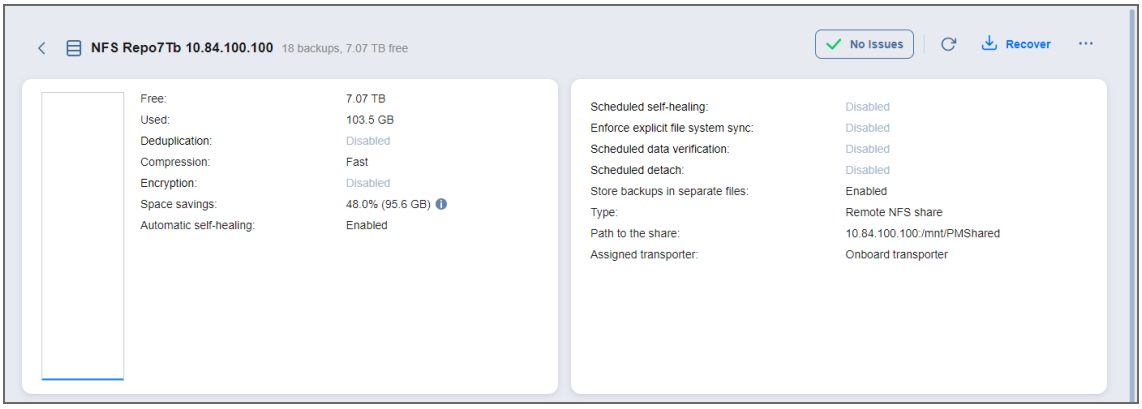
The following data is displayed:

- **Free**: The amount of free space available for the **Backup Repository**

> **Note**
> If the repositories are placed on the same same disk/share/etc., the free space metric can be incorrect.

- **Used**: The amount of space that the **Backup Repository** occupies on a disk. The amount of space that can be reclaimed is displayed in parentheses.
- **Deduplication**: The status of deduplication in the **Backup Repository**

- **Compression**: The compression level specified for the **Backup Repository**
- **Encryption**: The status of encryption in the **Backup Repository**
- **Space savings**: The estimated percentage and amount of space saved by compression and deduplication. For example, if 200 GB of data were backed up and the size of the backup was reduced to 50 GB, the ratio is calculated as 75%.
- **Automatic self-healing**: The state of the automatic self-healing option for the **Backup Repository**
- **Scheduled self-healing**: The state of the scheduled self-healing option for the **Backup Repository**
- **Enforce explicit file system sync**: The state of the enforce explicit file system sync option for the **Backup Repository**
- **Scheduled data verification**: The state of the scheduled data verification option for the **Backup Repository**
- **Scheduled space reclaiming**: The state of the scheduled space reclaiming option for the **Backup Repository**
- **Scheduled detach**: The state of the scheduled detach option for the **Backup Repository**
- **Store backups in separate files**: The behavior of the **Backup Repository** on backup data storage
- **Type**: The location of the **Backup Repository**, which can be one of the following:
    - **Local folder on assigned Transporter**
    - **Remote CIFS Share**
    - **Remote NFS Share**
    - **Amazon EC2**
    - **Microsoft 365**
    - **Microsoft Azure Blob Storage**
    - **Amazon S3**
    - **Generic S3-Compatible Storage**
    - **Wasabi**
    - **Backblaze B2 Cloud Storage**
    - **Deduplication Appliance**
- **Path to the folder**: The path to the **Backup Repository** folder
- **Assigned transporter**: The **Transporter** that manages the **Backup Repository** (that is, the **Transporter** that reads data from and writes data to the **Backup Repository**)
- **Backups**: List of available backups in the **Backup Repository**




**Note**

Standalone repositories used as federated repository members are not displayed in the **Repositories** table.

To view federated repository members, select the **Federated repository** to open the Federated repository details dashboard. In the **Storage** table, click the federated repository member name to open the **Backup Repository details** screen.

## Viewing Backup Details

Below, in the **Backups** table, you can view the details of the backups stored in the selected Backup Repository.

Large numbers of backups are split into pages. To find a specific backup, you can scroll through the pages manually or search for a specific backup by entering its name (or part of its name) into the **Search** box. The table dynamically changes to display the search results matching your query. Clicking the **Clear** button in the search box clears the query, and the table displays all backups.

The **Backups** table provides the following detailed information about each backup:

- **Name**: The name of the backup
- **Encryption**: The state of encryption of the backup object, which can be one of the following:
  - **Not encrypted**: None of the recovery points of the backup object are encrypted.
  - **Encrypted**: All the recovery points of the backup object are encrypted.
  - **Partially encrypted**: Some recovery points of the backup object are encrypted.
- **Password**: The state of the password of the backup object, which can be one of the following:
  - **Available**: All or some recovery points of the backup object are encrypted, and the password hash is available for all encrypted recovery points.

> **Note**
> In case of a self-backup, the status is displayed if all or some recovery points of the backup object are encrypted and the product database contains details of encryption for all encrypted recovery points.

- **Not available**: All or some recovery points of the backup object are encrypted, and the password hash is not available for all encrypted recovery points.

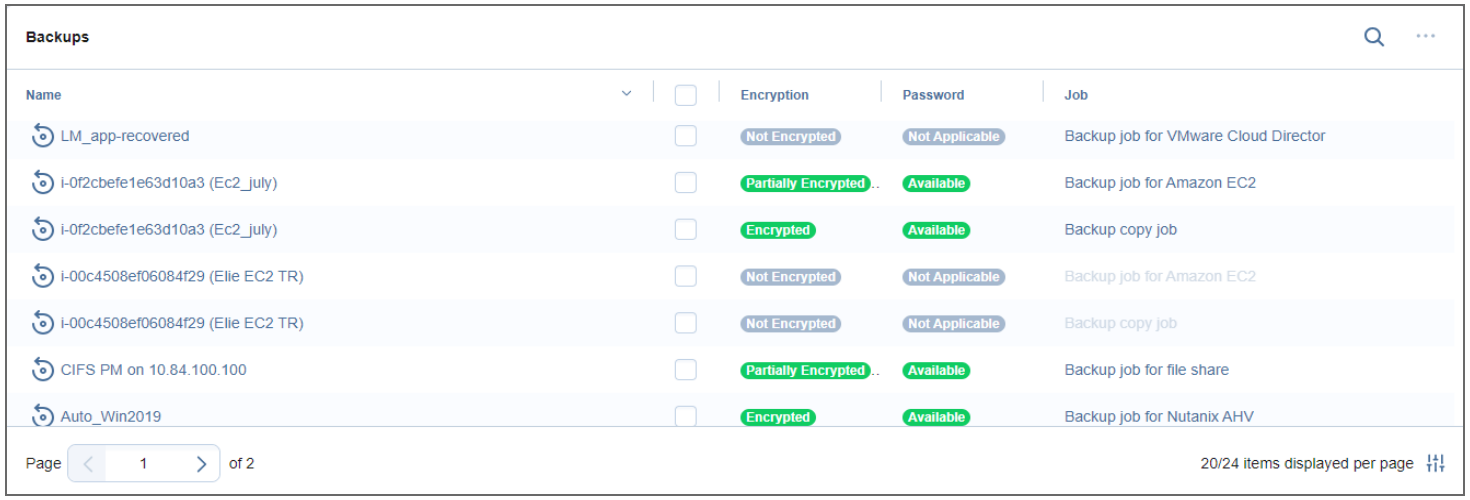
> **Note**
> In case of a self-backup, the status is displayed if all or some recovery points of the backup object are encrypted, and the product database does not contain the encryption details for all encrypted recovery points.

- **Partially available**: All or some recovery points of the backup object are encrypted, and the password hash is not available for some encrypted recovery points.

> **Note**
> In case of a self-backup, the status is displayed if all or some recovery points of the backup object are encrypted, and the product database does not contain the encryption details for some encrypted recovery points.

- **Not applicable**: The backup object does not contain encrypted recovery points.

- **Job**: The job type that created this backup
- **Size**: The total size of the backup




You can also initiate the recovery, verification, deletion, or password management process from here. Hover over the name of a backup and click the ellipsis **Manage** button on the right side to select one of the following options:

- **Recover**: Select this option to proceed with recovery.
- **Verify**: Select this option to verify the backup.
- **Provide password**: Select this option to manually provide the password if the currently selected backup is encrypted (some or all recovery points of the backup are encrypted) and the corresponding password hash is not available.

  Do the following:

  1. In the dialog box that appears, select the needed password or create a new one.

     Optionally, you can click the **Manage passwords** link to manage the existing or add a new password to the list of passwords.
  2. Click **Apply** to proceed.

  **Notes**

  - Optionally, you can select the checkboxes next to the backup names you want to manage, click the ellipsis **Manage** button, and then select **Provide password** in the drop-down list.
  - The **Provide password** option is disabled if the selected backups are not encrypted (none of the recovery points are encrypted).
  - The **Provide password** option is disabled if the selected backups are encrypted (some or all recovery points are encrypted) and the password hash is available for all encrypted recovery points.
  - It's recommended that you enable the **(AWS) Key Management Service** in the **Settings > General > System Settings > Encryption** tab. If **AWS** is enabled, all backup encryption passwords are encrypted with the Key Management Service cryptographic key to be available for recovery in case of product re-installation. For more information, refer to Enabling KMS.
- **Delete**: Select this option to delete the backup from the repository. Refer to "Deleting Backups" on page 924 for more details.

Clicking a backup name opens the **Backup Details** page where you can view the backup information and see all recovery points available for this backup.

The **Backup Details** section provides the following information about the backup:

- **Name**: The name of the backup item
- **Type**: The type of the backup
- **Tapes**: The number of tape cartridges that the backup occupies
- **Points**: The number of recovery points within the backup
- **Last point**: The date of the most recent recovery point of the backup
- **Size**: The total size of the backup

- **Job name**: The name of the job associated with the backup




## Viewing Recovery Point Details

You can view the details of a recovery point in the lower part of the screen. To find a recovery point for a specific date, you can use the **Search** bar on the right. The following information is displayed:

- **Date**: The date when the recovery point was created
- **Size**: The size of the recovery point
- **Type**: Type of backup used to create the recovery point
- **Schedule**: If applicable, the schedule that was used to create the recovery point
- **Encryption**: The state of encryption of the backup object, which can be one of the following:
  - **Not encrypted**: None of the recovery points of the backup object are encrypted.
  - **Encrypted**: All the recovery points of the backup object are encrypted.
- **Password**: The state of the backup object password, which can be one of the following:
  - **Available**:The recovery point is encrypted and the corresponding password hash required for its decryption is available.

    > **Note**
    > In case of a self-backup, the status is displayed if the recovery point is encrypted and the product database contains encryption details for this recovery point.

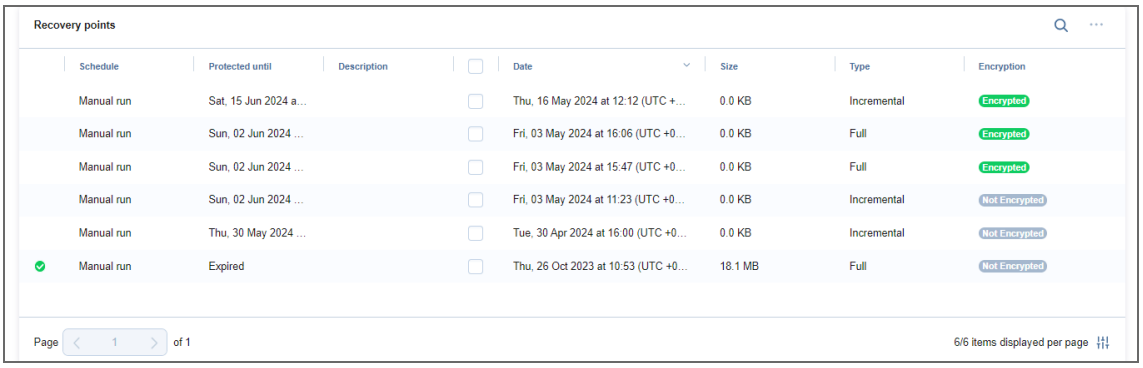
  - **Not available**: The recovery point is encrypted, and the corresponding password hash required for its decryption is not available.

**Notes**

- In case of a self-backup, the status is displayed if the recovery point is encrypted and the product database does not contain encryption details for this recovery point.
- In case of a self-backup, the status is displayed if you are upgrading your product from the version that does not support backup encryption to the version that supports it (for example, from 10.11 to 11.0).
- After a system configuration backup, the encrypted recovery points passwords that were created and saved in the password manager become unavailable and providing them manually is required.

    - **Not applicable**: The recovery point is not encrypted.

- **Immutable until**: If applicable, the date when the recovery point immutability expires
- **Protected until**: The date until which the recovery point is retained, displayed only for recovery points belonging to jobs that use the schedule retention approach
- **Description**: The description of the recovery point if one was provided




**Notes**

- The **Size**, **Type**, and **Immutable until** details are displayed only if the **Store backups in separate files** option (under **Storage Savings & Encryption**) is selected when creating or editing a Backup Repository.
- For recovery points belonging to jobs using legacy retention settings, **Use job retention** is displayed under **Protected until** instead.

**Date**, **Type**, and **Description** can also be viewed when selecting recovery points in the **Recovery Job Wizard**. Hover over the name of the recovery point and click the ellipsis **Manage** button on the right side to select one of the following options:

- **Recover**: Select this option to proceed with recovery.
- **Edit**: Select this option to edit the recovery point. Do the following:

- Optionally, you can add a **Description** to your recovery point.
- Choose the date until which the recovery point should be kept. The following options are available:
  - **Use job retention**: Choose this option to use the retention settings configured in the job for this recovery point.
  - **Keep forever**: Choose this option to keep this recovery point forever.
  - **Protect until**: Choose this option to keep this recovery point until a specific date. After selecting this option, choose the date in the calendar pop-up.

- **Provide password**: Select this option to manually provide the password if the currently selected recovery point is encrypted and the corresponding password hash is not available.

  Do the following:
  - In the dialog box that appears, select the needed password or create a new one.
  - Optionally, you can click the **Manage passwords** link to manage the existing or add a new password to the list of passwords.
  - Click **Apply** to proceed

  **Notes**
  - Optionally, you can select the checkboxes next to the recovery points you want to manage, click the ellipsis **Manage** button, and then select **Provide password** in the drop-down list.
  - The **Provide password** option is disabled if the selected recovery points are not encrypted (none of the recovery points are encrypted).
  - The **Provide password** option is disabled if the selected recovery points are encrypted (some or all recovery points are encrypted) and the password hash is available for all encrypted recovery points.
  - The **Provide password** option is disabled for self-backup recovery points created from an older version that does not support backup encryption.
  - It's recommended that you enable the **(AWS) Key Management Service** in the **Settings > General > System Settings > Encryption** tab. If **AWS** is enabled, all backup encryption passwords are encrypted with the Key Management Service cryptographic key to be available for recovery in case of product re-installation. For more information, refer to Enabling KMS.

- **Delete**: Select this option to delete the recovery point from the repository.

**Note**

To learn about the limitation on the chain of incremental recovery points, refer to the Knowledge Base article.

# Federated Repositories

A **Federated repository** is a backup repository that consists of one or more existing backup repositories (members) with identical attributes. A **Federated repository** can be scaled horizontally to consume more backup repositories. Thus, backup jobs can continue to run even if one of the members is unavailable or has insufficient storage space. For detailed information, refer to Federated Repository.

This section covers federated repository-related topics such as creation, management, etc. of **Federated repositories** and contains the following articles:

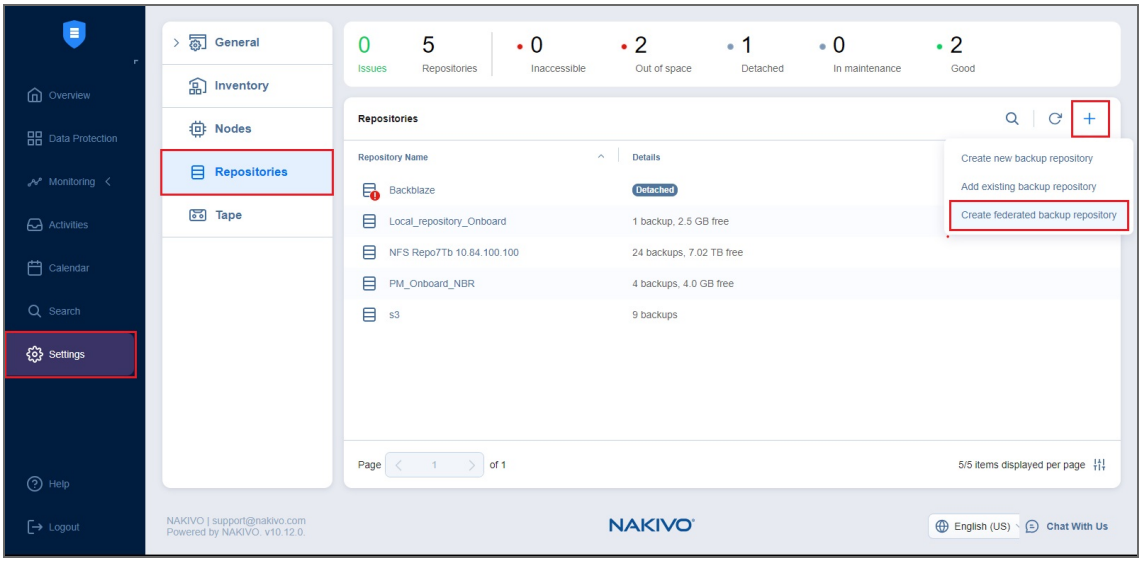
- Creating Federated Repositories
- Viewing Federated Repository Details
- Managing Federated Repositories

## Creating Federated Repositories

NAKIVO Backup & Replication allows you to create federated repositories for storing backups.

To create a new federated repository, follow the steps below.

1. In NAKIVO Backup & Replication, navigate to **Settings**.
2. Go to the **Repositories** tab and click **+**.
3. Click **Create federated backup repository**.



**Important**

- The **Federated repository** feature supports backing up data to **Incremental with full backups** repositories only.
- The **Create federated backup repository** option is disabled if no backup repositories of the supported type have been added to the product.

4. Proceed by completing the **Create Federated Backup Repository** wizard as described in the following sections:
    - Creating Federated Repository: Members
    - Creating Federated Repository: Options

# Creating Federated Backup Repository: Members

On the **Members** page of the **Create Federated Backup Repository** wizard, use the table of available standalone repositories to choose and define one or more existing backup repositories (members) to be added to a federated repository.

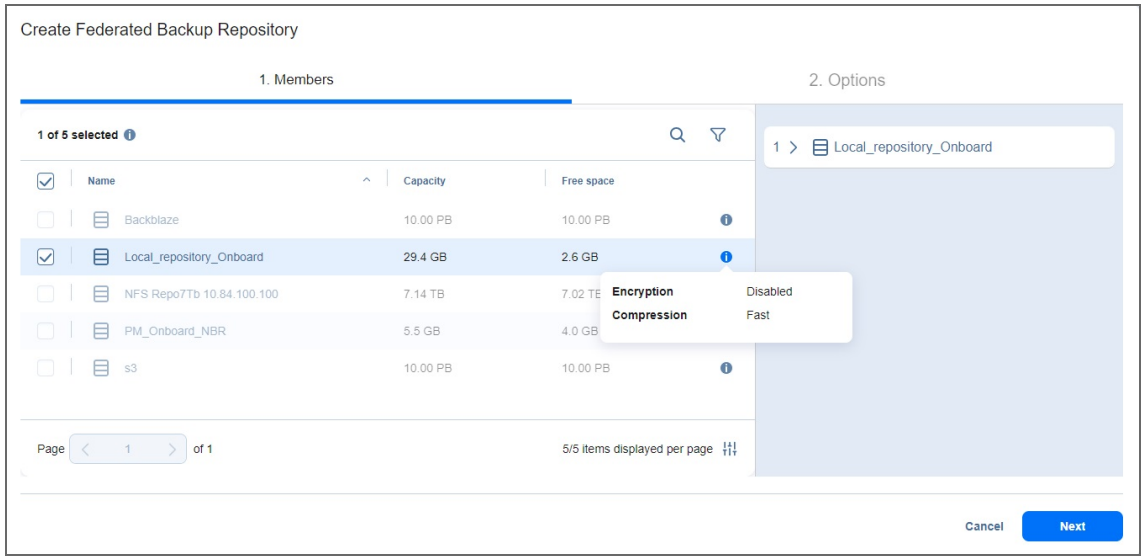
The table displays all available repositories with the following columns:

- **Repository name**: The name of the repository
- **Capacity**: The repository capacity in GB (gigabytes)
- **Free space**: The repository free space in GB (gigabytes)

Click the **Info** icon next to the repository to open a popup displaying the following information about the following supported federated repository member attributes:

**Encryption** (*Enabled* or *Disabled*)

**Compression** (*Disabled*, *Fast*, *Medium*, or *Best*)



**Notes**

- Only **Incremental with full backups** repositories are available for selection.
- Repositories that are already used as members of other federated repository(ies), repositories (and their members) containing a self-backup or selected as a self-backup storage, and those assigned to tenants are disabled.
- You cannot simultaneously select:
  - the members that support immutability and contain at least one immutable object or are associated with job(s) configured to create an immutable recovery point(s) and
  - the members that do not support immutability
- Selecting either type disables the member(s) of the other type.
- Federated repository members can use different transporters.

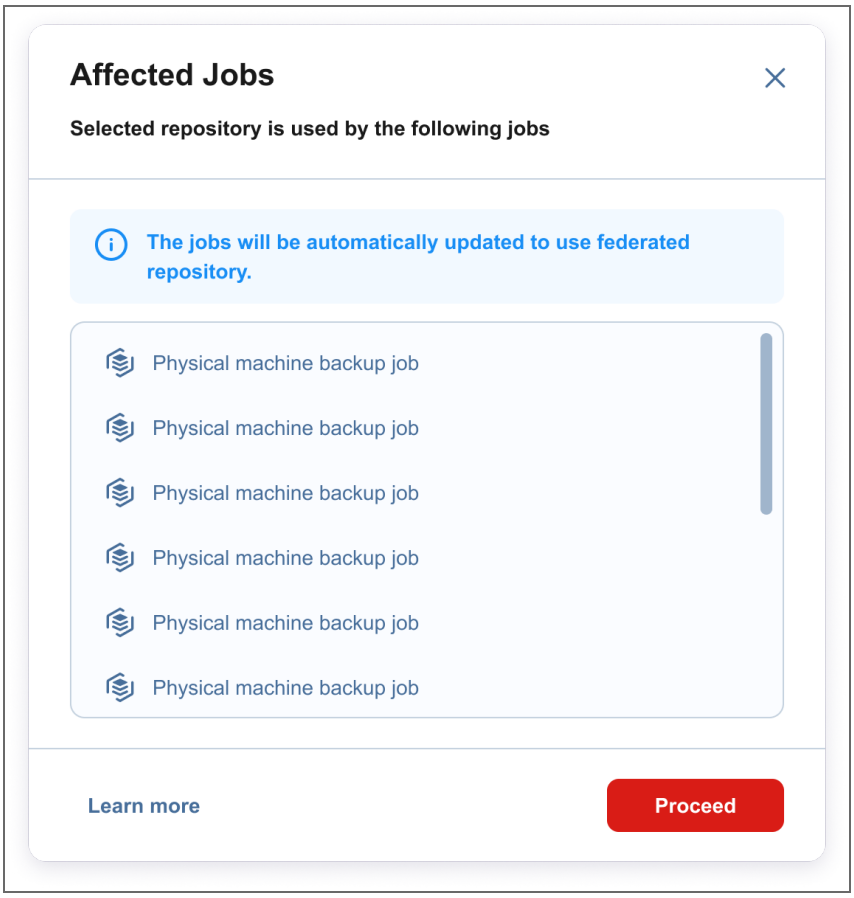
To choose and/or define one or more existing backup repositories (members) to be added to your federated repository, follow the steps below:

1. Select the needed repository item(s) by placing a checkmark to the left of the item(s).
2. Optionally, click the **Select All** checkbox to select all the enabled repositories listed.
3. The selected item(s) appear in the right pane of the page. You can remove a selected repository member in one of the following ways:
   a. Deselect the item in the left pane. This removes the repository member from the right pane; OR
   b. In the right pane, hover the pointer over the item you want to remove and click the "**X**" icon. This deselects the repository member in the left pane.

**Notes**

- If a federated repository is selected as the destination, the product automatically assigns an available federated repository member during the job run.
- Optionally, you can search members by clicking the **Search** button. This displays the search field where you can enter the partial or full name of the member.
- Click on the **Filter** button to filter the repository tree items by **Name** or **Type** (*Local folde*r, *NFS share*, or *CIFS share*).
- Federated repository members can be later added/removed to/from the federated backup repository or managed/refreshed individually.

4. If the selected member is used as a destination for existing jobs, the **Affected Jobs** dialogue is displayed with the list of jobs using the selected repository as a destination or source. Click **Proceed** to lock the selected repository before adding it as a new member and update the jobs after the federated repository is created.



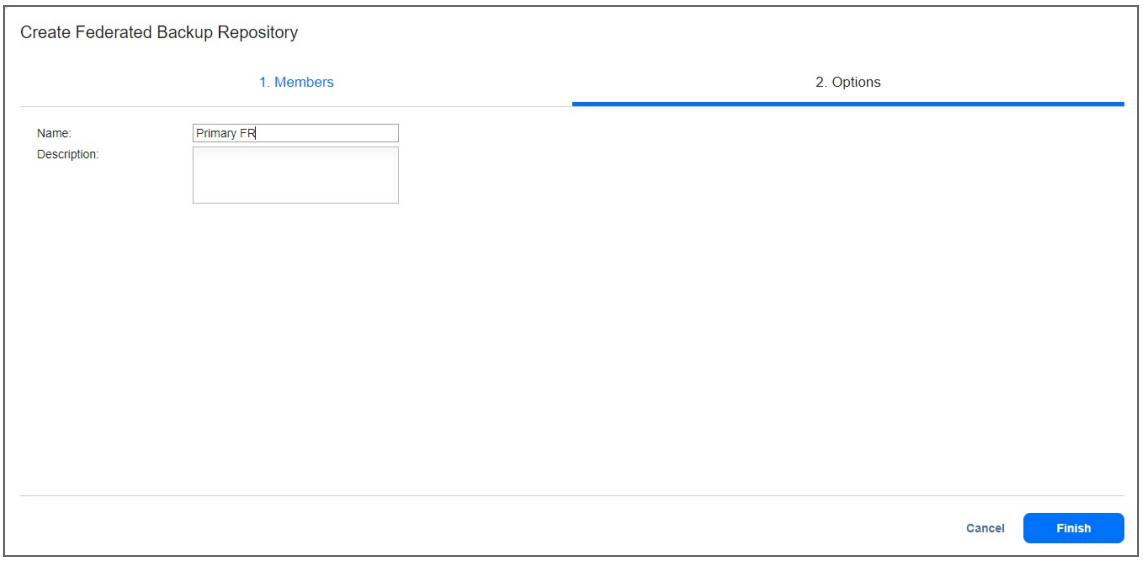

**Notes**

- After a backup repository is added as a member to a federated repository, it cannot be selected as a target for a new backup/backup copy job.
- After a federated repository is saved, the jobs using the repositories now selected to be part of the federated repository are automatically reconfigured to use this federated repository as a target.

5. Click **Next** to move to the next page of the wizard.

## Creating Federated Backup Repository: Options

1. On the **Options** page, provide the name and description (optional) for the federated backup repository.
2. Click **Finish** to confirm saving the federated repository.




You can now view the federated repository by going to the **Settings > Repositories** tab.

**Note**

After a federated repository is saved, the jobs using the repositories now selected to be part of the federated repository are automatically reconfigured to use this federated repository as a target.

# Managing Federated Repositories

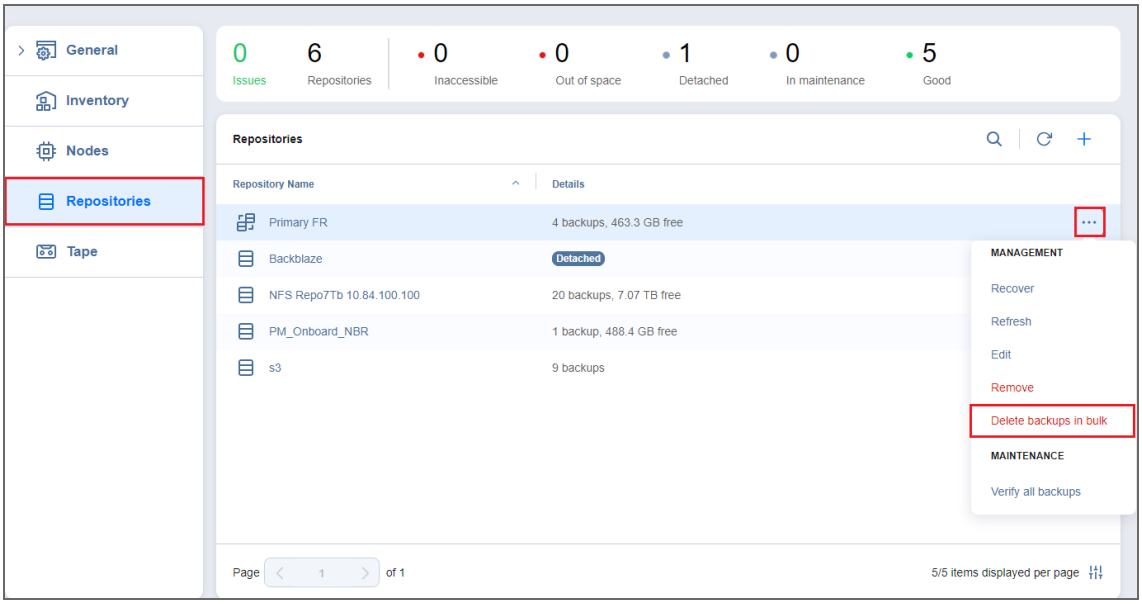
For more detailed information about managing federated repositories, refer to the following topics:

- Refreshing Federated Repositories
- Editing Federated Repositories
- Removing Federated Repositories
- Deleting Backups in Bulk
- Verifying All Backups
- Managing Federated Repository Members
- Migrating Backups between Federated Repository Members

## Deleting Backups in Bulk

To permanently delete several backups that match specific criteria, follow the steps below:

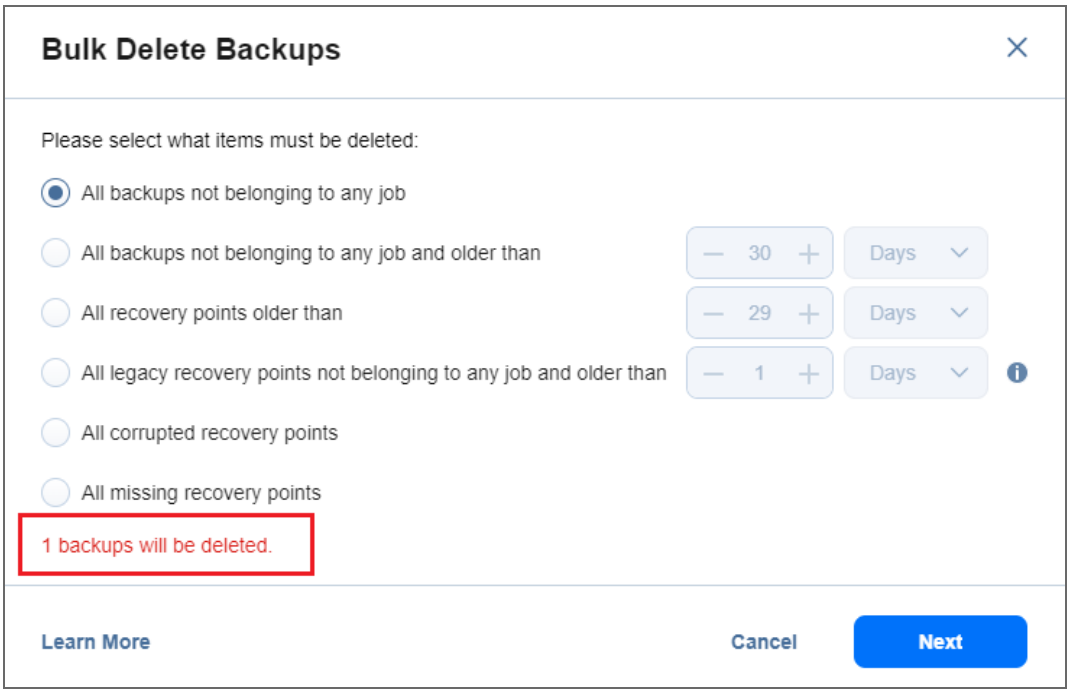
1. In the main menu of NAKIVO Backup & Replication, click **Settings**.
2. Go to the **Repositories** tab.
3. Hover over a **Federated repository**.
4. On the right side, click the ellipsis **Manage** button and then click **Delete backups in bulk**.



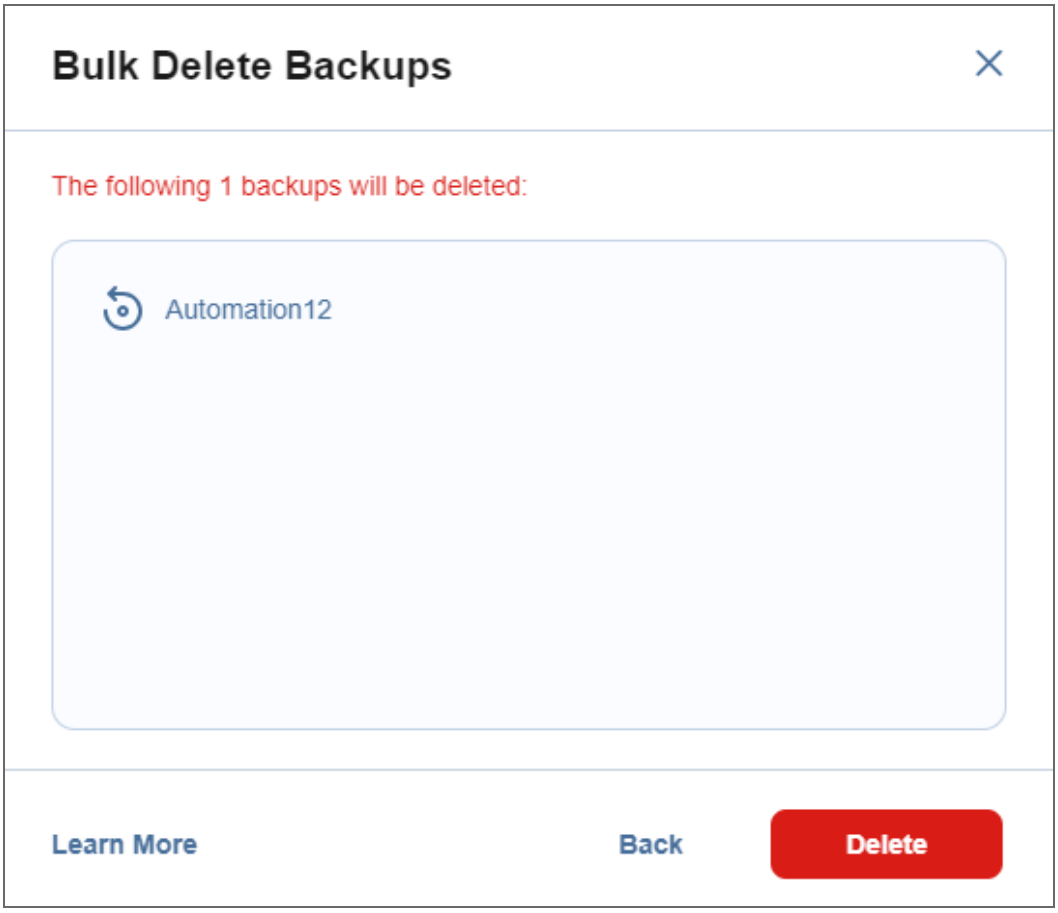

5. In the **Bulk Delete Backups** dialog box that opens, select one of the available options:

- **All backups not belonging to any job**
- **All backups not belonging to any job and older than** X <time_units>, where X is an integer and <time_units> is either days, weeks, or months
- **All recovery points older than** X <time_units>, where X is an integer and <time_units> is either days, weeks, or months
- **All legacy recovery points not belonging to any job and older than** X <time_units>, where X is an integer and <time_units> is either days, weeks, or months
- **All corrupted recovery points**
- **All missing recovery points**

6. The dialog shows the number of backups to be deleted.




7. Click **Next**.
8. The **Bulk Delete Backups** dialog box opens displaying the list of backups to be deleted.
9. Click **Delete** to confirm the deletion of backups.

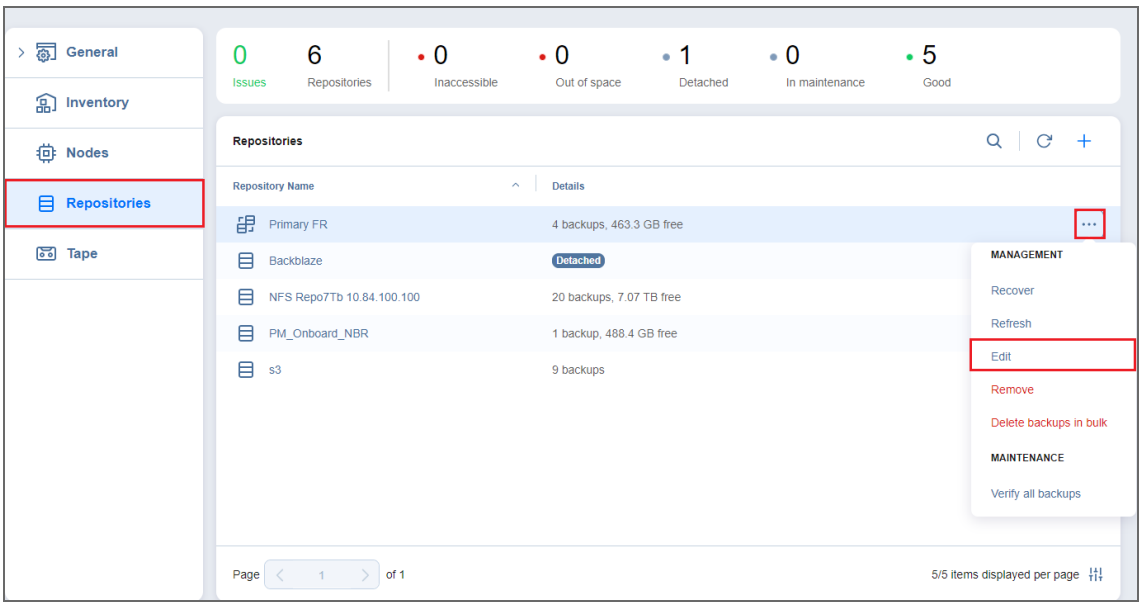


# Editing Federated Repositories

To edit the settings of an existing federated repository, follow the steps below:

In the main menu of NAKIVO Backup & Replication, click **Settings**.

1. Go to the **Repositories** tab.
2. Hover over a **Federated repository**.
3. On the right side, click **Manage**.
4. Click **Edit**.




Update the fields as necessary.

In the **Edit** mode, you can deselect federated repository members, remove them from the federated repository, or add new members.

5. Click **Apply**. Changes you have made are applied and the federated repository update starts.

**Note**

A federated repository cannot be edited while a backup job with this federated repository as target is running.

# Managing Federated Repository Members

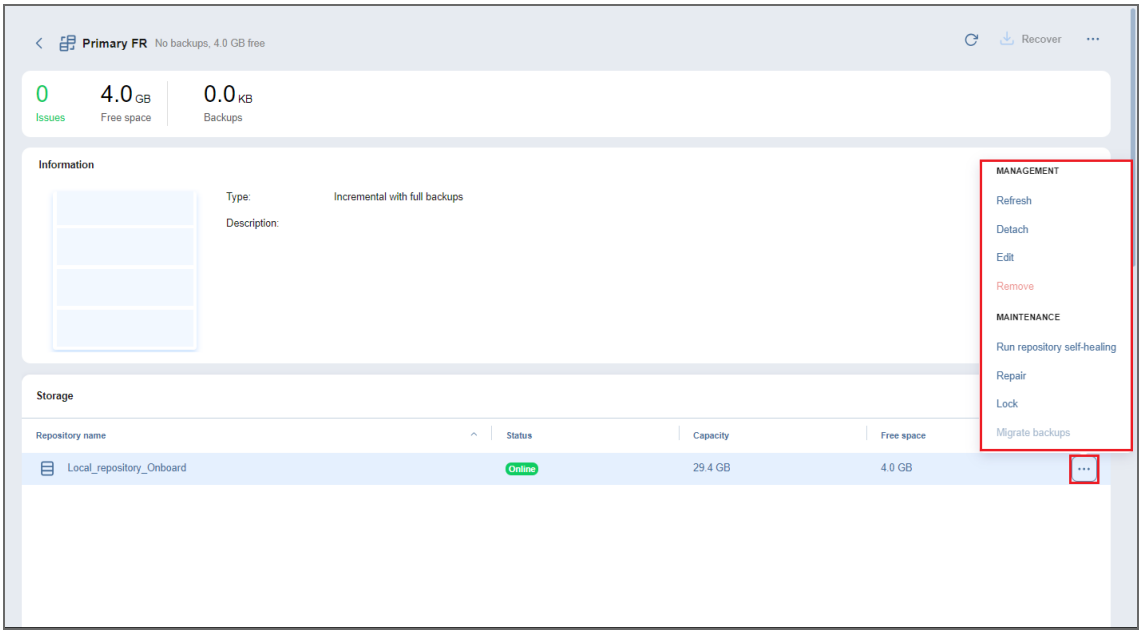
If needed, you can manage the standalone repositories used as members of federated repositories, add new members to, or remove them from the federated repository.

## Managing Members in a Federated Repository

To manage/maintain members of a federated repository, follow the steps below:

1. In the main menu of NAKIVO Backup & Replication, click **Settings**.
2. Go to the **Repositories** tab.

3. Click a **Federated repository** to open the **Storage** table.
4. In the **Storage** table, hover over the needed federated repository member.
5. On the right side, click the ellipsis **Manage** button.




6. Select the needed option from the dropdown.

The following action buttons are available:

- **Refresh**: Opens the **Refresh repository** dialog box
- **Detach/Attach**: Opens the **Detach/Attach** dialog box

> **Note**
> If the federated repository member is currently detached, the **Attach** option is available. Select it to reattach this member.

- **Edit**: Brings up the option to edit the repository member
- **Remove**: Brings up the option to remove the repository member
- **Run repository self-healing**: Starts the repository self-healing process
- **Repair**: Opens the Repair Repository dialog box
- **Lock/Unlock**: Opens the Lock the member dialog box. Select it to suspend saving new recovery points to this repository member

> **Note**
> If the federated repository member is currently locked, the **Unlock** option is available. When you select **Unlock**, the **Unlock the member?** dialog box opens. Select it to resume saving new recovery points to this member.

- **Migrate Backups**: Brings up the option of migrating backups between federated repository members.

**Notes**

- Actions not applicable to a federated repository member are disabled.
- The option to migrate backups is disabled if the member is not locked or if the federated repository runs out of sufficient space for storing migrated data.

## Adding Members to Federated Repository

To add members to a federated repository, follow the steps below:

1. In the main menu of NAKIVO Backup & Replication, click **Settings**.
2. Go to the **Repositories** tab.
3. Hover over a **Federated repository**.
4. On the right side, click **Manage**.
5. Click **Edit**.
6. Add a repository member from the list of available standalone repositories.

**Notes**

- After the changes to the federated repository are saved, the jobs using the repositories and now selected to be part of the federated repository are automatically reconfigured to use this federated repository as a target.
- Closing the federated member edit wizard opens the parent federated repository details screen.

7. Save the changes.

## Deselecting Federated Repository Members

To deselect members in a federated repository, follow the steps below:

1. In the main menu of NAKIVO Backup & Replication, click **Settings**.
2. Go to the **Repositories** tab.
3. Hover over a **Federated repository**.
4. On the right side, click **Manage**.
5. Click **Edit**.
6. In the **Storage** table, deselect a federated repository member.

**Note**

A federated repository member cannot be deselected if it contains at least one backup object referenced by existing jobs.

## Removing Members from Federated Repository

To remove members from a federated repository, follow the steps below:

1. In the main menu of NAKIVO Backup & Replication, click **Settings**.
2. Go to the **Repositories** tab.
3. Hover over a **Federated repository**.
4. On the right side, click **Manage**.
5. Click **Edit**.
6. In the **Storage** table, select a federated repository member.
7. Click **Manage**, then click **Remove**.

**Notes**

- The federated repository member can only be removed at the federated repository level.
- Removing a member that is not referenced by any existing jobs but contains recovery points from a backup object chain distributed across several members may break these backup chain(s). Consider migrating the backup objects before removing this member.

# Migrating Backups between Federated Repository Members

Backups from one federated repository member can be moved to another available qualified federated repository member(s).

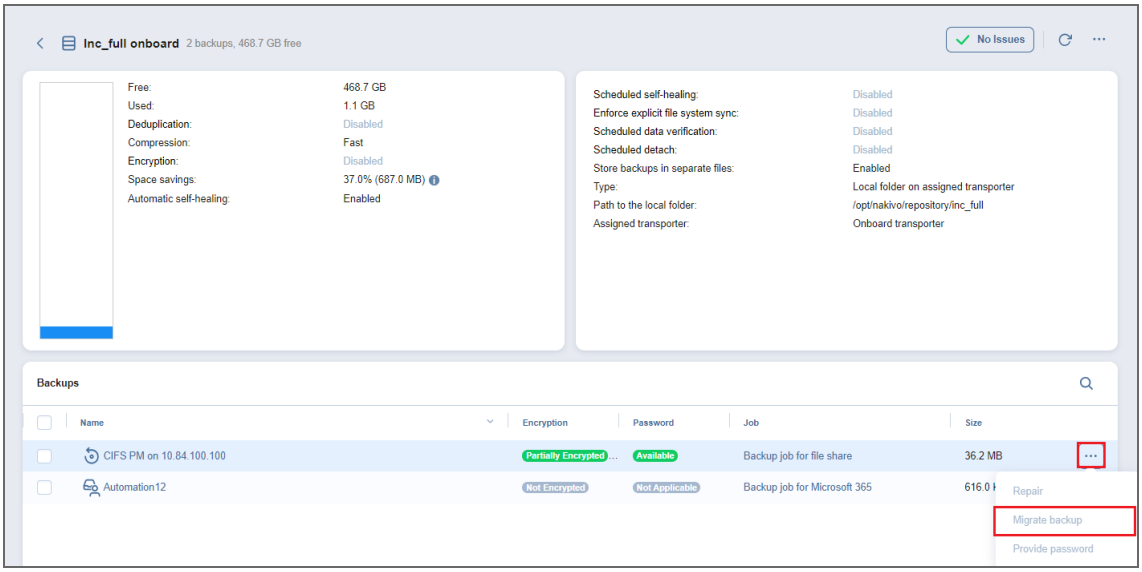
When considering backup migration, you can select one of two options:

- Migration of selected backups
- Migration of all backups

## Migration of Selected Backups

To migrate a selected backup from one federated repository member to another available federated repository member(s), follow the steps below:

1. Go to the main menu of NAKIVO Backup & Replication and click **Settings**.
2. Go to the **Repositories** tab.
3. Select the **Federated repository** to open the Federated repository details dashboard.
4. In the **Storage** table, select the federated repository member name to open the **Backup Repository details** screen.
5. In the **Backups** table, hover over a backup name and click the ellipsis **Manage** button.
6. Click **Migrate backup**.




**Notes**

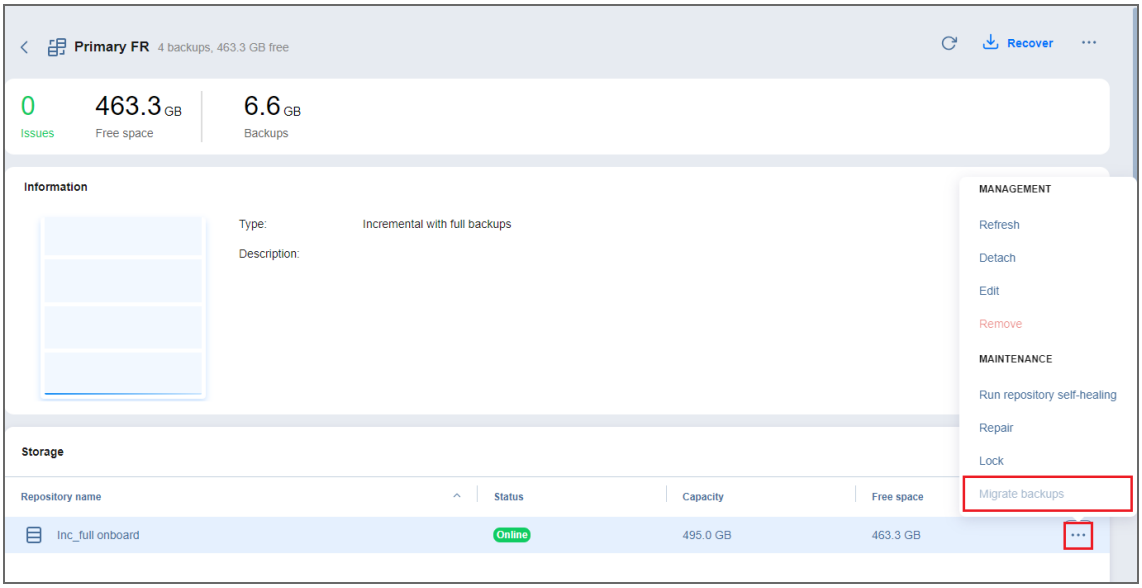
- The option of migrating backups is disabled if the member is not locked or if the federated repository runs out of sufficient space to store migrated data.
- The option is disabled if there are no other members with corresponding backup location policy.
- If a recovery point could not be cleaned up in the course of backup migration and it duplicates another recovery point, it is marked as inaccessible.
- Transporters assigned to both the source and the target federated repository members participate in the backup migration process.

## Migration of All Backups

To migrate all the backups saved on one federated repository member to another federated repository member(s), follow the steps below:

1. Go to the main menu of NAKIVO Backup & Replication and click **Settings**.
2. Go to the **Repositories** tab.
3. Select the **Federated repository** to open the Federated repository details dashboard.
4. In the **Storage** table, hover over a federated repository member name.
5. On the right side, click the ellipsis **Manage** button and then click **Lock the member**. Now, saving new recovery points to this repository member is suspended.
6. Click **Migrate backups**.

The process of migrating backups between federated repository members begins.




**Notes**

- The progress based on the size of migrated backups can be monitored in the Action Bar.
- In the course of backup migration, the product redistributes backups stored on the selected federated repository member to available qualified federated repository members. Then the product removes the successfully redistributed backups from the selected source federated repository member. However, the immutable backups are not removed.
- The following functionality is enabled by default (if applicable) in the course of backup migration:
  - Network acceleration
  - Network encryption
- If the cleanup of some recovery points/backup objects fail (Clean up is skipped for immutable backups.), these remaining recovery points are not considered part of the backup objects associated with existing jobs.
- Duplicates of existing recovery points are marked as inaccessible, and the corresponding notification is displayed.
- If backup migration fails, the product re-tries the operation.
- If migration has been stopped and some backups have not been migrated yet, they remain in the source member.
- In case of a backup migration failure, information is logged and a dialog with the alarms list is displayed in the Activities dashboard.

# Refreshing Federated Repositories

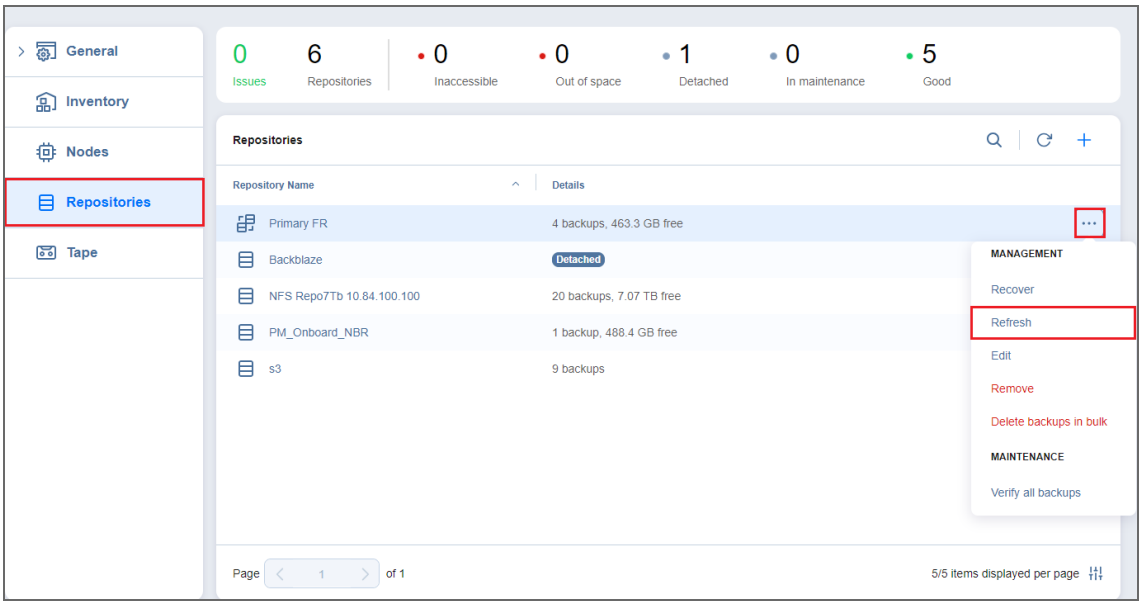
By default, NAKIVO Backup & Replication refreshes federated repository information hourly. During the refreshing process, the product collects all required information about federated repositories, such as the amount of available free space, number of backups, and number of recovery points.

Only one federated repository is refreshed at a time. Therefore, if you attempt to refresh multiple federated repositories, while one is being refreshed, all other federated repositories are queued.

To refresh a federated repository, follow the steps below:

In the main menu of NAKIVO Backup & Replication, click **Settings**.

1. Go to the **Repositories** tab.
2. Hover over a **Federated repository**.
3. On the right side, click **Manage**.
4. Click **Refresh**.




The federated repository refresh process begins.

# Removing Federated Repositories

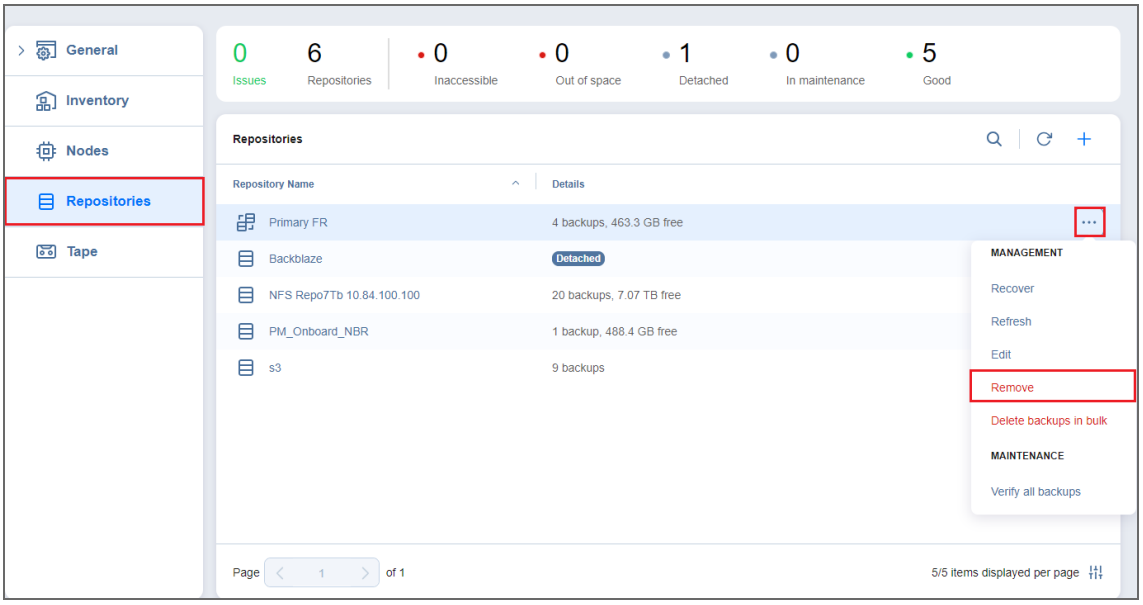
When you remove a federated repository, federated repository members and their contents are not removed and are treated as standalone backup repositories. However, the backup chains may be broken.

**Note**

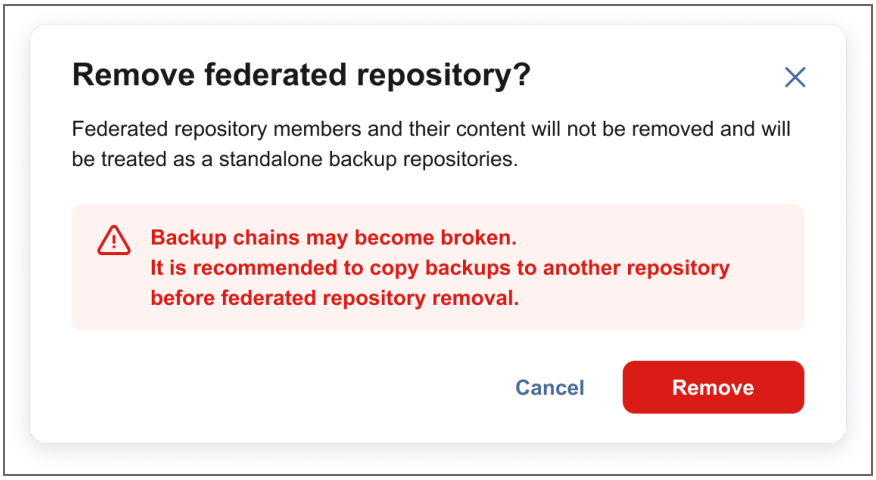
A federated repository cannot be removed while a backup job using this federated repository as a target is running or if it contains at least one backup object referenced by existing job(s). To remove such a federated repository, delete (or edit) the corresponding job so that no items are backed up to the aforementioned repository.

To remove a federated repository from the product, follow the steps below:

1. In the main menu of NAKIVO Backup & Replication, click **Settings**.
2. Go to the **Repositories** tab.
3. Hover over a **Federated repository**.
4. On the right side, click **Manage**.



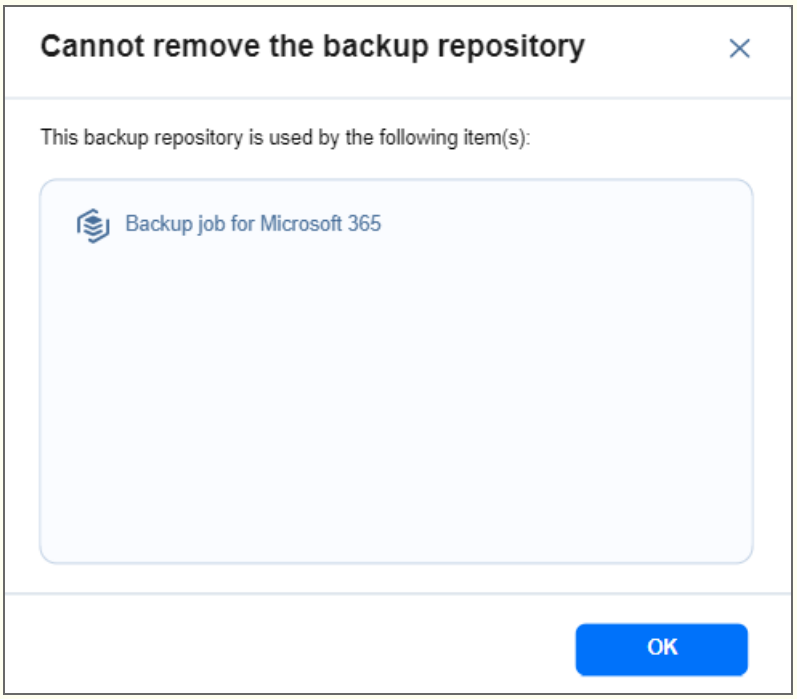

5. Click **Remove**.
6. Click **Remove** to confirm the action when the following confirmation message appears:




Make sure you copy the backups it contains to another repository.

**Note**

If a federated repository is used by one or more objects/processes or contains at least one backup object referenced by existing job(s), the following dialog box is displayed:




**Important**

You cannot remove the federated repository if it is used by one or more objects/processes.

## Verifying All Backups

Backup verification is a process of checking the integrity of a backup by verifying its recovery points.

The **Federated repository** feature introduces enhanced backup integrity verification:

- Performs comprehensive check of all recovery points within the repository
- Validates metadata completeness
- Ensures data integrity by comparing with checksums

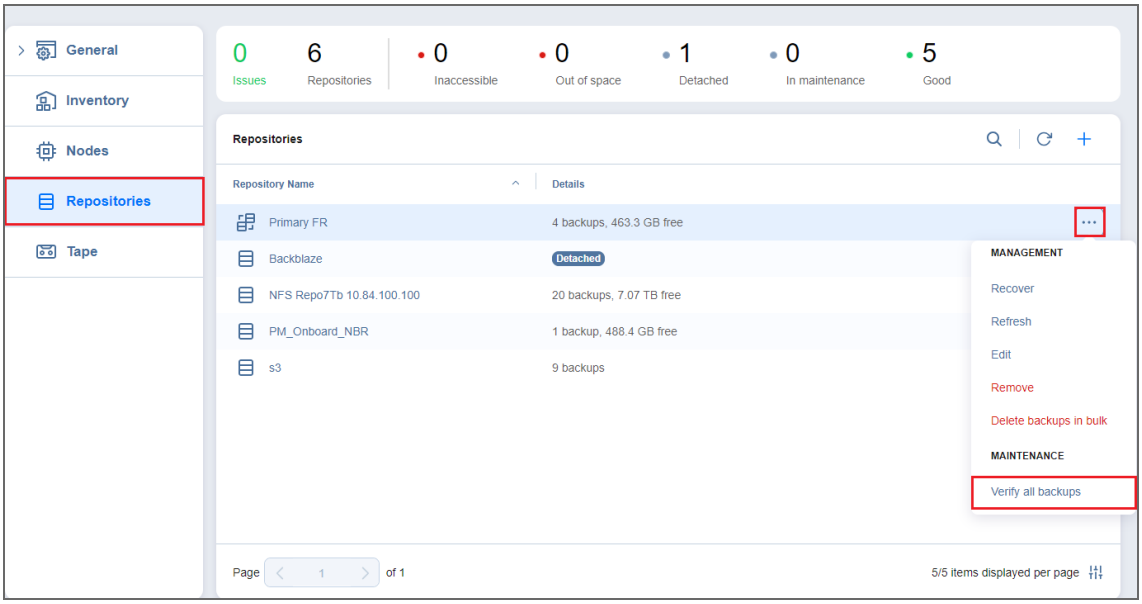
For specific recovery points, the following verification tasks will be done:

- Read timeline file:
  - Retrieves recovery point metadata from the timeline file
- Metadata verification:
  - Validates the presence and currency of metadata files specified in the timeline file
  - Compares timestamps for accuracy
- Incremental verification:
  - Searches for the parent FULL recovery point if the current recovery point is incremental
  - Validates the presence and correctness of parent recovery points, raw files, and metadata

- Raw file and metadata presence check:
  - Ensures the existence and correctness of raw files and their backup metadata (headers and records).

To start backup verification, follow the steps below:

1. In the main menu of NAKIVO Backup & Replication, click **Settings**.
2. Go to the **Repositories** tab.
3. Hover over a **Federated repository**.
4. On the right side, click the ellipsis **Manage** button and then click **Verify all backups**.




> **Note**
>
> The backup verification process cannot be started if a backup job using this Federated repository as a target is running.

5. To stop verification, hover over a federated repository, click the ellipsis **Manage** button and then click **Stop backup verification**.

# Viewing Federated Repository

To see information about specific federated backup repositories, their members, backups, and recovery points, check out the sections below.

Viewing Federated Repository Details

Viewing Federated Repository Member Details

Viewing Backup Details

To manage federated repositories, in the **Repositories** table, hover over the name of a federated repository and click the ellipsis **Manage** button on the right side. For more detailed information, refer to "Managing Federated Repositories" on page 700.

## Viewing Backup Details

To view the federated repository backup details, follow the steps below:

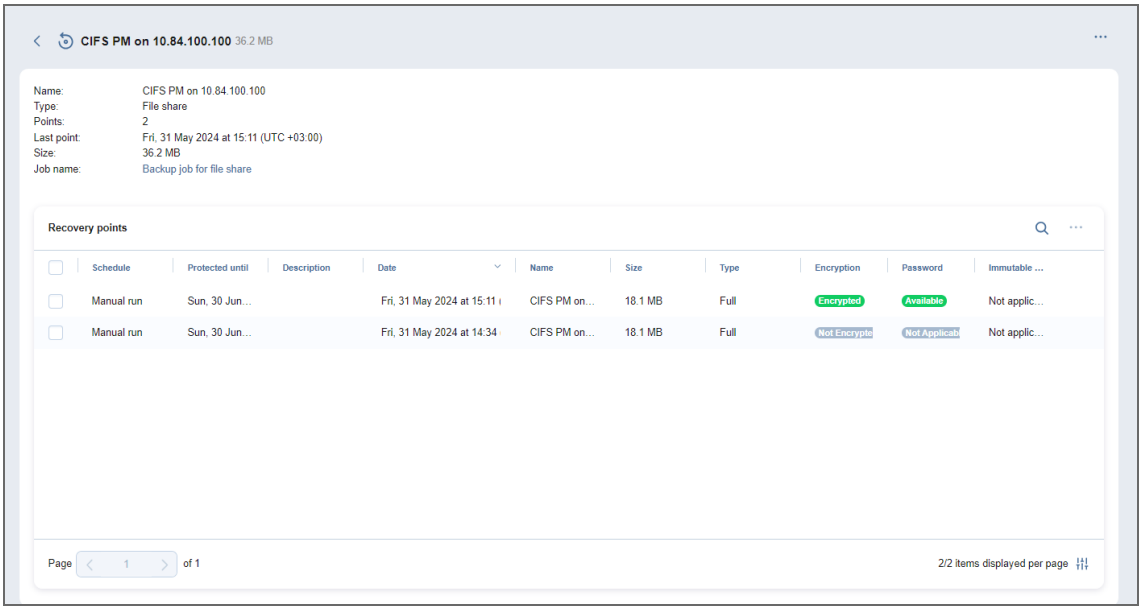
1. Go to the main menu of NAKIVO Backup & Replication and click **Settings**.
2. Go to the **Repositories** tab.
3. Select the **Federated repository**  to open the Federated repository details dashboard.
4. Scroll down to the **Backups** table and click the backup name to open the backup details screen in a new view.

The following information is displayed:

- **Name**: The name of the backup item
- **Type**: The type of job
- **Points**: The number of recovery points available
- **Last point**: The date of the latest recovery point
- **Size**: The total size of the backup
- **Job name**: The name of the job

**Notes**

- Backup details only contain the list of recovery points available in the selected backup object.
- Large numbers of backups are separated into pages to reduce clutter. To find a specific backup, you can scroll through the pages manually or simply look it up using the **Search** bar.




You can view the details of a recovery point in the lower part of the screen. To find a recovery point for a specific date, you can use the **Search** bar on the right. The following information is displayed:

- **Date**: The date when the recovery point was created
- **Size**: The size of the recovery point
- **Type**: Type of backup used to create the recovery point
- **Schedule**: If applicable, the schedule that was used to create the recovery point
- **Encryption**: The state of encryption of the backup object, which can be one of the following:
  - **Not encrypted**: None of the recovery points of the backup object are encrypted.
  - **Encrypted**: All the recovery points of the backup object are encrypted.
- **Password**: The state of the backup object password, which can be one of the following:
  - **Available**:The recovery point is encrypted and the corresponding password hash required for its decryption is available.
  - **Not available**: The recovery point is encrypted, and the corresponding password hash required for its decryption is not available.
  - **Not applicable**: The recovery point is not encrypted.
- **Immutable until**: If applicable, the date when the recovery point immutability expires
- **Protected until**: The date until which the recovery point is retained, displayed only for recovery points belonging to jobs that use the schedule retention approach
- **Description**: The description of the recovery point if one was provided

**Notes**

- For recovery points belonging to jobs using legacy retention settings, **Use job retention** is displayed under **Protected until** instead.
- **Date**, **Type**, and **Description** can also be viewed when selecting recovery points in the **Recovery Job Wizard**.

Optionally, hover over the name of the recovery point and click the ellipsis **Manage** button on the right side to select one of the following options:

- **Recover**: Select this option to proceed with recovery.
- **Edit**: Select this option to edit the recovery point. Do the following:
  - Optionally, you can add a **Description** to your recovery point.
  - Choose the date until which the recovery point should be kept. The following options are available:
    - **Use job retention**: Choose this option to use the retention settings configured in the job for this recovery point.
    - **Keep forever**: Choose this option to keep this recovery point forever.
    - **Protect until**: Choose this option to keep this recovery point until a specific date. After selecting this option, choose the date in the calendar pop-up.
- **Delete**: Select this option to delete the recovery point from the repository. Refer to "Deleting Recovery Points" on page 928 for more details.

## Viewing Federated Repository Details

To view the federated repository details, follow the steps below:

1. Go to the main menu of NAKIVO Backup & Replication and click **Settings**.
2. Open the **Repositories** tab.

The **Repositories** tab contains a **Summary** bar, which offers an overview of all backup repositories both federated and standalone (not used as members of federated repositories). The data displayed is as follows:

- **Issues**: Total number of issues/alarms related to repositories
- **Repositories**: Total number of repositories
- **Inaccessible**: Total number of inaccessible repositories
- **Out of Space**: Total number of repositories that are out of storage space
- **Detached**: Total number of detached repositories
- **In Maintenance**: Total number of repositories in maintenance
- **Good**: Total number of usable repositories

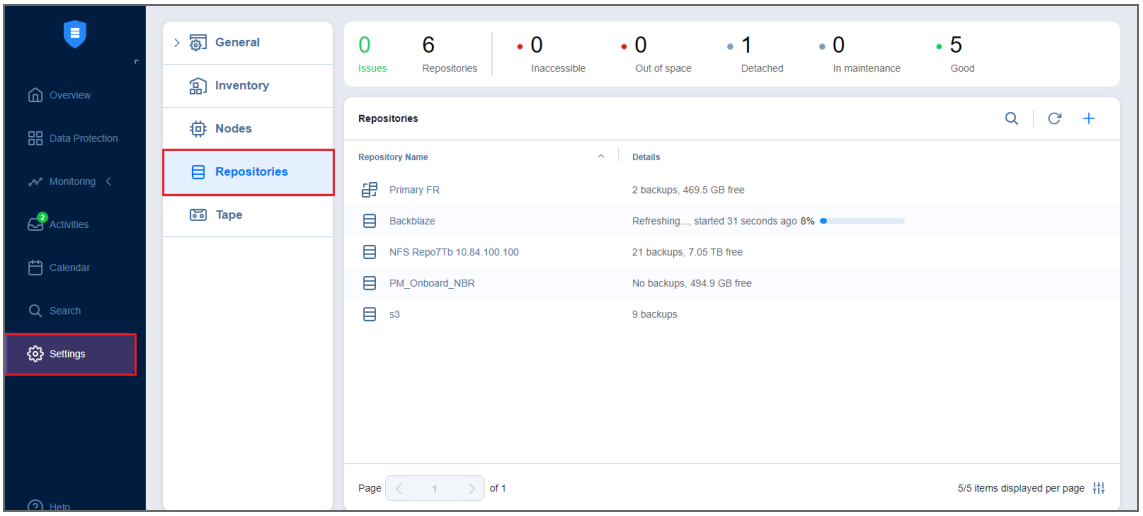
**Note**

In the **Repositories** table below, all repositories added to the product are listed according to the following priorities:

- Federated repositories
- Standalone repositories (not used as members of federated repositories)

The data in the **Repositories** table is displayed as follows:

- **Repository name**: The name of the selected repository and its icon
- **Details**: The number of repository backups, amount of free space, and repository status (displayed for detached and inaccessible repositories only)




**Notes**

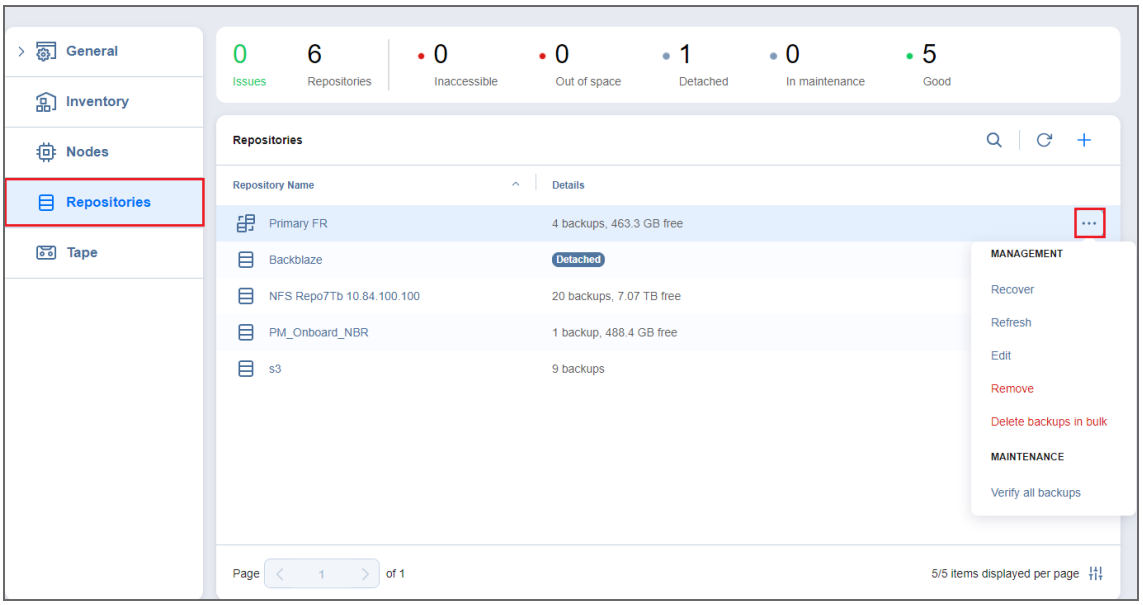
- If a repository status is displayed, it replaces the **Number of backups** and **Amount of free space data**.
- If a repository is both **Inaccessible** and **Detached**, the **Detached** status is displayed.
- If a repository is both **Out of space** and **Detached**, the **Detached** status is displayed.
- If some federated repository members are detached, other members are inaccessible, and there are no accessible attached members, the **Inaccessible** status is displayed for the federated repository.
- If any of the *in progress* statuses is in action (for example, *Self-healing*, *Verifying backups*), the progress bar is displayed.
- The progress bar is not displayed for **detached** and **inaccessible** repositories.

3. In the **Repositories** table, select the federated repository  to open the Federated repository details dashboard.

**Notes**

- Standalone repositories used as federated repository members are not displayed in the **Repositories** table. To view them, in the **Federated repository details** dashboard, proceed to the **Storage** table and click the federated repository member name to open the **Backup Repository details** screen.
- Optionally, you can search for standalone repositories by clicking the **Search** button. This displays the search field where you can enter the partial or full name of the repository.
- Click the **Refresh** button to refresh the list of items.
- Click the **+** button to create a new or add existing repository to your instance.
- To sort the table by **Repository name** click the head of the corresponding column.

Optionally, in the **Repositories** table, hover over the name of a federated repository and click the ellipsis **Manage** button on the right side to select one of the available options.




For more detailed information, refer to Managing Federated Repositories.

## Federated Repository Details Dashboard

The **Federated repository details** dashboard offers a detailed overview of the selected federated repository details.

For a detailed explanation of each component in the **Federated repository details** dashboard, see the sections below.


- "Action Bar" on the next page
- "Summary Bar" on the next page
- "Information Pane" on the next page
- "Storage Table" on page 719

- "Backups Table" on the next page
- "Events Table" on page 720

Action Bar

The **Federated repository** action bar contains the generic icon, the title of the federated repository, total amount of free space in the federated repository, and the following three action buttons:

**Refresh**: Opens the Refresh repository dialog box

**Recover**: Brings up a list of recovery options for the selected federated repository

**Manage**: Brings up the options to edit, remove, delete backups in bulk, or verify all backups.

For more information, see Managing Federated Repositories.

Summary Bar

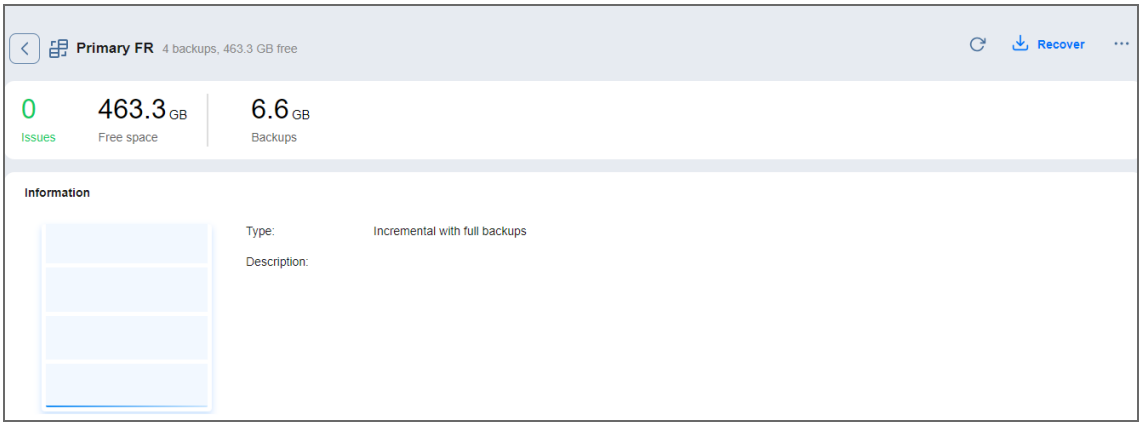
The **Federated backup repository** tab contains a **Summary** bar, which offers an overview of the selected federated repository. The data displayed is as follows:

- **Issues**: Total number of issues associated with the selected federated repository.
- **Free space**: Total amount of free space in the federated repository in GB (Gigabytes).
- **Backups**: Total amount of used space for backups in the federated repository in GB (Gigabytes).

Information Pane

This pane displays the following information about the selected repository:

- **Bar chart**: represents the amount of free/used space.
- **Type**
- **Description** (optional)




**Important**

The **Federated repository** feature supports backing up data to **Incremental with full backups** repositories only.

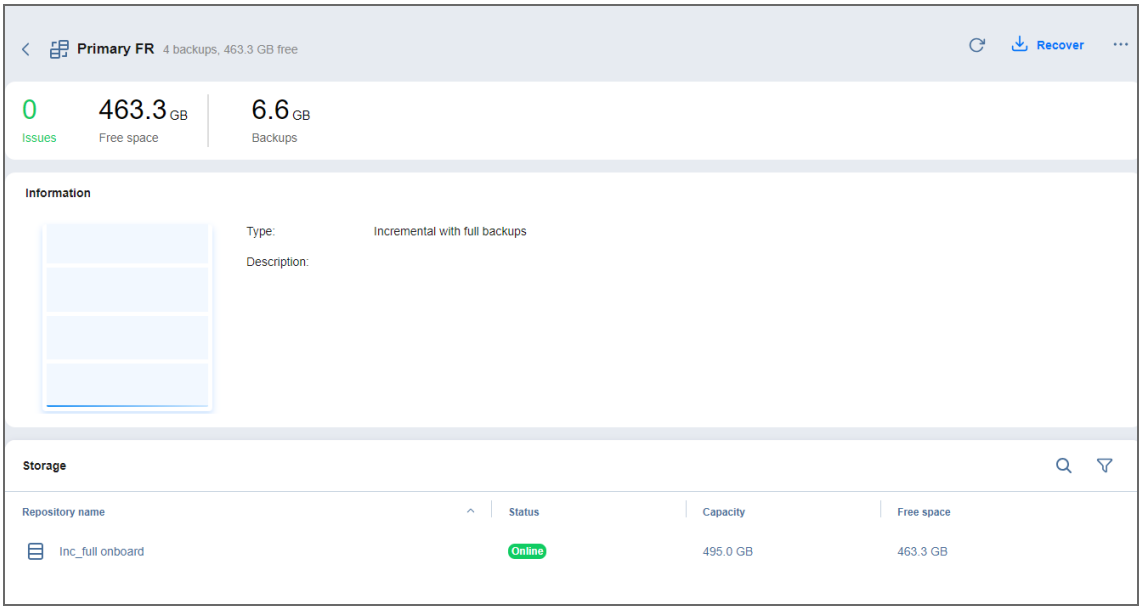
### Storage Table

The **Storage** table displays a list of federated repository members and the information about each of them in the following columns:

- **Repository name**: The name of the repository.
- **Status**: The repository status, such as *Online* (displayed when the repository is accessible, attached, consistent and not under maintenance), *Detached*, *Locked* (replaced by *Self-healing…* or *Verifying backups…* if those are the cases), *Inaccessible*, or *Inconsistent*).
- **Capacity**: The repository capacity in GB (Gigabytes).
- **Free space**: The repository free space in GB (Gigabytes).

**Notes**

- Optionally, you can search for members by clicking the **Search** button. This displays the search field where you can enter the partial or full name of the member.
- Click on the **Filter** button to filter the storage tree items by **Name**, **Status** (*Online*, *Detached*, *Locked*, *Inaccessible*, or *Inconsistent*) .




To manage federated repository members, in the **Storage** table, hover over the name of a federated repository member and click the ellipsis **Manage** button on the right side. For more detailed information about the available options, refer to Managing Federated Repository Members.

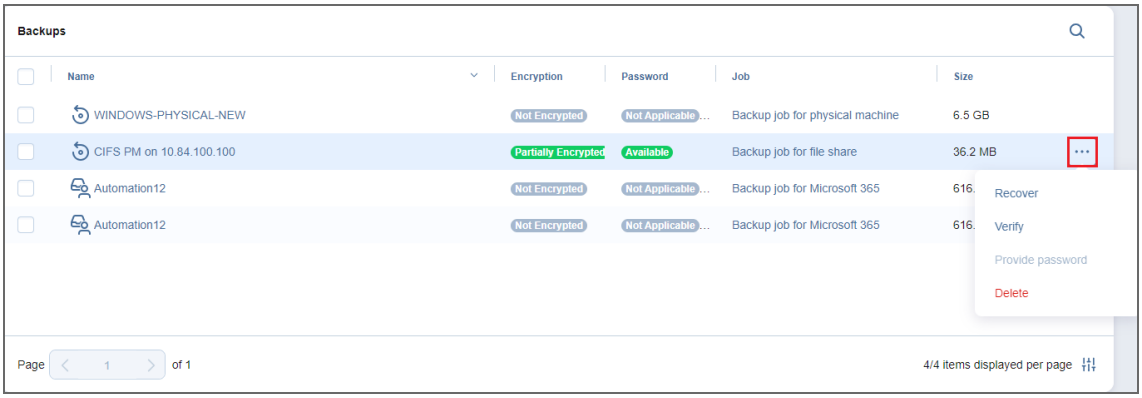
Click the federated repository member name to open the **Backup repository details** screen.

For more detailed information, refer to Viewing Federated Repository Details.

### Backups Table

You can view the details of the backups stored in the federated repository. The following information is displayed in the **Backups** table:

- **Backup name**: Name of the backup
- **Encryption**: The state of encryption of the backup object, which can be one of the following:
  - **Not encrypted**: None of the recovery points of the backup object are encrypted.
  - **Encrypted**: All recovery points of the backup object are encrypted.
  - **Partially encrypted**: Some recovery points of the backup object are encrypted.
- **Password**: The state of the password of the backup object, which can be one of the following:
  - **Available**: All or some recovery points of the backup object are encrypted, and the password hash is available for all encrypted recovery points.
  - **Not available**: All or some recovery points of the backup object are encrypted, and the password hash is not available for all encrypted recovery points.
  - **Partially available**: All or some recovery points of the backup object are encrypted, and the password hash is not available for some encrypted recovery points.
  - **Not applicable**: The backup object does not contain encrypted recovery points.
- **Job**: The job type that created this backup
- **Size**: The total size of the backup




**Notes**

- You can use backups stored in a federated repository to create replicas.
- Recovery point management is only available on the federated repository level.

Click the backup name to open the backup details screen in a new view.

**Note**

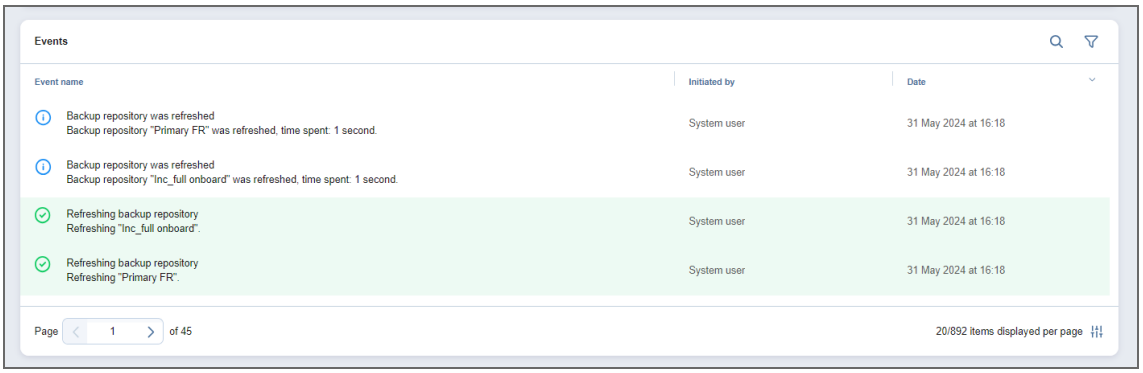
You can also migrate all the backup(s) from one federated repository member to the other available federated repository member(s). Refer to Migration of All Backups to learn more.

Events Table

The **Events** table displays a list of events associated with this federated repository in the following columns:

- **Event name**: The name of a given event in the group
- **Initiated by**: The initiator of the event

- **Date**: The date and time of the event



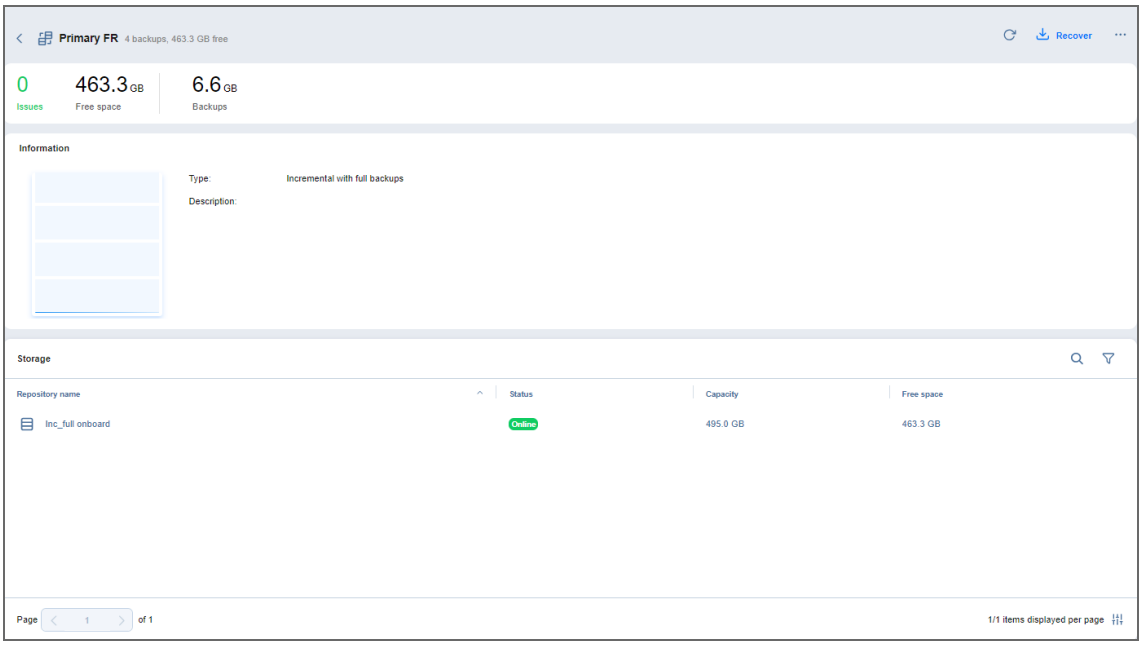

# Viewing Federated Repository Member Details

**Important**

Standalone repositories used as federated repository members are not displayed in the **Repositories** table.

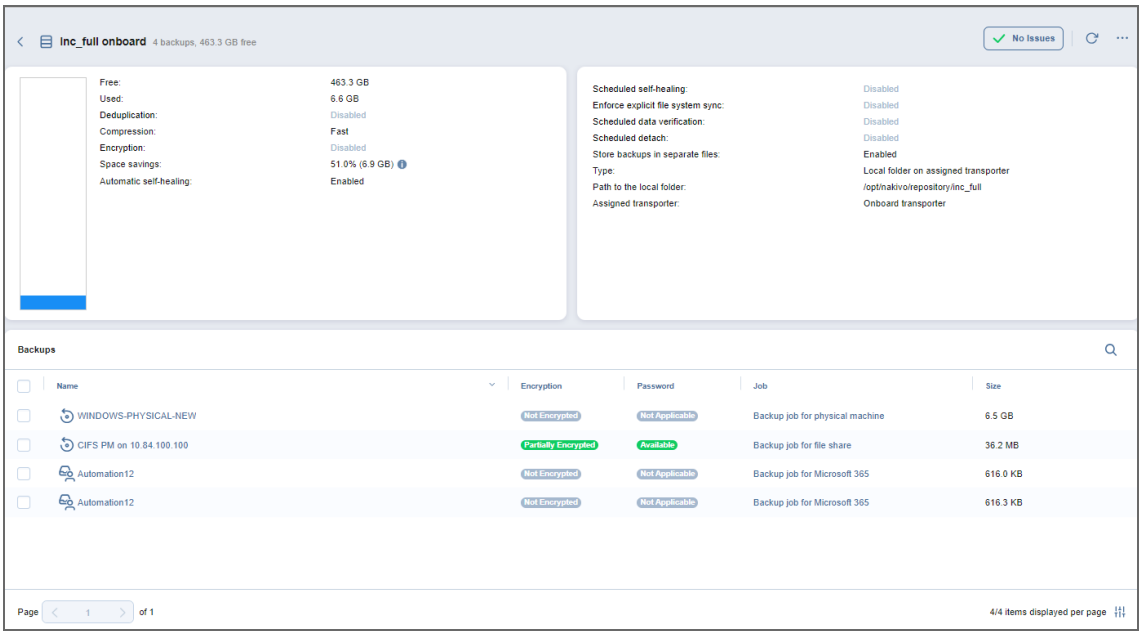
To view the federated repository member details, follow the steps below:

1. Go to the main menu of NAKIVO Backup & Replication and click **Settings**.
2. Go to the **Repositories** tab.
3. Select the **Federated repository** to open the **Federated repository details** dashboard.

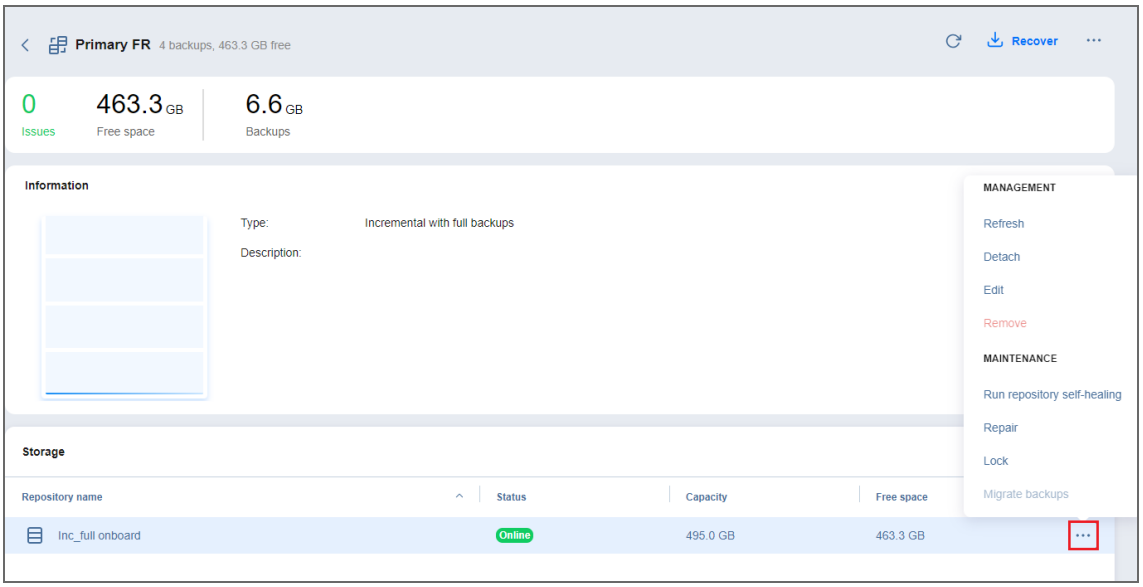



4. In the **Storage** table, click the federated repository member name to open the dashboard with the

federated repository member details.




To manage federated repository members, in the **Storage** table, hover over the name of a federated repository member and click the ellipsis **Manage** button on the right side.



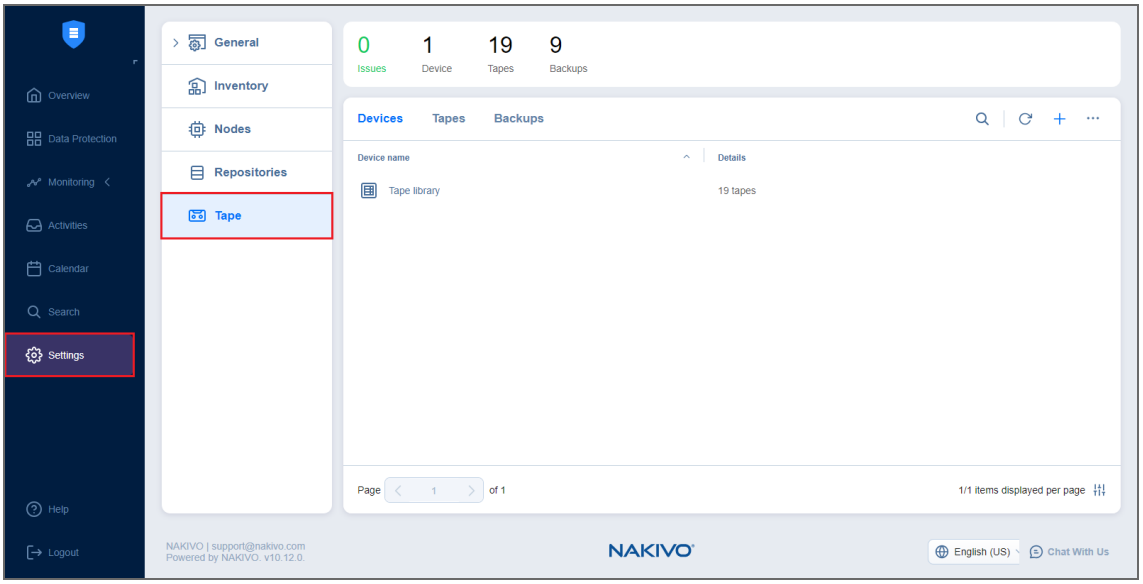

For more detailed information, refer to Managing Members in a Federated Repository.

**Notes**

- Backup details only contain the list of recovery points available in the selected member.
- No actions are available at this level.

# Tape

To start working with tape devices in NAKIVO Backup & Replication, you should first add and configure these devices on the **Tape** page of the **Settings** dashboard.




The **Summary** bar displays information about all tapes. The data displayed is as follows:

- **Issues**: Total number of issues/alarms related to tapes
- **Devices**: Total number of tape devices
- **Tapes**: Total number of tapes
- **Backups**: Total number of tape backups

The default view of the **Tape** page is set to the **Devices** tab. Once you add your tape devices, you can view and manage them on this page. To work with tapes and backups, click the **Tapes** and **Backups** tabs, respectively.

On the **Tape** page, you can perform the following operations:

- “Adding Robotic Tape Libraries or VTLs” on page 724
- “Adding Standalone Tape Drives” on page 731
- “Managing Backups” on page 735
- “Managing Locations” on page 743
- “Managing Media Pools” on page 746
- “Managing Tape Cartridges” on page 749
- “Managing Tape Devices” on page 760

# Adding Robotic Tape Libraries or VTLs

Make sure to observe the following prerequisites before adding Robotic Tape Libraries or Virtual Tape Libraries to **Inventory**:

- Vendor drivers should be installed on tape devices prior to adding them to NAKIVO Backup & Replication **Inventory**.
- To be able to work with AWS VTL and StorSafe VTL, you need to deploy a **Transporter** and manually mount VTL targets.

The process of adding a Robotic Tape Library or Virtual Tape Library to NAKIVO Backup & Replication includes the following steps:

- Launching Wizard
- Selecting Transporter
- Selecting Changers
- Selecting Drives
- Selecting Options
- Managing Added Tape Library

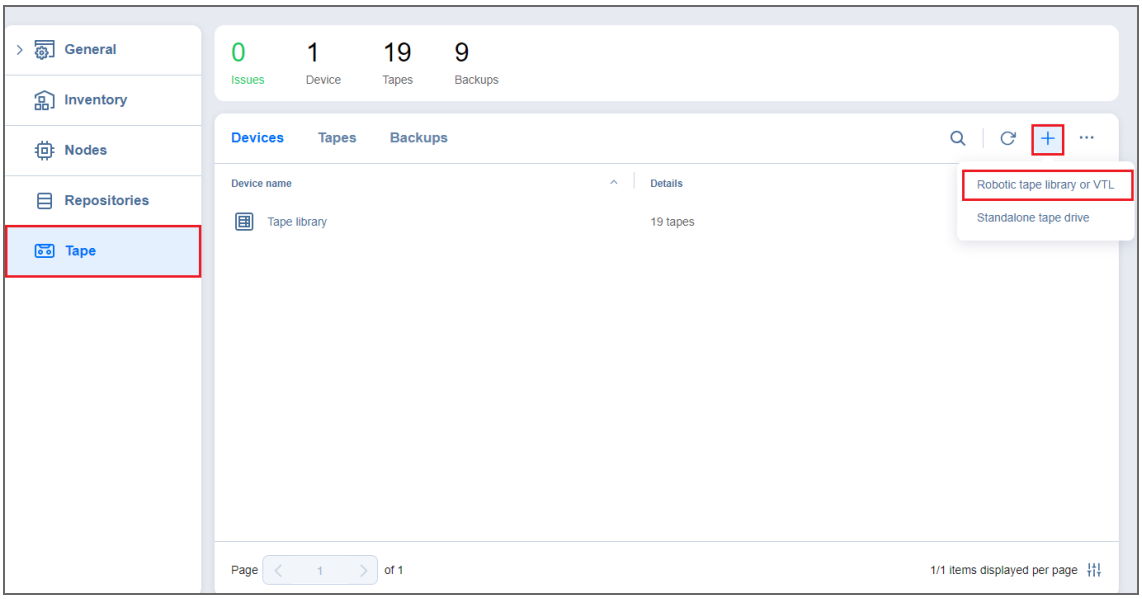
## Launching Wizard

Before adding a Robotic tape library or Virtual Tape Library, make sure that the on-premises VM or Amazon EC2 instance meets the necessary feature requirements.

To add a Robotic tape library or VTL to the system:

1. Go to **Settings** and click the **Tape** tab.
2. Go to the **Devices** tab.

3. Click the plus **Add** button and select **Robotic tape library or VTL**.




The **Add New Robotic Tape Library or Virtual Tape Library** wizard opens. Follow the steps below to add a new device.

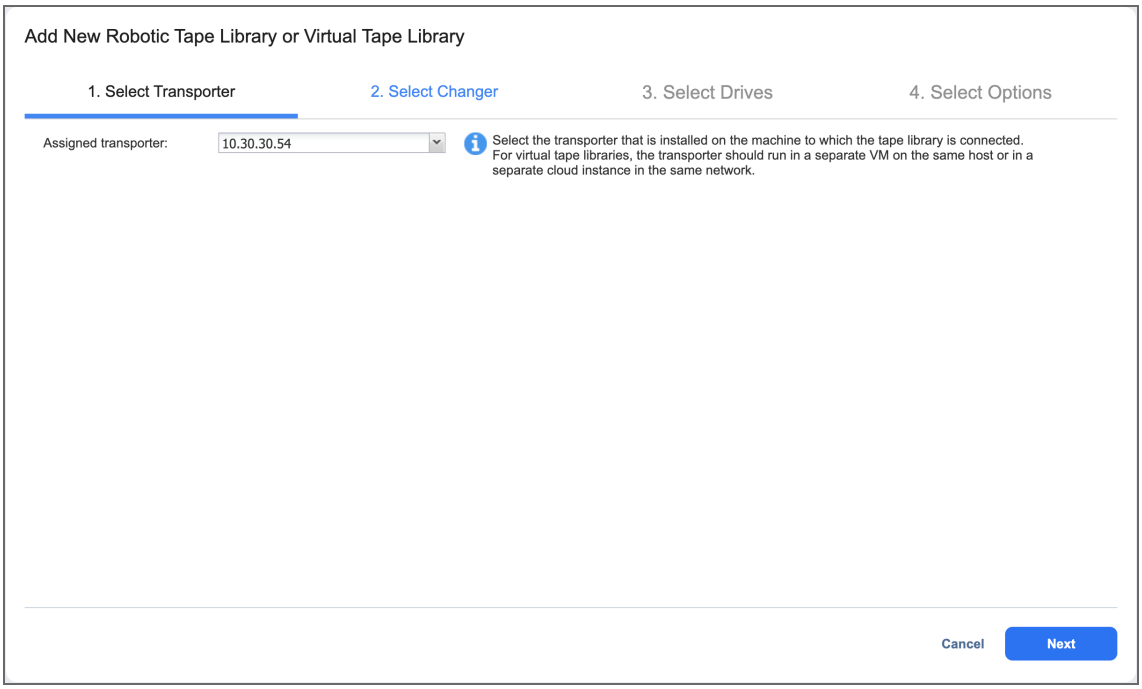
**Note**

Before adding a new tape device to NAKIVO Backup & Replication, you need to deploy or add an existing **Transporter** on a machine that is physically connected to the tape device. For virtual tape libraries, the transporter should run in a separate VM on the same host or in a separate cloud instance in the same network. For more information on **Transporter** deployment, refer to "Deploying Transporter as VMware Appliance" on page 575 and "Adding Existing Nodes" on page 565.

## Selecting Transporter

On the **Select Transporter** step, you need to specify a **Transporter** assigned to the device or VTL you want to add. This Transporter acts as a network appliance that manages traffic between the tape device and NAKIVO Backup & Replication.

1. From the **Assigned Transporter** drop-down list, select the relevant installed **Transporter**. You can also select a physical machine agent installed on the host connected to the tape device.



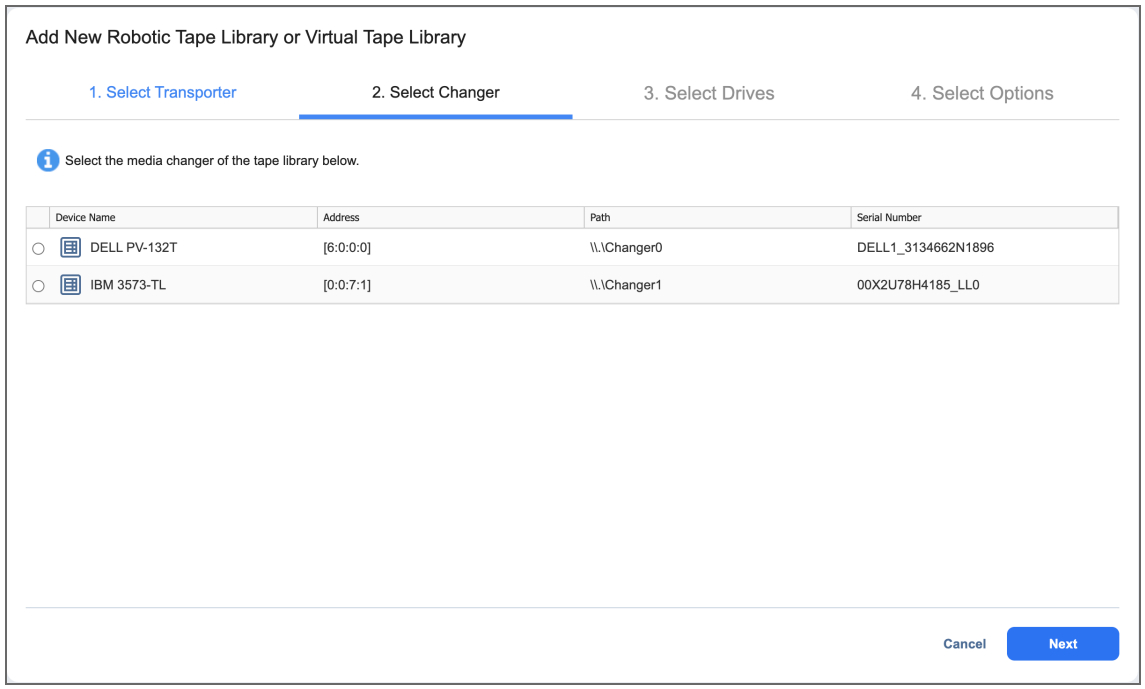

2. Click **Next**.

# Selecting Changers

The **Select Changers** page displays the list of media changers on the selected **Transporter**.

> **Note**
> If no media changers were found on the specified transporter, make sure the devices are connected, powered on, and the appropriate drivers are installed.

Select one media charger from the list. Media changers already being used in another discovered tape library are disabled.




The following information is displayed for each media changer to facilitate the selection:

- **Device name**: Indicates device's vendor and model, separated by space
- **Address**: Indicates the hardware address including the bus and node numbers
- **Path**: Indicates location in the operating system
- **Serial number**: Indicates the serial number of the device

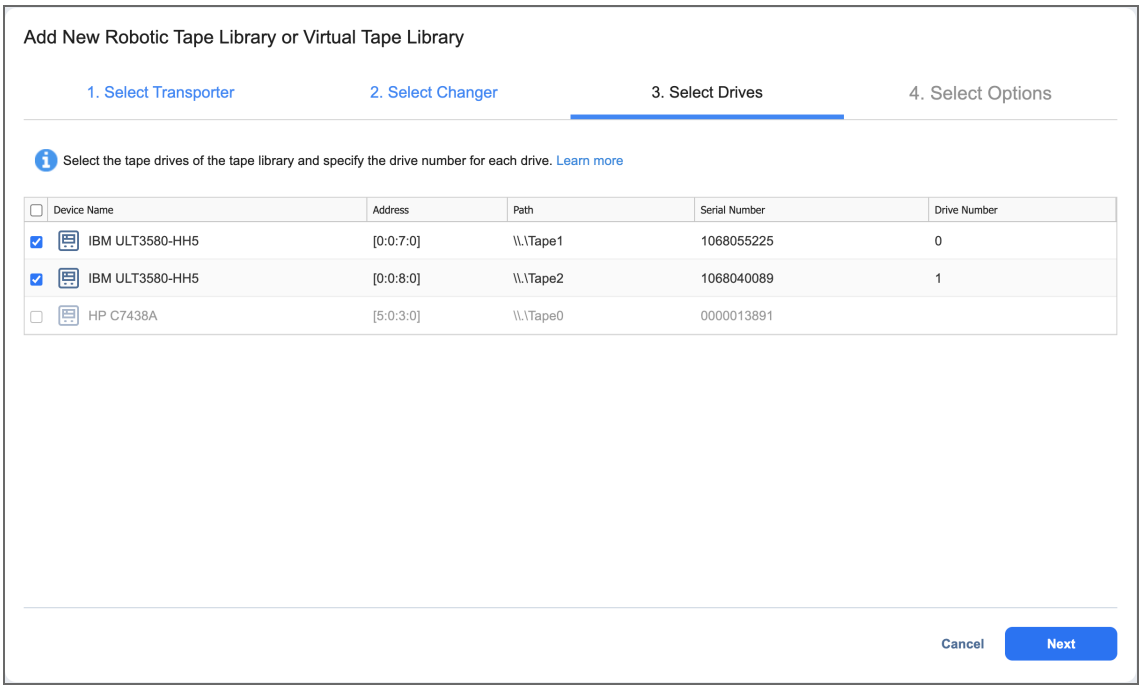
## Selecting Drives

On the **Select Drives** page, you can select tape drives from the tape library and specify the actual drive number for each drive. Drives already used in another discovered tape device are disabled and cannot be selected. The table provides the following information:

- **Device name**: Indicates the device's vendor and model.
- **Address**: Displays the hardware address including the bus and node numbers.
- **Path**: Shows the location in the operating system.
- **Serial number**: Shows the serial number of the drive.
- **Drive Number**: Indicates the drive number and allows changing it.  Changing the drive number may be required to address situations, where iSCSI targets are assigned incorrectly to the mounted drives.

**Note**

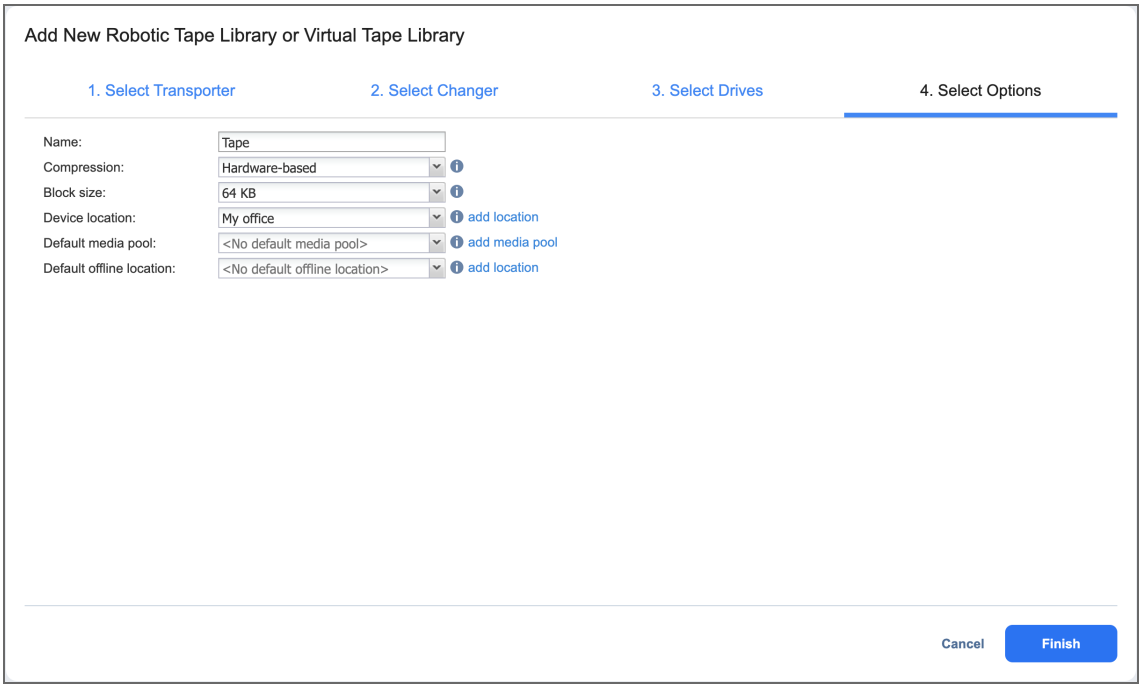
If more than one drive is selected, such drives should use the same host/buses.




Click **Next** to proceed to the next page.

# Selecting Options

The last step of adding a tape library is selecting its options.



1. Specify the following parameters:

- **Name**: Enter the name for the tape library
- **Compression**: Select a compression level of the tape device:
  - Software-based (fast) (default)
  - Hardware-based
  - Software-based (medium)
  - Software-based (best)

> **Note**
> Setting hardware compression is not recommended to avoid the issue of increased data size during transfer and long backup/recovery times. Having different types of compression for the tape device and a source/target **Backup Repository** during backup or recovery can also lead to this issue. For more information, see this article.

- **Block size**: Select the block size of the tape device:
  - 32 KB
  - 64 KB (default)
  - 128 KB
  - 256 KBs
  - 512 KB
  - 1 MB

> **Note**
> The system does not automatically detect the block size; make sure to use the correct block size when importing backups.

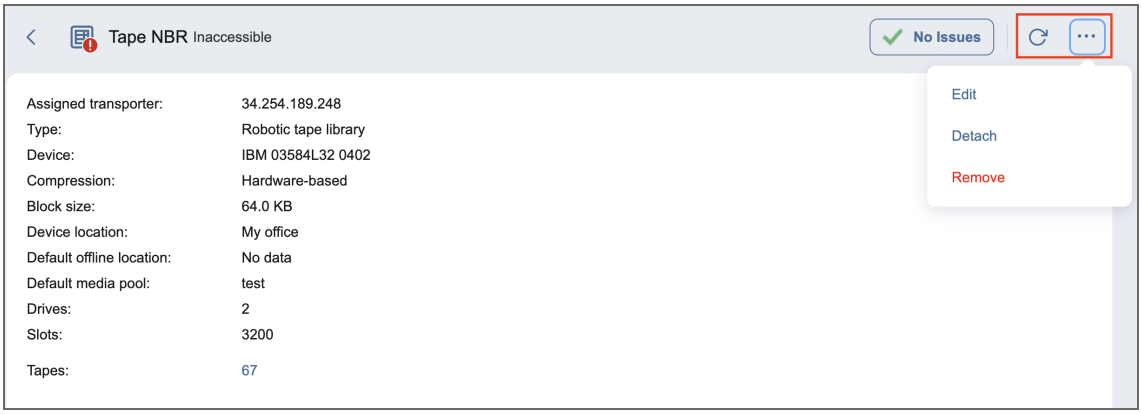
- **Device location**: Select the location of the device and all tapes inserted into this device. The automatically created **My office** location is selected by default. To create another location, click **Add Location**. For more information on locations management, refer to "Managing Locations" on page 743.
- **Default media pool**: Select a default media pool for all new tapes inserted into this device. To create another media pool, click **Add Media Pool**. For more information on media pools management, refer to "Managing Media Pools" on page 746.
- **Default offline location**: Select a default location for all tapes ejected from this device.

2. Click **Save** to start adding the tape library to NAKIVO Backup & Replication. After successful addition, the tape library will become available in the **Devices** tab.

# Managing Added Tape Library

Clicking the name of the tape library opens its **Parameters** page. In addition to giving the details on the selected tape library, the **Parameters** page provides the following options:

- **Refresh**: Allows for refreshing the device by initiating the process of updating information regarding the content of the tape device. Refreshing involves checking the tapes' barcodes and may include moving tape cartridges within the device
- **Manage**: Allows for performing the following actions with the tape library:
  - **Edit**: Selecting this option opens the same wizard as described in previous sections, but with all fields already predefined. All fields, apart from **Compression** and **Block size**, can be changed.
  - **Detach/Attach**: Allows performing manual tape library attach/detach. Tape cartridges contained in a detached tape device become offline.
  - **Remove**: Removes the device from NAKIVO Backup & Replication. This option is unavailable if the device is currently in use by a job or other process.
- **Tapes**: Clicking this link opens the **Tapes** screen where you can view and manage tape cartridges in the device.



# Adding Standalone Tape Drives

The process of adding a standalone tape drive to NAKIVO Backup & Replication includes the following steps:

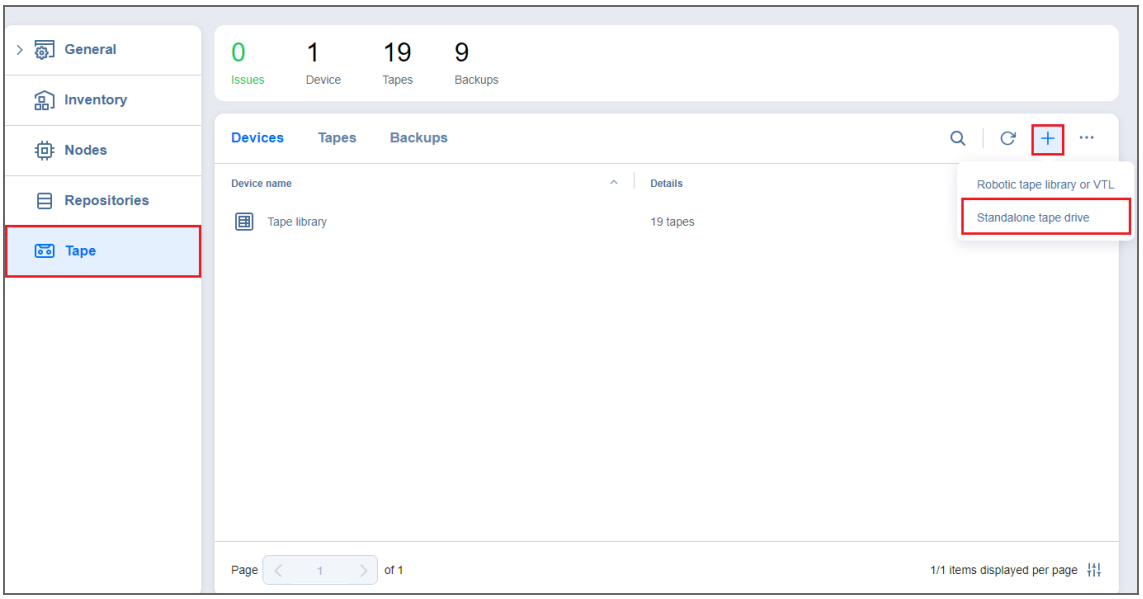
- Launching Wizard
- Selecting Transporter
- Selecting Options
- Managing Added Tape Drives

> **Note**
> Vendor drivers should be installed on tape devices prior to adding them to the NAKIVO Backup & Replication **Inventory**.

## Launching Wizard

To add a standalone tape drive to the system:

1. Go to **Settings** and click the **Tape** tab.
2. Go to the **Devices** tab.
3. Click the **Add (+)** button and select **Standalone tape drive**.



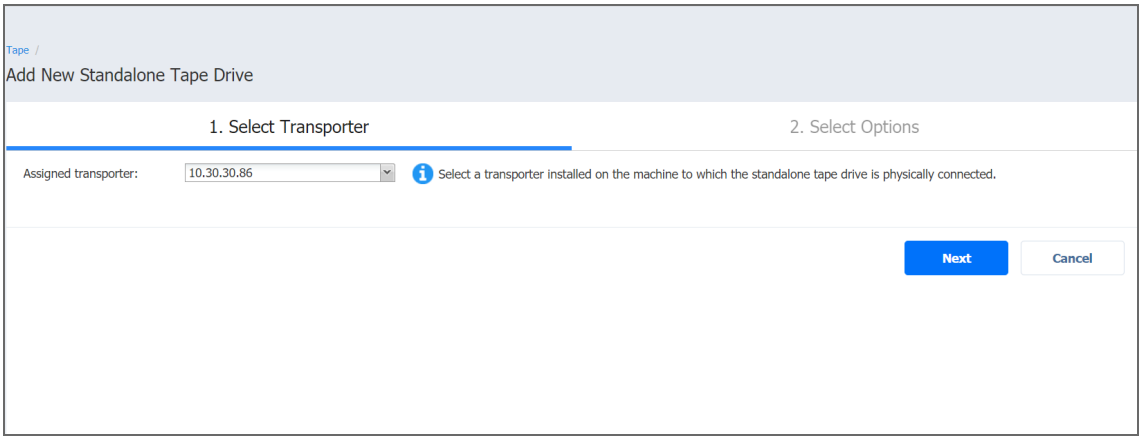

The **Add New Standalone Tape Drive** wizard opens. Follow the steps below to add a new tape drive.

> **Note**
> Before adding a new tape drive to NAKIVO Backup & Replication, you need to deploy or add an existing **Transporter** on a machine that is physically connected to the tape drive. For more information on the**Transporter** deployment, refer to "Deploying Transporter as VMware Appliance" on page 575 and "Adding Existing Nodes" on page 565.

# Selecting Transporter

During the **Select Transporter** step, you need to specify a **Transporter** assigned to the drive that you would like to add. This **Transporter** acts as a network appliance that manages traffic between the tape drive and NAKIVO Backup & Replication.

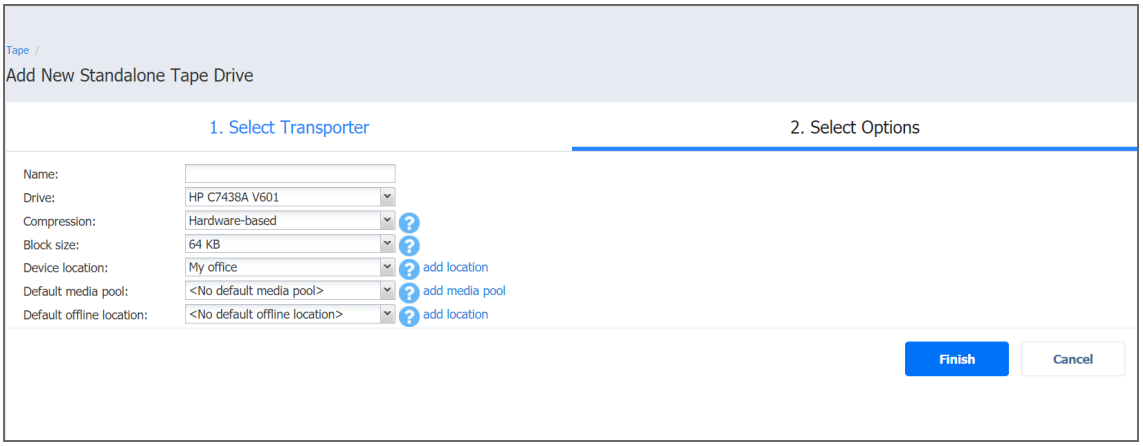
1. From the **Assigned Transporter** drop-down list, select the relevant installed **Transporter**. You can also select a physical machine agent installed on the host connected to the tape device.




2. Click **Next**.

# Selecting Options

The last step of adding a tape drive is selecting its options.




1. Specify the following parameters:

- **Name**: Enter the name of the tape library
- **Drive**: Select one of the standalone tape drives on the assigned transporter
- **Compression**: Select a compression level of the tape device:
    - Software-based (fast) (default)
    - Hardware-based
    - Software-based (medium)

- Software-based (best)

> **Note**
> Setting hardware compression is not recommended to avoid the issue of increased data size during transfer and long backup/recovery times. Having different types of compression for the tape device and a source/target **Backup Repository** during backup or recovery can also lead to this issue. For more information, see this article.

- **Block size**: Select the block size of the tape device:
  - 32 KB
  - 64 KB (default)
  - 128 KB
  - 256 KB
  - 512 KB
  - 1 MB

> **Note**
> The system does not automatically detect the block size; make sure to use the correct block size when importing backups.

- **Device location**: Select the location of the device and all tapes inserted into this device. The automatically created **My office** location is selected by default. To create another location, click **Add Location**. For more information on locations management, refer to "Managing Locations" on page 743.
- **Default media pool**: Select a default media pool for all new tapes inserted into this device. Optionally, you can select **No default media pool** if you want to skip this step. To create another media pool, click **Add Media Pool**. For more information on media pools management, refer to "Managing Media Pools" on page 746.
- **Default offline location**: Select a default location for all tapes ejected from this device. Optionally, you can select **No default offline location** if you want to skip this step.

2. Click **Save** to start adding the tape drive to NAKIVO Backup & Replication. After successful addition, the tape drive will become available in the **Devices** tab.

## Managing Added Tape Drives

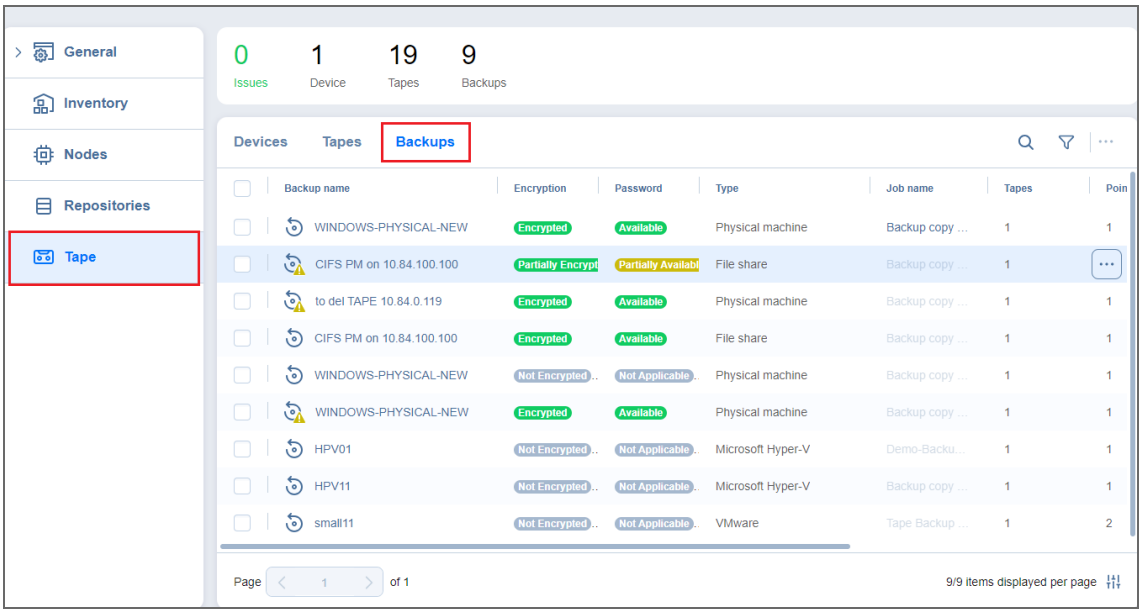
Clicking the name of the tape drive opens its **Parameters** page. Apart from giving details on the selected tape drive, the **Parameters** tab provides the following functionality:

- **Refresh**: Allows for refreshing the device by initiating the process of updating information regarding the content of the tape device.

- **Manage**: Allows for performing the following actions with the tape drive:
  - **Edit**: Selecting this option opens the same wizard as described in previous sections, but with all fields already predefined. All fields, apart from **Compression** and **Block size**, can be changed.
  - **Detach**/**Attach**: Allows for performing manual tape library attach/detach. Tape cartridges contained in a detached tape device become offline.
  - **Remove**: Removes the device from NAKIVO Backup & Replication. This option is unavailable in case the device is currently in use by a job or other process.
- **Tapes**: Clicking this link opens the **Tapes** screen where you can view and manage tape cartridges in the device.

# Managing Backups

From the **Tape** tab, you can also manage all backups stored on tape cartridges by clicking the **Backups** tab.




In this view, you can search for backups and filter them, recover from backups, and view backup details.

- Searching for Backups
- Filtering Backups
- Backups Table
- Recovering from Backups
- “Tape Backup Details” on page 740

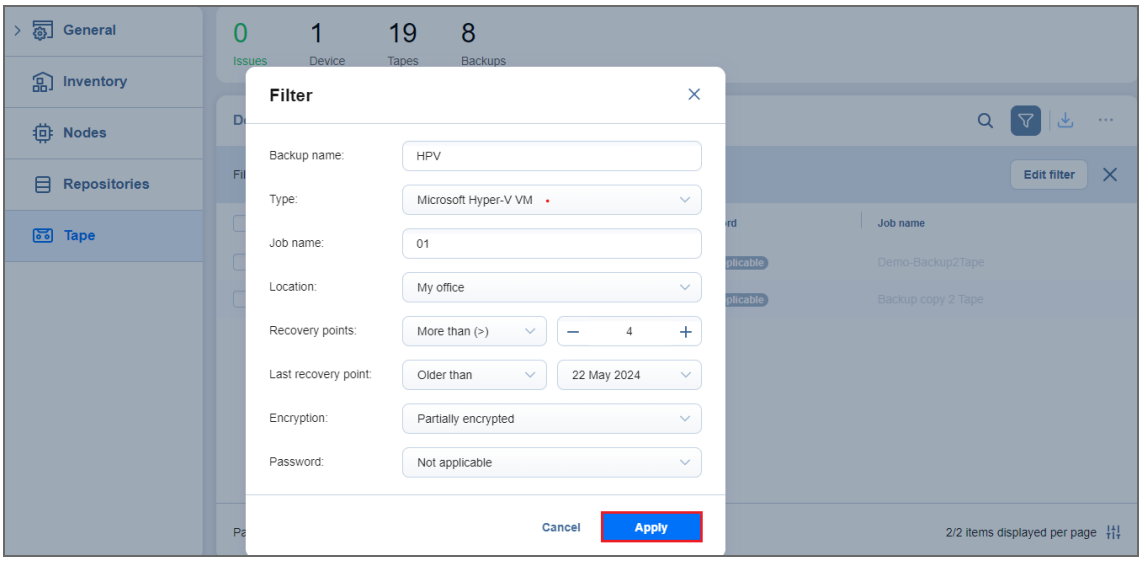
## Searching for Backups

You can search for a specific backup by entering its name (or part of its name) into the **Search** box. The table will dynamically change to display the search results matching your query.

Clicking the **Clear** button in the search box clears the query, and the table displays all backups.

## Filtering Backups

The **Backups** tab also provides sophisticated filtering options that can be applied to search for particular backups. To access filtering options, click the **Filter** icon in the top right corner. In the dialog box that opens, you can select one or several filtering criteria that are applied with the AND statement.




You can apply the following filtering criteria:

- **Backup name**: The backups with the provided name are displayed. Part of the name can be entered.
- **Type**: Specify the type of backups to be displayed:
  - Any
  - VMware VM
  - Hyper-V VM
  - EC2 instance
  - Nutanix AHV VM
  - Physical machine
- **Location**: Only the backups from the tape cartridges of the specified device location are displayed.
- **Job name**: Only the jobs with the specified string in their name are displayed.
- **Recovery points**: Only the backups with less or more recovery points are displayed.
- **Last recovery point**: Only the backups with the last recovery point created on/newer/later than the date specified are displayed.
- **Encryption**: Specify the type of encryption to be displayed:
  - Encrypted
  - Partially encrypted
  - Not encrypted
- **Password**: Specify the password option to be displayed:
  - Available
  - Partially available
  - Not available
  - Not applicable

**Note**

The **Search** and **Filter** features can only be applied separately; you cannot simultaneously search for a tape cartridge by name and select filtering options.

## Backups Table

The **Backups** table provides detailed information about each backup:

- **Backup name**: Displays the name of the backup. Clicking the name opens the **Recovery** screen.
- **Encryption**: The state of encryption of the backup object, which can be one of the following:
  - **Not encrypted**: None of the recovery points of the backup object are encrypted.
  - **Encrypted**: All recovery points of the backup object are encrypted.
  - **Partially encrypted**: Some recovery points of the backup object are encrypted.
- **Password**: The state of the password of the backup object, which can be one of the following:
  - **Available**: All or some recovery points of the backup object are encrypted, and the password hash is available for all encrypted recovery points.

    **Note**

    In case of a self-backup, the status is displayed if all or some recovery points of the backup object are encrypted and the product database contains details of encryption for all encrypted recovery points.

  - **Not available**: All or some recovery points of the backup object are encrypted, and the password hash is not available for all encrypted recovery points.

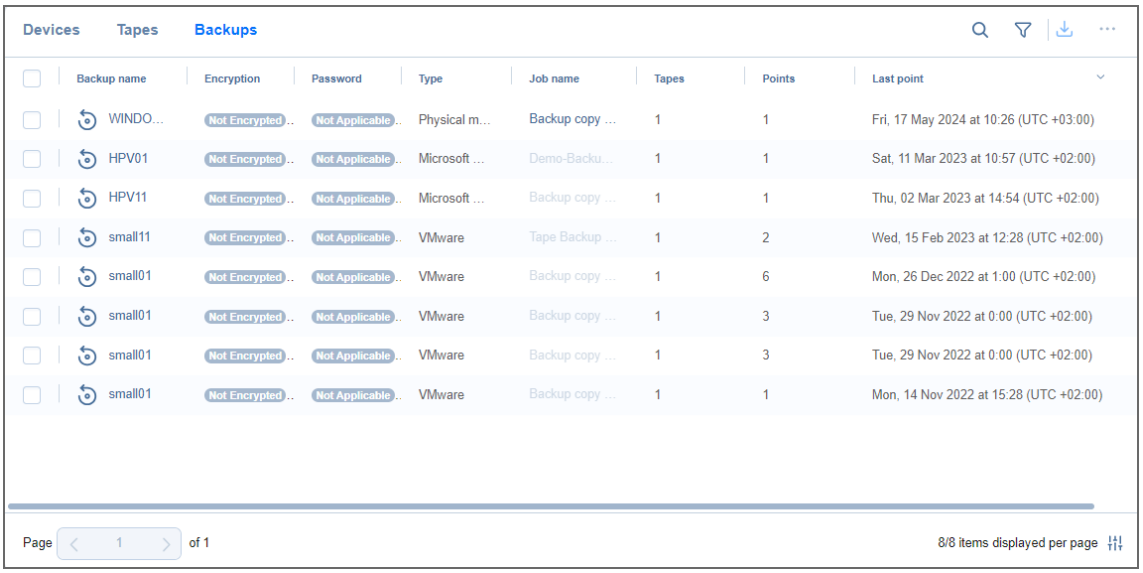
    **Notes**

    - In case of a self-backup, the status is displayed if the recovery point is encrypted and the product database does not contain encryption details for this recovery point.
    - In case of a self-backup, the status is displayed if you are upgrading your product from the version that does not support backup encryption to the version that supports it (for example, from 10.11 to 11.0).
    - After a system configuration backup, the encrypted recovery points passwords that were created and saved in the password manager become unavailable and providing them manually is required.

  - **Partially available**: All or some recovery points of the backup object are encrypted, and the password hash is not available for some encrypted recovery points.

> **Note**
>
> In case of a self-backup, the status is displayed if all or some recovery points of the backup object are encrypted and the product database contains details of encryption for all encrypted recovery points.

  - **Not applicable**: The backup object does not contain encrypted recovery points.
- **Type**: The type of the source object
- **Job name**: The name of the job associated with the backup
- **Tapes**: The number of tape cartridges that the backup occupies
- **Points**: The number of recovery points the backup has
- **Last point**: The date of the last recovery point on the backup
- **Size**: Summarized original size of all recovery points




## Managing Passwords

If the currently selected tape backup is encrypted (some or all recovery points of the backup are encrypted) and the corresponding password hash is not available, you can provide the password manually. To provide the password, do the following:

1. Click **Settings** in the left pane of the product.
2. Go to **Settings > Tape**.
3. In the **Backups** tab, select the checkboxes next to the backup names you want to manage, click the ellipsis **Manage** button, and then select **Provide password** in the drop-down list.
4. Optionally, hover over the backup you want to manage, and on the right side, click **Manage**, then select **Provide password** in the drop-down list.
5. In the dialog box that appears, select the needed password or create a new one.

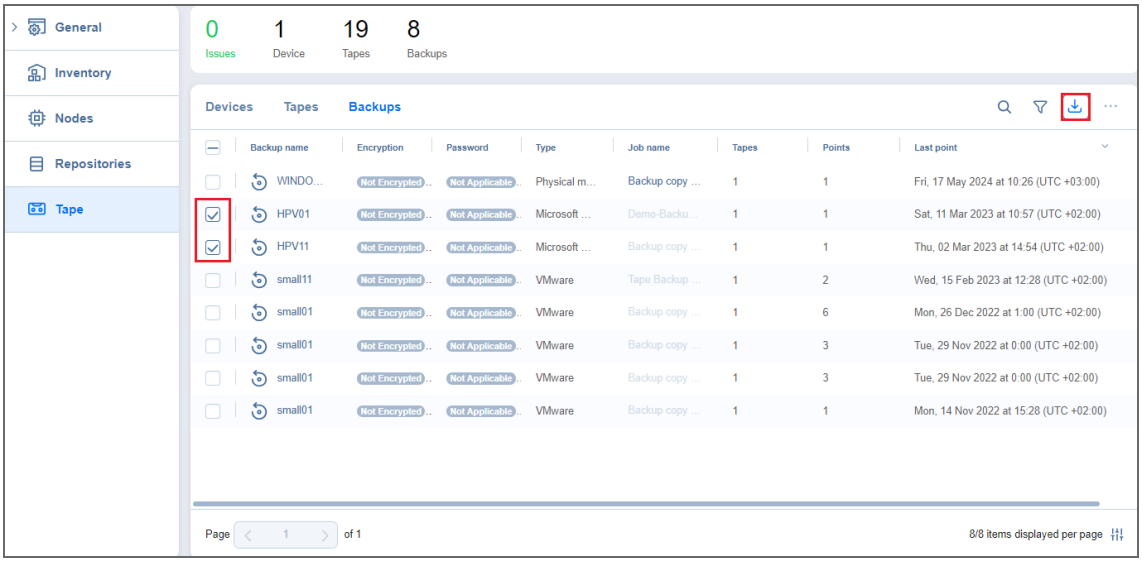
6. Optionally, you can click the **Manage passwords** link to manage the existing or add a new password to the list of passwords.
7. Click **Apply**.

**Notes**

- The **Provide password** option is disabled if the selected backups are not encrypted (none of the recovery points are encrypted)
- The **Provide password** option is disabled if the selected backups are encrypted (some or all recovery points are encrypted) and the password hash is available for all encrypted recovery points.
- The **Provide password** option is disabled for self-backup recovery points created from an older version that does not support backup encryption.
- It's recommended that you enable the **(AWS) Key Management Service** in the **Settings > General > System Settings > Encryption** tab. If **AWS** is enabled, all backup encryption passwords are encrypted with the Key Management Service cryptographic key to be available for recovery in case of product re-installation. For more information, refer to Enabling KMS.

## Recovering from Backups

You can initiate the recovery process from the **Tape > Backups** tab by selecting the checkboxes next to backup names, clicking the **Manage** button, and then selecting **Recover** in the drop-down list.




Optionally, hover over the backup you want to recover, and on the right side, click **Manage**, then select **Recover** in the drop-down list.

**Note**

If you are recovering from multiple backups, you may only select backups of the same **Type**. If you select multiple backup types, the **Recover** button is disabled.

The **New Recovery Job Wizard** opens with the specified backups and their latest recovery points selected.

## Tape Backup Details

Clicking a tape backup name opens the **Backup Details** page where you can view the backup information and see all recovery points available for this backup. You can also initiate the recovery or password management process from here.

The **Backup Details** section provides the following information about the backup:

- **Name**: The name of the backup
- **Type**: The type of backup: VMware VM, Hyper-V VM, EC2 instance or physical machine
- **Tapes**: The number of tape cartridges this backup is stored on
- **Job name**: The name of the job the backup belongs to
- **Recovery points**: The number of recovery points within the backup
- **First recovery point**: The date of the latest recovery point of the backup
- **Last recovery point**: The date of the most recent recovery point of the backup
- **Location(s**): The location the backup is assigned to

You can view the details of a recovery point in the lower part of the screen. To find a recovery point for a specific date, use the **Search** bar on the right.

The **Recovery points** table lists all the recovery points available for the current backup and provides the following information:

- **Date**: The date the recovery point was created. Clicking this parameter initiates recovery for this recovery point.
- **Type**: The type of backup: **Full** or **Incremental**
- **Encryption**: The state of encryption of the backup object, which can be one of the following:
  - **Not encrypted**: None of the recovery points of the backup object are encrypted.
  - **Encrypted**: All recovery points of the backup object are encrypted.
- **Password**: The state of the password of the backup object, which can be one of the following:
  - **Available**: All or some recovery points of the backup object are encrypted, and the password hash is available for all encrypted recovery points.
  - **Not available**: All or some recovery points of the backup object are encrypted, and the password hash is not available for all encrypted recovery points.
  - **Not applicable**: The backup object does not contain encrypted recovery points.
- **Tape**: The name of the tape cartridge the backup is stored on

- **Size**: The total size of the backup
- **Protected until**: The date when the recovery point expires

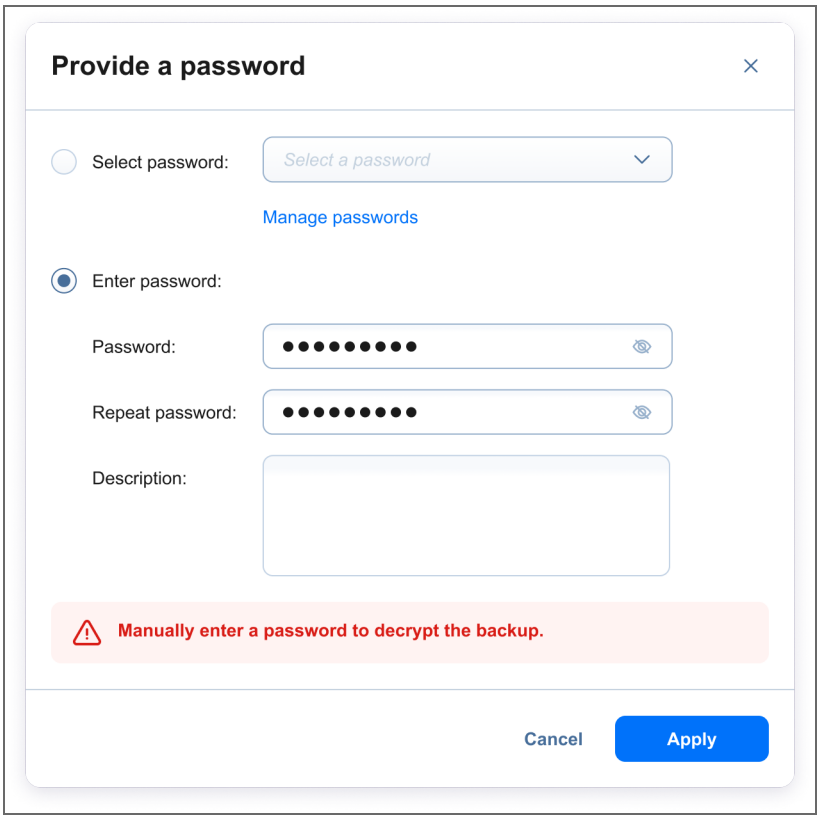
You can initiate the recovery process from the **Recovery points** table. To start recovery, hover over the recovery point you want to recover from, and on the right side, click **Manage**, then select **Recover** in the drop-down list.

This opens the **Recovery from Tape Wizard** for the selected recovery point. For more information about recovering from a tape backup, refer to Starting Recovery from Tape.

If the currently selected recovery point is encrypted and the corresponding password hash is not available, you can provide the password manually.

To provide the password, do the following:

1. Click **Settings** in the left pane of the product.
2. Go to **Settings > Tape**.
3. In the **Backups** tab, click a tape backup name.
4. In the **Recovery points** table, hover over the recovery point you want to manage, and on the right side, click the ellipsis **Manage** button, then select **Provide password** in the drop-down list.
   Optionally, you can select the checkboxes next to the recovery points you want to manage, and click the **Manage** button.
5. In the dialog box that appears, select the needed password or create a new one.
6. Optionally, you can click the **Manage passwords** link to manage the existing or add a new password to the list of passwords.




7. Click **Apply** to proceed.

**Notes**

- The **Provide password** option is disabled if the selected recovery points are not encrypted (none of the recovery points are encrypted).
- The **Provide password** option is disabled if the selected recovery points are encrypted (some or all recovery points are encrypted) but the password hash is available for all encrypted recovery points.
- The **Provide password** option is disabled for the self-backup recovery point created from an older version.
- It's recommended that you enable the **(AWS) Key Management Service** in the **Settings > General > System Settings > Encryption** tab. If **AWS** is enabled, all backup encryption passwords are encrypted with the Key Management Service cryptographic key to be available for recovery in case of product re-installation. For more information, refer to Enabling KMS.

# Managing Locations

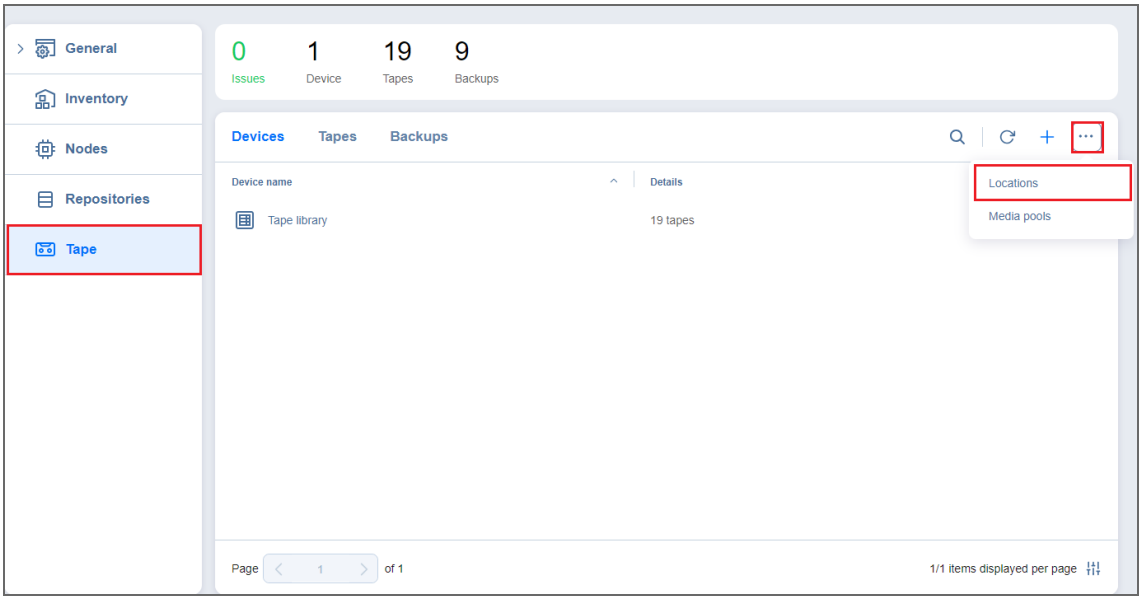
Device location is a logical container representing a geographical place where the tape devices are located. Larger companies can have their tape devices in different locations, e.g. the UK, USA, Australia, etc. By default, the system automatically creates the **My Office** device location, but you can create more device locations if necessary. Refer to these sections for details:

- Adding Device Locations
- Managing Device Locations

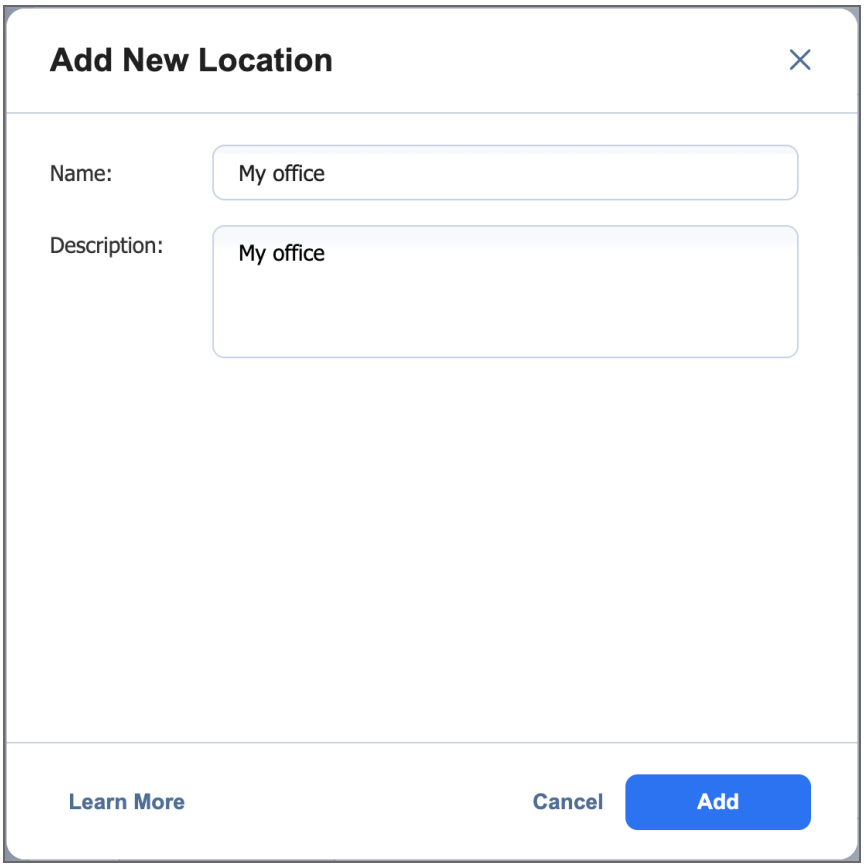
## Adding Device Locations

To add a Location:

1. Go to **Settings** > **Tape**.
2. In the **Devices** or **Tapes** tab, click the ellipsis **Manage** button and select **Locations**. The **Location Management** dialog box opens.



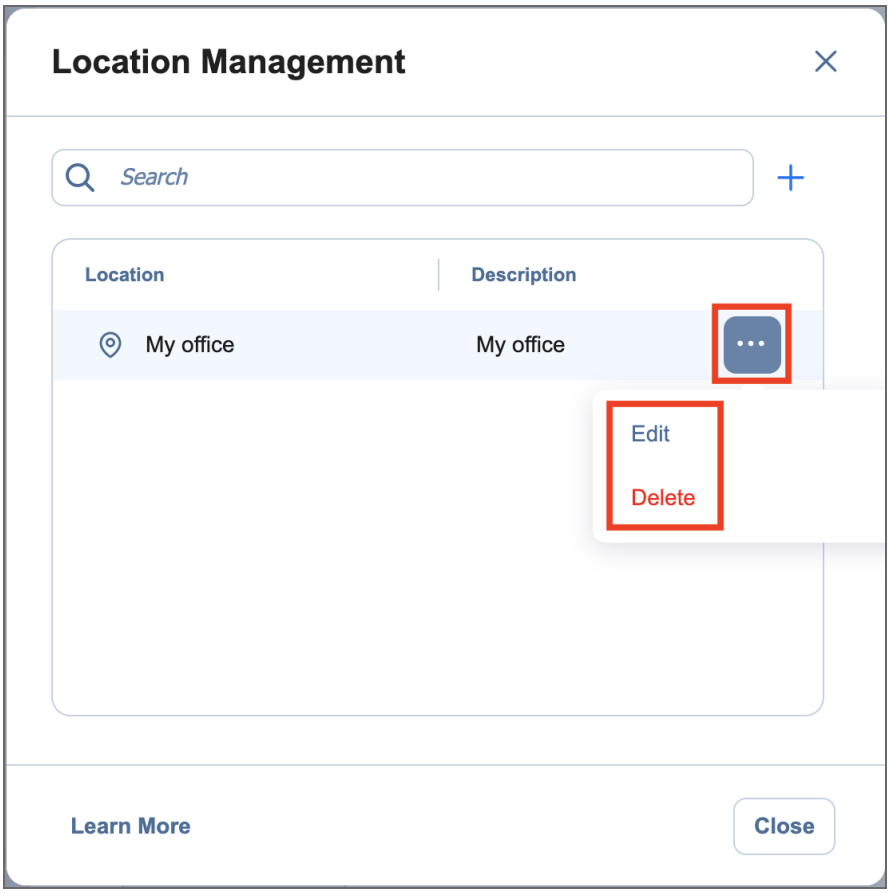

3. Click the plus **Add New Location** button.
4. In the **Add New Location** dialog box, specify a name for the device location and provide a description (optional).




5. Click **Add**. The new device location is added to the list.

## Managing Device Locations

From the **Location Management** screen, you can also edit or delete Locations by using the corresponding buttons or search for the location by entering a location name (or a part of its name) into the **Search** box.


Location Management
Search
Location
Description
My office
My office
Edit
Delete
Learn More
Close

# Managing Media Pools

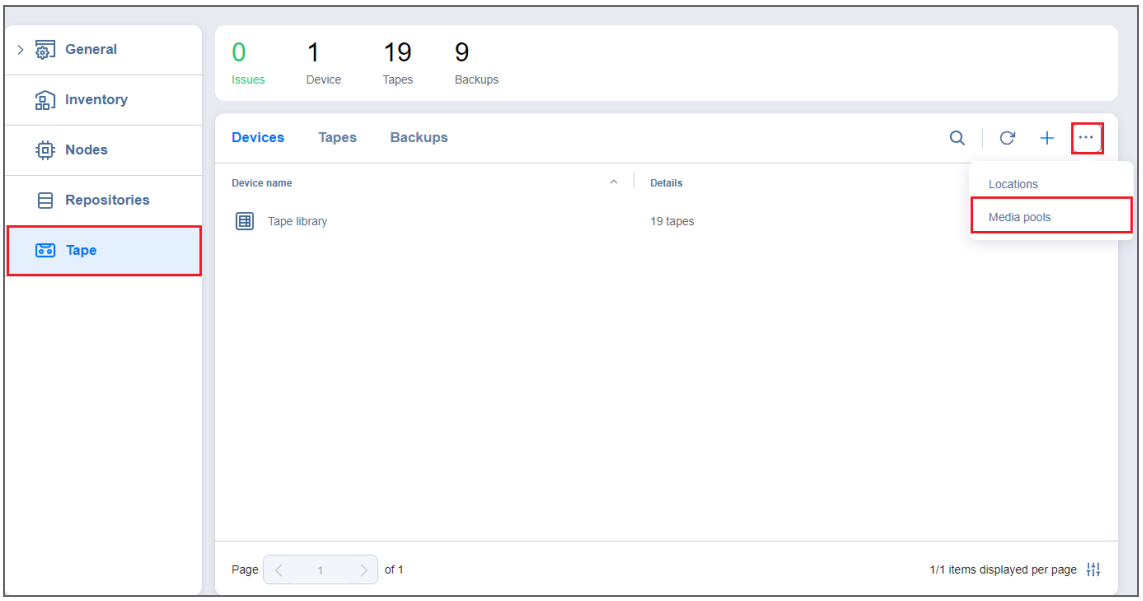
Media pools are logical containers created in NAKIVO Backup & Replication to organize and manage tape cartridges. No Media Pools are created by default, but you can create new ones if necessary. for details, refer to the following sections:

- Adding Media Pools
- Managing Media Pools

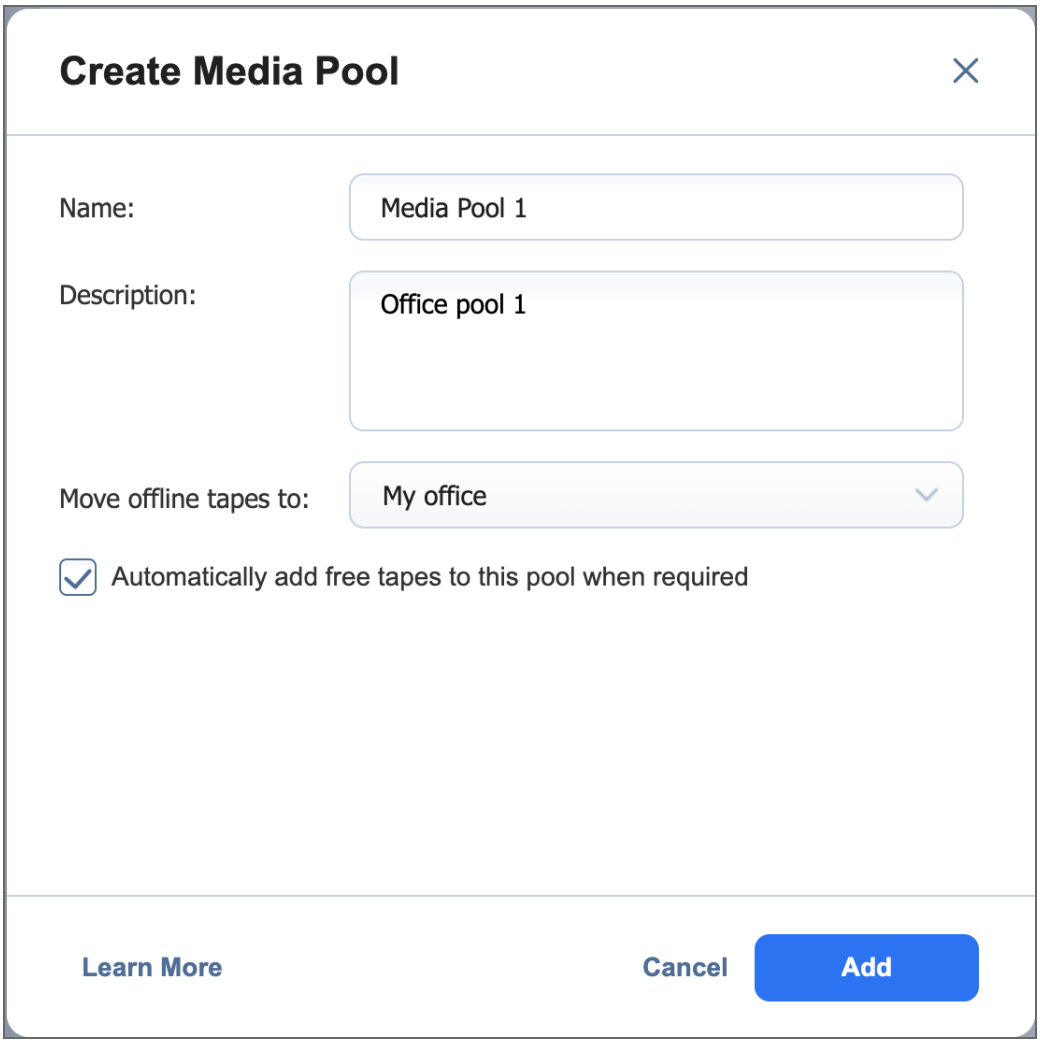
## Adding Media Pools

To create a Media Pool:

1. Go to **Settings** > **Tape**.
2. In the **Devices** or **Tapes** tab, click the ellipsis **Manage** button and select **Media Pools**. The **Media Pool Management** dialog box opens.



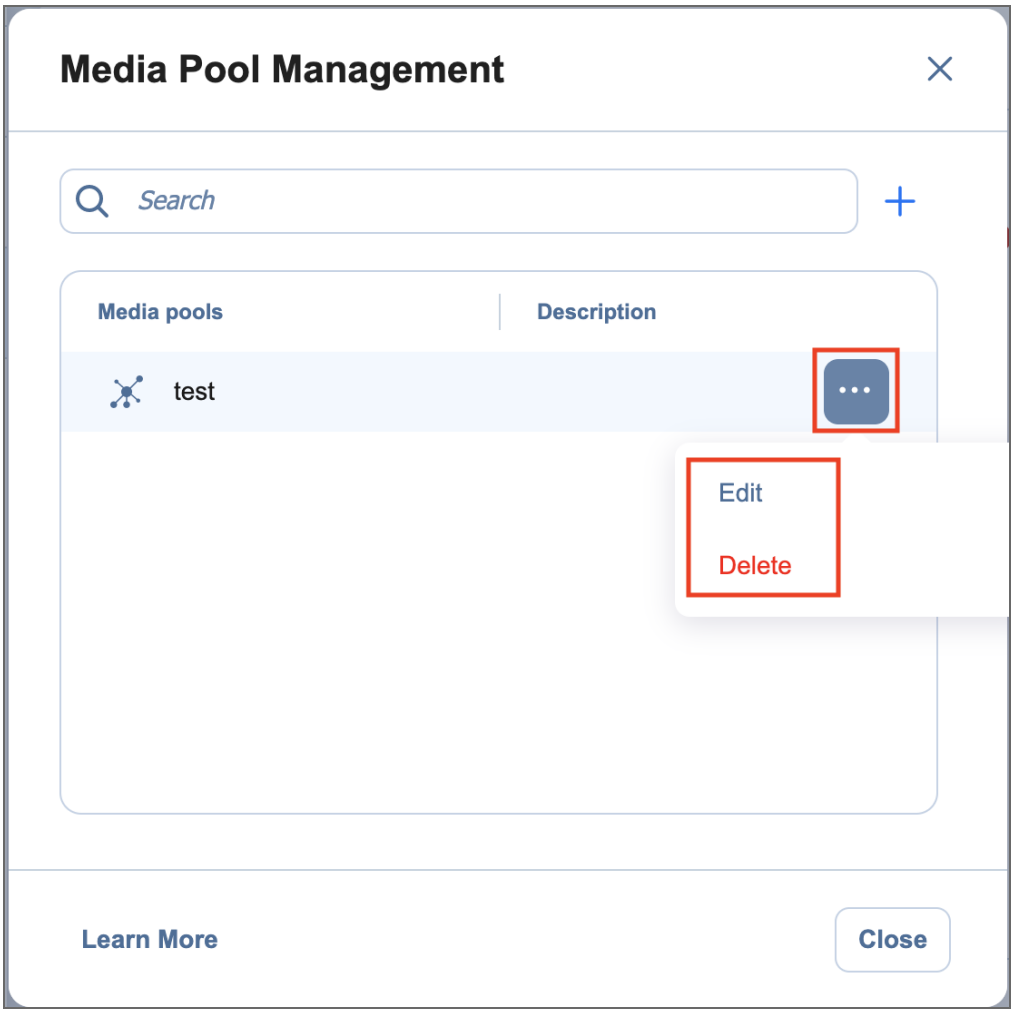

3. Click the plus **Create Media Pool** button.
4. In the **Create Media Pool** dialog box, specify the name for the Media Pool and provide a description (optional).
5. From the **Move Offline Tapes To** drop-down list, select a device location to determine which location is automatically set for all offline tapes from this media pool. If the tape cartridge goes online again, it will return to the initial device location.
6. Select the **Automatically add free tapes to this pool when required** checkbox to automatically add one of the empty available tape cartridges to this media pool if the media pool does not have available tape cartridges.




7. Click **Add**. The new Media Pool is created.

## Managing Media Pools

From the **Media Pool Management** screen, you can also edit or delete Media Pools by using the corresponding buttons or search for the media pool by entering its name or a part of its name into the **Search** box.


Media Pool Management
Search
Media pools
Description
test
Edit
Delete
Learn More
Close

# Managing Tape Cartridges

The **Tapes** view allows you to view and manage all tape cartridges registered in the system. This section covers the following topics:


- Viewing Tapes
- Searching for Tape Cartridges
- Filtering Tape Cartridges
- Tape Cartridge Management Page
    - Manage Options
    - Details Pane
    - Tape Contents table
    - Backup Details
- Bulk Tape Cartridge Management


## Viewing Tapes

To navigate to the **Tapes** menu, go to **Settings** > **Tapes** and click the **Tapes** tab.

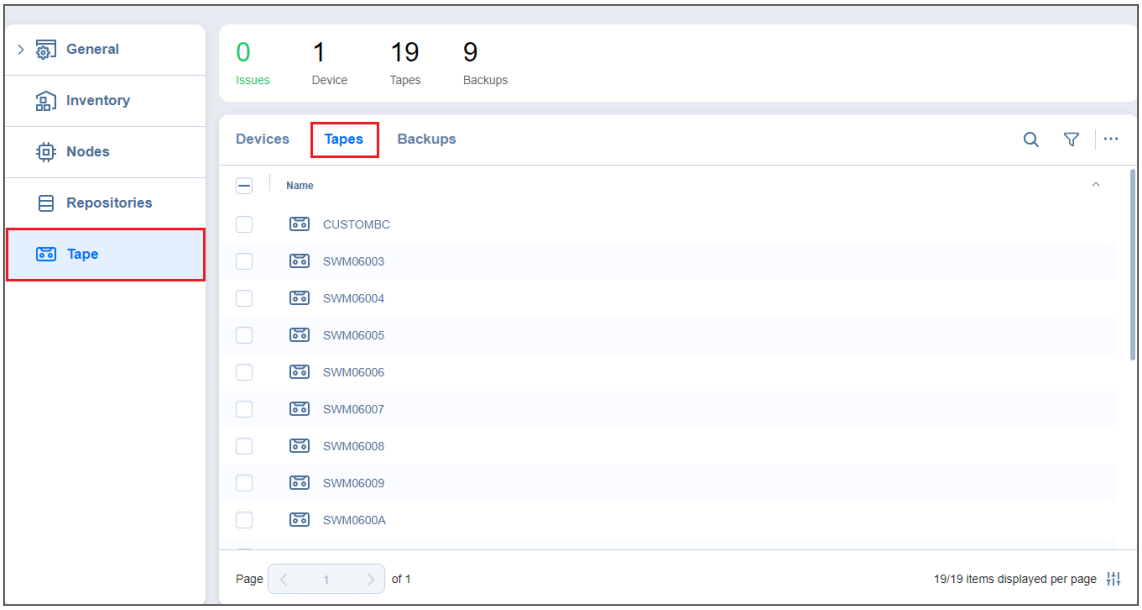
The **Tapes** menu provides you with the following information about the tape cartridges in the table:

- **Name**: Displays the tape cartridge name. Clicking the name opens the tape cartridge management page. For more information, see Tape Cartridge Management Page.
- **Label**: Displays the label assigned to the tape cartridge ("none" for tape cartridges without labels)
- **Status**: Displays the current status of the tape cartridge—Scanning / Online / Reading / Writing / Erasing / Warning / Error / Offline
- **Device**: Displays the name of the tape device that contains the tape cartridge
- **Slot/Drive**: Displays the slot/drive number of the tape cartridge
- **Last Written**: Displays the date of the last recording on the tape cartridge
- **Overwritable**: Displays the date when all recovery points on this tape cartridge will expire
- **Media Pool**: Displays the name of the media pool that the tape cartridge belongs to
- **Location**: Displays the name of the device location that the tape cartridge belongs to
- **Contents**: Indicates the contents of the tape cartridge (e.g. number of backups on the tape cartridge)
    - **Capacity**: Displays the amount of free space relative to the total tape capacity. Hovering over this row also reveals the amount of used space. In case the capacity cannot be retrieved, **Not available** will be displayed instead.

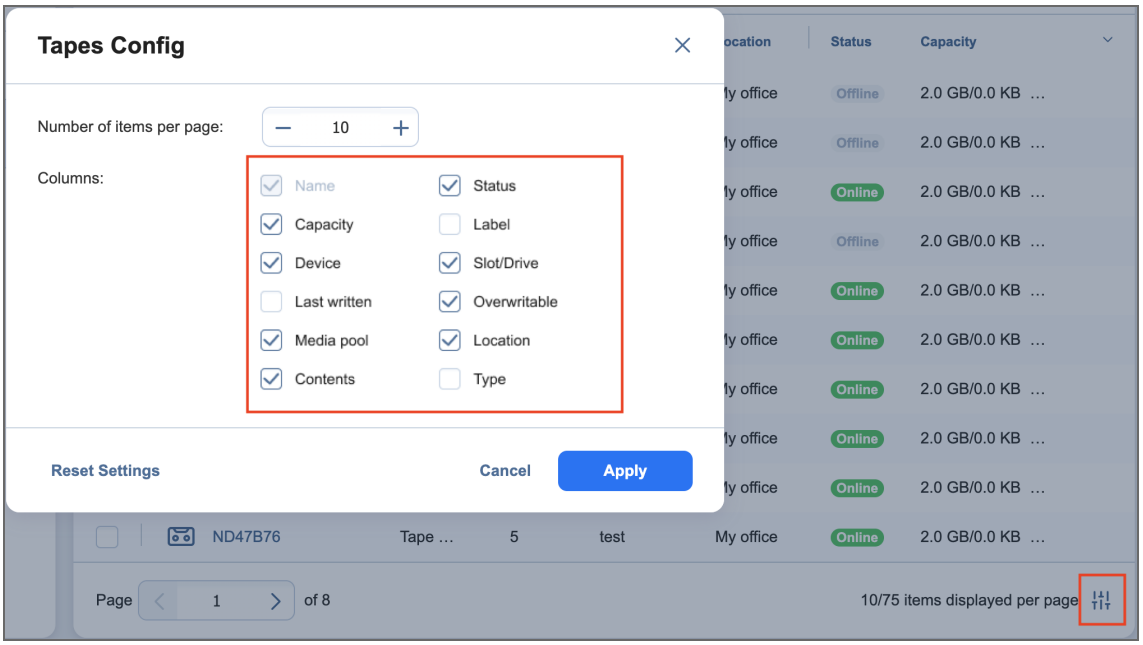
**Note**

If hardware compression is enabled when writing data to a tape cartridge, NAKIVO Backup & Replication may display twice the amount of total, free, and used space. For instance, an LTO-6 cartridge has a native capacity of 2.5 TB, while the 6.25 TB shown in some device UIs reflects compressed capacity based on a 2.5:1 ratio. See also LTO Specifications.

- **Type**: Displays the type of the tape cartridge:
  - Read/Write Tape
  - Write Protected Tape
  - Cleaning Tape

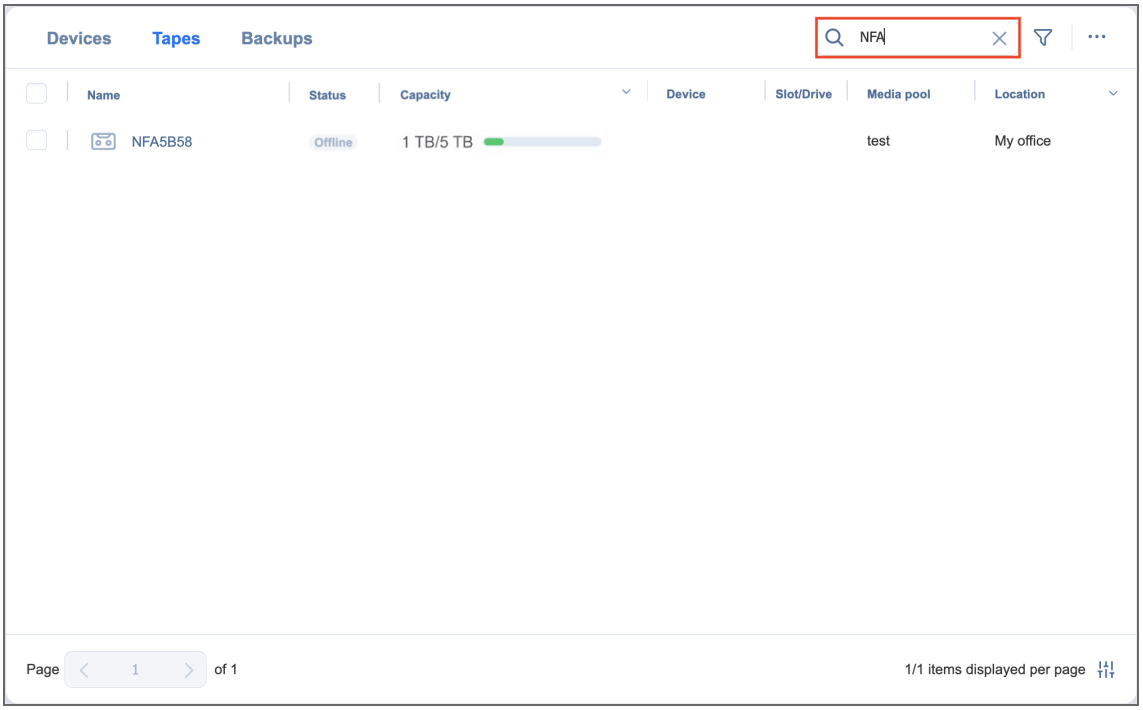



The column availability in the table can be managed by clicking the small config button in the bottom right of the table and checking/unchecking the boxes next to the column names.




# Searching for Tape Cartridges

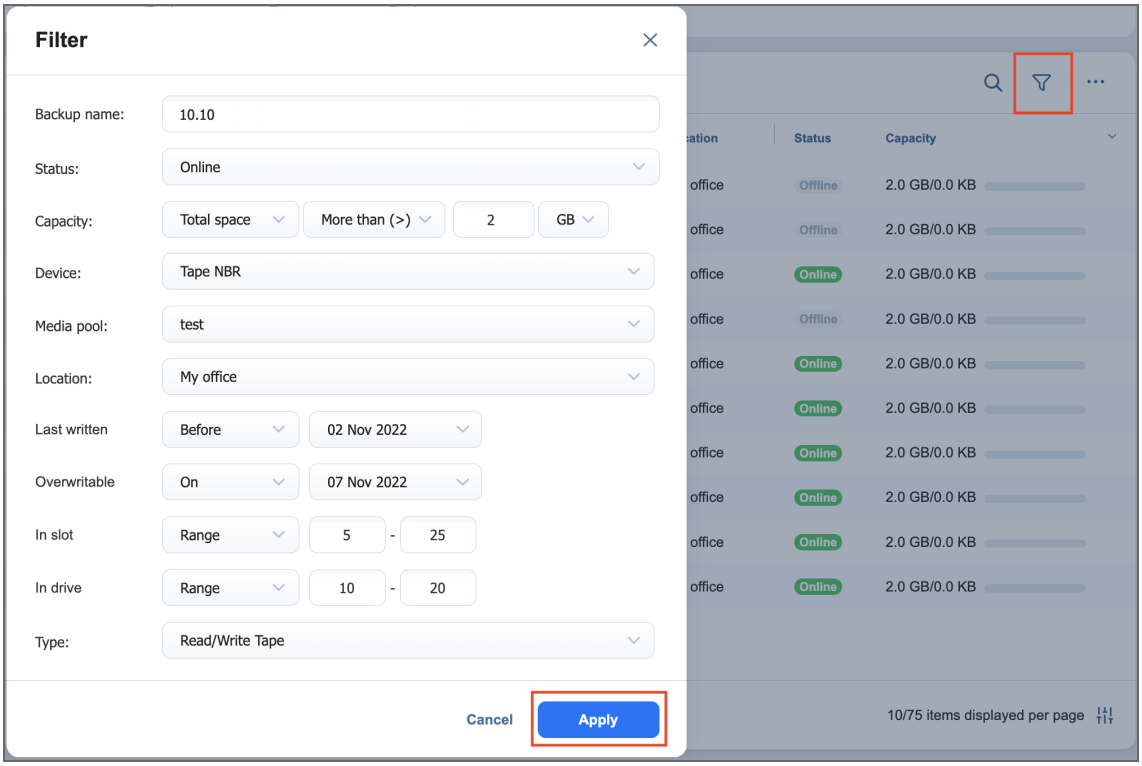
You can search for a specific tape cartridge by entering its name (or part of its name) into the **Search** box. The table will dynamically change to display the search results matching your query.




Clicking the **Clear** button in the search field will clear the query and the table will display all tape cartridges.

# Filtering Tape Cartridges

The **Tapes** view also provides sophisticated filtering options that can be applied to search for particular tape cartridges. To access filtering options, click the **Filter** button next to the **Search** box. In the **Filter** pop-up, select one or several filtering criteria that will be applied with the AND statement.




You can apply the following filtering criteria:

- **Backup name**: Tape cartridges containing the backups with the provided name will be displayed.
- **Status**: Tape cartridges in one of the following statuses will be displayed:
  - Offline
  - Online
  - Scanning...
  - Erasing...
  - Cleaning...
  - Reading...
  - Writing...
  - Moving...
  - Warning
  - Error
- **Capacity**: Filter by capacity by configuring the following options:
  - Select one of **Total space**, **Free space**, or **Used space**.
  - Select one of **More than (>)**, **Less than (<)**, or **Equal to (=)**.

    - Enter a value corresponding to the desired capacity in GB or TB.
    - Select either **GB** or **TB**.
- **Device**: Only the tape cartridges from the specified type device will be displayed.
- **Media Pool**: Only the tape cartridges from the specified media pool will be displayed.
- **Location**: Only the tape cartridges from the specified device location will be displayed.
- **Last Written**: Displays the tape cartridges that have the data written to it on/before/after a specified date
- **Overwritable**: Displays the date when the tape cartridge can be overwritten (calculated using the age and retention of all recovery points on this tape cartridge)
- **In Slot**: Displays the tape cartridges in a specified slot or range of slots
- **In Drive**: Displays the tape cartridges in a specified tape drive or range of tape drives
- **Type**: Displays the tape cartridges according to their type:
    - Read/Write Tape
    - Write Protected Tape
    - Cleaning Tape

**Note**
The Search and Filter features can only be applied separately; you cannot simultaneously search for a tape cartridge by name and select filtering options.

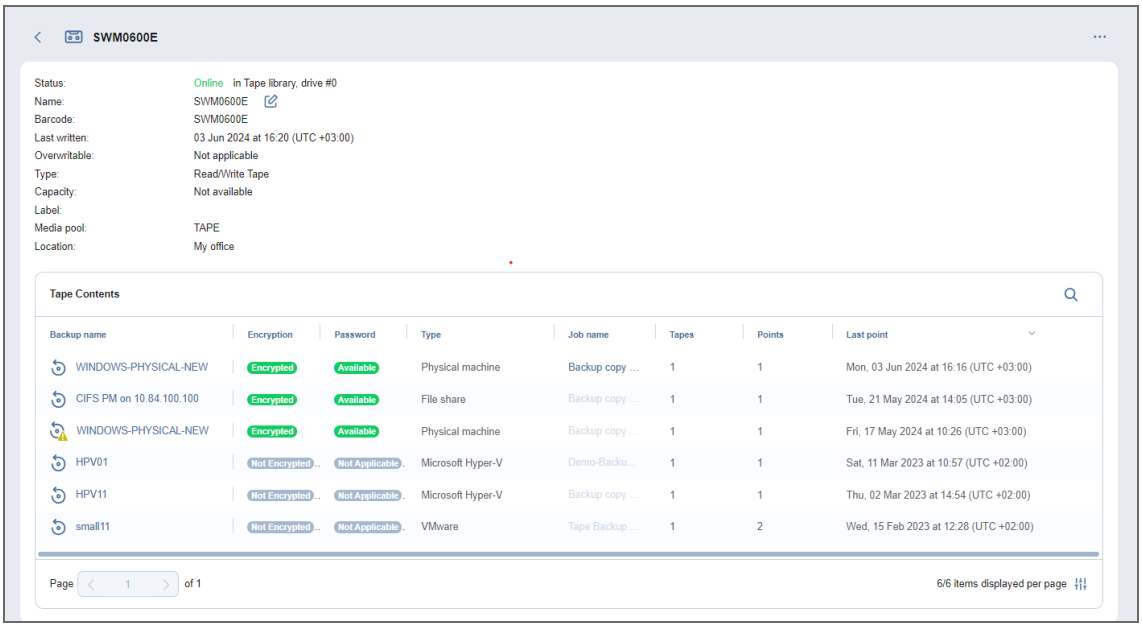
## Tape Cartridge Management Page

Clicking a tape cartridge name opens the tape cartridge management page where you can apply certain actions to the tape cartridge or get extensive information about it.

The tape cartridge management page consists of the following functional blocks:

- **Manage** button
- Detailed tape cartridge information

- **Tape Contents** table




## Manage Options

The ellipsis **Manage** button allows you to perform particular actions with the tape cartridge. Depending on the tape cartridge state, type, status, etc., the button's availability may vary. The button can be disabled if a certain action cannot be applied to the tape cartridge. Hovering over the disabled button displays a tooltip describing the reason for action unavailability.

Some of the actions can be applied to several tapes at once. For more information, refer to "Bulk Tape Cartridge Management" on page 758.

The following actions are available:

- **Scan**: Scans the tape cartridge for its contents. The system recognizes the contents to be:
  - Known NAKIVO Backup & Replication backups: Such content requires no scanning. The backups contained on this tape cartridge are displayed in the Tape Cartridge Contents Table and can be used for VM restoring.
  - Unknown NAKIVO Backup & Replication backups: The system recognizes the contents as created by NAKIVO Backup & Replication (that is, on another product instance) but cannot be used for VM restores until scanned.
  - Empty: The tape cartridge contains no data and is ready to be used for backup.
  - Third Party Data: The tape cartridge contains some third party data that cannot be recognized by NAKIVO Backup & Replication. Such tape cartridges cannot be used unless their contents are erased.

    - Incomplete Backups: The tape cartridge contains incomplete backup(s), the result of an inappropriately finished backup job (for example, in the case that a backup copy job was stopped by the user and the backup copy was not completed). Incomplete backups cannot be used for recovery.
    - Unidentified: The contents of a newly introduced tape cartridge is unknown to the system and must be scanned first.

When you insert new tape cartridges into the tape device, and these tape cartridges contain backups created using another instance of NAKIVO Backup & Replication, the application opens the **Scan new tape cartridges?** notification box, asking if you'd like to scan the new tape cartridges. Clicking the **Scan all** link initiates the scanning action for all newly discovered tape cartridges.

- **Edit**: Clicking the **Edit** button opens the **Edit Tape** dialog where you can:
    - Create or change a label for the tape for easier tape identifying
    - Assign the tape to a pre-created media pool
    - Allocate the tape to a pre-created location

  The newly added details are displayed in the **Options** pane.
- **Move**: This action allows you to move the tape cartridge to an available drive slot or tape drive. Occupied drive slots or tape drives are disabled in the menu.
- **Protect**: Applying this action to the tape cartridge protects it from data overwriting. This action is only available on tape cartridges that contain recovery points. Recovery from protected tape cartridges is available. Protected tape cartridges can be reverted by clicking the **Unprotect** button. Clicking the **Protect** or **Unprotect** button requires confirmation.
- **Mark as free**: Marking the tape cartridge as free makes it eligible for writing backups to it. Marking the tape cartridge as free does not erase the data right away: the next time the product needs a tape cartridge for writing data, it can take this tape cartridge and do a quick-erase before writing new data to it. The button is not available in case the tape cartridge is protected or empty. Marking the tape cartridge as free requires confirmation. The confirmation box displays detailed information about the data that is about to be deleted. This action cannot be undone.
- **Mark as cleaning**: Specialized tape cartridges designed for tape drive cleaning need to be marked as cleaning tapes. For tape cartridges that have been marked as cleaning tapes, this option is replaced by the **Mark as data** button. Selecting **Mark as data** reverts a cleaning tape cartridge to a regular data tape cartridge.

**Important**

Currently, cleaning tapes inserted into a device are not automatically recognized by the system as cleaning. Instead, the system identifies the tapes to contain third party data. It is the user's responsibility to mark the tape as cleaning once the tape is inserted into the device and discovered. Otherwise, the cartridge will perform cleaning of the drive automatically every time the library is refreshed.

- **Retire**: The tape cartridges marked as retired will not be used for new backups. Recovery from retired tape cartridges is still available. The action is not available for tape cartridges marked as free or do not contain recovery points. This action requires confirmation.
- **Erase**: The contents of the tape can be erased using:
  - Quick erase: The data is marked as deleted without actual data deletion. Such data can still be recovered if necessary.
  - Full erase: Deletes the data from the tape forever.

  Keep in mind that both methods can be very time-consuming.

  Clicking the **Erase** button opens the **Erase selected tape?** dialog providing detailed information about the data that is about to be deleted and allows choosing the erase method.
- **Clean drive**: This action is only available for cartridges marked as *cleaning*. Select a drive from the drop-down menu and click **Clean Drive** to initiate the drive cleaning cycle and move the cleaning cartridge to the selected drive for cleaning.
- **Remove**: Clicking this button will physically remove the tape cartridge from the tape device. The button is only available for the offline tape cartridges. The action requires your confirmation.

## Details Pane

The **Details** pane provides full information about the tape:

- **Status**: Displays the status of the tape cartridge and the current tape device name and drive slot/tape drive number. The tape cartridge can be in one of the following statuses: Scanning, Online, Reading, Writing, Erasing, Warning, Moving, Error, or Offline.
- **Name**: Displays the name of the tape cartridge; can be modified by clicking the **Edit** button
- **Barcode**: Displays the tape cartridge barcode if available
- **Last written**: Displays the date of the last write operation
- **Overwritable**: Displays the date when all recovery points on this tape cartridge expire
- **Type**: Displays the type of the tape cartridge: Read/Write Tape, WORM Tape, Write Protected Tape, or Cleaning Tape
- **Capacity**: Displays the amount of free space relative to the total tape capacity. Hovering over this row also reveals the amount of used space. In case the capacity cannot be retrieved, **Not available** will be displayed instead.

- **Label**: Displays the tape cartridge label, if any
- **Media pool**: Displays the assigned media pool, if any
- **Location**: Displays the assigned location, if any

## Tape Contents Table

The **Tape Contents** table provides information about the backups residing on the tape cartridge and allows for recovering VMs from backups right from the table. In case the tape cartridge contains no backups or has not yet been scanned, the table displays generic information about the tape cartridge contents, such as:

- "This tape contains third party data."
- "This tape cannot be identified. Scan the tape in order to discover its content."
- "This tape is empty."
- "This tape contains backups. Scan the tape to view the list of backups."

If the tape cartridge contains backups and has been scanned already, the **Tape Contents** table displays the backups and provides the following information:

- **Backup name**: Displays the name of the backup. Clicking the name of the backup opens the **Recovery** page.
- **Type**: Displays the type of a backup: VMware VM, Hyper-V VM or EC2 instance
- **Job**: Displays the name of the last known job
- **Tapes**: Shows the number of tape cartridges this backup is stored on
- **Points**: Displays the number of recovery points in the backup
- **Last point**: Displays the date of the most recent recovery point in the backup

The **Tape Contents** table can be modified to display the column you need by clicking the arrow icon in the table header and selecting the required columns.

Clicking the column's header sorts the contents of the column.

## Backup Details

Clicking a backup name in the **Tape Cartridge Contents** table opens the **Backup Details** page where you can view the backup information and see all recovery points available for this backup. You can also initiate the recovery process from here.

The **Backup Details** section provides the following information about the backup:

- **Name**: Displays the name of the backup
- **Type**: Displays the type of backup: VMware VM, Hyper-V VM, EC2 instance or physical machine
- **Tapes**: Displays the number of tape cartridges this backup is stored on
- **Recovery points**: Displays the number of recovery points within the backup
- **First recovery point**: Displays the date of the latest recovery point of the backup
- **Last recovery point**: Displays the date of the most recent recovery point of the backup

- **Location**: Displays the location the backup is assigned to
- **Job name**: Displays the name of the job the backup belongs to

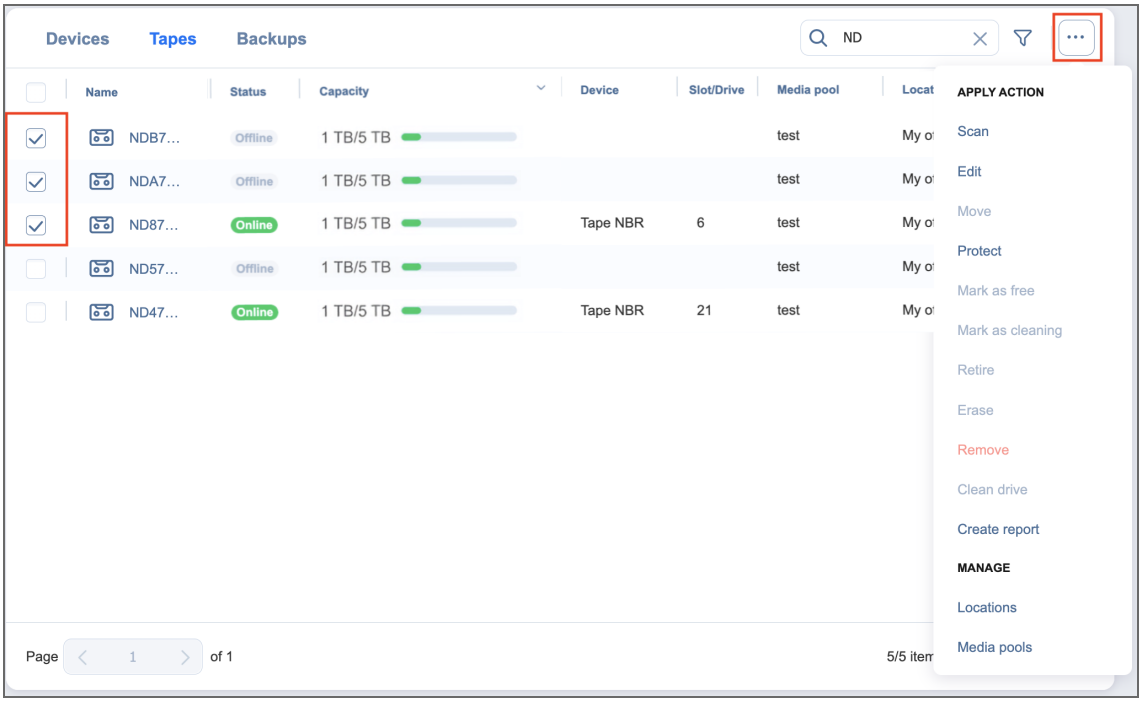
The **Recovery points** table lists all the recovery points available for the current backup and provides the following information:

- **Date**: Indicates the date the recovery point was created. Clicking this parameter initiates recovery for this recovery point.
- **Type**: Indicates the type of backup: Full or Incremental
- **Tape**: Indicates the name of the tape cartridge the backup is stored on
- **Protected until**: Indicates the date when the recovery point expires

Clicking either the **Recover** button or the date of the recovery point in the table opens the Recovery from Tape Wizard for the selected backup object. For more information about recovering from a tape cartridge, refer to “Starting Recovery from Tape” on page 1093 .

## Bulk Tape Cartridge Management

Certain actions can be applied to several tape cartridges simultaneously. While in the **Tapes** tab, select the checkbox next to the tape cartridges you need to apply an action to and click the ellipsis **Manage** button. In the dialog box that opens, select an action to apply. Note that the availability of actions depends on various factors, so not all actions may be available. For action descriptions, refer to the **Manage Options** section above.




The **Create report** action is unique to the bulk tape cartridge management and is used to generate reports about selected tape cartridges. The report is created as a PDF file and is stored locally on your computer.

# 2 tapes

## 56

| | |
|---|---|
| Barcode: | none |
| Last written: | Not applicable |
| Tape Label: | none |
| Media pool: | none |
| Location: | My office |
| Capacity: | 780.5 GB |
| Contents: | Incomplete backup(s) |

### Alarms & Notifications

No alarms or notifications

## 45

| | |
|---|---|
| Barcode: | none |
| Last written: | Wed, 30 Nov at 22:12 |
| Tape Label: | none |
| Media pool: | none |
| Location: | My office |
| Capacity: | 780.5 GB |
| Contents: | 2 backups |

| Name | Date | Type | Expires |
|---|---|---|---|
| VM1 | 02 Nov 2022 at 11:00 | Full | 02 Dec 2022 |
| VM1 | 11 Oct 2022 at 11:31 | Full | 10 Nov 2022 |

### Alarms & Notifications

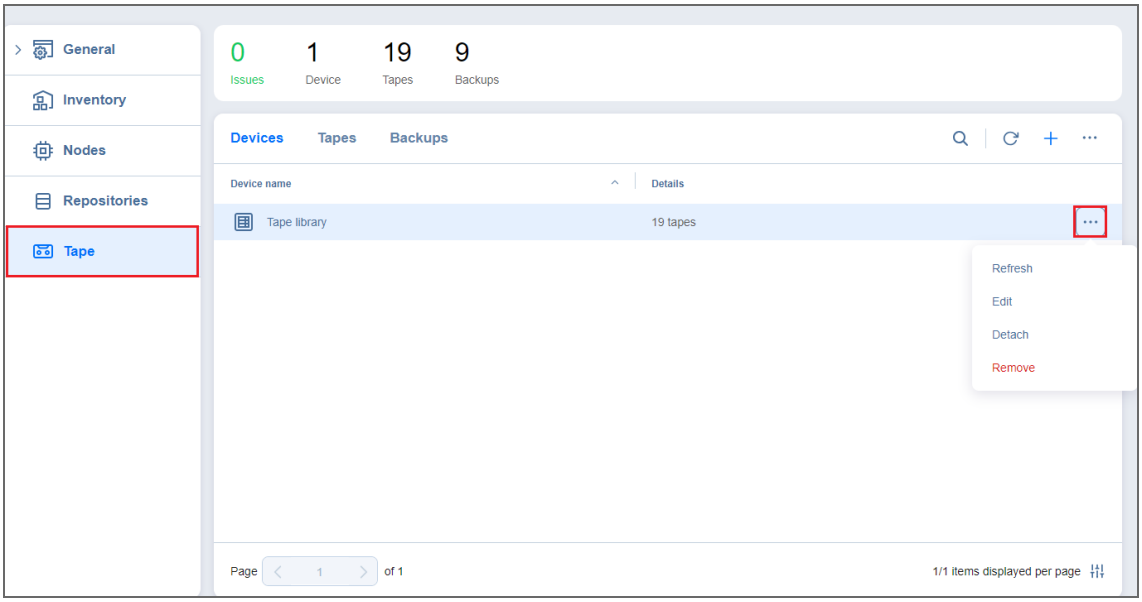
No alarms or notifications

# Managing Tape Devices

Once the tape devices are added to the system, you can view and manage them in the **Devices** tab.

Hovering the mouse cursor over the device name opens the management controls:

- **Manage**: opens the following options:
  - **Refresh**: Refresh action initiates the process of updating information regarding content of the tape device.
  - **Edit**: Opens the **Edit Tape Library** or **Edit Standalone Tape Drive** wizard, depending on the type of the device, where you can change the device's properties. Detached devices are greyed out in the interface and cannot be interacted with
  - **Detach**: Detaching a tape device saves the device's data and metadata in a consistent state and then stops the product's interaction with the device (such as read and write of data and metadata, and so on). You may want to detach a tape device to move it to a different location or to put it on maintenance.
  - **Remove**: Removes the tape device from the **Inventory**. The device may be then added again, for example, if you need to change the block size or compression type of the device.




- Clicking the name of the tape device opens the device's details window where you can manage it and

view the device's detailed information.



# Expert Mode

For advanced NAKIVO Backup & Replication configuration, you can enable the Expert mode.

To do this, take the following steps:

1. Log in to your NAKIVO Backup & Replication instance.
2. Add the word “expert” to the URL parameters of the **Settings** page.

   **Examples**:
   `https://localhost:4443/c/configuration?expert` or
   `https://localhost:4443/c/configuration?action=&targetId=&backUrl=&wizard=false&expert`

3. Click the **Expert** tab.

## Configuring Settings

To configure advanced product settings, make the necessary changes in the following parameters:

| Parameters | Description | Possible Values |
|---|---|---|
| **system.email.smtp.localhost.mode** | Specifies how to determine the name of the localhost that is used in the SMTP HELO or EHLO commands. | • Default<br>• Use DNS<br>• Provide custom hostname |
| **system.email.smtp.localhost.name** | Specifies the name of the localhost that is used in the SMTP HELO or EHLO commands. This setting is valid for custom hostname resolution mode only. | |
| **system.email.smtp.tls.version** | Specifies the TLS version to use for SMTP server communication when TLS is configured in the Email Settings. | • Default<br>• TLS10<br>• TLS11<br>• TLS12<br>• TLS13 |

| | | |
|---|---|---|
| **system.email.notifications.skip.event.list** | List of event names to skip when creating an email digest. Use space or "," or ";" as separators. The event names can be found in events.log. | Event names (example: error60) |
| **system.vmware.esxi.ssh.port** | For VMware only. Specifies the SSH port to connect to ESXi (global setting). | • Default value: 22<br>• Minimum value: 1<br>• Maximum value: 65535 |
| **system.vmware.skip.outdated.tools.checking** | For VMware only. When enabled, the system does not check VMware Tools outdated status when creating quiescing snapshot. | • Unchecked (default)<br>• Checked |
| **system.vmware.skip.tag.discovery** | VMware only. When enabled, the system does not discover VMware Tags. This is applied to all tenants. | • Unchecked (default)<br>• Checked |
| **system.debug.mode.log.vmware.api.incoming.requests** | VMware only. When enabled, the incoming message will be printed for VIJAVA API received response.<br>The option only works if **system.debug.mode.enabled** is checked. | • Unchecked (default)<br>• Checked |
| **system.debug.mode.log.vmware.api.outgoing.requests** | VMware only. When enabled, the outgoing message will be printed for VIJAVA API sent request.<br>The option only works if **system.debug.mode.enabled** is checked. | • Unchecked (default)<br>• Checked |

| | | |
|---|---|---|
| **http.max.upload.size** | Specifies the max upload size for file upload operations, **bytes** (global setting). If multiple files are uploaded, this is the total size. Use -1 for unlimited.<br>Example: 200MB: 200000000 | • Default value: 1073741824<br>• Minimum value: 1<br>• Maximum value: 9999999999 |
| **system.auth.use.lockout** | Enables or disables the login lockout feature. When enabled, the offending IP address is not allowed to login after several failed attempts. | • Unchecked (default)<br>• Checked |
| **system.auth.max.login.attempt.count** | Specifies the maximum number of failed login attempts to trigger the login lockout feature for the offending IP. | • Default value: 5<br>• Minimum value: 1<br>• Maximum value: 9999 |
| **system.auth.lockout.timeout** | Specifies the timeout (minutes) for the login lockout feature. The offending IP is allowed to login again after the timeout expires. | • Default value: 15<br>• Minimum value: 1<br>• Maximum value: 9999 |
| **system.auth.login.history.period** | Specifies the period (minutes) to calculate the maximum number of failed login attempts for the login lockout feature. | • Default value: 5<br>• Minimum value: 1<br>• Maximum value: 9999 |

| | | |
|---|---|---|
| **system.auth.ad.integration.follow.referrals** | Defines LDAP/Active Directory behavior for referrals. When set to follow, all referrals are resolved (can be slow); otherwise they are ignored.<br>What are the implications of the ignore option?<br>* If you only have one domain, there should be no effects.<br>* If you have multiple domains joined in a forest, then any cross-domain memberships will not be resolved.<br>More info: https://docs.oracle.com/javase/jndi/tutorial/ldap/referral/jndi.html | • follow (default)<br>• ignore |
| **system.auth.ad.integration.connect.timeout** | Specifies the timeout (miliseconds) for connecting LDAP/Active Directory. | • Default value: 2000<br>• Minimum value: 2000<br>• Maximum value: 100000 |
| **system.auth.ad.integration.read.timeout** | Specifies the timeout (miliseconds) for reading LDAP/Active Directory operations. | • Default value: 10000<br>• Minimum value: 10000<br>• Maximum value: 100000 |
| **system.auth.max.login.2fa.attempt.count** | Specifies the attempts for the login lockout feature. The offending IP is allowed to login again after the timeout expires. | • Default value: 5<br>• Minimum value: 1<br>• Maximum value: 9999 |

| | | |
|---|---|---|
| **system.auth.lockout.2fa.timeout** | Specifies the timeout (minutes) for the login lockout feature. The offending IP is allowed to login again after the timeout expires. | • Default value: 5<br>• Minimum value: 1<br>• Maximum value: 9999 |
| **system.job.block.size** | Select block size for processing data.<br><br>**Notes**<br>• Deduplication can only be efficient with recovery points using the same block size.<br>• Once the value is changed, the existing backup jobs, previously using a different block size, will produce a full backup on the next run.<br>• Mapping to a backup with a different block size will be skipped | • 4 MB (default)<br>• 2 MB<br>• 1 MB<br>• 512 KB<br>• 256 KB<br>• 128 KB<br>• 64 KB<br>• 32 KB<br>• 16 KB<br>• 8 KB<br>• 4 KB |
| **system.job.map.new.source.item.scope** | The scope to search for the existing backup when adding a new source item to the job. | • Default location (default)<br>• Default transporter's locations<br>• All locations |
| **system.job.pool.queue.length** | Specifies the length of the job queue. A job is placed in a queue before execution. Requires restart. | • Default value: 200<br>• Minimum value: 10<br>• Maximum value: 9999 |

| | | |
|---|---|---|
| **system.job.pool.thread.min** | Specifies the minimum thread pool size for jobs. A job requires 1 thread from the job pool to start running. Requires restart. | • Default value: 30<br>• Minimum value: 10<br>• Maximum value: 9999 |
| **system.job.pool.thread.max** | Specifies the maximum thread pool size for jobs. A job requires 1 thread from the job pool to start running. When the pool thread limit is reached, the job is placed in the job queue. Requires restart.<br>If using Linux and systemd, please add the following to the service startup script: TasksMax=infinity | • Default value: 200<br>• Minimum value: 10<br>• Maximum value: 9999 |
| **system.job.resolve.host.hostname.on.transporter** | If set, sends the source and/or target host hostname as is to Transporter. Transporter will resolve the hostname to the IP address(es) and check if it is reachable. This is done during the Transporter to host checks on a job run.<br>The default behavior is to do the resolution locally and send the IP addresses to Transporter. This can be a problem in complex network topologies (VPN, etc). | • Unchecked (default)<br>• Checked |

| | | |
|---|---|---|
| **system.job.resolve.transporter.hostname.on.transporter** | If set, sends the source and/or target Transporter hostname as is to Transporter. Transporter will resolve the hostname to the IP address(es) and check if it is reachable. This is done during Transporter to Transporter checks on job run.<br>The default behavior is to do the resolution locally, get hostnames for all resolved IP addresses, and then send them to Transporter. This can be a problem in complex network topologies (VPN, etc). | • Unchecked (default)<br>• Checked |
| **system.job.bandwidth.throttling.source** | If set, applies bandwidth throttling for data reading from source. | • Checked (default)<br>• Unchecked |
| **system.job.bandwidth.throttling.target** | If set, applies bandwidth throttling for data writing to target. | • Checked (default)<br>• Unchecked |
| **system.job.bandwidth.throttling.network** | If set, applies bandwidth throttling for data transfer between source and target. | • Checked (default)<br>• Unchecked |
| **system.job.ict.skip.new.disk** | If set, new disks added to the source item will not be added to the job automatically. | • Checked (default)<br>• Unchecked |
| **system.job.replica.vm.suffix** | The default suffix to append to replica VMs.<br>This setting is global and can only be changed inside the master tenant. | Can be between 1 and 20 characters ("-replica" by default) |

| | | |
|---|---|---|
| **system.job.recovered.vm.suffix** | The default suffix to append to recovered/flash-booted VMs.<br>This setting is global and can only be changed inside the master tenant. | Can be between 1 and 20 characters ("-recovered" by default) |
| **system.job.skip.manual.transporter.data.path.validation** | If set, transporter data path validation will be skipped for manually configured transporters. | • Unchecked (default)<br>• Checked |
| **system.metadata.disable.ec2.instance.id.update** | Disables EC2 instance ID detection on product startup. The detection is performed through an HTTP request to the AWS instance metadata service. This process is required for proper product functioning in the AWS cloud. | • Unchecked (default)<br>• Checked |
| **system.task.pool.queue.length** | Specifies the length of the task queue. A task is placed in the queue before execution. Requires restart. | • Default value: 200<br>• Minimum value: 10<br>• Maximum value: 9999 |
| **system.task.pool.thread.min** | Specifies the minimum thread pool size for tasks. A task requires 1 thread from the task pool to start running. Task example: repository refresh, Transporter refresh, support bundle creation. Requires restart. | • Default value: 30<br>• Minimum value: 10<br>• Maximum value: 9999 |
| **system.task.pool.thread.max** | Specifies the maximum thread pool size for tasks. A task requires 1 thread from the task pool to start running. When the pool thread limit is reached, the task is placed in the task queue. Task example: repository refresh, Transporter refresh, support bundle creation. Requires restart. | • Default value: 200<br>• Minimum value: 10<br>• Maximum value: 9999 |

| | | |
|---|---|---|
| **system.repository.min.free.space.byte** | Specifies the minimum free space (bytes) for the repository. If the free space goes below this value, an alarm is generated. | • Default value: 5368709120<br>• Minimum value: 1024<br>• Maximum value: 1099511627776 |
| **system.repository.min.free.space.percent** | Specifies the minimum free space (percent) for the Backup repository. If the free space goes below this value, an alarm is generated. | • Default value: 5<br>• Minimum value: 1<br>• Maximum value: 99 |
| **system.repository.ec2.min.free.space.resize.percent** | In case the free space is less than the set percentage of the total current storage, one minimum chunk will be added to the storage. | • Default value: 10<br>• Minimum value: 1<br>• Maximum value: 100 |
| **system.repository.ec2.max.free.space.resize.percent** | In case the free space is more than the set percentage of the total current storage, one minimum chunk will be removed from the storage. | • Default value: 15<br>• Minimum value: 1<br>• Maximum value: 100 |
| **system.repository.maintenance.interrupt.timeout.seconds** | Specifies the timeout (seconds) to wait for repository maintenance stop during job run. | • Default value: 300<br>• Minimum value: 1<br>• Maximum value: 86400 |

| | | |
|---|---|---|
| **system.repository.refresh.backup.size.calculation** | Specifies the backup size calculation on the repository refreshing. True: Always calculates backup size. False: Skips backup size calculation and only calculates backup size with necessary backups. | • Checked (default)<br>• Unchecked |
| **system.repository.refresh.timeout.seconds** | Specifies the timeout (seconds) to wait for repository refresh. | • Default value: 600<br>• Minimum value: 1<br>• Maximum value: 86400 |
| **system.repository.remove.backups.used.by.job** | The setting allows to remove backup objects associated with existing jobs, and remove the last RP of a backup object in case such RP is due to be removed according to the retention policy. If enabled, removal of the aforesaid objects can be done manually or automatically, in accordance with the configured retention policy. | • Unchecked (default)<br>• Checked |
| **system.product.skip.update.server.ssl.certificate.verification** | The product update check process requires the remote server certificate to be trusted. This parameter disables such check. It can be useful when secure (SSL/TLS) connections are being intercepted by third-party software.A product restart is required to apply. | • Unchecked (default)<br>• Checked |

| | | |
|---|---|---|
| **system.debug.mode.enabled** | The debug mode prints more information into the logs, including some sensitive one (hardware UUIDs, MAC addresses, etc). The passwords are not printed unless they are present in raw communication dumps (e.g., SOAP/XML/JSON). | • Unchecked (default)<br>• Checked |
| **system.debug.mode.log.passwords** | When debug mode is enabled, also log passwords. This can be a security risk. | • Unchecked (default)<br>• Checked |
| **system.debug.mode.log.api.requests** | When debug mode is enabled, also log product API requests/responses. The data is logged as is and will contain plaintext passwords. This can be a security risk. | • Unchecked (default)<br>• Checked |
| **system.hyperv.optimize.queries** | Hyper-V only. Instructs to use a faster query method to read VM and host information. This will speed up the refresh process in large environments. | • Checked (default)<br>• Unchecked |
| **system.hyperv.discovery.host.thread.count** | Hyper-V only. Sets the max parallel threads to run when refreshing cluster hosts during discovery. Each cluster host can be refreshed separately. This will speed up the refresh process in large environments. | • Default value: 2<br>• Minimum value: 1<br>• Maximum value: 20 |
| **system.hyperv.discovery.vm.thread.count** | Hyper-V only. Sets the max parallel threads to run when refreshing host VMs during discovery. When increasing the setting value, make sure to test its impact on host CPU usage during refresh. This will speed up the refresh process in large environments. | • Default value: 2<br>• Minimum value: 1<br>• Maximum value: 10 |

| | | |
|---|---|---|
| **system.database.scheduled.backup.path** | Specifies the target path for database backups. The tenant databases will be stored in subfolders, if present. The path can be local or absolute. The folder will be created automatically if it does not exist. | |
| **system.database.scheduled.backup.max.count** | Specifies the maximum number of files for periodic database backups. The number is applied separately to each tenant database. The master and tenants product databases are backed up each day. | • Default value: 5<br>• Minimum value: 0<br>• Maximum value: 365 |
| **system.logging.max.index** | Specifies the maximum index of log files. This works globally for all log files. Set 0 to use default value (configured in log4j.xml). | • Default value: 0<br>• Minimum value: 0<br>• Maximum value: 999 |
| **system.product.min.free.space.byte** | Specifies the minimum free space (bytes) for the product installation folder. If the free space goes below this value, an alarm is generated. | • Default value: 2147483648<br>• Minimum value: 10485760<br>• Maximum value: 10737418240 |
| **system.product.free.memory.threshold** | Specifies the minimum ratio for JVM free memory. If the free JVM memory goes below this value, an alarm is generated. | • Default value: 0.1<br>• Minimum value: 0.01<br>• Maximum value: 0.9 |

| | | |
|---|---|---|
| **system.nutanix.discovery.vm.thread.count** | Nutanix AHV only. Sets the max parallel threads to run when refreshing host VMs during discovery. When increasing the setting value, make sure to test its impact on host CPU usage during refresh. This will speed up the refresh process in large environments. | • Default value: 2<br>• Minimum value: 1<br>• Maximum value: 10 |
| **system.aws.discovery.region.thread.count** | AWS only. Sets the max parallel threads to run when refreshing the AWS Regions during discovery. When increasing the setting value, make sure to test its influence on host CPU usage during refresh. This will speed up the refresh process in large environments. | • Default value: 4<br>• Minimum value: 1<br>• Maximum value: 10 |
| **system.aws.discovery.other.thread.count** | AWS only. Sets the max parallel threads to run when refreshing other AWS entities inside the Region during discovery. When increasing the setting value, make sure to test its influence on host CPU usage during refresh. This will speed up the refresh process in large environments. | • Default value: 4<br>• Minimum value: 1<br>• Maximum value: 10 |
| **system.plugin.flr.operation.timeout.seconds** | Specifies the timeout (seconds) to wait for plugin session FLR/OLR. This is a low-level setting that is only sent to Transporter and used during iSCSI interaction. | • Default value: 900<br>• Minimum value: 1<br>• Maximum value: 86400 |
| **system.physical.skip.os.checking** | Physical Windows host discovery only. When enabled, the system will not check the supported OS version. | • Checked<br>• Unchecked (default) |

| | | |
|---|---|---|
| **system.transporter.agent.injection.skip.vc.redist** | When enabled, the system will not automatically install VC redistributable during Transporter/agent injection. | • Checked<br>• Unchecked (default) |
| **system.transporter.load.max.time.created.state.hours** | Specifies the timeout (hours) to wait for getting Transporter load request. Default is 5 hours. | • Default value: 5<br>• Minimum value: 1<br>• Maximum value: 72 |
| **system.transporter.modern.min.heap.size.megabyte** | Megabytes.<br>The -Xms option sets the initial and minimum Java heap size. The Java heap (the “heap”) is the part of the memory where blocks of memory are allocated to objects and freed during garbage collection.<br>**Note**: Transporter restart is required to apply the setting. | • Default value: 512<br>• Minimum value: 256<br>• Maximum value: 65536 |
| **system.transporter.modern.max.heap.size.megabyte** | Megabytes.<br>This option sets the maximum Java heap size. The Java heap (the “heap”) is the part of the memory where blocks of memory are allocated to objects and freed during garbage collection. Depending on the kind of operating system you are running, the maximum value you can set for the Java heap can vary.<br>**Notes**:<br>-Xmx does not limit the total amount of memory that the JVM can use.<br>Transporter restart is required to apply the setting. | • Default value: 3072<br>• Minimum value: 256<br>• Maximum value: 65536 |

| | | |
|---|---|---|
| **system.transporter.modern.thread.stack.size.kilobyte** | Kilobytes.<br>-Xss sets the thread stack size. Thread stacks are memory areas allocated for each Java thread for their internal use. This is where the thread stores its local execution state.<br>**Note**: Transporter restart is required to apply the setting. | • Default value: 512<br>• Minimum value: 64<br>• Maximum value: 2048 |
| **system.transporter.modern.job.handler.max.thread.count** | Specifies the job thread count for modern Transporter.<br>**Notes:**<br>• 1 job thread equals ~200MB of memory, consider changing the related setting.<br>• Transporter restart is required to apply the setting. | • Default value: 10<br>• Minimum value: 1<br>• Maximum value: 128 |
| **system.transporter.modern.service.handler.max.thread.count** | Specifies the service thread count for modern Transporter.<br>**Note**: Transporter restart is required to apply the setting. | • Default value: 10<br>• Minimum value: 1<br>• Maximum value: 128 |
| **system.transporter.jvm.ram.requirement** | Bytes.<br>For NASes only.<br>Specifies the minimal ram required on NASes to create a SaaS repository. | • Default value: 4294967296<br>• Minimum value: 0<br>• Maximum value: 10995 11627776 |

| | | |
|---|---|---|
| **system.transporter.modern.thread.pool.size** | Specifies the session factory thread pool size for modern Transporter.<br>**Note**: Transporter restart is required to apply the setting. | • Default value: 1000<br>• Minimum value: 100<br>• Maximum value: 1000 |
| **system.deleted.users.groups.remove.frequency** | Specifies the scheduled time for removing unnecessary deleted users, groups (in second). | • Default value: 86400<br>• Minimum value: 300<br>• Maximum value: 1.7976931348623157e+308 |
| **system.inventory.allow.duplicated** | Microsoft 365 and physical machines only. When enabled, the system allows duplicated discovery items. | • Unchecked (default)<br>• Checked |
| **system.inventory.optimize.discovery.time** | Microsoft 365 (SharePoint Online) only. When enabled, the system skips some attributes to optimize the discovery time. | • Unchecked (default)<br>• Checked |
| **system.o365.suppress.throttling.event** | Suppress throttling warning. | • Unchecked (default)<br>• Checked |
| **system.event.skip.creating.event.list** | List of event/alarm/notification names to skip when creating an event. The event is still logged and handled. Use space or , or ; as separators. The names can be found in events.log. | Event names (example: error60) |

| | | |
|---|---|---|
| **system.events.use.windows.event.integration** | Use Windows Event log integration. Some product events will also be created in the Application log.<br>This setting is global and can only be changed inside the master tenant. | • Unchecked (default)<br>• Checked |
| **system.exchange.enable.direct.recovery** | When enabled, you can recover Exchange items without using a recovery server. For example, you can download items to the browser or forward them to a certain email. To do this, select **Download items or forward via email** on the **Destination** page of the job wizard and then select the appropriate recovery type on the **Options** page.<br>Note that Google limits the total size of attachments within a message to 25 MB. Forwarding messages containing attachments that exceed this limit will fail. | • Unchecked (default)<br>• Checked |
| **system.olr.dsamain.mount.port** | TCP port where DSAMAIN mounts NTDS.dit (AD database) for. | • Default value: 5000<br>• Minimum value: 1<br>• Maximum value: 65535 |
| **system.product.register.disable.periodic.data.collection** | When enabled, the product will not send data bundles every 30 days. | • Unchecked (default)<br>• Checked |
| **system.repository.skip.periodic.refresh.on.transporter.busy.with.job** | When enabled and any Transporter repository is locked by a running job, the product skips periodic refresh for this Transporter repository. | • Unchecked (default)<br>• Checked |

| | | |
|---|---|---|
| **system.pql.custom.file.name** | PQL file name in the userdata folder. Empty by default. If empty, the file will be downloaded from web. Example: pql.json. | • Default value: empty<br>• Minimum value: 1<br>• Maximum value: 2048 |
| **system.pql.cache.ttl.hours** | Time to keep PQL file cache, in hours. Use 0 to disable the cache. | • Default value: 8<br>• Minimum value: 0<br>• Maximum value: 72 |
| **system.transporter.allow.new** | Allows using newer Transporter versions. | • Unchecked (default)<br>• Checked |
| **system.transporter.allow.old** | Allows using older (outdated) Transporter versions. | • Unchecked (default)<br>• Checked |
| **system.transporter.modern.idle.timeout** | Specifies the timeout (milliseconds) for modern Transporter IDLE. If you set it to 0, it will be an unlimited timeout, meaning the transporter can only be stopped manually.<br>**Note**: Transporter restart is required to apply the setting. | • Default value: 3600000<br>• Minimum value: 0<br>• Maximum value: 86400000 |
| **system.volatile.object.processing.type** | **Default**: try to remove the volatile objects periodically until their time to live (fixed) is reached. **Alternative**: fine-tune the settings. See the other system.volatile.object variables. The setting is global and can be changed inside the master tenant only. | • Default (default option)<br>• Alternative |

| | | |
|---|---|---|
| **system.volatile.object.retry.count** | Alternative processing type only. The maximum number of retries for volatile objects removal. 0 means no retries, so only one removal attempt will happen. The setting is global and can be changed inside the master tenant only. | • Default value: 7<br>• Minimum value: 0<br>• Maximum value: 256 |
| **system.volatile.object.retry.interval** | Alternative processing type only. Minutes. The desired delay between each removal retry. The real delay depends on the queue and on the exponential retry factor (configurable). The setting is global and can be changed inside the master tenant only. | • Default value: 60<br>• Minimum value: 5<br>• Maximum value: 14400 |
| **system.volatile.object.exponential.retry.interval.factor** | Alternative processing type only. The ratio to use when calculating the delay time for the next retry. The next delay equals interval * (factor^retry). Example: the interval is 60 minutes, the factor is 2. The first retry will happen in +60 minutes, the second in +240 minutes , ... . The setting is global and can be changed inside the master tenant only. | • Default value: 2<br>• Minimum value: 1<br>• Maximum value: 10 |
| **system.visual.notification.service.disable** | Disables the visual notification service. This can speed up the UI when the database contains many event entries. This setting is global and can be changed inside the master tenant only. | • Unchecked (default)<br>• Checked |
| **system.msp.console.listening.port** | TCP port used by the MSP product for listening to remote tenants. Port 6702 is used by default. | • Default: 6702<br>• Minimum value: 1<br>• Maximum value: 65535 |

| | | |
|---|---|---|
| **system.events.use.windows.event.integration** | Use Windows Event log integration. Some product events will also be created in the Application log.<br>The setting is global and can be changed inside the master tenant only. | • Unchecked (default)<br>• Checked |
| **system.transporter.load.path.cost.variation.percent** | Percent. Specifies the allowed data path cost variation.<br>During the job run, automatic transporter selection may happen. The first step is to choose the top N (by cost in milliseconds) data paths. The second step is to choose the best data path based on the lowest transporter load.<br>For example, 20 paths were found based on round trip time between source and target host.<br>If the setting is 10%, the best path cost is 2 (N=2), then only paths with costs in range 2 ... 2.2 (2+0.2) will be chosen on the first step.<br>The setting is global and can be changed inside the master tenant only. | • Default value: 10<br>• Minimum value: 1<br>• Maximum value: 10000 |
| **system.vmware.discovery.vm.detect.ipaddress.by.dns.skip** | The setting is applicable only to VMware vSphere infrastructure.<br>If enabled, the detection of VM IP address via DNS resolution will be skipped.<br>Note: Detection of VM IP address via DNS resolution is applied in case VMware Tools are not installed on the VM. | • Checked (default)<br>• Unchecked |

| | | |
|---|---|---|
| **system.job.default.retention.approach** | Schedule-retention fusion: New backup and backup copy jobs will use the new schedule-retention step.<br>Legacy: New backup and backup copy jobs will use the legacy schedule and retention steps. | • Schedule-retention fusion<br>• Legacy |
| **system.job.run.skip.infrastructure.refresh** | If set, infrastructure refresh will be skipped during job run. | • Unchecked (default)<br>• Checked |
| **system.job.ict.skip.new.source.item** | If set, new source items will not be added to the job automatically. | • Unchecked (default)<br>• Checked |
| **system.transporter.min.connect.timeout.seconds** | Specifies the minimum connect timeout (seconds) for transporter connection. Set 0 to use default value in the product (90 seconds). | • Default value: 0<br>• Minimum value: 0<br>• Maximum value: 86400 |
| **system.transporter.min.write.timeout.seconds** | Specifies the minimum write timeout (seconds) for transporter command. Set 0 to use default value in the product (60 seconds). | • Default value: 0<br>• Minimum value: 0<br>• Maximum value: 86400 |
| **system.transporter.min.idle.timeout.seconds** | Specifies the minimum read idle timeout (seconds) for transporter command. Set 0 to use default value in the product (300 seconds). | • Default value: 0<br>• Minimum value: 0<br>• Maximum value: 86400 |

| | | |
|---|---|---|
| **system.events.use.syslog.integration** | Use Syslog log integration. The product events will be copied to the remote log. The setting is global and can be changed inside the master tenant only | • Unchecked (default)<br>• Checked |
| **system.events.syslog.host** | Specifies the Syslog server hostname. The setting is global and can be changed inside the master tenant only. | • IP/hostname field (empty by default) |
| **system.events.syslog.port** | Specifies the Syslog server port. syslog-udp usually uses port 514, syslog-tls usually uses port 6514. The setting is global and can be changed inside the master tenant only. | • Default value: 514<br>• Minimum value: 1<br>• Maximum value: 65535 |
| **system.events.syslog.transport** | Specifies the transport for sending Syslog messages. The setting is global and can be changed inside the master tenant only. | • UDP (default)<br>• TCP<br>• TLS |
| **system.events.syslog.connect.timeout** | Milliseconds. Specifies the connect timeout when connecting to the Syslog server, TCP/TLS case. The setting is global and can be changed inside the master tenant only. | • Default value: 1000<br>• Minimum value: 10<br>• Maximum value: 180000 |

| | | |
|---|---|---|
| **system.events.syslog.message.facility** | Specifies the facility to be used in Syslog messages. Facility code Keyword Description<br>0 kern Kernel messages<br>1 user User-level messages<br>2 mail Mail system<br>3 daemon System daemons<br>4 auth Security/authentication messages<br>5 syslog Messages generated internally by syslogd<br>6 lpr Line printer subsystem<br>7 news Network news subsystem<br>8 uucp UUCP subsystem<br>9 cron Cron subsystem<br>10 authpriv Security/authentication messages<br>11 ftp FTP daemon<br>12 ntp NTP subsystem<br>13 security Log audit<br>14 console Log alert<br>15 solaris-cron Scheduling daemon<br>16–23 local0 – local7 Locally used facilities<br>The setting is global and can be changed inside the master tenant only. | • Default value: 1<br>• Minimum value: 0<br>• Maximum value: 23 |
| **system.events.syslog.message.app.name** | Specifies a custom source application name to be used in Syslog messages. Leave blank for default. The setting is global and can be changed inside the master tenant only. | • Empty (default) |

| | | |
|---|---|---|
| **system.events.syslog.message.hostname** | Specifies a custom source host name to be used in Syslog messages. Some Syslog cloud services may use this field to transmit a secret key. Leave blank for default. The setting is global and can be changed inside the master tenant only. | • Empty (default) |
| **system.events.syslog.message.format** | Specifies the message format to be used in Syslog messages. RFC 5424 has more fields and is used by rsyslog by default. The setting is global and can be changed inside the master tenant only. | • RFC_3164<br>• RFC_5424 (default)<br>• RFC_5424 |
| **system.events.syslog.message.style** | Specifies the message style to be used for Syslog messages. The setting is global and can be changed inside the master tenant only. | • TEXT<br>• JSON (default) |
| **system.inventory.o365.skip.throttled.items** | Microsoft 365 only. When enabled, the system will skip throttled items. | • Unchecked (default)<br>• Checked |
| **system.job.run.test.treat.immutable.days.as.minutes** | Test only. Treat 'Immutable for X days' job schedule option as 'Immutable for X minutes'. The setting is global and can be changed inside the master tenant only. | • Unchecked (default)<br>• Checked |
| **system.job.run.test.treat.keep.backup.days.as.minutes** | Test only. Treat 'Keep backups for X days' job schedule option as 'Keep backups for X minutes'. The setting is global and can be changed inside the master tenant only. | • Unchecked (default)<br>• Checked |
| **system.debug.config** | Predefined debug logging configuration. Default values will be used if input invalid JSON. Set to empty to reset default values. Should be edited by Dev only to debug specific feature. | • Format value : JSON |

## Configuring Actions View

Click the **Actions** tab to configure the following actions:

- **Remove all events**: By clicking the link, you can remove all events/alarms/etc for the current tenant.
- **Forget all passwords (except users)**: By clicking the link, you can set the stored passwords to “” for the current tenant items. The only exception is user passwords; they must be set manually.
- **Clean up job history**: By clicking the link, you can immediately apply the configured **Store job history for the last** setting.

In the text box, you can see the report on the actions.

**Example 1**

Request 1: sending (Remove all events)...

Request 1: success=true (Remove all events).

**Example 2**

Request 1: sending (Forget all passwords (except users))...

Request 1: success=true (Forget all passwords (except users)).

**Example 3**

Request 1: sending (Clean up job history)...

Request 1: success=true (Clean up job history).

## Packages

By clicking the **Packages** tab, you can see the following information:

- **Base local path:** packages. Location of packages in product installation directory
- List of **Existing packages**
- List of **Supported packages**

# Maintenance Mode

Your NAKIVO Backup & Replication instance can enter the Maintenance mode in the following cases:

- The product is updating
- External database migration is in progress
- The Director is being rebooted
- The Director is running out of space.

After the product enters the Maintenance mode, all the activities are stopped, and the maintenance page is displayed with the corresponding information about the maintenance activity. Below is an example of a maintenance page for system update.




**Note**

The informational text displayed on the maintenance page also depends on the following:

- Maintenance activities
- The instance deployment type (for example, if you are using a standalone solution or have been added as a tenant to the multi-tenant solution of NAKIVO Backup & Replication)
- The User type (Master Admin, local or remote tenant, etc.)

# Director is Running out of Space

The product checks the free disk space every 30 seconds.

If less than 1 GB of free disk space is available:

- The instance user is disabled.
- NAKIVO Backup & Replication enters Maintenance mode.
- All activities stop, and the maintenance page is displayed.

To check the current free space, click **Check Now**.




If more than 2 GB of free disk space is available:

- The product exits Maintenance mode and resumes normal operation.
- All Backup Repositories and Tape Devices detached during maintenance mode are reattached.
- All Jobs that were disabled when entering maintenance mode are enabled.

# Virtual Appliance Configuration

This section covers the following topics:


- “Virtual Appliance Interface” below
- “Configuring Network Settings of Virtual Appliance” on page 811
- “Increasing Backup Repository Size on Virtual Appliance” on page 812
- “Removing Disk with Backup Repository from Virtual Appliance” on page 813


## Virtual Appliance Interface

The NAKIVO Backup & Replication Virtual Appliance Interface allows users to configure certain features of the NAKIVO Backup & Replication VM Appliances run in text mode.

It is displayed automatically when the VM console is launched.

It does not require installing any additional software.

**Notes**

- The VA Interface can also be made available by logging in remotely (for example, via SSH).
- If needed, the VA Appliance Interface can be installed or updated/disabled.
- Refer to Updating Virtual Appliance to learn more about how you can update your virtual appliance (VA).

The NAKIVO Backup & Replication Virtual Appliance Interface allows the user to configure the below features of the VM Appliance. Refer to the following topics for more information:


- “Network Settings” on page 791
- “Security Settings” on page 795
- “Time and Time Zone” on page 798
- “System Performance” on page 801
- “Manage NAKIVO Backup & Replication Services” on page 803
- “Exit to System Console” on page 809


### Screen Modes and Navigation

All screens of the NAKIVO Backup & Replication Virtual Appliance Interface are available in three basic modes and follow similar navigation rules:

## List mode

In this mode, a list of available items is displayed with the first list item highlighted by default. You can navigate between list items and choose the needed item by using the *<Up/Down>* arrow keys.
To open the screen that corresponds to the highlighted item, press *<Enter>*.
Pressing the *<Esc>* key opens the higher-level screen (if any).
Pressing the *<F5>* key refreshes the list.
Press the *<F10>* key to save and exit.

## View mode

In this screen mode, particular VA information is displayed.
Pressing the appropriate key (screen specified) performs the required action.
Pressing the *<Esc>* key opens the higher-level screen (if any).

## Wizard mode

In this mode, the user is guided by text instructions through action steps. You can enter data or choose options (screen specified).
Pressing the *<Esc>* key will open the confirmation screen. If the action is confirmed, the higher-level screen is opened.

# Main Screen

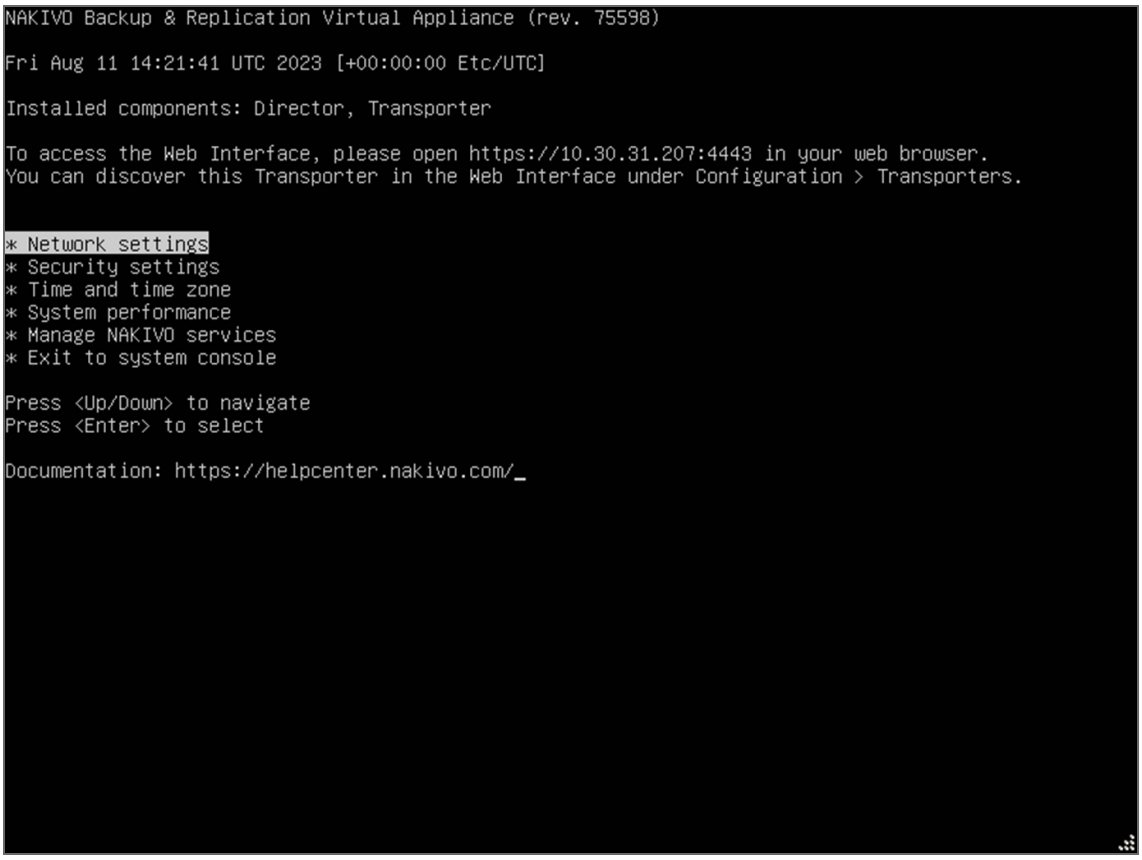
The main screen opens in the **List** mode displaying the information on the NAKIVO Backup & Replication VM Appliance interface revision version, its build number and date.

> **Important**
> The **NAKIVO Backup & Replication command console** option is displayed only if the **Director** component is present.

The installed components are listed below: "Director", "Transporter", or "Director, Transporter" (depending on the installed components).

> **Note**
> If the machine has no configured IP addresses, the following message is displayed:
> *Networking is not configured. Please open the Network settings and configure the networking.*

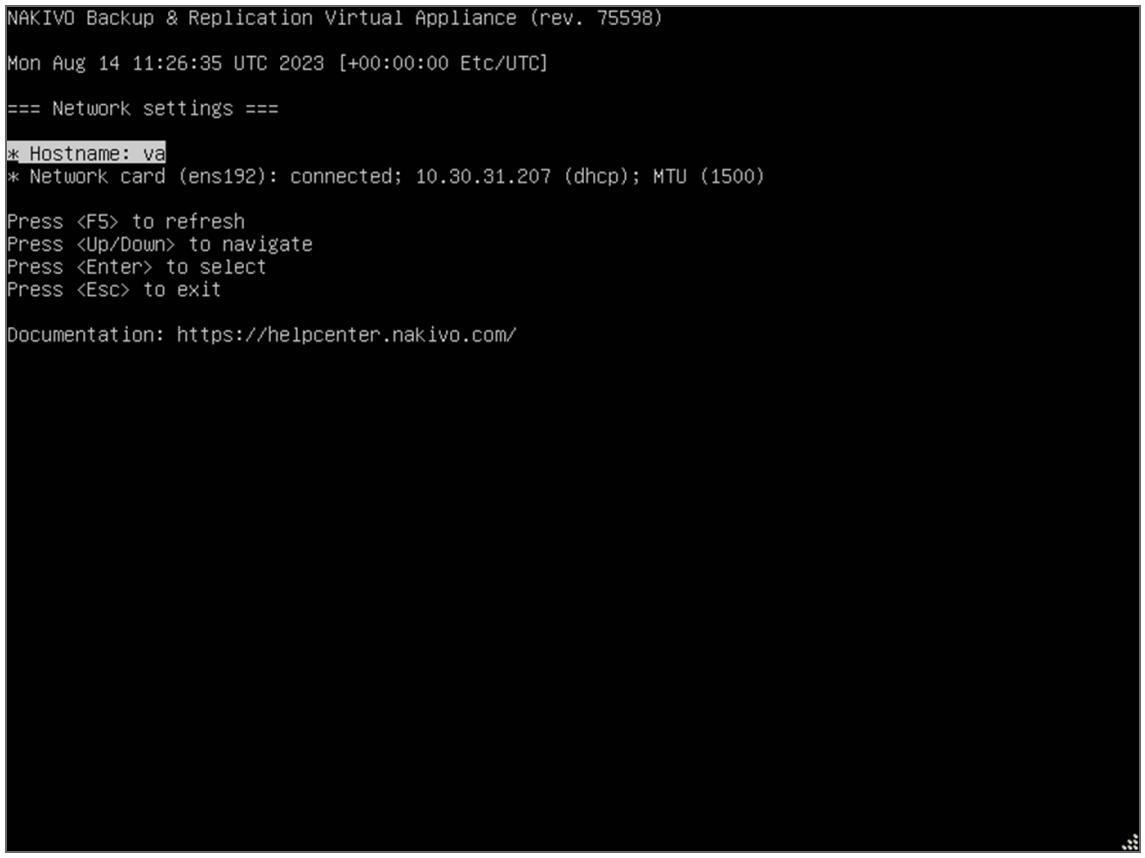
If the machine has at least one IP address configured and the **Director** component is present, the NAKIVO Backup & Replication command console option becomes available and the following message is displayed:
*To access the Web Interface, please open http://<current_ip> in your web browser.*
If there are multiple IP addresses, the first configured IP address is displayed.




The main screen is a starting point from where access to the VM Appliance configuration settings is granted. By default, the first list item, **Network settings**, is highlighted but you are free to navigate the list and select any other option to proceed with configuring your VM.

## Network Settings

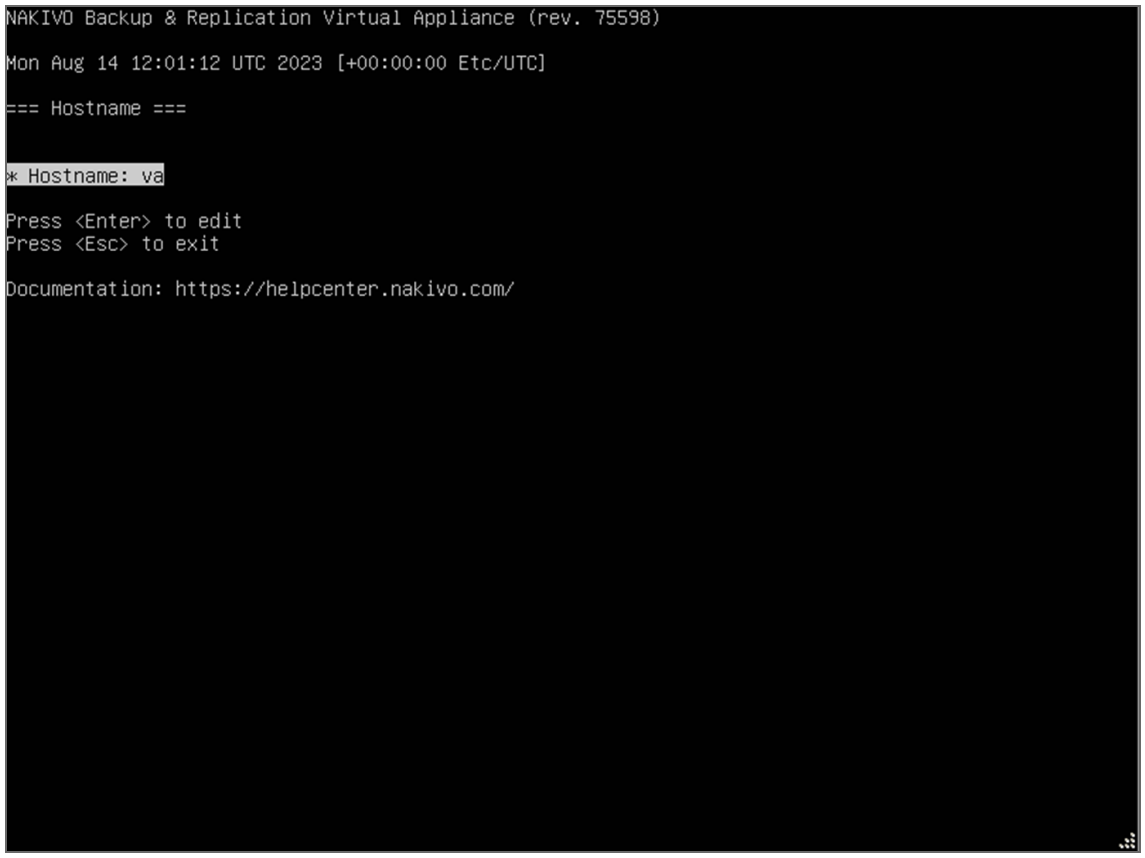
In this screen menu, you can select between two options, Hostname and Network card.




## Hostname

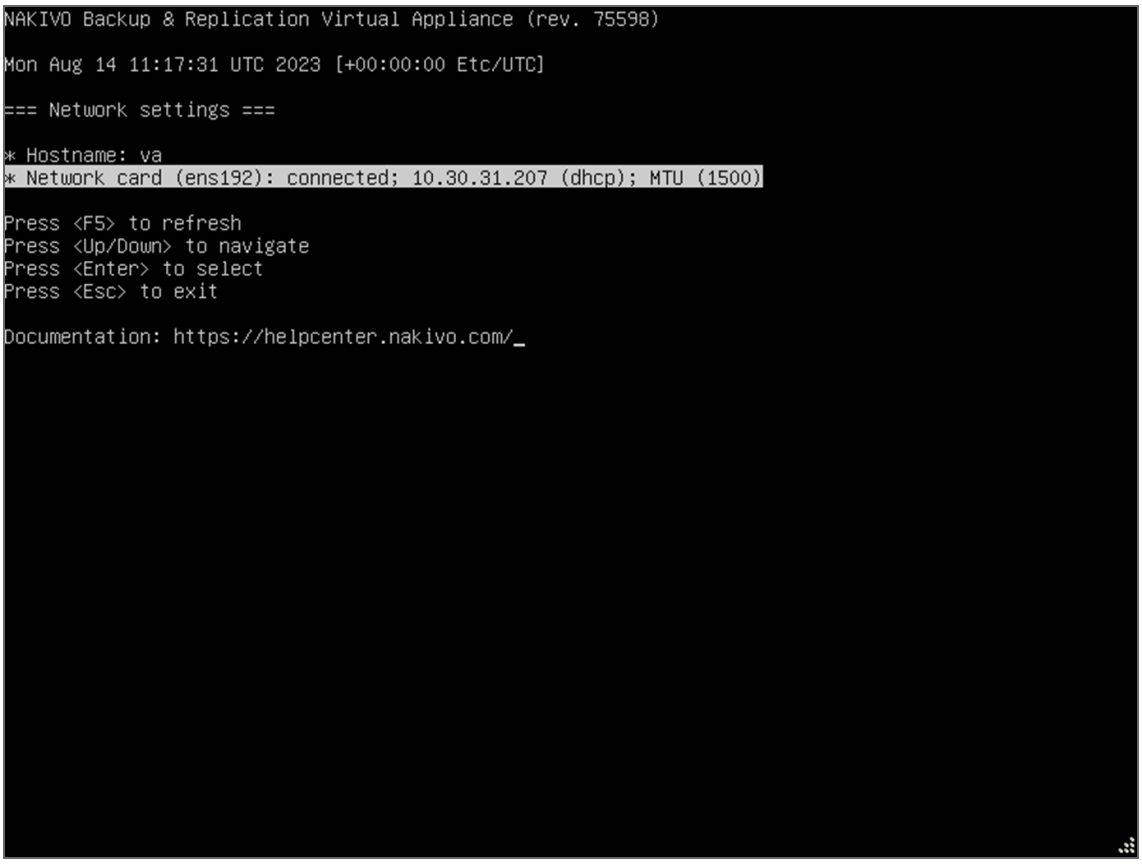
Select to view or edit the VM Appliance hostname and domain.

To edit the VM Appliance hostname, select it by pressing *<Enter>* and then follow the instructions on the screen.




## Network card

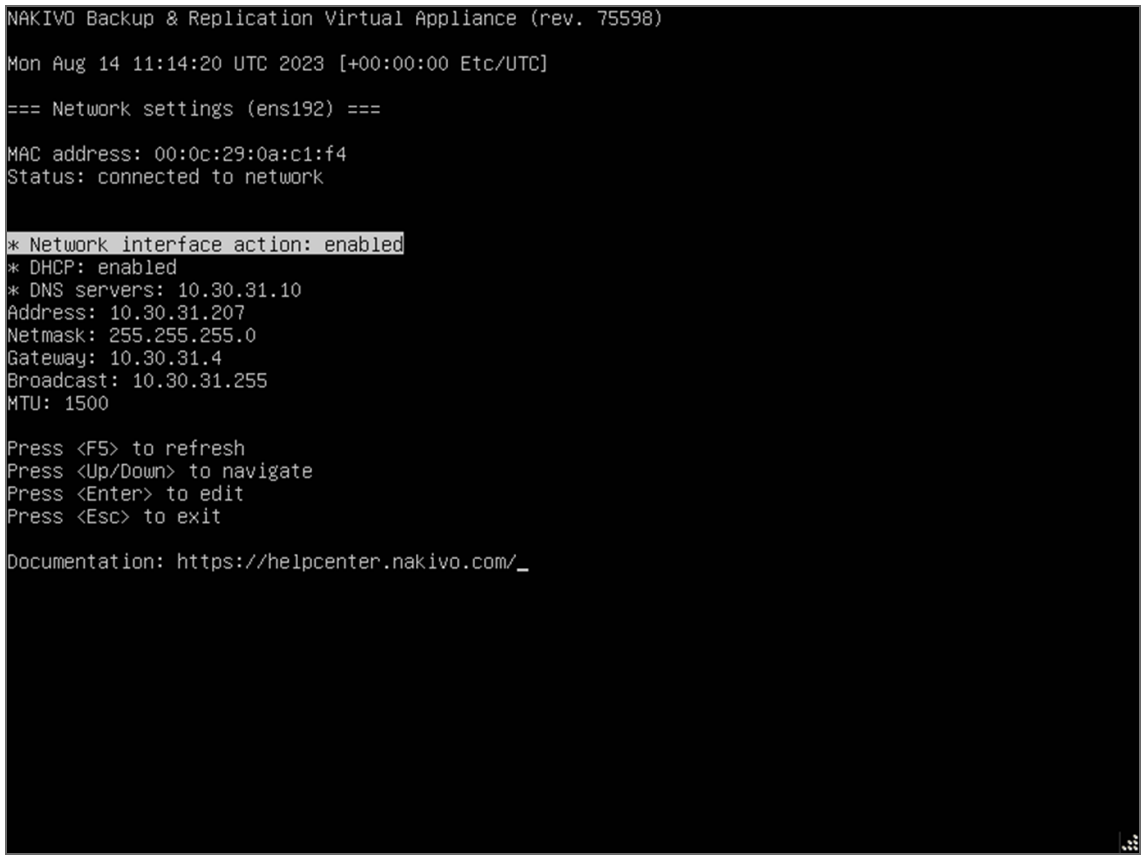
Select to view or edit the settings of the network adapters of the VM Appliance.




To open the network card menu data, select the needed item by pressing <*Enter*>.

Three options become available for editing:

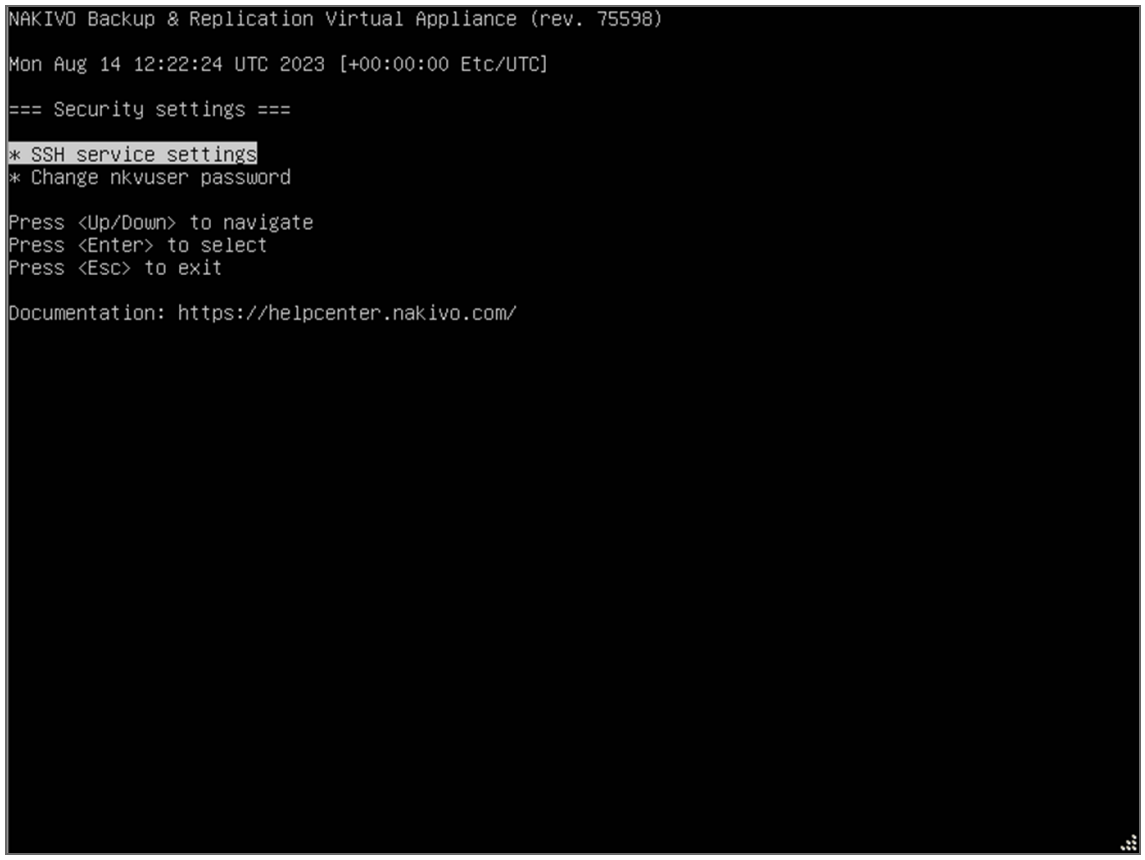
- **Network interface actions** (*enabled* or *disabled*).
- **DHCP** (*enabled* or *disabled*).
- **DNS servers**. Multiple DNS servers are supported but only valid values are accepted after entering and submitting.




To edit the VM Appliance network card settings, select it by pressing <Enter> and then follow the instructions on the screen.

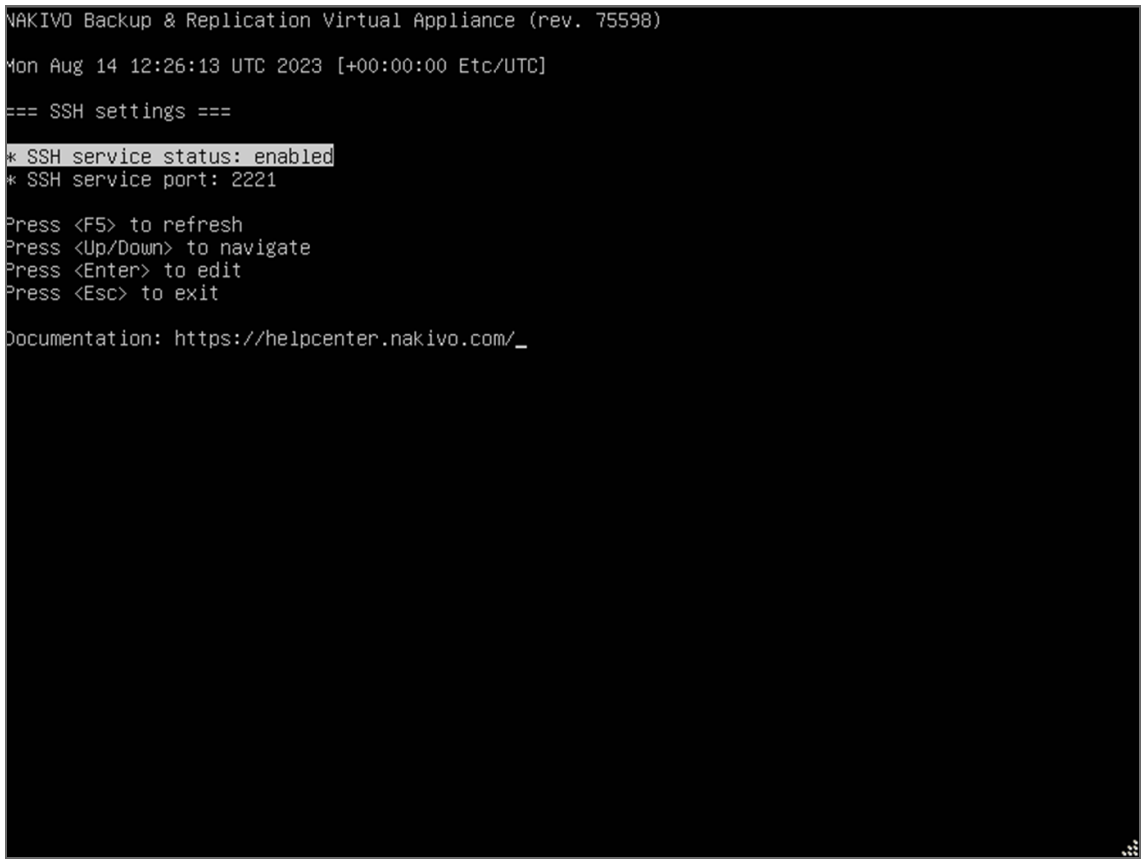
## Security Settings

In this screen menu, you can select between two options, SSH service settings and Change [username] password.




## SSH service settings

Select to view or edit the SSH service status and SSH service port by following the instructions on the screen.

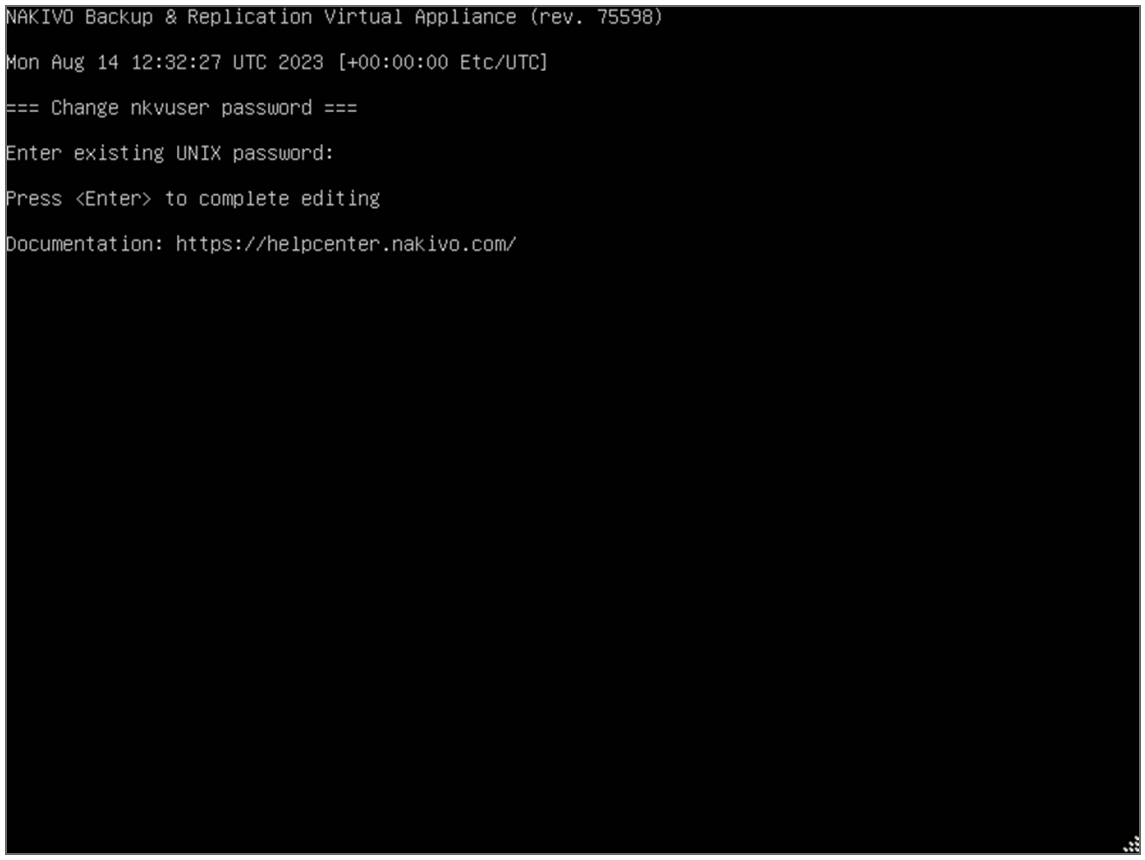



## Change [username] password

Select to view or change the [username] password.

To apply the password change, entering the existing UNIX password is required.

If access to password change is authorized, enter the new password, repeat it, and then press *<Enter>* to confirm.

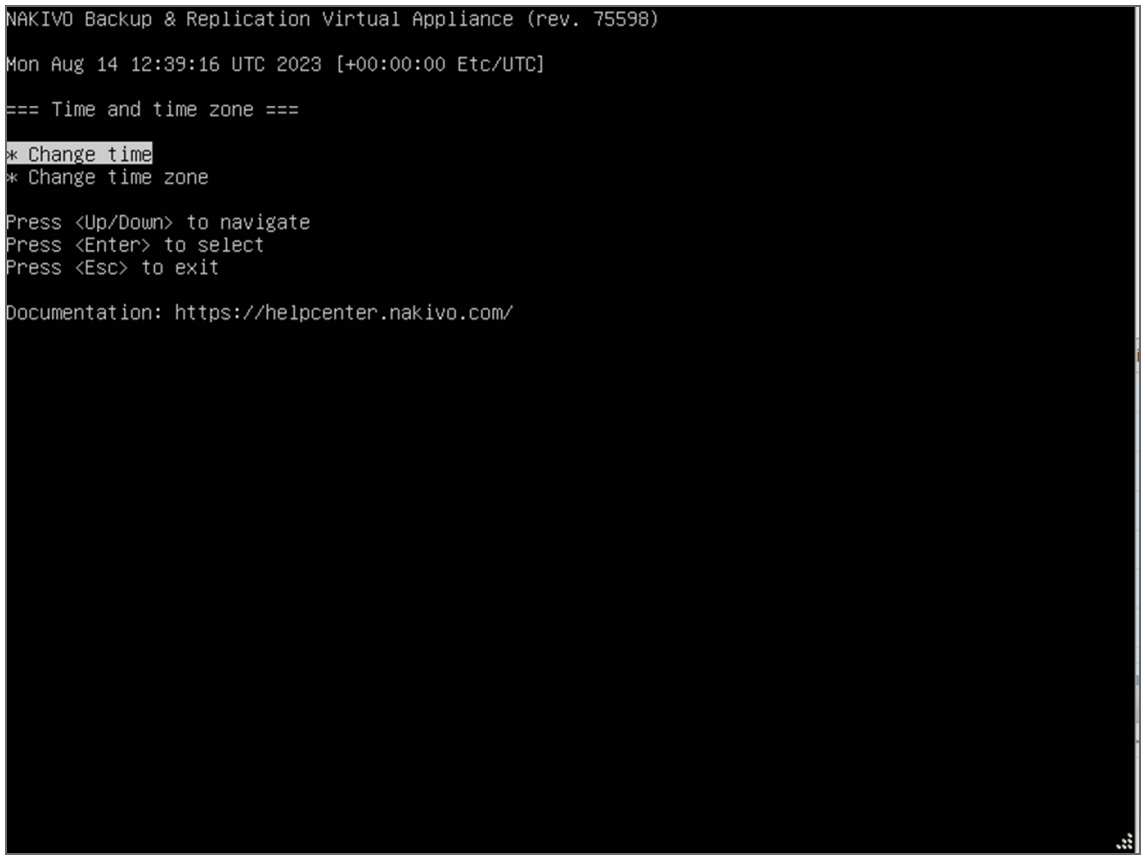



**Note**

To exit this screen, press *<Esc>*.

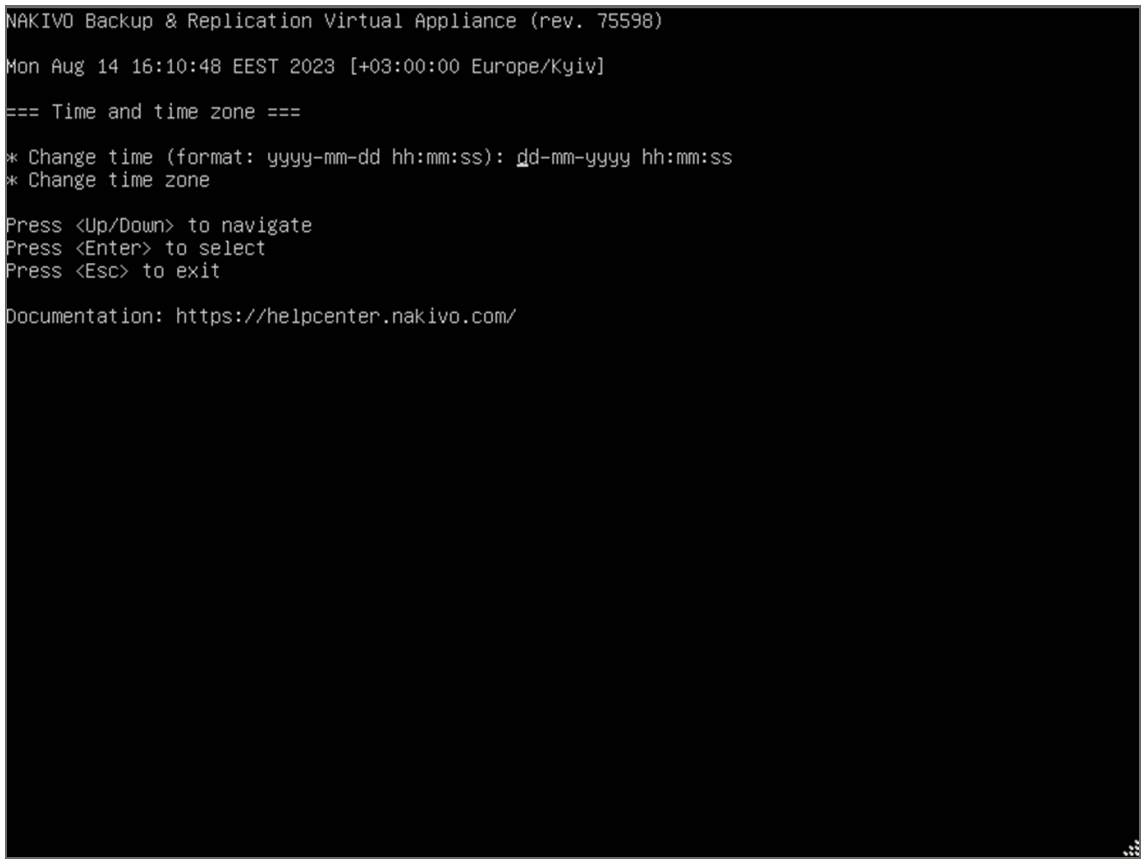
## Time and Time Zone

In this screen menu, you can select between two options, Change time and Change time zone.



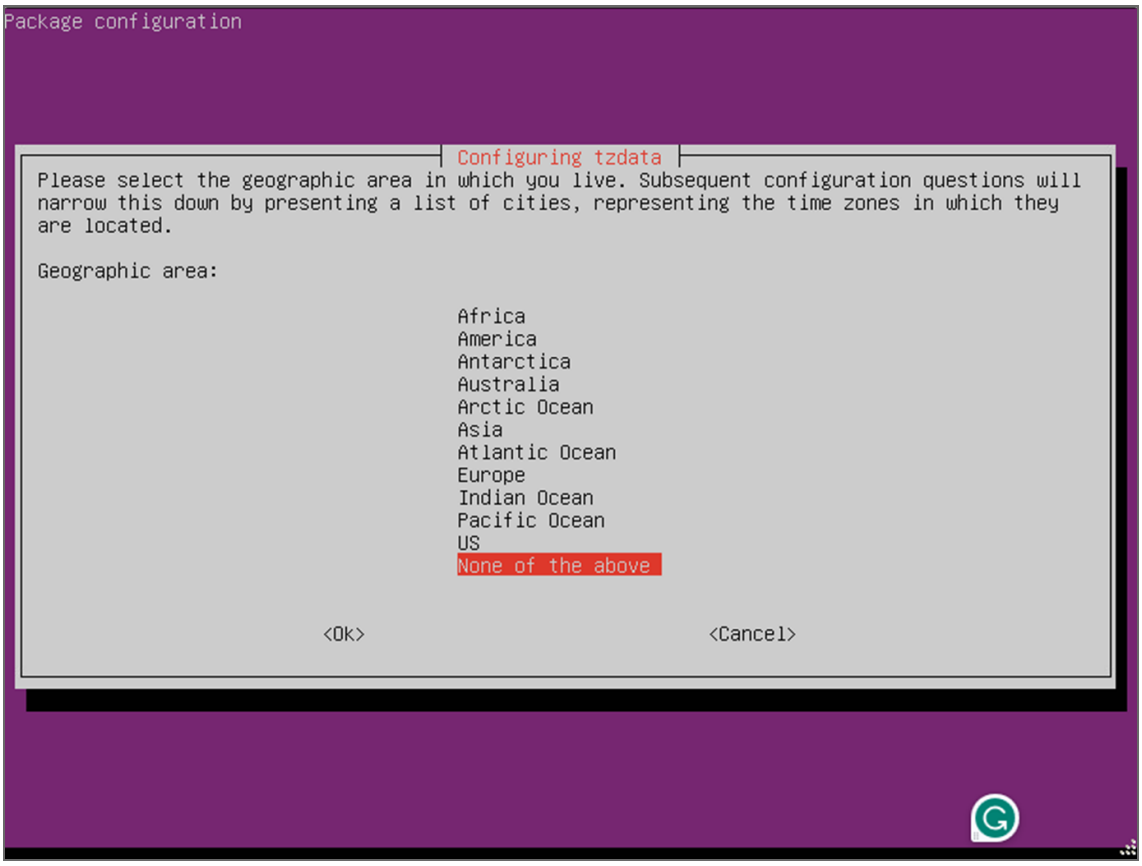

## Change time

Select to view or change the VA time by following the instructions on the screen.




## Change time zone

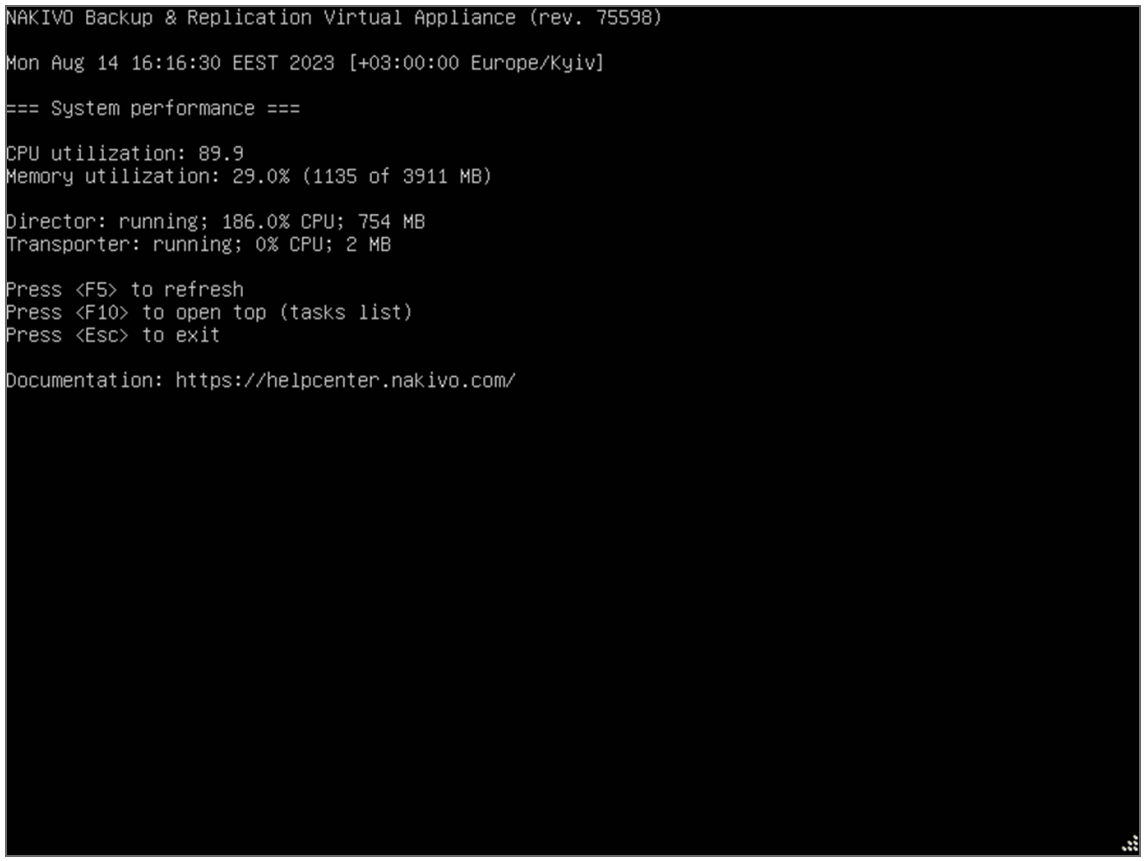
Select to view or change the VA time zone configuration by following the instructions on the screen.




# System Performance

In this screen, you can view the current system consumption and the consumption of the following system resources:

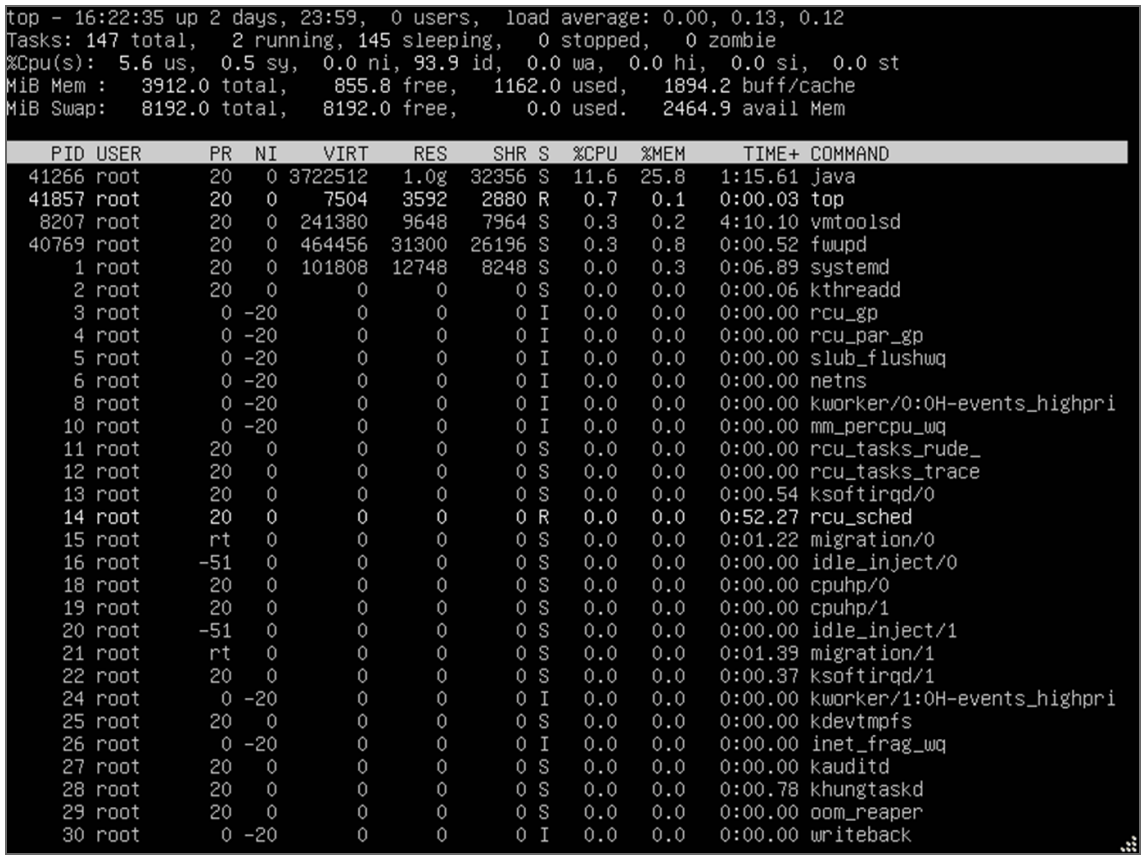
- CPU utilization by **Director** and **Transporter**
- Memory utilization by **Director** and **Transporter**




**Notes**

- The "Director" item is displayed only if the **Director** component is present.
- The "Transporter" item is displayed only if the **Transporter** component is present.
- The "Transporter" item accounts for all **Transporte**r processes.

To open the top tasks list, press *<F10>*.



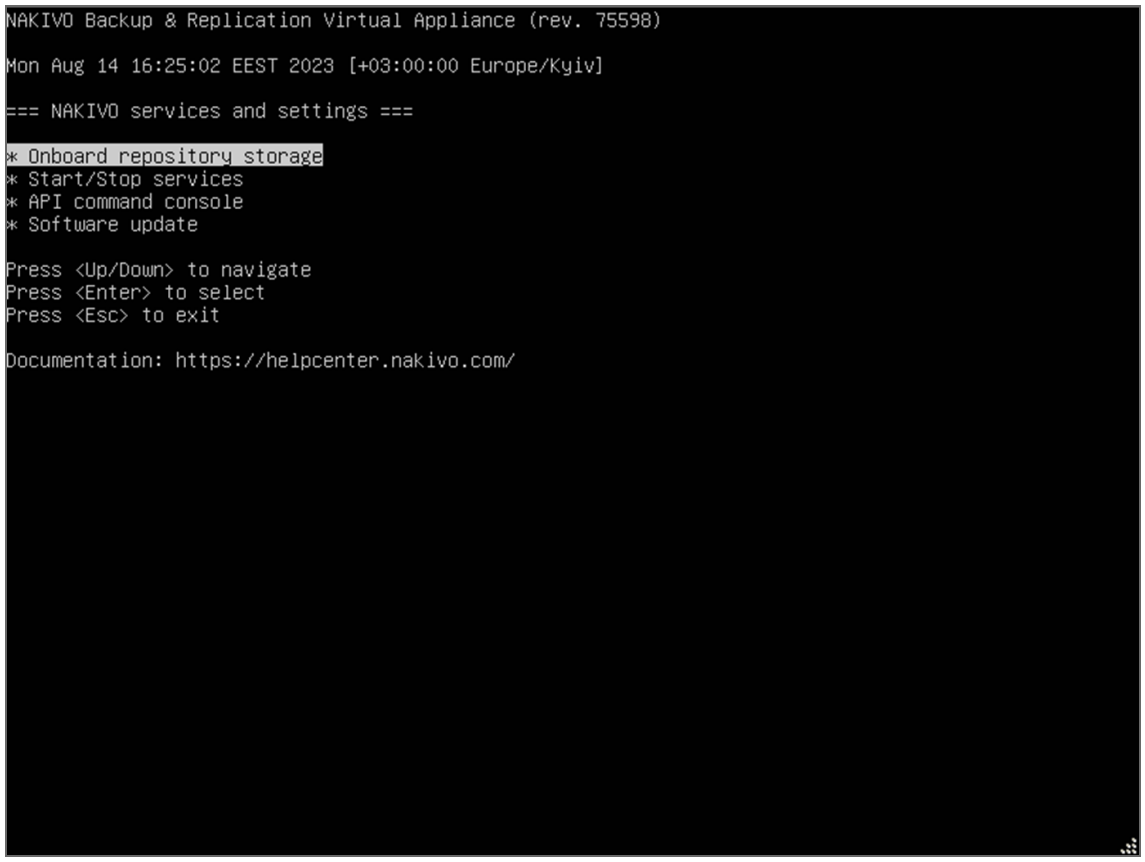

**Notes**

- Press *<F5>* to refresh.
- To exit this screen, press *<Esc>*.

## Manage NAKIVO Backup & Replication Services

Click on **Manage NAKIVO services** opens the NAKIVO services and settings menu where you can select between four options:

- “Onboard repository storage” on the next page
- “Start/Stop services” on page 806
- “API command console” on page 807
- “Software update” on page 808




## Onboard repository storage

Select to view information about the VM Appliance backup repository disk(s) kept on a logical volume (spread across multiple physical volumes).

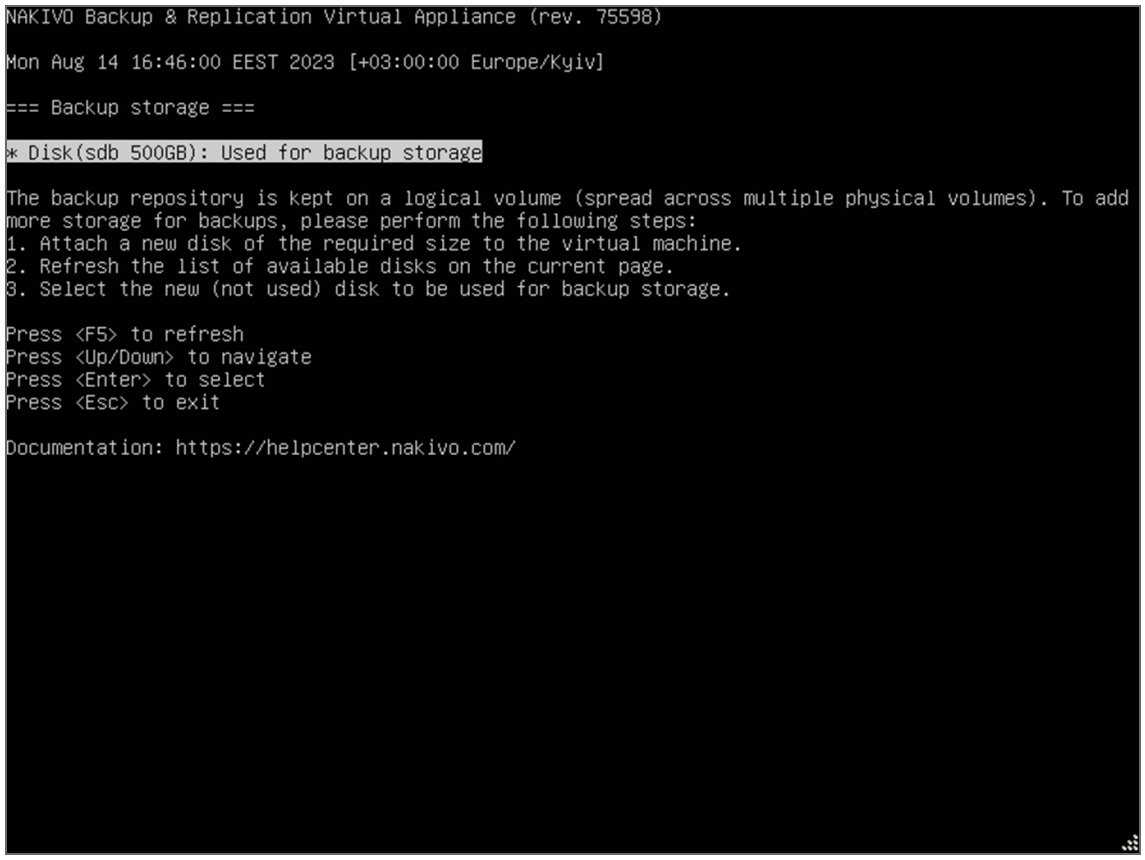
In this screen, you can also attach a new disk to your VA and then configure it for further use as backup storage by following the instructions on the screen.

To add more storage for backups, perform the following steps:

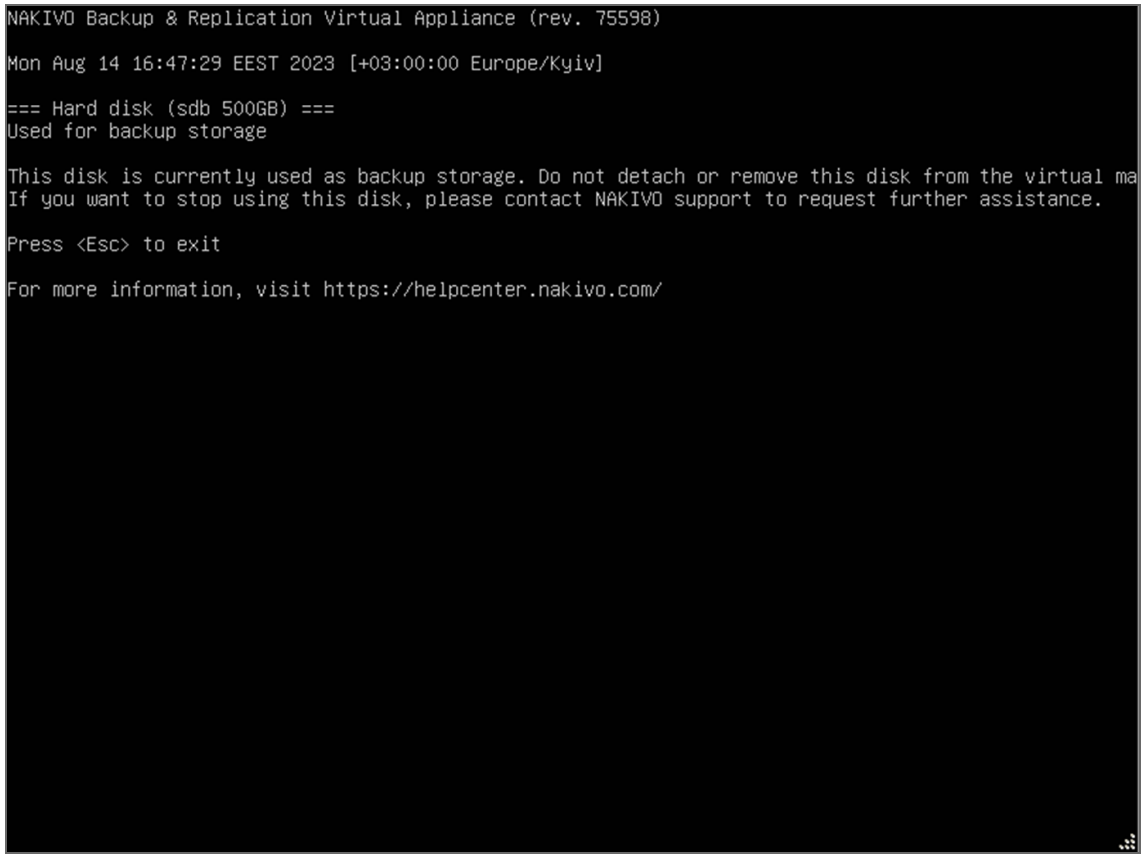
1. Attach a new disk of the required size to the virtual machine.
2. Refresh the list of available disks on the current page.
3. Select the new (not used) disk to be used for backup storage.

**Notes**

- The screen is not available if the **Transporter** component is not installed.
- Creation of a logical structure for backup storage is supported.




For example, if the **Disk** option is selected, the screen with the hard disk details opens.



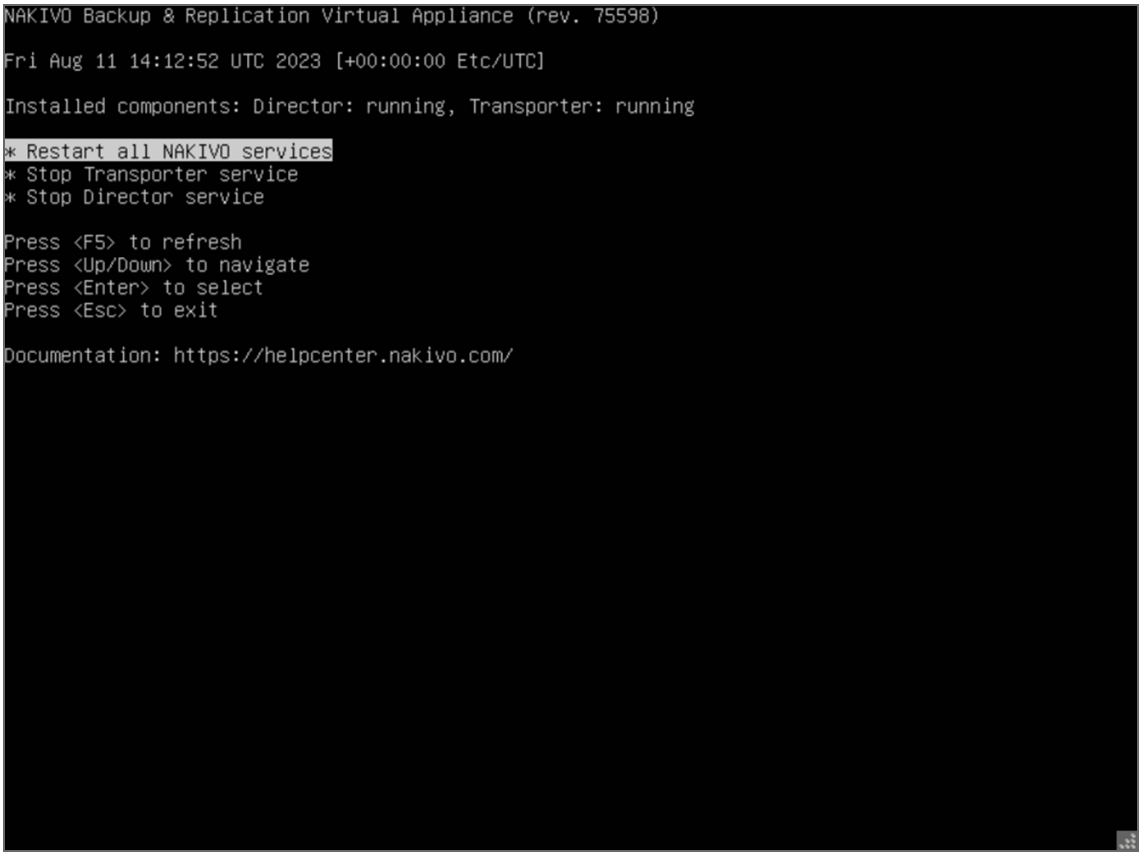
**Important**

- Do not detach or remove the disk currently used for backup storage from the virtual machine.
- If you want to stop using this disk, please contact NAKIVO support to request further assistance.
- If you select to use the disk that is not currently used for backup storage, it will be formatted and appended to the existing logical volume.
- If there is no logical volume structure yet, it will be created automatically.

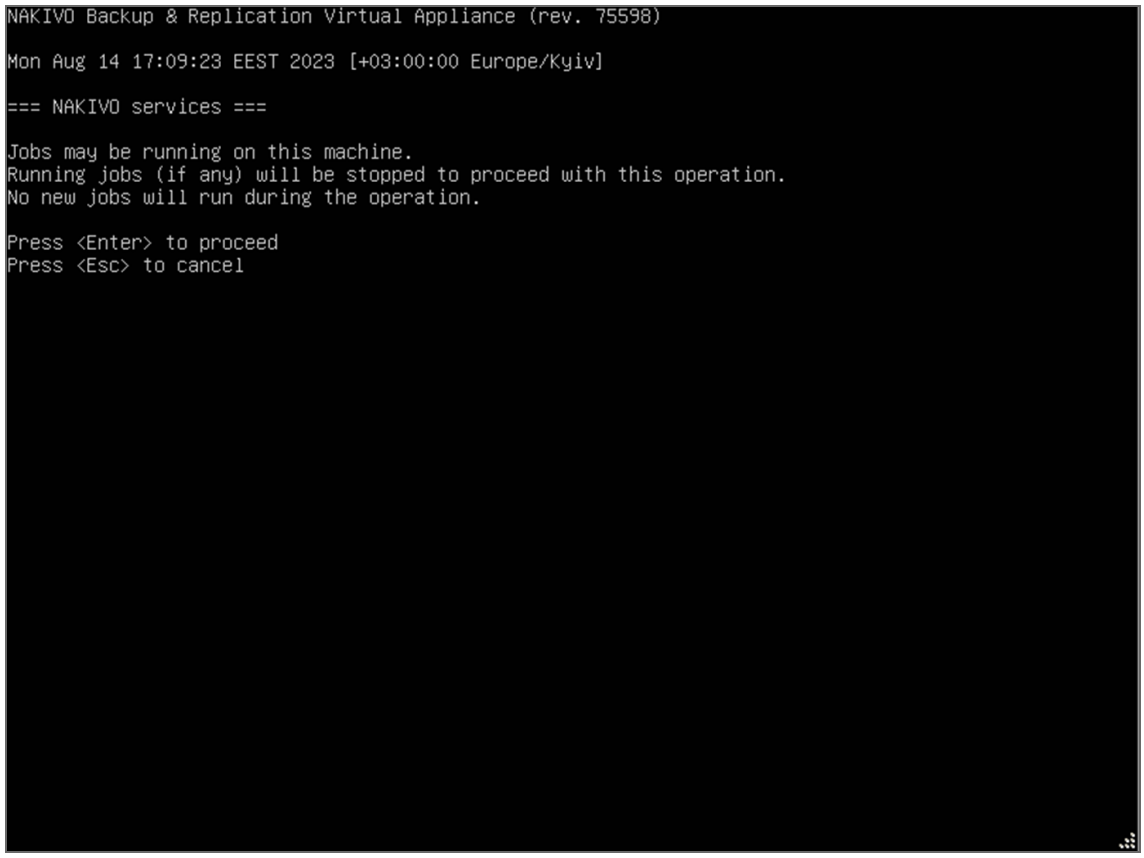
## Start/Stop services

Select to manage the NAKIVO Backup & Replication Virtual Appliance services. The menu offers to select between three options:

1. Restart all NAKIVO services
2. Stop Transporter service
3. Stop Director service




Note that if you decide to select any option, you are asked to confirm your selection.




## API command console

Select to open the system console and run commands in the **interactive mode** from the Nakivo Command line interface.




To exit this screen, type *Exit* and click *<Enter>*.

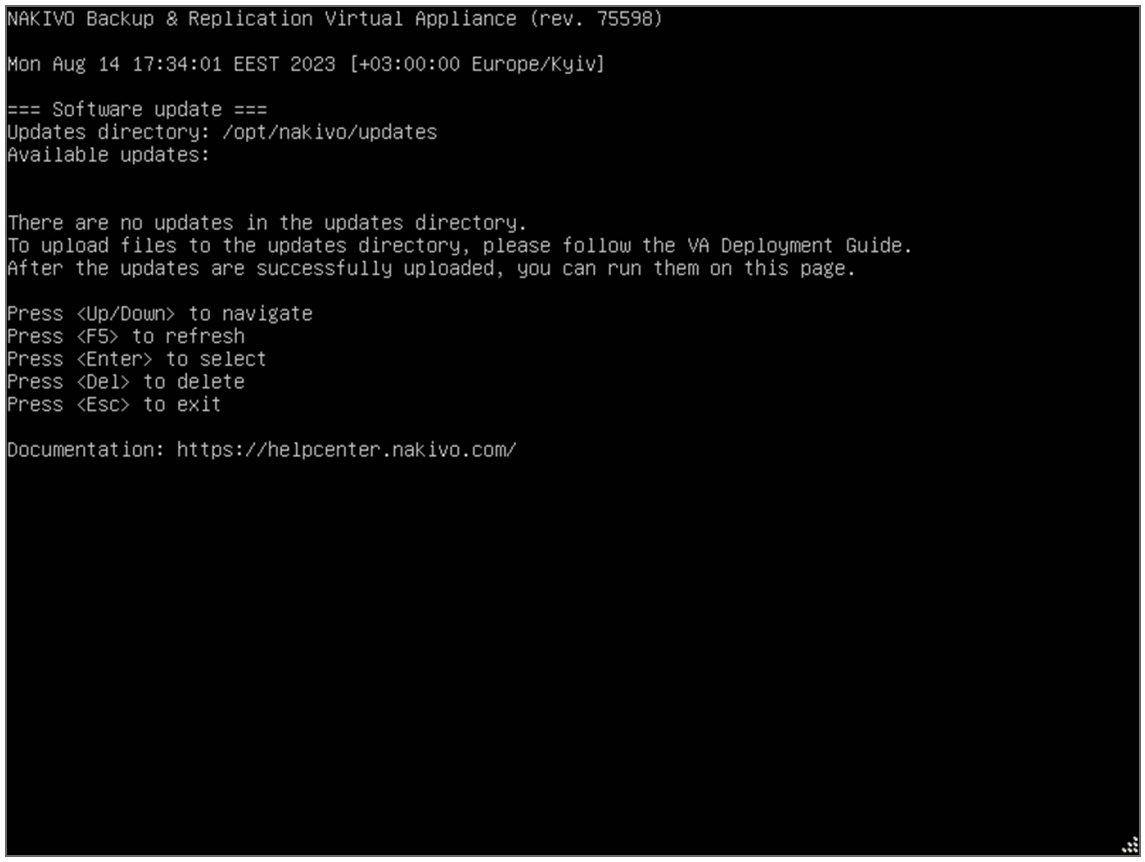
**Notes**

- The screen is not available if the **Director** component is not installed.
- If user credentials are configured in the **Director**, providing the credentials is required.

## Software update

Select to check for available software updates in a certain directory and manage the available updates (if any).

If you decide to start the software update, upload files to the updates directory by following the VA Deployment Guide (available at www.nakivo.com/documentation.htm).




After the updates are uploaded, you can apply them on the current page.

## Exit to System Console

Select to close the VM Appliance Interface and exit to the system console.


va login: _

# Configuring Network Settings of Virtual Appliance

To configure networking on the Virtual Appliance (VA), follow the steps below:

1. Open the VA console.
2. On the main menu, select the **Network Settings** option and press **Enter**.
3. Do either of the following:
   - To change the Virtual Appliance hostname, select the **Hostname** option, press **Enter**, enter a new hostname, and press **Enter** again.
   - To configure a network card, select it and press **Enter**. Press **Enter** to switch between DHCP and manual network settings. If you set the **DHCP** option to **disabled**, you can manually set up network settings by selecting an option, pressing **Enter**, entering a new value, and pressing **Enter** again. Press **F10** to save your changes and exit.

   For more details, check the “Network Settings” on page 791 page.

# Increasing Backup Repository Size on Virtual Appliance

A **Backup Repository** on a Virtual Appliance (VA) is located in a logical volume (that can spread across multiple physical volumes). To extend the **Backup Repository** size on the VA, you need to add a new disk to the VA and then use the VA console to extend the **Backup Repository** to the new disk.

The **Backup Repository** size on the VA cannot be increased by extending existing VA disks.

The backup repository size on the VA cannot be increased by extending existing VA disks. To increase the size of the backup repository on the Virtual Appliance, follow the steps below:

1. Attach a new disk to the VA.
2. Open the VA console in your hypervisor's client.
3. Run the following commands in the VA console depending on the NAKIVO Backup & Replication version you use:
    - For the product Version 8.1 and higher:
        a. Select **Manage NAKIVO services** in the main menu and press **Enter**.
        b. Select **Onboard repository storage** and press **Enter**.
    - For earlier product versions, select **Backup storage** in the main menu and press **Enter**.
4. Refresh the list of disks by pressing **F5**.
5. Select the disk that you have created and press **Enter**.
6. Press **Enter** again to confirm the procedure. The disk is formatted and added to the **Backup Repository** on the VA.

For more details, check the “Manage NAKIVO Backup & Replication Services” on page 803 page.

# Removing Disk with Backup Repository from Virtual Appliance

The Virtual Appliance (VA) comes with a 500 GB disk on which a **Backup Repository** is created. If you have deployed the Virtual Appliance disks using the **Thin Provision** option, then the disk does not consume 500 GB of space on your datastore – only the space occupied by VM backups is consumed.

If you still would like to delete the 500GB disk after you have deployed the Virtual Appliance, follow the steps below:

1. Log in to NAKIVO Backup & Replication.
2. Go to the **Configuration** > **Repositories** tab.
3. Click **Onboard repository**
4. Click **Manage** and choose **Remove** from the menu.
5. In the message that opens, click the **Remove Repository and Delete Backups** button.
6. Click **Remove** to confirm that you wish to remove the **Backup Repository**.
7. Open the vSphere client and launch the console of the VA.
8. In the Virtual Appliance interface, select the **Exit to system console** option and press **Enter**.
9. Enter a login and password (default are `root/roo`t).
10. Run the following command to unmount the volume on which the **Backup Repository** is located: `umount /opt/nakivo/repository`
11. Open the configuration file with the `nano` editor by running the following command: `nano/etc/fstab`
12. In the editor, delete the line: `dev/mapper/Volume_Group_Backup_Repository_ 500GB/Logical_Volume_Backup_Repository_500GB /opt/nakivo ext4 defaults 0 2`
13. Save changes by pressing **Ctrl+O**, and then pressing **Enter**.
14. Exit the editor by pressing **Ctrl+X**.
15. Power off the VA and delete the 500 GB disk.

# Multi-Tenant Mode Configuration

This section covers the following topics:


- “Changing Login and Password in Multi-Tenant Mode” on page 815
- “Configuring Branding Settings in Multi-Tenant Mode” on page 816
- “Configuring Email Notifications in Multi-Tenant Mode” on page 818
- “Configuring Email Settings in Multi-Tenant Mode” on page 819
- “Configuring System Settings in Multi-Tenant Mode” on page 820
- “Exporting and Importing Configuration in Multi-Tenant Mode” on page 822
- “Viewing Tenant Licensing Details in Multi-Tenant Mode” on page 823

# Changing Login and Password in Multi-Tenant Mode

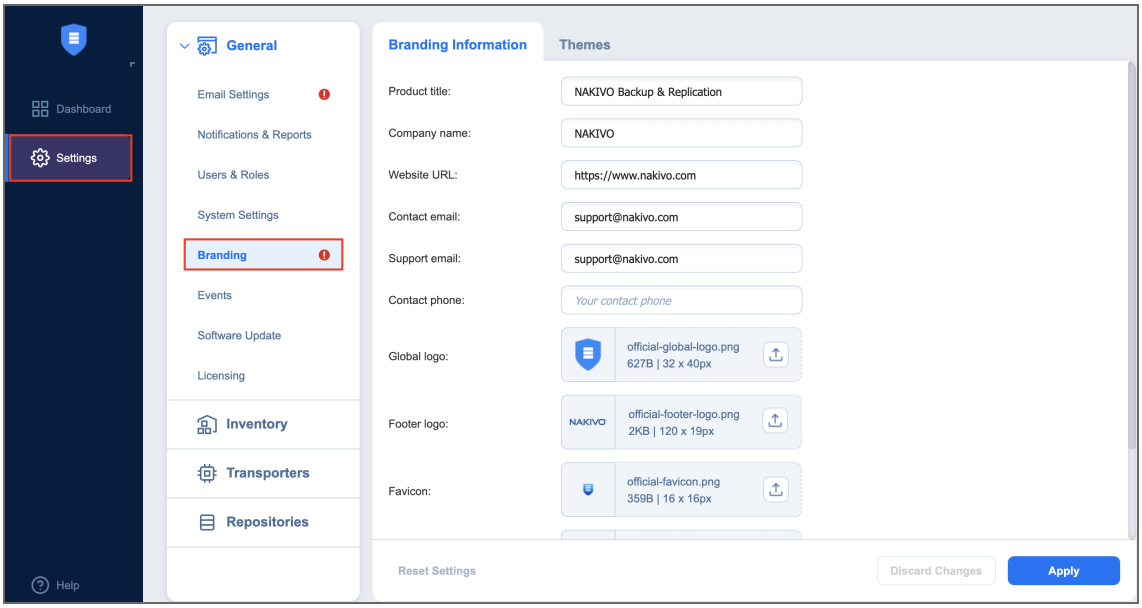
To change the login and password of the Master Admin, follow the steps below:

1. Log in to NAKIVO Backup & Replication as a Master Admin.
2. Click **Configuration** in the upper right corner of the product.
3. Go to the **General** tab and click **Users and Roles**.
4. In the list of users that opens, click the Master Admin user.
5. For the Master Admin, enter data in the **Login, Password**, **Confirm Password**, and **Admin email** boxes and click **Apply.**

# Configuring Branding Settings in Multi-Tenant Mode

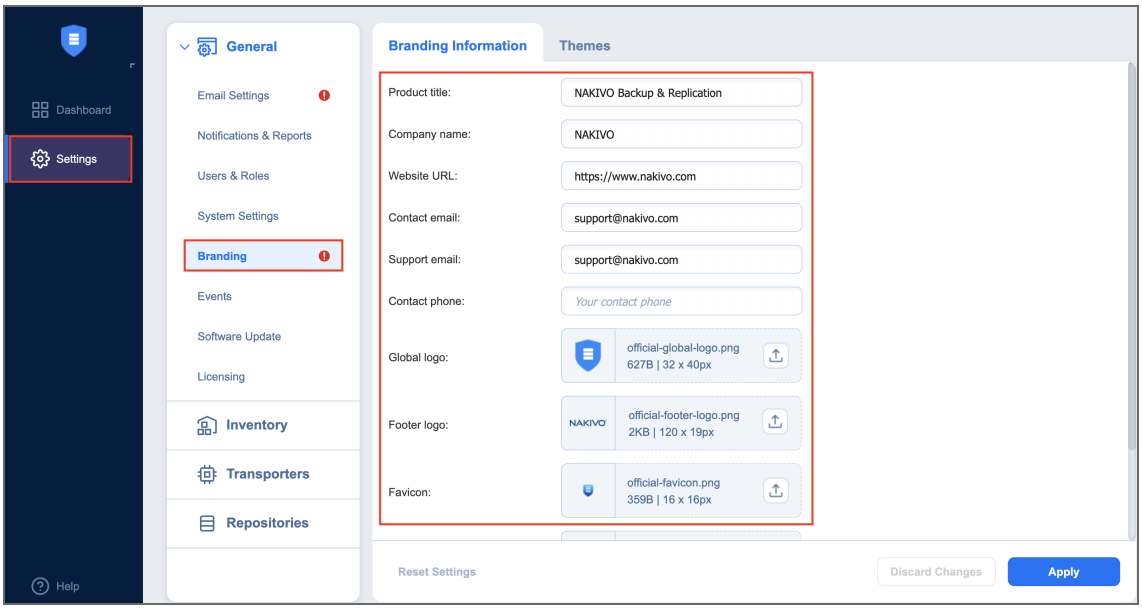
In the multi-tenant mode, you can change the product branding settings such as product name, logo, background, and so on. To configure the system settings, follow the steps below:

1. Log in to NAKIVO Backup & Replication as a Master Admin.
2. Click **Settings** in the left pane of the product.
3. Go to the **General** tab and click **Branding**.




4. Do the following:
   - To change the product title, company name, website URL, contact email, support email, and contact phone, type a new value in the appropriate field
   - To change the product logo, background, and default tenant logo, click **Change** click on the

appropriate box, select a new image, and click **Open**.




5. Click **Apply**.

**Note**

During upload, the logo and bookmark icon images are resized internally while preserving the aspect ratio. The background image is used as it is. To get the best image quality, follow the recommendations below:

| Image | Best format | Best resolution |
|---|---|---|
| Global logo | .png | 40x40 |
| Page background | .jpeg | 1920x1440 |
| Bookmark icon | .png | 16x16 |
| Default tent logo | .png | 120x95 |

# Configuring Email Notifications in Multi-Tenant Mode

NAKIVO Backup & Replication can send notifications and reports over email. To configure the email notifications, follow the steps below:

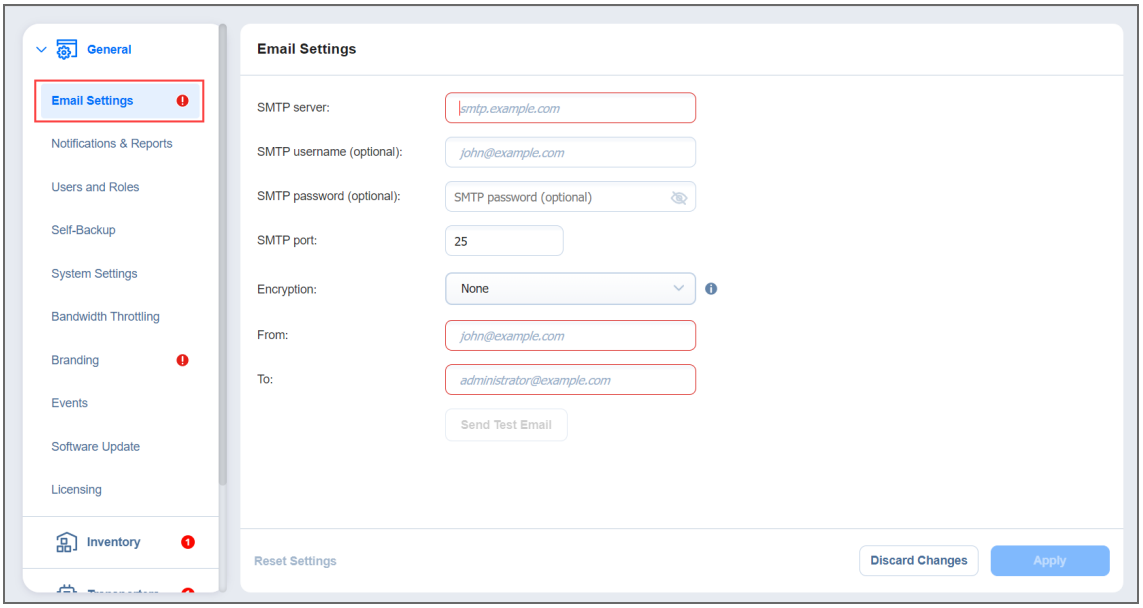
1. Make sure you have configured your email settings.
2. Log in to NAKIVO Backup & Replication as a Master Admin.
3. Click **Settings** in the left pane of the product and go to the **General** tab.
4. Click **Email settings**.
5. In the **Email Notifications** section, select the options as appropriate:
    a. **Send alarm (error) notifications**: If selected, this will send notifications about a job, repository, infrastructure, connection, and other failures to email addresses specified in the text field. Use a semicolon to separate multiple email addresses.
    b. **Send warning notifications**: If selected, this will send warning notifications on non-critical events, such as infrastructure change, to email addresses specified in the text field. Use a semicolon to separate multiple email addresses.
    c. **Limit email notification frequency to**: Set a limit to how often email notifications are sent.
6. In the **Automatic Reports** section, select or deselect the following automatic reports options:
    - **Attach PDF copy to automatic reports**: Specify whether you wish to include a copy of the PDF report with notifications.
    - **Send tenant Overview reports on schedule to**: If this option is selected, NAKIVO Backup & Replication will generate an Overview report (which includes information about all jobs and groups in the product) on the date and time specified in the scheduler and will send the report to the recipients specified in the text field. Use a semicolon to separate multiple email addresses.
    - **Send tenant Protection Coverage reports on schedule to**: If this option is selected, NAKIVO Backup & Replication will generate the Protection Coverage report (which includes information about all VMs & instances protected by backup and/or replication jobs as well as the information about all unprotected VMs & instances) on the date and time specified in the scheduler and will send the report to the recipients specified in the text field. Use a semicolon to separate multiple email addresses.
    - Click **Apply**.

# Configuring Email Settings in Multi-Tenant Mode

Configure email settings so that NAKIVO Backup & Replication can send email notifications as well as reports over email. If email settings are not configured, tenants will not be able to configure email notifications for their jobs. To configure email settings**,** follow the steps below:

1. Log in to NAKIVO Backup & Replication as a Master Admin.
2. Click **Settings** in the left pane of the product.
3. Go to the **General** tab and click **Email notifications**.
4. In the **Email Settings** section, enter data in the boxes, and click **Send Test Email** to verify the settings are correct.

After the email settings are configured, you can configure the product email notifications.



# Configuring System Settings in Multi-Tenant Mode

To configure the system settings, follow the steps below:

1. Log in to NAKIVO Backup & Replication as a Master Admin.
2. Click **Settings** in the left pane of the product.
3. Go to the **General** tab and click **System settings**.
4. Select or deselect the following options:
    - **Store system events for the last X days**: This option specifies the time period (from 10 to 365 days) during which the application events will be kept. Older events are automatically deleted.
    - **Auto log out after X minutes of inactivity**: If this option is selected, the current user will be automatically logged out of the product after the specified period of inactivity.
    - **Auto upload support bundles to support team server**: If this option is enabled, NAKIVO Backup & Replication will automatically create, encrypt, and upload support bundles once a day to a NAKIVO support server during the evaluation period. The NAKIVO Support team may use this information to improve the product experience and will be able to identify and resolve product issues faster.
    - **Enable built-in support chat:** If selected, this will allow you to chat with the NAKIVO support team.
    - **Display special offers**: If selected, this will show a toolbar with special offers in the GUI.
    - **Continue product update if self-backup fails**: If selected, product update will proceed even if automatic self-backup cannot be performed.
    - **Tape options**: These present you with setting options for tape devices:
        - **Auto erase expired tapes**: If selected, expired tape cartridges will be erased automatically.
        - **Wait for next tape for**: Specify for how long the system needs to wait for the next tape cartridge if there is no appropriate one. Select the **Send email notification** checkbox to allow you to receive email notifications.
        - **Auto refresh tapes every**: Select how often the contents of tape cartridges are to be refreshed in minutes or hours. Deselect if no refreshing is required.
    - **Regional options**: Set the clock format, short date format, long date format, first day of the week, decimal symbol, and default time zone in the corresponding fields.
- In the **Web Interface TLS/SSL Certificate** section, you can either**:**
    - **View current certificate**: A dialog containing the current certificate information opens.
    - **Install new certificate**: A dialog opens, allowing you to install a new TLS/SSL certificate for the NAKIVO Backup & Replication web interface. Certificates are generated either internally or through certification authorities. Proceed as follows to install a new certificate:

- Click **Browse** and navigate to the location of either of the following certificate file types:
  - **Private key**: A file in the *.key format.
  - **Private key password (optional)**: A password for your private key.
  - **Certificate file**: A file in the *.pem, *.crt, *.cer, *.p7b, or *.p7s format.
  - **Intermediate certificate (optional)**: A file in the *.pem, *.crt, *.cer, *.p7b, *.p7s format.
- Click **Install**.

**Note**

In the Web Interface TLS/SSL Certificate section, you can see a notification about imminent TLS/SSL certificate expiration in 30 days and onwards. If your certificate has expired, you will be asked to install a valid certificate.

# Exporting and Importing Configuration in Multi-Tenant Mode

System configuration export and import are recommended for easy migration to new product deployment. System configuration, such as jobs, user credentials, inventory items, Transporter and Backup Repository settings, is all exported into a single export bundle.

The export bundle can be applied to a new deployment.

To export system configuration from the old deployment, follow the steps below:

1. Open **Settings** in the old deployment.
2. Go to the **General** tab and click **System migration**.
3. Click **Export system configuration**.
4. In the dialog box that opens, click **Export**.
5. Click **Proceed** to confirm the operation.

> **Note**
> All activities in the old deployment (such as jobs and recovery sessions) will be automatically stopped and disabled.

6. Wait until the export is completed, and download the export bundle.

To import system configuration into the new deployment, follow the steps below:

1. Open **Settings** in the new deployment.
2. Go to the **General** tab and click **System migration**.
3. Click **Import system configuration**.
4. In the dialog window that appears, locate the export bundle using the **Browse** button.
5. Click **Import**.
6. Click **Proceed** to confirm the operation.

> **Note**
> If there is any existing data in the new deployment, it will be overwritten with the import operation.

7. Wait until the import is completed, and close the dialog box.

**Notes**

1. Data contained in backup repositories is not migrated to the new location automatically. If you are using a locally attached Backup Repository, the physical data must be copied or moved to the new location manually.
   After moving the files you may need to edit the Backup Repository settings in the new deployment so that the new settings refer to the actual Backup Repository location.
2. If a custom TLS/SSL certificate of the Web server was used in the old deployment, a manual service restart will be required in the new deployment.

# Viewing Tenant Licensing Details in Multi-Tenant Mode

**Important**

This tab is only displayed for users with an *MSP license*, *Beta instance*, *Promo license*, or *Trial license*.

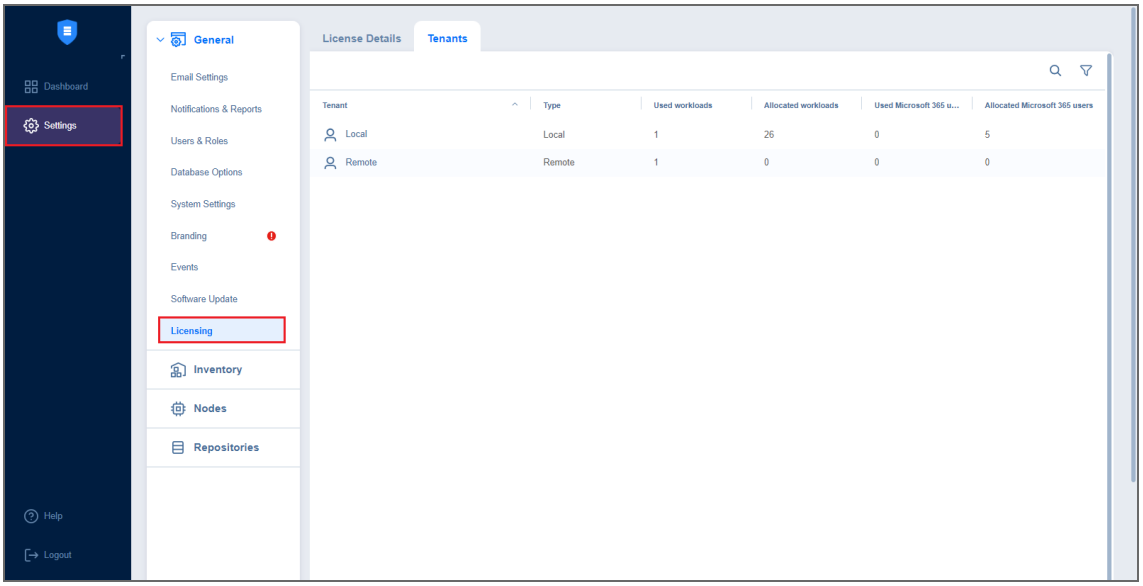
To check the license details of all local and remote tenants connected to a managed service provider (MSP) through the **MSP Console**, follow these steps:

1. Go to the main menu of NAKIVO Backup & Replication and click **Settings** > **General**.
2. Go to the **Licensing** > **Tenants** tab.

   **Note**

   To display the **MSP Console**, you must have the multi-tenancy mode installed.

3. The tab opens the page with the **Tenants** table that allows viewing the licensing data of all existing tenants.

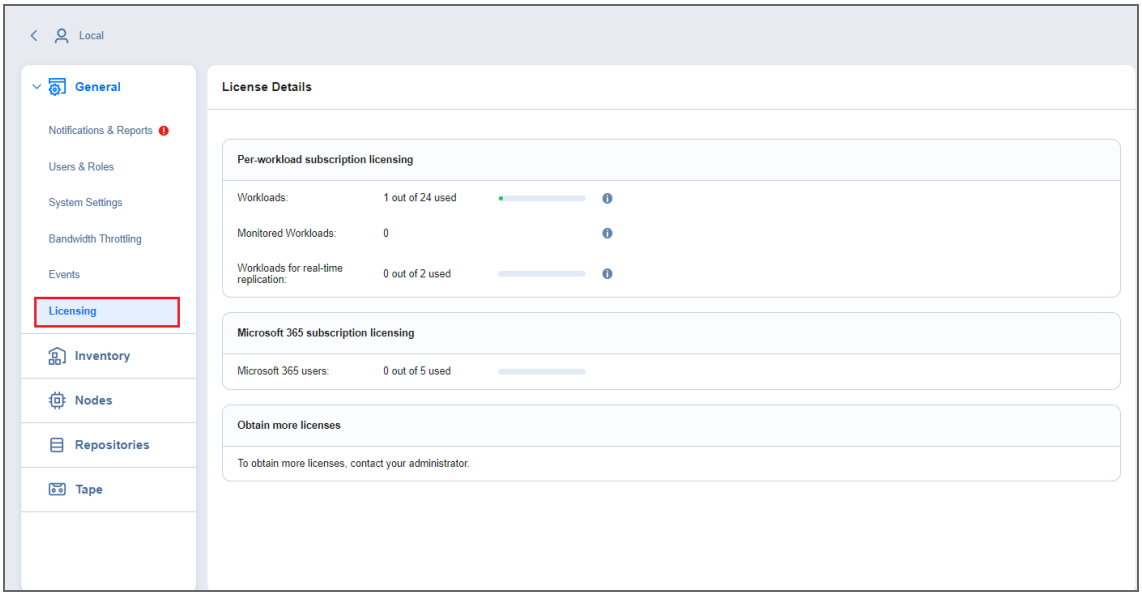



The table has the following columns:

- **Tenant**: The name of the tenant.

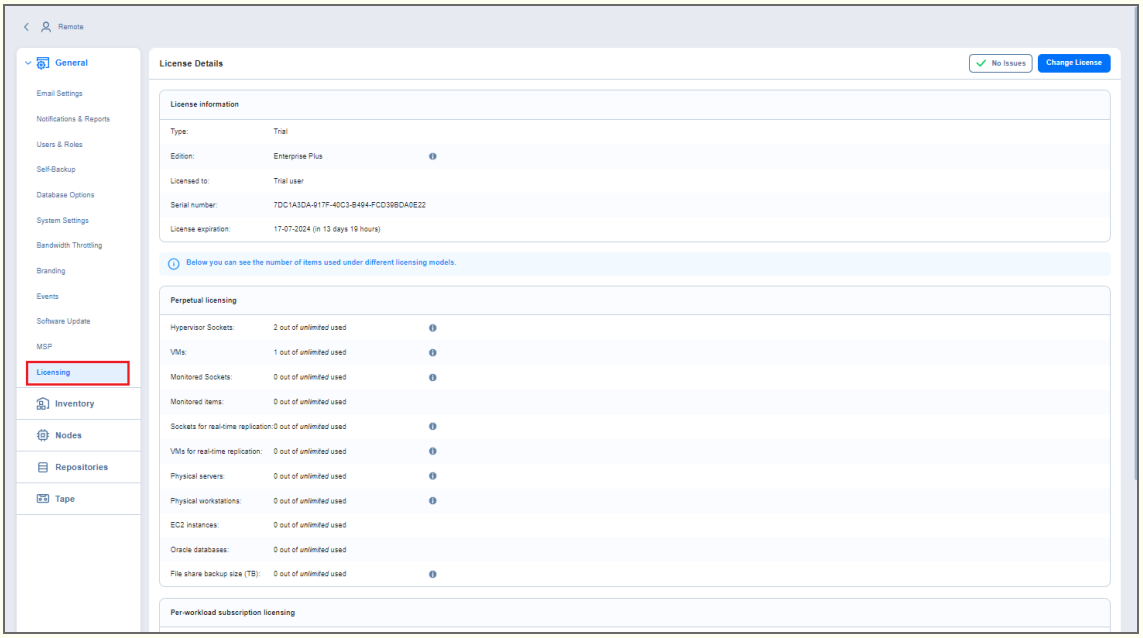
> **Notes:**
>
> - Clicking the name opens the **Licensing** tab of the corresponding tenant.
> - *Disabled*, *Disconnected*, and *Inaccessible* tenants are disabled.
> - Hovering over a tenant's data point (for example, tenant's name, type, workloads, etc.) displays the text with the time since the tenant's data was last updated.

- **Type**: The tenant's connection type – **Local** or **Remote**.
- **Used workloads**: The number of a tenant's used workloads.



**Note**

For the tenants with a perpetual license, the **Perpetual** link is displayed instead; clicking the link opens the **Tenant licensing details** popup displaying the **Perpetual licensing block** in the **Licensing** tab of the corresponding tenant.




- **Allocated workloads**:
  - For remote tenants, displays the number of workloads available to the tenant.
  - For local tenants, displays the number of workloads allocated to the tenant.

> **Note:**
> For the tenants with a perpetual license, the **Perpetual** link is displayed instead. Clicking the link opens the **Tenant licensing details** popup displaying the **Perpetual licensing block** in the **Licensing** tab of the corresponding tenant.

- **Used Microsoft 365 users**: The number of Microsoft 365 users used by the tenant.
- **Allocated Microsoft 365 users**:
  - For local tenants, displays the number of Microsoft 365 users allocated to the tenant.
  - For remote tenants, displays the number of Microsoft 365 users available to the tenant.

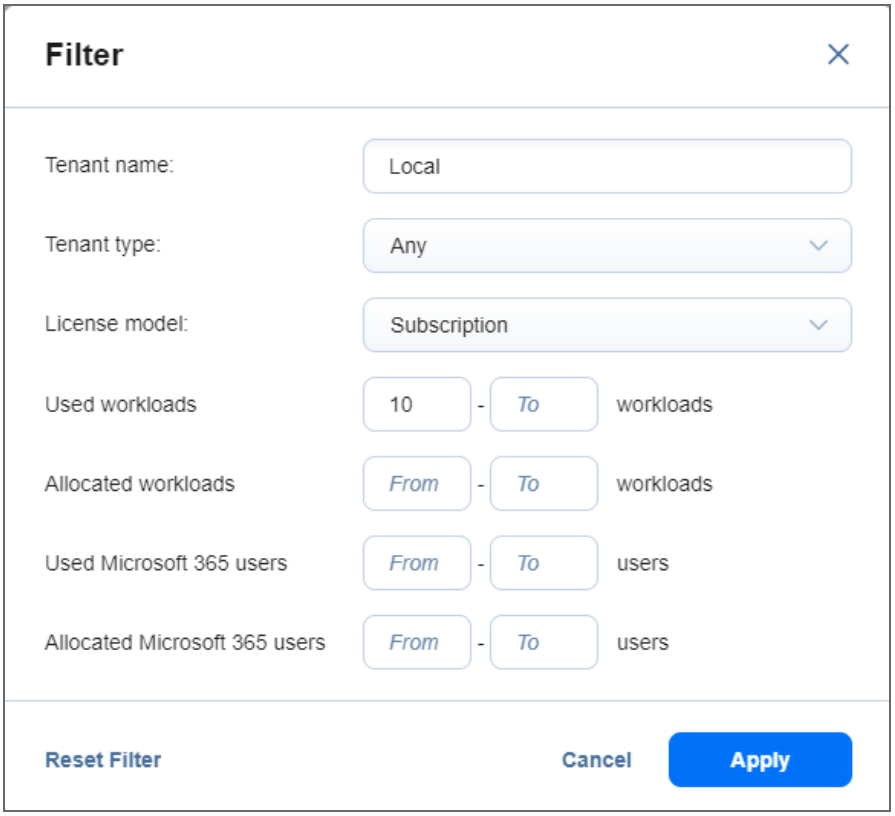
## Using Filtering

You can search or filter tenants from the **Tenants** table.

Use the **Search** field to browse or search for a specific tenant. Search can be performed on the **Name** column only.

To access filtering options, click the **Filter** icon in the top right corner. In the dialog box that opens, you can select one or more filtering criteria. The following filtering options are available:

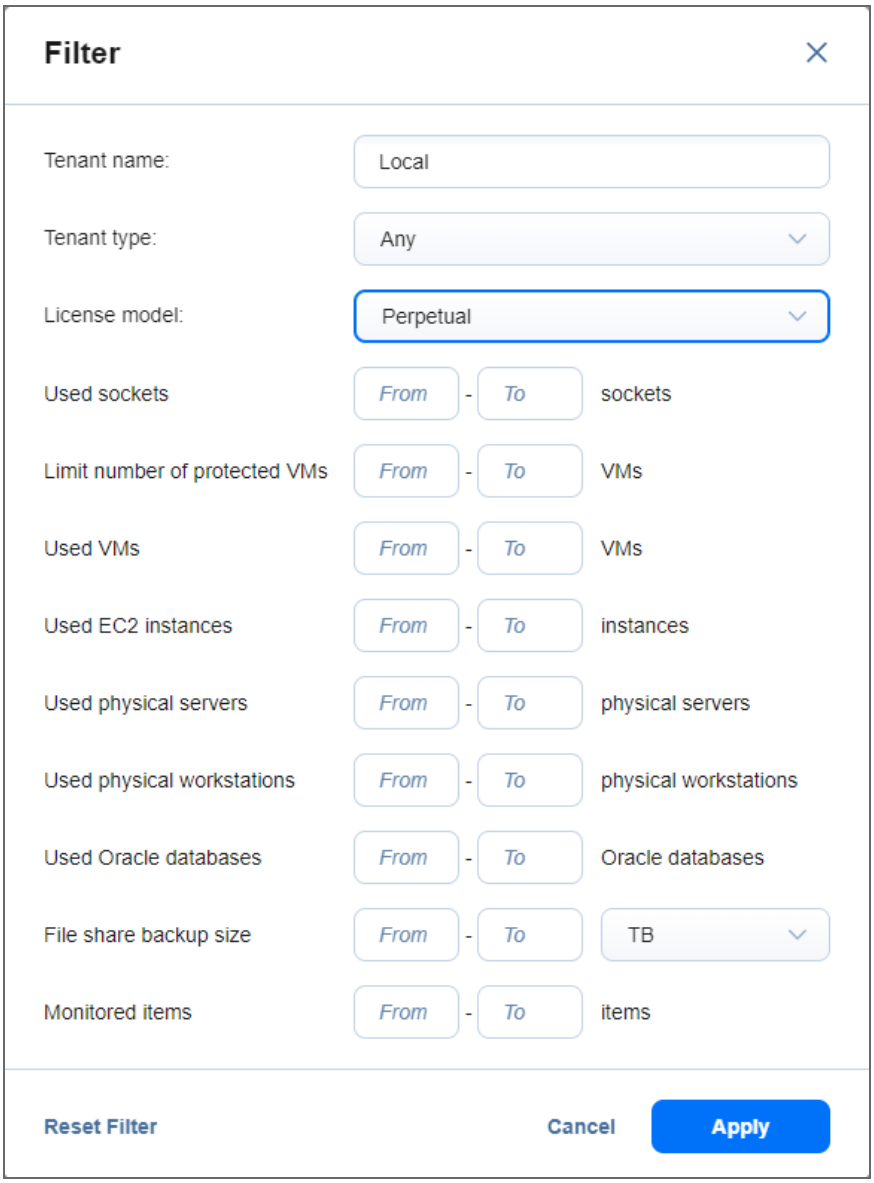
- **Tenant name**: Allows you to filter by tenant name.
- **Tenant type**: Allows you to filter by the following options:
  - **Any** (default): Both **Local** and **Remote** tenant types
  - **Local**
  - **Remote**
- **License model**: Allows you to filter by the following options:
  - **Any** (default): If selected, no extra fields are displayed.
  - **Subscription**: If selected, the following range input fields are displayed:
    - **Used workloads**
    - **Allocated workloads**
    - **Used Microsoft 365 users**
    - **Allocated Microsoft 365 users**




> **Note**
> The **Allocated workloads** and **Allocated Microsoft 365** users fields are not displayed if **Remote** is selected in the **Type** dropdown.

  - **Perpetual**: If selected, the following range input fields are displayed:
    - **Used sockets**
    - **Limit number of protected VMs**

- **Used VMs**
- **Used EC2 instances**
- **Used physical servers**
- **Used physical workstations**
- **Used Oracle databases**
- **File share backup size** with two dropdown options: **TB** (default) and **GB**
- **Monitored items**




**Note**

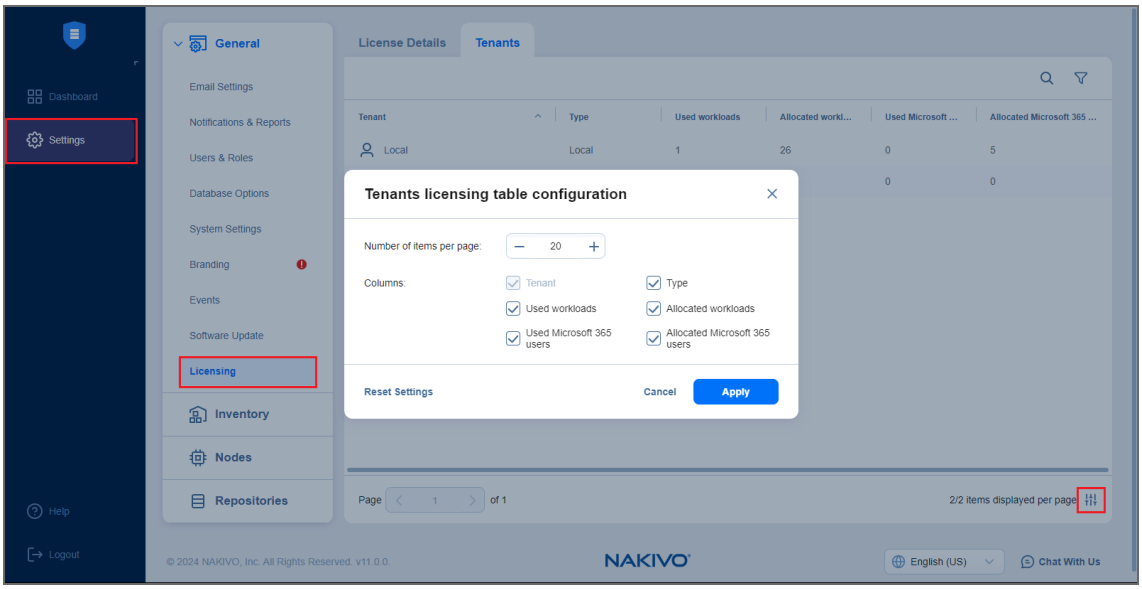
All range inputs can accept numeric values from *0* to *9999*.

Click **Apply** to close the dialog box and apply the filtering.

Click **Reset Filter** to reset the fields in the **Filter** dialog box to the default.

Click the **Cancel** or **X** button to close the dialog box without applying any changes.

Optionally, you can show/hide columns or modify the number of items per page in the **Tenants licensing table configuration** table. In the lower right corner, click the controls icon. In the dialog window that opens, select/deselect checkboxes and click **Apply**.




Click **Reset Settings** to reset the configuration settings to the default.

Click the **Cancel** or **X** button to close the dialog box without applying any changes.

# Support Bundles

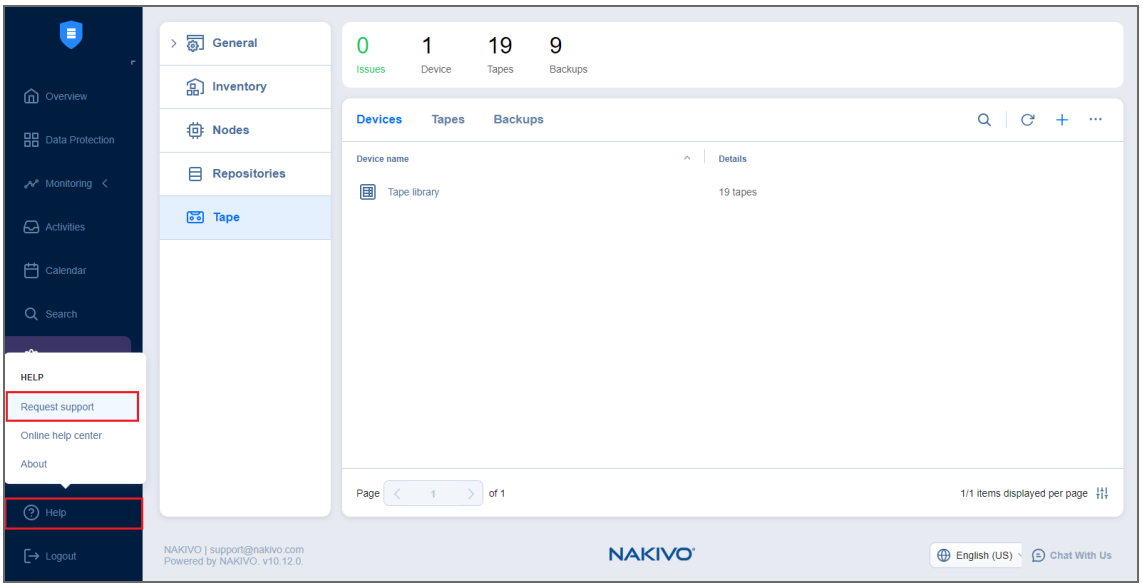
NAKIVO Backup & Replication provides you with the ability to create support bundles – a zipped collection of the product logs and system information. Sending a support bundle to the NAKIVO Support Team allows them to quickly identify the root cause of issues and suggest a proper solution.

- Creating Support Bundles
- Sending Support Bundles

## Creating Support Bundles

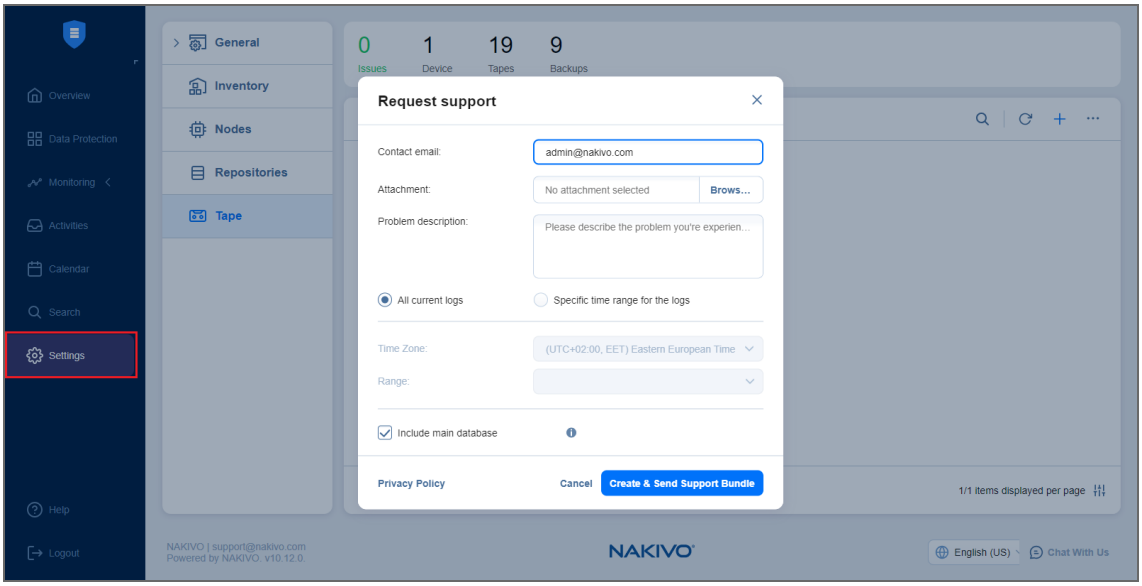
Before creating a support bundle, make sure Email settings are configured. To create a support bundle, follow these steps:

1. Click the **"?" (Help)** icon in the lower-left corner of the web UI.
2. Select and click **Request support**. The dialog box will appear.




3. Enter a description of your problem in the **Please describe the problem you're experiencing** box.
4. Enter your email address in the **Contact email** box.
5. If necessary, upload an attachment by clicking **Browse**.
6. Select **All current logs** if you wish to include all current log files to the support bundle.
7. Select **Specific time range for the logs** if you want to include the log files within the specific time range.
8. To specify the time range, click the **Range** dropdown and use the **From** and **To** selectors to set the time range; click **Apply**.
9. Optionally, click **Include main database** (recommended).

10. Click **Create & Send Support Bundle** to send the support bundle to NAKIVO Support Team. You will receive an answer from the NAKIVO Support Team within one business day.




# Sending Support Bundles Manually

Some support bundles may become overly large in size. This can occur due to large log files or file dumps. In such cases, it is recommended to upload these files manually.

To do this, follow these steps:

1. Open the **Upload Files to NAKIVO Support** page.
2. In the *Files* section, click **Browse** and select up to three files. You can select more than three files by clicking **Add Another File**.

> **Note**
> You can upload any files relevant to your issue: logs, file dumps, or the support bundles that you have manually downloaded from the product's UI.

3. Enter your email address in the **Contact email** field.
4. You can also enter the ID of your support ticket in the **Ticket ID** field if you have one opened.
5. Optionally, enter a description in the **Description** field.
6. Click **Upload** when you're done uploading the file(s).

> **Note**
> Wait for a successful upload notification before closing the page.

# Built-in Support Chat

You have the possibility to contact a NAKIVO representative via chat in the NAKIVO Backup & Replication interface.

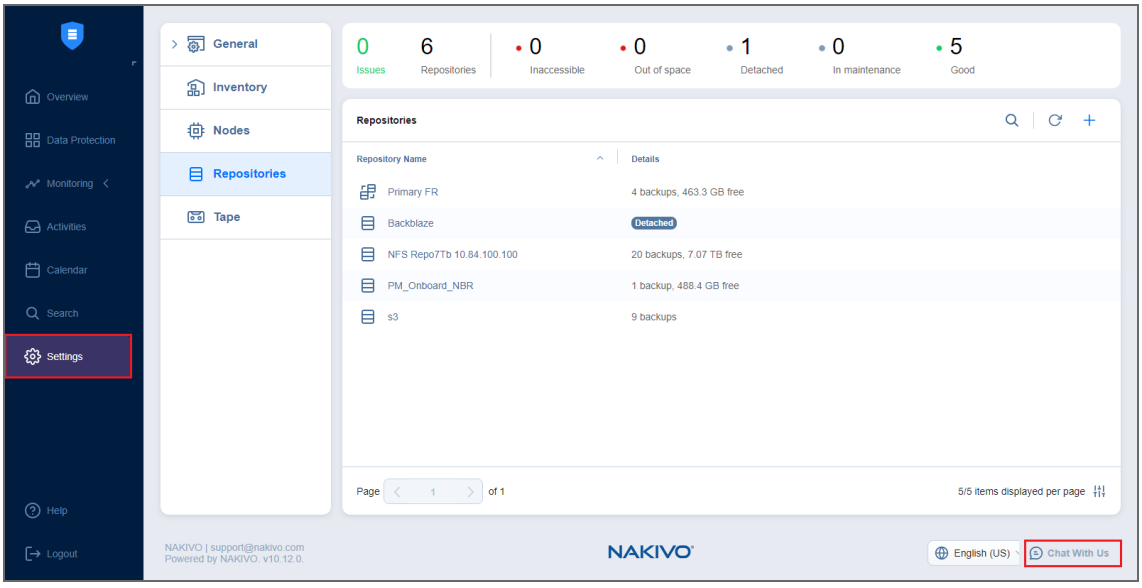
The built-in **Nakivo Support** chat is an interface where you can get live professional support. For a detailed explanation of each component, see the sections below:


- Opening Built-in Support Chat
- Sending Files in Built-in Support Chat
- Sending Email Transcript of Built-in Support Chat
- Editing Contact Details
- Disabling/Enabling Sound Notifications
- Disabling Built-in Support Chat
- Sending Feedback to Built-in Support Chat


## Opening Built-in Support Chat

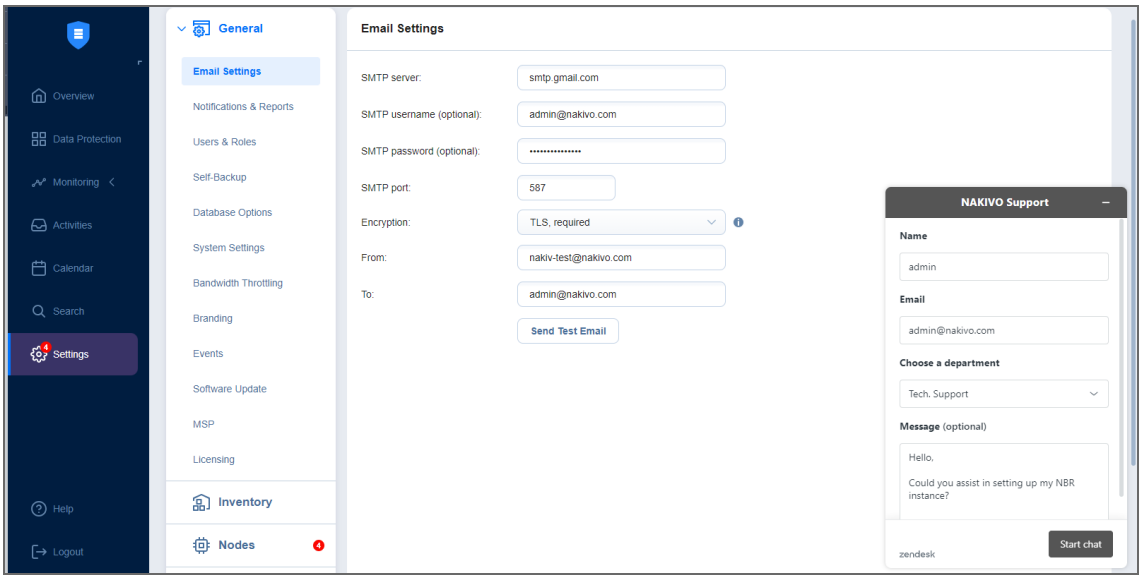
To open the built-in **Live Support Chat**, follow the steps below:

1. In the lower right corner of the NAKIVO Backup & Replication interface, click the **Chat With Us** button.

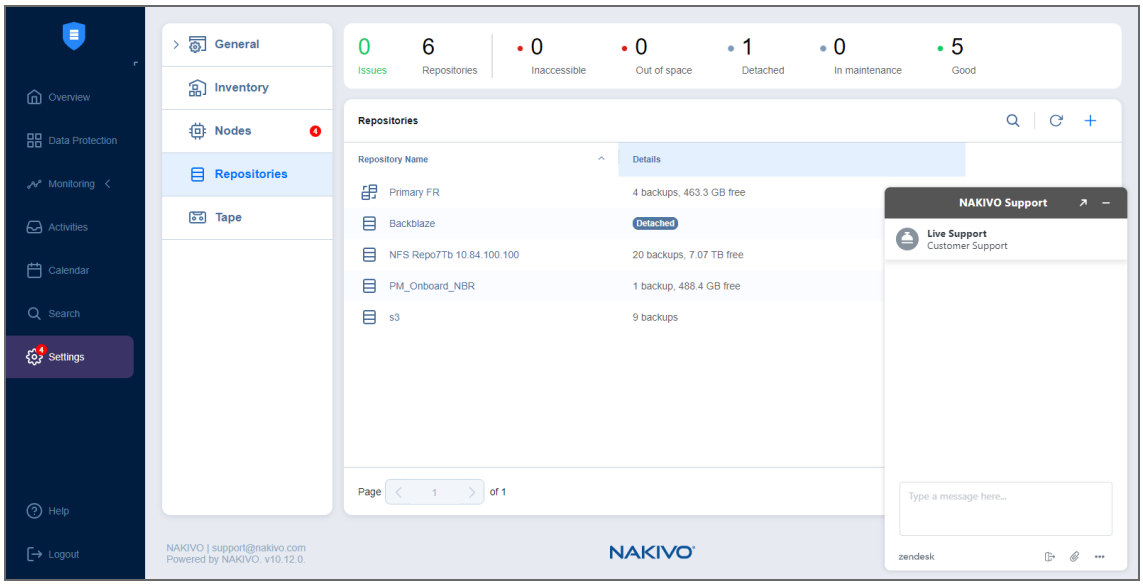



2. The **NAKIVO Support** dialog box opens. Introduce yourself by providing the following information:
   a. In the upper box of the dialog box, enter your name.
   b. In the box below, enter your email address.
3. Choose a department from the list of available departments.

4. Enter your message text and click **Start chat**.

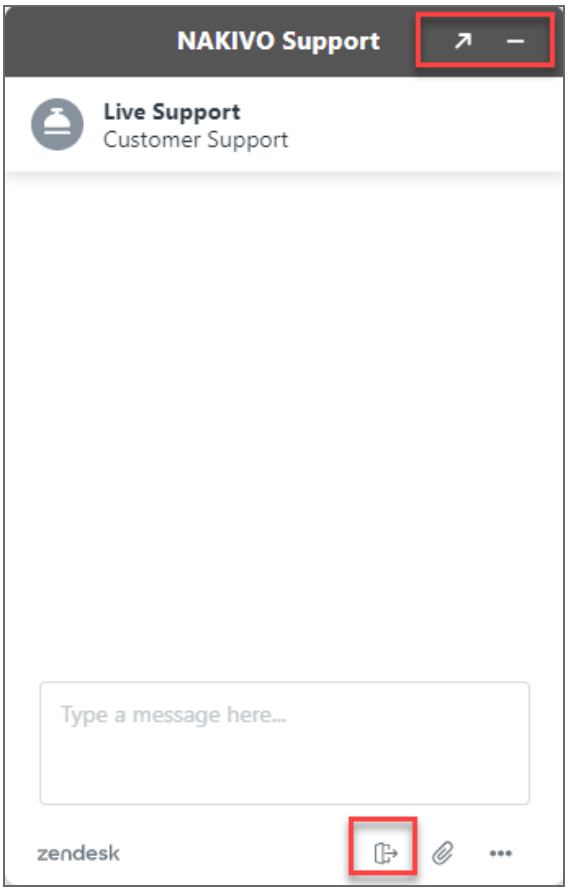



5. Your message is sent to a NAKIVO representative and the built-in **Live Support Chat** dialog box opens.




6. Optionally, you can open the **Web Widget Live Chat** by clicking the **Up right arrow** ↗ button in the chat header.

7. To leave the chat, click the minus (-)button in the chat header or click the **Exit** button in the lower right corner of the built-in **Live Support Chat** dialog box.

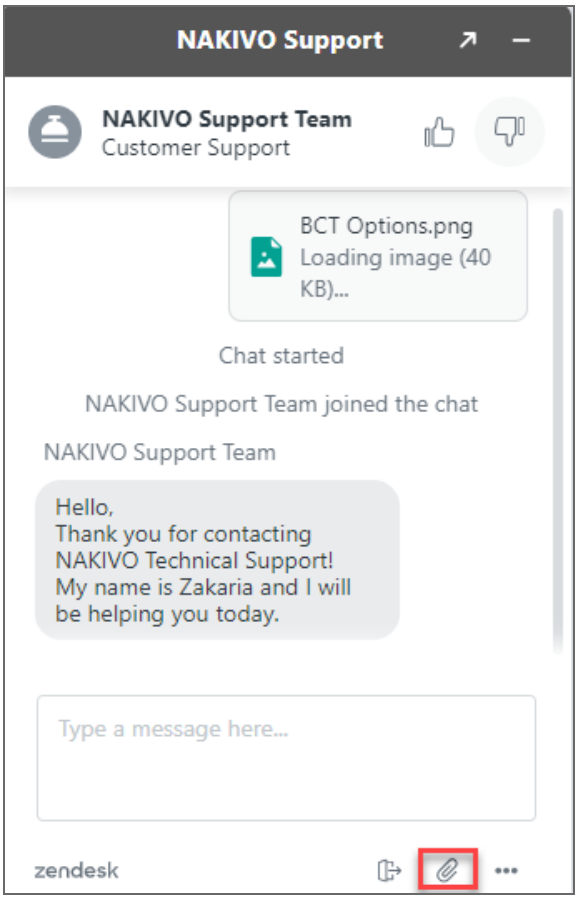



# Sending Files in Built-in Support Chat

Please use either of the following ways to send your files in the built-in **Live Support Chat**:

- Drag and drop: open **Windows File Explorer**, select necessary files, and then drag them and drop to the chat dialog.
- The built-in **Live Support Chat** interface:
  1. In the lower right corner of the built-in **Live Support Chat** dialog box, click the **Attach** button.
  2. In the window that opens, select the needed file to attach to the chat.
  3. If needed, you can provide additional details in the message field.

4. Your message is sent to a NAKIVO representative and will be processed as soon as possible.




**Note**

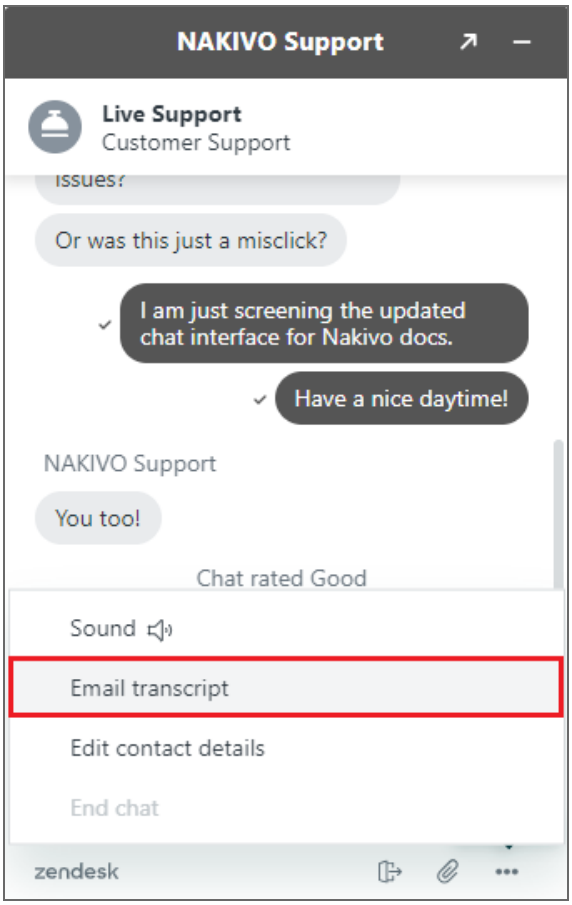
The following file formats are allowed: `.pdf`, `.png`, `.jpeg`, `.gif`, `.txt`. The maximum file size is 20 MB.

# Sending Email Transcript of Built-in Support Chat

Follow the steps below to send the transcript of your built-in **Live Support Chat** session:

1. In the lower right corner of the built-in **Live Support Chat** dialog, click the **...** (ellipsis) button.
2. In the dialog box that opens, click **Email transcript**.
3. In the dialog box that opens, make sure the email address of the recipient is correct, and then click **Send**.

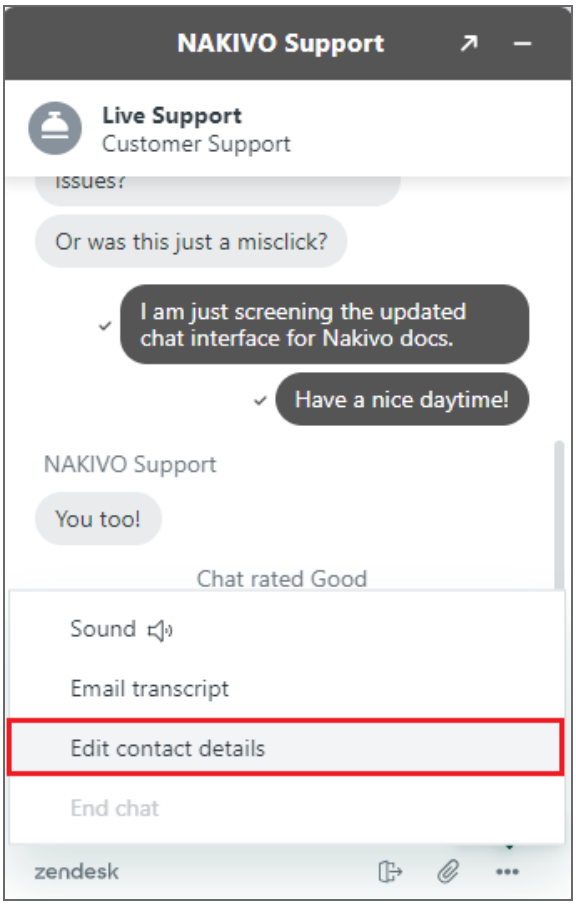
Your built-in **Live Support Chat** transcript will be sent to the specified email recipient.




# Editing Contact Details

Follow the steps below to edit the contact details of your built-in **Live Support Chat** session:

1. In the lower right corner of the built-in **Live Support Chat** dialog, click the **...** (ellipsis) button.
2. In the dialog box that opens, click **Edit contact details**.
3. In the dialog box that opens, edit the contact name and the email address, and then click **Save**.



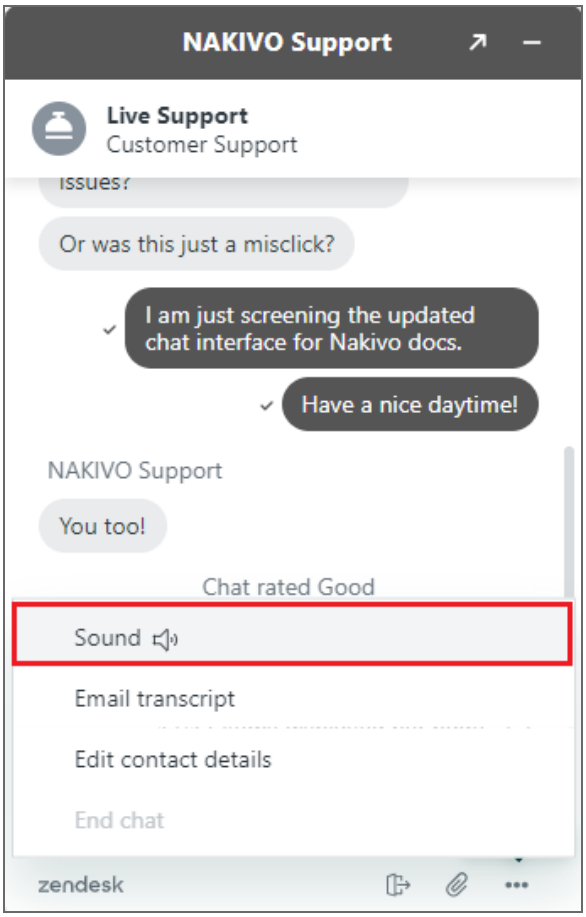

# Disabling/Enabling Sound Notifications

By default, sound notifications are enabled for the built-in **Live Support Chat**.

Do the following to disable sound notifications:

1. In the lower right corner of the built-in **Live Support Chat** dialog, click the **...** (ellipsis) button.
2. In the dialog box that opens, click **Sound**.
3. Close the options dialog.

Sound notifications will be disabled for the built-in **Live Support Chat**.

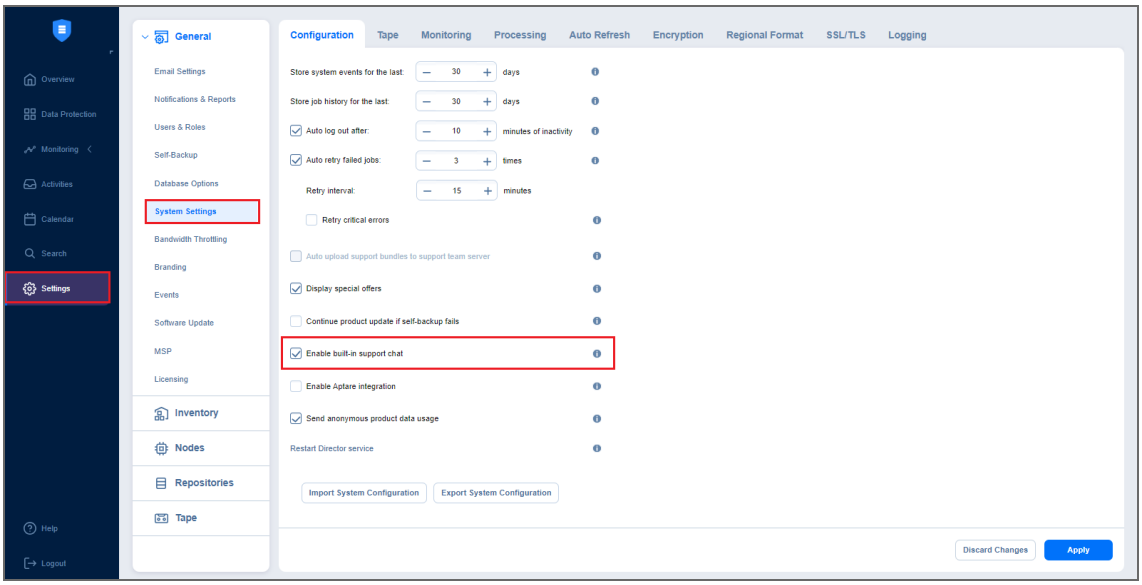



# Disabling Built-in Support Chat

By default, the built-in support chat is enabled in your instance of NAKIVO Backup & Replication.

Do the following to disable the built-in **Live Support Chat**:

1. Go to **Settings** > **General** > **System Settings**.
2. Deselect the **Enable built-in support chat** checkbox.
3. Click the **Apply** button.




**Note**

When disabled, the **Built-in Support Chat** is not available in all tenants of the NAKIVO Backup & Replication instance in the Multi-tenant mode.

# Sending Feedback to Built-in Support Chat

You have the possibility of sending feedback to built-in **Live Support Chat**: in the upper right corner of the dialog, click **Good** or **Bad**, as you deem appropriate.

If appropriate, leave a comment for the NAKIVO Support Team: click **Leave a comment** and in the text box that opens, enter your comment about the chat service. Then click **Send**.

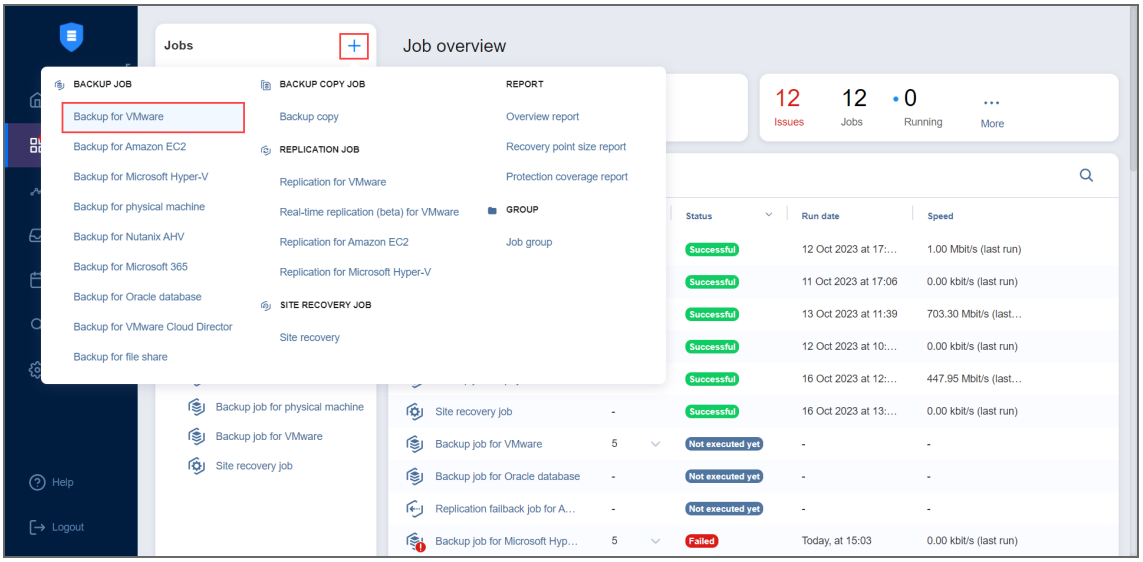
# Backup

This section contains the following topics:


- “Creating VMware Backup Jobs” below
- “Creating Backup Copy Jobs” on page 872
- “Backing Up to Tape” on page 901
- “Staging (Seeding) Initial Backup” on page 923
- “Deleting Backups” on page 924


# Creating VMware Backup Jobs

With NAKIVO Backup & Replication, you can back up VMware VMs by creating a backup job that specifies which VMs should be backed up, where the backups should be located, how often the backup should be run, and what backup options should be used. To create a backup job, click the plus **Create** button in the **Data Protection** menu and then click **Backup for VMware**.




The **New Backup Job Wizard for VMware** opens. Complete the wizard as described in the sections below:


- “Backup Job Wizard for VMware: Source” on the next page
- “Backup Job Wizard for VMware: Destination” on page 843
- “Backup Job Wizard for VMware: Schedule” on page 846
- “Backup Job Wizard for VMware: Retention” on page 856
- “Backup Job Wizard for VMware: Options” on page 857

# Backup Job Wizard for VMware: Source

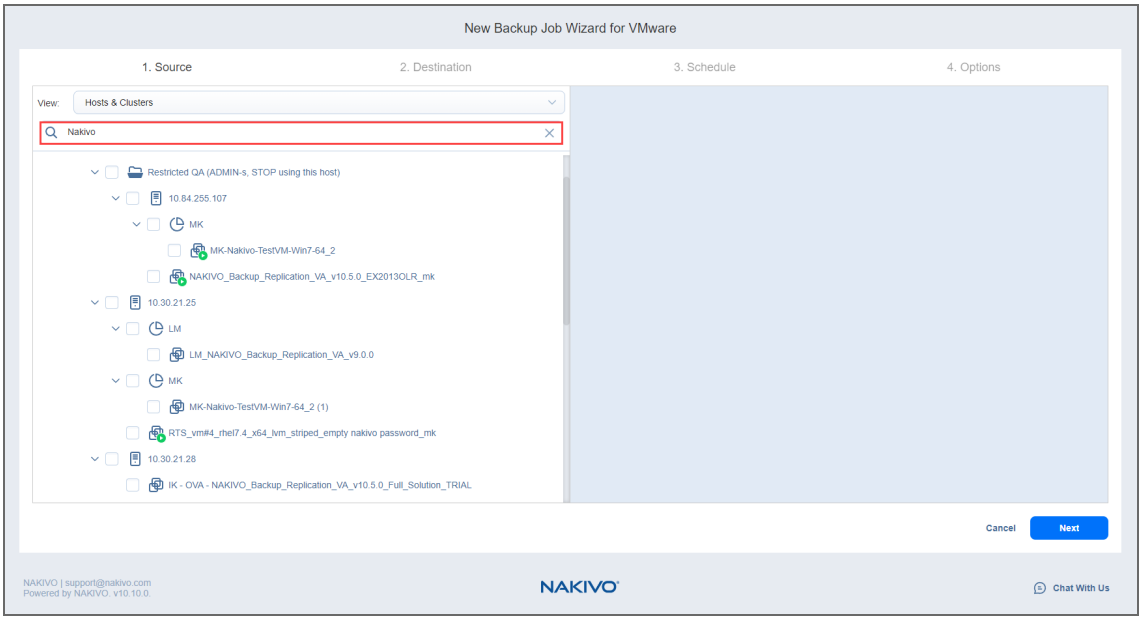
On the **Source** page in the wizard, you can add VMware VMs to your backup job by using one of the inventory views. Proceed as described in the sections below:

- Hosts and Clusters
- VMs and Templates
- Policy

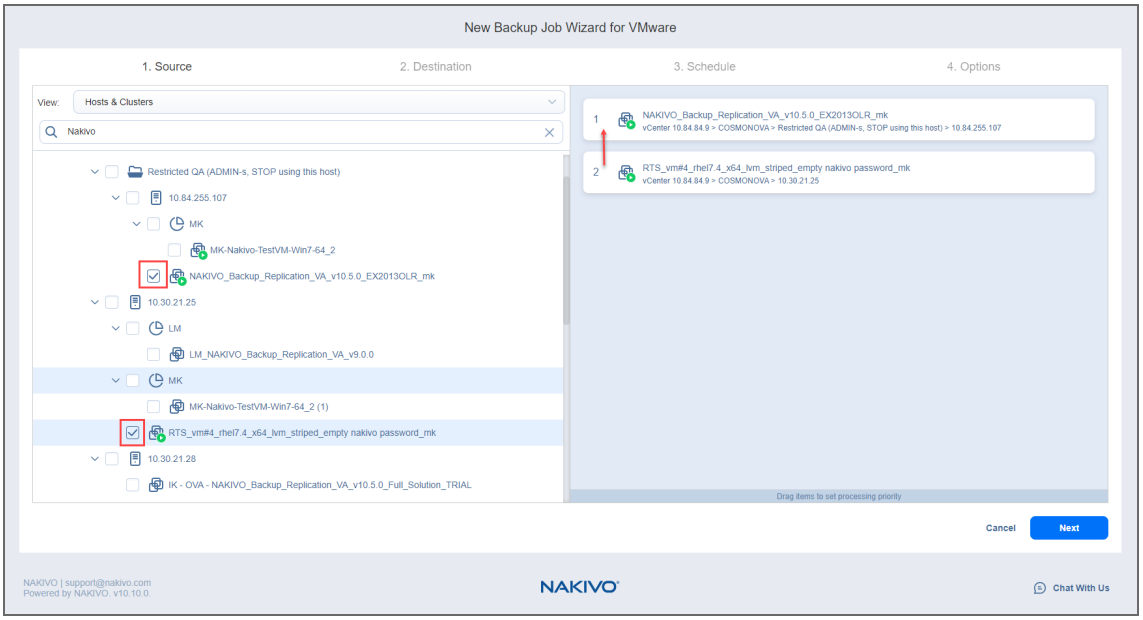
## Hosts and Clusters

When **Hosts & Clusters** view is selected, the inventory tree opens in the left pane and displays all VMware items: clusters, hosts, folders, resource pools, and VMs. Proceed as follows:

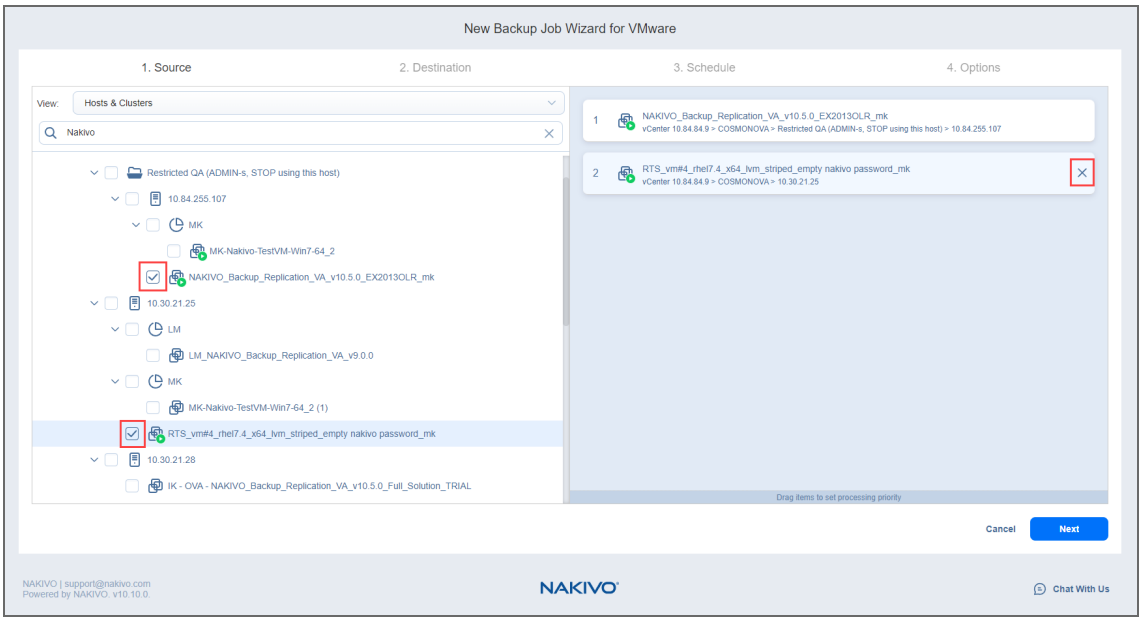
1. Optionally, you can filter the inventory tree by entering a string into the **Search** box. You can enter a part of the entire name of the item.




2. Select VMware items by placing a checkmark to the left of each item.
3. The selected items appear in the right pane of the page. You can drag a VM or a container with the pointer to a new position to reorder the selected items. By doing that, you can specify to back up the most important VMs first.

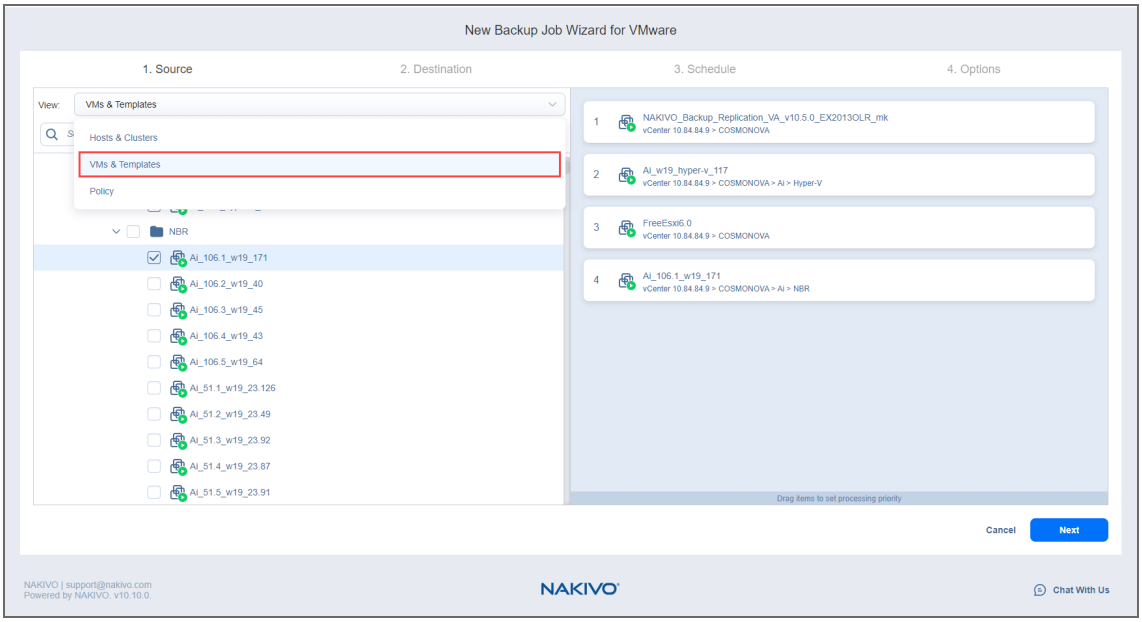



4. Review the list of selected VMware items. You can remove a selected VM or a container from the backup job in one of the following ways:
    - Deselect the object in the left pane. This will remove the object from the right pane; OR
    - In the right pane, hover the pointer over the item you want to remove and click "**X**" icon. This will deselect the object in the left pane.




## VMs and Templates

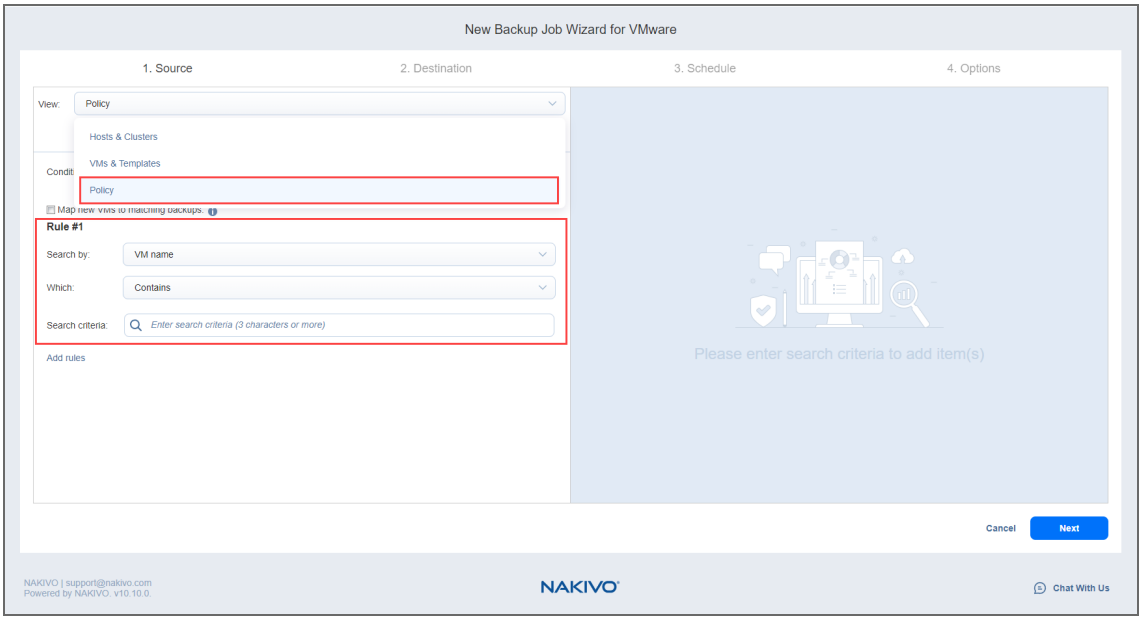
When the **VMs & Templates** view is selected, the inventory tree displays VMware hosts, VMs, and VM templates. Proceed as it was described for the **Hosts & Clusters** view above.




# Policy

When the **Policy** view is selected, job policies can be used. Refer to “Managing Job Policies” on page 382 for details. Follow the steps below:

1. If the items are selected in alternate views, a dialog box opens warning you that switching to the **Policy** view will reset your current selection. Click **Switch View** to confirm switching to the **Policy** view.
2. Add at least one rule to the job policy. Refer to “Managing Policy Rules” on page 385 for details.
3. Click **Next** to confirm adding selected VMs to the backup job. The wizard will display the next page.



**Notes**

- If you cannot find a VM or a container, try the following:
  - Make sure that the corresponding vCenter or ESXi host has been added to the inventory.
  - "Refreshing Inventory" on page 552.
- By adding a VMware container to the job, you ensure that important VMs are always protected. If you add a VMware container to the job:
  - All VMs currently available in the selected container will be backed up.
  - All new VMs that are created in (or moved to) the container in the future will be automatically added to the job and backed up.
- The order in which VMs are backed up is important if the Transporter performing the backup cannot process all VMs of the job simultaneously – either because the Transporter is processing other jobs at the same time or because the number of VM disks in the job exceeds the Transporter's maximum Load specified during the Transporter creation.

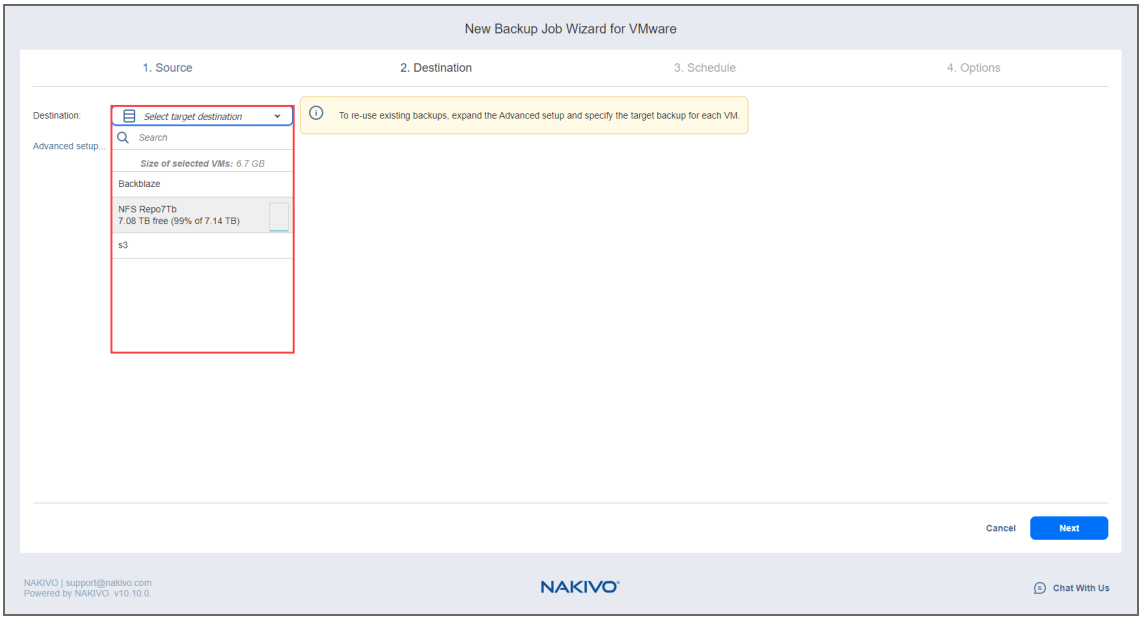
# Backup Job Wizard for VMware: Destination

On the **Destination** page of the wizard, you can select one or different Backup Repositories to back up your VMs. Proceed as described in the sections below:

- Setting a Single Backup Repository for All VMs
- Setting Different Backup Repositories for VMs
- Mapping Source VMs to Existing Backups
- Excluding VM Disks from the Backup Job

## Setting a Single Backup Repository for All VMs

To back up the VMs selected on the previous page to a single Backup Repository, choose a Backup Repository from the **Backup repository** drop-down list.



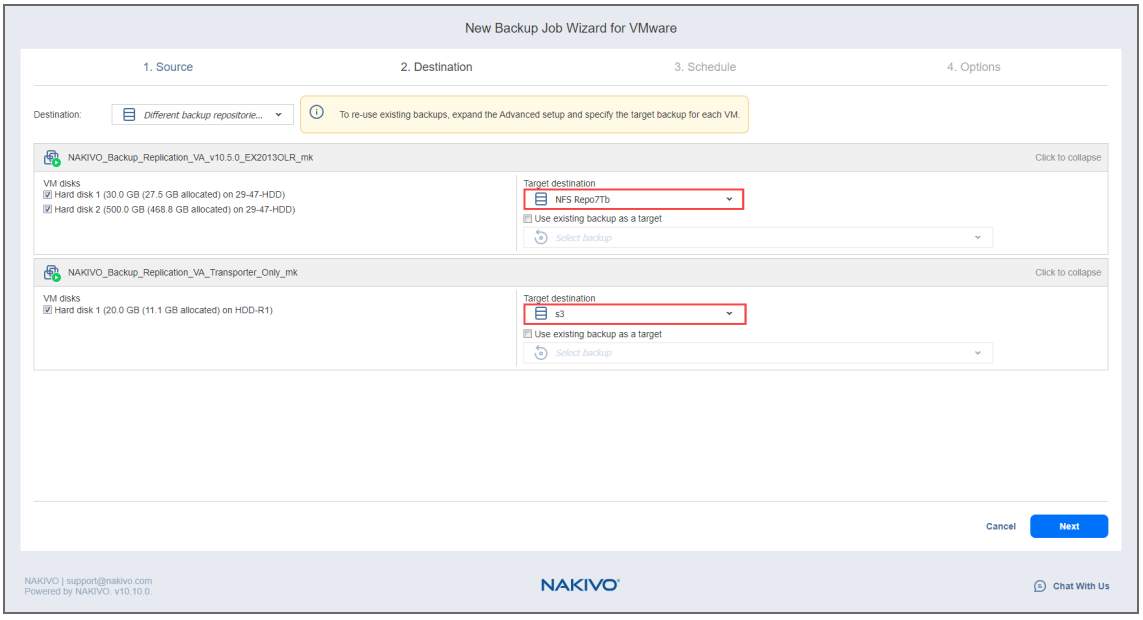

Both federated and standalone (not used as members of federated repositories) can be selected as a destination.

If a federated repository is selected as the destination, the product automatically selects an available qualified federated repository member during the job run.

# Setting Different Backup Repositories for VMs

To back up the selected VMs to different Backup Repositories, follow the steps below:

1. Click **Advanced setup**
2. In the VM boxes, choose a Backup Repository for each VM individually from the Target destination drop-down list.

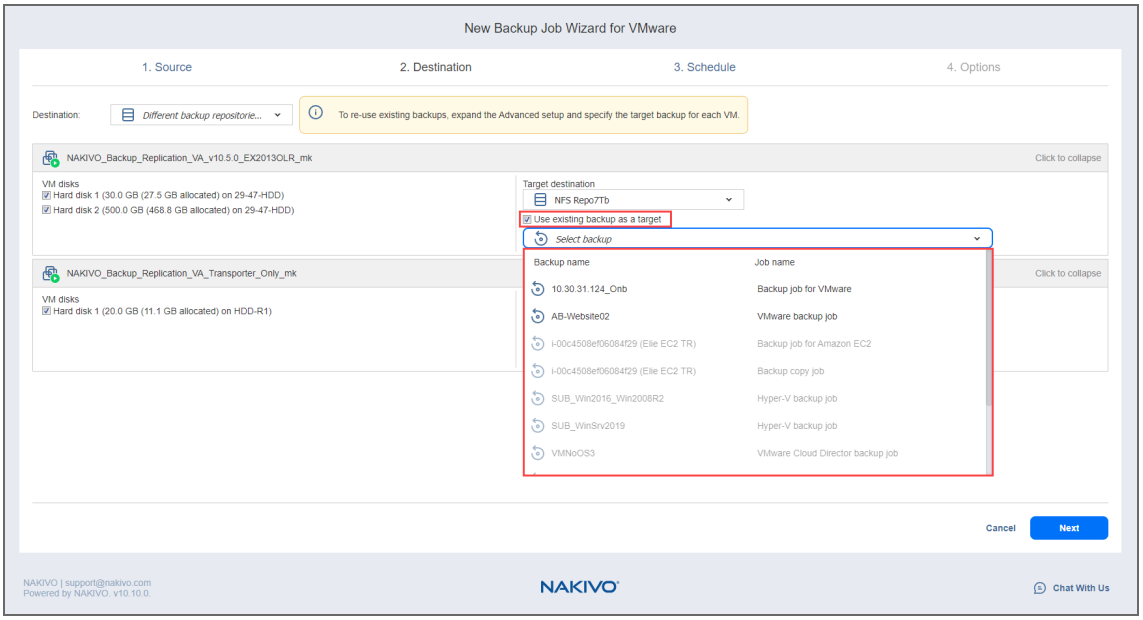


# Mapping Source VMs to Existing Backups

If you have previously backed up a VM and then lost the backup job due to accidental job deletion or a need to recreate jobs in a new copy of the product, you can map source VMs to existing backups in order to avoid running full VM backups again.

To map source VMs to existing backups, follow the steps below:

1. Click **Advanced setup**.
2. From the **Backup repository** drop-down list, choose a Backup Repository that contains an existing VM backup.
3. Select the **Use existing backup as a target** option and choose an existing backup from the drop-down list.




When you run the job, the product will analyze the target VM you have selected, determine how it is different from the source VM, and transfer only the differential data.

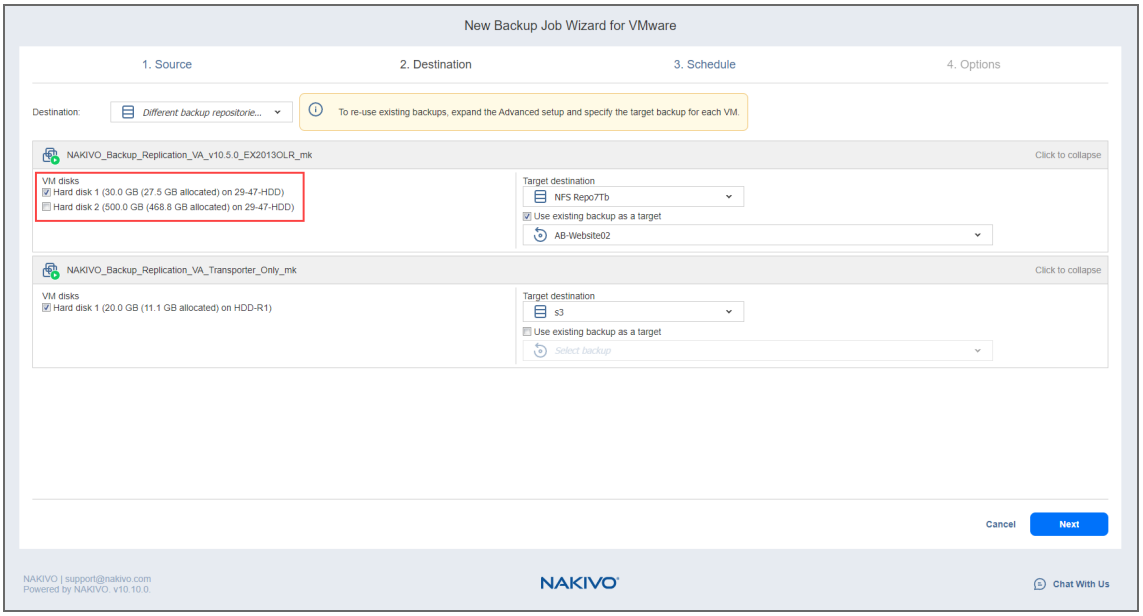
VM backup mapping can be a time-consuming process that can be equal to the time required to create a full VM backup.

After the job is completed, a new recovery point will be created and existing recovery points will not be changed or overwritten.

# Excluding VM Disks from the Backup Job

If you do not want to back up certain VM disks, you can exclude those disks from the backup job by following the steps below:

1. Click **Advanced setup**
2. Cancel the selection of the VM disks that you do not want to back up.




# Backup Job Wizard for VMware: Schedule

On the **Schedule** page of the wizard, select to run the backup job manually or schedule the job to run on a regular basis.

Proceed as described in the sections below:

- Switching to Improved Retention Approach
- Creating New Schedules
  - Weekly
  - Monthly
  - Yearly
  - Periodical
  - After another job
- Creating Legacy Schedules
  - Daily or Weekly Backup
  - Monthly or Yearly Backup
  - Periodic Backup
  - Chained Job

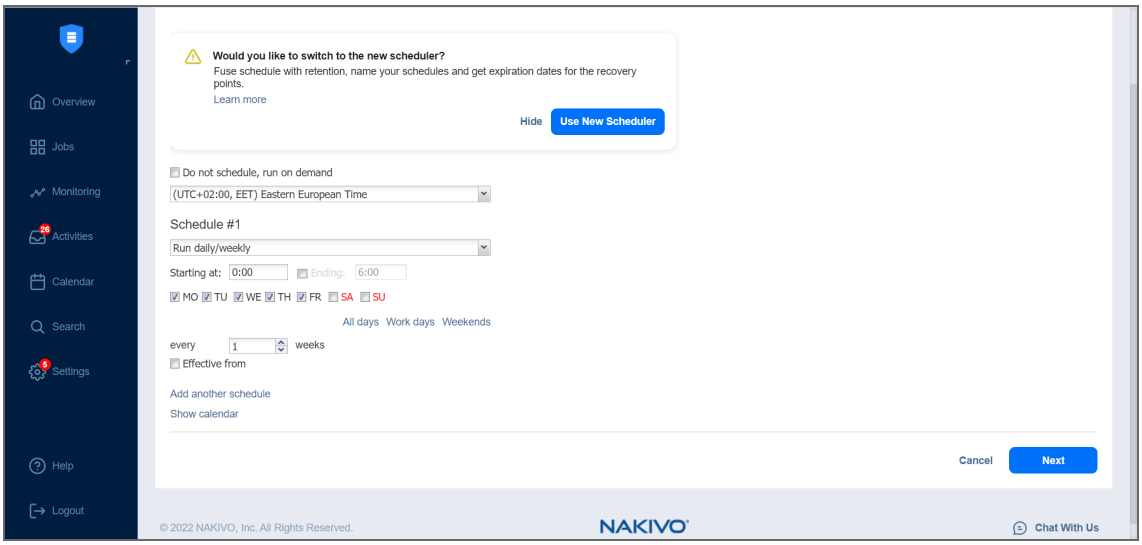
## Switching to Improved Retention Approach

NAKIVO Backup & Replication offers two approaches to retention and scheduling: the legacy or the improved approach. To learn more about how the legacy and improved approaches work, go here. If you create a new job or edit the existing one that uses the legacy approach, a popup appears offering that you to switch to the improved retention approach in the following cases:

- You have updated your instance of the product to v10.8 or later from an older version.
- You have imported a configuration to an instance of NAKIVO Backup & Replication v10.8 or later from an older version.

**Note**

If you install NAKIVO Backup & Replication v10.8 or higher, the new approach is enabled by default.




After the popup appears, do one of the following actions:

- If you do not want to switch to the new scheduler, click **Hide** to close the popup. You can later click **Use New Scheduler** on the **Schedule** page to proceed with the change if you change your mind.
- Alternatively, click **Use New Scheduler** in the popup. Next, choose one of the following options:
  - **MIGRATE SETTINGS**: When you select this option, the existing schedules are automatically converted to new schedules and the existing retention settings are mapped to the new schedules.
  - **CREATE NEW SCHEDULES**: When you select this option, you can create new schedules using the existing retention settings. Old schedules will be deleted.
  - **CONFIGURE SETTINGS ANEW**: Select this option to reset all existing schedules and retention settings and configure them from scratch.

**Notes**

- After switching to the new scheduler, the legacy schedule and retention settings are displayed on the right side of the page.
- After switching to the new scheduler, reverting to the legacy schedule and retention settings is impossible.
- You can learn how expiration dates are assigned to recovery points after migration here.

# Creating New Schedules

Before creating a new schedule, you can optionally enable the following settings:

- **Do not schedule, run on demand**: Enable this option if you want to start the job manually.
- **Prioritize schedules**: When this option is selected, NAKIVO Backup & Replication starts treating schedules based on their priority. The **Yearly** schedule will have higher priority than the **Monthly** schedule, etc. In case 2 or more schedules overlap, the schedules with lower priority will be skipped.
- Optionally, when creating any type of schedule, click **Show Calendar** to show the calendar or **Hide Calendar** to hide it.

When **Amazon EC2**, **Amazon S3**, **Wasabi**, **Azure Blob Storage**, **Backblaze B2 Cloud Storage**, or **Local Folder** is selected as the Backup Repository Type and the only backup destination, you can make recovery points in these repositories immutable during schedule creation. With immutability enabled, the recovery points are immutable and stored using the write-once-read-many (WORM) model. With **Amazon EC2**, **Amazon S3**, **Wasabi**, **Azure Blob Storage**, or **Backblaze B2 Cloud Storage** types of Backup Repository, immutable recovery points cannot be overwritten, deleted, or changed by the root user until the specified period has expired. For Local Folder type of Backup Repository, the root user can still clear immutability.

**Notes**

For the Immutability section to be available, the following conditions must be met:

- **Amazon EC2**, **Amazon S3**, **Wasabi**, **Azure Blob Storage**, **Backblaze B2 Cloud Storage**, or **Local Folder** must be selected for Backup Repository Type on the Destination page of the wizard.
- When **Amazon EC2**, **Amazon S3**, **Wasabi**, **Azure Blob Storage**, or **Backblaze B2 Cloud Storage** is selected as the Backup Repository Type, the bucket or blob container with the repository must have Object Lock or version-level immutability enabled respectively as well as Versioning.
- For **Local Folder** type of Backup Repository, see feature requirements.

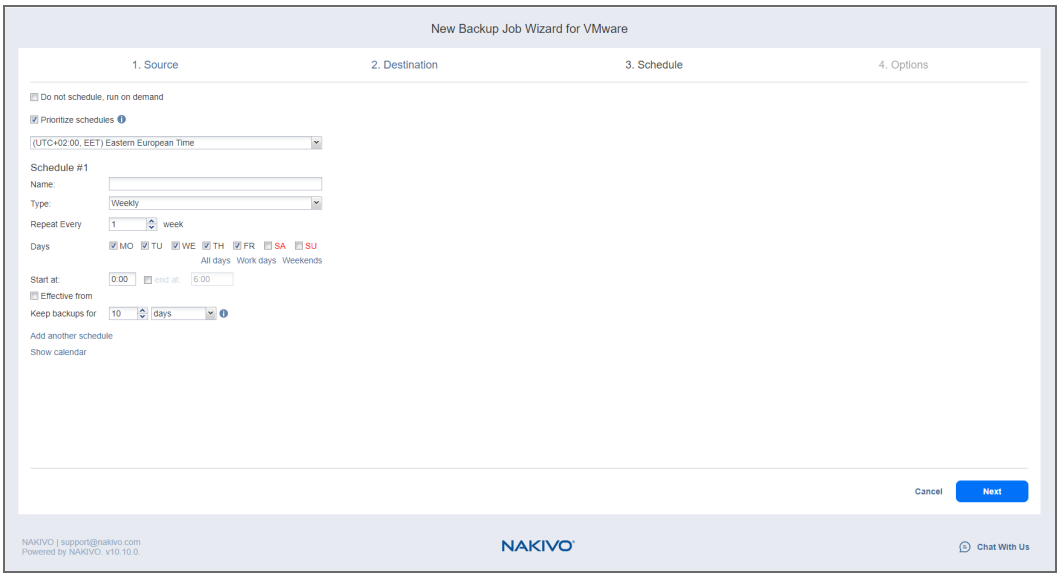
When creating the schedules, you can create schedules of the following types:

## Weekly

You can configure the following options for this schedule type:

- **Schedule title**: Designates the number of a schedule.
- **Name**: Enter the name of your schedule.
- **Repeat every X weeks**: Indicates how often the schedule is repeated.
- **Days**: Select specific days when the schedule executes the job.
- **Start at**: Specify the time when the job should start.
- **End at**: Specify the time when the job should end.

- Optionally, select the **Effective from** checkbox and choose the date when the schedule should come into effect.
- **Keep backups for**: Specify how many days, months, or years NAKIVO Backup & Replication should retain the backups.
- **Immutable for X days**: Enabling this option makes the recovery points immutable for the specified number of days.
- Optionally, click **Add another schedule** if you want to add more than one schedule.




## Monthly

You can configure the following options for this schedule type:

- **Schedule title**: Designates the number of a schedule.
- **Name**: Enter the name of your schedule.
- **Repeat every X months**: Indicates how often the schedule is repeated.
- **Run every**: Select specific days of the month when NAKIVO Backup & Replication executes the job.
- **Start at**: Specify the time when the job should start.
- **End at**: Specify the time when the job should end.
- Optionally, select the **Effective from** checkbox and choose the date when the schedule should come into effect.
- **Keep backups for**: Specify how many days, months, or years NAKIVO Backup & Replication should keep the backups.
- **Immutable for X days**: Enabling this option makes the recovery points immutable for the specified number of days.

- Optionally, click **Add another schedule** if you want to add more than one schedule.




## Yearly

You can configure the following options for this schedule type:

- **Schedule title**: Designates the number of a schedule.
- **Name**: Enter the name of your schedule.
- **Run every**: Select specific days of the specific month when NAKIVO Backup & Replication executes the job.
- Optionally, select the **Effective from** checkbox and choose the date when the schedule should come into effect.
- **Keep backups for**: Specify how many days, months, or years NAKIVO Backup & Replication should keep the backups.
- **Immutable for X days**: Enabling this option makes the recovery points immutable for the specified number of days.

- Optionally, click **Add another schedule** if you want to add more than one schedule.

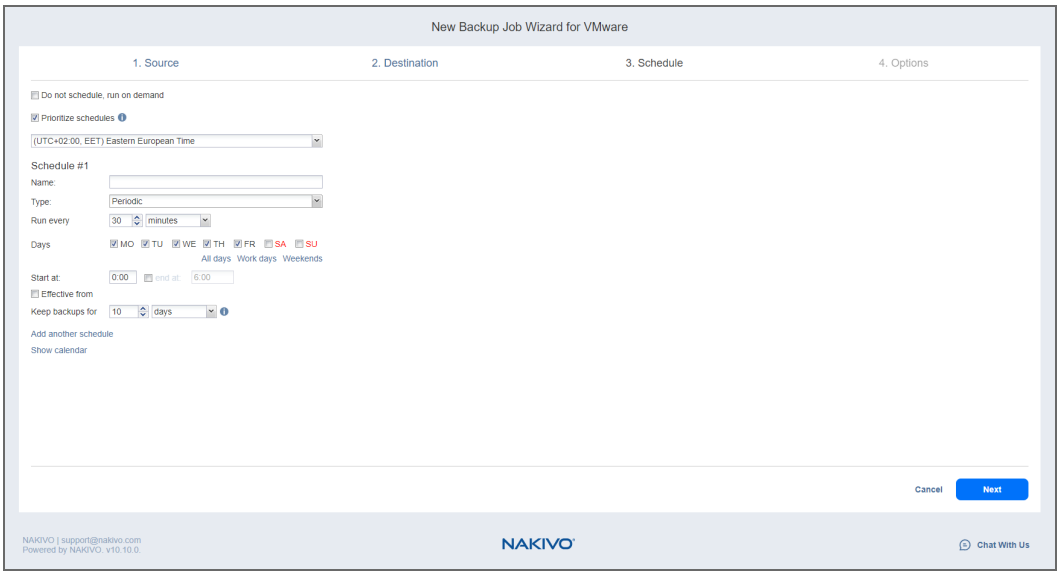



## Periodical

You can configure the following options for this schedule type:

- **Schedule title**: Designates the number of a schedule.
- **Name**: Enter the name of your schedule.
- **Run every**: Select the period measured in minutes, hours, or days when NAKIVO Backup & Replication executes the job.
- **Start at**: Specify the time when the job should start.
- **End at**: Specify the time when the job should end.
- **Days**: Select specific days when the schedule executes the job.
- Optionally, select the **Effective from** checkbox and choose the date when the schedule should come into effect.
- **Keep backups for**: Specify how many days, months, or years NAKIVO Backup & Replication should keep the backups.
- **Immutable for X days**: Enabling this option makes the recovery points immutable for the specified number of days.
- Optionally, click **Add another schedule** if you want to add more than one schedule.
- Optionally, when creating any type of schedule, click **Show Calendar** to show the calendar or **Hide**

**Calendar** to hide it.



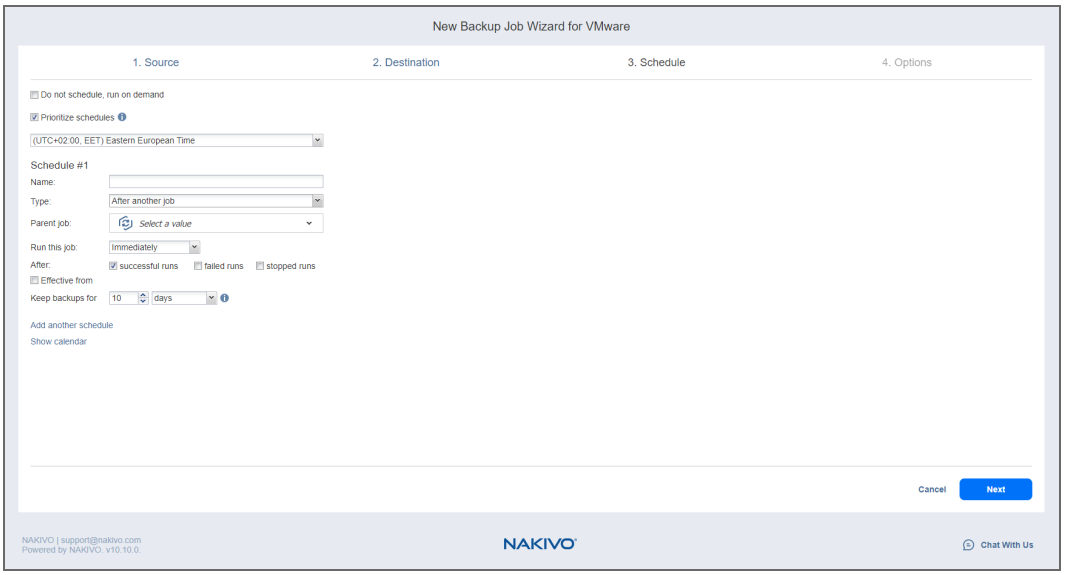

## After Another Job

You can configure the following options for this schedule type:

> **Note**
> This option is disabled if there are no other jobs.

- **Schedule title**: Designates the number of a schedule.
- **Name**: Enter the name of your schedule.
- **Parent job**: Select the job after which this job starts running.
- **Run this job**: Select one of the following options:
  - **Immediately**: The schedule starts right after the parent job is completed.
  - **Delayed**: The schedule starts after the specified number of **minutes** or **hours** following parent job completion.
- Optionally, select the **Effective from** checkbox and choose the date when the schedule should come into effect.
- **Keep backups for**: Specify how many days, months, or years NAKIVO Backup & Replication should keep the backups.
- **Immutable for X days**: Enabling this option makes the recovery points immutable for the specified number of days.
- Optionally, click **Add another schedule** if you want to add more than one schedule.

- Optionally, click **Show Calendar** to show the calendar or **Hide Calendar** to hide it.




# Creating Legacy Schedules

Before creating a new schedule, you can optionally enable the following settings:

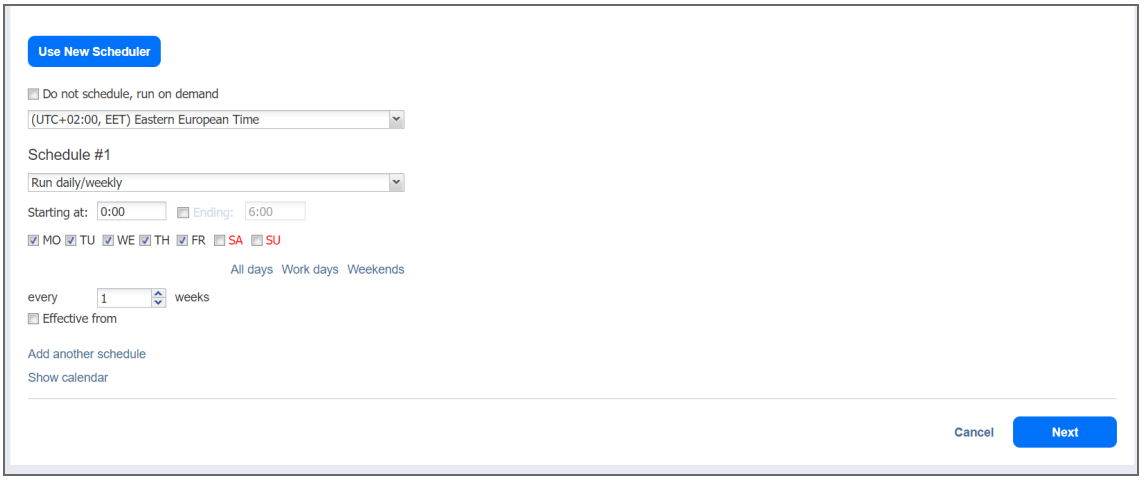
- Click **Use New Scheduler** to switch to the Improved retention approach.
- **Do not schedule, run on demand**: Enable this option if you want to start the job manually.
- Optionally, when creating any type of schedule, click **Show Calendar** to show the calendar or **Hide Calendar** to hide it.
- Optionally, click **Add another schedule** if you want to add more than one schedule.

## Daily or Weekly Backup

To run the job once a day, choose **Run daily/weekly** from the schedule drop-down list and do the following:

- Specify the time when the job should be started in the **Starting** box.
- Specify the end time for the job in the **Ending** box. If the job has not completed by the time specified, the job will be stopped.
- Choose a time zone that should be used for the job start and end times from the time zone drop-down list.
- Select the days of the week during which the job will be started.

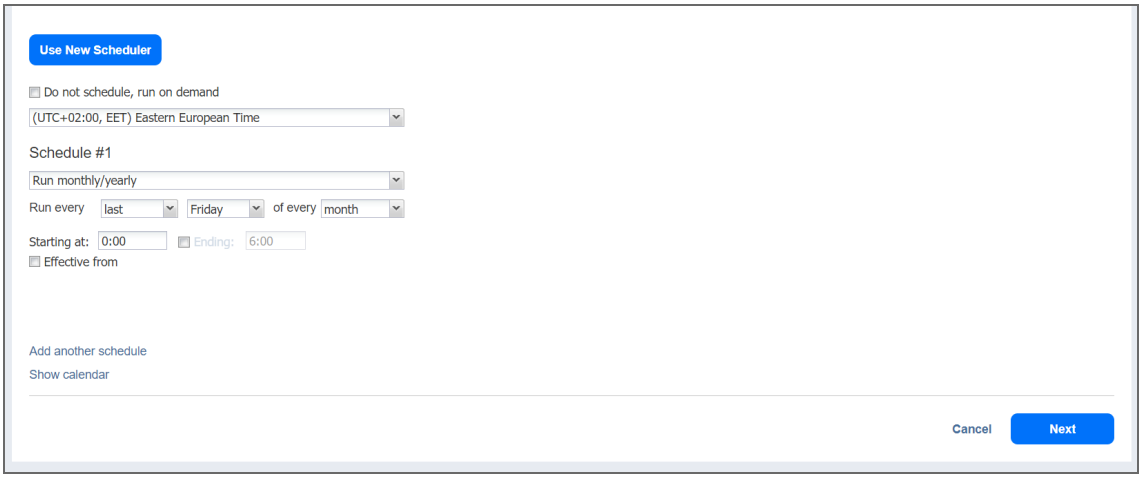
- If necessary, select the **Effective from** checkbox and pick the date when the schedule comes into effect.




## Monthly or Yearly Backup

To run the job monthly or yearly, choose **Monthly/yearly** from the schedule drop-down list and do the following:

- Specify the job start schedule in the appropriate boxes.
- Specify the day and month when the job should be started in the **Run every** boxes.
- Specify the time when the job should be started in the **Starting** box.
- Specify the end time for the job in the **Ending** box. If the job has not completed by the time specified, the job will be stopped.
- If necessary, select the **Effective from** checkbox and pick the date when the schedule comes into effect.




## Periodic Backup

To run the job multiple times per day, choose **Run periodically** from the schedule drop-down list and then choose a time period from the appropriate boxes:

- Specify the time when the job should be started in the **Starting** box.
- Specify the end time for the job in the **Ending** box. If the job has not completed by the time specified, the job will be stopped.
- Select the days of the week during which the job will be started.
- If necessary, select the **Effective from** checkbox and pick the date when the schedule comes into effect.




## Chained Job

To run the job after a previous one has completed, choose **Run after another job** from the schedule drop-down list and set the options as follows:

- **After the job**: Select a job after which the current job will be started.
- **Run this job**: Choose whether to run the current job immediately after the previous one has completed or within a delay.
- **After successful runs**: If selected, the job will run if the previous one has completed successfully.
- **After failed runs**: If selected, the job will run if the previous one has failed.
- **After stopped runs**: If selected, the job will run if the previous one has been stopped.
- **Effective from**: If selected, the schedule will come into effect on the date picked.



# Backup Job Wizard for VMware: Retention

**Important**
This page is not displayed if the new scheduler is enabled.

After each job run, NAKIVO Backup & Replication creates a recovery point for each VM in the Backup Repository. A recovery point represents the backed up VM as of a particular moment in time and allows you to recover individual files, application objects, or the entire VM from the Backup Repository. You can specify how many recovery points you want to be retained in the Backup Repository using the Grandfather-Father-Son (GFS) backup rotation scheme.

When **Amazon EC2**, **Amazon S3**, **generic S3-compatible storage**, **Wasabi**, **Azure Blob Storage**, **Backblaze B2 Cloud Storage**, **Dell EMC DataDomain**, **NEC HYDRAstor**, **NEC HYDRAstor**, or **Local Folder** is selected as the Backup Repository Type for the only backup destination, you can make recovery points in these repositories immutable. With immutability enabled, the recovery points are immutable and stored using the *write-once-read-many* (WORM) model. In case of **Amazon EC2**, **Amazon S3**, **generic S3-compatible storage**, **Wasabi**, **Azure Blob Storage**, **Dell EMC DataDomain**, **NEC HYDRAstor**, **NEC HYDRAstor**, or **Backblaze B2 Cloud Storage** types of Backup Repository, immutable recovery points cannot be overwritten, deleted, or changed by the root user, until the specified period has expired. For **Local Folder** type of Backup Repository, the root user can still clear immutability.

## Retention Settings

Here you can set the retention settings for the backup job. Set the following options:

- **Keep x last recovery points**: Retains the specified number of last recovery points for each VM in the job.
- **Keep one recovery point per day for x days**: Retains one last recovery point per day for the specified number of days.
- **Keep one recovery point per week for x weeks**: Retains the last available backup of every week for the specified number of weeks.
- **Keep one recovery point per month for x months**: Retains the last available backup of every month for the specified number of months.
- **Keep one recovery point per year for x years**: Retains the last available backup of every year for the specified number of years.

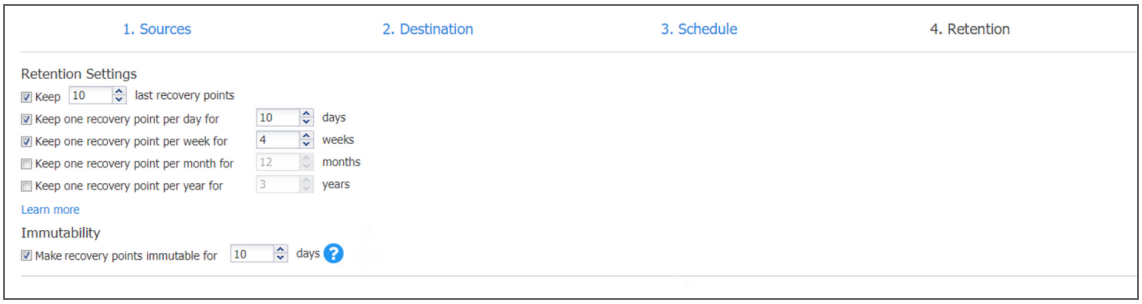
## Immutability

In this section, you can configure the **Make recovery points immutable for x days** option. The recovery points remain immutable for the specified number of days.

**Note**

For the *Immutability* section to be available, the following conditions must be met:

- **Amazon EC2**, **Amazon S3**, **generic S3-compatible storage**, **Wasabi**, **Azure Blob Storage**, **Backblaze B2 Cloud Storage**, **Dell EMC DataDomain**, **NEC HYDRAstor**, **NEC HYDRAstor**, or **Local Folder** must be selected for Backup Repository Type on the Destination page of the wizard.
- If **Amazon EC2**, **Amazon S3**, **generic S3-compatible storage**, **Wasabi**, **Azure Blob Storage**, **Dell EMC DataDomain**, **NEC HYDRAstor**, **NEC HYDRAstor**, or **Backblaze B2 Cloud Storage** is selected as the Backup Repository type, Object Lock or version-level immutability support and Versioning must be enabled bucket or blob container respectively where your Backup Repository is located.
- For **Local Folder** type of Backup Repository, see feature requirements.




For more details and an example of job retention settings, refer to the Keeping Recovery Points article in the Knowledge Base.

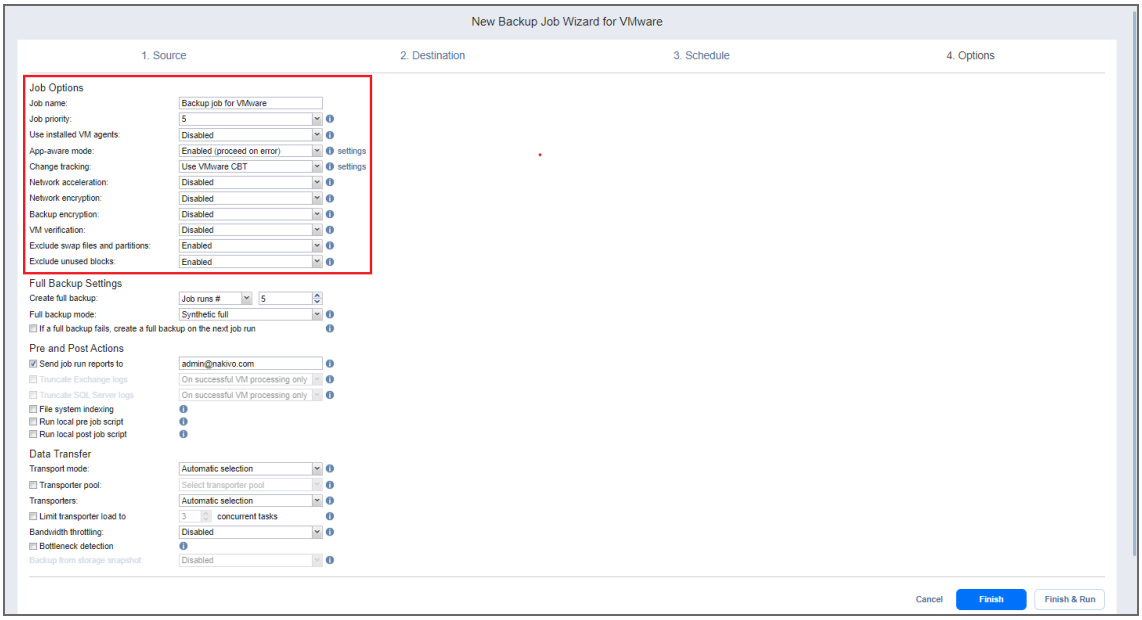
# Backup Job Wizard for VMware: Options

On the **Options** page of the wizard, you can set up job options. Proceed as described in the sections below:

- Job Options
- Full Backup Settings
- Pre and Post Job Actions
- Data Transfer
- Completing the New Backup Job Wizard for VMware

## Job Options

In this section, you can specify a name and a priority level for the backup job and enable/disable app-aware mode, change tracking, network acceleration, encryption, VM Verification, and other options. Proceed as described below.




## Job Name

Enter a name for the backup job in the appropriate box.

## Job Priority

Select a job priority level between 1 and 5, with 1 being the highest priority. Jobs with higher priority levels are prioritized by Transporters during job processing.

> **Note**
> This option is only available in the Enterprise, Enterprise Essentials, Enterprise Plus, MSP Enterprise, and MSP Enterprise Plus editions.

## Use Installed VM Agents

Enable this option if you wish to use installed VM agents for this job. Proceed as follows:

1. With the option enabled, click **settings** to open the **Manage VM Agents** pop-up menu.
2. Click **Scan All** to scan each VM in the job for installed VM agents.
3. Once the scan is complete, select the master password you wish to use for the discovered VM agents from the **Select credentials** drop-down list.
4. In the **Certificate** column, click **Verify** to verify the validity of a VM agent's certificate.

For more information on setting up VM agents, see "Using a VM Agent" on page 604. For more information on setting up master passwords, see "Managing Credentials" on page 559.

## App-Aware Mode

When the app-aware mode option is enabled, VM backup is performed using VMware Guest OS quiescing, which relies on Microsoft VSS to ensure that application data is consistent.

## Change Tracking

Select one of the options from the **Change tracking** drop-down list:

- **Use VMware CBT**: When selected, NAKIVO Backup & Replication enables the VMware Changed Block Tracking feature for source VMs. This feature enables the product to quickly identify the data blocks that have changed since the last job run, which significantly increases the job speed. Click the **error handling** link to specify the following options:
- **On error**: Choose one of the following job behaviors in relation to CBT error:
    - **switch to proprietary method**: If VMware CBT fails to provide data on changed blocks for a VM and this option is selected, NAKIVO Backup & Replication performs an incremental backup of the VM using the NAKIVO proprietary change tracking technology.
    - **reset CBT**: If VMware CBT fails to provide data on changed blocks for a VM and this option is selected, NAKIVO Backup & Replication resets VMware CBT for the VM.
    - **fail VM processing**: If VMware CBT fails to provide data on changed blocks for a VM and this option is selected, NAKIVO Backup & Replication does not process the VM and states job failure (other VMs in the job will be processed).
- **Double-check changed blocks provided by CBT**: When selected, NAKIVO Backup & Replication runs a check on data blocks provided by VMware CBT to ensure that VMware CBT does not overstate the amount of changed data.
- **Use proprietary method**: When selected, NAKIVO Backup & Replication performs incremental backups using the NAKIVO proprietary change tracking technology. This feature requires reading the contents of all VM disks to determine which data blocks have changed since the last job run.
- **No change tracking (always full)**: When selected, NAKIVO Backup & Replication always performs a full VM backup of all source VMs.

## Network Acceleration

When enabled, NAKIVO Backup & Replication uses compression and traffic reduction techniques to speed up data transfer. Select this option if you plan to back up over WAN or slow LAN links.

> **Note**
> The **Network acceleration** option is not available if the **Backup encryption** option is enabled.

## Network Encryption

When enabled, VM data is protected with AES 256 encryption while traveling over the network.

- Data encryption increases the backup time and CPU load on machines running Transporters. Select this option if you back up over WAN without a VPN connection.
- You need at least one **Transporter** at source and target sites to enable network encryption.

## Backup Encryption

When enabled, backup data is protected with AES 256 block cipher encryption with a 256-bit key length. You can protect the backup file by creating a new password or selecting an existing one. For more information, refer to Enabling Backup Encryption.

**Notes**

- This option is available only if the **Disk** or **Tape** destination type was chosen on the **Destination** page of the wizard.
- The **Backup encryption** option is not available if the **Network acceleration** option is enabled.
- If enabled, the created recovery points are encrypted.
- The **Backup encryption** option is not displayed for a backup job where forever incremental repositories are selected as the only target repositories.
- The **Backup encryption** option cannot be enabled if multiple targets with a mix of supported and unsupported (SaaS repositories or forever incremental repositories) repositories are selected as destinations.
- It's recommended that you enable the **(AWS) Key Management Service**. If **KMS** is enabled, all backup encryption passwords encrypted with the Key Management Service cryptographic key are available for recovery in case of product re-installation. For more information, refer to "Enabling KMS" on page 429.

## VM Verification

When VM verification is enabled, the NAKIVO Backup & Replication checks the integrity of the backup by starting it and interacting with it. You can choose one of the following VM verification options:

- **Disabled**: VM Verification is disabled.
- **Screenshot verification**: When enabled, all VM backups created by the job are verified: After a backup of a VM is completed, the VM is recovered from the backup using Flash boot (and is disconnected from networks) and a screenshot of the recovered VM is taken once the VM OS has booted, after which the VM is discarded. VM screenshots are included in email notifications (if they have been configured) and displayed on the **Dashboard**.
- **Boot verification**: When enabled, all VM backups created by the job are verified as follows. After a VM backup is completed, NAKIVO Backup & Replication recovers the VM using Flash boot, disables networking to prevent network connections, and verifies that system start is successful.

**Important**

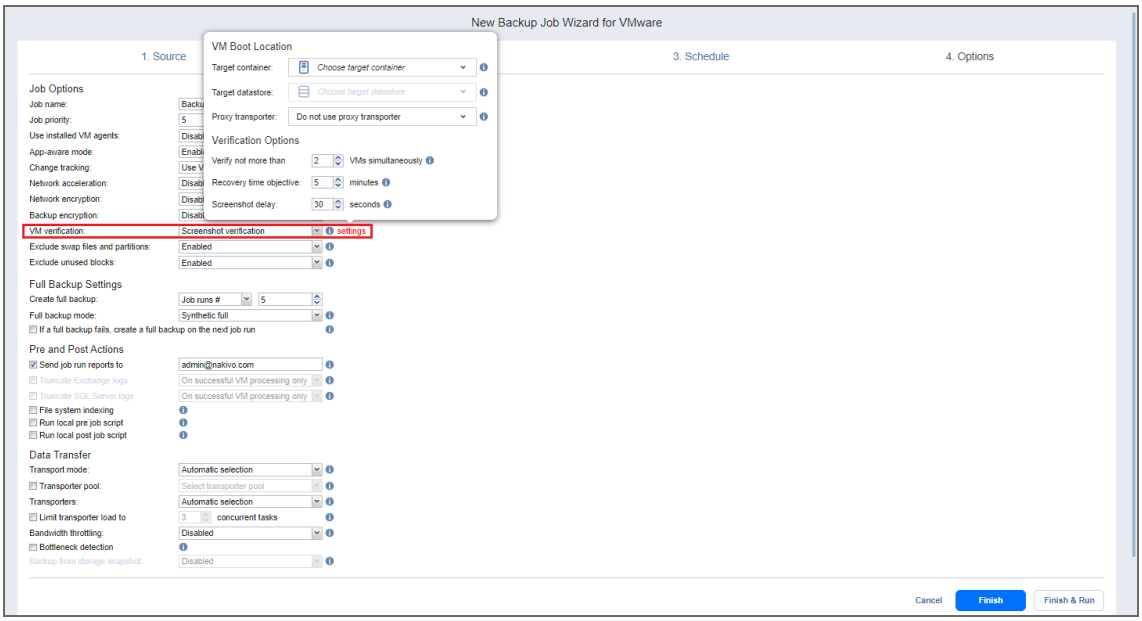
VM verification requires VMware Tools to be installed on all VMs.

After choosing **Screenshot verification**, do the following in the dialog box that opens:

1. Provide a location of the VMs to be booted:
    a. **Target Container**: Choose a target container (cluster, host, or resource pool) where VMs are run using Flash boot.
    b. **Target Datastore**: Choose a datastore that hosts changes to the recovered VMs.
    c. **Proxy transporter**: Choose a proxy **Transporter** from the list of available **Transporters**.

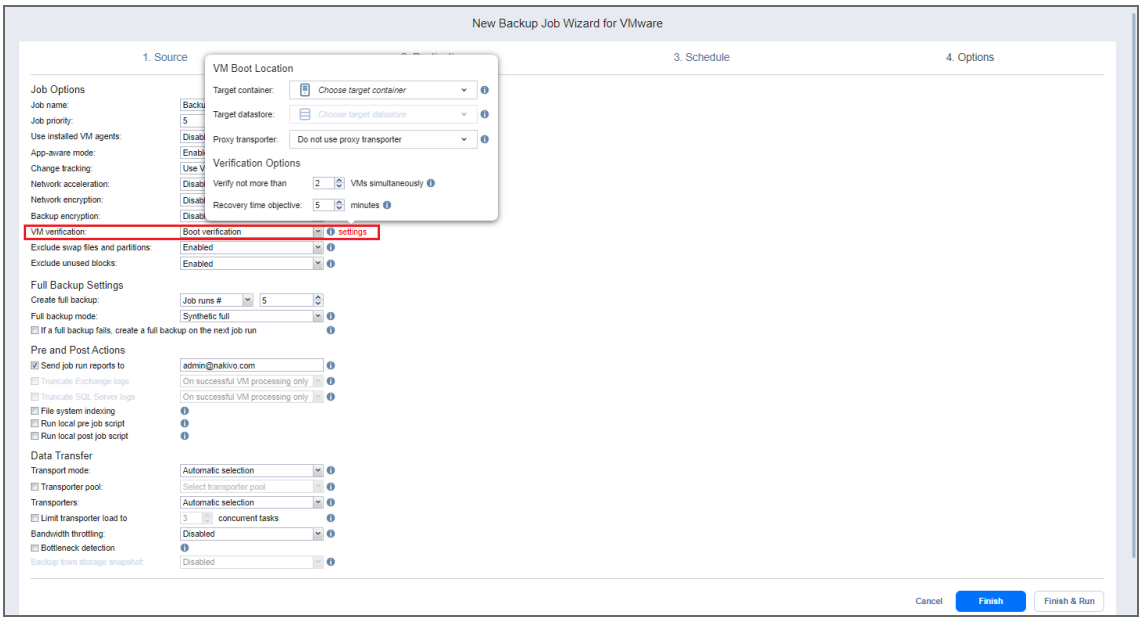
> **Notes**
>
> NAKIVO Backup & Replication uses a proxy **Transporter** in the following cases:
>
> - The Transporter assigned to the **Backup Repository** cannot use iSCSI port 3260 because it is occupied by other services.
> - iSCSI packages are missing on the **Transporter** assigned to the B**ackup Repository**.

2. Set verification options:
    a. **Verify not more than x VMs simultaneously**: Specify the maximum number of VMs that can be started on the target container simultaneously.
    b. **Recovery time objective**: Specify the amount of time allocated for verification of each VM backup. If a VM OS does not start within the specified amount of time, verification is considered failed.
    c. **Screenshot delay**: The amount of time that the product should wait after the guest OS start before making a screenshot. The specified time must be sufficient to fully start the VM OS. Try increasing this amount if the default amount is insufficient.




After choosing **Boot verification**, do the following in the dialog box that opens:

1. Provide a location of the VMs to be booted as described for the **Screenshot verification** option.
2. Set verification options:
    a. **Verify not more than x VMs simultaneously**: Specify the maximum number of VMs that can be started on the target container simultaneously.
    b. **Recovery time objective**: Specify the amount of time allocated for the verification of each VM backup. If a VM OS does not start within the specified amount of time, verification is considered failed.




## Exclude Swap Files and Partitions

With this option enabled, NAKIVO Backup & Replication automatically excludes swap files and partitions during the backup process.

## Exclude Unused Blocks

With this option enabled, NAKIVO Backup & Replication automatically excludes unused disk blocks and blocks occupied by deleted files during processing of source objects running Windows OS. This feature allows for reducing backup storage space and object processing time.

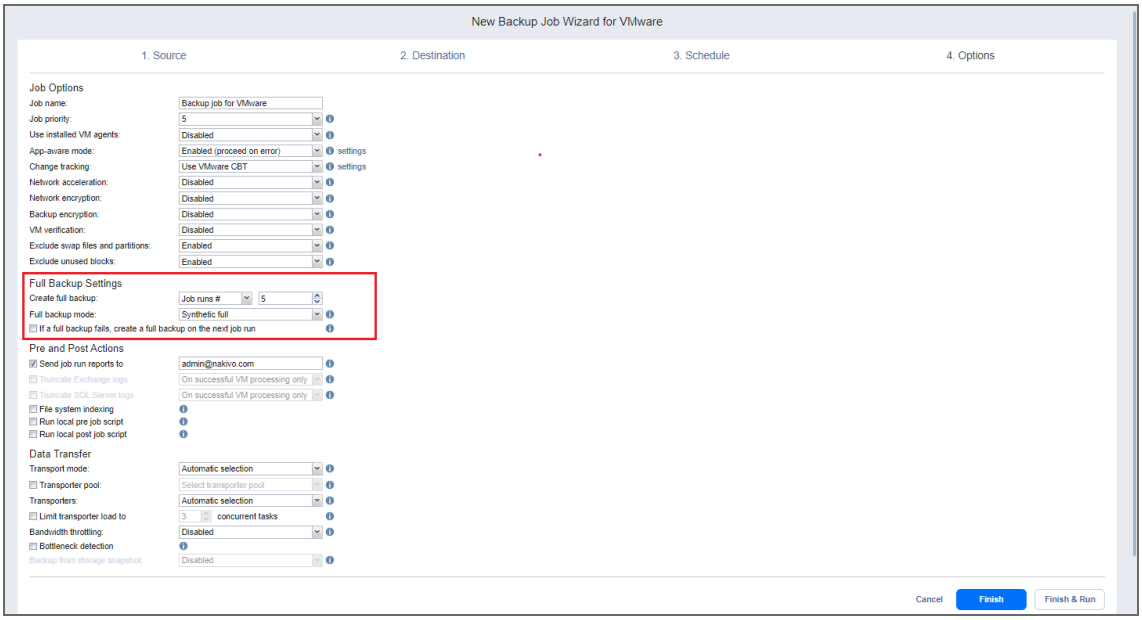
# Full Backup Settings

If the type of the **Backup Repository** that you've selected on the **Destination** page of the wizard is set to **Incremental with full backups** (**Store backups in separate files** option is selected), you can specify the following options:

- **Create full backup**: Specify how often full backups should be created.
- **Full backup mode**: Specify how the full backup should be created. You can choose one of the following options:

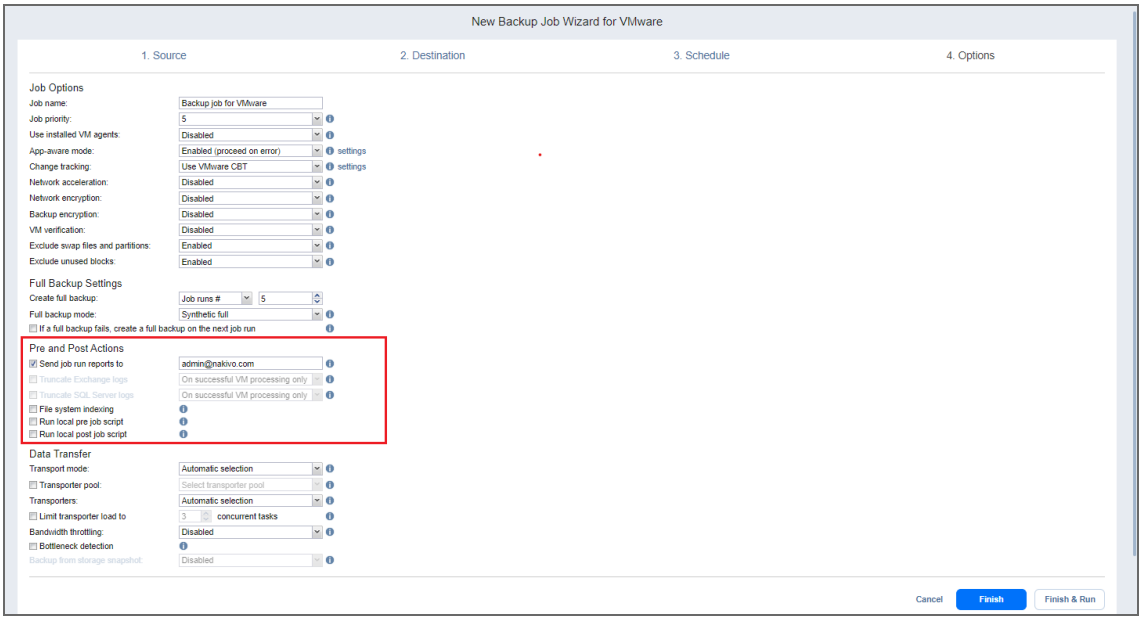
- **Synthetic Full**: When this option is selected, NAKIVO Backup & Replication first performs an incremental backup (that is, transfers only the data that has changed since the last backup) and then transforms the available data into a full backup file. The benefits of this approach are:
  - The Synthetic Full backup is usually faster than the Active Full backup.
  - The load on the network is lower as less data is transferred.
  - The load on the source datastores running your production VMs is lower.
- **Active Full**: When this option is selected, NAKIVO Backup & Replication reads all VM data from the source datastore and transfers it to the Backup Repository.

- **If a full backup fails, create a full backup on the next job run**: With this option selected, the next job run creates a full backup if the current job run fails to do so.




# Pre and Post Job Actions

NAKIVO Backup & Replication allows you to enable certain actions before a backup job begins and after it has completed. You can choose to send job run reports, truncate Microsoft Exchange and Microsoft SQL Server logs on specified VMs, and run local pre- and post- job scripts.




## Email Notifications

NAKIVO Backup & Replication can send email notifications about job completion status to specified recipients. This feature complements global notifications and allows you to configure notifications on a per-job level.

To enable this option, configure your Email settings. To send email notifications, select **Send job run reports to** and specify one or more email addresses in the text box. Use semicolons to separate multiple email addresses.

## Truncation of Microsoft Exchange Server Transaction Logs

Microsoft Exchange Server database transaction logs record all changes to an Exchange Server database. Over time, these log files accumulate and can consume all of the available disk space if not periodically removed. NAKIVO Backup & Replication provides an option to delete (aka truncate) Microsoft Exchange Server logs on the source VMs after job completion.

The transaction logs are deleted after the job is completed so that the log files are available in the VM backup. Note that the product deletes only those transaction logs which are already committed to (available in) the Microsoft Exchange database.

> **Note**
>
> Truncation is not supported for items added using **Direct Connect**.

To set up Microsoft Exchange log truncation, do the following:

1. Select the **Truncate Exchange logs** option.
2. In the dialog box that opens, select the checkboxes next to the VMs running Microsoft Exchange and then select the credentials next to each VM. These credentials are used to log in to the VMs that you have selected.

## Truncation of Microsoft SQL Server Transaction Logs

Microsoft SQL Server database transaction logs record all changes to an SQL Server database. Over time, these logs accumulate and can consume all of the available disk space if not periodically removed. NAKIVO Backup & Replication provides the option to delete (or truncate) Microsoft SQL Server logs on the source VMs after job completion.

The transaction logs are deleted after the job is completed so that the original log records are available in the VM backup. Note that the product deletes only those transaction logs that are already committed to (available in) the Microsoft SQL Server database.

> **Note**
>
> Truncation is not supported for items added using **Direct Connect**.

To set up Microsoft SQL Server log truncation, do the following:

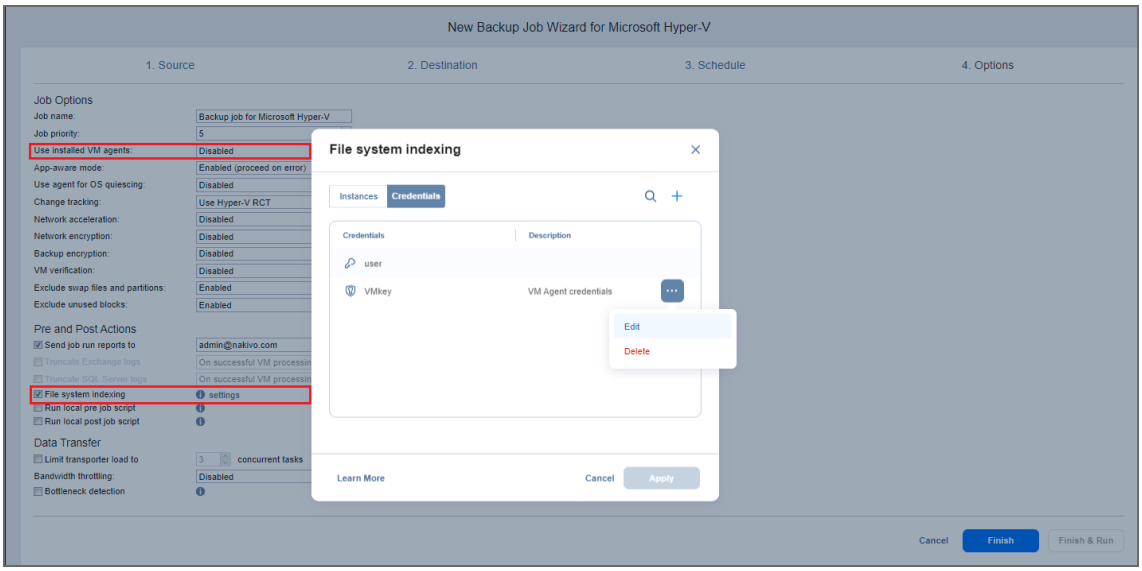
1. Select the **Truncate SQL Server logs** option.
2. In the dialog box that opens, select the checkboxes next to the VMs running Microsoft SQL Server and then select credentials next to each VM. These credentials are used to log in to the VMs that you have selected.

## File System Indexing

Select this option to create an index of files and folders during the backup job run.

To use a temporary agent for file system indexing:

1. Make sure that the **Use installed VM agents** option is disabled in the *Job Options* section.
2. Enable the **File system indexing** option and click the **settings** link to open the **File system indexing** popup.
3. In the **Instances** tab, select the appropriate credentials for each VM for which you would like to enable file system indexing in the Select credentials column. You may also manage your existing credentials or create new ones in the **Credentials** tab. See “Managing Credentials” on page 559 for more information.




**Notes**

- The source VM must be turned on during backup for data to be indexed.
- If the source VM contains files/folders a user has no access to, data about such files/folders is not indexed.
- If the backup job fails or is stopped, or the existing recovery point is removed, the index is deleted.
- If the original backup is removed but a copy is retained, the indexed data is retained.
- If the existing backup job is removed but the backup is retained, the indexed data is also retained.
- Indexing starts with VM processing within a running backup job.
- Indexed data is not duplicated when a backup copy to tape job is created.
- If the range of backups subject to scanning contains an encrypted recovery point that cannot be decrypted automatically, the **Cannot automatically scan backups** dialog box is displayed before scanning.
    - Click the **Cancel** or **X** button to close the dialog box without applying any changes.
    - Click **Proceed** to start scanning.

  Refer to Managing Backup Encryption for more details.

## Pre Job Script

To run a script before the product begins backing up VMs, do the following:

1. Place a script file on the machine on which the **Director** is installed.
2. Select the **Run local pre job script** option.
3. Specify the following options in the dialog box that opens:

- **Script path**: Specify a local path to the script on the machine on which the **Director** is installed. A script interpreter should be specified.
  **Example (Windows)**: *cmd.exe /c D:\script.bat*
  **Example (Linux)**: *bash /root/script.sh*
- **Job behavior**: Choose one of the following job behaviors in relation to script completion:
  - **Wait for the script to finish**: When this option is selected, VM backup is only started after the script is completed.
  - **Do not wait for the script to finish**: When this option is selected, the product runs the script and starts backing up VMs at the same time.
- **Error handling**: Choose one of the following job behaviors in relation to script failure:
  - **Continue the job on script failure**: When this option is selected, the job performs VM backup even if the script has failed.
  - **Fail the job on script failure**: When this option is selected and the script fails, the job is failed and VM backup is not performed.

## Post Job Script

To run a script after the product has finished backing up all VMs, do the following:

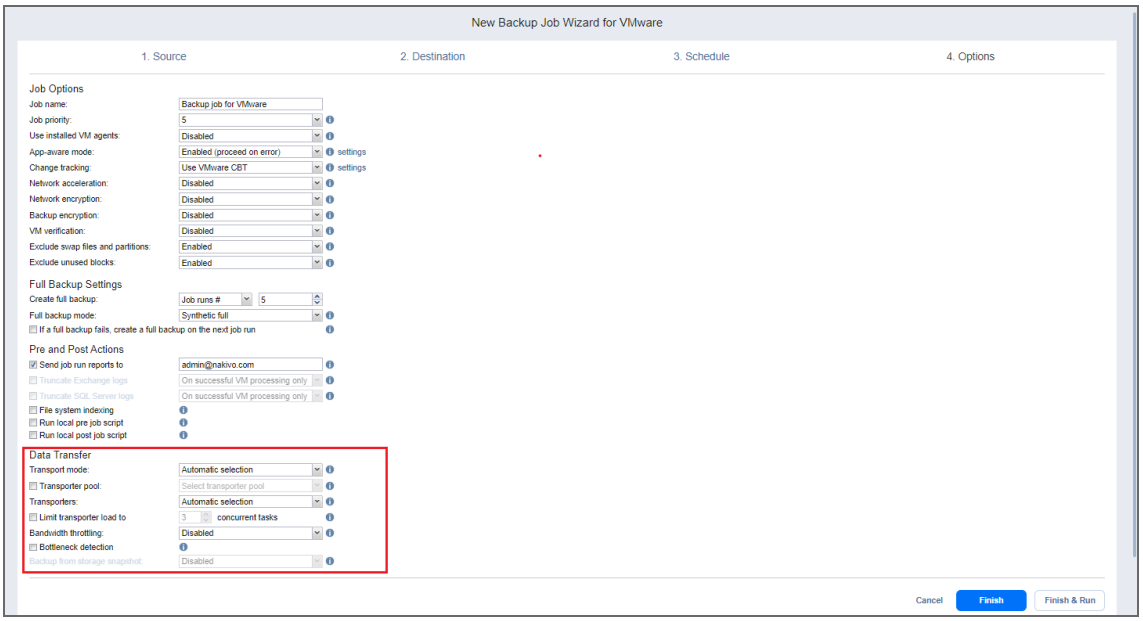
- Place a script file on the machine on which the **Director** is installed.
- Select the **Run local post job script** option.
- Specify the following options in the dialog box that opens:
  - **Script path**: Specify a local path to the script on the machine on which the **Director** is installed. A script interpreter should be specified.
    **Example (Windows)**: *cmd.exe /c D:\script.bat*
    **Example (Linux)**: *bash /root/script.sh*
  - **Job behavior**: Choose one of the following job behaviors in relation to script completion:
    - **Wait for the script to finish**: When this option is selected, the job is in the “running” state until the script is completed.
    - **Do not wait for the script to finish**: When this option is selected, the job is completed even if the script execution is still in progress.
  - **Error handling**: Choose one of the following job behaviors in relation to script failure:
    - **Continue the job on script failure**: When this option is selected, script failure does not influence the status of the job.
    - **Fail the job on script failure**: When this option is selected and the script fails, the job status is set to “failed” even if VM backup is successful.

**Important**

Pre- and post-job scripts can be executed only on the machine on which the **Director** is installed.

# Data Transfer

In the *Data Transfer* section of the **Options** page, you can choose a transport mode, select a **Transporter** to be used for reading data from source VMs, specify a b load, and configure bandwidth throttling.




## Transport Mode

In the *Data Transfer* section, choose a transport mode for retrieving VM data:

- **Automatic selection**: When this option is selected, NAKIVO Backup & Replication automatically selects the best transport mode available:
  - When the source **Transporter** is installed on a VM, NAKIVO Backup & Replication tries to use transport modes in the following order: Hot Add > SAN > LAN.
  - When the source **Transporter** is installed on a physical machine, NAKIVO Backup & Replication tries to use transport modes in the following order: SAN > LAN. Hot Add is not supported for physical machines.
- **SAN only**: When this option is selected, NAKIVO Backup & Replication only uses direct SAN access to retrieve VM data. If direct SAN access to VM data is not available, the job fails.
- **Hot Add only**: When this option is selected, NAKIVO Backup & Replication only uses Hot Add to retrieve VM data. If direct Hot Add is not available, the job fails.

**Notes**

- When hot-add cannot be enabled for at least one disk of a source VM (even if the disk is deselected in the job), Hot Add is unavailable for all disks of the VM.
- Do not create a manual snapshot of the transporter VM housing NAKIVO Backup & Replication while the Hot Add disks are attached it. In case you need to create a snapshot or restart a VM, make sure to stop all jobs and then detach all Hot Add disks in NAKIVO Backup & Replication.

- **LAN only**: When this option is selected, NAKIVO Backup & Replication only retrieves VM data over LAN.

## Transporter Pool

If this option is enabled, only the transporters that belong to the selected transporter pool shall be used during the job run.

## Transporters

In NAKIVO Backup & Replication, **Transporters** perform all of the heavy lifting: reading data from the source VM, compressing the data, transferring the data over the network, performing data deduplication, and so on. In large and geographically distributed environments, multiple **Transporters** can be deployed to distribute the data protection workload, optimize network traffic, and improve data transfer speeds. Thus, if more than one **Transporter** is deployed for NAKIVO Backup & Replication, it is important to determine which should be used to read data from a particular source VM (the target Transporter is always the one that is assigned to the **Backup Repository**).

By default, the product automatically determines which **Transporter** should be used to read data from the source VM. However, you can manually specify which **Transporters** should be used for the job. To manually specify **Transporters**, proceed as follows:

In the *Data Transfer* section, choose one of the following options:

- **Automatic selection**: The product automatically determines which Transporters are the closest to source hosts (the hosts that run selected VMs) and uses those Transporters to retrieve data from source VMs.
- **Manual - configured for all VMs**: Select this option to manually specify a single **Transporter** to be used to retrieve data from source VMs.
- **Manual - configured per host**: Select this option to manually specify which **Transporter** should be used to retrieve data from each source host.

Replacement Transporters

When configuring NAKIVO Backup & Replication, you can manually select a primary and replacement Transporter for each separate host or set them to be selected automatically.

**Note**
This option becomes available only if you select **Manual - configured per host** from the **Transporters** drop-down list.

You have the following options:

- **Select automatically**: When selected, the replacement **Transporters** are selected automatically for each host.
- **Use only primary transporters**: When selected, only primary **Transporters** are used during job execution.
- **Select manually per host**: When selected, the **Transporters** can be selected manually or can be set to be chosen automatically for each separate host.

## Transporter Load

You can limit the maximum number of **Transporter** tasks used by the job. By default, this number is set to 3 concurrent tasks.

To change the default number of tasks, do the following:

1. In the *Data Transfer* section, select the **Limit transporter load to** checkbox.
2. Specify the number of concurrent tasks in the corresponding box.

## Bandwidth Throttling

Follow the steps below to regulate the speed of data transfer over the network for your backup job:

1. For the **Bandwidth throttling** option, choose **Enabled**.

   **Note**
   If bandwidth throttling is disabled for the current job, global bandwidth rules may still apply to your job. Refer to "Bandwidth Throttling" on page 407 for details.

2. Click the **settings** link that becomes available.
3. The **Job Bandwidth Rules** dialog box opens, displaying the list of available rules. You have the following options:
   - Create a new bandwidth rule for your backup job:
     a. Click the **Create New Rule** button.
     b. The **New Bandwidth Rule** dialog box opens. Refer to the "Bandwidth Throttling" on page 407 topic for details on creating a bandwidth rule.
     c. Click **Save**.
   - Activate an existing bandwidth rule for your job. Select the checkbox to the left of the required bandwidth rule. To deactivate a bandwidth rule for your job, deselect the corresponding checkbox.

- Edit a bandwidth rule. Click the **Edit** link for a bandwidth rule and modify it in the **Edit Bandwidth Rule** dialog box that opens.
- Disable a bandwidth rule. Click the **Disable** link. The bandwidth rule is disabled for all jobs.
- Remove a bandwidth rule. Click the **Remove** link and then click **Delete** to confirm your operation.

### Backup from Storage Snapshots

With NAKIVO Backup & Replication, you can enable backup from storage snapshots. This approach can be helpful when you back up large volumes of data. Enabling this option can significantly reduce the load on the production environment. For more information, refer to "Backup from Storage Snapshots" on page 69.

**Notes**

This option is disabled in the following cases:

- No supported storage devices were added to the **Inventory**.
- The selected VMs and its disks don't reside on the supported storage devices.
- The selected VMs and its disks reside on more than 1 datastore/virtual volume.
- The selected VMs reside on multiple ESXi hosts.

Proceed by choosing one of the following options:

- **Enabled (proceed on error)**: With this option selected, NAKIVO Backup & Replication proceeds even if an error is encountered.
- **Enabled (fail on error)**: With this option selected, NAKIVO Backup & Replication automatically fails the job if an error is encountered.
- **Disabled**: Selecting this option disables backup from storage snapshots.

### Bottleneck Detection

When the **Bottleneck detection** option is enabled, additional information is collected and recorded in NAKIVO Backup & Replication logs in the course of data transfer for the purpose of bottleneck detection. Check this option to enable the **Bottleneck detection** capability of the **Transporters** engaged in the job.

## Completing the New Backup Job Wizard for VMware

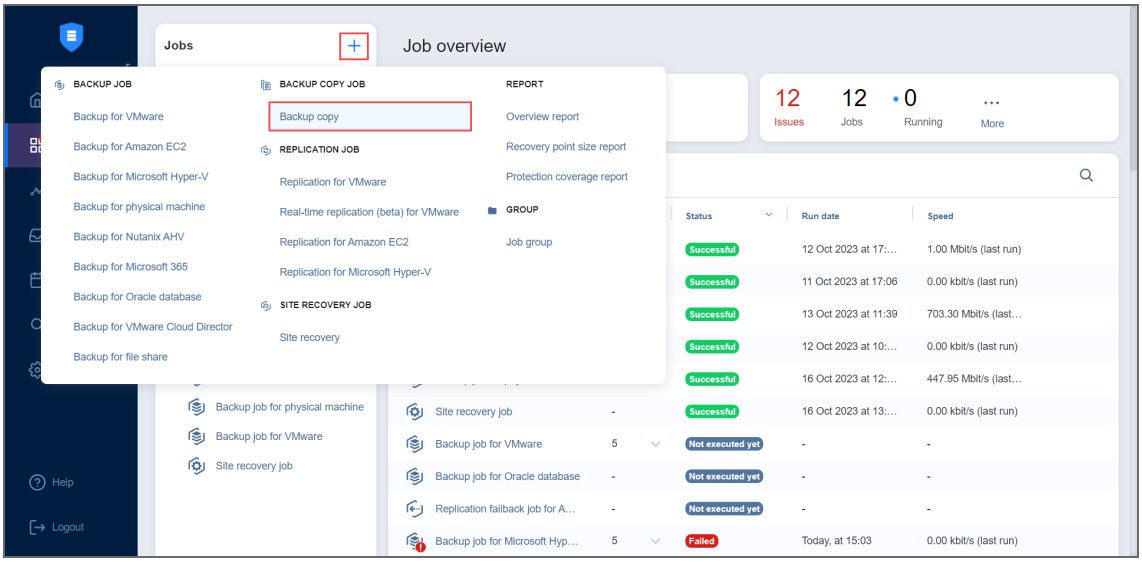
Click **Finish** or **Finish & Run** to complete job creation.

**Note**

If you click **Finish & Run**, you will have to define the scope of your job. Refer to "Running Jobs on Demand" on page 364 for details.

# Creating Backup Copy Jobs

To create a backup copy job, click the plus **Create (+)** button in the **Data Protection** menu, and then click **Backup copy**.




The **New Backup Copy Job Wizard** opens. Complete the wizard as described in the sections below:

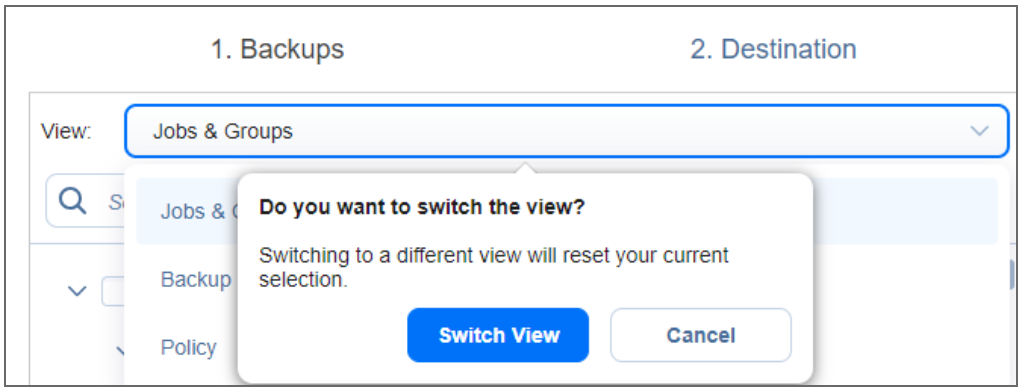
- "Backup Copy Job Wizard: Backups" below
- "Backup Copy Job Wizard: Destination" on page 875
- "Backup Copy Job Wizard: Schedule" on page 878
- "Backup Copy Job Wizard: Retention" on page 889
- "Backup Copy Job Wizard: Options" on page 891

## Backup Copy Job Wizard: Backups

On the **Backups** page of the wizard, you can add items to your backup copy job using one of the inventory views. Proceed as described in the sections below:

- Creating Backup Copies Using Jobs and Groups
- Creating Backup Copies Using Backup Repositories
- Creating Backup Copies Using Policies

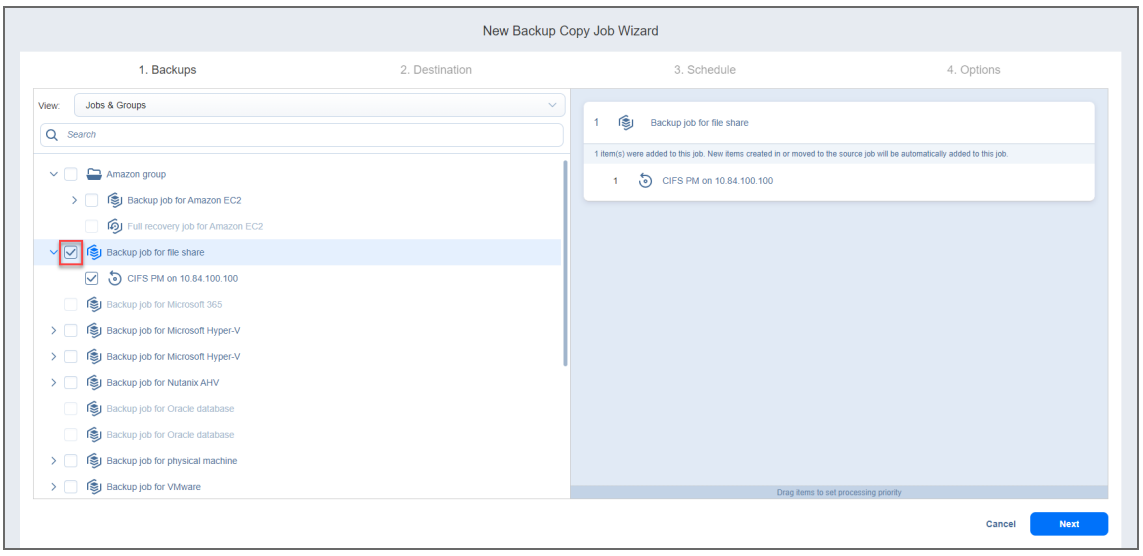
> **Note**
>
> When switching between the **Policy** view and **Jobs & Groups** or **Backup Repositories** views, while some workloads are selected, the following dialog is displayed. Click **Switch View** to switch the current view to the selected one and reset the selection or **Cancel** (or anywhere outside the pop-up notification) to dismiss the pop-up.




## Creating Backup Copies Using Jobs and Groups

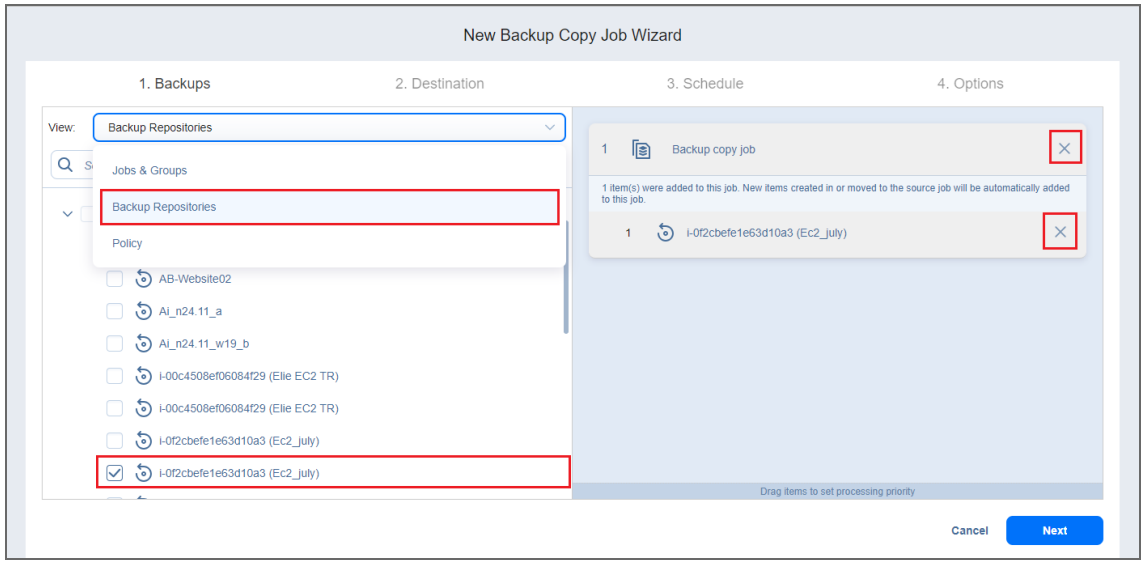
In the left pane of the page, select the **Jobs & Groups** view to use existing backup jobs and groups. The inventory tree opens in the left pane and displays the backup groups along with backups. Proceed as follows:

1. Optionally, you can filter the inventory tree by entering a string into the **Search** box. You can enter a part of the entire name of the item.
2. Select backup items by selecting the checkbox next to the item.
3. The selected items appear in the right pane of the page. If necessary, reorder the selected items by dragging them to a new position. By doing so, you can specify what items you wish to back up first.
4. Review the list of the selected items. If needed, remove a selected backup from the backup copy job in either of the following ways:
    - Cancel the selection of the item(s) in the left pane. This removes the item(s) from the right pane.
    - In the right pane, hover over the item you wish to remove and click the “X” to the right. This cancels the selection of the item(s) in the left pane.



# Creating Backup Copies Using Backup Repositories

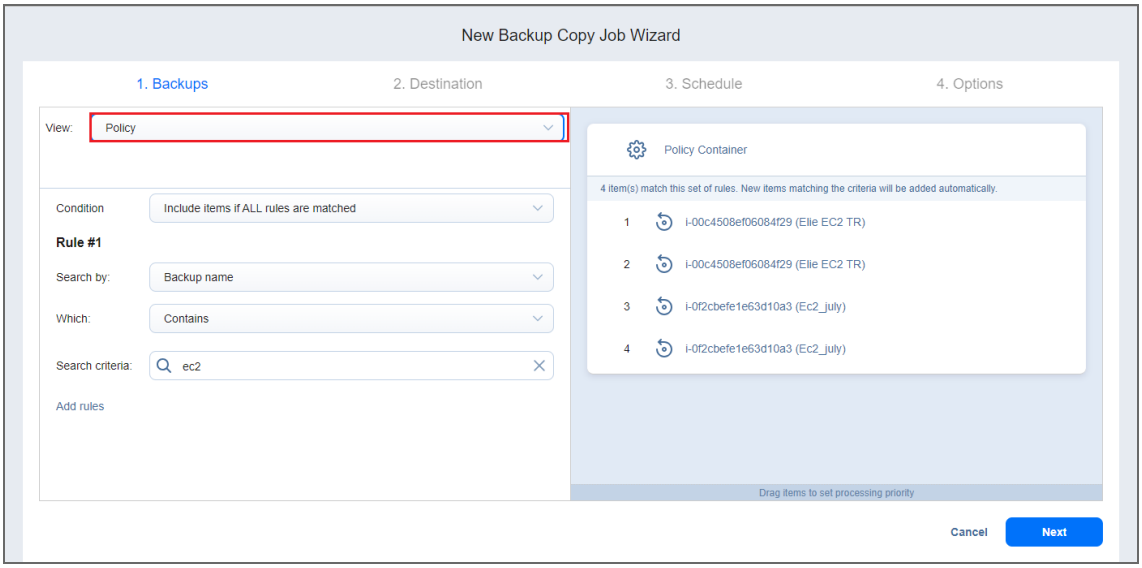
When the **Backup Repositories** view is selected, the inventory tree displays the Backup Repositories along with backups. Proceed as described for the **Jobs & Groups** view above.




# Creating Backup Copies Using Policies

When the **Policy** view is selected, it allows you to use job policies; refer to “Managing Job Policies” on page 382 for details. Follow the steps below:

1. When the items are selected in alternate views, a dialog box opens, warning you that switching to the **Policy** view resets your current selection. Click **Switch View** to confirm that you wish to switch to the **Policy** view.
2. Make sure that at least one item matches the available set of policy rules. Refer to “Managing Policy Rules” on page 385 for details.



Click **Next** to confirm that you wish to add selected items to the backup copy job. The wizard displays the next page.

**Notes**

- When you add a container—a group, job, or Backup Repository—to the backup copy job, the following happens:
  - All backups currently available in the selected container will be backed up.
  - All new backups that will be created in (or moved to) the container in the future will be automatically added to the job and backed up.
- The order in which backups are copied is important if the Transporter running the job cannot process all items simultaneously: either because the Transporter is processing other tasks at the same time or because the number of backups in the job exceeds the Transporter's Maximum Load specified during Transporter creation.
- A backup object containing encrypted recovery points with no corresponding password hash required for its decryption cannot be added to the job. To add such backup to the job provide the password for the corresponding recovery point manually. For more information, refer to Password Management.

# Backup Copy Job Wizard: Destination

On the **Destination** page of the wizard, select a target location for backup copies.

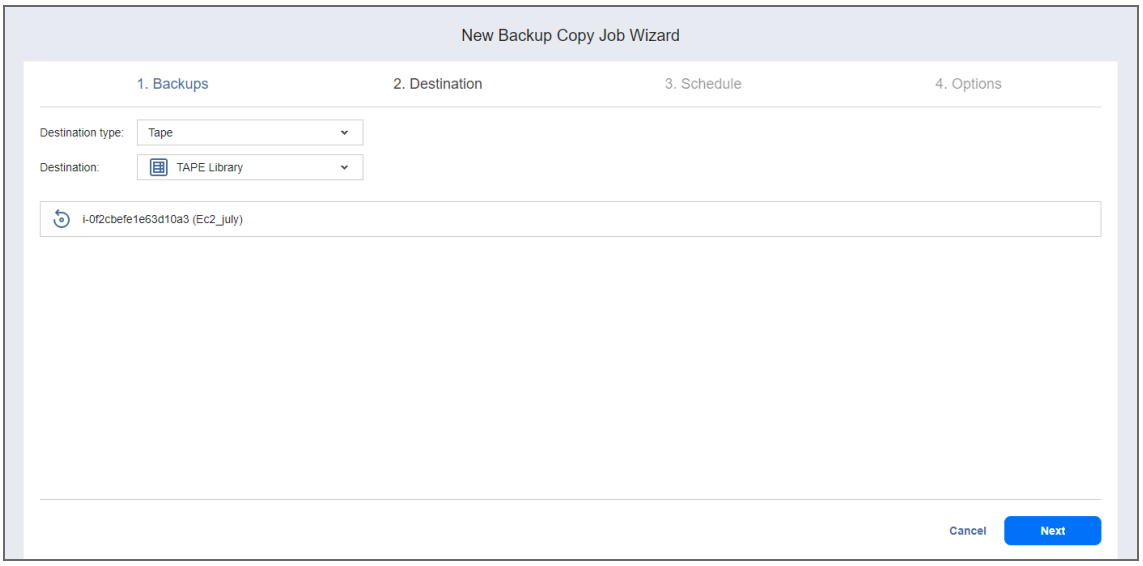
- Selecting a Tape Storage
- Selecting A Target Backup Repository
- Mapping Source Backups to Existing Backups

**Note**

You can create a copy job of a Proxmox VE VM backup to supported backup repositories (except SaaS) and tape media. Supported backup repositories include forever incremental wih full type of repositories with immutability supported as well as forever incremental repositories.

## Selecting a Tape Storage

The **Backup Copy Job Wizard** allows you to copy backups to tape devices or to media pools. To do this, select **Tape** from the **Destination type** drop-down list.




**Notes**

- To enable this option, at least one tape device must be added.
- It is not possible to send backup copies from SaaS repositories to tape.

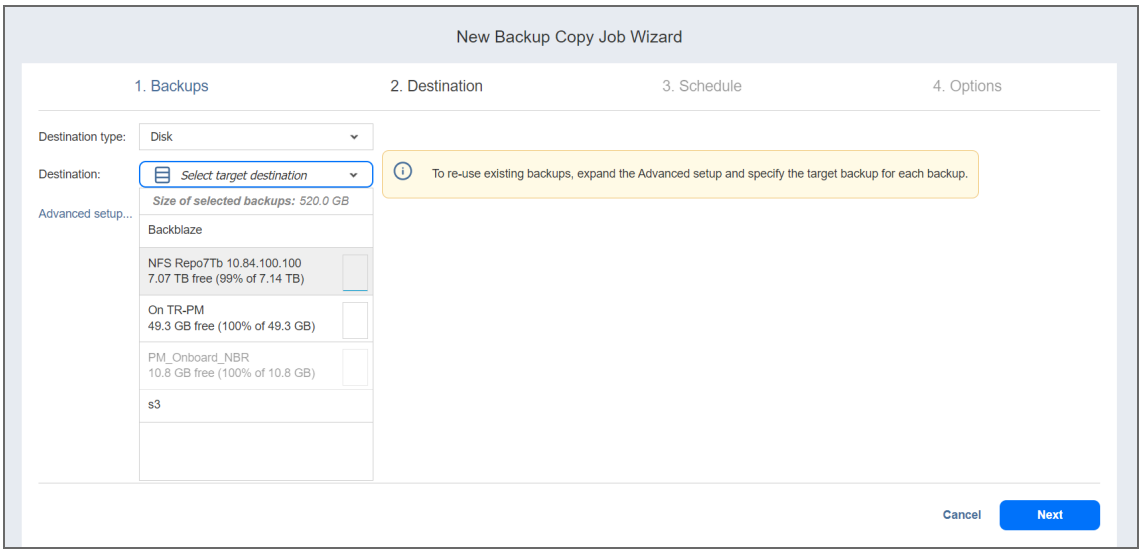
## Selecting a Target Backup Repository

Backup Copy jobs can copy backups from one Backup Repository to another. Select a target Backup Repository as described below:

- To copy all backups you have selected on the Backups page to a single Backup Repository, select **Disk** from the **Destination type** drop-down list and then select a Backup Repository from the **Destination** drop-down list.

**Notes**

- Both federated and standalone (not used as members of federated repositories) can be selected as a destination.
- If a federated repository is selected as the destination, the product automatically selects an available qualified federated repository member during the job run.
- SaaS backup repositories are not available for selection at this step.
- If some but not all backups that are selected in Backups step are Microsoft 365 backups, in **Destination** dropdown menu, unsupported Microsoft 365 locations are disabled.
- If Microsoft 365 backups are automatically added to jobs where the selected destination does not support Microsoft 365 backups, such backups are skipped.

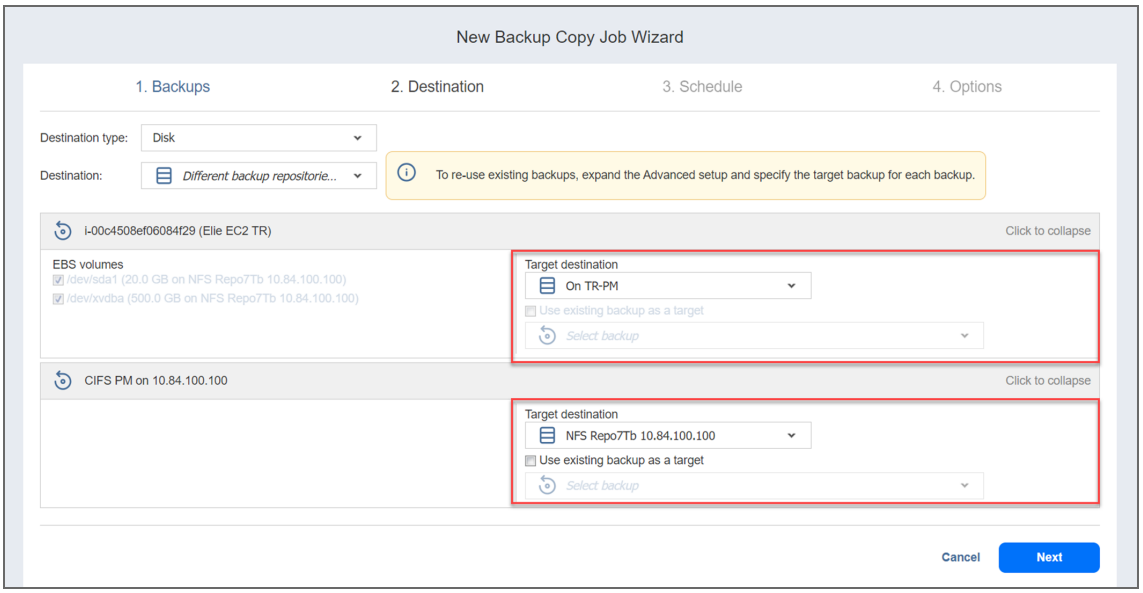



- To copy backups to different Backup Repositories, follow the steps below:

a. Click **Advanced setup...**.

b. For each backup, select a target Backup Repository.

c. Select the **Use existing backup as a target** option and select the existing backup copy from the drop-down list.

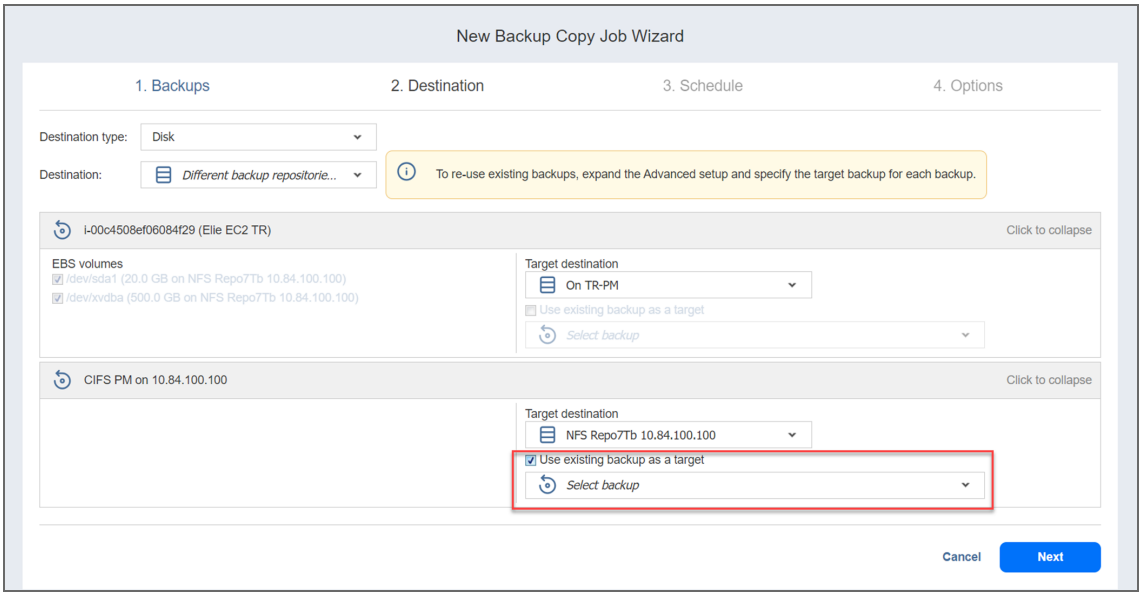



# Mapping Source Backups to Existing Backups

If you lose the Backup Copy job previously created due to accidental job deletion or need to recreate jobs in a new instance of the product, you can map source backups to existing backups in the target Backup Repository to avoid transferring all backup data again.

To map source backups to existing backups in a target Backup Repository, follow the steps below:

1. Click **Advanced setup...**.
2. From the **Backup repository** drop-down list, choose a Backup Repository that contains a copy of the source backup.
3. Select the **Use existing backup as a target** option and select the existing backup copy from the drop-down list.




When running the job, the product analyzes the existing backup copy you have selected, determines how it differs from the source backup, and transfers only the differential data.

> **Note**
> The backup object containing encrypted recovery points with no corresponding password hash required for its decryption cannot be added to the job. To add such backup to the job provide the password for the corresponding recovery point. For more information, refer to Password Management.

# Backup Copy Job Wizard: Schedule

On the **Schedule** page of the wizard, select to run the backup job manually or schedule the job to run on a regular basis.

Proceed as described in the sections below:

- Switching to the improved retention approach
- Creating New Schedules
  - Weekly
  - Monthly
  - Yearly

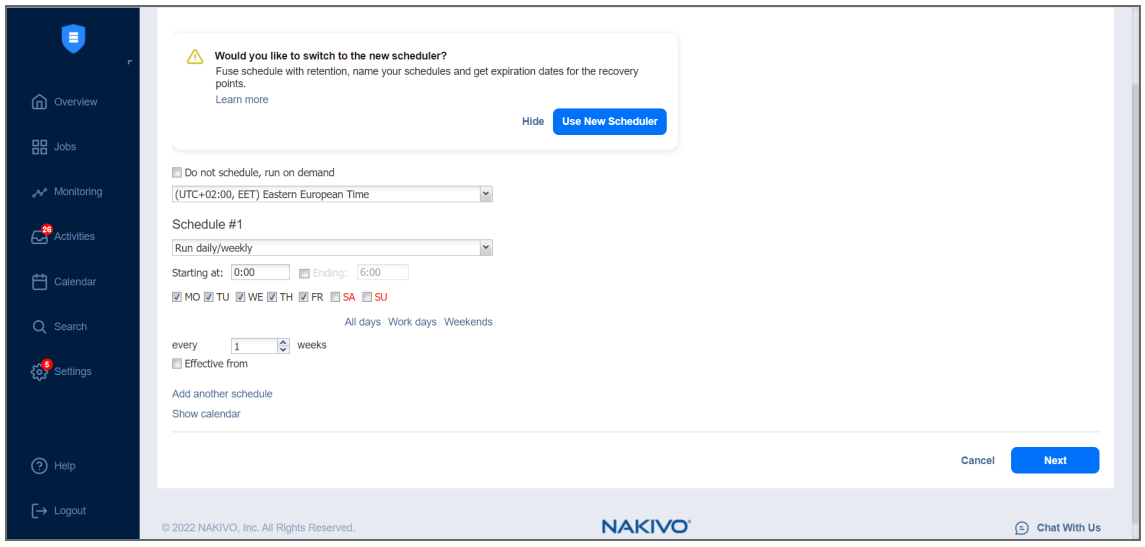
    - Periodical
    - After another job
- Creating Legacy Schedules
    - Daily or Weekly Backup
    - Monthly or Yearly Backup
    - Periodic Backup
    - Chained Job


## Switching to Improved Retention Approach

NAKIVO Backup & Replication offers two approaches to retention and scheduling: the legacy or the improved approach. To learn more about how the legacy and improved approaches work, go here. If you create a new job or edit the existing one that uses the legacy approach, a popup appears offering that you to switch to the improved retention approach in the following cases:

- You have updated your instance of the product to v10.8 or later from an older version.
- You have imported a configuration to an instance of NAKIVO Backup & Replication v10.8 or later from an older version.

> **Note**
> If you install NAKIVO Backup & Replication v10.8 or higher, the new approach is enabled by default.




After the popup appears, do one of the following actions:

- If you do not want to switch to the new scheduler, click **Hide** to close the popup. You can later click **Use New Scheduler** on the **Schedule** page to proceed with the change if you change your mind.
- Alternatively, click **Use New Scheduler** in the popup. Next, choose one of the following options:

- **MIGRATE SETTINGS**: When you select this option, the existing schedules are automatically converted to new schedules and the existing retention settings are mapped to the new schedules.
- **CREATE NEW SCHEDULES**: When you select this option, you can create new schedules using the existing retention settings. Old schedules will be deleted.
- **CONFIGURE SETTINGS ANEW**: Select this option to reset all existing schedules and retention settings and configure them from scratch.

**Notes**

- After switching to the new scheduler, the legacy schedule and retention settings are displayed on the right side of the page.
- After switching to the new scheduler, reverting to the legacy schedule and retention settings is impossible.
- After switching to the new scheduler, the **Schedule Retention Approach** feature becomes available. With this feature, you can select the **Synchronize recovery points and apply custom retention** and **Copy 1 latest recovery point only and apply custom retention** options during the Backup Copy job creation process.
- You can learn how expiration dates are assigned to recovery points after migration here.

## Creating New Schedules

Before creating a new schedule, you can optionally enable the following settings:

- **Do not schedule, run on demand**: Enable this option if you want to start the job manually.
- **Prioritize schedules**: When this option is selected, NAKIVO Backup & Replication starts treating schedules based on their priority. The **Yearly** schedule has higher priority than the **Monthly** schedule, etc. In case 2 or more schedules overlap, the schedules with lower priority will be skipped.
- **Maintain exact copy of the source backup**:This option is displayed only if the backup repository on **Disk** is selected on the **Destination** step of the job creation process. When this option is selected, the backup copy job creates and maintains an exact copy of the source backup and recovery points. Recovery points retention of the source backup is applied to all schedules. Enabling this option disables the **Copy 1 latest recovery point only and apply custom retention** option for all schedules.
- **Keep X last recovery points**: Select this option to keep the specified number of the latest recovery points. Note that selecting this option disables the **Keep backups for** option and enables the **Copy 1 latest recovery point only and applies custom retention** option for all schedules.
- **Keep all recovery points forever**: This option is displayed only if the backup repository on **Tape** is selected on the **Destination** step of the job creation process. With this option selected, all recovery points are kept forever. Enabling this option disables the **Copy 1 latest recovery point only and apply custom retention** option for all schedules.

- **Synchronize recovery points and apply custom retention**: With this option selected, the recovery points kept in the target are synchronized with the recovery points kept in the source. Enabling this option disables the **Copy 1 latest recovery point only and apply custom retention** option for all schedules and enables the **Keep backups for** option.
- **Copy 1 latest recovery point only and apply custom retention**: With this option selected, only 1 latest recovery point is copied from the source when the backup copy job runs on the corresponding schedule.
- Optionally, when creating any type of schedule, click **Show Calendar** to show the calendar or **Hide Calendar** to hide it.

> **Notes**
>
> - With the **Maintain exact copy of the source backup** option selected, all recovery points manually deleted from a source repository are also deleted from a target repository. If deselected, on each job run, the Backup copy job copies and stores recovery points according to its retention policy.
> - The **Copy 1 latest recovery point only and apply custom retention** option is displayed only if the **Keep X last recovery points** option is selected.

When **Amazon EC2**, **Amazon S3**, **Wasabi**, **Azure Blob Storage**, **Backblaze B2 Cloud Storage**, or **Local Folder** is selected as the Backup Repository Type and the only backup destination, you can make recovery points in these repositories immutable during schedule creation. With immutability enabled, the recovery points are immutable and stored using the write-once-read-many (WORM) model. With **Amazon EC2**, **Amazon S3**, **Wasabi**, **Azure Blob Storage**, or **Backblaze B2 Cloud Storage** types of Backup Repository, immutable recovery points cannot be overwritten, deleted, or changed by the root user until the specified period has expired. For Local Folder type of Backup Repository, the root user can still clear immutability.

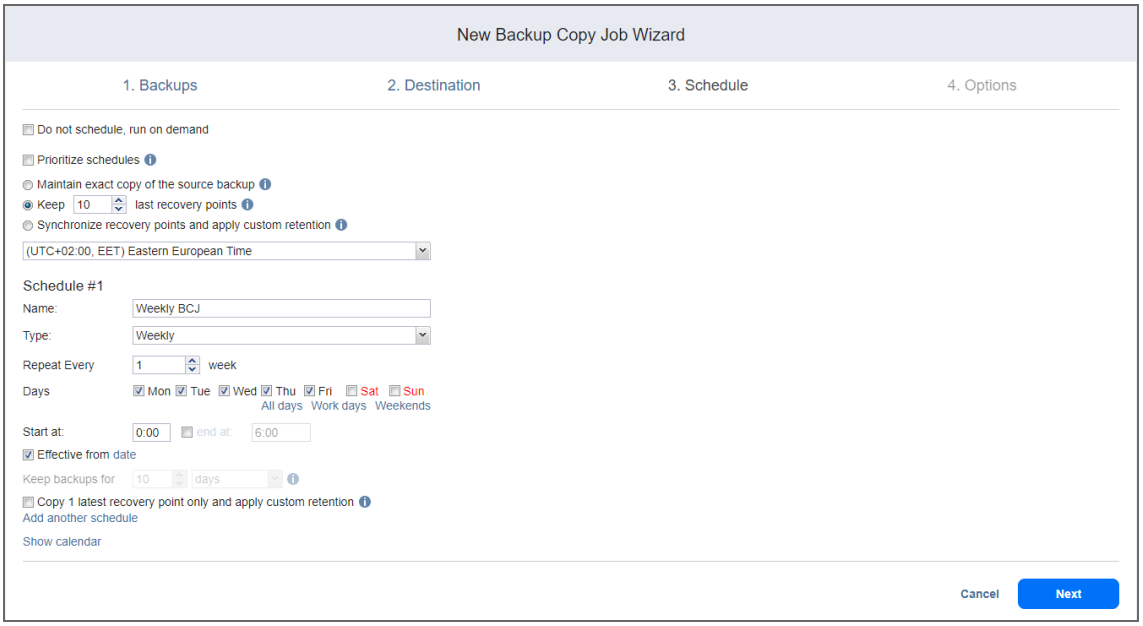
> **Notes**
>
> For the *Immutability* section to be available, the following conditions must be met:
>
> - **Amazon EC2**, **Amazon S3**, **Wasabi**, **Azure Blob Storage**, **Backblaze B2 Cloud Storage**, or **Local Folder** must be selected for Backup Repository Type on the Destination page of the wizard.
> - When **Amazon EC2**, **Amazon S3**, **Wasabi**, **Azure Blob Storage**, or **Backblaze B2 Cloud Storage** is selected as the Backup Repository Type, the bucket or blob container with the repository must have Object Lock or version-level immutability enabled respectively as well as Versioning.
> - For **Local Folder** type of Backup Repository, see feature requirements.

When creating the schedules, you can create schedules of the following types:

## Weekly

You can configure the following options for this schedule type:

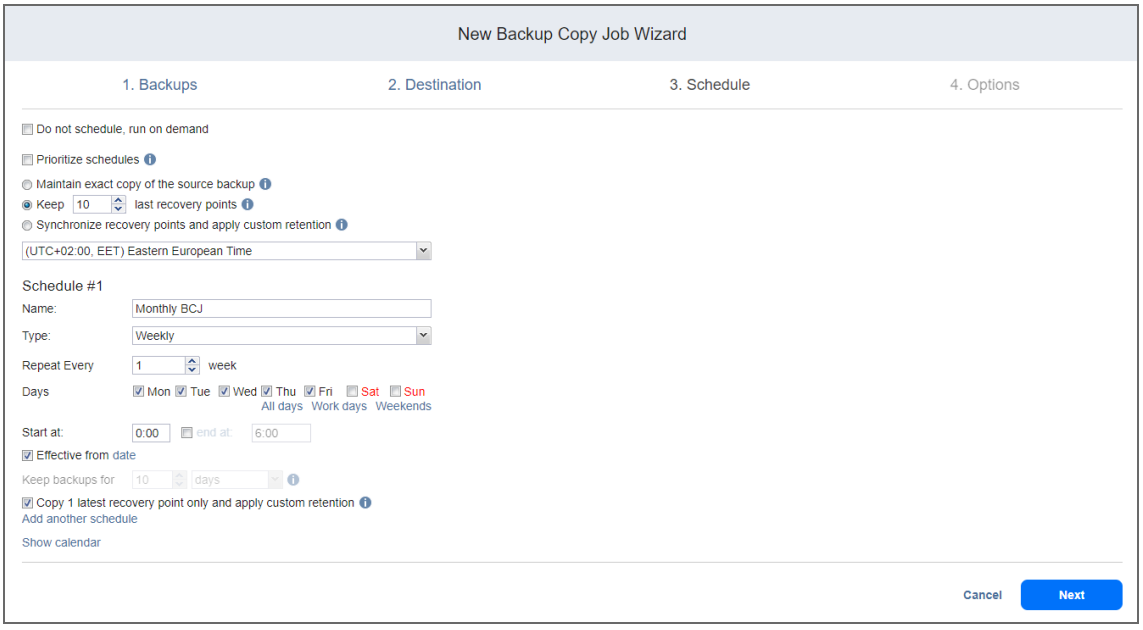
- **Schedule title**: Designates the number of a schedule.
- **Name**: Enter the name of your schedule.
- **Repeat every X weeks**: Indicates how often the schedule is repeated.
- **Days**: Select specific days when the schedule executes the job.
- **Start at**: Specify the time when the job should start.
- **End at**: Specify the time when the job should end.
- Optionally, select the **Effective from** checkbox and choose the date when the schedule should come into effect.
- **Keep backups for**: Specify how many days, months, or years NAKIVO Backup & Replication should retain the backups. This option is enabled if **Synchronize recovery points and apply custom retention** is selected.
- **Copy 1 latest recovery point only and apply custom retention**: With this option selected, only 1 latest recovery point will be copied from the source when the backup copy job runs on the corresponding schedule. This option is enabled if **Keep X last recovery points** is selected.
- **Immutable for X days**: Enabling this option makes the recovery points immutable for the specified number of days. This option is disabled if **Maintain exact copy of the source backup** is selected. This option is not displayed for Oracle database jobs.
- Optionally, click **Add another schedule** if you want to add more than one schedule.




## Monthly

You can configure the following options for this schedule type:

- **Schedule title**: Designates the number of a schedule.
- **Name**: Enter the name of your schedule.
- **Repeat every X months**: Indicates how often the schedule is repeated.
- **Run every**: Select specific days of the month when NAKIVO Backup & Replication executes the job.
- **Start at**: Specify the time when the job should start.
- **End at**: Specify the time when the job should end.
- Optionally, select the **Effective from** checkbox and choose the date when the schedule should come into effect.
- **Keep backups for**: Specify how many days, months, or years NAKIVO Backup & Replication should keep the backups. This option is enabled if **Synchronize recovery points and apply custom retention** is selected.
- **Copy 1 latest recovery point only and apply custom retention**: With this option selected, only 1 latest recovery point will be copied from the source when the backup copy job runs on the corresponding schedule. This option is enabled if **Keep X last recovery points** is selected.
- **Immutable for X days**: Enabling this option makes the recovery points immutable for the specified number of days. This option is disabled if **Maintain exact copy of the source backup** is selected. This option is not displayed for Oracle database jobs.
- Optionally, click **Add another schedule** if you want to add more than one schedule.

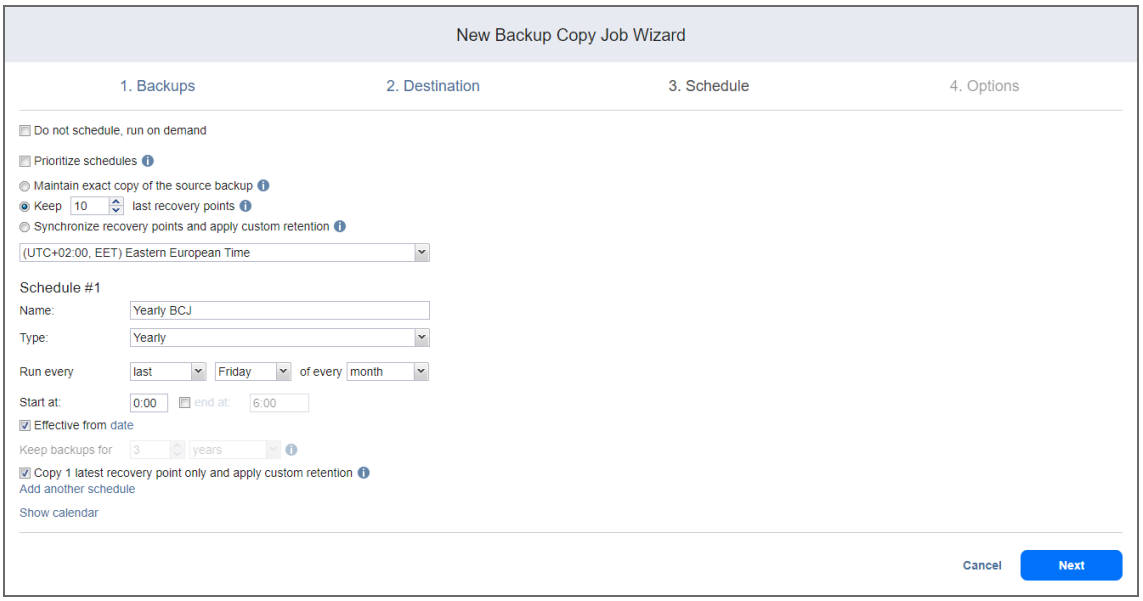



## Yearly

You can configure the following options for this schedule type:

- **Schedule title**: Designates the number of a schedule.
- **Name**: Enter the name of your schedule.

- **Run every**: Select specific days of the specific month when NAKIVO Backup & Replication executes the job.
- Optionally, select the **Effective from** checkbox and choose the date when the schedule should come into effect.
- **Keep backups for**: Specify how many days, months, or years NAKIVO Backup & Replication should keep the backups. This option is enabled if **Synchronize recovery points and apply custom retention** is selected.
- **Copy 1 latest recovery point only and apply custom retention**: With this option selected, only 1 latest recovery point will be copied from the source when the backup copy job runs on the corresponding schedule. This option is enabled if **Keep X last recovery points** is selected.
- **Immutable for X days**: Enabling this option makes the recovery points immutable for the specified number of days. This option is disabled if **Maintain exact copy of the source backup** is selected. This option is not displayed for Oracle database jobs.
- Optionally, click **Add another schedule** if you want to add more than one schedule.

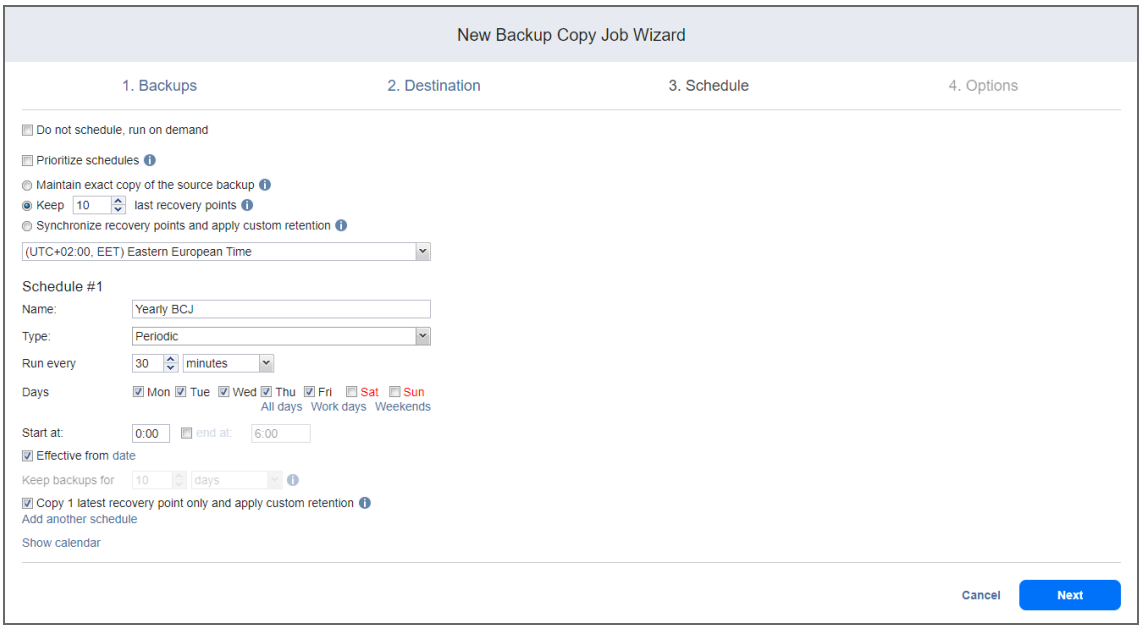



## Periodical

You can configure the following options for this schedule type:

- **Schedule title**: Designates the number of a schedule.
- **Name**: Enter the name of your schedule.
- **Run every**: Select the period measured in minutes, hours, or days when NAKIVO Backup & Replication executes the job.
- **Start at**: Specify the time when the job should start.
- **End at**: Specify the time when the job should end.
- **Days**: Select specific days when the schedule executes the job.

- Optionally, select the **Effective from** checkbox and choose the date when the schedule should come into effect.
- **Keep backups for**: Specify how many days, months, or years NAKIVO Backup & Replication should keep the backups. This option is enabled if **Synchronize recovery points and apply custom retention** is selected.
- **Copy 1 latest recovery point only and apply custom retention**: With this option selected, only 1 latest recovery point will be copied from the source when the backup copy job runs on the corresponding schedule. This option is enabled if **Keep X last recovery points** is selected.
- **Immutable for X days**: Enabling this option makes the recovery points immutable for the specified number of days. This option is disabled if Maintain exact copy of the source backup is selected. This option is not displayed for Oracle database jobs.
- Optionally, click **Add another schedule** if you want to add more than one schedule.
- Optionally, when creating any type of schedule, click **Show Calendar** to show the calendar or **Hide Calendar** to hide it.



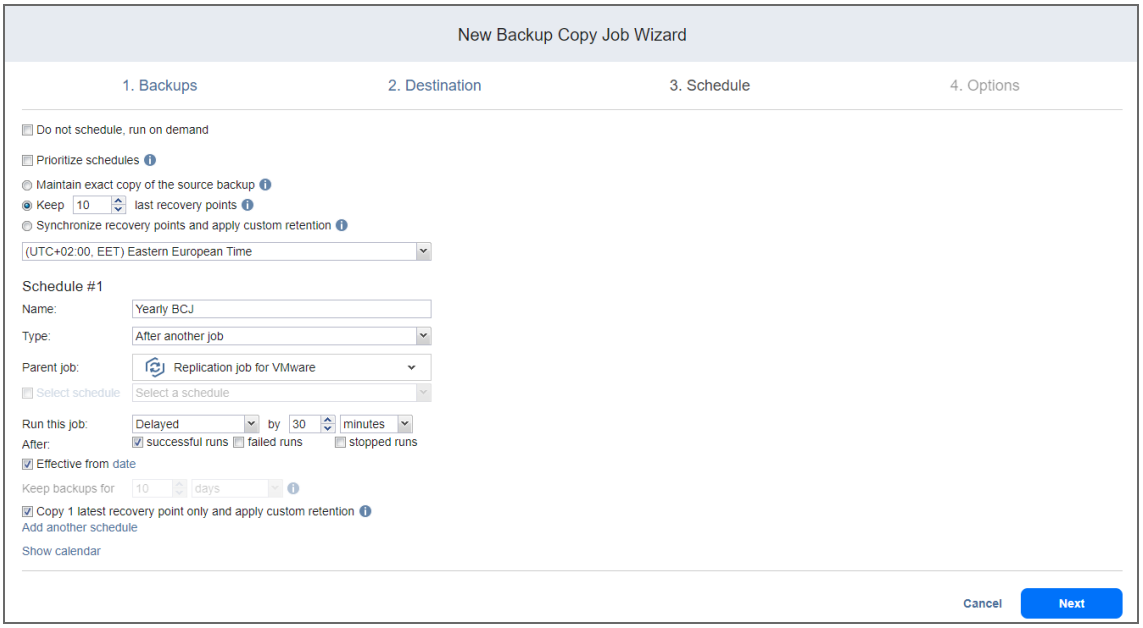

## After Another Job

You can configure the following options for this schedule type:

> **Note**
> This option is disabled if there are no other jobs.

- **Schedule title**: Designates the number of a schedule.
- **Name**: Enter the name of your schedule.
- **Parent job**: Select the job after which this job starts running.

- **Run this job**: Select one of the following options:
  - **Immediately**: The schedule starts right after the parent job is completed.
  - **Delayed**: The schedule starts after the specified number of **minutes** or **hours** following parent job completion.
- Optionally, select the **Effective from** checkbox and choose the date when the schedule should come into effect.
- **Keep backups for**: Specify how many days, months, or years NAKIVO Backup & Replication should keep the backups. This option is enabled if **Synchronize recovery points and apply custom retention** is selected.
- **Copy 1 latest recovery point only and apply custom retention**: With this option selected, only 1 latest recovery point will be copied from the source when the backup copy job runs on the corresponding schedule. This option is enabled if **Keep X last recovery points** is selected.
- **Immutable for X days**: Enabling this option makes the recovery points immutable for the specified number of days. This option is not displayed for Oracle database jobs.
- Optionally, click **Add another schedule** if you want to add more than one schedule.
- Optionally, click **Show Calendar** to show the calendar or **Hide Calendar** to hide it.




# Creating Legacy Schedules

Before creating a new schedule, you can optionally enable the following settings:

- Click **Use New Scheduler** to switch to the Improved retention approach.
- **Do not schedule, run on demand**: Enable this option if you want to start the job manually.
- Optionally, when creating any type of schedule, click **Show Calendar** to show the calendar or **Hide**

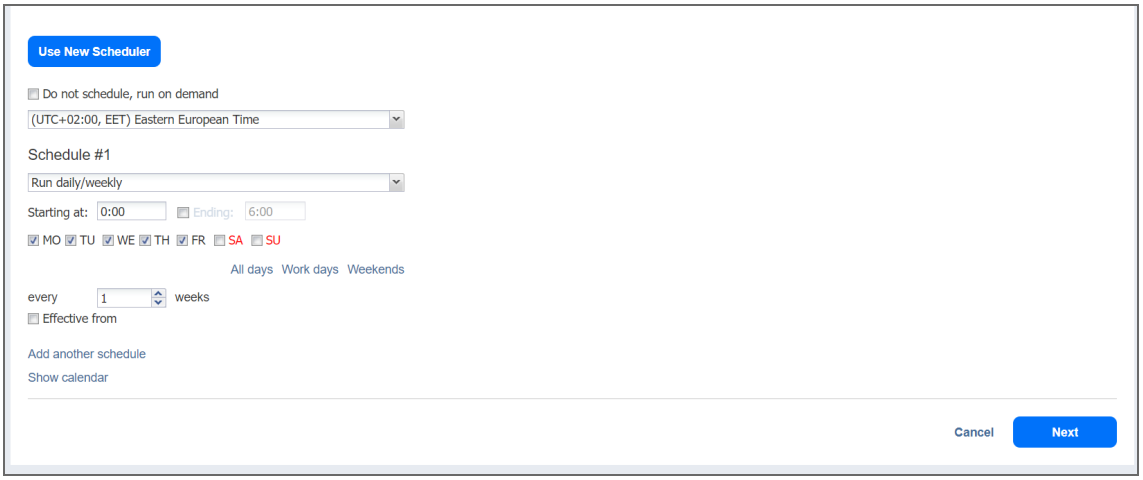
**Calendar** to hide it.

- Optionally, click **Add another schedule** if you want to add more than one schedule.

## Daily or Weekly Backup

To run the job once a day, choose **Run daily/weekly** from the schedule drop-down list and do the following:

- Specify the time when the job should be started in the **Starting** box.
- Specify the end time for the job in the **Ending** box. If the job has not completed by the time specified, the job will be stopped.
- Choose a time zone that should be used for the job start and end times from the time zone drop-down list.
- Select the days of the week during which the job will be started.
- If necessary, select the **Effective from** checkbox and pick the date when the schedule comes into effect.

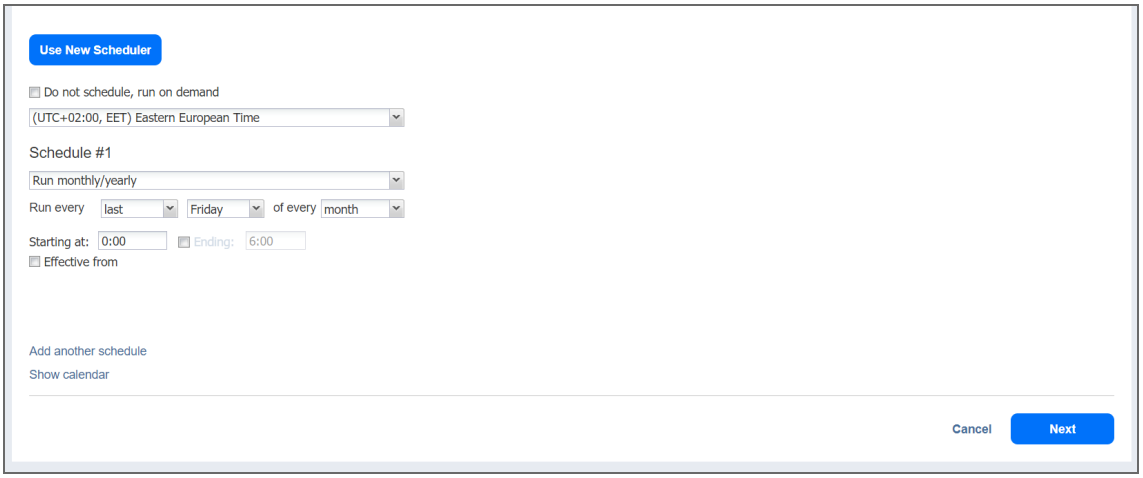



## Monthly or Yearly Backup

To run the job monthly or yearly, choose **Monthly/yearly** from the schedule drop-down list and do the following:

- Specify the job start schedule in the appropriate boxes.
- Specify the day and month when the job should be started in the **Run every** boxes.
- Specify the time when the job should be started in the **Starting** box.
- Specify the end time for the job in the **Ending** box. If the job has not completed by the time specified, the job will be stopped.

- If necessary, select the **Effective from** checkbox and pick the date when the schedule comes into effect.




## Periodic Backup

To run the job multiple times per day, choose **Run periodically** from the schedule drop-down list and then choose a time period from the appropriate boxes:

- Specify the time when the job should be started in the **Starting** box.
- Specify the end time for the job in the **Ending** box. If the job has not completed by the time specified, the job will be stopped.
- Select the days of the week during which the job will be started.
- If necessary, select the **Effective from** checkbox and pick the date when the schedule comes into effect.




## Chained Job

To run the job after a previous one has completed, choose **Run after another job** from the schedule drop-down list and set the options as follows:

- **After the job**: Select a job after which the current job will be started.
- **Run this job**: Choose whether to run the current job immediately after the previous one has completed or within a delay.
- **After successful runs**: If selected, the job will run if the previous one has completed successfully.
- **After failed runs**: If selected, the job will run if the previous one has failed.
- **After stopped runs**: If selected, the job will run if the previous one has been stopped.
- **Effective from**: If selected, the schedule will come into effect on the date picked.




# Backup Copy Job Wizard: Retention

**Important**
This page is not displayed if the new scheduler is enabled.

After each job run, NAKIVO Backup & Replication creates a recovery point in the Backup Repository for each instance. A recovery point represents the backed-up instance as of a particular moment in time and allows you to recover individual files, application objects, or the entire instance from the Backup Repository. You can specify how many recovery points to retain in the Backup Repository. The recovery points are retained based on the grandfather-father-son (GFS) backup rotation scheme.

When **Amazon EC2**, **Amazon S3**, **generic S3-compatible storage**, **Wasabi**, **Azure Blob Storage**, **Backblaze B2 Cloud Storage**, **Dell EMC DataDomain**, **NEC HYDRAstor**, or **Local Folder** is selected as the Backup Repository Type for the only backup destination, you can make recovery points in these repositories immutable. With immutability enabled, the recovery points are immutable and stored using the *write-once-read-many* (WORM) model. In case of **Amazon EC2**, **Amazon S3**, **generic S3-compatible storage**, **Wasabi**, **Azure Blob Storage**, **Dell EMC DataDomain**, **NEC HYDRAstor**, or **Backblaze B2 Cloud Storage** types of Backup Repository, immutable recovery points cannot be overwritten, deleted, or changed by the root user, until the specified period has expired. For **Local Folder** type of Backup Repository, the root user can still clear immutability.

## Retention Settings

Here you can set the retention settings for the backup job. Set the following options:

- **Maintain exact copy of the source backup**: When this option is selected, the backup copy job creates and maintains an exact copy of the source backup and recovery points. To set a different retention policy, deselect this option and choose one of the options below. This option only appears if **Disk** was selected in the **Destination** step.
- **Keep all recovery points forever**: When this option is selected, the backup copy job keeps all available recovery points until they are manually removed. To set a different retention policy, deselect this option and choose one of the options below. This option only appears if **Tape** was selected in the **Destination** step.
- **Synchronize recovery points and apply custom retention**: With this option selected, the recovery points kept in the target are synchronized with the recovery points kept in the source.
- **Copy 1 latest recovery point only and apply custom retention**: With this option selected, only 1 latest recovery point is copied from the source when the backup copy job runs on the corresponding schedule.
- **Keep x last recovery points**: Keeps the specified number of last recovery points for each VM in the job.
- **Keep one recovery point per day for x days**: Retains one last recovery point per day for the specified number of days.
- **Keep one recovery point per week for x weeks**: Retains the last available backup of every week for the specified number of weeks.
- **Keep one recovery point per month for x months**: Retains the last available backup of every month for the specified number of months.
- **Keep one recovery point per year for x years**: Retains the last available backup of every year for the specified number of years.

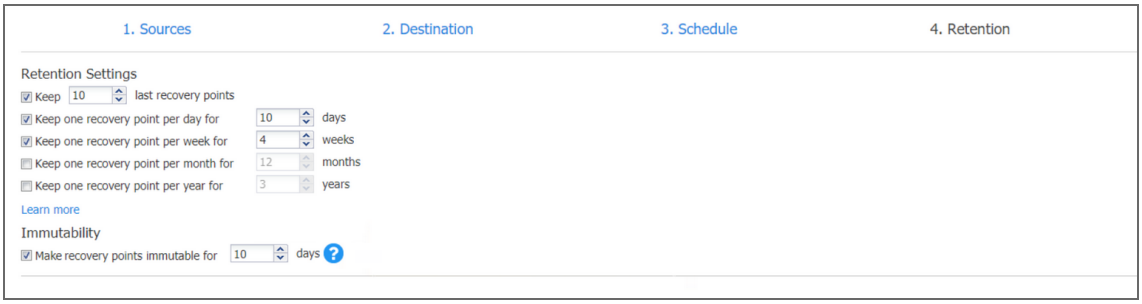
## Immutability

In this section, you can configure the **Make recovery points immutable for x days** option. The recovery points remain immutable for the specified number of days.

**Note**

For the *Immutability* section to be available, the following conditions must be met:

- **Amazon EC2**, **Amazon S3**, **generic S3-compatible storage**, **Wasabi**, **Azure Blob Storage**, **Backblaze B2 Cloud Storage**, **Dell EMC DataDomain**, **NEC HYDRAstor**, or **Local Folder** must be selected for Backup Repository Type on the Destination page of the wizard.
- If **Amazon EC2**, **Amazon S3**, **generic S3-compatible storage**, **Wasabi**, **Azure Blob Storage**, **Dell EMC DataDomain**, **NEC HYDRAstor**, or **Backblaze B2 Cloud Storage** is selected as the Backup Repository type, Object Lock or version-level immutability support and Versioning must be enabled bucket or blob container respectively where your Backup Repository is located.
- For **Local Folder** type of Backup Repository, see feature requirements.




For more details and an example of job retention settings, refer to the Keeping Recovery Points article in the Knowledge Base.

# Backup Copy Job Wizard: Options

On the **Options** page of the wizard, you can set up job options. Proceed as described in these sections:

- Job Options
  - Job Name
  - Job Priority
  - Network Acceleration
  - Backup Encryption
  - Encryption
  - VM Verification
- Full Backup Settings
- Pre and Post Actions
  - Email Notifications
  - Pre Job Script
  - Post Job Script
- Data Transfer
  - Transporter Load
  - Bandwidth Throttling
- Completing the New Backup Copy Job Wizard

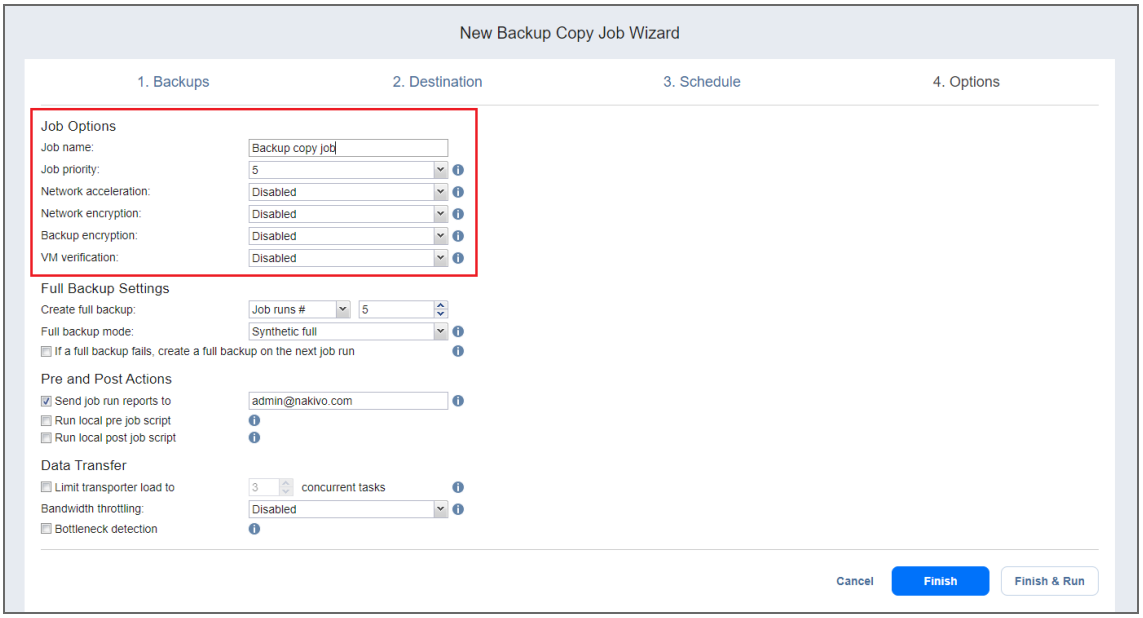
**Note**

If multiple backup types are selected during **Backups** step, including a Microsoft 365 backup, options that are not supported for Microsoft 365 backups become disabled.

If a file-level backup job is selected at the **Backups** step, the **VM Verification** option is unavailable.

# Job Options

In this section, you can give a name to the backup copy job and enable/disable network acceleration, change tracking, set encryption and VM Verification. Proceed as described below.




## Job Name

Specify a name for the backup copy job in the **Job Name** box.

## Job Priority

Select a job priority level between 1 and 5, with 1 being the highest priority. Jobs with higher priority levels are prioritized by Transporters during job processing.

**Note**

This option is only available in the Enterprise, Enterprise Essentials, Enterprise Plus, MSP Enterprise, and MSP Enterprise Plus editions.

## Network Acceleration

If network acceleration is enabled, NAKIVO Backup & Replication uses compression and traffic reduction techniques to speed up data transfer. Select this option if you plan to back up over WAN or slow LAN links.

**Note**

The **Network acceleration** option is not available if the **Backup encryption** option is enabled.

## Network Encryption

If the **Network Encryption** option is selected, backup data will be protected with AES 256 encryption while traveling over the network. Data encryption increases the backup time and CPU load on machines running Transporters. Select this option if you are backing up over WAN without a VPN connection.

**Note**

You need at least one Transporter at the source and target sites to enable network encryption.

## Backup Encryption

When enabled, backup data is protected with AES 256 block cipher encryption with a 256-bit key length. The following three options are available:

- **Disabled (default)**
- **Enabled on source**: If enabled, the single-time-use AES key is passed to the Transporter which reads the source data
- **Enabled on target**: If enabled, the single-time-use AES key is passed to the Transporter which writes the data to target

You can protect the backup file by creating a new password or selecting an existing one. For more information, refer to Enabling Backup Encryption.

**Notes**

- **Backup encryption** is available only if the **Disk** or **Tape** destination type was chosen on the **Destination** page of the wizard.
- The **Backup encryption** dropdown is disabled if the **Network acceleration** option is enabled.
- The **Backup encryption** options cannot be enabled if all source backups are encrypted or multiple targets where supported and unsupported (SaaS repositories and forever incremental) repositories are selected as destinations.
- The **Backup encryption** options are not displayed for a backup job where forever incremental repositories are selected as the only target repositories.
- You cannot copy the encrypted recovery points for which no corresponding password hash is available.
- It's recommended that you enable the **(AWS) Key Management Service**. If **KMS** is enabled, all backup encryption passwords encrypted with the Key Management Service cryptographic key are available for recovery in case of product re-installation. For more information, refer to "Enabling KMS" on page 429.

## VM Verification

VM Verification allows you to check the integrity of the backup by starting it and interacting with it. For more details, refer to "VM Verification" on page 66

You can choose one of the following VM verification options:

- **Disabled:** VM verification is disabled.
- **Screenshot verification**: When enabled, all VM backups created by the job are verified: After a backup of a VM is completed, the VM will be recovered from the backup using Flash boot (and will be disconnected from networks) and a screenshot of the recovered VM will be taken once the VM OS has booted, after which the VM will be discarded. VM screenshots will be included in email notifications (if they're configured) and displayed on the **Dashboard**.
- **Boot verification**: When enabled, all VM backups created by the job are verified as follows. After a VM backup is completed, NAKIVO Backup & Replication recovers the VM using Flash boot, disables networking to prevent network connections, and verifies that system start is successful.

After choosing **Screenshot verification**, provide the following information in the dialog box that opens:

1. Provide a location of the VMs that need to be booted:
    a. **Target Container**: Choose a target container (cluster, host, or resource pool) where VMs will be run using Flash boot.
    b. **Target Datastore**: Choose a datastore that will host changes to the recovered VMs.
    c. **Proxy transporter**: Choose a proxy transporter from the list of available Transporters.

**Note**

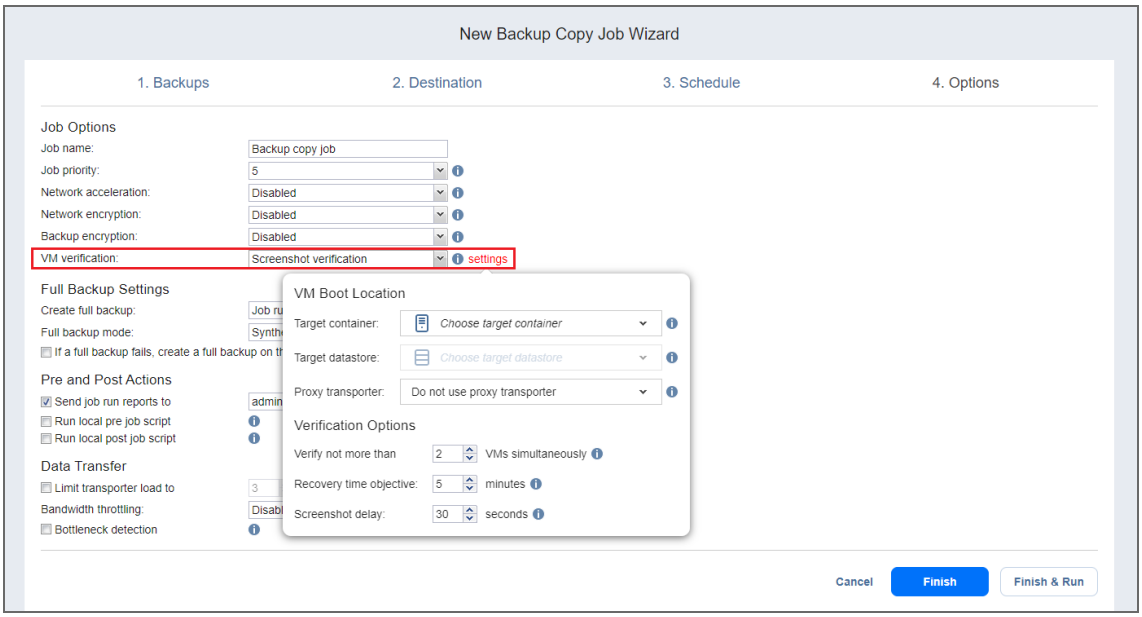
NAKIVO Backup & Replication will use a proxy Transporter in the following cases:

The Transporter assigned to the Backup Repository cannot use iSCSI port 3260 because it is occupied by other services.

iSCSI packages are missing on the Transporter assigned to the Backup Repository.

2. Set verification options:
   - **Verify not more than X VMs simultaneously**: Specify the maximum number of VMs that can be started on the Target Container simultaneously.
   - **Recovery time objective**: Specify the amount of time allocated for the verification of each VM backup. If a VM OS does not start within the specified amount of time, verification will be considered failed.
   - **Screenshot delay**: The amount of time that the product should wait after the guest OS starts before taking a screenshot.

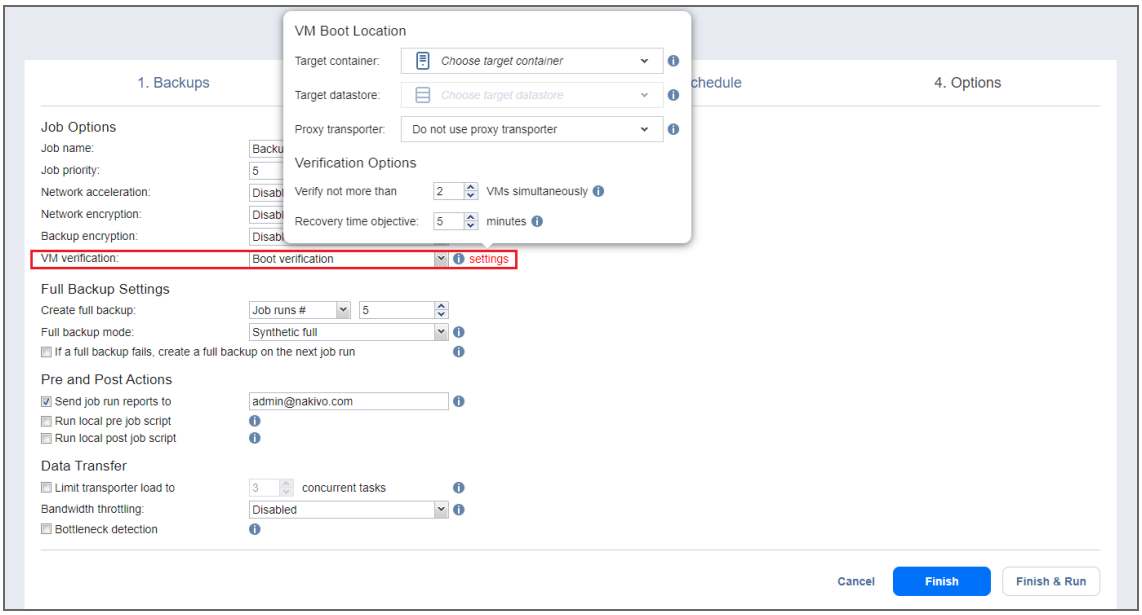
     The specified time must be sufficient to fully start the VM OS. Try increasing this amount if the default amount is not sufficient.




After selecting **Boot verification**, do the following in the dialog box that opens:

1. Provide the location of the VMs to be booted as described for the **Screenshot verification** option.
2. Set verification options:
   - **Verify not more than x VMs simultaneously**: Specify the maximum number of VMs that can be started on the Target Container simultaneously.
   - **Recovery time objective**: Specify the amount of time allocated for verification of each VM backup. If a VM OS does not start within the specified amount of time, verification will be

considered failed.

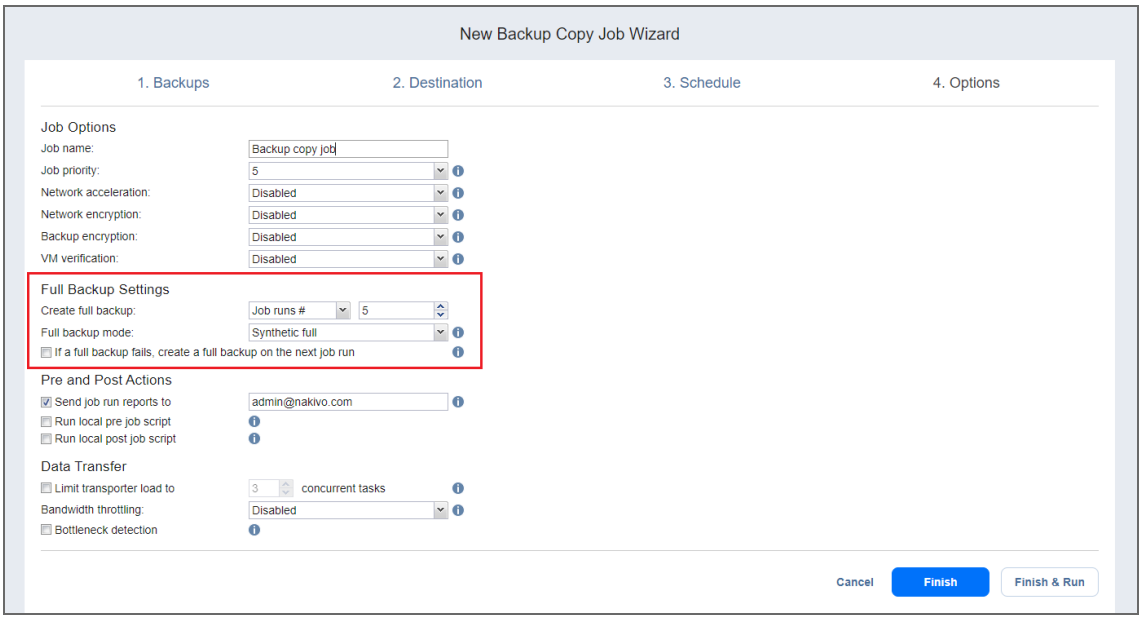



# Full Backup Settings

If the type of the Backup Repository that you selected on the Destination page is set to **Incremental with full backups** (**Store backups in separate files** option is selected), you can specify the following options:

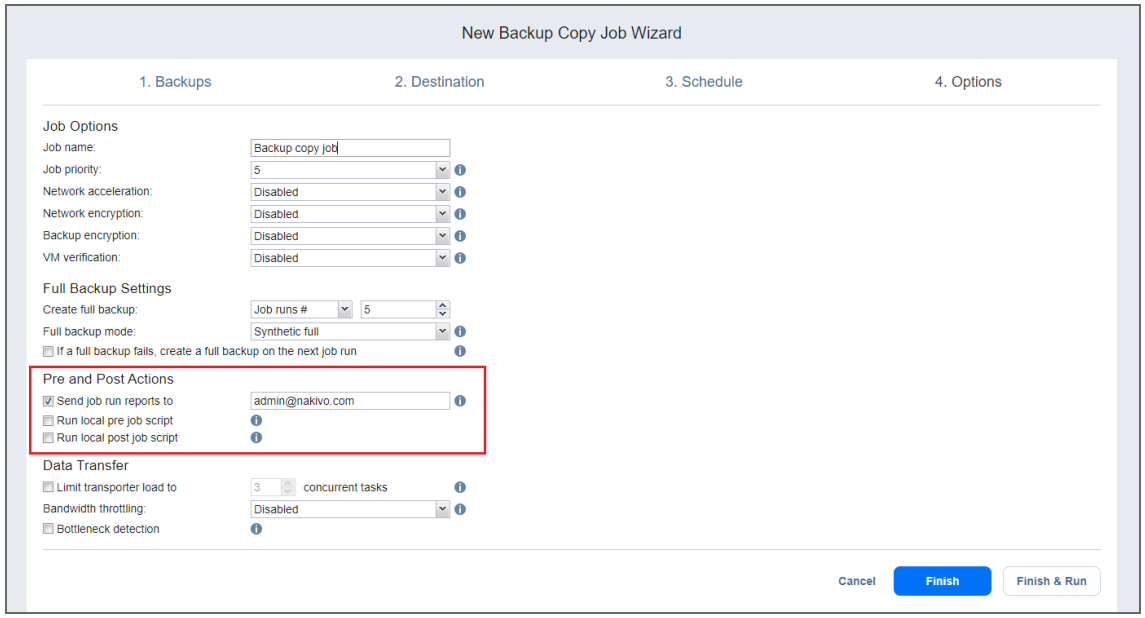
- **Create full backup**: Specify how often full backups should be created.
- **Full backup mode**: Specify how the full backup should be created. You can choose between the following options:
  - **Synthetic full**: If this option is selected, NAKIVO Backup & Replication will first perform an incremental backup (that is, will transfer only the data that changed since the last backup) and will then transform the available data into a full backup file. This approach has the following benefits:
    - The synthetic full backup is usually faster than the active full backup.
    - The load on the network is lower, as less data is transferred.
    - The load on the source datastores running your production VMs is lower.
  - **Active full**: If this option is selected, NAKIVO Backup & Replication will read all data from the source and transfer it to the Backup Repository.
- **If a full backup fails, create a full backup on the next job run**: With this option selected, the next job

run creates a full backup if the current job run fails to do so.




# Pre and Post Actions

NAKIVO Backup & Replication allows you to set up certain actions before a backup copy job begins and after it has completed. You can choose to send job run reports to the email provided and run local pre and post job scripts.




# Email Notifications

NAKIVO Backup & Replication can send email notifications on job completion status to specified recipients. This feature complements global notifications and provides you with the ability to configure notifications on a per-job level.

To enable this option, configure your Email settings.

To send email notifications, do the following:

1. In the *Pre and Post Actions* section, select the **Send job run reports to** option.
2. Specify one or more email addresses in the text box. Separate multiple email addresses with a semicolon.

## Pre Job Script

To run a script before the product begins copying backups, do the following:

1. Place a script file on the machine where the Director is installed.
2. In the *Pre and Post Actions* section, select the **Run local pre job script** option and click the **settings** link. Specify the following parameters in the dialog box that opens:

- **Script path**: Specify a local path to the script on the machine where the Director is installed. Script interpreter should be specified.
  - **Example (Windows)**: *cmd.exe /c D:\script.bat*
  - **Example (Linux)**: *bash /root/script.sh*
- **Job behavior**: Choose either of the following job behaviors in relation to script completion:
  - **Wait for the script to finish**: If this option is selected, the backup copy will not be started until the script is completed.
  - **Do not wait for the script to finish**: If this option is selected, the product will run the script and will start copying backups at the same time.
- **Error handling**: Choose either of the following job behaviors in relation to script failure:
  - **Continue the job on script failure**: If this option is selected, the job will perform backup copy even if the script has failed.
  - **Fail the job on script failure**: If this option is selected and the script fails, the job will be failed and the backup copy will not be performed.

## Post Job Script

To run a script after the product has finished copying all backups, do the following:

1. Place a script file on the machine on which the Director is installed.
2. In the *Pre and Post Actions* section, select the **Run local post job script** option and click the **settings** link. Specify the following parameters in the dialog box that opens:

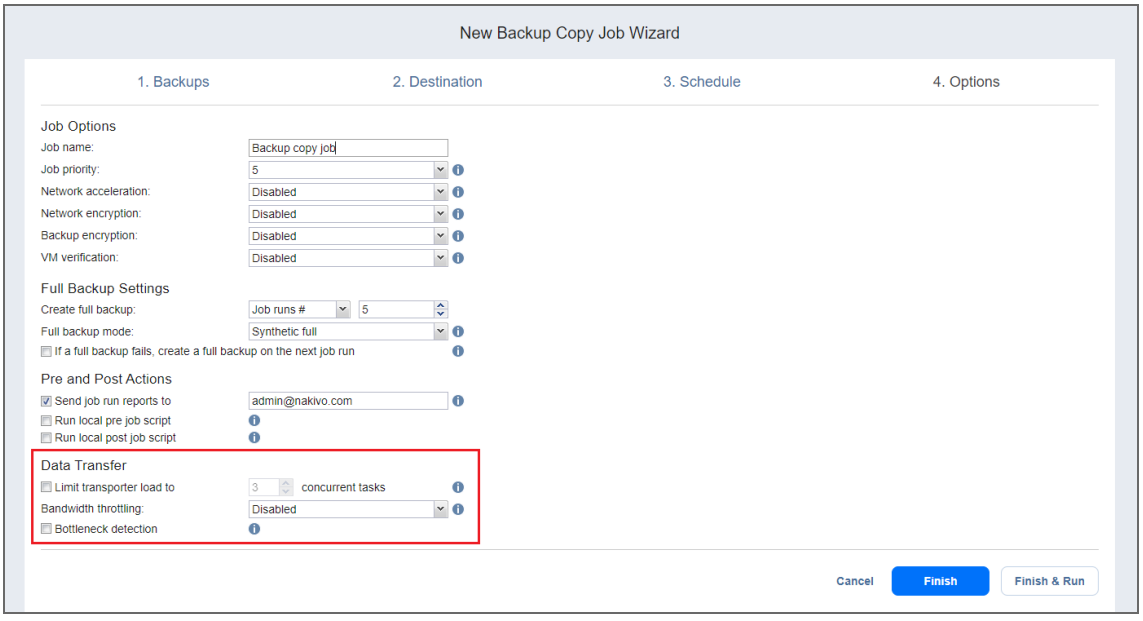
- **Script path**: Specify a local path to the script on the machine on which the Director is installed. Script interpreter should be specified.
  - **Example (Windows)**: *cmd.exe /c D:\script.bat*
  - **Example (Linux)**: *bash /root/script.sh*
- **Job behavior**: Choose either of the following job behaviors in relation to script completion:
  - **Wait for the script to finish**: If this option is selected, the job will be in the "running" state until the script is completed.
  - **Do not wait for the script to finish**: If this option is selected, the job will be completed even if the script execution is still in progress.
- **Error handling**: Choose either of the following job behaviors in relation to script failure.
  - **Continue the job on script failure**: If this option is selected, script failure will not influence the status of the job.
  - **Fail the job on script failure**: If this option is selected and the script has failed, the job status will be set to "failed" even if VM backup has been successful.

**Notes**

- Pre- and post-job scripts can be executed only on the machine on which the Director is installed.
- When Integration Services are used on Hyper-V 2016 and above, custom pre/post scripts are unavailable for Windows VMs.

## Data Transfer

In the *Data Transfer* section of the **Options** page, you can specify a Transporter load and configure bandwidth throttling.



## Transporter Load

You can limit the maximum number of Transporter tasks used by the job. By default, it is set to 3 concurrent tasks.

To change the default number of tasks, do the following:

1. In the *Data Transfer* section, select the **Limit transporter load to checkbox**.
2. Specify the number of concurrent tasks in the corresponding box.

## Bandwidth Throttling

Follow the steps below to regulate the speed of data transfer over the network for your backup copy job:

1. For the **Bandwidth throttling** option, choose **Enabled**.

   > **Note**
   > If bandwidth throttling is disabled for the current job, global bandwidth rules may still apply to your job.

2. Click the **settings** link that becomes available.
3. The **Job Bandwidth Rules** dialog box opens displaying you the list of available rules. You have the following options:
   - Create a new bandwidth rule for your backup copy job:
     a. Click the **Create New Rule** button.
     b. The **New Bandwidth Rule** dialog box opens. Refer to "Bandwidth Throttling" on page 407 for details on creating a bandwidth rule.
     c. Click **Save**.
   - Activate an existing bandwidth rule for your job. Select the checkbox to the left of the necessary bandwidth rule. To deactivate a bandwidth rule for your job, clear the corresponding checkbox.
   - Edit a bandwidth rule. Click the **Edit** link for a bandwidth rule and modify it in the **Edit Bandwidth Rule** dialog box that opens.
   - Disable a bandwidth rule. Click the **Disable** link. The bandwidth rule will be disabled for all jobs.
   - Remove a bandwidth rule. Click the **Remove** link and then click **Delete** to confirm your operation.

## Bottleneck detection

When the **Bottleneck detection** option is enabled, additional information is collected and recorded in NAKIVO Backup & Replication logs in the course of data transfer for the purpose of bottleneck detection. Check this option to enable the **Bottleneck detection** capability of the Transporters engaged in the job.

**Note**

This option is available only if the **Disk** destination type was chosen on the **Destination** page of the wizard.

## Completing the New Backup Copy Job Wizard

Click **Finish** or **Finish & Run** to complete the job creation.

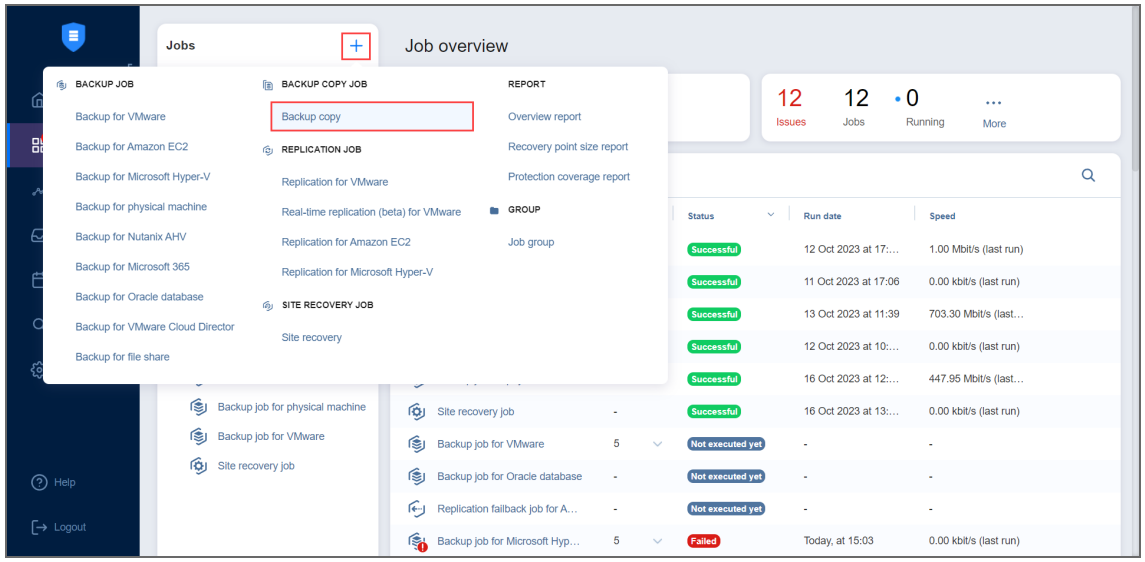
**Note**

If you click **Finish & Run**, you will have to define the scope of your job. Please refer to "Running Jobs on Demand" on page 364 for details.

# Backing Up to Tape

Backing up to tape is, in essence, performing a backup copy job with the destination set to a tape device or media pool. Currently, the direct backing up to tape is not supported, instead, it is done in stages: the backup is first put into a Backup Repository and then moved to tape via a Backup copy job.

To create a backup copy job, click **Create** in the **Data Protection** menu and then click **Backup copy**.

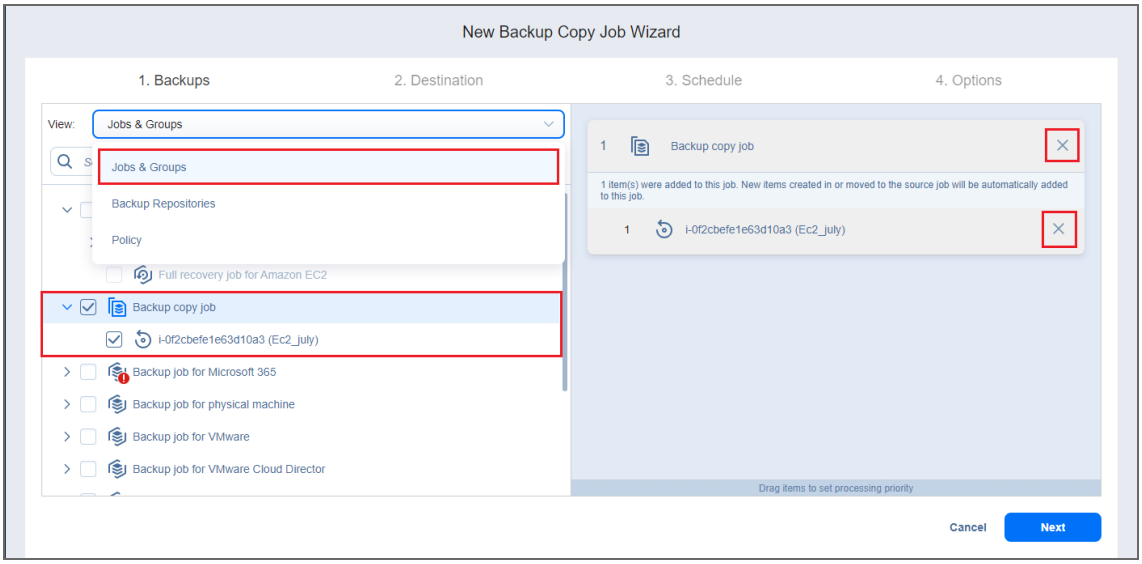



- "Tape Backup Wizard: Backups" below
- "Tape Backup Wizard: Destination" on page 904
- "Tape Backup Wizard: Schedule" on page 905
- "Tape Backup Wizard: Retention" on page 914
- "Tape Backup Wizard: Options" on page 914

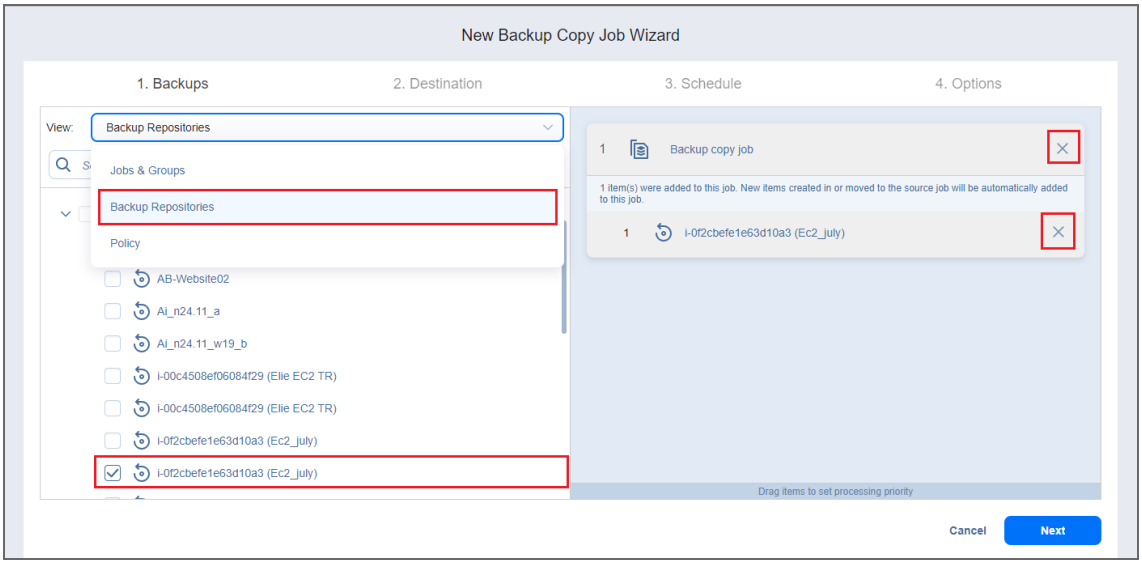
## Tape Backup Wizard: Backups

On the **Backups** page of the wizard, you can add items to your backup copy job. Proceed as follows:

1. In the left pane of the page, choose either of the following inventory views:
   - **Jobs & Groups**: If chosen, the inventory tree opens in the left pane and shows the backup groups along with backups. Proceed as follows:
     a. Optionally, filter the inventory tree by entering a string to the **Search** box. You can enter a part or the entire item name.
     b. Select backup items by selecting the checkbox next to the them.
     c. The selected items appear in the right pane of the page. If necessary, reorder the selected items by dragging them to a new position. By doing so, you can specify the order in which the items should be backed up.
     d. Review the list of the selected items. If needed, remove a selected backup in the backup copy job in either of the following ways:
        - Cancel the selection of the item(s) in the left pane. This removes the item(s) from the right pane.
        - In the right pane, hover over the item you want to remove and click the red "X" to the right. This will cancel the selection of the item(s) in the left pane.

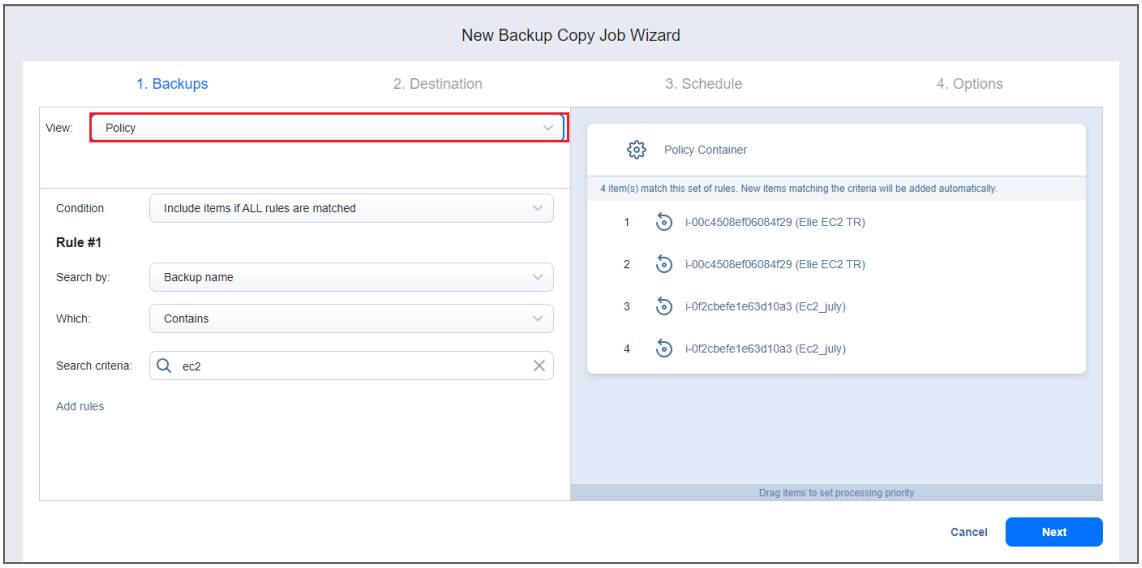



   - **Backup Repositories**: If chosen, the inventory tree shows available Backup Repositories along with the backups in them. Proceed as it is described for the **Jobs & Groups** view above.




Both federated and standalone (not used as members of federated repositories) backup repositories can be selected.

- **Policy**: If selected, this allows you to use job policies; refer to "Managing Job Policies" on page 382 for details. Please follow the steps below:
  a. If items were selected in alternate views, a dialog box opens warning you that switching to the **Policy** view will reset your current selection. Click **Switch View** to confirm switching to the **Policy** view.
  b. Make sure that at least one item matches the available set of policy rules. Refer to "Managing Policy Rules" on page 385 for details.




2. Click **Next** to confirm adding selected items to the backup copy job.

The wizard will display the next page.

**Notes**

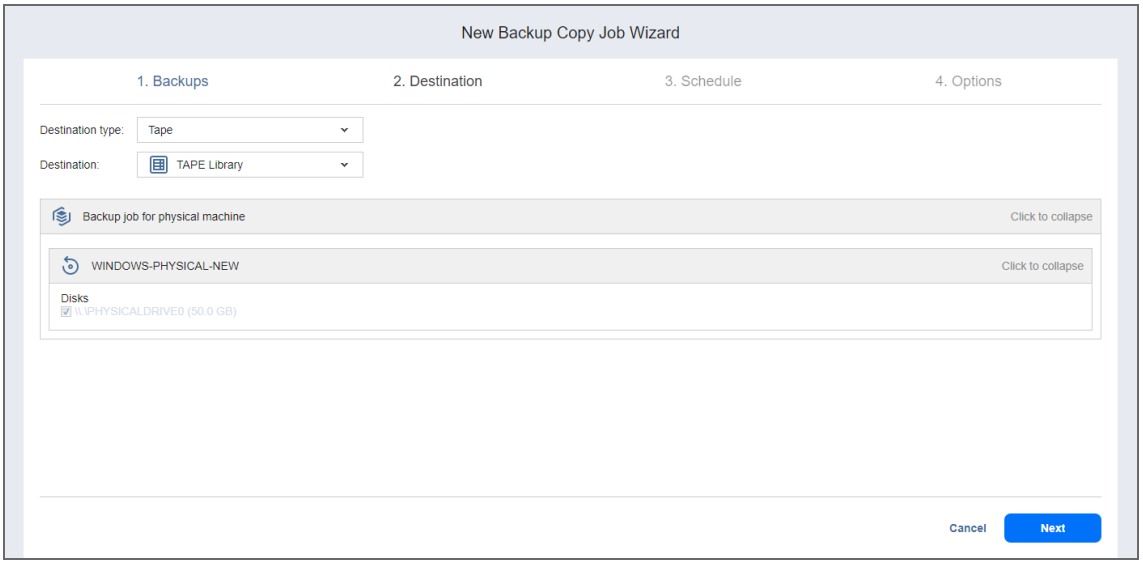
- If you add a container—a group, job, or Backup Repository—to the backup copy job, the following actions will occur:
  - All backups currently available in the selected container will be backed up.
  - All new backups that are created in (or moved to) the container in the future will be automatically added to the job and backed up.
- The order in which backups are copied is important if the Transporter that is running the job cannot process all items simultaneously: either because the Transporter is processing other tasks at the same time or because the number of backups in the job exceeds the Transporter's Maximum Load specified during the Transporter creation.

# Tape Backup Wizard: Destination

On the **Destination** page, you can specify where the backup will be stored. You can select a device or a media pool. The job allows for copying a backup from a Backup Repository to tape cartridges or a virtual tape library.

To specify a destination for the selected backups:

1. From the **Destination type** drop-down list, select **Tape**.
2. From the **Destination** drop-down list, select one of the configured devices or media pools.
3. To see the space and disks the individual backups take, click the name of the job and expand the backups in the list.




4. Click **Next** to proceed to the next page.

# Tape Backup Wizard: Schedule

On the **Schedule** page of the wizard, select to run the backup job manually or schedule the job to run on a regular basis.

Proceed as described in the sections below:


- Switching to the improved retention approach
- Creating New Schedules
  - Weekly
  - Monthly
  - Yearly
  - Periodical
  - After another job
- Creating Legacy Schedules
  - Daily or Weekly Backup
  - Monthly or Yearly Backup
  - Periodic Backup
  - Chained Job

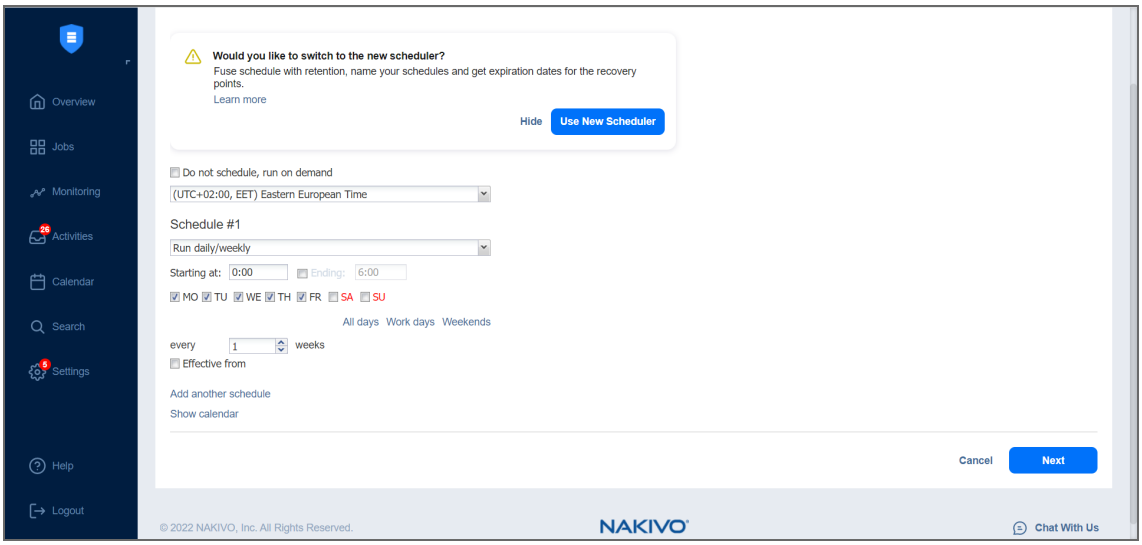

## Switching to Improved Retention Approach

NAKIVO Backup & Replication offers two approaches to retention and scheduling: the legacy or the improved approach. To learn more about how the legacy and improved approaches work, go here. If you create a new job or edit the existing one that uses the legacy approach, a popup appears offering that you to switch to the improved retention approach in the following cases:

- You have updated your instance of the product to v10.8 or later from an older version.
- You have imported a configuration to an instance of NAKIVO Backup & Replication v10.8 or later from an older version.

**Note**

If you install NAKIVO Backup & Replication v10.8 or higher, the new approach is enabled by default.




After the popup appears, do one of the following actions:

- If you do not want to switch to the new scheduler, click **Hide** to close the popup. You can later click **Use New Scheduler** on the **Schedule** page to proceed with the change if you change your mind.
- Alternatively, click **Use New Scheduler** in the popup. Next, choose one of the following options:
  - **MIGRATE SETTINGS**: When you select this option, the existing schedules are automatically converted to new schedules and the existing retention settings are mapped to the new schedules.
  - **CREATE NEW SCHEDULES**: When you select this option, you can create new schedules using the existing retention settings. Old schedules will be deleted.
  - **CONFIGURE SETTINGS ANEW**: Select this option to reset all existing schedules and retention settings and configure them from scratch.

**Notes**

- After switching to the new scheduler, the legacy schedule and retention settings are displayed on the right side of the page.
- After switching to the new scheduler, reverting to the legacy schedule and retention settings is impossible.
- You can learn how expiration dates are assigned to recovery points after migration here.

# Creating New Schedules

Before creating a new schedule, you can optionally enable the following settings:

- **Do not schedule, run on demand**: Enable this option if you want to start the job manually.
- **Keep all recovery points forever**: When this option is selected, the job keeps all available recovery points until they are manually removed.

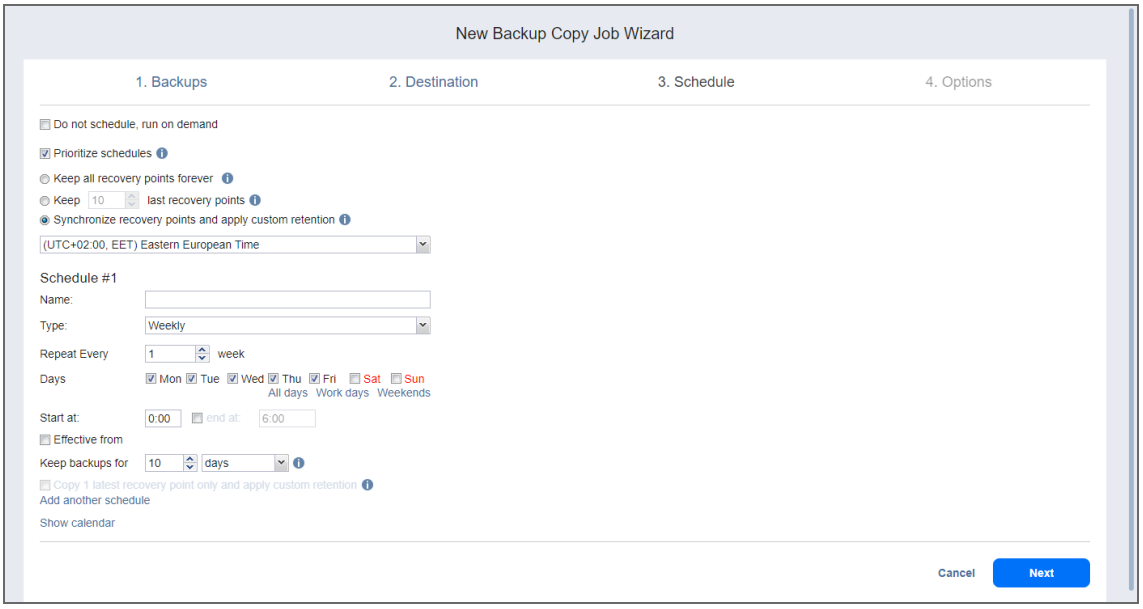
- **Prioritize schedules**: When this option is selected, NAKIVO Backup & Replication starts treating schedules based on their priority. The **Yearly** schedule will have higher priority than the **Monthly** schedule, etc. In case 2 or more schedules overlap, the schedules with lower priority will be skipped.
- Optionally, when creating any type of schedule, click **Show Calendar** to show the calendar or **Hide Calendar** to hide it.

When creating the schedules, you can create schedules of the following types:

## Weekly

You can configure the following options for this schedule type:

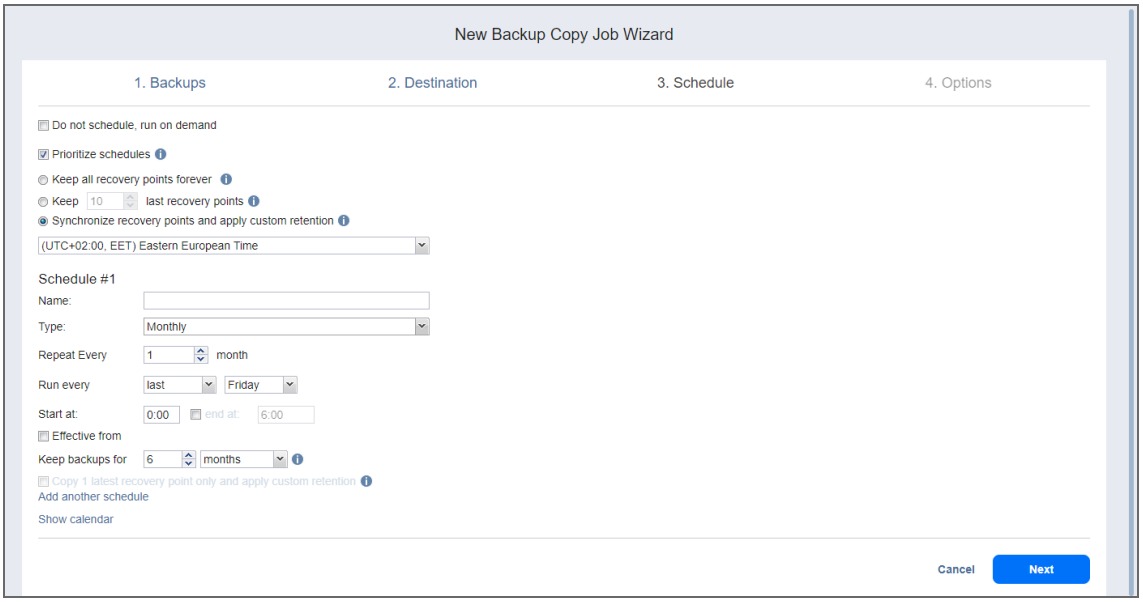
- **Schedule title**: Designates the number of a schedule.
- **Name**: Enter the name of your schedule.
- **Repeat every X weeks**: Indicates how often the schedule is repeated.
- **Days**: Select specific days when the schedule executes the job.
- **Start at**: Specify the time when the job should start.
- **End at**: Specify the time when the job should end.
- Optionally, select the **Effective from** checkbox and choose the date when the schedule should come into effect.
- **Keep backups for**: Specify how many days, months, or years NAKIVO Backup & Replication should retain the backups.
- Optionally, click **Add another schedule** if you want to add more than one schedule.




## Monthly

You can configure the following options for this schedule type:

- **Schedule title**: Designates the number of a schedule.
- **Name**: Enter the name of your schedule.
- **Repeat every X months**: Indicates how often the schedule is repeated.
- **Run every**: Select specific days of the month when NAKIVO Backup & Replication executes the job.
- **Start at**: Specify the time when the job should start.
- **End at**: Specify the time when the job should end.
- Optionally, select the **Effective from** checkbox and choose the date when the schedule should come into effect.
- **Keep backups for**: Specify how many days, months, or years NAKIVO Backup & Replication should keep the backups.
- Optionally, click **Add another schedule** if you want to add more than one schedule.

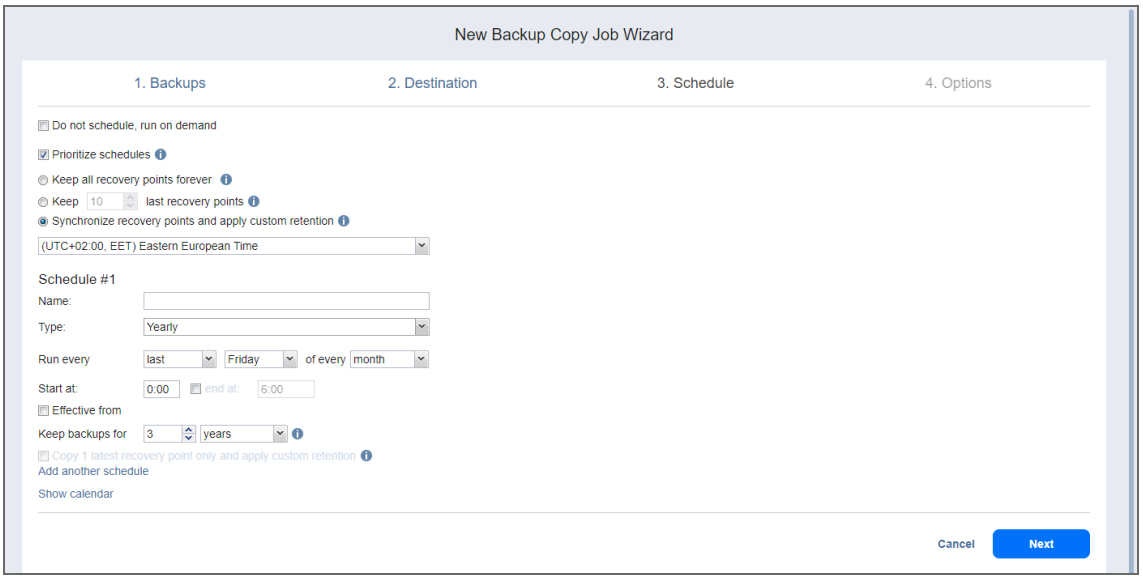



## Yearly

You can configure the following options for this schedule type:

- **Schedule title**: Designates the number of a schedule.
- **Name**: Enter the name of your schedule.
- **Run every**: Select specific days of the specific month when NAKIVO Backup & Replication executes the job.
- Optionally, select the **Effective from** checkbox and choose the date when the schedule should come into effect.
- **Keep backups for**: Specify how many days, months, or years NAKIVO Backup & Replication should keep the backups.

- Optionally, click **Add another schedule** if you want to add more than one schedule.

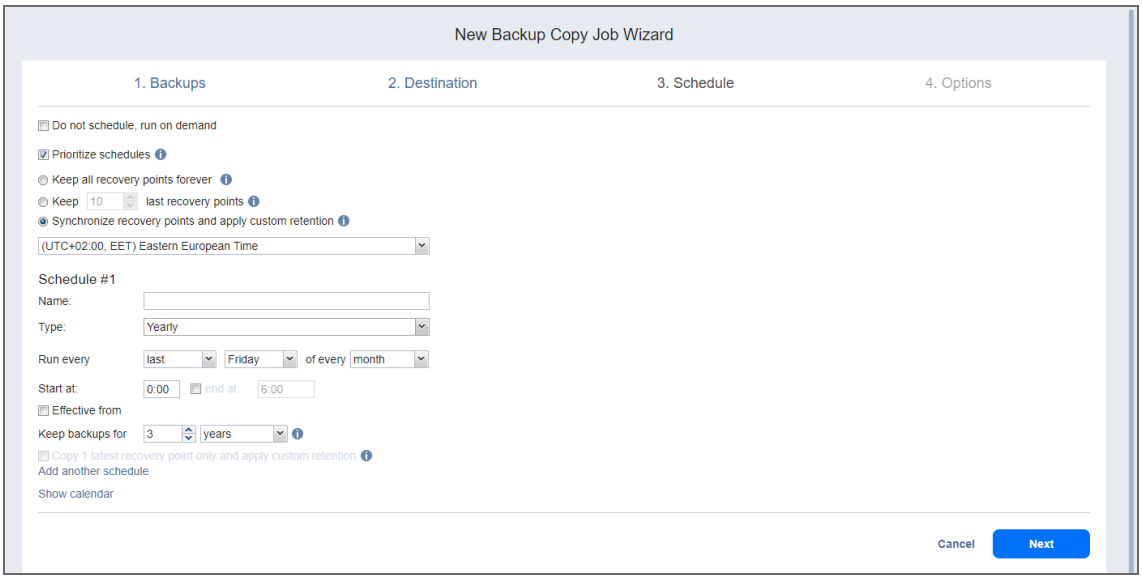



## Periodical

You can configure the following options for this schedule type:

- **Schedule title**: Designates the number of a schedule.
- **Name**: Enter the name of your schedule.
- **Run every**: Select the period measured in minutes, hours, or days when NAKIVO Backup & Replication executes the job.
- **Start at**: Specify the time when the job should start.
- **End at**: Specify the time when the job should end.
- **Days**: Select specific days when the schedule executes the job.
- Optionally, select the **Effective from** checkbox and choose the date when the schedule should come into effect.
- **Keep backups for**: Specify how many days, months, or years NAKIVO Backup & Replication should keep the backups.
- Optionally, click **Add another schedule** if you want to add more than one schedule.
- Optionally, when creating any type of schedule, click **Show Calendar** to show the calendar or **Hide**

**Calendar** to hide it.



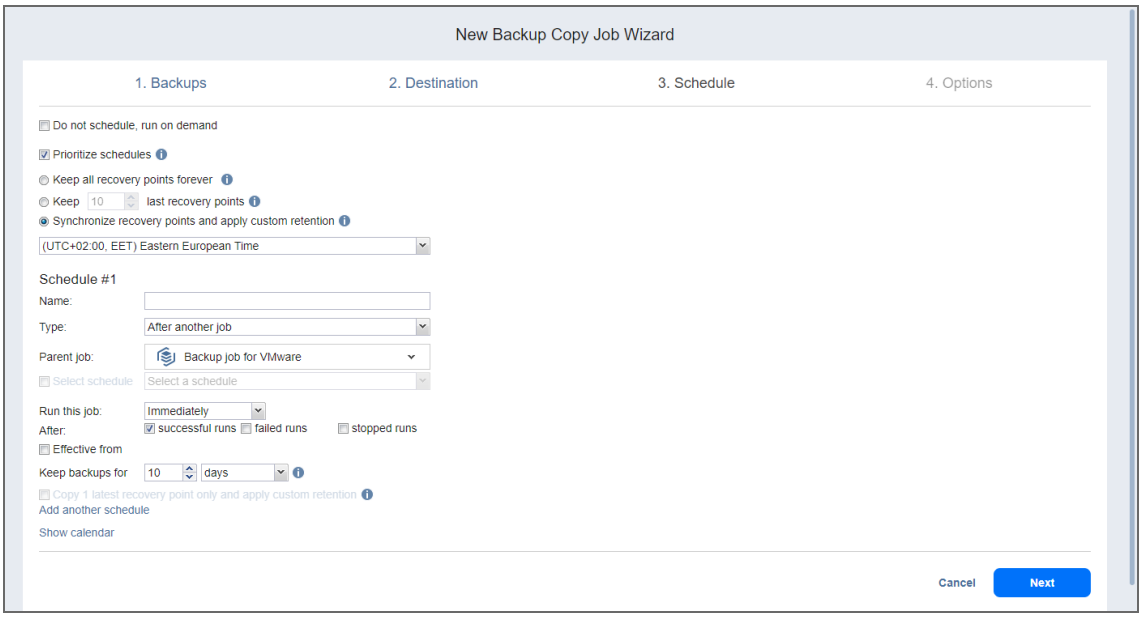

## After Another Job

You can configure the following options for this schedule type:

> **Note**
> This option is disabled if there are no other jobs.

- **Schedule title**: Designates the number of a schedule.
- **Name**: Enter the name of your schedule.
- **Parent job**: Select the job after which this job starts running.
- **Run this job**: Select one of the following options:
  - **Immediately**: The schedule starts right after the parent job is completed.
  - **Delayed**: The schedule starts after the specified number of **minutes** or **hours** following parent job completion.
- Optionally, select the **Effective from** checkbox and choose the date when the schedule should come into effect.
- **Keep backups for**: Specify how many days, months, or years NAKIVO Backup & Replication should keep the backups.
- Optionally, click **Add another schedule** if you want to add more than one schedule.

- Optionally, click **Show Calendar** to show the calendar or **Hide Calendar** to hide it.




# Creating Legacy Schedules

Before creating a new schedule, you can optionally enable the following settings:

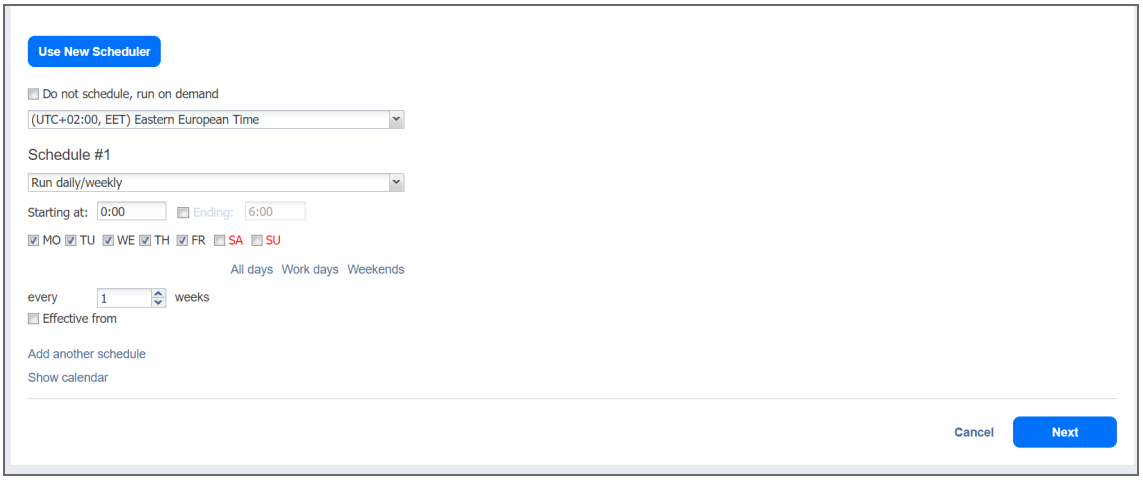
- Click **Use New Scheduler** to switch to the Improved retention approach.
- **Do not schedule, run on demand**: Enable this option if you want to start the job manually.
- Optionally, when creating any type of schedule, click **Show Calendar** to show the calendar or **Hide Calendar** to hide it.
- Optionally, click **Add another schedule** if you want to add more than one schedule.

## Daily or Weekly Backup

To run the job once a day, choose **Run daily/weekly** from the schedule drop-down list and do the following:

- Specify the time when the job should be started in the **Starting** box.
- Specify the end time for the job in the **Ending** box. If the job has not completed by the time specified, the job will be stopped.
- Choose a time zone that should be used for the job start and end times from the time zone drop-down list.
- Select the days of the week during which the job will be started.

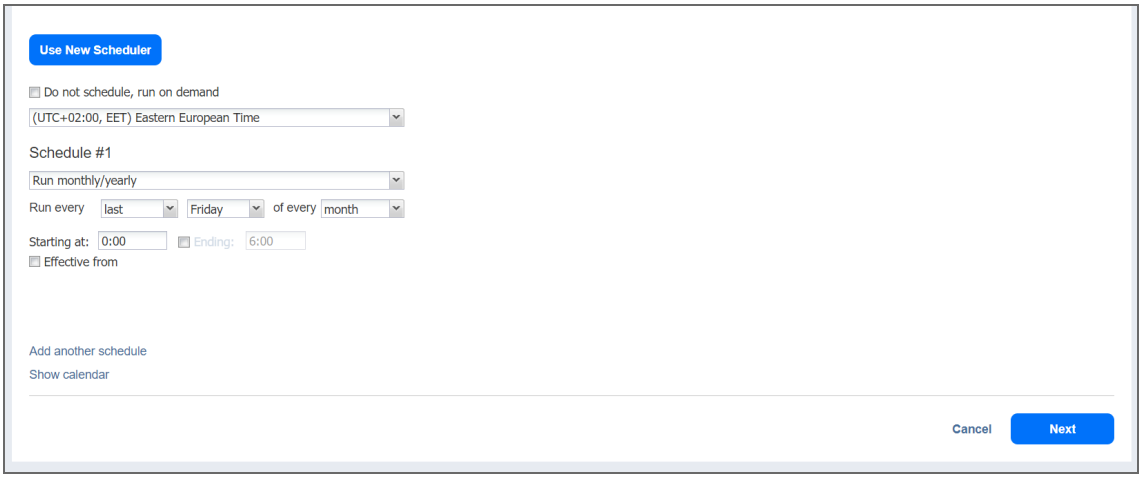
- If necessary, select the **Effective from** checkbox and pick the date when the schedule comes into effect.




## Monthly or Yearly Backup

To run the job monthly or yearly, choose **Monthly/yearly** from the schedule drop-down list and do the following:

- Specify the job start schedule in the appropriate boxes.
- Specify the day and month when the job should be started in the **Run every** boxes.
- Specify the time when the job should be started in the **Starting** box.
- Specify the end time for the job in the **Ending** box. If the job has not completed by the time specified, the job will be stopped.
- If necessary, select the **Effective from** checkbox and pick the date when the schedule comes into effect.




## Periodic Backup

To run the job multiple times per day, choose **Run periodically** from the schedule drop-down list and then choose a time period from the appropriate boxes:

- Specify the time when the job should be started in the **Starting** box.
- Specify the end time for the job in the **Ending** box. If the job has not completed by the time specified, the job will be stopped.
- Select the days of the week during which the job will be started.
- If necessary, select the **Effective from** checkbox and pick the date when the schedule comes into effect.




## Chained Job

To run the job after a previous one has completed, choose **Run after another job** from the schedule drop-down list and set the options as follows:

- **After the job**: Select a job after which the current job will be started.
- **Run this job**: Choose whether to run the current job immediately after the previous one has completed or within a delay.
- **After successful runs**: If selected, the job will run if the previous one has completed successfully.
- **After failed runs**: If selected, the job will run if the previous one has failed.
- **After stopped runs**: If selected, the job will run if the previous one has been stopped.
- **Effective from**: If selected, the schedule will come into effect on the date picked.

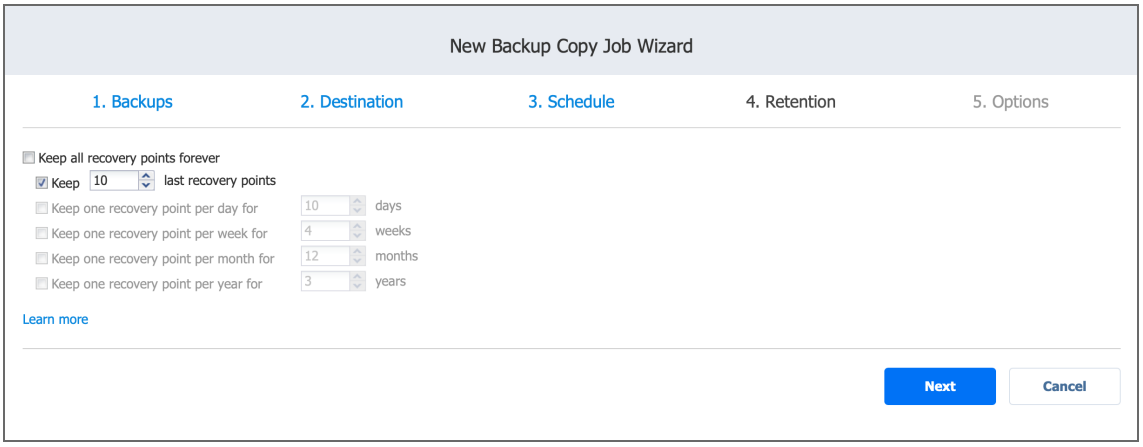


# Tape Backup Wizard: Retention

**Important**
This page is not displayed if the new scheduler is enabled.

Specify how many recovery points you need to be copied by the job. Use the following options:

- **Keep all recovery points forever**: When this option is selected, the job keeps all available recovery points until they are manually removed. To set a different retention policy, deselect this option and choose one of the options below.
- **Keep X last recovery points**: Keeps the specified number of last recovery points for each machine in the job.
- **Keep one recovery point per day for X days**: Keeps one last recovery point per day for the specified number of days.
- **Keep one recovery point per week for X weeks**: Keeps the last available backup of every week for the specified number of weeks.
- **Keep one recovery point per month for X months**: Keeps the last available backup of every month for the specified number of months.
- **Keep one recovery point per year for X years**: Keeps the last available backup of every year for the specified number of years.




For an example of job retention settings and further explanation, refer to the Keeping Recovery Points Knowledge Base article.

# Tape Backup Wizard: Options

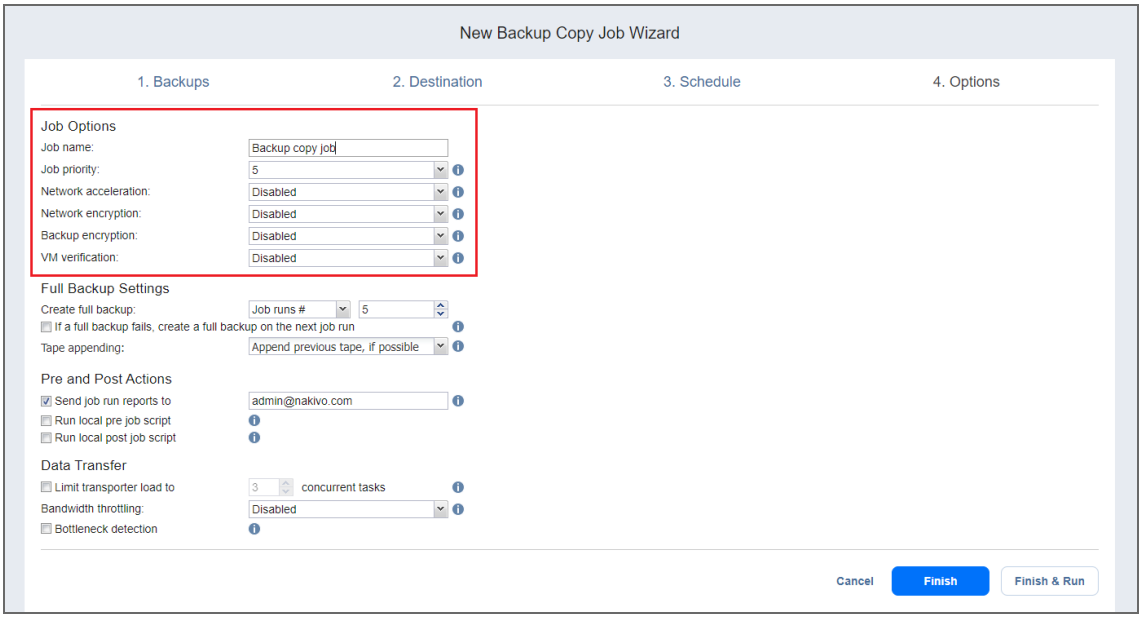
On the **Options** page of the wizard, you can name the job and enable/disable network acceleration and encryption. In addition, you can set up pre and post actions and enable bandwidth throttling.

- Job Options
  - Job Name
  - Job Priority

    - Network Acceleration
    - “Network Encryption” on the next page
    - “Backup Encryption” on the next page
- Tape Backup Settings
    - Create Full Backup
    - Tape Appending
- Pre and Post Job Actions
    - Email Notifications
    - Pre Job Script
    - Post Job Script
- Data Transfer
    - Bandwidth Throttling
    - Bottleneck Detection
- Completing Tape Backup Wizard


# Job Options

In this section, you can give a name to the tape backup job and enable/disable app-aware mode, change tracking, network acceleration, and encryption.




## Job Name

Specify a name for the tape backup job in the **Job name** box.

## Job Priority

Select a job priority level between 1 and 5, with 1 being the highest priority. Jobs with higher priority levels are prioritized by Transporters during job processing.

**Notes**

- This option is only available in the Enterprise, Enterprise Essentials, Enterprise Plus, MSP Enterprise, and MSP Enterprise Plus editions.
- If a tape backup job contains multiple backup objects, setting its **Job Priority** to 1 does not necessarily mean that the job finishes before lower priority jobs. This is because backup objects in Tape backup jobs are processed one by one, meaning that a Transporter is released after processing a given backup object. After the target tape cartridge for a subsequent backup object is found, a Transporter is assigned to the backup object with the highest priority.

## Network Acceleration

Once Network Acceleration is enabled, NAKIVO Backup & Replication uses compression and traffic reduction techniques to speed up data transfer. Select this option if you are backing up over WAN without a VPN connection. For more information, refer to Network Acceleration. If the source backup is already compressed, it is recommended to enable Network Acceleration and make sure that the source Backup Repository and target tape device have the same compression level. For more information, see this article.

**Note**

The **Network acceleration** option is not available if the **Backup encryption** option is enabled.

## Network Encryption

If the **Network Encryption** option is selected, backup data is protected with AES 256 encryption while traveling over the network. Data encryption increases the backup time and CPU load on machines running Transporters. Select this option if you back up over WAN without a VPN connection. For more information, refer to “Encryption in Flight and at Rest” on page 56.

**Note**

You need at least one Transporter at the source and target sites to enable network encryption.

## Backup Encryption

When enabled, backup data is protected with AES 256 block cipher encryption with a 256-bit key length. You can protect the backup file by creating a new password or selecting an existing one.

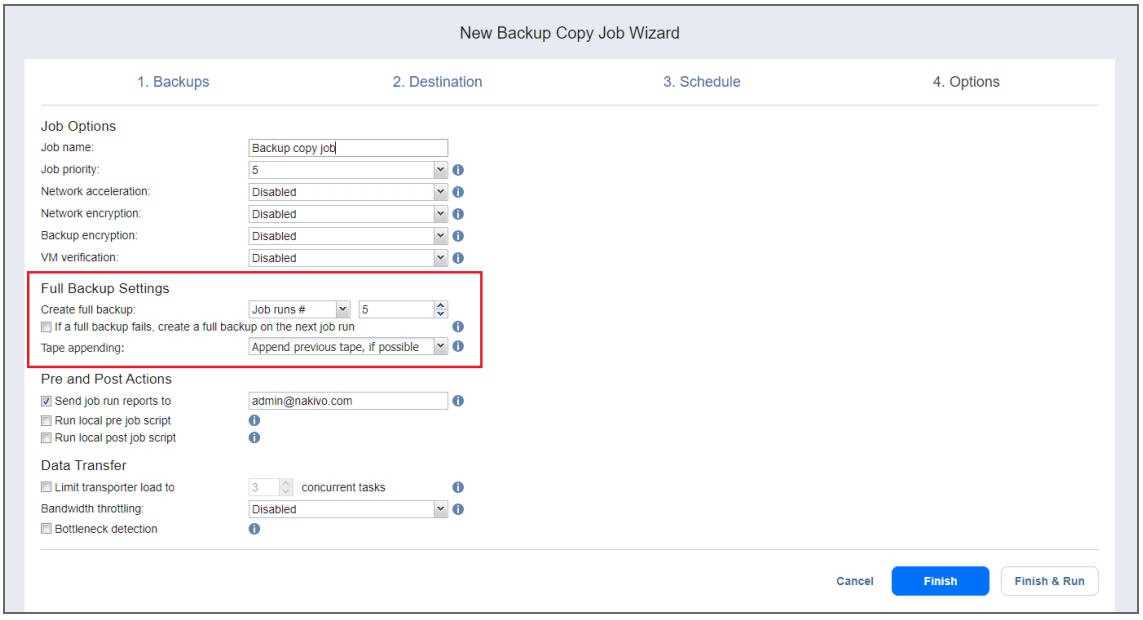
For more information, refer to Enabling Backup Encryption.

**Notes**

- This option is available only if the **Disk** or **Tape** destination type was chosen on the **Destination** page of the wizard.
- The **Backup encryption** option is not available if the **Network acceleration** option is enabled.
- If enabled, the created recovery points are encrypted.
- The **Backup encryption** option is not displayed for a backup job where forever incremental repositories are selected as the only target repositories.
- The **Backup encryption** option cannot be enabled if multiple targets with a mix of supported and unsupported (SaaS repositories or forever incremental repositories) repositories are selected as destinations.
- It's recommended that you enable the **(AWS) Key Management Service**. If **KMS** is enabled, all backup encryption passwords encrypted with the Key Management Service cryptographic key are available for recovery in case of product re-installation. For more information, refer to "Enabling KMS" on page 429.

## Tape Backup Settings

In this section, select when you want the system to create full backups and set up the rules for data appending.




## Creating Full Backup

With the **Create full backup** list, you can specify how often the system should perform a full (not incremental) backup to tape. The following options are available:

- **Always**: Every backup will be full
- **Every**: Select a day of the week. The full backup will be performed once a week on a specified day
- **Every 2nd**: Select a day of the week. The full backup will be performed once in 2 weeks on a specified day
- **First**: Select a day of the week. The full backup will be performed once per month starting with the first specified day of the month
- **Second**: Select a day of the week. The full backup will be performed once per month starting with the second specified day of the month
- **Third**: Select a day of the week. The full backup will be performed once per month starting with the third specified day of the month
- **Fourth**: Select a day of the week. The full backup will be performed once per month starting with the fourth specified day of the month
- **Last**: Select a day of the week. The full backup will be performed once per month starting with the last specified day of the month
- **Day #**: Select the day number. The full backup will be performed once per month on the specified day number
- **Job runs #**: Specify the number of backup jobs to pass before running a full backup

With the **If a full backup fails, create a full backup on the next job run** option selected, the next job run creates a full backup if the current job run fails to do so.

> **Note**
> Only **Always** and **Job runs #** options are always available for selection. The rest of the options' availability depends on the scheduling settings specified on the **Schedule** page.

## Tape Appending

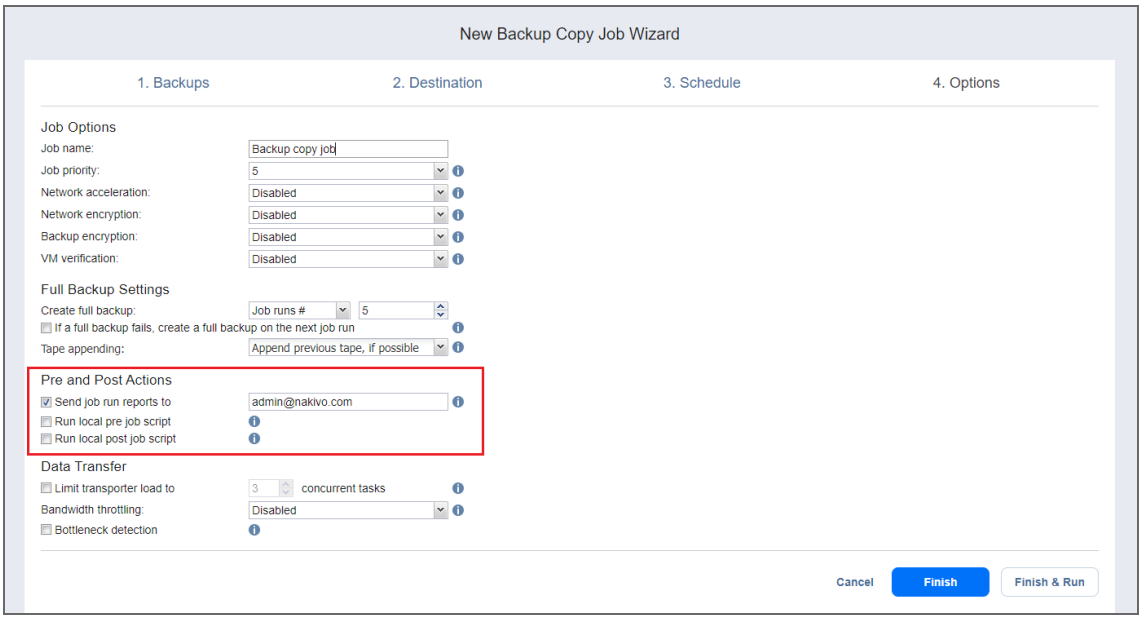
The **Tape appending** feature allows you to set up the rules for data appending. The following options are available:

- **Append previous tape if possible**:
  - The job run appends data to the last tape cartridge.
  - If the tape cartridge that was last written during the last job run is not available in the device or is full, the job starts with an empty cartridge:
    - All job objects within the job run are written to the selected tape cartridge one by one.
    - If the selected tape cartridge runs out of space, the next empty tape cartridge is selected and the process repeats until all job objects are written.
- **Start full backup with an empty tape**:

    - In case the backup modes of all job objects within the job run are defined as **full**:
        - The job run starts with an empty tape cartridge.
        - All job objects within the job run are written to the selected tape cartridge one by one.
        - If the selected tape cartridge runs out of space, the next empty tape cartridge is selected and the process repeats until all job objects are written.
    - In case the backup modes of all job objects within the job run are defined as **incremental**:
        - The job run appends data to the last tape cartridge.
        - All job objects within the job run are written to the selected tape cartridge one by one.
        - If the selected tape cartridge runs out of space, the next empty tape cartridge is selected and the process repeats until all job objects are written.
    - In case the backup modes of all job objects within the job run are defined as a **mix of full and incremental** modes:
        - The job run appends data to the last tape cartridge.
        - All job objects within the job run are written to the selected tape cartridge one by one.
        - If the selected tape cartridge runs out of space, the next empty tape cartridge is selected and the process repeats until all job objects are written.
- **Always start with an empty tape**:
    - The job starts with an empty tape cartridge.
    - All job objects within the job run are written to the selected tape cartridge one by one.
    - In case the selected tape cartridge runs out of space, the next empty tape cartridge is selected and the process repeats until all job objects are written.

## Pre and Post Job Actions

NAKIVO Backup & Replication provides you with the ability to enable certain actions before a tape backup job begins and after it has completed. You can choose to send job run reports to the email provided and run local pre and post job scripts.




## Email Notifications

NAKIVO Backup & Replication can send email notifications on job completion status to specified recipients. This feature complements global notifications and provides you with the ability to configure notifications on a per-job level.

To enable this option, configure your Email settings.

To send email notifications, do the following:

- In the *Pre- and Post- Actions* section, select the **Send job run reports to** option and specify one or more email addresses in the text box. Separate multiple email addresses with a semicolon.

## Pre Job Script

To run a script before the product begins copying backups, do the following:

1. Place a script file on the machine on which the Director is installed.
2. In the *Pre and Post Actions* section, select the **Run local pre job script** option and click the **settings** link. Specify the following parameters in the dialog that opens:

- **Script path**: Specify a local path to the script on the machine on which the Director is installed. Script interpreter should be specified.
  Example (Windows): *cmd.exe /c D:\script.bat*
  Example (Linux): *bash /root/script.sh*
- **Job behavior**: Choose either of the following job behaviors in relation to script completion:
  - **Wait for the script to finish**: If this option is selected, tape backup will not be started until the script is completed.
  - **Do not wait for the script to finish**: If this option is selected, the product will run the script and will start copying backups at the same time.
- **Error handling**: Choose either of the following job behaviors in relation to script failure:
  - **Continue the job on script failure**: If this option is selected, the job will perform tape backup even if the script has failed.
  - **Fail the job on script failure**: If this option is selected and the script fails, the job will be failed and tape backup will not be performed.

## Post Job Script

To run a script after the product has finished copying all backups, do the following:

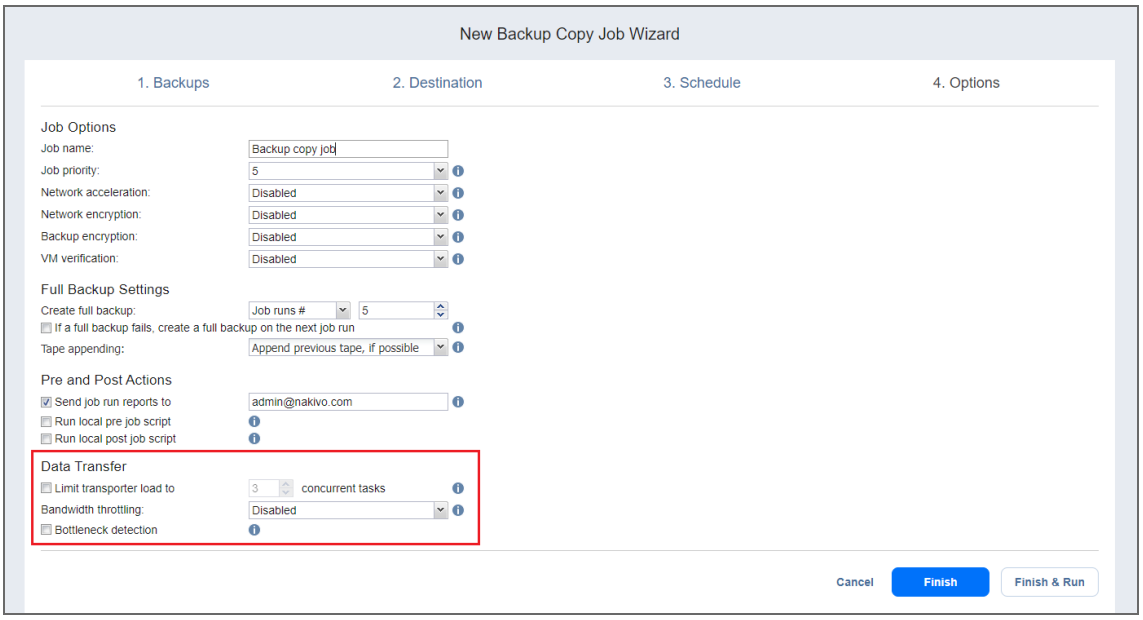
1. Place a script file on the machine on which the Director is installed.
2. In the *Pre and Post Actions* section, select the **Run local post job script** option and click the **settings** link. Specify the following parameters in the dialog box that opens:

- **Script path**: Specify a local path to the script on the machine on which the Director is installed. Script interpreter should be specified.
  Example (Windows): *cmd.exe /c D:\script.bat*
  Example (Linux): *bash /root/script.sh*
- **Job behavior**: Choose either of the following job behaviors in relation to script completion:
  - **Wait for the script to finish**: If this option is selected, the job is in the “running” state until the script is completed.
  - **Do not wait for the script to finish**: f this option is selected, the job is completed even if the script execution is still in progress.
- **Error handling**: Choose either of the following job behaviors in relation to script failure**:**
  - **Continue the job on script failure**: If this option is selected, script failure does not influence the status of the job.
  - **Fail the job on script failure**: If this option is selected and the script has failed, the job status is set to “failed” even if VM backup has been successful.

**Important**
Pre and post job scripts can be executed only on the machine where the Director is installed.

# Data Transfer

In the *Data Transfer* section of the Options page, you can configure bandwidth throttling.




## Bandwidth Throttling

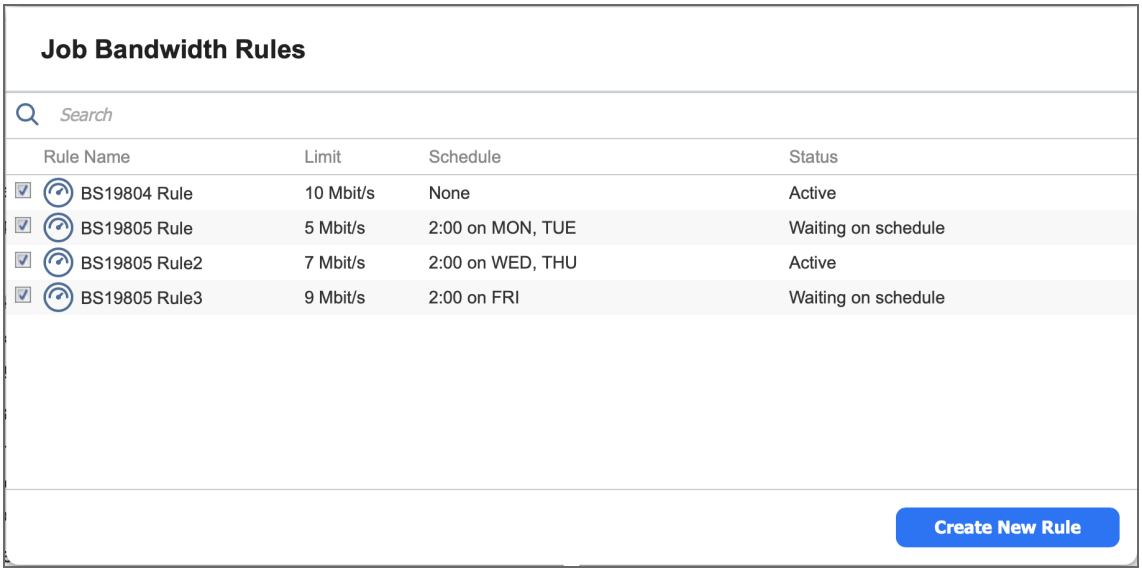
Follow the steps below to regulate the speed of data transfer over the network for your tape backup job:

1. For the **Bandwidth throttling** option, choose **Enabled**.

> **Note**
> If bandwidth throttling is disabled for the current job, global bandwidth rules may still apply to your job. Refer to "Bandwidth Throttling" on page 407 for details.

2. Click the **Settings** link that becomes available.
3. The *Job Bandwidth Rules* dialog box opens displaying you the list of available rules. You have the following options:
   - Create a new bandwidth rule for your tape backup job:
     1. Click the **Create New Rule** button.
     2. The *New Bandwidth Rule* dialog box opens. Refer to "Bandwidth Throttling" on page 407 for details on creating a bandwidth rule.
     3. Click **Save**.
   - Activate an existing bandwidth rule for your job. Select the checkbox to the left of the necessary bandwidth rule. To deactivate a bandwidth rule for your job, clear the corresponding checkbox.
   - Edit a bandwidth rule. Click the **Edit** link for a bandwidth rule and modify it in the *Edit Bandwidth Rule* dialog box that opens.
   - Disable a bandwidth rule. Click the **Disable** link. The bandwidth rule is disabled for all jobs.

- Remove a bandwidth rule. Click the **Remove** link and then click **Delete** to confirm your operation.




## Bottleneck detection

When the **Bottleneck detection** option is enabled, additional information is collected and recorded in NAKIVO Backup & Replication logs in the course of data transfer for the purpose of bottleneck detection. Check this option to enable the **Bottleneck detection** capability of the physical machine agent engaged in the job.

## Completing Tape Backup Wizard

Click **Finish** or **Finish & Run** to complete the job creation.

> **Note**
> If you click **Finish & Run**, you have to define the scope of your job. Refer to "Running Jobs on Demand" on page 364 for details.

# Staging (Seeding) Initial Backup

Since initial backups are often large, the first backup job run can be slow and time-consuming when done over WAN, and it can also put an undesirable load on the network when done over LAN. To speed up the initial backup and save network bandwidth, you can perform a staged backup (seed backups): run the initial backup on a removable media device (such as an external USB hard drive), transfer the media to a new location, and then run an incremental backup to the new location.
To perform a staged backup, follow the steps below:

1. Create a new Backup Repository.
2. Create and run a new backup job.
3. After the job has completed, move the Backup Repository to a new location.
4. If required, edit the backup job and specify a schedule for the backup job.

# Deleting Backups

With NAKIVO Backup & Replication, you can permanently delete a backup with all of its recovery points if this backup is available in a Backup Repository. You can also delete specific recovery points in a backup without affecting any of the other recovery points. The option to delete a specific recovery point can be used if you get an alert about corrupted recovery points in a backup.

> **Note**
> You can delete a backup only if you have deleted the corresponding backup job or edited the backup job to not include the backup's source VM or physical machine.

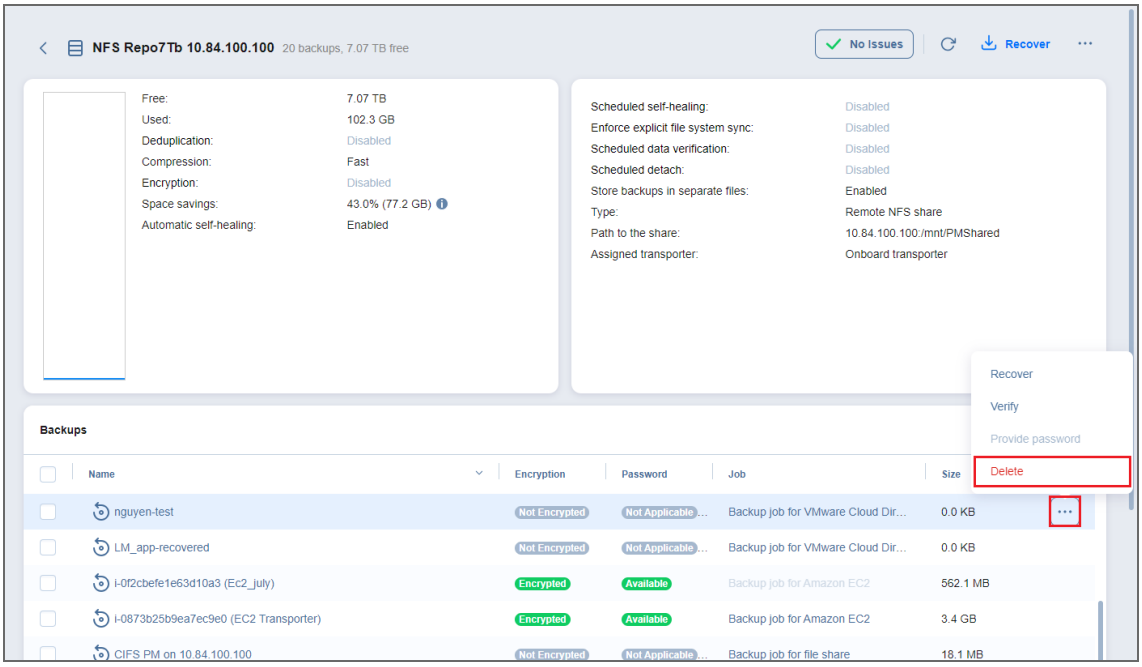
Refer to one of the following sections:

- Deleting a Single Backup
- Deleting Backups in Bulk
- Deleting Recovery Points
  - Deleting a Single Recovery Point
  - Bulk Recovery Points Deletion

## Deleting a Single Backup

To delete a backup permanently, follow the steps below:

1. In the main menu, click **Settings**.
2. Go to the **Repositories** tab and click the Backup Repository you need.
3. Hover over the backup you want to delete, and on the right side, click **Delete**.




4. Click **Delete** in the dialog box that opens.

**Notes**

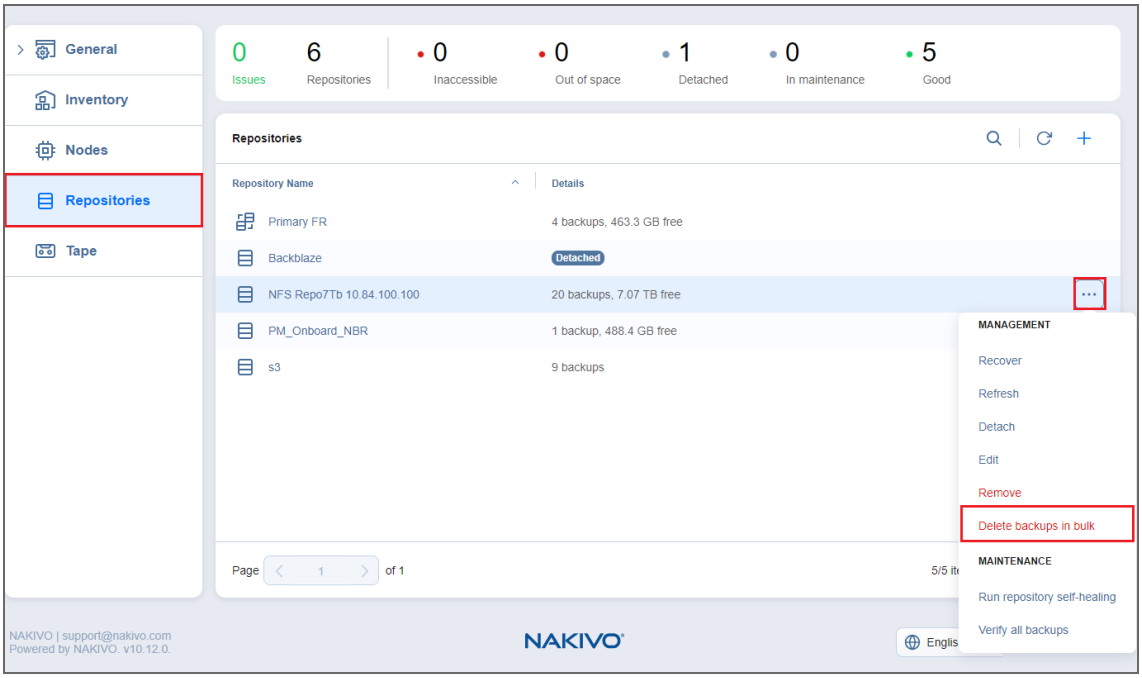
- You cannot delete a backup if there is a job that backs up to it. To remove such a backup, delete (or edit) the corresponding jobs so no items are backed up to the aforementioned backup.
- For a **Forever-incremental** Backup Repository (that is, when the **Store backups in separate files** option is not selected), the space that was occupied by the deleted backup is marked “free” and reused by new data blocks on subsequent job runs. However, the actual size of the Backup Repository may not change. To free up the space that was occupied by the deleted backup, you can reclaim the free space.

For SaaS Backup Repositories, manually removing backup data may not return space to the operating system correctly.

# Deleting Backups in Bulk

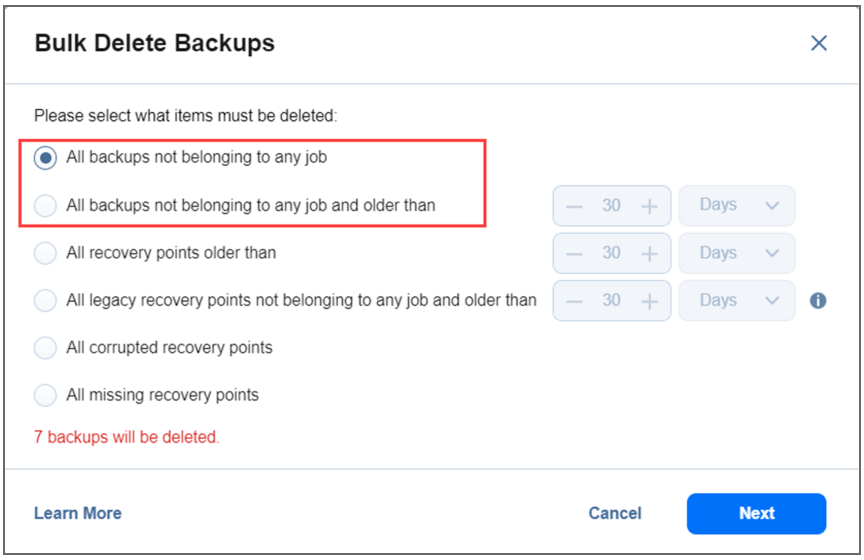
To permanently delete several backups that match specific criteria, follow the steps below:

1. In the main menu, click **Settings**.
2. Go to the **Repositories** tab and hover over the Backup Repository you need.
3. Click the ellipsis **Manage** button and then click **Delete backups in bulk**.

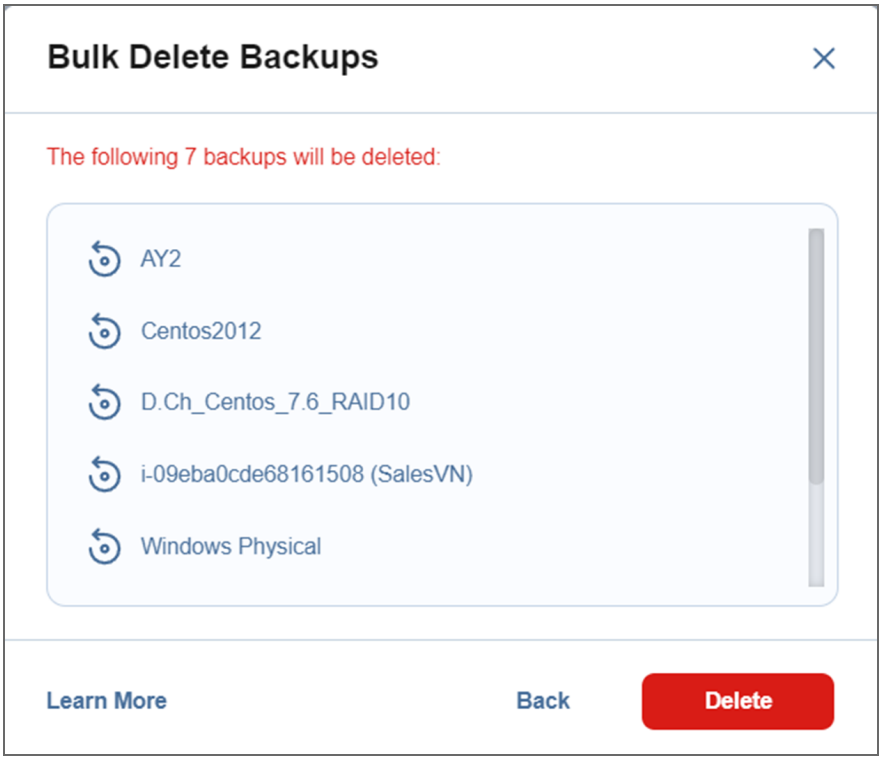



4. In the **Bulk Delete Backup** dialog box that opens, select one of the available options:
   - **All backups not belonging to any job**
   - **All backups not belonging to any job and older than** X <time_units>, where X is an integer and <time_units> is either days, weeks, or months
   - **All recovery points older than** X <time_units>, where X is an integer and <time_units> is either days, weeks, or months
   - **All legacy recovery points not belonging to any job and older than** X <time_units>, where X is an integer and <time_units> is either days, weeks, or months
   - **All corrupted recovery points**
   - **All missing recovery points**

   The dialog also shows the number of backups to be deleted.




5. Click **Next**.
6. The **Bulk Delete Backups** dialog box opens displaying the list of backups to be deleted. Click **Delete** to confirm the deletion.



**Note**

For a **Forever-incremental** Backup Repository (that is, when the **Store backups in separate files** option is not selected), the space that was occupied by the deleted backup is marked "free" and reused by new data blocks on subsequent job runs. However, the actual size of the Backup Repository may not change. To free up the space that was occupied by the deleted backup, you can reclaim the free space.

For SaaS Backup Repositories, manually removing backup data may not return space to the operating system correctly.

# Deleting Recovery Points

You can select to delete a single recovery point, all corrupted recovery points, or all recovery points older than a specified number of days.

**Note**

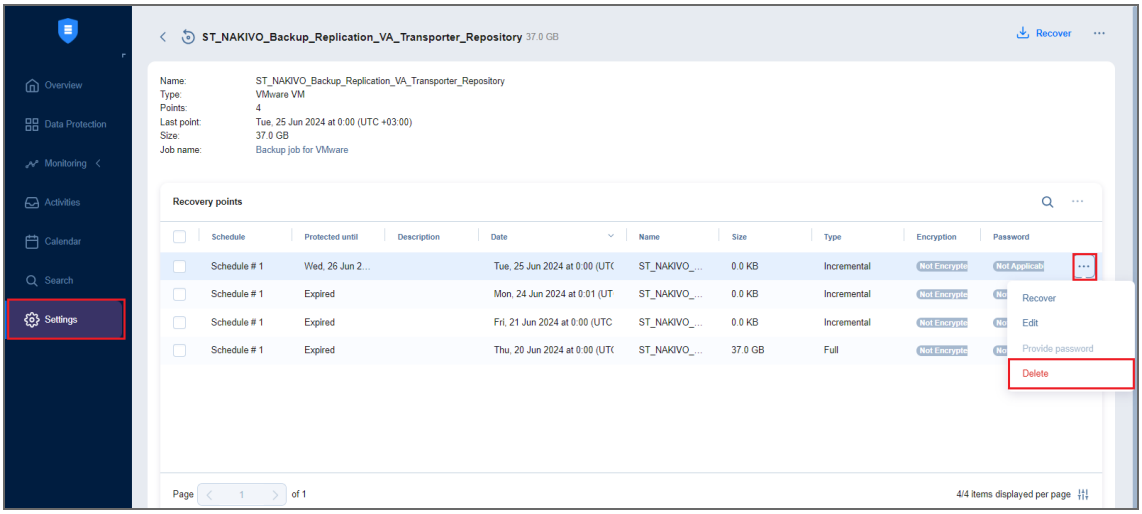
The **Delete** action is disabled:

- For VMware Cloud Director VMs underlying vApps
- For the only remaining and/or uncorrupted recovery point, except in the cases when the recovery point is set to be removed automatically according to the configured retention policy. In this case, removal of the last recovery point removes the corresponding backup object and data pertaining thereto.
- For the recovery point if there is at least one other recovery point depending on the current recovery point
- For inaccessible backup objects
- For pending removal recovery points
- For inconsistent repositories
- For the recovery point that is currently in use.

## Deleting a Single Recovery Point

To delete a single recovery point in response to a corruption alert or for functional requirements, do the following:

1. In the main menu, click **Settings**.
2. Go to the **Repositories** tab and click the Backup Repository you need.
3. Click the backup with the recovery point you want to delete.
4. Hover over the recovery point that you want to delete, and on the right side, click **Delete**.



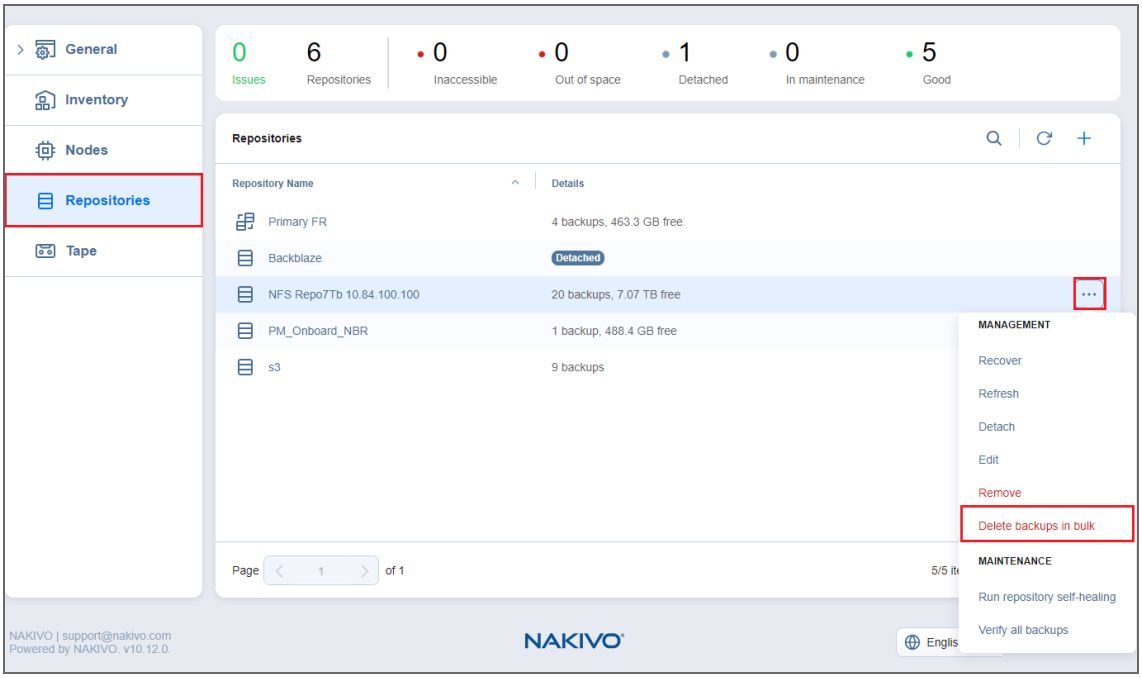

5. Click **Delete** in the dialog box that opens.

**Note**

For a **Forever-incremental** Backup Repository (that is, when the **Store backups in separate files** option is not selected), the space that was occupied by the deleted recovery point is marked “free” and reused by new data blocks on subsequent job runs. However, the actual size of the Backup Repository may not change. To free up the space that was occupied by the deleted recovery point, you can reclaim the free space.

# Bulk Recovery Points Deletion

1. In the main menu, click **Settings**.
2. Go to the **Repositories** tab and click the Backup Repository you need.
3. Click **Manage** and then click **Delete backups in bulk**.




4. In the **Bulk Delete Backups** dialog box that opens, select criteria for recovery points to be deleted:
   - **All recovery points older than** X <time_units>, where X is an integer and <time_units> is either days, weeks, or months. When selected, the recovery points that are older than the specified time interval are deleted.

     **Notes**

     The following deletion exclusions are applicable:

     - For **Forever-incremental** repositories (that is, when the **Store backups in separate files** option is not selected): If all recovery points of a backup match the deletion criteria, the latest recovery point – whether corrupted or not – is not deleted.
     - For **Incremental-with-full-backups** repositories (that is, when the **Store backups in separate files** option is selected):
       - Recovery points that are older than the end of the time interval that have dependent recovery points that are newer than the beginning of the time interval are not deleted.
       - If all recovery points in a backup match the deletion criteria, the latest full recovery point – whether corrupted or not – is not deleted.

   - **All corrupted recovery points**: When this option is selected, all recovery points that are corrupted are deleted. Recovery point selection criteria include the following:

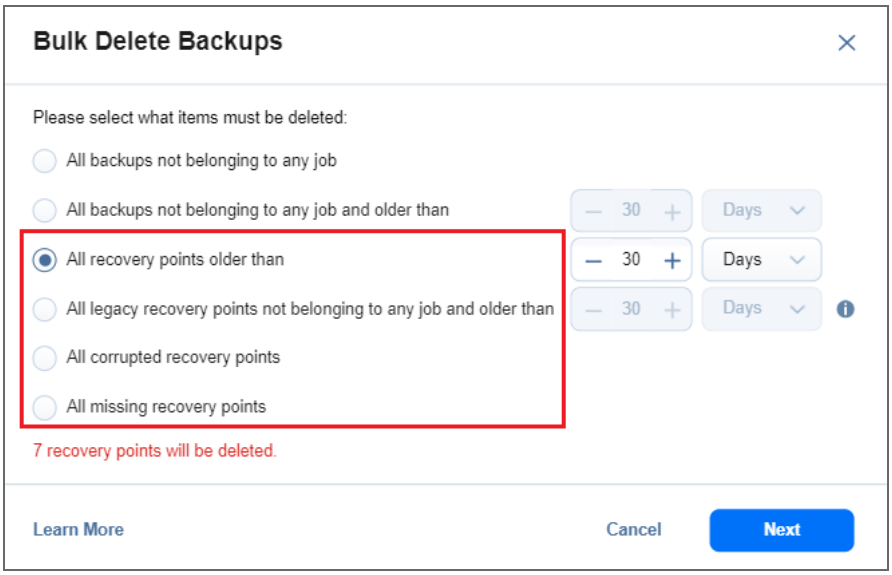
- For **Forever-incremental** repositories (that is, when the **Store backups in separate files** option is not selected), if a backup is used by a backup job and all its recovery points are corrupted, the latest recovery point is not deleted.
- For **Incremental-with-full-backups** repositories (that is, when the **Store backups in separate files** option is selected), this option also deletes all recovery points that are dependent on corrupted recovery points. If all recovery points in a backup are corrupted or depend on a corrupted recovery point and match the deletion criteria, the latest full recovery point is not deleted.

> **Note**
> This option is not available for Microsoft 365 backups.

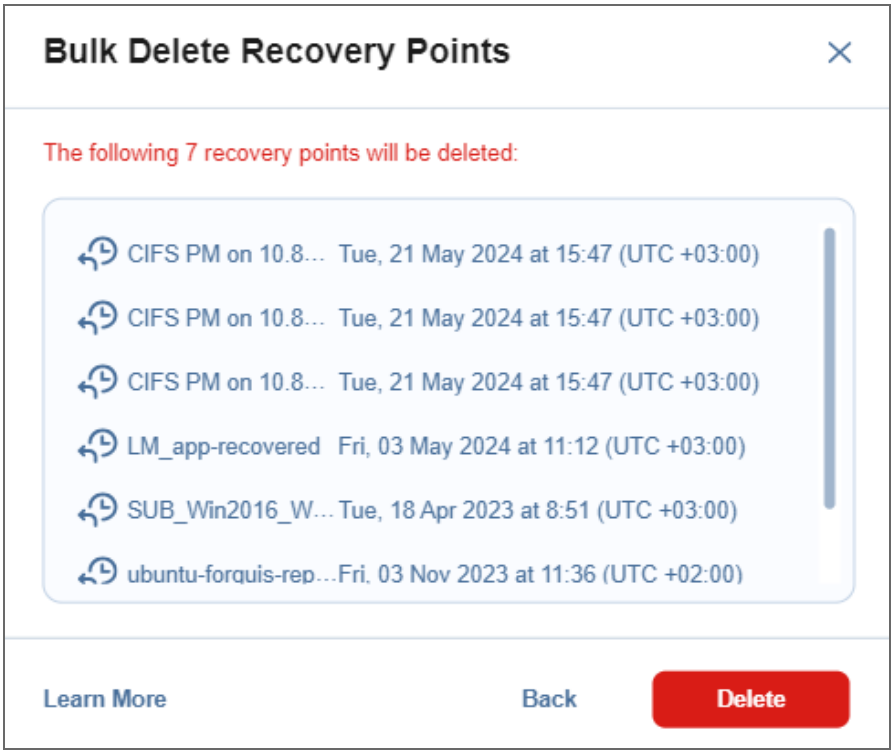
- **All missing recovery points**: When selected, all missing recovery points are deleted. Recovery point selection criteria include the following:
  - For **Forever-incremental** repositories (that is, when the **Store backups in separate files** option is not selected), this option deletes all missing recovery points. If all recovery points in a backup are missing, the latest recovery point is not deleted.
  - For **Incremental-with-full-backups** repositories (that is, when the **Store backups in separate files** option is selected), this option deletes all missing recovery points and any recovery points that are dependent on them. If all recovery points in a backup are missing or depend on missing recovery points, the latest full recovery point is not deleted.

> **Note**
> The recovery points of a missing backup object are not accounted as missing recovery points.

The dialog box shows the number of recovery points to be deleted.



5. The **Bulk Delete Recovery Points** dialog box opens displaying the list of recovery points to be deleted. Click **Delete** to confirm deleting the recovery points.



# Replication

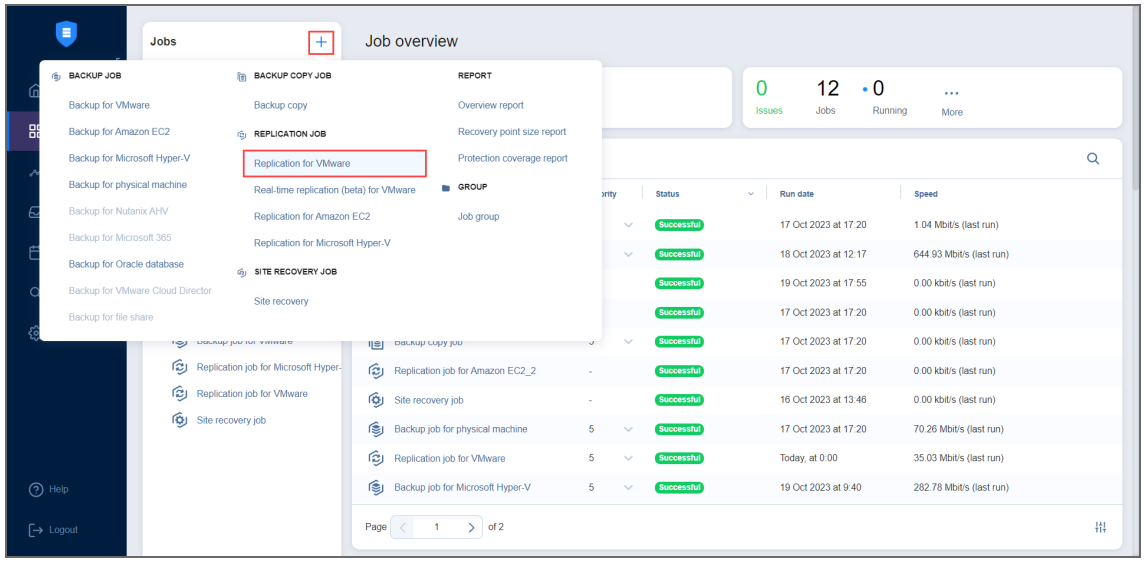
With NAKIVO Backup & Replication, you can perform replication of virtual machines. Replication creates and maintains an identical copy of the source VM at the target location.

Refer to the following topics for more information:


- “Creating VMware Replication Jobs” on page 934
- “Creating VMware Real-Time Replication Jobs” on page 968
- “Staging (Seeding) VM Replication” on page 989

# Creating VMware Replication Jobs

With NAKIVO Backup & Replication, you can replicate VMware VMs using the workflow with a wide array of available options. To create a replication job, click the plus **Create** button in the **Data Protection** menu and then click **Replication for VMware**.





- “Replication Job Wizard for VMware: Source” on page 935
- “Replication Job Wizard for VMware: Destination” on page 940
- “Replication Job Wizard for VMware: Networks” on page 944
- “Replication Job Wizard for VMware: Re-IP” on page 947
- “Replication Job Wizard for VMware: Schedule” on page 950
- “Replication Job Wizard for VMware: Retention” on page 955
- “Replication Job Wizard for VMware: Options” on page 956

# Replication Job Wizard for VMware: Source

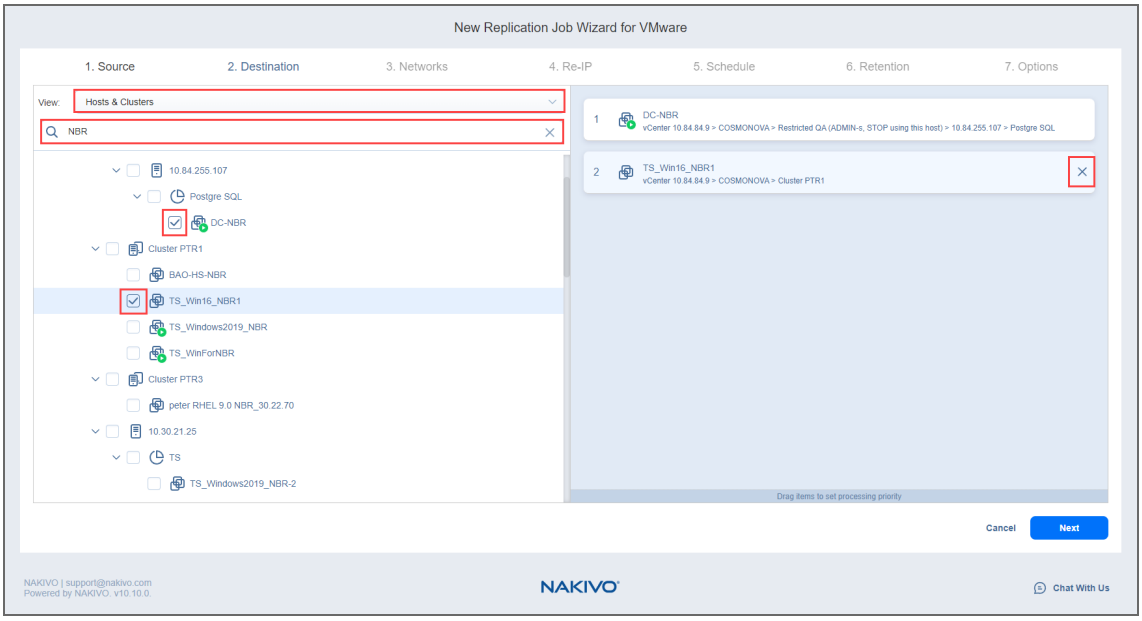
On the **Source** page of the wizard, select one of the views to add VMware items to your replication job.

- Host and Clusters
- VMs and Templates
- Jobs and Groups
- Backup Repositories
- Policy

## Host and Clusters

When the **Host & Clusters** view is chosen, the inventory tree opens in the left pane and displays all VMware items: clusters, hosts, folders, resource pools, and VMs. Proceed as follows:

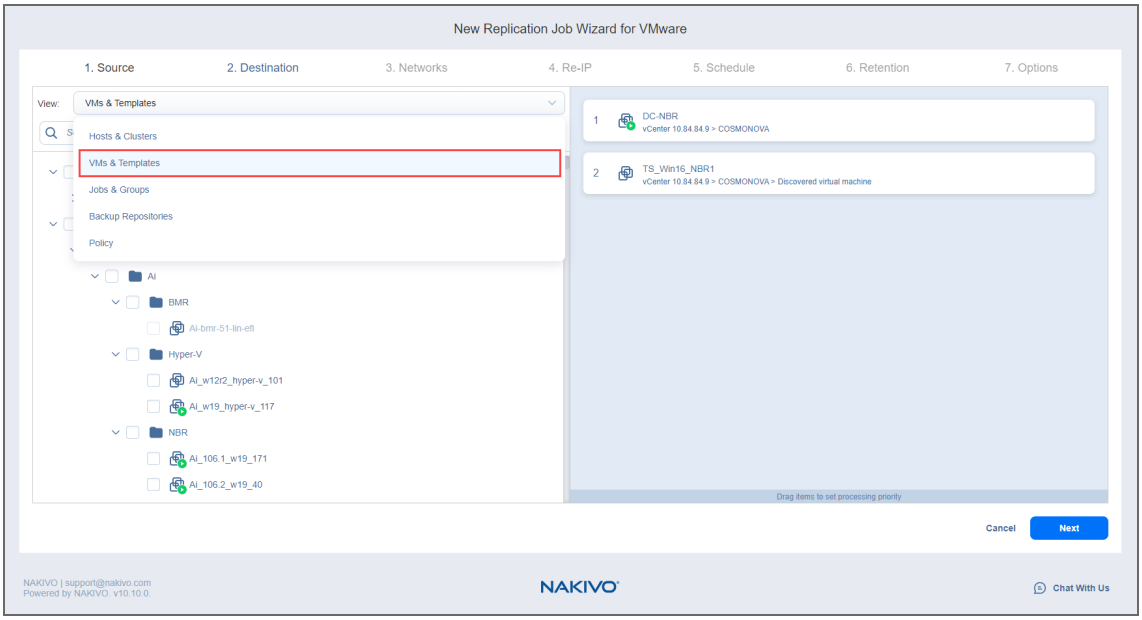
1. Optionally, you can filter the inventory tree by entering a string into the **Search** box. You can enter a part of or the entire name of the item.
2. Select VMware items by placing a checkmark to the left of each item.
3. The selected items appear in the right pane of the page. You can reorder the selected items by dragging a VM or a container with the pointer to a new position. By doing that, you can specify to replicate the most important VMs first.
4. Review the list of selected VMware items. You can remove a selected VM or a container from the replication job in one of the following ways:
    - Deselect the item in the left pane. This will remove the item from the right pane; OR
    - In the right pane, hover the pointer over the item you want to remove and click **x** on the right. This will deselect the item in the left pane.



## VMs and Templates

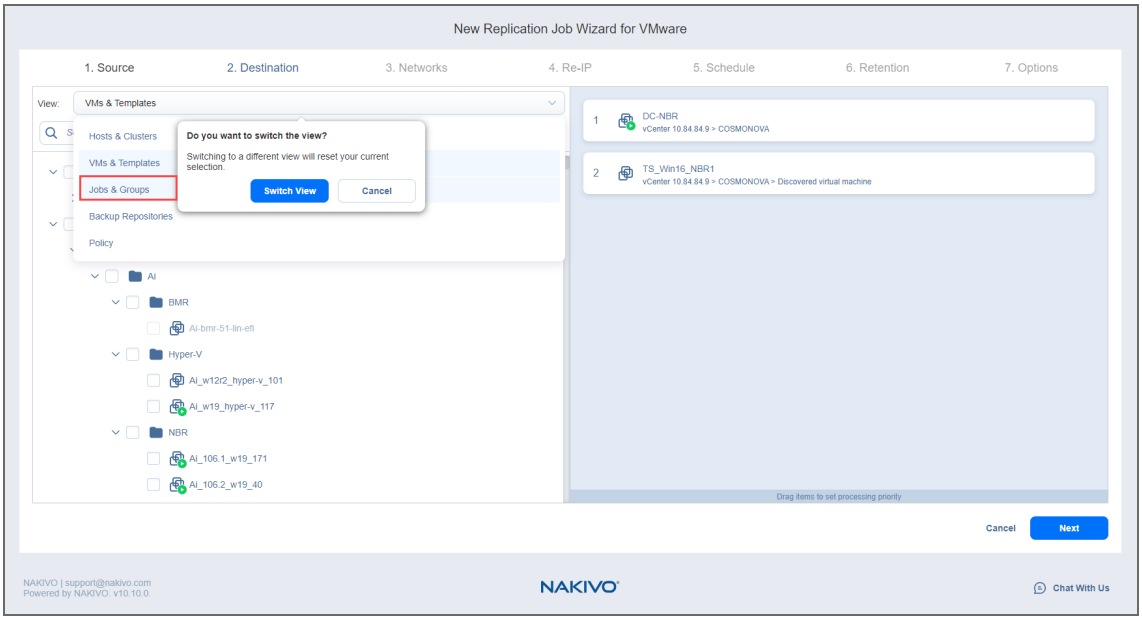
When the **VMs & Templates** view is selected, the inventory tree displays VMware hosts, VMs, and VM templates. Proceed as described for the **Hosts & Clusters** view above.

> **Note**
>
> Replication of VM templates is not supported.




## Jobs and Groups

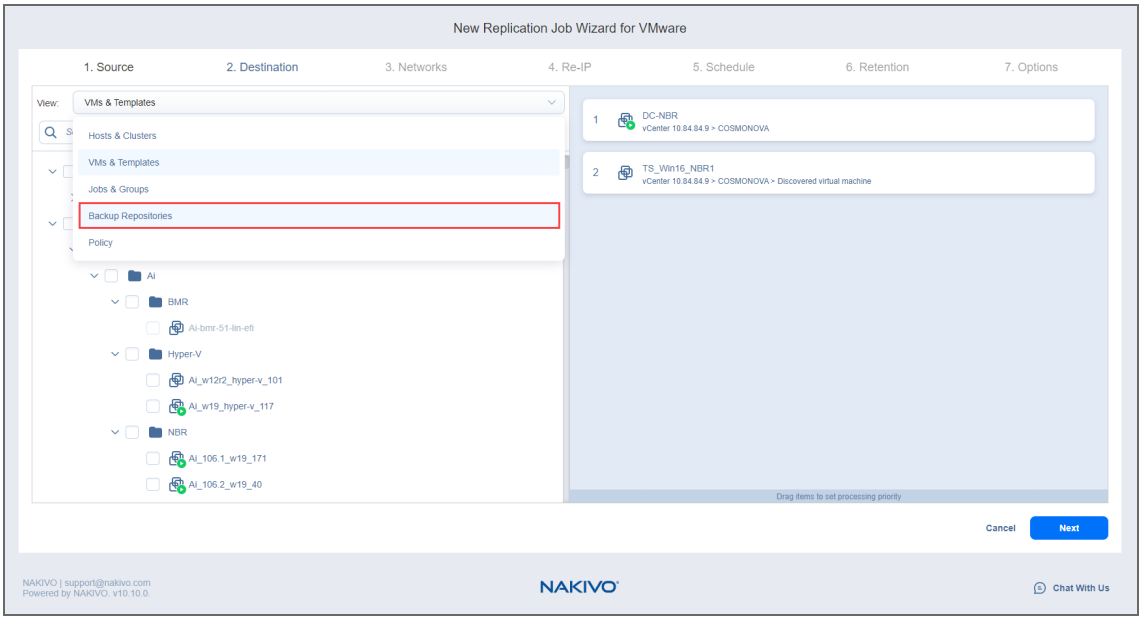
When the **Jobs & Groups** view is selected, the inventory tree displays groups, jobs, and backups of the appropriate hypervisor. Selecting a backup from the list allows you to replicate VMs directly from the backup (Refer to “Replication From Backup” on page 37). Proceed as described for the **Hosts & Clusters** view above.




# Backup Repositories

When the **Backup Repositories** view is selected, the inventory tree displays backup repositories that contain backups of the appropriate hypervisor. Proceed as described for the **Hosts & Clusters** view above.

> **Note**
>
> Replication of VM templates is not supported.




Both federated and standalone (not used as members of federated repositories) backup repositories can be selected.

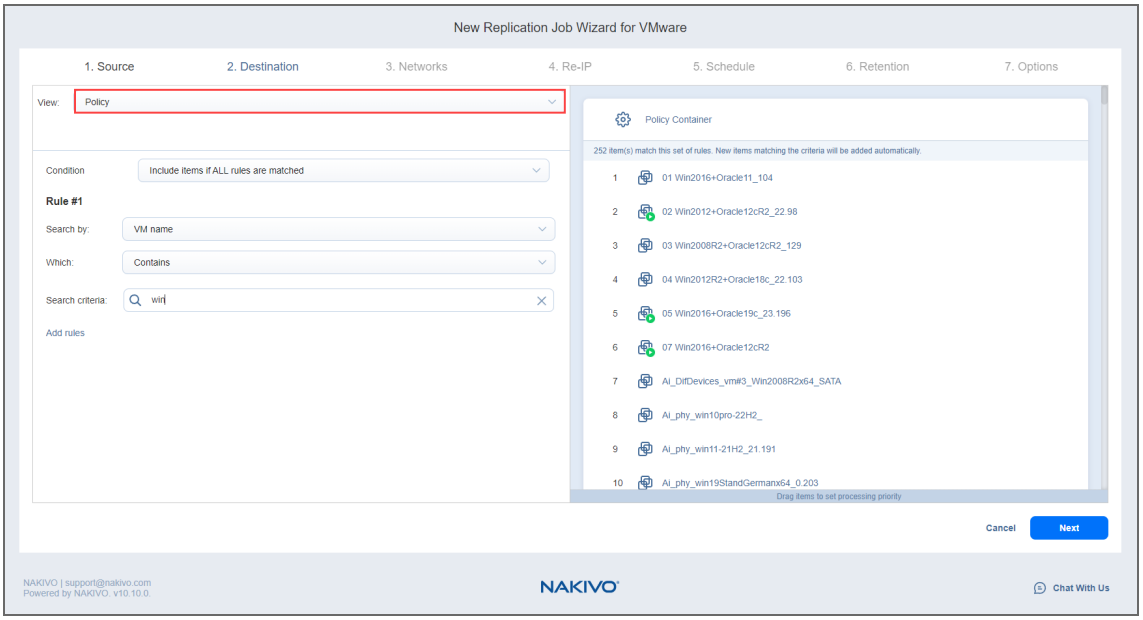
**Note**

You cannot select a federated repository member as a source for replication jobs.

Selecting an item in the tree adds the parent item and all children items to the right pane.

## Policy

When the **Policy** view selected, job policies can be used. Refer to “Managing Job Policies” on page 382 for details. If the items were selected in alternate views, a dialog box opens warning you that switching to the **Policy** view will reset your current selection. Click **Switch View** to confirm switching to the **Policy** view. Make sure that at least one item matches the available set of policy rules. Refer to “Managing Policy Rules” on page 385 for details.




Click **Next** to confirm adding selected VMs to the replication job. The wizard will display the next page.

**Notes**

- If you cannot find a VM or a container:
  - Make sure the corresponding vCenter or ESX(i) host has been added to the inventory.
  - Refresh the Inventory.
- By adding a VMware container to the job, you ensure that important VMs are always protected. If you add a VMware container to the job:
  - All VMs currently available in the selected container will be replicated.
  - All new VMs that will be created in (or moved to) the container in the future will be automatically added to the job and replicated.
- The order in which VMs are replicated is important if the **Transporter** performing replication cannot process all VMs of the job simultaneously — either because the Transporter is processing other jobs at the same time or because the job contains more VM disks than a **Transporter**'s maximum load specified during the **Transporter** creation.
- If all disks of a VM are unsupported (such as RDM disks in physical compatibility mode), this VM will be disabled in the inventory tree and it will not be possible to select it. Refer to "Supported Platforms" on page 131 for details.
- VMs marked as primary are protected by VMware Fault Tolerance. Disabled VMs are Fault Tolerance shadow VMs.

# Replication Job Wizard for VMware: Destination

On the **Destination** page, select a location for your replica(s).


- "Setting the Same Host, Datastore, and Network for All Replicas" below
- "Setting the Default Destination for Replicas" on the next page
- "Setting Different Options for Replicas" on the next page
- "Mapping Source VMs to Existing Replicas" on page 942
- "Excluding a VM Disk from the Replication Job" on page 943

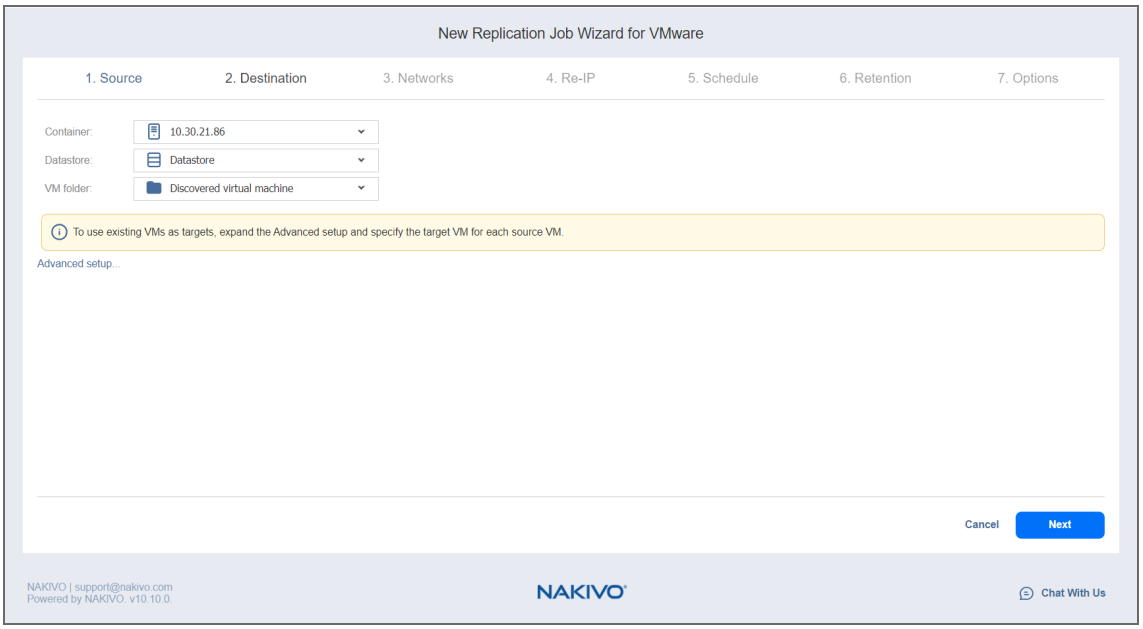

**Note**

When replicating an ESXi 8.0 VM to an ESXi 7.0 host, the **VirtualMachineConfigSpec** > **VirtualMachineSgxInfo** property is skipped.

## Setting the Same Host, Datastore, and Network for All Replicas

To replicate all VMs to the same container and datastore, and to connect all replicas to the same network, follow the steps below:

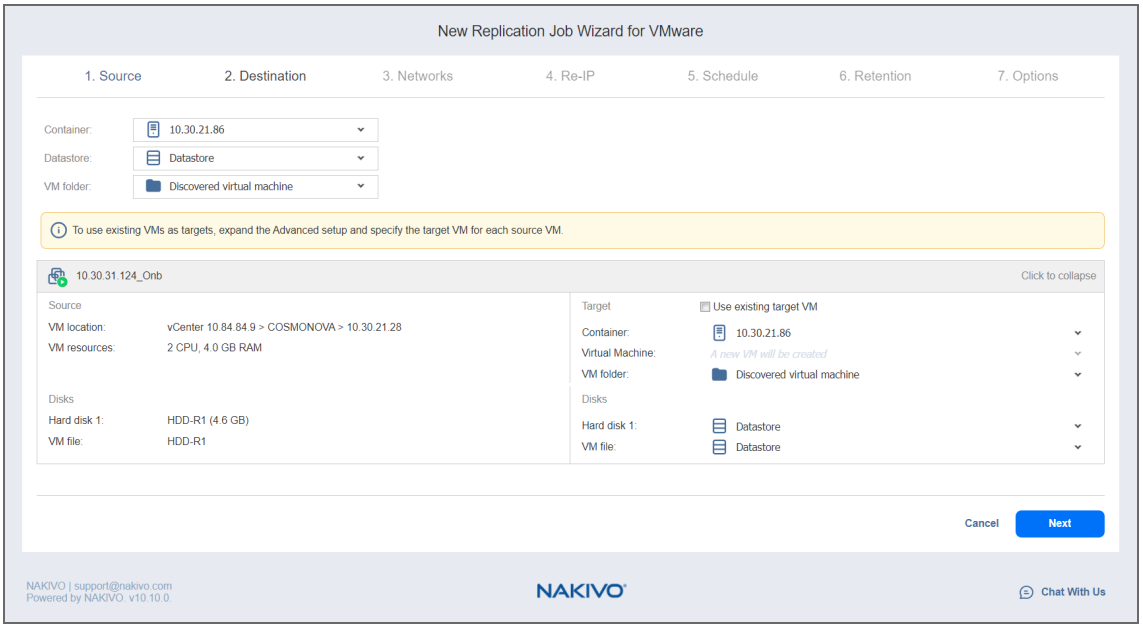
1. Choose a cluster, host, or resource pool from the **Container** drop-down list.
2. Choose a datastore from the **Datastore** drop-down list.
3. If you are creating a Replication job from an existing Backup job (**Jobs & Groups** or **Backup repository** view), select a target network from the **Network** drop-down list.
4. Optionally, you can choose a target VM folder from the **VM folder** drop-down list.



# Setting the Default Destination for Replicas

If you have chosen a host, cluster, folder, or a resource pool as a source for your replication job on the **Source** wizard page, you can set the default container, datastore, and VM folder for replicas. To do this, follow the steps below:

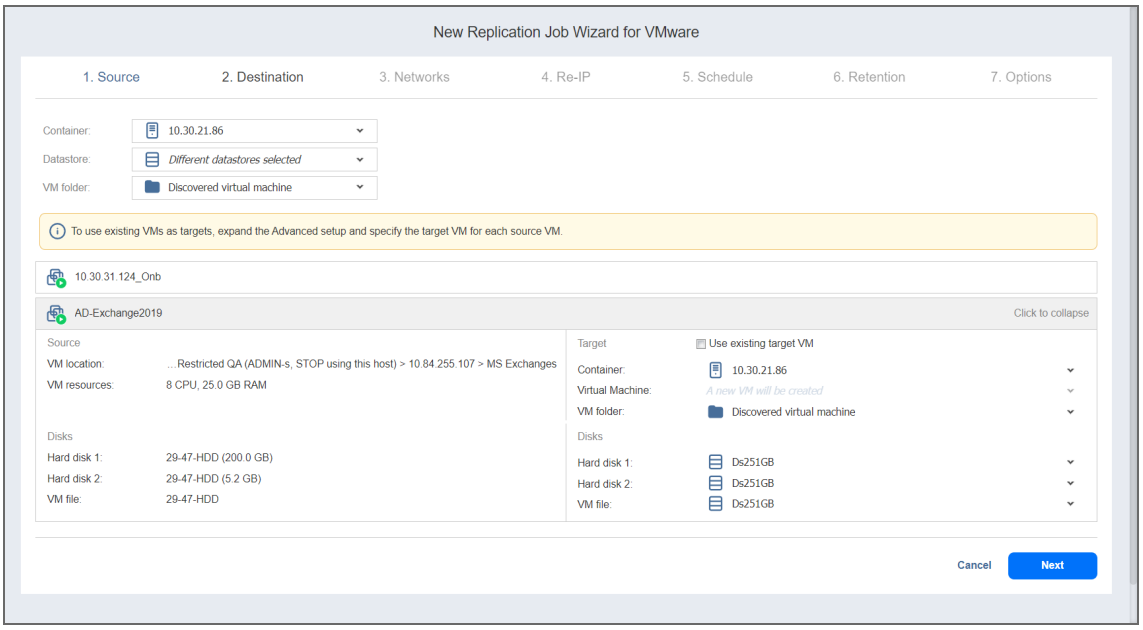
1. Click **Advanced setup** and then click on the name of the chosen host, cluster, folder, or a resource pool.
2. Choose a **Default container**.
3. If you have chosen the backup job on the **Source** page, you can choose a **Default Network**.
4. Optionally, you can also choose a **Default VM folder.**




# Setting Different Options for Replicas

To specify different replication options for VMs, follow the steps below:

1. Click **Advanced setup**.
2. Choose a target container, target VM, and target datastore for each VM.

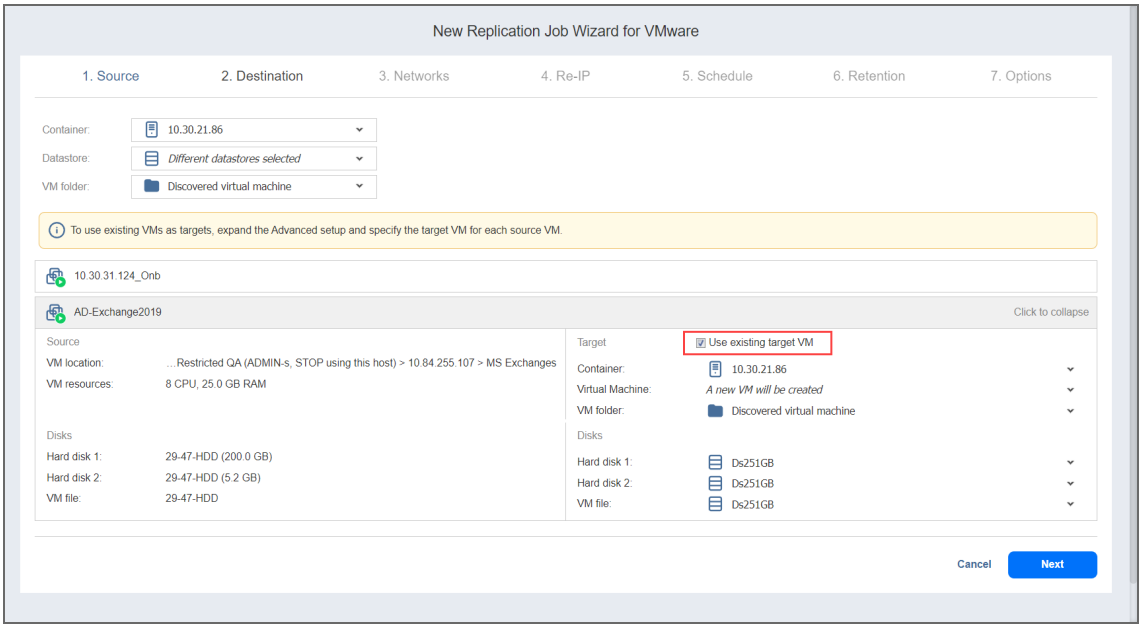



# Mapping Source VMs to Existing Replicas

If you want to perform staged replication or if you lose the replication job (due to accidental job deletion or because you need to recreate jobs in a new copy of the product), you can map source VMs to existing replicas to avoid running full VM replication again.

To map source VMs to existing VMs, follow the steps below:

1. Click **Advanced setup**.
2. Open the target VM drop-down list and select the **Use existing target VM** option.
3. Select the VM that should be used as a target for replication in the **Virtual Machine** drop-down list.



When you run the job, the product analyzes the target VM you have selected, determines how it is different from the source VM, and transfers only the differential data.

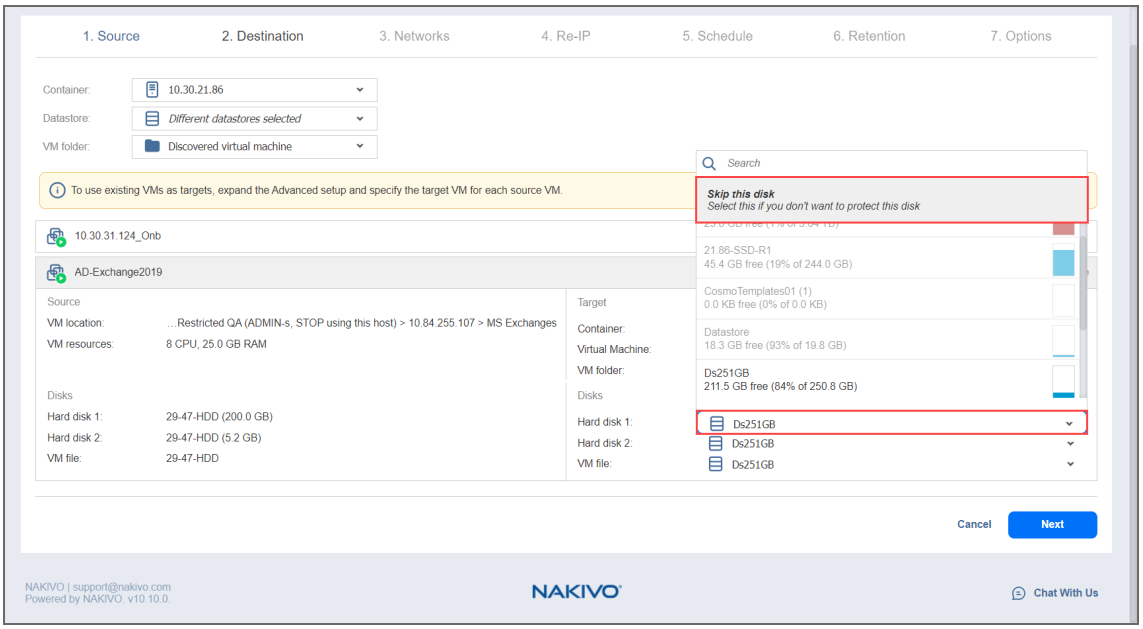
VM replication mapping can be a time-consuming process that can be equal to the time required to create a full VM replication.

After the job completion, a new recovery point will be created and existing recovery points will not be changed or overwritten.

## Excluding a VM Disk from the Replication Job

If you do not want to replicate some disks of a VM, you can exclude those disks from the replication job. Follow the steps below:

1. Click **Advanced setup**.
2. Open the target VM drop-down list.
3. Click the drop-down list next to the disk that you want to skip and select the **Skip this disk** option.
4. After you have specified replica location options, click **Next** to go to the next page.

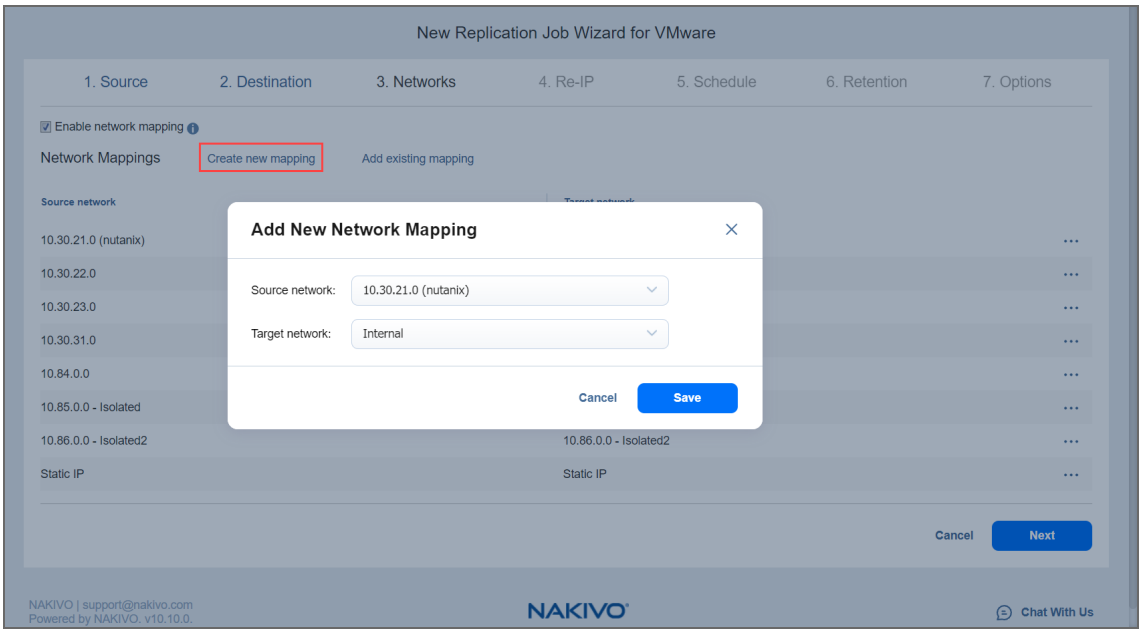


# Replication Job Wizard for VMware: Networks

> **Note**
> The **Networks** page is skipped if you have chosen an existing backup as the source for your replication job on the **Source** wizard page.

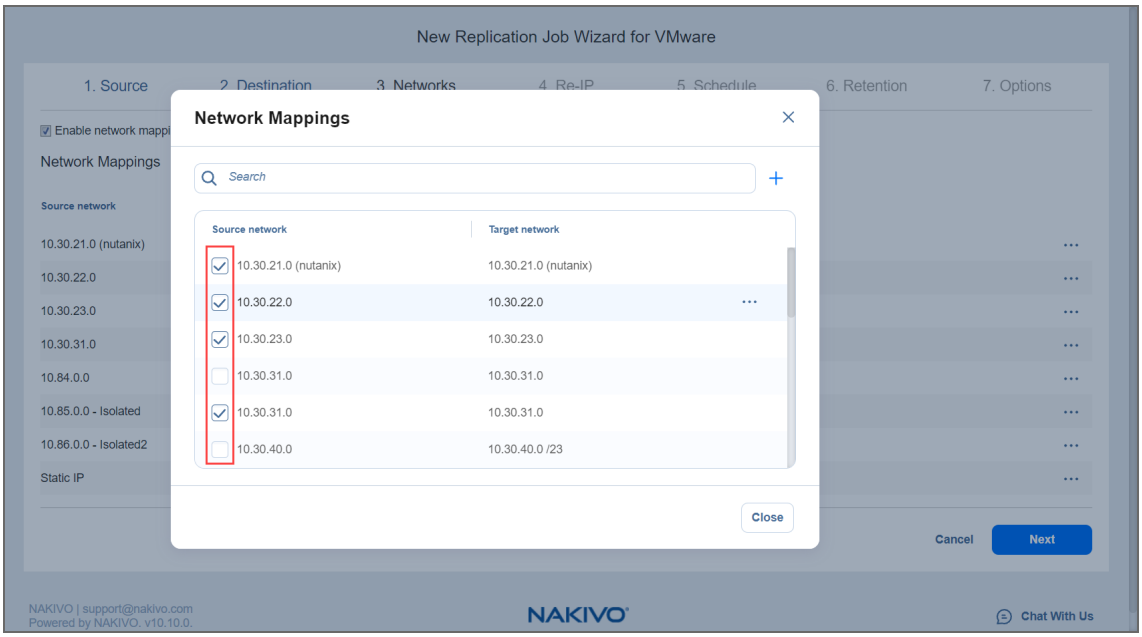
To map source VM virtual networks to appropriate target virtual networks, do the following on the **Networks** page in the wizard:

1. Select **Enable network mapping**.
2. The **Network Mappings** section opens. You have the following options:
   - Create a new mapping:
     a. Click **Create new mapping**.
     b. The **New Network Mapping** dialog box opens. Choose a source network and a target network and click **Save**.

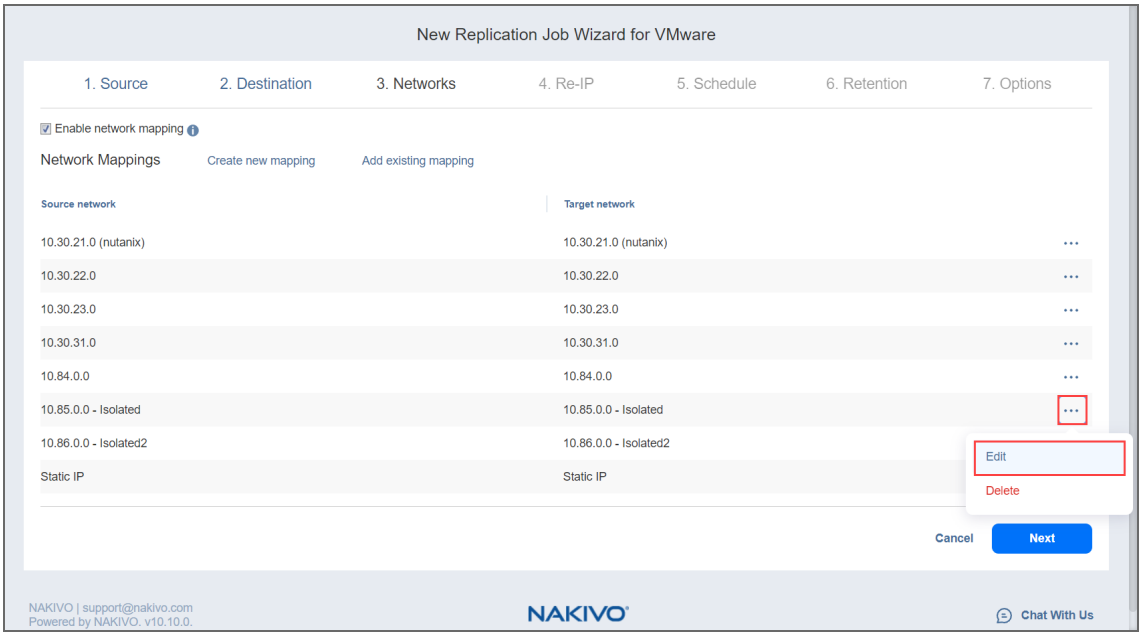



   - Add an existing mapping:
     a. Click **Add existing mapping**.
     b. The **Network Mappings** dialog box opens. Choose one or more appropriate network

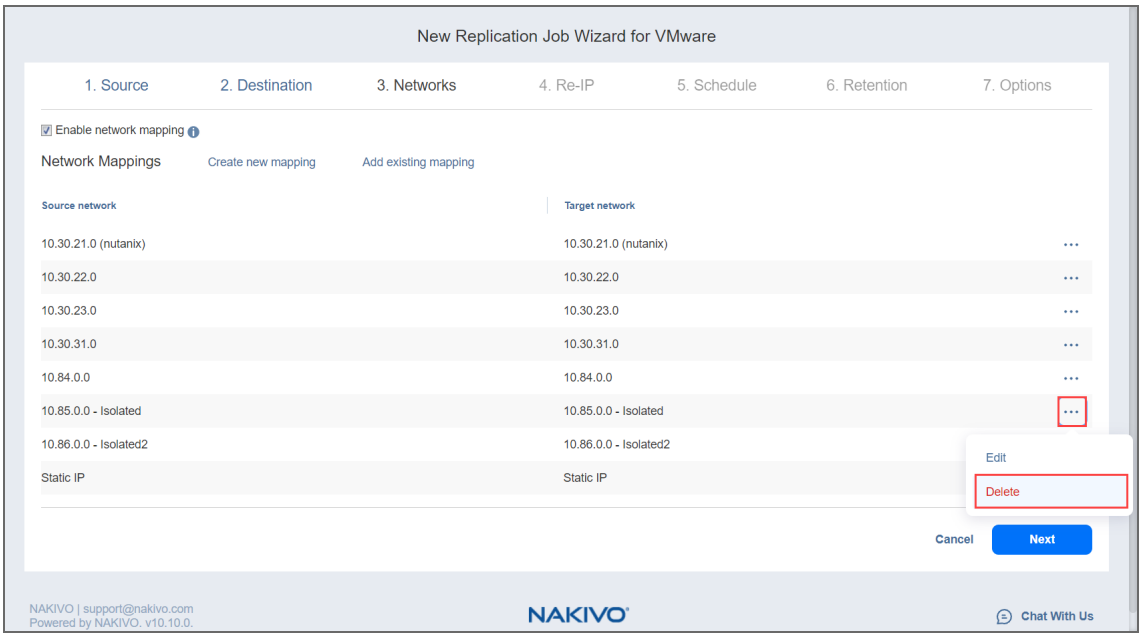
mappings and close the dialog box.




- Edit an existing mapping:
  1. Hover the pointer over the required item in the **Network Mappings** list and click the ellipsis **Manage** button to the right of the item. Then, click **Edit**.
  2. The **Edit Network Mapping** dialog box opens. Choose the required item from the **Target network** drop-down list and click **Save**.




- Delete an existing mapping: Hover the pointer over the required item in the **Network Mappings** list and click the ellipsis **Manage** button to the right of the item. Then, click **Delete**.




3. Click **Next** to go to the next page of the wizard.

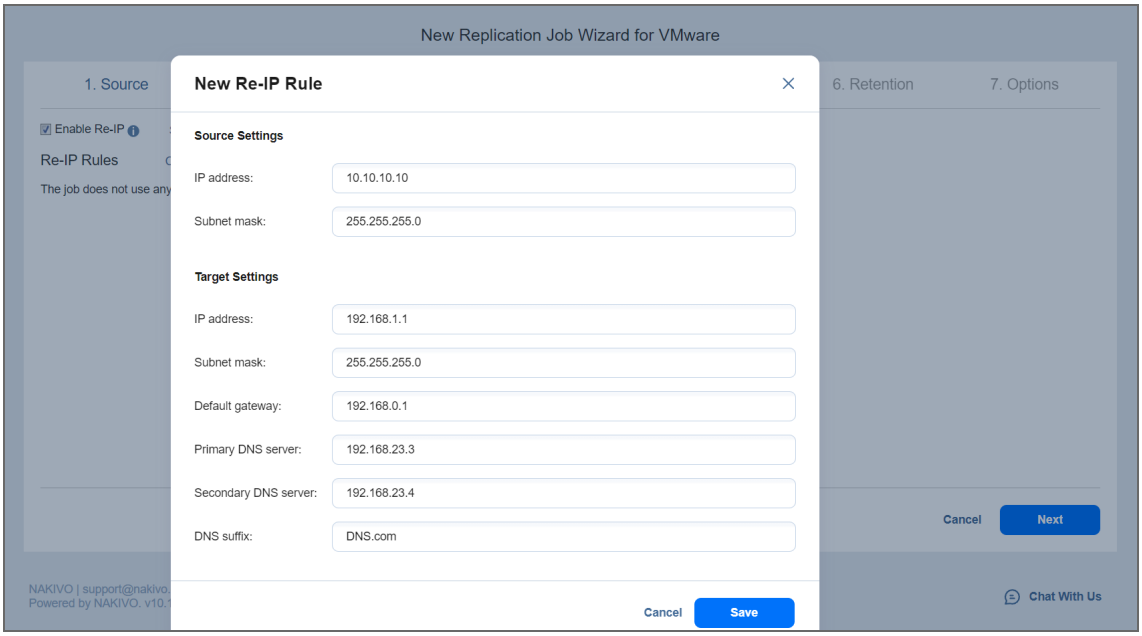
# Replication Job Wizard for VMware: Re-IP

**Note**

The **Re-IP** page is skipped if you have chosen an existing backup as the target for your replication job on the **Source** wizard page.

To enable Re-IP rules for your replication job, do the following on the **Re-IP** page of the wizard:

1. Select **Enable Re-IP**.
2. Click the **Select VMs** link.
3. The **Re-IP** dialog box opens. In the list of your source VMs, select at least one, and close the dialog box.
4. You have the following options:
   - Create a new rule:
     a. Click **Create new rule**.
     b. The **New Re-IP Rule** dialog box opens. Enter source and target settings for the Re-IP rule and click **Save**.



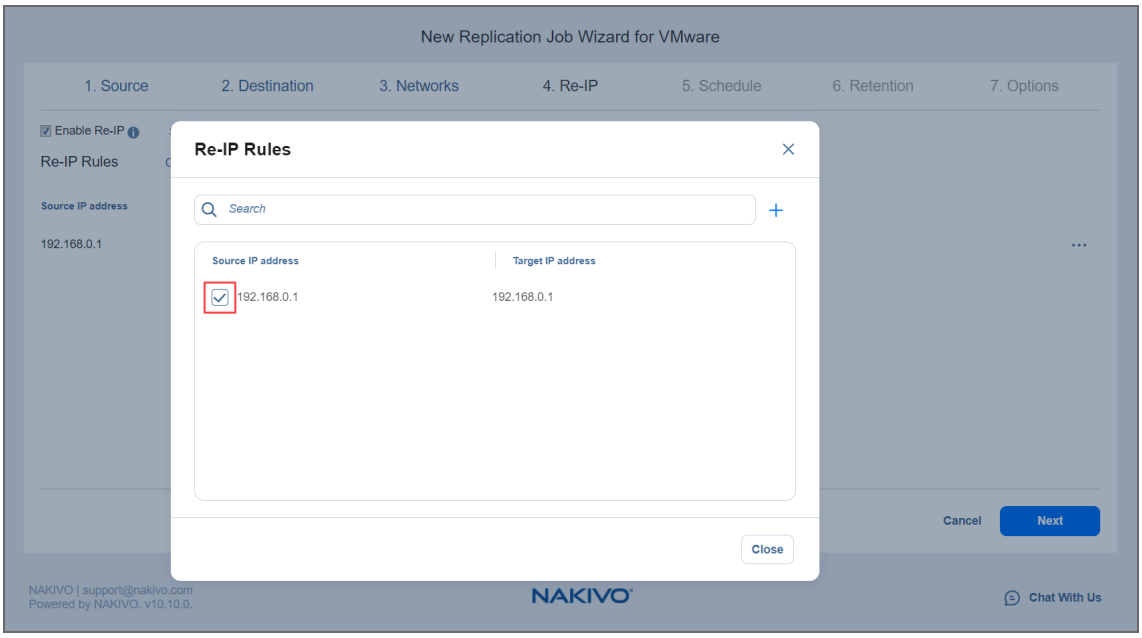

**Note**

You can use wildcards for IP addresses. Refer to the “Failover Job Wizard for VMware: Re-IP” on page 1113 topic for details.

   - Add an existing rule:

a. Click **Add existing rule**.

b. The **Re-IP Rules** dialog box opens. Select one or more appropriate Re-IP rules and close the dialog box.




- Edit an existing Re-IP rule:

a. Hover the pointer over the required item in the **Re-IP Rules** list and click the ellipsis **Manage** button to the right of the item. Then, click **Edit**.

b. The **Edit Re-IP Rule** dialog box opens. Edit the required properties of the Re-IP rule and click **Save**.




- Delete an existing mapping: Hover the pointer over the required item in the **Re-IP Rules** list and click the **Delete** icon to the right of the item.




- To leave the list of existing Re-IP rules intact, go to the next page of the wizard.

5. Click **Next** to go to the next page of the wizard.

**Notes**

- Re-IP Rules enabled for replication jobs are only stubs for failover jobs. They do not work at the replication stage.
- Re-IP rules that are enabled for your replication job can be used for creating the corresponding failover jobs. Refer to "Failover Job Wizard for VMware: Re-IP" on page 1113 for details.

# Replication Job Wizard for VMware: Schedule

On the **Schedule** page of the wizard, select to run the replication job manually or schedule the job to run on a regular basis.


- "Disabling Scheduled Job Execution" below
- "Daily or Weekly Replication" below
- "Monthly or Yearly Replication" on the next page
- "Periodic Replication" on page 952
- "Chained Job" on page 953
- "Add Another Schedule" on page 954


## Disabling Scheduled Job Execution

If you wish to start the job manually (without any schedule), select the **Do not schedule, run on demand** checkbox.

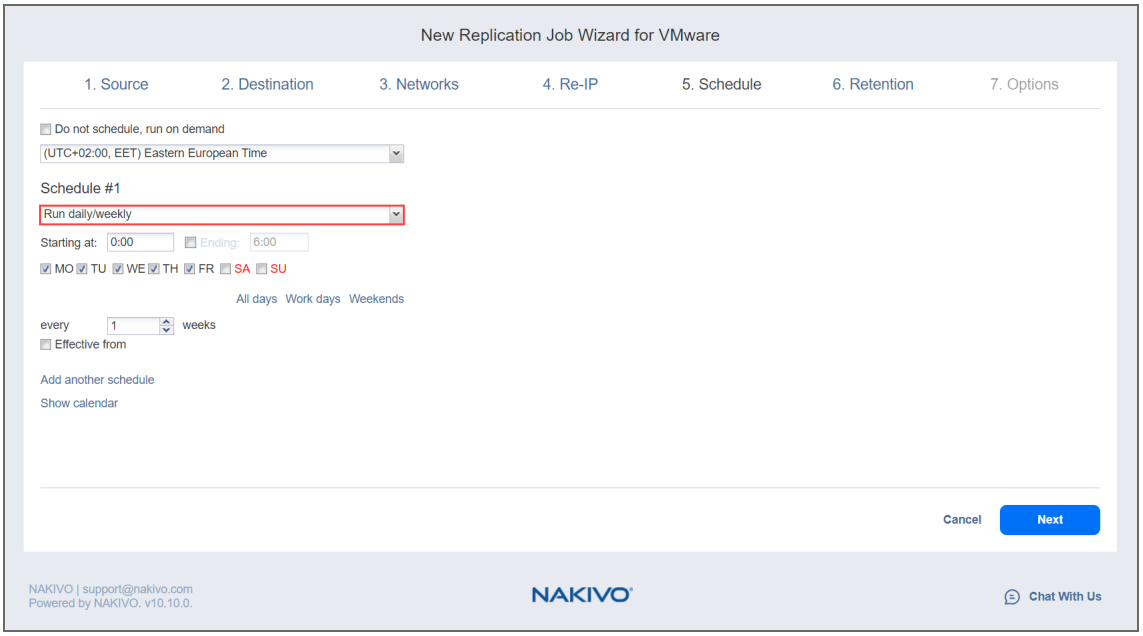



## Daily or Weekly Replication

To run the job once a day, choose **Run daily/weekly** from the schedule drop-down list and do the following:

- Choose a time zone that should be used for the job start and end times from the time zone drop-down list.
- Specify the time when the job should be started in the **Starting at** box.
- Specify the end time for the job in the **Ending** box. If the job has not completed by the time specified, the job will be stopped.
- Select the days of the week during which the job will be started.

- If necessary, select the **Effective from** checkbox and pick the date when the schedule comes into effect.

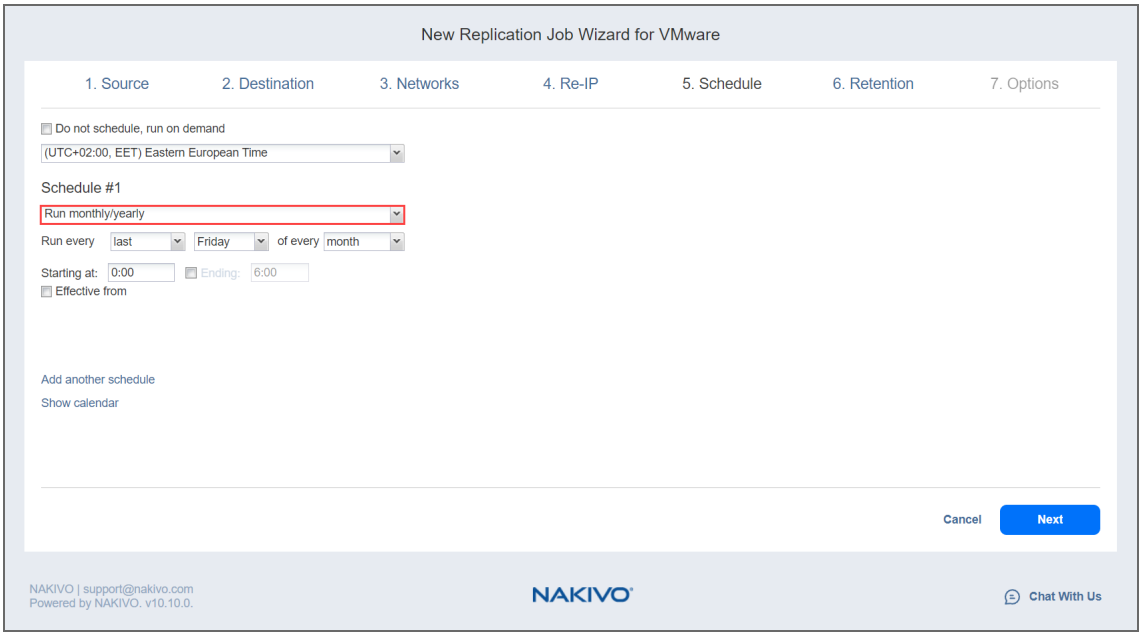



## Monthly or Yearly Replication

To run the job monthly or yearly, choose **Run monthly/yearly** from the schedule drop-down list and do the following:

- Specify the job start schedule in the appropriate boxes.
- Specify the day and month when the job should be started in the **Run every** boxes.
- Specify the time when the job should be started in the **Starting at** box.
- Specify the end time for the job in the **Ending** box. If the job has not completed by the time specified, the job will be stopped.

- If necessary, select the **Effective from** checkbox and pick the date when the schedule comes into effect.




## Periodic Replication

To run the job multiple times per day, choose **Run periodically** from the schedule drop-down list and then choose a time period from the appropriate boxes:

- Specify the time when the job should be started in the **Starting at** box.
- Specify the end time for the job in the **Ending** box. If the job has not completed by the time specified, the job will be stopped.
- Select the days of the week during which the job will be started.

- If necessary, select the **Effective from** checkbox and pick the date when the schedule comes into effect.

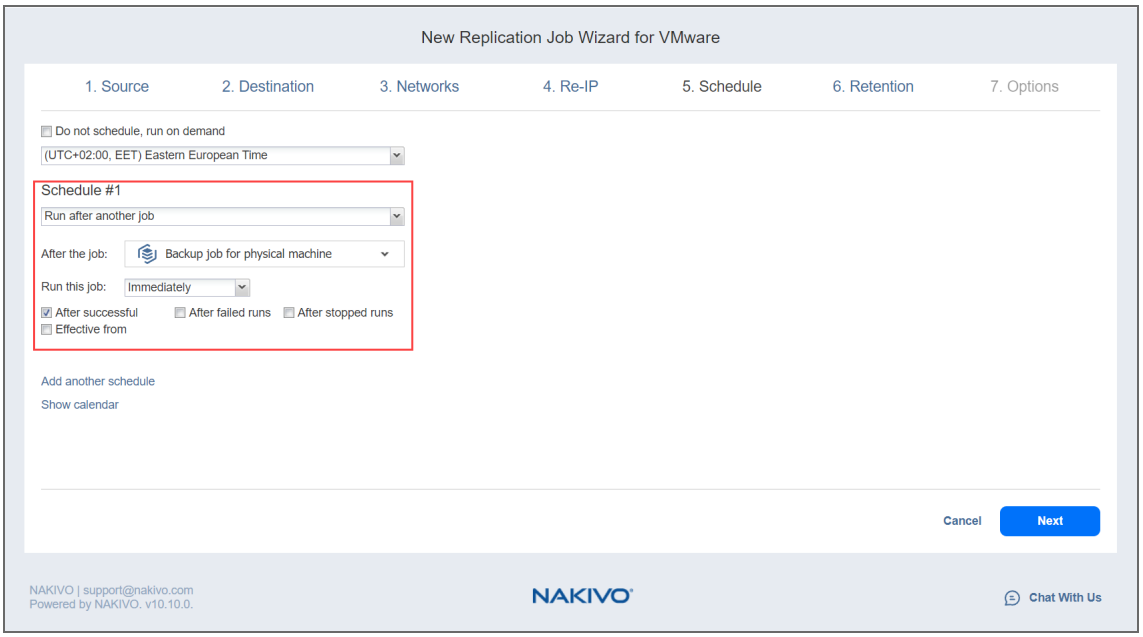



## Chained Job

To run the job after a previous one has completed, choose **Run after another job** from the schedule drop-down list and set the options as follows:

- **After the job**: Select a job after which the current job will be started.
- **Run this job**: Choose whether to run the current job immediately after the previous one has completed or within a delay.
- **After successful runs**: If selected, the job will run if the previous one has completed successfully.
- **After failed runs**: If selected, the job will run if the previous one has failed.
- **After stopped runs**: If selected, the job will run if the previous one has been stopped.

- **Effective from**: If selected, the schedule will come into effect on the date picked.




# Add Another Schedule

If you want to have more than one schedule for your job, click **Add another schedule** and set it up as has been described above.



# Replication Job Wizard for VMware: Retention

**Important**

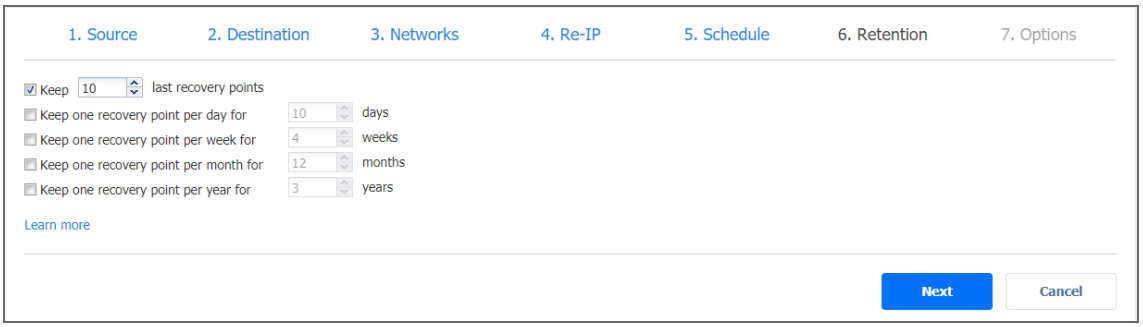
This page is not displayed if the new scheduler is enabled.

NAKIVO Backup & Replication can create a recovery point (snapshot) on the replica VM after each job run. You can specify the number of recovery points to be retained using the Grandfather-Father-Son (GFS) backup rotation scheme.

Up to 30 recovery points in total can be created on a replica VM.

Use the following options to specify a retention policy:

- **Keep x last recovery points**: Retains the specified number of last recovery points for each VM in the job.
- **Keep one recovery point per day for x days:** Retains one last recovery point per day for the specified number of days.
- **Keep one recovery point per week for x weeks**: Retains the last available backup of every week for the specified number of weeks.
- **Keep one recovery point per month for x months**: Retains the last available backup of every month for the specified number of months.
- **Keep one recovery point per year for x years**: Retains the last available backup of every year for the specified number of years.



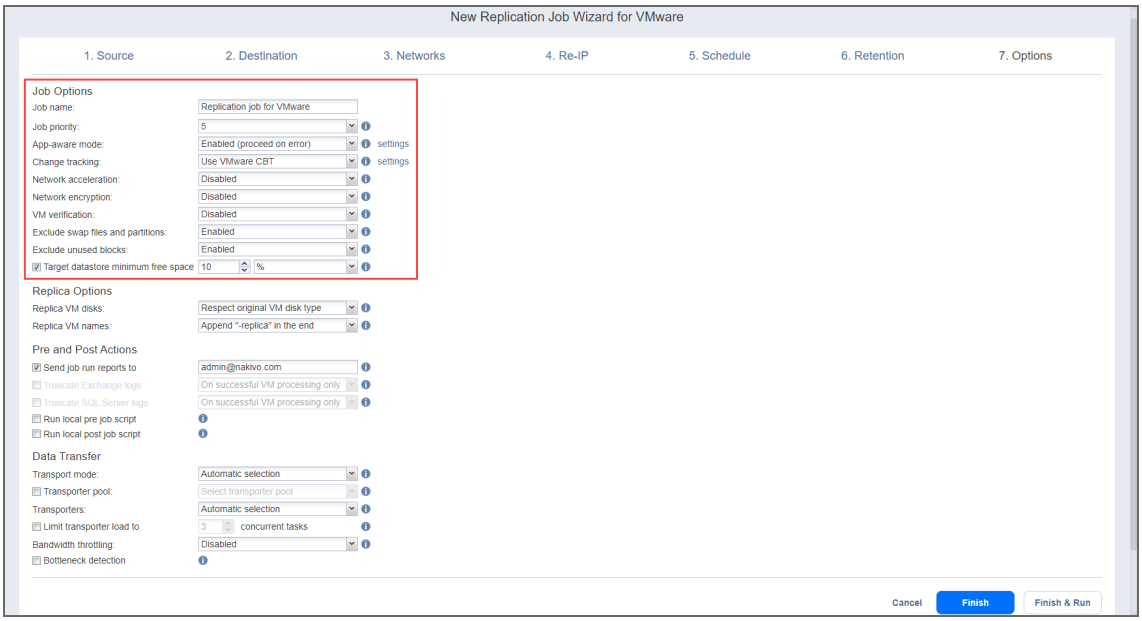
# Replication Job Wizard for VMware: Options

On the **Options** page, set up replication job options as described in the sections below:


- "Job Options" below
  - "Job Name" on the next page
  - "Job Priority" on the next page
  - "App-aware Mode" on the next page
  - "Change Tracking" on page 958
  - "Network Acceleration" on page 958
  - "Network Encryption" on page 959
  - "VM Verification" on page 959
  - "Exclude Swap Files and Partitions" on page 961
  - "Exclude Unused Blocks" on page 961
  - "Target Datastore Free Space" on page 961
- "Replica Options" on page 961
  - "Replica VM Disks" on page 961
  - "Replica VM Names" on page 962
- "Pre and Post Actions" on page 962
  - "Email Notifications" on page 962
  - "Microsoft Exchange Server Logs Truncation" on page 963
  - "Microsoft SQL Server Logs Truncation" on page 963
  - "Pre Job Script" on page 963
  - "Post Job Script" on page 964
- "Data Transfer" on page 965
  - "Transport Mode" on page 965
  - "Transporter Pool" on page 966
  - "Transporters" on page 966
  - "Transporter Load" on page 966
  - "Bandwidth Throttling" on page 966
- "Completing the New Replication Job Wizard for VMware" on page 967


## Job Options

In this section, you can specify a name and a priority level for your replication job, and enable/disable app-aware mode, change tracking, network acceleration, encryption, VM Verification, and excluding swap files, partitions and unused blocks.




## Job Name

Specify a name for the replication job.

## Job Priority

Select a job priority level between *1* and *5*, with *1* being the highest priority. Jobs with higher priority levels are prioritized by Transporters during job processing.

> **Note**
> This option is only available in the Enterprise, Enterprise Essentials, Enterprise Plus, MSP Enterprise, and MSP Enterprise Plus editions.

## App-aware Mode

With the **App-aware mode** selected, VM replication is performed using VMware Guest OS quiescing (which in turn relies on Microsoft VSS) to ensure the consistency of application data. Select one of the options from the **Change tracking** drop-down list:

- **Enabled (proceed on error)**: With this option selected, NAKIVO Backup & Replication proceeds even if an application quiescing error is encountered.
- **Enabled (fail on error):** With this option selected, NAKIVO Backup & Replication automatically fails the job if an application quiescing error is encountered.
- **Disabled:** Selecting this option disables the app-aware mode.

**Note**
This option is not available for existing backups chosen as the target of replication on the **Source** wizard page.

## Change Tracking

Select one of the options from the **Change tracking** drop-down list:

- **Utilize VMware CBT:** With this option selected, NAKIVO Backup & Replication enables the Change tracking feature for source VMs. This feature quickly identifies which data blocks have changed since the last job run, significantly increasing job speed. Click **settings** to specify the following options:
  - **On error:**
    - **switch to proprietary method immediately/on the next retry/on the last retry:** If VMware CBT fails to provide data on changed blocks for a VM with this option selected, NAKIVO Backup & Replication performs an incremental backup of the VM using the proprietary change tracking technology.
    - **reset CBT on the next retry/on the last retry:** If VMware CBT fails to provide data on changed blocks for a VM with this option selected, NAKIVO Backup & Replication resets VMware CBT for the VM.
    - **fail VM processing immediately/on the next retry/on the last retry:** If VMware CBT fails to provide data on changed blocks for a VM with this option selected, NAKIVO Backup & Replication does not process the VM and states job failure (other VMs in the job are processed).
    - **Double-check changed blocks provided by CBT**: With this option selected, NAKIVO Backup & Replication runs a check on data blocks provided by VMware CBT to ensure that VMware CBT does not overstate the amount of changed data.
- **Use proprietary method:** With this option selected, NAKIVO Backup & Replication performs incremental backups using the proprietary change tracking technology. This feature requires reading the contents of all VM disks to determine the data blocks that have changed since the last job run.
- **No change tracking (always full):** With this option selected, NAKIVO Backup & Replication always performs a full VM backup of all source VMs.

**Note**
This option is not available for existing backups chosen as the target of replication on the **Source** wizard page.

## Network Acceleration

With Network acceleration enabled, NAKIVO Backup & Replication uses compression and traffic reduction techniques to speed up data transfer. Select this option if you plan to replicate over WAN or slow LAN links.

## Network Encryption

If Network Encryption is enabled VM data is protected with AES 256 encryption while traveling over the network. Data encryption increases the replication time and CPU load on machines running Transporters. Select this option when replicating over WAN without a VPN connection.

> **Note**
> You need at least one Transporter at source and target sites to enable encryption.

## VM Verification

VM Verification allows you to check the integrity of the backup by starting it and interacting with it. For more details, refer to the "VM Verification" on page 66 article.

You can choose one of the following **VM Verification** options:

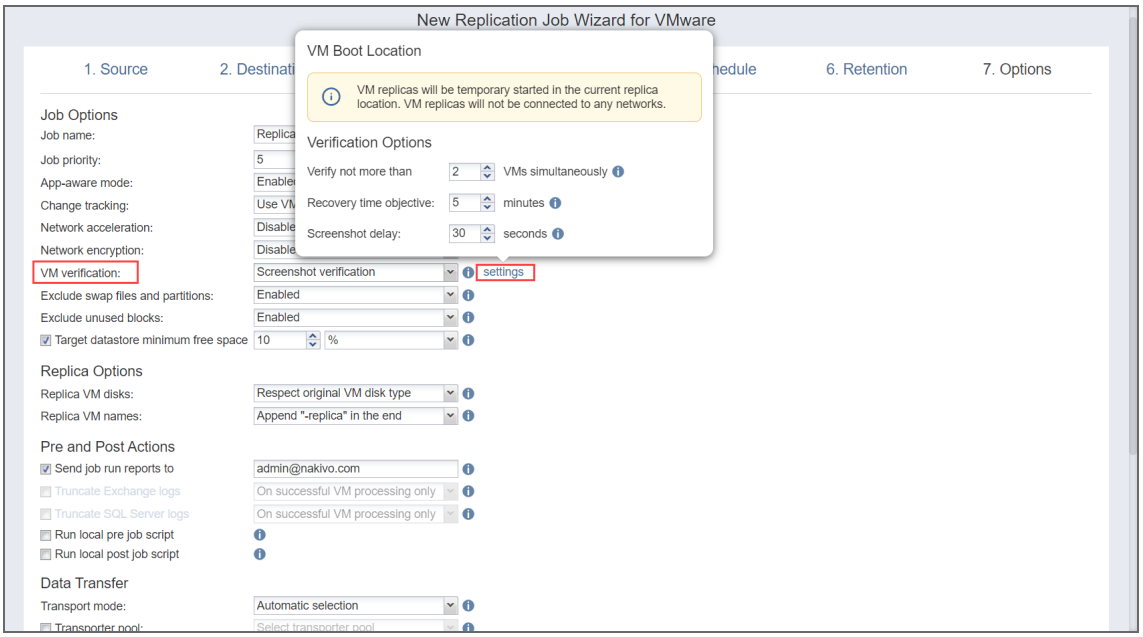
- **Disabled**: VM Verification is disabled.
- **Screenshot verification**: When enabled, the VM replica created by the job is verified: NAKIVO Backup & Replication powers on this replica with networking turned off, takes a screenshot of the OS, then powers off the VM replica. The VM screenshot will be included in email notifications (if they're configured. See "Notifications & Reports" on page 444.) and displayed on the **Dashboard**.
- **Boot verification**: When enabled, the VM replica created by the job is verified: After VM replication is completed, NAKIVO Backup & Replication recovers the VM using Flash boot, disables networking to prevent network connections, and verifies that system start is successful.

> **Important**
> VM verification requires VMware Tools to be installed on all VMs.

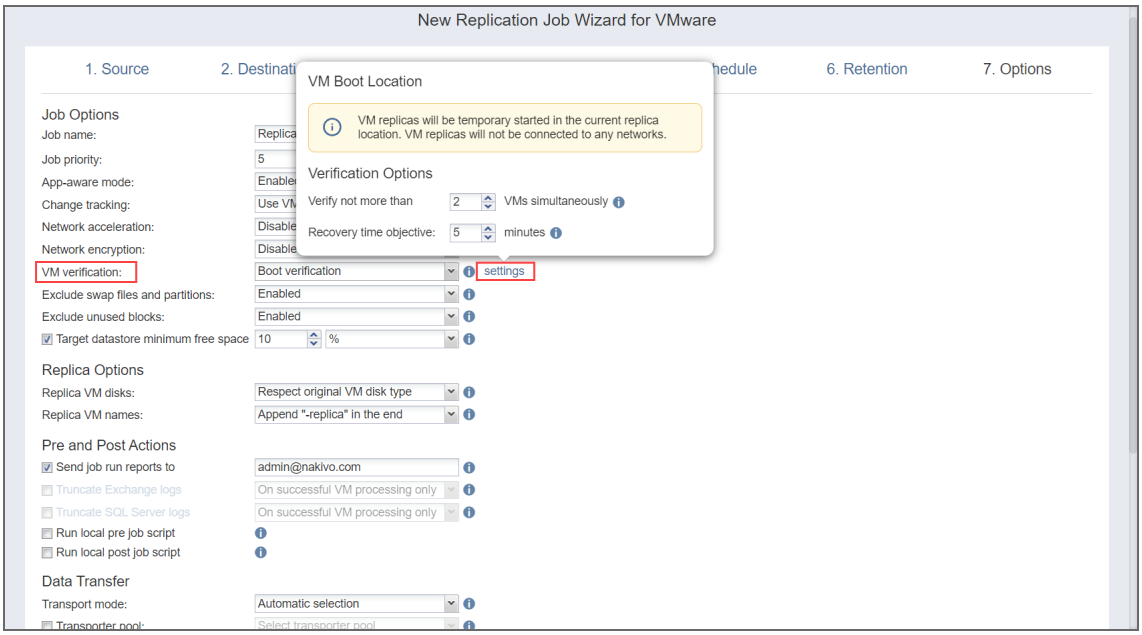
After selecting the **Screenshot verification** option, do the following in the dialog box that opens:

- **Verify not more than x VMs simultaneously**: Specify the maximum number of VMs that can be started on the target container simultaneously.
- **Recovery time objective x minutes**: Specify the amount of time allocated for verification of each VM backup. If a VM OS does not start within the specified amount of time, verification will be considered failed.
- **Screenshot delay x seconds**: The amount of time that the product should wait after the Guest OS starts before taking a screenshot.

> **Note**
> The specified time must be sufficient to fully start the VM OS. Try increasing this amount if the default amount is not sufficient.




1. Provide a location of the VMs to be booted as described for the **Screenshot verification** option.
2. Set verification options:
    - **Verify not more than x VMs simultaneously**: Specify the maximum number of VMs that can be started on the Target Container simultaneously.
    - **Recovery time objective**: Specify the amount of time allocated for verification of each VM backup. If a VM OS does not start within the specified amount of time, verification will be considered failed.



## Exclude Swap Files and Partitions

When this option is enabled, NAKIVO Backup & Replication automatically excludes swap files and partitions during the backup process.

> **Note**
> This feature is not available for replication from backup jobs.

## Exclude Unused Blocks

When this option is enabled, NAKIVO Backup & Replication automatically excludes unused disk blocks and blocks occupied by deleted files during processing of source objects running Windows OS. This feature allows for reducing backup storage space and object processing time.

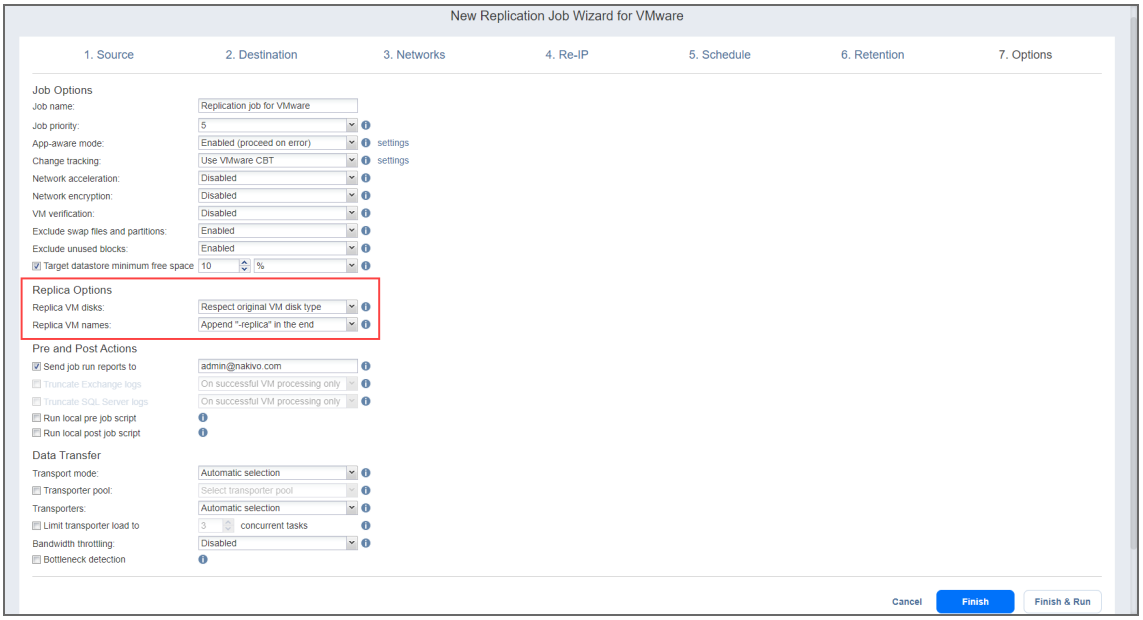
> **Note**
> This feature is not available for replication from backup jobs.

## Target Datastore Free Space

With this option enabled, specify the minimum free space for the target datastore during incremental job run. In case the free space on the target datastore is below the set value, the job automatically fails.

# Replica Options

In this section, you can choose a VM disk type and VM name. Proceed as described below.




## Replica VM Disks

Choose one of the following options:

- **Respect original VM disk type**: With this option enabled, the created disk will be of the same type as that of the source VM.
- **Create only thin disks on target VMs**: With this option enabled, only thin disks are created on replicas, regardless of the disk types of the original source VM.

## Replica VM Names

NAKIVO Backup & Replication allows you to change VM Replica names to easily distinguish between the VM Replicas and the source VMs. By default, the text "-replica" is appended to the end of the VM Replica name.

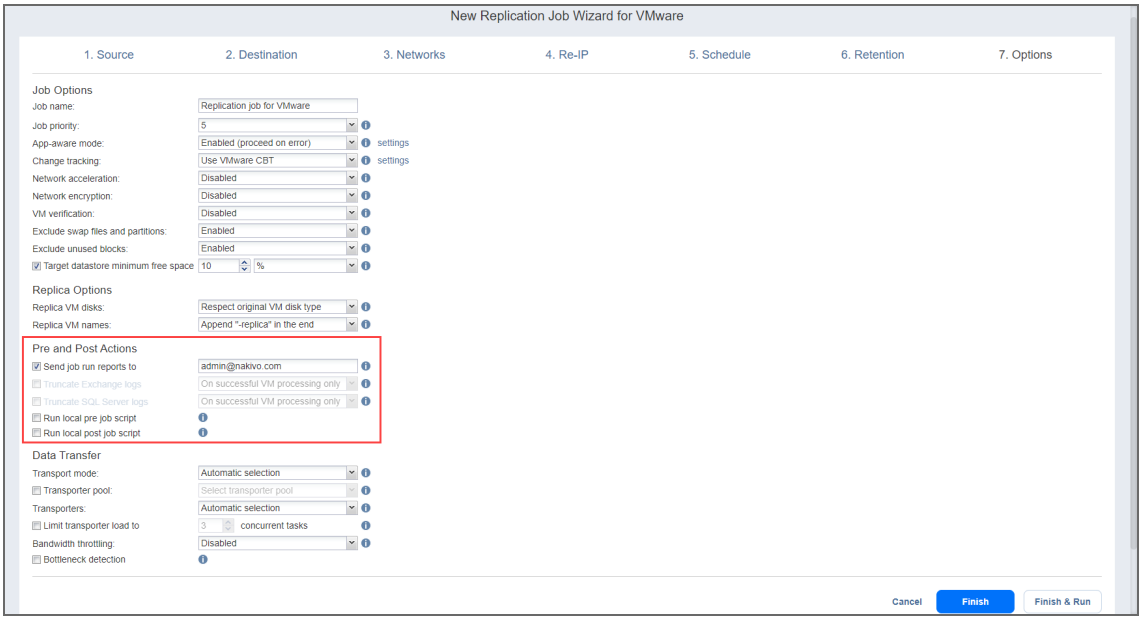
To change VM Replica names:

In the *Replica Options* section, choose one of the following Replica VM names options:

- **Append "-replica" in the end**: Source VM names are used for replica names and "-replica" are added to the replica name.
- **Leave replica names as is**: Replica names will be identical to the source VM names.
- **Enter custom replica names**: Enter custom names for replicas.

# Pre and Post Actions

In the *Pre and Post Actions* section, you can set up email notifications, Exchange and SQL Server logs truncation, pre and post job scripts.




## Email Notifications

NAKIVO Backup & Replication can send email notifications about the job completion status to specified recipients. This feature complements global notifications and allows you to configure notifications on a per-job level.

**Note**

To enable this option, make sure your Email settings are configured.

To send email notifications, do the following:

In the *Pre and Post Actions* section:

1. Select **Send job run reports to**.
2. Specify one or more email addresses in the text field. Use semicolons to separate multiple email addresses.

### Microsoft Exchange Server Logs Truncation

NAKIVO Backup & Replication allows you to delete (aka truncate) Microsoft Exchange Server logs on the source VMs after job completion.

To set up Microsoft Exchange log truncation, do the following:

1. In the **Pre and Post Actions** section, select the **Truncate Exchange logs** option.
2. In the **Exchange Log Truncation** dialog box that opens, select the checkboxes next to the VMs running Microsoft Exchange and then select the credentials next to each VM. These credentials are used to log into the VMs that you have selected.

**Note**

This option is not available for existing backups chosen as the target of replication on the **Source** wizard page.

### Microsoft SQL Server Logs Truncation

NAKIVO Backup & Replication allows you to delete (aka truncate) Microsoft SQL Server logs on the source VMs after job completion.

To set up Microsoft SQL log truncation, do the following:

1. In the **Pre and Post Actions** section, select the **Truncate SQL Server logs** option.
2. In the **SQL Server Log Truncation** dialog box that opens, select the checkboxes next to the VMs running Microsoft SQL Server and then select the credentials next to each VM. These credentials are used to log into the VMs that you have selected.

**Note**

This option is not available for existing backups chosen as the target of replication on the **Source** wizard page.

## Pre Job Script

To run a script before the product begins replicating VMs:

1. Place a script file on the machine where the Director is installed.
2. In the *Pre and Post Actions* section, select the **Run local pre job script** option.
    - **Script path**: Specify a local path to the script on the machine where the **Director** is installed. Script interpreter should be specified.
    - **Example (Windows)**: *cmd.exe /c D:\script.bat*
      **Example (Linux)**: *bash /root/script.sh*
    - **Job behavior**: Choose one of the following job behaviors in relation to script completion:
        - **Do not wait for the script to finish**: With this option selected, the product runs the script and starts replicating VMs at the same time.
        - **Wait for the script to finish**: With this option selected, VM replication is started only after the script is completed.
    - **Error handling**: Choose one of the following job behaviors in relation to script failure:
        - **Fail the job on script failure**: With this option selected, the job is failed and VM replication is not performed if the script has failed.
        - **Continue the job on script failure**: With this option selected, the job performs VM replication even if the script has failed.
3. Specify the following parameters in the dialog box that opens:

## Post Job Script

To run a script after the product has finished backing up all VMs:

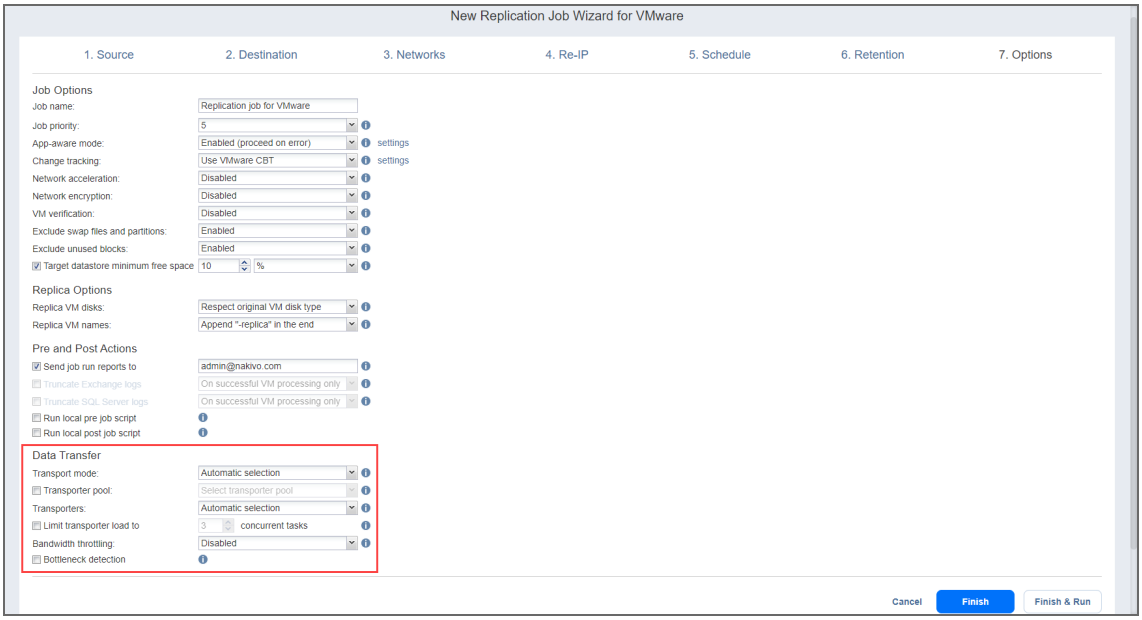
1. Place a script file on the machine where the **Director** is installed.
2. In the *Pre and Post Actions* section, select the **Run local post job script** option.
3. Specify the following parameters in the dialog box that opens:
    - **Script path**: Specify a local path to the script on the machine where the **Director** is installed. Script interpreter should be specified.

      **Example (Windows)**: *cmd.exe /c D:\script.bat*
      **Example (Linux)**: *bash /root/script.sh*
    - **Job behavior**: Choose one of the following job behaviors in relation to script completion:
        - **Wait for the script to finish**: With this option selected, the job is in the “running” state until the script is completed.
        - **Do not wait for the script to finish**: With this option selected, the job is completed even if the script execution is still in progress.
    - **Error handling**: Choose one of the following job behaviors in relation to script failure:

- **Continue the job on script failure**: With this option selected, script failure does not influence the status of the job.
- **Fail the job on script failure**: With this option selected and the script has failed, the job status will be set to “failed” even if VM replication has been successful.

# Data Transfer

In the *Data Transfer* section, you can choose a transport mode and a **Transporter** to be used for reading data from source VMs, and configure bandwidth throttling. Proceed as described below.




## Transport Mode

To select a transport mode, in the *Data Transfer* section, choose a transport mode for retrieving VM data:

- **Automatic selection**: With this option selected, NAKIVO Backup & Replication automatically selects the best transport mode available:
  - If the source **Transporter** is installed on a VM, NAKIVO Backup & Replication tries to use transport modes in the following order: Hot Add > SAN > LAN.
  - If the source **Transporter** is installed on a physical machine, NAKIVO Backup & Replication tries to use transport modes in the following order: SAN > Hot Add > LAN.
- **SAN only**: With this option selected, NAKIVO Backup & Replication only uses direct SAN access to retrieve VM data. If direct SAN access to VM data is not available, the job will fail.
- **Hot Add only**: With this option selected, NAKIVO Backup & Replication only uses Hot Add to retrieve VM data. If direct Hot Add is not available, the job will fail.
- **LAN only**: With this option selected, NAKIVO Backup & Replication only retrieves VM data via LAN.

**Note**
This option is not available for existing backups chosen as the target of replication on the **Source** wizard page.

## Transporter Pool

If this option is enabled, only the transporters that belong to the selected transporter pool shall be used during the job run.

## Transporters

By default, the product automatically determines the **Transporter** to be used to read data from the source VM. However, you can manually specify the **Transporters** to be used for the job by choosing one of the following **Transporters** options in the **Data Transfer** section:

- **Automatic selection**: The product automatically determines the **Transporters** that are the closest to source and target hosts.
- **Manual - configured for all VMs**: Select this option to manually specify a single source and a single target **Transporter** to be used for data transfer by the job.
- **Manual - configured per host**: Select this option to manually specify **Transporters** for all source and target hosts.

## Transporter Load

You can limit the maximum number of transporter tasks used by the job. By default, it is set to 3 concurrent tasks.

To change the default number of tasks, do the following:

1. In the **Data Transfer** section, select the checkbox next to **Limit transporter load to**.
2. Specify the number of concurrent tasks in the corresponding field.

## Bandwidth Throttling

To regulate the speed of data transfer over the network for your replication job:

1. For the **Bandwidth throttling** option, choose **Enabled**.

   **Note**
   If bandwidth throttling is disabled for the current job, global bandwidth rules may still apply to your job. Refer to "Bandwidth Throttling" on page 407 for details.

2. Click the **settings** link that becomes available.
3. The **Job Bandwidth Rules** dialog box opens displaying the list of available rules. You have the following options:

- Create a new bandwidth rule for your replication job:
  a. Click the **Create New Rule** button.
  b. The **New Bandwidth Rule** dialog box opens. Refer to the "Bandwidth Throttling" on page 407 topic for details on creating a bandwidth rule.
  c. Click **Save**.
- Activate an existing bandwidth rule for your job: Select the checkbox to the left of the required bandwidth rule. To deactivate a bandwidth rule for your job, deselect the corresponding checkbox.
- Edit a bandwidth rule: Click the **Edit** link for a bandwidth rule and modify it in the **Edit Bandwidth Rule** dialog box that opens.
- Disable a bandwidth rule: Click the **Disable** link. The bandwidth rule will be disabled for all jobs.
- Remove a bandwidth rule: Click the **Remove** link and then click **Delete** to confirm your operation.

### Bottleneck detection

When the **Bottleneck detection** option is enabled, additional information is collected and recorded in NAKIVO Backup & Replication logs in the course of data transfer for the purpose of bottleneck detection. Check this option to enable the **Bottleneck detection** capability of the **Transporters** engaged in the job.

## Completing the New Replication Job Wizard for VMware

Click **Finish** or **Finish & Run** to complete job creation.

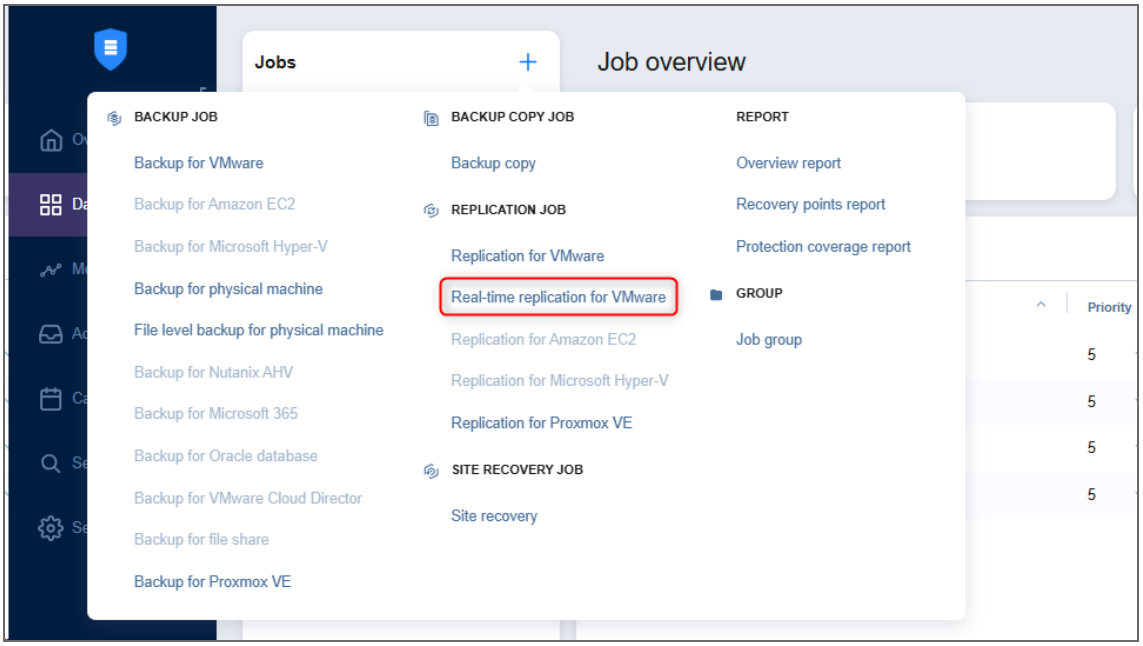
> **Note**
> If you click **Finish & Run**, you will have to define the scope of your job. Please refer to "Running Jobs on Demand" on page 364 for details.

# Creating VMware Real-Time Replication Jobs

With NAKIVO Backup & Replication, you can create real-time replication (RTR) jobs for VMware. Unlike regular replication jobs, real-rime replication jobs for VMware continuously copy data from the source host as it is being generated.

Before creating RTR jobs for VMware, ensure you have installed the **I/O Filter** on the source host and the **Transporter** (VMware vSphere appliance) is deployed on the target host/cluster. For more information on the prerequisites to creating RTR jobs for VMware, see the **Real-Time Replication for VMware** section in "Feature Requirements" on page 169.

To create a real-time replication job for VMware, click the plus **Create** button in the **Data Protection** menu and then select **Real-time replication for VMware**.




**Notes**

- VMware VM involved in a RTR job cannot be concurrently processed by other jobs.
- If no VMware clusters have the **I/O Filter** installed, the **Real-time replication services are required** dialog box opens warning that you need to install the **I/O Filter** first. Click **Open** to access the Setup Real-Time Replication Service Wizard.

The **New Real-Tme Replication Job Wizard for VMware** opens. Complete the wizard to create a real-time replication job for VMware VMs.


- "Real-Time Replication Job Wizard for VMware: Source" on the next page
- "Real-Time Replication Job Wizard for VMware: Destination" on page 972
- "Real-Time Replication Job Wizard for VMware: Networks" on page 976
- "Real-Time Replication Job Wizard for VMware: Re-IP" on page 978

- "Real-Time Replication Job Wizard for VMware: Retention" on page 981
- "Real-Time Replication Job Wizard for VMware: Options" on page 985

# Real-Time Replication Job Wizard for VMware: Source

On the **Source** page of the wizard, select one of the views to add VMware items to your Real-time replication job for VMware.

**Notes**

- It is not possible to select VMs with no defined VM storage policy.
- To configure the VM Storage Policy and the I/O Filter, click the **Setup Real-Time Replication Services** link and navigate to Setup Real-Time Replication Services Wizard.
- It is not possible to select VMs from a cluster with an unsupported ESXi version.
- Templates are not supported.

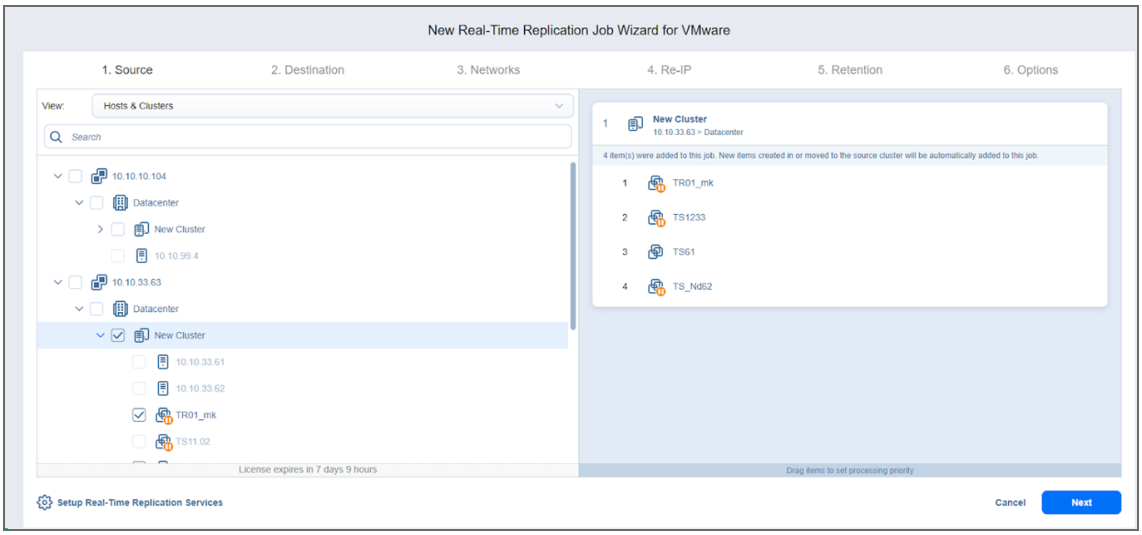
- Host and Clusters
- VMs and Templates
- Policy

## Host and Clusters

When the **Host & Clusters** view is chosen, the inventory tree opens in the left pane and displays all VMware items: clusters, hosts, folders, resource pools, and VMs. Proceed as follows:

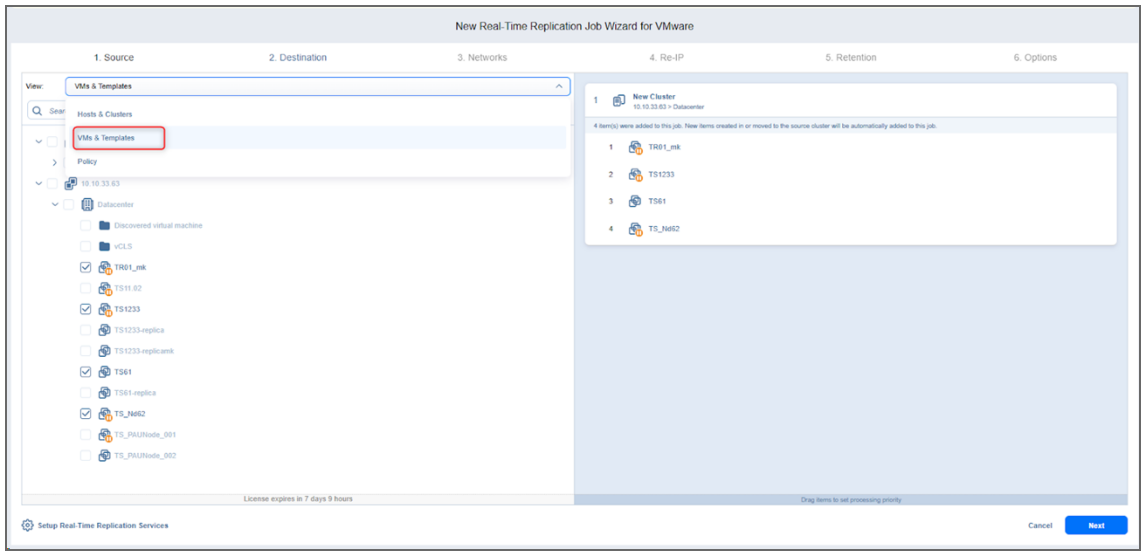
1. Optionally, you can filter the inventory tree by entering a string into the **Search** box. You can enter a part of or the entire name of the item.
2. Select VMware items by placing a checkmark to the left of each item. You can select only the VMs with assigned VM Storage policy. For more information on assigning the storage policy to VMs see I/O Filter Installation.
3. The selected items appear in the right pane of the page. You can reorder the selected items by dragging a VM or a container with the pointer to a new position. By doing that, you can specify to replicate the most important VMs first.
4. Review the list of selected VMware items. You can remove a selected VM or a container from the replication job in one of the following ways:
    - Deselect the item in the left pane. This will remove the item from the right pane; OR
    - In the right pane, hover the pointer over the item you want to remove and click the "x" on the

right. This will deselect the item in the left pane.




## VMs and Templates

When the **VMs & Templates** view is selected, the inventory tree displays VMware hosts, VMs, and VM templates. Proceed as described for the **Hosts & Clusters** view above.

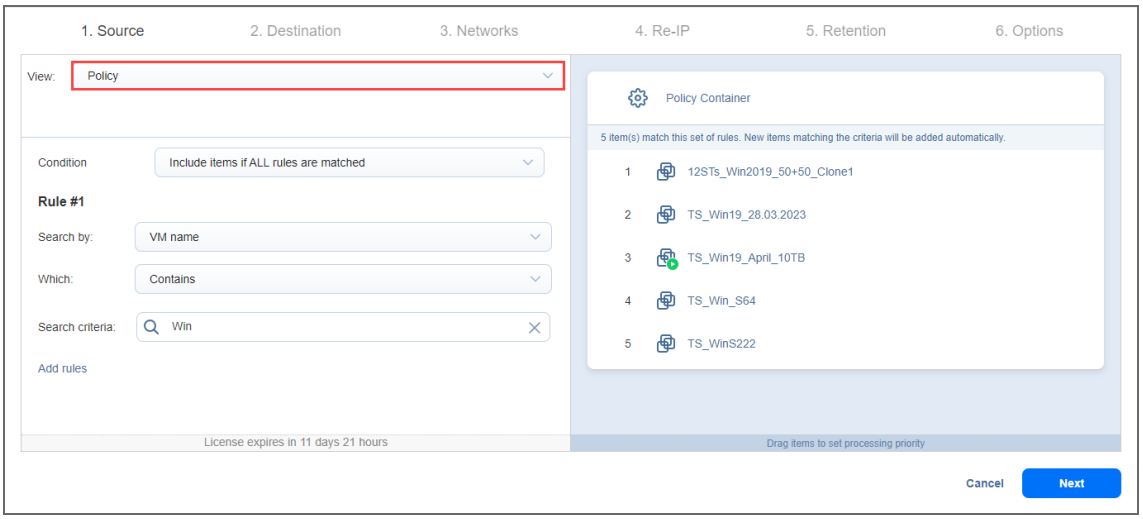



**Note**

For one real-time replication job for VMware, the Transporter's maximum capacity is limited to 11 source VMs in case there is one disk per VM or up to 50 disks per source VM. For more information, see this article.

# Policy

When the **Policy** view is selected, job policies can be used. Refer to “Managing Job Policies” on page 382 for details. If the items were selected in alternate views, a dialog box opens warning you that switching to the **Policy** view will reset your current selection. Click **Switch View** to confirm switching to the **Policy** view. Make sure that at least one item matches the available set of policy rules. Refer to “Managing Policy Rules” on page 385 for details.




Click **Next** to confirm adding selected VMs to the replication job. The wizard will display the next page.

**Notes**

- If you cannot find a VM or a container:
  - Ensure the corresponding vCenter has been added to the inventory.
  - Refresh the Inventory.
  - Ensure that the VM Storage policy is assigned to the source VM.
- By adding a VMware container to the job, you ensure that important VMs are always protected. If you add a VMware container to the job:
  - All VMs with defined VM storage policy that are currently available in the selected container will be replicated.
  - All new VMs with defined VM storage policy that will be created in (or moved to) the container in the future will be automatically added to the job and replicated.
- The order in which VMs are replicated is important if the **Transporter** performing replication cannot process all VMs of the job simultaneously — either because the **Transporter** is processing other jobs at the same time or because the job contains more VM disks than a **Transporter**'s maximum load specified during the **Transporter** creation.
- If all disks of a VM are unsupported (such as RDM disks in physical compatibility mode), this VM will be disabled in the inventory tree and it will not be possible to select it. Refer to "Supported Platforms" on page 131 for details.
- VMs marked as primary are protected by VMware Fault Tolerance. Disabled VMs are Fault Tolerance shadow VMs.

# Real-Time Replication Job Wizard for VMware: Destination

On the **Destination** page, select a location for your replica(s).

**Note**

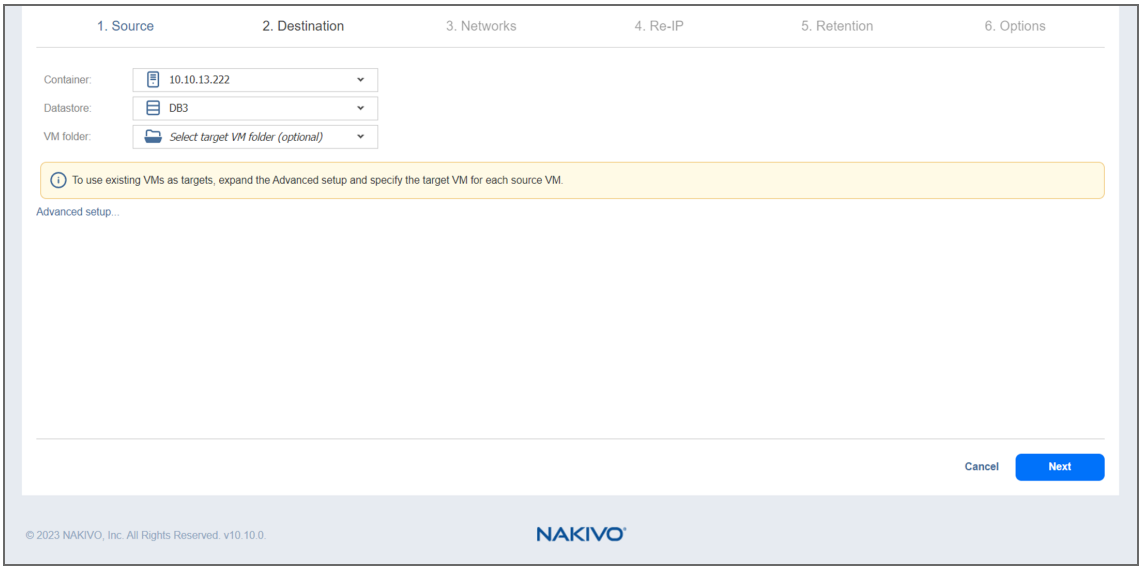
The target host/cluster must have a **Transporter** appliance and **Journal Service** installed. For more information, see "Deploying Transporter as VMware Appliance" on page 575.


- "Setting the Same Host, Datastore, and Network for All Replicas" on the next page
- "Setting the Default Destination for Replicas" on the next page
- "Setting Different Options for Replicas" on page 974
- "Mapping Source VMs to Existing Replicas" on page 974
- "Excluding a VM Disk from the Replication Job" on page 975

## Setting the Same Host, Datastore, and Network for All Replicas

To replicate all VMs to the same container and datastore, and to connect all replicas to the same network, follow the steps below:

1. Choose a cluster, host, or resource pool from the **Container** drop-down list.
2. Choose a datastore from the **Datastore** drop-down list.
3. Optionally, you can choose a target VM folder from the **VM folder** drop-down list.

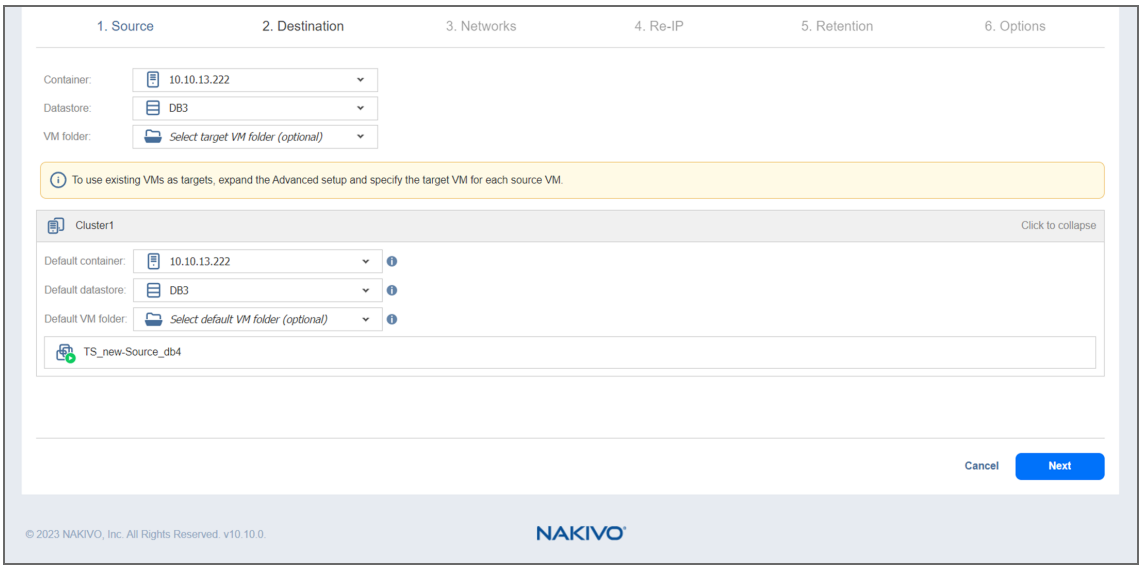



## Setting the Default Destination for Replicas

If you have chosen a host, cluster, folder, or a resource pool as a source for your replication job on the **Source** wizard page, you can set the default container, datastore, and VM folder for replicas. To do this, follow the steps below:

1. Click **Advanced setup** and then click on the name of the chosen host, cluster, folder, or a resource pool.
2. Choose a **Default container**.

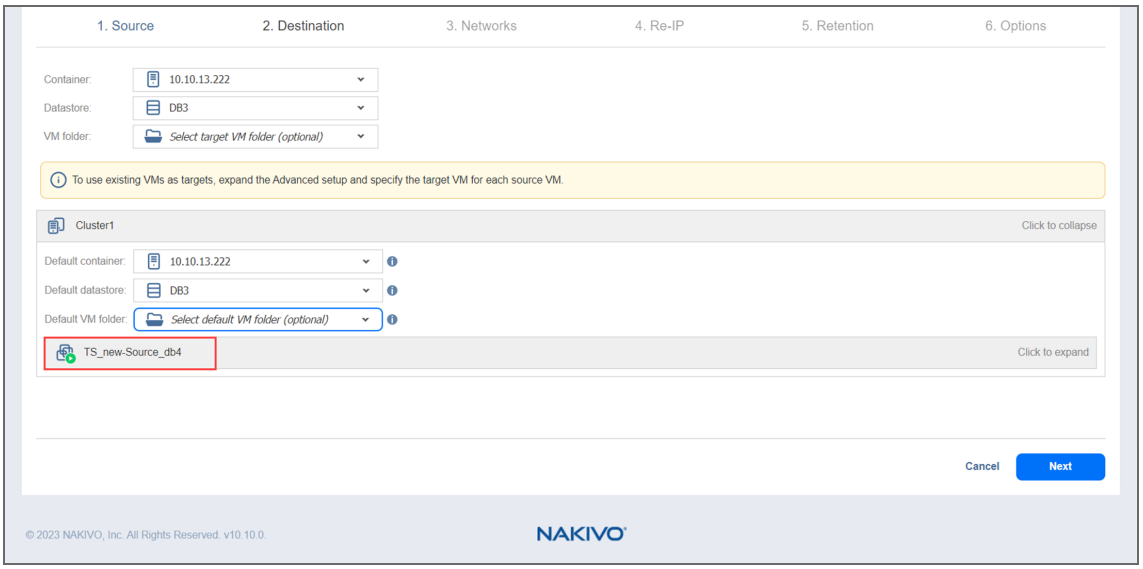
3. Optionally, you can also choose a **Default VM folder.**




## Setting Different Options for Replicas

To specify different replication options for VMs, follow the steps below:

1. Click **Advanced setup**.
2. Choose a target container, target VM, and target datastore for each VM.

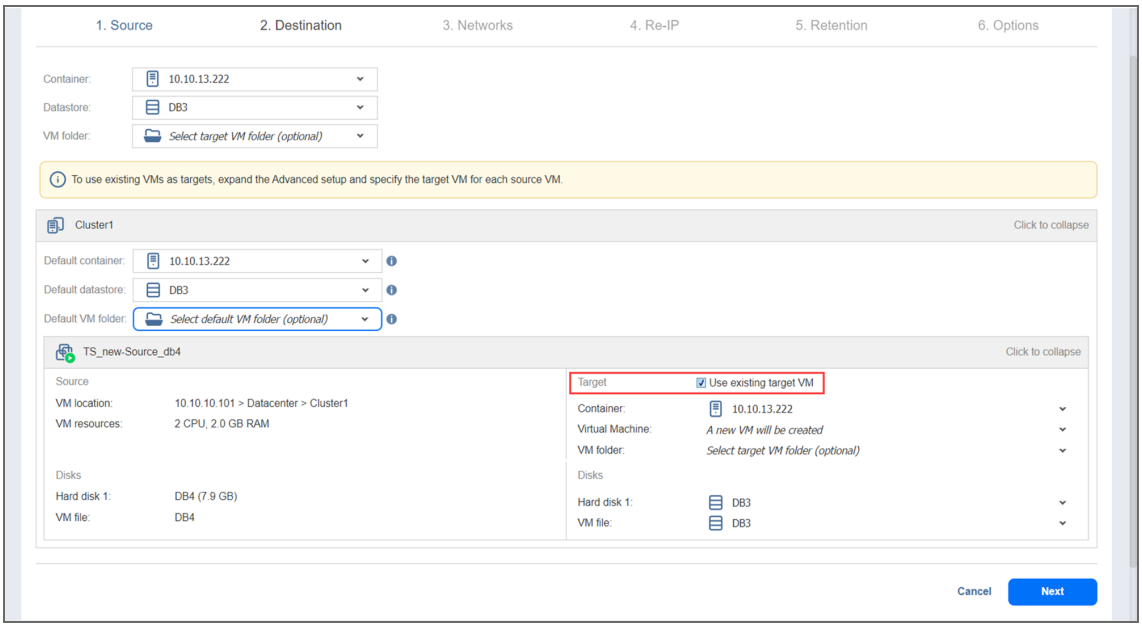



## Mapping Source VMs to Existing Replicas

If you want to perform staged replication or if you lose the replication job (due to accidental job deletion or because you need to recreate jobs in a new copy of the product), you can map source VMs to existing replicas to avoid running full VM replication again.

To map source VMs to existing VMs, follow the steps below:

1. Click **Advanced setup**.
2. Open the target VM drop-down list and select the **Use existing target VM** option.
3. Select the VM that should be used as a target for replication in the **Virtual Machine** drop-down list.




When you run the job, the product analyzes the target VM you have selected, determines how it is different from the source VM, and transfers only the differential data.

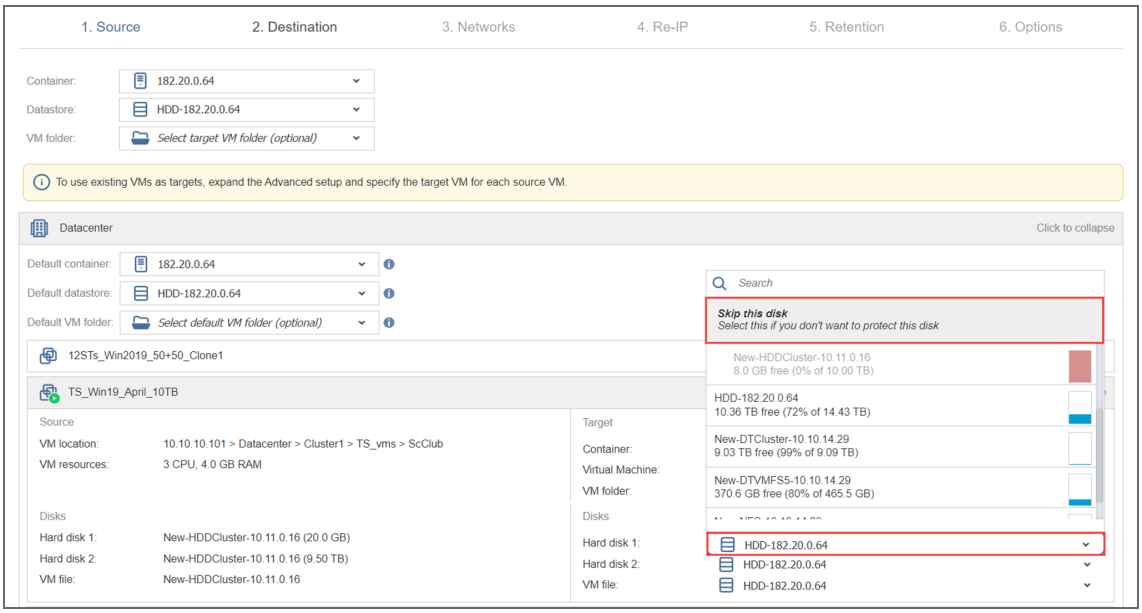
VM replication mapping can be a time-consuming process that can be equal to the time required to create a full VM replication.

## Excluding a VM Disk from the Replication Job

If you do not want to replicate some disks of a VM, you can exclude those disks from the replication job. Follow the steps below:

1. Click **Advanced setup**.
2. Open the target VM drop-down list.
3. Click the drop-down list next to the disk that you want to skip and select the **Skip this disk** option.

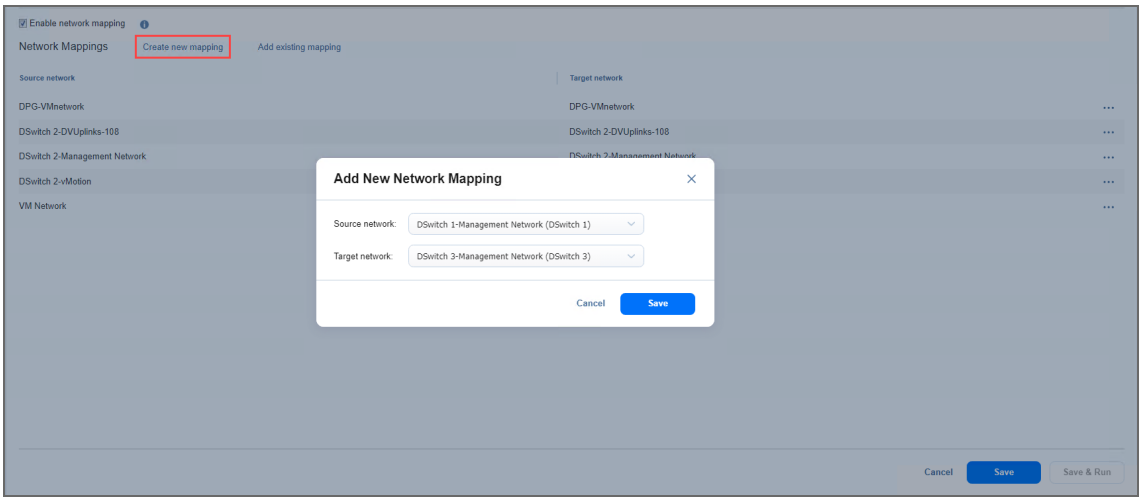
4. After you have specified replica location options, click **Next** to go to the next page.




# Real-Time Replication Job Wizard for VMware: Networks

To map source VM virtual networks to appropriate target virtual networks, do the following on the **Networks** page in the wizard:

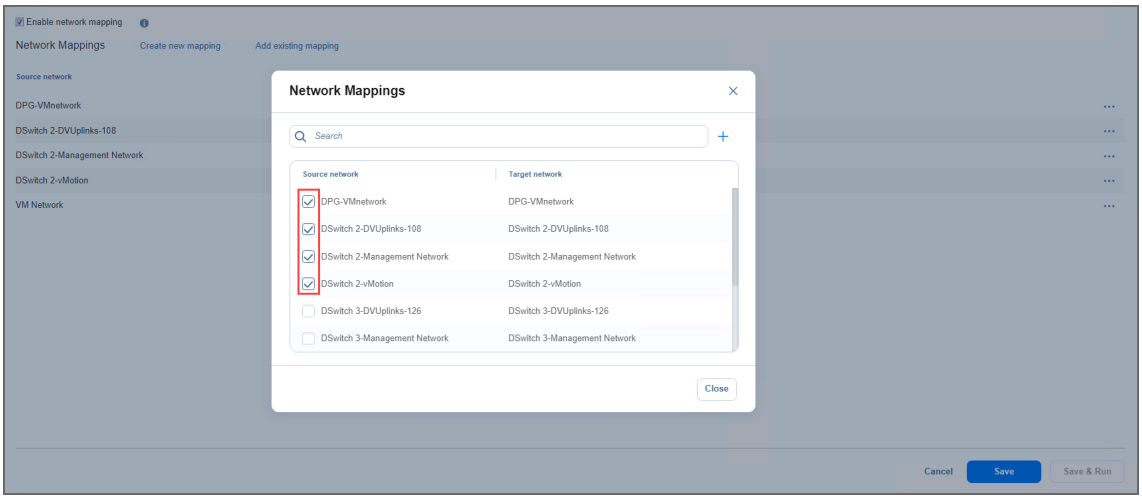
1. Select **Enable network mapping**.
2. The **Network Mappings** section opens. You have the following options:
   - Create a new mapping:
     a. Click **Create new mapping**.
     b. The **Add New Network Mapping** dialog box opens. Choose a source network and a target network and click **Save**.




   - Add an existing mapping:

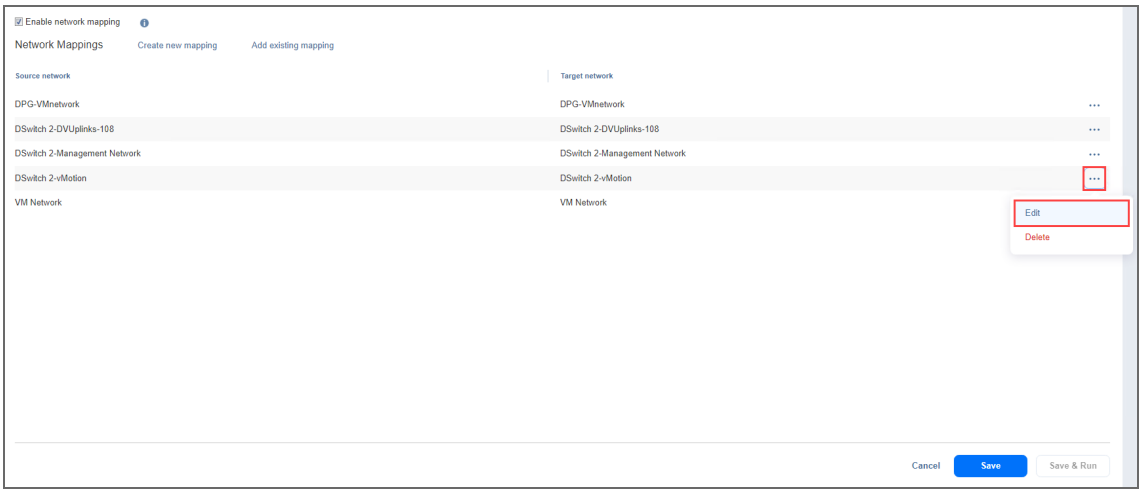
a. Click **Add existing mapping**.

b. The **Network Mappings** dialog box opens. Choose one or more appropriate network mappings and close the dialog box.

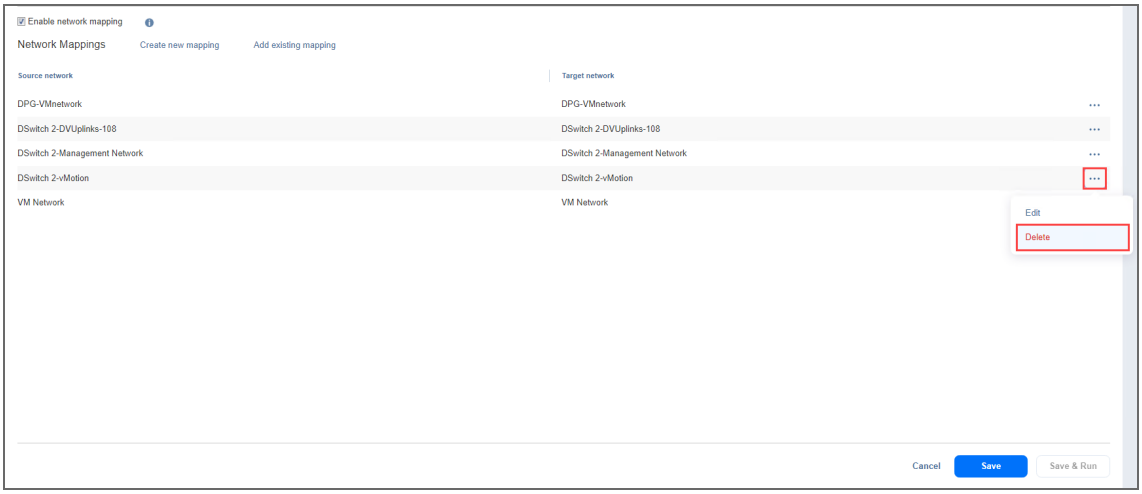



- Edit an existing mapping:

1. Hover the pointer over the required item in the **Network Mappings** list and click the ellipsis **Manage** button to the right of the item. Then, click **Edit**.
2. The **Edit Network Mapping** dialog box opens. Choose the required item from the **Target network** drop-down list and click **Save**.




- Delete an existing mapping: Hover the pointer over the required item in the **Network Mappings** list and click the ellipsis **Manage** button to the right of the item. Then, click **Delete**.




3. Click **Next** to go to the next page of the wizard.

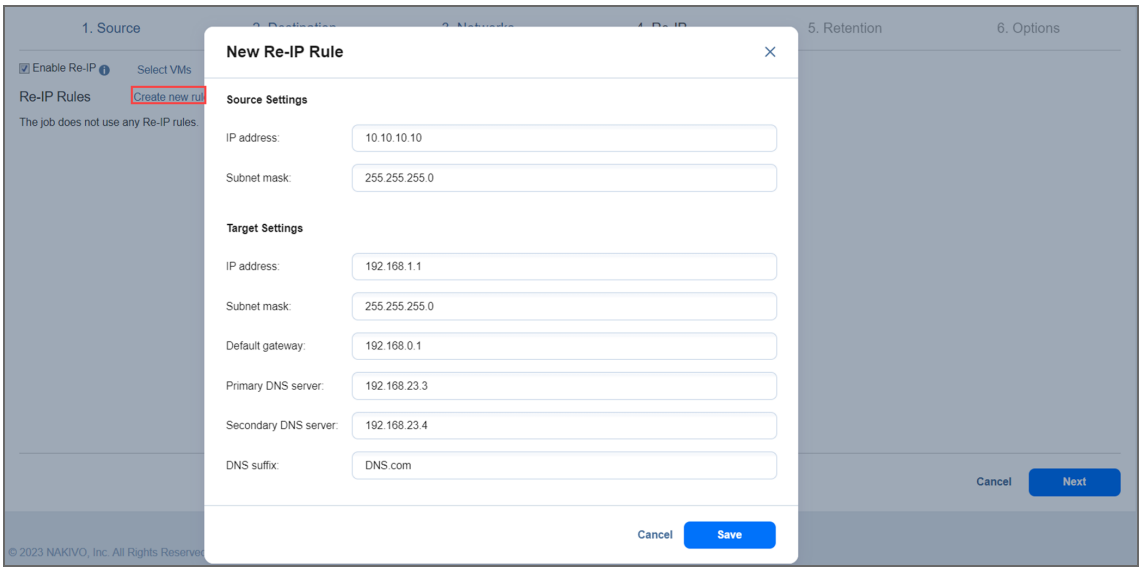
# Real-Time Replication Job Wizard for VMware: Re-IP

> **Note**
> Re-IP rules will be applied only to VMs that have a static IPv4 address configured.

To enable Re-IP rules for your replication job, do the following on the **Re-IP** page of the wizard:

1. Select **Enable Re-IP**.
2. Click the **Select VMs** link.
3. The **Re-IP** dialog box opens. In the list of your source VMs, select at least one, and close the dialog box.
4. You have the following options:
    - Create a new rule:
        a. Click **Create new rule**.
        b. The **New Re-IP Rule** dialog box opens. Enter source and target settings for the Re-IP rule

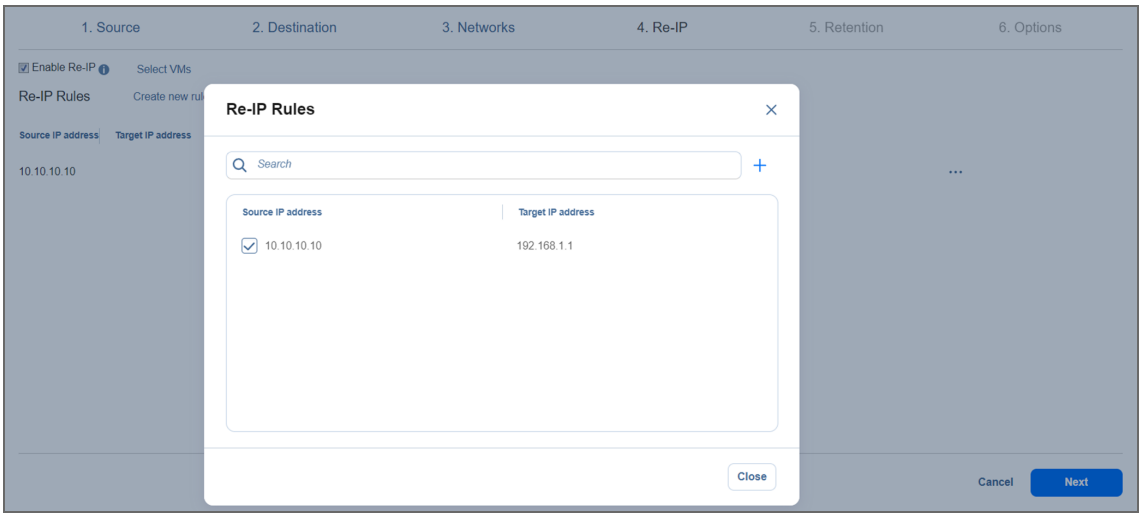
and click **Save**.




**Note**

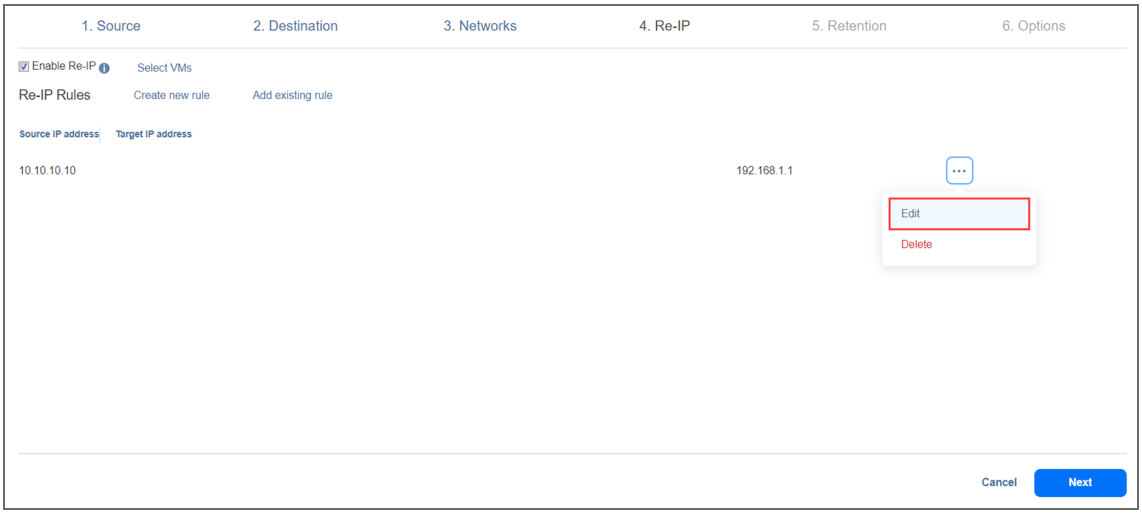
You can use wildcards for IP addresses. Refer to the "Real-Time Replication Failover Job Wizard for VMware: Re-IP" on page 1141 topic for details.

- Add an existing rule:
  a. Click **Add existing rule**.
  b. The **Re-IP Rules** dialog box opens. Select one or more appropriate Re-IP rules and close the dialog box.

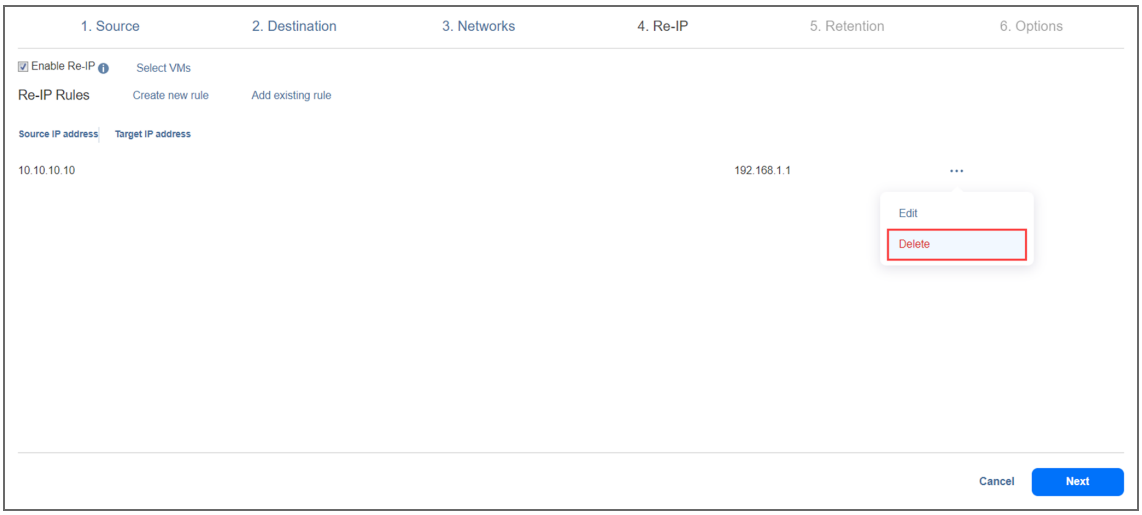



- Edit an existing Re-IP rule:
  a. Hover the pointer over the required item in the **Re-IP Rules** list and click the ellipsis **Manage** button to the right of the item. Then, click **Edit**.

b. The **Edit Re-IP Rule** dialog box opens. Edit the required properties of the Re-IP rule and click **Save**.




- Delete an existing mapping: Hover the pointer over the required item in the **Re-IP Rules** list and click the **Delete** icon to the right of the item.




- To leave the list of existing Re-IP rules intact, go to the next page of the wizard.
- Click **Next** to go to the next page of the wizard.

**Notes**

- Re-IP Rules enabled for replication jobs are only stubs for failover jobs. They do not work at the replication stage.
- Re-IP rules that are enabled for your replication job can be used for creating the corresponding failover jobs. Refer to "Real-Time Replication Failover Job Wizard for VMware: Re-IP" on page 1141 for details.

# Real-Time Replication Job Wizard for VMware: Retention

To allow replica rollback to a point in time your product keeps an **I/O Journal** of all I/O activities. This can be defined on the **Retention** page of the wizard by configuring the Recovery Point Objective (**RPO**) and **I/O Journal** settings as described in the sections below:


- RPO
- I/O Journal


## RPO

In the **RPO** section, configure the following parameters:

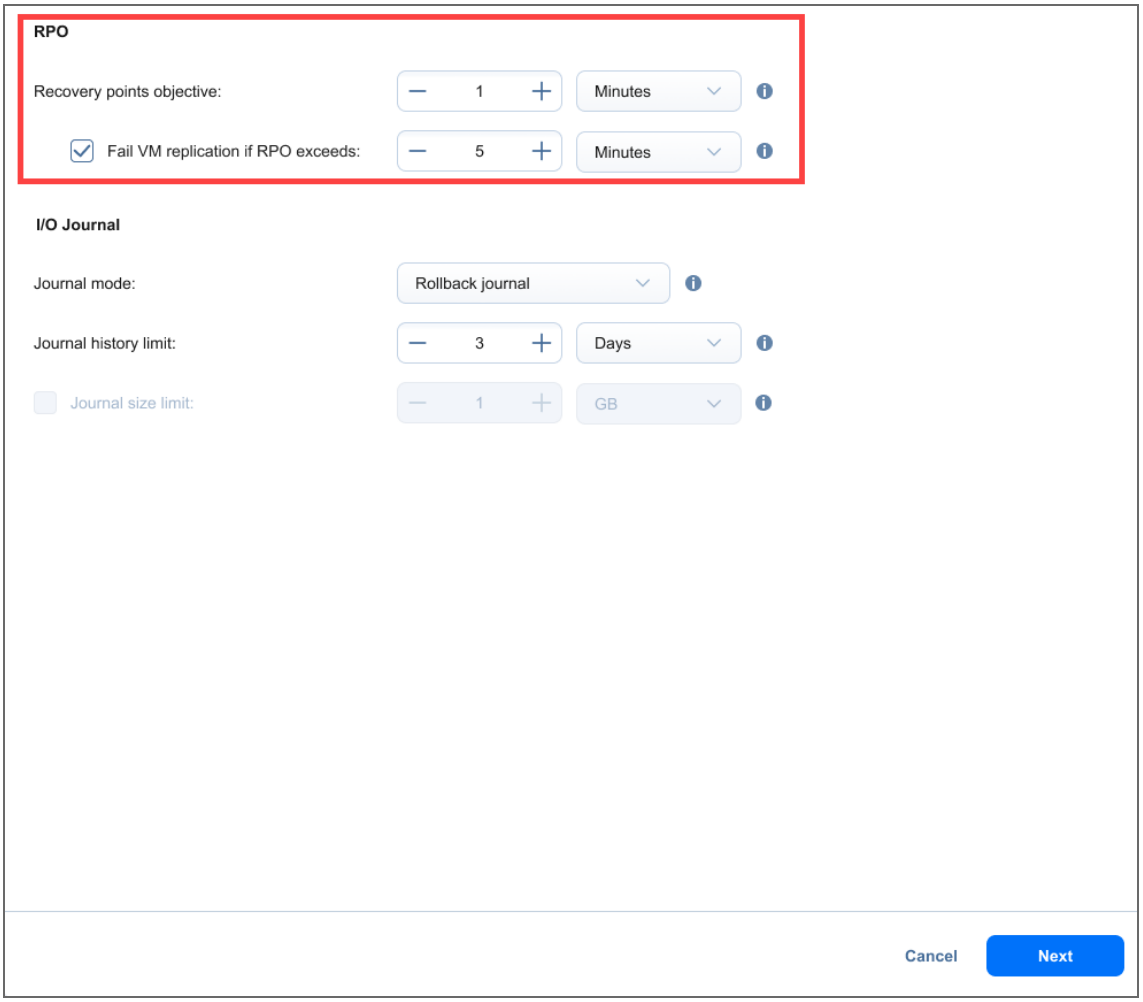
- **Recovery point objective**: Set the target amount of time between recovery points creation (between 1 second and 60 minutes) for the creation of subsequent recovery points.

  **Notes**

  The **RPO** elements are:

  - Now - the current time. Current **RPO** is calculated on the target side (**Journal Service**) as time between the last source data capture and the previous source data capture
  - Last received time - the time when the last heartbeat was received by the **Journal Service**.
  - Last captured data time - the time when the last portion of source VM data was captured by the **I/O Filter Daemon** (a single instance of the **I/O Filter** software which is running on a specific ESXi host of the cluster).
  - Last recovery point time - the time when the last recovery point was created in the **I/O Journal**.

- **Fail VM replication if RPO exceeds**: Optionally, select this option to fail and retry real-time replication job for VMware if the **RPO** exceeds a specified period of time. The set time cannot be lower than the selected Recovery point objective and cannot exceed 60 minutes.




Failed VMs and VMs with missed **RPO** will appear in red on top of the list.

> **Important**
>
> Make sure the time on the source ESXi host(s) and the target Transporter VA is synchronized.
>
> Otherwise, the **RPO** value and/or recovery point date-time data may be incorrect. Also, the **Real-Time Replication Failover Job Wizard for VMware** process may yield the incorrect recovery points.

If the job is not running (is created/stopped), the "No data" message will be displayed.

If the job is running:

- During initial and incremental replication, the "No data" message will be displayed.

> **Notes**
>
> - **Initial replication** - snapshot based full regular replication.
> - **Incremental replication** - snapshot based incremental regular replication.

- During real-time replication process and when the **RPO** data is available, the current **RPO** will be displayed.

> **Note**
> Current **RPO** is calculated as time between the last source data capture and the previous source data capture.

> **Notes**
> The **RPO** data will be displayed as:
> - If time span <= 1 minute: "Xs".
> - If time span =< 1 hour and > 1 minute: "Xm Ys".
> - If time span =< 1 day and > 1 hour: "Xh Ym Zs".

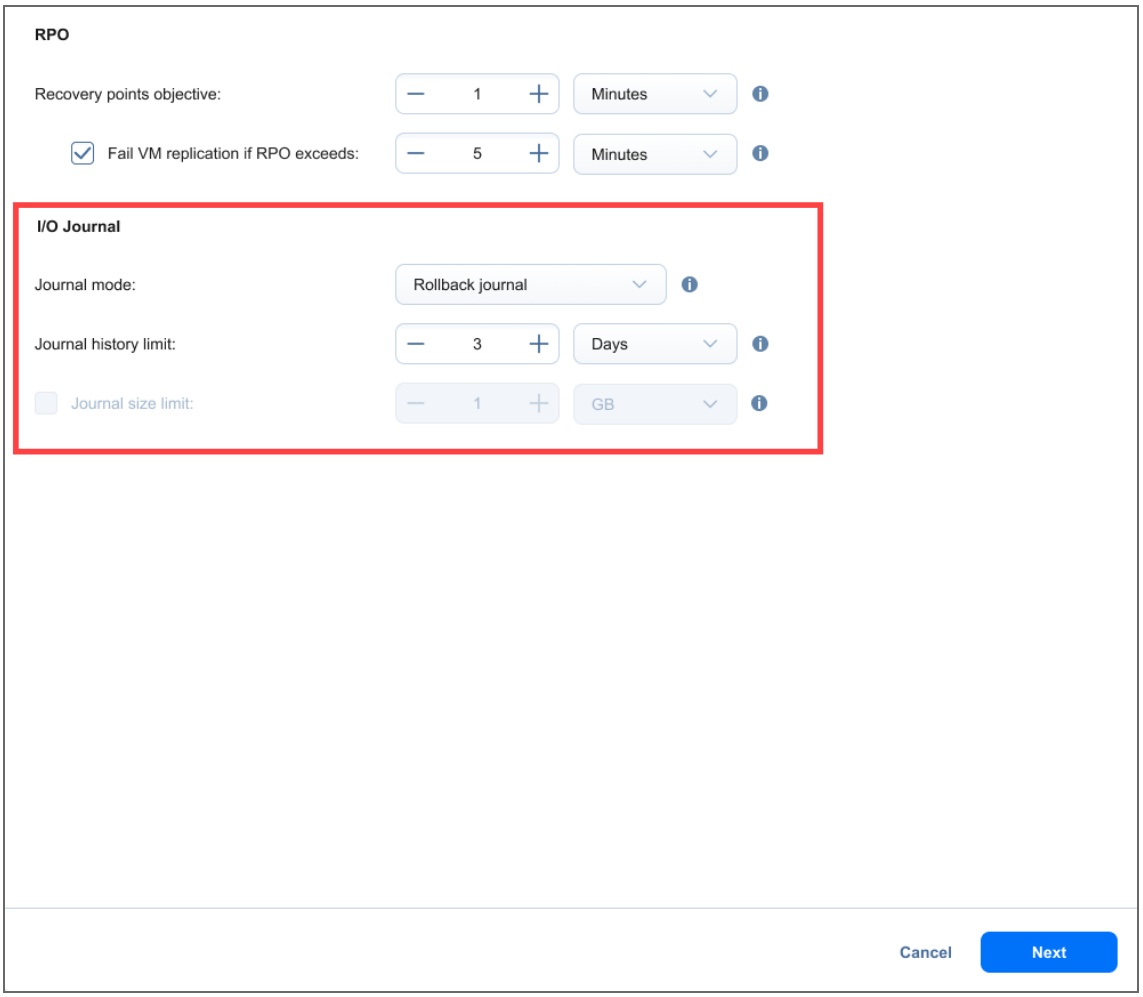
# I/O Journal

**I/O Journal** is a per disk set of records of I/O operations that keeps recovery points of a replica VM and can be used for failover to a point in time.
For every real-time replication run, a new **Journal Extent** is created in the **I/O Journal**. Each **Journal Extent** contains one or more **Frames** (internal logical parts of the **Journal Extent** containing up to 256 MB of data).
Configuration changes made to the source VM (if any) are checked every 15 minutes, and saved to the replica configuration journal. If no data changes on the Source VM are detected, the source **I/O Filter Daemon** will keep sending heartbeat messages to the target **Journal Service**.

> **Notes**
> If the connection from the source **I/O Filter** to the **Journal Service** (target) is lost:
> - **I/O Filter** keeps capturing the data on the source and caching/storing it locally.
> - **I/O Filter** keeps trying to re-connect to the **Journal Service** (target).
> - After the connection is restored, the data is sent to the **Journal Service** (target).

In the **I/O Journal** section, configure the following parameters:

- **Journal mode**: In this dropdown menu you can select one of the following:
  - **Rollback journal**: With this option selected, new data changes are saved to the replica directly, and the old data in the replica is saved to the journal. Old data in the journal is removed based on the journal settings.
  - **Roll forward journal**: With this option selected, new data changes are saved to the journal, and the old data is merged to the replica based on the journal settings.
- **Journal history limit**: Optionally, select this option to set a limit for journal history, which is between 1 hour and 30 days.
- **Journal size limit**: Set a limit for the journal size per disk between 1 GB and 20 TB.




**Note**

The journal size will be limited to the set value per VM. If the existing journal size is smaller than the limit specified in the job, it can be extended. Changing the existing journal to a smaller size is not applicable.

When the journal history limit is reached or when the **I/O Journal** has under 512 MB of free space, the following occurs:

- If the replica **I/O Journal** contains multiple **Journal Extents**, the oldest **Journal Extent** is removed.
- If the replica **I/O Journal** contains only one **Journal Extent**, its oldest **Frame** is applied to the replica. Afterwards, the oldest frame is overwritten by a new **Frame**.
- If the replica **I/O Journal** contains only one **Journal Extent** with only one **Frame**, the **Frame** is not applied to the replica.

**Note**

For example, if the size of the single **Journal Extent Frame** (256 GB) is not reached within 1 hour (the specified journal history limit), the data allocated within this **Frame** will not be merged and the recovery points of a replica VM that were created according to the **RPO** setting will not be deleted.

In this case, the data merge process to the replica will continue until the **I/O Journal** frame size (256 GB) is reached during the subsequent time frame.

In case the size of the single **Journal Extent Frame** (256 GB) is reached within the next time frame, a new **Frame** will be generated and the recovery points created before the second **Frame** generation will be deleted.

For more information on the above terminology, see "Installing Journal Service" on page 590.

# Real-Time Replication Job Wizard for VMware: Options

On the **Options** page of the wizard, configure job, replica, and data transfer settings as described in the sections below:

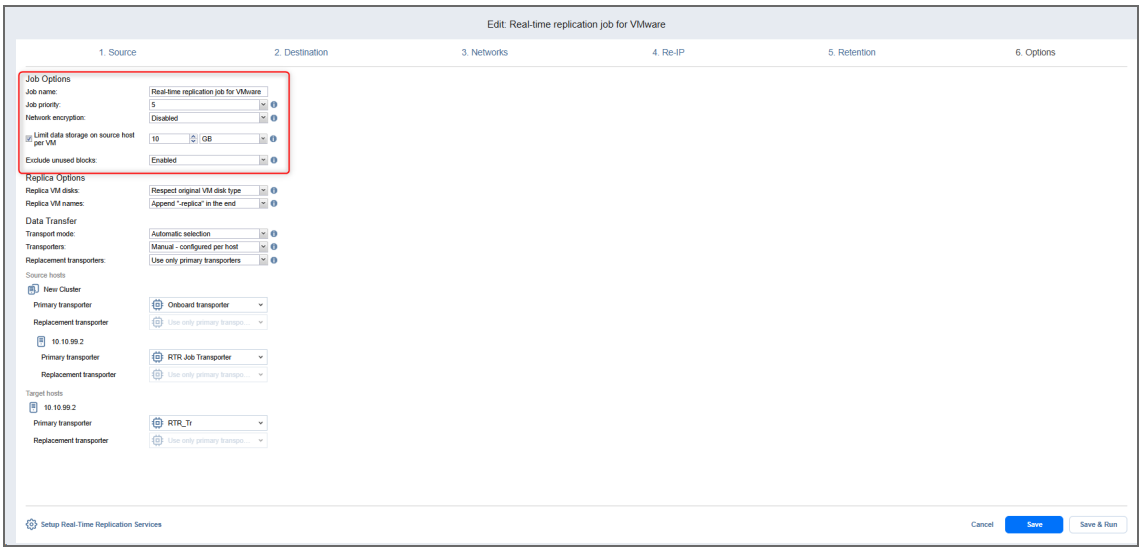
- Job Options
- Replica Options
- Data Transfer

## Job Options

In the **Job Options** section, configure the following:

- **Job name**: Set a name for the real-time replication job.
- **Job priority**: Select a job priority level between 1 and 5, with 1 being the highest priority. Jobs with higher priority levels are prioritized during job processing.
- **Limit data storage on source host per VM**: Optionally, select this option to limit the data stored on the source host before transferring to the target host to the set size per VM. If the size limit is exceeded for a VM, replication for that VM fails.
- **Exclude unused blocks**: Select this option to automatically exclude unused disk blocks and blocks occupied by deleted files while processing source objects running Windows OS. This allows you to reduce backup storage space and object processing time.
- **Network encryption**: Optionally, enable this option to encrypt the traffic between **I/O Filter daemo**n and **Journal service** (over port 33092). The encryption uses the certificate on the **Transporter**, which

establishes a secure connection to the **Journal Service**.




# Replica Options

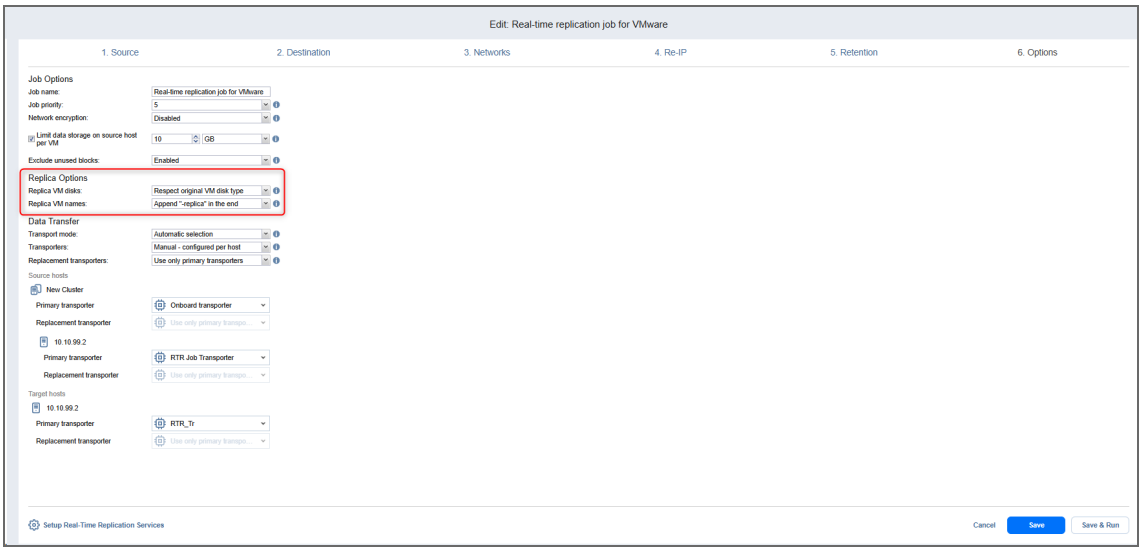
In the *Replica Options* section (available for VMware jobs only), configure the following:

- **Replica VM disks**: Choose one of the following options from the drop-down menu:
  - **Respect original VM disk type**: With this option enabled, the created disk will be of the same type as that of the source VM.
  - **Create only thin disks on target VMs** (VMware vSphere): With this option enabled, only thin disks will be created on target VMs. Use this option to save space on target datastore.

> **Note**
>
> The **Replica VM disks** menu is disabled if target objects already exist for all source objects of the replication job and the user will receive the "Disk type cannot be changed for the existing replica." message.

- **Replica VM names**: Choose one of the following replica VM names options:
  - **Append “-replica” in the end**: Source VM names are used for replica names and “-replica” is added to the replica name.
  - **Leave replica names as is**: Replica names will be identical to the source VM names.

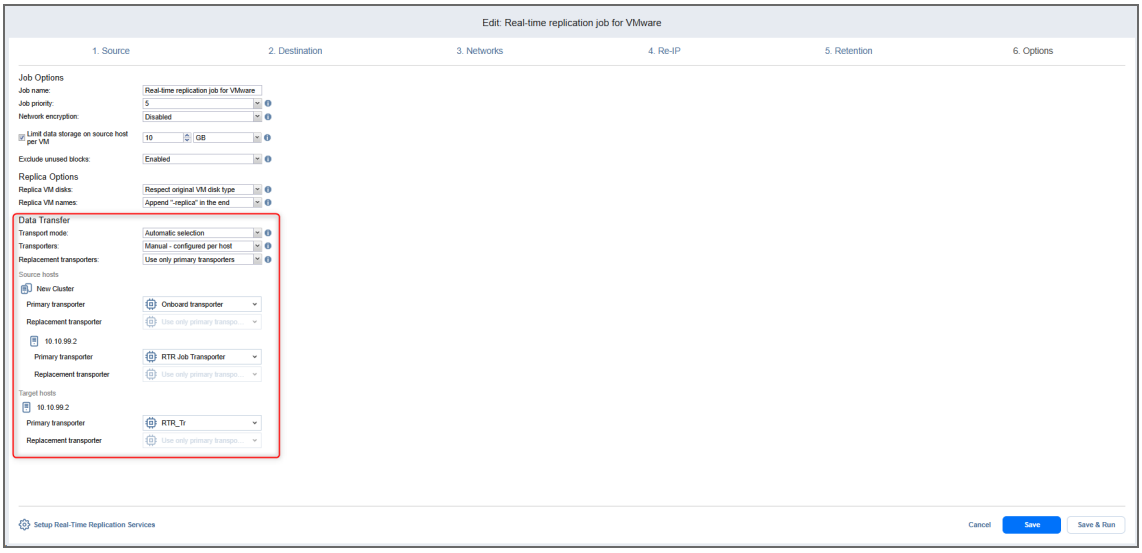
- **Enter custom replica names**: Enter custom names for replicas.




# Data Transfer

In the *Data Transfer* section, configure the following:

- **Transport mode**: Choose one of the following transport modes for the real-time replication job:
  - **SAN only**: With this option selected, NAKIVO Backup & Replication only uses direct SAN access to retrieve VM data. If direct SAN access to VM data is not available, the job will fail.
  - **Hot Add only**: With this option selected, NAKIVO Backup & Replication only uses Hot Add to retrieve VM data. If direct Hot Add is not available, the job will fail.
  - **LAN only**: With this option selected, NAKIVO Backup & Replication only retrieves VM data via LAN.
- **Transporters**: Select the **Transporters** to be used for the Real-Time Replication job. With this option, manually specify a **Transporter** to be used for each source and target host.




- **Replacement transporters**: Select a primary **Transporter** for each separate host.

# Staging (Seeding) VM Replication

With VMs usually being large, the initial (full) VM replication can be slow and time-consuming, and can put an undesirable load on the network. Perform staged replication to speed up the initial VM replication and save network bandwidth. Staging requires the transfer of VMs to the target site using a removable medium (such as an external USB hard drive). You can then create a new replication job that will use the transferred VMs as a target and perform only incremental replication.

To stage VM replication, follow the steps below:

1. Put VMs that you want to replicate on a removable medium (such as an external USB hard drive) using backup, replication, or any other method.
2. Transfer the medium to the target location.
3. Add (recover) the VMs to the desired server and datastore.
4. Create a new replication job and map the source VMs to the transferred VMs.

# Recovery

During outage events that threaten business continuity, NAKIVO Backup & Replications offers multiple recovery options allowing you to resume normal business operations swiftly.

This section covers the following recovery topics:


- “Granular Recovery” on page 991
- “Full Recovery” on page 1060
- “Planning Disaster Recovery” on page 1106
- “Starting Recovery from Tape” on page 1093
- “Performing Flash Boot Recovery” on page 1185

# Granular Recovery

The granular recovery technology allows you to instantly recover specific files and objects from image-based backups. With this technology, you can easily recover corrupted or accidentally deleted files or objects without fully restoring a VM first. With NAKIVO Backup & Replication you can recover files for physical machines, VMware, Microsoft Hyper-V, AWS, and Nutanix virtual environments. You can also recover Microsoft Exchange emails, and Microsoft Active Directory and Microsoft SQL Server objects directly from compressed and deduplicated backups. In addition, multiple recovery jobs and/or users may access the same recovery point even if it is currently in use by an existing recovery job/session.

Before you start the recovery process, verify that:

- The target VM/instance/physical machine is powered on.
- The target VM/instance/physical machine has enough space. The required minimum of free space is equal to the size of the recovered object + 1 GB.
- The target VM/instance/physical machine is accessible over the network.

For more details, refer to the corresponding articles below:


- "File Recovery" on page 992
- "Object Recovery for Microsoft Exchange" on page 1012
- "Object Recovery for Microsoft Active Directory" on page 1025
- "Importing Recovered Objects to Active Directory" on page 1039
- "Object Recovery for Microsoft SQL Server" on page 1040
- "Performing Universal Object Recovery" on page 1051

# File Recovery

With NAKIVO Backup & Replication, you can recover files or folders directly from compressed and deduplicated backups. Refer to “Instant File Recovery to Source” on page 30 for more information.

**Notes**

- File recovery is restricted to supported disk types and file systems.
- You can recover files or folders from compressed and deduplicated Proxmox VE VM backups directly to source or custom VMs.

Refer to the following topics to learn more:


- “Opening File Recovery Wizard” on page 993
- “File Recovery Wizard: Backup” on page 995
- “File Recovery Wizard: Recovery Method” on page 998
- “File Recovery Wizard: Files” on page 1004
- “File Recovery Wizard: Options” on page 1007
- “File Recovery Wizard: Finish” on page 1011

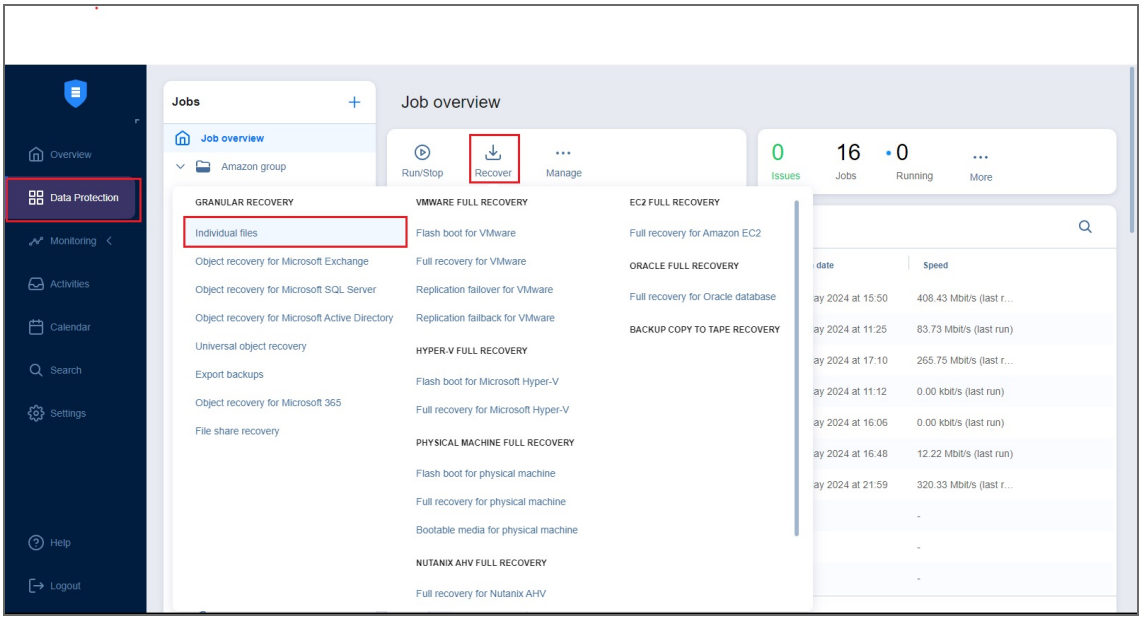
# Opening File Recovery Wizard

You can start the recovery process either from the **Data Protection** menu, by using the Search function, or from the **Repositories** tab in **Settings** (for example, if you no longer have a backup job, but still have the backup). Refer to the following sections for more details:

- Starting File Recovery from Data Protection Menu
- Starting File Recovery from Backup Repository

## Starting File Recovery from Data Protection Menu

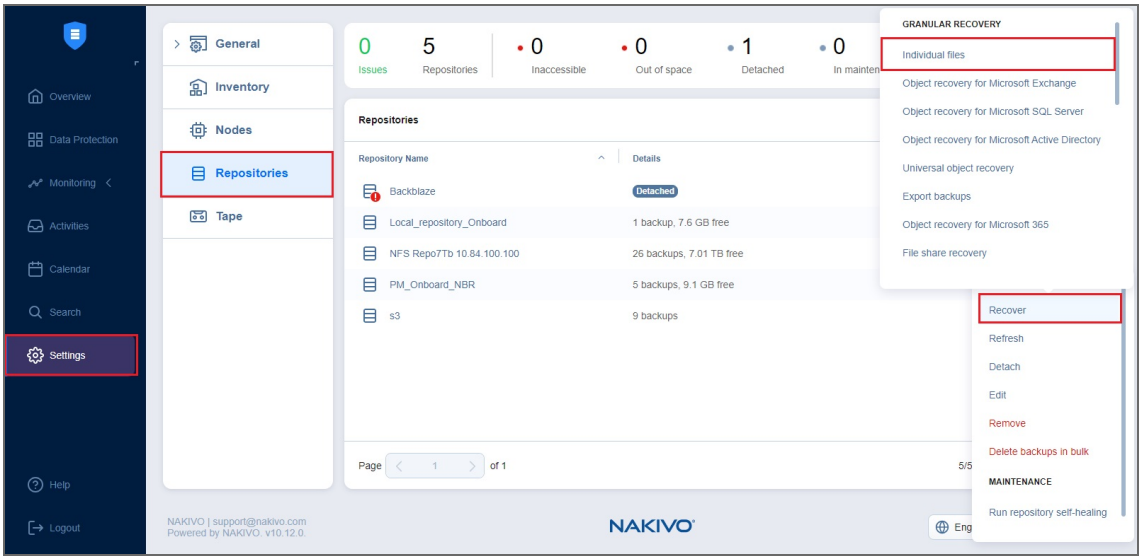
To start file recovery from the **Data Protection** menu, click **Recover** and then click **Individual Files**.




## Starting File Recovery from Backup Repository

To start file recovery from a Backup Repository, do the following:

1. From the main menu of NAKIVO Backup & Replication, click **Settings**.
2. Go to the **Repositories** tab and hover over the **Backup Repository** containing the required backup.


Overview
Data Protection
Monitoring
Activities
Calendar
Search
Settings
Help
Logout
General
Inventory
Nodes
Repositories
Tape
0 Issues
5 Repositories
0 Inaccessible
0 Out of space
1 Detached
0 In mainten
Repositories
Repository Name
Details
Backblaze
Detached
Local_repository_Onboard
1 backup, 7.6 GB free
NFS Repo7Tb 10.84.100.100
26 backups, 7.01 TB free
PM_Onboard_NBR
5 backups, 9.1 GB free
s3
9 backups
Page 1 of 1
5/5
GRANULAR RECOVERY
Individual files
Object recovery for Microsoft Exchange
Object recovery for Microsoft SQL Server
Object recovery for Microsoft Active Directory
Universal object recovery
Export backups
Object recovery for Microsoft 365
File share recovery
Recover
Refresh
Detach
Edit
Remove
Delete backups in bulk
MAINTENANCE
Run repository self-healing
NAKIVO | support@nakivo.com
Powered by NAKIVO. v10.12.0
NAKIVO

# File Recovery Wizard: Backup

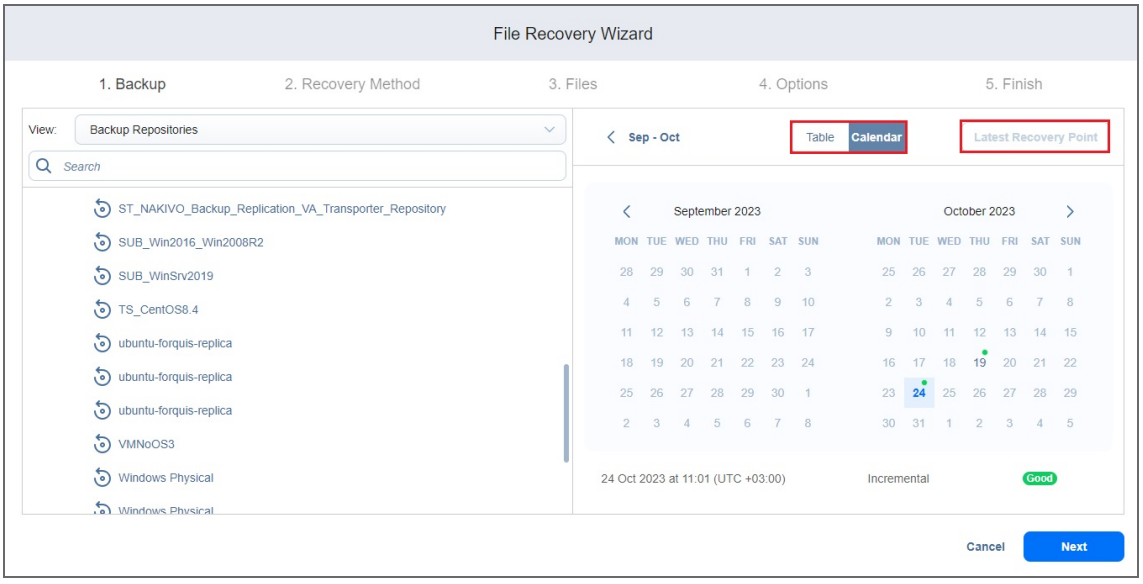
1. On the **Backup** page of the wizard, select a backup using the **Backup Repositories** or **Jobs & Groups** view in the left pane.

Both federated and standalone (not used as members of federated repositories) backup repositories can be selected.

**Notes**

- You cannot select a federated repository member as a source for recovery jobs.
- You cannot perform recovery from recovery points located in an inaccessible federated repository member or those that depend on inaccessible recovery points.
- You can search for a specific backup, job, group, or repository (depending on the selected view) by entering its name into the **Search** field.
- Selecting a single backup object opens the **Calendar** view of the date picker if this view was the last user-selected view. Selecting a single backup object opens the **Table** view if this was the last view the user selected.
- You can select a backup object with the corrupted, inaccessible, or pending removal recovery points, but you cannot select the backup objects of the detached repository.

2. The latest recovery point is selected by default.



**Notes**

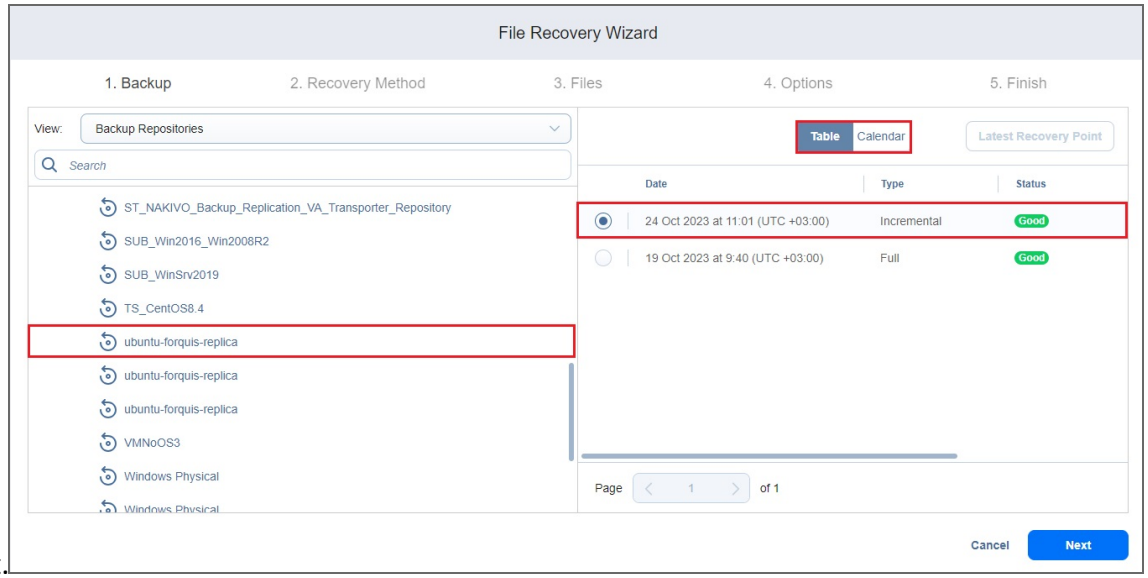
- The selected date is highlighted.
- If a recovery point is selected in the **Calendar** or **Table** view, this recovery point is selected by default the next time you open the **Calendar** or **Table** view.
- The selected view, either **Calendar** or **Table**, is saved on a per-user basis.

3. If necessary, toggle between the **Calendar** and **Table** views:
   - In the **Calendar** view, select a date or click **Latest Recovery Point** to select the date with the latest recovery point.

**Notes**

- The **Latest Recovery Point** option is disabled if:
  - The latest recovery point is currently selected.
  - All the recovery points of the backup object are corrupted, inaccessible, or pending removal.
- A small dot is displayed on top of the date if there is at least one recovery point created on that date.
- The selected recovery point has a preview displayed under the calendar showing the recovery point status:
  - **green** – verified recovery points
  - **red** – inaccessible/corrupted/pending removal recovery points
  - **dark grey** – unverified recovery points
- Corrupted, inaccessible, or pending removal recovery points cannot be recovered.

   - In the **Table** view, select a date or click **Latest Recovery Point** to select the date with the latest recovery point.



**Note**

Clicking the **Latest Recovery Point** option selects the date with the latest recovery point that is not corrupted, inaccessible, or pending removal.

4. Click **Next** to go to the next page of the wizard.

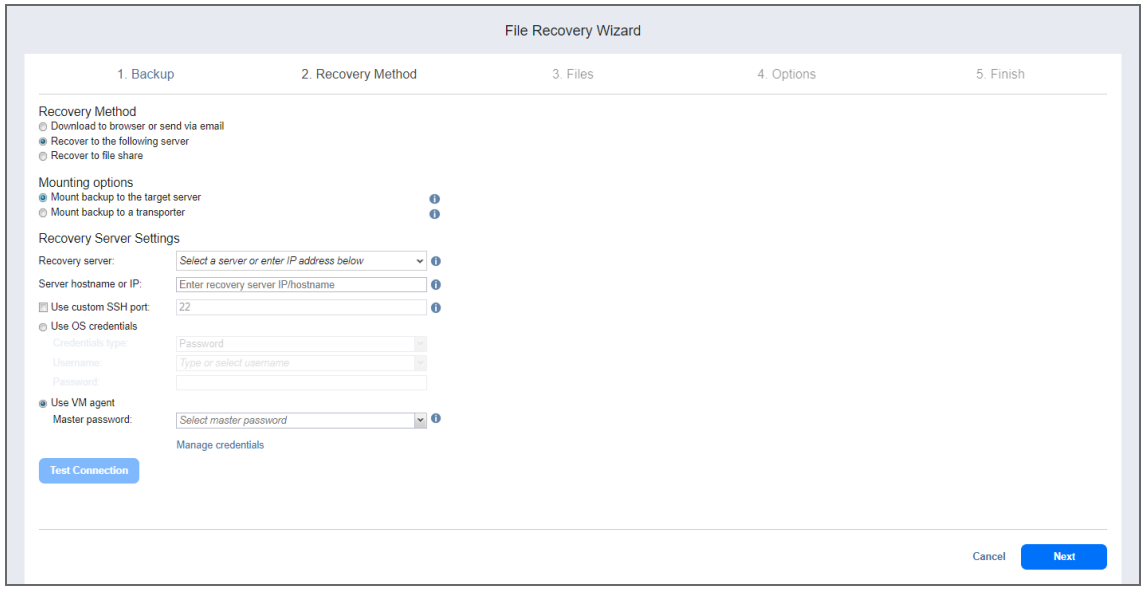
To learn about the limitation on the chain of incremental recovery points, refer to the Knowledge Base article.

# File Recovery Wizard: Recovery Method

In this page of the wizard, choose one of the following recovery methods:

- Recovering Files to a Server
- Downloading Files to Browser or Sending Files via Email
- Recovering Files to a File Share

## Recovering Files to a Server




When recovering files to the server, you can select one of the following options:

- Mounting a backup to the target server
- Mounting a backup to a transporter

### Mounting Backup To the Target Server

With this option, the selected backup is directly mounted to the recovery server for data processing. To recover files to the target server, follow the steps below:

1. In the **Recovery Method** section, choose **Recover to the following server**.
2. The **Mounting Options** section opens.
3. Select **Mount backup to the target server** (default) to mount the selected backup(s) for data processing directly to the recovery server.
4. In the **Recovery Server Settings** section, set the following options:

a. **Recovery server**: Choose the target server from the drop-down list.

> **Notes**
>
> - NAKIVO Backup & Replication tries to auto-detect the IP address automatically.
> - File recovery to the original location is executed via a system account.

b. **Server hostname or IP**: Enter the IP address of the recovery server if it is not detected by the application based on the recovery server name.

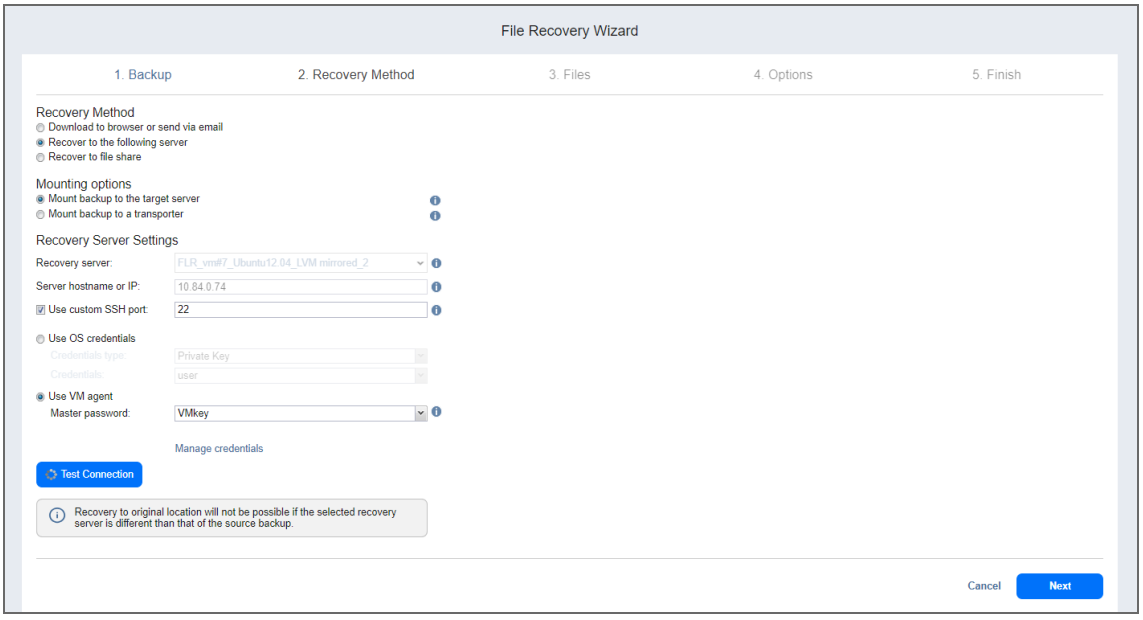
c. **Use custom SSH port**: If necessary, enter an SSH port to be used for connecting to the recovery server. The default value is `22`.

d. **Credentials type**: Choose your preferred option and enter your respective credentials:

   a. **Password**: Enter a username with administrative privileges for the file share entered above and your password.

   b. **Private key**: Select your private key from the drop-down list.

5. Select the **Use VM Agent** option to recover the files through the Permanent Agent. The product automatically fills out the Master password if the Permanent Agent is detected from the selected recovery server.
6. Click the **Test Connection** button to test your credentials for the specified recovery server. If your credentials are correct, a checkmark appears to the right of the button.
7. Click **Next**.




After NAKIVO Backup & Replication prepares a recovery point, the next page of the wizard opens.

Mounting Backup To a Transporter

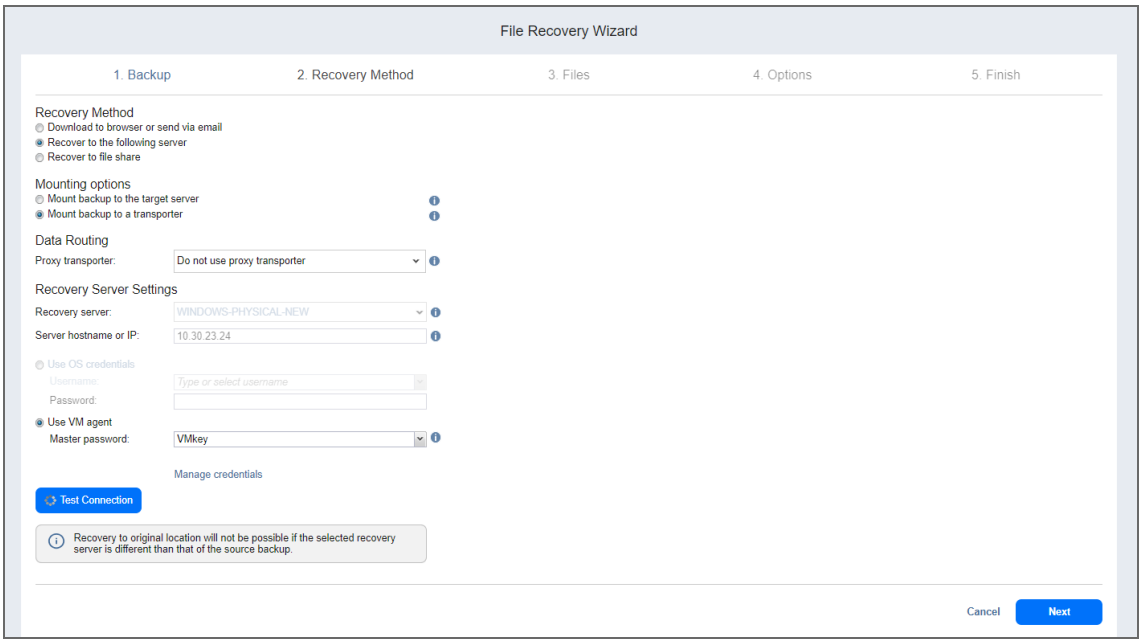
With the **File Restore through Permanent Agent** feature, you can manually select a proxy transporter instead of the source transporter chosen by default to mount the backups and communicate with the auto-detected permanent virtual machine agent (VMA) at the target destination. Do the following:

1. In the **Recovery Method** section, choose **Recover to the following server**.
2. The **Mounting Options** section opens.
3. Select **Mount backup to a transporter** to select the proxy transporter.
4. The **Data Routing** section is displayed. In the **Proxy transporter** drop-down list, select the target server where the recovery point is mounted. This proxy transporter is used instead of the source transporter to communicate with the assigned transporter at the recovery server for data processing of the selected backup.
5. Optionally, keep the default **Do not use proxy transporter** option to mount the recovery point to the source transporter.

   **Notes**

   - Recovering files to the Permanent Agent using the NAS transporter as a proxy transporter is not supported.
   - The **Do not use proxy transporter** option is disabled if the source transporter is NAS.

6. In the **Recovery Server Settings** section, choose the target server from the **Recovery server** drop-down list. Enter the IP address of the recovery server if it is not detected by the application based on the recovery server name.
7. Select the **Use VM Agent** option to recover the files through the Permanent Agent. The product automatically fills out the **Master password** if the Permanent Agent is detected from the selected recovery server.
8. Click the **Test Connection** button to test your credentials for the specified recovery server. If your credentials are correct, a checkmark appears to the right of the button.
9. Click **Next**.




**Note**

If **Mount backup to a transporter** is selected in the **Mounting Options** section, the **Use OS credentials** option is disabled.

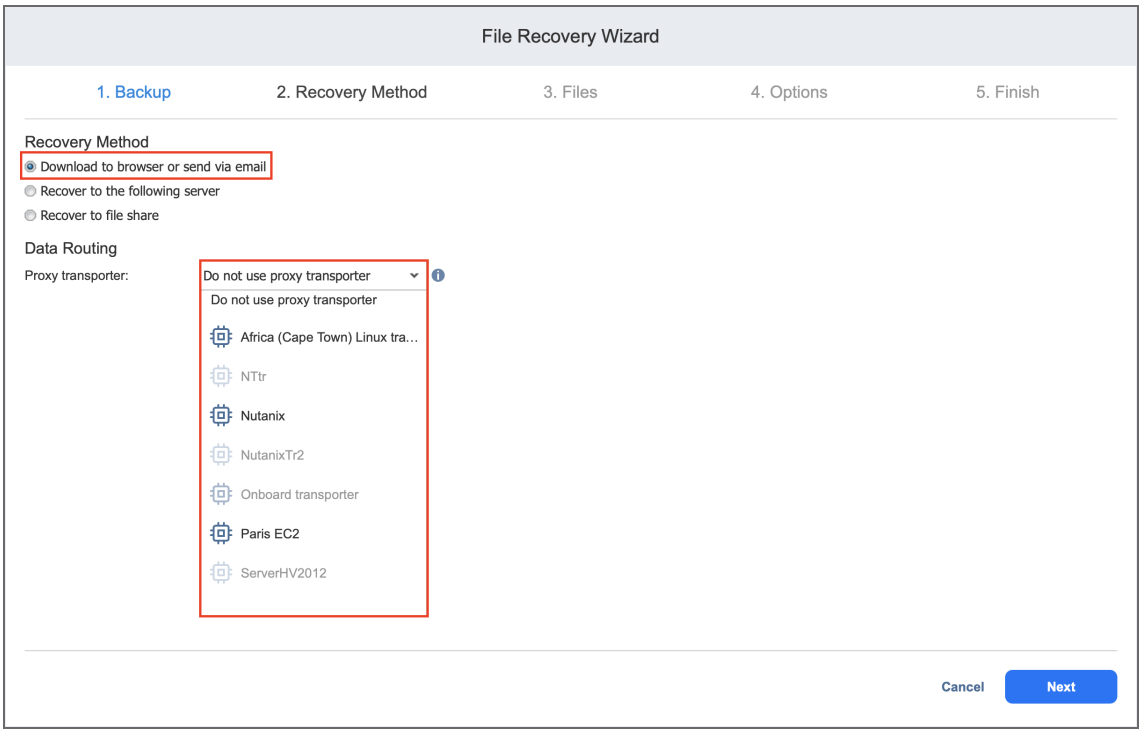
After NAKIVO Backup & Replication prepares a recovery point, the next page of the wizard opens.

## Downloading Files to Browser or Sending Files via Email

To download files to your browser or send them via email, follow the steps below:

1. In the *Recovery Method* section, choose **Download to browser or send via email**.
2. The *Data Routing* section opens. In the **Proxy transporter list**, the **Do not use proxy transporter option** is chosen by default. You can also choose a proxy transporter from the list of available transporters.

3. Click **Next**.




**Notes**

NAKIVO Backup & Replication will use a proxy transporter in the following cases:

- The Transporter assigned to the backup repository is missing support for some file systems.
- The Transporter assigned to the backup repository is missing iSCSI packages.

NAKIVO Backup & Replication starts preparing a recovery point for the recovery. After the recovery point is prepared successfully, the next page of the wizard opens.

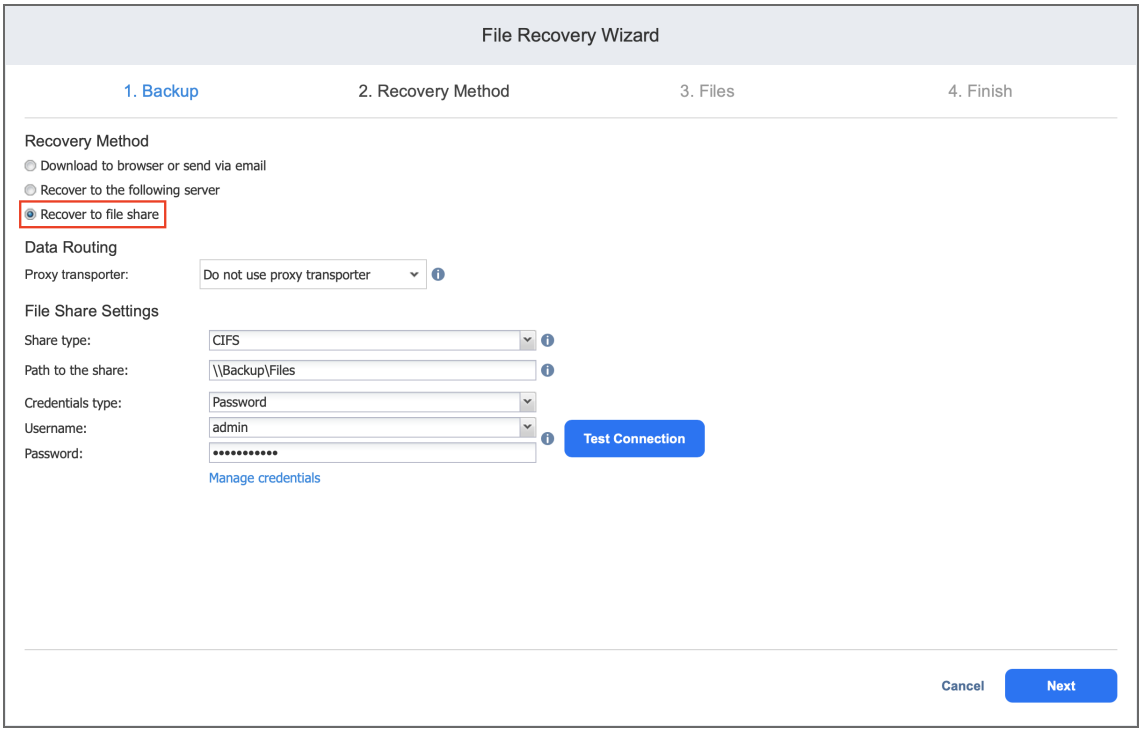
## Recovering Files to a File Share

To recovery files to a file share, do the following:

1. In the *Recovery Method* section, choose **Recover to file share**.
2. The *Data Routing* section opens. In the **Proxy transporter list**, the **Do not use proxy transporter option** is chosen by default. You can also choose a proxy transporter from the list of available transporters.
3. In the *File Share Settings* section, set the following options:
    a. **Share type**: Choose the type of file share.
    b. **Path to the share**: Enter the path to the file share.
    c. **Credentials type**: Choose your preferred option and enter your respective credentials:

a. **Password**: Enter a username with administrative privileges for the file share entered above and your password.
b. **Private key**: Enter your private key.

4. Click the **Test Connection** button to test your credentials for the specified file share. If your credentials are correct, a checkmark appears to the right of the button.
5. Click **Next**.




**Notes**

File-level recovery to CIFS share may fail if network credentials are shared across multiple open CIFS connections. If you encounter issues while testing connection, try adding local host domain to the used credentials (i.e. localhost\Administrator).

In addition, NAKIVO Backup & Replication will use a proxy transporter in the following cases:

- The **Transporter** assigned to the backup repository is missing support for some file systems.
- The **Transporter** assigned to the backup repository is missing iSCSI packages.

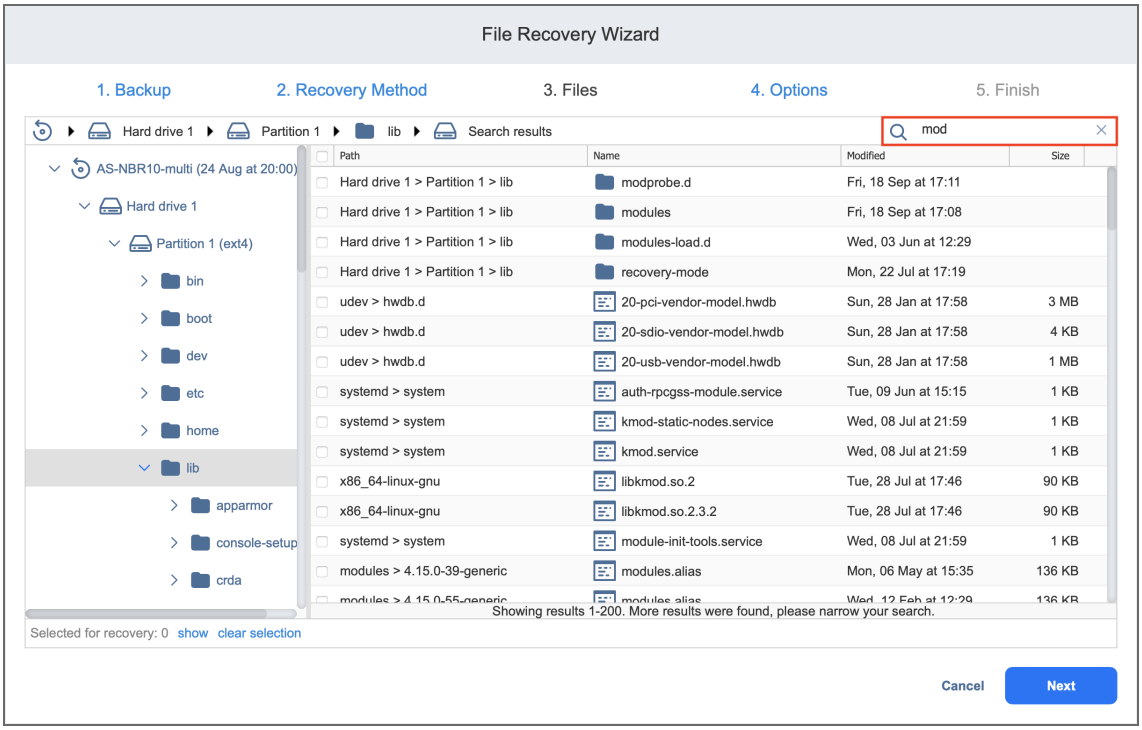
# File Recovery Wizard: Files

On this page of the wizard, select files for recovery.

- Searching for Files and Folders
- Browsing Files and Folders
- Selecting Files and Folders for Recovery

## Searching for Files and Folders

To search for a file or a folder, enter a part of or the entire name of the item into the **Search** box and press **Enter**.

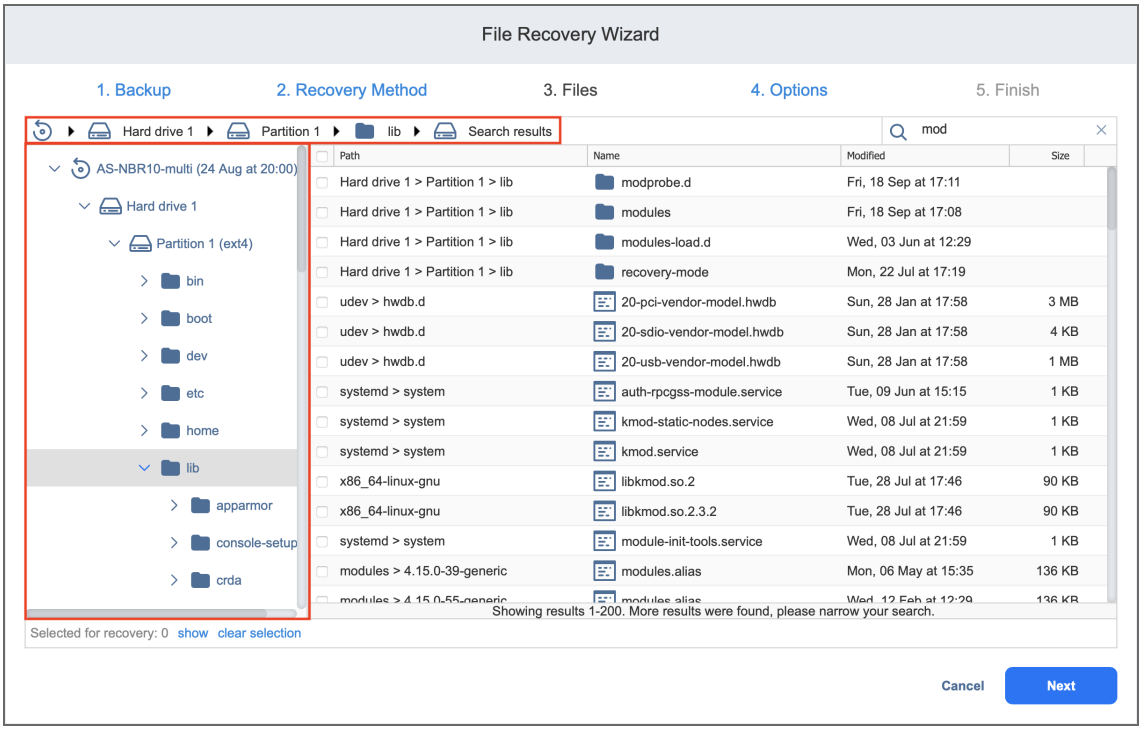



**Notes**

- NFS-mounted folders appear in the file tree as empty and the wizard does not recover the content of these folders.
- The search is performed starting from the point selected in the navigation pane. For example, if you select **Hard drive 1 > Disk 1 > Program Files**, the search will be performed only inside the **Program Files** folder.

## Browsing Files and Folders

You can browse the files and folders of a VM backup using the navigation pane:



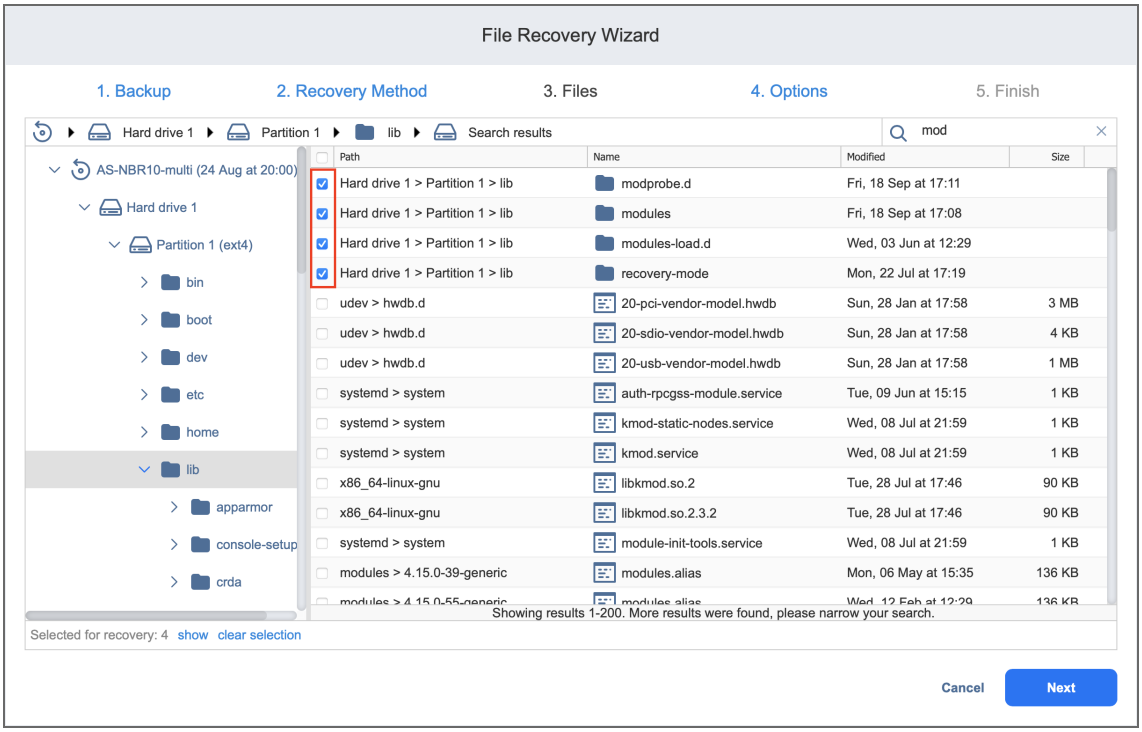

If a VM backup contains Linux LVM volumes or Windows dynamic disks, the navigation pane will display these logical groups in addition to all hard drives available in the VM backup. If a hard drive does not contain any partitions and servers as a part of a Linux LVM volume or a Windows dynamic disk, this hard drive will appear as empty.

You can also quickly move between folders by using the navigation bar above the navigation pane.

## Selecting Files and Folders for Recovery

After locating the item you want to recover, select the checkbox next to it. The number of items selected for recovery is displayed at the bottom of the wizard page. You can also do the following:

- Click **show** to view the list of all items selected for recovery.
- Click **clear selection** to clear the list of items selected for recovery.




Click **Next** to go to the next page of the wizard.

# File Recovery Wizard: Options

On this page of the wizard, you can choose one of the following recovery types:

- Recovering Files via Recovery Server
  - Recovering Files to the Original Location
  - Recovering Files to a Custom Location
- Downloading Files to Browser or Sending Files via Email
  - Downloading Files
  - Forwarding Files via Email

## Recovering Files via Recovery Server

If you have chosen the **Recover to the following server** recovery method, on the **Recovery Method** page of the wizard, proceed as follows.

**Important**

- File recovery is not possible if a backup contains an incomplete set of disks that are a part of the spanned volume/dynamic disks/LVM/RAID software or any other disk structures.
- Servers added using Direct Connect are not supported.

Recovering Files to the Original Location

To recover files to original location:

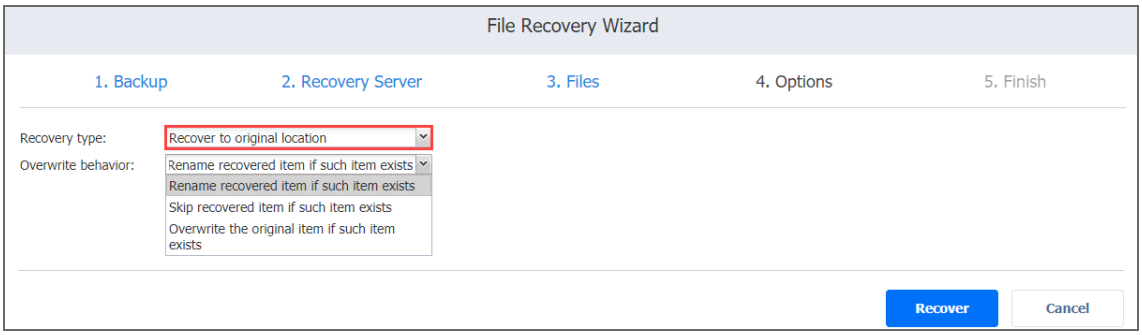
1. In the **Recovery type** list, choose **Recover to original location**.

   **Note**

   This option is disabled if the **Mount backup to a transporter** option in the **Recovery Method** page of the wizard has been selected.

2. The **Overwrite behavior** list opens. Please choose one of the following:
   - **Rename recovered item if such item exists**: Choose the necessary server from the drop-down list.
   - **Skip recovered item if such item exists**
   - **Overwrite the original item if such item exists**

3. **Click Recover** to start recovering files to original location.




Recovering Files to a Custom Location

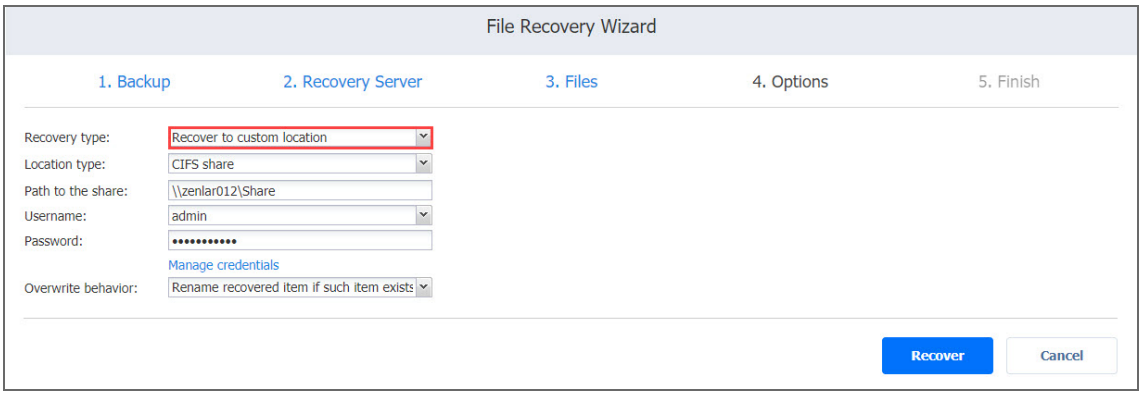
To recover files to a custom location:

1. In the **Recovery type** list, choose **Recover to custom location**.
2. A number of boxes open to let you set the options for a custom location. Do the following:
    a. In the **Location type** box, choose **Local folder on Recovery Server**.

> **Notes**
>
> - The **CIFS share** and **NFS share** options are disabled if the **Mount backup to a transporter** option in the **Recovery Method** page of the wizard has been selected.
> - If the selected archive is deleted from the share during the recovery process to CIFS share, the archive may still reappear in the folder and is deleted after the job is completed. Note that in such case the job is still marked as completed.

    b. For **Path to the local folder**, enter the local path to be used for file recovery.
    c. In the **Overwrite behavior** box, choose of of the following:
        - **Rename recovered item if such item exists**
        - **Skip recovered item is such item exists**
        - **Overwrite the original item if such item exists**
    d. In the **Username** and **Password** boxes, enter the credentials required for accessing the CIFS share location you specified above.

3. Click **Recover**.



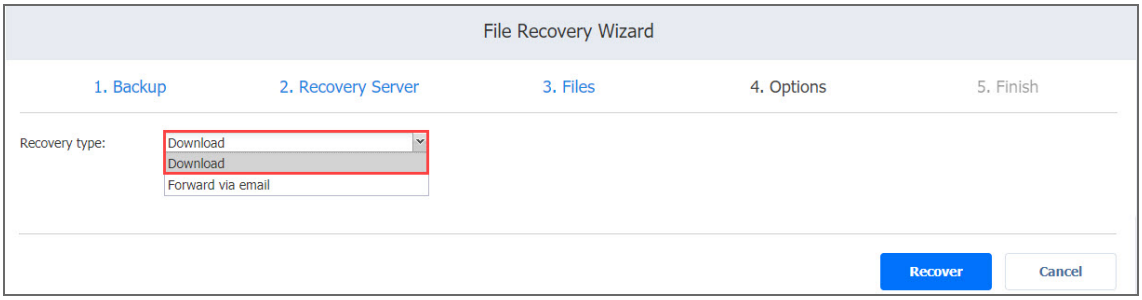

## Downloading Files to Browser or Sending Files via Email

If you have chosen the **Download to browser or send via email** recovery method, on the **Recovery Method** page of the wizard, proceed as follows.

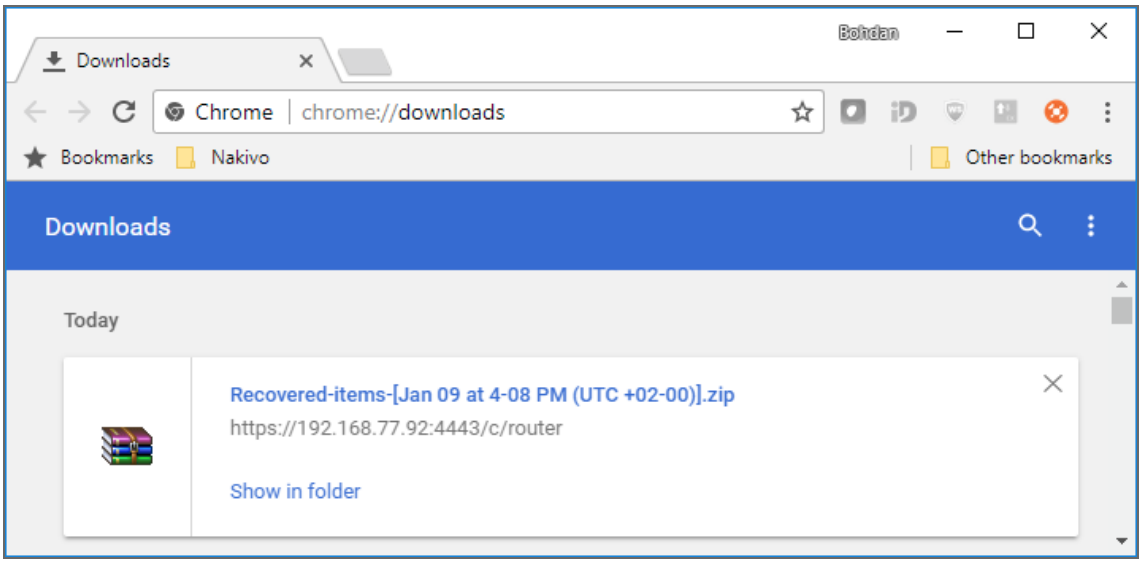
### Downloading Files

Please do the following to download files for recovery:

1. In the **Recovery Type** drop-down list, select **Download**.
2. Click **Recover**.




When the download has finished successfully, the archive with the recovered items appears in the browser downloads folder.




Forwarding Files via Email

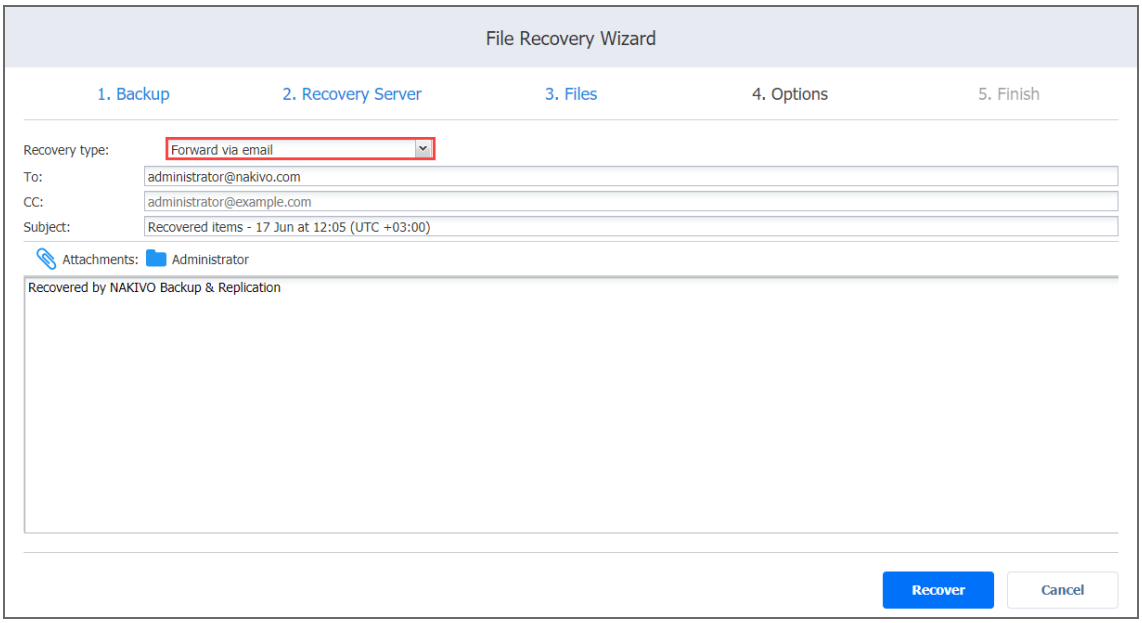
**Note**

To use this recovery type, your Email settings must be properly configured in the NAKIVO Backup & Replication Configuration. Refer to "Email Settings" on page 419 for details.

Please do the following to forward recovered files via Email:

1. In the **Recovery type** list, choose **Forward via email**.
2. A number of boxes open to set the options required for forwarding recovery files via email. Do the following:
   a. In the **To** box, enter one or more email addresses to be primary recipients of the recovery files. Use semicolons to separate multiple email addresses. The recipient's email address is mandatory.
   b. Optionally, in the **CC** box, you can enter one or more email addresses of secondary recipients.
3. Optionally, you can enter a subject in the **Subject** box.

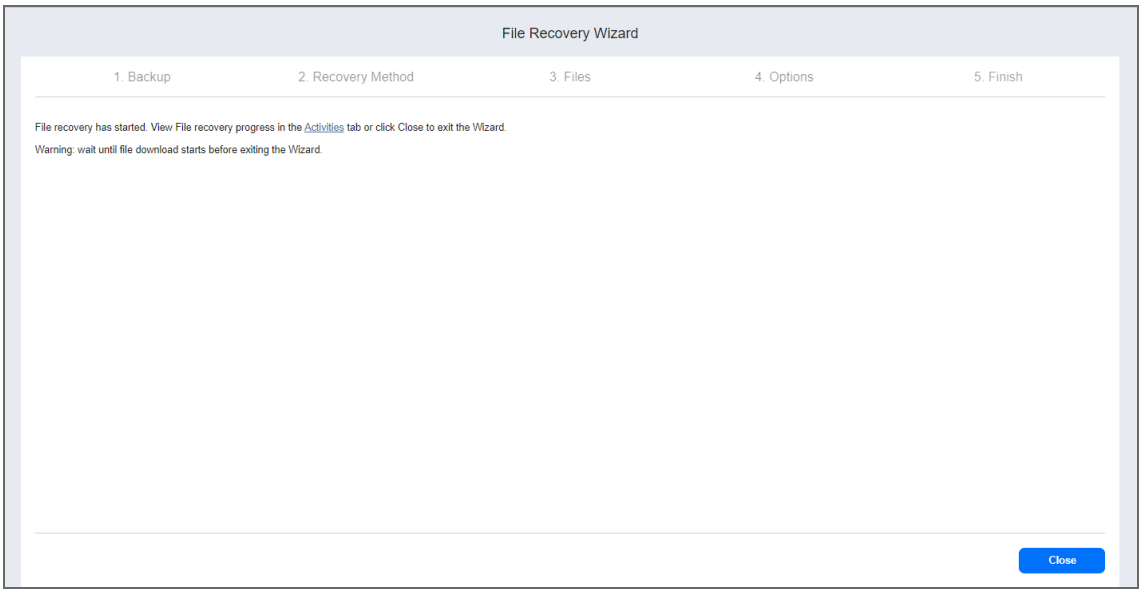
4. Click **Recover**.




# File Recovery Wizard: Finish

Click the **Activities** link to open the Activities tab.

Click **Close** to close the wizard.




**Note**

Once in this step, you cannot return to the previous steps.

# Object Recovery for Microsoft Exchange

The object recovery feature in NAKIVO Backup & Replication allows you to browse, search, and recover Microsoft Exchange emails directly from compressed and deduplicated backups. Recovery can also be performed back to the source or any other location including CIFS share. The Object Recovery for Microsoft Exchange feature is agentless, works right out of the box, and does not require creating a special lab or running a special backup type.

Refer to the following topics for more information:


- "Starting Object Recovery for Microsoft Exchange" on page 1013
- "Object Recovery Wizard for Microsoft Exchange: Backup" on page 1015
- "Object Recovery Wizard for Microsoft Exchange: Recovery Method" on page 1018
- "Object Recovery Wizard for Microsoft Exchange: Objects" on page 1020
- "Object Recovery Wizard for Microsoft Exchange: Options" on page 1022

# Starting Object Recovery for Microsoft Exchange

You can start the recovery process either from the **Data Protection** menu, by using the search function, or from the **Repositories** tab in **Settings** (for example, if you no longer have a backup job but still have the backup).

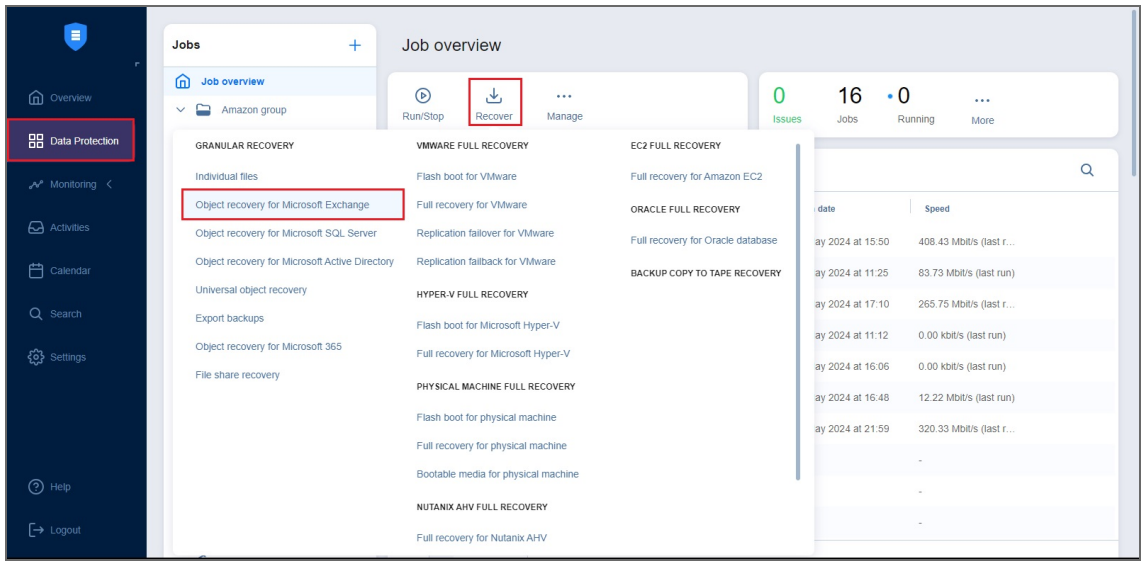
> **Important**
> The recovery process may result in additional load and memory usage on the target server. Therefore, make sure that the server has enough memory.

Refer to the following sections to learn how to start the object recovery process for Microsoft Exchange:

- Starting Object Recovery for Microsoft Exchange from Data Protection Menu
- Starting Object Recovery for Microsoft Exchange from Backup Repository

## Starting Object Recovery for Microsoft Exchange from Data Protection Menu

To start object recovery for Microsoft Exchange from the **Data Protection** menu , click **Recover** and then click **Object recovery for Microsoft Exchange**.

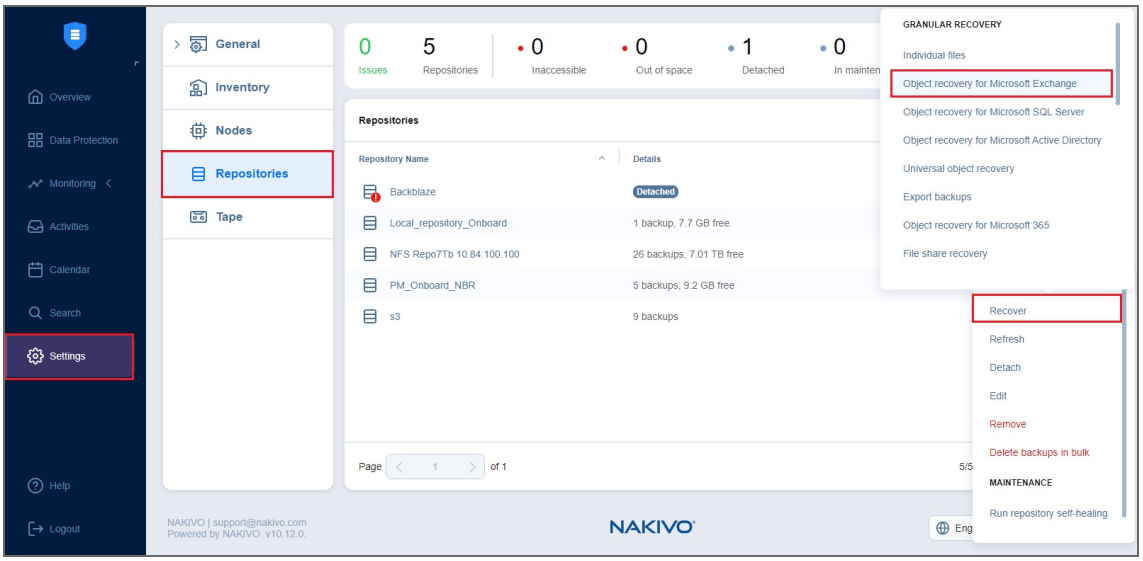



## Starting Object Recovery for Microsoft Exchange from Backup Repository

To start object recovery for Microsoft Exchange from a **Backup Repository**, do the following:

1. Click **Settings** in the main menu of the product.
2. Go to the **Repositories** tab and hover over the **Backup Repository** containing the required backup.

3. Click the ellipsis **Manage** button, click **Recover**, and select **Object recovery for Microsoft Exchange**.




The **Object Recovery Wizard for Microsoft Exchange** opens.

# Object Recovery Wizard for Microsoft Exchange: Backup

1. On the **Backup** page of the wizard, select a backup using either a **Backup Repositories** or **Jobs & Groups** view in the left pane.

Both federated and standalone (not used as members of federated repositories) backup repositories can be selected.

**Notes**

- You can recover individual objects, such as emails from Microsoft Exchange to a custom destination directly from the Proxmox VE VM backups.
- You cannot select a federated repository member as a source for recovery jobs.
- You cannot perform recovery from recovery points located in an inaccessible federated repository member or those that depend on inaccessible recovery points.
- You can search for a specific backup, job, group, or repository (depending on the selected view) by entering its name into the **Search** field.
- Selecting a single backup object opens the **Calendar** view of the date picker if this view was the last user-selected view. Selecting a single backup object opens the **Table** view if this was the last view the user selected.
- You can select a backup object with the corrupted, inaccessible, or pending removal recovery points, but you cannot select the backup objects of the detached repository.
- By default, NAKIVO Backup & Replication automatically searches the selected recovery point for Microsoft Exchange databases (files with `.edb` extension) from which application objects can be recovered. This process can take a few minutes. If you want to manually specify the location of the database file, deselect the **Automatically locate application databases** option.

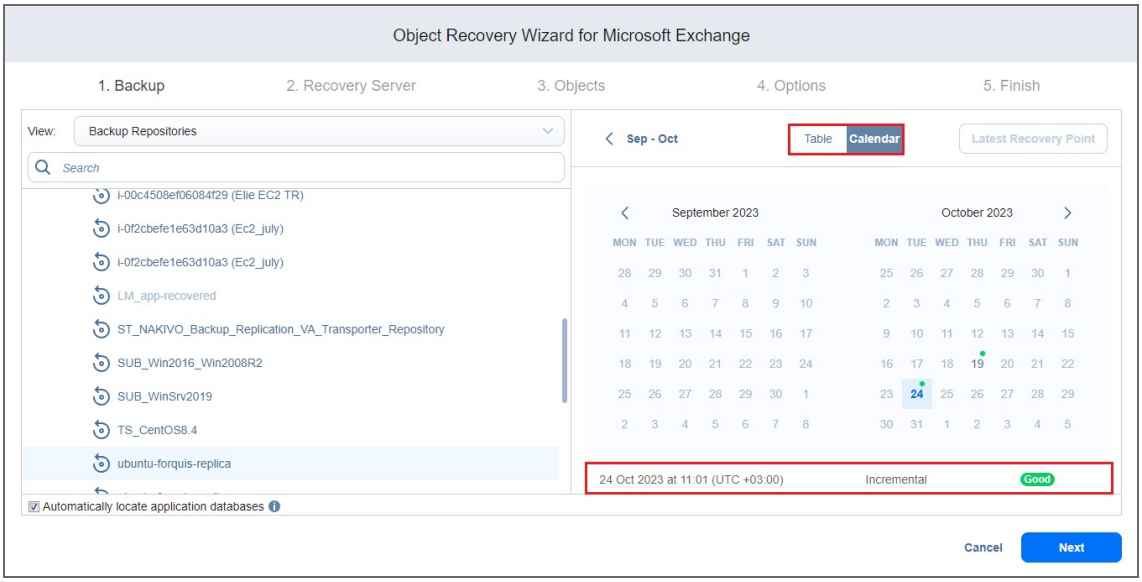
2. The latest recovery point is selected by default.

**Notes**

- The selected date is highlighted.
- If a recovery point is selected in the **Calendar** or **Table** view, this recovery point is selected by default the next time you open the **Calendar** or **Table** view.
- The selected view, either **Calendar** or **Table**, is saved on a per-user basis.

3. If necessary, toggle between the **Calendar** and **Table** views:

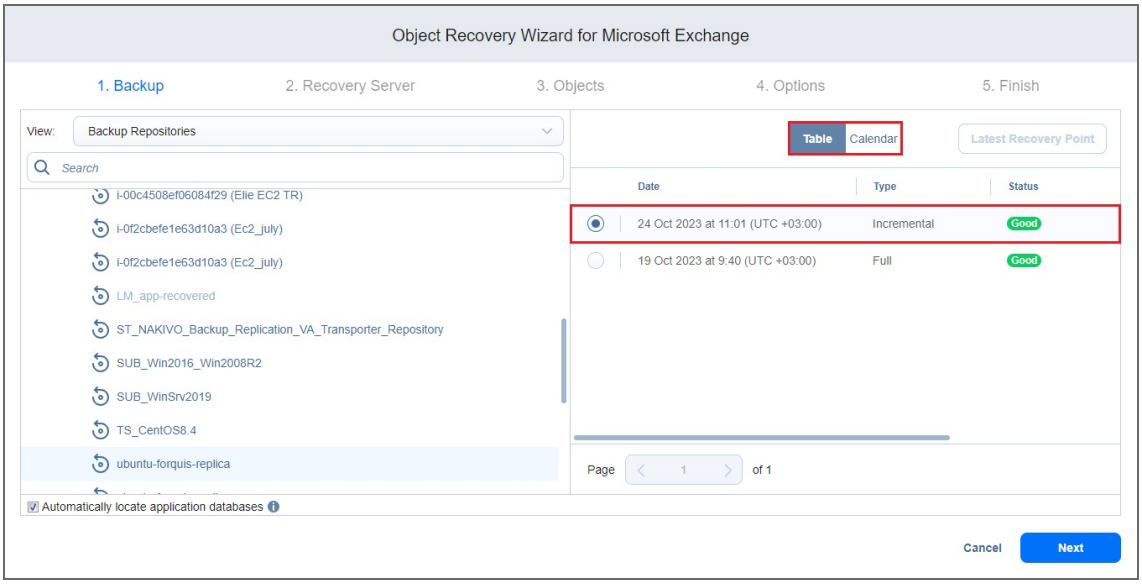
- In the **Calendar** view, select a date or click **Latest Recovery Point** to select the date with the latest recovery point.




**Notes**

- The **Latest Recovery Point** option is disabled if:
  - The latest recovery point is currently selected.
  - All the recovery points of the backup object are corrupted, inaccessible, or pending removal.
- A small dot is displayed on top of the date if there is at least one recovery point created on that date.
- The selected recovery point has a preview displayed under the calendar showing the recovery point status:
  - **green** – verified recovery points
  - **red** – inaccessible/corrupted/pending removal recovery points
  - **dark grey** – unverified recovery points
- Corrupted, inaccessible, or pending removal recovery points cannot be recovered.

- In the **Table** view, select a date or click **Latest Recovery Point** to select the date with the latest recovery point.




**Note**

Clicking the **Latest Recovery Point** option selects the date with the latest recovery point that is not corrupted, inaccessible, or pending removal.

4. Click **Next** to go to the next page of the wizard.

To learn about the limitation on the chain of incremental recovery points, refer to the Knowledge Base article.

# Object Recovery Wizard for Microsoft Exchange: Recovery Method

On the **Recovery Method** page, select the Exchange Server you want to recover to and provide authentication information:

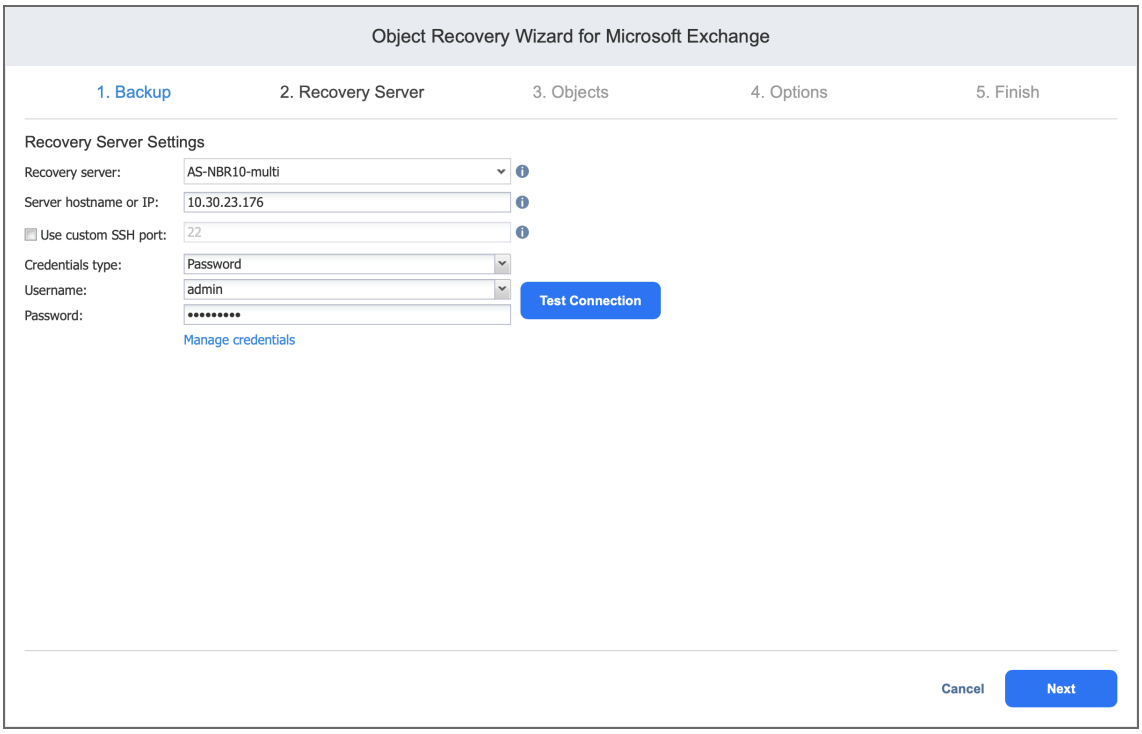
- **Recovery server**: From the drop-down list, select the Exchange Server instance to which the objects must be recovered. The original VM is selected by default. The selection functionality lets you switch views to display the platform where the required VM resides: VMware vSphere, Microsoft Hyper-V, Amazon EC2, Nutanix AHV, or a physical machine. You can also search for the VM by its name. You can skip this parameter altogether and enter the VM's IP address manually in the next field.

  **Notes**

  - Selecting a different recovery server or entering the IP address of a different server may be blocked in case the user has insufficient permissions.
  - Servers added using **Direct Connect** are not supported.

- **Server IP address**: Displays the automatically detected IP address of the server to which the objects must be recovered. You'll need to enter the IP address of the recovery server manually if autodetection fails or if you did not select anything in the previous parameter.
- **Use custom SSH port** (for Linux objects only): Put a checkmark and enter the port number to be used for SSH connection. When the **Use custom SSH port** checkbox is not checked, the default value is used for SSH connections.
- **Credentials type**: Choose your preferred option and enter your respective credentials:
  a. **Password**: Enter a username with administrative privileges for the file share entered above and your password.
  b. **Private key**: Select your private key from the drop-down list.
- **Test Connection**: Click this button to verify the credentials. You won't be able to proceed until after the connection has been successfully established.

- Click **Next** to proceed to the next step.




**Info**

To download items to a browser or forward them via email, enable the **system.exchange.enable.direct.recovery** setting in the Expert tab. Note that contacts and calendar items will not be recoverable with this enabled setting.

# Object Recovery Wizard for Microsoft Exchange: Objects

On the **Objects** page of the wizard, select Microsoft Exchange objects for recovery. Proceed as described in the following sections:

- Searching for Microsoft Exchange Objects
- Browsing Microsoft Exchange Objects
- Viewing Microsoft Exchange Objects
- Selecting Microsoft Exchange Objects to Recover

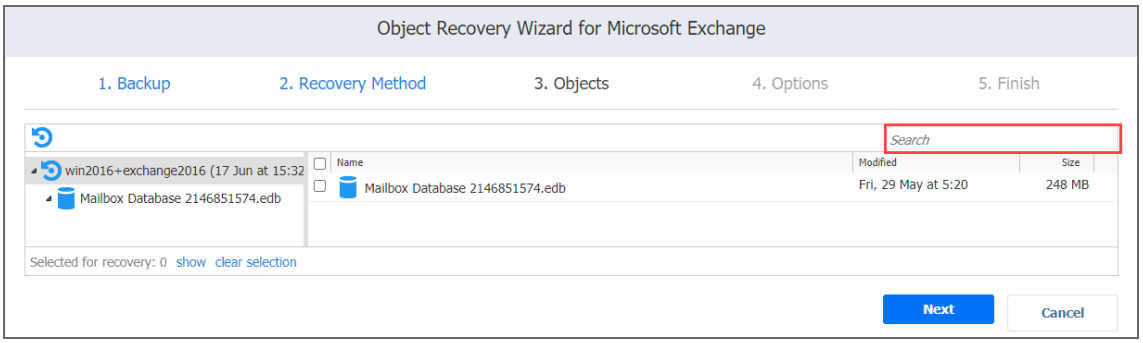
## Searching for Microsoft Exchange Objects

NAKIVO Backup & Replication allows you to search for emails. The search functionality, however, has the following limitations:

- The product can search for emails only by email subject or email body
- If text formatting (such as “bold text”) is applied to a keyword that is searched for, the search may not find the keyword due to formatting conversion issues.
- The product does not create or maintain an index of the Exchange database contents. The search is performed on the fly and can take a long time to complete.

To speed up the search, perform the search within a particular folder, rather than in a mailbox.

To search for an email by its subject or body, type a word in the **Search** field and press **Enter**.




The search is performed starting from the point selected in the left (navigation) pane. For example, if you have selected Mailbox Database > John Smith, the search will be performed only inside the John Smith mailbox.

## Browsing Microsoft Exchange Objects

NAKIVO Backup & Replication scans the selected recovery point for Microsoft Exchange databases (files with the `.edb` extension) and displays the list of found databases in the left pane.

Not all of the found database files contain Microsoft Exchange objects that can be recovered by the product.

To browse Microsoft Exchange objects, expand the appropriate database in the left pane.

## Viewing Microsoft Exchange Objects

To view a Microsoft Exchange object such as an email, click the object. Object contents will be displayed.

> **Note**
> Emails may be blocked from reading in case the user has insufficient permissions.

## Selecting Microsoft Exchange Objects to Recover

In the right pane, select checkboxes next to files and folders you want to recover. The number of items selected for recovery is displayed at the bottom of the wizard page. You can also:

- Click **show** to view the list of all items selected for recovery.
- Click **clear selection** to clear the list of items selected for recovery.
- Click **hide** to hide the list of items selected for recovery.

> **Important**
> For successful recovery of databases, make sure that the Exchange Server license supports the number of databases you plan to recover.

After selecting objects for recovery, click **Next** to go to the next page of the wizard.

# Object Recovery Wizard for Microsoft Exchange: Options

On the **Options** page, specify the location for recovered objects and define overwriting options and naming conventions.

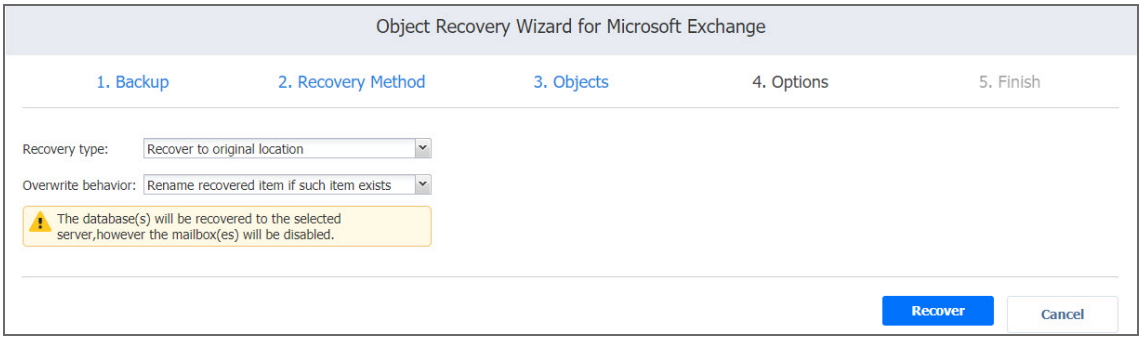
> **Info**
> To download items to a browser or forward them via email, enable the **system.exchange.enable.direct.recovery** setting in the Expert tab. Note that contacts and calendar items will not be recoverable with this enabled setting.

- Recovering to the Original Location
- Recovering to a Custom Location
- Exporting to a Custom Location
- Overwriting Behavior

## Recovering to the Original Location

In the **Recovery type** drop-down list, select **Recover to original location** to recover objects to their original location on the recovery VM.

> **Note**
> Recovering to the original location is not supported when a whole mailbox is selected for recovery.

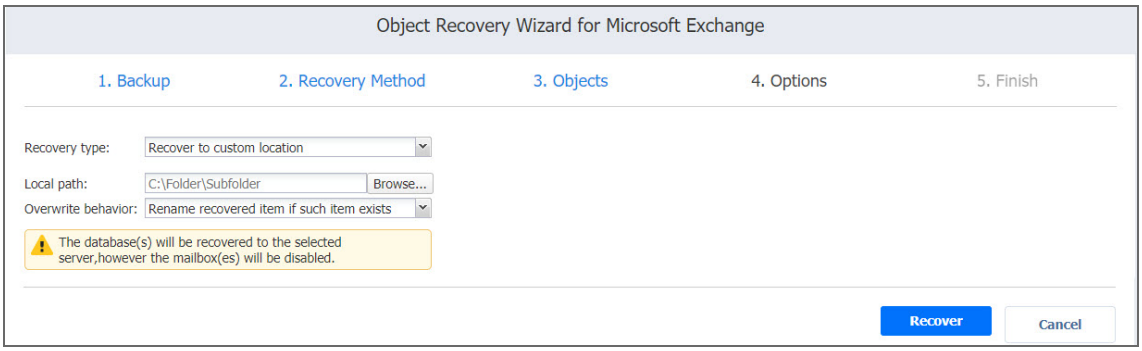



## Recovering to a Custom Location

In the **Recovery type** drop-down list, select **Recover to custom location** to recover objects to a custom location on the recovery VM. Specify the recovery location in the **Local path field** or browse to find it.

**Notes**

- Recovering to a custom location may be blocked in case the user has insufficient permissions.
- Recovering to a custom location is not supported when a whole mailbox is selected for recovery.




## Exporting to a Custom Location

In the **Recovery type** drop-down list, you can choose **Export** to export Microsoft Exchange. You can choose the following locations for the export:

- **Local folder**: After selecting this option, enter the local path to the folder where the recovered objects should be stored.
- **CIFS share**: After selecting this option, provide the path to the file share and enter the necessary credentials.

**Notes**

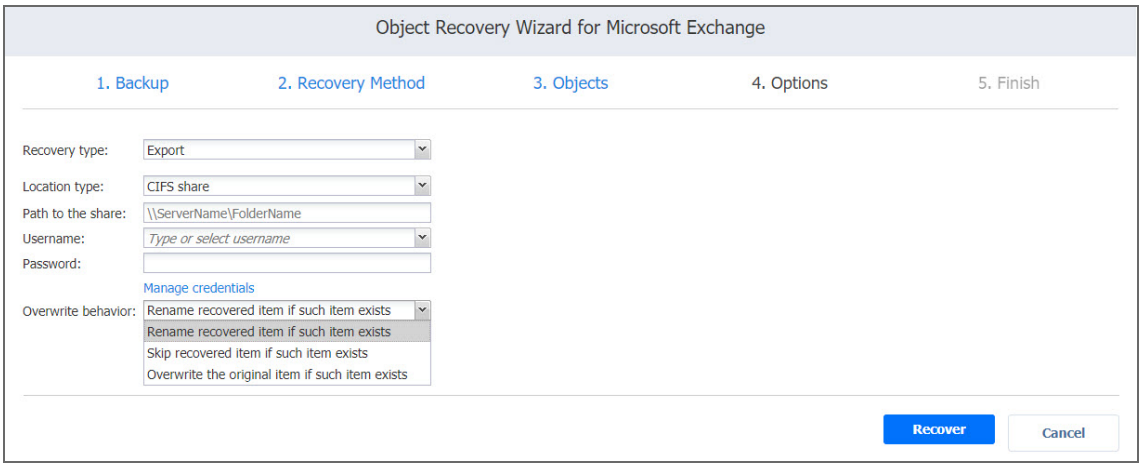
- This option is not supported for databases.
- When this option is selected, some object types are not exported to .pst files:
  - Contacts are exported to .vcf files.
  - Calendar objects are exported to .vcs files.
  - Emails are exported to .eml files.

## Overwriting Behavior

Specify the naming convention for the recovered folders by choosing one of the following options from the **Overwrite behavior** drop-down list:

- **Rename recovered item if such an item exists**
- **Skip recovered item if such an item exists**

- **Overwrite the original item if such an item exists**




Click **Recover** to proceed with the recovery process. The **Finish** page is displayed. You cannot return to the previous pages of the wizard at this point, however, you can check the progress of the job execution by clicking the **Activities** link.

# Object Recovery for Microsoft Active Directory

The instant object recovery feature allows you to browse, search and recover Microsoft Active Directory objects directly from compressed and deduplicated backups. This feature is agentless, works right out of the box, and does not require you to create a special lab or run a special type of backup. Microsoft Active Directory objects can be recovered in `.ldif` format and then be imported back to the Active Directory Server.

Refer to the following topics for more information:


- "Starting Object Recovery for Microsoft Active Directory" on page 1026
- "Object Recovery Wizard for Microsoft AD Server: Backup" on page 1028
- "Object Recovery Wizard for Microsoft AD Server: Recovery Server" on page 1031
- "Object Recovery Wizard for Microsoft AD Server: Objects" on page 1033
- "Object Recovery Wizard for Microsoft AD Server: Options" on page 1036

# Starting Object Recovery for Microsoft Active Directory

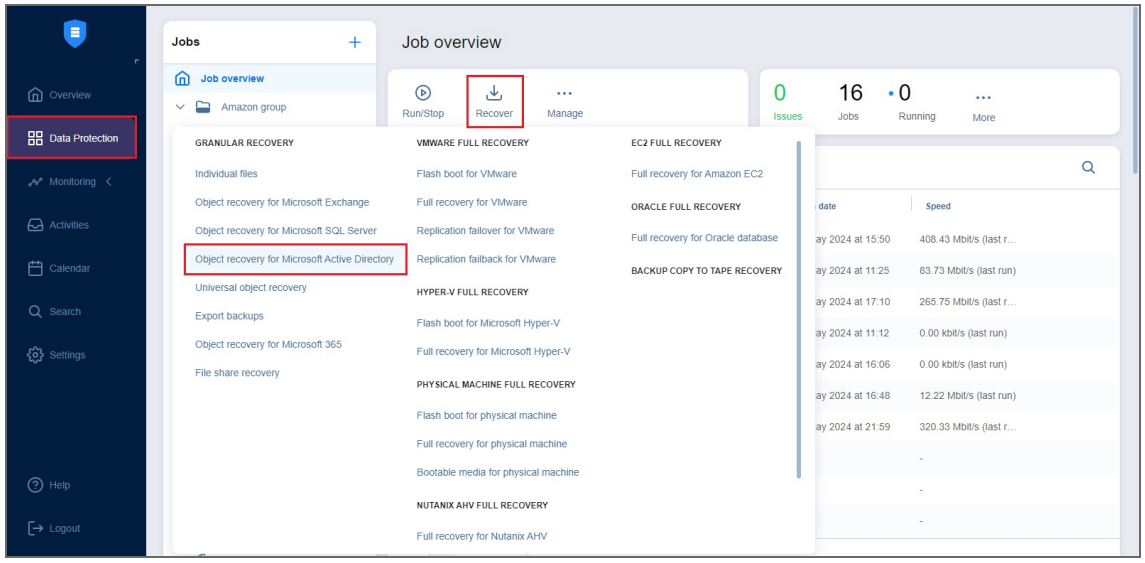
You can start the recovery process either from the **Data Protection** menu, by using the Search function, or from the **Repositories** tab in the **Settings** (for example, if you no longer have a backup job, but still have the backup). Refer to the following sections for more details:

- Starting Active Directory Object Recovery from Data Protection Menu
- Starting Active Directory Object Recovery from a Backup Repository

> **Note**
>
> During the Active Directory object recovery process, the system shall generate a new objectSID and assign it to the recovered object. Due to Windows limitations, once the recovered object has a new SID, it is not possible to change it back to the original SID.

## Starting Active Directory Object Recovery from Data Protection Menu

To start Active Directory Object Recovery from the **Data Protection** menu, click **Recover** and then choose **Object recovery for Microsoft Active Directory**.

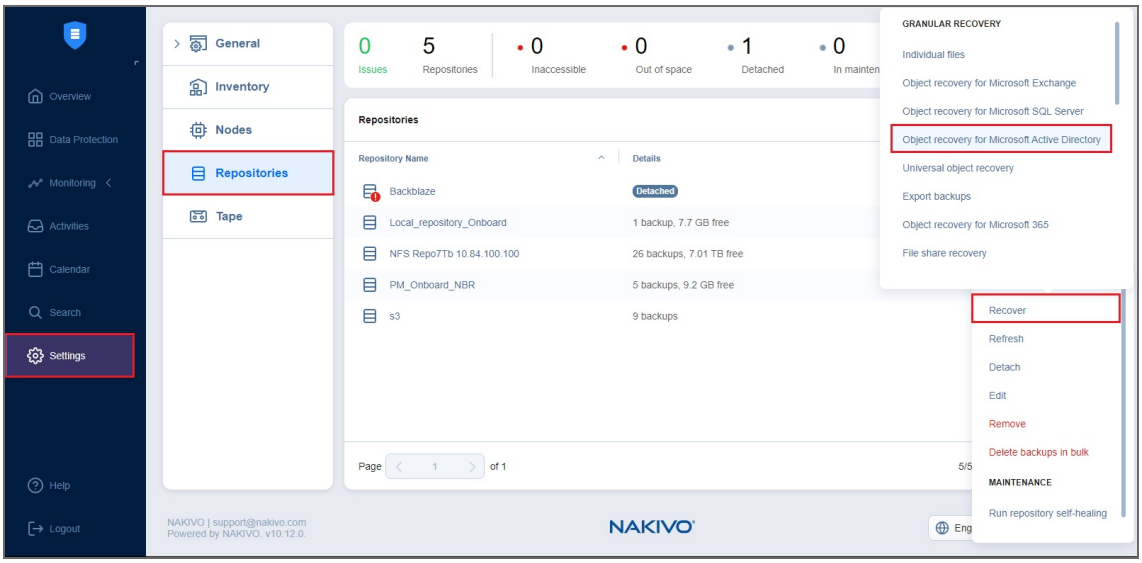



## Starting Active Directory Object Recovery from a Backup Repository

To start Active Directory Object Recovery from a **Backup Repository**, do the following:

1. Go to the main menu of the product and click **Settings**.
2. Go to the **Repositories** tab and hover over the **Backup Repository** containing the required backup.

3. Click the ellipsis **Manage** button, click **Recover**, and select **Microsoft Active Directory objects**.




The **Object Recovery Job Wizard for Microsoft AD Server** opens.

# Object Recovery Wizard for Microsoft AD Server: Backup

1. On the **Backup** page of the wizard, select a backup of a VM with the Microsoft Active Directory server using either the **Backup Repositories** or **Jobs & Groups** view in the left pane.
Both federated and standalone (not used as members of federated repositories) backup repositories can be selected.

**Notes**

- You can recover individual objects, such as entries from Microsoft Active Directory to a custom destination directly from the Proxmox VE VM backups.
- You cannot select a federated repository member as a source for recovery jobs.
- You can search for a specific backup, job, group, or repository (depending on the selected view) by entering its name into the **Search** field.
- Selecting a single backup object opens the **Calendar** view of the date picker if this view was the last user-selected view. Selecting a single backup object opens the **Table** view if this was the last view the user selected.
- You can select a backup object with the corrupted, inaccessible, or pending removal recovery points, but you cannot select the backup objects of the detached repository.
- You cannot perform recovery from recovery points located in an inaccessible federated repository member or those that depend on inaccessible recovery points.
- By default, NAKIVO Backup & Replication automatically searches the selected recovery point for Microsoft Active Directory database from which application objects can be recovered. This process can take a few minutes. If you want to manually specify the location of the database file, deselect the **Automatically locate application databases** option.

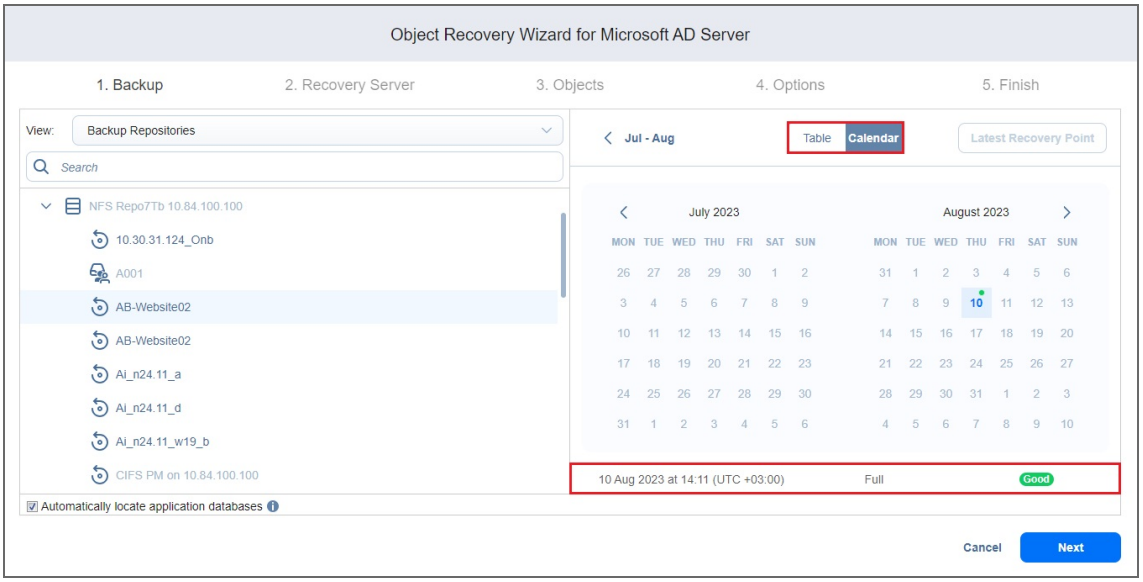
2. The latest recovery point is selected by default.

**Notes**

- The selected date is highlighted.
- If a recovery point is selected in the **Calendar** or **Table** view, this recovery point is selected by default the next time you open the **Calendar** or **Table** view.
- The selected view, either **Calendar** or **Table**, is saved on a per-user basis.

3. If necessary, toggle between the **Calendar** and **Table** views:

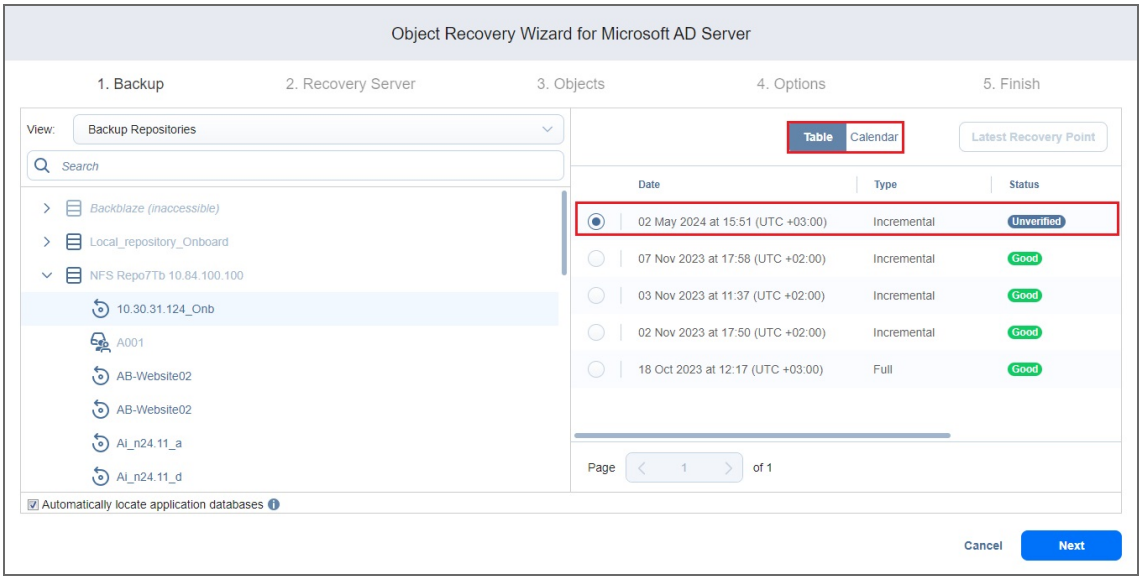
- In the **Calendar** view, select a date or click **Latest Recovery Point** to select the date with the latest recovery point.




**Notes**

- The **Latest Recovery Point** option is disabled if:
  - The latest recovery point is currently selected.
  - All the recovery points of the backup object are corrupted, inaccessible, or pending removal.
- A small dot is displayed on top of the date if there is at least one recovery point created on that date.
- The selected recovery point has a preview displayed under the calendar showing the recovery point status:
- **green** – verified recovery points
- **red** – inaccessible/corrupted/pending removal recovery points
- **dark grey** – unverified recovery points
- Corrupted, inaccessible, or pending removal recovery points cannot be recovered.

- In the **Table** view, select a date or click **Latest Recovery Point** to select the date with the latest recovery point.




**Note**

Clicking the **Latest Recovery Point** option selects the date with the latest recovery point that is not corrupted, inaccessible, or pending removal.

4. Click **Next** to go to the next page of the wizard.

To learn about the limitation on the chain of incremental recovery points, refer to the Knowledge Base article.

# Object Recovery Wizard for Microsoft AD Server: Recovery Server

On the **Recovery Server** page of the wizard, set up a Microsoft Active Directory server to which objects will be recovered.

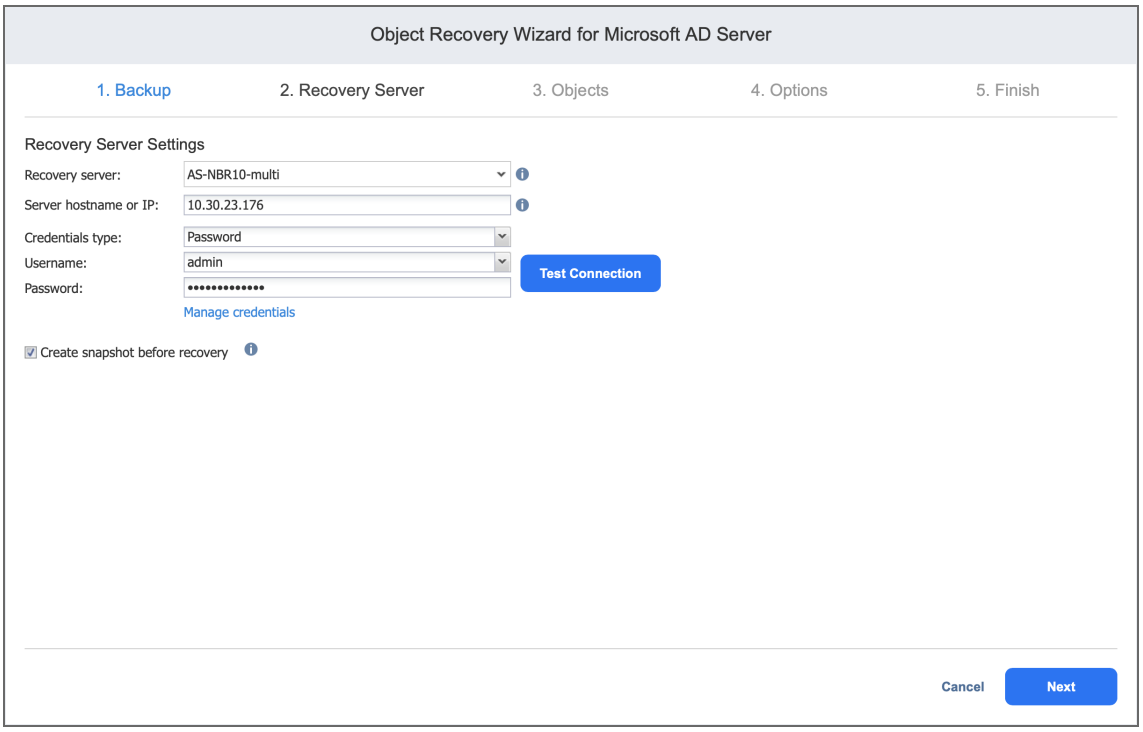
**Important**

- The ISCSI Initiator service must be running on the recovery server.
- The `vc_redist.x86.exe` (v.2015) file must be installed on the recovery server. Refer to the Microsoft article for installation details.
- Servers added using Direct Connect are not supported.

Set up a Microsoft Active Directory server the following way:

- In the **Recovery Server** drop-down list, select a recovery server name.
- In the **Server IP address** box, enter the IP address of the recovery server. This is necessary if the application has not detected the IP address based on the recovery server name.
- **Credentials type**: Choose your preferred option and enter your respective credentials:
  a. **Password**: Enter a username with administrative privileges for the file share entered above and your password.
  b. **Private key**: Select your private key from the drop-down list.
- Click the **Test Connection** button to test your credentials for the specified recovery server. If your credentials are correct, a checkmark appears to the right of the button.
- **Create snapshot before recovery:** When selected, a snapshot of the VM will be taken if recovery fails, and the VM will be reverted to this snapshot.

- Click **Next** to go to the next page of the wizard.



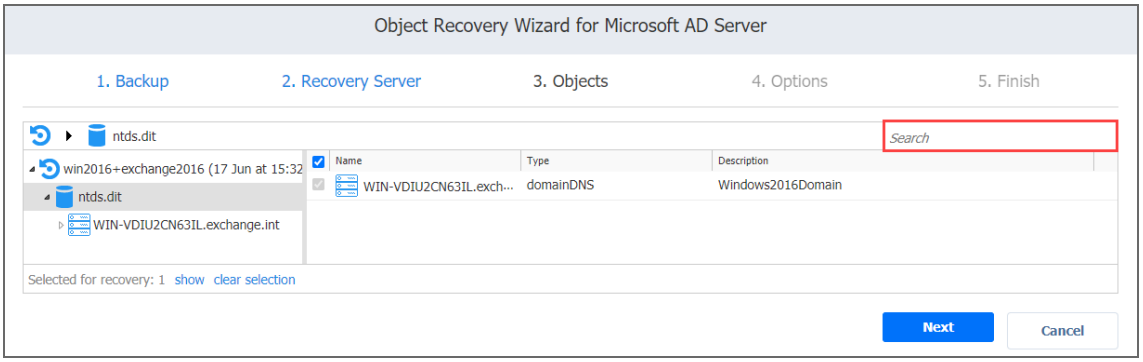
# Object Recovery Wizard for Microsoft AD Server: Objects

On the **Objects** page of the wizard, select Active Directory objects you want to recover.

- Searching for Active Directory Objects
- Browsing Active Directory Objects
- Viewing Active Directory Objects
- Selecting Active Directory Objects to Recover

## Searching for Active Directory Objects

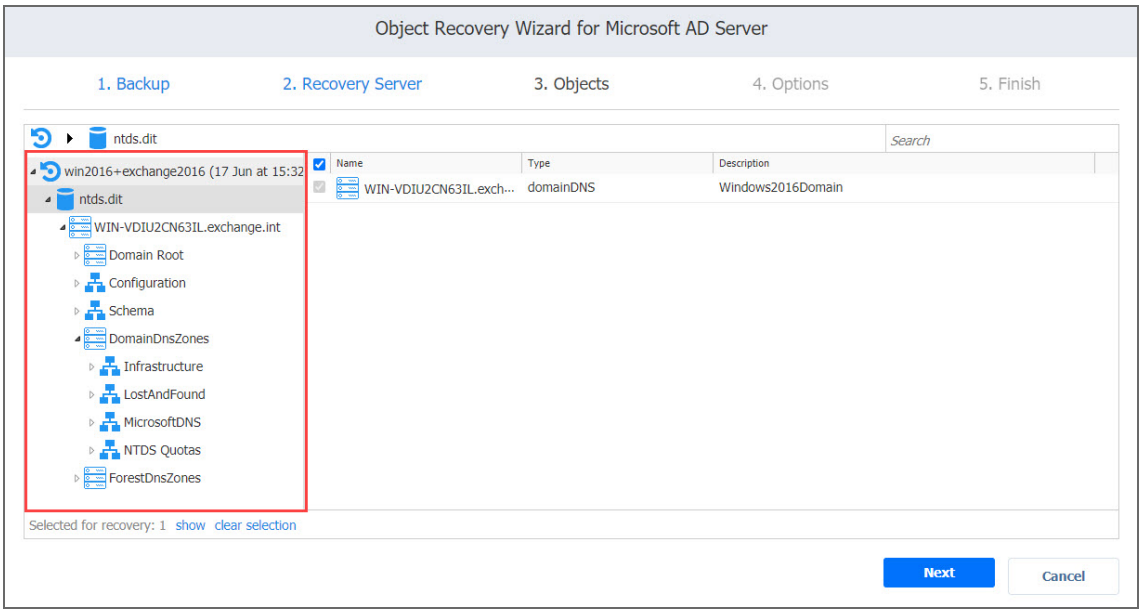
NAKIVO Backup & Replication allows you to search Active Directory objects by name. To find an object by its name, enter a word in the **Search** box and press **Enter**




The search is performed starting from the point selected in the left (navigation) pane. For example, if you have selected the **Users** group, the search will only be performed inside the **Users** group.

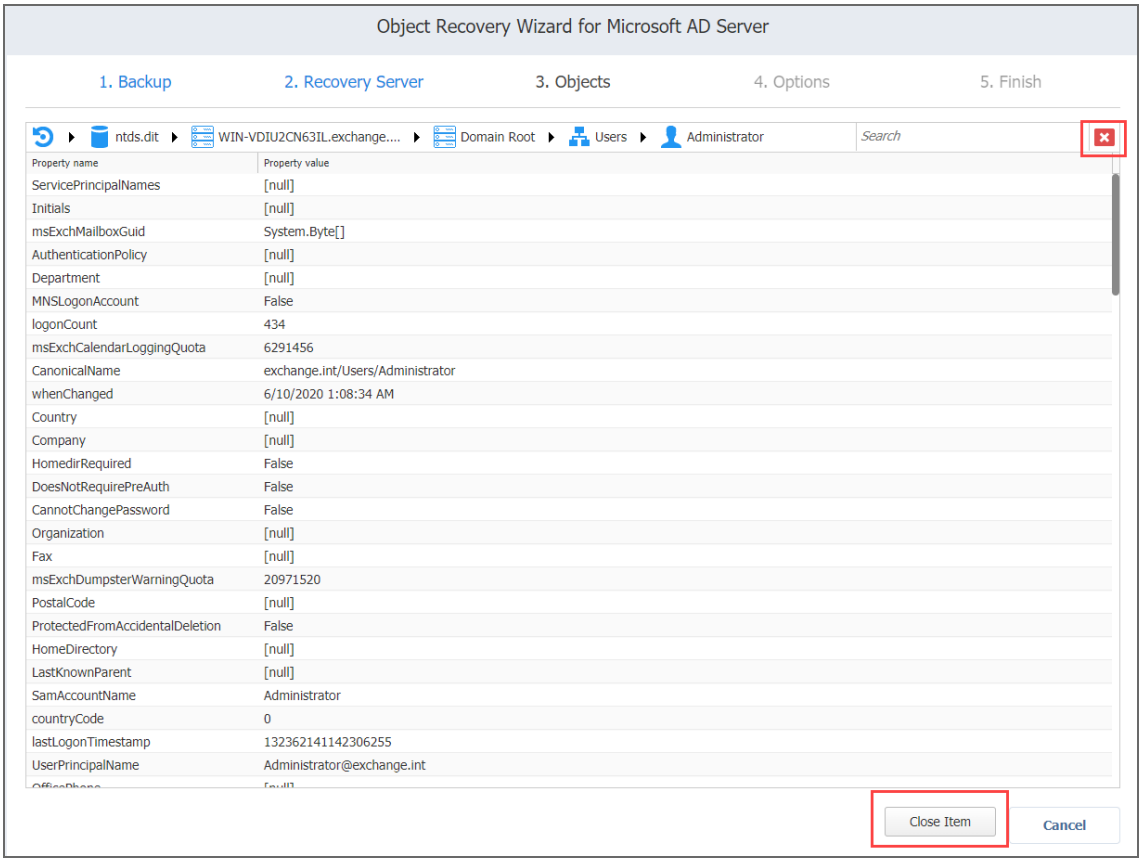
## Browsing Active Directory Objects

NAKIVO Backup & Replication scans the selected recovery point for Active Directory databases (files with ".edb" extension) and displays the list of identified databases in the left (navigation) pane. To browse Microsoft Active Directory objects, simply expand the appropriate database in the left pane. You can also browse the tree by using the scroll bar.




## Viewing Active Directory Objects

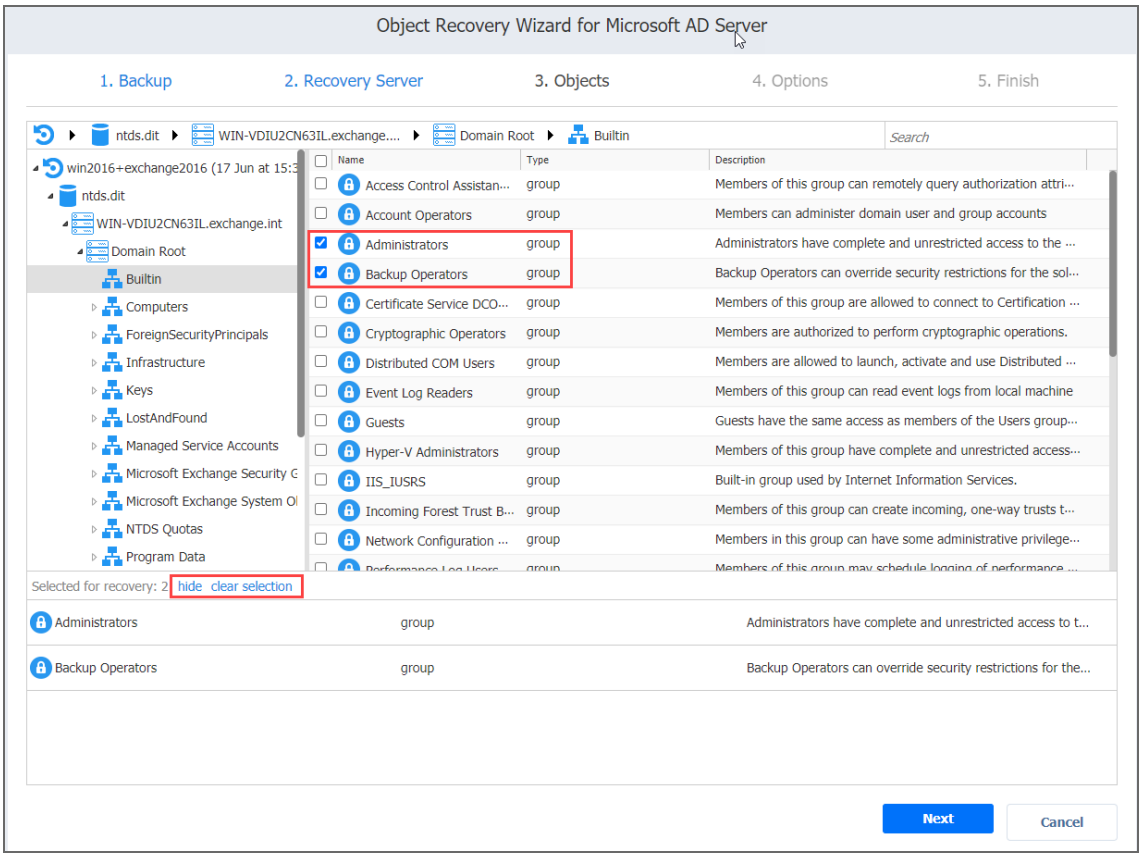
To view a Microsoft Active Directory object, click the object. The object contents will be displayed. Use the close buttons to close the item.



# Selecting Active Directory Objects to Recover

In the **Contents** pane to the right, select a checkbox next to the items you want to recover. The number of items selected for recovery is displayed at the bottom of the wizard page. You can also:

- Click **show** to view the list of all items selected for recovery.
- Click **clear selection** to clear the list of items selected for recovery.
- Click **hide** to hide the list of items selected for recovery.




When ready with selecting Microsoft Active Directory objects for recovery, click **Next** to go to the next page of the wizard

# Object Recovery Wizard for Microsoft AD Server: Options

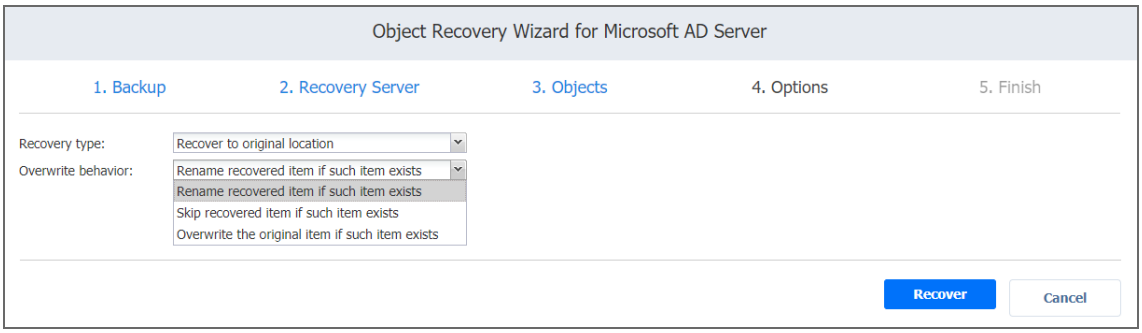
On the **Options** page of the wizard, you can set up the following options for your object recovery job:

- Recovering Objects to the Original Location
- Exporting Active Directory Objects

## Recovering Objects to the Original Location

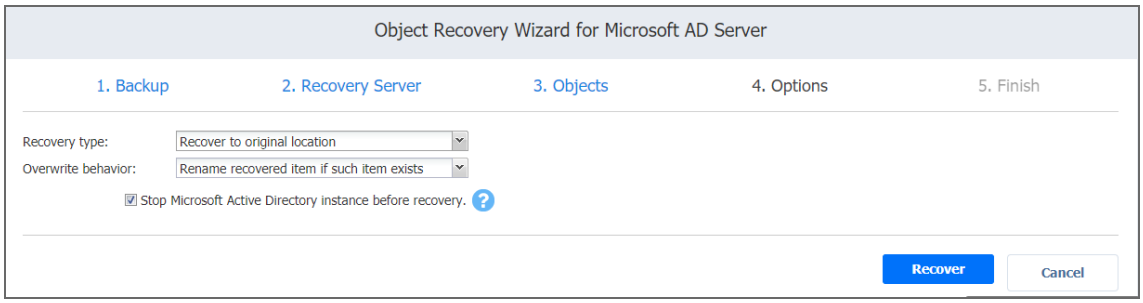
Follow the steps below to recover objects of your Microsoft Active Directory server to the original location:

1. In the **Recovery type** list, select **Recover to original location**.
2. If you have selected multiple objects or container(s) that include one or more “user” objects, the **Recover of user object** list becomes available. Select either of the following options:
   - **User will be disabled**: If this option is selected, NAKIVO Backup & Replication disables all recovered “user” objects and the corresponding user accounts are disabled after importing these objects to Active Directory.
   - **User must change password at next log on**: If this option is selected, NAKIVO Backup & Replication generates a new password for each recovered “user” object. The `passwords.txt` file is added to the ZIP archive with recovered objects, and it contains the new passwords. After importing the “user” objects to Active Directory, corresponding users are forced to change the password on the next log on.




3. In the **Overwrite behavior** list, select what you wish to do if the recovered item conflicts with an existing one:
   - **Rename recovered item if such item exists**
   - **Skip recovered item if such item exists**
   - **Overwrite the original item if such item exists**
4. If you have chosen to recover a full database (`ntds.dit` file) on the **Recovery Server** pageput a check-mark in the**Stop Microsoft Active Directory instance before recovery** checkbox to stop the instance before the recovery process begins.This option is recommended for the safe recovery of Active Dir-

ectory objects.




5. Click **Recover**.

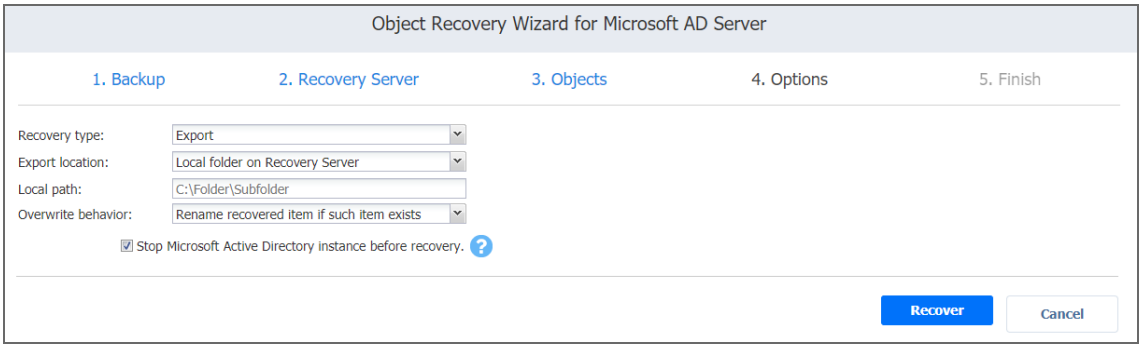
**Notes**

- Some attributes may be skipped for the selected object(s) depending on the Active Directory system settings.
- In case the recovery process fails, theVM will be reverted to the snapshot taken on the **Recovery Server** page of the wizard

## Exporting Active Directory Objects

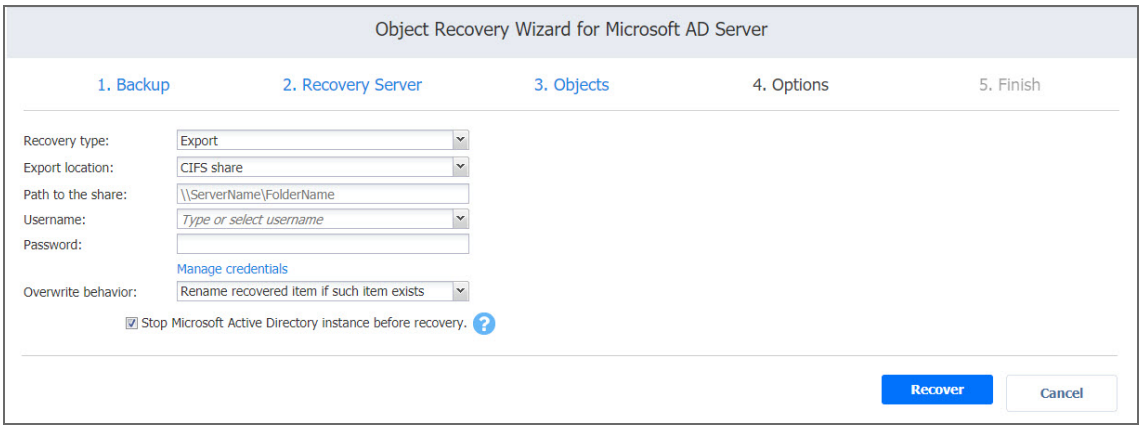
Follow the steps below to export recovered objects of your Microsoft Active Directory server to a custom location:

1. In the **Recovery type** list, select **Export**. A number of options become available for setting up a custom location.
2. In the **Export location** list, select the appropriate location type:
   - **Local folder on Recovery Server:** If this option is selected, you will have to enter the path to a local folder on the recovery server in the **Local path** field.




   - **CIFS share**: If this option is selected, enter the following values:
     1. **Path to the share**
     2. **Username**

3. **Password**




3. In the **Overwrite behavior** list, select what needs to be done if the recovered item conflicts with an existing item. Refer to the section above for an explanation.
4. Click **Recover**.

The **Finish** page of the wizard opens informing you that Microsoft Active Directory object recovery has started. To view the object recovery progress, open the **Activities** tab.

To close the wizard, click **Close**.

# Importing Recovered Objects to Active Directory

Refer to the sections below for information on how to import recovered objects in Active Directory.

- Importing Non-User Objects
- Importing User Objects

## Importing Non-User Objects

If Active Directory objects or containers that you have recovered do not contain "User" objects, follow the steps below to import the objects in Active Directory:

1. On the Active Directory machine, run command line as an administrator.
2. Run the following command: `ldifde -i -k -f filename -j logfolder`, where "`filename.ldif`" is the path to the recovered `ldif` file, and "`logfolder`" is the path to the folder where import logs will be saved.

## Importing User Objects

If you have recovered one or more "User" objects or if you have recovered containers that include one or more "User" objects, follow the steps below to import the objects in Active Directory:

1. On the Active Directory machine, run command line as an administrator.
2. Enable a secure LDAP connection on the Active Directory machine:
    a. Log on to the server and open the Server Manager tool.
    b. Add the Active Directory Certificate Services role. On the **Role services** page of the **Add Roles and Features** wizard, select a Certification Authority.
    c. When configuring the Active Directory Certificate service on the destination server, use proper credentials to configure the service, choose the **Enterprise CA** setup type, and choose a **Root CA for CA Type**.
    d. Follow the rest of wizard instructions to complete adding the Active Directory Certificate Services role.
3. Run the following command: `ldifde -i -t 636 -f filename.ldif -k -j logfolder`, where "`filename.ldif`" is the path to the recovered `ldif` file, and "`logfolder`" is the path to the folder where import logs will be saved.
4. Edit the group policy by adding imported users. After importing one or more users, you may need to verify password options via user logon.

# Object Recovery for Microsoft SQL Server

The instant object recovery feature in NAKIVO Backup & Replication allows you to browse, search, and recover Microsoft SQL Server objects directly from compressed and deduplicated backups. This out-of-the-box feature is agentless, and it does not require creating a special lab or running a special backup type. Microsoft SQL Server objects can be recovered to a source or another VM.

Refer to the following topics for more information:


- "Starting Object Recovery for Microsoft SQL Server" on page 1041
- "Object Recovery Wizard for Microsoft SQL Server: Backup" on page 1043
- "Object Recovery Wizard for Microsoft SQL Server: Recovery Server" on page 1046
- "Object Recovery Wizard for Microsoft SQL Server: Objects" on page 1048
- " Object Recovery Wizard for Microsoft SQL Server: Options" on page 1049

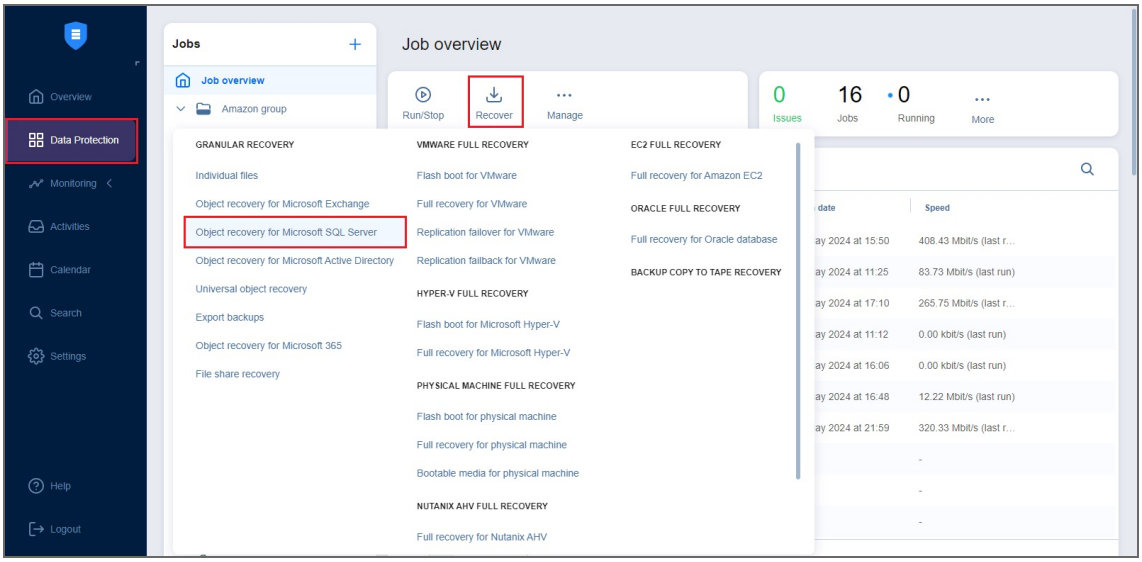
# Starting Object Recovery for Microsoft SQL Server

You can start the recovery process either from the **Data Protection** menu, by using the Search function, or from the **Repositories** tab in the **Settings** (for example, if you no longer have a backup job but still have the backup). Refer to the following sections for more details:

- Starting SQL Server Object Recovery from Data Protection Menu
- Starting SQL Server Object Recovery from Backup Repository

## Starting SQL Server Object Recovery from Data Protection Menu

To start Microsoft SQL Server object recovery from the **Data Protection** menu, click **Recover** and choose **Object recovery for Microsoft SQL Server**.

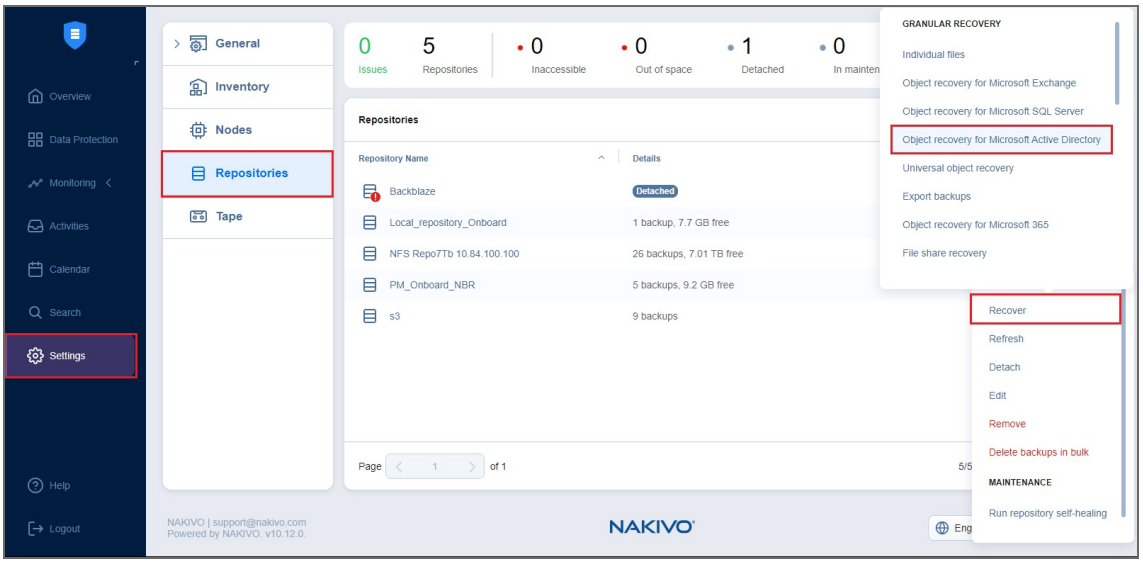



## Starting SQL Server Object Recovery from Backup Repository

To start SQL Server object recovery from a **Backup Repository:**

1. Click **Settings** in the main menu of the product.
2. Go to the **Repositories** tab and hover the cursor over the **Backup Repository** containing the required backup.

3. Click the ellipsis **Manage** button, click **Recover**, and select **Object recovery for Microsoft SQL Server**.




The **New Object Recovery Wizard for Microsoft SQL Server** opens.

# Object Recovery Wizard for Microsoft SQL Server: Backup

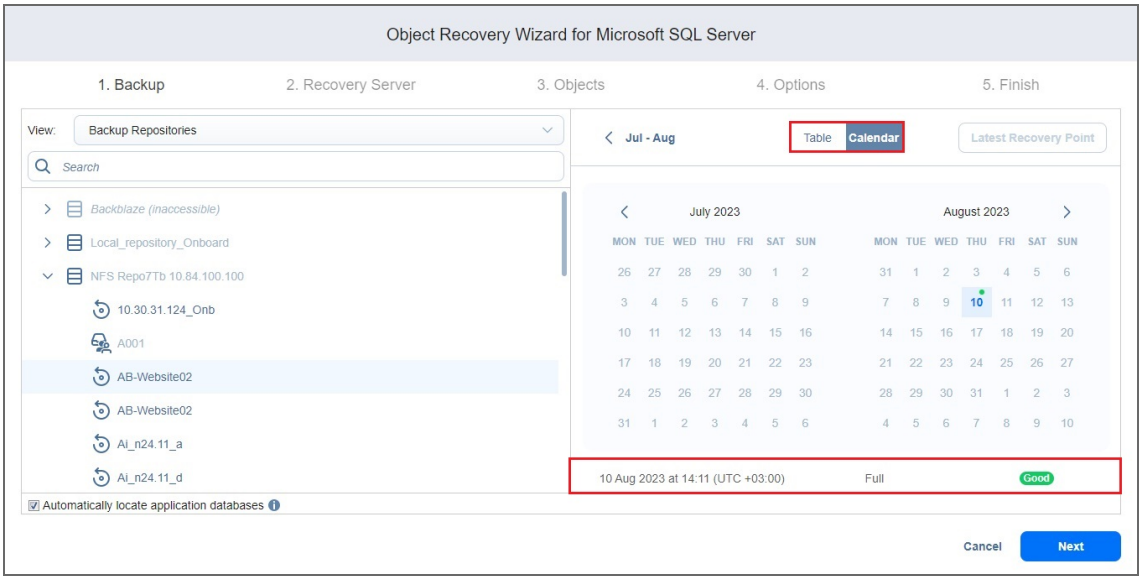
1. On the **Backup** page of the wizard, select a backup of a VM with Microsoft SQL using either the **Backup Repositories** or **Jobs & Groups** view in the left pane.

Both federated and standalone (not used as members of federated repositories) backup repositories can be selected.

**Notes**

- You can recover individual objects, such as data from Microsoft SQL Server to a custom destination directly from the Proxmox VE VM backups.
- You cannot select a federated repository member as a source for recovery jobs.
- You cannot perform recovery from recovery points located in an inaccessible federated repository member or those that depend on inaccessible recovery points.
- You can search for a specific backup, job, group, or repository (depending on the selected view) by entering its name into the **Search** field.
- Selecting a single backup object opens the **Calendar** view of the date picker if this view was the last user-selected view. Selecting a single backup object opens the **Table** view if this was the last view the user selected.
- You can select a backup object with the corrupted, inaccessible, or pending removal recovery points, but you cannot select the backup objects of the detached repository.
- By default, NAKIVO Backup & Replication automatically searches the selected recovery point for Microsoft SQL database from which objects can be recovered. This process can take a few minutes. If you want to manually specify the location of the database file, deselect the **Automatically locate application databases** option.

2. The latest recovery point is selected by default.




**Notes**

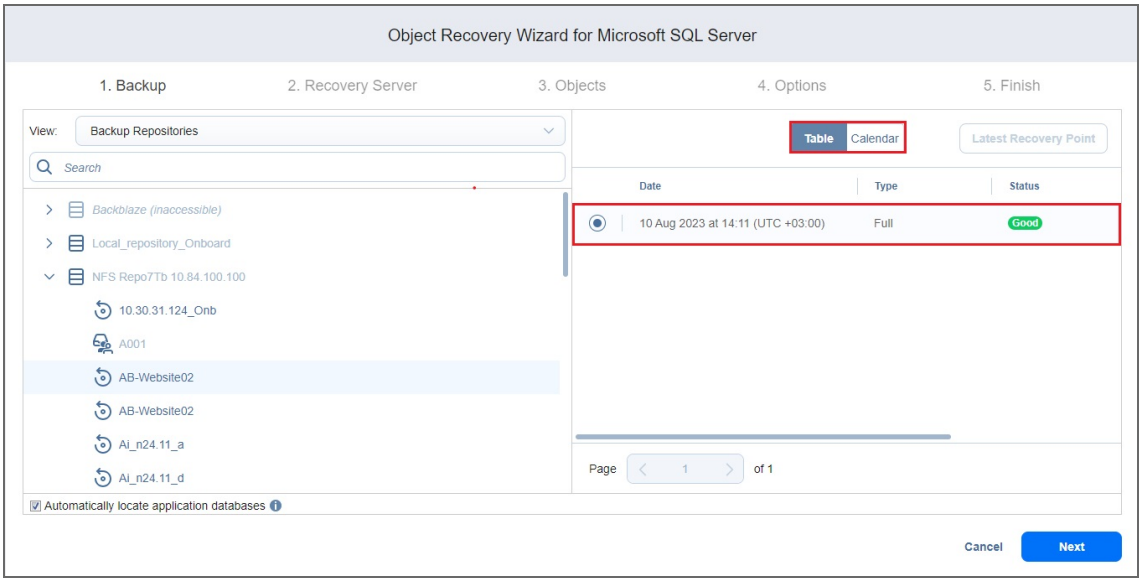
- The selected date is highlighted.
- If a recovery point is selected in the **Calendar** or **Table** view, this recovery point is selected by default the next time you open the **Calendar** or **Table** view.
- The selected view, either **Calendar** or **Table**, is saved on a per-user basis.

3. If necessary, toggle between the **Calendar** and **Table** views:
   - In the **Calendar** view, select a date or click **Latest Recovery Point** to select the date with the latest recovery point.

**Notes**

- The **Latest Recovery Point** option is disabled if:
  - The latest recovery point is currently selected.
  - All the recovery points of the backup object are corrupted, inaccessible, or pending removal.
- A small dot is displayed on top of the date if there is at least one recovery point created on that date.
- The selected recovery point has a preview displayed under the calendar showing the recovery point status:
  - **green** – verified recovery points
  - **red** – inaccessible/corrupted/pending removal recovery points
  - **dark grey** – unverified recovery points
- Corrupted, inaccessible, or pending removal recovery points cannot be recovered.

- In the **Table** view, select a date or click **Latest Recovery Point** to select the date with the latest recovery point.




**Note**

Clicking the **Latest Recovery Point** option selects the date with the latest recovery point that is not corrupted, inaccessible, or pending removal.

4. Click **Next** to go to the next page of the wizard.

To learn about the limitation on the chain of incremental recovery points, refer to the Knowledge Base article.

# Object Recovery Wizard for Microsoft SQL Server: Recovery Server

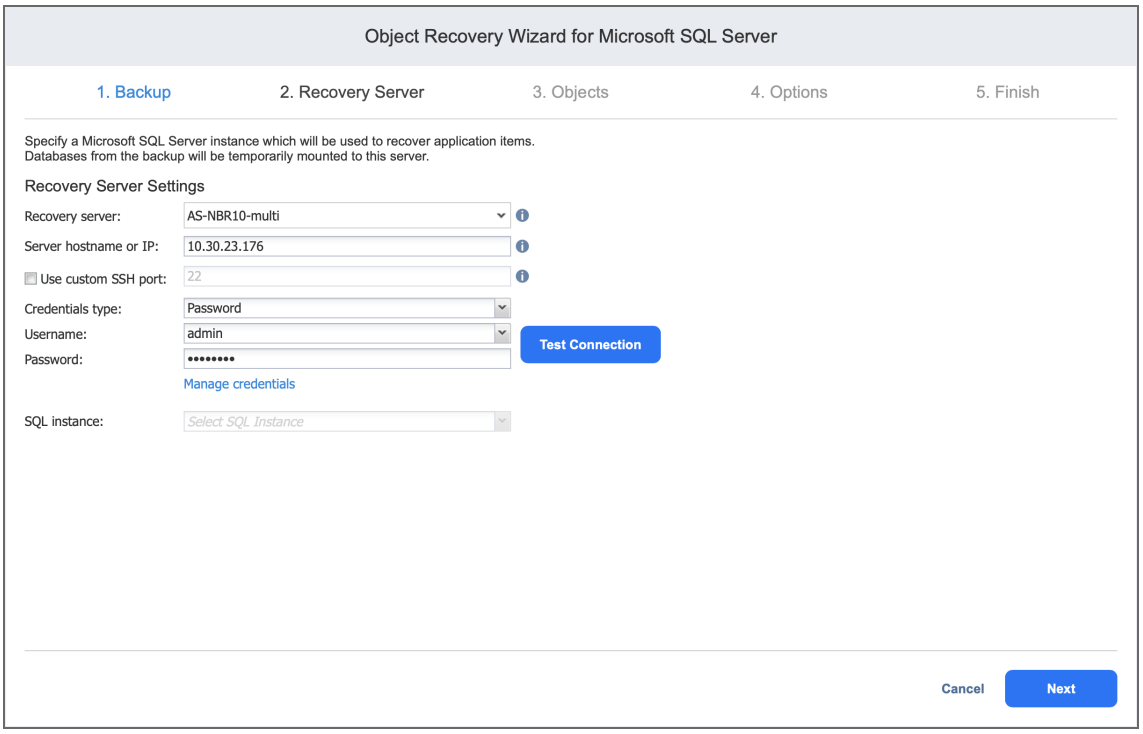
To set up a recovery server for Microsoft SQL Server objects:

1. The **Recovery Server Settings** section opens. Please enter the following values:
   - **Recovery server**: Choose the target server from the drop-down list.

     **Notes**
     - NAKIVO Backup & Replication will try to auto-detect the IP address automatically.
     - Servers added using **Direct Connect** are not supported.

   - **Server IP address**: Enter the IP address of the recovery server if it is not detected by the application based on the recovery server name.
   - **Use custom SSH port**: If necessary, enter an SSH port to be used for connecting to the recovery server. The default value is *22*.
   - **Credentials type**: Choose your preferred option and enter your respective credentials. Refer to "Requirements for Microsoft SQL Server Object Recovery" on page 206 for a full list of requirements for recovering files to server.
     a. **Password**: Enter a username with administrative privileges for the file share entered above and your password.
     b. **Private key**: Select your private key from the drop-down list.
   - **SQL instance**: Select a target SQL instance.
2. Click the **Test Connection** button to test your credentials for the specified recovery server. If your credentials are correct, a checkmark appears to the right of the button.

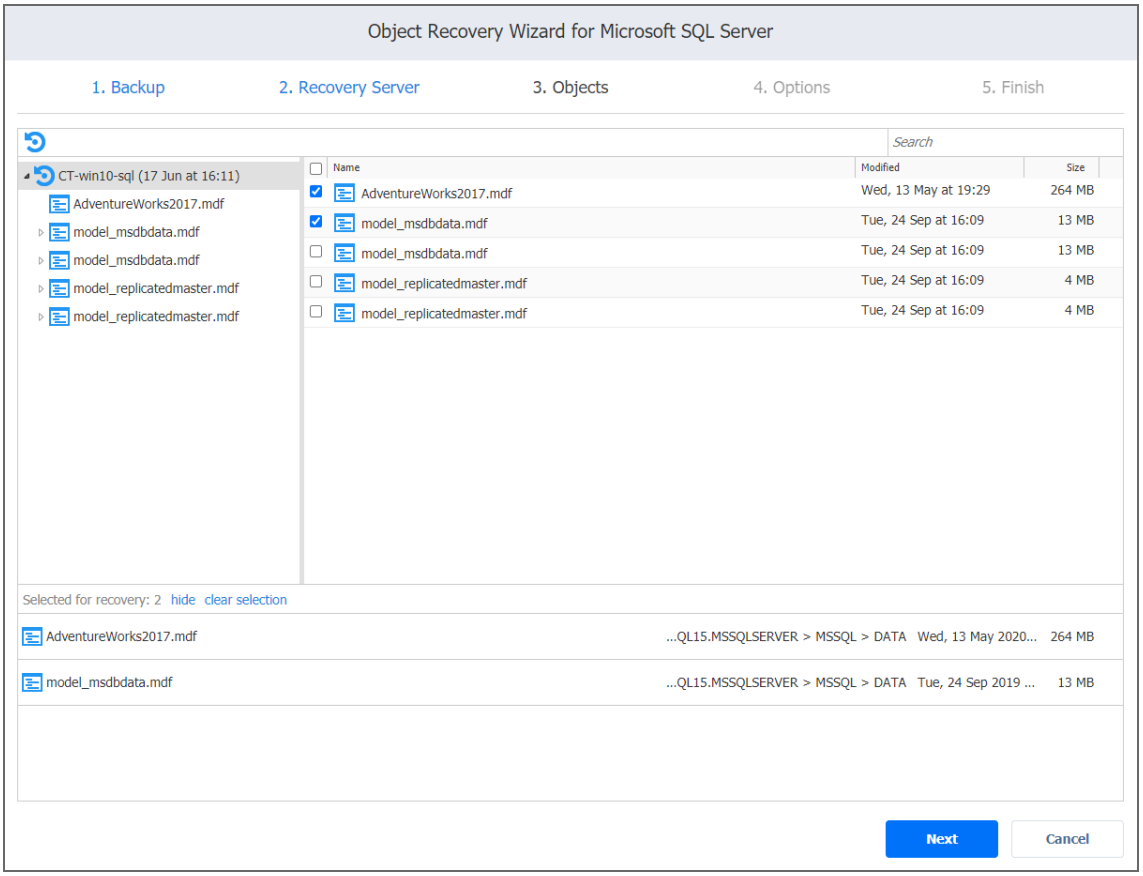
3. Click **Next**.


Object Recovery Wizard for Microsoft SQL Server
1. Backup
2. Recovery Server
3. Objects
4. Options
5. Finish
Specify a Microsoft SQL Server instance which will be used to recover application items.
Databases from the backup will be temporarily mounted to this server.
Recovery Server Settings
Recovery server:
AS-NBR10-multi
Server hostname or IP:
10.30.23.176
Use custom SSH port:
22
Credentials type:
Password
Username:
admin
Password:
Test Connection
Manage credentials
SQL instance:
Select SQL Instance
Cancel
Next

# Object Recovery Wizard for Microsoft SQL Server: Objects

On the **Objects** page of the wizard, select objects for recovery. You can select either entire databases or individual objects for recovery.

1. Select the database in the left pane.
2. Select the objects in the right pane. If you want to restore an entire database, select all objects in this pane.
3. Alternatively, you can look for objects using the **Search** bar.
4. When you are done, click **Next**.



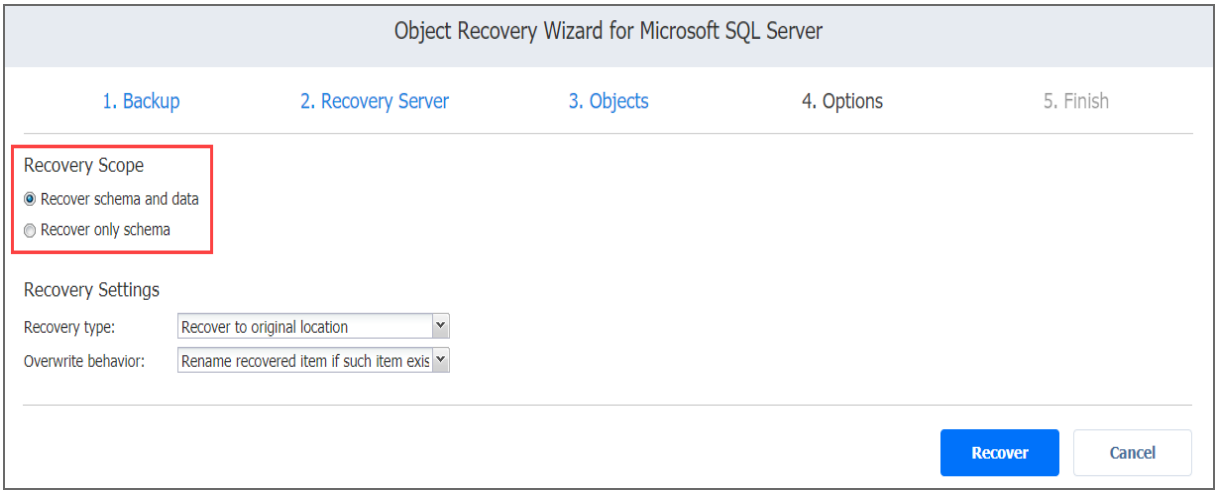
# Object Recovery Wizard for Microsoft SQL Server: Options

On the **Options** page of the wizard, set the options for the recovery job.

- Recovery Scope
- Recovery Settings
- Overwrite Behavior

## Recovery Scope

Set the recovery scope by selecting either **Recover schema and data** or **Recover only schema**.



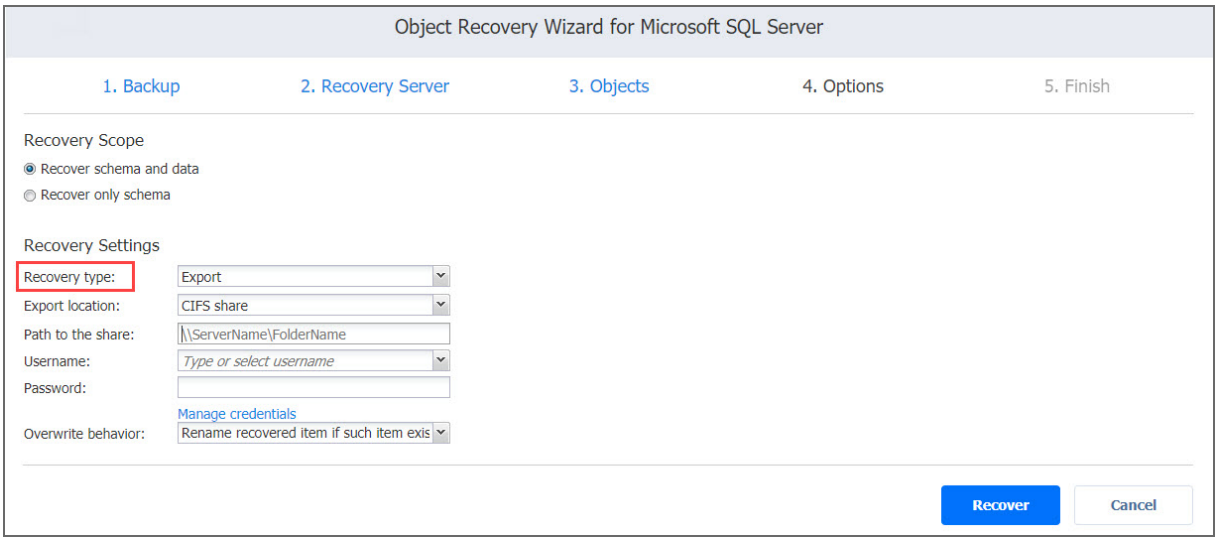

## Recovery Settings

Set up the recovery type and overwrite behavior.

Recovery Type

- **Recover to original location**: Recover objects to the same server and SQL instance where they were originally located.
- **Recover to custom location**: Recover objects to a different instance.
- **Export** : Export objects as files to a specified location.
  - **SQL instance**: Select the target SQL instance.
  - **Target database**: Select the target database of the selected instance.
  - **Local folder on Recovery Server**: Specify a path to save objects.
  - **CIFS share**: Specify a remote CIFS (Windows) file share and your credentials for it (or select them

from the **Manage credentials** list).

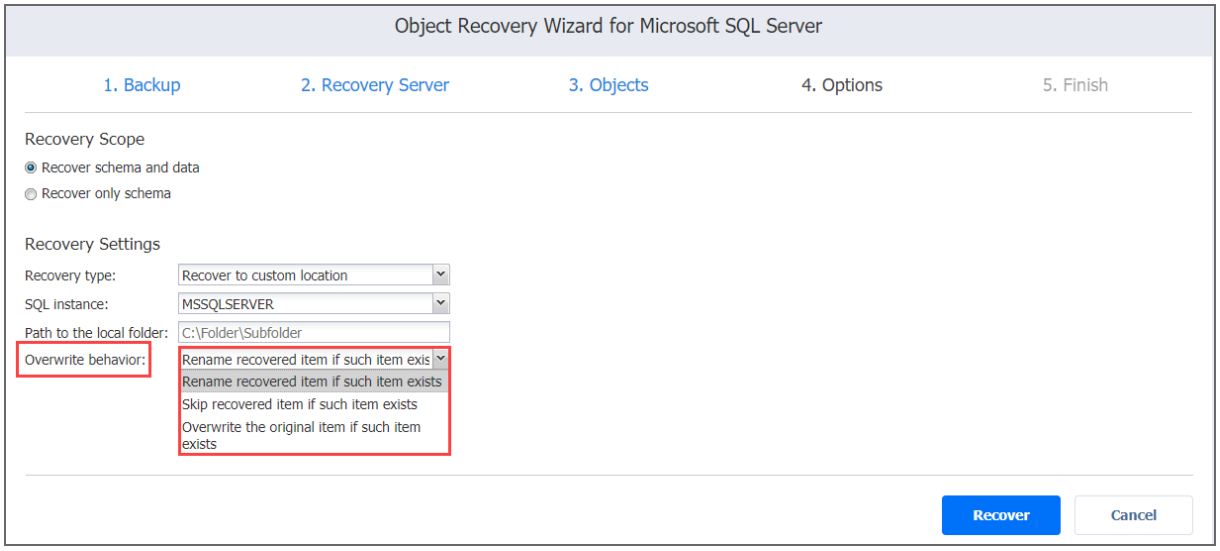



If you are using a domain name, enter it in the following format: domain\username

## Overwrite Behavior

Select what to do if the recovered item conflicts with an existing one:

- **Rename recovered item if such item exists**
- **Skip recovered item if such item exists**
- **Overwrite the original item if such item exists**




Click **Recover** to start the object recovery process. The **Finish** page opens.

# Performing Universal Object Recovery

With Universal Object Recovery you can choose a disk from a VM recovery point and mount it to a target machine. This will allow you to recover backup data located on the mounted disk. Before creating a Universal Object Recovery job, make sure the System Requirements for recovering files to a server are met.

Please refer to the following topics for creating a Universal Object Recovery job:


- "Opening Universal Object Recovery Wizard" on page 1052
- "Universal Object Recovery Wizard: Backup" on page 1054
- "Universal Object Recovery Wizard: Disks" on page 1057
- "Universal Object Recovery Wizard: Options" on page 1058

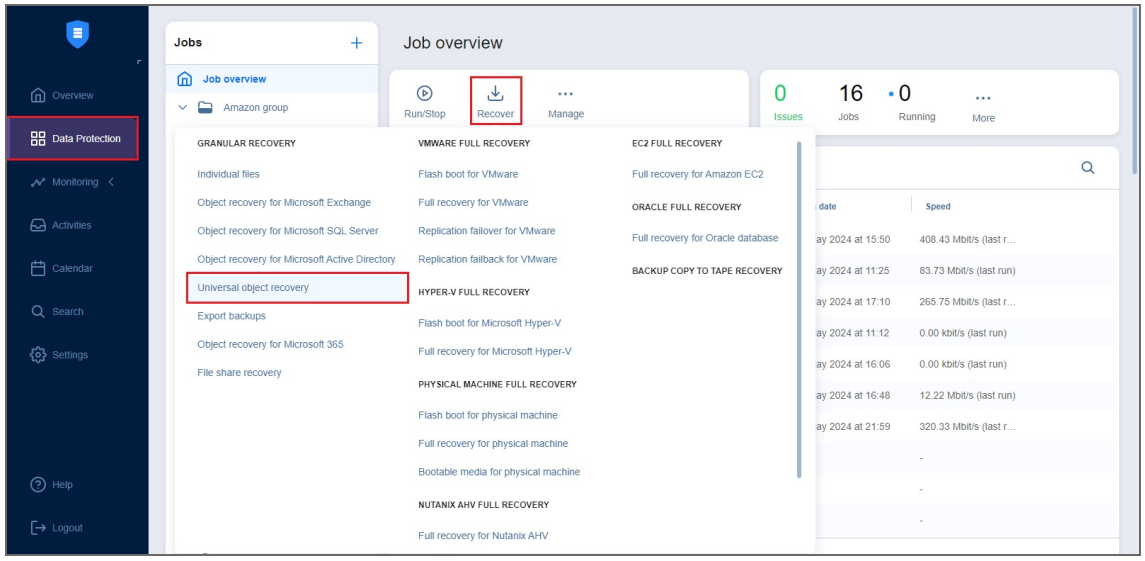
# Opening Universal Object Recovery Wizard

You can start the recovery process either from the **Data Protection** menu, by using the Search function, or from the **Repositories** page in **Settings** (for example, if you no longer have a backup job but still have the backup). Refer to the following sections for more details:

- Starting Universal Object Recovery from Data Protection Menu
- Starting Universal Object Recovery from a Backup Repository

## Starting Universal Object Recovery from Data Protection Menu

To start Universal object recovery from the **Data Protection** menu, click **Recover** and choose **Universal object recovery**.

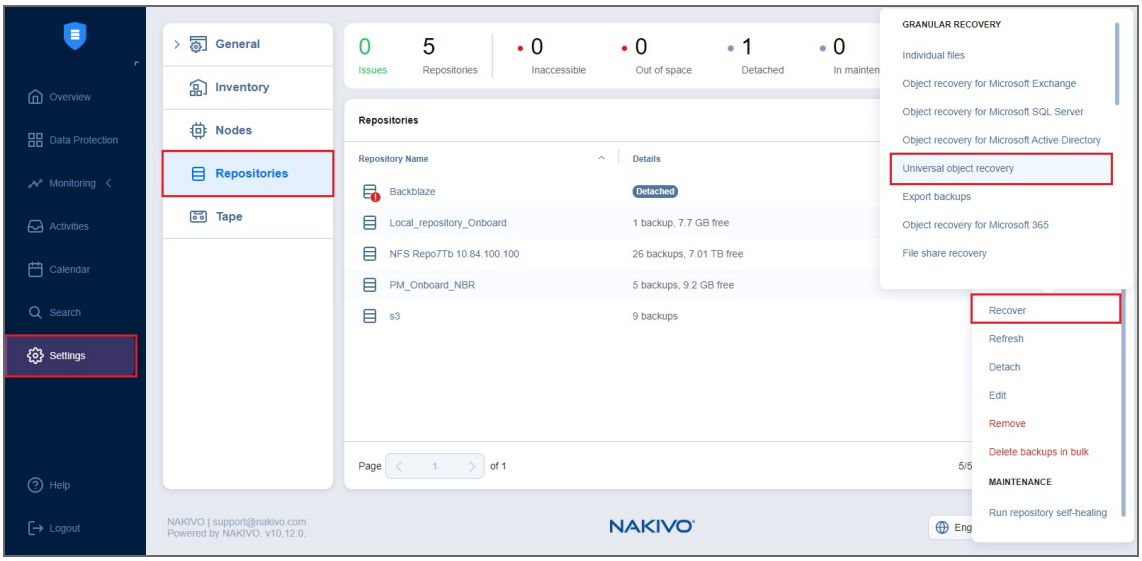



## Starting Universal Object Recovery from a Backup Repository

To start Universal object recovery from a **Backup Repository**:

1. Click **Settings** in the main menu of NAKIVO Backup & Replication.
2. Go to the **Repositories** tab and hover the cursor over the **Backup Repository** containing the required backup.

3. Click the ellipsis **Manage** button, click **Recover**, and select **Universal Object Recovery**.




The new **Universal Recovery Job Wizard** opens.

# Universal Object Recovery Wizard: Backup

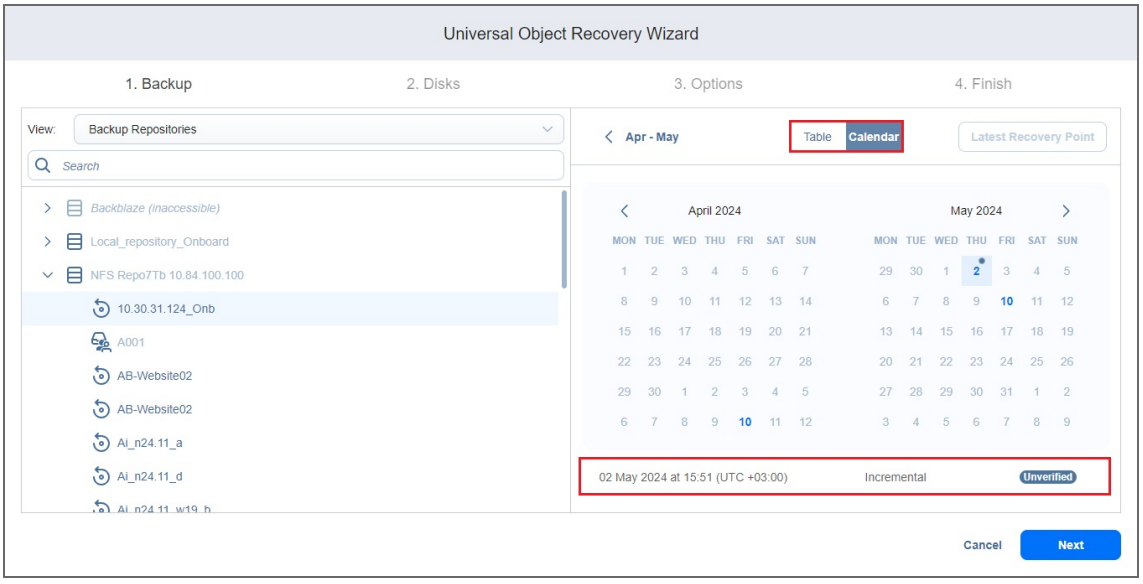
1. On the **Backup** page of the wizard, select a backup using either the **Backup Repositories** or **Jobs & Groups** view in the left pane.

Both federated and standalone (not used as members of federated repositories) backup repositories can be selected.

**Notes**

- You can perform Universal Object Recovery from Proxmox VE VM backups.
- You cannot select a federated repository member as a source for recovery jobs.
- You can search for a specific backup, job, group, or repository (depending on the selected view) by entering its name into the **Search** field.
- Selecting a single backup object opens the **Calendar** view of the date picker if this view was the last user-selected view. Selecting a single backup object opens the **Table** view if this was the last view the user selected.
- You can select a backup object with the corrupted, inaccessible, or pending removal recovery points, but you cannot select the backup objects of the detached repository.
- You cannot perform recovery from recovery points located in an inaccessible federated repository member or those that depend on inaccessible recovery points.

2. The latest recovery point is selected by default.



**Notes**

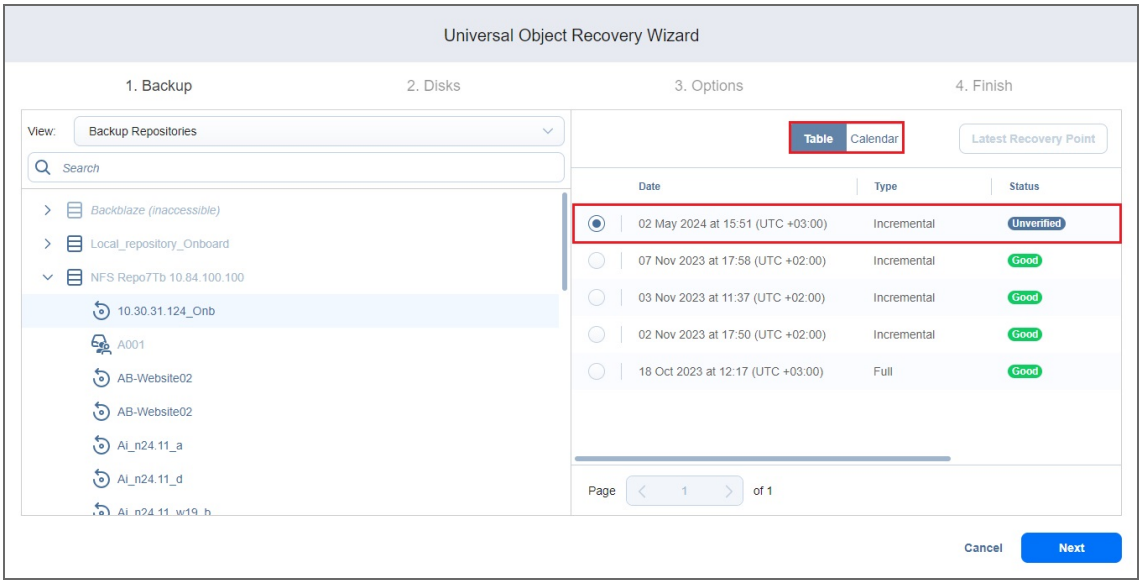
- The selected date is highlighted.
- If a recovery point is selected in the **Calendar** or **Table** view, this recovery point is selected by default the next time you open the **Calendar** or **Table** view.
- The selected view, either **Calendar** or **Table**, is saved on a per-user basis.

3. If necessary, toggle between the **Calendar** and **Table** views:
   - In the **Calendar** view, select a date or click **Latest Recovery Point** to select the date with the latest recovery point.

**Notes**

- The **Latest Recovery Point** option is disabled if:
  - The latest recovery point is currently selected.
  - All the recovery points of the backup object are corrupted, inaccessible, or pending removal.
- A small dot is displayed on top of the date if there is at least one recovery point created on that date.
- The selected recovery point has a preview displayed under the calendar showing the recovery point status:
- **green** – verified recovery points
- **red** – inaccessible/corrupted/pending removal recovery points
- **dark grey** – unverified recovery points
- Corrupted, inaccessible, or pending removal recovery points cannot be recovered.

   - In the **Table** view, select a date or click **Latest Recovery Point** to select the date with the latest recovery point.



**Note**

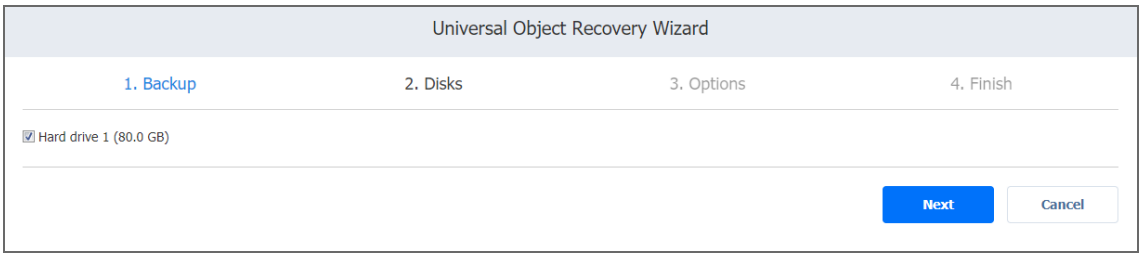
Clicking the **Latest Recovery Point** option selects the date with the latest recovery point that is not corrupted, inaccessible, or pending removal.

4. Click **Next** to go to the next page of the wizard.

To learn about the limitation on the chain of incremental recovery points, refer to the Knowledge Base article.

# Universal Object Recovery Wizard: Disks

On the **Disks** page of the wizard, choose one or more disks from the list of disks. Click **Next** to go to the next page of the wizard.



# Universal Object Recovery Wizard: Options

In the **Options** page of the wizard:

1. Specify mount location options:
    - **Mount location**: Choose the mount location from the drop-down list.

      > **Note**
      > Servers added using Direct Connect are not supported.

    - **Location IP address**: Enter the IP address of the server to which the disks will be mounted if it is not detected by the application based on the Mount location value. Here you can enter an IP address of any virtual or physical machine.
    - **Use custom SSH port**: To recover to a Linux server, select this option to enter a custom SSH port to be used for connecting to the recovery server. The default value is *22.*
    - **Credentials type**: Choose your preferred option and enter your respective credentials:
        - **Password**: Enter a username with administrative privileges for the file share entered above and your password.
        - **Private key**: Select your private key from the drop-down list.
2. Click the **Test Connection** button to test your credentials for the specified recovery server. If your credentials are correct, a checkmark appears to the right of the button.
3. **Malware detection**: With this option enabled, the backups are scanned for malware using the configured antivirus software on the scan server.
4. Optionally, if you have selected **Enabled** for the **Malware detection** option, click the **settings** link to configure the following options:
    - **Scan server**: Select a specific scan server for the job or leave the **Default** setting. If **Default** is selected, the Transporter is used as the scan server and can support a maximum of 2 concurrent scan tasks.

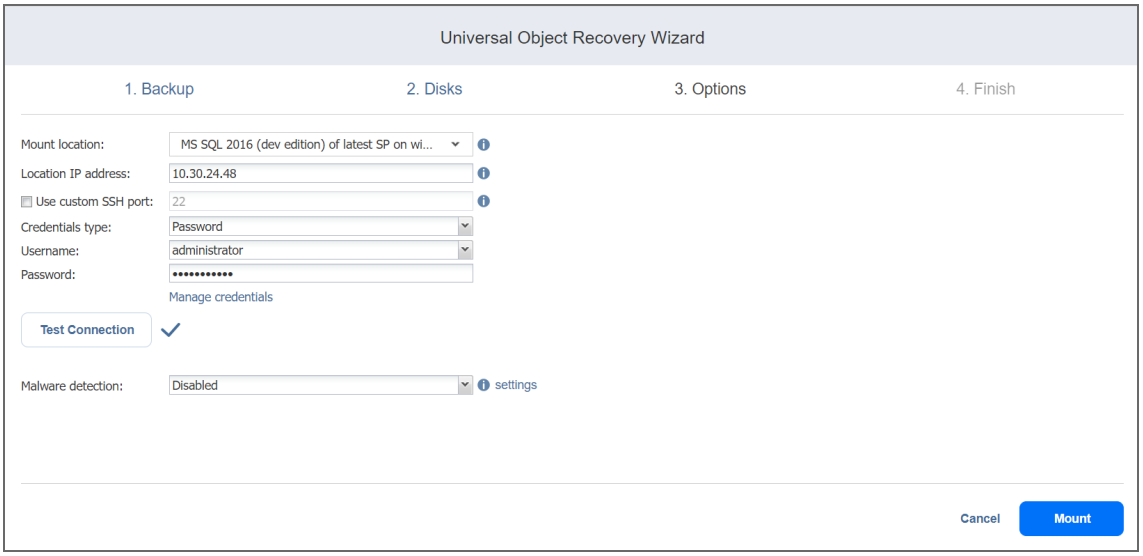
      > **Notes**
      > - For the **Default** option, if the **Repository Transporter** is the installed **Transporter**, it require the master password to function as the scan server.
      > - For more details on the requirements for Scan Server, refer to the Feature Requirements.

    - **Scan type**: Choose between the **Deep scan** and the **Quick scan**:
        - **Deep scan**: When this option is selected, the antivirus software scans the entire backup and may take longer to complete.

- **Quick scan**: When this option is selected, the antivirus software scans only OS disks in the backup.

- **If malware is detected**: Choose the behavior if malware is detected:
  - **Fail recovery**: With this option, the recovery process fails in case the job has only one disks. If the job has several disks, the infected disks are skipped and the job continues to run.
  - **Continue recovery**: When this option is selected, the recovery job completes the scanning process and recovers the infected disks.
- **Scan timeout**: Specify the timeout for the malware detection process. If the specified amount of time is exceeded, the recovery job fails.
- Click **Apply** when you're done.

5. Click **Mount** to confirm mounting your disks to the selected recovery server.




The Universal Object Recovery is started and the **Finish** page of the wizard opens.

6. Click the **Activities** link to go to the **Activities** page if you want to view the progress of the Universal Object Recovery.
7. Click **Close** to close the Universal Object Recovery Wizard. Upon successful Universal Object Recovery, the disks are mounted to the recovery server.

# Full Recovery

With NAKIVO Backup & Replication, you can recover an entire VM in case of hardware or VM failure. The VM is recovered in the same state as it was during the backup and it will appear on the host selected for recovery. You can also perform cross-platform recovery that allows you to export virtual disks from VM backups to different formats for further manual recovery of the VMs in different virtual environments.

The Flash boot feature allows you to recover an entire VM from the backup in seconds. This feature makes it possible to boot a VM directly from a compressed and deduplicated backup without recovering the entire VM. The VM can be started in a few seconds, and it can run directly from the backup.

With the Backup Copy to Tape Recovery feature, you can back up/send backup copies of selected folders or send immutable recovery points to repositories, then recover from these repositories and send the backup copies to tape.

Use the Proxmox VE VM Recovery feature to recover an entire VM in case of disaster. The VM is recovered in the same state as it was during the backup and appears on the host selected for recovery.

In addition, multiple recovery jobs and/or users may access the same recovery point even if it is currently in use by an existing recovery job/session.

For more details, refer to the corresponding articles below:


- “Performing Flash Boot Recovery” on page 1185
- “VMware VM Recovery” on page 1061
- “Starting Recovery from Tape” on page 1093
- “Performing Cross-Platform Recovery” on page 1096

# VMware VM Recovery

With NAKIVO Backup & Replication, you can recover full VMs from backups. Multiple VMs can be recovered within a single recovery job—the system allows you to select separate resources (containers, datastores, networks, and folders) for different VMs or to restore all VMs using the same setup. When you run VM recovery, a new VM is created; the source VM is not reverted to a previous state or replaced with the new VM.

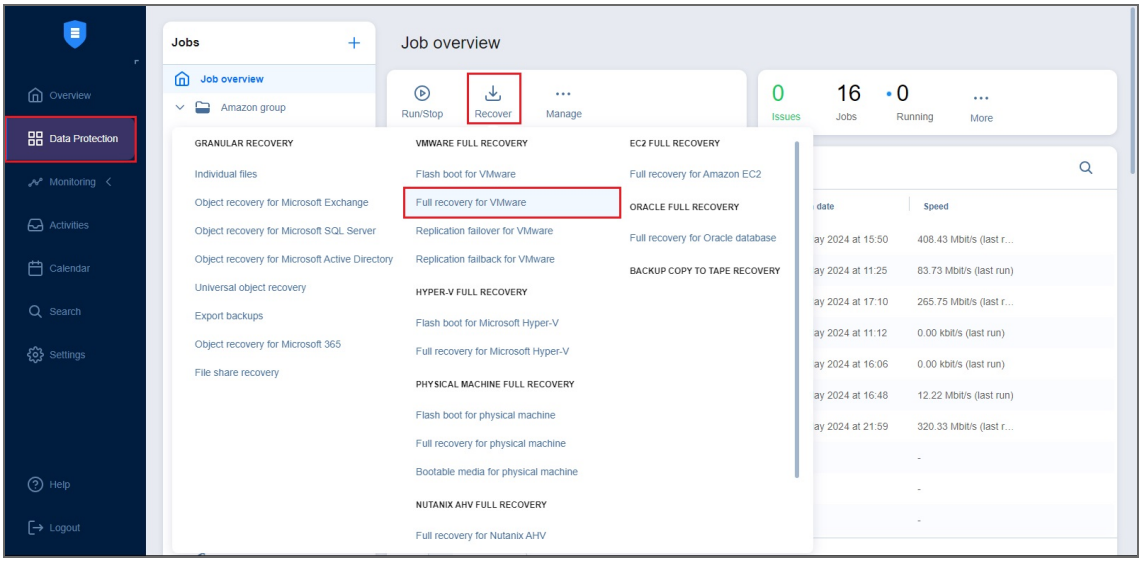
Refer to the following topics for more information:


- "Starting VMware VM Recovery" on page 1062
- "Recovery Job Wizard for VMware: Backups" on page 1064
- "Recovery Job Wizard for VMware: Destination" on page 1068
- "Recovery Job Wizard for VMware: Options" on page 1072

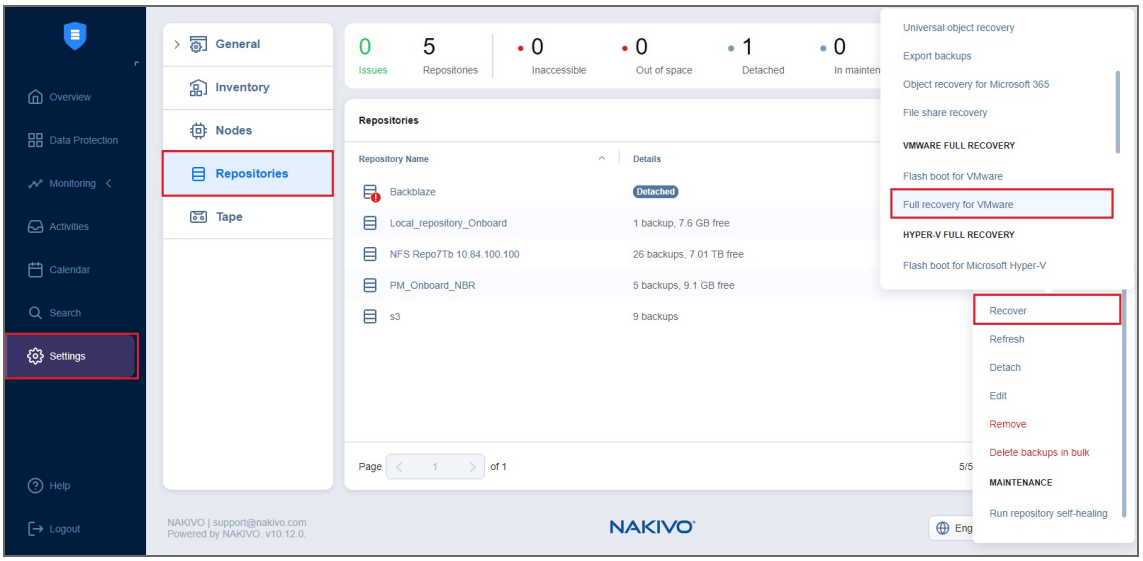
# Starting VMware VM Recovery

To recover entire VMware VMs from backups, do one of the following:

- Start recovery from the **Data Protection** menu by clicking **Recover** and then clicking **Full recovery for VMware**.




- Open the **New Full Recovery Job Wizard for VMware** from the **Repositories** tab by following the steps below:
  1. From the main menu of NAKIVO Backup & Replication, click **Settings**.
  2. Go to the **Repositories** tab.
  3. Hover over the **Backup Repository** containing the needed backup, click the ellipsis **Manage** button, and click **Recover**.
  4. Click **Full recovery for VMware** under **VMware Full Recovery**.




- Open the **New Full Recovery Job Wizard** from the **Tape** tab by following the steps below:

1. From the main menu of NAKIVO Backup & Replication, click **Settings**.
2. Go to the **Tape** tab and click **Backups**.
3. Select the necessary VMware VM backups.
4. Click the **Recover** button.

- Alternatively, the recovery can be performed by using the Search function.

The **Full Recovery Job Wizard for VMware** opens.

# Recovery Job Wizard for VMware: Backups

1. On the **Backups** page of the wizard, select a backup using either a **Backup Repositories**, **Jobs & Groups**, or **Tape** view in the left pane.

Both federated and standalone (not used as members of federated repositories) backup repositories can be selected.

> **Notes**
>
> - You cannot select a federated repository member as a source for recovery jobs.
> - You cannot perform recovery from recovery points located in an inaccessible federated repository member or those that depend on inaccessible recovery points.
> - File share backup and Microsoft 365 backup options are disabled when recovering backups from tape.

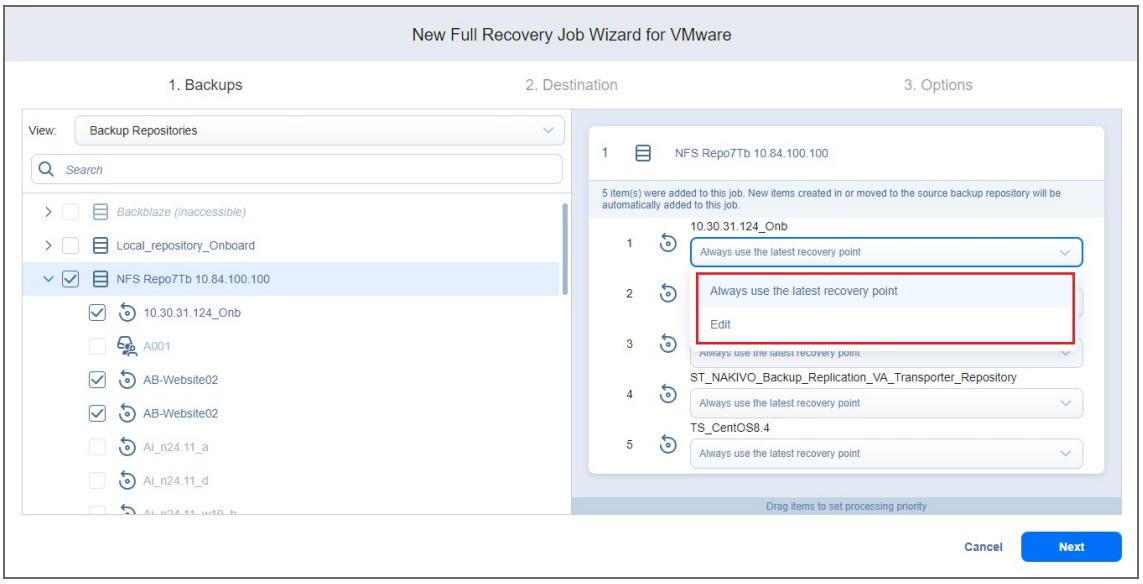
Selecting an item in the tree adds the parent item and all children items to the right pane.

> **Notes**
>
> - You can search for a specific backup, job, group, or repository (depending on the selected view) by entering its name into the **Search** field.
> - Selecting a single backup object opens the **Calendar** view of the date picker if this view was the last user-selected view. Selecting a single backup object opens the **Table** view if this was the last view the user selected.
> - You can select a backup object with the corrupted, inaccessible, or pending removal recovery points, but you cannot select the backup objects of the detached repository.

2. In the drop-down list of the selected item on the right pane, select one of the options:

- **Always use the latest recovery point** (displayed by default).

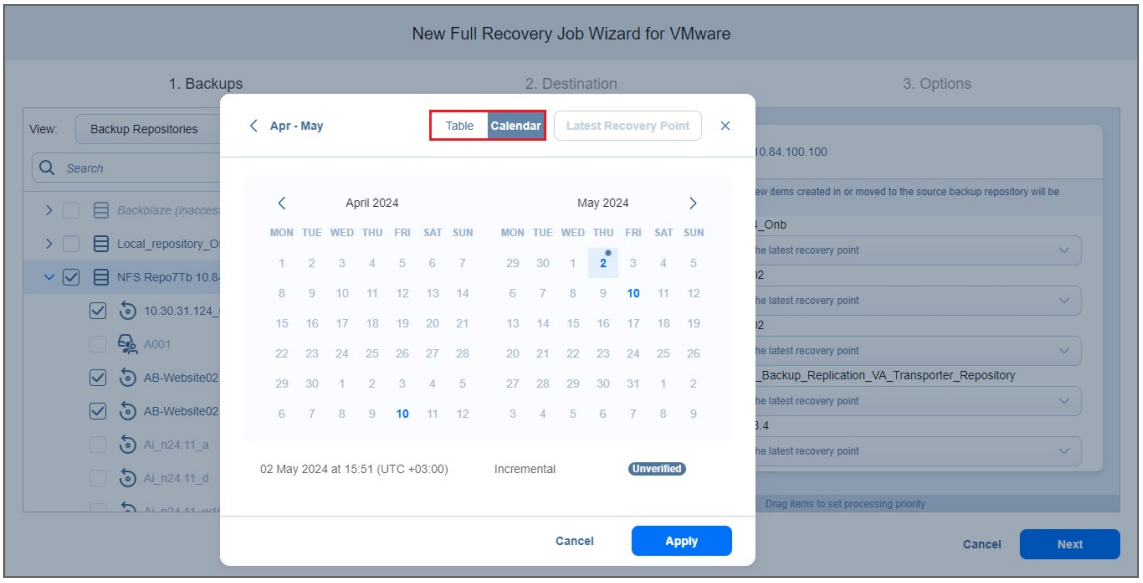
- **Edit** to open the **Recovery points picker** dialog box in the **Calendar** or **Table** view and pick a date.




**Notes**

- If a recovery point is selected in the **Calendar** or **Table** view, this recovery point is selected by default the next time you open the **Calendar** or **Table** view.
- The selected view, either **Calendar** or **Table**, is saved on a per-user basis.

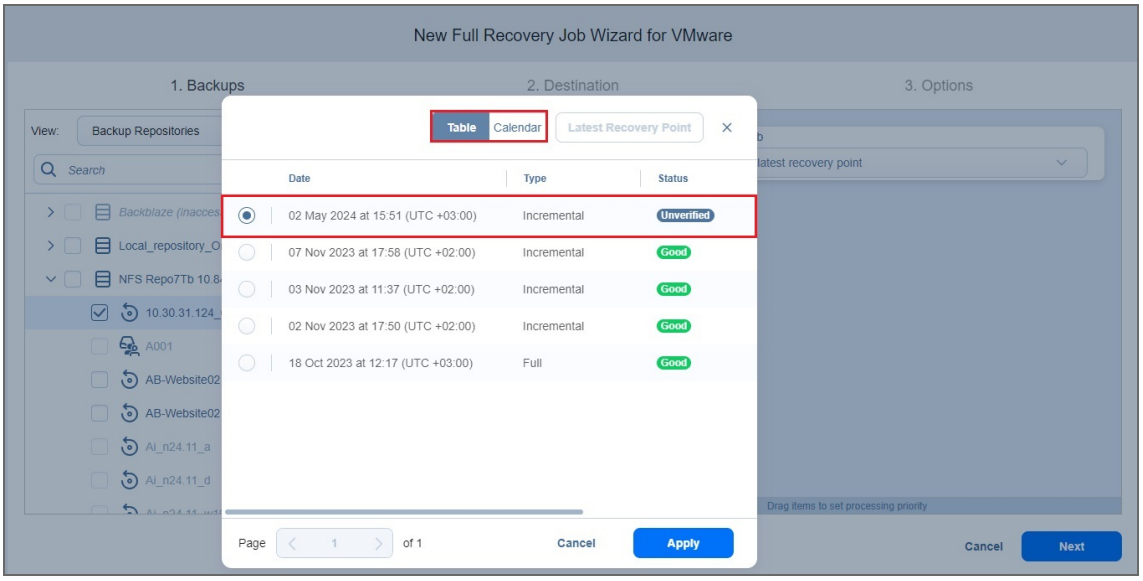
3. If necessary, toggle between the **Calendar** and **Table** views:
   - In the **Calendar** view, select a date or click **Latest Recovery Point** to select the date with the latest recovery point.



**Notes**

- A small dot is displayed on top of the date if there is at least one recovery point created on that date.
- The selected recovery point has a preview displayed under the calendar showing the recovery point status:
- **green** – verified recovery points
- **red** – inaccessible/corrupted/pending removal recovery points
- **dark grey** – unverified recovery points
- Corrupted, inaccessible, or pending removal recovery points cannot be recovered.
- The selected date is highlighted.

- In the **Table** view, select a date or click **Latest Recovery Point** to select the date with the latest recovery point.



**Notes**

- Clicking the **Latest Recovery Point** option selects the date with the latest recovery point that is not corrupted, inaccessible, or pending removal.
- The **Latest Recovery Point** option is disabled if:
  - The latest recovery point is currently selected.
  - All the recovery points of the backup object are corrupted, inaccessible, or pending removal.
- If the selected recovery point is encrypted and the corresponding password hash is not available, the **Always use the latest recovery point** or the recovery point timestamp is highlighted in red with a red lock icon. To add such a recovery point to the job, you need to provide the password manually. See Providing Passwords for encrypted recovery points.

4. Click **Next** to go to the next page of the wizard.

**Note**

You cannot proceed to the next step if there is at least one selected recovery point to which a password has not been provided.

# Recovery Job Wizard for VMware: Destination

On the **Destination** page of the wizard, select a location for the recovered VMs or VM templates. Proceed as described in these sections:

- Setting the Same Host, Datastore, and Network for All Recovered VMs
- Setting Original Location for All Recovered VMs
- Setting Backup Repository as Destination
- Setting the Default Destination for Recovered VMs
- Setting Different Options for Recovered VMs

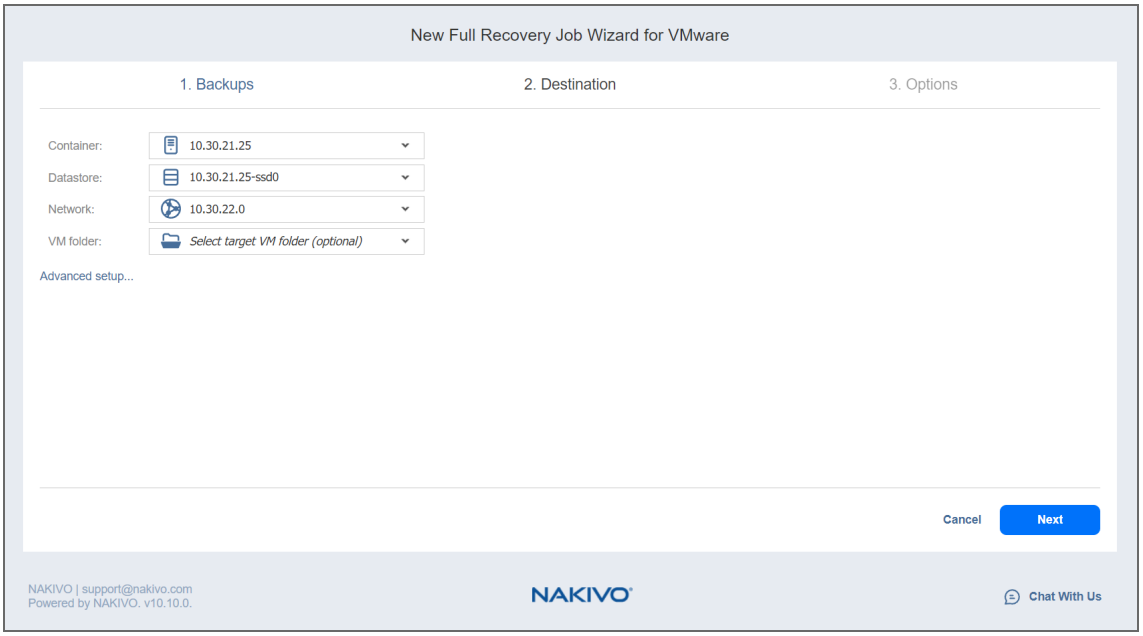
**Note**

When recovering a VMware VM from an ESXi 8.0 host to an ESXi 7.0 host, the **VirtualMachineConfigSpec** > **VirtualMachineSgxInfo** property is skipped.

## Setting the Same Host, Datastore, and Network for All Recovered VMs

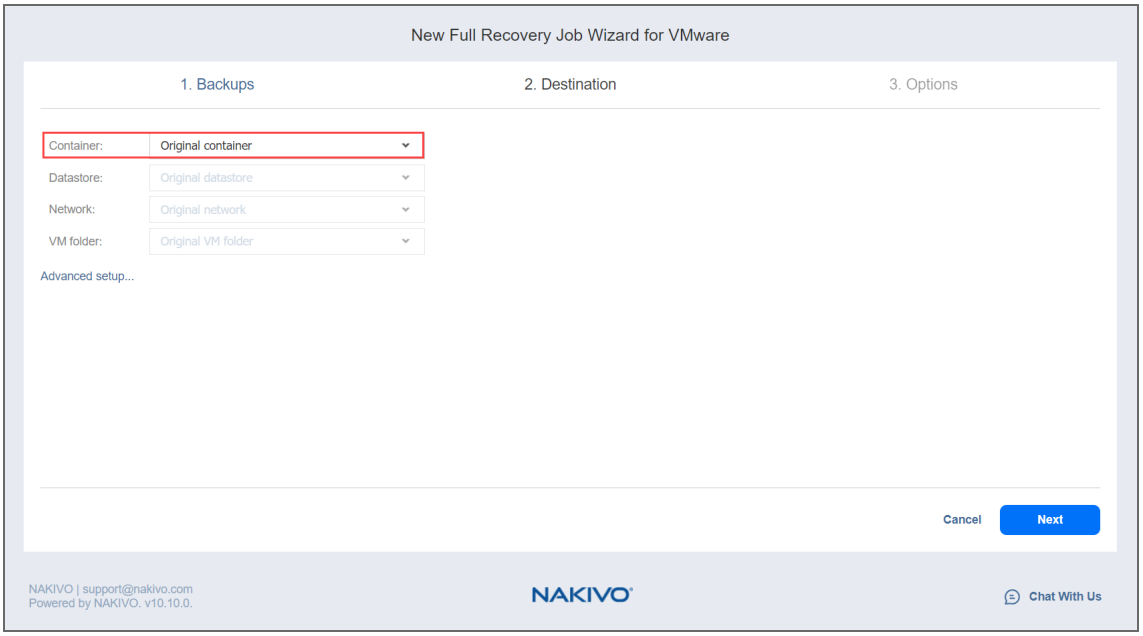
To recover all VMs or VM templates to the same container/folder and datastore, and to connect all recovered VMs to the same networks, follow the steps below:

1. If you selected a tape backup on the **Backups** page, select **New VM(s)** in the **Recover to** drop-down list. Otherwise, proceed directly to step 2.
2. Choose a cluster, host, or resource pool from the **Container** drop-down list.
3. Choose a datastore from the **Datastore** drop-down list.
4. Choose a network from the **Network** drop-down list.
5. Optionally, you can choose a folder if there is one in the container.



## Setting Original Location for All Recovered VMs

To recover VMs or VM templates to their original location, select **Original container** from the **Container** drop-down list and click **Next**.




**Important**

If the location of the VMs or VM templates is unknown or unavailable, you will have to configure it manually via **Advanced options**.

## Setting Backup Repository as Destination

If you selected a tape backup on the **Backups** page, you have the additional option of recovering to an existing **Backup Repository**. To do so, select **Backup Repository** from the **Recover to** drop-down menu and choose the appropriate repository from the **Repository** drop-down menu below.

**Notes**

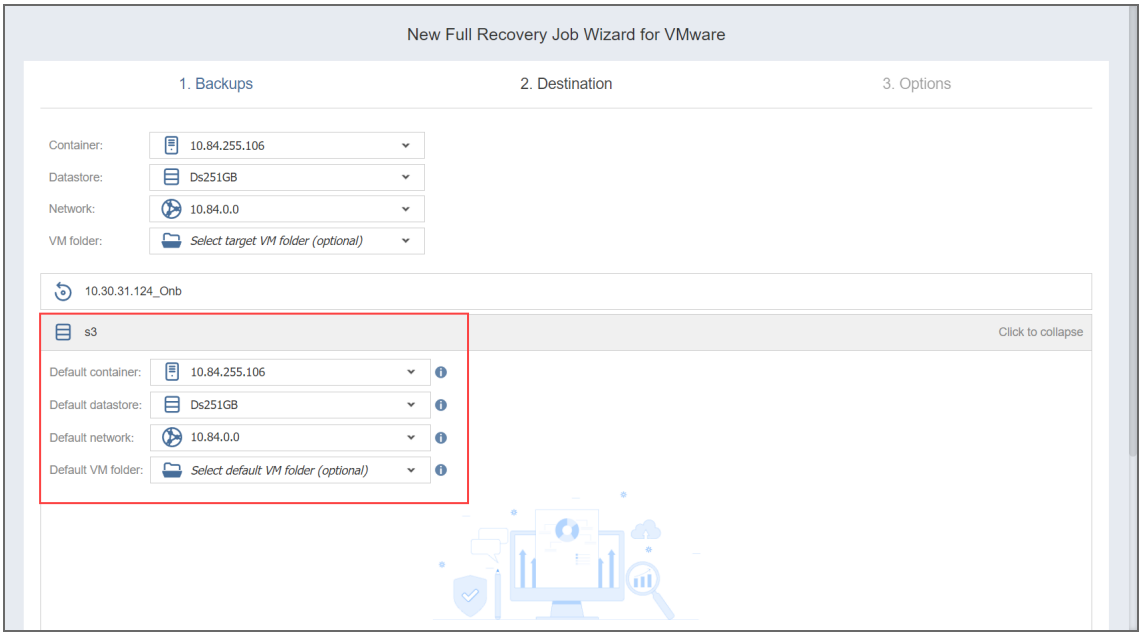
- You cannot perform recovery from recovery points located in an inaccessible federated repository member or those that depend on inaccessible recovery points.
- You cannot select a federated repository as a destination for recovery from tape media.

## Setting the Default Destination for Recovered VMs

If you have chosen a host, cluster, folder, or a **Backup Repository** as a source for your recovery job on the **Backups** page, you can set the default container, datastore, and VM folder for the recovered VMs or VM templates. To do this, follow the steps below:

1. Click **Advanced setup** and then click on the name of the chosen host, cluster, folder, or a resource pool.

2. If you selected a tape backup on the **Backups** page, select **New VM(s)** in the **Recover to** drop-down list. Otherwise, proceed to the next step.
3. Choose a **Default container**.
4. If you have chosen the backup job on the **Source** page, you can choose a **Default Network**.
5. Optionally, you can also choose a **Default VM folder.**

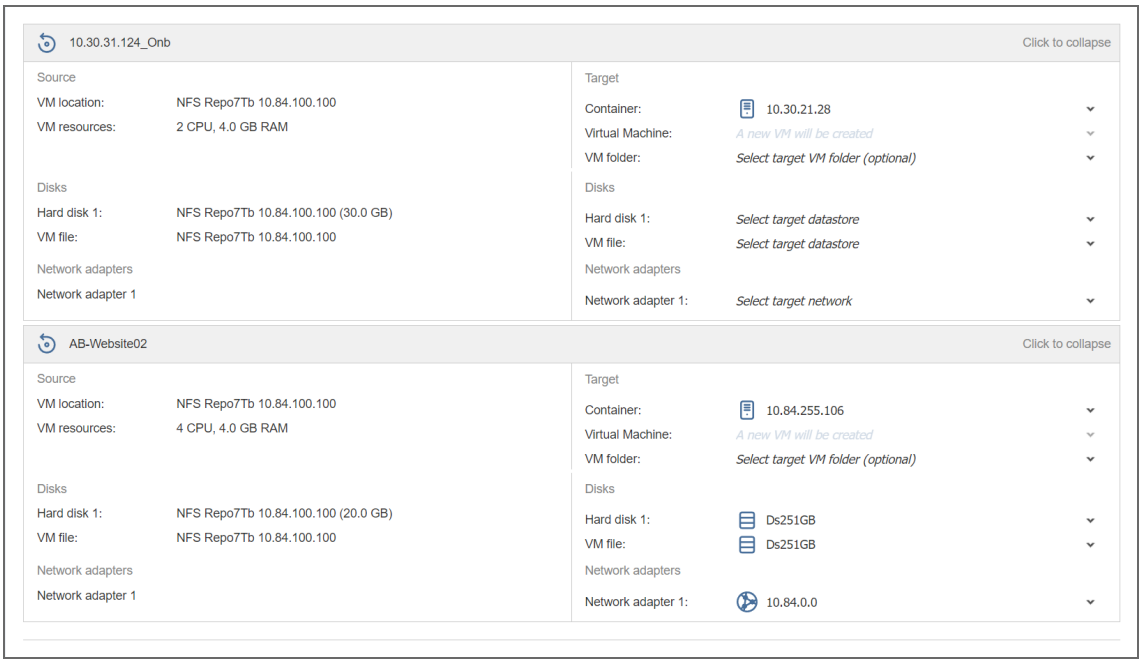



## Setting Different Options for Recovered VMs

To specify different options for recovered VMs or VM templates, follow the steps below:

1. Click **Advanced setup**.
2. Click on the backup to expand its recovery options.
3. If you selected a tape backup on the **Backups** page, select **New VM(s)** or **Backup Repository** in the **Recover to** drop-down list. Otherwise, proceed to the next step.
4. Choose a target location, target datastore, and target network for each VM.
5. To get additional information about the source and target VMs within a backup, click on its name.

6. Click **Next** to go to the next page of the wizard.



# Recovery Job Wizard for VMware: Options

On the **Options** page of the wizard, set the recovery job options.


- "Job Options" below
- "Recovered VM Options" on page 1074
- "Pre and Post Actions" on page 1075
  - "Setting a Pre-Job Script" on page 1075
  - "Setting a Post-Job Script" on page 1076
  - "Email Notifications" on page 1077
- "Data Transfer" on page 1077
  - "Transport Mode" on page 1078
  - "Transporter Pool" on page 1078
  - "Transporters" on page 1078
  - "Transporter Load" on page 1079
  - "Bandwidth Throttling" on page 1079
  - "Multi-Channel Processing" on page 1079
  - "Bottleneck Detection" on page 1080
- "Completing the New Recovery Job Wizard for VMware" on page 1080


## Job Options

Specify the general options as follows:

1. **Job name**: Specify a name for the recovery job.
2. **Job priority**: Select a job priority level between *1* and *5*, with *1* being the highest priority. Jobs with higher priority levels are prioritized by **Transporters** during job processing.

   > **Note**
   > This option is only available in the Enterprise, Enterprise Essentials, Enterprise Plus, MSP Enterprise, and MSP Enterprise Plus editions.

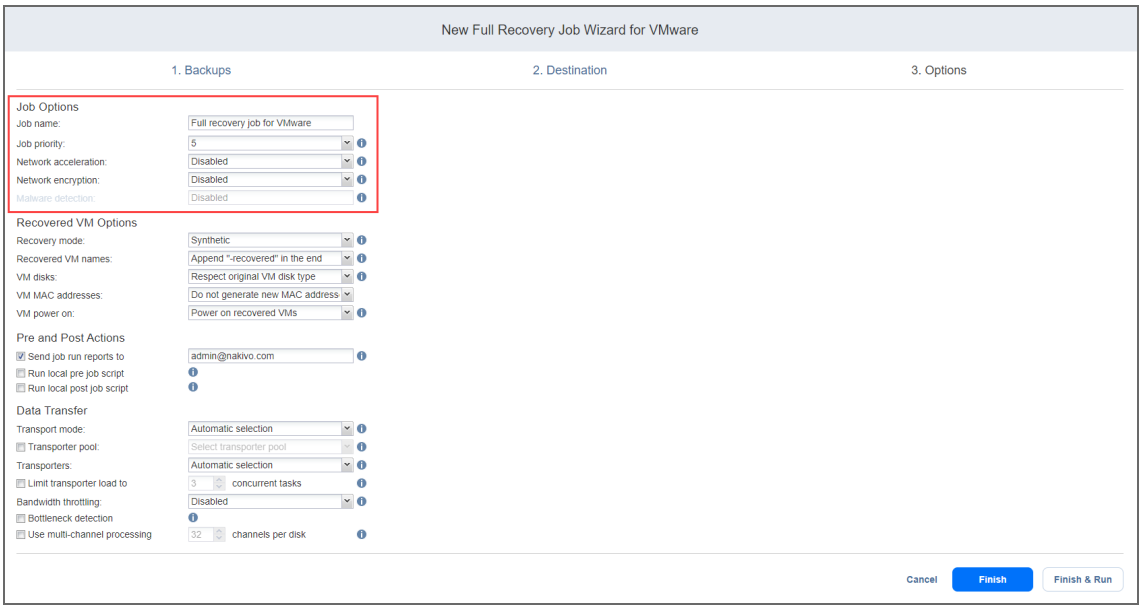
3. **Network acceleration**: With network acceleration enabled, NAKIVO Backup & Replication uses compression and traffic reduction techniques to speed up data transfer. Enable this option if you plan to recover VMs over WAN or slow LAN links.
4. **Network encryption**: With network encryption enabled, VM data is protected with AES 256 encryption while traveling over the network. Data encryption increases the backup time and CPU load on machines running **Transporters**. Select this option when recovering over WAN without a VPN connection.

5. **Malware detection**: With this option enabled, the backups are scanned for malware using the configured antivirus software on the scan server.
6. Optionally, if you have selected **Enabled** for the **Malware detection** option, click the **settings** link to configure the following options:
    - **Scan server**: Select a specific scan server for the job or leave the **Default** setting. If **Default** is selected, the Transporter is used as the scan server and can support a maximum of 2 concurrent scan tasks.

      **Notes**
      - For the **Default** option, if the **Repository Transporter** is the installed **Transporter**, it require the master password to function as the scan server.
      - For more details on the requirements for Scan Server, refer to the Feature Requirements.

    - **Scan type**: Choose between the **Deep scan** and the **Quick scan**:
        - **Deep scan**: When this option is selected, the antivirus software scans the entire backup and may take longer to complete.
        - **Quick scan**: When this option is selected, the antivirus software scans only OS disks in the backup.
    - **If malware is detected**: Choose the behavior if malware is detected:
        - **Fail the recovery job**: With this option, the recovery process fails in case the job has only one VM. If the job has several VMs, the infected VMs are skipped and the job continues to run.
        - **Continue and recover to isolated network**: When this option is selected, the recovery job completes the scanning process and recovers the infected VMs to a temporary isolated network.
    - **Scan timeout**: Specify the timeout for the malware detection process. If the specified amount of time is exceeded, the recovery job fails.

- Click **Apply** when you're done.



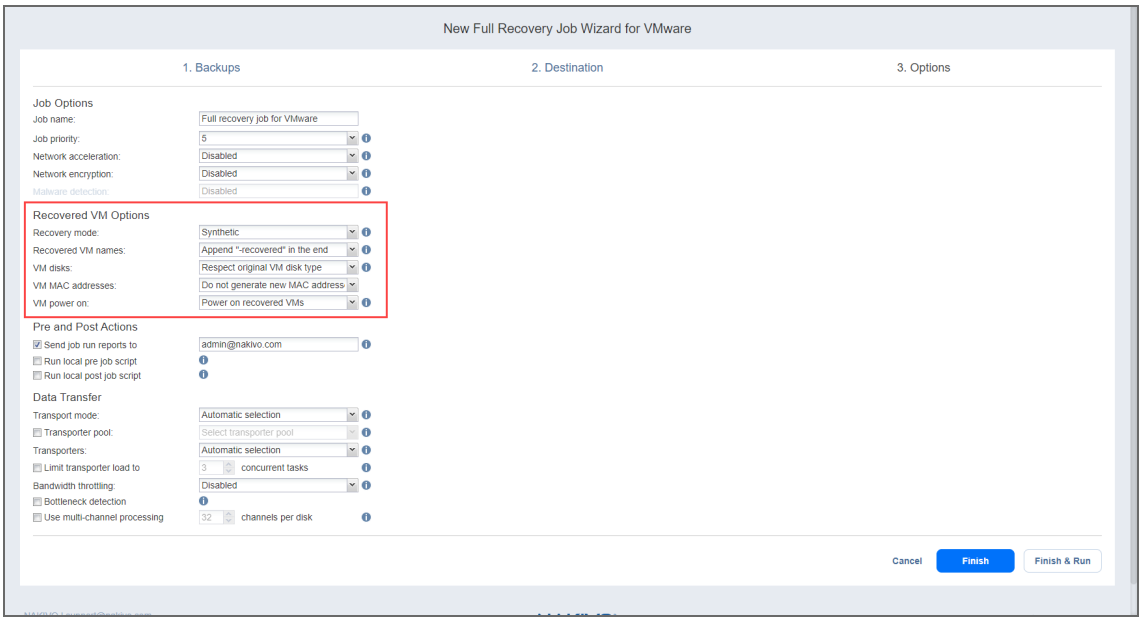

## Recovered VM Options

Specify the recovery options as follows:

1. **Recovery mode**: Choose one of the following:
   - **Synthetic**: With this recovery mode, the VMs are recovered with the environmental dependencies (such as CPU affinity) removed. Select this option when recovering VMs to a new location.
   - **Production**: With this recovery mode, environment dependencies are preserved on recovered VMs. Make sure the location to which the VMs are recovered does not contain the original VMs, otherwise, UUID and MAC address conflicts may occur.
2. **Recovered VM names**: Choose one of the following:
   - **Append "-recovered" in the end**: Source VM names are used for recovered VM names and "-recoverd" is added after the name of the recovered VMs.
   - **Leave recovered VM names as is**: Recovered VM names are identical to the source VM names.
   - **Enter custom recovered VM names**: You can enter custom names for recovered VMs.
3. **VM disks**: Choose one of the following:
   - **Respect original VM disk type**: When specified, disks that respect the original VM disk type are created on target VMs. Select this option to recover VMs to their original location.
   - **Create only thin disks on target VMs**: When specified, only thin disks are created on target VMs. Use this option to save space on the target datastore.
4. **VM MAC addresses**: Choose one of the following:
   - **Do not generate new MAC addresses**: When this option is selected, the recovered VM has the same MAC address as the source VM.

- **Generate new MAC addresses**: When this option is selected, a new MAC address is generated for the recovered VM.

5. **VM power on**: When the **Power on recovered VMs** option is selected, the recovered VMs are powered on.




## Pre and Post Actions

NAKIVO Backup & Replication allows you to run a script before VM recovery begins (a pre-job script) and after the recovery of all VMs in the job has completed (a post-job script). The scripts can only be executed on the machine on which the **Director** is installed.

### Setting a Pre-Job Script

To run a script before the product begins recovering VMs, do the following:

1. Place a script file on the machine on which the **Director** is installed.
2. Select the **Run local pre job script** option and click the **settings** link.
3. Specify the following parameters in the dialog box that opens:
   - **Script path**: Specify a local path to the script on the machine on which the **Director** is installed. A script interpreter should be specified.

> **Example (Windows)**: *cmd.exe /c D:\script.bat*
> **Example (Linux)**: *bash /root/script.sh*

   - **Job behavior**: Choose one of the following job behaviors in relation to script completion:
     - **Wait for the script to finish**: With this option selected, VM backup is started only after the script is completed.

- **Do not wait for the script to finish**: With this option selected, the product runs the script and starts backing up VMs at the same time.

- **Error handling**: Choose one of the following job behaviors in relation to script failure:
  - **Continue the job on script failure**: With this option selected, the job performs VM backup even if the script has failed.
  - **Fail the job on script failure**: With this option selected, if the script fails, the job is failed and VM backup is not performed.

Setting a Post-Job Script

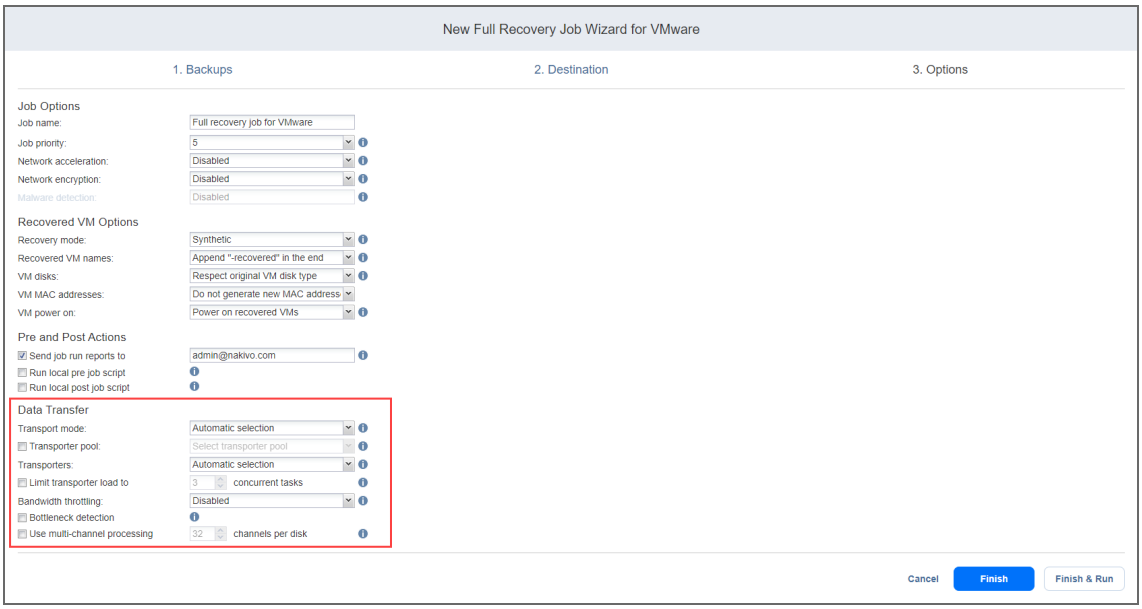
To run a script after the product has finished backing up all VMs, do the following:

1. Place a script file on the machine on which the **Director** is installed.
2. Select the **Run local post job script** option and click the **settings** link.
3. Specify the following parameters in the dialog box that opens:
   - **Script path**: Specify a local path to the script on the machine on which the **Director** is installed. A script interpreter should be specified.

     **Example (Windows)**: *cmd.exe /c D:\script.bat*
     **Example (Linux)**: *bash /root/script.sh*

   - **Job behavior**: Choose one of the following job behaviors in relation to script completion:
     - **Wait for the script to finish**: With this option selected, the job is in the “running” state until the script is completed.
     - **Do not wait for the script to finish**: With this option selected, the job is completed even if the script execution is still in progress.
   - **Error handling**: Choose one of the following job behaviors in relation to script failure:
     - **Continue the job on script failure**: With this option selected, script failure does not influence the status of the job.
     - **Fail the job on script failure**: With this option selected, if the script fails, the job

status is set to “failed” even if VM backup is successful.




### Email Notifications

NAKIVO Backup & Replication can send email notifications about job completion status to specified recipients. This feature complements global notifications and allows you to configure notifications on a per-job level.

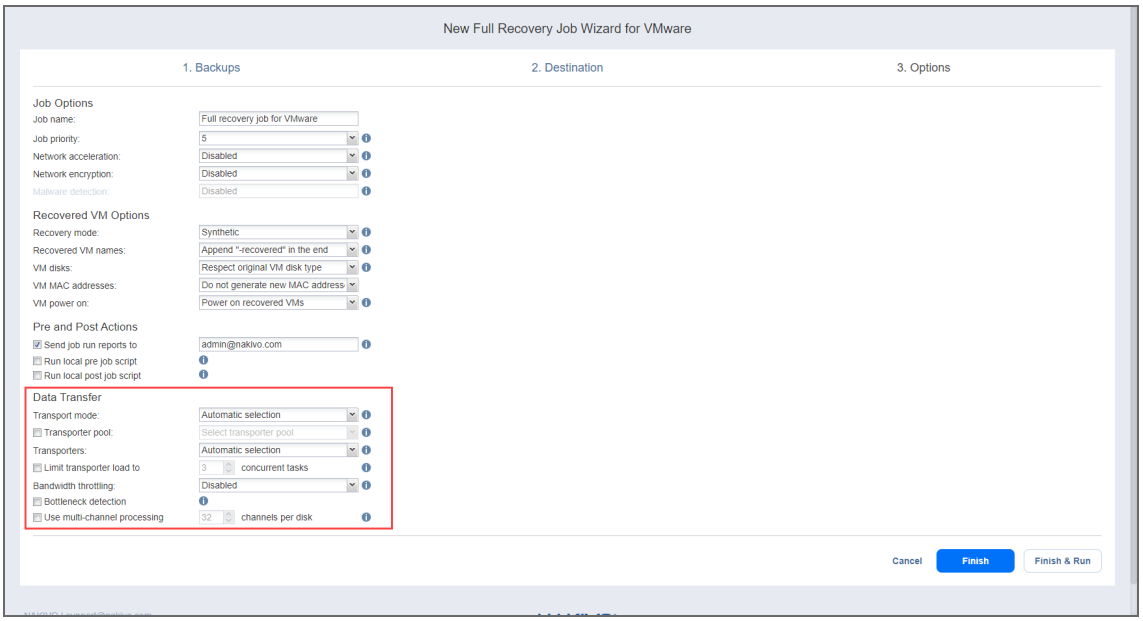
To enable this option, make sure that Email settings are configured.

To send email notifications, do the following:

1. In the *Pre and Post Actions* section, select the **Send job run reports to.**
2. Specify one or more email addresses in the text box. Use semicolons to separate multiple email addresses.

## Data Transfer

In the *Data Transfer* section, choose a transport mode for writing VM data, select which **Transporter** to be used for reading data from the source VM, and set bandwidth throttling.




### Transport Mode

NAKIVO Backup & Replication provides the following transport modes for writing VM data:

- **Automatic**: When this option is chosen, Hot Add mode is used where possible. If the product cannot use Hot Add, LAN mode is used.
- **Hot Add only**: NAKIVO Backup & Replication can write data directly to the datastore, bypassing the network, which can significantly improve job performance. This is achieved with the help of VMware's Hot Add technology. In order for the Hot Add feature to work, the target **Transporter** (the one that will be writing data ) should run on a host that has access to the target datastore(s).
- **LAN only**: Data will be written over LAN.

### Transporter Pool

If this option is enabled, only the transporters that belong to the selected transporter pool shall be used during the job run.

### Transporters

By default, the product automatically determines which **Transporter** should be used to read data from the source VM. However, you can manually specify which **Transporters** should be used for the job:

- **Automatic selection**: The product automatically determines the **Transporters** that are closest to the source and target hosts.
- **Manual - configured for all VMs**: Select this option to manually specify a single source and a single target **Transporter** to be used for data transfer by the job.
- **Manual - configured per host**: Select this option to manually specify **Transporters** for all source and target hosts.

### Transporter Load

You can limit the maximum number of transporter tasks used by the job. By default, this number is set to 3 concurrent tasks.

To change the default number of tasks, do the following:

1. In the **Data Transfer** section, select the **Limit transporter load to** checkbox.
2. Specify the number of concurrent tasks in the corresponding box.

### Bandwidth Throttling

Follow the steps below to set the speed of data transfer over the network for your recovery job:

1. For the **Bandwidth throttling** option, choose **Enabled**.

> **Note**
> If bandwidth throttling is disabled for the current job, global bandwidth rules may still apply to your job. Refer to Bandwidth Throttling for details.

2. Click the **Settings** link that becomes available.
3. The **Job Bandwidth Rules** dialog box opens displaying the list of available rules. You have the following options:
   - Create a new bandwidth rule for your recovery job:
     1. Click the **Create New Rule** button.
     2. The **New Bandwidth Rule** dialog box opens. Refer to the "Bandwidth Throttling" on page 407 topic for details on creating a bandwidth rule.
     3. Click **Save**.
   - **Activate an existing bandwidth rule for your job**: Select the checkbox to the left of the required bandwidth rule. To deactivate a bandwidth rule for your job, deselect the corresponding checkbox.
   - **Edit a bandwidth rule:** Click the **Edit** link for a bandwidth rule and modify it in the **Edit Bandwidth Rule** dialog box that opens.
   - **Disable a bandwidth rule:** Click the **Disable** link. The bandwidth rule is disabled for all jobs.
   - **Remove a bandwidth rule**: Click the **Remove** link and then click **Delete** to confirm your operation.

### Multi-Channel Processing

When this option is enabled, NAKIVO Backup & Replication performs recovery in multiple channels simultaneously, which can increase recovery speed. Specify the number of threads in the # channels per disk field.

**Important**

- The **Transporter** needs to have at least 2 CPU cores and 8 GB RAM available to perform recovery in multiple channels.
- Multi-channel processing is not supported for items added using **Direct Connect**.

Bottleneck Detection

When the **Bottleneck detection** option is enabled, additional information is collected and recorded in NAKIVO Backup & Replication logs in the course of data transfer for the purpose of bottleneck detection. Check this option to enable the **Bottleneck detection** capability of the **Transporters** engaged in the job.

### Completing the New Recovery Job Wizard for VMware

Click **Finish** or **Finish & Run** to complete the job creation.

**Note**

If you click **Finish & Run**, you will have to define the scope of your job. Refer to "Running Jobs on Demand" on page 364 for details.

## VMware Cloud Director Recovery

With NAKIVO Backup & Replication, you can recover VMware Cloud Director objects from backups. The recovery can be performed either for individual VMs or for vApps which may contain one or more VMs. When you run VM recovery, a new VM is created; the source VM is not reverted to a previous state or replaced with the new VM.

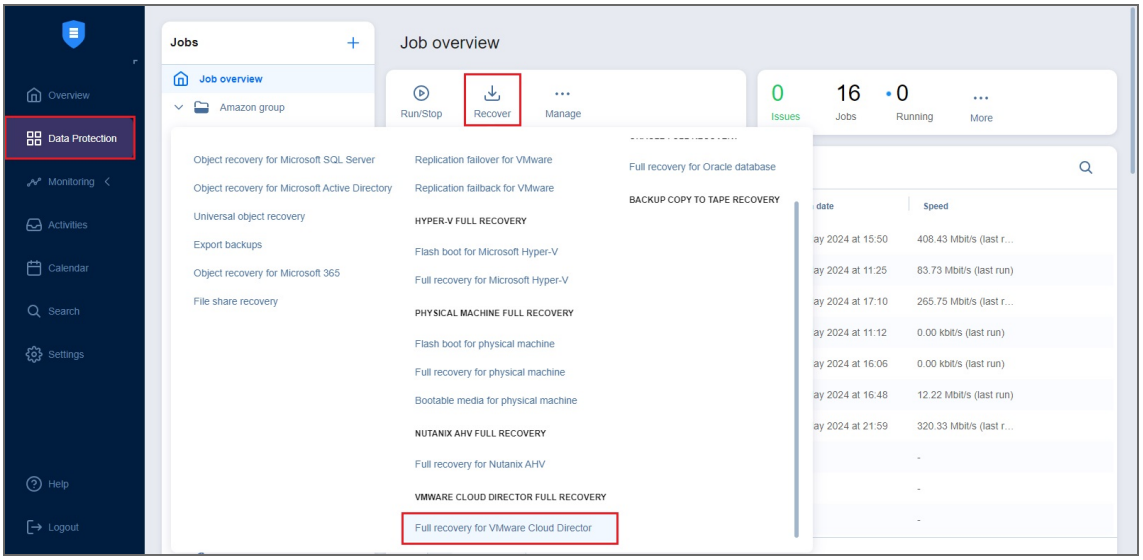
Refer to the following topics for more information:

- "Starting VMware Cloud Director Recovery" below
- "Recovery Job Wizard for VMware Cloud Director: Object" on page 1085
- "Recovery Job Wizard for VMware Cloud Director: Backups" on page 1082
- "Recovery Job Wizard for VMware Cloud Director: Destination" on page 1084
- "Recovery Job Wizard for VMware Cloud Director: Options" on page 1086

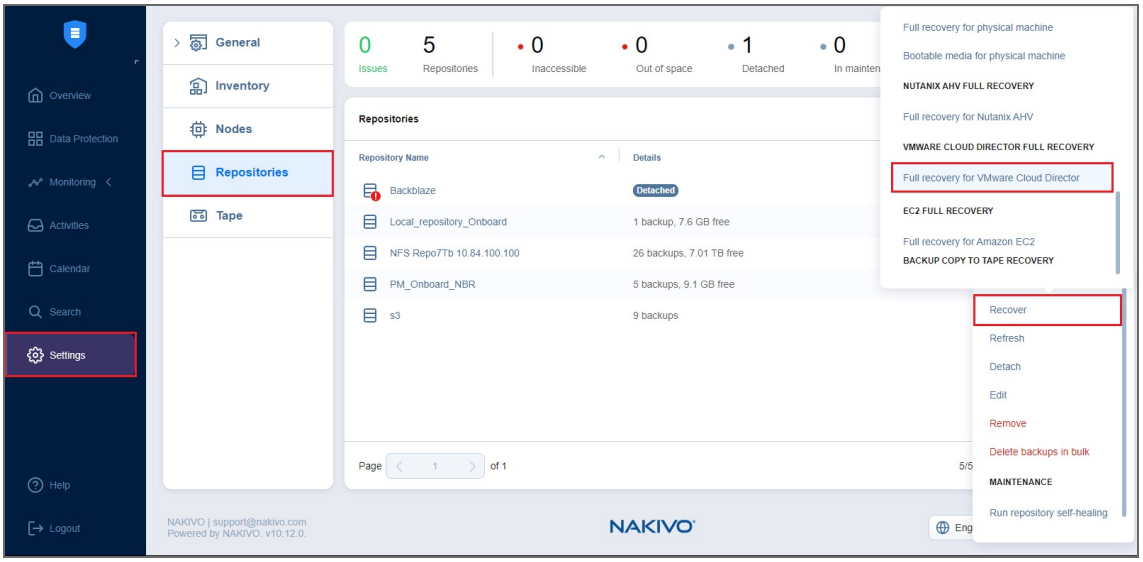
### Starting VMware Cloud Director Recovery

To recover all VMware Cloud Director objects from backups, do one of the following:

- Start recovery from the **Data Protection** menu by clicking **Recover** and then clicking **Full recovery for VMware Cloud Director**.




- Open the **New Full Recovery Job Wizard** from the **Repositories** tab by following the steps below:
    1. From the main menu of NAKIVO Backup & Replication, click **Settings**.
    2. Go to the **Repositories** tab.
    3. Hover over the **Backup Repository** containing the needed backup, click the ellipsis **Manage** button, and click **Recover**.
    4. Click **Full recovery for VMware Cloud Director** under **VMware Cloud Director Full Recovery**.




- Open the **New Full Recovery Job Wizard** from the **Tape** tab by following the steps below:
    1. From the main menu of NAKIVO Backup & Replication, click **Settings**.
    2. Go to the **Tape** tab and click **Backups**.
    3. Select the necessary VMware Cloud Director backups.
    4. Click the **Recover** button.
- Alternatively, the recovery can be performed by using the search function.

The **New Full Recovery Job Wizard for VMware Cloud Director** opens.

# Recovery Job Wizard for VMware Cloud Director: Backups

1. On the **Backups** page of the wizard, select a backup using either a **Backup Repositories** or **Jobs & Groups** view in the left pane.

Both federated and standalone (not used as members of federated repositories) backup repositories can be selected.

**Notes**

- You cannot select a federated repository member as a source for recovery jobs.
- You cannot perform recovery from recovery points located in an inaccessible federated repository member or those that depend on inaccessible recovery points.

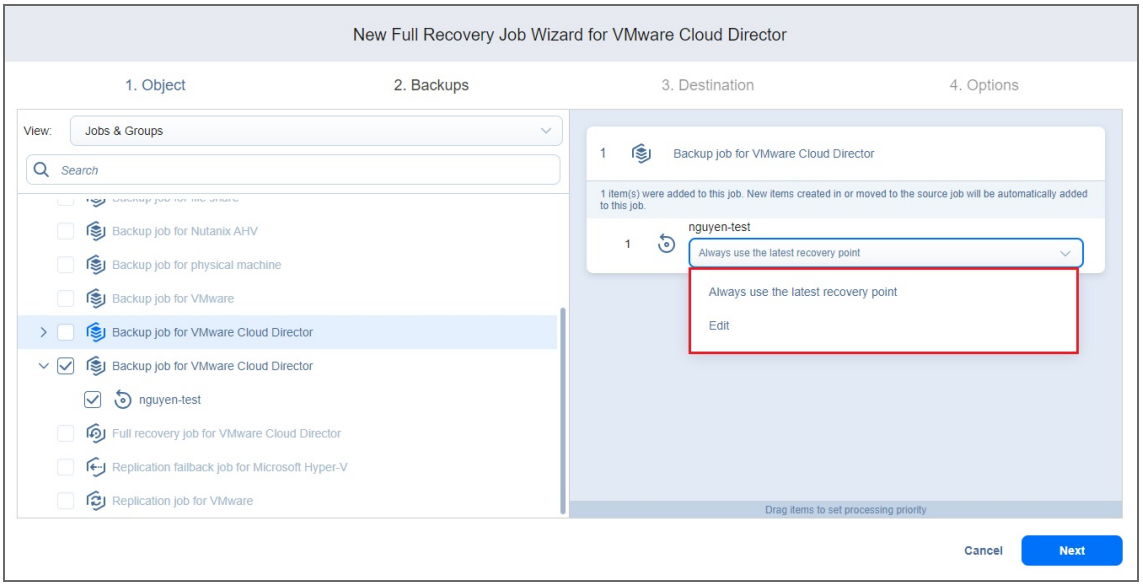
Selecting an item in the tree adds the parent item and all children items to the right pane.

**Notes**

- You can search for a specific backup, job, group, or repository (depending on the selected view) by entering its name into the **Search** field.
- Selecting a single backup object opens the **Calendar** view of the date picker if this view was the last user-selected view. Selecting a single backup object opens the **Table** view if this was the last view the user selected.
- You can select a backup object with the corrupted, inaccessible, or pending removal recovery points, but you cannot select the backup objects of the detached repository.

2. In the drop-down list of the selected item on the right pane, select one of the options:

- **Always use the latest recovery point** (displayed by default).
- **Edit** to open the **Recovery points picker** dialog box in the **Calendar** or **Table** view and pick a date.

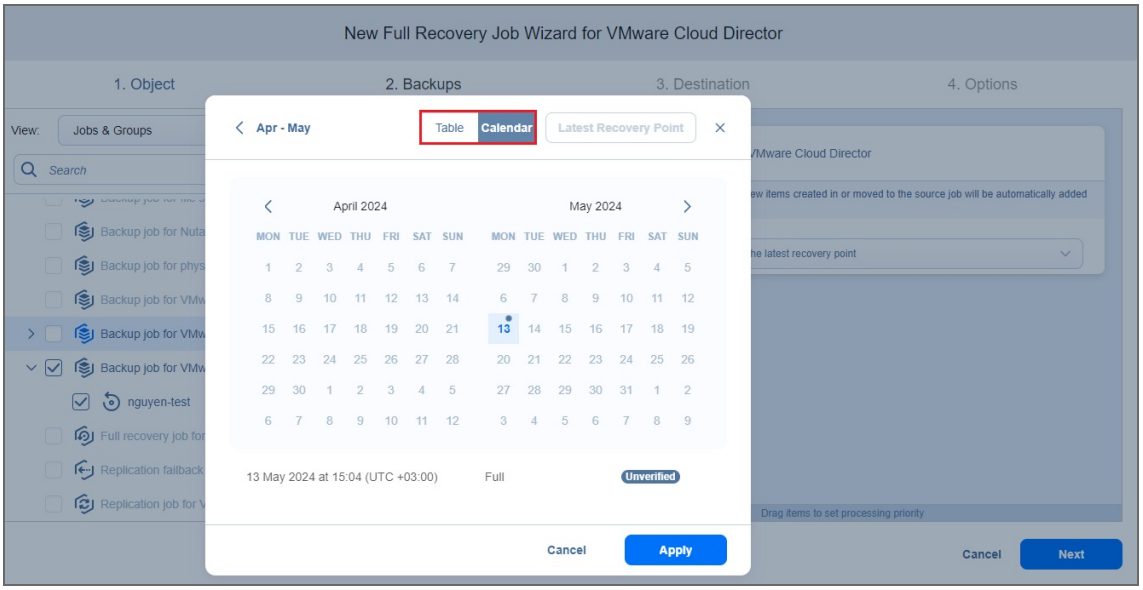


**Notes**

- If a recovery point is selected in the **Calendar** or **Table** view, this recovery point is selected by default the next time you open the **Calendar** or **Table** view.
- The selected view, either **Calendar** or **Table**, is saved on a per-user basis.

3. If necessary, toggle between the **Calendar** and **Table** views:

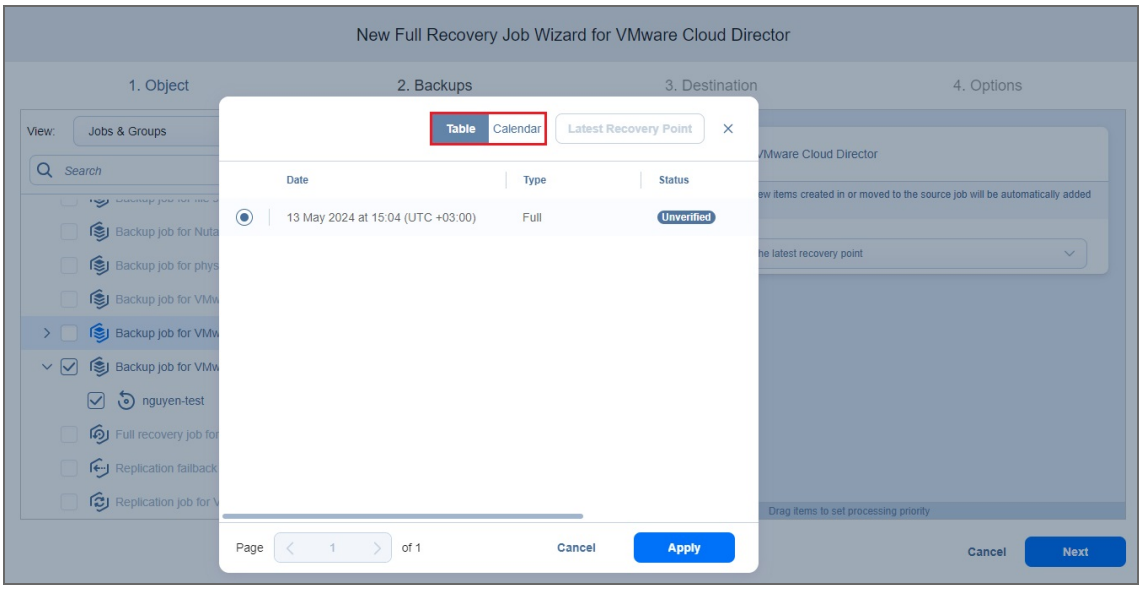
- In the **Calendar** view, select a date or click **Latest Recovery Point** to select the date with the latest recovery point.




**Notes**

- A small dot is displayed on top of the date if there is at least one recovery point created on that date.
- The selected recovery point has a preview displayed under the calendar showing the recovery point status:
- **green** – verified recovery points
- **red** – inaccessible/corrupted/pending removal recovery points
- **dark grey** – unverified recovery points
- Corrupted, inaccessible, or pending removal recovery points cannot be recovered.
- The selected date is highlighted.

- In the **Table** view, select a date or click **Latest Recovery Point** to select the date with the latest recovery point.




**Notes**

- Clicking the **Latest Recovery Point** option selects the date with the latest recovery point that is not corrupted, inaccessible, or pending removal.
- The **Latest Recovery Point** option is disabled if:
  - The latest recovery point is currently selected.
  - All the recovery points of the backup object are corrupted, inaccessible, or pending removal.
- If the selected recovery point is encrypted and the corresponding password hash is not available, the **Always use the latest recovery point** or the recovery point timestamp is highlighted in red with a red lock icon. To add such a recovery point to the job, you need to provide the password manually. See Providing Passwords for encrypted recovery points.

4. Click **Next** to go to the next page.

**Note**

You cannot proceed to the next step if there is at least one selected recovery point to which a password has not been provided.

# Recovery Job Wizard for VMware Cloud Director: Destination

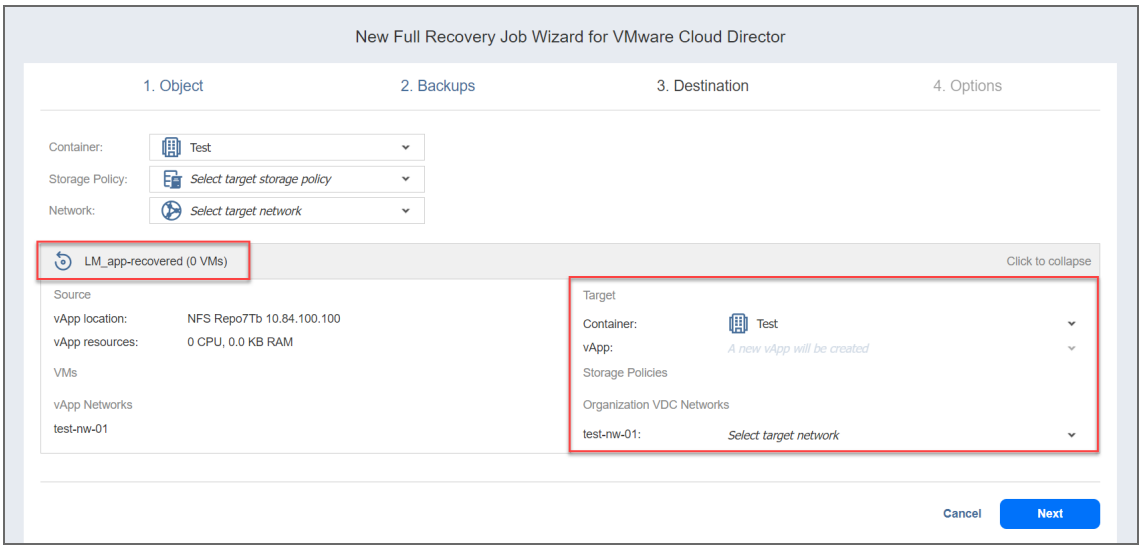
On the **Destination** page of the wizard, select a location for the recovered objects. Proceed as follows:

1. Choose the VMware Cloud Director server, underlying Organization or Virtual Datacenter from the **Container** drop-down list.

**Note**

If you chose a tape backup in the **Backups** step of the wizard, the **Backup Repository** is automatically selected in the **Recover to** drop-down list. Choose the required repository in the **Repository** drop-down list below.

2. Choose a vApp from the **vApp** drop-down list. Note that this option is disabled if you selected **vApp** on the Recovery Job Wizard for VMware Cloud Director: Object page.
3. Choose a network from the **Network** drop-down list.
4. Select the **Storage Policy** from the drop-down list.
5. Optionally, select **Advanced setup** to configure the destination options for each separate recovered object. Proceed as follows:
   - If you selected **Virtual machine** on the **Recovery Job Wizard for VMware Cloud Director: Object** page, configure the options described above for each individual VM.
   - If you selected **vApp** on the **Recovery Job Wizard for VMware Cloud Director: Object** page, select the following:
     a. Select the container for each vApp.
     b. Choose the target organization VCD network for each source organization VCD network.
     c. If the vApp contains VMs, choose a storage policy for each individual VM.
6. Click **Next** to proceed.



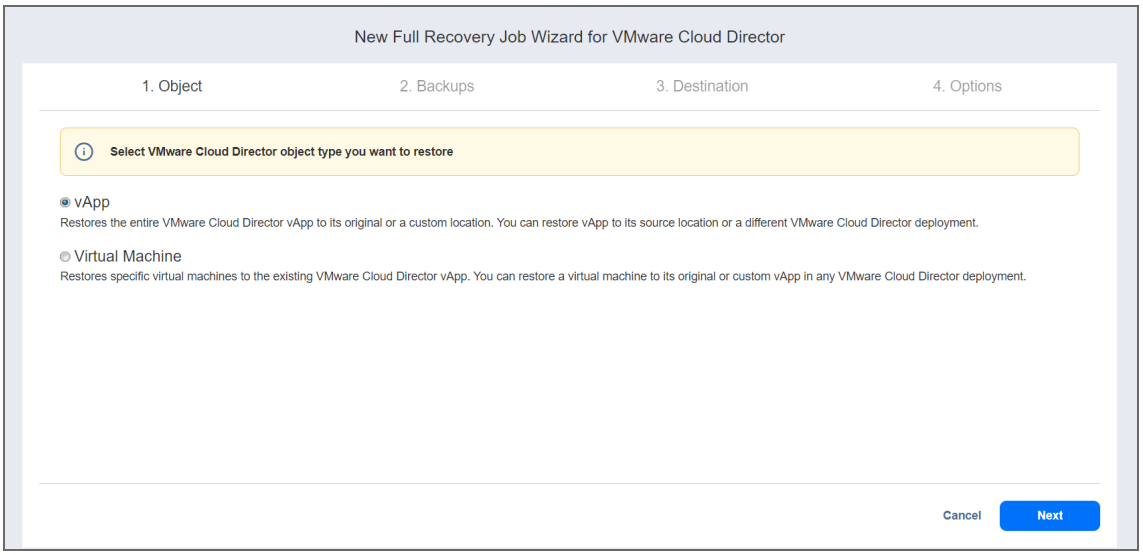

# Recovery Job Wizard for VMware Cloud Director: Object

On this page of the wizard you can select which VMware Cloud Director object type you want to recover. Choose one of the following:

- **vApp**: Choosing this option allows you to recover the entire VMware Cloud Director vApp to its original or a custom location.

- **Virtual machine**: Choosing this option allows you to recover specific VMs to their original or a custom location.

Click **Next** to proceed.




# Recovery Job Wizard for VMware Cloud Director: Options

On the **Options** page of the wizard, set the recovery job options. Proceed as described in these sections:

- Job Options
- Recovered Object Options
- Pre and Post Job Actions
- Data Transfer
- “Completing the New Full Recovery Job Wizard for VMware” on page 1093

## Job Options

Specify the general options as follows:

- **Job name**: Specify a name for the recovery job.
- **Job priority**: Select a job priority level between *1* and *5*, with *1* being the highest priority. Jobs with higher priority levels are prioritized by **Transporters** during job processing.

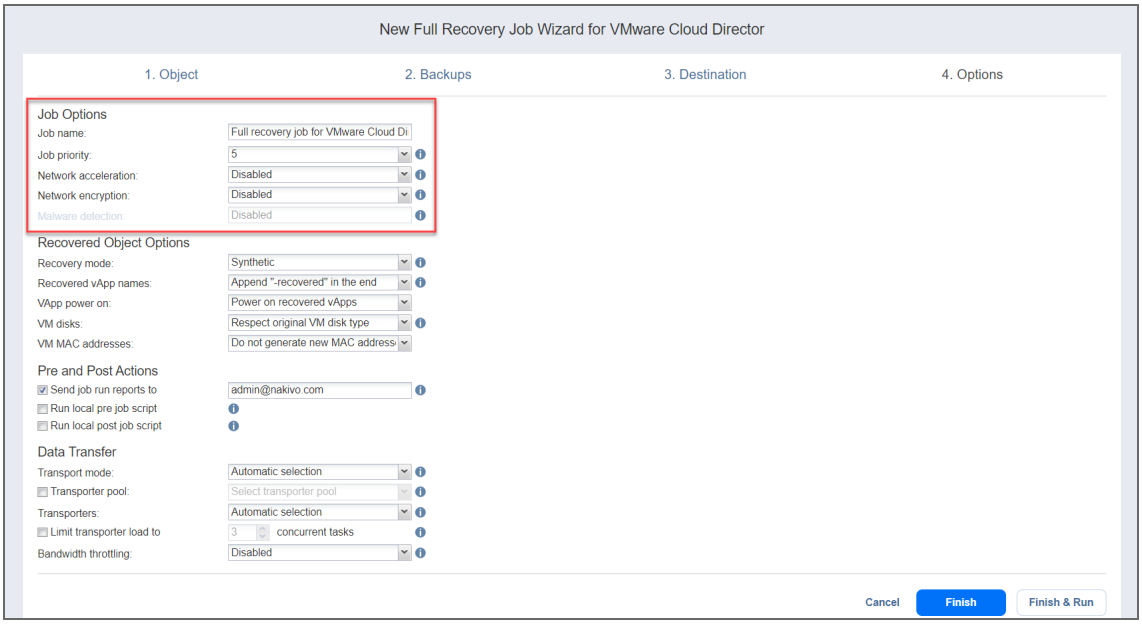
> **Note**
> This option is only available in the Enterprise, Enterprise Essentials, Enterprise Plus, MSP Enterprise, and MSP Enterprise Plus editions.

- **Network acceleration**: With network acceleration enabled, NAKIVO Backup & Replication uses compression and traffic reduction techniques to speed up data transfer. Enable this option if you plan to recover objects over WAN or slow LAN links.

- **Network encryption**: With network encryption enabled, the data is protected with AES 256 encryption while traveling over the network. Data encryption increases the backup time and CPU load on machines running **Transporters**. Select this option when recovering over WAN without a VPN connection.
- **Malware detection** With this option enabled, the backups are scanned for malware using the configured antivirus software on the scan server.
- Optionally, if you have selected **Enabled** for the **Malware detection** option, click the **settings** link to configure the following options:
    - **Scan server**: Select a specific scan server for the job or leave the **Default** setting. If **Default** is selected, the Transporter is used as the scan server and can support a maximum of 2 concurrent scan tasks.

      **Notes**
      - For the **Default** option, if the **Repository Transporter** is the installed **Transporter**, it require the master password to function as the scan server.
      - For more details on the requirements for Scan Server, refer to the Feature Requirements.

    - **Scan type**: Choose between the **Deep scan** and the **Quick scan**:
        - **Deep scan**: When this option is selected, the antivirus software scans the entire backup and may take longer to complete.
        - **Quick scan**: When this option is selected, the antivirus software scans only OS disks in the backup.
    - **If malware is detected**: Choose the behavior if malware is detected:
        - **Fail the recovery job**: With this option, the recovery process fails in case the job has only one VM/vApp. If the job has several VMs/vApps, the infected VMs/vApps are skipped and the job continues to run.
        - **Continue and recover to isolated network**: When this option is selected, the recovery job completes the scanning process and recovers the infected VMs/vApps to a temporary isolated network.
    - **Scan timeout**: Specify the timeout for the malware detection process. If the specified amount of time is exceeded, the recovery job fails.
    - Click **Apply** when you're done.



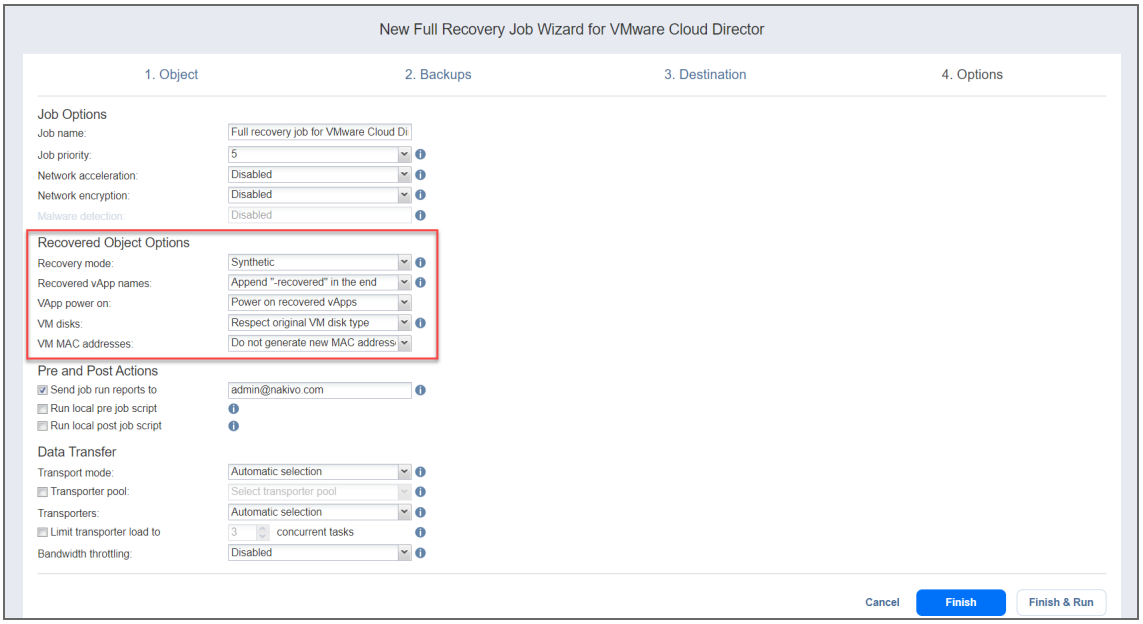

## Recovered Object Options

Specify the recovery options as follows:

1. **Recovery mode**: Choose one of the following:
    - **Synthetic**: With this recovery mode, the objects are recovered with the environmental dependencies (such as CPU affinity) removed. Select this option when recovering objects to a new location.
    - **Production**: With this recovery mode, environment dependencies are preserved on recovered objects. Make sure the location to which the objects are recovered does not contain the original objects, otherwise, UUID and MAC address conflicts may occur.
2. **Recovered VM/vApp names**: Choose one of the following:
    - **Append "-recovered" in the end**: Source VM/vApp names are used for recovered VM/vApp names and “-recoverd” is added after the name of the recovered VMs/vApps.
    - **Leave recovered VM/vApp names as is**: Recovered VM/vApp names are identical to the source VM/vApp names.
    - **Enter custom recovered VM names**: You can enter custom names for recovered VMs/vApps.
3. **VM disks**: Choose one of the following:
    - **Respect original VM disk type**: When specified, disks that respect the original VM disk type are created on target VMs. Select this option to recover VMs to their original location.
    - **Create only thin disks on target VMs**: When specified, only thin disks are created on target VMs. Use this option to save space on the target datastore.
4. **VM MAC addresses**: Choose one of the following:

- **Do not generate new MAC addresses**: When this option is selected, the recovered VM has the same MAC address as the source VM.
- **Generate new MAC addresses**: When this option is selected, a new MAC address is generated for the recovered VM.

5. **VM/vApp power on**: When the **Power on recovered VM/vApp** option is selected, the recovered VMs/vApps are powered on.




## Pre and Post Job Actions

NAKIVO Backup & Replication allows you to run a script before recovery begins (a pre-job script) and after the recovery of all objects in the job has completed (a post-job script). The scripts can only be executed on the machine on which the **Director** is installed. For more information about pre and post job scripts, refer to "Pre and Post Job Scripts" on page 91.

### Setting a Pre Job Script

To run a script before NAKIVO Backup & Replication begins recovering objects, do the following:

1. Place a script file on the machine on which the **Director** is installed.
2. Select the **Run local pre job script** option and click the **settings** link.
3. Specify the following parameters in the dialog box that opens:
   - **Script path**: Specify a local path to the script on the machine on which the **Director** is installed. A script interpreter should be specified.

     **Example (Windows)**: cmd.exe /c D:\script.bat
     **Example (Linux)**: bash /root/script.sh

   - **Job behavior**: Choose one of the following job behaviors in relation to script completion:

- **Wait for the script to finish**: With this option selected, object recovery is started only after the script is completed.
- **Do not wait for the script to finish**: With this option selected, the product runs the script and starts recovering objects at the same time.

- **Error handling**: Choose one of the following job behaviors in relation to script failure:
  - **Continue the job on script failure**: With this option selected, the job performs object recovery even if the script has failed.
  - **Fail the job on script failure**: With this option selected, if the script fails, the job is failed and object recovery is not performed.

Setting a Post Job Script

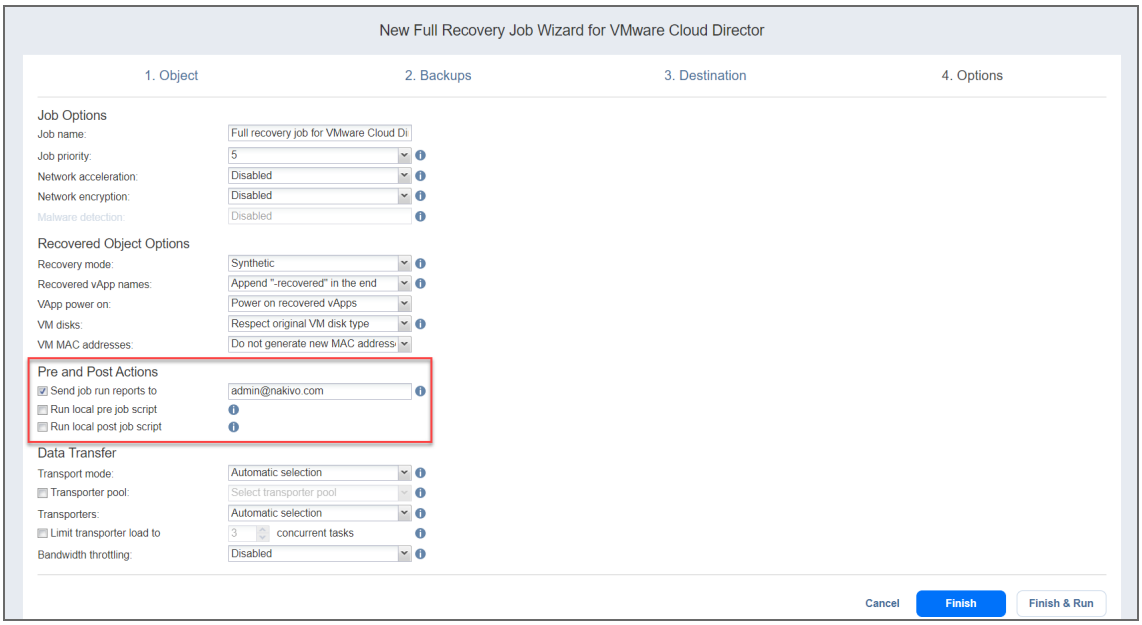
To run a script after NAKIVO Backup & Replication finishes recovering all objects, do the following:

1. Place a script file on the machine on which the **Director** is installed.
2. Select the **Run local post job script** option and click the **settings** link.
3. Specify the following parameters in the dialog box that opens:
   - **Script path**: Specify a local path to the script on the machine on which the Director is installed. A script interpreter should be specified.

     **Example (Windows)**: cmd.exe /c D:\script.bat
     **Example (Linux)**: bash /root/script.sh

   - **Job behavior**: Choose one of the following job behaviors in relation to script completion:
     - **Wait for the script to finish**: With this option selected, the job is in the “running” state until the script is completed.
     - **Do not wait for the script to finish**: With this option selected, the job is completed even if the script execution is still in progress.
   - **Error handling**: Choose one of the following job behaviors in relation to script failure:
     - **Continue the job on script failure**: With this option selected, script failure does not influence the status of the job.
     - **Fail the job on script failure**: With this option selected, if the script fails, the job status is

set to “failed” even if object recovery is successful.




### Email Notifications

NAKIVO Backup & Replication can send email notifications about job completion status to specified recipients. This feature complements global notifications and allows you to configure notifications on a per-job level.

> **Note**
> To enable this option, make sure that Email settings are configured. Refer to “Email Settings” on page 419.

To receive email notifications, do the following:

1. In the *Pre and Post Actions* section, select the **Send job run reports** to.
2. Specify one or more email addresses in the text box. Use semicolons to separate multiple email addresses.

## Data Transfer

In the *Data Transfer* section, choose a transport mode for writing data, select which Transporter to be used for reading data from the source VM, and set bandwidth throttling.

### Transport Mode

NAKIVO Backup & Replication provides the following transport modes for writing VM/vApp data:

- **Automatic**: When this option is chosen, Hot Add mode is used where possible. If the product cannot use Hot Add, LAN mode is used.

- **HotAdd only**: NAKIVO Backup & Replication can write data directly to the datastore, bypassing the network, which can significantly improve job performance. This is achieved with the help of VMware's HotAdd technology. In order for the Hot Add feature to work, the target **Transporter** (the one that will be writing data ) should run on a host that has access to the target datastore(s).
- **LAN only**: Data will be written over LAN.

Transporter Pool

If this option is enabled, only the transporters that belong to the selected transporter pool shall be used during the job run.

Transporters

By default, NAKIVO Backup & Replication automatically determines which **Transporter** should be used to read data from the source VM. However, you can manually specify which **Transporters** should be used for the job:

- **Automatic selection**: The product automatically determines the **Transporters** that are closest to the source and target hosts.
- **Manual - configured for all VMs**: Select this option to manually specify a single source and a single target **Transporter** to be used for data transfer by the job.
- **Manual - configured per host**: Select this option to manually specify **Transporters** for all source and target hosts.

Transporter Load

You can limit the maximum number of **Transporter** tasks assigned to it within the job. By default, this number is set to 3 concurrent tasks.

To change the default number of tasks, do the following:

- In the *Data Transfer* section, select the **Limit transporter load** to checkbox.
- Specify the number of concurrent tasks in the corresponding box.

Bandwidth Throttling

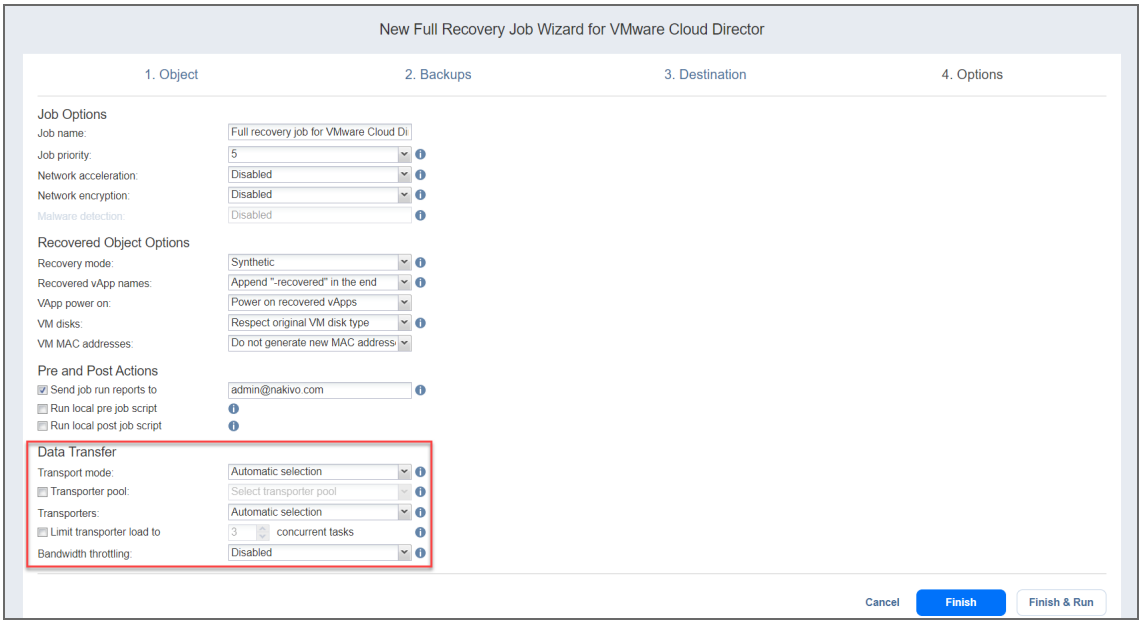
Follow the steps below to set the speed of data transfer over the network for your recovery job:

1. For the **Bandwidth throttling** option, choose **Enabled**.

> **Note**
> If bandwidth throttling is disabled for the current job, global bandwidth rules may still apply to your job.

2. Click the **settings** link that becomes available. The **Job Bandwidth Rules** dialog box opens displaying the list of available rules. You have the following options:

- Create a new bandwidth rule for your recovery job:
  a. Click the **Create New Rule** button.
  b. The New Bandwidth Rule dialog box opens. For details on creating a bandwidth rule, refer to "Bandwidth Throttling" on page 407.
  c. Click **Save**.
- Activate an existing bandwidth rule for your job: Select the checkbox to the left of the required bandwidth rule. To deactivate a bandwidth rule for your job, deselect the corresponding check-box.
- Edit a bandwidth rule: Click the **Edit** link for a bandwidth rule and modify it in the **Edit Bandwidth Rule** dialog box that opens.
- Disable a bandwidth rule: Click the **Disable** link. The bandwidth rule is disabled for all jobs.
- Remove a bandwidth rule: Click the **Remove** link and then click **Delete** to confirm your operation.




## Completing the New Full Recovery Job Wizard for VMware

Click **Finish** or **Finish & Run** to complete the job creation.

> **Note**
>
> If you click **Finish & Run**, you will have to define the scope of your job. Fore details, refer to "Running Jobs on Demand" on page 364.

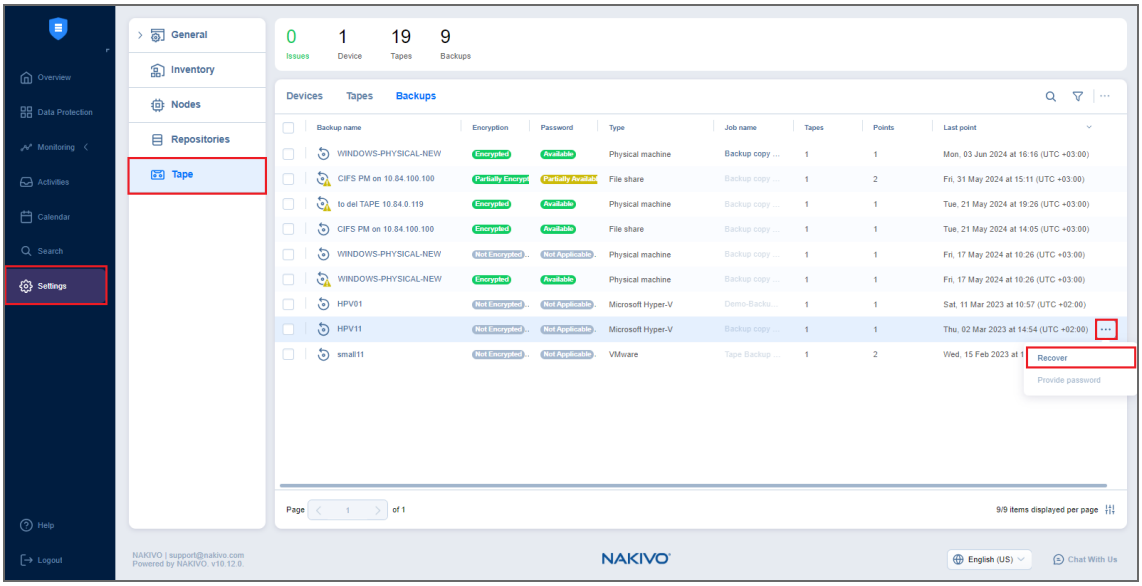
# Starting Recovery from Tape

To initiate recovery from a tape backup, do the following:

# Starting Recovery from the Tape Tab

1. Go to **Settings**, click the **Tape** tab, and open the **Backups** view.
   - In the **Backups** table, select one or several backups that you want to recover, click the ellipsis **Manage** button and select **Recover**. This opens the recovery wizard with specified backups and their latest recovery point selected.

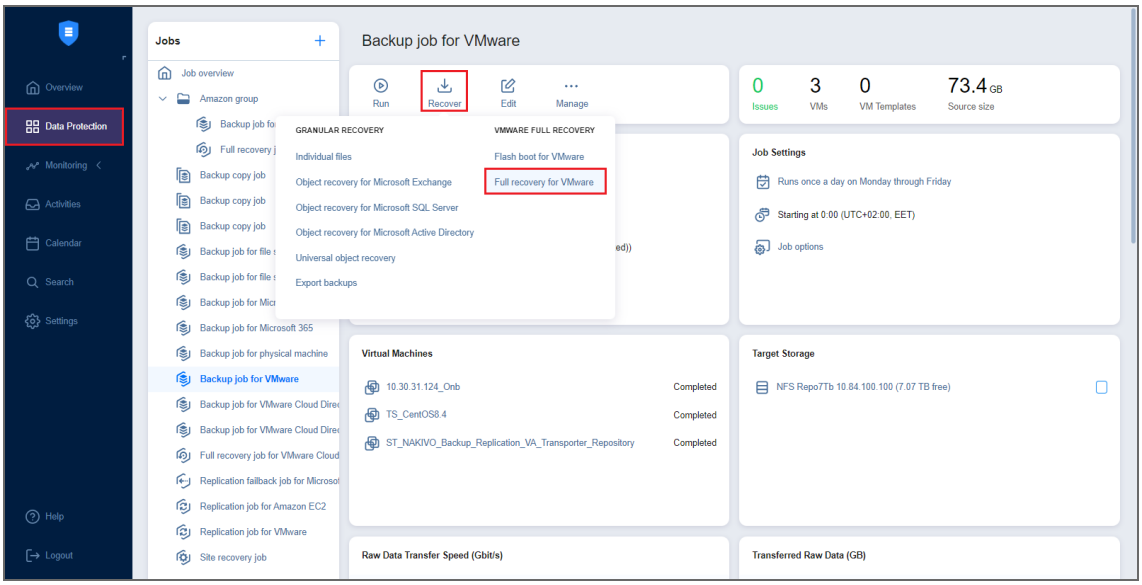
> **Note**
> Only backups of the same type can be selected. That is, you cannot select VMware and Hyper-V type backups and launch the **Recovery wizard**.




2. Alternatively, go to the “Managing Tape Cartridges” on page 749 page, select a backup in the **Tape contents** pane and then click the recovery point you want to restore from.

# Starting Recovery from the Data Protection Dashboard

You may also initiate recovery from a tape backup via the **Data Protection** dashboard. To do so, click **Recover** and select the appropriate recovery object and type. For example, to launch the VMware recovery wizard, click **Full recovery for VMware** under **VMware Full Recovery**.




Then, proceed as described in the following topics:

- “Recovery Job Wizard for VMware: Backups” on page 1064
- “Recovery Job Wizard for VMware Cloud Director: Backups” on page 1082

# Performing Cross-Platform Recovery

With the Cross-Platform Recovery feature of NAKIVO Backup & Replication, you can export backups to standard formats that are compatible with various platforms. The following formats are supported:

- VMDK for disk(s) of VMware VMs, Proxmox VE VMs, Nutanix AHV VMs, and AWS EC2 instances
- VHD and VHDX for disk(s) of Hyper-V VMs and Nutanix AHV VMs

To export your backup for subsequent recovery on the same platform or a different one, use the Backup Export Wizard in NAKIVO Backup & Replication. Refer to "Feature Requirements" on page 169 for the supported scenarios for cross-platform recovery.

NAKIVO Backup & Replication does not run VM preparation when exporting the backups into a specific format. If you plan to import the VM into a different platform and VM preparation is required, prepare your VM in advance.

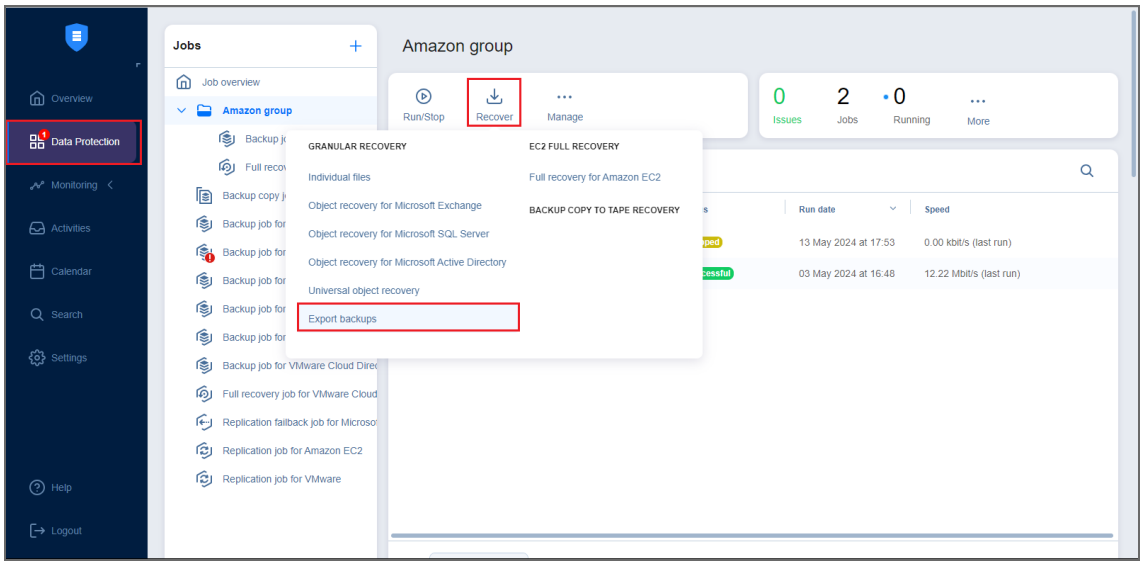
This section includes the following topics:


- "Opening Backup Export Wizard" on page 1097
- "Backup Export Wizard: Backups" on page 1098
- "Backup Export Wizard: Disks" on page 1102
- "Backup Export Wizard: Options" on page 1103
- "Backup Export Wizard: Finish" on page 1105

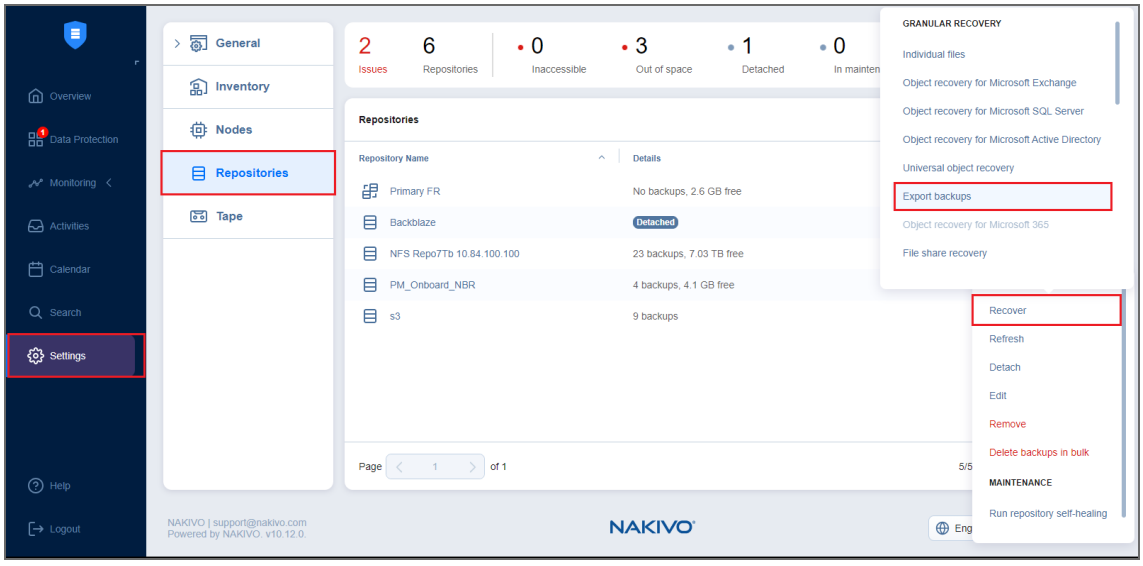
# Opening Backup Export Wizard

Open **Backup Export Wizard** using one of the following ways:

- Navigate to the **Data Recovery** menu, click **Recover** and then click **Export backups**.




- On the **Settings** page:
  1. Click the **Repositories** tab.
  2. In the list of repositories, hover over a repository and click the ellipsis **Manage** button.
  3. Click **Recover** and select **Export backups**.




- Alternatively, the Backup Export can be performed by using by using the Search function.

# Backup Export Wizard: Backups

On the **Backups** page of the wizard:

1. In the left pane, select one or more backups using either a **Backup Repositories** or **Jobs & Groups** view in the left pane.

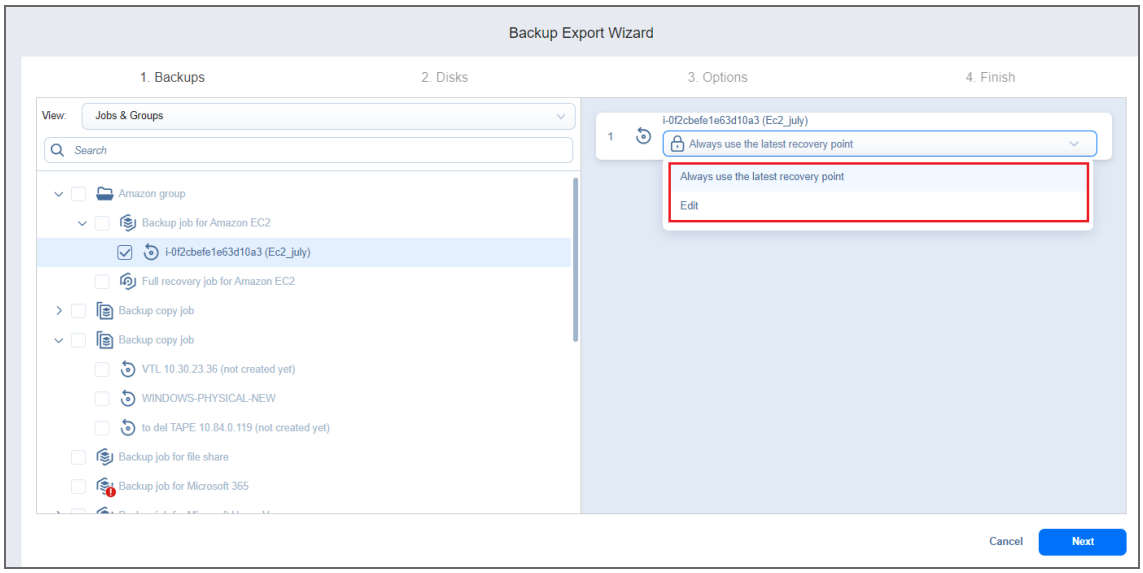
   Both federated and standalone (not used as members of federated repositories) backup repositories can be selected.

   > **Notes**
   >
   > - You cannot select a federated repository member as a source for recovery jobs. You cannot perform recovery from recovery points located in an inaccessible federated repository member or those that depend on inaccessible recovery points.

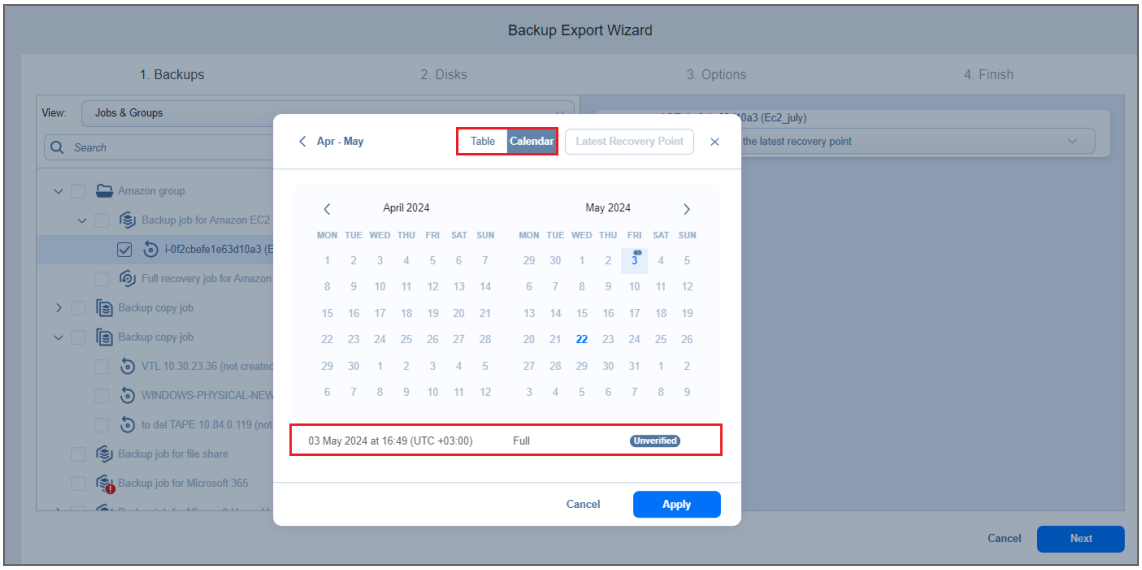
   Selecting an item in the tree adds the parent item and all children items to the right pane.

   > **Notes**
   >
   > - You can search for a specific backup, job, group, or repository (depending on the selected view) by entering its name into the **Search** field.
   > - Selecting a single backup object opens the **Calendar** view of the date picker if this view was the last user-selected view. Selecting a single backup object opens the **Table** view if this was the last view the user selected.
   > - You can select a backup object with the corrupted, inaccessible, or pending removal recovery points, but you cannot select the backup objects of the detached repository.

2. In the drop-down list of the selected item on the right pane, select one of the options:
   - **Always use the latest recovery point** (displayed by default).
   - **Edit** to open the **Recovery points picker** dialog box in the **Calendar** or **Table** view and pick a date.



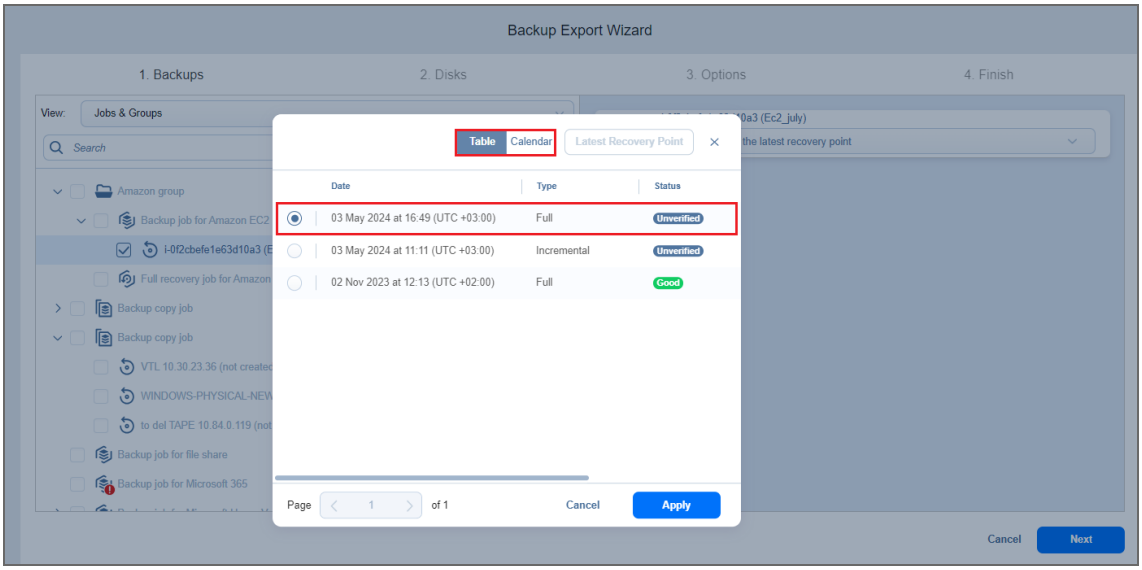

3. If necessary, toggle between the **Calendar** and **Table** views:
    - In the **Calendar** view, select a date or click **Latest Recovery Point** to select the date with the latest recovery point.



**Notes**

- A small dot is displayed on the top of the date if there is at least one recovery point on that date.
- The selected recovery point has a preview displayed under the calendar showing the recovery point status:
  - **green** – verified recovery points
  - **red** – inaccessible/corrupted/pending removal recovery points
  - **dark grey** – unverified recovery points
- The corrupted, inaccessible, or pending removal recovery points cannot be recovered.
- The selected date becomes highlighted.

- In the **Table** view, select a date or click **Latest Recovery Point** to select the date with the latest recovery point.



**Notes**

- Clicking the **Latest Recovery Point** option selects the date with the latest recovery point that is not corrupted, inaccessible, or pending removal.
- The **Latest Recovery Point** option is disabled if:
  - The latest recovery point is currently selected.
  - All the recovery points of the backup object are corrupted, inaccessible, or pending removal.
- If the selected recovery point is encrypted and the corresponding password hash is not available, the **Always use the latest recovery point** or the recovery point timestamp is highlighted in red with a red lock icon. To add such a recovery point to the job, you need to provide the password manually. See Providing Passwords for encrypted recovery points.

4. Click **Next** to go to the next page of the wizard.

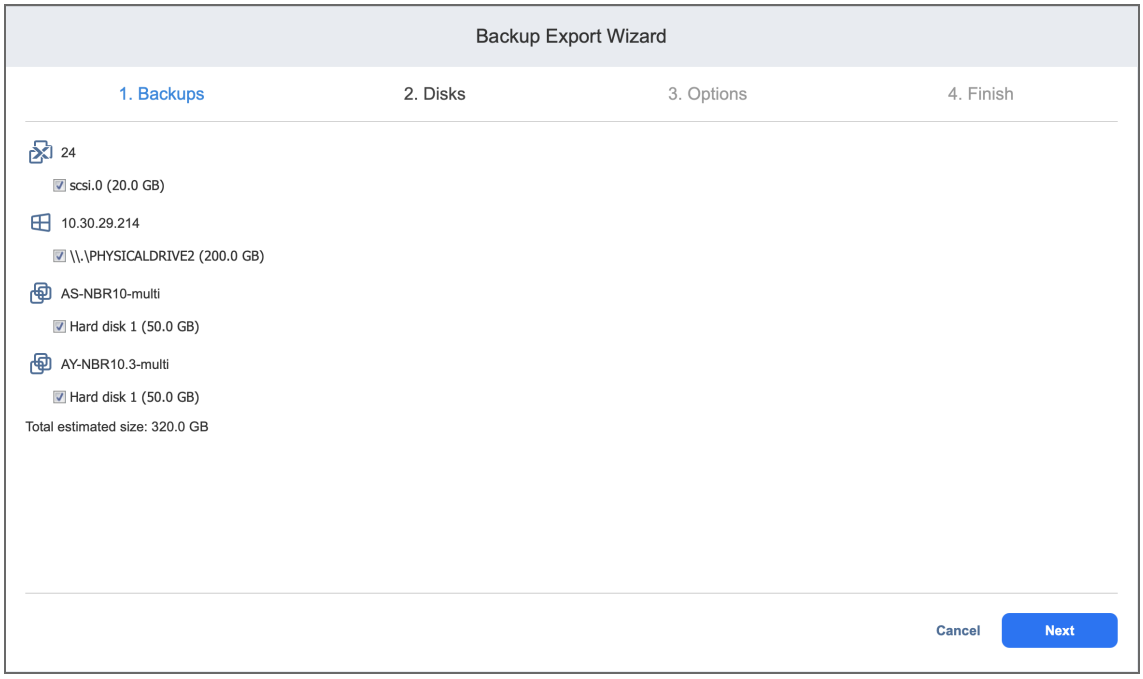
**Note**

You cannot proceed to the next step if there is at least one selected recovery point to which a password has not been provided.

# Backup Export Wizard: Disks

On the **Disks** page of the wizard:

1. Select one or more disks under each backup.
2. Click **Next** to go to the next page of the wizard.



# Backup Export Wizard: Options

On the **Options** page of the wizard, specify options for exporting your backups:

1. **Export format**: Choose one of the following:
    - VMDK
    - VHD
    - VHDX

    **Note**

    VMDK disks are always pre-allocated with the thick provisioning type of storage.

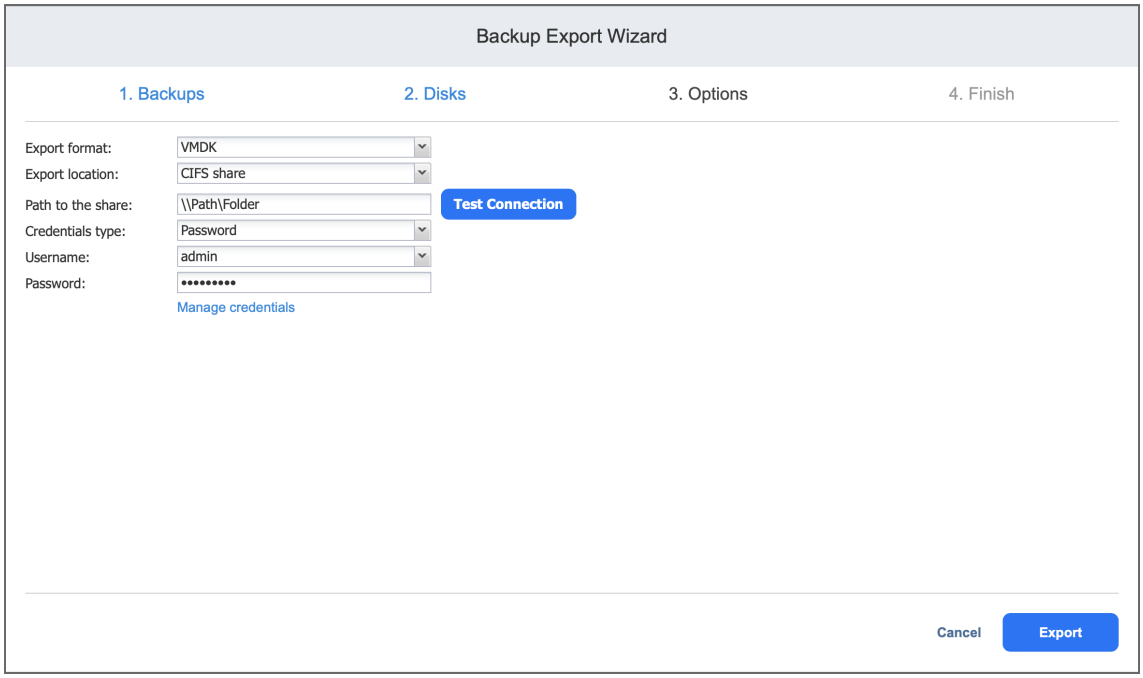
2. **Export location:** Choose one of the following:
    - **Local folder on assigned Transporter**: With this option selected, you have to specify a path to the local folder to which the backups will be exported.
    - **CIFS share**: With this option selected, proceed as follows:
        a. Enter the following:
            - Path to the share
            - Username and Password or Private Key

            **Note**

            Provide the full domain credentials, for instance, “example.com\\smithjohn”.

        b. Click **Test Connection** to check your credentials for the specified share.
    - **NFS share**: With this option selected, proceed as follows:
        a. Enter Path to the share.
        b. Click **Test Connection** to check the connection to the specified share.

3. Click **Export** to go to the next page of the wizard.



# Backup Export Wizard: Finish

The **Finish** page of the wizard informs you that your backup export has started. To view the status of your backup export, go to **Activities**.

To view the backup export progress, go to **Settings** > **General** > **Events**.

To close the **Backup Export Wizard**, click **Close**.

# Planning Disaster Recovery

With NAKIVO Backup & Replication, you can automate disaster recovery workflows for VMware, Microsoft Hyper-V, and Amazon EC2 environments to meet all major disaster recovery planning needs. For more information about disaster recovery (DR), check the “Disaster Recovery” on page 36 section.

The solution allows you to protect VMs running within a cluster, replicate VMs, including creation of real-time replication jobs, perform failover to replicas and replica failback operations.

You can include up to 200 actions in a single Site Recovery job, including failover, failback, start/stop VMs and instances, run/stop jobs, run scripts, attach or detach repositories, send emails, wait, and check conditions. You can check the Site Recovery page to learn more.

In case of failure, you can switch your VMware VM to its latest real-time replica by using the Real-Time Replication Failover feature. For more information on how to manage real-time replica failover jobs, see .

Any disaster recovery job can be created by combining actions and conditions into an automated algorithm. For more details, refer to the corresponding sections below:

- “Failover to Replica for VMware” on page 1107
- “Replica Failback for VMware” on page 1118
- “Real-Time Replication Failover for VMware” on page 1136
- “Site Recovery Job” on page 1146

# Failover to Replica for VMware

With NAKIVO Backup & Replication, you can switch your VMware VM to the VM's latest replica in case of failure.

Refer to the following topics for more information:

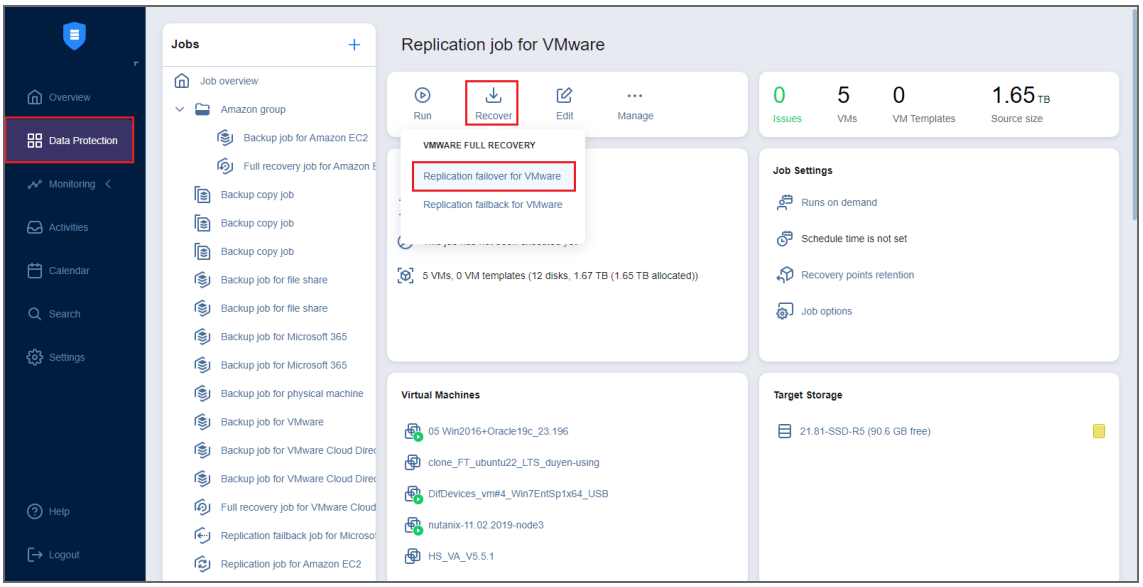

- "Starting VMware Failover to Replica" on page 1108
- "Failover Job Wizard for VMware: Source" on page 1109
- "Failover Job Wizard for VMware: Networks" on page 1110
- "Failover Job Wizard for VMware: Re-IP" on page 1113
- "Failover Job Wizard for VMware: Options" on page 1116

# Starting VMware Failover to Replica

Before starting a replica failover job for VMware, make sure that you have an active VMware replication job available. Refer to “Creating VMware Replication Jobs” on page 934 for details.

To start the procedure of switching a VMware system workload to a backup VM, do the following:

1. Click **Recover** in the **Data Protection** menu.
2. Click **Replication failover for VMware** in the **VMware Full Recovery** section. Note that the VMware replication job must have been completed successfully for this option to be available.
3. Alternatively, select the relevant replication job and click **Recover** > **Replication failover for VMware**.
4. Additionally, the failover can be performed by using the search function.

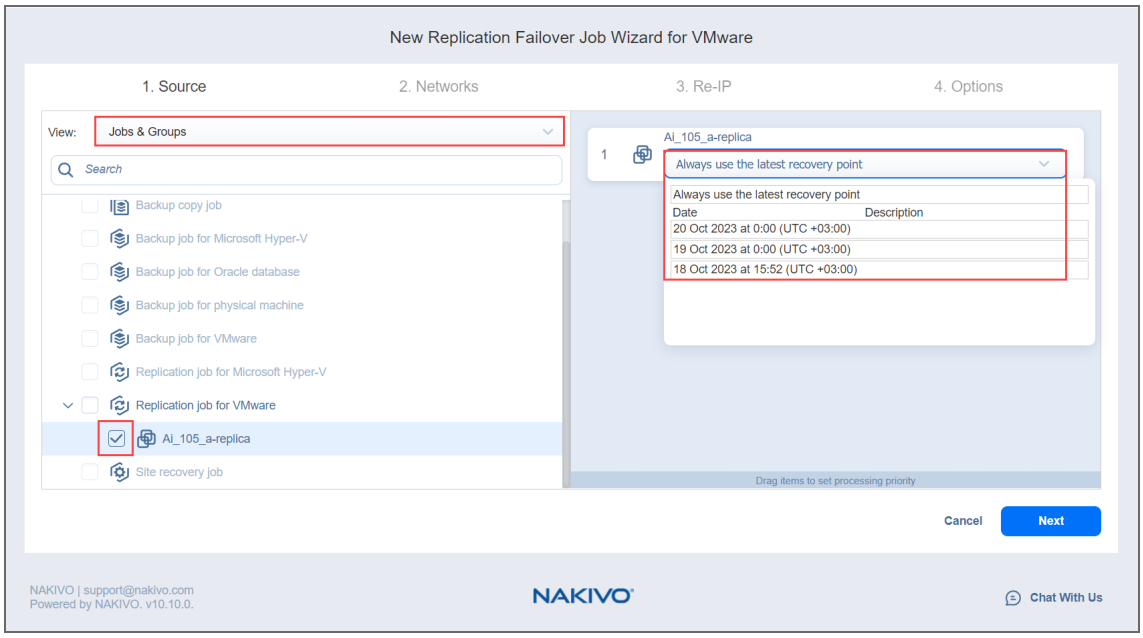



The **New Replication Failover Job Wizard for VMware** opens.

# Failover Job Wizard for VMware: Source

On the **Source** page of the wizard, do the following:

1. In the left pane of the page, choose either of the following inventory views:
   - **Jobs & Groups**: When selected, the inventory tree displays groups, jobs, and backups of the appropriate hypervisor. You can select the required replica from the list of replication jobs.
   - **Policy**: When selected, job policies can be used. Refer to "Managing Job Policies" on page 382 for details. If the items were selected in alternate views, a dialog box opens warning you that switching to the **Policy** view will reset your current selection. Click **Switch View** to confirm switching to the **Policy** view. Make sure that at least one item matches the available set of policy rules. Refer to "Managing Policy Rules" on page 385 for details.
2. Select one or more source VMware VMs in the left pane and then select a recovery point for each source VM in the right pane.




3. Click **Next** to go to the next page of the wizard.

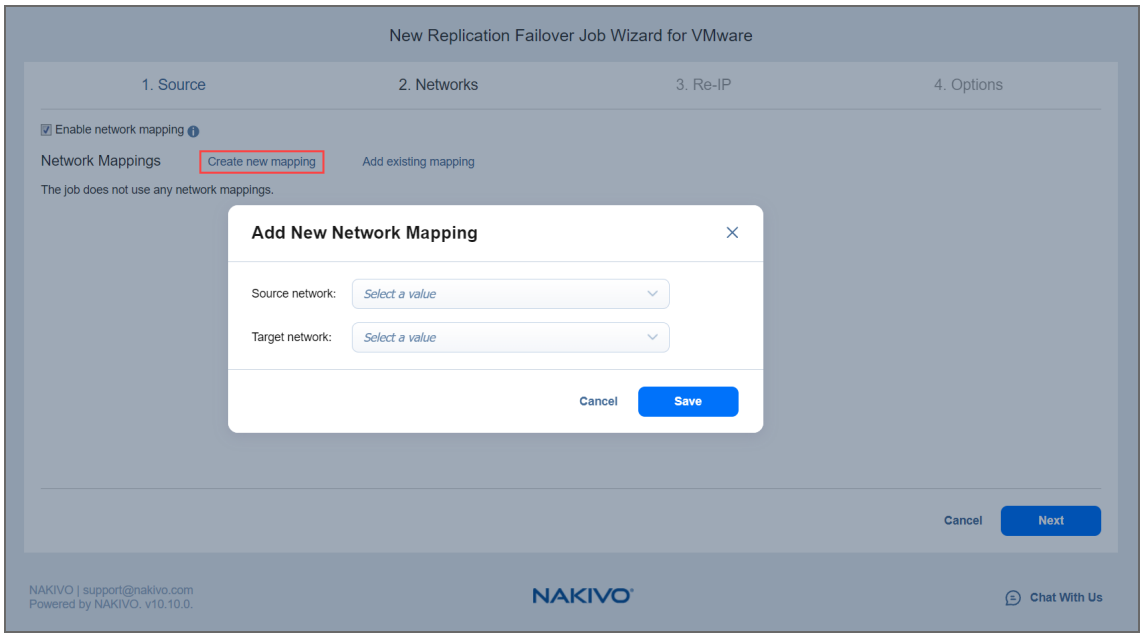
# Failover Job Wizard for VMware: Networks

When the replica network – or target VMware virtual network, – differs from the network address where the source VMs are deployed – or source VMware virtual network, – your failover job needs a relation between these networks to be performed successfully. Such relation is called a network mapping.

> **Note**
> When replicating an ESXi 8.0 VM to an ESXi 7.0 host, the **VirtualMachineConfigSpec** > **VirtualMachineSgxInfo** property is skipped.

To map source VMware virtual networks to appropriate target virtual networks, please do the following on the **Networks** page of the wizard:

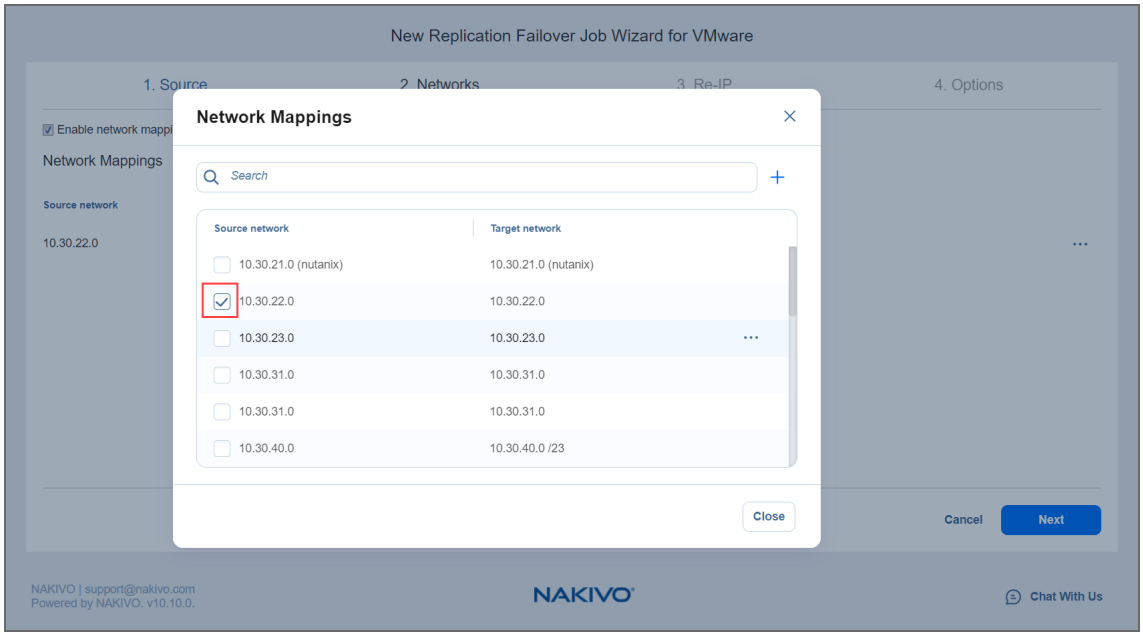
1. Select **Enable network mapping**.
2. The **Network Mappings** dialog box opens. You have the following options:
    - Create a new network mapping:
        a. Click **Create new mapping**.
        b. The **New Network Mapping** dialog opens. Choose a source network and a target network and click **Save**.




    - Add an existing network mapping:

a. Click **Add existing mapping**.

b. The **Network Mappings** dialog opens. Choose one or more appropriate network mappings and close the dialog box.




- Edit an existing network mapping:

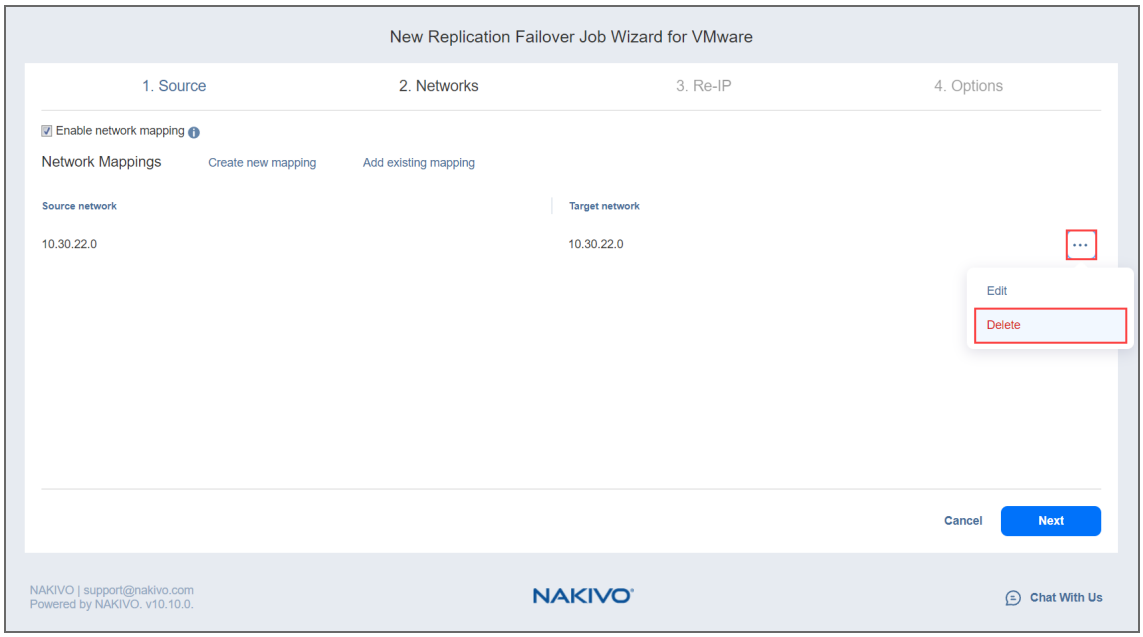
a. Hover over the necessary item in the **Network Mappings** list and then click the ellipsis **Manage** button to the right of the item. Then, click **Edit**.

b. The **Edit Network Mapping** dialog box opens. Choose an appropriate item from the

**Target network** drop-down list and then click **Save**.




- Delete an existing mapping: hover over the necessary item in the **Network Mappings** list and then click the ellipsis **Manage** button to the right of the item. Then, click **Delete**.




- To leave the list of existing network mappings intact, go to the next page of the wizard.

3. Click **Next** to go to the next page of the wizard.

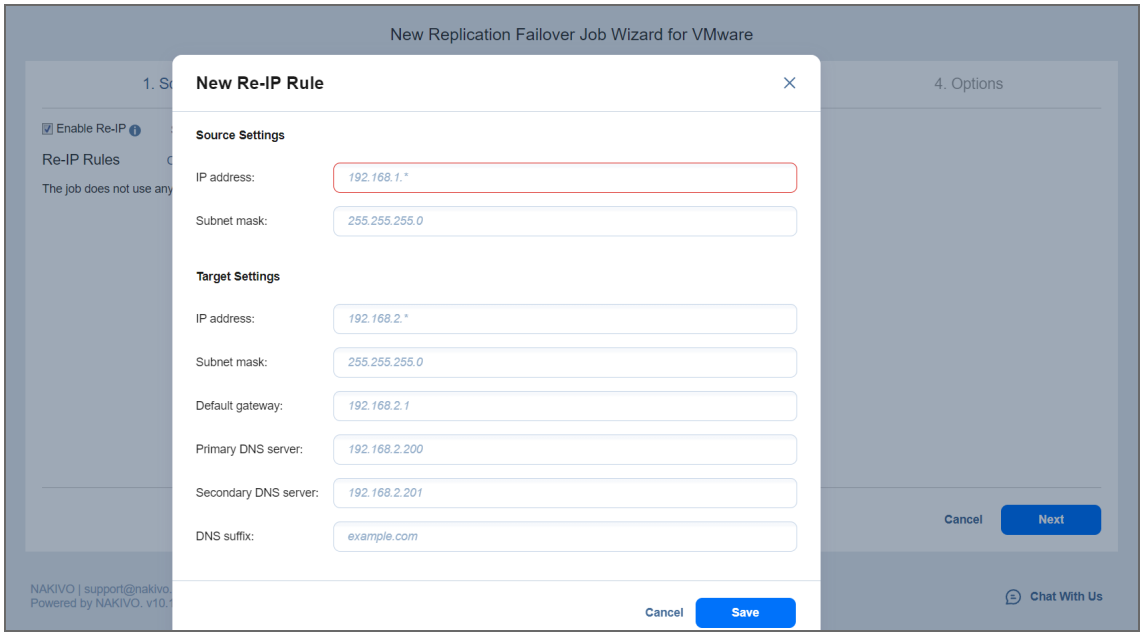
# Failover Job Wizard for VMware: Re-IP

When the IP addressing scheme for the target replica differs from the IP addressing scheme for the VMware source VMs, your failover job needs relations between source VMs addresses and the target replica VMs addresses to be performed successfully. These relations are called re-IP rules.

**Warning**
VMware Tools must be running on source VMs to successfully enable re-IP for your failover job.

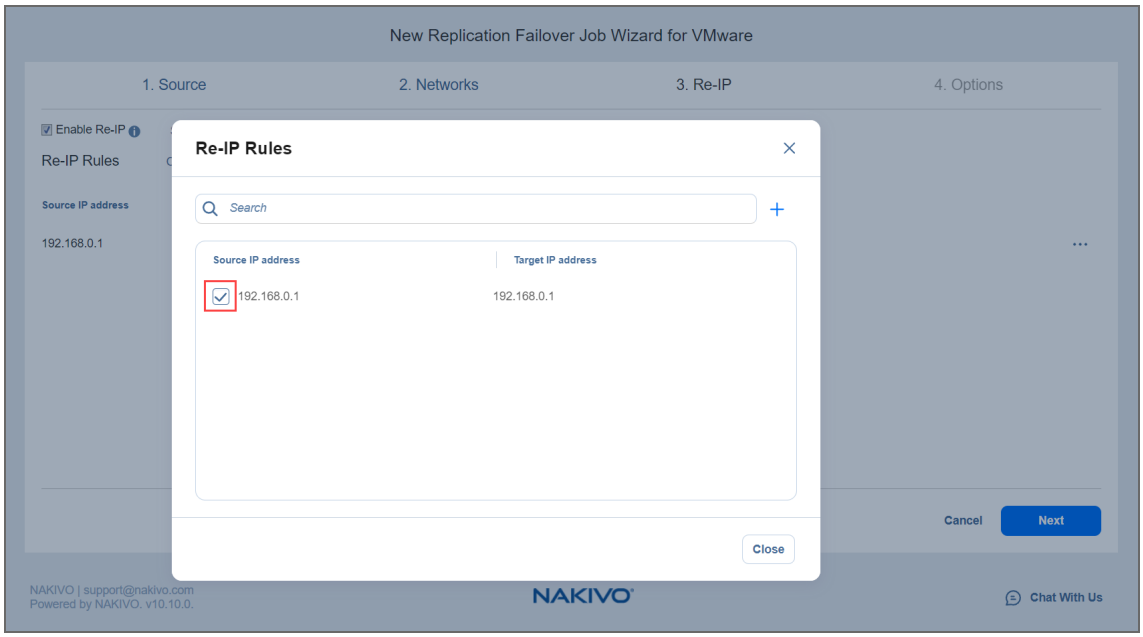
To map a source VM IP address to a specific target IP address, do the following on the Re-IP page of the wizard:

1. Select **Enable Re-IP**.
2. The **Re-IP Rules** section opens. Click the **Select VMs** link.
3. The **Re-IP dialog** box opens. In the list of your source VMs, select at least one VM and close the dialog box.
4. You have the following options:
    - Create a new rule:
        a. Click **Create new rule**.
        b. The **New Re-IP Rule** dialog opens. Enter source and target settings for the Re-IP rule and click **Save**.



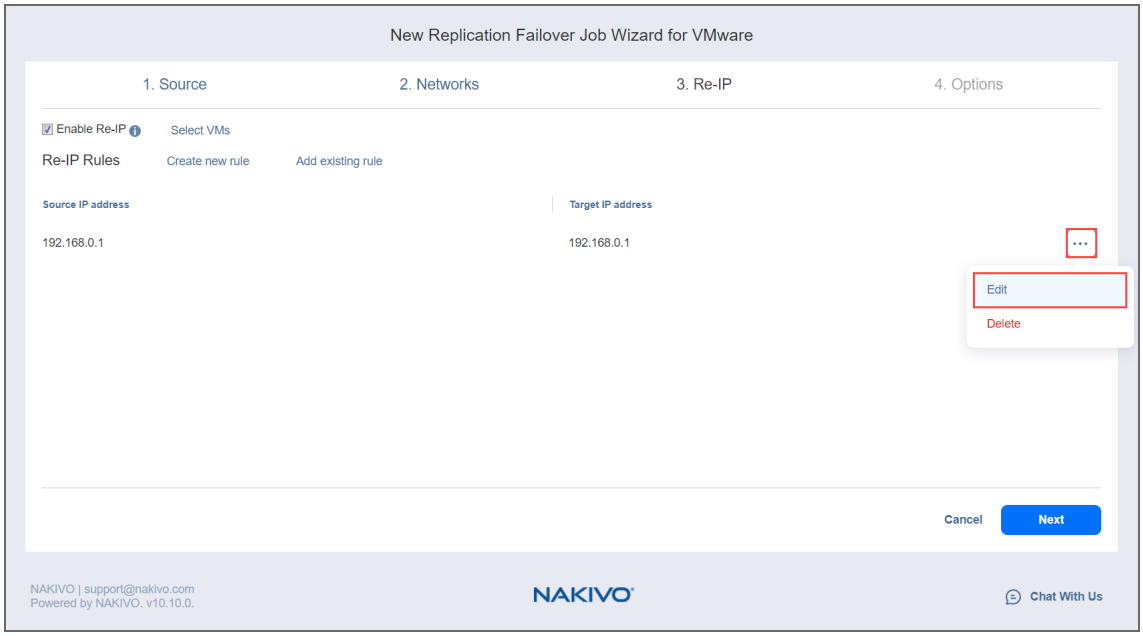
Note that you can use wildcards for IP addresses. For example, if you enter `192.168.1.* -> 10.30.30.*` for the re-IP rule, a source VM IP address such as `192.168.1.50` will be changed to the `10.30.30.50` IP address for your failover job.

- Add an existing rule:
    a. Click **Add existing rule**.
    b. The **Re-IP Rules** dialog opens. Select one or more appropriate Re-IP rules and close the dialog.




- Edit an existing Re-IP rule:
    a. Hover over the required item in the Re-IP Rules list and then click the ellipsis **Manage** button to the right of the item. Then, click **Edit**.
    b. The **Edit Re-IP Rule** dialog box opens. Edit the required properties of the Re-IP rule and

click **Save**.




- Delete an existing mapping: Hover over the required item in the Re-IP Rules list and then click the ellipsis **Manage** button to the right of the item. Then, click **Delete**.




- To leave the list of existing Re-IP rules intact, go to the next page of the wizard.

5. Click **Next** to go to the next page of the wizard.

# Failover Job Wizard for VMware: Options

On this page of the wizard, set the options for the VMware failover job:


- "Job Options" below
- "Pre and Post Actions" below
- "Completing the New Replication Failover Job Wizard for VMware" on the next page


## Job Options

The following failover job options are available for editing:

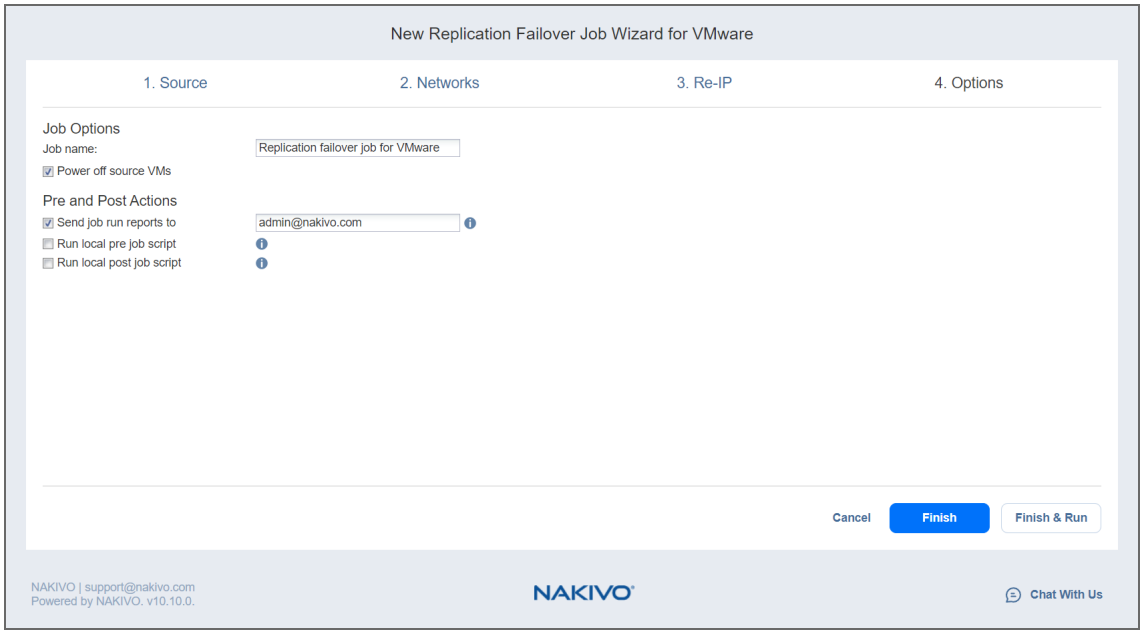
1. **Job name**: A string of 50 or fewer characters specifying the name of the failover job.
2. **Power off source VMs**: Select this checkbox to power off source VMs when the failover job has completed.

## Pre and Post Actions

The following pre- and post-job actions can be set up:

1. **Send job run reports to**: When selected, a job report is sent to specified recipients each time the failover job has completed. This overrides the default setting in the NAKIVO Backup & Replication "Settings" on page 405.
2. **Run local pre job script**: When selected, a dialog box opens in which you can specify the following options of the pre-job script to be executed:
    - **Script path**: A local path to the script on the machine where the **Director** is installed. A script interpreter should be specified as well.
    - **Job behavior**: The following options are available:
        - **Wait for the script to finish**: This is the default option.
        - **Do not wait for the script to finish**.
    - **Error handling**: The following options are available:
        - **Continue the job on script failure**: This is the default option.
        - **Fail the job on script failure**.
3. **Run local post job script**: When selected, a dialog box opens in which you can specify options of the post-job script to be executed on the machine where the product is installed. The options are

similar to the ones of the pre-job script.




## Completing the New Replication Failover Job Wizard for VMware

Click **Finish** or **Finish & Run** to complete the job creation.

> **Note**
>
> If you click **Finish & Run**, you will have to define the scope of your job. Refer to "Running Jobs on Demand" on page 364 for details.

# Replica Failback for VMware

With NAKIVO Backup & Replication, you can switch your VMware VM replicas back to the source or to a new location.

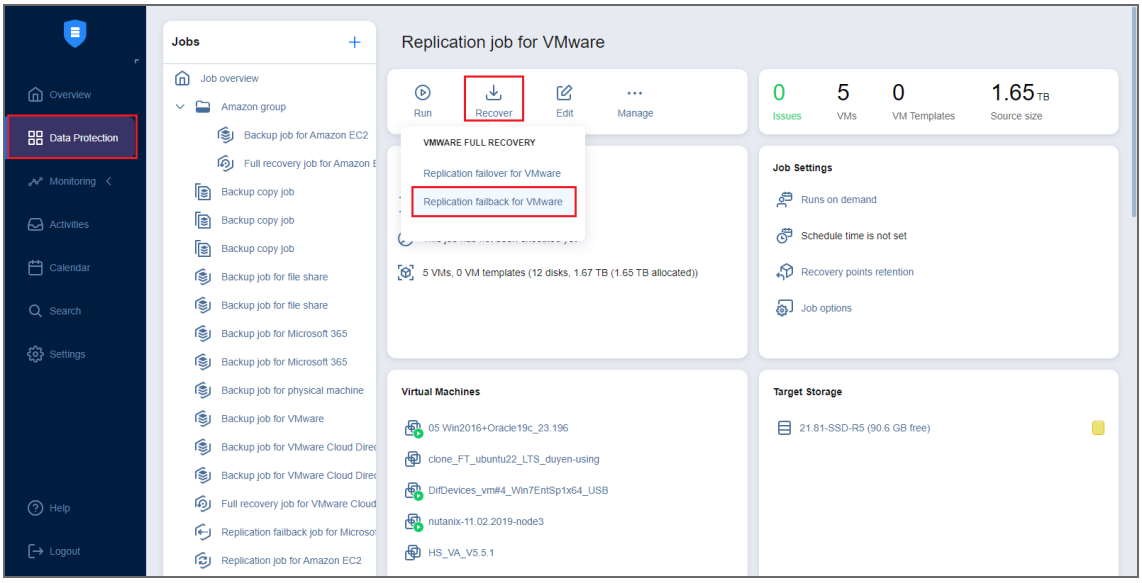
Refer to the following topics for more information:


- "Starting Replica Failback for VMware" on page 1119
- "Failback Job Wizard for VMware: Source" on page 1120
- "Failback Job Wizard for VMware: Destination" on page 1122
- "Failback Job Wizard for VMware: Networks" on page 1125
- "Failback Job Wizard for VMware: Re-IP" on page 1128
- "Failback Job Wizard for VMware: Options" on page 1131

# Starting Replica Failback for VMware

Before starting replica failback for VMware vSphere, make sure that you have switched the replicas to the “Failover” state with a VMware Failover Job. Refer to “Failover to Replica for VMware” on page 1107 for details.

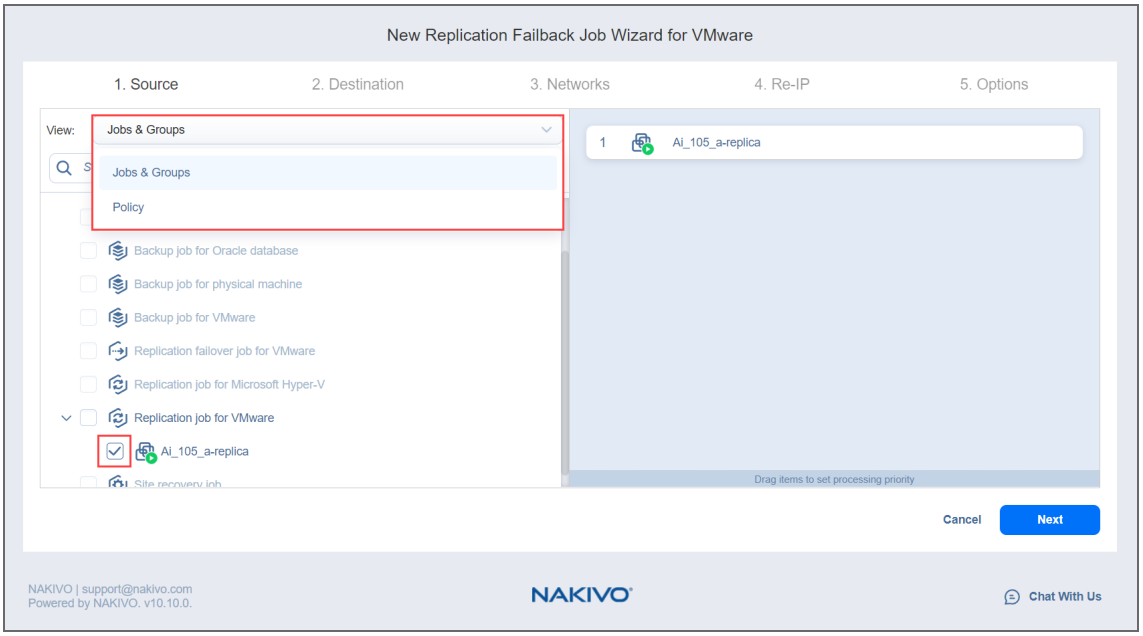
To start the failback procedure, click **Recover** in the **Data Protection** menu and then click **Replication failback for VMware** in the **VMware Full Recovery** section. Alternatively, the failback can be performed by using by using the search function.



# Failback Job Wizard for VMware: Source

On the **Source** page of the wizard, do the following:

1. In the left pane of the page, choose either of the following inventory views:
   - **Jobs & Groups**: When selected, the inventory tree displays groups, jobs, and backups of the appropriate hypervisor. You can select the required replica from the list of replication jobs.
   - **Policy**: When selected, job policies can be used. Refer to “Managing Job Policies” on page 382 for details. If the items were selected in alternate views, a dialog box opens warning you that switching to the **Policy** view will reset your current selection. Click **Switch View** to confirm switching to the **Policy** view. Make sure that at least one item matches the available set of policy rules. Refer to “Managing Policy Rules” on page 385 for details.
2. Select one or more replica VMware VMs in the left pane of the page.




Click **Next** to confirm adding selected replicas to the failback job. The wizard will display the next page.

**Notes**

- If you cannot find a replica or a container:
  - Make sure the corresponding vCenter or ESX(i) host has been added to the inventory.
  - Refresh inventory.
- Adding a VMware container to the job ensures that important replicas are always protected. If you add a VMware container to the job:
  - All replicas in the "Failover" state that are available in the selected container will be processed.
  - All new replicas that will be created in (or moved to) the container in the future will be automatically added to the job and processed.
- The order in which replicas are processed is important if the Transporter performing failback cannot process all replicas of the job simultaneously — either because the Transporter is processing other jobs at the same time or because the job contains more VM disks than the Transporter's maximum load specified during the Transporter creation.

# Failback Job Wizard for VMware: Destination

On this page of the wizard, you can specify a failback destination for your VMware replicas. Failback to the original location will update/replace the existing source VM in the original location. Failback to the new location will create a new VM.

> **Note**
> When recovering a VMware VM from an ESXi 8.0 host to an ESXi 7.0 host, the **VirtualMachineConfigSpec** > **VirtualMachineSgxInfo** property is skipped.

Refer to the following subsections on providing a destination for your VMware vSphere Replica Failback job:

- “Using Original VMware Failback Location” below
- “Creating New VMware Failback Location” below
- “Creating a Different VMware Failback Location for All Replicas” on the next page

## Using Original VMware Failback Location

To use the original location for your VMware failback job, follow the steps below:

1. In the **Failback location** list, choose **Original Location**.
2. Click **Next** to go to the **Networks** page of the wizard.

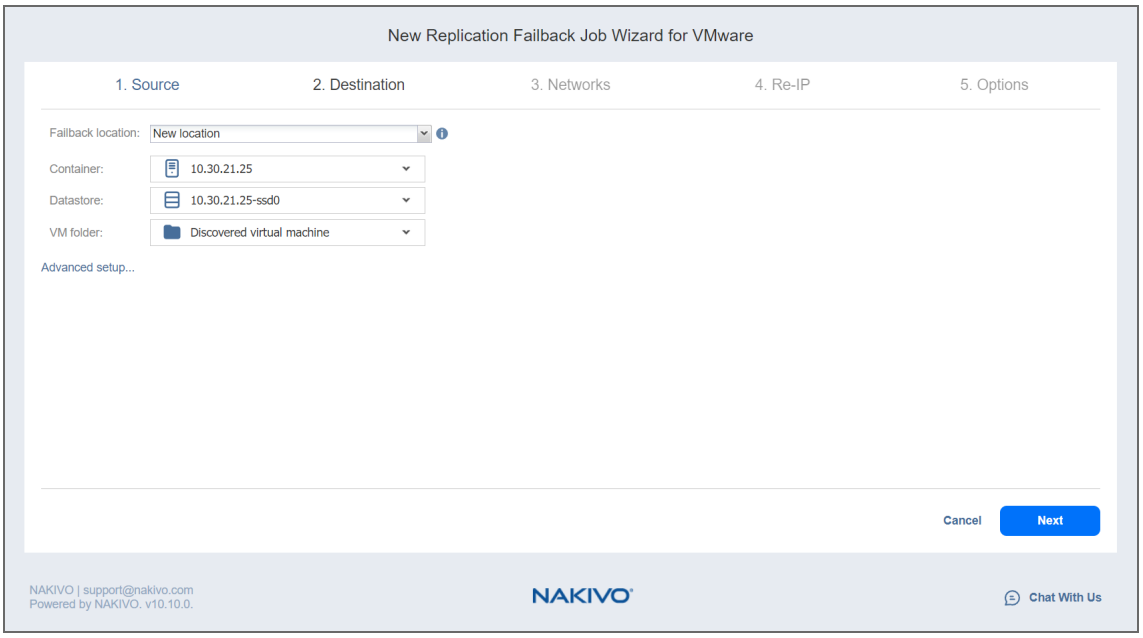



## Creating New VMware Failback Location

To create a new location for your VMware failback job, follow the steps below:

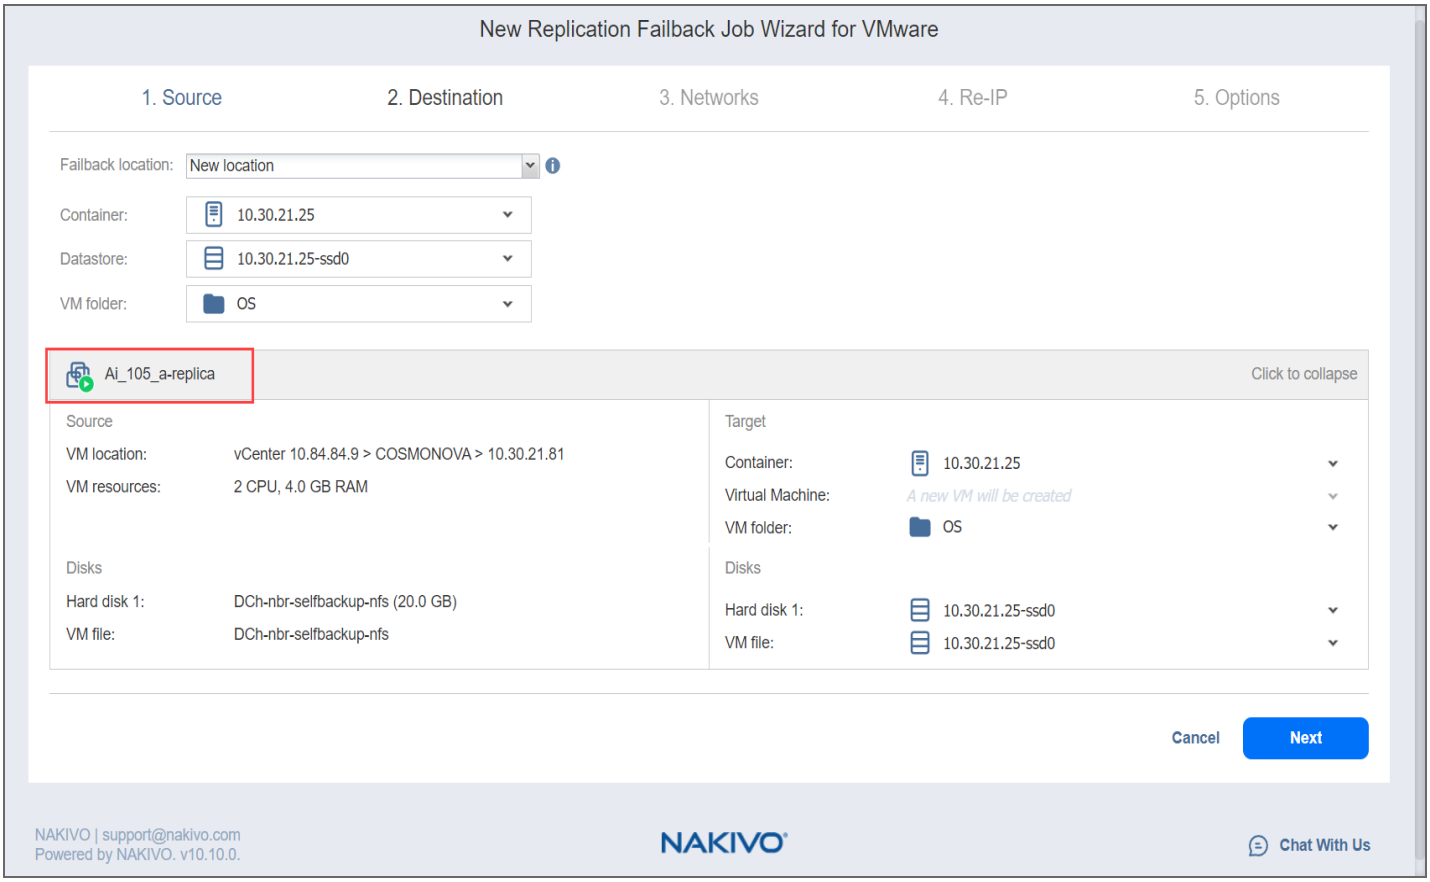
1. In the **Failback location** list, choose **New Location**.
2. In the **Container** list, choose a container for your failback location.

3. In the **Datastore** list, choose a datastore for your failback location.
4. Optionally, you can choose a target VM folder from the **VM folder** drop-down list.
5. Click **Next** to go to the **Networks** page of the wizard.




## Creating a Different VMware Failback Location for All Replicas

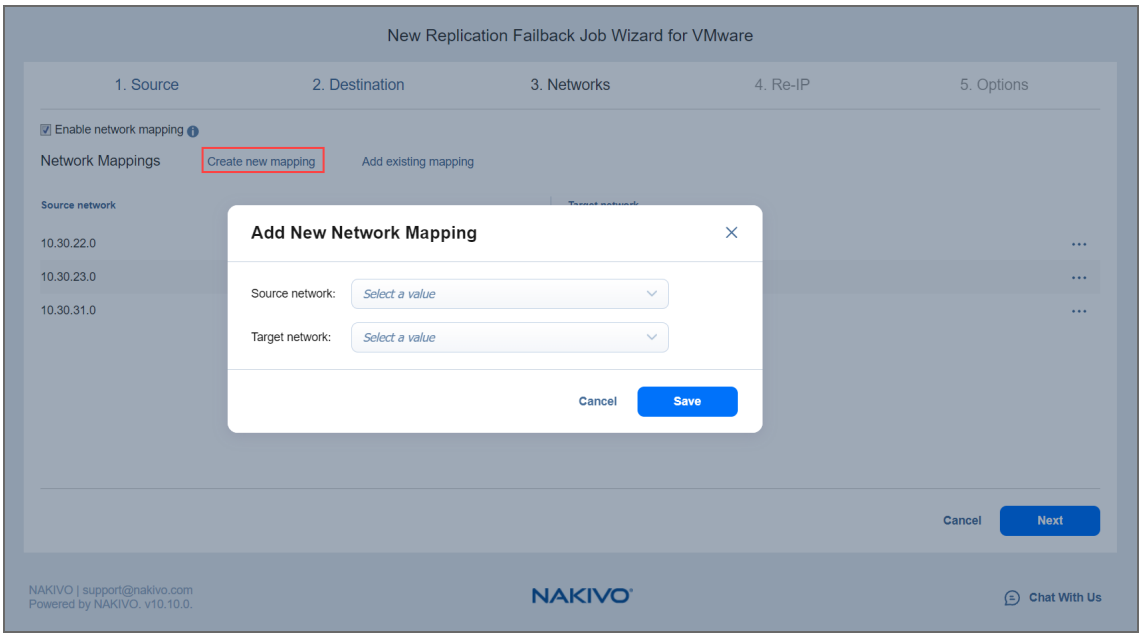
To create a new failover location for every replica you've added to the job, click **Advanced options** and set up the options as described for **New Location**.


New Replication Failback Job Wizard for VMware
1. Source
2. Destination
3. Networks
4. Re-IP
5. Options
Failback location:
New location
Container:
10.30.21.25
Datastore:
10.30.21.25-ssd0
VM folder:
OS
Ai_105_a-replica
Click to collapse
Source
VM location:
vCenter 10.84.84.9 > COSMONOVA > 10.30.21.81
VM resources:
2 CPU, 4.0 GB RAM
Disks
Hard disk 1:
DCh-nbr-selfbackup-nfs (20.0 GB)
VM file:
DCh-nbr-selfbackup-nfs
Target
Container:
10.30.21.25
Virtual Machine:
A new VM will be created
VM folder:
OS
Disks
Hard disk 1:
10.30.21.25-ssd0
VM file:
10.30.21.25-ssd0
Cancel
Next
NAKIVO | support@nakivo.com
Powered by NAKIVO. v10.10.0.
NAKIVO
Chat With Us

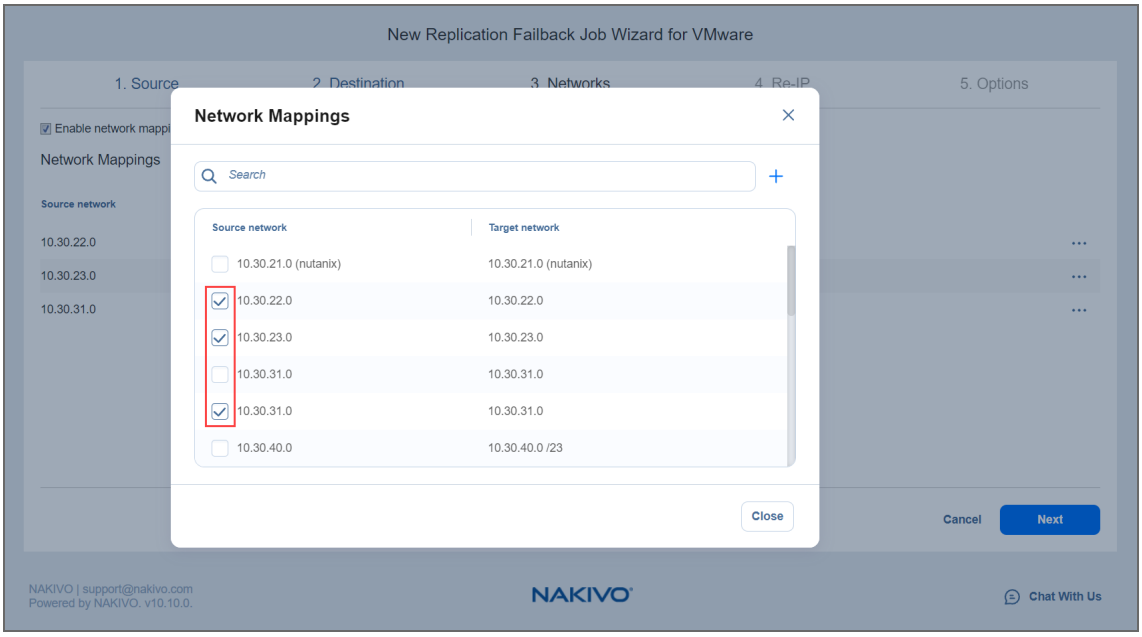
# Failback Job Wizard for VMware: Networks

To map source VM virtual networks to appropriate target virtual networks, take the following actions on the **Networks** page of the wizard:

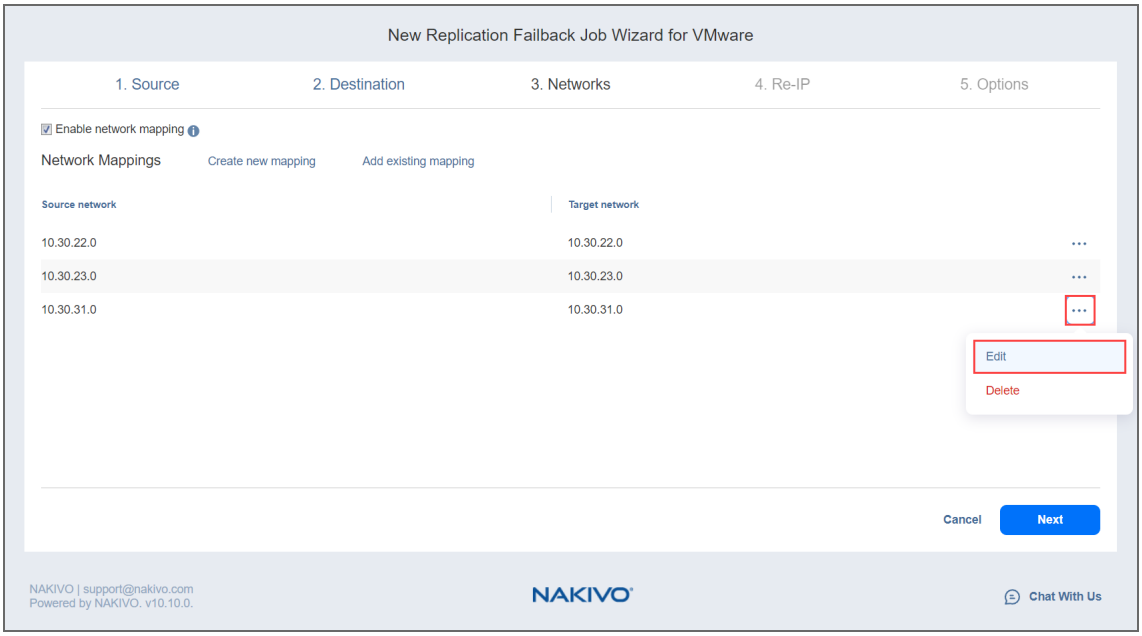
1. Select **Enable network mapping**.
2. The **Network Mappings** section opens. You have the following options:
    - Create a new mapping:
        a. Click **Create new mapping**.
        b. The **Add New Network Mapping** dialog opens. Choose a source network and a target network and click **Save**.



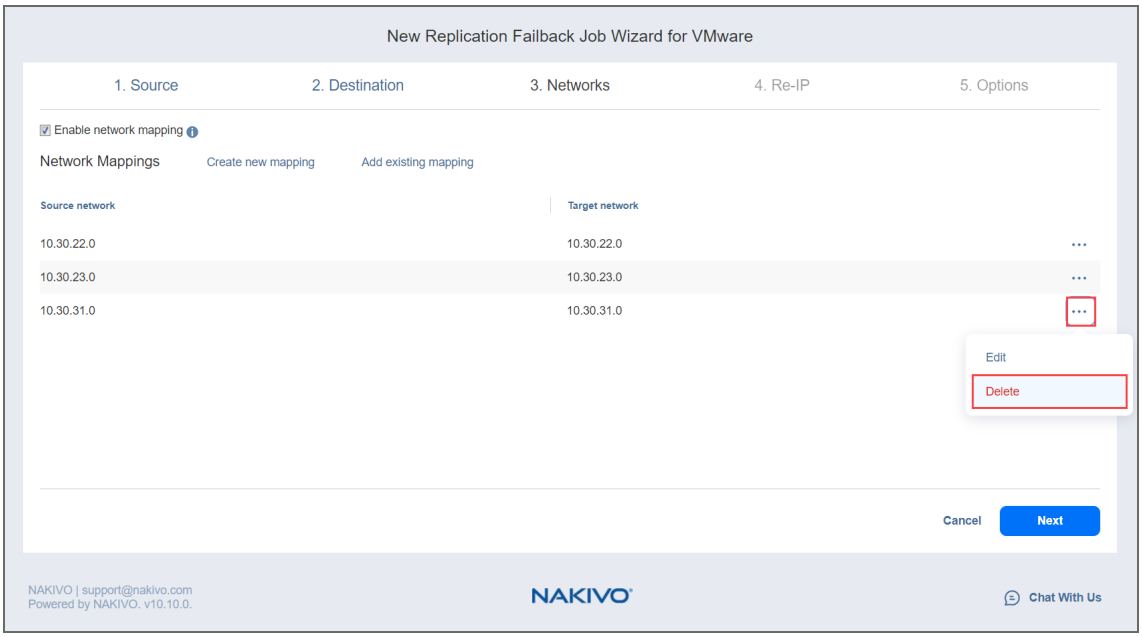

    - Add an existing mapping:
        a. Click **Add existing mapping**.
        b. The **Network Mappings** dialog opens. Choose one or more appropriate network mappings and close the dialog box.




- Edit an existing mapping:
  a. Hover the pointer over the necessary item in the **Network Mappings** list.
  b. A toolbar with icons opens to the right of the item. Click the **Edit** button.
  c. The **Edit Network Mapping** dialog opens. Choose an appropriate item from the **Target network** drop-down list and then click **Save**.




- Delete an existing mapping:
  a. Hover the pointer over the necessary item in the **Network Mappings** list.
  b. A toolbar with icons opens to the right of the item. Click the **Delete** icon.




- To leave the list of existing network mappings intact, go to the next page of the wizard.

3. Click **Next** to go to the next page of the wizard.



# Failback Job Wizard for VMware: Re-IP

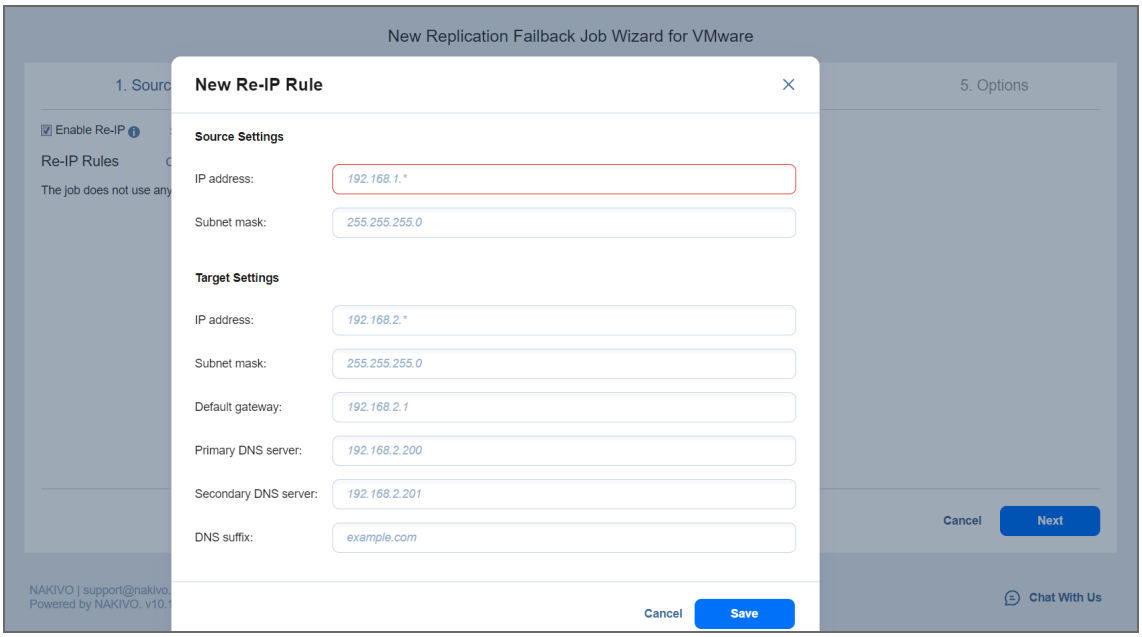
On the **Re-IP** page of the wizard, you can map a source VM IP address to a specific target IP address.

Please proceed as follows:

1. Select **Enable Re-IP**.
2. The **Re-IP Rules** section opens. Click the **Select VMs** link**.**
3. The **Re-IP** dialog box opens. In the list of your source VMs, select at least one VM. Select the credentials to be used for each VM and close the dialog box.

> **Note**
> Re-IP rules will be applied only to VMs that have a static IPv4 address configured.

4. You have the following options:
   - Create a new rule:
     a. Click **Create new rule**.
     b. The **New Re-IP Rule** dialog opens. Enter source and target settings for the Re-IP rule and click **Save**.

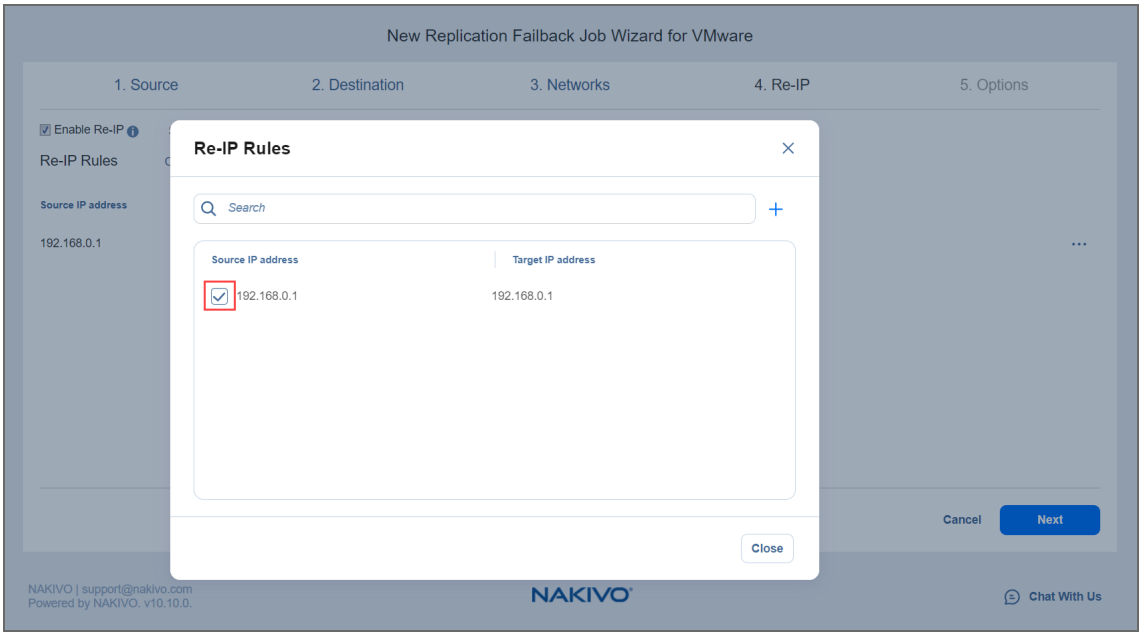



> **Note**
> You can use wildcards for IP addresses.

**Example**

When you enter 192.168.1.* -> 10.30.30.* Re-IP rule, the source VM IP address such as192.168.1.50 will be changed to the 10.30.30.50 IP address for your replica failback job. When there are several Re-IP rules applicable to your source VM, the application will define the most suitable one and apply it to the source VM IP address.

- Add an existing rule:
  a. Click **Add existing rule**.
  b. The **Re-IP Rules** dialog opens.
  c. Choose an appropriate Re-IP rule and close the dialog.




- Edit an existing Re-IP rule:
  a. Hover the pointer over the required item in the **Re-IP Rules** list.
  b. A toolbar with icons opens to the right of the item. Click the **Edit** button.
  c. The **Edit Re-IP Rule** dialog opens. Edit the required properties of the Re-IP rule and click **Save**.
- Delete an existing mapping:
  a. Hover the pointer over the required item in the **Re-IP Rules** list.
  b. A toolbar with icons opens to the right of the item. Click the **Delete** icon.
- To leave the list of existing Re-IP rules intact, go to the next page of the wizard.

5. Click **Next** to go to the next page of the wizard.


New Replication Failback Job Wizard for VMware
1. Source
2. Destination
3. Networks
4. Re-IP
5. Options
Enable Re-IP
Select VMs
Re-IP Rules
Create new rule
Add existing rule
Source IP address
Target IP address
192.168.0.1
192.168.0.1
Cancel
Next
NAKIVO | support@nakivo.com
Powered by NAKIVO. v10.10.0.
NAKIVO
Chat With Us

# Failback Job Wizard for VMware: Options

On the **Options** page of the wizard, specify options for the VMware Failback job as described in the following sections:

- "Job Options" below
- "Pre and Post Actions" below
- "Data Transfer" on the next page

## Job Options

1. In the **Job name** box, enter a string of 50 or fewer characters specifying the name of your failback job.
2. Select a **Job Priority** level between *1* and *5*, with *1* being the highest priority. Jobs with higher priority levels are prioritized by **Transporters** during job processing.

> **Note**
> This option is only available in the Enterprise, Enterprise Essentials, Enterprise Plus, MSP Enterprise, and MSP Enterprise Plus editions.

3. Select **Power off replica VMs** to power off the replica workloads during the failback.




## Pre and Post Actions

If needed, enable pre and post actions:

1. **Send job run reports to**: When selected, a job report is sent to specified recipients each time the failback job has completed. This overrides the default setting in the NAKIVO Backup & Replication "Settings" on page 405 dashboard.

2. **Run local pre job script:** When selected, a dialog box open in which you can specify the following options of the pre-job script to be run:
    - **Script path**: A local path to the script on the machine where the **Director** is installed. A script interpreter should be specified as well.
    - **Job behavior:** The following options are available:
        - **Wait for the script to finish**: This is the default option.
        - **Do not wait for the script to finish.**
    - **Error handling**: The following options are available:
        - **Continue the job on script failure**: This is the default option.
        - **Fail the job on script failure.**
3. **Run local post job script**: When selected, a dialog box opens in which you can specify options of the post-job script to be executed on the machine where the product is installed. Proceed with setting the options as described for **Run local pre job script**.




## Data Transfer

In the *Data Transfer* section of the **Options** page, you can select a transport mode and select a **Transporter** to be used for reading data. Proceed as described in the sections below.

### Transport Mode

Transport mode defines the method of VM data retrieval by a source **Transporter**. Select one of the following:

- **Automatic selection**: When this option is selected, NAKIVO Backup & Replication automatically selects the best transport mode available:

    - When the source **Transporter** is installed on a VM, NAKIVO Backup & Replication tries to use transport modes in the following order: Hot Add > SAN > LAN.
    - When the source **Transporter** is installed on a physical machine, NAKIVO Backup & Replication tries to use transport modes in the following order: SAN > LAN. Hot Add is not supported for physical machines.
- **SAN only**: When this option is selected, NAKIVO Backup & Replication only uses direct SAN access to retrieve VM data. If direct SAN access to VM data is not available, the job fails.
- **Hot Add only**: When this option is selected, NAKIVO Backup & Replication only uses Hot Add to retrieve VM data. If direct Hot Add is not available, the job fails.
- **LAN only**: When this option is selected, NAKIVO Backup & Replication only retrieves VM data over LAN.

Transporter Pool

If this option is enabled, only the transporters that belong to the selected transporter pool shall be used during the job run.

Transporters

By default, the product automatically determines which **Transporter** should be used to read data from the source VM. However, you can manually specify which **Transporters** should be used for the job. Here's how:
In the *Data Transfer* section, choose one of the following options:

- **Automatic selection**: The product will automatically determine which **Transporters** are the closest to source hosts (the hosts that run selected VMs) and will use those **Transporters** to retrieve data from source VMs.
- **Manual - configured for all VMs**: Select this option to manually specify a single **Transporter** that will be used to retrieve data from source VMs.
- **Manual - configured per host**: Select this option to manually specify which **Transporter** should be used to retrieve data from each source host. When selected, the **Replacement Transporter** drop-down list becomes available. In case a primary Transporter is unavailable, a replacement **Transporter** will be used. Note that the product selects what **Transporter** to use (primary or replacement) at the beginning of the job run, not while the job is running.

Transporter Load

You can limit the maximum number of **Transporter** tasks used by the job. By default, this number is set to 3 concurrent tasks.
To change the default number of tasks, do the following:

1. In the *Data Transfer* section, select the **Limit transporter load to** checkbox.
2. Specify the number of concurrent tasks in the corresponding box.

Bandwidth Throttling

Follow the steps below to regulate the speed of data transfer over the network for your backup job:

1. For the **Bandwidth throttling** option, choose **Enabled**.

> **Note**
> If bandwidth throttling is disabled for the current job, global bandwidth rules may still apply to your job. Refer to "Bandwidth Throttling" on page 407 for details.

2. Click the **settings** link that becomes available.
3. The **Job Bandwidth Rules** dialog box opens, displaying the list of available rules. You have the following options:
   - Create a new bandwidth rule for your backup job:
     a. Click the **Create New Rule** button.
     b. The **New Bandwidth Rule** dialog box opens. Refer to the "Bandwidth Throttling" on page 407 topic for details on creating a bandwidth rule.
     c. Click **Save**.
   - Activate an existing bandwidth rule for your job. Select the checkbox to the left of the required bandwidth rule. To deactivate a bandwidth rule for your job, deselect the corresponding checkbox.
   - Edit a bandwidth rule. Click the **Edit** link for a bandwidth rule and modify it in the **Edit Bandwidth Rule** dialog box that opens.
   - Disable a bandwidth rule. Click the **Disable** link. The bandwidth rule is disabled for all jobs.
   - Remove a bandwidth rule. Click the **Remove** link and then click **Delete** to confirm your operation.

Bottleneck detection

When the **Bottleneck detection** option is enabled, additional information is collected and recorded in NAKIVO Backup & Replication logs in the course of data transfer for the purpose of bottleneck detection. Check this option to enable the **Bottleneck detection** capability of the Transporters engaged in the job.

Click **Finish** or **Finish & Run** to complete the job creation.




**Note**

If you click **Finish & Run**, you will have to define the scope of your job. Refer to "Running Jobs on Demand" on page 364 for details.

# Real-Time Replication Failover for VMware

With NAKIVO Backup & Replication, you can switch your VMware VM to its latest real-time replica in case of failure.

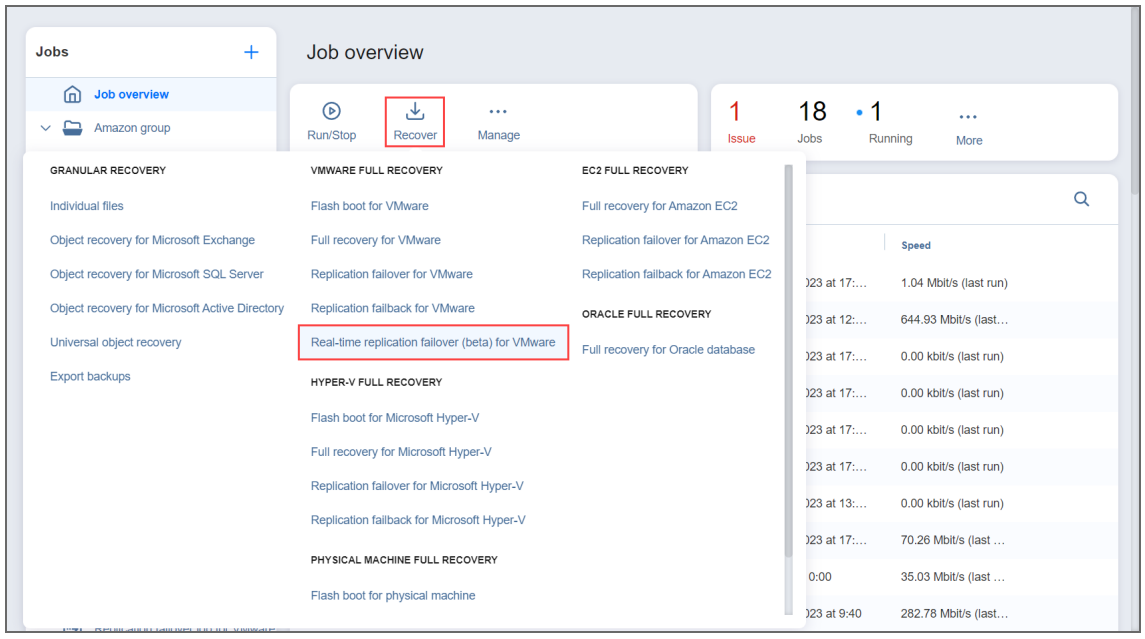
Refer to the following topics for more information:

- "Starting Real-Time Replication Failover for VMware" below
- "Real-Time Replication Failover Job Wizard for VMware: Source" on the next page
- "Real-Time Replication Failover Job Wizard for VMware: Networks" on page 1138
- "Real-Time Replication Failover Job Wizard for VMware: Re-IP" on page 1141
- "Real-Time Replication Failover Job Wizard for VMware: Options" on page 1143

## Starting Real-Time Replication Failover for VMware

Before starting a real-time replication failover job for VMware, make sure that you have an active real-time replication job for VMware available. Refer to "Creating VMware Real-Time Replication Jobs" on page 968 for details.

To start the real-time replication failover process, select the relevant real-time replication job for VMware and click **Recover** > **Real-time replication failover for VMware**. Note that the real-time replication job for VMware must be running successfully (or was running successfully before and now is stopped) for this option to be available.

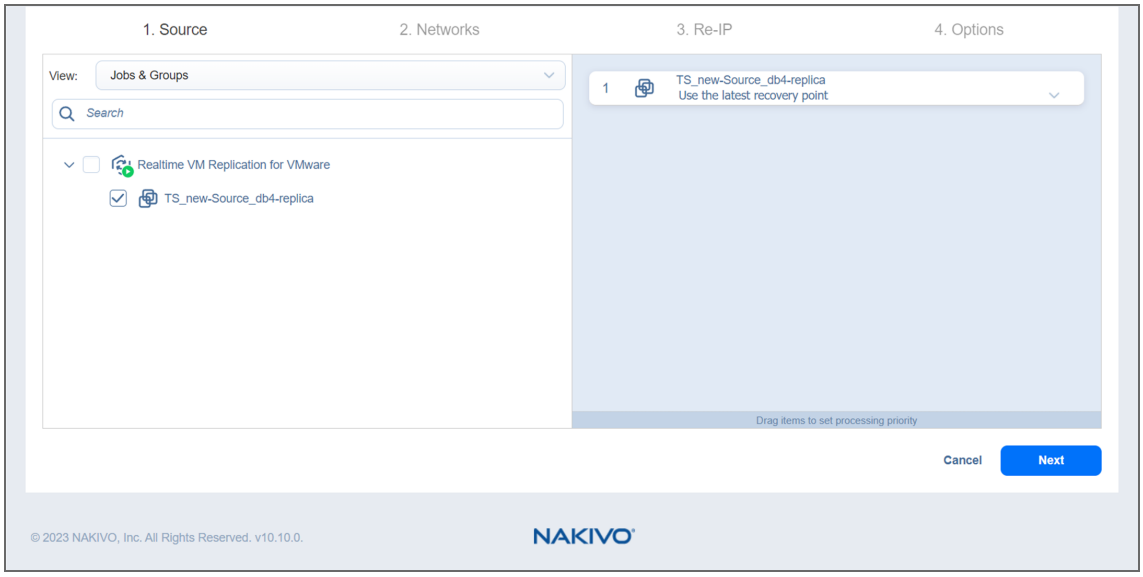



The **Real-Time Replication Failover Job Wizard for VMware** opens.

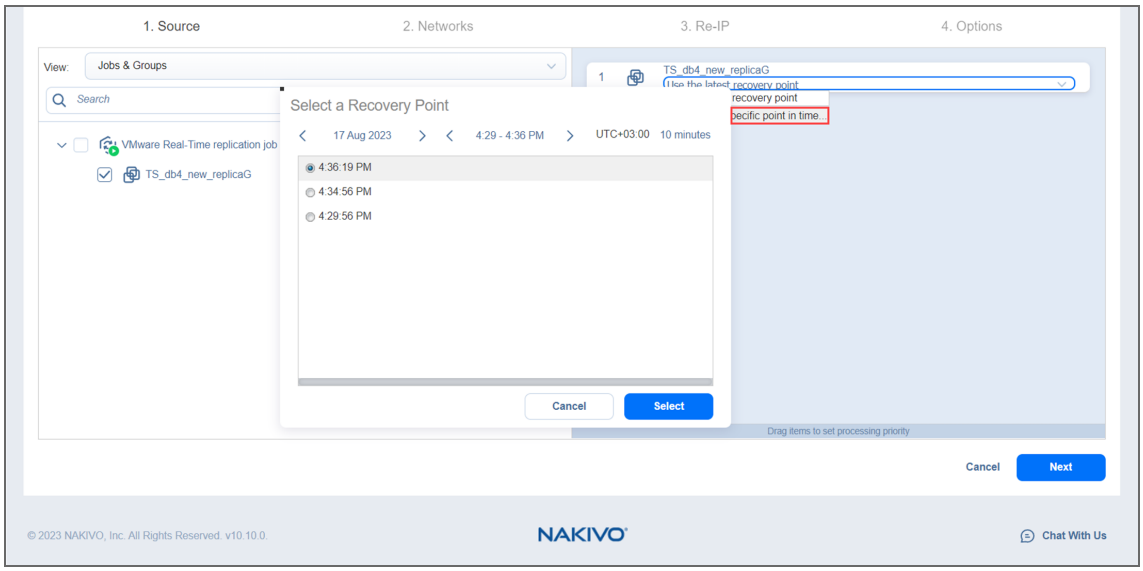
# Real-Time Replication Failover Job Wizard for VMware: Source

The **Real-Time Replication Failover Job Wizard for VMware** allows you to restore a real-time replica to a specific point in time. On the **Source** page of the wizard, do the following:

1. In the left pane of the page, choose either of the following inventory views:
   - **Jobs & Groups**: When selected, the inventory tree displays groups, jobs, and backups of the appropriate hypervisor. You can select the required replica from the list of replication jobs.
   - **Policy**: When selected, job policies can be used. Refer to “Managing Job Policies” on page 382 for details. If the items were selected in alternate views, a dialog box opens warning you that switching to the Policy view will reset your current selection. Click **Switch View** to confirm switching to the **Policy** view. Make sure that at least one item matches the available set of policy rules. Refer to “Managing Policy Rules” on page 385 for details.
2. Select one or more source VMware VMs in the left pane and then select a recovery point for each source VM in the right pane. By default, the option **Use the latest recovery point** is selected.




3. To restore to a specific recovery point, click the drop-down below a chosen object and select **Restore to a specific point in time**. In the **Select a Recovery Point** dialog box, choose a recovery point from the list and click **Select**.




**Notes**

- If the selected date has no recovery points, the text "No items available." will be displayed in the **Select a Recovery Point** dialog box.
- If the replica VM is powered on, the error message will be displayed with the text "Cannot open the journal. The replica is running."
- For one real-time replication failover job for VMware, the Transporter's maximum capacity is limited to 11 source VMs in case there is one disk per VM or up to 50 disks per source VM. For more information, see this article.

4. Click **Next** to go to the next page of the wizard.

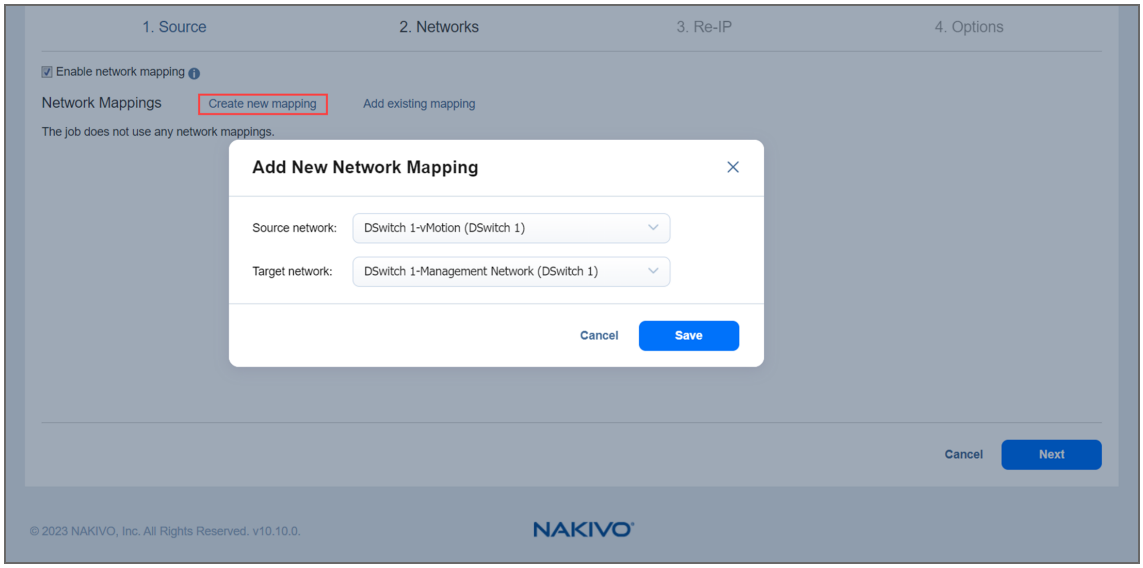
## Real-Time Replication Failover Job Wizard for VMware: Networks

When the real-time replica network – or target VMware virtual network – differs from the network address where the source VMs are deployed – or source VMware virtual network – your failover job needs a relation between these networks to be performed successfully. Such relation is called a network mapping.

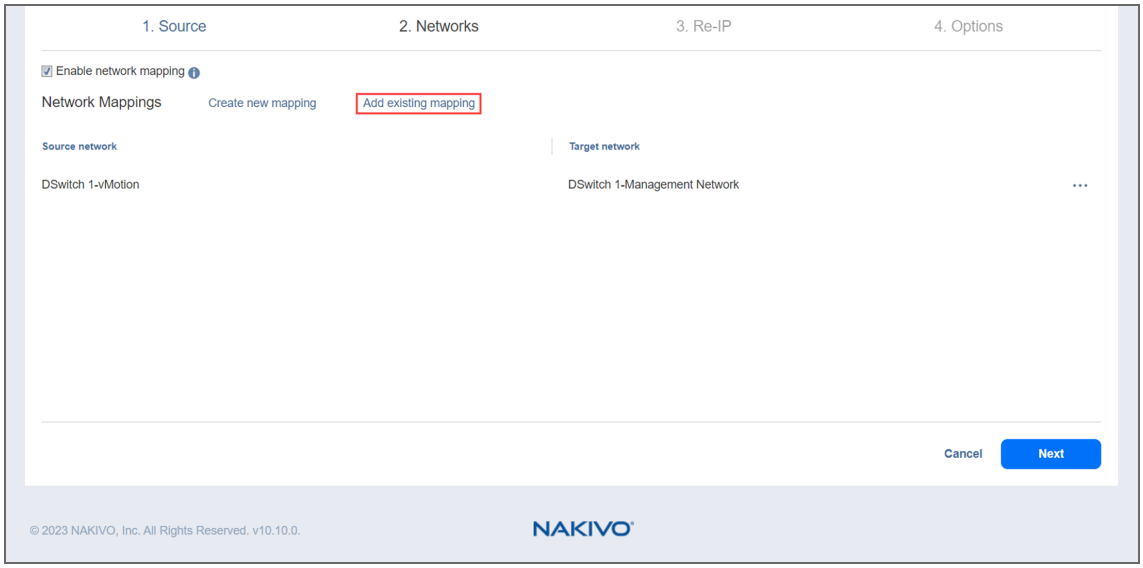
To map source VMware virtual networks to appropriate target virtual networks, please do the following on the **Networks** page of the wizard:

1. Select **Enable network mapping**.
2. The **Network Mappings** dialog box opens. You have the following options:

- Create a new network mapping:
    a. Click **Create new mapping**.
    b. The **Add New Network Mapping** dialog opens. Choose a source network and a target network and click **Save**.



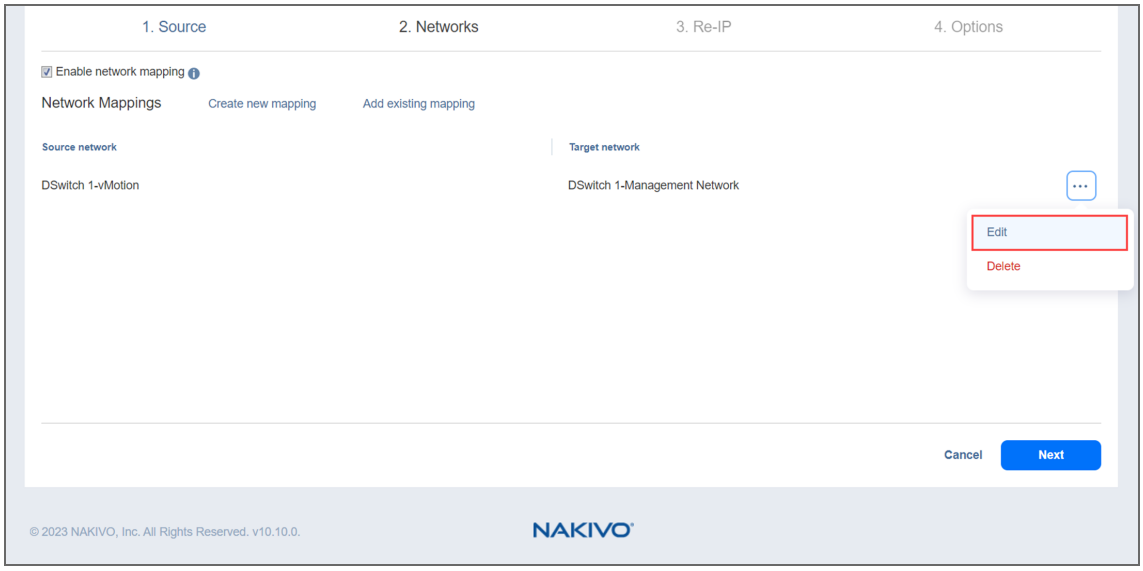

- Add an existing network mapping:
    a. Click **Add existing mapping**.
    b. The **Network Mappings** dialog opens. Choose one or more appropriate network mappings and close the dialog box.



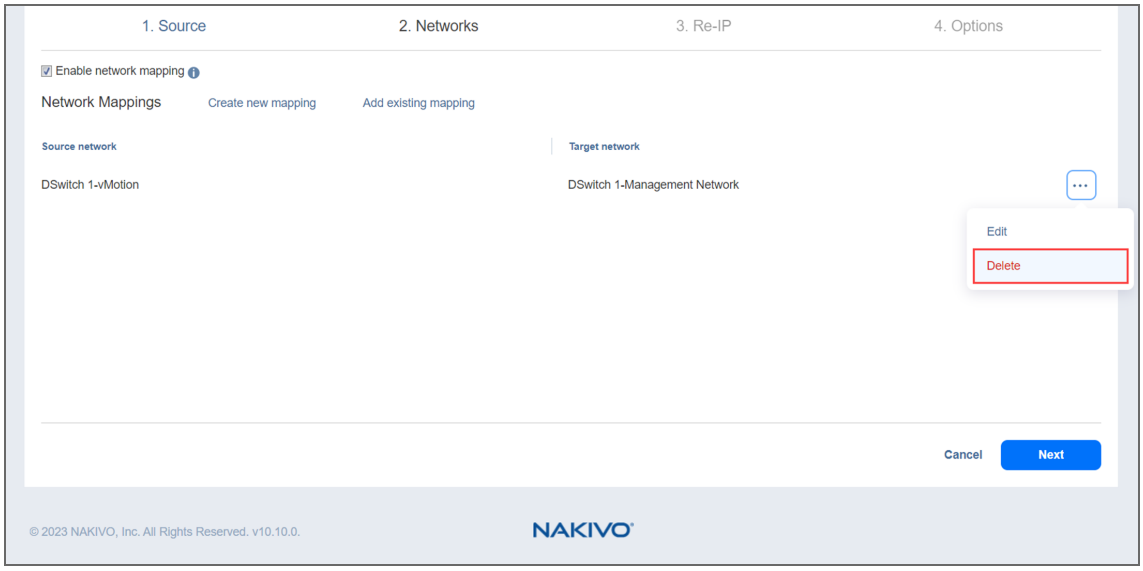

- Edit an existing network mapping:

a. Hover over the necessary item in the **Network Mappings** list and then click the ellipsis **Manage** button to the right of the item. Then, click **Edit**.

b. The **Edit Network Mapping** dialog box opens. Choose an appropriate item from the **Target network** drop-down list and then click **Save**.




- Delete an existing mapping: hover over the necessary item in the **Network Mappings** list and then click the ellipsis **Manage** button to the right of the item. Then, click **Delete**.




- To leave the list of existing network mappings intact, go to the next page of the wizard.

3. Click **Next** to go to the next page of the wizard.

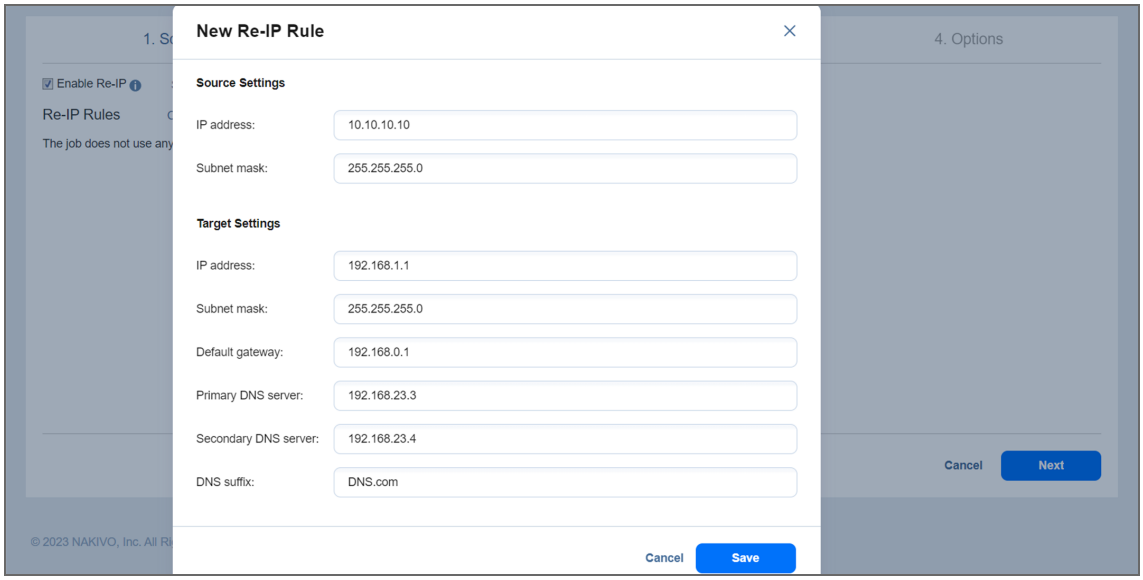
# Real-Time Replication Failover Job Wizard for VMware: Re-IP

When the IP addressing scheme for the target real-time replica differs from the IP addressing scheme for the VMware source VMs, your failover job needs relations between source VMs addresses and the target replica VMs addresses to be performed successfully. These relations are called re-IP rules.

> **Note**
> VMware Tools must be running on source VMs to successfully enable re-IP for your failover job.

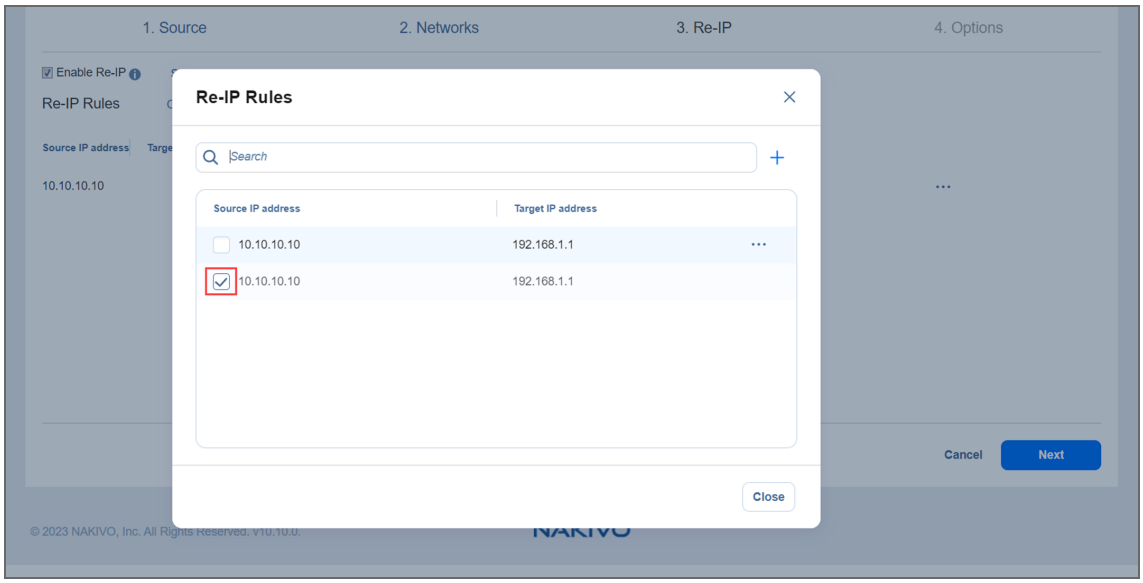
To map a source VM IP address to a specific target IP address, do the following on the Re-IP page of the wizard:

1. Select **Enable Re-IP**.
2. The **Re-IP Rules** section opens. Click the **Select VMs** link.
3. The **Re-IP dialog** box opens. In the list of your source VMs, select at least one VM and close the dialog box.
4. You have the following options:
    - Create a new rule:
        a. Click **Create new rule**.
        b. The **New Re-IP Rule** dialog opens. Enter source and target settings for the Re-IP rule and click **Save**.



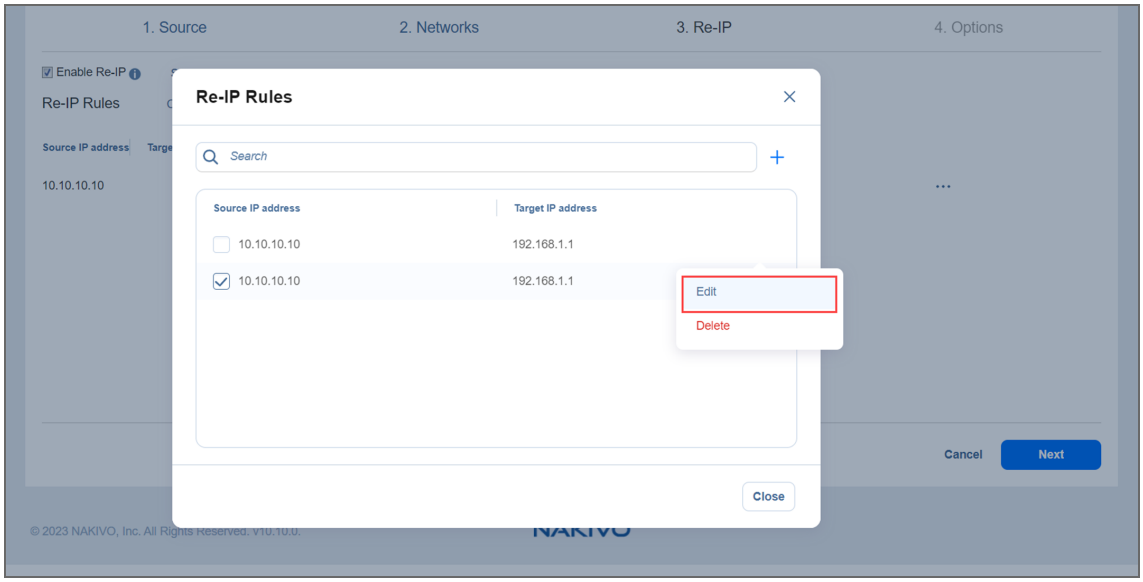
Note that you can use wildcards for IP addresses. For example, if you enter *192.168.1.** -> *10.30.30.** for the re-IP rule, a source VM IP address such as *192.168.1.50* will be changed to the *10.30.30.50* IP address for your failover job.

- Add an existing rule:
    a. Click **Add existing rule**.
    b. The **Re-IP Rules** dialog opens. Select one or more appropriate Re-IP rules and close the dialog.

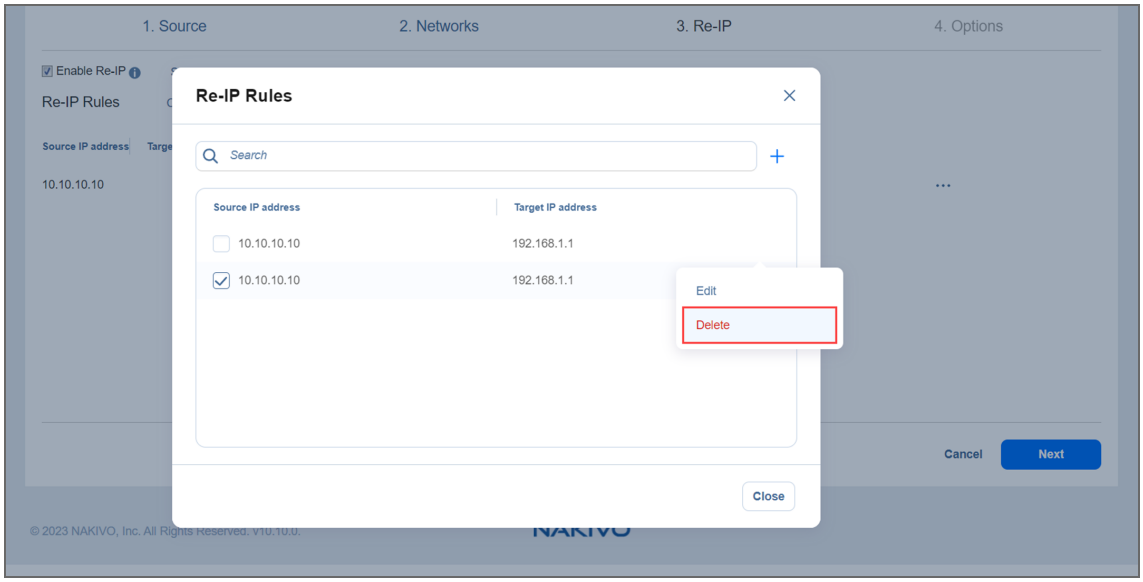



- Edit an existing Re-IP rule:
    a. Hover over the required item in the Re-IP Rules list and then click the ellipsis **Manage** button to the right of the item. Then, click **Edit**.
    b. The **Edit Re-IP Rule** dialog box opens. Edit the required properties of the Re-IP rule and

click **Save**.




- Delete an existing mapping: Hover over the required item in the Re-IP Rules list and then click the ellipsis **Manage** button to the right of the item. Then, click **Delete**.




- To leave the list of existing Re-IP rules intact, go to the next page of the wizard.

5. Click **Next** to go to the next page of the wizard.

# Real-Time Replication Failover Job Wizard for VMware: Options

On this page of the wizard, set the options for the VMware real-time replication failover job for VMware:

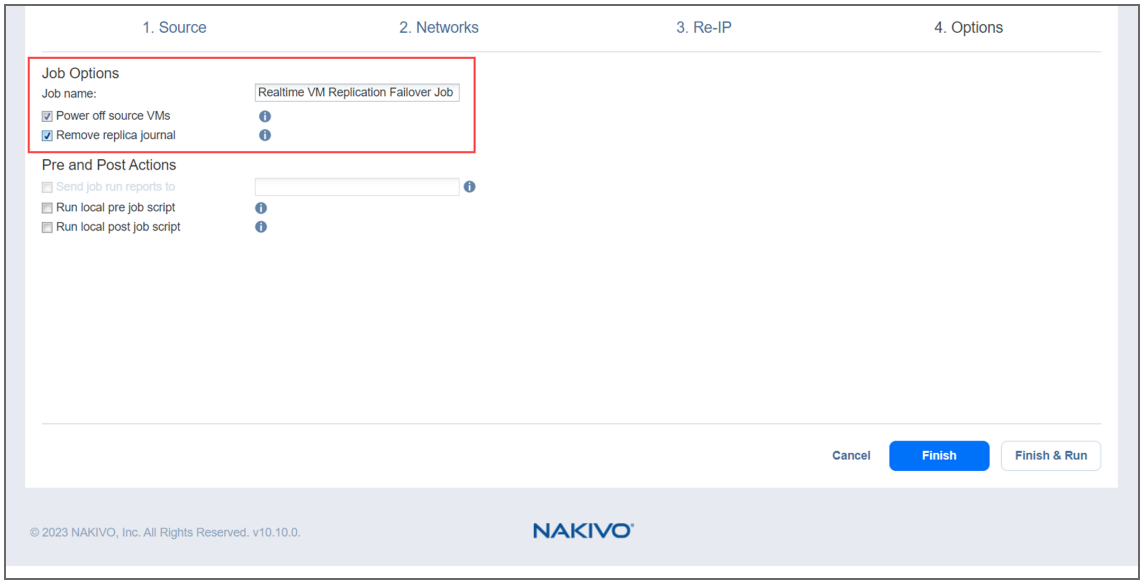
- “Job Options” on the next page
- “Pre and Post Actions” on the next page

- “Completing the New Real-Time Replication Failover Job Wizard for VMware ” on the next page

## Job Options

The following failover job options are available for editing:

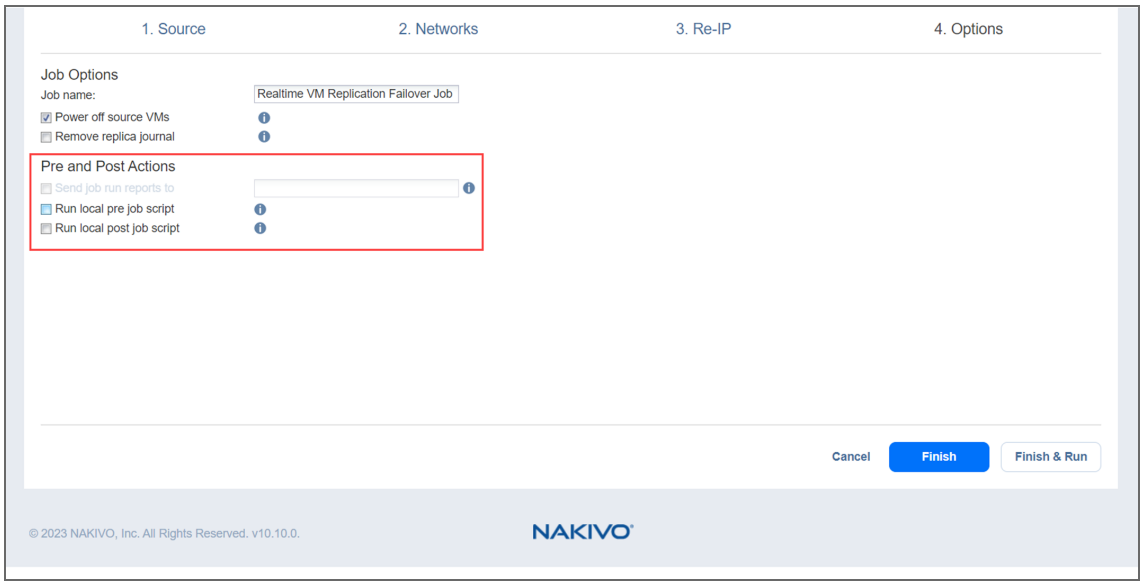
1. **Job name**: A string of 50 or fewer characters specifying the name of the failover job.
2. **Power off source VMs**: Select this checkbox to power off source VMs when the failover job has completed.
3. **Remove replica journal**: Select this checkbox to remove the replica journal during failover.




## Pre and Post Actions

The following pre- and post-job actions can be set up:

1. **Send job run reports to**: When selected, a job report is sent to specified recipients each time the failover job has completed. To enable this option, configure your Email settings. This overrides the default setting in the NAKIVO Backup & Replication “Settings” on page 405.




2. **Run local pre job script**: When selected, a dialog box opens in which you can specify the following options of the pre-job script to be executed:
    - **Script path**: A local path to the script on the machine where the **Director** is installed. A script interpreter should be specified as well.
    - **Job behavior**: The following options are available:
        - **Wait for the script to finish**: This is the default option.
        - **Do not wait for the script to finish**.
    - **Error handling**: The following options are available:
        - **Continue the job on script failure**: This is the default option.
        - **Fail the job on script failure**.
3. **Run local post job script**: When selected, a dialog box opens in which you can specify options of the post-job script to be executed on the machine where the product is installed. The options are similar to the ones of the pre-job script.

## Completing the New Real-Time Replication Failover Job Wizard for VMware

Click **Finish** or **Finish & Run** to complete the job creation.

> **Note**
> If you click **Finish & Run**, you will have to define the scope of your job. Refer to "Running Jobs on Demand" on page 364 for details.

# Site Recovery Job

Site recovery job is a special job that automates execution of one or multiple site recovery actions. You can execute your site recovery job on demand or on schedule. The list of available actions can be found under Site Recovery Job Wizard: Actions.

Creating a site recovery job is done with a wizard and includes the following steps:

- “Creating Site Recovery Job” on page 1147
- “Running Site Recovery Job” on page 1182

# Creating Site Recovery Job

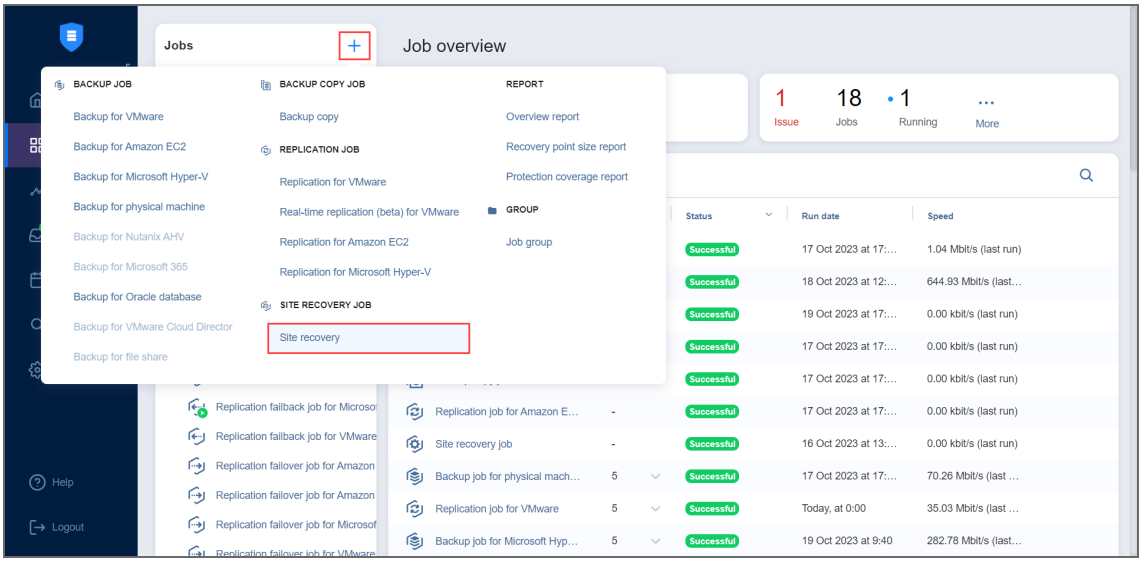
Creating a site recovery job is done with a wizard and includes the following steps:


- "Starting Site Recovery Job Wizard" on page 1148
- "Site Recovery Job Wizard: Actions" on page 1149
- "Site Recovery Job Wizard: Networks" on page 1170
- "Site Recovery Job Wizard: Re-IP" on page 1173
- "Site Recovery Job Wizard: Test Schedule" on page 1176
- "Site Recovery Job Wizard: Options" on page 1181

# Starting Site Recovery Job Wizard

Follow the steps below to start creating a site recovery job:

1. Go to the **Data Protection** menu and click the plus **Create** button.
2. In the menu that opens, click **Site recovery**.




The **New Site Recovery Job Wizard** opens.

## Site Recovery Job Wizard: Actions

On the **Actions** page of the **Site Recovery Job wizard**, you can add one or more actions to a Site Recovery job. Refer to the subsections below for details:

- Actions Available for Site Recovery Job
- Managing Actions of Site Recovery Job
- Options Common to Most Actions

When finished with adding actions to a Site Recovery job, click **Next** to go to the **Networks** page of the **Site Recovery Job Wizard**.

### Actions Available for Site Recovery Job

The list of actions available for a Site Recovery job is available in the left pane. It includes the following:

- "Failover VMware VMs Action" on page 1151
- "Failover Real-Time Replication VMware VMs Action" on page 1152
- "Failback VMware VMs Action" on page 1155
- "Start VMs / Instances Action" on page 1159
- "Stop VMs / Instances Action" on page 1161
- "Run / Stop Jobs Action" on page 1163
- "Run Script Action" on page 1164
- "Attach / Detach Repository Action" on page 1165
- "Send Email Action" on page 1166
- "Wait Action" on page 1167
- "Check Condition Action" on page 1168

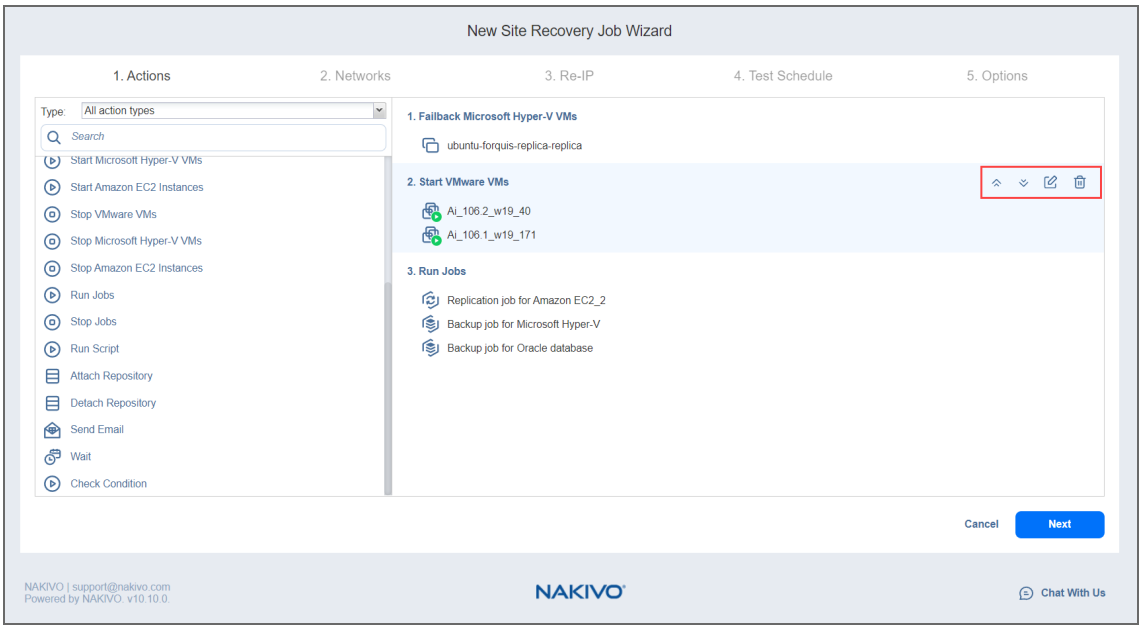
To add an action to a Site Recovery job, click the corresponding item in the actions list and follow the instructions of the wizard that opens.

### Managing Actions of Site Recovery Job

The actions list of a Site Recovery job is located in the right pane and contains the actions you add.
For a selected action, a toolbar with buttons is available allowing you to perform the following commands:

- **Move up / Move down**: Moves your action up/down in the list of Site Recovery job actions.
- **Edit**: A page opens allowing you to make changes to your Site Recovery job action. When finished with editing the Site Recovery job action, click **Save**.
- **Remove**: A dialog box opens asking you to confirm removing the Site Recovery job action. Click

the **Remove** button to confirm your operation.




Options Common to Most Actions

Most of the actions you add to your Site Recovery job will have the following common options:

- **Run this action in**: This option allows you to choose how to run the action:
  - **Run this action in both testing and production mode**.
  - **Run this action in production mode only**.
  - **Run this action in testing mode only**.
- **Waiting behavior**: This option allows you to choose one of the following:
  - **Wait for this action to complete**: The Site Recovery job will wait for the action to be completed before continuing to run the job.
  - **Start next action immediately**: The Site Recovery job will continue running while the action is in progress.
- **Error handling**: This option allows you to choose one of the following:
  - **Stop and fail the site recovery job if this action fails**: The Site Recovery job will stop and fail if the action fails.
  - **Continue the site recovery job if the action fails**: The Site Recovery job will continue if the action fails.

# Failover VMware VMs Action

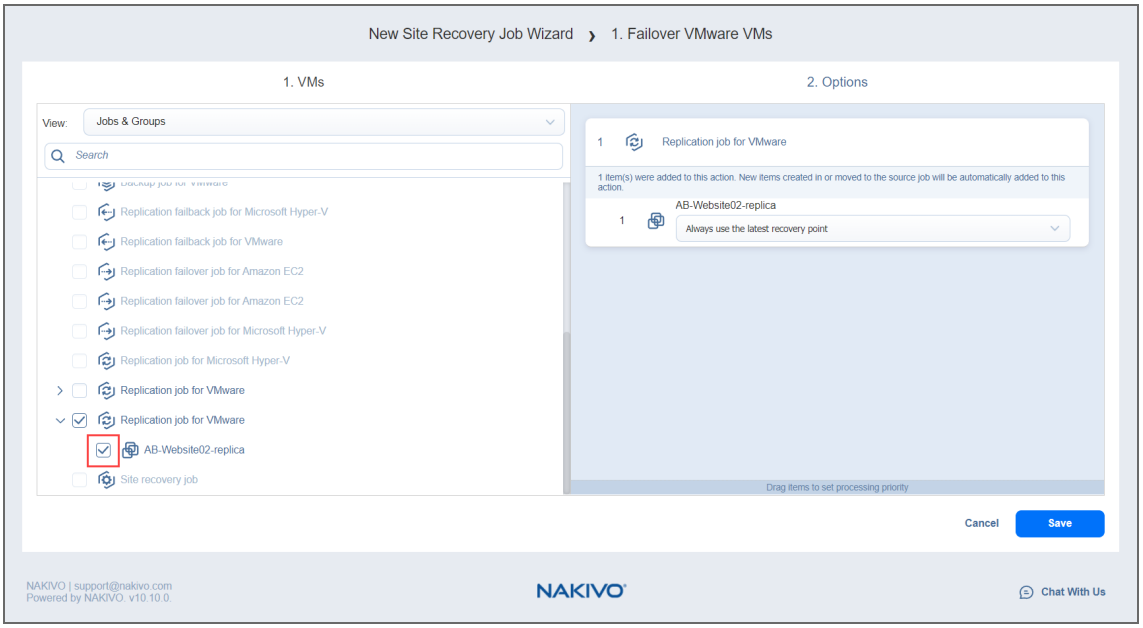
Please refer to the following sections on adding a **Failover VMware VMs** action to your Site Recovery job:

- Failover VMware VMs: VMs
- Failover VMware VMs: Options

Failover VMware VMs: VMs

On this step of the wizard, do the following:

1. In the left pane of the page, choose either of the following inventory views:
   - **Jobs & Groups**: When selected, the inventory tree displays groups, jobs, and backups of the appropriate hypervisor. You can select the required replica from the list of replication jobs.
   - **Policy**: When selected, job policies can be used. Refer to "Managing Job Policies" on page 382 for details. If the items were selected in alternate views, a dialog box opens warning you that switching to the **Policy** view will reset your current selection. Click **Switch View** to confirm switching to the **Policy** view. Make sure that at least one item matches the available set of policy rules. Refer to "Managing Policy Rules" on page 385 for details.
2. Select one or more source VMware VMs in the left pane and then select a recovery point for each source VM in the right pane.
3. Click **Next** to go to the **Options** page.

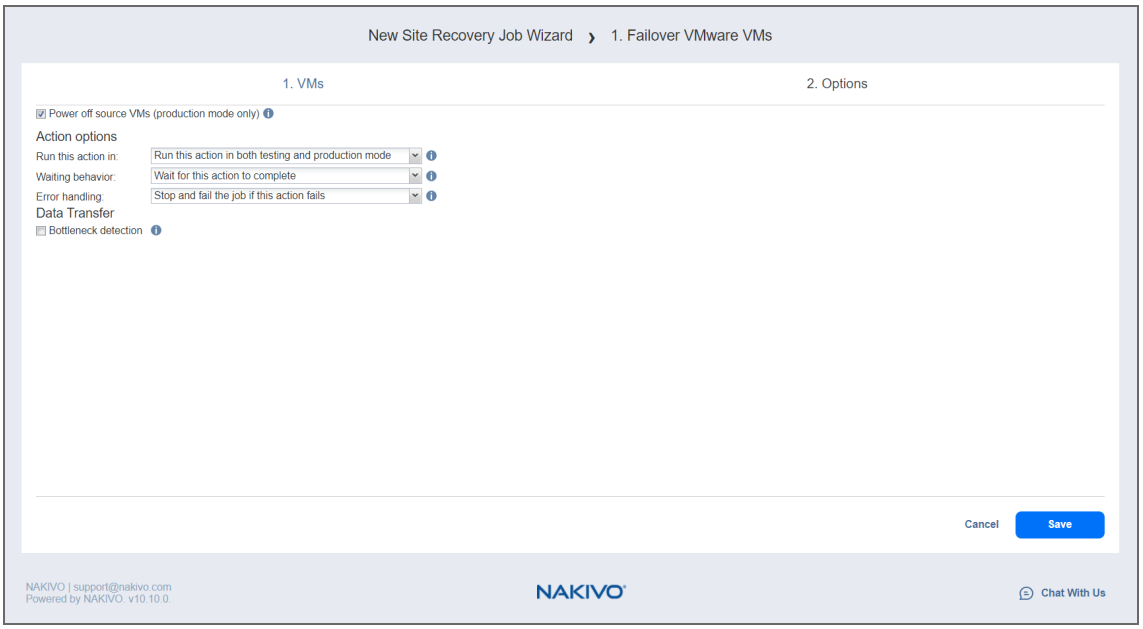



Failover VMware VMs: Options

On this page of the wizard, do the following:

1. Select **Power off source VMs** to power off the production workloads during the failover. Note that this will take place only in production mode; no workloads will be powered on in test mode.

2. In the *Action options* section, set the options for your action. Refer to "Site Recovery Job Wizard: Actions" on page 1149 for details.
3. Click **Save**.



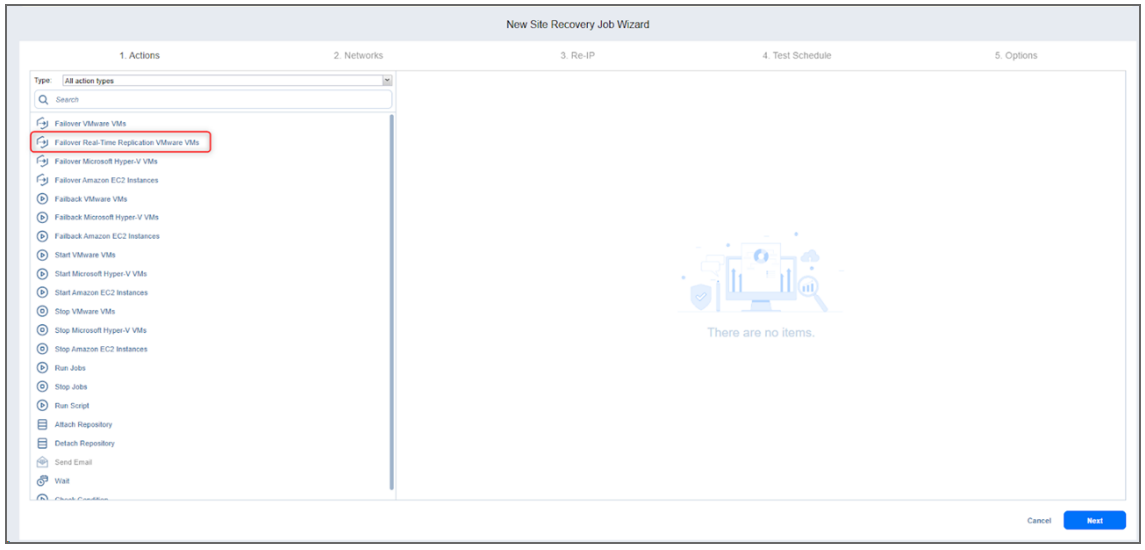

The **Options** page closes and your **Failover VMware VMs** action is added to the Site Recovery Job.

## Failover Real-Time Replication VMware VMs Action

Please refer to the following sections on adding a **Failover Real-Time Replication VMware VMs action** to your **Site Recovery job**:

"Failover Real-Time Replication VMware VMs: VMs" on the next page

"Failover Real-Time Replication VMware VMs: Options" on the next page



# Failover Real-Time Replication VMware VMs: VMs

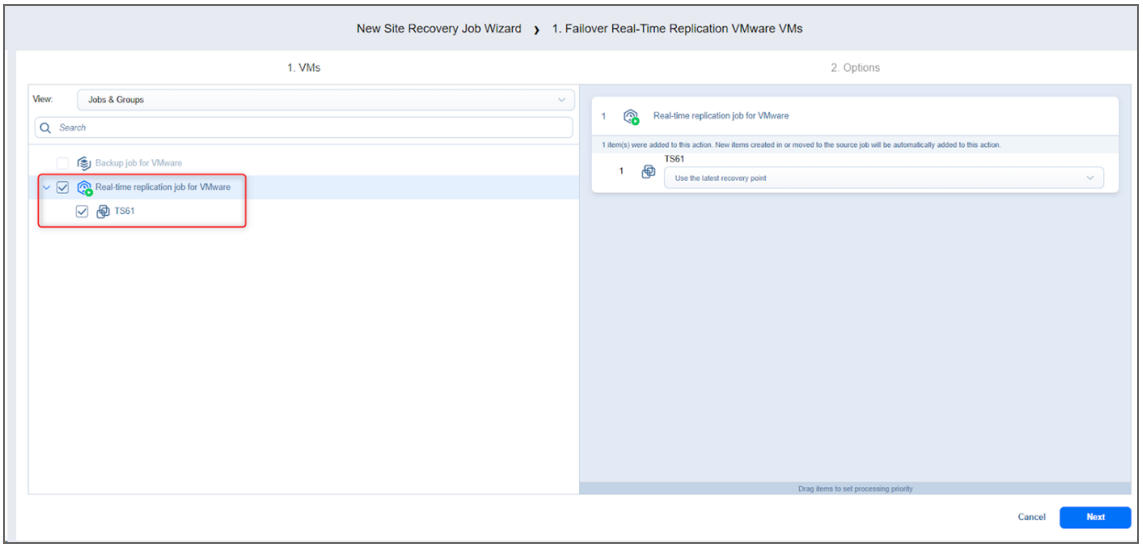
On this step of the wizard, do the following:

1. In the left pane of the page, choose either of the following inventory views:
   - **Jobs & Groups**: When selected, the inventory tree displays groups, jobs, and backups of the appropriate hypervisor. You can select the required replica from the list of real-time replication jobs.
   - **Policy**: When selected, job policies can be used. Refer to "Managing Job Policies" on page 382 for details. If the items were selected in alternate views, a dialog box opens warning you that switching to the **Policy** view will reset your current selection.

     Click **Switch View** to confirm switching to the **Policy** view. Make sure that at least one item matches the available set of policy rules. Refer to "Managing Policy Rules" on page 385 for details.
2. Select one or more source RTR Replicas or RTR jobs in the left pane.

   > **Note**
   > The latest recovery point is used by default.

3. Optionally select a recovery point for each source VM in the right pane, if you want to restore to a specific point.
4. Click **Next** to go to the **Options** page.




# Failover Real-Time Replication VMware VMs: Options

On this page of the wizard, do the following:

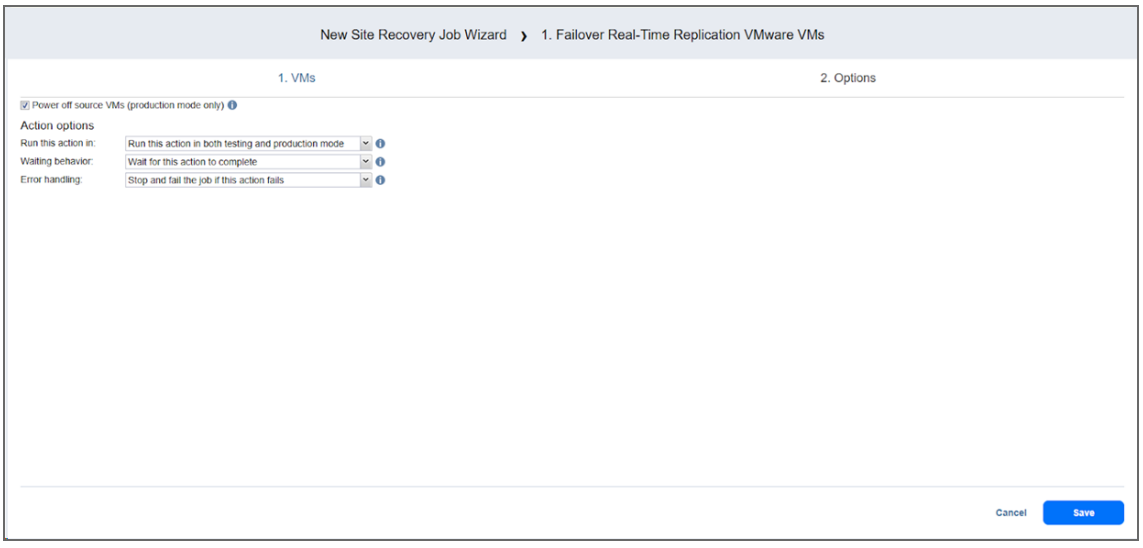
1. Ensure the **Power off source VMs** option is selected to power off the production workloads during the failover (available in production mode only).

2. In **the Action options** section, set the options for your action. Refer to "Site Recovery Job Wizard: Actions" on page 1149 for details.

**Notes**

- A Site Recovery job can be executed in test and production modes.
  - The Test mode is intended for verifying the Site Recovery job workflow and results. The Test mode allows you to run the job on demand or by scheduling it.
  - The Production mode is used for recovering the environment in case of a disaster. Customers can only execute a Site Recovery job in production mode on demand.
- The **Real-Time Replication Failover** action will operate in **Test Failover** mode if configured to run in **testing mode** or **both testing and production modes**.
- The **Real-Time Replication Failover** action will operate in **Failover mode** if configured to run in **production mode** or **both testing and production modes**.
- When the action **in Test failover mode** is done, the RTR job will keep running. As for production failover, the RTR job will be stopped and the RTR replica will enter into the failover state. To reset the failover state, go to the **Data Protection** > **Job overview** tab, and select the RTR job. The **Reset replica failover state** dialog box opens. Click **Reset & Run** to start the job.

3. Click **Save**.




The **Options** page closes and your **Failover Real-Time Replication VMware VMs** action is added to the Site Recovery Job.

# Failback VMware VMs Action

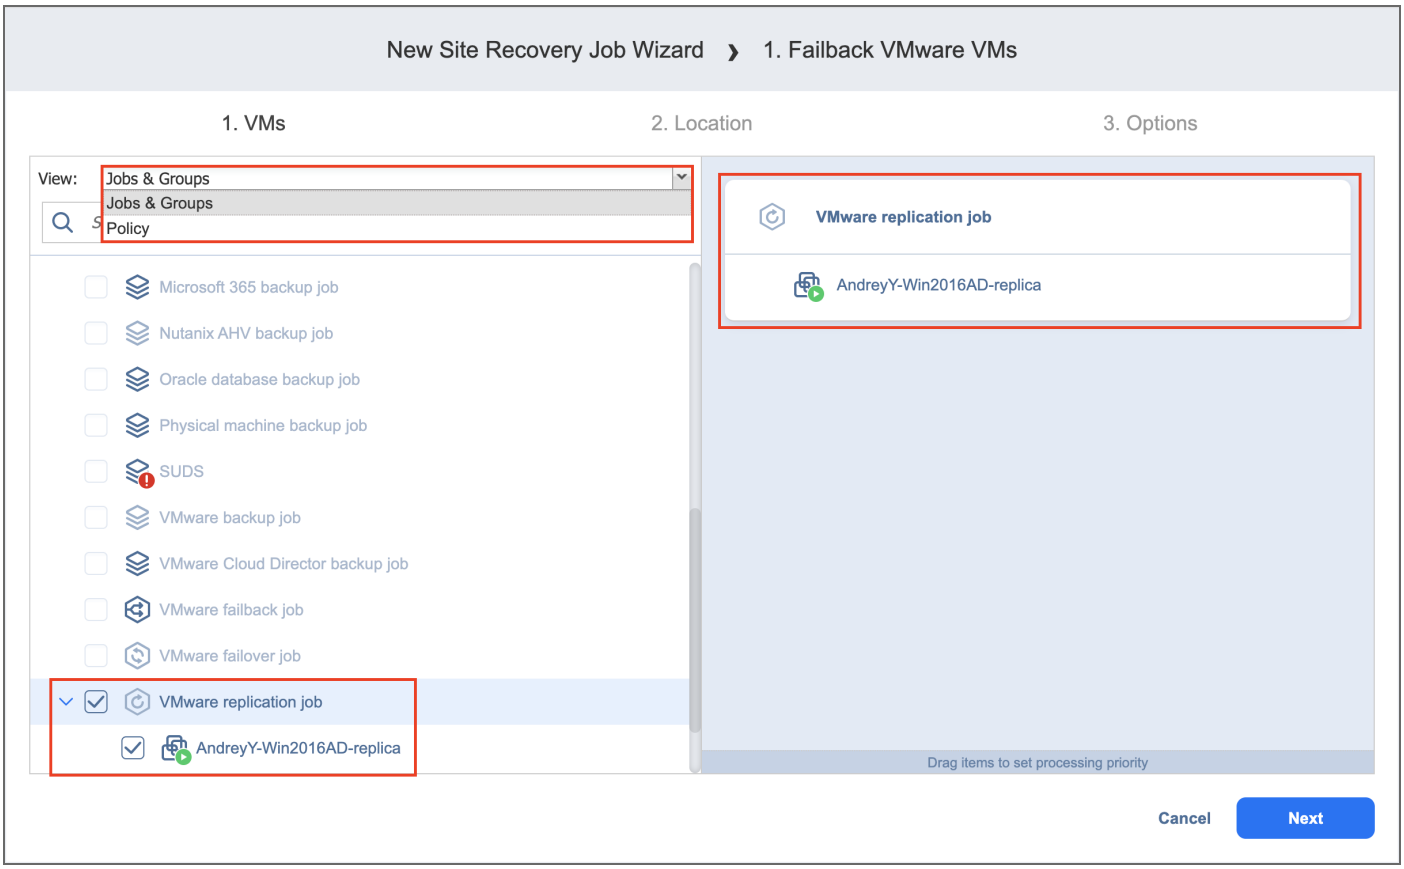
Please refer to the following sections on adding a Failback VMware VMs action to your Site Recovery job:

- Failback VMware VMs: VMs
- Failback VMware VMs: Location
- Failback VMware VMs: Options

Failback VMware VMs: VMs

On this page of the wizard, do the following:

1. In the left pane of the page, choose either of the following inventory views:
   - **Jobs & Groups**: When selected, the inventory tree displays groups, jobs, and backups of the appropriate hypervisor. You can select the required replica from the list of replication jobs.
   - **Policy**: When selected, job policies can be used. Refer to "Managing Job Policies" on page 382 for details. If the items were selected in alternate views, a dialog box opens warning you that switching to the **Policy** view will reset your current selection. Click **Switch View** to confirm switching to the **Policy** view. Make sure that at least one item matches the available set of policy rules. Refer to "Managing Policy Rules" on page 385 for details.
2. Select one or more replica VMware VMs in the left pane of the page.
3. Click **Next** to go to the **Location** page of the wizard.



Failback VMware VMs: Location

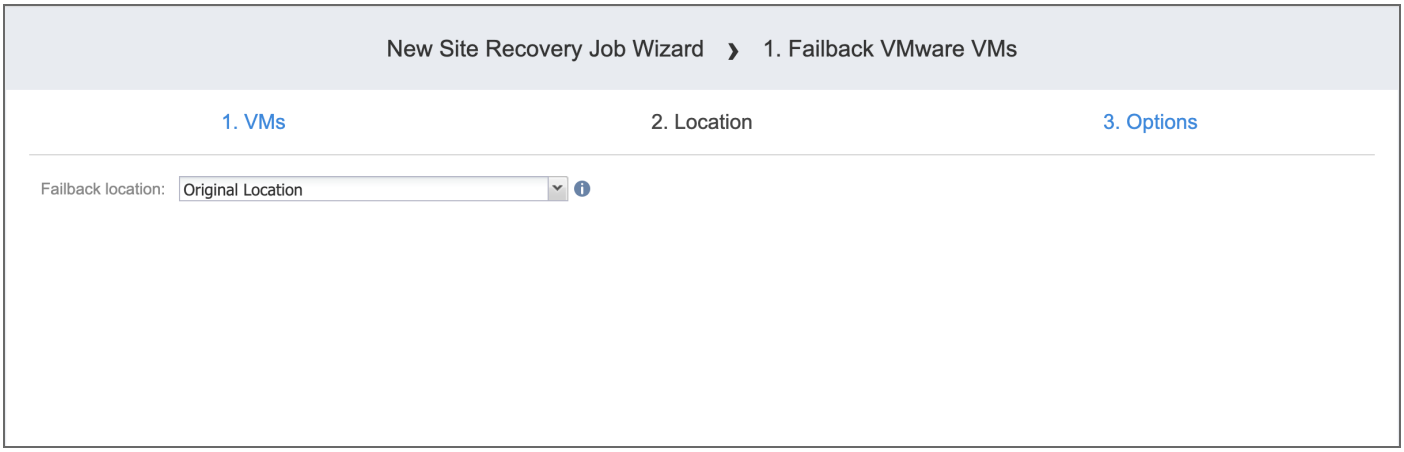
Please refer to the following subsections on providing a location for your Failback VMware VMs action:

- “Using Original VMware Failback Location” below
- “Creating New VMware Failback Location” below
- “Reusing Existing VMware Failback Location” on the next page

# Using Original VMware Failback Location

To use the original location for your **Failback VMware VMs** action, follow the steps below:

1. In the **Failback location** list, choose **Original Location**.
2. Click **Next** to go to the **Options** page.

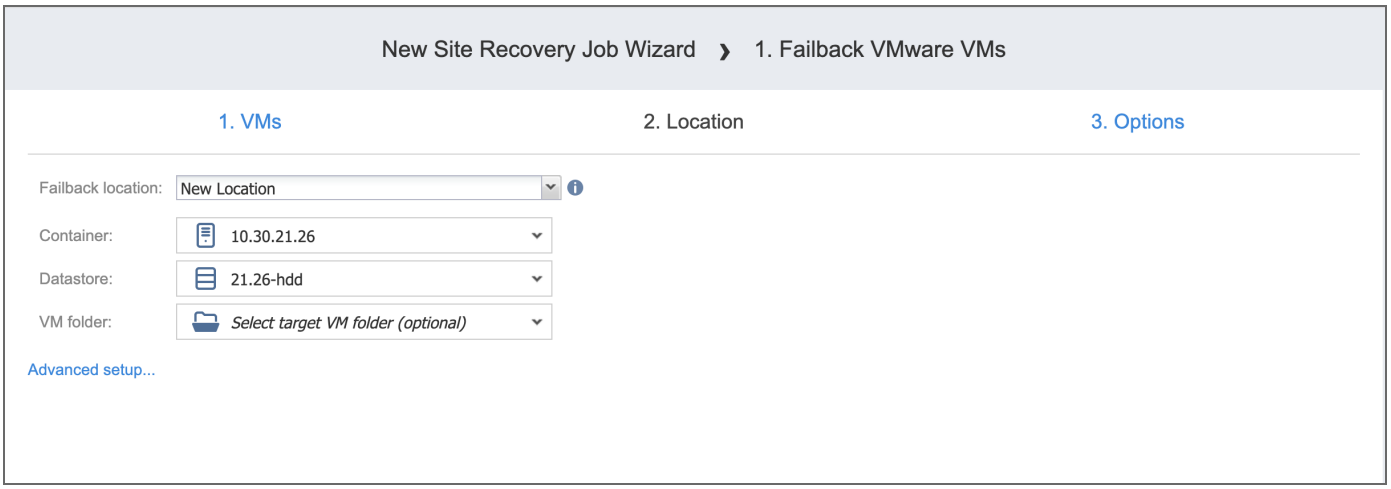



# Creating New VMware Failback Location

To create a new location for your **Failback VMware VMs** action, follow the steps below:

1. In the **Failback location** list, choose **New Location**.
2. In the **Container** list, choose a container for your failback location.
3. In the **Datastore** list, choose a datastore for your failback location.

4. Click **Next** to go to the **Options** page.




# Reusing Existing VMware Failback Location

To reuse existing locations for your **Failback VMware VMs** action, follow the steps below:

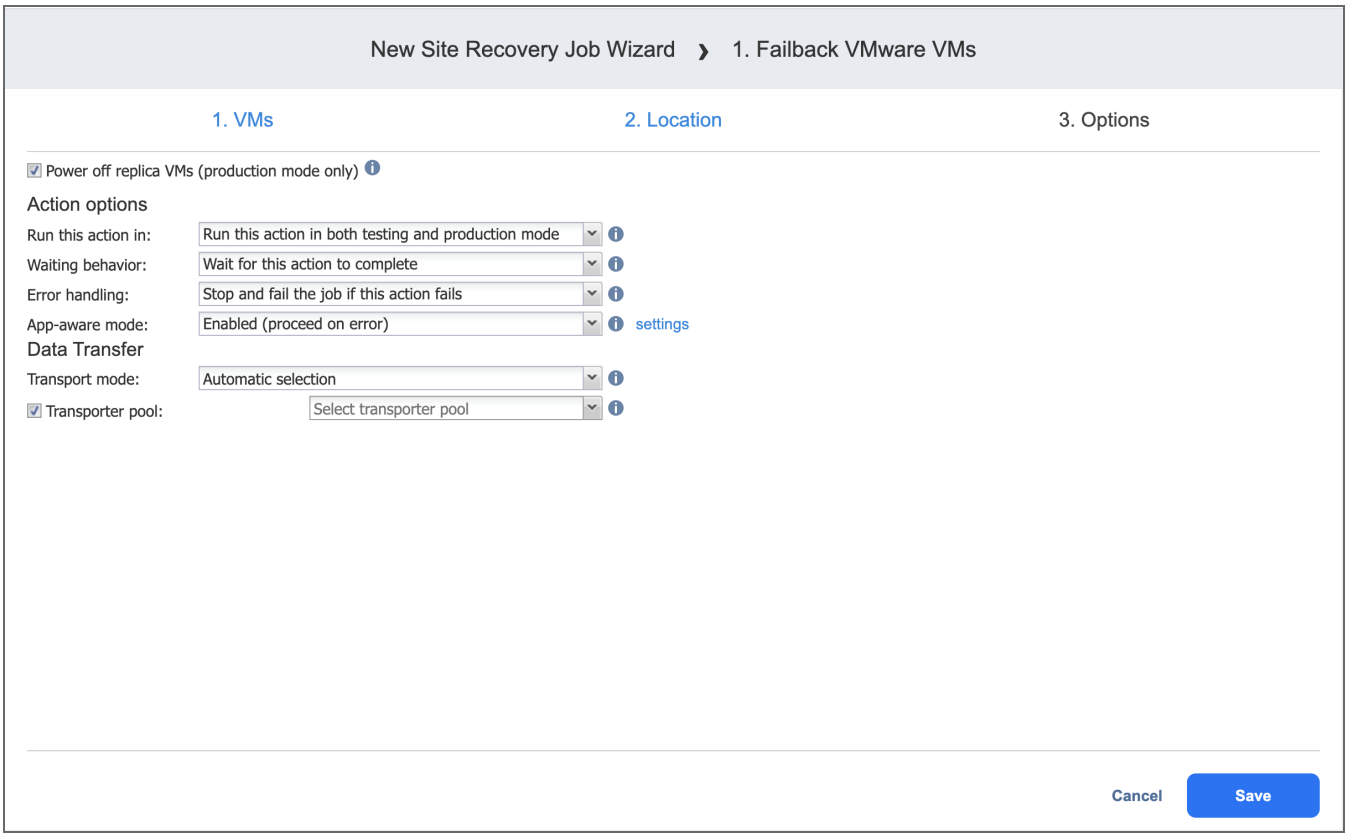
1. In the **Failback location** list, choose **New Location**.
2. Click **Advanced Options**.
3. Click on the necessary source VMs to expand them.
4. For the required VM, click **Use existing target VM** to select it.
5. The **Virtual Machine** list updates to include VMs. Select the VM that should be used as a target.
6. After setting the location advanced options, click **Next** to go to the **Options** page.

Failback VMware VMs: Options

On this page of the wizard, do the following:

1. Select **Power off replica VMs** to power off the production workloads during the failover. Please note that this will take place only in production mode; no workloads will be powered on in test mode.
2. In the *Action options* and *Data Transfer* sections, set the options for your action. Refer to "Site Recovery Job Wizard: Actions" on page 1149 for details.

3. Click **Save**.




The **Options** page closes and your **Failback VMware VMs** action is added to the Site Recovery job.

# Start VMs / Instances Action

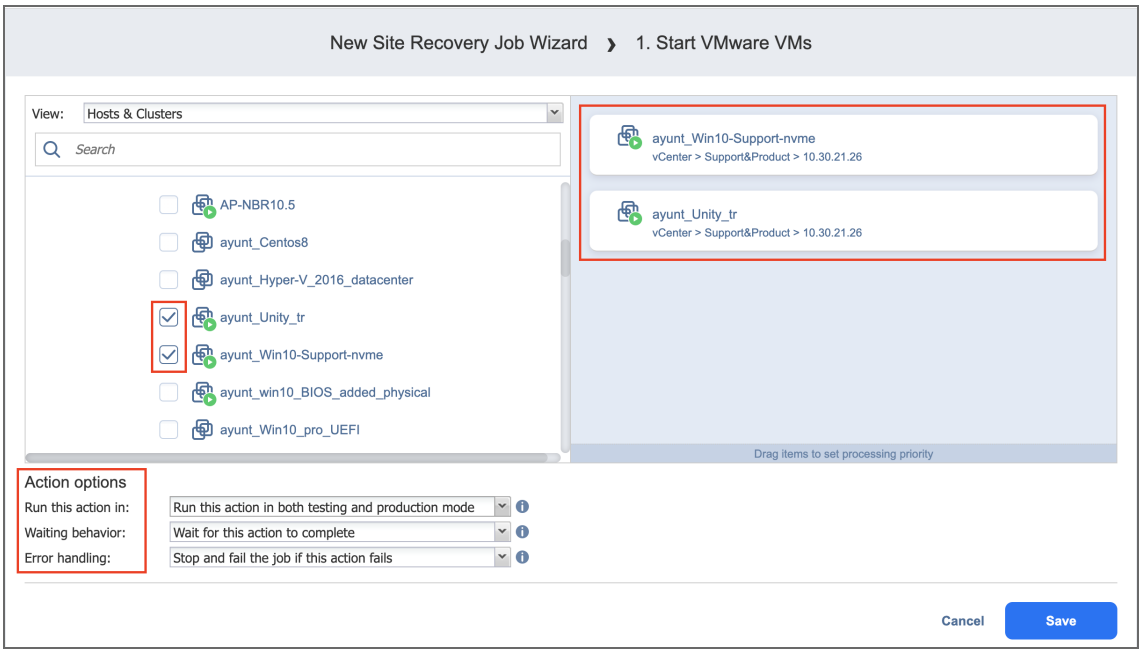
The topic covers the following actions of a Site Recovery Job:

- Start VMware VMs
- Start Microsoft Hyper-V VMs
- Start Amazon EC2 Instances

On the wizard's page that opens, do the following:

1. In the left pane of the page, choose either of the following inventory views:
   - **Hosts & Clusters**: Not available for AWS EC2. When chosen, the inventory tree displays all containers and VMs.
   - **VMs & Templates**. Available for VMware only. When chosen, the inventory tree displays VMware hosts, VMs, and VM templates.
   - **AWS Accounts**: Available for AWS EC2 only. When chosen, the inventory tree displays all AWS EC2 accounts along with their regions and available instances.
   - **Policy**: When selected, job policies can be used. Refer to Managing Job Policies for details.

   **Note**
   Switching to an alternative view resets your selection in the current view.




2. Optionally, you can filter the inventory tree by entering a string into the **Search** box. You can enter a part of or the entire name of the item.
3. Select at least one VMware/Hyper-V VM or Amazon EC2 instance in the inventory tree.
4. Set the action options. Refer to "Site Recovery Job Wizard: Actions" on page 1149 for details.

5. Click **Save**. The page of the wizard closes and your **Start VMs / Instances** action is added to the Site Recovery Job.

# Stop VMs / Instances Action

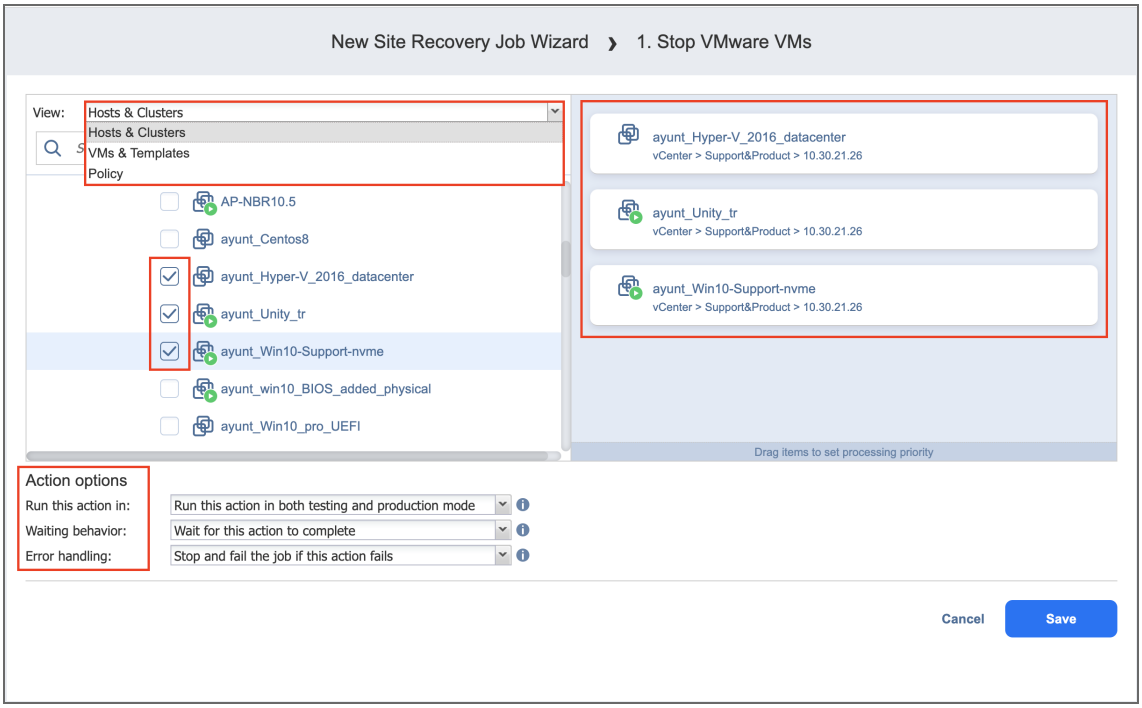
The topic covers the following actions of your site recovery job:

- Stop VMware VMs
- Stop Microsoft Hyper-V VMs
- Stop Amazon EC2 Instances

On the **Stop VMs / Instances** page that opens from the **Actions** page of the Site Recovery Job Wizard, do the following:

1. In the left pane of the page, choose one of the following inventory views:
   - **Hosts & Clusters**: Not available for AWS EC2. When chosen, the inventory tree displays all containers and VMs.
   - **VMs & Templates view**: Available for VMware only. When chosen, the inventory tree displays VMware hosts, VMs, and VM templates.
   - **AWS Accounts**: Available for AWS EC2 only. When chosen, the inventory tree displays all AWS EC2 accounts along with their regions and available instances.
   - **Policy**: When selected, job policies can be used. Refer to "Managing Job Policies" on page 382 for details.

   **Note**
   Switching to an alternative view resets your selection in the current view.




2. Optionally, you can filter the inventory tree by entering a string into the **Search** box. You can enter a part of the entire name of the item.

3. Select at least one VMware/Hyper-V VM or Amazon EC2 instance in the inventory tree.
4. Set action options. Refer to “Site Recovery Job Wizard: Actions” on page 1149 for details.
5. Click **Save**.

The page of the wizard closes, and your **Stop VMs / Instances** action is added to the site recovery job.

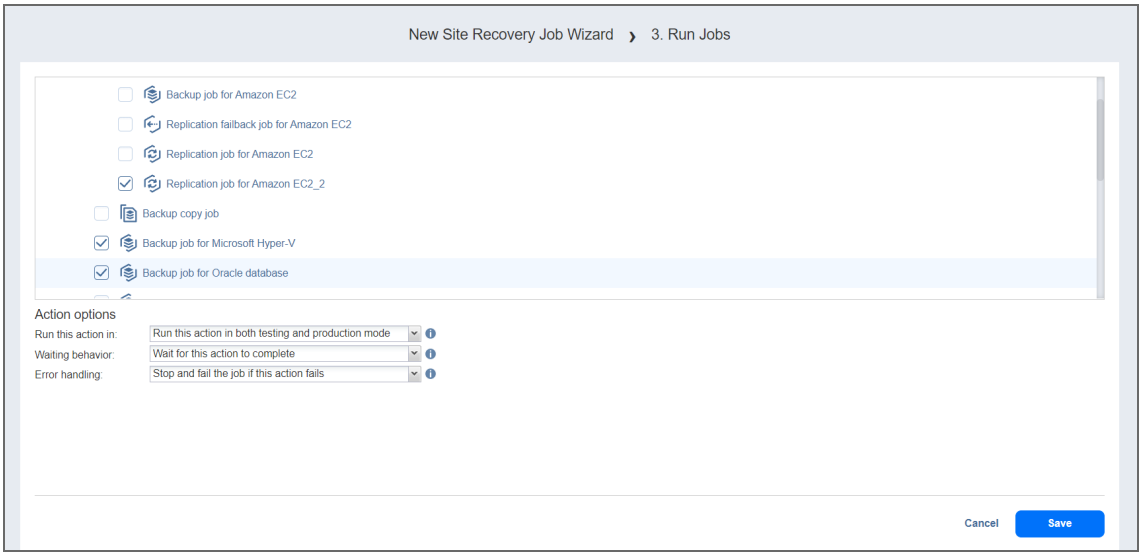
# Run / Stop Jobs Action

The topic covers the following actions of a Site Recovery Job:

- Run Jobs
- Stop Jobs

On the **Run / Stop Jobs** page that opens from the **Actions** page of the Site Recovery Job Wizard, do the following:

1. Select at least one item in the list of jobs.
2. Configure action options. Refer to "Site Recovery Job Wizard: Actions" on page 1149 for details.
3. Click **Save**.

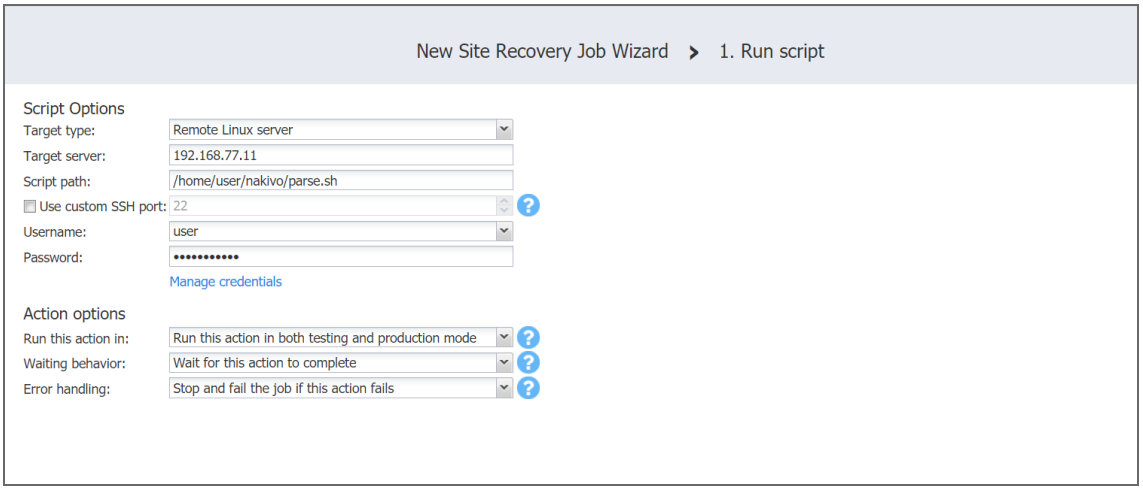



The page of the wizard closes and your **Run / Stop Jobs** action is added to the Site Recovery Job.

## Run Script Action

To add the action to a Site Recovery Job, on the **Run Script** page of the wizard, do the following:

1. **Target type**: Choose one of the following target types for your script:
   - **This server (Director)**: The script will run on the machine where the Director is deployed. Provide the following options:
     1. **Script path**: A path to the script.
     2. **Username / Password**: Credentials for running your script on the machine.
   - **Remote Windows / Linux server**: The script will run on a remote Windows / Linux server. Provide the following options:
     1. **Target server**: The hostname or the IP address of the remote server.
     2. **Script path**: A path to the script on the remote server.
     3. **Username / Password**: Credentials for running your script on the remote server.
   - **VMware / Hyper-V VM / EC2 instance**: The script will run on a VMware or Hyper-V VM or an EC2 instance. Provide the following options:
     1. **Target VM / Instance**: Choose the required item from the inventory tree.
     2. **Script path**: A path to the script on the VMware or Hyper-V VM or EC2 instance.
     3. **Username / Password**: Credentials for running your script on the VMware or Hyper-V VM or EC2 instance.
2. Set the **Action options**. Refer to "Site Recovery Job Wizard: Actions" on page 1149 for details.
3. Click **Save**.




The page of the wizard closes and your **Run Script** action is added to the Site Recovery Job.

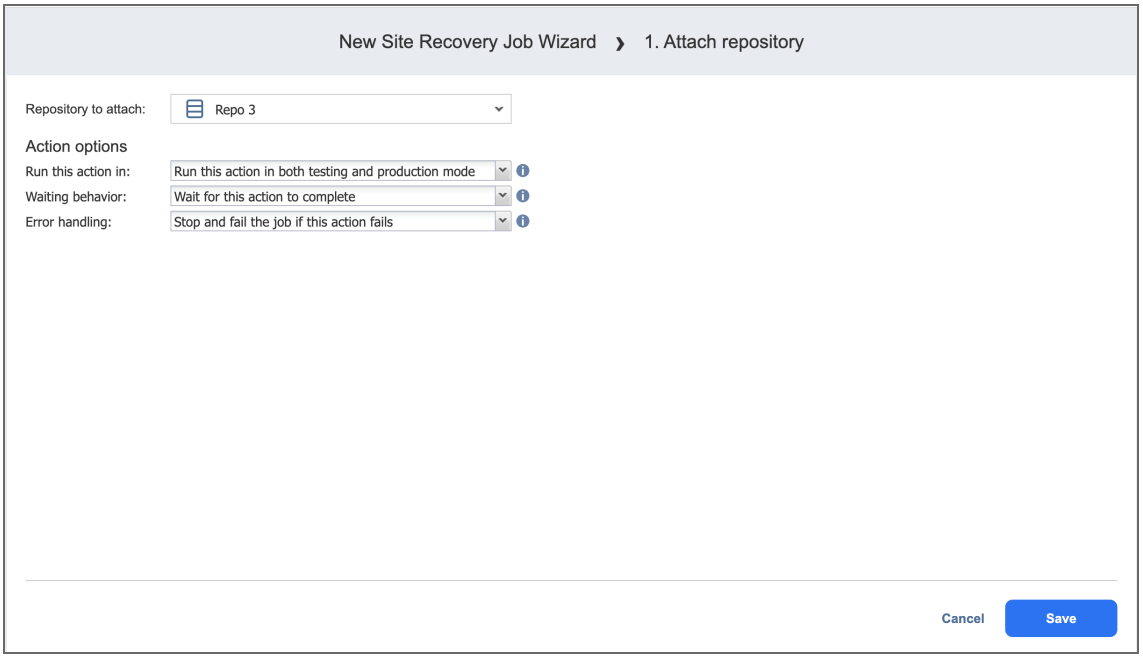
# Attach / Detach Repository Action

The topic covers the following actions of a Site Recovery Job:

- Attach Repository
- Detach Repository

On the **Attach / Detach Repository** page that opens from the **Actions** page of the Site Recovery Job Wizard, do the following:

1. Select a repository in the list of repositories.
2. Specify action options. Refer to Options Common to Most Actions for details.
3. Click **Save**.

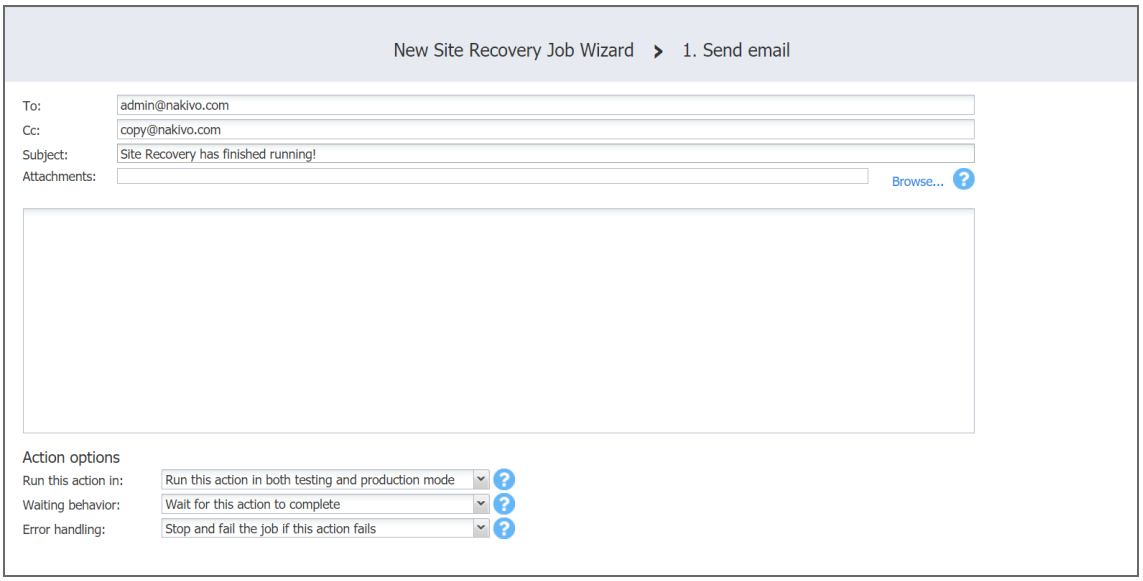



The page of the wizard closes and your **Attach / Detach Repository** action is added to the site recovery job.

## Send Email Action

On the **Send Email** page of the wizard, do the following to add the action to your site recovery job:

1. In the **To** box, enter a valid email address for the recipient.
2. In the **Cc** box, optionally you can enter an email address of the carbon copy recipient.
3. In the **Subject** box, optionally you can enter a subject of the mail.
4. Optionally, you can add attachments to your mail with the **Browse** button.
5. Enter your message text in the email body box.
6. Set up action options. Refer to "Site Recovery Job Wizard: Actions" on page 1149 for details.
7. Click **Save**.




The page of the wizard closes, and your **Send Email** action will be added to the site recovery job.

## Wait Action

To add a wait action to your Site Recovery job, on the **Wait** page of the wizard, do the following:

1. Enter the time to wait in minutes or hours.
2. Set the action options. Refer to "Site Recovery Job Wizard: Actions" on page 1149 for details.
3. Click **Save**.




The page of the wizard closes and your **Wait** action is added to the Site Recovery job.

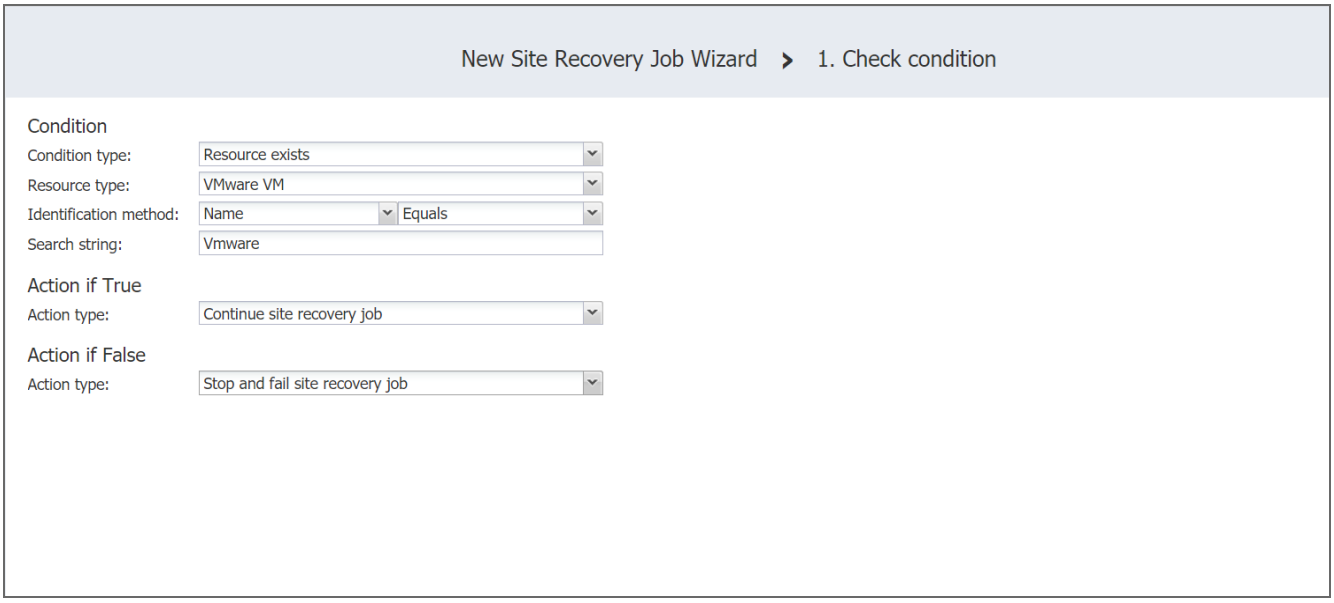
## Check Condition Action

On the **Check Condition** page of the wizard, do the following to add the action to your site recovery job:

1. Choose a condition type from the list:
    - **Resource exists**: This condition checks whether the specified resource exists. With this option, you have to set the following:
        a. Choose a resource type from the list:
            - **VMware VM**
            - **Hyper-V VM**
            - **EC2 instance**
        b. Define your identification method using two lists:
            i. In the first list, choose either **Name** of **ID**.
            ii. In the second list, choose one of the following:
                - **Equals**
                - **Contains**
                - **Starts with**
                - **Ends with**
        c. Enter your condition criterion in the **Search string** box.
    - **Resource is running**: This condition checks whether the specified resource is running. With this option, you have to set the options as described for the **Resource exists** option above.
    - **IP/hostname is reachable**: This condition checks whether the specified IP/hostname is reachable. With this option, you have to enter the following:
        a. Choose the source endpoint among the following:
            - **This server (Director)**: With this option, the condition checks whether the Director can reach the **IP/hostname** specified in the box below.
            - **Remote transporter**: With this option, the condition checks whether one of your remote Transporters can reach the **IP/hostname** specified in the box below.
        b. In the **IP/hostname** box, enter an IP address or a host name of the resource to be checked for reachability.
2. In the *Action if True* section of the page, choose an action to be taken if the condition criterion is satisfied:
    - **Continue site recovery job**: Your Site Recovery job will be continued.
    - **Stop and fail site recovery job**: Your Site Recovery job will be stopped as failed.
    - **End site recovery job successfully**: Your Site Recovery job will be ended as successful.

- **Go to another site recovery Job action**: Another action of your Site Recovery job will be initiated. If you choose this action type, a new box will open to allow you to choose the necessary action.

3. In the *Action if False* section of the page, choose an action to be taken if the condition criterion is not satisfied. Available options are similar to those described in the *Action if True* section above.
4. Click **Save**.




The page of the wizard closes and your **Check Condition** action is added to the Site Recovery job.

## Site Recovery Job Wizard: Networks

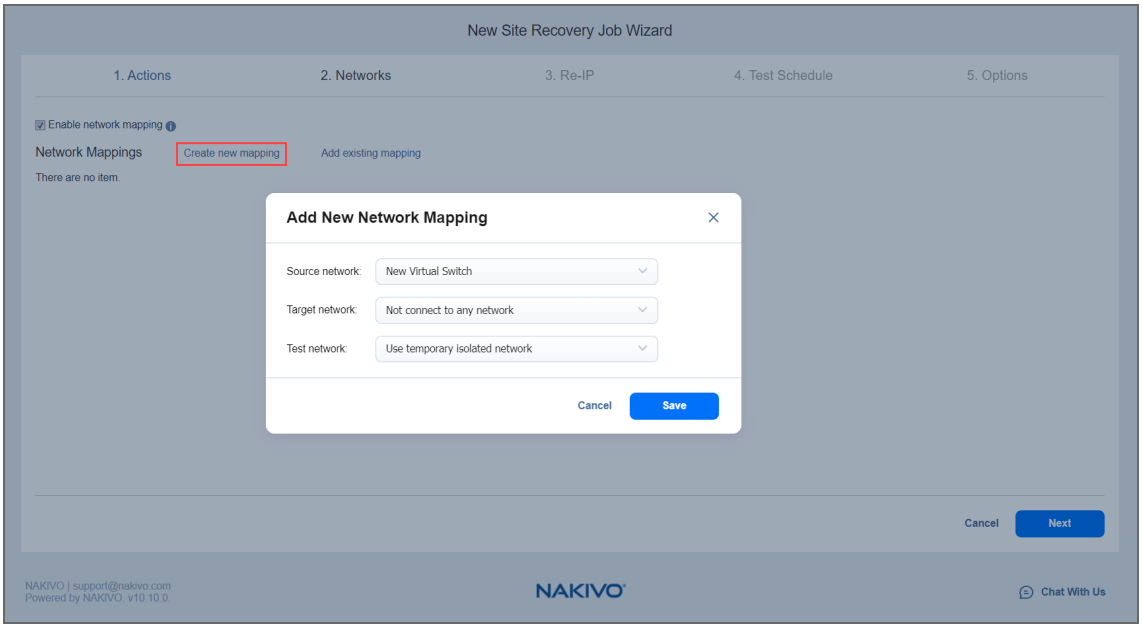
On the **Networks** page of the **Site Recovery Job Wizard** you can map source VM virtual networks to appropriate target virtual networks and test networks.

Please proceed as follows:

1. Select **Enable network mapping**.

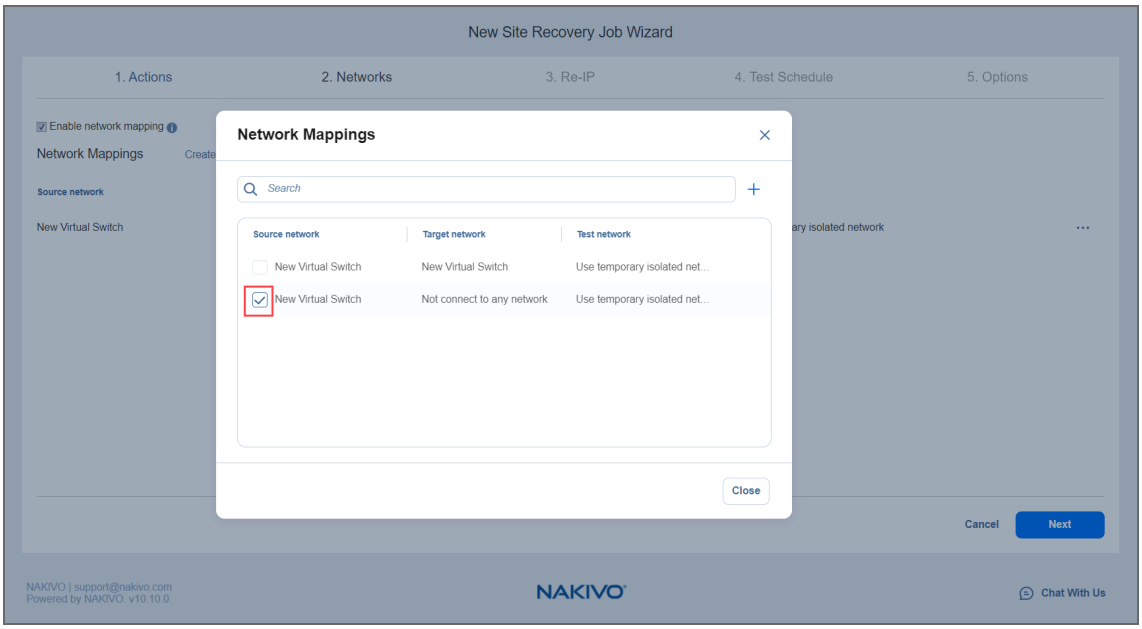
> **Note**
> A failover or a failback action needs to be on the actions list to allow enabling network mapping for your site recovery job.

2. The *Network Mapping* section opens. You have the following options:
   - Create a new network mapping:
     a. Click **Create new mapping**.
     b. The **Add New Network Mapping** dialog opens. Choose a source network, a target network and a test network, and click **Save**.




   - Add an existing network mapping:
     a. Click **Add existing mapping**.
     b. The **Network Mappings** dialog opens. Choose an appropriate network mapping and close

the dialog box.




- Edit an existing network mapping:
  a. Hover the pointer over the necessary item in the **Network Mappings** list and then click the **Edit** button to the right of the item.
  b. The **Edit Network Mapping** dialog box opens. Choose appropriate items from the **Target network** and the **Test network** lists and then click **Save**.




- Delete an existing mapping: hover the pointer over the necessary item in the **Network Mappings** list and then click the **Delete** icon to the right of the item.




- To leave the list of existing network mappings intact, go to the next page of the wizard.

3. Click **Next** to go to the next page of the wizard.

# Site Recovery Job Wizard: Re-IP

In the **Re-IP** page of the **Site Recovery Job Wizard** you can map a source VM IP address to a specific target IP address.

Please proceed as follows:

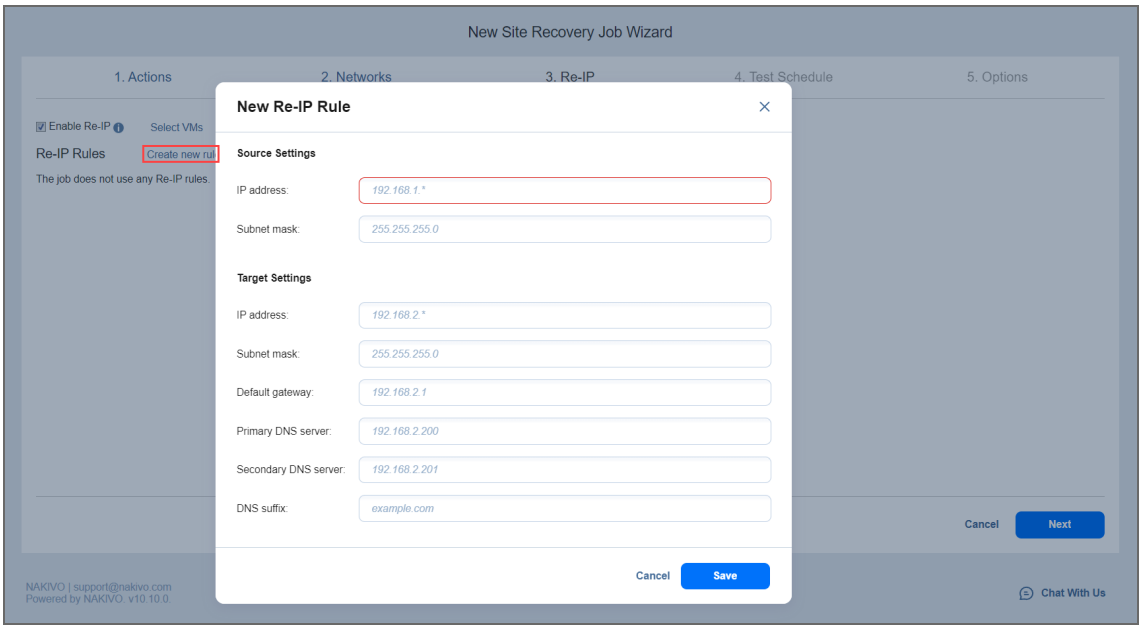
1. Select **Enable Re-IP**.

> **Note**
> A failover or a failback action needs to be in the actions list to allow enabling Re-IP for your site recovery job.

2. The *Re-IP Rules* section opens. Click the **Select VMs** link.
3. The **Re-IP** dialog opens. In the list of your source VMs, select at least one, select the credentials to be used for each VM, and close the dialog.

> **Note**
> Re-IP rules will be applied only to VMs that have a static IPv4 address configured.

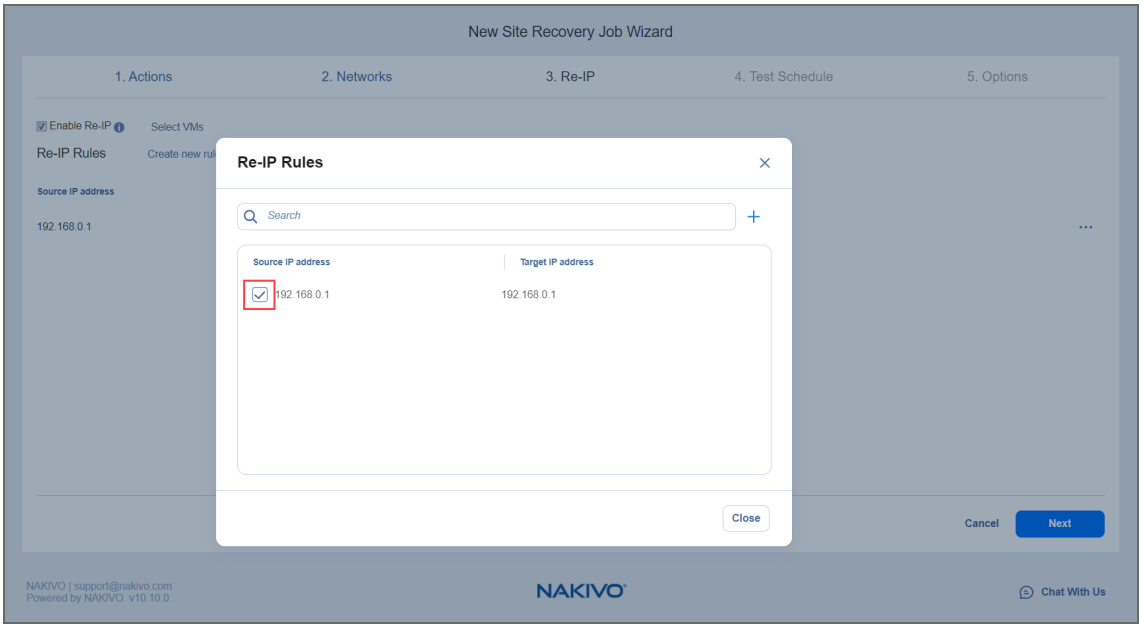
4. You have the following options:
   - Create a new rule:
     a. Click **Create new rule**.
     b. The **New Re-IP Rule** dialog opens. Enter source and target settings for the Re-IP rule and click **Save**.



**Note**

You can use wildcards for IP addresses. For example, when the 192.168.1.* -> 10.30.30.* Re-IP rule is available, it will change the source VM IP address like 192.168.1.50 to the 10.30.30.50 IP address, for your site recovery job. When there are several Re-IP rules applicable to your source VM, the application will define the most suitable one and apply it to the source VM IP address.

- Add an existing rule:
  a. Click **Add existing rule**.
  b. The **Re-IP Rules** dialog box opens. Choose an appropriate Re-IP rule and close the dialog.




- Edit an existing Re-IP rule:
  a. Hover the pointer over the necessary item in the **Re-IP Rules** list and then click the **Edit** button to the right of the item.
  b. The **Edit Re-IP Rule** dialog opens. Edit the necessary properties of the Re-IP rule and then

click **Save**.




- Delete an existing mapping: hover the pointer over the necessary item in the **Re-IP Rules** list and then click the **Delete** icon to the right of the item.




- To leave the list of existing Re-IP rules intact, go to the next page of the wizard.

5. Click **Next** to go to the next page of the wizard.

# Site Recovery Job Wizard: Test Schedule

On the **Test Schedule** page of the **Site Recovery Job Wizard** you can schedule testing your site recovery job.

Please refer to the following sections for details:


- "Disabling Site Recovery Job Test Schedule" below
- "Daily Site Recovery Job Testing" below
- "Monthly or Yearly Site Recovery Job Testing" on the next page
- "Periodic Site Recovery Job Testing" on page 1178
- "Chained Site Recovery Job" on page 1179
- "Additional Schedule" on page 1180


## Disabling Site Recovery Job Test Schedule

If you only want to start the site recovery job manually (without any test schedule), select the **Do not schedule, test on demand** check box.



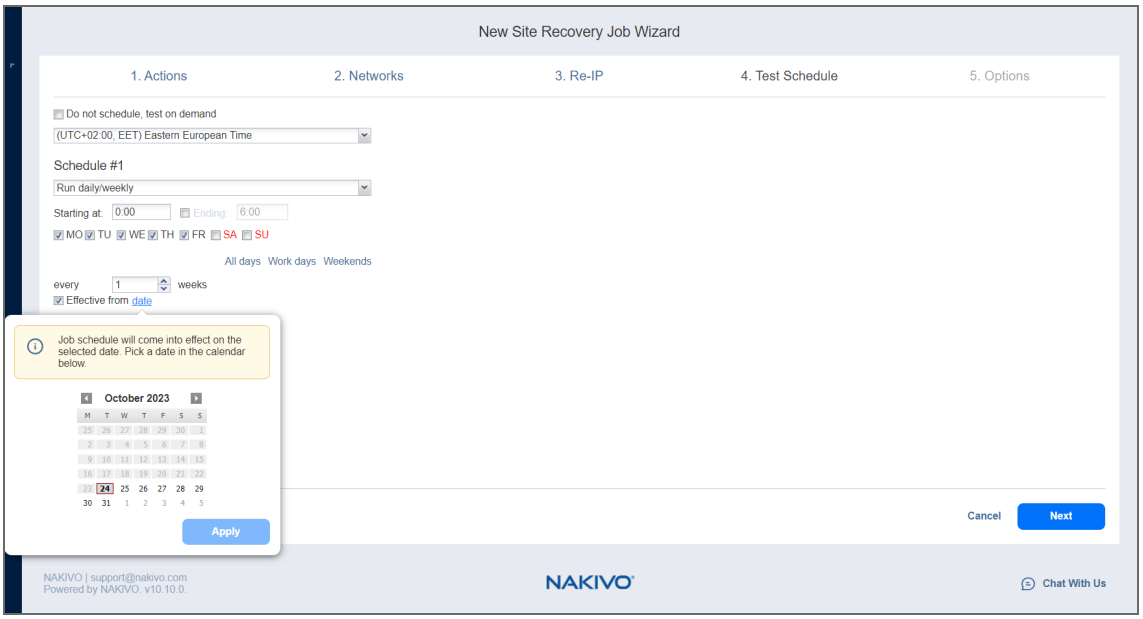

Then click **Next** to go to the **Options** page of the **Site Recovery Job Wizard**.

## Daily Site Recovery Job Testing

To test your site recovery job once a day, do the following:

- Choose a time zone that should be used for the site recovery job start and end times from the time zone list.
- Choose **Run daily/weekly** from the **Schedule #1** list.
- Specify the time when the site recovery job should be started in the **Starting at** box.
- Specify the end time for the site recovery job in the **Ending** box. If the site recovery job has not completed by the time specified, the site recovery job will be stopped.

- Select the days of the week during which the site recovery job will be started.
- To specify a date when the job test schedule comes into effect, click **Effective from,** click **date** and then pick a date in the calendar that opens.
- Click **Next** to go to the **Options** page of the **Site Recovery Job Wizard**.




Monthly or Yearly Site Recovery Job Testing

To test your site recovery job monthly or yearly, do the following:

- Choose **Monthly/yearly** from the schedule list.
- Choose a time zone that should be used for the job start and end times, in the list of available time zones.
- Specify a site recovery job start schedule in the appropriate boxes.
- Specify the time when the site recovery job should be started, in the **Starting at** box.
- Specify the end time for the site recovery job, in the **Ending** box. If the site recovery job has not completed by the time specified, the site recovery job will be stopped.
- Select the days of the week during which the job will be started.
- To specify a date when the job test schedule comes into effect, click **Effective from,** click **Date** and then pick a date in the calendar that opens.

- Click **Next** to go to the **Options** page of the **Site Recovery Job Wizard**.




Periodic Site Recovery Job Testing

To test your site recovery job multiple times per day, do the following:

- Choose a time zone that should be used for the site recovery job start and end times from the list of time zones.
- Choose **Run periodically** from the **Schedule #1** list and then choose a time period from the appropriate boxes.
- Specify the time when the site recovery job should be started in the **Starting at** box.
- Specify the end time for the job in the **Ending** box. If the site recovery job has not completed by the time specified, the site recovery job will be stopped.
- To specify a date when the job test schedule comes into effect, click **Effective from,** click **Date** and then pick a date in the calendar that opens.

- Click **Next** to go to the **Options** page of the **Site Recovery Job Wizard**.

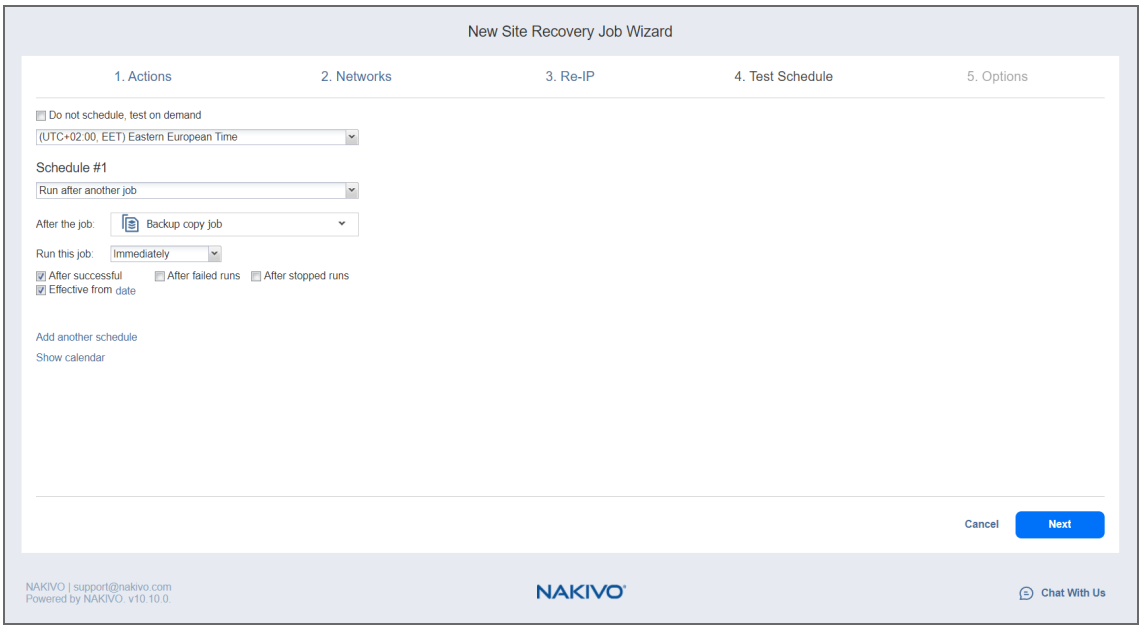



Chained Site Recovery Job

To run the site recovery job after a previous job has completed, do the following:

1. Choose **Run after another job** from the **Schedule #1** list.
2. Set the options as follows:
    - **After the job**: select a job after which the current site recovery job will be started.
    - **Run this job**: Choose whether to run the current site recovery job immediately after the previous job has completed, or specify a delay.
    - **After successful runs**: If selected, the site recovery job will run if the previous job has completed successfully.
    - **After failed runs**: If selected, the site recovery job will run if the previous job has failed.
    - **After stopped runs**: If selected, the site recovery job will run if the previous job has been stopped.
3. To specify a date when the job test schedule comes into effect, click **Effective from**, click **date** and then pick a date in the calendar that opens.

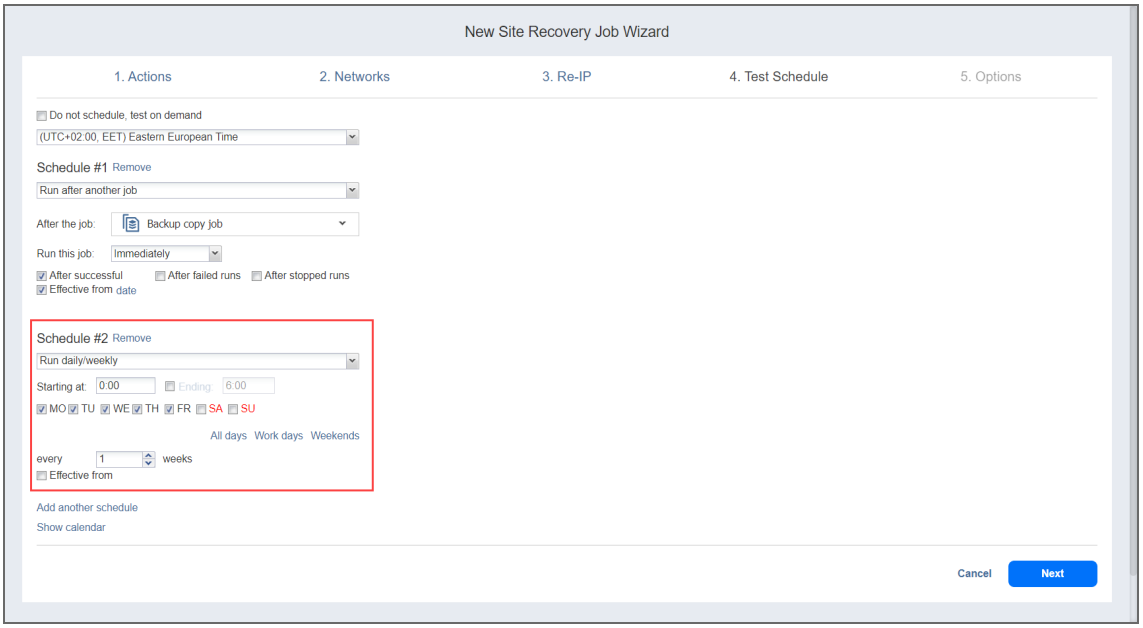
4. Click **Next** to go to the **Options** page of the **Site Recovery Job Wizard**.




## Additional Schedule

To add more than one schedule to your site recovery job, do the following:

1. Click **Add another schedule**.
2. The *Schedule #2* section opens. Proceed with instructions provided in the sections above.
3. When ready with adding an additional schedule to your site recovery job, click **Next**.



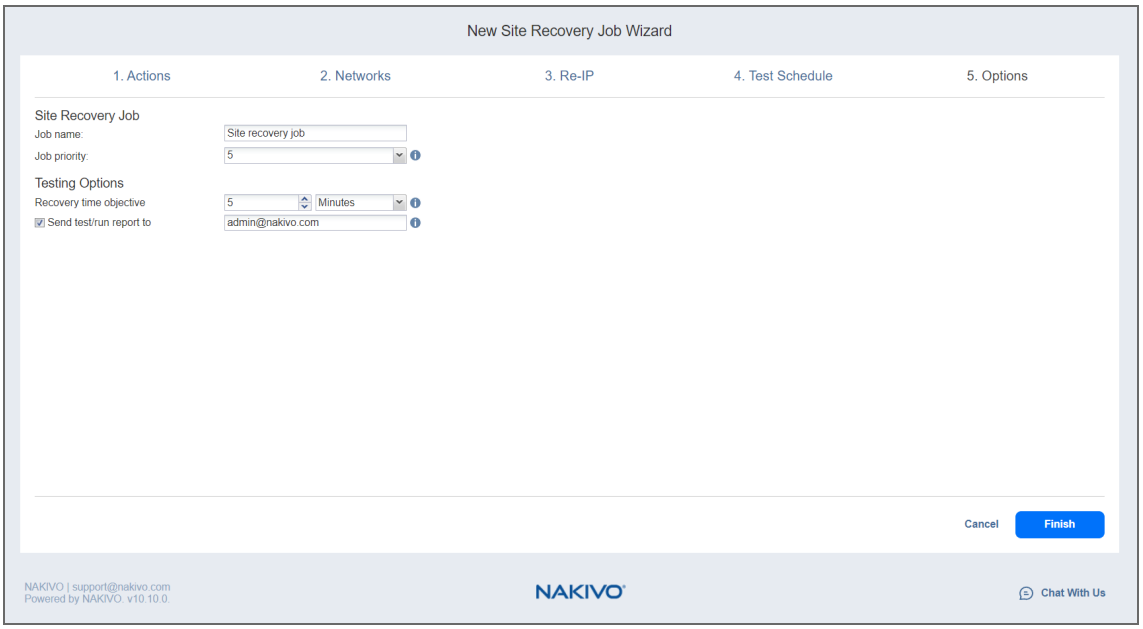
## Site Recovery Job Wizard: Options

On the **Options** page of the **Site Recovery Job Wizard** you can specify the options of the site recovery job. Proceed as follows:

1. In the *Site Recovery Job* section, specify a name for the site recovery job. If the site recovery job contains failback actions that require a Transporter, you may also specify a job priority level from *1* to *5*, with *1* being the highest priority. Jobs with higher priority levels are prioritized by **Transporters** during job processing.

   > **Note**
   > The **Job Priority** option is only available in the Enterprise, Enterprise Essentials, Enterprise Plus, MSP Enterprise, and MSP Enterprise Plus editions.

2. In the *Testing Options* section:
   - **Recovery time objective**: Enter the amount of time allowed for the site recovery job test to complete. The report will inform you of whether this objective has been met.
   - **Send test/run report to**: When selected, this option enables sending email notifications to the specified recipients. Use the semi-colon character to separate multiple email addresses.
3. Click **Finish** to complete creating the site recovery job.




The **Site Recovery Job Wizard** closes and the site recovery job appears in the list of NAKIVO Backup & Replication jobs.

# Running Site Recovery Job

The section includes the following topics:

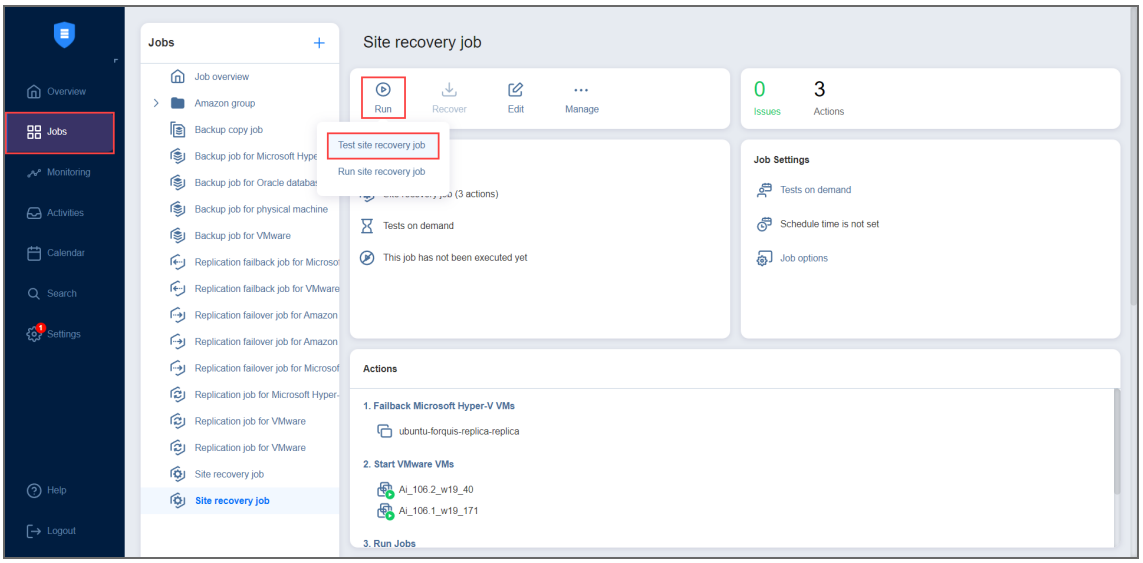
- Running Site Recovery Job in Test Mode
- Running Site Recovery Job in Production Mode

## Running Site Recovery Job in Test Mode

Running your site recovery job in the test mode allows you to verify the site recovery workflow and results.

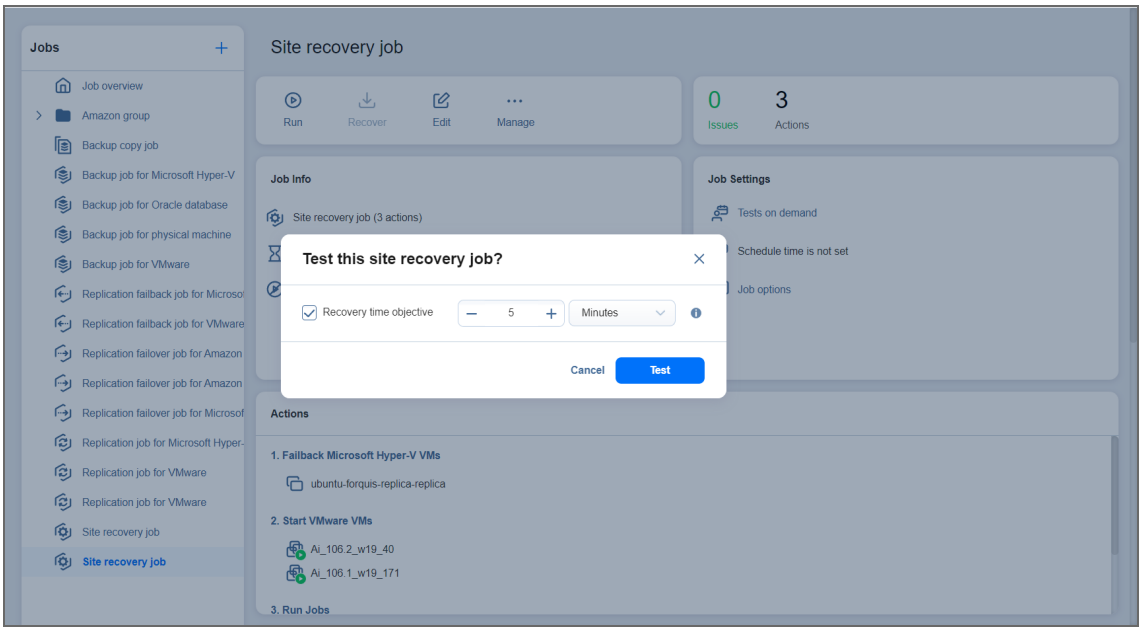
Please follow the steps below to run your site recovery Job in the test mode:

1. In the **Jobs** menu, select your site recovery job and then click the **Run** button.
2. In the dialog that opens, click **Test site recovery job**.




3. The **Recovery time objective** dialog opens. Here you can:
    - Disable/enable the **Recovery time objective** option.
    - If the **Recovery time objective** is enabled, modify the amount of time allowed for the job to be

completed.




4. Click **Test** when ready. The site recovery job starts running in the test mode.

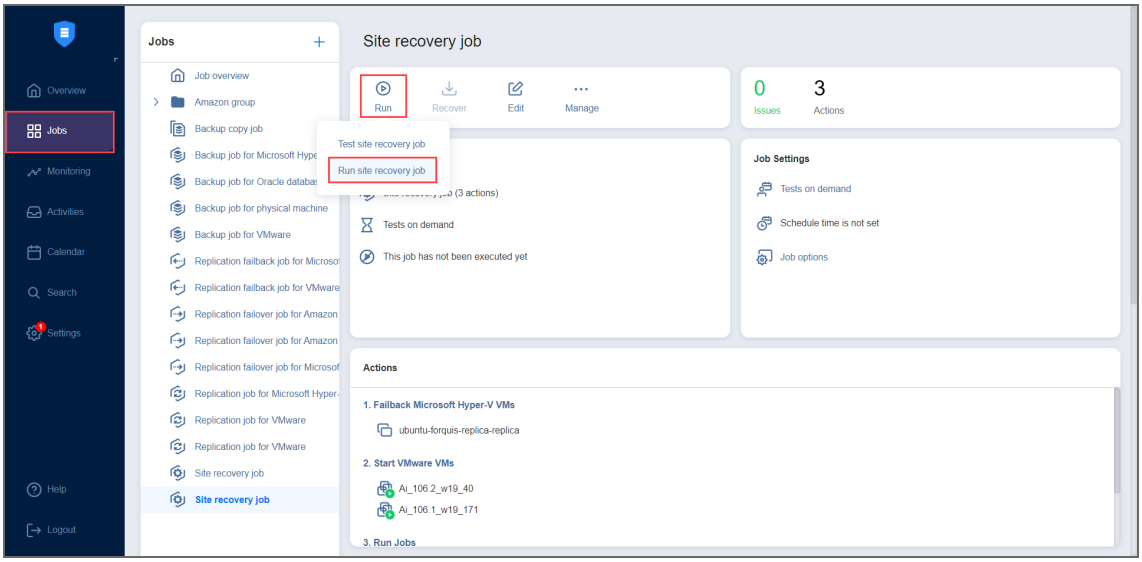
> **Note**
> In addition to testing the site recovery job on demand, testing can also be scheduled. Refer to Site Recovery Job Wizard: Test Schedule for details.

## Running Site Recovery Job in Production Mode

Running your site recovery job in the production mode allows you to recover your environment from disaster.

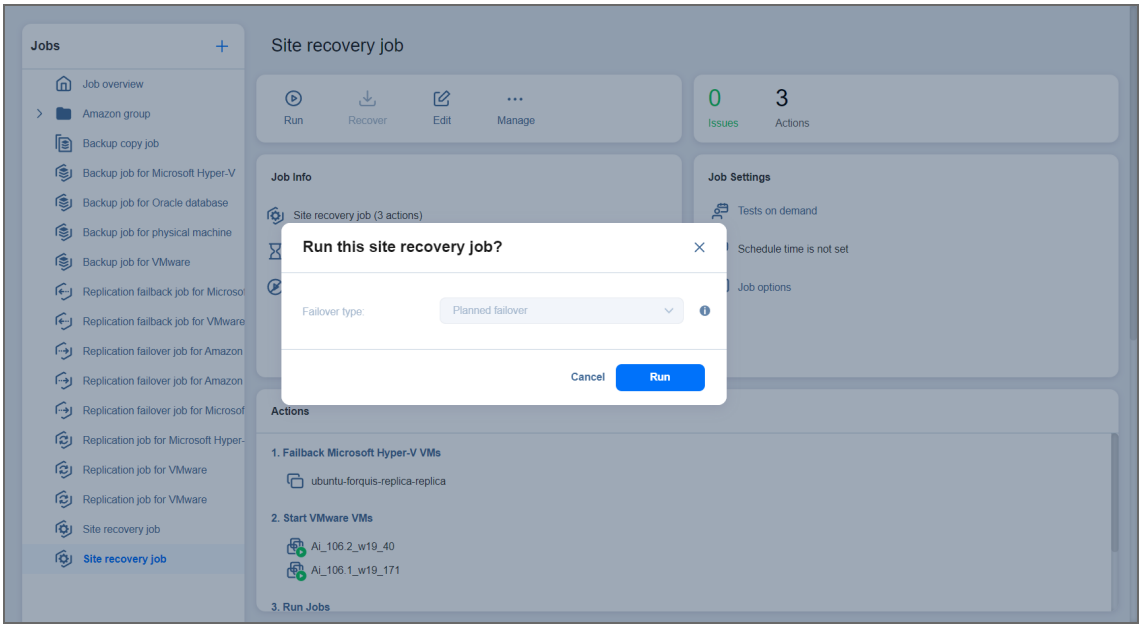
Please follow the steps below to run your site recovery job in the production mode:

1. In the **Jobs** menu, select your site recovery job and then click the **Run Job** button.
2. In the dialog that opens, click **Run site recovery job**.




3. The **Failover type** dialog opens. Choose either of the following failover types:
   - **Planned failover**: The application will sync replica data with the source VM before switching workloads to the replica.
   - **Emergency failover**: The application will switch workloads from the source VM to the replica immediately.

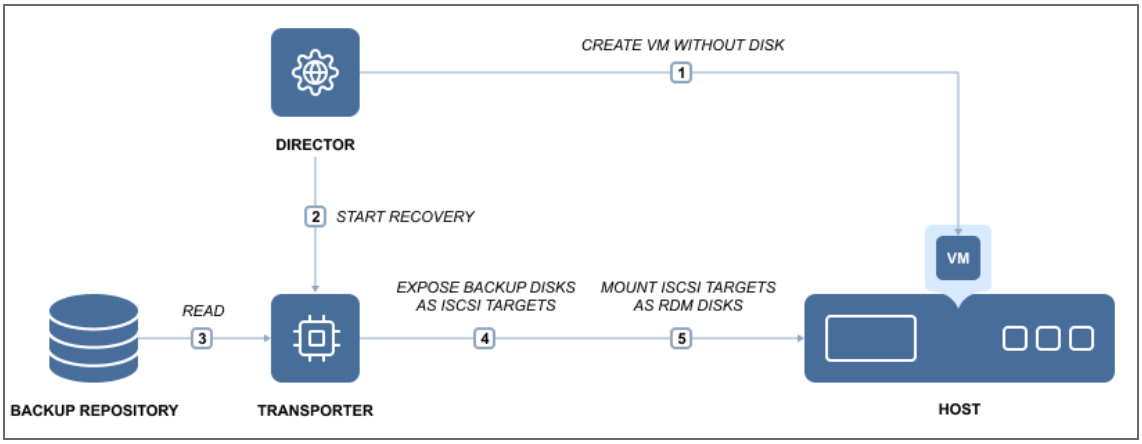
> **Note**
> The **Failover type** option is only available for site recovery jobs containing a **Failover** action.




4. Click **Run**. The site recovery job starts running in the production mode.

# Performing Flash Boot Recovery

The Flash boot feature allows you to run (boot) VMware and Hyper-V VMs directly from compressed and deduplicated VM backups, without recovering entire VMs first. When you boot a VM from a backup, NAKIVO Backup & Replication creates a new VM on the target server.



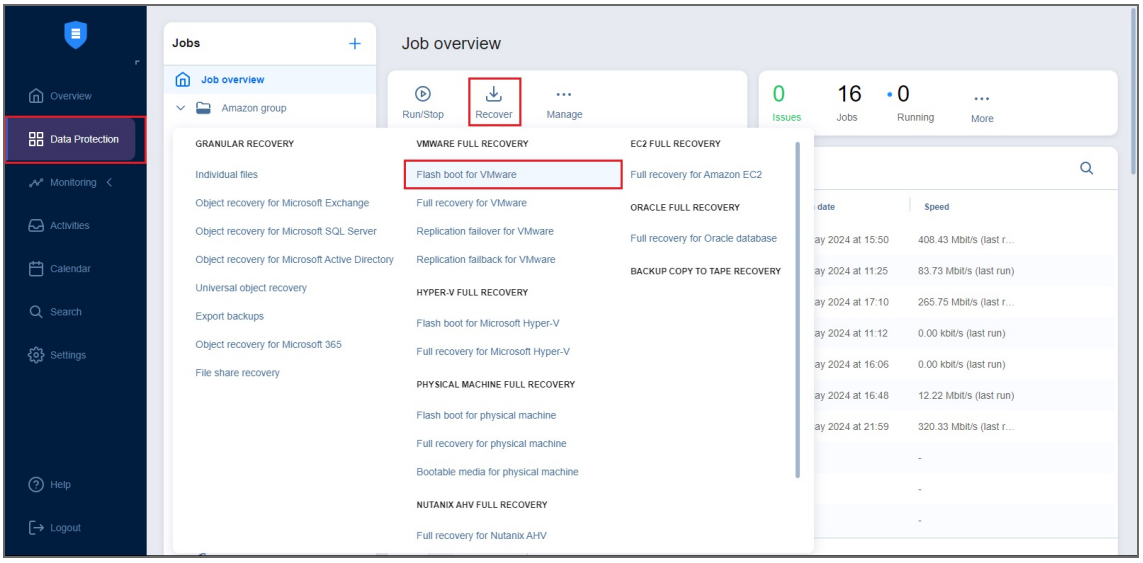

When a VMware VM is created, NAKIVO Backup & Replication takes a snapshot of the VM: this way all changes that occur to the VM are temporarily stored in the snapshot and discarded when you stop the job. When a Hyper-V VM is created, the application temporarily stores the changes to the VM in a disk-based write cache in the Backup Repository; changes are discarded when the job is stopped. For more information, refer to the following topics:

- “Creating VMware Flash Boot Jobs” on page 1186
- “Migrating Recovered VMs Using Flash Boot” on page 1204

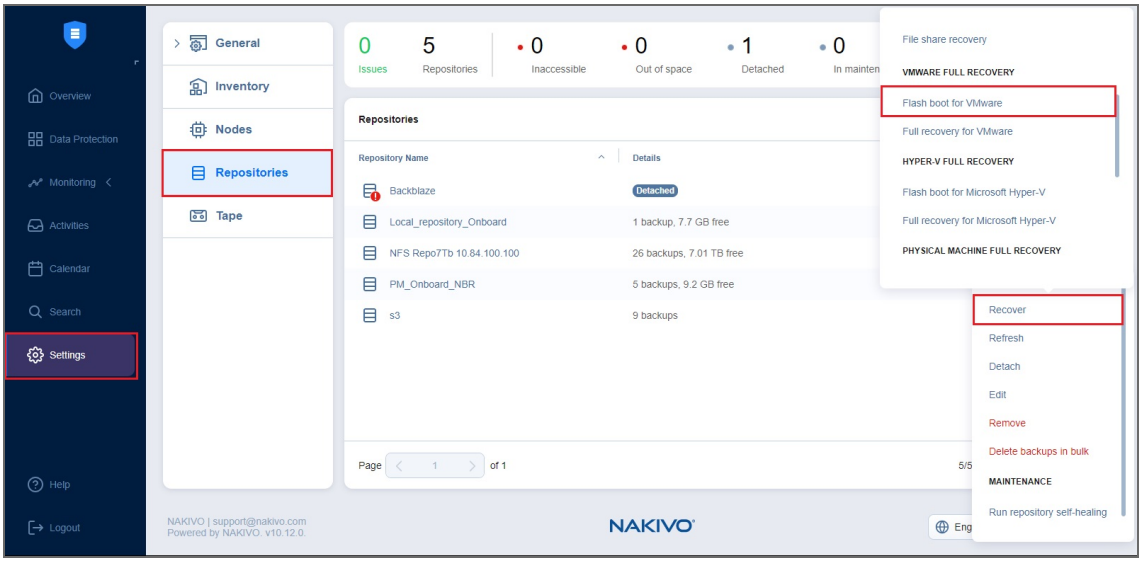
# Creating VMware Flash Boot Jobs

To create a VMware Flash Boot job, do one of the following:

- Open the Flash Boot wizard from the **Data Protection** menu by clicking **Recover** and then clicking **Flash boot for VMware**.




- Open the Flash Boot wizard from the **Repositories** tab in **Settings** by following the steps below:
  1. Go to the main menu and click **Settings**.
  2. Go to the **Repositories** tab and hover over the **Backup Repository** containing the backup.
  3. Click the ellipsis **Manage** button, click **Recover**, and select **Flash boot for VMware** under **VMware Full Recovery**.




- Alternatively, the recovery can be performed by using by using the Search function.

The **New Flash Boot Job Wizard** opens.

- “VMware Flash Boot Job Wizard: Backups” on page 1188
- “VMware Flash Boot Job Wizard: Destination” on page 1191
- “VMware Flash Boot Job Wizard: Schedule” on page 1193
- “VMware Flash Boot Job Wizard: Options” on page 1197

# VMware Flash Boot Job Wizard: Backups

1. On the **Backups** page of the wizard, select a backup using either a **Backup Repositories** or **Jobs & Groups** view in the left pane.

Both federated and standalone (not used as members of federated repositories) backup repositories can be selected.

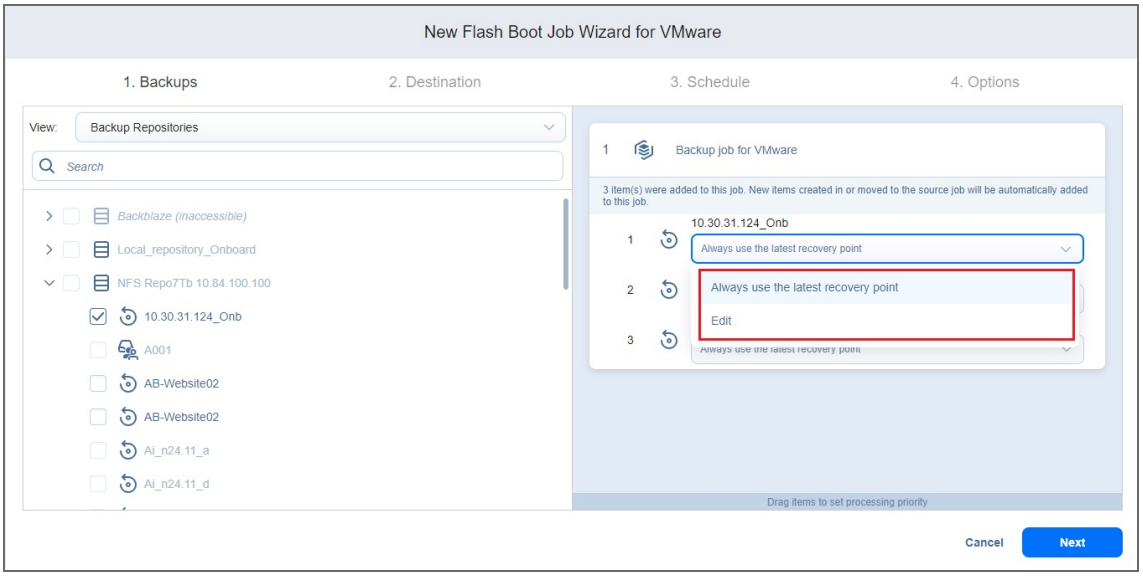
> **Notes**
>
> - You cannot select a federated repository member as a source for recovery jobs.
> - You cannot perform recovery from recovery points located in an inaccessible federated repository member or those that depend on inaccessible recovery points.

Selecting an item in the tree adds the parent item and all children items to the right pane.

> **Notes**
>
> - You can search for a specific backup, job, group, or repository (depending on the selected view) by entering its name into the **Search** field.
> - Selecting a single backup object opens the **Calendar** view of the date picker if this view was the last user-selected view. Selecting a single backup object opens the **Table** view if this was the last view the user selected.
> - You can select a backup object with the corrupted, inaccessible, or pending removal recovery points, but you cannot select the backup objects of the detached repository.

2. In the drop-down list of the selected item on the right pane, select one of the options:

- **Always use the latest recovery point** (displayed by default).
- **Edit** to open the **Recovery points picker** dialog box in the **Calendar** or **Table** view and pick a date.

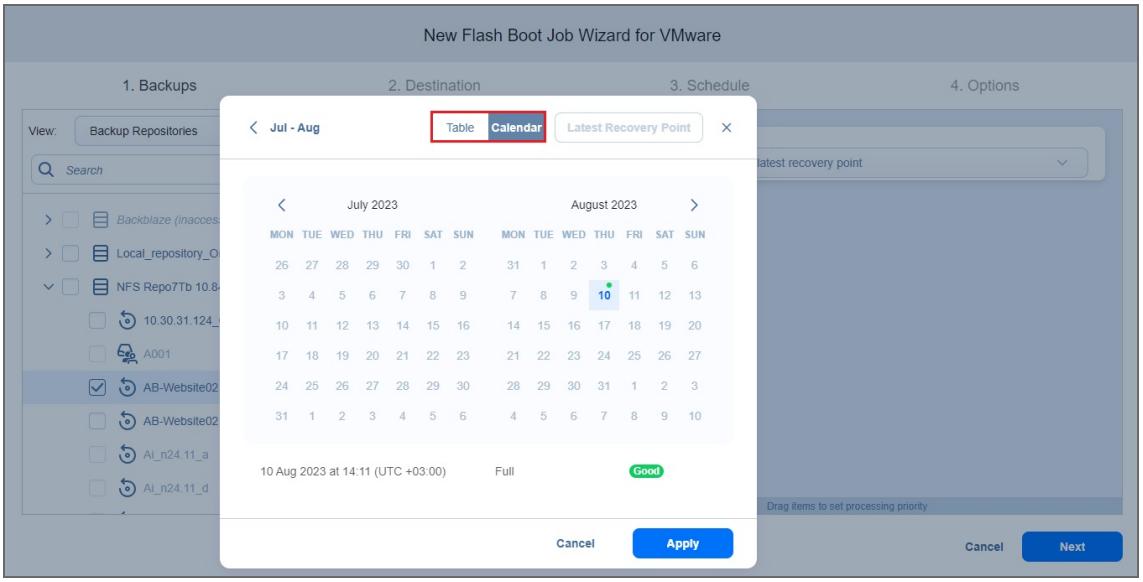


**Notes**

- If a recovery point has been preselected (shown in the dropdown), it is selected by default when the **Calendar** or **Table** view is opened.
- The selected view, either **Calendar** or **Table**, is saved on a per-user basis.

3. If necessary, toggle between the **Calendar** and **Table** views:

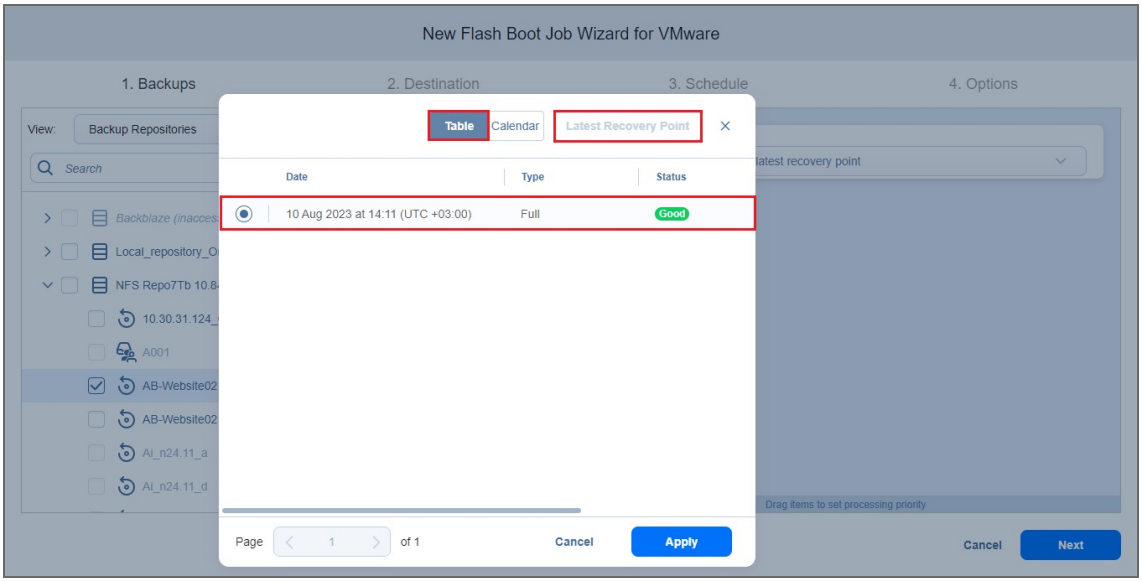
- In the **Calendar** view, select a date or click **Latest Recovery Point** to select the date with the latest recovery point.




**Notes**

- A small dot is displayed on the top of the date if there is at least one recovery point on that date.
- The selected recovery point has a preview displayed under the calendar showing the recovery point status:
  - **green** – verified recovery points
  - **red** – inaccessible/corrupted/pending removal recovery points
  - **dark grey** – unverified recovery points
- The corrupted, inaccessible, or pending removal recovery points cannot be recovered.
- The selected date becomes highlighted.

- In the **Table** view, select a date or click **Latest Recovery Point** to select the date with the latest recovery point.




**Notes**

- Clicking the **Latest Recovery Point** option selects the date with the latest recovery point that is not corrupted, inaccessible, or pending removal.
- The **Latest Recovery Point** option is disabled if:
  - The latest recovery point is currently selected.
  - All the recovery points of the backup object are corrupted, inaccessible, or pending removal.
- If the selected recovery point is encrypted and the corresponding password hash is not available, the **Always use the latest recovery point** or the recovery point timestamp is highlighted in red with a red lock icon. To add such a recovery point to the job, you need to provide the password manually. See Providing Passwords for encrypted recovery points.

To learn about the limitation on the chain of incremental recovery points, refer to the Knowledge Base article.

4. Click **Next** to go to the next page of the wizard.

**Note**

You cannot proceed to the next step if there is at least one selected recovery point to which a password has not been provided.

# VMware Flash Boot Job Wizard: Destination

On the **Destination** page, select a destination for the recovered VMs.

**Note**

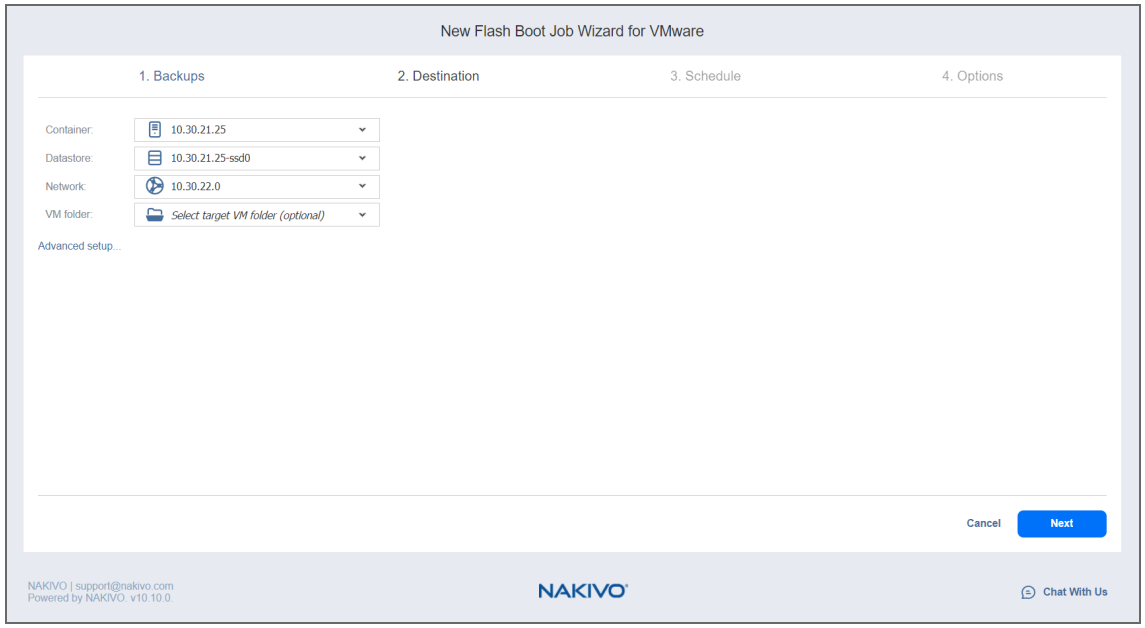
- Due to vSphere storage limitations, only 256 VM disks can be flash-booted per host.

Proceed as described in these sections:

- "Setting the Same Host, Datastore, and Network for all VMs" below
- "Setting Different Options for Recovered VMs" below

## Setting the Same Host, Datastore, and Network for all VMs

To run all VMs on the same host (cluster, or resource pool), datastore and VM folder (optional), and to connect all recovered VMs to the same network, choose a target container, datastore, VM folder and network from the appropriate drop-down lists.

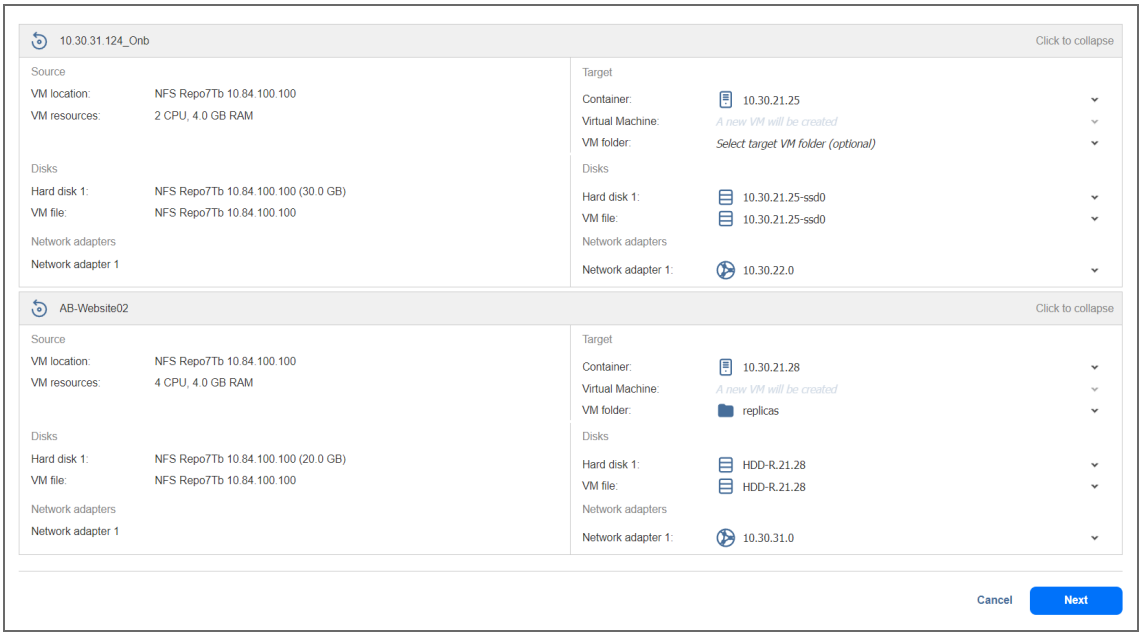



If you choose the **Connect to temporary isolated network** option from the **Network** drop-down list, NAKIVO Backup & Replication will create a new vSwitch and a new network on all hosts where VMs will be recovered. Recovered VMs will be connected to the newly created network.

## Setting Different Options for Recovered VMs

To specify different options for VMs, follow the steps below:

1. Click **Advanced setup**.
2. Choose a target container, target datastore, and target network for each individual VM.




If you choose the **Connect to temporary isolated network** option for a NIC, NAKIVO Backup & Replication will create a new vSwitch and a new network on the host where the VM will be recovered. The recovered VM will be connected to the network.

**Note**

- When performing full recovery from a Proxmox VM agent-based backup to a different hypervisor (VMware) with the Proxmox physical sources configured with *VirtIO SCSI*, make sure to change target to *SATA*.

3. Click **Next**.

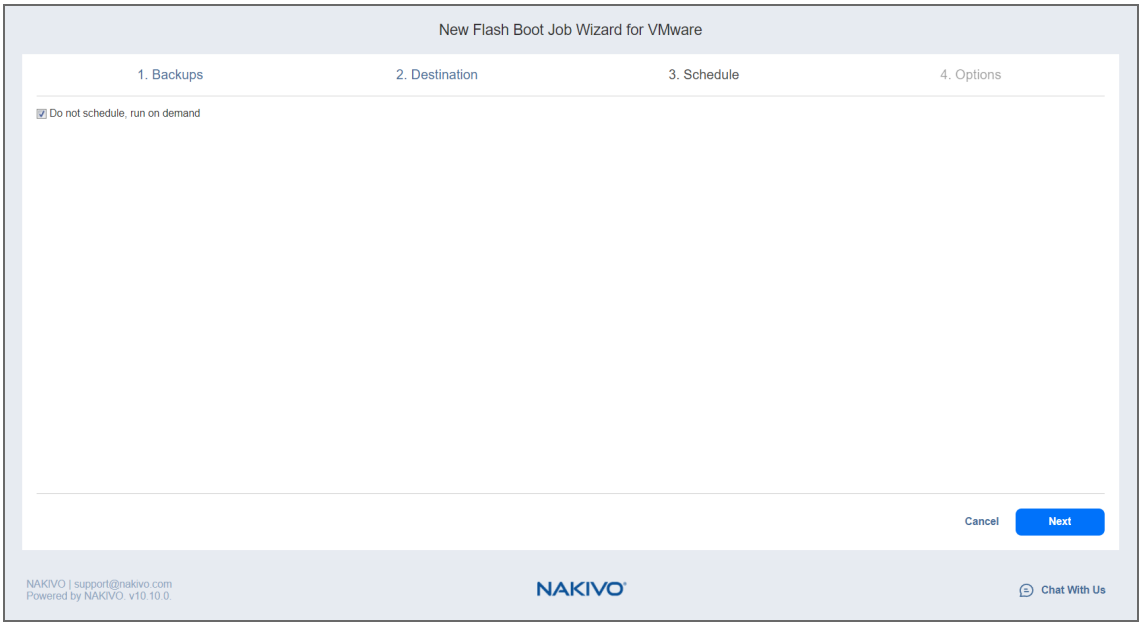
# VMware Flash Boot Job Wizard: Schedule

On the **Schedule** page of the wizard, select to run the recovery job manually or schedule the job to run on a regular basis. Proceed as described in the sections below:


- "Disabling Scheduled Job Execution" below
- "Daily Job Execution" below
- "Monthly or Yearly Job Execution" on the next page
- "Periodic Job Execution" on page 1195
- "Chained Job" on page 1195
- "Additional Schedule" on page 1196


## Disabling Scheduled Job Execution

If you want to start the job manually only (without scheduling), select the **Do not schedule, run on demand** checkbox:

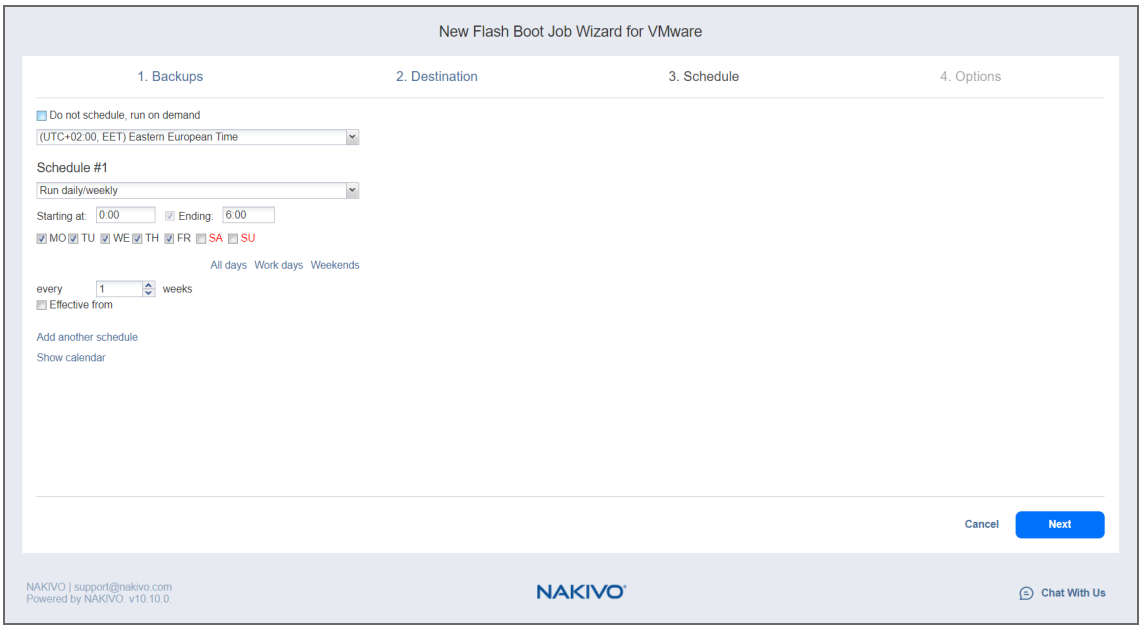



## Daily Job Execution

To run a job once a day, choose **Run daily/weekly** from the schedule drop-down list:

- Choose a time zone to be used for the job start and end times from the time zone drop-down list.
- Choose **Run daily/weekly** from the **Schedule** drop-down list.
- Specify the time when the job should be started in the **Starting at** field.
- Specify the end time for the job in the **Ending** field. If the job has not completed by the time specified, the job will be stopped.

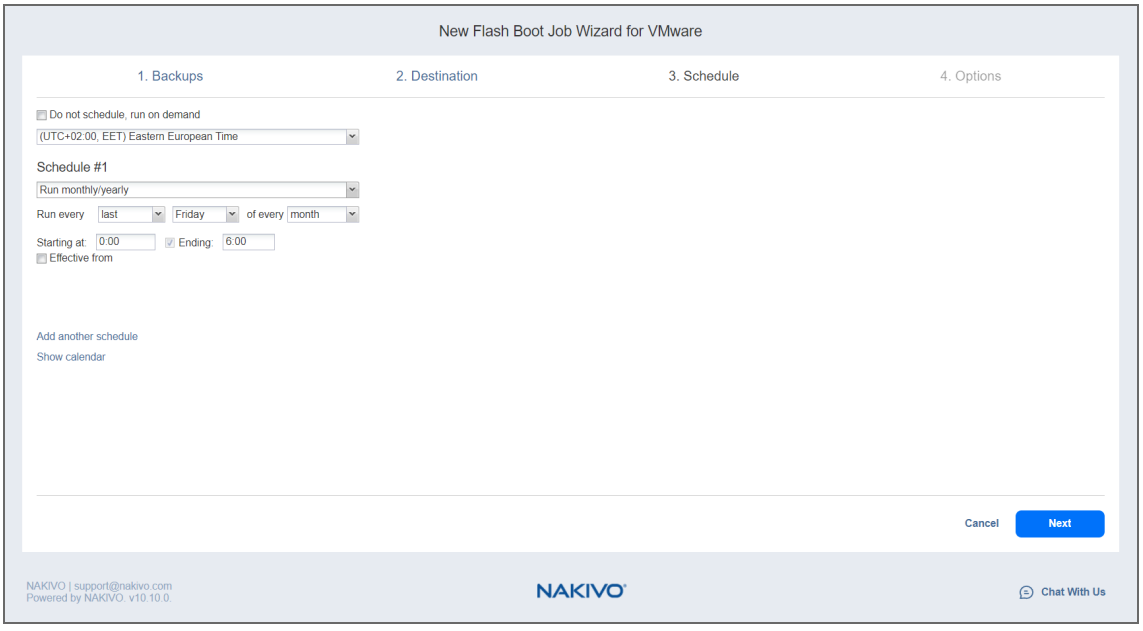
- Select the days of the week on which the job will be started.




# Monthly or Yearly Job Execution

To run a job monthly or yearly, choose **Run monthly/yearly** from the schedule drop-down list:

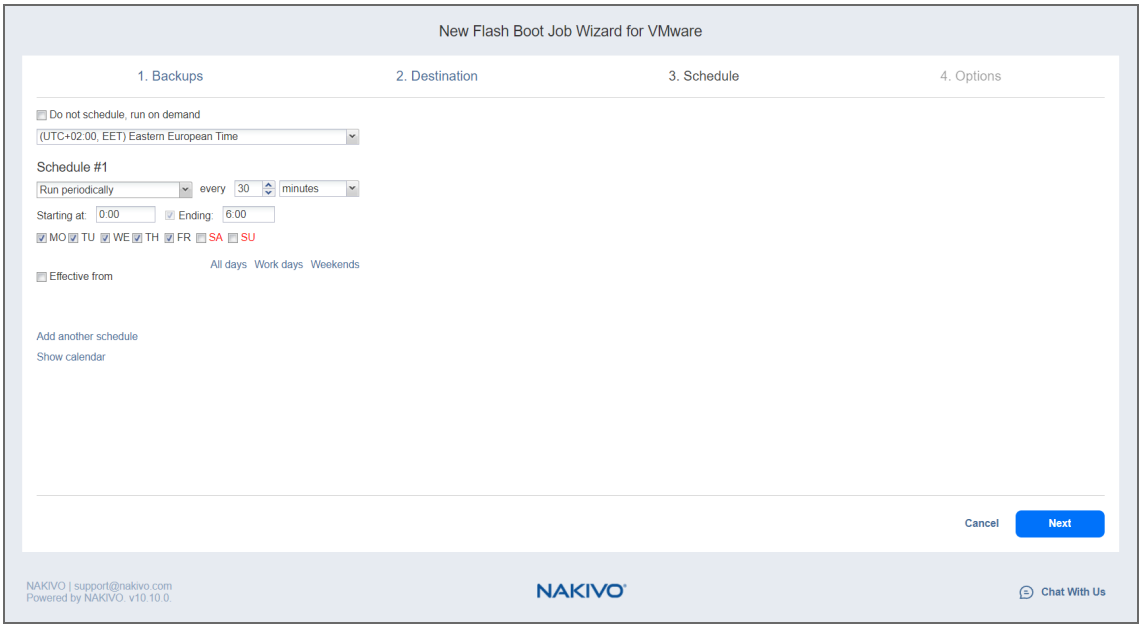
- Specify the job start schedule in the appropriate fields.
- Specify the time when the job should be started in the **Starting at** field.
- Specify the end time for the job in the **Ending** field. If the job has not completed by the time specified, the job will be stopped.
- Choose a time zone to be used for the job start and end times from the time zone drop-down list.
- Select the days of the week on which the job will be started.



## Periodic Job Execution

To run a job multiple times per day, choose **Run periodically** from the schedule drop-down list and then choose a time period from the appropriate fields:

- Specify the time when the job should be started in the **Starting at** field.
- Specify the end time for the job in the **Ending** field. If the job has not completed by the time specified, the job will be stopped.
- Choose a time zone to be used for the job start and end times from the time zone drop-down list.

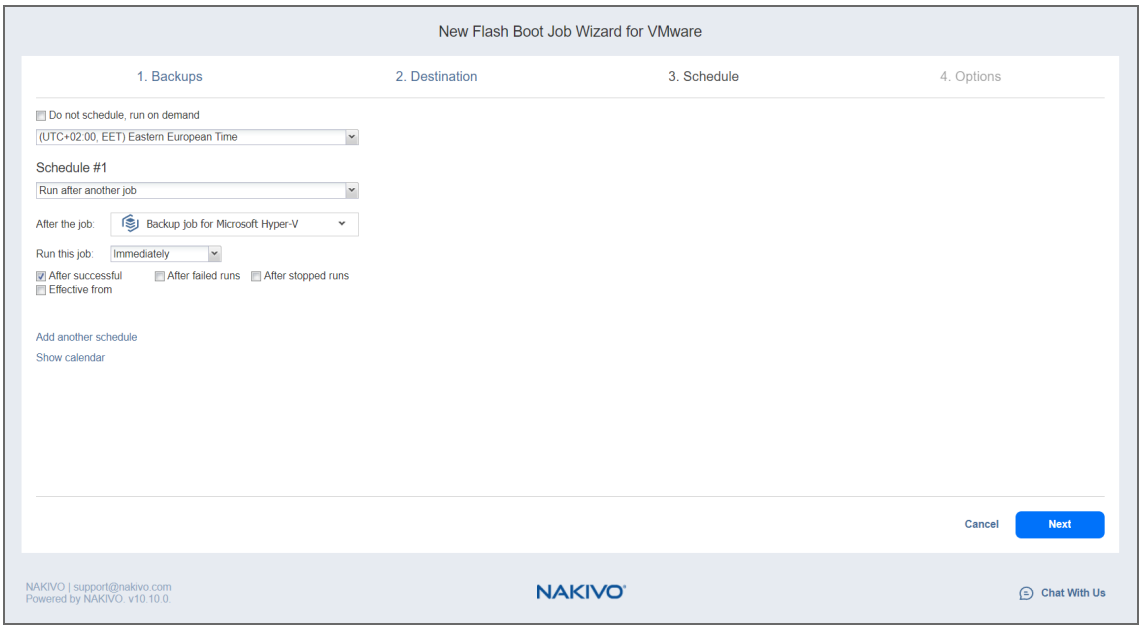



## Chained Job

To run the job after a previous one has completed, choose **Run after another job** from the schedule drop-down list and set the options as follows:

- **After the job**: Select a job after which the current job will be started.
- **Run this job**: Choose whether to run the current job immediately after the previous one has completed or specify a delay.
- **After successful runs**: When selected, the job will run if the previous one has completed successfully.
- **After failed runs**: When selected, the job will run if the previous one has failed.

- **After stopped runs**: When selected, the job will run if the previous one has been stopped.




# Additional Schedule

If you need to add an additional schedule, click **Add another schedule** and configure it as described above.

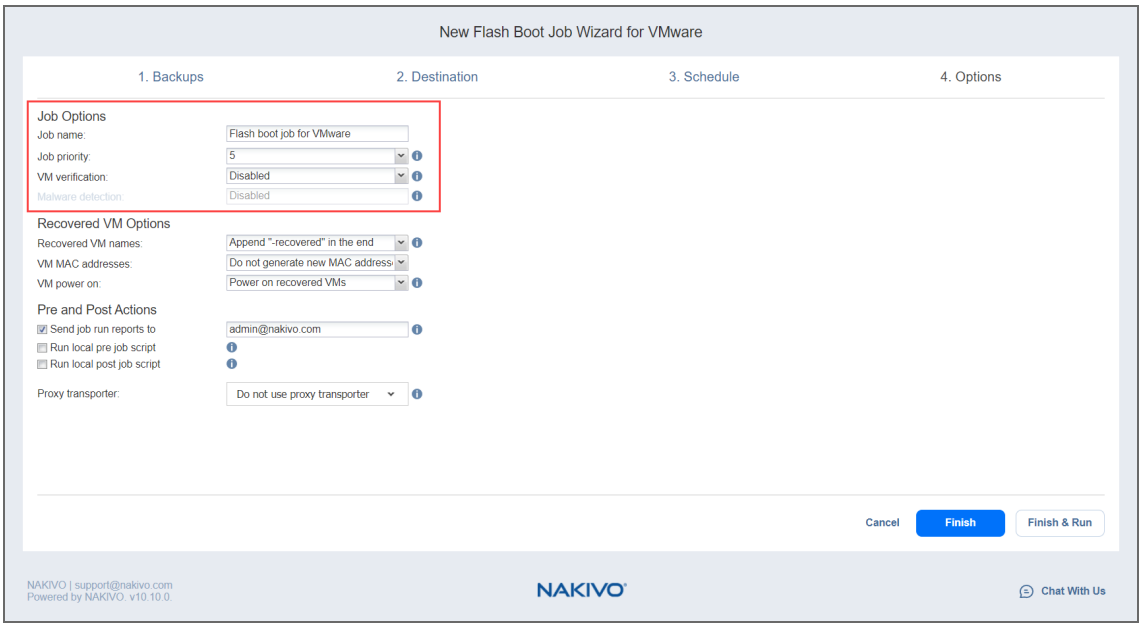
# VMware Flash Boot Job Wizard: Options

On the **Options** page, set up job options as described in these sections:


- "Job Options" below
  - "Job Name" on the next page
  - "Job Priority" on the next page
  - "VM Verification" on the next page
  - Malware Detection
- "Recovered VM Options" on page 1199
  - "Specifying VM Names" on page 1200
  - "Generating VM MAC Addresses" on page 1200
  - "Powering Recovered VMs" on page 1200
- "Pre and Post Actions" on page 1201
  - "Setting up Email Notifications for the Job" on page 1201
  - "Setting Up a Pre-Job Script" on page 1201
  - "Setting Up a Post-Job Script" on page 1202
- "Data Routing" on page 1203
- "Completing the New Flash Boot Job Wizard" on page 1203


## Job Options

In this section, specify a job name and a priority level, select a VM verification type, and enable malware detection.



## Job Name

Enter a name for your job in the **Job name** box.

## Job Priority

Select a job priority level between *1* and *5*, with *1* being the highest priority. Jobs with higher priority levels are prioritized by **Transporters** during job processing.

> **Note**
> This option is only available in the Enterprise, Enterprise Essentials, Enterprise Plus, MSP Enterprise, and MSP Enterprise Plus editions.

## VM Verification

VM Verification allows you to check the integrity of the backup by starting it and interacting with it. For more details, refer to the VM Verification article.

You can choose one of the following VM verification options:

- **Screenshot verification**: When enabled, all VM backups created by the job will be verified as follows: After a backup of a VM is completed, the VM is recovered from the backup using Flash boot (and is disconnected from networks), a screenshot of the recovered VM is taken once the VM OS has booted, after which the VM is discarded. VM screenshots will be included in email notifications (if they are configured) and displayed on the main Dashboard.
- **Boot verification**: When enabled, all VM backups created by the job will be verified as follows: After a VM backup is completed, NAKIVO Backup & Replication recovers the VM using Flash boot, disables networking to prevent network connections, and verifies successful system start.

If you select the **Screenshot verification** option, provide verification options in the **VM Boot Location** dialog box that opens:

- **Verify not more than x VMs simultaneously**: Specify the maximum number of VMs that can be started on the target container simultaneously.
- **Recovery time objective**: Specify an amount of time allocated for verification of each VM backup. If the VM OS does not start within the specified amount of time, verification will be considered failed.
- **Screenshot delay**: Specify an amount of time that the product should wait after a Guest OS start before taking a screenshot.

> **Note**
> The specified time must be sufficient to fully start the VM OS. Try increasing this amount if the default amount is not sufficient.

When **Boot verification** is selected, specify verification options in the dialog box that opens:

- **Verify not more than x VMs simultaneously**: Specify the maximum number of VMs that can be started on the target container simultaneously.
- **Recovery time objective**: Specify the amount of time allocated for verification of each VM backup. If a VM OS does not start within the specified amount of time, verification will be considered failed.

### Malware Detection

With this option enabled, the backups are scanned for malware using the configured antivirus software on the scan server. Optionally, if you have selected **Enabled** for the **Malware detection** option, click the **settings** link to configure the following options:

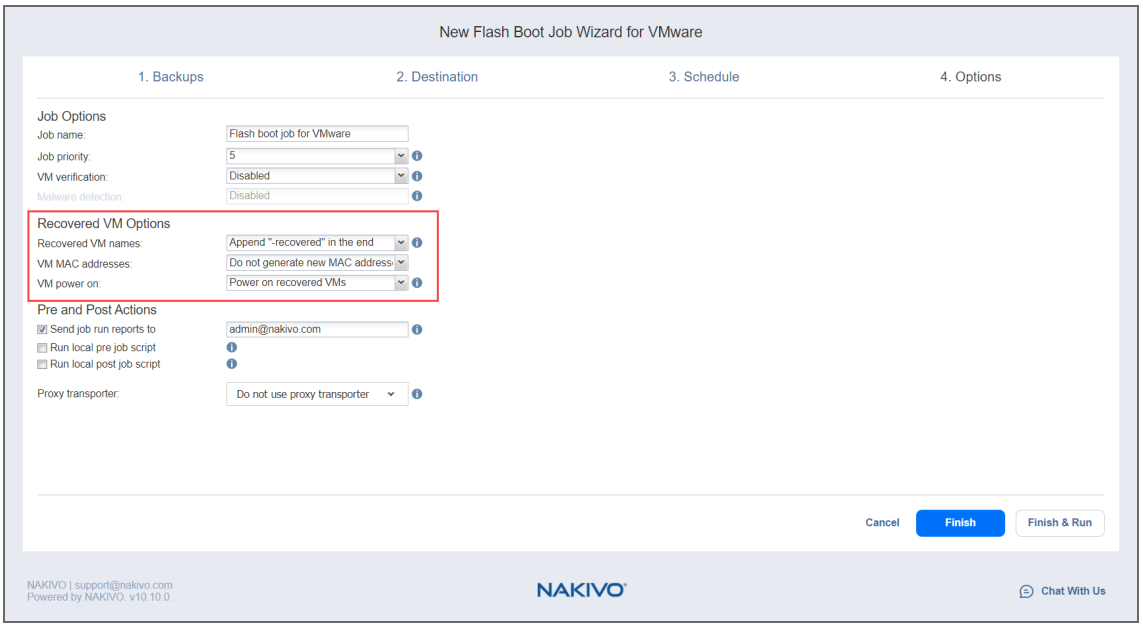
- **Scan server**: Select a specific scan server for the job or leave the **Default** setting. If **Default** is selected, the Transporter is used as the scan server and can support a maximum of 2 concurrent scan tasks.

  **Notes**
  - For the **Default** option, if the **Repository Transporter** is the installed **Transporter**, it require the master password to function as the scan server.
  - For more details on the requirements for Scan Server, refer to the Feature Requirements.

- **Scan type**: Choose between the **Deep scan** and the **Quick scan**:
  - **Deep scan**: When this option is selected, the antivirus software scans the entire backup and may take longer to complete.
  - **Quick scan**: When this option is selected, the antivirus software scans only OS disks in the backup.
- **If malware is detected**: Choose the behavior if malware is detected:
  - **Fail the recovery job**: With this option, the recovery process fails in case the job has only one VM. If the job has several VMs, the infected VMs are skipped and the job continues to run.
  - **Continue and recover to isolated network**: When this option is selected, the recovery job completes the scanning process and recovers the infected VMs to a temporary isolated network.
- **Scan timeout**: Specify the timeout for the malware detection process. If the specified amount of time is exceeded, the recovery job fails.
- Click **Apply** when you're done.

## Recovered VM Options

In this section, specify VM names, generate VM MAC addresses, and choose whether you want to power on recovered VMs or not.




## Specifying VM Names

NAKIVO Backup & Replication allows you to change the names of recovered VMs so you can distinguish between recovered VMs and source VMs. By default, the text “- recovered” is appended to the end of the recovered VM name.

To change VM replica names, choose one of the following **Recovered VM names** options in the Recovered VM Options section:

- **Append “-recovered” in the end**: Source VM names are used for recovered VM names and “-recovered” is added after the recovered VM name.
- **Leave recovered VM names as is**: Recovered VM names are identical to the source VM names.
- **Enter custom recovered VM names**: You can enter custom names for recovered VMs.

## Generating VM MAC Addresses

In the **Recovered VM Options** section, you can choose one of the following options in relation to recovered VM MAC addresses:

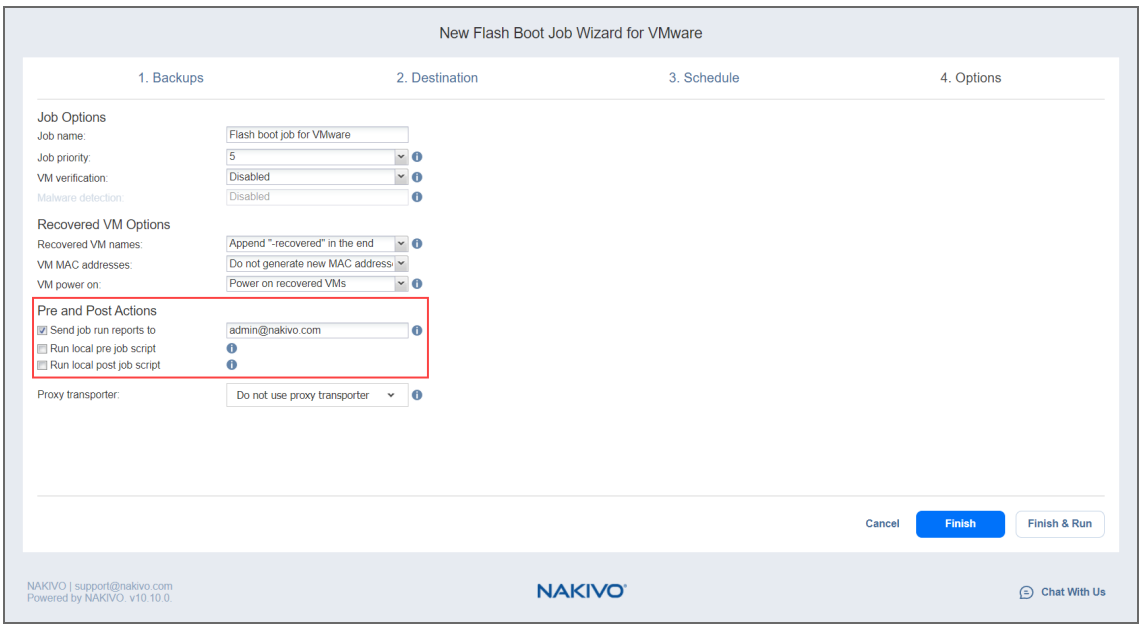
- **Do not generate new MAC addresses**: The recovered VMs will have the same MAC address as the source VMs
- **Generate new MAC addresses**: A new MAC address will be generated for each recovered VM.

## Powering Recovered VMs

To power on the recovered VMs, choose the **VM power on** option.

# Pre and Post Actions

NAKIVO Backup & Replication allows you to run a script before Flash boot begins (a pre-job script) and after the boot of all VMs in the job has completed (a post-job script). The scripts can only be executed on the machine where the **Director** is installed. You can also set up email notifications for the job.




## Setting up Email Notifications for the Job

NAKIVO Backup & Replication can send email notifications about job completion status to specified recipients. This feature complements global notifications and allows you to configure notifications on a per-job level.

To enable this option, make sure Email settings are configured.

To send email notifications, do the following:

1. In the **Pre and Post Actions** section, select the **Send job run reports to** option.
2. In the text box that becomes enabled, specify one or more email addresses. Use semicolons to separate multiple email addresses.

## Setting Up a Pre-Job Script

To run a script before the product begins recovering VMs, do the following:

1. Place a script file on the machine where the **Director** is installed.
2. In the **Pre and Post Actions** section, select the **Run local pre job script** option and click **settings**.
    - **Script path**: Specify a local path to the script on the machine where the **Director** is installed. A script interpreter should be specified.
      Example (Windows): *cmd.exe /c D:\script.bat*

Example (Linux)*: bash /root/script.sh*

- **Job behavior**: Choose one of the following job behaviors in relation to script completion:
  - **Wait for the script to finish**:With this option selected, the job will be in the "running" state until the script is completed.
  - **Do not wait for the script to finish**: With this option selected, the job will be completed even if the script execution is still in progress.
- **Error handling**: Choose one of the following job behaviors in relation to script failure:
  - **Continue the job on script failure**: With this option selected, script failure will not influence the status of the job.
  - **Fail the job on script failure**: With this option selected, if the script fails, the job status will be set to "failed" even if VM backup has been successful.

3. Specify the following parameters in the dialog that opens:

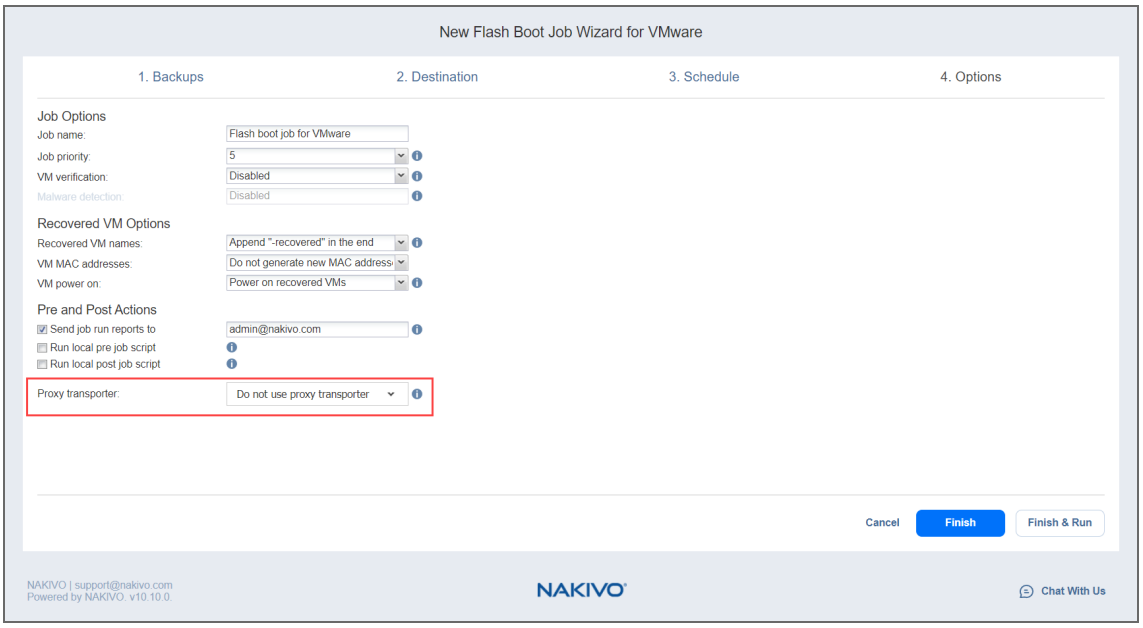
## Setting Up a Post-Job Script

To run a script after the product has finished recovering all VMs, do the following:

1. Place a script file on the machine where the **Director** is installed.
2. In the **Pre and Post Actions** section, select the **Run local post job script** option and click **settings**.
   - **Script path**: Specify a local path to the script on the machine where the **Director** is installed. A script interpreter should be specified.
     Example (Windows): *cmd.exe /c D:\script.bat*
     Example (Linux): *bash /root/script.sh*
   - **Job behavior**: Choose one of the following job behaviors in relation to script completion:
     - **Wait for the script to finish**: With this option selected, the job will be in the "running" state until the script is completed.
     - **Do not wait for the script to finish**: With this option selected, the job will be completed even if the script execution is still in progress.
   - **Error handling**: Choose one of the following job behaviors in relation to script failure:
     - **Continue the job on script failure**: With this option selected, script failure will not influence the status of the job.
     - **Fail the job on script failure**: With this option selected, if the script fails, the job status will be set to "failed" even if VM backup has been successful.
3. Specify the following parameters in the dialog that opens.

## Data Routing

If the **Transporter** assigned to the **Backup Repository** cannot use iSCSI port *3260* because it is occupied by other services, you can set data routing: a proxy transporter can be used to forward iSCSI target exposed from the **Backup Repository** to the target host. To set data routing, go to the *Data routing* section and choose a proxy transporter from the list of available **Transporters**.




## Completing the New Flash Boot Job Wizard

Click **Finish** or **Finish & Run** to complete the job creation.

> **Note**
>
> If you click **Finish & Run**, you will have to define the scope of your job. Please refer to "Running Jobs on Demand" on page 364 for details.

# Migrating Recovered VMs Using Flash Boot

Using Flash boot, you can migrate the recovered VMs to another location. To do this, follow the instructions below:

## Migrating Recovered VMware VMs Using Flash boot

To migrate a booted VM from one datastore to another, do the following while the virtual machine is running:

1. Open vSphere Client.
2. In the vSphere Client inventory, select the virtual machine recovered with Flash boot that you want to migrate.
3. Right-click on the VM and select Migrate.
4. Select **Change datastore** and click **Next**.
5. Select the specific format of your virtual disk.

> **Note**
> When selecting virtual disk format, do not leave the **Same format as source** option selected as the VM will be discarded when clicking the **Discard VMs**button in the Flash boot job.

6. Select the datastore where you want to store the VM.
7. Review the page and click **Finish**.
8. After migration, verify that the VM retains its original state (running or powered off).

> **Note**
> After migrating the booted VMs, you can click the **Discard VMs** button in the Flash boot job. The job detects that the VMs were migrated and does not discard them.

Replicate booted VMs as described below:

1. Refresh the Inventory in Settings (so that the VMs created by Flash boot are added to the product).
2. Create a new replication job for the VMs that you want to permanently recover.
3. Run the replication job.
4. After replication is finished, run the job once again. This is required because the recovered VMs were running (and obtaining some changes) while the initial replication was in progress.
5. To ensure zero data loss and maintain the VM state, power off the VMs created by Flash Boot in the VMware vSphere client and run the VM replication job one last time.
6. Discard the VMs created by Flash boot.

7. Power on the VMs created by the replication job.
8. After replication, check that the resulting VM retains the same state (running or powered off).

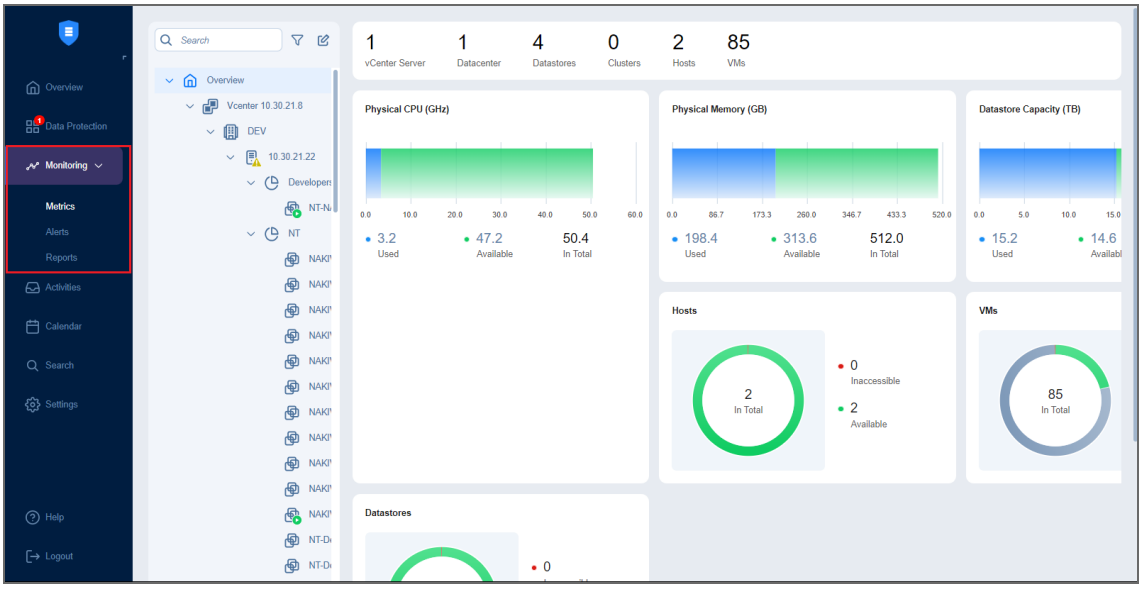
# Monitoring

Monitoring in NAKIVO Backup & Replication allows you to track VMware vSphere infrastructure resources used by hosts, datastores, and VMs. You can monitor the key metrics for CPU, memory, and disk usage in real time and over a specified period of time to make sure the infrastructure has enough resources to meet organizational demands.

**Monitoring** also allows you to configure alerts and automated reports to track CPU usage, memory usage, and datastore performance. These reports help you detect unusual activity that can signal ransomware or other malware in your VMware vSphere infrastructure.

Refer to the following topics to learn how to use **Monitoring**:

- “Metrics” below
- “Alerts” on page 1211
- “Reports” on page 1220




# Metrics

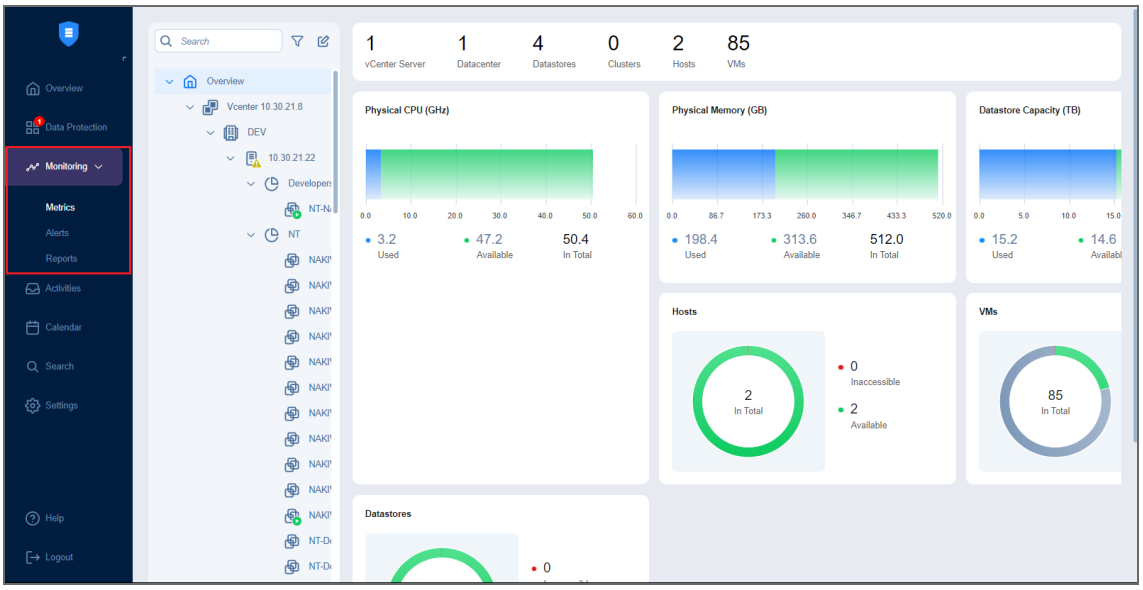
The **Metrics** dashboard displays the key indicators for monitored items. When you go to **Monitoring**, the **Metrics** dashboard is displayed by default. Refer to the following topics to learn how to use the **Metrics** dashboard:

- “Metrics” above
- “Managing Monitored Items” on page 1208

## Metrics Overview

In the left pane of the **Metrics** dashboard, you can find the following:

- **Search** box: Enter the partial or full name of an item in the search box to find it in your inventory items.
- **Filter** button: Click the button to filter the items displayed by type. You can filter the displayed items by:
    - Hosts
    - Datastores
    - VMs
- **Edit** button: Manually select items from your Inventory to monitor or create a policy rule for items to be added automatically.
- **Overview**: Monitored items are displayed in tree view when you expand **Overview**. Select a monitored item to view the corresponding metrics.




In the right pane of the **Metrics** dashboard, you can see the number of issues and the corresponding performance bar and line charts for the monitored item.

The following information is displayed for each type of selected item:

- VMs
    - **CPU Load**: Shows the CPU load in percentage in real time.
    - **Memory Load**: Shows the memory usage in percentage in real time.
    - **Disk Usage**: Shows the disk usage in gigabytes in real time.
    - **CPU Usage (%)**: Shows the CPU load in percentage over a specified period of time.
    - **Memory Usage (GB)**: Shows the memory usage in gigabytes over a specified period of time.
    - **Disk Usage (GB)**: Shows the disk usage in gigabytes over a specified period of time.
- Hosts
    - **CPU Load**: Shows the CPU load in percentage in real time.
    - **Memory Load**: Shows the memory usage in gigabytes in real time.

    - **CPU Usage (%)**: Shows the CPU load in percentage over a specified period of time.
    - **Memory Usage (GB)**: Shows the memory usage in gigabytes over a specified period of time.
- Datastores
    - **Disk Usage**: Shows the disk usage in terabytes in real time.
    - **Disk Usage (TB)**: Shows the disk usage in terabytes over a specified period of time.

# Managing Monitored Items

In this section, you can find out how to manage the items monitored in NAKIVO Backup & Replication, add items to the monitoring dashboard, specify time frames for charts, and understand data collection and data aggregation.


- "Adding Items to Monitoring Dashboard" below
- "Setting Up Chart Time Frames" on page 1210
- "Data Collection" on page 1211
- "Data Aggregation" on page 1211


## Adding Items to Monitoring Dashboard

The items added to the **Monitoring** dashboard depend on the NAKIVO Backup & Replication license.

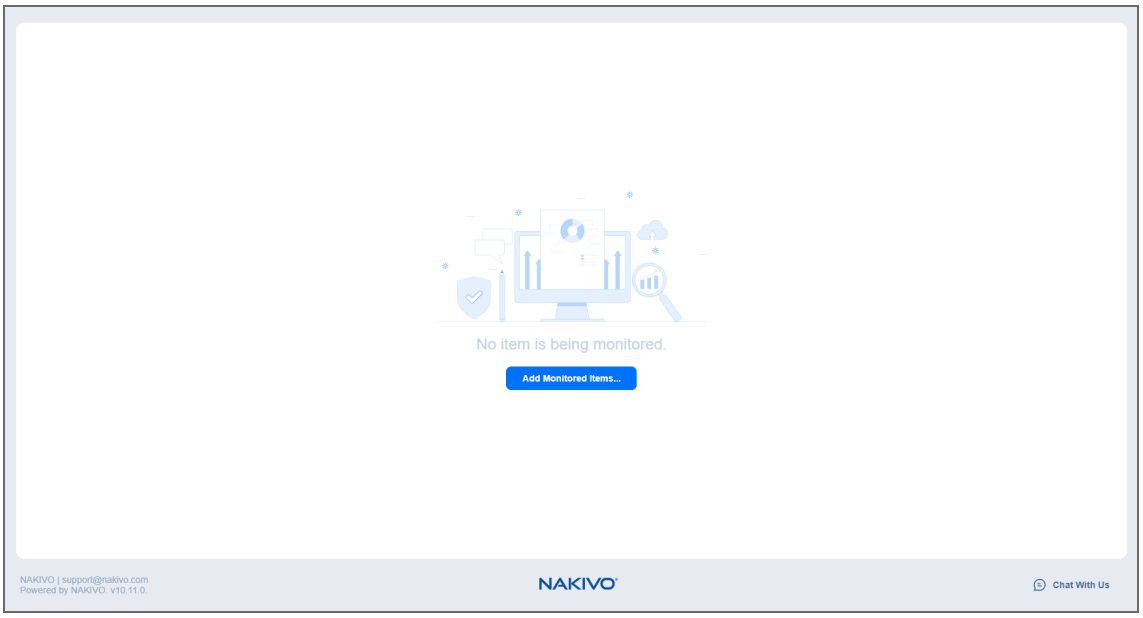
**Note**

To use **Monitoring**, the following conditions must be met:

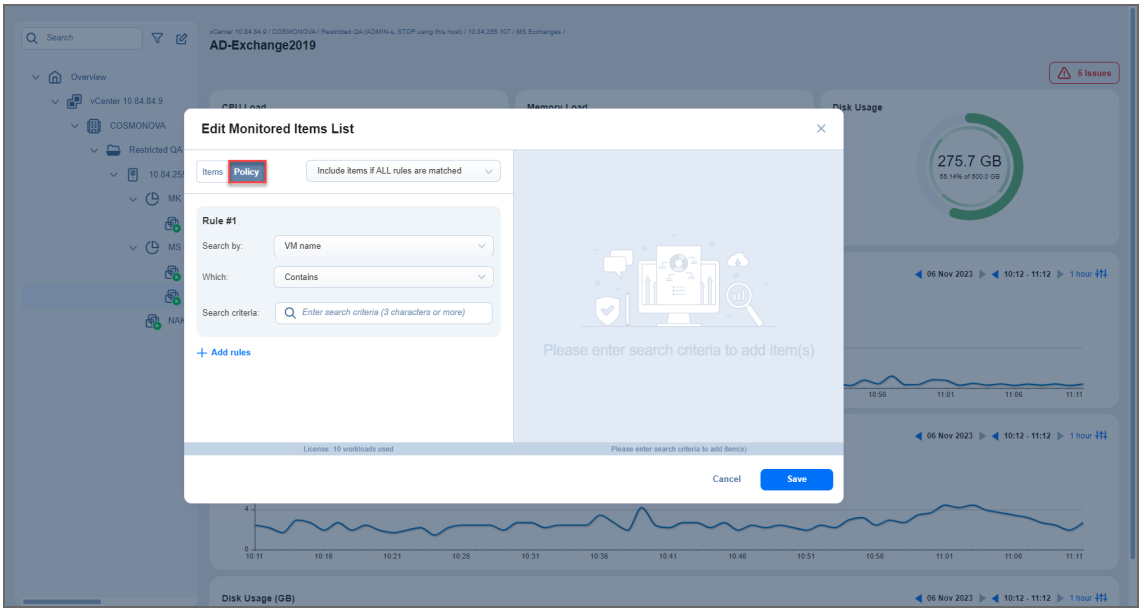
- You must have a license for **Monitoring**. When using a non-trial license, you can add supported **Inventory** items to be monitored and edit those items manually.
- Items of the supported platform must be added to the **Inventory**.
- If a Free Trial license is installed, the product automatically adds top-level VMware vSphere containers from the **Inventory** to the list of monitored items when an **Inventory** item is initially discovered.

When first accessing the dashboard, you see the **No item is being monitored** message.

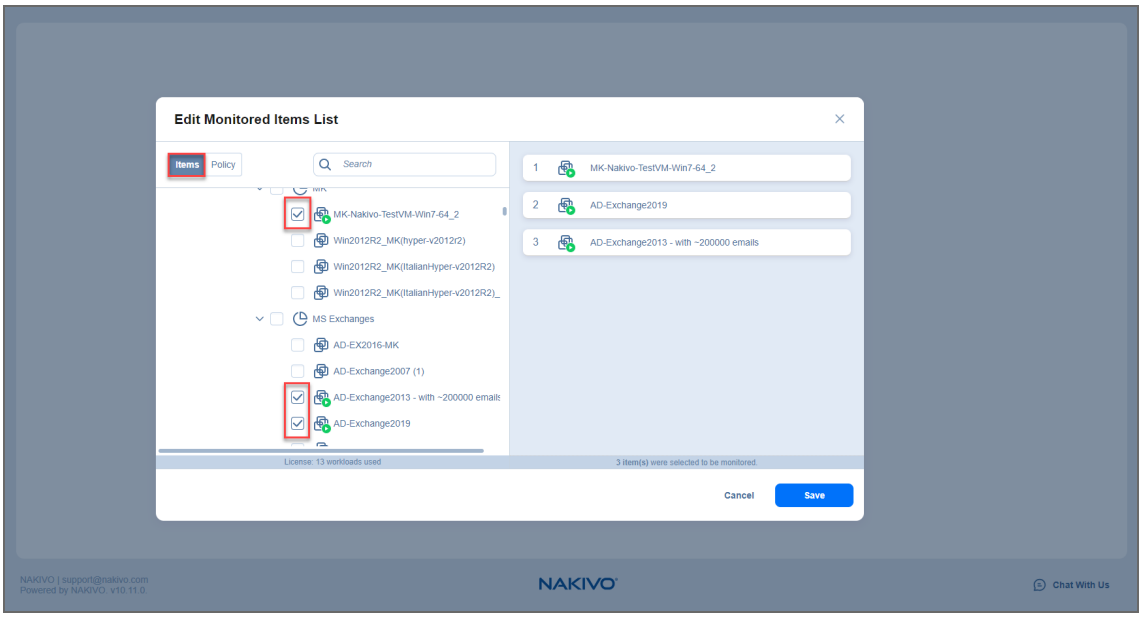
1. Click **Add Monitored Items...** to open the **Edit Monitored Items List** dialog window:




2. On the left side, choose the items you want to add for monitoring. Alternatively, you can set up policy rules to automatically add inventory items for monitoring.

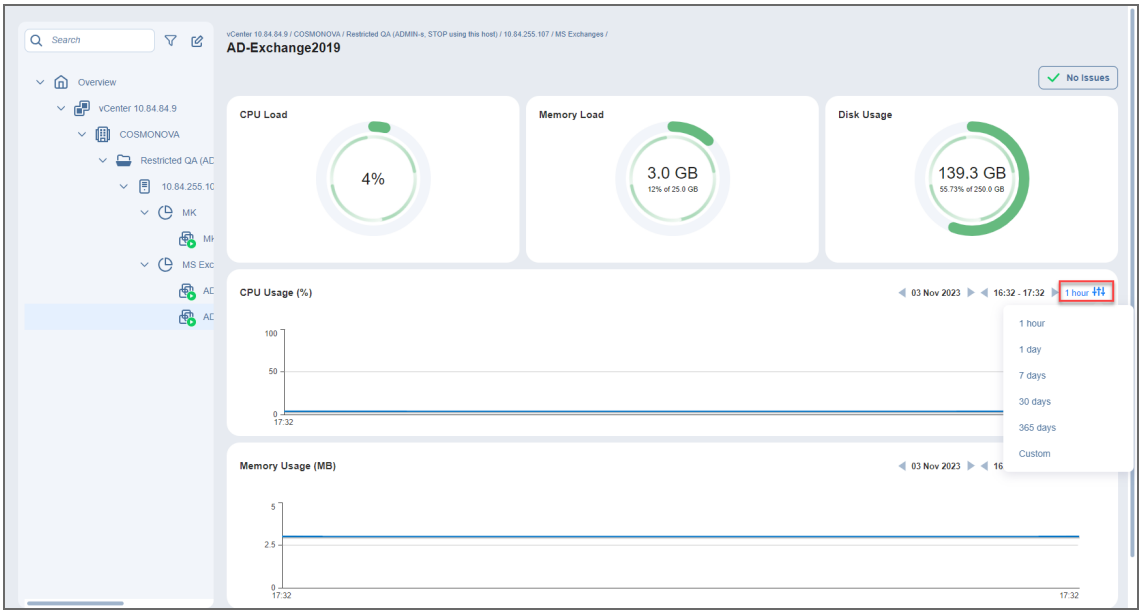



3. Click **Save** when you are done.




# Setting Up Chart Time Frames

You can specify time frames for the displayed historical CPU, memory, and disk usage line charts. To do so, click on the controls icon on the right for the needed chart and choose one of the following options:




**1 hour** – to display data for a 1-hour time frame. The hour can then be specified.

**1 day** – to display data for a 1-day time frame. The day can then be specified.

**7 days** – to display data for a 7-day time frame. The 7-day period can then be specified.

**30 days** – to display data for a 30-day time frame. The 30-day period can then be specified.

**365 days** – to display data for a 365-day time frame. The 365-day period can then be specified.

**Custom** – when you select this option, the **Calendar** popup opens, and you can specify a custom time range for the chart. Click **Apply** when you are done.

## Data Collection

NAKIVO Backup & Replication collects and updates data about monitored items every minute.

For VM CPU load, VM memory load, host CPU load, and host memory load, the following exceptions apply:

- If the product cannot collect data every 1 minute for 5 minutes, data collection switches to a 5-minute interval. This mode collects 5 data points at once.
- If the product cannot collect data every 5 minutes for 10 minutes, data collection switches to a 10-minute interval. This mode collects 10 data points at once.

For VM disk usage, VM memory usage, and datastore disk usage, if the product cannot collect data every 1 minute for 5 minutes, multiple real-time data points are collected but only the latest single data point within the 5-minute interval is retained and displayed.

> **Note**
> It may take up to 3-5 minutes for the widgets of the **Overview** tab to finish collecting data and display the latest information.

## Data Aggregation

NAKIVO Backup & Replication creates a data point every minute. As the archive grows, the data points are aggregated as follows:

- 1-minute data points are aggregated into 10-minute data points every 10 minutes and calculated as the average value of all 1-minute points. 1-minute data points that are older than 1 day are automatically removed.
- 10-minute data points are aggregated into 1-hour data points and calculated as the average value of all 10-minute points. 10-minute data points that are older than 1 week are automatically removed.
- After a year, 1-hour data points are removed.

# Alerts

With **Monitoring** in NAKIVO Backup & Replication, you can set up alerts to be triggered for hosts, VMs, and datastores when configured monitoring rules are met. You can create custom alerts using alert templates to help you detect malicious changes in the infrastructure.

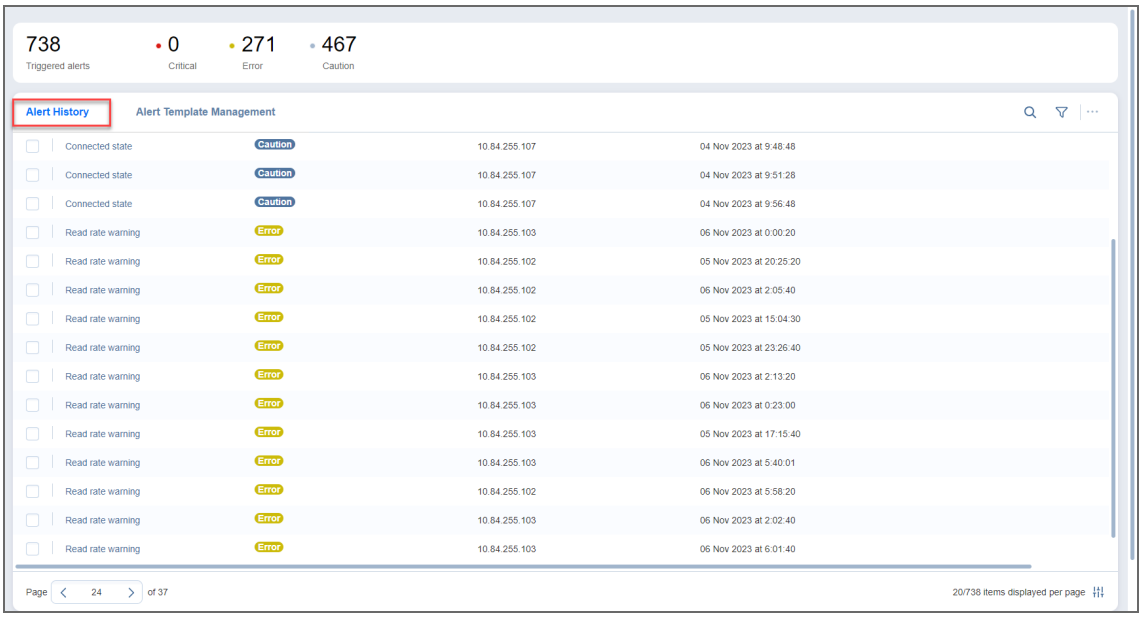
The **Alerts** dashboard displays the list of all alerts triggered when the rules for enabled alert templates are met and allows you to create or edit any alert templates. A summary bar is displayed at the top of the **Alerts** dashboard with the total number of triggered alerts and alert templates. The summary bar displays the total number of triggered alerts categorized by three severity levels: *critical*, *error*, and *caution*.

Refer to the following topics to learn how to view triggered alerts and configure alert templates:

- "Alert History" below
- "Alert Template Management" on the next page


# Alert History

The **Alert History** tab displays the list of triggered alerts. When the rules are met in an enabled alert template, the configured alert is triggered and displayed in the list.




In the **Alert History** tab, you can do the following:

- View **Alert History**: View the triggered alerts displayed in the alert history table.
- Search alerts by entering the partial or full name of an alert, source, or description.
- Filter the alerts by **Severity** (*Critical*, *Error*, *Caution*) or **Triggered Time**.
- Dismiss alerts by selecting an alert, clicking the **...** (ellipsis) button, and then selecting **Dismiss** in the dialog window.
- View alert settings by right-clicking the triggered alert to open the dialog window and selecting **View Settings**.

The **Alert History** table displays all triggered alerts with the following columns:

- **Alert name**: The name of the alert.
- **Severity**: The severity of the alert, which can be *Caution*, *Error*, or *Critical*.
- **Source**: The host, VM, or datastore that the alert was triggered for.
- **Triggered time**: The date and time when the alert was triggered.
- **Description**: The description of the alert as defined in the alert template. For details, refer to Creating Alert Templates.

To customize the sorting of the table, click the head of the column you wish to sort by.

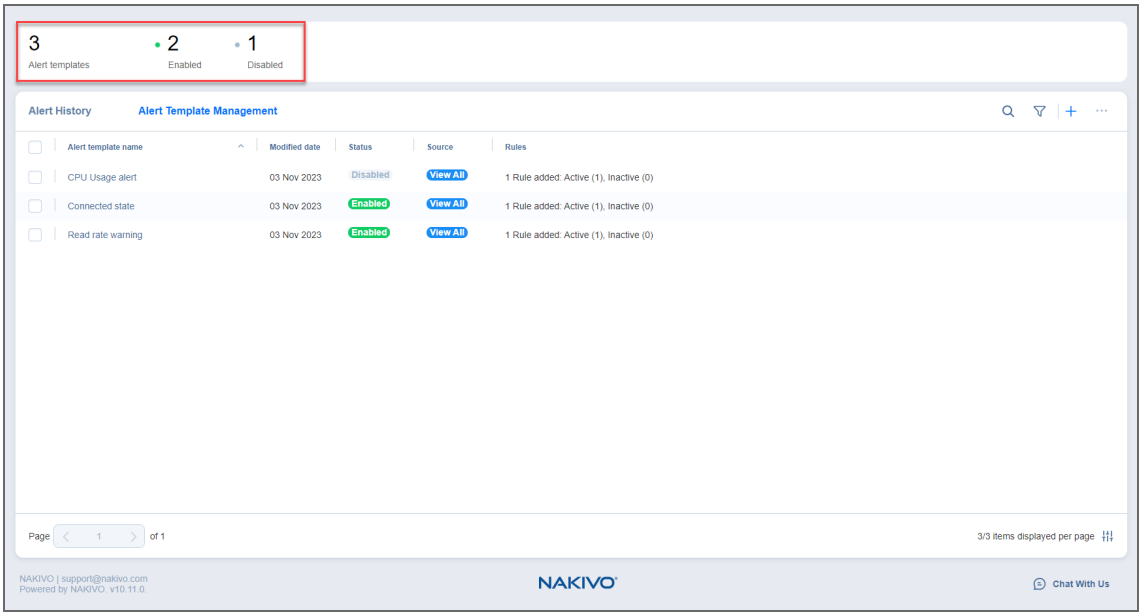
**Note**

The **Alert History** table can display up to 1,000 triggered alert records. After reaching this limit, the oldest alert is replaced automatically with the latest alert.

# Alert Template Management

The **Alert Template Management** tab displays all configured alert templates. Alert templates are composed of specific rules that apply to predefined monitored items. When these rules are met in an enabled alert template, an alert is triggered. Disabled alert templates do not trigger alerts. You can configure custom alert templates for specific metrics and thresholds.

The summary bar displays the total number of alert templates and the number of enabled and disabled templates.




In the **Alert Template Management** tab, you can do the following:

- View alert templates: View all enabled and disabled alert templates. You can view details by clicking on an alert template. For details, refer to "Viewing Alert Template Details" on page 1215.
- Search alert templates by entering the partial or full name of the alert.
- Filter alerts by **Status** (*enabled*/*disabled*), **Severity** (*Critical*, *Error*, *Caution*), or **Modified Date**.
- Add a new alert template by clicking the **+** (add) button and completing the **Create New Alert Template** wizard. For details, refer to "Creating Alert Templates" on page 1216.
- Edit or delete the alert template by clicking the **...** (ellipsis) button.

The **Alert Template Management** table displays all configured alert templates with the following columns:

- **Alert template name**: The name of the alert template.
- **Modified date**: The date when the alert template was last modified.

- **Status**: The status of the alert template. Click the **Enabled** button to change the status to **Disabled** and vice versa.
- **Severity**: The severity of the alert when triggered. The alert severity can be *Caution*, *Error*, or *Critical*.
- **Source**: Click **View All** to open the **All Sources** dialog window that displays a list of the VMs, datastores, and hosts selected for the alert template.
- **Rules**: Total number of active and inactive rules in the alert template.

To sort the table by **Name**, **Modified date**, **Status**, or **Severity**, click the head of the column you want to sort by.

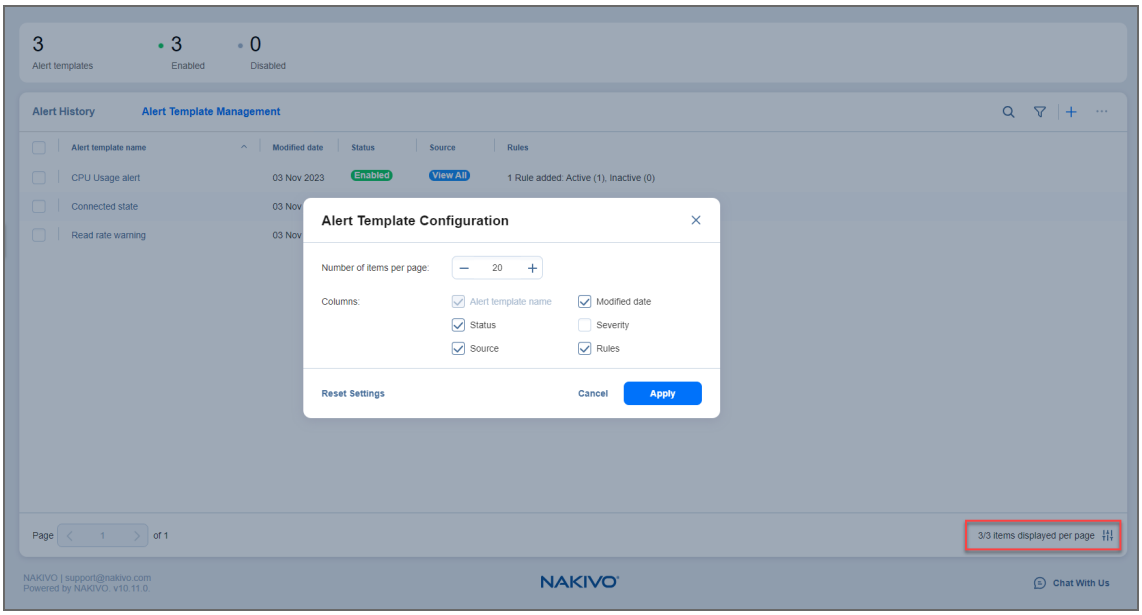
> **Note**
> The **Alert Template Management** table can display up to 100 records for configured alert templates. When this limit is reached, the **+** (add) button will be disabled.

## Configuring Table Columns

You can show/hide columns or modify the number of items per page in the **Alert Template Management** table. In the lower right corner, click the controls icon. In the dialog window that opens, select/deselect checkboxes and click **Apply**. The following options are available:

- Alert template name
- Status
- Source
- Modified date
- Severity
- Rules

> **Note**
> You cannot modify the **Alert template name** column in this dialog window. Refer to "Editing Alert Template" on page 1219 to learn more about editing alert templates.



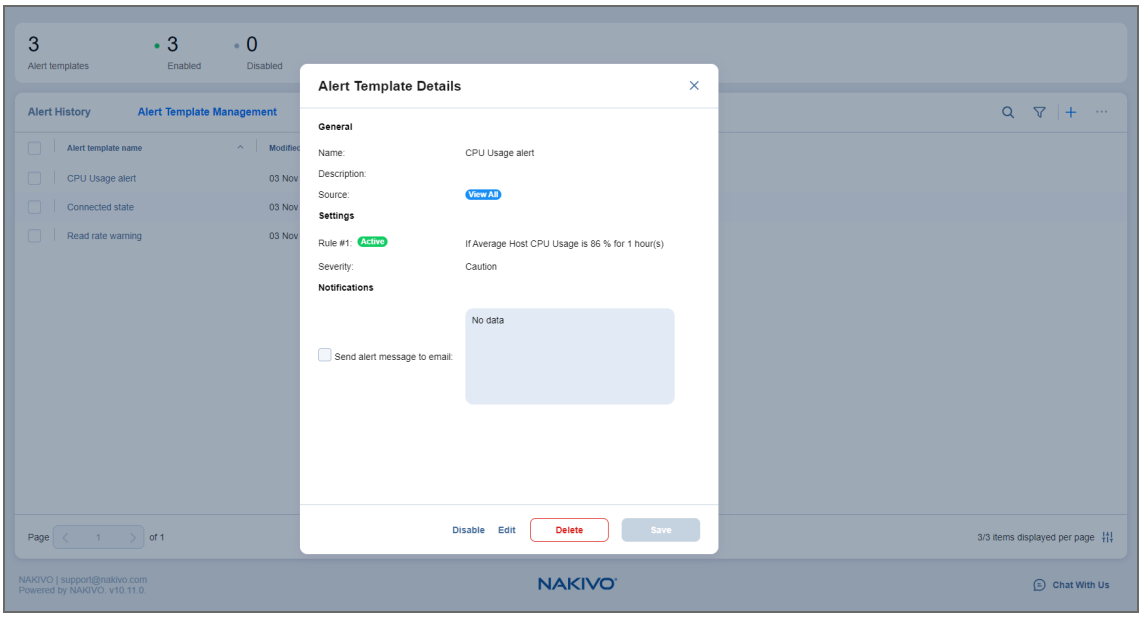

Click **Cancel** to exit or **Reset Settings** to reset to default settings.

## Viewing Alert Template Details

Click an alert template to view the details that were configured in the "Creating Alert Templates" on the next page wizard, including:

- **General** section: Displays the name and description of the alert. Click **View All** to display all source VMs, hosts, and datastores that were selected for this alert.
- **Settings** section: Lists the defined rules and severity of the alert. Click the **Active** button to change the status of a rule to *Inactive* and vice versa. Alerts will be triggered according to the active rules.
- **Notifications** section: Displays all the email addresses that will receive a notification when the alert is triggered.
- **Enable/Disable** button: Click to enable/disable the alert template.
- **Edit**: Click to open the **Edit Alert Template Wizard**. For details, refer to "Editing Alert Template" on page 1219.
- **Delete**: Click to delete the alert template.

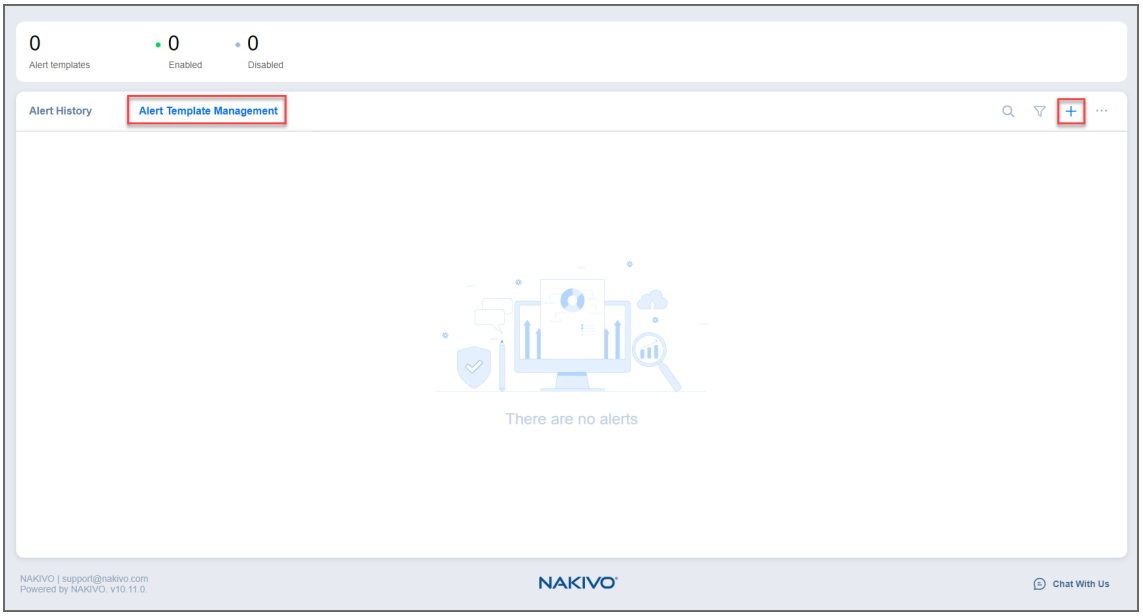
- **Save**: Click to save changes.




# Creating Alert Templates

To create an alert template:

1. Go to **Monitoring > Alerts** and click the **Alert Template Management** tab.
2. Click the **+** (add) button. The **Create New Alert Template** wizard opens.

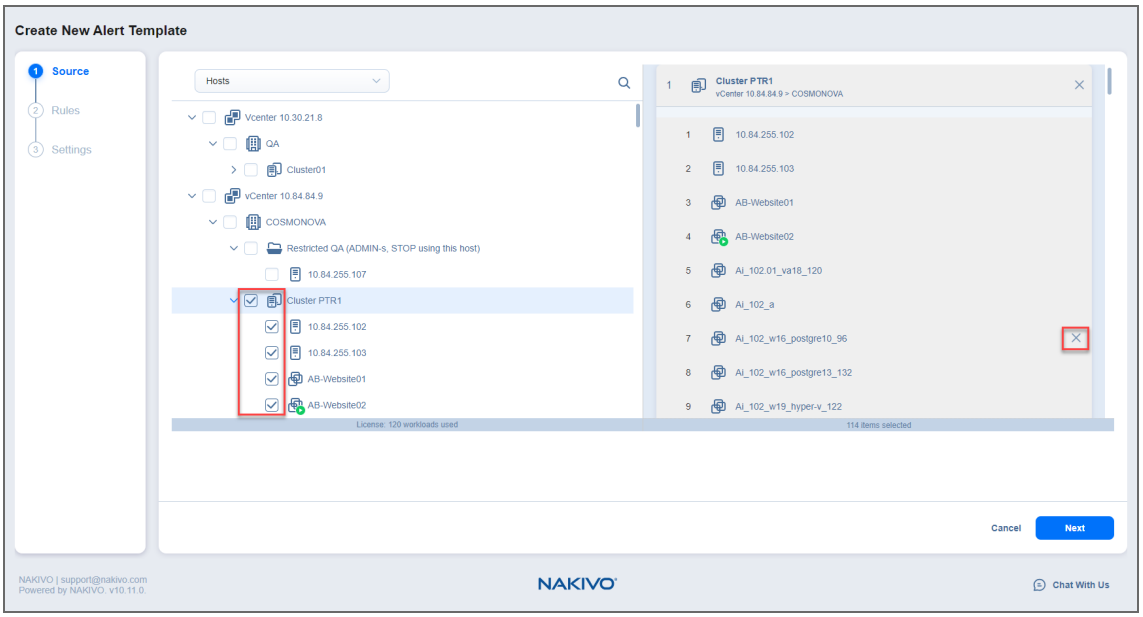



3. Complete the wizard as described in the sections below:

- “Create New Alert Template Wizard: Source” on the next page
- “Create New Alert Template Wizard: Rules” on the next page
- “Create New Alert Template Wizard: Settings” on page 1218

## Create New Alert Template Wizard: Source

1. At the **Source** step of the wizard, you can select the items to be covered by the alert template:
2. Select a view in the drop-down list:
    - **Hosts**: Displays all hosts in the **Inventory**.
    - **Datastores**: Displays all datastores in the **Inventory**.
    - **VM**: Displays all VMs in the **Inventory**.
3. Select the needed item(s). These items will appear in the right pane. You can click the **X** button to the right of an item to remove it. When done, click **Next**.

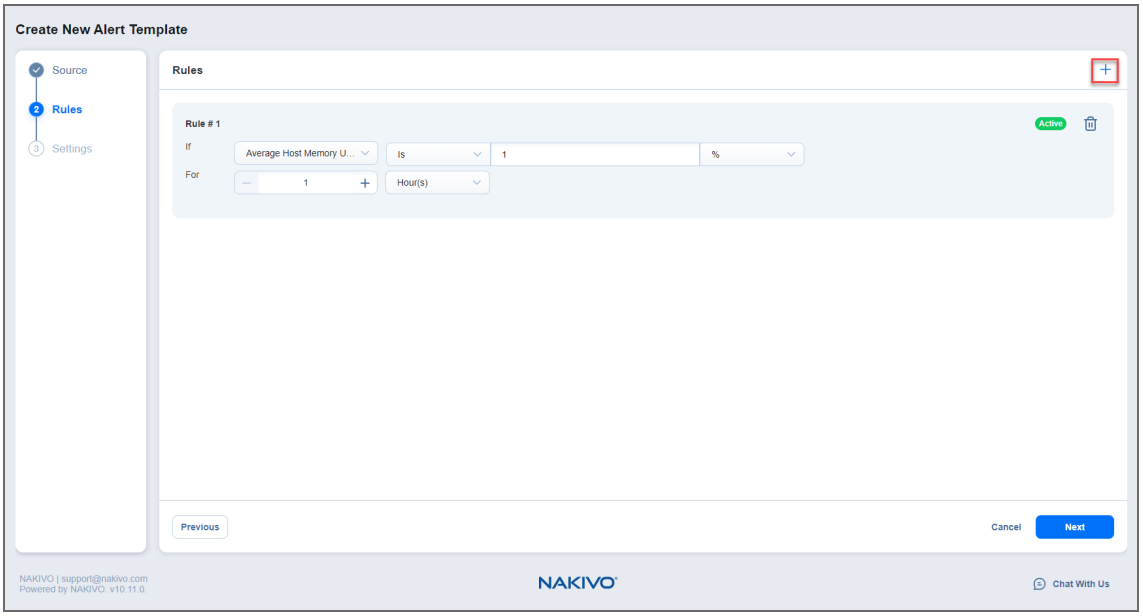



**Note**

You can select only one monitored items type view at a time. If you switch to another view, items selected in the previous view are deselected.

## Create New Alert Template Wizard: Rules

At the **Rules** step of the wizard, you can create the rules for triggering an alert for the items selected at the **Source** step. To add a rule, click the **+** (add) button.




Each rule includes the following components:

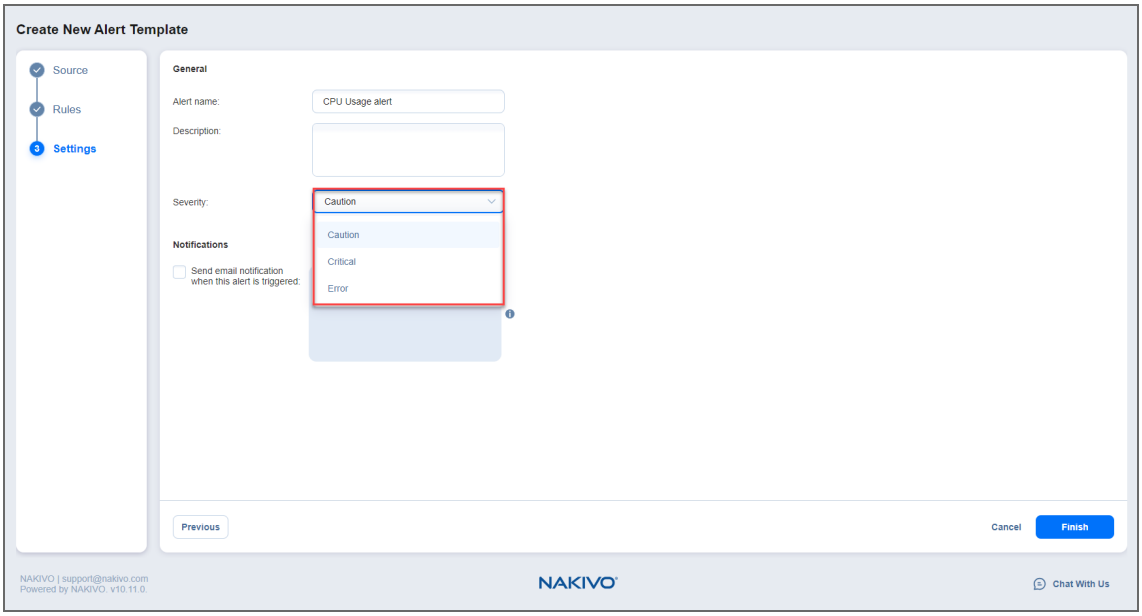
- Rule order number.
- **Active/Inactive** button: Click the button to change the status of the rule.
- **Delete** icon: Click the icon to delete a rule.
- **If** row, including the following components:
  - **Metrics**: Displays a list of metrics relevant to the item type selected in the **Source** step.
  - **Conditions**: Displays a list of conditions relevant to the selected metrics in the **Metrics** dropdown list.
  - **Metric value and unit:** Displays the value field and units button.
- **For** row: Configure a duration by selecting a unit of time and inputting a number.

You can create additional rules if needed. The maximum number of rules that you can add per template is 10. When done, click **Next**.

## Create New Alert Template Wizard: Settings

At the **Settings** step of the wizard, you can specify the alert name and severity. You can also enable email notifications for triggered alerts:

1. Enter the name, description, and severity of the alert that will be displayed in the **Alert History** and **Alert Template Management** tables. For details, refer to “Viewing Alert Template Details” on page 1215.
2. To receive alert notifications by email, select the checkbox in the **Notifications** section and specify the email addresses. You cannot enter more than 100 email addresses.
3. Click **Finish** to close the wizard. The configured alert template will be added to the **Alert Template Management** table.



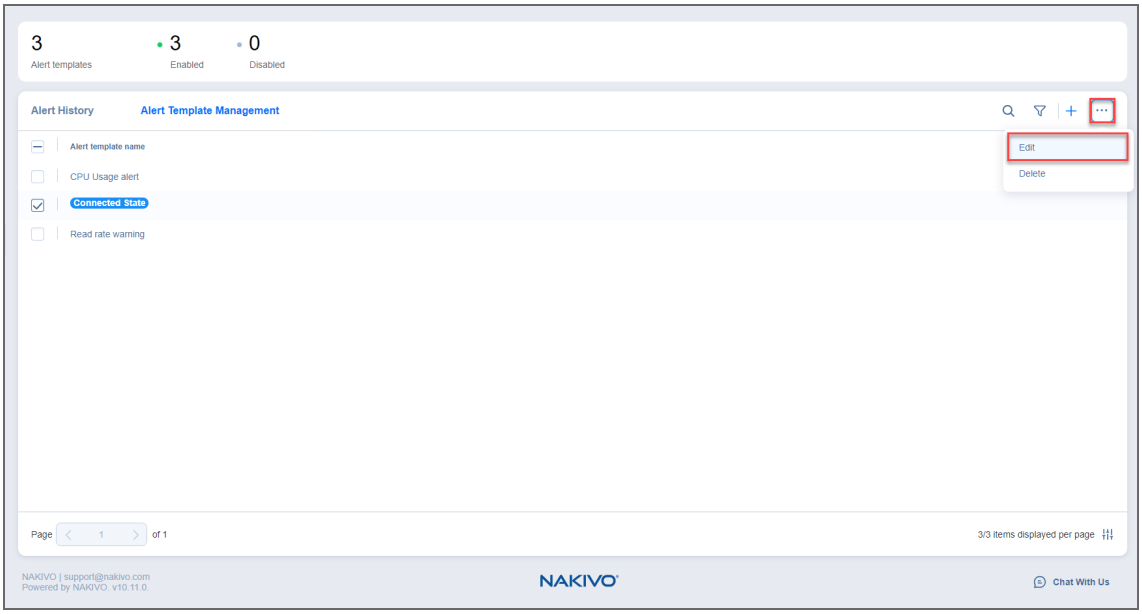

## Editing Alert Template

To edit an alert template:

Go to **Monitoring > Alerts** and click the **Alert Template Management** tab.

Select the checkbox next to the template that you want to edit, click the **...** (ellipsis) button, and then select **Edit** in the drop-down list.

The **Edit Alert Template** wizard opens. Follow the same procedure as described in the "Creating Alert Templates" on page 1216 section.




**Note**

You cannot modify the Item types in the dropdown list of the **Source** tab.

# Reports

The **Reports** dashboard allows you to add and view different types of reports about the monitored items in your VMware vSphere infrastructure. You can group reports, export reports to *PDF* or *CSV* format, as well as edit and delete them.

Refer to the following topics to learn how to use the **Reports** feature:

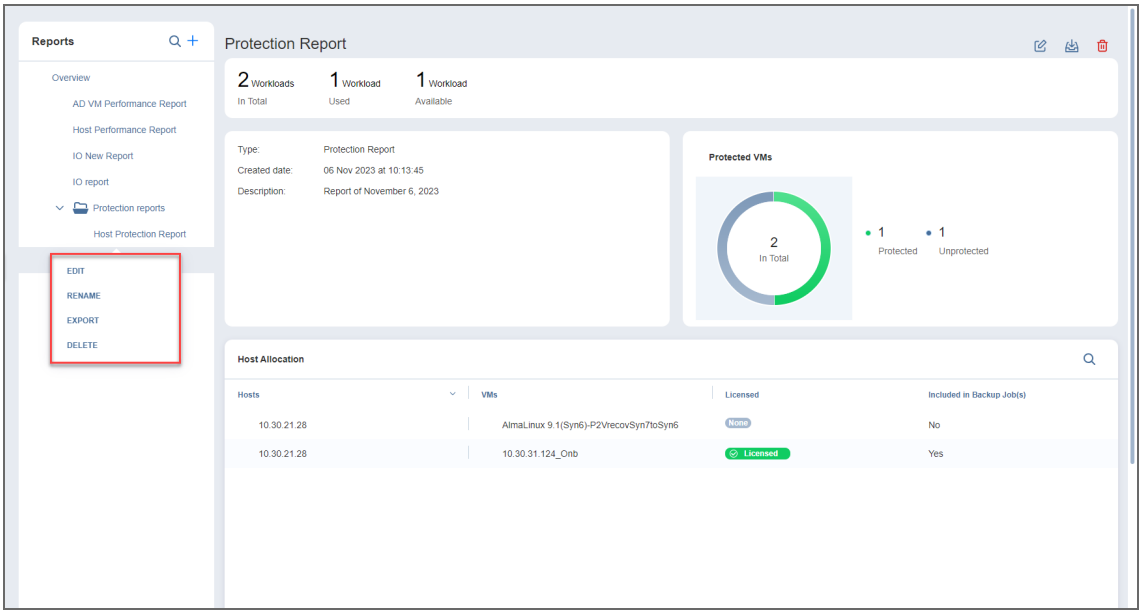
"Reports" above
"Report Types" on the next page
Creating Reports
"Editing Reports" on page 1231
"Exporting Reports" on page 1232

## Reports Overview

The **Reports** dashboard displays all reports that have been added.

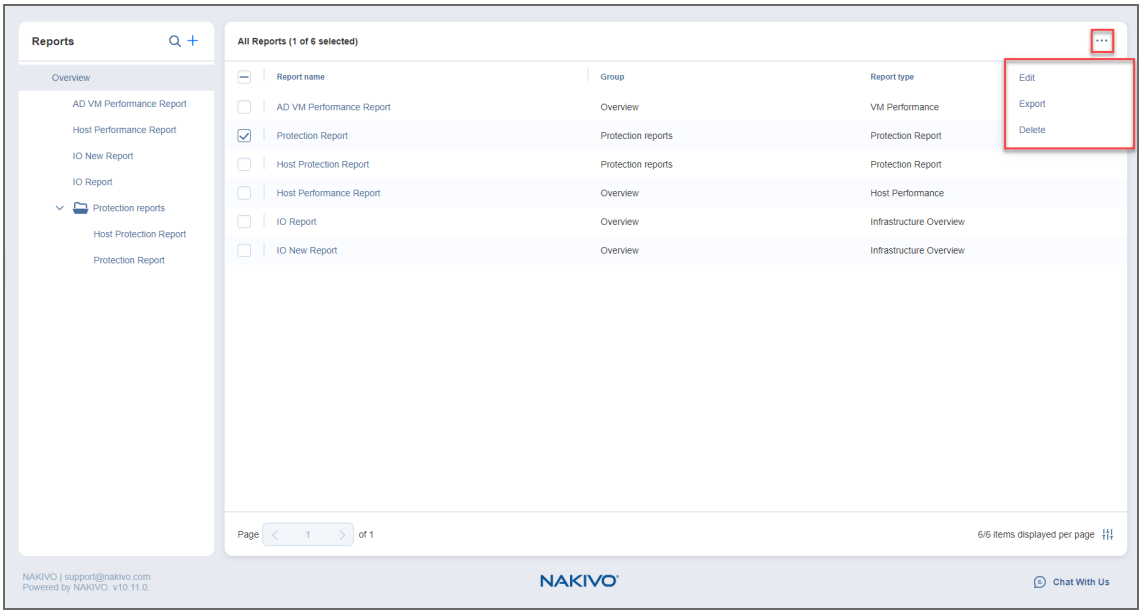
The left pane of the **Reports** dashboard includes the following options:

- **Search** button: Click the button and enter the partial or full name of a report/group in the search bar to find it.
- **+** (add) button: Click to add a new report or create a report group. For details, refer to Creating Reports.
- **Overview** list: Displays all the added reports and/or report groups. Select a report to display its details in the right panel. Right-click a report to open the dialog window with the following options: **EDIT**, **RENAME**, **EXPORT**, and **DELETE**.




In the right pane of the **Reports** dashboard, the **All Reports** table is displayed with the following options:

- **...** (ellipsis) button in the top right corner:
  - **Edit**: Select a report and click **Edit** to open the **Edit Report** wizard.
  - **Export**: Select a report and click **Export** to export it in *PDF* or *CSV* format.
  - **Delete**: Select a report and click **Delete** to remove it from the **All Reports** table.
- **Checkbox**: Select a report to edit, download, or delete reports as described above.
- **Report name**: Name of the configured report.
- **Group**: Displays the name of the group the report belongs to. If a report is not part of a specific group, **Overview** is displayed.
- **Report type**: Displays the report type selected in the **Add New Report** wizard. For details, refer to "Report Types" below and Creating Reports.




**Note**

The **Reports** dashboard can display up to 100 created reports. When this limit is reached, the **+** (add) button will be disabled.

# Report Types

The following report types are available:

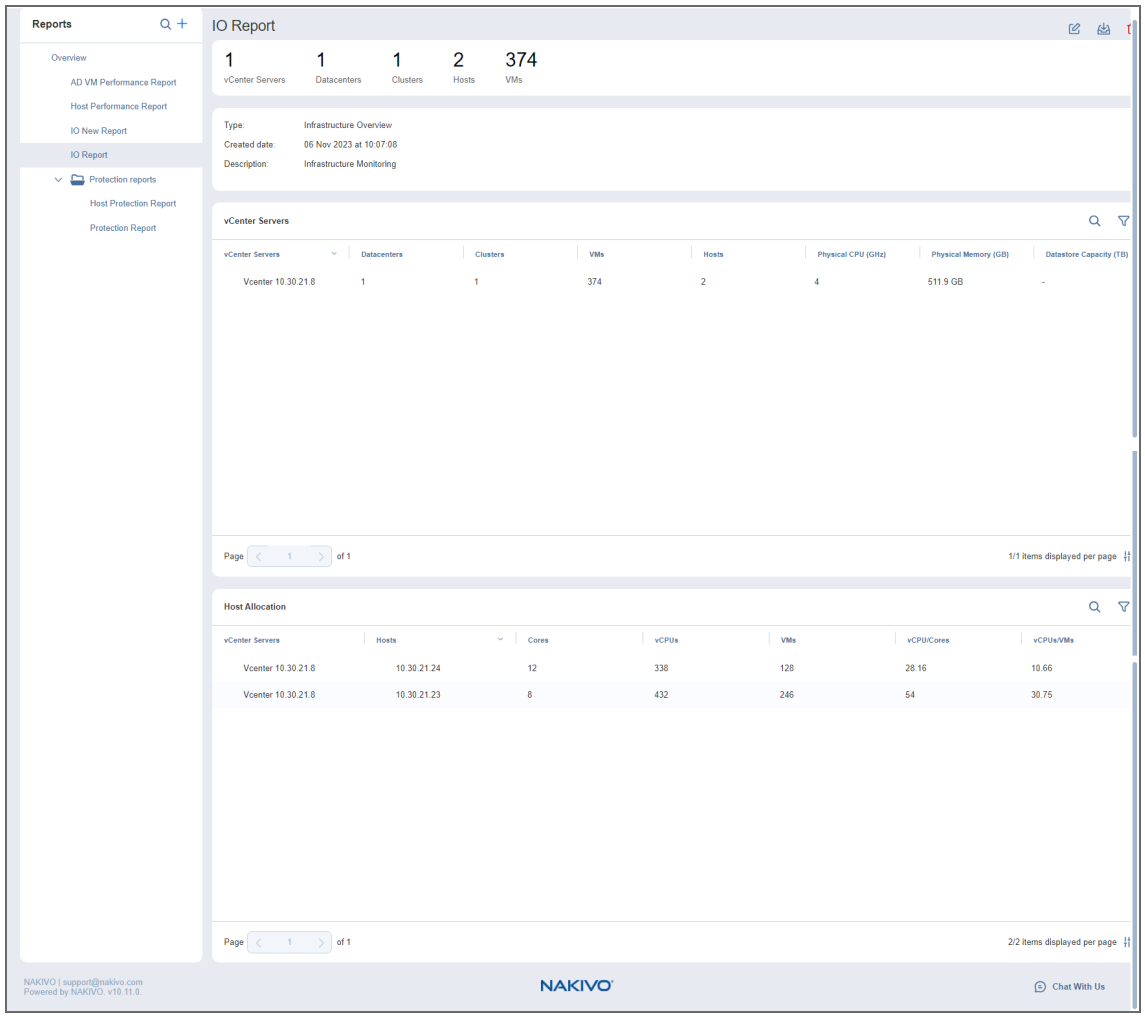

- "Infrastructure Overview" on the next page
- "VM Performance" on page 1223
- "Datastore Capacity" on page 1224
- "Host Performance" on page 1225
- "Protection Report" on page 1226

**Note**

The report content will vary depending on the report type and configurations defined in the **Add New Report** wizard or **Edit Report** wizard.

# Infrastructure Overview

Displays an overview of selected vCenters, child hosts, and standalone hosts.




- **Summary bar**: Displays the total number of infrastructure items including:
  - vCenter Servers
  - Datacenters
  - Clusters
  - Hosts
  - VMs
- **Report Settings**: Displays the report type, date the report was created, and report description.

- **vCenter Servers** table: Displays the name of all selected vCenter Servers and their respective numbers of datacenters, clusters, VMs, and hosts, as well as physical CPU (GHz), physical memory (GB), datastore capacity (TB).
- **Host Allocation** table: Displays host allocation details of the selected items. The table displays the following columns:
    - Hosts
    - vCenter Servers
    - Cores
    - vCPUs. Virtual CPUs configured for VMs
    - VMs
    - vCPU/Cores. The number of vCPU divided by the number of Cores
    - vCPUs/VMs. The number of vCPU divided by the number of VMs

**Notes**

- If added separately, child hosts and/or standalone hosts will not be displayed in the **vCenter Servers** table.
- The **Host Allocation** table will display **-** (hyphen) for child hosts and/or standalone hosts listed in the **vCenter Servers** column.

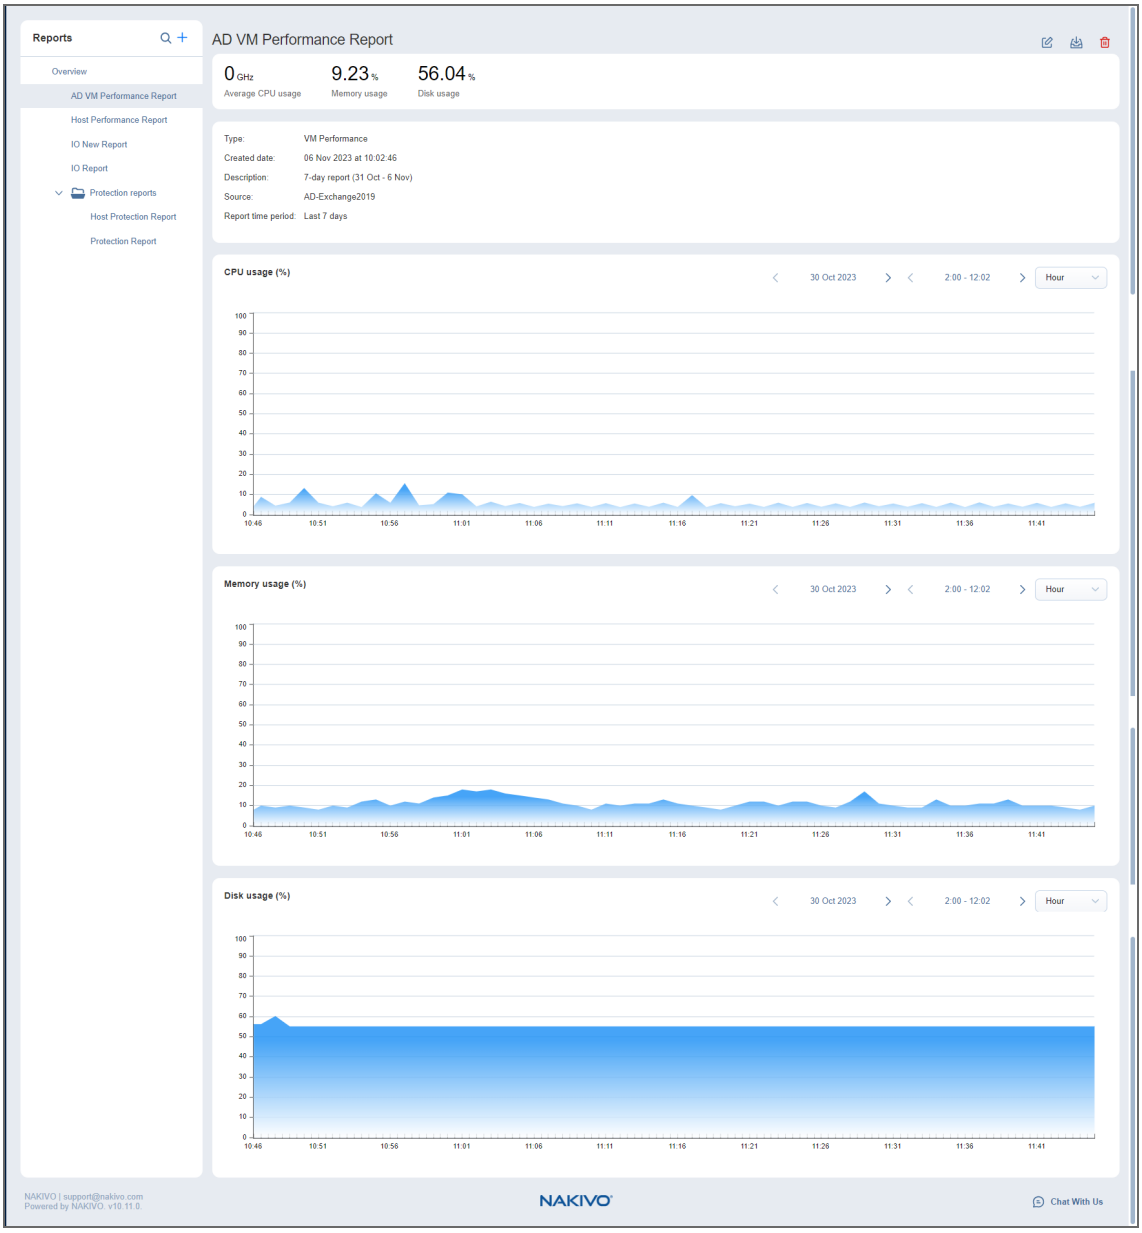
Use the **search** and **filter** buttons in the top right corner of the table to narrow down the list shown.

**Note**

Corresponding search results will be displayed for the selected table: for the **vCenter Servers** table, search will return the vCenter servers list while search in the **Host Allocation** table will return hosts list.

## VM Performance

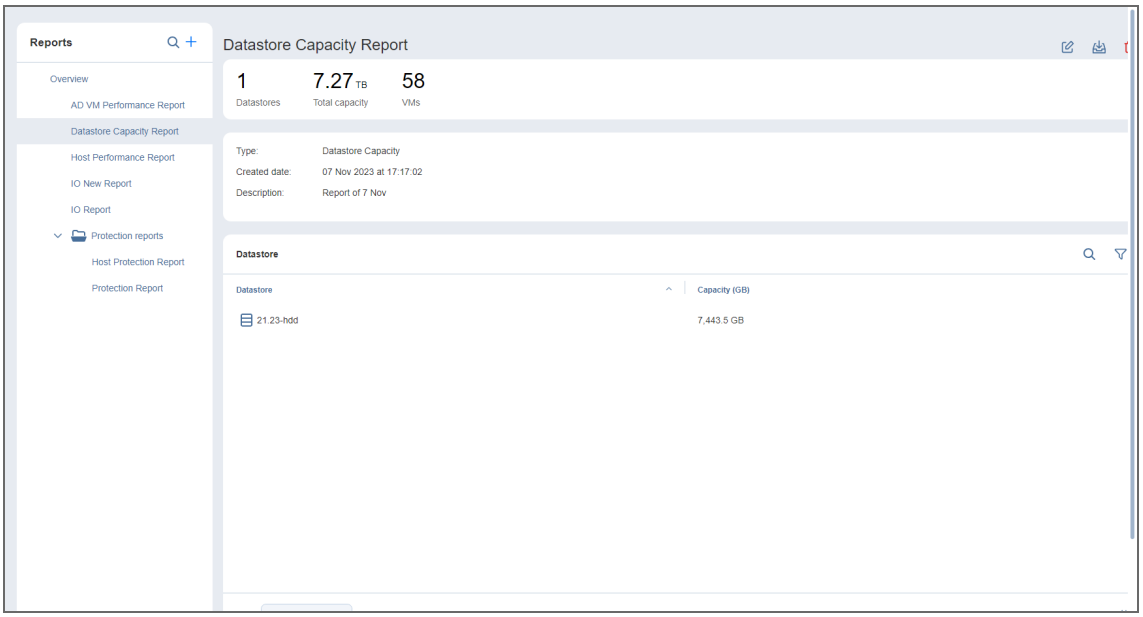
Displays the changes in the average VM CPU usage, memory usage, and disk usage for the items.




- **Summary bar**: Displays the following:
  - Average CPU Usage in GHz
  - Memory Usage in %
  - Disk Usage in %
  - **Report Settings**: Displays the report type, date the report was created, description of the report, source items, and report time period.
- **CPU usage (%)**, **Memory usage (%)**, and **Disk usage (%)** graphs: Display the average CPU usage, memory usage, and disk usage of selected sources.

# Datastore Capacity

Displays information about the capacity of monitored datastores.

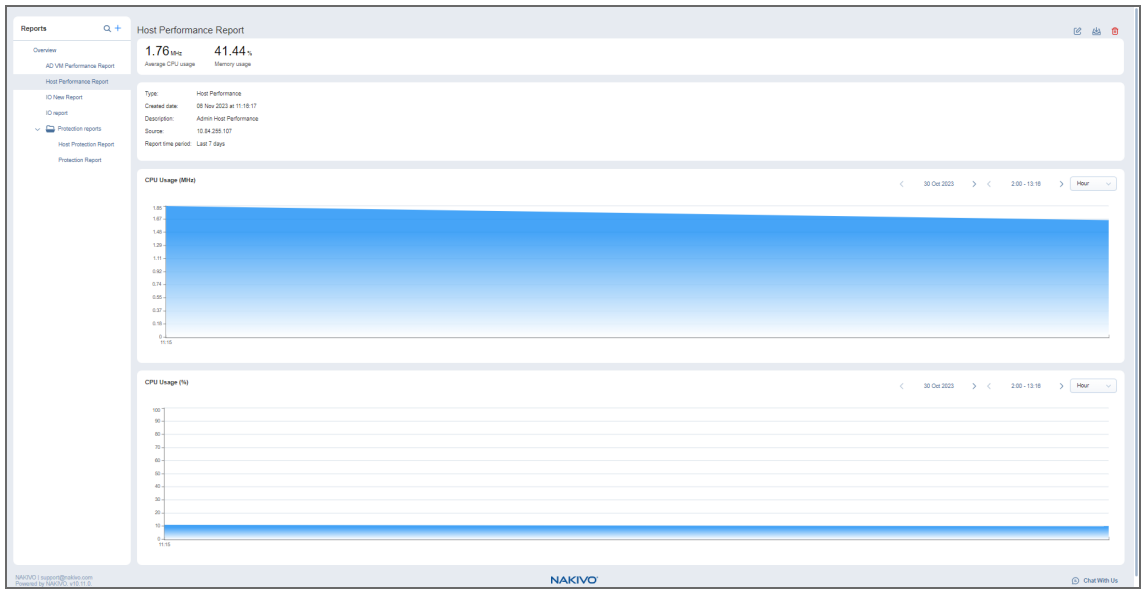



- **Summary bar**: Displays the following:
  - Number of datastores
  - Total capacity for all datastores in TB
  - Total number of VMs stored on these datastores
- **Report Settings**: Displays the report type, the date the report was created, and the description of the report.
- **Datastore** table: Displays the search and filter buttons and the following columns:
  - Datastore
  - Capacity (GB)
  - Free space (GB)
  - Used space
  - VM count: Displays the number of VMs stored on the datastores.

If needed, use the **search** and **filter** buttons in the top right corner of the table to narrow down the list shown.

# Host Performance

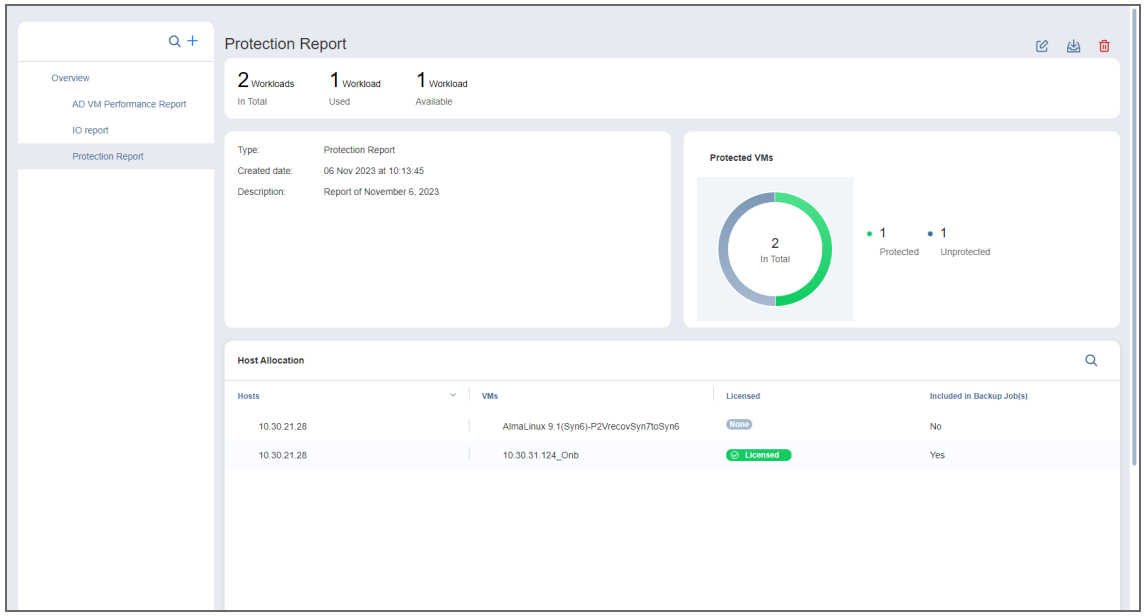
Displays the changes in the average CPU usage and memory usage metrics from the selected sources.




- **Summary bar**: Displays the following:
  - Average CPU usage in MHz
  - Memory usage in % unit
- **Report Settings**: Displays the report type, date the report was created, description of the report, source items, and report time period.
- **CPU usage (MHz)** and **CPU Usage (%)** graphs: Display the average CPU usage of selected sources.

# Protection Report

Captures and displays the current (by the time the report is created) state/information of applied NAKIVO Backup & Replication licenses for selected hosts and VMs.



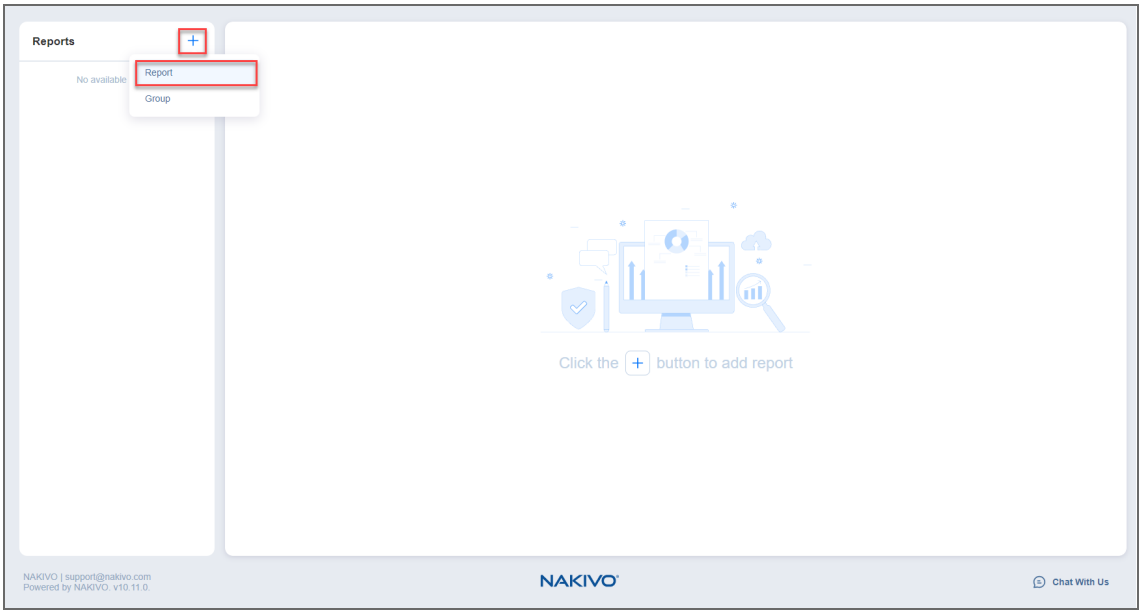
- **Summary bar**: Depending on the license type, the summary bar will display the following:
  - A total number of discovered NAKIVO Backup & Replication workloads/sockets
  - A total number of used NAKIVO Backup & Replication workloads/sockets
  - A total number of available workloads/sockets
- **Report Settings**: Displays the report type, the date when the report was created, and the description of the report.
- **Protected VMs** chart: Displays the number of protected (licensed) and unprotected (unlicensed) VMs from selected source items.
- **Host Allocation** table: Displays host allocation details of the selected source items. The table displays the following columns:
  - **Hosts**
  - **VMs**
  - **Licensed**: Displays **Licensed** if a NAKIVO Backup & Replication license is applied to the monitored item for backup jobs.
  - **Included in Backup Job(s)**: Displays **Yes** if the VMs have been backed up in a backup job.

If needed, use the **search** button in the top right corner of the table to narrow down the list shown.

# Adding Reports

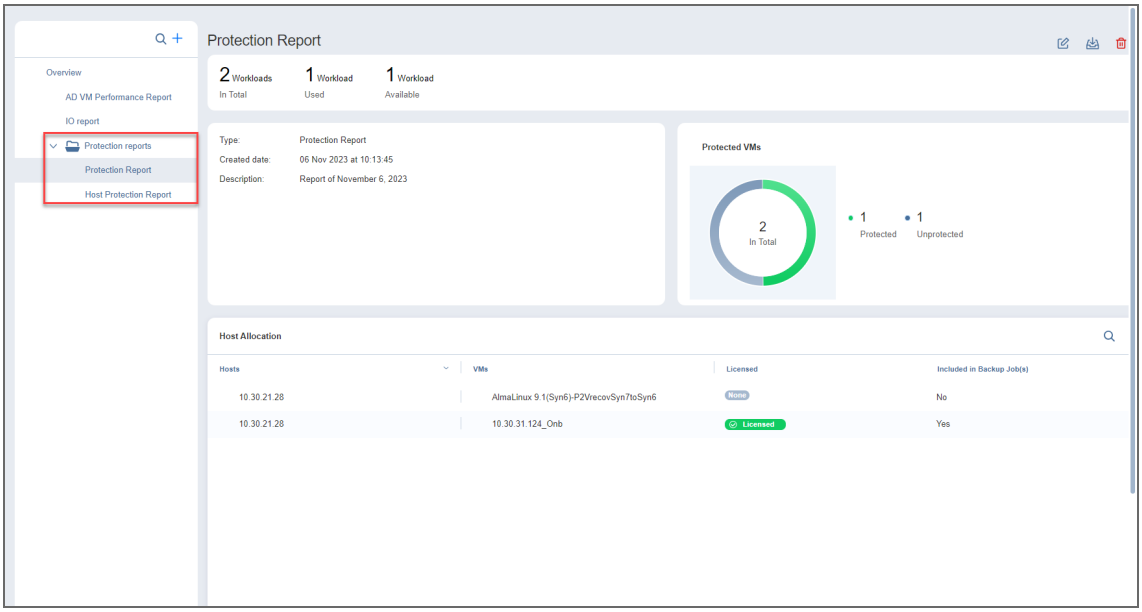
To create a report:

1. Go to **Monitoring > Reports**.
2. In the **Reports** dashboard, click the **+** (add) button. The **Add New Report** wizard opens.



**Note**

You can also select **Group** to add a report group and then drag the needed report to it in the **Reports Overview** pane.




Complete the wizard as described in the sections below:

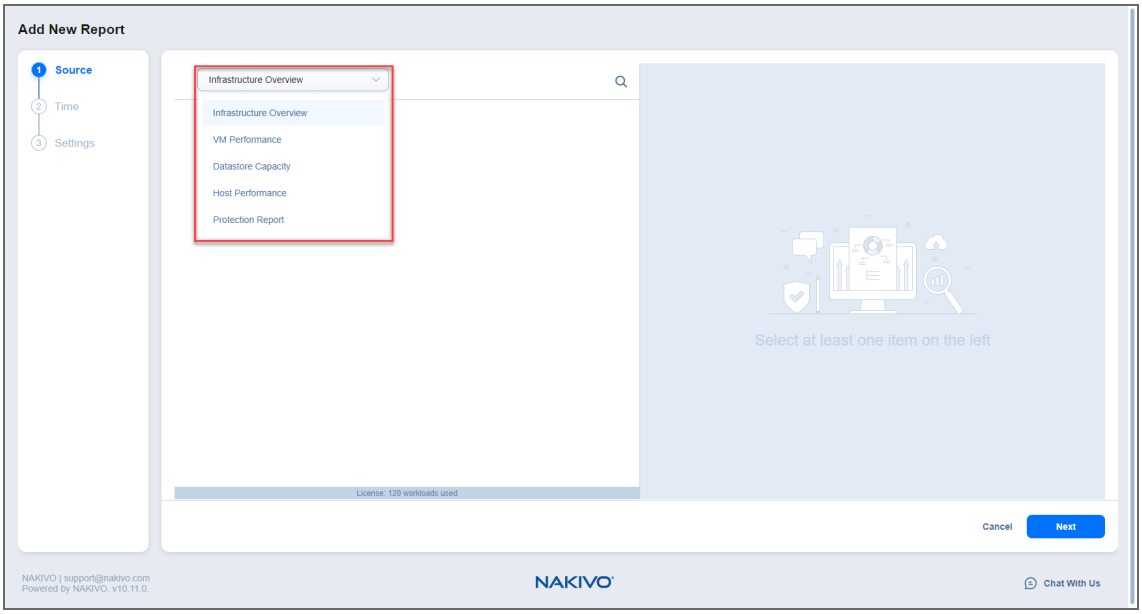
- “Add New Report Wizard: Source” below
- “Create New Report Wizard: Time” on the next page
- “Create New Report Wizard: Settings” on page 1230

# Add New Report Wizard: Source

At the **Source** step of the wizard, you can….

1. Select the reports type in the drop-down list:
    - **Infrastructure Overview**: Displays all discovered vCenters, child hosts, and standalone hosts.
    - **VM Performance**: Displays all VMs with hosts and clusters, which were added to monitoring items.
    - **Datastore Capacity**: Displays all the datastores that were added to monitoring items.
    - **Host Performance**: Displays all the hosts that were added to monitoring items.
    - **Protection Report**: Displays all discovered hosts and all child VMs.

    For details, refer to “Report Types” on page 1221.




2. Select the needed items and then:
    - Click **Next** to proceed to the next step.
    - Click **Cancel** to close the wizard.

If necessary, click the **X** button to the right of an item in the right pane to remove it.

> **Note**
> You cannot select multiple report types.

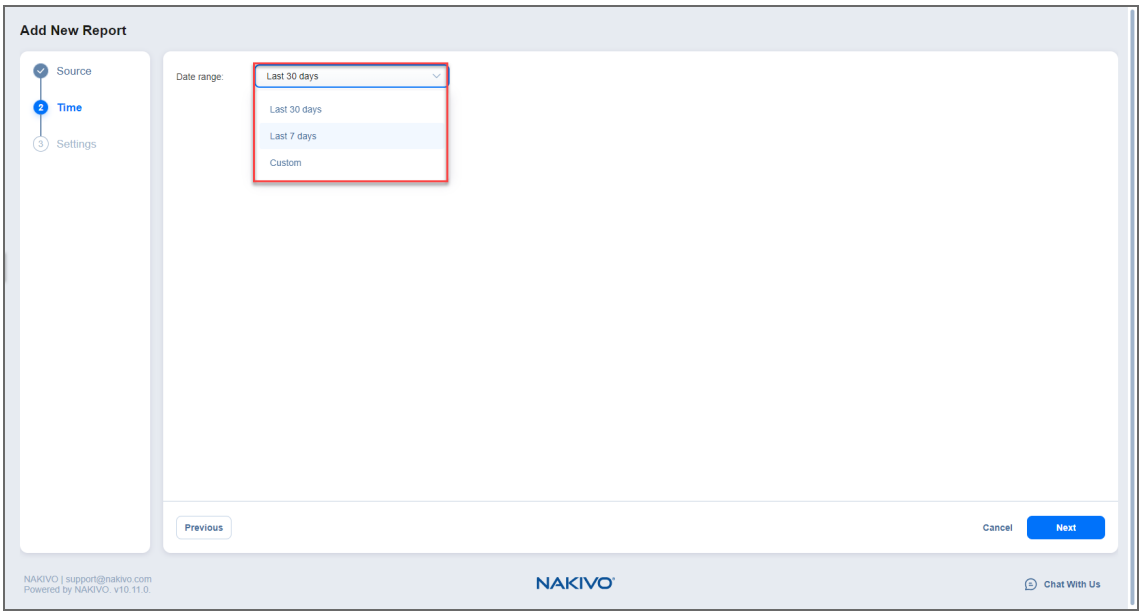
If necessary, use the search box to display items based on the selected report.

## Create New Report Wizard: Time

At the **Time** step, you can specify the date range for a report.

Select the time period the report should cover and then click **Next**. The following options are available:

- **Last 30 days**
- **Last 7 days**
- **Custom:** Select to pick a custom date range




> **Note**
>
> If you select **Infrastructure Overview**, **Datastore Capacity**, or **Protection Report** at the **Source** step, you can only generate a report for the same day.

Click **Cancel** to close the wizard or click on a previous step to view or make changes.

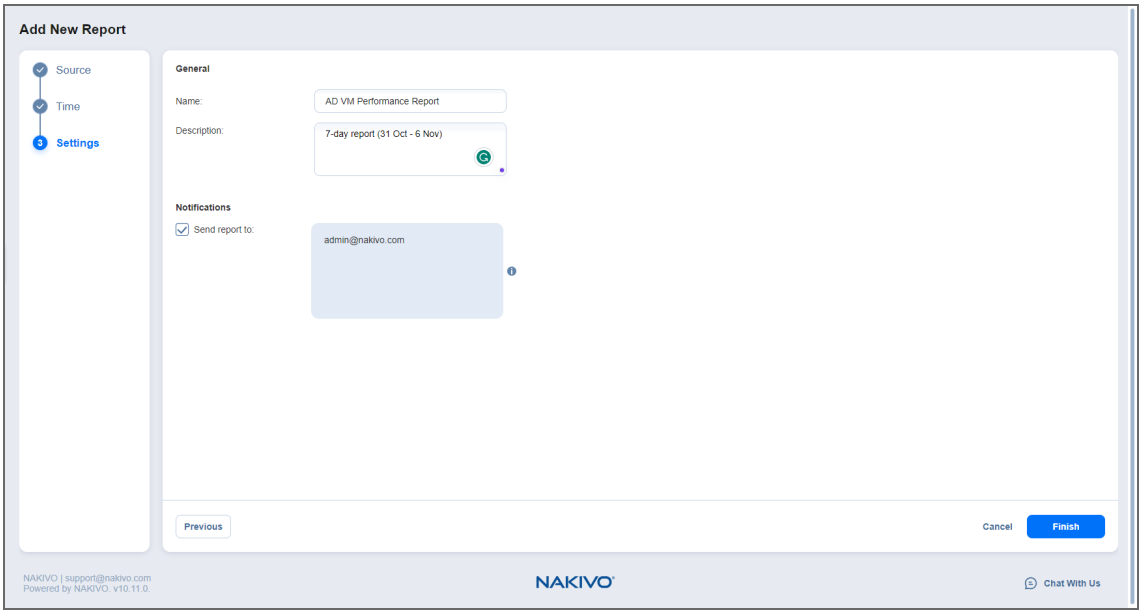
## Create New Report Wizard: Settings

At the **Settings** step:

1. In the **General** section: Provide a name and description (optional) of the report.
2. In the **Notifications** section: Select the **Send report to** checkbox and enter the email recipient(s) you want to send the report to. For an email address to be saved for the report, press **Enter** to turn the entered email address into a tag.
3. Click **Finish** to create the report. The configured report will be added to the **Reports** table and the report will be sent to the specified email recipients if the option was selected.

> **Note**
>
> You cannot enter more than 100 email addresses.

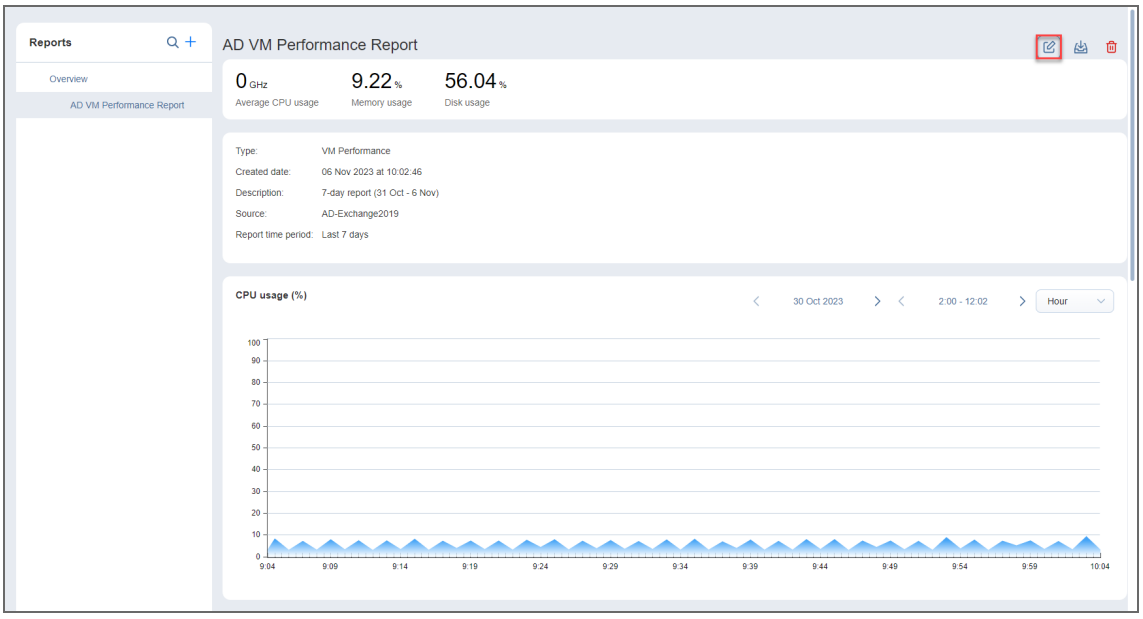
Click **Cancel** to close the wizard or click on a previous step to view or edit it.




# Editing Reports

To edit a report:

1. Go to **Monitoring > Reports** and click **Overview**.
2. In the **All Reports** table, select a report, click the **...** (ellipsis) button, and then select **Edit** in the dropdown list. Alternatively, select a report in the left panel and click the **Edit** icon in the top right of the panel. The **Edit Report** wizard opens.




3. Follow the procedure described in “Adding Reports” on page 1227.

**Note**

The **Item types** in the dropdown list of the **Source** tab cannot be modified in the **Edit Report** wizard.

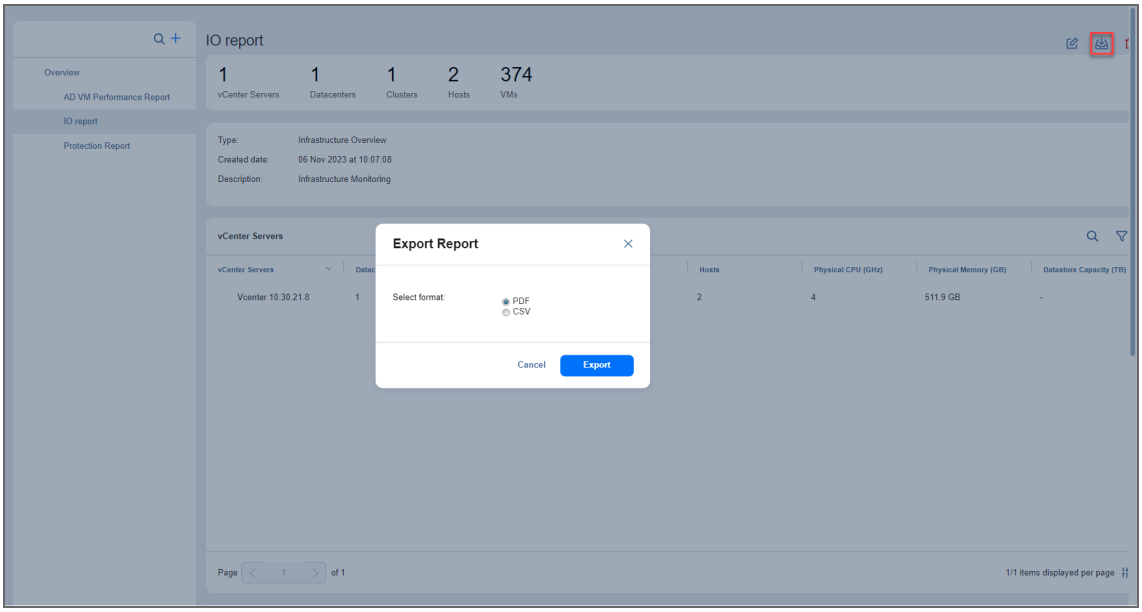
# Exporting Reports

**Note**

The report content will vary depending on the report type and configurations defined in the Add New Report wizard or Edit Report wizard.

To export a report:

1. Go to **Monitoring > Reports** and click **Overview**.
2. In the **All Reports** table, select a report, click the **...** (ellipsis) button, and then select **Export** in the dropdown list. Alternatively, select a report in the left panel and click the **Export** icon in the top right of the panel. The **Export Report** wizard opens.
3. Select the format: *PDF* or *CSV*, and click **Export**.




**Note**

**Host Performance** and **VM Performance** reports can be exported to the *PDF* format only.

# Integration and Automation

This section contains the following topics:


- "Command Line Interface" on page 1234
- "Automation with HTTP API" on page 1248
- "Aptare IT Analytics Integration" on page 1249

# Command Line Interface

This section covers the following topics:


- “Using Command Line Interface” on page 1235
- “Available Commands” on page 1237
- “Exit Codes” on page 1247

# Using Command Line Interface


- "Operation Modes of Command Line Interface" below
- "Using Command Line Interface Locally" below
- "Using Command Line Interface Remotely" below
- "Using Command Line Interface in Multi-Tenant Mode" on the next page


NAKIVO Backup & Replication allows you running actions from the product's command line interface (CLI). In case credentials are configured for the product, running an action via CLI requires providing administrator credentials as arguments, namely, `--username [login] --password [password]`, where `[login]` is the administrator user name and `[password]` is the administrator password.

## Operation Modes of Command Line Interface

You can run CLI in either of the following modes:

- *Interactive mode.* This allows you to use a single login for a session. When opened in the interactive mode, CLI allows you executing commands without dashes.
  To open CLI in the interactive mode, enter `cli.bat --interactive --username [login] --password [password]` and press **Enter**. To exit the CLI interactive mode, enter **Ctrl-C**.
- *Non-interactive mode.* This requires entering your credentials for each command. You will have to enter dashes before commands. For example: `cli.bat --username [login] --password [password] --inventory-list`

## Using Command Line Interface Locally

To use CLI on the machine where NAKIVO Backup & Replication Director is installed, follow the steps below:

1. Run the CLI executable:
   - If NAKIVO Backup & Replication is installed on a Windows OS, run the `cli.bat` file located in the `bin` folder inside the product installation folder ("`C:\Program Files\NAKIVO Backup & Replication`" by default).
   - If NAKIVO Backup & Replication is installed on a Linux OS, run the `cli.sh` file located in the `director/bin` folder inside the product installation folder (`/opt/nakivo/` by default).
2. Run available commands.

## Using Command Line Interface Remotely

To use CLI from a remote machine, follow the steps below:

1. Copy the CLI executable and `.jar` files to the machine from where you plan to use the CLI:
   - If NAKIVO Backup & Replication is installed on a Windows OS, copy the `cli.bat` and `cli.jar` files located in the `bin` folder inside the product installation folder ("`C:\Program Files\NAKIVO Backup & Replication`" by default).
   - If NAKIVO Backup & Replication is installed on a Linux OS, copy the `cli.sh` and `cli.jar` files located in the `director/bin` folder inside the product installation folder (`/opt/nakivo/` by default).
2. On the machine from where you plan to use the CLI, configure the `PATH` system variable as described at http://java.com/en/download/help/path.xml
3. Run commands using the following format: `<command> <host> <port> <username> <password>`

**Example**

To get a list of jobs of the product which is installed on the machine with the `192.168.10.10` IP address, uses the `4443` port number for the **Director Web HTTPS port**, and has "`admin`" as login and password for the product's web UI, run the following command: `--job-list --host 192.168.10.10 --port 4443 --username admin --password admin`

## Using Command Line Interface in Multi-Tenant Mode

Triggering an action inside a tenant in the multi-tenant mode via command line interface requires providing a tenant ID as an argument:

```
cli.bat --repository-detach [repo_id] --username [login] --password [password] --tenant [tenant-id]
```

# Available Commands

You can run CLI commands either in interactive or non-interactive mode. Refer to the *Operation Modes of Command Line Interface* subsection of the "Using Command Line Interface" on page 1235 topic.

Use either long or short form of the commands*.

| Command | Long form | Short form | Output |
|---|---|---|---|
| Help | | | |
| General help | cli.bat --help | cli.bat -h | • Command name<br>• Description |
| Job Management | | | |
| List all jobs | cli.bat --job-list | cli.bat -jl | • Job ID<br>• Job name<br>• Current job status<br>• Job last run result |
| Start a job | cli.bat --job-start [job_id]<br>Options:<br>• --backup-type<br>  • Arguments:<br>  ◦ full<br>  ◦ incremental<br>• --full-backup-mode<br>  • Arguments:<br>  ◦ active<br>  ◦ synthetic<br>Example: --job-start 1 --backup-type full --full-backup-mode active | cli.bat -jr [job_id] | |
| Stop a job | cli.bat --job-stop [job_id] | cli.bat -js [job_id] | |
| Disable a job | cli.bat --job-disable [job_id] | cli.bat -jd [job_id] | |

| Command | Long form | Short form | Output |
|---|---|---|---|
| Disable multiple jobs | cli.bat --job-disable [job_id1] [job_id2] [job_id3] ... [job_idX] | cli.bat -jd [job_id1] [job_id2] [job_id3] ... [job_idX] | |
| Enable a job | cli.bat --job-enable [job_id] | cli.bat -je [job_id] | |
| Enable multiple jobs | cli.bat --job-enable [job_id1] [job_id2] [job_id3] ... [job_idX] | cli.bat -je [job_id1] [job_id2] [job_id3] ... [job_idX] | |
| Generate a report for a job in PDF format | cli.bat --job-report [job_id]<br>• The command with no arguments creates the job report and saves it to the current directory.<br>• To save the report to other directory: cli.bat --job-report [job_id] --save-to [dir_path]<br>• To send the report to default email(s): cli.bat --job-report [job_id] --send-by-email<br>• To send the report to other email: cli.bat --job-report [job_id] --send-by-email [email_address] | cli.bat -jp [job_id]<br>• The command with no arguments creates the job report and saves it to the current directory.<br>• To save the report to other directory: cli.bat -jp [job_id] -f [dir_path]<br>• To send the report to default email(s): cli.bat -jp [job_id] -eml<br>• To send the report to other email: cli.bat -jp [job_id] -eml [email_address] | |

| Command | Long form | Short form | Output |
|---|---|---|---|
| Generate a report for a job in CSV format | cli.bat --job-report [job_id] --csv<br>• The command with no arguments creates the job report and saves it to the current directory.<br>• To save the report to other directory: cli.bat --job-report [job_id] --save-to [dir_path] --csv<br>• To send the report to default email(s): cli.bat --job-report [job_id] --send-by-email --csv<br>• To send the report to other email: cli.bat --job-report [job_id] --send-by-email [email_address] --csv | cli.bat -jp [job_id] --csv<br>• The command with no arguments creates the job report and saves it to the current directory.<br>• To save the report to other directory: cli.bat -jp [job_id] -f [dir_path] --csv<br>• To send the report to default email(s): cli.bat -jp [job_id] -eml --csv<br>• To send the report to other email: cli.bat -jp [job_id] -eml [email_address] --csv | |
| Return information about a job | cli.bat --job-info [job_id] | cli.bat -ji [job_id] | • Job ID<br>• Job name<br>• Current job status<br>• Job last run result |
| Inventory | | | |

| Command | Long form | Short form | Output |
|---|---|---|---|
| List all inventory items | cli.bat --inventory-list | cli.bat -il | • Item ID<br>• Item IP/host name<br>• Item type (host/vCenter)<br>• Item children count (X hosts, E VMs)<br>• Item current state<br>• Item current status |
| Update all inventory items | cli.bat --inventory-update | cli.bat -iu | |
| Update an inventory item | cli.bat --inventory-update [item_id] | cli.bat -iu [item_id] | |
| Return information about an inventory item | cli.bat --inventory-info [item_id] | cli.bat -ii [item_id] | • Item ID<br>• Item IP/host name<br>• Item type (host/vCenter)<br>• Item children count (X hosts, E VMs)<br>• Item current state<br>• Item current status |
| Transporters | | | |

| Command | Long form | Short form | Output |
|---|---|---|---|
| List all transporters | cli.bat --transporter-list | cli.bat -trl | • Transporter ID<br>• Transporter IP/host name<br>• Transporter current load<br>• Transporter maximum load<br>• Transporter current state<br>• Transporter current status |
| Update all transporters | cli.bat --transporter-update | cli.bat -tru | |
| Update a transporter | cli.bat --transporter-update [transporter_id] | cli.bat -tru [transporter_id] | |
| Return information about a transporter | cli.bat --transporter-info [transporter_id] | cli.bat -tri [transporter_id] | • Transporter ID<br>• Transporter IP/host name<br>• Transporter current load<br>• Transporter maximum load<br>• Transporter current state<br>• Transporter current status |
| Repositories | | | |

| Command | Long form | Short form | Output |
|---|---|---|---|
| List all repositories | cli.bat --repository-list | cli.bat -rl | • Repository ID<br>• Repository name<br>• Assigned transporter<br>• Backup count<br>• Free space<br>• Attached or detached<br>• Consistent or inconsistent<br>• Repository current state<br>• Repository current status |
| Update all repositories | cli.bat --repository-update | cli.bat -ru | |
| Update a repository | cli.bat --repository-update [repo_id] | cli.bat -ru [repo_id] | |
| Detach a repository | cli.bat --repository-detach [repo_id] | cli.bat -rd [repo_id] | |
| Attach a repository | cli.bat --repository-attach [repo_id] | cli.bat -ra [repo_id] | |
| Start repository maintenance | cli.bat --repository-maintenance [repo_id] [parameter]<br>Parameters:<br>• --selfheal<br>• --verify<br>• --spacereclaim | cli.bat -rm [repo_id] [parameter]<br>Parameters:<br>• --selfheal<br>• --verify<br>• --spacereclaim | |
| Stop repository maintenance | cli.bat --repository-maintenance-stop [repo_id] | cli.bat -rms [repo_id] | |

| Command | Long form | Short form | Output |
| --- | --- | --- | --- |
| Return information about a repository | cli.bat --repository-info [repo_id] | cli.bat -ri [repo_id] | • Repository ID<br>• Repository name<br>• Assigned transporter<br>• Backup count and free space<br>• Attached or detached<br>• Consistent or inconsistent<br>• Repository current state<br>• Repository current status |
| Support | | | |

| Command | Long form | Short form | Output |
|---|---|---|---|
| Generate a support bundle | cli.bat --bundle-create<br>• The command with no parameters will create a support bundle and save it in the current directory.<br>• To save the bundle to other directory: cli.bat --bundle-create --save-to [dir_path]<br>• To send the bundle to support over email: cli.bat --bundle-create --send-to-support<br>• To send the bundle to other email: cli.bat --bundle-create --send-by-email [email_address] | cli.bat -bc<br>• The command with no parameters will create a support bundle and save it in the current directory.<br>• To save the bundle to other directory: cli.bat -bc -f [dir_path]<br>• To send the bundle to support over email: cli.bat -bc -sup<br>• To send the bundle to other email: cli.bat -bc -eml [email_ address] | |
| Licensing | | | |
| Get the current license information | cli.bat --license-info | cli.bat -li | |
| Replace the current license with a new license file | cli.bat --license-replace [file_ path] | cli.bat -lin [file_path] | |
| Multi-Tenancy | | | |

| Command | Long form | Short form | Output |
|---|---|---|---|
| List all tenants | cli.bat --tenant-list | cli.bat -tl | • Tenant ID<br>• Tenant name<br>• Allocated items type and count<br>• Tenant status<br>• Enabled or disabled |
| Disable a tenant | cli.bat --tenant-disable [tenant_id] | cli.bat -td [tenant_id] | |
| Enable a tenant | cli.bat --tenant-enable [tenant_id] | cli.bat -te [tenant_id] | |
| Return information about a tenant | cli.bat --tenant-info [tenant_id] | cli.bat -ti [tenant_id] | • Tenant ID<br>• Tenant Account ID<br>• Tenant name<br>• Allocated items type and count<br>• Tenant status<br>• Enabled or disabled |
| Accept the new certificate of master tenant | --msp-certificate-accept | -mca | |
| Create a support bundle for master admin level | • Generate the support bundle for master level only: cli.bat --bundle-create<br>• Generate the support bundle with all tenants logs: cli.bat --bundle-create --include-tenants | • Generate the support bundle for master level only: cli.bat -bc<br>• Generate the support bundle with all tenants logs: cli.bat -bc -ite | |
| Miscellaneous | | | |

| Command | Long form | Short form | Output |
|---|---|---|---|
| Get the CLI version | cli.bat --version<br>The command returns the CLI version which is equal to the full version of NAKIVO Backup & Replication. | - | |
| Run a command in the debug mode | cli.bat --repository-info [repo_ id] --debug<br>This is an option that can be added to any other CLI command.<br>With the debug mode turned on, the commands will return the full error text. | cli.bat -ri [repo_id] --debug | |

*Examples are given for Windows OS.

# Exit Codes

NAKIVO Backup & Replication CLI provides the following exit codes:

- **0:** Normal
- **1:** Unknown command
- **2:** Cannot login
- **3:** Command failed
- **4:** Local failure
- **5:** No arguments

# Automation with HTTP API

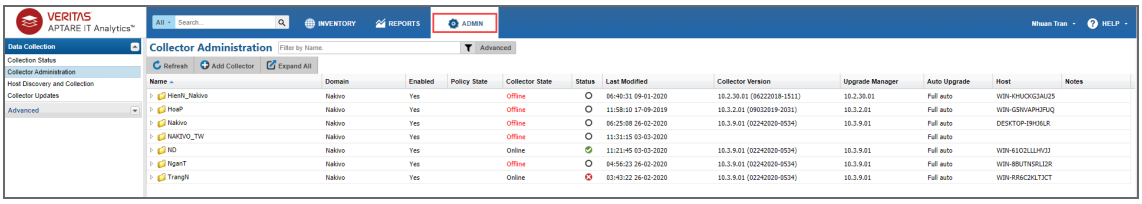
HTTP API allows you to run common NAKIVO Backup & Replication commands outside of the product web interface.

The API is JSON-RPC based. For detailed request and response syntax, refer to API Reference.

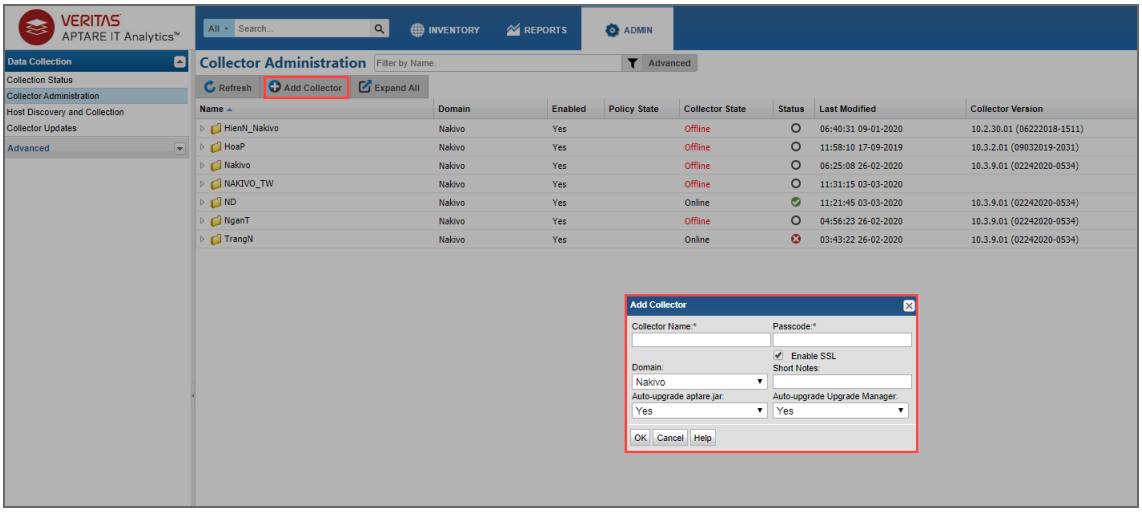
# Aptare IT Analytics Integration

APTARE IT Analytics is a storage resource management platform for integrating storage and backup solutions. The integration with NAKIVO Backup & Replication is based on an APTARE data collector that sends storage component information to the system's platform. The steps for integrating NAKIVO Backup & Replication with APTARE IT Analytics are as follows:

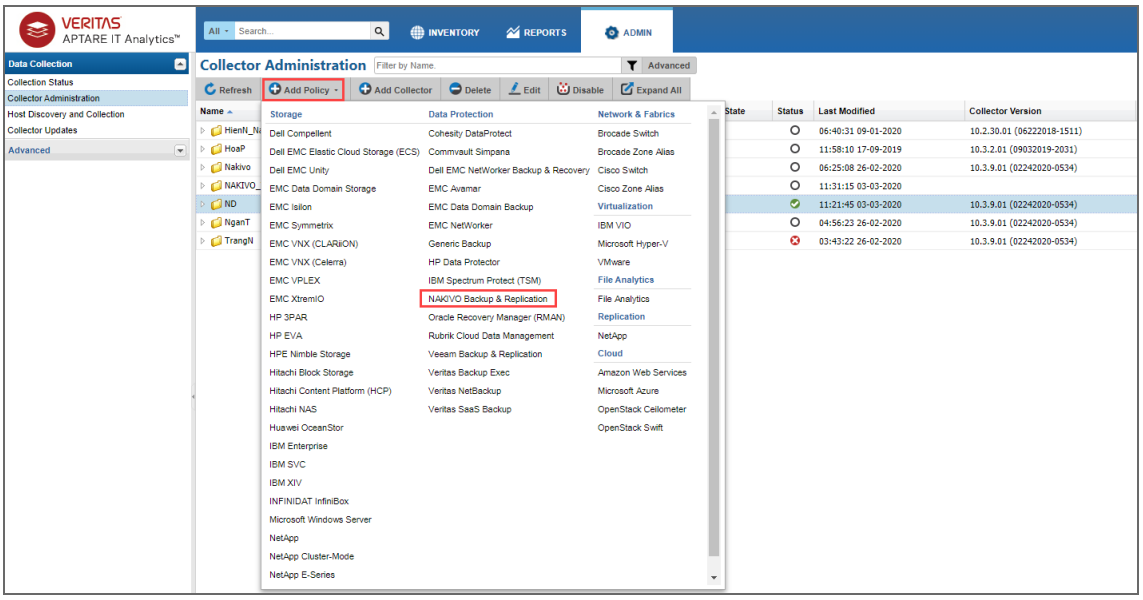
1. On the machine where NAKIVO Backup & Replication is deployed, do the following:
   a. Install APTARE StorageConsole Data Collector with NAKIVO connector.
   b. When the installation has been successfully completed, make sure that the APTARE Agent service is running.
2. Open your NAKIVO Backup & Replication instance and run your backup jobs.
3. Log in to the APTARE portal.
4. Go to the **ADMIN** tab and take the following steps:

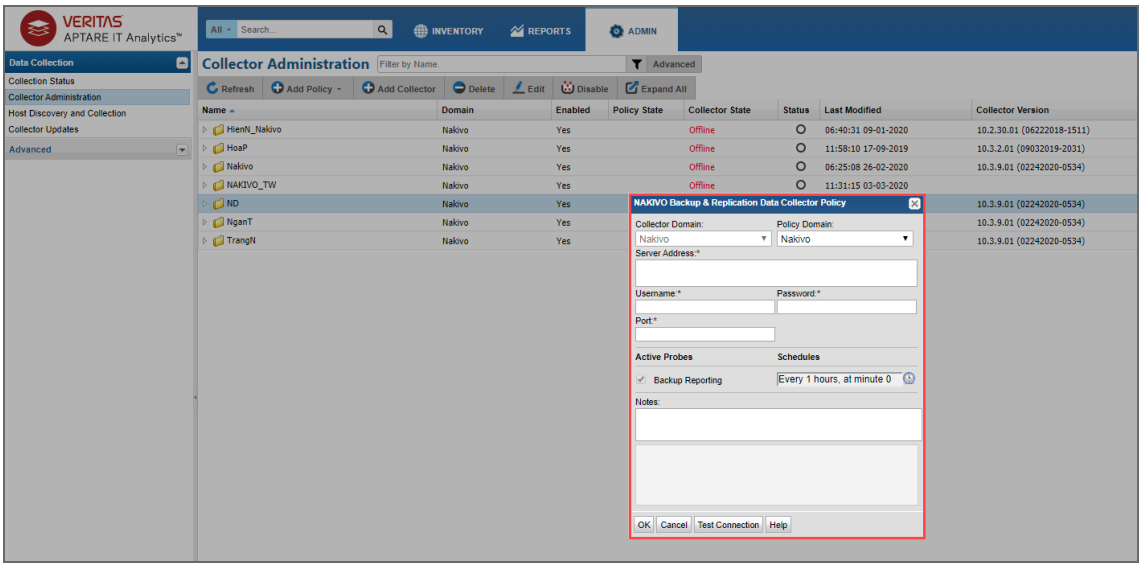



   a. Add a Collector. For details, refer to the Managing and Monitoring Data Collection subsection of the APTARE IT Analytics User Guide.

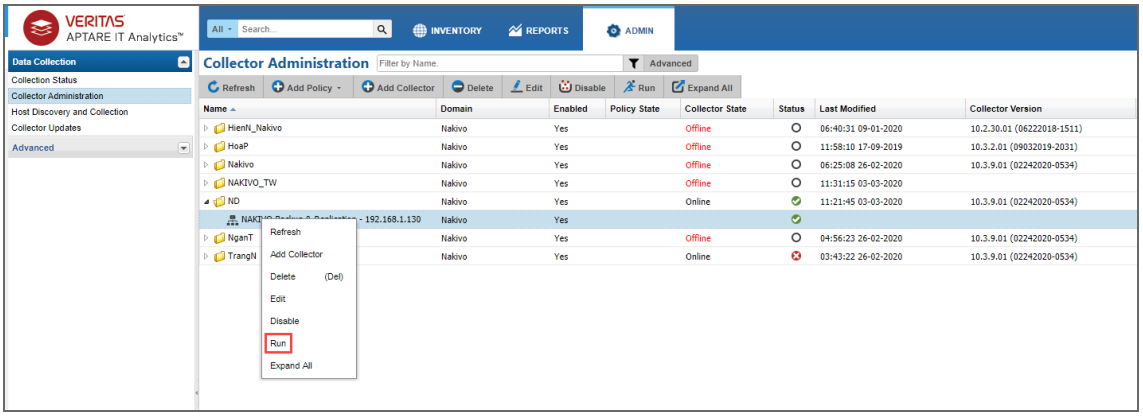



   b. Add a NAKIVO Backup & Replication data protection policy with a connection to your NAKIVO Backup & Replication instance. For details, refer to the Pre-Installation Setup for Generic Backup subsection of the APTARE IT Analytics User Guide.

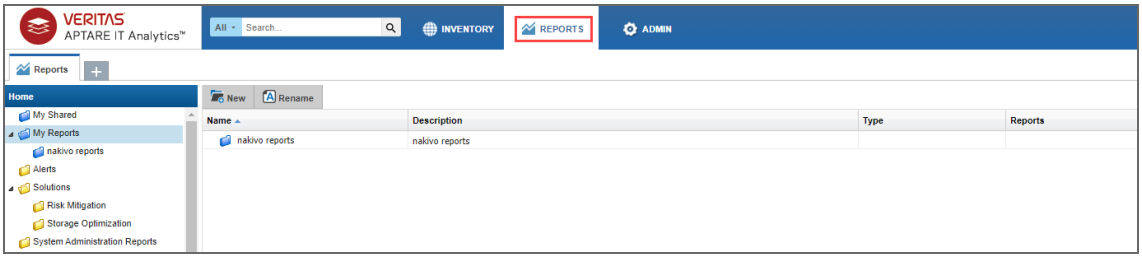



c. Run your policy.

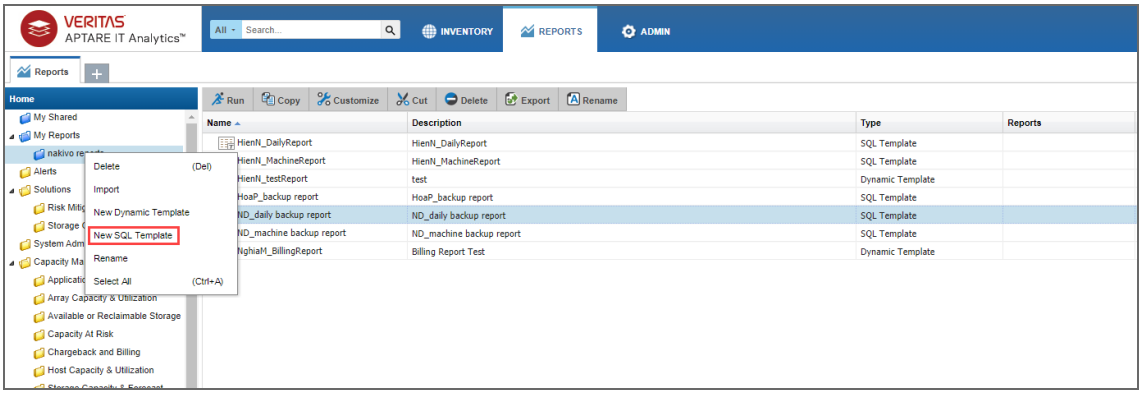



5. Go to the **REPORTS** tab in the APTARE portal and take the following actions:

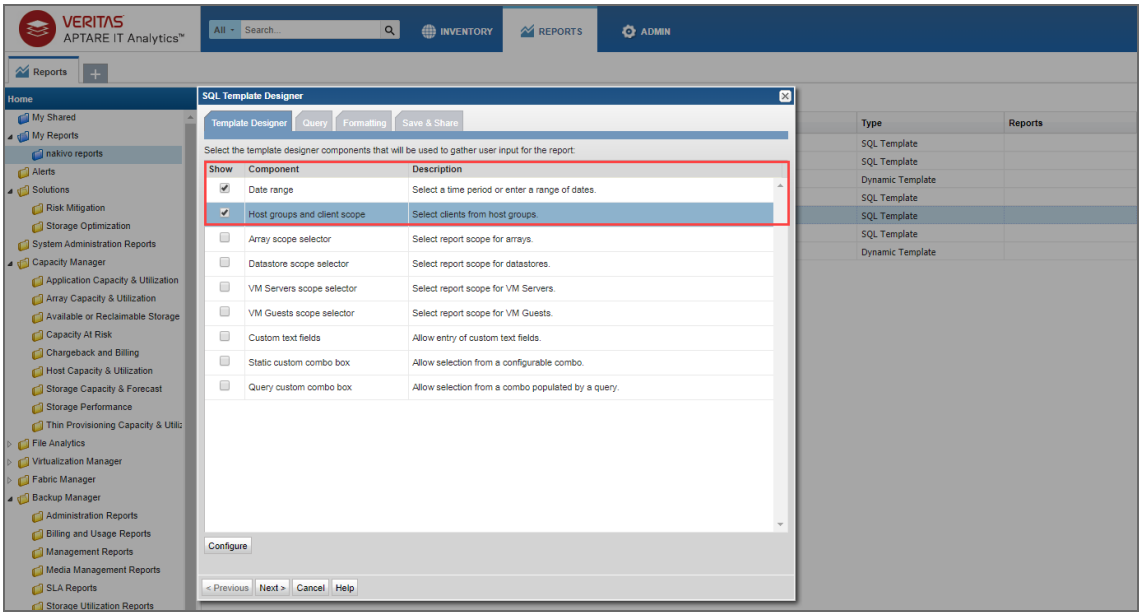



a. Create and configure the report for your backup job the following way:

i. Right-click on your report folder and select **New SQL Template**.




ii. Select the template designer that will be used to gather user input for the report.



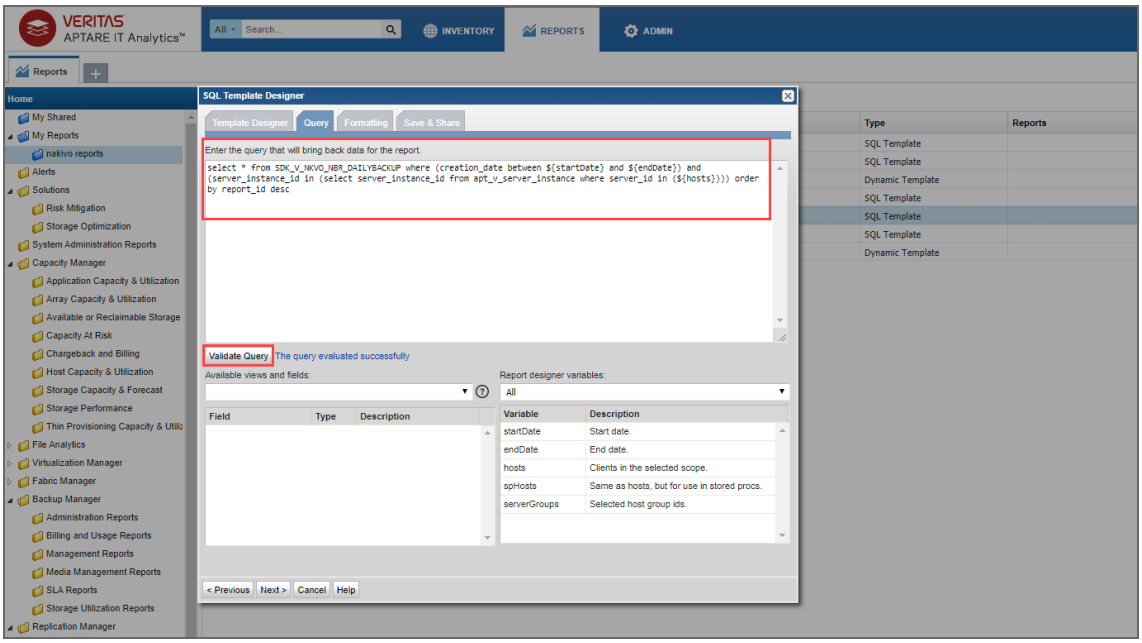

iii. Build an SQL query for your machine backup report or daily backup report, use either:

**Daily Backup Report**

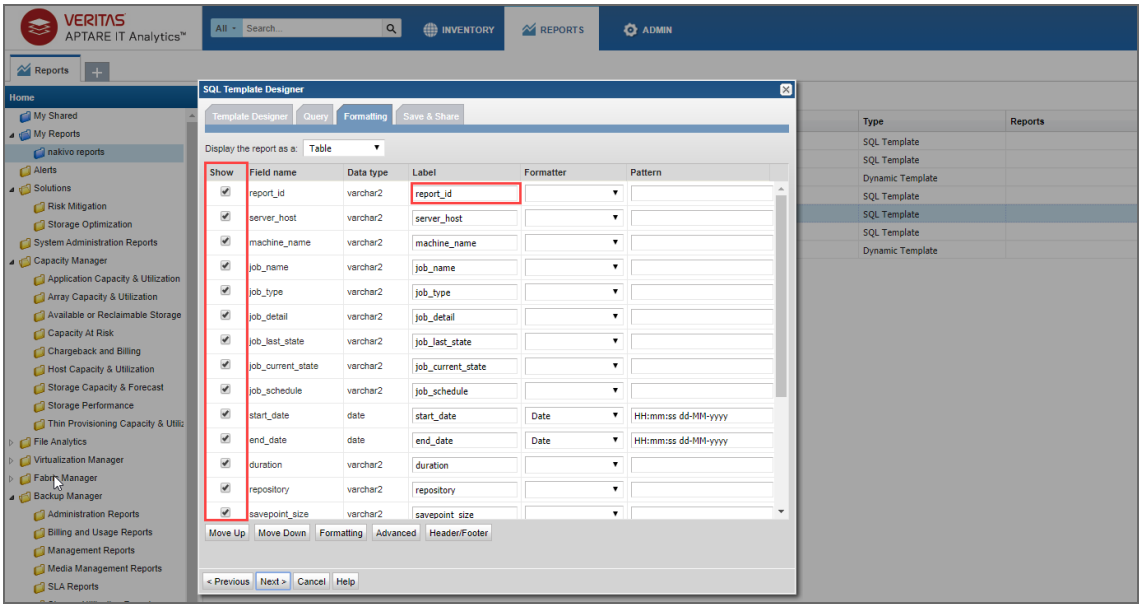
```
select * from SDK_V_NKVO_NBR_DAILYBACKUP where (creation_date
between ${startDate} and ${endDate}) and (server_instance_id
in (select server_instance_id from apt_v_server_instance
where server_id in (${hosts}))) order by report_id desc
```

**Machine Backup Report**

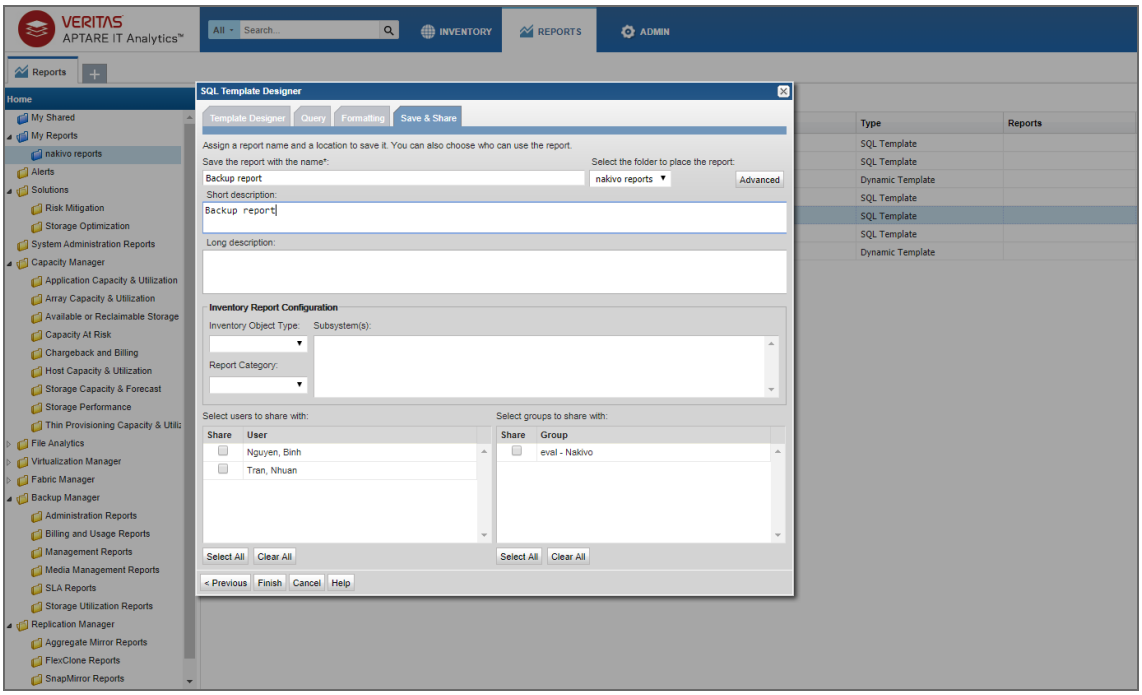
```
select * from SDK_V_NKVO_NBR_MACHINEBACKUP where (creation_
date between ${startDate} and ${endDate}) and (server_
instance_id in (select server_instance_id from apt_v_server_
instance where server_id in (${hosts}))) order by report_id
desc
```



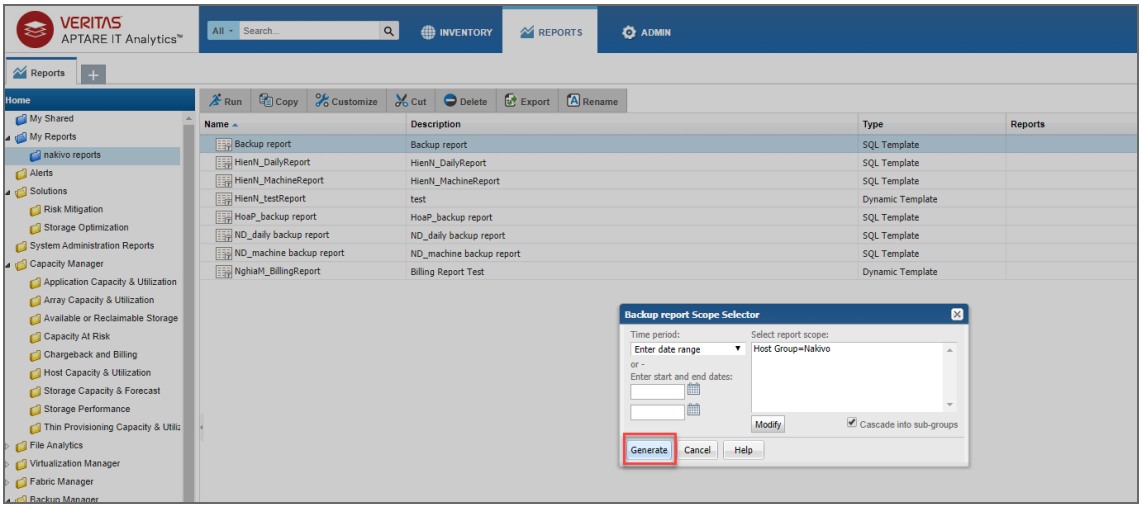

iv. Change formatting options if necessary.




v. Provide a name and description of the report and select users to share it with. Click **Finish**.




b. Double-click on the report. In the dialog box that opens, enter the necessary time and report scope. Click **Generate** to generate your report. For details, refer to the Generating and Maintaining Reports subsection of the APTARE IT Analytics User Guide.




To know more about APTARE IT Analytics, refer to the APTARE IT Analytics User Guide.

# Multi-Tenant Mode

This section covers the following topics:

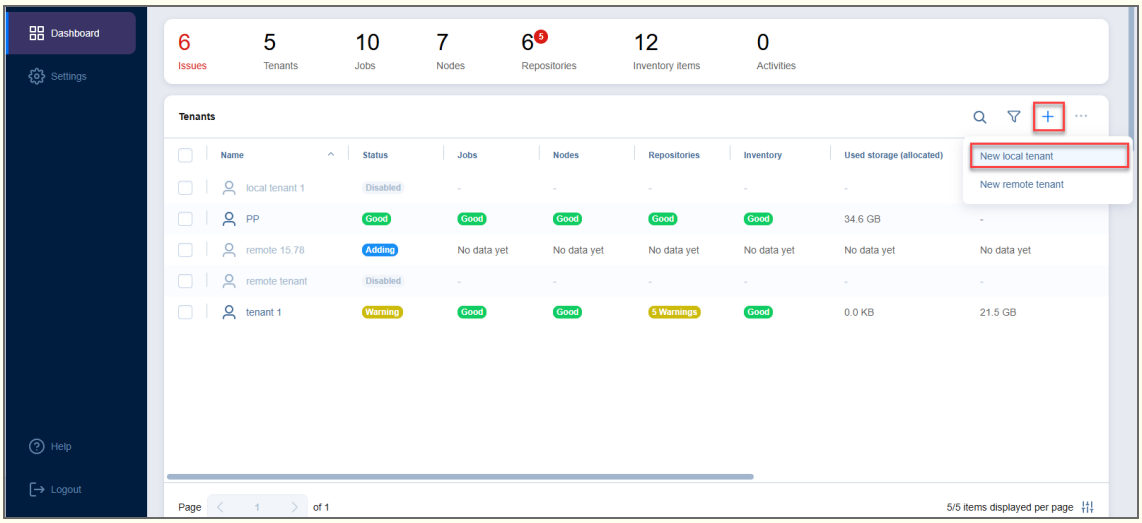

- "Creating a Local Tenant" below
- "Creating a Remote Tenant" on page 1261
- "Using the MSP Console" on page 1265
- "Using the MSP Dashboard" on page 1278
- "Granting Self-Service Access" on page 1289


# Creating a Local Tenant

This section covers the topics describing the local tenant creation process in NAKIVO Backup & Replication. The data protection resources (Inventory items, Backup Repositories, and Nodes) of a local tenant account can only be added and edited by the Master tenant.

**Notes**

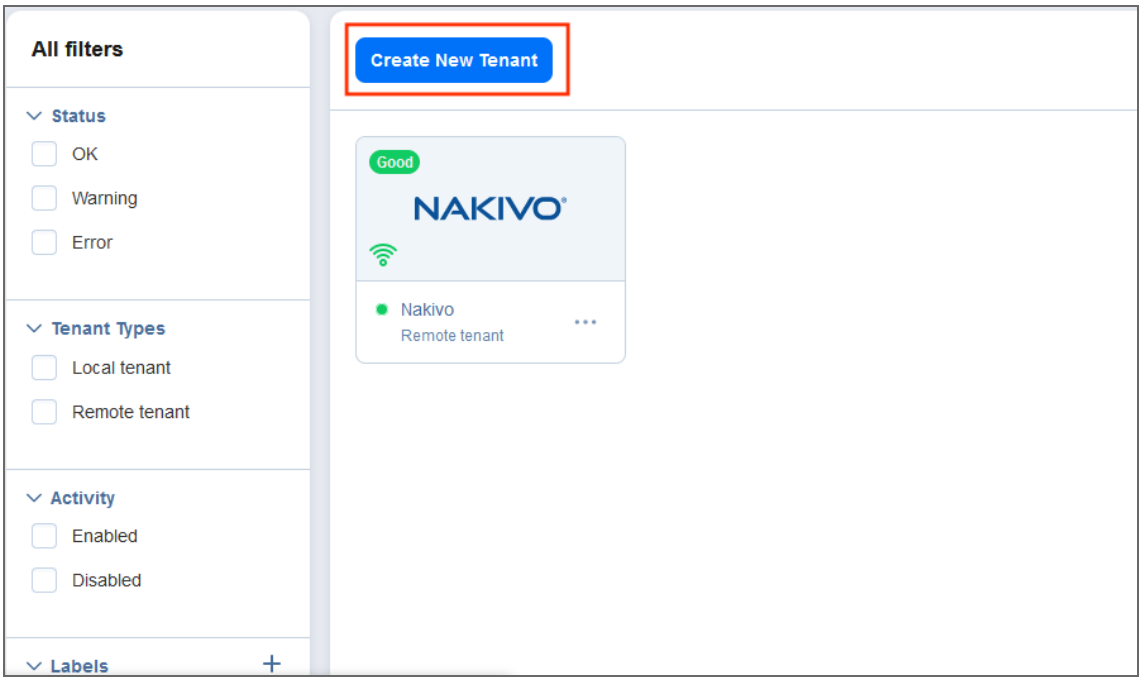
- You can create new local tenants by using the **MSP Console**. To learn how to do it, check here.




- The **MSP Console** is only available for users with an *MSP license*, *Beta instance*, *Promo license*, or *Trial license*. Users with multi-tenant instances of NAKIVO Backup & Replication without these license types have access to the MSP Dashboard.

To create a new local tenant, follow the steps below:

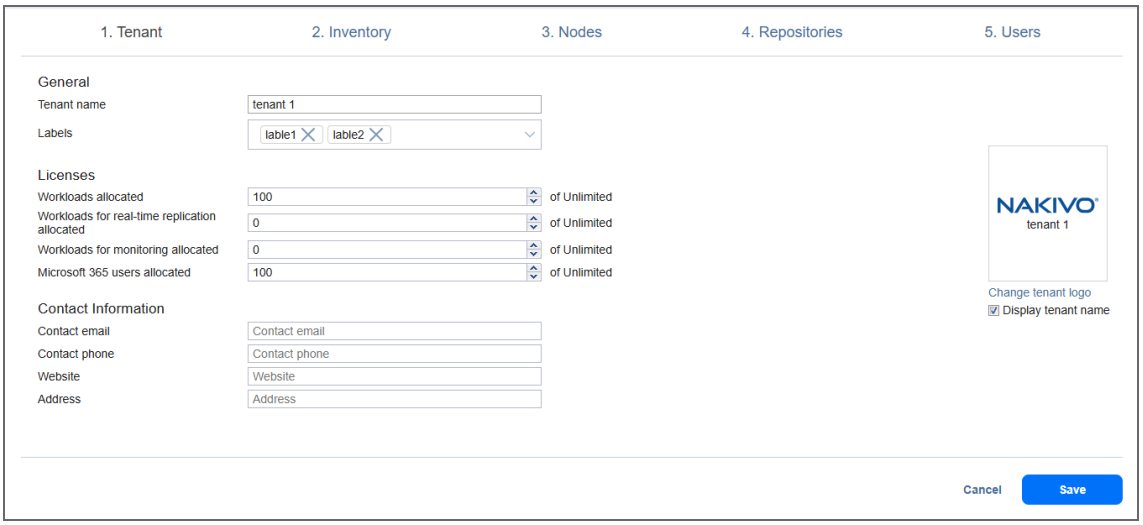
1. Log in to NAKIVO Backup & Replication as a Master Admin.
2. Click **Create New Tenant**.




3. In the popup, select **New local tenant**.
4. Complete the wizard as described in the topics below to finish the tenant creation process.
    - “Local Tenant Creation Wizard: Tenant” below
    - “Local Tenant Creation Wizard: Inventory” on page 1257
    - “Local Tenant Creation Wizard: Nodes” on page 1258
    - “Local Tenant Creation Wizard: Repositories” on page 1259
    - “Local Tenant Creation Wizard: Users” on page 1260
    - “Local Tenant Configuration” on page 1261

# Local Tenant Creation Wizard: Tenant

On this page of the wizard, you can provide a name for the local tenant, assign licenses to the local tenant, and enter or edit contact information for the local tenant. Additionally, the master tenant can enable VM Limitation for the new local tenant. When this option is enabled, the tenant cannot exceed the number of allocated VMs for the purpose of backup and replication. Tenants can see the number of allocated and used VMs in the licensing tab and in the job creation wizard.




Proceed as follows:

1. To add a tenant logo, click **Change tenant logo**, navigate to a new image, select it, and click **Open**. The uploaded image is resized and displayed on the right side of the page.
2. In the **Tenant name** field, enter a name for the local tenant. By default, the tenant name is displayed under the tenant logo. If you do not want the tenant name to be displayed, deselect the **Display tenant name** checkbox.
3. Optionally, in the **Labels** field, select the tags you want to assign to the tenant. Additionally, you can enter the name of the new label in the field and click **Create new label** to create and add it to the **Labels** field automatically.
4. In case the Trial or Subscription license is installed, do the following:
    a. In the **Workloads allocated** field, enter the number of workloads you want to assign to the local tenant.
    b. In the **Microsoft 365 users allocated** field, enter the number of Microsoft 365 users you want to assign to the local tenant.
5. In case a Perpetual license is installed, do the following:
    a. In the **Sockets allocated** field, enter the number of sockets you want to assign to the local tenant.
        i. Optionally, enable the **Limit number of protected VMs** option.

        > **Note**
        > In case the option is not available, make sure that the feature requirements are met.

        ii. Enter the number of protected VMs for the tenant.

        > **Note**
        > Even with VM limitation enabled, licenses are counted on a per-socket basis.

b. In the **Physical servers allocated** field, enter the number of physical server licenses you want to assign to the local tenant.
c. In the **Physical workstations allocated** field, enter the number of physical workstation licenses you want to assign to the local tenant.
d. In the **Microsoft 365 users** allocated field, enter the number of Microsoft 365 users you want to assign to the local tenant.
e. In the **Oracle databases** allocated field, enter the number of Oracle Database licenses you want to assign to the local tenant.

6. Optionally, in the **Contact email** field, enter the email address of the local tenant.
7. Optionally, in the **Contact phone** field, enter the phone number of the local tenant.
8. Optionally, in the **Website field**, enter the website URL of the local tenant.
9. Optionally, in the **Address** field, enter the address of the local tenant.
10. Click **Next** to proceed to the **Inventory** page or **Save** to save the changes.

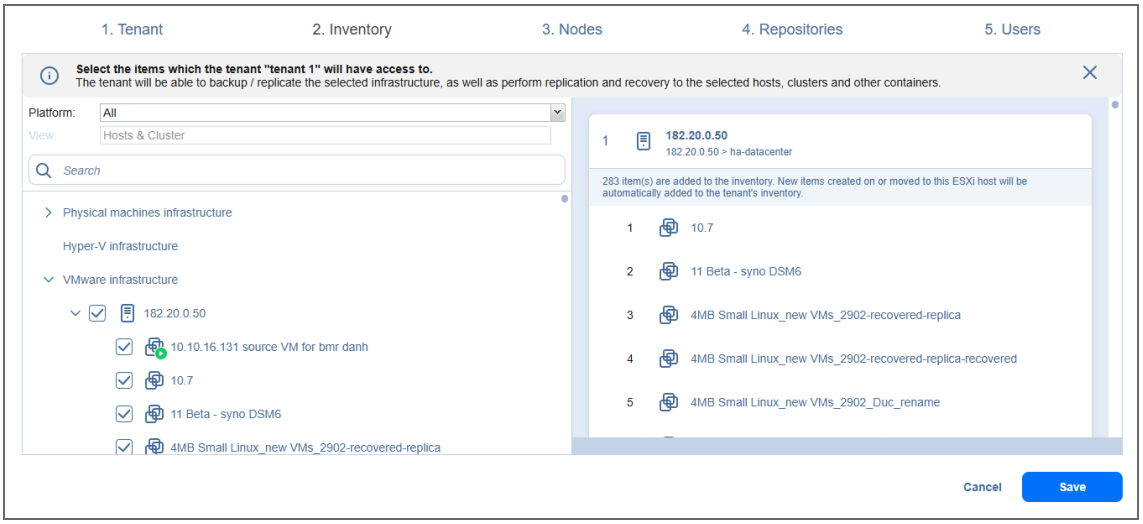
# Local Tenant Creation Wizard: Inventory

On this page, you can assign inventory items to the local tenant or edit the items assigned to this tenant. Proceed as follows:

1. Choose the platform to display the items added to the inventory. **All** is selected by default.

> **Note**
> Items that are assigned to other tenants are visible, but cannot be selected.

2. Optionally, you can filter the Inventory tree by entering a string into the **Search** box. You can enter either a part or the entire name of the item.
3. Select the items you want to assign to the local tenant. The selected items appear in the right pane.




4. Click **Next** to proceed to the **Transporters** page or **Save** to save the changes.

# Local Tenant Creation Wizard: Nodes

On this page of the wizard, you can assign/edit the **Transporters** that the local tenant will be able to use for backup, recovery, and replication jobs. Proceed as follows:

1. In the **Search** field, you can enter the name or part of the name of the **Transporter** to find the specific ones you need.

   **Notes**

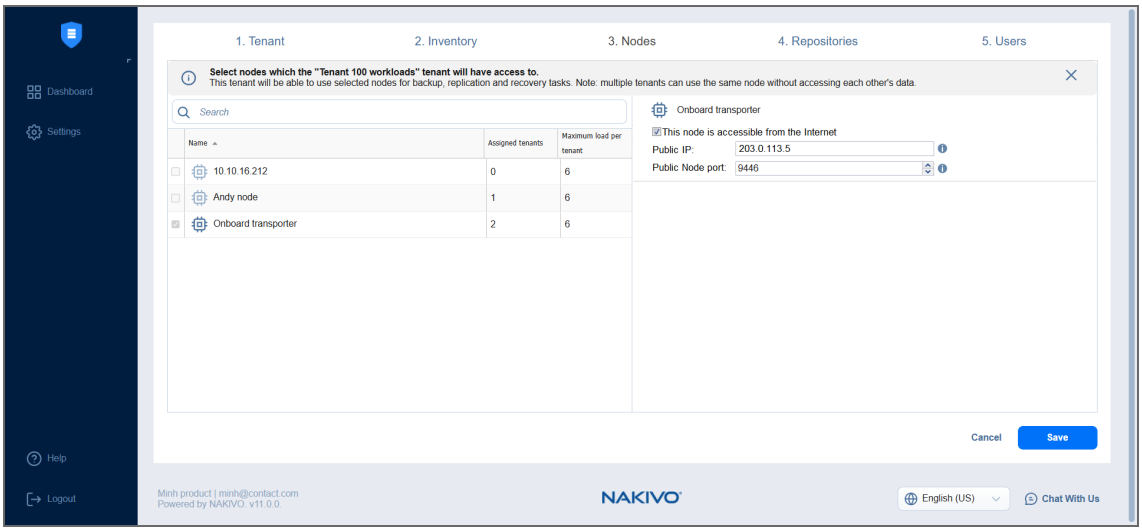
   - When you assign an Inventory item with a dependent **Transporter** to the local tenant on the Inventory page of the wizard, that Transporter is selected automatically and cannot be deselected. If an Inventory item with a dependent **Transporter** was not assigned to the local tenant, that **Transporter** cannot be selected on this page.
   - The **Transporter** deployed in the virtual appliance cannot be assigned to multiple tenants.

2. On the left pane of the screen, you can select the **Transporters** to be assigned to the tenant. The following information is available:
   - **Name**: Name of the **Transporter**.
   - **Assigned tenants**: The number of tenants assigned to the **Transporter**. Multiple tenants can use the same **Transporter** without accessing each others data.
   - **Maximum load per tenant**: The maximum number of tasks that the **Transporter** can perform at the same time per each assigned tenant.
3. The selected **Transporters** appear in the right pane of the screen, where you can edit the following **Transporter** details:
   - **This node is accessible from the Internet**: Select this checkbox to enable **Public IP** and **Public Node** port fields and establish **Direct Connect** communication via specified IP address and port accessible from the internet.
   - **Public IP**: The external **MSP Transporter** IP address reachable from the internet and used for communication with the **Direct Connect Transporter**.
   - **Public Node port**: The external **MSP Transporter** port reachable from the internet and used for communication with the **Direct Connect Transporter**. Ensure that the entered combination of the port and IP address is not used by other transporters.

     **Note**

     These fields are displayed for each **MSP Transporter** with enabled **Direct Connect** option.

4. Click **Next** to proceed to the **Repositories** page or **Save** to save the changes.




# Local Tenant Creation Wizard: Repositories

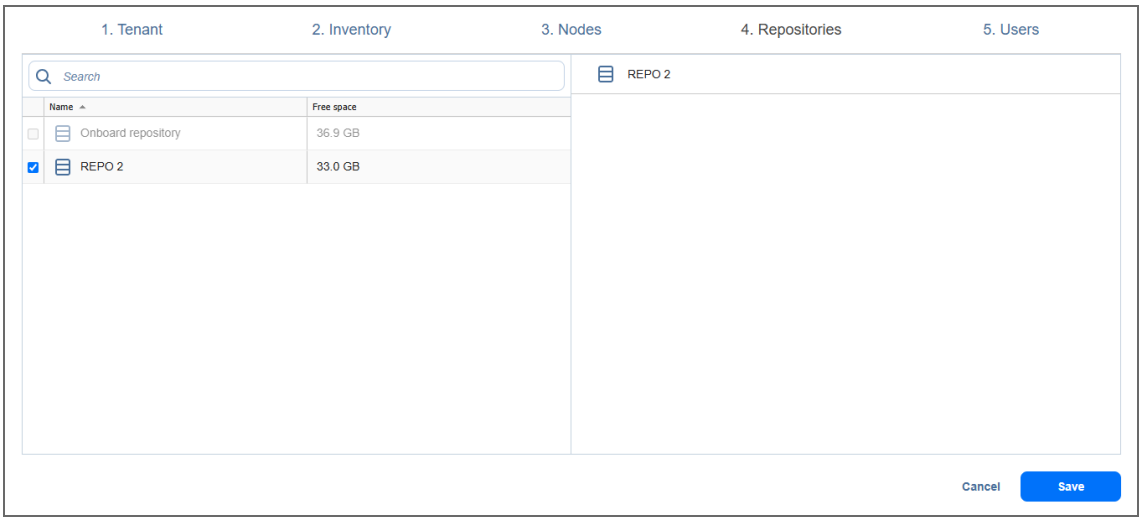
On this page of the wizard, you can assign/edit Backup Repositories that the local tenant will be able to use for backup, recovery, and replication jobs. Note that a single repository cannot be used by multiple tenants. Proceed as follows:

1. In the **Search** field, you can enter either a part or the entire name of the Backup Repository to find the specific ones you need.

> **Notes**
>
> - If the dependent Transporter was not chosen on the **Nodes** page of the wizard, the Backup Repositories assigned to this Transporter would not be available for selection.
> - A federated repository can be assigned to a tenant. A federated repository member cannot be assigned.

2. On the left pane of the screen, you can select the Backup Repositories to be assigned to the local tenant. The following information is available
   - **Name**: Name of the Backup Repository.
   - **Free space**: The amount of free space available on the Backup Repository.

The selected Backup Repositories appear in the right pane.

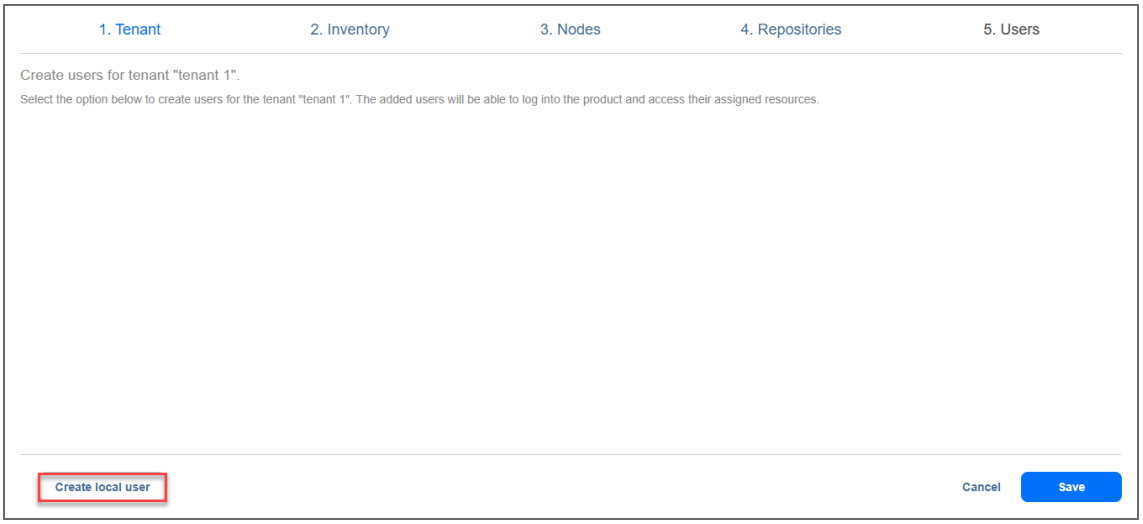



3. Click **Next** to proceed to the next page of the wizard or **Save** to save the changes.

# Local Tenant Creation Wizard: Users

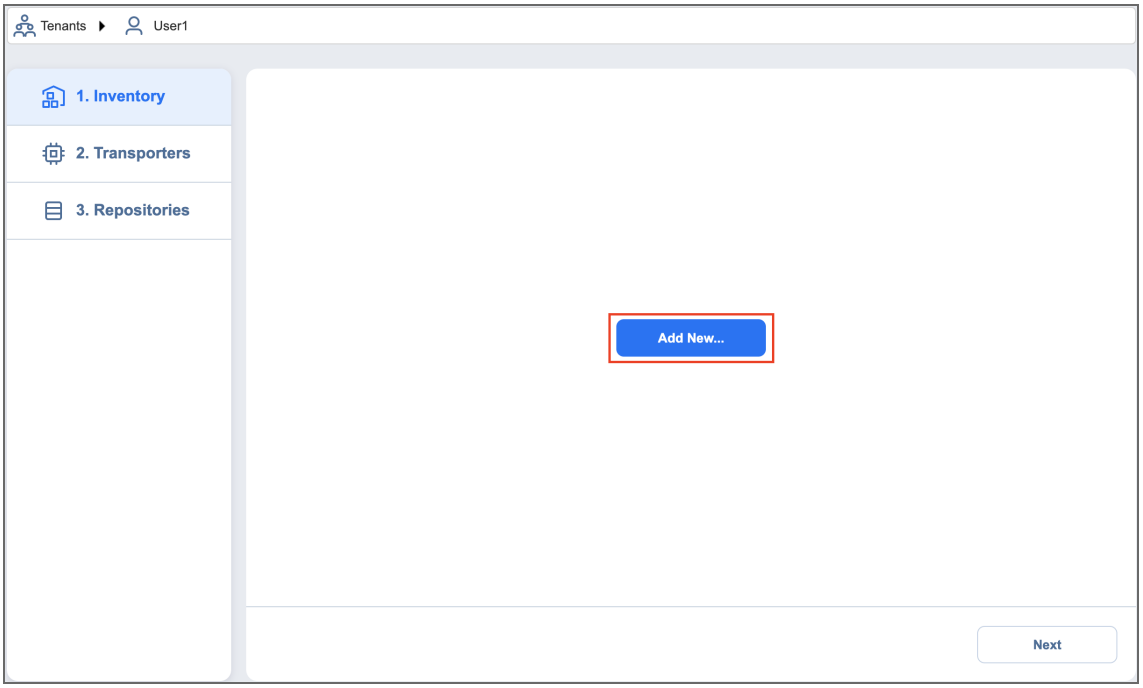
On this page of the wizard, you can create/edit local users or import Active Directory users for the tenant. The added users can use the product and have access to the allocated resources. Do the following:

1. In the lower-left pane of the screen, click **Create local user** to create a new local user for the tenant.
2. If you have successfully configured AD integration, you can click **Add AD user** to import AD user for the tenant.
3. Once you're done, click **Finish** to complete the Local Tenant Creation Wizard or **Save** to save the changes.



## Local Tenant Configuration

After creating a new tenant, click the tenant to open the initial Tenant Configuration Wizard which will guide you through the tenant setup process. Refer to “First Steps with NAKIVO Backup & Replication” on page 341 for a description of the initial configuration wizard.

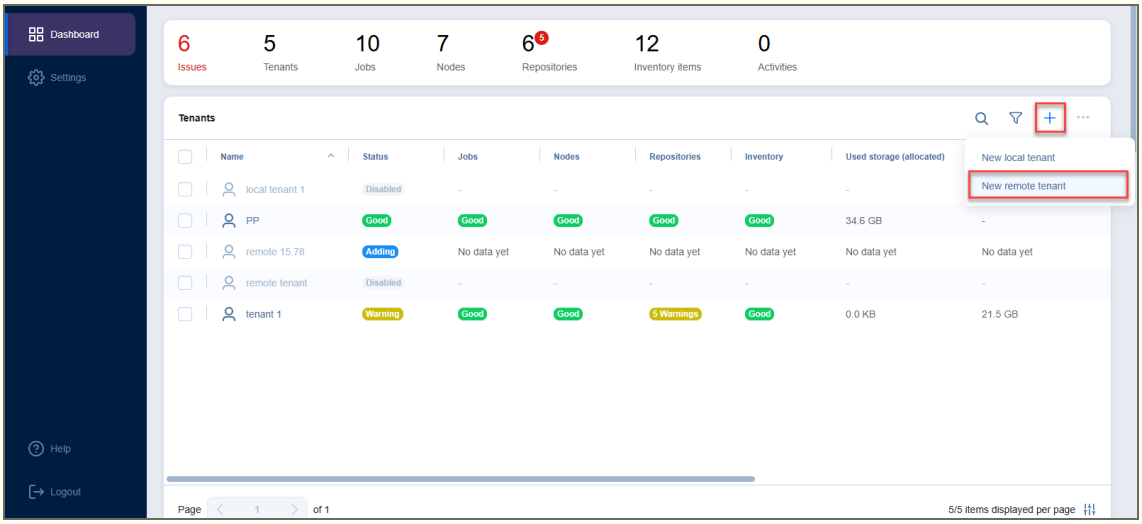



# Creating a Remote Tenant

This section covers the topics describing the remote tenant creation process. Creating and configuring a Remote tenant allows a Master tenant to monitor a standalone instance of NAKIVO Backup & Replication. The remote tenant account retains the ability to manage the resources in the client’s data protection infrastructure.

**Notes**

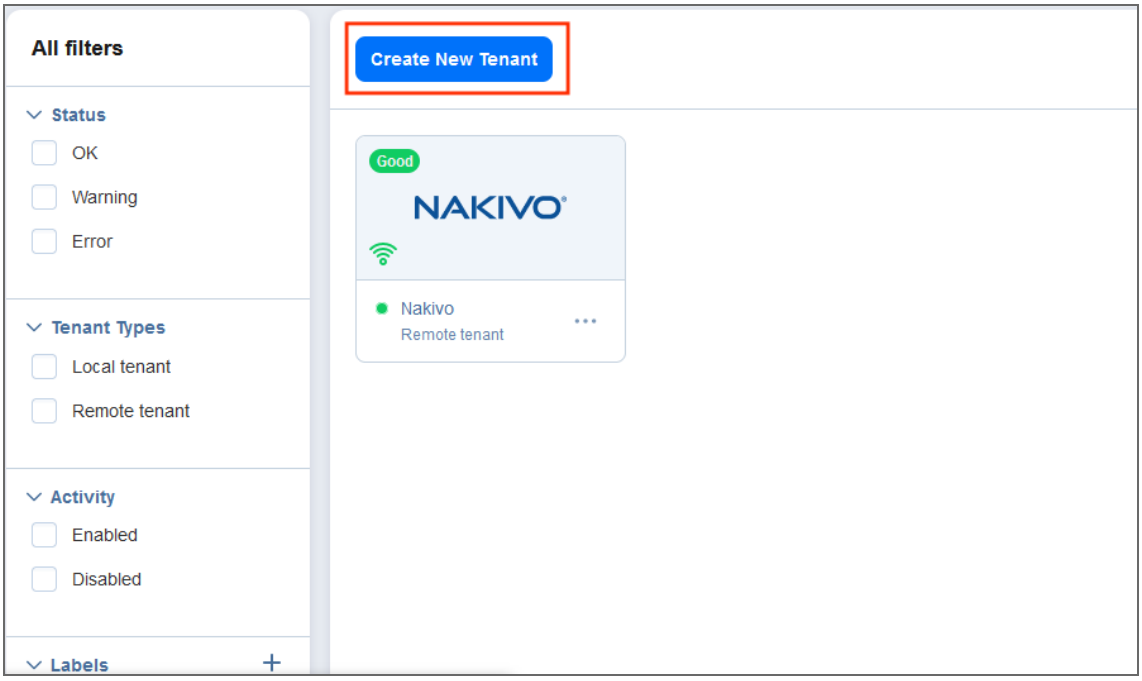
- You can create new remote tenants by using the **MSP Console**. To learn how to do it, refer to Creating New Tenants.




- The **MSP Console** is only available for users with an *MSP license*, *Beta instance*, *Promo license*, or *Trial license*. Users with multi-tenant instances of NAKIVO Backup & Replication without these license types have access to the MSP Dashboard.

To create a new remote tenant, follow the steps below:

1. Log in to NAKIVO Backup & Replication as a Master Admin.
2. Click **Create New Tenant**.




3. In the popup, select **New remote tenant**.
4. Complete the remote tenant creation process as described in the topics below:

- “Remote Tenant Creation Wizard: Tenant” below
- “Remote Tenant Creation Wizard: User” below
- “Remote Tenant Configuration” on the next page

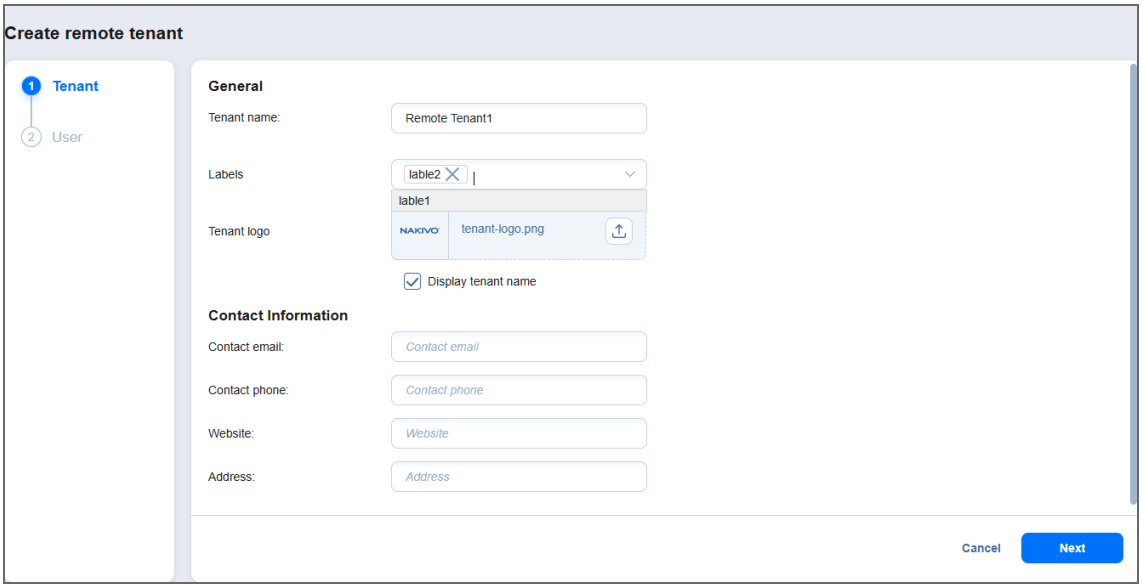
# Remote Tenant Creation Wizard: Tenant

Complete the **Tenant** section of the **Remote Tenant Creation** wizard by configuring the following fields:

1. **Tenant name**: Specify the name of the remote tenant.
2. **Labels**: Optionally, you can create a new tag or assign existing tags to the remote tenant using the drop-down menu.
3. **Tenant logo**: Upload a logo to be displayed for the remote tenant in the Multi-tenancy Dashboard. The photo is automatically resized and a preview is generated.
4. **Display tenant name**: Enable this option if you want the tenant name to be displayed in the **Master Tenant Dashboard**.

Optionally, add contact information for the remote tenant by filling in the following fields:

- **Contact email**
- **Contact phone**
- **Website**
- **Address**

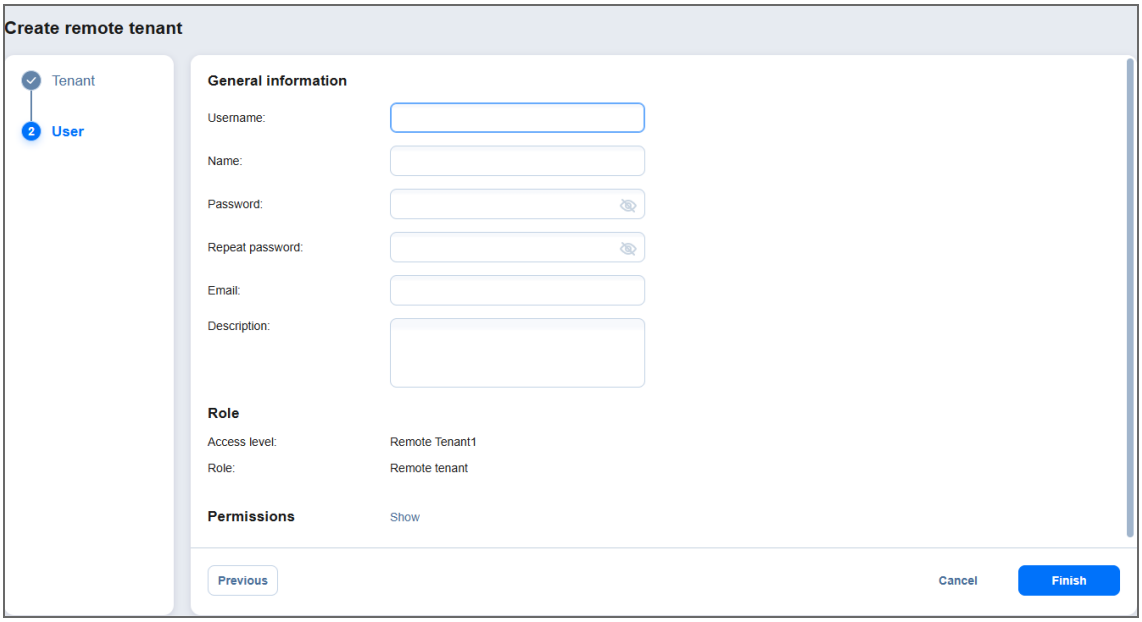



When you’re done, click **Next** to move to the next page of the wizard.

# Remote Tenant Creation Wizard: User

Complete the **User** section of the **Remote Tenant Creation** wizard by configuring the following fields:

1. **Username**: Specify a username for the remote tenant **User**.
2. **Name**: Specify the remote tenant display name.
3. **Password**: Create a password for this user and repeat it in the **Repeat Password** field below.
4. **Email**: Enter the user's email address.
5. **Description**: Optionally, you can add a description for this **User**.




The **Remote tenant** role and its permissions are added to the **User** automatically. Click **Finish** to confirm the creation of the remote tenant.

# Remote Tenant Configuration

When a remote tenant is created, it is automatically added to the list of tenants in the **Master Tenant Dashboard**. To connect a remote tenant to your multi-tenant installation of NAKIVO Backup & Replication, follow the steps below:

1. Provide the remote tenant with the credentials created for the remote tenant user, as well as your hostname/IP address and **Director** port number (*4443* by default). In addition, you will need to open a separate listening port for communication with the remote tenant's instance (port *6702* is used by default). For more information on the required TCP ports, see the **MSP Console** section in "Feature Requirements" on page 169.
2. The remote tenant must go to **Settings** > **MSP** in their own instance of NAKIVO Backup & Replication and add the MSP using the above information, and then click **Add**.
3. A popup with certificate details appears. The remote tenant should click **Apply** to add the MSP to the **MSP** tab.
4. In your **Master Tenant Dashboard**, the remote tenant should now have a green **Connected** icon on the tenant card. Clicking on the remote tenant's name allows you to drill down and monitor their instance.

For more information on tenant-side **MSP Console** configuration, refer to “Adding an MSP” on page 437.

# Using the MSP Console

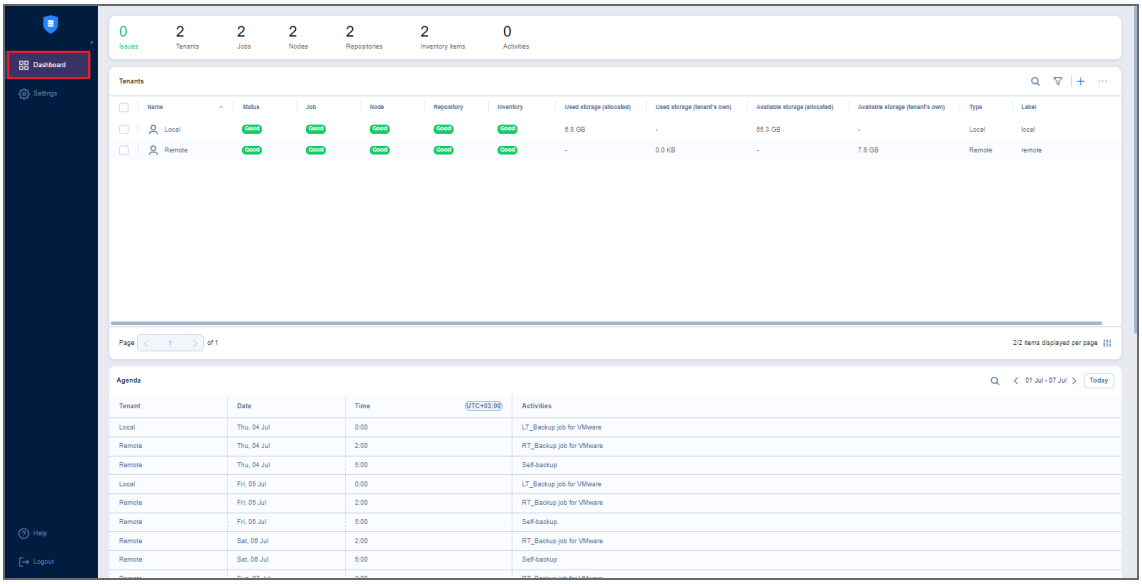
The **MSP Console** feature allows users to connect their standalone NAKIVO Backup & Replication instance to a managed service provider (MSP) and vice versa. Conversely, using the **MSP Console**, MSPs can create, add, and connect to local and remote tenants’ environments to monitor a remote tenant's instance of NAKIVO Backup & Replication once a connection has been established on both sides.

**Important**

The **MSP Console** is only available for users with an *MSP license*, *Beta instance*, *Promo license*, or *Trial license*. Users with multi-tenant instances of NAKIVO Backup & Replication without these license types have access to the MSP Dashboard.

The **MSP Console** displays the key statistics for each local and remote tenant in the MSP’s instance of NAKIVO Backup & Replication.

Using the **MSP Console**, MSPs can save time on tenants’ management activities monitoring their statuses without drilling down into each tenant once a connection has been established on both sides.




See the topics below for more information:

- “MSP Architecture” on page 439
- “Opening MSP Console” on the next page
- “Creating New Tenants” on page 1267
- “Viewing Tenants Details” on page 1267
- “Managing Tenants Details” on page 1273
- “Using Search / Filter Options” on page 1276

# Opening MSP Console

As a managed service provider (MSP), you can use the **MSP Console** to monitor and manage all standalone instances of NAKIVO Backup & Replication connected to your instance of the solution and the data protection activities of the tenants you create.

**Note**

To use the MSP Console feature as an MSP (Master tenant), you must have NAKIVO Backup & Replication installed in the multi-tenant mode.

To establish a connection with a standalone instance of the solution, the MSP must first create a remote tenant account for the client and share the credentials with the client. The remote tenant must then use these to connect their instance to the MSP's instance of NAKIVO Backup & Replication.

To use the MSP Console to create and manage local tenants, a managed service provider (MSP) can enable Direct Connect to establish a connection with client remote resources.

**Important**

To use the **MSP Console**, the managed service provider (MSP) needs to configure the following TCP ports:

- **MSP Director port**: This is the TCP port used by the Director for the MSP's instance of NAKIVO Backup & Replication. By default, this is TCP port *4443*. The MSP must provide a remote tenant with the Director port number during configuration. The remote tenant needs to enter this port number when adding the MSP to their standalone instance of NAKIVO Backup & Replication.
- **Listening port**: The MSP must have a port open for listening to the remote tenant. By default, TCP port *6702* is used. The MSP may change the listening port used by changing the ***system.msp.console.listening.port*** parameter in Expert settings.

Once the standalone NAKIVO Backup & Replication instance is added to the **MSP Console**, it appears as a remote tenant on the Tenants table and the Licensing > Tenants tab.

**Note**

Only users with an *MSP license*, *Beta instance*, *Promo license*, or *Trial license* can access the MSP Console and the Licensing > Tenants tab. This allows them to efficiently oversee all independent instances of NAKIVO Backup & Replication associated with a managed service provider (MSP) as well as local tenants from a unified interface, eliminating the need to navigate through individual tenants.

Users with multi-tenant instances of NAKIVO Backup & Replication without these license types only have access to the MSP Dashboard.

# Creating New Tenants

The MSP Console allows MSPs to create, add, and connect to Local and Remote tenant environments.

1. To create a new tenant, in the top right corner of the Tenants table, click **+** (plus) > **Create New Tenant**.
2. In the dialog box that opens, select one of the options: **New local tenant** or **New remote tenant**, then proceed to create a remote or local tenant.




3. Complete the wizard as described in the topics below to finish the tenant creation process.
   - "Local Tenant Creation Wizard: Tenant" on page 1255
   - "Local Tenant Creation Wizard: Inventory" on page 1257
   - "Local Tenant Creation Wizard: Nodes" on page 1258
   - "Local Tenant Creation Wizard: Repositories" on page 1259
   - "Local Tenant Creation Wizard: Users" on page 1260
   - "Local Tenant Configuration" on page 1261

**Notes**

- You cannot assign any resources to a remote tenant.
- All the resources (inventory items, transporters, repositories) in the added single tenant NAKIVO Backup & Replication instance are treated as local tenant resources.
- Added standalone NAKIVO Backup & Replication instances are visible in the Tenants Dashboard as remote tenants. To learn how to create new remote tenants by using the **Tenants Dashboard**, refer to "Creating a Remote Tenant" on page 1261.

# Viewing Tenants Details

In the MSP Console, tenant data is displayed in the following widgets:

- Summary bar: Displays the number of tenant-related issues (errors and notifications), jobs, transporters, repositories, monitored items, and running activities. This widget includes search and filter functions that simplify finding specific issues by tenant, category, type, or date range.
- Tenants table: Allows MSPs to view the real-time status of all existing tenants, tenant-related jobs, nodes, repositories, inventory items, allocated resources, etc. You can also create new local and remote tenants, enable/disable existing tenants, and edit or delete them.

- Agenda widget: Allows MSPs to view information about the running and scheduled activities of all available tenants. By default, this widget displays the current week.
- Events widget: Lists all events, including errors and/or warnings of all available tenants, and is sorted by date by default.

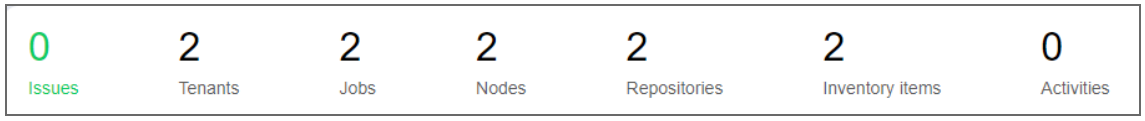
**Notes**

- The Master tenant collects data from tenants every 1 minute and stores this data in the master tenant's database.
- The **Agenda** and **Events** widgets do not collect data from *Disabled*, *Inaccessible*, and *Disconnected* tenants.

See the sections below for more information on monitoring and managing tenants.


- "Summary Bar" below
  - "Alarms" on the next page
- "Tenants Table" on page 1270
- "Agenda" on page 1271
- "Events" on page 1272


## Summary Bar




The **Summary bar** displays the following data collected from tenants: total number of issues (errors and notifications), tenants, jobs, nodes, repositories, monitored items, and running activities.
The data displayed is as follows:

- **Issues**: Total number of alarms/notifications related to all the available tenants (including the Master tenant). Clicking the number displays the Alarms dialog box.

  **Note**

  Alarms and notifications of *Disabled*, *Inaccessible*, and *Disconnected* tenants are not displayed, except for the reasons why the tenants are inaccessible/disconnected.

- **Tenants**: Total number of existing tenants (excluding Master tenant).
- **Jobs**: Total number of jobs inside tenants.
- **Nodes**: Total number of nodes inside tenants (including Master tenant). The nodes include allocated transporters, independently added/deployed transporters inside the tenant, and VMA and PMA.
- **Repositories**: Total number of tenants' repositories (including Master tenant). The repositories include both allocated and independently added/created repositories.

- **Inventory items**: Total number of tenants' inventory (including Master tenant). The inventory includes both allocated inventory and independently added inventory.
- **Activities**: Total number of running activities inside tenants (excluding Master tenant).

**Notes**

- Data is not collected for *Disabled*, *Inaccessible*, and *Disconnected* tenants.
- The small number inside the red circle displays the total number of alarms and notifications related to the corresponding data section: nodes, repositories, or inventory items of all tenants (including the Master tenant).

## Alarms

The **Alarms** dialog box displays all the alarms and notifications relating to existing tenants.

Here, you can browse or search for a specific issue in the **Search** field.

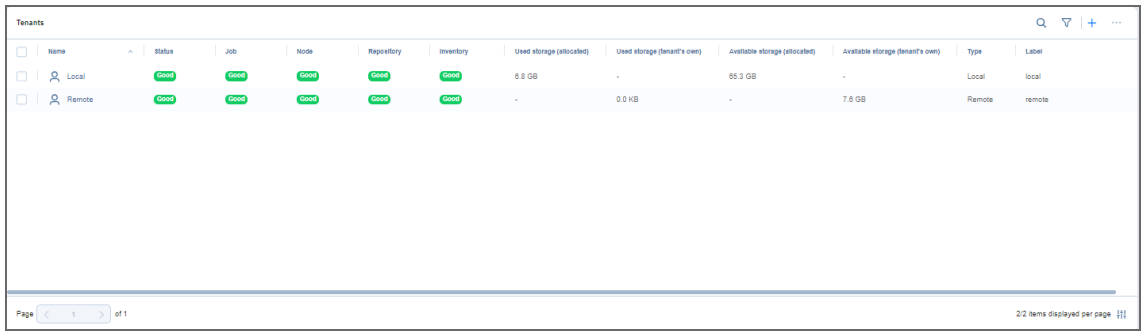
Optionally, click the **Filter** button to display issues based on their characteristics. The following filters are available:

- **Tenant**: Select to filter the alarms by tenant name. Multiple selections are supported.
- **Category:** Click to filter by selecting the following options inside the dropdown:
  - **Job**
  - **Inventory**
  - **Node**
  - **Repository**
  - **Tape**
  - **Monitoring**
  - **Users & Roles**
  - **Auto Update**
  - **Licensing**
- **Type**: Select to filter the alarms and notifications by type. The following types are available:
  - **Warning**
  - **Error**
  - **Dismissible**
  - **Non-dismissible**
- **Date**: Click to filter by date.

Once the filtering options are set, click **Apply** to start filtering.

Click the **Cancel** or **X** button to close the dialog box without applying any changes.

# Tenants Table




The **Tenants** table gives MSPs an overview of all local and remote tenants:

- **Name**: The icon (**active** or **disabled**) and the name of a tenant. Clicking the tenant's name opens the Overview dashboard of the corresponding tenant.

  > **Note**
  > *Disabled*, *Inaccessible*, and *Disconnected* tenants are not clickable.

- **Status**: The status of a tenant:
  - **Adding**: A remote tenant has been added to the MSP console but the first data collection has not been completed yet.
  - **Good**: The tenant is functional.
  - **Warning**: The tenant has notifications only.
  - **Error**: The tenant has alarms.
  - **Inaccessible**: The tenant is not functional.
  - **Disconnected**: The remote tenant is disconnected from the MSP. You cannot drill down to a disconnected tenant.
  - **Disabled**: The tenant is disabled. Disabled tenants still store all the data but cannot log in and perform any activities.

  > **Note**
  > If a tenant has multiple statuses, the priorities of displaying these statuses are: *Disabled* > *Inaccessible/Disconnected* > *Good/Warning/Error*.

- **Job**: The status of a tenant's jobs. Clicking the status indicator opens the **Job Overview Dashboard** of the corresponding tenant. *Adding*, *Disabled*, *Inaccessible*, and *Disconnected* tenants are not clickable.
- **Node**: The status of the tenant's nodes. Clicking the status indicator opens the **Nodes** tab of the corresponding tenant. *Adding*, *Disabled*, *Inaccessible*, and *Disconnected* tenants are not clickable.

- **Repository**: The status of a tenant's repositories. Clicking the status indicator opens the **Repositories** tab of the corresponding tenant . *Adding*, *Disabled*, *Inaccessible*, and *Disconnected* tenants are not clickable.
- **Inventory**: The status of the tenant's inventory. Clicking the status indicator opens the **Inventory** tab of the corresponding tenant . *Adding*, *Disabled*, *Inaccessible*, and *Disconnected* tenants are not clickable.
- **Used storage (allocated)**: The amount of allocated storage by the Master tenant that has been used up by the tenant. For *Adding* tenants, **No data yet** is displayed.
- **Used storage (tenant's own)**: The amount of the tenant's own storage that has been used by the tenant. For *Adding* tenants, **No data yet** is displayed.
- **Available storage (allocated)**: The amount of allocated storage by the Master tenant that is available to the tenant. For *Adding* tenants, **No data yet** is displayed.
- **Available storage (tenant's own)**: The amount of the tenant's own storage that is currently available. For *Adding* tenants, **No data yet** is displayed.

> **Note**
> If the repositories are placed on the same same disk/share/etc., the free space metric can be incorrect.

- **Label**: The tenant's labels.
- **Type**: The tenant's connection type – **Local** or **Remote**.

## Agenda

The **Agenda** widget allows MSPs to view information about all the running and scheduled activities of tenants (except for *Disabled*, *Inaccessible*, and *Disconnected* tenants).

Clicking the tenant name, date, time, or activity name opens the **Calendar dashboard** of the corresponding tenant.

The **Agenda** widget allows you to search for **Tenant** name on the currently displayed week.




# Events

The **Events** widget gives an overview of all **Error** and **Warning** events for all tenants (except for *Disabled*, *Inaccessible*, and *Disconnected* tenants) and provides information about each event in the following columns:

- **Tenant**: Displays information about all tenants and their events. Clicking the tenant name opens the **Events** page of the corresponding tenant.
- **Event name**: Displays the icon and the message of the event.
- **Category**: Displays one of the following tenant event categories:
  - **Job**
  - **Inventory**
  - **Node**
  - **Repository**
  - **Tape**
  - **Monitoring**
  - **User & Roles**
  - **Auto Update**
  - **Licensing**
- **Initiated by**: Displays the name of the tenant that initiated the event.
- **Date**: Displays the date and time of the event.

You can browse or search for a specific event in the **Search** field. Search can be performed on the **Tenant** column only.

You can sort the events by clicking the respective name of the column.

To access filtering options, click the **Filter** icon in the top right corner. In the dialog box that opens, you can select one or several filtering criteria.

The following filtering options are available (multiple selections are supported):

- **Initiated by**: Allows you to filter by tenant name.
- **Event type**: Allows you to filter by the following options:
  - **Info**
  - **Error**
  - **Warning**
  - **Debug**
- **Date**: Click to select the date range.

Optionally, you can show/hide columns or modify the number of items per page in the **Events table configuration** table. In the lower right corner, click the controls icon. In the dialog box that opens, select/deselect checkboxes and click **Apply**.

Click **Reset Settings** to reset the configuration settings to the default.

Click the **Cancel** or **X** button to close the dialog box without applying any changes.

# Managing Tenants Details

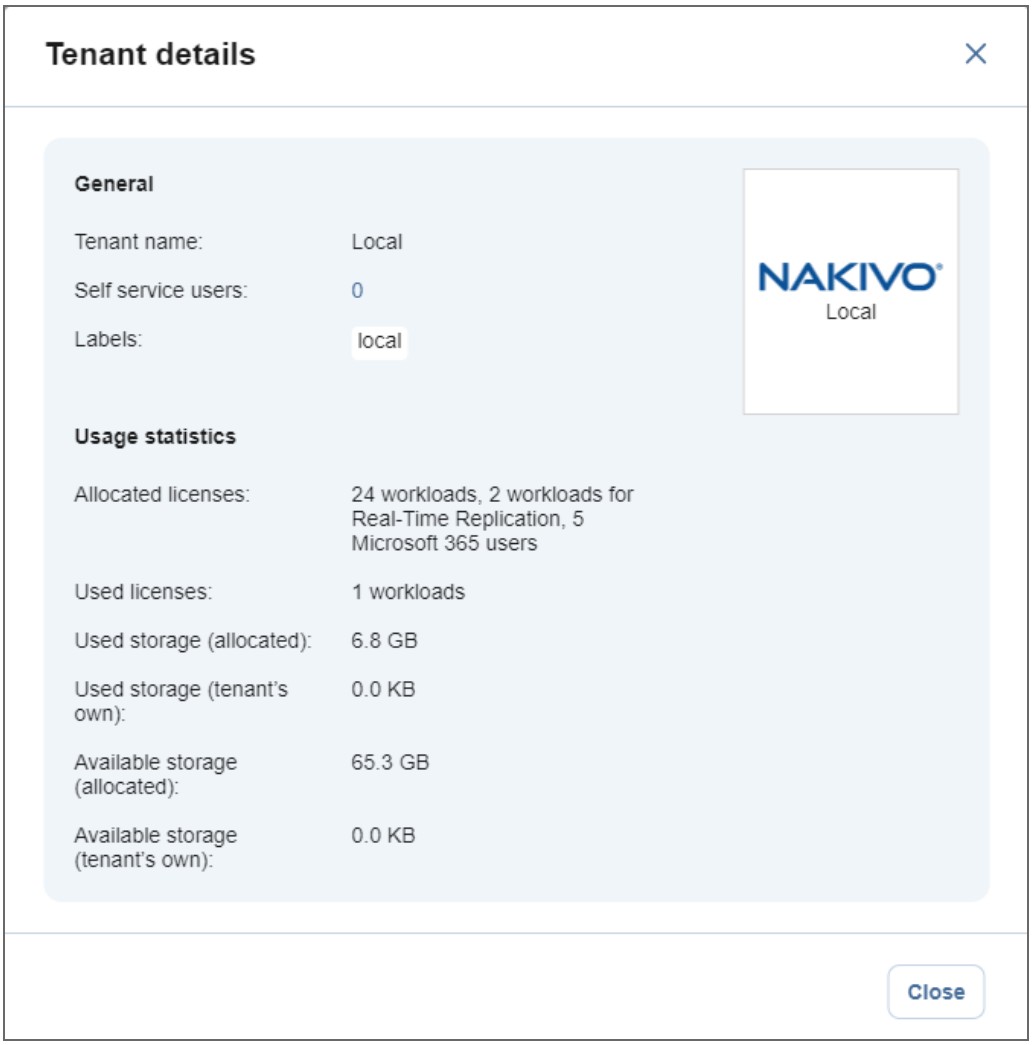
In the Tenants table, you can edit, disable, enable, or delete tenants.

1. To open the **Tenant details** popup window, hover over the corresponding row and click the ellipsis **Manage** button.

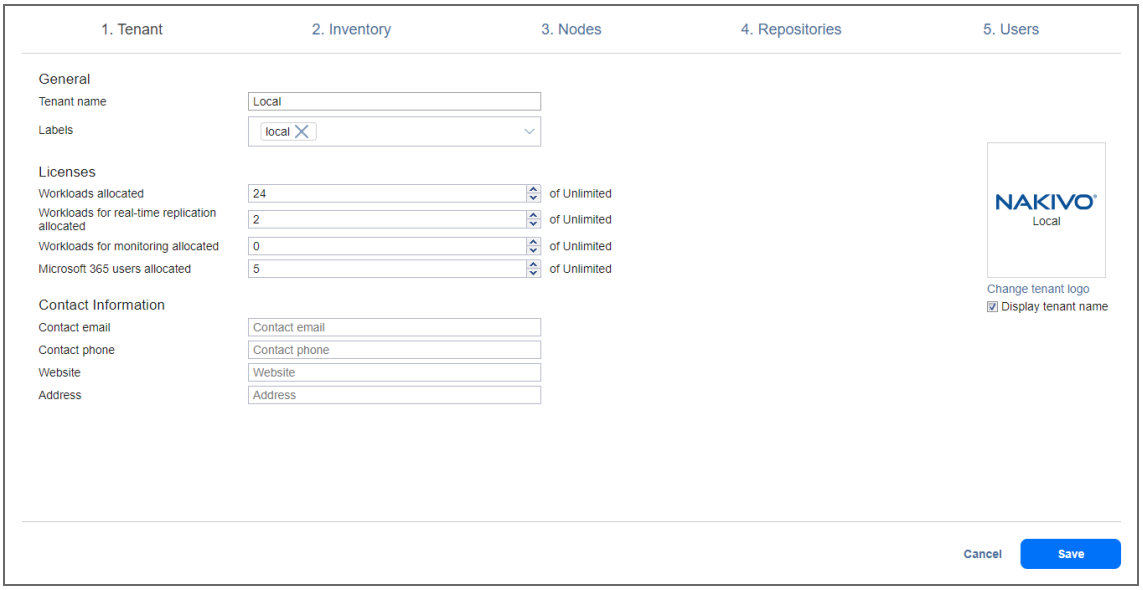



2. Then select the corresponding item in the list of actions:

- **View details**: Click to open the **Tenant details** popup.




- **Edit**: Click to open the tenant editor.




- **Enable/Disable**: Click to enable/disable the selected tenants.

> **Note**
>
> The option is not displayed if the tenant is disabled/enabled.

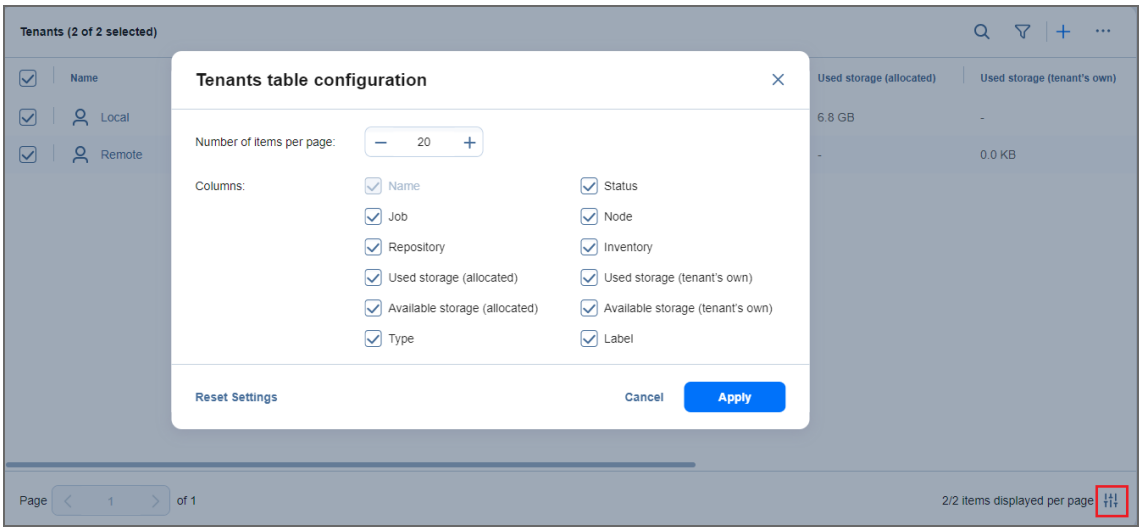
- **Delete**: Click to delete the selected tenants. You can select to delete the tenants and keep the tenant data or completely delete the tenant and all related data.

3. In the confirmation dialog box that opens, confirm or cancel the selected action.

These actions, except **Edit**, can also be done in bulk by selecting at least 2 items in the table or checking the box in the upper left pane to select all users and clicking the **...** (ellipsis) button.




Optionally, you can show/hide columns or modify the number of items per page in the **Tenants table configuration** table. In the lower right corner, click the controls icon. In the dialog box that opens, select/deselect checkboxes and click **Apply**.




Click **Reset Settings** to reset the configuration settings to the default.

Click the **Cancel** or **X** button to close the dialog box without applying any changes.

# Using Search / Filter Options

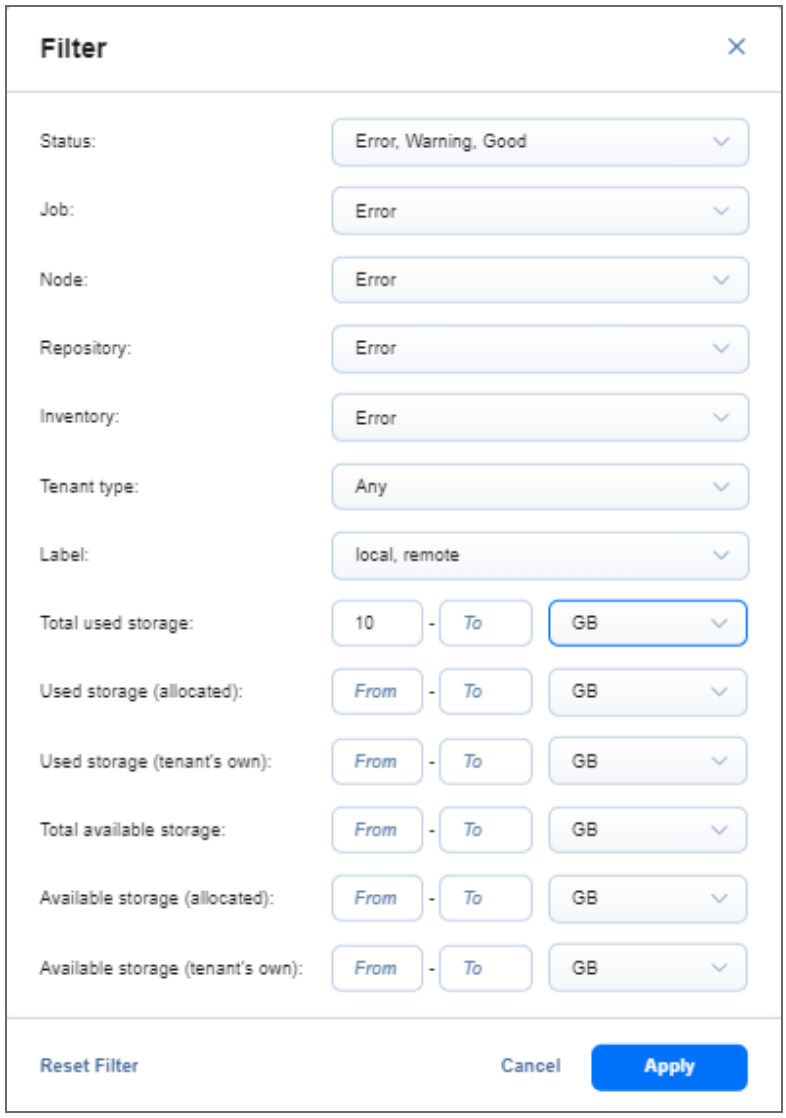
You can search or filter tenants from the Tenants table.

Use the **Search** field to browse or search for a specific tenant. Search can be performed on the **Name** column only.

To access filtering options, click the **Filter** icon in the top right corner. In the dialog box that opens, you can select one or several filtering criteria.

Click **Apply** to filter the search results.

Click the **Cancel** or **X** button to close the dialog box without applying any changes.




The following filtering options are available:

- **Status**: Allows you to filter by the following tenant status options:
  - **Good**
  - **Warning**
  - **Error**

    - **Inaccessible**
    - **Disconnected**
    - **Disabled**
- **Job**: Allows you to filter by the following job options:
    - **Good**
    - **Warning**
    - **Error**
- **Node**: Allows you to filter by the following node options:
    - **Good**
    - **Warning**
    - **Error**
- **Repository**: Allows you to filter by the following repository options:
    - **Good**
    - **Warning**
    - **Error**
- **Inventory**: Allows you to filter by the following inventory options:
    - **Good**
    - **Warning**
    - **Error**
- **Tenant type**: Allows you to filter by the following tenant type options:
    - **Any** (default)
    - **Local**
    - **Remote**
- **Label**: Allows you to filter by the following label options:
    - **Good**
    - **Warning**
    - **Error**
- **Total used storage**: From the dropdown, select to filter the items displayed by:
    - **MB**
    - **GB** (default)
    - **TB**
- **Used storage (allocated)**: From the dropdown, select to filter the items displayed by:
    - **MB**
    - **GB** (default)
    - **TB**
- **Used storage (tenant's own)**: From the dropdown, select to filter the items displayed by:

    - **MB**
    - **GB** (default)
    - **TB**
- **Total available storage**: From the dropdown, select to filter the items displayed by:
    - **MB**
    - **GB** (default)
    - **TB**
- **Available storage (allocated)**: From the dropdown, select to filter the items displayed by:
    - **MB**
    - **GB** (default)
    - **TB**
- **Available storage (tenant's own)**: From the dropdown, select to filter the items displayed by:
    - **MB**
    - **GB** (default)
    - **TB**

# Using the MSP Dashboard

Tenant management for a multi-tenancy product with a non-MSP license is available from the MSP (Master Tenant) Dashboard.

This section covers the following topics:

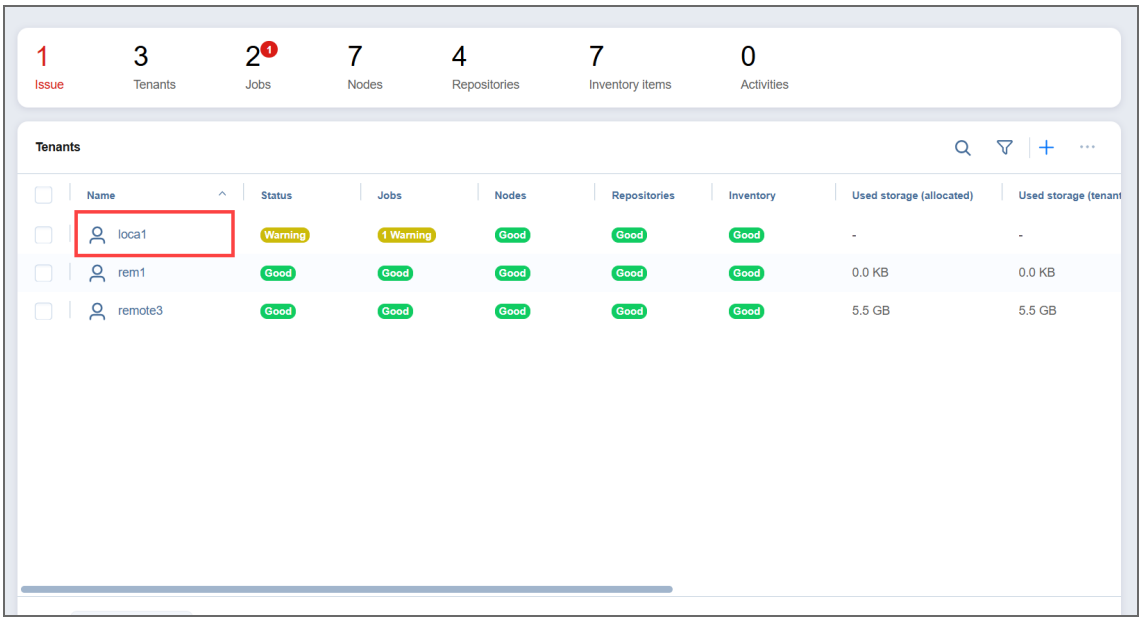

- "Opening Tenant Dashboard" on the next page
- "Viewing Tenant Information" on page 1280
- "Using Filters" on page 1282
- "Using Labels" on page 1283
- "Editing Tenants" on page 1286
- "Disabling Tenants" on page 1287
- "Deleting Tenants" on page 1288

**Notes**

- Users with an *MSP license*, *Beta instance*, *Promo license*, or *Trial license* can additionally access the MSP Console and the Licensing > Tenants tab. This allows them to efficiently oversee all independent instances of NAKIVO Backup & Replication associated with a managed service provider (MSP) as well as local tenants from a unified interface, eliminating the need to navigate through individual tenants.
- With the MSP Console feature, the standalone user is referred to as a remote tenant. Unlike local tenants in the traditional Multi-Tenancy workflow, remote tenants retain the ability to manage their resources in their data protection infrastructure.

# Opening Tenant Dashboard

To work with a tenant's instance, you should open the tenant's **Dashboard**. For local tenants, this allows you to configure the tenant, create jobs and groups, and perform recovery. For remote tenants, this allows you to monitor their instance of NAKIVO Backup & Replication. To open a tenant's **Dashboard**, simply click the tenant.

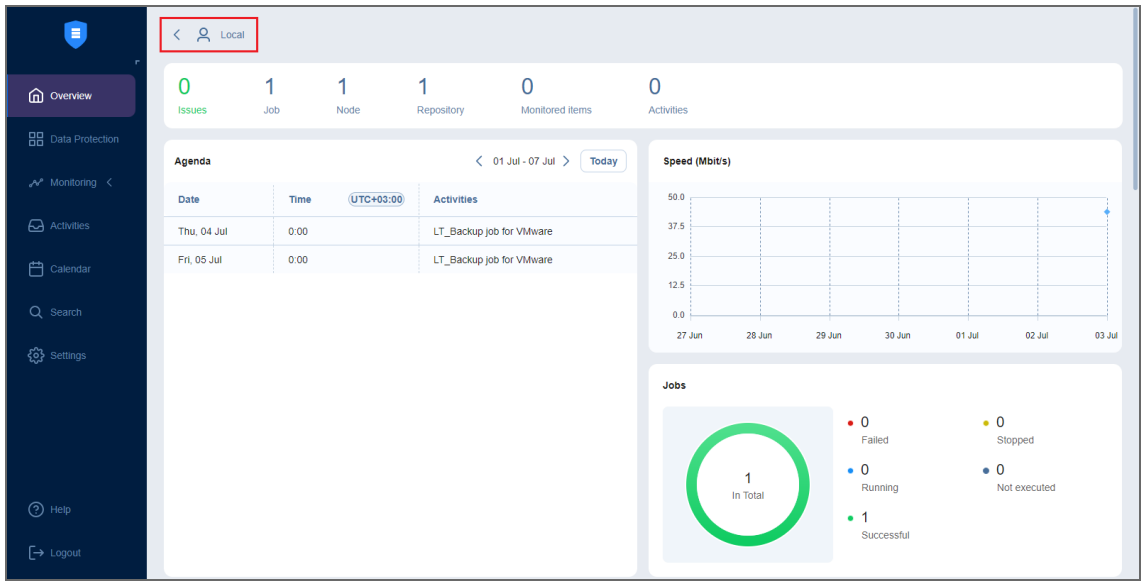


**Note**

Users with an *MSP license*, *Beta instance*, *Promo license*, or *Trial license* can additionally access the MSP Console and the Licensing > Tenants tab. This allows them to efficiently oversee all independent instances of NAKIVO Backup & Replication associated with a managed service provider (MSP) as well as local tenants from a unified interface, eliminating the need to navigate through individual tenants.

## Returning to Master Admin Dashboard

To return to the **Master Tenant Dashboard** from a local tenant's instance, click **Tenants** in the navigation bar. To return from a remote tenant's instance, click the arrow to the left of the tenant name.

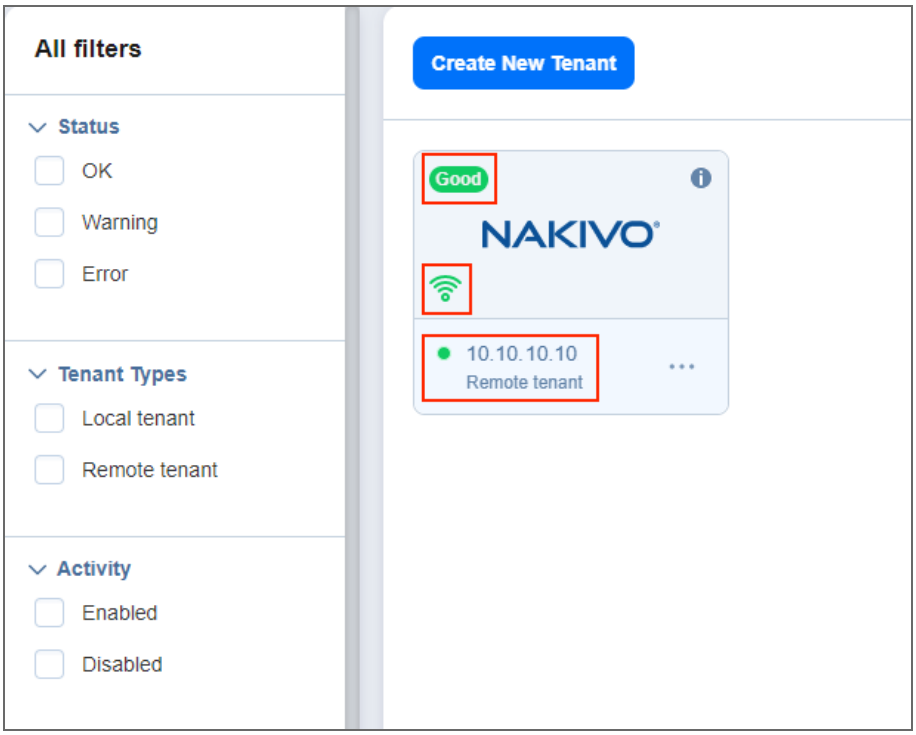



# Viewing Tenant Information

On the **Master Tenant Dashboard**, you can view information about each tenant's instance of NAKIVO Backup & Replication. The information readily displayed on a tenant's card is as follows:

- **Tenant status**: The color and content of this indicator gives an overview of the tenant instance's alarms and notifications status. The given number reflects the number of alarms and/or notifications present at the remote tenant. A green **Good** indicator means there are no outstanding alarms and notifications. Other colors represent the following:
  - **Yellow**: There are outstanding notifications.
  - **Red**: There are outstanding alarms.
  - **Grey**: The tenant is disabled.

- **Connection** (remote tenants only): A green signal icon on a remote tenant's card indicates that a connection has been established between the remote tenant and Master Tenant instances; that is to say, a green signal icon will appear on the remote tenant's card once they have successfully added the MSP. A red signal icon means the connection could not be established or has been interrupted.
- **Accessibility**: A green circle icon next to the tenant's name indicates that the tenant is currently accessible by the Master Tenant.
- **Tenant name and type**: Lastly, the tenant card indicates the name and type of a given tenant.

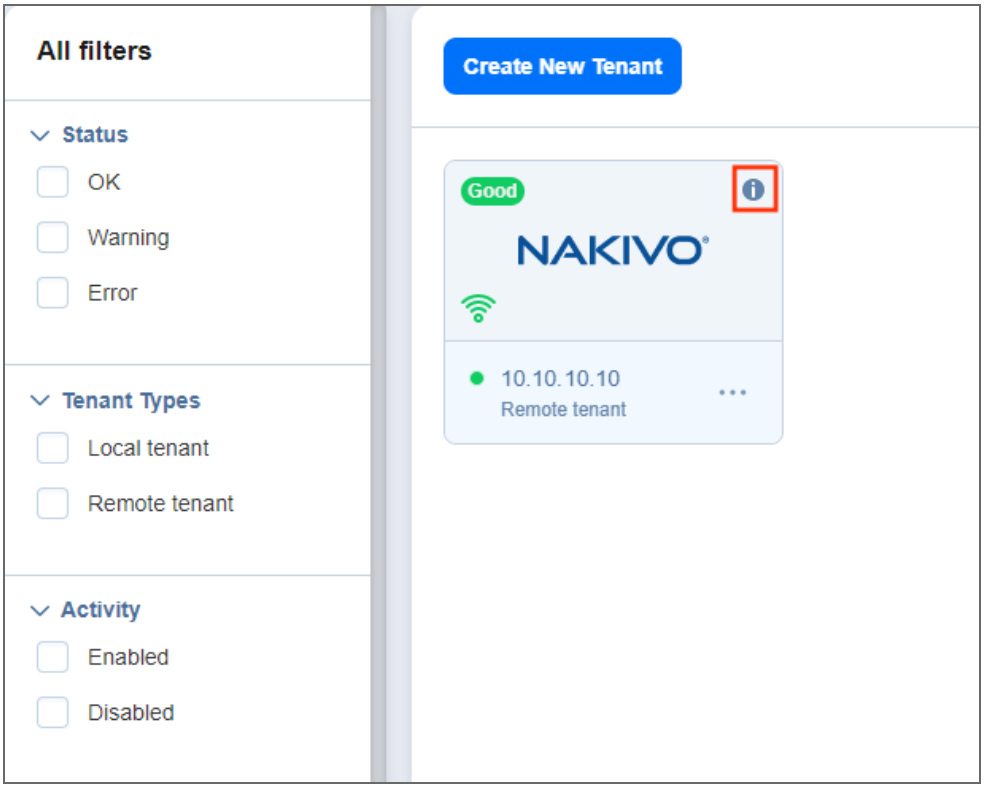



**Note**

Users with an *MSP license*, *Beta instance*, *Promo license*, or *Trial license* can additionally access the MSP Console and the Licensing > Tenants tab. This allows them to efficiently oversee all independent instances of NAKIVO Backup & Replication associated with a managed service provider (MSP) as well as local tenants from a unified interface, eliminating the need to navigate through individual tenants.

To learn more, refer to "Using the MSP Console" on page 1265.

For more tenant information, hover over the tenant card and click on the **Info** button. A pop-up window opens with general tenant details and usage statistics.




# Using Filters

- About Filters
- Applying Filters

## About Filters

The **Master Tenant Dashboard** has 5 filter categories, which allow you to quickly display tenants based on their characteristics. The following filters are available:

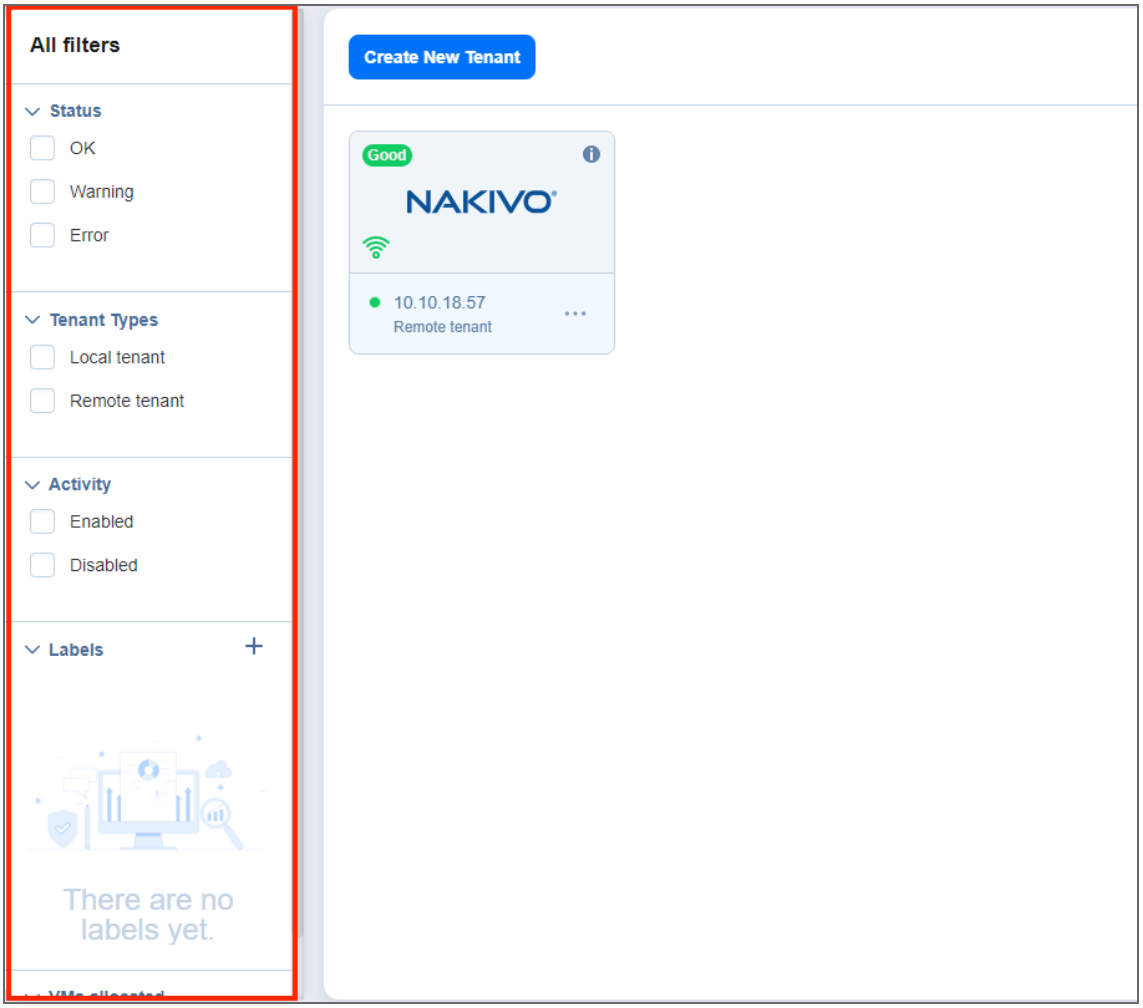
- **Status**:
  - **OK**: Displays tenants that have no errors and notifications
  - **Warning**: Displays only tenants that have notifications
  - **Error**: Displays only tenants that have errors
- **Type**:
  - **Local**: Displays only local tenants
  - **Remote**: Displays only remote tenants
- **Activity**:
  - **Enabled**: Displays only enabled tenants
  - **Disabled**: Displays only disabled tenants
- **Labels**: Filters tenants by labels assigned to them
- **VMs Allocated**: Filters tenants by the number of VMs allocated to them

**Note**

Users with an *MSP license*, *Beta instance*, *Promo license*, or *Trial license* can additionally access the MSP Console and the Licensing > Tenants tab. This allows them to efficiently oversee all independent instances of NAKIVO Backup & Replication associated with a managed service provider (MSP) as well as local tenants from a unified interface, eliminating the need to navigate through individual tenants.

## Applying Filters

To apply a filter, check the box to the left of a filter name.




To dismiss a filter, simply uncheck the box to the left of the name of an active filter.

# Using Labels

- About Labels
- Creating Labels
- Assigning Labels to Tenants

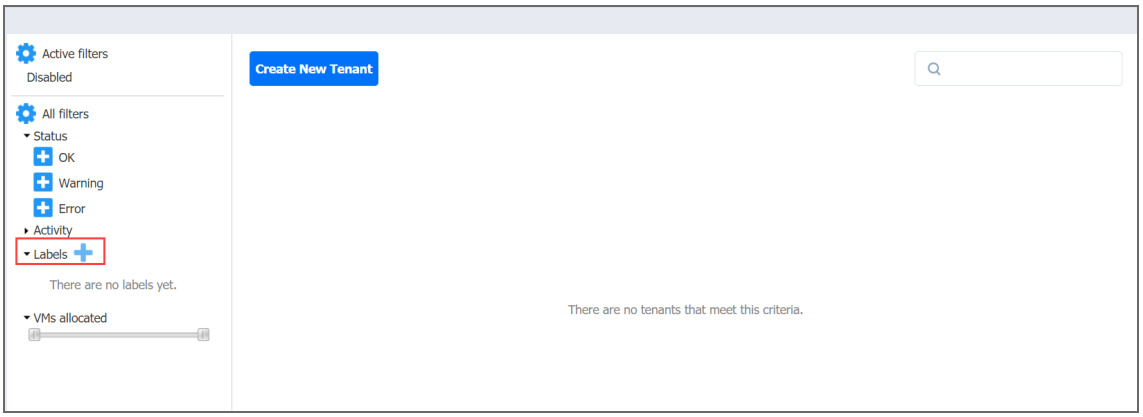
- Editing Label Names
- Deleting Label


## About Labels

With NAKIVO Backup & Replication, you can create custom labels and assign them to tenants. Assigning a label to a tenant allows you to quickly sort existing tenants into different categories, such as location, SLA level, etc.

## Creating Labels

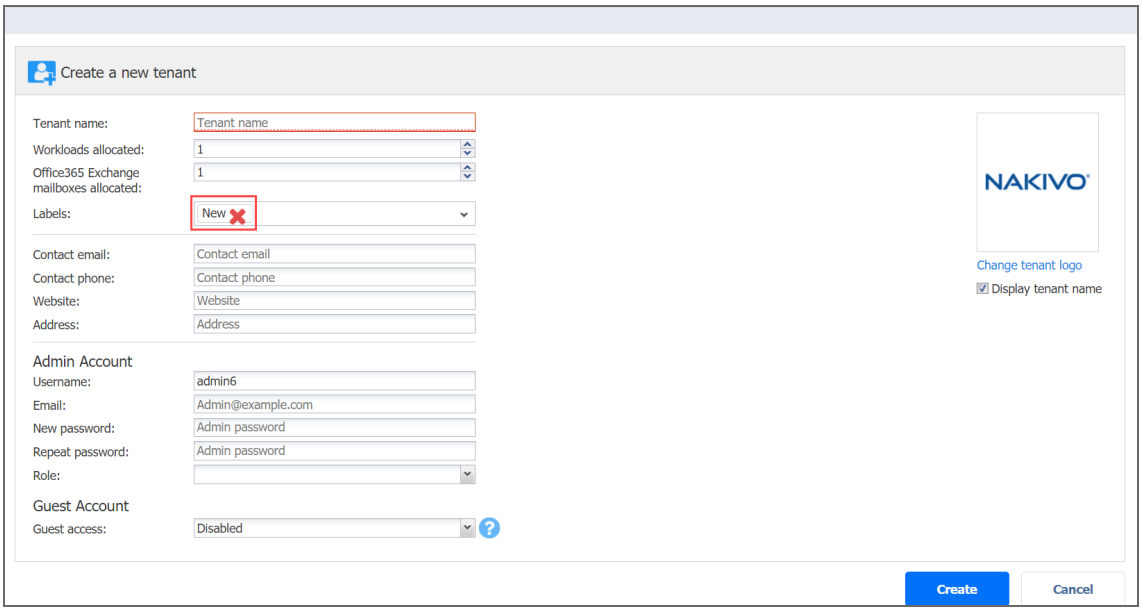
To create a new label, click the **Plus** icon next to **Labels** and enter a name for the new label, and press the **Enter** key.




You can also create a new label when creating a new tenant.

# Assigning Labels to Tenants

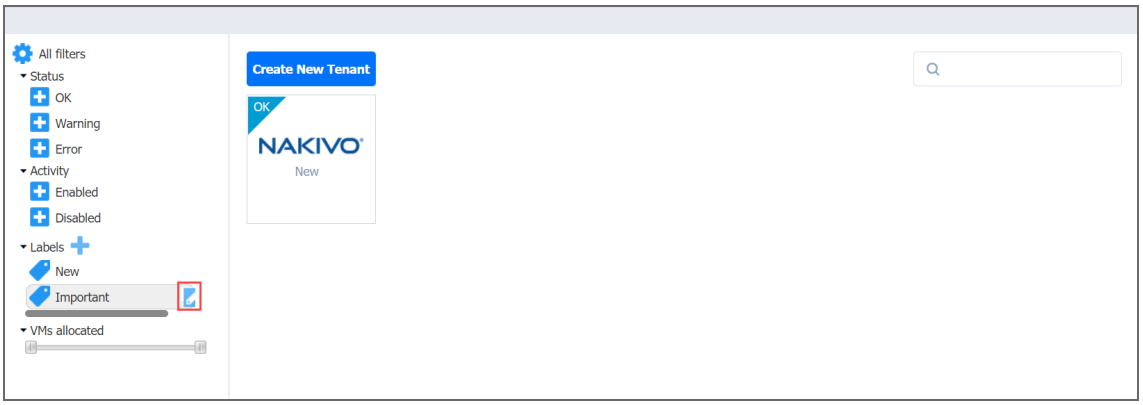
You can assign a label to a tenant either during the tenant creation or by editing the tenant.




# Editing Label Names

To change a label name, do the following:

1. Hover over the label.
2. Click the **Edit** icon.




3. Enter the new label name and press the **Enter** key.

# Deleting Labels

To permanently delete a label, do the following:

1. Hover the mouse pointer over a label.
2. Click the **Delete** icon.

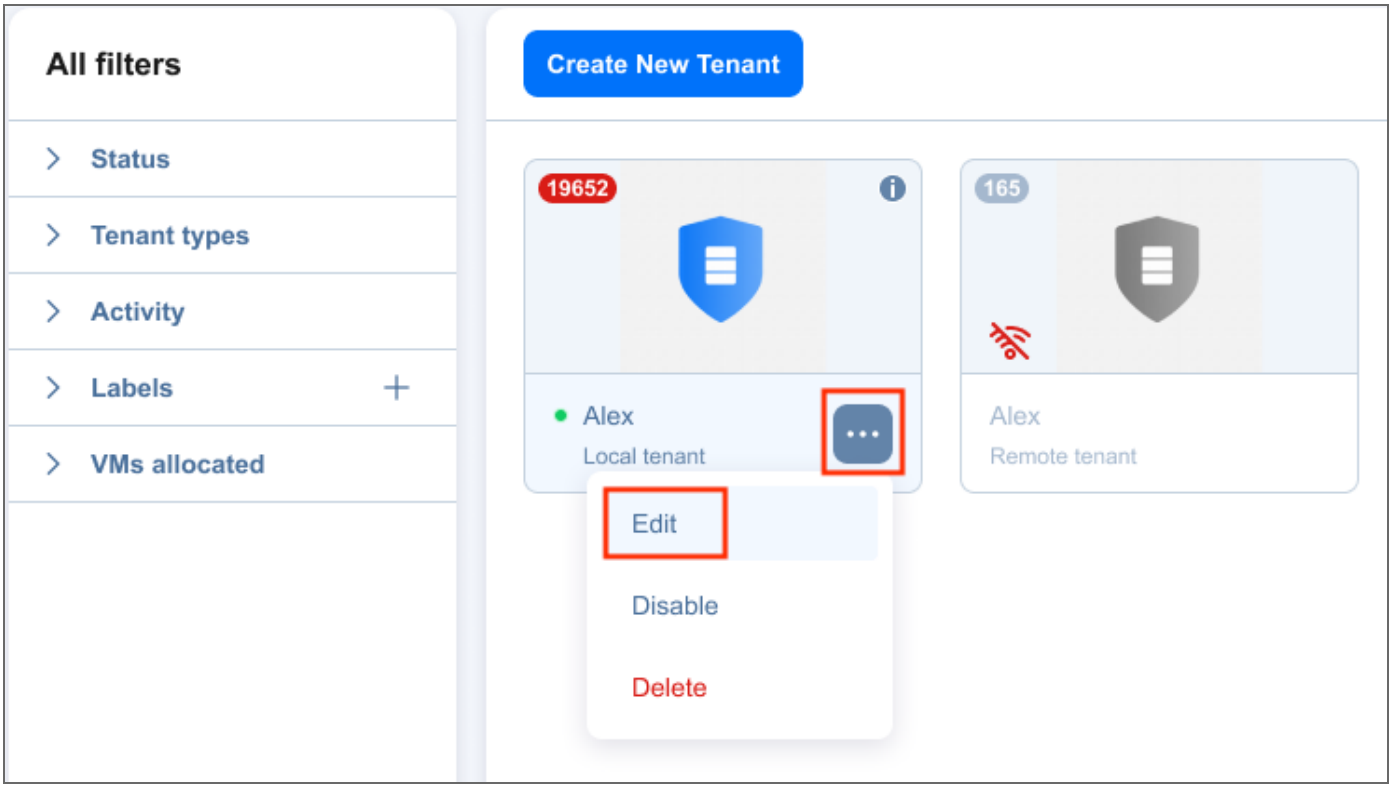
3. In the dialog box that opens, click **Delete** to confirm that you wish to permanently delete the label




# Editing Tenants

To edit a tenant, do the following:

1. Hover over the tenant card and click the ellipsis **Manage** button, then click **Edit**.




2. In the **Edit** dialog box that opens, make the required changes and click **Save**.

**Note**

Users with an *MSP license*, *Beta instance*, *Promo license*, or *Trial license* can additionally access the MSP Console and the Licensing > Tenants tab. This allows them to efficiently oversee all independent instances of NAKIVO Backup & Replication associated with a managed service provider (MSP) as well as local tenants from a unified interface, eliminating the need to navigate through individual tenants.

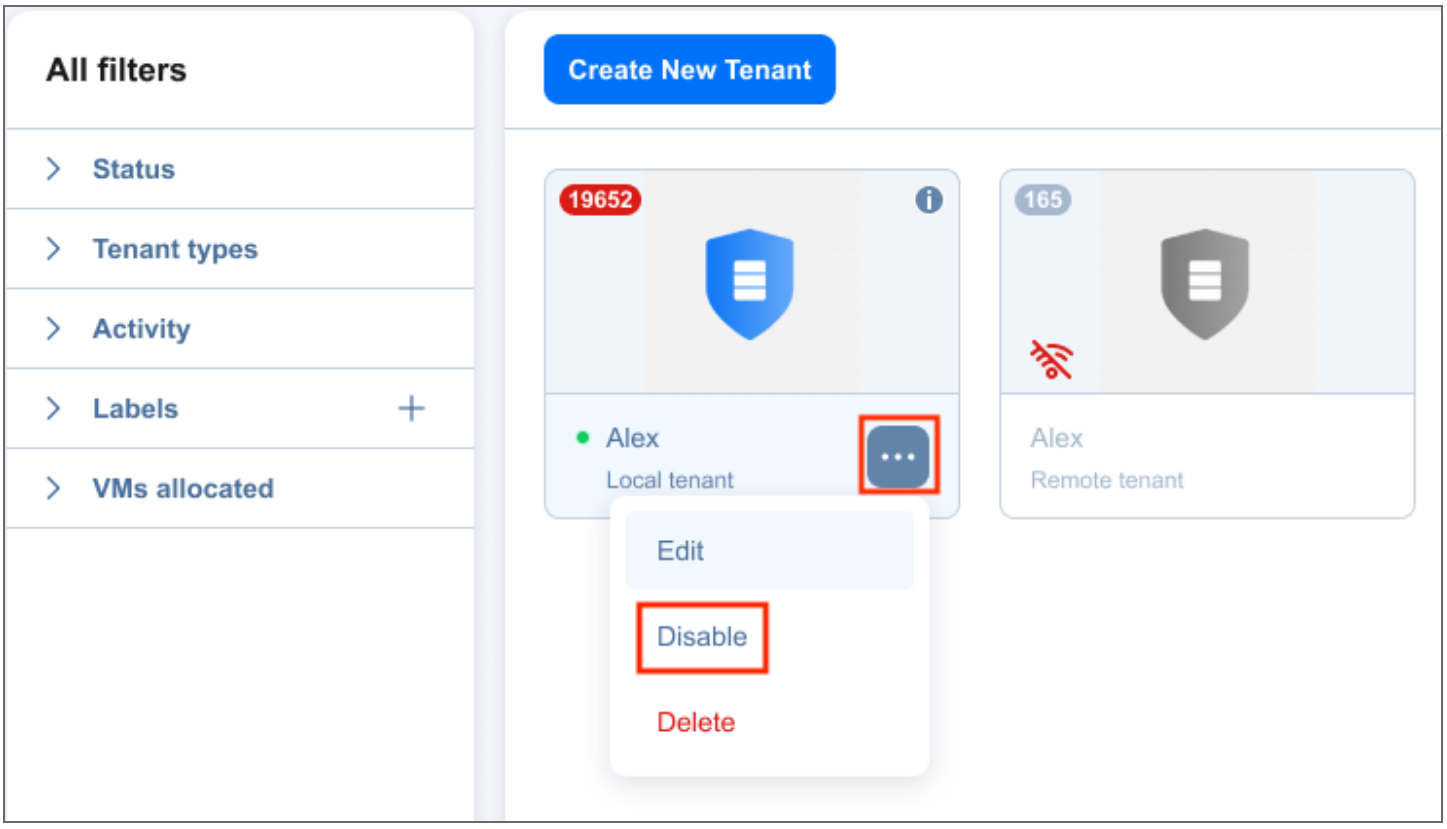
To learn more, refer to "Using the MSP Console" on page 1265.

# Disabling Tenants

In multi-tenant mode, Master Admin can disable a tenant to temporarily stop delivering backup, replication, and recovery services for that tenant. After disabling a tenant:

- Tenant admin and tenant guest will not be able to log in to the self-service interface. A message saying that the service has been disabled will be displayed after login attempts.
- Existing jobs will not be run on schedule.
- All currently running jobs will be allowed to complete.

To disable a tenant, hover over the tenant card and click the ellipsis **Manage** button, then click **Disable**.



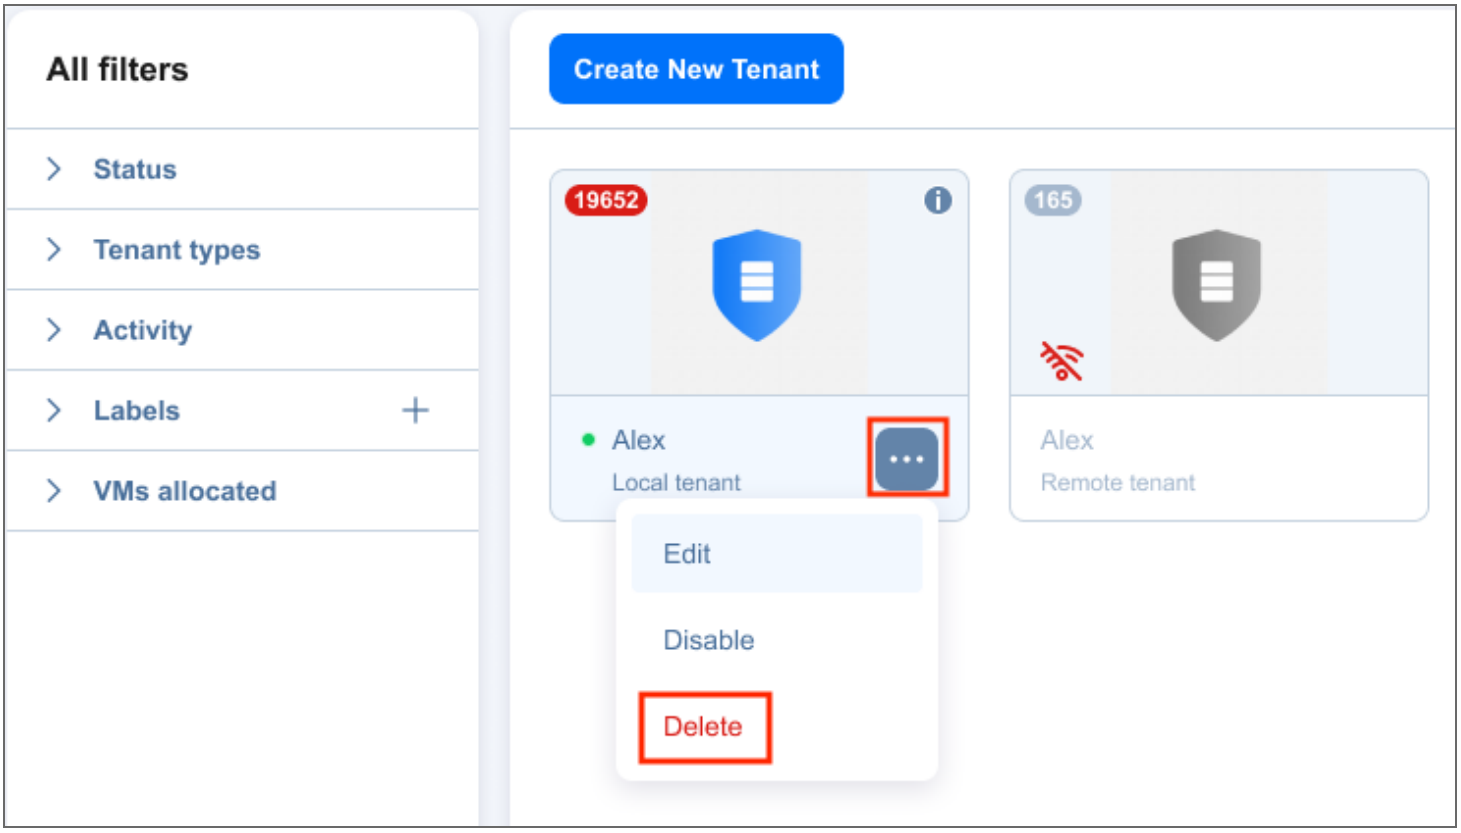
**Note**

Users with an *MSP license*, *Beta instance*, *Promo license*, or *Trial license* can additionally access the MSP Console and the Licensing > Tenants tab. This allows them to efficiently oversee all independent instances of NAKIVO Backup & Replication associated with a managed service provider (MSP) as well as local tenants from a unified interface, eliminating the need to navigate through individual tenants.

To learn more, refer to "Using the MSP Console" on page 1265.

# Deleting Tenants

To permanently delete a tenant from the product, hover over the tenant card and click the ellipsis **Manage** button, then click **Delete**.




The tenant will be permanently deleted from NAKIVO Backup & Replication.

Tenant Transporters are not uninstalled and the Tenant Backup Repositories are not removed.

**Note**

Users with an *MSP license*, *Beta instance*, *Promo license*, or *Trial license* can additionally access the MSP Console and the Licensing > Tenants tab. This allows them to efficiently oversee all independent instances of NAKIVO Backup & Replication associated with a managed service provider (MSP) as well as local tenants from a unified interface, eliminating the need to navigate through individual tenants.

To learn more, refer to "Using the MSP Console" on page 1265.

# Granting Self-Service Access

In the Multi-tenant mode, you can provide local tenants with access to their dashboards. By default, a tenant admin account is automatically created when you create a new local tenant. The tenant admin has full control over the product features inside the tenant dashboard (such as the ability to edit and update tenant inventory, transporters, and backup repositories, and create and manage jobs and groups). For each local tenant, one guest account can also be created. The tenant guest has limited permissions inside the tenant instance and can only generate job and group reports by default. To provide a local tenant with access to the self-service interface, send the following information to the tenant:

- Link to NAKIVO Backup & Replication **Director**
- Tenant login
- Tenant password