

# TOC

---

<b>NAKIVO Backup &amp; Replication Overview</b> .....	<b>6</b>
Data Protection .....	7
Data Recovery .....	24
Disaster Recovery .....	30
Reliability .....	37
Performance .....	49
Administration .....	64
Automation .....	67
Integration .....	69
BaaS .....	72
NAKIVO Licensing Policy .....	77
<b>Deployment</b> .....	<b>87</b>
Architecture .....	87
System Requirements .....	96
Deployment Scenarios .....	137
Installing NAKIVO Backup & Replication .....	148
Updating NAKIVO Backup & Replication .....	214
Uninstalling NAKIVO Backup & Replication .....	235
<b>Getting Started</b> .....	<b>239</b>
Logging in to NAKIVO Backup & Replication .....	239
First Steps with NAKIVO Backup & Replication .....	245
Web Interface Components .....	249



---

Managing Jobs and Activities .....	255
<b>Settings .....</b>	<b>284</b>
General .....	284
Inventory .....	334
Transporters .....	365
Backup Repositories .....	388
Tape .....	445
Expert Mode .....	472
Virtual Appliance Configuration .....	489
Multi-Tenant Mode Configuration .....	491
Support Bundles .....	497
Built-in Support Chat .....	500
<b>Backup .....</b>	<b>507</b>
Creating VMware Backup Jobs .....	507
Creating VMware Cloud Director Backup Jobs .....	533
Creating Hyper-V Backup Jobs .....	551
Creating Amazon EC2 Backup Jobs .....	573
Creating Nutanix AHV Backup Jobs .....	597
Creating Physical Machine Backup Jobs .....	615
Creating Microsoft 365 Backup Jobs .....	634
Creating File Share Backup Jobs .....	646
Creating Oracle Database Backup Jobs .....	657
Creating Backup Copy Jobs .....	668
Backing Up to Tape .....	689
Staging (Seeding) Initial Backup .....	705

---

Deleting Backups .....	706
<b>Replication .....</b>	<b>714</b>
Creating VMware Replication Jobs .....	714
Creating Hyper-V Replication Jobs .....	745
Creating Amazon EC2 Replication Jobs .....	773
Staging (Seeding) VM Replication .....	786
<b>Recovery .....</b>	<b>787</b>
Granular Recovery .....	788
Full Recovery .....	847
Planning Disaster Recovery .....	971
Failover VMware VMs: VMs .....	1022
Failover VMware VMs: Options .....	1023
Failover Hyper-V VMs: VMs .....	1023
Failover Hyper-V VMs: Options .....	1024
Failover Amazon EC2 Instances: Instances .....	1025
Failover Amazon EC2 Instances: Destination .....	1026
Failover Amazon EC2 Instances: Options .....	1029
Failback VMware VMs: VMs .....	1030
Failback VMware VMs: Location .....	1031
Failback VMware VMs: Options .....	1033
Failback Hyper-V VMs: VMs .....	1034
Failback Hyper-V VMs: Location .....	1034
Failback Hyper-V VMs: Options .....	1036
Failback Amazon EC2 Instances: Instances .....	1037
Failback Amazon EC2 Instances: Location .....	1038

---

Failback EC2 Instances: Options .....	1041
<b>Monitoring .....</b>	<b>1062</b>
Navigating The Monitoring Dashboard .....	1062
Managing The Monitored Items .....	1063
<b>Integration and Automation .....</b>	<b>1067</b>
Command Line Interface .....	1067
Automation with HTTP API .....	1075
Aptare IT Analytics Integration .....	1075
<b>Multi-Tenant Mode .....</b>	<b>1082</b>
Tenant Creation .....	1082
Tenant Configuration .....	1087
Tenant Management .....	1088
Granting Self-Service Access .....	1094

# NAKIVO Backup & Replication v10.6

## **USER GUIDE**

# 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” below](#)
- [“Backup Copy” on page 9](#)
- [“Backup to Cloud” on page 14](#)
- [“Container Protection” on page 15](#)
- [“Microsoft 365 Backup” on page 16](#)
- [“Oracle Database Backup” on page 17](#)
- [“Physical Machine Backup” on page 18](#)
- [“Backup to Tape” on page 20](#)
- [“Virtual Machine Backup” on page 22](#)
- [“File Share Backup” on page 23](#)

## 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

An Amazon EC 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 EC2 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

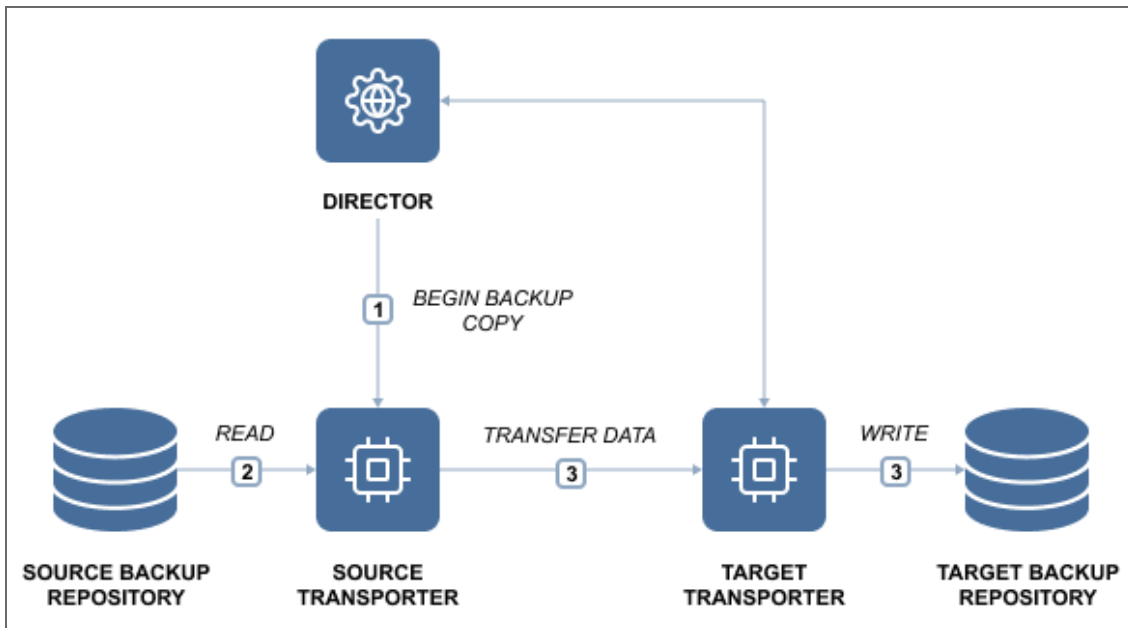
Backups can be lost on account of a number of reasons, so having more than one copy of your business-critical backups is vital for ensuring that your data can be recovered in case of disaster. Backup Copy jobs provide a simple yet powerful way to create and maintain copies of your backups. Backup copy 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](#)



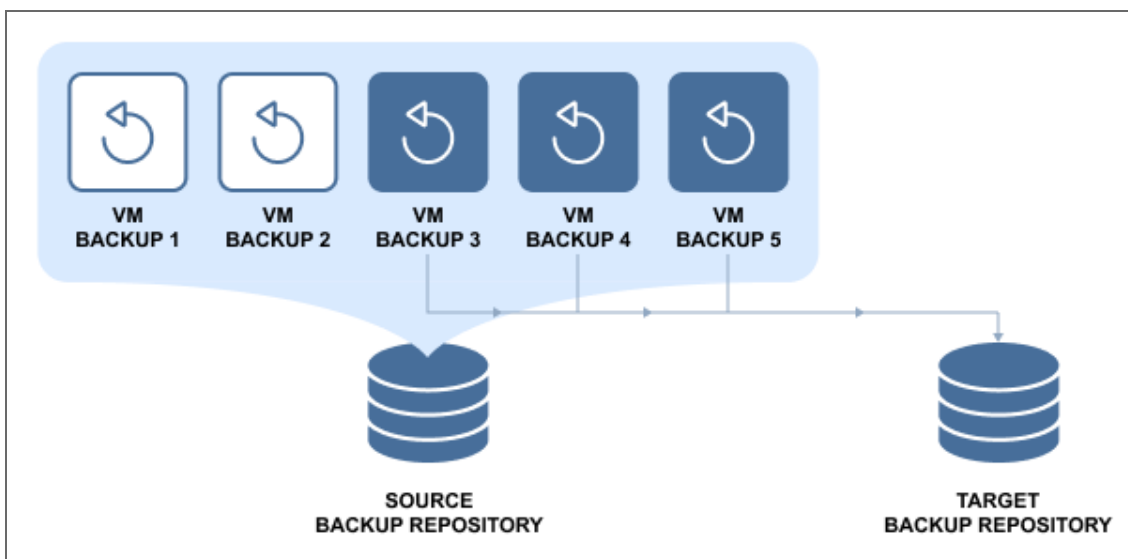
## Create Mirrored Copy of your Backup Repository

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

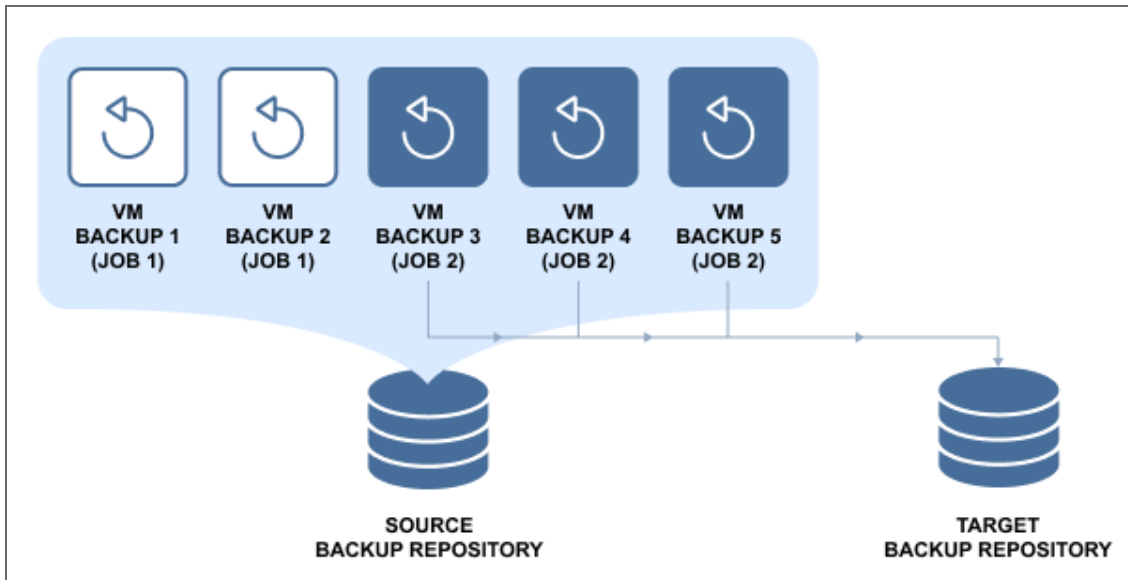
To save storage space on your secondary Backup Repository and to speed up data transfer, you can choose to create a Backup Copy job for only the most important backups:



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.

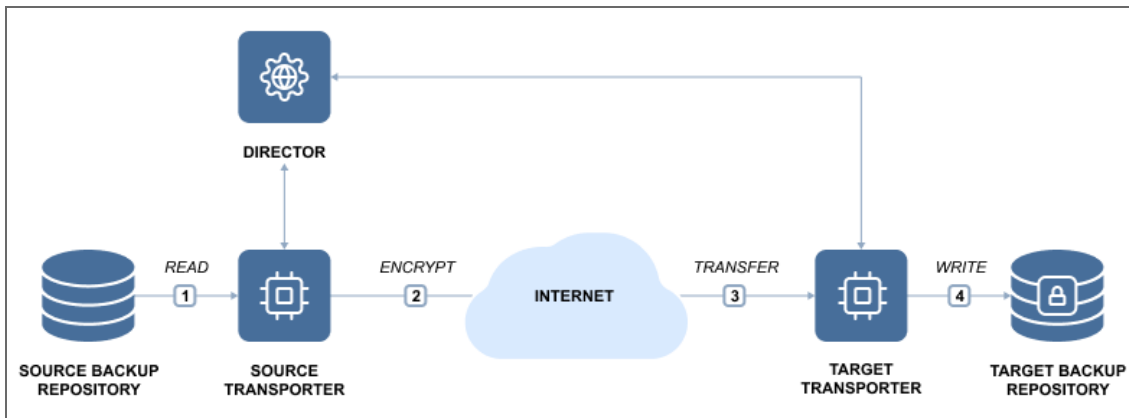
## Resource Efficiency and Variable Data Compression

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.

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.

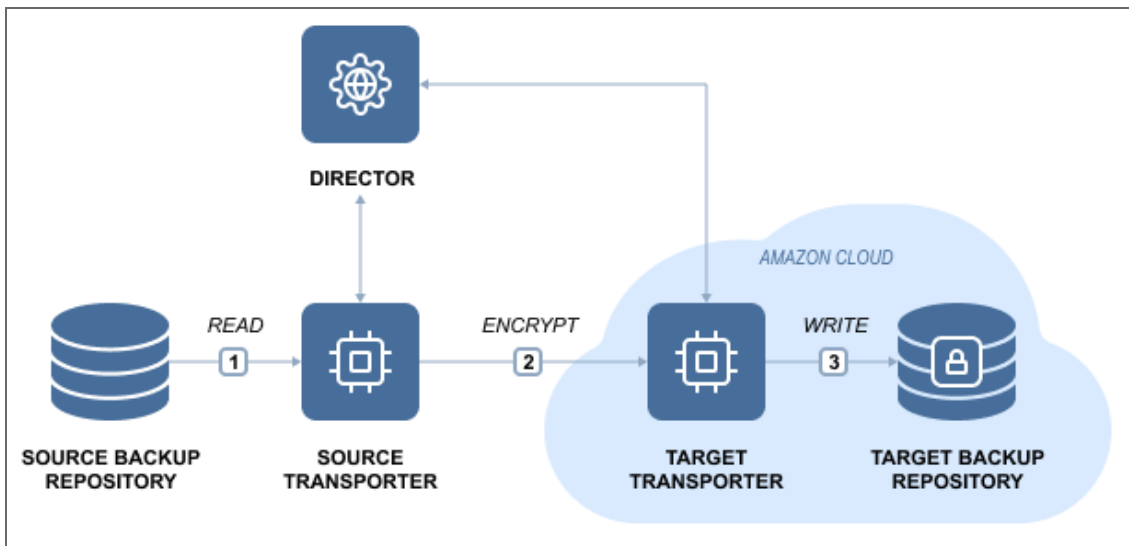
## Copy Backups Offsite

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.

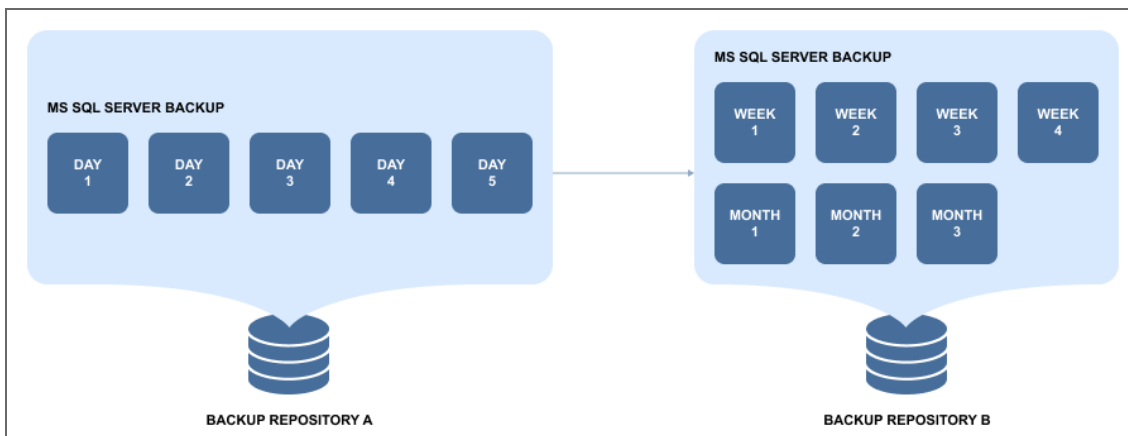


## Copy Recovery Points that You Need

Each backup can contain up to 4,000 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.

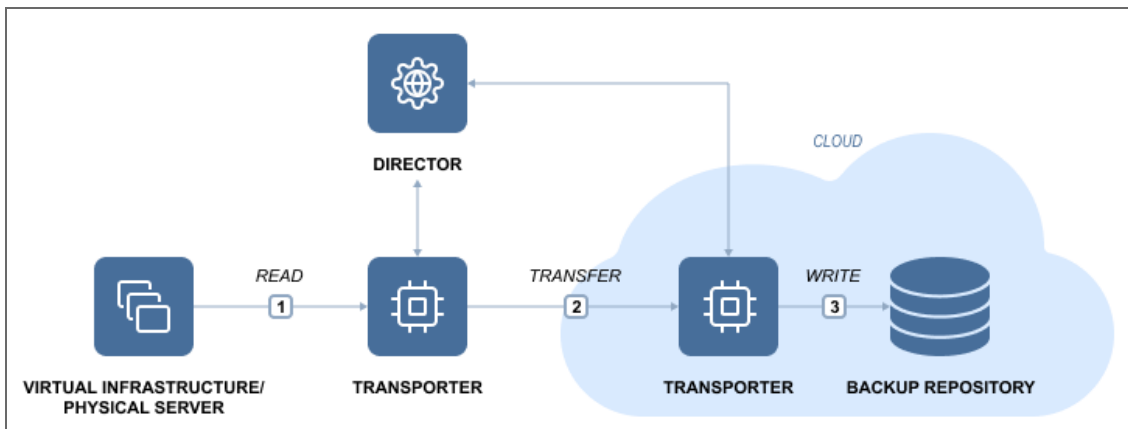
## Schedule Backup Copy to Suit Your Needs

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 668](#).

# Backup to Cloud

NAKIVO Backup & Replication provides a great way for safeguarding your business-critical data by letting you send backup copies to Amazon EC2, Amazon S3, 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 local infrastructure becomes unavailable.
- Flexible Backup Storage – cloud environments allow for expanding storage space 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 and configuring 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 for scheduling regular "backup to cloud" jobs.

While cloud providers offer cloud storage at an affordable price, NAKIVO Backup & Replication helps reduce offsite backup costs with features like incremental backup, exclude swap files and partitions, backup deduplication, backup compression, and other. Using NAKIVO Backup & Replication, you can keep the entire environment in the cloud or only use Amazon EC2, Amazon S3, or Wasabi as a storage for backups.

## Technology Behind "Backup to Cloud"

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 a destination designated for data storage. NAKIVO Backup & Replication allows you to send backups or their copies to private/public clouds such as Amazon EC2, Amazon S3, 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 machine.
2. The data that was changed (since the last backup) is 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 run backup copy jobs, which allow you to create and manage copies of your VMware, Hyper-V, Amazon EC2 or physical machine backups. Creating copies of critical backups serves as an additional level of data protection.

NAKIVO Backup & Replication enables you to copy backups from one Backup Repository to another without touching the source hosts or VMs. This decreases backup time and reduces network load. The process is entirely automatic, meaning that you are only required to create and set up a backup copy job. After the initial configuration, your secondary Backup Repository is automatically updated with all backups and recovery points from the primary Backup Repository.

NAKIVO Backup & Replication 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.

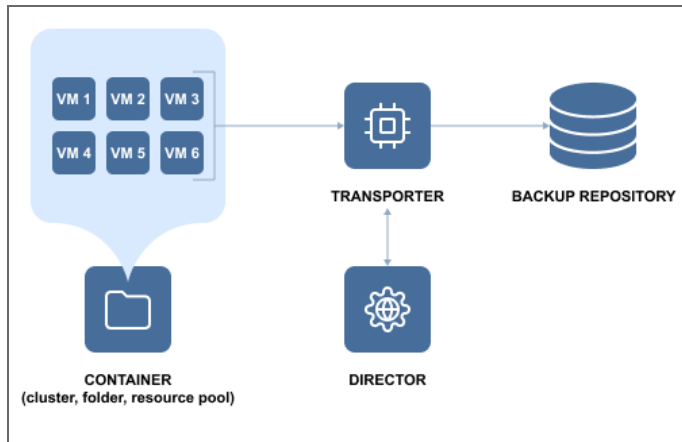
For more details, refer to the following topics:

- [“Creating Amazon EC2 Backup Jobs” on page 573](#)
- [“Backup Repository in Amazon EC2” on page 391](#)
- [“Backup Repository in Amazon S3” on page 396](#)
- [“Backup Repository in Wasabi Hot Cloud Storage” on page 399](#)


## 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

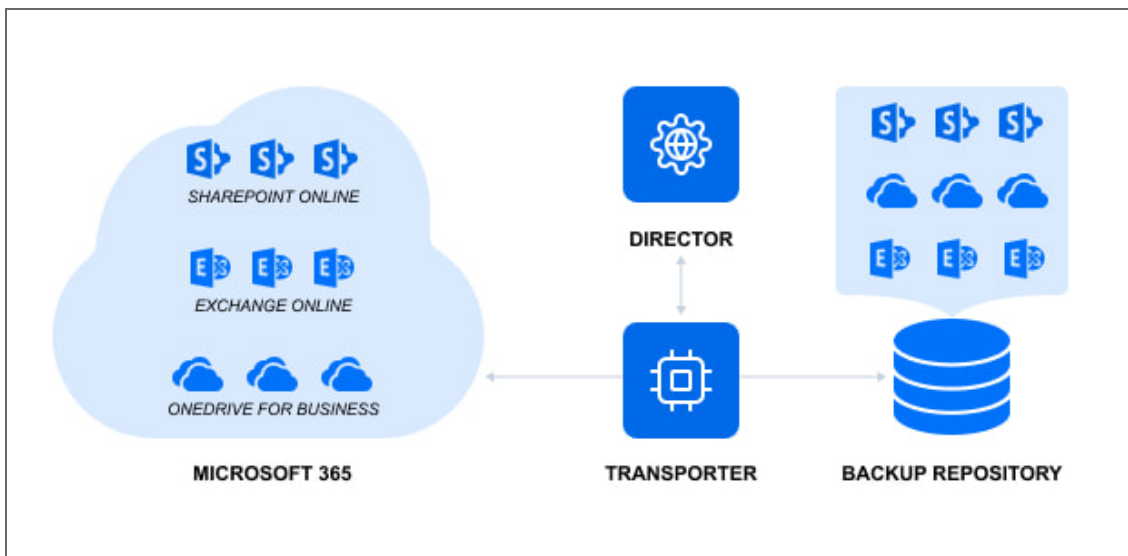


- Day 2



## Microsoft 365 Backup

With NAKIVO Backup & Replication, you can back up your organization's entire Microsoft 365 account or individual users who have access to the following services: Exchange Online, OneDrive for Business, and SharePoint Online. The backups are stored in a SaaS Backup Repository for further recovery of Exchange Online, OneDrive for Business, and SharePoint Online data including user, shared, and group mailboxes, contacts, calendar events, emails, drives, team, communication, and personal sites, subsites, document libraries, lists, list items, and individual files and folders.



A common misconception among Microsoft 365 users is that data stored in the cloud is safe and that it is not necessary to back it up. However, under Microsoft’s “shared responsibility model”, Microsoft is responsible for reliable uptime and the physical and virtual infrastructure. Beyond simple geo-redundancy – which is not the same as a backup – ensuring data safety is the responsibility of users. To address this gap, NAKIVO Backup & Replication offers Backup for Microsoft 365 to ensure that data is accessible and recoverable at all times. By backing up Microsoft 365 application data, you ensure that if data loss occurs, the necessary items can be easily recovered to the original or a custom location.

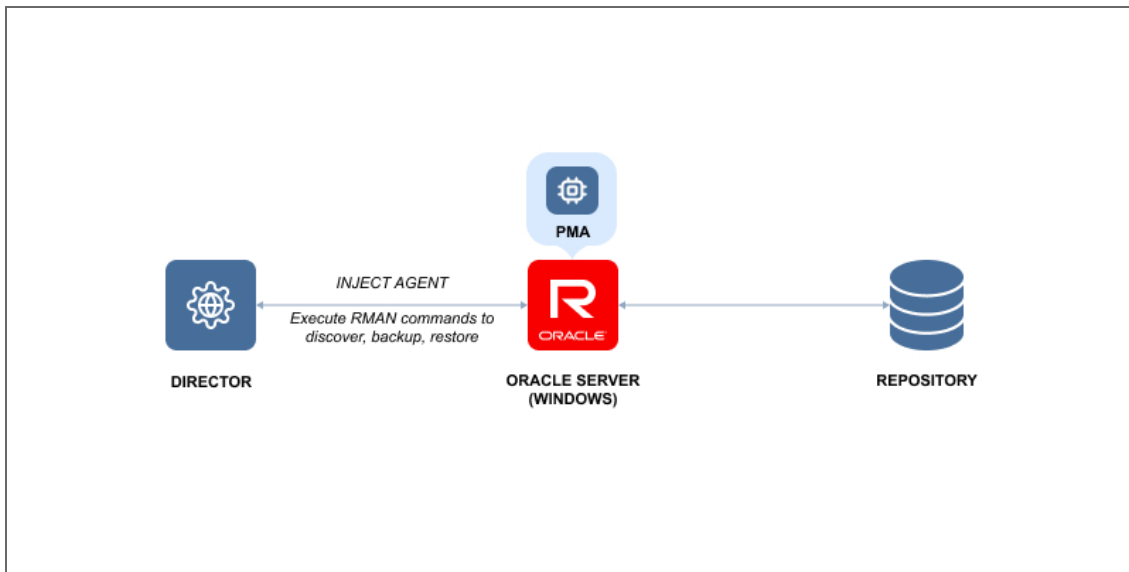
To start using this feature, take the following preliminary steps:

- [Add your Microsoft 365 accounts to the inventory](#)
- [Create a SaaS Backup Repository](#)
- [Create a Microsoft 365 job](#)

## Oracle Database Backup

With NAKIVO Backup & Replication and the help of the Oracle Recovery Manager (RMAN), you can successfully protect your Oracle databases. The product uses an agent to run scripts/execute RMAN commands for backup and recovery operations. To start the backup process, you need to pre-configure specific settings in the [RMAN configuration](#) and then add your Oracle database to the product [inventory](#). Using Dashboard, you can create Oracle backup jobs and recover entire databases to the destination pre-configured in RMAN.





To start using this feature, take the following steps:

- [Add your Oracle database to the product inventory](#)
- [Create an Oracle database backup job](#)
- [Recover the Oracle database](#)

## Physical Machine Backup

NAKIVO Backup & Replication offers capabilities for data protection of [physical infrastructures](#) with strict data protection requirements. By using the product, you can seamlessly perform physical machine backups ensuring the consistency of applications and databases.

- [Physical Machine Backup Feature](#)
- [How It Works](#)

### Physical Machine Backup Feature

NAKIVO Backup & Replication performs incremental backup jobs using a proprietary change tracking method which allows you to save time and storage resources as it transfers only changed blocks of data to the backup repository. To protect your backups, the product utilizes AES 256-bit encryption to transform data into an unreadable ciphertext to prevent unauthorized access to it. To improve backup performance, you can enable network acceleration to speed up data transfer over LAN and WAN networks. The product also allows you to recover files and application objects as well as to recover a physical machine to VMware and Microsoft Hyper-V virtual machines. This way, you can easily perform the migration of physical workloads to virtual environments.

## How it Works

To ensure your physical machines are successfully backed up and recovered, you need to [add a new physical machine](#) to the product inventory first. Support for physical machines is done via the Physical Machine Agent (PMA) deployed in the physical machine OS. Communication between the Director and PMA is secured by means of a Certificate and Pre-shared Key generated by the Director and then injected into the PMA. During a physical machine discovery, NAKIVO Backup & Replication checks to see if the PMA is already installed on the physical machine. If the PMA is detected, the product adds the physical machine to the inventory. If the PMA is outdated, it is automatically updated. If the PMA or a Transporter is not detected, the product:

1. Auto-generates a Self-signed Certificate.
2. Auto-generates a Pre-shared Key.
3. Installs the PMA and injects a Self-signed Certificate as well as a Pre-shared Key into the PMA.

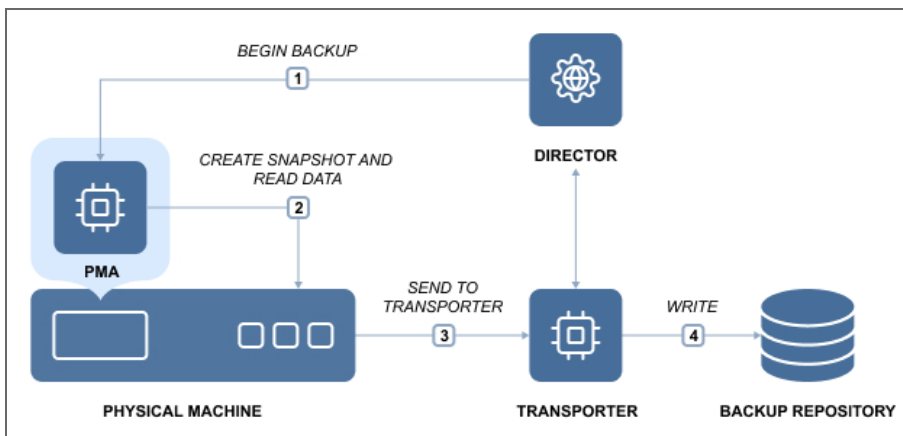
### Notes

- Manual installation of the PMA is not supported.
- The installation path is the same as the one used for the Transporter on the corresponding OS.
- The communication ports are the same as those used for the Transporter on the corresponding OS.

Creating a physical machine backup job is a five-step process:

1. Identify a physical machine that you need to backup.
2. Choose a backup repository for storing backups.
3. Set the backup job schedule.
4. Specify your retention policy.
5. Configure the backup job options.

Once the physical machine backup job has started, NAKIVO Backup & Replication captures the necessary data blocks from the physical machine and sends them to the selected backup repository for storage. The backed up data then can be accessed and recovered whenever needed.




# 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
<b>Tape Library</b>	A storage device that includes one or more tape drives, a number of slots and a media changer (robot).
<b>Tape Drive</b>	A device component (or a standalone device) used to read and write the tape cartridge.
<b>Slot</b>	A place in the tape library designed to hold a single cartridge.
<b>Mail Slot</b>	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.
<b>Media Changer</b>	A device component used to move a single tape cartridge between slots and load/unload the cartridge to/from the tape drive.
<b>Tape Cartridge (Tape)</b>	A unit of sequential magnetic medium and an optional barcode used for identification.
<b>Media Pool</b>	A logical container that contains tape cartridges.
<b>Backup (Tape)</b>	A logical entity containing one or more recovery points on one or more tape cartridge(s) that belong to a single source object.


Term	Description
<b>Recovery Point (Tape)</b>	A complete or incomplete data set required to rebuild a VM or instance as of a particular moment in time.

## 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 507](#).

## File Share Backup

NAKIVO Backup & Replication allows you to back up and recover data stored on NFS and CIFS shares. The product allows you to backup entire file shares as well as separate folders within the file shares selected for backup. The data can be backed up to both **Incremental with full backups** and **Forever incremental** backup repository types. When needed, recovery of files and folders from these repositories can be carried out to a custom share. Additionally, the recovered files can be forwarded via email or downloaded to the browser.



To protect your data stored on NFS and CIFS shares, take the following actions:

1. [Add](#) the file share to NAKIVO Backup & Replication inventory.
2. [Create](#) a file share backup job.
3. [Recover](#) file shares using the File Share Recovery Wizard when needed.

# 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:

- [“Instant File Recovery to Source” below](#)
- [“Instant Object Recovery” on the next page](#)
- [“Instant VM Recovery - Flash VM Boot” on page 26](#)
- [“Physical to Virtual Machine Recovery” on page 27](#)
- [“Universal Object Recovery” on page 27](#)
- [“Cross-Platform Recovery” on page 28](#)

## 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 788](#).

## 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 788](#).





## Instant VM Recovery - Flash VM Boot

The Flash VM 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.

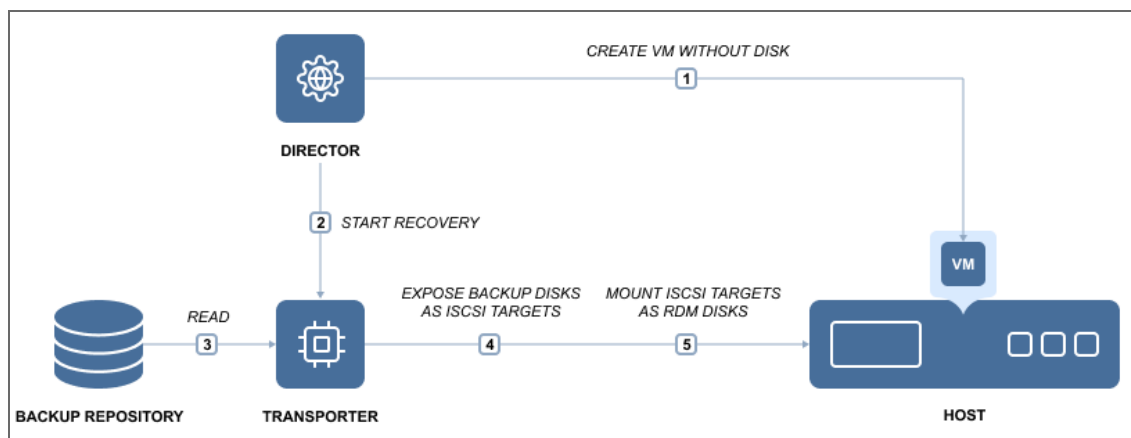
Once the machine is running, you can migrate it to production for permanent recovery. Note that the backup from which the VM is booted is not affected. Changes you make to the running VM will not modify or remove the data in your VM backup. In addition to the VM recovery capabilities, the Flash VM Boot 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 VM 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 VM 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.



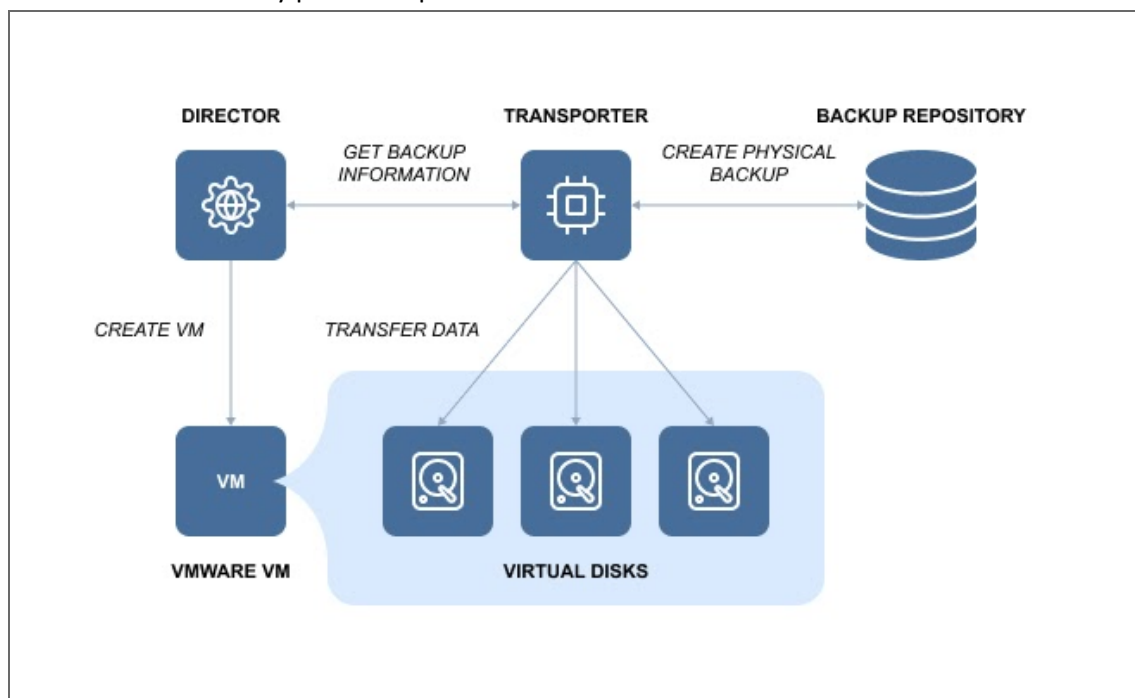
This process is fully automated and takes mere seconds to complete, after which the machine OS boot is started. Once booted, the machine can be migrated to the production environment using the hypervisor's native live migration feature.

With NAKIVO Backup & Replication you can also perform Flash VM 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 VM Boot feature, refer to [“Performing Flash VM Boot Recovery” on page 847](#)

## Physical to Virtual Machine Recovery

To protect mixed physical and virtual IT environments, NAKIVO Backup & Replication offers the Physical to Virtual Machine Recovery feature. To recover a physical machine with NAKIVO Backup & Replication, first, add the physical machine and VMware vCenter/ESXi host to the [inventory](#). Then, run a [physical machine backup job](#) and [recover the backup to a VMware VM](#) either via the **Dashboard** or the **Repositories** page in [Settings](#). See the diagram below to know how the recovery process is performed.



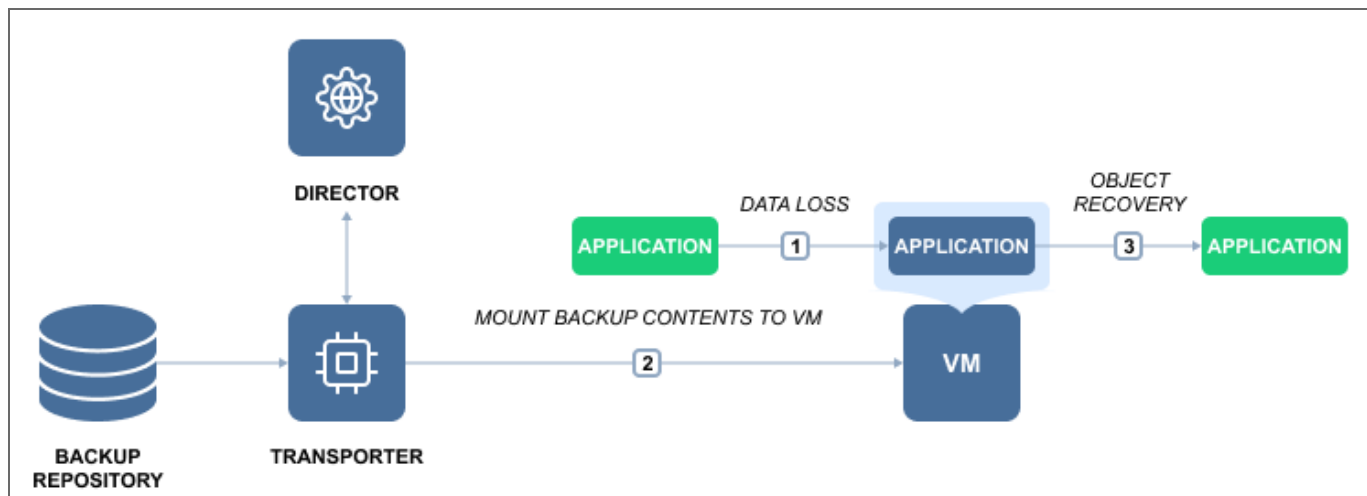
## Universal Object Recovery

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 788](#) section.

## Cross-Platform Recovery

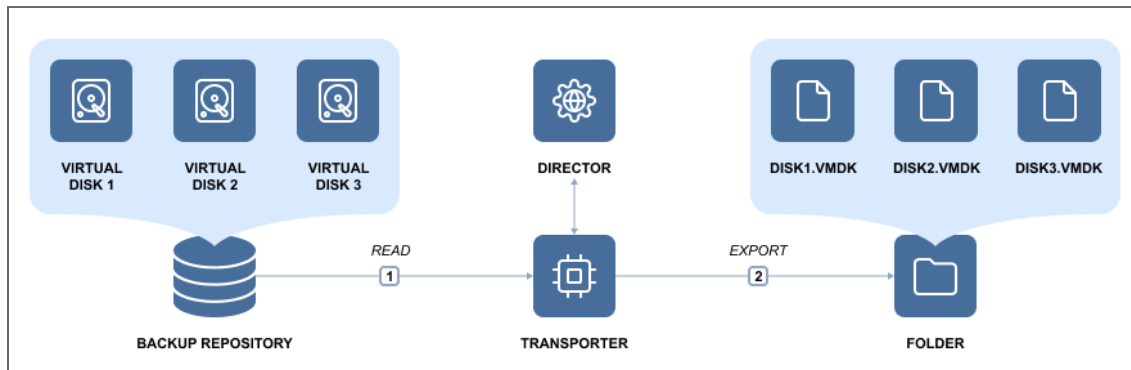
With Cross-Platform Recovery, you can seamlessly protect VM/physical machine data across multiple platforms and virtualized environments. You can also benefit from the following other advantages:

- **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 or physical server).
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.

# 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 application allows you to protect VMs running within a cluster, replicate VMs, and fail over to replicas.

When utilizing Site Recovery, 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.


This section contains the following topics:

- [“Replication From Backup” below](#)
- [“Replication Types” on the next page](#)
- [“Site Recovery” on page 33](#)

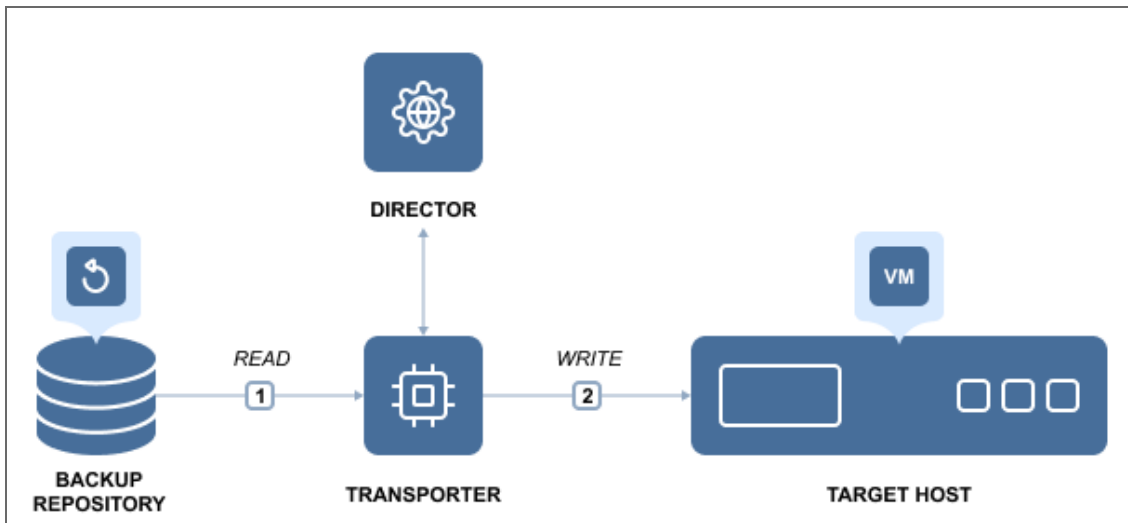
## Replication From Backup

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 714](#).

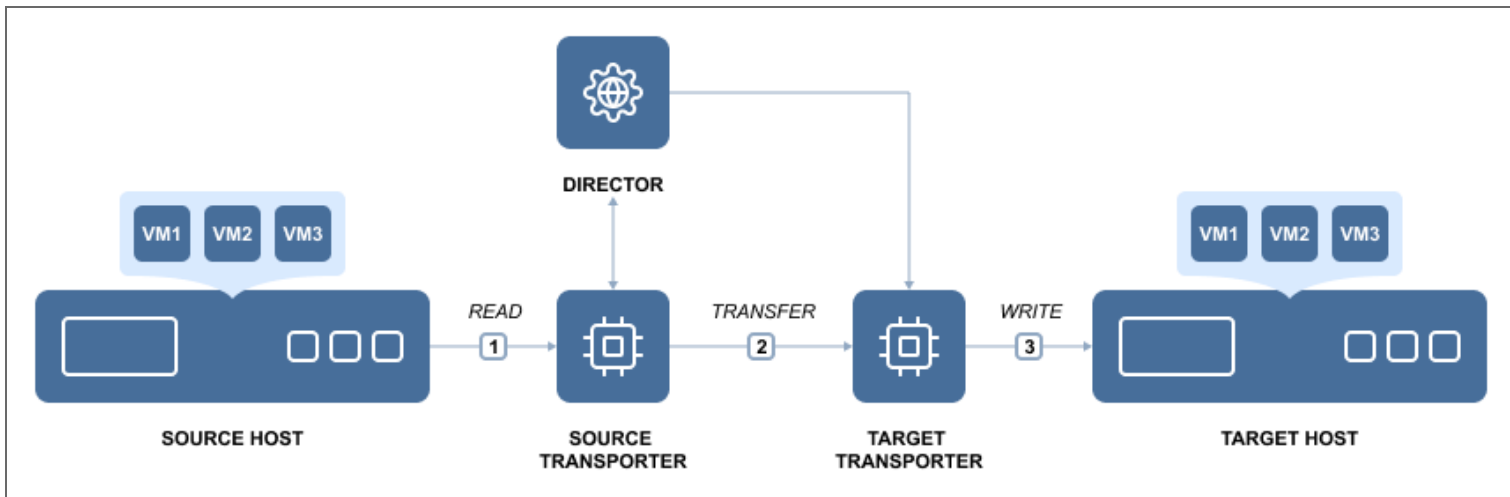
## Replication Types

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.



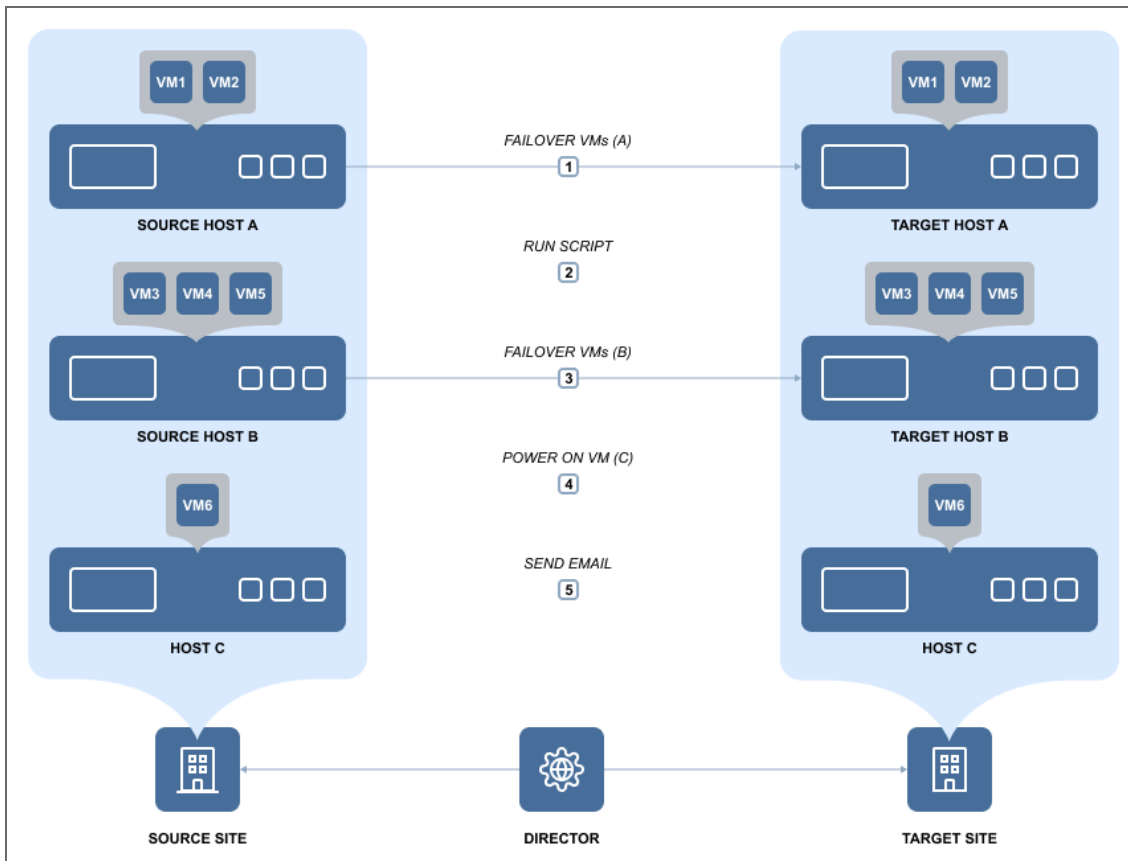
In the case of same region replication, NAKIVO Backup & Replication first initiates the creation of snapshots of selected source volumes and then creates AMIs from the source instance configuration and created snapshots. In case of cross-region replication, NAKIVO Backup & Replication initiates the creation 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 714](#).

## Site Recovery

With [Site Recovery Jobs](#) introduced, NAKIVO Backup & Replication allows you to automate the execution of one or more actions. An action refers to a single task that can be included in a Site Recovery Job. Refer to [“Site Recovery Job Wizard: Actions” on page 1019](#) for the list of available actions.





Special Actions that are 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 1058](#) 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 1060](#) 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 :

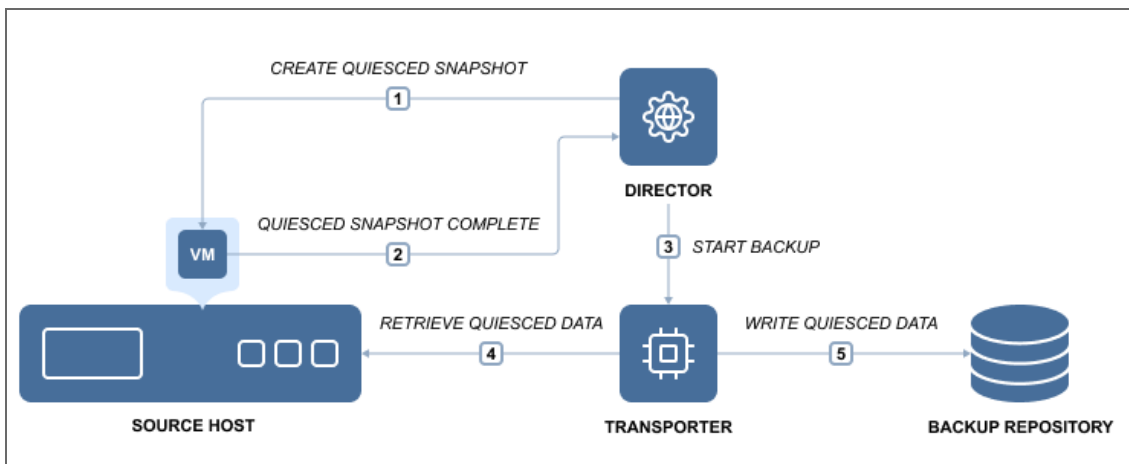
- [“Application and Database Support” below](#)
- [“Backup Immutability” on page 39](#)
- [“Backup Size Reduction” on page 39](#)
- [“Direct Connect” on page 42](#)
- [“Encryption in Flight and at Rest” on page 42](#)
- [“External Product Database Support” on page 44](#)
- [“Log Truncation” on page 44](#)
- [“Self-Backup Feature” on page 46](#)
- [“Two-Factor Authentication” on page 47](#)
- [“VM Verification” on page 48](#)

## 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](#).



## Application Awareness for Amazon EC2 Instances

NAKIVO Backup & Replication provides the ability to create application-consistent replicas of Amazon EC2 Instances. The application awareness approach is different for Windows-based and Linux-based Instances. Refer to the sections below for more information.

### Application Awareness for Windows-based Amazon EC2 Instances

NAKIVO Backup & Replication provides application awareness for supported Windows operating systems by using Microsoft Volume Shadow Copy (VSS) service. If application awareness is enabled for a Windows-based Instance, NAKIVO Backup & Replication does the following:

1. Connects to the source Instance using credentials provided.
2. Automatically deploys a NAKIVO VSS Driver into the Instance OS.
3. With the help of the VSS driver, creates application-consistent snapshots of the Instance's EBS volumes.
4. Automatically removes the VSS driver from the Instance OS.

### Application Awareness for Linux-based Amazon EC2 Instances

NAKIVO Backup & Replication provides the ability to run custom pre-freeze and post-thaw scripts on Linux-based Instances to enable application awareness. A pre-freeze script is executed before a snapshot of an Instance is taken, and post-thaw script is executed after the snapshot has been taken.

## Pre-Conditions for Application Awareness on Linux-based Instances

- Create pre-freeze and post-thaw scripts for your Linux-based Instances.
- Make sure that SSH is enabled on the Instances.
- Place the scripts in the appropriate folders on each Instance:
  - **Pre-freeze:** `/usr/sbin/pre-freeze-script`
  - **Post-thaw:** `/usr/sbin/post-thaw-script`

## Backup Immutability

When creating a backup job and selecting the **Amazon S3** 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 S3** type of Backup Repository, S3 Object Lock should be enabled in Amazon S3 for the bucket used to store backups. This type of immutability cannot be shortened or lifted, not even by the root user.

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 the immutability cannot be lifted or changed by anyone, not even the root user.

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

- [“Backup Job Wizard for VMware: Retention” on page 520](#)
- [“Backup Job Wizard for VMware Cloud Director: Retention” on page 541](#)
- [“Backup Job Wizard for Hyper-V: Retention” on page 563](#)
- [“Backup Job Wizard for Nutanix AHV: Retention” on page 606](#)
- [“Backup Job Wizard for Physical Machine: Retention” on page 626](#)
- [“Backup Copy Job Wizard: Retention” on page 679](#)
- [“Deploying VMware Virtual Appliance” on page 149](#)
- [“Deploying Nutanix AHV Virtual Appliance” on page 166](#)

For more details on the requirements for this feature, refer to [this page](#) in the User Guide.

## Backup Size Reduction

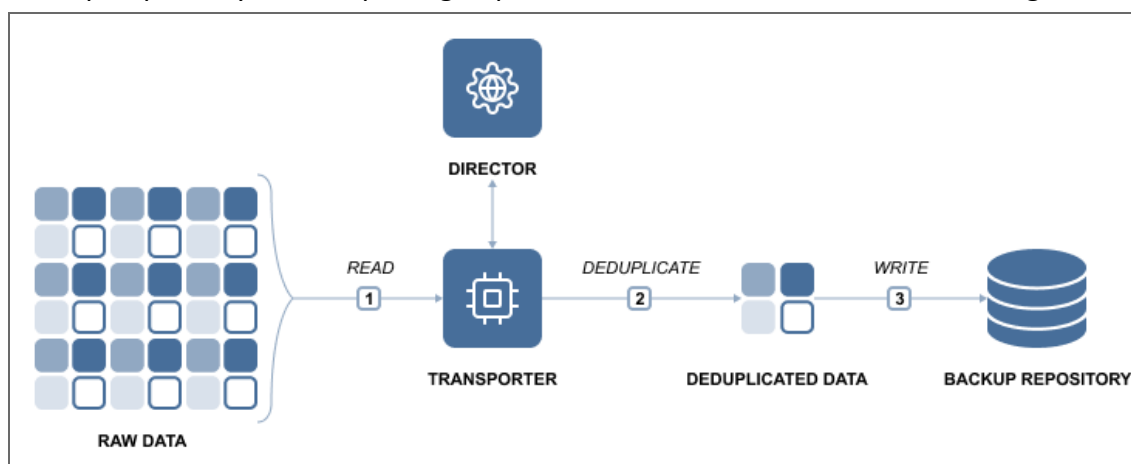
NAKIVO Backup & Replication utilizes multiple methods, such as deduplication and compression, to optimize the size of stored backups. The main purpose of these methods is to reach the correct balance between the amount of data read and transferred during backup.

This section contains the following topics:

- [“Global Data Deduplication and Compression” below](#)
- [“Excluding Swap Files and Partitions” on the next page](#)
- [“Excluding Unused Blocks” on the next page](#)

## Global Data Deduplication and Compression

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.

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.

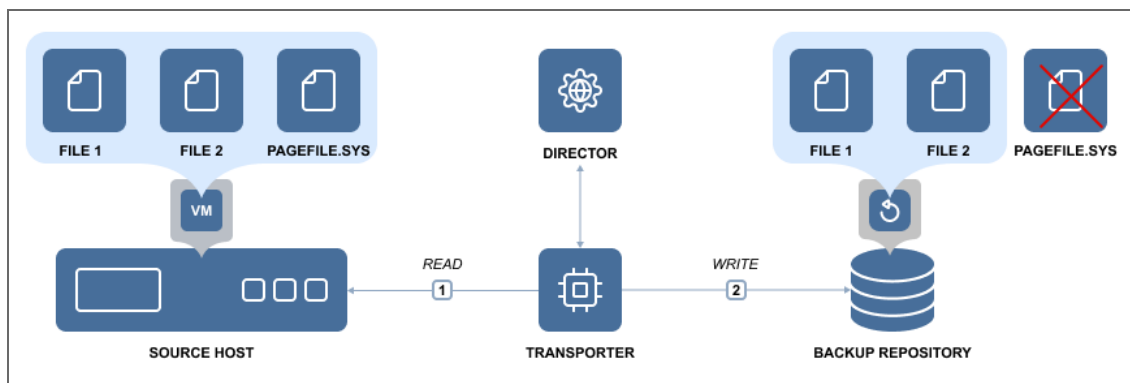
More efficient disk space utilization allows for storing more recovery points per VM backup. In addition, lower storage space requirements save money on direct storage costs (as fewer disks are needed to store the same amount of information) and on related costs (such as cooling, electricity, and maintenance).

## 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.

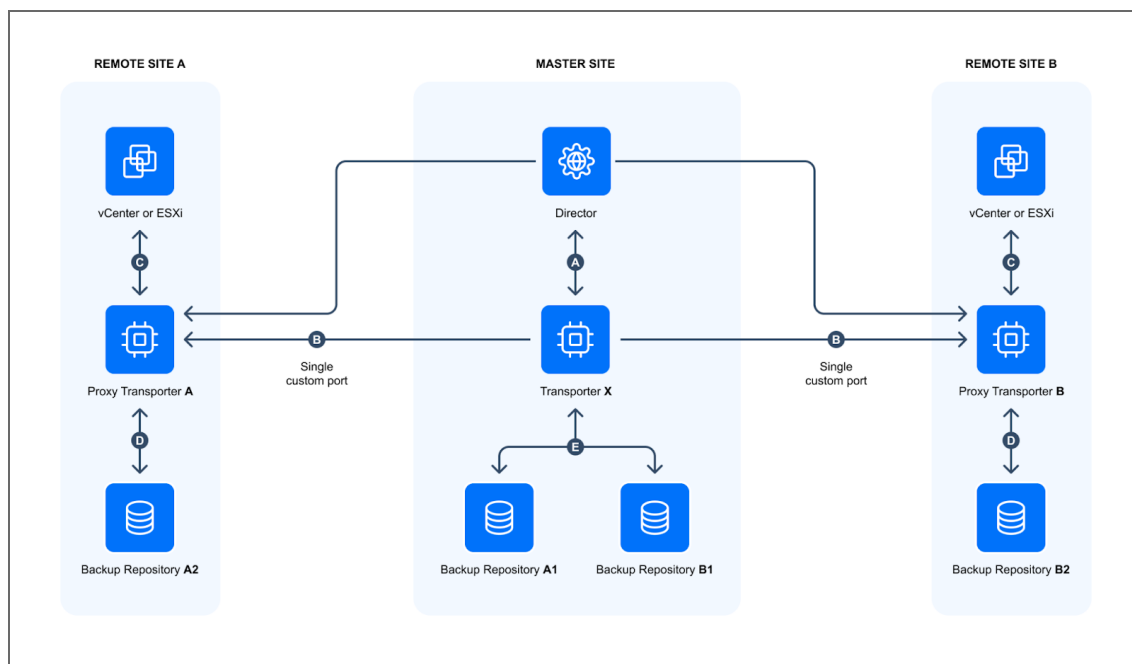


This option can be configured on a per-job basis on the **Options** page of backup and replication jobs and is enabled by default. The feature supports processing source objects running on Windows OS. It is available for the NTFS file system.

## Direct Connect

NAKIVO Backup & Replication now allows access to remote resources via a single port connection without the need to establish a VPN connection.

This feature is especially useful for MSPs providing services to remote clients in case having a persistent VPN connection between the MSP datacenter and the client site is not desired.




- **Proxy Transporter** (used within the context of this feature) - a Transporter deployed at a remote environment that has Direct Connect feature enabled.

The feature can be enabled by following the instructions in [this article](#). Additionally, make sure to check [feature requirements](#).

## 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.

## 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 418](#)
- [“Backup Repository on CIFS Share” on page 402](#)
- [“Backup Repository on NFS Share” on page 413](#)
- [“Backup Repository in Amazon EC2” on page 391](#)
- [“Backup Repository on Deduplication Appliance” on page 407](#)

# 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. For more information, refer to the following articles:

- [Database Options](#)
- [Troubleshooting External Database Connection Issues](#)

## Log Truncation

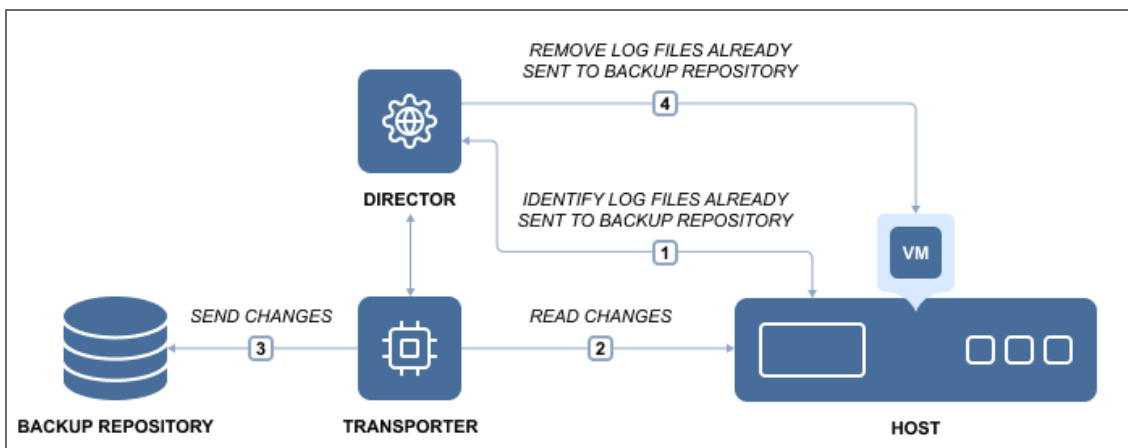
With NAKIVO Backup & Replication, you can remove (truncate) transaction log files of Microsoft Exchange and Microsoft SQL servers which will allow you to reduce the size of backups and, as a result, to optimize the use of storage space. Log truncation can be enabled on the **Options** page of backup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

# 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 297](#).

## 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:

- [“Configuring Two-Factor Authentication” on page 315](#)
- [“Logging in to NAKIVO Backup & Replication” on page 239](#)

# 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 VM 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 VM 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 VM Boot Job](#)
- [Hyper-V VM Backup Job](#)
- [Hyper-V VM Replication Job](#)

- [Hyper-V Flash VM 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” below](#)
- [“Backup from HPE 3PAR and HPE Nimble Storage Snapshots” on page 55](#)
- [“Deduplication Appliance Support” on page 56](#)
- [“Full Synthetic Data Storage” on page 58](#)
- [“Incremental Jobs” on page 59](#)
- [“Jobs and Concurrent Tasks” on page 60](#)
- [“LAN-Free Data Transfer” on page 61](#)
- [“Network Acceleration” on page 63](#)

## 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](#)
- [Distributing Bandwidth Between Tasks](#)


## About Bandwidth Rules

A bandwidth rule specifies the bandwidth amount that can be used by one job, by multiple jobs, or by all applicable jobs.


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. In case multiple per job rules are active for the same job, the bandwidth rule with the lowest bandwidth amount will be applied. In case there are multiple global rules – and no per job bandwidth rules, – the global rule with the lowest bandwidth amount will be applied. When a NAKIVO Backup & Replication job is running and a bandwidth rule is applied to this job, the job will get bandwidth amount that is allowed by the bandwidth rule (for example 10 Mbit/s).



When a NAKIVO Backup & Replication job is running with a bandwidth rule applied and the bandwidth rule becomes disabled for this job – and there are no other bandwidth rules applied to the job, – the job will get unlimited bandwidth.



Bandwidth rules may be always active, active on schedule, or disabled. Refer to [“Bandwidth Throttling” on page 284](#) for details.

When a job containing multiple VMs starts running with a bandwidth rule active, the rule divides bandwidth between tasks. Incremental backup tasks receive significantly less bandwidth than full backup tasks; this ensures that no tasks receive too little bandwidth to be processed in a reasonable time. When the Transporter is ready and there is enough unallocated bandwidth, the tasks start to be processed. Any change to the bandwidth amount will only be applied to the tasks not yet started for processing. Once started for processing, the tasks do not change the consumed bandwidth amount. It means there will be no dynamic change in the bandwidth amount for the tasks already being processed.

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

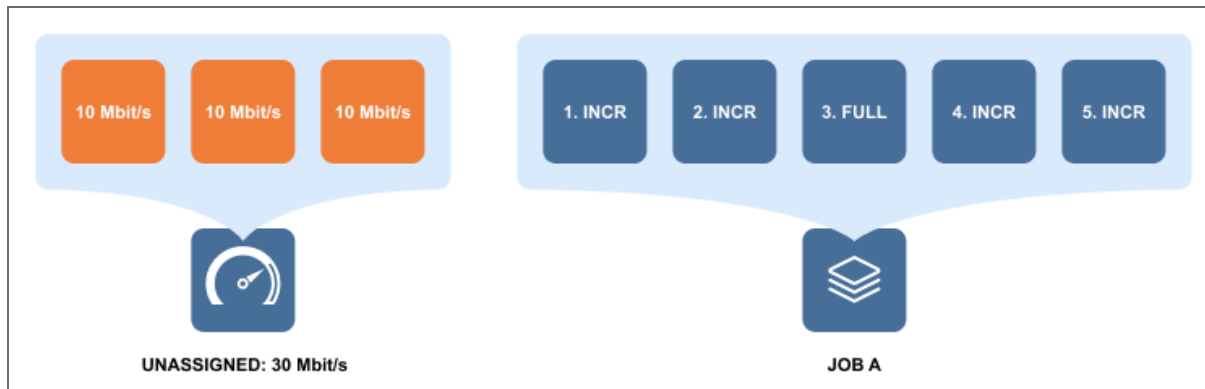
- Backup Job
- Backup Copy Job
- Replication Job, except Amazon EC2 Replication Job
- Recovery Job
- Failover Job

## Distributing Bandwidth Between Tasks


To illustrate distribution of bandwidth between tasks, one can take a backup job – Job A, – of 5 VMs; the 3<sup>rd</sup> VM backup is a full backup and the rest are incremental backups.

Job A starts running with the 30 Mbit/s bandwidth rule activated as follows:

1. The bandwidth amount is split into 3 chunks 10 Mbit/s each.



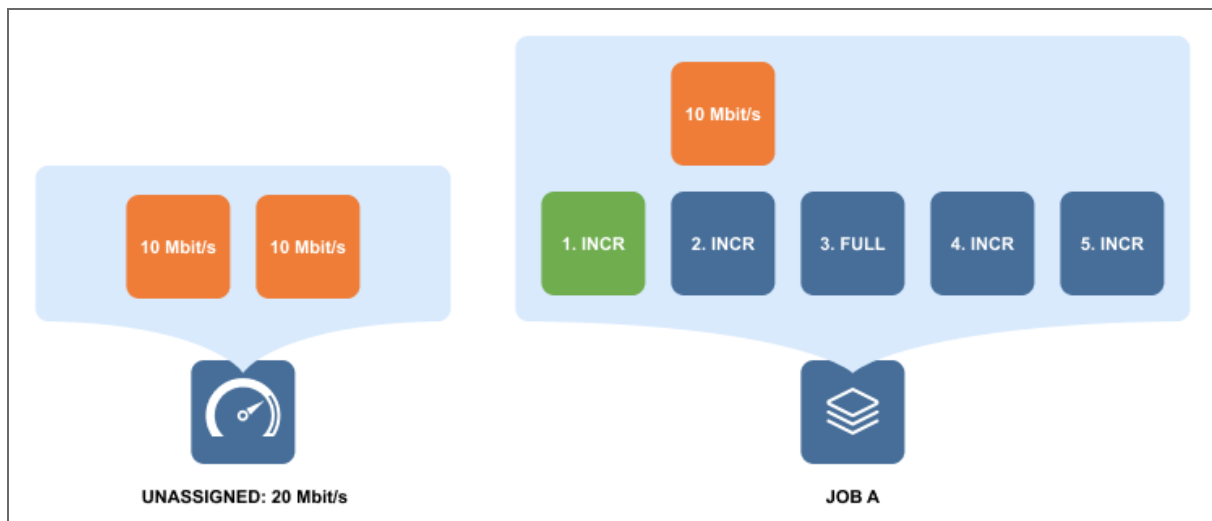
2. VM 1 and VM 2 backups receive 10 Mbit/s each. One bandwidth chunk remains unassigned since the full backup usually requires all the bandwidth to start.
3. The remaining bandwidth is distributed from the start of the queue, so VM 1 backup receives additional 10 Mbit/s.
4. VM 1 backup and VM 2 backup start running.



**Note**

The Transporter can process a [limited number](#) of concurrent tasks.

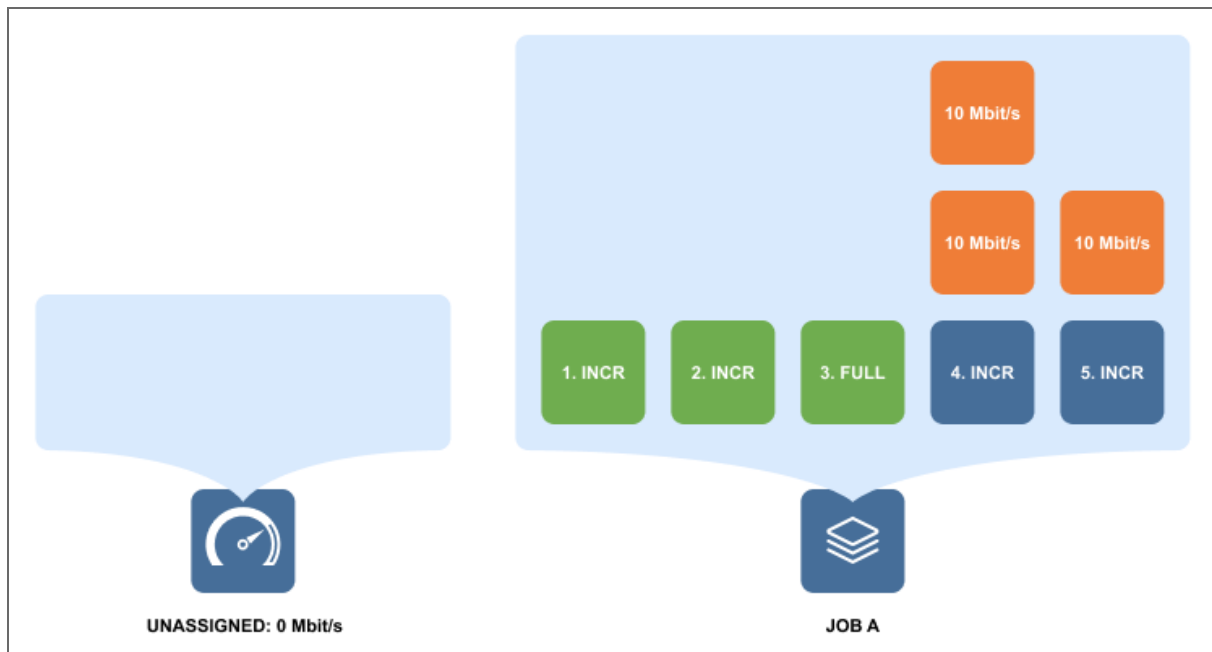
- 5. When VM 1 backup finishes execution, it frees two bandwidth chunks 10 Mbit/s each. However, VM 3 full backup still cannot start because it requires all the available bandwidth to start running. Hence, these two bandwidth chunks are left idle.



- 6. When VM 2 backup finishes running, it frees another bandwidth chunk, and full backup of VM 3 starts running with all the bandwidth assigned.

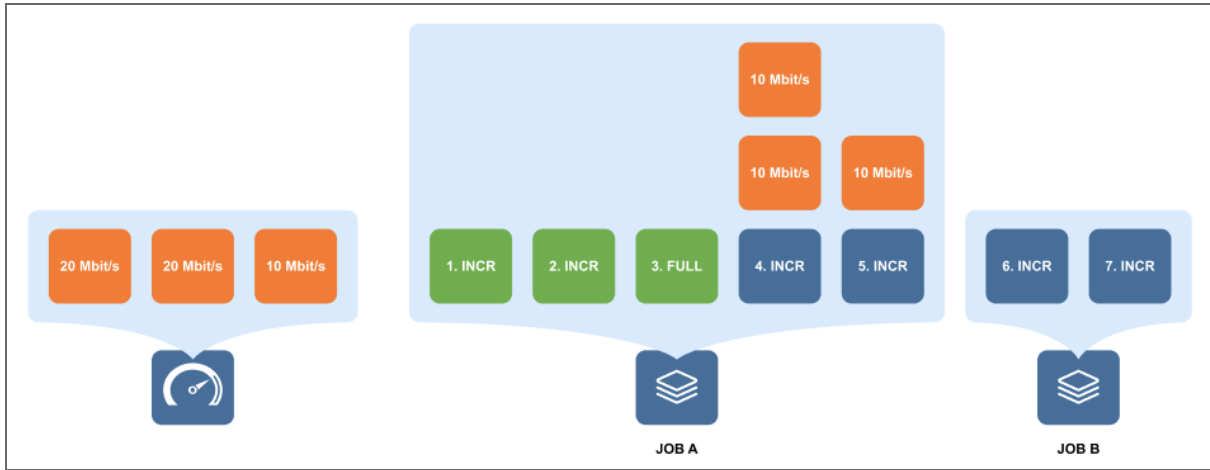


7. When full backup of VM 3 is finished, three bandwidth chunks are now available for the two remaining VM backups.
8. VM 4 backup receives the 20 Mbit/s bandwidth in total and VM 5 backup receives a 10 Mbit/s bandwidth chunk.




When the rule changes the bandwidth to 80 Mbit/s and is also activated for another Job B consisting of two VM incremental backups, the Transporter starts distributing bandwidth as follows:

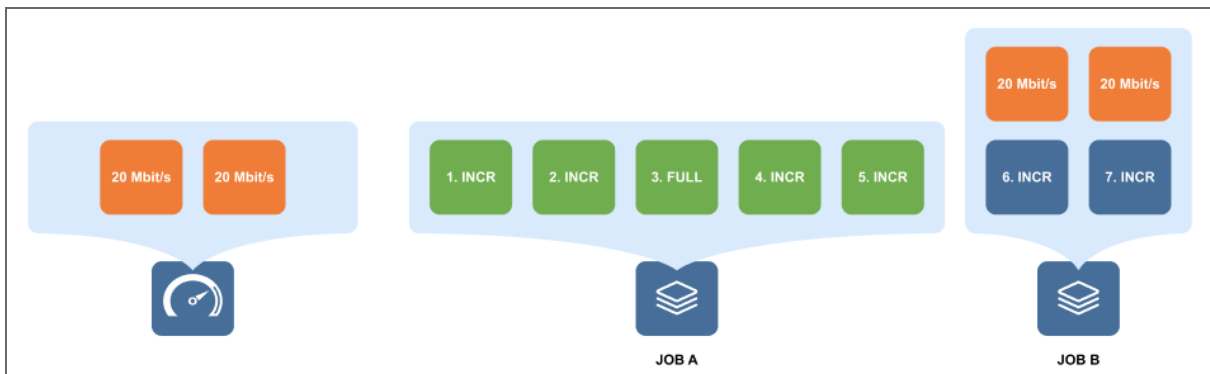
1. The 80 Mbit/s amount is split into 4 chunks of 20 Mbit/s.



2. VM 6 backup and VM 7 backup of Job B receive a 20 Mbit/s bandwidth chunk each and start running, with 10 Mbit/s remaining unassigned.




3. When VM 4 backup and VM 5 backup of Job A are finished, two 20 Mbit/s bandwidth chunks are freed. However there are no queued tasks to assign them to, so the bandwidth is left idle.



When the bandwidth rule changes the bandwidth amount back to 30 Mbit/s and is also activated for another Job C consisting of one VM incremental backup, the Transporter starts distributing bandwidth as follows:


1. The 30 Mbit/s amount is split into three chunks of 10 Mbit/s.
2. The currently running tasks occupy 40 Mbit/s of bandwidth, which is three 10 Mbit/s bandwidth chunks and one 10 Mbit/s bandwidth chunk over the limit. Therefore, there is no free bandwidth for VM 8 backup of Job C to use.



**Note**

Jobs and tasks may wait for a long time until bandwidth is available for them to start.

3. When VM 6 backup is finished, freeing up 20 Mbit/s of bandwidth, of which 10 Mbit/s was exceeding the 30 Mbit/s limit, VM 8 backup of Job C starts executing using another 10 Mbit/s bandwidth chunk.

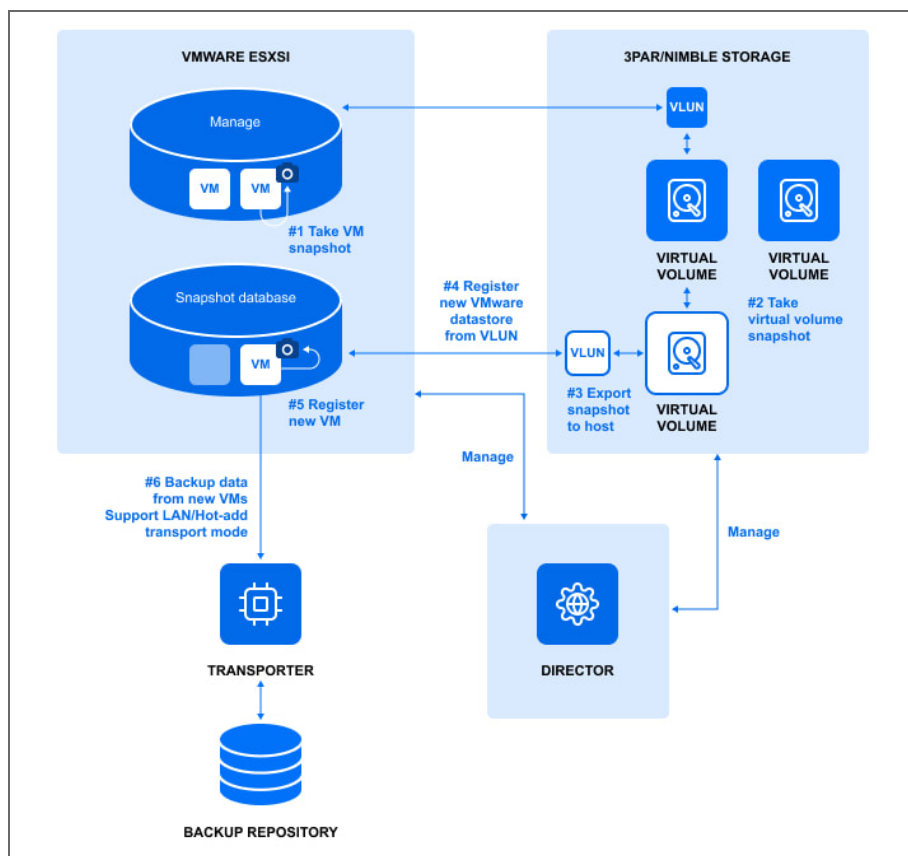


## Backup from HPE 3PAR and HPE Nimble Storage Snapshots

NAKIVO Backup & Replication allows you to use HPE 3PAR StoreServ and HPE Nimble 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. The backup from storage snapshot process 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 storage snapshots, do the following:

1. Add the [HPE 3PAR storage device](#) or [HPE Nimble storage device](#) to the inventory.
2. During [VMware Backup Job creation](#), add VMs that have their disks residing on the supported HPE 3PAR/Nimble 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

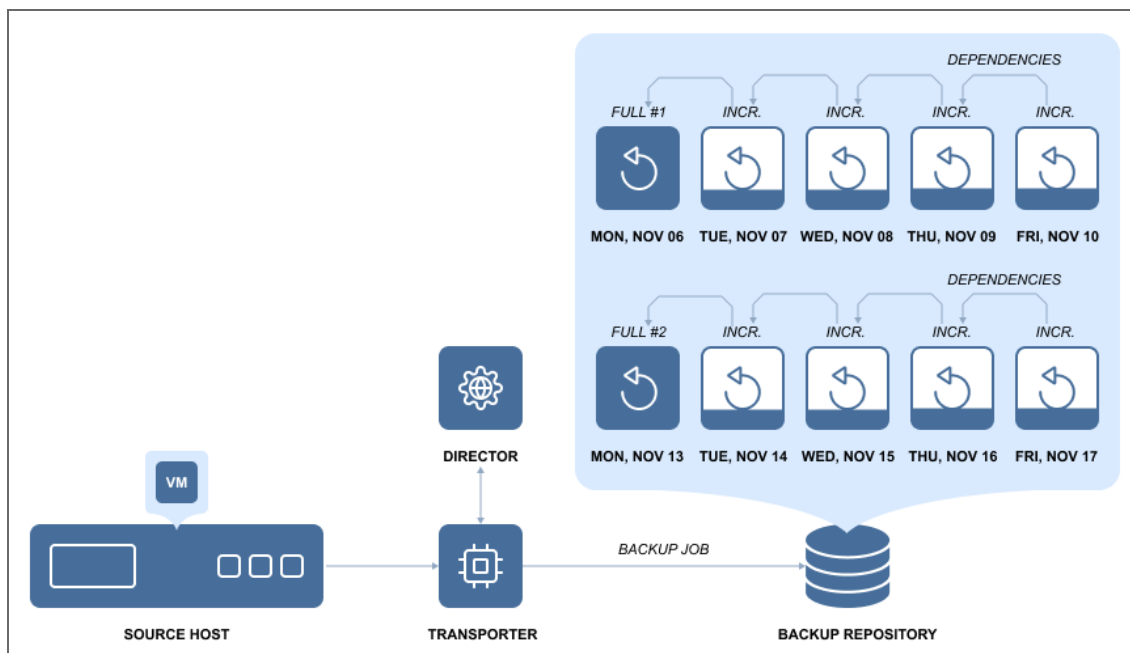
Deduplication appliances are solutions that implement specialized data reduction techniques to eliminate duplicate copies of repeated data. Deduplication appliances are leveraged across a range of data protection solutions, regardless of whether network-attached storage, disk, and/or tape is used. The 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 Transporters” on page 381](#) 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” below](#)
- [“Hot Add for VMware” on the next page](#)

## 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.

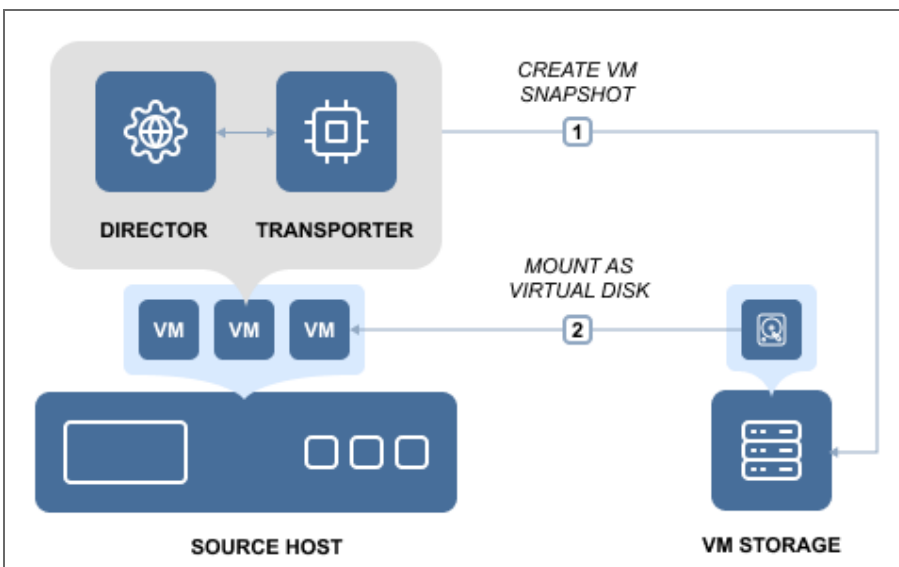


NAKIVO Backup & Replication relies on the VMware VDDK library that provides SAN support and does not make any write/modification operations on the SAN LUNs. That is to say, the product will not compromise data integrity or interfere with the data of running VMs on your SAN datastores.

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.

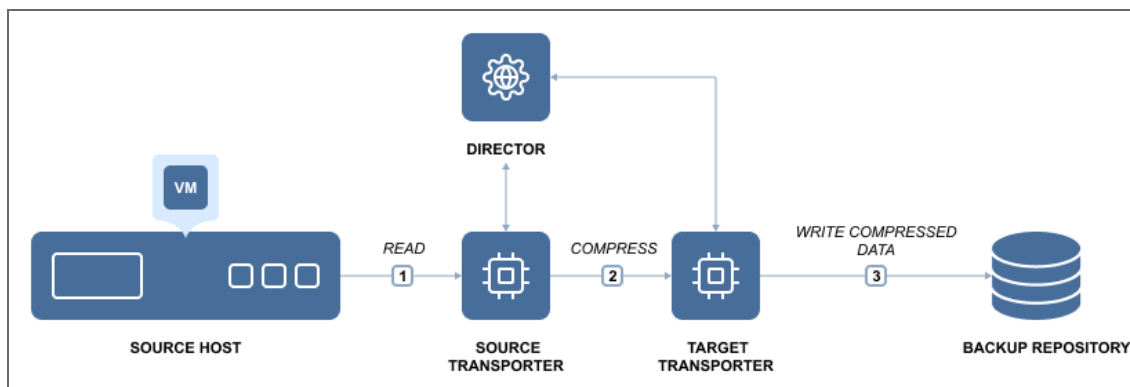
# Network Acceleration

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.

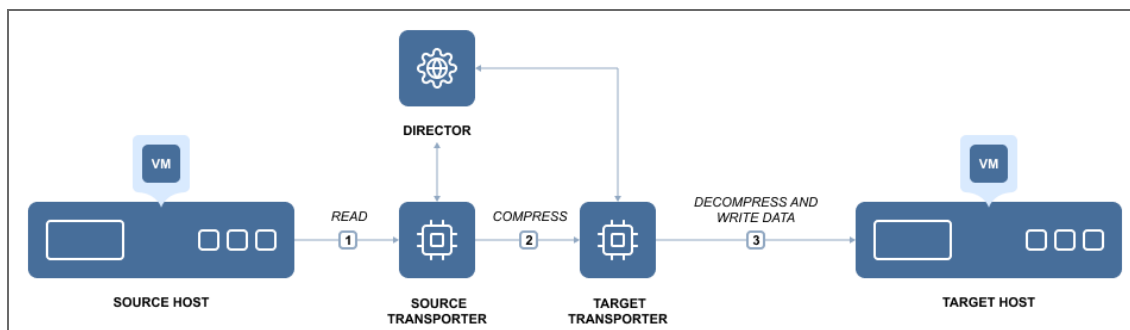
Network acceleration is achieved by the use of two instances of [Transporter](#). Transporter is the product component that performs all data protection and recovery tasks, such as backup, replication, recovery, encryption, and so on. To simplify deployment and configuration, one instance of Transporter 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



# 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” below](#)
- [“Global Search” on the next page](#)
- [“Monitoring” on page 66](#)
- [“Policy-Based Data Protection” on page 66](#)

## 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.



## 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

The monitoring dashboard allows you to track the resources of your VMware hosts, datastores, and VMs. This allows for monitoring of CPU, memory, and disk usage over a specified period of time to make sure the infrastructure has enough resources to meet organizational demands. Refer to the following topics for additional details:

- [Navigating The Monitoring Dashboard](#)
- [Managing The Monitored Items](#)

## 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 266](#)
- [“Managing Policy Rules” on page 269](#)

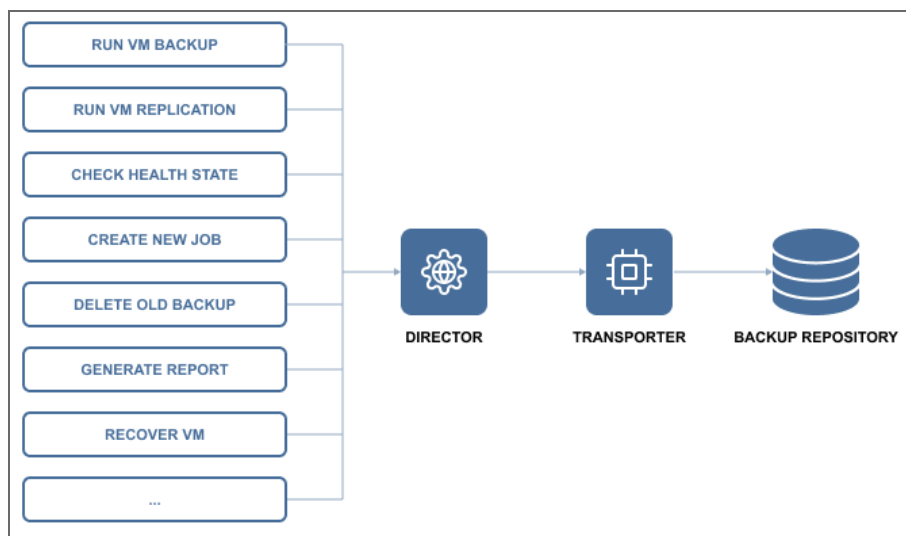
## Automation

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

- [“HTTP APIs” below](#)
- [“Job Chaining” on the next page](#)
- [“Pre and Post Job Scripts” on the next page](#)

## 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.



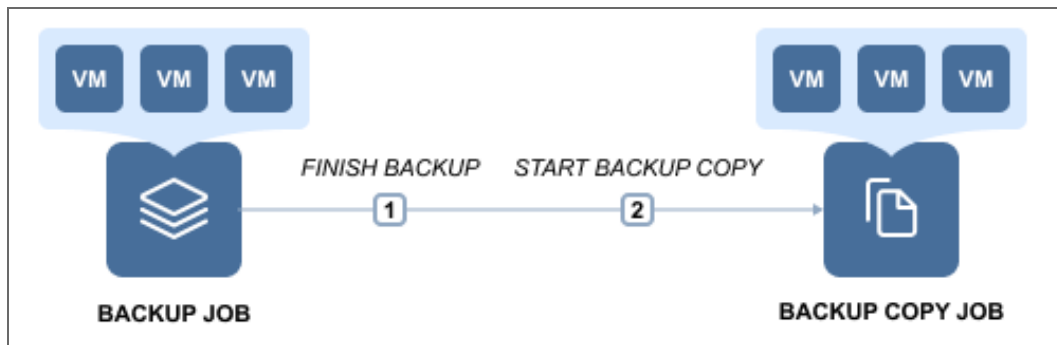
The API allows you to easily integrate NAKIVO Backup & Replication with monitoring, automation, and orchestration solutions to reduce time spent on backup management and reduce data protection costs. To speed up integration time, the API comes as part of an Integration Kit, which includes API documentation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

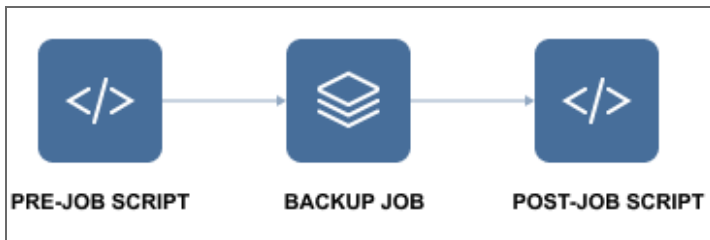
Job Chaining allows you to link jobs so that they run one directly after another. For example, you can set up a VM backup job that saves backups locally and then starts a [Backup Copy job](#) that copies the newly created backups to Amazon cloud.



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 HYDRAsstor. 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 HYDRAsstor 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” below](#)
- [“EMC DD Boost” below](#)
- [“HPE 3PAR” on the next page](#)
- [“HPE Nimble Storage” on the next page](#)
- [“HPE StoreOnce Catalyst” on the next page](#)
- [“NEC HYDRAsstor” on page 72](#)

## 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 318](#)
- [“Managing Active Directory Users” on page 317](#)

## 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 102](#)
- [Integrating with EMC DD Boost](#)
- [“Backup Repository on Deduplication Appliance” on page 407](#)

## HPE 3PAR

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 HPE 3PAR and HPE Nimble Storage Snapshots” on page 55.](#)

## 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 HPE 3PAR and HPE Nimble Storage Snapshots” on page 55.](#)

## 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 56](#)
- [“Storage Integration Requirements” on page 102](#)
- [“Backup Repository on Deduplication Appliance” on page 407](#)

## NEC HYDRAsstor

HYDRAsstor 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 HYDRAsstor 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 HYDRAsstor 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.

HYDRAsstor 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.

HYDRAsstor 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 HYDRAsstor, refer to the [NEC official website](#).

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

- [“Storage Integration Requirements” on page 102](#)
- [Integrating with NEC HYDRAsstor](#)
- [“Backup Repository on Deduplication Appliance” on page 407](#)

## BaaS

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

This section contains the following topics:

- [“Branding” on the next page](#)
- [“License Delegation” on the next page](#)

- [“Multi-Tenancy” on the next page](#)
- [“Self-Service” on page 75](#)

## 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.

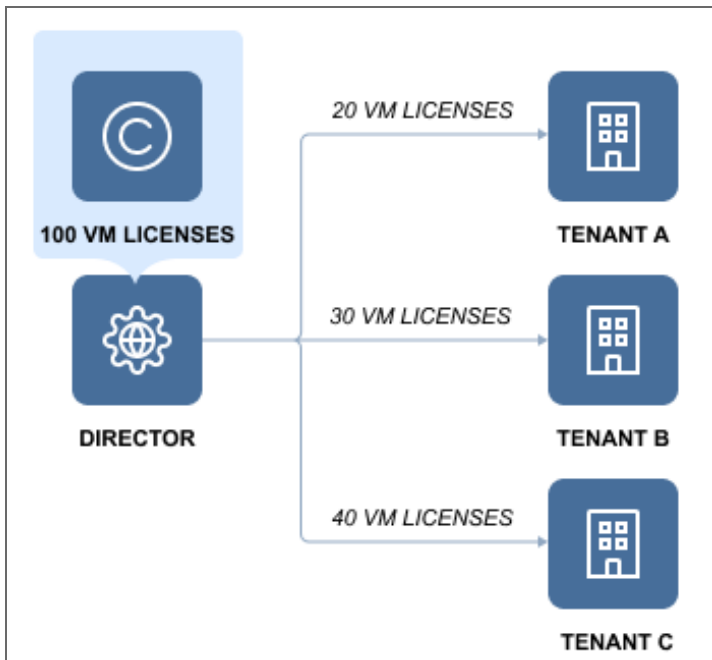
For information on branding configuration, refer to [“Branding” on page 287](#).

## 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.

## Multi-Tenancy

Multi-tenancy enables you to create and manage up to 1,000 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](#)
- [Licensing for Backup for Microsoft 365](#)
- [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, and Nutanix AHV virtual machines (VMs), 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 three 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.

**Note**

Valid support is required to receive product updates.

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
<b>Pro Essentials</b>	VMware vSphere	Min. 2 sockets Max. 6 sockets	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		
	Windows/Linux Physical Machines	Min. 10 servers Max. 50 servers	
		Min. 10 workstations Max. 150 workstations	
NAS	Min. 1 TB Max. 50 TB		
<b>Enterprise Essentials</b>	VMware vSphere	Min. 2 sockets 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		
	Windows/Linux Physical Machines	Min. 10 servers Max. 50 servers	
		Min. 10 workstations Max. 150 workstations	
NAS	Min. 1 TB Max. 50 TB		

<b>Pro</b>	VMware vSphere	No limits	Includes most product features with limitations on backup to the cloud, administrative tools, and BaaS
	Nutanix AHV		
	Microsoft Hyper-V		
	Windows/Linux Physical Machines		
	NAS		
<b>Enterprise</b>	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		
	Windows/Linux Physical Machines		
	NAS		
<b>Enterprise Plus</b>	VMware vSphere	No limits	The most complete edition of NAKIVO Backup & Replication
	Nutanix AHV		
	Microsoft Hyper-V		
	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, 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
<b>Pro Essentials</b>	VMware vSphere	Min. 5 workloads 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		
	Amazon EC2		
	Windows/Linux Physical Machines		
	NAS		
<b>Enterprise Essentials</b>	VMware vSphere	Min. 5 workloads 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		
	Amazon EC2		
	Windows/Linux Physical Machines		
	NAS		
<b>Pro</b>	VMware vSphere	No limits	Includes most product features with limitations on backup to the cloud, administrative tools, and BaaS
	Nutanix AHV		
	Microsoft Hyper-V		
	Amazon EC2		
	Windows/Linux Physical Machines		
	NAS		

<b>Enterprise</b>	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		
	Amazon EC2		
	Windows/Linux Physical Machines		
	NAS		
<b>Enterprise Plus</b>	VMware vSphere	No limits	The most complete edition of NAKIVO Backup & Replication
	Nutanix AHV		
	Microsoft Hyper-V		
	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		
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

Platform	License Unit (Workload)	Pro Essentials/Enterprise Essentials Editions Limits
VMware vSphere	1 VM	Minimum of 5 workloads Maximum of 50 workloads
Microsoft Hyper-V		
Nutanix AHV		
Amazon EC2	1 Instance	
Windows/Linux Physical Server	1 Server	
Windows/Linux Workstation	3 Workstations	
NAS File Share	0.5 Terabyte	
Oracle Database	1 Database	

### 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**.

## 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, and/or SharePoint Online. 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.

## 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.

## 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
- Number of license units (sockets or VMs)
- The support end date
- The type of support (Standard or 24/7)
- License unit limits for the Pro Essentials/Enterprise Essentials editions: 2-6 sockets

## 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.

## 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” below](#)
- [“System Requirements” on page 96](#)
- [“Deployment Scenarios” on page 137](#)
- [“Installing NAKIVO Backup & Replication” on page 148](#)
- [“Updating NAKIVO Backup & Replication” on page 214](#)
- [“Uninstalling NAKIVO Backup & Replication” on page 235](#)

## 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” below](#)
- [“Transporter” on page 89](#)
- [“Backup Repository” on page 94](#)

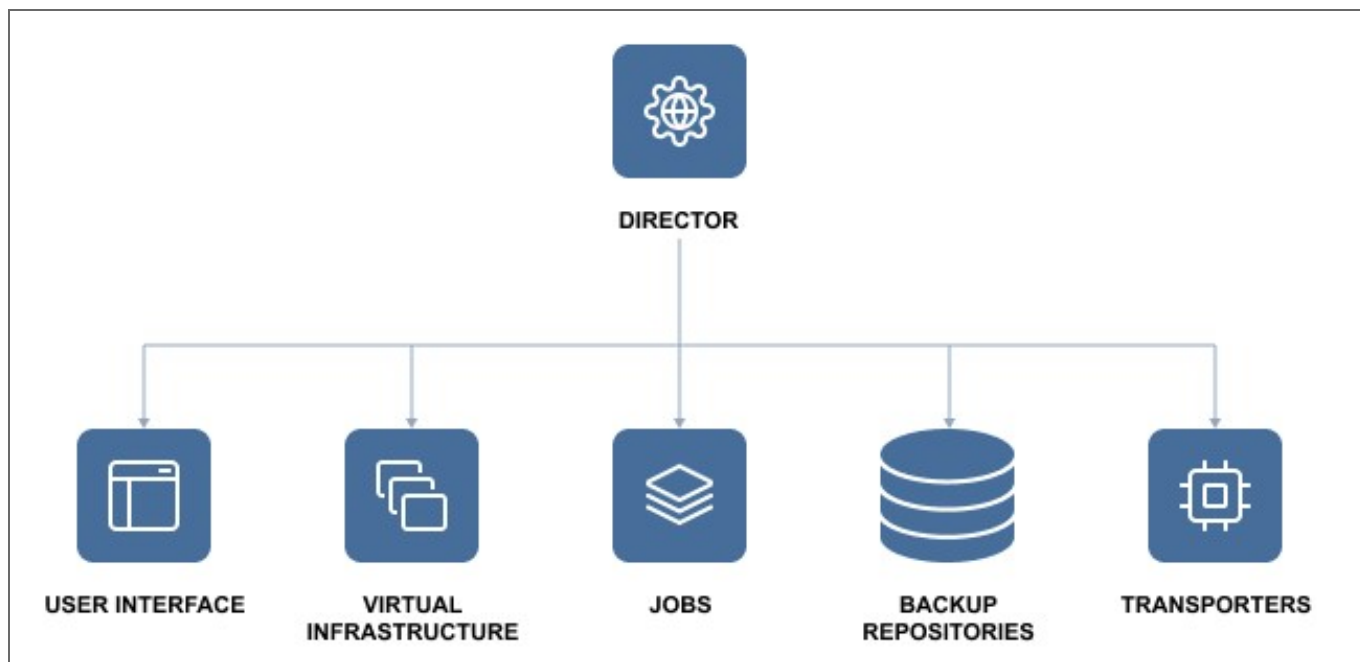
All components can be installed on a single machine or can be distributed across multiple machines and geographical locations.

## 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.




## Transporter

- [What is Transporter?](#)
- [How many Transporters Should be Deployed?](#)
- [How Transporters are Selected for Jobs](#)
- [Transporter Security](#)
- [Transporter as a VMware Appliance](#)
- [Transporter as a Nutanix AHV Appliance](#)

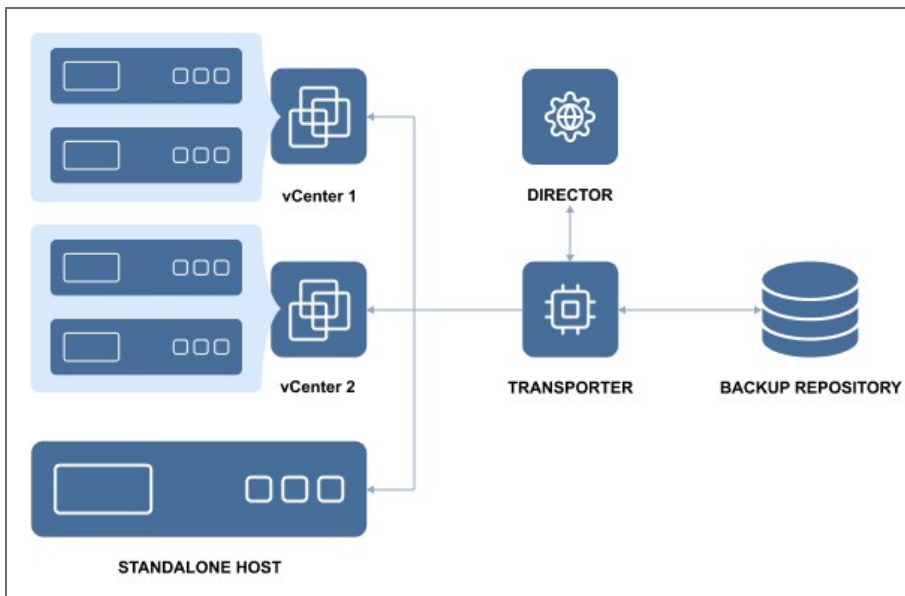
## What is Transporter?

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.

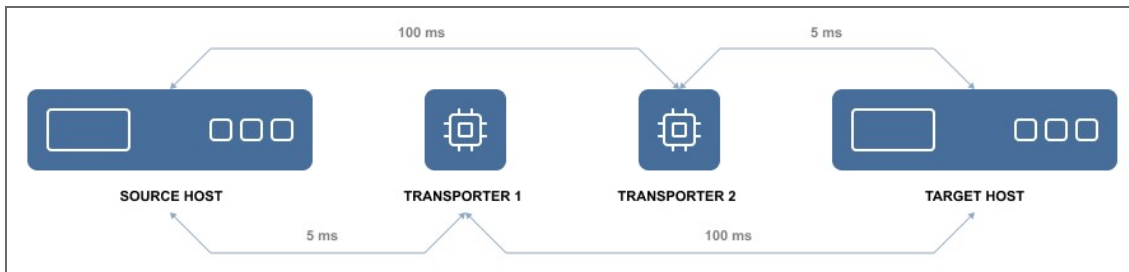


If you plan to transfer data over WAN without a VPN connection from your source site to the target site, make sure the source and target Transporters are added to the product using external IP addresses or DNS names that can be properly resolved in WAN, so that the two Transporters can connect to each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

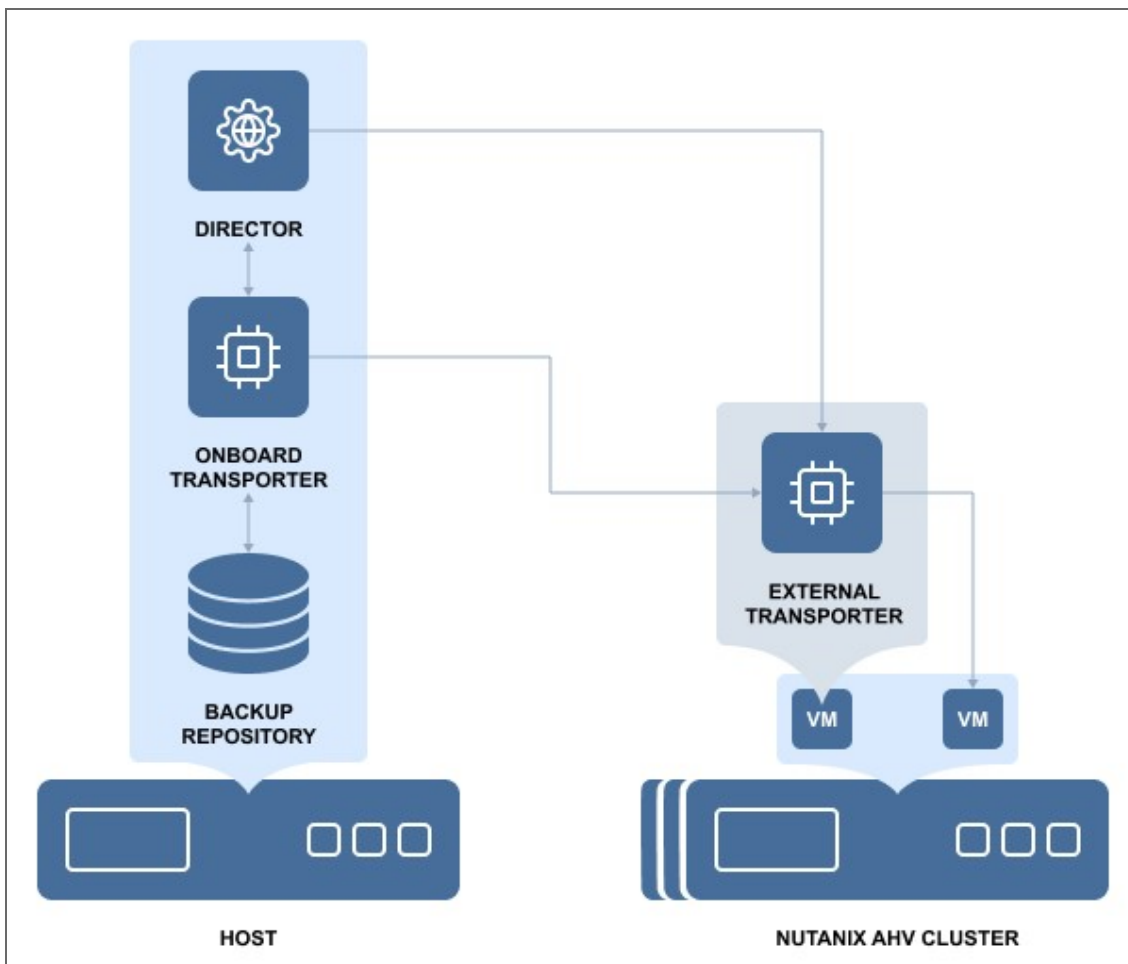
In this case, you will need to [deploy an additional Transporter as a VMware appliance](#) to allow NAKIVO Backup & Replication to work with VMware vCenters and ESXi hosts, and protect your virtual infrastructure.



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

## Transporter as a Nutanix AHV Appliance

Nutanix AHV support is done via deploying a Nutanix Transporter to a Nutanix cluster to read/write data from/to VMs in Nutanix cluster. For more information, refer to Adding Installed Transporters and Deploying Transporter as Nutanix AHV Appliance sections.



## 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?

NAKIVO Backup & Replication can store up to 128 TB of data in a single Backup Repository. The number of Backup Repositories 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.

# 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 the next page](#)
- [“Storage Integration Requirements” on page 102](#)
- [“Deployment Requirements” on page 104](#)
- [“Feature Requirements” on page 124](#)

# Supported Platforms

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

- VMware vSphere v5.5 - v7.0.3
- VMware Cloud Director v10.2.1 - v10.2.2
- Microsoft Server with a Hyper-V role 2022 (21H2), 20H2, 2019, 2016, 2012R2, 2012, 20H1
- Amazon EC2
- Nutanix AHV v5.10, v5.15, 5.20 (LTS)
- Microsoft Windows Server 2022, 20H2, 2019, 2016, 2012R2, 2012, 2008R2
- Microsoft Windows 10 Pro, Windows 10 Home, Windows 11
- Linux Servers and workstations (see [Physical machine requirements](#))
- Microsoft 365 (Exchange Online, OneDrive for Business, SharePoint Online)
- Oracle Database (RMAN)
- File shares with the following protocols: NFS 3.x, SMB 2.x, SMB 3.x

## Notes

- To learn about 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 the NAKIVO Backup & Replication, make sure that your system has been updated with the latest patch and all the necessary requirements are met.

Find the necessary requirements below:

- [Hypervisor Requirements](#)
  - [VMware vSphere](#)
  - [Microsoft Hyper-V](#)
  - [Nutanix AHV](#)
- [Physical Machine Requirements](#)
- [Public Cloud Requirements](#)
- [Microsoft 365 Requirements](#)
- [Oracle Database 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.

### Microsoft Hyper-V

- Full GUI installation of Windows Server is required.
- Full administrative permissions are required.
- Default administrative shares must be enabled on Hyper-V hosts.



- The "File server" role must be enabled on Hyper-V hosts.
- The SMB 2 protocol should be enabled on Hyper-V hosts.

## Nutanix AHV

- User admin or cluster admin permissions are required.

## Physical Machine Requirements

To provide data protection for Windows physical machines, make sure they meet the following hardware and software requirements:

### Hardware

- **CPU:** x86\_64
- **RAM:** At least 1 GB
- **Firmware:** BIOS or UEFI
- **Partition table:** MBR or GPT

### Software

- Supported Windows operating systems:
  - Microsoft Windows Server 2022 (21H2) (x64)
  - Microsoft Windows Server 20H2 (x64)
  - Microsoft Windows Server 2019 (x64)
  - Microsoft Windows Server 2016 (x64)
  - Microsoft Windows Server 2012R2 (x64)
  - Microsoft Windows Server 2012 (x64)
  - Microsoft Windows Server 2008R2 (x64)
- Supported workstations:
  - Windows 11 (x64)
  - Microsoft Windows 10 Pro (x64)
  - Microsoft Windows 10 Home (x64)
- Supported Windows File Systems:
  - NTFS
  - ReFS
- Physical machines should be accessible over the network.
- Administrative credentials should be provided to the physical machine.
- PowerShell must be installed.
- SMBv2 or higher version of SMB protocol must be enabled. In case a firewall is enabled, the corresponding rule for SMB-in needs to be enabled too.
- Selected users should have permissions to "Log on as a batch job".
- Default administrative shares must be enabled on the physical machine and accessible over the network.
- Supported Linux operating systems:
  - Ubuntu v16.04-20.04 (x64)
  - RHEL v7.4-8.5 (x64)

- SLES v12 SP3 - v15 SP2 (x64)
- CentOS v7.0-8.4 (x64)
- CentOS Stream 8 (x64)
- Supported workstations:
  - Ubuntu 18.04 (Desktop) LTS, 64-bit
  - Ubuntu 20.04 (Desktop) LTS, 64-bit
- Supported Linux File Systems:
  - NTFS
  - Ext2
  - Ext3
  - Ext4
  - FAT32
  - XFS
  - Linux SWAP
  - ReiserFS
- 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.
  - 9446 – Used by NAKIVO Backup & Replication to communicate with the server.
- `openssh-server` package should be installed.
- `sshd` service should be running.
- `root` login over `ssh` should be enabled if you use the `root` user. Check the `/etc/ssh/sshd_config` file to have a line: `PermitRootLogin yes`. Refer to Linux vendor documentation for more details.
- `/etc/pam.d/bhsvc` file with special permissions provided to the Transporter service is required. Refer to the [Required Permissions for Linux Recovery Server](#) knowledge base article for details.
- `selinux` configuration should be set to disabled if present. Refer to the [File Recovery: Empty Disk](#) knowledge base article for details.
- `PasswordAuthentication` should be set to "Yes".
- If you add a physical machine to the NAKIVO Backup & Replication inventory with a non-root account, the following is required:
  - `sudo` must be installed on the physical machine.
  - Disable `requiretty` for non-root user accounts. Update the `etc/sudoers` file to have a line: `Defaults !requiretty`

## Public Cloud Requirements

Refer to [Required AWS IAM Permissions for Amazon EC2](#) for details.

### Note

There is also an option for granting full IAM permissions for the application.

# Microsoft 365 Requirements

To provide data protection for your Microsoft 365, the following requirements must be met:

- Exchange Online must be purchased as a part of the Microsoft 365 plan or a standalone service.
- The required API permissions must be provided to NAKIVO Backup & Replication. Refer to [Required API Permissions for Microsoft 365](#).
- For automatic enabling of the "Site Collection Administrator" role, a Microsoft 365 account user must have the "SharePoint admin" or "Global admin" role. This is required for backing up and recovery of SharePoint Online sites.
- Two-factor/multi-factor authentication must be disabled in SharePoint Online administrator account.
- Microsoft 365 account must be accessible over the network.
- Conditional Access must be disabled in Azure portal. Alternatively, the account must be added as "excluded" from the Conditional Access.
- To backup and restore Microsoft 365 services, a Transporter should be installed on one of the following operating systems:

## Windows

- Windows Server 2022 (21H2) (x64)
- Windows Server 20H2 (x64)
- Windows Server 2019 Standard (x64)
- Windows Server 2016 Standard (x64)
- Windows Server 2012 R2 Standard (x64)
  - Update for Windows Server 2012 R2 (KB3179574)
- Windows Server 2012 Standard (x64)
- Windows Server 2008 R2 Standard (x64)
- Windows 11 (x64)
- Windows 10 Pro (x64)
- Windows 8 Professional (x64)
- Windows 7 Professional (x64)

## Linux

- Ubuntu 20.04 Server (x64)
- Ubuntu 18.04 Server (x64)
- Ubuntu 16.04 Server (x64)
- SUSE Linux Enterprise Server 12 SP3 – 15 SP2 (x64)
- Red Hat Enterprise Linux v7.4 – v8.5 (x64)
- CentOS v7.0 – v8.4 (x64)

## NAS

Refer to [“Deployment Requirements” on page 104](#) to see the list of supported NAS devices.

## Note

ARM-based NAS devices and FreeNAS devices are not supported for backup and recovery of Microsoft 365 accounts.

## Oracle Database Requirements

To provide data protection for your Oracle database, make sure of the following:

- Oracle database is accessible over the network.
- Provided Oracle database credentials have sufficient permissions to be used by the product.
- Provided OS credentials for the Oracle database have sufficient permissions to be used by the product.
- Oracle database must be set to ARCHIVELOG mode to create backups.

Supported operating systems:

- Windows Server 2016 Standard (x64)
- Windows Server 2012 R2 Standard (x64)
  - Update for Windows Server 2012 R2
- Windows Server 2012 Standard (x64)
- Windows Server 2008 R2 Standard (x64)

Supported Oracle Database release versions:

- Oracle Database 19c
- Oracle Database 18c
- Oracle Database 12c R2
- Oracle Database 11g R2

## Cloud Region Requirements

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

- US East (Ohio)
- US East (N. Virginia)
- US West (N. California)
- US West (Oregon)
- Africa (Cape Town)
- Asia Pacific (Hong Kong)
- Asia Pacific (Mumbai)
- Asia Pacific (Seoul)
- Asia Pacific (Singapore)
- Asia Pacific (Sydney)
- Asia Pacific (Tokyo)
- Canada (Central)
- EU (Frankfurt)
- EU (Ireland)
- EU (London)
- EU (Milan)
- EU (Paris)

- EU (Stockholm)
- Middle East (Bahrain)
- 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 Central 1 (Amsterdam)
- Wasabi EU Central 2 (Frankfurt)
- Wasabi AP Northeast 1 (Tokyo)
- Wasabi AP Northeast 2 (Osaka)
- 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](#)

### Dell-EMC Data Domain

#### Supported versions:

- Dell-EMC Data Domain 6.1
- Dell-EMC Data Domain 6.2

- Dell-EMC Data Domain 7.0

## NEC HYDRastor

### Supported systems:

- NEC HYDRastor v5.5.1-5.6.0
- NEC Storage HS Universal Express I/O Module Version v1.8.0-1.8.4

## HPE StoreOnce with Catalyst Support

### Supported versions:

- HPE StoreOnce 3.18.18
- HPE StoreOnce 4.2.3

### 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
<b>VSA</b>			
VSA Gen 4 (128+sessions)	128-256	6	7
<b>HPE ProLiant Gen 10 (StoreOnce 4.2.3)</b>			
3620	128	6	7
3640	192	6	14
5200	512	10	21
5250	512	10	21
5650	1024	16	30
<b>HPE ProLiant Gen 9 (StoreOnce 3.18.18)</b>			
3500	192	6	14
5100	320	10	14
5500	1000	16	30
6600	1024	16	30
<b>HPE ProLiant Gen 8 (StoreOnce 3.18.18)</b>			

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 [“Backup from HPE 3PAR and HPE Nimble Storage Snapshots” on page 55.](#)

### 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 [“Backup from HPE 3PAR and HPE Nimble Storage Snapshots” on page 55.](#)

### Supported versions:

- HPE Nimble OS 5.0.2 and above

## 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.

- [Hardware](#)
  - [VM or Physical Machine](#)
  - [Network Attached Storage](#)
- [Operating Systems](#)
- [Networking Requirements](#)
  - [Required TCP Ports](#)
  - [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 100 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 100 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
  - For SaaS Backup Repository-related activities:
    - minimum total RAM: 4 GB
    - additional 100 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 100 MB for each concurrent Java Transporter task
- **Free space:** 5 GB

### Note

[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 120](#)
- **QNAP:** For a full list of supported models, refer to [“Supported QNAP NAS Devices” on page 111](#)
- **ASUSTOR:** For a full list of supported models, refer to [“Supported ASUSTOR NAS Devices” on page 110](#)
- **NETGEAR:** For a full list of supported. For a full list of supporter models, refer to [“Supported NETGEAR NAS Devices” on page 117.](#)



- **Western Digital:** For a full list of supported models, refer to [“Supported Western Digital NAS Devices” on page 123](#).

## Generic ARM-based NAS devices

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

- Single-board computer with ARMv7/ARMv8 CPU (e.g. Raspberry Pi 3 Model B+)
- 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

## Raspberry Pi

NAKIVO Backup & Replication can be installed on a Raspberry Pi 3 Model B+ machine with the following minimum hardware characteristics:

Director and Onboard Transporter:

- **RAM:** 1 GB + 250 MB for each concurrent task
- **Free space:** 16 GB

Transporter only:

- **RAM:** 512 MB + 250 MB for each concurrent task
- **Free space:** 16 GB

## 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.

### Windows

- 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 Server 2012 R2 Standard (x64)
- Microsoft Windows Server 2012 Standard (x64)
- Microsoft Windows Server 2008 R2 Standard (x64)
- Microsoft Windows 11 (x64)
- Microsoft Windows 10 Home (x64)

- Microsoft Windows 10 Professional (x64)
- Microsoft Windows 8 Professional (x64)
- Microsoft Windows 7 Professional (x64)

## **Linux**

- Ubuntu 20.04 Server (x64)
- Ubuntu 18.04 Server (x64)
- Ubuntu 16.04 Server (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 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 8 (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)

## **NAS**

- ASUSTOR ADM v3.5
- ASUSTOR ADM v4.0
- FreeNAS 11.3
- Netgear ReadyNAS OS v6.10.3
- Netgear ReadyNAS OS v6.9
- Synology DSM v6.0-v7.1
- QNAP QTS v4.3-v5.0
- QNAP QuTS Hero h4.5.3-v5.0
- WD MyCloud v5
- TrueNAS CORE 12

### 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
<b>NAKIVO Backup &amp; Replication</b>		
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 Transporters for cross-Transporter data transfer. Must be opened on the Transporter machine.
<b>VMware</b>		
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 hosts	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 VM Boot. Must be opened on the Transporter machine and the ESXi host used as the target for a Flash VM Boot.
<b>Hyper-V</b>		
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. Must be opened on Used by Transporters for cross-Transporter data transfer. Must be opened on the Transporter machine. 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.
<b>Physical machine (Windows)</b>		
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.
<b>Physical machine (Linux)</b>		
22	Linux machine	Used by Director to access a Linux physical machine via SSH.
9446 (opens auto-	Linux machine	Used to create the Transporter installed by default.

atically)		
-----------	--	--

## 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
- Mozilla Firefox: Version 74

## 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

#### Transporter Only

- AS1002T
- AS1002T v2
- AS1004T
- AS1004T v2

For minimum hardware requirements, refer to [“Network Attached Storage” on page 105](#).

#### 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-563
- TS-653 Pro
- TS-653A
- TS-653B
- TS-653B
- TS-673
- 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-473A
- TS-673A
- TS-873A
- TS-EC879U-RP
- TVS-675

### Transporter Only

- TS-131P
- TS-231P
- TS-431P
- TS-431X

For minimum hardware requirements, refer to [“Network Attached Storage” on page 105](#).

### 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

#### 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 105](#)

## 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+
- FS3400
- FS3600
- FS6400
- RS819
- RS820+
- RS820RP+
- SA3200D

- SA3400
- SA3600
- DS1520+
- DS1621+
- DS1621xs+
- DS1821+
- DS220
- RS1221+
- RS1221RP+
- RS2421+
- RS2421RP+
- RS4021xs+

### 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 105](#).

### 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 105](#).

# Feature Requirements

Some NAKIVO Backup & Replication features require certain conditions in order to function 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.

- [Hot Add](#)
- [File Recovery](#)
- [File Share Backup](#)
- [Object Recovery and Log Truncation for Microsoft Exchange Server](#)
- [Object Recovery and Log Truncation for Microsoft SQL Server](#)
- [Object Recovery for Microsoft Active Directory](#)
- [Site Recovery](#)
- [Physical Machine Backup And Recovery](#)
- [Cross-Platform Recovery](#)
- [App-Aware Mode](#)
- [Encrypted Backup Repository](#)
- [Direct Connect](#)
- [Native Tape Support](#)
- [Backup Immutability](#)
- [VM Limitation for Multi-Tenancy](#)
- [Monitoring](#)
- [External Database](#)

## 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.

## File Recovery

Recovered files can be downloaded or sent via email. They can also be recovered to a server or file share. 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

If the Transporter assigned to the backup repository is installed on Linux:

- NTFS
- FAT32
- EXT3
- EXT4
- XFS

- For the ReiserFS file system, it is necessary to install the `linux-image-extra-virtual` package:  
`apt-get -y install linux-image-extra-virtual` for Ubuntu.
- 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**
  - Microsoft Windows Server 2022 (x64)
  - Microsoft Windows Server 2019 Standard (x64)
  - Microsoft Windows Server 2016 Standard (x64)
  - Microsoft Windows Server 2012 R2 Standard (x64)
  - Microsoft Windows Server 2012 Standard (x64)
  - Microsoft Windows Server 2008 R2 Standard (x64)
  - Microsoft Windows Server 2008 R2 Standard 32-bit
  - Microsoft Windows 11 (x64)
  - Microsoft Windows 10 Home (x64)
  - Microsoft Windows 10 Professional (x64)
  - Microsoft Windows 8 Professional (x64)
  - Microsoft Windows 8 Professional (x32)
  - Microsoft Windows 7 Professional (x64)
  - Microsoft Windows 7 Professional (x32)
- **Linux**
  - Ubuntu 20.04 Server (x64)
  - Ubuntu 18.04 Server (x64)
  - Ubuntu 16.04 Server (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 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 8 (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)

### **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.

### **Services and packages**

The following packages and services should be installed/running:

## Microsoft Windows

- Net Security package should be installed
- Microsoft iSCSI Initiator service should be installed and running
- `net . exe` utility should be installed
- SMB 2 / CIFS File Sharing Support feature should be turned on
- PowerShell should have 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)

## 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)

## 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

## 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 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, or Linux 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.
- File shares with the following protocols are supported:
  - NFS 3.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

## 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 2019
- Microsoft SQL Server 2017
- Microsoft SQL Server 2016
- Microsoft SQL Server 2014
- Microsoft SQL Server 2012
- Microsoft SQL Server 2008 R2
- Microsoft SQL Server 2008

### 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 2 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 9445 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 2 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 2019
- Windows Server 2016
- Windows Server 2012 R2 Windows Server 2012
- Windows Server 2008 R2

### 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.

## 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 2 / 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.

## Physical Machine Backup And Recovery

The following requirement must be met:

- `$ExecutionContext.SessionState.LanguageMode` in PowerShell must be set to `FullLanguage`.

### Supported Operating Systems for Physical Machine Recovery

- Microsoft Windows Server 2019 (x64)
- Microsoft Windows Server 2016 (x64)
- Microsoft Windows 10 Pro (x64)

### Supported Recovery Destinations

- VMware vSphere version 6.7 or later

## Cross-Platform Recovery

The following scenarios are supported if a VM is [exported from backup](#) and imported into a different hypervisor:

Source Platforms	Target Platforms	
	VMware vSphere 6.7	Microsoft Hyper-V 2016/2019
VMware vSphere 6.7	<ul style="list-style-type: none"> <li>• Windows Server 2016/2019</li> <li>• Ubuntu Server 18.04</li> <li>• RHEL 7</li> </ul>	<ul style="list-style-type: none"> <li>• Windows Server 2016/2019</li> <li>• Ubuntu Server 18.04</li> <li>• RHEL 7*</li> </ul>
Microsoft Hyper-V 2016/2019	<ul style="list-style-type: none"> <li>• Windows Server 2016/2019</li> <li>• Ubuntu Server 18.04</li> <li>• RHEL 7</li> </ul>	<ul style="list-style-type: none"> <li>• Windows Server 2016/2019</li> <li>• Ubuntu Server 18.04</li> <li>• RHEL 7</li> </ul>
Physical Machines	<ul style="list-style-type: none"> <li>• Windows Server 2016/2019</li> <li>• Ubuntu Server 18.04</li> <li>• RHEL 7</li> </ul>	<ul style="list-style-type: none"> <li>• Windows Server 2016/2019</li> <li>• Ubuntu Server 18.04</li> <li>• RHEL 7</li> </ul>

\* To run a VM with RHEL 7 on Microsoft Hyper-V 2016/2019, 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)
```

## 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.

### Microsoft Hyper-V

- Target system must be in the running state and custom OS quiescing must be enabled for it.
- Only the following target systems are supported for custom OS quiescing:
  - Windows Server 2008 R2 Standard (x64)
  - Windows Server 2012 Standard (x64)
  - Windows Server 2012 R2 Standard (x64)
  - Windows Server 2019 Standard (x64)
- Refer to the Knowledge Base [article](#) for more details.

### Nutanix AHV

- Nutanix Guest Tools should be installed and running on all source VMs.
- The necessary permissions should be granted. Refer to the Knowledge Base [article](#) for more details.

### Amazon EC2

To enable application awareness for AWS EC2 instances running on Windows, make sure you meet the following requirements:

#### Supported Operating Systems

- Windows Server 2019 Standard (x64)
- Windows Server 2016 Standard (x64)
- Windows Server 2012 R2 Standard (x64)
- Windows Server 2012 Standard (x64)
- Windows Server 2008 R2 Standard (x64)

#### Requirements for Windows-based Amazon EC2 Instances

Default administrative shares in Windows-based AWS EC2 Instances should be enabled in order for application awareness to work.

#### Requirements for Security Groups of Source EC2 Instances

Ports 445, and 9445, and 9446 should be added to the security group of source AWS EC2 Instances to be backed up or replicated.

#### Requirements for the Director Machine

If the Director is installed on an Amazon EC2 Instance, the following requirements should be met:

- OS firewall should allow access to the TCP port 9446 on the machine where the Director is installed.
- The security group of the Instance where the Director is installed should allow access to TCP ports 445, 9445, and 9446.

## Physical Machines

- Transporter should be available and configured on the source machine.
- All source objects of the job should be running OS that is supported for OS quiescing.
- Microsoft VSS should be available and configured on Windows-based source objects of the job.
- Custom OS quiescing should be enabled on Linux-based source objects of the job.
- The physical source machine should contain [supported volumes](#).

## 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](#).

## Direct Connect

The following platform is supported:

- VMware vSphere

### Note

The free version of VMware vSphere is not supported.

## Requirements

Direct connect supports the following Transporter operating systems:

- Windows
- Linux

Direct connect supports the following Transporter types:

- Onboard transporter
- Installed service
- VMware vSphere appliance

The following deployment scenarios are supported:

- Director and Transporter(s) installed at the MSP's site and a single Proxy Transporter installed at each client site.

- Primary repository at client site (managed by client's Proxy Transporter) and a secondary repository at the MSP's site.

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 Transporter.

## Supported Features

Direct connect supports the following product features:

- Inventory discovery/refresh
- Repository creation/management
- VM backup
- VM replication
- Full VM recovery
- VM verification
- Flash VM boot (from client repository to client infrastructure/from MSP repository to MSP infrastructure)
- Backup copy
- File recovery to browser
- Auto update

## 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).
- Only VTLs connected to Linux are currently supported.
- The gateway VM deployed on-premises needs to have the following minimum resources:
  - **CPU:** x86-64, 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.

## Backup Immutability

To make backups immutable in a **Amazon S3** type of Backup Repository, the following options must be enabled for the Amazon S3 bucket where the repository is located:

- Object Lock
- Versioning

### Note

Disable Object Lock retention mode and retention period on the S3 bucket where the repository is located, as retention settings are set by NAKIVO Backup & Replication during job creation.

For **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.

For the feature to be available on the FreeNAS\TrueNAS, the following 2 settings must be configured:

- *allow\_chflags = yes*
- *seclabel = 0*

### Note

Only Linux OS and NAS OS specified in [system requirements](#) are supported.

## 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

## 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 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](#).

## External Database

The following external databases are supported:

- PostgreSQL v10-14

Make sure to adhere to the following system requirements for the machine housing external database:

- **CPU:** x86-64, 4+ cores
- **RAM:** 4-8+ GB
- **Free Space:** 50 GB. SSD is highly recommended.
- **OS:** Windows and Linux [operating systems](#).

### Note

The external database can be created on either the physical machine, VM, or stored in a container.

# 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” below](#)
- [“Multi-Tenant Deployment” on the next page](#)
- [“Single Site Deployment” on page 148](#)

## 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 a Remote Site to a Master Site” below](#)
- [“Replication from a Remote Site to a Master Site” on page 140](#)
- [“Local Backup at Remote Site” on page 142](#)
- [“Local Replication at Remote Site” on page 144](#)
- [“Backup at Master Site” on page 145](#)
- [“Replication at Master Site” on page 147](#)
- [“Multi-Tenant Mode” on page 1082](#)

### Backup from a Remote Site to a Master Site

- [Deployment Scenario](#)
- [Deployment Diagram](#)
- [Deployment Steps](#)
- [Connections](#)

## Deployment Scenario

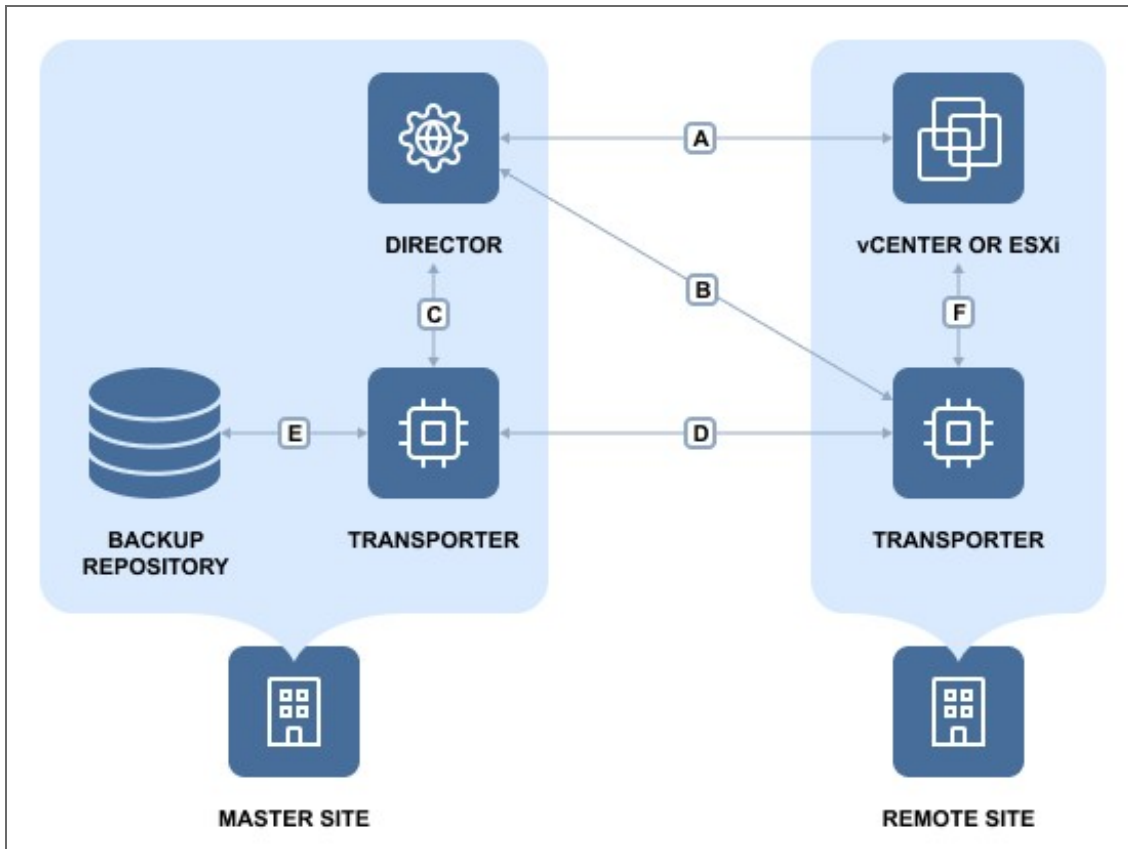
In this scenario, tenant VMs are running at remote sites and are backed up to a single master site.

### 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
<b>A</b>	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>B</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.
<b>C</b>	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.
<b>D</b>	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.
<b>E</b>	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.
<b>F</b>	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.
<b>G</b>	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 a Remote Site to a 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
<b>A</b>	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>B</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.
<b>C</b>	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.

<b>D</b>	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.
<b>E</b>	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.
<b>F</b>	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.
<b>G</b>	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
<b>A</b>	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>B</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.
<b>C</b>	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

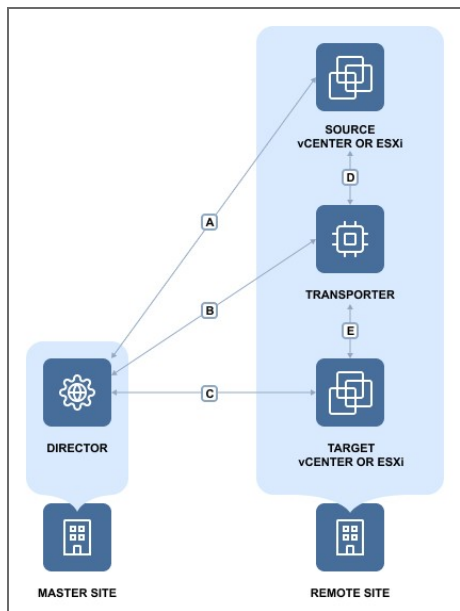
In this scenario, tenant VMs are running and replicated locally at the remote sites.

#### 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
<b>A</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.
<b>B</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.
<b>C</b>	Connection from the machine on which the Transporter is installed to vCenter servers and ESXi hosts running source VMs.
<b>D</b>	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

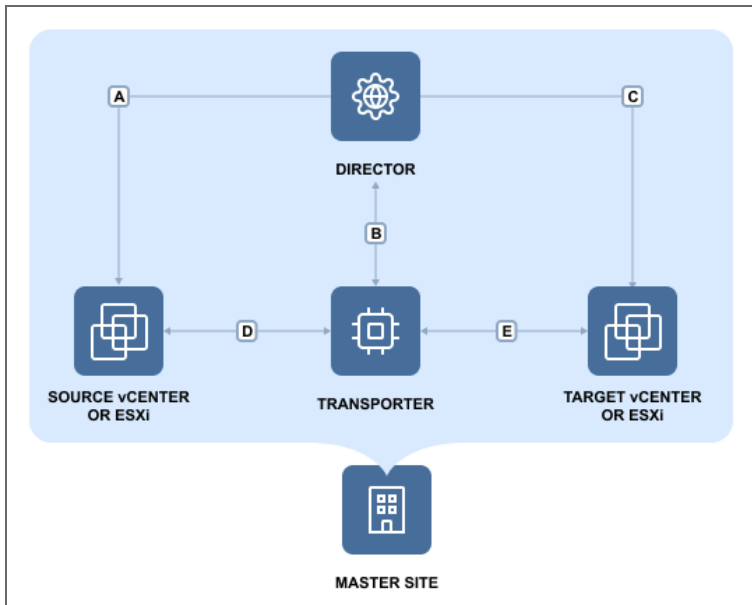
In this scenario, tenant VMs are running at the Master site and the replication of tenant VMs is also performed at the Master site.

#### 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
------------	-------------

<b>A</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.
<b>B</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.
<b>C</b>	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.
<b>D</b>	Connection from the machine on which the Transporter is installed to vCenter servers and ESXi hosts running source VMs.
<b>E</b>	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 the next page](#)
- [“Deploying Nutanix AHV Virtual Appliance” on page 166](#)

- [“Installing on Windows” on page 177](#)
- [“Installing on Linux” on page 186](#)
- [“Installing on Synology NAS” on page 194](#)
- [“Installing on QNAP NAS” on page 199](#)
- [“Installing on Western Digital NAS” on page 207](#)
- [“Installing on ASUSTOR NAS” on page 203](#)
- [“Installing on NETGEAR ReadyNAS” on page 209](#)
- [“Installing on Generic ARM-based Device” on page 212](#)
- [“Deploying Amazon Machine Image in Amazon EC2” on page 174](#)
- [“Installing on FreeNAS/TrueNAS” on page 212](#)
- [“Installing on Raspberry Pi” on page 213](#)

## Deploying VMware Virtual Appliance

- [Deploying Virtual Appliance with vSphere Web Client](#)
- [Deploying Virtual Appliance with vSphere 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

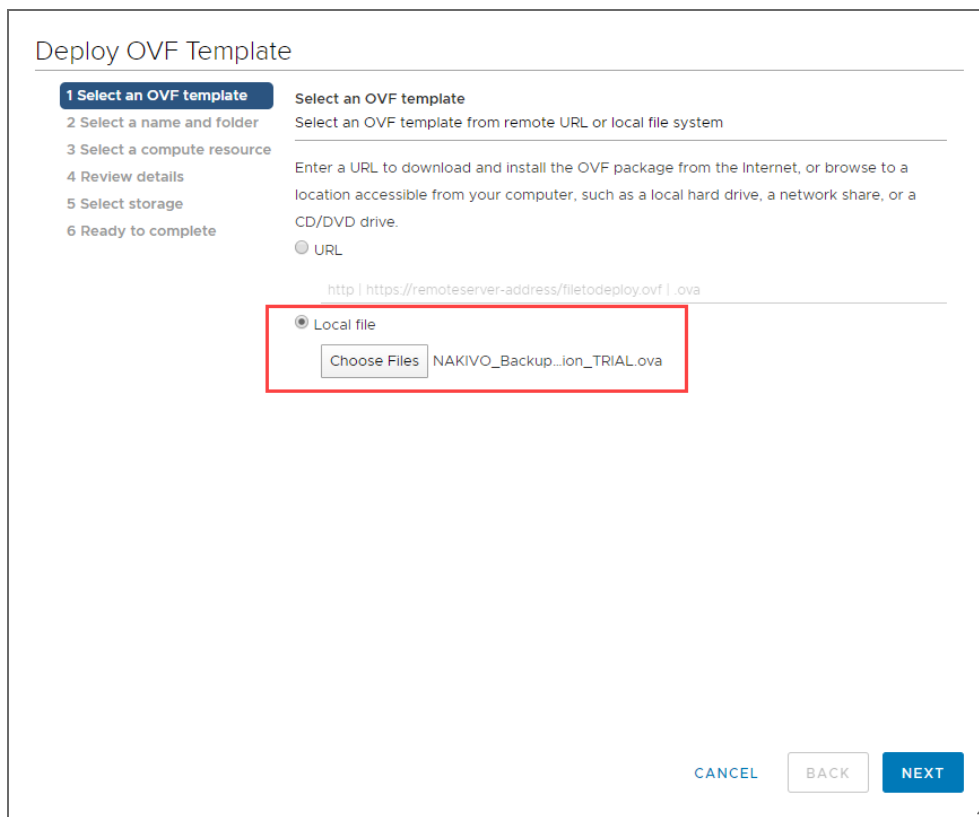
The Virtual Appliance (VA) has two disks: the first (30 GB) contains a Linux OS with NAKIVO Backup & Replication, and the second (500 GB) is used as a Backup Repository. If you deploy the Virtual Appliance disks using the **Thin Provision** option, then the disks will not reserve space on your datastore 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**.

## Deploy OVF Template

- ✓ 1 Select an OVF template
- 2 Select a name and folder**
- 3 Select a compute resource
- 4 Review details
- 5 Select storage
- 6 Ready to complete

Select a name and folder  
Specify a unique name and target location

Virtual machine name: NAKIVO - TW (PRO)

Select a location for the virtual machine.

- 10.30.21.8
  - Support&Product**
  - TEST
  - vSAN cluster

CANCEL BACK 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**.

### Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- 3 Select a compute resource**
- 4 Review details
- 5 Select storage
- 6 Ready to complete

**Select a compute resource**  
Select the destination compute resource for this operation

- ▼ Support&Product
- > Cluster
- > 10.30.21.26

Compatibility

✓ Compatibility checks succeeded.

CANCEL BACK NEXT

7. On the **Review details** page, review the template details and click **Next**.

### Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- 4 Review details**
- 5 License agreements
- 6 Select storage
- 7 Select networks
- 8 Ready to complete

**Review details**  
Verify the template details.

Publisher	No certificate present
Product	NAKIVO Backup and Replication
Version	9.2
Description	Ubuntu 18.04 Server VA with NAKIVO Backup and Replication 9.2 preinstalled VA login: root VA password: QEXS-6b%3D Product URL: https://<SERVER_IP>:4443
Download size	1.0 GB
Size on disk	2.5 GB (thin provisioned)
	525.0 GB (thick provisioned)

CANCEL BACK 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**.

Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Review details
- 5 License agreements**
- 6 Select storage
- 7 Select networks
- 8 Ready to complete

**License agreements**  
The end-user license agreement must be accepted.

Read and accept the terms for the license agreement.

END USER LICENSE AGREEMENT (EULA)  
(03/12/2018)

BY OPENING THE PACKAGE, INSTALLING, PRESSING "AGREE", OR "YES", OR "ACCEPT", OR USING THE PRODUCT, THE ENTITY OR INDIVIDUAL ENTERING INTO THIS AGREEMENT AGREES TO BE BOUND BY THE FOLLOWING TERMS. YOU ALSO ACKNOWLEDGE THAT YOU HAVE READ AND ACCEPTED OUR PRODUCT PRIVACY POLICY [www.nakivo.com/support/product-privacy-policy/](http://www.nakivo.com/support/product-privacy-policy/). IF YOU DO NOT AGREE WITH ANY OF THESE TERMS OR OUR PRIVACY POLICY, DO NOT INSTALL OR USE THE PRODUCT, PROMPTLY RETURN THE PRODUCT TO NAKIVO OR YOUR NAKIVO RESELLER. IF YOU REJECT THIS AGREEMENT, YOU WILL NOT ACQUIRE ANY LICENSE TO USE THE PRODUCT.

I accept all license agreements.

CANCEL BACK 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.



### Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Review details
- ✓ 5 License agreements
- 6 Select storage**
- 7 Select networks
- 8 Ready to complete

**Select storage**  
Select the storage for the configuration and disk files

Encrypt this virtual machine (No encryption policies available)

Select virtual disk format: Thin Provision

VM Storage Policy: [Empty]

Name	Capacity	Provisioned	Free	Type
21.26-hdd	7.27 TB	12.17 TB	756.57 GB	VM
CosmoTemplates01	37.93 TB	56.23 GB	37.88 TB	NF
VMTemplates03	7.28 TB	2.43 TB	4.84 TB	NF

Compatibility

✓ Compatibility checks succeeded.

CANCEL BACK NEXT

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**.

### Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Review details
- ✓ 5 License agreements
- ✓ 6 Select storage
- 7 Select networks**
- 8 Ready to complete

**Select networks**  
Select a destination network for each source network.

Source Network	Destination Network
192.168.77.0	10.30.22.0

1 items

**IP Allocation Settings**

IP allocation: Static - Manual

IP protocol: IPv4

CANCEL BACK NEXT

11. On the **Ready to complete** page, review the summary of the setups you have configured and click **Finish** to complete deployment.

Deploy OVF Template

Ready to complete  
Click Finish to start creation.

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Review details
- ✓ 5 License agreements
- ✓ 6 Select storage
- ✓ 7 Select networks
- 8 Ready to complete**

Provisioning type	Deploy from template
Name	NAKIVO - TW (PRO)
Template name	NAKIVO_Backup_Replication_VA_v9.2.1_Full_Solution_TRIAL
Download size	1.0 GB
Size on disk	2.5 GB
Folder	Support&Product
Resource	10.30.21.26
Storage mapping	1
All disks	Datastore: CosmoTemplates01; Format: Thin provision
Network mapping	1
192.168.77.0	10.30.22.0
IP allocation settings	
IP protocol	IPV4
IP allocation	Static - Manual

CANCEL BACK FINISH

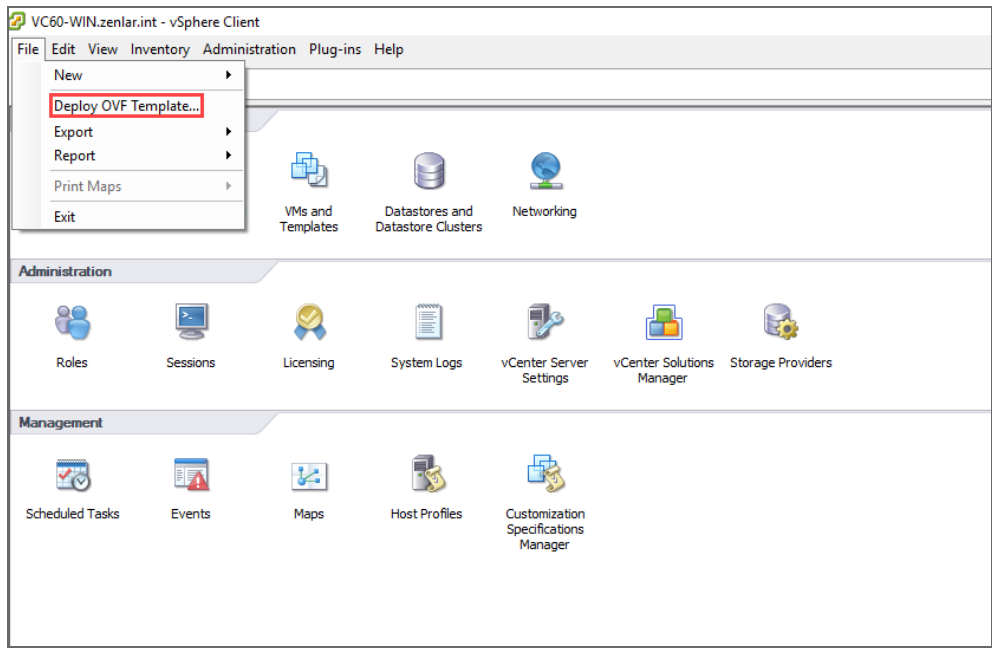
After the Virtual Appliance is deployed, you may need to [configure](#) it.

### 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

## Deploying Virtual Appliance with vSphere Client


1. [Download NAKIVO Backup & Replication VA.](#)
2. Log in to your vSphere vCenter with the vSphere Client, go to **File** in the top menu and select **Deploy OVF Template**.



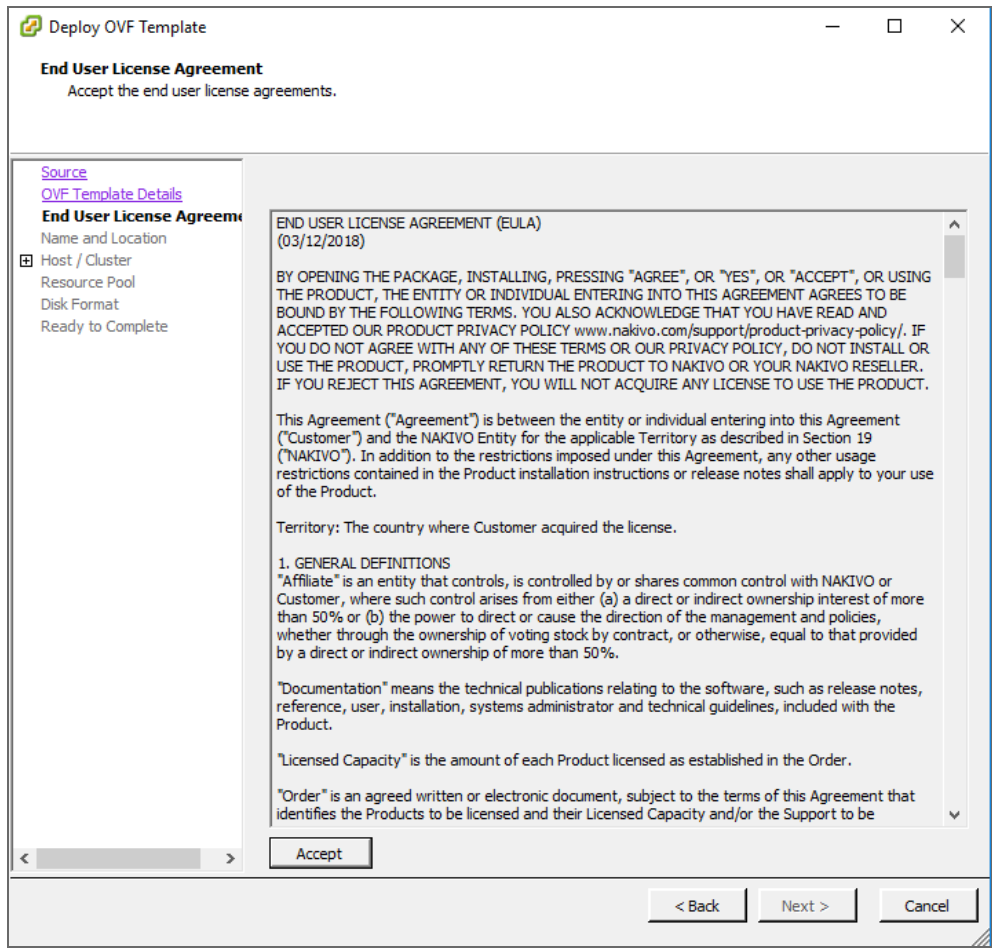
3. On the **Source** page of the **Deploy OVF Template** wizard, select and locate the file with the template. Click **Next**.




4. On the **OVF Template Details** page, review the template details and click **Next**.



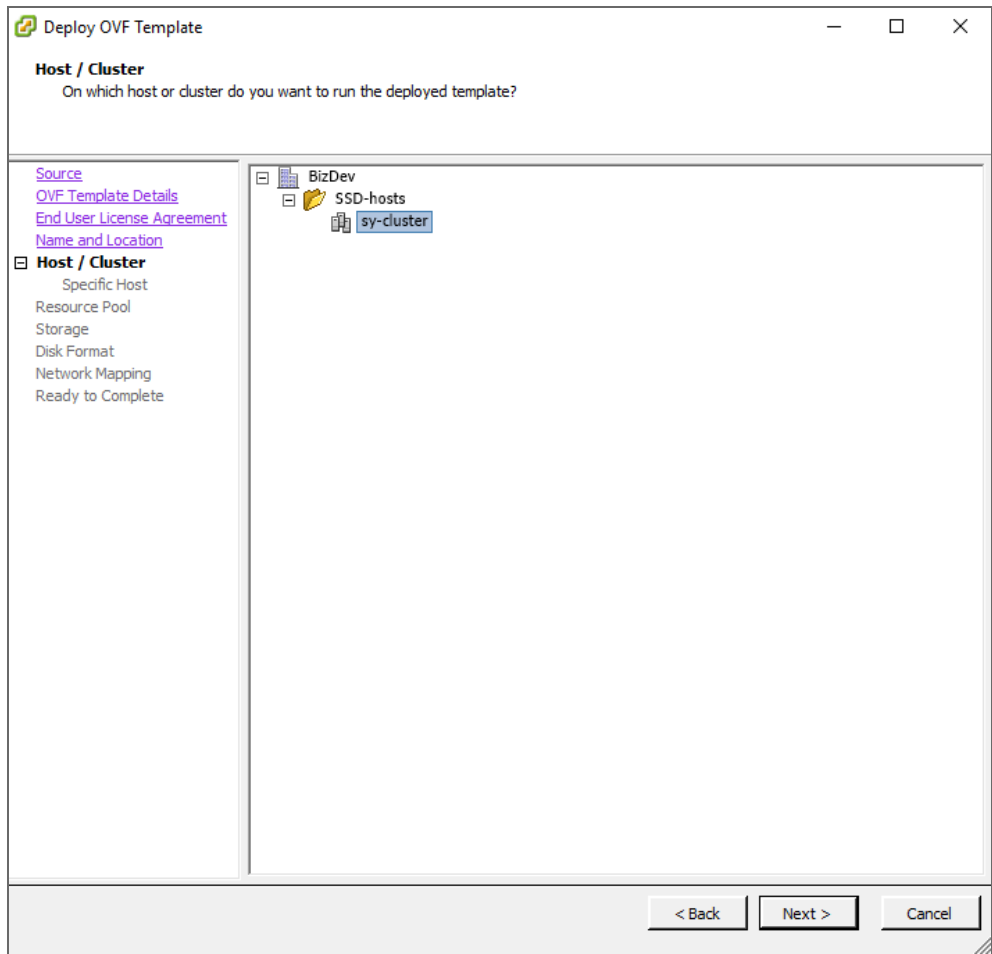
5. On the **End User License Agreement** page, read the license agreement. If you agree to its terms, click **Accept** and then click **Next**.



6. On the **Name and Location** page, specify a name and location for the deployed VA and click **Next**.




7. On the **Host/Cluster** page, select the host or cluster on which you wish to run the deployed template and click **Next**.



8. On the **Storage** page, select a datastore where you would like to keep the VA disk and click **Next**.






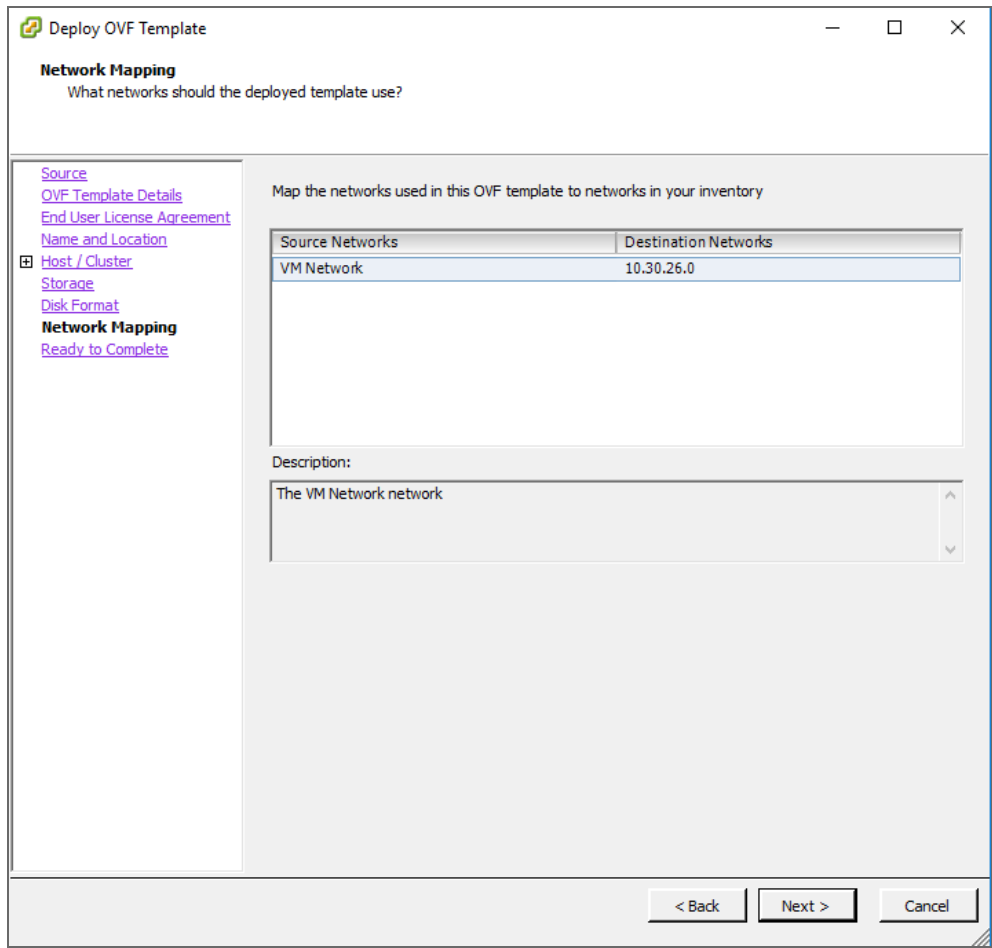
9. On the **Disk Format** page, select a virtual disk format (**Thin Provision** is recommended) and click **Next**.

**Important**

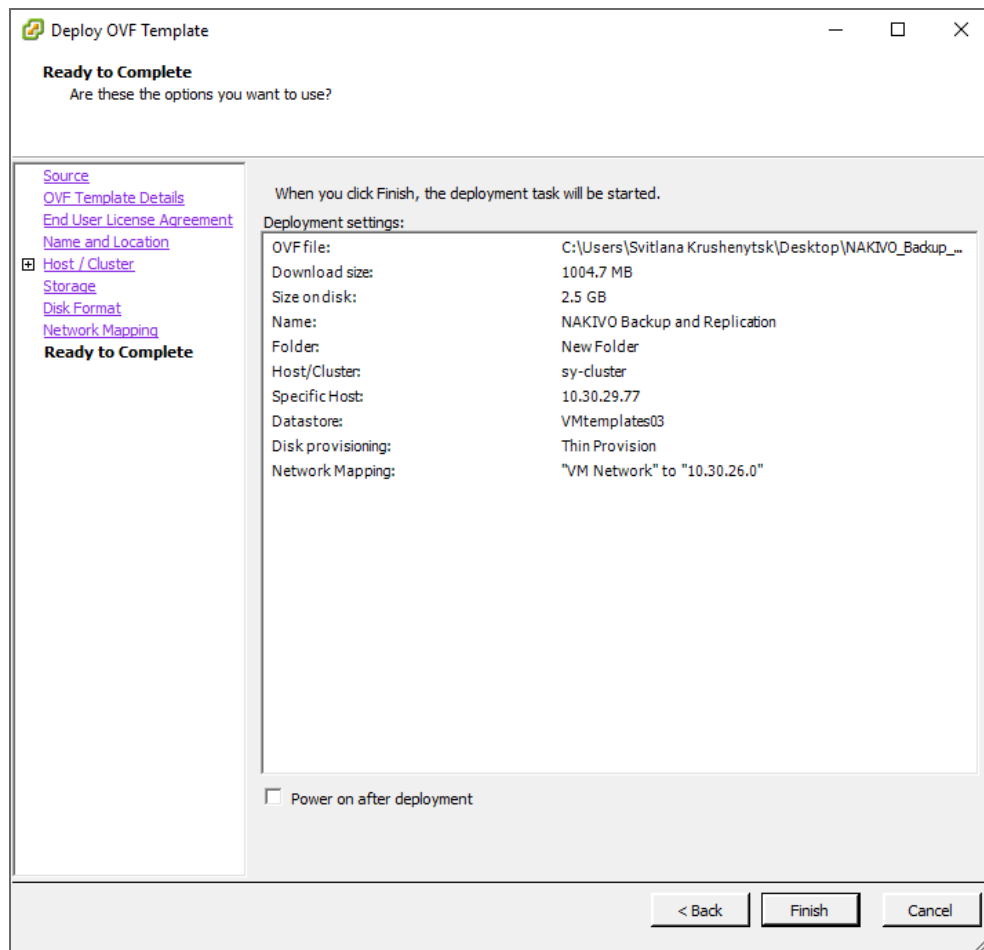
If you wish to select one of the **Thick Provision** options instead of **Thin Provision**, keep in mind that NAKIVO Backup & Replication can take 0,5 TB of data. Check to see if it is 0,5 TB by default for all cases.



10. On the **Network Mapping** page, select a network to which the VA will be connected. It is recommended that you choose a network with DHCP and Internet access. Click **Next**.



11. On the **Ready to Complete** page, review the summary of the options you have configured and select the **Power on after deployment** option.



12. Click **Finish** to complete the deployment.
13. After the Virtual Appliance is deployed, [configure](#) it if necessary.

## Virtual Appliance OS, Credentials, and Security

The appliance runs Ubuntu 20.04, 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 an **Amazon S3** or **Local Folder** type 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 the following URL to access the product's web interface of the VA: `https://Appliance_VM_IP:4443`.

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

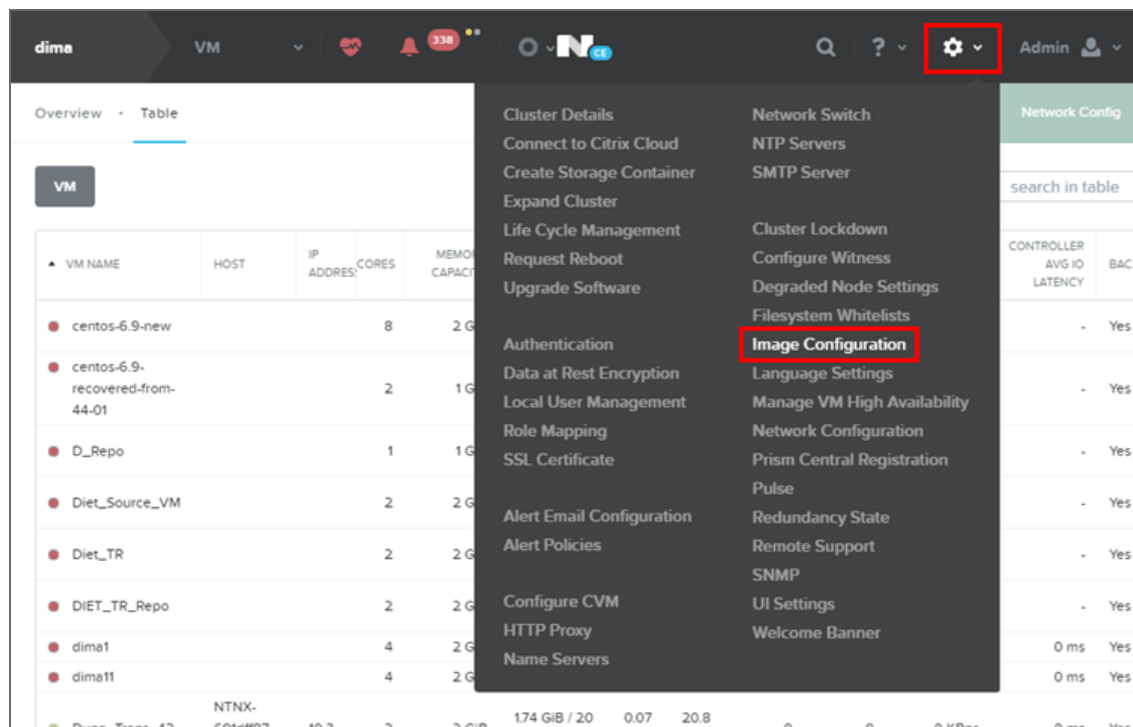
The NAKIVO Backup & Replication instance must be deployed in a Nutanix AHV cluster in order to enable backup and recovery functions.

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**.



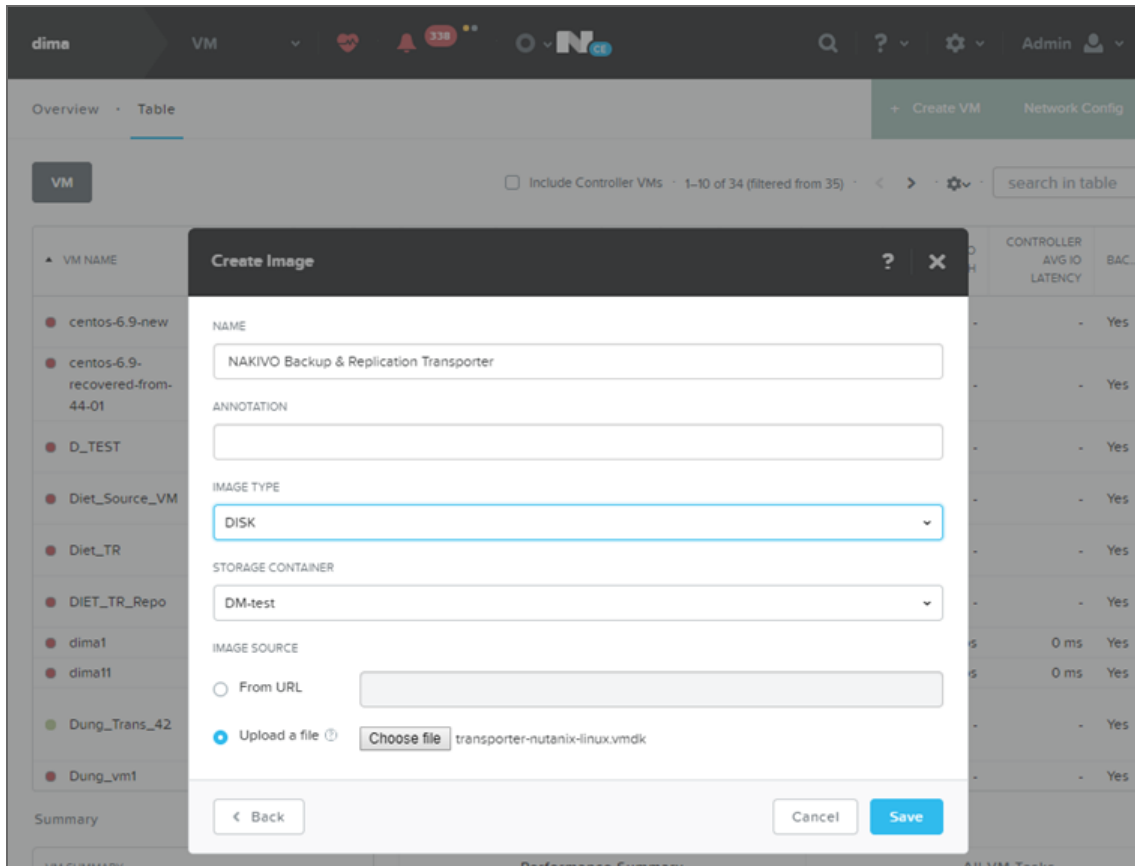
4. In the **Image Configuration** dialog, click **Upload Image**.



5. In the **Create Image** dialog, fill in the following options:

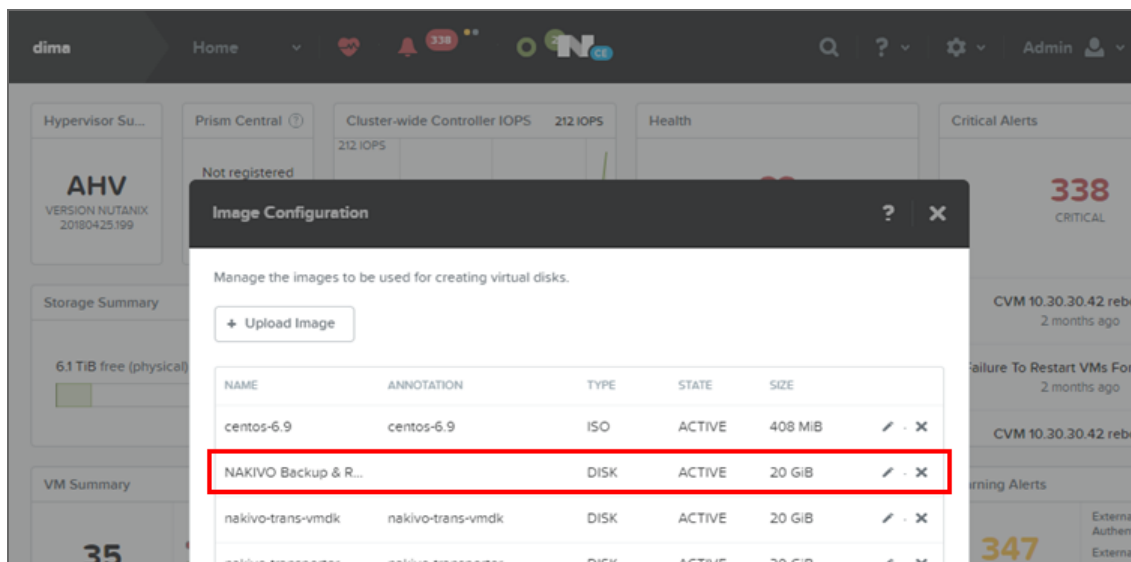
- **Name:** Enter a name for the new image.
- **Image Type:** From the drop-down list, select **DISK**.
- **Storage Container:** Select the storage container you wish to use from the drop-down list. The list includes all storage containers created for this cluster. If there are no storage containers 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.

**Add Disk** ? X

TYPE  
DISK

OPERATION  
Clone from Image Service

BUS TYPE  
SCSI


IMAGE ?  
NAKIVO Backup & Replication Transporter

SIZE (GIB)  
Please note that changing the size of an image is not allowed.  
20


Cancel Add

- 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 20.04, 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 an **Amazon S3** or **Local Folder** type 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 the following URL to access the product's web interface of the VA: `https://Appliance_VM_IP:4443`.

Refer to [“Getting Started” on page 239](#) 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.
3. **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 20.04 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 239](#) 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 39](#). For this feature to be available, the backups must be stored in an **Amazon S3** or **Local folder** type 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 94](#) 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 **snappd**
- 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](#)

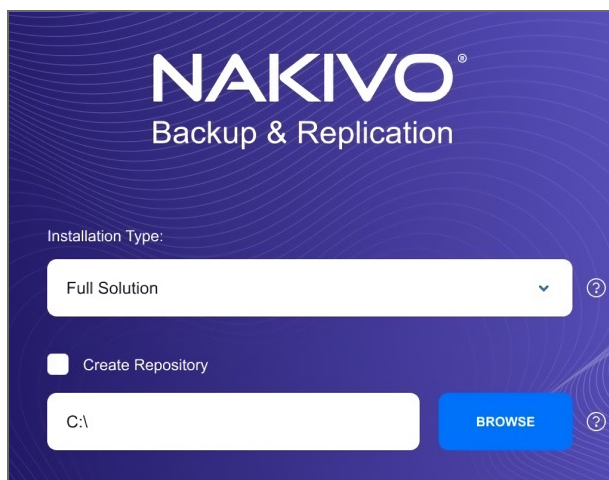
## Installing Full Solution on Windows

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.

You can also change the installation options as follows:

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
```

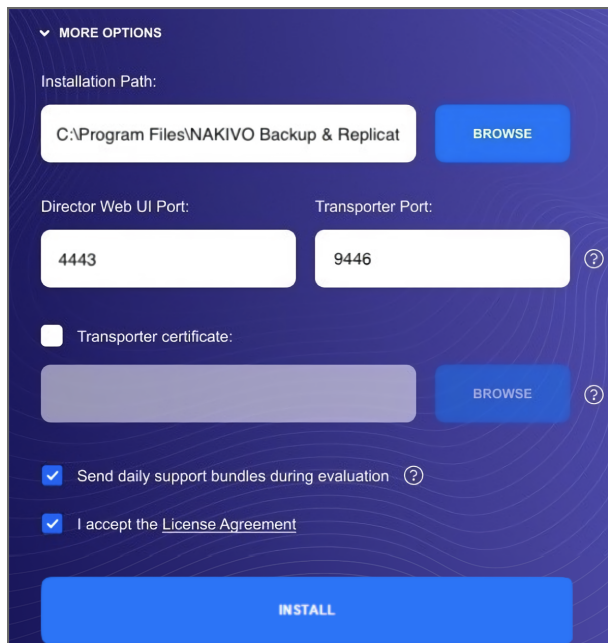
- Use the following command for Linux OS:

```
installer.sh --cert /tmp/certificate.pem --eula-accept
```

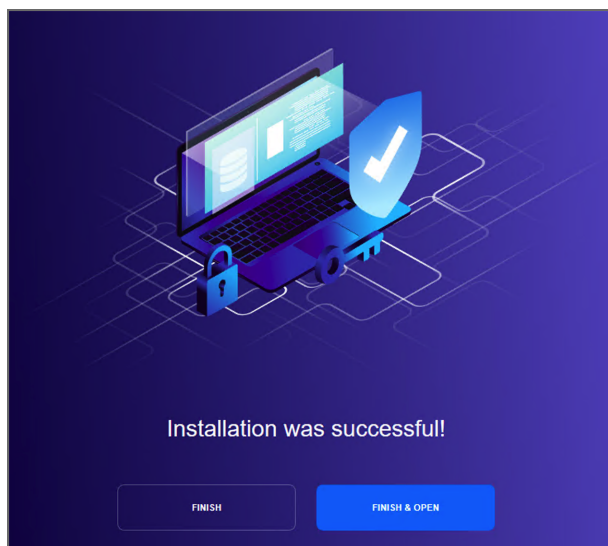
- **Send daily support bundles during evaluation:** When this option is selected, NAKIVO Backup & Replication automatically creates, encrypts, and uploads support bundles once a day to a NAKIVO support server during the evaluation period. The NAKIVO support team may use this information to improve the 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. For more details, refer to [“Logging in to NAKIVO Backup & Replication” on page 239](#) .

## 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:

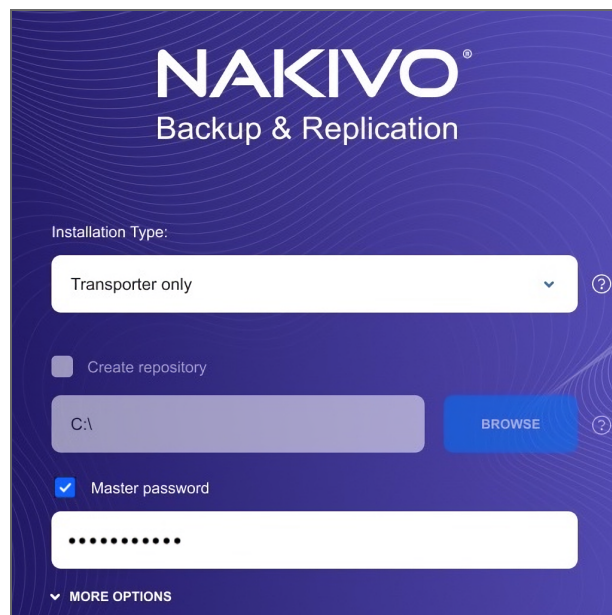
## Transporter Installation Prerequisites

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 Repository (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**.

The screenshot shows the 'MORE OPTIONS' section of the NAKIVO Backup & Replication installation wizard. It features a dark blue background with white and light blue text and buttons. The 'Installation path' is set to 'C:\Program Files\NAKIVO Backup & Replical' with a 'BROWSE' button. The 'Director web UI port' is '4443' and the 'Transporter port' is '9446'. There is an unchecked checkbox for 'Transporter certificate' with a 'BROWSE' button. An unchecked checkbox for 'Send daily support bundles during evaluation' is also present. A checked checkbox for 'I accept the License Agreement' is at the bottom. A large blue 'INSTALL' button is at the very bottom.

7. When the installation is complete the **Transporter installation was successful** notification appears.



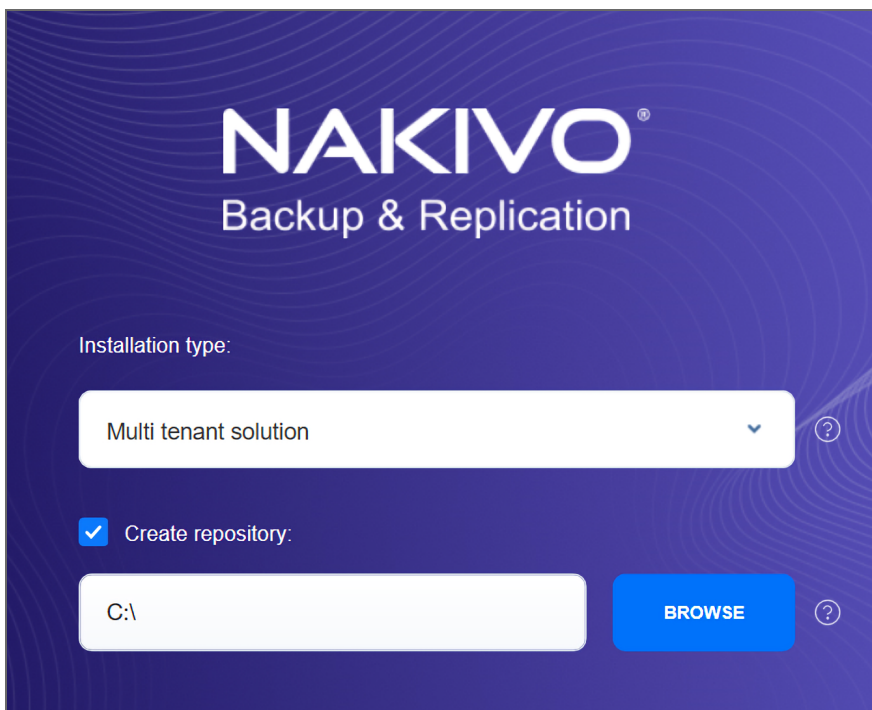
8. Add the **Transporter** to NAKIVO Backup & Replication.

## 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.

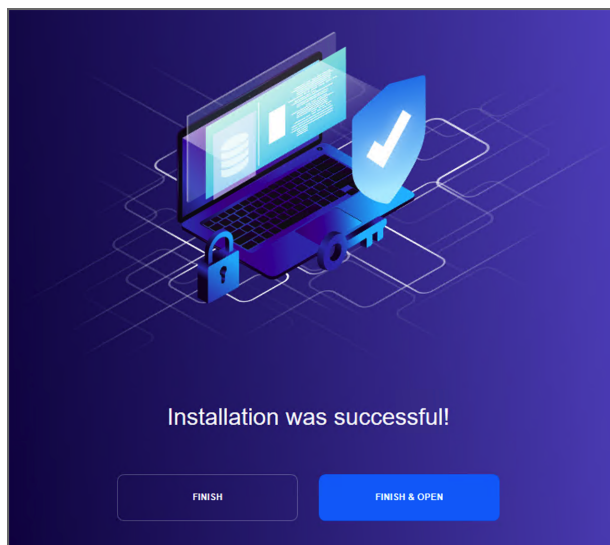
4. Click **Install**.

The screenshot shows the 'MORE OPTIONS' section of the installation wizard. It includes the following fields and options:

- Installation Path:** A text box containing 'C:\Program Files\NAKIVO Backup & Replicat' and a 'BROWSE' button.
- Director Web UI Port:** A text box containing '4443'.
- Transporter Port:** A text box containing '9446' with a help icon.
- Transporter certificate:** A checkbox (unchecked), a text box, and a 'BROWSE' button with a help icon.
- Send daily support bundles during evaluation:** A checked checkbox with a help icon.
- I accept the License Agreement:** A checked checkbox.
- INSTALL:** A large blue button at the bottom.



- 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.

- Create an account by completing the form. For details, refer to [“Logging in to NAKIVO Backup & Replication” on page 239](#).

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.



<b>--eula-accept, -ea</b>	Indicates that you have read and agree to the <a href="#">End User License Agreement</a> .
<b>-f</b>	Performs the silent installation of the full solution (Director and Transporter).
<b>-t</b>	Performs the silent installation of Transporter only.
<b>-m</b>	Performs the silent installation of the full solution in multi-tenant mode.
<b>-u</b>	Performs the silent update of the installed product components.
<b>--release-notes, -n</b>	Indicates the user has read the release notes for the new release during an update.
<b>-sii</b>	Performs the silent install or update ignoring the single installer instance check.
<b>--ignore-pre-install-action-failures, -ipiaf</b>	All pre-install action failures are ignored.
<b>--cert</b>	Allows to set up a custom Transporter certificate.
<b>--master-pass</b> (short version: <b>-b</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 requirements are met:

- On Ubuntu and SLES, NAKIVO Backup & Replication relies on the following packages:
  - cifs-utils
  - open-iscsi

- ntfs-3g
- On RedHat Enterprise Linux, NAKIVO Backup & Replication relies on the following packages:
  - cifs-utils
  - iscsi-initiator-utils
  - ntfs-3g

## 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
<b>-h, -help, help</b>	Display the list of available arguments without starting the installation.
<b>--eula-accept, -ea</b>	Indicates that you have read and agree to the <a href="#">End User License Agreement</a> .
<b>-f</b>	Shall perform the silent installation of the full solution (Director and Transporter).
<b>-t</b>	Shall perform the silent installation of Transporter only.
<b>-m</b>	Shall perform the silent installation of the full solution in <a href="#">multi-tenant mode</a> .
<b>-u</b>	Shall perform the silent update of the installed product components.
<b>-e</b>	Shall install Transporter on Amazon EC2, or update Transporter installed on Amazon EC2. Refer to <a href="#">Updating on Amazon EC2</a> for details.
<b>-a</b>	Shall enable uploading support bundles to support team server (Call Home). Refer to <a href="#">System Settings</a> for details.
<b>-y</b>	Shall accept limitations silently.
<b>-i &lt;install_path&gt;</b>	Shall install to the specified installation path.
<b>-d &lt;director_port&gt;</b>	Shall provide a custom Director port.
<b>-p &lt;transporter_port&gt;</b>	Shall provide a custom Transporter port.

Argument	Description
<b>-r &lt;port1&gt;-&lt;port2&gt;</b>	Shall provide a custom transporter data ports range.
<b>-C</b>	Shall suppress creating the repository.
<b>-c &lt;repo_path&gt;</b>	Shall create the repository. The <b>&lt;repo_path&gt;</b> parameter is optional.
<b>--rt &lt;repo_type&gt;</b>	Shall create a repository of the specified type. The <b>&lt;repo_type&gt;</b> parameter may accept the following values: <b>1</b> – "Forever incremental with deduplication"; <b>2</b> – "Forever incremental without deduplication"; <b>3</b> – "Incremental with full backups (deduplication devices)".
<b>--rc &lt;compress_level&gt;</b>	Shall specify the repository compression level. The parameter may accept the following values: Disabled; Fast; Medium; Best. Refer to <a href="#">Creating Backup Repositories</a> for details.
<b>--pnp-cleanup</b>	Shall clean up the database of the device manager for the Linux kernel.
<b>--cert</b>	Allows to set up a custom Transporter certificate.
<b>--master-pass</b> (short version: <b>-b</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**.
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 following URL in your web browser: `https://machine_IP_or_DNS:director_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.

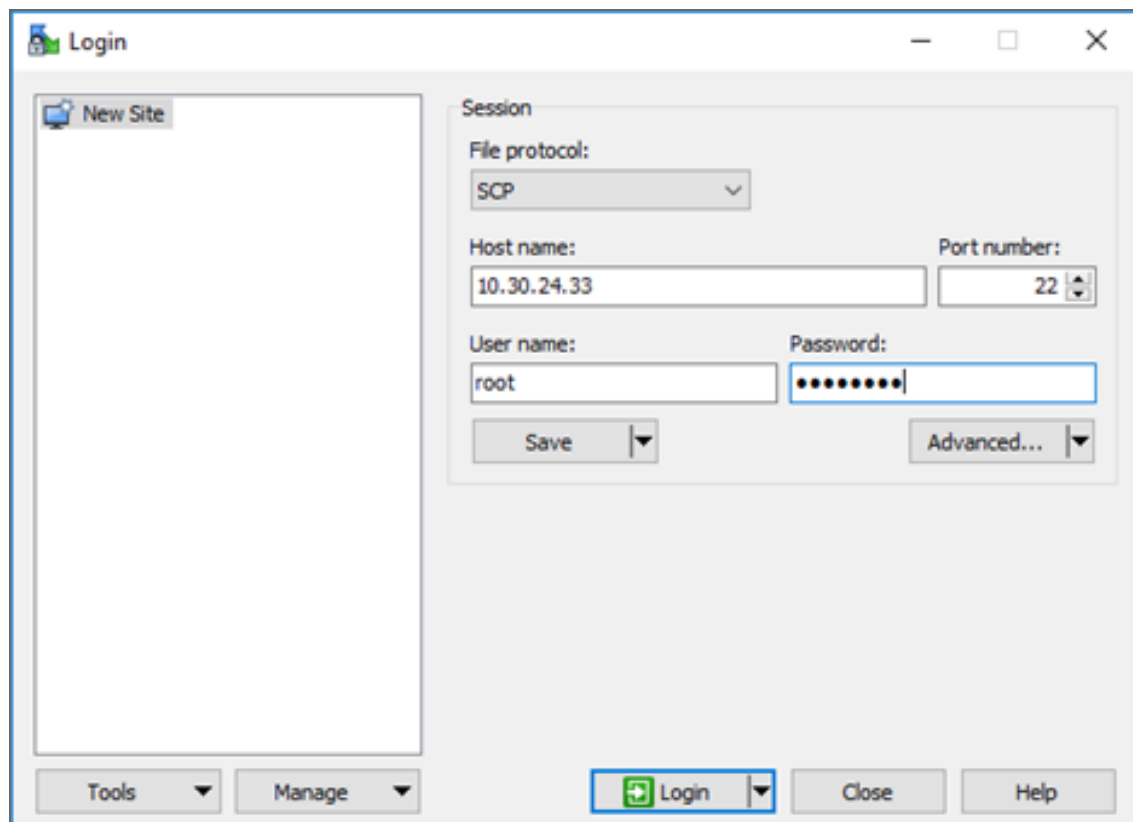
After the installation is complete, you can log in to NAKIVO Backup & Replication by going to the following URL in your web browser: `https://machine_IP_or_DNS:director_https_port`.

Refer to [“Getting Started” on page 239](#) 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**.



6. Click **Yes** in the dialog box that opens.
7. In the left pane, find the folder that contains the Linux installer, in the right pane, go up to the root folder.
8. Drag and drop the installer from left to the right pane.
9. Choose **Binary** from the **Transfer settings** drop-down list in the Copy dialog box that opens.





10. 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 the next page](#)
- [“Installing on Synology NAS Manually” on page 196](#)


### 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 [“Adding Installed Transporters” on page 365](#).

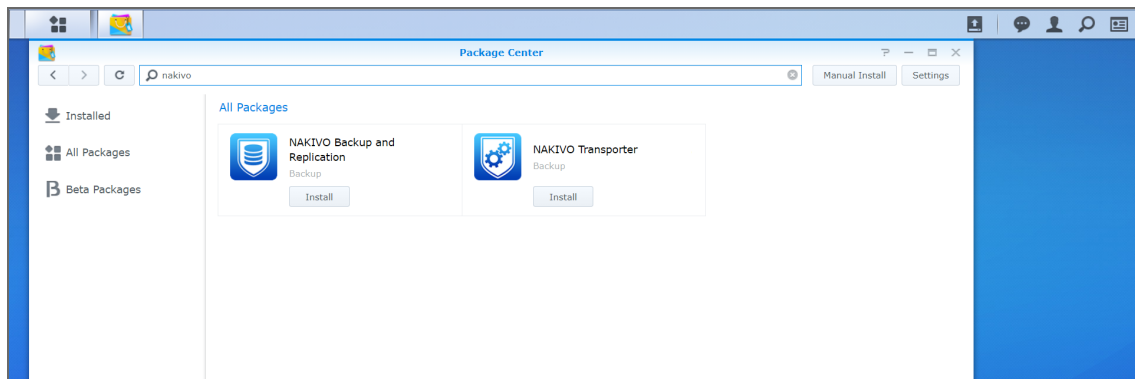
## Installing on Synology NAS via Package Center

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 334](#) topic for details.

Refer to [“Getting Started” on page 239](#) 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](#).
2. Log in to your Synology account and open the **Package Center** in the management interface.



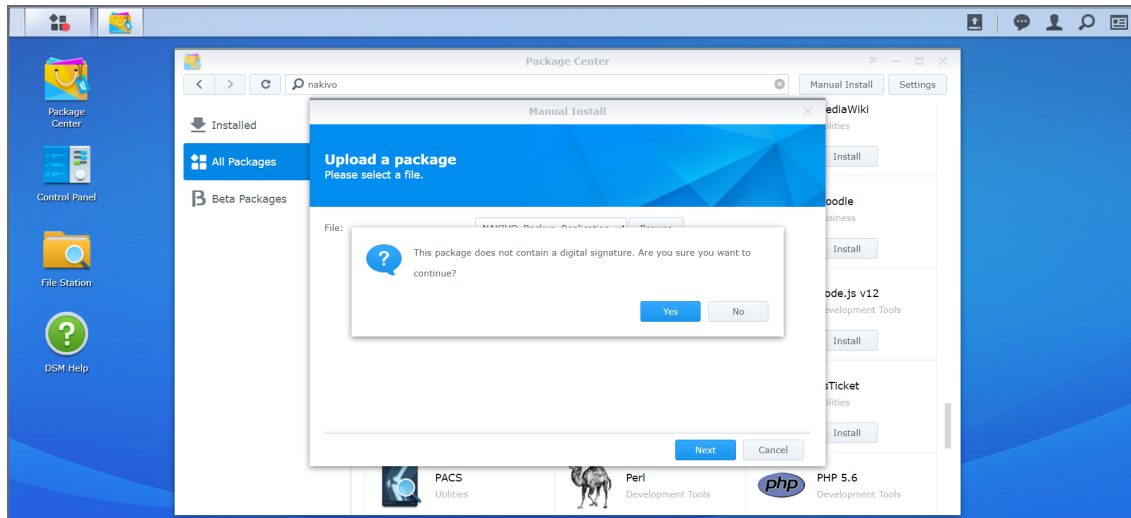
3. Click **Manual Install**.




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**.

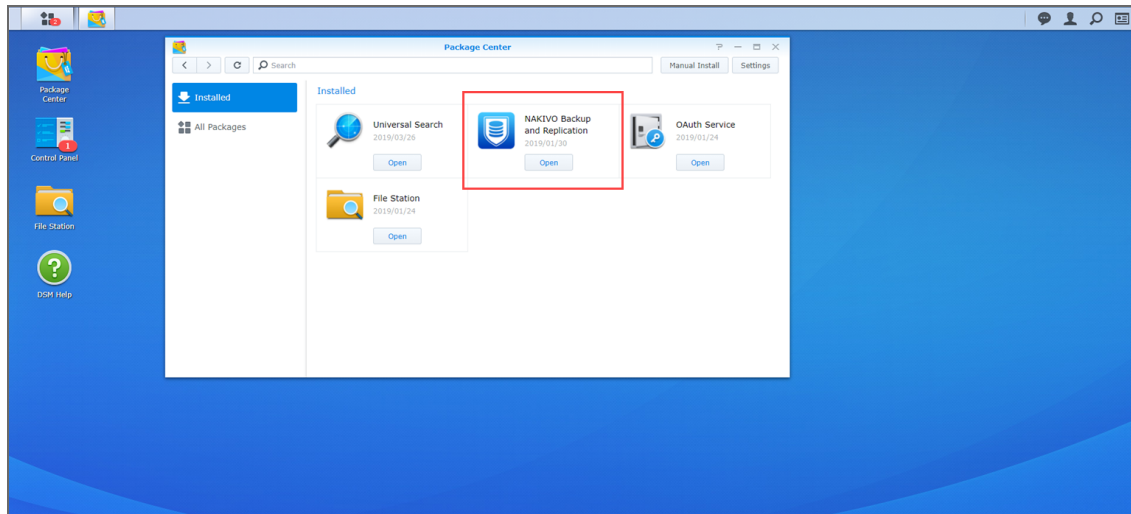


7. Optionally check **Run after installation** to start NAKIVO Backup & Replication immediately after the install process is finished. Click **Apply**.





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 334](#) topic for details.

Refer to [“Getting Started” on page 239](#) 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” below](#)
- [“Installing on QNAP NAS Manually” on the next page](#)

**Note**

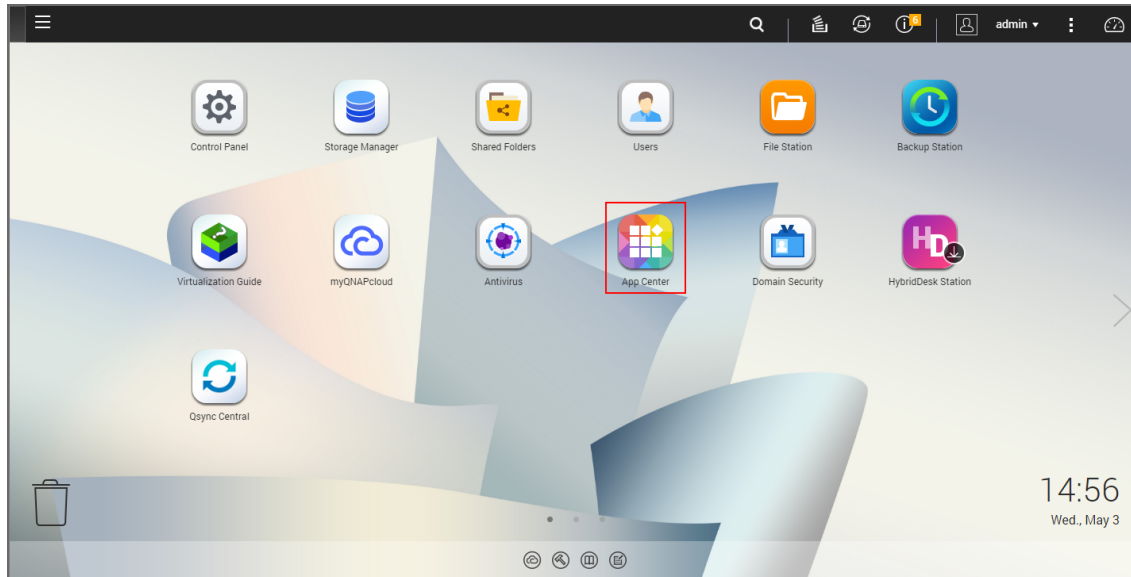
A pre-shared key is not created during Transporter-only installation. When adding this Transporter to NAKIVO Backup & Replication, filling out the master password field is not required. The master password can be manually set and reset later. For details, refer to [“Adding Installed Transporters” on page 365](#).

## 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 334](#) topic for details. Refer to [“Getting Started” on page 239](#) 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.

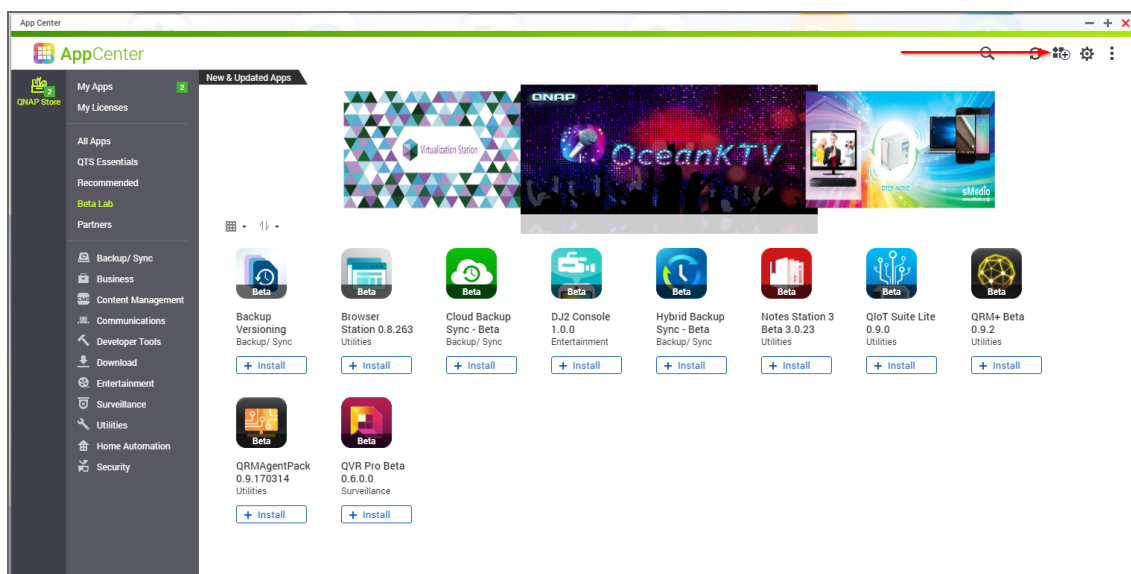
To install NAKIVO Backup & Replication on a NAS:



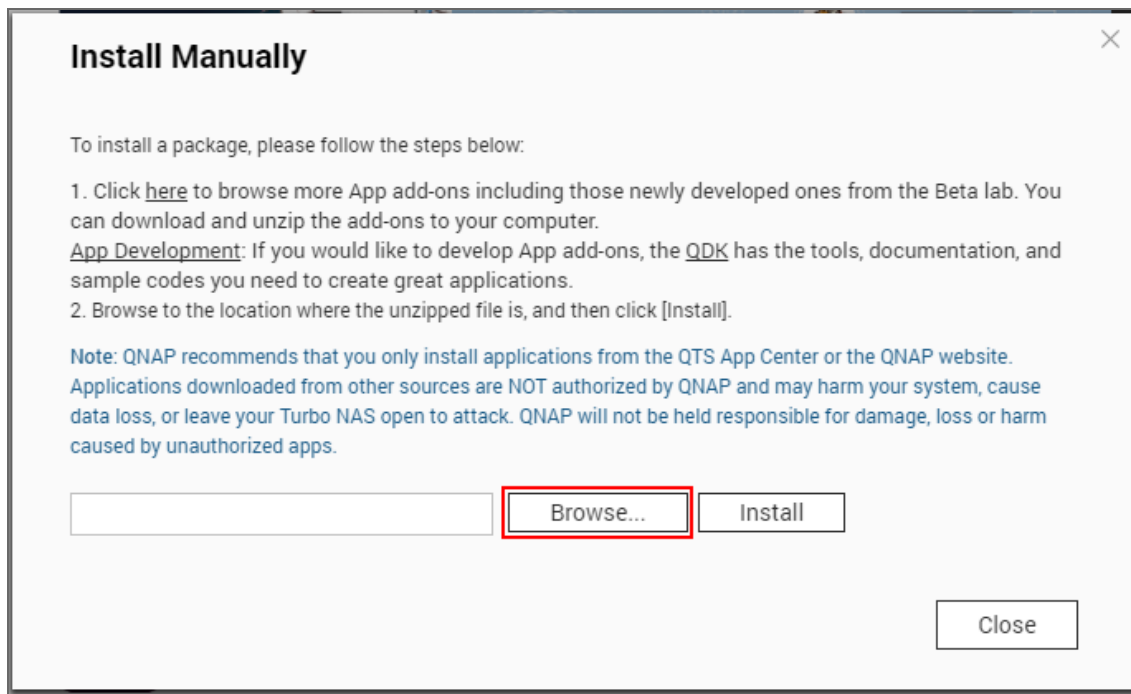
1. Open the QNAP Desktop in your browser by entering the IP address of your QNAP NAS.



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](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 239](#) 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 the next page](#)
- [“Installing on ASUSTOR NAS Manually” on page 205](#)

**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 [“Adding Installed Transporters” on page 365](#).

## Installing on ASUSTOR NAS via App Central

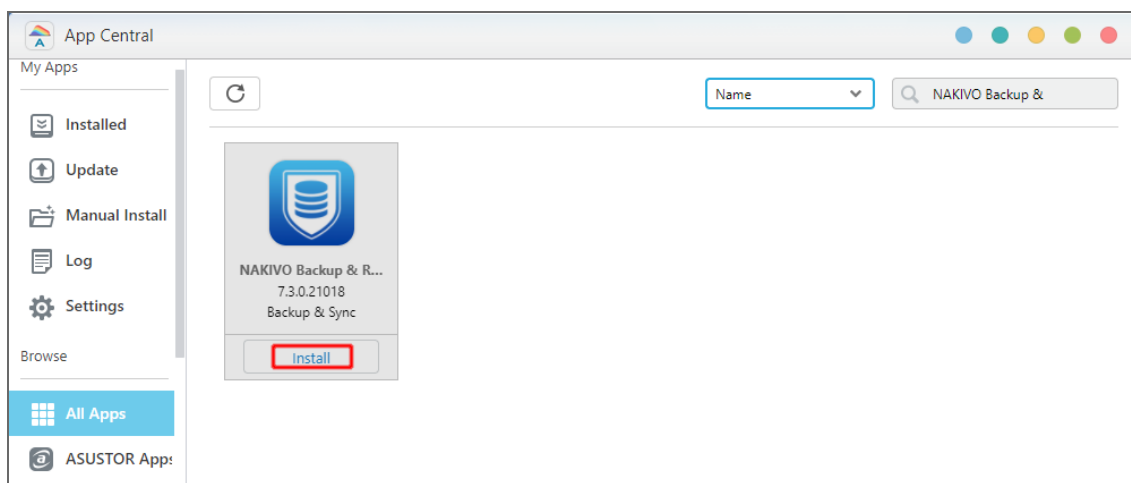
Before you begin installing NAKIVO Backup & Replication on a NAS make sure your NAS model is [supported](#).

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 334](#) topic for details.

Refer to [“Getting Started” on page 239](#) to understand better how to continue working with NAKIVO Backup & Replication.

## Installing on ASUSTOR NAS Manually

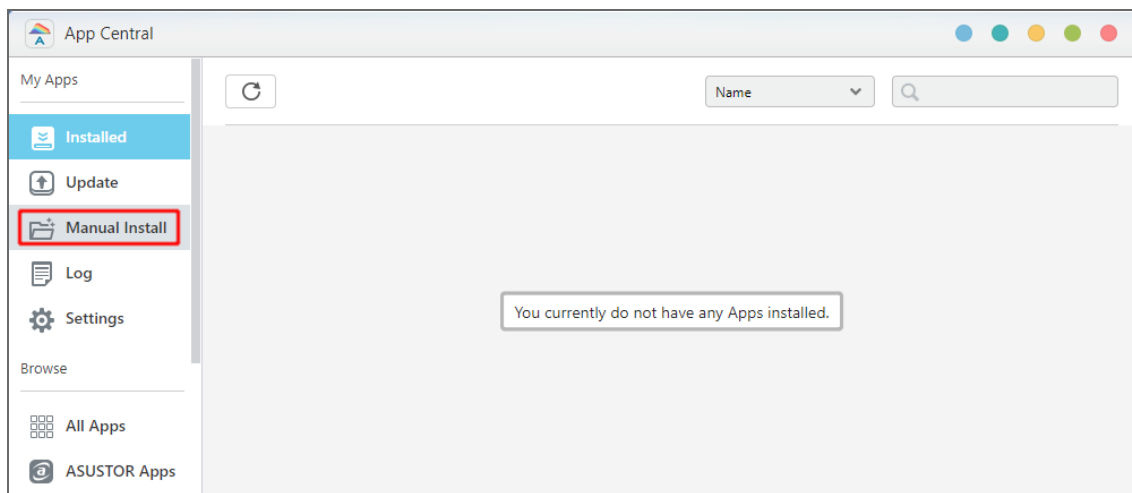
Before you begin installing NAKIVO Backup & Replication on a NAS, make sure your NAS model is [supported](#) and you have downloaded an installer (.apk file) for ASUSTOR NAS.

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 **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 334](#) topic for details.

Refer to [“Getting Started” on page 239](#) to understand better 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 [“Adding Installed Transporters” on page 365](#).

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 239](#) to better understand 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” below](#)
- [“Installing on NETGEAR ReadyNAS Manually” on the next page](#)

### 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 [“Adding Installed Transporters” on page 365](#).

## Installing on NETGEAR ReadyNAS via Available Apps

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 239](#) to know how to continue working with NAKIVO Backup & Replication.

## Installing on NETGEAR ReadyNAS Manually

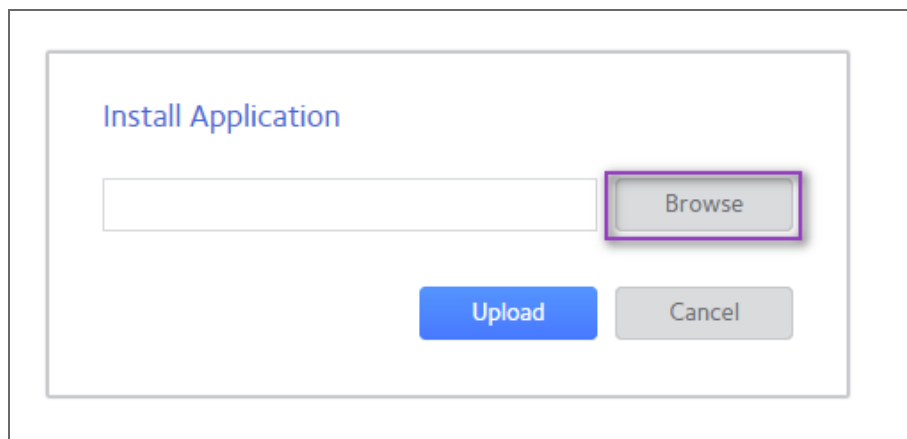
Before you begin installing NAKIVO Backup & Replication or NAKIVO Transporter on a NETGEAR ReadyNAS device, make sure your NAS model is [supported](#) and you have downloaded a relevant installer (.deb file) for NETGEAR ReadyNAS.

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 239](#) to understand better how to continue working with NAKIVO Backup & Replication.

# Installing on Generic ARM-based Device

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](#).
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

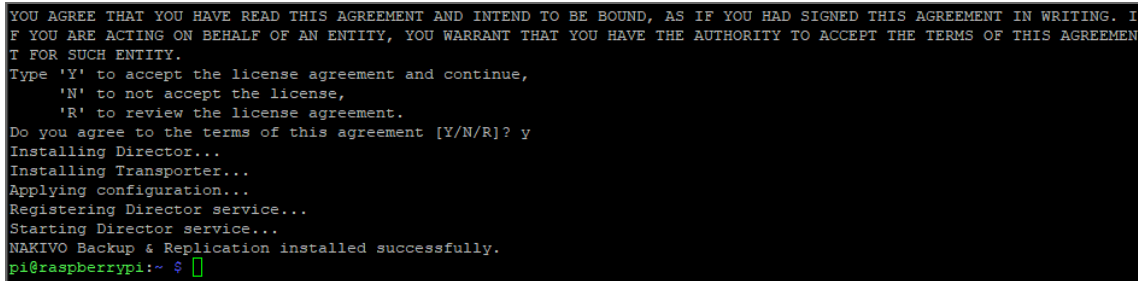
```
chmod +x ./NAKIVO Backup_&_Replication v8.5.0.30224 Installer-NAS-ARM-TRIAL.sh
```

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.



```
YOU AGREE THAT YOU HAVE READ THIS AGREEMENT AND INTEND TO BE BOUND, AS IF YOU HAD SIGNED THIS AGREEMENT IN WRITING. IF YOU ARE ACTING ON BEHALF OF AN ENTITY, YOU WARRANT THAT YOU HAVE THE AUTHORITY TO ACCEPT THE TERMS OF THIS AGREEMENT FOR SUCH ENTITY.
Type 'Y' to accept the license agreement and continue,
    'N' to not accept the license,
    'R' to review the license agreement.
Do you agree to the terms of this agreement [Y/N/R]? y
Installing Director...
Installing Transporter...
Applying configuration...
Registering Director service...
Starting Director service...
NAKIVO Backup & Replication installed successfully.
pi@raspberrypi:~$
```

Refer to [“Getting Started” on page 239](#) to better understand how to continue working with NAKIVO Backup & Replication.

# Installing on FreeNAS/TrueNAS

Make sure the following prerequisites are met:

1. You have access to the FreeNAS/TrueNAS system.
2. Your FreeNAS/TrueNAS system meets [system requirements](#) for installing NAKIVO Backup & Replication.
3. The `ioage` jail/container manager is installed on your FreeNAS/TrueNAS system. Refer to the [ioage README page](#) for a description.

4. A storage pool is created on your FreeNAS/TrueNAS system. Make sure the pool has enough storage for all NAKIVO Backup & Replication functionality. Refer to [FreeNAS User Guide](#) for more details on creating storage pools.

Follow the steps below to install NAKIVO Backup & Replication on a FreeNAS/TrueNAS system:

1. Log in to the FreeNAS/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 v12.2:  
`wget https://github.com/NAKIVO/iocage-plugin-nbr/raw/master/nbr.json`
  - for the full NAKIVO Backup & Replication installation on a FreeNAS v11.3:  
`wget https://github.com/NAKIVO/iocage-plugin-nbr/raw/11.3-RELEASE/nbr.json`
  - for the NAKIVO Backup & Replication Transporter installation on a TrueNAS v12.2:  
`wget https://github.com/NAKIVO/iocage-plugin-nbr-transporter/raw/master/nbr-transporter.json`
  - for the NAKIVO Backup & Replication Transporter installation on a FreeNAS v11.3:  
`wget https://github.com/NAKIVO/iocage-plugin-nbr-transporter/raw/11.3-RELEASE/nbr-transporter.json`

**Note**

If a utility for downloading files like `wget` or `curl` is missing on your FreeNAS/TrueNAS system, you can first download the necessary file to your local machine and then upload it to FreeNAS 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 FreeNAS/TrueNAS system.

- For the full NAKIVO Backup & Replication installation on a FreeNAS/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 FreeNAS/TrueNAS:  
`iocage fetch -P nbr-transporter.json vnet="off" ip4="inherit" ip4_addr="em0|x.x.x.x/24"`

5. For the NAKIVO Backup & Replication Transporter installation, add the Transporter to the Director. Refer to [“Adding Installed Transporters” on page 365](#) for details.

## Installing on Raspberry Pi

NAKIVO Backup & Replication can be installed on a Raspberry Pi computer.

- For system requirements, refer to [“Generic ARM-based NAS devices” on page 106](#).
- For the installation procedure, refer to [“Installing on Generic ARM-based Device” on the previous page](#).

Refer to [“Getting Started” on page 239](#) to better understand how to continue working with NAKIVO Backup & Replication.

# 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.

To manually update any copy of NAKIVO Backup & Replication, go to the download page with updaters.

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:

- [“Software Update” on page 300](#)
- [“Updating Virtual Appliance” below](#)
- [“Updating on Windows” on page 218](#)
- [“Updating on Linux” on page 220](#)
- [“Updating on Synology NAS” on page 220](#)
- [“Updating on Western Digital NAS” on page 222](#)
- [“Updating on Amazon EC2” on page 223](#)
- [“Updating on QNAP NAS” on page 228](#)
- [“Updating on ASUSTOR NAS” on page 231](#)
- [“Updating on NETGEAR ReadyNAS” on page 232](#)
- [“Updating on FreeNAS” on page 234](#)
- [“Updating on Generic ARM-based Device” on page 234](#)


## 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.

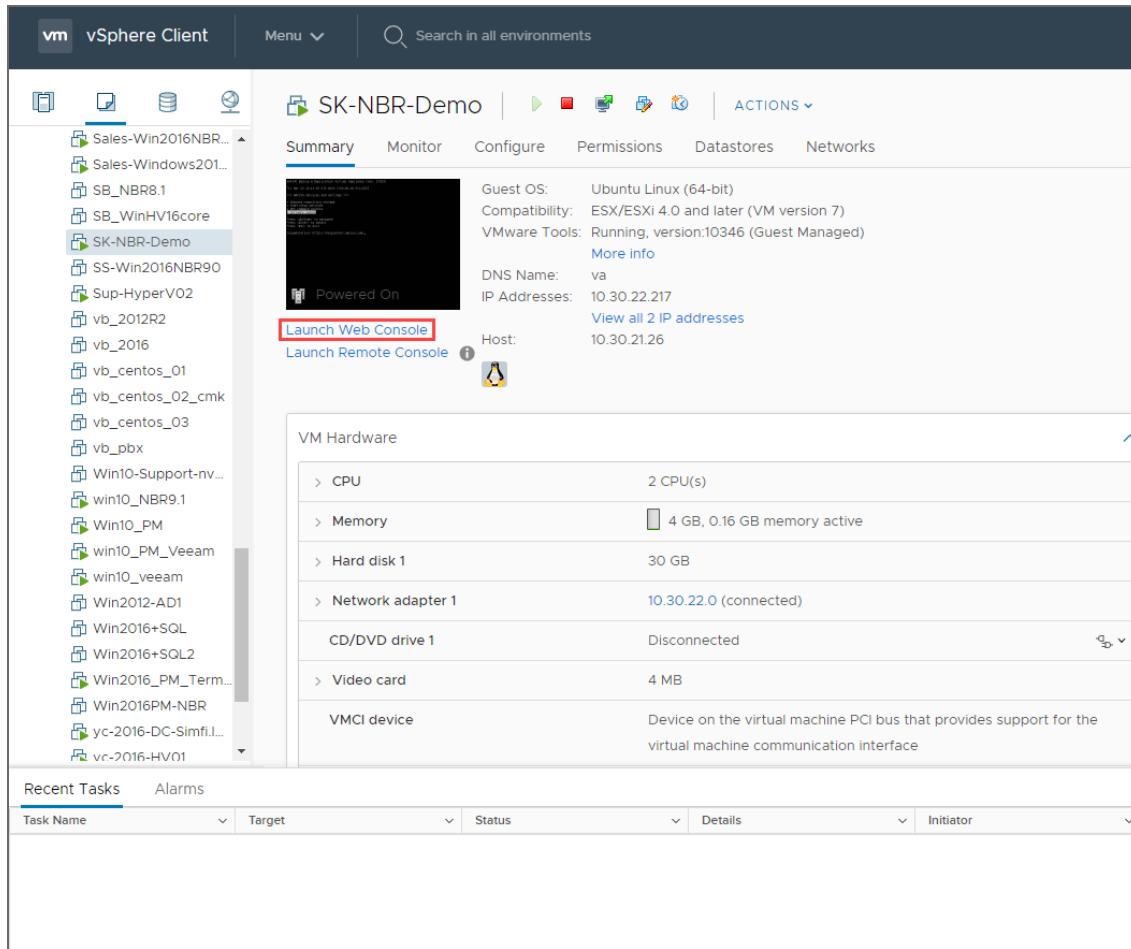


2. Download the latest VA and Linux updater from [www.nakivo.com/resources/download/update/](http://www.nakivo.com/resources/download/update/).
3. Change the directory to `/opt/nakivo/updates` and locate the updater.



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**.

```
NAKIVO Backup & Replication Virtual Appliance (rev. 27323)
Fri Mar 20 12:54:52 UTC 2020 [+00:00:00 Etc/UTC]
Installed components: Director, Transporter

To access the Web Interface, please open https://10.30.21.129:4443 in your web browser.
You can discover this Transporter in the Web Interface under Configuration > Transporters.

* Network settings
* Security settings
* Time and time zone
* System performance
* Manage NAKIVO services
* Exit to system console

Press <Up/Down> to navigate
Press <Enter> to select

Documentation: https://helpcenter.nakivo.com/
```

2. In the menu that opens, select **Software update** and press **Enter**.

```
NAKIVO Backup & Replication Virtual Appliance (rev. 27323)
Fri Mar 20 12:11:02 UTC 2020 [+00:00:00 Etc/UTC]

=== NAKIVO services and settings ===

* Onboard repository storage
* Start/Stop services
* APT command console
* Software update

Press <Up/Down> to navigate
Press <Enter> to select
Press <Esc> to exit

Documentation: https://helpcenter.nakivo.com/
```

3. Select the updater that you have downloaded and press **Enter**.

```
NAKIVO Backup & Replication Virtual Appliance (rev. 46186)
Tue Jan 5 12:01:40 UTC 2021 [+00:00:00 Etc/UTC]

=== Software update ===
Updates directory: /opt/nakivo/updates
Available updates:

* NAKIVO_Backup_&_Replication_v10.2.0.49807_Updater.sh

Press <Up/Down> to navigate
Press <F5> to refresh
Press <Enter> to select
Press <Del> to delete
Press <Esc> to exit

Documentation: https://helpcenter.nakivo.com/
```



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.

```
machine, such as a Unix or Intel based server. A mainframe machine would be an individual mainframe
computer having single or multiple processors or engines.

"Enterprise" is the environment consisting of all hardware owned or leased by Customer in the Territ
ory.

b. LICENSE RESTRICTIONS. The following restrictions apply to certain Products. Each "NAKIVO Backup &
Replication" License is limited for use per CPU - Subcapacity or per Computer - Subcapacity.

c. UNITS OF MEASUREMENT. The following units of measurement apply to certain Products.

per CPU - Full Capacity: A license is required for the total number of active, physical CPUs in each
Computer upon which the Product is performing backup or replication tasks, either remotely or local
ly. "CPU" means a physical processor or central unit in a designated Computer containing the logic c
ircuitry that performs the instructions of a Computer's programs and refers to the "socket" which ca
n contain one or more processor cores.

per CPU - Subcapacity: A license is required for all active, physical CPUs upon which the Product is
performing backup or replication tasks, either remotely or locally. "CPU" means a physical processo
r or central unit in a designated Computer containing the logic circuitry that performs the instruct
ions of a Computer's programs and refers to the "socket" which can contain one or more processor cor
es.

per Computer - Full Capacity: A license is required for all active Computers (either virtual or phys
ical) upon which the Product is upon which the Product is performing backup or replication tasks, ei
ther remotely or locally.

per Computer - Subcapacity: A license is required for all active Computers upon which the Product is
performing backup or replication tasks, either remotely or locally.

YOU AGREE THAT YOU HAVE READ THIS AGREEMENT AND INTEND TO BE BOUND, AS IF YOU HAD SIGNED THIS AGREEM
ENT IN WRITING. IF YOU ARE ACTING ON BEHALF OF AN ENTITY, YOU WARRANT THAT YOU HAVE THE AUTHORITY TO
ACCEPT THE TERMS OF THIS AGREEMENT FOR SUCH ENTITY.
Type 'Y' to accept the license agreement and continue,
      'N' to not accept the license,
      'R' to review the license agreement.
Do you agree to the terms of this agreement [Y/N/R]? Y_
```

- 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 [www.nakivo.com/resources/download/update/](http://www.nakivo.com/resources/download/update/).
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 [www.nakivo.com/resources/download/update/](http://www.nakivo.com/resources/download/update/).
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 89](#).

# Updating on Synology NAS

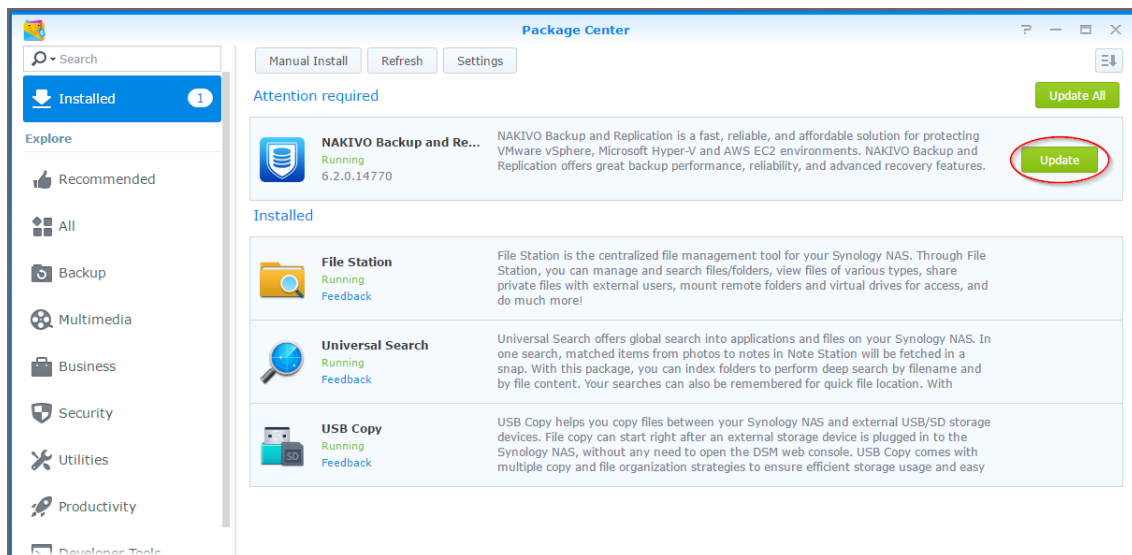
- [Updating via Synology Package Center](#)
- [Updating Manually](#)

## 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.



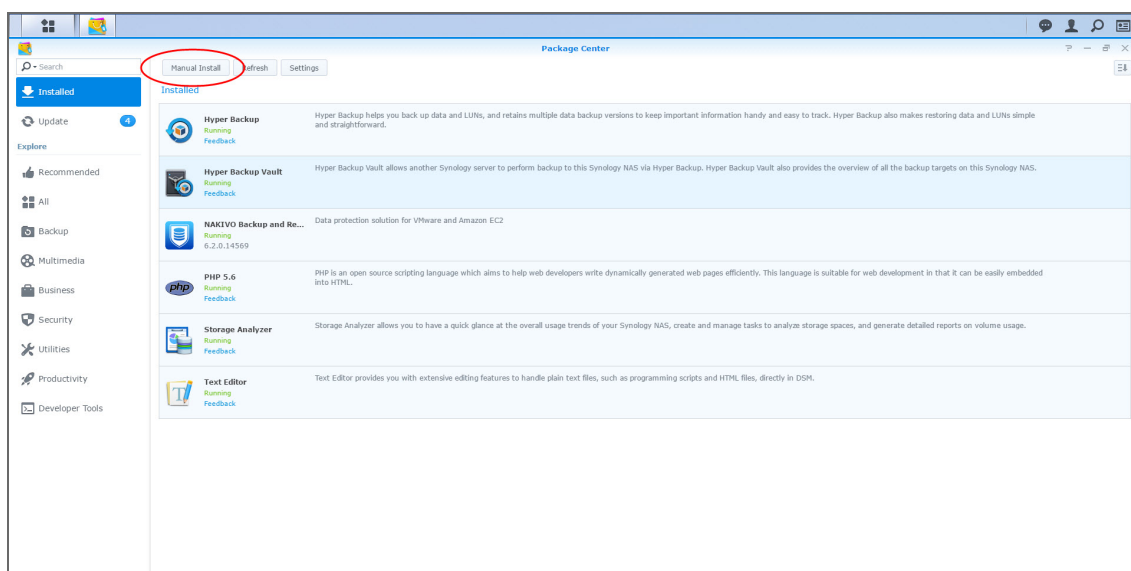
5. Click **Update**.
6. Wait until the update is complete.
7. Repeat these steps on all Synology NAS where you have also installed a Transporter.

## Updating Manually

1. Download the latest Synology NAS updater from [www.nakivo.com/resources/download/update/](http://www.nakivo.com/resources/download/update/).
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.

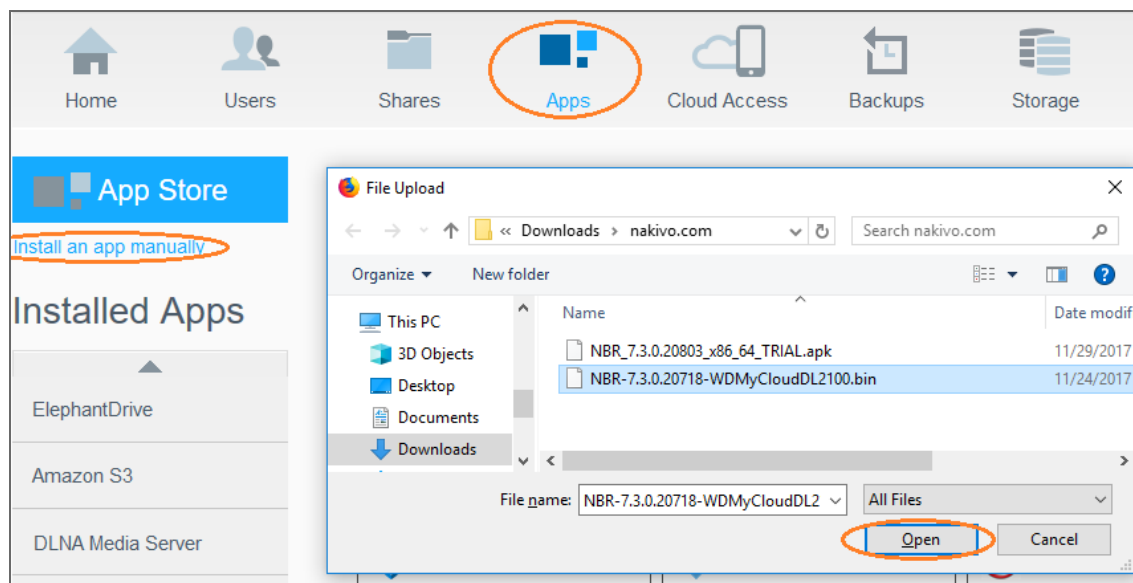
## Updating on Western Digital NAS

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. In 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 the NAKIVO Backup & Replication installer for Western Digital NAS and click **Open**. The update progress bar opens.



4. 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

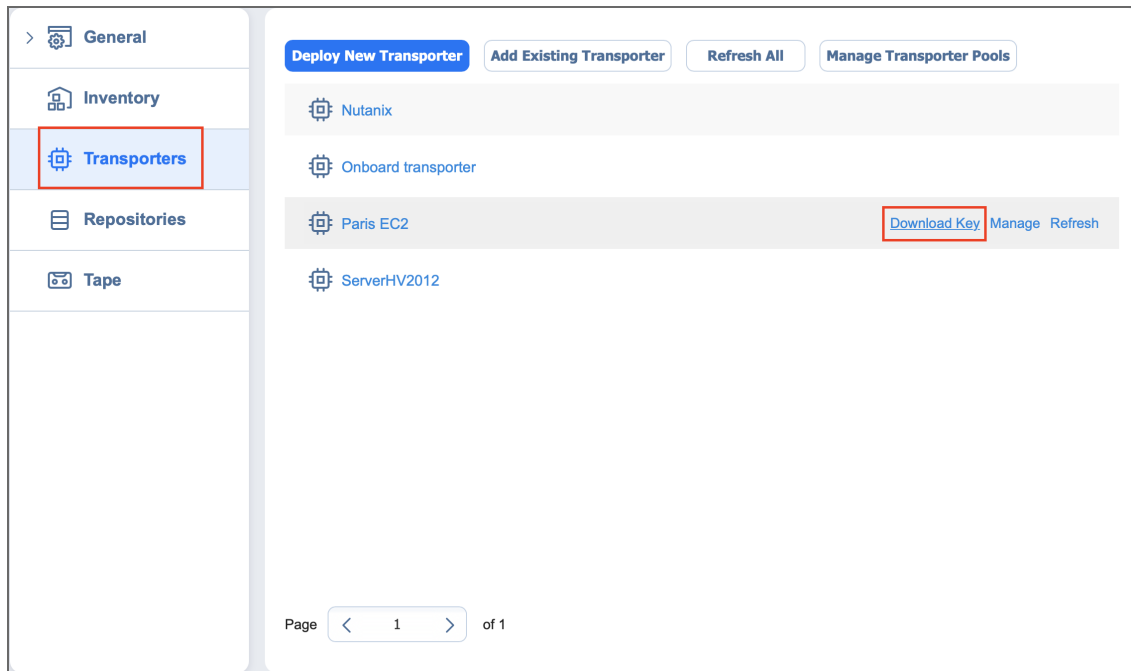
- You have to apply the **-e** argument for executing the installer, in order to avoid changing the Amazon EC2 Transporter with the regular Linux Transporter. Refer to [“Installing on Linux” on page 186](#) for a description of the available arguments.
- Only the main installation of NAKIVO Backup & Replication needs to be updated manually. Transporters installed on Amazon EC2 instances are updated automatically.

## Connecting to an Amazon EC2 Instance 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*.



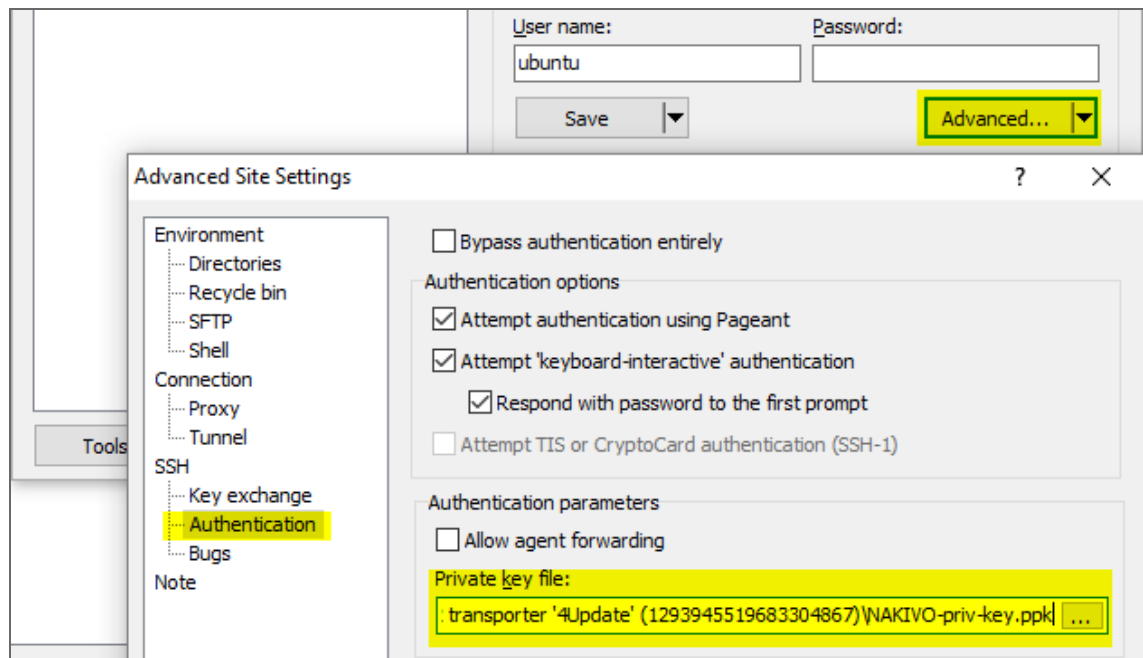


2. Locate the `SSH_key.pem` you just downloaded and unzipped. If you don't see it in the **Open...** dialogue box, change the file type to **All files**.



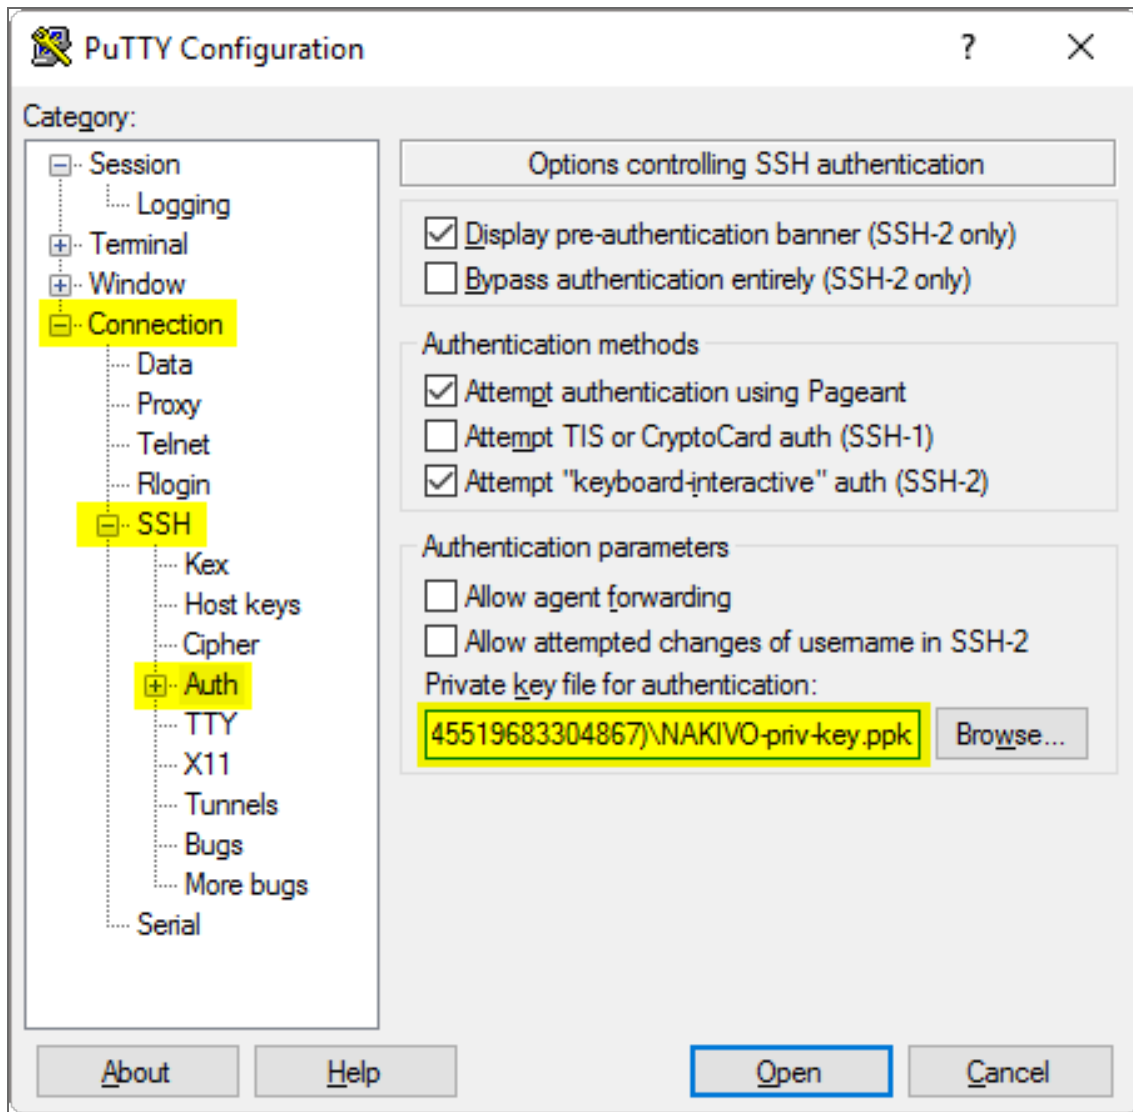
3. Click on **Save private key**. If PuTTYgen asks you to save the key without a passphrase, click **Yes**.
7. Open WinSCP.
8. Create a new session:
  1. Add the hostname or IP address of your Amazon instance you received on step 4 into the **Host Name** box.
  2. In the **Username** box, enter `ubuntu`.
  3. Leave the **Password** box empty.
  4. 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**.
5. Click **Login**.
6. Upload the updater file.
7. Open PuTTY.
8. Enter the IP-address or hostname of the Amazon EC2 instance.

9. Go to *Connection > SSH > Auth* and add the private key in *Private key file for authentication:* box.



10. Click **Open**.
  11. In the command line prompt that opens: log in to the Amazon EC2 instance:
    1. For **login**, enter `ubuntu`
    2. For **password**, leave a blank line.
9. Update NAKIVO Backup & Replication following the [instructions](#).

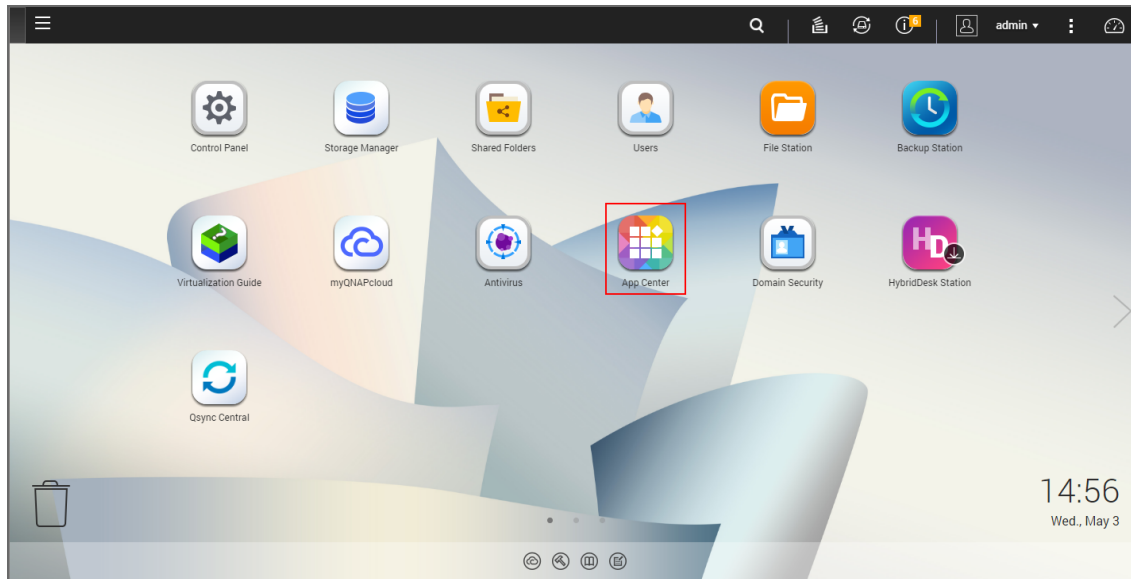
## Updating on QNAP NAS

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. Go to **App Center**.
3. 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".
4. If the new version of NAKIVO Backup & Replication is available in the QNAP App Center, you will see a green **Update** button.




5. Click the **Update** button and wait till update finishes.

## Updating Manually

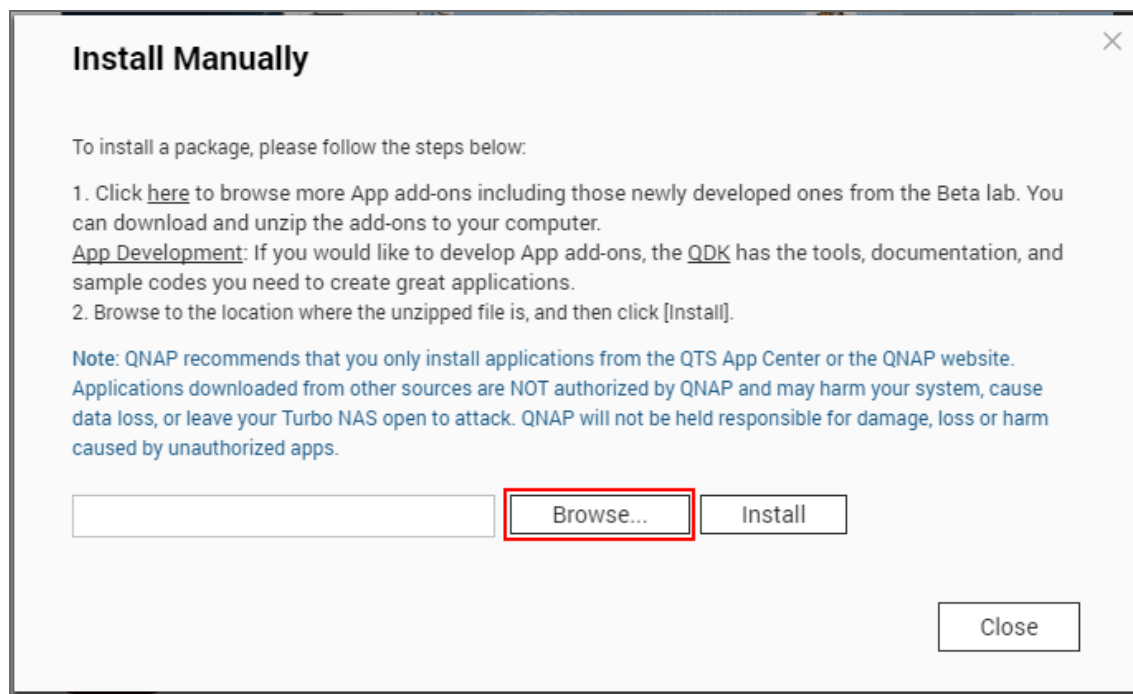
1. Download the update package from [www.nakivo.com/resources/download/update/](http://www.nakivo.com/resources/download/update/)
2. Open the QNAP Desktop in your browser by entering the IP address of your QNAP NAS.



3. Go to **App Center**.
4. Click the **Install Manually** icon.



5. Click **Browse**. In the window appears, locate the installer (.pkg 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.

Follow the steps below to update NAKIVO Backup & Replication on ASUSTOR NAS manually:

1. Open the **App Central** from the ASUSTOR NAS **Desktop**.
2. Click **Management** in the bottom left corner and click **Manual Install**.
3. The **Manual Install** pane opens to the right of the **App Central**. Click **Browse**.
4. The **Open** dialog box opens. Locate your copy of NAKIVO Backup & Replication installer for ASUSTOR NAS and click the **Open** button.
5. The **Open** dialog closes, and the **Upload** button becomes enabled. Click the **Upload** button.
6. When the upload finishes, the **About This App** dialog opens. If you are sure the requirements are met, click the **Next** button.

7. 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**.
8. The **About This App** dialog closes, and the **Installed** pane of the **App Central** opens.
9. 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](#)

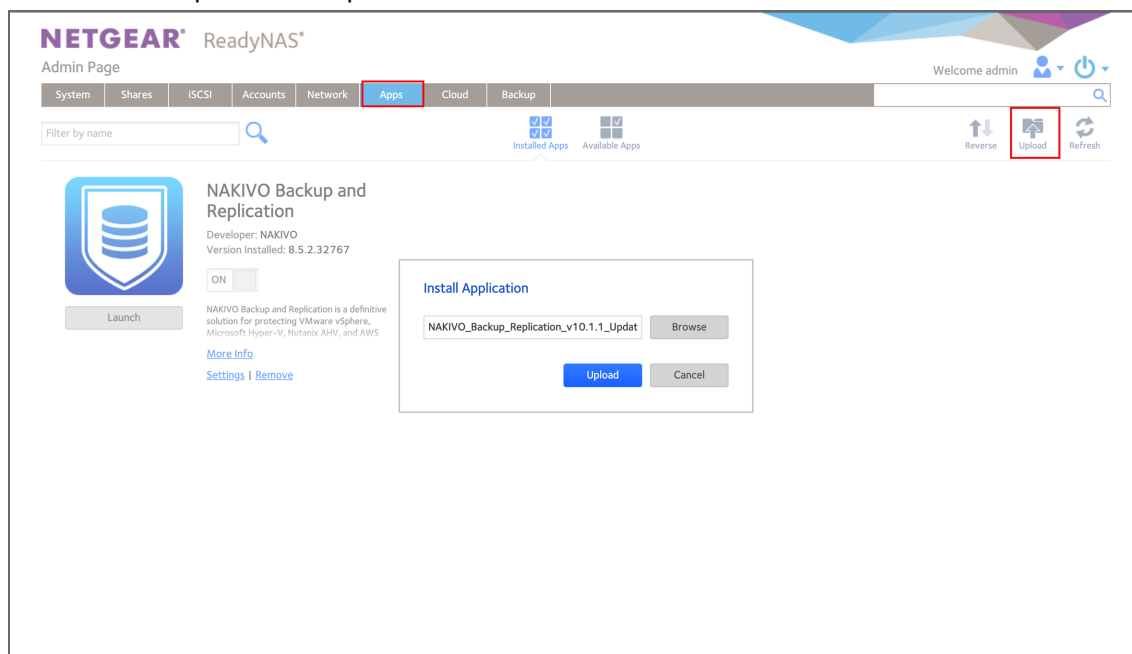
## Updating on NETGEAR ReadyNAS Manually

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. 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 update is completed.



## Updating on NETGEAR ReadyNAS via Available Apps

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 FreeNAS

## Prerequisites:

- You are logged in to the FreeNAS system with the FreeNAS GUI.
- The **Shell** button is enabled in the interface.

Follow the steps below to update NAKIVO Backup & Replication on your FreeNAS system:

1. Navigate to the **Jails** page of the FreeNAS GUI and click the jail of the NAKIVO Backup & Replication plugin to select it.
2. Click the **Shell** button to open a web shell.
3. 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://d96i82q710b04.cloudfront.net/res/product/NAKIVO_Backup_Replication_vX.X.X.XXXXX_Updater.sh
```

4. Change the updater file permission with the `chmod` command:

```
chmod +x NAKIVO_Backup_Replication_vX.X.X.XXXXX_Updater.sh
```

5. Run the updater in silent mode:

```
./NAKIVO_Backup_Replication_vX.X.X.XXXXX_Updater.sh -y -u --eula-accept
```

# Updating on Generic ARM-based Device

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 [www.nakivo.com/resources/download/update/](http://www.nakivo.com/resources/download/update/).
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](#).

## 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 FreeNAS](#)

### 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 Ddirector 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 FreeNAS

Uninstalling a plugin deletes the associated FreeNAS 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 FreeNAS:

1. Log in to the FreeNAS system using the FreeNAS GUI.
2. In the left pane of the FreeNAS 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 uninstallation 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” below](#)
- [“First Steps with NAKIVO Backup & Replication” on page 245](#)
- [“Web Interface Components” on page 249](#)
- [“Managing Jobs and Activities” on page 255](#)

## Logging in to NAKIVO Backup & Replication

- [Getting to the Login Page](#)
- [Creating a User Account](#)
- [Changing Password](#)
- [Default Password in Amazon EC2](#)
- [Passing Verification](#)

### Getting to the Login Page

To go to the NAKIVO Backup & Replication login page, open the following URL in your web browser: `https://machine_IP_or_DNS:4443`.

#### Note

If you selected a custom HTTPS port during installation, replace 4443 with the custom value.

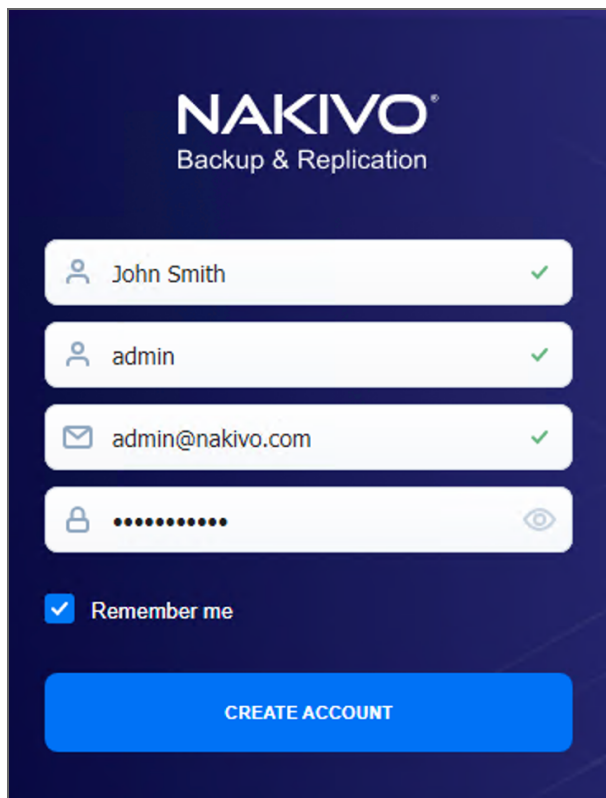
### Creating a User Account

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.



The image shows a dark blue background with the NAKIVO logo and 'Backup & Replication' text at the top. Below the logo are four white input fields, each with a green checkmark on the right. The first field contains 'John Smith' with a person icon. The second field contains 'admin' with a person icon. The third field contains 'admin@nakivo.com' with an envelope icon. The fourth field contains a masked password '.....' with a lock icon and an eye icon. Below the fields is a checked checkbox labeled 'Remember me'. At the bottom is a blue button with the text 'CREATE ACCOUNT'.


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.

## Changing Password

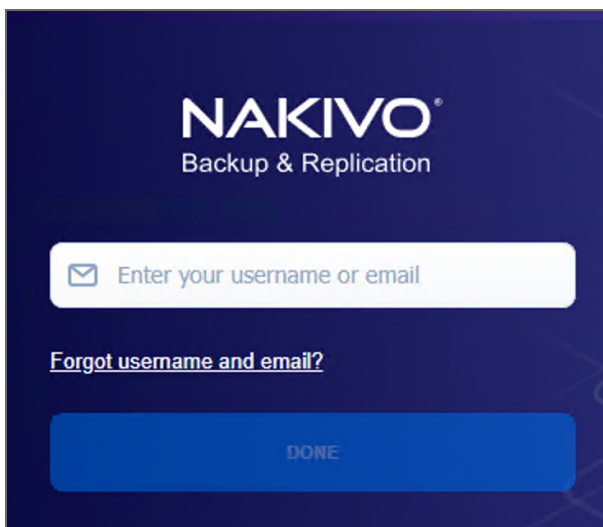
If you forget the password used to log in to NAKIVO Backup & Replication, you can restore it by following the steps below:

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**.

The screenshot shows a web interface for a user named John Smith. A modal dialog box is open for changing the password. It contains three input fields: 'Current password:', 'New password:', and 'Repeat new password:'. Below these fields are two buttons: a blue 'Change' button and a grey 'Cancel' button. In the bottom left corner of the interface, there is a 'Logout' button with a red box around it. At the bottom center, there is a copyright notice: '© 2021 NAKIVO, Inc. All Rights Reserved.'

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:

- Enter your username and click **Done**.
- 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.

- Paste the security string from the file in the appropriate field.
- 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.
- 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 FreeNAS (inside the jail): `/usr/local/nakivo/director`
  - For NETGEAR NAS: `/apps/nbr`
  - For QNAP NAS: `/share/CACHEDEV1_DATA/.qpkg/NBR`
  - For Raspberry PI: `/opt/nakivo/director`
  - 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
- Cloud
- File Share
- Physical
- Application



**1. Inventory**

**2. Transporters**

**3. Repositories**

**What is a Transporter?**

A transporter is the service that runs the actual backup, replication, and recovery tasks. Transporters are also responsible for backup repository management.

One transporter called the "Onboard transporter" is already included in your installation. If you run a large environment or have multiple sites, you may need additional transporters.

**How do I add more transporters?**

**Generic Transporter**

- Download the transporter installer or virtual appliance from <https://www.nakivo.com>.
- Deploy the virtual appliance or install the transporter on an existing machine.
- To add the installed transporter, go to Settings > Transporters.

**VMware vSphere Transporter**

- To add your VMware vCenter or standalone ESXi host, go to Settings > Inventory.
- To deploy a new VMware vSphere transporter, go to Settings > Transporters.

**Nutanix AHV Transporter**

- To add your Nutanix AHV cluster, go to Settings > Inventory.
- To deploy a new Nutanix AHV transporter, go to Settings > Transporters.

**Amazon EC2 Transporter**

- To add your AWS account, go to Settings > Inventory.
- To deploy a new Amazon EC2 transporter, go to Settings > Transporters.


**Microsoft Hyper-V Transporter**

- When you add a Microsoft Hyper-V host to Inventory, a transporter is automatically installed on that host.
- You do not need to install an additional Microsoft Hyper-V transporter manually.

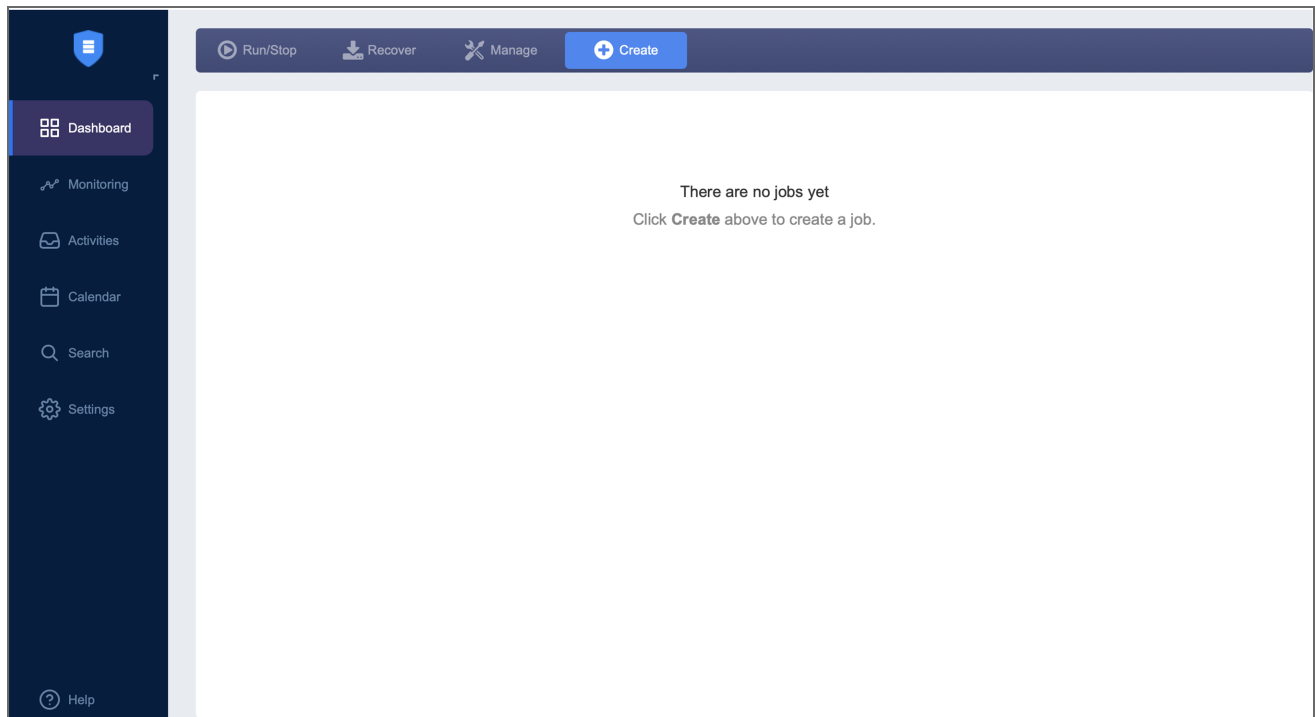
[Learn more](#) **Got It**

[Import system configuration](#) **Back** **Next**

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 **Dashboard** of the application opens. Proceed with creating your **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.

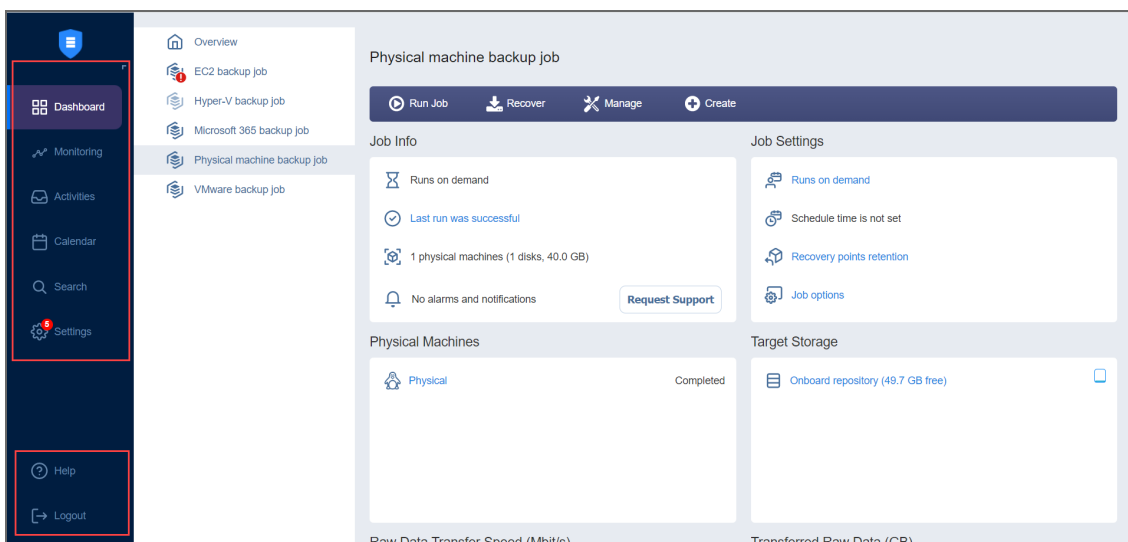
## Web Interface Components

The interface of NAKIVO Backup & Replication consists of the following components:

- [Main Menu](#)
- [Dashboard](#)
- [Monitoring](#)
- [Activities](#)
- [Calendar](#)
- [Search](#)
- [Settings](#)
- [Help Menu](#)
- [Online Chat Dialog](#)
- [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 jobs dashboard, activities, calendar, global search, and product settings. It also contains the **Help** menu and **Log Out** button.

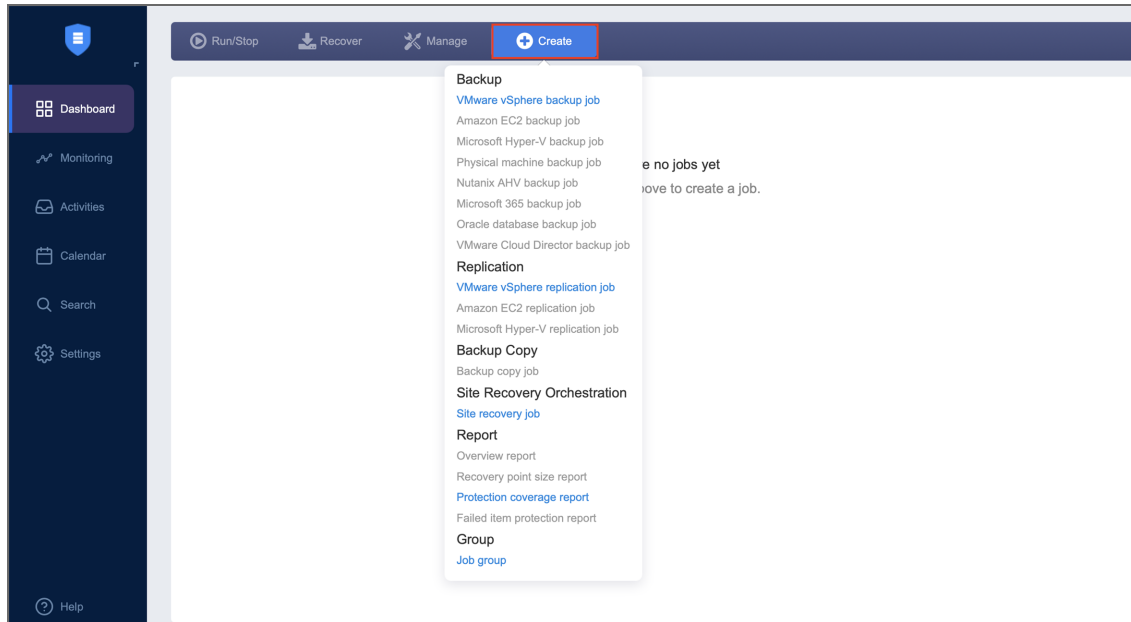




# Dashboard

Using the **Dashboard**, you can:

- View, run, and stop jobs on demand
- Recover files, objects and entire sites
- Manage jobs
- Create backup and replication jobs
- 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.

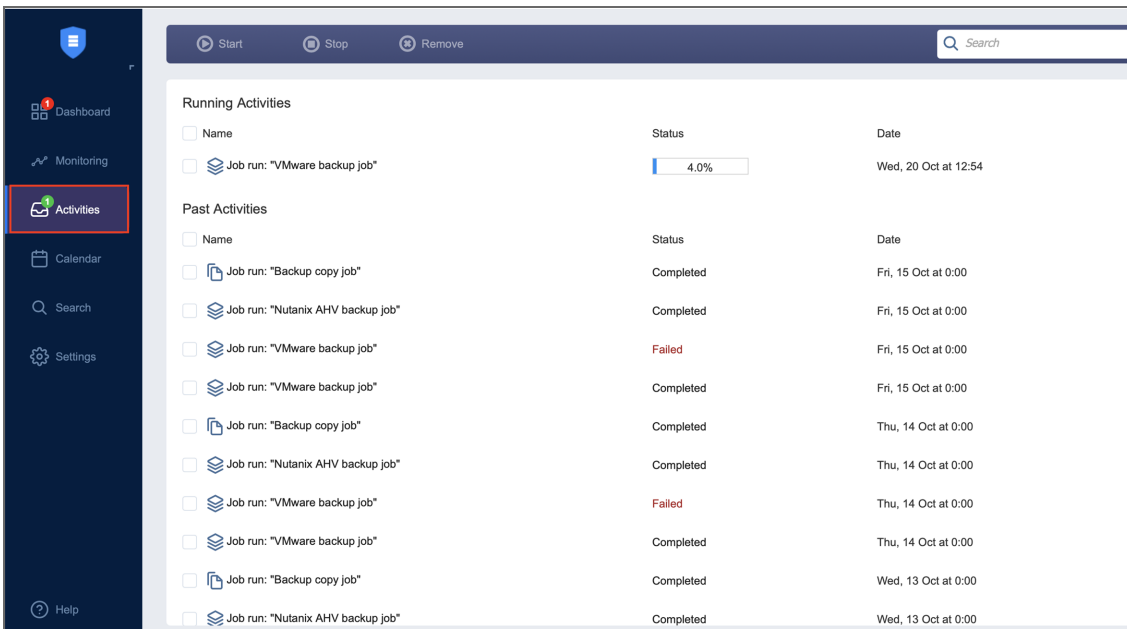


## Activities

The **Activities** page displays a list of all running and past activities, such as:

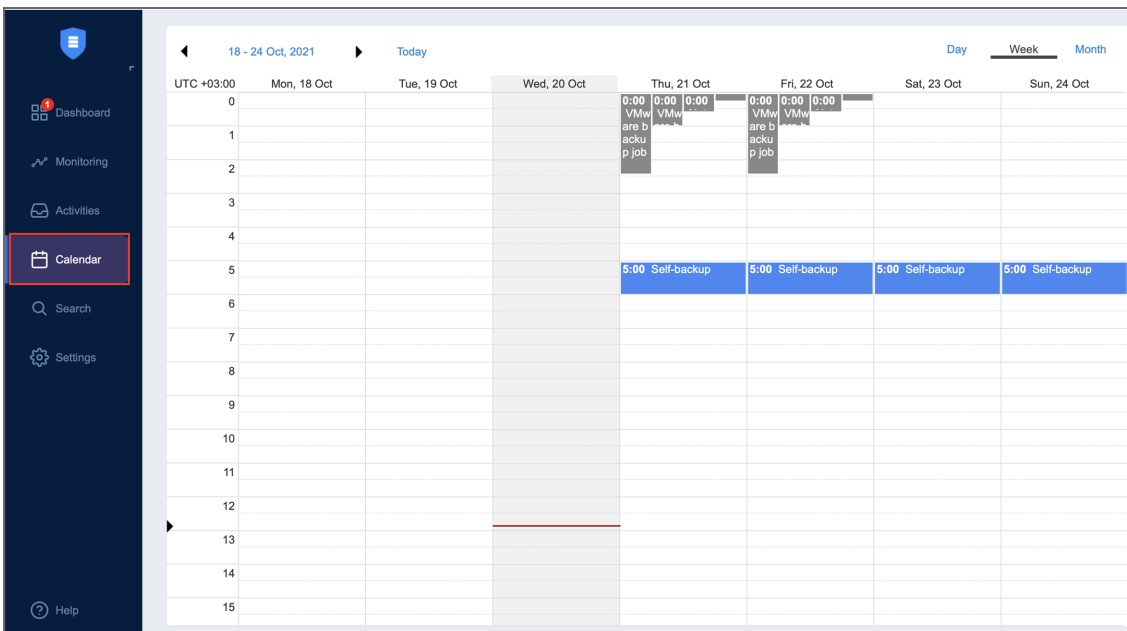
- Job run
- Repository Self-Backup
- File download
- Application object download
- Universal object recovery
- Repository [space reclaim](#)
- Repository [self-healing](#)
- Repository backup verification
- [Backup export](#)
- [Tape-specific activities](#), namely: scanning, erasing, and reading Tape
- Other

For further details and information, refer to [“Managing Activities” on page 274](#).



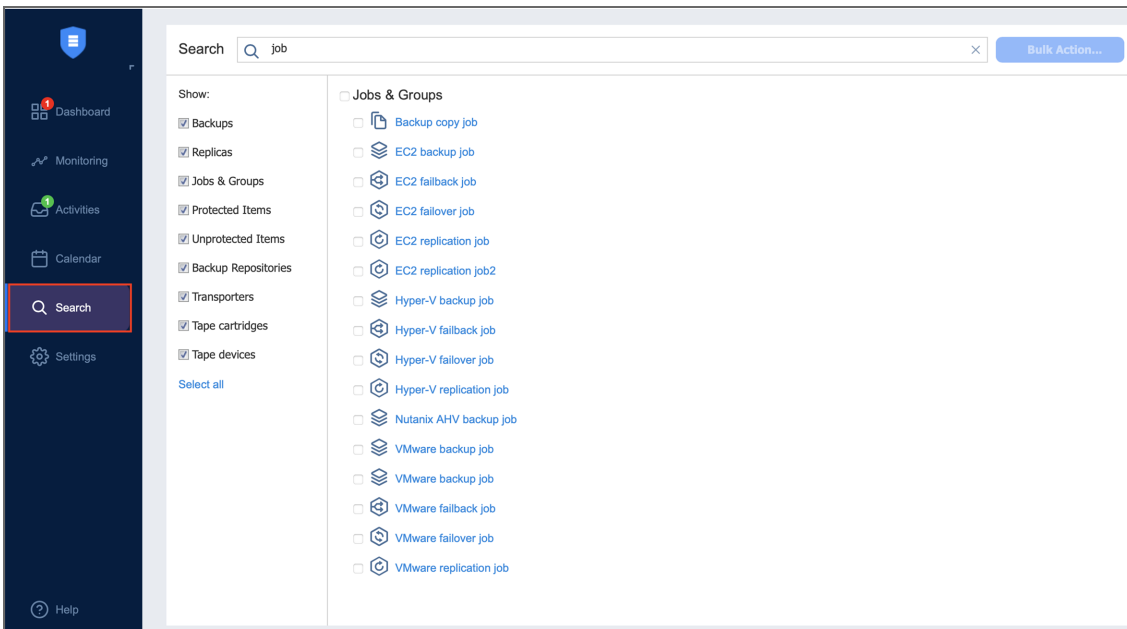
## Calendar

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 278](#).



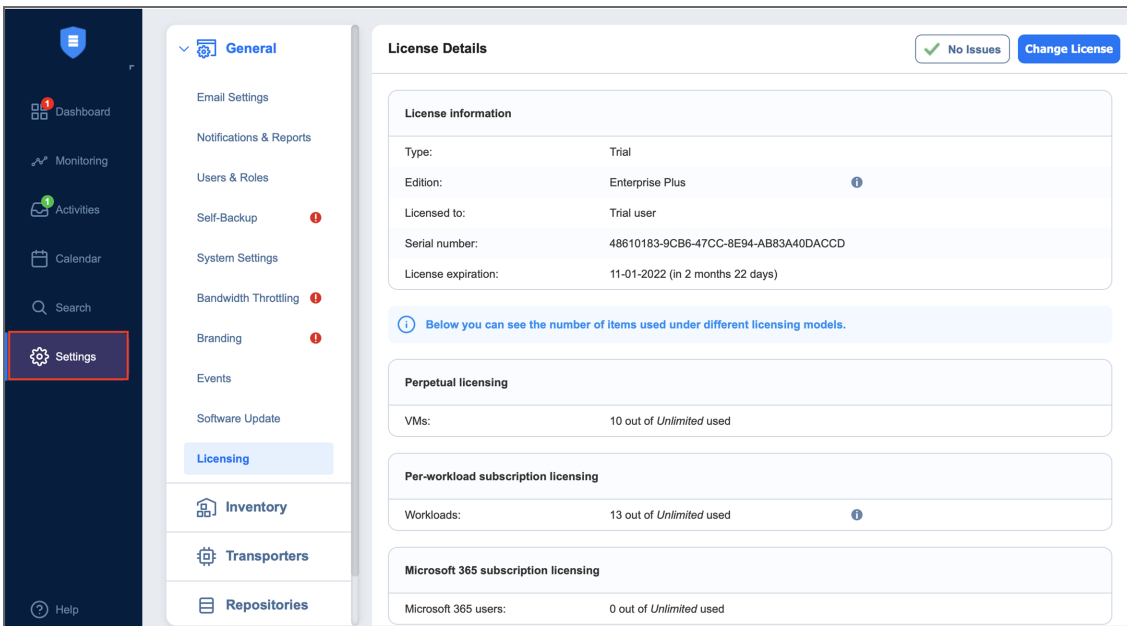
## Search

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 280](#).



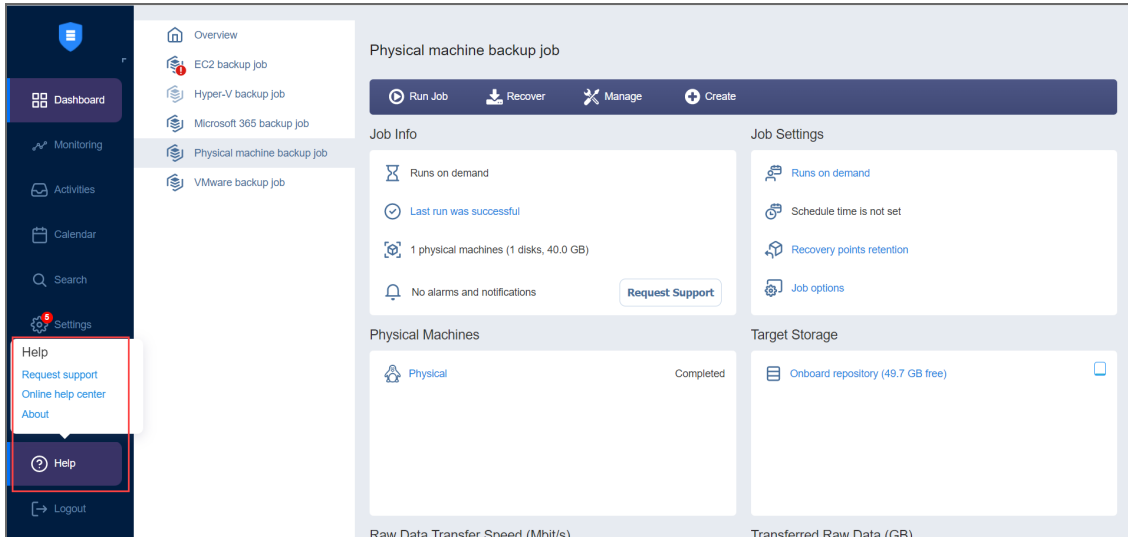
## Settings

On the **Settings** page, you can configure NAKIVO Backup & Replication [General](#), [Inventory](#), [Transporters](#), [Repositories](#), and [Tape](#) settings. Refer to [“Settings” on page 284](#) for more detailed information.



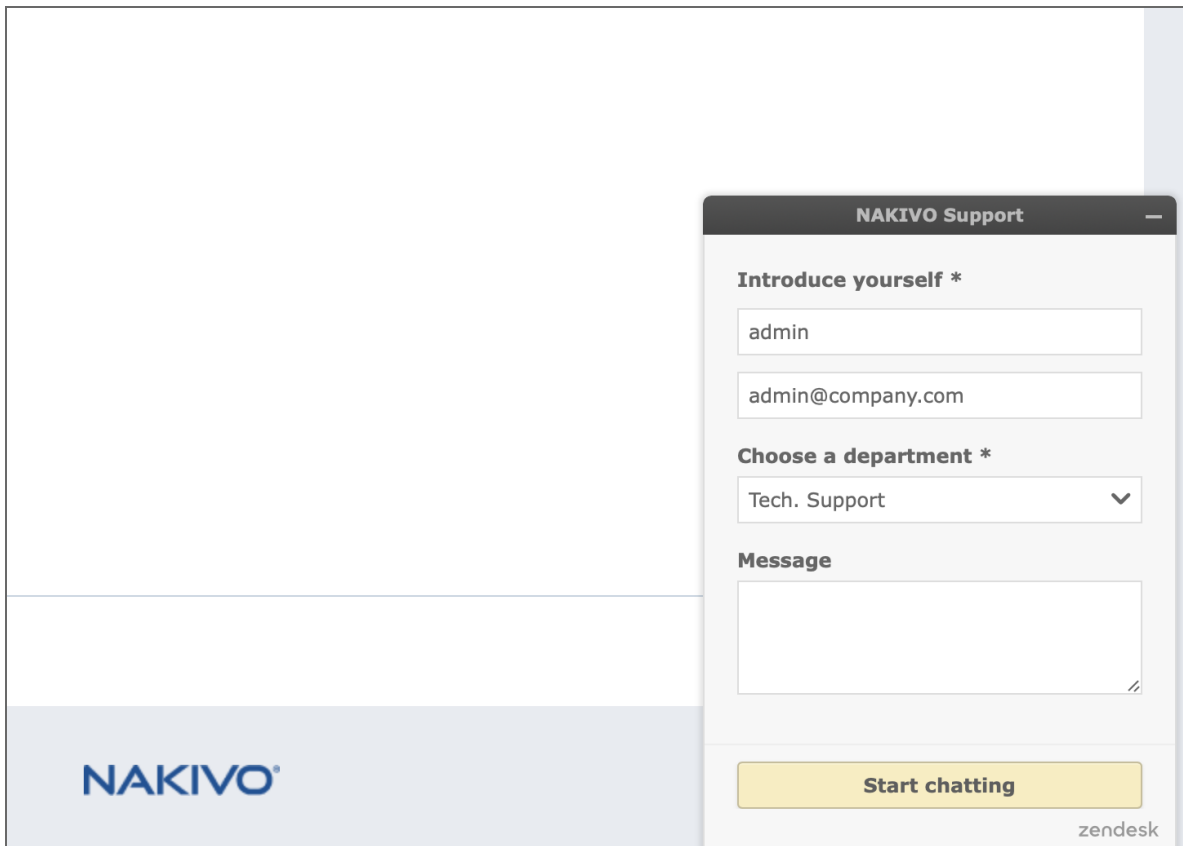
# Help Menu

Use the **Help** menu to request technical support and access the NAKIVO online [help center](#). If you are evaluating NAKIVO Backup & Replication, you may also use 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.



## 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 302](#) 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:

- [“Running Jobs on Demand” on the next page](#)
- [“Managing Jobs” on page 259](#)
- [“Managing Job Policies” on page 266](#)
- [“Managing Policy Rules” on page 269](#)

- [“Job Alarms and Notifications” on page 273](#)
- [“Managing Activities” on page 274](#)
- [“Using Calendar” on page 278](#)
- [“Using Global Search” on page 280](#)

## Running Jobs on Demand

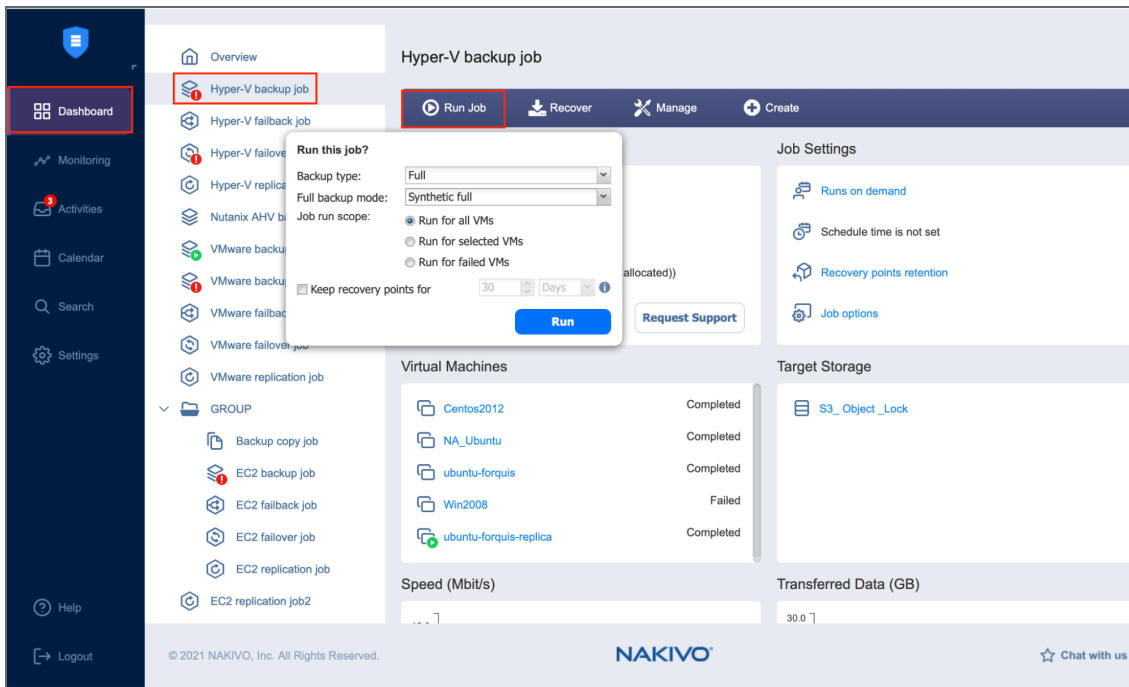
Use the **Dashboard** to start and stop jobs on demand.

- [Starting Jobs](#)
- [Stopping Jobs](#)
- [Managing Grouped Jobs](#)

### Starting Jobs

To start a job, follow the steps below:

1. Go to the **Dashboard**, select the job from the list of jobs, and click **Run Job**.
  - a. Choose one of the following options:
    - **Run for all VMs/backups/physical machines/items:** The job will run for all VMs/backups/physical machines/ items.
    - **Run for selected VMs/backups/physical machines/items:** The job will run for the VMs/backups/physical machines/items you select.
    - **Run for failed VMs/backups/physical machines/items:** If applicable, the job will run for previously failed VMs/backups/physical machines/ items only.
    - **Run for stopped VMs/backups/physical machines/items:** If applicable, the job will run for previously stopped VMs/backups/physical machines/items only.
  - b. If backups in the Backup Repository selected for your job are stored in separate files, you will have to choose between the following backup types:
    - **Incremental:** Your job will create an incremental backup.
    - **Full:** Your job will create a full backup. If you choose this option, please choose the full backup mode:
      - **Synthetic full:** The application will first perform an incremental backup – that is, will transfer only the VM data that changed since the last backup, – and will then transform the available data into a full backup.
      - **Active full:** Will read all source VM data and transfer it to the backup repository.
2. You can modify the retention period for job recovery points. The recovery points created as a result of this particular job run are kept for the specified period of time and then expire. The expired recovery points are removed during the following job run. Select **Keep recovery points for** and specify the relevant time period.
3. Click the **Run** button to confirm your operation.



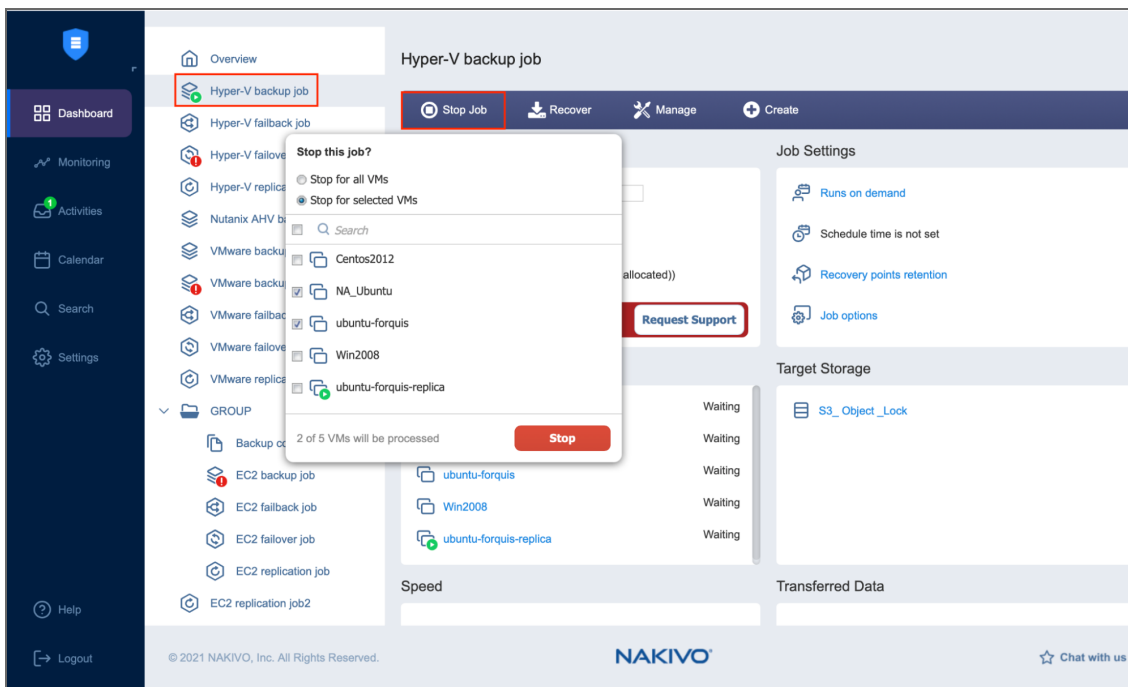
The product will close the dialog box and start running your job.

## Stopping Jobs

To stop a job that is currently running, follow the steps below:

1. Go to the **Dashboard**, select the job from the list of jobs, and click **Stop Job**.
2. In the dialog that opens, choose either of the following:
  - **Stop for all VMs/backups/physical machines/items:** Your job will stop for all VMs/backups/physical machines/items.
  - **Stop for selected VMs/backups/physical machines/items:** Your job will stop for the VMs/backups/physical machines/items you select.
3. Click the **Stop** button in the dialog to confirm your operation.






The product will close the dialog box and stop 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 **Dashboard**, select the desired job group and click **Run/Stop**. To manage all jobs and groups at once, select **Overview** 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. If you're running backup jobs, specify the retention period for your recovery points with the **Keep recovery points** for option.
  - 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 operation.



## Managing Jobs

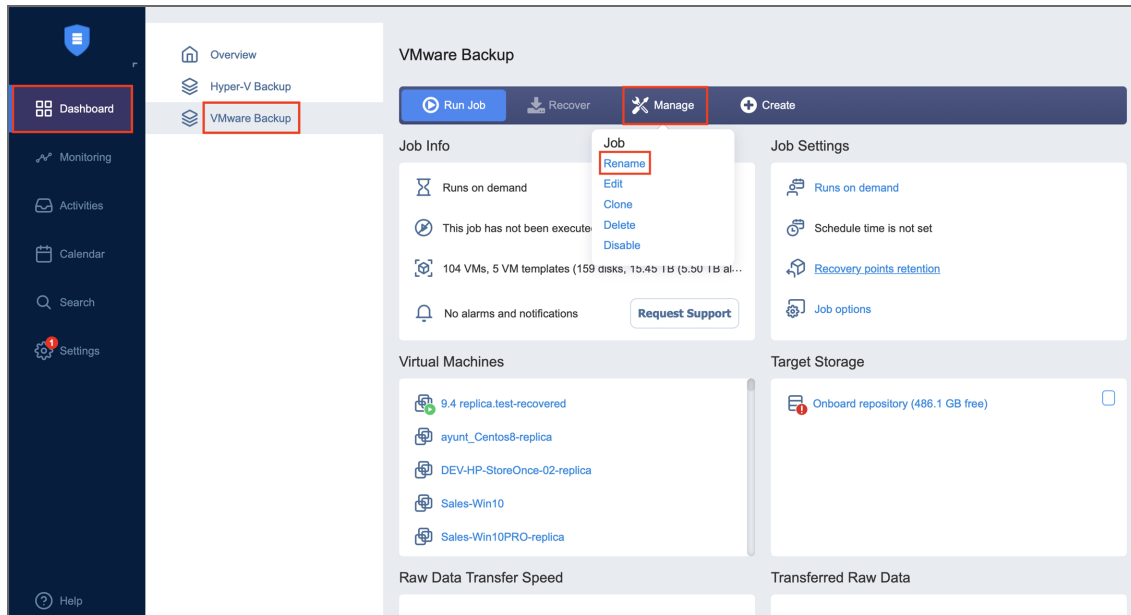
Using the **Dashboard**, you can easily manage your jobs. Go to the **Manage** menu to rename, edit, delete and enable/disable jobs.

- [Renaming Jobs](#)
- [Editing Jobs](#)
- [Cloning Jobs](#)
- [Deleting Jobs](#)
- [Disabling and Enabling Jobs](#)
- [Grouping Jobs](#)
  - [Creating Groups](#)
- [Creating Job Reports](#)

### Renaming Jobs

1. From the list of jobs, select the job you wish to rename.
2. On the **Dashboard**, click **Manage**.
3. Click **Rename**.
4. In the dialog box that opens, specify the new name for the job.

## 5. Click **Rename**.



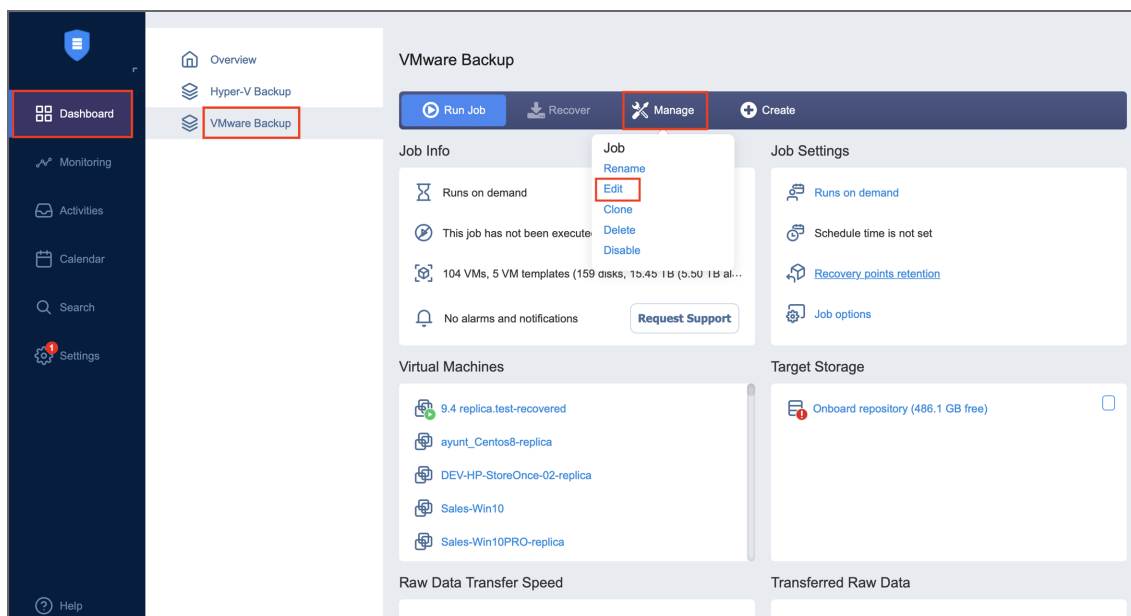
### Note

You can also rename jobs by right-clicking on a job and selecting **Rename** from the **Manage Job** menu.

## Editing Jobs

To edit a job, follow the steps below:

1. Select the job you wish to edit from the list of jobs.
2. On the **Dashboard**, click **Manage**.
3. Click **Edit**.



4. In the **Edit** wizard, click the necessary page to open it for editing.
5. Make the required changes and then 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 right-clicking on a job and selecting **Edit** from the **Manage Job** menu.

## Cloning Jobs

To clone a job, follow the steps below:

1. Select the job you would like to clone from the list of jobs.
2. On the **Dashboard**, click **Manage**.
3. Click **Clone**.



## Note

You can also clone jobs by right-clicking on a job and selecting **Clone** from the **Manage Job** menu.

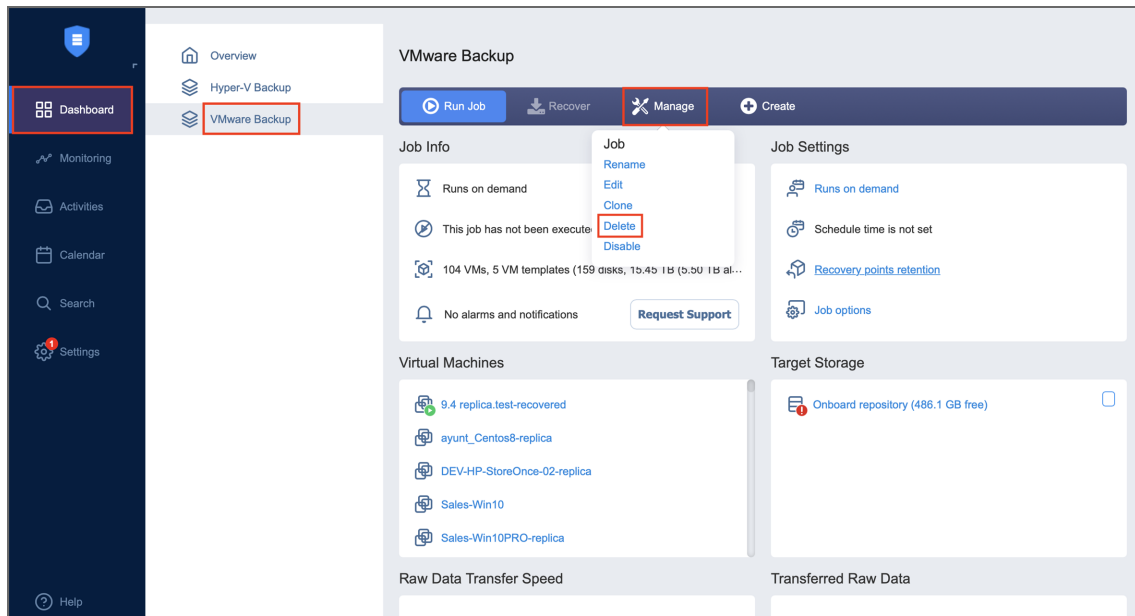
## Deleting Jobs

To delete a job follow the steps below:

1. Select the job you want to delete from the list of jobs.
2. On the **Dashboard**, click **Manage**.
3. Click **Delete**.
4. From the dialog box that opens, select one of the following:
  - **Delete job and keep backups**
  - **Delete job and keep backups**
5. Click **Delete**

## Notes

- You can also delete jobs by right-clicking on a job and selecting **Delete** from the **Manage Job** menu.
- Backups can also be **deleted** from Backup Repositories.

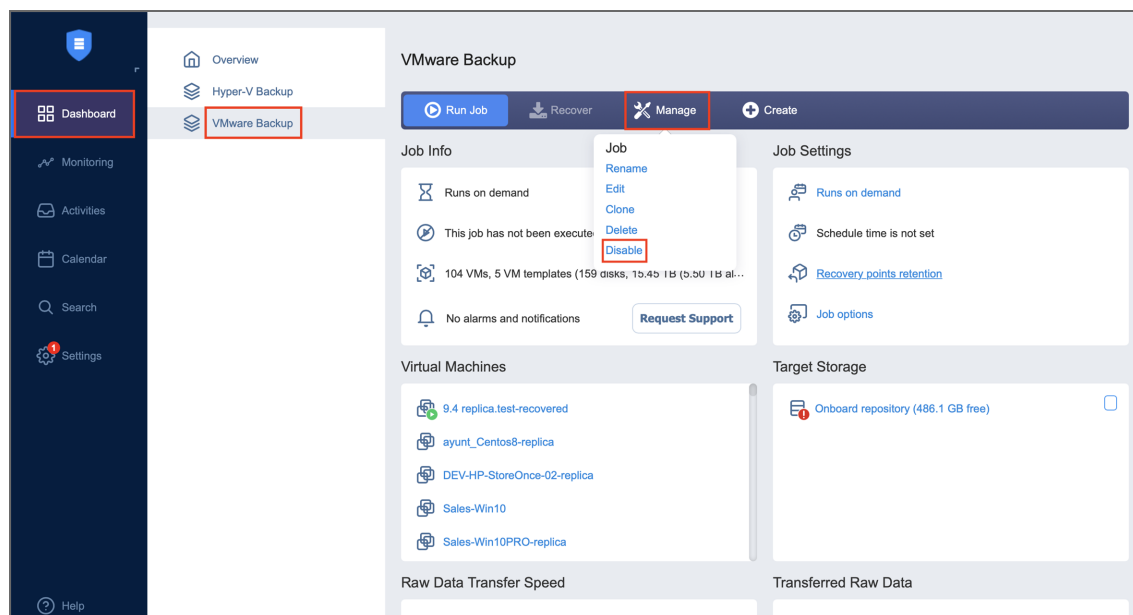


## Disabling and Enabling Jobs

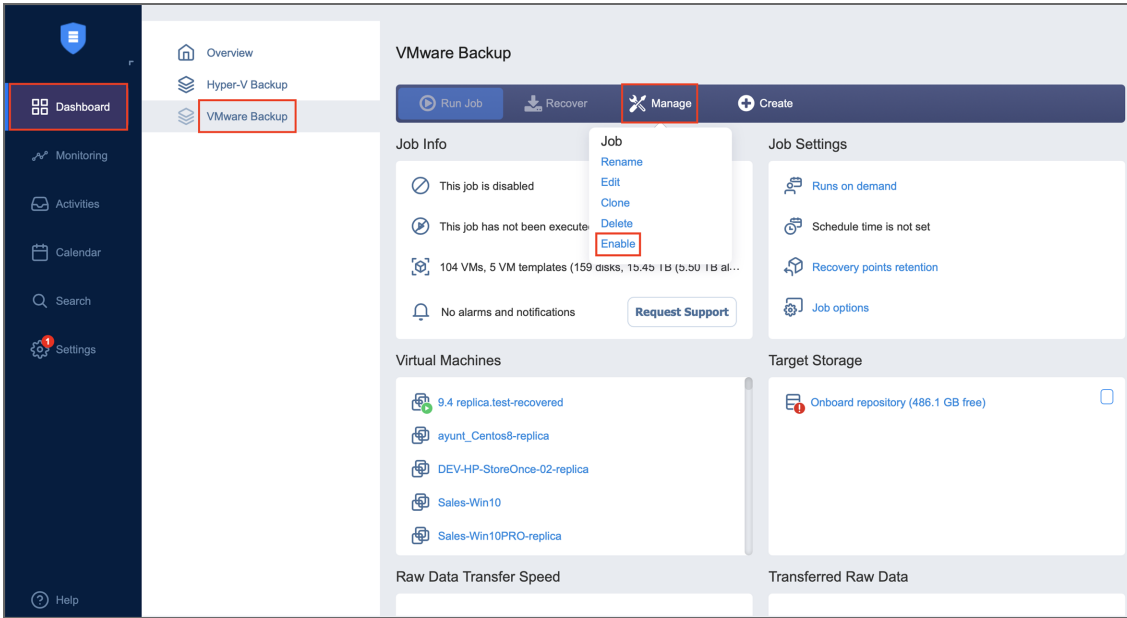
NAKIVO Backup & Replication provides you with the ability to disable jobs. A disabled job does not run on a schedule, nor can it be run on demand.

To disable a job, follow the steps below

1. From the list of jobs, select the job you want to disable.
2. On the **Dashboard**, click **Manage**.
3. Click **Disable**.

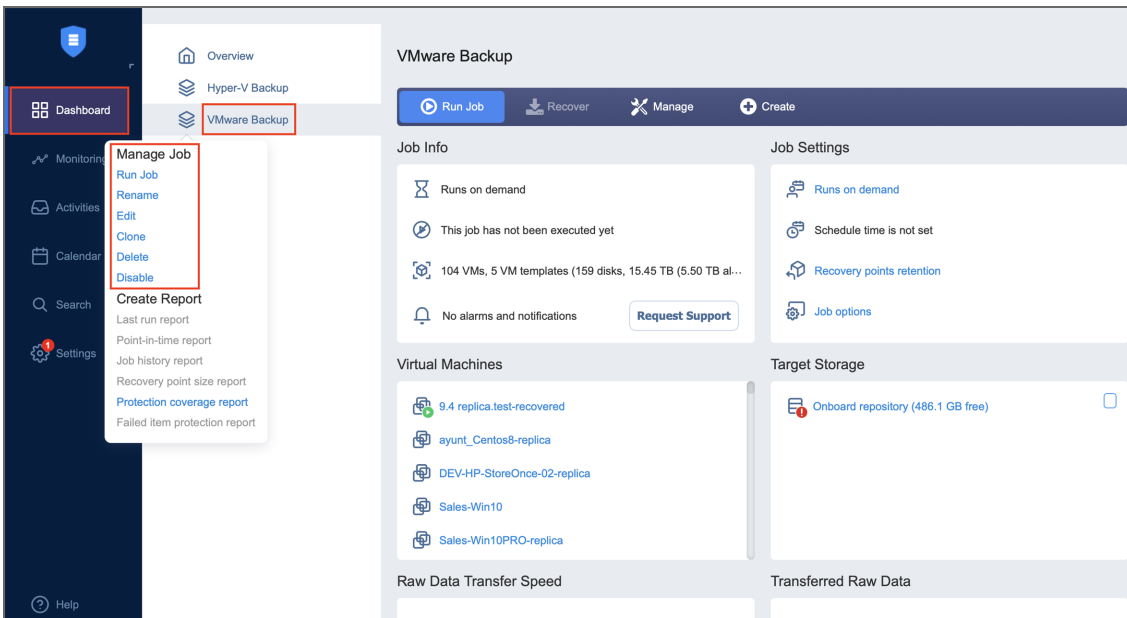


To enable a job, select **Enable** from the **Manage** menu.



### Note

You can also manage jobs by right-clicking on a job and selecting the desired action from the **Manage Job** menu.



## Grouping Jobs

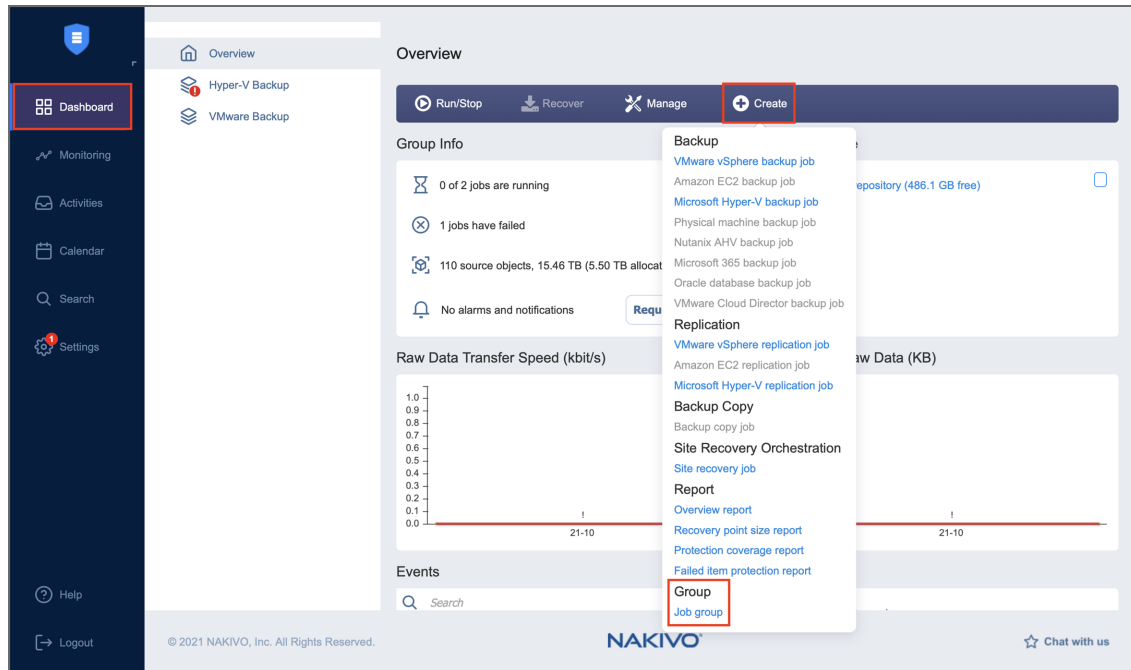
Groups are folders which allow you to:

- 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. On the **Dashboard**, click **Create** and then click **Job group**.
2. Type in the group name in the dialog box that appears and click **OK**.



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 appears. Confirm the group deletion when prompted to do so. Note that when deleting a group, its jobs are not deleted and are moved to the parent group (or 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 appears, select the jobs you wish 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 appears, select the jobs you would like to stop and click **Stop Jobs**.

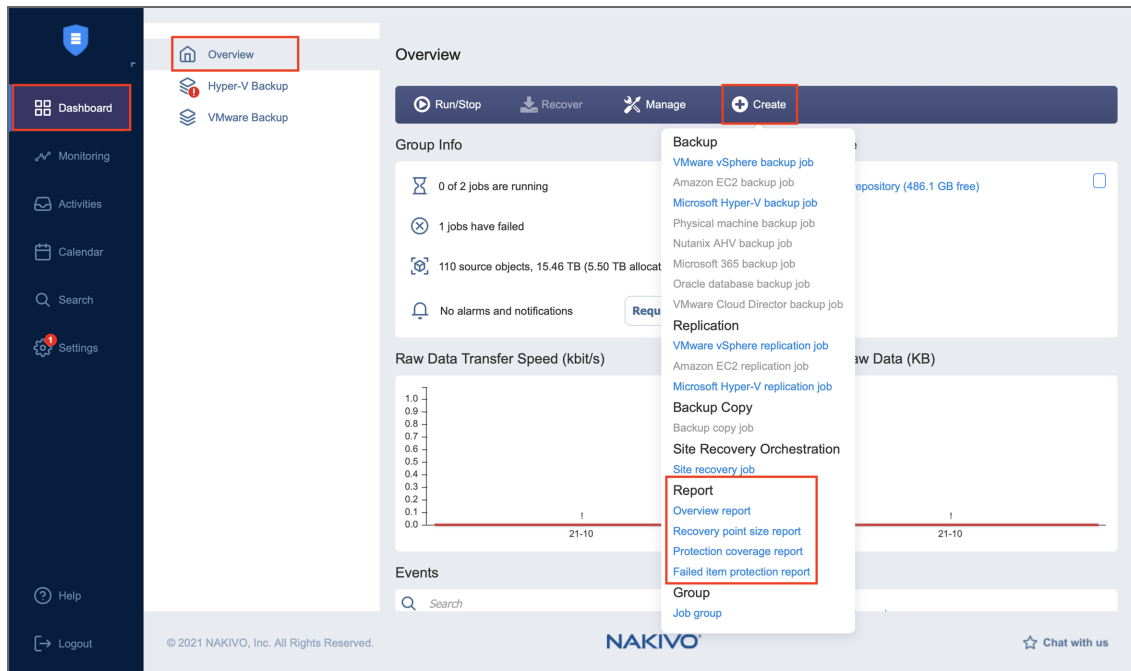
## Creating Job Reports

To create a general report for all your jobs:

1. Select **Overview** on the **Dashboard**.
2. Click **Create**.
3. Choose one of the following reports in the **Report** section:

- **Overview report:** Contains information about the status and errors of all your jobs.
- **Recovery point size report:** Contains information regarding the sizes of recovery points of 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 PDF or CSV formats for your **Protection coverage report** and click **Create**.
- **Failed item protection report:** Contains information about all VMs and instances which had failed to be protected by backup and/or replication jobs, and the error message. Select the date range for your **Failed item protection report** and click **Create**.

4. Choose a location to save the report and click **Save**.

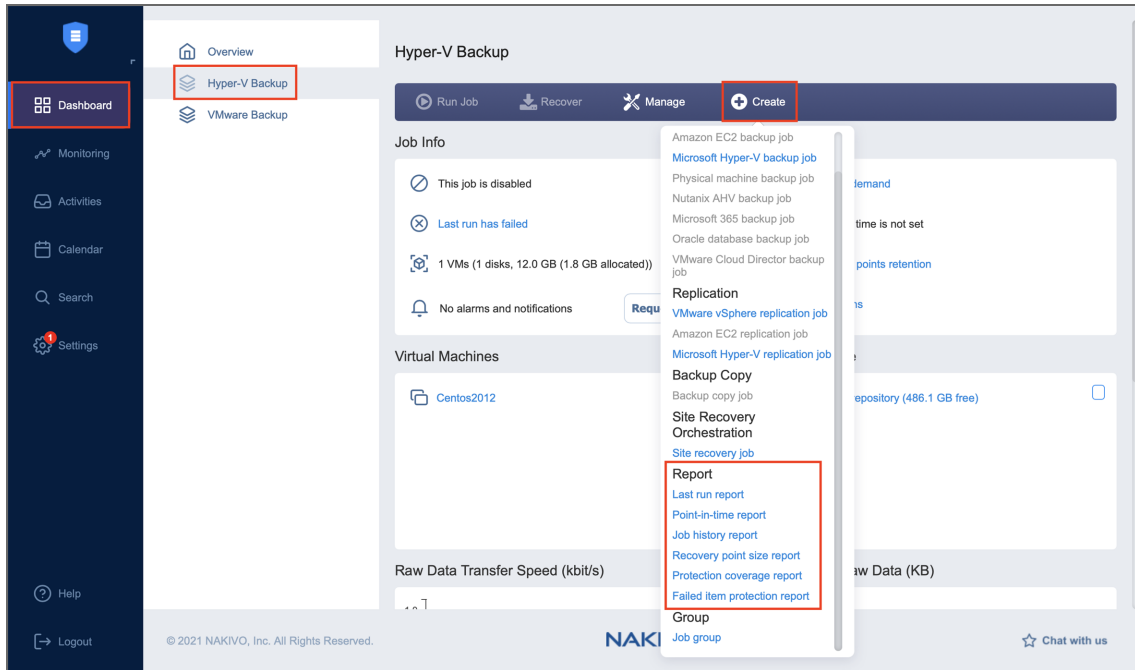


To generate reports from for an individual job, do the following:

1. Go to the list of jobs.
2. Select the job that you need to generate a report for and right-click on 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.
  - **Point-in-time Report:** Provides data on a particular job run. To generate a report, pick a date in the resulting pop-up and click **Create**.
  - **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 resulting pop-up and click **Create**.
  - **Recovery point size 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.
- **Failed item protection report:** Contains information about job objects for which processing had 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.



## 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 269](#) 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 the next page](#) for details. Make sure that at least one item matches the available set of policy rules.
5. Save your job.

The screenshot displays a configuration window with five tabs: 1. Source, 2. Destination, 3. Schedule, 4. Retention, and 5. Options. The 'View' dropdown is set to 'Policy'. Below this, there are three radio buttons for rule matching: 'Include items if ANY rule is matched' (selected), 'Include items if ALL rules are matched', and 'Include items if ANY rule is matched'. The 'Search by:' dropdown is set to 'VM name', and the 'Which:' dropdown is set to 'Contains'. The 'Search criteria:' input field contains '2012'. A '+ Add rules' link is visible. On the right, a 'Policy Container' is shown with a list item 'Centos2012'. At the bottom, there are 'Cancel', 'Save', and 'Save & Run' buttons. A license expiration notice 'License expires in 2 months 21 days' is at the bottom left, and a prompt 'Drag items to set processing priority' is at the bottom right.

## Editing Job Policy

Follow the steps below to edit a job policy:

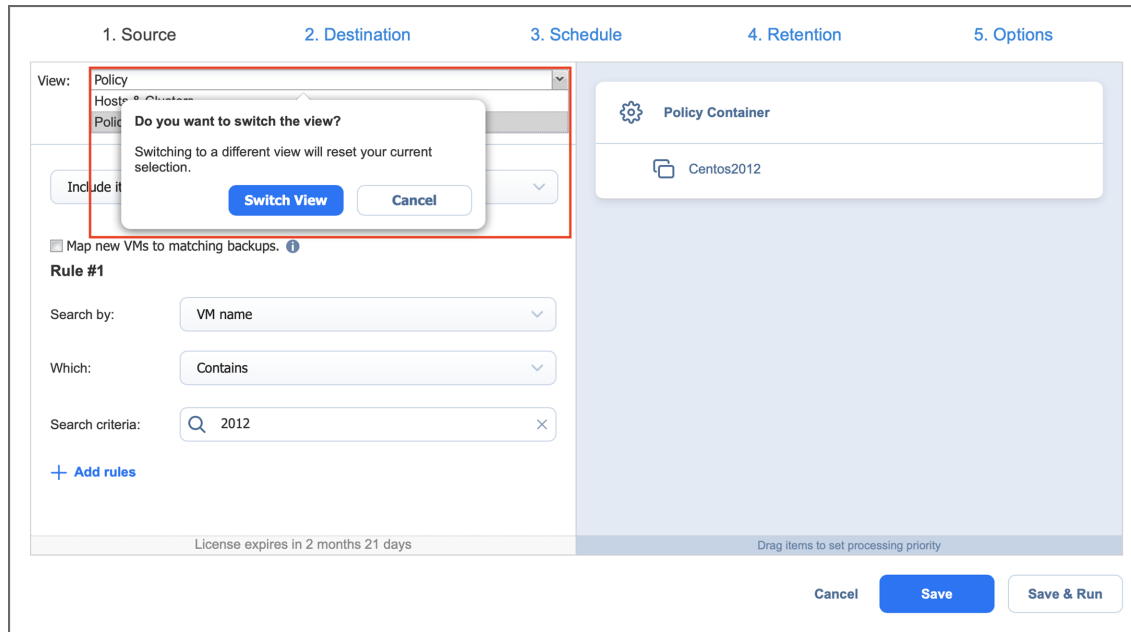
1. Make sure your job is opened in the **Policy** view.
2. Change the necessary parameters of your job policy:
  1. **Condition.**
  2. Add, edit or delete policy rules. Refer to [“Managing Policy Rules” on the next page](#) for details.
3. Save your job.

## 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:

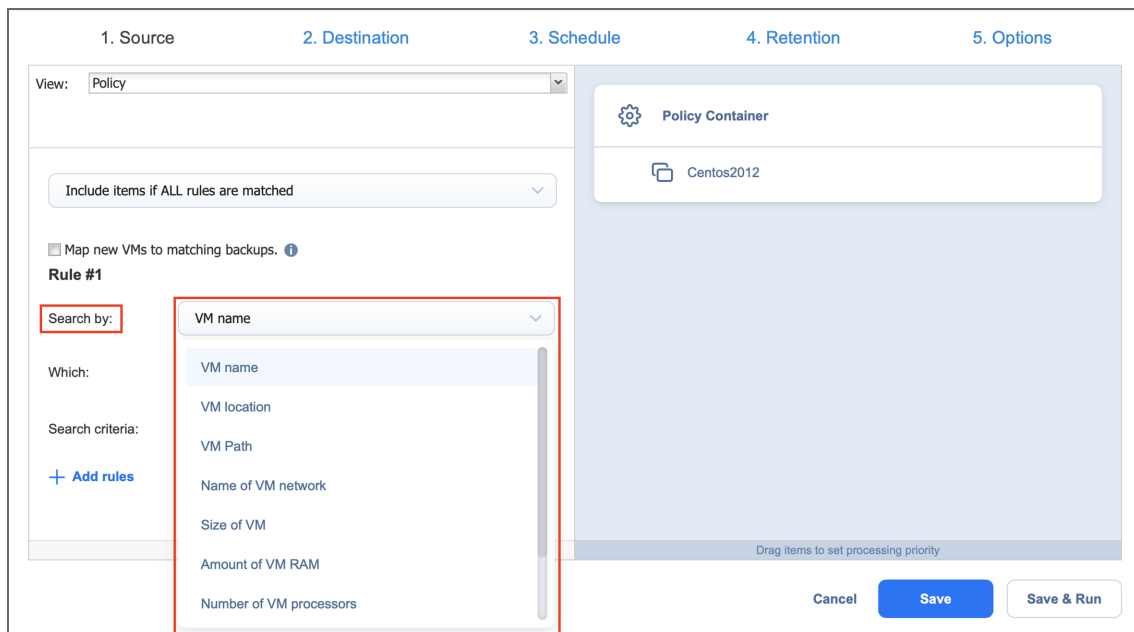
1. **Search by:** A drop-down list with the following search criteria:
  - **VM/VM Template/Instance/Backup/Replica/Machine name:** The rule is to be applied based on the name of the object.
  - **VM/Instance tag:** The rule is to be applied based on the tag of the object.
  - **VM/VM Template/Instance/Replica/Backup 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**

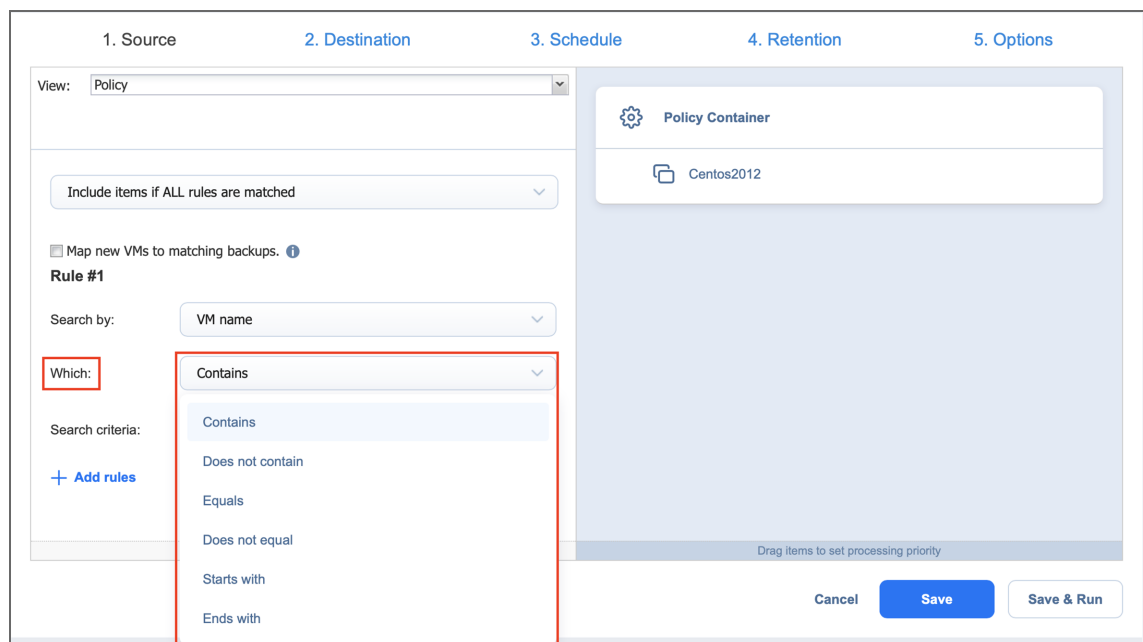
The objects corresponding to the above criteria are as follows:

- **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 name, Name of VM network/Name of VM Template network/Name of subnet, Name of VM datastore/Name of VM Template datastore/VM 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 RAM, Number of VM processors/CPU sockets/Instance virtual CPUs, and Size of VM/VM Template/Instance**, 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 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.

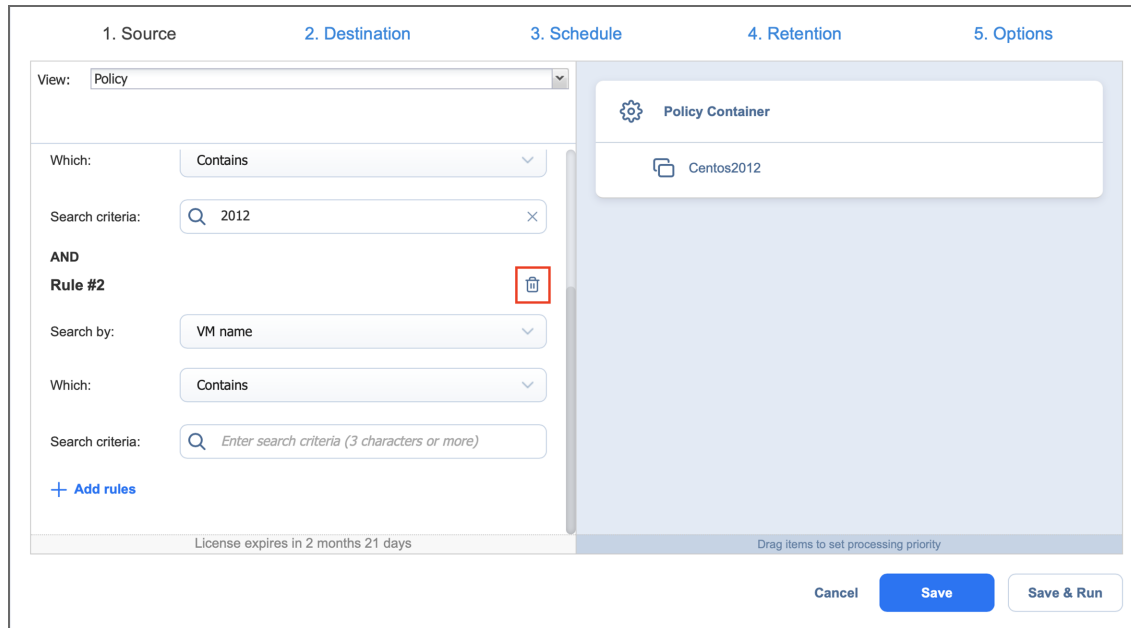
The screenshot shows a multi-step wizard for configuring a policy rule. The steps are: 1. Source, 2. Destination, 3. Schedule, 4. Retention, and 5. Options. The current step is '3. Schedule'. The 'View' is set to 'Policy'. The configuration includes a dropdown for 'Include items if ALL rules are matched', a checkbox for 'Map new VMs to matching backups', and a section for 'Rule #1'. Under 'Rule #1', there are fields for 'Search by' (set to 'VM name'), 'Which' (set to 'Contains'), and 'Search criteria' (set to '2012'). A red box highlights the '+ Add rules' button. At the bottom, there are 'Cancel', 'Save', and 'Save & Run' buttons. A license expiration notice is visible at the bottom left, and a drag instruction is at the bottom right.

## 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

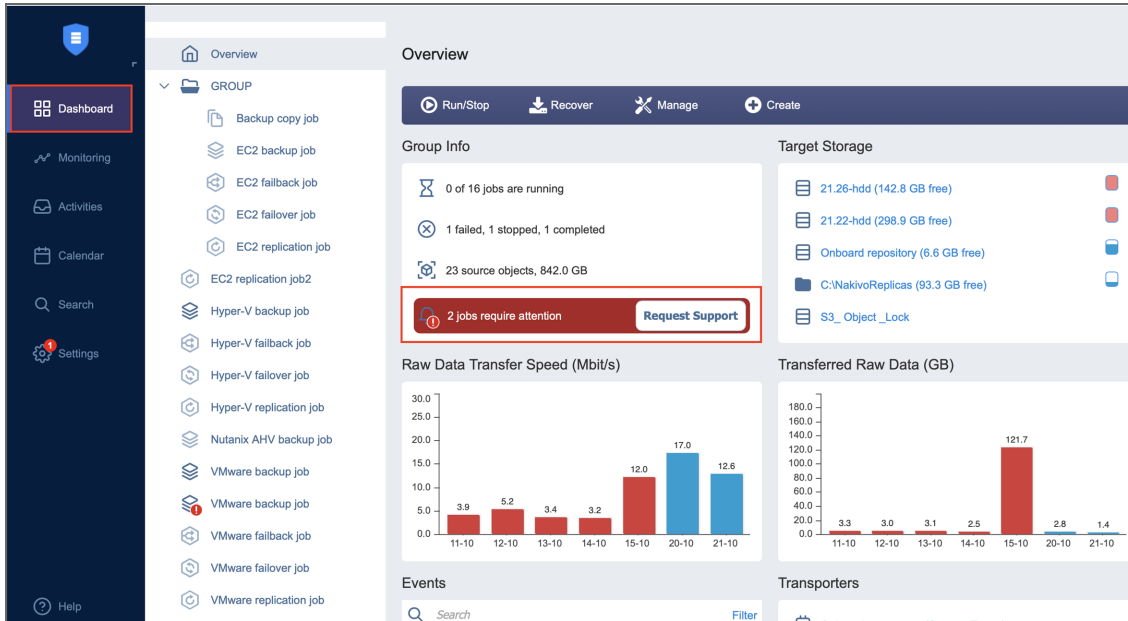
For details, refer to the following sections:

- [Viewing Alarms and Notifications](#)
- [Dismissing Alarms and Notifications](#)



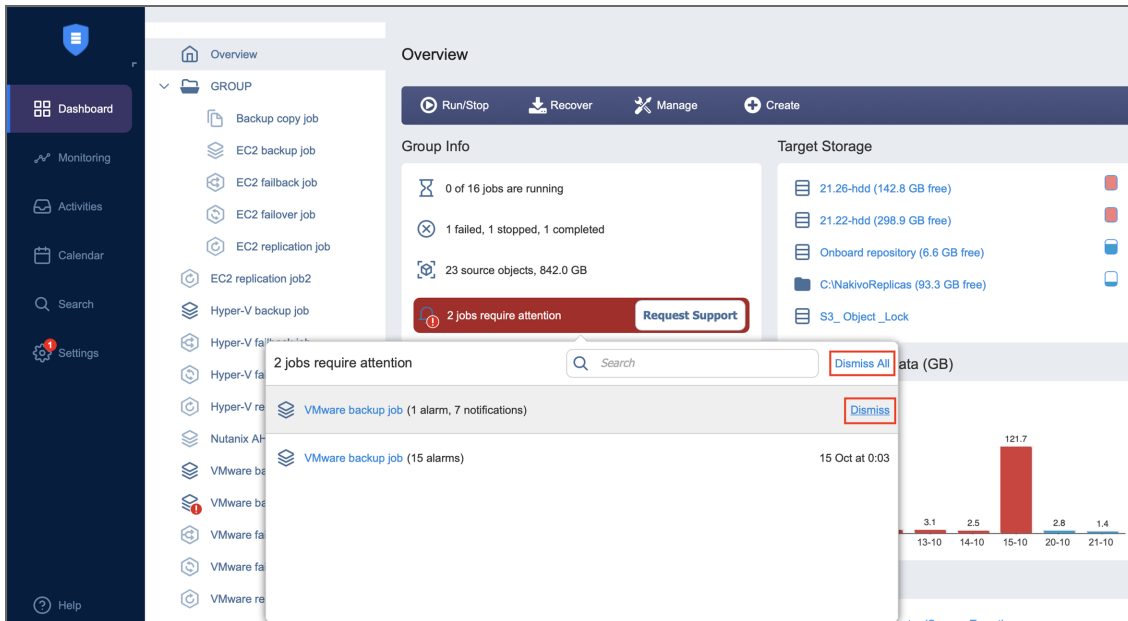
## Viewing Alarms and Notifications

To view alarms and notifications, click the red/yellow box in the Job/Group Info widget.



## Dismissing Alarms and Notifications

To dismiss all alarms and notifications in a job, click **Dismiss All**. To dismiss an individual alarm or notification, hover the mouse pointer over the alarm or notification and click **Dismiss**.



## 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.

The screenshot shows the 'Activities' dashboard interface. On the left is a dark sidebar with navigation options: Dashboard, Monitoring, **Activities** (highlighted with a red box), Calendar, Search, Settings, and Help. The main content area has a top bar with 'Start', 'Stop', and 'Remove' buttons, and a search field. Below this, there are two sections: 'Running Activities' and 'Past Activities'. The 'Running Activities' section is currently empty, displaying 'There are no running activities.' The 'Past Activities' section contains a table with the following data:

Name	Status	Date
Job run: "VMware backup job"	Completed	Thu, 21 Oct at 0:00
Job run: "VMware backup job"	Completed	Wed, 20 Oct at 12:54
Job run: "Backup copy job"	Completed	Fri, 15 Oct at 0:00
Job run: "Nutanix AHV backup job"	Completed	Fri, 15 Oct at 0:00
Job run: "VMware backup job"	Failed	Fri, 15 Oct at 0:00
Job run: "VMware backup job"	Completed	Fri, 15 Oct at 0:00
Job run: "Backup copy job"	Completed	Thu, 14 Oct at 0:00
Job run: "Nutanix AHV backup job"	Completed	Thu, 14 Oct at 0:00
Job run: "VMware backup job"	Failed	Thu, 14 Oct at 0:00
Job run: "VMware backup job"	Completed	Thu, 14 Oct at 0:00


## 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.



## 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

<input checked="" type="checkbox"/> Name	Status	Date	
<input checked="" type="checkbox"/> Job run: "Hyper-V backup job"	96.4%	Thu, 21 Oct at 13:46	

Past Activities

<input type="checkbox"/> Name	Status	Date
<input type="checkbox"/> Job run: "VMware backup job"	Completed	Thu, 21 Oct at 0:00
<input type="checkbox"/> Job run: "VMware backup job"	Completed	Wed, 20 Oct at 12:54
<input type="checkbox"/> Job run: "Backup copy job"	Completed	Fri, 15 Oct at 0:00
<input type="checkbox"/> Job run: "Nutanix AHV backup job"	Completed	Fri, 15 Oct at 0:00
<input type="checkbox"/> Job run: "VMware backup job"	Failed	Fri, 15 Oct at 0:00
<input type="checkbox"/> Job run: "VMware backup job"	Completed	Fri, 15 Oct at 0:00
<input type="checkbox"/> Job run: "Backup copy job"	Completed	Thu, 14 Oct at 0:00
<input type="checkbox"/> Job run: "Nutanix AHV backup job"	Completed	Thu, 14 Oct at 0:00
<input type="checkbox"/> Job run: "VMware backup job"	Failed	Thu, 14 Oct at 0:00
<input type="checkbox"/> Job run: "VMware backup job"	Completed	Thu, 14 Oct at 0:00

## 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.

Past Activities

<input type="checkbox"/> Name	Status	Date	
<input type="checkbox"/> Job run: "Hyper-V backup job"	Failed	Thu, 21 Oct at 13:46	
<input checked="" type="checkbox"/> Job run: "VMware backup job"	Completed	Thu, 21 Oct at 0:00	
<input checked="" type="checkbox"/> Job run: "VMware backup job"	Completed	Wed, 20 Oct at 12:54	
<input checked="" type="checkbox"/> Job run: "Backup copy job"	Completed	Fri, 15 Oct at 0:00	
<input type="checkbox"/> Job run: "Nutanix AHV backup job"	Completed	Fri, 15 Oct at 0:00	
<input type="checkbox"/> Job run: "VMware backup job"	Failed	Fri, 15 Oct at 0:00	
<input type="checkbox"/> Job run: "VMware backup job"	Completed	Fri, 15 Oct at 0:00	
<input type="checkbox"/> Job run: "Backup copy job"	Completed	Thu, 14 Oct at 0:00	
<input type="checkbox"/> Job run: "Nutanix AHV backup job"	Completed	Thu, 14 Oct at 0:00	
<input type="checkbox"/> Job run: "VMware backup job"	Failed	Thu, 14 Oct at 0:00	

## Removing Activities

To remove activities from the list, tick the checkbox next to each desired activity and click **Remove** in the toolbar above. To remove all activities from the list at once, tick the **Select/Deselect all** checkbox at the top and click **Remove**. You can also remove a single activity by clicking the **Remove** icon that appears when you hover over a specific activity.



## 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** menu will appear.

Using this menu, 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.



## 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](#)
- [Viewing Object Info](#)

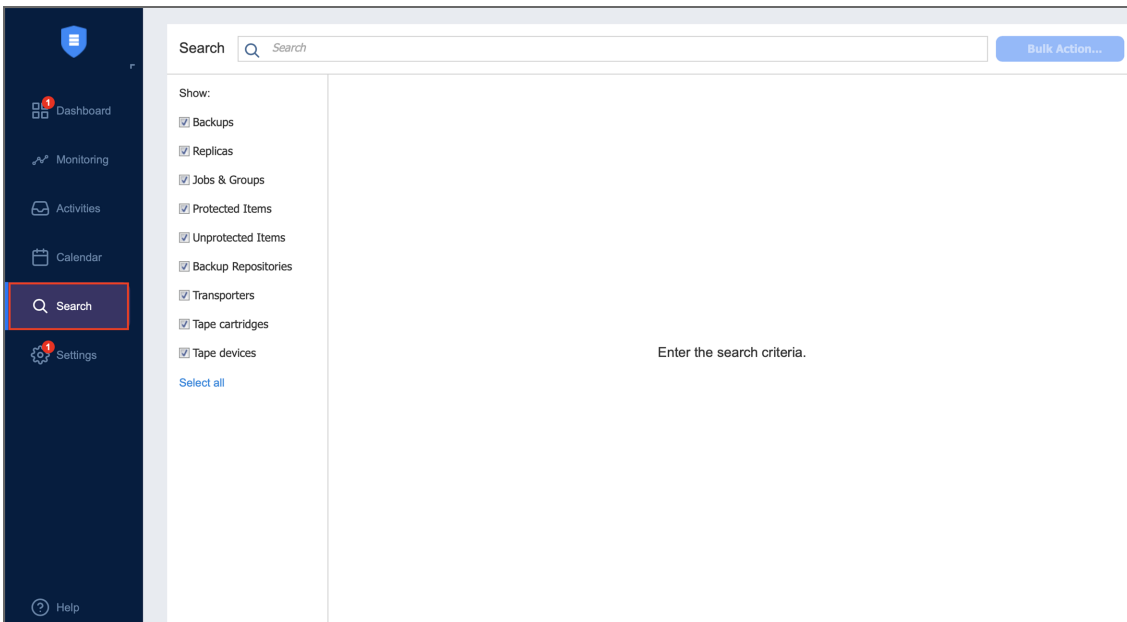
### Note

When the multi-tenant mode is enabled, Global Search will operate 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 138](#)
- [“Multi-Tenancy” on page 74](#)
- [“Multi-Tenant Mode” on page 1082](#)

## Opening Global Search

To open **Global Search**, click the **Search** icon in the main toolbar of the application.

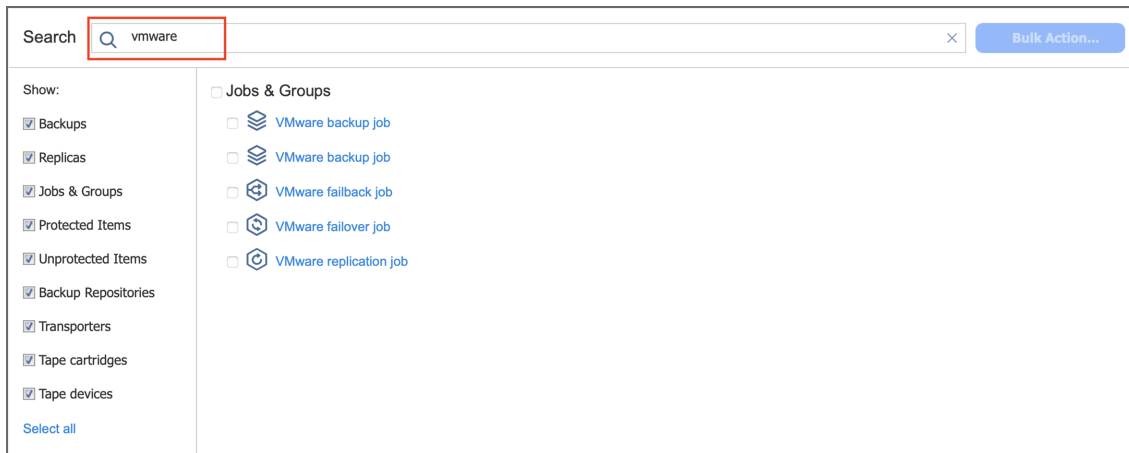


## Running Global Search

When the **Global Search** dashboard opens, you can enter your search string into the search box. The string you have entered will be immediately followed by a display of the search results in the form of a list.

To help you fine-tune your search, the following wildcards are applicable:

- "?" representing a single character.
- "\*" representing zero or more characters.



Please note the following:

- Search is case insensitive.
- Search results are grouped by categories.



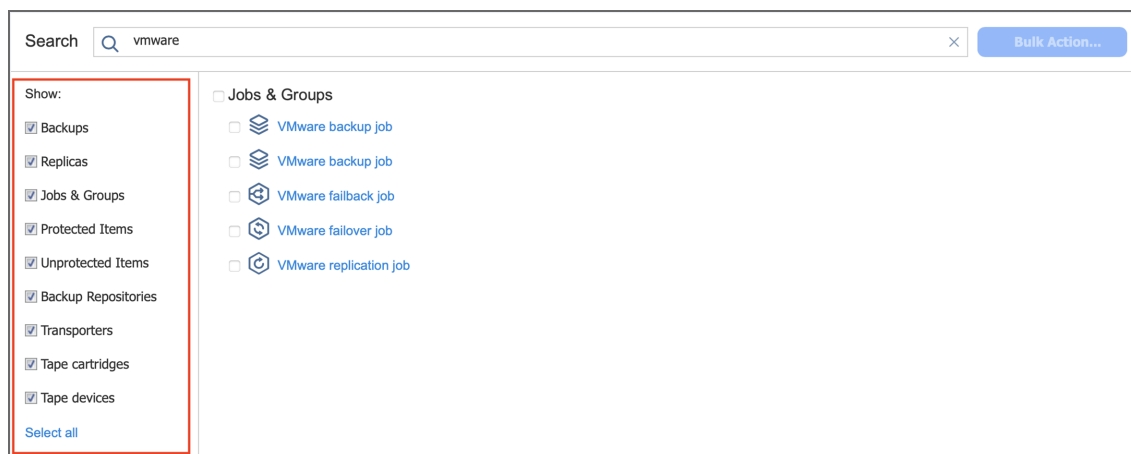
## Filtering Search Results

By default, your search results are unfiltered. This means that the search is applied to all categories of NAKIVO Backup & Replication objects.

To narrow your search results, deselect some categories in the categories list:

- Backups
- Replicas
- Jobs & Groups
- Protected Items
- Unprotected Items
- Backup Repositories
- Transporters

The filtered search results will be displayed immediately in the search results list.



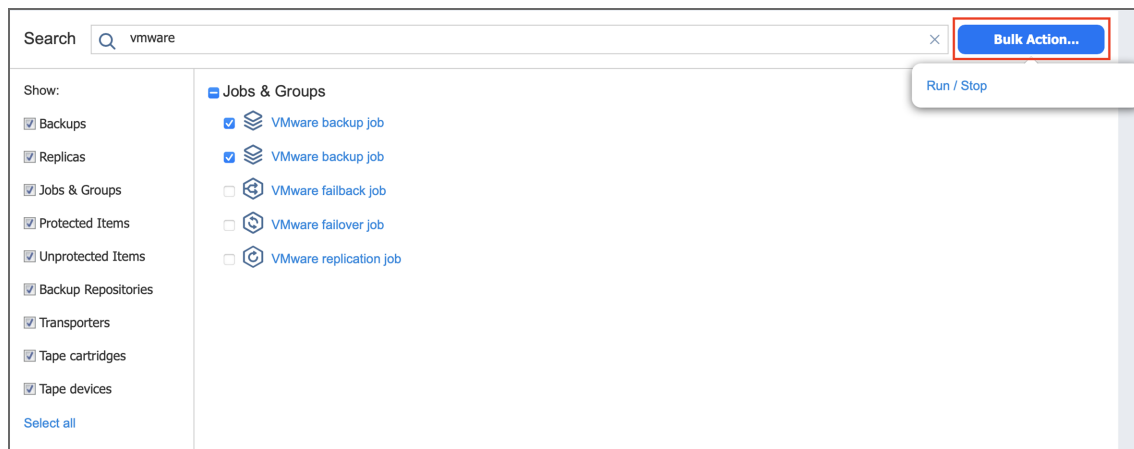
To get back to the default filtering settings, click **Select all** below the categories list.

## 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**.



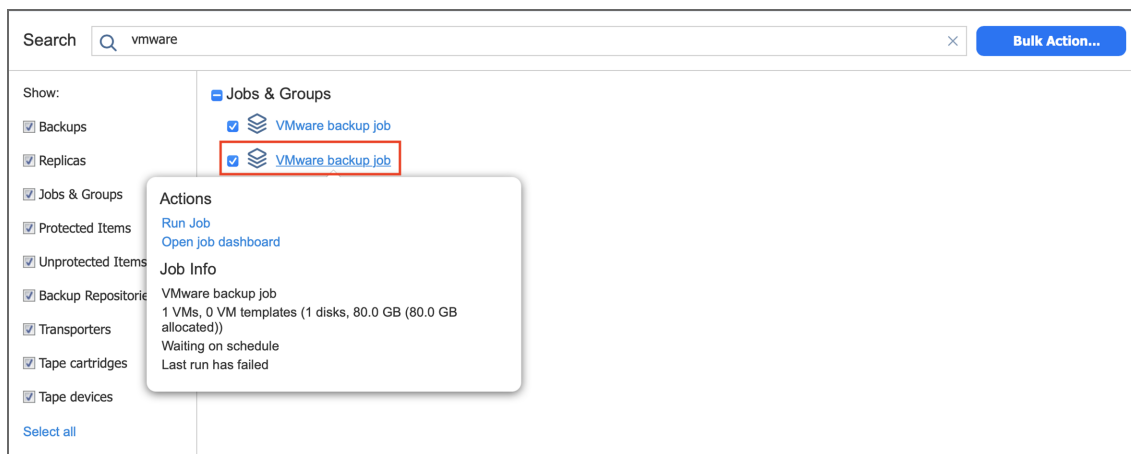
A dialog opens with the list of actions applicable to the selected items. To proceed with the necessary action, click the corresponding item in the list of actions.

**Note**

Bulk actions are not applicable to NAKIVO Backup & Replication dissimilar objects.

## Viewing Object Info

To view info on a specific object available in the search result, click the object.



A dialog opens displaying object info, along with the list of typical actions applicable to the object.

# Settings

This section covers the following topics:

- [“General” below](#)
- [“Inventory” on page 334](#)
- [“Transporters” on page 365](#)
- [“Backup Repositories” on page 388](#)
- [“Tape” on page 445](#)
- [“Expert Mode” on page 472](#)
- [“Virtual Appliance Configuration” on page 489](#)
- [“Multi-Tenant Mode Configuration” on page 491](#)
- [“Support Bundles” on page 497](#)
- [“Built-in Support Chat” on page 500](#)

## General

This section contains the following topics:

- [“Bandwidth Throttling” below](#)
- [“Branding” on page 287](#)
- [“Configuring Events” on page 289](#)
- [“Email Settings” on page 292](#)
- [“Database Options” on page 289](#)
- [“Licensing” on page 293](#)
- [“Notifications & Reports” on page 294](#)
- [“Self-Backup” on page 297](#)
- [“Software Update” on page 300](#)
- [“System Settings” on page 302](#)
- [“Users and Roles” on page 312](#)

## 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. Refer to [“Advanced Bandwidth Throttling” on page 49](#) for a description of bandwidth rules.

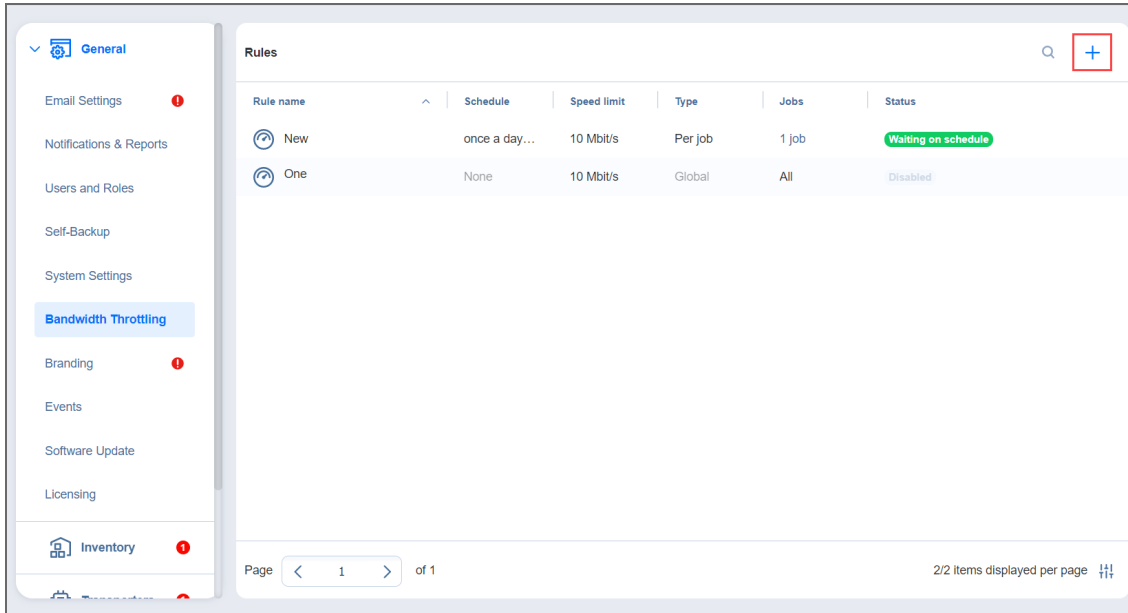
This topic contains the following instructions:

- [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:

- **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 the bandwidth rule to.

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: kbit/s, 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 **Save**.

**Create Bandwidth Rule**

Type

2 Settings

Name:

Throttle bandwidth to:     
 Equals 1.25 MB/s or 14 minutes to transfer 1GB of data

Rule schedule:

Starting at:  :

Ending at:  :

Days:

Every:  weeks

Time Zone:

## 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.
- **Remove:** 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 507](#)
- [“Creating Hyper-V Backup Jobs” on page 551](#)
- [“Creating Nutanix AHV Backup Jobs” on page 597](#)
- [“Creating Physical Machine Backup Jobs” on page 615](#)
- [“Creating Backup Copy Jobs” on page 668](#)
- [“Creating VMware Replication Jobs” on page 714](#)
- [“Creating Hyper-V Replication Jobs” on page 745](#)
- [“VMware VM Recovery” on page 888](#)
- [“Hyper-V VM Recovery” on page 912](#)
- [“Amazon EC2 Instance Recovery” on page 923](#)
- [“Nutanix AHV VM Recovery” on page 945](#)

## 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**
- **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.

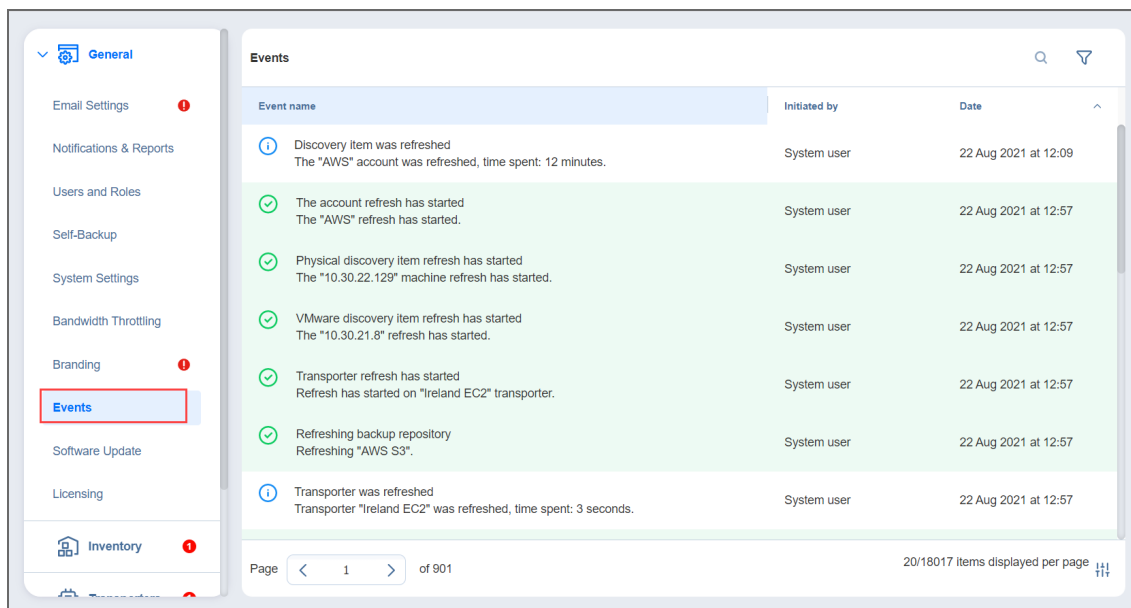
Image	Best format	Best resolution
Global logo	.png	32x40
Footer logo	.png	32x40
Favicon	.png	16x16

# 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 dropdown
  - **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

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.
- NAKIVO Backup & Replication may automatically recommend migrating to an external database if the current workload exceeds the limit recommended for the internal H2 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**.

## 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
```

Note that method (md5) may be different for some versions of PostgreSQL. Check the respective method for your version of PostgreSQL before applying the changes.
- 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.

# 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.

The screenshot displays the 'Email Settings' configuration page. On the left sidebar, the 'General' tab is active, and 'Email Settings' is highlighted with a red box. The main content area contains the following fields:

- SMTP server:
- SMTP username (optional):
- SMTP password (optional):
- SMTP port:
- Encryption:
- From:
- To:

At the bottom of the form, there is a 'Send Test Email' button, a 'Reset Settings' link, and 'Discard Changes' and 'Apply' buttons.

## 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. As an example, use instructions for Google accounts provided in the [Create & use App Passwords](#) article. When configuring email setting of the product, enter this password in the **SMTP password** box.

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.

### Note

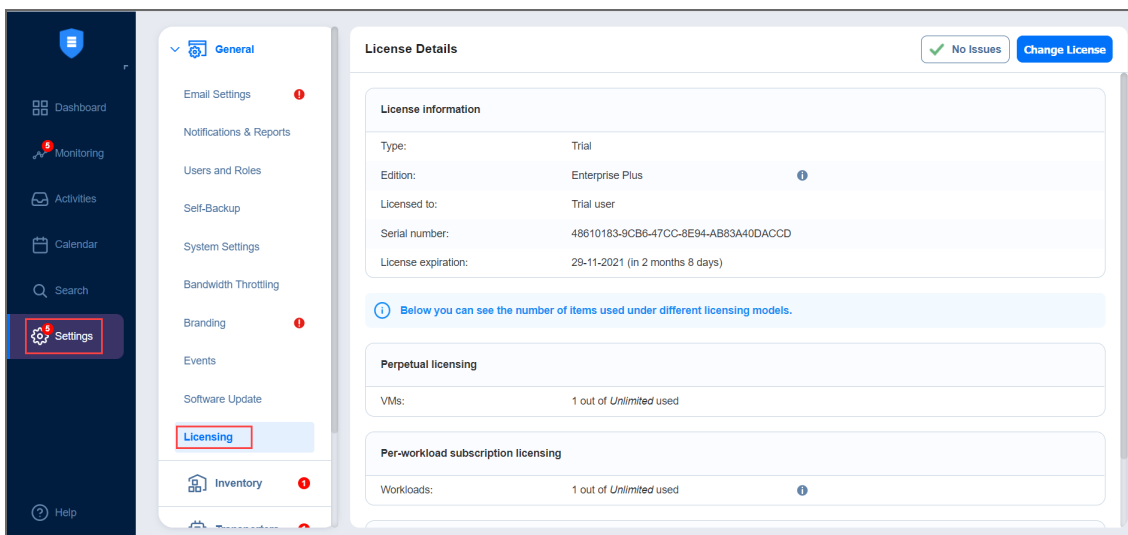
If you want to use a Gmail account to receive email notifications, turn on the **Less secure apps access** setting by navigating to **Manage your Google Account > Security** in your Google account.

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.

## Licensing

To check all 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 current license details.



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 license becomes expired

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 **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 current 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 current license type is **Free** and the **Trial** license has not yet been applied to the current 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 goes back to using your **Free** license.

## Notifications & Reports

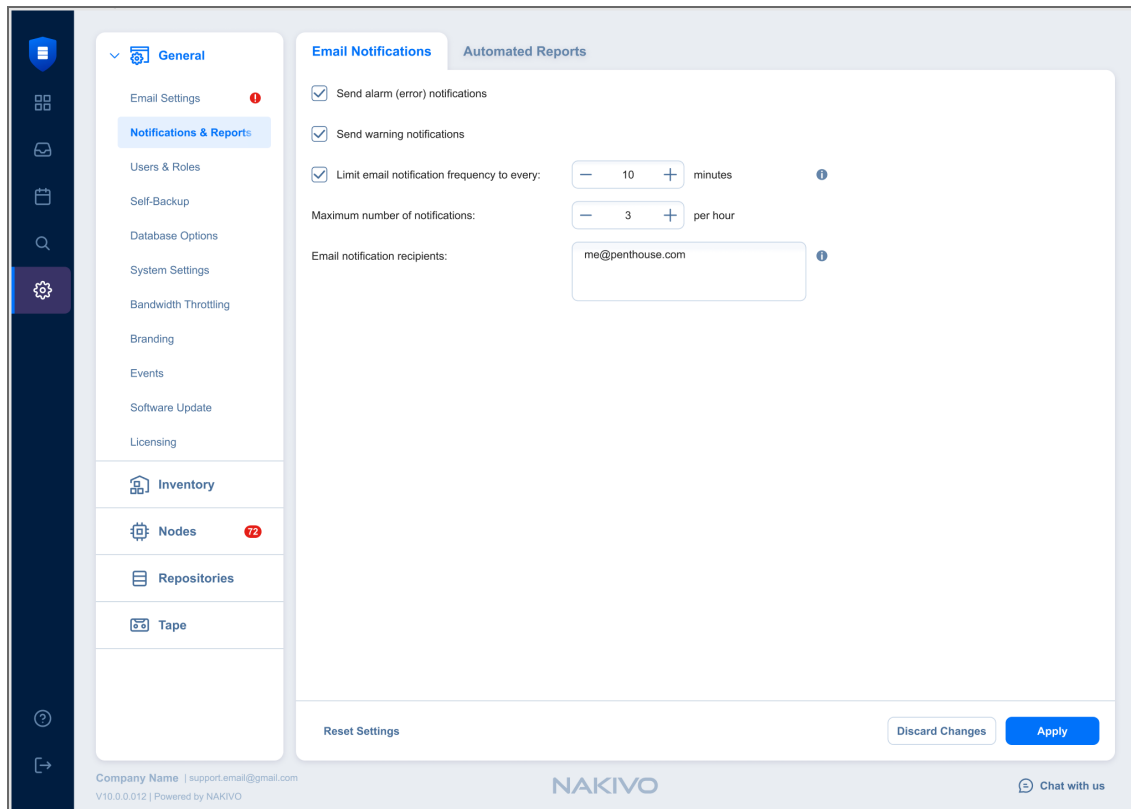
NAKIVO Backup & Replication can send notifications and reports over email.

- [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.



## 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 46](#).

### Note

Self-backup is not supported for the multi-tenant configuration.

To configure self-backup options, proceed as described in the following sections:

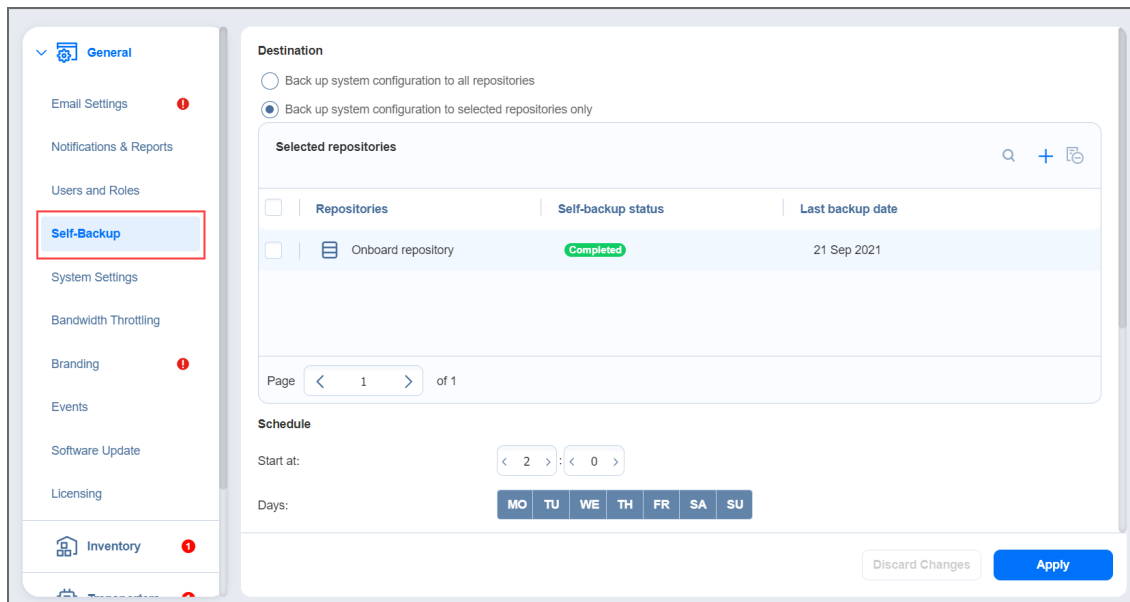
- [Accessing Self-Backup Options](#)
- [Setting Up Self-Backup Destination](#)
- [Self-Backup Schedule](#)
- [Self-Backup Options](#)
- [Self-Backup Encryption](#)
- [Recovering from Self-Backup](#)



## 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 will be backed up. If deselected, you can remove specific repositories from the list.

### Important

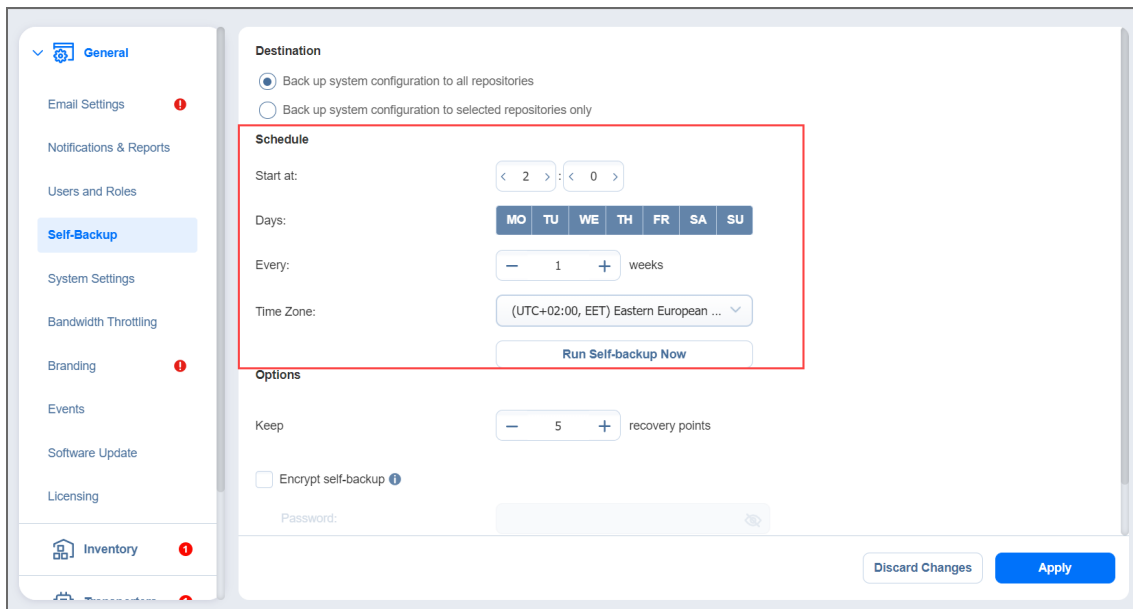
- Backing up your NAKIVO Backup & Replication system configuration to a [DD Boost storage unit Backup Repository](#) causes the DD Boost storage unit to be unmounted. Therefore, to avoid re-adding the DD Boost storage as an existing Backup Repository manually, exclude DD Boost storage unit repositories 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.

## Self-Backup Schedule

To configure the self-backup schedule, follow the steps below:

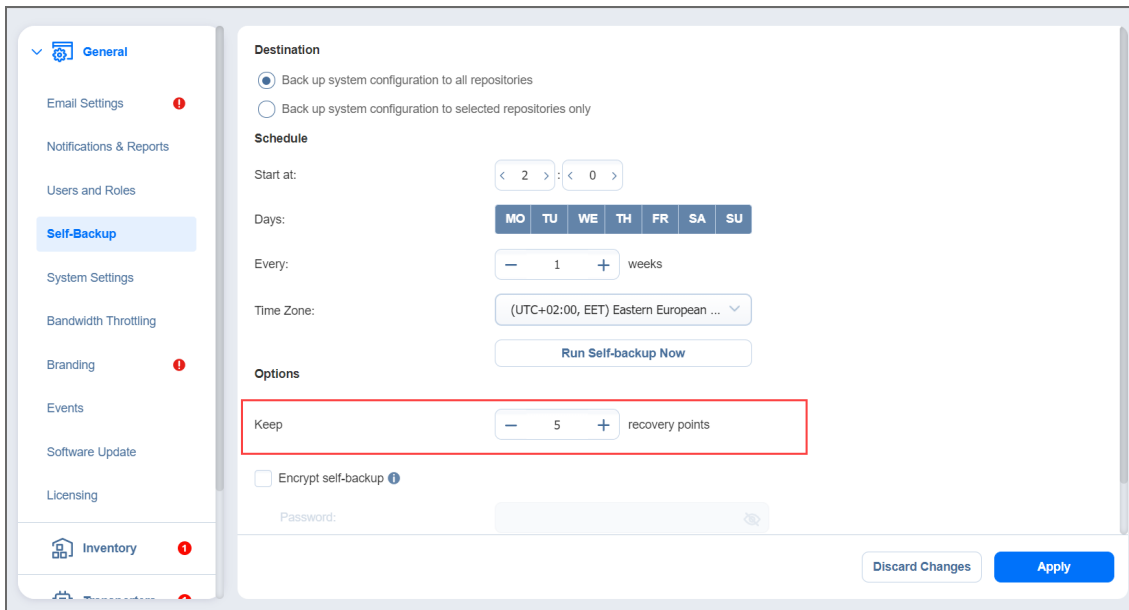
1. In the **Schedule section**, enter time to trigger starting the self-backup. You can choose a specific time zone from the list, enter the hours and minutes of the day, and select the necessary days of the week.
2. If you need to start the self-backup immediately, click **Run Self-backup Now**.
3. When ready with configuring the self-backup schedule, click **Apply**.



The screenshot displays the 'Self-Backup' configuration page. The left sidebar contains navigation options: General, Email Settings, Notifications & Reports, Users and Roles, Self-Backup (selected), System Settings, Bandwidth Throttling, Branding, Events, Software Update, and Licensing. The main content area is divided into sections: Destination, Schedule, and Options. The 'Schedule' section is highlighted with a red border and contains the following controls: 'Start at:' with a time picker set to 2:00; 'Days:' with a row of buttons for MO, TU, WE, TH, FR, SA, and SU; 'Every:' with a spinner set to 1 and the unit 'weeks'; and 'Time Zone:' with a dropdown menu showing '(UTC+02:00, EET) Eastern European ...'. Below the 'Schedule' section is the 'Options' section, which includes a 'Keep' spinner set to 5 'recovery points', an 'Encrypt self-backup' checkbox, and a 'Password:' field. At the bottom right, there are 'Discard Changes' and 'Apply' buttons.

## 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

Select **Encrypt self-backup** to encrypt your backup for additional security. Afterwards, enter the password in the **Password** and **Confirm Password** fields which are required to recover from the self-backup.

## 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. Select a recovery point and click **Restore**.
5. Wait while the system configuration is restored. When the self-backup recovery process is completed, a message announcing success appears.

### Note


If a selected recovery point was created from an encrypted self-backup, you will have to enter the password to it.

## Software Update

- [Download & Update Option](#)
- [Download Option](#)

If the full solution of NAKIVO Backup & Replication is installed on a Windows or Linux machine, you can download product updates and install them using NAKIVO Backup & Replication interface. To check if the upgrade 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.



## Note

If you are using a multi-tenant solution, only master-tenant users who have appropriate permissions will be able to see and manage this button.

## Product Auto-Updating Prerequisites

- At least 1GB of free space must be available on the machine on which the full solution is installed.
- Make sure your Maintenance & Support period is active. You can verify this on the product [Licensing](#) page.

## 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. Tick the **I have read the Release Notes** box.
3. Click **Download & update**.
4. Click **Update Now** to confirm stopping all current activities and start downloading the update. When the download is complete, the product updating process will begin.

The product will download the update to the Director first. When the Director is updated, the update will be downloaded to the Transporters that in turn will be updated simultaneously. If some Transporters are not updated, you can update them outside the product. Refer to the corresponding [articles](#) for details.

Updating the product will conduct self-backup and stop all current activities including running jobs, recovery jobs, repository maintenance, etc.

## Notes

- Only the following NAKIVO Backup & Replication Transporters can be auto-updated:
  - Windows including Hyper-V Transporters
  - Linux including Hyper-V Transporters
  - Amazon EC2 Transporters
  - VMware Transportes
- Only 20 Transporters can be updated simultaneously. All other Transporters will be sent to a queue and updated when their turn comes.

## Download Option

If you wish to postpone updating or schedule it for a certain period of time, take the following steps to download the update only:

1. Optionally, click **Release Notes** to see features and improvements implemented in the new product version.
2. Tick the **I have read the Release Notes** box.
3. Click **Download**.
4. After the download is finished, do either 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 on a schedule:
    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 the notification in the product menu with the **Update Reminder** dialog box. Do any of the following:
      - a. Click **Reschedule** if you want to reschedule the update and pick a different time.
      - b. Click **Cancel update** to cancel updating of 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) will be 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) will be deleted.

#### Note

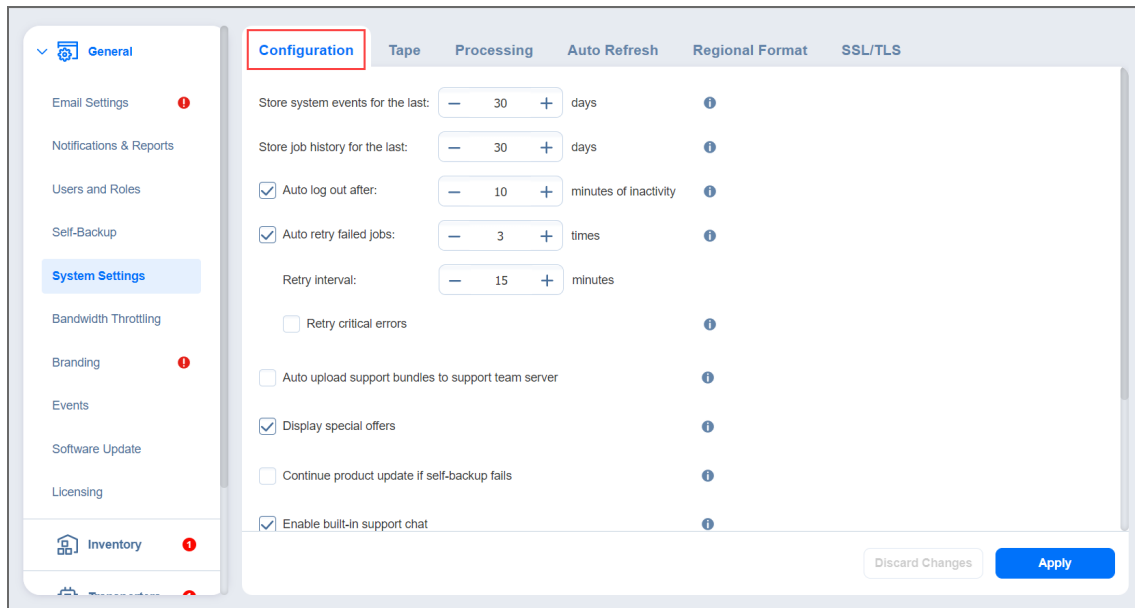
This option is not displayed for Master tenant in Multi-tenancy mode.

- **Auto log out after x minutes of inactivity:** When this option is selected, the current user will be 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 will be 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 the persistent errors which are unlikely to change without any additional intervention, i.e. hardware failure.
  - The term **non-critical error** refers to the non-persistent errors which are likely to change without any additional intervention, i.e. 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 APTARE storage resource management platform with NAKIVO Backup & Replication. For integration details, refer to [Aptare IT Analytics Integration](#).
- 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](#).




- In the **Tape** tab:
  - **Auto erase expired tapes:** When this option is selected, expired tapes are erased automatically.
 

**Important**

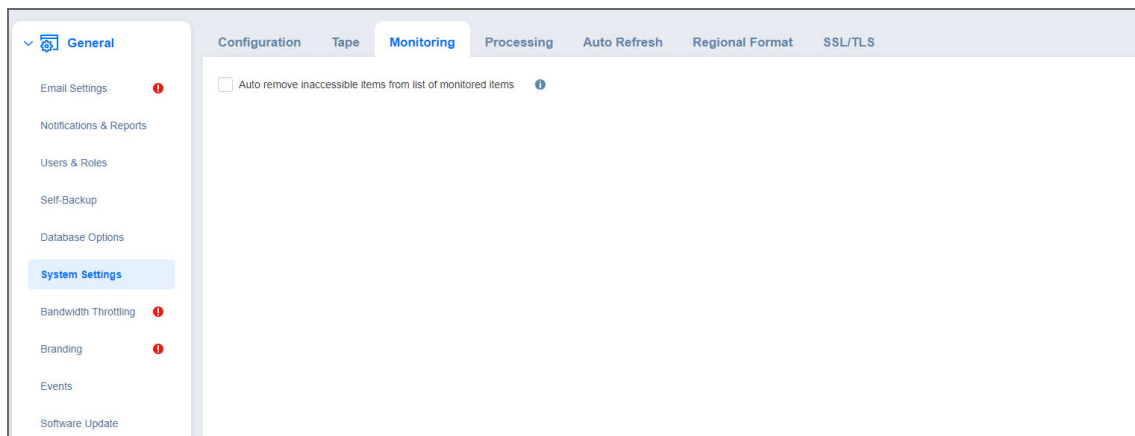
If this option is selected, the following prerequisites must be met for a cartridge to be erased:

    - All recovery points within the tape cartridge are 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.

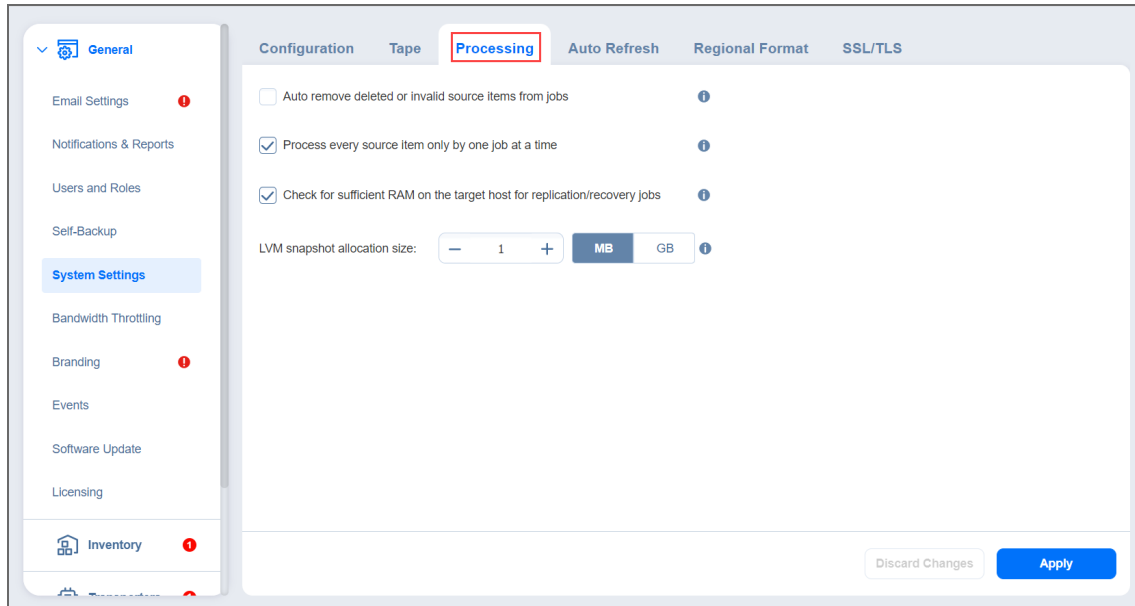


- In the **Processing** tab:

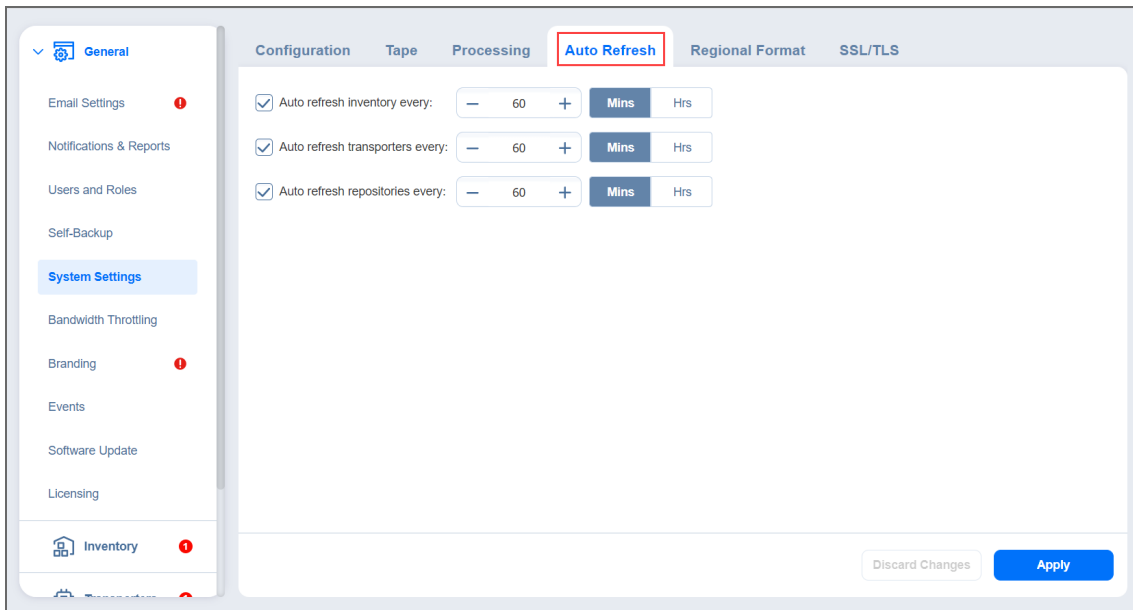
- **Auto remove deleted or invalid source items from jobs:** This option applies to a protected container (such as a VMware cluster or EC2 region). When this option is selected, if NAKIVO Backup & Replication discovers (during the [inventory refresh](#)) that a VM(s) and/or EC2 instance(s) is no longer available in the protected container, NAKIVO Backup & Replication automatically removes these VMs and EC2 instances 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 will not be affected after changing this setting. For physical servers, this option is always enabled.



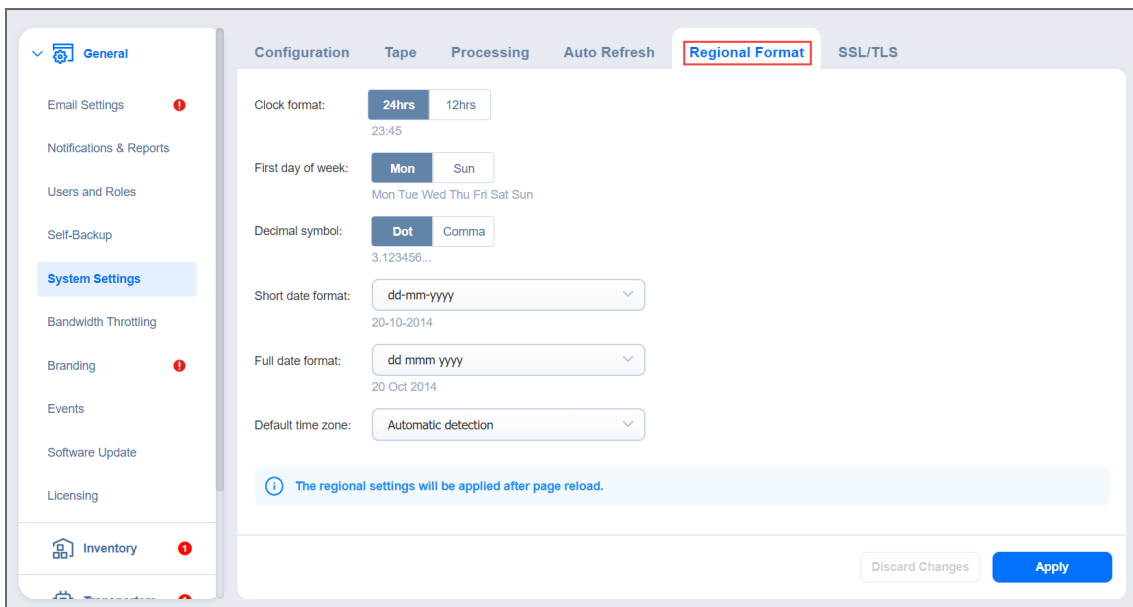
- **Check for sufficient RAM on the target host for replication/recovery jobs:** When this option is deselected, NAKIVO Backup & Replication does not check whether the amount of RAM on the target host is sufficient for replication and recovery jobs.
- **LVM snapshot allocation size:** This option allows 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.



- 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 inventories to be refreshed.



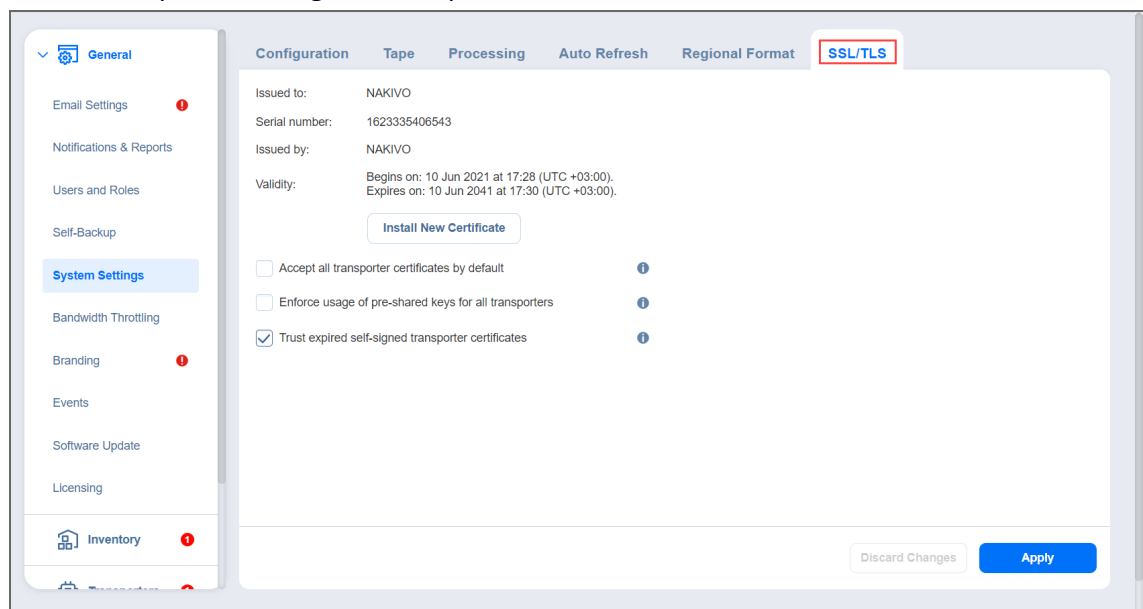
- In the **Regional Format** tab, set:
  - **Clock format**
  - **First day of week**
  - **Decimal symbol**
  - **Short date format**
  - **Full date format**
  - **Default time zone**



**Note**

If any time zone other than **(UTC+00:00, UTC) Coordinated Universal Time** is chosen, daylight savings times are honored.

- 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.



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 will result in failed jobs in the new instance after the migration. All jobs will have to be recreated, and full initial job run will be required.

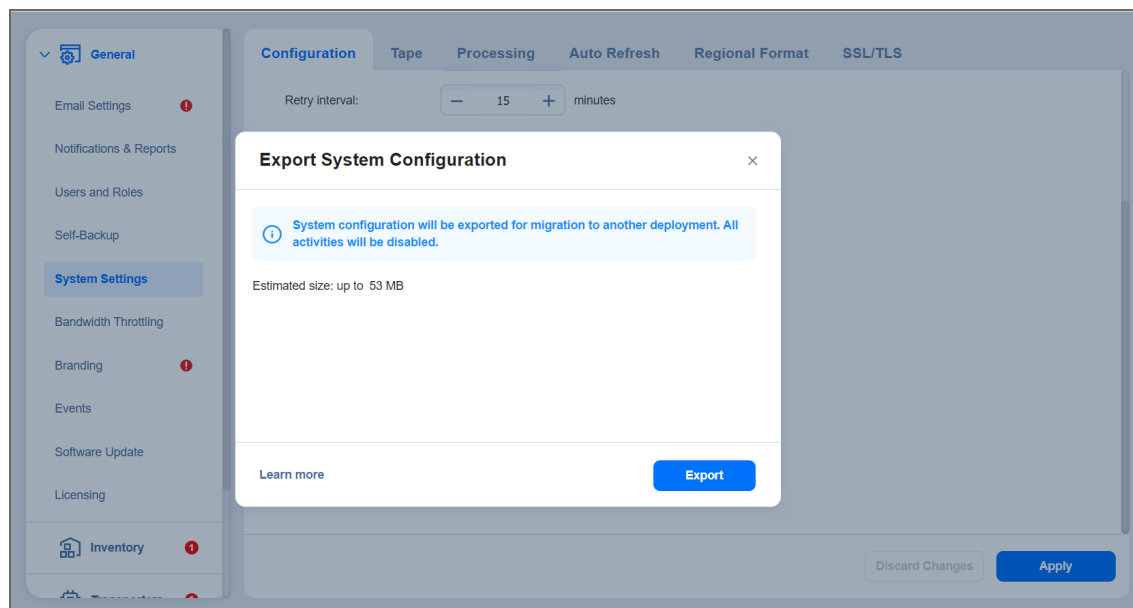
See the topics below for more information:

- [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**.



5. Click **Proceed** to confirm the operation.

### Note

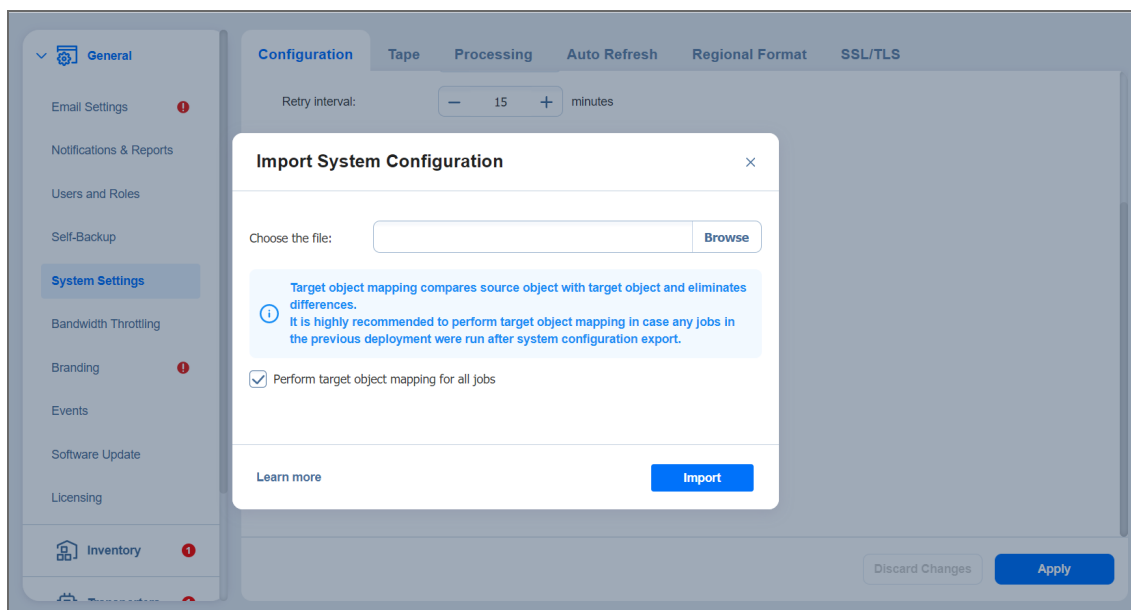
All activities in the old instance (such as jobs and recovery sessions) will be automatically stopped and all jobs will be disabled.

6. Wait until the export is completed and download the export bundle.
7. 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 window that appears, locate the system configuration bundle using the **Browse** button.



5. Click **Import**.
6. 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.
- 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.

7. 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, s 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 will be required in the new deployment.

# Users and Roles

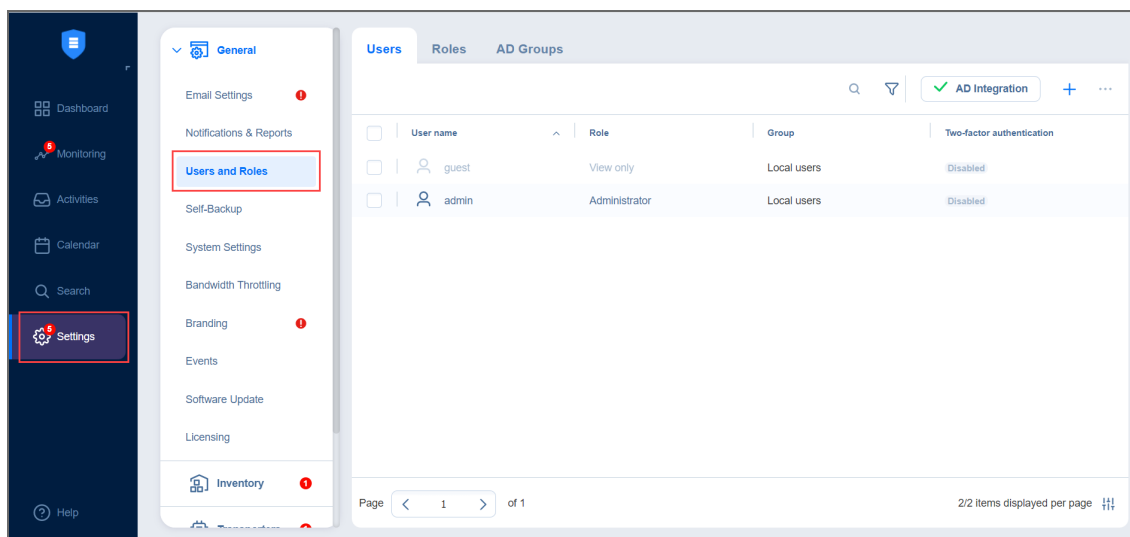
Accessing NAKIVO Backup & Replication is possible either with a user account created in the product or with an account added to the product from Active Directory. Each user in the product is assigned a role, which is a set of specific permissions.

- [Managing Users and Roles](#)
- [Navigating Users View](#)
- [Navigating Roles View](#)
- [Navigating AD 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 and Roles**.



## Navigating Users View

To see the list of all local 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 local users added to NAKIVO Backup & Replication.
- Sort the list by **Name**, **Role**, **2FA**, **Access level**, or **Group** 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 assigned to the user.

- 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.
  - Clicking **Filter** opens a new window that allows you to filter the list of local users according to **User name, Role, State, and Group**.
- Add a new local user by clicking “+” icon.
- Integrate Active Directory account by clicking **AD Integration**.
- 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 "**ellipsis**" icon.

### Note

When selecting all local users to apply a bulk action, NAKIVO Backup & Replication selects only those users that are displayed on the screen.

- Edit the role assigned to the local use by clicking on the name of the role in the respective column.

## 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.
  - 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 local user by clicking the “+” icon.
- 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 "**ellipsis**" icon.

Note



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.

## 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 the **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.
  - 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 “+” icon.
- Integrate Active Directory account by clicking **AD Integration**.
- 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 “ellipsis” icon.

### 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 317](#)
- [“Managing Local Users” on page 320](#)
- [“Managing User Roles” on page 327](#)
- [“Configuring Two-Factor Authentication” on the next page](#)

## 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

- 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**.

Proceed with configuring two-factor authentication:

1. Click **Continue** in the dialogue window that appears.
2. Click **Continue** in the **Verify your Email Address** popup that appears.
  - Optionally, click on the **change your email** link to enter the new email address for the user
  - Select **Continue** to proceed with 2FA configuration.
3. Enter the verification code that was sent to the specified email address, and click **Continue**.
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 by clicking **change email** link and add an alternative email by clicking **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.
4. 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.
5. 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**.

6. 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.

## Managing Active Directory Users

With NAKIVO Backup & Replication, you can configure Active Directory integration at any time. You can also freely add, edit, disable, delete AD users, or assign a role to them. For details, refer to the topics below:

- [“Adding Active Directory User” below](#)
- [“Assigning Role to Active Directory User” on the next page](#)
- [“Configuring Active Directory Integration” on the next page](#)
- [“Deleting Active Directory User” on page 319](#)
- [“Disabling Active Directory User” on page 319](#)
- [“Editing Active Directory User” on page 320](#)

### 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 "+" symbol. 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**. Active Directory users appear in the NAKIVO Backup & Replication list of users.

## 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 and Roles**.
2. The **Users and Roles** page opens in the **Users** view. Hover over the Active Directory user, and then click **ellipsis** symbol 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 and Roles**.
2. The **Users & Roles** page opens. Click the **Configure AD Integration** button.
3. The **Active Directory Configuration Wizard** opens on the **Settings** page. 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 **Preferred prioritized 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. **Refresh AD information every**: Specify a periodicity of refreshing Active Directory information.
- g. Click the **Test Integration** button to verify the successful integration with Active Directory.

- h. If Active Directory integration is tested successfully, a checkmark appears beside the **Test Integration** button. Then click **Next** to go to the next page of the wizard. If you fail to connect to the AD domain, refer to the [Knowledge Base](#) article for possible causes.
- i. On the **Users** page of the wizard, proceed with [adding an Active Directory user](#).

When the wizard closes, the **Users & Roles** page opens, displaying the newly-added Active Directory users in the list of users.

## Deleting Active Directory User

Follow the steps below to delete an Active Directory user:

1. Go to **Settings > General > Users and Roles**.
2. The **Users and Roles** page opens in the **Users** view. Hover over the Active Directory user you want to delete, and then click **ellipsis** icon in the rightmost cell of the row.
3. In the menu that opens, click **Delete**.
4. In the dialog box that opens, click **Delete** to confirm that you wish to delete the AD user.

The Active Directory user disappears from the list of users.

## Disabling Active Directory User

Follow the steps below to disable an Active Directory user:

1. Go to **Settings > General > Users and Roles**.
2. The **Users and Roles** page opens in the **Users** view. Hover over the Active Directory user you want to disable, and then click **ellipsis** icon 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.

## Editing Active Directory User

Follow the steps below to edit an Active Directory user:

1. Go to **Settings > General > Users and Roles**.
2. The **Users and 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 **ellipsis** icon in the rightmost column of the row.
  - c. Click **Edit**.
3. 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 Users” below](#)
- [“Assigning Role to Local User” on page 323](#)
- [“Deleting Local User” on page 323](#)
- [“Disabling Local User” on page 324](#)
- [“Editing Local User” on page 325](#)

The application has the following built-in local users:

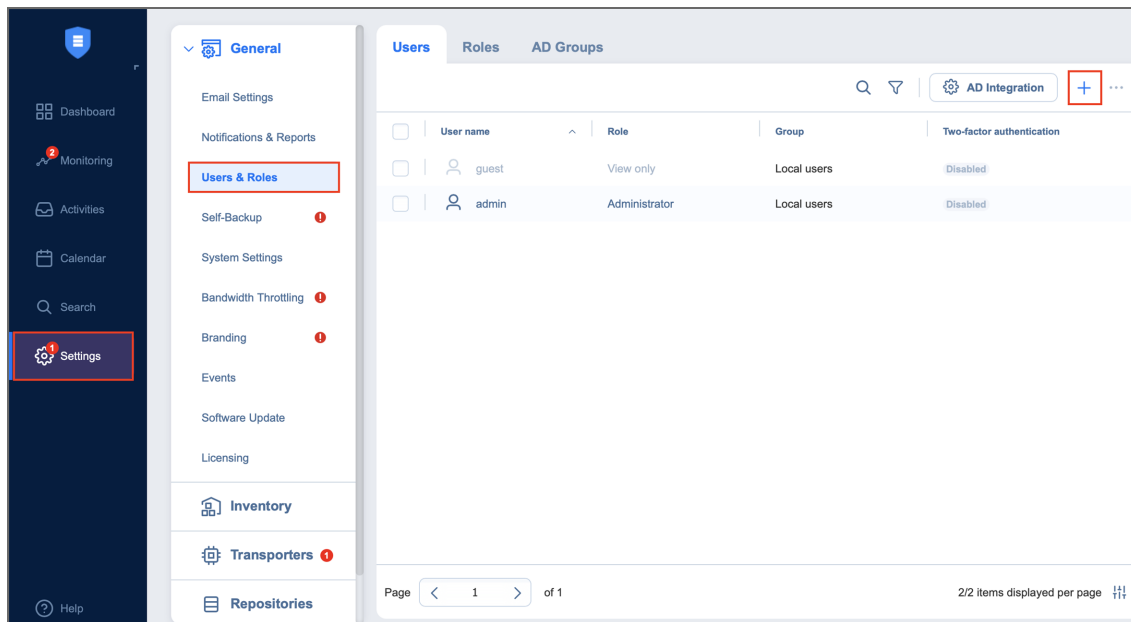
- **admin**: This user has the **Administrator** role assigned. You cannot delete it, disable it, or assign another role.
- **guest**: This user has the **View only** role assigned, with configurable file and object recovery permissions. By default, the account is disabled.

## Adding Local Users

Follow the steps below to add a local user:

1. Go to **Settings > General > Users and Roles**
2. The **Users and Roles** page opens on the **Users** tab.

3. Click the + symbol.



4. The **Add Local User** page opens. Proceed as follows:
- In the **Username** box, enter the user name.
  - In the **Name** box, enter the user's real name.
  - 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**.
  - In the **Repeat password** box, re-enter the user password.
  - In the **Email** box, enter the user's email address.
  - In the **Description** box, optionally enter a user description.



**Add Local User**

1 General

2 Role

Username:

Name:

Password:

Repeat password:

Generate password and send by email

Email:

Description:

Cancel **Next**

- 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 327](#) 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.

**Add Local User**

General

2 Role

Role: **Administrator**

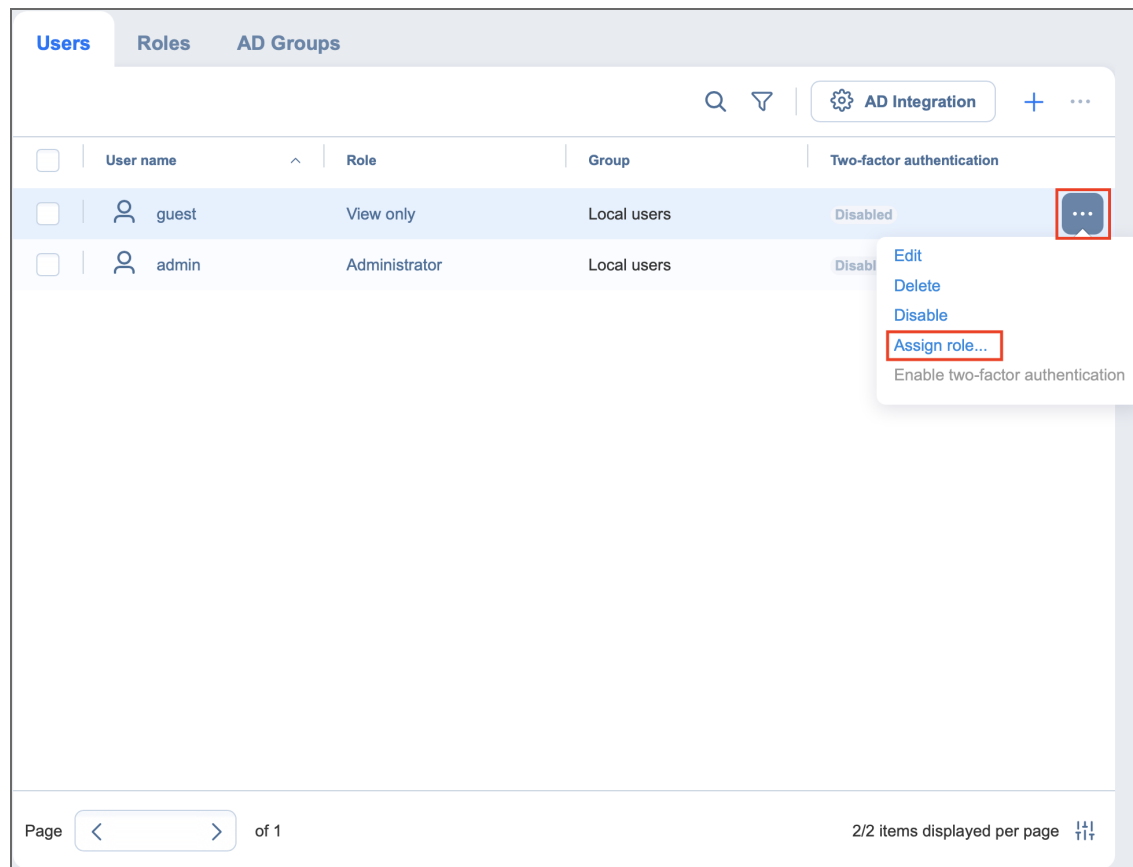
Calendar	Administrator	Full access
Activities	Backup operator	Full access
Global Search	Recovery operator	Full access
	View only	
	Run/stop job	Full access
Configuration	Create job	Full access
<b>Jobs</b>	Edit job	Full access
User profile	Edit job source	Full access
Help and Support	Edit job target	Full access
Aptare Report Generation	Edit job schedule	Full access
Monitoring	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 and Roles**.
2. The **Users and Roles** page opens in the **Users** view. Hover over the local user, and then click the **Ellipsis** icon in the rightmost cell of the row.
3. In the resulting menu, click **Assign role**.
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 and Roles**.
2. The **Users and 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 and Roles**.
2. The **Users and Roles** page opens in the **Users** view. Hover over the local user you wish to be disabled, and then click the **Ellipsis** icon 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.




## Editing Local User

Please follow the steps below to edit a local user:

1. Go to **Settings > General > Users and Roles**.
2. The **Users and 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** icon 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 **Name** box, edit the user name.
  - b. In the **Password** box, edit the user password.
  - c. If you edited the user password, re-enter the user password in the **Repeat password** box.
  - d. In the **Email** box, edit the user's email address.
  - e. Optionally, enable **Two-factor authentication**.

#### Note

This feature is disabled when no email address has been provided for the user.

- f. In the **Description** box, edit the user description.
- g. In the **Role** tab, edit the user's role.

- h. Click **Save** to save your modifications to the local user.

**Edit User:**

**General**

**Role**

Username: guest

Name: New Guest

Password: .....

Repeat password: .....

Email: guest@email.com

Two-factor authentication: Not configured [Configure](#) ⓘ

Description:

Cancel **Save**

## 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 332](#)
- [“Adding User Role” below](#)
- [“Editing User Role” on page 330](#)
- [“Cloning User Role” on the next page](#)
- [“Deleting User Role” on page 329](#)

## Adding User Role

Follow the steps below to add a user role:

1. Go to **Settings > General > Users and Roles**.
2. On the **Users and Roles** page, switch to the **Roles** tab.
3. Click + symbol and then select **Add Role**.
4. The **Add Role** page opens. Proceed as follows:
  - a. In the **Role name** box, enter the 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.

**Add Role**

1 General

2 Permission

Role name:

Description:

Cancel **Next**

- d. Click **Next** to proceed to **Permission** tab.
- e. A list of permissions opens. Specify necessary permissions for the user role.

**Add Role**

General

2 **Permission**

Calendar >

Activities >

Global Search >

Configuration >

**Jobs** >

User profile >

Help and Support >

Aptare Report Generation >

Monitoring >

No access  View only  Run only  Full access  Custom

No access  View only  Run only  Full access  Custom

No access  View only  Run only  Full access  Custom

View job

Run/stop job

Create job

Edit job

Edit job source

Edit job target

Edit job schedule

Previous Cancel **Finish**

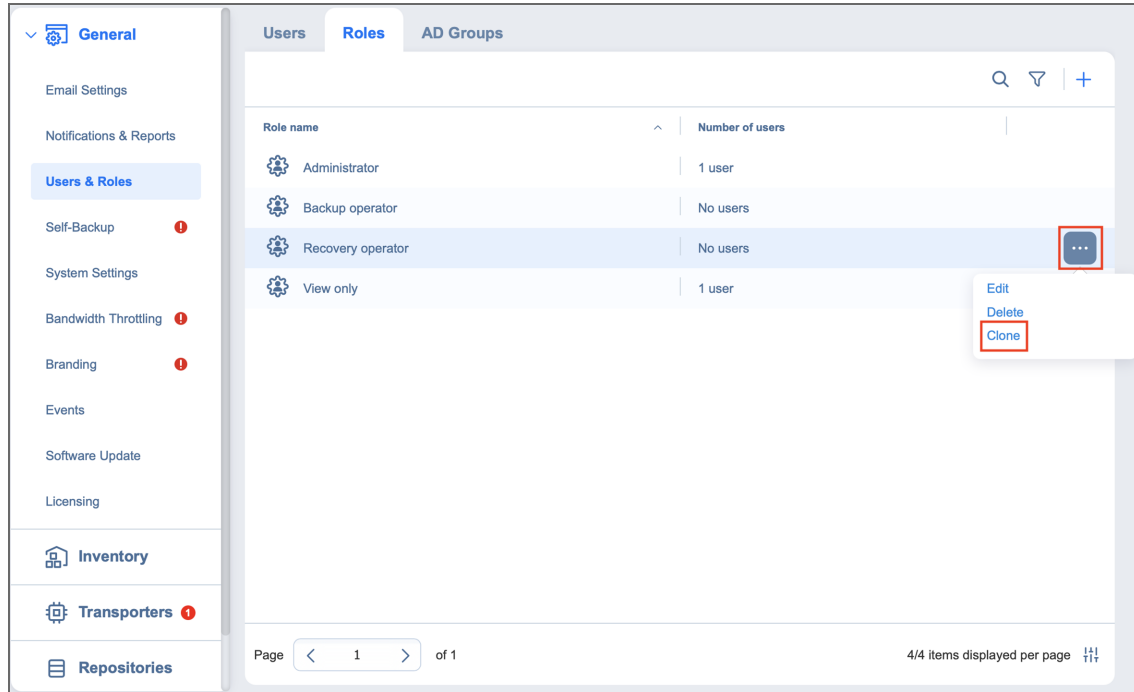
- f. Click **Finish** in the lower right corner of the page.

The user role appears in the list of roles.

## Cloning User Role

Follow the steps below to clone a user role:

1. Go to **Settings > General > Users and Roles**.
2. On the **Users and Roles** page, switch to the **Roles** tab.
3. Hover over the user role, click the **Ellipsis** icon 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 and Roles**.
2. On the **Users and Roles** page, switch to the **Roles** tab.
3. Hover over the user role, click the **Ellipsis** icon 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.



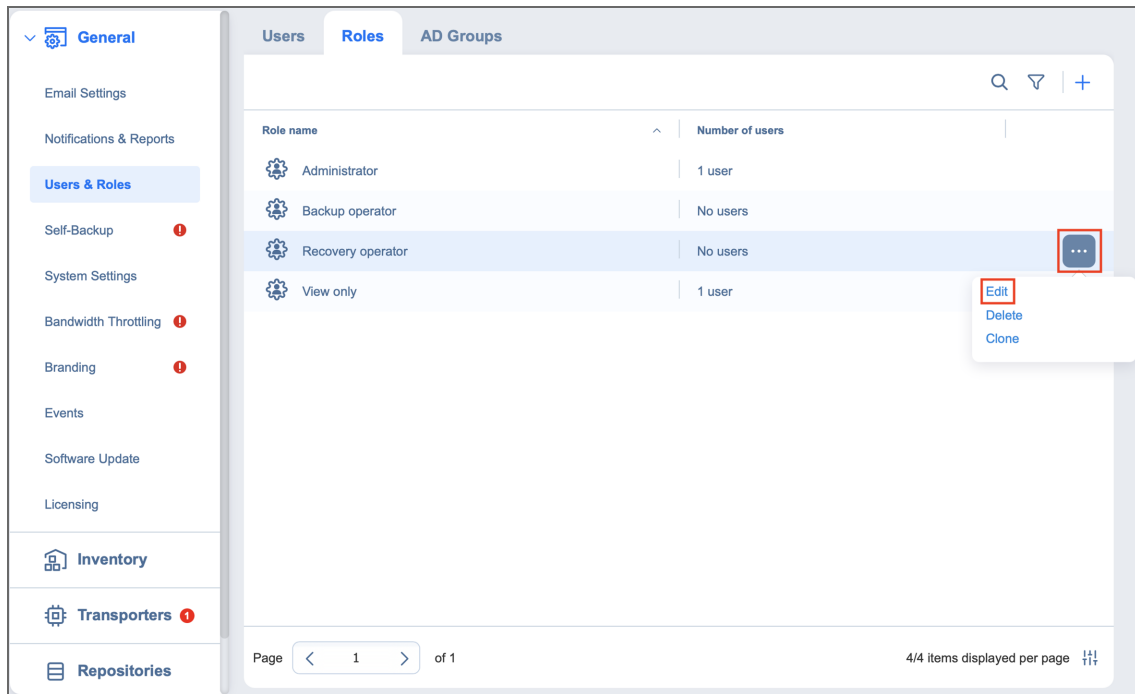
The user role disappears from the list of roles.

## Editing User Role

Follow the steps below to edit a user role:

1. Go to **Settings > General > Users and Roles**.
2. On the **Users and 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** icon 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.

The screenshot shows the 'Add Role' configuration page. The 'Permission' tab is active, and the 'Jobs' category is selected. The 'View only' permission level is chosen for the main category. Specific permissions for 'Jobs' are set to 'Full access'.

## 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 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

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.

# 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. Refer to the following sections to learn more:

- [“Adding VMware vCenters and ESXi hosts” below](#)
- [“Adding VMware Cloud Director Server” on page 337](#)
- [“Adding Microsoft Hyper-V Servers” on page 338](#)
- [“Adding Amazon EC2 Accounts” on page 339](#)
- [“Adding File Shares to Inventory” on page 353](#)
- [“Adding Wasabi Accounts” on page 340](#)
- [“Adding Nutanix AHV Clusters” on page 343](#)
- [“Adding Physical Machines to Inventory” on page 344](#)
- [“Adding Microsoft 365 Accounts” on page 345](#)
- [“Adding HPE 3PAR Storage Devices to Inventory” on page 355](#)
- [“Adding HPE Nimble Storage Devices to Inventory” on page 356](#)
- [“Adding Oracle Database to Inventory” on page 356](#)
- [“Managing Credentials” on page 361](#)
- [“Managing Inventory” on page 358](#)

## Adding VMware vCenters and ESXi hosts

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 **Add New**.
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.

Add Inventory Item

1. Platform      2. Type      3. Options

VMware vCenter or ESXi host

Microsoft Hyper-V host or cluster

Nutanix AHV cluster

VMware Cloud Director server

Cancel    Next

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.

**Note**

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.

- 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 use the **Direct Connect** functionality. To do this:
  - i. Select **Use Direct Connect**.
  - ii. In the **Assigned transporter** drop-down list, choose the Transporter with the enabled Direct Connect functionality.

**Notes**

- This option is available only if there is a Transporter with Direct Connect enabled. For details, refer to [“Adding Installed Transporters” on page 365](#).
- 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 373](#) for more details about deploying additional Transporters that support VMware vSphere.

# Adding VMware Cloud Director Server

To add a VMware Cloud Director server to Inventory in NAKIVO Backup & Replication, follow the steps below:

1. Navigate to the main menu of NAKIVO Backup & Replication and click **Settings**.
2. Go to the **Inventory** page and click **Add New**.
3. On the **Platform** page, select Virtual and then click **Next** to proceed.
4. On the **Type** page, select **VMware Cloud Director server** and then click **Next** to proceed.
5. Complete the wizard as described in the sections below.
  - [Adding VMware Cloud Director Server](#)
  - [Providing vCenter Credentials](#)

## Adding VMware Cloud Director Server

On the **VMware Cloud Director Details** page of the **Add New VMware Cloud Director Server** wizard, enter the following information:

1. **Display name:** Specify a name for the VMware Cloud Director server. This name will be displayed in the inventory. If you leave the field empty, the contents of the hostname or IP are displayed.
2. **Hostname or IP:** Specify the hostname or IP address of the VMware Cloud Director server that you want to add to the inventory.
3. **Username:** Provide a username for the VMware Cloud Director server.
4. **Password:** Provide a password to the VMware Cloud Director server.
5. **Port:** Specify the port for the VMware Cloud Director server.
6. Click **Next** to proceed to the next page.

The screenshot shows a web-based wizard titled "Add Inventory Item" with four steps: 1. Platform, 2. Type, 3. VMware Cloud Director Details (current step), and 4. vCenter Credentials. The form fields are as follows:

Display name:	<input type="text" value="VMware Cloud Director"/>
Hostname or IP:	<input type="text" value="10.30.22.444"/>
Username:	<input type="text" value="admin"/>
Password:	<input type="password" value="*****"/>
Web services port:	<input type="text" value="443"/>

At the bottom right, there are two buttons: "Cancel" and "Next". The "Next" button is highlighted in blue.



## Providing vCenter Credentials

On the **vCenter Credentials** page of the wizard, provide credentials for the selected vCenters. Proceed as follows:

1. On the left side of the screen, select the checkboxes of the vCenters to be added to the inventory.
2. In the **Username** field, provide the username.
3. In the **Password** field, enter the password.
4. Click **Finish**.

## Adding Microsoft Hyper-V Servers

To add Microsoft Hyper-V servers to the product, follow the steps below:

1. In the main menu, click **Settings**.
2. Go to the **Inventory** page and click **Add New**.
3. On the **Platform** page, select **Virtual** and then click **Next** to proceed.
4. On the **Type** page, select **Microsoft Hyper-V host or cluster** and then click **Next** to proceed.
5. The **Options** page opens. Fill out the fields as follows:
  - **Display name:** Enter a name for the host. This name will be displayed in the inventory.
  - **Type:** Select whether you want to add a standalone Microsoft Hyper-V server or a cluster. For Hyper-V Failover Clusters, it's enough to add any cluster member by an IP address or domain name. The list of cluster members will be retrieved automatically.
  - **Hostname or IP:** Specify the hostname or IP address of the Microsoft Hyper-V server.
  - **Username and Password:** Specify the credentials of the Microsoft Hyper-V server that you want to add to the inventory. The credentials you specify should have full administrative privileges to the Hyper-V server
  - **WS MAN port:** Specify the port number for WS MAN.

## 6. Click **Finish**.

The screenshot shows a dialog box titled "Add Inventory Item" with three tabs: "1. Platform", "2. Type", and "3. Options". The "1. Platform" tab is active. The form contains the following fields:

- Display name: Hyper-V
- Type: Standalone Hyper-V host
- Hostname or IP: 10.22.41.11
- Username: admin
- Password: [masked]
- WS-MAN port: 5986

At the bottom right, there are "Cancel" and "Finish" buttons.

After the process has completed successfully, you can exit **Settings** and create jobs with the newly discovered VMs.

### Notes

- When adding a new Microsoft Hyper-V server to the inventory, NAKIVO Backup & Replication automatically installs the [Transporter](#) service onto the server. This service is used to read data from source VMs during backup and replication.
- Multiple Directors are not supported. If you add the same Microsoft Hyper-V host to the inventory of an additional NAKIVO Backup & Replication instance, Transporter installation fails.

## 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” on the next page](#)

### 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 **Add New**.
3. On the **Platform** page of the wizard, select **Cloud** and click **Next** to proceed.
4. 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.
5. Click **Finish** when you're done.

The screenshot shows a wizard window titled "Add Inventory Item" with two steps: "1. Platform" and "2. Options". The "2. Options" step is active. The form contains the following fields:

- Display name:
- Type:  (dropdown menu)
- Region(s):  (dropdown menu)
- Access key ID:
- Secret access key:

At the bottom right, there are "Cancel" and "Finish" buttons.

## Adding Wasabi Accounts

Add a Wasabi account to NAKIVO Backup & Replication as described in the sections below.

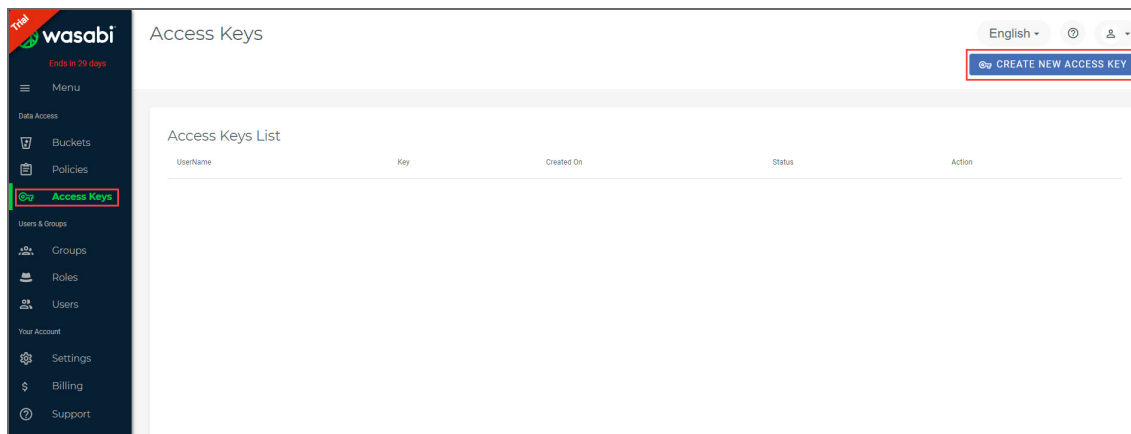
- [Creating Wasabi Access Key ID and Secret Access Key](#)
- [Adding a Wasabi Account to Inventory](#)

## 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/](https://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** 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.

### Note

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?](#)

Create New Access Key

Note: Download your key file now, which contains your new access key and secret access key. If you do not download the key file now, you will not be able to retrieve your secret access key again.

Root User  Sub-User

Assign to a user:  
SearchForUser

Liana

Liana REMOVE

CANCEL CREATE

6. Click **Download CSV** and save the file with generated keys in a secure location. Keep the Access key confidential in order to protect your account.

## Adding a Wasabi Account to Inventory

To add an Wasabi account to NAKIVO Backup & Replication, follow the steps below:

1. Click **Settings**.
2. Go to the **Inventory** page and click **Add New**.
3. On the **Platform** page of the wizard, select **Cloud** and click **Next** to proceed.
4. 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 account from the **Type** drop-down list.
  - c. Select the Wasabi 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.
5. Click **Finish**.

Add Inventory Item

1. Platform 2. Options

Display name: Wasabi

Type: Wasabi account

Region(s): Wasabi US West 1 (Oregon)

Access key ID: key

Secret access key: .....

Cancel Finish

# Adding Nutanix AHV Clusters

To add a Nutanix AHV cluster to NAKIVO Backup & Replication, follow the steps below:

1. Make sure the AOS on your AHV cluster is a [supported](#) version.
2. Click **Settings** in the left pane of NAKIVO Backup & Replication.
3. Go to the **Inventory** page and click **Add New**.
4. On the **Platform** page, select **Virtual** and then click **Next** to proceed.
5. On the **Type** page, select **Nutanix AHV** cluster and then click **Next** to proceed.
6. The **Options** page opens. Set the following parameter page:
  - **Display name:** Specify a name for the cluster. This name will be displayed in the inventory. If you leave the name field empty, the system will fill it in automatically with the contents of the Cluster hostname or IP field during the save process.
  - **Cluster hostname or IP:** Specify the hostname or IP address of the Nutanix AHV cluster that you want to add to the inventory.
  - **Username and Password:** Specify the credentials of the Nutanix AHV cluster that you want to add to the inventory.
  - **Cluster port:** Specify the port of the cluster (9440 by default).

Add Inventory Item

1. Platform      2. Type      3. Options

Display name:

Cluster hostname or IP:

Username:

Password:

Cluster port:

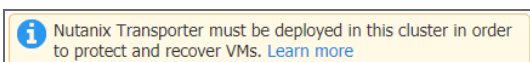
**i** Nutanix Transporter must be deployed in this cluster in order to protect and recover VMs. [Learn more](#)

Cancel      Finish

7. Click **Finish**. The added cluster is displayed on the **Inventory** tab.

## Important

The Nutanix Transporter must be deployed in the cluster to enable backup/recovery operations. You will still be able to add the cluster to the inventory, but Nutanix operations will be disabled in the menu until the Transporter is deployed. The system will notify you with a warning message if the Transporter is missing.



Refer to [“Deploying Transporter as Nutanix AHV Appliance” on page 376](#) for more details about deploying Nutanix AHV transporters.

## Adding Physical Machines to Inventory

To add a Windows or Linux physical machine to NAKIVO Backup & Replication, follow the steps below:

1. Make sure that a physical machine is a supported version and all prerequisites are met before proceeding. For more details, refer to [“Supported Platforms” on page 97](#).
2. Click **Settings** in the left pane of the product.
3. Go to the **Inventory** page and click **Add New**.
4. On the **Platform** page of the wizard, select **Physical** and click **Next** to proceed.
5. On the **Options** page of the wizard, fill in the following fields.
  - For Windows physical machines:
    - **Display name:** Specify a name for the physical machine. This name will be displayed in the inventory.
    - **Type:** Display the type of machine that you add to the inventory.
    - **Hostname(s) or IP(s):** Specify the hostname or IP address of the physical machine that you want to add to the inventory. To add multiple physical machines at once, use commas to separate hostnames or IP addresses, or use a dash for an IP address range.
    - **Username:** Provide a username for the physical machine.
    - **Password:** Provide a password for the physical machine.

The screenshot shows a web-based wizard titled "Add Inventory Item". It has two steps: "1. Platform" and "2. Options". The "1. Platform" step is active. The form contains the following fields:

- Display name: Text input with "Physical" entered.
- Type: Dropdown menu with "Microsoft Windows" selected.
- Hostname(s) or IP(s): Text input with "192.11.11.11" entered.
- Username: Dropdown menu with "admin" selected.
- Password: Password input field with masked characters.

Below the password field is a link labeled "Manage credentials". A yellow information box with an 'i' icon contains the text: "A transporter will be installed on each machine, unless there is a transporter already installed." At the bottom right of the form are two buttons: "Cancel" and "Finish".

- For Linux physical machines:
  - **Display name:** Specify a name for the physical machine. This name will be displayed in the inventory.

- **Type:** Display the type of machine that you add to the inventory.
- **Hostname(s) or IP(s):** Specify the hostname or IP address of the physical machine that you want to add to the inventory. To add multiple physical machines at once, use commas to separate hostnames or IP addresses, or use a dash for an IP address range.
- **SSH port:** Specify the SSH port number to provide access to the physical machine. The default SSH port number is 22.
- **Credentials Type:** Select the type of credentials used to access the physical machine.
  - **Password**
    - **Username:** Provide a username for the physical machine.
    - **Password:** Provide a password for the physical machine.
  - **Private Key**
    - **Credentials:** Select the saved private key-based credentials. Refer to [“Managing Credentials” on page 361](#) for details.

The screenshot shows a web-based form titled "Add Inventory Item" with two tabs: "1. Platform" (active) and "2. Options". The form contains the following fields and values:

- Display name: Physical
- Type: Linux
- Hostname(s) or IP(s): 192.11.11.11
- Credentials type: Password
- Username type: admin
- Username: admin
- Password: [masked]
- SSH port: 22

Below the fields is a yellow information box with a blue 'i' icon: "A transporter will be installed on each machine, unless there is a transporter already installed." At the bottom right of the form are "Cancel" and "Finish" buttons.

### Note

You will not be able to change the type of key credentials through the [Manage Credentials](#) option.

6. Click **Finish**. The successfully added physical machine is displayed on the inventory list.

## Adding Microsoft 365 Accounts

Before you start [backing up](#) items from your organization’s Microsoft 365 account, you need to add the Microsoft 365 account to Inventory in NAKIVO Backup & Replication.

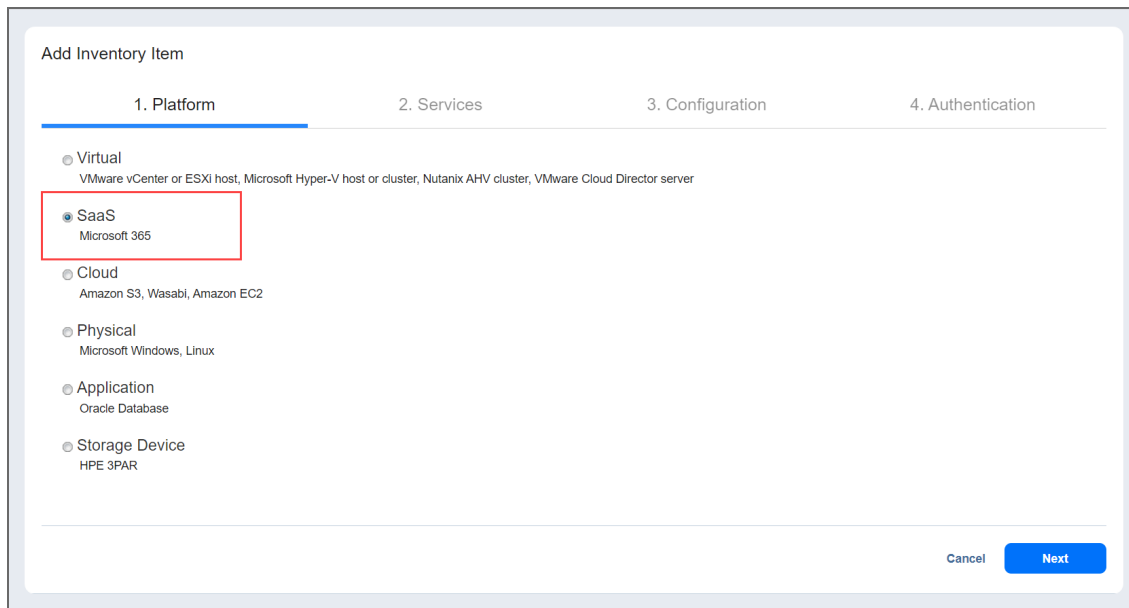
- [Adding Microsoft 365 Account to Inventory](#)
- [Obtaining Microsoft 365 Credentials](#)



## Adding Microsoft 365 Account to Inventory

To add a Microsoft 365 account to Inventory, do the following:

1. Check if you meet the [requirements for Backup for Microsoft 365](#).
2. Click **Settings** in the left pane.
3. Go to the **Inventory** tab and click **Add New**.
4. On the **Platform** page of the wizard, select **SaaS**. Click **Next** to proceed.



The screenshot shows the 'Add Inventory Item' wizard with four steps: 1. Platform, 2. Services, 3. Configuration, and 4. Authentication. Step 1 is active. Under 'Platform', several options are listed with radio buttons: Virtual (VMware vCenter or ESXI host, Microsoft Hyper-V host or cluster, Nutanix AHV cluster, VMware Cloud Director server), SaaS (Microsoft 365), Cloud (Amazon S3, Wasabi, Amazon EC2), Physical (Microsoft Windows, Linux), Application (Oracle Database), and Storage Device (HPE 3PAR). The 'SaaS' option is selected and highlighted with a red rectangular box. At the bottom right, there are 'Cancel' and 'Next' buttons.

5. On the **Services** page, provide the following information:
  - **Display name:** Enter a name for the account.
  - **Services:** Select the Microsoft 365 services that you want to back up:
    - **Exchange Online Users**
    - **Exchange Online Groups**
    - **OneDrive for Business**
    - **SharePoint Online**
  - **Username:** Provide the administrator username required for **SharePoint Online** and **Exchange Online Groups**.
  - **Password:** Provide the administrator password required for **SharePoint Online** and **Exchange Online Groups**.
6. Click **Next** to proceed.

Add Inventory Item

1. Platform      2. Services      3. Configuration      4. Authentication

Display name: M365

Services: SharePoint Online ⓘ

Username: admin ⓘ

Password: ..... ⓘ

Cancel      Next

7. On the **Configuration** page, choose one of the following:

- **Automatically register a new Azure AD application:** When this option is selected, NAKIVO Backup & Replication automatically registers the Microsoft 365 account as a new application in Azure Active Directory and grants the required permissions after authentication.
- **Use an existing Azure AD application:** When this option is selected, you can manually add an existing Microsoft 365 application already registered in Azure Active Directory.

### Note

In order to add a Microsoft 365 account to NAKIVO Backup & Replication, the account has to be registered in Azure Active Directory.

8. Click **Next** to proceed.

9. The **Credentials** page opens if you have chosen **Use an existing Azure AD application** option on the **Configuration** page. Provide the following information:

- **Tenant ID:** Enter your Azure Tenant ID created when registering your Microsoft 365 account in the Azure Portal.
- **Azure Client ID:** Enter your Azure Client ID created when registering your Microsoft 365 account in the Azure Portal.
- **Azure Client Secret:** Enter your Azure Client Secret obtained from the Azure Portal. For more information on obtaining the Azure credentials, refer to the [Obtaining Microsoft 365 Credentials](#) below.

10. Optionally, select **Automatically grant required permissions**. This allows NAKIVO Backup & Replication to automatically check the required API permissions for the selected services of your organization's Microsoft 365 account and to add them if they are missing. If you do not select this option, you have to manually grant the required permissions. If you have selected this option, click **Next**. Alternatively, click **Finish** to complete adding the item to Inventory.
11. If you have selected **Automatically grant required permissions** on the **Credentials** page or **Automatically register a new Azure AD** application on the **Configuration** page, the **Authentication** page opens. Do the following:
  - Enter a name for the application in the **Application name** field.
  - The **Code** field shows the code that needs to be entered at <https://microsoft.com/devicelogin>.
  - Clicking the link will open the page in a new tab and you will be required to log in to the Microsoft Azure Cross-Platform Command-Line Interface (xplat-cli) using the provided code.
  - If the provided credentials are correct, the account will be authenticated and you will be able to continue with the discovery process.

## Notes

- The Global Administrator role must be assigned to the account in the Microsoft 365 admin center for the authentication process to be successful.
- If **SharePoint Online** was selected on **Services** page, the SharePoint Administrator role must be assigned to the account in Microsoft 365 admin center for the authentication process to be successful.

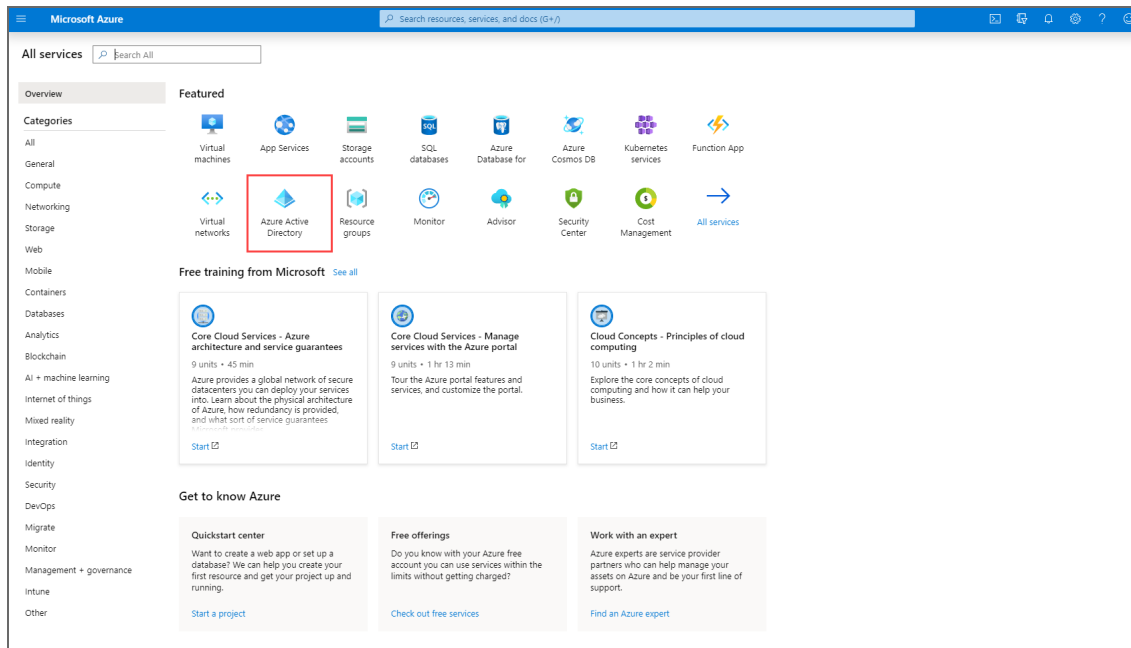
- If NAKIVO Backup & Replication is updated from a version that did not include support for Microsoft 365 Groups to a version that does, it is possible for the group mailboxes and group sites to be discovered automatically for existing Microsoft 365 accounts:
  - If you have selected Automatically grant required permissions on the Credentials page, the group mailboxes and group sites will be added automatically with all of the required permissions granted.
  - If you have selected Automatically grant required permissions on the Credentials page, the group mailboxes and group sites will remain undiscovered by the solution.

12. Click **Finish** to complete adding the item to Inventory.


## Obtaining Microsoft 365 Credentials

To obtain the credentials required to add a Microsoft 365 account to Inventory in NAKIVO Backup & Replication, follow the steps below:


1. Open the Azure Portal by going to [portal.azure.com](https://portal.azure.com)
2. Sign in to Microsoft Azure with your Microsoft 365 account credentials.
3. Select **Azure Active Directory** from the Dashboard or from the Portal Menu.



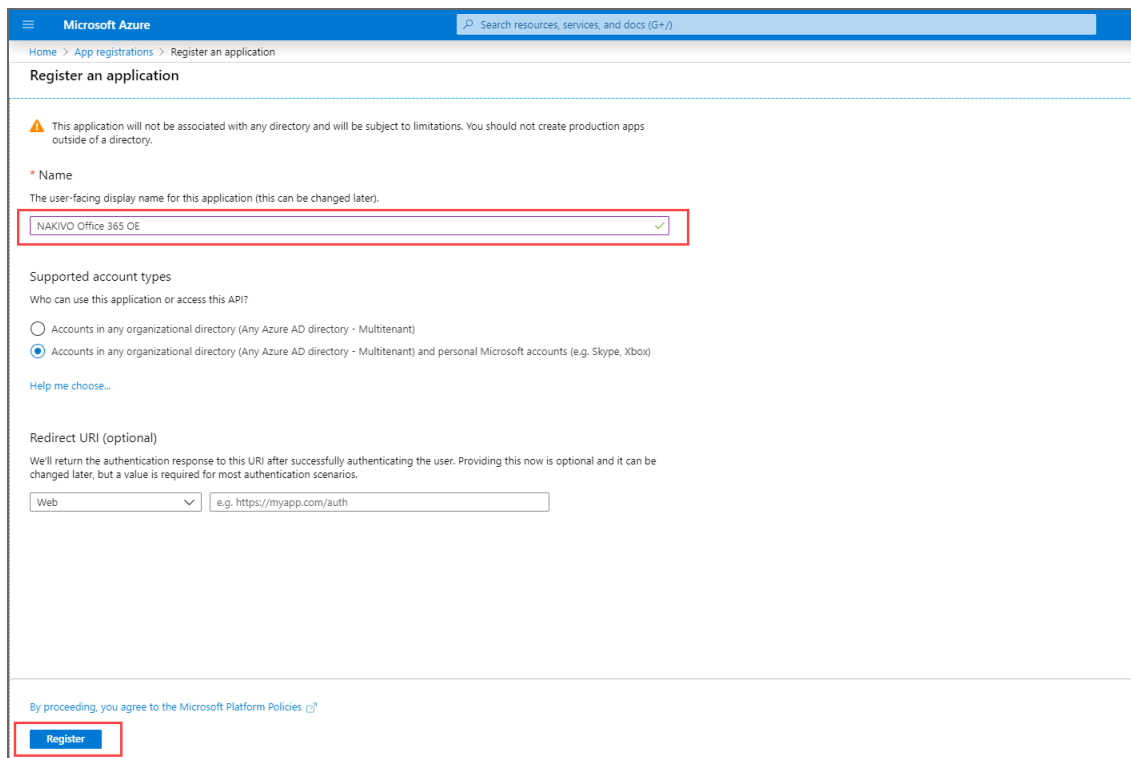
4. In the left menu, click **App registrations**.



5. Click **New registration** on the **App registrations** page.




6. On the **Register an application** page, enter a name for the application and click **Register**.



The application has been successfully registered and **Tenant ID** and **Azure Client ID** are displayed: **Directory (tenant) ID** and **Application (client) ID** respectively.

7. Click **View API Permissions** to add the necessary permissions.

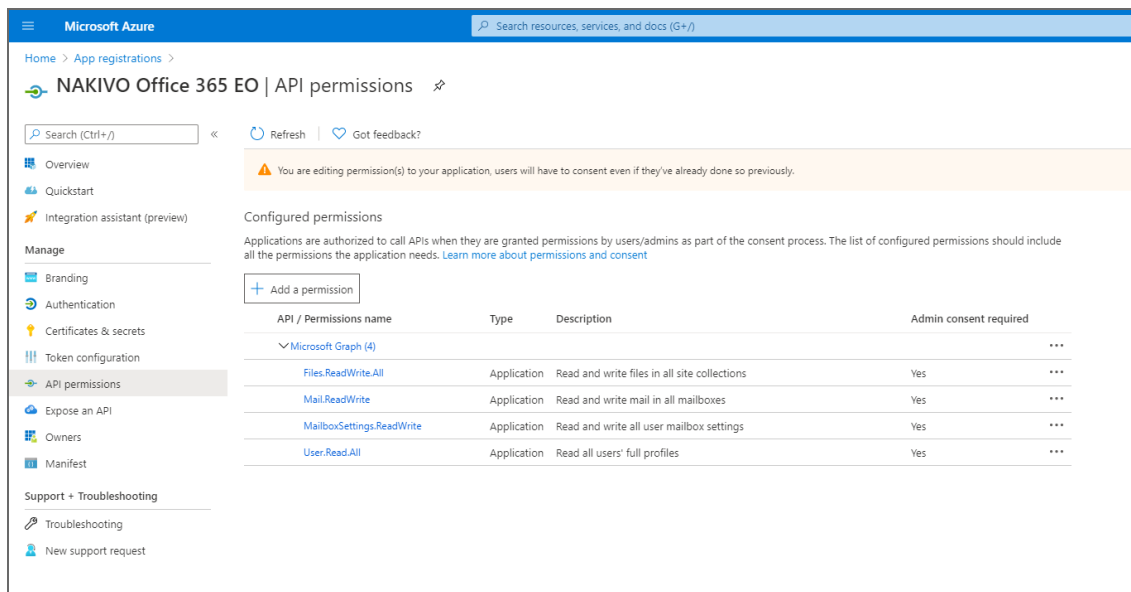


8. Click **Microsoft Graph**.
9. Click the **Application permissions** tab.
10. Provide the necessary API Permissions. Refer to [Required API Permissions from Microsoft 365](#) for details.

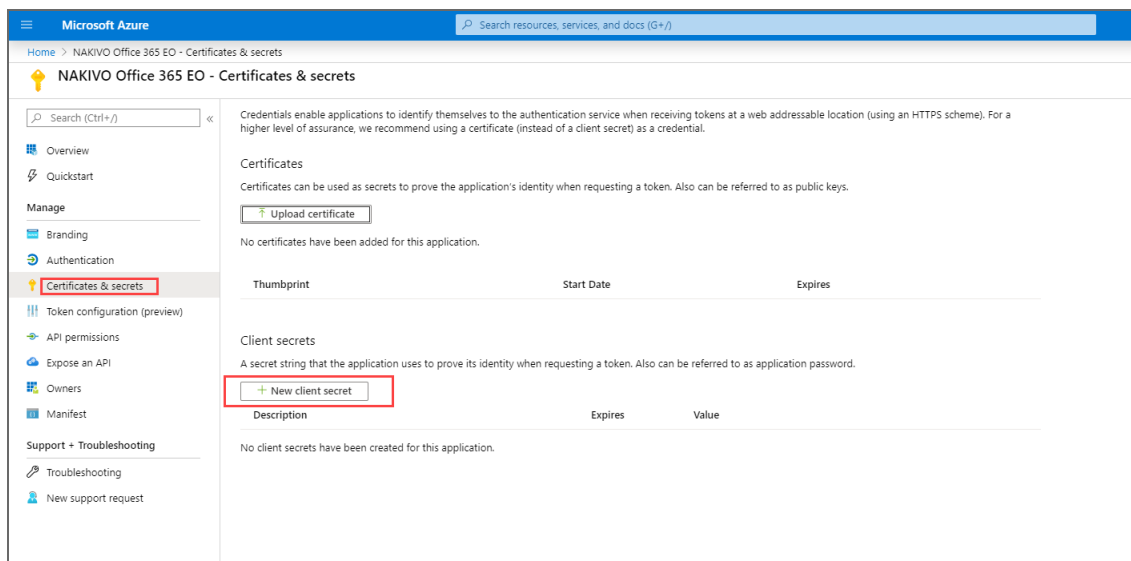
## Notes

- To skip discovering Exchange Online mailboxes, OneDrives or SharePoint sites in inventory, disable the API permissions for the corresponding service.
- If the necessary Microsoft Exchange Online/OneDrive/Sharepoint API permissions are not provided, the corresponding service will not be discovered by NAKIVO Backup & Replication.
- If the necessary API permissions for Microsoft Exchange Online contact and calendar items are not provided, the items will not be supported for backup and recovery operations.
- To recover messages and contacts containing a lot of content, you also need to enable full\_access\_as\_app for Office 365 Exchange Online in **APIs my organization uses**.


### 11. Click **Update Permissions**.




### 12. Click **Certificates & secrets > New client secret** to create a new client secret for your app.



### 13. Enter a description for the client secret, select the expiration period, and click **Add**.



The new **Client secret** is generated.



Make sure to save the client secret ID in a safe location. If you lose it, you will need to generate a new one.

## Adding File Shares to Inventory

To back up NFS/CIFS shares located on a NAS device, you first need to add these shares to the NAKIVO Backup & Replication Inventory.

To add an NFS/CIFS share located on a NAS device to the NAKIVO Backup & Replication Inventory, do the following:

1. Verify that all the [necessary requirements](#) have been met.
2. In the left pane, click **Settings > Inventory**.
3. Click **Add New**. The **Add Inventory Item** wizard opens.
4. On the **Platform** tab, select **File Share** and click **Next** to proceed.
5. On the **Options** tab, provide the following information:



- **Display name:** Enter a name for the folder.
- **Connection type:** Select the appropriate connection type. Choose either CIFS or NFS.
- **Assigned transporter:** Select the transporter that supports the chosen connection type.
- **Path to the share:** Provide the path to the CIFS or NFS share.
- **Username:** For a CIFS connection type, provide the username to the CIFS folder.
- **Password:** For a CIFS connection type, provide the password to the CIFS folder.
- Select **Advanced mount options** if needed.

Click **Finish** to add the file share to the Inventory.

#### **Note**

- You can add a file share to the Inventory only once.
- After a NFS or CIFS share is discovered in the Inventory, the ability to read its state may not be available for a few minutes. That is to say, refreshing the Inventory item during this time may not return accurate data.
- Due to Microsoft limitations, files or folders whose names end in a space or period cannot be opened or created. If a Windows-based transporter is used to discover a file share whose path ends in a space or period, discovery will fail.

# Adding HPE 3PAR Storage Devices to Inventory

To add an HPE 3PAR storage device to 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 **Add New**.
4. On the **Platform** page of the wizard, select **Storage Device** and click **Next** to proceed.
5. On the **Type** page of the wizard, select **HPE 3PAR storage device** 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.
  - **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.
  - Optionally, enable **Direct Connect** for this device by doing the following:
    - Select **Use Direct Connect**.
    - In the **Assigned Transporter** drop-down list, choose a transporter with **Direct Connect** enabled.

## Note

This option is only available if there is a Transporter with **Direct Connect** enabled. For details, refer to [“Adding Installed Transporters” on page 365](#).

7. Click **Finish** when you are done.

Add Inventory Item

1. Platform      2. Type      3. Options

---


Display name:

Hostname or IP:

Username:

Password:

Web services port:

Use Direct Connect 

Assigned transporter:

# Adding HPE Nimble Storage Devices to Inventory

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 **Add New**.
4. On the **Platform** page of the wizard, select **Storage Device** and click **Next** to proceed.
5. On the **Type** page of the wizard, select **HPE Nimble storage device** 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.
  - **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.

The screenshot shows a web interface titled "Add Inventory Item" with three steps: 1. Platform, 2. Type, and 3. Options. Step 3 is active. The form contains the following fields:

Display name:	<input type="text" value="HPE Nimble"/>
Hostname or IP:	<input type="text" value="10.10.10.10"/>
Username:	<input type="text" value="admin"/>
Password:	<input type="password" value="*****"/>
REST API port:	<input type="text" value="5392"/>

# Adding Oracle Database to Inventory

To add an Oracle Database to the inventory, follow the steps below:

1. Make sure you meet the necessary [requirements](#).
2. Click **Settings** in the left pane of NAKIVO Backup & Replication.
3. Go to the **Inventory** page and click **Add New**.
4. On the **Platform** page of the wizard, select **Application** and **click** Next to proceed.
5. On the **Options** page of the wizard, provide the following information:

1. **Display name:** Enter the preferred name for the database.
  2. **Hostname or IP:** Specify the hostname or IP address of the Oracle database you want to add.
  3. **OS username:** Enter your OS username.
  4. **OS password:** Enter your OS password.
  5. **Database SID:** Enter your database SID, that is, the Oracle System Identifier specific to that database.
  6. **Oracle username:** If the OS credentials do not belong to special Oracle groups (OSDBA/OSOPER), enter your Oracle username. Otherwise, leave the field empty.
  7. **Oracle password:** If the OS credentials do not belong to special Oracle groups (OSDBA/OSOPER), enter your Oracle password. Otherwise, leave the field empty.
6. Click **Finish**. The Oracle database is added to the inventory.

The screenshot shows a dialog box titled "Add Inventory Item" with two tabs: "1. Platform" (selected) and "2. Options". The "1. Platform" tab contains the following fields:


Display name:	Oracle
Hostname or IP:	192.33.11.11
OS username:	admin
OS password:	*****
Database SID:	oracle
Oracle username:	admin2
Oracle password:	*****

At the bottom right of the dialog, there are two buttons: "Cancel" and "Finish". The "Finish" button is highlighted in blue.

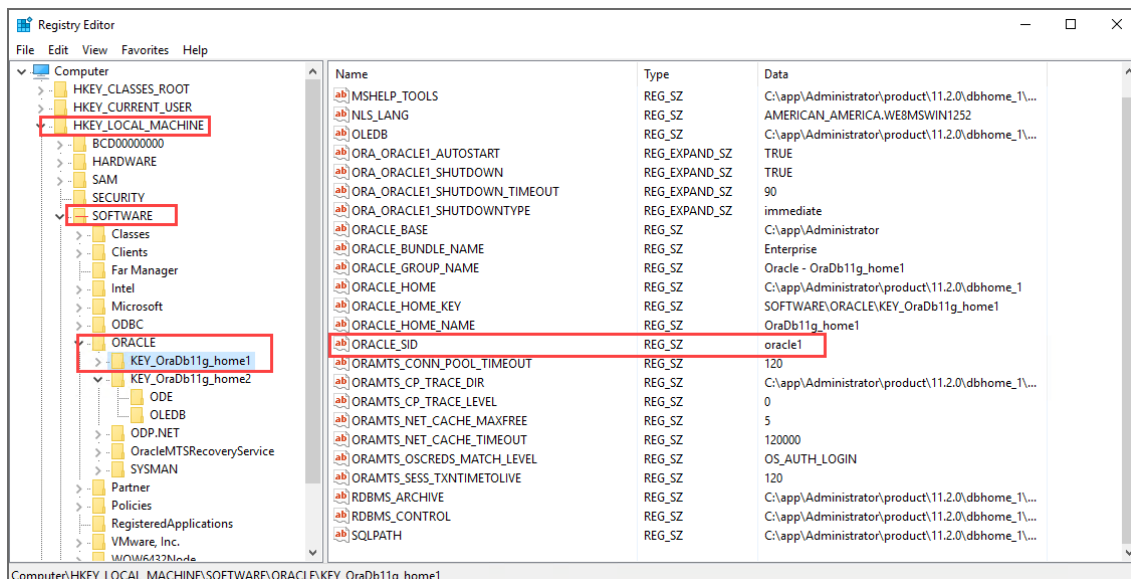
## Obtaining Oracle Database SID

To obtain your Oracle Database SID, follow the steps below:

1. Open the **Run** command box (Windows key+R key).
2. To open the registry editor, type `regedit` in the **Run** command dialog box and click **OK**.



3. Navigate to the **Key Registry** of your Oracle Home. The location path is the following: `HKEY_LOCAL_MACHINE > SOFTWARE > ORACLE >`



4. The `ORACLE_SID` entry shows your SID.

## Managing Inventory

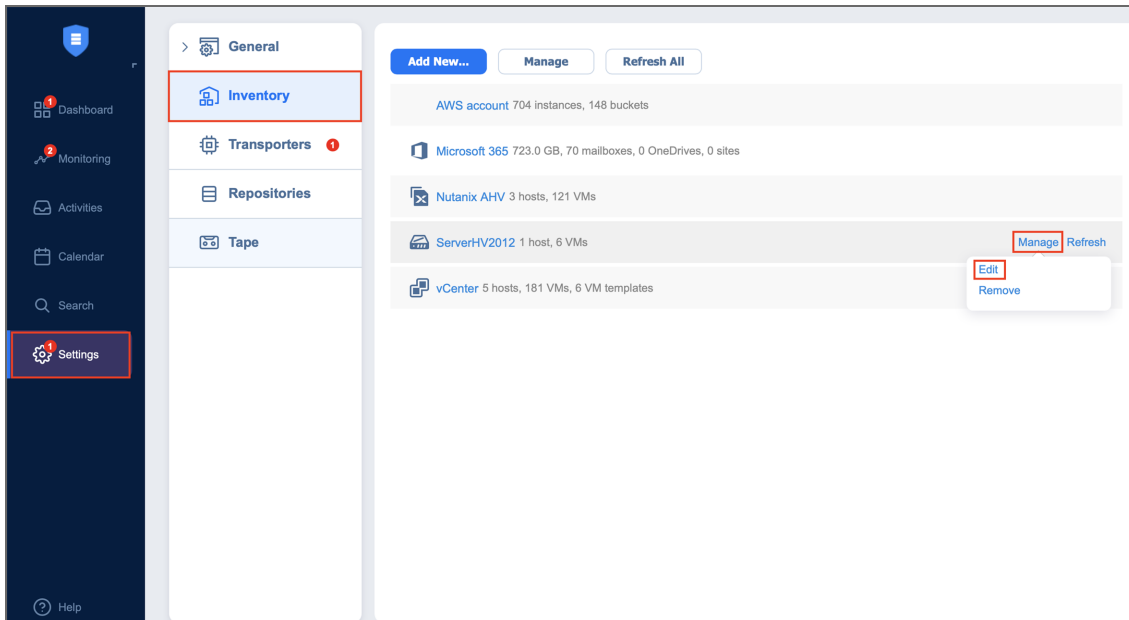
Refer to the following topics:

- [“Refreshing Inventory” on the next page](#)
- [“Editing Inventory Items” on the next page](#)
- [“Removing Items from Inventory” on page 360](#)

## 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 **Manage** on the right side and then click **Edit**.



5. Update the appropriate fields and click **Apply**.

## 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

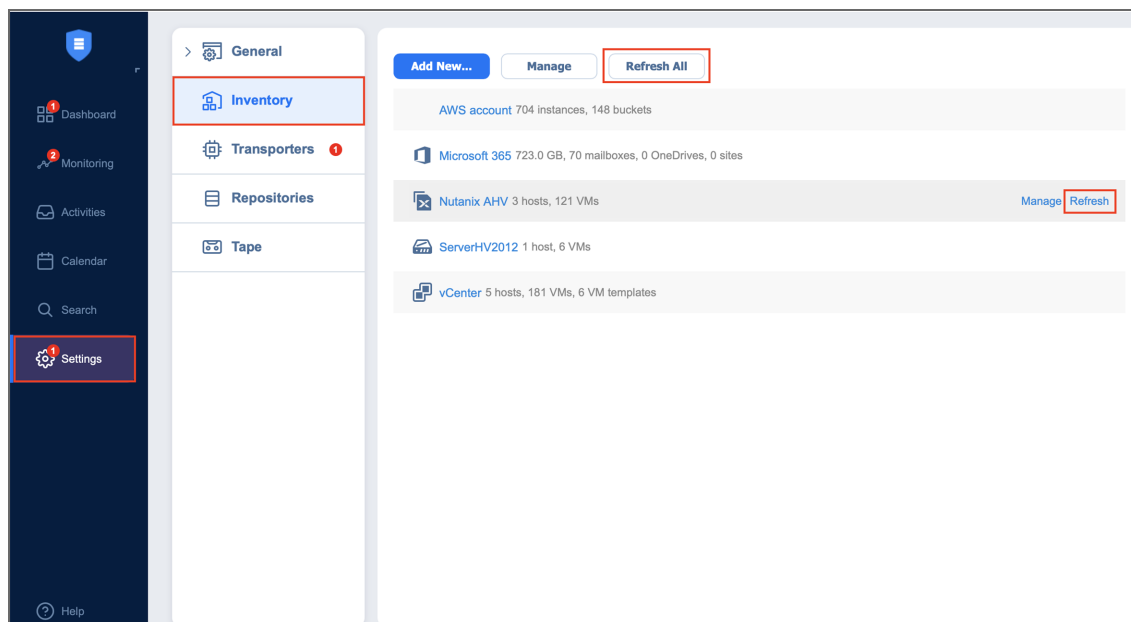
To refresh all inventory items, follow the steps below:

1. Click **Settings** in the left pane of the product and go to the **Inventory** tab.
2. Click **Refresh All**.

## 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 item that you would like to refresh.
3. Click **Refresh** on the right side.

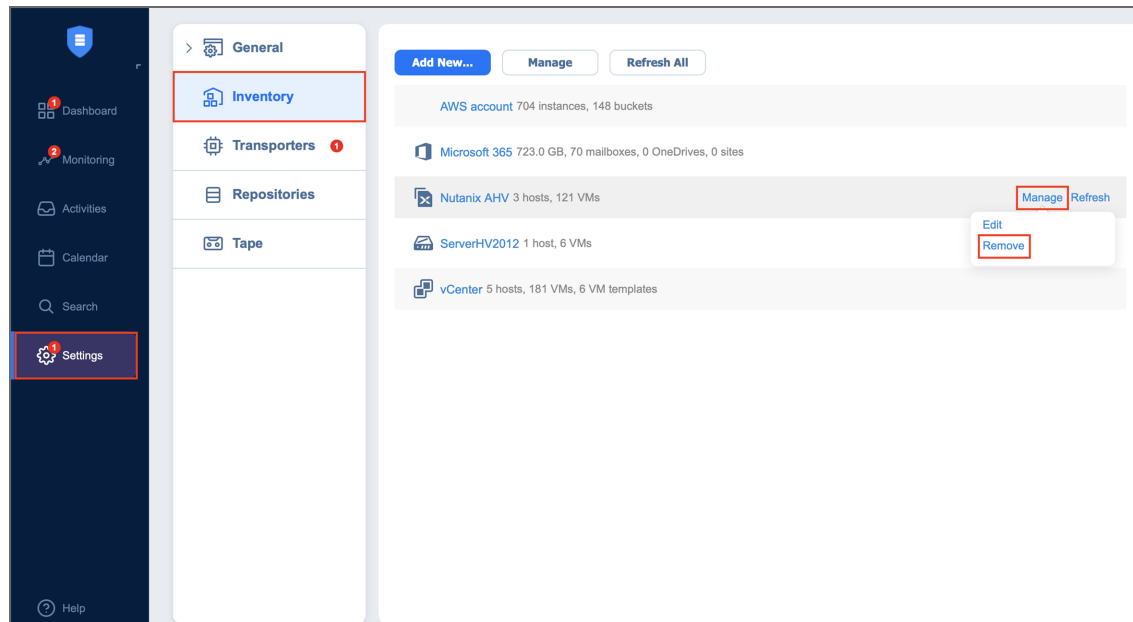


## Removing Items from Inventory

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 Credentials

NAKIVO Backup & Replication provides you with the ability to store your OS login and password, Amazon EC2 instance private keys or 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 that opens, click **Manage Credentials**.



5. In the **Manage Credentials** dialog that opens, click **Add Credentials**.

Manage Credentials

Type: Password

Username: Password

Password: Private Key

Repeat password:

Description:

[Learn more](#) **Save** Cancel

6. Then do the following:

- **Type:** Select the type of credentials:

- To add a username and password, fill out the **Username**, **Password**, and **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 a 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.

**Information**

Supported key formats: RSA, DSA

Supported file extensions: no extension, .pem, .key, .cer, .der, .txt

- f. Fill out the **Description** box.

g. Click **Save**.

Manage Credentials

Type: Private Key

Username: linux1

Password: .....

Repeat password: ..... 🔑

Private Key: Please upload the key

Description:

[Learn more](#)

You can now assign the credentials while creating jobs.

## 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 credentials**.
4. Hover the mouse pointer over the record that you would like to edit, and click **Edit**.

Manage Credentials

*Search*

admin	
nakivo	<input type="button" value="Edit"/> <input type="button" value="Delete"/>

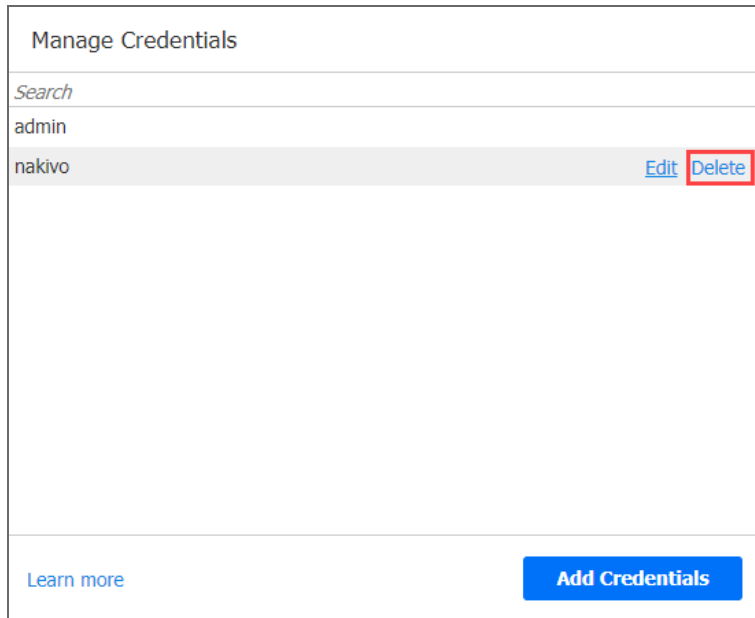
[Learn more](#)

5. Make any required changes, and then click **Save**.

## Deleting Credentials

Do the following:

1. Click **Settings** in the left pane of the product.
2. Go to the **Inventory** tab.
3. Click **Manage credentials**.
4. Hover the mouse pointer over the record that you would like to delete, and click **Delete**.



5. Click **Delete** in the confirmation dialog that opens.

# Transporters

The **Transporter** is one of NAKIVO Backup & Replication component that does all of the heavy-lifting: it performs backup, replication, and recovery, as well as data compression, deduplication, and encryption. To learn how to add an additional Transporter and how to manage it, refer to the topics below:

- [“Adding Installed Transporters” below](#)
- [“Deploying Transporter as VMware Appliance” on page 373](#)
- [“Deploying Transporter as Nutanix AHV Appliance” on page 376](#)
- [“Deploying Transporters in Amazon EC2” on page 377](#)
- [“Enabling Direct Connect” on page 380](#)
- [“Managing Transporters” on page 381](#)

## Adding Installed Transporters

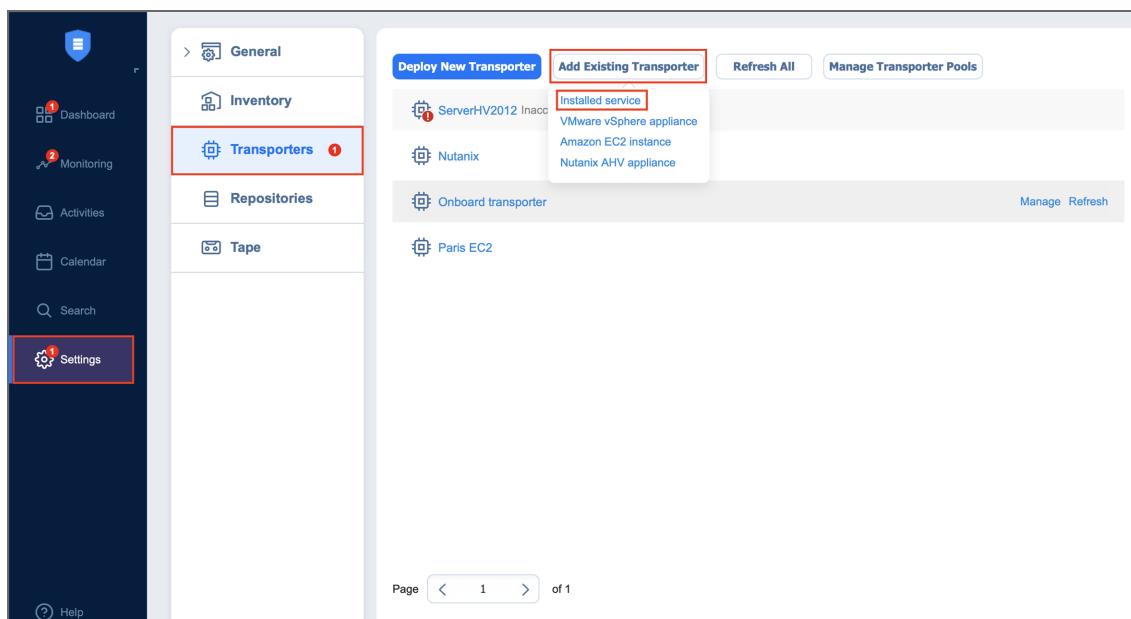
After you have installed a Transporter, you need to add it to NAKIVO Backup & Replication so that the Transporter can be used for backup, replication, and recovery tasks. Refer to the following topics:

- [Installed Service](#)
- [VMware Appliance](#)
- [Amazon EC2 Instance](#)
- [Nutanix AHV Appliance](#)

## Installed Service

Follow the steps below to add a Transporter that is installed as a service:

1. Click **Settings** in the left pane of the product and go to the **Transporters** tab.
2. Click **Add Existing Transporter** and then click **Installed service** in the dialog that opens.



3. The **Add Existing Transporter - Installed Service** dialog opens. In the **Hostname or IP** box, enter the IP address or hostname of the machine on which the Transporter is installed.

#### Note

If you are adding the Transporter by a DNS name, make sure this DNS name can be resolved on the machines on which the Director and any other Transporters (which you plan to use in conjunction with the current one) are installed.

4. Click **More options** to reveal and edit the following fields:
  - 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. Note that setting a master password is required when the **Enable Direct Connect** for this Transporter 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. Follow these steps:
  - Enter the following command `bhsvc -b P@ssword123`
  - [Restart](#) the Transporter service.
- 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. 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 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.
- **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.

5. Click **Add**. The Transporter is added to the product and can be used for backup, replication, and recovery jobs.

**Add Existing Transporter - Installed Service**

Hostname or IP:  ⓘ

**Networking**

Transporter port:  ⓘ

Data transfer ports:  ⓘ

**Settings**

Transporter name:

Maximum load:  concurrent tasks ⓘ

Additional load for recovery jobs:  concurrent tasks ⓘ

Enable Direct Connect for this transporter (requires master password) ⓘ

Enable debug logging for this transporter ⓘ

**Security**

Master password:  ⓘ Connect

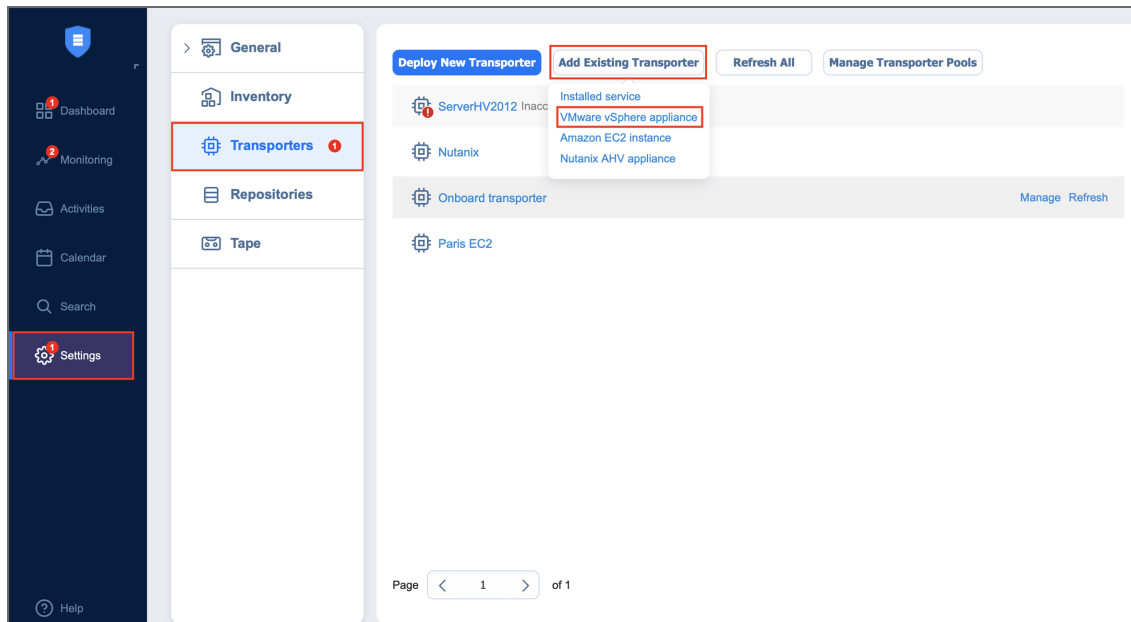
---

Cancel Add

## 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 dashboard and go to the **Transporters** tab.
2. Click **Add Existing Transporter** and then click **VMware vSphere appliance** in the dialog that opens.



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.

Add Existing Transporter - VMware vSphere Appliance

Host or cluster:

Virtual machine:

OS username:

OS password:

SSH port:

Networking

Transporter port:

Data transfer ports:

Settings

Transporter name:

Maximum load:  concurrent tasks

Additional load for recovery jobs:  concurrent tasks

Enable debug logging for this transporter

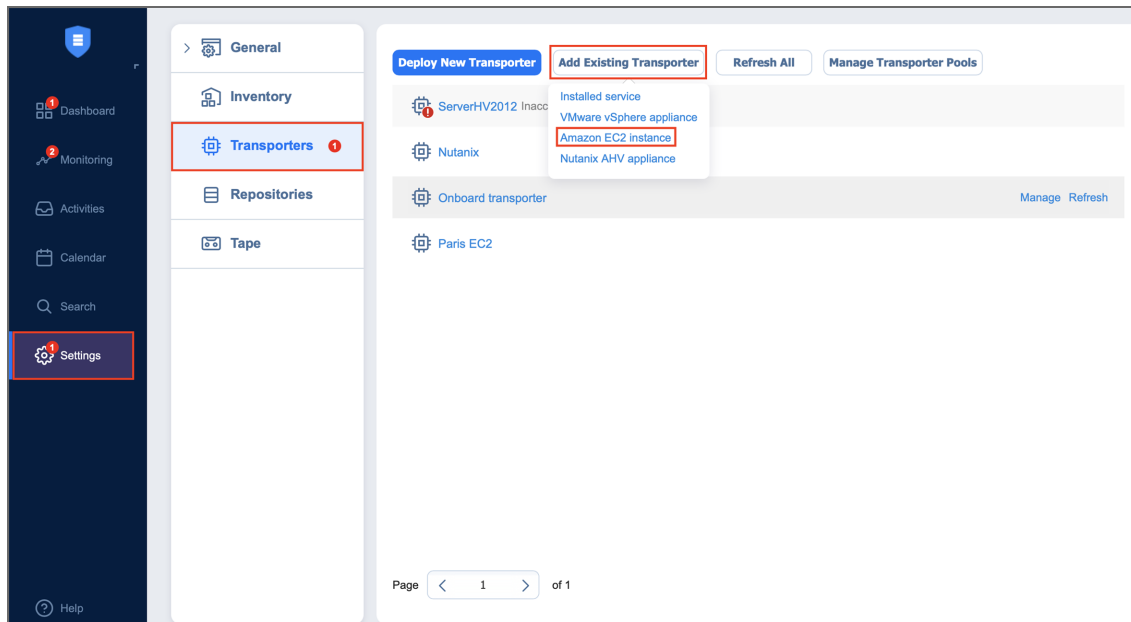
Enable Direct Connect for this transporter

## 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 and go to the **Transporters** tab.
2. Click **Add Existing Transporter** and then click **Amazon EC2 instance** in the pop-up that opens.





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.

Add Existing Transporter - Amazon EC2 Instance

AWS account:

Region:

EC2 instance:

Private key: Please upload the key

**Networking**

Transporter port:

Data transfer ports:

**Settings**

Operation mode:

Transporter name:

Maximum load:  concurrent tasks

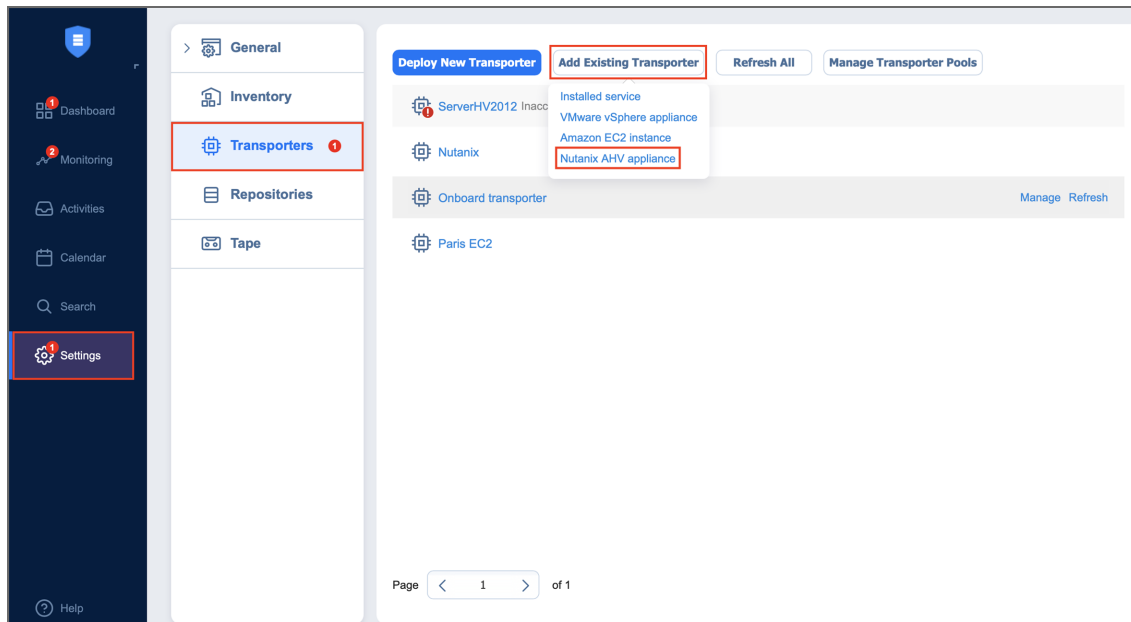
Additional load for recovery jobs:  concurrent tasks

Enable debug logging for this transporter

## 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 and go to the **Transporters** tab.
2. Click **Add Existing Transporter** and then 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.

Add Existing Transporter - Nutanix AHV Appliance

Cluster: Nutanix AHV

Virtual machine: 24

OS username: user

OS password: \*\*\*\*\*

SSH port: 2221

Transporter port: 9446

Data transfer ports: 9448-10000

Settings

Transporter name: Nutanix

Maximum load: 6 concurrent tasks

Additional load for recovery jobs: 2 concurrent tasks

Enable debug logging for this transporter

Cancel Add

4. Click **Add**. The Transporter is added to the product and can be used for backup, replication, and recovery jobs.

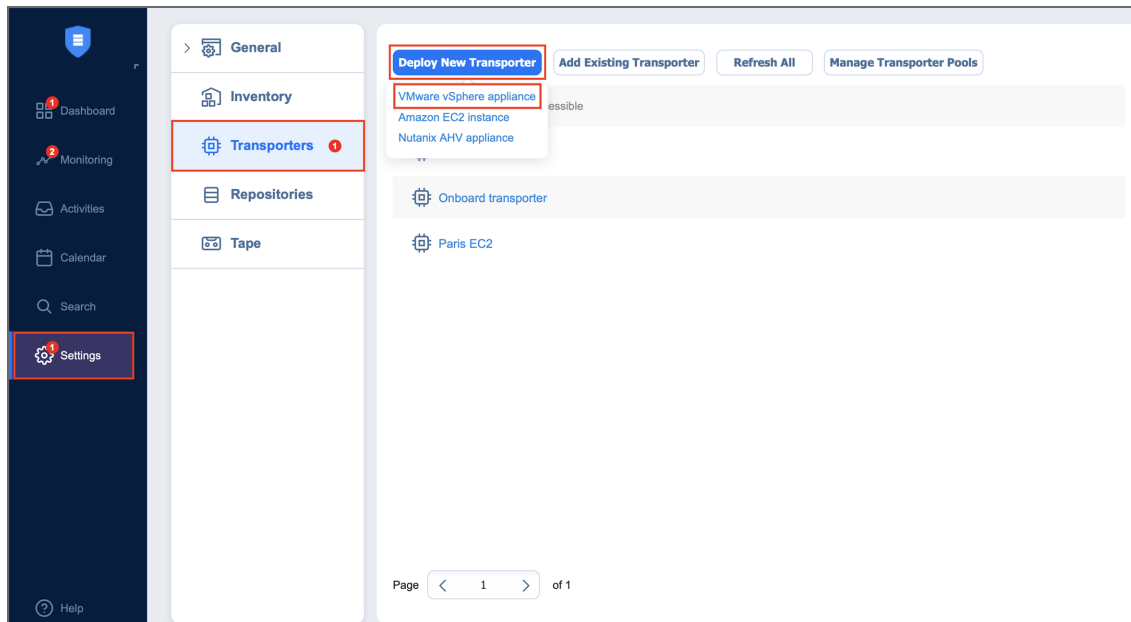
## Deploying Transporter as VMware Appliance

### 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. Go to the **Settings > Transporters** tab and click **Deploy New Transporter**.
2. In the dialog that opens, 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.

Deploy New Transporter - VMware vSphere Appliance

Transporter name:

Host or cluster:  ⓘ

Datastore:  ⓘ

Virtual network:  ⓘ

**Networking**

IP configuration:  ⓘ

IP address:

Subnet mask:

Default gateway:

DNS configuration:  ⓘ

Primary DNS:

Secondary DNS:

Transporter port:  ⓘ

Data transfer ports:  ⓘ

**Settings**

Maximum load:  concurrent tasks ⓘ

Additional load for recovery jobs:  concurrent tasks ⓘ

Enable debug logging for this transporter ⓘ

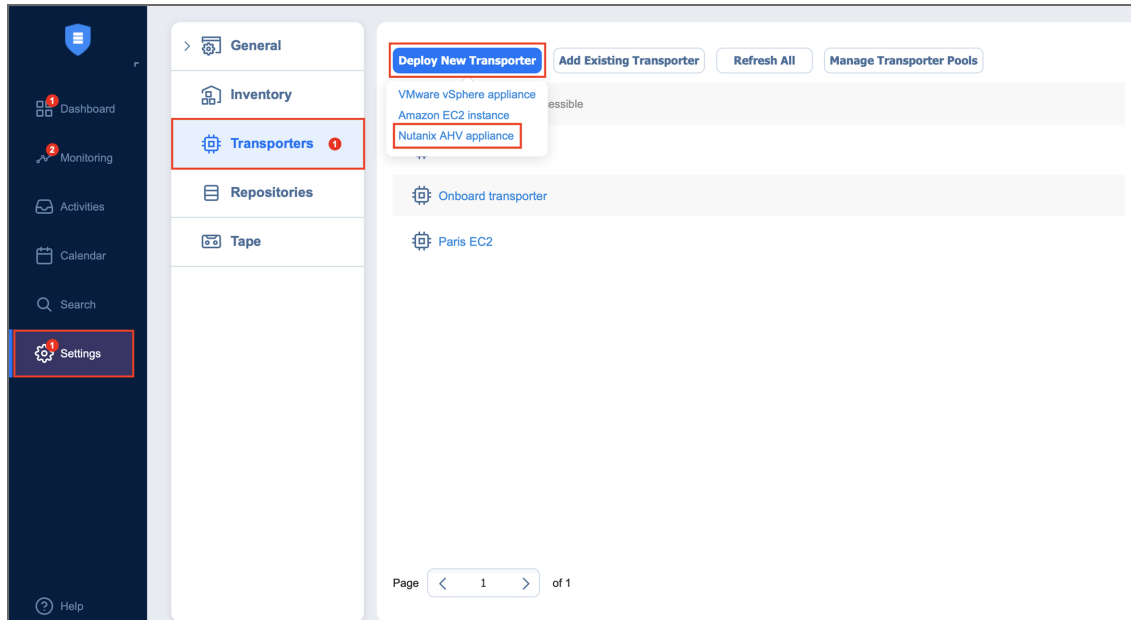
Enable Direct Connect for this transporter ⓘ

# 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. Go to the **Settings > Transporters** tab and click **Deploy New Transporter**.
2. In the dialog that opens, 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.

Deploy New Transporter - Nutanix AHV Appliance

Transporter name:

Cluster:  Nutanix AHV ?

Storage container:  NutanixManagementShare ?

Virtual network:  77 ?

**Networking**

IP configuration:  Automatic setup (DHCP) ?

IP address:

Subnet mask:

Default gateway:

DNS configuration:  Automatic setup (DHCP) ?

Primary DNS:

Secondary DNS:

Transporter port:  9446 ?

Data transfer ports:  ?

**Settings**

Maximum load:  6 concurrent tasks ?

Additional load for recovery jobs:  2 concurrent tasks ?

Enable debug logging for this transporter ?

6. Click **Deploy** to begin the deployment process. Successfully deployed Transporters are displayed in the **Transporters** tab.

## 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 and go to the **Transporters** tab.
2. Click **Deploy New Transporter** and click **Amazon EC2 instance** in the resulting drop-down list.





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.
  - 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.
  - **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.

#### 4. Click **Deploy**.

Deploy New Transporter - Amazon EC2 Instance

Transporter name:

Region:

Instance type:

**Networking**

Automatically configure VPC for this transporter ?

Network:

Subnet:

Allowed traffic from:  ?

Transporter port:  ?

Data transfer ports:  ?

**Settings**

Operation mode:

Platform:

Maximum load:  concurrent tasks ?

Additional load for recovery jobs:  concurrent tasks ?

Enable debug logging for this transporter ?

#### 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.

## Enabling Direct Connect

To start using this feature, take the following steps:

1. Download NAKIVO Backup & Replication installer or a VA with Transporter.
2. [Install Transporter](#) at the remote environment and set the master password during the installation. Alternatively, deploy [Transporter VA](#) at the remote environment. The master password for the VA can be set after the deployment.
3. 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"
```
4. [Restart the Transporter service](#) after generating the password to apply the change.
5. Expose the Transporter port on the local machine to be externally available via the Internet.
6. Add the remote transporter to NAKIVO Backup & Replication by entering the Transporter details.
7. Select **Enable Direct Connect for this transporter**.
8. Enter master password.
9. Click **Connect**.
10. View Transporter certificate information.
11. Accept received certificate.
12. Confirm adding the Transporter.

Hostname or IP:  ?

Networking

Transporter port:  ?

Data transfer ports:  ?

Settings

Transporter name:

Maximum load:  concurrent tasks ?

Additional load for recovery jobs:  concurrent tasks ?

Enable Direct Connect for this transporter (requires master password) ?

Enable debug logging for this transporter ?

Security

Master password:  ?

## Managing Transporters

Refer to the following topics:

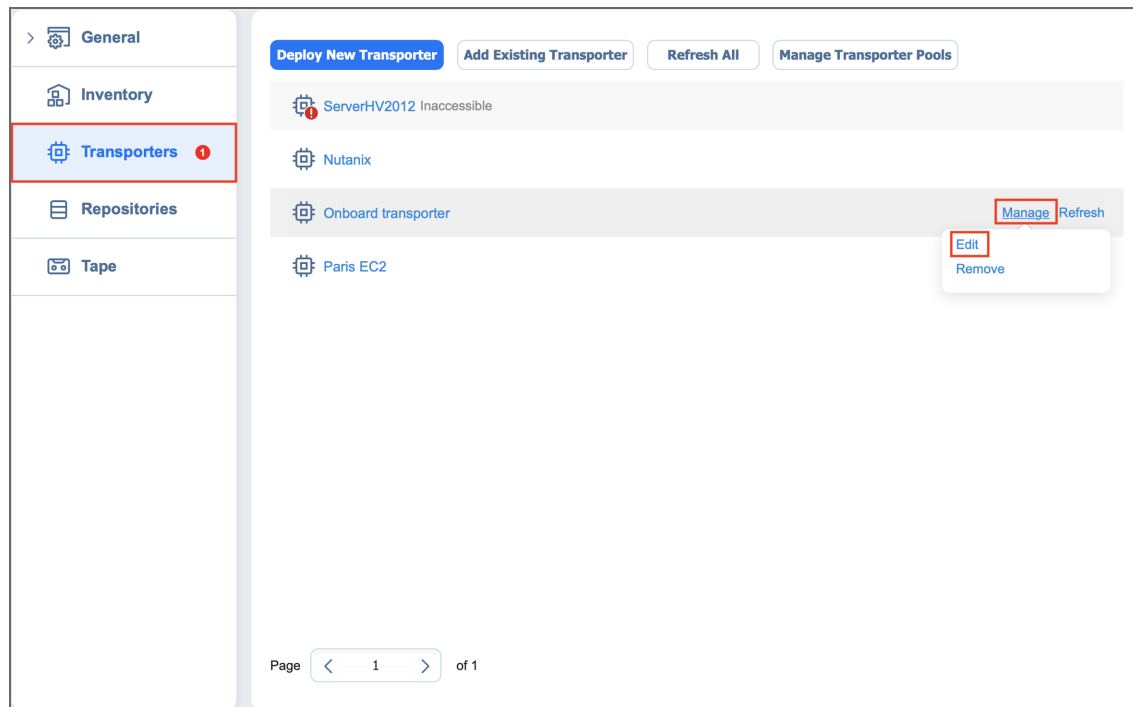
- [“Editing Transporters” below](#)
- [“Downloading Transporter's Credentials” on page 386](#)
- [Managing Transporter Pools](#)
- [“Refreshing Transporter Details” on page 384](#)
- [“Removing \(Deleting\) Transporters” on page 386](#)

## Editing Transporters

To modify the settings of an existing Transporter, follow the steps below:

1. Click **Settings** in the left pane of the product.
2. Go to the **Transporters** tab and hover over the Transporter you would like to edit.

3. On the right side, click **Manage** and then click **Edit**.



4. A dialog opens for editing the Transporter settings. Edit the settings as required:

- **Hostname or IP:** Here you can edit the IP address or hostname of the machine on which the Transporter is installed.
  - In the *Networking* section:
    - **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:
    - **Transporter name:** Edit the name of your Transporter.
    - **Maximum load:** Edit the number of tasks concurrently processed by the Transporter.
    - **Additional load for recovery jobs:** If selected, the specified amount of tasks will be added to the 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 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.
- **Enable debug logging for this transporter:** Enable/disable debug level logging for the Transporter. Having this option enabled on a permanent basis is not recommended.

5. Click **Apply** to save your changes.

Edit: Onboard transporter

Hostname or IP:  ⓘ

Networking

Transporter port:  ⓘ

Data transfer ports:  ⓘ

Settings

Transporter name:

Maximum load:  concurrent tasks ⓘ

Additional load for recovery jobs:  concurrent tasks ⓘ

Enable Direct Connect for this transporter ⓘ

Enable debug logging for this transporter ⓘ

Security

Master password:  ⓘ Connect

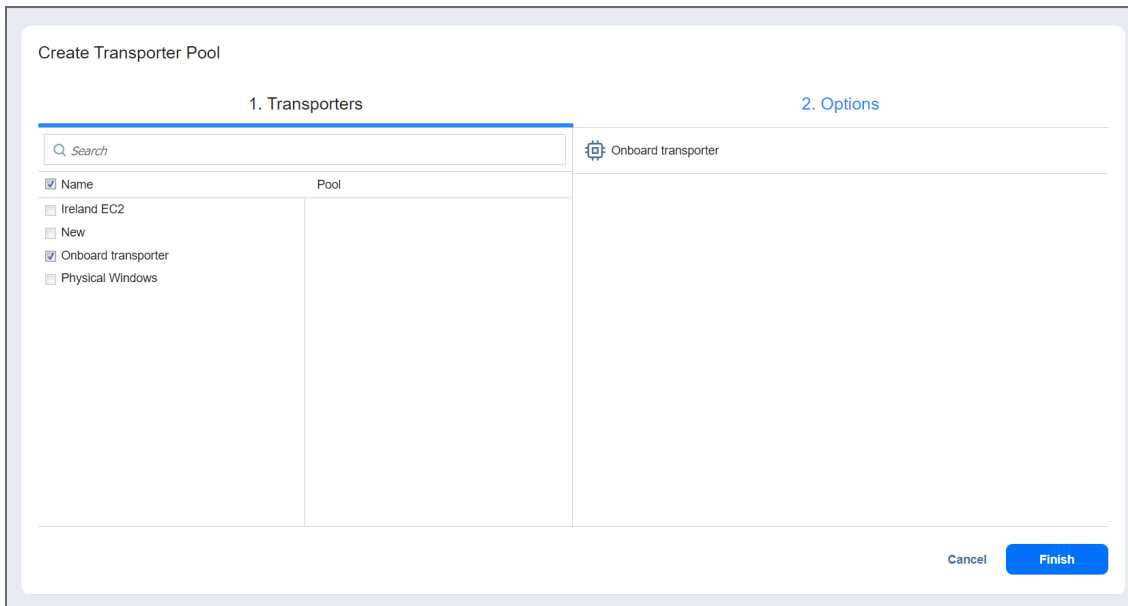
---

Cancel Apply

## Managing Transporter Pools

NAKIVO Backup & Replication allows you to group VMware Transporters into pools to optimize VMware backup, replication, and recovery jobs. To create a Transporter pool, take the following actions:

- Navigate to **Settings**.
- Click the **Transporters** tab.
- Click **Manage Transporter Pools**. The Manage Transporter Pools dialog box opens.
- Click Create Transporter Pool.
- Complete the Transporter Pool Creation wizard and click **Finish**.



A Transporter pool can be selected in the *Data Transfer* section on the **Options** page of VMware backup, replication, and recovery jobs. A Transporter can be included in only one pool. To move a Transporter from one pool to another, you need to remove it from the original pool first.

## Refreshing Transporter Details

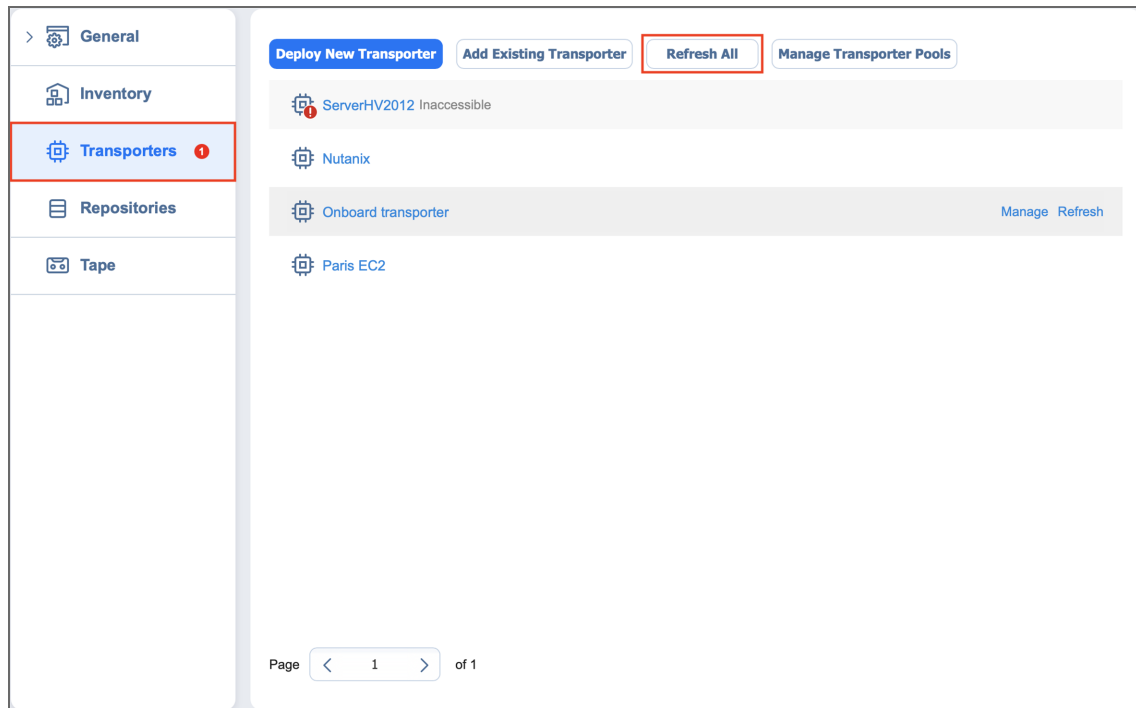
By default, NAKIVO Backup & Replication refreshes the information about Transporters every hour. During the refreshing process, the product collects all the required information about all Transporters. Only one Transporter can be refreshed at a time. If you have more than one Transporter, all others will remain in the queue until they are able to be refreshed.

- [Manually Refreshing All Transporters](#)
- [Manually Refreshing a Single Transporter](#)

## Manually Refreshing All Transporters

To refresh all Transporters, follow the steps below:

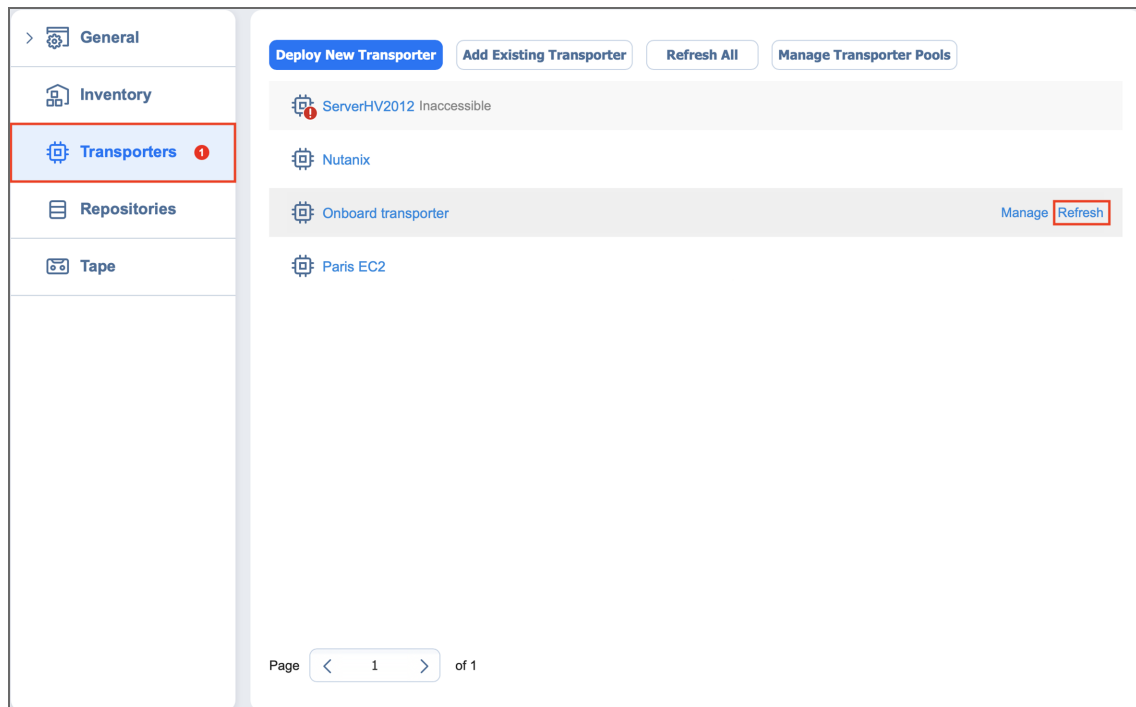
1. Click **Settings** in the left pane of the product and go to the **Transporters** tab.
2. Click **Refresh All**.



## Manually Refreshing a Single Transporter

To refresh a single Transporter, follow the steps below:

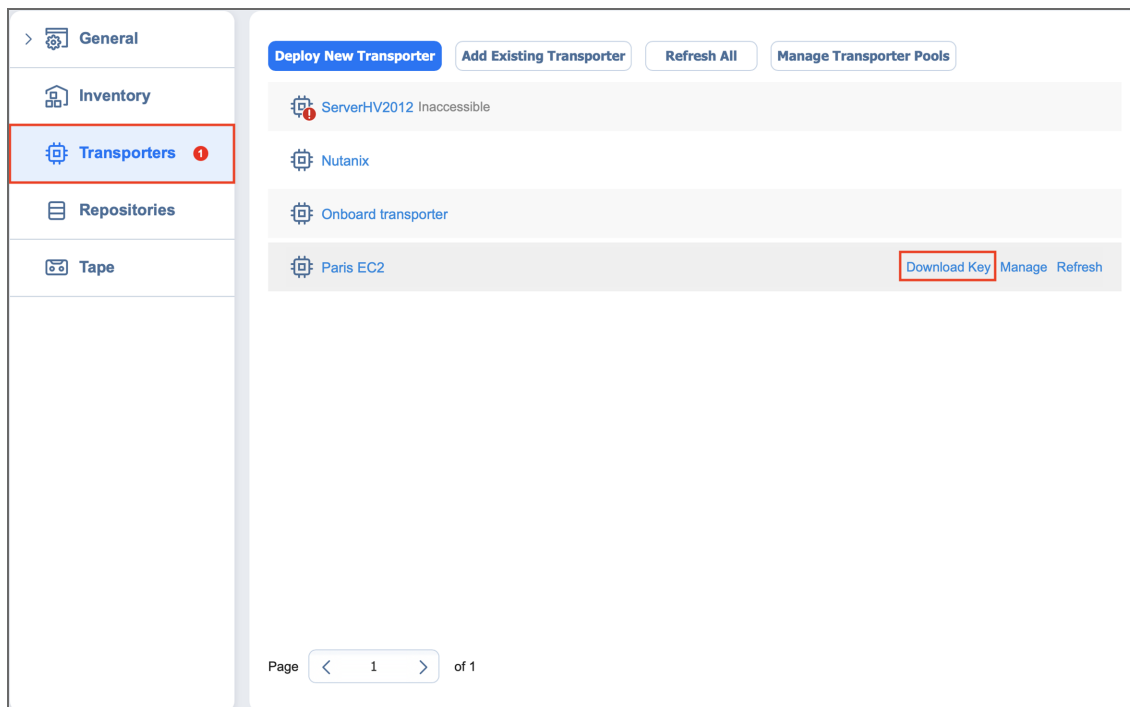
1. Click **Settings** in the left pane of the product.
2. Go to the **Transporters** tab.
3. Hover over the Transporter you would like to refresh.
4. On the right side, click **Refresh**.





## Downloading Transporter's 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 **Download Key** button on the right. This begins the download of a ZIP file containing the Transporter's credentials.

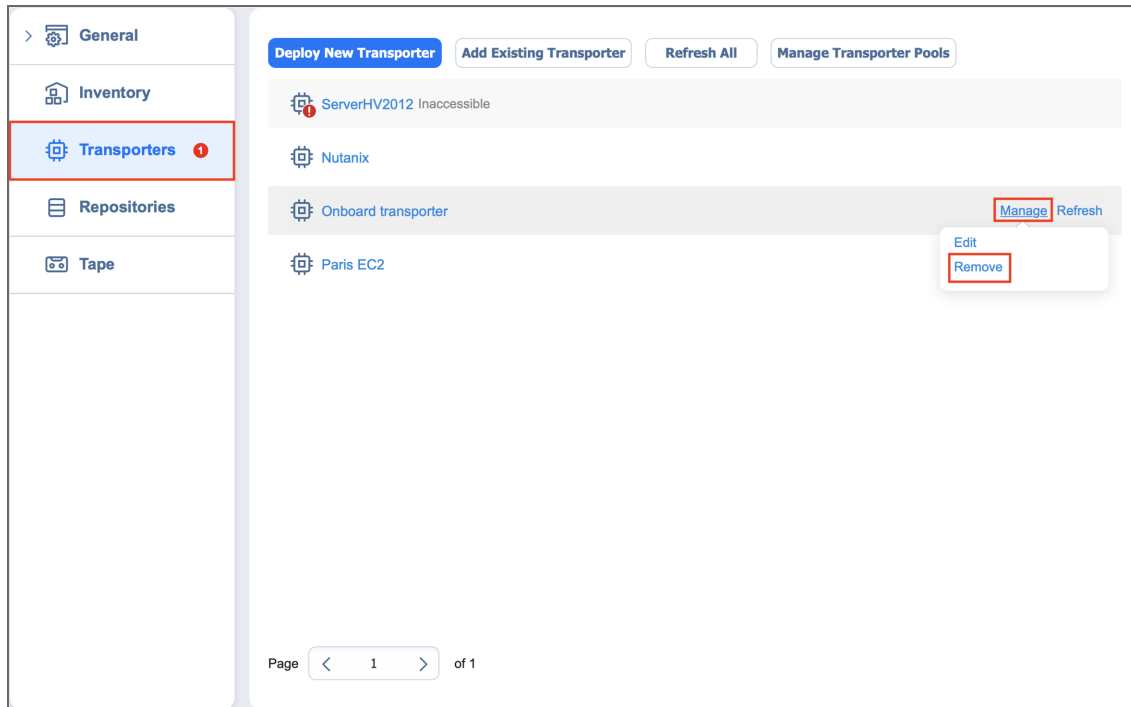


## Removing (Deleting) Transporters

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 **Transporters** tab.
3. Hover over the Transporter you would like to remove.

4. On the right side, click **Manage** and then click **Remove**.



### Note

The following Transporters cannot be removed:

- The Onboard Transporter (which is installed with the [“Director” on page 87](#) by default)
- Transporters manually assigned to a job
- Transporters assigned to the backup repositories in Amazon Cloud.

# 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 94](#).

This section covers repository-related topics such as creation, management, etc. of Backup Repositories and contains the following articles:

- [“Creating Backup Repositories” on page 390](#)
- [“Adding Existing Backup Repositories” below](#)
- [“Viewing Backup Repository Details” on page 440](#)
- [“Managing Backup Repositories” on page 424](#)

## Adding Existing Backup Repositories

NAKIVO Backup & Replication provides you with the ability to add an existing Backup Repository to a new copy of the product.

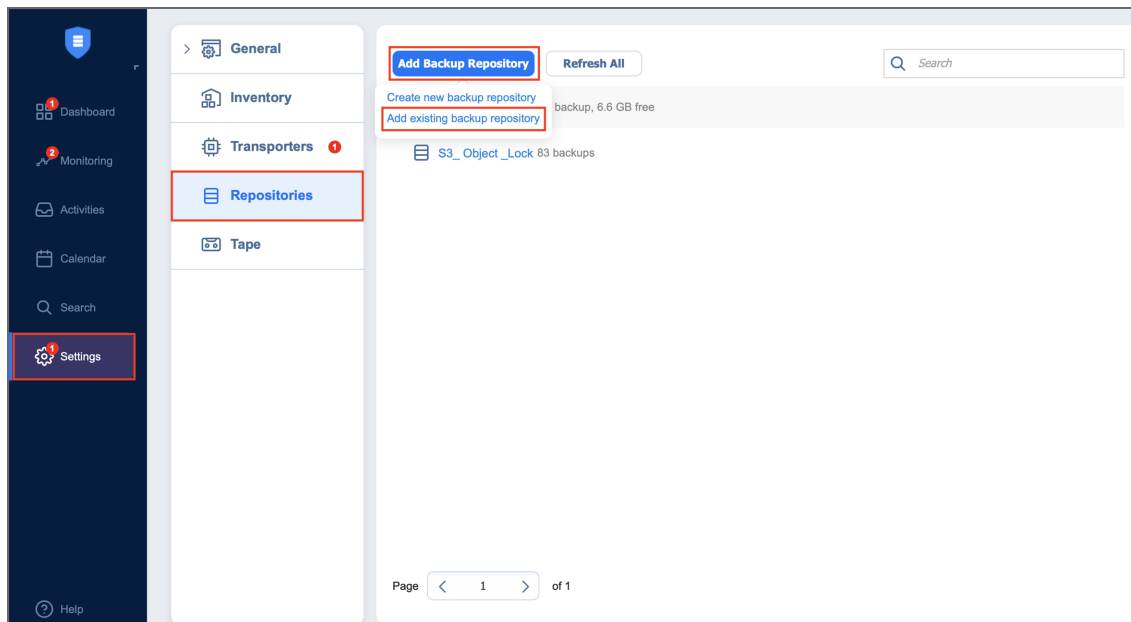
### Note

During the import process, NAKIVO Backup & Replication searches for the "NakivoBackup" folder in the specified location, so if your Backup Repository is located in `E:\backup\NakivoBackup`, you need to specify the following path: `E:\backup`

To import an existing Backup Repository, do the following:

1. From the main menu, click **Settings**.
2. Go to the **Repositories** tab and click **Add Backup Repository**.

3. Click **Add existing backup repository** in the dialog 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**
- **Amazon EC2**
- **Amazon S3**
- **Wasabi**
- **SaaS**
- **Deduplication Appliance**

5. 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.

6. 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 repository self-healing on schedule:** If required, select this checkbox to additionally 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.

- **Run full data verification on schedule:** If selected, the product 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 process. This way the product will verify each recovery points 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.
- **Reclaim unused space on schedule:** If required, select this option to run the Backup Repository [space reclaim](#) process 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 Saturday at 12 PM.
- **Enforce explicit file system sync:** If selected, explicit sync with the file system will be 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 on a schedule. Detaching a Backup Repository saves the Backup Repository data and metadata in a consistent state and stops the product interaction with the Backup Repository (so that the Backup Repository 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.

7. Click **Finish**. The Backup Repository is imported to the list.


## Creating Backup Repositories

NAKIVO Backup & Replication allows you to create additional Backup Repositories for storing and restoring your backup. You can use a local share, NFS share, CIFS share, EBS/Wasabi storage, or a deduplication appliance as a Backup Repository location. To create a new Backup Repository, follow the steps below:

### Important

Do not create Backup Repositories 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 **Add Backup Repository**.
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 418](#)
- [“Backup Repository on CIFS Share” on page 402](#)
- [“Backup Repository on NFS Share” on page 413](#)
- [“Backup Repository in Amazon EC2” below](#)
- [“Backup Repository in Amazon S3” on page 396](#)
- [“Backup Repository in Wasabi Hot Cloud Storage” on page 399](#)
- [“SaaS Backup Repository” on page 422](#)
- [“Backup Repository on Deduplication Appliance” on page 407](#)

## 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 the antivirus software running on the machine on which the NAKIVO Backup Repository is set up.
- You may be additionally charged for using a 3rd-party resource. Please refer to the 3rd-party resource provider documentation for details.

To create a repository on an Amazon EC2, 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 **Amazon EC2** and click **Next** to move to the next page of the wizard.

Repositories / Create Backup Repository

1. Type      2. Name & Location      3. Options

Local Folder  
Directly attached storage or locally mounted share on an assigned transporter.

CIFS Share  
Network share used by transporter through CIFS Interface.

NFS Share  
Network share used by transporter through NFS Interface.

Amazon EC2  
EBS storage attached to transporter running in Amazon EC2 instance.

Amazon S3  
Highly scalable AWS object storage.

Wasabi  
Cost effective cloud based object storage.

SaaS  
Directly attached storage or locally mounted share on an assigned transporter used for storing Microsoft 365 objects.

Deduplication Appliance  
Integration with Dell EMC Data Domain Boost, HPE StoreOnce Catalyst and NEC HYDRAsTOR.

**Next**      Cancel

## 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. Click **Next** to go to the next page of the wizard.

Repositories / Create Backup Repository

1. Type      2. Name & Location      3. Options

Name:

Assigned transporter:  ?

**Next**      Cancel

## 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 the 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:** A Backup Repository in Amazon EC2 is created by using multiple EBS Volumes (chunks). The maximum size of the Backup Repository is limited by 50 EBS Volumes (chunks) or 16 TB (whatever comes 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 16000 GB, it is recommended that you have 400 GB storage chunk or bigger. Storage chunk cannot be changed later.
- **Automatically resize storage:** If this option is selected, the cloud storage will be automatically increased and reduced as required.

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.
- **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.

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

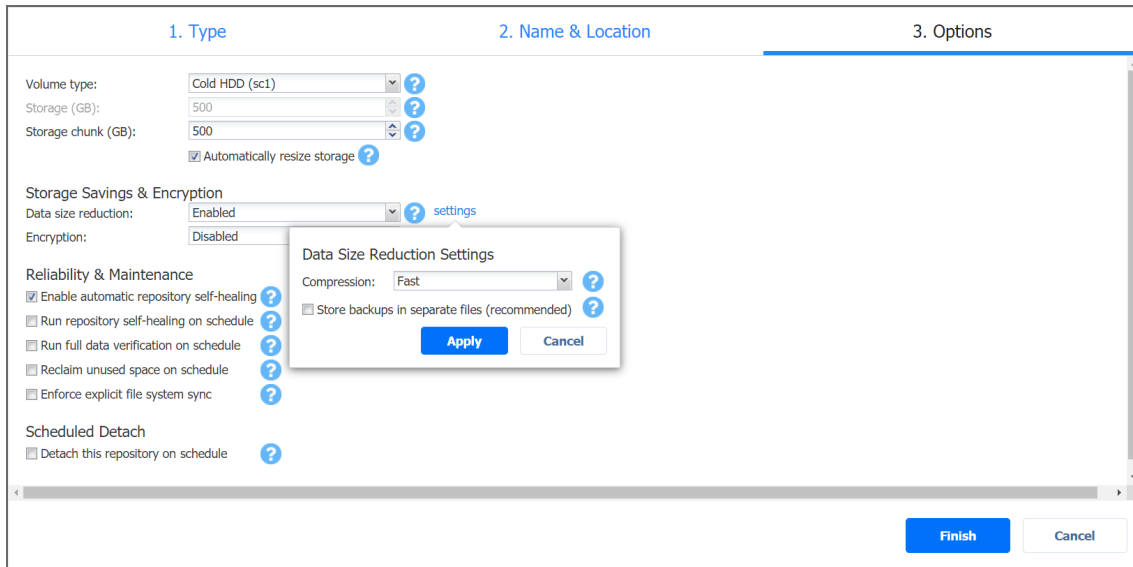
Do not reboot/disconnect the "null" Transporter and storage device while space reclaim is in progress to avoid Backup Repository corruption.

- **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.

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 the Backup Repository data and metadata in a consistent state and then stops the product's interaction with the Backup Repository (so that the Backup Repository 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.

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 Wasabi](#) 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 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: Name and Location](#)
- [Create Backup Repository: Options](#)

### Create Backup Repository: Type

On the **Type** page of the **Create Backup Repository** wizard, select **Amazon S3** and click **Next** to go to the next page of the wizard.

Repositories / Create Backup Repository

1. Type      2. Name & Location      3. Options

- Local Folder  
Directly attached storage or locally mounted share on an assigned transporter.
- CIFS Share  
Network share used by transporter through CIFS Interface.
- NFS Share  
Network share used by transporter through NFS Interface.
- Amazon EC2  
EBS storage attached to transporter running in Amazon EC2 instance.
- Amazon S3  
Highly scalable AWS object storage.
- Wasabi  
Cost effective cloud based object storage.
- SaaS  
Directly attached storage or locally mounted share on an assigned transporter used for storing Microsoft 365 objects.
- Deduplication Appliance  
Integration with Dell EMC Data Domain Boost, HPE StoreOnce Catalyst and NEC HYDRAsTOR.

## 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. 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.

Repositories / Create Backup Repository

1. Type      2. Name & Location      3. Options

Name:

Assigned transporter:  ?

Account:  ?

AWS Region:  ?

Bucket:  ?

## 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 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). 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 Wasabi Hot Cloud Storage

Select the **Wasabi** option if you want to create a Backup Repository in Wasabi. Before creating a repository, grant the required Wasabi access permissions to NAKIVO Backup & Replication. For details, refer to [Required AWS IAM Permissions for Amazon S3 and 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.
- Wasabi object lock and bucket versioning are not supported by NAKIVO Backup & Replication.

To create a Backup Repository in a Wasabi bucket, 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 **Wasabi** and click **Next** to go to the next page of the wizard.

Repositories / Create Backup Repository

1. Type      2. Name & Location      3. Options

- Local Folder  
Directly attached storage or locally mounted share on an assigned transporter.
- CIFS Share  
Network share used by transporter through CIFS Interface.
- NFS Share  
Network share used by transporter through NFS Interface.
- Amazon EC2  
EBS storage attached to transporter running in Amazon EC2 instance.
- Amazon S3  
Highly scalable AWS object storage.
- Wasabi  
Cost effective cloud based object storage.
- SaaS  
Directly attached storage or locally mounted share on an assigned transporter used for storing Microsoft 365 objects.
- Deduplication Appliance  
Integration with Dell EMC Data Domain Boost, HPE StoreOnce Catalyst and NEC HYDRAsTOR.

## 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. 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.

Repositories / Create Backup Repository

1. Type      2. Name & Location      3. Options

Name:

Assigned transporter:  ?

Account:  ?

Wasabi region:  ?

Bucket:  ?

## 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 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.

The screenshot shows the 'Create Backup Repository' wizard in the 'Options' step. The wizard has three steps: 1. Type, 2. Name & Location, and 3. Options. The 'Options' step is currently active. It contains the following settings:

- Storage Savings & Encryption:** Compression is set to 'Fast'.
- Reliability & Maintenance:**
  - Run full data verification on schedule
  - Enforce explicit file system sync
- Scheduled Detach:**
  - Detach this repository on schedule

At the bottom right, there are two buttons: 'Finish' (highlighted in blue) and 'Cancel'.

## 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.

Repositories /

## Create Backup Repository

1. Type      2. Name & Location      3. Options

Local Folder  
Directly attached storage or locally mounted share on an assigned transporter.

CIFS Share  
Network share used by transporter through CIFS Interface.

NFS Share  
Network share used by transporter through NFS Interface.

Amazon EC2  
EBS storage attached to transporter running in Amazon EC2 instance.

Amazon S3  
Highly scalable AWS object storage.

Wasabi  
Cost effective cloud based object storage.

SaaS  
Directly attached storage or locally mounted share on an assigned transporter used for storing Microsoft 365 objects.

Deduplication Appliance  
Integration with Dell EMC Data Domain Boost, HPE StoreOnce Catalyst and NEC HYDRAsTOR.

## 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.

The screenshot shows the 'Create Backup Repository' wizard in the 'Name & Location' step. The form includes the following fields and options:

- Name: CIFS
- Assigned transporter: Onboard transporter
- Path to the share: \\10.30.30.61\ayunt\_cifs1
- Username: administrator
- Password: masked with asterisks
- Advanced mount options: unchecked

At the bottom right, there are 'Next' and 'Cancel' buttons. The 'Next' button is highlighted in blue.

## 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.

- **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

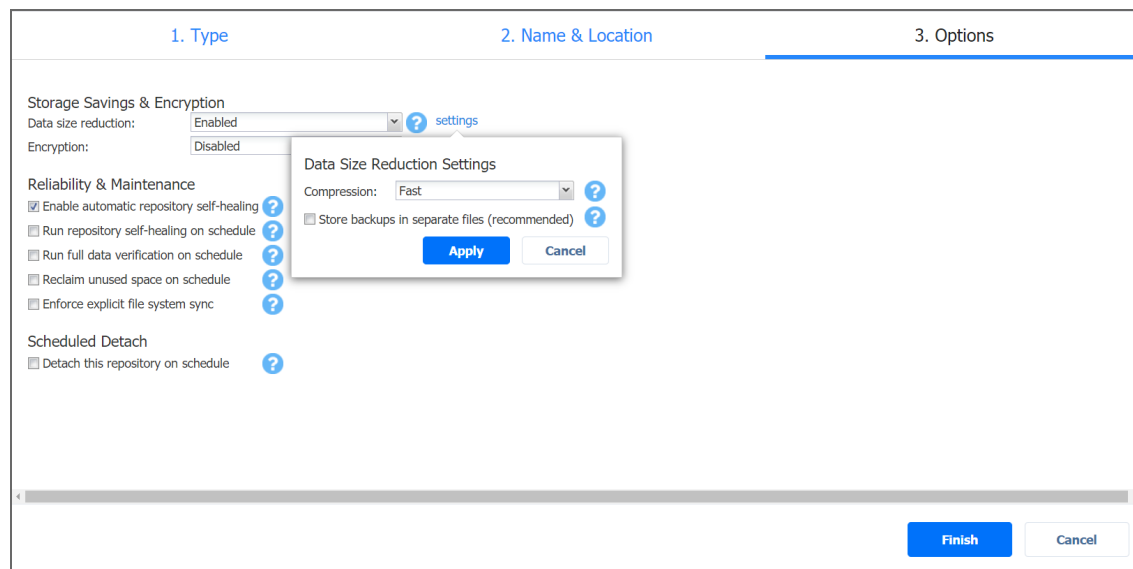
Do not reboot/disconnect the "null" Transporter and storage device while space reclaim is in progress to avoid Backup Repository corruption.

- **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 the Backup Repository data and metadata in a consistent state and then stops the product's interaction with the Backup Repository (so that the Backup Repository 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 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 HYDRAsstor, you need to configure the NEC HYDRAsstor and the machine on which NAKIVO Transporter is installed. Refer to [Integrating with NEC HYDRAsstor](#) for details.
- To create a Backup Repository on other deduplication appliances, refer to [“Backup Repository on NFS Share” on page 413](#).

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 102](#) to see the list of supported advanced deduplication appliances.

Repositories / Create Backup Repository

1. Type      2. Device      3. Name & Location      4. Options

Local Folder  
Directly attached storage or locally mounted share on an assigned transporter.

CIFS Share  
Network share used by transporter through CIFS Interface.

NFS Share  
Network share used by transporter through NFS Interface.

Amazon EC2  
EBS storage attached to transporter running in Amazon EC2 instance.

Amazon S3  
Highly scalable AWS object storage.

Wasabi  
Cost effective cloud based object storage.

SaaS  
Directly attached storage or locally mounted share on an assigned transporter used for storing Microsoft 365 objects.

**Deduplication Appliance**  
Integration with Dell EMC Data Domain Boost, HPE StoreOnce Catalyst and NEC HYDRAsstor.

**Next**      Cancel

## Create Backup Repository: Device

- On the **Device** page, select one of the devices:
  - Dell EMC Data Domain Boost**
  - HP StoreOnce Catalyst**
  - NEC HYDRAsstor**
- Click **Next** to go to the next page of the wizard.

Repositories / Create Backup Repository

1. Type      **2. Device**      3. Name & Location      4. Options

Dell EMC Data Domain Boost  
Use NAKIVO Backup & Replication along with the source-side deduplication of Dell/EMC Data Domain Boost.

HPE StoreOnce Catalyst  
Use NAKIVO Backup & Replication along with the source-side deduplication of HPE StoreOnce Catalyst.

NEC HYDRAsstor  
Use NAKIVO Backup & Replication along with the deduplication of NEC HYDRAsstor.

**Next**      Cancel

## 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.

Repositories / Create Backup Repository

1. Type      2. Device      3. Name & Location      4. Options

Dell EMC Data Domain Boost  
Use NAKIVO Backup & Replication along with the source-side deduplication of Dell/EMC Data Domain Boost.

HPE StoreOnce Catalyst  
Use NAKIVO Backup & Replication along with the source-side deduplication of HPE StoreOnce Catalyst.

NEC HYDRAsstor  
Use NAKIVO Backup & Replication along with the deduplication of NEC HYDRAsstor.

Next      Cancel

- **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.

Repositories / Create Backup Repository

1. Type      2. Device      **3. Name & Location**      4. Options

Name:

Assigned transporter:  ?

Connection type:  ?

Server name:  ?

Catalyst store name:  ?

Username:

Password:

• **NEC HYDRAsstor**

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`.

Repositories / Create Backup Repository

1. Type      2. Device      **3. Name & Location**      4. Options

Name:

Assigned transporter:  ?

Path to the mount point:  ?

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.
- **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**

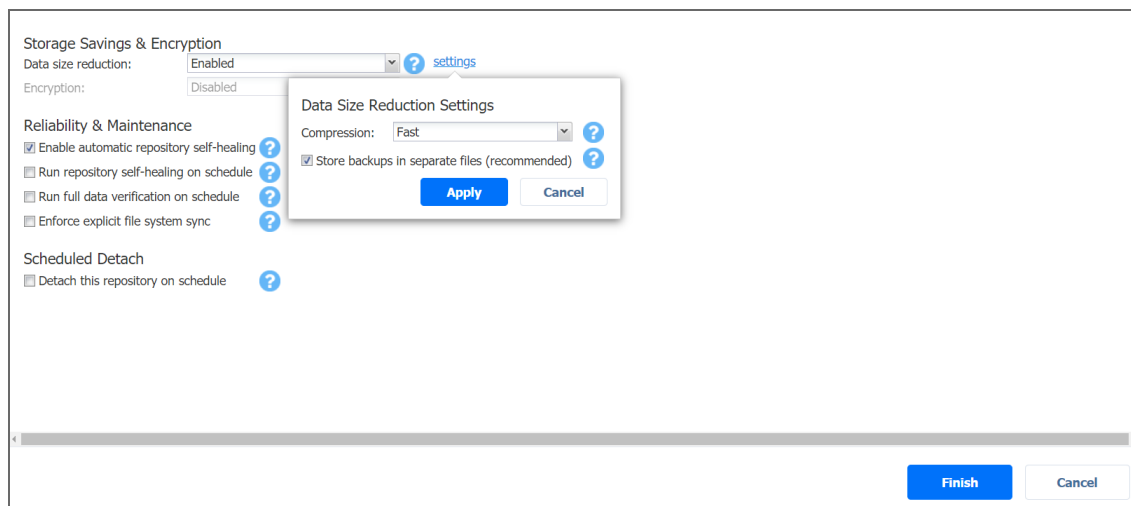
Backup verification is a time-consuming process and consumes CPU of the Transporter assigned to the Backup Repository. It is recommended that you schedule backup verification during non-working hours.

- **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 the Backup Repository data and metadata in a consistent state and then stops the product's interaction with the Backup Repository (so that the Backup Repository 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.

Repositories /

## Create Backup Repository

1. Type      2. Name & Location      3. Options

- Local Folder  
Directly attached storage or locally mounted share on an assigned transporter.
- CIFS Share  
Network share used by transporter through CIFS Interface.
- NFS Share  
Network share used by transporter through NFS Interface.
- Amazon EC2  
EBS storage attached to transporter running in Amazon EC2 instance.
- Amazon S3  
Highly scalable AWS object storage.
- Wasabi  
Cost effective cloud based object storage.
- SaaS  
Directly attached storage or locally mounted share on an assigned transporter used for storing Microsoft 365 objects.
- Deduplication Appliance  
Integration with Dell EMC Data Domain Boost, HPE StoreOnce Catalyst and NEC HYDRAStor.

## 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

FreeNAS share path: 192.168.3.2:/mnt/NFS\_dataset/nfs01

### 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.

Repositories / Create Backup Repository

1. Type      2. Name & Location      3. Options

Name:

Assigned transporter:  ?

Path to the share:  ?

Advanced mount options:  ?

## 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.
- **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 `encryptfs` 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

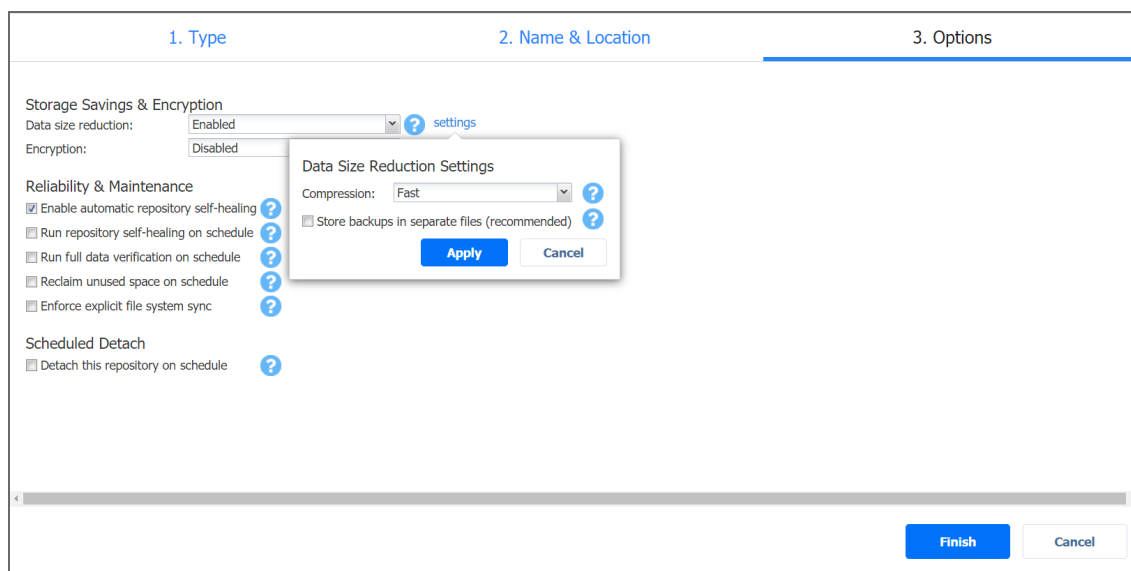
This option is available only if **Store backups in separate files** is not enabled.

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 the Backup Repository data and metadata in a consistent state and then stops the product's interaction with the Backup Repository (so that the Backup Repository 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.





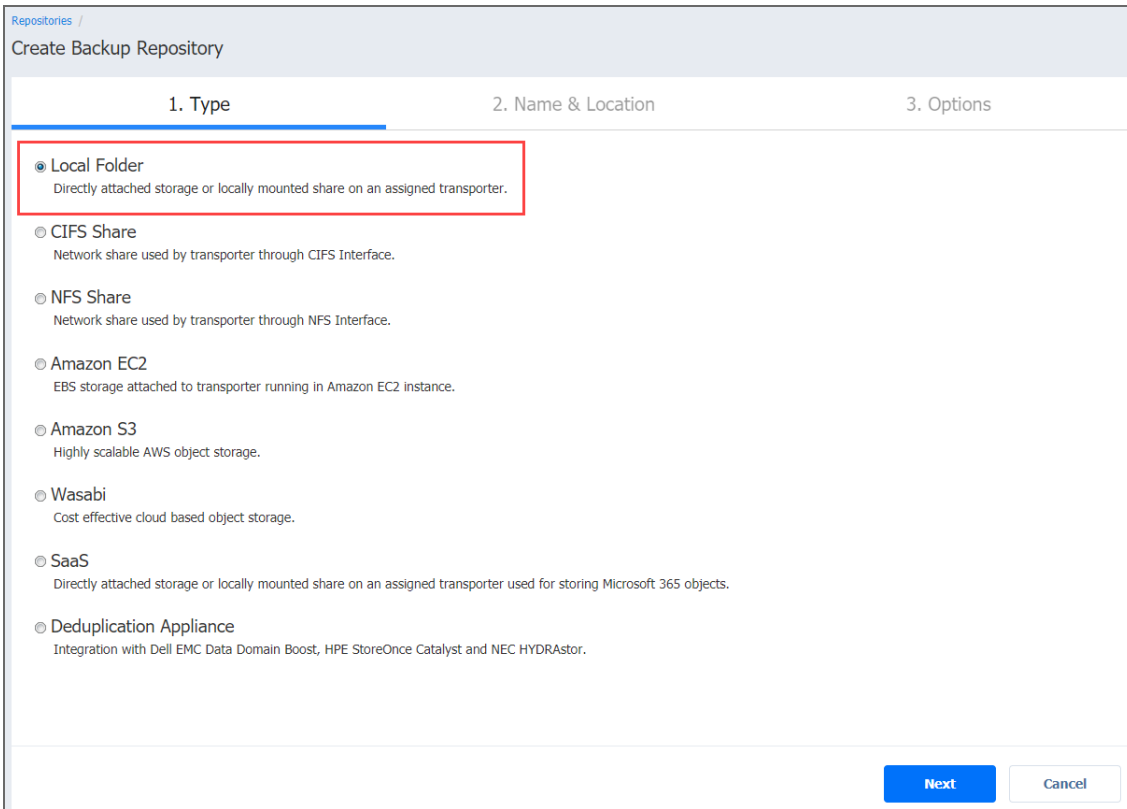
# Local Backup Repository

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.

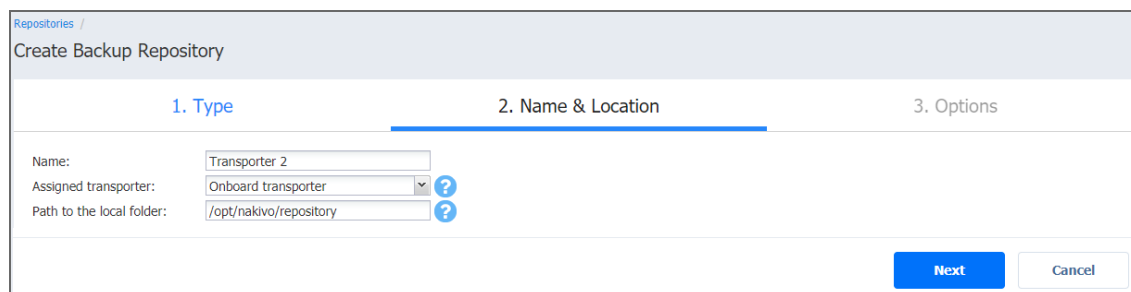
### Example

`/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.



The screenshot shows a web-based wizard for creating a backup repository. The title is 'Create Backup Repository'. It has three steps: 1. Type, 2. Name & Location (current step), and 3. Options. The 'Name' field contains 'Transporter 2'. The 'Assigned transporter' dropdown menu is set to 'Onboard transporter'. The 'Path to the local folder' field contains '/opt/nakivo/repository'. There are question mark icons next to the dropdown menus. At the bottom right, there are 'Next' and 'Cancel' buttons.

## 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.

- **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

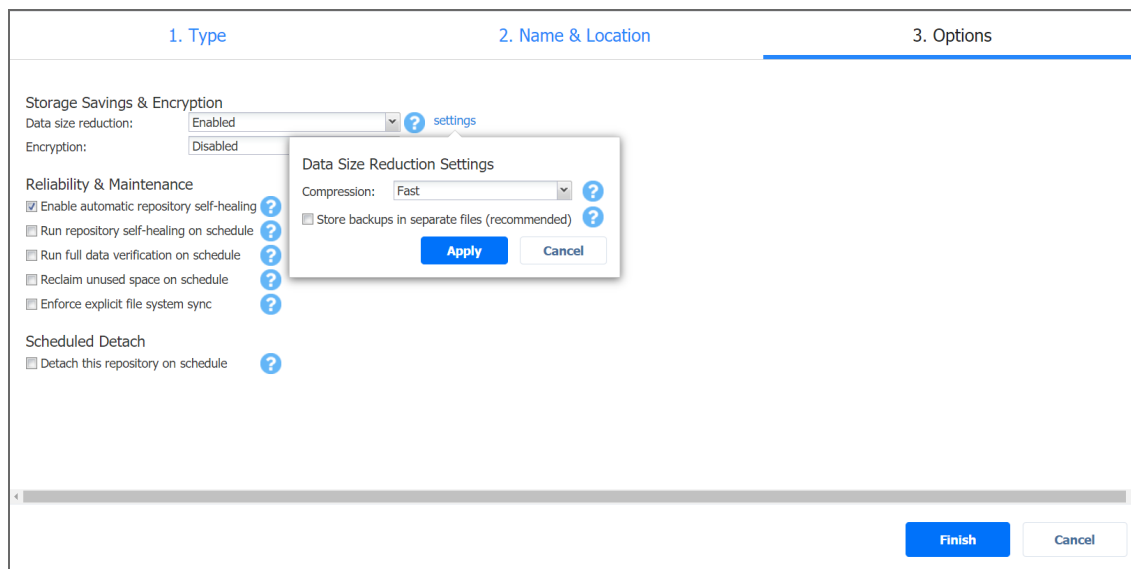
Do not reboot/disconnect the "null" Transporter and storage device while space reclaim is in progress to avoid Backup Repository corruption.

- **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 the Backup Repository data and metadata in a consistent state and then stops the product's interaction with the Backup Repository (so that the Backup Repository 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.



## SaaS Backup Repository

Choose this type of Backup Repository for all your Microsoft 365 related activities.

### Note

- For SaaS repositories, only local folders are supported as a location.
- Before creating a SaaS repository, provide read and write permissions to the local folder where the repository will be located.
- For Office 365 SaaS repositories, manual removal of backup data may not return space to the operating system correctly.
- Before creating SaaS Backup Repository on NAS device, sudo package must be installed on the device beforehand.

To create a SaaS Backup Repository, 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 **SaaS** and click **Next** to go to the next page of the wizard.

Repositories / Create Backup Repository

1. Type      2. Name & Location      3. Options

- Local Folder  
Directly attached storage or locally mounted share on an assigned transporter.
- CIFS Share  
Network share used by transporter through CIFS Interface.
- NFS Share  
Network share used by transporter through NFS Interface.
- Amazon EC2  
EBS storage attached to transporter running in Amazon EC2 instance.
- Amazon S3  
Highly scalable AWS object storage.
- Wasabi  
Cost effective cloud based object storage.
- SaaS  
Directly attached storage or locally mounted share on an assigned transporter used for storing Microsoft 365 objects.
- Deduplication Appliance  
Integration with Dell EMC Data Domain Boost, HPE StoreOnce Catalyst and NEC HYDRAsTOR.

## 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 a path to the local folder in the corresponding box.
4. Click **Next** to go to the next page of the wizard.

Repositories / Create Backup Repository

1. Type      2. Name & Location      3. Options

Name:

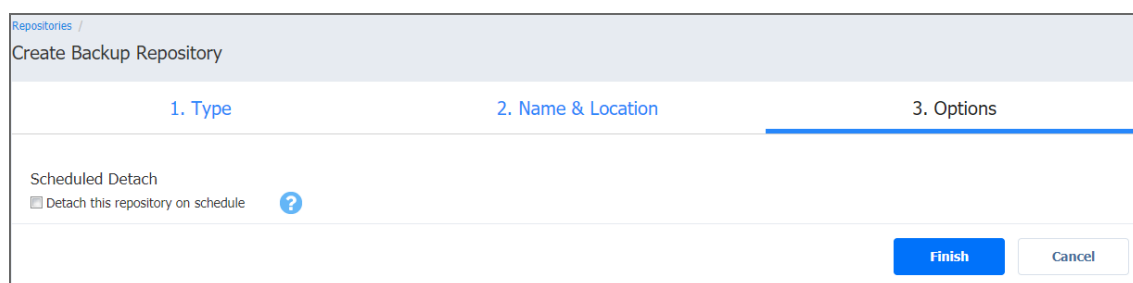
Assigned transporter:  ?

Path to the local folder:  ?

## Create Backup Repository: Options

On the **Options** page:

1. 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 the Backup Repository data and metadata in a consistent state and then stops the product's interaction with the Backup Repository (so that the Backup Repository 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.
2. Click **Finish** to finish creating the Backup Repository.



Repositories / Create Backup Repository

1. Type      2. Name & Location      3. Options

Scheduled Detach

Detach this repository on schedule ?

Finish      Cancel

## Managing Backup Repositories

Refer to the following topics:

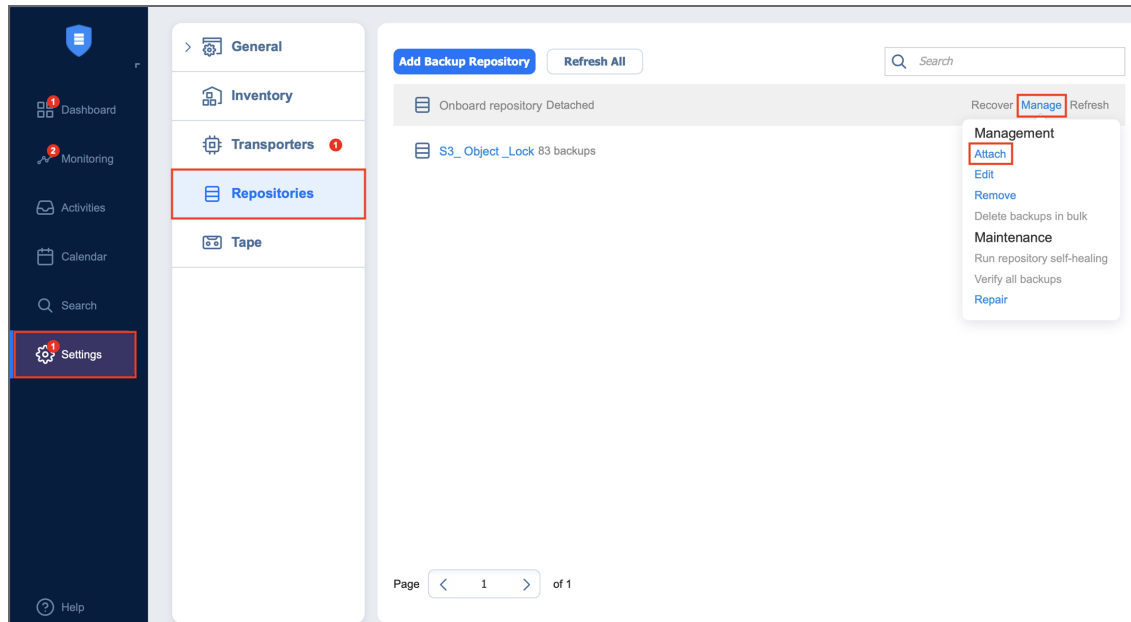
- [“Attaching Backup Repositories” below](#)
- [“Detaching Backup Repositories” on the next page](#)
- [“Editing Backup Repositories” on page 426](#)
- [“How to Copy Backup Repository to Tape” on page 427](#)
- [“Reclaiming Backup Repository Space” on page 428](#)
- [“Refreshing Backup Repositories” on page 429](#)
- [“Removing and Deleting Backup Repositories” on page 431](#)
- [“Repairing Backup Repository” on page 432](#)
- [“Running Backup Repository Self-Healing” on page 435](#)
- [“Running Block-Level Backup Verification” on page 437](#)

## 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 **Manage** 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 the Backup Repository 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 **Manage** and then click **Detach**.





## Note

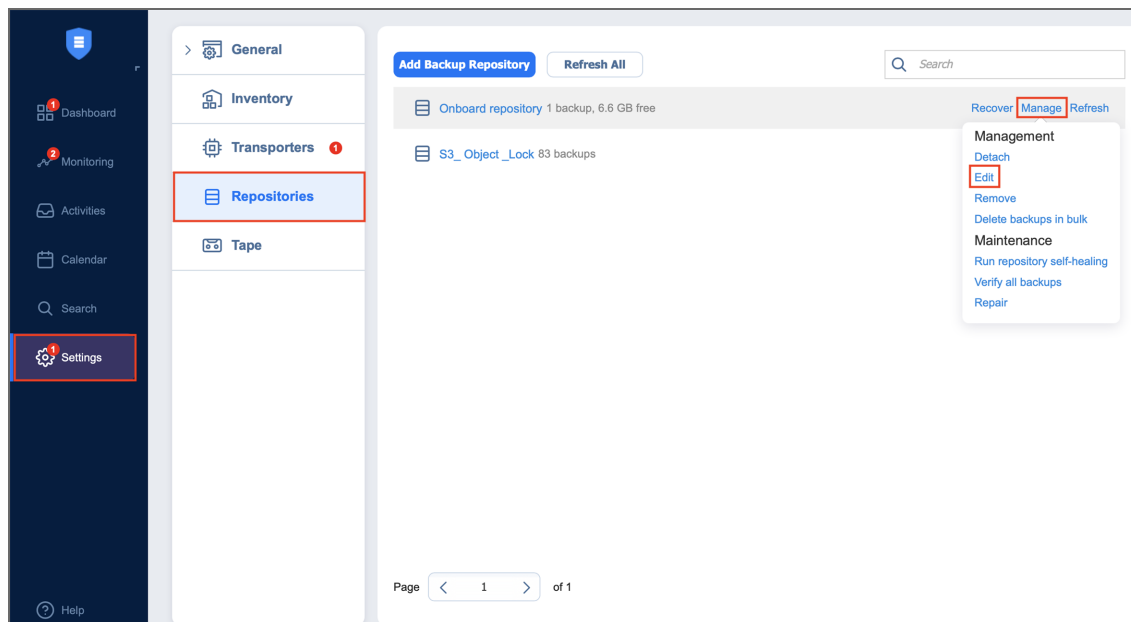
A Backup Repository cannot be detached if a job that backs up to this Backup Repository is running.

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 **Manage** 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](#)

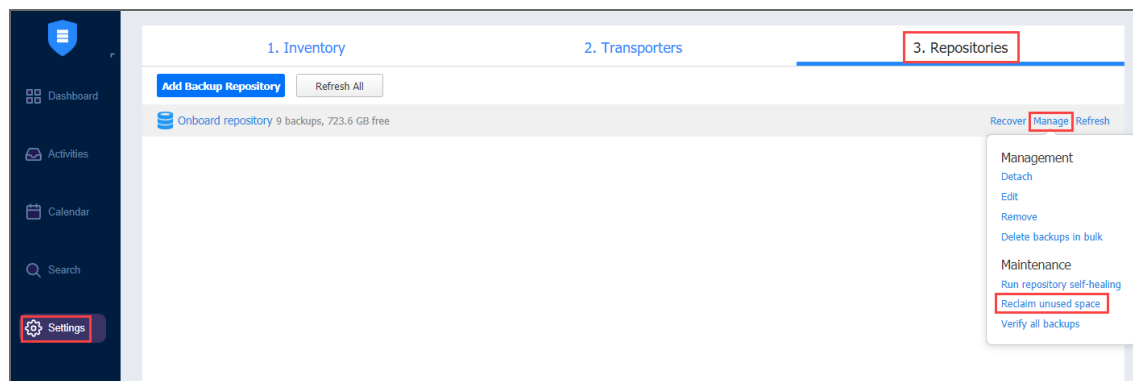
## Starting the Space Reclaiming Process

### 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 **Manage** 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.

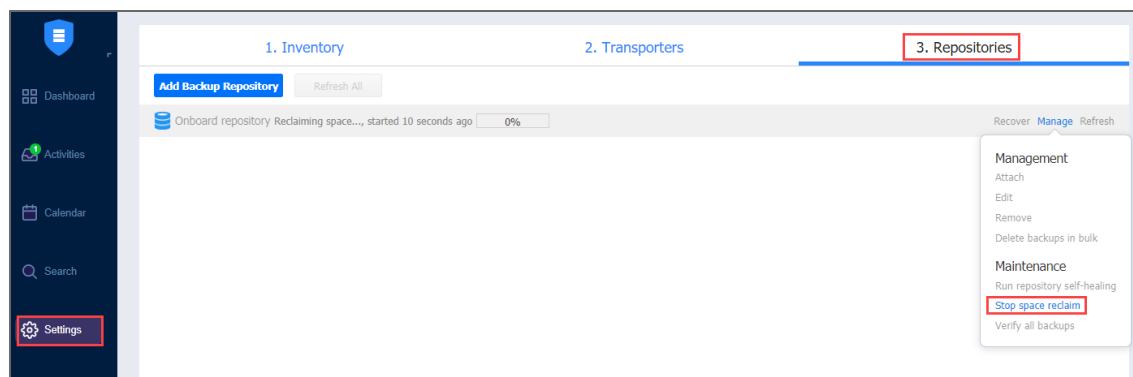
## Stopping the Space Reclaiming Process

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 **Manage** 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.

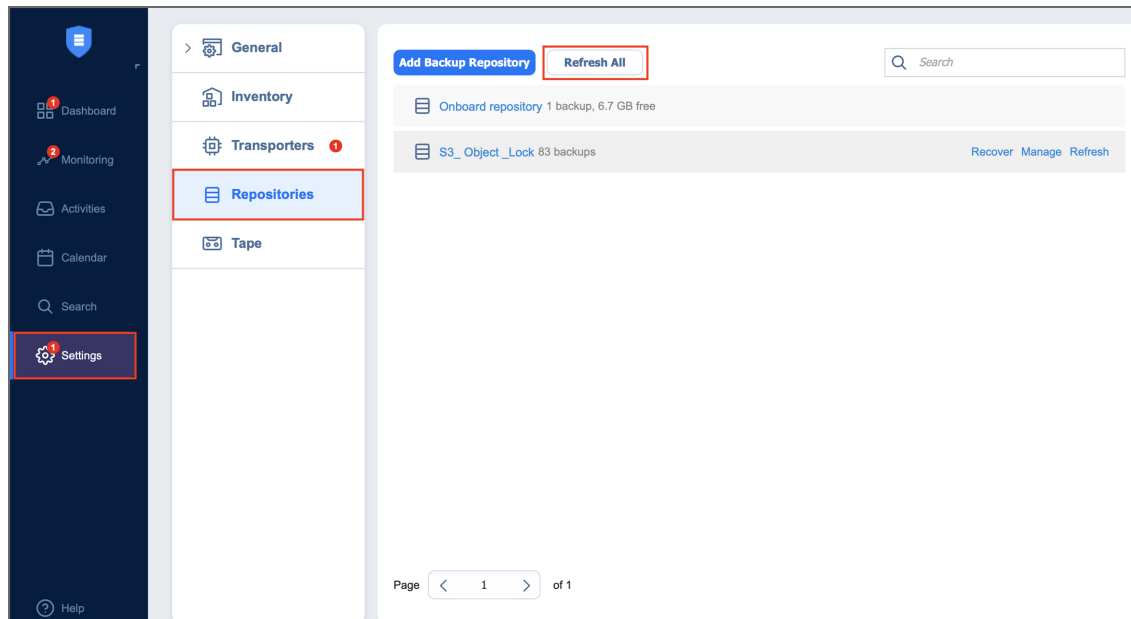
Only one Backup Repository is refreshed at a time. Therefore, if you attempt to refresh multiple Backup Repositories, all but one will be added to a queue.

- [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 **Refresh All**.



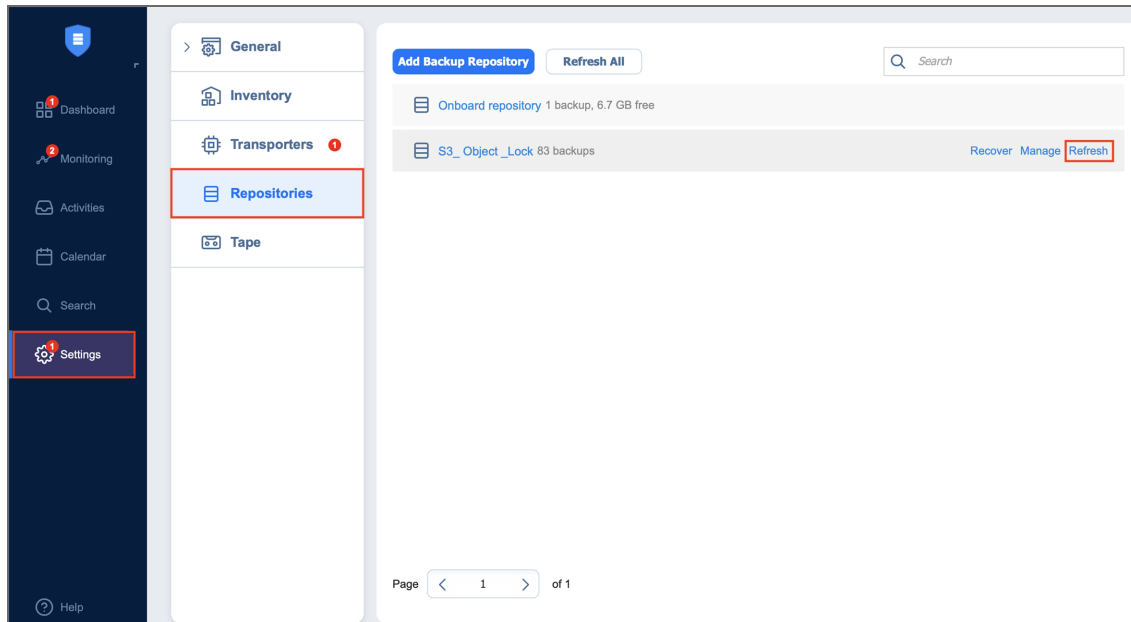
The Backup Repository refresh process begins.

## 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.

4. On the right side, 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.
- Deleting Backup Repository located in Amazon S3 automatically deletes all data from S3 bucket.

## 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 **Manage** 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** Backup 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.

### Repair Repository

Please select repair options for the ransomware-proof local repository.  
If no options are selected, corrupted repository metadata will be overwritten.

Overwrite repository metadata i

Overwrite backup objects i

Verify backup objects i

[Learn more](#) Repair

#### 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.

### Repair Repository

The repository has become corrupted as the result of space-reclaim interruption. Click Repair to try to repair the repository to an uncorrupted state. New jobs will not start while the repair is running.

**Do not reboot/disconnect the "10.30.31.32" transporter and storage device while the repair is in progress.**

[Learn more](#) Repair

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**.

Name	Job	Size	
24	Nutanix AHV backup job	5.7 GB	
AD-Exchange2019_ping1	VMware backup job	Inaccessible	
AD-Exchange2019_ping1	Backup copy job 123	Inaccessible	
All2016	VMware backup job	20.6 GB	Recover Verify <b>Repair</b> Delete

[Close](#)

### 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.

### Repair Backup Object

Clicking Repair will overwrite the backup object and repair only the immutable data.

Verify backup object i

[Learn more](#)

[Repair](#)

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.

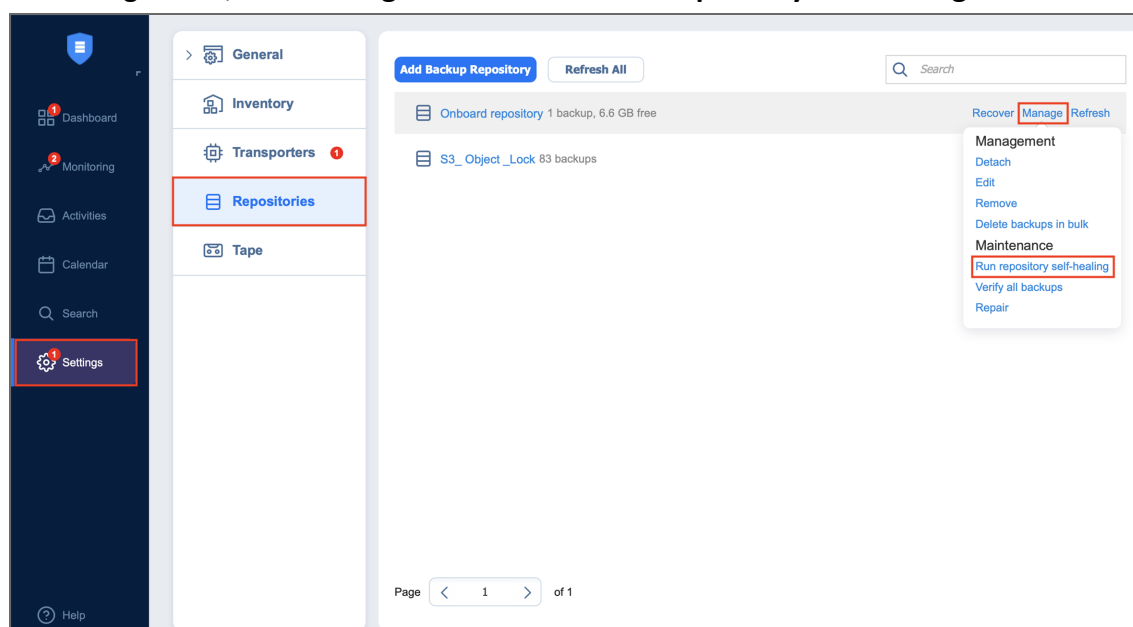
Refer to the following topics to learn more:

- [“Starting the Self-Healing Process” below](#)
- [“Stopping the Self-Healing Process” below](#)

## 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 **Manage** and then click **Run repository self-healing**.



4. In the dialog box that appears, click **Start**. The self-healing process begins.

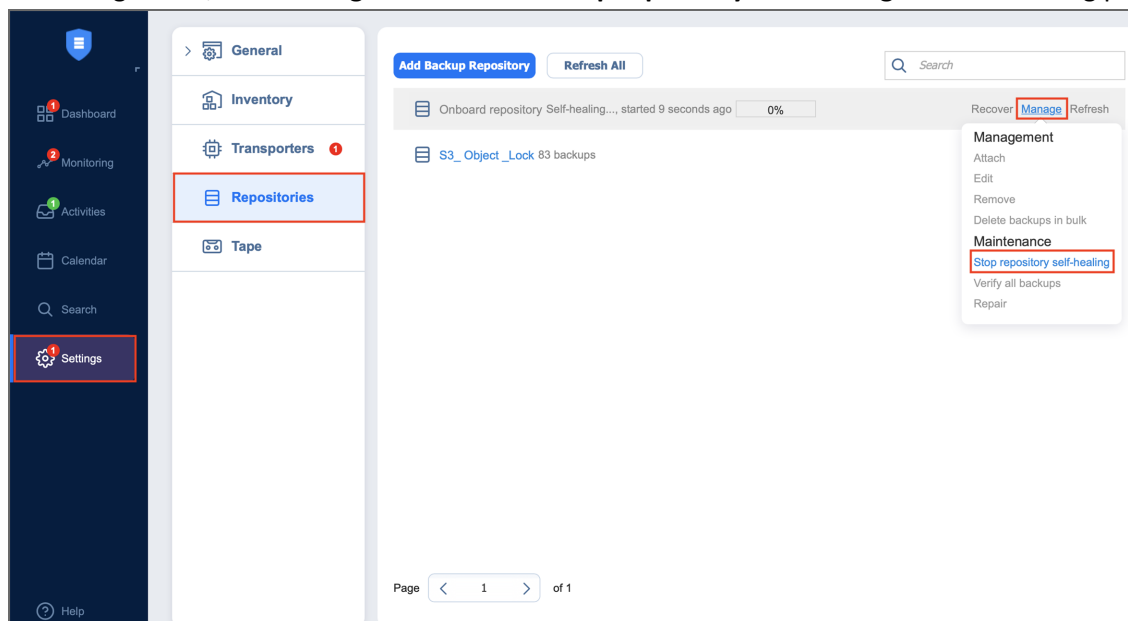
## Stopping the Self-Healing Process

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” below](#)
  - [“Verifying All VM Backups” on the next page](#)
  - [“Verifying a Single Backup” on the next page](#)
- [“Stopping the Backup Verification Process” on page 439](#)
  - [“Stopping Backup Verification for a Backup Repository” on page 439](#)
  - [“Stopping Backup Verification for a Single Backup” on page 439](#)

## Verifying Backups

### Note

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 will fail while the backup verification process is in progress.

## Verifying All VM Backups

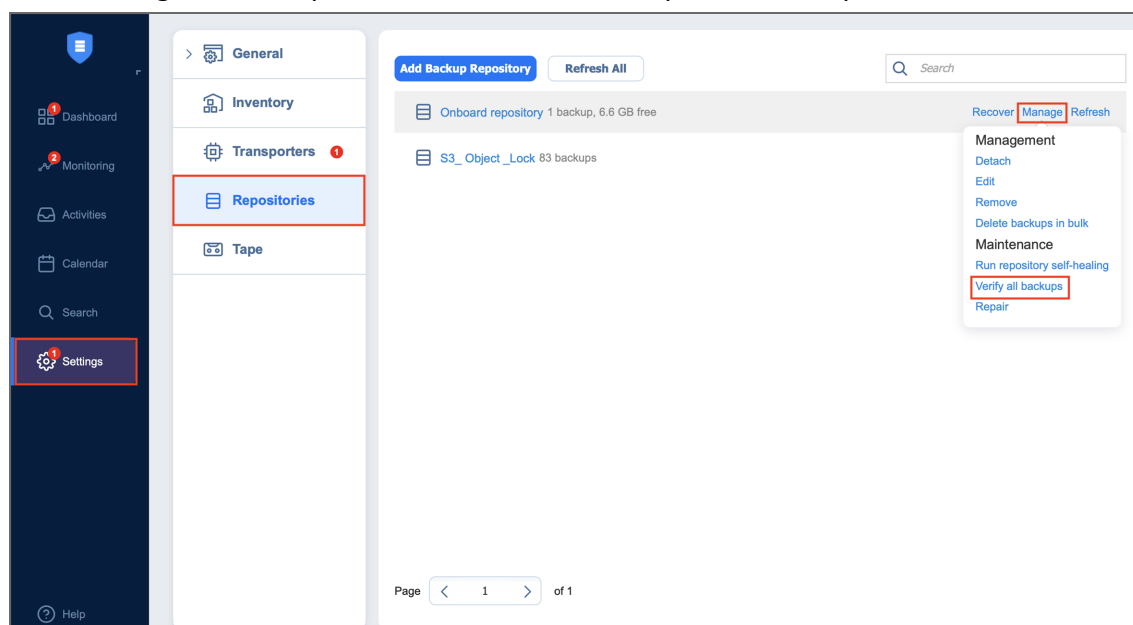
To verify all VM backups in a repository, follow the steps below:

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.

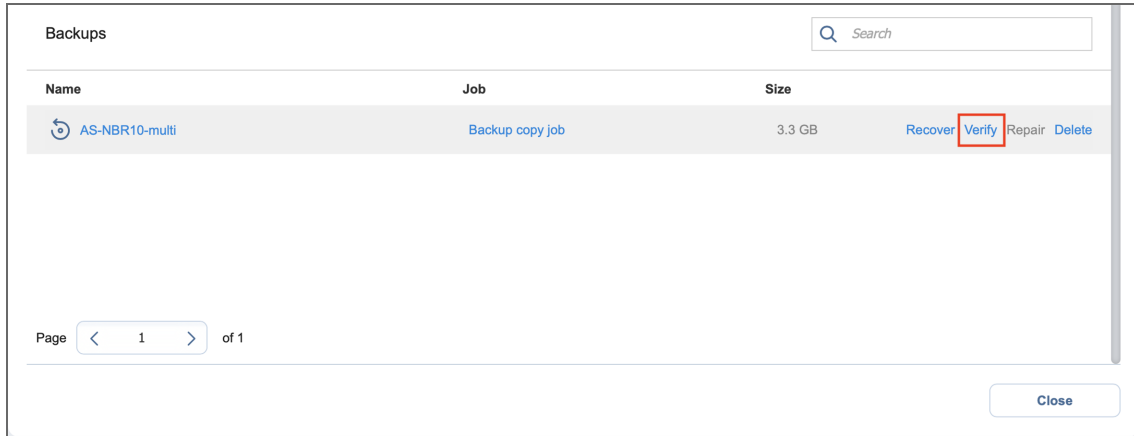


## Verifying a Single Backup

To verify a single VM 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 and click the **Verify** button on the right side.



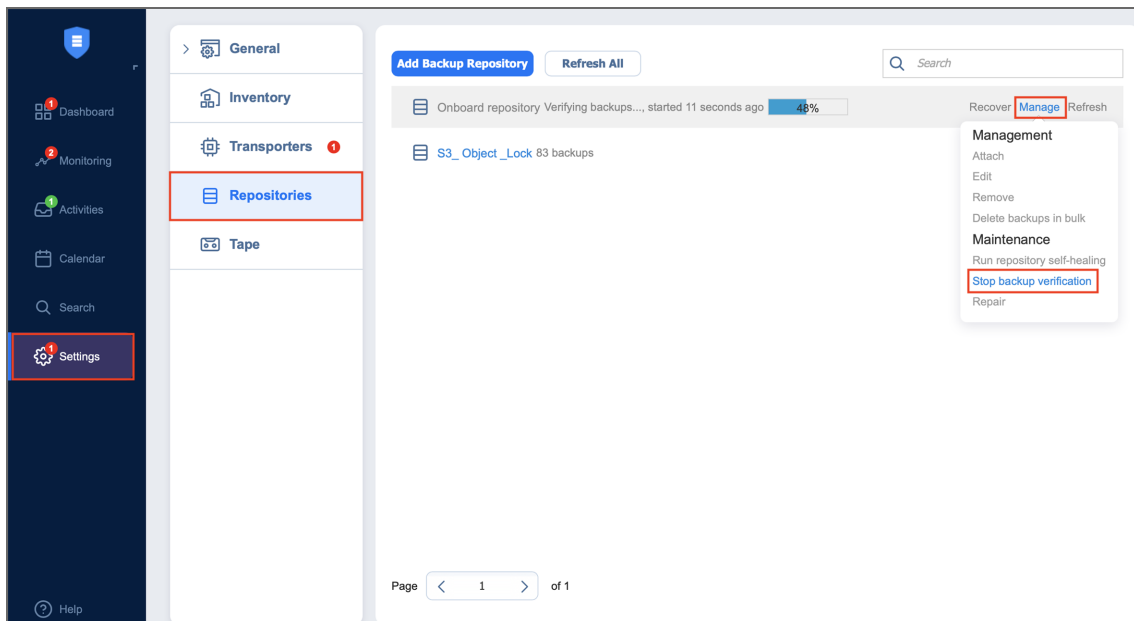
### Stopping the Backup Verification Process

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:

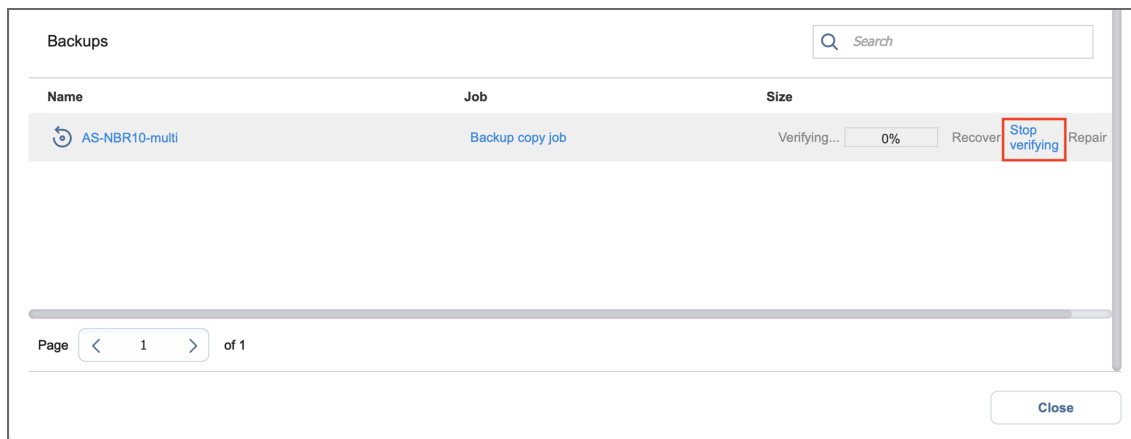
1. From the main menu of NAKIVO Backup & Replication, click **Settings**.
2. Go to the **Repositories** tab and hover over a Backup Repository.
3. On the right side, click **Manage** and then click **Stop backup verification**. The backup verification process is stopped.



#### 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 and click **Stop verifying** on the right side.



For near-instant backup verification, refer to the [“VM Verification” on page 48](#) feature.

## Viewing Backup Repository Details

To view Backup Repository 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. Click a Backup Repository.

4. The following data is displayed:

The screenshot displays the 'Local repo' management interface. At the top, it shows 'Local repo' with a hamburger menu icon, '1 backup, 5.0 GB free', and three buttons: 'Recover' (highlighted in blue), 'Manage', and 'Refresh'. Below the buttons is a vertical bar chart showing 5.0 GB free space (white) and 10.1 GB used space (blue). To the right of the chart is a list of configuration options:

Free:	5.0 GB
Used:	10.1 GB
Deduplication:	Disabled
Compression:	Fast
Encryption:	Disabled
Space savings:	49.0% (9.7 GB) ⓘ
Automatic self-healing:	Enabled
Scheduled self-healing:	Disabled
Enforce explicit file system sync:	Disabled
Scheduled data verification:	Disabled
Scheduled detach:	Disabled
Store backups in separate files:	Enabled
Type:	Local folder on assigned transporter
Path to the local folder:	/home/repo
Assigned transporter:	Onboard transporter

- **Free:** The amount of free space currently available for the Backup Repository.
- **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 current Backup Repository.
- **Compression:** The compression level specified for the current Backup Repository.
- **Encryption:** The status of encryption in the current 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 current state of the automatic self-healing option for the Backup Repository.
- **Scheduled self-healing:** The current state of the scheduled self-healing option for the Backup Repository.
- **Enforce explicit file system sync:** The current state of the enforce explicit file system sync option for the Backup Repository.
- **Scheduled data verification:** The current state of the scheduled data verification option for the Backup Repository.
- **Scheduled space reclaiming:** The current state of the scheduled space reclaiming option for the Backup Repository.
- **Scheduled detach:** The current state of the scheduled detach option for the Backup Repository.
- **Store backups in separate files:** The current 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**
  - **SaaS**
  - **Amazon S3**
  - **Wasabi**
  - **Deduplication Appliance**
- **Path to the folder:** The path to the Backup Repository folder.
- **Assigned transporter:** The Transporter that manages the Backup Repository (i.e. reads data from and writes data to the Backup Repository).
- **Backups:** List of available backups in the Backup Repository.


## Viewing Backup Details

Here you can view the details of the backups stored in the Backup Repository. The following information is displayed:

- **Name:** Name of the backup.
- **Job:** The job type that created this backup.
- **Size:** The total size of the backup.

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 via the **Search** bar. Hover over the name of a backup to select one of the following options that appear on the right side of the screen:


- **Recover:** Select this option to proceed with recovery.
- **Verify:** Select this option to verify the backup.
- **Repair:** If the backup is corrupted, this option will attempt to restore it to an uncorrupted state.
- **Delete:** Select this option to delete the backup from the repository.

Backups		<input type="text" value="Search"/>	
Name	Job	Size	
 AS-NBR10-multi	Backup copy job	3.3 GB	<a href="#">Recover</a> <a href="#">Verify</a> <a href="#">Repair</a> <a href="#">Delete</a>

Page  of 1

Click on a backup name to view more information about the backup and see the recovery points available. The following information is displayed:

- **Name:** Name of the job.
- **Type:** Type of the job.
- **Points:** Number of recovery points available.
- **Last point:** Date of the latest recovery point.
- **Size:** The total size of the backup.
- **Job name:** Name of the job.

 AS-NBR10-multi 3.3 GB

[Recover](#) [Verify](#) [Repair](#) [Delete](#)

Name:	AS-NBR10-multi
Type:	VMware VM
Points:	6
Last point:	Wed, 25 Aug 2021 at 20:00 (UTC +03:00)
Size:	3.3 GB
Job name:	<a href="#">Backup copy job</a>

## Viewing Recovery Point Details

You can view the details of the 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:** Date of the recovery point.
- **Size:** Size of the recovery point.
- **Type:** Type of the backup used for this recovery point.
- **Immutable until:** If applicable, the date that recovery point immutability expires.
- **Protected until:** The date until which the recovery point remains protected.

- **Description:** Description of the recovery point.

Recovery points Q Search

Date	Size	Type	Immutable until	Protected until	Description
Wed, 25 Aug 2021 at 20:00 ...	62 MB	Incremental	Not applicable	Use job retention	<a href="#">Recover</a> <a href="#">Edit</a> <a href="#">Delete</a>
Tue, 24 Aug 2021 at 20:00 (...)	59 MB	Incremental	Not applicable	Use job retention	
Mon, 23 Aug 2021 at 20:00 ...	59 MB	Incremental	Not applicable	Use job retention	
Sun, 22 Aug 2021 at 20:00 ...	65 MB	Incremental	Not applicable	Use job retention	
Wed, 18 Aug 2021 at 20:00 ...	63 MB	Incremental	Not applicable	Use job retention	
Sun, 15 Aug 2021 at 20:00 ...	3.0 GB	Full	Not applicable	Use job retention	

[Close](#)

### Note


**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.

**Date**, **Type**, and **Description** can also be viewed when selecting recovery points in Recovery Job Wizard. Hover over the name of the recovery point to select one of the following options that appear on the right side of the screen:

- **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 remain protected. The following options are available:
    - **Use job retention:** Choose this option to use the retention settings selected in the job for this recovery point.
    - **Keep forever:** Choose this option to keep this recovery point forever.
    - **Protect until:** Choose this option to protect this recovery point until a specific date. After selecting this option, choose the date in the calendar.
- **Delete:** Select this option to delete the recovery point from the repository.

# Tape

To start working with tape devices in NAKIVO Backup & Replication, you first need to add and configure them on the **Tape** page of the **Settings** dashboard.



The default view of the **Tape** page is set to **Devices**, and once you add your tape devices, you will be able to view and manage them here. Also, by selecting different views from the **View** drop-down list you will be able to work with tapes and backups.

On the **Tape** page, you can perform the following operations:

- [“Adding Robotic Tape Libraries or VTLs” below](#)
- [“Adding Standalone Tape Drives” on page 451](#)
- [“Managing Backups” on page 454](#)
- [“Managing Locations” on page 457](#)
- [“Managing Media Pools” on page 459](#)
- [“Managing Tape Devices” on page 470](#)
- [“Managing Tape Cartridges” on page 461](#)

## 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, 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:

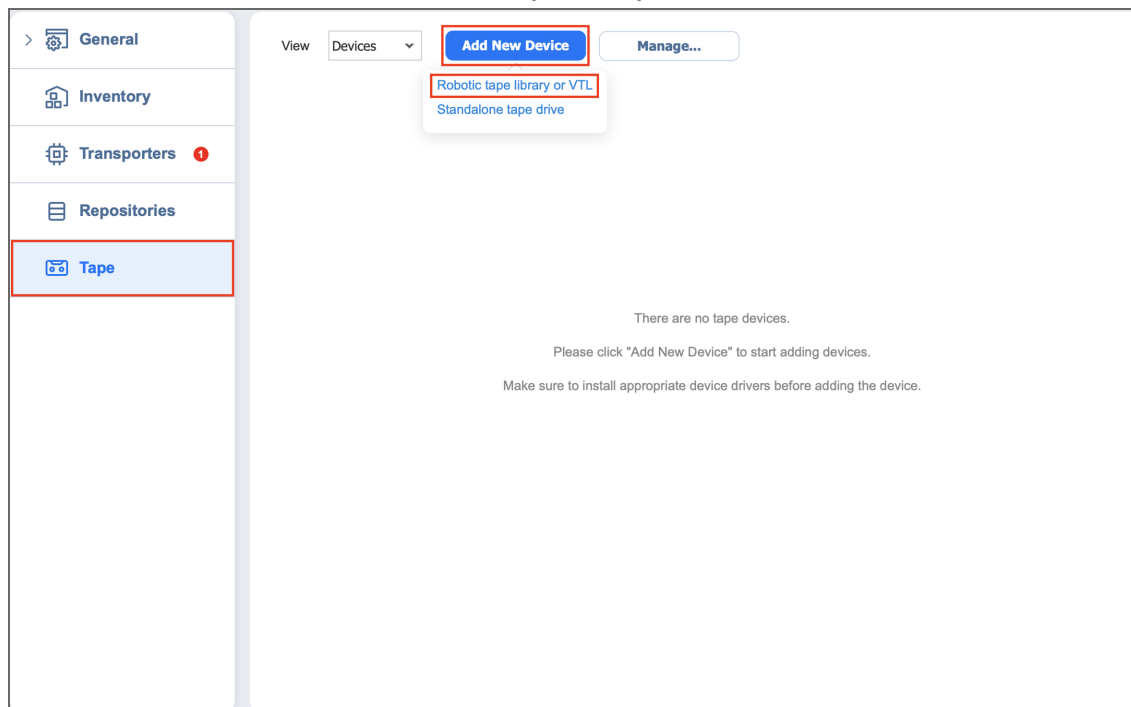
- [Launching Wizard](#)
- [Selecting Transporter](#)
- [Selecting Changers](#)
- [Selecting Drives](#)
- [Selecting Options](#)
- [Managing Added Tape Library](#)

## 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. Select **Devices** from the **View** drop-down list.
3. Click **Add New Device** 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

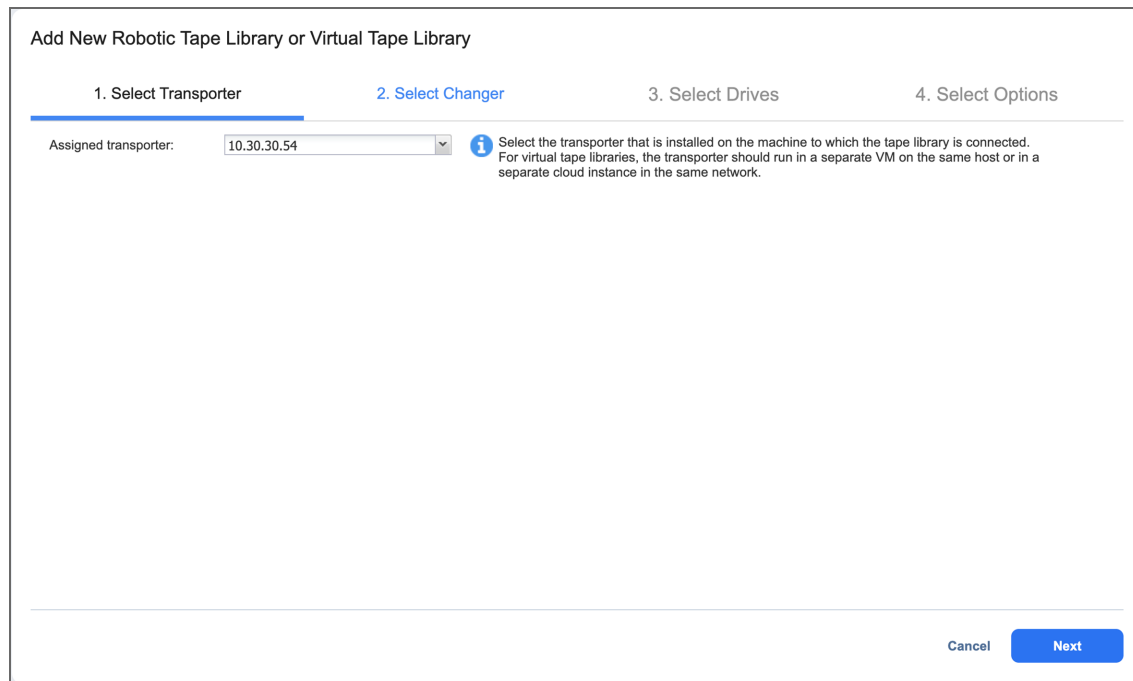
Before adding a new tape device to NAKIVO Backup & Replication, you need to deploy or add an existing Transporter on a machine that is physically connected to the tape device. For virtual tape libraries, the transporter should run in a separate VM on the same host or in a separate cloud instance in the same network.

For more information on Transporter deployment, refer to [“Deploying Transporter as VMware Appliance” on page 373](#) and [“Adding Installed Transporters” on page 365](#).

## 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.



The screenshot shows a wizard window titled "Add New Robotic Tape Library or Virtual Tape Library". It has four steps: "1. Select Transporter", "2. Select Changer", "3. Select Drives", and "4. Select Options". Step 1 is currently active. Under "Assigned transporter:", there is a dropdown menu with the IP address "10.30.30.54" selected. To the right of the dropdown is an information icon and a note: "Select the transporter that is installed on the machine to which the tape library is connected. 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." At the bottom right, there are "Cancel" and "Next" buttons.

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.

Add New Robotic Tape Library or Virtual Tape Library

1. Select Transporter      2. Select Changer      3. Select Drives      4. Select Options

**i** Select the media changer of the tape library below.

Device Name	Address	Path	Serial Number
<input type="radio"/>  DELL PV-132T	[6:0:0:0]	\\.\Changer0	DELL1_3134662N1896
<input type="radio"/>  IBM 3573-TL	[0:0:7:1]	\\.\Changer1	00X2U78H4185_LL0

Cancel      **Next**

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.

Add New Robotic Tape Library or Virtual Tape Library

1. Select Transporter      2. Select Changer      **3. Select Drives**      4. Select Options

**i** Select the tape drives of the tape library and specify the drive number for each drive. [Learn more](#)

<input type="checkbox"/>	Device Name	Address	Path	Serial Number	Drive Number
<input checked="" type="checkbox"/>	IBM ULT3580-HH5	[0:0:7:0]	\\.\Tape1	1068055225	0
<input checked="" type="checkbox"/>	IBM ULT3580-HH5	[0:0:8:0]	\\.\Tape2	1068040089	1
<input type="checkbox"/>	HP C7438A	[5:0:3:0]	\\.\Tape0	0000013891	

Cancel **Next**

Click **Next** to proceed to the next page.

## Selecting Options

The last step of adding a tape library is selecting its options.

Add New Robotic Tape Library or Virtual Tape Library

1. Select Transporter      2. Select Changer      3. Select Drives      **4. Select Options**

Name:

Compression:  **i**

Block size:  **i**

Device location:  **i** [add location](#)

Default media pool:  **i** [add media pool](#)

Default offline location:  **i** [add location](#)

Cancel **Finish**

1. Specify the following parameters:

- **Name:** Enter the name for the tape library
- **Compression:** Select a compression level of the tape device:



- Hardware-based (default)
- Software-based (fast)
- Software-based (medium)
- Software-based (best)

Note, that combining hardware compression with software compression is not recommended.

- **Block size:** Select the block size of the tape device:

- 32 KB
- 64 KB (default)
- 128 KB
- 256 KBs
- 512 KB
- 1 MB

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 457](#)
- **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 459](#)
- **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** view of the **Tape** 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

- **View all:** Clicking the **view all** link opens the **Tapes** screen where you can view and manage tape cartridges in the device.

 **Tape**

[Refresh](#) [Manage...](#)

Assigned transporter:	10.30.30.54
Type:	Standalone tape drive
Device:	HP C7438A
Compression:	Hardware-based
Block size:	64 KB
Device location:	My office
Default offline location:	No data
Default media pool:	Pool
Tape:	<a href="#">939ee6d8-4150-4f5b-bbb7-64b40d02d9d0</a>

## 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

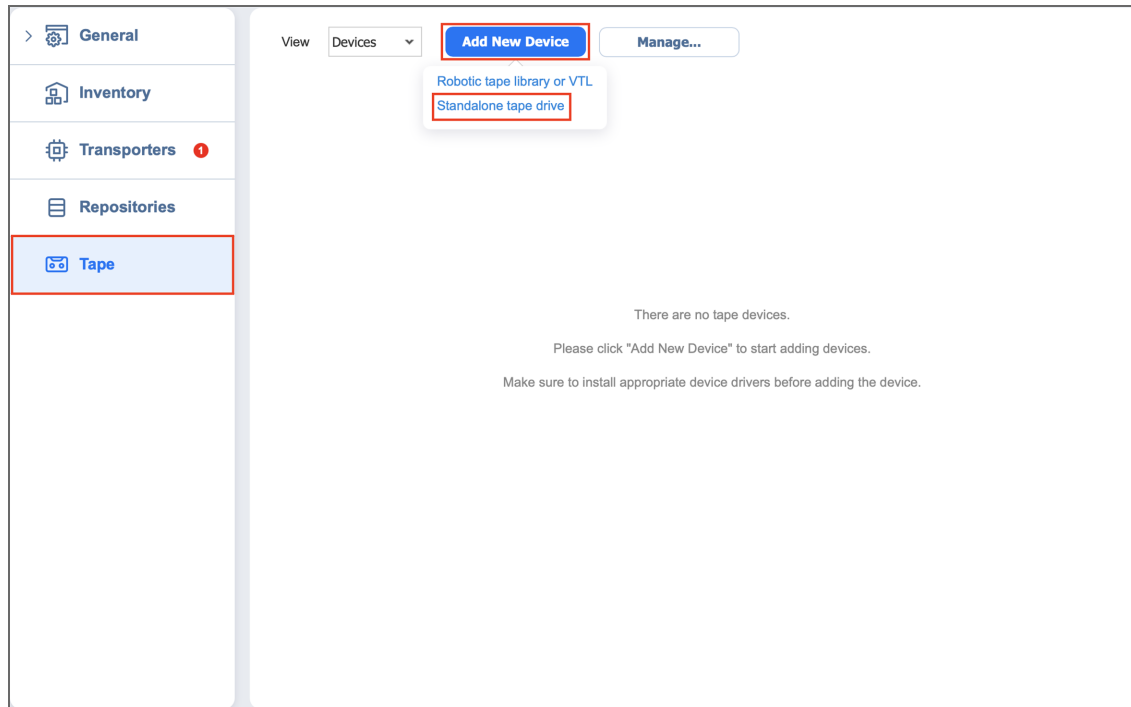
Vendor drivers should be installed on tape devices prior to adding them to NAKIVO Backup & Replication inventory.

## Launching Wizard

To add a standalone tape drive to the system:

1. Go to **Settings** and click the **Tape** tab.
2. Select **Devices** from the **View** drop-down list

3. Click **Add New Device** 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 Transporter deployment, refer to [“Deploying Transporter as VMware Appliance” on page 373](#) and [“Adding Installed Transporters” on page 365](#).

## 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.

The screenshot shows the 'Add New Standalone Tape Drive' wizard. The title bar reads 'Tape / Add New Standalone Tape Drive'. The main area is divided into two steps: '1. Select Transporter' (active) and '2. Select Options'. Under '1. Select Transporter', there is a label 'Assigned transporter:' followed by a dropdown menu showing '10.30.30.86'. To the right of the dropdown is an information icon (i) and a text note: 'Select a transporter installed on the machine to which the standalone tape drive is physically connected.' At the bottom right, there are two buttons: 'Next' (highlighted in blue) and 'Cancel'.

2. Click **Next**.

## Selecting Options

The last step of adding a tape drive is selecting its options.

The screenshot shows the 'Add New Standalone Tape Drive' wizard, Step 2: Select Options. The title bar reads 'Tape / Add New Standalone Tape Drive'. The main area is divided into two steps: '1. Select Transporter' and '2. Select Options' (active). The form contains several fields with dropdown menus and help icons (i):  
- Name: [text input]  
- Drive: HP C7438A V601  
- Compression: Hardware-based  
- Block size: 64 KB  
- Device location: My office (with 'add location' link)  
- Default media pool: <No default media pool> (with 'add media pool' link)  
- Default offline location: <No default offline location> (with 'add location' link)  
At the bottom right, there are two buttons: 'Finish' (highlighted in blue) and 'Cancel'.

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:
  - Hardware-based (default)
  - Software-based (fast)
  - Software-based (medium)
  - Software-based (best)

Note, that combining hardware compression with software compression is not recommended

- **Block size:** Select the block size of the tape device:
  - 32 KB
  - 64 KB (default)
  - 128 KB
  - 256 KB

- 512 KB
- 1 MB

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 457](#)
- **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 459](#)
- **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** view of the **Tape** 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.
- **View tapes:** Clicking the **view all** 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 selecting the **Backups** option from the **View** drop-down list.



From the **Backups** view, you can search for backups, recover from backups, and view backups' details.

- [Searching for Backups](#)
- [Filtering Backups](#)
- [Backups Table](#)
- [Recovering from Backups](#)

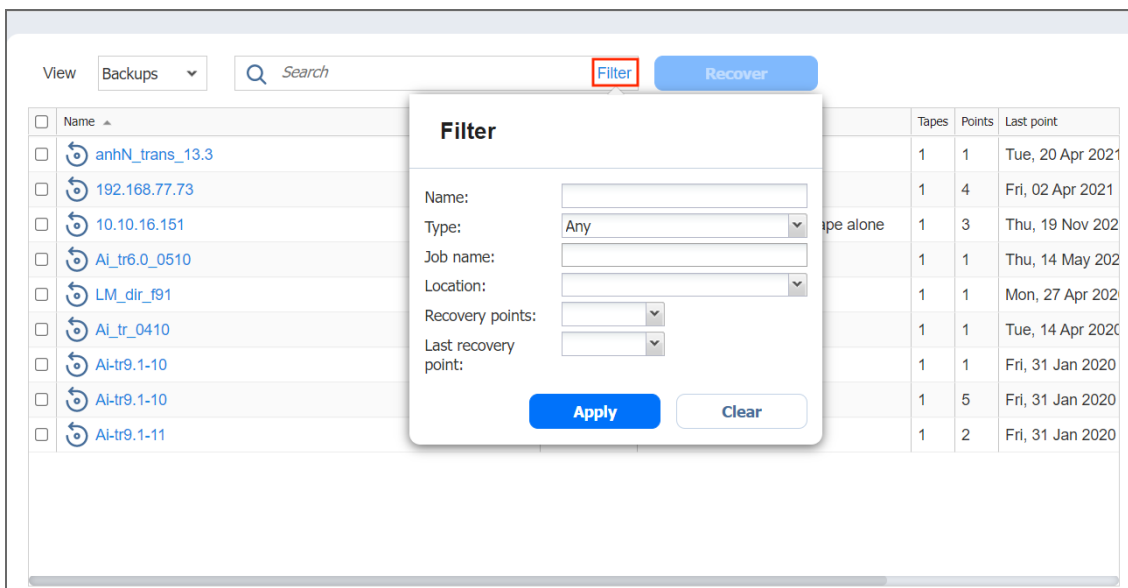
## Searching for Backups

You can search for particular backup(s) by entering its name (or part of it) 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 will clear the query and the table will display all backups.

## Filtering Backups

The Backups view also provides sophisticated filtering options that can be applied to search for particular backups. To access filtering options, click the **Filter** link in the Search box. In the dialog that opens, you can select one or several filtering criteria that will be applied with the AND statement.



You can apply the following filtering criteria:

- **Backup name:** The backups with the name provided will be displayed. Part of the name can be entered.
- **Status:** Specify the type of backups to be displayed:
  - VMware VM
  - Hyper-V VM
  - Amazon EC2 instance
  - Any
- **Location:** Only the backups from the tape cartridges of the specified device location will be displayed.
- **Recovery points:** Only the backups with less or more recovery points will be displayed.
- **Last recovery point:** Only the backups with the last recovery point created on/newer/older than the date specified will be displayed.

### Note

The Search and Filter features can only be applied separately; that is, you cannot simultaneously enter a search string and select filtering options.

## Backups Table

The **Backups** table provides the detailed information about each backup:

- **Name:** Displays the name of the backup. Clicking on the name opens the **Recovery screen**.
- **Type:** Displays the type of backup.
- **Job Name:** Displays the name of the job associated with the backup.
- **Tapes:** Displays how many tape cartridges the backup occupies.
- **Points:** Displays how many recovery points the backup has.
- **Last point:** Displays the date of the last recovery point on the backup.
- **Location:** Displays the location the tape(s) with the backup belongs to.

## Recovering from Backups

You can initiate the recovery process from the **Backups** view by selecting the checkboxes next to the backups' names and clicking the **Recover** button.

View Backups  Filter Recover

<input type="checkbox"/>	Name	Type	Job name	Tapes	Points	Last point
<input type="checkbox"/>	anhN_trans_13.3	VMware	Backup copy job tape	1	1	Tue, 20 Apr 2021
<input checked="" type="checkbox"/>	192.168.77.73	Physical	Backup copy #1	1	4	Fri, 02 Apr 2021
<input checked="" type="checkbox"/>	10.10.16.151	Physical	Backup copy physical to tape alone	1	3	Thu, 19 Nov 2020
<input type="checkbox"/>	AI_tr6.0_0510	VMware	Unknown	1	1	Thu, 14 May 2020
<input type="checkbox"/>	LM_dir_f91	VMware	Unknown	1	1	Mon, 27 Apr 2020
<input type="checkbox"/>	AI_tr_0410	VMware	Unknown	1	1	Tue, 14 Apr 2020
<input type="checkbox"/>	AI-tr9.1-10	VMware	Unknown	1	1	Fri, 31 Jan 2020
<input type="checkbox"/>	AI-tr9.1-10	VMware	Unknown	1	5	Fri, 31 Jan 2020
<input type="checkbox"/>	AI-tr9.1-11	VMware	Unknown	1	2	Fri, 31 Jan 2020

The **New Recovery Job Wizard** opens with the specified backups and their latest recovery points selected.

## 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:

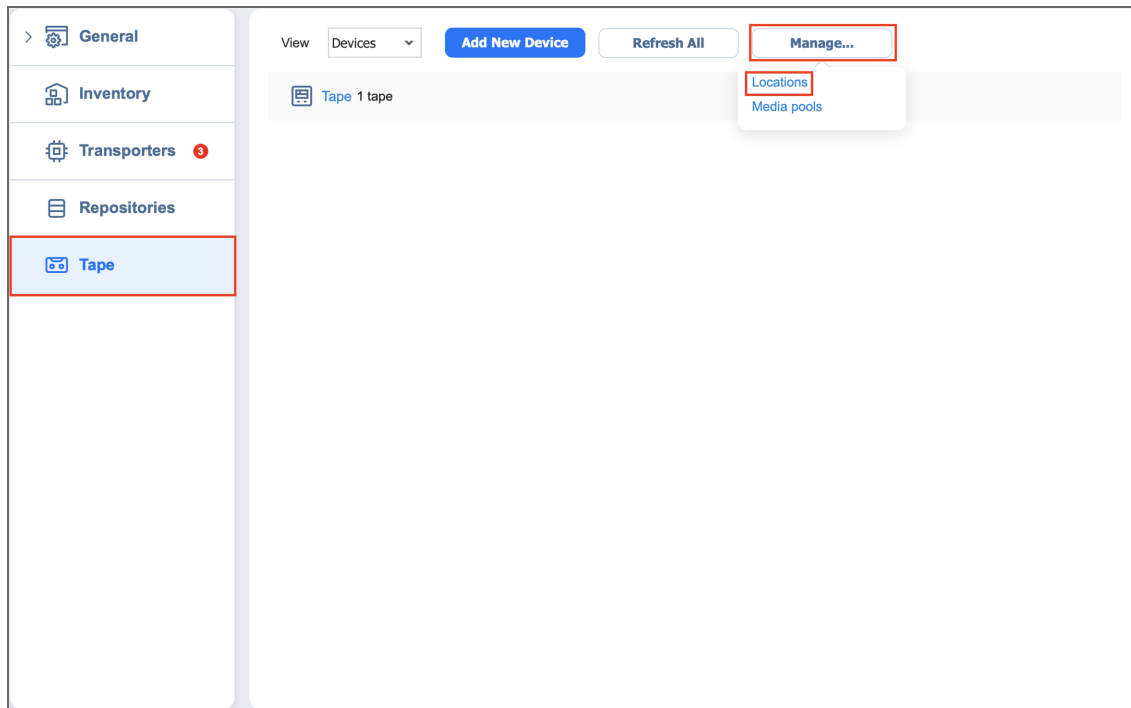
- [Adding Device Locations](#)
- [Managing Device Locations](#)

### Adding Device Locations


To add a Location:

1. Go to **Settings > Tape**.
2. With the **Devices** or **Tapes** view selected, click **Manage** and select **Locations**. The **Location Management** dialog opens.





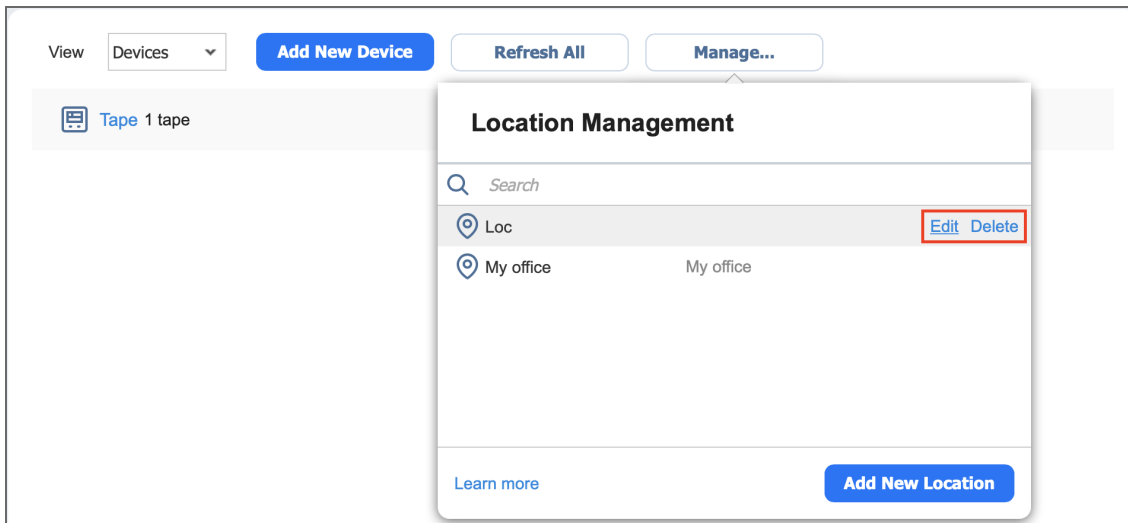
3. Click **Add New Location**.
4. In the **Add New Location** dialog box, specify a name for the device location and provide its description (optionally).



5. Click **Save**. 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.



## 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](#)

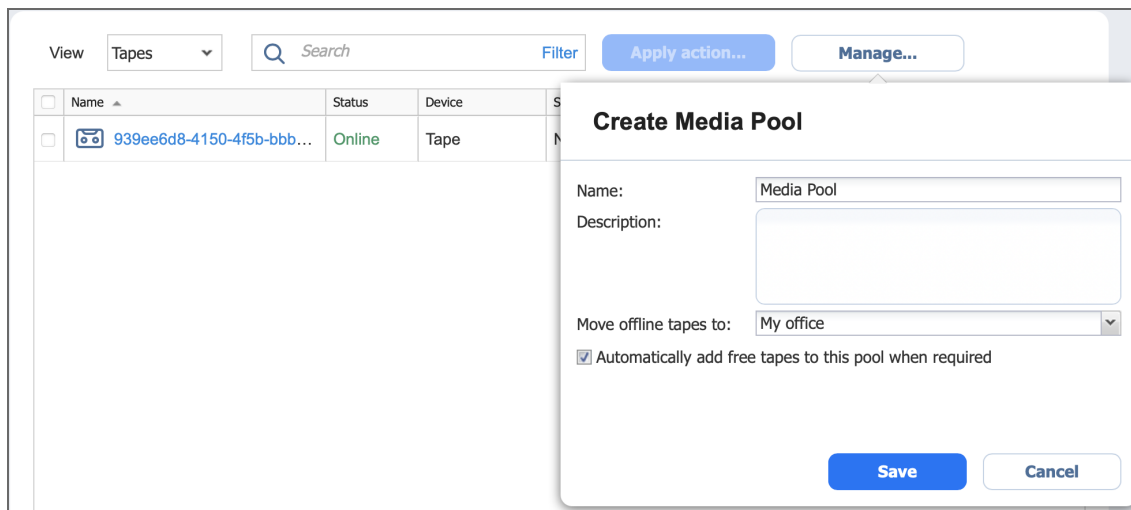
### Adding Media Pools

To create a Media Pool:

1. Go to **Settings > Tape**.
2. With the **Devices** or **Tapes** view selected, click **Manage** and select **Media Pools**. The **Media Pool Management** dialog box opens.




3. Click **Create Media Pool**.
4. In the **Create Media Pool** dialog box, specify the name for the Media Pool and provide its description (optionally).
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 **Save**. 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 it into the **Search** box.



## 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](#)
  - [Action Buttons](#)
  - [Details Pane](#)
  - [Options Pane](#)
  - [Tape Cartridge Contents Table](#)
  - [Recovery Page](#)
- [Bulk Tape Cartridge Management](#)

### Viewing Tapes

To navigate to the **Tapes** view, go to **Settings > Tapes** and select **Tapes** from the View drop-down list.

The **Tapes** view 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 in which the tape cartridge currently is.
- **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: number of backups present on the tape cartridge, the tape cartridge is empty, unidentified or contains the 3rd party data.
- **Type:** Displays the type of the tape cartridge:
  - Read/Write Tape
  - Write Protected Tape
  - Cleaning Tape

The columns availability within the table can be managed by clicking the small arrow in any column header and selecting/clearing the checkboxes next to the column name.

The screenshot shows a web interface for managing tape cartridges. On the left is a sidebar with navigation options: General, Inventory, Transporters, Repositories, and Tape (highlighted). The main area has a 'View' dropdown set to 'Tapes', a search box, and buttons for 'Filter', 'Apply action...', and 'Manage...'. Below this is a table with the following columns: Name, Status, Device, Slot #/Drive #, Media pool, Location, and Contents. A 'Columns' dropdown menu is open over the table, listing various columns with checkboxes: Label, Status, Device, Slot #/Drive #, Last written, Overwritable, Media pool, Location, Contents, and Type. The table currently displays one row with the name '939ee6d8-4150-4f5b-bbb...' and 9 backups.

## Searching for Tape Cartridges

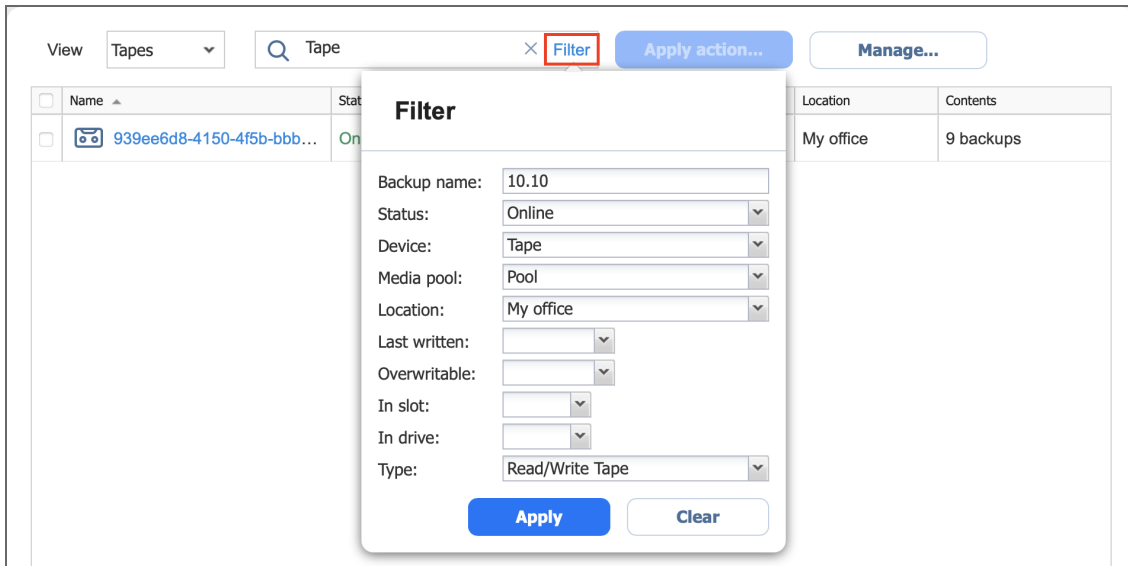
You can search for particular tape cartridge(s) by entering its name (or part of it) into the **Search** box. The table will dynamically change to display the search results matching your query.

This screenshot shows the same interface as the previous one, but with the search box containing the text 'Tape'. The table now displays a single row with the name '939ee6d8-4150-4f5b-bbb...', status 'Online', device 'Tape', slot/drive information 'No data', media pool 'Pool', location 'My office', and '9 backups'.

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** link in the **Search** box. In the dialog that opens you can 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
  - Reading
  - Writing
  - Warning
  - Moving
  - Error
- **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; that is, you cannot simultaneously enter a character's string and select filtering options.

## Tape Cartridge Management Page

Clicking on the 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:

- Action Buttons
- Details Pane
- Options Pane
- Tape Cartridge Contents Table

939ee6d8-4150-4f5b-bbb7-64b40d02d9d0

Scan
Move
Manage

**Details**

Status: Online in Tape

Name: 939ee6d8-4150-4f5b-bbb7-64b40d02d9d0

Barcode: none

Last written: Not applicable

Overwritable: Not applicable

Type: Read/Write Tape

**Options**

Label:

Media pool: Pool

Location: My office

**Tape Contents**

Name	Type	Job name	Tapes	Points	Last point
<a href="#">anhN_trans_13.3</a>	VMware	Backup copy job tape	1	1	Tue, 20 Apr 2021 at 19:41 (UTC +03:00)
<a href="#">192.168.77.73</a>	Physical	Backup copy #1	1	4	Fri, 02 Apr 2021 at 17:47 (UTC +03:00)
<a href="#">10.10.16.151</a>	Physical	Backup copy physical to tape alone	1	3	Thu, 19 Nov 2020 at 23:30 (UTC +02:00)
<a href="#">Ai_tr6.0_0510</a>	VMware	Unknown	1	1	Thu, 14 May 2020 at 14:49 (UTC +03:00)
<a href="#">LM_dir_f91</a>	VMware	Unknown	1	1	Mon, 27 Apr 2020 at 20:13 (UTC +03:00)
<a href="#">Ai_tr_0410</a>	VMware	Unknown	1	1	Tue, 14 Apr 2020 at 18:46 (UTC +03:00)
<a href="#">Ai-tr9.1-10</a>	VMware	Unknown	1	1	Fri, 31 Jan 2020 at 15:15 (UTC +02:00)

Close

## Action Buttons

The action buttons allow you to perform particular actions with the tape cartridge. Depending on the state of the tape cartridge, its type, status, etc., the button's availability may vary; the button can be disabled in case a certain action cannot be applied to the tape cartridge. Hovering over the disabled button opens 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 469](#).

The following actions can be applied to the tape cartridge:

- **Scanning:** Scanning of the tape cartridge implies checking 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 (i.e. on another product instance) but cannot be used for VM restoring until scanned.
  - Empty: The tape cartridge contains no data and is ready to be used for backing up.
  - 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 (e.g. 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?** message box, asking you to scan all tape cartridges. Clicking the **Scan all** link will initiate the scanning action for all newly discovered tape cartridges.

- **Editing:** 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.

- **Moving:** 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.
- **Protecting:** Applying this action to the tape cartridge makes it protected 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.



- **Marking 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.
- **Marking as cleaning:** Specialized tape cartridges designed for tape drive cleaning need to be marked as cleaning tapes. For the tape cartridges that have been marked as cleaning tapes, though still need to be reverted to normal tapes, the **Mark as cleaning** button is substituted with the **Mark as data** button.

### **Important**

Currently, the cleaning tapes inserted into the device are not automatically recognized by the system as cleaning. Instead, the system identifies the tapes to contain a 3rd party data. It is a user's responsibility to mark the tape as cleaning once the tape is inserted into the device and discovered. Otherwise, the cartridge will be performing the drive's cleaning automatically every time the library is refreshed

- **Retiring:** 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 your confirmation.
- **Erasing:** 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
- **Removing:** 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 #.  
The tape cartridge can be in one of the following statuses: Scanning / Online / Reading / Writing / Erasing / Warning/ Moving / Error / 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 will be expired.

- **Type:** Displays the type of the tape cartridge: Read/Write Tape / WORM Tape / Write Protected Tape / Cleaning Tape.

## Options Pane

In the **Options** pane, you can view the tape cartridge label, media pool and device location where the tape cartridge belongs. Changing this information is available via the **Edit** button.

## Tape Cartridge 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, the table displays generic information about the tape cartridge contents:

- "This tape contains third party data."
- "This tape cannot be identified due to a lack of barcode. Please scan the tape in order to discover its content."
- "This tape is empty."
- "This tape contains backups.": The tape contains backups created by NAKIVO Backup & Replication but has not been scanned yet.
- A grid of tape backups in case the tape contains backups.

If the tape cartridge contains backups and has been scanned already, the **Tape Contents** table displays the backups in the grid and provides the following information:

- **Name:** Displays the name of the backup. Clicking on the name of the backup opens the **Recovery** page.
- **Type:** Displays the type of a backup: VMware VM, Hyper-V VM or EC2 instance.
- **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.
- **Location:** Displays the location the tape cartridge is assigned to.

The **Tape Contents** table can be modified to display the column you need by clicking on the arrow icon in the table header, and selecting the required columns.

Clicking the column's header sorts the contents of the column.

## Recovery Page

Clicking on a backup name in the **Tape Cartridge Contents** table opens the **Recovery** page where you can view the backup information as well as see all recovery points available for this backup. From here, you can also initiate the recovery process.

10.10.16.151

Recover

**Backup Details**

Name: 10.10.16.151  
 Type: Physical machine  
 Tapes: 1  
 Recovery points: 3  
 First recovery point: Thu, 19 Nov 2020 at 13:49 (UTC +02:00)  
 Last recovery point: Thu, 19 Nov 2020 at 23:30 (UTC +02:00)  
 Location: My office  
 Job name: Backup copy physical to tape alone

**Recovery points**

Date	Type	Tape	Expiration Date
Thu, 19 Nov 2020 at 23:30 (UTC +02:00)	Incremental	939ee6d8-4150-4f5b-bbb7-64b...	Not applicable
Thu, 19 Nov 2020 at 17:46 (UTC +02:00)	Incremental	939ee6d8-4150-4f5b-bbb7-64b...	Not applicable
Thu, 19 Nov 2020 at 13:49 (UTC +02:00)	Full	939ee6d8-4150-4f5b-bbb7-64b...	Not applicable

Close

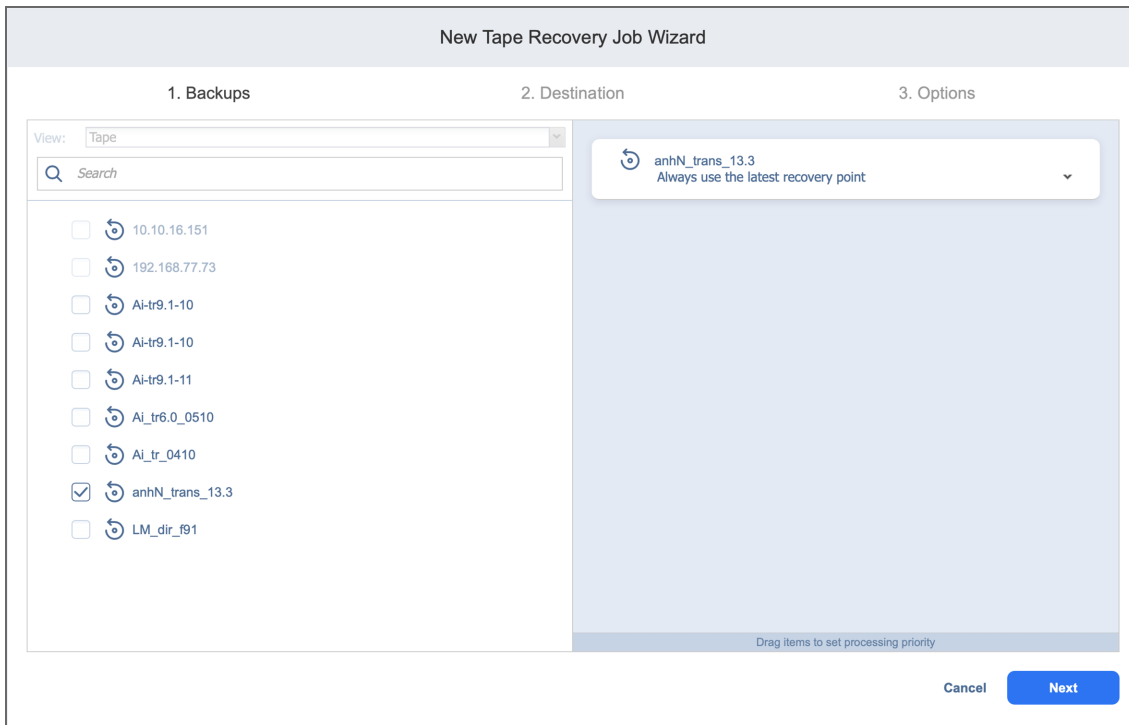
The **Backup Details** section provides the following information about the backup:

- **Name:** Shows the name of the backup.
- **Type:** Shows the type of backup: VMware, Hyper-V VM, EC2 instance or physical machine.
- **Tapes:** Shows the number of tape cartridges this backup is stored on.
- **Recovery points:** Shows the number of recovery points within the backup.
- **First recovery point:** Shows the date of the latest recovery point of the backup.
- **Last recovery point:** Shows the date of the most recent recovery point of the backup.
- **Location:** Shows the location the backup is assigned 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 on the recovery point, you can start the Recovery Wizard with the current backup and recovery point selected.
- **Type:** Indicates the type of backup: Full or Incremental.
- **Tape:** Indicates the name of the tape cartridge the backup is stored on.
- **Expiration date:** Indicates the date when the recovery point expires.

Clicking either the **Recover** button or the name of the recovery point opens the Recovery Job Wizard. For more information about recovering from tape cartridge, refer to [“Recovery From Tape” on page 956](#) .



## Bulk Tape Cartridge Management

Certain actions can be applied to several tape cartridges simultaneously. While being on the **Tapes** view, select the checkbox next to the tape cartridges you need to apply an action to, and click **Apply action**. In the dialog that opens, select an action to apply. Note that the availability of actions depends on various factors, thus not all actions may be available. For actions description, refer to Action Buttons.



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.

1 tape

■ 939ee6d8-4150-4f5b-bbb7-64b40d02d9d0

Barcode:	none
Last written:	Not applicable
Tape Label:	none
Media pool:	Pool
Location:	My office
Contents:	9 backups

Name	Date	Type	Expires
anhN_trans_13.3	20 Apr 2021 at 18:41	Full	Not applicable
192.168.77.73	02 Apr 2021 at 16:47	Incremental	Not applicable
192.168.77.73	02 Apr 2021 at 16:43	Incremental	Not applicable
192.168.77.73	02 Apr 2021 at 16:39	Incremental	Not applicable
192.168.77.73	02 Apr 2021 at 16:35	Full	Not applicable
10.10.16.151	19 Nov 2020 at 23:30	Incremental	Not applicable
10.10.16.151	19 Nov 2020 at 17:46	Incremental	Not applicable
10.10.16.151	19 Nov 2020 at 13:49	Full	Not applicable
AI_tr6.0_0510	14 May 2020 at 13:49	Full	Not applicable
LM_dir_f91	27 Apr 2020 at 19:13	Full	Not applicable
AI_tr_0410	14 Apr 2020 at 17:46	Full	Not applicable
AI-tr9.1-10	31 Jan 2020 at 15:15	Full	Not applicable
AI-tr9.1-10	31 Jan 2020 at 15:15	Incremental	Not applicable
AI-tr9.1-10	24 Jan 2020 at 14:22	Incremental	Not applicable
AI-tr9.1-10	20 Jan 2020 at 10:36	Incremental	Not applicable
AI-tr9.1-10	20 Jan 2020 at 10:26	Incremental	Not applicable
AI-tr9.1-10	17 Jan 2020 at 14:03	Full	Not applicable
AI-tr9.1-11	31 Jan 2020 at 14:48	Incremental	Not applicable
AI-tr9.1-11	24 Jan 2020 at 14:41	Full	Not applicable


**Alarms & Notifications**  
No alarms or notifications

## Managing Tape Devices

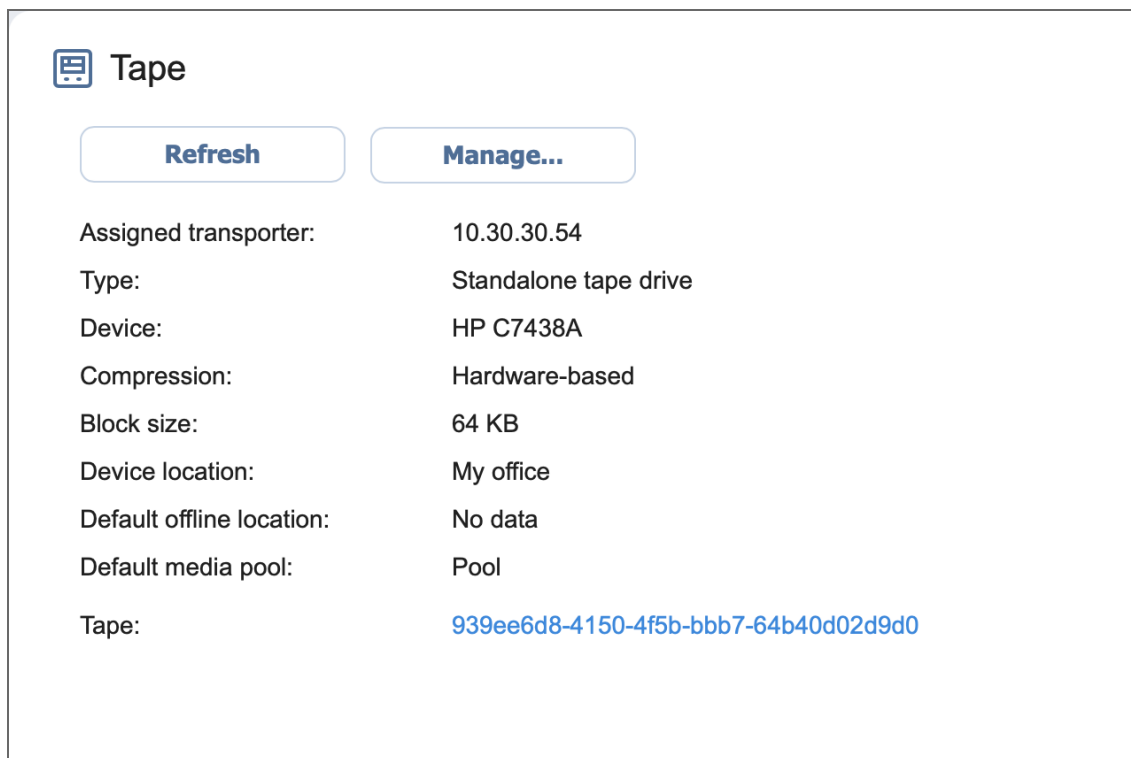
Once the tape devices are added to the system, you can view and manage them on the **Devices** view of the **Tape** page.

Hovering the mouse cursor over the device name opens the management controls:

- **Manage:** opens the following options:
  - **Edit:** Opens the Add New Robotic Tape Library or Virtual Tape Library or Add New 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.
- **Refresh:** Refresh action shall initiate the process of updating information regarding content of the tape 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
<b>system.email.smtp.localhost.mode</b>	Specifies how to determine the name of the localhost that is used in the SMTP HELO or EHLO commands.	<ul style="list-style-type: none"><li>• Default</li><li>• Use DNS</li><li>• Provide custom hostname</li></ul>
<b>system.email.smtp.localhost.name</b>	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.	
<b>system.email.smtp.tls.version</b>	Specifies the TLS version to use for SMTP server communication when TLS is configured in the Email Settings.	<ul style="list-style-type: none"><li>• Default</li><li>• TLS10</li><li>• TLS11</li><li>• TLS12</li><li>• TLS13</li></ul>
<b>system.email.notifications.skip.event.list</b>	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)

<b>system.vmware.esxi.ssh.port</b>	For VMware only. Specifies the SSH port to connect to ESXi (global setting).	<ul style="list-style-type: none"> <li>• Default value: 22</li> <li>• Minimum value: 1</li> <li>• Maximum value: 65535</li> </ul>
<b>system.vmware.skip.outdated.tools.checking</b>	For VMware only. When enabled, the system does not check VMware Tools outdated status when creating quiescing snapshot.	<ul style="list-style-type: none"> <li>• Unchecked (default)</li> <li>• Checked</li> </ul>
<b>system.vmware.skip.tag.discovery</b>	VMware only. When enabled, the system does not discover VMware Tags. This is applied to all tenants.	<ul style="list-style-type: none"> <li>• Unchecked (default)</li> <li>• Checked</li> </ul>
<b>http.max.upload.size</b>	Specifies the max upload size for file upload operations, <b>bytes</b> (global setting). If multiple files are uploaded, this is the total size. Use -1 for unlimited. Example: 200MB: 200000000	<ul style="list-style-type: none"> <li>• Default value: 1073741824</li> <li>• Minimum value: 1</li> <li>• Maximum value: 999999999999</li> </ul>
<b>system.auth.use.lockout</b>	Enables or disables the login lockout feature. When enabled, the offending IP address is not allowed to login after several failed attempts.	<ul style="list-style-type: none"> <li>• Unchecked (default)</li> <li>• Checked</li> </ul>
<b>system.auth.max.login.attempt.count</b>	Specifies the maximum number of failed login attempts to trigger the login lockout feature for the offending IP.	<ul style="list-style-type: none"> <li>• Default value: 5</li> <li>• Minimum value: 1</li> <li>• Maximum value: 9999</li> </ul>



<b>system.auth.lockout.timeout</b>	Specifies the timeout (minutes) for the login lockout feature. The offending IP is allowed to login again after the timeout expires.	<ul style="list-style-type: none"> <li>• Default value: 15</li> <li>• Minimum value: 1</li> <li>• Maximum value: 9999</li> </ul>
<b>system.auth.login.history.period</b>	Specifies the period (minutes) to calculate the maximum number of failed login attempts for the login lockout feature.	<ul style="list-style-type: none"> <li>• Default value: 5</li> <li>• Minimum value: 1</li> <li>• Maximum value: 9999</li> </ul>
<b>system.auth.ad.integration.follow.referrals</b>	<p>Defines LDAP/Active Directory behavior for referrals. When set to follow, all referrals are resolved (can be slow); otherwise they are ignored.</p> <p>What are the implications of the ignore option?</p> <p>* If you only have one domain, there should be no effects.</p> <p>* If you have multiple domains joined in a forest, then any cross-domain memberships will not be resolved.</p> <p>More info:  <a href="https://docs.oracle.com/javase/jndi/tutorial/ldap/referral/jndi.html">https://docs.oracle.com/javase/jndi/tutorial/ldap/referral/jndi.html</a></p>	<ul style="list-style-type: none"> <li>• follow (default)</li> <li>• ignore</li> </ul>
<b>system.auth.ad.integration.connect.timeout</b>	Specifies the timeout (milliseconds) for connecting LDAP/Active Directory.	<ul style="list-style-type: none"> <li>• Default value: 2000</li> <li>• Minimum value: 2000</li> <li>• Maximum value: 100000</li> </ul>

<b>system.auth.ad.integration.read.timeout</b>	Specifies the timeout (milliseconds) for reading LDAP/Active Directory operations.	<ul style="list-style-type: none"> <li>• Default value: 10000</li> <li>• Minimum value: 10000</li> <li>• Maximum value: 100000</li> </ul>
<b>system.auth.max.login.2fa.attempt.count</b>	Specifies the attempts for the login lockout feature. The offending IP is allowed to login again after the timeout expires.	<ul style="list-style-type: none"> <li>• Default value: 100</li> <li>• Minimum value: 1</li> <li>• Maximum value: 9999</li> </ul>
<b>system.auth.login.2fa.timeout</b>	Specifies the timeout (minutes) for the login lockout feature. The offending IP is allowed to login again after the timeout expires.	<ul style="list-style-type: none"> <li>• Default value: 5</li> <li>• Minimum value: 1</li> <li>• Maximum value: 9999</li> </ul>
<b>system.job.block.size</b>	<p>Select block size for processing data.</p> <p><b>Notes:</b></p> <ul style="list-style-type: none"> <li>• Deduplication can only be efficient with recovery points using the same block size.</li> <li>• Once the value is changed, the existing backup jobs, previously using a different block size, will produce a full backup on the next run.</li> <li>• Mapping to a backup with a different block size will be skipped.</li> </ul>	<ul style="list-style-type: none"> <li>• 4 MB (default)</li> <li>• 2 MB</li> <li>• 1 MB</li> <li>• 512 KB</li> <li>• 256 KB</li> <li>• 128 KB</li> <li>• 64 KB</li> </ul>
<b>system.job.map.new.source.item.scope</b>	The scope to search for the existing backup when adding a new source item to the job.	<ul style="list-style-type: none"> <li>• Default location (default)</li> </ul>

		<ul style="list-style-type: none"> <li>• Default transporter's locations</li> <li>• All locations</li> </ul>
<b>system.job.pool.queue.length</b>	Specifies the length of the job queue. A job is placed in a queue before execution. Requires restart.	<ul style="list-style-type: none"> <li>• Default value: 200</li> <li>• Minimum value: 10</li> <li>• Maximum value: 9999</li> </ul>
<b>system.job.pool.thread.min</b>	Specifies the minimum thread pool size for jobs. A job requires 1 thread from the job pool to start running. Requires restart.	<ul style="list-style-type: none"> <li>• Default value: 30</li> <li>• Minimum value: 10</li> <li>• Maximum value: 9999</li> </ul>
<b>system.job.pool.thread.max</b>	<p>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.</p> <p>If using Linux and systemd, please add the following to the service startup script: TasksMax=infinity</p>	<ul style="list-style-type: none"> <li>• Default value: 200</li> <li>• Minimum value: 10</li> <li>• Maximum value: 9999</li> </ul>
<b>system.job.resolve.host.hostname.on.transporter</b>	<p>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.</p>	<ul style="list-style-type: none"> <li>• Unchecked (default)</li> <li>• Checked</li> </ul>

	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).	
<b>system.job.resolve.transporter.hostname.on.transporter</b>	<p>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.</p> <p>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).</p>	<ul style="list-style-type: none"> <li>• Unchecked (default)</li> <li>• Checked</li> </ul>
<b>system.job.bandwidth.throttling.source</b>	If set, applies bandwidth throttling for data reading from source.	<ul style="list-style-type: none"> <li>• Checked (default)</li> <li>• Unchecked</li> </ul>
<b>system.job.bandwidth.throttling.target</b>	If set, applies bandwidth throttling for data writing to target.	<ul style="list-style-type: none"> <li>• Checked (default)</li> <li>• Unchecked</li> </ul>
<b>system.job.bandwidth.throttling.network</b>	If set, applies bandwidth throttling for data transfer between source and target.	<ul style="list-style-type: none"> <li>• Checked (default)</li> <li>• Unchecked</li> </ul>
<b>system.job.ict.skip.new.disk</b>	If set, new disks added to the source item will not be added to the job automatically.	<ul style="list-style-type: none"> <li>• Checked (default)</li> <li>• Unchecked</li> </ul>
<b>system.job.skip.manual.transporter.data.path.validation</b>	If set, transporter data path validation will be skipped for manually configured transporters.	<ul style="list-style-type: none"> <li>• Unchecked (default)</li> <li>• Checked</li> </ul>

<b>system.metadata.disable.ec2.instance.id.update</b>	Disables EC2 instance ID detection on product startup. The detection is done via a HTTP request to <a href="http://169.254.169.254/latest/meta-data/instance-id">http://169.254.169.254/latest/meta-data/instance-id</a> This is required for proper product functioning in the AWS cloud.	<ul style="list-style-type: none"> <li>• Unchecked (default)</li> <li>• Checked</li> </ul>
<b>system.task.pool.queue.length</b>	Specifies the length of the task queue. A task is placed in the queue before execution. Requires restart.	<ul style="list-style-type: none"> <li>• Default value: 200</li> <li>• Minimum value: 10</li> <li>• Maximum value: 9999</li> </ul>
<b>system.task.pool.thread.min</b>	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.	<ul style="list-style-type: none"> <li>• Default value: 30</li> <li>• Minimum value: 10</li> <li>• Maximum value: 9999</li> </ul>
<b>system.task.pool.thread.max</b>	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.	<ul style="list-style-type: none"> <li>• Default value: 200</li> <li>• Minimum value: 10</li> <li>• Maximum value: 9999</li> </ul>
<b>system.repository.min.free.space.byte</b>	Specifies the minimum free space (bytes) for the repository. If the free space goes below this value, an alarm is generated.	<ul style="list-style-type: none"> <li>• Default value: 5368709120</li> <li>• Minimum value: 1024</li> </ul>

		<ul style="list-style-type: none"> <li>• Maximum value: 1099511627776</li> </ul>
<b>system.repository.min.free.space.percent</b>	Specifies the minimum free space (percent) for the Backup repository. If the free space goes below this value, an alarm is generated.	<ul style="list-style-type: none"> <li>• Default value: 5</li> <li>• Minimum value: 1</li> <li>• Maximum value: 99</li> </ul>
<b>system.repository.maintenance.interrupt.timeout.seconds</b>	Specifies the timeout (seconds) to wait for repository maintenance stop during job run.	<ul style="list-style-type: none"> <li>• Default value: 300</li> <li>• Minimum value: 1</li> <li>• Maximum value: 86400</li> </ul>
<b>system.repository.refresh.backup.size.calculation</b>	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.	<ul style="list-style-type: none"> <li>• Checked (default)</li> <li>• Unchecked</li> </ul>
<b>system.repository.refresh.timeout.seconds</b>	Specifies the timeout (seconds) to wait for repository refresh.	<ul style="list-style-type: none"> <li>• Default value: 600</li> <li>• Minimum value: 1</li> <li>• Maximum value: 86400</li> </ul>

<p><b>system.repository.remove.backups.used.by.job</b></p>	<p>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.</p>	<ul style="list-style-type: none"> <li>• Unchecked (default)</li> <li>• Checked</li> </ul>
<p><b>system.product.skip.update.server.ssl.certificate.verification</b></p>	<p>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.</p>	<ul style="list-style-type: none"> <li>• Unchecked (default)</li> <li>• Checked</li> </ul>
<p><b>system.debug.mode.enabled</b></p>	<p>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).</p>	<ul style="list-style-type: none"> <li>• Unchecked (default)</li> <li>• Checked</li> </ul>
<p><b>system.debug.mode.log.passwords</b></p>	<p>When debug mode is enabled, also log passwords. This can be a security risk.</p>	<ul style="list-style-type: none"> <li>• Unchecked (default)</li> <li>• Checked</li> </ul>
<p><b>system.debug.mode.log.api.requests</b></p>	<p>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.</p>	<ul style="list-style-type: none"> <li>• Unchecked (default)</li> <li>• Checked</li> </ul>

<b>system.hyperv.optimize.queries</b>	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.	<ul style="list-style-type: none"> <li>• Checked (default)</li> <li>• Unchecked</li> </ul>
<b>system.hyperv.discovery.host.thread.count</b>	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.	<ul style="list-style-type: none"> <li>• Default value: 2</li> <li>• Minimum value: 1</li> <li>• Maximum value: 20</li> </ul>
<b>system.hyperv.discovery.vm.thread.count</b>	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.	<ul style="list-style-type: none"> <li>• Default value: 2</li> <li>• Minimum value: 1</li> <li>• Maximum value: 10</li> </ul>
<b>system.database.scheduled.backup.path</b>	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.	
<b>system.database.scheduled.backup.max.count</b>	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.	<ul style="list-style-type: none"> <li>• Default value: 5</li> <li>• Minimum value: 0</li> <li>• Maximum value: 365</li> </ul>
<b>system.logging.max.index</b>	Specifies the maximum index of log files. This works globally for all log files. Set 0 to use default value (configured in log4j.xml).	<ul style="list-style-type: none"> <li>• Default value: 0</li> <li>• Minimum value: 0</li> </ul>



		<ul style="list-style-type: none"> <li>• Maximum value: 999</li> </ul>
<b>system.product.min.free.space.byte</b>	Specifies the minimum free space (bytes) for the product installation folder. If the free space goes below this value, an alarm is generated.	<ul style="list-style-type: none"> <li>• Default value: 2147483648</li> <li>• Minimum value: 10485760</li> <li>• Maximum value: 10737418240</li> </ul>
<b>system.product.free.memory.threshold</b>	Specifies the minimum ratio for JVM free memory. If the free JVM memory goes below this value, an alarm is generated.	<ul style="list-style-type: none"> <li>• Default value: 0.1</li> <li>• Minimum value: 0.01</li> <li>• Maximum value: 0.9</li> </ul>
<b>system.nutanix.discovery.vm.thread.count</b>	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.	<ul style="list-style-type: none"> <li>• Default value: 2</li> <li>• Minimum value: 1</li> <li>• Maximum value: 10</li> </ul>
<b>system.plugin.flr.operation.timeout.seconds</b>	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.	<ul style="list-style-type: none"> <li>• Default value: 900</li> <li>• Minimum value: 1</li> <li>• Maximum value: 86400</li> </ul>
<b>system.physical.skip.os.checking</b>	Physical Windows host discovery only. When enabled, the system will not check the supported OS version.	<ul style="list-style-type: none"> <li>• Checked</li> <li>• Unchecked (default)</li> </ul>

<p><b>system.transporter.agent.injection.skip.vc.redist</b></p>	<p>When enabled, the system will not automatically install VC redistributable during Transporter/agent injection.</p>	<ul style="list-style-type: none"> <li>• Checked</li> <li>• Unchecked (default)</li> </ul>
<p><b>system.transporter.load.max.time.created.state.hours</b></p>	<p>Specifies the timeout (hours) to wait for getting Transporter load request. Default is 5 hours.</p>	<ul style="list-style-type: none"> <li>• Default value: 5</li> <li>• Minimum value: 1</li> <li>• Maximum value: 72</li> </ul>
<p><b>system.transporter.modern.min.heap.size.megabyte</b></p>	<p>Megabytes.</p> <p>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.</p> <p><b>Note:</b> Transporter restart is required to apply the setting.</p>	<ul style="list-style-type: none"> <li>• Default value: 512</li> <li>• Minimum value: 256</li> <li>• Maximum value: 65536</li> </ul>
<p><b>system.transporter.modern.max.heap.size.megabyte</b></p>	<p>Megabytes.</p> <p>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.</p> <p><b>Notes:</b></p> <p>-Xmx does not limit the total amount of memory that the JVM can use.</p> <p>Transporter restart is required to apply the setting.</p>	<ul style="list-style-type: none"> <li>• Default value: 3072</li> <li>• Minimum value: 256</li> <li>• Maximum value: 65536</li> </ul>

<p><b>system.transporter.modern.thread.stack.size.kilobyte</b></p>	<p>Kilobytes.</p> <p>-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.</p> <p><b>Note:</b> Transporter restart is required to apply the setting.</p>	<ul style="list-style-type: none"> <li>• Default value: 512</li> <li>• Minimum value: 64</li> <li>• Maximum value: 2048</li> </ul>
<p><b>system.transporter.modern.job.handler.max.thread.count</b></p>	<p>Specifies the job thread count for modern Transporter.</p> <p><b>Notes:</b></p> <ul style="list-style-type: none"> <li>• 1 job thread equals ~200MB of memory, consider changing the related setting.</li> <li>• Transporter restart is required to apply the setting.</li> </ul>	<ul style="list-style-type: none"> <li>• Default value: 10</li> <li>• Minimum value: 1</li> <li>• Maximum value: 128</li> </ul>
<p><b>system.transporter.modern.service.handler.max.thread.count</b></p>	<p>Specifies the service thread count for modern Transporter.</p> <p><b>Note:</b> Transporter restart is required to apply the setting.</p>	<ul style="list-style-type: none"> <li>• Default value: 10</li> <li>• Minimum value: 1</li> <li>• Maximum value: 128</li> </ul>
<p><b>system.transporter.jvm.ram.requirement</b></p>	<p>Bytes.</p> <p>For NASes only.</p> <p>Specifies the minimal ram required on NASes to create a SaaS repository.</p>	<ul style="list-style-type: none"> <li>• Default value: 4294967296</li> <li>• Minimum value: 0</li> <li>• Maximum value: 1099511627776</li> </ul>
<p><b>system.transporter.modern.thread.pool.size</b></p>	<p>Specifies the session factory thread pool size for modern Transporter.</p> <p><b>Note:</b> Transporter restart is required to apply the setting.</p>	<ul style="list-style-type: none"> <li>• Default value: 1000</li> <li>• Minimum value: 100</li> </ul>

		<ul style="list-style-type: none"> <li>• Maximum value: 1000</li> </ul>
<b>system.deleted.users.groups.remove.frequency</b>	Specifies the scheduled time for removing unnecessary deleted users, groups (in second).	<ul style="list-style-type: none"> <li>• Default value: 86400</li> <li>• Minimum value: 300</li> <li>• Maximum value: 1.7976931348623157e+308</li> </ul>
<b>system.inventory.allow.duplicated</b>	Microsoft 365 and physical machines only. When enabled, the system allows duplicated discovery items.	<ul style="list-style-type: none"> <li>• Unchecked (default)</li> <li>• Checked</li> </ul>
<b>system.inventory.optimize.discovery.time</b>	Microsoft 365 (SharePoint Online) only. When enabled, the system skips some attributes to optimize the discovery time.	<ul style="list-style-type: none"> <li>• Unchecked (default)</li> <li>• Checked</li> </ul>
<b>system.o365.suppress.throttling.event</b>	Suppress throttling warning.	<ul style="list-style-type: none"> <li>• Unchecked (default)</li> <li>• Checked</li> </ul>
<b>system.olr.dsamain.mount.port</b>	TCP port where DSAMAIN mounts NTDS.dit (AD database) for.	<ul style="list-style-type: none"> <li>• Default value: 5000</li> <li>• Minimum value: 1</li> <li>• Maximum value: 65535</li> </ul>
<b>system.product.register.disable.periodic.data.collection</b>	When enabled, the product will not send data bundles every 30 days.	<ul style="list-style-type: none"> <li>• Unchecked (default)</li> <li>• Checked</li> </ul>

<p><b>system.exchange.enable.direct.recovery</b></p>	<p>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 <b>Download items or forward via email</b> on the <b>Destination</b> page of the job wizard and then select the appropriate recovery type on the <b>Options</b> page.</p> <p>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.</p>	<ul style="list-style-type: none"> <li>• Unchecked (default)</li> <li>• Checked</li> </ul>
<p><b>system.repository.skip.periodic.refresh.on.transporter.busy.with.job</b></p>	<p>When enabled and any Transporter repository is locked by a running job, the product skips periodic refresh for this Transporter repository.</p>	<ul style="list-style-type: none"> <li>• Unchecked (default)</li> <li>• Checked</li> </ul>
<p><b>system.pql.custom.file.name</b></p>	<p>PQL file name in the userdata folder. Empty by default. If empty, the file will be downloaded from web.</p>	
<p><b>system.pql.cache.ttl.hours</b></p>	<p>Time to keep PQL file cache, in hours. Use 0 to disable the cache.</p>	<ul style="list-style-type: none"> <li>• Default value: 8</li> <li>• Minimum value: 0</li> <li>• Maximum value: 72</li> </ul>
<p><b>system.transporter.allow.new</b></p>	<p>Allows using newer Transporter versions.</p>	<ul style="list-style-type: none"> <li>• Unchecked (default)</li> <li>• Checked</li> </ul>

<b>system.event.skip.creating.event.list</b>	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)
<b>system.transporter.allow.old</b>	Allows using older (outdated) Transporter versions.	
<b>system.transporter.modern.idle.timeout</b>	Specifies the timeout (milliseconds) for modern Transporter IDLE. <b>Note:</b> Transporter restart is required to apply the setting.	<ul style="list-style-type: none"> <li>• Default value: 3600000</li> <li>• Minimum value: 0</li> <li>• Maximum value: 86400000</li> </ul>

## 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.

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)).

## 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**

# Virtual Appliance Configuration

This section covers the following topics:

- [“Configuring Network Settings of Virtual Appliance” below](#)
- [“Increasing Backup Repository Size on Virtual Appliance” below](#)
- [“Removing the Disk with Backup Repository from Virtual Appliance” on the next page](#)

## 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.

## 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.



# Removing the 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/root`).
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” below](#)
- [“Configuring Branding Settings in Multi-Tenant Mode” below](#)
- [“Configuring Email Notifications in Multi-Tenant Mode” on page 493](#)
- [“Configuring Email Settings in Multi-Tenant Mode” on page 494](#)
- [“Configuring System Settings in Multi-Tenant Mode” on page 494](#)
- [“Exporting and Importing Configuration in Multi-Tenant Mode” on page 495](#)

## 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:

<b>Image</b>	<b>Best format</b>	<b>Best resolution</b>
--------------	--------------------	------------------------

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](#).

The screenshot shows a configuration window for email settings. The fields are as follows:

- SMTP server: smtp.nakivo.com
- SMTP username: admin
- SMTP password: [masked]
- SMTP port: 25
- Encrypted connection: TLS, if possible
- From: info@nakivo.com
- To: admin@nakivo.com (highlighted with a red box)

Buttons: Send Test Email, Apply, Cancel.

## 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**

- 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.
- 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.

# 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](#)

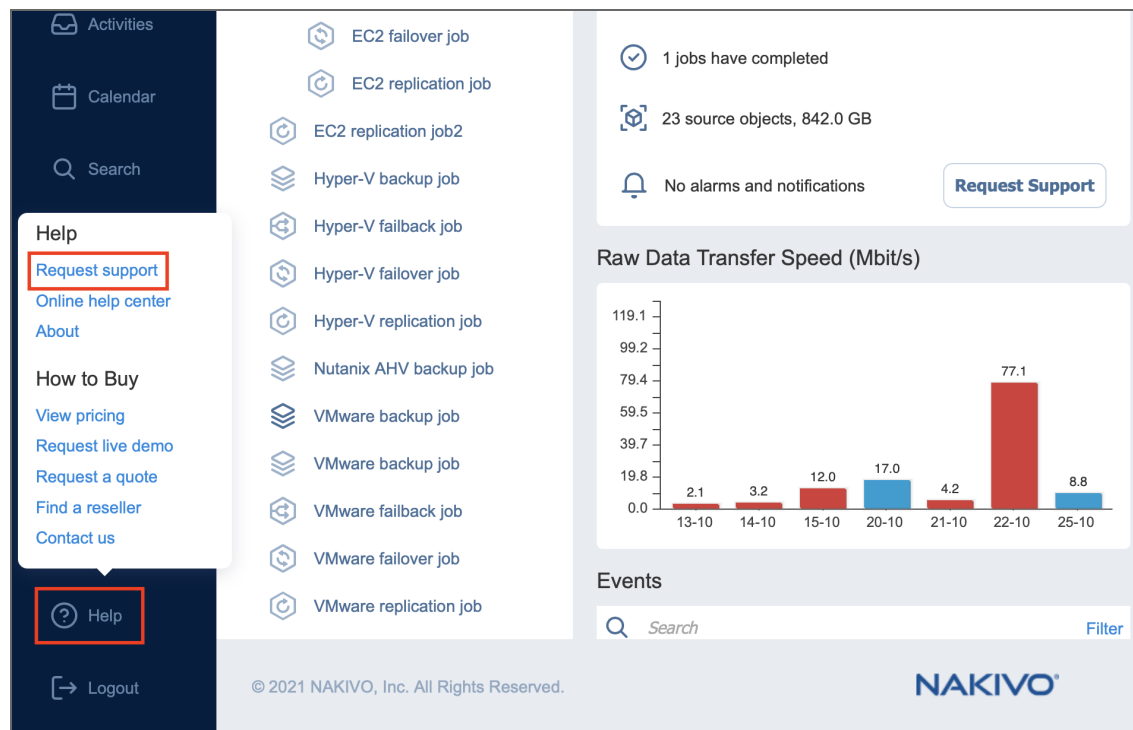
## Creating Support Bundles

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 **Include logs of all tenants** if you wish to include log files of all tenants to the support bundle.
7. Select **Include main database** if you want to include your main database.
8. Select **Include tenant databases** if you wish to include tenant databases containing most of the tenant configuration, including inventory, transporters, repositories, and jobs.
9. 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.



- Optionally, click **Download** to save the support bundle on your machine.



## 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:

- Open the [Upload Files to NAKIVO Support](#) page.
- 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.

- Enter your email address in the **Contact email** field.
- You can also enter the ID of your support ticket in the **Ticket ID** field if you have one opened.
- Optionally, enter a description in the **Description** field.
- 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.

- [Opening Built-in Support Chat](#)
- [Sending Files in Built-in Support Chat](#)
- [Sending Feedback to Built-in Support Chat](#)
- [Sending Email Transcript of Built-in Support Chat](#)
- [Disabling/Enabling Sound Notifications](#)
- [Disabling Built-in Support Chat](#)

## Opening Built-in Support Chat

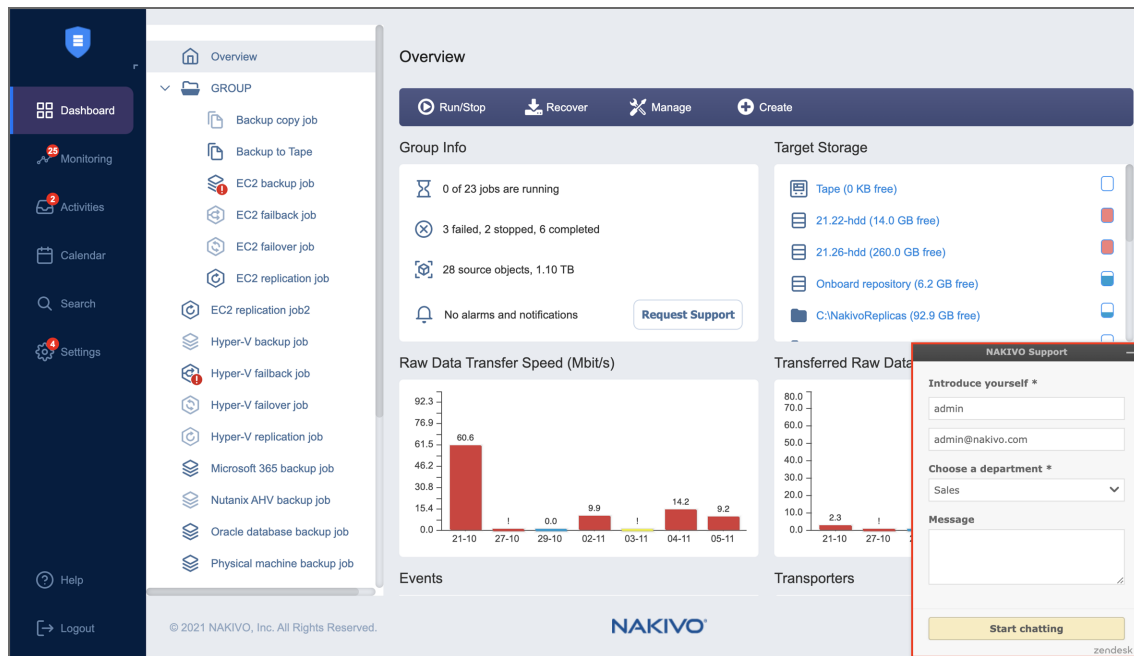
To open Built-in Support Chat, follow the steps below:

1. In the lower right corner of the NAKIVO Backup & Replication interface, click the chat 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 Chatting**.



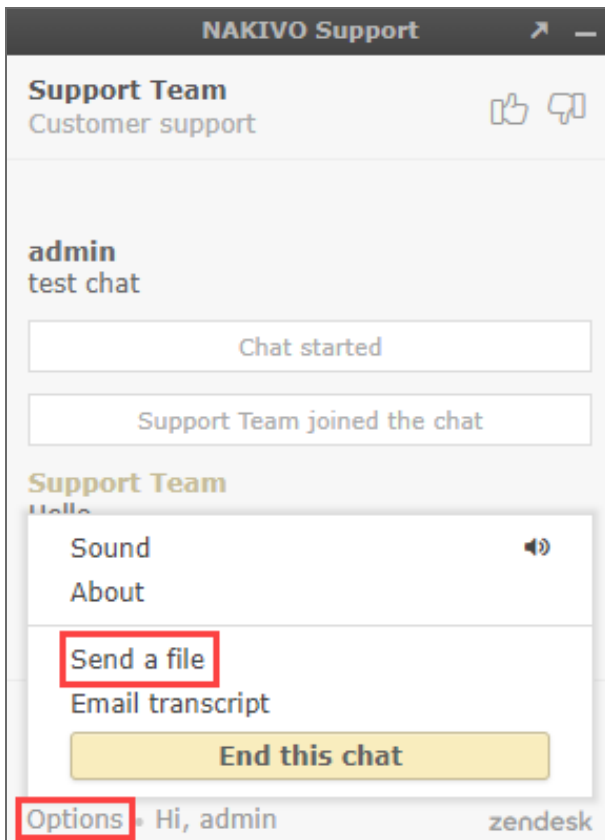
5. Your message is sent to a NAKIVO representative and will be processed as soon as possible. If needed, click the **Send Another** button to proceed with sending another chat message.

## Sending Files in Built-in Support Chat

Please use either of the following ways to send your files in Built-in Support Chat:

- Drag and drop: open **Windows File Explorer**, select necessary files, and then drag them and drop to the chat dialog.
- Built-in Support Chat interface:
  1. In the upper left corner of the Built-in Support Chat dialog, click **Options**.
  2. In the dialog that opens, click **Send a file**.

3. The **Open** dialog opens. Navigate to the location of your files, select them and then click **Open**.




**Note**

The following file formats are allowed: .pdf, .png, .jpeg, .gif, .txt. The maximum file size is 20 MB.

## Sending Feedback to Built-in Support Chat

You have the possibility of sending feedback to Built-in Support Chat: in the upper right corner of the dialog, click **Good** or **Bad**, as you deem appropriate.

If appropriate, leave a comment for NAKIVO Support Team: click **Leave a comment** and in the text box that opens, enter your comment about the chat service. Then click **Send**.

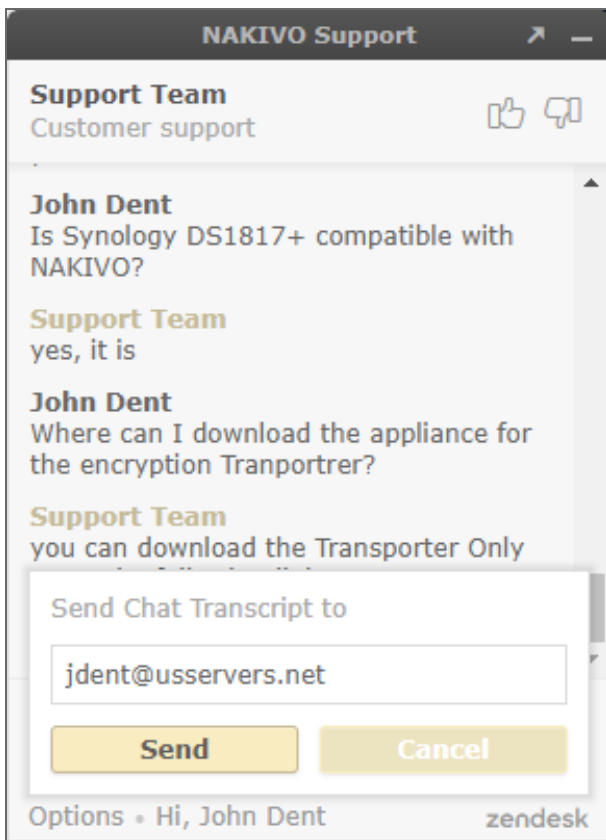


## Sending Email Transcript of Built-in Support Chat

Follow the steps below to send the transcript of your Built-in Support Chat session:

1. In the upper left corner of the Built-in Support Chat dialog, click **Options**.
2. In the dialog that opens, click **Email transcript**.
3. In the dialog that opens, make sure the email address of the recipient is correct, and then click **Send**.

Your Built-in Support Chat transcript will be sent to the specified email recipient.




## Disabling/Enabling Sound Notifications

By default, sound notifications are enabled for Built-in Support Chat.

Do the following to disable sound notifications in Built-in Support Chat:

1. In the upper left corner of the Built-in Support Chat dialog, click **Options**.
2. In the dialog that opens, click **Sound**.
3. Close the options dialog.

Sound notifications will be disabled for Built-in Support Chat.



## Disabling Built-in Support Chat

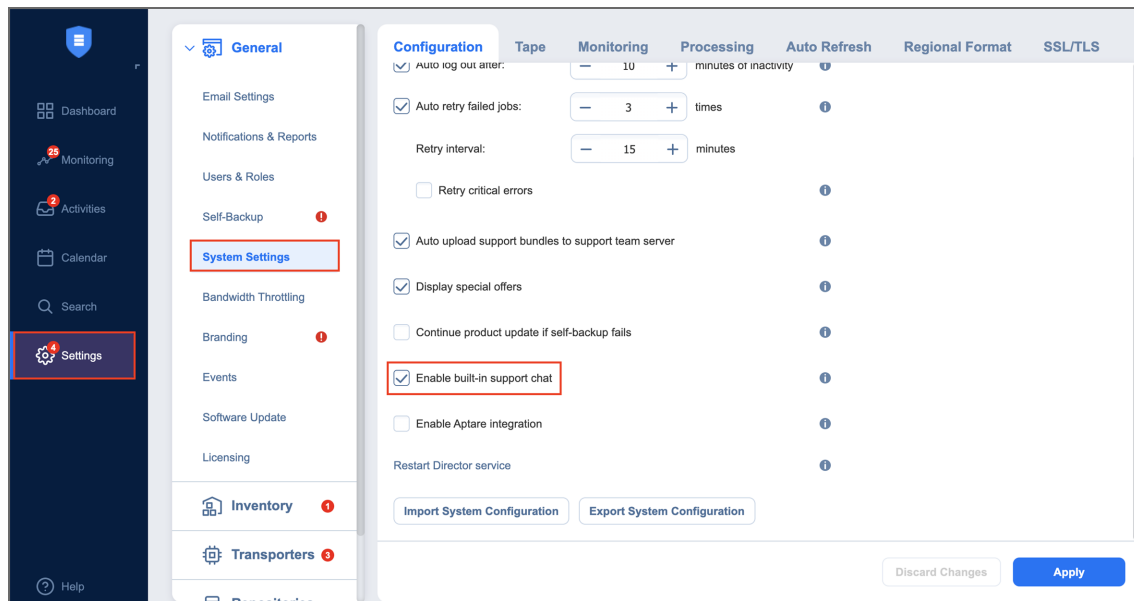
By default, the built-in support chat is enabled in your instance of NAKIVO Backup & Replication.

Do the following to disable built-in support chat:

1. Go to **Settings > General > System settings**.
2. Click **Edit** to make system settings editable and then deselect the **Enable built-in support chat** checkbox.



3. Click the **Apply** button.



### Note

When disabled, the Built-in Support Chat will not be available in all tenants of the NAKIVO Backup & Replication instance in multi-tenant mode.

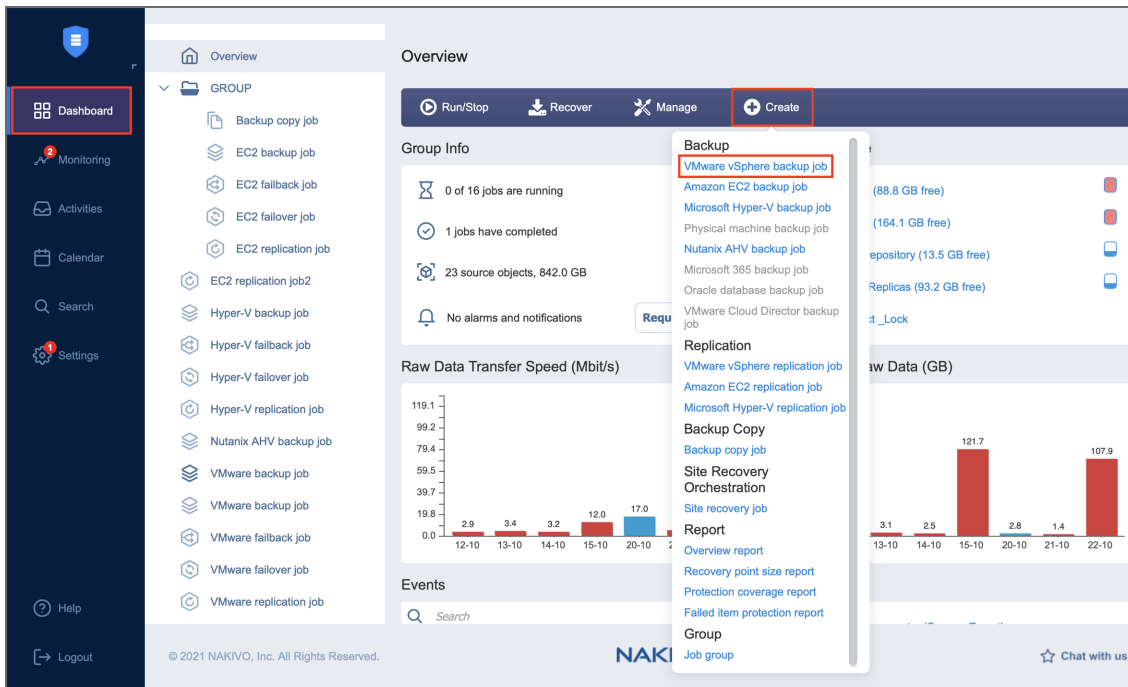
# Backup

This section contains the following topics:

- [“Creating VMware Backup Jobs” below](#)
- [“Creating VMware Cloud Director Backup Jobs” on page 533](#)
- [“Creating Hyper-V Backup Jobs” on page 551](#)
- [“Creating Amazon EC2 Backup Jobs” on page 573](#)
- [“Creating Nutanix AHV Backup Jobs” on page 597](#)
- [“Creating Physical Machine Backup Jobs” on page 615](#)
- [“Creating Microsoft 365 Backup Jobs” on page 634](#)
- [“Creating File Share Backup Jobs” on page 646](#)
- [“Creating Oracle Database Backup Jobs” on page 657](#)
- [“Creating Backup Copy Jobs” on page 668](#)
- [“Backing Up to Tape” on page 689](#)
- [“Staging \(Seeding\) Initial Backup” on page 705](#)
- [“Deleting Backups” on page 706](#)

## 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 **Create** on the **Dashboard** and then click **VMware vSphere backup job**.



The **New Backup Job Wizard for VMware vSphere** opens. Complete the wizard as described in the sections below:

- [“Backup Job Wizard for VMware: Source”](#) below
- [“Backup Job Wizard for VMware: Destination”](#) on page 512
- [“Backup Job Wizard for VMware: Schedule”](#) on page 516
- [“Backup Job Wizard for VMware: Retention”](#) on page 520
- [“Backup Job Wizard for VMware: Options”](#) on page 521

## 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](#)

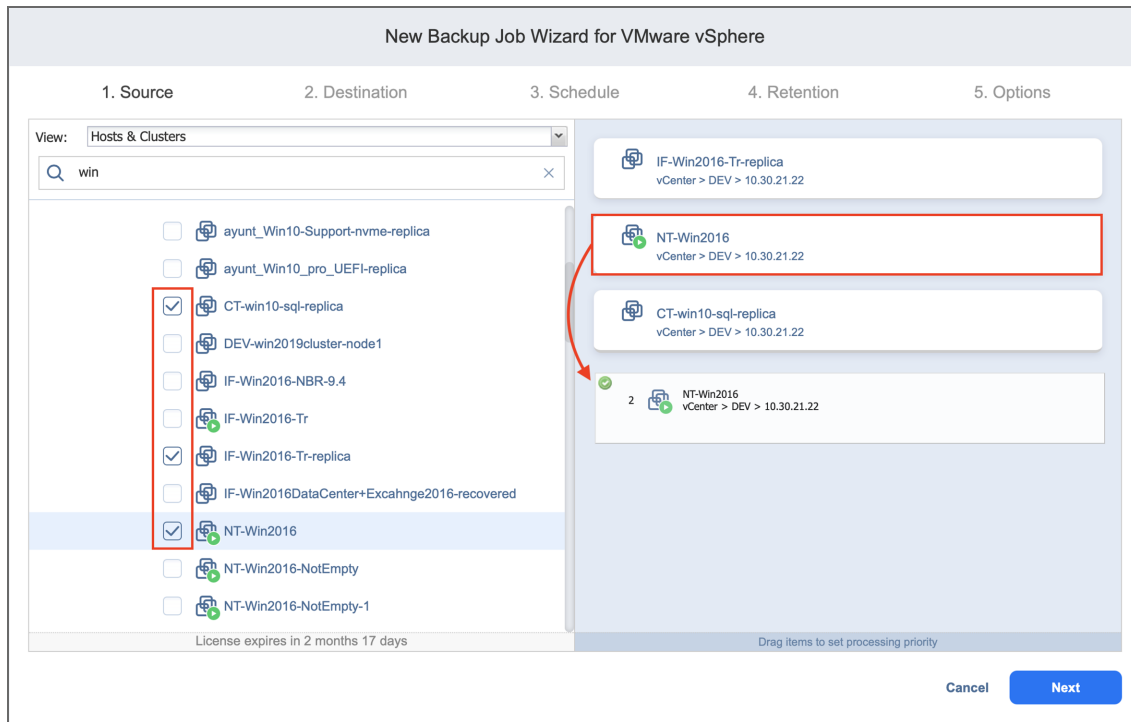
### Hosts and Clusters

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.

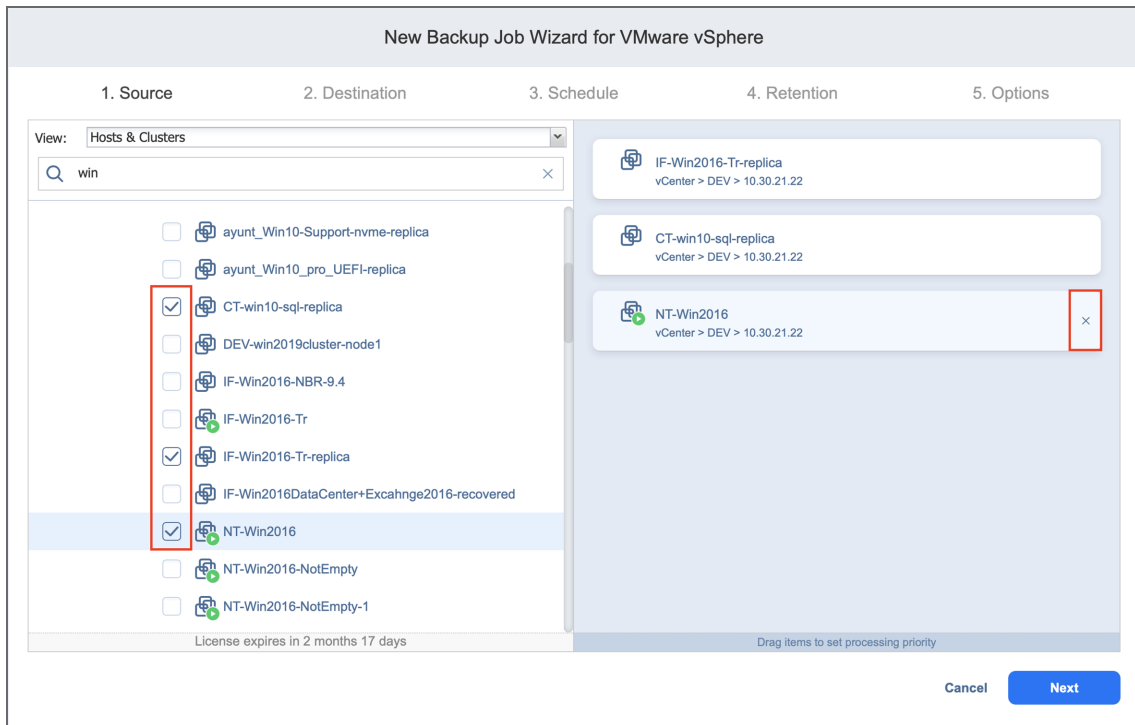


2. Select VMware items by placing a checkmark to the left of each item.
3. The selected items appear in the right pane of the page. You can drag a VM or a container with the pointer to a new position to reorder the selected items. By doing that, you can specify to back up the most important VMs first.




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.

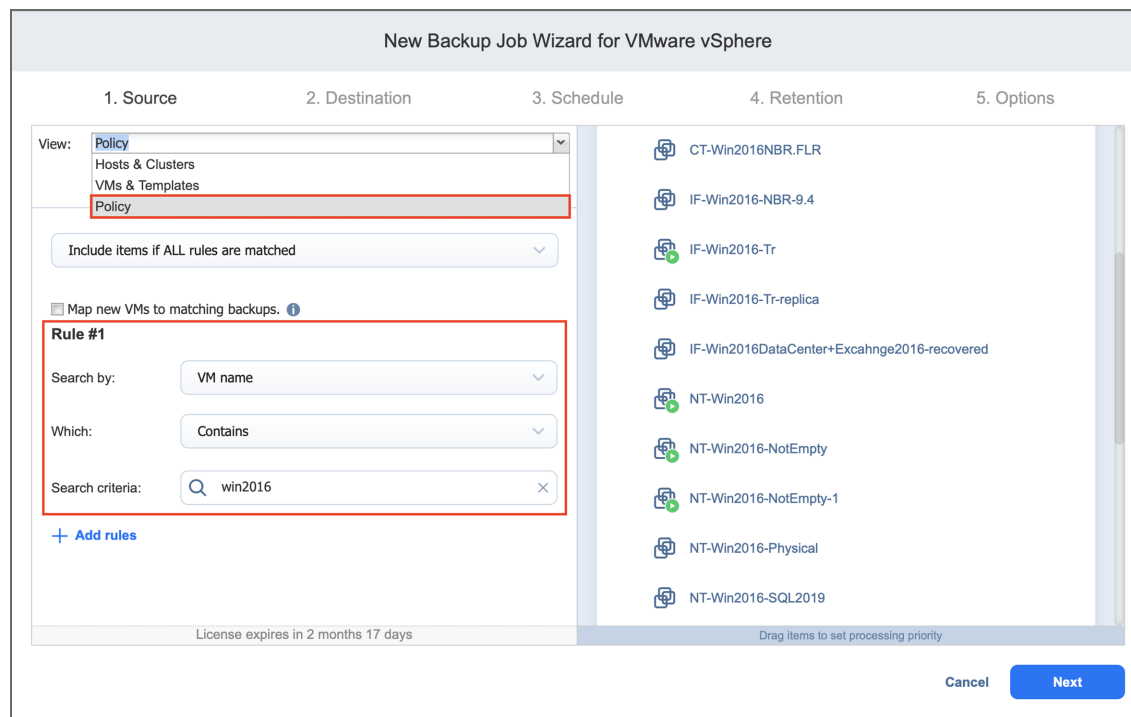


## Policy

When the **Policy** view is selected, job policies can be used. Refer to [“Managing Job Policies” on page 266](#) 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 269](#) 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 359.](#)
- 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.

New Backup Job Wizard for VMware vSphere

1. Source      2. Destination      3. Schedule      4. Retention      5. Options

Destination: Select target destination

Advanced setup... Search

Size of selected VMs: 360.0 GB

Onboard repository  
13.4 GB free (62% of 21.5 GB)

S3\_Object\_Lock  
10,240.00 TB free (100% of 10,240.00 TB)

To re-use existing backups, expand the Advanced setup and specify target backup for each VM.

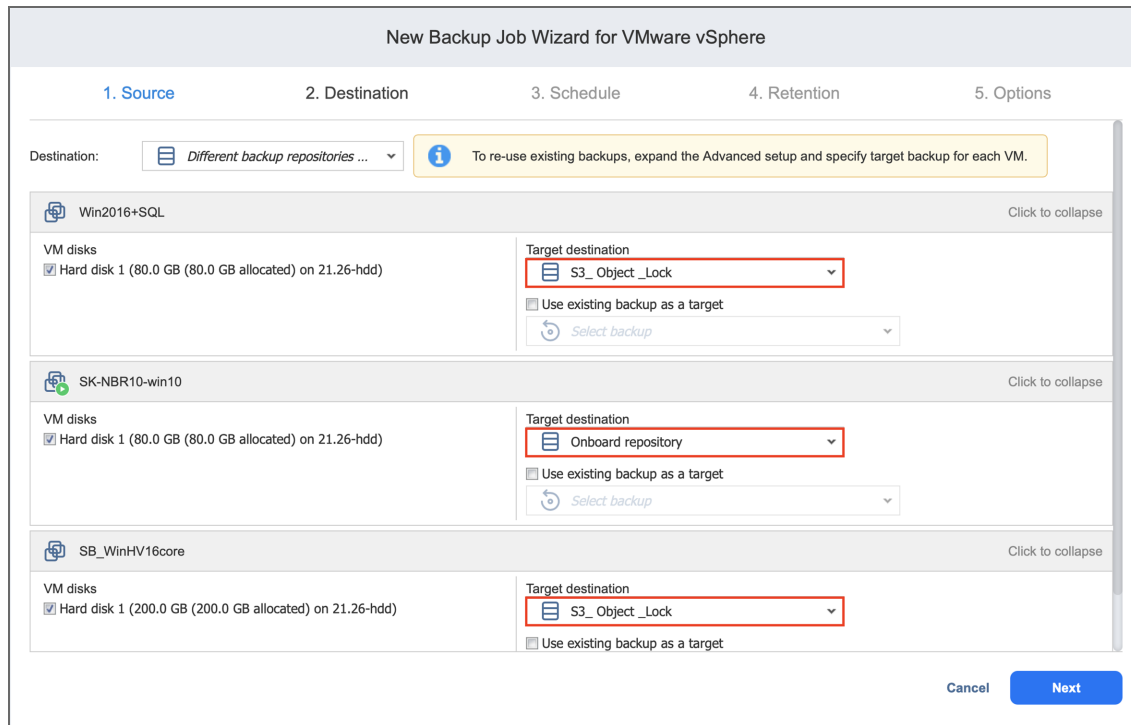
Cancel      Next

## 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.

New Backup Job Wizard for VMware vSphere

1. Source      2. Destination      3. Schedule      4. Retention      5. Options

Destination: ☰ Different backup repositories ... ℹ To re-use existing backups, expand the Advanced setup and specify target backup for each VM.

**Win2016+SQL** Click to collapse

VM disks  
 Hard disk 1 (80.0 GB (80.0 GB allocated) on 21.26-hdd)

Target destination  
☰ S3\_Object\_Lock  
 Use existing backup as a target  
🔄 Select backup

**SK-NBR10-win10** Click to collapse

VM disks  
 Hard disk 1 (80.0 GB (80.0 GB allocated) on 21.26-hdd)

Target destination  
☰ Onboard repository  
 Use existing backup as a target  
🔄 Select backup

**SB\_WinHV16core** Click to collapse

VM disks  
 Hard disk 1 (200.0 GB (200.0 GB allocated) on 21.26-hdd)

Target destination  
☰ S3\_Object\_Lock  
 Use existing backup as a target

Cancel Next

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.

The screenshot shows the 'New Backup Job Wizard for VMware vSphere' in the 'Source' step. It lists three VMs: SK-NBR10-win10, SB\_WinHV16core, and Win10\_PM. For each VM, there is a 'VM disks' section with checkboxes for 'Hard disk 1' and 'Hard disk 2'. In the Win10\_PM section, 'Hard disk 1' is unchecked and 'Hard disk 2' is checked, with a red box highlighting these two entries. To the right of each VM is a 'Target destination' dropdown menu and a checkbox for 'Use existing backup as a target'. At the bottom right, there are 'Cancel' and 'Next' buttons.

## 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:

- [Disabling Scheduled Job Execution](#)
- [Daily or Weekly Backup](#)
- [Monthly or Yearly Backup](#)
- [Periodic Backup](#)
- [Chained Job](#)
- [Additional Schedule](#)

### Disabling Scheduled Job Execution

If you want to start the job manually (without any schedule), select the **Do not schedule, run on demand** checkbox.

New Backup Job Wizard for VMware vSphere

1. Source      2. Destination      3. Schedule      4. Retention      5. Options

---

Do not schedule, run on demand

(UTC+02:00, EET) Eastern European Time

Schedule #1

Run daily/weekly

Starting at: 0:00      Ending: 6:00

Mon    Tue    Wed    Thu    Fri    Sat    Sun

[All days](#)   [Work days](#)   [Weekends](#)

every 1 weeks

Effective from

[Add another schedule](#)

[Show calendar](#)

**Next**      Cancel

## 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:

- 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 should come into effect.

New Backup Job Wizard for VMware vSphere

1. Source      2. Destination      3. Schedule      4. Retention      5. Options

---

Do not schedule, run on demand

(UTC+02:00, EET) Eastern European Time

Schedule #1

Run daily/weekly

Starting at: 0:00      Ending: 6:00

Mon    Tue    Wed    Thu    Fri    Sat    Sun

[All days](#)   [Work days](#)   [Weekends](#)

every 1 weeks

Effective from

[Add another schedule](#)

[Show calendar](#)

**Next**      Cancel

## Monthly or Yearly Backup

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.

The screenshot shows the 'New Backup Job Wizard for VMware vSphere' interface, specifically the '3. Schedule' step. The wizard has five steps: 1. Source, 2. Destination, 3. Schedule, 4. Retention, and 5. Options. In the 'Schedule' step, there are several options and settings:

- Do not schedule, run on demand
- Time zone: (UTC+02:00, EET) Eastern European Time
- Schedule #1**: Run monthly/yearly (highlighted with a red box)
- Run every: last Friday of every month
- Starting at: 0:00
- Ending: 6:00
- Effective from
- Links: [Add another schedule](#), [Show calendar](#)
- Buttons: **Next** (blue), **Cancel** (white)

## 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 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.

New Backup Job Wizard for VMware vSphere

1. Source      2. Destination      3. Schedule      4. Retention      5. Options

Do not schedule, run on demand  
 (UTC+02:00, EET) Eastern European Time

Schedule #1

**Run periodically** every 30 minutes

Starting at: 0:00    Ending: 6:00

Mon    Tue    Wed    Thu    Fri    Sat    Sun

All days   Work days   Weekends

Effective from

[Add another schedule](#)  
[Show calendar](#)

**Next**    Cancel

## 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.

New Backup Job Wizard for VMware vSphere

1. Source      2. Destination      3. Schedule      4. Retention      5. Options

Do not schedule, run on demand  
 (UTC+02:00, EET) Eastern European Time

Schedule #1

**Run after another job**

After the job: Physical machine backup job

Run this job: Immediately

After successful runs    After failed runs    After stopped runs

Effective from

[Add another schedule](#)  
[Show calendar](#)

**Next**    Cancel

## Additional 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.

The screenshot shows the 'New Backup Job Wizard for VMware vSphere' with five steps: 1. Source, 2. Destination, 3. Schedule, 4. Retention, and 5. Options. Step 3 is active. It includes a checkbox for 'Do not schedule, run on demand', a time zone dropdown set to '(UTC+02:00, EET) Eastern European Time', and a 'Schedule #1' section. This section has a dropdown for 'Run monthly/yearly', 'Run every' dropdowns for 'last', 'Friday', and 'of every month', and 'Starting at' and 'Ending at' time fields set to 0:00 and 6:00 respectively. There is also an 'Effective from' checkbox. A red box highlights the 'Add another schedule' button, with a 'Show calendar' link below it. 'Next' and 'Cancel' buttons are at the bottom right.

## Backup Job Wizard for VMware: Retention

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 S3** 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 S3** type 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:

- Only **Amazon S3** or **Local Folder** must be selected for Backup Repository Type on the Destination page of the wizard.
- If **Amazon S3** is selected as the Backup Repository, Object Lock must be enabled for the Amazon S3 bucket 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 for the backup job and enable/disable [app-aware mode](#), change tracking, [network acceleration](#), [encryption](#), [VM Verification](#), and other options. Proceed as described below.

The screenshot shows the 'New Backup Job Wizard for VMware vSphere' interface, specifically the '5. Options' step. The 'Job Options' section is highlighted with a red box. It includes the following settings:

- Job name: VMware backup job
- App-aware mode: Enabled (proceed on error) [settings]
- Change tracking: Use VMware CBT [settings]
- Network acceleration: Disabled [i]
- Network encryption: Disabled [i]
- VM verification: Disabled [i]
- Exclude swap files and partitions: Enabled [i]
- Exclude unused blocks: Enabled [i]

Below the Job Options section are the following sections:

- Full Backup Settings**
  - Create full backup: Every [Friday]
  - Full backup mode: Synthetic full [i]
  - If a full backup fails, create a full backup on the next job run [i]
- Pre and Post Actions**
  - Send job run reports to [i]
  - Truncate Exchange logs [On successful VM processing only] [i]
  - Truncate SQL Server logs [On successful VM processing only] [i]
  - Run local pre job script [i]
  - Run local post job script [i]
- Data Transfer**

At the bottom right, there are three buttons: 'Cancel', 'Finish', and 'Finish & Run'.

### Job Name

Enter a name for the backup job in the appropriate box.

### 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.

## 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 encryption.

## 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 VM 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 VM 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:

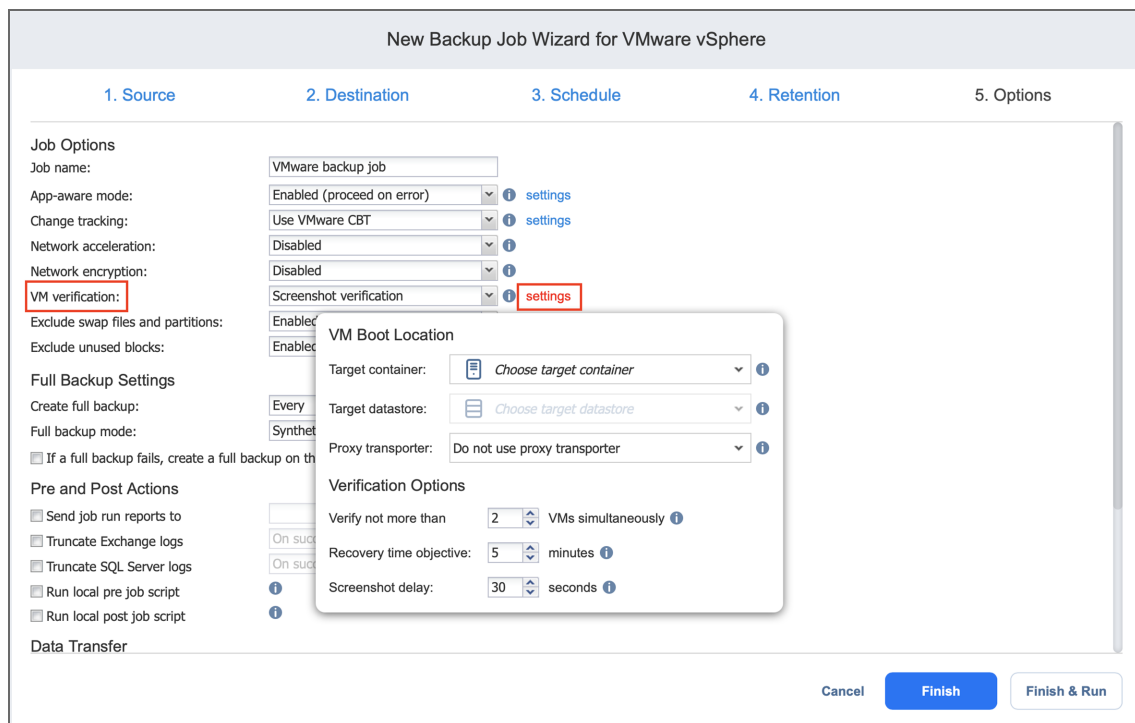
1. Provide a location of the VMs to be booted:
  - a. **Target Container:** Choose a target container (cluster, host, or resource pool) where VMs will be run using Flash VM Boot.
  - b. **Target Datastore:** Choose a datastore that will host changes to the recovered VMs.
  - c. **Proxy transporter:** Choose a proxy Transporter from the list of available Transporters.

**Note**

NAKIVO Backup & Replication will use a proxy Transporter in the following cases:

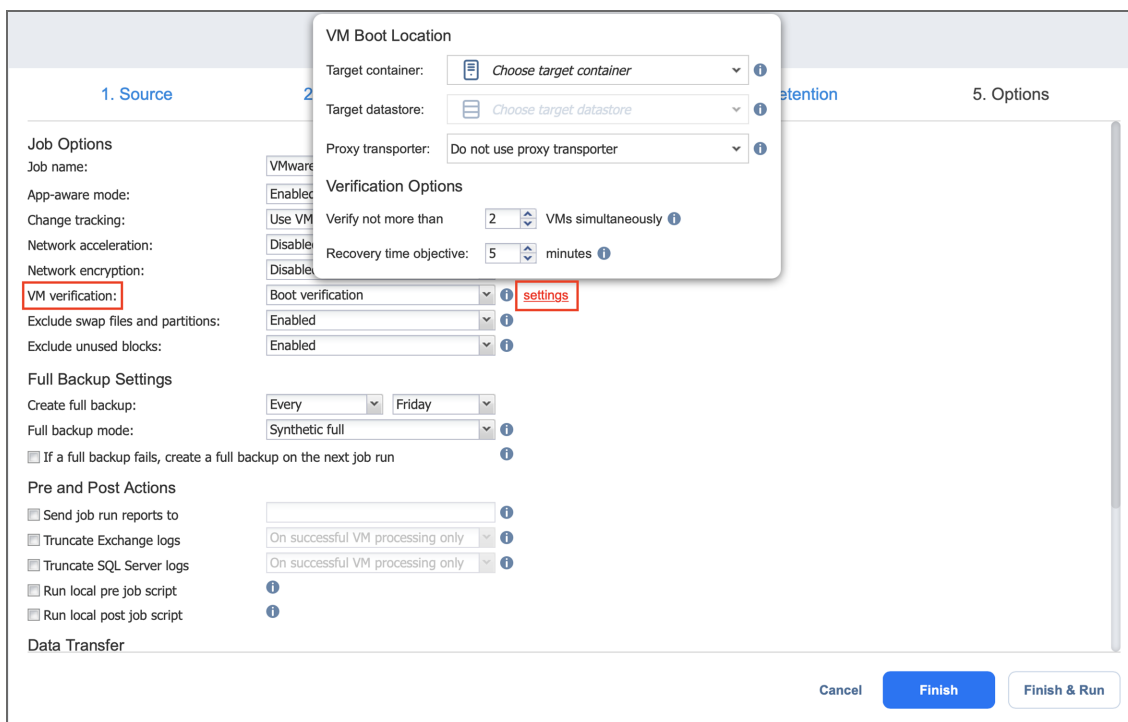
- The Transporter assigned to the Backup Repository cannot use iSCSI port 3260 because it is occupied by other services.
- iSCSI packages are missing on the Transporter assigned to the Backup Repository.

2. Set verification options:
  - 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 will be 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.

The screenshot shows the 'New Backup Job Wizard for VMware vSphere' with five steps: 1. Source, 2. Destination, 3. Schedule, 4. Retention, and 5. Options. The 'Options' step is active. Under 'Job Options', various settings are listed, including 'App-aware mode', 'Change tracking', 'Network acceleration', 'Network encryption', 'VM verification', 'Exclude swap files and partitions', and 'Exclude unused blocks'. The 'Full Backup Settings' section is highlighted with a red box and contains:
 

- 'Create full backup' set to 'Every Friday'
- 'Full backup mode' set to 'Synthetic full'
- A checked checkbox for 'If a full backup fails, create a full backup on the next job run'

 Below this are 'Pre and Post Actions' with checkboxes for sending reports, truncating logs, and running scripts. At the bottom right are 'Cancel', 'Finish', and 'Finish & Run' buttons.

## 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.

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 will be 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.

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 will be used to log in to the VMs that you have selected.

## 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.

The screenshot shows the 'New Backup Job Wizard for VMware vSphere' dialog box, specifically the 'Options' tab. The wizard is divided into five steps: 1. Source, 2. Destination, 3. Schedule, 4. Retention, and 5. Options. The 'Options' tab is currently selected. The 'Pre and Post Actions' section is highlighted with a red box and contains the following options:

- Send job run reports to
- Truncate Exchange logs
- Truncate SQL Server logs
- Run local pre job script
- Run local post job script

The 'Data Transfer' section is also visible and contains the following options:

- Transport mode: Automatic selection
- Transporter pool: Select transporter pool
- Transporters: Automatic selection
- Limit transporter load to: 3 concurrent tasks
- Bandwidth throttling: Disabled
- Bottleneck detection
- Backup from storage snapshot: Disabled

At the bottom of the dialog, there are three buttons: 'Cancel', 'Finish', and 'Finish & Run'.

## 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 Transporter load, and configure [bandwidth throttling](#).

## Transport Mode

To select a transport mode, do the following:

1. 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 do a manual snapshot of the transporter VM housing NAKIVO Backup & Replication while the hot-add disks are attached it. In case you need to do 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 284](#) 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 284](#) 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 HPE 3PAR and HPE Nimble Storage Snapshots” on page 55](#).

### Note

This option is disabled in the following cases:

- No supported storage devices were added to the Inventory.
- The selected VMs/disks don't reside on the supported storage devices.
- The selected VMs/disks reside on more than 1 datastore/virtual volume.

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

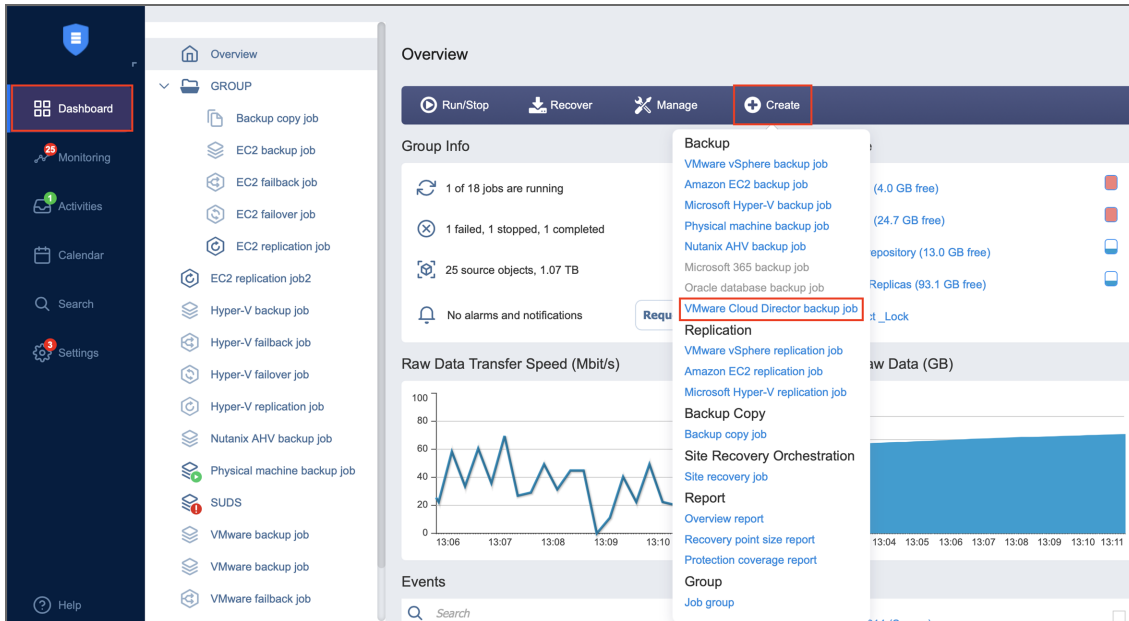
Click **Finish** or **Finish & Run** to complete job creation.

### Note

If you click **Finish & Run**, you will have to define the scope of your job. Refer to [“Running Jobs on Demand” on page 256](#) for details.

# Creating VMware Cloud Director Backup Jobs

With NAKIVO Backup & Replication, you can back up VMware Cloud Director objects, including vApps and individual VMs. To do this, create a backup job that specifies which object(s) should be backed up, where the backups should be located, how often the backup jobs should be run, and what backup options should be used. To create a backup job, click **Create** on the Dashboard and then click **VMware Cloud Director backup job**.



The **New Backup Job Wizard for VMware Cloud Director** opens. Complete the wizard as described in the sections below:


- [“Backup Job Wizard for VMware Cloud Director: Source”](#) below
- [“Backup Job Wizard for VMware Cloud Director: Destination”](#) on the next page
- [“Backup Job Wizard for VMware Cloud Director: Schedule”](#) on page 537
- [“Backup Job Wizard for VMware Cloud Director: Retention”](#) on page 541
- [“Backup Job Wizard for VMware Cloud Director: Options”](#) on page 542

## Backup Job Wizard for VMware Cloud Director: Source

On the **Source** page of the wizard, add VMware Cloud Director objects to your backup job. In the left pane of the wizard, the inventory tree displays all VMware Cloud Director objects added to the inventory and available for backup. Proceed as follows:

1. Optionally, you can filter the inventory tree by typing a string into the **Search** box. You can enter a part or the entire name of the object.

2. Select VMware Cloud Director objects by placing a checkmark next to the items you want to back up. When vApp is selected, all underlying VMs will be selected and backed up automatically.
3. The selected objects appear in the right pane of the page. You can remove a selected VMware Cloud Director object from the backup job in one of the following ways:
  - Deselect the object in the left pane. This will removes 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.
4. Click **Next** to proceed.




## Backup Job Wizard for VMware Cloud Director: Destination

On the **Destination** page of the wizard, select one or several Backup Repositories as the target for your backups of VMware Cloud Director objects. Proceed as described in the sections below:

- [Setting a Single Backup Repository for All Objects](#)
- [Setting Different Backup Repositories for Objects](#)
- [Mapping Source Objects to Existing Backups](#)
- [Excluding VM Disks from the Backup Job](#)

### Setting a Single Backup Repository for All Objects

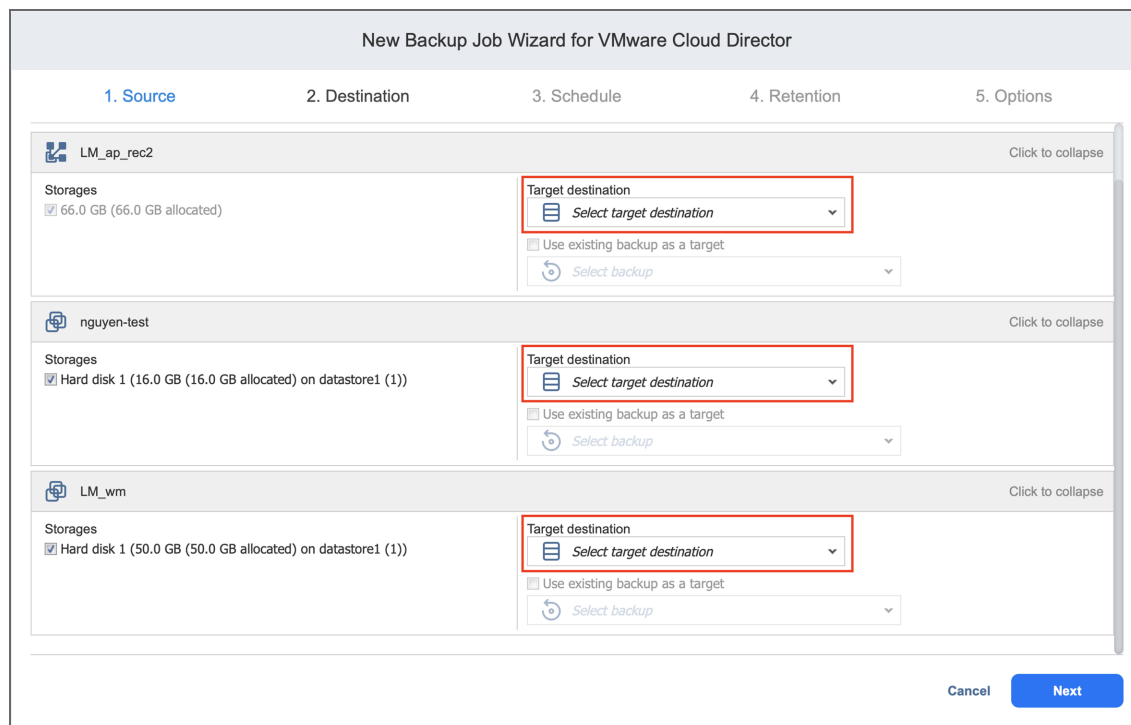
To back up the objects selected on the previous page to a single Backup Repository, choose a Backup Repository from the **Backup repository** drop-down list.



## Setting Different Backup Repositories for Objects

To back up the selected objects to different Backup Repositories, follow the steps below:

1. Click **Advanced setup**.
2. In the boxes, choose a Backup Repository for each object individually from the **Target destination** drop-down list.



## Mapping Source Objects to Existing Backups

If you have previously backed up VMware Cloud Director objects and then lost the backup job due to accidental job deletion, or if you need to recreate jobs in a new copy of NAKIVO Backup & Replication, you can map source objects to existing backups in order to avoid running full object backups again.

To map source objects 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 object 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, NAKIVO Backup & Replication analyzes the target object you have selected, determines how it is different from the source object, and transfers only the differential data.

### Note

Object backup mapping can be a time-consuming process that can take as much time as creating a full object backup.

After the job is completed, a new recovery point is created and existing recovery points are not be changed or overwritten.

The screenshot displays the 'New Backup Job Wizard for VMware Cloud Director' interface, specifically the 'Options' step. It shows three source objects being configured for backup:

- LM\_ap\_rec2**: Storage is 66.0 GB (66.0 GB allocated). Target destination is 'Onboard repository'. The checkbox 'Use existing backup as a target' is checked, and the 'Select backup' dropdown menu is highlighted with a red box.
- nguyen-test**: Storage is 'Hard disk 1 (16.0 GB (16.0 GB allocated) on datastore1 (1))'. Target destination is 'S3\_Object\_Lock'. The checkbox 'Use existing backup as a target' is unchecked.
- LM\_wm**: Storage is 'Hard disk 1 (50.0 GB (50.0 GB allocated) on datastore1 (1))'. Target destination is 'Onboard repository'. The checkbox 'Use existing backup as a target' is checked, and the 'Select backup' dropdown menu is highlighted with a red box.

At the bottom right, there are 'Cancel' and 'Next' buttons.

## Excluding 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.

The screenshot shows the 'New Backup Job Wizard for VMware Cloud Director' in the '1. Source' step. It lists three VMs: LM\_ap\_rec2, nguyen-test, and LM\_wm. Each VM has a 'Storages' section with a list of disks and a 'Target destination' section with a dropdown menu and a checkbox for 'Use existing backup as a target'. The 'Hard disk 1' for each VM is highlighted with a red box, indicating it is selected for backup.

VM Name	Storage Name	Storage Size	Storage Allocation	Target Destination	Use Existing Backup as a Target
LM_ap_rec2	66.0 GB	66.0 GB	66.0 GB allocated	Onboard repository	<input checked="" type="checkbox"/>
nguyen-test	Hard disk 1	16.0 GB	16.0 GB allocated on datastore1 (1)	S3_Object_Lock	<input type="checkbox"/>
LM_wm	Hard disk 1	50.0 GB	50.0 GB allocated on datastore1 (1)	Onboard repository	<input checked="" type="checkbox"/>

## Backup Job Wizard for VMware Cloud Director: 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:

- [Disabling Scheduled Job Execution](#)
- [Daily or Weekly Backup](#)
- [Monthly or Yearly Backup](#)
- [Periodic Backup](#)
- [Chained Job](#)
- [Additional Schedule](#)

### Disabling Scheduled Job Execution

If you do not want the job to run based on a schedule, select the **Do not schedule, run on demand** checkbox.



## Daily or Weekly Backup

To run the job on a daily basis, 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 start in the **Starting at** box.
- Optionally, you can specify the end time for the job in the **Ending** box. If the job is not completed by the time specified, the job is stopped.
- Select the days of the week on which the job should be started.
- If necessary, select the **Effective from** checkbox and pick the date when the schedule should come into effect.

The screenshot shows the 'New Backup Job Wizard for VMware Cloud Director' interface. The '3. Schedule' step is active. The 'Do not schedule, run on demand' checkbox is selected. The time zone is set to '(UTC+02:00, EET) Eastern European Time'. Under 'Schedule #1', the 'Run daily/weekly' option is selected in a dropdown menu, which is highlighted with a red box. Below this, the 'Starting at' time is 0:00 and the 'Ending' time is 6:00. The days of the week are checked: Mon, Tue, Wed, Thu, Fri, Sat, and Sun. There are also links for 'All days', 'Work days', and 'Weekends'. The frequency is set to 'every 1 weeks'. There is an 'Effective from' checkbox which is currently unchecked. At the bottom right, there are 'Next' and 'Cancel' buttons.

## Monthly or Yearly Backup

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.
- Optionally, specify the end time for the job in the **Ending** box. If the job is not completed by the time specified, the job is stopped.
- If necessary, select the **Effective from** checkbox and pick the date when the schedule should come into

effect.

The screenshot shows the 'New Backup Job Wizard for VMware Cloud Director' interface. The wizard has five steps: 1. Source, 2. Destination, 3. Schedule, 4. Retention, and 5. Options. The 'Schedule' step is currently active. It features a checkbox for 'Do not schedule, run on demand'. Below this is a time zone dropdown menu set to '(UTC+02:00, EET) Eastern European Time'. The 'Schedule #1' section has a dropdown menu with 'Run monthly/yearly' selected and highlighted by a red rectangular box. Below the dropdown are fields for 'Run every' (set to 'last'), 'of every' (set to 'month'), and 'of every' (set to 'month'). There are also 'Starting at' and 'Ending' time fields, both set to '0:00' and '6:00' respectively. An 'Effective from' checkbox is present but unchecked. At the bottom right, there are 'Next' and 'Cancel' buttons.

## 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 at** box.
- Specify the end time for the job in the **Ending** box. If the job is not completed by the time specified, the job is stopped.
- Select the days of the week on during which the job should be started.
- If necessary, select the **Effective from** checkbox and pick the date when the schedule should come into

effect.

New Backup Job Wizard for VMware Cloud Director

1. Source      2. Destination      3. Schedule      4. Retention      5. Options

Do not schedule, run on demand

(UTC+02:00, EET) Eastern European Time

Schedule #1

Run periodically every 30 minutes

Starting at: 0:00    Ending: 6:00

Mon    Tue    Wed    Thu    Fri    Sat    Sun

All days   Work days   Weekends

Effective from

[Add another schedule](#)

[Show calendar](#)

**Next**    Cancel

## Chained Job

To run the job after a previous one completes, 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 should be started.
- **Run this job:** Choose whether to run the current job immediately after the previous one completes or after a delay.
- **After successful runs:** When selected, the job runs if the previous one has completed successfully.
- **After failed runs:** When selected, the job runs if the previous one has failed.
- **After stopped runs:** When selected, the job runs if the previous one has been stopped.

- **Effective from:** When selected, the schedule comes into effect on the date selected.

The screenshot shows the 'New Backup Job Wizard for VMware Cloud Director' interface. The wizard is currently on the '3. Schedule' step. The 'Do not schedule, run on demand' option is selected. The time zone is set to '(UTC+02:00, EET) Eastern European Time'. Under 'Schedule #1', the dropdown menu is open and 'Run after another job' is selected, highlighted with a red box. Below this, 'After the job' is set to 'VMware Cloud Director backup job', and 'Run this job' is set to 'Immediately'. There are checkboxes for 'After successful runs' (checked), 'After failed runs', and 'After stopped runs'. The 'Effective from' checkbox is also present but unchecked. At the bottom right, there are 'Next' and 'Cancel' buttons.

## Additional 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.

## Backup Job Wizard for VMware Cloud Director: Retention

After each job run, NAKIVO Backup & Replication creates a recovery point for each object in the Backup Repository. A recovery point represents the backed up object as of a particular moment in time and allows you to recover individual files, application objects, or the entire object from the Backup Repository. You can specify how many recovery points you want to be retained in the Backup Repository based on the grandfather-father-son (GFS) backup rotation scheme.

When **Amazon S3** 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 S3** type 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. To set retention options, refer to the following sections:

- [Retention Settings](#)
- [Immutability](#)

## Retention Settings

Set the retention settings for the backup job:

- **Keep X last recovery points:** Retains the specified number of last recovery points for each object 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.

For more details and an example of job retention settings, refer to the [Keeping Recovery Points](#) article in the Knowledge Base.

## 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 that the recovery points are made immutable when the entire vApp backup is complete.

### Note

For the *Immutability* section to be available, the following conditions must be met:

- Only **Amazon S3** or **Local Folder** must be selected for Backup Repository Type on the Destination page of the wizard.
- If **Amazon S3** is selected as the Backup Repository, Object Lock must be enabled for the Amazon S3 bucket where your Backup Repository is located.
- For **Local Folder** type of Backup Repository, see [feature requirements](#).

The screenshot shows the 'Retention Settings' section of the Backup Job Wizard. It features five retention options, each with a checkbox and a dropdown menu for the number of days, weeks, months, or years. The 'Keep 10 last recovery points' and 'Make recovery points immutable for 10 days' options are checked. Below the retention settings is a 'Learn more' link and an 'Immutability' section with a checked option for 'Make recovery points immutable for 10 days'.

## Backup Job Wizard for VMware Cloud Director: Options

On the **Options** page of the wizard, you can set up various job options. Proceed as described in the sections below:

- [Job Options](#)
- [Full Backup Settings](#)
- [Pre and Post Job Actions](#)
- [Data Transfer](#)

## Job Options

In this section, you can specify a name for the backup job and enable/disable app-aware mode, change tracking, network acceleration, encryption, VM Verification, and other job options. Proceed as described below.

### Job Name

Enter a name for the backup job in the appropriate box.

### App-Aware Mode

When the app-aware mode option is enabled, object backup is performed using VMware Guest OS quiescing which relies on Microsoft VSS to ensure that application data is consistent. For details, refer to [“Application and Database Support” on page 37](#).

### 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 speeds up the job. Click the **error handling** link to specify how errors should be handled.

**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 transfers. Select this option when you back up over WAN or slow LAN links. For more information about this feature, refer to [“Network Acceleration” on page 63](#).

## Exclude Swap Files and Partitions

With this option enabled, NAKIVO Backup & Replication automatically excludes swap files and partitions during the backup process. For more information about the feature, refer to [“Excluding Swap Files and Partitions” on page 41](#).

## 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 you to reduce backup storage space and object processing time. For more information, refer to [“Excluding Unused Blocks” on page 41](#).

The screenshot displays the 'New Backup Job Wizard for VMware Cloud Director' interface, specifically the 'Options' step. The wizard is divided into five steps: 1. Source, 2. Destination, 3. Schedule, 4. Retention, and 5. Options. The 'Options' step is currently active and contains several configuration sections:

- Job Options:** This section is highlighted with a red box. It includes:
  - Job name: VMware Cloud Director backup job
  - App-aware mode: Enabled (proceed on error)
  - Change tracking: Use VMware CBT
  - Network acceleration: Disabled
  - Network encryption: Disabled
  - Exclude swap files and partitions: Enabled
  - Exclude unused blocks: Enabled
- Full Backup Settings:**
  - Create full backup: Every Friday
  - Full backup mode: Synthetic full
  - If a full backup fails, create a full backup on the next job run
- Pre and Post Actions:**
  - Send job run reports to
  - Truncate Exchange logs: On successful VM processing only
  - Truncate SQL Server logs: On successful VM processing only
  - Run local pre job script
  - Run local post job script
- Data Transfer:**
  - Transport mode: Automatic selection

At the bottom of the wizard, there are three buttons: 'Cancel', 'Finish', and 'Finish & Run'.

## Full Backup Settings

When you set the type of Backup Repository on the **Destination page** of the wizard 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 the data 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.

The screenshot shows the 'New Backup Job Wizard for VMware Cloud Director' interface, specifically the 'Options' step. The wizard is divided into five steps: 1. Source, 2. Destination, 3. Schedule, 4. Retention, and 5. Options. The 'Options' step is currently active. The 'Job Options' section includes fields for 'Job name' (VMware Cloud Director backup job), 'App-aware mode' (Enabled (proceed on error)), 'Change tracking' (Use VMware CBT), 'Network acceleration' (Disabled), 'Network encryption' (Disabled), 'Exclude swap files and partitions' (Enabled), and 'Exclude unused blocks' (Enabled). The 'Full Backup Settings' section is highlighted with a red box and includes 'Create full backup' (Every Friday), 'Full backup mode' (Synthetic full), and a checkbox for 'If a full backup fails, create a full backup on the next job run'. Below this is the 'Pre and Post Actions' section with checkboxes for 'Send job run reports to', 'Truncate Exchange logs', 'Truncate SQL Server logs', 'Run local pre job script', and 'Run local post job script'. The 'Data Transfer' section includes 'Transport mode' (Automatic selection). At the bottom right, there are 'Cancel', 'Finish', and 'Finish & Run' buttons.

## 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 specific 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. Refer to [“Notifications & Reports” on page 294](#).

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 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 (that is 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 that are already committed to (available in) the Microsoft Exchange database.

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 will be 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 (that is 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.

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 will be used to log in to the VMs that you have selected

## Pre Job Script

To run a script before NAKIVO Backup & Replication begins backing up 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.
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, object 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 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:** When this option is selected, the job performs object 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 object backup is not performed.

## Post Job Script

To run a script after NAKIVO Backup & Replication has finished backing up 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.
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, 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 affect 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 object backup is successful.

### Important

Pre- and post-job scripts can be executed only on the machine on which the Director is installed.

The screenshot shows the 'New Backup Job Wizard for VMware Cloud Director' with the 'Options' tab selected. The 'Pre and Post Actions' section is highlighted with a red box and contains the following options:

- Send job run reports to
- Truncate Exchange logs
- Truncate SQL Server logs
- Run local pre job script
- Run local post job script

Other visible options in the 'Options' tab include:

- Network acceleration: Unaccelerated
- Network encryption: Disabled
- Exclude swap files and partitions: Enabled
- Exclude unused blocks: Enabled
- Full Backup Settings: Create full backup: Every Friday; Full backup mode: Synthetic full
- Data Transfer: Transport mode: Automatic selection; Transporter pool: Select transporter pool; Transporters: Automatic selection; Limit transporter load to: 3 concurrent tasks; Bandwidth throttling: Disabled

## Data Transfer

In the *Data Transfer* section of the **Options** page, you can choose the transport mode, select a Transporter to be used for reading data from source objects, specify the Transporter load, and configure [bandwidth throttling](#).

## Transport Mode

To select a transport mode, do the following:

1. In the *Data Transfer* section, choose the transport mode for retrieving object 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: HotAdd > 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. HotAdd is not supported for physical machines.
  - **SAN only:** When this option is selected, NAKIVO Backup & Replication uses direct SAN access only to retrieve VM data. If direct SAN access to VM data is not available, the job fails.

- **HotAdd only:** When this option is selected, NAKIVO Backup & Replication uses HotAdd only to retrieve VM data. If direct HotAdd is not available, the job fails.

**Note**

When HotAdd cannot be enabled for at least one disk of a source VM (even if the disk is deselected in the job), HotAdd is unavailable for all disks of the VM.

- **LAN only:** When this option is selected, NAKIVO Backup & Replication retrieves VM data over LAN only.

## 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 object, 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. When more than one Transporter is deployed for NAKIVO Backup & Replication, it is important to determine which of these Transporters 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.

**Note**

The target Transporter for the backup job will always be the Transporter assigned to the Backup Repository.

## 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.
- **Select manually per host:** When selected, the Transporters can be selected manually or can be set to be chosen automatically for each separate host.
- **Use primary transporters only:** When selected, only primary Transporters are used during job execution.

## Transporter Load

You can limit the maximum number of Transporter tasks assigned to 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](#) 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](#) 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.

## Completing the New Backup Job Wizard for VMware Cloud Director

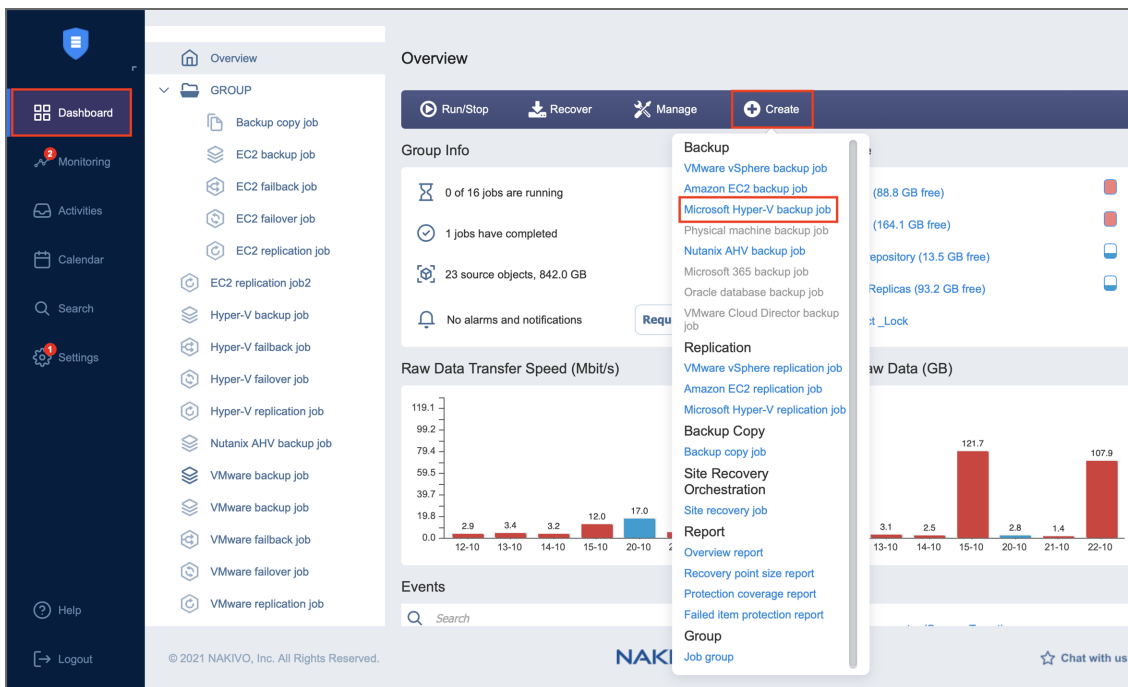
Click **Finish** or **Finish & Run** to complete job creation.

### Note

If you click **Finish & Run**, you will have to define the scope of your job. For details, refer to [“Running Jobs on Demand” on page 256](#).

## Creating Hyper-V Backup Jobs

With NAKIVO Backup & Replication, you can [back up](#) Hyper-V 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 occur, and what backup options should be used. To create a Hyper-V backup job, click **Create** on the **Dashboard** and then click **Microsoft Hyper-V backup job**.




The **New Backup Job Wizard for Microsoft Hyper-V** opens. Complete the wizard as described in the sections below:

- [“Backup Job Wizard for Hyper-V: Source” below](#)
- [“Backup Job Wizard for Hyper-V: Destination” on page 556](#)
- [“Backup Job Wizard for Hyper-V: Schedule” on page 559](#)
- [“Backup Job Wizard for Hyper-V: Retention” on page 563](#)
- [“Backup Job Wizard for Hyper-V: Options” on page 564](#)

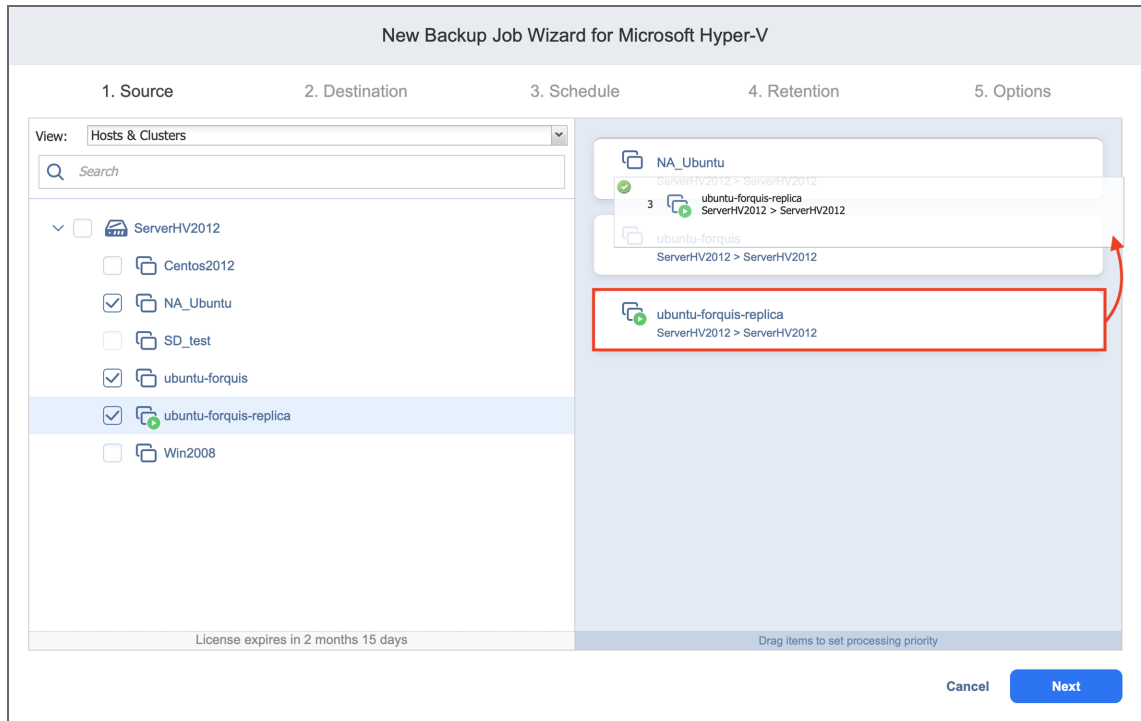
## Backup Job Wizard for Hyper-V: Source

On the **Source** page of the wizard, you can add Hyper-V VMs to your backup job. Proceed as follows:

1. In the left pane of the page, choose either of the following inventory views:
  - **Hosts & Clusters:** If chosen, the inventory tree opens in the left pane and shows all of the Hyper-V items: clusters, servers, and VMs. Proceed as follows:



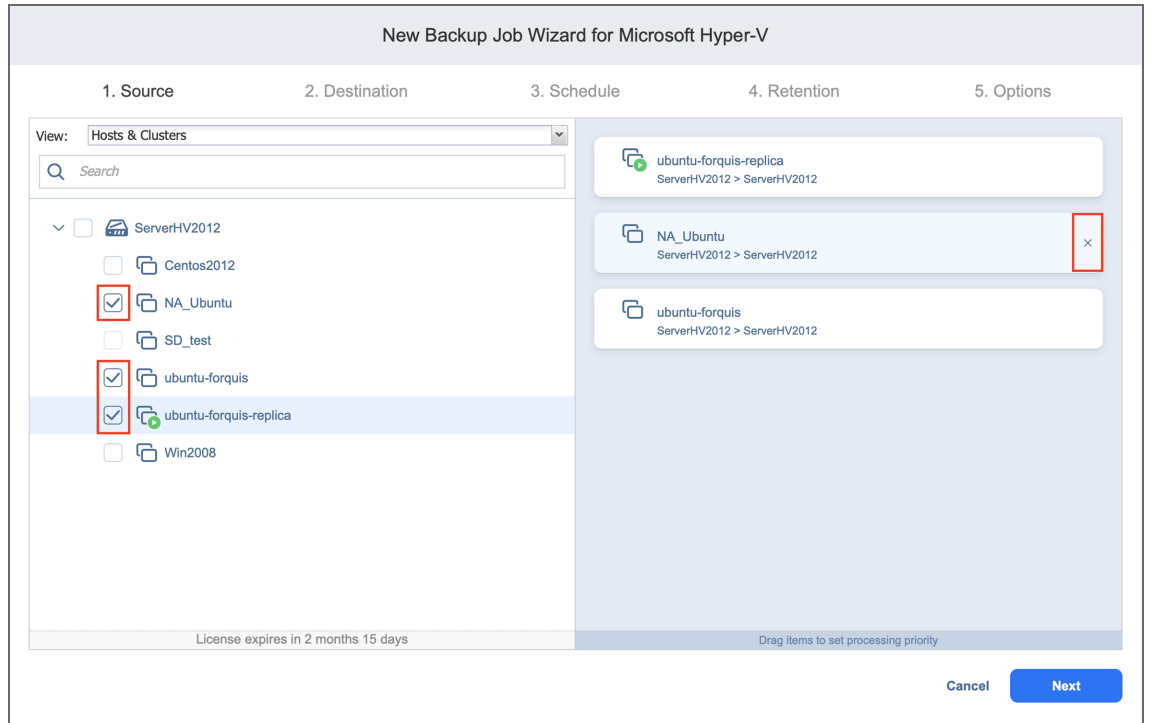
- Optionally, filter the inventory tree by entering a string to the **Search** box. You can enter a part or the whole item name.
- Select Hyper-V items by selecting the checkbox next to the item.
- The selected items appear in the right pane of the page. If necessary, reorder the selected items by dragging a VM or a container to a new position. By doing so, you can specify that you wish to back up the most important VMs first.





d. Review the list of selected Hyper-V items. If needed, remove a selected VM or a container from the backup job in either of the following ways:

- Cancel the selection of the item in the left pane. This will remove the item from the right pane.
- In the right pane, hover over the item you wish to remove and click the red “X” to the right. This will cancel the selection of the item in the left pane.



- **Policy:** If selected, this allows you to use job policies; for details, refer to [“Managing Job Policies” on page 266](#). Please follow the steps below:
  - a. If items are selected in the alternate view, a dialog 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. For details, refer to

[“Managing Policy Rules” on page 269](#) .

The screenshot shows the 'New Backup Job Wizard for Microsoft Hyper-V' interface. The wizard is on the '1. Source' step. A 'View:' dropdown menu is open, showing 'Policy' selected and highlighted with a red border. Below it, there are search filters: 'Include items if ALL rules are matched', a checkbox for 'Map new VMs to matching backups.', and 'Rule #1' with search criteria 'VM name', 'Contains', and 'win'. On the right, a 'Policy Container' is shown with 'Win2008' listed below it. At the bottom, there are 'Cancel' and 'Next' buttons.

2. Click **Next** to confirm that you wish to add the selected VMs to the backup job.

The wizard will display the next page.

## Notes

- If you cannot find a particular VM or container, try the following:
  - Make sure the corresponding Hyper-V server or cluster is added to the [inventory](#).
  - [Refresh the Inventory](#).
- If you add a Hyper-V server or clusters to the job:
  - All VMs currently available on the selected server or cluster will be backed up
  - All new VMs that will be created in (or moved to) the server or cluster in the future will be automatically added to the job and backed up.
  - All VMs are tracked during backup jobs to prevent failures resulting from live migrations within Hyper-V clusters.
- 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.
- If all disks of a given VM are unsupported, said VM will be disabled in the inventory tree, and it will not be possible to select that VM.


# Backup Job Wizard for Hyper-V: Destination

On the **Destination** page of the wizard, you can select one or multiple different Backup Repositories to back up your Hyper-V VMs.

- [Setting 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 Single Backup Repository for All VMs

To back up all VMs to a single Backup Repository, choose a Backup Repository from the **Destination** drop-down list.



## 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.

New Backup Job Wizard for Microsoft Hyper-V

1. Source      2. Destination      3. Schedule      4. Retention      5. Options

Destination: ☰ Different backup repositories ... ℹ To re-use existing backups, expand the Advanced setup and specify target backup for each VM.

**Win2008** Click to collapse

VM disks  
 win2008-quietse.vhdx: 40.0 GB (9.7 GB allocated)

Target destination  
☰ Onboard repository  
 Use existing backup as a target  
🔄 Select backup

**ubuntu-forquis-replica** Click to collapse

VM disks  
 e71abb5d-b86b-4c48-a0f6-4d6f4e072601.vhdx: 25.0 GB (5.8 GB allocated)

Target destination  
☰ S3\_Object\_Lock  
 Use existing backup as a target  
🔄 Select backup

**ubuntu-forquis** Click to collapse

VM disks  
 ubuntu-hv.vhdx: 25.0 GB (5.8 GB allocated)

Target destination  
☰ Onboard repository  
 Use existing backup as a target  
🔄 Select backup

Cancel Next

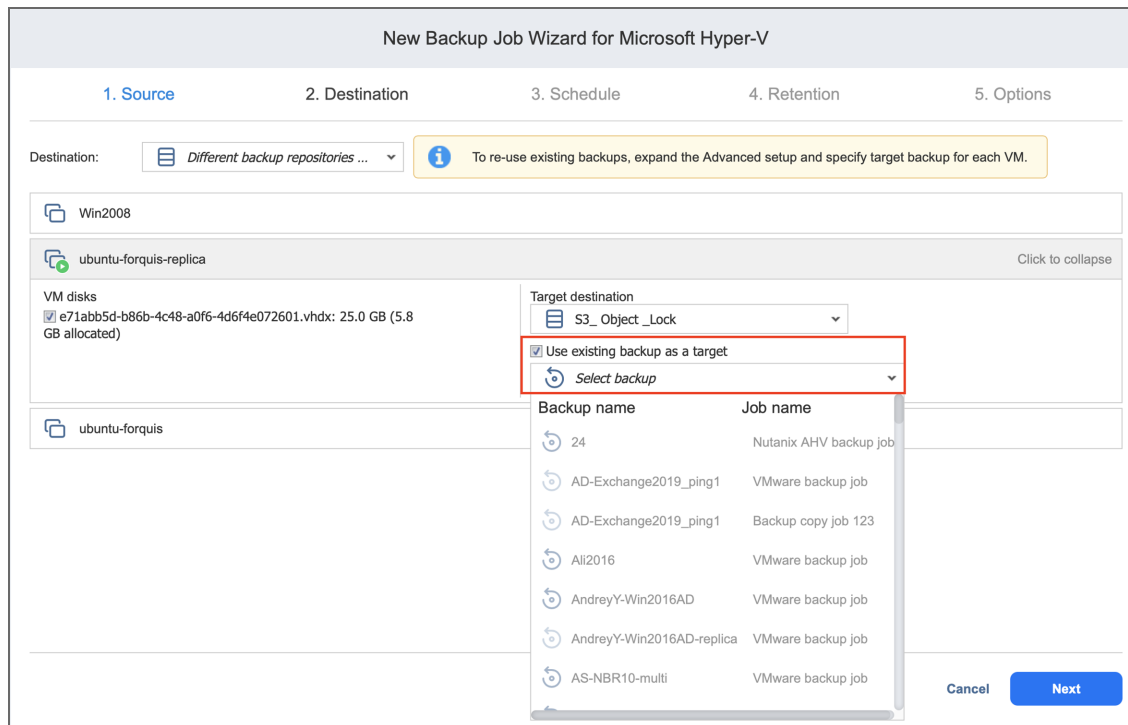
## Mapping Source VMs to Existing Backups

If you have previously backed up a VM and then lost the backup job due to accidental job deletion or because you needed to recreate jobs in a new copy of the product, you can map the source VMs to existing backups 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 **Destination** 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.

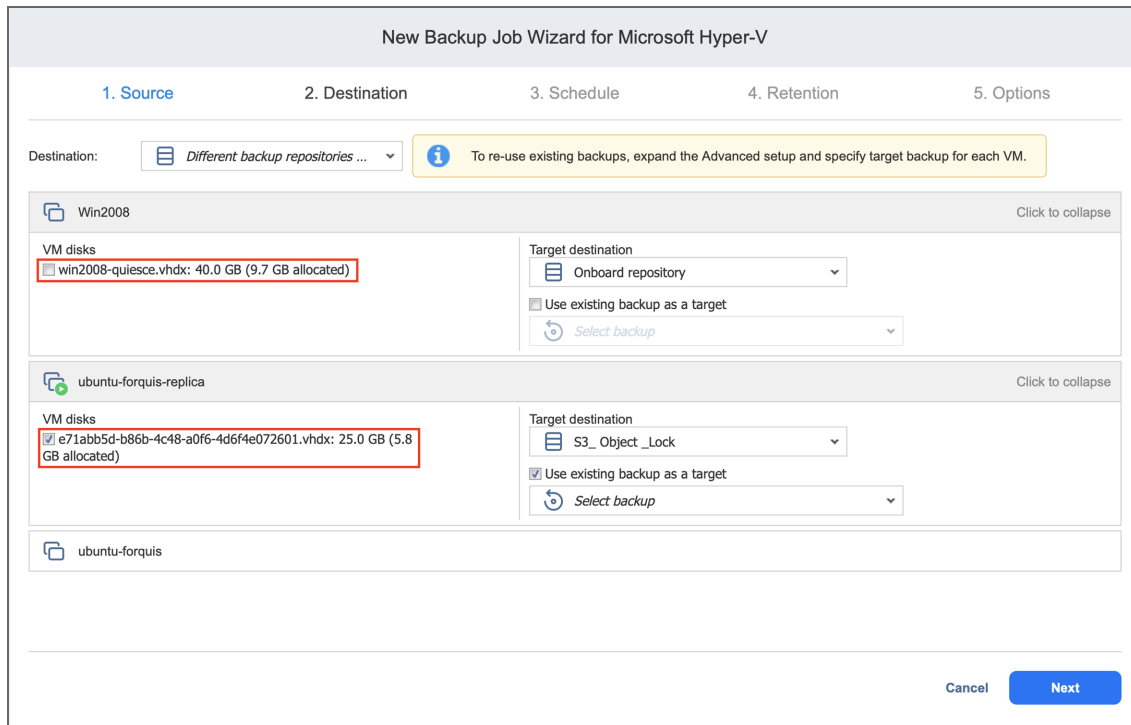


When you run the job, the product will analyze the target VM you have selected, determine how it is different from the source VM, and transfer only the differential data. VM backup mapping can be a time-consuming process that can be equal to the time required to create a full VM backup. After the job 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 would not like to back up.




## Backup Job Wizard for Hyper-V: 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.

- [Disabling Scheduled Job Execution](#)
- [Daily or Weekly Backup](#)
- [Monthly or Yearly Backup](#)
- [Periodic Backup](#)
- [Chained Backup](#)
- [Additional Schedule](#)

### 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 Backup

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.

The screenshot shows the 'New Backup Job Wizard for Microsoft Hyper-V' interface, specifically the '3. Schedule' step. The wizard has five steps: 1. Source, 2. Destination, 3. Schedule, 4. Retention, and 5. Options. In the 'Schedule' step, there are several options for scheduling the backup job. The 'Do not schedule, run on demand' option is selected. Below this, there is a time zone dropdown menu set to '(UTC+02:00, EET) Eastern European Time'. The 'Schedule #1' dropdown menu is highlighted with a red box and shows 'Run daily/weekly' selected. Below the schedule dropdown, there are fields for 'Starting at: 0:00' and 'Ending: 6:00'. There are also checkboxes for days of the week: Mon, Tue, Wed, Thu, Fri, Sat, and Sun. Below these are links for 'All days', 'Work days', and 'Weekends'. There is a field for 'every 1 weeks' and an 'Effective from' checkbox. At the bottom right, there are 'Next' and 'Cancel' buttons.

## Monthly or Yearly Backup

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.

New Backup Job Wizard for Microsoft Hyper-V

1. Source      2. Destination      3. Schedule      4. Retention      5. Options

Do not schedule, run on demand  
 (UTC+02:00, EET) Eastern European Time

Schedule #1

Run every: last Friday of every month

Starting at: 0:00    Ending: 6:00

Effective from

[Add another schedule](#)  
[Show calendar](#)

**Next**    Cancel

## 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 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.

New Backup Job Wizard for Microsoft Hyper-V

1. Source      2. Destination      3. Schedule      4. Retention      5. Options

Do not schedule, run on demand  
 (UTC+02:00, EET) Eastern European Time

Schedule #1  
 every 30 minutes

Starting at: 0:00    Ending: 6:00

Mon    Tue    Wed    Thu    Fri    Sat    Sun

[All days](#)   [Work days](#)   [Weekends](#)

Effective from

[Add another schedule](#)  
[Show calendar](#)

**Next**    Cancel

## 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.

The screenshot shows the 'New Backup Job Wizard for Microsoft Hyper-V' at the '3. Schedule' step. The wizard has five steps: 1. Source, 2. Destination, 3. Schedule, 4. Retention, and 5. Options. The 'Do not schedule, run on demand' checkbox is unchecked. The time zone is set to '(UTC+02:00, EET) Eastern European Time'. Under 'Schedule #1', the 'Run after another job' dropdown is highlighted with a red box. Below it, 'After the job:' is set to 'Physical machine backup job' and 'Run this job:' is set to 'Immediately'. There are checkboxes for 'After successful runs' (checked), 'After failed runs', and 'After stopped runs'. There is also an 'Effective from' checkbox which is unchecked. At the bottom right, there are 'Next' and 'Cancel' buttons.

## Additional Schedule

If you want to have more than one schedule for your job, click **Add another schedule** and set it up as it has been described above.

The screenshot shows the 'New Backup Job Wizard for Microsoft Hyper-V' at the '3. Schedule' step. The 'Do not schedule, run on demand' checkbox is unchecked. The time zone is set to '(UTC+02:00, EET) Eastern European Time'. Under 'Schedule #1', the 'Run periodically' dropdown is selected, with 'every 30 minutes'. The 'Starting at:' is set to '0:00' and 'Ending:' is set to '6:00'. The days of the week are checked: Mon, Tue, Wed, Thu, Fri, Sat, and Sun. There are links for 'All days', 'Work days', and 'Weekends'. There is an 'Effective from' checkbox which is unchecked. At the bottom left, the 'Add another schedule' link is highlighted with a red box. At the bottom right, there are 'Next' and 'Cancel' buttons.

# Backup Job Wizard for Hyper-V: Retention

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 S3** 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 S3** type 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:

- Only **Amazon S3** or **Local Folder** must be selected for Backup Repository Type on the Destination page of the wizard.
- If **Amazon S3** is selected as the Backup Repository, Object Lock must be enabled for the Amazon S3 bucket where your Backup Repository is located.
- For **Local Folder** type of Backup Repository, see [feature requirements](#).

Retention Settings

- Keep 10 last recovery points
- Keep one recovery point per day for 10 days
- Keep one recovery point per week for 4 weeks
- Keep one recovery point per month for 12 months
- Keep one recovery point per year for 3 years

[Learn more](#)

Immutability

- Make recovery points immutable for 10 days

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 Hyper-V: Options

On the **Options** page of the wizard, you can set up backup job options. Proceed as described in the sections below:

- [Job Options](#)
- [Full Backup Settings](#)
- [Pre and Post Job Actions](#)
- [Completing New Backup Job Wizard for Hyper-V](#)

### Job Options

In the *Job Options* section, you can give a name to the backup job and enable/disable [app-aware mode](#), change tracking, [network acceleration](#), [encryption](#), [VM verification](#), and other.

New Backup Job Wizard for Microsoft Hyper-V

1. Source    2. Destination    3. Schedule    4. Retention    5. Options

**Job Options**

Job name: Hyper-V backup job

App-aware mode: Enabled (proceed on error) [settings](#)

Use agent for OS quiescing: Disabled [i](#)

Change tracking: Use Hyper-V RCT [settings](#)

Network acceleration: Disabled [i](#)

Network encryption: Disabled [i](#)

VM verification: Disabled [i](#)

Exclude swap files and partitions: Enabled [i](#)

Exclude unused blocks: Enabled [i](#)

**Full Backup Settings**

Create full backup: Every Friday

Full backup mode: Synthetic full [i](#)

If a full backup fails, create a full backup on the next job run [i](#)

**Pre and Post Actions**

Send job run reports to [i](#)

Truncate Exchange logs: On successful VM processing only [i](#)

Truncate SQL Server logs: On successful VM processing only [i](#)

Run local pre job script [i](#)

Run local post job script [i](#)

Cancel    **Finish**    Finish & Run

## Job Name

Enter a name for the backup job.

## App-Aware Mode

If the **App-aware mode** option is enabled, NAKIVO Backup & Replication will trigger the Volume Shadow Copy Service (VSS) inside guest OS of source VMs prior to making a VM snapshot. The VSS service will instruct VSS-aware applications and databases to flush data from memory to disk and save data in a consistent state. Thus, the VM snapshot taken after triggering the VSS service will contain consistent data.

We also recommend reading the support articles about the App-aware mode requirements:

- [Hyper-V Server configuration](#)
- [Linux and FreeBSD guest OS configuration.](#)

Select one of the following options from the **App-aware mode** drop-down list:

- **Disabled:** VM backup will be performed without application awareness.
- **Enabled (fail on error):** In the dialog box that opens, select the checkboxes next to the VMs for which you want to create application-aware backups, and then select credentials next to each VM. These credentials will be used to log in to VMs you have selected and trigger the VSS service. When this option is selected, NAKIVO Backup & Replication will fail the backup if the app-aware mode fails (for example, due to wrong credentials).
- **Enabled (proceed on error):** In the dialog box that opens, select checkboxes next to the VMs for which you want to create application-aware backups, and then select credentials next to each VM. These credentials will be used to log in to the VMs you have selected and trigger the VSS service. When this option is selected, NAKIVO Backup & Replication will proceed with the backup even if the app-aware mode fails (for example, due to wrong credentials).
- Under **Settings**, you can select VMs for which the App-aware mode will be enabled. You need to specify credentials for them.

## Use Agent for OS Quiescing

When this option is enabled, NAKIVO Backup & Replication injects an agent into VMs to perform OS quiescing. This option is only available with App-aware mode enabled.

### Note

Enabling this option may generate additional data transfer for the creation of shadow copies during job runs.

## Change Tracking

Select one of the options from the **Change tracking** drop-down list:

- **Use Hyper-V RCT:** Native Hyper-V change tracking method. This option enables the fastest incremental backup.
- **Use proprietary method:** When selected, NAKIVO Backup & Replication reads the entire contents of all VM disks in order to determine what data has 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

If the **Network Acceleration** option 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.

## Encryption

If the **Encryption** option is enabled, VM 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.

You need at least one Transporter at the 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 [“VM Verification” on page 48](#).

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 VM 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. See [“Notifications & Reports” on page 294](#)) 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 VM Boot, disables networking to prevent network connections, and verifies that system start is successful.

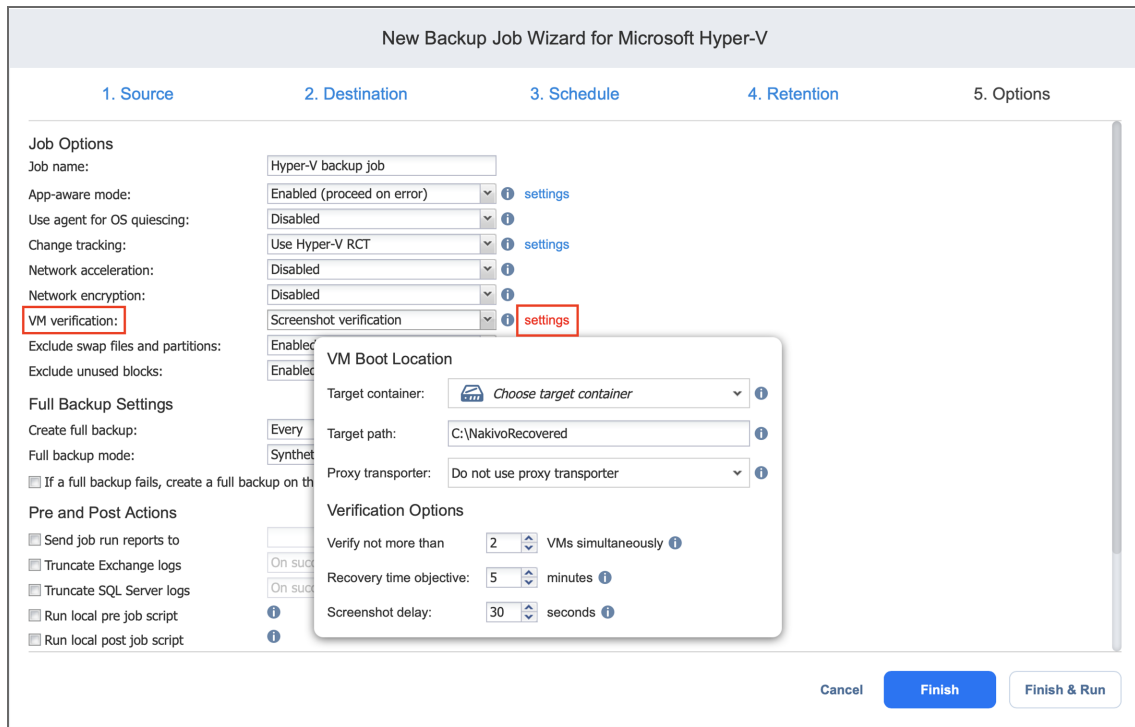
## Important

Hyper-V Integration Services must be running on the source VMs to successfully enable screenshot verification for your backup job.

After choosing **Screenshot verification**, do the following in the dialog box that opens:

1. Provide a location of the VMs to be booted:
  - **Target Container:** Choose a target container (cluster, host, or resource pool) where VMs will be run using Flash VM Boot.
  - **Target Datastore:** Choose a datastore that will host changes to the recovered VMs.
  - **Proxy transporter:** Choose a proxy Transporter from the list of available Transporters.
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.
- **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 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

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 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 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.

## Note

On a full active backup, the entire source disks will be read even if Hyper-V RCT is selected.

The screenshot shows the 'New Backup Job Wizard for Microsoft Hyper-V' dialog box, specifically the 'Options' tab. The 'Full Backup Settings' section is highlighted with a red box. It includes the following settings:

- Change tracking: Use Hyper-V RCT
- Network acceleration: Disabled
- Network encryption: Disabled
- VM verification: Disabled
- Exclude swap files and partitions: Enabled
- Exclude unused blocks: Enabled
- Full Backup Settings:**
  - Create full backup: Every Friday
  - Full backup mode: Synthetic full
  - If a full backup fails, create a full backup on the next job run
- Pre and Post Actions:**
  - Send job run reports to
  - Truncate Exchange logs (On successful VM processing only)
  - Truncate SQL Server logs (On successful VM processing only)
  - Run local pre job script
  - Run local post job script
- Data Transfer:**
  - Limit transporter load to 3 concurrent tasks
  - Bandwidth throttling: Disabled
  - Bottleneck detection

Buttons at the bottom: Cancel, Finish, Finish & Run.

## Pre and Post Job Actions

NAKIVO Backup & Replication provides you with the ability to enable certain actions before a backup job begins and after it has completed. You can choose to send job run reports, [truncate Exchange and SQL Server logs](#) on specified VMs, 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. The semi-colon character should be used to separate multiple email addresses.



## Truncation of Microsoft Exchange Server Transaction Logs

Microsoft Exchange Server database transaction logs record all changes to a Microsoft 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 you with the option of deleting (or truncating) 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 transactions which have already been committed to (available in) the Microsoft Exchange database.

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 dialog box that opens, select the checkboxes next to the VMs running Microsoft Exchange and then select credentials next to each VM. These credentials will be used to log in to VMs you have selected.

## Truncation of Microsoft SQL Server Transaction Logs

Microsoft SQL Server database transaction logs record all changes to a Microsoft 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 you with the option for deleting (or truncating) 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.

To set up Microsoft SQL Server log truncation, do the following:

1. In the *Pre and Post Actions* section, select the **Truncate SQL 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 will be used to log in to the VMs you have selected.

## 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. In the *Pre and Post Actions* section, select the **Run local pre 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 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, VM 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 backing up VMs 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 VM 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 VM backup will not be performed.

## 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. 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 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 one of the following job behaviors in relation to script completion:
    - **Wait for the script to finish:** When selected, the job will be in the "running" state until the script is completed.
    - **Do not wait for the script to finish:** When 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:** When selected, script failure will not influence the status of the job.
    - **Fail the job on script failure:** When selected and the script has failed, the job status will be set to "failed" even if the VM backup has been successful.

## Notes

1. Pre- and post-job scripts can be executed only on the machine on which the Director is installed.
2. When Integration Services are used on Hyper-V 2016 and above, custom pre/post scripts are unavailable for Windows VMs.

New Backup Job Wizard for Microsoft Hyper-V

1. Source      2. Destination      3. Schedule      4. Retention      5. Options

Change tracking: Use Hyper-V KLI settings

Network acceleration: Disabled

Network encryption: Disabled

VM verification: Disabled

Exclude swap files and partitions: Enabled

Exclude unused blocks: Enabled

**Full Backup Settings**

Create full backup: Every Friday

Full backup mode: Synthetic full

If a full backup fails, create a full backup on the next job run

**Pre and Post Actions**

Send job run reports to

Truncate Exchange logs Always

Truncate SQL Server logs Always

Run local pre job script

Run local post job script

**Data Transfer**

Limit transporter load to 3 concurrent tasks

Bandwidth throttling: Disabled

Bottleneck detection

Cancel    **Finish**    Finish & Run

## 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 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 284](#) 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 backup job:
  - a. Click the **Create New Rule** button.
  - b. The **New Bandwidth Rule** dialog box opens. Refer to [“Bandwidth Throttling” on page 284](#) 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.

## Completing New Backup Job Wizard for Microsoft Hyper-V

Click **Finish** or **Finish & Run** to complete the job creation.

### Note

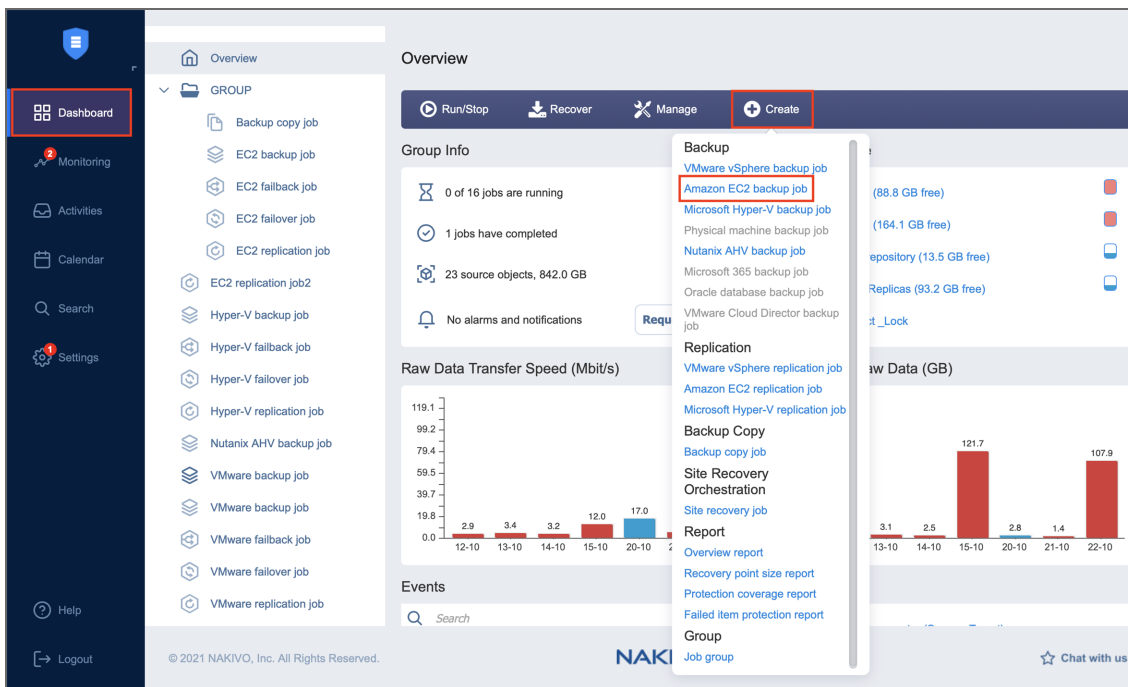
If you click **Finish & Run**, you will have to define the scope of your job. Please refer to [“Running Jobs on Demand” on page 256](#) for details.

# Creating Amazon EC2 Backup Jobs

With NAKIVO Backup & Replication, you can [back up Amazon EC2](#) instances by creating a backup job that specifies which instances should be backed up, where the backups should be located, how often the backup should occur, and what backup options should be used. To create a backup job, click **Create** on the **Dashboard**, and then click **Amazon EC2 backup job**.

### Note

You must deploy a Transporter in Amazon EC2. Refer to [“Deploying Transporters in Amazon EC2” on page 377](#).



The **New Backup Job Wizard for Amazon EC2** opens. Complete the wizard as described in the sections below:

- [“Backup Job Wizard for Amazon EC2: Source” below](#)
- [“Backup Job Wizard for Amazon EC2: Destination” on page 578](#)
- [“Backup Job Wizard for Amazon EC2: Schedule” on page 581](#)
- [“Backup Job Wizard for Amazon EC2: Retention” on page 585](#)
- [“Backup Job Wizard for Amazon EC2: Options” on page 586](#)

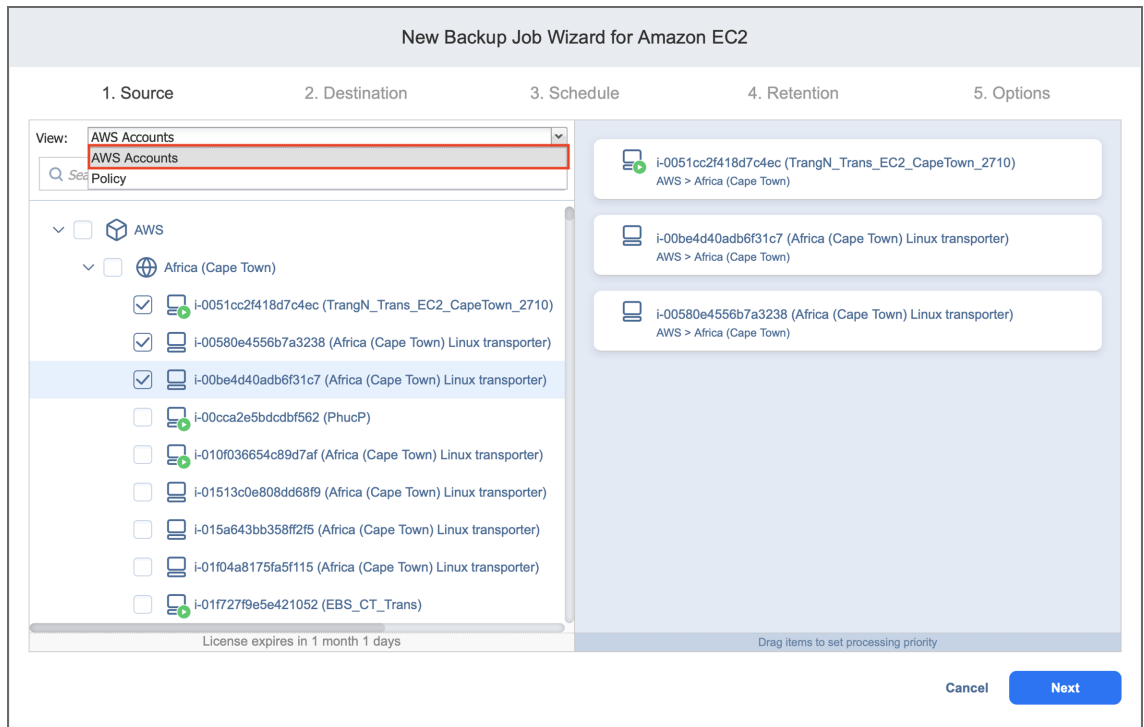
## Backup Job Wizard for Amazon EC2: Source

On the **Source** page of the wizard, you can add EC2 Instances to your backup job. Proceed as follows:

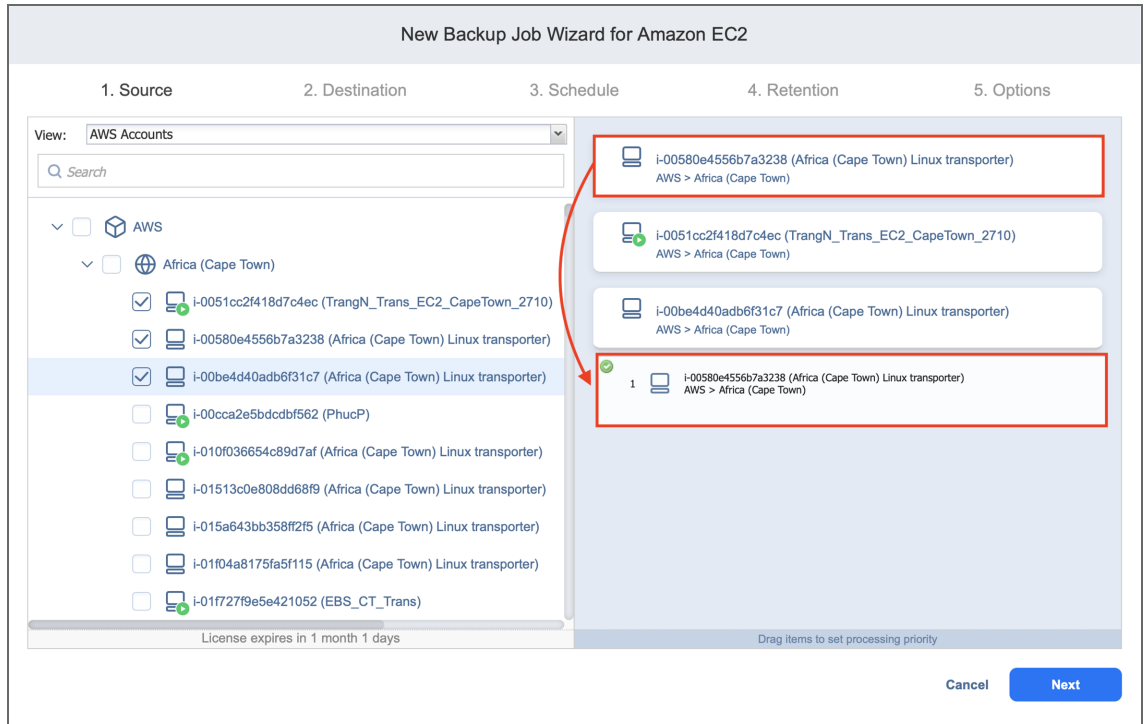
1. In the left pane of the page, choose either of the following inventory views:
  - **AWS Accounts:** If chosen, the inventory tree opens in the left pane and shows you AWS accounts along with their Regions and available Amazon EC2 instances. Proceed as follows:
    1. Optionally, filter the inventory tree by entering a string to the **Search** box. You can enter a part or the entire item name.
    2. Select Amazon EC2 items by ticking the checkbox next to the item.

### Note

If you add a paid AWS EC2 instance, you may be double-charged for that instance for the duration of the backup. For more information, refer to the [Knowledge Base article](#).

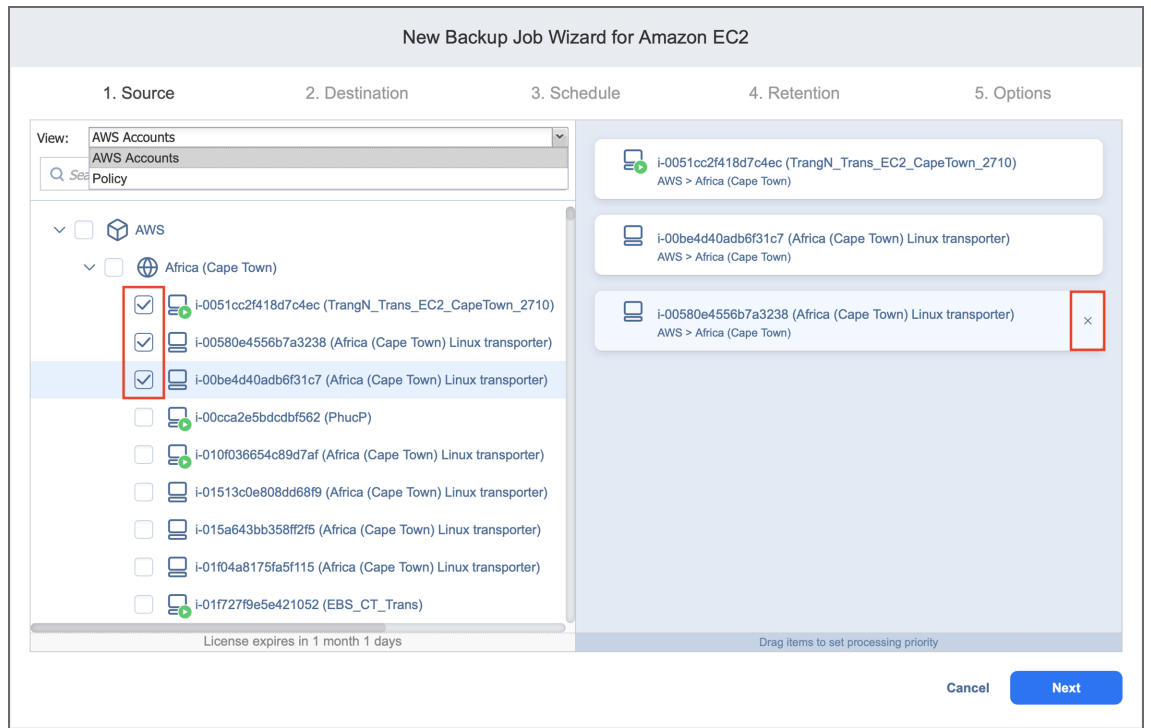


3. The selected items appear in the right pane of the page. If necessary, reorder the selected items by dragging an Instance, Region, or Account to a new position. By doing so, you can specify that you wish to back up the most important Instances first.



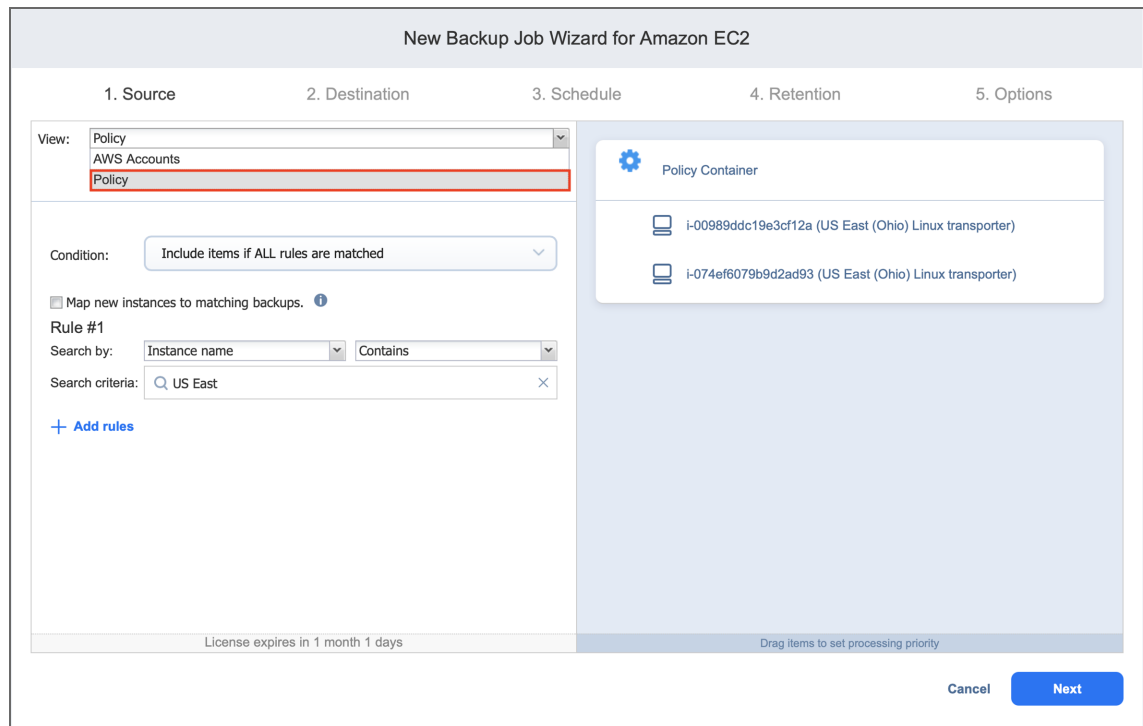
4. Review the list of selected Amazon EC2 items. If needed, remove a selected Image or a Region from the backup job in either of the following ways:

- Deselect the item in the left pane. This will remove the item from the right pane.
- In the right pane, hover over the item you want to remove and click the “x” to the right. This will deselect the item in the left pane.



- **Policy:** If selected, this allows you to use job policies; refer to [“Managing Job Policies” on page 266](#) for details. Please follow the steps below:
  - a. If items were selected in the alternate view, a dialog 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 269](#) for details.



2. Click **Next** to confirm that you wish to add selected EC2 items to the backup job.

The wizard displays the next page.

## Notes

- If you cannot find an Instance, try the following:
  - Make sure that an appropriate Amazon account (that which was used to create the Instance) is added to the [inventory](#).
  - [Refresh](#) the Inventory.
- If you add a Region or an Account to the job, the following will occur:
  - All Instances currently available in the selected Region/Account will be backed up by the job.
  - All new instances that are created in (or moved to) the Region/Account in the future will be automatically added to the backup job and backed up.
- The order in which Instances are backed is important if the Transporter performing the backup cannot process all Instances of the job simultaneously – either because the Transporter is loaded or because the number of disks in the job exceeds the Transporter’s maximum load specified during the Transporter’s creation.
- To back up Instances in a Region, at least one [Transporter must be installed](#) in that region. If there are no Transporters in a particular Region, the product will suggest that you automatically install a Transporter with the default options.



# Backup Job Wizard for Amazon EC2: Destination

On the **Destination** page, select a Backup Repository for backing up all AWS EC2 instances or select a different repository for each AWS EC2 instance.

- [Setting the Same Backup Repository for All Amazon EC2 Instances](#)
- [Setting a Different Backup Repository for Each Amazon EC2 Instance](#)
- [Mapping Instances to Existing Backups](#)
- [Excluding EBS Volumes from Backup Job](#)

## Setting the Same Backup Repository for All Amazon EC2 Instances

To back up all Amazon EC2 instances selected on the previous wizard page to the same Backup Repository, choose a Backup Repository from the **Destination** drop-down list.

The screenshot shows the 'New Backup Job Wizard for Amazon EC2' interface. The 'Destination' step is active, and the 'Destination' dropdown menu is open, showing a list of backup repositories. The 'Destination' dropdown is highlighted with a red box. The list includes:

- AWS S3: 10,240.00 TB free (100% of 1...)
- CIFS: 61.5 GB free (4% of 1.66 TB)
- Onboard repository: 135.2 GB free (68% of 200.0 ...)
- SaaS: 135.2 GB free (97% of 139.5 ...)

A yellow information box contains the text: "To re-use existing backups, expand the Advanced setup and specify target backup for each instance." The 'Next' button is highlighted in blue.

## Setting a Different Backup Repository for Each Amazon EC2 Instance

To set up different Backup Repositories for individual Amazon EC2 instances, follow the steps below:

1. Click **Advanced setup**.
2. In the instance boxes, choose a Backup Repository for each instance.

New Backup Job Wizard for Amazon EC2

1. Source      2. Destination      3. Schedule      4. Retention      5. Options

Destination: ☰ Different backup repositories ... ℹ To re-use existing backups, expand the Advanced setup and specify target backup for each instance.

**i-0051cc2f418d7c4ec (TrangN\_Trans\_EC2\_CapeTown\_2710)** Click to collapse

EBS volumes

/dev/sda1 (vol-0586d448fb2fddb84): 20.0 GB

Target destination: ☰ CIFS

Use existing backup as a target

↻ Select backup

**i-00580e4556b7a3238 (Africa (Cape Town) Linux transporter)** Click to collapse

EBS volumes

/dev/sda1 (vol-02cc813ec0014e641): 8.0 GB

Target destination: ☰ AWS S3

Use existing backup as a target

↻ Select backup

**i-00be4d40adb6f31c7 (Africa (Cape Town) Linux transporter)** Click to collapse

EBS volumes

/dev/sda1 (vol-0761875525e7bb148): 8.0 GB

Target destination: ☰ Onboard repository

Use existing backup as a target

↻ Select backup

Cancel Next

## Mapping Instances to Existing Backups

If you have previously backed up an instance and then lost the backup job due to accidental job deletion or because you need to recreate jobs in a new copy of the product, you can map the instance to its existing backup so as to avoid running the full backup again.

To map instances to existing backups, follow the steps below:

1. Click **Advanced setup**.
2. From the **Target destination** drop-down list, choose a Backup Repository that contains an existing instance backup.

3. Select the **Use existing backup as a target** option and choose an existing backup from the drop-down list.

New Backup Job Wizard for Amazon EC2

1. Source      2. Destination      3. Schedule      4. Retention      5. Options

Destination: ☰ Different backup repositories ... ℹ To re-use existing backups, expand the Advanced setup and specify target backup for each instance.

**i-0051cc2f418d7c4ec (TrangN\_Trans\_EC2\_CapeTown\_2710)** Click to collapse

EBS volumes

/dev/sda1 (vol-0586d448fb2fddb84): 20.0 GB

Target destination: ☰ CIFS

Use existing backup as a target

↻ Select backup

**i-00580e4556b7a3238 (Africa (Cape Town) Linux transporter)** Click to collapse

EBS volumes

/dev/sda1 (vol-02cc813ec0014e641): 8.0 GB

Target destination: ☰ AWS S3

Use existing backup as a target

↻ Select backup

**i-00be4d40adb6f31c7 (Africa (Cape Town) Linux transporter)** Click to collapse

EBS volumes

/dev/sda1 (vol-0761875525e7bb148): 8.0 GB

Target destination: ☰ Onboard repository

Use existing backup as a target

↻ Select backup

Cancel Next

When you run the job, the product will analyze the backup you have selected, determine how different it is from the instance, and transfer only the differential data.

### Note

Instance backup mapping can be a time-consuming process that can be equal to the time required to create a full backup.

After job completion, a new recovery point will be created and existing recovery points will not be changed or overwritten.

## Excluding EBS Volumes from the Backup Job

If you do not want to back up certain EBS volumes, you can exclude them from the backup job by following the steps below:

1. Click **Advanced setup**.
2. In the instance box, deselect the checkbox next to the EBS volumes that you do not need to back up.

## Backup Job Wizard for Amazon EC2: 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.

- [“Disabling Scheduled Job Execution” below](#)
- [“Daily or Weekly Backup” on the next page](#)
- [“Monthly or Yearly Backup” on the next page](#)
- [“Periodic Backup” on page 583](#)
- [“Chained Job” on page 583](#)
- [“Additional Schedule” on page 584](#)

### Disabling Scheduled Job Execution

If you wish to start the job manually (without any schedule), select the **Do not schedule, run on demand** checkbox.

New Backup Job Wizard for Amazon EC2

1. Source      2. Destination      3. Schedule      4. Retention      5. Options

Do not schedule, run on demand

[Next](#) [Cancel](#)

## 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:

- 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.

New Backup Job Wizard for Amazon EC2

1. Source      2. Destination      3. Schedule      4. Retention      5. Options

Do not schedule, run on demand

(UTC+02:00, EET) Eastern European Time

Schedule #1

Run daily/weekly

Starting at: 0:00    Ending: 6:00

Mon    Tue    Wed    Thu    Fri    Sat    Sun

All days   Work days   Weekends

every 1 weeks

Effective from

[Add another schedule](#)

[Show calendar](#)

[Next](#) [Cancel](#)

## Monthly or Yearly Backup

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.

The screenshot shows the 'New Backup Job Wizard for Amazon EC2' interface. It has five tabs: 1. Source, 2. Destination, 3. Schedule, 4. Retention, and 5. Options. The '3. Schedule' tab is active. Under this tab, there is a checkbox for 'Do not schedule, run on demand' which is unchecked. Below it is a time zone dropdown menu set to '(UTC+02:00, EET) Eastern European Time'. The 'Schedule #1' section has a dropdown menu with 'Run monthly/yearly' selected, which is highlighted by a red rectangular box. Below this dropdown are three more dropdown menus: 'Run every' set to 'last', 'of every' set to 'month', and 'Friday'. There are also input fields for 'Starting at: 0:00' and 'Ending: 6:00', and an unchecked 'Effective from' checkbox. At the bottom right, there are 'Next' and 'Cancel' buttons.

## 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 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.

The screenshot shows the 'New Backup Job Wizard for Amazon EC2' interface, similar to the previous one but with the 'Run periodically' option selected in the 'Schedule #1' dropdown menu, which is highlighted by a red rectangular box. The 'Run every' dropdown is set to 'every', followed by a spinner box containing '30' and a dropdown menu set to 'minutes'. The 'Starting at: 0:00' and 'Ending: 6:00' fields are present. Below these are checkboxes for days of the week: Mon, Tue, Wed, Thu, Fri, Sat, and Sun. Underneath, there are links for 'All days', 'Work days', and 'Weekends'. The 'Effective from' checkbox is unchecked. The 'Next' and 'Cancel' buttons are at the bottom right.

## 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.

New Backup Job Wizard for Amazon EC2

1. Source      2. Destination      3. Schedule      4. Retention      5. Options

Do not schedule, run on demand

(UTC+02:00, EET) Eastern European Time

Schedule #1

**Run after another job**

After the job: Physical machine backup job

Run this job: Immediately

After successful runs    After failed runs    After stopped runs

Effective from

[Add another schedule](#)

[Show calendar](#)

**Next**      Cancel

## Additional 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.

New Backup Job Wizard for Amazon EC2

1. Source      2. Destination      3. Schedule      4. Retention      5. Options

Do not schedule, run on demand

(UTC+02:00, EET) Eastern European Time

Schedule #1

Run after another job

After the job: Physical machine backup job

Run this job: Immediately

After successful runs    After failed runs    After stopped runs

Effective from

**Add another schedule**

[Show calendar](#)

**Next**      Cancel

# Backup Job Wizard for Amazon EC2: Retention

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 S3** 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 S3** type 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 instance 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:

- Only **Amazon S3** or **Local Folder** must be selected for Backup Repository Type on the Destination page of the wizard.
- If **Amazon S3** is selected as the Backup Repository, Object Lock must be enabled for the Amazon S3 bucket where your Backup Repository is located.
- For **Local Folder** type of Backup Repository, see [feature requirements](#).



1. Sources	2. Destination	3. Schedule	4. Retention
Retention Settings			
<input checked="" type="checkbox"/> Keep 10 last recovery points			
<input checked="" type="checkbox"/> Keep one recovery point per day for 10 days			
<input checked="" type="checkbox"/> Keep one recovery point per week for 4 weeks			
<input type="checkbox"/> Keep one recovery point per month for 12 months			
<input type="checkbox"/> Keep one recovery point per year for 3 years			
<a href="#">Learn more</a>			
Immutability			
<input checked="" type="checkbox"/> Make recovery points immutable for 10 days <a href="#">?</a>			

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 Amazon EC2: Options

On the **Options** page of the wizard, you can set up various options for your Amazon EC2 backup job. Proceed as described in these sections:

- [“Job Options” on the next page](#)
  - [“Job Name” on the next page](#)
  - [“Application Awareness” on the next page](#)
    - [“Application Awareness for Windows-based Amazon EC2 Instances” on the next page](#)
    - [“Application Awareness for Linux-based Amazon EC2 Instances” on page 588](#)
    - [“Pre-Conditions for Application Awareness on Linux-based Instances” on page 588](#)
    - [“Setting Application Awareness” on page 588](#)
  - [“Change Tracking” on page 588](#)
  - [“Network Acceleration” on page 589](#)
  - [“Encryption” on page 589](#)
  - [“Temporary Volumes” on page 589](#)
  - [“Exclude Swap Files and Partitions” on page 590](#)
  - [“Exclude Unused Blocks” on page 590](#)
- [“Full Backup Settings” on page 590](#)
- [“Pre and Post Actions” on page 592](#)
  - [“Email Notifications” on page 592](#)
  - [“Truncation of Microsoft Exchange Server Transaction Logs” on page 592](#)
  - [“Truncation of Microsoft SQL Server Transaction Logs” on page 593](#)
  - [“Pre and Post Job Scripts” on page 593](#)
    - [“Running a Pre-Job Script” on page 593](#)
    - [“Running a Post-Job Script” on page 594](#)
- [“Data Transfer” on page 595](#)
  - Transporter Pool

- [“Transporters” on page 595](#)
- [“Transporter Load” on page 596](#)
- [“Bandwidth Throttling” on page 596](#)
- [“Completing the New Backup Job Wizard for Amazon EC2” on page 597](#)

## Job Options

Set up general options for the job as described below.

The screenshot shows the 'New Backup Job Wizard for Amazon EC2' with five steps: 1. Source, 2. Destination, 3. Schedule, 4. Retention, and 5. Options. The '5. Options' step is active. A red box highlights the 'Job Options' section, which includes the following settings:

- Job name: EC2 backup job
- App-aware mode: Disabled
- Change tracking: Use proprietary method
- Network acceleration: Disabled
- Network encryption: Disabled
- Exclude swap files and partitions: Enabled
- Exclude unused blocks: Enabled
- Temporary volumes: Magnetic

Below the 'Job Options' section are the following settings:

- Full Backup Settings:**
  - Create full backup: Every Friday
  - Full backup mode: Synthetic full
  - If a full backup fails, create a full backup on the next job run
- Pre and Post Actions:**
  - Send job run reports to
  - Truncate Exchange logs
  - Truncate SQL Server logs
  - Run local pre job script
  - Run local post job script
- Data Transfer:**

At the bottom right, there are three buttons: 'Cancel', 'Finish', and 'Finish & Run'.

## Job Name

Enter a name for the backup job in the **Job name** box.

## Application Awareness

NAKIVO Backup & Replication provides you with the ability to create [application-consistent backups](#) of Amazon EC2 instances. The application awareness approach is different for Windows-based and Linux-based instances. Refer to the sections below for more information.

### Application Awareness for Windows-based Amazon EC2 Instances

NAKIVO Backup & Replication provides application awareness for [supported](#) Windows operating systems by using Microsoft Volume Shadow Copy (VSS) service. If application awareness is enabled for a Windows-based instance, NAKIVO Backup & Replication does the following:

1. Connects to the source instance using credentials provided in the backup job wizard.
2. Automatically deploys a NAKIVO VSS driver into the instance OS.

3. With the help of the VSS driver, it creates application-consistent snapshots of the instance EBS volumes.
4. Automatically removes the VSS driver from the instance OS.

## Application Awareness for Linux-based Amazon EC2 Instances

NAKIVO Backup & Replication provides you with the ability to run custom pre-freeze and post-thaw scripts on Linux-based instances in order to enable application awareness. A pre-freeze script is executed before a snapshot of an instance is taken, and the post-thaw script is executed after the snapshot has been taken.

### Pre-Conditions for Application Awareness on Linux-based Instances

- Create pre-freeze and post-thaw scripts for your Linux-based instances.
- Make sure that SSH is enabled on the instances.
- Place the scripts in the appropriate folders on each instance:

- **Pre-freeze:** `/usr/sbin/pre-freeze-script`
- **Post-thaw:** `/usr/sbin/post-thaw-script`

Make sure that you allow the execution of the script files.

For example: `chmod +x /usr/sbin/pre-freeze-script`

### Setting Application Awareness

Set up application awareness as follows:

1. Select one of the options from the **App-aware mode** drop-down list:
  - **Enabled (proceed on error):** If this option is selected, the instance backup will continue to run even if application awareness fails.
  - **Enabled (fail on error):** If this option is selected, instance backup will fail if application awareness fails.
  - **Disabled:** If this option is selected, application awareness will not be used for instance backup.
2. If the application awareness mode is enabled, specify the OS credentials that will be used to connect to instances:
  1. Click the **Settings** link.
  2. In the dialog box that opens, select a checkbox next to each instance which you would like to back up in application-aware mode.
  3. Select the appropriate credentials from the drop-down list next to each selected instance.

### Change Tracking

Select one of the options from the **Change tracking** drop-down list:

- **Use proprietary method:** NAKIVO Backup & Replication will read the entire contents of all VM disks in order to determine what data has changed since the last job run.
- **No change tracking (always full):** the application will always perform a full VM backup of all source VMs.

## Network Acceleration

If this option is selected, NAKIVO Backup & Replication will use compression and traffic reduction techniques to speed up the data transfer. Select this option if you plan to back up over WAN or slow LAN links. Refer to [“Network Acceleration” on page 63](#) for more details.

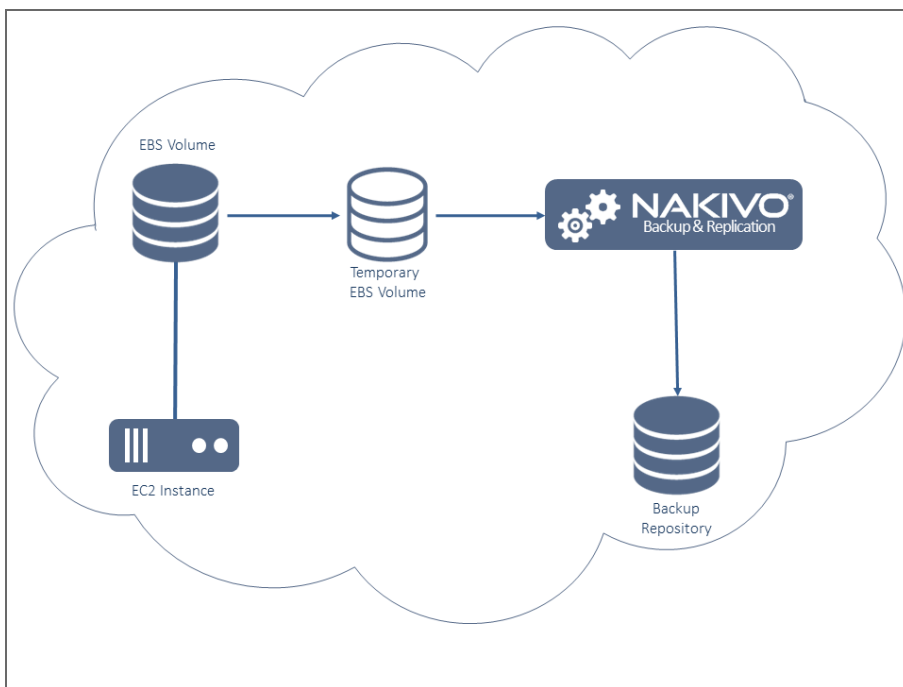
## Encryption

If 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 plan to back up over WAN without a VPN connection. For more information, refer to [“Encryption in Flight and at Rest” on page 42](#).

You need at least one Transporter at the source and target sites to enable encryption.

## Temporary Volumes

When you back up an EC2 instance, NAKIVO Backup & Replication creates temporary snapshots of the instance's EBS volumes, converts the snapshots into temporary EBS Volumes, and then reads the data from the temporary volumes and sends it to the backup repository.



By default, the type of temporary EBS volumes is set to "General Purpose SSD" to ensure higher performance. However, if you run backups frequently, this kind of storage may result in noticeably higher storage costs. To reduce costs, you can set the type of the temporary EBS Volumes to **Magnetic** disk. You can switch between those two options under the **Temporary volumes** drop-down list.

## 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 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 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 instances 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.

## New Backup Job Wizard for Amazon EC2

1. Source

2. Destination

3. Schedule

4. Retention

5. Options

### Job Options

Job name:	<input type="text" value="EC2 backup job"/>
App-aware mode:	<input type="text" value="Disabled"/> ⓘ
Change tracking:	<input type="text" value="Use proprietary method"/> ⓘ
Network acceleration:	<input type="text" value="Disabled"/> ⓘ
Network encryption:	<input type="text" value="Disabled"/> ⓘ
Exclude swap files and partitions:	<input type="text" value="Enabled"/> ⓘ
Exclude unused blocks:	<input type="text" value="Enabled"/> ⓘ
Temporary volumes:	<input type="text" value="Magnetic"/> ⓘ

### Full Backup Settings

Create full backup:	<input type="text" value="Every"/> <input type="text" value="Friday"/>
Full backup mode:	<input type="text" value="Synthetic full"/> ⓘ
<input type="checkbox"/> If a full backup fails, create a full backup on the next job run	ⓘ

### Pre and Post Actions

<input type="checkbox"/> Send job run reports to	<input type="text" value=""/>	ⓘ
<input type="checkbox"/> Truncate Exchange logs	<input type="text" value="On successful instance processing o"/>	ⓘ
<input type="checkbox"/> Truncate SQL Server logs	<input type="text" value="On successful instance processing o"/>	ⓘ
<input type="checkbox"/> Run local pre job script		ⓘ
<input type="checkbox"/> Run local post job script		ⓘ

### Data Transfer

Cancel

Finish

Finish & Run

## Pre and Post 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 instances, and run local [pre- and post- job scripts](#).

The screenshot shows the 'New Backup Job Wizard for Amazon EC2' interface. The 'Options' tab is selected, and the 'Pre and Post Actions' section is highlighted with a red box. The 'Pre and Post Actions' section includes the following options:

- Send job run reports to
- Truncate Exchange logs
- Truncate SQL Server logs
- Run local pre job script
- Run local post job script

The 'Data Transfer' section includes the following options:

- Transporter pool: Select transporter pool
- Transporters: Automatic selection
- Limit transporter load to: 3 concurrent tasks
- Bandwidth throttling: Disabled
- Bottleneck detection

At the bottom of the wizard, there are three buttons: 'Cancel', 'Finish', and 'Finish & Run'.

## Email Notifications

NAKIVO Backup & Replication can send email notifications on job completion status to specified recipients. This feature complements global notifications and provides the ability to configure notifications on the 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. The semi-colon character should be used to separate multiple email addresses.

## Truncation of Microsoft Exchange Server Transaction Logs

Microsoft Exchange Server database transaction logs record all changes to a Microsoft 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 (or truncate) Microsoft Exchange Server logs on the source instances after job completion.

The transaction logs are deleted after the job is completed so that the log files are available in the instance backup. Note that the product deletes only those transaction logs which are already committed to (available in) the Microsoft Exchange database.

To set up Microsoft Exchange log truncation, do the following:

1. In the *Pre and Post Action* section, select the **Truncate Exchange logs** option.
2. In the dialog box that opens, select the checkboxes next to the instances running Microsoft Exchange and then select credentials next to each instance. These credentials will be used to log in to instances you have selected.

## Truncation of Microsoft SQL Server Transaction Logs

Microsoft SQL Server database transaction logs record all changes to a Microsoft 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 you with the option of deleting (or truncating) Microsoft SQL Server logs on the source instances after job completion.

The transaction logs are deleted after the job is completed so that the original log records are available in the instance backup. Note that the product deletes only those transaction logs that are already committed to (available in) the Microsoft SQL Server database.

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 dialog box that opens, select checkboxes next to the instances running Microsoft SQL and then select credentials next to each instance. These credentials will be used to log in to instances you have selected.

## Pre and Post Job Scripts

NAKIVO Backup & Replication allows you to run a script before Instance backup begins (a pre-job script) and after the backup of all Instances in the job has completed (a post-job script). The scripts can be executed only on the machine on which the Director is installed. Do the following:

### Running a Pre-Job Script

To run a script before the product begins backing up Instances, 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.
3. In the dialog box that opens, specify the following parameters:
  - **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, Instance 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 backing up Instances at the same time.
- **Error handling:** Choose either of the following job behaviors in relation to scrip failure:
  - **Continue the job on script failure:** If this option is selected, the job will perform Instance 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 Instance backup will not be performed.

## Running a Post-Job Script

To run a script after the product has finished backing up all Instances, 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.
3. In the dialog box that opens, specify the following parameters:
  - **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 Instance backup has been successful.

## Data Transfer

In the *Data Transfer* section of the **Options** page, you can choose a Transporter to be used for reading data from source Instances, assign a Transporter load, and configure [bandwidth throttling](#).

The screenshot shows the 'New Backup Job Wizard for Amazon EC2' with the 'Options' tab selected. The 'Data Transfer' section is highlighted with a red box. The options in this section are:

- Transporter pool: Select transporter pool
- Transporters: Automatic selection
- Limit transporter load to: 3 concurrent tasks
- Bandwidth throttling: Disabled
- Bottleneck detection

At the bottom of the wizard, there are buttons for 'Cancel', 'Finish', and 'Finish & Run'.

## 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 Instance, 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 Instance (the target Transporter is always that which has been assigned to the Backup Repository).

By default, the product automatically determines which Transporter should be used to read data from the source Instance. However, you can manually specify which Transporters should be used for the job by choosing either of the following options for **Transporters** in the *Data Transfer* section:

- **Automatic selection:** The product will automatically determine which Transporter should be used to retrieve data from source Instances.

- **Manual - configured for all Instances:** Select this option to manually specify a single Transporter that will be used to retrieve data from all Instances.
- **Manual - configured per region:** Select this option to manually set a Transporter for each Region in Amazon EC2.

The target Transporter for the backup job will always be the Transporter assigned to the Backup Repository.

## 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 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 284](#) 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 backup job:
    1. Click the **Create New Rule** button.
    2. The **New Bandwidth Rule** dialog box opens. Refer to the [“Bandwidth Throttling” on page 284](#) 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 necessary 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 Backup Job Wizard for Amazon EC2

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 256](#) for details.

# Creating Nutanix AHV Backup Jobs

With NAKIVO Backup & Replication, you can [back up](#) Nutanix 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 occur, and what backup options should be used. To create a backup job, click **Create** on the **Dashboard** and select **Nutanix AHV backup job**.

### Important

To start creating Nutanix AHV backup jobs, you must [deploy a Transporter as Nutanix AHV Appliance](#).

The screenshot displays the NAKIVO Backup & Replication dashboard. On the left, a sidebar contains navigation options: Dashboard (highlighted with a red box), Monitoring, Activities, Calendar, Search, and Settings. The main area is titled 'Overview' and features a top navigation bar with 'Run/Stop', 'Recover', 'Manage', and 'Create' (highlighted with a red box). Below this, a 'Group Info' section shows '0 of 16 jobs are running' and '23 source objects, 842.0 GB'. A 'Raw Data Transfer Speed (Mbit/s)' bar chart shows values for 12-10, 13-10, 14-10, 15-10, and 20-10. A 'Backup' dropdown menu is open, listing various backup types, with 'Nutanix AHV backup job' highlighted in red. Other options include VMware vSphere, Amazon EC2, Microsoft Hyper-V, Physical machine, and Oracle database. The 'Replication' section lists VMware vSphere, Amazon EC2, and Microsoft Hyper-V replication jobs. The 'Backup Copy' section includes Backup copy job, Site Recovery Orchestration, and Site recovery job. The 'Report' section includes Overview report, Recovery point size report, Protection coverage report, Failed item protection report, and Group Job group. On the right, a 'Raw Data (GB)' bar chart shows values for 13-10, 14-10, 15-10, 20-10, 21-10, and 22-10. The bottom of the interface shows the NAKIVO logo, copyright information, and a 'Chat with us' button.

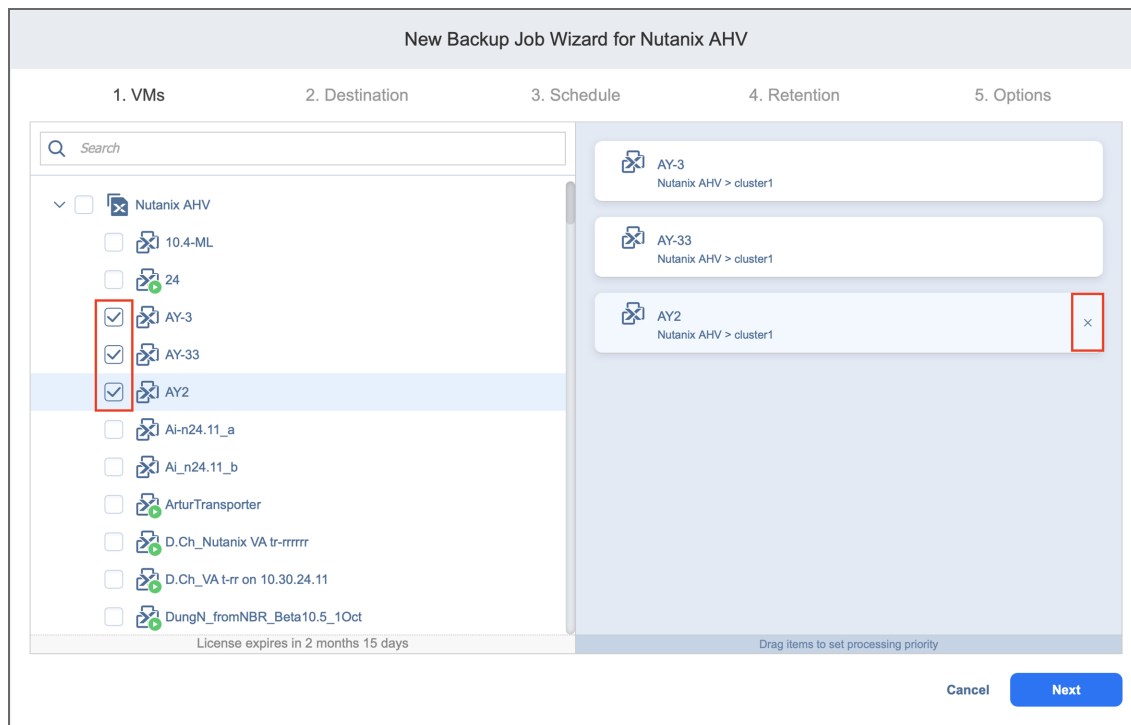
The **New Backup Job Wizard for Nutanix AHV** opens. Complete the wizard as described in the sections below:

- [“Backup Job Wizard for Nutanix AHV: VMs” below](#)
- [“Backup Job Wizard for Nutanix AHV: Destination” on the next page](#)
- [“Backup Job Wizard for Nutanix AHV: Schedule” on page 602](#)
- [“Backup Job Wizard for Nutanix AHV: Retention” on page 606](#)
- [“Backup Job Wizard for Nutanix AHV: Options” on page 607](#)

## Backup Job Wizard for Nutanix AHV: VMs

The first step in the creation of a backup job is selecting VMs you would like to back up. You can select individual VMs to be backed up by selecting the checkboxes next to their names, or you can select all VMs for backing up within the cluster by selecting a checkbox next to a cluster name. Adding a Nutanix cluster to the job ensures that important VMs are always protected. If you add a Nutanix cluster to the job, then all VMs currently available in the selected cluster will be backed up and all new VMs that are created in (or moved to) the cluster in the future will be automatically added to the job and backed up as well. If you cannot find a particular VM or cluster, make sure the corresponding Nutanix cluster is added to the Inventory or try refreshing the Inventory.

The selected VMs and clusters are displayed in the right pane of the page.



### Note

Not all VMs can be added to the selection. The following restrictions apply to:

- Some or all disks of a VM that belong to a Volume Group.
- VMs that have no disks at all.

Such VMs appear dimmed and are not available for backing up. Hovering over a VM displays the reason in a tooltip.

The following operations are available with the selected items:

- **Exclusion:** If necessary, you can exclude VMs from the backup by deselecting them in the left pane or by clicking the **Delete** icon to the right of the VM's name in the right pane.
- **Changing processing priority:** You can set the processing priority for the VM by dragging it to the higher or lower position—by doing so, you can specify to back up the most important VMs first. The processing priority of the whole cluster can also be changed by dragging. 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.

When you finish with VMs selection, click **Next** to proceed to the next page.

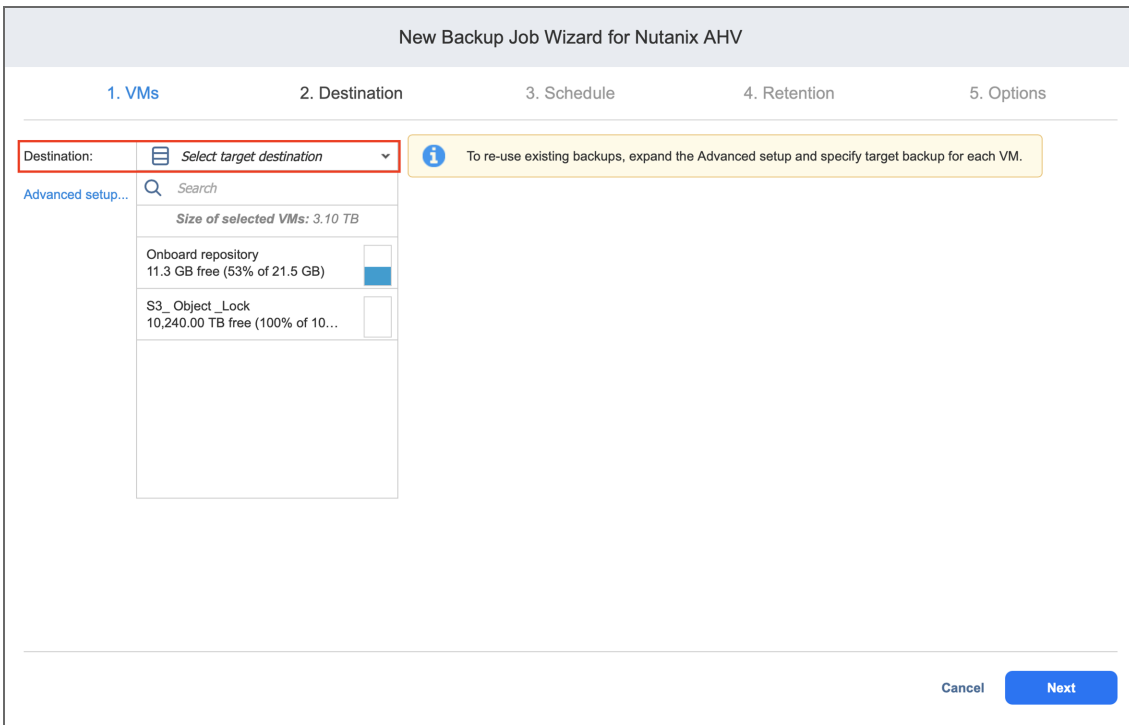
## Backup Job Wizard for Nutanix AHV: Destination

On the **Destination** page of the wizard, you can select one or multiple different Backup Repositories to backup your VMs.

- [“Setting a Single Backup Repository for All VMs” below](#)
- [“Setting Different Backup Repositories for VMs” on the next page](#)
- [“Mapping Source VMs to Existing Backups” on page 601](#)
- [“Excluding VM Disks from the Backup Job” on page 601](#)

### Setting a Single Backup Repository for All VMs

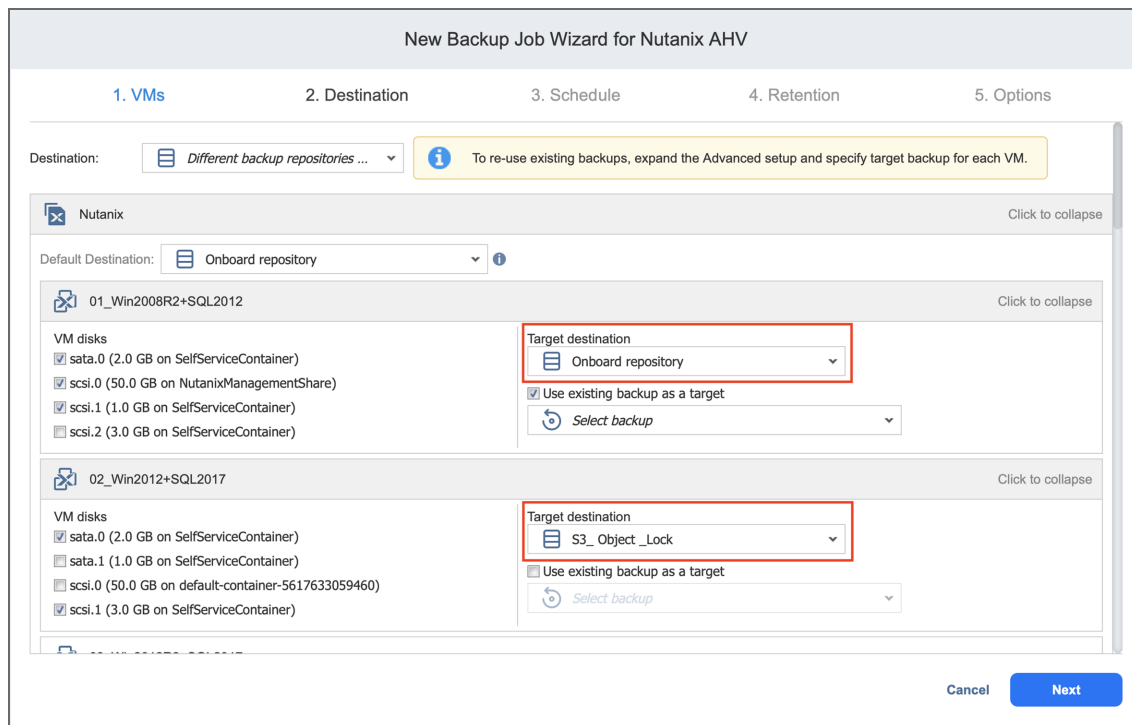
To back up all VMs selected on the previous page to a single Backup Repository, choose a Backup Repository from the **Destination** drop-down list.



## 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. Hover over the VMs and choose a **Target destination** for each VM individually.



## 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 because you needed to recreate jobs in a new copy of the product you can map source VMs to existing backups 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 **Target destination** 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.

The screenshot shows the 'New Backup Job Wizard for Nutanix AHV' interface, specifically the '2. Destination' step. The wizard is titled 'New Backup Job Wizard for Nutanix AHV' and has five steps: 1. VMs, 2. Destination, 3. Schedule, 4. Retention, and 5. Options. The 'Destination' step is active, showing a 'Destination:' dropdown set to 'Different backup repositories ...'. A yellow information box states: 'To re-use existing backups, expand the Advanced setup and specify target backup for each VM.' Below this, the 'Nutanix' section is expanded, showing a 'Default Destination:' dropdown set to 'Onboard repository'. Two VMs are listed: '01\_Win2008R2+SQL2012' and '02\_Win2012+SQL2017'. For '01\_Win2008R2+SQL2012', the 'VM disks' section lists: sata.0 (2.0 GB on SelfServiceContainer), scsi.0 (50.0 GB on NutanixManagementShare), scsi.1 (1.0 GB on SelfServiceContainer), and scsi.2 (3.0 GB on SelfServiceContainer). The 'Target destination' dropdown is set to 'Onboard repository', and the 'Use existing backup as a target' checkbox is checked. A red box highlights the 'Use existing backup as a target' option and the 'Select backup' dropdown. For '02\_Win2012+SQL2017', the 'VM disks' section lists: sata.0 (2.0 GB on SelfServiceContainer), sata.1 (1.0 GB on SelfServiceContainer), scsi.0 (50.0 GB on default-container-5617633059460), and scsi.1 (3.0 GB on SelfServiceContainer). The 'Target destination' dropdown is set to 'S3\_Object\_Lock', and the 'Use existing backup as a target' checkbox is unchecked. At the bottom right, there are 'Cancel' and 'Next' buttons.

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.

### Note

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's completion, 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 you do not want to back up.

## Backup Job Wizard for Nutanix AHV: 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.

- [“Disabling Scheduled Job Execution” below](#)
- [“Daily or Weekly Backup” on the next page](#)
- [“Monthly or Yearly Backup” on the next page](#)
- [“Periodic Backup” on page 604](#)
- [“Chained Job” on page 605](#)
- [“Additional Schedule” on page 606](#)

### Disabling Scheduled Job Execution

If you wish to start the job manually (without any schedule), select the **Do not schedule, run on demand** checkbox.

New Backup Job Wizard for Nutanix AHV

1. VMs      2. Destination      3. Schedule      4. Retention      5. Options

Do not schedule, run on demand

**Next**      Cancel

## 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:

- 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.

New Backup Job Wizard for Nutanix AHV

1. VMs      2. Destination      3. Schedule      4. Retention      5. Options

Do not schedule, run on demand

(UTC+02:00, EET) Eastern European Time

Schedule #1

**Run daily/weekly**

Starting at: 0:00      Ending: 6:00

Mon    Tue    Wed    Thu    Fri    Sat    Sun

All days   Work days   Weekends

every 1 weeks

Effective from

[Add another schedule](#)

[Show calendar](#)

**Next**      Cancel

## Monthly or Yearly Backup

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.

New Backup Job Wizard for Nutanix AHV

1. VMs      2. Destination      3. Schedule      4. Retention      5. Options

Do not schedule, run on demand  
 (UTC+02:00, EET) Eastern European Time

Schedule #1

**Run monthly/yearly**

Run every last Friday of every month

Starting at: 0:00    Ending: 6:00

Effective from

[Add another schedule](#)  
[Show calendar](#)

**Next**    Cancel

## 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 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.

New Backup Job Wizard for Nutanix AHV

1. VMs      2. Destination      3. Schedule      4. Retention      5. Options

Do not schedule, run on demand  
 (UTC+02:00, EET) Eastern European Time

Schedule #1

**Run periodically** every 30 minutes

Starting at: 0:00    Ending: 6:00

Mon    Tue    Wed    Thu    Fri    Sat    Sun  
 All days   Work days   Weekends

Effective from

[Add another schedule](#)  
[Show calendar](#)

**Next**    Cancel

## 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.

The screenshot shows the 'New Backup Job Wizard for Nutanix AHV' interface, specifically the '3. Schedule' step. The wizard has five steps: 1. VMs, 2. Destination, 3. Schedule, 4. Retention, and 5. Options. The '3. Schedule' step is active. The options are as follows:

- Do not schedule, run on demand
- (UTC+02:00, EET) Eastern European Time
- Schedule #1**
  - Run after another job** (highlighted with a red box)
- After the job: VMware backup job-test
- Run this job: Immediately
- After successful runs  After failed runs  After stopped runs
- Effective from

At the bottom, there are links for 'Add another schedule' and 'Show calendar', and buttons for 'Next' and 'Cancel'.

## Additional 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.

The screenshot shows the 'New Backup Job Wizard for Nutanix AHV' interface, specifically the '3. Schedule' step. The interface is divided into five tabs: 1. VMs, 2. Destination, 3. Schedule, 4. Retention, and 5. Options. The '3. Schedule' tab is active. It contains the following options:

- Do not schedule, run on demand
- Time zone: (UTC+02:00, EET) Eastern European Time
- Schedule #1
  - Run after another job: Run after another job
  - After the job: VMware backup job-test
  - Run this job: Immediately
  - After successful runs  After failed runs  After stopped runs
  - Effective from
- Add another schedule** (highlighted with a red box)
- Show calendar

At the bottom right, there are 'Next' and 'Cancel' buttons.

## Backup Job Wizard for Nutanix AHV: Retention

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 S3** 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 S3** type 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:

- Only **Amazon S3** or **Local Folder** must be selected for Backup Repository Type on the Destination page of the wizard.
- If **Amazon S3** is selected as the Backup Repository, Object Lock must be enabled for the Amazon S3 bucket 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 Nutanix AHV: Options

On the **Options** page of the wizard, set up job options. Proceed as described in the sections below.

- [“Job Options” on the next page](#)
  - [“Job Name” on the next page](#)
  - [“App-aware Mode” on the next page](#)
  - [“Change Tracking” on page 609](#)
  - [“Network Acceleration” on page 609](#)
  - [“Encryption” on page 610](#)
  - [“Exclude Swap Files and Partitions” on page 610](#)
- [“Full Backup Settings” on page 610](#)

- [“Pre and Post Actions” on page 611](#)
  - [“Email Notifications” on page 612](#)
  - [“Pre Job Script” on page 612](#)
  - [“Post Job Script” on page 612](#)
- [“Data Transfer” on page 613](#)
  - [“Transporter Pool” on page 613](#)
  - [“Transporters” on page 613](#)
  - [“Transporter Load” on page 614](#)
  - [“Bandwidth Throttling” on page 614](#)
- [“Completing the New Backup Job Wizard for Nutanix AHV” on page 615](#)

## Job Options

In this section, you can give a name to the backup job, enable/disable [app-aware mode](#), change tracking, [network acceleration](#), [encryption](#), and select to [exclude swap files](#), [partitions](#), and [unused blocks](#).

## Job Name

Enter the name of the backup job in the **Job name** box.

## App-aware Mode

If the **App-aware mode** option is enabled, VM backup will be performed using guest OS quiescing to ensure the consistency of application data. Please note that this feature requires the latest version of Nutanix Guest Tools to be installed on all source VMs. Processing of a VM can be set to continue or fail in case an application quiescing error is encountered.

Select one of the following options from the **App-aware mode** drop-down list:

- **Disabled:** VM backup will be performed without application awareness.
- **Enabled (fail on error):** Click **settings** and in the dialog box that opens, select checkboxes next to the VMs for which you want to create application-aware backups. When this option is selected, NAKIVO Backup & Replication will fail the backup if the app-aware mode fails.
- **Enabled (proceed on error):** Click **settings** and in the dialog box that opens, select check boxes next to the VMs for which you want to create application-aware backups. When this option is selected, NAKIVO Backup & Replication will proceed with the backup even if the app-aware mode fails.

## Change Tracking

Select one of the options from the **Change tracking** drop-down list:

- **No change tracking (always full):** If this option is selected, NAKIVO Backup & Replication will always perform a full VM backup of all source VMs.
- **Use Nutanix CRT:** If this option is selected, NAKIVO Backup & Replication will enable the Changed Region Tracking feature for source VMs. This feature enables the product to quickly identify which data regions have changed since the last job run, which significantly increases job speed. Click the **settings** link to specify additional options:
  - **On error:** Choose either of the following job behaviors in relation to Nutanix CRT error:
    - **Switch to proprietary method:** If Nutanix CRT fails to provide data on changed regions for a given VM and this option is selected, NAKIVO Backup & Replication will perform an incremental backup of that VM using the proprietary change tracking technology.
    - **Fail VM processing:** If Nutanix CRT fails to provide data on changed regions for a VM and this option is selected, NAKIVO Backup & Replication will not process that VM and will state job failure (other VMs in the job will be processed).
    - You can also specify when to switch or fail by selecting one of the available options: **immediately, on the next retry** or **on the last retry**.
  - **Double-check changed blocks provided by CRT:** If this option is selected, the changed data provided by the host will be additionally checked against the proprietary change tracking base. This approach is considered more reliable. It may perform more slowly, however, when compared to normal host-based change tracking.
  - **Use proprietary method:** If this option is selected, NAKIVO Backup & Replication will perform incremental backups using a proprietary change tracking technology. This feature requires the reading of contents of all VM disks to determine which data blocks have changed since the last job run.

## Network Acceleration

If the **Network Acceleration** option is enabled, job data will be transferred using compression and traffic reduction techniques, which can save the load on bandwidth while increasing the load on Transporters. Select this option if you transfer data over a slow WAN. Note that you need at least one Transporter on the source and target sites to enable network acceleration.



## Encryption

With the **Encryption** option 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 are planning to back up over WAN without a VPN connection.

You need at least one Transporter at source and target sites to enable encryption.

## Exclude Swap Files and Partitions

With this option enabled, NAKIVO Backup & Replication automatically excludes swap files and partitions during the backup process. This feature allows for creating smaller backups.

## Full Backup Settings

If the type of the Backup Repository that you've selected on the Destination wizard 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 VM data from the source datastore 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.

The screenshot displays the configuration window for a backup job. The 'Full Backup Settings' section is highlighted with a red box. It includes the following options:

- Job name: Nutanix AHV backup job
- App-aware mode: Enabled (proceed on error)
- Change tracking: Use Nutanix CRT
- Network acceleration: Disabled
- Network encryption: Disabled
- Exclude swap files and partitions: Enabled
- Full Backup Settings:
  - Create full backup: Every Friday
  - Full backup mode: Synthetic full
  - If a full backup fails, create a full backup on the next job run
- Pre and Post Actions:
  - Send job run reports to: admin@nakivo.com
  - Run local pre job script
  - Run local post job script
- Data Transfer:
  - Transporter pool: Select transporter pool
  - Transporters: Automatic selection
  - Limit transporter load to: 3 concurrent tasks
  - Bandwidth throttling: Disabled
  - Bottleneck detection

Buttons at the bottom: Cancel, Finish, Finish & Run.

## Pre and Post Actions

NAKIVO Backup & Replication allows you to email job run reports and [run scripts](#) before VM backup begins (a pre-job script) and after the backup of all VMs in the job has completed (a post-job script). The scripts can be executed only on the machine on which the Director is installed.

This screenshot is similar to the one above but highlights the 'Pre and Post Actions' section with a red box. The configuration is as follows:

- Job name: Nutanix AHV backup job
- App-aware mode: Enabled (proceed on error)
- Change tracking: Use Nutanix CRT
- Network acceleration: Disabled
- Network encryption: Disabled
- Exclude swap files and partitions: Enabled
- Full Backup Settings:
  - Create full backup: Every Friday
  - Full backup mode: Synthetic full
  - If a full backup fails, create a full backup on the next job run
- Pre and Post Actions:
  - Send job run reports to: admin@nakivo.com
  - Run local pre job script
  - Run local post job script
- Data Transfer:
  - Transporter pool: Select transporter pool
  - Transporters: Automatic selection
  - Limit transporter load to: 3 concurrent tasks
  - Bandwidth throttling: Disabled
  - Bottleneck detection

Buttons at the bottom: Cancel, Finish, Finish & Run.

## Email Notifications

NAKIVO Backup & Replication can send email notifications on job completion 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, select **Send job run reports to** and specify one or more email addresses in the text box. The semi-colon character should be used to separate multiple email addresses.

## 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 where the Director is installed.
2. Select the **Run local pre 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 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, VM 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 backing up VMs 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 VM 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 VM backup will not be performed.

## 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 where the Director is installed.
2. 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 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 the VM backup was successful.

## Important

Pre and post job scripts can be executed only on the machine on which the Director is installed.

## Data Transfer

In this section, you can choose a transport mode, select a transporter to be used for reading data from source VMs, specify a transporter load, and configure [bandwidth throttling](#).

The screenshot shows a configuration window for a backup job. The 'Data Transfer' section is highlighted with a red box and contains the following settings:

- Transporter pool: Select transporter pool
- Transporters: Automatic selection
- Limit transporter load to: 3 concurrent tasks
- Bandwidth throttling: Disabled
- Bottleneck detection

Other sections visible in the window include:

- Job Options:** Job name: Nutanix AHV backup job; App-aware mode: Enabled (proceed on error); Change tracking: Use Nutanix CRT; Network acceleration: Disabled; Network encryption: Disabled; Exclude swap files and partitions: Enabled.
- Full Backup Settings:** Create full backup: Every Friday; Full backup mode: Synthetic full;  If a full backup fails, create a full backup on the next job run.
- Pre and Post Actions:**  Send job run reports to: admin@nakivo.com;  Run local pre job script;  Run local post job script.

Buttons at the bottom: Cancel, Finish, Finish & Run.

## 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 Instance. However, you can manually specify which Transporters should be used for the job.

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 cluster:** Select this option to manually specify which Transporter should be used to retrieve data from each source cluster. You can select the **Automatic** option to allow the system to select the closest Transporter.
- For manually configured Transporters, clear the **Auto select replacement transporters** checkbox to explicitly lock the job to certain transporter(s)—otherwise, a replacement transporter will be picked automatically in case a Transporter is unavailable during the job start.

The target Transporter for the backup job will always be the Transporter assigned to the Backup Repository.

## 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, put the checkmark in the box next to **Limit transporter load to**.
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](#) 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](#) 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 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.

## Completing the New Backup Job Wizard for Nutanix AHV

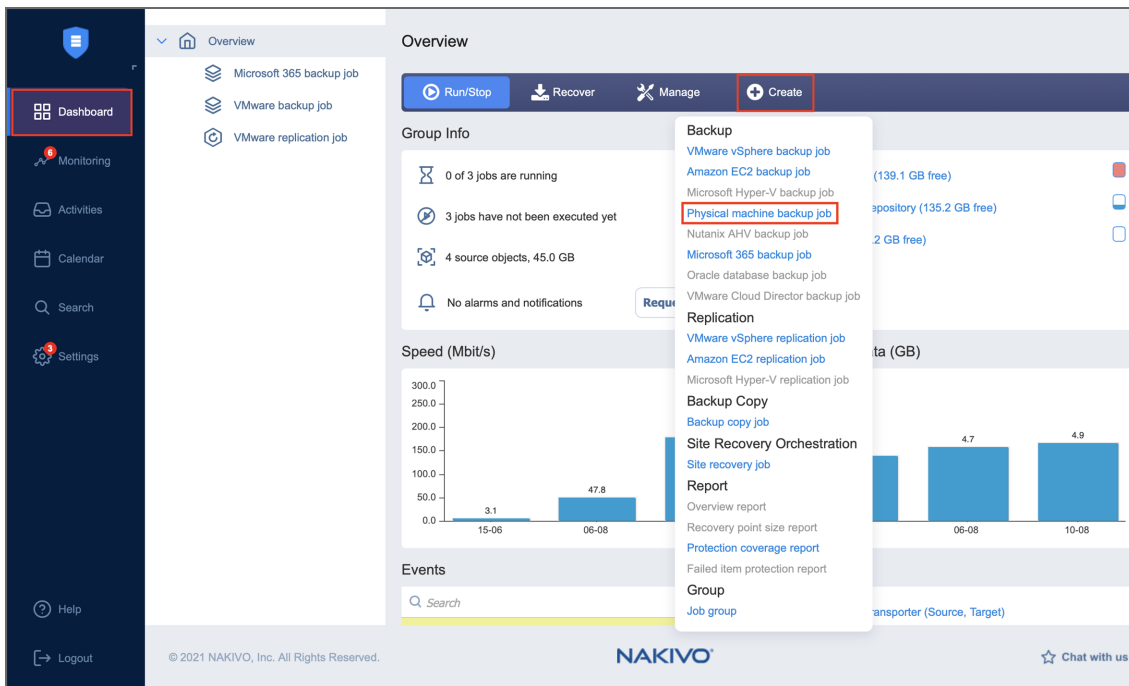
Click **Finish** or **Finish & Run** to complete the job creation.

### Note

If you click **Finish & Run**, you will have to define the scope of your job. Please refer to [Running Jobs on Demand](#) for details.

# Creating Physical Machine Backup Jobs

With NAKIVO Backup & Replication, you can **back up** both Windows and Linux physical machines. This can be done by creating a physical machine backup job that specifies which machines should be backed up, where the backups should be located, how often the backup should occur, and what backup options should be used. To create a physical machine backup job, click **Create** on the **Dashboard** and then click **Physical machine backup job**.



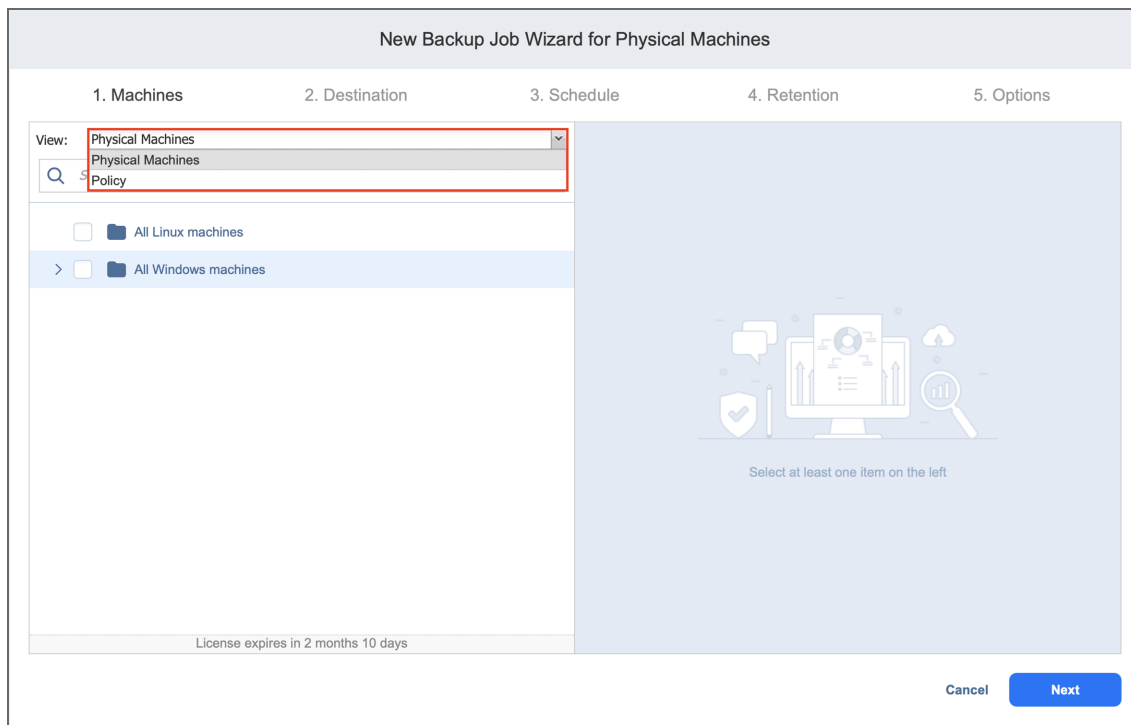
The **New Backup Job Wizard for Physical machines** opens. Complete the wizard as described in the sections below:

- [“Backup Job Wizard for Physical Machines: Machines”](#) below
- [“Backup Job Wizard for Physical Machines: Destination”](#) on page 618
- [“Backup Job Wizard for Physical Machines: Schedule”](#) on page 622
- [“Backup Job Wizard for Physical Machine: Retention”](#) on page 626
- [“Backup Job Wizard for Physical Machines: Options”](#) on page 627

## Backup Job Wizard for Physical Machines: Machines

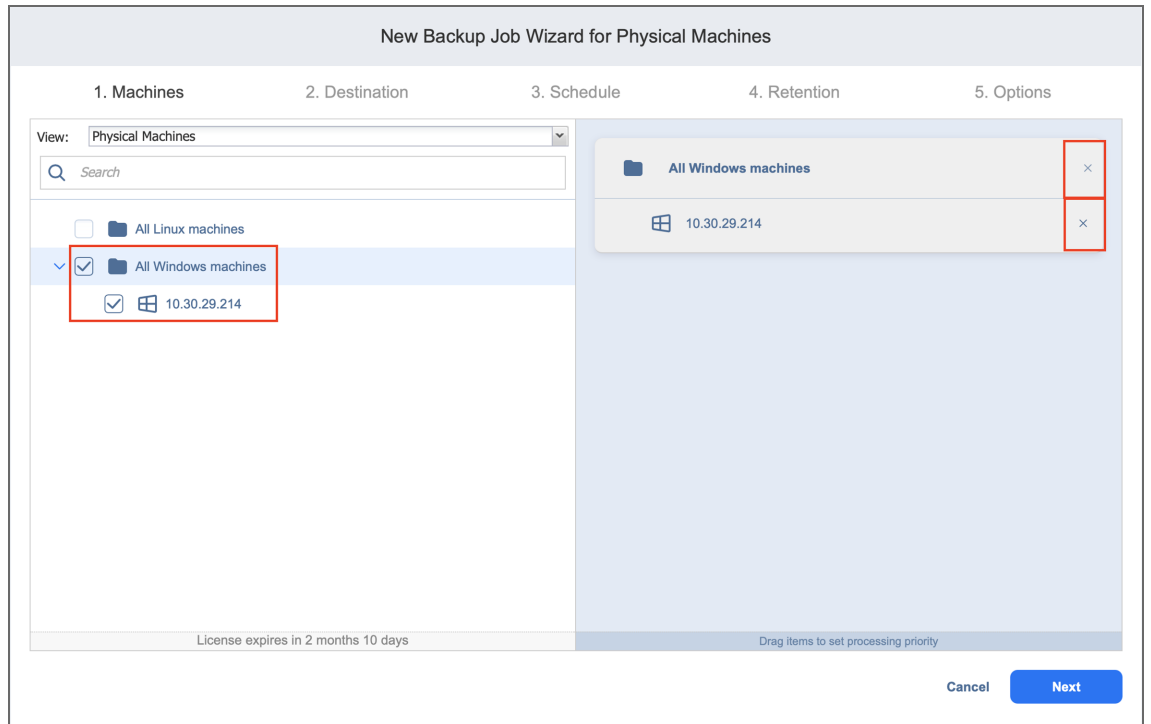
On the **Machines** page of the wizard, add a physical machine to your backup job. To do this, proceed as follows:

1. In the left pane of the page, choose either of the following inventory views:



- **Physical Machines:** If chosen, the inventory tree opens in the left pane and shows all physical machines that have been added to the inventory. Proceed as follows:
  - a. Select items by placing a checkmark next to them. The selected items will appear in the right pane of the page.
  - b. If necessary, reorder the selected items by dragging a machine to a new position. By doing so, you can specify which machines should be backed up first.
  - c. If needed, remove a selected machine from the backup job in either of the following ways:

- Cancel the selection of the object in the left pane. This will remove the object from the right pane.
- 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.



- Optionally, filter the inventory tree by entering a string to the Search box. You can enter a part or the entire item name.
- **Policy:** If selected, this allows you to use job policies; refer to [“Managing Job Policies” on page 266](#) for details. Please follow the steps below:
    - 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. Add at least one rule to the job policy. Refer to [“Managing Policy Rules” on page 269](#) for details.

The screenshot shows the 'New Backup Job Wizard for Physical Machines' interface. The wizard is currently on the '1. Machines' step. The 'View' dropdown menu is open, showing 'Policy' selected. Below the dropdown, there are options for 'Condition' (set to 'Include items if ALL rules are matched'), a checkbox for 'Map new physical machines to matching backups.', and a 'Rule #1' section. The rule is configured with 'Search by: Machine name', 'Which: Contains', and 'Search criteria: 10.'. There is an '+ Add rules' button. On the right side, there is a 'Policy Container' section with a grid icon and the IP address '10.30.29.214'. At the bottom right, there are 'Cancel' and 'Next' buttons. The 'Next' button is highlighted in blue.

2. Click **Next** to confirm that you wish to add the selected machines to the backup job.

## Notes

- If you cannot locate the necessary physical machine, try the following:
  - a. Make sure the corresponding physical machine has been added to the [inventory](#).
  - b. [Refresh](#) the inventory.
- Since Cluster Shared Volumes (CSV) are not supported, they appear dimmed and will be skipped during the job run time.

## Backup Job Wizard for Physical Machines: Destination

On the **Destination** page of the wizard, you can select one or multiple Backup Repositories to backup your physical machines.

- [Setting a Single Backup Repository for All Machines](#)
- [Setting Different Backup Repositories for Machines](#)
- [Mapping Source Machines to Existing Backups](#)
- [Excluding Physical Disks from Backup Job](#)

## Setting a Single Backup Repository for All Machines

To back up all selected machines to a single Backup Repository, choose a Backup Repository from the Backup Repository drop-down list.

New Backup Job Wizard for Physical Machines

1. Machines      2. Destination      3. Schedule      4. Retention      5. Options

Destination: Select target destination

Advanced setup... Search

Size of selected physical machines: 5.38 TB

Onboard repository  
13.0 GB free (60% of 21.5 GB)

S3\_Object\_Lock  
10,240.00 TB free (100% of 10...)

To re-use existing backups, expand the Advanced setup and specify target backup for each physical machine.

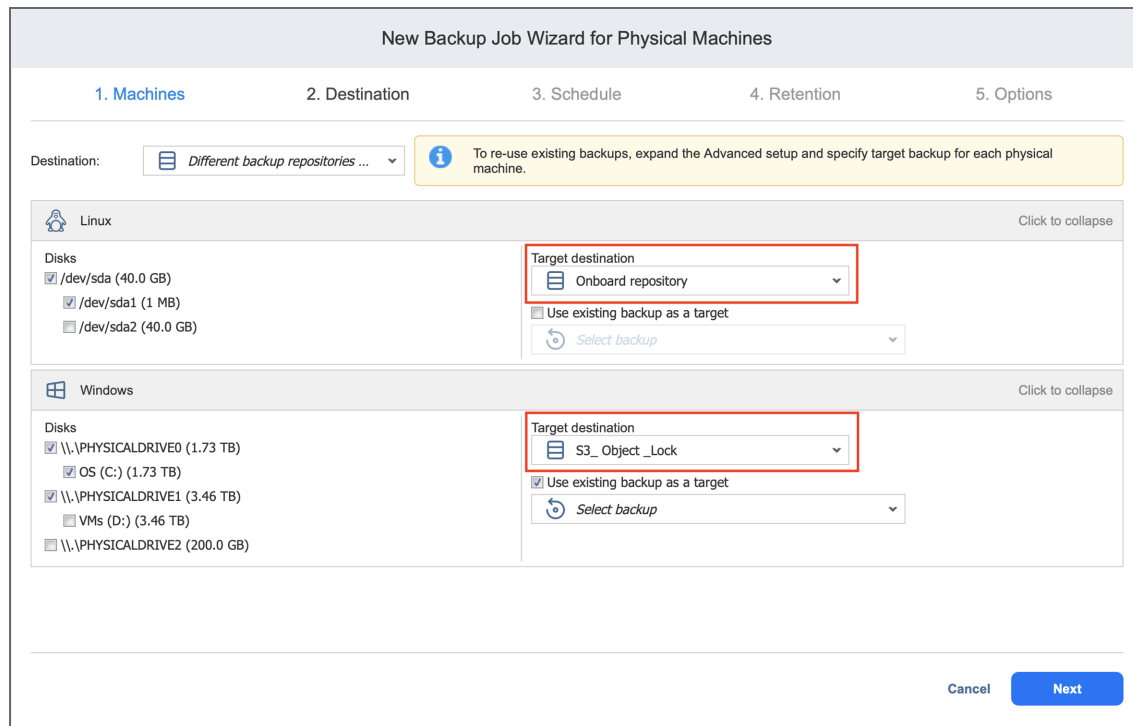
Cancel      Next

## Setting Different Backup Repositories for Machines

To back up physical machines to different Backup Repositories, follow the steps below:

1. Click **Advanced setup**.
2. Hover over a machine and click **Click to expand** to view machine details.

3. Choose the Backup Repository that you want to assign in the right pane and click **Next**.



## Mapping Source Machines to Existing Backups

If you have previously backed up a machine 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 machines to existing backups in order to avoid running full backups again.

To map source machines 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 backup.

3. Select the **Use existing backup as a target** option and choose an existing backup from the drop-down list.

New Backup Job Wizard for Physical Machines

1. Machines    2. Destination    3. Schedule    4. Retention    5. Options

Destination: ☰ Different backup repositories ... ℹ To re-use existing backups, expand the Advanced setup and specify target backup for each physical machine.

**Linux** Click to collapse

Disks

- /dev/sda (40.0 GB)
- /dev/sda1 (1 MB)
- /dev/sda2 (40.0 GB)

Target destination

☰ Onboard repository

Use existing backup as a target

↻ Select backup

**Windows** Click to collapse

Disks

- \\.\PHYSICALDRIVE0 (1.73 TB)
- OS (C:) (1.73 TB)
- \\.\PHYSICALDRIVE1 (3.46 TB)
- VMs (D:) (3.46 TB)
- \\.\PHYSICALDRIVE2 (200.0 GB)

Target destination

☰ S3\_Object\_Lock

Use existing backup as a target

↻ Select backup

Cancel Next

## Excluding Physical Disks from Backup Job

If you do not want to back up some physical disks, you can exclude them from the backup job by following the steps below:

1. Hover over a machine and click **Click to expand** to open advanced options.
2. In the machine box, clear the checkbox next to the disks you do not want to back up.

### 3. Click **Next**.

New Backup Job Wizard for Physical Machines

1. Machines    2. Destination    3. Schedule    4. Retention    5. Options

Destination: ☰ Different backup repositories ...

**i** To re-use existing backups, expand the Advanced setup and specify target backup for each physical machine.

**Linux** Click to collapse

Disks

- /dev/sda (40.0 GB)
- /dev/sda1 (1 MB)
- /dev/sda2 (40.0 GB)

Target destination

- ☰ Onboard repository
- Use existing backup as a target
- 🔄 Select backup

**Windows** Click to collapse

Disks

- \\.\PHYSICALDRIVE0 (1.73 TB)
- OS (C:) (1.73 TB)
- \\.\PHYSICALDRIVE1 (3.46 TB)
- VMs (D:) (3.46 TB)
- \\.\PHYSICALDRIVE2 (200.0 GB)

Target destination

- ☰ S3\_Object\_Lock
- Use existing backup as a target
- 🔄 Select backup

Cancel Next

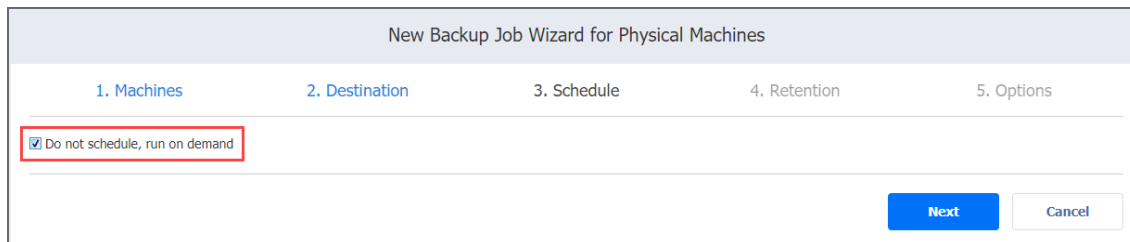
## Backup Job Wizard for Physical Machines: Schedule

On the **Schedule** page of the wizard, select whether you wish to run the backup job manually or schedule the job to run on a regular basis.

- [Disabling Scheduled Job Execution](#)
- [Daily or Weekly Backup](#)
- [Monthly or Yearly Backup](#)
- [Periodic Backup](#)
- [Chained Job](#)
- [Additional Schedule](#)

## Disabling Scheduled Job Execution

If you want to start the job manually (without any schedule), select the **Do not schedule, run on demand** checkbox.



New Backup Job Wizard for Physical Machines

1. Machines    2. Destination    3. Schedule    4. Retention    5. Options

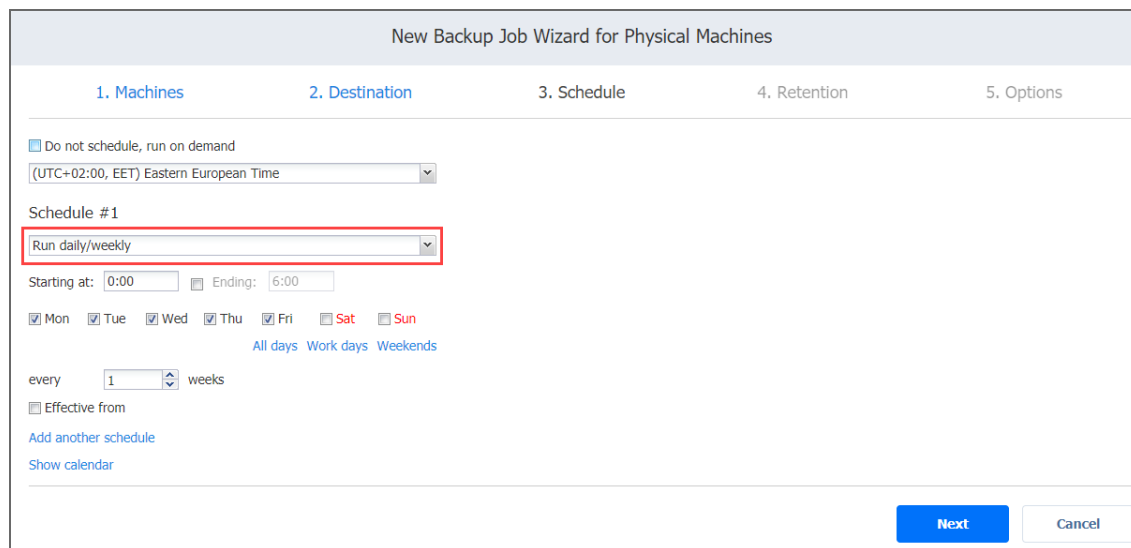
Do not schedule, run on demand

Next    Cancel

## 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:

- 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.



New Backup Job Wizard for Physical Machines

1. Machines    2. Destination    3. Schedule    4. Retention    5. Options

Do not schedule, run on demand

(UTC+02:00, EET) Eastern European Time

Schedule #1

Run daily/weekly

Starting at: 0:00    Ending: 6:00

Mon     Tue     Wed     Thu     Fri     Sat     Sun

All days    Work days    Weekends

every 1 weeks

Effective from

[Add another schedule](#)

[Show calendar](#)

Next    Cancel

## Monthly or Yearly Backup

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.

The screenshot shows the 'New Backup Job Wizard for Physical Machines' at step 3, 'Schedule'. The wizard has five steps: 1. Machines, 2. Destination, 3. Schedule, 4. Retention, and 5. Options. In the 'Schedule' step, there is a checkbox for 'Do not schedule, run on demand'. Below it is a time zone dropdown menu set to '(UTC+02:00, EET) Eastern European Time'. The 'Schedule #1' dropdown menu is highlighted with a red box and contains the option 'Run monthly/yearly'. Below this, there are dropdown menus for 'Run every' (set to 'last'), 'of every' (set to 'month'), and a day dropdown (set to 'Friday'). There are also input fields for 'Starting at' (0:00) and 'Ending' (6:00). At the bottom left, there is an 'Effective from' checkbox. At the bottom right, there are 'Next' and 'Cancel' buttons.

## 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 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.

New Backup Job Wizard for Physical Machines

1. Machines    2. Destination    3. Schedule    4. Retention    5. Options

Do not schedule, run on demand

(UTC+02:00, EET) Eastern European Time

Schedule #1

**Run periodically** every 30 minutes

Starting at: 0:00    Ending: 6:00

Mon    Tue    Wed    Thu    Fri    Sat    Sun

All days   Work days   Weekends

Effective from

[Add another schedule](#)  
[Show calendar](#)

**Next**    Cancel

## 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.

New Backup Job Wizard for Physical Machines

1. Machines    2. Destination    3. Schedule    4. Retention    5. Options

Do not schedule, run on demand

(UTC+02:00, EET) Eastern European Time

Schedule #1

**Run after another job**

After the job: Physical machine backup job

Run this job: Immediately

After successful runs    After failed runs    After stopped runs

Effective from

[Add another schedule](#)  
[Show calendar](#)

**Next**    Cancel



## Additional 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.

New Backup Job Wizard for Physical Machines

1. Machines 2. Destination 3. Schedule 4. Retention 5. Options

Do not schedule, run on demand

(UTC+02:00, EET) Eastern European Time

Schedule #1

Run after another job

After the job: Physical machine backup job

Run this job: Immediately

After successful runs  After failed runs  After stopped runs

Effective from

Add another schedule

Show calendar

Next Cancel

## Backup Job Wizard for Physical Machine: Retention

After each job run, NAKIVO Backup & Replication creates a recovery point in the Backup Repository for each physical machine. A recovery point represents the backed up physical machine as of a particular moment in time and allows you to recover individual files, application objects, or the entire machine from the Backup Repository. You can specify how many recovery points should be preserved in the Backup Repository by using the Grandfather-Father-Son (GFS) backup rotation scheme.

When **Amazon S3** 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 S3** type 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 machine 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:

- Only **Amazon S3** or **Local Folder** must be selected for Backup Repository Type on the Destination page of the wizard.
- If **Amazon S3** is selected as the Backup Repository, Object Lock must be enabled for the Amazon S3 bucket 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 Physical Machines: Options

On the **Options** page of the wizard, you can specify job options. Proceed as described in these sections:

- [Job Options](#)
- [Full Backup Settings](#)
- [Pre and Post Job Actions](#)
- [Data Transfer](#)

### Job Options

In this section, you can configure the following settings:

- **Job name:** Specify a name for the backup job.
- **App-aware mode:** If the mode is enabled, machine processing will be performed using guest OS quiescing to ensure that application data is consistent. Before enabling app-aware mode, make sure you meet the [feature requirements](#) for physical machines. In case of failure, the data gets automatically copied directly from source volumes. If the option is disabled, the product creates normal snapshots of source volumes.
- **Change tracking:** Select one of the options from the drop-down list:
  - **Use proprietary method:** If this option is selected, NAKIVO Backup & Replication will perform incremental backups using a proprietary change tracking technology. This feature requires the reading of contents of all VM disks to determine which data blocks have changed since the last job run.
  - **No change tracking (always full):** If this option is selected, NAKIVO Backup & Replication will always perform a full backup of all source machines.
- **Network acceleration:** Enable [network acceleration](#) if you transfer data over a slow WAN. Note that you need at least one Transporter on the source and target sites to enable network acceleration.
- **Encryption:** If the Encryption option is enabled, machine data will be protected with AES 256 encryption while traveling over the network. Data encryption increases the backup time and CPU load on the machines running Transporters. Select this option if you plan to back up over WAN without a VPN connection. For more details, refer to the [“Encryption in Flight and at Rest” on page 42](#) article.

The screenshot shows the 'New Backup Job Wizard for Physical Machines' interface, specifically the 'Options' step. The wizard is divided into five steps: 1. Machines, 2. Destination, 3. Schedule, 4. Retention, and 5. Options. The 'Job Options' section is highlighted with a red box and contains the following settings:

- Job name: Physical machine backup job
- App-aware mode: Disabled
- Change tracking: Use proprietary method
- Network acceleration: Disabled
- Network encryption: Disabled

Below the Job Options section, there are several other settings:

- Full Backup Settings:**
  - Create full backup: Every Friday
  - Full backup mode: Synthetic full
  - If a full backup fails, create a full backup on the next job run
- Pre and Post Actions:**
  - Send job run reports to
  - Truncate Exchange logs
  - Truncate SQL Server logs
  - Run local pre job script
  - Run local post job script
- Data Transfer:**
  - Limit transporter load to 3 concurrent tasks
  - Bandwidth throttling: Disabled
  - Bottleneck detection

At the bottom right of the wizard, there are three buttons: 'Cancel', 'Finish', and 'Finish & Run'.

## Full Backup Settings

If the 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 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 machines is lower.
  - **Active full:** If this option is selected, NAKIVO Backup & Replication will read all VM data from the source datastore 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.

The screenshot shows the 'New Backup Job Wizard for Physical Machines' at the 'Options' step. The 'Full Backup Settings' section is highlighted with a red box. The settings are as follows:

Section	Setting
Job Options	Job name: Physical machine backup job
Job Options	App-aware mode: Disabled
Job Options	Change tracking: Use proprietary method
Job Options	Network acceleration: Disabled
Job Options	Network encryption: Disabled
Full Backup Settings	Create full backup: Every Friday
Full Backup Settings	Full backup mode: Synthetic full
Full Backup Settings	<input checked="" type="checkbox"/> If a full backup fails, create a full backup on the next job run
Pre and Post Actions	<input type="checkbox"/> Send job run reports to
Pre and Post Actions	<input type="checkbox"/> Truncate Exchange logs
Pre and Post Actions	<input type="checkbox"/> Truncate SQL Server logs
Pre and Post Actions	<input type="checkbox"/> Run local pre job script
Pre and Post Actions	<input type="checkbox"/> Run local post job script
Data Transfer	<input type="checkbox"/> Limit transporter load to 3 concurrent tasks
Data Transfer	Bandwidth throttling: Disabled
Data Transfer	<input type="checkbox"/> Bottleneck detection

Buttons at the bottom: Cancel, Finish, Finish & Run

## Pre and Post Job Actions

NAKIVO Backup & Replication provides you with the ability 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 SQL Server logs](#) and run local [pre and post job scripts](#).

The screenshot shows the 'New Backup Job Wizard for Physical Machines' with the 'Options' tab selected. The 'Pre and Post Actions' section is highlighted with a red box and contains the following options:

- Send job run reports to
- Truncate Exchange logs
- Truncate SQL Server logs
- Run local pre job script
- Run local post job script

Other visible options in the wizard include:

- Job Options:** Job name (Physical machine backup job), App-aware mode (Disabled), Change tracking (Use proprietary method), Network acceleration (Disabled), Network encryption (Disabled).
- Full Backup Settings:** Create full backup (Every Friday), Full backup mode (Synthetic full), and a checkbox for 'If a full backup fails, create a full backup on the next job run'.
- Data Transfer:** Limit transporter load to (3 concurrent tasks), Bandwidth throttling (Disabled), and Bottleneck detection (unchecked).

Buttons at the bottom right: Cancel, Finish, Finish & Run.

## Email Notifications

NAKIVO Backup & Replication can send email notifications to specified recipients on job completion status. This feature complements global notifications and provides you with the ability to configure notifications on a per-job level.

### Note

To enable this option, configure your [Email settings](#).

## Truncation of Microsoft Exchange Server Transaction Logs

Microsoft Exchange Server database transaction logs record all changes in a Microsoft Exchange server database. Over time, these log files accumulate and can consume all of the available disk space if not removed periodically. NAKIVO Backup & Replication provides you with the option to delete (or truncate) Microsoft Exchange Server logs on the source machines after job completion.

The transaction logs are deleted after the job is completed so that the log files are available in the backup. Note that the product deletes only those transaction logs which are already committed to (available in) the Microsoft Exchange database.

To set up Microsoft Exchange log truncation, do the following:

1. Select the **Truncate Exchange logs** option.
2. Select one of the following options:
  - **On successful physical machine processing only**
  - **Always**
3. In the dialog box that opens, select the checkboxes next to the physical machines running Microsoft Exchange and then select the credentials next to each physical machine. These credentials will be used to log in to the physical machines that you have selected.

## Truncation of Microsoft SQL Server Transaction Logs

Microsoft SQL Server database transaction logs record all changes in a Microsoft SQL Server database. Over time, these logs accumulate and can consume all of the available disk space if not removed periodically. NAKIVO Backup & Replication provides you with the option to delete (or truncate) Microsoft SQL Server logs on the source machines after job completion. The transaction logs are deleted after job completion so that the original log records are available in the backup. Note that the product deletes only those transaction logs that are already committed to (available in) the Microsoft SQL Server database.

To set up Microsoft SQL Server log truncation, do the following:

1. Select the **Truncate SQL logs** option.
2. Select one of the following options:
  - **On successful physical machine processing only**
  - **Always**
3. In the dialog box that opens, select the checkboxes next to the physical machines running Microsoft SQL Server and then select the credentials next to each machine. These credentials will be used to log in to the physical machines that you have selected.

## Pre Job Script

To run a script before the product begins backing up the machines, 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.
3. Specify the following parameters in the dialog box that appears:
  - **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, physical machine 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 backing up machines 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 machine 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 physical machine backup will not be performed.

## Post Job Script

To run a script after the product has finished backing up all physical machines, 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.
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. 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 a physical machine 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 specify a transporter load and configure [bandwidth throttling](#).

The screenshot shows the 'New Backup Job Wizard for Physical Machines' with the 'Options' step selected. The 'Data Transfer' section is highlighted with a red box. It contains the following settings:

- Limit transporter load to: 3 concurrent tasks
- Bandwidth throttling: Disabled
- Bottleneck detection

Other sections visible include:

- Job Options:** Job name: Physical machine backup job; App-aware mode: Disabled; Change tracking: Use proprietary method; Network acceleration: Disabled; Network encryption: Disabled.
- Full Backup Settings:** Create full backup: Every Friday; Full backup mode: Synthetic full;  If a full backup fails, create a full backup on the next job run.
- Pre and Post Actions:**  Send job run reports to;  Truncate Exchange logs;  Truncate SQL Server logs;  Run local pre job script;  Run local post job script.

Buttons at the bottom: Cancel, Finish, Finish & Run.

## Transporter Load

In the *Data Transfer* section, 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. 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 284](#) 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:
    1. Click the **Create New Rule** button.
    2. The **New Bandwidth Rule** dialog box opens. Refer to the [“Bandwidth Throttling” on page 284](#) 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 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 physical machine agent engaged in the job.

## Completing the New Backup Job Wizard for Physical Machine

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 256](#) for details.

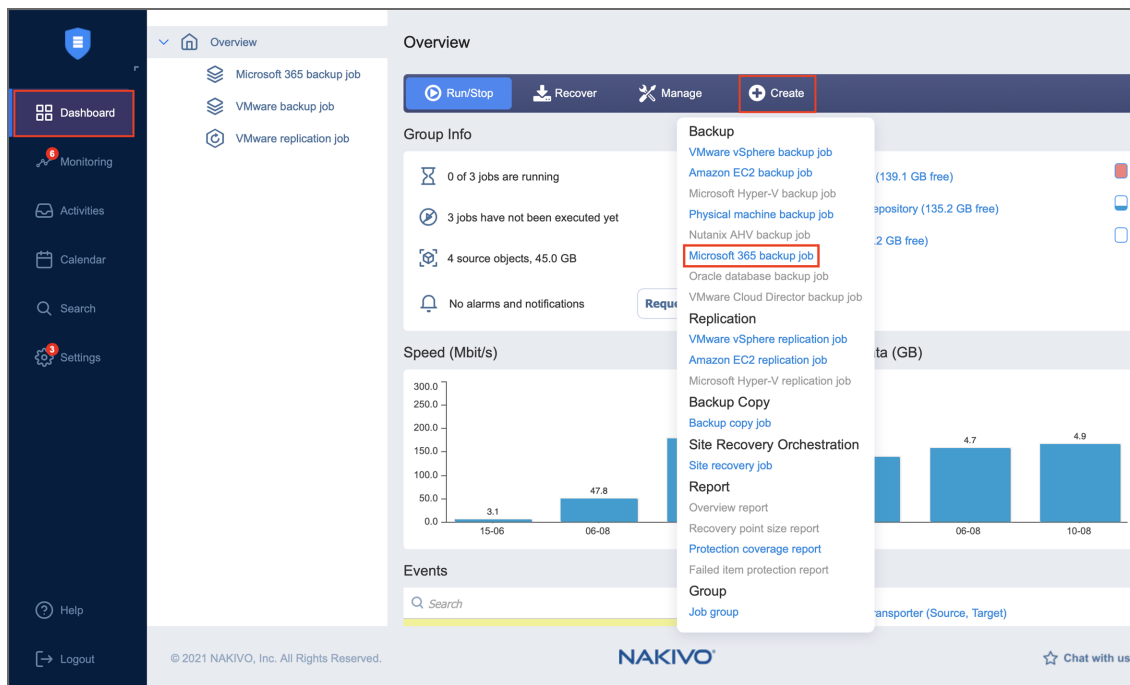
# Creating Microsoft 365 Backup Jobs

With NAKIVO Backup & Replication, you can back up an entire Microsoft 365 account or individual mailboxes, emails, folders, files, OneDrive instances, and SharePoint sites within that account. When creating a backup job for your Microsoft 365 account(s), you can specify which items to back up, where to store the backups, how often the backup job will run, and other backup options. To create a backup job, click **Create** on the **Dashboard**, and then click **Microsoft 365 backup job**.

### Note

- Before creating a Microsoft 365 backup job, you must add a [“SaaS Backup Repository” on page 422](#).
- Refer to [Required API Permissions for Microsoft 365](#) to see the list of required permissions for backing up Microsoft 365 objects.
- Refer to Microsoft 365 Platform Limitations of the latest [Release Notes](#) for the full list of Microsoft 365 items that are not supported.

- For mailbox backups, the size displayed in the Transferred Raw Data widget shows the amount of data that was actually transferred during backup. This size may differ from the total mailbox size displayed in the Job Info widget. Refer to [“Backup Job Wizard for Microsoft 365: Source”](#) below for a list of items that are skipped during the backup process.



The **New Backup Job Wizard for Microsoft 365** opens. Complete the wizard as described in the sections below:

- [“Backup Job Wizard for Microsoft 365: Source”](#) below
- [“Backup Job Wizard for Microsoft 365: Destination”](#) on the next page
- [“Backup Job Wizard for Microsoft 365: Schedule”](#) on page 638
- [“Backup Job Wizard for Microsoft 365: Retention”](#) on page 642
- [“Backup Job Wizard for Microsoft 365: Options”](#) on page 642

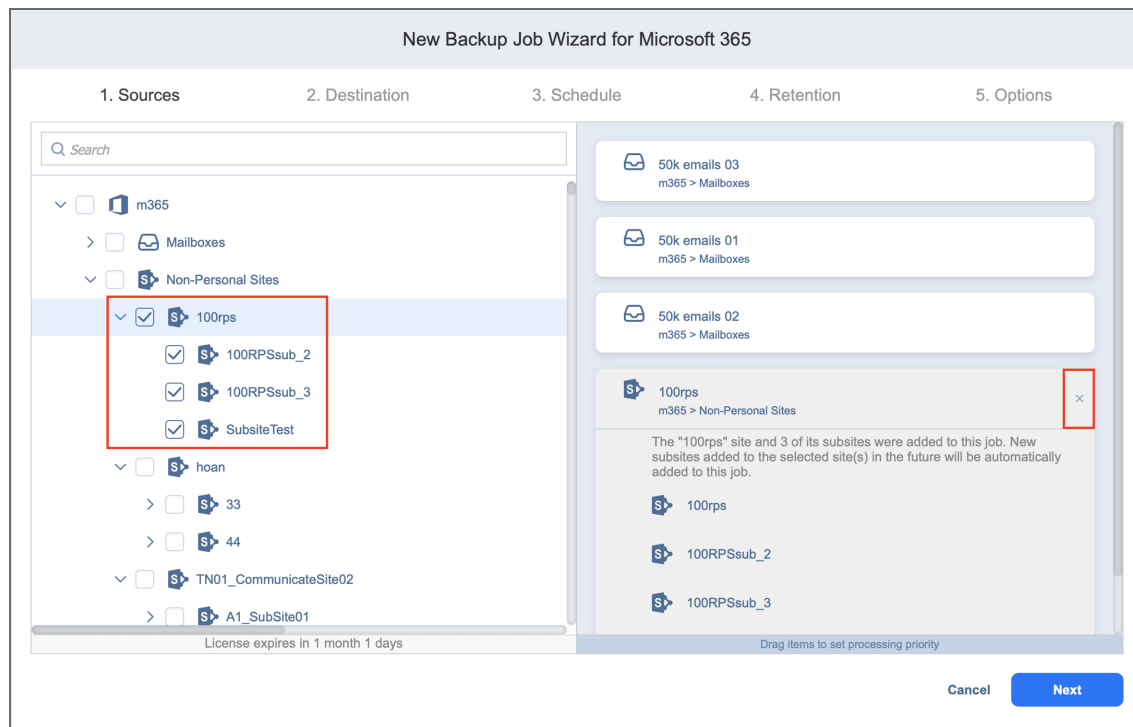
## Backup Job Wizard for Microsoft 365: Source

On the **Source** page of the wizard, add the Microsoft 365 account(s) or separate OneDrives, SharePoint sites, and mailboxes hosted in Exchange Online to your backup job. Proceed as follows:

1. In the left pane of the page, select the items you want to back up.
2. To quickly find an item, use the **Search** functionality. You can enter a part of or the full name of the item. Selected items are displayed in the right pane. You can remove items from the pane if necessary or change the priority of the item by dragging and dropping it in the required position. The priority determines the order in which the item will be processed during the job run.
3. Click **Next** to confirm the selection and go to the next page of the wizard.

### Notes

- If you select a Microsoft 365 account, all mailboxes, SharePoint sites, and OneDrives contained in that account are added to the backup job.
- If you select a group mailbox, only the group mailbox data is backed up. This does not include the data of the group members, such as user mailbox, user OneDrive, or personal site.
- If you select a group site, only the group site data is backed up. This does not include the data of the group members, such as user mailbox, user OneDrive, or personal site.
- In addition, the following mailbox items are not backed up:
  - Outbox folders
  - Calendar event messages
  - Event request messages
  - ReadOnly folders in calendars



## Backup Job Wizard for Microsoft 365: Destination

On the **Destination** page of the wizard, you can specify the storage location for the backup. You can specify a single location for all items in the backup or select different Backup Repositories for different items.

- [Setting a Single Backup Repository for All Items](#)
- [Setting a Different Backup Repository for Each Item](#)

## Setting a Single Backup Repository for All Items

To back up all items selected on the **Sources** page to a single Backup Repository, choose a Backup Repository from the **Destination** drop-down list.

The screenshot shows the 'New Backup Job Wizard for Microsoft 365' interface. The wizard is currently on step 2, 'Destination'. The 'Destination' dropdown menu is open, showing a list of backup repositories. The 'Saas' option is selected and highlighted with a red border. The list includes:

- AWS S3: 10,240.00 TB free (100% of 1...)
- CIFS: 61.5 GB free (4% of 1.66 TB)
- Onboard repository: 135.2 GB free (68% of 200.0 ...)
- Saas: 135.2 GB free (97% of 139.5 ...)

At the bottom right of the wizard, there are 'Cancel' and 'Next' buttons.

## Setting a Different Backup Repository for Each Item

To back up the items to different Backup Repositories, follow the steps below:

1. Click **Advanced setup** and do one of the following:
  - If you have selected a Microsoft 365 account(s) on the **Source** page, the account block is displayed.
    - a. Click the Microsoft 365 name to expand it and view all of the mailboxes, OneDrive instances, and SharePoint sites.
    - b. In the **Default Destination** drop-down list, select the Backup Repository for storing the backups of all services within the Microsoft account(s).
  - If you need a specific mailbox(es), OneDrive instance or SharePoint site to be stored in a different Backup Repository, click the name of the service and select a different location from the **Destination** drop-down list.

2. Click **Next** to go to the next page of the wizard.

New Backup Job Wizard for Microsoft 365

1. Sources      2. Destination      3. Schedule      4. Retention      5. Options

Destination:

50k emails 03 Click to collapse

Destination:

50k emails 01 Click to collapse

Destination:

50k emails 02 Click to collapse

Destination:

100rps Click to collapse

Destination:

## Backup Job Wizard for Microsoft 365: 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.

- [Disabling Scheduled Job Execution](#)
- [Daily or Weekly Backup](#)
- [Monthly or Yearly Backup](#)
- [Periodic Backup](#)
- [Chained Job](#)
- [Additional Schedule](#)

### Disabling Scheduled Job Execution

If you wish to start the job manually (without any schedule), select the **Do not schedule, run on demand** checkbox.

New Backup Job Wizard for Microsoft 365

1. Sources      2. Destination      3. Schedule      4. Retention      5. Options

Do not schedule, run on demand

## 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:

- 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.

The screenshot shows the 'New Backup Job Wizard for Microsoft 365' with five steps: 1. Sources, 2. Destination, 3. Schedule, 4. Retention, and 5. Options. In the '3. Schedule' step, the 'Do not schedule, run on demand' checkbox is checked. The time zone is set to '(UTC+02:00, EET) Eastern European Time'. Under 'Schedule #1', the 'Run daily/weekly' option is selected in a dropdown menu, which is highlighted with a red border. Below this, the 'Starting at' time is 0:00 and the 'Ending' time is 6:00. The days of the week are selected as Mon, Tue, Wed, Thu, and Fri, with Sat and Sun unselected. There are links for 'All days', 'Work days', and 'Weekends'. The frequency is set to 'every 1 weeks'. There is an 'Effective from' checkbox which is unchecked. At the bottom right, there are 'Next' and 'Cancel' buttons.

## Monthly or Yearly Backup

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.

The screenshot shows the 'New Backup Job Wizard for Microsoft 365' at step 3, 'Schedule'. The wizard has five steps: 1. Sources, 2. Destination, 3. Schedule, 4. Retention, and 5. Options. Under 'Do not schedule, run on demand', the time zone is set to '(UTC+02:00, EET) Eastern European Time'. In the 'Schedule #1' section, the dropdown menu is set to 'Run monthly/yearly', which is highlighted with a red box. Below this, 'Run every' is set to 'last' of every 'month' on 'Friday'. The 'Starting at' time is '0:00' and the 'Ending' time is '6:00'. The 'Effective from' checkbox is checked. At the bottom right, there are 'Next' and 'Cancel' buttons.

## 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 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.

The screenshot shows the 'New Backup Job Wizard for Microsoft 365' at step 3, 'Schedule'. The wizard has five steps: 1. Sources, 2. Destination, 3. Schedule, 4. Retention, and 5. Options. Under 'Do not schedule, run on demand', the time zone is set to '(UTC+02:00, EET) Eastern European Time'. In the 'Schedule #1' section, the dropdown menu is set to 'Run periodically', which is highlighted with a red box. Below this, 'every' is set to '30' 'minutes'. The 'Starting at' time is '0:00' and the 'Ending' time is '6:00'. The days of the week are selected as Mon, Tue, Wed, Thu, and Fri. The 'Effective from' checkbox is unchecked. At the bottom right, there are 'Next' and 'Cancel' buttons.

## 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.

The screenshot shows the 'New Backup Job Wizard for Microsoft 365' at the '3. Schedule' step. The wizard has five steps: 1. Sources, 2. Destination, 3. Schedule, 4. Retention, and 5. Options. The 'Do not schedule, run on demand' checkbox is unchecked. The time zone is set to '(UTC+02:00, EET) Eastern European Time'. Under 'Schedule #1', the 'Run after another job' dropdown is highlighted with a red box. Below it, 'After the job:' is set to 'Select a value', 'Run this job:' is set to 'Immediately', and the 'After successful runs' checkbox is checked. There are also options for 'After failed runs', 'After stopped runs', and 'Effective from'. At the bottom right, there are 'Next' and 'Cancel' buttons.

## Additional 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.

This screenshot is identical to the previous one, showing the 'New Backup Job Wizard for Microsoft 365' at the '3. Schedule' step. In this version, the 'Add another schedule' link is highlighted with a red box, indicating the next step in the process.



# Backup Job Wizard for Microsoft 365: Retention

After each job run, NAKIVO Backup & Replication creates a recovery point for each item in the Backup Repository. A recovery point represents the backed up Microsoft 365 account or mailboxes as of a particular moment in time and allows you to recover individual objects or the entire account from the Backup Repository. You can specify how many recovery points should be preserved in the Backup Repository using the Grandfather-Father-Son (GFS) backup rotation scheme. Use the following options:

- **Keep X last recovery points:** Keeps the specified number of last recovery points for each item 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.

New Backup Job Wizard for Microsoft 365

1. Sources    2. Destination    3. Schedule    4. Retention    5. Options

Keep 10 last recovery points

Keep one recovery point per day for 10 days

Keep one recovery point per week for 4 weeks

Keep one recovery point per month for 12 months

Keep one recovery point per year for 3 years

[Learn more](#)

[Next](#)    [Cancel](#)

# Backup Job Wizard for Microsoft 365: Options

On the **Options** page, you can specify the job's name, define pre- and post-job actions, and limit the Transporter load. Proceed as described in these sections:

- [Job Options](#)
- [Pre and Post Job Actions](#)
  - [Email Notifications](#)
  - [Pre Job Script](#)
  - [Post Job Script](#)
- [Data Transfer](#)
  - [Transporter Load](#)
- [Completing the New Backup Job Wizard for Microsoft 365](#)

## Job Options

Here you can do the following:

- Enter a name for the job to the Job name box.
- Optionally, select **Back up group mailbox files** option to back up the files connected to the Microsoft 365 Group. Note that doing this may cause some files to be duplicated if you are also backing up the connected Group Site.

### Note

- If NAKIVO Backup & Replication is updated from a version that did not include support for Microsoft 365 Groups to a version that does, the Back up Group mailboxes files option is not selected automatically.
- If other Microsoft 365 items were selected in addition to Exchange Online Group mailboxes , **Backup Group mailboxes files** option can still be selected and applies only to group mailboxes.

The screenshot shows the 'New Backup Job Wizard for Microsoft 365' interface, specifically the 'Options' step. The wizard is divided into five steps: 1. Sources, 2. Destination, 3. Schedule, 4. Retention, and 5. Options. The 'Options' step is currently active. In the 'Job Options' section, the 'Job name' field is set to 'Microsoft 365 backup job'. The 'Backup Group mailboxes files' checkbox is checked. Below this, the 'Pre and Post Actions' section includes options for sending job run reports to a specified email address, and running local pre and post job scripts. The 'Data Transfer' section includes a checkbox for 'Limit transporter load to' with a value of 3 concurrent tasks. At the bottom right, there are three buttons: 'Cancel', 'Finish', and 'Finish & Run'.

## Pre and Post Job Actions

With NAKIVO Backup & Replication, you can enable certain actions before a backup job begins and after it is completed. You can choose to send job run reports and run local pre- and post-job scripts. For more information, refer to [“Pre and Post Job Scripts” on page 68](#).

New Backup Job Wizard for Microsoft 365

1. Sources      2. Destination      3. Schedule      4. Retention      5. Options

Job Options  
Job name:

Backup Group mailboxes files ⓘ

Pre and Post Actions

Send job run reports to

Run local pre job script ⓘ

Run local post job script ⓘ

Data Transfer

Limit transporter load to  concurrent tasks ⓘ

Cancel   **Finish**   Finish & Run

## Email Notifications

NAKIVO Backup & Replication can send email notifications on job completion status to specified recipients. This feature complements global notification and allows you to configure notifications on a per-job level.

### Note

To enable this option, configure your email settings. For details, refer to [“Notifications & Reports” on page 294](#).

To send email notifications, select **Send job run reports to** and specify one or more email addresses in the text box. Use the semi-colon character (;) to separate multiple email addresses.

## Pre Job Script

To run a script before the product begins backing up your items, 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 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. The 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 backup is started only 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 your items 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 product performs the backup job even if the script has failed.
- **Fail the job on script failure:** When this option is selected and the script fails, the job fails and the backup is not performed.

## Post Job Script

To run a script after the product has finished backing up all items, 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.
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. The 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 remains 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 backup is successful.

### Note

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, you can specify the Transporter load.

New Backup Job Wizard for Microsoft 365

1. Sources      2. Destination      3. Schedule      4. Retention      5. Options

Job Options

Job name:

Backup Group mailboxes files ⓘ

Pre and Post Actions

Send job run reports to  ⓘ

Run local pre job script ⓘ

Run local post job script ⓘ

Data Transfer

Limit transporter load to  concurrent tasks ⓘ

Cancel   **Finish**   Finish & Run

## Transporter Load

You can limit the maximum number of Transporter tasks used by a job. By default, the number is set to 3 concurrent tasks. For a Microsoft 365 backup job, one task is equal to processing one mailbox or one OneDrive instance.

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.

## Completing the New Backup Job Wizard for Microsoft 365

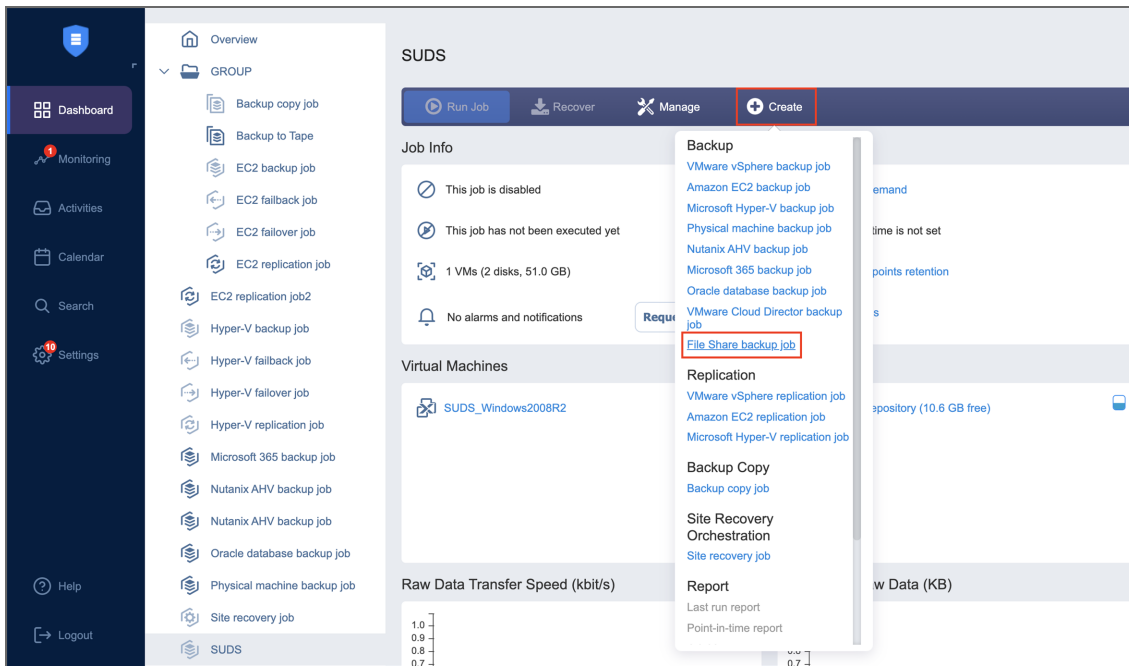
Click **Finish** or **Finish & Run** to complete the job creation.

### Note

If you click **Finish & Run**, you will have to define the scope of your job.

# Creating File Share Backup Jobs

With NAKIVO Backup & Replication, you can back up NFS and CIFS share folders located on supported NAS devices. To do this, you should create a file share backup job that specifies the file shares to back up, the target location for the backups, backup frequency, and backup options. To create a file share backup job, click **Create** on the **Dashboard** and then click **File Share backup job**.



The **New Backup Job Wizard for File Share** opens. Complete the wizard as described in the sections below:

- [“Backup Job Wizard for File Share: Source” below](#)
- [Backup Job Wizard for File Share: Destination](#)
- [“Backup Job Wizard for File Share: Schedule” on page 650](#)
- [“Backup Job Wizard for File Share: Retention” on page 655](#)
- [“Backup Job Wizard for File Share: Options” on page 655](#)

## Backup Job Wizard for File Share: Source

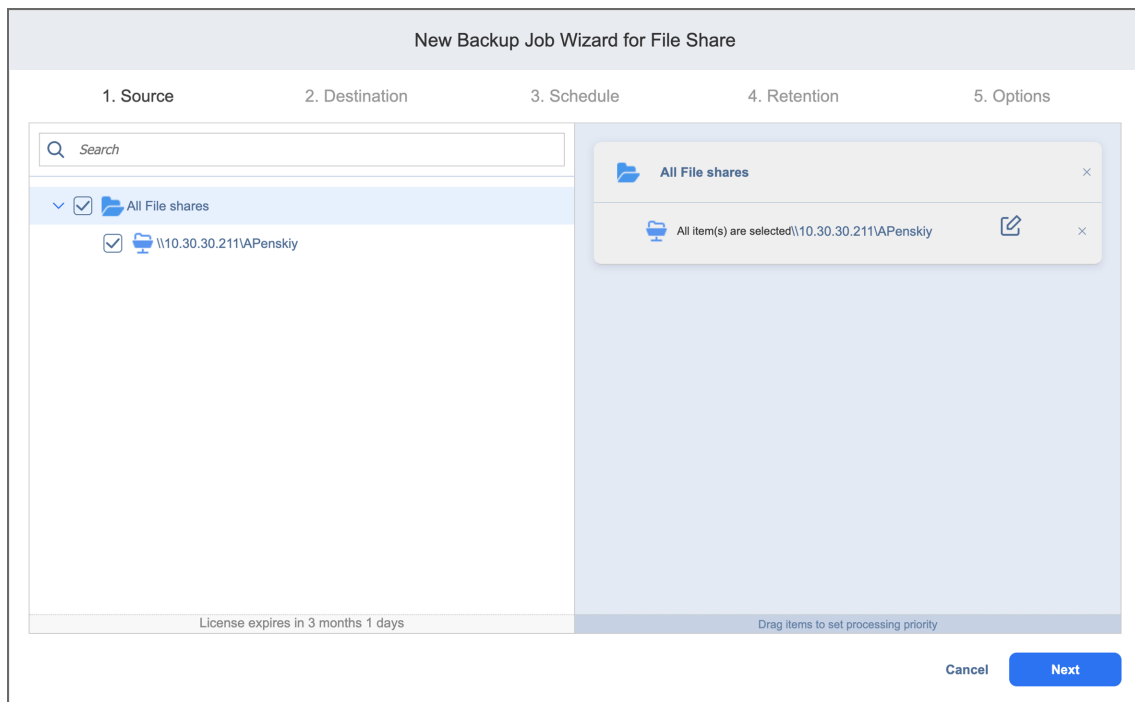
On the Source page of the wizard, add a file share to your backup job. To do this, proceed as follows:

1. In the left pane of the page, select the file shares you want to back up. Use the Search functionality to find shares by name.

### Notes

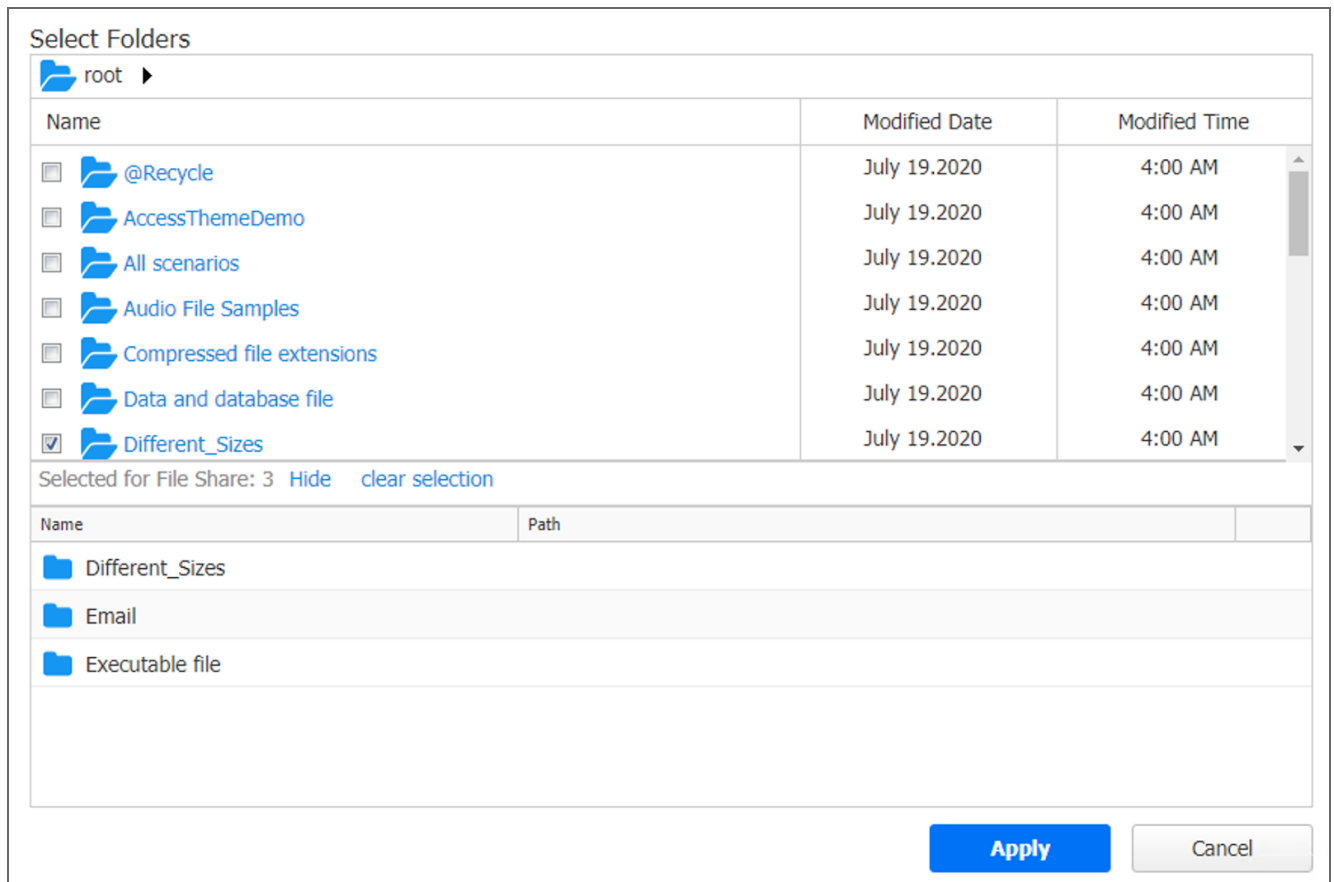
- When **All File Shares** is selected, all the file shares that were added to Inventory are added to the job. New items created in or moved to the source group are automatically added to this job.

- Due to Microsoft limitations, Windows files and folders whose names end in a space or period cannot be opened or created. If a Windows-based transporter is used to execute a file share backup job that includes file shares containing such files or folders, the job will fail.



2. Optionally, you can view and select separate folders of a share to back up. To do this, proceed as follows:
  - a. Hover over the selected share and click the edit icon. The **Select Folders** dialog box opens.
  - b. Select the folders you want to back up by selecting the checkbox next to them. To see the list of the selected folders, click **Show**. To cancel the selection, click **Clear selection**.

c. Click **Apply**.



3. Click **Next** to move to the next page of the wizard.

## Backup Job Wizard for File Share: Destination

On the Destination page, select a backup repository for storing the backups. You can back up all shares to one repository or select a different repository for each share.

### Note

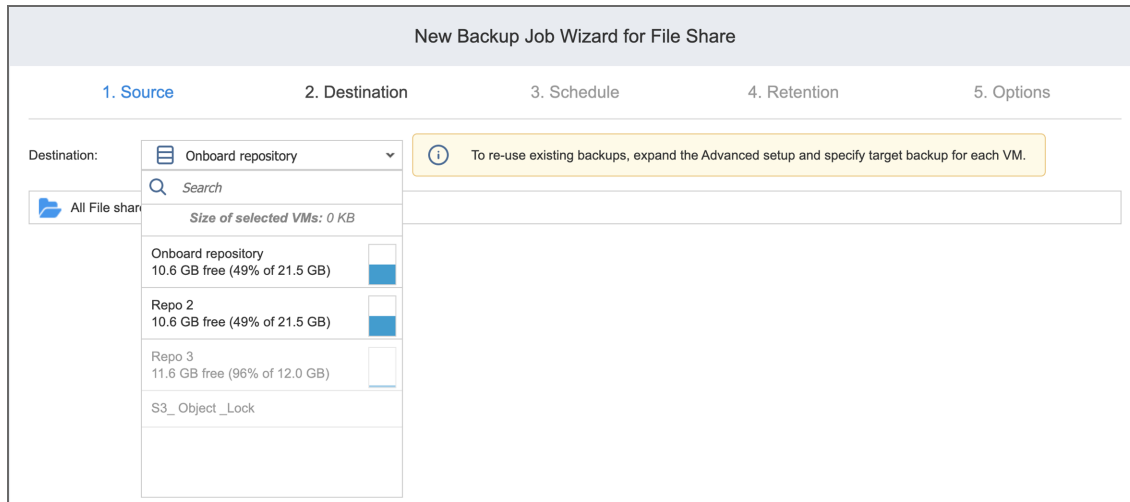
Backup of the file share is not possible to the following destinations:

- Backup repositories in Amazon S3
- **Local type** backup repositories that have backups with [immutability](#) enabled
- SaaS backup repositories
- Backup repositories in Wasabi Hot Cloud Storage
- Deduplication Appliance (HPE Catalyst, Dell EMC, NEC HYDRASTOR)



## Setting a Single Backup Repository for All Shares

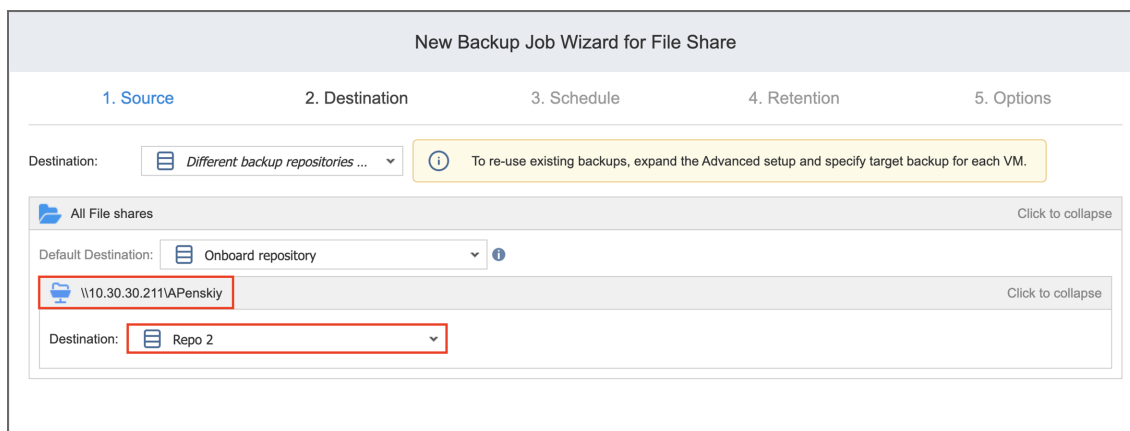
To back up selected shares to a single repository, choose a backup repository from the **Destination** drop-down list.



## Setting Different Backup Repositories for Shares

To back up the selected shares to different backup repositories, follow the steps below:

1. Click **Advanced setup**.
2. In the share boxes, choose a different backup repository for each share individually from the **Destination** drop-down list.



## Backup Job Wizard for File Share: 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.

- [Disabling Scheduled Job Execution](#)
- [Daily or Weekly Backup](#)
- [Monthly or Yearly Backup](#)
- [Periodic Backup](#)

- [Chained Backup](#)
- [Additional Schedule](#)

## Disabling Scheduled Job Execution

If you wish to start the job manually (without any schedule), select the **Do not schedule, run on demand** checkbox.

New Backup Job Wizard for File Share

1. Source
2. Destination
3. Schedule
4. Retention

Do not schedule, run on demand

(UTC+02:00, EET) Eastern European Time ▼

Schedule #1

Run daily/weekly ▼

Starting at:   Ending:

MO
  TU
  WE
  TH
  FR
  SA
  SU

[All days](#)
[Work days](#)
[Weekends](#)

every  weeks

Effective from

[Add another schedule](#)

[Show calendar](#)

## 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:

- 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.

New Backup Job Wizard for File Share

1. Source      2. Destination      3. Schedule      4. Retention

Do not schedule, run on demand

(UTC+02:00, EET) Eastern European Time

Schedule #1

**Run daily/weekly**

Starting at: 0:00    Ending: 6:00

MO  TU  WE  TH  FR     SA     SU

[All days](#)   [Work days](#)   [Weekends](#)

every 1 weeks

Effective from

[Add another schedule](#)

[Show calendar](#)

## Monthly or Yearly Backup

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.

New Backup Job Wizard for File Share

1. Source      2. Destination      3. Schedule      4. Retention

Do not schedule, run on demand

(UTC+02:00, EET) Eastern European Time

Schedule #1

**Run monthly/yearly**

Run every last Friday of every month

Starting at: 0:00    Ending: 6:00

Effective from

[Add another schedule](#)

[Show calendar](#)

## 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 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.

New Backup Job Wizard for File Share

1. Source      2. Destination      3. Schedule      4. Retention

Do not schedule, run on demand

(UTC+02:00, EET) Eastern European Time

Schedule #1

Run periodically every 30 minutes

Starting at: 0:00     Ending: 6:00

MO  TU  WE  TH  FR  SA  SU

[All days](#) [Work days](#) [Weekends](#)

Effective from

[Add another schedule](#)

[Show calendar](#)

## 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.

New Backup Job Wizard for File Share


1. Source      2. Destination      3. Schedule      4. Retention

---

Do not schedule, run on demand  
(UTC+02:00, EET) Eastern European Time

Schedule #1

Run after another job

After the job:  Microsoft 365 backup job

Run this job:

After successful runs    After failed runs    After stopped runs

Effective from

[Add another schedule](#)

[Show calendar](#)

## Additional Schedule

If you want to have more than one schedule for your job, click **Add another schedule** and set it up as it has been described above.

New Backup Job Wizard for File Share

1. Source      2. Destination      3. Schedule      4. Retention

---

Do not schedule, run on demand  
(UTC+02:00, EET) Eastern European Time

Schedule #1

Run daily/weekly

Starting at:   Ending:

MO  TU  WE  TH  FR  SA  SU

[All days](#)   [Work days](#)   [Weekends](#)

every  weeks

Effective from

[Add another schedule](#)

[Show calendar](#)

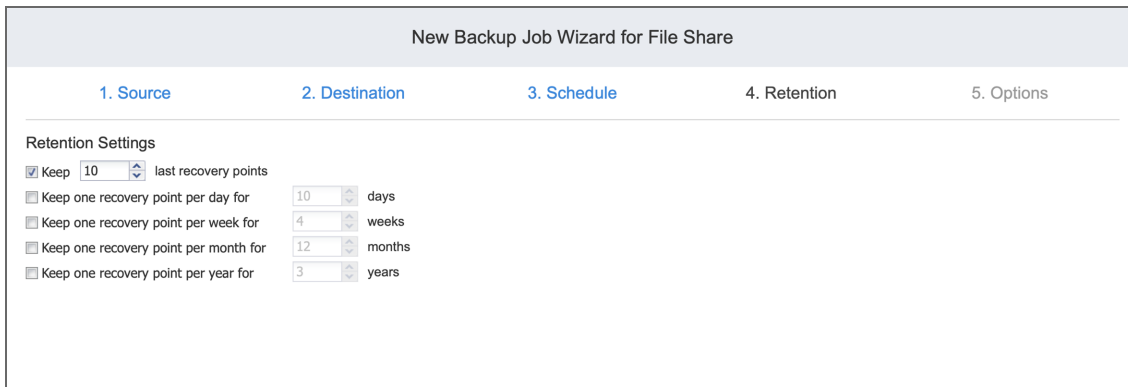
# Backup Job Wizard for File Share: Retention

After each job run, NAKIVO Backup & Replication creates a recovery point in the backup repository for each share. A recovery point represents the backed-up share at a particular moment in time and allows you to recover individual folders or an entire share from the backup repository. You can specify the number of recovery points to retain in the backup repository using the grandfather-father-son (GFS) backup rotation scheme.

## 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.



New Backup Job Wizard for File Share

1. Source      2. Destination      3. Schedule      4. Retention      5. Options

Retention Settings

Keep 10 last recovery points

Keep one recovery point per day for 10 days

Keep one recovery point per week for 4 weeks

Keep one recovery point per month for 12 months

Keep one recovery point per year for 3 years

## Backup Job Wizard for File Share: Options

On the **Options** page, set the job options and configure transporter load for the backup job.

### Job Options

In the **Options** section, do the following:

- **Job name:** Enter a name for the backup job.
- **Change tracking:** Choose one of the following options:

- **Use proprietary method:** When this option is selected, NAKIVO Backup & Replication performs incremental backups using proprietary change tracking technology. This feature requires the reading of contents of all file shares to determine which data blocks have changed since the last job run.
- **No change tracking (always full):** When this option is selected, the full data set is transferred on every job run.
- **Network acceleration:** When enabled, NAKIVO Backup & Replication uses compression and traffic reduction techniques to speed up data transfer. Select this option to back up over WAN or slow LAN links. For more information, refer to [“Network Acceleration” on page 63](#).
- **Network encryption:** When enabled, file share data is protected with AES 256 encryption while traveling over the network. For more information, refer to [“Encryption in Flight and at Rest” on page 42](#).

New Backup Job Wizard for File Share

1. Source      2. Destination      3. Schedule      4. Retention      5. Options

**Job Options**

Job name:

Change tracking:  ⓘ

Network acceleration:  ⓘ

Network encryption:  ⓘ

**Data Transfer**

Limit transporter load to  concurrent tasks ⓘ

## Note

You need at least one transporter at source and target sites to enable encryption.

## Data Transfer

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.

New Backup Job Wizard for File Share

1. Source      2. Destination      3. Schedule      4. Retention      5. Options

Job Options

Job name: File Share backup job

Change tracking: Use proprietary method ⓘ

Network acceleration: Disabled ⓘ

Network encryption: Disabled ⓘ

Data Transfer

Limit transporter load to 3 concurrent tasks ⓘ

## Creating Oracle Database Backup Jobs


With NAKIVO Backup & Replication, you can [back up Oracle Databases](#) by creating a backup job that specifies which databases should be backed up, how often the backup should be run, and what backup options should be used.

### Important

- To start the backup process, make sure the product is synced with the Oracle database and the following options are pre-configured in RMAN:
  - [Destination](#)
  - [Change-block tracking](#)
  - [Retention](#)
- Set an Oracle database to [ARCHIVELOG mode](#) to create backups.

To create a backup job, click **Create** on the **Dashboard** and then click **Oracle database backup job**.





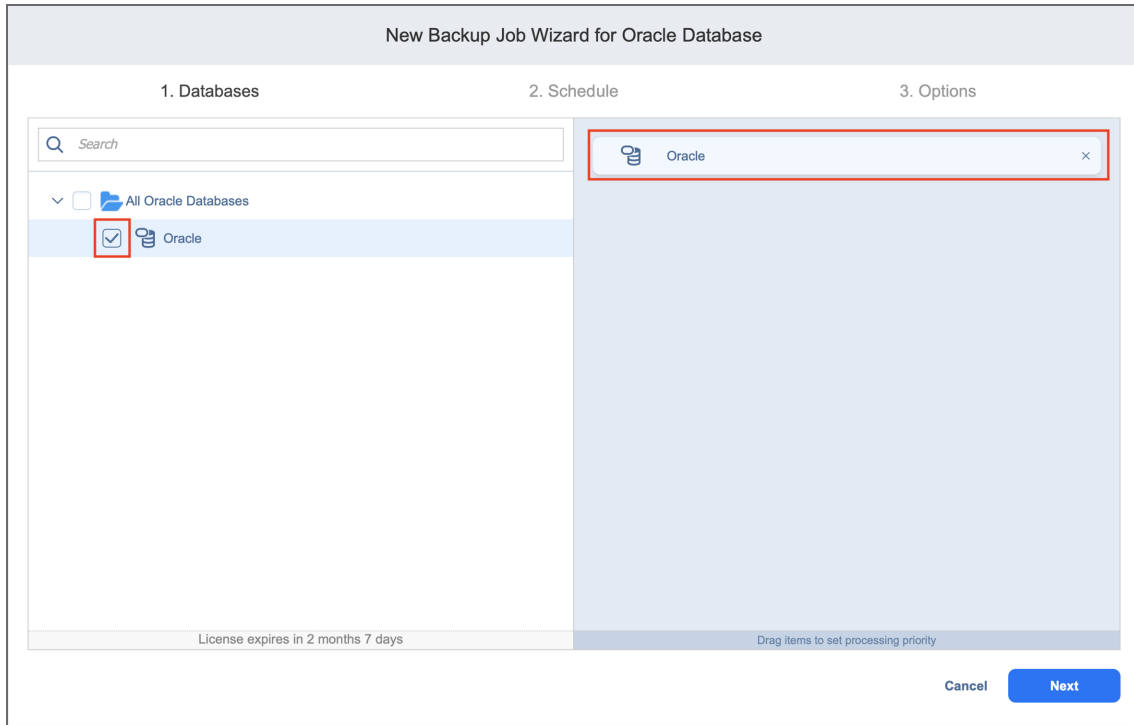
The **New Backup Job Wizard for Oracle Database** opens. Complete the wizard as described in the sections below:

- [“Backup Job Wizard for Oracle Database: Databases”](#) below
- [“Backup Job Wizard for Oracle Database: Schedule”](#) on page 660
- [“Backup Job Wizard for Oracle Database: Options”](#) on page 664

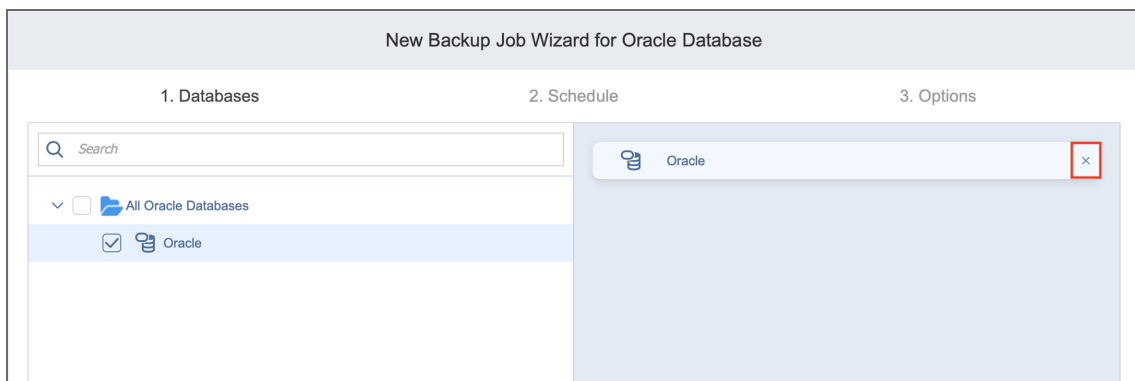
## Backup Job Wizard for Oracle Database: Databases

On the **Databases** page of the wizard, you can add Oracle databases to your backup job. Proceed as follows:

1. In the left pane of the page, you can select the checkbox of the items you want to back up. The selected items appear in the right pane of the page.



2. If necessary, reorder the selected items by dragging them to a new position. By doing so, you can specify which databases should be backed up first. If needed, remove the selected database from the backup job in either of the following ways:
  - Cancel the selection of the database in the left pane. This will remove it from the right pane.
  - In the right pane, hover the pointer over the database you wish to remove and click the remove button (x). This will remove it from the right pane.



3. Optionally, filter the inventory tree by entering a string in the Search box. You can enter a part or the entire item name.
4. Click **Next** to confirm adding the selected databases to the backup job.

# Backup Job Wizard for Oracle Database: 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.

- [“Disabling Scheduled Job Execution” below](#)
- [“Daily or Weekly Backup” below](#)
- [“Monthly or Yearly Backup” on the next page](#)
- [“Periodic Backup” on page 662](#)
- [“Chained Job” on page 663](#)
- [“Additional Schedule” on page 663](#)

## Disabling Scheduled Job Execution

If you wish to start the job manually (without any schedule), select the **Do not schedule, run on demand** checkbox.

The screenshot shows the 'New Backup Job Wizard for Oracle Database' interface. At the top, it says 'New Backup Job Wizard for Oracle Database'. Below that, there are three steps: '1. Databases', '2. Schedule', and '3. Options'. The '2. Schedule' step is currently active. In the '2. Schedule' section, there is a checkbox labeled 'Do not schedule, run on demand' which is checked and highlighted with a red border. At the bottom right of the wizard, there are two buttons: 'Next' (a blue button) and 'Cancel' (a white button with a grey border).

## 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:

- 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.

The screenshot shows the 'New Backup Job Wizard for Oracle Database' window, specifically the '2. Schedule' step. The interface includes three tabs: '1. Databases', '2. Schedule', and '3. Options'. Under the '2. Schedule' tab, there are several configuration options:

- A checkbox for 'Do not schedule, run on demand' is checked.
- A time zone dropdown menu is set to '(UTC+02:00, EET) Eastern European Time'.
- The 'Schedule #1' dropdown menu is highlighted with a red box and set to 'Run daily/weekly'.
- 'Starting at:' is set to '0:00' and 'Ending:' is set to '6:00'.
- Day selection checkboxes are checked for Mon, Tue, Wed, Thu, and Fri, and unchecked for Sat and Sun.
- Below the day selection, there are links for 'All days', 'Work days', and 'Weekends'.
- 'every' is set to '1' weeks.
- The 'Effective from' checkbox is unchecked.
- There are links for 'Add another schedule' and 'Show calendar'.
- At the bottom right, there are 'Next' and 'Cancel' buttons.

## Monthly or Yearly Backup

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.

New Backup Job Wizard for Oracle Database

1. Databases      2. Schedule      3. Options

Do not schedule, run on demand  
 (UTC+02:00, EET) Eastern European Time

Schedule #1  
Run monthly/yearly

Run every last Friday of every month

Starting at: 0:00    Ending: 6:00

Effective from

[Add another schedule](#)  
[Show calendar](#)

**Next**    Cancel

## 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 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.

New Backup Job Wizard for Oracle Database

1. Databases      2. Schedule      3. Options

Do not schedule, run on demand  
 (UTC+02:00, EET) Eastern European Time

Schedule #1  
Run periodically every 30 minutes

Starting at: 0:00    Ending: 6:00

Mon  Tue  Wed  Thu  Fri  Sat  Sun  
 All days Work days Weekends

Effective from

[Add another schedule](#)  
[Show calendar](#)

**Next**    Cancel

## 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.

The screenshot shows the 'New Backup Job Wizard for Oracle Database' interface, specifically the '2. Schedule' step. The wizard is divided into three sections: '1. Databases', '2. Schedule', and '3. Options'. In the '2. Schedule' section, the following options are visible:

- Do not schedule, run on demand
- (UTC+02:00, EET) Eastern European Time
- Schedule #1: Run after another job (highlighted with a red box)
- After the job: VMware backup job-test
- Run this job: Immediately
- After successful runs  After failed runs  After stopped runs
- Effective from

At the bottom of the wizard, there are two buttons: 'Next' (highlighted in blue) and 'Cancel'.

## Additional 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.

## Backup Job Wizard for Oracle Database: Options

On the **Options** page of the wizard, you can name the job, choose to back up and truncate Oracle database archived redo logs, set full backups as well as pre and post job actions.

- [“Job Options” below](#)
- [“Full Backup Settings” on the next page](#)
- [“Pre and Post Actions” on page 666](#)
  - [“Email Notifications” on page 667](#)
  - [“Pre Job Script” on page 667](#)
  - [“Post Job Script” on page 668](#)

### Job Options

In this section, you can set the following options:

- **Job name:** Specify a name for the backup job.
- **Back up Oracle database archived redo logs:** If selected, the Oracle database archived redo logs will be backed up on the specified databases.

- **Truncate Oracle database archived redo logs:** If selected, the Oracle database archived redo logs will be truncated on the specified databases immediately after backup is complete.

## Full Backup Settings

In this section, you can set time and frequency for full backups.

Select one of the following options:

- **Always:** If selected, the product will always make full backups
- **Every** Monday/Tuesday/Wednesday/Thursday/Friday/Saturday/Sunday.
- **Every 2nd** Monday/Tuesday/Wednesday/Thursday/Friday/Saturday/Sunday.
- **First** Monday/Tuesday/Wednesday/Thursday/Friday/Saturday/Sunday of the month.
- **Second** Monday/Tuesday/Wednesday/Thursday/Friday/Saturday/Sunday of the month.
- **Third** Monday/Tuesday/Wednesday/Thursday/Friday/Saturday/Sunday of the month.
- **Fourth** Monday/Tuesday/Wednesday/Thursday/Friday/Saturday/Sunday of the month.
- **Last** Monday/Tuesday/Wednesday/Thursday/Friday/Saturday/Sunday of the month.
- **Day#:** Specify the day when you would like the product to make a full backup.
- **Job run#:** Specify the number of the job run when you would like the product to make 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.



New Backup Job Wizard for Oracle Database

1. Databases                      2. Schedule                      3. Options

---

Job Options

Job name:

Back up Oracle database archived redo logs ⓘ

Truncate Oracle database archived redo logs ⓘ

**Full Backup Settings**

Create full backup:

If a full backup fails, create a full backup on the next job run ⓘ

Pre and Post Actions

Send job run reports to  ⓘ

Run local pre job script ⓘ

Run local post job script ⓘ

ⓘ All other settings for selected items should be pre-configured in the Oracle database RMAN configuration. [Learn more](#)

## Note

If your current job or the immediate parent job does not run daily/weekly at least once a week, only the **Always** and **Job run#** options will be available.

## Pre and Post Actions

NAKIVO Backup & Replication provides you with the ability to enable certain actions before a backup job begins and after it has been completed. You can choose to send job run reports, and run local [pre and post job scripts](#).

New Backup Job Wizard for Oracle Database

1. Databases                      2. Schedule                      3. Options

---

Job Options  
 Job name:

Back up Oracle database archived redo logs ⓘ  
 Truncate Oracle database archived redo logs ⓘ

Full Backup Settings  
 Create full backup:    
 If a full backup fails, create a full backup on the next job run ⓘ

**Pre and Post Actions**  
 Send job run reports to  ⓘ  
 Run local pre job script ⓘ  
 Run local post job script ⓘ

ⓘ All other settings for selected items should be pre-configured in the Oracle database RMAN configuration. [Learn more](#)

Cancel   **Finish**   Finish & Run

## Email Notifications

NAKIVO Backup & Replication can send email notifications on job completion status to specified recipients. This feature complements global notification and provides you with the ability to configure notifications on a per-job level.

### Note

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. The semi-colon character (;) should be used to separate multiple email addresses.

## Pre Job Script

To run a script before the product begins backing up databases, 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. Script interpreter should be specified.
  - **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 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 backing up databases 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 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 backup will not be performed.

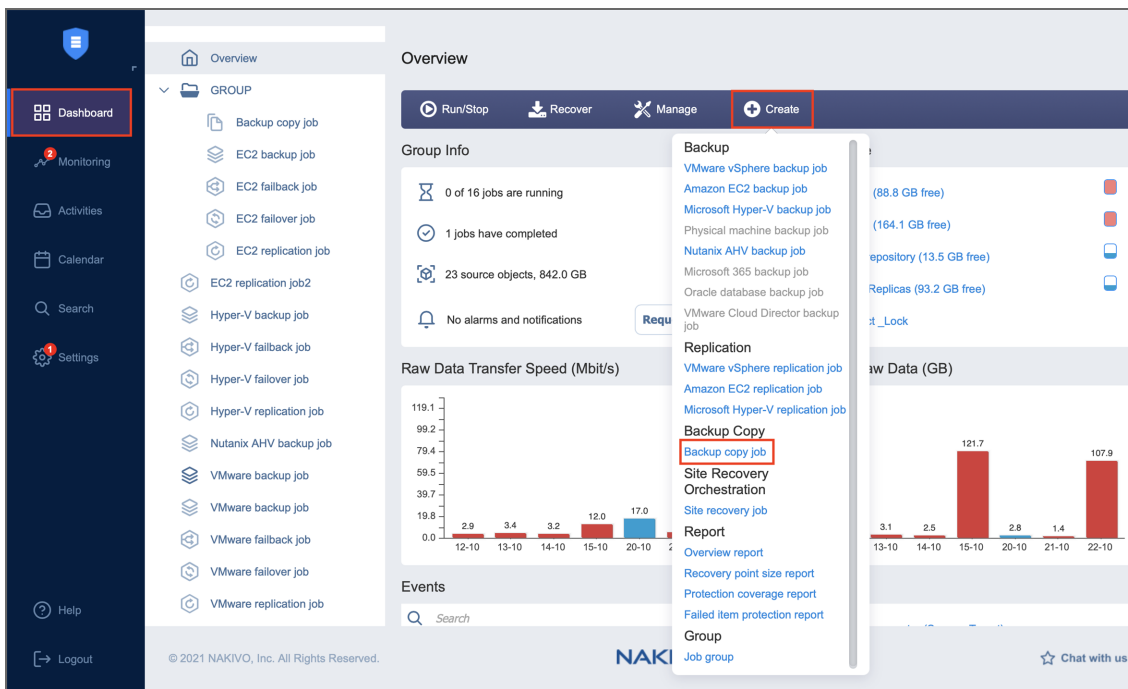
## Post Job Script

To run a script after the product has finished backing up all databases, 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.
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. Script interpreter should be specified.
  - **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 the backup has been successful.

# Creating Backup Copy Jobs

To create a [backup copy](#) job, click **Create** on the **Dashboard**, and then click **Backup copy job**.



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 672
- [“Backup Copy Job Wizard: Schedule”](#) on page 675
- [“Backup Copy Job Wizard: Retention”](#) on page 679
- [“Backup Copy Job Wizard: Options”](#) on page 680

## 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:

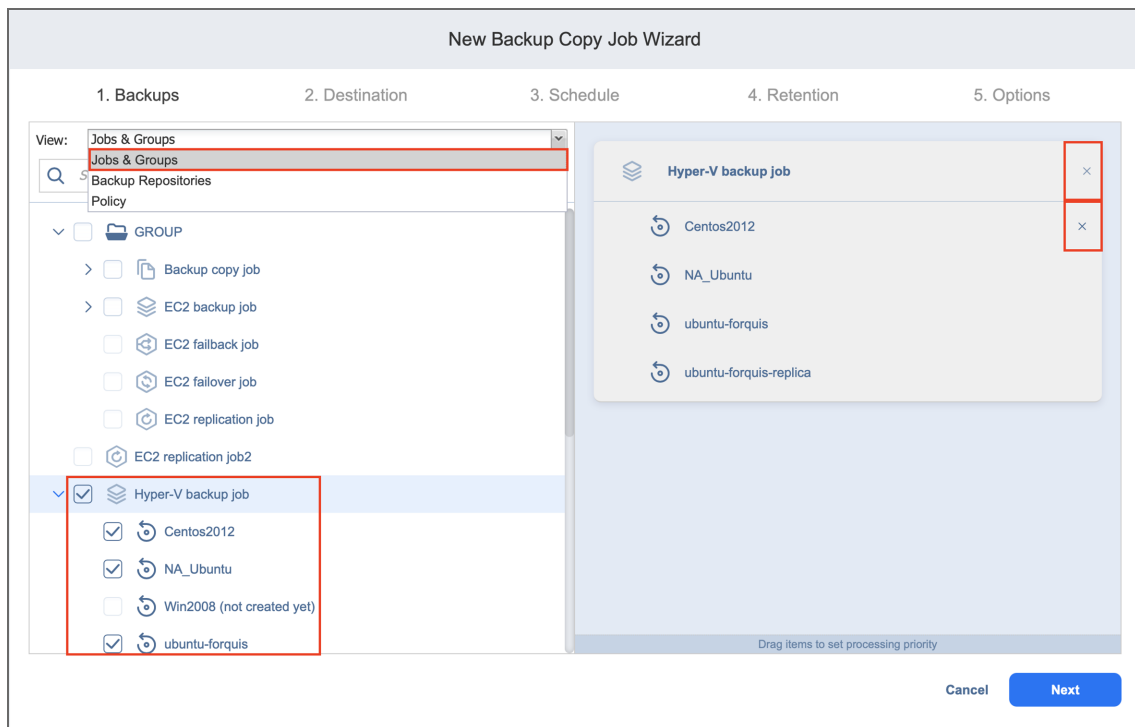
- [Creating Backup Copies Using Jobs and Groups](#)
- [Creating Backup Copies Using Backup Repositories](#)
- [Creating Backup Copies Using Policies](#)

### 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. If you select VMware Cloud Director objects, it will not be possible to select VMs separately if they are part of the vApp. In this case, the entire vApp needs to be selected to proceed.

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. The selected vApps containing VMs have a number of underlying VMs displayed in parenthesis next to its name. 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 will remove 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 will cancel 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 266](#) 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 will reset 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 269](#) for details.



Click **Next** to confirm that you wish to add selected items to the backup copy job. The wizard will display 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.

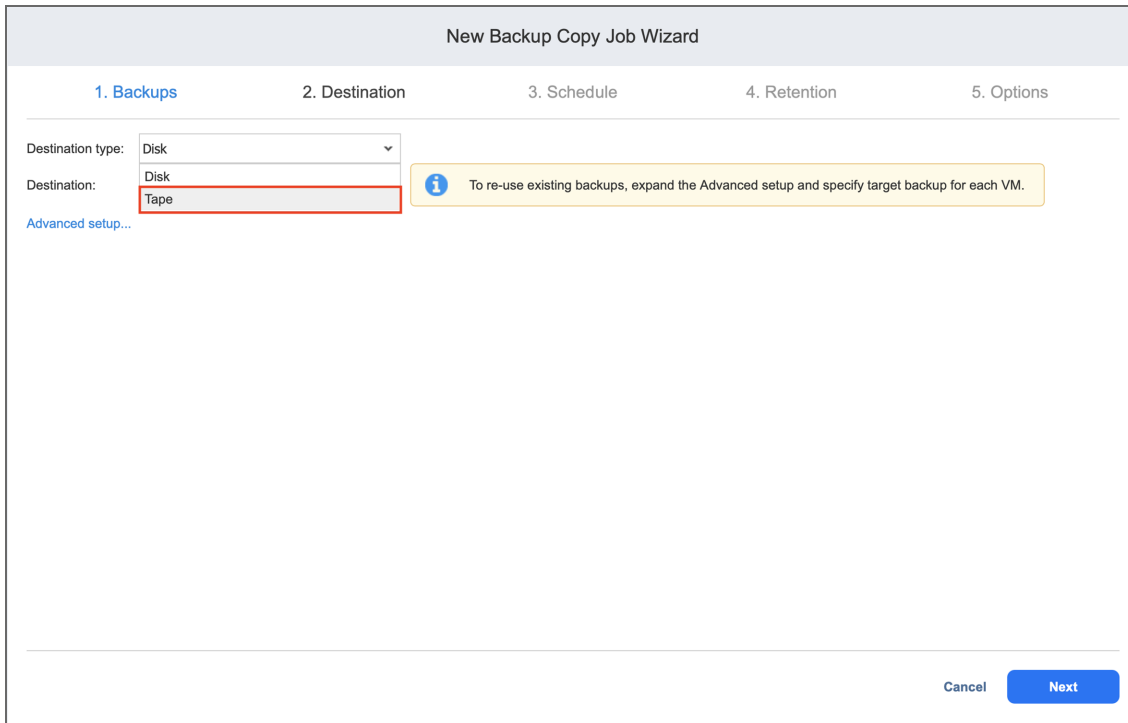
## Backup Copy Job Wizard: Destination

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](#)

## 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.



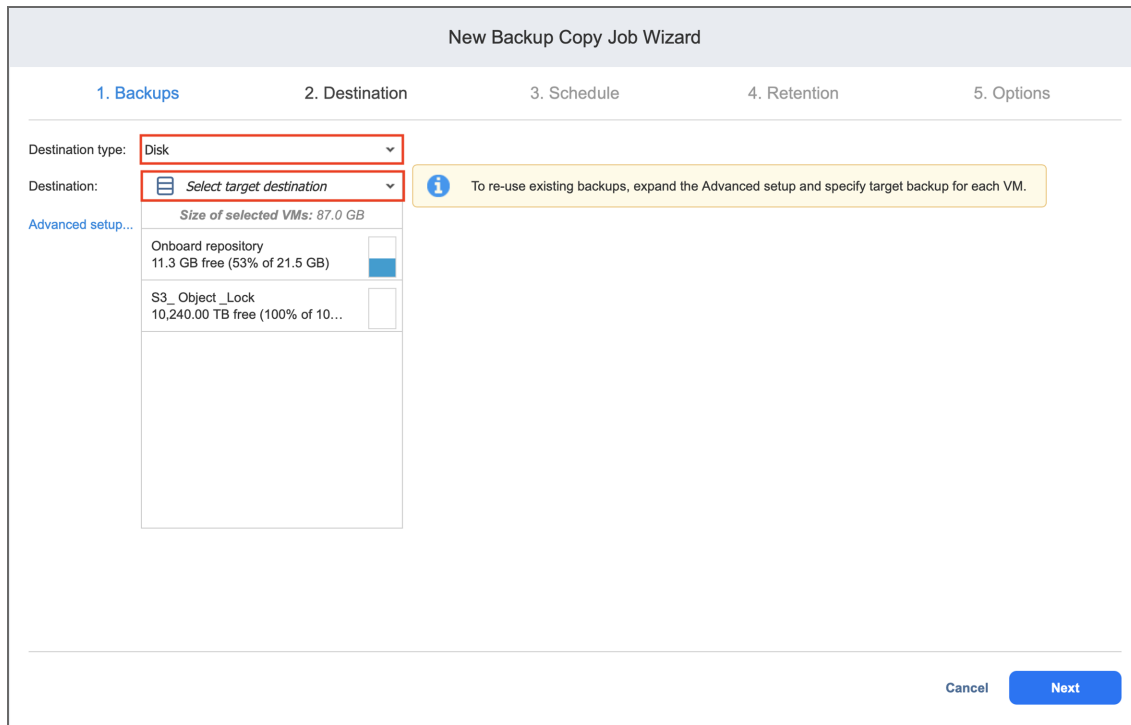
The screenshot shows the 'New Backup Copy Job Wizard' interface. The wizard has five steps: 1. Backups, 2. Destination, 3. Schedule, 4. Retention, and 5. Options. The 'Destination' step is currently active. In this step, there are two dropdown menus: 'Destination type' and 'Destination'. The 'Destination type' dropdown is set to 'Disk', and the 'Destination' dropdown is set to 'Tape'. A red box highlights the 'Tape' option in the 'Destination' dropdown. Below the dropdowns, there is a link for 'Advanced setup...'. To the right of the dropdowns, there is an information icon and a message: 'To re-use existing backups, expand the Advanced setup and specify target backup for each VM.' At the bottom right of the wizard, there are 'Cancel' and 'Next' buttons.

## Selecting a Target Backup Repository

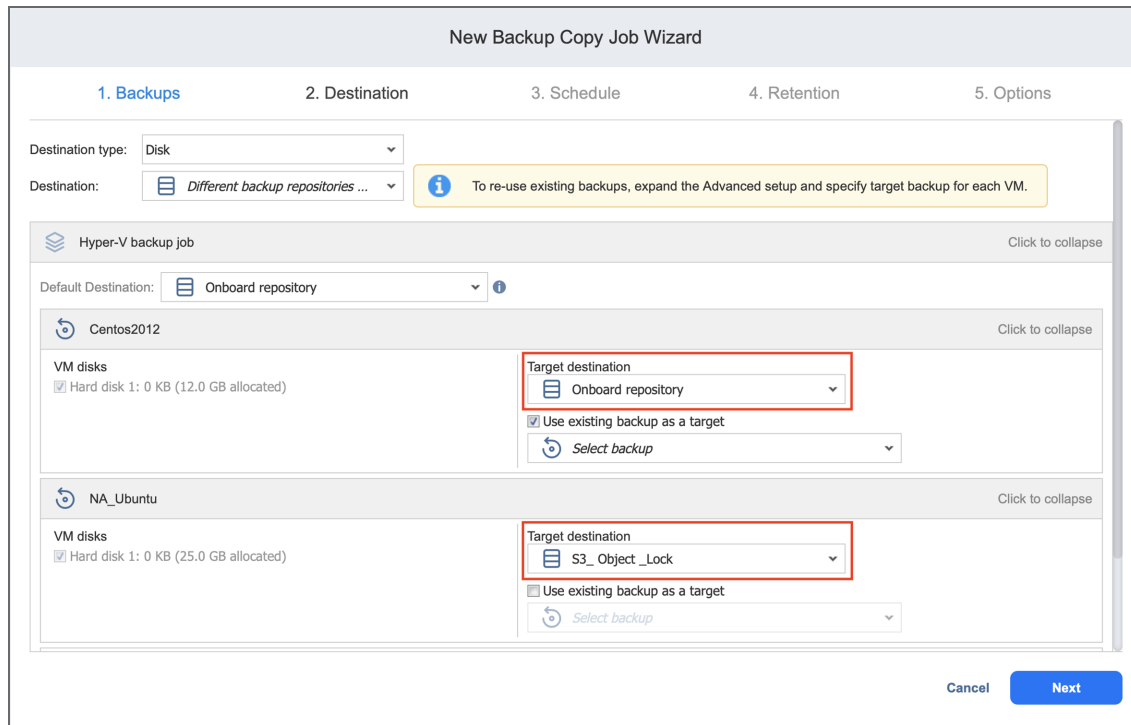
Backup Copy jobs can copy backups from one Backup Repository to another. Select a target Backup Repository as described below:

- 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.





- To copy backups to different Backup Repositories, follow the steps below:
  - a. Click **Advanced setup**.
  - b. For each backup, select a target Backup Repository.



## Mapping Source Backups to Existing Backups

If you have previously copied backups to a different Backup Repository and then lost the Backup Copy job (due to accidental job deletion or because you need to recreate jobs in a new copy 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.

The screenshot shows the 'New Backup Copy Job Wizard' interface. The wizard is divided into five steps: 1. Backups, 2. Destination, 3. Schedule, 4. Retention, and 5. Options. The current step is '1. Backups'. The 'Destination type' is set to 'Disk' and the 'Destination' is 'Different backup repositories ...'. A yellow information box states: 'To re-use existing backups, expand the Advanced setup and specify target backup for each VM.' The 'Hyper-V backup job' section is expanded, showing the 'Default Destination' as 'Onboard repository'. Underneath, two VMs are listed: 'Centos2012' and 'NA\_Ubuntu'. For 'Centos2012', the 'VM disks' section shows 'Hard disk 1: 0 KB (12.0 GB allocated)'. The 'Target destination' is 'Onboard repository', and the checkbox 'Use existing backup as a target' is checked. A red box highlights the 'Select backup' dropdown menu. For 'NA\_Ubuntu', the 'VM disks' section shows 'Hard disk 1: 0 KB (25.0 GB allocated)'. The 'Target destination' is 'S3\_Object\_Lock', and the checkbox 'Use existing backup as a target' is unchecked. The 'Select backup' dropdown menu is also visible. At the bottom right, there are 'Cancel' and 'Next' buttons.

When running the job, the product analyzes the existing backup copy you have selected, determines how it is different from the source backup, and transfers only the differential data.

## 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.

- [Disabling Scheduled Job Execution](#)
- [Daily or Weekly Backup](#)
- [Monthly or Yearly Backup](#)
- [Periodic Backup](#)

- [Chained Job](#)
- [Adding Additional Schedule](#)

## Disabling Scheduled Job Execution

If you wish to start the job manually (without any schedule), select the **Do not schedule, run on demand** checkbox.

The screenshot shows the 'New Backup Copy Job Wizard' interface. At the top, the title is 'New Backup Copy Job Wizard'. Below the title are five tabs: '1. Backups', '2. Destination', '3. Schedule', '4. Retention', and '5. Options'. The '3. Schedule' tab is active. In this tab, there is a checkbox labeled 'Do not schedule, run on demand' which is checked. At the bottom right of the wizard, there are two buttons: 'Next' (highlighted in blue) and 'Cancel'.

## 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:

- 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.

The screenshot shows the 'New Backup Copy Job Wizard' interface. At the top, the title is 'New Backup Copy Job Wizard'. Below the title are five tabs: '1. Backups', '2. Destination', '3. Schedule', '4. Retention', and '5. Options'. The '3. Schedule' tab is active. In this tab, there is a checkbox labeled 'Do not schedule, run on demand' which is unchecked. Below this, there is a time zone dropdown menu showing '(UTC+02:00, EET) Eastern European Time'. Underneath, there is a section for 'Schedule #1' with a dropdown menu showing 'Run daily/weekly' (highlighted with a red box). Below the dropdown, there are two time input fields: 'Starting at: 0:00' and 'Ending: 6:00'. Below these, there are checkboxes for days of the week: Mon, Tue, Wed, Thu, Fri, Sat, Sun. The 'All days' option is selected. Below the day checkboxes, there are links for 'All days', 'Work days', and 'Weekends'. Below the day checkboxes, there is a section for 'every 1 weeks' with a dropdown menu showing '1'. Below this, there is a checkbox labeled 'Effective from' which is unchecked. At the bottom left, there are links for 'Add another schedule' and 'Show calendar'. At the bottom right of the wizard, there are two buttons: 'Next' (highlighted in blue) and 'Cancel'.

## Monthly or Yearly Backup

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.

The screenshot shows the 'New Backup Copy Job Wizard' window, specifically the '3. Schedule' step. The window has a title bar and five tabs: '1. Backups', '2. Destination', '3. Schedule', '4. Retention', and '5. Options'. The '3. Schedule' tab is active. Below the tabs, there is a checkbox for 'Do not schedule, run on demand'. Below that is a time zone dropdown menu showing '(UTC+02:00, EET) Eastern European Time'. The 'Schedule #1' section has a dropdown menu with 'Run monthly/yearly' selected and highlighted by a red box. Below this is the 'Run every' section with dropdowns for 'last', 'Friday', 'of every', and 'month'. The 'Starting at' field is set to '0:00' and the 'Ending' field is set to '6:00'. There is an 'Effective from' checkbox which is unchecked. At the bottom left, there are links for 'Add another schedule' and 'Show calendar'. At the bottom right, there are 'Next' and 'Cancel' buttons.

## 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 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.

New Backup Copy Job Wizard

1. Backups 2. Destination 3. Schedule 4. Retention 5. Options

Do not schedule, run on demand  
(UTC+02:00, EET) Eastern European Time

Schedule #1

**Run periodically** every 30 minutes

Starting at: 0:00 Ending: 6:00

Mon  Tue  Wed  Thu  Fri  Sat  Sun  
All days Work days Weekends

Effective from

[Add another schedule](#)  
[Show calendar](#)

**Next** Cancel

## 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.

New Backup Copy Job Wizard

1. Backups 2. Destination 3. Schedule 4. Retention 5. Options

Do not schedule, run on demand  
(UTC+02:00, EET) Eastern European Time

Schedule #1

**Run after another job**

After the job: Hyper-V backup job

Run this job: Immediately

After successful runs  After failed runs  After stopped runs

Effective from

[Add another schedule](#)  
[Show calendar](#)

**Next** Cancel

## Adding Additional 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.

New Backup Copy Job Wizard

1. Backups 2. Destination 3. Schedule 4. Retention 5. Options

Do not schedule, run on demand

(UTC+02:00, EET) Eastern European Time

Schedule #1

Run daily/weekly

Starting at: 0:00 Ending: 6:00

Mon  Tue  Wed  Thu  Fri  Sat  Sun

All days Work days Weekends

every 1 weeks

Effective from

**Add another schedule**

Show calendar

Next Cancel

## Backup Copy Job Wizard: Retention

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 S3** 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 S3** type 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:** All available recovery points are copied by the job. 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 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:

- Only **Amazon S3** or **Local Folder** must be selected for Backup Repository Type on the Destination page of the wizard.
- If **Amazon S3** is selected as the Backup Repository, Object Lock must be enabled for the Amazon S3 bucket 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](#)
  - [Network Acceleration](#)
  - [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](#)

## Job Options

In this section, you can give a name to the backup copy job and enable/disable [network acceleration](#), change tracking, [encryption](#), and [VM Verification](#). Proceed as described below.

The screenshot shows the 'New Backup Copy Job Wizard' at step 5, 'Options'. The 'Job Options' section is highlighted with a red box and contains the following settings:

- Job name: Backup copy job
- Network acceleration: Disabled
- Network encryption: Disabled
- VM verification: Disabled

Other settings visible in the wizard include:

- Full Backup Settings:** Create full backup: Every Friday; Full backup mode: Synthetic full;  If a full backup fails, create a full backup on the next job run.
- Pre and Post Actions:**  Send job run reports to;  Run local pre job script;  Run local post job script.
- Data Transfer:**  Limit transporter load to 3 concurrent tasks; Bandwidth throttling: Disabled;  Bottleneck detection.

At the bottom right, there are buttons for 'Cancel', 'Finish', and 'Finish & Run'.

### Job Name

Specify a name for the backup copy job in the **Job Name** box.

### 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.

### Encryption

If the **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 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 [“VM Verification” on page 48](#) .

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 VM 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 VM 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 VM 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.

New Backup Copy Job Wizard

1. Backups 2. Destination 3. Schedule 4. Retention 5. Options

Job Options

Job name: Backup copy job

Network acceleration: Disabled

Network encryption: Disabled

VM verification: Screenshot verification **settings**

Full Backup Settings

Create full backup: Every

Full backup mode: Synthetic

If a full backup fails, create a full backup on the next backup

Pre and Post Actions

Send job run reports to

Run local pre job script

Run local post job script

Data Transfer

Limit transporter load to

Bandwidth throttling: Disabled

Bottleneck detection

VM Boot Location

Target container: Choose target container

Target datastore: Choose target datastore

Proxy transporter: Do not use proxy transporter

Verification Options

Verify not more than: 2 VMs simultaneously

Recovery time objective: 5 minutes

Screenshot delay: 30 seconds

Cancel Finish Finish & Run

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.

The screenshot shows the 'New Backup Copy Job Wizard' interface, specifically the 'Options' step. The wizard is divided into five sections: 1. Backups, 2. Destination, 3. Schedule, 4. Retention, and 5. Options. In the 'Options' section, the 'VM verification' option is highlighted with a red box. Below it, the 'Full Backup Settings' section is visible, including 'Create full backup' (set to 'Every'), 'Full backup mode' (set to 'Synthetic'), and 'If a full backup fails, create a full backup on the next job run' (checked). A 'VM Boot Location' dialog box is open, showing options for 'Target container' (Choose target container), 'Target datastore' (Choose target datastore), and 'Proxy transporter' (Do not use proxy transporter). The 'Verification Options' section in the dialog shows 'Verify not more than' set to 2 VMs simultaneously and 'Recovery time objective' set to 5 minutes. At the bottom right, there are 'Cancel', 'Finish', and 'Finish & Run' buttons.

## 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 VM data from the source datastore 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.

New Backup Copy Job Wizard

1. Backups 2. Destination 3. Schedule 4. Retention 5. Options

Job Options

Job name: Backup copy job

Network acceleration: Disabled

Network encryption: Disabled

VM verification: Boot verification settings

**Full Backup Settings**

Create full backup: Every Friday

Full backup mode: Synthetic full

If a full backup fails, create a full backup on the next job run

Pre and Post Actions

Send job run reports to

Run local pre job script

Run local post job script

Data Transfer

Limit transporter load to 3 concurrent tasks

Bandwidth throttling: Disabled

Bottleneck detection

Cancel Finish Finish & Run

## 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](#).

New Backup Copy Job Wizard

1. Backups 2. Destination 3. Schedule 4. Retention 5. Options

Job Options

Job name: Backup copy job

Network acceleration: Disabled

Network encryption: Disabled

VM verification: Boot verification settings

Full Backup Settings

Create full backup: Every Friday

Full backup mode: Synthetic full

If a full backup fails, create a full backup on the next job run

**Pre and Post Actions**

Send job run reports to

Run local pre job script

Run local post job script

Data Transfer

Limit transporter load to 3 concurrent tasks

Bandwidth throttling: Disabled

Bottleneck detection

Cancel Finish Finish & Run

## 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](#).

The screenshot shows the 'New Backup Copy Job Wizard' with five steps: 1. Backups, 2. Destination, 3. Schedule, 4. Retention, and 5. Options. The 'Options' step is active. The 'Data Transfer' section is highlighted with a red box and contains the following settings:

- Limit transporter load to: 3 concurrent tasks
- Bandwidth throttling: Disabled
- Bottleneck detection

Other visible settings include:

- Job Options:** Job name: Backup copy job; Network acceleration: Disabled; Network encryption: Disabled; VM verification: Boot verification.
- Full Backup Settings:** Create full backup: Every Friday; Full backup mode: Synthetic full.
- Pre and Post Actions:** Send job run reports to, Run local pre job script, Run local post job script.

At the bottom right, there are buttons for 'Cancel', 'Finish', and 'Finish & Run'.

## 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 284](#) 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.

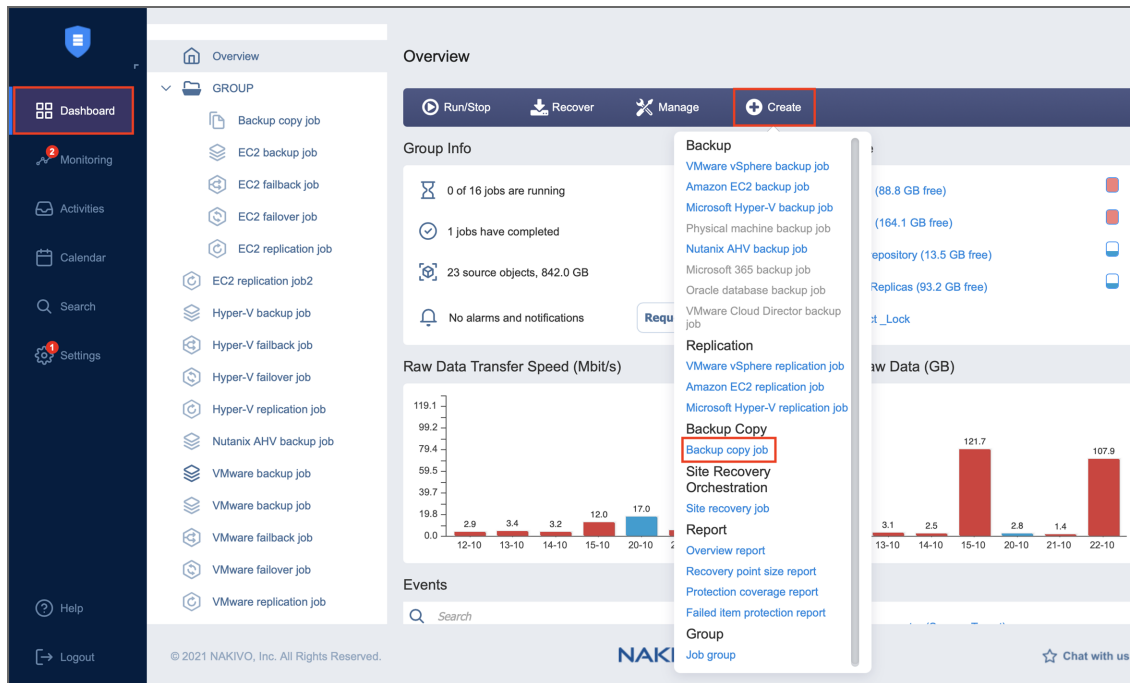
### Note

If you click **Finish & Run**, you will have to define the scope of your job. Please refer to [“Running Jobs on Demand” on page 256](#) 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** on the **Dashboard** and then click **Backup copy job**.



The **New Backup Copy Job Wizard** opens. Complete the wizard as described in the sections below:

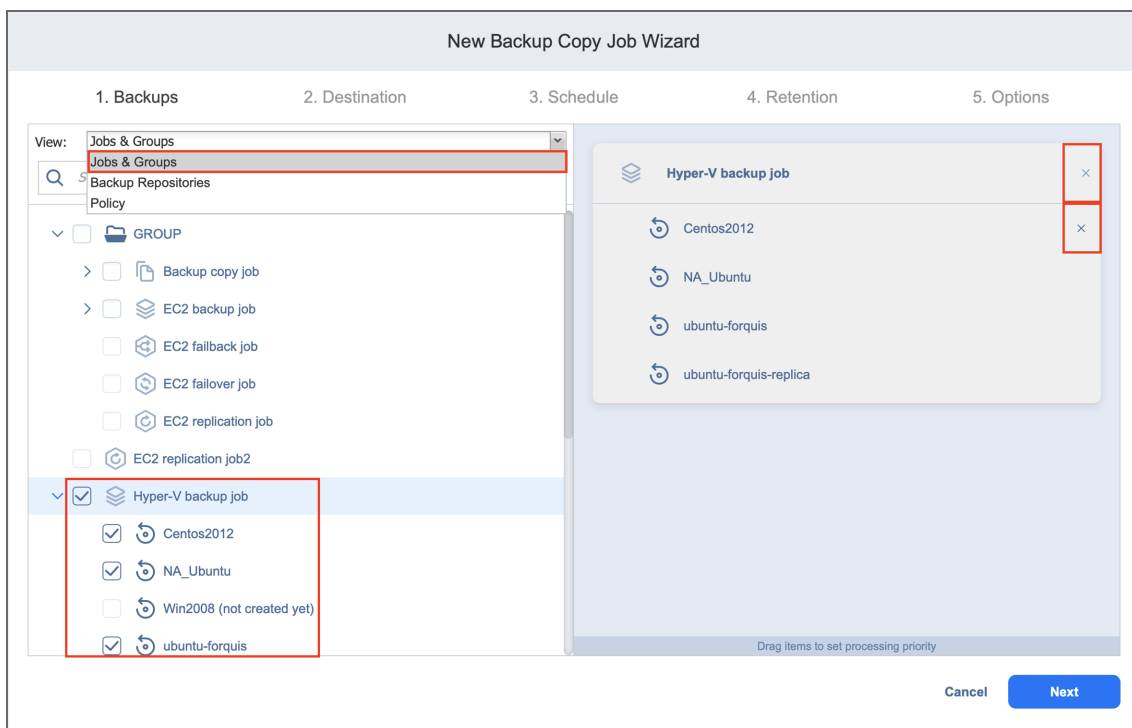
- [“Tape Backup Wizard: Backups” on the next page](#)
- [“Tape Backup Wizard: Destination” on page 692](#)
- [“Tape Backup Wizard: Schedule” on page 693](#)
- [“Tape Backup Wizard: Retention” on page 697](#)
- [“Tape Backup Wizard: Options” on page 698](#)




# 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 will remove 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.



- **Policy:** If selected, this allows you to use job policies; refer to [“Managing Job Policies” on page 266](#) 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 269](#) for details.



2. Click **Next** to confirm adding selected items to the backup copy job.

The wizard will display the next page.

## Notes

1. 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.
2. 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.

New Backup Copy Job Wizard

1. Backups      2. Destination      3. Schedule      4. Retention      5. Options

Destination type:

Destination:

**NT-dummy** Click to collapse

VM disks

- Hard disk 1 (0 KB (2.0 GB allocated) on S3\_ Object \_Lock)

**NA\_Ubuntu** Click to collapse

VM disks

- Hard disk 1: 0 KB (25.0 GB allocated)

**NBR10.4** Click to collapse

VM disks

- Hard disk 1 (0 KB (40.0 GB allocated) on S3\_ Object \_Lock)
- Hard disk 2 (0 KB (200.0 GB allocated) on S3\_ Object \_Lock)

**NT-SQL2-recovered** Click to collapse

VM disks

- Hard disk 1 (0 KB (5.0 GB allocated) on S3\_ Object \_Lock)

Cancel Next

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.

- [Disabling Scheduled Job Execution](#)
- [Daily or Weekly Backup](#)
- [Monthly or Yearly Backup](#)
- [Periodic Backup](#)
- [Chained Job](#)
- [Adding Another Schedule](#)

## Disabling Scheduled Job Execution

If you wish to start the job manually (without any schedule), select the **Do not schedule, run on demand** checkbox.

The screenshot shows the 'New Backup Copy Job Wizard' interface. At the top, it says 'New Backup Copy Job Wizard'. Below that, there are five tabs: '1. Backups', '2. Destination', '3. Schedule', '4. Retention', and '5. Options'. The '1. Backups' tab is selected. In this tab, there is a checkbox labeled 'Do not schedule, run on demand' which is checked. This checkbox is highlighted with a red rectangular box. At the bottom right of the wizard, there are two buttons: 'Next' (in blue) and 'Cancel' (in white with a grey border).

## 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.

The screenshot shows the 'New Backup Copy Job Wizard' interface, specifically the '3. Schedule' step. At the top, it says 'New Backup Copy Job Wizard'. Below that, there are five tabs: '1. Backups', '2. Destination', '3. Schedule', '4. Retention', and '5. Options'. The '3. Schedule' tab is selected. In this tab, there is a checkbox labeled 'Do not schedule, run on demand' which is unchecked. Below this, there is a time zone dropdown menu showing '(UTC+02:00, EET) Eastern European Time'. Underneath, there is a section for 'Schedule #1' with a dropdown menu showing 'Run daily/weekly', which is highlighted with a red rectangular box. Below the dropdown, there are two input fields: 'Starting at: 0:00' and 'Ending: 6:00'. There are also checkboxes for days of the week: Mon, Tue, Wed, Thu, Fri, Sat, and Sun. Below these are links for 'All days', 'Work days', and 'Weekends'. There is a field for 'every 1 weeks'. At the bottom, there is a checkbox for 'Effective from' and two links: 'Add another schedule' and 'Show calendar'. At the bottom right of the wizard, there are two buttons: 'Next' (in blue) and 'Cancel' (in white with a grey border).

## 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.

New Backup Copy Job Wizard

1. Backups      2. Destination      3. Schedule      4. Retention      5. Options

Do not schedule, run on demand  
 (UTC+02:00, EET) Eastern European Time

Schedule #1  
Run monthly/yearly

Run every  of every  of every

Starting at:  Ending:

Effective from

[Add another schedule](#)  
[Show calendar](#)

## 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.

New Backup Copy Job Wizard

1. Backups      2. Destination      3. Schedule      4. Retention      5. Options

Do not schedule, run on demand

(UTC+02:00, EET) Eastern European Time

Schedule #1

**Run periodically** every 30 minutes

Starting at: 0:00    Ending: 6:00

Mon    Tue    Wed    Thu    Fri    Sat    Sun

All days   Work days   Weekends

Effective from

[Add another schedule](#)

[Show calendar](#)

**Next**    Cancel

## 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.

New Backup Copy Job Wizard

1. Backups      2. Destination      3. Schedule      4. Retention      5. Options

Do not schedule, run on demand

(UTC+02:00, EET) Eastern European Time

Schedule #1

**Run after another job**

After the job: **Tape recovery job**

Run this job: Immediately

After successful runs    After failed runs    After stopped runs

Effective from

[Add another schedule](#)

[Show calendar](#)

**Next**    Cancel

## Adding 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.

The screenshot shows the 'New Backup Copy Job Wizard' window, specifically the '3. Schedule' step. The window has a light blue header and a white body. At the top, there are five tabs: '1. Backups', '2. Destination', '3. Schedule', '4. Retention', and '5. Options'. The '3. Schedule' tab is active. Below the tabs, there is a checkbox for 'Do not schedule, run on demand'. Below that is a dropdown menu for time zone, currently set to '(UTC+02:00, EET) Eastern European Time'. Underneath is a section for 'Schedule #1' with a dropdown menu set to 'Run daily/weekly'. Below this are two input fields for 'Starting at: 0:00' and 'Ending: 6:00'. There are checkboxes for days of the week: Mon, Tue, Wed, Thu, Fri, Sat, and Sun. Below these are three links: 'All days', 'Work days', and 'Weekends'. There is an input field for 'every 1 weeks'. Below that is a checkbox for 'Effective from'. At the bottom left, there is a link 'Add another schedule' which is highlighted with a red rectangular box. Below it is a link 'Show calendar'. At the bottom right, there are two buttons: 'Next' (blue) and 'Cancel' (white).

## Tape Backup Wizard: Retention

Specify how many recovery points you need to be copied by the job. Use the following options:

- **Keep all recovery points forever:** All available recovery points will be copied by the job. 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](#)
  - [Network Acceleration](#)
  - [Data Encryption](#)
- [Full Backup Settings](#)
  - [Create Full Backup](#)
  - [Tape Appending](#)
- [Pre and Post Job Actions](#)
  - [Email Notifications](#)
  - [Pre Job Script](#)
  - [Post Job Script](#)
- [Data Transfer](#)
  - [Bandwidth Throttling](#)
- [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.

New Backup Copy Job Wizard

1. Backups 2. Destination 3. Schedule 4. Retention 5. Options

**Job Options**

Job name: Backup copy job

Network acceleration: Disabled ⓘ

Network encryption: Disabled ⓘ

Full Backup Settings

Create full backup: Job runs # 5

If a full backup fails, create a full backup on the next job run ⓘ

Tape appending: Append previous tape, if possible ⓘ

Pre and Post Actions

Send job run reports to ⓘ

Run local pre job script ⓘ

Run local post job script ⓘ

Data Transfer

Bandwidth throttling: Disabled ⓘ

Cancel Finish Finish & Run

### Job Name

Specify a name for the tape backup job in the **Job name** box.

### Network Acceleration

Once Network Acceleration is enabled, NAKIVO Backup & Replication will use 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](#).

### Data Encryption

If the **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 back up over WAN without a VPN connection. For more information, refer to [“Encryption in Flight and at Rest” on page 42](#).

#### Note

You need at least one Transporter at the source and target sites to enable encryption.

## Full Backup Settings

In this section, select when you want the system to create full backups and set up the rules for data appending.

The screenshot shows the 'New Backup Copy Job Wizard' interface. At the top, there are five tabs: '1. Backups', '2. Destination', '3. Schedule', '4. Retention', and '5. Options'. The '1. Backups' tab is active. Below the tabs, there are several sections of configuration options:

- Job Options:** Job name: 'Backup copy job'; Network acceleration: 'Disabled'; Network encryption: 'Disabled'.
- Full Backup Settings (highlighted with a red box):** Create full backup: 'Job runs #' dropdown and '5' spinner; checkbox 'If a full backup fails, create a full backup on the next job run'; Tape appending: 'Append previous tape, if possible' dropdown.
- Pre and Post Actions:** checkboxes for 'Send job run reports to', 'Run local pre job script', and 'Run local post job script'.
- Data Transfer:** Bandwidth throttling: 'Disabled'.

At the bottom right, there are three buttons: 'Cancel', 'Finish', and 'Finish & Run'.

## 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](#).

The screenshot shows the 'New Backup Copy Job Wizard' with five steps: 1. Backups, 2. Destination, 3. Schedule, 4. Retention, and 5. Options. The 'Options' step is active. Under 'Job Options', 'Job name' is 'Backup copy job', 'Network acceleration' and 'Network encryption' are both 'Disabled'. Under 'Full Backup Settings', 'Job runs #' is set to 5, and 'Tape appending' is 'Append previous tape, if possible'. The 'Pre and Post Actions' section is highlighted with a red box and contains three unchecked checkboxes: 'Send job run reports to', 'Run local pre job script', and 'Run local post job script'. Below these is a text input field. Under 'Data Transfer', 'Bandwidth throttling' is 'Disabled'. At the bottom right are 'Cancel', 'Finish', and 'Finish & Run' buttons.

## 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 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.

## 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](#).

The screenshot shows the 'New Backup Copy Job Wizard' interface. At the top, there are five tabs: '1. Backups', '2. Destination', '3. Schedule', '4. Retention', and '5. Options'. The 'Options' tab is selected. Below the tabs, there are several sections of settings:

- Job Options:** Job name: Backup copy job; Network acceleration: Disabled; Network encryption: Disabled.
- Full Backup Settings:** Create full backup: Job runs # 5; If a full backup fails, create a full backup on the next job run: ; Tape appending: Append previous tape, if possible.
- Pre and Post Actions:** Send job run reports to: ; Run local pre job script: ; Run local post job script: .
- Data Transfer:** Bandwidth throttling: Disabled (highlighted with a red box).

At the bottom right, there are three buttons: 'Cancel', 'Finish', and 'Finish & Run'.

## 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 284](#) 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 284](#) 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 will be disabled for all jobs.

- Remove a bandwidth rule. Click the **Remove** link and then click **Delete** to confirm your operation.

Job Bandwidth Rules				
<input type="text" value="Search"/>				
Rule Name	Limit	Schedule	Status	
<input checked="" type="checkbox"/>  BS19804 Rule	10 Mbit/s	None	Active	
<input checked="" type="checkbox"/>  BS19805 Rule	5 Mbit/s	2:00 on MON, TUE	Waiting on schedule	
<input checked="" type="checkbox"/>  BS19805 Rule2	7 Mbit/s	2:00 on WED, THU	Active	
<input checked="" type="checkbox"/>  BS19805 Rule3	9 Mbit/s	2:00 on FRI	Waiting on schedule	

[Create New Rule](#)

## Completing Tape Backup 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. Refer to [“Running Jobs on Demand” on page 256](#) 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

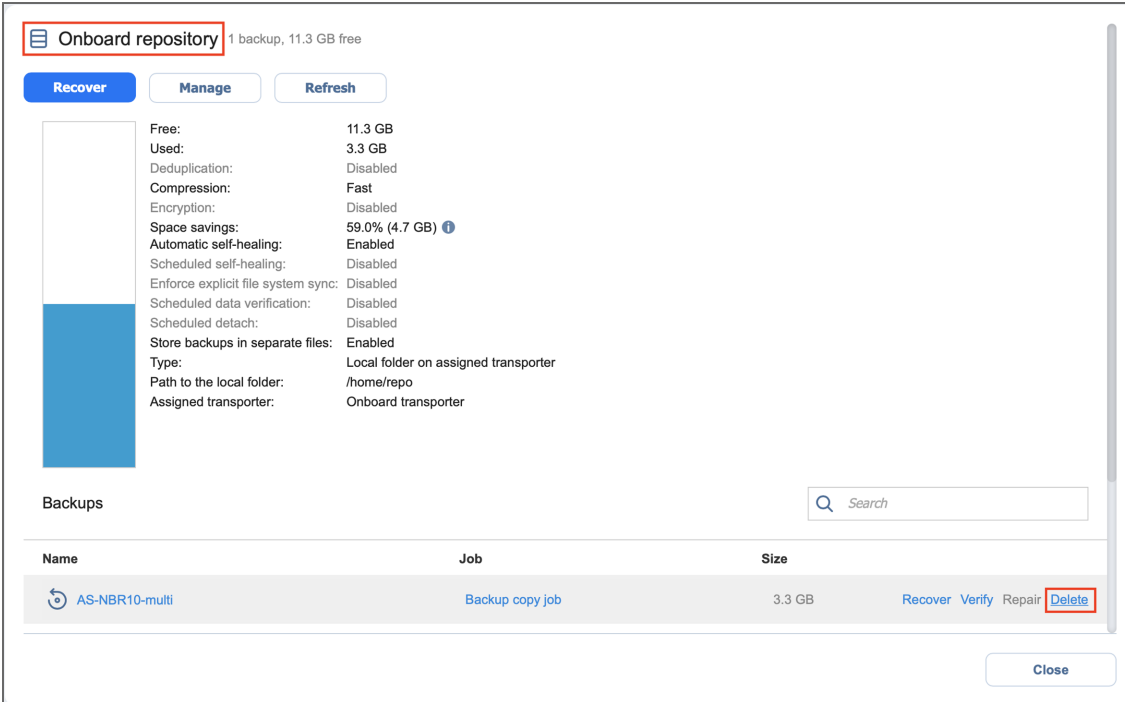
To delete a backup permanently, follow the steps below:

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.

## 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.



The screenshot shows the 'Onboard repository' management interface. At the top, there is a header 'Onboard repository' with a sub-header '1 backup, 11.3 GB free'. Below this are three buttons: 'Recover', 'Manage', and 'Refresh'. A large blue bar on the left indicates the repository's status. To the right, a list of settings is displayed:

- Free: 11.3 GB
- Used: 3.3 GB
- Deduplication: Disabled
- Compression: Fast
- Encryption: Disabled
- Space savings: 59.0% (4.7 GB) ⓘ
- Automatic self-healing: Enabled
- Scheduled self-healing: Disabled
- Enforce explicit file system sync: Disabled
- Scheduled data verification: Disabled
- Scheduled detach: Disabled
- Store backups in separate files: Enabled
- Type: Local folder on assigned transporter
- Path to the local folder: /home/repo
- Assigned transporter: Onboard transporter

Below the settings is a 'Backups' section with a search bar. A table lists the backups:

Name	Job	Size	Actions
AS-NBR10-multi	Backup copy job	3.3 GB	Recover Verify Repair <b>Delete</b>


A 'Close' button is located at the bottom right of the interface.

## Deleting Backups in Bulk

To permanently delete several backups that match specific criteria, follow the steps below:

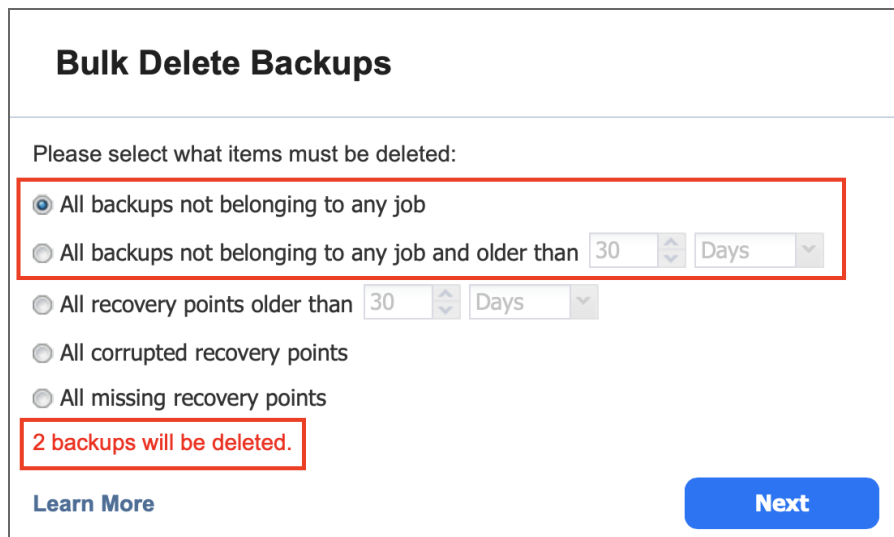
1. In the main menu, click **Settings**.
2. Go to the **Repositories** tab and hover over the Backup Repository you need.

3. Click **Manage** and then click **Delete backups in bulk**.



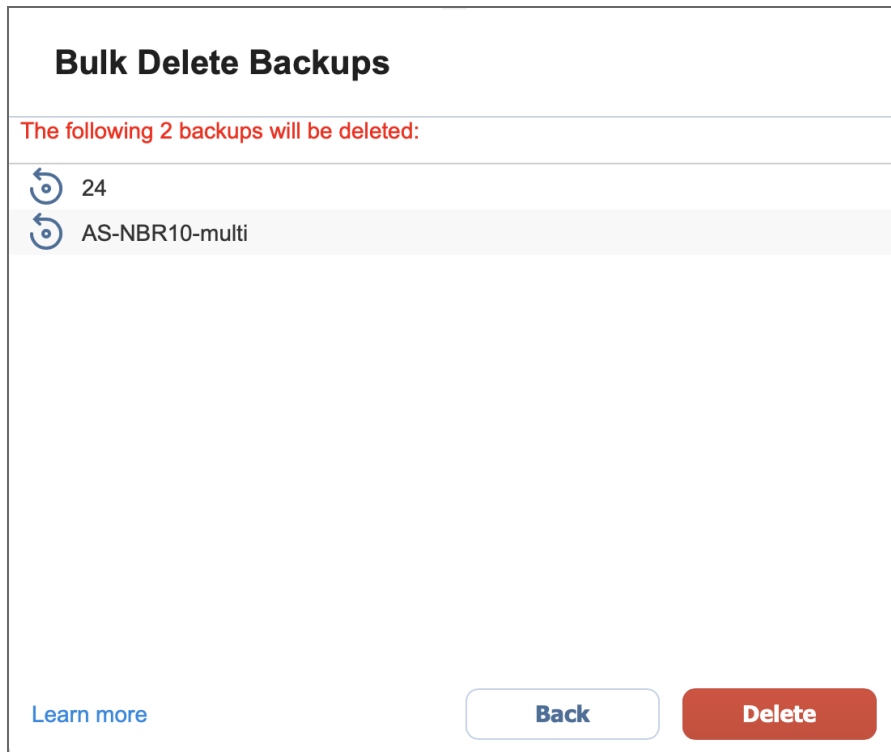
4. In the **Bulk Delete Backup** dialog box that opens, select one of the available options:
- **All backups not belonging to any job**
  - **All backups not belonging to any job and older than X <time\_units>**, where X is an integer and <time\_units> is either days, weeks, or months

The dialog shows the number of backups to be deleted.



5. Click **Next**.

- The **Bulk Delete Backups** dialog box opens displaying the list of backups to be deleted. Click **Delete** to confirm the deletion of backups.



### 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 delete a single recovery point, all corrupted recovery points, or all recovery points older than a specified number of days.

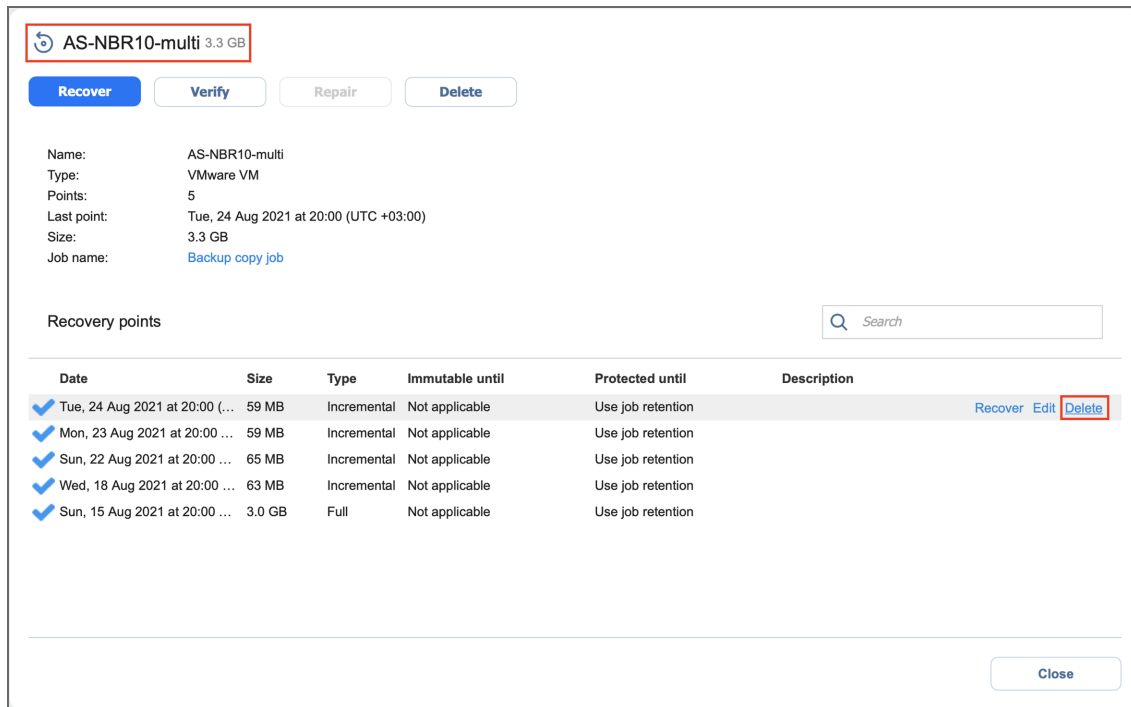
### Note

- If a backup is used by a VM/EC2 instance/physical machine backup job, you cannot delete the last recovery point in that backup.
- A recovery point cannot be deleted while the source VM/EC2 instance/physical machine backup job is running.

## 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**.



The screenshot shows a backup management interface for a repository named 'AS-NBR10-multi' (3.3 GB). It features buttons for 'Recover', 'Verify', 'Repair', and 'Delete'. Below these are details for the backup: Name (AS-NBR10-multi), Type (VMware VM), Points (5), Last point (Tue, 24 Aug 2021 at 20:00 (UTC +03:00)), Size (3.3 GB), and Job name (Backup copy job). A 'Recovery points' section contains a table with columns: Date, Size, Type, Immutable until, Protected until, and Description. The first row is highlighted, and its 'Delete' button is circled in red. A 'Close' button is at the bottom right.

Date	Size	Type	Immutable until	Protected until	Description
✓ Tue, 24 Aug 2021 at 20:00 (...)	59 MB	Incremental	Not applicable	Use job retention	Recover Edit <b>Delete</b>
✓ Mon, 23 Aug 2021 at 20:00 ...	59 MB	Incremental	Not applicable	Use job retention	
✓ Sun, 22 Aug 2021 at 20:00 ...	65 MB	Incremental	Not applicable	Use job retention	
✓ Wed, 18 Aug 2021 at 20:00 ...	63 MB	Incremental	Not applicable	Use job retention	
✓ Sun, 15 Aug 2021 at 20:00 ...	3.0 GB	Full	Not applicable	Use job retention	

5. Click **Delete Recovery Point** in the message 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](#).

For SaaS Backup Repositories, manually removing backup data may not return space to the operating system correctly.

## 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.

#### **Note**

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.

The dialog shows the number of recovery points to be deleted.

### Bulk Delete Backups

Please select what items must be deleted:

All backups not belonging to any job

All backups not belonging to any job and older than

All recovery points older than

All corrupted recovery points

All missing recovery points


1 recovery points will be deleted.

[Learn More](#)

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.

### Bulk Delete Recovery Points

The following 1 recovery points will be deleted:

 24	Wed, 22 Dec 2021 at 19:55 (UTC +02:00)
--	--

[Learn more](#) Back Delete

**Note**

For SaaS Backup Repositories, manual removal of backup data may not return space to the operating system correctly.



# Replication

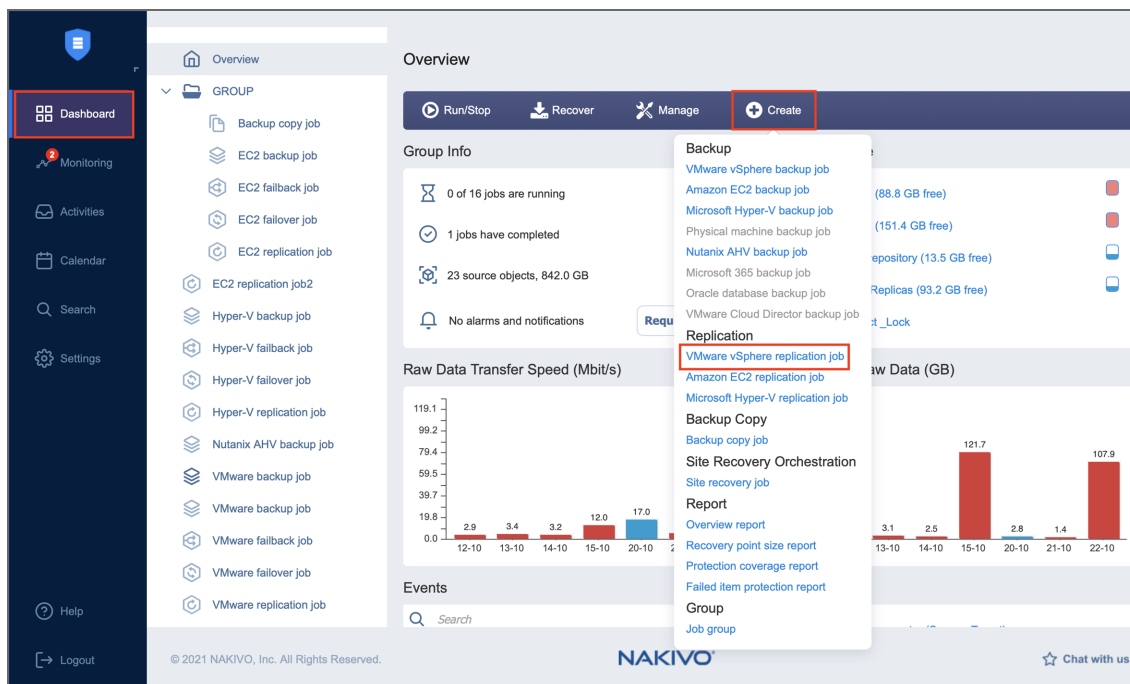
With NAKIVO Backup & Replication, you can perform replication of virtual machines. Replication creates and maintains an identical copy of the source VM/Instance at the target location.

Refer to the following topics for more information:

- [“Creating VMware Replication Jobs” below](#)
- [“Creating Hyper-V Replication Jobs” on page 745](#)
- [“Creating Amazon EC2 Replication Jobs” on page 773](#)
- [“Staging \(Seeding\) VM Replication” on page 786](#)

## 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 **Create** and then click **VMware vSphere replication job** on the **Dashboard**.



The **New Replication Job Wizard for VMware vSphere** opens. Complete the wizard to create a replication job.

- [“Replication Job Wizard for VMware: Source” on the next page](#)
- [“Replication Job Wizard for VMware: Destination” on page 719](#)
- [“Replication Job Wizard for VMware: Networks” on page 724](#)
- [“Replication Job Wizard for VMware: Re-IP” on page 726](#)
- [“Replication Job Wizard for VMware: Schedule” on page 729](#)

- [“Replication Job Wizard for VMware: Retention” on page 732](#)
- [“Replication Job Wizard for VMware: Options” on page 733](#)

## 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:

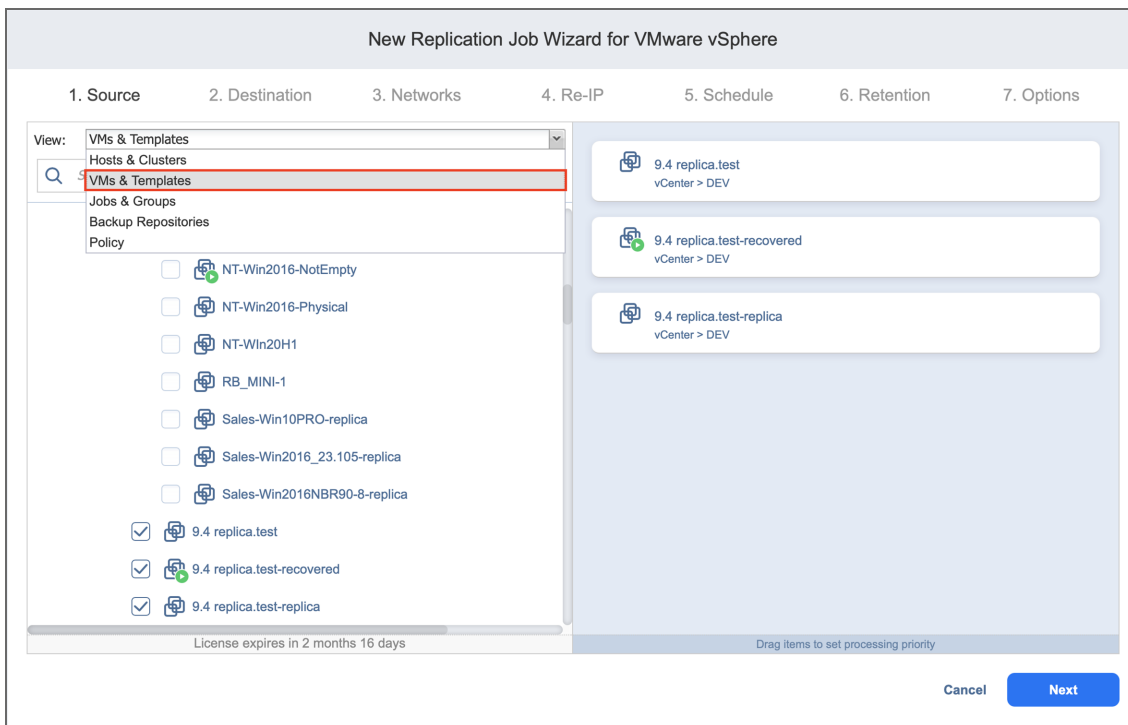
1. Optionally, you can filter the inventory tree by entering a string into the **Search** box. You can enter a part of or the entire name of the item.
2. Select VMware items by placing a checkmark to the left of each item.
3. The selected items appear in the right pane of the page. You can reorder the 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 the “x” on the right. This will deselect the item in the left pane.



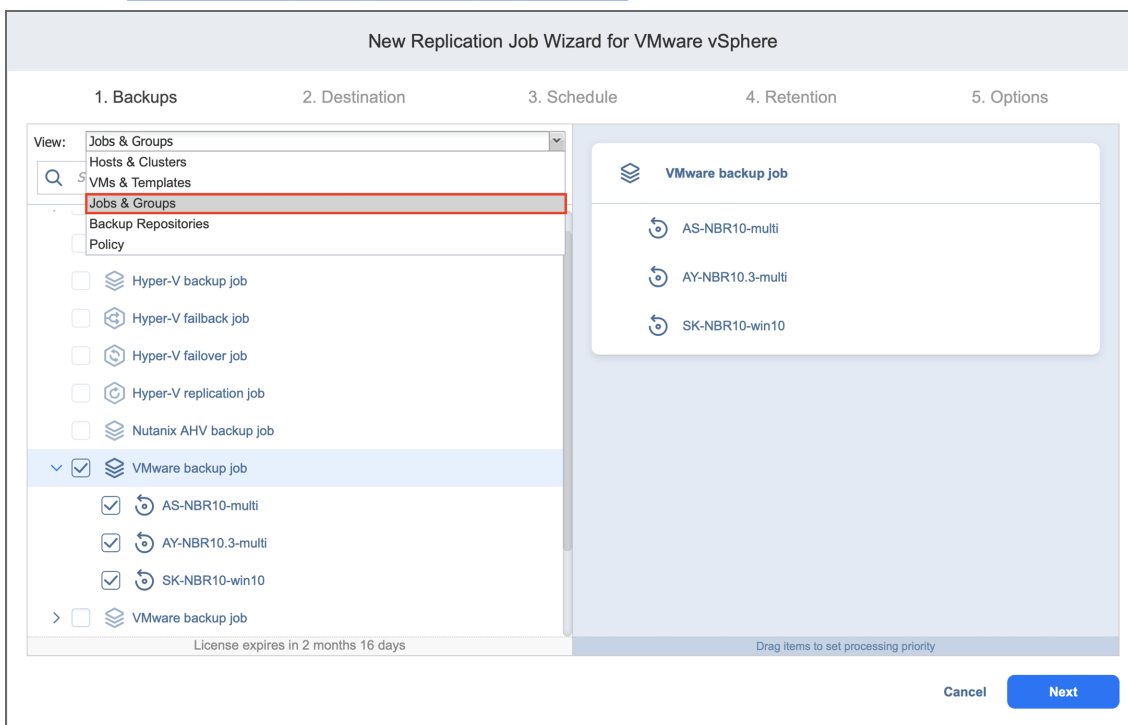
## VMs and Templates

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.



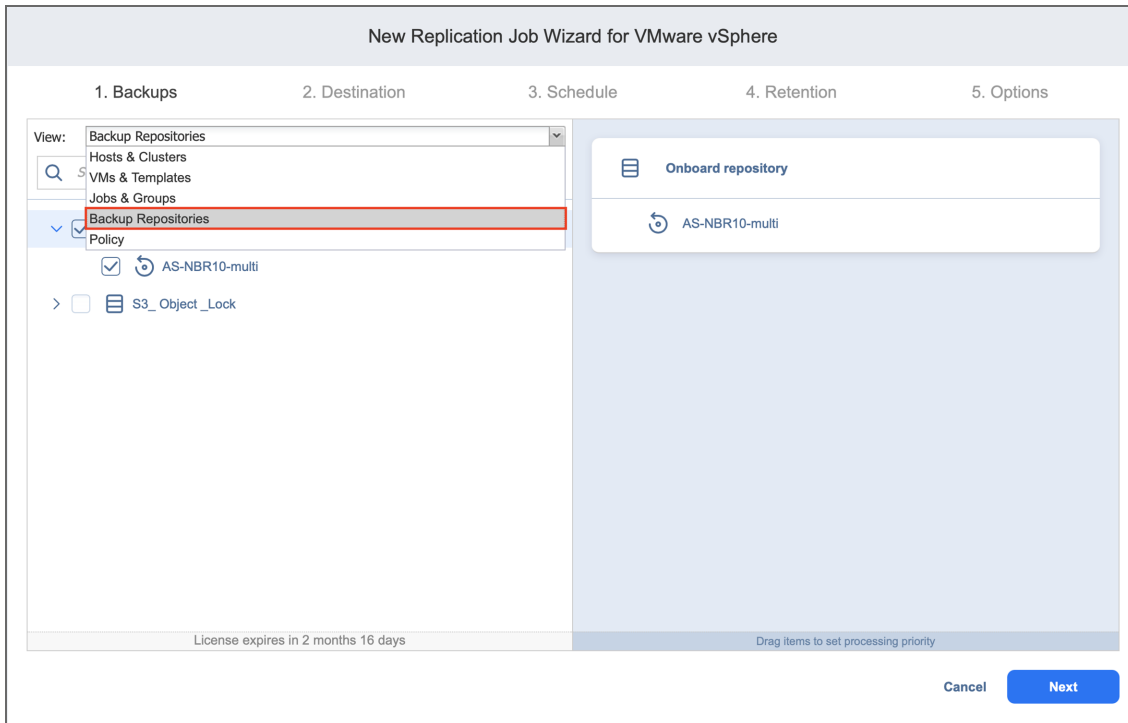
## Jobs and Groups

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 30](#)). 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.



## Policy

When the **Policy** view selected, job policies can be used. Refer to [“Managing Job Policies” on page 266](#) 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 269](#) 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 97](#) 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” below](#)
- [“Setting Different Options for Replicas” on the next page](#)
- [“Mapping Source VMs to Existing Replicas” on page 722](#)
- [“Excluding a VM Disk from the Replication Job” on page 723](#)

## 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.

The screenshot shows the 'New Replication Job Wizard for VMware vSphere' interface. The wizard is currently on the '1. Source' step. The 'Container' dropdown is set to '10.30.21.22', the 'Datastore' dropdown is set to '21.22-hdd', and the 'VM folder' dropdown is set to 'Discovered virtual machine'. A yellow information bar contains the text: 'To use existing VMs as targets, expand the Advanced setup and specify target VM for each source VM.' Below this bar is a link labeled 'Advanced setup...'. At the bottom right of the wizard, there are two buttons: 'Cancel' and 'Next'.

## 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**.

New Replication Job Wizard for VMware vSphere

1. Source    2. Destination    3. Networks    4. Re-IP    5. Schedule    6. Retention    7. Options

Container: 10.30.21.22  
Datastore: 21.22-hdd  
VM folder: Discovered virtual machine

**To use existing VMs as targets, expand the Advanced setup and specify target VM for each source VM.**

10.30.21.26 Click to collapse

Default container: 10.30.21.22  
Default datastore: 21.22-hdd  
Default VM folder: Discovered virtual machine

NBR10.4  
AP-NBR10.5

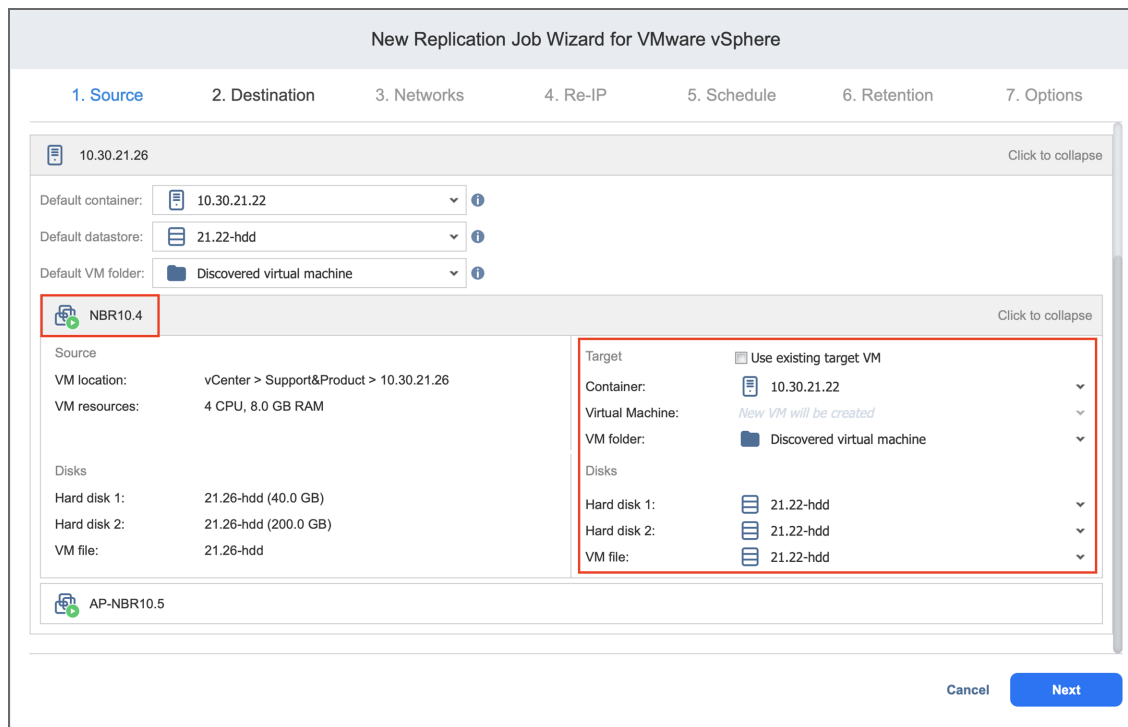
Cancel    Next

## 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.

The screenshot displays the 'New Replication Job Wizard for VMware vSphere' interface. It features a progress bar at the top with seven steps: 1. Source, 2. Destination, 3. Networks, 4. Re-IP, 5. Schedule, 6. Retention, and 7. Options. The 'Source' step is currently active.

The source configuration is for VM 'NBR10.4' with the following details:

Category	Value
VM location:	vCenter > Support&Product > 10.30.21.26
VM resources:	4 CPU, 8.0 GB RAM
Disks	
Hard disk 1:	21.26-hdd (40.0 GB)
Hard disk 2:	21.26-hdd (200.0 GB)
VM file:	21.26-hdd

The target configuration is for VM 'AP-NBR10.5' with the following details:

Category	Value
Target:	<input checked="" type="checkbox"/> Use existing target VM
Container:	10.30.21.22
Virtual Machine:	New VM will be created
VM folder:	Discovered virtual machine
Disks	
Hard disk 1:	21.22-hdd
Hard disk 2:	21.22-hdd
VM file:	21.22-hdd

At the bottom right, there are 'Cancel' and 'Next' buttons.

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.

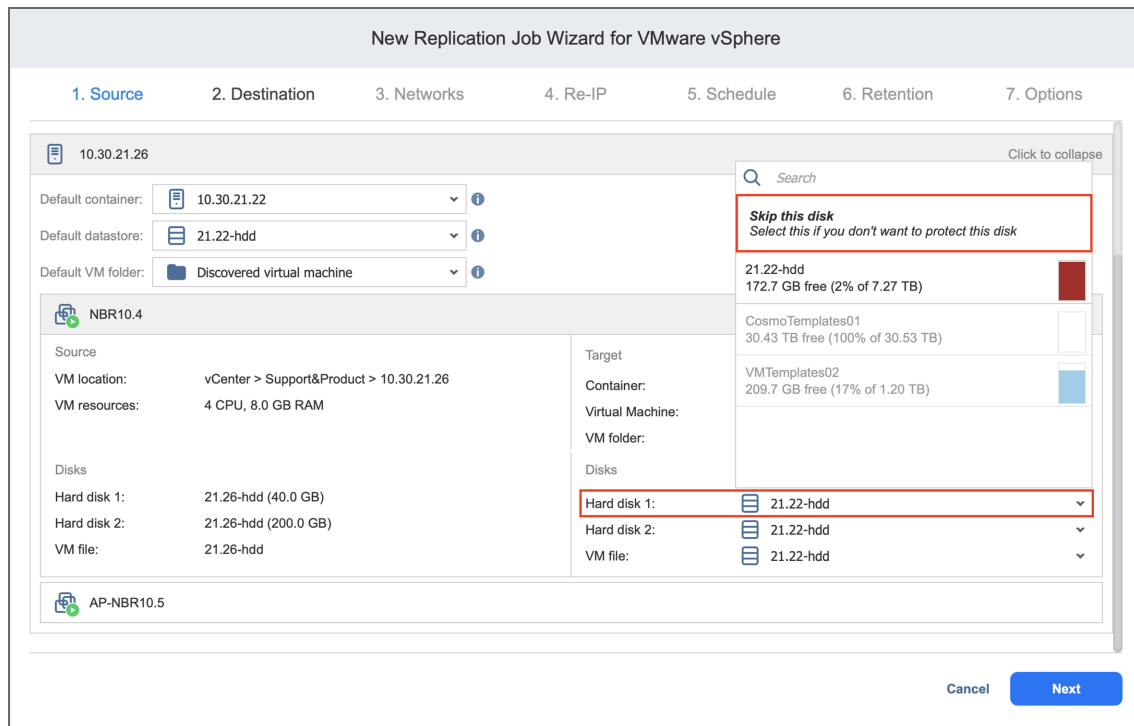
## Excluding a VM Disk from the Replication Job

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

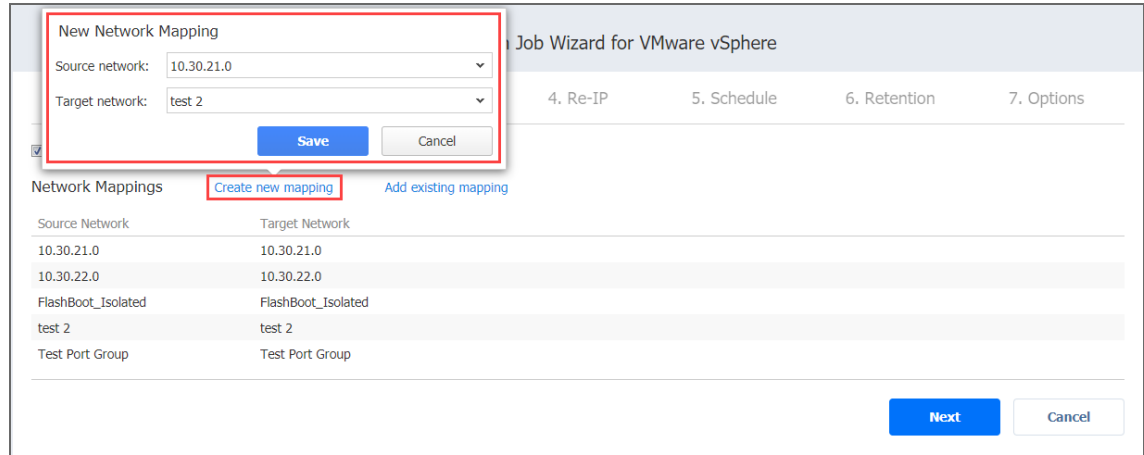
### Note

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.

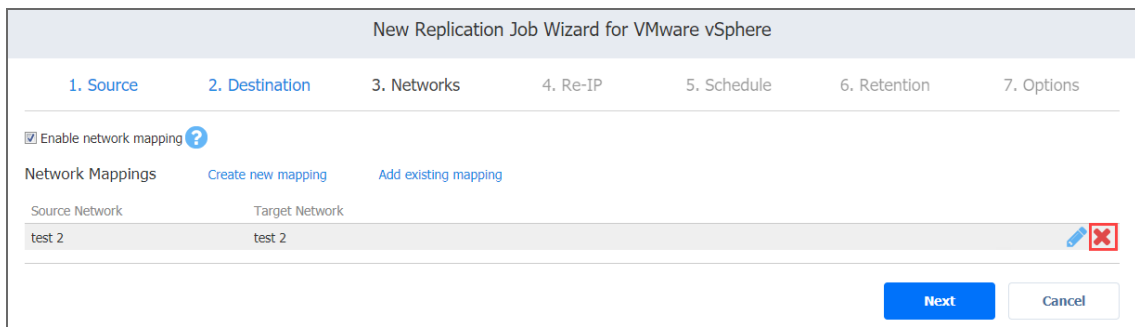


- Edit an existing mapping:
  1. Hover the pointer over the required item in the **Network Mappings** list and click the **Edit** button to the right of the item.

- 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 Delete icon (x) to the right of the item.



- Click **Next** to go to the next page of the wizard.

## Replication Job Wizard for VMware: Re-IP

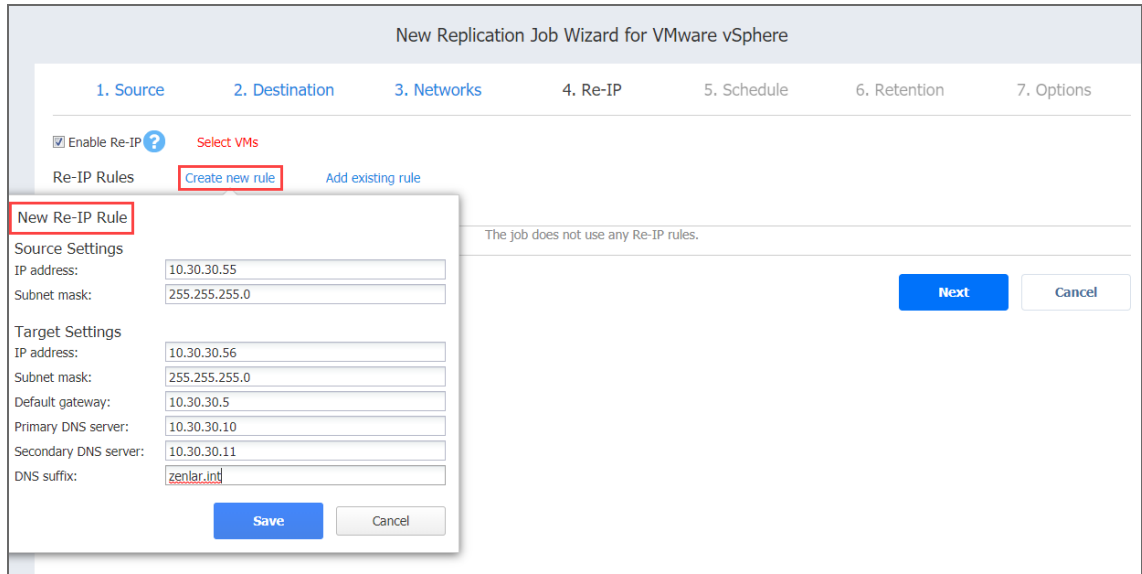
### Note

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:

- Select **Enable Re-IP**.
- Click the **Select VMs** link.
- The **Re-IP** dialog box opens. In the list of your source VMs, select at least one, and close the dialog box.
- You have the following options:

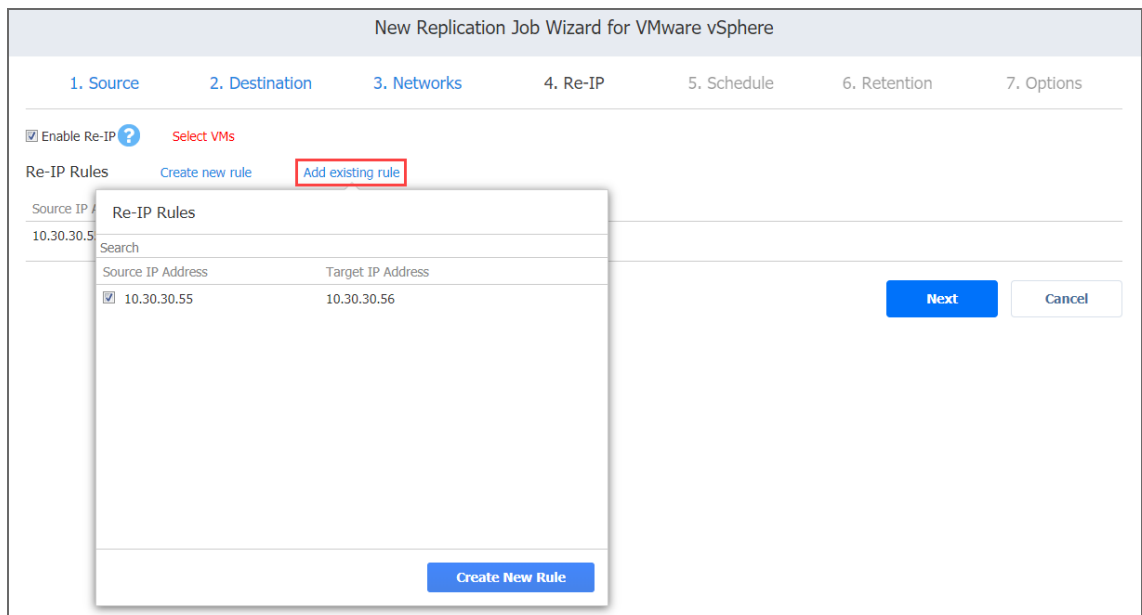
- Create a new rule:
  - a. Click **Create new rule**.
  - b. The **New Re-IP Rule** dialog box opens. Enter source and target settings for the Re-IP rule and click **Save**.



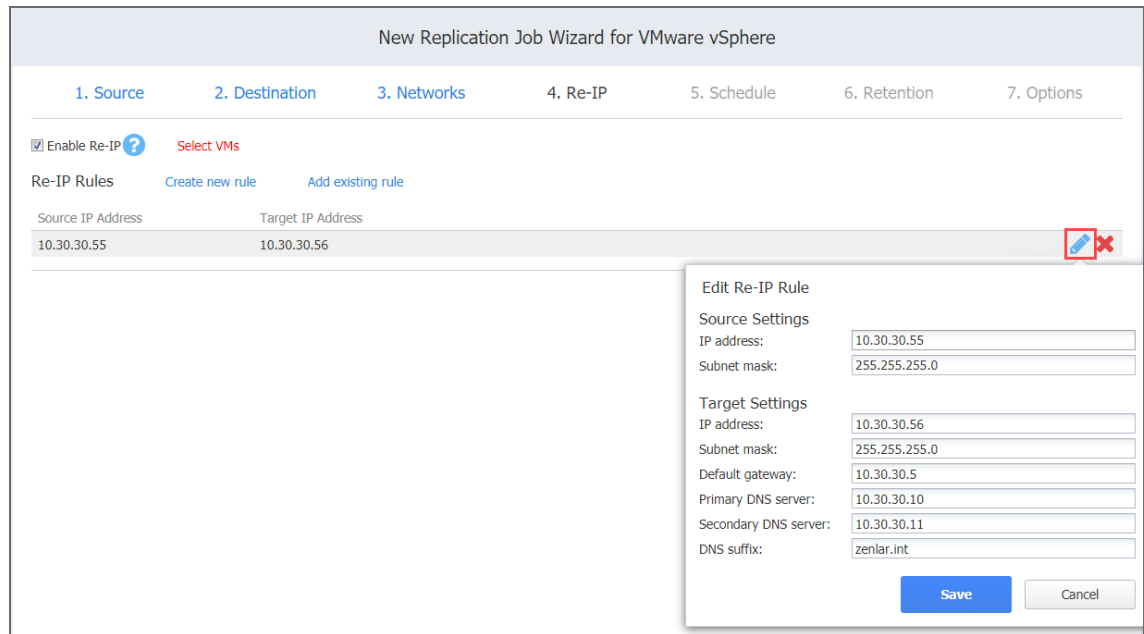
### Note

You can use wildcards for IP addresses. Refer to the [“Failover Job Wizard for VMware: Re-IP” on page 975](#) topic for details.

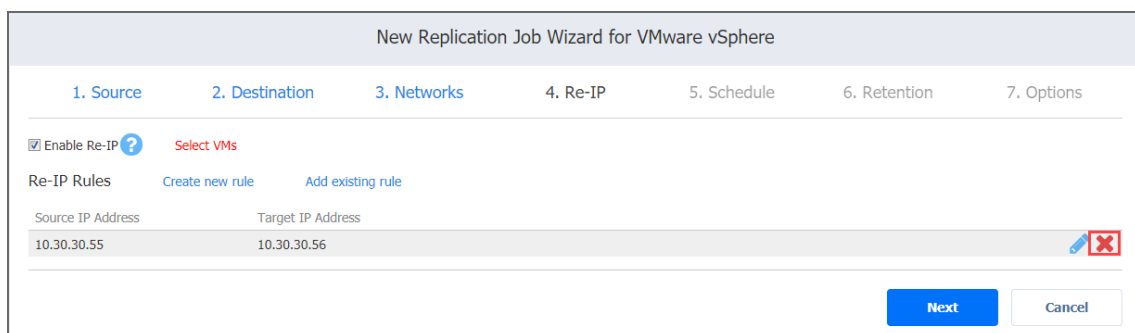
- Add an existing rule:
  - a. Click **Add existing rule**.
  - b. The **Re-IP Rules** dialog box opens. Select one or more appropriate Re-IP rules and close the dialog box.



- Edit an existing Re-IP rule:
  - a. Hover the pointer over the required item in the **Re-IP Rules** list and click the **Edit** button to the right of the item.
  - 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 975](#) for details.

# Replication 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.

- [“Disabling Scheduled Job Execution” below](#)
- [“Daily or Weekly Replication” below](#)
- [“Monthly or Yearly Replication” on the next page](#)
- [“Periodic Replication” on page 731](#)
- [“Chained Job” on page 731](#)
- [“Add Another Schedule” on page 732](#)

## Disabling Scheduled Job Execution

If you wish to start the job manually (without any schedule), select the **Do not schedule, run on demand** checkbox.

The screenshot shows the 'New Replication Job Wizard for VMware vSphere' interface. At the top, there are seven tabs: 1. Source, 2. Destination, 3. Networks, 4. Re-IP, 5. Schedule, 6. Retention, and 7. Options. The '5. Schedule' tab is active. Below the tabs, there is a checkbox labeled 'Do not schedule, run on demand' which is checked and highlighted with a red border. At the bottom right of the wizard, there are two buttons: 'Next' (a blue button) and 'Cancel' (a white button with a grey border).

## 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.

New Replication Job Wizard for VMware vSphere

1. Source 2. Destination 3. Networks 4. Re-IP 5. Schedule 6. Retention 7. Options

Do not schedule, run on demand  
 (UTC+02:00, EET) Eastern European Time

Schedule #1  
Run daily/weekly

Starting at: 0:00 Ending: 6:00

Mon  Tue  Wed  Thu  Fri  Sat  Sun  
 All days Work days Weekends

every 1 weeks

Effective from

[Add another schedule](#)  
[Show calendar](#)

**Next** Cancel

## 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.

New Replication Job Wizard for VMware vSphere

1. Source 2. Destination 3. Networks 4. Re-IP 5. Schedule 6. Retention 7. Options

Do not schedule, run on demand  
 (UTC+02:00, EET) Eastern European Time

Schedule #1  
Run monthly/yearly

Run every: last Friday of every month

Starting at: 0:00 Ending: 6:00

Effective from

[Add another schedule](#)  
[Show calendar](#)

**Next** Cancel

## 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.

The screenshot shows the 'New Replication Job Wizard for VMware vSphere' interface, specifically the '5. Schedule' step. The wizard has seven steps: 1. Source, 2. Destination, 3. Networks, 4. Re-IP, 5. Schedule, 6. Retention, and 7. Options. In the 'Schedule' step, there is a checkbox for 'Do not schedule, run on demand' which is unchecked. Below it is a time zone dropdown set to '(UTC+02:00, EET) Eastern European Time'. The 'Schedule #1' section features a dropdown menu with 'Run periodically' selected and highlighted by a red rectangle. To the right of this dropdown are the words 'every', a numeric input field containing '30', and a unit dropdown set to 'minutes'. Below these are 'Starting at' and 'Ending' time input fields, both set to '0:00' and '6:00' respectively. There are checkboxes for each day of the week: Mon, Tue, Wed, Thu, Fri (all checked), Sat, and Sun (unchecked). Below the day checkboxes are three links: 'All days', 'Work days', and 'Weekends'. There is also an 'Effective from' checkbox which is unchecked. At the bottom left, there are links for 'Add another schedule' and 'Show calendar'. At the bottom right, there are 'Next' and 'Cancel' buttons.

## 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.

The screenshot shows the '5. Schedule' step of the 'New Replication Job Wizard for VMware vSphere'. The wizard has seven steps: 1. Source, 2. Destination, 3. Networks, 4. Re-IP, 5. Schedule, 6. Retention, and 7. Options. In the '5. Schedule' step, there are several options:
 

- Do not schedule, run on demand
- A time zone dropdown menu set to '(UTC+02:00, EET) Eastern European Time'.
- 'Schedule #1' section:
  - A dropdown menu for 'Run after another job' is highlighted with a red box.
  - 'After the job:' dropdown menu set to 'Physical machine backup job'.
  - 'Run this job:' dropdown menu set to 'Immediately'.
  - Three checkboxes:  After successful runs,  After failed runs,  After stopped runs.
  - Effective from
- Links: 'Add another schedule' and 'Show calendar'.
- Buttons: 'Next' and 'Cancel'.

## 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.

This screenshot is identical to the one above, showing the '5. Schedule' step of the wizard. The 'Add another schedule' link is highlighted with a red box, indicating the next step in the process.

## Replication Job Wizard for VMware: Retention

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” on the next page](#)
  - [“Job Name” on the next page](#)
  - [“App-aware Mode” on page 735](#)
  - [“Change Tracking” on page 735](#)
  - [“Network Acceleration” on page 736](#)
  - [“Encryption” on page 736](#)
  - [“VM Verification” on page 736](#)
  - [“Exclude Swap Files and Partitions” on page 738](#)
  - [“Exclude Unused Blocks” on page 738](#)
  - [“Target Datastore Free Space” on page 738](#)
- [“Replica Options” on page 739](#)
  - [“Replica VM Disks” on page 739](#)
  - [“Replica VM Names” on page 739](#)
- [“Pre and Post Actions” on page 740](#)
  - [“Email Notifications” on page 740](#)
  - [“Microsoft Exchange Server Logs Truncation” on page 740](#)

- [“Microsoft SQL Server Logs Truncation” on page 741](#)
- [“Pre Job Script” on page 741](#)
- [“Post Job Script” on page 742](#)
- [“Data Transfer” on page 743](#)
  - [“Transport Mode” on page 743](#)
  - Transporter Pools
  - [“Transporters” on page 744](#)
  - [“Transporter Load” on page 744](#)
  - [“Bandwidth Throttling” on page 744](#)
- [“Completing the New Replication Job Wizard for VMware” on page 745](#)

## Job Options

In this section, you can specify a name for your replication job, and enable/disable [app-aware mode](#), change tracking, [network acceleration](#), [encryption](#), [“VM Verification” on page 48](#), and [excluding swap files](#), partitions and [unused blocks](#).

The screenshot shows the 'New Replication Job Wizard for VMware vSphere' interface, specifically the '7. Options' step. The 'Job Options' section is highlighted with a red box. The options are as follows:

- Job name:** VMware replication job
- App-aware mode:** Enabled (proceed on error)
- Change tracking:** Use VMware CBT
- Network acceleration:** Disabled
- Network encryption:** Disabled
- VM verification:** Disabled
- Exclude swap files and partitions:** Enabled
- Exclude unused blocks:** Enabled
- Target datastore minimum free space:** 10 %

Below the Job Options section, there are sections for 'Replica Options' and 'Pre and Post Actions':

- Replica Options:**
  - Replica VM disks: Respect original VM disk type
  - Replica VM names: Append "-replica" in the end
- Pre and Post Actions:**
  - Send job run reports to: (empty field)
  - Truncate Exchange logs: On successful VM processing only
  - Truncate SQL Server logs: On successful VM processing only
  - Run local pre job script: (checkbox)
  - Run local post job script: (checkbox)

At the bottom right, there are buttons for 'Cancel', 'Finish', and 'Finish & Run'.

## Job Name

Specify a name for the replication job.

## 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.

## Encryption

If 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 48](#) 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 294.](#)) 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 VM 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.

The screenshot shows a configuration window for VM backup with several tabs: 1. Source, 2. Destination, 6. Retention, and 7. Options. The 'Verification Options' dialog box is open, displaying the following settings:

- VM Boot Location:** VM replicas will be temporary started in the current replica location. VM replicas will not be connected to any networks.
- Verify not more than:** 2 VMs simultaneously
- Recovery time objective:** 5 minutes
- Screenshot delay:** 30 seconds
- Screenshot verification:** settings


The 'VM verification:' option in the main configuration is highlighted with a red box. At the bottom of the window, there are buttons for 'Cancel', 'Finish', and 'Finish & Run'.

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:
  - **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

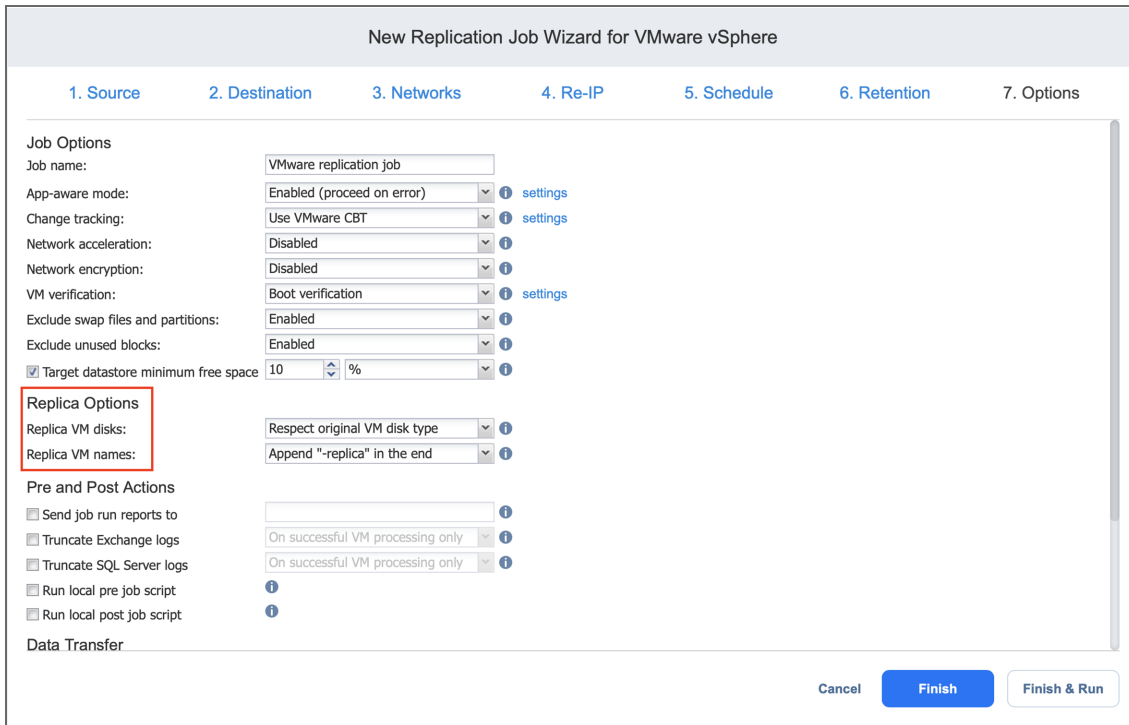
This feature is not available for replication from backup jobs.

## 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.

The screenshot shows the 'New Replication Job Wizard for VMware vSphere' interface, specifically the 'Options' step. The wizard is divided into seven steps: 1. Source, 2. Destination, 3. Networks, 4. Re-IP, 5. Schedule, 6. Retention, and 7. Options. The 'Options' step is currently active. The 'Pre and Post Actions' section is highlighted with a red box and contains the following options:

- Send job run reports to
- Truncate Exchange logs
- Truncate SQL Server logs
- Run local pre job script
- Run local post job script

Other options visible in the wizard include:

- VM verification: boot verification
- Exclude swap files and partitions: Enabled
- Exclude unused blocks: Enabled
- Target datastore minimum free space: 10 %
- Replica Options: Respect original VM disk type
- Replica VM names: Append "-replica" in the end
- Data Transfer: Transport mode: Automatic selection
- Limit transporter load to: 3 concurrent tasks
- Bandwidth throttling: Disabled

Buttons at the bottom right include 'Cancel', 'Finish', and 'Finish & Run'.

## 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.

The screenshot shows the 'New Replication Job Wizard for VMware vSphere' with seven steps: 1. Source, 2. Destination, 3. Networks, 4. Re-IP, 5. Schedule, 6. Retention, and 7. Options. The 'Options' step is active, and the 'Data Transfer' section is highlighted with a red box. The 'Data Transfer' section includes the following settings:

- Transport mode: Automatic selection
- Transporter pool: Select transporter pool
- Transporters: Automatic selection
- Limit transporter load to: 3 concurrent tasks
- Bandwidth throttling: Disabled
- Bottleneck detection:

Other settings in the wizard include VM verification (boot verification), Exclude swap files and partitions (Enabled), Exclude unused blocks (Enabled), Target datastore minimum free space (10%), Replica VM disks (Respect original VM disk type), Replica VM names (Append "-replica" in the end), and Pre and Post Actions (Send job run reports to, Truncate Exchange logs, Truncate SQL Server logs, Run local pre job script, Run local post job script).

## 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 284](#) 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 284](#) 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

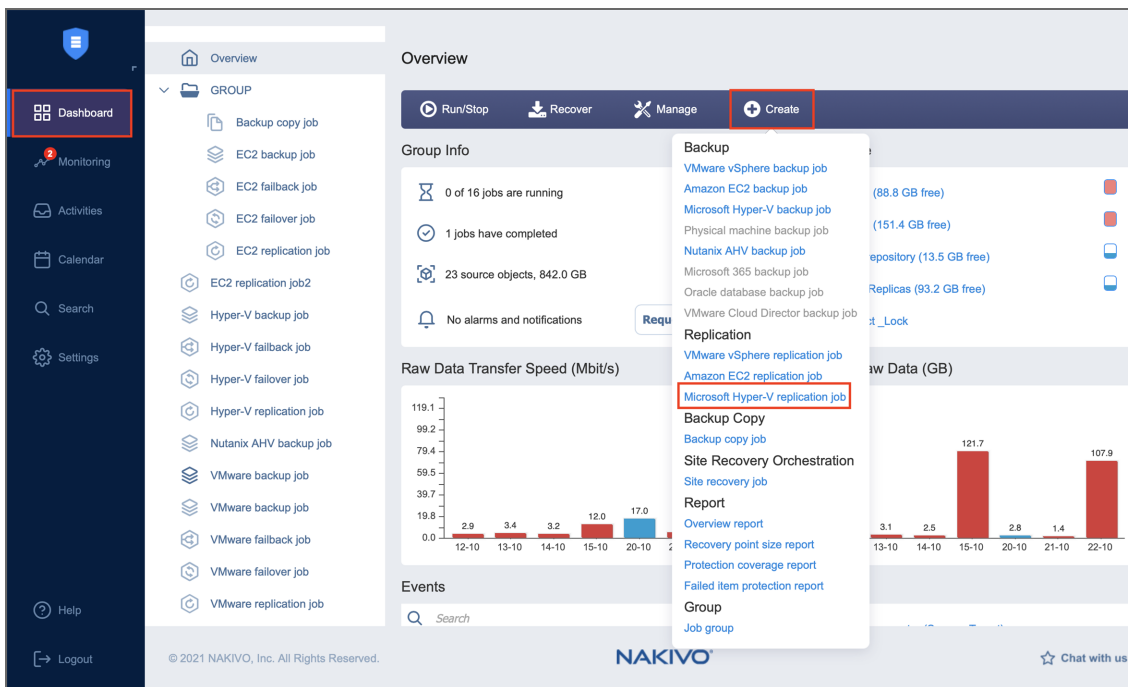
If you click **Finish & Run**, you will have to define the scope of your job. Please refer to [“Running Jobs on Demand” on page 256](#) for details.

## Creating Hyper-V Replication Jobs

With NAKIVO Backup & Replication, you can replicate Hyper-V VMs using the workflow with a broad range of options.

To create a replication job, click **Create** on the **Dashboard** and then click **Microsoft Hyper-V replication job**.





The **New Replication Job Wizard for Microsoft Hyper-V** opens. Complete the wizard to create a replication job.

- [“Replication Job Wizard for Hyper-V: Source” below](#)
- [“Replication Job Wizard for Hyper-V: Destination” on page 750](#)
- [“Replication Job Wizard for Hyper-V: Networks” on page 754](#)
- [“Replication Job Wizard for Hyper-V: Re-IP” on page 756](#)
- [“Replication Job Wizard for Hyper-V: Schedule” on page 758](#)
- [“Replication Job Wizard for Hyper-V: Retention” on page 762](#)
- [“Replication Job Wizard for Hyper-V: Options” on page 763](#)

## Replication Job Wizard for Hyper-V: Source


On the **Source** page in the wizard, you can add Hyper-V items to your replication using one of the inventory views. Proceed as described in these sections:

- [Hosts and Clusters](#)
- [Jobs and Groups](#)
- [Backup Repositories](#)
- [Policy](#)

### Host and Clusters


When the **Host & Clusters** view is selected, the inventory tree opens in the left pane and shows you all Hyper-V items: servers, clusters, 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 Hyper-V items by selecting the checkbox next to the 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 Hyper-V 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.
  - 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.
5. Click **Next** to confirm adding selected VMs to the replication job.



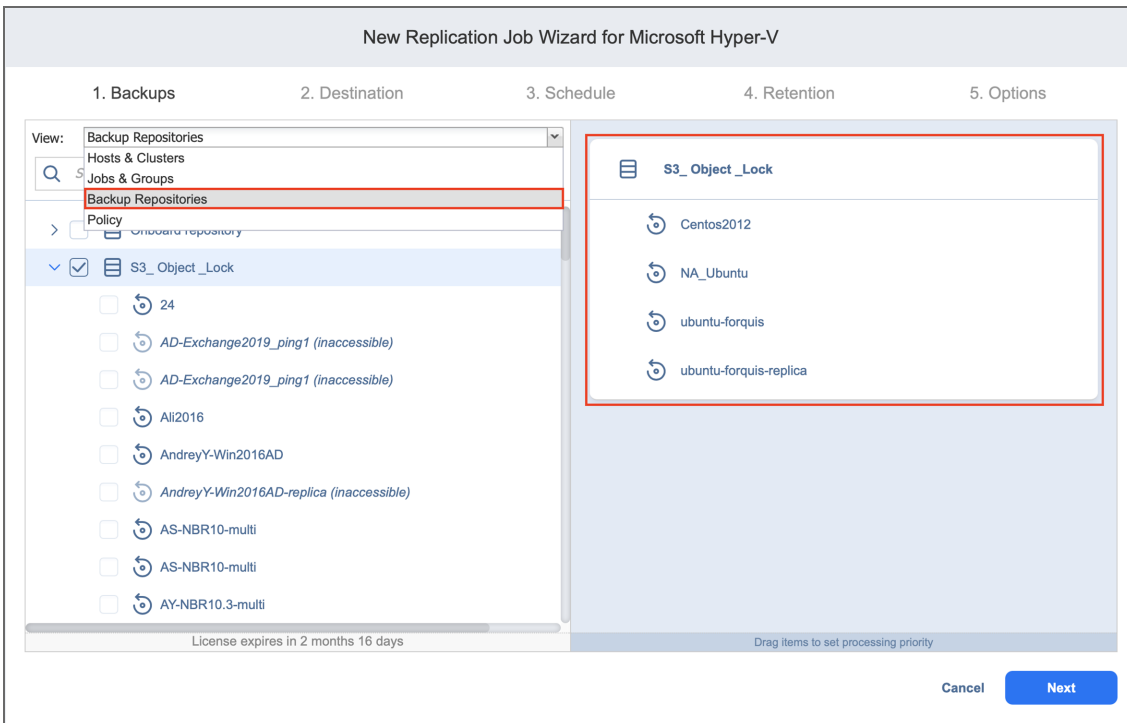
## Jobs and Groups

When the **Jobs & Groups** view is selected, the inventory tree shows groups, jobs, and backups of the appropriate hypervisor. Proceed as described for the **Hosts & Clusters** view above.




## Backup Repositories

When the **Backup Repositories** view is selected, the inventory tree shows the Backup Repositories that contain backups of the appropriate hypervisor. Proceed as 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 266](#) 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 269](#) for details.



## Notes

- If you cannot find a VM or a container:
  - Make sure the corresponding Hyper-V server or failover cluster has been added to the [inventory](#).
  - [Refresh the Inventory](#).
- If you add a Hyper-V server or cluster 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.
- If all disks of a VM are unsupported, this VM will be disabled in the inventory tree and it will not be possible to select it. Refer to [“Supported Platforms” on page 97](#) for details.
- 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.

# Replication Job Wizard for Hyper-V: Destination

On the **Destination** page, select a location for your Hyper-V replicas.

- [Setting the Same Host and Path for All Replicas](#)
- [Setting the Default Destination for Replicas](#)
- [Setting Different Options for Replicas](#)
- [Mapping Source VMs to Existing Replicas](#)
- [Excluding a VM Disk from the Replication Job](#)

## Setting the Same Host and Path for All Replicas

To replicate all VMs to the same server and location, and to connect all replicas to the same network:

1. Choose a server from the **Container** drop-down list.
2. Enter a path to the location where you want to store VM replicas in the **Path** field. It can be either a local or shared path.

### Notes

To connect to a shared path successfully, make sure that the following requirements are met:

- The shared path is created with the same credentials as the corresponding Hyper-V container. See [“Adding Microsoft Hyper-V Servers” on page 338](#) for details on adding Hyper-V containers to the inventory of NAKIVO Backup & Replication.
- The logon session in which you created the shared path has not ended. As a workaround, create a symbolic link to the shared path from the Hyper-V container. Refer to Step 7 in [High Availability of NAKIVO Backup & Replication](#) for details.

As a workaround, create a symbolic link to the shared path from the Hyper-V container. Refer to [High Availability of NAKIVO Backup & Replication – Step 7 of the procedure](#),– for an explanation.

3. If you are creating a Replication job from an existing Backup job, select the target network from the **Network** drop-down list. For more details, see [“Replication Job Wizard for Hyper-V: Source” on](#)

New Replication Job Wizard for Microsoft Hyper-V

1. Backups      2. Destination      3. Schedule      4. Retention      5. Options

Container:

Path:

Network:

**i** To use existing VMs as targets, expand the Advanced setup and specify target VM for each source VM.

[Advanced setup...](#)

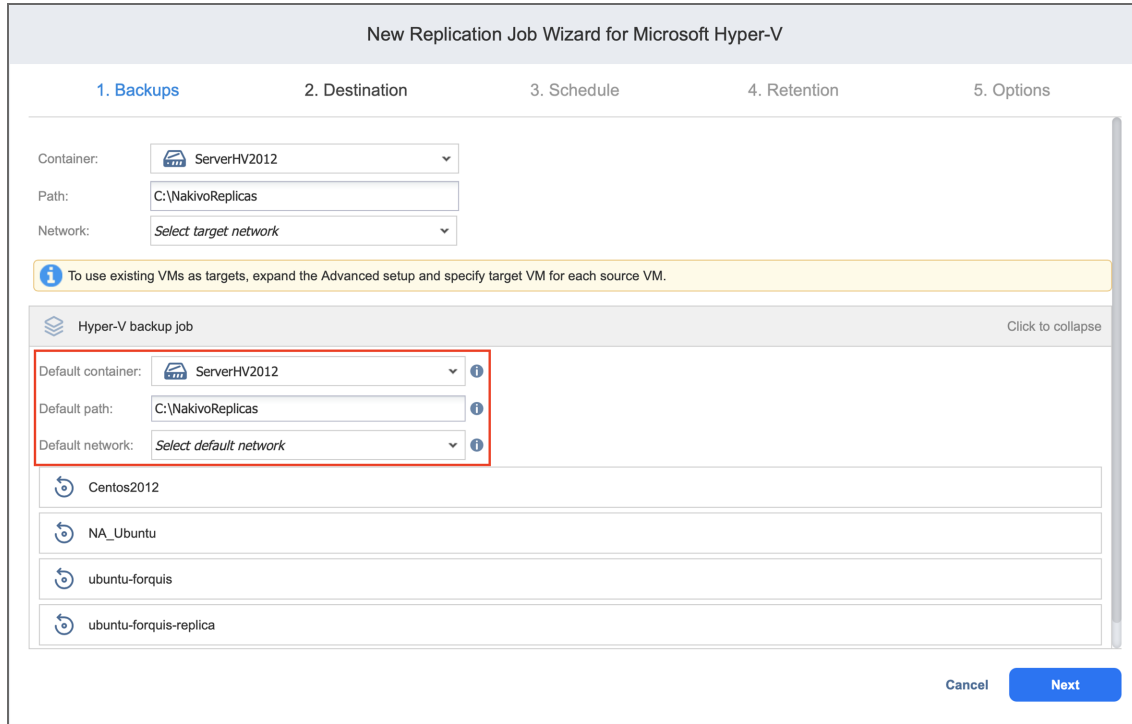
Cancel    **Next**

## 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 a default container, datastore, and VM folder for replicas. To do this:

1. Click **Advanced setup** and then click on the name of the chosen host, cluster, folder, or resource pool.
2. Choose a **Default container**.

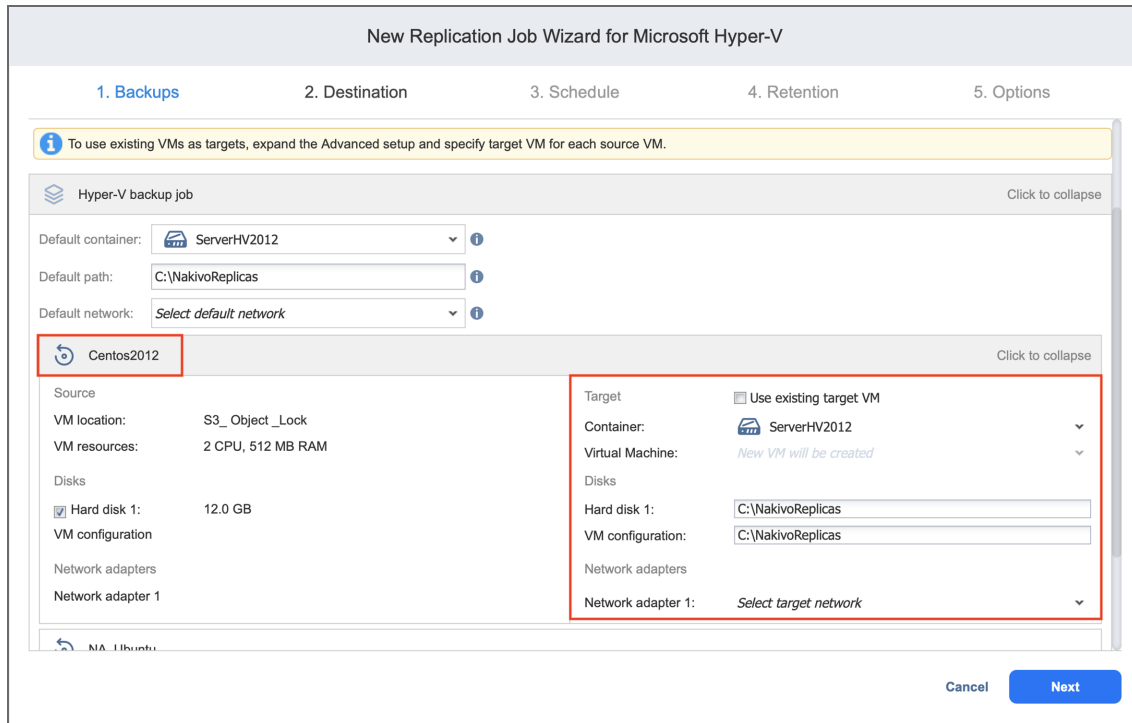
3. If you have chosen a Backup job on the **Source** page, you can choose a **Default Network**.



## Setting Different Options for Replicas

To specify different replication options for VMs, follow the steps below:

1. Click **Advanced setup**.
2. Set a target server, path, and network 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. Select the **Use existing target VM** option.
3. From the **Virtual Machine** drop-down list, select the VM that should be used as a target for replication.

The screenshot shows the 'New Replication Job Wizard for Microsoft Hyper-V' interface. The wizard is currently on step 1, 'Backups'. A yellow information banner at the top states: 'To use existing VMs as targets, expand the Advanced setup and specify target VM for each source VM.' Below this, the 'Hyper-V backup job' section is visible, with fields for 'Default container' (ServerHV2012), 'Default path' (C:\NakivoReplicas), and 'Default network' (Select default network). The 'Centos2012' source VM configuration is shown, including 'VM location' (S3\_Object\_Lock), 'VM resources' (2 CPU, 512 MB RAM), and 'Disks' (Hard disk 1: 12.0 GB). The 'Target' section is highlighted with a red box, showing the 'Use existing target VM' checkbox checked, the 'Container' dropdown set to 'ServerHV2012', and the 'Virtual Machine' dropdown set to 'Win2008'. Other target configuration fields include 'Hard disk 1' (C:\VM\), 'VM configuration' (C:\ProgramData\Microsoft\Windows\Hyper-V\Virtual Machine), and 'Network adapter 1' (New Virtual Switch). 'Cancel' and 'Next' buttons are at the bottom right.

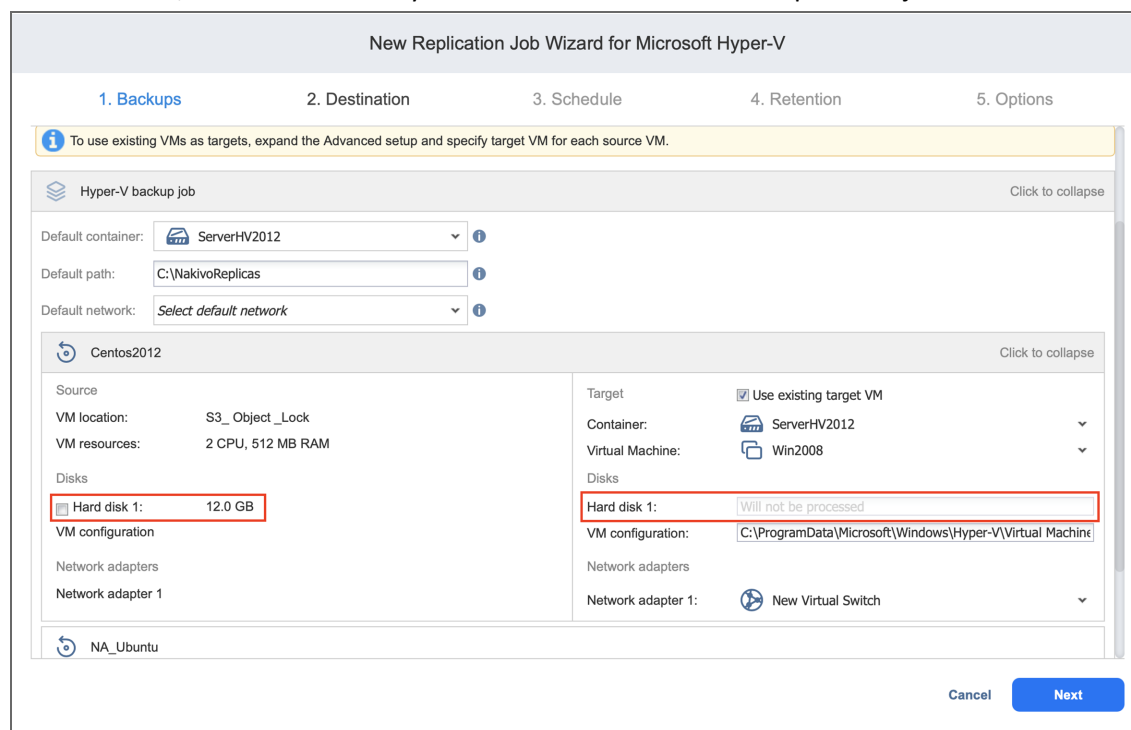
When you run the job, the product analyzes the target VM you have selected, determines how it is different from the source VM, and transfers only the differential data. VM replication mapping can be a time-consuming process that can be equal to the time required to create a full VM replication. After job completion, a new recovery point is created and existing recovery points are not 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. In the VM box, uncheck the disks you want to exclude from the replication job.




## Replication Job Wizard for Hyper-V: 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, please do the following on the **Networks** page of the wizard:

1. Select **Enable network mapping**.

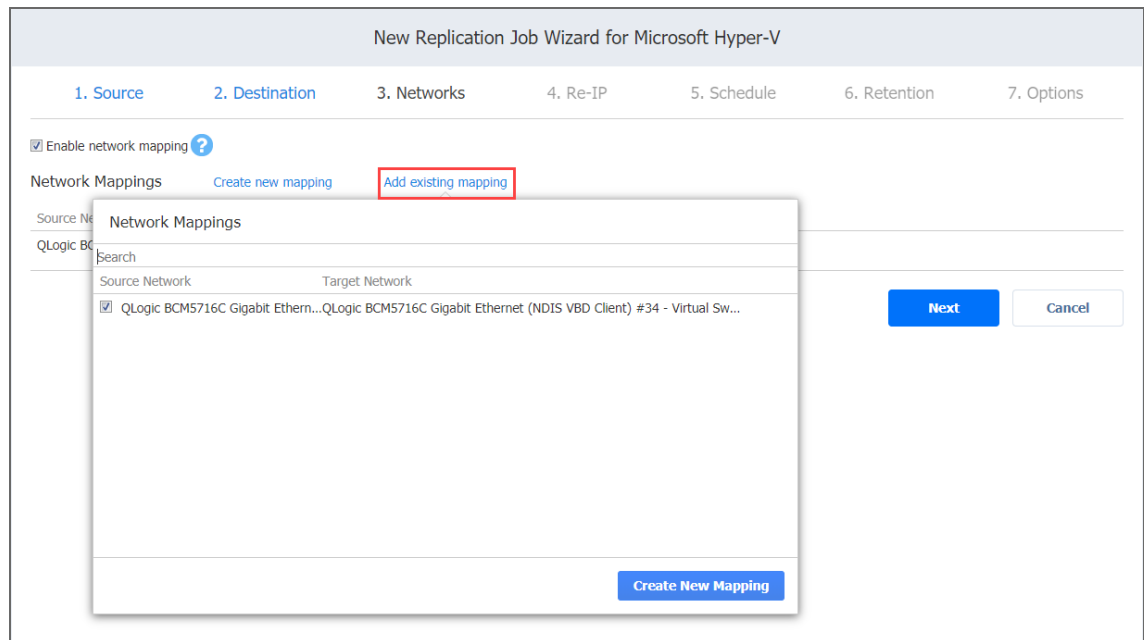


2. The **Network Mappings** section opens. You have the following options:

- Create a new mapping:
  - a. Click **Create new mapping**.
  - b. The **New Network Mapping** dialog box opens. Choose a source network and target network and click **Save**.

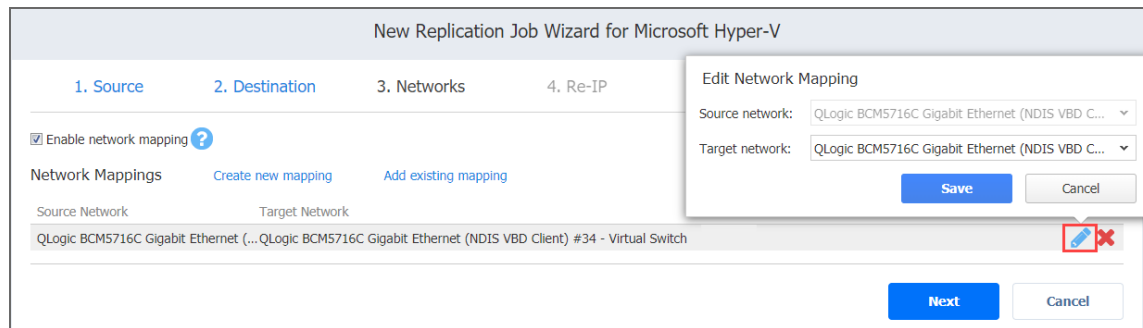


- Add an existing mapping:
  - a. Click **Add existing mapping**.
  - b. The **Network Mappings** dialog box opens. Choose one or more appropriate network mappings and close the dialog box.



- Edit an existing mapping:
  - a. Hover the pointer over the required item in the **Network Mappings** list and click the **Edit** button to the right of the item.
  - b. The **Edit Network Mapping** dialog box opens. Choose the appropriate 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 delete icon to the right of the item.
3. Click **Next** to go to the next page of the wizard.


## Replication Job Wizard for Hyper-V: Re-IP

### Notes

- The **Re-IP** page is skipped if you have chosen an existing backup as the target for your replication job on the **Source** page.
- 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 Microsoft Hyper-V: Re-IP” on page 982](#) for details.

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. 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, 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 Microsoft Hyper-V: Re-IP” on page 982](#) 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 **Edit** button to the right of the item.
- 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.

5. Click **Next** to go to the next page of the wizard.

## Replication Job Wizard for Hyper-V: 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.

- [Disabling Scheduled Job Execution](#)
- [Daily or Weekly Replication](#)
- [Monthly or Yearly Replication](#)
- [Periodic VM Replication](#)
- [Chained Job](#)
- [Additional Schedule](#)

### 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.

The screenshot shows the 'New Replication Job Wizard for Microsoft Hyper-V' at step 5, 'Schedule'. The wizard has seven steps: 1. Source, 2. Destination, 3. Networks, 4. Re-IP, 5. Schedule, 6. Retention, and 7. Options. In the 'Schedule' step, there is a checkbox for 'Do not schedule, run on demand' which is unchecked. Below it is a time zone dropdown menu set to '(UTC+02:00, EET) Eastern European Time'. Under 'Schedule #1', a dropdown menu is highlighted with a red box and contains the text 'Run daily/weekly'. Below this, there are input fields for 'Starting at: 0:00' and 'Ending: 6:00'. There are also checkboxes for days of the week: Mon, Tue, Wed, Thu, Fri, Sat, and Sun. The 'All days' option is selected. Below the day checkboxes, there is a section for 'every 1 weeks' with a dropdown for the number of weeks. There is also an 'Effective from' checkbox which is unchecked. At the bottom right, there are 'Next' and 'Cancel' buttons.

## 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.

New Replication Job Wizard for Microsoft Hyper-V

1. Source 2. Destination 3. Networks 4. Re-IP 5. Schedule 6. Retention 7. Options

Do not schedule, run on demand  
(UTC+02:00, EET) Eastern European Time

Schedule #1  
**Run monthly/yearly**

Run every last Friday of every month

Starting at: 0:00 Ending: 6:00

Effective from

[Add another schedule](#)  
[Show calendar](#)

**Next** Cancel

## Periodic VM 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** 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.

New Replication Job Wizard for Microsoft Hyper-V

1. Source 2. Destination 3. Networks 4. Re-IP 5. Schedule 6. Retention 7. Options

Do not schedule, run on demand  
(UTC+02:00, EET) Eastern European Time

Schedule #1  
**Run periodically** every 30 minutes

Starting at: 0:00 Ending: 6:00

Mon  Tue  Wed  Thu  Fri  Sat  Sun  
All days Work days Weekends

Effective from

[Add another schedule](#)  
[Show calendar](#)

**Next** Cancel

## 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.

The screenshot shows the 'New Replication Job Wizard for Microsoft Hyper-V' at the '5. Schedule' step. The wizard has seven steps: 1. Source, 2. Destination, 3. Networks, 4. Re-IP, 5. Schedule, 6. Retention, and 7. Options. The '5. Schedule' step is active. The options are as follows:

- Do not schedule, run on demand
- (UTC+02:00, EET) Eastern European Time
- Schedule #1: Run after another job (highlighted with a red box)
- After the job: Hyper-V backup job 2
- Run this job: Immediately
- After successful runs
- After failed runs
- After stopped runs
- Effective from

At the bottom, there are links for 'Add another schedule' and 'Show calendar', and buttons for 'Next' and 'Cancel'.

## Additional 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.



New Replication Job Wizard for Microsoft Hyper-V

1. Source   2. Destination   3. Networks   4. Re-IP   5. Schedule   6. Retention   7. Options

Do not schedule, run on demand

(UTC+02:00, EET) Eastern European Time

Schedule #1

Run after another job

After the job: Hyper-V backup job 2

Run this job: Immediately

After successful runs    After failed runs    After stopped runs

Effective from

Add another schedule

Show calendar

Next   Cancel

## Replication Job Wizard for Hyper-V: Retention

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.

New Replication Job Wizard for Microsoft Hyper-V

1. Source   2. Destination   3. Networks   4. Re-IP   5. Schedule   6. Retention   7. Options

Keep 10 last recovery points

Keep one recovery point per day for 10 days

Keep one recovery point per week for 4 weeks

Keep one recovery point per month for 12 months

Keep one recovery point per year for 3 years

[Learn more](#)

Next   Cancel

# Replication Job Wizard for Hyper-V: Options

On the **Options** page of the wizard, set up job options as described in the following sections:

- [“Job Options” below](#)
- [“Replica Options” on page 767](#)
- [“Pre and Post Actions” on page 768](#)
- [“Data Transfer” on page 770](#)
- [“Completing the New Replication Job Wizard for Microsoft Hyper-V” on page 772](#)

## Job Options

In this section, you can specify a name for your replication job, and enable/disable [app-aware mode](#), change tracking, [network acceleration](#), [encryption](#), [VM verification](#), [excluding swap files](#), partitions and [unused blocks](#).

The screenshot shows the 'New Replication Job Wizard for Microsoft Hyper-V' interface. The 'Options' step is selected, and the 'Job Options' section is highlighted with a red box. The 'Job name' is 'Hyper-V replication job'. 'App-aware mode' is 'Enabled (proceed on error)'. 'Use agent for OS quiescing' is 'Disabled'. 'Change tracking' is 'Use Hyper-V RCT'. 'Network acceleration' is 'Disabled'. 'Network encryption' is 'Disabled'. 'VM verification' is 'Disabled'. 'Exclude swap files and partitions' is 'Enabled'. 'Exclude unused blocks' is 'Enabled'. Below this, 'Replica Options' includes 'Respect original VM disk type' and 'Append "-replica" in the end'. 'Pre and Post Actions' includes 'Send job run reports to', 'Truncate Exchange logs', 'Truncate SQL Server logs', 'Run local pre job script', and 'Run local post job script'. 'Data Transfer' is also visible. At the bottom, there are 'Cancel', 'Finish', and 'Finish & Run' buttons.

## Job Name

Specify a name for the replication job.

## App-aware Mode

With the **App-aware mode** option enabled, VM processing is performed using guest OS quiescing to ensure the consistency of application data.

## Note

This option is not available for existing backups chosen as the target of replication on the **Source** wizard page.

We recommend reading the support articles about the requirements for App-aware mode:

- [Hyper-V Server configuration](#)
- [Linux and FreeBSD guest OS configuration.](#)

Select one of the following options from the **App-aware mode** drop-down list:

- **Enabled (proceed on error):** In the dialog box that opens, select the checkboxes next to the VMs for which you want to create application-aware replicas, and then select the credentials next to each VM. These credentials will be used to log into the VMs you have selected and trigger the VSS service. With this option selected, NAKIVO Backup & Replication proceeds with the replication even if the app-aware mode fails (for example, because of wrong credentials).
- **Enabled (fail on error):** In the dialog box that opens, select the checkboxes next to the VMs for which you want to create application-aware replicas, and then select the credentials next to each VM. These credentials will be used to log into the VMs you have selected and trigger the VSS service. With this option selected, NAKIVO Backup & Replication fails the replication if the app-aware mode fails (for example, because of wrong credentials).
- **Disabled:** VM replication is performed without application awareness.

## Use Agent for OS Quiescing

If this option is enabled, NAKIVO Backup & Replication injects an agent into VMs to perform OS quiescing. This option is only available with App-aware mode enabled.

## Note

Enabling this option may generate additional data transfer for the creation of shadow copies during job runs.

## Change Tracking

Select one of the options from the **Change tracking** drop-down list:

- **Use Hyper-V RCT:** It is the native Microsoft Hyper-V change block tracking method. You can select this option for the fastest incremental backup.
- **Use proprietary method:** With this option selected, NAKIVO Backup & Replication reads all the contents of all VM disks to determine which data has changed since the last job run.
- **No change tracking (always full):** With this option selected, NAKIVO Backup & Replication always performs a full VM replication 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.

## Encryption

With Encryption enabled, VM data is protected with AES 256 encryption while traveling over the network. Data encryption increases replication time and CPU load on machines running Transporters. Select this option when replicating over WAN without a VPN connection.

You need at least one Transporter at the 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.

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 294.](#)) 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 VM Boot, disables networking to prevent network connections, and verifies that system start is successful.

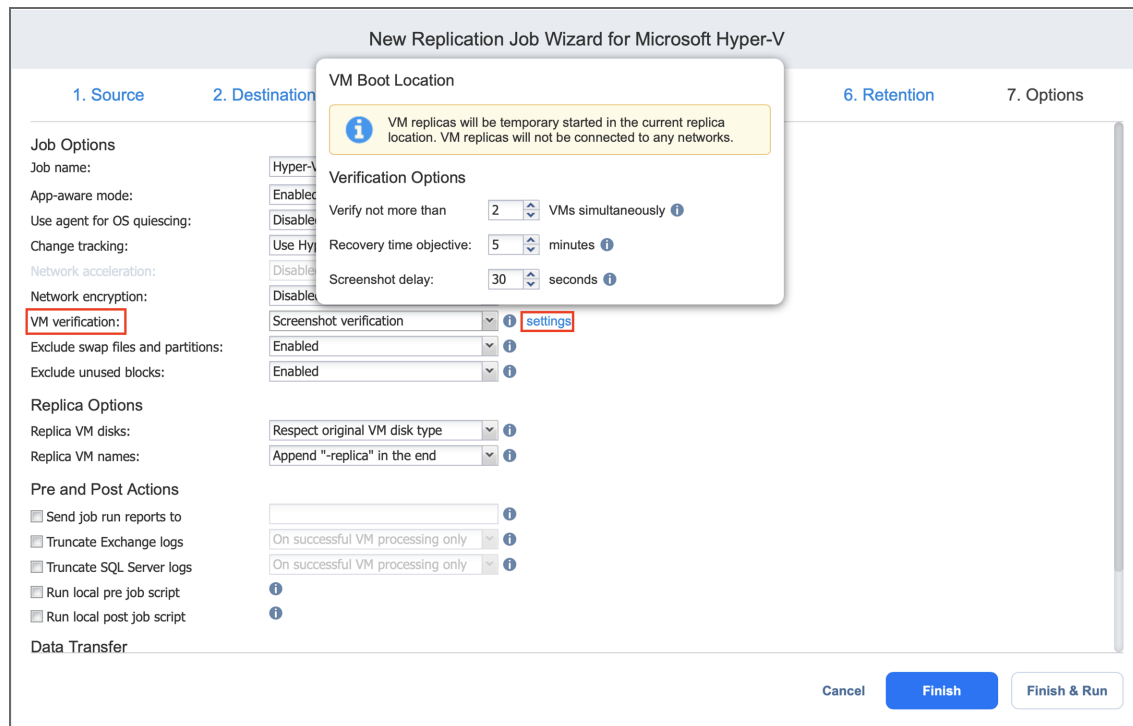
## Important

Hyper-V Integration Services must be running on the source VMs to successfully enable screenshot verification for your backup job.

After choosing **Screenshot verification** option, do the following in the dialog box that opens:

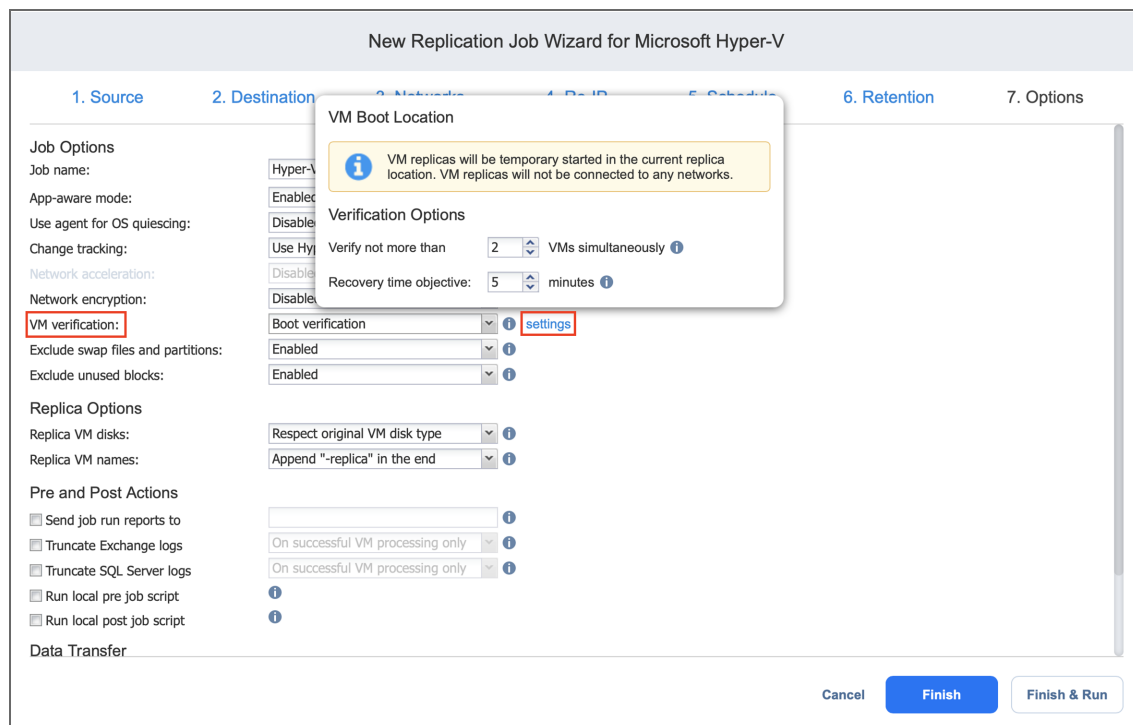
1. **Verify not more than x VMs simultaneously:** Specify the maximum number of VM replicas that can be started simultaneously.
2. **Recovery time objective:** Specify the amount of time allocated for the verification of each VM replica. If a VM OS does not start within the specified amount of time, verification will be considered failed. Specified time must be sufficient to fully start the VM OS. Try increasing this amount if the default amount is not sufficient.
3. **Screenshot delay:** Specify the amount of time that NAKIVO Backup & Replication waits before taking a

screenshot after the guest OS starts.



After choosing **Boot verification**, set verification options in the dialog box that opens:

1. **Verify not more than x VMs simultaneously:** Specify the maximum number of VMs that can be started on the target container simultaneously.
2. **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.



## Exclude Swap Files and Partitions

With this option enabled, NAKIVO Backup & Replication automatically excludes swap files and partitions during the backup process.

### Note

The feature is not available for replication from backup jobs.

## 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.

### Note

The feature is not available for replication from backup jobs.

## Replica Options

In the *Replica Options* section, you can choose a VM disk type and VM name.

The screenshot shows the 'New Replication Job Wizard for Microsoft Hyper-V' interface. At the top, there are seven tabs: 1. Source, 2. Destination, 3. Networks, 4. Re-IP, 5. Schedule, 6. Retention, and 7. Options. The 'Options' tab is selected. Below the tabs, there are several sections of configuration options. The 'Replica Options' section is highlighted with a red box. It contains two dropdown menus: 'Replica VM disks' set to 'Respect original VM disk type' and 'Replica VM names' set to 'Append "-replica" in the end'. Other options include 'Job name' (Hyper-V replication job), 'App-aware mode' (Enabled (proceed on error)), 'Use agent for OS quiescing' (Disabled), 'Change tracking' (Use Hyper-V RCT), 'Network acceleration' (Disabled), 'Network encryption' (Disabled), 'VM verification' (Disabled), 'Exclude swap files and partitions' (Enabled), and 'Exclude unused blocks' (Enabled). At the bottom, there are 'Pre and Post Actions' checkboxes and a 'Data Transfer' section. The bottom right corner has 'Cancel', 'Finish', and 'Finish & Run' buttons.

## 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 dynamic size disks on target VMs:** With this option enabled, dynamic size disks are created on replicas, occupying only space actually required by VM files, including OS ones and database records, regardless of the disk types of the original source VM.

## Replica VM Names

With NAKIVO Backup & Replication, you can change VM replica names to easily distinguish between 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, choose one of the following **Replica VM names** options in the **Replica Options** section:

- **Append "-replica" in the end:** Source VM names are used for replica names and “-replica” is added before or after 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, select to [truncate Microsoft Exchange and SQL Server logs](#), and run [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.

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

With NAKIVO Backup & Replication, you can 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 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.

## Truncation of Microsoft SQL Server Transaction Logs

With NAKIVO Backup & Replication, you can 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, do the following:

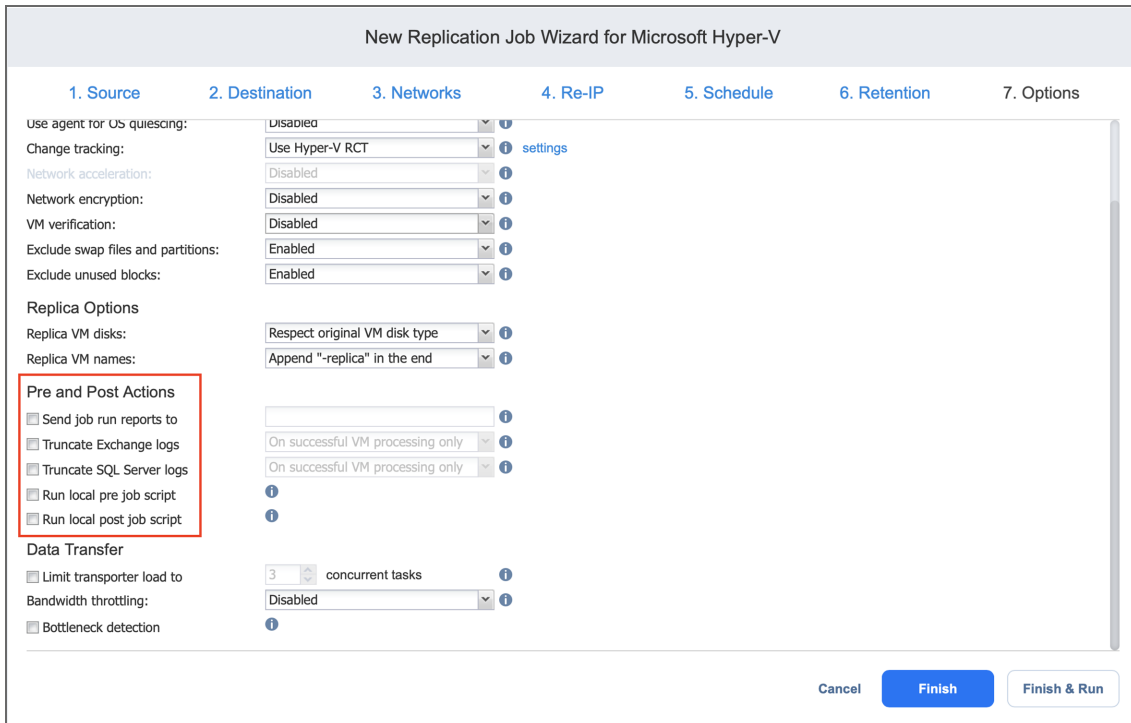
1. Place a script file on the machine where the Director is installed.
2. In the *Pre and Post Actions* section, select **Run local pre job script**.
3. Specify the following parameters in the dialog box that opens:
  - **Script path:** Specify a local path to the script on the machine where the Director 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, VM replication is only started after the script is completed.
    - **Do not wait for the script to finish:** With this option selected, the script runs and the VMs start replicating 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 replication even if the script has failed.
    - **Fail the job on script failure:** With this option selected, the job is failed and VM replication is not performed.

## 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 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, if 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 on the **Options** page, you can set the transporter load and configure [bandwidth throttling](#).

The screenshot shows the 'New Replication Job Wizard for Microsoft Hyper-V' with the 'Options' tab selected. The 'Data Transfer' section is highlighted with a red box. It contains the following options:

- Limit transporter load to
- Bandwidth throttling: Disabled
- Bottleneck detection

The 'concurrent tasks' field is set to 3. Other options in the wizard include 'Use agent for OS quiescing' (Disabled), 'Change tracking' (Use Hyper-V RCT), 'Network acceleration' (Disabled), 'Network encryption' (Disabled), 'VM verification' (Disabled), 'Exclude swap files and partitions' (Enabled), 'Exclude unused blocks' (Enabled), 'Replica VM disks' (Respect original VM disk type), and 'Replica VM names' (Append "-replica" in the end).

## 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 284](#) 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:
    1. Click the **Create New Rule** button.
    2. The **New Bandwidth Rule** dialog box opens. Refer to the [“Bandwidth Throttling” on page 284](#) topic for details about 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, 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 Microsoft Hyper-V

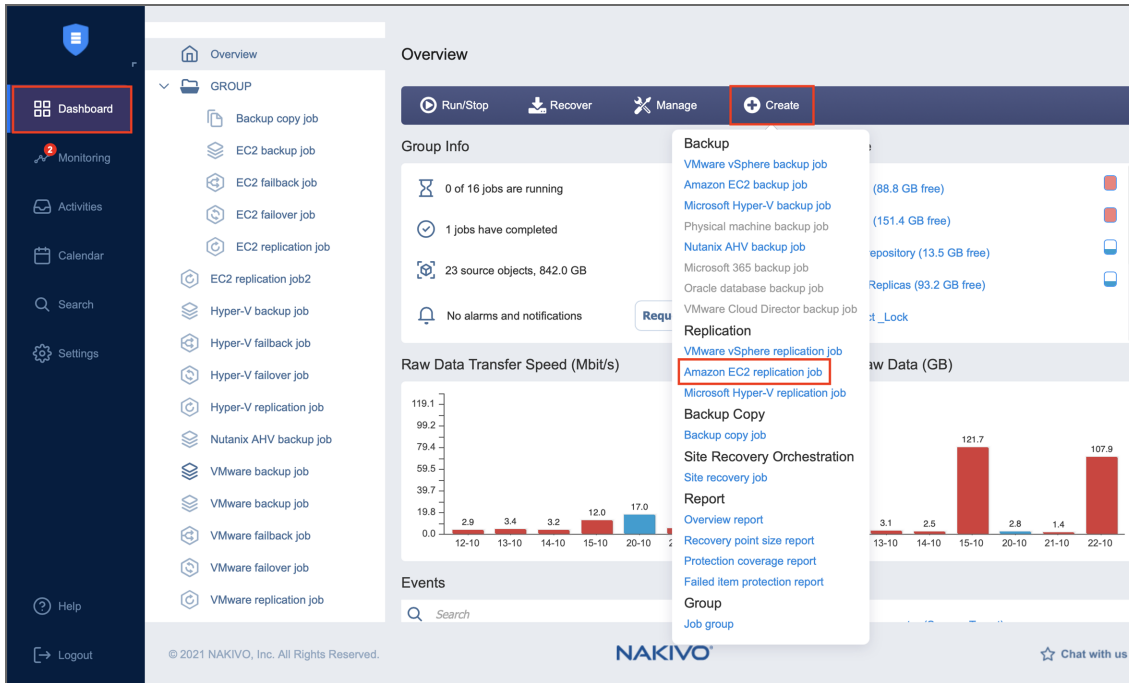
Click **Finish** or **Finish & Run** to complete the job creation.

### 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 256](#) for details.

# Creating Amazon EC2 Replication Jobs

With NAKIVO Backup & Replication, you can replicate Amazon EC2 instances by using the workflow with a broad range of options. To create a replication job, click **Create** on the **Dashboard** and then click **Amazon EC2 replication job**.



The **New Replication Job Wizard for Amazon EC2** opens. Complete the wizard to create a replication job.

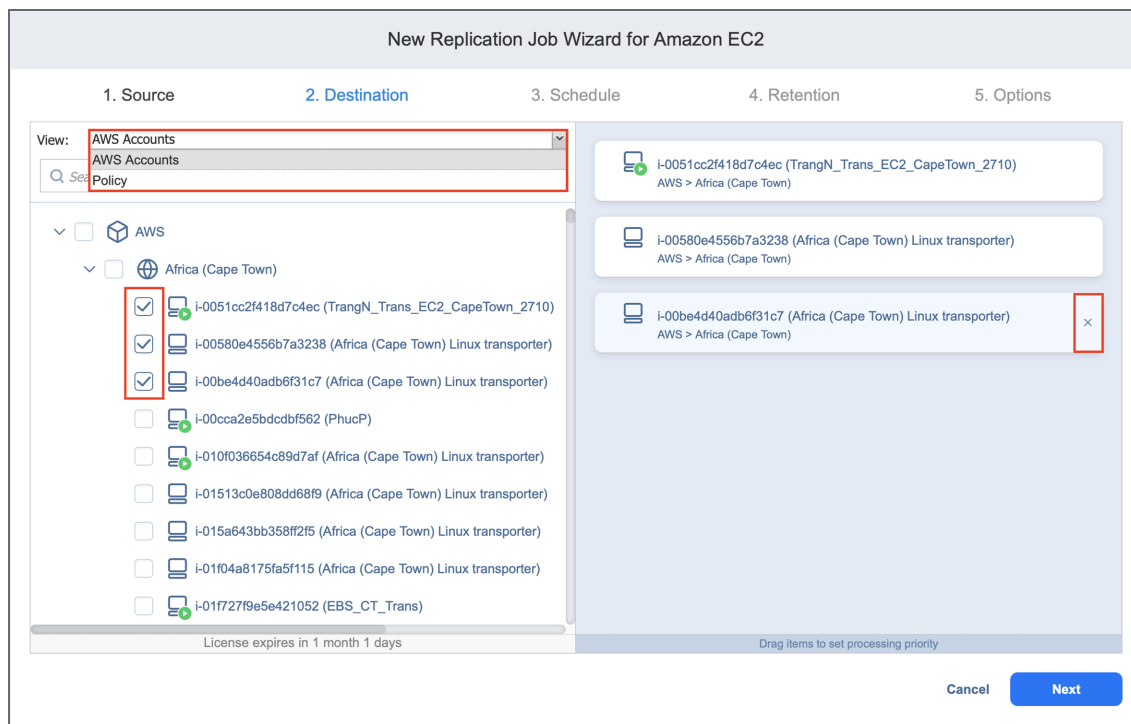
- [“Replication Job Wizard for Amazon EC2: Source” below](#)
- [“Replication Job Wizard for Amazon EC2: Destination” on page 775](#)
- [“Replication Job Wizard for Amazon EC2: Schedule” on page 777](#)
- [“Replication Job Wizard for Amazon EC2: Retention” on page 780](#)
- [“Replication Job Wizard for Amazon EC2: Options” on page 781](#)

## Replication Job Wizard for Amazon EC2: Source

On the **Source** page of the wizard, you can add EC2 items to your replication job. Proceed as follows:

1. In the left pane of the page, select one of the following inventory views:
  - **AWS Accounts:** If chosen, the inventory tree opens in the left pane and shows you all Amazon EC2 accounts along with their Regions and available Instances. Proceed as follows:
    - a. 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.
    - b. Select EC2 items by selecting the checkbox next to each item.

- c. 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 Instances first.
  - d. Review the list of the selected EC2 items. You can remove a selected item or a container from the replication job using one 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 want to remove and click the red “x” on the right. This will deselect the item in the left pane.
  - **Policy:** When selected, job policies can be used. Refer to [“Managing Job Policies” on page 266](#) 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 269](#) for details.
2. Click **Next** to confirm adding selected VMs to the replication job.



The wizard will display the next page.

### Notes

- If you cannot find an Instance, try the following:
  - Make sure that an appropriate AWS account (the one that was used to create the Instance) has been added to the [inventory](#).
  - [Refresh](#) the Inventory.

- After you add a Region or an Account to the job:
  - All Instances currently available in the selected Region/Account will be backed up by the job.
  - All new instances that will be created in (or moved to) the Region/Account in the future will be automatically added to the replication job and backed up.
- To replicate Instances in a region, at least one Transporter must be [installed](#) in that region. If there are no Transporters in a particular region, that region will be disabled.
- The order in which Instances are backed up is important if the Transporter performing the replication cannot process all Instances of the job simultaneously — either because the Transporter is loaded or because the number of disks in the job exceeds the Transporter’s maximum load specified during the Transporter creation.

## Replication Job Wizard for Amazon EC2: Destination

On the **Destination** page, select a target region for your replicas.

- [“Selecting a Target Region” below](#)
- [“Selecting a Different Region for Every Amazon EC2 Instance” below](#)
- [“Excluding an EBS Volume from the Replication Job” on the next page](#)

### Selecting a Target Region

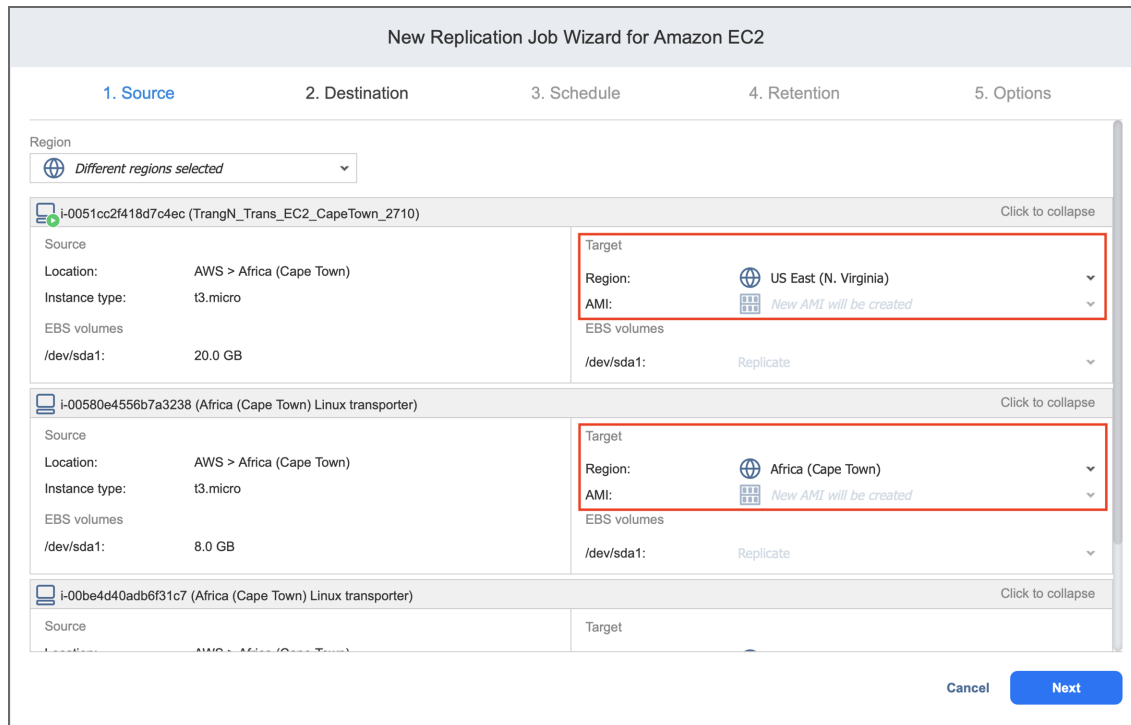
To replicate all Amazon EC2 Instances to the same region, select a target region from the **Region** drop-down list.

The screenshot shows the 'New Replication Job Wizard for Amazon EC2' interface. The 'Destination' step is active, and the 'Region' dropdown menu is open. The dropdown menu is highlighted with a red border. The dropdown menu shows a search bar and a list of regions under the 'AWS' category. The regions listed are: Africa (Cape Town), Asia Pacific (Hong Kong), Asia Pacific (Mumbai), Asia Pacific (Seoul), Asia Pacific (Singapore), Asia Pacific (Sydney), and Asia Pacific (Tokyo). The 'Next' button is visible at the bottom right.

### Selecting a Different Region for Every Amazon EC2 Instance

To replicate each Amazon EC2 Instance to a different region:

1. Click **Advanced setup**.
2. Select a region for each Amazon EC2 Instance.



## Note

You may be charged additionally for using a third-party resource. Please refer to the third-party resource provider documentation for details.

## Excluding an EBS Volume from the Replication Job

If you do not want to replicate some EBS volumes of an Amazon EC2 Instance, you can exclude those EBS volumes from the replication job. Follow the steps below to exclude the EBS volumes:

1. Click **Advanced setup**.
2. Expand the box of an Amazon EC2 Instance.

- From the drop-down list of an EBS Volume, choose the **Do not replicate** option.

New Replication Job Wizard for Amazon EC2

1. Source      2. Destination      3. Schedule      4. Retention      5. Options

Instance type: t3.micro  
EBS volumes  
/dev/sda1: 8.0 GB

Region: Asia Pacific (Hong Kong)  
AMI: New AMI will be created  
EBS volumes  
/dev/sda1: Replicate

i-00be4d40adb6f31c7 (Africa (Cape Town) Linux transporter)      Click to collapse

Source  
Location: AWS > Africa (Cape Town)  
Instance type: t3.micro  
EBS volumes  
/dev/sda1: 8.0 GB

Target  
Region: Asia Pacific (Hong Kong)  
AMI: New AMI will be created  
EBS volumes  
/dev/sda1: Replicate

i-0dc845fc34824bee7 (NA\_Win2016)      Click to collapse

Source  
Location: AWS > EU (Ireland)  
Instance type: t3.medium  
EBS volumes  
/dev/sda1: 30.0 GB  
xvdb: 100.0 GB

Target  
Region: Asia Pacific (Singapore)  
AMI: New AMI will be created  
EBS volumes  
/dev/sda1: Replicate  
xvdb: Do not replicate

Cancel      Next

## Replication Job Wizard for Amazon EC2: 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.

- [“Disabling Scheduled Job Execution” below](#)
- [“Daily or Weekly Replication” on the next page](#)
- [“Monthly or Yearly VM Replication” on the next page](#)
- [“Periodic Replication” on page 779](#)
- [“Chained Job” on page 780](#)
- [“Additional Schedule” on page 780](#)

## Disabling Scheduled Job Execution

If you wish to start the job manually (without any schedule), select the **Do not schedule, run on demand** checkbox.



New Replication Job Wizard for Amazon EC2

1. Source      2. Destination      3. Schedule      4. Retention      5. Options

Do not schedule, run on demand

[Next](#)      [Cancel](#)

## 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.

New Replication Job Wizard for Amazon EC2

1. Source      2. Destination      3. Schedule      4. Retention      5. Options

Do not schedule, run on demand

(UTC+02:00, EET) Eastern European Time

Schedule #1

Run daily/weekly

Starting at: 0:00      Ending: 6:00

Mon     Tue     Wed     Thu     Fri     Sat     Sun

All days    Work days    Weekends

every 1 weeks

Effective from

[Add another schedule](#)

[Show calendar](#)

[Next](#)      [Cancel](#)

## Monthly or Yearly VM 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.

The screenshot shows the 'New Replication Job Wizard for Amazon EC2' at the '3. Schedule' step. The interface includes a progress bar with five steps: 1. Source, 2. Destination, 3. Schedule, 4. Retention, and 5. Options. Below the progress bar, there are several options and settings:

- Do not schedule, run on demand
- (UTC+02:00, EET) Eastern European Time
- Schedule #1: Run monthly/yearly (highlighted with a red box)
- Run every: last Friday of every month
- Starting at: 0:00, Ending: 6:00
- Effective from
- Links: Add another schedule, Show calendar
- Buttons: Next, Cancel

## 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.

The screenshot shows the 'New Replication Job Wizard for Amazon EC2' at the '3. Schedule' step. The interface includes a progress bar with five steps: 1. Source, 2. Destination, 3. Schedule, 4. Retention, and 5. Options. Below the progress bar, there are several options and settings:

- Do not schedule, run on demand
- (UTC+02:00, EET) Eastern European Time
- Schedule #1: Run periodically (highlighted with a red box) every 30 minutes
- Starting at: 0:00, Ending: 6:00
- Mon  Tue  Wed  Thu  Fri  Sat  Sun
- Links: All days, Work days, Weekends
- Effective from
- Links: Add another schedule, Show calendar
- Buttons: Next, Cancel

## 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.

The screenshot shows the 'New Replication Job Wizard for Amazon EC2' interface. The wizard is currently on step 3, 'Schedule'. The previous steps are '1. Source', '2. Destination', '4. Retention', and '5. Options'. The 'Do not schedule, run on demand' checkbox is unchecked. The time zone is set to '(UTC+02:00, EET) Eastern European Time'. Under 'Schedule #1', the 'Run after another job' option is selected and highlighted with a red box. The 'After the job:' dropdown is set to 'Hyper-V backup job'. The 'Run this job:' dropdown is set to 'Immediately'. The 'After successful runs' checkbox is checked, while 'After failed runs' and 'After stopped runs' are unchecked. The 'Effective from' checkbox is also unchecked. There are links for 'Add another schedule' and 'Show calendar'. At the bottom right, there are 'Next' and 'Cancel' buttons.

## Additional 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 Amazon EC2: Retention

NAKIVO Backup & Replication can create a recovery point (snapshot) on the replica Amazon EC2 Instance after each job run. You can specify the number of recovery points to be retained. Use the following options to specify a retention policy:

- **Keep x last recovery points:** Retains the specified number of last recovery points for each Amazon EC2 Instance 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 Amazon EC2: Options

On the **Options** page of the wizard, set up job options as described in the following sections:

- [“Job Options” below](#)
  - [“Job Name” below](#)
  - [“App-aware Mode” below](#)
- [“Pre and Post Actions” on the next page](#)
  - [“Email Notifications” on the next page](#)
  - [“Microsoft Exchange Server Logs Truncation” on page 783](#)
  - [“Truncation of Microsoft SQL Transaction Logs” on page 783](#)
  - [“Running a Pre-Job Script” on page 784](#)
  - [“Running a Post-Job Script” on page 785](#)
- [“Completing the New Replication Job Wizard for Amazon EC2” on page 786](#)

### Job Options

In this section, you can specify a name for your replication job and enable [app-aware mode](#).

#### Job Name

Specify a name for the replication job:

#### App-aware Mode

With NAKIVO Backup & Replication, you can create application-consistent replicas of Amazon EC2 Instances.

To set up application awareness:

1. Select one of the App-aware modes:
  - **Enabled (proceed on error):** With this option selected, the Instance replication continues to run even if application awareness fails.
  - **Enabled (fail on error):** With this option selected, Instance replication fails if application awareness fails.
  - **Disabled:** With this option selected, application awareness is not used for Instance replication.
2. When enabling application-aware mode, specify the OS credentials that will be used to connect to Amazon EC2 Instances via SSH:
  - In the **Application-Aware Mode** dialog box that opens, select the checkbox next to each Instance that you want to replicate using the application-aware mode.
  - Select the appropriate credentials from the drop-down list next to each selected Amazon EC2 Instance.

The screenshot shows the 'New Replication Job Wizard for Amazon EC2' dialog box, specifically the '5. Options' step. The 'App-aware mode' dropdown menu is highlighted with a red box and is currently set to 'Disabled'. Below this, there are sections for 'Pre and Post Actions' with checkboxes for 'Send job run reports to', 'Truncate Exchange logs', 'Truncate SQL Server logs', 'Run local pre job script', and 'Run local post job script'. Each checkbox has a corresponding dropdown menu with a question mark icon. At the bottom right, there are three buttons: 'Finish' (blue), 'Finish & Run' (grey), and 'Cancel' (grey).

## Pre and Post Actions

In the *Pre and Post Actions* section, you can set up email notifications, [Exchange and SQL Server logs truncation](#), and [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.

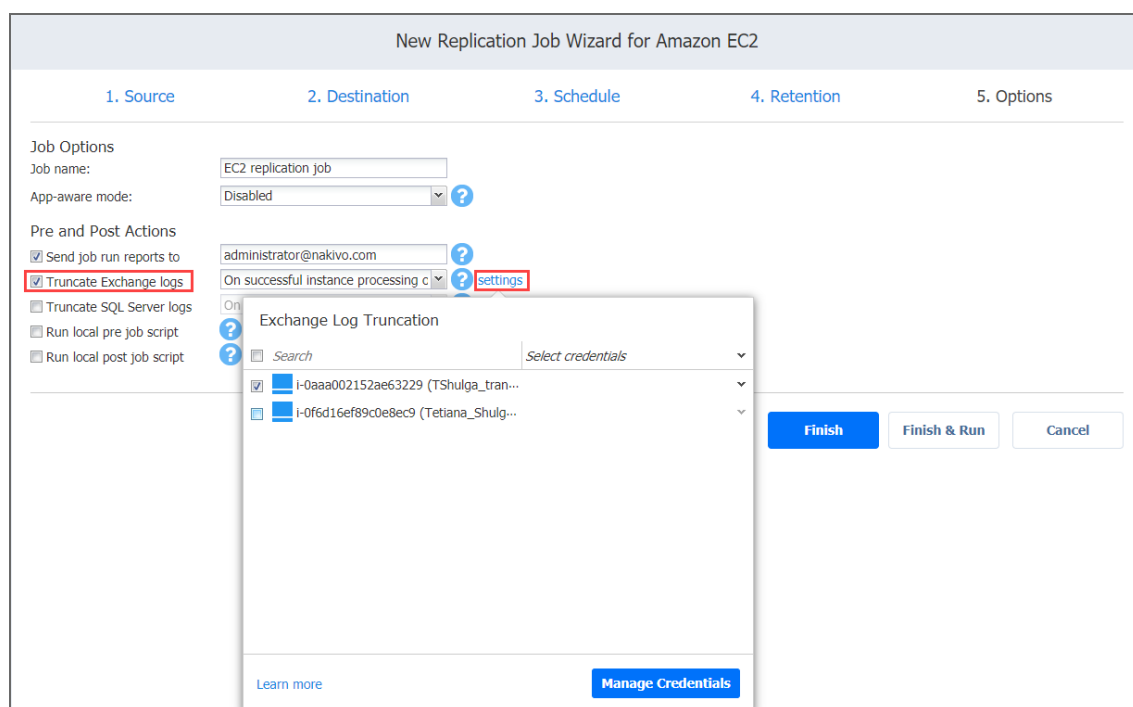
To send email notifications, select the **Send job run reports to** option in the *Pre and Post Actions* section, and specify one or more email addresses in the text field. Use semicolons to separate multiple email addresses.

## Microsoft Exchange Server Logs Truncation

NAKIVO Backup & Replication allows you to delete (aka truncate) Microsoft Exchange Server logs on the source instances after job completion.

To set up Microsoft Exchange Server log truncation, do the following:

1. In the **Pre and Post Actions** section, select **Truncate Exchange Logs**.
2. In the **Exchange Log Truncation** dialog box that opens, select the checkboxes next to the instances running Microsoft Exchange and then select the credentials next to each instance. These credentials are used to log into the instances that you have selected.




## Truncation of Microsoft SQL Transaction Logs

NAKIVO Backup & Replication allows you to delete (aka truncate) Microsoft SQL Server logs on the source instances after job completion.

To set up Microsoft SQL log truncation, do the following:

1. In the **Pre and Post Actions** section, select **Truncate SQL Server logs**.
2. In the **SQL Server Log Truncation** dialog box that opens, select the checkboxes next to the instances running Microsoft SQL and then select the credentials next to each instance. These credentials are used to log into the instances that you have selected.



## Running a Pre-Job Script

To run a script before the product begins replicating instances, do the following:

1. Place a script file on the machine where the Director is installed.
2. Select the **Run local pre 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. The 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, instance replication 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 replicating instances 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 instance replication even if the script has failed.
- **Fail the job on script failure:** With this option selected, the job is failed and instance replication is not performed.

The screenshot shows the 'New Replication Job Wizard for Amazon EC2' at the 'Options' step. The 'Job Options' section includes a 'Job name' field with 'EC2 replication job', a 'Script path' field with 'cmd.exe /c D:\script.bat', a 'Job behavior' dropdown set to 'Wait for the script to finish', and an 'Error handling' dropdown set to 'Fail the job on script failure'. Below these fields, there are two checkboxes: 'Run local pre job script' (checked and highlighted with a red box) and 'Run local post job script' (unchecked). A tooltip is displayed over the 'Run local pre job script' checkbox, showing the script path, job behavior, and error handling settings. At the bottom right, there are three buttons: 'Finish', 'Finish & Run', and 'Cancel'.

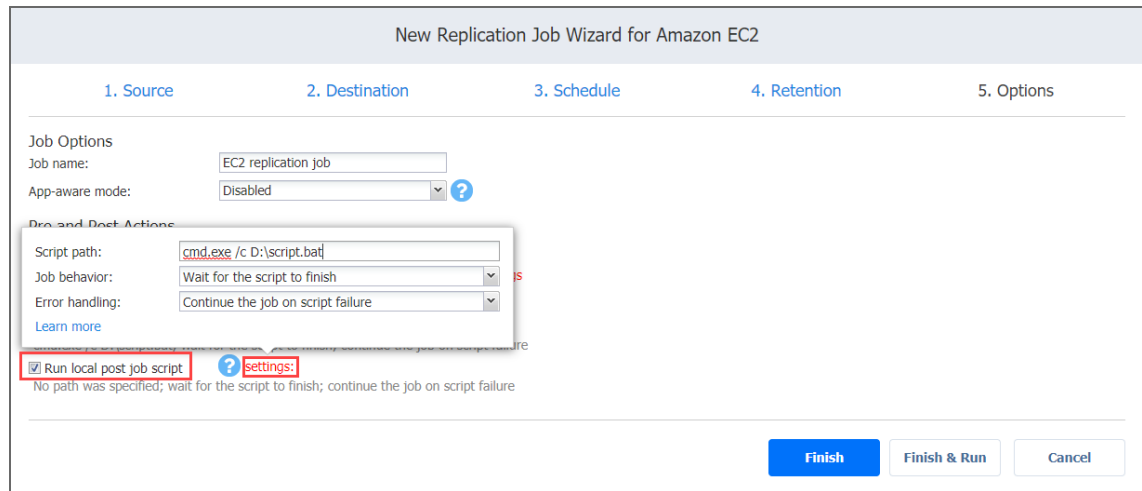
## Running a Post-Job Script

To run a script after the product has finished backing up all instances, do the following:

1. Place a script file on the machine where the Director is installed.
2. 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. The 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 has failed, the job status is set to “failed” even if instance replication has been successful.



## Completing the New Replication Job Wizard for Amazon EC2

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 256](#) for details.

## 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:

- [“Full Recovery” on page 847](#)
- [“Granular Recovery” on the next page](#)
- [“Planning Disaster Recovery” on page 971](#)

# 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” below](#)
- [“File Share Recovery” on page 801](#)
- [“Object Recovery for Microsoft Exchange” on page 807](#)
- [“Object Recovery for Microsoft Active Directory” on page 814](#)
- [“Importing Recovered Objects to Active Directory” on page 823](#)
- [“Object Recovery for Microsoft SQL Server” on page 824](#)
- [“Object Recovery for Microsoft 365” on page 831](#)
- [“Object Recovery for Oracle Database” on page 839](#)
- [“Performing Universal Object Recovery” on page 843](#)

## 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 24](#) for more information.

### Note

File recovery is restricted to [supported disk types and file systems](#).

Refer to the following topics to learn more:

- [“Opening File Recovery Wizard” on the next page](#)
- [“File Recovery Wizard: Backup” on page 790](#)
- [“File Recovery Wizard: Recovery Server” on page 791](#)
- [“File Recovery Wizard: Files” on page 794](#)
- [“File Recovery Wizard: Options” on page 797](#)

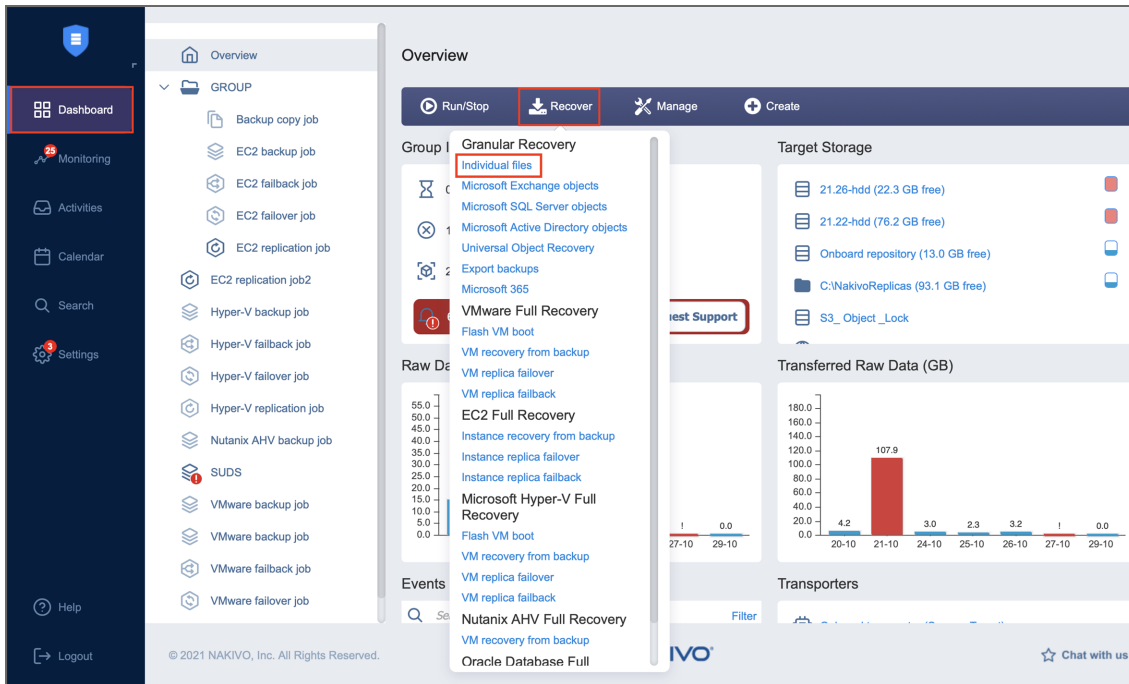
## Opening File Recovery Wizard

You can start the recovery process either from the **Dashboard**, 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 Dashboard](#)
- [Starting File Recovery from Backup Repository](#)

## Starting File Recovery from Dashboard

To start file recovery from the **Dashboard**, click **Recover** and then click **Individual Files**.

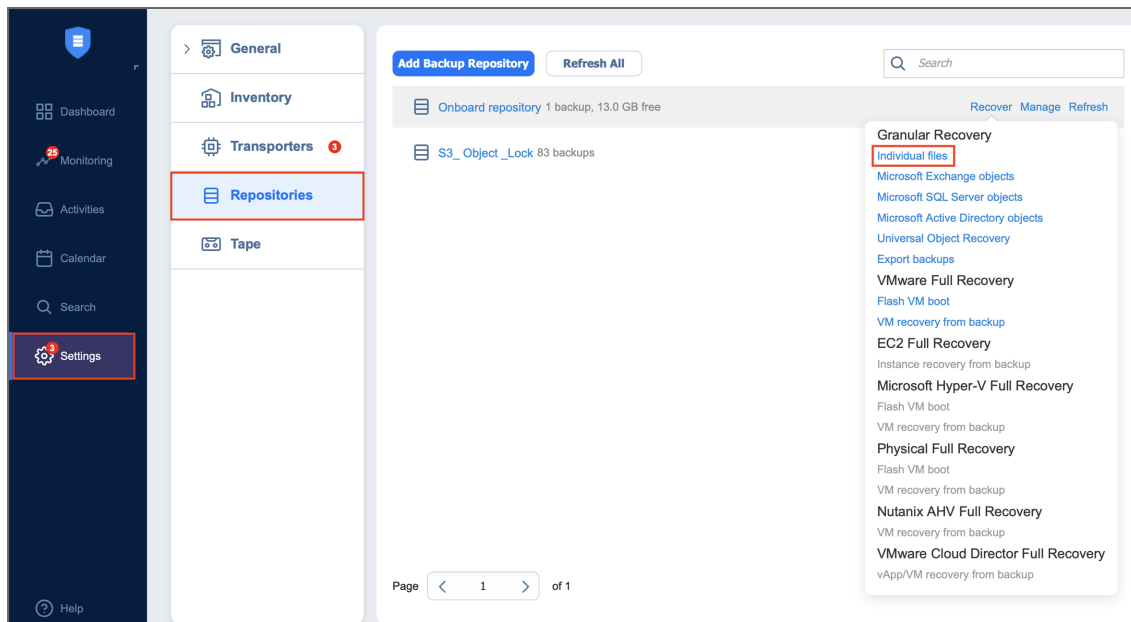


## Starting File Recovery from Backup Repository

To start file recovery from a Backup Repository, do the following:

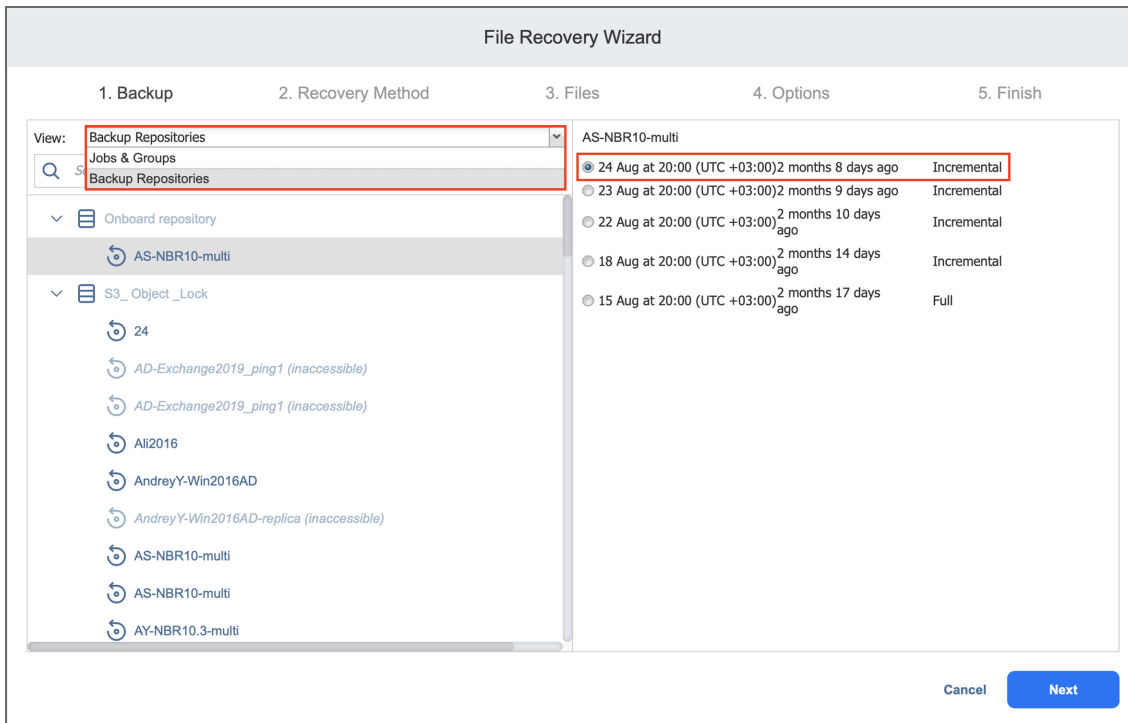
1. From the main menu of NAKIVO Backup & Replication, click **Settings**.
2. Go to the **Repositories** tab and hover over the Backup Repository containing the required backup.

3. Click the **Recover** button and then click **Individual Files**. The **File Recovery Wizard** opens.



## File Recovery Wizard: Backup

On the **Backup** page of the wizard, select a backup using either a **Backup Repository** or **Jobs & Groups** view in the left pane, and then select a recovery point in the right pane.



Click **Next** to go to the next page of the wizard.

## File Recovery Wizard: Recovery Server

In this page of the wizard, choose one of the following recovery methods:

- [Recovering Files to Server](#)
- [Downloading Files to Browser or Sending Files via Email](#)
- [Recovering Files to File Share](#)

### Recovering Files to Server

Please do the following to recover files to a server:

1. In the **Recovery Method** section, choose **Recover to the following server**.
2. The **Recovery Server Settings** section opens. 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 IP address:** 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.
3. 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.

4. Click **Next**.

File Recovery Wizard

1. Backup      2. Recovery Method      3. Files      4. Options      5. Finish

Recovery Method

Download to browser or send via email

Recover to the following server

Recover to file share

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: ●●●●●●●●

[Manage credentials](#)

[Test Connection](#)

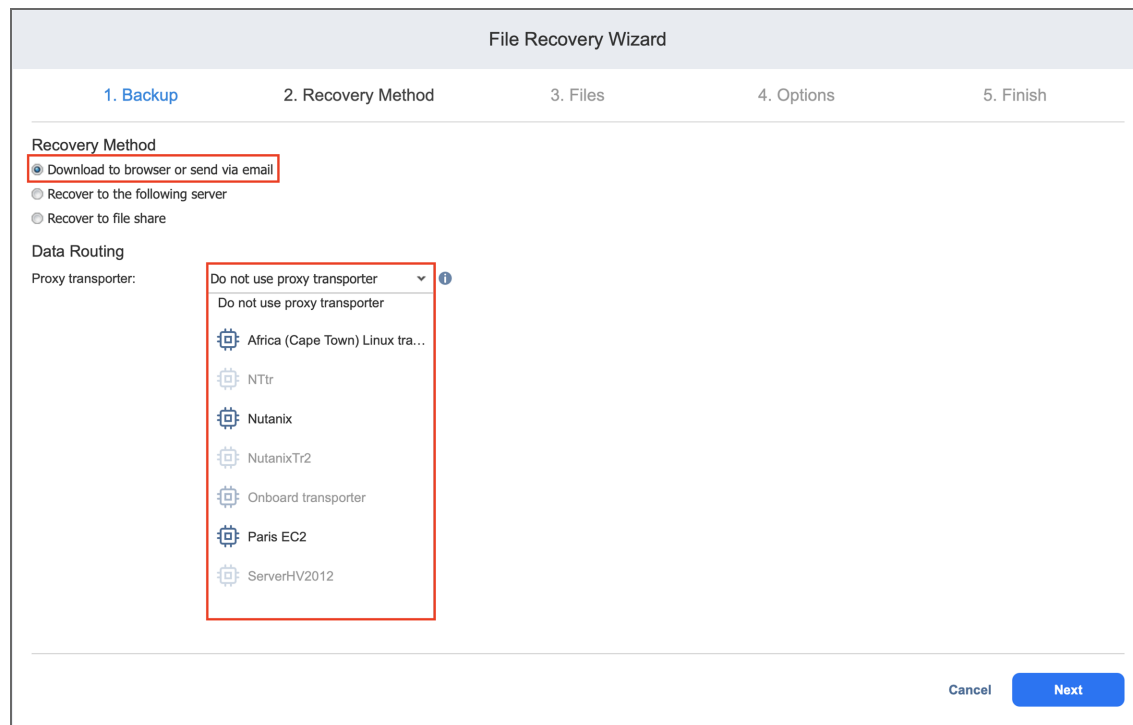
[Cancel](#) [Next](#)

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**.



## Note

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 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**.

File Recovery Wizard

1. Backup      2. Recovery Method      3. Files      4. Finish

Recovery Method

- Download to browser or send via email
- Recover to the following server
- Recover to file share

Data Routing

Proxy transporter: Do not use proxy transporter

File Share Settings

Share type: CIFS

Path to the share: \\Backup\Files

Credentials type: Password

Username: admin

Password: .....

Manage credentials

Test Connection

Cancel      Next

## Note

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 localhost 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.

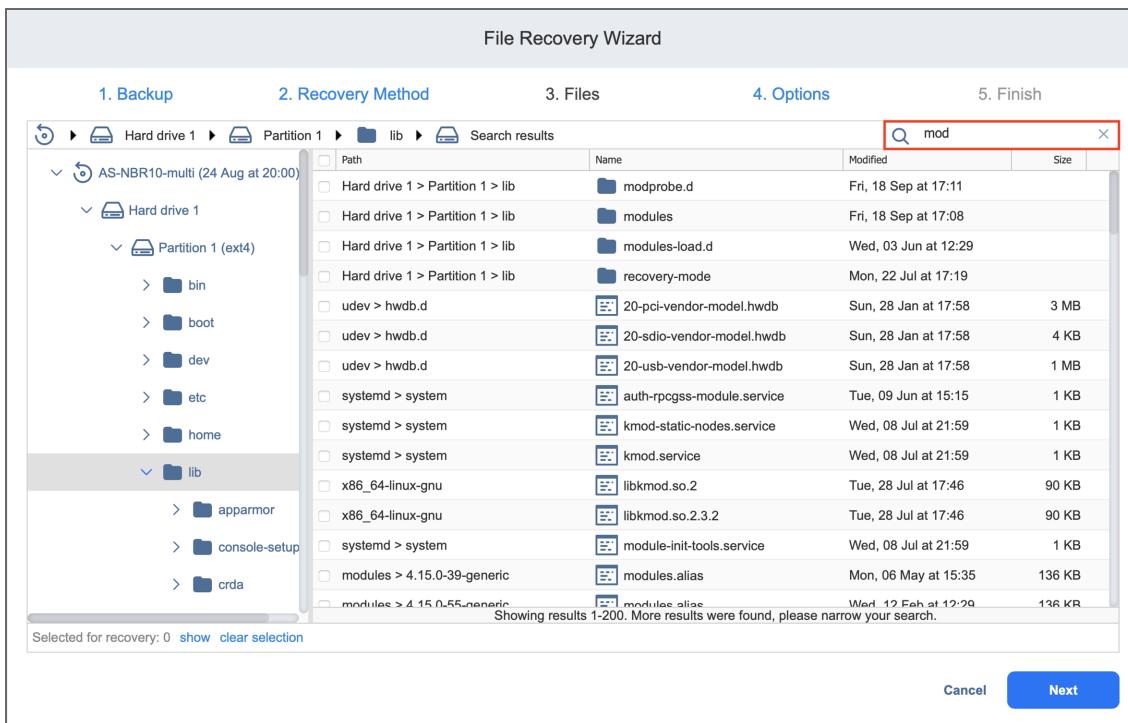
## File Recovery Wizard: Files

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**.

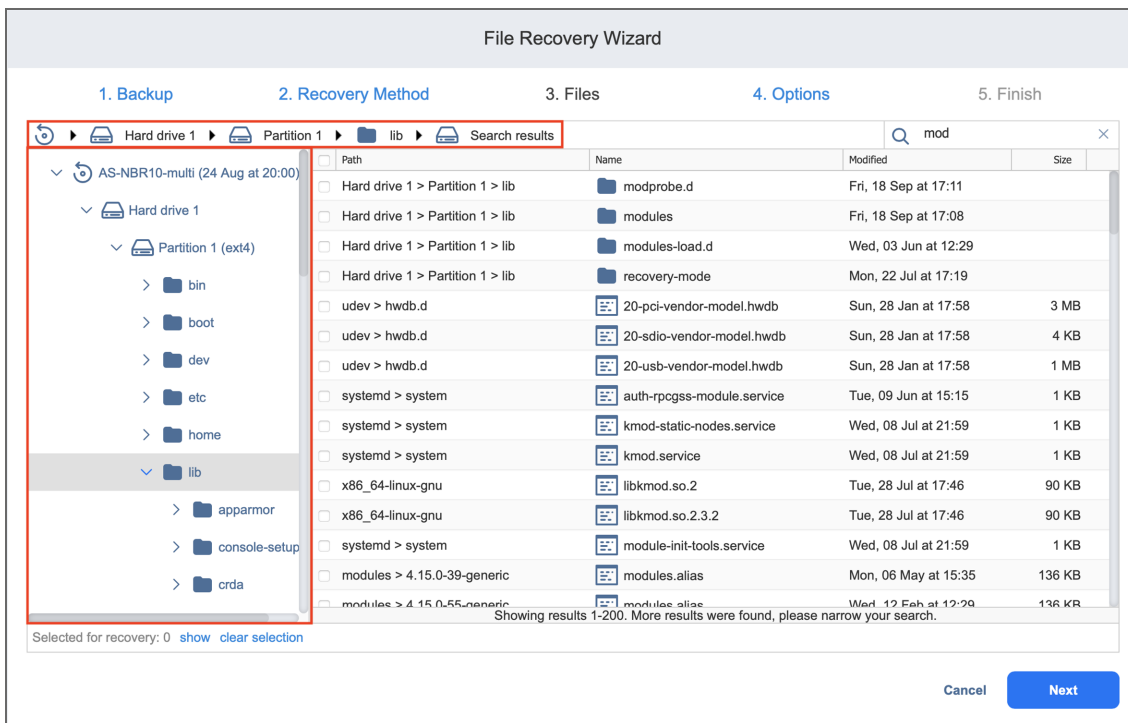


## Notes

- 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:



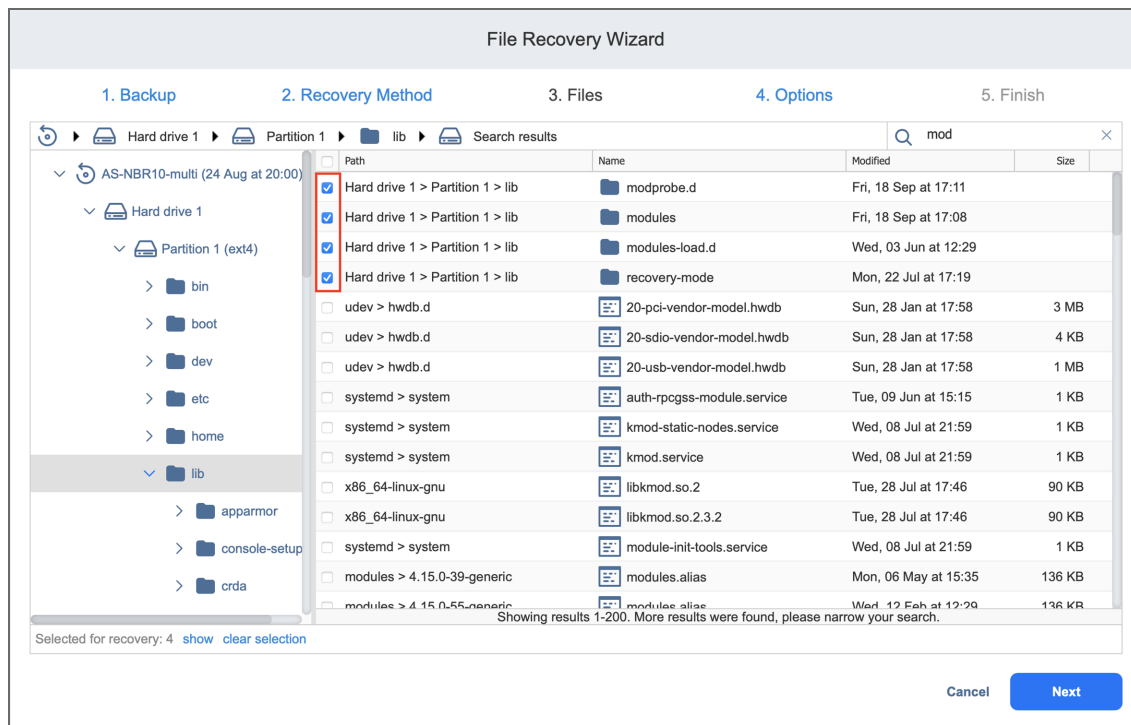
If a VM backup contains Linux LVM volumes or Windows dynamic disks, the navigation pane will display these logical groups in addition to all hard drives available in the VM backup. If a hard drive does not contain any partitions and servers as a part of a Linux LVM volume or a Windows dynamic disk, this 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 Server** page of the wizard, proceed as follows.

#### Warning

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.

### Recovering Files to the Original Location

To recover files to original location:

1. In the **Recovery type** list, choose **Recover to original location**.
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.

The screenshot shows the 'File Recovery Wizard' dialog box. At the top, there are five steps: 1. Backup, 2. Recovery Server, 3. Files (current step), 4. Options, and 5. Finish. Below the steps, there are two dropdown menus. The first is labeled 'Recovery type:' and is set to 'Recover to original location'. The second is labeled 'Overwrite behavior:' and is open, showing three options: 'Rename recovered item if such item exists', 'Skip recovered item if such item exists', and 'Overwrite the original item if such item exists'. At the bottom right, there are two buttons: 'Recover' (highlighted in blue) and 'Cancel'.

## 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 one of the following:
    - **Local folder on Recovery Server**
    - **CIFS share**
    - **NFS share**

### Note

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. In the **Location path/Path to share box**, enter the path to be used for file recovery:
    - A local path if you choose the **Local folder on Recovery Server** option.
    - A path to share on a remote server if you choose **CIFS share/NFS share**.
  - 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**.

## Note

File-level recovery to CIFS share may fail if network credentials are shared across multiple open CIFS connections. If you encounter issues while performing file recovery to CIFS share, try adding localhost domain to the used credentials (i.e. localhost\Administrator).

File Recovery Wizard

1. Backup    2. Recovery Server    3. Files    4. Options    5. Finish

Recovery type: Recover to custom location

Location type: CIFS share

Path to the share: \\zenlar012\Share

Username: admin

Password: \*\*\*\*\*

Overwrite behavior: Rename recovered item if such item exists

Manage credentials

Recover    Cancel

## 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 Server** 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**.


File Recovery Wizard

1. Backup    2. Recovery Server    3. Files    4. Options    5. Finish

Recovery type: Download

Recover    Cancel

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 [“Notifications & Reports” on page 294](#) 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**.

The screenshot shows the 'File Recovery Wizard' interface. At the top, there are five steps: 1. Backup, 2. Recovery Server, 3. Files, 4. Options, and 5. Finish. The 'Options' step is currently selected. Below the steps, there are several input fields: 'Recovery type:' with a dropdown menu set to 'Forward via email' (highlighted with a red box), 'To:' with the email 'administrator@nakivo.com', 'CC:' with 'administrator@example.com', and 'Subject:' with 'Recovered items - 17 Jun at 12:05 (UTC +03:00)'. Below these fields, there is an 'Attachments:' section with a blue square icon and the text 'Administrator'. A large text area below contains the text 'Recovered by NAKIVO Backup & Replication'. At the bottom right, there are two buttons: 'Recover' (blue) and 'Cancel' (grey).

## File Share Recovery

With NAKIVO Backup & Replication, you can recover NFS/CIFS shares from backups in a single ZIP archive. For details, refer to the following topics:

- [“Starting File Share Recovery” below](#)
- [“Recovery Job Wizard for File Share: Backup” on page 803](#)
- [“Recovery Job Wizard for File Share: Files” on page 803](#)
- [“Recovery Job Wizard for File Share: Options” on page 805](#)

## Starting File Share Recovery

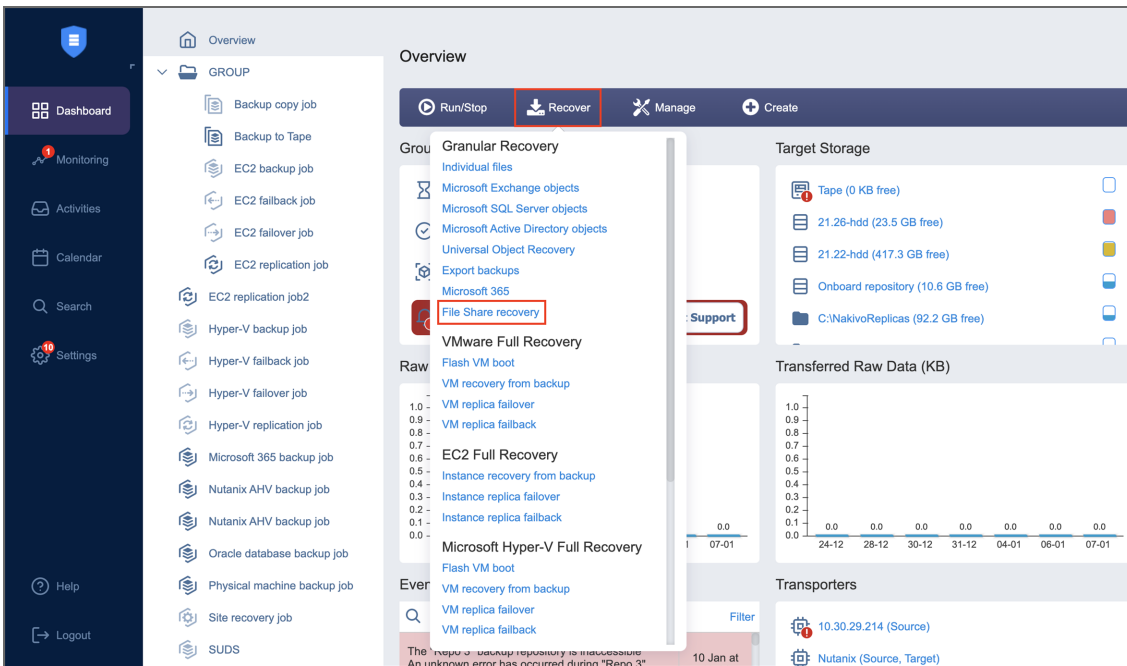
You can start the recovery process from the **Dashboard** or from the **Repositories** tab in **Settings** (for example, if you no longer have a backup job, but still have the backup). Refer to the following sections for more details:

- [Starting File Share Recovery from Dashboard](#)
- [Starting File Share Recovery from Backup Repository](#)

## Starting File Share Recovery from Dashboard

To start file share recovery from the **Dashboard**, click **Recover** and then click **File share recovery**.

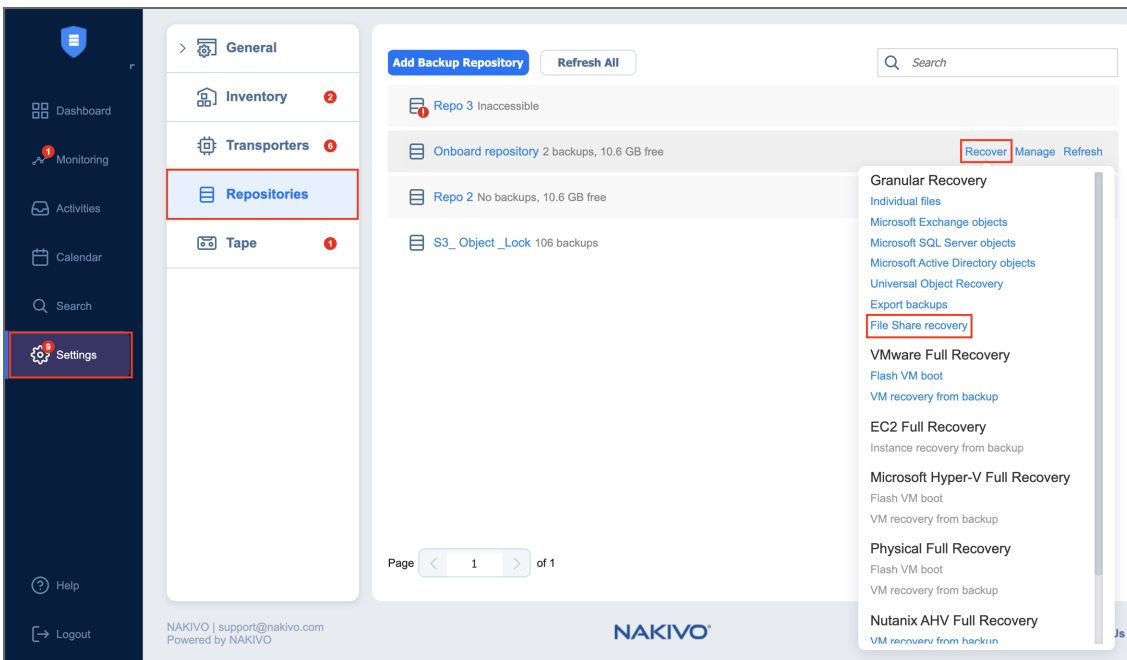




## Starting File Share Recovery from Backup Repository

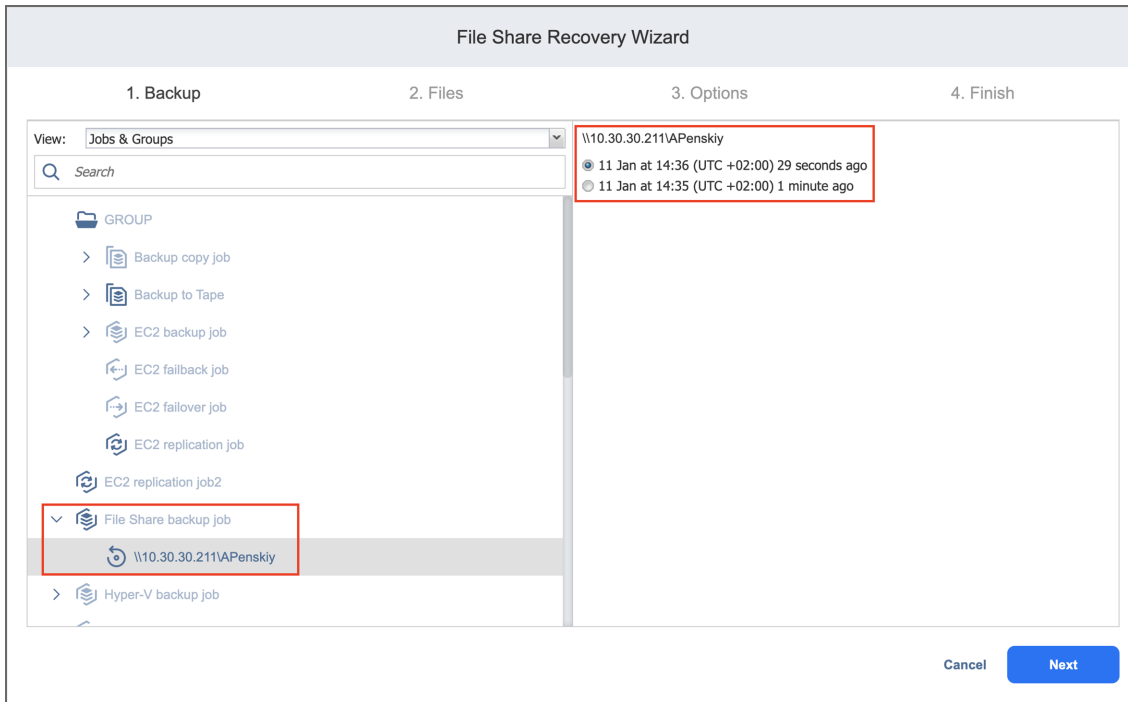
To start file recovery from a Backup Repository, do the following:

1. In the left pane, click **Settings**.
2. Go to the **Repositories** tab and hover over the backup repository containing the backup you need.
3. Click the **Recover** button and then click **File Share recovery**. The **File Share Recovery Wizard** opens.



## Recovery Job Wizard for File Share: Backup

On the **Backup** page of the wizard, select a backed-up file share using the **Backup Repository** view or the **Jobs & Groups** view in the left pane, and then select a recovery point in the right pane.



Click **Next** to go to the next page of the wizard.


## Recovery Job Wizard for File Share: 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](#)

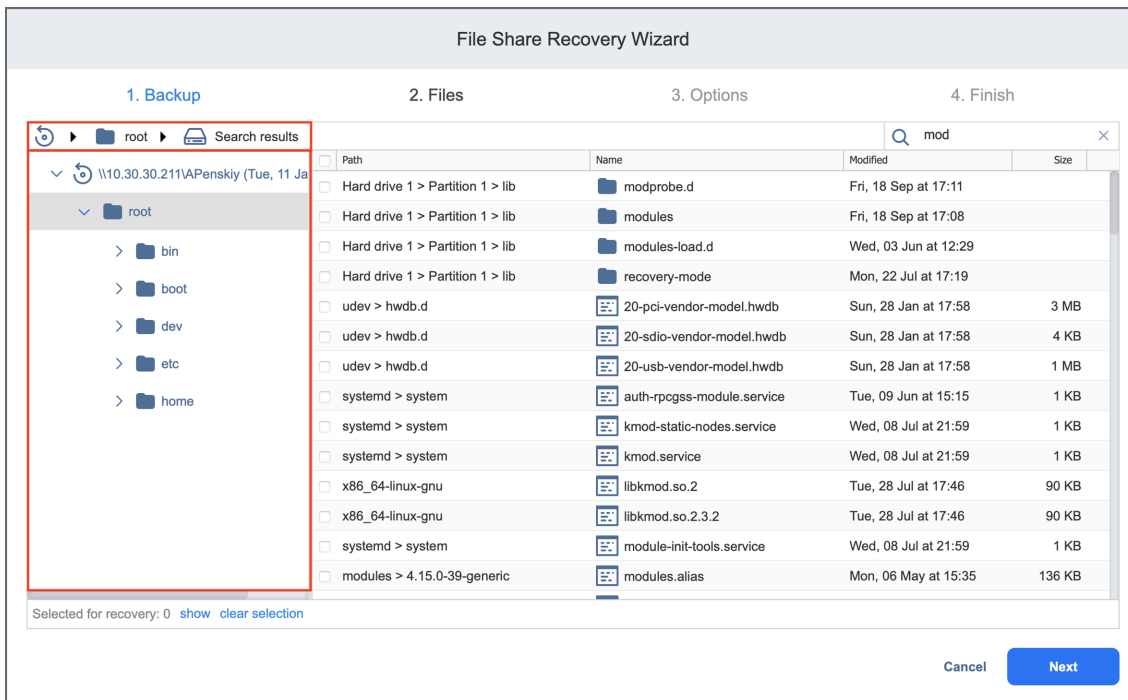
### Searching for Files and Folders

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**.



## Browsing Files and Folders

You can browse the files and folders within a backup using 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. 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.



## Recovery Job Wizard for File Share: Options

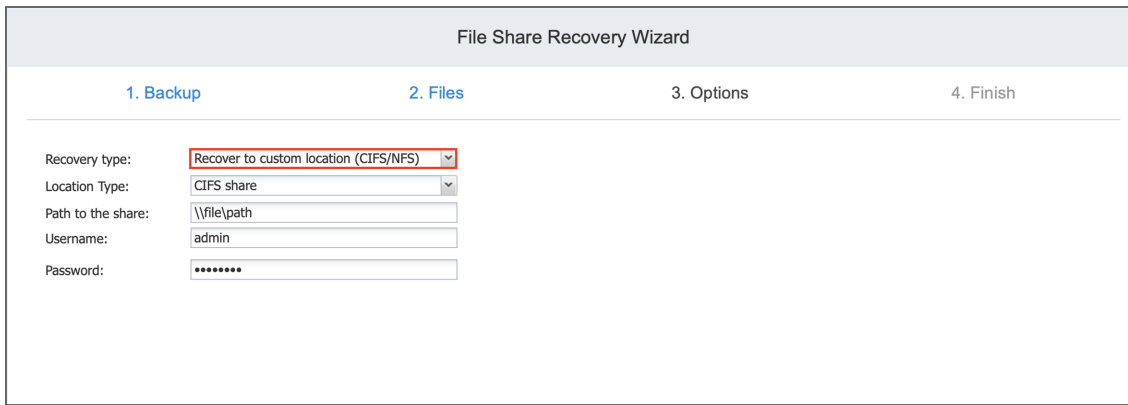
On the Options page, select a recovery and location type. Proceed as described in these sections:

- [Recovering to a Custom Location](#)
- [Downloading to Browser](#)
- [Forwarding via Email](#)

### Recovering to a Custom Location

To recover file shares to a custom location, do the following:

1. From the **Recovery type** drop-down list, select **Recover to custom location**.
2. Select a location type from the **Location type** drop-down list:
  - a. **CIFS share**
  - b. **NFS share**
3. Enter a path to the share.
4. If a **CIFS share** location type is selected, enter the username and password to the share.
5. Click **Recover**.

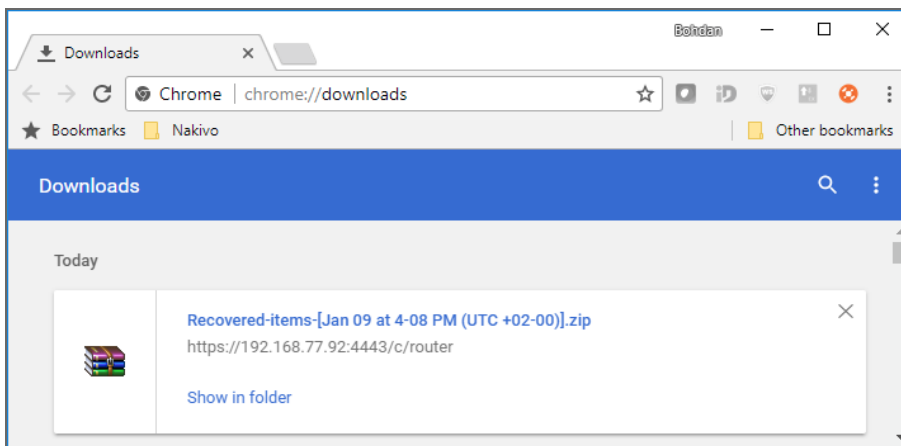


## Downloading to Browser

To download file shares for recovery to browser, do the following:

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 via Email

To use this recovery type, your Email settings must be properly configured in the NAKIVO Backup & Replication Configuration. Refer to [“Notifications & Reports” on page 294](#) for details.

To forward file shares for recovery via email, do the following:

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 the 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, enter one or more email addresses of secondary recipients.
3. Optionally, enter a subject in the **Subject** box.
4. Click **Recover**.

File Share Recovery Wizard

1. Backup
2. Files
3. Options
4. Finish

Recovery type: Forward via email

To:

CC:

Subject:

Attachments: root

Recovered by NAKIVO Backup & Replication

Cancel Recover

## 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” below](#)
- [“Object Recovery Wizard for Microsoft Exchange: Backup” on page 809](#)
- [“Object Recovery Wizard for Microsoft Exchange: Recovery Method” on page 810](#)
- [“Object Recovery Wizard for Microsoft Exchange: Objects” on page 811](#)
- [“Object Recovery Wizard for Microsoft Exchange: Options” on page 812](#)

## Starting Object Recovery for Microsoft Exchange

You can start the recovery process either from the **Dashboard**, 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

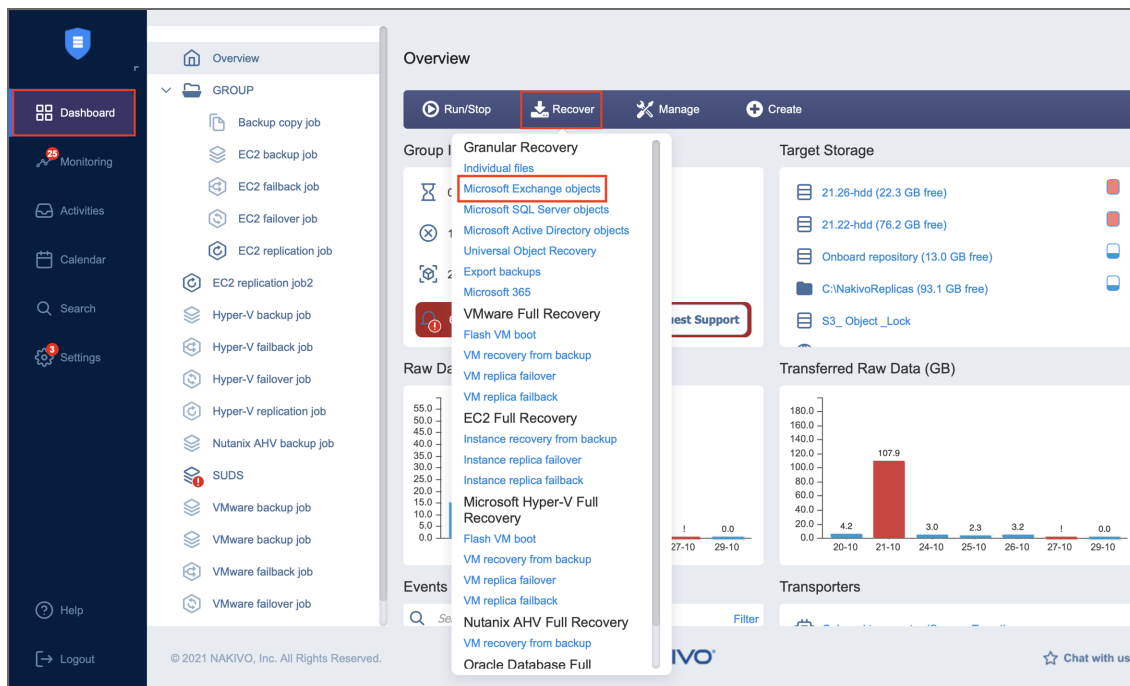
The recovery process may result in additional load and memory usage on the target server. Therefore, make sure that the server has enough memory.

Refer to the following sections to learn how to start the object recovery process for Microsoft Exchange:

- Starting Object Recovery for Microsoft Exchange from Dashboard
- Starting Object Recovery for Microsoft Exchange from Backup Repository

## Starting Object Recovery for Microsoft Exchange from Dashboard

To start object recovery for Microsoft Exchange from the **Dashboard**, click **Recover** and then click **Microsoft Exchange Objects**.




## 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.

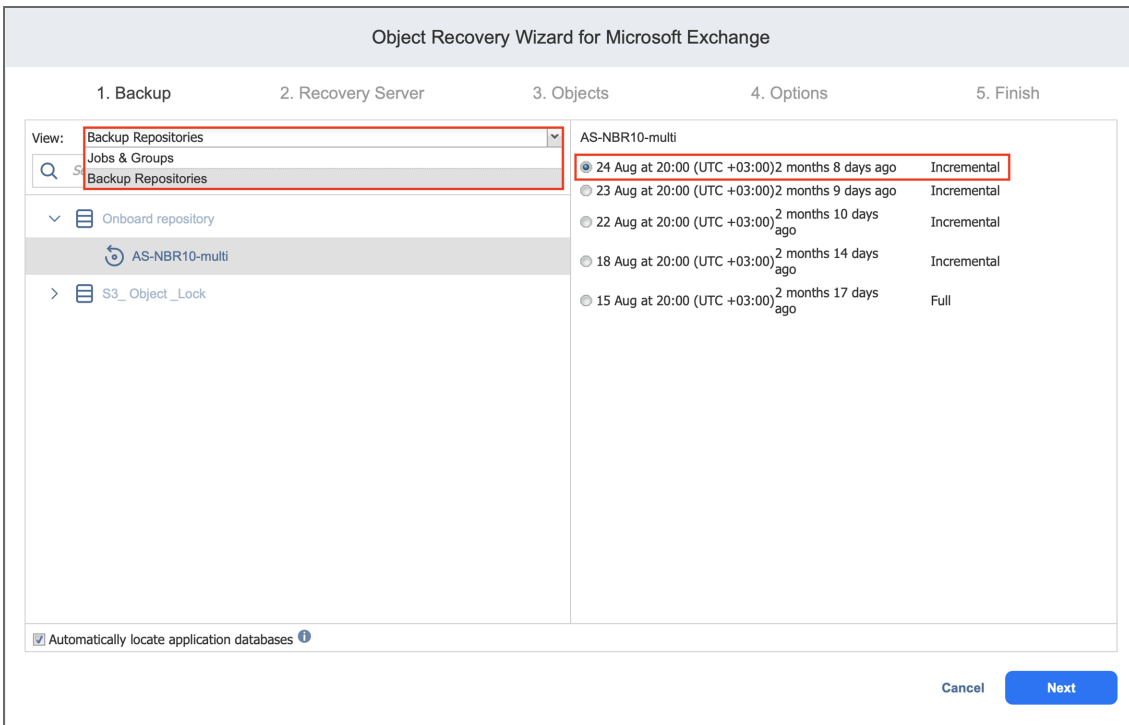
3. Click the **Recover** button and then click **Microsoft Exchange Objects**.



The **Object Recovery Wizard for Microsoft Exchange** opens.

## Object Recovery Wizard for Microsoft Exchange: Backup

On the **Backups** page of the wizard, select a backup using either a **Backup Repository** or **Jobs & Groups** view in the left pane, and then select a recovery point in the right pane.





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.

Click **Next** to go to the next page of the wizard.

## 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.
- **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.

Object Recovery Wizard for Microsoft Exchange

1. Backup      2. Recovery Server      3. Objects      4. Options      5. Finish

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: ●●●●●●

[Manage credentials](#)

[Test Connection](#)

Cancel   [Next](#)

### 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 “.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.

## 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 the objects to their original location on the recovery VM.

Object Recovery Wizard for Microsoft Exchange

1. Backup      2. Recovery Method      3. Objects      4. Options      5. Finish

Recovery type:

Overwrite behavior:

**⚠** The database(s) will be recovered to the selected server, however the mailbox(es) will be disabled.

## Recovering to a Custom Location

In the **Recovery type** drop-down list, select **Recover to custom location** to recover the objects to a custom location on the VM. You can select the recovery location by specifying it in the **Local path field** or by browsing to it.

Object Recovery Wizard for Microsoft Exchange

1. Backup      2. Recovery Method      3. Objects      4. Options      5. Finish

Recovery type:

Local path:

Overwrite behavior:

**⚠** The database(s) will be recovered to the selected server, however the mailbox(es) will be disabled.

## 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.

## 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**

The screenshot shows the 'Object Recovery Wizard for Microsoft Exchange' at step 2, 'Recovery Method'. The wizard has five steps: 1. Backup, 2. Recovery Method, 3. Objects, 4. Options, and 5. Finish. The 'Recovery type' is set to 'Export'. The 'Location type' is 'CIFS share'. The 'Path to the share' is '\\ServerName\FolderName'. The 'Username' is 'Type or select username'. The 'Password' field is empty. The 'Overwrite behavior' dropdown is open, showing three options: 'Rename recovered item if such item exists', 'Skip recovered item if such item exists', and 'Overwrite the original item if such item exists'. The 'Recover' button is highlighted in blue, and the 'Cancel' button is greyed out.

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 .Idif 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” below](#)
- [“Object Recovery Wizard for Microsoft AD Server: Backup” on page 816](#)
- [“Object Recovery Wizard for Microsoft AD Server: Recovery Server” on page 817](#)
- [“Object Recovery Wizard for Microsoft AD Server: Objects” on page 818](#)
- [“Object Recovery Wizard for Microsoft AD Server: Options” on page 821](#)

## Starting Object Recovery for Microsoft Active Directory

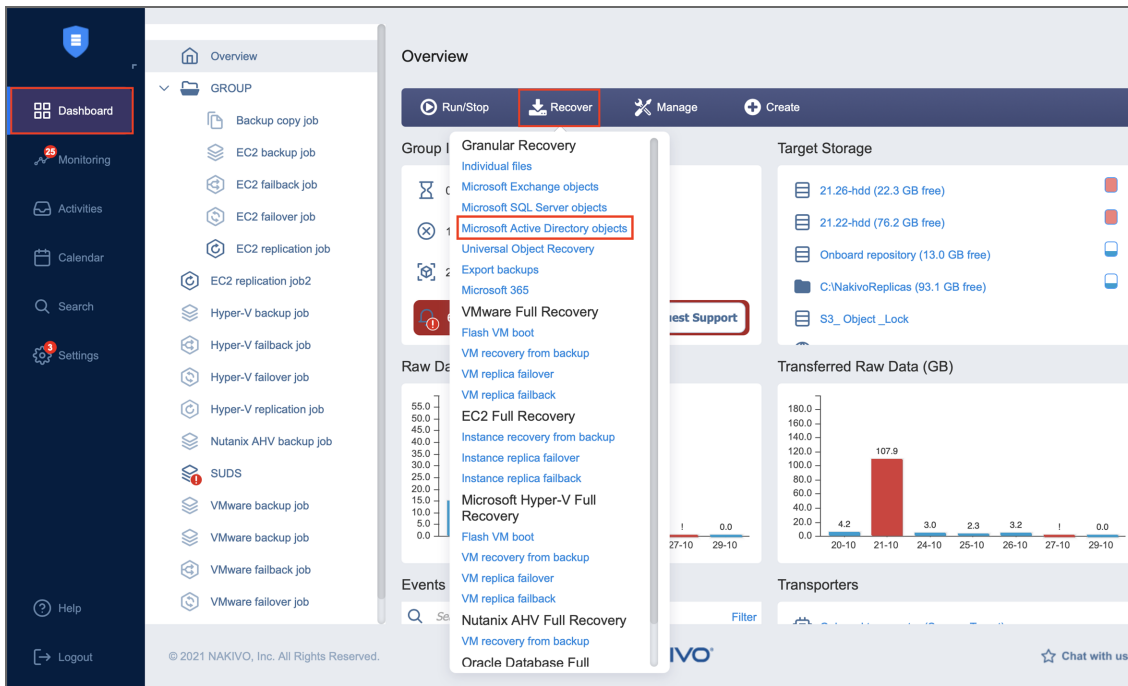
You can start the recovery process either from the **Dashboard**, by using the [search](#) function, or from the **Repositories** tab in the **Settings** (for example, if you no longer have a backup job, but still have the backup).

Refer to the following sections for more details:

- Starting Active Directory Object Recovery from Dashboard
- Starting Active Directory Object Recovery from a Backup Repository

## Starting Active Directory Object Recovery from Dashboard

To start Active Directory Object Recovery from the **Dashboard**, click **Recover** and then choose **Microsoft Active Directory objects**.

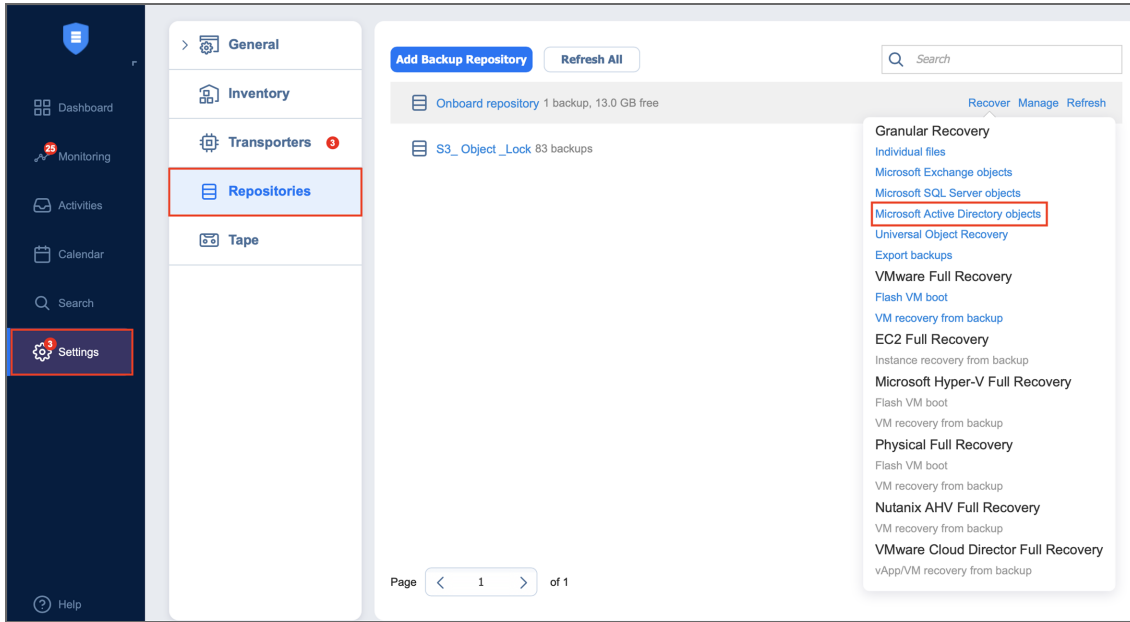


## Starting Active Directory Object Recovery from a Backup Repository

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.

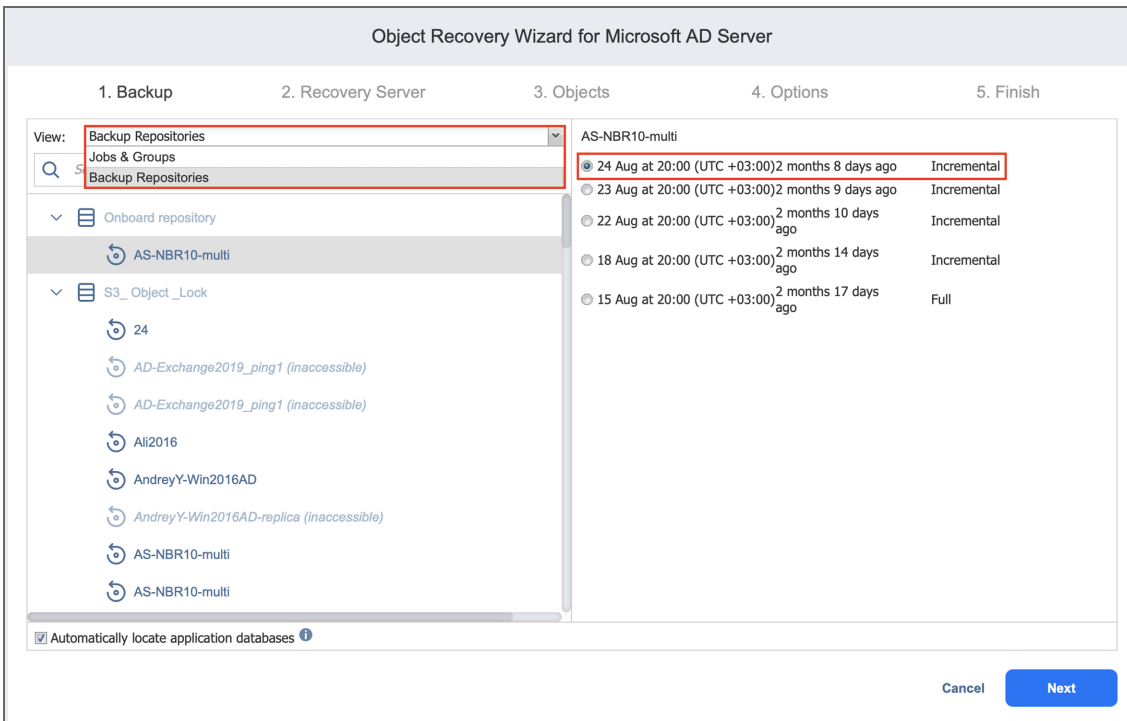
3. Click the **Recover** button and then click **Microsoft Active Directory objects**.



The **Object Recovery Job Wizard for Microsoft AD Server** opens.

## Object Recovery Wizard for Microsoft AD Server: Backup

On the **Backup** page of the wizard, select a backup of a VM with the Microsoft Active Directory server in the left pane and then select a recovery point in the right pane. You can choose a backup from either a Backup Repository or jobs and groups you've created.



By default, NAKIVO Backup & Replication automatically searches the selected recovery point for the 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** checkbox.

Click **Next** to go to the next page of the wizard.

## 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.

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

1. Backup      2. Recovery Server      3. Objects      4. Options      5. Finish

Recovery Server Settings

Recovery server: AS-NBR10-multi ⓘ

Server hostname or IP: 10.30.23.176 ⓘ

Credentials type: Password

Username: admin

Password: ●●●●●●●●

[Manage credentials](#)

**Test Connection**

Create snapshot before recovery ⓘ

Cancel      **Next**

## 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**

Object Recovery Wizard for Microsoft AD Server

1. Backup      2. Recovery Server      3. Objects      4. Options      5. Finish

ntds.dit

win2016+exchange2016 (17 Jun at 15:32)

ntds.dit

WIN-VDIU2CN63IL.exchange.int

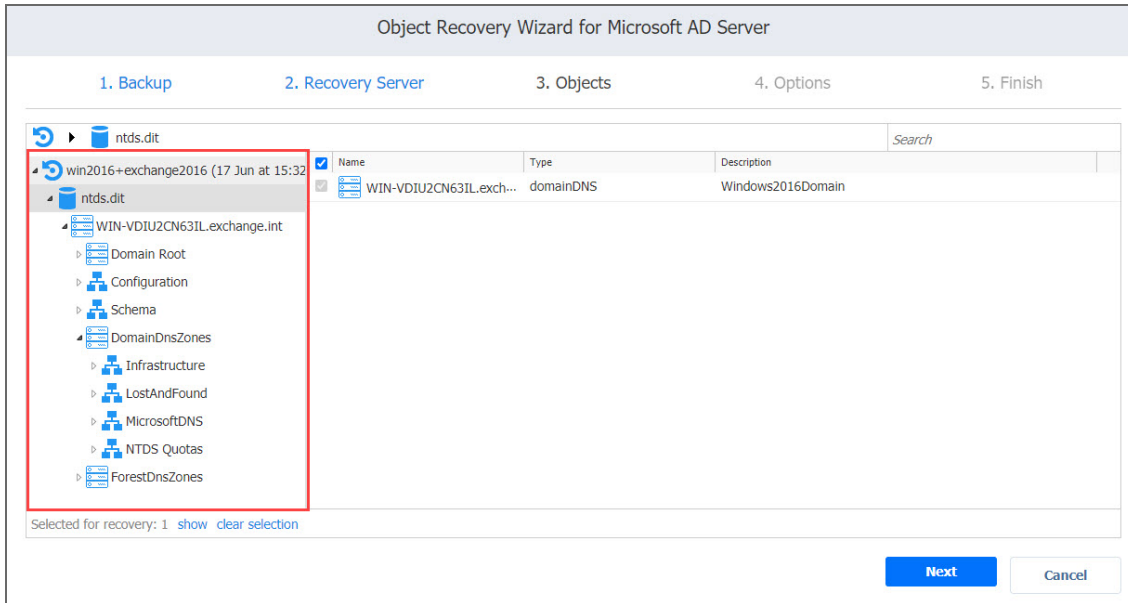
Selected for recovery: 1 [show](#) [clear selection](#)

**Next**      Cancel

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** page put a checkmark 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 Directory 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, the VM 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, go to the Activities page.

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 the next page](#)
- [“Object Recovery Wizard for Microsoft SQL Server: Backup” on page 826](#)
- [“Object Recovery Wizard for Microsoft SQL Server: Recovery Server” on page 827](#)
- [“Object Recovery Wizard for Microsoft SQL Server: Objects” on page 828](#)
- [“Object Recovery Wizard for Microsoft SQL Server: Options” on page 829](#)

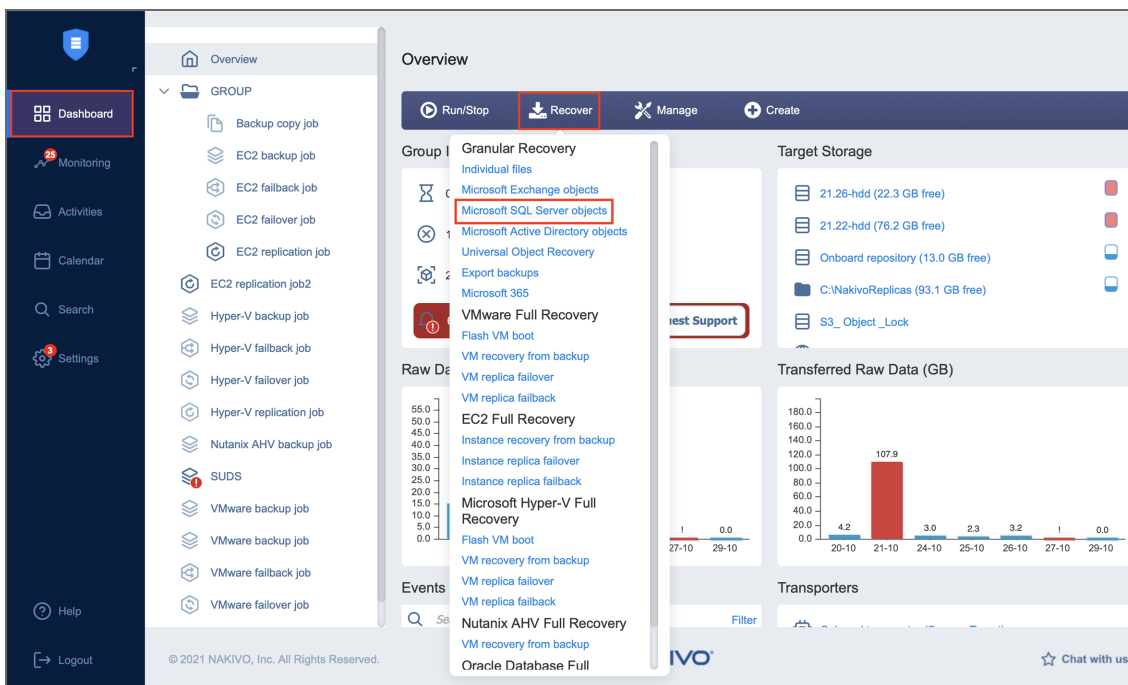
## Starting Object Recovery for Microsoft SQL Server

You can start the recovery process either from the **Dashboard**, 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 Dashboard](#)
- [Starting SQL Server Object Recovery from a Backup Repository](#)

## Starting SQL Server Object Recovery from Dashboard

To start Microsoft SQL Server object recovery from the **Dashboard**, click **Recover** and choose **Microsoft SQL Server objects**.



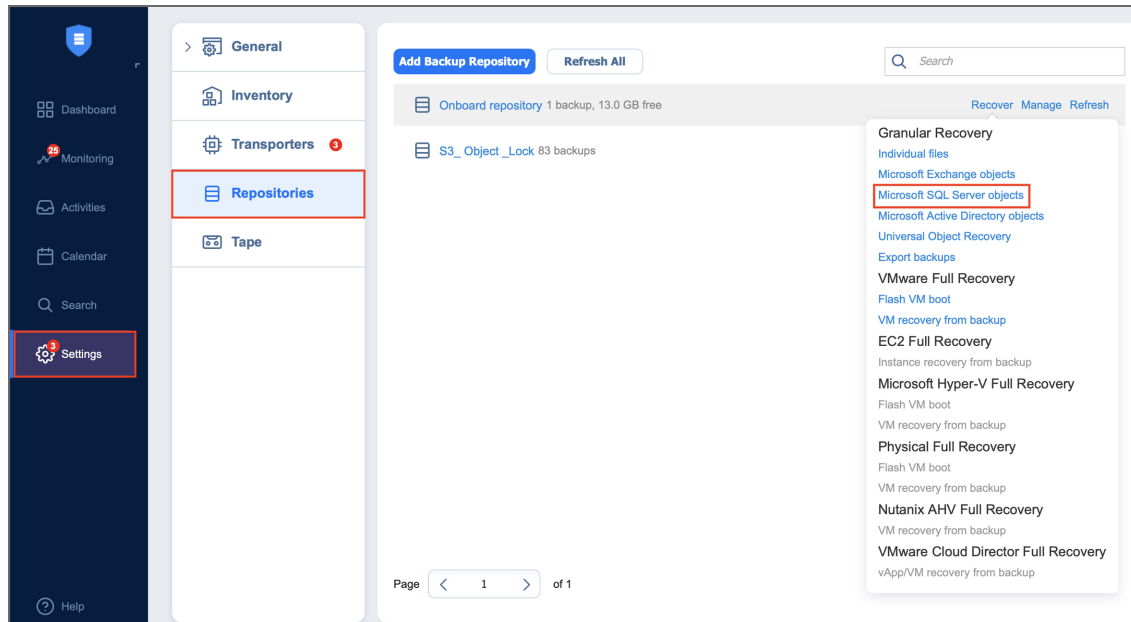
## Starting SQL Server Object Recovery from Backup Repository

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 **Recover** button and then click Microsoft SQL Server objects.

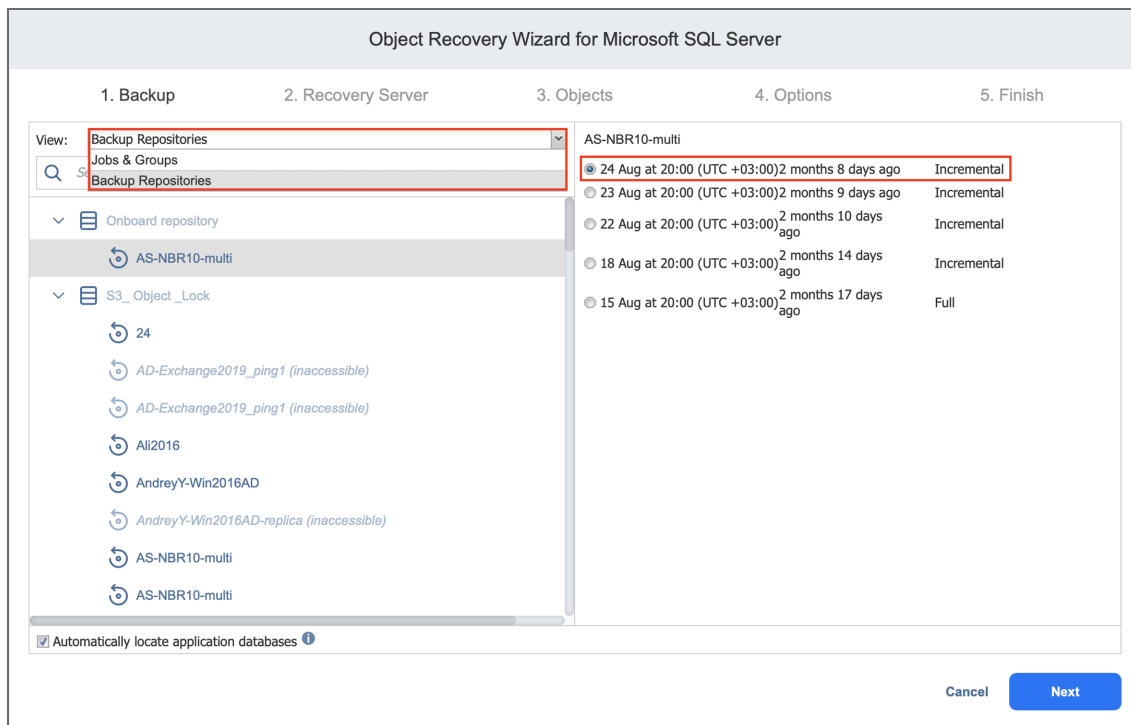


The **New Object Recovery Wizard for Microsoft SQL Server** opens.

## Object Recovery Wizard for Microsoft SQL Server: Backup

On the **Backup** page of the wizard:

1. Select a backup of a VM with Microsoft SQL in the left pane using either the **Backup Repositories** or **Jobs & Groups** view.
2. Select a recovery point in the right pane.



### 3. Click **Next**.

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.

## 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:

- a. **Recovery server:** Choose the target server from the drop-down list.

#### **Note**

NAKIVO Backup & Replication will try to auto-detect the IP address automatically.

- b. **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.
- 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. Refer to [“Requirements for Microsoft SQL Server Object Recovery” on page 130](#) 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.
- e. **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: ●●●●●●

[Manage credentials](#)

SQL instance: Select SQL Instance

Test Connection

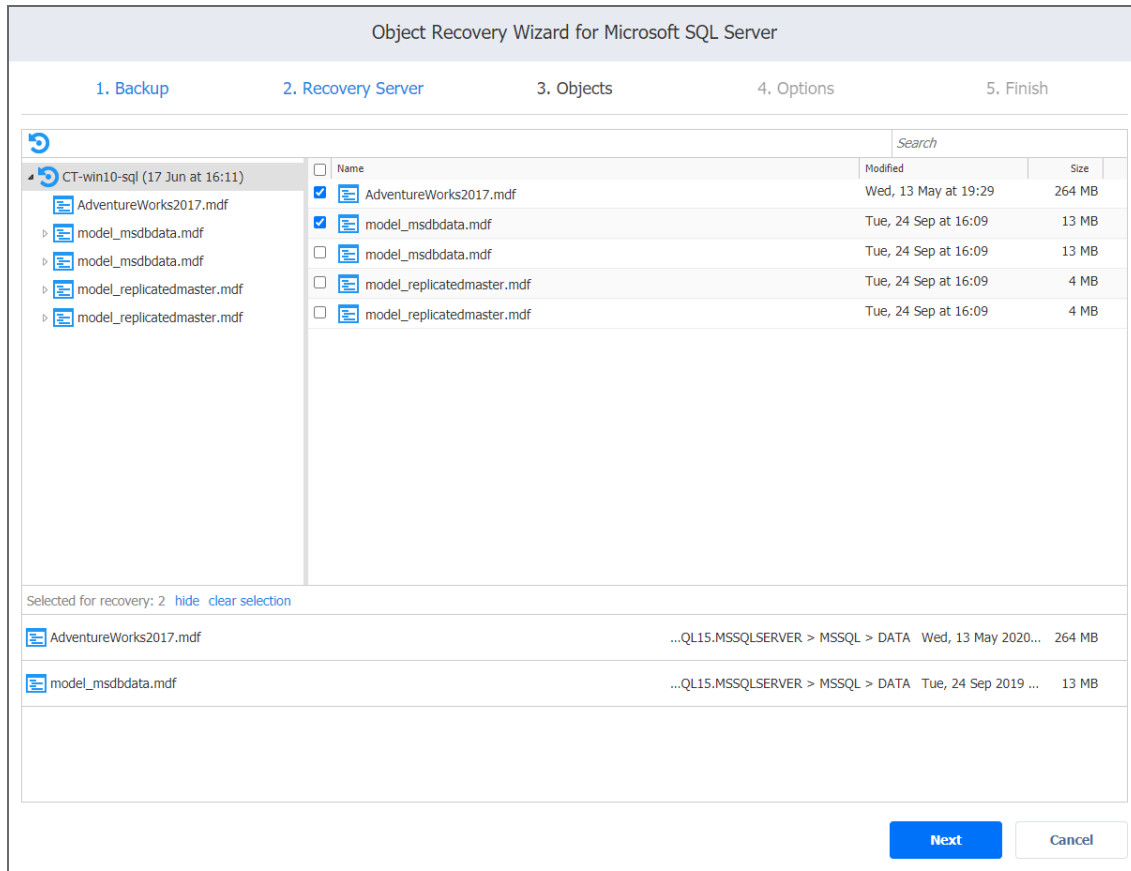
Cancel      Next

## 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**.

Object Recovery Wizard for Microsoft SQL Server

1. Backup      2. Recovery Server      3. Objects      4. Options      5. Finish

Recovery Scope

Recover schema and data

Recover only schema

Recovery Settings

Recovery type: Recover to original location

Overwrite behavior: Rename recovered item if such item exists

Recover      Cancel

## 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).

Object Recovery Wizard for Microsoft SQL Server

1. Backup      2. Recovery Server      3. Objects      4. Options      5. Finish

Recovery Scope

Recover schema and data

Recover only schema

Recovery Settings

Recovery type: Export

Export location: CIFS share

Path to the share: \\ServerName\FolderName

Username: Type or select username

Password: [Manage credentials](#)

Overwrite behavior: Rename recovered item if such item exists

Recover      Cancel

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**

The screenshot shows the 'Object Recovery Wizard for Microsoft SQL Server' at step 4, 'Options'. The wizard has five steps: 1. Backup, 2. Recovery Server, 3. Objects, 4. Options, and 5. Finish. Under 'Recovery Scope', 'Recover schema and data' is selected. Under 'Recovery Settings', 'Recovery type' is 'Recover to custom location', 'SQL instance' is 'MSSQLSERVER', and 'Path to the local folder' is 'C:\Folder\Subfolder'. The 'Overwrite behavior' dropdown menu is highlighted with a red box and shows the following options: 'Rename recovered item if such item exists', 'Rename recovered item if such item exists', 'Skip recovered item if such item exists', and 'Overwrite the original item if such item exists'. At the bottom right, there are 'Recover' and 'Cancel' buttons.

Click **Recover** to start the object recovery process. The **Finish** page opens.

## Object Recovery for Microsoft 365

The Object Recovery for Microsoft 365 feature enables browsing, searching, and recovering Exchange Online mailboxes, OneDrive for Business instances, and SharePoint Online sites directly from backups. This feature is agentless, works right out of the box, and does not require creating a special lab or running a special backup type.

### Important

- Refer to Microsoft 365 Platform Limitations of the latest [Release Notes](#) for the full list of Microsoft 365 items that are not supported.
- Refer to [Required API Permissions for Microsoft 365](#) to see the list of required permissions for recovering Microsoft 365 objects.

Refer to the following topics for more information:

- [“Starting Object Recovery for Microsoft 365” on the next page](#)
- [“Microsoft 365 Object Recovery Wizard: Backup” on page 833](#)
- [“Microsoft 365 Object Recovery Wizard: Recovery Account” on page 834](#)
- [“Microsoft 365 Object Recovery Wizard: Objects” on page 834](#)
- [“Microsoft 365 Object Recovery Wizard: Options” on page 836](#)

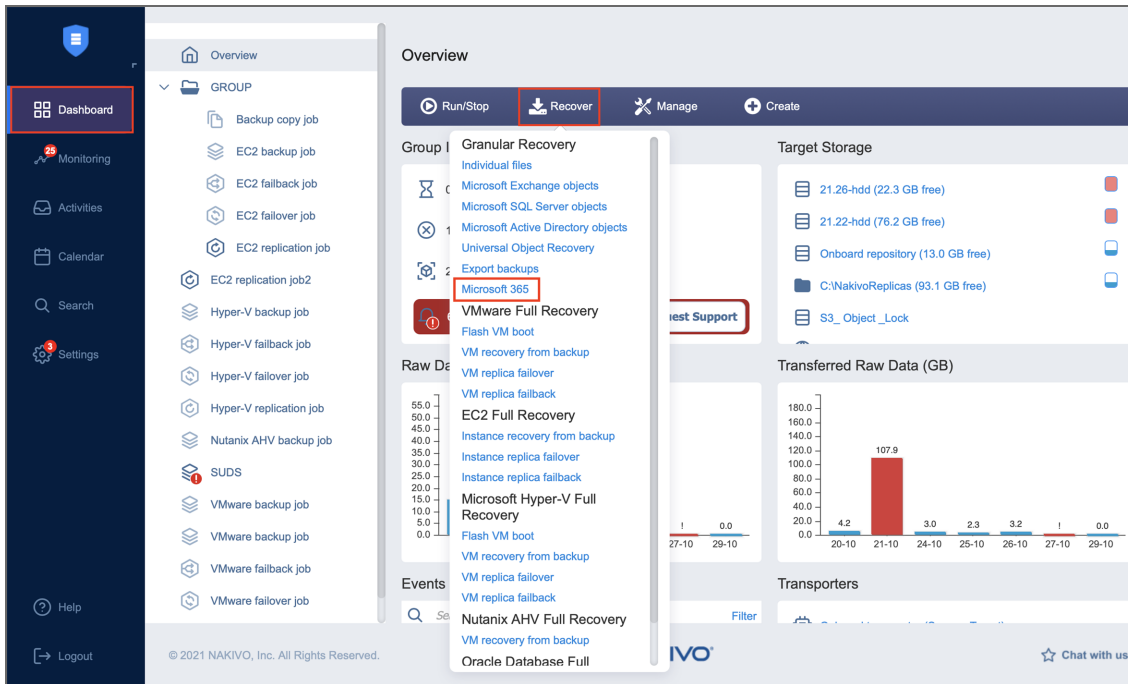
## Starting Object Recovery for Microsoft 365

You can start the recovery process either from the **Dashboard**, by using the [search](#) function [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 Object Recovery for Microsoft 365 from Dashboard](#)
- [Starting Object Recovery for Microsoft 365 from a Backup Repository](#)

## Starting Object Recovery for Microsoft 365 from Dashboard

To start object recovery job for Microsoft 365 from the **Dashboard**, click **Recover** and then click **Microsoft 365**.

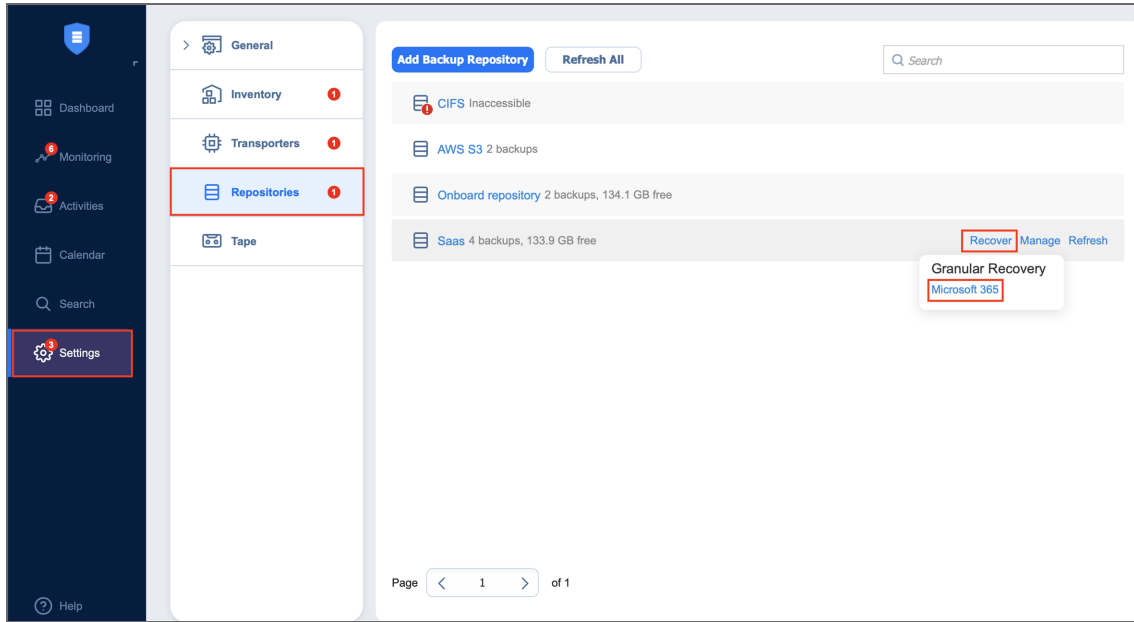


## Starting Object Recovery for Microsoft 365 from a Backup Repository

To start object recovery for Microsoft 365 from a Backup Repository, do the following:

1. Go to the main menu of NAKIVO Backup & Replication and click **Settings**.
2. Go to the **Repositories** tab and hover the cursor over the Backup Repository containing the required backup.

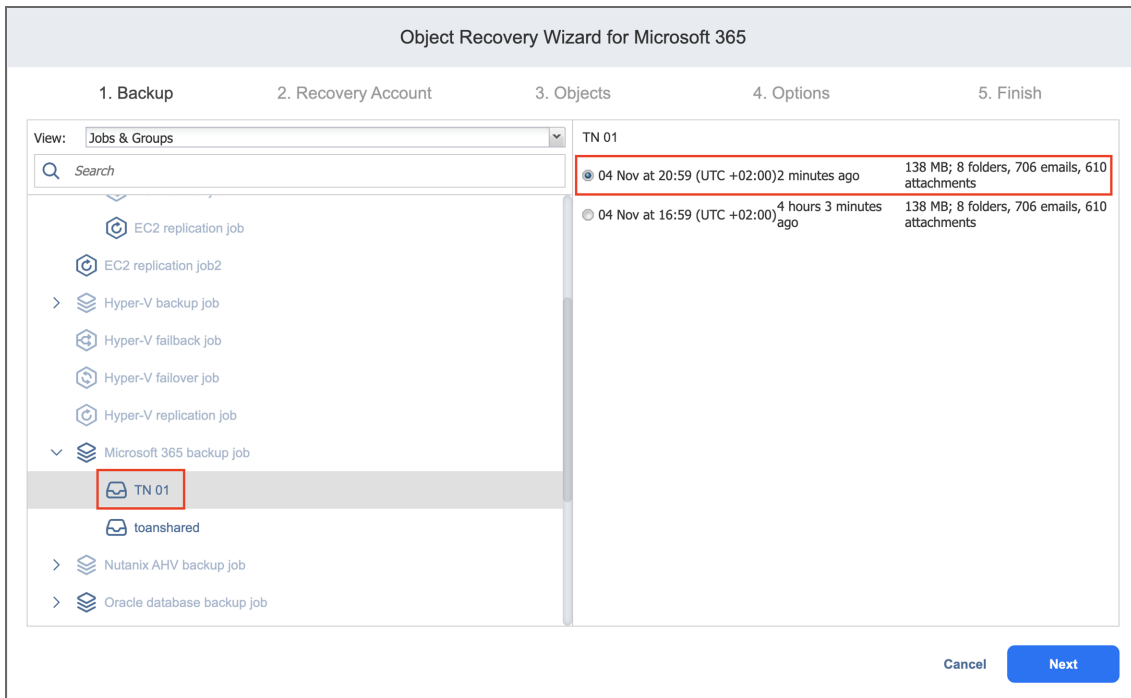
3. Click the **Recover** button and then click **Microsoft 365**.



The **Object Recovery Wizard for Microsoft 365** opens.

## Microsoft 365 Object Recovery Wizard: Backup

On the **Backup** page of the Object Recovery Wizard for Microsoft 365, select a backed up Exchange Online mailbox, OneDrive for Business, or a SharePoint Online site using either a **Backup Repository** or **Jobs & Groups** view in the left pane. Then select a recovery point in the right pane. Note, that only one backup can be selected at a time.

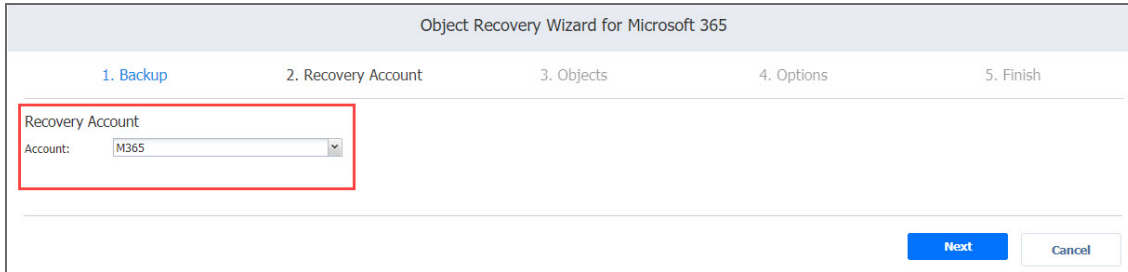


Click **Next** to proceed to the next page of the wizard.



## Microsoft 365 Object Recovery Wizard: Recovery Account

On the **Recovery Account** page, specify the Microsoft 365 account to which you want to recover your items. You can choose to recover to a specific Online Exchange mailbox, OneDrive for Business, or a SharePoint site later in the wizard. Select the required account from the **Account** drop-down list. The list contains all Microsoft 365 accounts added to the inventory. Refer to [“Adding Microsoft 365 Accounts” on page 345](#) for details.



Object Recovery Wizard for Microsoft 365

1. Backup    2. Recovery Account    3. Objects    4. Options    5. Finish

Recovery Account

Account: M365

Next    Cancel

Click **Next** to proceed to the next page of the wizard.

## Microsoft 365 Object Recovery Wizard: Objects

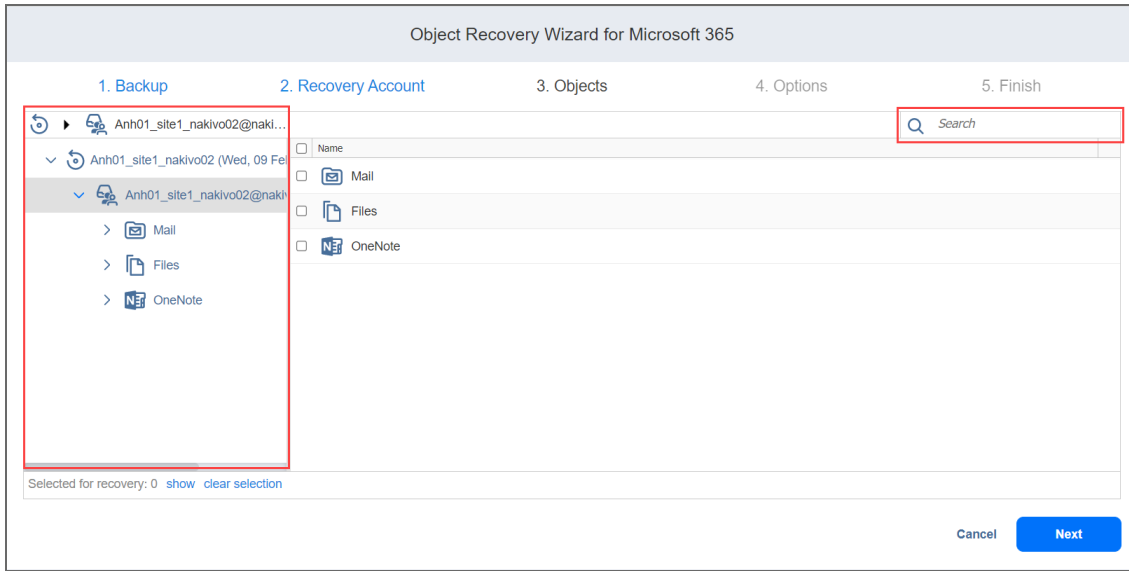
On the **Objects** page, the following items can be selected for recovery:

- Mailboxes, folders, contacts, calendar events, files, OneNote items, and emails
- Drives, folders, and files
- Sites, subsites, document libraries, lists, list items, individual files, and folders.
- **Notes**
  - Linked contacts are displayed as separate contact items.
  - The lookup data of SharePoint Online lists displayed in the recovery wizard is inaccurate, and this data can't be recovered.
  - Read-only calendars and calendars added from the directory can't be recovered.
  - Refer to the Microsoft 365 Platform Limitations of the latest [Release Notes](#) for the full list of Microsoft 365 items that are not supported.

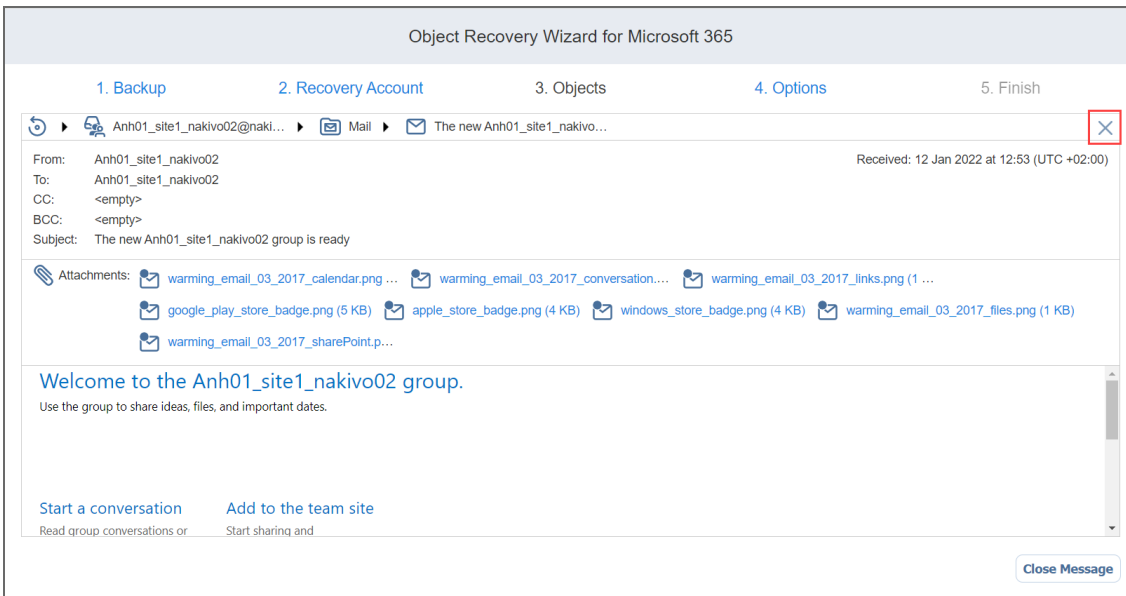
The **Objects** page contains:

- **Navigation:** Use the Navigation pane on the left to locate the items you need to restore. Selecting the container in the Navigation pane loads the container contents in the right pane. If a container has sub-folders, they are also displayed in the right pane.
- **Search:** The search box allows you to search for the objects that you want to recover. The search is performed within:
  - Mailboxes, folders, contacts, calendar items, files, OneNote items, and emails;
  - Drives, folders, and individual files;
  - Sites, document libraries, folders, files, lists, and list items.

Enter the word or a part of the word into the search box and hit Enter on the keyboard. All items matching the criteria will be displayed in the Contents pane on the right.



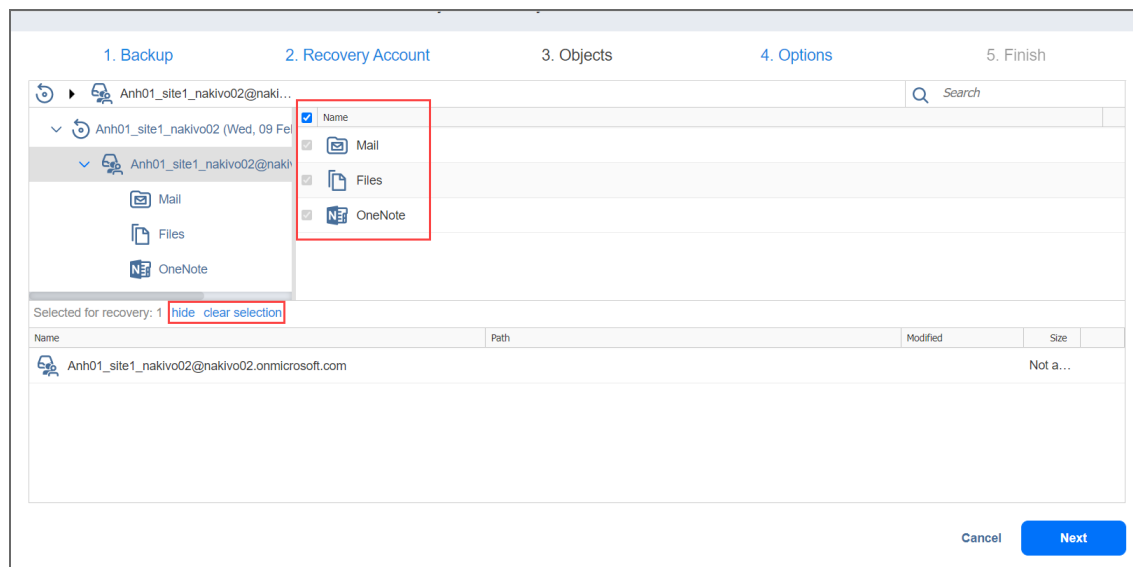
When recovering emails, contacts, or calendar events, you can also preview the contents of these items by clicking the item name in the right pane. To close the email message, contact, or calendar event preview, click the “Close” button at the bottom or click the “X” button above the item body.



After locating the items you need to recover, select the checkboxes next to their names. You can select different object types for recovery at the same time. 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.



Click **Next** to proceed to the next page of the wizard.

## Microsoft 365 Object Recovery Wizard: Options

On the **Options** page, specify the location to which the items will be restored and overwrite behavior for folders. Proceed as described in the sections below:

- [Recovering to the Original Location](#)
- [Recovering to Mailbox or to Group](#)
- [Recovering to OneDrive](#)
- [Recovering to Site](#)

### Recovering to the Original Location

To recover Microsoft 365 items to the original location, take the following steps:

#### Note

- Group mailboxes can only be recovered to the original location or other group mailbox types.
- User mailboxes can be recovered to the original location or other group mailbox types and cannot be recovered to the group mailbox type.
- Personal site non-personal site items can be recovered to group site type and vice versa.

1. Select **Recover to original location** from the **Recovery type** drop-down list. All selected items will be recovered to their original locations within the Microsoft 365 account. This option is disabled if a mailbox, OneDrive, or site no longer exists (deleted from the product, for example) or if you have selected a recovery account on the **Recovery Account** page of the wizard that is different from the original account.
2. Choose the naming convention for the recovered items by selecting one of the following **Overwrite behavior** options:
  - **Rename recovered item if such item exists**
  - **Skip recovered item if such item exists**
  - **Overwrite the original item if such item exists**
3. Click **Recover**.

Object Recovery Wizard for Microsoft 365

1. Backup      2. Recovery Account      3. Objects      4. Options      5. Finish

Recovery type:

Overwrite behavior:

## Recovering to Mailbox or to Group

If you want to recover items to a group or to a specific mailbox, take the following steps:

1. Select **Recover to mailbox** from the **Recovery type** drop-down list. Note that **Recover to group** is displayed if group mailbox items were selected.
2. Select the required mailbox from the **Mailbox** drop-down list. You may also search for the required mailbox by typing its name or part of its name into the search bar.

### Note

- **Group** is displayed if group mailbox items were selected.
  - The following may occur if recovering Group messages to a different Group mailbox:
    - The original sender of the message may be changed to the name of the user used for account discovery.
    - The original receiver of the message may be changed to the target user group mailbox name.
    - The timestamp of the message may be changed.
3. Choose the overwrite behavior for the recovered items by selecting one of the following **Overwrite behavior** options:

- **Rename recovered item if such item exists**
- **Skip recovered item if such item exists**
- **Overwrite the original item if such item exists**

4. Click **Recover**.

Object Recovery Wizard for Microsoft 365

1. Backup      2. Recovery Account      3. Objects      4. Options      5. Finish

Recovery type: Recover to mailbox

Mailbox: Automation06

Overwrite behavior:
 

- Rename recovered item if such item exists
- Rename recovered item if such item exists
- Skip recovered item if such item exists
- Overwrite the original item if such item exists

Recover      Cancel

## Recovering to OneDrive

To recover OneDrive items to a specific OneDrive account, take the following steps:

1. Select **Recover to OneDrive** from the **Recovery type** drop-down list.
2. Select the required OneDrive from the **OneDrive** drop-down list. You may also search for the required OneDrive by typing its name or part of its name into the search bar.
3. Choose the overwrite behavior for the recovered items by selecting one of the following **Overwrite behavior** options:
  - **Rename recovered item if such item exists**
  - **Skip recovered item if such item exists**
  - **Overwrite the original item if such item exists**

4. Click **Recover**.

Object Recovery Wizard for Microsoft 365

1. Backup      2. Recovery Account      3. Objects      4. Options      5. Finish

Recovery type: Recover to OneDrive

OneDrive: nt

Overwrite behavior:
 

- Rename recovered item if such item exists
- Rename recovered item if such item exists
- Skip recovered item if such item exists
- Overwrite the original item if such item exists

Recover      Cancel

## Recovering to Site

To recover SharePoint objects to a specific site, take the following steps:

1. Select **Recover to site** from the **Recovery type** drop-down list.
2. Select the required site from the **Site** drop-down list. You may also search for the required site by typing its name or part of its name into the search bar.
3. Choose the overwrite behavior for the recovered items by selecting one of the following **Overwrite behavior** options:
  - **Rename recovered item if such item exists**
  - **Skip recovered item if such item exists**
  - **Overwrite the original item if such item exists**
4. Click **Recover**.

Object Recovery Wizard for Microsoft 365

1. Backup      2. Recovery Account      3. Objects      4. Options      5. Finish

Recovery type:

Site:

Overwrite behavior:

When the recovery process is completed, the **Finish** page is displayed. You cannot return to the previous pages at this point; however, you can check the progress of the job by clicking the **Activities** link.

## Object Recovery for Oracle Database

NAKIVO Backup & Replication [object recovery](#) feature enables browsing, searching, and recovering Oracle databases directly from compressed and deduplicated backups. Refer to the following topics for more information:

- [“Starting Object Recovery for Oracle Database” below](#)
- [“Recovery Job Wizard for Oracle Database: Backups” on the next page](#)
- [“Recovery Job Wizard for Oracle Database: Options” on page 841](#)

## Starting Object Recovery for Oracle Database

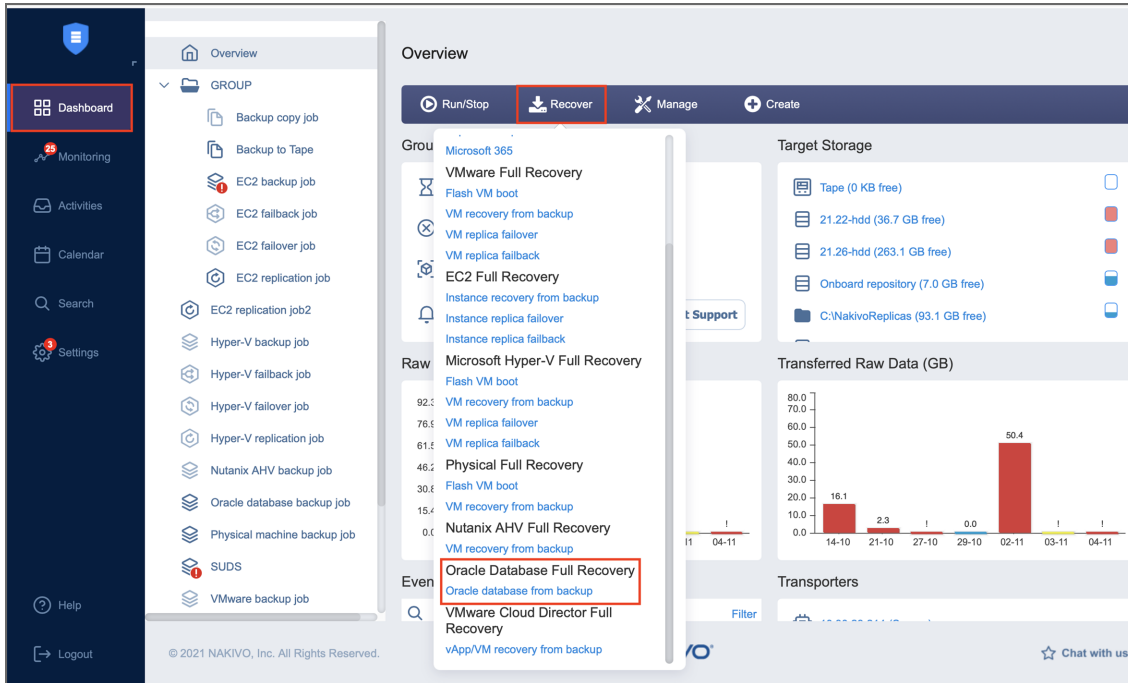
You can start the recovery process either from the **Dashboard**, by using the [search](#) function, or from the **Repositories** tab in **Settings** (for example, if you no longer have a backup job, but still have the backup).

Refer to the following sections for more details:

- [“Starting Object Recovery for Oracle Database from Dashboard ” on the next page](#)
- [“Starting Object Recovery for Oracle Database from Backup Repository” on the next page](#)

## Starting Object Recovery for Oracle Database from Dashboard

To start object recovery for Oracle database from the **Dashboard**, click **Recover** and then click **Oracle database from backup**.



## Starting Object Recovery for Oracle Database from Backup Repository

To start object recovery for Oracle database from a Backup Repository, do the following:

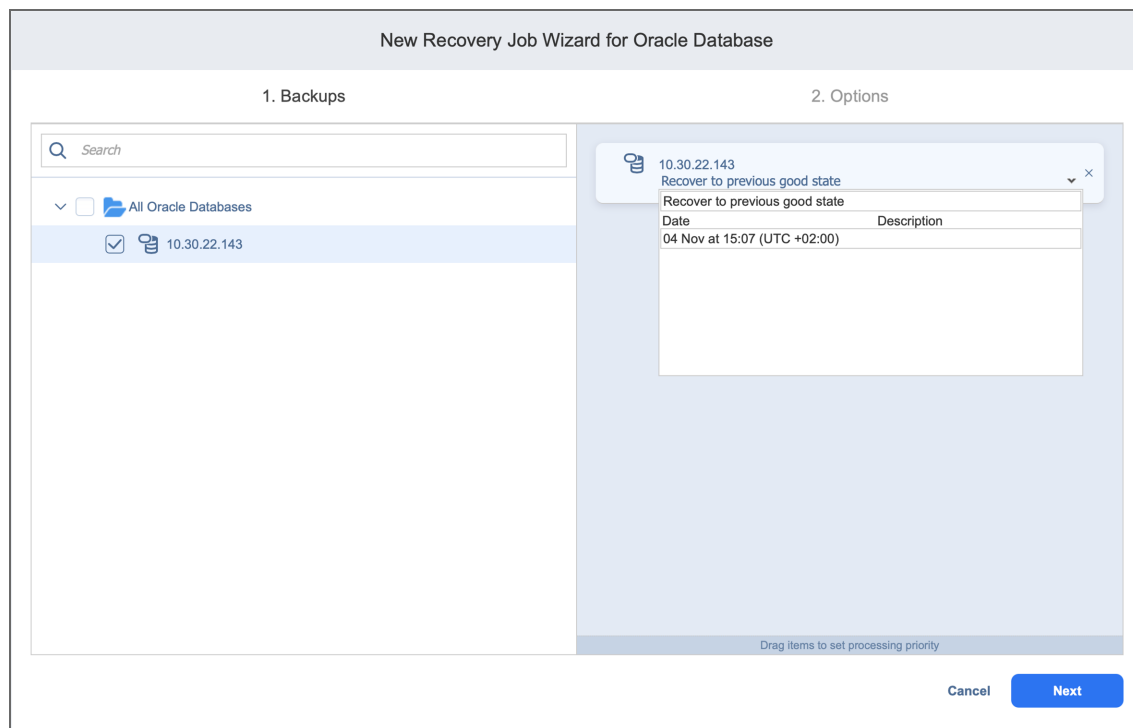
1. Click **Settings** in the upper right corner of the product.
2. Go to the **Repositories** tab and hover over the Backup Repository containing the required backup.
3. Click the **Recover** button and then click **Oracle database from backup**.

## Recovery Job Wizard for Oracle Database: Backups

On the first page of the wizard,

1. Select a backup with Oracle database in the left pane.
2. Select a recovery point in the right pane.

3. Click **Next** to move to the next page.



## Recovery Job Wizard for Oracle Database: Options

On the **Options** page, you can name the recovery job and set pre and post job actions.

- [Job Options](#)
- [Pre and Post Actions](#)
  - [Email Notifications](#)
  - [Pre Job Script](#)
  - [Post Job Script](#)

### Job Options

**Job name** : Enter a name for the job into the Job name box.

### Pre and Post Actions

NAKIVO Backup & Replication provides you with the ability to enable certain actions before recovery begins and after it has been completed. You can choose to send job run reports, 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 notification and provides you with the ability to configure notifications on a per-job level. To send email notifications, select **Send job reports to** and specify one or more email addresses in the text box. The semi-colon character (;) should be used to separate multiple email addresses.



## Note

To enable this option, configure your [email settings](#).

## Pre Job Script

To run a [script](#) before the product begins recovering databases, 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. Script interpreter should be specified.
  - **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, recovery 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 recovering databases 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 recovery 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 recovery will not be performed.

## Post Job Script

To run a script after the product has finished recovering all databases, 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.
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. Script interpreter should be specified.
  - **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 recovery has been successful.

When all options are set, click **Finish & Run** to start the recovery process.

New Recovery Job Wizard for Oracle Database

1. Backups 2. Options

Job Options  
Job name:

Pre and Post Actions

Send job run reports to  ?

Run local pre job script ?

Run local post job script ?

**i** All other settings for selected items should be pre-configured in the Oracle database RMAN configuration. [Learn more](#)

## 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” below](#)
- [“Universal Object Recovery Wizard: Backup” on the next page](#)
- [“Universal Object Recovery Wizard: Disks” on page 845](#)
- [“Universal Object Recovery Wizard: Options” on page 846](#)

## Opening Universal Object Recovery Wizard

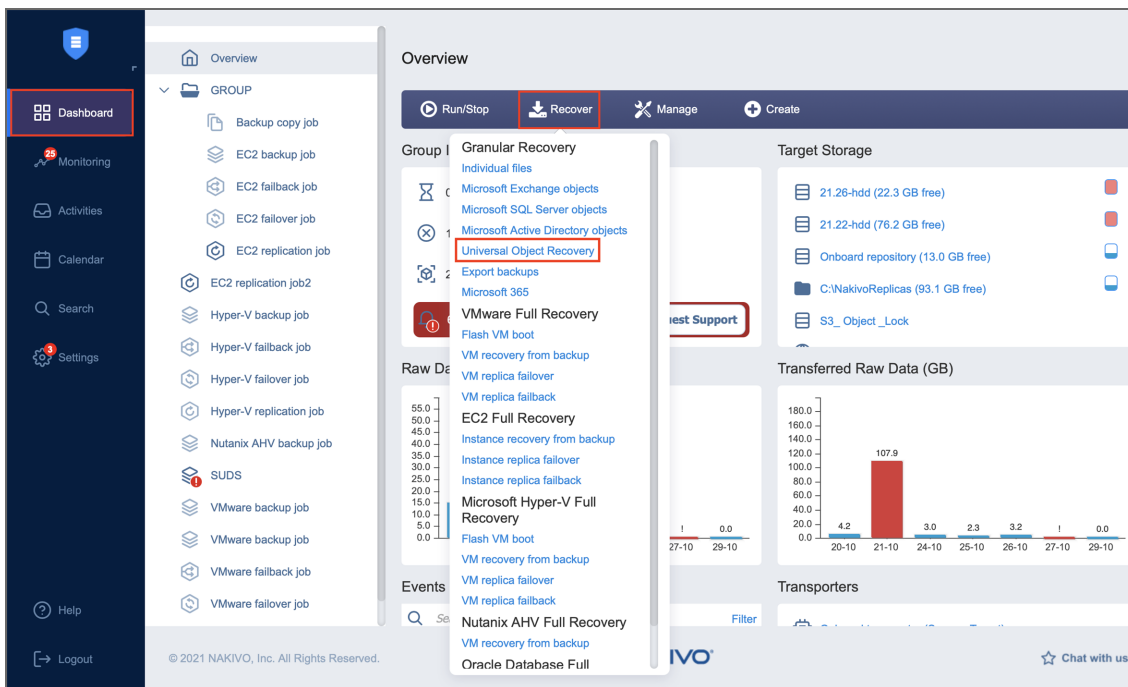
You can start the recovery process either from the **Dashboard**, by using the [search](#) function, or from the **Repositories** page in **Settings** (for example, if you no longer have a backup job but still have the backup).

Refer to the following sections for more details:

- [Starting Universal Object Recovery from Dashboard](#)
- [Starting Universal Object Recovery from a Backup Repository](#)

## Starting Universal Object Recovery from Dashboard

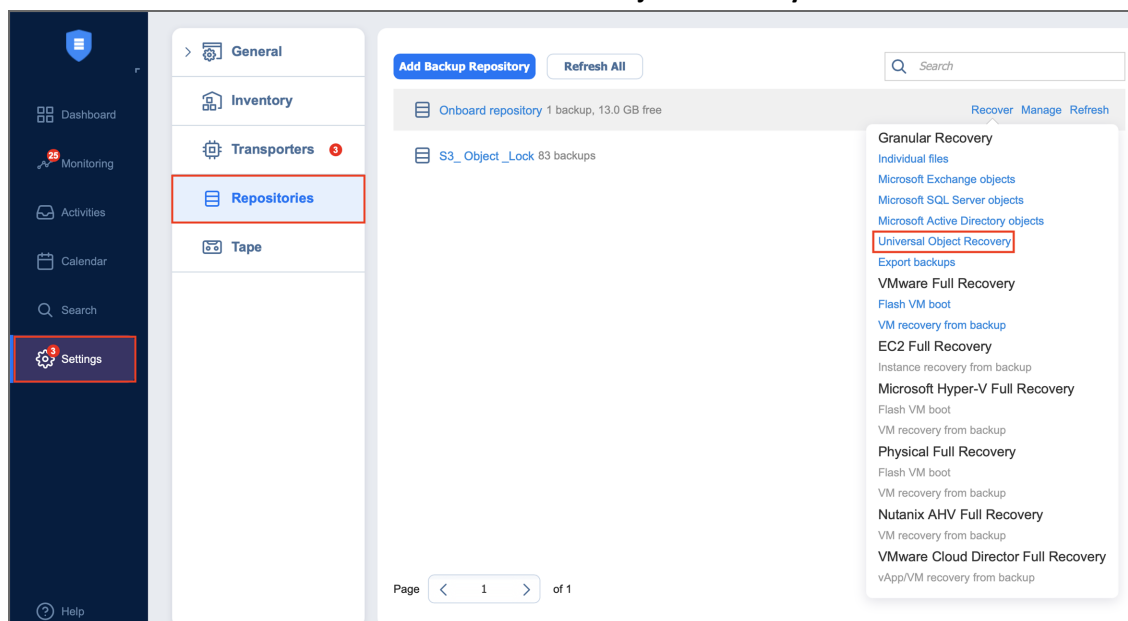
To start Universal object recovery from the **Dashboard**, click **Recover** and choose **Universal Object Recovery**.



## 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 **Recover** button and then click **Universal Object Recovery**.

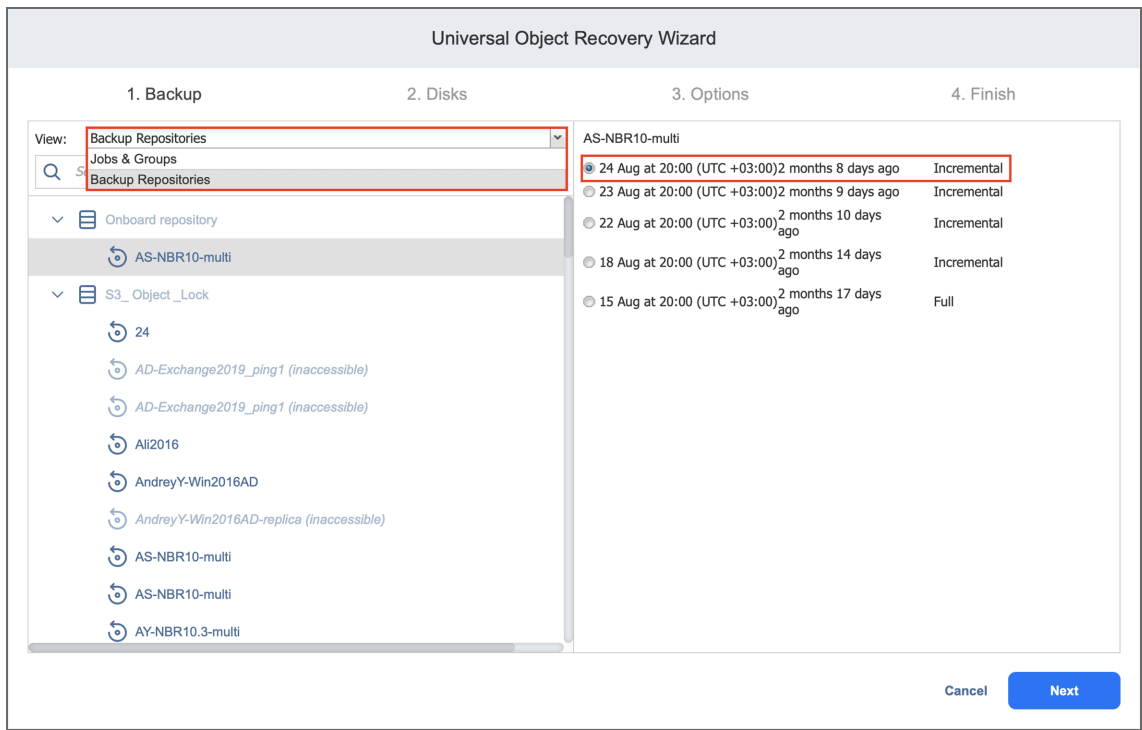


The new **Universal Recovery Job Wizard** opens.

## Universal Object Recovery Wizard: Backup

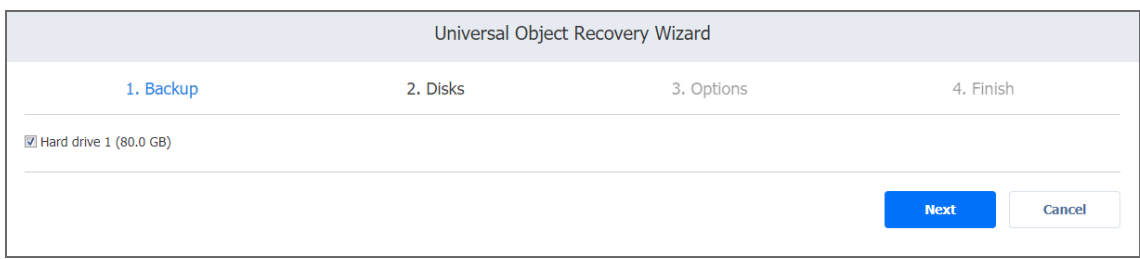
On the **Backup** page of the wizard, do the following:

1. Optionally, you can filter the items tree by entering a string into the **Search** box. You can enter a part of or the entire name of the item.
2. Choose either of the following item views:
  - **Backup Repositories:** When chosen, the Backup Repositories tree opens in the left pane. Proceed as follows:
    - a. Expand a repository by clicking the icon to the left of the repository.
    - b. Choose a backup in the left pane and then choose a recovery point in the right pane.
  - **Jobs & Groups:** When chosen, the jobs' tree opens in the left pane. Proceed as follows:
    - a. Expand a job by clicking the icon to the left of the job.
    - b. Choose a backup in the left pane and then choose a recovery point in the right pane.
3. Click **Next** to go to the next page of the wizard.



## 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.
  - **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. Click **Mount** to confirm mounting your disks to the selected recovery server.

The screenshot shows the 'Options' step of the 'Universal Object Recovery Wizard'. The wizard has four steps: 1. Backup, 2. Disks, 3. Options (current), and 4. Finish. The 'Options' step contains the following fields and controls:

- Mount location:** A dropdown menu with 'AS-NBR10-multi' selected.
- Location IP address:** A text input field containing '10.30.23.176'.
- Use custom SSH port:** A checkbox that is checked, with a text input field containing '22'.
- Credentials type:** A dropdown menu with 'Password' selected.
- Username:** A text input field containing 'admin'.
- Password:** A text input field with masked characters (dots).
- Test Connection:** A blue button.
- Manage credentials:** A blue link below the password field.
- Cancel:** A grey button at the bottom right.
- Mount:** A blue button at the bottom right.

The Universal Object Recovery is started and the **Finish** page of the wizard opens.

4. Click the **Activities** link to go to the **Activities** page if you want to view the progress of the Universal Object Recovery.
5. Click **Close** to close the Universal Object Recovery Wizard. Upon successful Universal Object Recovery, the disks are mounted to the recovery server.

# Full Recovery

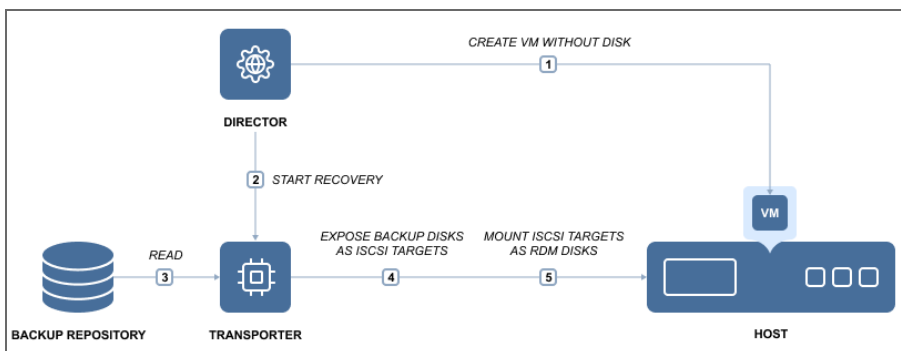
With NAKIVO Backup & Replication, you can recover an entire VM in case of hardware or VM failure. The VM is recovered in the same state as it was during the backup and it will appear on the host selected for recovery. You can also perform cross-platform recovery that allows you to export virtual disks from VM backups to different formats for further manual recovery of the VMs in different virtual environments. The Flash VM 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. 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 VM Boot Recovery” below](#)
- [“VMware VM Recovery” on page 888](#)
- [“VMware Cloud Director Recovery” on page 902](#)
- [“Hyper-V VM Recovery” on page 912](#)
- [“Amazon EC2 Instance Recovery” on page 923](#)
- [“Physical Machine Recovery” on page 934](#)
- [“Nutanix AHV VM Recovery” on page 945](#)
- [“Recovery From Tape” on page 956](#)
- [“Performing Cross-Platform Recovery” on page 964](#)

## Performing Flash VM Boot Recovery

The [Flash VM 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.



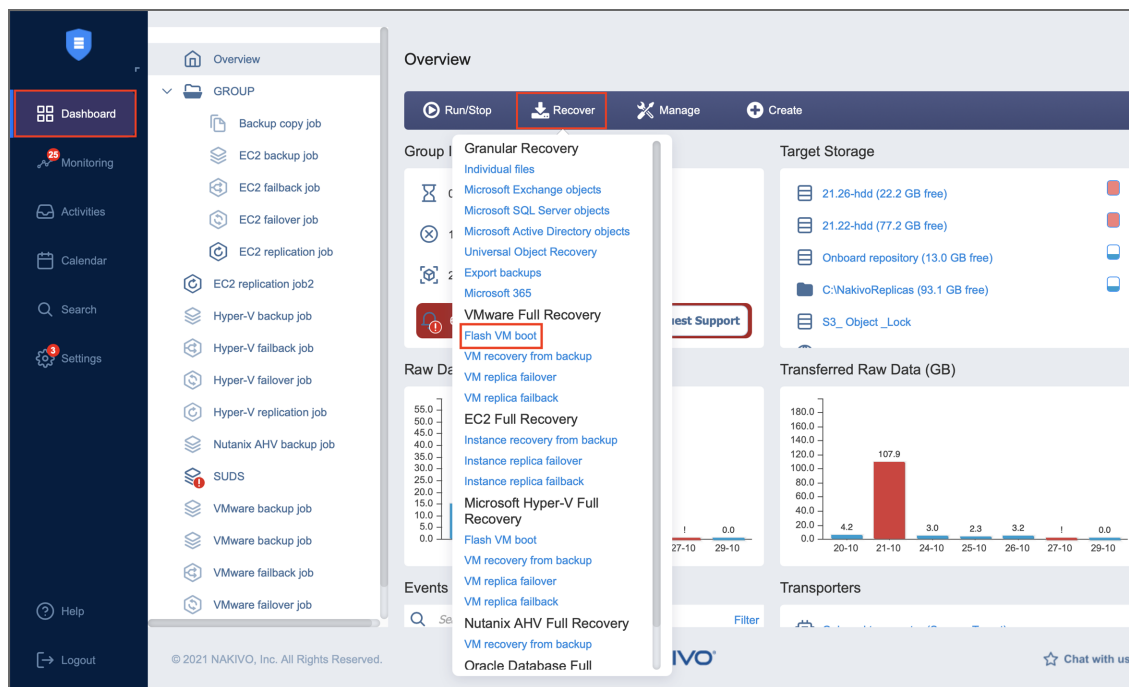
When the VMware VM is created, NAKIVO Backup & Replication takes a snapshot of the VM: this way all changes that occur to the VM are temporarily stored in the snapshot and discarded when you stop the job. When the Hyper-V VM is created, the application temporarily stores the changes to the VM in a disk-based 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 VM Boot Jobs” below](#)
- [“Creating Hyper-V Flash VM Boot Jobs” on page 861](#)
- [“Creating Flash Boot Jobs for Physical Machines” on page 873](#)
- [“Migrating Recovered VMs Using Flash Boot” on page 886](#)

## Creating VMware Flash VM Boot Jobs

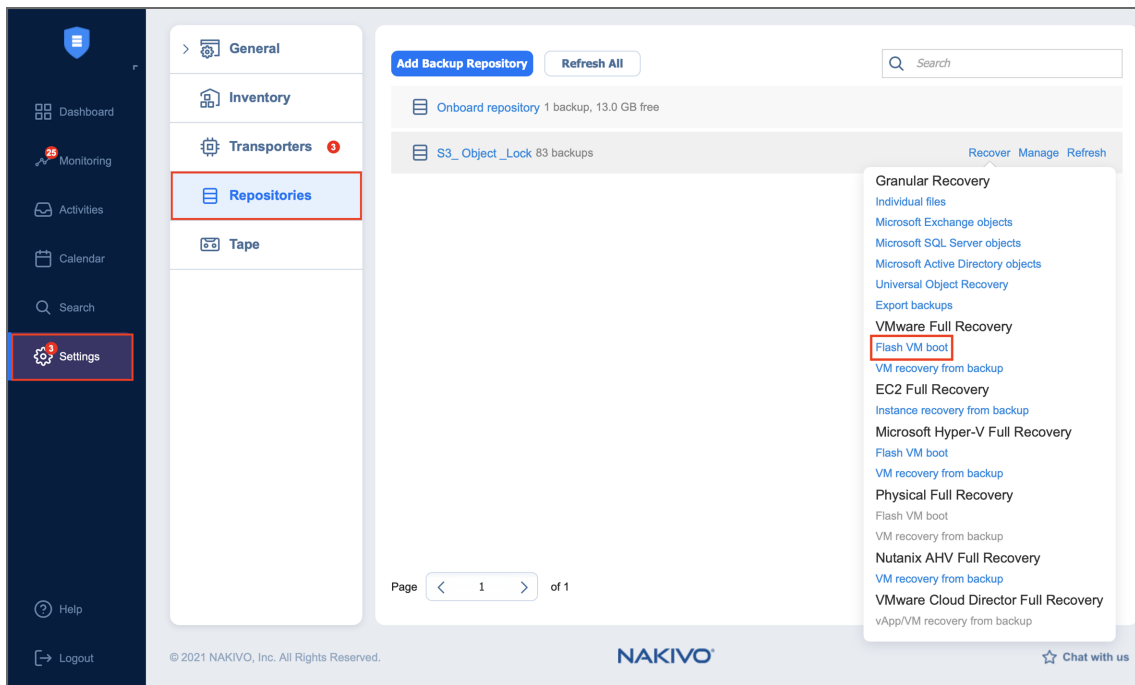
To create a VMware Flash VM Boot job, do one of the following:

- Open the Flash VM Boot wizard from the **Dashboard** by clicking **Recover** and then clicking **Flash VM Boot**.



- Open the Flash VM 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 select a Backup Repository.

3. In the Backup Repository title, click **Recover** and then click **Flash VM Boot**.



- Alternatively, the recovery can be performed by using the [search](#) function.

The **New Flash Boot Job Wizard** opens.

- [“VMware Flash VM Boot Job Wizard: Backups”](#) below
- [“VMware Flash VM Boot Job Wizard: Destination”](#) on the next page
- [“VMware Flash VM Boot Job Wizard: Schedule”](#) on page 852
- [“VMware Flash VM Boot Job Wizard: Options”](#) on page 855


## VMware Flash VM Boot Job Wizard: Backups

On the **Backups** page of the wizard, select backups for recovery by taking the following actions:

1. Select one of the views:
  - **Jobs & Groups:** Select one or more VM backups in the left pane and then select a recovery point for each backup in the right pane.
  - **Backup Repositories:** Select one or more Backup Repositories in the left pane and then select a



recovery point for each backup in the right pane.



2. Click **Next** to go to the next page of the wizard.

## VMware Flash VM 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” on the next page](#)

### 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.

New Flash Boot Job Wizard for VMware vSphere

1. Backups      2. Destination      3. Schedule      4. Options

Container: 10.30.21.22

Datastore: 21.22-hdd

Network: 10.30.21.0

VM folder: Select target VM folder (optional)

Advanced setup...

Cancel Next

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.

New Flash Boot Job Wizard for VMware vSphere

1. Backups      2. Destination      3. Schedule      4. Options

AS-NBR10-multi Click to collapse

<p>Source</p> <p>VM location: Onboard repository</p> <p>VM resources: 2 CPU, 4.0 GB RAM</p> <p>Disks</p> <p>Hard disk 1: Onboard repository (50.0 GB)</p> <p>VM file: Onboard repository</p> <p>Network adapters</p> <p>Network adapter 1</p>	<p>Target</p> <p>Container: 10.30.21.22</p> <p>Virtual Machine: <i>New VM will be created</i></p> <p>VM folder: <i>Select target VM folder (optional)</i></p> <p>Disks</p> <p>Hard disk 1: 21.22-hdd</p> <p>VM file: 21.22-hdd</p> <p>Network adapters</p> <p>Network adapter 1: 10.30.21.0</p>
---	---

AY-NBR10.3-multi Click to collapse

<p>Source</p> <p>VM location: S3_Object_Lock</p> <p>VM resources: 2 CPU, 4.0 GB RAM</p> <p>Disks</p> <p>Hard disk 1: S3_Object_Lock (50.0 GB)</p> <p>VM file: S3_Object_Lock</p> <p>Network adapters</p> <p>Network adapter 1</p>	<p>Target</p> <p>Container: 10.30.21.22</p> <p>Virtual Machine: <i>New VM will be created</i></p> <p>VM folder: <i>Select target VM folder (optional)</i></p> <p>Disks</p> <p>Hard disk 1: 21.22-hdd</p> <p>VM file: 21.22-hdd</p> <p>Network adapters</p> <p>Network adapter 1: 10.30.21.0</p>
---	---

Cancel    **Next**

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.

3. Click **Next**.

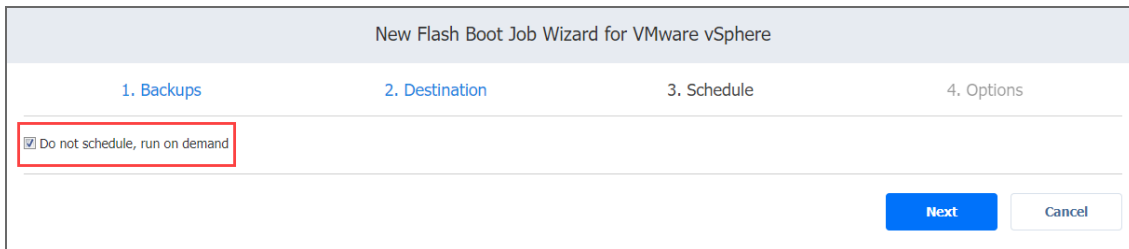
## VMware Flash VM Boot Job Wizard: Schedule

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” on the next page](#)
- [“Daily Job Execution” on the next page](#)
- [“Monthly or Yearly Job Execution” on the next page](#)
- [“Periodic Job Execution” on page 854](#)
- [“Chained Job” on page 855](#)
- [“Additional Schedule” on page 855](#)

## Disabling Scheduled Job Execution

If you want to start the job manually only (without scheduling), select the **Do not schedule, run on demand** checkbox:



New Flash Boot Job Wizard for VMware vSphere

1. Backups 2. Destination 3. Schedule 4. Options

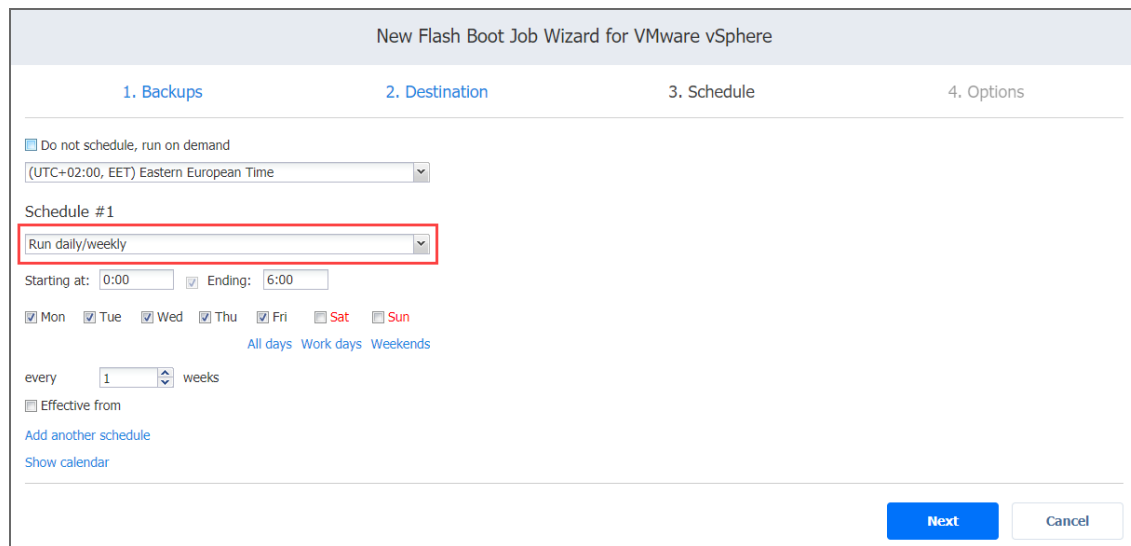
Do not schedule, run on demand

Next Cancel

## 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.



New Flash Boot Job Wizard for VMware vSphere

1. Backups 2. Destination 3. Schedule 4. Options

Do not schedule, run on demand

(UTC+02:00, EET) Eastern European Time

Schedule #1

Run daily/weekly

Starting at: 0:00 Ending: 6:00

Mon  Tue  Wed  Thu  Fri  Sat  Sun

All days Work days Weekends

every 1 weeks

Effective from

[Add another schedule](#)

[Show calendar](#)

Next Cancel

## 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.

New Flash Boot Job Wizard for VMware vSphere

1. Backups 2. Destination 3. Schedule 4. Options

Do not schedule, run on demand  
 (UTC+02:00, EET) Eastern European Time

Schedule #1  
 Run monthly/yearly

Run every last Friday of every month

Starting at: 0:00 Ending: 6:00

Effective from

[Add another schedule](#)  
[Show calendar](#)

Next Cancel

## 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.

New Flash Boot Job Wizard for VMware vSphere

1. Backups 2. Destination 3. Schedule 4. Options

Do not schedule, run on demand  
 (UTC+02:00, EET) Eastern European Time

Schedule #1  
 Run periodically every 30 minutes

Starting at: 0:00 Ending: 6:00

Mon  Tue  Wed  Thu  Fri  Sat  Sun  
 All days Work days Weekends

Effective from

[Add another schedule](#)  
[Show calendar](#)

Next Cancel

## 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.

The screenshot shows the 'New Flash Boot Job Wizard for VMware vSphere' interface. The wizard is currently on the '3. Schedule' step. The '1. Backups' step is highlighted in blue. The '2. Destination' step is also highlighted in blue. The '4. Options' step is not highlighted. The 'Do not schedule, run on demand' checkbox is unchecked. The time zone is set to '(UTC+02:00, EET) Eastern European Time'. The 'Schedule #1' dropdown menu is open, and 'Run after another job' is selected. The 'After the job:' dropdown menu is set to 'Hyper-V backup job'. The 'Run this job:' dropdown menu is set to 'Immediately'. The 'After successful runs' checkbox is checked, while 'After failed runs' and 'After stopped runs' are unchecked. The 'Effective from' checkbox is unchecked. There are links for 'Add another schedule' and 'Show calendar'. At the bottom right, there are 'Next' and 'Cancel' buttons.

## Additional Schedule

If you need to add an additional schedule, click **Add another schedule** and configure it as described above.

## VMware Flash VM Boot Job Wizard: Options

On the **Options** page, set up job options as described in these sections:

- [“Job Options” on the next page](#)
  - [“Job Name” on the next page](#)
  - [“VM Verification” on the next page](#)
- [“Recovery Options” on page 857](#)
  - [“Specifying VM Names” on page 858](#)
  - [“Generating VM MAC Addresses” on page 858](#)
  - [“Powering Recovered VMs” on page 858](#)
- [“Pre and Post Actions” on page 859](#)

- [“Setting up Email Notifications for the Job” on page 859](#)
- [“Setting Up a Pre-Job Script” on page 859](#)
- [“Setting Up a Post-Job Script” on page 860](#)
- [“Data Routing” on page 861](#)
- [“Completing the New Flash VM Boot Job Wizard” on page 861](#)

## Job Options

In this section, specify a job name and select a VM verification type.

The screenshot shows the 'New Flash Boot Job Wizard for VMware vSphere' interface, specifically the 'Options' step. The 'Job Options' section is highlighted with a red box. It includes the following fields and options:

- Job name:** VMware flash VM boot job
- VM verification:** Disabled
- Recovered VM Options:**
  - Recovered VM names: Append "-recovered" in the end
  - VM MAC addresses: Do not generate new MAC address
  - VM power on: Power on recovered VMs
- Pre and Post Actions:**
  - Send job run reports to
  - Run local pre job script
  - Run local post job script
- Data routing:** Do not use proxy transporter

Buttons at the bottom right include 'Finish', 'Finish & Run', and 'Cancel'.

## Job Name

Enter a name for your job in the **Job name** box.

## 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 VM 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 VM 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.

### Recovery Options

In this section, specify VM names, generate VM MAC addresses, and choose whether you want to power on recovered VMs or not.

New Flash Boot Job Wizard for VMware vSphere

1. Backups      2. Destination      3. Schedule      4. Options

Job Options

Job name: VMware flash VM boot job

VM verification: Disabled

**Recovered VM Options**

Recovered VM names: Append "-recovered" in the end

VM MAC addresses: Do not generate new MAC address

VM power on: Power on recovered VMs

Pre and Post Actions

Send job run reports to

Run local pre job script

Run local post job script

Data routing

Proxy transporter: Do not use proxy transporter

Finish      Finish & Run      Cancel



## 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 VM 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.

The screenshot shows the 'New Flash Boot Job Wizard for VMware vSphere' interface, specifically the '4. Options' step. The 'Pre and Post Actions' section is highlighted with a red box. It contains the following options:

- Send job run reports to [text box]
- Run local pre job script
- Run local post job script

Other visible options include:

- Job name: VMware flash VM boot job
- VM verification: Disabled
- Recovered VM names: Append "-recovered" in the end
- VM MAC addresses: not generate new MAC addresses
- VM power on: Power on recovered VMs
- Proxy transporter: Do not use proxy transporter

Buttons at the bottom: Finish, Finish & Run, Cancel.

## 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.

New Flash Boot Job Wizard for VMware vSphere

1. Backups      2. Destination      3. Schedule      4. Options

Job Options  
Job name: VMware flash VM boot job  
VM verification: Disabled

Recovered VM Options  
Recovered VM names: Append "-recovered" in the end  
VM MAC addresses: Do not generate new MAC address  
VM power on: Power on recovered VMs

Pre and Post Actions  
 Send job run reports to  
 Run local pre job script  
 Run local post job script

**Data routing**  
Proxy transporter: Do not use proxy transporter

- 10.30.22.197
- 10.30.30.86
- Hyper-V
- Nutanix Transporter
- Onboard transporter
- Wind2

Finish      Finish & Run      Cancel

## Completing the New Flash VM 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 256](#) for details.

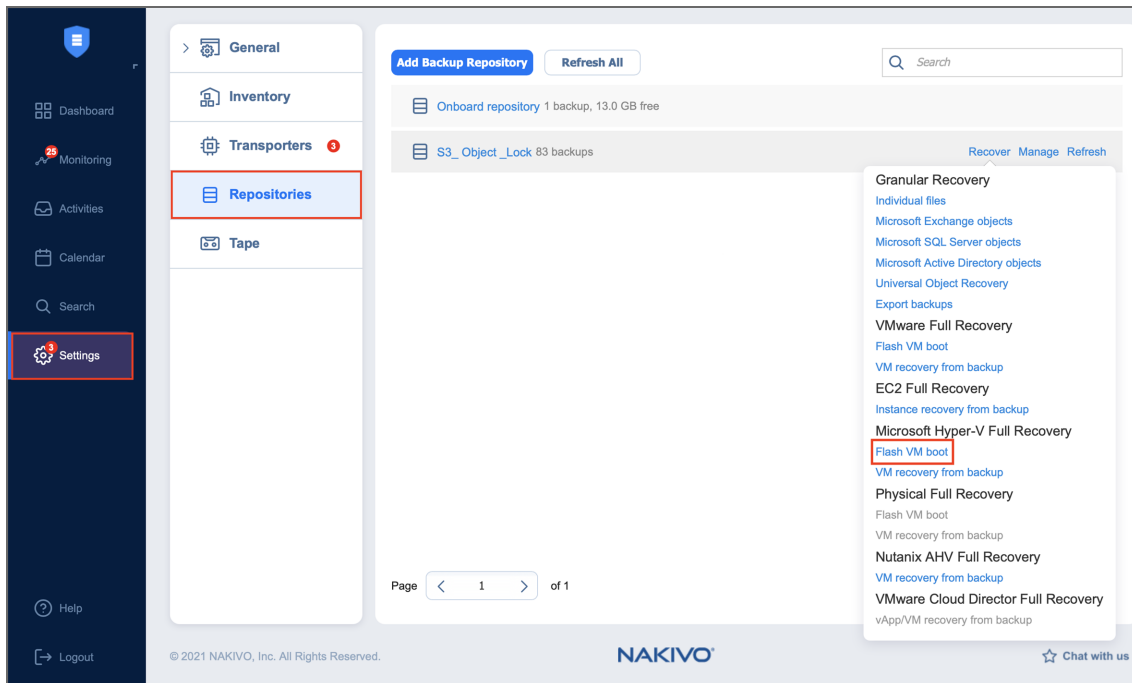
## Creating Hyper-V Flash VM Boot Jobs

To create a Hyper-V Flash VM Boot job, do one of the following:

- Open the **New Flash Boot Job Wizard** from the **Dashboard** by clicking **Recover** and then clicking **Flash VM Boot**.



- Open the Flash VM Boot wizard from the **Repositories** tab by following the steps below:
  1. Click **Settings** in the main menu of the product.
  2. Go to the **Repositories** tab.
  3. Hover over the name of the Backup Repository, click **Recover** and then click **Flash VM Boot**.



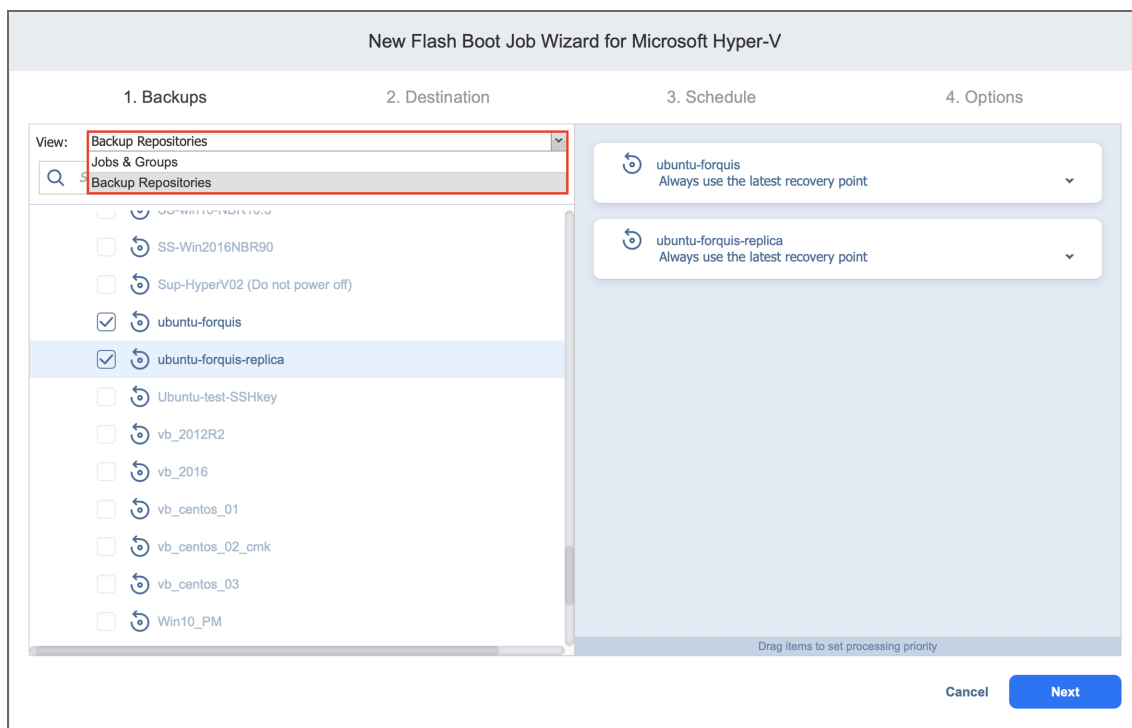
- Alternatively, the recovery can be performed by using the **search** function. The **New Flash Boot Job Wizard** opens.

- [“Hyper-V Flash Boot Job Wizard: Backups” below](#)
- [“Hyper-V Flash Boot Job Wizard: Destination” below](#)
- [“Hyper-V Flash Boot Job Wizard: Schedule” on page 865](#)
- [“Hyper-V Flash Boot Job Wizard: Options” on page 868](#)

## Hyper-V Flash Boot Job Wizard: Backups

On the **Backups** page of the wizard, proceed as follows:

1. Select one of the views:
  - **Jobs & Groups:** When selected, perform the following:
    - a. Select one or more VM backups in the left pane.
    - b. Select a recovery point for each backup in the right pane.
  - **Backup Repositories:** When selected, perform the following:
    - a. Select one or more Backup Repositories in the left pane.
    - b. Select a recovery point for each backup in the right pane



2. Select a recovery point for each backup in the right pane.
3. Click **Next** to go to the next page.

## Hyper-V Flash Boot Job Wizard: Destination

On the **Destination** page, select a destination for the recovered VMs.

## Note

Due to Microsoft iSCSI Target Server limits on Windows Server, only 256 iSCSI target instances (VM disks) can be flash-booted per iSCSI Target Server.

Proceed as described in the sections below.

- [Setting the Same Container, Path, and Network for All VMs](#)
- [Setting Different Options for Recovered VMs](#)

### Setting the Same Container, Path, and Network for All VMs

To run all VMs on the same container and path, and to connect all recovered VMs to the same network, choose a target container, path, and network from the appropriate drop-down lists.

New Flash Boot Job Wizard for Microsoft Hyper-V

1. Backups      2. Destination      3. Schedule      4. Options

Container:

Path:

Network:

[Advanced setup...](#)

Cancel      Next

## Note

If you choose the **Connect to temporary isolated network** option from the **Network** drop-down list, NAKIVO Backup & Replication will create a new Hyper-V virtual switch 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 and network adapter for each individual VM.

### 3. Click **Next**.

The screenshot shows the 'New Flash Boot Job Wizard for Microsoft Hyper-V' in the '1. Backups' step. The wizard is configured for two backup jobs: 'ubuntu-forquis' and 'ubuntu-forquis-replica'. Both jobs are set to be recovered to a new Hyper-V virtual switch named 'New Virtual Switch'. The 'Target' section for each job is highlighted with a red box, showing the container 'ServerHV2012', a new virtual machine, and the VM configuration 'C:\NakivoRecovered'. The 'Next' button is visible at the bottom right.

#### Note

If you choose the **Connect to temporary isolated network** option for a NIC, NAKIVO Backup & Replication will create a new Hyper-V virtual switch and a new network on the host where the VM will be recovered. The recovered VM will be connected to the network.

#### Hyper-V 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](#)
- [Daily Job Execution](#)
- [Monthly or Yearly Job Execution](#)
- [Periodic Job Execution](#)
- [Chained Job](#)
- [Additional Schedule](#)

#### Disabling Scheduled Job Execution

If you want to start the job manually only (without any schedule), select the **Do not schedule, run on demand** checkbox and click **Next**.



New Flash Boot Job Wizard for Microsoft Hyper-V

1. Backups      2. Destination      3. Schedule      4. Options

Do not schedule, run on demand

[Next](#)   [Cancel](#)

## Daily Job Execution

To run a job once a day, do the following:

- 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:** 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 on which the job will be started.
- If necessary, select the **Effective from** checkbox and pick the date when the schedule should come into effect.
- Click **Next**.

New Flash Boot Job Wizard for Microsoft Hyper-V

1. Backups      2. Destination      3. Schedule      4. Options

Do not schedule, run on demand

(UTC+02:00, EET) Eastern European Time

Schedule #1

**Run daily/weekly**

Starting at: 0:00    Ending: 6:00

Mon    Tue    Wed    Thu    Fri    Sat    Sun

All days   Work days   Weekends

every 1 weeks

Effective from

[Add another schedule](#)

[Show calendar](#)

[Next](#)   [Cancel](#)

## Monthly or Yearly Job Execution

To run a job monthly or yearly, do the following:

- Choose a time zone to be used for the job start and end times from the time zone drop-down list.
- Choose **Run monthly/yearly** from the schedule drop-down list.
- Specify a job start schedule in 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:** field. 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 should come into effect.
- Click **Next**.

New Flash Boot Job Wizard for Microsoft Hyper-V

1. Backups      2. Destination      3. Schedule      4. Options

Do not schedule, run on demand  
 (UTC+02:00, EET) Eastern European Time

Schedule #1  
Run monthly/yearly

Run every last Friday of every month

Starting at: 0:00  Ending: 6:00

Effective from

[Add another schedule](#)  
[Show calendar](#)

**Next**    Cancel

## Periodic Job Execution

To run a job multiple times per day, do the following:

- Choose a time zone to be used for the job start and end times from the time zone drop-down list.
- 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 on which the job will be started.
- If necessary, select the **Effective from** checkbox and pick the date when the schedule should come into effect.
- Click **Next**.

New Flash Boot Job Wizard for Microsoft Hyper-V

1. Backups      2. Destination      3. Schedule      4. Options

Do not schedule, run on demand  
 (UTC+02:00, EET) Eastern European Time

Schedule #1  
Run periodically every 30 minutes

Starting at: 0:00  Ending: 6:00

Mon  Tue  Wed  Thu  Fri  Sat  Sun  
 All days    Work days    Weekends

Effective from

[Add another schedule](#)  
[Show calendar](#)

**Next**    Cancel

## 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.
- If necessary, select the **Effective from** checkbox and pick the date when the schedule should come into effect.
- Click **Next**.

The screenshot shows the 'New Flash Boot Job Wizard for Microsoft Hyper-V' interface, specifically the '3. Schedule' step. The wizard has four steps: 1. Backups, 2. Destination, 3. Schedule, and 4. Options. The 'Schedule #1' section is active, and the 'Run after another job' option is selected in the dropdown menu, which is highlighted with a red box. Other options include 'Do not schedule, run on demand', 'After the job: Physical machine backup job-2(D)', 'Run this job: Immediately', and checkboxes for 'After successful runs', 'After failed runs', 'After stopped runs', and 'Effective from'. The 'Next' button is visible at the bottom right.

## Additional Schedule

If you need to add an additional schedule, click **Add another schedule** and configure it as described above.

## Hyper-V Flash Boot Job Wizard: Options

On the **Options** page, set up job options as described in these sections:

- [Job Options](#)
  - [Job Name](#)
  - [VM Verification](#)
- [Recovered VM Options](#)
- [Pre and Post Actions](#)
  - [Setting Up Email Notifications for the Job](#)
  - [Setting Up a Pre Job Script](#)
  - [Setting Up a Post Job Script](#)
- [Data Routing](#)
- [Completing the New Flash VM Boot Job Wizard](#)

## Job Options

In this section, specify a job name and select the VM verification type. Proceed as described below:

## Job Name

In the *Job Options* section, enter the name for the job.

## 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 48](#).

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 VM 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 (see [“Notifications & Reports” on page 294](#)) and displayed on the **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 VM 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.

After selecting **Boot verification**, specify the verification options in the dialog box that opens:

1. **Verify not more than x VMs simultaneously:** Specify the maximum number of VMs that can be started on the target container simultaneously.
2. **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.

The screenshot shows the 'New Flash Boot Job Wizard for Microsoft Hyper-V' with four steps: 1. Backups, 2. Destination, 3. Schedule, and 4. Options. The 'Options' step is active. The 'Job Options' section is highlighted with a red box and contains the following fields: 'Job name' (text input: 'Hyper-V flash VM boot job'), 'VM verification' (dropdown: 'Disabled'), 'Recovered VM Options' (dropdown: 'Append "-recovered" in the end'), 'VM MAC addresses' (dropdown: 'Do not generate new MAC addresses'), and 'VM power on' (dropdown: 'Power on recovered VMs'). The 'Pre and Post Actions' section has three checkboxes: 'Send job run reports to' (unchecked), 'Run local pre job script' (unchecked), and 'Run local post job script' (unchecked). The 'Data routing' section has a dropdown for 'Proxy transporter' set to 'Do not use proxy transporter'. At the bottom right are three buttons: 'Finish' (blue), 'Finish & Run' (grey), and 'Cancel' (grey).

## Recovered VM Options

In the *Recovered VM Options* section, you can specify names for recovered VMs so that you can distinguish between the recovered VMs and the source VMs. Also, you can set up the VM MAC addresses option and choose to either power on or not the VM after recovery. Proceed the following way:

1. In the **Recovered VM names** list, choose one of the following:
  - **Append “-recovered” in the end:** This is the default option. VM names are used for recovered VMs and “-recovered” is appended to the end of the name.
  - **Leave recovered VM names as is:** Recovered VM names are identical to the source VM names.
  - **Enter custom recovered VM names:** When selected, the **Custom Recovered VM Names** box can be opened and you can enter a name for each VM to be recovered.
2. Choose whether a new MAC address will be generated for the recovered VM. When you select **Generate new MAC address in the VM MAC addresses**, a new MAC address is generated for each recovered VM. If you choose the **Do not generate new MAC addresses** option, the recovered VMs will have the same MAC address as the source VMs.
3. In addition, you can choose whether the recovered VMs will be powered on. In the **VM power on** list, choose one of the following:
  - **Power on recovered VMs**

- **Do not power on recovered VMs**

### Pre and Post Actions

NAKIVO Backup & Replication allows you to run a script before Flash VM 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 be executed only on the machine where the Director is installed. Refer to [“Pre and Post Job Scripts” on page 68](#) for details.

## Setting up Email Notifications for the Job

1. Select the **Send job run reports to** checkbox.
2. The text box opens for editing next to the checkbox. Specify one or more email addresses in the text box. Use semicolons to separate multiple email addresses.

### Note

By default, the text box contains email recipients that are specified in the **Email notifications** section of the NAKIVO Backup & Replication.

## 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. Select the **Run local pre job script** option. In the *Settings* dialog box that opens, specify the following parameters:
  - **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, 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 will not be performed.

## Setting Up a Post Job Scripts

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. Select the **Run local post job script** option and in the **settings** dialog box that opens, specify the parameters as described for pre-job scripts.

The screenshot shows the 'New Flash Boot Job Wizard for Microsoft Hyper-V' dialog box, specifically the '4. Options' step. The 'Pre and Post Actions' section is highlighted with a red box. It contains three checkboxes: 'Send job run reports to' (unchecked), 'Run local pre job script' (unchecked), and 'Run local post job script' (checked). Other options in the wizard include 'Job name' (Hyper-V flash VM boot job), 'VM verification' (Disabled), 'Recovered VM names' (Append "-recovered" in the end), 'VM MAC addresses' (Do not generate new MAC address), 'VM power on' (Power on recovered VMs), and 'Data routing' (Proxy transporter: Do not use proxy transporter). At the bottom right, there are three buttons: 'Finish', 'Finish & Run', and 'Cancel'.

### 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 the 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.

New Flash Boot Job Wizard for Microsoft Hyper-V

1. Backups      2. Destination      3. Schedule      4. Options

---

**Job Options**

Job name:

VM verification:  ?

**Recovered VM Options**

Recovered VM names:  ?

VM MAC addresses:

VM power on:

**Pre and Post Actions**

Send job run reports to

Run local pre job script ?

Run local post job script ?

**Data routing**

Proxy transporter:  ?

---

## Completing the New Flash VM Boot Job Wizard

Click **Finish** or **Finish & Run** to complete the job creation.

### Note

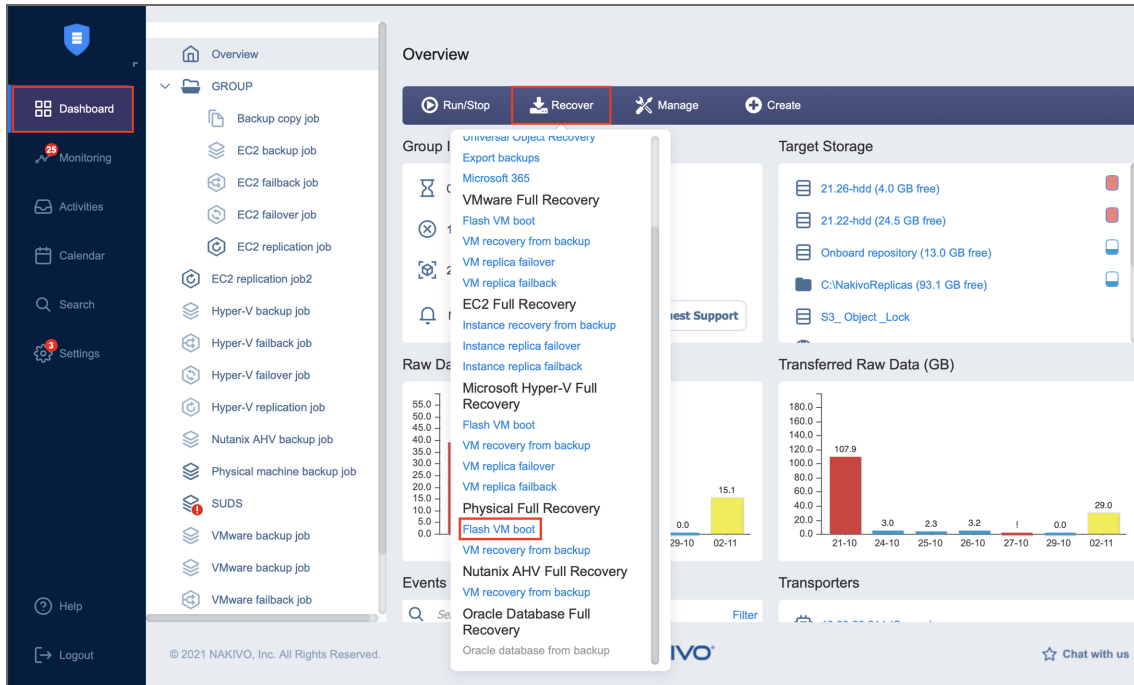
If you click **Finish & Run**, you will have to define the scope of your job. Please refer to [“Running Jobs on Demand” on page 256](#) for details.

## Creating Flash Boot Jobs for Physical Machines

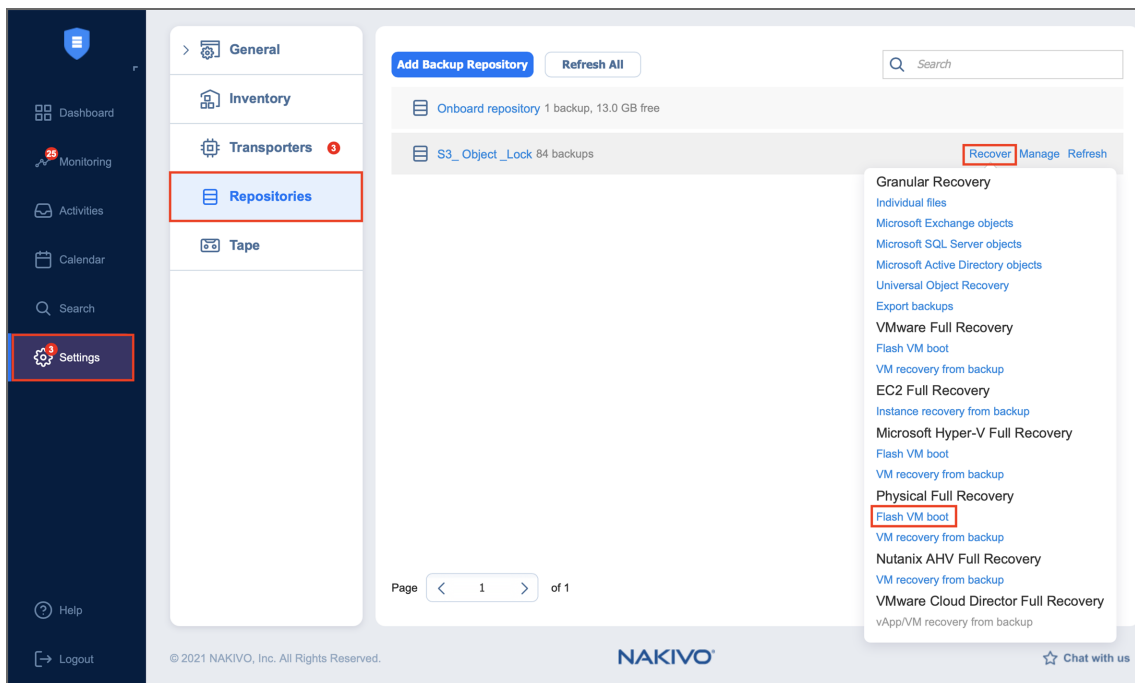
To create a Flash VM Boot job for physical machines, do one of the following:



- Create a Flash VM Boot job from the **Dashboard** by clicking **Recover** and then selecting **Flash VM Boot** for physical machines.



- Create a Flash VM Boot job from the **Repositories** tab in **Settings** by following the steps below:
  1. Go to the main menu of NAKIVO Backup & Replication and click **Settings**.
  2. Go to the **Repositories** tab and hover over a Backup Repository containing the required backup.
  3. Click **Recover** and then click **Flash VM Boot** for physical machines.



- Alternatively, the recovery can be performed by using by using the [search](#) function.

The **New Flash Boot Job Wizard for Physical Machines** opens.

### Note

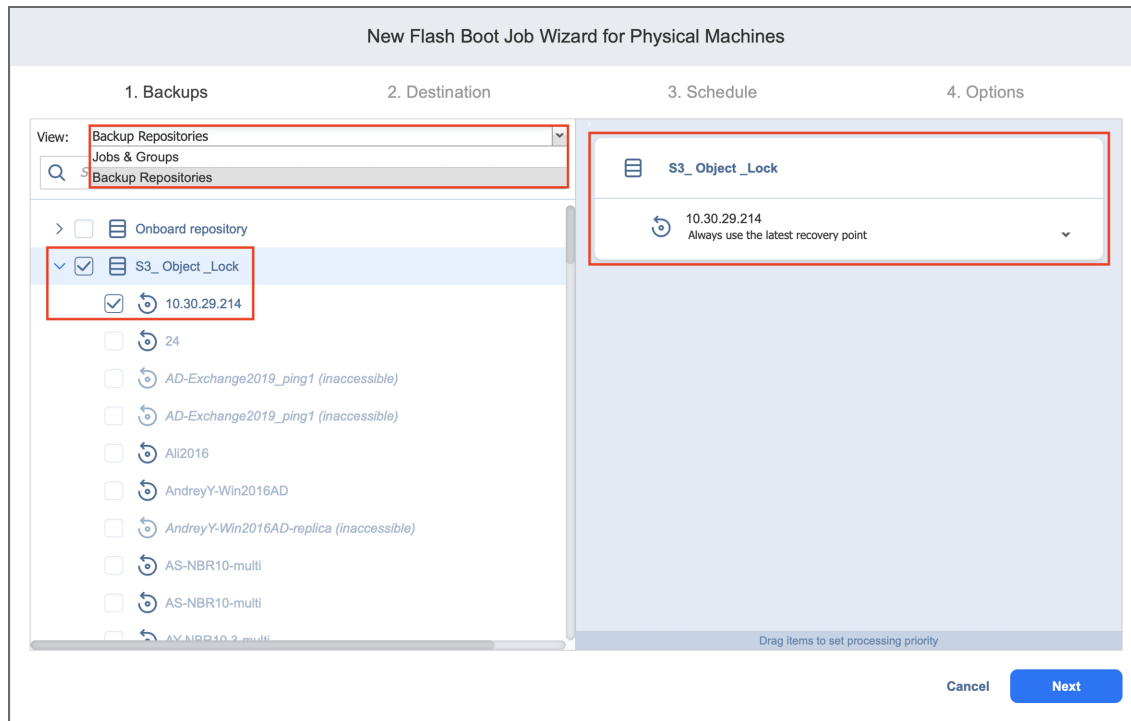
Free ESXi is not supported for physical to virtual recovery.

- [“Flash Boot Job Wizard for Physical Machines: Backups” below](#)
- [“Flash Boot Job Wizard for Physical Machines: Destination” below](#)
- [“Flash Boot Job Wizard for Physical Machines: Schedule” on page 878](#)
- [“Flash Boot Job Wizard for Physical Machines: Options” on page 882](#)

## Flash Boot Job Wizard for Physical Machines: Backups

On the **Backups** page of the wizard:

1. In the left pane, select one or more physical machine backups using one of the following views:
  - **Jobs & Groups**
  - **Backup Repositories**
2. In the left pane and then select a recovery point for each backup in the right pane.
3. Click **Next** to go to the next page of the wizard.



## Flash Boot Job Wizard for Physical Machines: Destination

On the **Destination** page, select a location for the recovered physical machines.

- [“Setting the Same Host, Datastore, and Network for All VMs” below](#)
- [“Setting the Default Destination for Recovered Machines” below](#)
- [“Setting Different Options for VMs” on the next page](#)

### Setting the Same Host, Datastore, and Network for All VMs

To run all machines on the same host, container, and datastore, and to connect all VMs to the same network, follow the steps below:

1. Choose a cluster, host, or a 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.
4. Click **Next**.

The screenshot shows a wizard window titled "New Flash Boot Job Wizard for Physical Machines". It has four tabs: "1. Backups", "2. Destination", "3. Schedule", and "4. Options". The "1. Backups" tab is active. It contains four dropdown menus: "Container" with value "10.30.21.22", "Datastore" with value "21.22-hdd", "Network" with value "10.30.21.0", and "VM folder" with value "Select target VM folder (optional)". Below these is a link "Advanced setup...". At the bottom right are "Cancel" and "Next" buttons.

### Note

By default, NAKIVO Backup & Replication sets the same amount of resources for the created VMs as that set for the source physical machine, including the same number of CPUs, cores per socket, and amount of RAM.

### Setting the Default Destination for Recovered Machines

If you have chosen a Backup Repository or a folder as a source for your recovery job on the **Backups** page, you can set the default container, datastore, and VM folder for the recovered machines. 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**.

The screenshot shows the 'New Flash Boot Job Wizard for Physical Machines' interface. The wizard is currently on the '1. Backups' step. The configuration options are as follows:

- Container: 10.30.21.22
- Datastore: 21.22-hdd
- Network: 10.30.21.0
- VM folder: Select target VM folder (optional)

Below these options is a section titled 'S3\_Object\_Lock' with a 'Click to collapse' link. This section contains the following options, which are highlighted with a red box in the image:

- Default container: 10.30.21.22
- Default datastore: 21.22-hdd
- Default network: 10.30.21.0
- Default VM folder: Select default VM folder (optional)

At the bottom right of the wizard, there are 'Cancel' and 'Next' buttons.

## Setting Different Options for VMs

You can customize VM options for every machine. To do this, follow the steps below:

1. Click **Advanced setup**.
2. Hover over a VM and click **Click to expand**. Choose a target container, virtual machine, and VM folder in the corresponding boxes.
3. Set the number of virtual CPUs, cores per socket, and the amount of RAM in the corresponding boxes.

### Note

The number of virtual CPUs and cores per socket and the amount of RAM cannot exceed the maximum value available on the destination host.

4. Select a target datastore for each disk and for each disk and for each VM file from the drop-down list.
5. Keep the source disk configuration by selecting **Keep source configuration**, or choose one of the following disk controller types:
  - **SCSI LSI Logic SAS**
  - **SCSI LSI Logic Parallel**
  - **SCSI VMware Paravirtual**
  - **SCSI BusLogic Parallel**

- IDE
- SATA
- NVMe

6. Select network adapters from the **Network adapters** drop-down list

7. Click **Next**.

The screenshot displays the 'New Flash Boot Job Wizard for Physical Machines' in the '1. Backups' step. The interface is divided into four sections: 1. Backups, 2. Destination, 3. Schedule, and 4. Options. The '1. Backups' section is active and contains two main areas: 'Source' and 'Target'. The 'Source' area shows a backup location of 'S3\_Object\_Lock' and physical server resources: 2 CPUs, 2 cores per socket, and 128.0 GB RAM. It also lists disks: \\.\PHYSICALDRIVE2: 200.0 GB with a SCSI disk controller. Network adapters include Intel(R) I350 Gigabit N... and TAP-Windows Adapter... The 'Target' area shows a container '10.30.21.22' where a new VM will be created. It specifies 2 virtual CPUs, 2 cores per socket, and 128 GB RAM. The target disk is \\.\PHYSICALDRIVE2: 21.22-hdd with a 'Keep source configuration' controller. Network adapters for the target include Intel(R) I350 Gigabit N... and TAP-Windows Adapter... The wizard has 'Cancel' and 'Next' buttons at the bottom right.

## Flash Boot Job Wizard for Physical Machines: 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.

- [“Disabling Scheduled Job Execution” on the next page](#)
- [“Daily or Weekly Run” on the next page](#)
- [“Monthly or Yearly Run” on the next page](#)
- [“Periodic Run” on page 880](#)
- [“Chained Job” on page 881](#)
- [“Adding Another Schedule” on page 882](#)

## Disabling Scheduled Job Execution

If you want to start the job manually (without scheduling), select the **Do not schedule, run on demand** checkbox.

New Flash Boot Job Wizard for Physical Machines

1. Backups 2. Destination 3. Schedule 4. Options

Do not schedule, run on demand

Next Cancel

## Daily or Weekly Run

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 to 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 on which the job will be started.
- Specify what weeks you want the job to be executed.
- If necessary, select the **Effective from** checkbox and set the date when the schedule comes into effect.

New Flash Boot Job Wizard for Physical Machines

1. Backups 2. Destination 3. Schedule 4. Options

Do not schedule, run on demand  
(UTC+02:00, EET) Eastern European Time

Schedule #1  
 Run daily/weekly

Starting at: 0:00 Ending: 6:00

Mon  Tue  Wed  Thu  Fri  Sat  Sun  
All days Work days Weekends

every 1 weeks

Effective from  
[Add another schedule](#)  
[Show calendar](#)

Next Cancel

## Monthly or Yearly Run

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 set the date when the schedule comes into effect.

The screenshot shows the 'New Flash Boot Job Wizard for Physical Machines' interface, specifically the '3. Schedule' step. The wizard has four steps: 1. Backups, 2. Destination, 3. Schedule, and 4. Options. In the '3. Schedule' step, there are several options and fields:

- A checkbox for 'Do not schedule, run on demand' is unchecked.
- A time zone dropdown menu is set to '(UTC+02:00, EET) Eastern European Time'.
- Under 'Schedule #1', a dropdown menu is set to 'Run monthly/yearly', which is highlighted with a red rectangular box.
- The 'Run every' section is set to 'last' of every 'month' on 'Friday'.
- The 'Starting at' field is set to '0:00' and the 'Ending' field is set to '6:00'.
- The 'Effective from' checkbox is unchecked.
- There are links for 'Add another schedule' and 'Show calendar'.
- At the bottom right, there are 'Next' and 'Cancel' buttons.

## Periodic Run

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 how often the job should be executed in **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.
- Select the days of the week on which the job will be started.

- If necessary, select the **Effective from** checkbox and set the date when the schedule comes into effect.

New Flash Boot Job Wizard for Physical Machines

1. Backups      2. Destination      3. Schedule      4. Options

Do not schedule, run on demand  
 (UTC+02:00, EET) Eastern European Time

Schedule #1

**Run periodically** every 30 minutes

Starting at: 0:00 Ending: 6:00

Mon  Tue  Wed  Thu  Fri  Sat  Sun  
 All days Work days Weekends

Effective from

[Add another schedule](#)  
[Show calendar](#)

**Next** Cancel

### 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.
- **Effective from:** When selected, the schedule will come into effect on the set date.

New Flash Boot Job Wizard for Physical Machines

1. Backups      2. Destination      3. Schedule      4. Options

Do not schedule, run on demand  
 (UTC+02:00, EET) Eastern European Time

Schedule #1

**Run after another job**

After the job: AD+Exchange

Run this job: Immediately

After successful runs  After failed runs  After stopped runs

Effective from

[Add another schedule](#)  
[Show calendar](#)

**Next** Cancel



## Adding Another Schedule

If you want to have more than one schedule for your job, click **Add another schedule** and set it as described above.

## Flash Boot Job Wizard for Physical Machines: Options

On the **Options** page of the wizard, specify a job name, set up recovered VM options, and choose data routing.

- [“Job Options ” below](#)
  - [“Recovered VM Options” on the next page](#)
  - [“Specifying VM Names” on the next page](#)
  - [“Generating VM MAC Addresses” on the next page](#)
  - [“Powering Recovered VMs” on the next page](#)
- [“Pre and Post Actions” on page 884](#)
  - [“Setting Up Email Notifications ” on page 884](#)
  - [“Setting Up a Pre Job Script” on page 884](#)
  - [“Setting Up a Post Job Script” on page 885](#)
- [“Data Routing” on page 886](#)
- [“Completing Flash Boot Job Wizard for Physical Machines” on page 886](#)

## Job Options

In this section, specify a job name.

New Flash Boot Job Wizard for Physical Machines

1. Backups      2. Destination      3. Schedule      4. Options

**Job Options**

Job name:

Recovered VM Options

Recovered VM names:  ?

VM MAC addresses:

VM power on:

Pre and Post Actions

Send job run reports to  ?

Run local pre job script ?

Run local post job script ?

Data routing

Proxy transporter:  ?

## 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 that you can distinguish between recovered VMs and the source physical machines. By default, the text “- recovered” is appended to the end of the recovered VM name.

To change VM names, choose one of the following options in the *Recovered VM Options* section:

- **Append “-recovered” in the end:** Source machine 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 machine names.
- **Enter custom recovered VM names:** Allows you to enter custom names for recovered VMs.

## Generating VM MAC Addresses

In the *Recovered VM Options* section, choose one of the following options in relation to recovered VM MAC addresses:

- **Generate new MAC addresses:** A new MAC address is generated for each recovered VM.
- **Do not generate new MAC addresses:** The recovered VMs have the same MAC address as the source machines.

## Powering Recovered VMs

To power on the recovered VMs, choose the **VM power on** option.

New Flash Boot Job Wizard for Physical Machines

1. Backups 2. Destination 3. Schedule 4. Options

Job Options

Job name: Physical machine flash boot job

**Recovered VM Options**

Recovered VM names: Append "-recovered" in the end ?

VM MAC addresses: Generate new MAC addresses

VM power on: Power on recovered VMs

Pre and Post Actions

Send job run reports to ?

Run local pre job script ?

Run local post job script ?

Data routing

Proxy transporter: Do not use proxy transporter ?

Finish Finish & Run Cancel

## Pre and Post Actions

NAKIVO Backup & Replication allows you to run a [script](#) before Flash VM 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. Also, you can set up email notifications for the job. Refer to [“Notifications & Reports” on page 294](#) for details.

## Setting Up 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.

## Setting Up a 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. 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:
    - **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:

## Setting Up a 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. 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 has failed, the job status will be set to “failed” even if VM replication has been successful.

The screenshot shows the 'New Flash Boot Job Wizard for Physical Machines' dialog box, step 4: Options. The 'Pre and Post Actions' section is highlighted with a red box. The 'Run local post job script' option is selected. The 'Data routing' section shows 'Do not use proxy transporter' selected. The 'Job Options' section shows 'Physical machine flash boot job' in the 'Job name' field. The 'Recovered VM Options' section shows 'Append "-recovered" in the end' for 'Recovered VM names', 'Generate new MAC addresses' for 'VM MAC addresses', and 'Power on recovered VMs' for 'VM power on'. The 'Data routing' section shows 'Do not use proxy transporter' for 'Proxy transporter'. The 'Finish', 'Finish & Run', and 'Cancel' buttons are visible at the bottom right.

4.

## Data Routing

In case the Transporter assigned to a Backup Repository cannot use iSCSI port 3260 because the port is occupied by other services, you can set data routing: a proxy Transporter can be used to forward the 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.

The screenshot shows the 'New Flash Boot Job Wizard for Physical Machines' at step 4, 'Options'. The wizard has four steps: 1. Backups, 2. Destination, 3. Schedule, and 4. Options. The 'Data routing' section is highlighted with a red box. It contains a 'Proxy transporter' dropdown menu set to 'Do not use proxy transporter'. Other options include 'Job name' (Physical machine flash boot job), 'Recovered VM names' (Append "-recovered" in the end), 'VM MAC addresses' (Generate new MAC addresses), 'VM power on' (Power on recovered VMs), and 'Pre and Post Actions' (Send job run reports to, Run local pre job script, Run local post job script).

## Completing Flash Boot Job Wizard for Physical Machines

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 256](#) for details.

## Migrating Recovered VMs Using Flash Boot

Using Flash VM Boot, you can migrate the recovered VMs to another location. To do this, follow the instructions below:

### Migrating Recovered VMware VMs Using Flash VM 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 VM 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 VM Boot job.

6. Select the datastore where you want to store the VM.
7. Review the page and click **Finish**.

## Note

After migrating the booted VMs, you can click the **Discard VMs** button in the Flash VM 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 VM 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, power off the VMs created by Flash VM Boot in the VMware vSphere client and run the VM replication job one last time.
6. Discard the VMs created by Flash VM Boot.
7. Power on the VMs created by the replication job.

## Migrating Recovered Hyper-V VMs Using Flash VM Boot to Production

Follow the steps below to migrate a Hyper-V VM recovered using Flash Boot to production:

1. Open the Hyper-V Manager and shut down the recovered VM.
2. Go to the VM settings and open the Hard Drive properties.
3. Make a note of the name of the physical disk from which the VM was booted. This is typically Disk N where N is the index number of your physical disk.
4. In the *Media section* of the Hard Drive properties, click **Virtual hard disk** and then click **New**.
5. The **New Virtual Hard Disk** wizard opens. Proceed as follows:
  - a. On the **Choose Disk Format** step of the wizard, choose VHDX as the format of your virtual hard disk. This format supports virtual disks up to 64 TB. Click Next.

### Important

On Hyper-V 2019, adding a new virtual hard disk may fail with an error message "Failed to convert the virtual disk".

- b. On the **Choose Disk Type** step of the wizard, choose **Dynamically expanding**. This initially creates a virtual hard disk of a small size that changes as data is added. Click **Next**.
- c. On the **Specify Name and Location** step of the wizard, specify a name and a location for the virtual hard disk file and click **Next**.

- d. On the **Configure Disk** step of the wizard, click **Copy the contents** of the specified physical disk and choose the disk you previously noted from the list of physical hard disks. Click **Next**.
- e. On the **Summary** step of the wizard, review the information and click **Finish**.
- f. Wait till the wizard creates the new virtual hard disk and closes. Click **Finish**.

Now the VM is ready for use in production.

### Note

If you apply the Discard VMs command to the VM recovered via flash boot, the corresponding job will be stopped and the VM will NOT be deleted.

Alternatively, you can move your VM to another computer. Please follow the steps below:

1. In the Hyper-V Manager, right-click on your VM and choose **Move** in the context menu.
2. The **Move VM Wizard** opens. Proceed as follows:
  - a. On the **Choose Move Type** step, choose **Move virtual machine** and click **Next**.
  - b. On the **Specify Destinations** page, choose the name of the computer to which to move the VM and click **Next**.
  - c. On the **Choose Move Option** step, choose **Move the virtual machine's data to a single location** and click **Next**.
  - d. On the **Virtual Machine** step, specify the location of the VM on the destination computer. Then click **Next**.
  - e. On the **Summary** step, check the configuration of your moving operation and click **Finish**.

When the VM has moved to the destination computer successfully, you can discard the VM with NAKIVO Backup & Replication.

## 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.

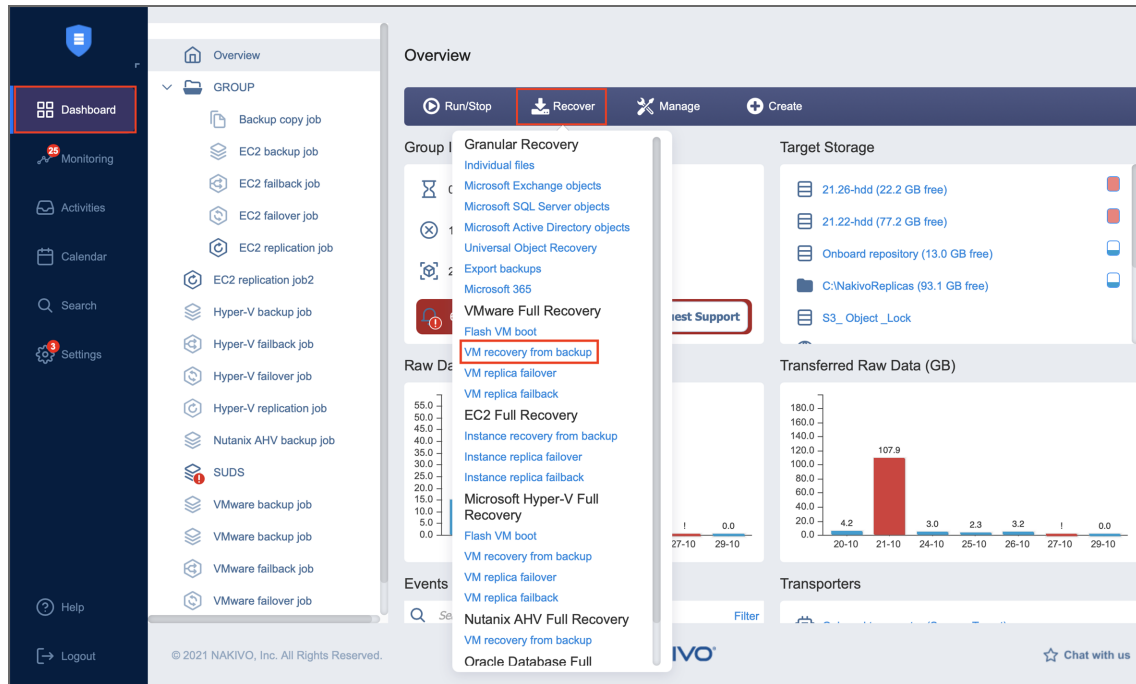
Refer to the following topics for more information:

- [“Starting VMware VM Recovery” below](#)
- [“Recovery Job Wizard for VMware: Backups” on page 890](#)
- [“Recovery Job Wizard for VMware: Destination” on page 890](#)
- [“Recovery Job Wizard for VMware: Options” on page 894](#)

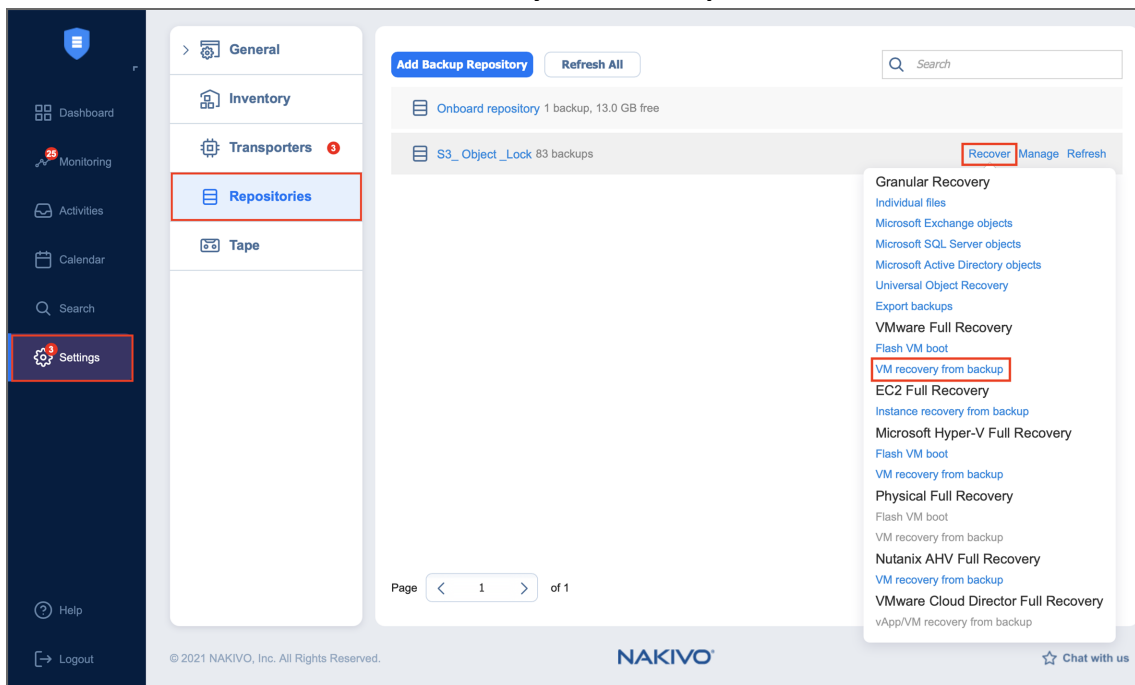
## Starting VMware VM Recovery

To recover entire VMware VMs from backups, do one of the following:

- Start recovery from the **Dashboard** by clicking **Recover** and then clicking **VM recovery from backup**



- Open the **New Recovery Job Wizard** from the **Repositories** tab by following the steps below:
  - Go to the main menu of NAKIVO Backup & Replication and click **Settings**.
  - Go to the **Repositories** tab and hover over a Backup Repository containing the necessary backup.
  - Click **Recover** and then click **VM recovery from backup**.



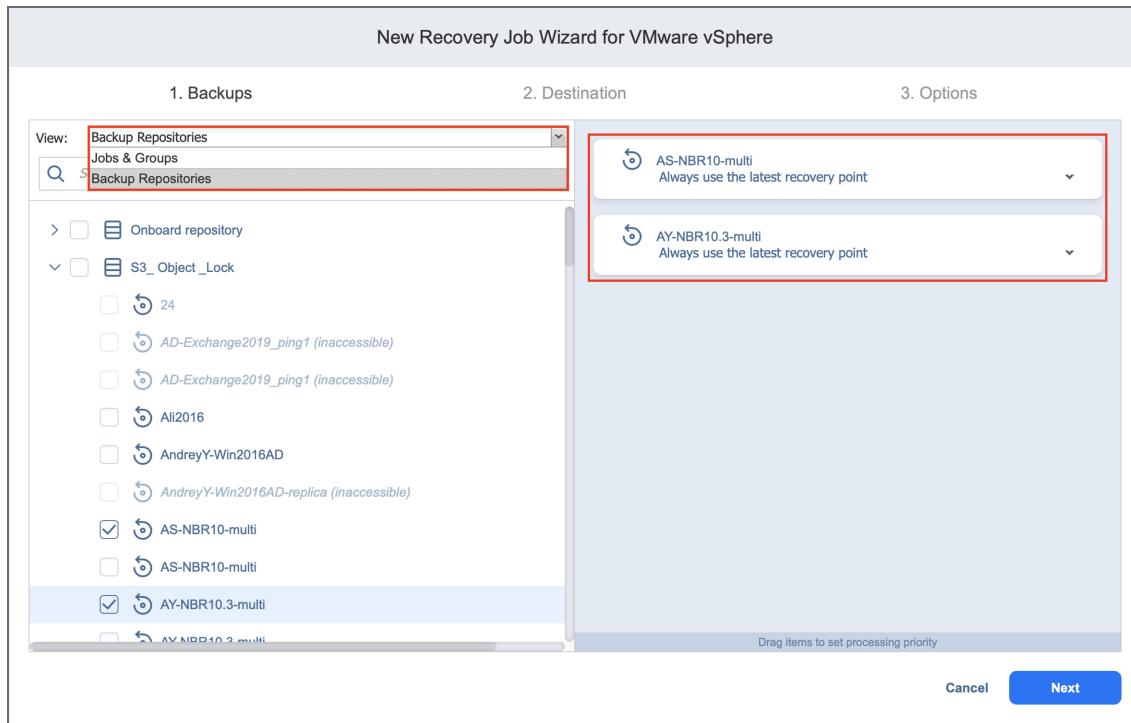
- Alternatively, the recovery can be performed by using the [search](#) function. The **New Recovery Job Wizard for VMware vSphere** opens.



## Recovery Job Wizard for VMware: Backups

On the **Backups** page of the wizard, proceed as follows:

1. Select one of the views:
  - **Jobs & Groups:** When selected, perform the following:
    - a. Select one or more VM or VM template backups in the left pane.
    - b. Select a recovery point for each backup in the right pane.
  - **Backup Repositories:** When selected, perform the following:
    - a. Select one or more Backup Repositories in the left pane.
    - b. Select a recovery point for each backup in the right pane.



2. Click **Next** to go to the next page.

## 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” below](#)
- [“Setting Original Location for All Recovered VMs” on the next page](#)
- [“Setting the Default Destination for Recovered VMs” on page 892](#)
- [“Setting Different Options for Recovered VMs” on page 893](#)

### 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. Choose a cluster, host, or resource pool from the **Container** drop-down list.
2. Choose a datastore from the **Datastore** drop-down list.
3. Choose a network from the **Network** drop-down list.
4. Optionally, you can choose a folder if there is one in the container.

New Recovery Job Wizard for VMware vSphere

1. Backups      2. Destination      3. Options

Container: 10.30.21.22

Datastore: 21.22-hdd

Network: 10.30.21.0

VM folder: Select target VM folder (optional)

[Advanced setup...](#)

Cancel    Next

## 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**.

New Recovery Job Wizard for VMware vSphere

1. Backups                      2. Destination                      3. Options

---

Container: Original container ▼

Datastore: Original datastore ▼

Network: Original network ▼

VM folder: Original VM folder ▼

[Advanced setup...](#)

---

[Cancel](#) [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 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 options** 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**.

- Optionally, you can also choose a **Default VM folder**.

The screenshot shows the 'New Recovery Job Wizard for VMware vSphere' interface. The wizard is currently on the '1. Backups' step. The '2. Destination' and '3. Options' steps are visible but not active. The 'Container' is set to '10.30.21.22', 'Datastore' to '21.22-hdd', and 'Network' to '10.30.21.0'. The 'VM folder' dropdown is set to 'Select target VM folder (optional)'. Below this, there is a section titled 'S3\_Object\_Lock' with a 'Click to collapse' link. This section contains four dropdown menus: 'Default container: 10.30.21.22', 'Default datastore: 21.22-hdd', 'Default network: 10.30.21.0', and 'Default VM folder: Select default VM folder (optional)'. The 'Default VM folder' dropdown is highlighted with a red box. Below these dropdowns, there are two buttons: 'AS-NBR10-multi' and 'AY-NBR10.3-multi'. At the bottom right, there are 'Cancel' and 'Next' buttons.

## Setting Different Options for Recovered VMs

To specify different options for recovered VMs or VM templates, follow the steps below:

- Click **Advanced options**.
- Click on the backup to expand its recovery options.
- Choose a target location, target datastore, and target network for each VM.
- To get additional information about the source and target VMs within a backup, click on its name.

5. Click **Next** to go to the next page of the wizard.

The screenshot displays the 'New Recovery Job Wizard for VMware vSphere' interface. The wizard is currently on the '2. Destination' step. It shows two job configurations: 'AS-NBR10-multi' and 'AY-NBR10.3-multi'. Each configuration has a 'Source' section and a 'Target' section. The 'Source' section includes 'VM location', 'VM resources', 'Disks', and 'Network adapters'. The 'Target' section includes 'Container', 'Virtual Machine', 'VM folder', 'Disks', and 'Network adapters'. The 'Next' button is highlighted in blue.

## Recovery Job Wizard for VMware: Options

On the **Options** page of the wizard, set the recovery job options.

- [“Job Options” below](#)
- [“Recovery Options” on the next page](#)
- [“Pre and Post Actions” on page 896](#)
  - [“Setting a Pre-Job Script” on page 896](#)
  - [“Setting a Post-Job Script” on page 897](#)
  - [“Email Notifications” on page 898](#)
- [“Data Transfer” on page 898](#)
  - [“Transport Mode” on page 899](#)
  - Transporter Pool
  - [“Transporters” on page 899](#)
  - [“Transporter Load” on page 900](#)
  - [“Bandwidth Throttling” on page 900](#)
  - Multi-Channel Processing
  - Bottleneck Detection
- [“Completing the New Recovery Job Wizard for VMware” on page 901](#)

### Job Options

Specify the general options as follows:

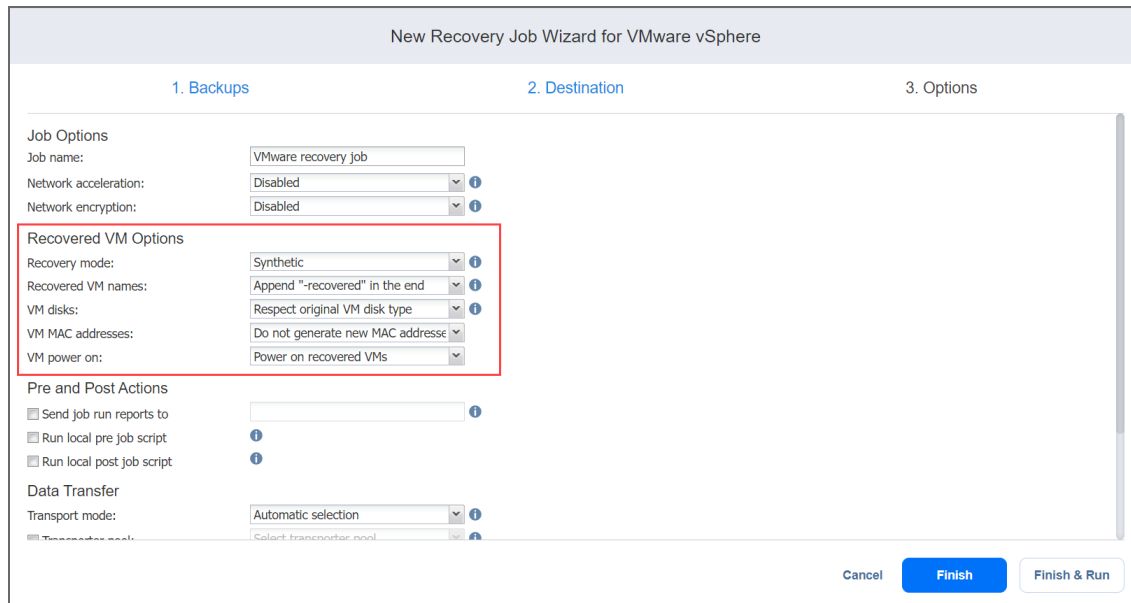
1. **Job name:** Specify a name for the recovery job.
2. **Network acceleration:** With [network acceleration](#) enabled, NAKIVO Backup & Replication uses compression and traffic reduction techniques to speed up data transfer. Enable this option if you plan to recover VMs over WAN or slow LAN links.
3. **Encryption:** With [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.

## Recovery 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 "-recovered" 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.

New Recovery Job Wizard for VMware vSphere

1. Backups 2. Destination 3. Options

Job Options

Job name: VMware recovery job

Network acceleration: Disabled

Network encryption: Disabled

Recovered VM Options

Recovery mode: Synthetic

Recovered VM names: Append "-recovered" in the end

VM disks: Respect original VM disk type

VM MAC addresses: Do not generate new MAC addresses

VM power on: Power on recovered VMs

Pre and Post Actions

Send job run reports to VMware recovery job

Run local pre job script

Run local post job script

Data Transfer

Transport mode: Automatic selection

Transporter pool: Select transporter pool

Cancel Finish Finish & Run

## 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](#).

New Recovery Job Wizard for VMware vSphere

1. Backups
2. Destination
3. Options

---

Recovery mode: Synthetic ⓘ

Recovered VM names: Append "-recovered" in the end ⓘ

VM disks: Respect original VM disk type ⓘ

VM MAC addresses: Do not generate new MAC address

VM power on: Power on recovered VMs

Pre and Post Actions

Send job run reports to  ⓘ

Run local pre job script ⓘ

Run local post job script ⓘ

Data Transfer

Transport mode: Automatic selection ⓘ

Transporter pool: Select transporter pool ⓘ

Transporters: Automatic selection ⓘ

Limit transporter load to 3 concurrent tasks ⓘ

Bandwidth throttling: Disabled ⓘ

Bottleneck detection ⓘ

Use multi-channel processing 32 channels per disk ⓘ

Cancel
Finish
Finish & Run

## Transport Mode

NAKIVO Backup & Replication provides the following transport modes for writing VM data:

- **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.
- **Automatic:** When this option is chosen, Hot Add mode is used where possible. If the product cannot use Hot Add, LAN mode is used.

## 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 284](#) 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.

## 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 256](#) 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.

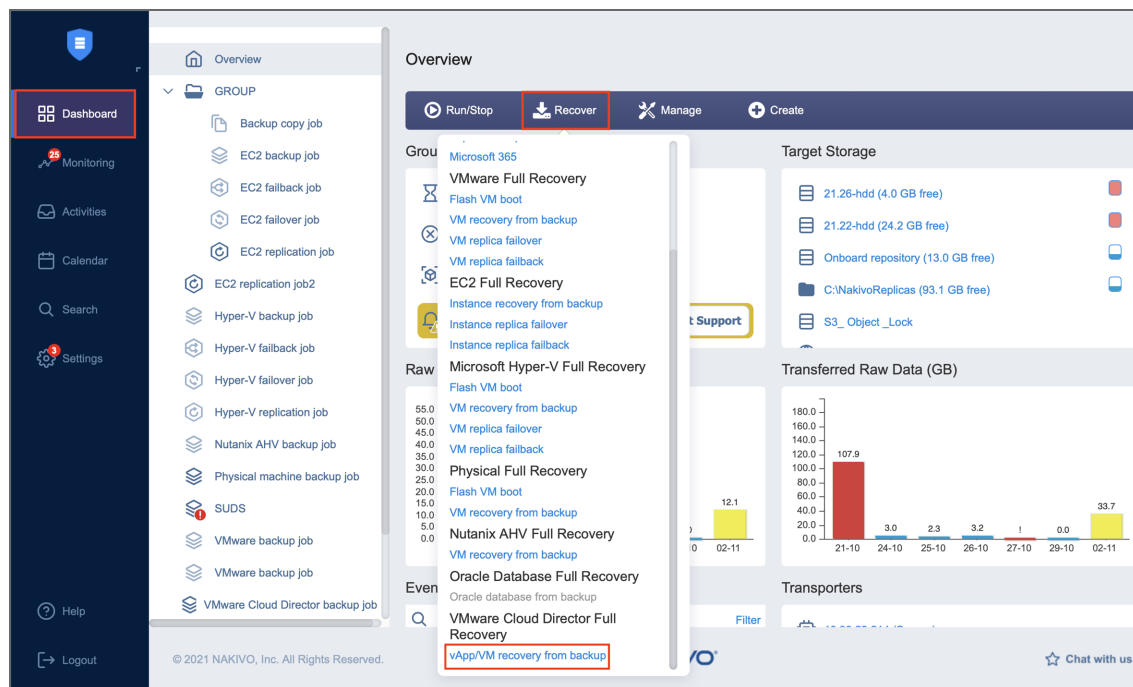
Refer to the following topics for more information:

- [“Starting VMware Cloud Director Recovery” below](#)
- [“Recovery Job Wizard for VMware Cloud Director: Object” on page 905](#)
- [“Recovery Job Wizard for VMware Cloud Director: Backups” on the next page](#)
- [“Recovery Job Wizard for VMware Cloud Director: Destination” on page 904](#)
- [“Recovery Job Wizard for VMware Cloud Director: Options” on page 906](#)

## Starting VMware Cloud Director Recovery

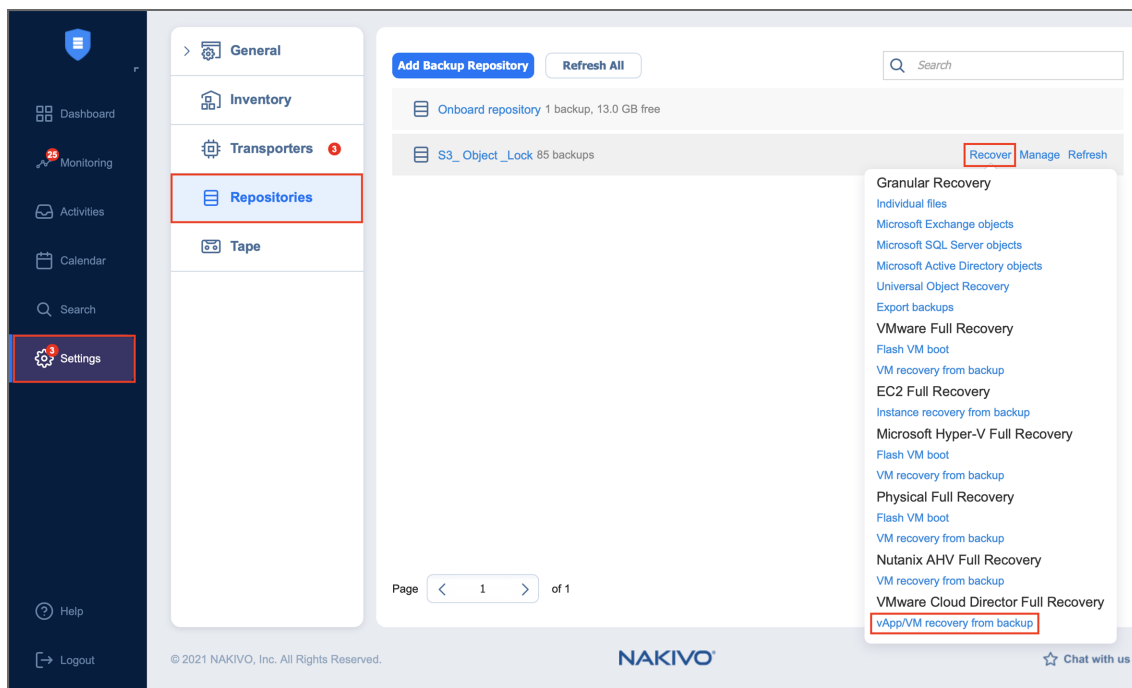
To recover all VMware Cloud Director objects from backups, do one of the following:

- Start recovery from the Dashboard by clicking **Recover** and then clicking **vApp/VM recovery from backup**.



- Open the **New Recovery Job Wizard for VMware Cloud Director** from the **Repositories** tab by following the steps below:
  1. Go to the main menu of NAKIVO Backup & Replication and click **Settings**.
  2. Go to the **Repositories** tab and hover over a Backup Repository containing the necessary backup.
  3. Click **Recover**.

4. Select **vApp/VM recovery from backup**. The **New Recovery Job Wizard for VMware Cloud Director** opens.



- Alternatively, the recovery can be performed by using by using the [search](#) function.

## Recovery Job Wizard for VMware Cloud Director: Backups

On the **Backups** page of the wizard, 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 one of the views:
  - **Jobs & Groups:** When selected, perform the following:
    - a. Select one or more VMware Cloud Director object backups in the left pane.
    - b. Select a recovery point for each backup in the right pane.

### Note

The vApps cannot be selected if Virtual machine option was selected on the [Recovery Job Wizard for VMware Cloud Director: Object](#) page.

- **Backup Repositories:** When selected, perform the following:
  - a. Select one or more Backup Repositories in the left pane.
  - b. Select a recovery point for each backup in the right pane.

3. Click **Next** to go to the next page.



## 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.
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.

New Recovery Job Wizard for VMware Cloud Director

1. Object      2. Backups      3. Destination      4. Options

Container: Test

vApp: LM\_app

Storage Policy: \*

Network: test-nw-01

LM\_wm Click to collapse

Source	Target
VM location: S3_Object_Lock	Container: Test
VM resources: 2 CPU, 4.0 GB RAM	Virtual Machine: <i>New VM will be created</i>
	vApp: LM_app
Disks	Storage Policies
Hard disk 1: S3_Object_Lock (50.0 GB)	Storage Policy: *
Network adapters	vApp Networks
Network adapter 1	Network adapter 1: test-nw-01

Cancel      **Next**

## 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.



New Recovery Job Wizard for VMware Cloud Director

1. Object                      2. Backups                      3. Destination                      4. Options

---

**i** Select VMware Cloud Director object type you want to restore

vApp  
Restores the entire VMware Cloud Director vApp to its original or a custom location. You can restore vApp to its source location or a different VMware Cloud Director deployment.

Virtual Machine  
Restores specific virtual machines to the existing VMware Cloud Director vApp. You can restore a virtual machine to its original or custom vApp in any VMware Cloud Director deployment.

---

## 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 Objects Options](#)
- [Pre and Post Job Actions](#)
- [Data Transfer](#)

### Job Options

Specify the general options as follows:

- **Job name:** Specify a name for the recovery job.
- **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.
- **Encryption:** With [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.

The screenshot shows the 'New Recovery Job Wizard for VMware Cloud Director' interface. The wizard is in the 'Options' step, with four tabs: '1. Object', '2. Backups', '3. Destination', and '4. Options'. The 'Job Options' section is highlighted with a red box. The 'Job name' is 'VMware Cloud Director recovery job'. 'Network acceleration' and 'Network encryption' are both set to 'Disabled'. Under 'Recovered Object Options', 'Recovery mode' is 'Synthetic', 'Recovered VM names' is 'Append "-recovered" in the end', 'VM disks' is 'Respect original VM disk type', 'VM MAC addresses' is 'Do not generate new MAC addresses', and 'VM power on' is 'Power on recovered VMs'. Under 'Pre and Post Actions', 'Send job run reports to' is empty, and 'Run local pre job script' and 'Run local post job script' are unchecked. Under 'Data Transfer', 'Transport mode' is 'Automatic selection', 'Transporter pool' is 'Select transporter pool', 'Transporters' is 'Automatic selection', 'Limit transporter load to' is '3 concurrent tasks', and 'Bandwidth throttling' is 'Disabled'. At the bottom right, there are 'Cancel', 'Finish', and 'Finish & Run' buttons.

## 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 "-recovered" 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 68](#).

### 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.

The screenshot shows the 'New Recovery Job Wizard for VMware Cloud Director' interface. The wizard is divided into four steps: 1. Object, 2. Backups, 3. Destination, and 4. Options. The 'Options' step is active, showing various configuration options. The 'Pre and Post Actions' section is highlighted with a red box, containing three checkboxes: 'Send job run reports to' (checked), 'Run local pre job script', and 'Run local post job script'. Below this is a text input field for email addresses. Other sections include 'Job Options', 'Recovered Object Options', and 'Data Transfer'.

## 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 [“Notifications & Reports” on page 294](#).

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:

- **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.
- **Automatic:** When this option is chosen, Hot Add mode is used where possible. If the product cannot use Hot Add, LAN mode is used.

### 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 284](#).
  - 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.

The screenshot shows the 'New Recovery Job Wizard for VMware Cloud Director' dialog box, specifically the '4. Options' step. The 'Data Transfer' section is highlighted with a red box. It contains the following settings:

- Transport mode: Automatic selection
- Transporter pool: Select transporter pool
- Transporters: Manual - configured for all VMs
- Target transporter: Automatic
- Replacement transporter: Automatic
- Limit transporter load to: 3 concurrent tasks
- Bandwidth throttling: Disabled

At the bottom right of the dialog, there are three buttons: 'Cancel', 'Finish', and 'Finish & Run'.

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. For details, refer to [“Running Jobs on Demand” on page 256](#).

## Hyper-V VM Recovery

With NAKIVO Backup & Replication, you can recover full VMs from backups. 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.

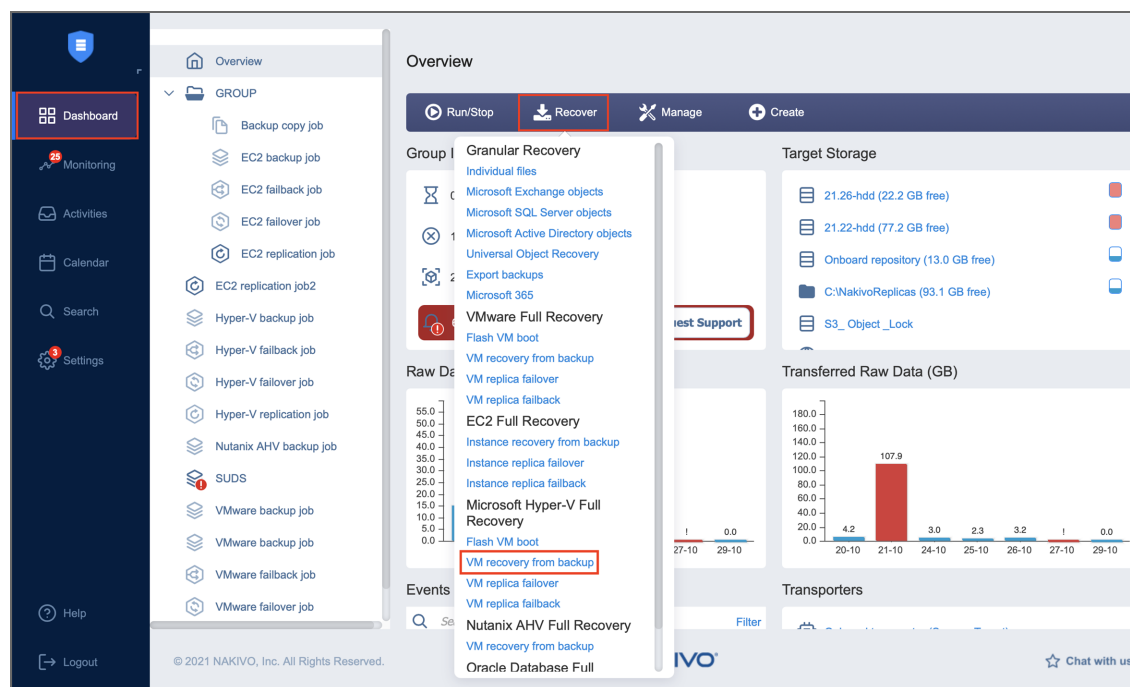
Refer to the following topics for more information:

- [“Starting Hyper-V VM Recovery” below](#)
- [“Recovery Job Wizard for Hyper-V: Backups” on the next page](#)
- [“Recovery Job Wizard for Hyper-V: Destination” on page 915](#)
- [“Recovery Job Wizard for Hyper-V: Options” on page 918](#)

## Starting Hyper-V VM Recovery


To recover entire Hyper-V VMs from backups, do one of the following:

- Start the recovery process from the **Dashboard** by clicking **Recover** and then clicking **VM recovery from backup**.



- Open the **New Recovery Job Wizard** from the **Repositories** tab following the steps below:
  1. Go to the **Repositories** tab.
  2. Click **Settings** in the main menu of NAKIVO Backup & Replication.
  3. Hover over the Backup Repository containing the necessary backup, click **Recover** and then click **VM recovery from backup**.





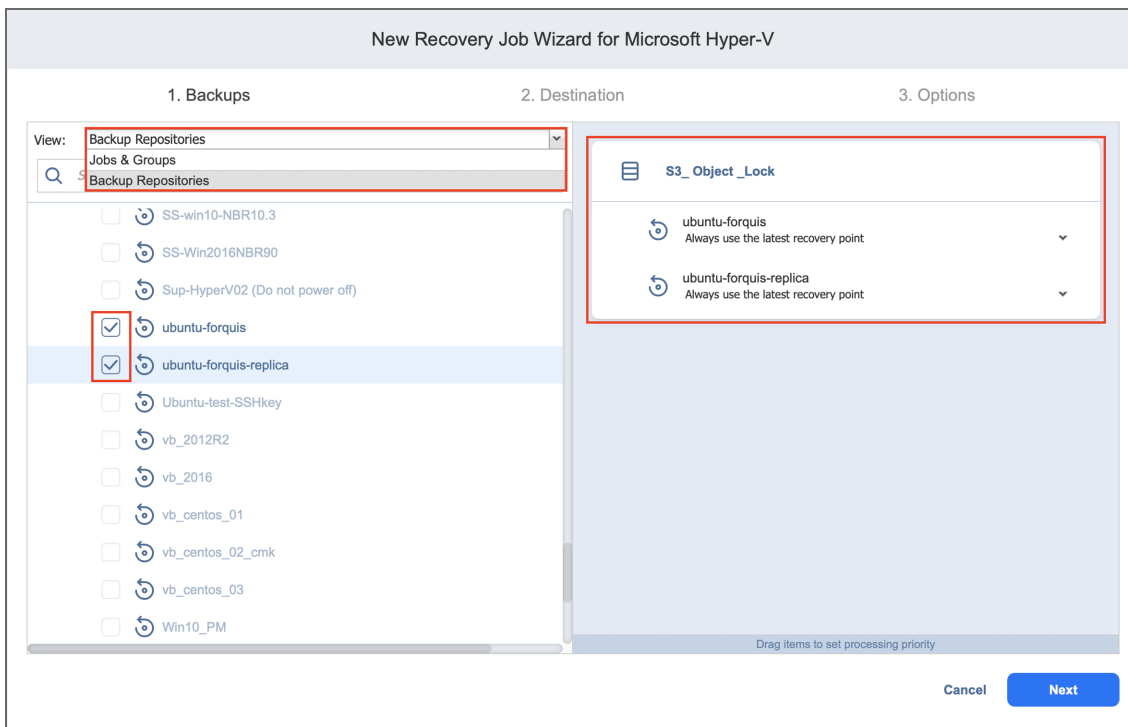
- Alternatively, the recovery can be performed by using the [search](#) function.

The **New Recovery Job Wizard for Hyper-V** opens.

## Recovery Job Wizard for Hyper-V: Backups

On the **Backups** page of the wizard:

1. In the left pane, select one or more VM backups using one of the following views:
  - **Jobs & Groups**
  - **Backup Repositories**
2. Select a recovery point for each backup in the right pane.



3. Click **Next** to go to the next step.

## Recovery Job Wizard for Hyper-V: Destination

On the **Destination** page, select a location for the recovered VMs.

- [Setting the Same Host, Datastore, and Network for All Recovered VMs](#)
- [Setting the Default Destination for Recovered VMs](#)
- [Setting Different Options for Recovered VMs](#)

### Setting the Same Host, Datastore, and Network for All Recovered VMs

To recover all VMs to the same server and location and to connect all recovered VMs to the same networks, follow the steps below:

1. Choose a server from the **Container** drop-down list.
2. Type a location in the **Path** field. It can be a local or a shared path.

3. Choose a network from the **Network** drop-down list.

New Recovery Job Wizard for Microsoft Hyper-V

1. Backups      2. Destination      3. Options

Container: ServerHV2012

Path: C:\NakivoRecovered

Network: New Virtual Switch

[Advanced setup...](#)

Cancel    Next

## Note

To connect to a shared path successfully, the following conditions must be met:

1. The shared path is created with the same credentials as the corresponding Hyper-V container that was added. See [Adding Hyper-V Servers](#) for details on adding Hyper-V containers to NAKIVO Backup & Replication.
2. The logon session in which you created the shared path has not been ended.

As a workaround, create a symbolic link to the shared path from the Hyper-V container. Refer to Step 7 of [High Availability of NAKIVO Backup and Replication](#) for details.

## 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. To do this, follow the steps below:

1. Click **Advanced options** and then click on the name of the chosen host, cluster, or folder.
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**.

New Recovery Job Wizard for Microsoft Hyper-V

1. Backups      2. Destination      3. Options

Container: ServerHV2012

Path: C:\NakivoRecovered

Network: New Virtual Switch

S3\_Object\_Lock Click to collapse

Default container: ServerHV2012

Default path: C:\NakivoRecovered

Default network: New Virtual Switch

ubuntu-forquis

ubuntu-forquis-replica


Cancel    Next

## Setting Different Options for Recovered VMs

To specify different options for recovered VMs, follow the steps below:

1. Click **Advanced setup**.
2. Choose a target container, disk, and network adapter for each VM.

3. Click **Next** to go to the next step.



## Recovery Job Wizard for Hyper-V: Options

On the **Options** page, you can set up recovery options for the job.

- [“Job Options” below](#)
- [“Recovery Options” on the next page](#)
- [“Pre and Post Job Actions” on page 920](#)
  - [“Setting Up a Pre Job Script” on page 920](#)
  - [“Setting Up a Post Job Script” on page 921](#)
  - [“Email Notifications” on page 922](#)
- [“Data Transfer” on page 922](#)
  - [“Transporter Load” on page 922](#)
  - [“Bandwidth Throttling” on page 922](#)
  - [Bottleneck Detection](#)
- [“Completing the New Recovery Job Wizard for Microsoft Hyper-V” on page 923](#)

## Job Options

In the *Job Options* section, specify the following:

1. **Job name:** Specify a name for the recovery job.
2. **Network acceleration:** When [network acceleration](#) is 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.

3. **Encryption** When [encryption](#) is 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 recovering over WAN without a VPN connection.

The screenshot shows the 'New Recovery Job Wizard for Microsoft Hyper-V' interface, specifically the 'Options' step. The 'Job Options' section is highlighted with a red box and contains the following settings:

- Job name: Hyper-V recovery job
- Network acceleration: Disabled
- Network encryption: Disabled

Other sections visible in the wizard include:

- Recovered VM Options:**
  - Recovery mode: Synthetic
  - Recovered VM names: Append "-recovered" in the end
  - VM disks: Respect original VM disk type
  - VM power on: Power on recovered VMs
  - VM MAC addresses: Do not generate new MAC address
- Pre and Post Actions:**
  - Send job run reports to: [empty field]
  - Run local pre job script: [checkbox]
  - Run local post job script: [checkbox]
- Data Transfer:**
  - Limit transporter load to: 3 concurrent tasks
  - Bandwidth throttling: Disabled

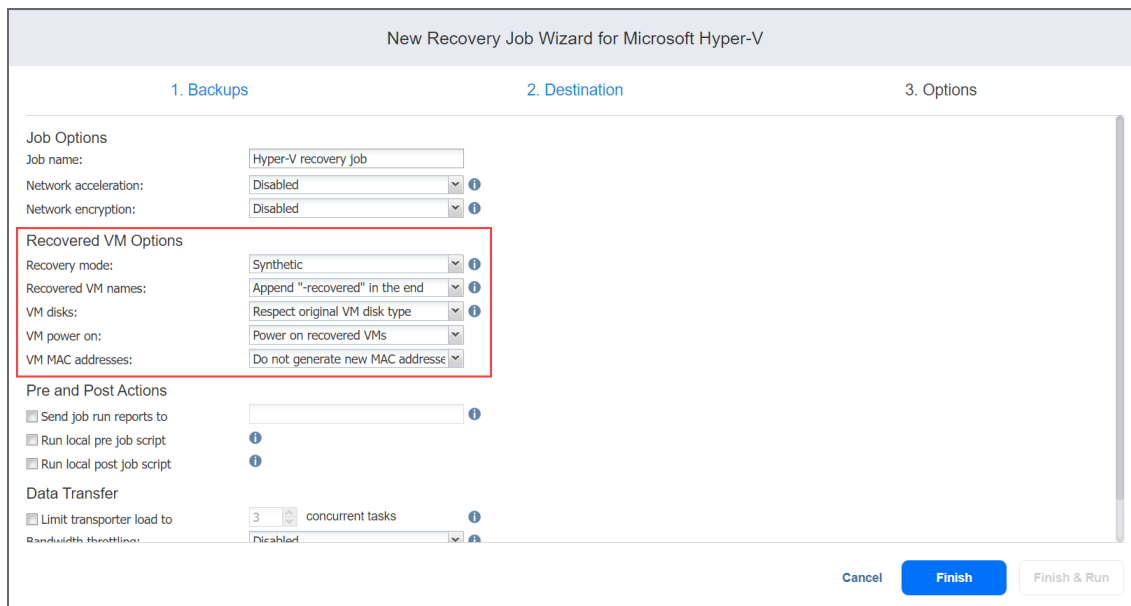
At the bottom right, there are three buttons: 'Cancel', 'Finish', and 'Finish & Run'.

## Recovery Options

In the *Recovery Options* section, specify the following:

- **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 will be preserved on the recovered VM (s). Make sure the location to which the VM(s) will be recovered does not contain the original VM(s), otherwise UUID and MAC address conflicts may occur.
- **Recovered VM names:** Choose one of the following:
  - **Append "-recovered" in the end:** Source VM names are used for recovered VM names and "-recovered" are 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 VM.
- **VM disks:** Choose one of the following:
  - **Respect original VM disk type:** When specified, disks that respect original VM disk type will be 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 will be created on target VMs. Use this option to save space on target datastore.
- **VM MAC addresses:** Choose one of the following:
  - **Do not generate new MAC addresses:** When this option is chosen, the recovered VM will have the same MAC address as the source VM.

- **Generate new MAC addresses:** When this option is chosen, a new MAC address will be generated for the recovered VM.
- **VM power on:** When the **Power on recovered VMs** option is chosen, the recovered VMs will be powered on.



## Pre and Post Job 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 where the Director is installed. Also, you can set up email notifications to be sent to the specified recipients on job completion.

### 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. 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 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, VM backup will only be started after the script is completed.
    - **Do not wait for the script to finish:** With this option selected, the product will run the script and will start 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 will perform VM backup even if the script has failed.
  - **Fail the job on script failure:** With this option selected, if the script fails, the job will be failed and VM backup will not be performed.

## Setting Up 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 where 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 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.

The screenshot shows the 'New Recovery Job Wizard for Microsoft Hyper-V' dialog box, specifically the '3. Options' tab. The 'Pre and Post Actions' section is highlighted with a red box. It contains three options:

- Send job run reports to: [Text input field]
- Run local pre job script
- Run local post job script

Other options visible in the dialog include:

- Job Options:** Job name: Hyper-V recovery job; Network acceleration: Disabled; Network encryption: Disabled.
- Recovered VM Options:** Recovery mode: Synthetic; Recovered VM names: Append "-recovered" in the end; VM disks: Respect original VM disk type; VM power on: Power on recovered VMs; VM MAC addresses: Do not generate new MAC address.
- Data Transfer:** Limit transporter load to: 3 concurrent tasks; Bandwidth throttling: Disabled.

Buttons at the bottom: Cancel, Finish, Finish & Run.



## 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:

- In the *Pre and Post Actions* section, select the **Send job run reports to** option
- Specify one or more email addresses in the text box. Use semicolons to separate multiple email addresses.

## Data Transfer

In the *Data Transfer* section, specify the Transporter load for the recovery job and set up [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. Select the **Limit transporter load to** checkbox.
2. Specify the number of concurrent tasks in the corresponding box.

### Bandwidth Throttling

Please follow the steps below to regulate 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” on page 284](#) 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 recovery job:
    - a. Click the **Create New Rule** button.
    - b. The **New Bandwidth Rule** dialog box opens. Refer to [“Bandwidth Throttling” on page 284](#) 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, 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.

The screenshot shows the 'New Recovery Job Wizard for Microsoft Hyper-V' dialog box, specifically the 'Options' tab. The 'Data Transfer' section is highlighted with a red box. It contains the following settings:

- Limit transporter load to: 3 concurrent tasks
- Bandwidth throttling: Disabled
- Bottleneck detection

Other settings visible in the dialog include:

- Network acceleration: Disabled
- Network encryption: Disabled
- Recovered VM Options:
  - Recovery mode: Synthetic
  - Recovered VM names: Append "-recovered" in the end
  - VM disks: Respect original VM disk type
  - VM power on: Power on recovered VMs
  - VM MAC addresses: Do not generate new MAC address
- Pre and Post Actions:
  - Send job run reports to
  - Run local pre job script
  - Run local post job script

Buttons at the bottom include 'Cancel', 'Finish', and 'Finish & Run'.

## Completing the New Recovery Job Wizard for Microsoft Hyper-V

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 256](#) for details.

## Amazon EC2 Instance Recovery

With NAKIVO Backup & Replication, you can easily recover Amazon EC2 instances from their backups. When you run the recovery, a new Amazon EC2 instance is created, that is, the original Amazon EC2 instance (if it exists in the recovery location) is not reverted to a previous state or replaced with the new instance.

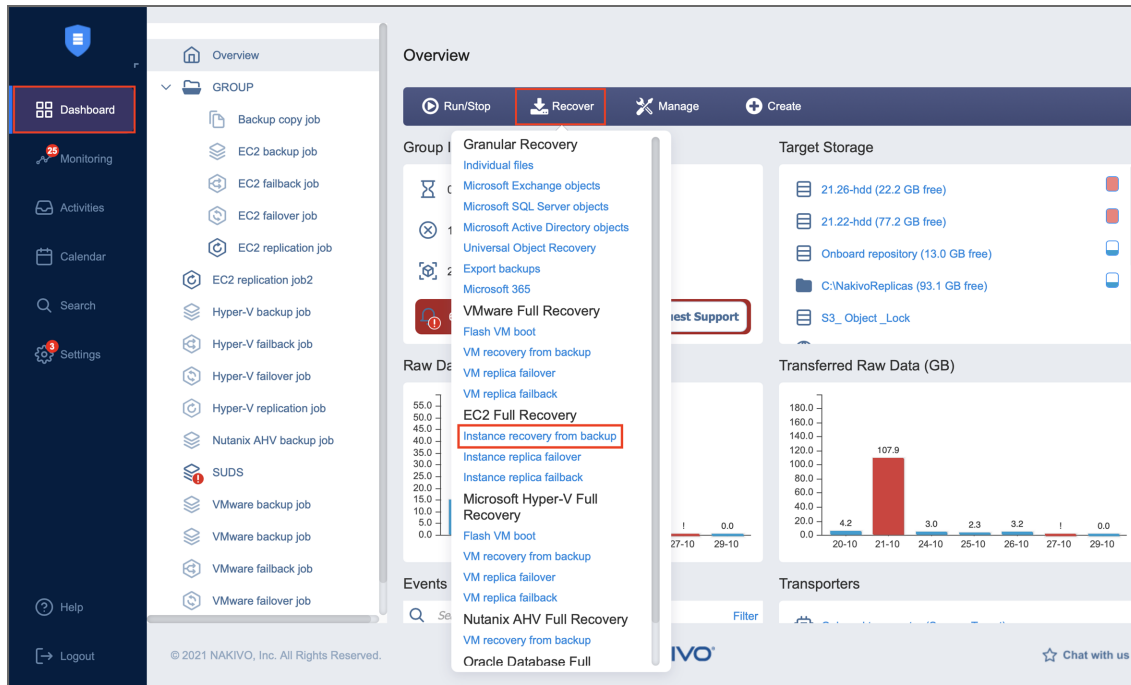
Refer to the following topics for more information:

- [“Starting Amazon EC2 Instance Recovery” below](#)
- [“Recovery Job Wizard for Amazon EC2: Backups” on page 925](#)
- [“Recovery Job Wizard for Amazon EC2: Destination” on page 925](#)
- [“Recovery Job Wizard for Amazon EC2: Options” on page 928](#)

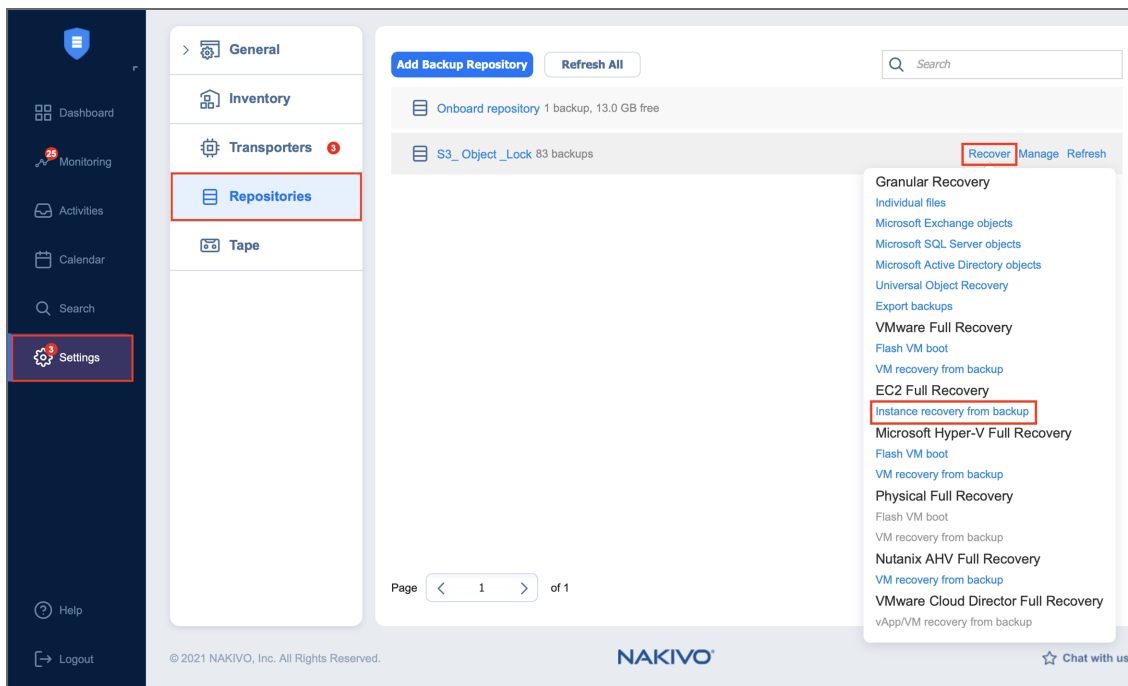
## Starting Amazon EC2 Instance Recovery

To recover Amazon EC2 instances from backups, do one of the following:

- Start recovery from the **Dashboard** by clicking **Recover>Instance recovery from backup**.



- Start recovery from a Backup Repository by following the steps below:
  1. Click **Settings** in the main menu of NAKIVO Backup & Replication.
  2. Go to the **Repositories** tab.
  3. Hover over a Backup Repository that contains Amazon EC2 backups and click **Recover>Instance recovery from backup**.

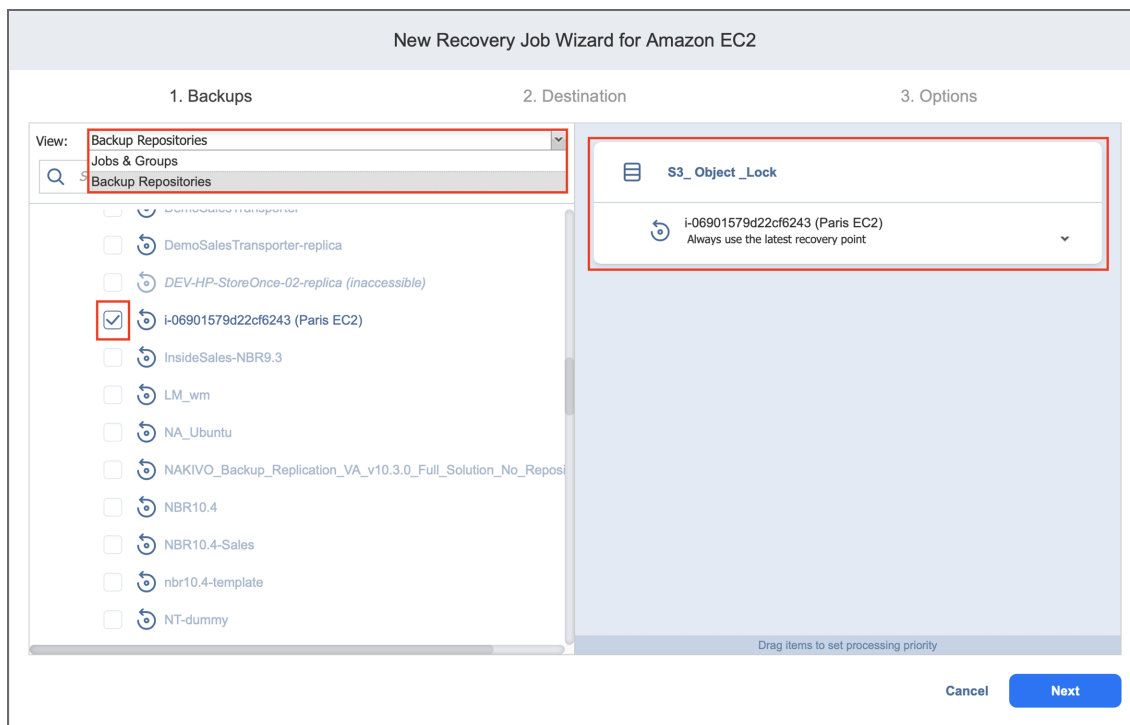


- Alternatively, the recovery can be performed by using the **search** function.

The **New Recovery Job Wizard for Amazon EC2** opens.

## Recovery Job Wizard for Amazon EC2: Backups

On the **Backups** page of the wizard, select one or more Amazon EC2 backups in the left pane and then select a recovery point for each backup in the right pane.



Click **Next** to go to the next page.

## Recovery Job Wizard for Amazon EC2: Destination

On the **Destination** page, choose a recovery location for the selected Amazon EC2 instances. Proceed as described in the sections below:

- [“Setting the Same Recovery Location for All Amazon EC2 Instances” below](#)
- [“Setting Different Recovery Locations for Amazon EC2 Instances” on page 927](#)

### Note

You may be additionally charged for using a third-party resource. Refer to the third-party resource provider documentation for details.

## Setting the Same Recovery Location for All Amazon EC2 Instances

To recover all selected Amazon EC2 instances to the same Amazon region and with the same destination options, set up the following:

- **Region:** Select a target Amazon EC2 [region](#) to which all Amazon EC2 instances will be recovered.
- **Network:** Select an Amazon EC2 [VPC](#) to which all Amazon EC2 instances will be connected.
- **Subnet:** Select an Amazon EC2 [Subnet](#) to which all Amazon EC2 instances will be connected.
- **Security group:** Select an Amazon EC2 [Security Group](#) for all Amazon EC2 instances.
- **Instance access:** Specify Amazon EC2 [Key Pair](#) options as follows:
  - **Use existing key pair:** Choose an existing Amazon EC2 Key Pair that you want use to connect to the recovered instances.
  - **Create a new key pair:** Select this option to create a new Amazon EC2 Key Pair and assign it to all recovered Amazon EC2 instances. Type in a name for the new Amazon EC2 Key Pair in the **Key pair name** field and click **Download**.

#### **Note**

This is your only chance to save the private key file, so make sure you save the private key file in a safe place. You will need to provide the name of your key pair when you start the recovered Amazon EC2 instances and the corresponding private key each time you connect to the instances.

- **Proceed without key pair:** Select this option if you do not want to assign an Amazon Key Pair to the recovered Amazon EC2 instances.

#### **Note**

A key pair enables you to connect to Linux instances via SSH. If you recover Linux instances without a key pair, you will not be able to connect to them.



## Note

This is your only chance to save the private key file, so make sure you save the private key file in a safe place. You will need to provide the name of your key pair when you start the recovered Amazon EC2 instance and the corresponding private key each time you connect to the instance.

- **Proceed without key pair:** Select this option if you do not want to assign an Amazon Key Pair to the recovered Amazon EC2 instance.

## Note

A key pair enables you to connect to Linux instances via SSH. If you recover Linux instances without a key pair, you will not be able to connect to them.

New Recovery Job Wizard for Amazon EC2

1. Backups 2. Destination 3. Options

Key pair:

i-06901579d22cf6243 (Paris EC2) Click to collapse

Source	Target
Location: S3_Object_Lock	Region: EU (Ireland)
Instance type: t3.micro	Instance type: t3.micro
EBS volumes	EBS volumes: Recover
/dev/sda1: 8.0 GB	/dev/sda1: Recover
Networking	Networking
Network: vpc-ddeca7b4	Network: sy-vpc (10.30.122.0/24)
Network adapter 1	Network adapter 1
Subnet: subnet-cf7707a6	Subnet: sy-test-pub (10.30.122.0/25)
Security group 1: sg-06fd8e9ee8b5489ee	Security group: amazon_084ead86-f3d9-4ab3-8790-70fd571699d3
Primary IP: 10.0.253.236	Primary IP: Enter primary IP (optional)
Instance access	Instance access
Key pair name: No data	Instance access: Use existing key pair
	Key pair: Select target key pair

Cancel Finish Finish & Run

Click **Next** to go to the next page of the wizard.

## Recovery Job Wizard for Amazon EC2: Options

On the **Options** page, set up recovery job options as described in the sections below:

- [“Job Options” on the next page](#)
- [“Recovered Instance Options” on the next page](#)
- [“Pre and Post Actions” on page 931](#)
  - [“Running a Pre Job Script” on page 931](#)
  - [“Running a Post Job Script” on page 931](#)

- [“Email Notifications” on page 932](#)
- [“Data Transfer” on page 932](#)
  - Transporter Pool
  - [“Transporters” on page 933](#)
  - [“Transporter Load” on page 933](#)
  - [“Bandwidth Throttling” on page 933](#)
  - Bottleneck Detection
- [“Completing the New Recovery Job Wizard for Amazon EC2” on page 934](#)

## Job Options

In this section, specify general options of your Amazon EC2 recovery job:

- **Job Name:** Specify a name for the recovery job.
- **Network Acceleration:** When network acceleration is enabled, NAKIVO Backup & Replication will use compression and traffic reduction techniques to speed up data transfer. Select this option if you plan to recover VMs over WAN or slow LAN links. Refer to [“Network Acceleration” on page 63](#) for details.
- **Encryption:** When this option is 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. Refer to [“Encryption in Flight and at Rest” on page 42](#) for more information.

The screenshot shows the 'New Recovery Job Wizard for Amazon EC2' interface, specifically the '3. Options' step. The 'Job Options' section is highlighted with a red box and contains the following fields:

- Job name: EC2 recovery job
- Network acceleration: Disabled
- Network encryption: Disabled

Below the 'Job Options' section, there are several other sections:

- Recovered Instance Options:**
  - Recovery mode: Synthetic
  - Recovered instance names: Append "-recovered" in the end
  - EBS volumes: Respect original volume type
  - Instance power on: Power on recovered instances
- Pre and Post Actions:**
  - Send job run reports to
  - Run local pre job script
  - Run local post job script
- Data Transfer:**
  - Transport pool: Select transporter pool
  - Transporters: Automatic selection
  - Limit transporter load to 2 concurrent tasks

At the bottom right of the wizard, there are three buttons: 'Cancel', 'Finish', and 'Finish & Run'.

## Recovered Instance Options

In this section, you can specify the options for your recovered instance:



- **Recovery mode:** Choose one of the following recovery modes:
  - **Synthetic:** The location and environment dependencies will be removed from recovered Amazon EC2 instances. All instances will be fully functional and will contain all their data.
  - **Production:** The environment dependencies will be preserved on the recovered Amazon EC2 instances. Make sure the location to which the instances will be recovered does not contain the original instances, otherwise conflicts may occur.
- **Recovered instance names:** NAKIVO Backup & Replication allows you to change recovered Amazon EC2 instance names. By default, the text “- recovered” is appended to the end of recovered instance names. To change Amazon EC2 instance names, choose one of the following:
  - **Append "recovered" in the end:** Source instance names are used for recovered instance names and the specified text is added after the name of the recovered instance.
  - **Leave recovered instance names as is:** Recovered Amazon EC2 instances names are identical to the source instance names.
  - **Enter custom recovered instance names:** You can enter custom names for recovered Amazon EC2 instances.
- **EBS volumes:** Choose one of the following options:
  - **Use Magnetic type for all volumes:** When this option is selected, NAKIVO Backup & Replication restores all EBS Volumes as Magnetic, regardless of the source EBS volume types.
  - **Respect original volume type:** When this option is selected, the type of source EBS volumes is preserved.
- **Instance power on:** Select one of the following options:
  - **Power on recovered instance:** When this option is selected, the instances are powered on after recovery.
  - **Do not power on recovered instance:** When this option is selected, the recovered instances are not powered on after recovery.

The screenshot shows the 'New Recovery Job Wizard for Amazon EC2' interface, specifically the '3. Options' step. The 'Recovered Instance Options' section is highlighted with a red box. The settings are as follows:

- Job name: EC2 recovery job
- Network acceleration: Disabled
- Network encryption: Disabled
- Recovered Instance Options:
  - Recovery mode: Synthetic
  - Recovered instance names: Append "-recovered" in the end
  - EBS volumes: Respect original volume type
  - Instance power on: Power on recovered instances
- Pre and Post Actions:
  - Send job run reports to: [empty]
  - Run local pre job script: [checkbox]
  - Run local post job script: [checkbox]
- Data Transfer:
  - Transport pool: Select transporter pool
  - Transporters: Automatic selection
  - Limit transporter jobs to: 2 concurrent tasks

Buttons at the bottom: Cancel, Finish, Finish & Run.

## Pre and Post Actions

NAKIVO Backup & Replication allows you to run a [script](#) before recovery begins (a pre job script) and after the recovery has completed (a post job script). The scripts can only be executed on the machine where the Director is installed. Also, you can set up email notifications to be sent to the specified recipients about job completion status. See the details below.

### Running a Pre Job Script

To run a script before the product begins recovery, do the following:

1. Place a script file on the machine where 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 that opens:
  1. **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`
  2. **Job behavior:** Choose either of the following job behaviors in relation to scrip completion:
    - **Wait for the script to finish:** With this option selected, recovery is started only 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 instances at the same time.
  3. **Error handling:** Choose either of the following job behaviors in relation to scrip failure:
    - **Continue the job on script failure:** With this option selected, the job will perform recovery even if the script has failed.
    - **Fail the job on script failure:** With this option selected, if the script fails, the job will fail and recovery will not be performed.

### Running a Post Job Script

To run a script after the product has finished recovery, do the following:

1. Place a script file on the machine where the Director is installed.
2. Select the **Run local post job script** option and click the **settings** link.
3. Specify the following options in the dialog box that opens:
  1. **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`
  2. **Job behavior:** Choose either of the following job behaviors in relation to scrip 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.
3. **Error handling:** Choose either of the following job behaviors in relation to scrip 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 recovery has been successful.

## Email Notifications

NAKIVO Backup & Replication can send [email notifications](#) about job completion status to specified recipients.

To enable this option, make sure that [Email settings](#) are configured.

To send email notifications:

1. Select the **Send job run reports to** option.
2. Specify one or more email addresses in the text field. Use semicolons to separate multiple email addresses.

The screenshot shows the 'New Recovery Job Wizard for Amazon EC2' with three steps: 1. Backups, 2. Destination, and 3. Options. The '3. Options' step is active. Under 'Job Options', 'Job name' is 'EC2 recovery job', 'Network acceleration' is 'Disabled', and 'Network encryption' is 'Disabled'. Under 'Recovered Instance Options', 'Recovery mode' is 'Synthetic', 'Recovered instance names' is 'Append "-recovered" in the end', 'EBS volumes' is 'Respect original volume type', and 'Instance power on' is 'Power on recovered instances'. The 'Pre and Post Actions' section is highlighted with a red box and contains three items: 'Send job run reports to' (checked), 'Run local pre job script' (unchecked), and 'Run local post job script' (unchecked). The 'Data Transfer' section has 'Transport pool' set to 'Select transporter pool' and 'Transporters' set to 'Automatic selection'. At the bottom right are 'Cancel', 'Finish', and 'Finish & Run' buttons.

## Data Transfer

In this section, you can specify data transfer options for the recovered instance.

### 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: read data, compress data, transfer data over the network, perform data deduplication, and so on. If more than one Transporter is deployed for NAKIVO Backup & Replication, it is important to determine which ones should be used to read data from a particular instance and which ones should be used to write data to the target AWS Region.

Choose one of the following options:

- **Automatic selection:** The product automatically determines which Transporters to use.
- **Manual - configured for all instances:** Select this option to manually specify a single target Transporter to be used for writing data to target regions.
- **Manual - configured per region:** Select this option to manually specify a Transporter for each region.

## 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. Select the **Limit transporter load to** checkbox.
2. Specify the number of concurrent tasks in the corresponding box.

## Bandwidth Throttling

Please follow the steps below to regulate 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:
    - a. Click the **Create New Rule** button.
    - b. The **New Bandwidth Rule** dialog box opens. Refer to [“Bandwidth Throttling” on page 284](#) 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, 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.

The screenshot shows the 'New Recovery Job Wizard for Physical Machines' dialog box, specifically the 'Options' tab. The dialog is divided into three sections: '1. Backups', '2. Destination', and '3. Options'. The 'Options' section contains several settings:

- Recovered VM Options:**
  - Recovered VM names: Append "--recovered" in the end
  - VM disks: Create only thin disks on target VM
  - VM MAC addresses: Generate new MAC addresses
  - VM power on: Power on recovered VMs
- Pre and Post Actions:**
  - Send job run reports to: (empty field)
  - Run local pre job script: (checkbox)
  - Run local post job script: (checkbox)
- Data Transfer:**
  - Transport mode: Automatic selection
  - Transporter pool: Select transporter pool
  - Transporters: Automatic selection
  - Limit transporter load to: 3 concurrent tasks
  - Bandwidth throttling: Disabled
  - Bottleneck detection: (checkbox, highlighted with a red box)**

At the bottom right, there are three buttons: 'Cancel', 'Finish', and 'Finish & Run'.

## Completing the New Recovery Job Wizard for Amazon EC2

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 256](#) for details.

## Physical Machine Recovery

With NAKIVO Backup & Replication, you can perform a [full recovery of a physical machine](#) to a VMware virtual machine. This feature allows you to protect mixed IT environments.

### Note

Free ESXi is not supported for physical to virtual recovery.

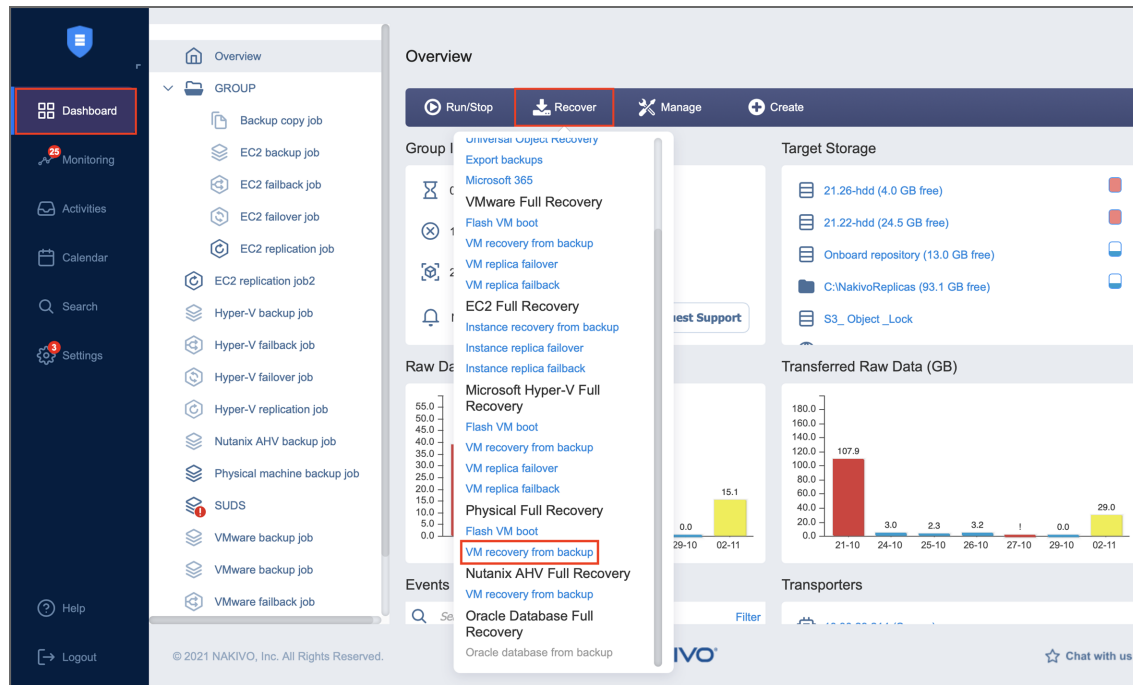
Refer to the following topics to learn how to perform a full recovery of a physical machine to a VMware VM:

- [“Starting Physical Machine Recovery” on the next page](#)
- [“Recovery Job Wizard for Physical Machines: Backups” on page 936](#)
- [“Recovery Job Wizard for Physical Machines: Destination” on page 937](#)
- [“Recovery Job Wizard for Physical Machines: Options” on page 939](#)

## Starting Physical Machine Recovery

To recover a physical machine to a VMware VM, take one of the following actions:

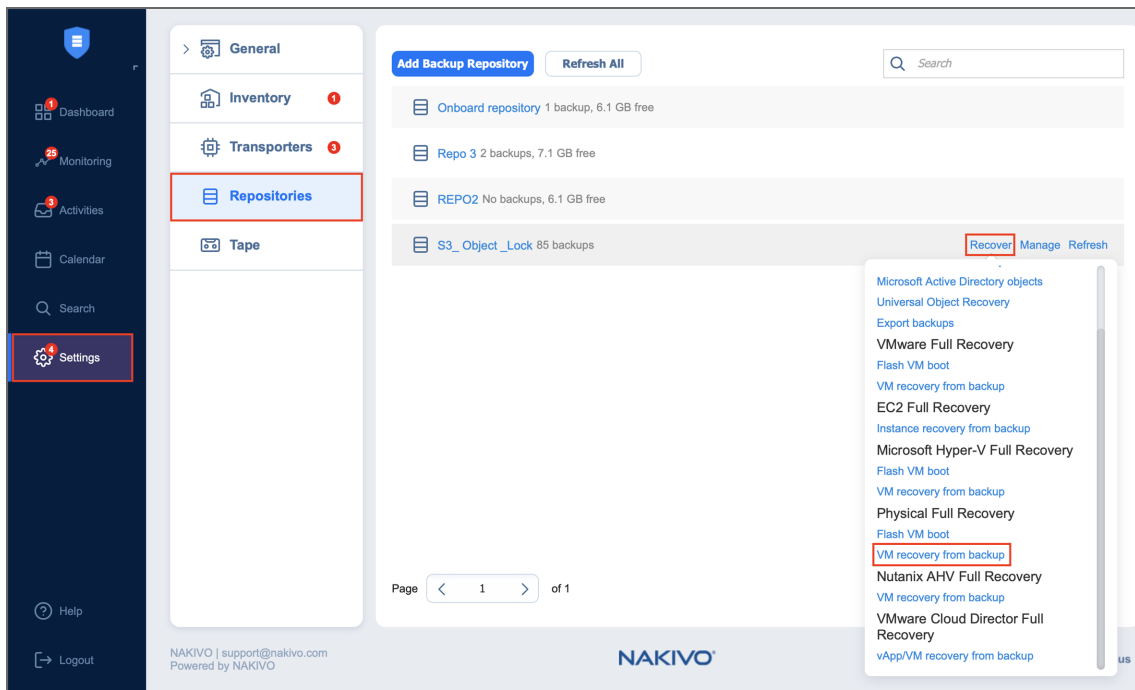
- Go to Dashboard, click **Recover**, and select **Physical Full Recovery > VM recovery from backup**.



The **New Recovery Job Wizard for Physical Machines** opens.

- Navigate to **Settings** and do the following:
  1. Click **Repositories**.
  2. Hover over the repository with the backup you want to recover.

3. Select **Physical Full Recovery > VM recovery from backup.**

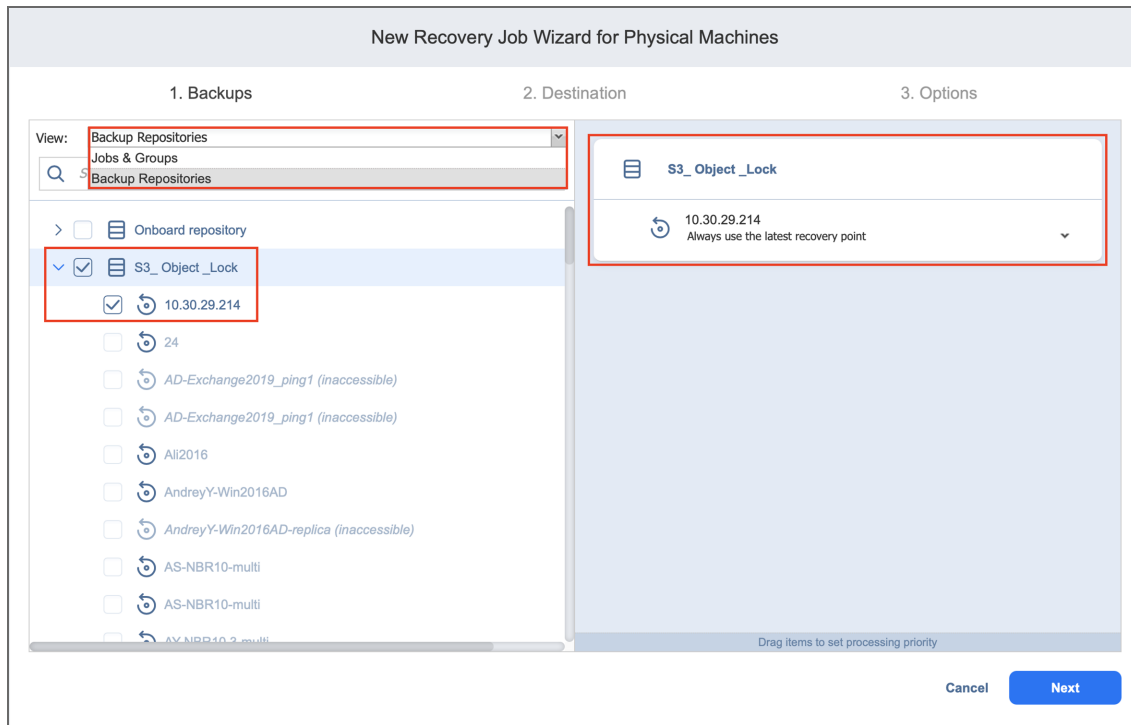


- Alternatively, the recovery can be performed by using by using the [search](#) function.  
The **New Recovery Job Wizard for Physical Machines** opens.

## Recovery Job Wizard for Physical Machines: Backups

On the **Backups** page of the **Recovery Job Wizard for Physical Machines**:

1. Select one of the views:
  - **Jobs & Groups**: Select one or more backups in the left pane and then select a recovery point for each backup in the right pane.
  - **Backup Repositories**: Select one or more backups from the Backup Repositories in the left pane and then select a recovery point for each backup in the right pane.
2. Click **Next** to move to the next page.



## Recovery Job Wizard for Physical Machines: Destination

Choose the location for storing the recovered physical machines.

- [“Setting the Same Host, Datastore, and Network for All Recovered VMs” below](#)
- [“Setting Different Options for Each Recovered VM” on the next page](#)

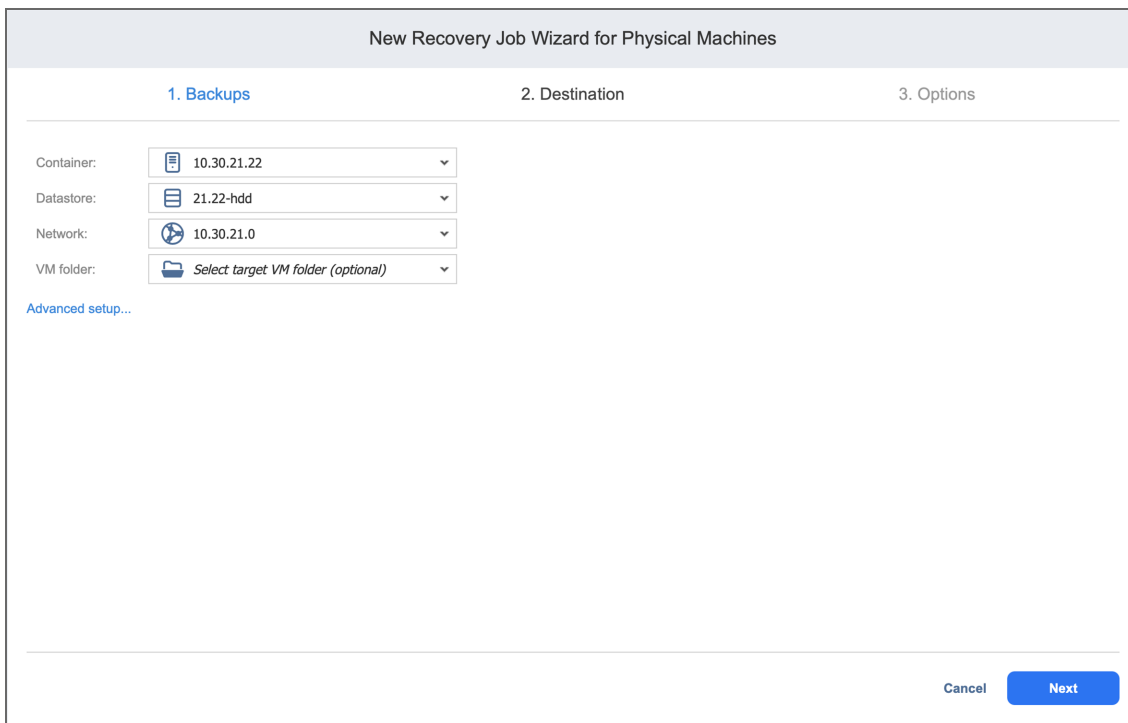
### Setting the Same Host, Datastore, and Network for All Recovered VMs

To recover all machines to the same container/folder and datastore, and to connect all recovered VMs to the same networks, 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. Choose a network from the **Network** drop-down list.

Optionally, you can choose a folder from the **VM folder** drop-down list if there is one in the container.





## Setting Different Options for Each Recovered VM

To specify different options for each recovered physical machine, follow the steps below:

1. Click **Advanced options**.
2. Click on the backup to expand its recovery options.
3. Choose a target location by selecting the necessary container, virtual machine, and folder.
4. Configure VM resources:
  - **Virtual CPU**
  - **Cores per socket**
  - **RAM**

By default, the displayed VM resources correspond to the source physical machine configuration.

### **Important**

If the default CPU configuration has been changed, the target VM might become unstable. In addition, the modified configuration might not comply with the licensing policy of the Guest OS.

5. Select a disk from the **Disks** drop-down list.
6. Keep a source disk controller configuration by selecting **Keep source configuration** from the **Disk Controller** drop-down list (recommended option) or select one of these types of disk controllers:
  - **SCSI LSI Logic SAS**
  - **SCSI LSI Logic Parallel**
  - **SCSI VMware Paravirtual**
  - **SCSI BusLogic Parallel**

- IDE
- SATA
- NVME

### Important

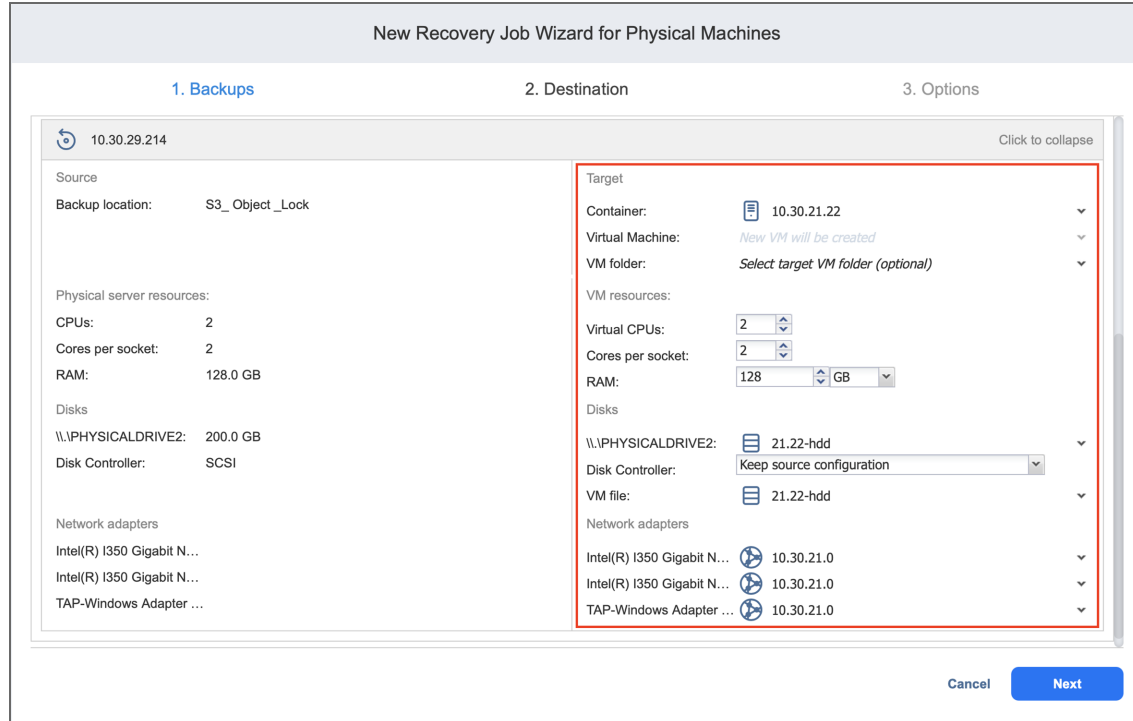
If for the target VM you select a disk controller type that differs from a source machine, the recovery may fail with an error and the emergency mode will be turned on on the recovered machine.

7. Select a VM file from the **VM file** drop-down list.
8. Select network adapters from the **Network adapters** drop-down list. For each physical network adapter, a virtual network adapter is created. The other available options are:
  - Skip this network adapter
  - Not connect to any virtual network
  - Connect to temporary isolated network

### Important

VMware allows you to assign no more than 10 network adapters (NICs) per one VM. Therefore, if the source machine has more than 10 NICs, you need to manually skip some NICs to make sure that the machine has no more than 10 NICs.

9. Click **Next** to go to the next page of the wizard.



## Recovery Job Wizard for Physical Machines: Options

On the **Options** page, set the options for the physical machine recovery job.

- [“Job Options” below](#)
- [“Recovered VM Options” below](#)
- [“Pre and Post Actions” on the next page](#)
- [“Data Transfer” on page 943](#)
  - [“Transport Mode” on page 943](#)
  - Transporter Pool
  - [“Transporters” on page 943](#)
  - [“Transporter Load” on page 944](#)
  - [“Bandwidth Throttling” on page 944](#)
  - Bottleneck Detection

## Job Options

In the *Job Options* section, set the following:

- **Job name:** Enter the name for the recovery job.
- **Network acceleration:** Enable [network acceleration](#) if you transfer data over a slow WAN. Note that you need at least one Transporter on source and target sites for this feature to work.
- **Encryption:** Enable [encryption](#) to protect your data while transferring it over a WAN without VPN. Job data will be encrypted during the transfer that will increase the load on the Transporter(s).

The screenshot shows the 'New Recovery Job Wizard for Physical Machines' interface. The wizard is divided into three steps: 1. Backups, 2. Destination, and 3. Options. The 'Job Options' section is highlighted with a red box and includes the following fields:

- Job name: Physical machine recovery job
- Network acceleration: Disabled
- Network encryption: Disabled

Below the 'Job Options' section are the following sections:

- Recovered VM Options:**
  - Recovered VM names: Append "-recovered" in the end
  - VM disks: Create only thin disks on target VM:
  - VM MAC addresses: Generate new MAC addresses
  - VM power on: Power on recovered VMs
- Pre and Post Actions:**
  - Send job run reports to: [empty field]
  - Run local pre job script: [checkbox]
  - Run local post job script: [checkbox]
- Data Transfer:**
  - Transport mode: Automatic selection
  - Transporter pool: Select transporter pool
  - Transporters: Automatic selection

The bottom right corner of the wizard has three buttons: Cancel, Finish, and Finish & Run.

## Recovered VM Options

In the *Recovered VM Options* section, set the following:

- **Recovered VM names:** Select one of the following VM name options:
  - **Append "-recovered" in the end:** Select this option to use the name of the source physical machine for the recovered VM name with **-recovered** added to the in the end.

- **Leave recovered VM names as is:** Select this option to retain the name of the source physical machine for the recovered VM name.
- **Enter custom recovered VM names:** Select this option specify a custom name for the recovered VM.
- **VM disks:** Select one of the following VM disk types:
  - **Respect original VM disk type:** Select this option to keep the same disk type as the source machine for the recovered VM.
  - **Create only thin disks on target VMs:** Select this option to create thin disks on your target VM.
- **VM MAC addresses:** Select one of the following actions for the recovered VM:
  - **Generate new MAC addresses**
  - **Do not generate new MAC addresses**
- **VM power on:** Select one of the following options:
  - **Power on recovered VMs**
  - **Do not power on recovered VMs**

The screenshot shows the 'New Recovery Job Wizard for Physical Machines' interface, specifically the '3. Options' step. The 'Recovered VM Options' section is highlighted with a red box. The settings in this section are:

- Job name: Physical machine recovery job
- Network acceleration: Disabled
- Network encryption: Disabled
- Recovered VM names: Append "-recovered" in the end
- VM disks: Create only thin disks on target VMs
- VM MAC addresses: Generate new MAC addresses
- VM power on: Power on recovered VMs

Other sections visible include 'Pre and Post Actions' with checkboxes for 'Send job run reports to', 'Run local pre job script', and 'Run local post job script'. The 'Data Transfer' section includes 'Transport mode' (Automatic selection) and 'Transporter pool' (Select transporter pool). At the bottom right, there are 'Cancel', 'Finish', and 'Finish & Run' buttons.

## Pre and Post Actions

In the *Pre and Post Actions* section, set the following actions after the recovery job is completed:

- **Send job run reports to:** Enter one or more email addresses in the text field. Use semicolons to separate multiple email addresses.

### Note

To enable this option, make sure that [Email settings](#) are configured.

- **Run local pre job script:** To run a [script](#) after the product has finished the recovery job, 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 options in the dialog box that appears:

- **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*

- **Job behavior:** Select one of the following job behaviors in relation to script completion:
  - **Wait for the script to finish:** When selected, the job remains in the “running” state until the script is executed.
  - **Do not wait for the script to finish:** When selected, the product runs the script and starts the recovery process at the same time.
- **Error handling:** Select one of the following job behaviors in relation to script failure:
  - **Continue the job on script failure:** When selected, script failure does not influence the status of the job.
  - **Fail the job on script failure:** When selected and the script fails, the job status is set to “failed” and recovery is not performed.

- **Run local post job script: To run a script after the product has finished the recovery process, 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 options in the dialog box that appears:

- **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*

- **Job behavior:** Select one of the following job behaviors in relation to script completion:
  - **Wait for the script to finish:** When selected, the job remains in the “running” state until the script is executed.
  - **Do not wait for the script to finish:** When selected, the job is completed even if the script execution is still in progress.
- **Error handling:** Select one of the following job behaviors in relation to script failure:
  - **Continue the job on script failure:** When selected, script failure does not influence the status of the job.
  - **Fail the job on script failure:** When selected and the script execution fails, the job status is

set to “failed” even if the recovery process is successful.

The screenshot shows the 'New Recovery Job Wizard for Physical Machines' interface. The wizard is on the '2. Destination' step. The 'Pre and Post Actions' section is highlighted with a red box. It contains three checkboxes: 'Send job run reports to' (unchecked), 'Run local pre job script' (checked), and 'Run local post job script' (checked). Other sections include 'Job Options', 'Recovered VM Options', and 'Data Transfer'.

## Data Transfer

### Transport Mode

NAKIVO Backup & Replication provides the following transport modes for writing VM data:

- **Hot Add only:** NAKIVO Backup & Replication can write data directly to the datastore bypassing the network, which can significantly increase the 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.
- **Automatic:** When this option is chosen, Hot Add mode is used where possible. If the product cannot use Hot Add, LAN mode is used.

### 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 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 that will 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” on page 284](#) 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:
    - a. Click the **Create New Rule** button.
    - b. The *New Bandwidth Rule* dialog opens. Refer to [“Bandwidth Throttling” on page 284](#) 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.

The screenshot shows the 'New Recovery Job Wizard for Physical Machines' interface, specifically the '3. Options' tab. The 'Data Transfer' section is highlighted with a red box. The settings in this section are:

- Transport mode: Automatic selection
- Transporter pool: Select transporter pool
- Transporters: Automatic selection
- Limit transporter load to: 3 concurrent tasks
- Bandwidth throttling: Disabled
- Bottleneck detection:

## Nutanix AHV VM Recovery

With NAKIVO Backup & Replication, you can recover full VMs from backups. Multiple VMs can be recovered by a single recovery job. The system allows selecting separate resources (clusters, storage containers, and networks) for different VM.

### Important

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:

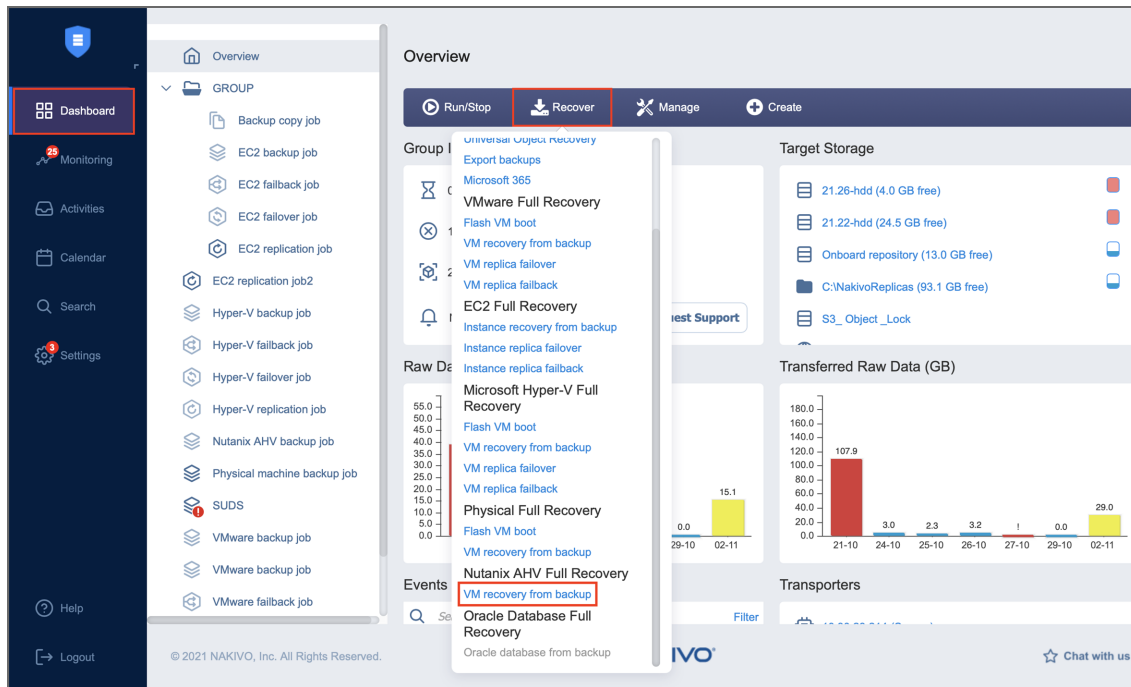
- [“Starting Nutanix AHV VM Recovery” below](#)
- [“Recovery Job Wizard for Nutanix AHV: Backups” on page 947](#)
- [“Recovery Job Wizard for Nutanix AHV: Destination” on page 947](#)
- [“Recovery Job Wizard for Nutanix AHV: Options” on page 950](#)

## Starting Nutanix AHV VM Recovery

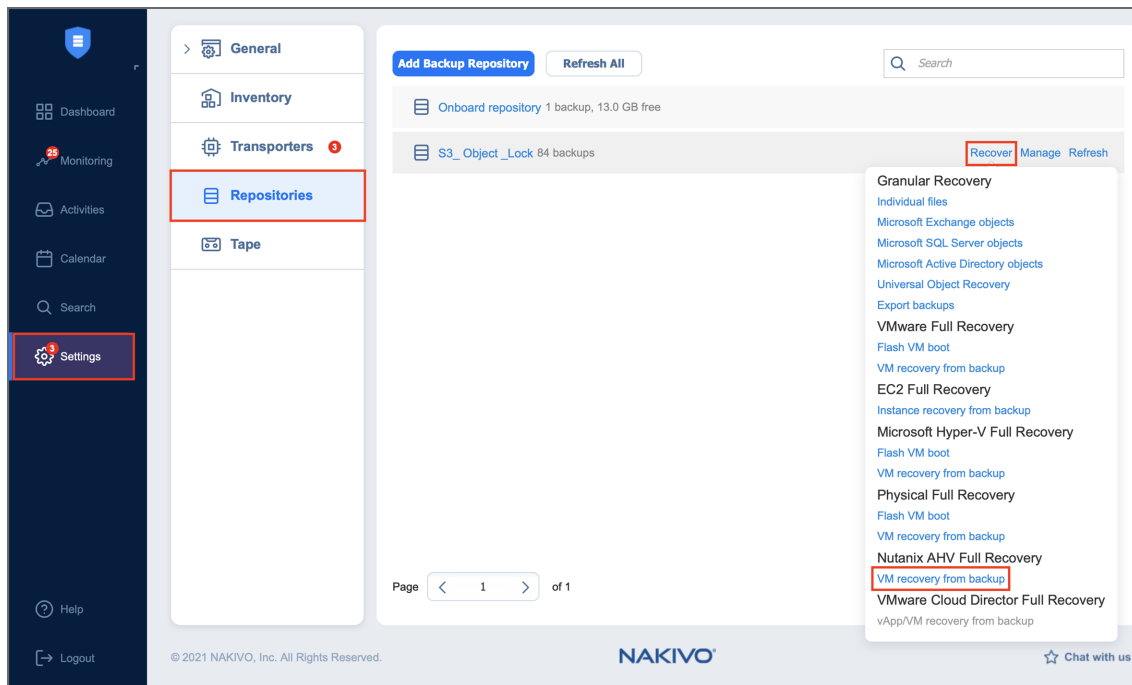
To recover entire Nutanix AHV VMs from backups, do one of the following:

- Start recovery from the **Dashboard** by clicking **Recover** and then clicking **VM recovery from backup** under the **Nutanix AHV Full Recovery** section.





- Open the New Recovery Job Wizard from the **Repositories** tab in **Settings** by following the steps below:
  1. Click **Settings** in the main menu of NAKIVO Backup & Replication.
  2. Go to the **Repositories** tab.
  3. Hover over the Backup Repository containing the necessary backup, click **Recover** and select **VM recovery from backup**.



- Alternatively, the recovery can be performed by using the [search](#) function. The new **Recovery Job Wizard for Nutanix AHV** opens.

## Recovery Job Wizard for Nutanix AHV: Backups

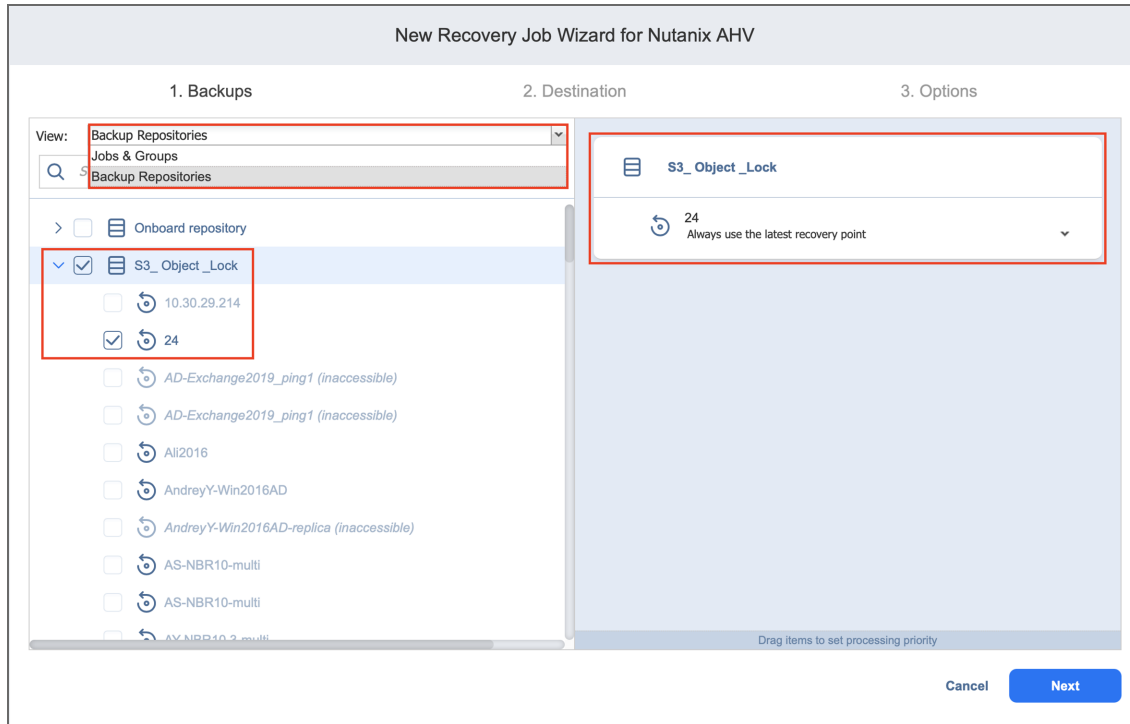
In the first step of the wizard, select one or more VM backups in the left pane and then select a recovery point for each backup in the right pane.

Use the **View** drop-down list to find the backups that you want to recover from:

- **Backup Repositories:** Displays all available repositories and backups stored within them.
- **Jobs & Groups:** Displays all jobs/groups and backups within them.

You can also search for a specific backup, job or container by entering its name into the **Search** field.

Click **Next** to go to the next step.



## Recovery Job Wizard for Nutanix AHV: Destination

On the **Destination** page, select a location for the recovered VMs. Proceed as described in these sections:

- [“Setting the Same Cluster, Datastore and Network for All Recovered VMs” below](#)
- [“Setting Different Destination Options for Recovered VMs” on the next page](#)

### Setting the Same Cluster, Datastore and Network for All Recovered VMs

To recover all VMs to the same cluster or storage container, and to connect all recovered VMs to the same networks, follow the steps below:

1. From the **Cluster** drop-down list, select a cluster.
2. From the **Datastore** drop-down list, select a datastore.

3. From the **Network** drop-down list, select a network.

New Recovery Job Wizard for Nutanix AHV

1. Backups      2. Destination      3. Options

Cluster: Nutanix

Datastore: NutanixManagementShare

Network: vlan0

[Advanced setup...](#)

Cancel    Next

With this setup, all VMs are recovered to the same cluster and storage container and they use the same network.

### Setting Different Destination Options for Recovered VMs

To specify different options for recovered VMs, follow the steps below:

1. Click **Advanced setup**.
2. Click on the VM to expand its recovery options.

3. Choose a target cluster, disks and network adapters for each VM.

New Recovery Job Wizard for Nutanix AHV

1. Backups      2. Destination      3. Options

Datstore:

Network:

---

S3\_Object\_Lock Click to collapse

Default cluster:  ⓘ

Default storage container:  ⓘ

Default network:  ⓘ

---

24 Click to collapse

Source		Target	
VM location:	S3_Object_Lock	Container:	<input type="text" value="Nutanix"/> ▼
VM resources:	2 CPU, 2.0 GB RAM	Virtual Machine:	<i>New VM will be created</i> ▼
Disks		Disks	
scsi.0:	S3_Object_Lock (20.0 GB)	scsi.0:	<input type="text" value="NutanixManagementShare"/> ▼
Network adapters		Network adapters	
24:		24:	<input type="text" value="vlan0"/> ▼

Cancel

If you selected groups or Backup Repositories on the **Backups** page, the advanced options allow you to specify the default destination options for all VMs contained in those groups or repositories.

The screenshot shows the 'New Recovery Job Wizard for Nutanix AHV' interface. At the top, there are three tabs: '1. Backups', '2. Destination', and '3. Options'. The '1. Backups' tab is active. Below the tabs, there are three dropdown menus: 'Cluster' (Nutanix), 'Datastore' (NutanixManagementShare), and 'Network' (vian0). Below these, there is a section titled 'S3\_Object\_Lock' with a 'Click to collapse' link. This section contains three dropdown menus: 'Default cluster' (Nutanix), 'Default storage container' (NutanixManagementShare), and 'Default network' (vian0). Below these, there is a '24' icon. At the bottom right, there are 'Cancel' and 'Next' buttons.

You can still define destination options individually for each VM within every group or repository by clicking the VM's name.

## Recovery Job Wizard for Nutanix AHV: Options

On the **Options** page, set up the recovery job options.

- [“General Options” on the next page](#)
- [“Recovered VM Options” on the next page](#)
- [“Pre and Post Actions” on page 952](#)
  - [“Email Notifications” on page 952](#)
  - [“Setting Up a Pre Job Script” on page 952](#)
  - [“Setting Up a Post Job Script” on page 953](#)
- **Data Transfer**
  - **Transporter Pool**
  - [“Transporters” on page 954](#)
  - [“Bandwidth Throttling” on page 954](#)
  - **Bottleneck Detection**
- [“Completing the New Recovery Job Wizard for Nutanix AHV” on page 956](#)

## General Options

Specify the general options as follows:

1. **Job name:** Specify a name for the recovery job.
2. **Network acceleration:** When [network acceleration](#) is enabled, NAKIVO Backup & Replication uses compression and traffic reduction techniques to speed up data transfer. Enable this option if recovering VMs over WAN or slow LAN links.
3. **Encryption:** When [encryption](#) is enabled, VM data is protected with AES 256 encryption while traveling over the network. Data encryption increases the backup time and CPU load on the machines running Transporters. Select this option when recovering over WAN without a VPN connection.

The screenshot shows the 'New Recovery Job Wizard for Nutanix AHV' interface, specifically the '3. Options' step. The 'Job Options' section is highlighted with a red box and contains the following fields:

- Job name: Nutanix AHV recovery job
- Network acceleration: Disabled
- Network encryption: Disabled

Other sections visible in the wizard include:

- Recovered VM Options:**
  - Recovery mode: Synthetic
  - Recovered VM names: Append "-recovered" in the end
  - VM power on: Power on recovered VMs
  - VM MAC addresses: Do not generate new MAC address
- Pre and Post Actions:**
  - Send job run reports to: [empty field]
  - Run local pre job script: [checkbox]
  - Run local post job script: [checkbox]
- Data Transfer:**
  - Transporter pool: Select transporter pool
  - Transporters: Automatic selection
  - Link transporter load to concurrent tasks: [checkbox]

At the bottom right, there are buttons for 'Cancel', 'Finish', and 'Finish & Run'.

## Recovered VM Options

Specify the recovery options as follows:

1. **Recovery mode:** Choose one of the following:
  - **Synthetic:** With the synthetic recovery mode selected, 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 the production recovery mode selected, environment dependencies are preserved on the recovered VM(s). Make sure the location where the VM(s) will be recovered does not contain the original VM(s), 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 "-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:** Allows you to enter custom names for recovered VMs.
3. **VM power on:** When the **Power on recovered VMs** option is chosen, the recovered VMs are powered on.

#### 4. VM MAC addresses: Choose one of the following:

- **Do not generate new MAC addresses:** When this option is chosen, the recovered VM will have the same MAC address as the source VM.
- **Generate new MAC addresses:** When this option is chosen, a new MAC address will be generated for the recovered VM.

The screenshot shows the 'New Recovery Job Wizard for Nutanix AHV' interface, specifically the '3. Options' step. The 'Recovered VM Options' section is highlighted with a red box. It contains the following settings:

- Job name: Nutanix AHV recovery job
- Network acceleration: Disabled
- Network encryption: Disabled
- Recovery mode: Synthetic
- Recovered VM names: Append "-recovered" in the end
- VM power on: Power on recovered VMs
- VM MAC addresses: Do not generate new MAC address

Other sections visible include 'Pre and Post Actions' with checkboxes for 'Send job run reports to', 'Run local pre job script', and 'Run local post job script'. The 'Data Transfer' section includes 'Transporter pool' (Select transporter pool) and 'Transporters' (Automatic selection). At the bottom right, there are 'Cancel', 'Finish', and 'Finish & Run' buttons.

## 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 where the Director is installed.

## 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.

To enable this option, make sure that [Email settings](#) are configured.

To send email notifications, in the *Pre and Post Actions* section:

1. Select the **Send job run reports to** option.
2. Specify one or more email addresses in the text box. 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. 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 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, VM backup only starts 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 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 will recover VMs even if the script has failed.
  - **Fail the job on script failure:** With this option selected, if the script fails, the job will be failed and VM recovery will not be performed.

### 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. Select the **Run local post job script** option and click the **settings** link.

3. Specify the following parameters in the dialog that appears:

- **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 recovery has been successful.

The screenshot shows the 'New Recovery Job Wizard for Nutanix AHV' interface. The wizard is on the '2. Destination' step. The 'Pre and Post Actions' section is highlighted with a red box, showing three checkboxes: 'Send job run reports to' (checked), 'Run local pre job script' (unchecked), and 'Run local post job script' (unchecked). Other sections include 'Job Options', 'Recovered VM Options', and 'Data Transfer'.

## 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](#).

## 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 also specify manually which Transporters should be used for a job.

From the **Transporters** drop-down list, select one of the following options:

- **Automatic selection:** The product automatically determines which Transporters 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 cluster:** Select this option to manually specify Transporters for all source and target clusters.

## Bandwidth Throttling

Please follow the steps below to regulate 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” on page 284](#) 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:
    - a. Click the **Create New Rule** button.
    - b. The **New Bandwidth Rule** dialog box opens. Refer to the [Bandwidth Throttling](#) 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 necessary 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 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.

The screenshot shows the 'New Recovery Job Wizard for Nutanix AHV' dialog box, specifically the 'Options' tab. The 'Data Transfer' section is highlighted with a red box and contains the following settings:

- Transporter pool: Select transporter pool
- Transporters: Automatic selection
- Limit transporter load to: 3 concurrent tasks
- Bandwidth throttling: Disabled
- Bottleneck detection

At the bottom of the dialog, there are three buttons: 'Cancel', 'Finish', and 'Finish & Run'.

## Completing the New Recovery Job Wizard for Nutanix AHV

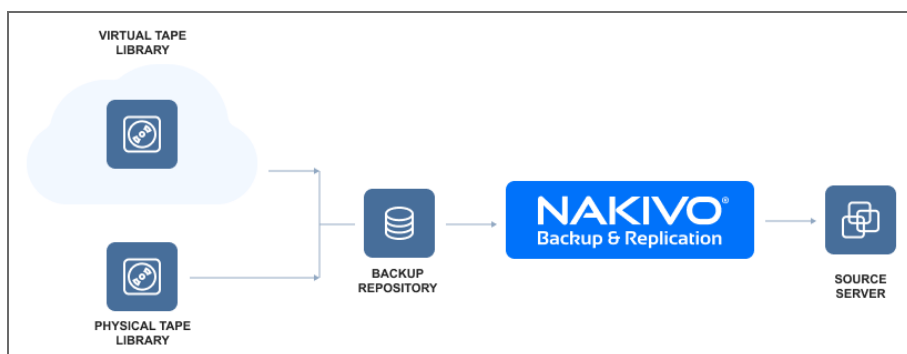
Click **Finish** or **Finish & Run** to complete the job creation.

### Note

If you click **Finish & Run**, you will have to define the scope of your job. Refer to [“Running Jobs on Demand” on page 256](#) for details.

## Recovery From Tape

To recover backups from tape, move the backed up data from a tape cartridge to a backup repository. Once the data is in the repository, you can restore the contents using the standard NAKIVO Backup & Replication tools.



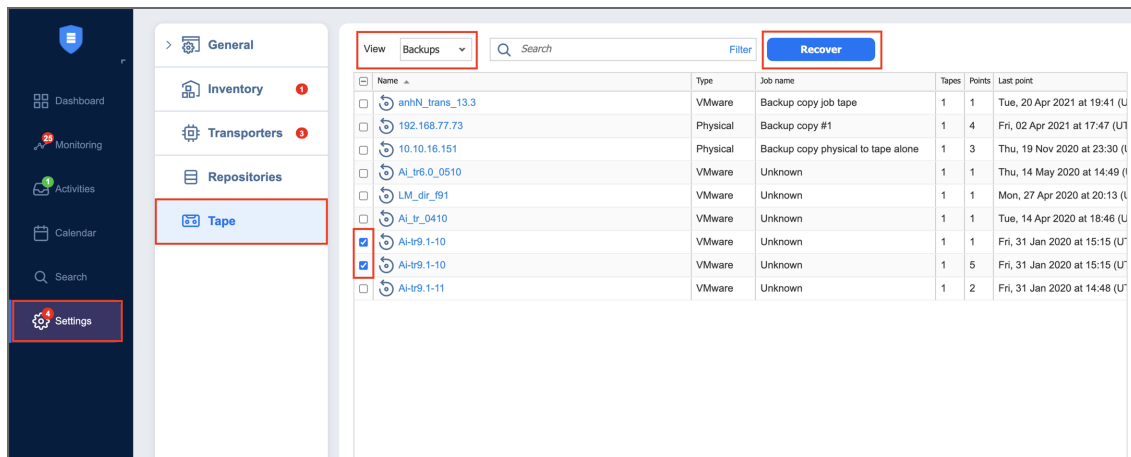
Refer to the following topics for more information:

- [“Starting Recovery from Tape Wizard” below](#)
- [“Recovery from Tape Wizard: Backups” on the next page](#)
- [“Recovery from Tape Wizard: Destination” on page 958](#)
- [“Recovery from Tape Wizard: Options” on page 959](#)

## Starting Recovery from Tape Wizard

To launch the recovery from the Tape wizard, do the following:

1. Go to **Settings**, click the **Tape** tab, and select **Backups** from the **View** drop-down list.



2. In the **Backups** table, do one of the following:

- Select the checkbox next to one or several backups that you want to recover and click the **Recover** button. 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.

- Click the name of the backup to go to the **Recovery** of the **Tape Cartridge Management** page where you can launch the Recovery wizard.

The **New Tape Recovery Job Wizard** opens.

3. Alternatively, go to the [“Managing Tape Cartridges” on page 461](#) page, select a backup in the **Tape contents** pane and then click the recovery point you want to restore from.

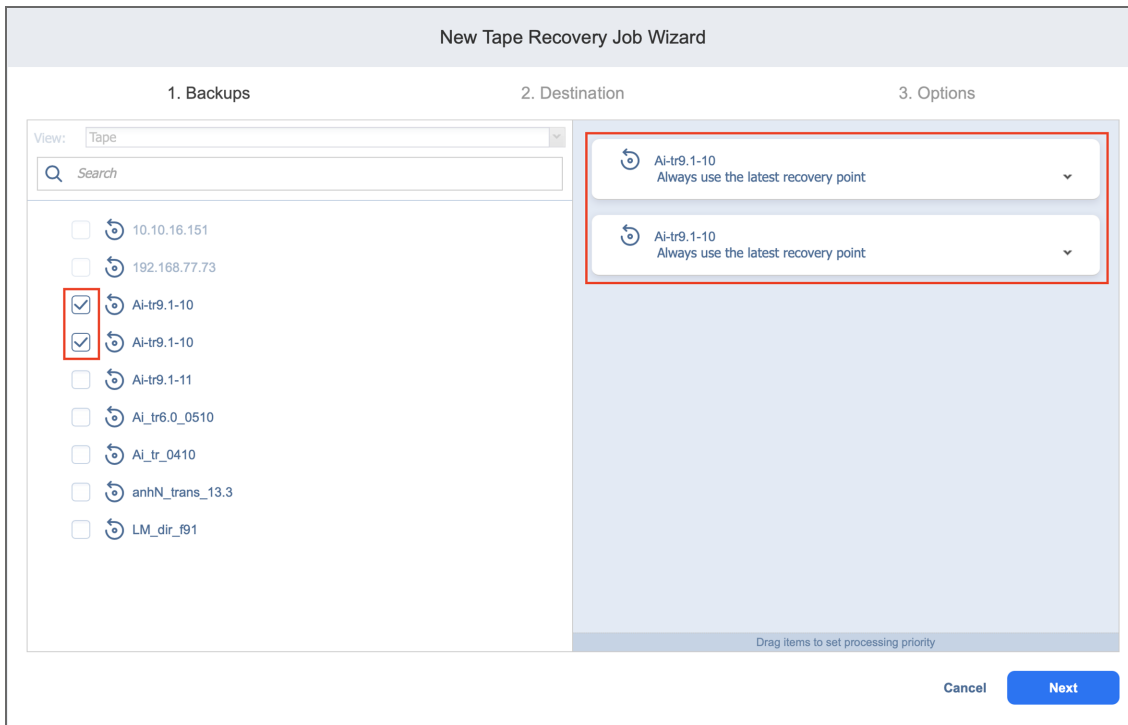
## Recovery from Tape Wizard: Backups

The first page of the Recovery Wizard is **Backups**. The number of backups and recovery points present in the table depends on the backups and recovery points you selected when launching the wizard. However, during this step, you can add or delete the backups and select different recovery points of the same type (hypervisor). You can also search for backups by entering a name (or part of it) into the Search box and group the backups by media pools, device locations, or tape devices.

#### Note

- If a selected recovery point of the job object is a full recovery point, NAKIVO Backup & Replication will recover the selected recovery point.

- If a selected recovery point of the job is incremental, NAKIVO Backup & Replication will recover the chain of recovery points starting with the full recovery point that is the ancestor to the selected recovery point and finishing with the selected incremental recovery point.



After you are done, click **Next**.

## Recovery from Tape Wizard: Destination

On the **Destination page**, you define the Backup Repository where the backup will be placed by selecting an option from the **Container** drop-down list. You can also select which VM disks to recover by clicking **Advanced options**.

New Tape Recovery Job Wizard

1. Backups
2. Destination
3. Options

Recovery to:

Container:

<span style="font-size: 1.2em;">↺</span> Ai-tr9.1-10 <span style="float: right; font-size: 0.8em;">Click to collapse</span>
VM disks <input checked="" type="checkbox"/> Hard disk 1 (0 KB (20.0 GB allocated))

<span style="font-size: 1.2em;">↺</span> Ai-tr9.1-10 <span style="float: right; font-size: 0.8em;">Click to collapse</span>
VM disks <input checked="" type="checkbox"/> Hard disk 1 (0 KB (20.0 GB allocated))

Cancel
Next

Click **Next** to proceed to the next page.

## Recovery from Tape Wizard: Options

- [“General Options” below](#)
- [“Pre and Post Actions” on the next page](#)
  - [“Setting Up a Pre-Job Script” on page 961](#)
  - [“Setting Up a Post Job Script” on page 962](#)
  - [“Email Notifications” on page 962](#)
- [“Data Transfer” on page 963](#)
  - [“Bandwidth Throttling” on page 964](#)
  - [“Multi-Channel Processing” on page 964](#)

### General Options

Specify the general options as follows:

1. **Job name:** Specify a name for the recovery job.
2. **Network acceleration:** When network acceleration is 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. For more information, refer to [“Network Acceleration” on page 63](#).

3. **Encryption:** When encryption is 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 recovering over WAN without a VPN connection. For details, refer to [“Encryption in Flight and at Rest” on page 42.](#)

New Tape Recovery Job Wizard

1. Backups      2. Destination      3. Options

**Job Options**

Job name:

Network acceleration:  ⓘ

Network encryption:  ⓘ

**Pre and Post Actions**

Send job run reports to  ⓘ

Run local pre job script ⓘ

Run local post job script ⓘ

**Data Transfer**

Bandwidth throttling:  ⓘ

Use multi-channel processing  channels per disk ⓘ

Cancel **Finish** Finish & Run

## 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 where the Director is installed. Refer to [“Pre and Post Job Scripts” on page 68](#) for details.

New Tape Recovery Job Wizard

1. Backups                      2. Destination                      3. Options

Job Options

Job name:

Script path:

Job behavior:

Error handling:

Run local pre job script [settings](#)

No path was specified; wait for the script to finish; continue the job on script failure

Run local post job script

Data Transfer

Bandwidth throttling:

Use multi-channel processing  channels per disk

Cancel    **Finish**    Finish & Run

## 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. Select the **Run local pre job script** option and click the **settings** link.
3. Specify the following parameters in the dialog that appears:
  - **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, VM recovery is only started 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 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 will perform VM recovery even if the script has failed.
    - **Fail the job on script failure:** With this option selected, if the script fails, the job will be failed and VM recovery will not be performed.



## Setting Up 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 where 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 that appears:
  - **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, VM recovery is only started 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 VMs at the same time.
  - **Error handling:** Choose one of the following job behaviors in relation to scrip failure:
    - **Continue the job on script failure:** With this option selected, the job will perform VM recovery even if the script has failed.
    - **Fail the job on script failure:** With this option selected, if the script fails, the job will be failed and VM recovery will not be performed.

## Email Notifications

NAKIVO Backup & Replication can send email notifications about job completion status to specified recipients. This feature complements the global notifications feature and allows you to configure notifications on a per-job level.

To send email notifications, select the **Send job run reports to** option in the **Pre and Post Actions** section and specify one or more email addresses in the text field. The semicolon character should be used to separate multiple email addresses. To enable this option, make sure that your email setting are configured. Refer to [“Notifications & Reports” on page 294](#) for details.

New Tape Recovery Job Wizard

1. Backups                      2. Destination                      3. Options

---

**Job Options**

Job name:

Network acceleration:  ⓘ

Network encryption:  ⓘ

**Pre and Post Actions**

Send job run reports to  ⓘ

Run local pre job script ⓘ

Run local post job script ⓘ

**Data Transfer**

Bandwidth throttling:  ⓘ

Use multi-channel processing  channels per disk ⓘ

---

Cancel   **Finish**   Finish & Run

## Data Transfer

In the **Data Transfer** section of the **Options** page, you can set or configure bandwidth throttling rules and multi-channel processing.

New Tape Recovery Job Wizard

1. Backups                      2. Destination                      3. Options

---

**Job Options**

Job name:

Network acceleration:  ⓘ

Network encryption:  ⓘ

**Pre and Post Actions**

Send job run reports to  ⓘ

Run local pre job script ⓘ **settings:**  
No path was specified; wait for the script to finish; continue the job on script failure

Run local post job script ⓘ

**Data Transfer**

Bandwidth throttling:  ⓘ **settings**

Use multi-channel processing  channels per disk ⓘ

---

Cancel   **Finish**   Finish & Run

## 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 284](#) for details.

2. Click the **Settings** link that becomes available.
3. The **Job Bandwidth Rules** dialog opens displaying you 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 opens. Refer to the [“Bandwidth Throttling” on page 284](#) 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 necessary 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.

## 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.

### Note

The Transporter needs to have at least 2 CPU cores and 8 GB RAM available to perform recovery in multiple channels.

## 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 virtual machine(s)
- VHD for disk(s) of Hyper-V virtual machine(s)
- VHDX for disk(s) of Hyper-V virtual machine(s)

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 124](#) 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.

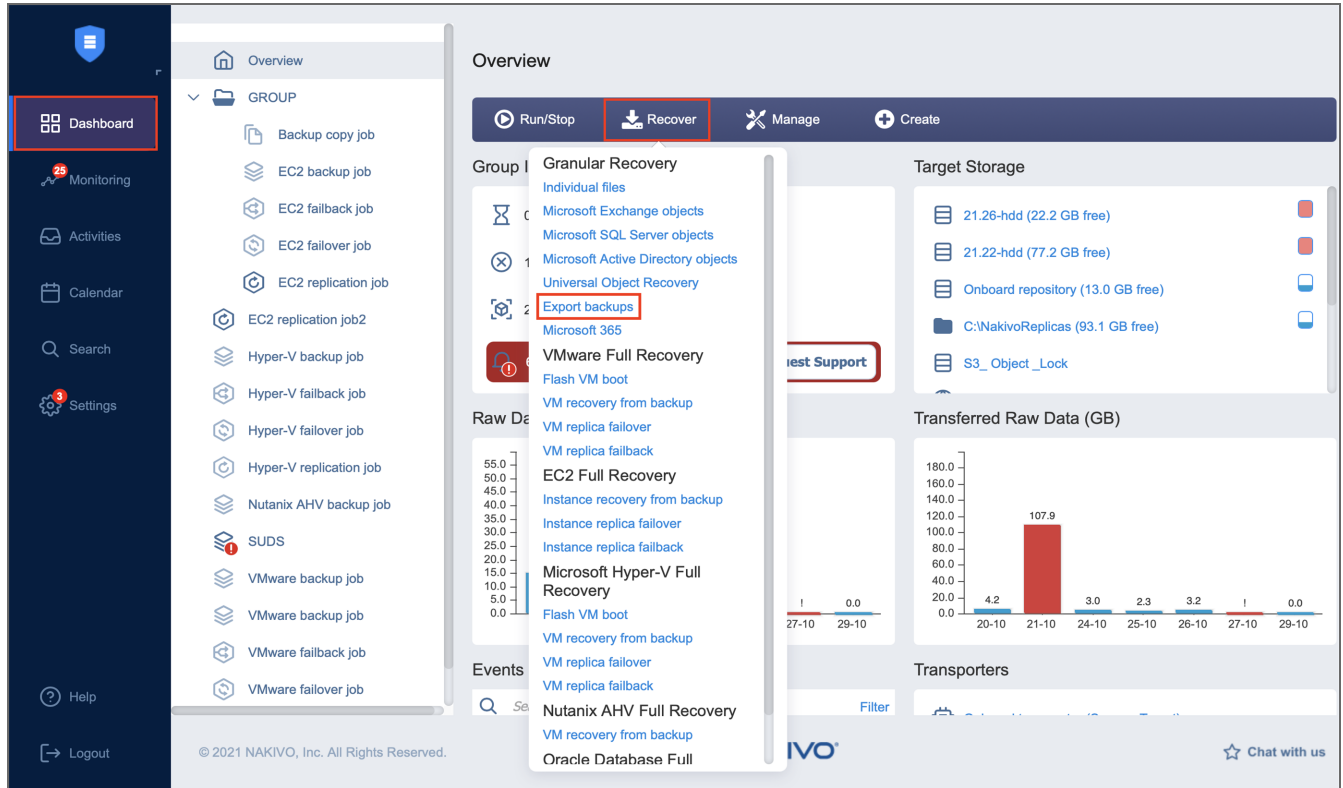
This section includes the following topics:

- [“Opening Backup Export Wizard” below](#)
- [“Backup Export Wizard: Backups” on page 967](#)
- [“Backup Export Wizard: Disks” on page 968](#)
- [“Backup Export Wizard: Options” on page 969](#)
- [“Backup Export Wizard: Finish” on page 970](#)

## Opening Backup Export Wizard

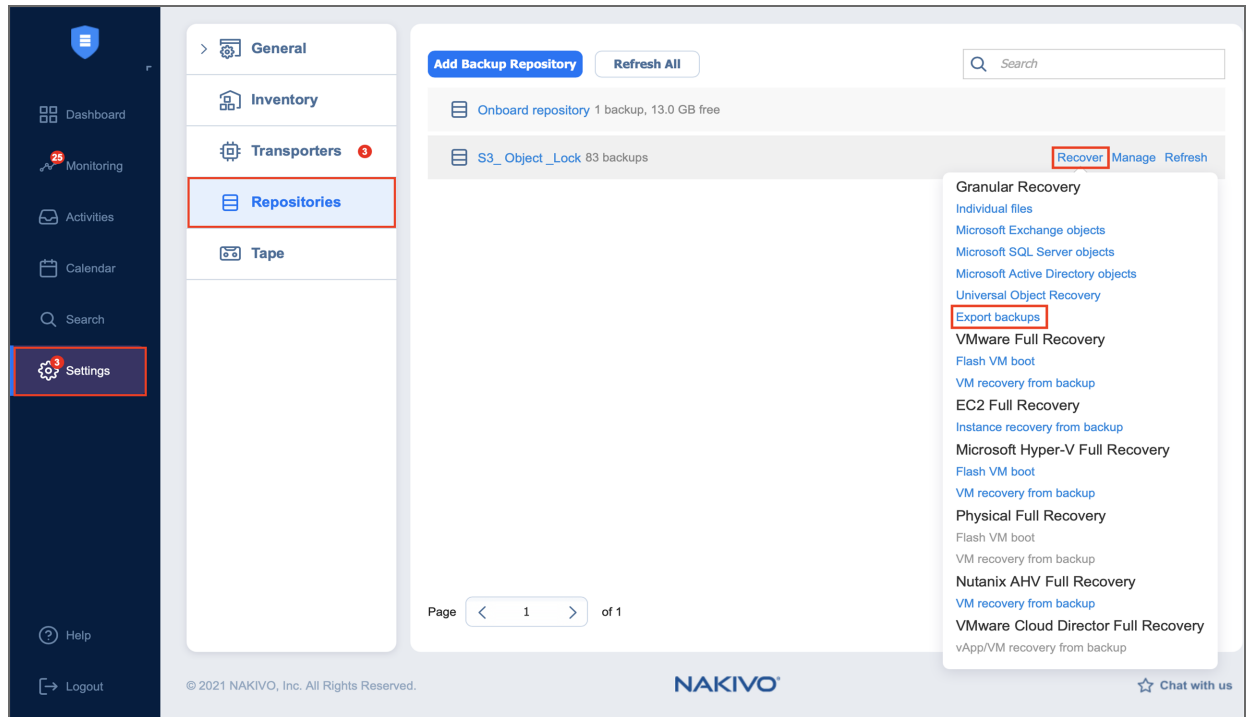
Open **Backup Export Wizard** using one of the following ways:

- Navigate to the **Dashboard**, 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 **Recover**.

3. In the menu that opens, click **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 one of these views:
  - **Jobs & Groups**
  - **Backup Repositories**
2. Select a recovery point for each backup in the right pane.


3. Click **Next** to go to the next page of the wizard.



## 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
  - 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

1. Backups      2. Disks      3. Options      4. Finish

Export format: VMDK

Export location: CIFS share

Path to the share: \\Path\Folder **Test Connection**

Credentials type: Password

Username: admin

Password: ••••••••

[Manage credentials](#)

Cancel      **Export**

## 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

NAKIVO Backup & Replication allows you to address all major [disaster recovery](#) (DR) planning points by creating automated DR workflows for VMware, Microsoft Hyper-V, and Amazon EC2 environments. The application allows you to protect VMs running within a cluster, replicate VMs, failover to replicas, and replica failback.

When using [Site Recovery](#), you can include up to 200 actions in a single job, including failover, failback, start/stop VMs and instances, run/stop jobs, run script, attach or detach repository, send an email, wait, and check condition. By arranging actions and conditions into one automated algorithm, you can create disaster recovery jobs of any complexity.

For more details, refer to the corresponding articles below:

- [“Failover to Replica for VMware” below](#)
- [“Failover to Replica for Microsoft Hyper-V” on page 978](#)
- [“Failover to Replica for Amazon EC2” on page 985](#)
- [“Replica Failback for VMware vSphere” on page 991](#)
- [“Replica Failback for Microsoft Hyper-V” on page 1002](#)
- [“Replica Failback for Amazon EC2” on page 1012](#)
- [“Site Recovery Job” on page 1018](#)

## 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:

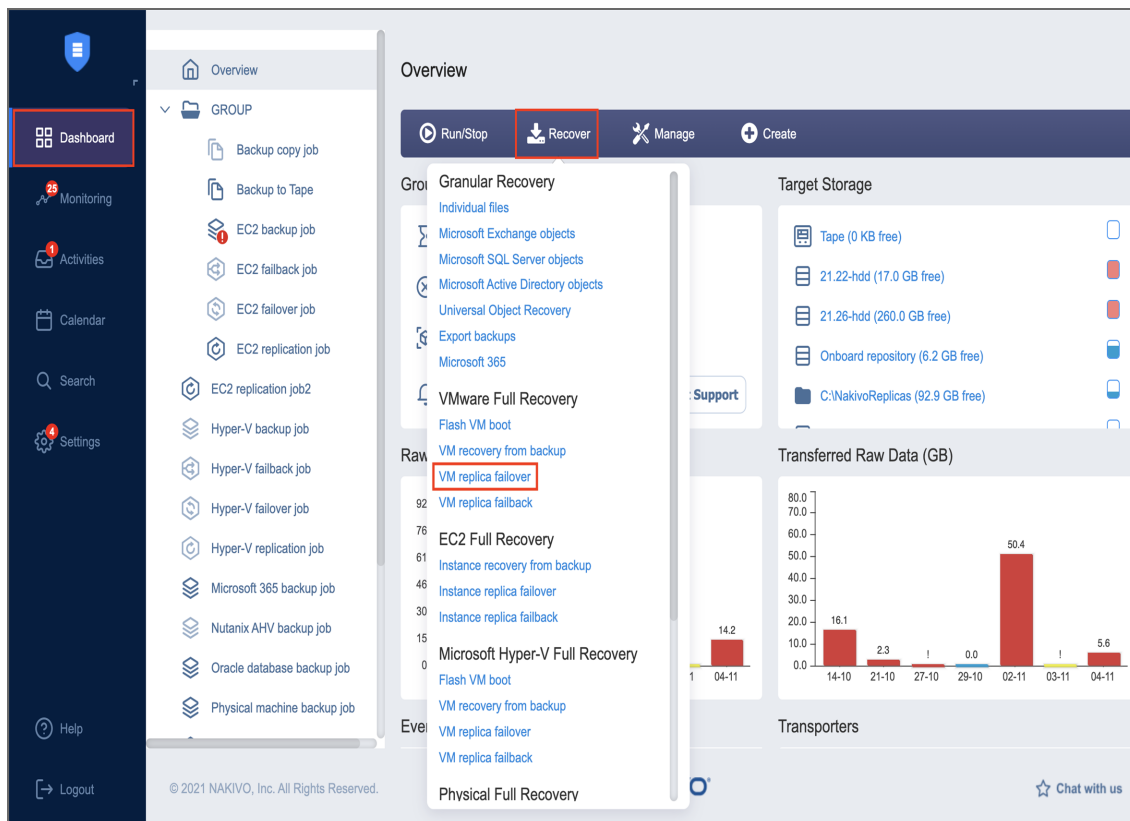
- [“Starting VMware Failover to Replica” below](#)
- [“Failover Job Wizard for VMware: Source” on the next page](#)
- [“Failover Job Wizard for VMware: Networks” on page 973](#)
- [“Failover Job Wizard for VMware: Re-IP” on page 975](#)
- [“Failover Job Wizard for VMware: Options” on page 977](#)

## Starting VMware Failover to Replica

Before starting a replica failover job for VMware, make sure that you have a VMware replication job with at least one backup available. Refer to [“Creating VMware Replication Jobs” on page 714](#) for details.

To start the procedure of switching a VMware system workload to a backup VM, do the following:

1. Click **Recover** in the NAKIVO Backup & Replication **Dashboard**
2. Click **VM replica failover** 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 > VM replica failover**.
4. Additionally, the failover can be performed by using by using the [search](#) function.




The **New Failover Job Wizard for VMware vSphere** 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 266](#) 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 269](#) 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.

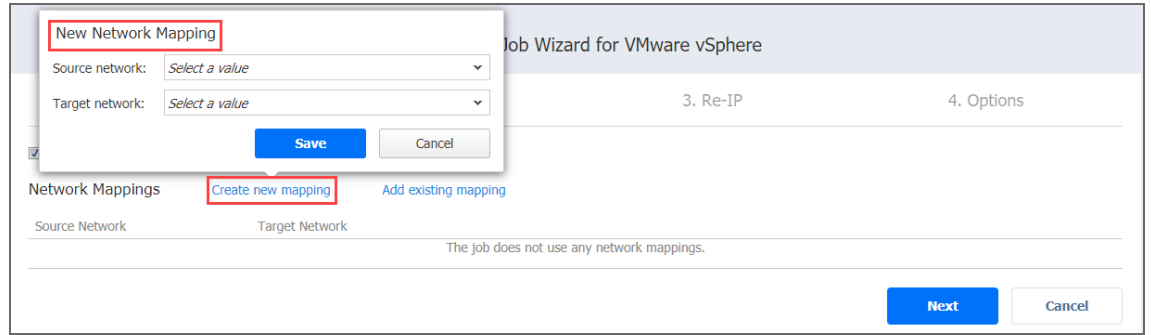
## Failover Job Wizard for VMware: Networks

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.


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 **Edit** button to the right of the item.
  - 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 X 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.

## 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 click the **Edit** button to the right of the item.
  - 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 click the X 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.

## 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 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 284](#).



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.

New Failover Job Wizard for VMware vSphere

1. Source      2. Networks      3. Re-IP      4. Options

Job Options

Job name:

Power off source VMs

Pre and Post Actions

Send job run reports to

Run local pre job script ?

Run local post job script ?

## Completing the New 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 256](#) for details.

## Failover to Replica for Microsoft Hyper-V

With NAKIVO Backup & Replication, you can switch your Hyper-V VM to the VM’s latest replica in case of failure.

Refer to the following topics for more information:

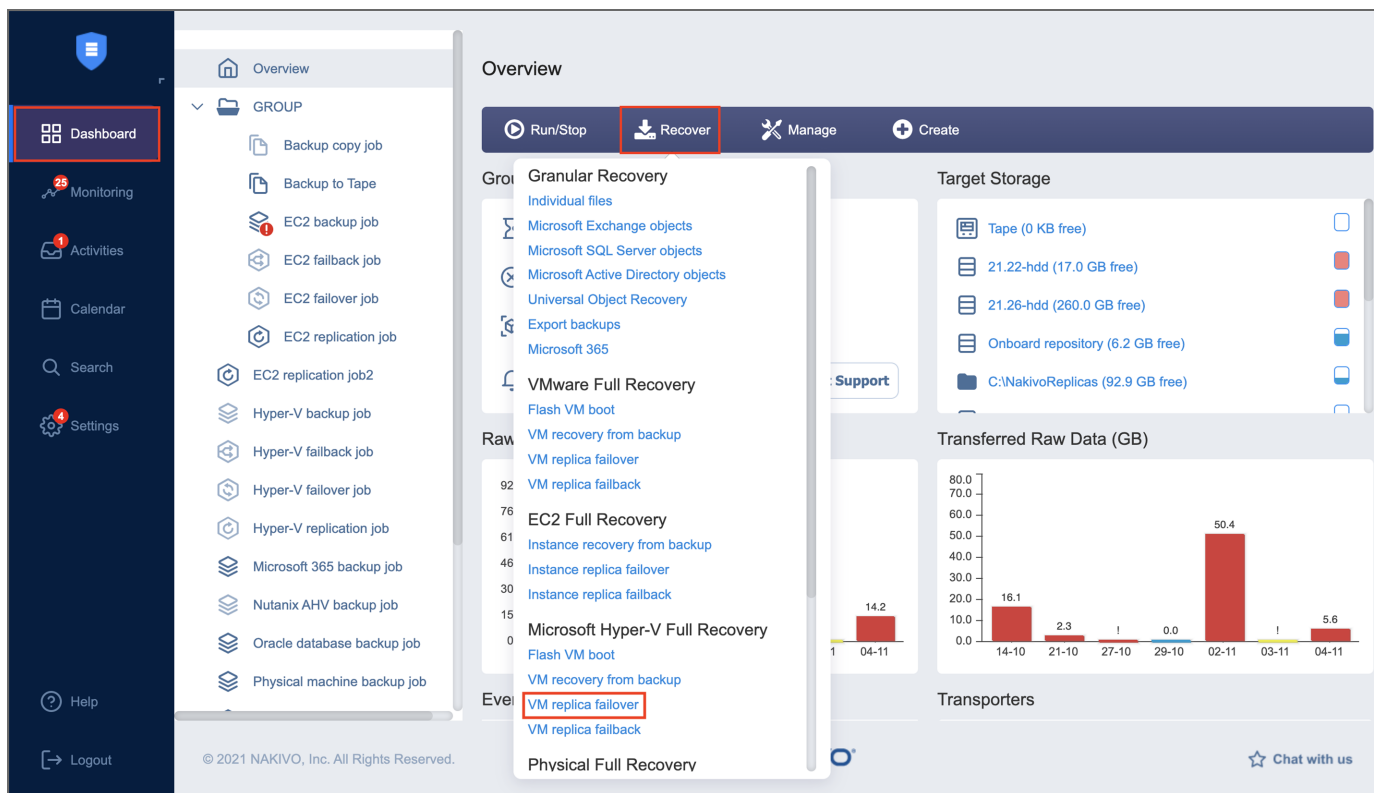
- [“Starting Microsoft Hyper-V Failover to Replica” below](#)
- [“Failover Job Wizard for Microsoft Hyper-V: Source” below](#)
- [“Failover Job Wizard for Microsoft Hyper-V: Networks” on the next page](#)
- [“Failover Job Wizard for Microsoft Hyper-V: Re-IP” on page 982](#)
- [“Failover Job Wizard for Microsoft Hyper-V: Options” on page 984](#)

## Starting Microsoft Hyper-V Failover to Replica

Before starting a replica failover job for Hyper-V, make sure that you have a Hyper-V replication job with at least one backup available. Refer to [“Creating Hyper-V Replication Jobs” on page 745](#) for details.

To start the procedure of switching a Hyper-V system workload to a backup VM, do the following:

1. Click **Recover** in the NAKIVO Backup & Replication **Dashboard**
2. Click **VM replica failover** in the **Microsoft Hyper-V Full Recovery** section. Note that the Hyper-V replication job must have been completed successfully for this option to be available.
3. Alternatively, select the relevant replication job and click **Recover > VM replica failover**.
4. Additionally, the failover can be performed by using by using the [search](#) function.

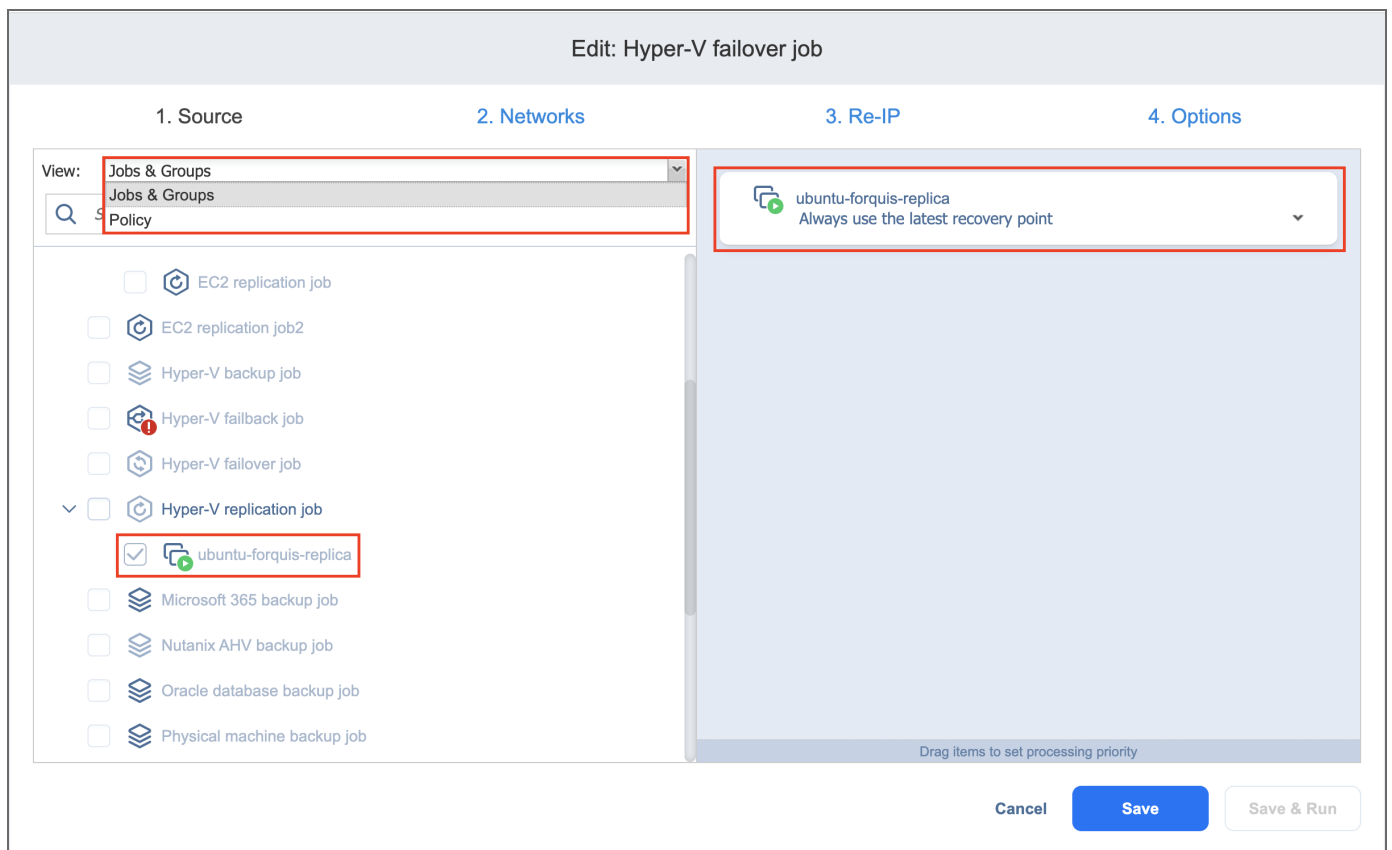


The **New Failover Job Wizard for Microsoft Hyper-V** opens.

## Failover Job Wizard for Microsoft Hyper-V: Source

On **Source** page of the wizard, do the following:

- 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 Hyper-V VM from the list of replication jobs.
  - Policy:** When selected, job policies can be used. Refer to [“Managing Job Policies” on page 266](#) 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 269](#) for details.
- Select one or more source Hyper-V VMs in the left pane and then select a recovery point for each source VM in the right pane.
- Click **Next** to go to the next page.




## Failover Job Wizard for Microsoft Hyper-V: Networks


When the replica network – or target Hyper-V virtual network – differs from the network address where the source VMs are deployed – or source Hyper-V virtual network – your failover job needs a relation between these networks to be performed successfully. This relation is called a network mapping.

To map source VM virtual networks to appropriate target virtual networks, do the following on the **Networks** page of the wizard:

1. Select **Enable network mapping**.
2. The **Network Mapping section** opens. The following options are available:
  - Create a new network 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 network mapping:
  - a. Click **Add existitng mapping**.
  - b. The **Network Mappings** dialog box opens. Choose an appropriate network mapping and close the dialog box.



- Edit an existing network mapping:
  - a. Hover over the required item in the **Network Mappings** list and click **Edit** to the right of the item.

- b. 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 over the required item in the **Network Mappings** list and click the "X" 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.

## Failover Job Wizard for Microsoft Hyper-V: Re-IP

When the IP addressing scheme for the target replica differs from the IP addressing scheme for the Hyper-V source VMs, your failover job needs to have relations between your source VMs addresses and target replica VMs addresses so that the job is performed successfully. Such relations are called Re-IP rules.

### Warning

Hyper-V Integration Services must be running on source VMs to enable successful 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 and close the dialog box.
4. The following options are available:
  - 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.

New Failover Job Wizard for Microsoft Hyper-V

1. Source      2. Networks      3. Re-IP      4. Options

Enable Re-IP ?    [Select VMs](#)

Re-IP Rules    [Create new rule](#)    [Add existing rule](#)

**New Re-IP Rule**

Source Settings

IP address:

Subnet mask:

Target Settings

IP address:

Subnet mask:

Default gateway:

Primary DNS server:

Secondary DNS server:

DNS suffix:

## Note

You can use wildcard characters 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. If several re-IP rules are applicable to your source VM, the product will define the most suitable rule 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 box.

New Failover Job Wizard for Microsoft Hyper-V

1. Source      2. Networks      3. Re-IP      4. Options

Enable Re-IP ?    [Select VMs](#)

Re-IP Rules    [Create new rule](#)    [Add existing rule](#)

Source IP:

Search

Source IP Address	Target IP Address
<input checked="" type="checkbox"/> 10.30.30.55	10.30.30.56

- Edit an existing re-IP rule:

- a. Hover over the required item in the **Re-IP Rules** list and click the **Edit** button to the right of the item.
  - b. The **Edit Re-IP Rule** dialog box opens. Edit the required properties of the re-IP rule and click **Save**.
- Delete an existing rule: Hover over the required item in the **Re-IP Rules** list and click the **X** icon to the right of the item.

The screenshot shows the 'New Failover Job Wizard for Microsoft Hyper-V' at the '3. Re-IP' step. It features a progress bar with four steps: 1. Source, 2. Networks, 3. Re-IP (current), and 4. Options. Below the progress bar, there is a checkbox for 'Enable Re-IP' which is checked, and a 'Select VMs' link. Under 'Re-IP Rules', there are links for 'Create new rule' and 'Add existing rule'. A table lists Re-IP rules with columns for 'Source IP Address' and 'Target IP Address'. One rule is shown with Source IP 10.30.30.55 and Target IP 10.30.30.56. To the right of this rule, there are icons for editing (a pencil) and deleting (an 'X'), both enclosed in a red box. At the bottom right, there are 'Next' and 'Cancel' buttons, with 'Next' being highlighted in blue.

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 Microsoft Hyper-V: Options

On the **Options** page of the wizard, specify the options for the Microsoft Hyper-V failover job.

- [Job Options](#)
- [Pre and Post Actions](#)
- [Completing the New Failover Job Wizard for Microsoft Hyper-V](#)

### Job Options

In this section, you can specify the following failover job options:

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

You can enable the following pre- and post-actions:

1. **Send job run reports to:** When selected, a job report is sent to the specified recipients each time the failover job has completed. This option overrides the default setting in the NAKIVO Backup & Replication [Settings](#).
2. **Run local pre-job script:** When selected, a dialog box opens. You can specify the following options of the [pre-job script](#) to be executed in this dialog box:
  - **Script path:** A local path to the script on the machine where the Director is installed. A script interpreter should also be specified.
  - **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 the options of the post-job script to be executed on the machine where the product is installed. Proceed as for the pre-job script.

The screenshot shows the 'New Failover Job Wizard for Microsoft Hyper-V' dialog box, specifically the '4. Options' step. The wizard has four steps: 1. Source, 2. Networks, 3. Re-IP, and 4. Options. In the 'Job Options' section, the 'Job name' is 'Hyper-V failover job'. The checkbox for 'Power off source VMs' is checked. In the 'Pre and Post Actions' section, the checkbox for 'Send job run reports to' is checked, with the email address 'administrator@nakivo.com' entered in the adjacent text box. The checkboxes for 'Run local pre job script' and 'Run local post job script' are unchecked. At the bottom right, there are three buttons: 'Finish' (highlighted in blue), 'Finish & Run', and 'Cancel'.

## Completing the New Failover Job Wizard for Microsoft Hyper-V

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](#) for details.

## Failover to Replica for Amazon EC2

With NAKIVO Backup & Replication, you can switch your EC2 instance to the EC2 instance’s latest replica in case of failure.

Refer to the following topics for more information:

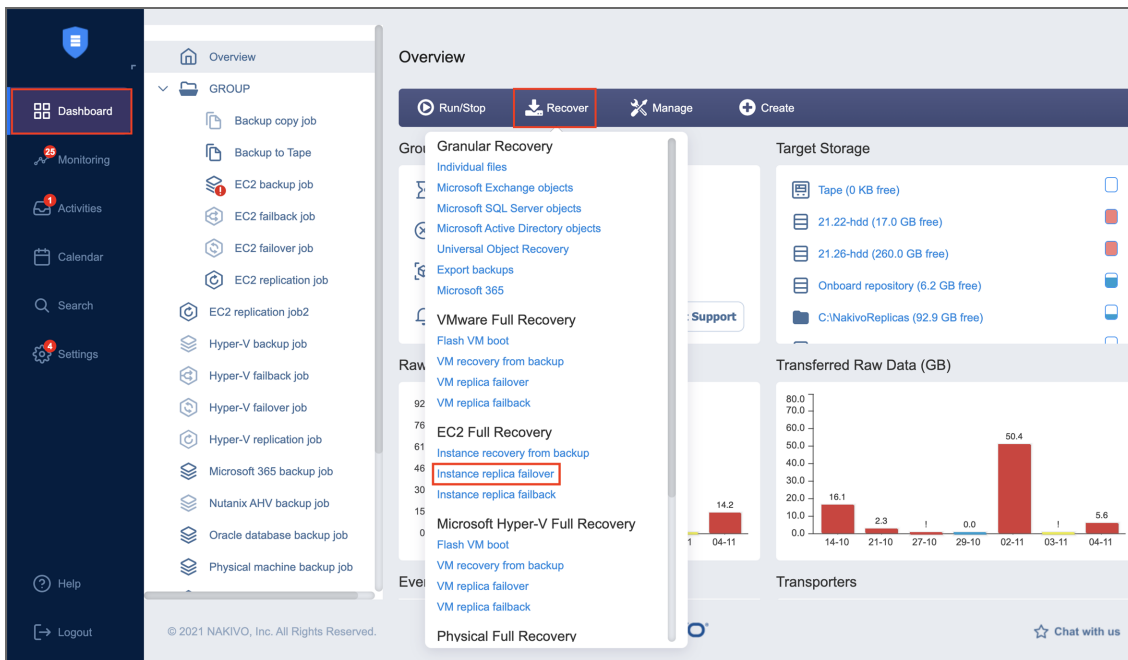
- [“Starting Amazon EC2 Failover to Replica” on the next page](#)
- [“Failover Job Wizard for Amazon EC2: Source” on the next page](#)
- [“Failover Job Wizard for Amazon EC2: Destination” on page 987](#)
- [“Failover Job Wizard for Amazon EC2: Options” on page 990](#)



## Starting Amazon EC2 Failover to Replica

Before starting a replica failover job for Amazon EC2, make sure that you have an EC2 replication job with at least one backup available. Refer to [“Creating Amazon EC2 Replication Jobs” on page 773](#) for details.

To start the procedure of switching an Amazon EC2 system workload to a backup EC2 Instance, click **Recover** in the NAKIVO Backup & Replication **Dashboard**, and then click **Instance replica failover** in the **EC2 Full Recovery** section. Alternatively, the failover can be performed by using the [search](#) function.

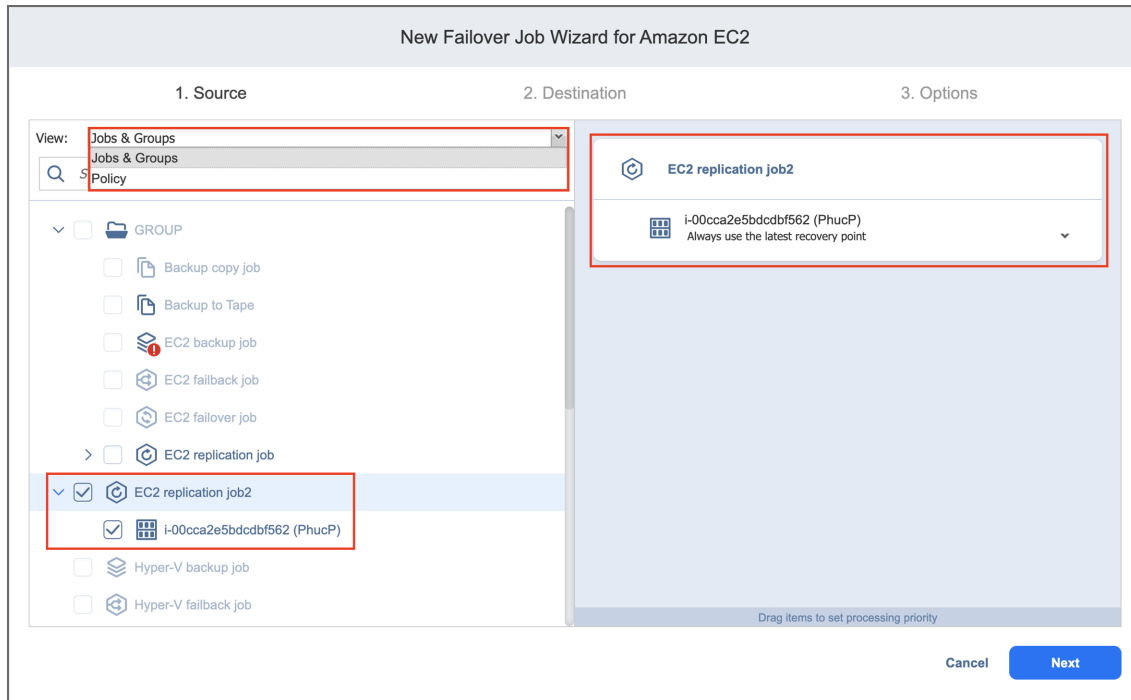


## Failover Job Wizard for Amazon EC2: Source

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 AWS account. 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 266](#) 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 269](#) for details.
2. Select one or more source Amazon EC2 instances from the inventory tree and then select a recovery point for each source Amazon EC2 instance in the right pane.

3. Click **Next** to go to the next page of the wizard.



## Failover Job Wizard for Amazon EC2: Destination

- [“Single Destination for Amazon EC2 Instances” below](#)
- [“Different Destinations for Amazon EC2 Instances” on the next page](#)

On the **Destination** page of the wizard, you can specify a failover destination for your Amazon EC2 instances.

### Important

You may be additionally charged for using a third-party resource. Refer to the 3rd-party resource provider documentation for details.

### Single Destination for Amazon EC2 Instances

To fail over all Amazon EC2 instances defined in the previous page in the wizard to the same AWS region and with the same options, set up the following:

1. **Network:** Select an Amazon EC2 [VPC](#) to which all Amazon EC2 instances will be connected.
2. **Subnet:** Select an Amazon EC2 [Subnet](#) to which all Amazon EC2 instances will be connected.
3. **Security group:** Select an Amazon EC2 [Security Group](#) for all Amazon EC2 instances.
4. **Instance access:** Specify Amazon EC2 [Key Pair](#) options by choosing one of the following:

- **Use existing key pair:** An existing Amazon EC2 Key Pair will be used to connect to the recovered instances.
- **Create a new key pair:** A new Amazon EC2 Key Pair will be created and assigned to all recovered Amazon EC2 instances. Enter a name for the new Amazon EC2 Key Pair in the **Key pair name** box and click **Download**.

### Important

This is your only chance to save the private key file, so make sure you save the private key file in a safe place. You will need to provide the name of your key pair when you start recovered Amazon EC2 instances and the corresponding private key each time you connect to the instances.

- **Proceed without key pair:** No AWS Key Pair will be assigned to the recovered Amazon EC2 instances.

### Important

A key pair allows you to connect to Linux instances via SSH. If you recover Linux instances without a key pair, you will not be able to connect to them.

5. Click **Next** to go to the next page in the wizard.

The screenshot shows the 'New Failover Job Wizard for Amazon EC2' with three steps: 1. Source, 2. Destination, and 3. Options. Step 1 is active. The configuration options are:

- Network: NAKIVO vpc (10.0.0.0/16)
- Subnet: NAKIVO subnet (10.0.0.0/16)
- Security group: amazon EC2\_\_30a997e0-6ba5-48e1-9...
- Instance access: Use existing key pair
- Key pair: Select target key pair

A warning message is displayed: "Target region cannot be changed for replicas. You can only re-select networks and security groups to which replicas will be connected." At the bottom right, there are 'Cancel' and 'Next' buttons.

## Different Destinations for Amazon EC2 Instances

To set different destination options for individual Amazon EC2 instances, follow the steps below:

1. Click **Advanced options**.
2. Click an EC2 instance to expand it. Set the options as follows:

- a. **Region:** Select a target [region](#) for the Amazon EC2 instance.
- b. **Instance type:** Select a [type](#) for the recovered Amazon EC2 instance.
- c. **EBS Volumes:** Choose whether to recover or skip recovering individual AWS [EBS Volumes](#).
- d. **Network:** Select an Amazon EC2 [VPC](#) to which the Amazon EC2 instance will be connected.
- e. **Network adapter(s):** Specify the Amazon EC2 [network adapter](#) options for every network adapter:

- **Subnet:** Select an Amazon EC2 [subnet](#) to which the network adapter will be connected.
- **Security group:** Select an Amazon EC2 [security group](#) for the Amazon EC2 instance.
- **Primary IP:** Assign a valid primary IP to a current adapter.

**Note**

If a source instance has more network adapters than a target one, the target network adapters exceeding the supported number will be disabled.

- f. **Instance access:** Specify Amazon EC2 [Key Pair](#) options as follows:
  - **Use existing key pair:** Choose an existing Amazon EC2 Key Pair to use to connect to the recovered instance.
  - **Create a new key pair:** Select this option to create a new Amazon EC2 Key Pair and assign it to the recovered Amazon EC2 instance. Enter a name for the new Amazon EC2 Key Pair in the **Key pair name** box and click **Download**.


**Important**

This is your only chance to save the private key file, so make sure you keep it in a safe place. You will need to provide the name of your key pair when you launch the recovered EC2 instance and the corresponding private key each time you connect to the instance.

- **Proceed without key pair:** Select this option if you do not want to assign an AWS Key Pair to the recovered Amazon EC2 instance.

**Important**

A key pair enables you to connect to a Linux instance via SSH. If you recover Linux instances without a key pair, you will not be able to connect to them.



3. Click **Next** to go to the next page in the wizard.

## Failover Job Wizard for Amazon EC2: Options

On the **Options** page of the wizard, specify the options for the Amazon EC2 failover job.

- [“Job Options” below](#)
- [“Pre and Post Actions” below](#)
- [“Completing the New Failover Job Wizard for Amazon EC2” on the next page](#)

### Job Options

You can specify the following failover job options:

1. **Job name:** A string of 50 or fewer characters specifying the name of the failover job.
2. **Power off source instances:** Select this checkbox to power off source Amazon EC2 instances when the failover job has completed.

### Pre and Post Actions

You can enable the following pre- and post-actions:

1. **Send job run reports to:** When selected, a job report is sent to the specified recipients each time the failover job has completed. This option overrides the default setting in the NAKIVO Backup & Replication **Settings**.
2. **Run local pre job script:** When selected, a dialog box opens. You can specify the following options of the [pre-job script](#) to be executed in this dialog box:
  - **Script path:** A local path to the script on the machine where the Director is installed. A script interpreter should also be specified.

- **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 the options of the [post-job script](#) to be executed on the machine where the product is installed. Proceed as for the pre - job script.

## Completing the New Failover Job Wizard for Amazon EC2

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 256](#) for details.

## Replica Failback for VMware vSphere

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:

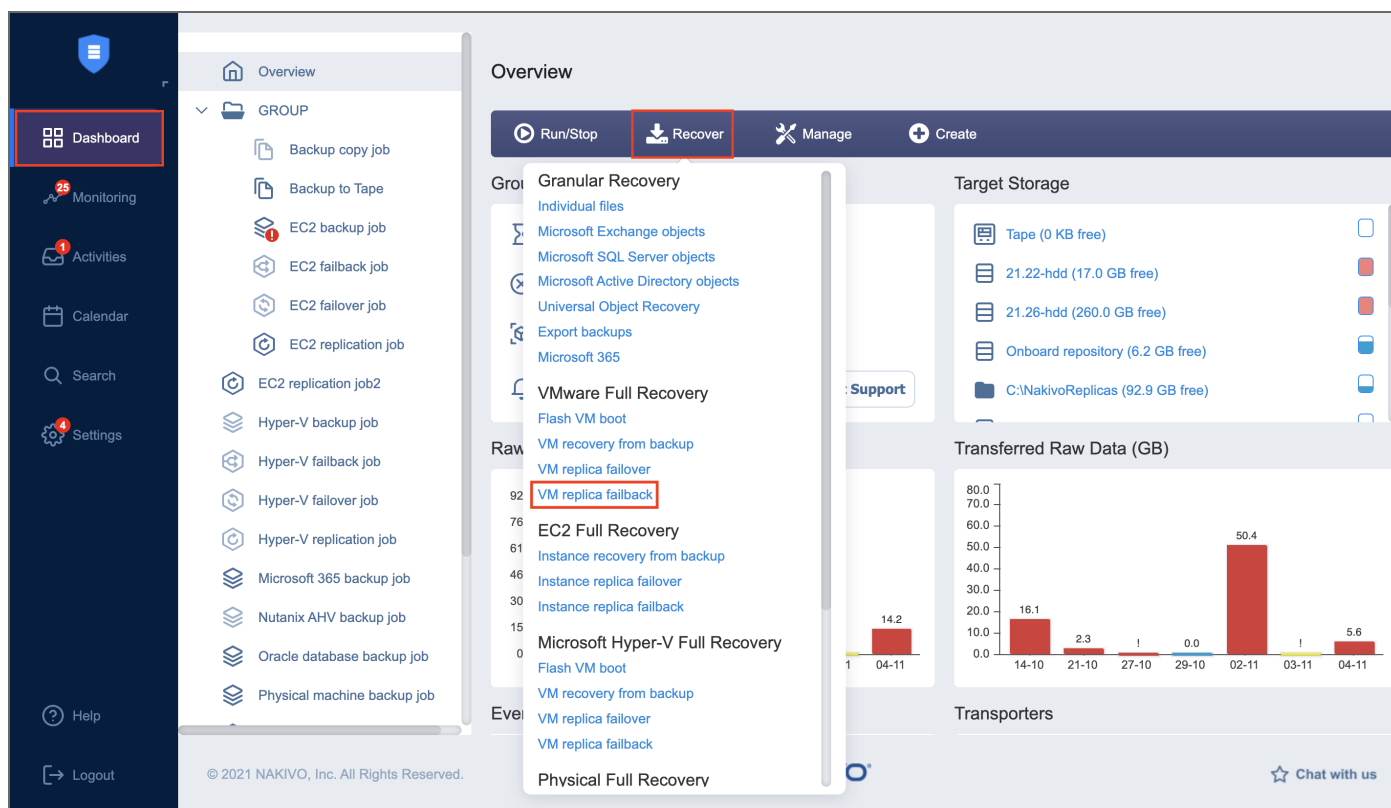
- [“Starting Replica Failback for VMware vSphere” on the next page](#)
- [“Failback Job Wizard for VMware vSphere: Source” on the next page](#)
- [“Failback Job Wizard for VMware vSphere: Destination” on page 994](#)
- [“Failback Job Wizard for VMware vSphere: Networks” on page 995](#)

- [“Failback Job Wizard for VMware vSphere: Re-IP” on page 996](#)
- [“Failback Job Wizard for VMware vSphere: Options” on page 997](#)

## Starting Replica Failback for VMware vSphere

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 971](#) for details.

To start the failback procedure, click **Recover** in the **Dashboard** of NAKIVO Backup & Replication and then click **VM replica failback** in the **VMware Full Recovery** section. Alternatively, the failback can be performed by using the [search](#) function.



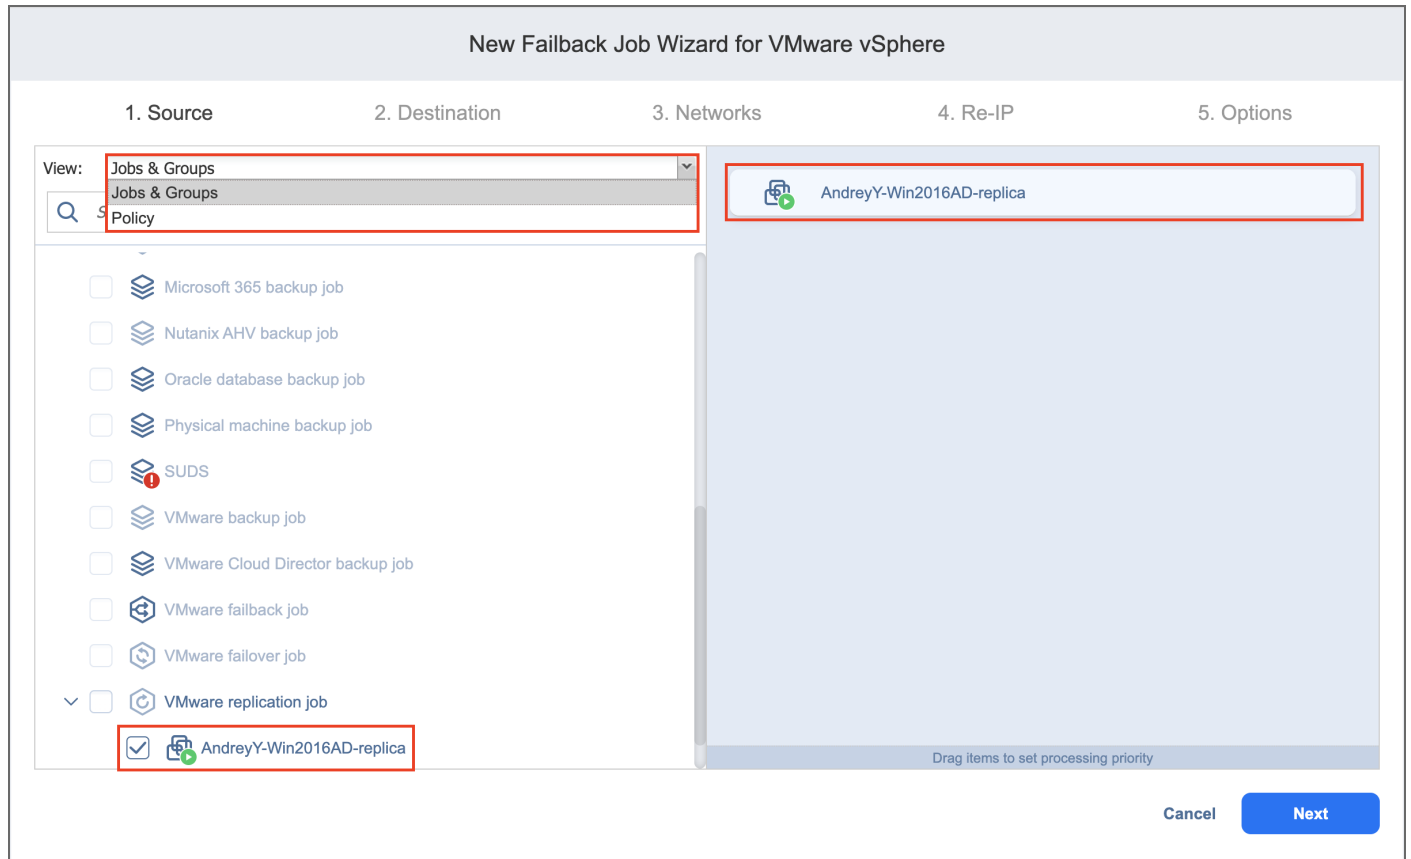
## Failback Job Wizard for VMware vSphere: Source

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 266](#) 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 269](#) 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 vSphere: 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.

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.

New Failback Job Wizard for VMware vSphere

1. Source      2. Destination      3. Networks      4. Re-IP      5. Options

---

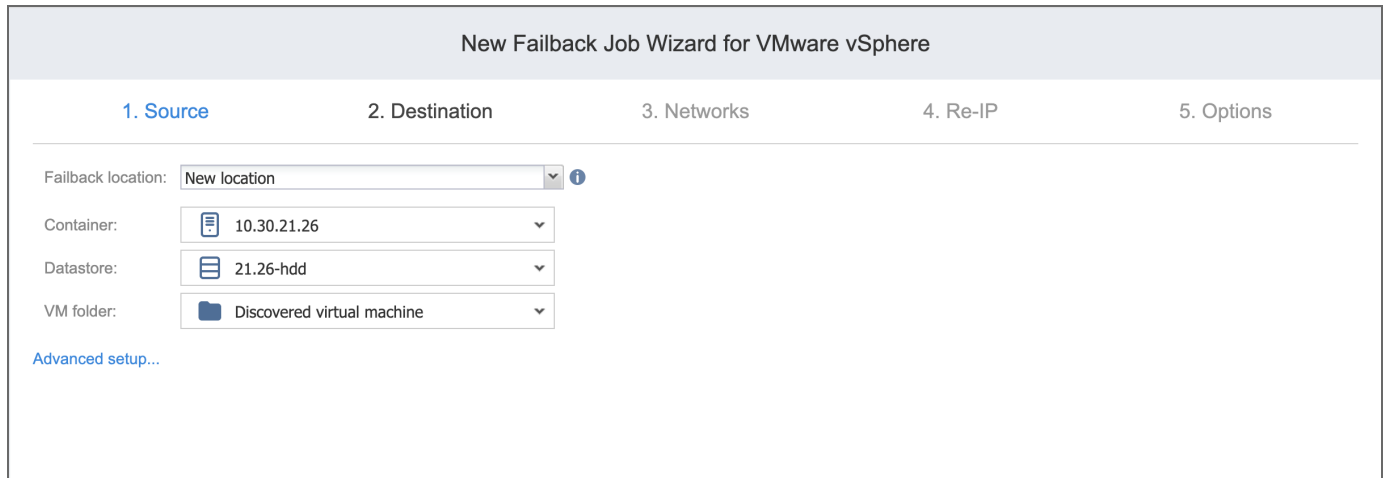
Failback location:  ⓘ

### 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**.



## Failback Job Wizard for VMware vSphere: 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 **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.

New Failback Job Wizard for VMware vSphere

1. Source      2. Destination      3. Networks      4. Re-IP      5. Options

Enable network mapping ?

Network Mappings    [Create new mapping](#)    [Add existing mapping](#)

Source Network	Target Network
10.30.21.0	10.30.21.0
10.30.22.0	10.30.22.0

## Failback Job Wizard for VMware vSphere: 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.

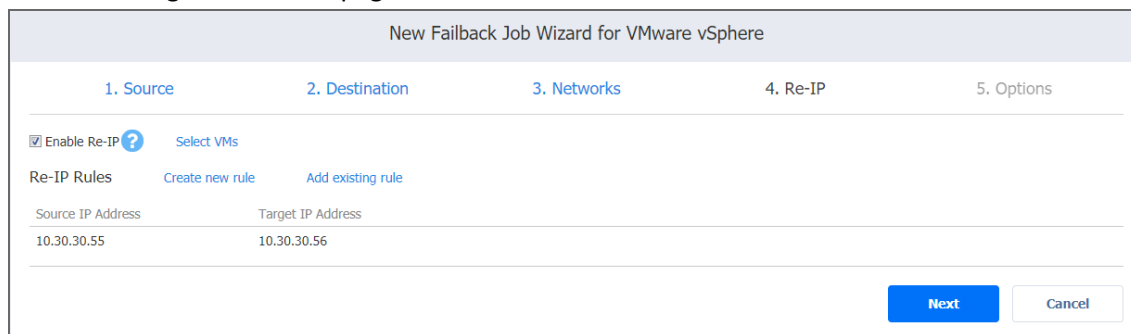
**Example**

When you enter 192.168.1.\* -> 10.30.30.\* Re-IP rule, the source VM IP address such as 192.168.1.50 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.



## Failback Job Wizard for VMware vSphere: 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 **Power off replica VMs** to power off the replica workloads during the failback.

The screenshot shows the 'New Failback Job Wizard for VMware vSphere' with five steps: 1. Source, 2. Destination, 3. Networks, 4. Re-IP, and 5. Options. The 'Options' step is active. The 'Job Options' section is highlighted with a red box and contains the following settings:

- Job name:** VMware failback job
- App-aware mode:** Enabled (proceed on error) [settings]
- Power off replica VMs**

Other sections include:

- Pre and Post Actions:**
  - Send job run reports to [ ] [i]
  - Run local pre job script [i]
  - Run local post job script [i]
- Data Transfer:**
  - Transport mode:** Automatic selection [i]
  - Transporter pool:** Select transporter pool [i]
  - Transporters:** Automatic selection [i]
  - Limit transporter load to:** 3 concurrent tasks [i]
  - Bandwidth throttling:** Disabled [i]
  - Bottleneck detection** [i]

At the bottom right, there are three buttons: 'Cancel', 'Finish', and 'Finish & Run'.

## 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 284](#) 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**.

The screenshot shows the 'New Failback Job Wizard for VMware vSphere' dialog box, specifically the '5. Options' step. The 'Pre and Post Actions' section is highlighted with a red box. The options in this section are:

- Send job run reports to
- Run local pre job script
- Run local post job script

Other sections visible in the wizard include:

- Job Options:** Job name: VMware failback job; App-aware mode: Enabled (proceed on error);  Power off replica VMs.
- Data Transfer:** Transport mode: Automatic selection; Transporter pool: Select transporter pool; Transporters: Automatic selection; Limit transporter load to: 3 concurrent tasks; Bandwidth throttling: Disabled; Bottleneck detection: .

Buttons at the bottom right: Cancel, Finish, Finish & Run.

## 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. Do one of the following:

- Select **Automatic selection** to allow the product to choose the best transport mode available. If the product cannot use SAN or Hot Add, LAN mode will be used.
- Manually select one of the available transport mode options:
  - [SAN](#)

- [Hot Add](#)
- **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 284](#) 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 284](#) 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.

New Failback Job Wizard for VMware vSphere

1. Source
2. Destination
3. Networks
4. Re-IP
5. Options

**Job Options**

Job name:

App-aware mode:  [settings](#)

Power off replica VMs

**Pre and Post Actions**

Send job run reports to

Run local pre job script

Run local post job script

**Data Transfer**

Transport mode:

Transporter pool:

Transporters:

Limit transporter load to  concurrent tasks

Bandwidth throttling:

Bottleneck detection

Cancel
Finish
Finish & Run



## Note

If you click **Finish & Run**, you will have to define the scope of your job. Refer to [“Running Jobs on Demand” on page 256](#) for details.

## Replica Failback for Microsoft Hyper-V

With NAKIVO Backup & Replication, you can switch your Hyper-V VM replicas to source or to a new location.


Refer to the following topics for more information:

- [“Starting Replica Failback for Microsoft Hyper-V” below](#)
- [“Failback Job Wizard for Microsoft Hyper-V: Source” on the next page](#)
- [“Failback Job Wizard for Microsoft Hyper-V: Destination” on page 1004](#)
- [“Failback Job Wizard for Microsoft Hyper-V: Networks” on page 1006](#)
- [“Failback Job Wizard for Microsoft Hyper-V: Re-IP” on page 1007](#)
- [“Failback Job Wizard for Microsoft Hyper-V: Options” on page 1008](#)

## Starting Replica Failback for Microsoft Hyper-V

Before starting replica failback for Microsoft Hyper-V, make sure that you have switched the replicas to the "Failover" state with a Hyper-V Failover Job. Refer to [“Failover to Replica for Microsoft Hyper-V” on page 978](#) for details.

To start the failback procedure, click **Recover** in the **Dashboard** of NAKIVO Backup & Replication and then click **VM replica failback** in the **Microsoft Hyper-V Full Recovery** section. Alternatively, the failback can be performed by using the [search](#) function.




## Failback Job Wizard for Microsoft Hyper-V: Source

On the **Source** page of the wizard:

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. For details, refer to [“Managing Job Policies” on page 266](#). 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. For details, refer to [“Managing Policy Rules” on page 269](#).
2. Select one or more replica Hyper-V VMs in the left pane of the page.

3. Click **Next** to confirm adding the selected replicas to the failback job.



## Failback Job Wizard for Microsoft Hyper-V: Destination

On the **Destination** page of the wizard, you can specify a failback destination for your Hyper-V replicas. Failback to the original location will update or replace the existing source VM in the original location. Failback to a new location will create a new VM. Refer to the following subsections on providing a destination for your Microsoft Hyper-V Replica Failback job:

- [Using Original Hyper-V Failback Location](#)
- [Creating New Hyper-V Failback Location](#)
- [Creating a Different Hyper-V Failback Location for All Replicas](#)

### Using Original Hyper-V Failback Location

To use the original location for your Hyper-V failback job, follow the steps below:

1. In the **Failback location** list, choose **Original Location**.
2. Click **Next** to go to the **Networks** page.

New Failback Job Wizard for Microsoft Hyper-V

1. Source      2. Destination      3. Networks      4. Re-IP      5. Options

---

Failback location:  ⓘ

## Creating New Hyper-V Failback Location

To create a new location for your Hyper-V 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 **Path** box, enter a path to the location of the Hyper-V VM replica.
4. Click **Next** to go to the next page of the wizard.

New Failback Job Wizard for Microsoft Hyper-V

1. Source      2. Destination      3. Networks      4. Re-IP      5. Options

---

Failback location:  ⓘ

Container:

Path:

[Advanced setup...](#)

## Creating a Different Hyper-V Failback Location for All Replicas

To create a new failover location for every replica you've added to the job:

1. Click **Advanced setup**.
2. Select a target container and target disks.
3. Click **Next** to go to the next page.

**New Failback Job Wizard for Microsoft Hyper-V**

1. Source      2. Destination      3. Networks      4. Re-IP      5. Options

---

Failback location:  ⓘ

Container:

Path:

ubuntu-forquis-replica Click to collapse

Source	Target
VM location: ServerHV2012 > ServerHV2012	Container: <input type="text" value="ServerHV2012"/>
VM resources: 2 CPU, 1.0 GB RAM	Virtual Machine: <input type="text" value="New VM will be created"/>
Disks	Disks
<input checked="" type="checkbox"/> e71abb5d-b86b-4...	<input type="text" value="e71abb5d-b86b-4c48-a..."/>
5.8 GB	<input type="text" value="C:\NakivoReplicas"/>
VM configuration	VM configuration: <input type="text" value="C:\NakivoReplicas"/>

## Failback Job Wizard for Microsoft Hyper-V: Networks

When the replica network address differs from the network address where the source VMs are deployed, your failback job needs a relation between these networks to be performed successfully. This relation is called a network mapping.

To map source VM virtual networks to the appropriate target virtual networks, please do the following 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 network 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 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 over the required item in the **Network Mappings** list.
  - b. A toolbar with icons opens to the right of the item. Click the **Edit** icon.
  - c. The **Edit Network Mapping** dialog opens. Choose an appropriate item from the **Target network** drop-down list and click **Save**.
  - Delete an existing mapping:
    - a. Hover over the required 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 Microsoft Hyper-V: Re-IP

When the IP addressing scheme for the replica differs from the IP addressing scheme for the Hyper-V source VMs, your failback job needs relations between source VMs addresses and the replica VMs addresses to be performed successfully. These relations are called Re-IP rules.

### Warning

Hyper-V Integration Services must be running on source VMs to enable successful re-IP for your failback 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 opens. In the list of your source VMs, select at least one and close the dialog.
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 the 192.168.1.\* -> 10.30.30.\* 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 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. Choose an appropriate Re-IP rule and close the dialog box.
  - Edit an existing Re-IP rule:
    - a. Hover 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** icon.
    - c. The **Edit Re-IP Rule** dialog box opens. Edit the required properties of the Re-IP rule and click **Save**.
  - Delete an existing mapping:
    - a. Hover the pointer over the necessary item in the **Re-IP Rules**
    - 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 Failback Job Wizard for Microsoft Hyper-V

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
10.30.30.55	10.30.30.56

[Next](#)    [Cancel](#)

## Failback Job Wizard for Microsoft Hyper-V: Options

On the **Options** page of the wizard, specify the options for the Microsoft Hyper-V failback job as described in the following sections:

- [Job Options](#)
- [Pre and Post Actions](#)

### 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 **Power off replica VMs** to power off the replica workloads during the failback.

New Failback Job Wizard for Microsoft Hyper-V

1. Source      2. Destination      3. Networks      4. Re-IP      5. Options

---

**Job Options**  
 Job name:   
 App-aware mode:  [settings](#)  
 Power off replica VMs

**Pre and Post Actions**

Send job run reports to  [i](#)

Run local pre job script [i](#)

Run local post job script [i](#)

**Data Transfer**

Limit transporter load to  concurrent tasks [i](#)

Bandwidth throttling:  [i](#)

Bottleneck detection [i](#)

---

## Pre and Post Actions

If needed, enable pre and post actions for the failback job:

1. **Send job run reports to:** When selected, a job report is sent to the specified recipients each time the failback job is completed. This overrides the default setting specified in the NAKIVO Backup & Replication **Settings**.
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 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 run on the machine where the product is installed. Proceed to set the options as described for pre-



job scripts.

The screenshot shows the 'New Failback Job Wizard for Microsoft Hyper-V' interface, specifically the 'Options' step. The wizard is divided into five steps: 1. Source, 2. Destination, 3. Networks, 4. Re-IP, and 5. Options. The 'Options' step is currently active. The 'Job Options' section includes a 'Job name' field with the value 'Hyper-V failback job', an 'App-aware mode' dropdown set to 'Enabled (proceed on error)', and a checked checkbox for 'Power off replica VMs'. The 'Pre and Post Actions' section is highlighted with a red box and contains three unchecked checkboxes: 'Send job run reports to', 'Run local pre job script', and 'Run local post job script'. The 'Data Transfer' section includes a checked checkbox for 'Limit transporter load to', a spinner box set to '3' with the label 'concurrent tasks', a 'Bandwidth throttling' dropdown set to 'Disabled', and an unchecked checkbox for 'Bottleneck detection'. At the bottom right, there are three buttons: 'Cancel', 'Finish', and 'Finish & Run'.

## Data Transfer

In the *Data Transfer* section of the **Options** page, you can configure Transporter load, bandwidth load, and other options. Proceed as described in the sections below.

### 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 284](#) 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 284](#) 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.

New Failback Job Wizard for Microsoft Hyper-V

1. Source
2. Destination
3. Networks
4. Re-IP
5. Options

**Job Options**

Job name:

App-aware mode:  [settings](#)

Power off replica VMs

**Pre and Post Actions**

Send job run reports to

Run local pre job script

Run local post job script

**Data Transfer**

Limit transporter load to  concurrent tasks

Bandwidth throttling:

Bottleneck detection

Cancel
Finish
Finish & Run

# Replica Failback for Amazon EC2

With NAKIVO Backup & Replication, you can switch your Amazon EC2 instance replica back to the source or to a new location.

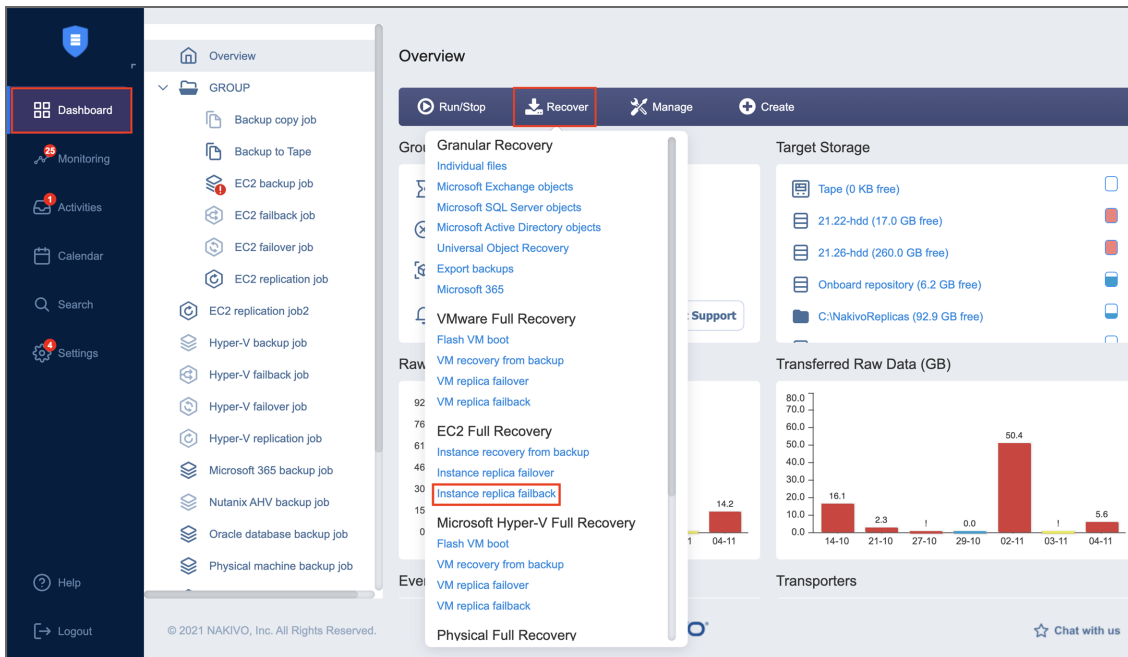
Refer to the following topics for more information:

- [“Starting Replica Failback for Amazon EC2” below](#)
- [“Failback Job Wizard for Amazon EC2: Source” below](#)
- [“Failback Job Wizard for Amazon EC2: Destination” on the next page](#)
- [“Failback Job Wizard for Amazon EC2: Options” on page 1016](#)

## Starting Replica Failback for Amazon EC2

Before starting replica failback for Amazon EC2, make sure that you switched the replicas to the "Failover" state with an Amazon EC2 Failover Job. Refer to [“Failover to Replica for Amazon EC2” on page 985](#) for details.

To start the failback procedure, click **Recover** in the **Dashboard** of NAKIVO Backup & Replication and then click **Instance replica failback** in the **EC2 Full Recovery** section. Alternatively, the failback can be performed by using the [search](#) function.

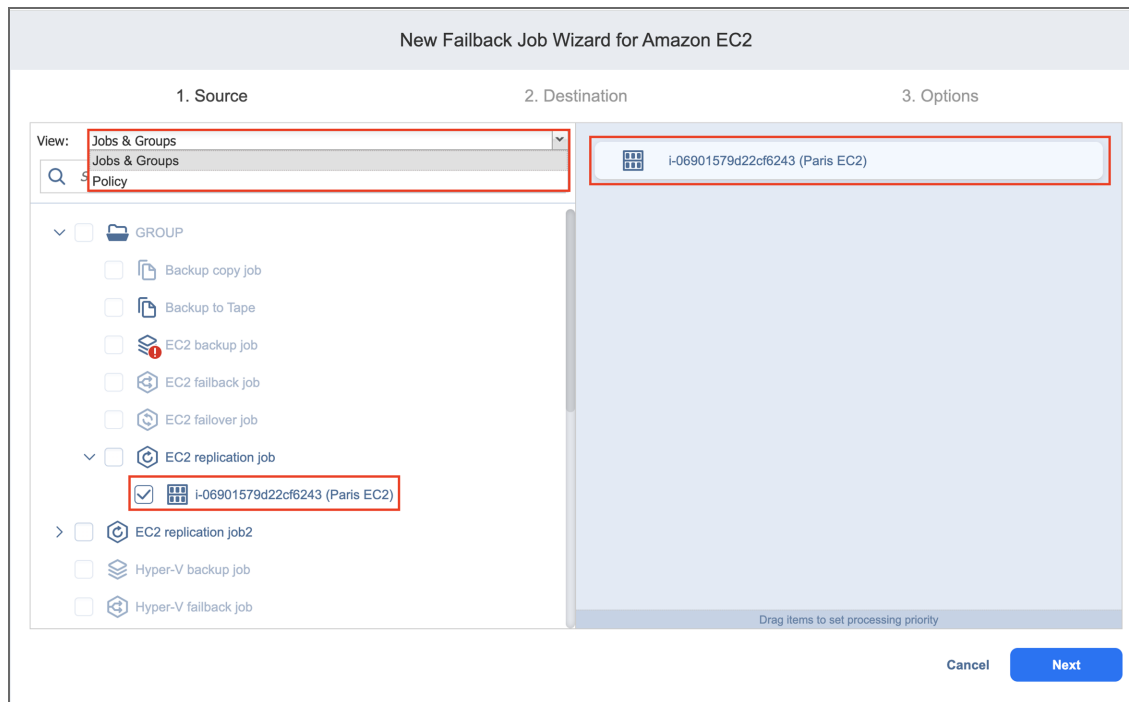


## Failback Job Wizard for Amazon EC2: Source

On the **Source** page of the wizard, select replicas you want to add to the failover job by using one of the 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 266](#) 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 269](#) for details.



Click **Next** to go to the **Destination** page.

## Failback Job Wizard for Amazon EC2: Destination

On this page of the wizard, you can specify a failback destination for your Amazon EC2 instances. Failback to the original location will update/replace the existing source instance in the original location; failback to the new location will create a new instance. Refer to the following subsections on providing a destination for your Amazon EC2 replica failback job:

- [“Setting Original EC2 Failback Destination” below](#)
- [“Setting Single New EC2 Failback Destination” on the next page](#)
- [“Setting Multiple New Failback Destinations for EC2 Instances” on page 1015](#)

### Setting Original EC2 Failback Destination

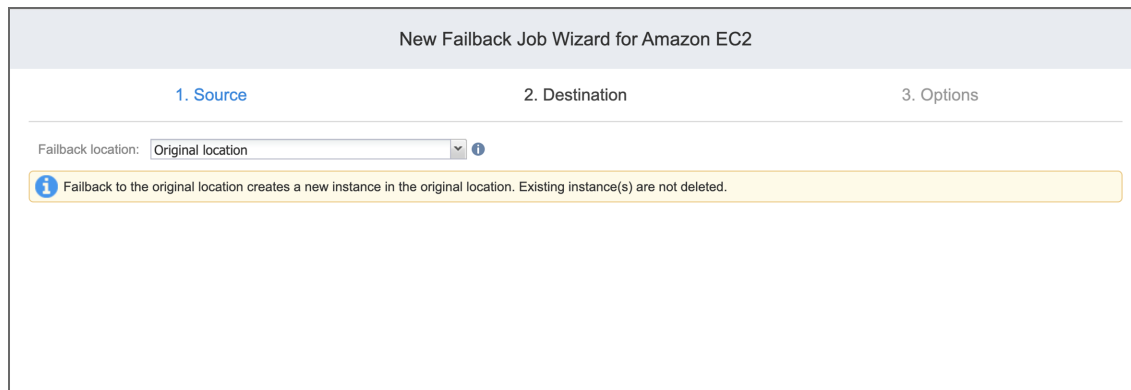
To use the original location as a destination for your Amazon EC2 failback job, follow the steps below:

1. In the **Failback location** list, choose **Original Location**.

#### **Important**

To successfully failback an Amazon EC2 instance to its original location, the source instance should not exist.

2. Click **Next** to go to the **Options** page.



## Setting Single New EC2 Failback Destination

To create a new destination for all instances of your Amazon EC2 failback job, follow the steps below:

1. In the **Failback location** list, choose **New Location**.
2. A set of fields opens for editing:
  - a. **Region:** Select an Amazon EC2 [region](#) from the list of available regions.
  - b. **Network:** Select an Amazon EC2 [VPC](#) to which all EC2 instances will be connected.
  - c. **Subnet:** Select an Amazon EC2 [Subnet](#) to which all EC2 instances will be connected.
  - d. **Security group:** Select an Amazon EC2 [Security Group](#) for all EC2 instances.
  - e. **Instance access:** Specify Amazon EC2 [Key Pair](#) options by choosing one of the following follows:
    - **Use existing key pair:** An existing Amazon EC2 Key Pair will be used to connect to the instances.
    - **Create a new key pair:** A new Amazon EC2 Key Pair will be created and assigned to all recovered EC2 instances. Type in a name for the new Amazon EC2 Key Pair in the **Key pair name** field and click **Download**.

### **Important**

This is your only chance to save the private key file, so make sure you save the private key file in a safe place. You will need to provide the name of your key pair when you start recovered EC2 instances and the corresponding private key each time you connect to the instances.

- **Proceed without key pair:** No AWS Key Pair will be assigned to the recovered EC2 instances.

### **Important**

A key pair enables you to connect to Linux Instances via SSH. If you recover Linux Instances without a key pair, you will not be able to connect to them.

3. Click **Next**.

New Failback Job Wizard for Amazon EC2

1. Source      2. Destination      3. Options

Failback location: **New location** ⓘ

Region: Asia Pacific (Seoul)

Network: NAKIVO vpc (10.0.0.0/16)

Subnet: NAKIVO subnet (10.0.0.0/16)

Security group: Asia Pacific (Seoul) Linux transporter\_...

Instance access: Use existing key pair

Key pair: Select target key pair

[Advanced setup...](#)

Cancel    Next

## Setting Multiple New Failback Destinations for EC2 Instances

To set multiple destinations for the instances of your Amazon EC2 failback job, follow the steps below:

1. In the **Failback location** list, choose **New Location**.
2. A list of options opens for setting up. Click on the required instances to expand them.
3. For the required instance, set the options as follows:
  - a. **Region:** Select an Amazon EC2 [region](#) to which the Amazon EC2 instance will failback.
  - b. **Instance type:** Select a [type](#) for your failback EC2 instance.
  - c. **EBS Volumes:** Choose whether to fail back or skip individual Amazon [EBS Volumes](#).
  - d. **Network:** Select an Amazon EC2 VPC to which the EC2 instance will fail back.
  - e. **Network adapter(s):** Specify the [Amazon EC2 network adapter](#) options for every network adapter:
    - **Subnet:** Select an Amazon EC2 [subnet](#) to which the network adapter will be connected.
    - **Security group:** Select an Amazon EC2 [security group](#) for the Amazon EC2 instance.
    - **Primary IP:** Assign a valid primary IP to a current adapter.
    - **Instance access:** Specify Amazon EC2 [key pair](#) options as follows:
      - **Use existing key pair:** Choose an existing Amazon EC2 Key Pair that you want to use to connect to the failback EC2 instance.

- **Create a new key pair:** Select this option to create a new Amazon EC2 Key Pair and assign it to the failback EC2 instance. Enter a name for the new Amazon EC2 Key Pair in the key pair name box and click **Download**.

**Important**

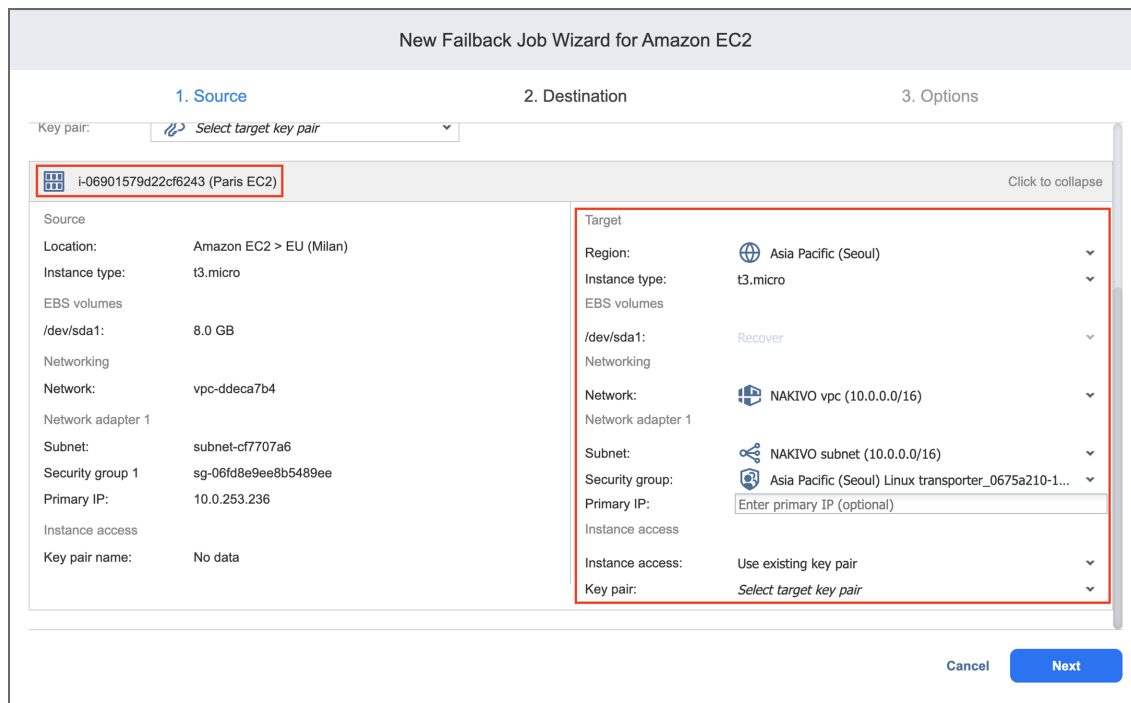
This is your only chance to save the private key file, so make sure to save the private key file in a safe place. You will need to provide the name of your key pair when you launch the recovered Amazon EC2 instance and the corresponding private key each time you connect to the instance.

- **Proceed without key pair:** Select this option if you do not want to assign an AWS Key Pair to the failback EC2 instance.

**Important**

A key pair enables you to connect to Linux Instances via SSH. If you recover Linux Instances without a key pair, you will not be able to connect to them.

4. Click **Next** to go to the **Options** page.



## Failback Job Wizard for Amazon EC2: Options

On the **Options** page of the wizard, specify options for the Amazon EC2 failback job as described in the sections below:

- [“Job Options” on the next page](#)
- [“Pre and Post Actions” on the next page](#)

## Job Options

- In the **Job name** box, enter a string of 50 or fewer characters to specify the name of your failback job.
- Optionally, enable **App-aware mode**. Choose among the following options:
  - **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.
- Select **Power off replica instances** to power off the replica workloads during the failback.

The screenshot shows the 'New Failback Job Wizard for Amazon EC2' at the 'Options' step. The 'Job Options' section is highlighted with a red box. It contains the following fields and options:

- Job name:** EC2 failback job
- App-aware mode:** Disabled
- Power off replica instances**

Below the 'Job Options' section is the 'Pre and Post Actions' section, which includes three checkboxes:

- Send job run reports to
- Run local pre job script
- Run local post job script

At the bottom of the wizard, there are three buttons: 'Cancel', 'Finish', and 'Finish & Run'.

## 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 is completed. This overrides the default setting specified in the NAKIVO Backup & Replication settings.
2. **Run local pre job script:** When selected, a dialog box opens in which you can set up 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 run on the machine where the product is installed. Proceed with the options as described for **Run local pre job script**.

The screenshot shows the 'New Failback Job Wizard for Amazon EC2' dialog box, specifically the '3. Options' step. The 'Job Options' section includes a 'Job name' field with 'EC2 failback job', an 'App-aware mode' dropdown set to 'Disabled', and a checked checkbox for 'Power off replica instances'. The 'Pre and Post Actions' section, highlighted with a red box, contains three checkboxes: 'Send job run reports to' (unchecked), 'Run local pre job script' (unchecked), and 'Run local post job script' (checked). At the bottom right, there are 'Cancel', 'Finish', and 'Finish & Run' buttons.

#### Note

If you click **Finish & Run**, you will have to define the scope of your job. Refer to [“Running Jobs on Demand” on page 256](#) 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.

Creating a Site Recovery Job is done with a wizard and includes the following steps:

- [“Creating Site Recovery Job” below](#)
- [“Running Site Recovery Job” on page 1058](#)

## Creating Site Recovery Job

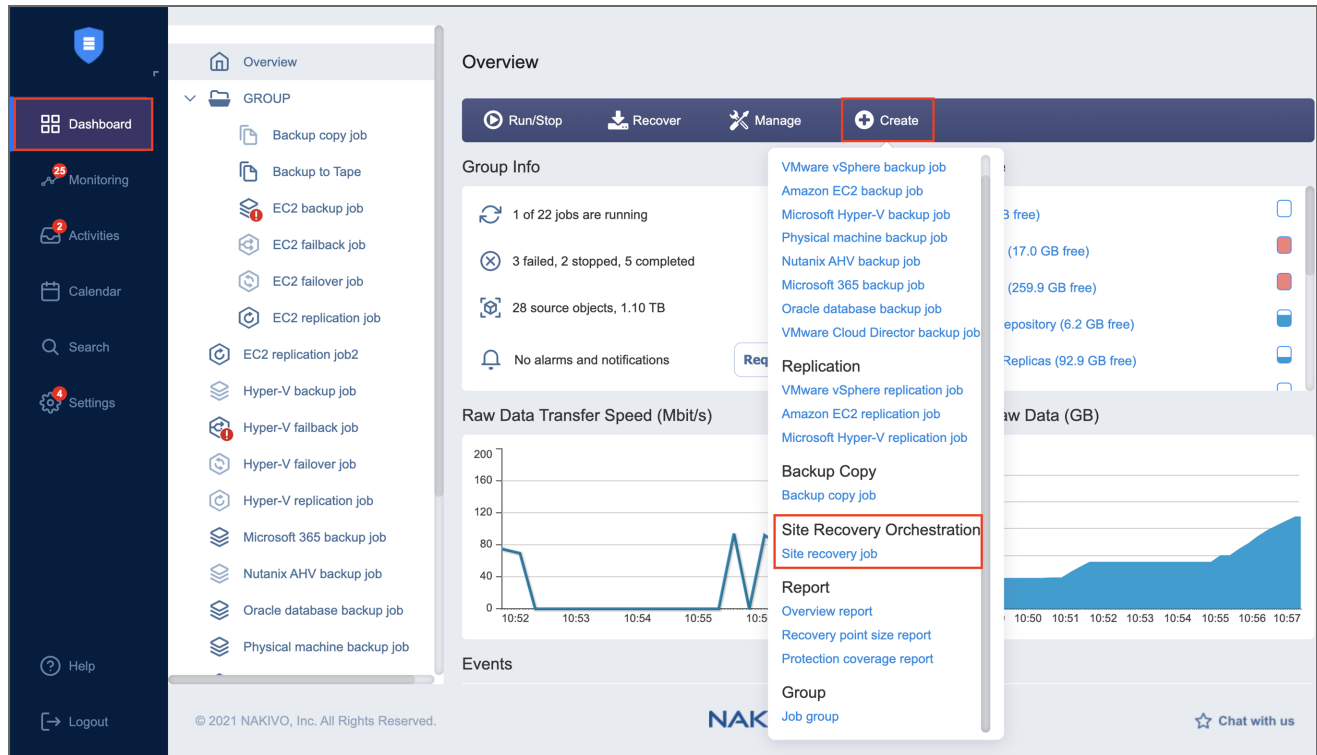
Creating a Site Recovery Job is done with a wizard and includes the following steps:

- [“Starting Site Recovery Job Wizard” on the next page](#)
- [“Site Recovery Job Wizard: Actions” on the next page](#)
- [“Site Recovery Job Wizard: Networks” on page 1049](#)
- [“Site Recovery Job Wizard: Re-IP” on page 1051](#)
- [“Site Recovery Job Wizard: Test Schedule” on page 1053](#)
- [“Site Recovery Job Wizard: Options” on page 1057](#)

## Starting Site Recovery Job Wizard

Follow the steps below to start creating a Site Recovery Job:

1. Go to the **Dashboard** and open the **Create** menu.
2. In the **Create** menu, click **Site recovery job**.



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: VMs” on page 1022](#)
- [“Failover Hyper-V VMs: VMs” on page 1023](#)
- [“Failover Amazon EC2 Instances: Instances” on page 1025](#)
- [“Failback VMware VMs: VMs” on page 1030](#)
- [“Failback Hyper-V VMs: VMs” on page 1034](#)

- [“Failback Amazon EC2 Instances: Instances” on page 1037](#)
- [“Start VMs / Instances Action” on page 1042](#)
- [“Stop VMs / Instances Action” on page 1043](#)
- [“Run / Stop Jobs Action” on page 1044](#)
- [“Run Script Action” on page 1045](#)
- [“Attach / Detach Repository Action” on page 1046](#)
- [“Send Email Action” on page 1047](#)
- [“Wait Action” on page 1047](#)
- [“Check Condition Action” on page 1048](#)

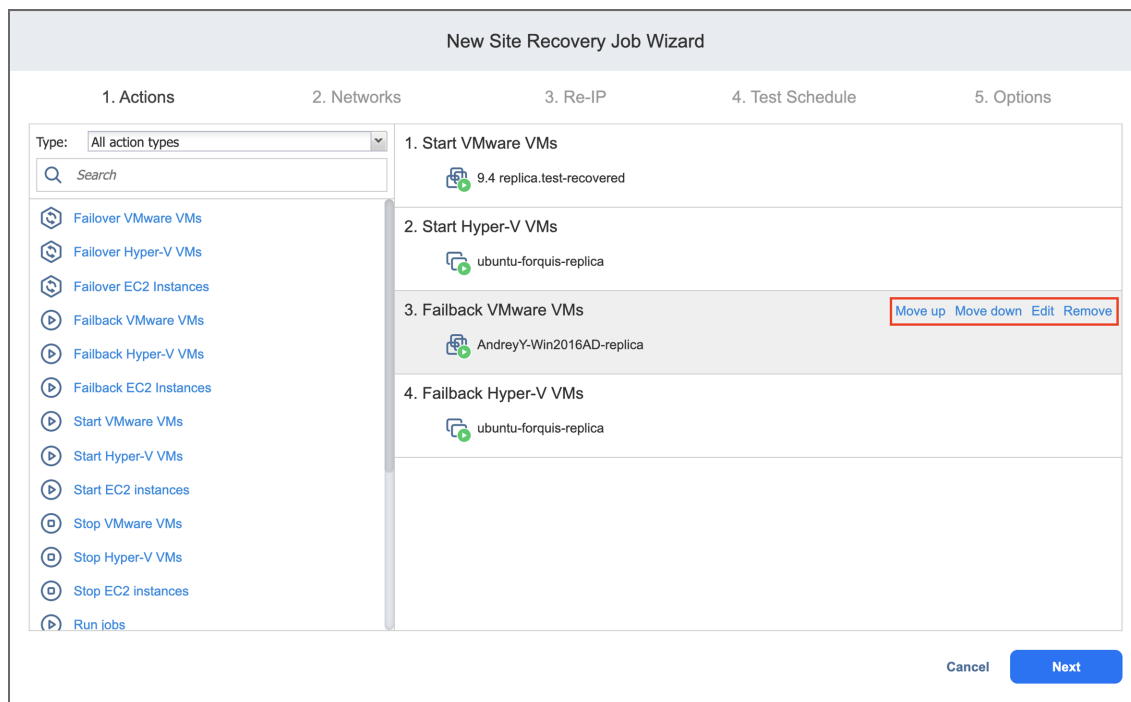
To add an action to a Site Recovery job, click the corresponding item in the actions list and follow the instructions of the wizard that opens.

### 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.
- **Transporter Pool:** If this option is enabled, only the transporters that belong to the selected transporter pool shall be used during the job run.
- **Transport mode:** This option defines the method of VM data retrieval by a source Transporter. Do one of the following:
  - Select **Automatic selection** to allow the product to choose the best transport mode available. If the product cannot use SAN or Hot Add, LAN mode will be used.
  - Manually select one of the available transport mode options:
    - **SAN**
    - **Hot Add**
    - **LAN**
- **Transporters:** This option allows you to specify which Transporters should be used for the job. 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.

**Note**

The **Transport mode** and **Transporters** options are available for VMware failback options action only.

## 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 266](#) 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 269](#) 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.

New Site Recovery Job Wizard > 1. Failover VMware VMs

1. VMs | 2. Options

View: Jobs & Groups

- Nutanix AHV backup job
- Oracle database backup job
- Physical machine backup job
- SUDS
- VMware backup job
- VMware Cloud Director backup job
- VMware failback job
- VMware failover job
- VMware replication job
  - ayunt\_Centos8-replica-replica

VMware replication job

ayunt\_Centos8-replica-replica  
Always use the latest recovery point

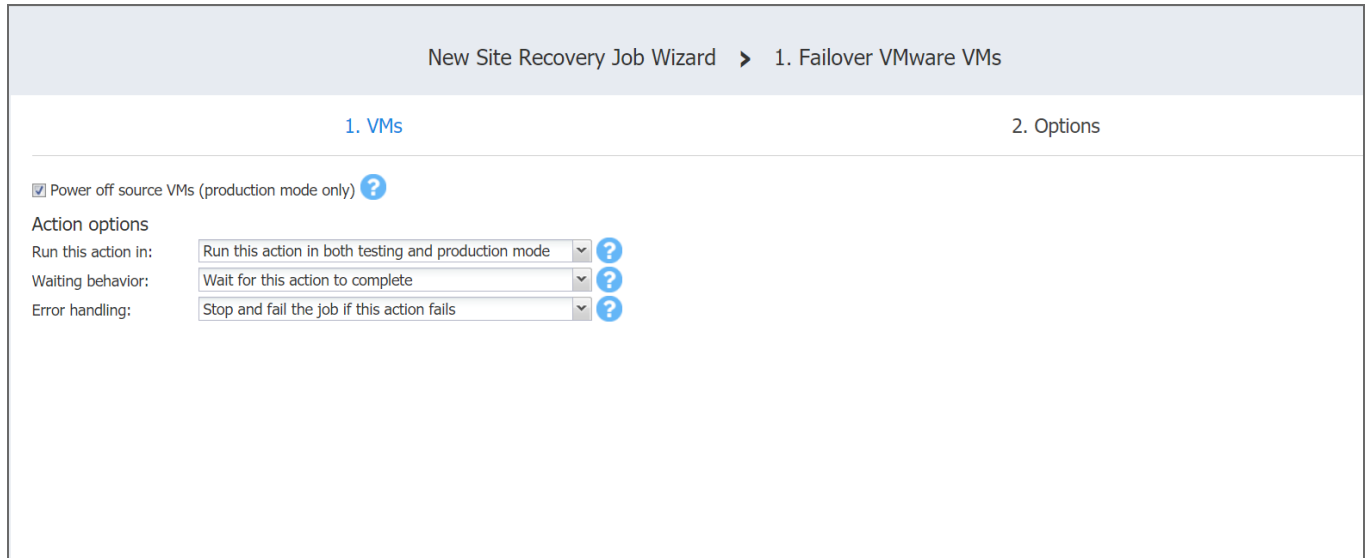
Drag items to set processing priority

Cancel Next

# 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 1019](#) for details.
3. Click **Save**.



New Site Recovery Job Wizard > 1. Failover VMware VMs

1. VMs 2. Options

Power off source VMs (production mode only) ?

Action options

Run this action in: Run this action in both testing and production mode ?

Waiting behavior: Wait for this action to complete ?

Error handling: Stop and fail the job if this action fails ?

The **Options** page closes and your **Failover VMware VMs** action is added to the Site Recovery Job.

## Failover Hyper-V VMs Action

Refer to the following sections on adding a **Failover Hyper-V VMs** action to your Site Recovery job:

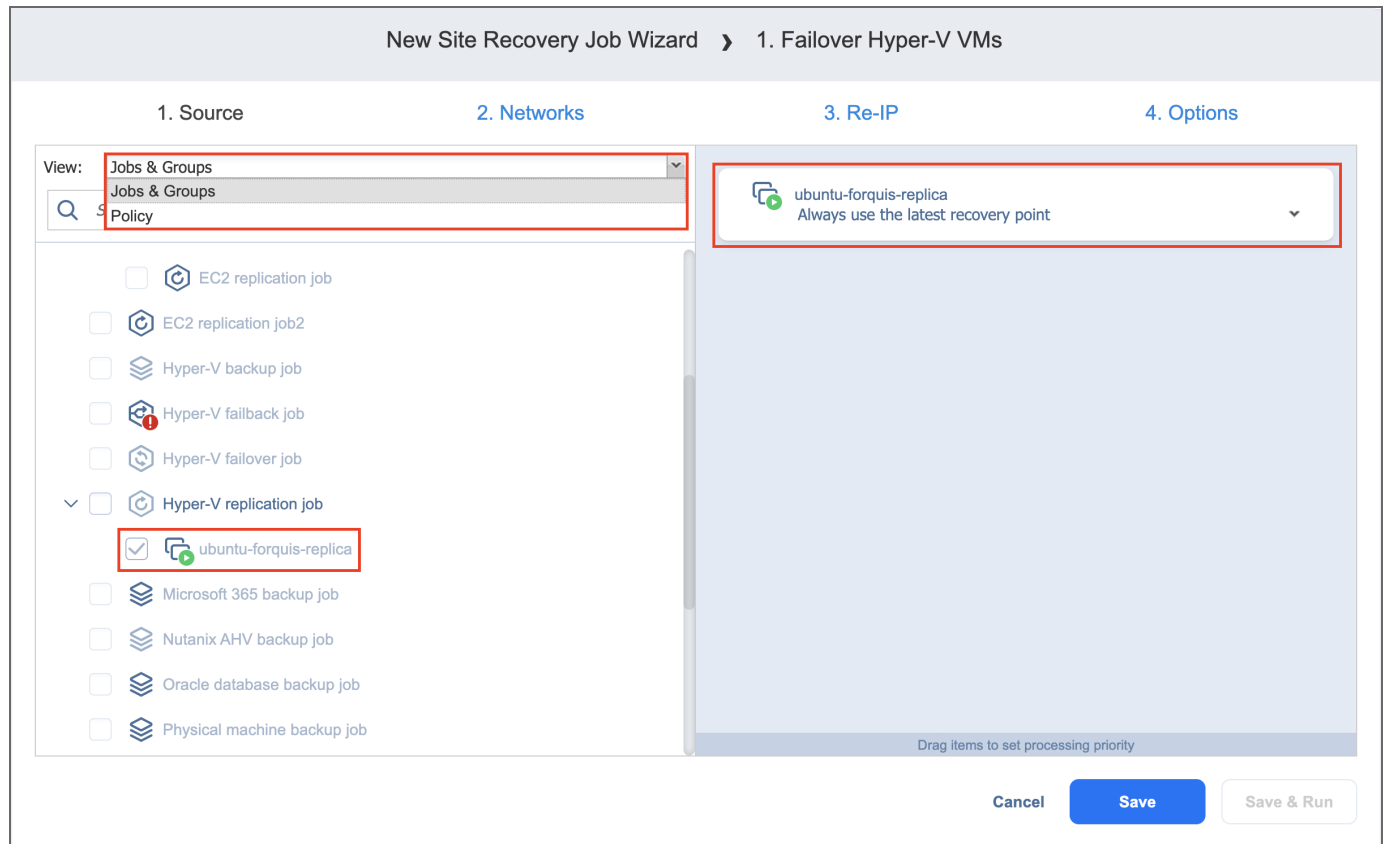
- [Failover Hyper-V VMs: VMs](#)
- [Failover Hyper-V VMs: Options](#)

# Failover Hyper-V 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 Hyper-V VM from the list of replication jobs.
  - **Policy:** When selected, job policies can be used. Refer to [“Managing Job Policies” on page 266](#) 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 269](#) for details.

2. Select one or more source Hyper-V 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.

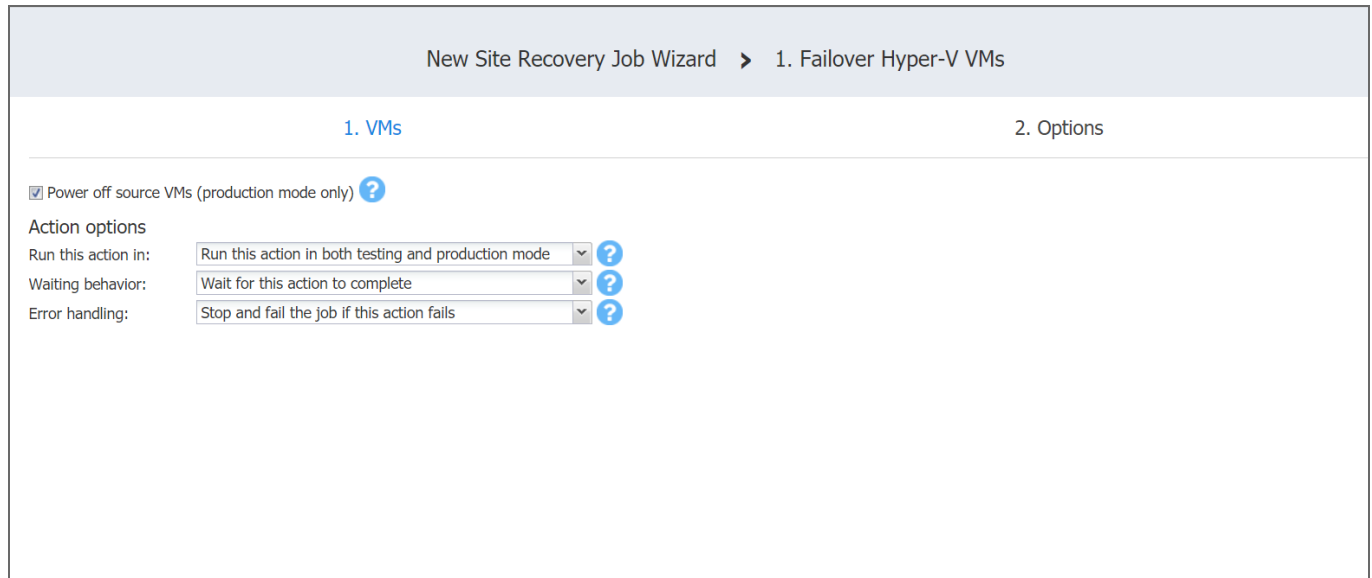


## Failover Hyper-V VMs: Options

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 1019](#) for details.

3. Click **Save**.



The **Options** page closes and your **Failover Hyper-V VMs** action is added to the Site Recovery Job.

### Failover Amazon EC2 Instances Action

Refer to the following sections on adding a **Failover EC2 Instances** action to your Site Recovery job:

- [“Failover Amazon EC2 Instances: Instances” below](#)
- [“Failover Amazon EC2 Instances: Destination” on the next page](#)
  - [“Single Destination for Amazon EC2 Instances” on the next page](#)
  - [“Different Destinations for Amazon EC2 Instances” on page 1027](#)
- [“Failover Amazon EC2 Instances: Options” on page 1029](#)

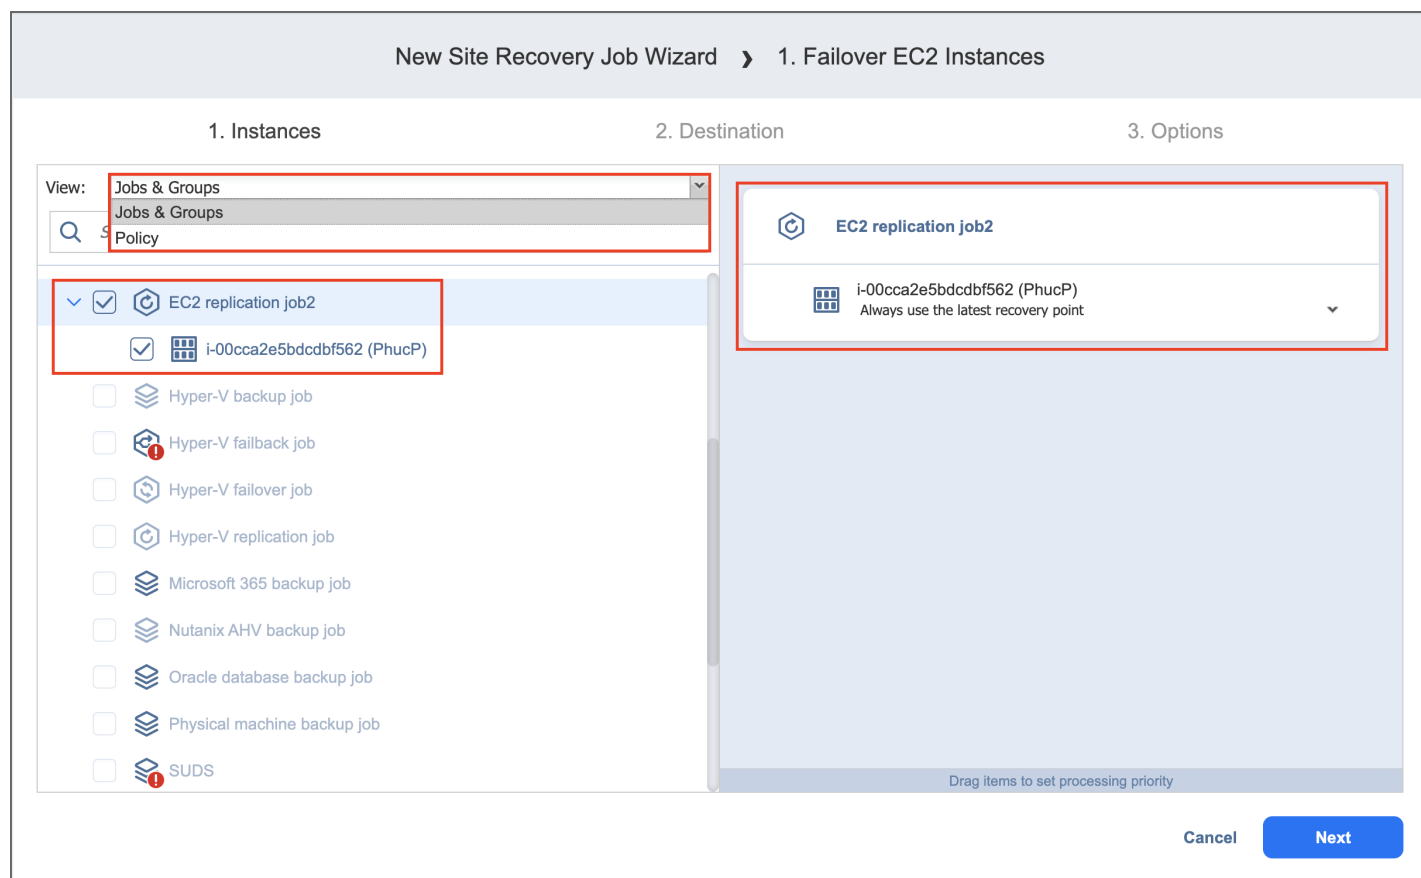
## Failover Amazon EC2 Instances: Instances

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 AWS account. 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 266](#) 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 269](#) for details.
2. Select one or more source Amazon EC2 instances from the inventory tree and then select a recovery point for each source Amazon EC2 instance in the right pane.



3. Click **Next** to go to the **Destination** page.



## Failover Amazon EC2 Instances: Destination

On this page of the wizard, you can specify a failover destination for your Amazon EC2 instances.

### Single Destination for Amazon EC2 Instances

To fail over all Amazon EC2 instances defined in the previous page of the wizard to the same AWS region and with the same options, do the following:

1. **Network:** Select an Amazon EC2 [VPC](#) to which all Amazon EC2 instances will be connected.
2. **Subnet:** Select an Amazon EC2 [Subnet](#) to which all Amazon EC2 instances will be connected.
3. **Security group:** Select an Amazon EC2 [Security Group](#) for all Amazon EC2 instances.
4. **Instance access:** Specify Amazon EC2 [Key Pair](#) options by choosing one of the following:
  - **Use existing key pair:** An existing Amazon EC2 Key Pair will be used to connect to the recovered instances.
  - **Create a new key pair:** A new Amazon EC2 Key Pair will be created and assigned to all recovered EC2 instances. Type in a name for the new Amazon EC2 Key Pair in the **Key pair name** field and click **Download**.

### Important

This is your only chance to save the private key file, so make sure you save the private key file in a safe place. You will need to provide the name of your key pair when you start recovered EC2 instances and the corresponding private key each time you connect to the instances.

- **Proceed without key pair:** No AWS Key Pair will be assigned to the recovered Amazon EC2 instances.

### Important

A key pair enables you to connect to Linux Instances via SSH. If you recover Linux Instances without a key pair, you will not be able to connect to them.

5. Click **Next** to go to the next page of the wizard.

New Site Recovery Job Wizard > 1. Failover EC2 Instances

1. Instances      2. Destination      3. Options

Network: NAKIVO vpc (10.0.0.0/16)

Subnet: NAKIVO subnet (10.0.0.0/16)

Security group: amazon EC2\_\_30a997e0-6ba5-48e1-9...

Instance access: Use existing key pair

Key pair: Select target key pair

Advanced setup...

Target region cannot be changed for replicas. You can only re-select networks and security groups to which replicas will be connected.

Cancel      Next

## Different Destinations for Amazon EC2 Instances

To set different destination options for individual Amazon EC2 instances, follow the steps below:

1. Click **Advanced options**.
2. Click an Amazon EC2 Instance to expand it and set the options as follows:
  - a. **Region:** Select a target Amazon EC2 **Region** for the Amazon EC2 instance.
  - b. **Instance type:** Select a **type** for the recovered Amazon EC2 instance.

- c. **EBS Volumes:** Choose whether to recover or skip recovering individual AWS [EBS Volumes](#).
- d. **Network:** Select an Amazon EC2 [VPC](#) to which the Amazon EC2 instance will be connected.
- e. **Subnet:** Select an Amazon EC2 [Subnet](#) to which the Amazon EC2 instance will be connected.
- f. **Security group:** Select an Amazon EC2 [Security Group](#) for the Amazon EC2 instance.
- g. **Instance access:** Specify Amazon EC2 [Key Pair](#) options as follows:
  - **Use existing key pair:** Choose an existing Amazon EC2 Key Pair that you want to use to connect to the recovered instance.
  - **Create a new key pair:** Select this option to create a new Amazon EC2 Key Pair and assign it to the recovered Amazon EC2 instance. Enter a name for the new Amazon EC2 Key Pair in the **Key pair name** box and click **Download**.

**Important**


This is your only chance to save the private key file, so make sure you save the private key file in a safe place. You will need to provide the name of your key pair when you start recovered EC2 instances and the corresponding private key each time you connect to the instances.

- **Proceed without key pair:** Select this option if you do not want to assign an Amazon Key Pair to the recovered Amazon EC2 instance.

**Important**

A key pair enables you to connect to Linux Instances via SSH. If you recover Linux Instances without a key pair, you will not be able to connect to them.

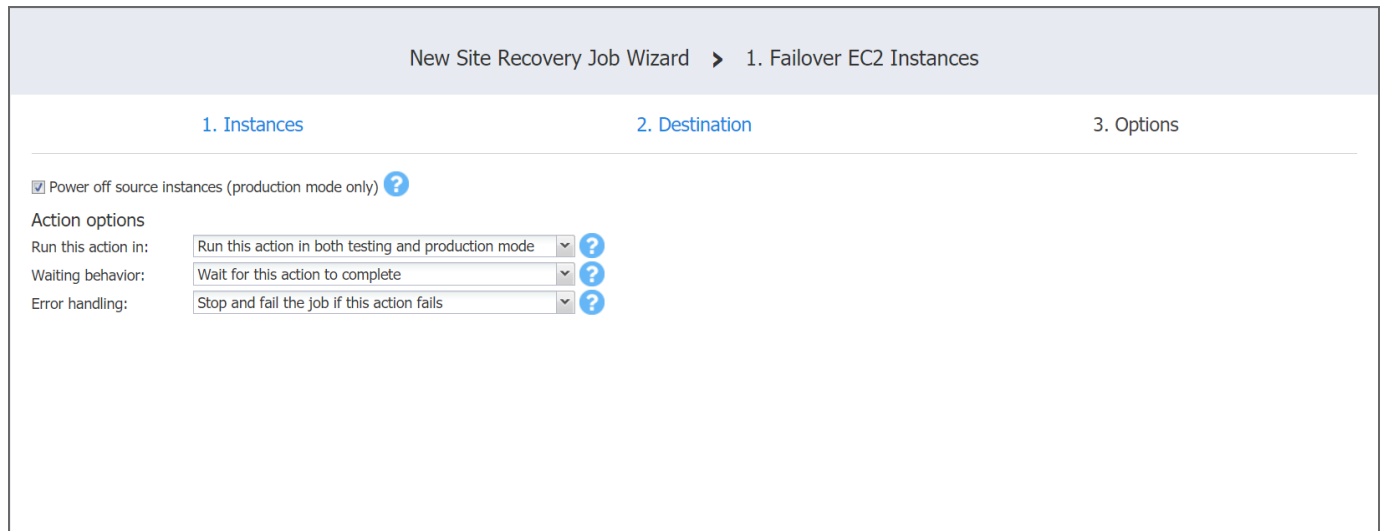
3. Click **Next** to go to the next page of the wizard.



## Failover Amazon EC2 Instances: Options

On this page of the wizard, do the following:

1. Select **Power off source Instances** 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** section, set the options for your action. Refer to [“Site Recovery Job Wizard: Actions” on page 1019](#) for details.
3. Click **Save**.



The **Options** page closes and your **Failover EC2 Instances** 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:

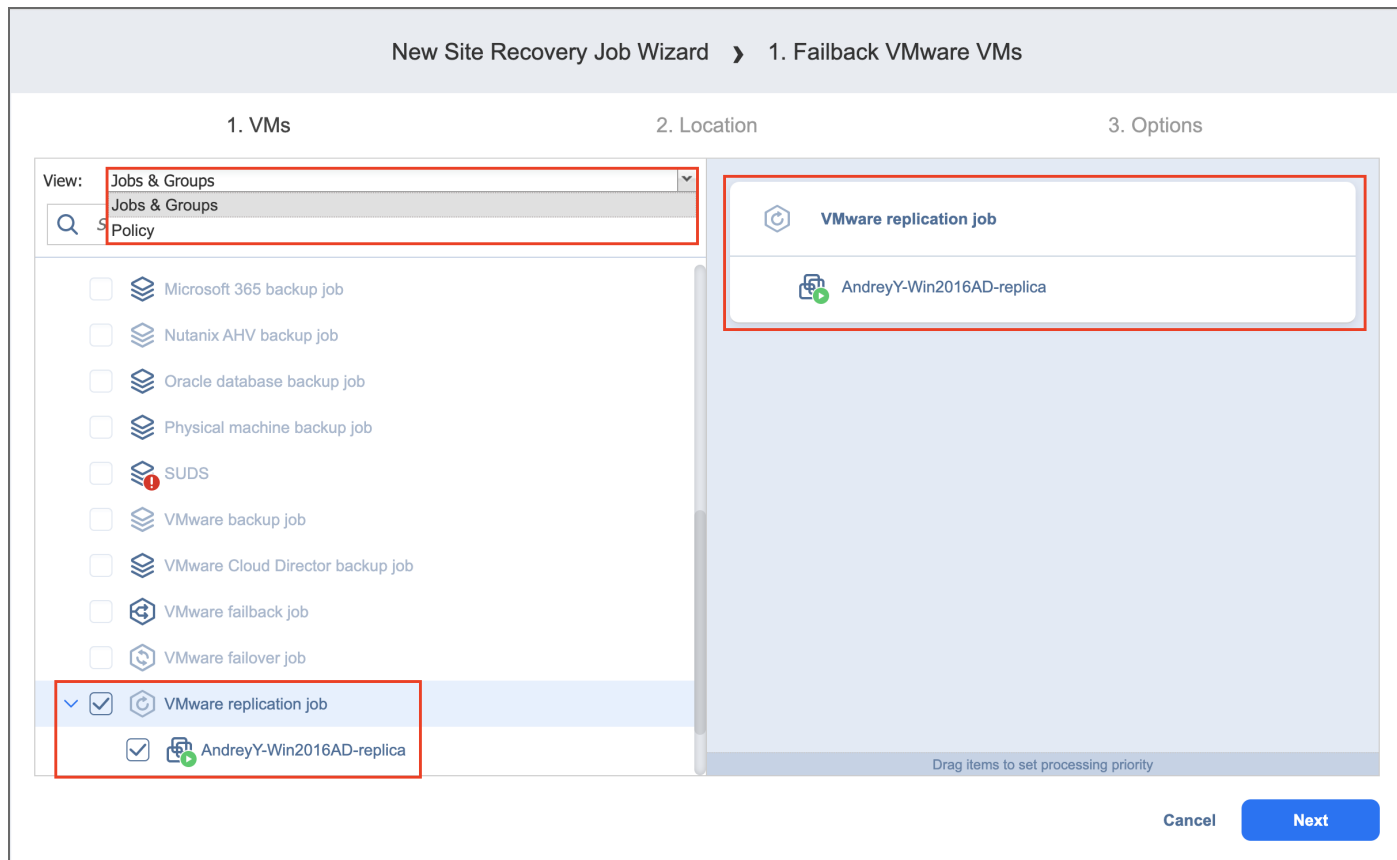
- [Failback VMware VMs: VMs](#)
- [Failback VMware VMs: Location](#)
- [Failback VMware VMs: Options](#)

## 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 266](#) 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 269](#) 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

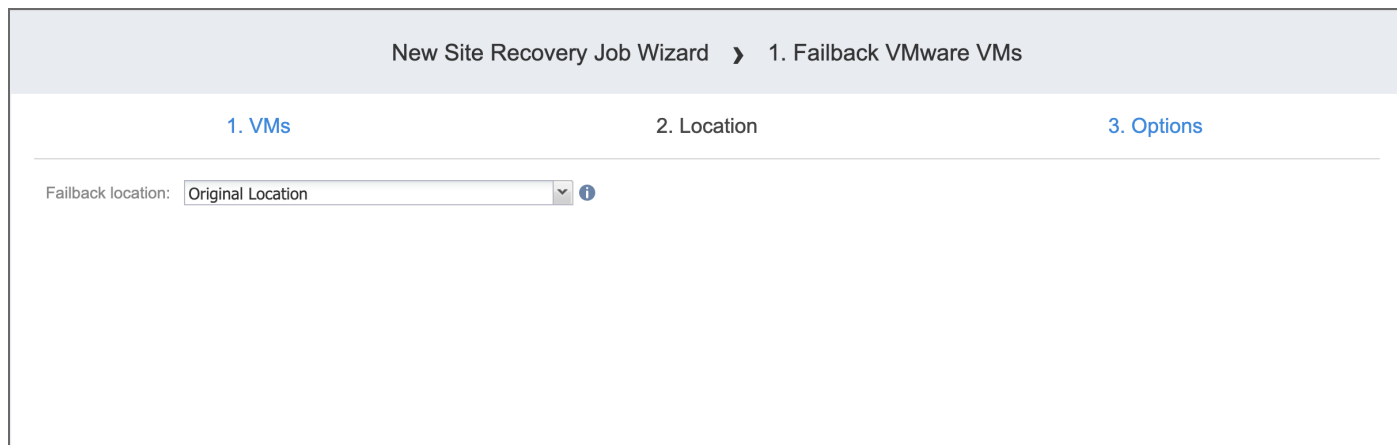
Please refer to the following subsections on providing a location for your Failback VMware VMs action:

- [“Using Original VMware Failback Location” below](#)
- [“Creating New VMware Failback Location” on the next page](#)
- [“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:

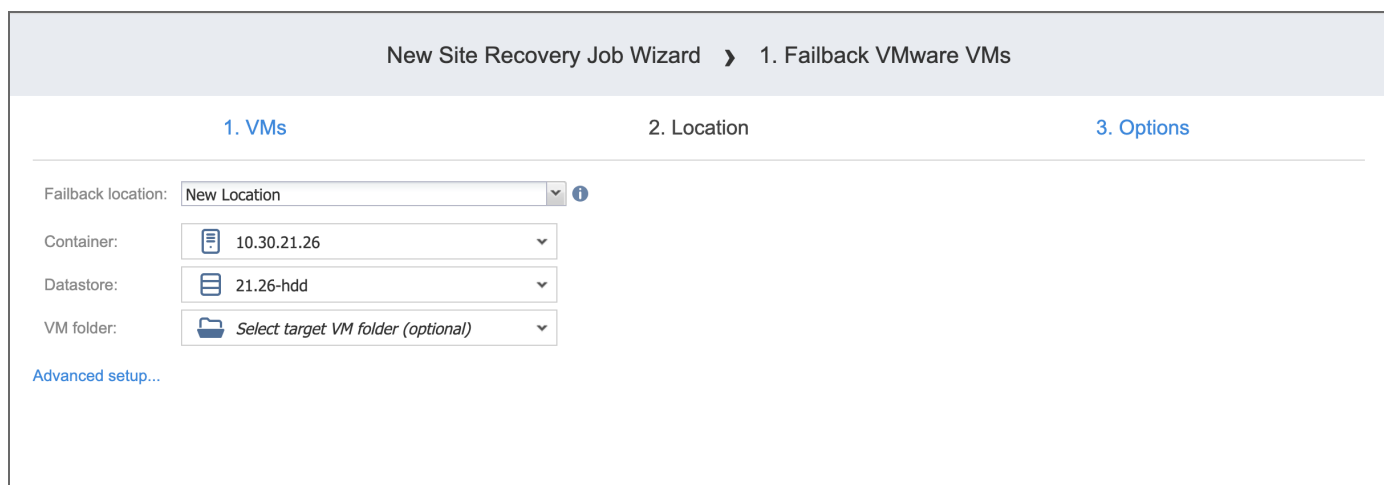
1. In the **Failback location** list, choose **Original Location**.
2. Click **Next** to go to the **Options** page.



## 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 1019](#) for details.
3. Click **Save**.

New Site Recovery Job Wizard > 1. Failback VMware VMs

1. VMs 2. Location 3. Options

Power off replica VMs (production mode only) ⓘ

Action options

Run this action in: Run this action in both testing and production mode ⓘ

Waiting behavior: Wait for this action to complete ⓘ

Error handling: Stop and fail the job if this action fails ⓘ

App-aware mode: Enabled (proceed on error) ⓘ settings

Data Transfer

Transport mode: Automatic selection ⓘ

Transporter pool: Select transporter pool ⓘ

Cancel Save

The **Options** page closes and your **Failback VMware VMs** action is added to the Site Recovery job.

### Failback Hyper-V VMs Action

Refer to the following sections on adding a *Failback Hyper-V VMs* action to your Site Recovery job:

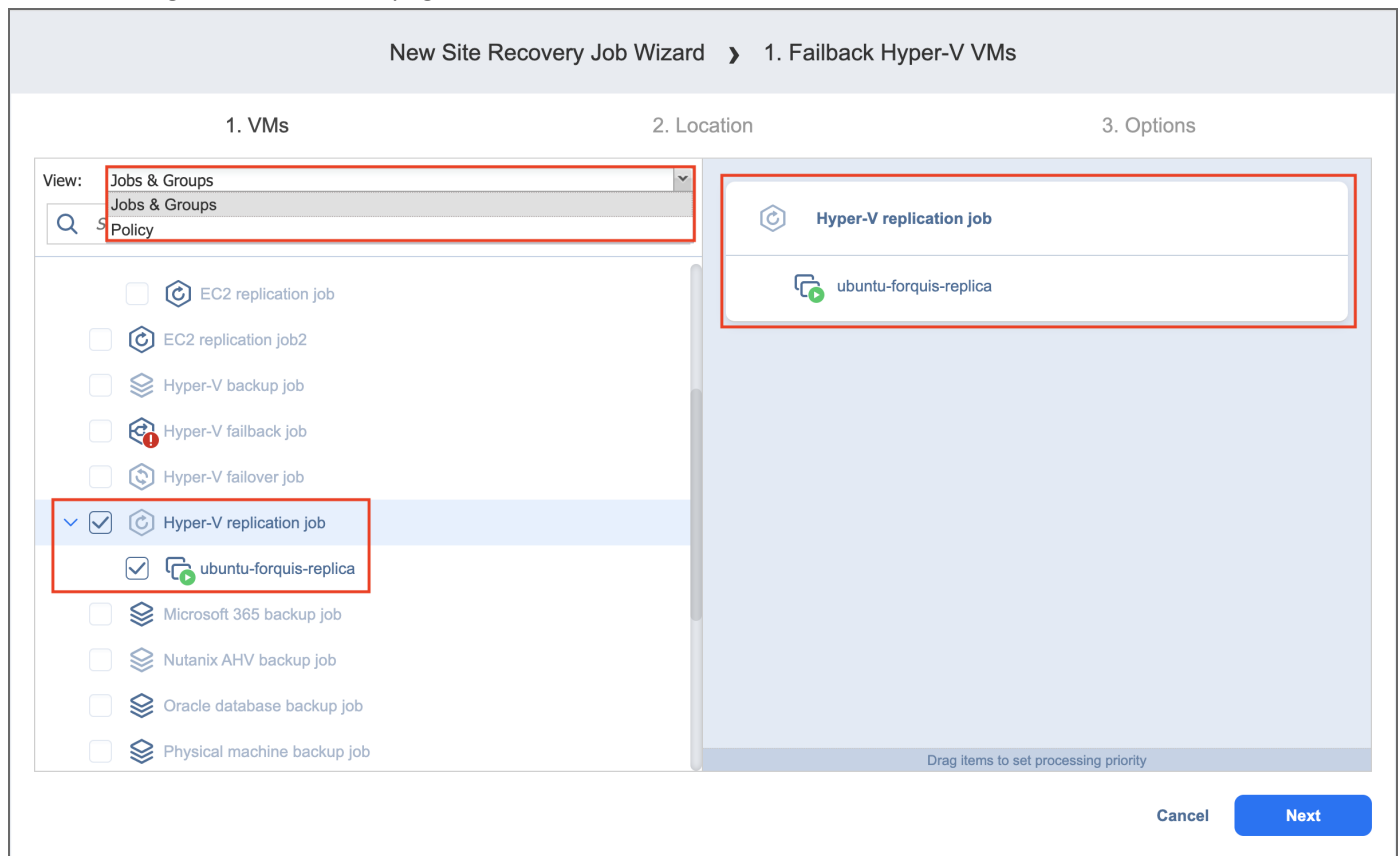
- [Failback Hyper-V VMs: VMs](#)
- [Failback Hyper-V VMs: Location](#)
- [Failback Hyper-V VMs: Options](#)



# Failback Hyper-V VMs: VMs

On the **VMs** 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 266](#) 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 269](#) for details.
2. Select one or more replica Hyper-V VMs in the left pane of the page.
3. Click **Next** to go to the **Location** page.



# Failback Hyper-V VMs: Location

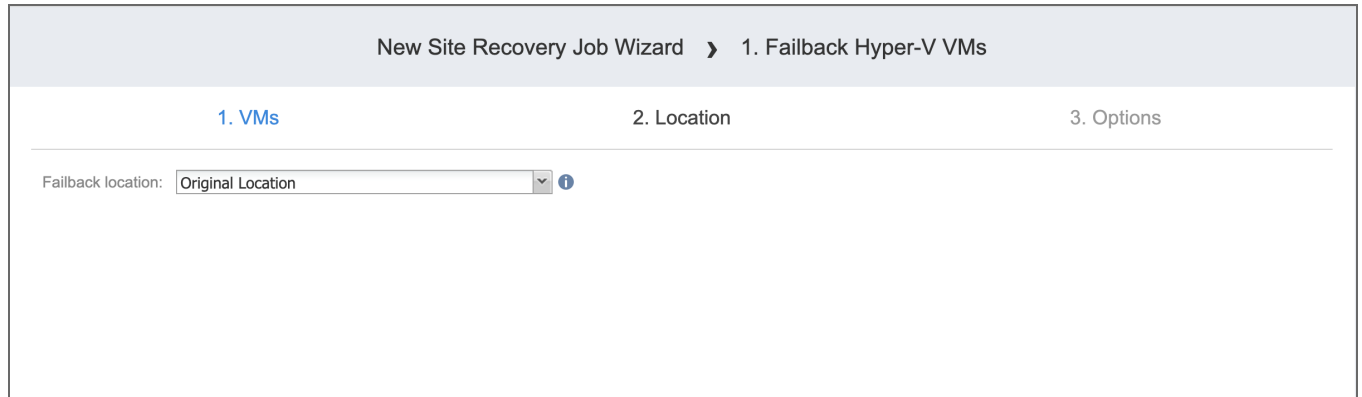
Refer to the following subsections on providing a location for your **Failback Hyper-V VMs** action:

- [“Using Original Hyper-V Failback Location” below](#)
- [“Creating New Hyper-V Failback Location” below](#)
- [“Reusing Existing Hyper-V Failback Location” on the next page](#)

## Using Original Hyper-V Failback Location

To use the original location for your **Failback Hyper-V VMs** action, follow the steps below:

1. In the **Failback location** list, choose **Original Location**.
2. Click **Next** to go to the **Options** page.



The screenshot shows a wizard interface titled "New Site Recovery Job Wizard" with a breadcrumb "1. Failback Hyper-V VMs". Below the title is a progress bar with three steps: "1. VMs" (highlighted in blue), "2. Location", and "3. Options". Underneath the progress bar, there is a label "Failback location:" followed by a dropdown menu currently showing "Original Location". To the right of the dropdown is a small blue information icon.

## Creating New Hyper-V Failback Location

To create a new location for your **Failback Hyper-V 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 **Path** box, enter a path to the location of the Hyper-V VM replica.

4. Click **Next** to go to the **Options** page.

New Site Recovery Job Wizard > 1. Failback Hyper-V VMs

1. VMs      2. Location      3. Options

Failback location: New Location ⓘ

Container: ServerHV2012

Path: C:\NakivoReplicas

ⓘ To use existing VMs as targets, expand the Advanced setup and specify target VM for each source VM.

Hyper-V replication job Click to collapse

Default container: ServerHV2012 ⓘ

Default path: C:\NakivoReplicas ⓘ

ubuntu-forquis-replica Click to collapse

Source		Target	
VM location:	ServerHV2012 > ServerHV2012	Container:	ServerHV2012
VM resources:	2 CPU, 1.0 GB RAM	Virtual Machine:	New VM will be created
Disks		Disks	
<input checked="" type="checkbox"/> e71abb5d-b86b-4...	5.8 GB	e71abb5d-b86b-4c48-a...	C:\NakivoReplicas
VM configuration		VM configuration:	C:\NakivoReplicas

Cancel Save

## Reusing Existing Hyper-V Failback Location

To reuse existing locations for your **Failback Hyper-V VMs** action, follow the steps below:

1. In the **Failback location** list, choose **New Location**.
2. Click **Advanced options**.
3. Click 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 to be used as a target.
6. After setting the location advanced options, click **Next** to go to the **Options** page.

## Failback Hyper-V VMs: Options

On the **Options** page of the wizard, do the following:

1. Select **Power off replica 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 up options for your action. Refer to [“Site Recovery Job Wizard: Actions” on page 1019](#) for details.

### 3. Click **Save**.

New Site Recovery Job Wizard > 1. Failback Hyper-V VMs

1. VMs                                      2. Location                                      3. Options

Power off replica VMs (production mode only) ⓘ

Action options

Run this action in: Run this action in both testing and production mode ⓘ

Waiting behavior: Wait for this action to complete ⓘ

Error handling: Stop and fail the job if this action fails ⓘ

App-aware mode: Enabled (proceed on error) ⓘ settings

Cancel      Save

The **Options** page closes, and your Failback Hyper-V VMs action is added to the Site Recovery job.

#### Failback Amazon EC2 Instances Action

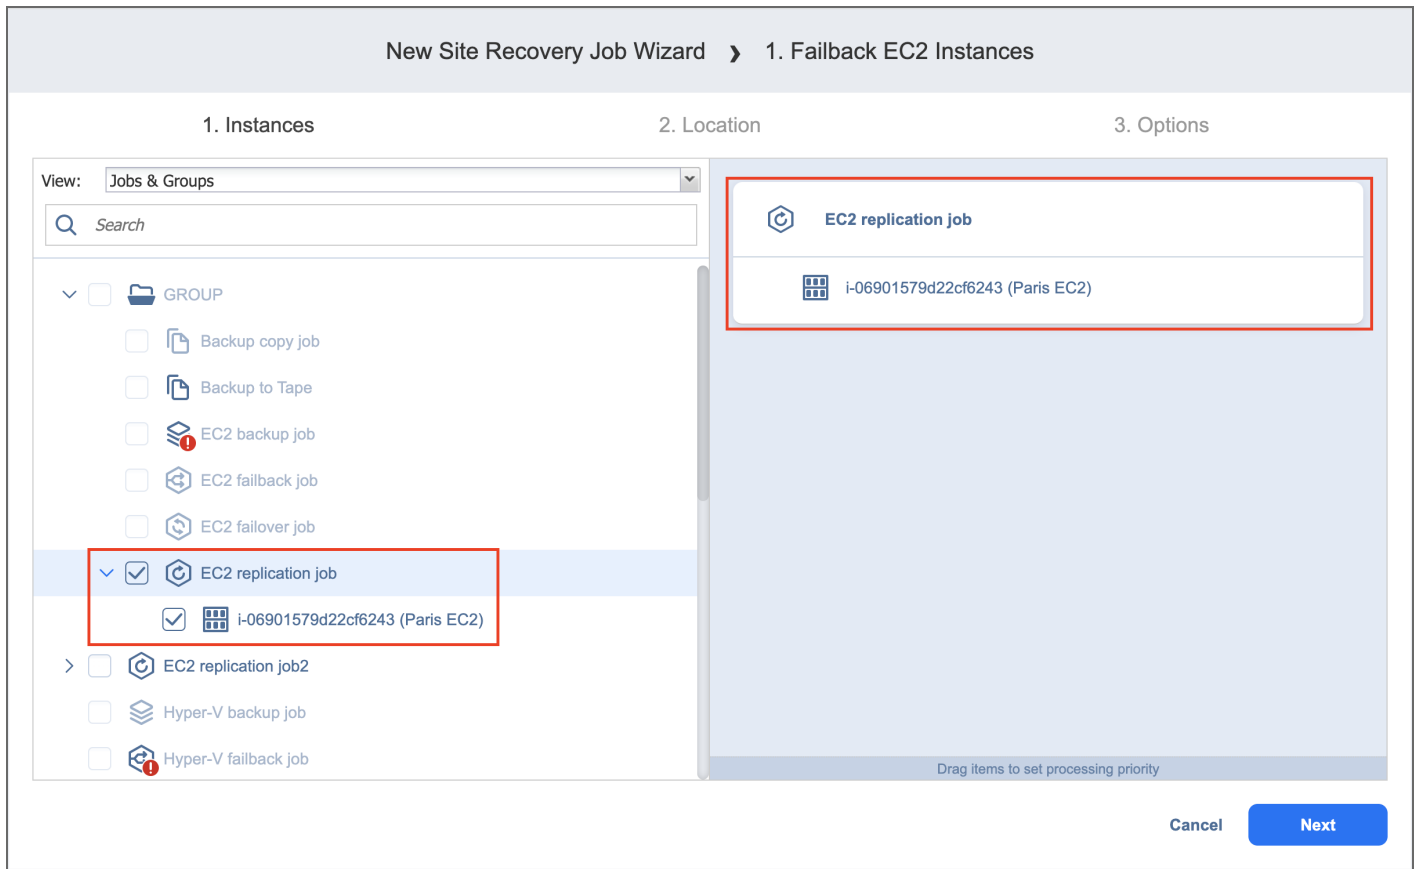
Refer to the following sections on adding a **Failback EC2 Instances** action to your Site Recovery job:

- [“Failback Amazon EC2 Instances: Instances” below](#)
- [“Failback Amazon EC2 Instances: Location” on the next page](#)
- [“Failback EC2 Instances: Options” on page 1041](#)

## Failback Amazon EC2 Instances: Instances

On this page of the wizard, do the following:

1. Select one or more replica Amazon EC2 Instances in the left pane of the page.
2. Click **Next** to go to the **Location** page.



## Failback Amazon EC2 Instances: Location

On this page of the wizard, you can specify failback location for your Amazon EC2 instances:

- [“Setting Original Amazon EC2 Failback Location” below](#)
- [“Setting Single New Amazon EC2 Failback Location” on the next page](#)
- [“Setting Multiple New Failback Locations for EC2 Instances” on page 1040](#)

### Setting Original Amazon EC2 Failback Location

To use the original location for your **Failback EC2 Instances** action, follow the steps below:

1. In the **Failback location** list, choose **Original Location**.

#### **Important**

To successfully failback an Amazon EC2 instance to its original location, the source instance should not exist.



7. Click **Next** to go to the **Options** page.

New Site Recovery Job Wizard > 1. Failback EC2 Instances

1. Instances      2. Location      3. Options

Failback location: New Location ⓘ

Region: Africa (Cape Town) ▾

Network: NAKIVO VPC (10.0.0.0/16) ▾

Subnet: NAKIVO subnet (10.0.0.0/16) ▾

Security group: Africa (Cape Town) Linux transporter\_... ▾

Instance access: Use existing key pair ▾

Key pair: Select target key pair ▾

EC2 replication job Click to collapse

Region: Africa (Cape Town) ▾

Instance type: Select target instance type ▾

Network: NAKIVO VPC (10.0.0.0/16) ▾

Subnet: NAKIVO subnet (10.0.0.0/16) ▾

Security group: Africa (Cape Town) Linux transporter\_... ▾

Cancel    Next

## Setting Multiple New Failback Locations for EC2 Instances

To set multiple new locations for your **Failback EC2 Instances** action, follow the steps below:

1. In the **Failback location** list, choose **New Location**.
2. Click **Advanced Options**.
3. Click on the required source instances to expand them.
4. For a required instance, set up the options as follows:
  - a. **Region**: Select an EC2 region to which the Amazon EC2 instance will fail back.
  - b. **Instance type**: Select a type for your failback EC2 instance.
  - c. **EBS Volumes**: Choose whether to fail back or skip individual Amazon EBS volumes.
  - d. **Network**: Select an Amazon EC2 VPC to which the EC2 instance will fail back.
  - e. **Subnet**: Select an Amazon EC2 Subnet to which the Amazon EC2 Instance will fail back.
  - f. **Security group**: Select an Amazon EC2 Security Group for the Amazon EC2 Instance.
  - g. **Instance access**: Specify Amazon EC2 Key Pair options as follows:
    - **Use existing key pair**: Choose an existing Amazon EC2 Key Pair that you want to use to connect to the failback EC2 instance.

- **Create a new key pair:** Select this option to create a new Amazon EC2 Key Pair and assign it to the failback EC2 instance. Enter a name for the new Amazon EC2 Key Pair in the **Key pairname** box and click **Download**.

**Important**

This is your only chance to save the private key file, so make sure you save the private key file in a safe place. You will need to provide the name of your key pair when you start recovered EC2 instances and the corresponding private key each time you connect to the instances.

- **Proceed without key pair:** Select this option if you do not want to assign an Amazon Key Pair to the failback Amazon EC2 instance.

**Important**

A key pair enables you to connect to Linux Instances via SSH. If you recover Linux Instances without a key pair, you will not be able to connect to them.

5. After setting the location advanced options, click **Next** to go to the **Options** page.

## Failback EC2 Instances: Options

On this page of the wizard, do the following:

1. Select **Power off replica Instances** 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 [Options Common to Most Actions](#) for details.



3. Click **Save**.

New Site Recovery Job Wizard > 1. Failback EC2 Instances

1. Instances      2. Location      3. Options

Power off replica VMs (production mode only) ⓘ

Action options

Run this action in: Run this action in both testing and production mode ⓘ

Waiting behavior: Wait for this action to complete ⓘ

Error handling: Stop and fail the job if this action fails ⓘ

App-aware mode: Disabled ⓘ

Cancel      Save

The **Options** page closes and your **Failback EC2 Instances** 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 Hyper-V VMs
- Start 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 EC2 instance in the inventory tree.
4. Set the action options. Refer to [“Site Recovery Job Wizard: Actions” on page 1019](#) 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 Hyper-V VMs
- Stop 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 266](#) 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 EC2 instance in the inventory tree.
4. Set action options. Refer to [“Site Recovery Job Wizard: Actions” on page 1019](#) 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**

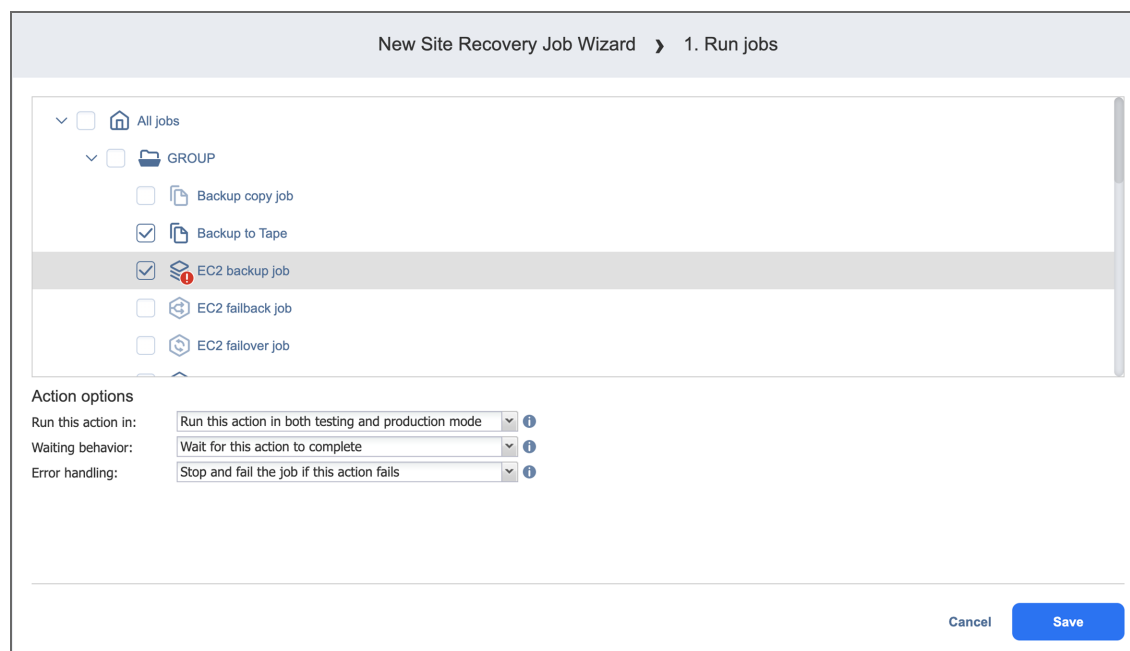
The topic covers the following actions of a Site Recovery Job:

- 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 1019](#) 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 1019](#) for details.

3. Click **Save**.

The screenshot shows the 'New Site Recovery Job Wizard' at step '1. Run script'. The 'Script Options' section includes: 'Target type' (Remote Linux server), 'Target server' (192.168.77.11), 'Script path' (/home/user/nakivo/parse.sh), 'Use custom SSH port' (checked, 22), 'Username' (user), and 'Password' (masked). The 'Action options' section includes: 'Run this action in' (both testing and production mode), 'Waiting behavior' (Wait for this action to complete), and 'Error handling' (Stop and fail the job if this action fails). A 'Manage credentials' link is visible below the password field.

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 screenshot shows the 'New Site Recovery Job Wizard' at step '1. Attach repository'. The 'Repository to attach' dropdown is set to 'Repo 3'. The 'Action options' section includes: 'Run this action in' (both testing and production mode), 'Waiting behavior' (Wait for this action to complete), and 'Error handling' (Stop and fail the job if this action fails). 'Cancel' and 'Save' buttons are at the bottom right.

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 1019](#) for details.
7. Click **Save**.

New Site Recovery Job Wizard > 1. Send email

To:

Cc:

Subject:

Attachments:  [Browse...](#)

Action options

Run this action in:

Waiting behavior:

Error handling:

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 1019](#) for details.

3. Click **Save**.

New Site Recovery Job Wizard > 1. Wait

Time to wait: 5 Minutes

Action options

Run this action in: Run this action in both testing and production mode ?

Waiting behavior: Wait for this action to complete ?

Error handling: Stop and fail the job if this action fails ?

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**.

New Site Recovery Job Wizard > 1. Check condition

Condition	
Condition type:	<input type="text" value="Resource exists"/>
Resource type:	<input type="text" value="VMware VM"/>
Identification method:	<input type="text" value="Name"/> <input type="text" value="Equals"/>
Search string:	<input type="text" value="Vmware"/>
Action if True	
Action type:	<input type="text" value="Continue site recovery job"/>
Action if False	
Action type:	<input type="text" value="Stop and fail site recovery job"/>

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:



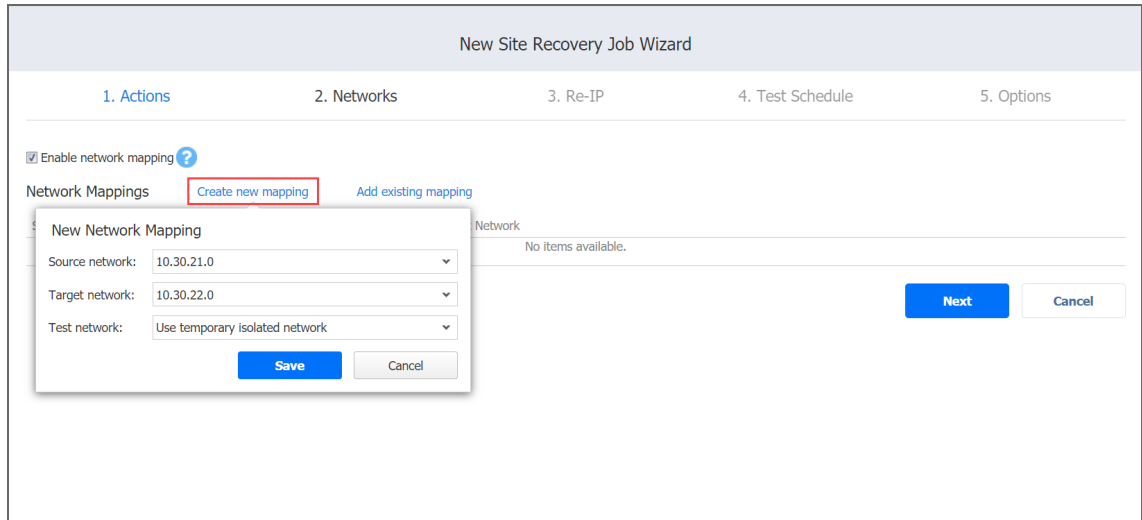
1. Select **Enable network mapping**.

**Note**

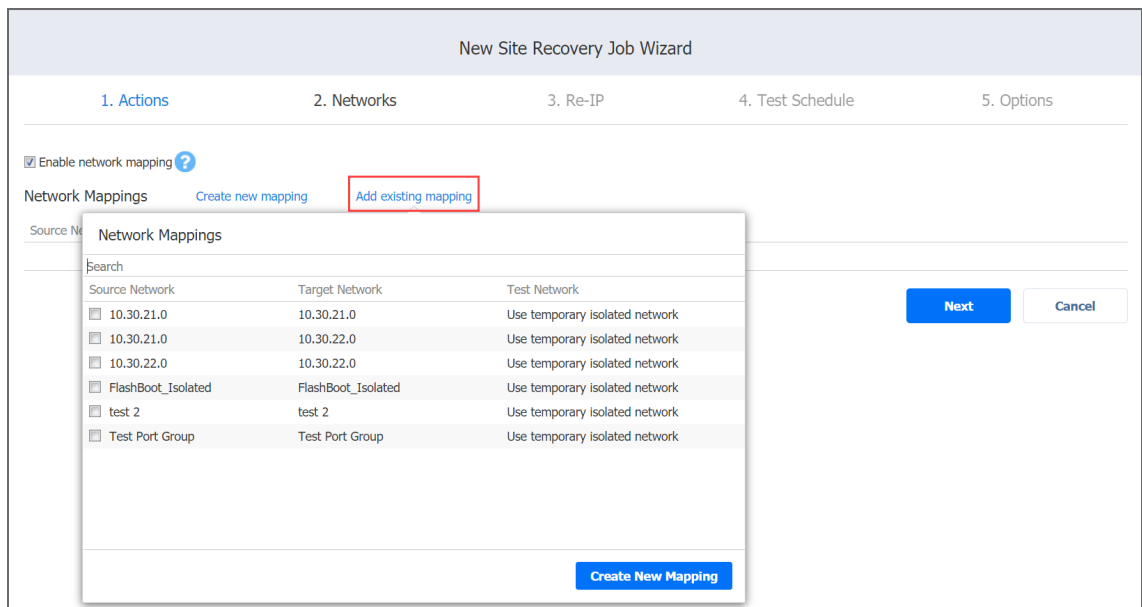
A failover or a failback action needs to be on the actions list to allow enabling network mapping for your site recovery job.

2. The *Network Mapping* section opens. You have the following options:

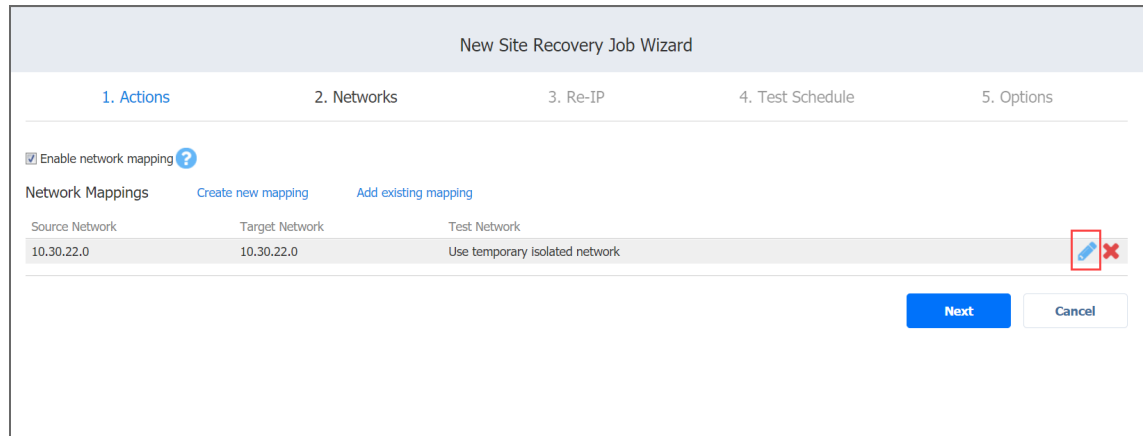
- Create a new network mapping:
  - a. Click **Create new mapping**.
  - b. The **New Network Mapping** dialog opens. Choose a source network, a target network and a test network, and click **Save**.



- 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

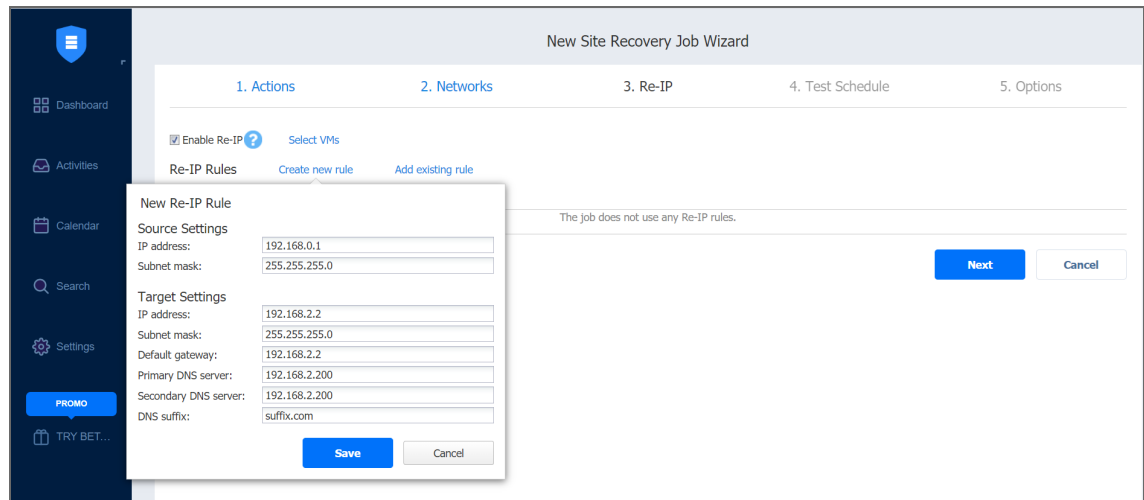
A failover or a failback action needs to be in the actions list to allow enabling Re-IP for your site recovery job.

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**.

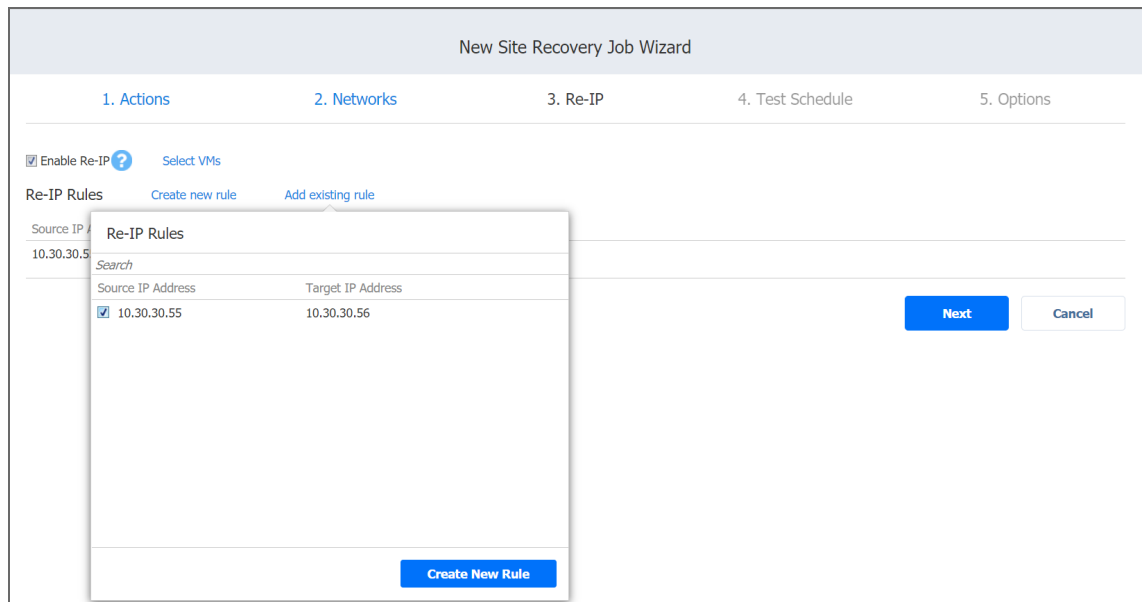


## Note

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**.

The screenshot shows the 'New Site Recovery Job Wizard' at step 3, 'Re-IP'. The 'Enable Re-IP' checkbox is checked. Below it, a table lists Re-IP Rules with columns for 'Source IP Address' and 'Target IP Address'. A rule is listed with source IP 10.30.30.55 and target IP 10.30.30.56. To the right, the 'Edit Re-IP Rule' dialog is open, showing the following settings:

Setting	Value
Source Settings	
IP address:	10.30.30.55
Subnet mask:	255.255.255.0
Target Settings	
IP address:	10.30.30.56
Subnet mask:	255.255.255.0
Default gateway:	10.30.30.5
Primary DNS server:	10.30.30.10
Secondary DNS server:	10.30.30.11
DNS suffix:	zenlar.int

Buttons for 'Save' and 'Cancel' are visible at the bottom of the dialog.

- 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.

The screenshot shows the 'New Site Recovery Job Wizard' at step 3, 'Re-IP'. The 'Enable Re-IP' checkbox is checked. Below it, a table lists Re-IP Rules with columns for 'Source IP Address' and 'Target IP Address'. A rule is listed with source IP 10.30.30.55 and target IP 10.30.30.56. To the right, the 'Edit Re-IP Rule' dialog is open, showing the following settings:

Setting	Value
Source Settings	
IP address:	10.30.30.55
Subnet mask:	255.255.255.0
Target Settings	
IP address:	10.30.30.56
Subnet mask:	255.255.255.0
Default gateway:	10.30.30.5
Primary DNS server:	10.30.30.10
Secondary DNS server:	10.30.30.11
DNS suffix:	zenlar.int

Buttons for 'Next' and 'Cancel' are visible at the bottom of the dialog.

- 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

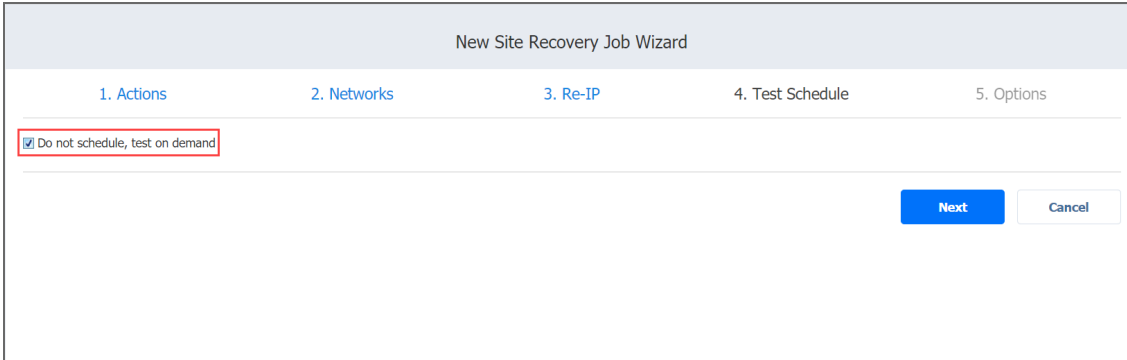
On the **Test Schedule** page of the **Site Recovery Job Wizard** you can schedule testing your site recovery job.

Please refer to the following sections for details:

- [“Disabling Site Recovery Job Test Schedule” on the next page](#)
- [“Daily Site Recovery Job Testing” on the next page](#)
- [“Monthly or Yearly Site Recovery Job Testing” on the next page](#)
- [“Periodic Site Recovery Job Testing” on page 1055](#)
- [“Chained Site Recovery Job” on page 1056](#)
- [“Additional Schedule” on page 1057](#)

## 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.



New Site Recovery Job Wizard

1. Actions    2. Networks    3. Re-IP    4. Test Schedule    5. Options

Do not schedule, test on demand

Next    Cancel

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**.



New Site Recovery Job Wizard

2. Networks    3. Re-IP    4. Test Schedule    5. Options

Job schedule will come into effect on the selected date. Pick a date in the calendar below.

June 2020

M	T	W	T	F	S	S
1	2	3	4	5	6	7
8	9	10	11	12	13	14
15	16	17	18	19	20	21
22	23	24	25	26	27	28
29	30	1	2	3	4	5
6	7	8	9	10	11	12

Apply

Effective from date

Add another schedule

Show calendar

Next    Cancel

## 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**.

New Site Recovery Job Wizard

1. Actions      2. Networks      3. Re-IP      4. Test Schedule      5. Options

Do not schedule, test on demand

(UTC+02:00, EET) Eastern European Time

Schedule #1

Run monthly/yearly

Run every last Friday of every month

Starting at: 0:00    Ending: 6:00

Effective from

[Add another schedule](#)

[Show calendar](#)

**Next**    Cancel

## 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**.

New Site Recovery Job Wizard

1. Actions    2. Networks    3. Re-IP    4. Test Schedule    5. Options

Do not schedule, test on demand  
(UTC+02:00, EET) Eastern European Time

Schedule #1

Run periodically every 50 minutes

Starting at: 0:00 Ending: 6:00

Mon  Tue  Wed  Thu  Fri  Sat  Sun

[All days](#) [Work days](#) [Weekends](#)

Effective from

[Add another schedule](#)  
[Show calendar](#)

**Next** Cancel

## 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**.

New Site Recovery Job Wizard

1. Actions    2. Networks    3. Re-IP    4. Test Schedule    5. Options

Do not schedule, test on demand  
(UTC+02:00, EET) Eastern European Time

Schedule #1

Run after another job

After the job: EC2 backup job

Run this job: Immediately

After successful runs    After failed runs    After stopped runs

Effective from

[Add another schedule](#)  
[Show calendar](#)

**Next**    Cancel

## 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**.

Do not schedule, test on demand  
(UTC+02:00, EET) Eastern European Time

Schedule #1 [Remove](#)

Run after another job

After the job: EC2 backup job

Run this job: Immediately

After successful runs    After failed runs    After stopped runs

Effective from

Schedule #2 [Remove](#)

Run daily/weekly

Starting at: 0:00    Ending: 6:00

Mon    Tue    Wed    Thu    Fri    Sat    Sun

[All days](#)   [Work days](#)   [Weekends](#)

every 1 weeks

Effective from

[Add another schedule](#)  
[Show calendar](#)

**Next**    Cancel

## Site Recovery Job Wizard: Options

On the **Options** page of the **Site Recovery Job Wizard** you can specify options of your Site Recovery Job.

Proceed as follows:

1. In the *Site Recovery Job* section, specify a name for your Site Recovery Job.
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, it enables sending Email notifications to the specified recipients. The semi-colon character should be used to separate multiple email addresses.

3. Click **Finish** to complete creating your Site Recovery Job.

New Site Recovery Job Wizard

1. Actions      2. Networks      3. Re-IP      4. Test Schedule      5. Options

Site Recovery Job  
Job name:

Testing Options  
Recovery time objective:     
 Send test/run report to:

The **Site Recovery Job Wizard** will close and your Site Recovery Job will appear in the list of NAKIVO Backup & Replication jobs.

## Running Site Recovery Job

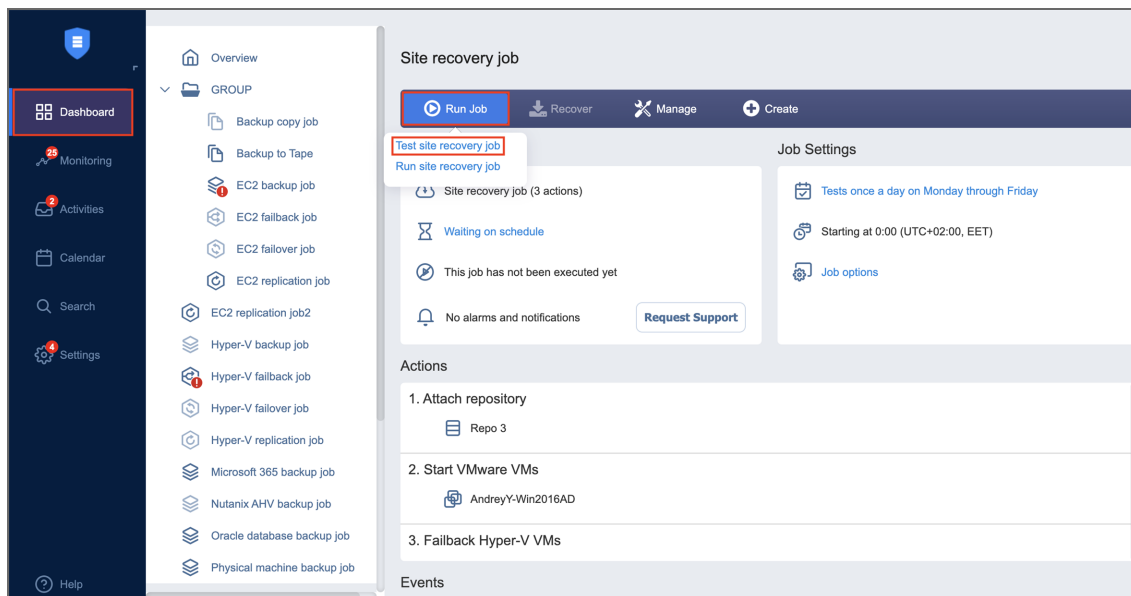
The section includes the following topics:

- [Running Site Recovery Job in Test Mode](#)
- [Running Site Recovery Job in Production Mode](#)

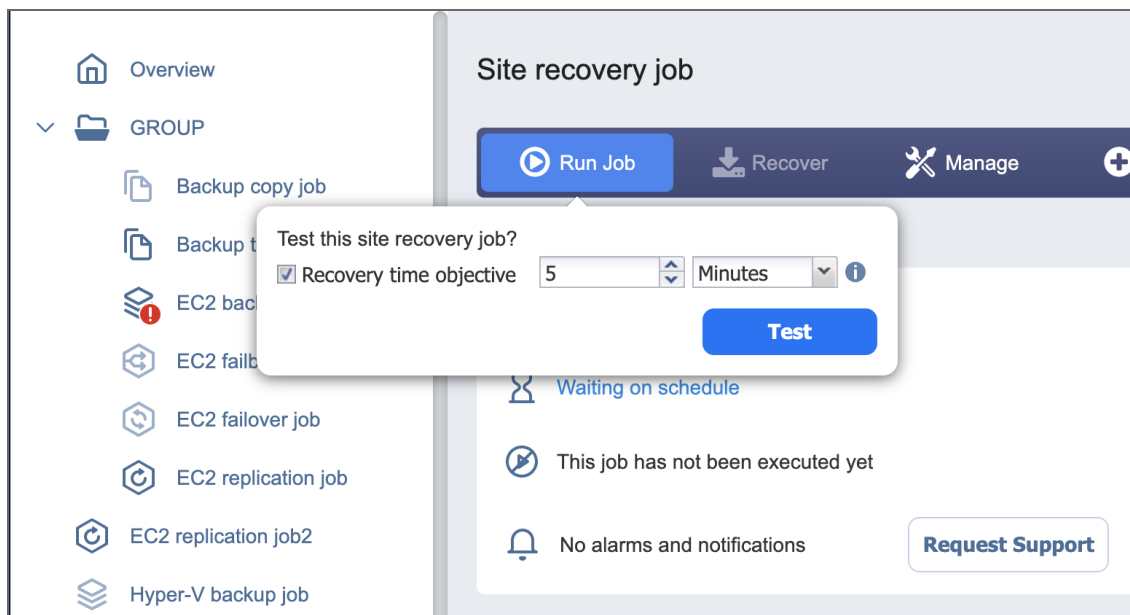
### Running Site Recovery Job in Test Mode

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. On the Dashboard, select your Site Recovery Job and then click the **Run Job** 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

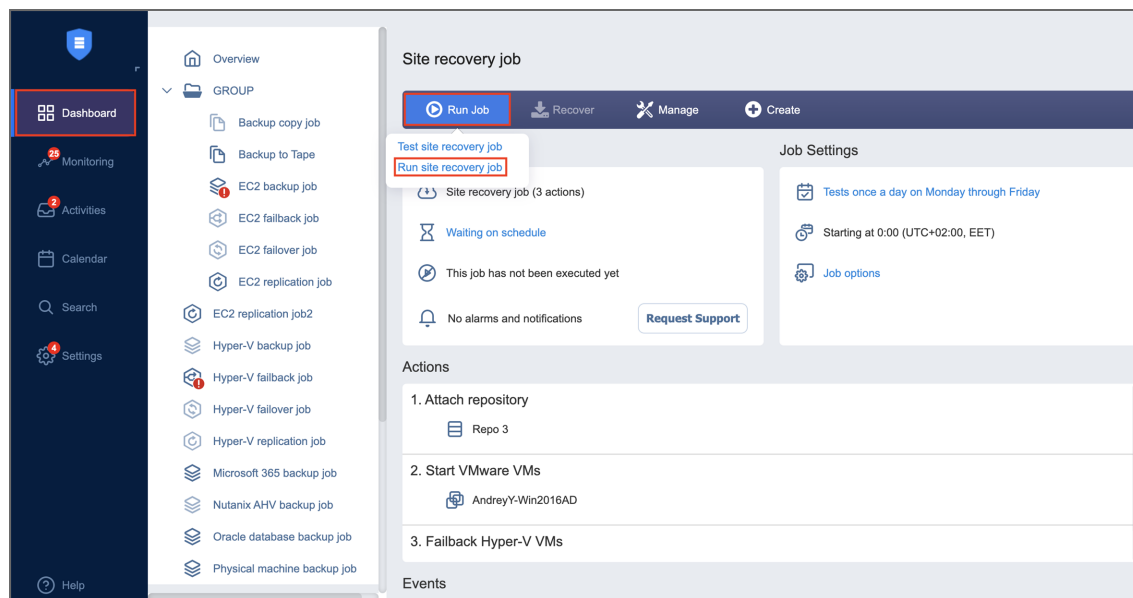
In addition to testing the site recovery job on demand, testing can also be scheduled. Refer to [Site Recovery Job Wizard: Test Schedule](#) for details.

## 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:

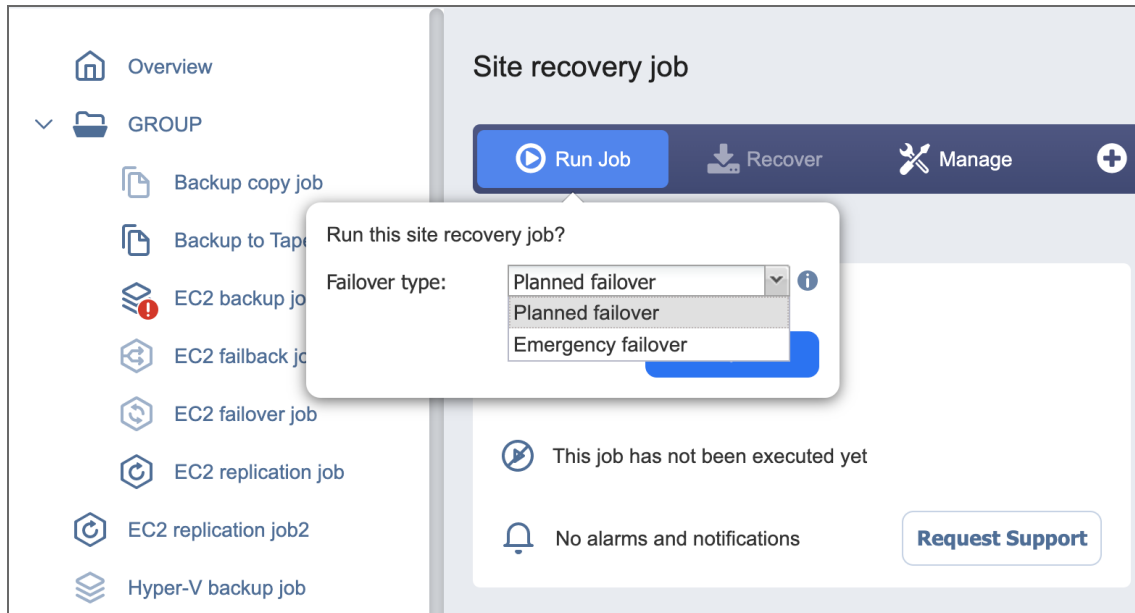
1. In the **Jobs** dashboard, select your Site Recovery Job and then click the **Run Job** button.
2. In the dialog that opens, click **Run site recovery job**.



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.

# Monitoring

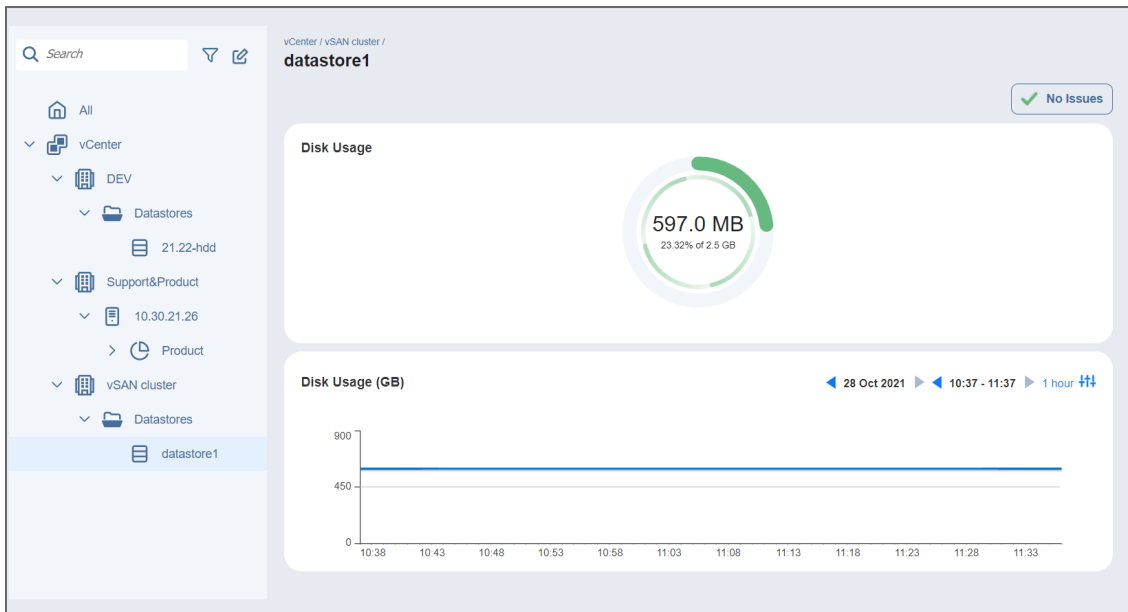
Refer to the following topics to learn how to use Monitoring feature:

- [“Navigating The Monitoring Dashboard” below](#)
- [“Managing The Monitored Items” on the next page](#)

## Navigating The Monitoring Dashboard

On the left side of the page, you can find the following:

- **Search bar:** Enter the partial or full name of an item to search for it among your inventory items.
- **Filter:** Click the button to filter the type of items you want to be displayed. These include the following:
  - Hosts
  - Datastores
  - VMs
- **Edit:** Manually select the items you want to be monitored among your inventory or add them according to a [specific policy](#).
- **Inventory tree:** Your monitored items are displayed here. Click on any monitored item of a supported type to view its information.



On the right side of the Monitoring dashboard, you can find the notification bar and the relevant performance charts for the monitored inventory items. The displayed information is as follows:

- VMs
  - **CPU Load:** Displays the CPU load of the selected monitored item in percentage at the current time.

- **Memory Load:** Displays the memory usage of the selected monitored item in percentage at the current time.
- **Disk Usage:** Displays the disk usage of the selected monitored item in gigabytes at the current time.
- **CPU Usage (%):** Displays the CPU load of the selected monitored item in percentage over a specified period of time.
- **Memory Usage (GB):** Displays the memory usage of the selected monitored item in gigabytes over a specified period of time.
- **Disk Usage (GB):** Displays the disk usage of the selected monitored item in gigabytes over a specified period of time.
- Hosts
  - **CPU Load:** Displays the CPU load of the selected monitored item in percentage at the current time.
  - **Memory Load:** Displays the memory usage of the selected monitored item in gigabytes at the current time.
  - **CPU Usage (%):** Displays the CPU load of the selected monitored item in percentage over a specified period of time.
  - **Memory Usage (GB):** Displays the memory usage of the selected monitored item in gigabytes over a specified period of time.
- Datastores
  - **Disk Usage:** Displays the disk usage of the selected monitored item in gigabytes at the current time.
  - **Disk Usage (GB):** Displays the disk usage of the selected monitored item in gigabytes over a specified period of time.

## Managing The Monitored Items

On this page, you can find out how to manage the items you are monitoring. Learn how to add items to the monitoring dashboard, set up chart time frames, and how data collection and data aggregation works.

- [Adding items to the Monitoring dashboard](#)
- [Setting up chart time frames](#)
- [Data collection](#)
- [Data aggregation](#)

## Adding Items to the Monitoring Dashboard

Items are added to the **Monitoring** dashboard depending on the installed license.

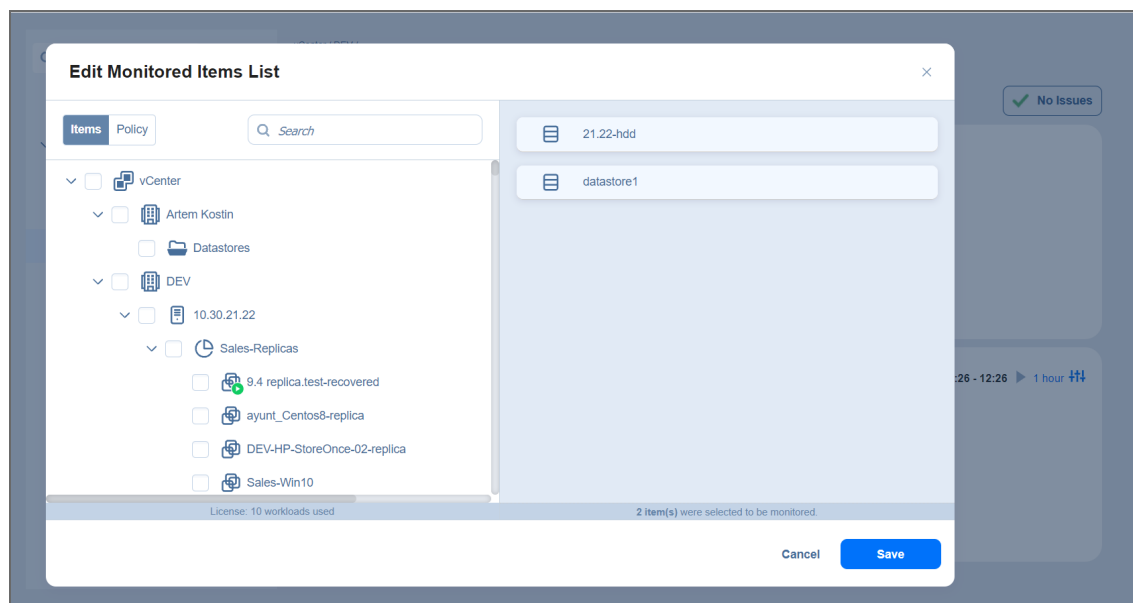
## Note

To use the monitoring function, the following conditions must be met:

- Monitoring must be enabled by the current license;
  - The items of the [supported platform](#) must be [added to the inventory](#).
- 
- If a Trial license is installed, the product automatically adds top-level containers of supported platforms in the Inventory to the list of monitored items on the first discovery of an Inventory item.
  - If a non-Trial license is installed, the product allows for adding and editing items to the list of monitored items manually.


When first accessing the dashboard, a **No item is being monitored message** is displayed. Click **Add monitored items...** to open the **Edit** popup:

1. On the left side, choose the items you want to add to the monitoring dashboard.
2. Alternatively, you can set up [specific policies](#) for the automated addition of the inventory items to the dashboard.
3. Click **Save** once you're finished.



## Setting Up Chart Time Frames

It is possible to select specific time frames for displaying historical CPU, memory, and disk usage charts. Click on the rightmost icon on the chart and choose among the following options:



**1 hour:** The data is displayed for the 1-hour time frame. The specific hour can then be chosen on the **Monitoring** dashboard.

**1 day:** The data is displayed for the 1-day time frame. Select the specific day on the Monitoring dashboard.

**7 days:** The data is displayed for the 7-days time frame. Select the specific days on the Monitoring dashboard.

**30 days:** The data is displayed for the 30-days time frame. Select the specific days on the Monitoring dashboard.

**365 days:** The data is displayed for the 365-days time frame. Select the specific days on the Monitoring dashboard.

**Custom:** Selecting this option opens the **Calendar** popup where you can select the specific time frame for the chart. Click **Apply** once you're done.

## Data Collection

NAKIVO Backup & Replication collects data every minute. The following exceptions apply:

- If the product cannot collect data every 1 minute for 5 minutes, data collection switches to an interval of every 5 minutes. This mode collects 5 data points at once.
- If the product cannot collect data every 5 minutes for 10 minutes, data collection switches to an interval of every 10 minutes. This mode collects 10 data points at once.

For VM disk usage, VM memory usage, and datastore disk usage, multiple data points are created by copying a single data point.

## Data Aggregation

NAKIVO Backup & Replication creates a data point every minute. As the length of the archive increases, the data points are aggregated in the following way:



- 1-minute data points are aggregated into 10-minute data points.
- 10-minute data points are aggregated into 1-hour data points.
- After a year, 1-hour data points are removed.

# Integration and Automation

This section contains the following topics:

- [“Command Line Interface” below](#)
- [“Automation with HTTP API” on page 1075](#)
- [“Aptare IT Analytics Integration” on page 1075](#)

## Command Line Interface

This section covers the following topics:

- [“Using Command Line Interface” below](#)
- [“Available Commands” on the next page](#)
- [“Exit Codes” on page 1075](#)

## Using Command Line Interface

- [“Operation Modes of Command Line Interface” below](#)
- [“Using Command Line Interface Locally” below](#)
- [“Using Command Line Interface Remotely” on the next page](#)
- [“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 the previous page](#) topic.

Use either long or short form of the commands\*.

Command	Long form	Short form	Output
Help			

Command	Long form	Short form	Output
General help	cli.bat --help	cli.bat -h	<ul style="list-style-type: none"> <li>• Command name</li> <li>• Description</li> </ul>
<b>Job Management</b>			
List all jobs	cli.bat --job-list	cli.bat -jl	<ul style="list-style-type: none"> <li>• Job ID</li> <li>• Job name</li> <li>• Current job status</li> <li>• Job last run result</li> </ul>
Start a job	cli.bat --job-start [job_id]	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]	
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	cli.bat --job-report [job_id] <ul style="list-style-type: none"> <li>• The command with no arguments creates the job report and saves it to the current directory.</li> <li>• To save the report to other directory: cli.bat --job-report [job_id] --save-to [dir_path]</li> <li>• To send the report to default email(s): cli.bat --job-report [job_id] --send-by-email</li> </ul>	cli.bat -jp [job_id] <ul style="list-style-type: none"> <li>• The command with no arguments creates the job report and saves it to the current directory.</li> <li>• To save the report to other directory: cli.bat -jp [job_id] -f [dir_path]</li> <li>• To send the report to default email(s): cli.bat -jp [job_id] -eml</li> </ul>	

Command	Long form	Short form	Output
	<ul style="list-style-type: none"> <li>To send the report to other email: cli.bat. --job-report [job_id] --send-by-email [email_address]</li> </ul>	<ul style="list-style-type: none"> <li>To send the report to other email: cli.bat. -jp [job_id] -eml [email_address]</li> </ul>	
Return information about a job	cli.bat --job-info [job_id]	cli.bat -ji [job_id]	<ul style="list-style-type: none"> <li>Job ID</li> <li>Job name</li> <li>Current job status</li> <li>Job last run result</li> </ul>
<b>Inventory</b>			
List all inventory items	cli.bat --inventory-list	cli.bat -il	<ul style="list-style-type: none"> <li>Item ID</li> <li>Item IP/host name</li> <li>Item type (host/vCenter)</li> <li>Item children count (X hosts, E VMs)</li> <li>Item current state</li> <li>Item current status</li> </ul>
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]	<ul style="list-style-type: none"> <li>Item ID</li> <li>Item IP/host name</li> <li>Item type (host/vCenter)</li> <li>Item children count (X hosts, E VMs)</li> <li>Item current state</li> <li>Item current status</li> </ul>
<b>Transporters</b>			
List all transporters	cli.bat --transporter-list	cli.bat -trl	<ul style="list-style-type: none"> <li>Transporter ID</li> </ul>

Command	Long form	Short form	Output
			<ul style="list-style-type: none"> <li>• Transporter IP/host name</li> <li>• Transporter current load</li> <li>• Transporter maximum load</li> <li>• Transporter current state</li> <li>• Transporter current status</li> </ul>
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]	<ul style="list-style-type: none"> <li>• Transporter ID</li> <li>• Transporter IP/host name</li> <li>• Transporter current load</li> <li>• Transporter maximum load</li> <li>• Transporter current state</li> <li>• Transporter current status</li> </ul>
<b>Repositories</b>			
List all repositories	cli.bat --repository-list	cli.bat -rl	<ul style="list-style-type: none"> <li>• Repository ID</li> <li>• Repository name</li> <li>• Assigned transporter</li> <li>• Backup count</li> <li>• Free space</li> <li>• Attached or detached</li> <li>• Consistent or inconsistent</li> </ul>

Command	Long form	Short form	Output
			<ul style="list-style-type: none"> <li>Repository current state</li> <li>Repository current status</li> </ul>
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] Parameters: <ul style="list-style-type: none"> <li>--selfheal</li> <li>--verify</li> <li>--spacereclaim</li> </ul>	cli.bat -rm [repo_id] [parameter] Parameters: <ul style="list-style-type: none"> <li>--selfheal</li> <li>--verify</li> <li>--spacereclaim</li> </ul>	
Stop repository maintenance	cli.bat --repository-maintenance-stop [repo_id]	cli.bat -rms [repo_id]	
Return information about a repository	cli.bat --repository-info [repo_id]	cli.bat -ri [repo_id]	<ul style="list-style-type: none"> <li>Repository ID</li> <li>Repository name</li> <li>Assigned transporter</li> <li>Backup count and free space</li> <li>Attached or detached</li> <li>Consistent or inconsistent</li> <li>Repository current state</li> </ul>

Command	Long form	Short form	Output
			<ul style="list-style-type: none"> <li>Repository current status</li> </ul>
<b>Support</b>			
Generate a support bundle	<p>cli.bat --bundle-create</p> <ul style="list-style-type: none"> <li>The command with no parameters will create a support bundle and save it in the current directory.</li> <li>To save the bundle to other directory: cli.bat --bundle-create --save-to [dir_path]</li> <li>To send the bundle to support over email: cli.bat --bundle-create --send-to-support</li> <li>To send the bundle to other email: cli.bat --bundle-create --send-by-email [email_address]</li> </ul>	<p>cli.bat -bc</p> <ul style="list-style-type: none"> <li>The command with no parameters will create a support bundle and save it in the current directory.</li> <li>To save the bundle to other directory: cli.bat -bc -f [dir_path]</li> <li>To send the bundle to support over email: cli.bat -bc -sup</li> <li>To send the bundle to other email: cli.bat -bc -eml [email_address]</li> </ul>	
<b>Licensing</b>			
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]	
<b>Multi Tenancy</b>			
List all tenants	cli.bat --tenant-list	cli.bat -tl	<ul style="list-style-type: none"> <li>Tenant ID</li> <li>Tenant name</li> </ul>



Command	Long form	Short form	Output
			<ul style="list-style-type: none"> <li>Allocated items type and count</li> <li>Tenant status</li> <li>Enabled or disabled</li> </ul>
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]	<ul style="list-style-type: none"> <li>Tenant ID</li> <li>Tenant Account ID</li> <li>Tenant name</li> <li>Allocated items type and count</li> <li>Tenant status</li> <li>Enabled or disabled</li> </ul>
Create a support bundle for master admin level	<ul style="list-style-type: none"> <li>Generate the support bundle for master level only: cli.bat --bundle-create</li> <li>Generate the support bundle with all tenants logs: cli.bat --bundle-create --include-tenants</li> </ul>	<ul style="list-style-type: none"> <li>Generate the support bundle for master level only: cli.bat -bc</li> <li>Generate the support bundle with all tenants logs: cli.bat -bc -ite</li> </ul>	
<b>Miscellaneous</b>			
Get the CLI version	cli.bat --version The command returns the CLI version which is equal to the full version of NAKIVO Backup & Replication.	-	

Command	Long form	Short form	Output
Run a command in the debug mode	<pre>cli.bat --repository-info [repo_id] --debug</pre> <p>This is an option that can be added to any other CLI command.</p> <p>With the debug mode turned on, the commands will return the full error text.</p>	<pre>cli.bat -ri [repo_id] -- debug</pre>	

\*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:

Name	Domain	Enabled	Policy State	Collector State	Status	Last Modified	Collector Version	Upgrade Manager	Auto Upgrade	Host	Notes
HienH_Nakivo	Nakivo	Yes		Offline	Offline	06:40:31 09-01-2020	10.2.30.01 (06222018-1511)		Full auto	WIN-KPHCKG2HAU25	
Hoop	Nakivo	Yes		Offline	Offline	11:58:10 17-09-2019	10.3.2.01 (09032019-2031)		Full auto	WIN-GSBJUWPKPQ3	
Nakivo	Nakivo	Yes		Offline	Offline	06:25:08 26-02-2020	10.3.9.01 (02242020-0534)	10.3.9.01	Full auto	DESKTOP-DMHJLR	
NAKIVO_TW	Nakivo	Yes		Offline	Offline	11:31:15 03-03-2020			Full auto		
ND	Nakivo	Yes		Online	Online	11:21:45 03-03-2020	10.3.9.01 (02242020-0534)	10.3.9.01	Full auto	WIN-6102LLHVFJ	
NganT	Nakivo	Yes		Offline	Offline	04:56:23 26-02-2020	10.3.9.01 (02242020-0534)	10.3.9.01	Full auto	WIN-88U7HSRLD2R	
TrangH	Nakivo	Yes		Online	Offline	03:43:22 26-02-2020	10.3.9.01 (02242020-0534)	10.3.9.01	Full auto	WIN-FRAC2H1CT	



- a. Add a Collector. For details, refer to the [Managing and Monitoring Data Collection](#) subsection of the APTARE IT Analytics User Guide.

The screenshot shows the 'Collector Administration' page with the 'Add Collector' dialog box open. The dialog box contains the following fields and options:


- Collector Name: \*
- Passcode: \*
- Domain: Nakivo
- Auto-upgrade aptare.jar: Yes
- Auto-upgrade Upgrade Manager: Yes
- Enable SSL:
- Short Notes: \*

Buttons: OK, Cancel, Help


- 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.

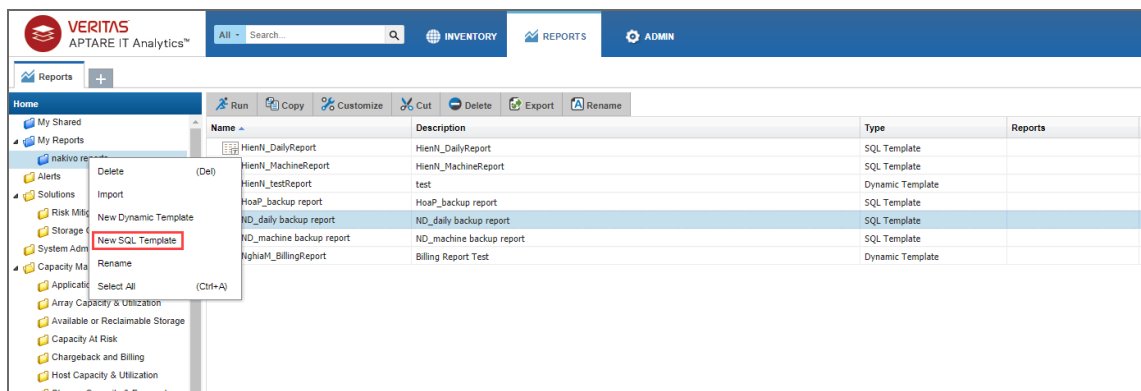


5. Go to the **REPORTS** tab in the APTARE portal and take the following actions:




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.



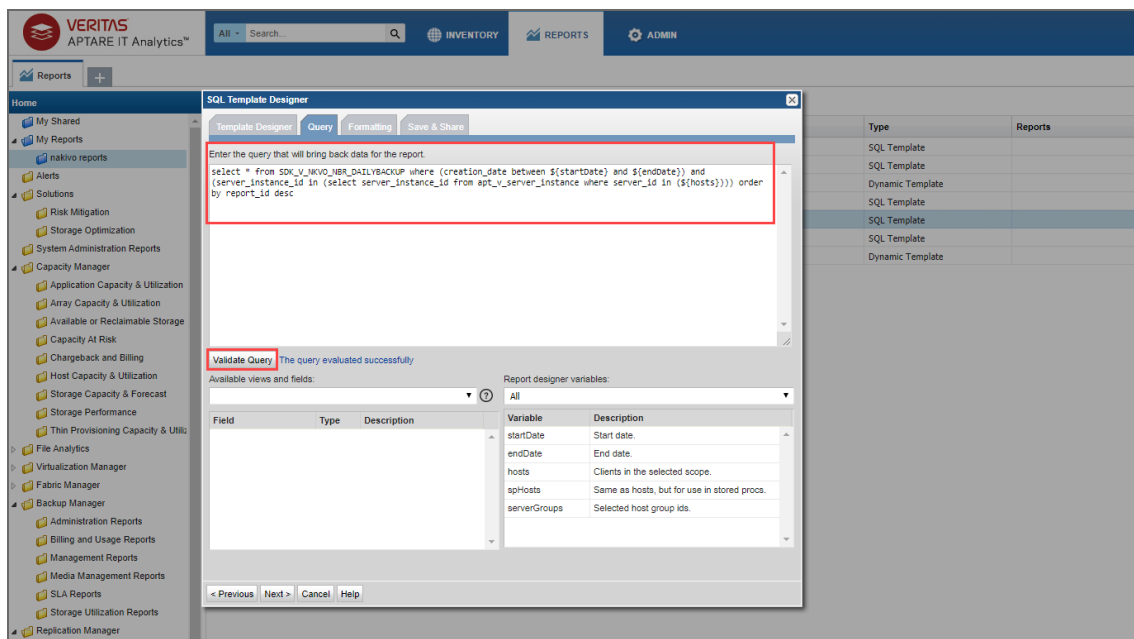
- iii. Build an SQL query for your machine backup report or daily backup report, use either:

### Daily Backup Report

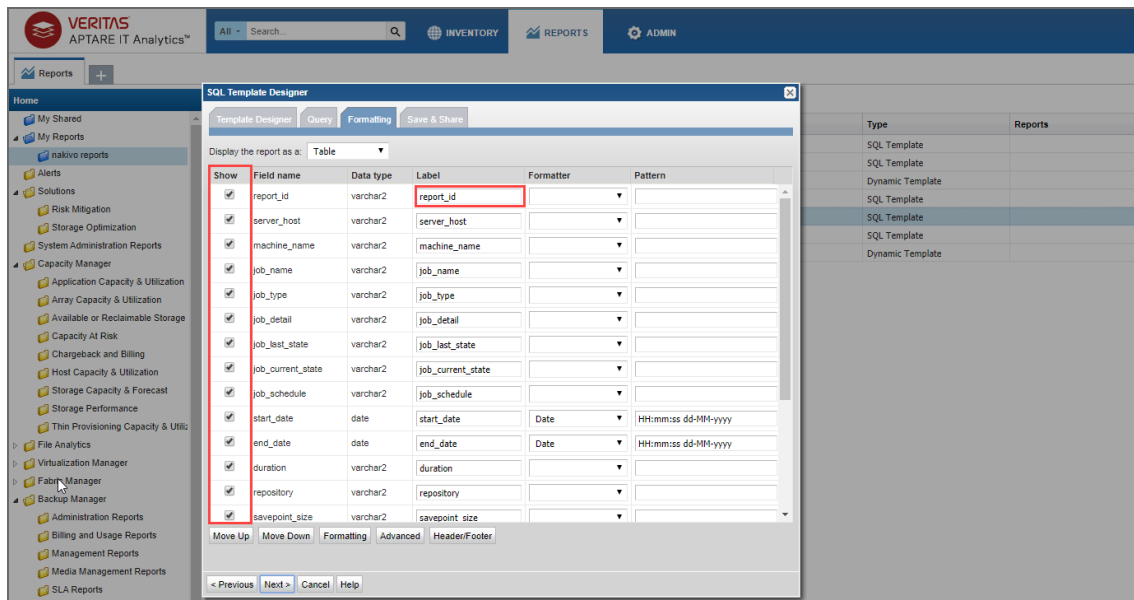
```
select * from SDK_V_NKVO_NBR_DAILYBACKUP where (creation_date
between ${startDate} and ${endDate}) and (server_instance_id in
(select server_instance_id from apt_v_server_instance where server_
id in (${hosts}))) order by report_id desc
```

### 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:


- [“Tenant Creation” below](#)
- [“Tenant Configuration” on page 1087](#)
- [“Tenant Management” on page 1088](#)
- [“Granting Self-Service Access” on page 1094](#)

## Tenant Creation

This section covers the topics describing the tenant creation process in NAKIVO Backup & Replication.

To create a new tenant, follow the steps below:

1. Log in to NAKIVO Backup & Replication as a Master Admin.
2. Click **Create New Tenant**.



3. Complete the wizard as described in the topics below to finish the tenant creation process:
  - [“Tenant Creation Wizard: Tenant” below](#)
  - [“Tenant Creation Wizard: Inventory” on the next page](#)
  - [“Tenant Creation Wizard: Transporters” on page 1084](#)
  - [“Tenant Creation Wizard: Repositories” on page 1085](#)
  - [“Tenant Creation Wizard: Users” on page 1086](#)

## Tenant Creation Wizard: Tenant

On this page of the wizard, you can provide a name for the tenant, assign licenses to the tenant, and enter contact information for the tenant. Additionally, master tenant can enable VM Limitation for the new tenants. 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 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 tenant.
  - b. In the **Microsoft 365 users allocated** field, enter the number of Microsoft 365 users you want to assign to the tenant.
5. In case the Perpetual license is installed, do the following:
  - a. In the **Sockets allocated** field, enter the number of sockets you want to assign to the tenant.
    - a. 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.
    - b. Enter the number of protected VMs for the tenant.

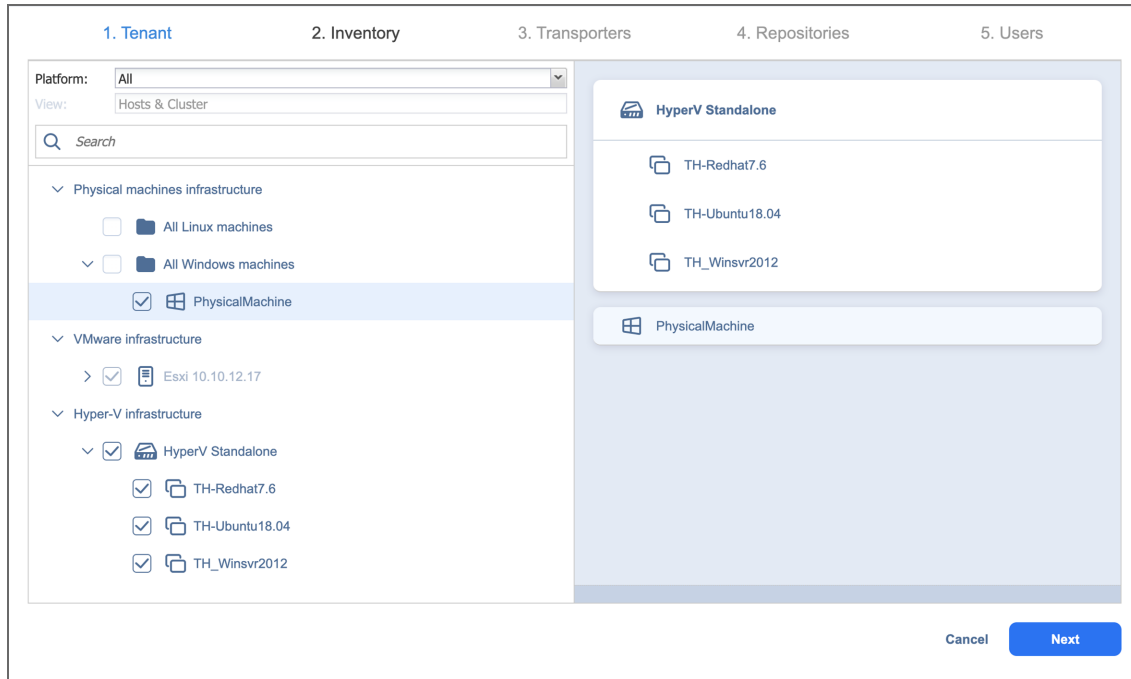
**Note**

Even with VM limitation enabled, the 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 tenant.
  - c. In the **Physical workstations allocated** field, enter the number of physical workstation licenses you want to assign to the tenant.
  - d. In the **Microsoft 365 users** allocated field, enter the number of Microsoft 365 users you want to assign to the tenant.
  - e. In the **Oracle databases** allocated field, enter the number of Oracle Database licenses you want to assign to the tenant.
6. Optionally, in the **Contact email** field, enter the email address of the tenant.
7. Optionally, in the **Contact phone** field, enter the phone number of the tenant.
8. Optionally, in the **Website field**, enter the website URL of the tenant.
9. Optionally, in the **Address** field, enter the address of the tenant.
10. Click **Next** to proceed to the **Inventory** page.

## Tenant Creation Wizard: Inventory

On this page, you can assign inventory items to the 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 be assigned to the tenant. The selected items appear in the right pane.



4. Click **Next** to proceed to the **Transporters** page.

## Tenant Creation Wizard: Transporters

On this page of the wizard, you can assign the Transporters that the tenant will be able to use for backup, recovery, and replication jobs. Proceed as follows:

1. In the **Search** field, you can enter either a part or the entire name of the Transporter to find the specific ones you need.

### Note

If you assigned an inventory item with the dependant Transporter to the tenant on the Inventory page of the wizard, that Transporter would not be selected automatically, and it cannot be deselected. If an inventory item with the dependant Transporter was not assigned to a tenant, that Transporter cannot be selected on this page.

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 other's data.

- **Maximum load per tenant:** The maximum number of tasks that the Transporter is able to perform at the same time per each assigned tenant.

3. The selected Transporters appear in the right pane. Click **Next** to proceed.

The screenshot shows the '3. Transporters' step of the wizard. It features a search bar at the top left. Below it is a table with three columns: 'Name', 'Assigned tenants', and 'Maximum load per tenant'. The table lists three transporters: 'HyperV Standalone' (checked), 'Onboard transporter' (unchecked), and 'PhysicalMachine' (checked). To the right of the table is a list of selected transporters, which currently shows 'HyperV Standalone' and 'PhysicalMachine'. At the bottom right, there are 'Cancel' and 'Next' buttons.

Name	Assigned tenants	Maximum load per tenant
<input checked="" type="checkbox"/> HyperV Standalone	0	6
<input type="checkbox"/> Onboard transporter	1	6
<input checked="" type="checkbox"/> PhysicalMachine	0	6

## Tenant Creation Wizard: Repositories

On this page of the wizard, you can assign Backup Repositories that the 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.

### Note

If the dependent Transporter was not chosen on the Transporters page of the wizard, the Backup Repositories assigned to this Transporter would not be available for selection.

2. On the left pane of the screen, you can select the Backup Repositories to be assigned to the 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.

## Tenant Creation Wizard: Users

On this page of the wizard, you can create 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 Tenant Creation Wizard.

User name	Role	Group	Two-factor authentication	
Local	Self-service user	Local users	disabled	<a href="#">Edit</a> <a href="#">Delete</a>

[Create local user](#) Cancel [Finish](#)

## 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 245](#) for a description of the initial configuration wizard.

Tenants ▶ User1

- 1. Inventory
- 2. Transporters
- 3. Repositories

[Add New...](#)

[Next](#)

# Tenant Management

This section covers the following topics:

- [“Using Filters” below](#)
- [“Using Labels” on the next page](#)
- [“Viewing Tenant Information” on page 1091](#)
- [“Opening Tenant Dashboard” on page 1091](#)
- [“Disabling Tenants” on page 1092](#)
- [“Editing Tenants” on page 1093](#)
- [“Deleting Tenants” on page 1093](#)

## Using Filters

- [About Filters](#)
- [Applying Filters](#)
- [Dismissing Filters](#)

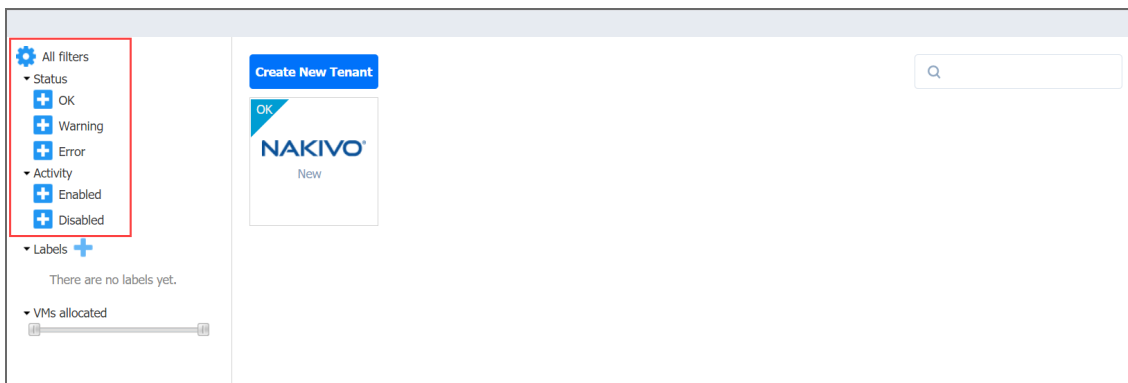
## About Filters

NAKIVO Backup & Replication comes with four built-in filters that allow you to quickly display tenants according to their state. The following filters are available:

- **OK:** Displays tenants that have no errors and notifications.
- **Warning:** Displays only tenants that have notifications.
- **Error:** Displays only tenants that have errors.
- **Enabled:** Displays only enabled tenants.
- **Disabled:** Displays only disabled tenants.

## Applying Filters


To apply a filter, click on the filter name.



The filters that are currently applied are displayed under the **Active Filters**.

## Dismissing Filters

To dismiss a filter, click the filter name under **Active filters**.



## Using Labels

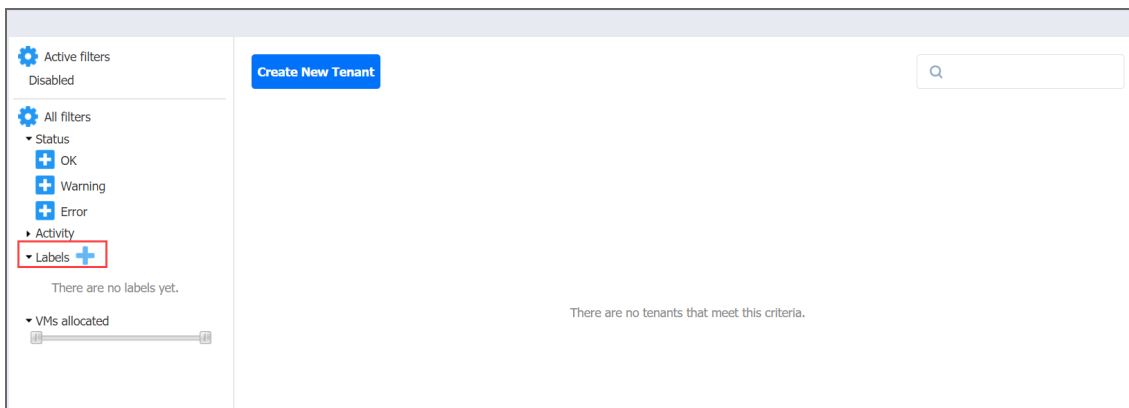
- [About Labels](#)
- [Creating Labels](#)
- [Assigning Labels to Tenants](#)
- [Editing Label Names](#)
- [Deleting Label](#)

### 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.

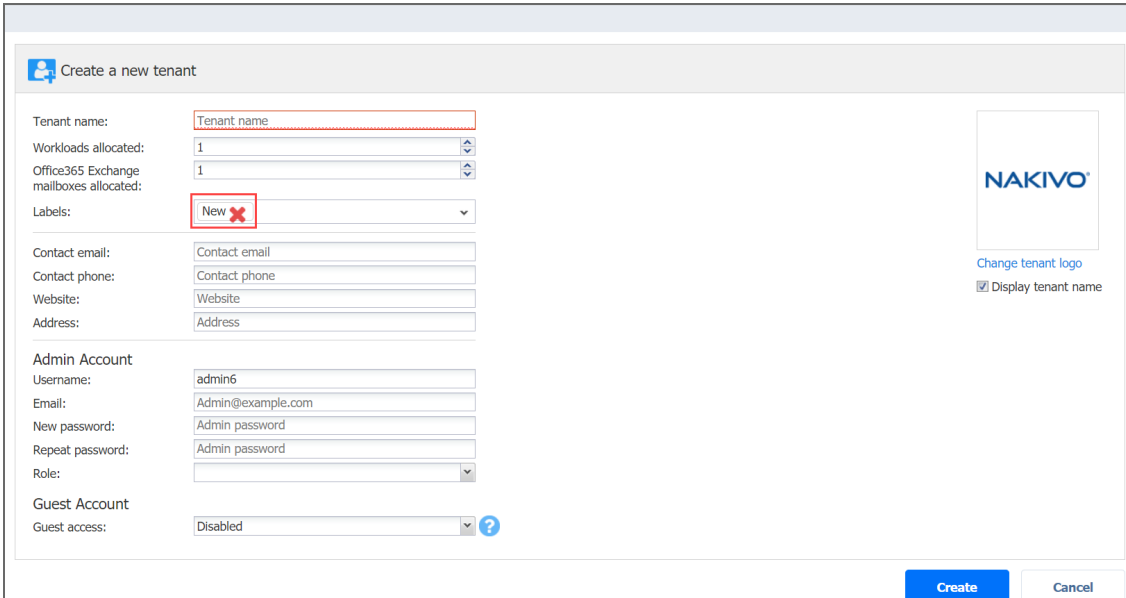


You can also create a new label when creating a new tenant.



## Assigning Labels to Tenants

You can assign a label to a tenant either during the tenant creation or by [editing the tenant](#).

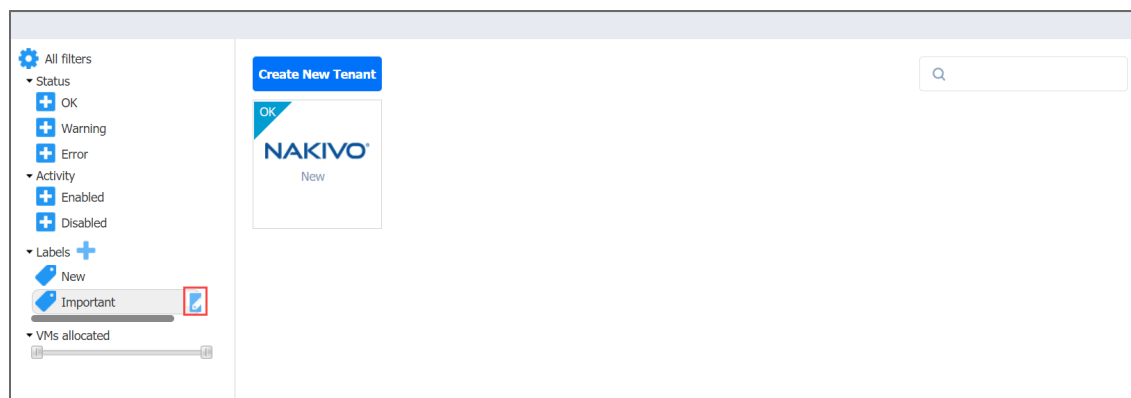


The screenshot shows the 'Create a new tenant' form. The 'Labels' dropdown menu is open, and the 'New' option is selected and highlighted with a red box. The form includes fields for Tenant name, Workloads allocated, Office365 Exchange mailboxes allocated, Contact email, Contact phone, Website, Address, Admin Account (Username, Email, New password, Repeat password, Role), and Guest Account (Guest access). A 'Create' button is at the bottom right.

## 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:

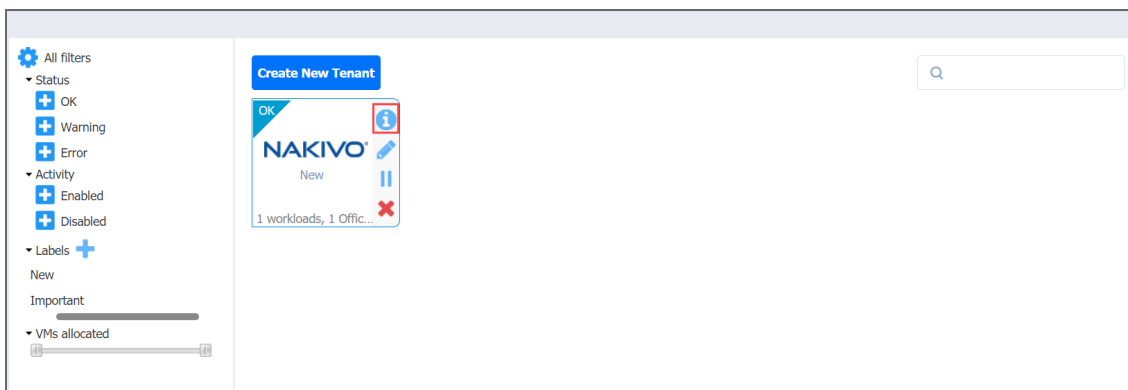
1. Hover the mouse pointer over a label.
2. Click the **Delete** icon.

3. In the dialog box that opens, click **Delete** to confirm that you wish to permanently delete the label



## Viewing Tenant Information

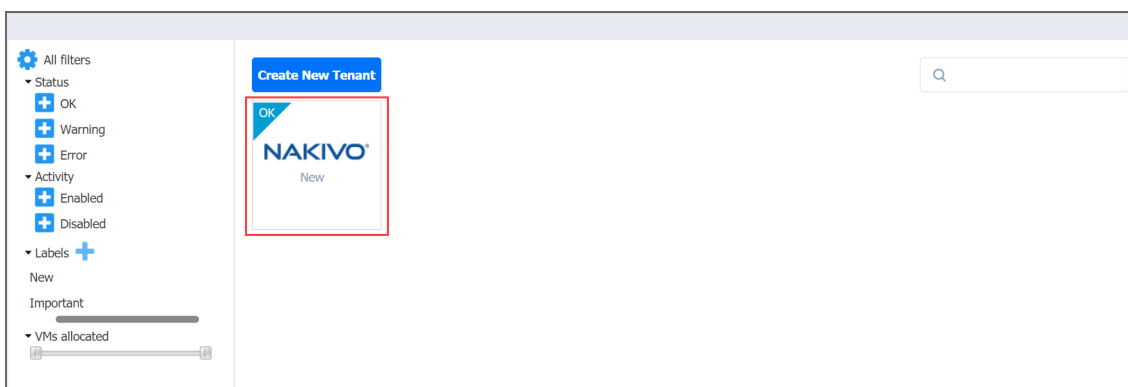
To view tenant information, hover over the tenant and click on the **Info** button.



The tenant information is displayed.

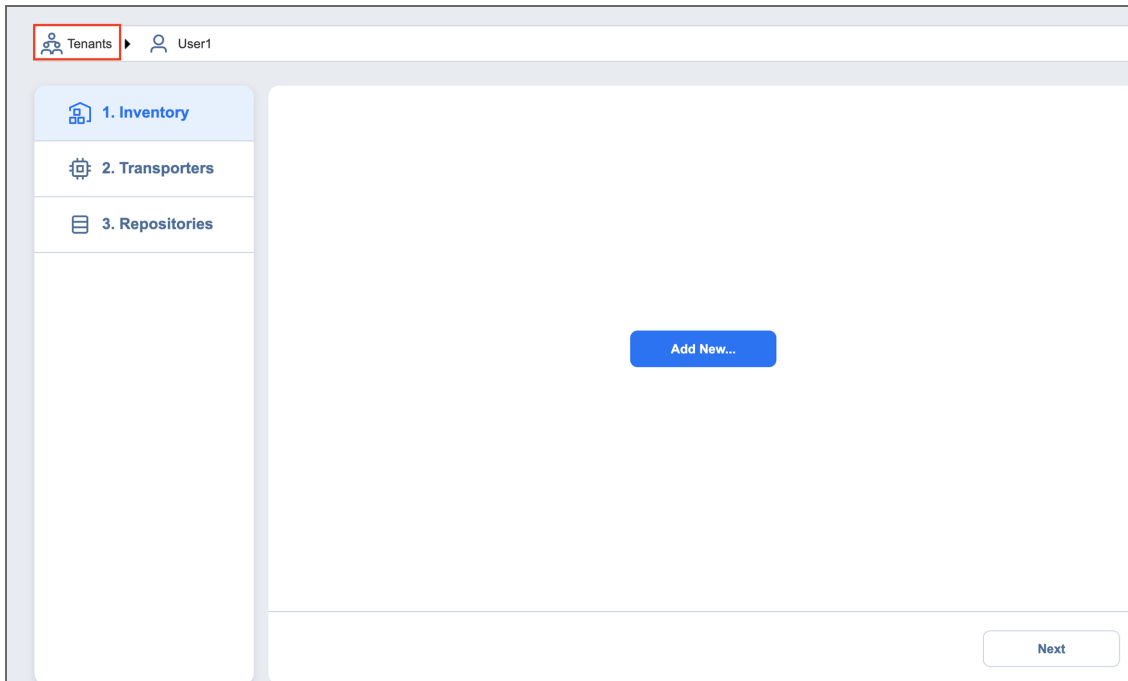
## Opening Tenant Dashboard

In the multi-tenant mode, you need to open the tenant dashboard to perform tenant configuration, create jobs and groups for the tenant, and recover files and emails. To open a tenant dashboard, simply click the tenant.



## Returning to Master Admin Dashboard

To return to the **Master Admin** dashboard, click **Tenants** in the navigation bar.

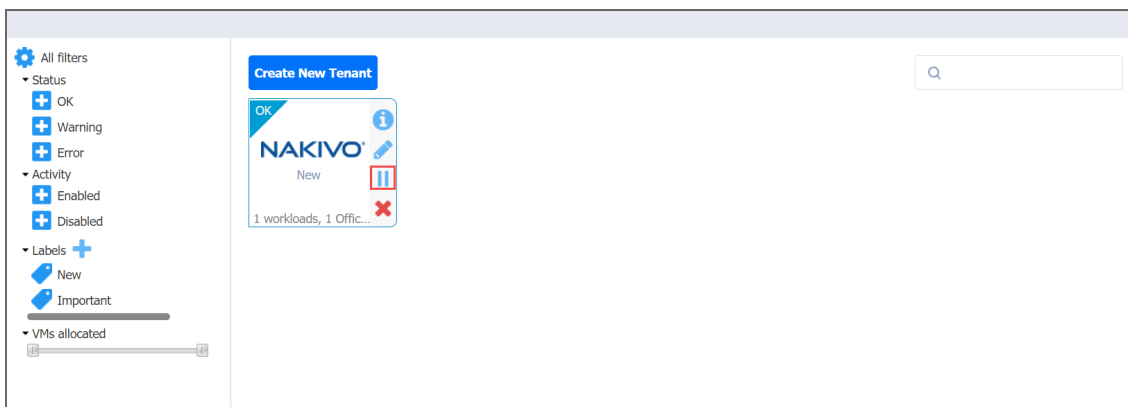


## Disabling Tenants

In multi-tenant mode, you 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.

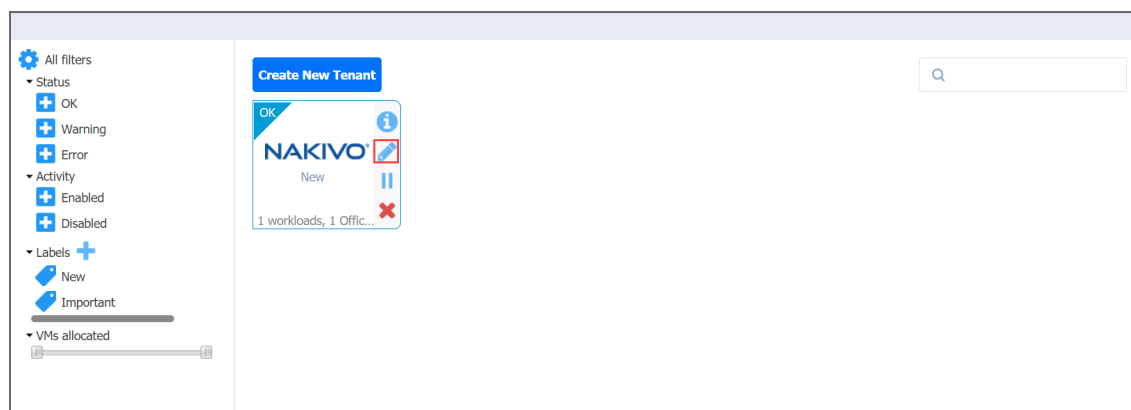
To disable a tenant, hover over the tenant and click the **Disable** button.



# Editing Tenants

To edit a tenant, do the following:

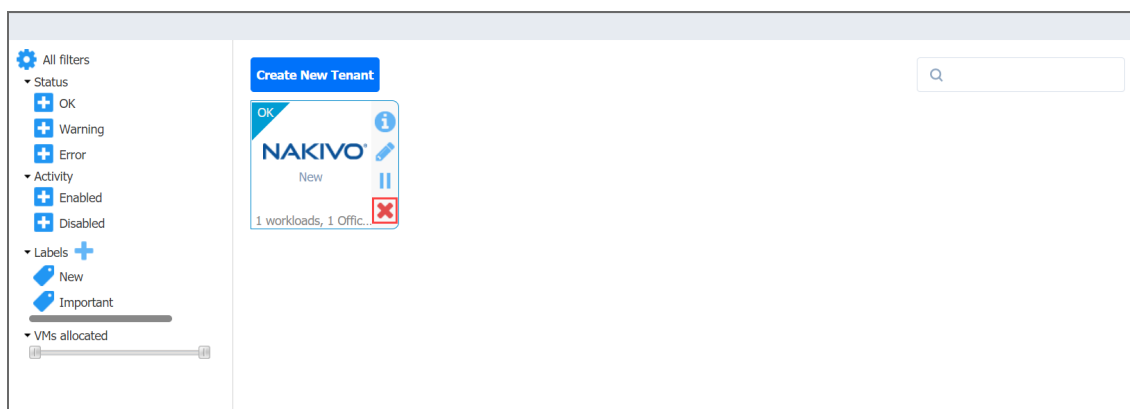
1. Hover over a tenant box and click the **Edit** icon.



2. In the **Edit** dialog that opens, make the required changes and click **Save**.

# Deleting Tenants

To permanently delete a tenant from the product, hover over a tenant and click the **Delete** icon.



The tenant will be permanently deleted from NAKIVO Backup & Replication.

Tenant Transporters are not uninstalled and the Tenant Backup Repositories are not removed.

# Granting Self-Service Access

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. The tenant admin has full control over the product features inside the tenant dashboard (such as edit and update tenant inventory, Transporters, and Backup Repositories, and create and manage jobs and groups). For each tenant, one guest account can also be created. The tenant guest has limited permissions inside the tenant and can only generate job and group reports by default. To provide a 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