NAKIVO Backup & Replication v10.4

User Guide for Hyper-V

Table of Contents

NAKIVO Backup & Replication Overview **5**

Deployment Options 6

Data Protection 7

Data Recovery 22

Disaster Recovery 30

Backup Size Reduction 36

Reliability 40

Performance 48

Administration 63

Automation 67

Integration 71

BaaS 76

Licensing 81

Deployment **84**

Architecture 85

System Requirements 94

Installing NAKIVO Backup & Replication 129

Updating NAKIVO Backup & Replication 199

Uninstalling NAKIVO Backup & Replication 230

Getting Started **234**

Logging in to NAKIVO Backup & Replication 235

First Steps with NAKIVO Backup & Replication 241

Web Interface Components 244
Managing Jobs and Activities 249

Settings **279**

General 280
Inventory 329
Transporters 344
Backup Repositories 369
Tape 432
Virtual Appliance Configuration 459
Multi-Tenant Mode Configuration 463
Support Bundles 472
Built-in Support Chat 474
Replacing License 481
Expert Mode 482

Replication **496**

Creating Hyper-V Replication Jobs 497
Staging (Seeding) VM Replication 526

Backup **527**

Creating Hyper-V Backup Jobs 528
Creating Backup Copy Jobs 552
Backing Up to Tape 574
Deleting Backups 591
Staging (Seeding) Initial Backup 596

Recovery **597**

Granular Recovery 598
Full Recovery 647
Planning Disaster Recovery 694

Integration and Automation 750

Command Line Interface 751
Automation with HTTP API 762
Aptare IT Analytics Integration 763

Multi-Tenant Mode 769

Tenant Creation 770
Tenant Configuration 777
Tenant Management 778
Granting Self-Service Access 789

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.

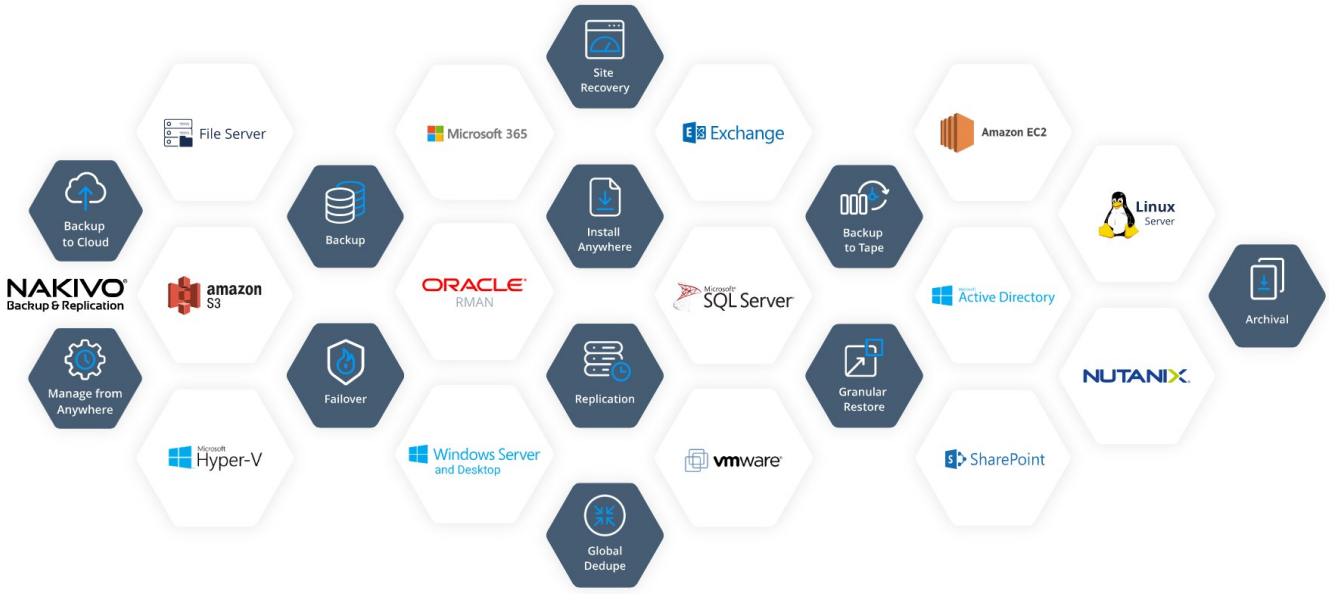

The product provides image-based, application-aware, incremental backup and replication. You can easily schedule jobs using the calendar in the product's web interface and save up to 1,000 recovery points for each backup, rotating them on a GFS basis. You can also protect your VMs and instances more 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.

Deployment Options

NAKIVO Backup & Replication is a versatile solution that can be installed on most modern operating systems and hardware solutions. For details, refer to the following topics:

- "Installing on Windows" on page 154
- "Installing on Linux" on page 165
- "Deploying Amazon Machine Image in Amazon EC2" on page 153
- "Installing on FreeNAS" on page 196
- "Installing on Synology NAS" on page 174
- "Installing on QNAP NAS" on page 181
- "Installing on Western Digital NAS" on page 191
- "Installing on ASUSTOR NAS" on page 186
- "Installing on NETGEAR ReadyNAS" on page 193
- "Installing on Generic ARM-based Device" on page 197
- "Installing on Raspberry Pi" on page 198

For the full list of supported systems and devices, refer to "Deployment Requirements" on page 98.

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:

- "Backup Copy" on page 10
- "Backup to Cloud" on page 15
- "Container Protection" on page 9
- "Backup to Tape" on page 19
- "Virtual Machine Backup" on page 8

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:

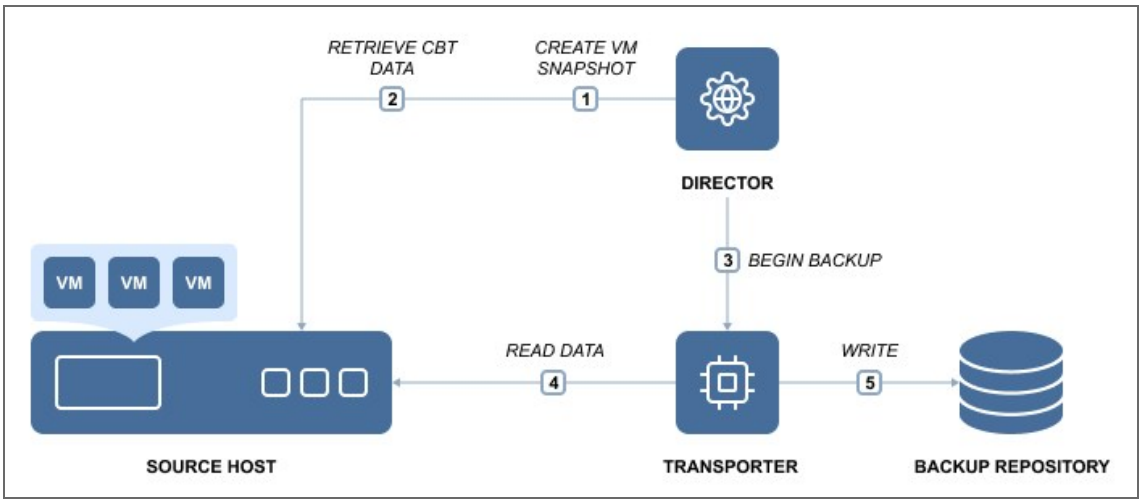

To learn how to create VM backups with NAKIVO Backup & Replication, refer to "Backup" on page 527.

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

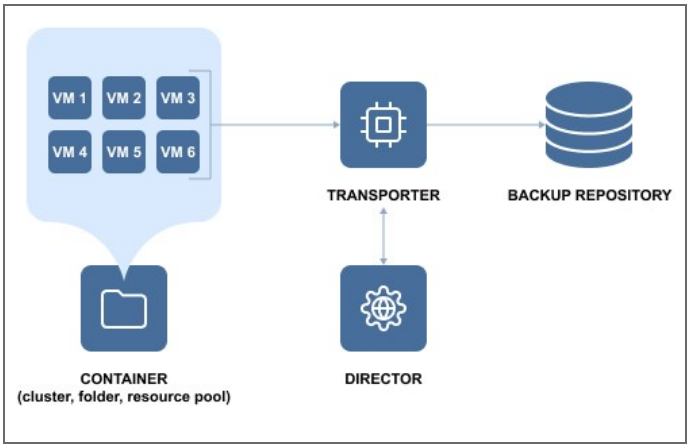

- Day 2

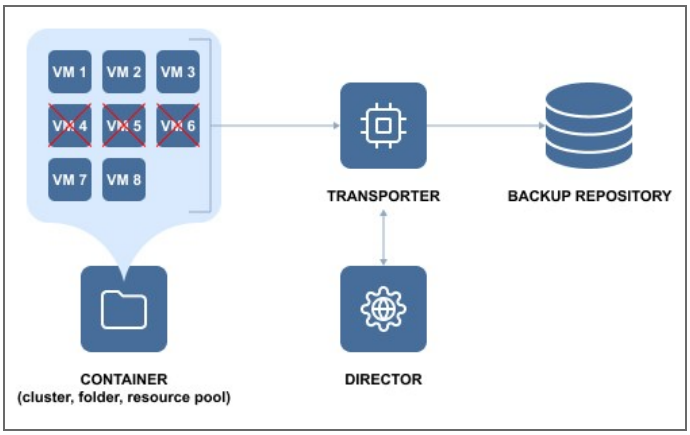

Backup Copy

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
- Save Storage Space with 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:

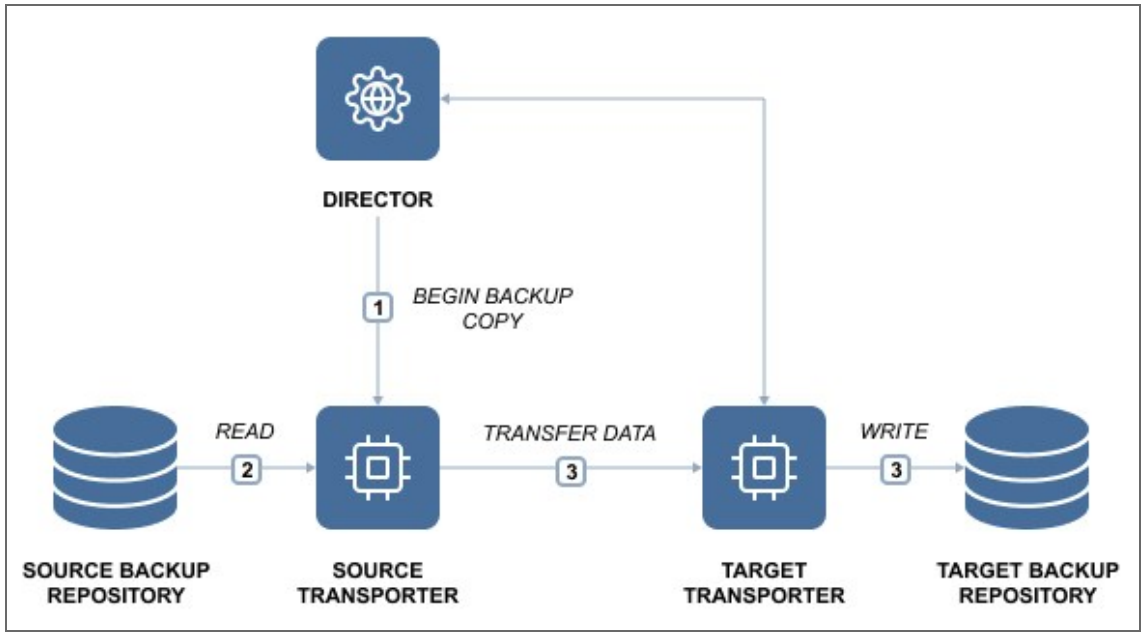

Copy Most Important Backups

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:

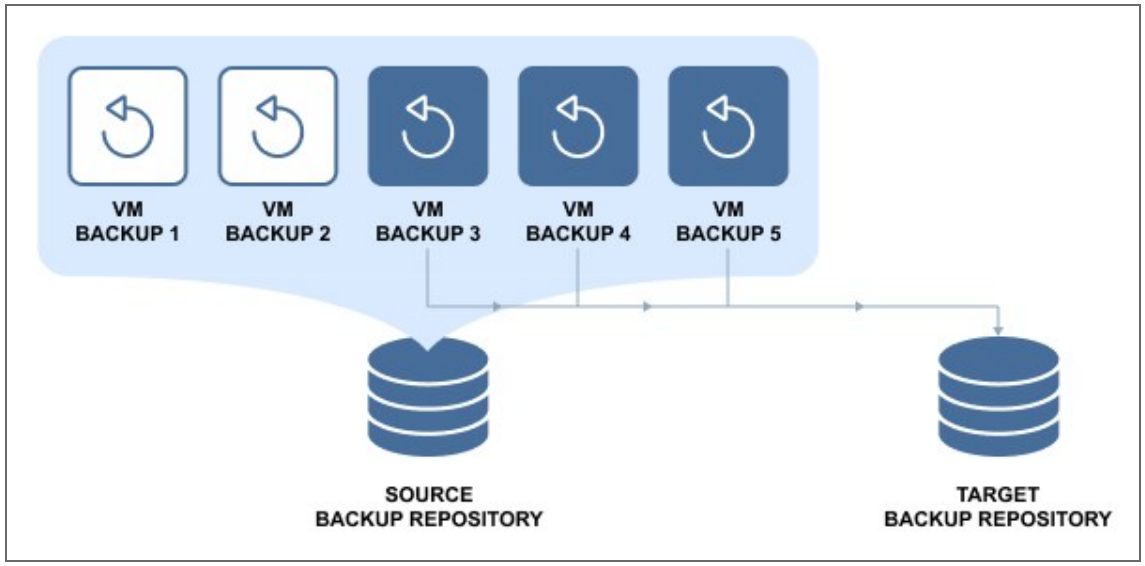

This way, only the selected backups (and their recovery points) will be transferred to the secondary Backup Repository.

Copy Backups Created by Particular Backup Jobs

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

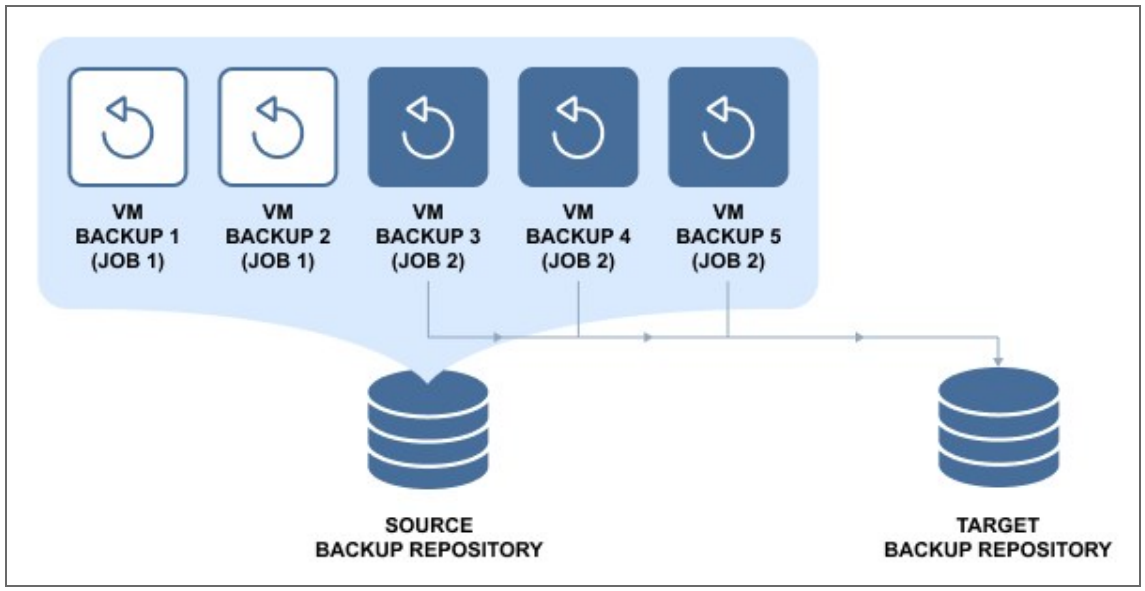

This way, you can ensure that all backups created by important Backup jobs are copied to a secondary Backup Repository.

Save Storage Space with 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.

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.

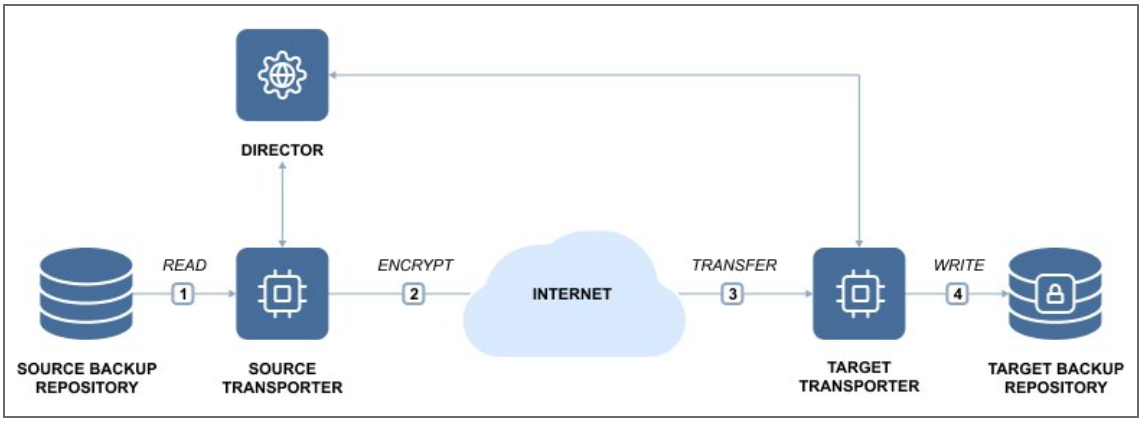

Copy Backups to Amazon Cloud

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.

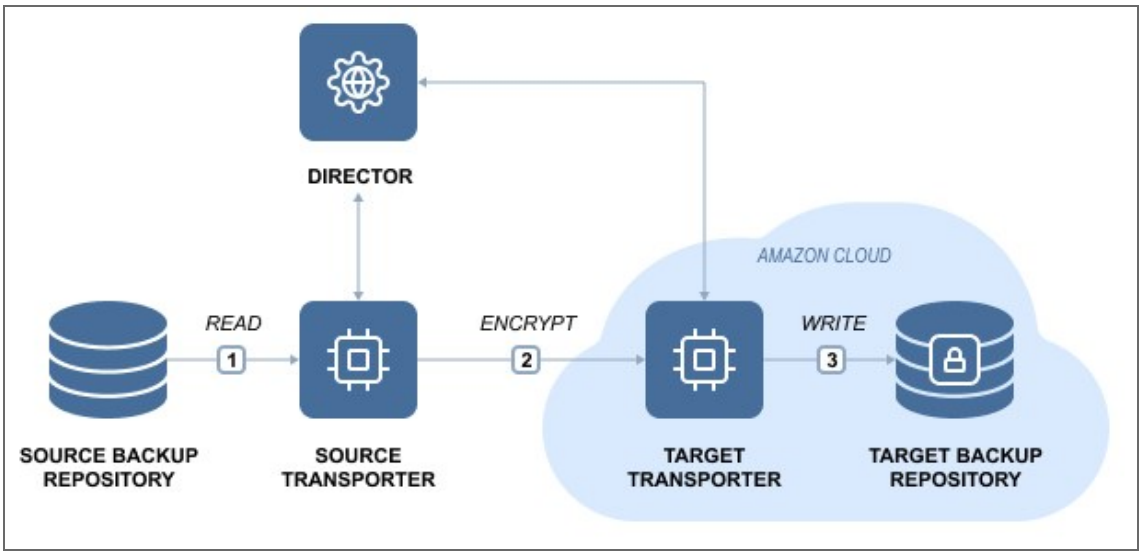

Copy Recovery Points that You Need

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.

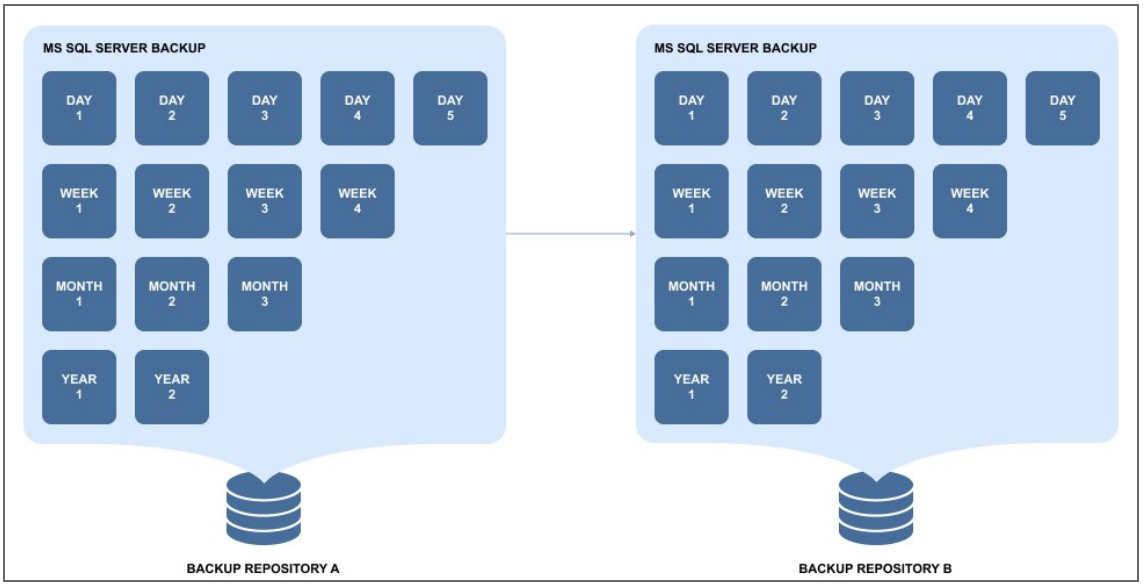

However, Backup and Backup Copy are different jobs, so you can set different retention policies for your primary backups and their copies in a different Backup Repository. This way, for example, you can store several daily backups onsite, and keep (archive) weekly, monthly, and yearly copies of backups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

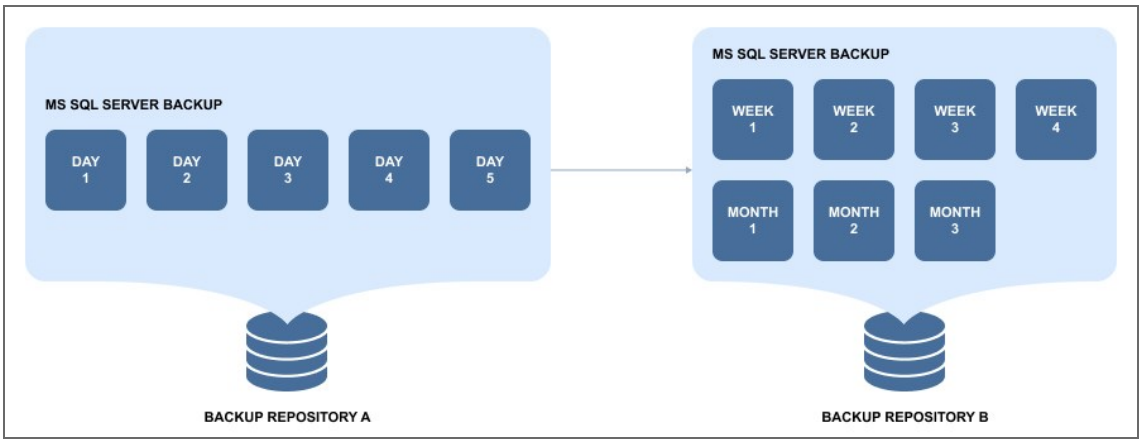

Also, you can use fast storage for a subset of backups and use slower, but more reliable storage for long-term archiving.

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 552.

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.

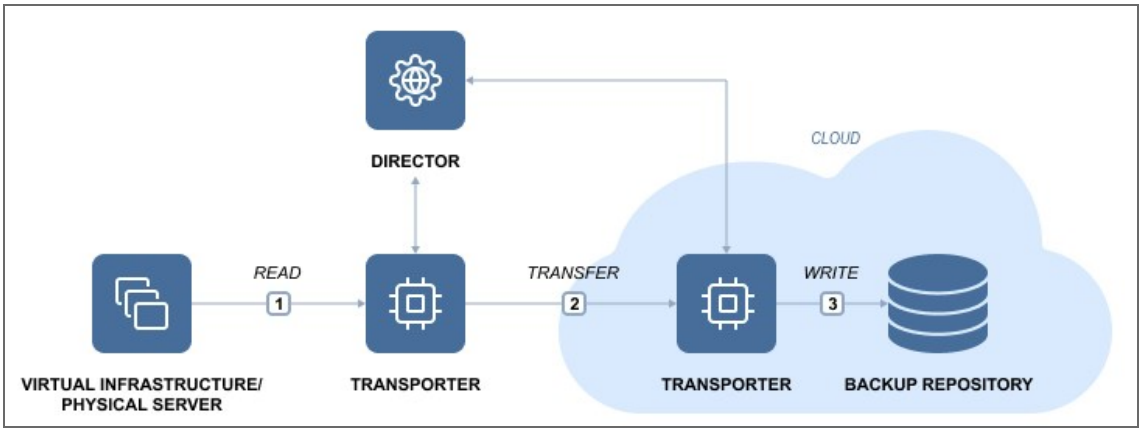

Keeping backups in the cloud provides a number of benefits, including:

- 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, skip 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:

- "Backup Repository in Amazon EC2" on page 389
- "Backup Repository in Amazon S3" on page 394
- "Backup Repository in Wasabi Hot Cloud Storage" on page 398

Amazon EC2 Concepts

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

Instance

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

EBS Volume

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

Region

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

Availability Zone

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

VPC

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

Subnet

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

Security Group

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

Key Pair

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

Elastic Network Adapter

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

Backup to Tape

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

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

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

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

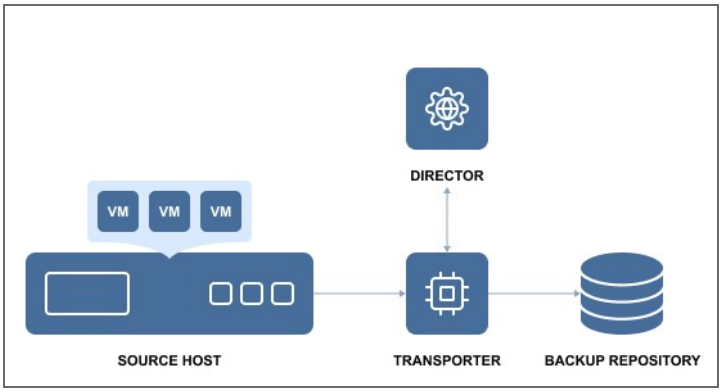

Stage 2 – copying backups from the repository to the tape library:

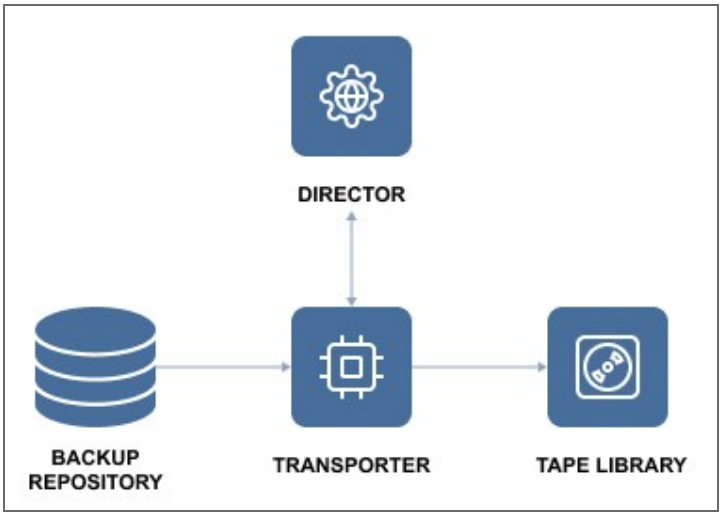

Recovering from tape is the reverse of backing up: the backups stored on the tape cartridges are first recovered to the Backup Repository and then recovered using NAKIVO Backup & Replication's tools.

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

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

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

Note

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

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

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

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

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

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

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" on page 23
- "Instant Object Recovery" on page 24
- "Instant VM Recovery - Flash VM Boot" on page 25
- "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.

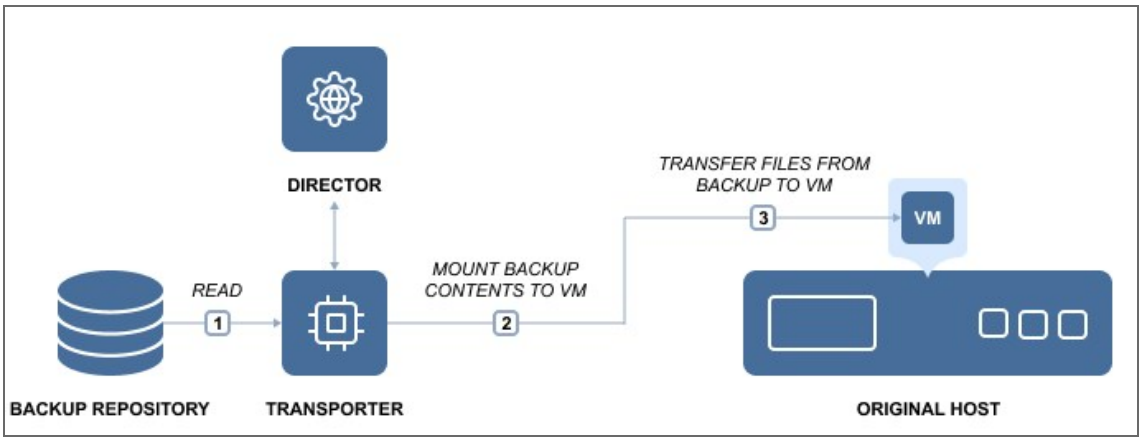

To learn how to recover files with NAKIVO Backup & Replication, refer to "File Recovery" on page 599.

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 598.

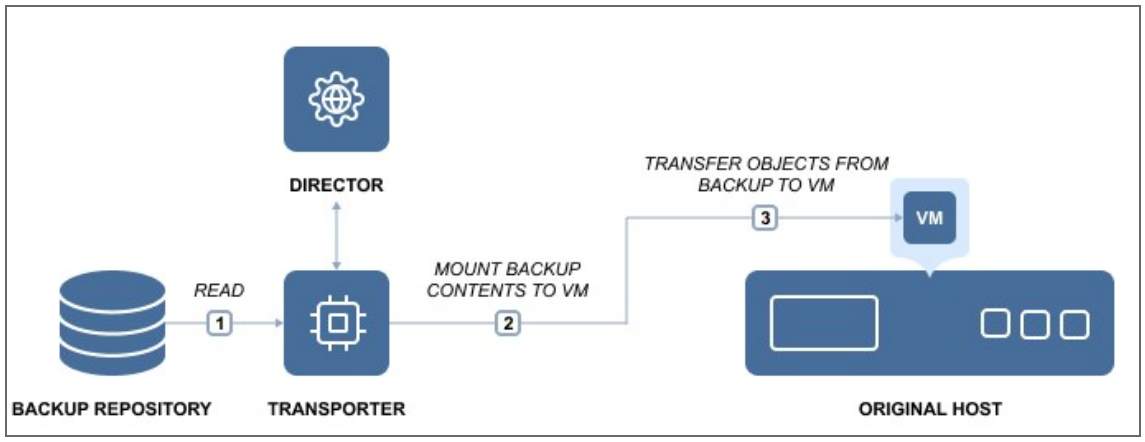

Instant VM Recovery - Flash VM Boot

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.

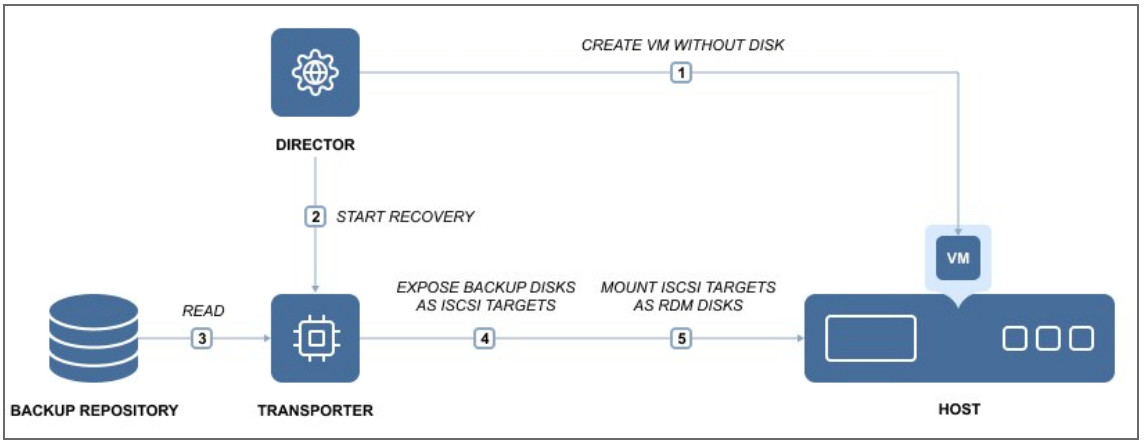

This process is fully automated and takes mere seconds to complete, after which the machine OS boot is started. Once booted, the machine can be migrated to the production environment using the hypervisor's native live migration feature.

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 648

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.

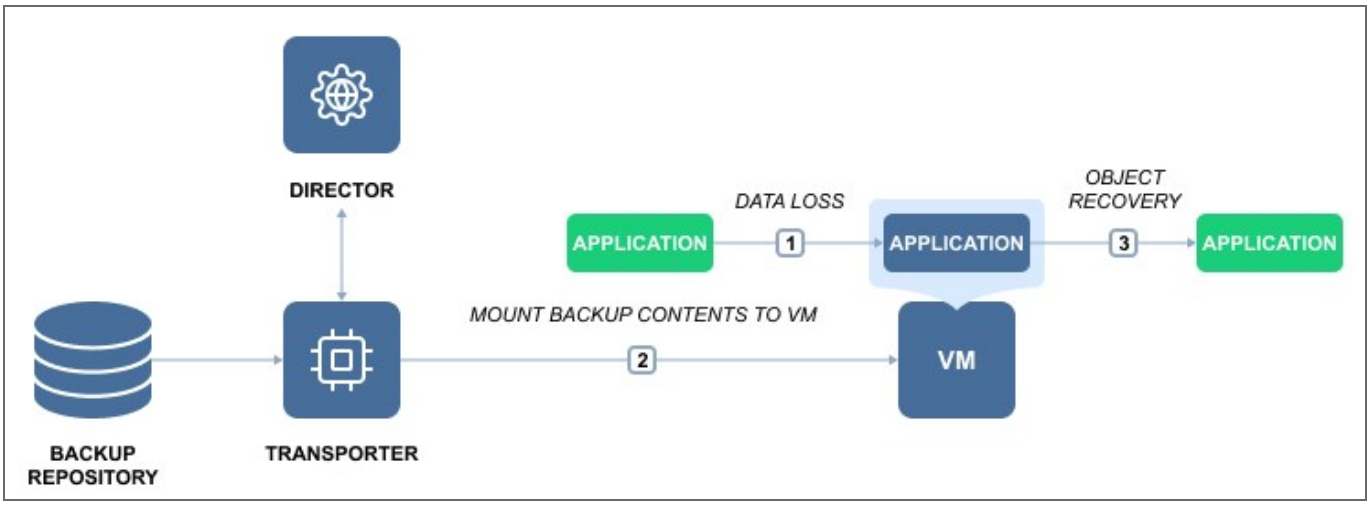

To learn how to create object recovery jobs with NAKIVO Backup & Replication, refer to the corresponding topics of the “Granular Recovery” on page 598 section.

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.

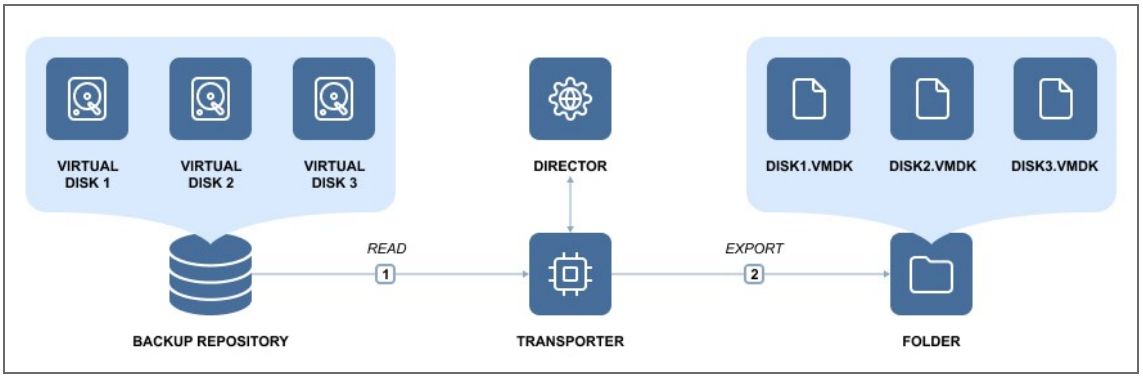

Once exported, the files can be used for recovery or long-term storage. Cross-Platform Recovery allows for unrestricted data protection across different hypervisors, physical machines and cloud platforms. Whether one of your hypervisors or physical machines is down or you need to migrate data from one 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" on page 31
- "Site Recovery" on page 32

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

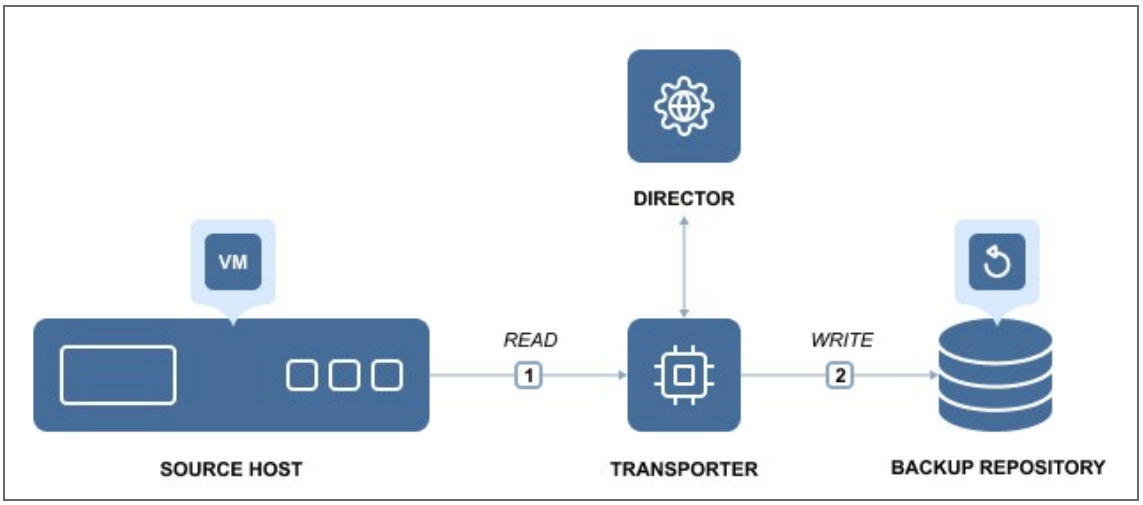

Step 2 - Replicate VM from backup

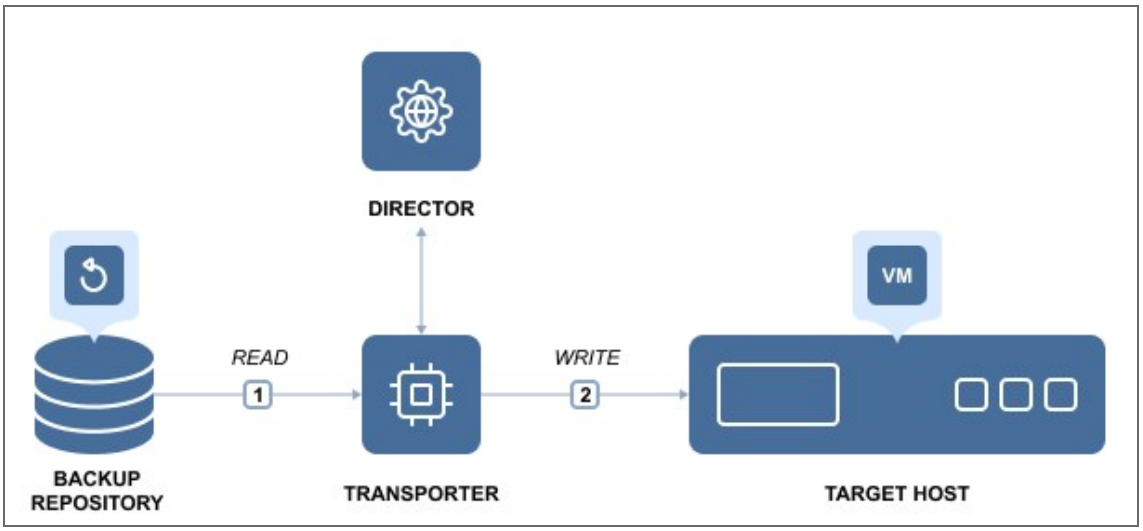

Setting up a replication from backup job for VMware and Hyper-V environments is no more time-consuming than setting up a traditional replication job. Once you launch a new replication job wizard and select VM backup as the source, NAKIVO Backup & Replication proceeds to read the data from the repository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

To learn how to create replication jobs with NAKIVO Backup & Replication, refer to “Replication” on page 496.

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 717 for the list of available actions.

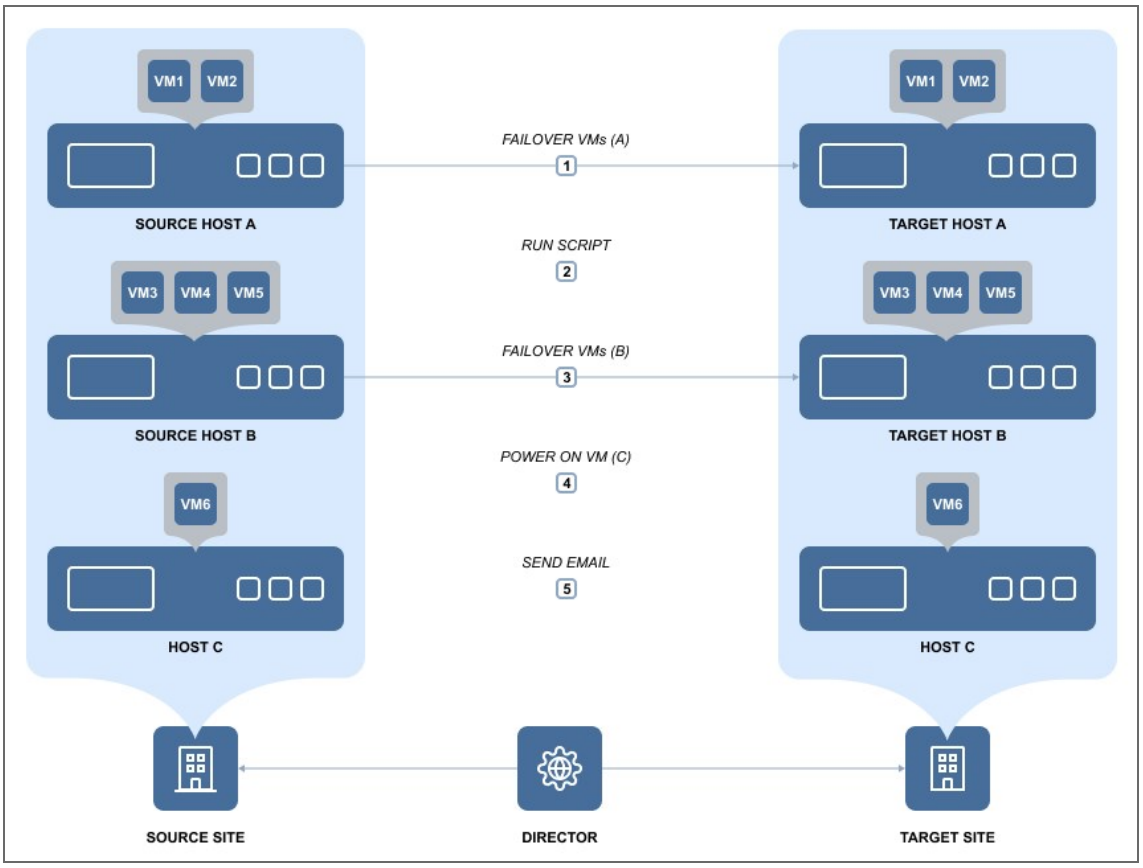

Special Actions that are used in recovering your IT environment with a Site Recovery Job are Failover and Failback:

- **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 746 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 747 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.

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" on page 37
- "Skipping Swap Files and Partitions" on page 38
- "Skipping Unused Blocks" on page 39

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.

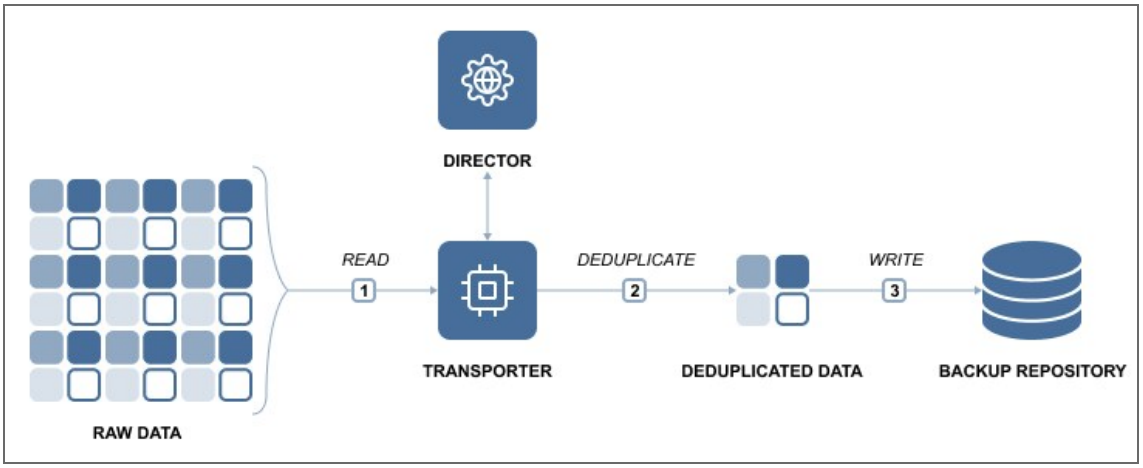

NAKIVO Backup & Replication automatically deduplicates all backups across an entire Backup Repository. This means that all data blocks are taken into account by backup deduplication, even if you back up your VMware VMs, Hyper-V VMs, and Amazon EC2 instances to the same backup repository. While backup deduplication is turned on by default, you can disable it if, for example, you want to use a hardware-based data deduplication device such as an EMC Data Domain.

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).

Skipping 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 skips 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.

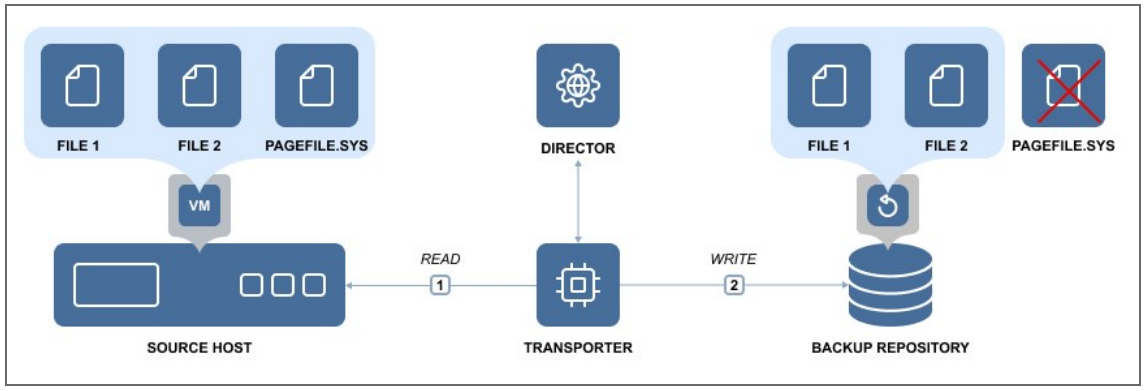

Skipping Unused Blocks

In addition to skipping swap files and partitions, NAKIVO Backup & Replication allows you to skip 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. Skipping 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.

Reliability

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

This section contains the following topics:

- "Application and Database Support" on page 41
- "Encryption in Flight and at Rest" on page 42
- Backup Immutability
- Direct Connect
- "Log Truncation" on page 44
- "Self-Backup Feature" on page 46
- "VM Verification" on page 47

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.

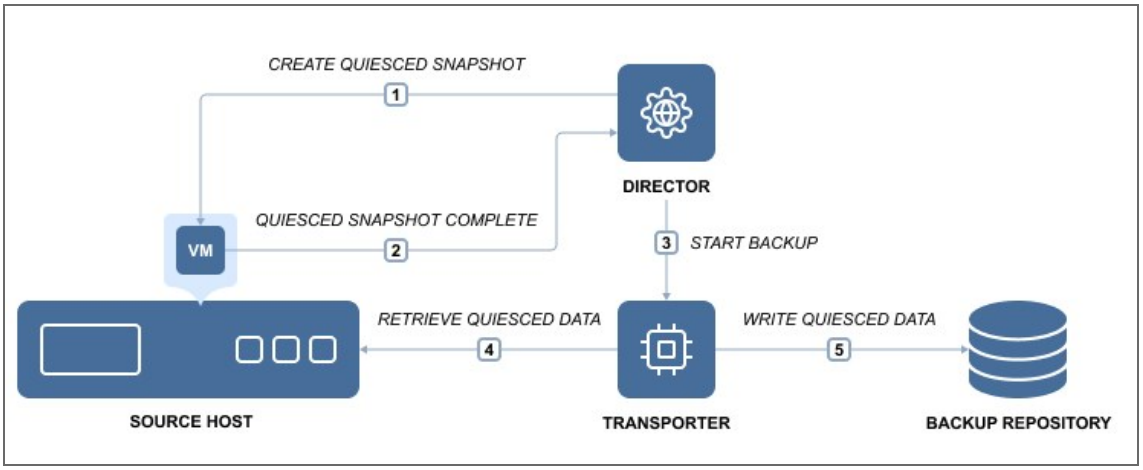

Encryption in Flight and at Rest

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).

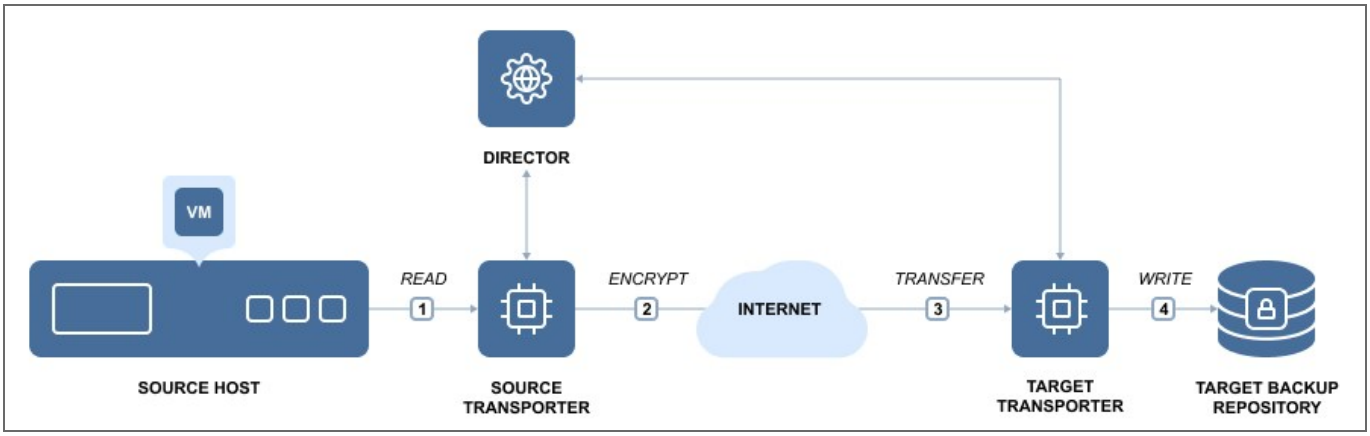

NAKIVO Backup & Replication uses AES 256 encryption to protect VM backups, which is the de facto worldwide encryption standard that secures online information and transactions by financial institutions, banks, and e-commerce sites.

- 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 374
- “Backup Repository on CIFS Share” on page 384
- “Backup Repository on NFS Share” on page 379
- “Backup Repository in Amazon EC2” on page 389
- “Backup Repository on Deduplication Appliance” on page 401

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 2016, 2013, 2010, and 2007 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.

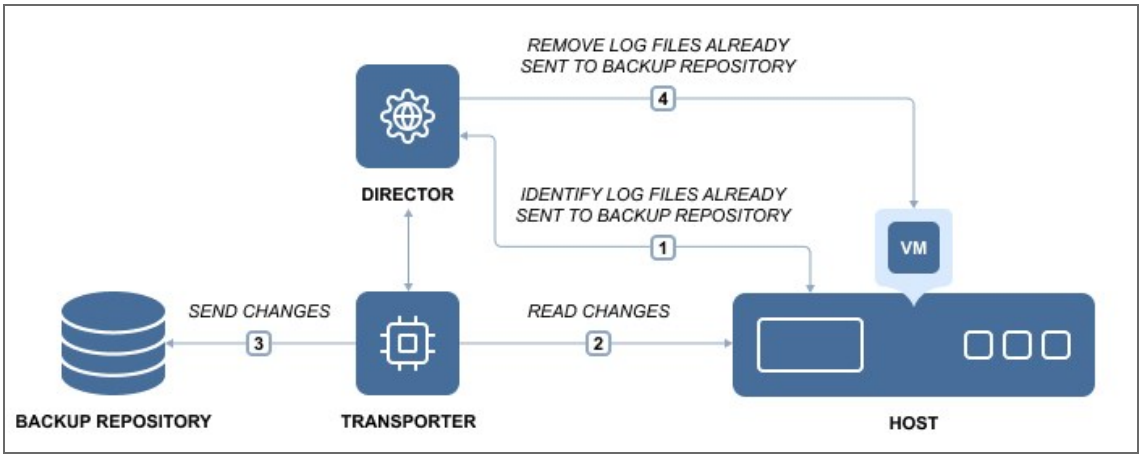

As a result, NAKIVO Backup & Replication creates regular, application-consistent backups of your Microsoft Exchange server and also removes the transaction log files so they don't consume all free disk space on the server.

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.

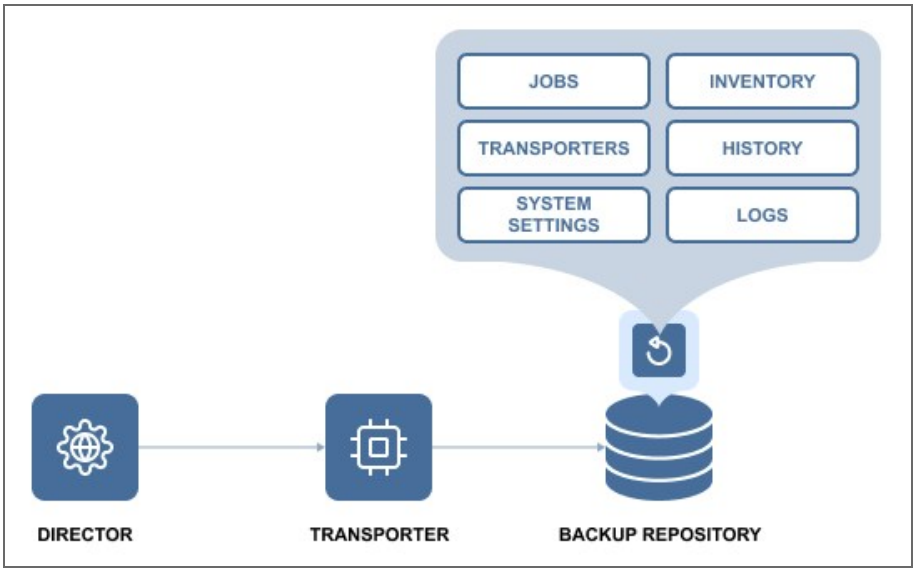

If you accidentally make some undesired changes in the product, you can easily roll back to a previous system state from the backup. Migrating the system configuration to a new product instance is simple: just install a new copy of NAKIVO Backup & Replication, import a Backup Repository that contains 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 Configuration” on page 291.

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:

- 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" on page 49
- "Deduplication Appliance Support" on page 56
- "Full Synthetic Data Storage" on page 58
- "Incremental Jobs" on page 60
- "Jobs and Concurrent Tasks" on page 61
- "Network Acceleration" on page 62

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.

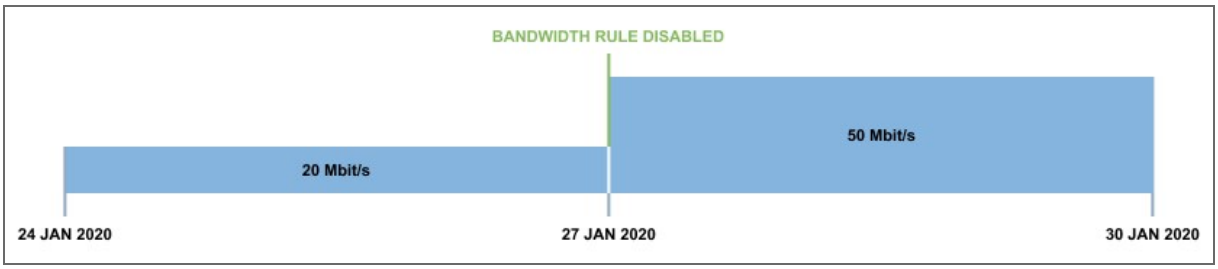

Bandwidth rules may be always active, active on schedule, or disabled. Refer to "Bandwidth Throttling" on page 281 for details.

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
- 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 3rd 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.

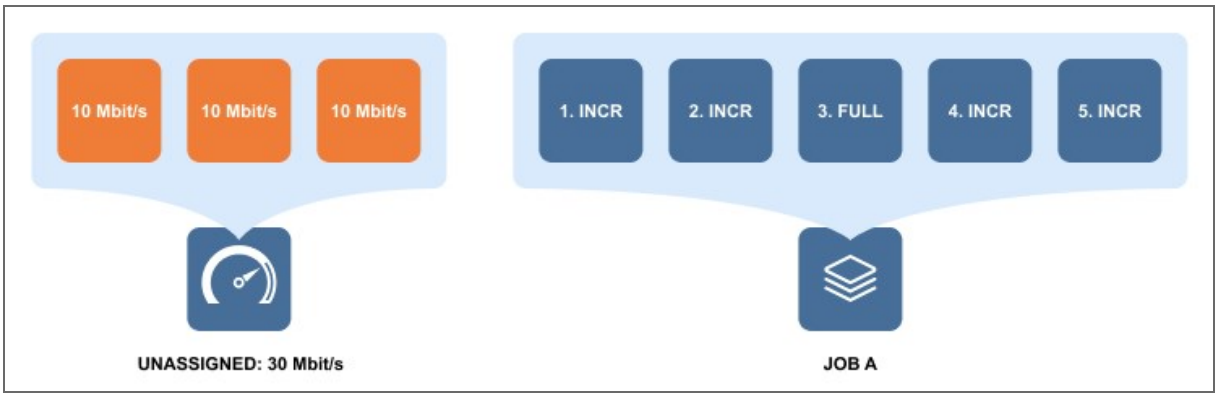

2. VM 1 and VM 2 backups receive 10 Mbit/s each. One bandwidth chunk remains unassigned since the full backup usually requires all the bandwidth to start.
3. The remaining bandwidth is distributed from the start of the queue, so VM 1 backup receives additional 10 Mbit/s.
4. VM 1 backup and VM 2 backup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

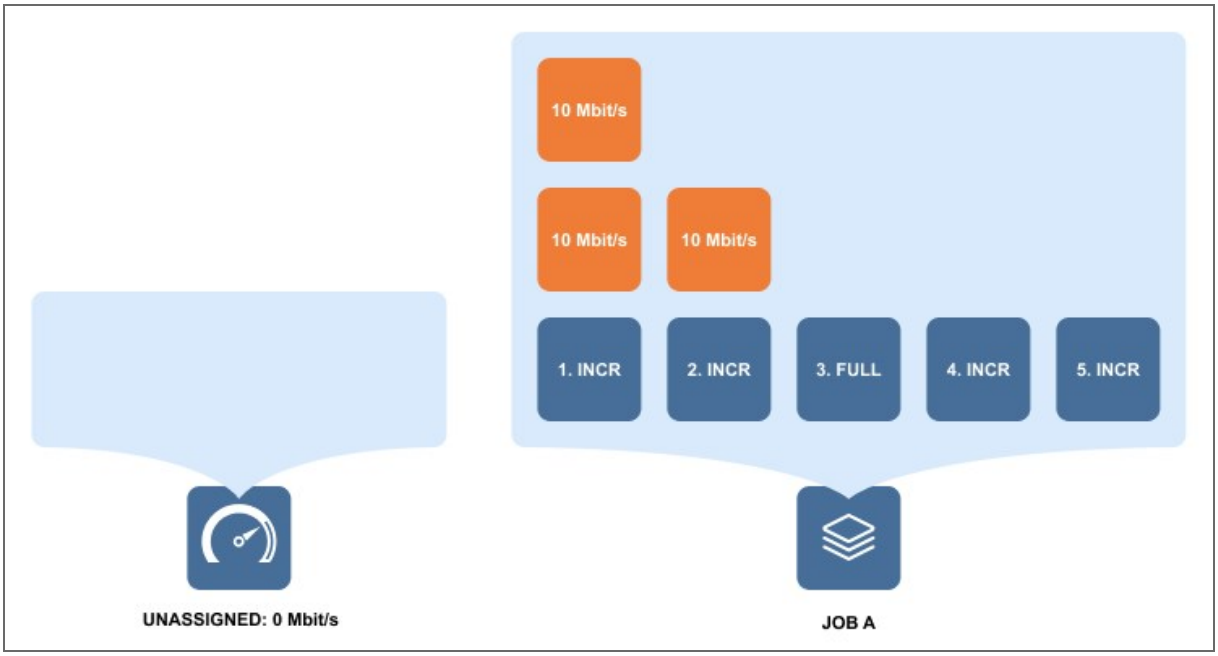

Note

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.

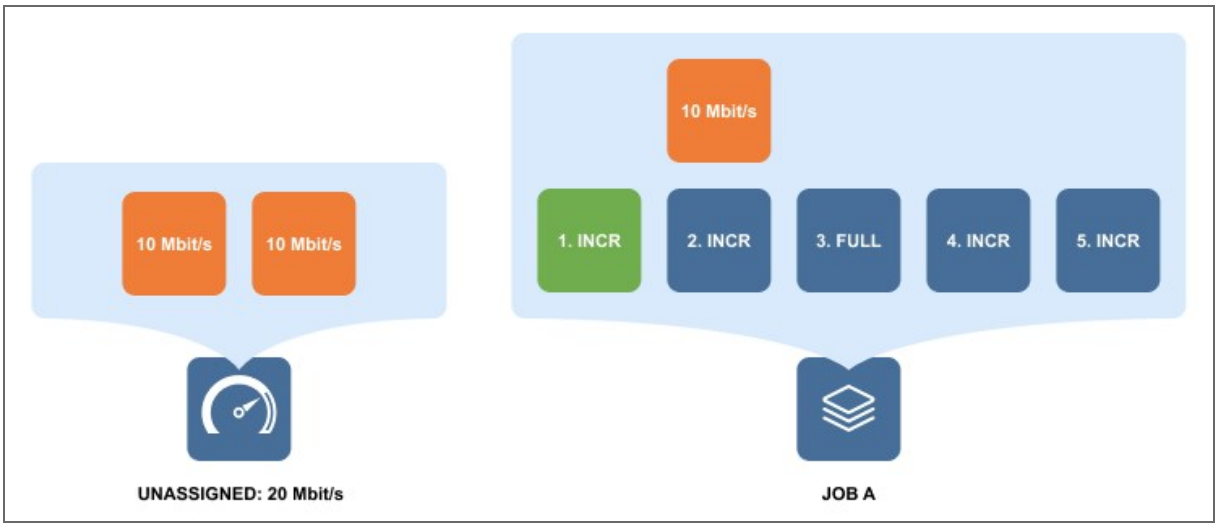

6. When VM 2 backup finishes running, it frees another bandwidth chunk, and full backup of VM 3 starts running with all the bandwidth assigned.

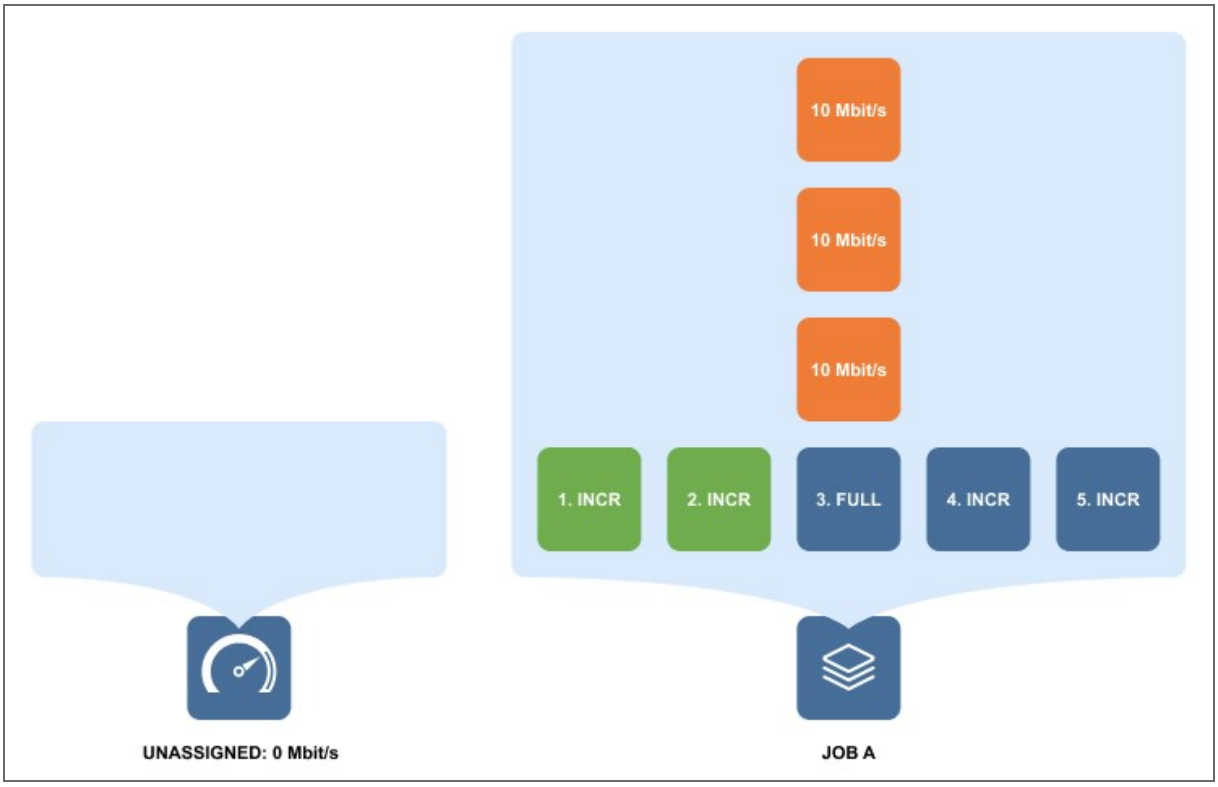

7. When full backup of VM 3 is finished, three bandwidth chunks are now available for the two remaining VM backups.

8. VM 4 backup receives the 20 Mbit/s bandwidth in total and VM 5 backup receives a 10 Mbit/s bandwidth chunk.

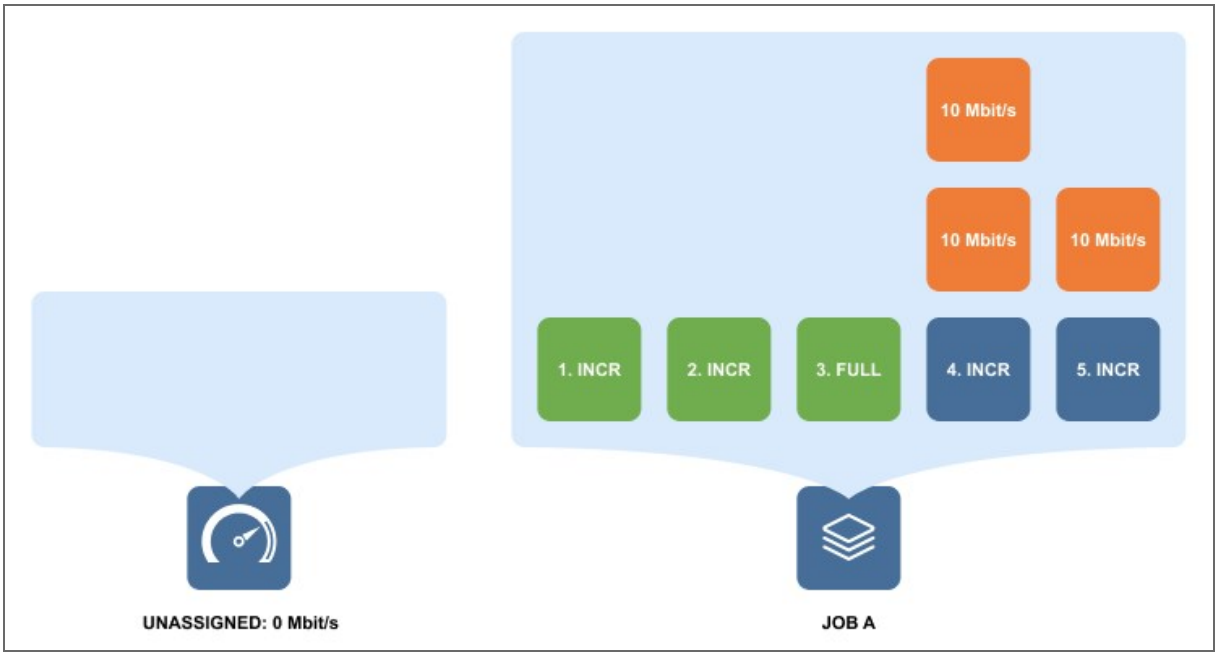

When the rule changes the bandwidth to 80 Mbit/s and is also activated for another Job B consisting of two VM incremental backups, the Transporter starts distributing bandwidth as follows:

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

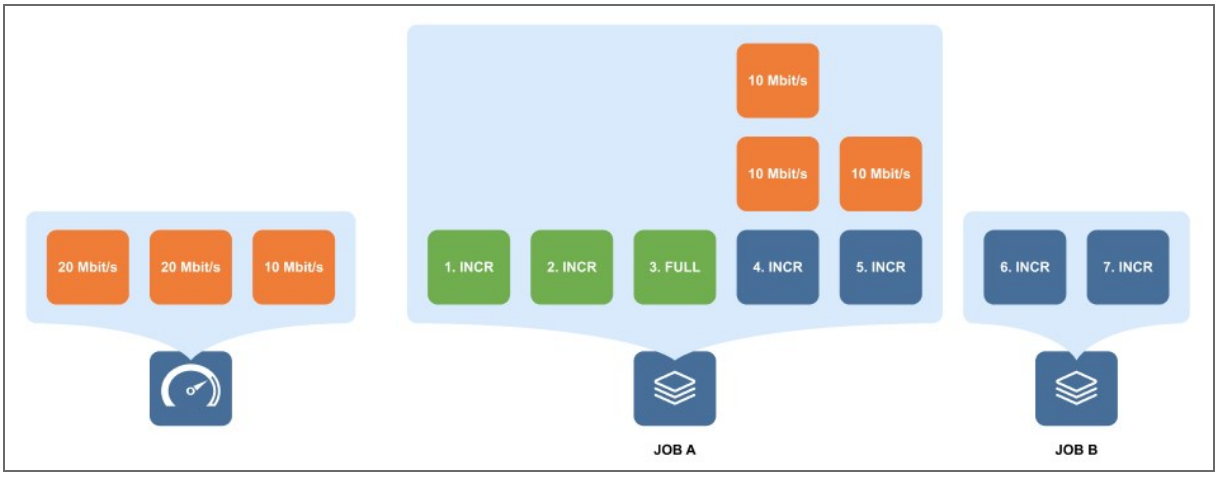

2. VM 6 backup and VM 7 backup of Job B receive a 20 Mbit/s bandwidth chunk each and start running, with 10 Mbit/s remaining unassigned.

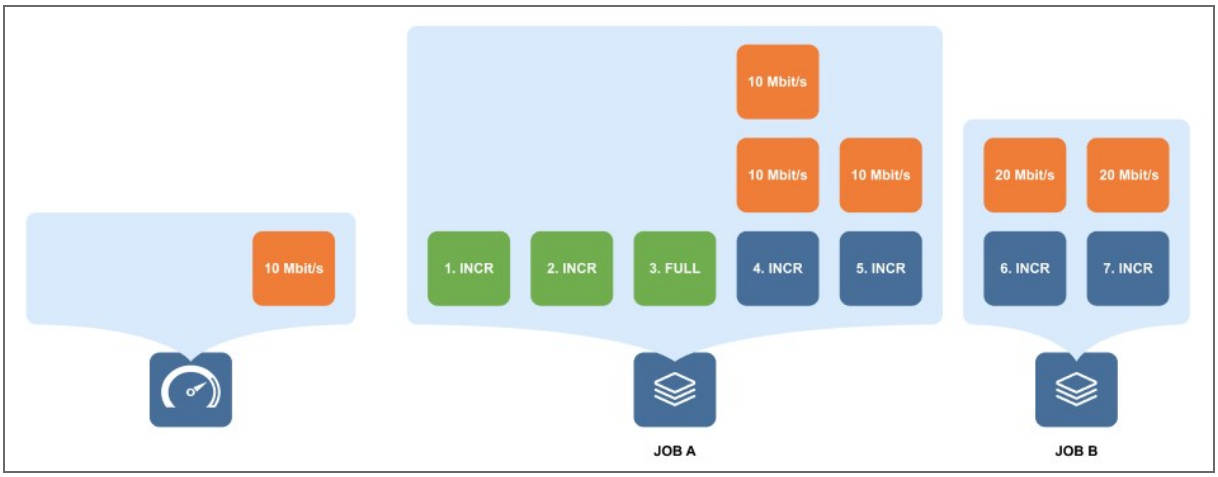

3. When VM 4 backup and VM 5 backup of Job A are finished, two 20 Mbit/s bandwidth chunks are freed. However there are no queued tasks to assign them to, so the bandwidth is left idle.

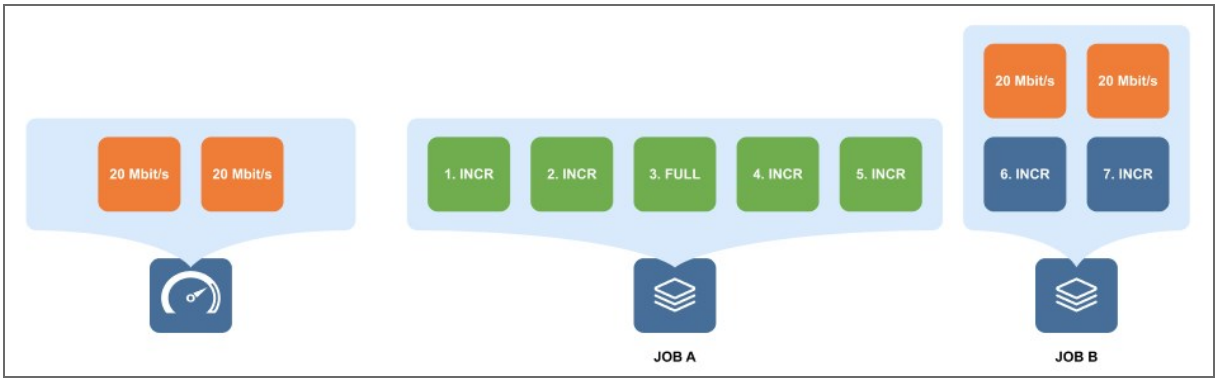

When the bandwidth rule changes the bandwidth amount back to 30 Mbit/s and is also activated for another Job C consisting of one VM incremental backup, the Transporter starts distributing bandwidth as follows:

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.

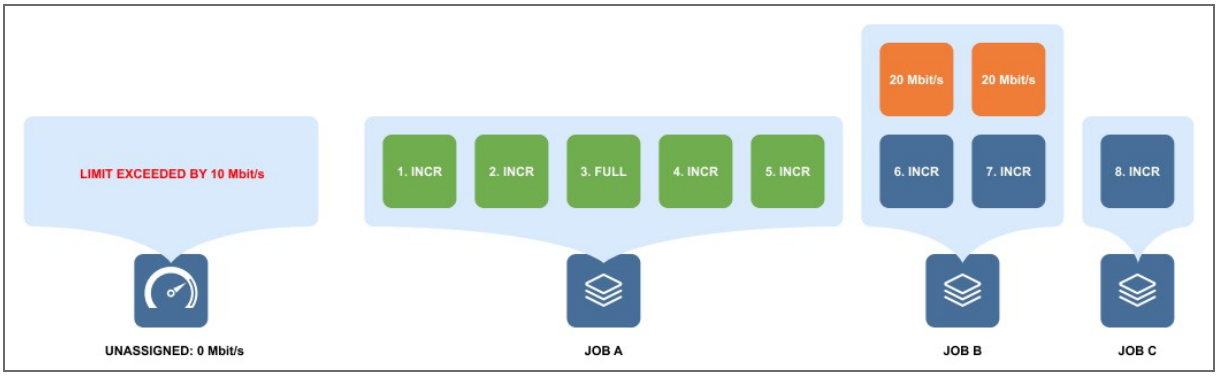

Note

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.

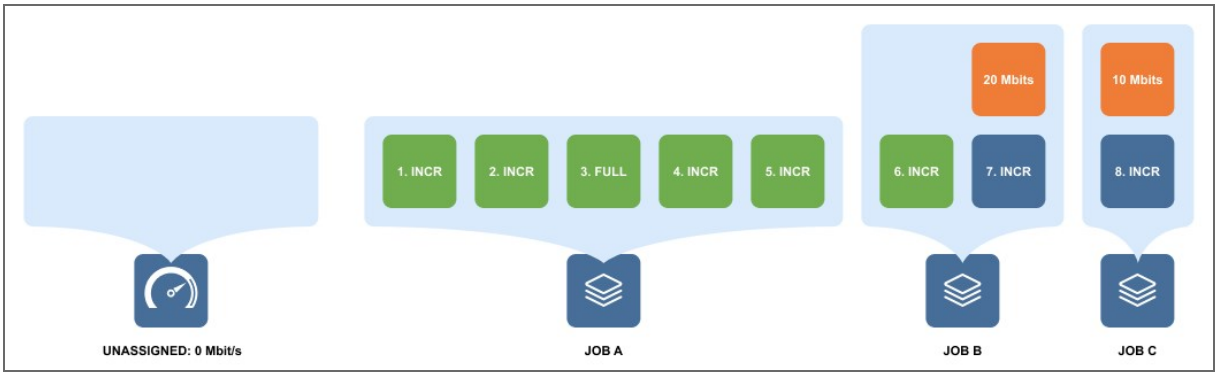

Deduplication Appliance Support

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.

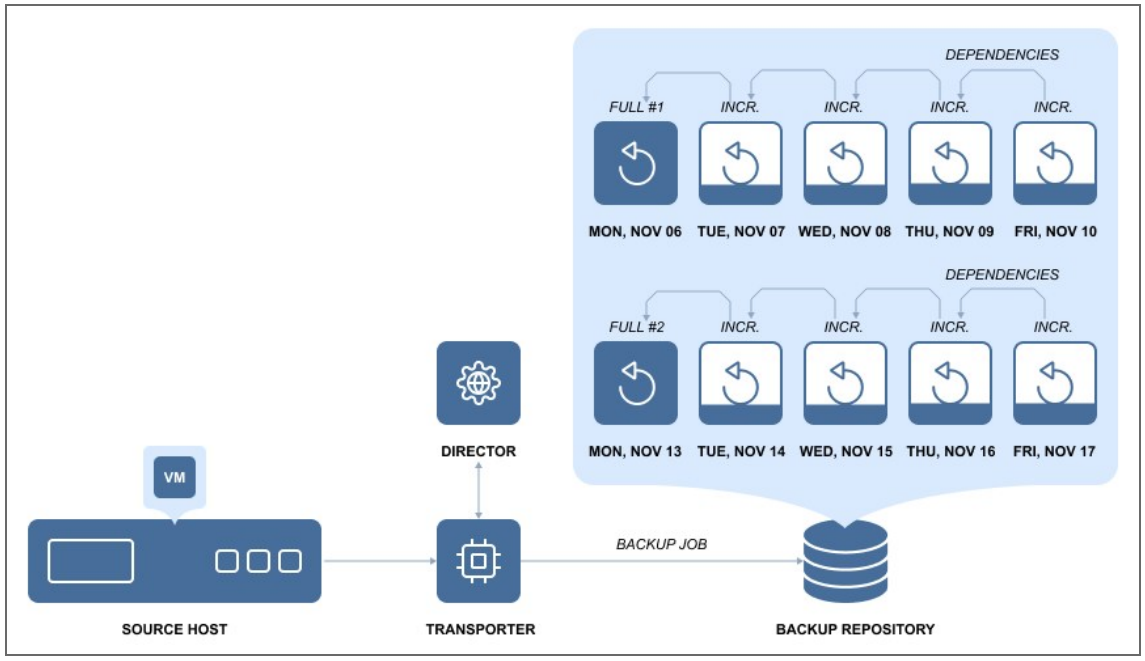

In terms of integration with deduplication appliances, a stream repository:

- 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** 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:

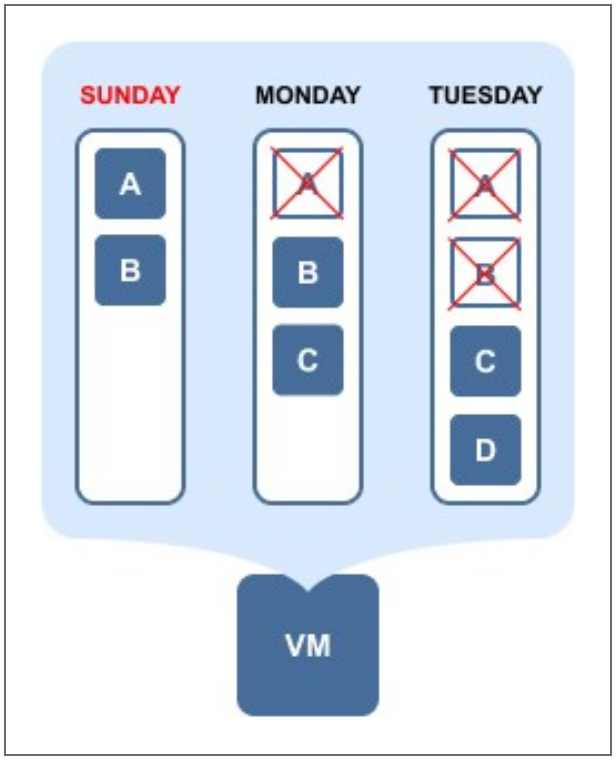

And here's how the data will be stored in the **forever incremental Backup Repository** if the job is set to keep 3 or more recovery points:

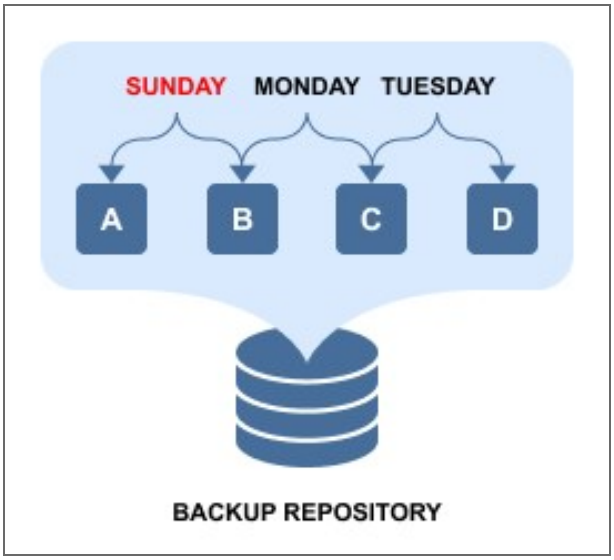

As you can see from above, each unique data block is stored only once to save space, while recovery points are just references to data blocks that are required to reconstruct the VM as of a particular moment in time. If, for example, you delete Monday's recovery point, then no actual data removal will 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** 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** 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.

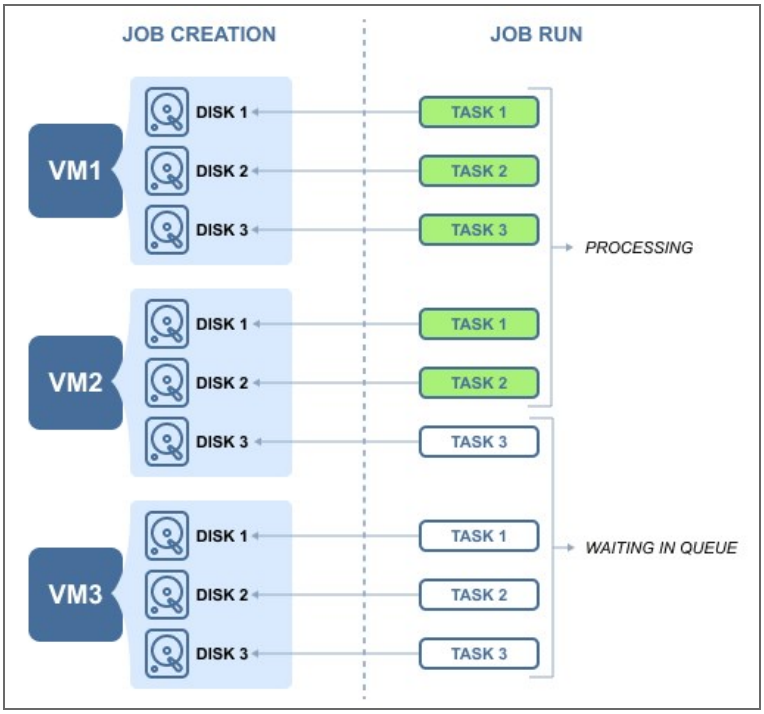

Each job object may consist of one or more machine disks, Oracle databases, Exchange Online mailboxes, OneDrive for Business instances or SharePoint Online sites that have to be processed within a job run. Data processing that is related to a specific VM disk or service constitutes a single task, in 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 364 to learn how to change the Transporter maximum load.

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

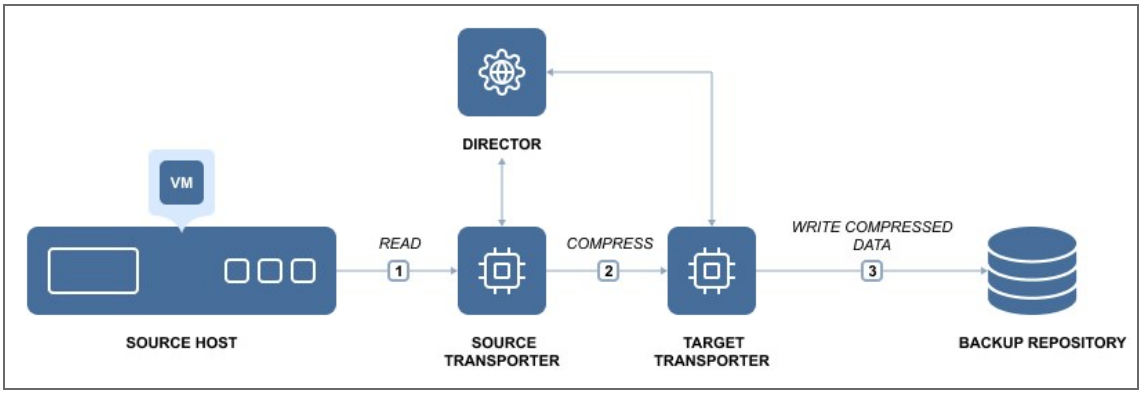

Network Acceleration for Replication

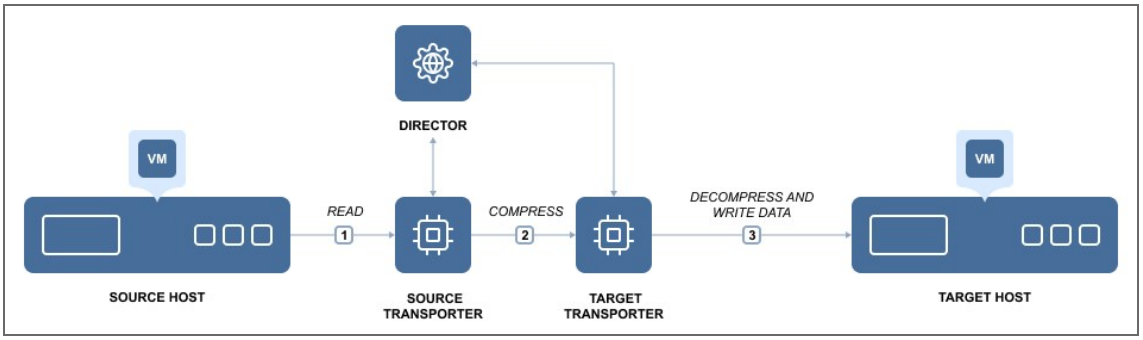

Administration

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” on page 64
- “Global Search” on page 65
- “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.

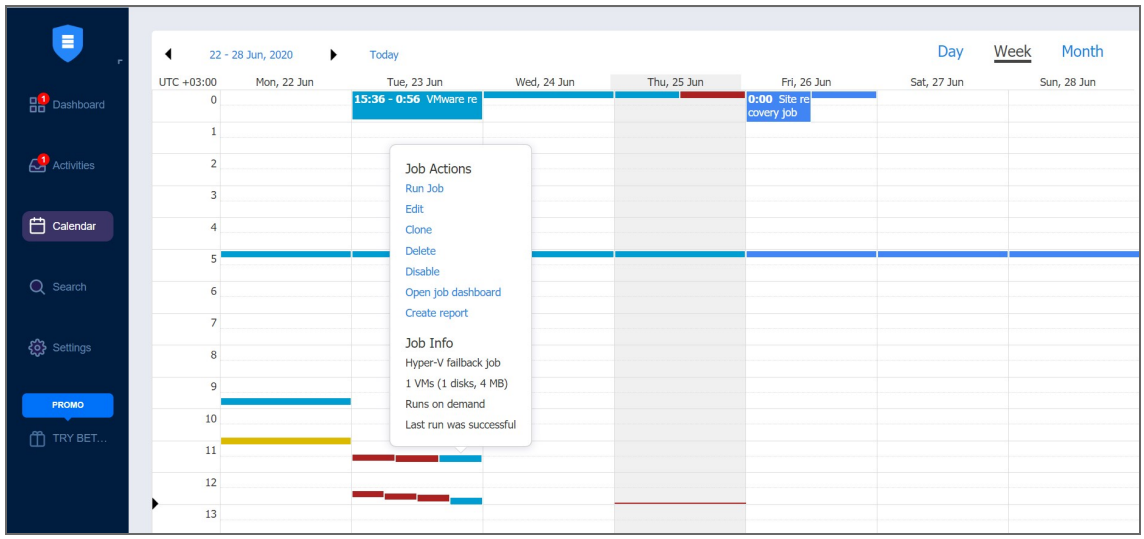

Global Search

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.

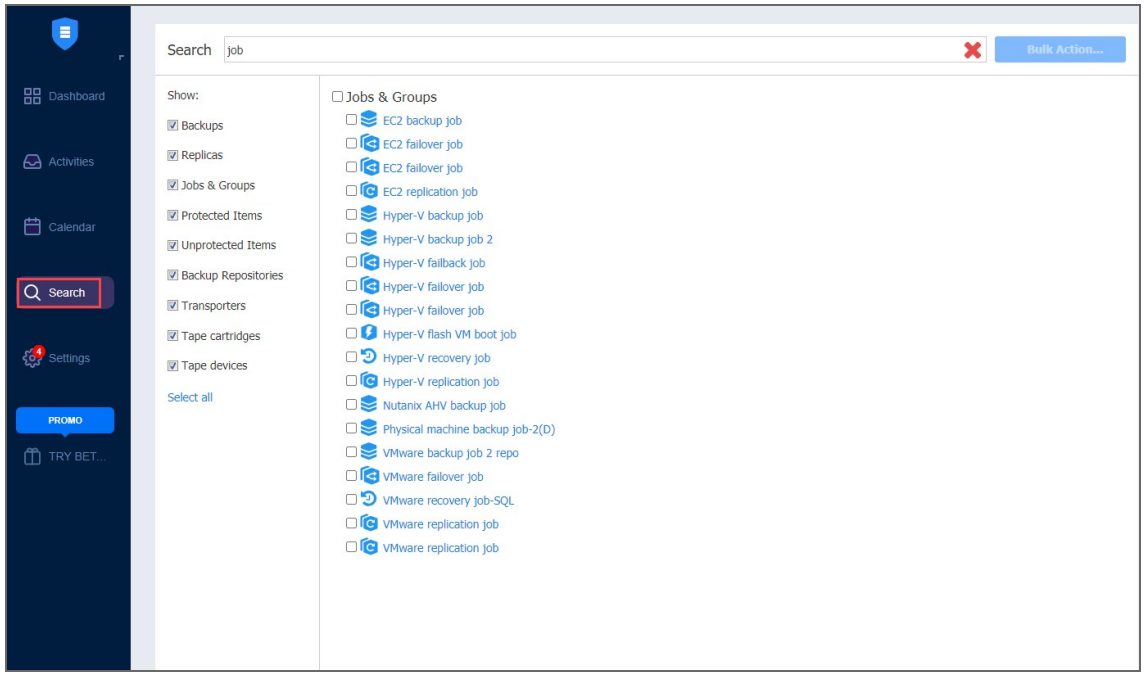

Policy-Based Data Protection

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 262
- “Managing Policy Rules” on page 265

Automation

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

- "HTTP APIs" on page 68
- "Job Chaining" on page 69
- "Pre and Post Job Scripts" on page 70

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.

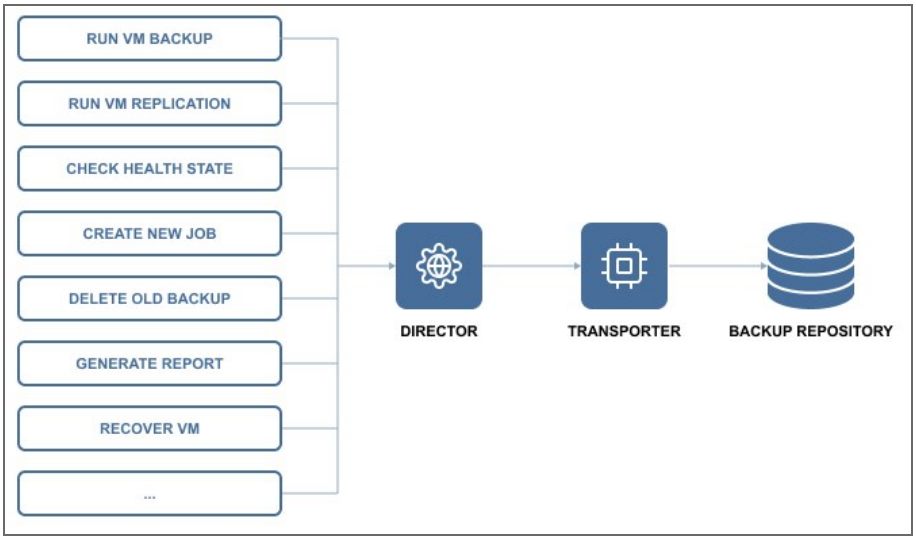

The API allows you to easily integrate NAKIVO Backup & Replication with monitoring, automation, and orchestration solutions to reduce time spent on backup management and reduce data protection costs. To speed up integration time, the API comes as part of an Integration Kit, which includes API documentation 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, which saves backups locally and then starts a Backup Copy job, which copies the newly created backups to Amazon cloud.

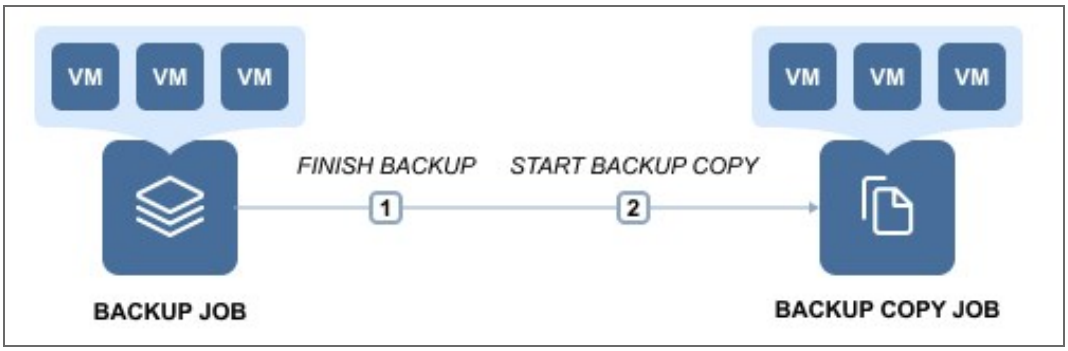

You can link any type of jobs together – backup, backup copy, replication and recovery – and add any number of jobs to the chain. For instance, you can set up a series of backup jobs that trigger one another in the order of priority, or set up a series of Backup Copy jobs, which first send weekly backups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

Pre and Post Job Scripts

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



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

Integration

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

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

- The regular Backup Repository, which is optimized for generic storage systems and performs forever-incremental 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”, and the name truly speaks for itself. 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 with its recovery points is stored in a separate folder for easier manageability.

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

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

- “Active Directory” on page 72
- “EMC DD Boost” on page 73
- “HPE StoreOnce Catalyst” on page 74
- “NEC HYDRAstor” on page 75

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 307
- "Managing Active Directory Users" on page 304

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 96
- Integrating with EMC DD Boost
- “Backup Repository on Deduplication Appliance” on page 401

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 96
- "Backup Repository on Deduplication Appliance" on page 401

NEC HYDRAstor

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

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

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

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

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

- “Storage Integration Requirements” on page 96
- Integrating with NEC HYDRAstor
- “Backup Repository on Deduplication Appliance” on page 401

BaaS

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

This section contains the following topics:

- "Branding" on page 77
- "License Delegation" on page 78
- "Multi-Tenancy" on page 79
- "Self-Service" on page 80

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.
- **Look and feel**: Bookmark icon and page background.

For information on branding configuration, refer to “Branding Configuration” on page 284.

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.

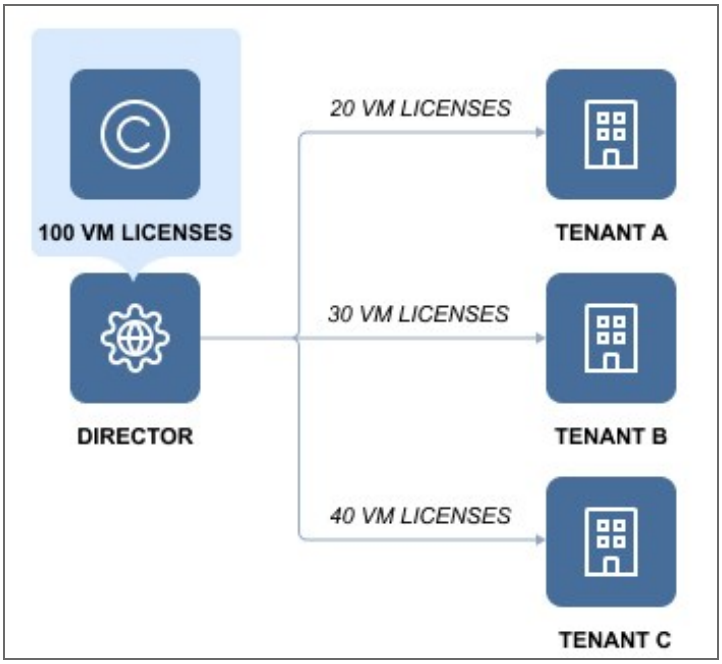

At any moment, the Master Admin can redistribute licenses: revoke any number of licenses from any tenant, which will return them to the Master License Pool, and add licenses to another tenant. The License Delegation feature makes license management simple and manageable in large and distributed environments.

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.

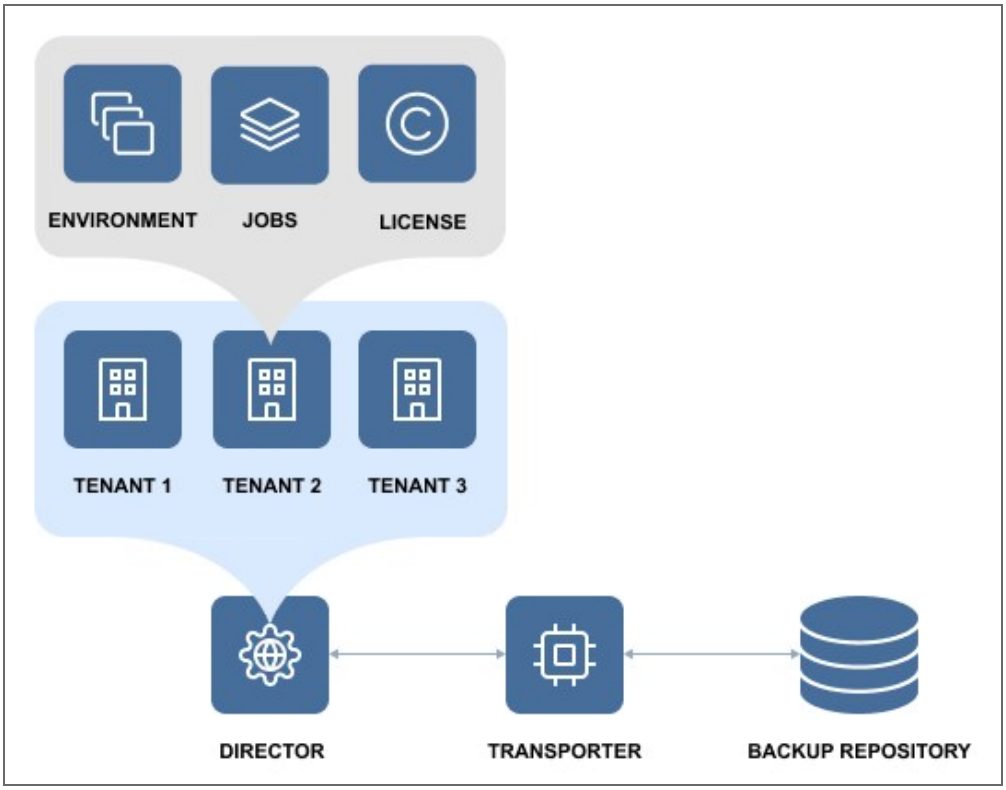

In Multi-tenant mode, each tenant can access their own environment through a self-service portal, and perform all data protection and recovery tasks. At the same time, tenants are isolated from each other and cannot access the environment and jobs of other tenants.

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

Licensing

Choose the licensing type that best suits your business needs. For details, refer to the topics below:

- Licensing Types
  - Perpetual Licensing
  - Per-Workload Licensing
  - Subscription Licensing for Microsoft 365
- Licensing Rules
- NAKIVO Support

Licensing Types

NAKIVO Backup & Replication offers the following licensing types:

Perpetual Licensing

For VMware, Hyper-V, and Nutanix AHV infrastructures, NAKIVO Backup & Replication can be licensed on a per-socket basis. A license is required for each socket on a host where you plan to back up or replicate VMs. Licenses are required only for the source side of backup and replication, that is, you do not need to license target servers for your replicas or the servers on which you want to recover VMs.

Perpetual licenses are also available for physical machines on a per server or workstation basis and Oracle databases (Enterprise Plus edition only) on a per Oracle database basis.

Notes

- Perpetual licenses for physical servers can't be applied for physical workstations and vice versa.
- Physical machines with unsupported OS are treated as physical servers.

Per-workload Subscription Licensing

NAKIVO Backup & Replication can be licensed on a per-workload basis. A workload can be a VMware VM, Microsoft Hyper-V VM, Nutanix AHV VM, physical machine (1 physical server or 3 workstations), Oracle database (Enterprise Plus edition only), or Amazon EC2 instance that you plan to back up or replicate. Regardless of the type, each item is counted as one workload. Licenses are required only for the source side of backup and replication, that is, you do not need to license target servers for your replicas or the servers on which you want to recover workloads. Subscription licenses include 24/7 Support.

Subscription Licensing for Microsoft 365

Backup and Recovery for Microsoft 365 is licensed on a per-user, per-month basis. The license can be purchased together with any edition (Basic, Pro Essentials, Enterprise Essentials, Pro, Enterprise, or Enterprise Plus) and combined with any license type (Perpetual or Subscription). Subscription licenses for Backup and Recovery for Microsoft 365 include 24/7 Support. The minimum number of licenses per order is 10.

When combined with a perpetual license, the support end date of the perpetual license and subscription license for Backup and Recovery for Microsoft 365 must be aligned. The support level may be Standard for perpetual sockets and 24/7 for Microsoft 365 subscription. Optionally, the support level for perpetual sockets can be upgraded to 24/7 Support.

The table below provides information on licensing options.

License Type	VMware	Hyper-V	Nutanix AHV	Amazon EC2	Physical Machine	Microsoft 365	Oracle Database
Perpetual	+	+	+	-	+	-	+
Subscription	+	+	+	+	+	+	+

Licensing Rules

- Perpetual and subscription licenses cannot be combined in one license.
- Subscription license for Backup and Recovery for Microsoft 365 can be combined with both perpetual and subscription (per-workload) licenses.
- Shared mailboxes do not require a license for backup and recovery.

Note

For the most recent information about licensing, refer to the NAKIVO's pricing page.

NAKIVO Support

NAKIVO Backup & Replication offers two levels of technical support:

- Standard Support
- 24/7 Support

Standard Support provides coverage from Monday to Friday during business hours as defined in the Customer Support Policy. One year of Standard Support is included in all new perpetual license purchases.
24/7 Support provides 24/7/365 coverage via phone, chat or email. To switch from Standard to 24/7 Support, you need to purchase a support Upgrade. Customers who upgrade to a higher-tier edition and have purchased additional years of support are required to upgrade their support too.
If a support agreement has been expired for more than 1 month, it can be extended with Expired Support Renewal.
For more information about the terms and conditions of using NAKIVO Customer Support services, refer to the following resources:

- NAKIVO Customer Support Policy
- NAKIVO Customer Support Agreement
- End-User License Agreement

Deployment

This section contains the following topics :

- “Architecture” on page 85
- “System Requirements” on page 94
- “Installing NAKIVO Backup & Replication” on page 129
- “Updating NAKIVO Backup & Replication” on page 199
- “Uninstalling NAKIVO Backup & Replication” on page 230

Architecture

- What is NAKIVO Backup & Replication?
- Solution Components

What is NAKIVO Backup & Replication?

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

Solution Components

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

- “Director” on page 86
- “Transporter” on page 88
- “Backup Repository” on page 92

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

Director

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

What is Director?

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

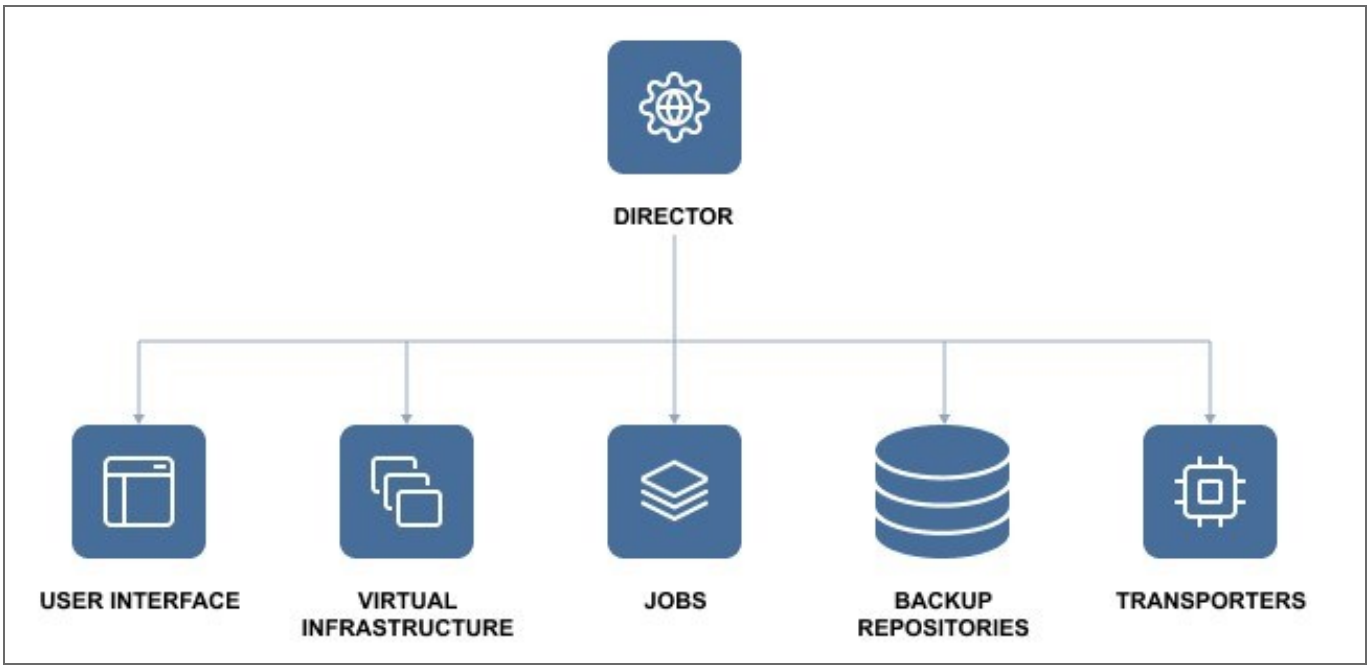

How Many Directors Should be Deployed

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

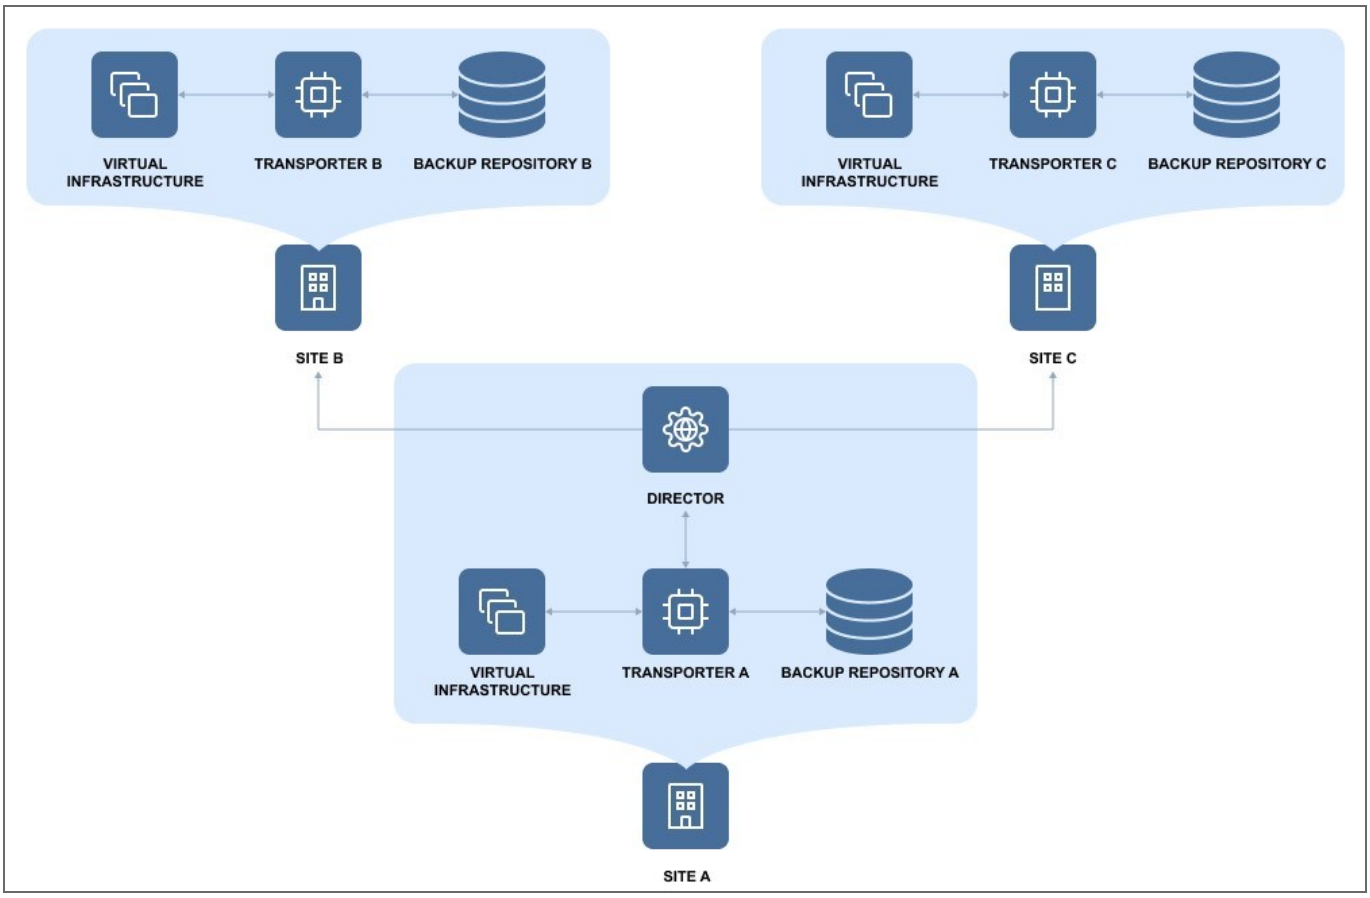
VIRTUAL INFRASTRUCTURE
TRANSPORTER B
BACKUP REPOSITORY B
VIRTUAL INFRASTRUCTURE
TRANSPORTER C
BACKUP REPOSITORY C
SITE B
SITE C
DIRECTOR
VIRTUAL INFRASTRUCTURE
TRANSPORTER A
BACKUP REPOSITORY A
SITE A

Transporter

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

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.

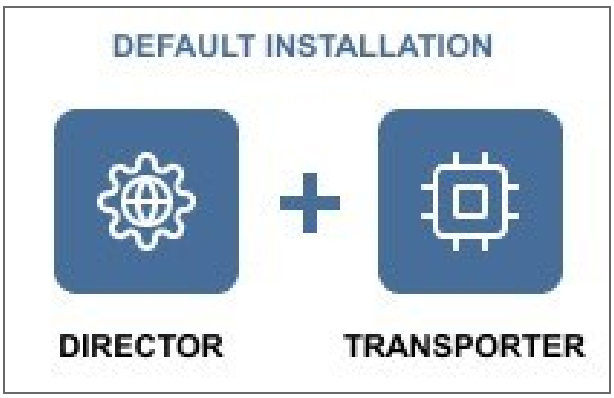

A single Transporter can back up, replicate, and recover multiple VMs and cloud instances.

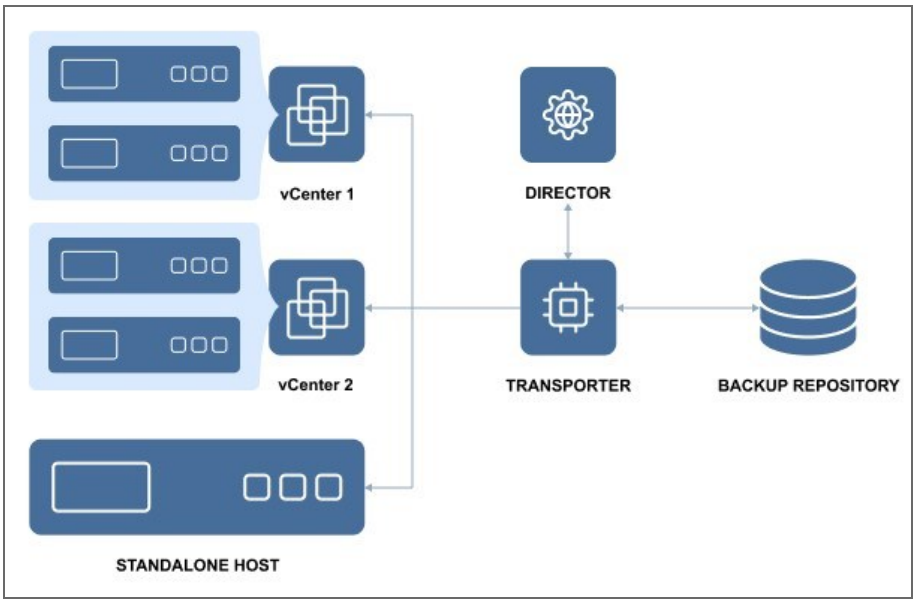

One Transporter can simultaneously process multiple source disks (6 by default) during backup, replication, and recovery. If jobs contain more disks than the Transporter is set to process simultaneously, the disks will be put in a queue and will be processed once the Transporter frees up.

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.

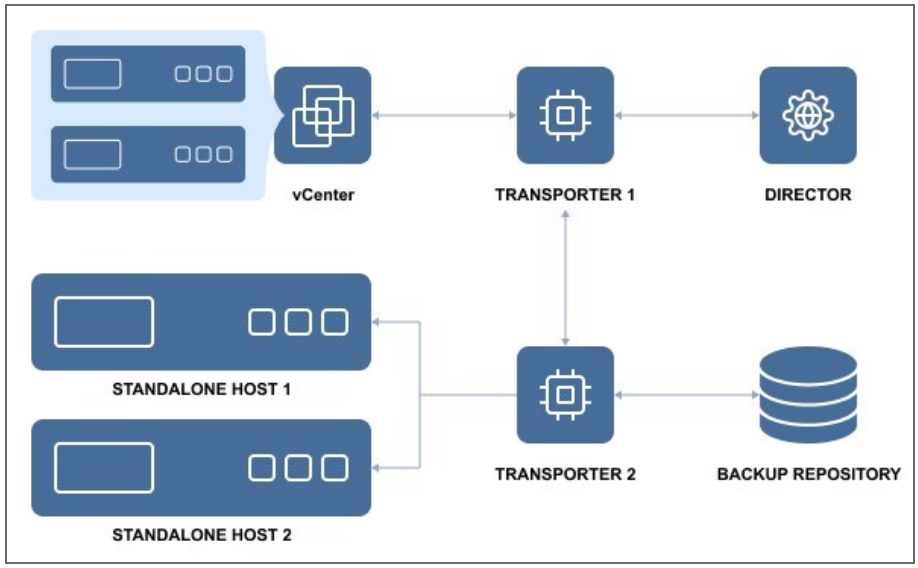

Deploying multiple Transporters also enables network acceleration and AES 256 encryption of traffic between a pair of Transporters. For example, if VMs are replicated over WAN between two sites, the Transporter installed in the source site can compress and encrypt data before transferring it over WAN,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

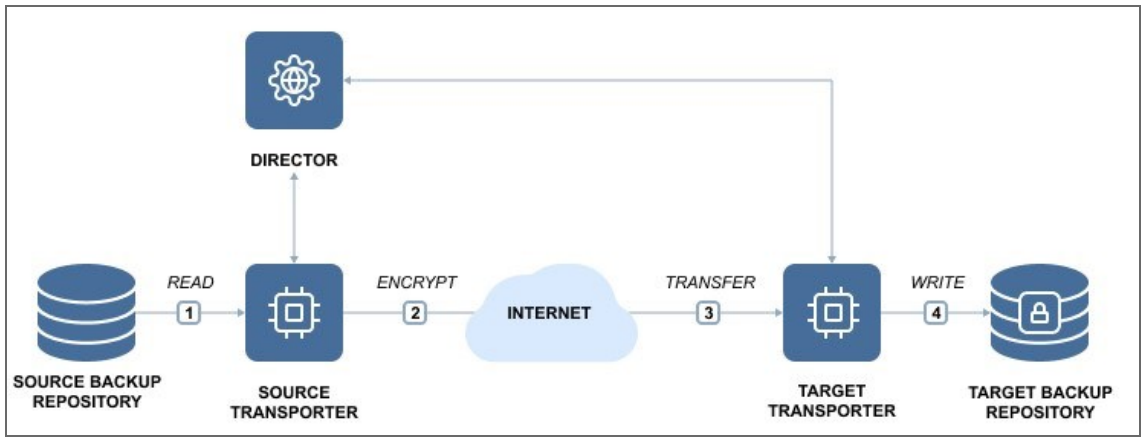

If you plan to transfer data over WAN without a VPN connection from your source site to the target site, make sure the source and target Transporters are added to the product using external IP addresses or DNS names that can be properly resolved in WAN, so that the two Transporters can connect to each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

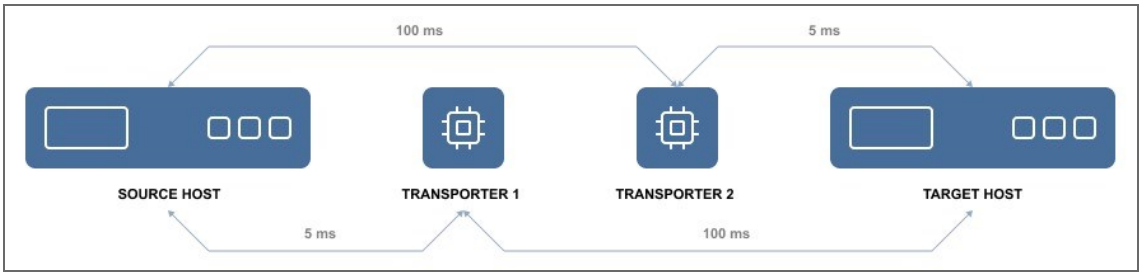

In the example above, Transporter 1 will be selected to read data from the Source ESXi, and Transporter 2 will be selected to write data to the Target ESXi.
The Transporter selection can also be configured manually during job creation.

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.

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. By default, backups are compressed and deduplicated at the block level across the entire Backup Repository 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.

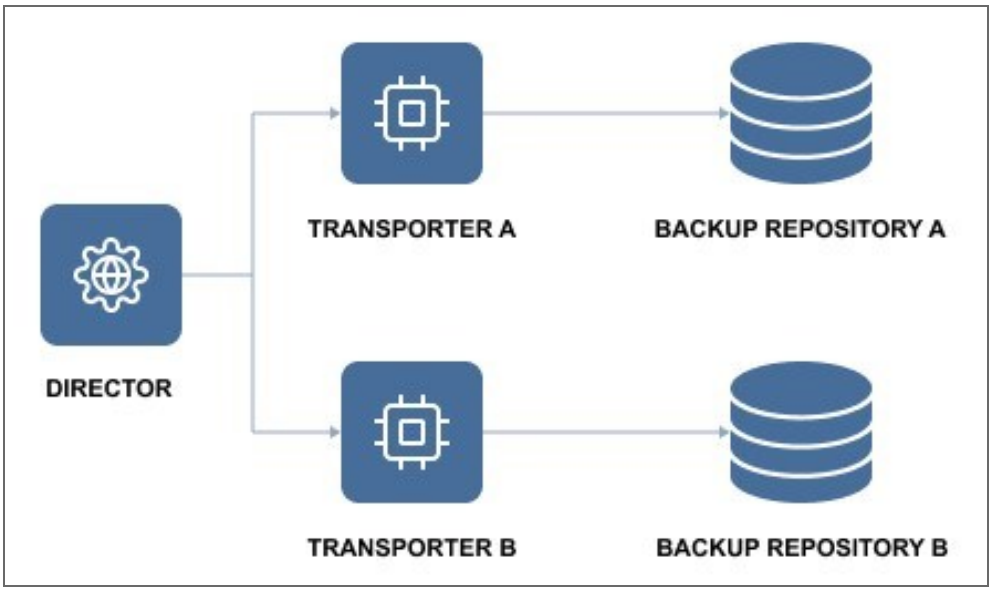

The Assigned Transporter is responsible for all interaction with its Backup Repository. A single Transporter can be assigned to and manage multiple Backup Repositories.

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 page 95
- "Storage Integration Requirements" on page 96
- "Deployment Requirements" on page 98
- "Feature Requirements" on page 120

Supported Platforms

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

- Microsoft Server with a Hyper-V role 2019, 2016, 2012R2, 2012

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
  - Microsoft Hyper-V

Hypervisor Requirements

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

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.

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 a stream 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 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, you will 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

Dell-EMC Data Domain

Supported versions:

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

NEC HYDRAstor

Supported systems:

- NEC HYDRAstor v 5.5.1
- NEC Storage HS Universal Express I/O Module Version 1.8.0

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

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

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 of 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 115
- **QNAP**: For a full list of supported models, refer to "Supported QNAP NAS Devices" on page 109
- **ASUSTOR**: For a full list of supported models, refer to "Supported ASUSTOR NAS Devices" on page 104
- **NETGEAR**: For a full list of supported. For a full list of supporter models, refer to "Supported NETGEAR NAS Devices" on page 106.
- **Western Digital**: For a full list of supported models, refer to "Supported Western Digital NAS Devices" on page 119.

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:

Windows

- 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)
- Windows 10 Professional (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 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 8.3 (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 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
- FreeNAS 11.3
- Netgear ReadyNAS OS v6.10.3
- Netgear ReadyNAS OS v6.9
- Synology DSM 6.2.3
- Synology DSM v6.2
- Synology DSM v6.1
- Synology DSM v6.0
- QNAP QTS v4.5.1
- QNAP QTS v4.4
- QNAP QTS v4.3
- WD MyCloud v3
- 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:

TC Port (Default)	Where	Description
NAKIVO Backup & Replication		
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.
VMware		
443	vCenter Server, ESXi host	Used by Director and Transporters to access VMware infrastructure. Must be opened on vCenter Servers and ESXi hosts.
902	ESXi hosts	Used by Transporters to access VMware infrastructure. Must be opened on ESXi hosts.
Hyper-V		
137 - 139	Hyper-V hosts	Used by Director to upload files and install configuration service. Must be opened on Hyper-V servers.
445	Hyper-V hosts	Used by Director to upload files and install configuration service.
5986 (opens automatically)	Hyper-V hosts	Used by Transporter to add a host to inventory and establish a connection with it.
9445 (opens automatically)	Hyper-V hosts	Used by Director to upload files and install configuration service. Must be opened on Hyper-V host if NAKIVO Backup & Replication is installed on a host and this host is added to inventory simultaneously.
9446 (opens automatically)	Hyper-V hosts	Used by Director and Transporters to communicate with the Transporter. 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.
Physical machine (Windows)		
445	Windows machine	Used by Director to upload files and install configuration service via SMB.

9446 (opens automatically)	Windows machine	Used to create the Transporter installed by default.
Physical machine (Linux)		
22	Linux machine	Used by Director to access a Linux physical machine via SSH.
9446 (opens automatically)	Linux machine	Used to create the Transporter installed by default.

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

Transporter Only

- AS1002T
- AS1002T v2
- AS1004T
- AS1004T v2

For minimum hardware requirements, refer to “Network Attached Storage” on page 98.

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 98

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

Transporter Only

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

For minimum hardware requirements, refer to "Network Attached Storage" on page 98.

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+
- 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+

Transporter Only

- RS217
- RS214
- DS416slim
- DS416j
- DS414slim
- DS414j
- DS218j
- DS216
- DS216j
- DS215j
- DS214
- DS213j
- DS115
- DS114
- DS220j
- DS419slim

Important

Backup and recovery of Microsoft 365 accounts is not supported on Synology NAS devices with ARMv8 CPU architecture.

For minimum hardware requirements, refer to "Network Attached Storage" on page 98.

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 98.

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.

- File Recovery
- 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
- Cross-Platform Recovery
- App-Aware Mode
- Encrypted Backup Repository
- Native Tape Support
- Backup Immutability

File Recovery

Recovered files can be downloaded or sent via email. They can also be recovered to a server. Below are the requirements which must be met for every 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
  - 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)
  - Windows Server 2008 R2 Standard 32-bit
  - Windows 10 Professional (x64)
  - Windows 8 Professional (x64)
  - Windows 8 Professional (x32)
  - Windows 7 Professional (x64)
  - Windows 7 Professional (x32)
- Linux
  - Ubuntu 18.04. Server (x64)
  - Ubuntu 16.04 Server (x64)
  - Ubuntu 16.04 Server (x32)
  - Ubuntu 12.04 Server (x64)
  - Ubuntu 12.04 Server (x32)
  - SUSE Linux Enterprise Server 12 SP3 (x64)
  - SUSE Linux Enterprise Server 12 SP2 (x64)
  - SUSE Linux Enterprise Server 12 SP1 (x64)
  - SUSE Linux Enterprise Server 11 SP4 (x64)
  - SUSE Linux Enterprise Server 11 SP4 (x32)
  - SUSE Linux Enterprise Server 11 SP3 (x64)
  - SUSE Linux Enterprise Server 11 SP3 (x32)
  - Red Hat Enterprise Linux 7.4 (x64)
  - Red Hat Enterprise Linux 7.3 (x64)
  - Red Hat Enterprise Linux 7.2 (x64)
  - Red Hat Enterprise Linux 7.1 (x64)
  - Red Hat Enterprise Linux 7.0 (x64)
  - CentOS 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.

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".
- 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.

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.

Cross-Platform Recovery

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

	Target Platforms	
Source Platforms	**VMware vSphere 6.7**	**Microsoft Hyper-V 2016/2019**
VMware vSphere 6.7	• Windows Server 2016/2019 • Ubuntu Server 18.04 • RHEL 7	• Windows Server 2016/2019 • Ubuntu Server 18.04 • RHEL 7*

Microsoft Hyper-V 2016/2019	• Windows Server 2016/2019 • Ubuntu Server 18.04 • RHEL 7	• Windows Server 2016/2019 • Ubuntu Server 18.04 • RHEL 7
Physical Machines	• Windows Server 2016/2019 • Ubuntu Server 18.04 • RHEL 7	• Windows Server 2016/2019 • Ubuntu Server 18.04 • RHEL 7

* 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:

Microsoft Hyper-V

- Target system should be in the running state and custom OS quiescing is enabled for it.
- Connection between the product and target system should be established.
- Refer to the Knowledge Base article for more details.

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 feature supports the following Transporter operating systems:

- Windows
- Linux

Additionally, the following requirements must be met:

- 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.

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:
  - Four virtual processors.
  - 16 GiB of reserved RAM.
  - 80 GiB of disk space for the installation of VM image and system data.
- 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.

To make backups immutable, the following condition must be met for Amazon S3:

- Object Lock must be enabled for the Amazon S3 bucket where your Backup Repository is located.

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.
- Only Linux OS and NAS OS, except FreeNAS/TrueNAS, specified in system requirements are supported.
- The Backup Repository type must be **Local Folder**.
- The Backup Repository data storage type must be set to **Incremental with full backups**.

Installing NAKIVO Backup & Replication

Refer to the sections below to learn how to install NAKIVO Backup & Replication:

- "Deploying VMware Virtual Appliance" on page 130
- "Deploying Nutanix AHV Virtual Appliance" on page 145
- "Installing on Windows" on page 154
- "Installing on Linux" on page 165
- "Installing on Synology NAS" on page 174
- "Installing on QNAP NAS" on page 181
- "Installing on Western Digital NAS" on page 191
- "Installing on ASUSTOR NAS" on page 186
- "Installing on NETGEAR ReadyNAS" on page 193
- "Installing on Generic ARM-based Device" on page 197
- "Deploying Amazon Machine Image in Amazon EC2" on page 153
- "Installing on FreeNAS" on page 196
- "Installing on Raspberry Pi" on page 198

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.

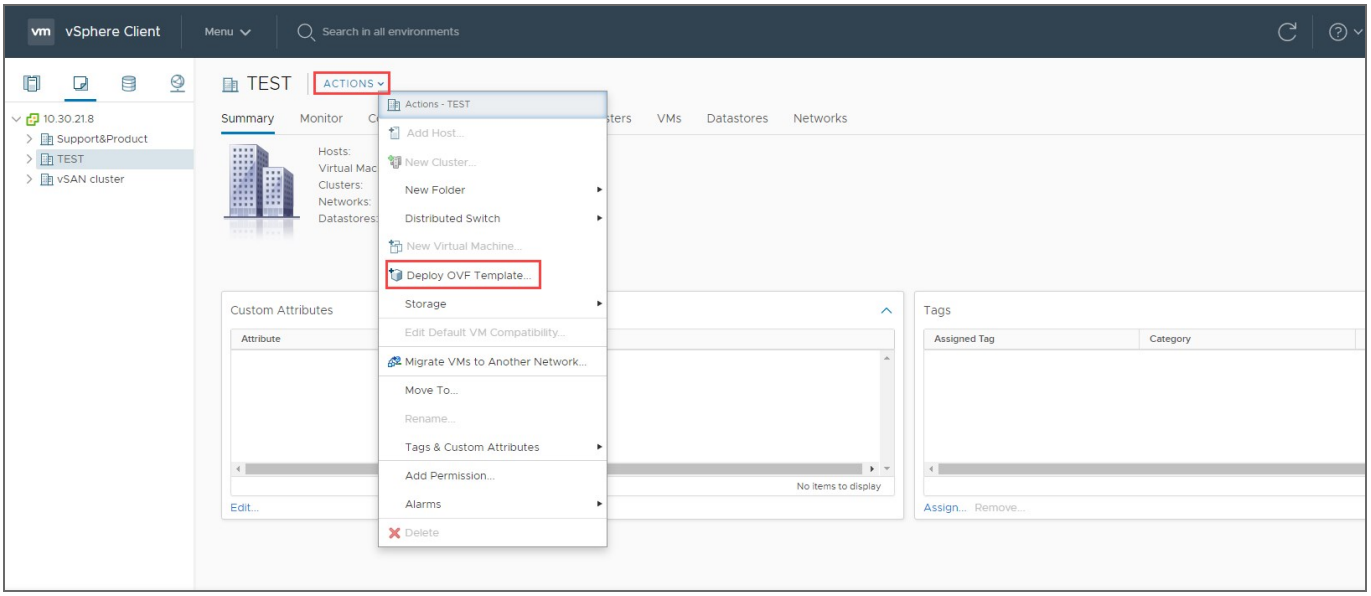

4. On the **Select an OVF template** page of the **Deploy OVF Template** wizard, select **Local file** and upload the VA file (.ova) you've downloaded. Click **Next**.

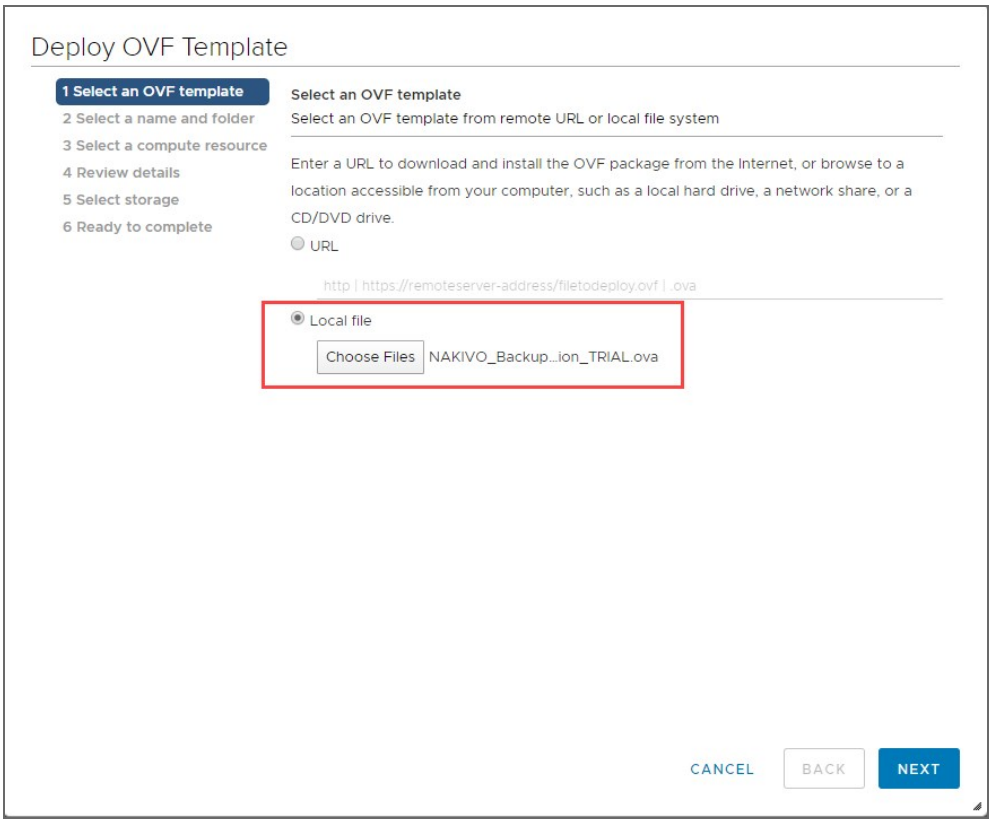

5. On the **Select a name and folder** page, specify a unique name and target location for the Virtual Appliance. Click **Next**.

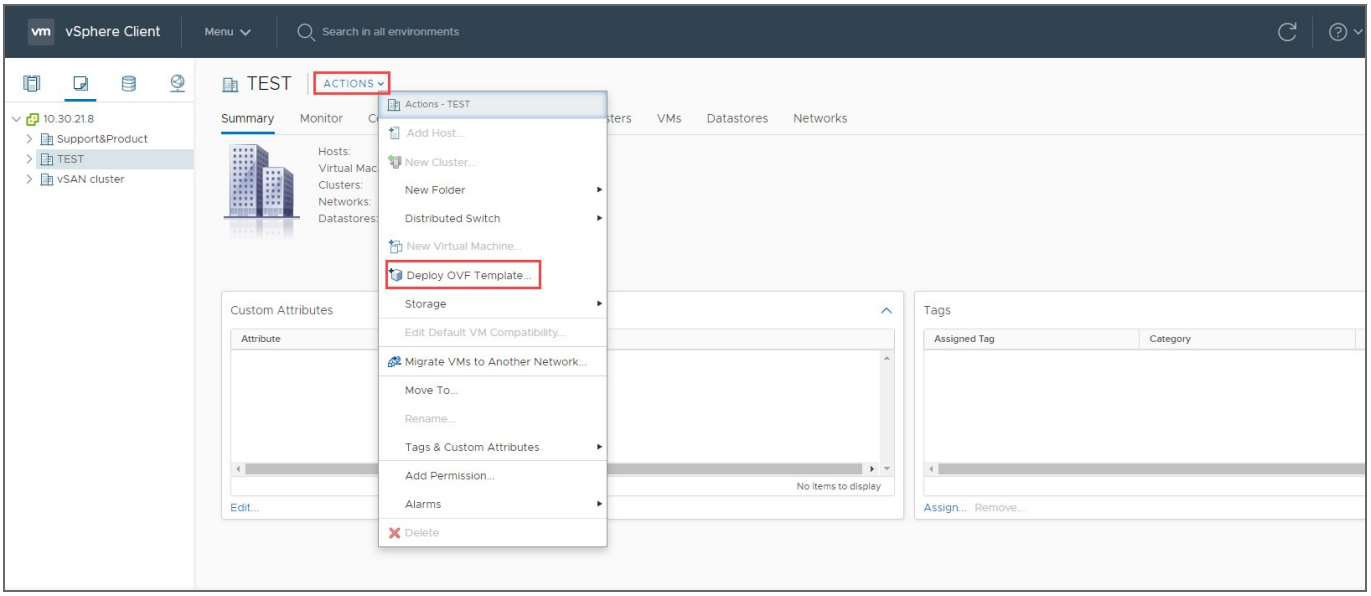

6. On the **Select a computer resource** page, select the resource pool within which you would like to deploy the Virtual Appliance and click **Next**.

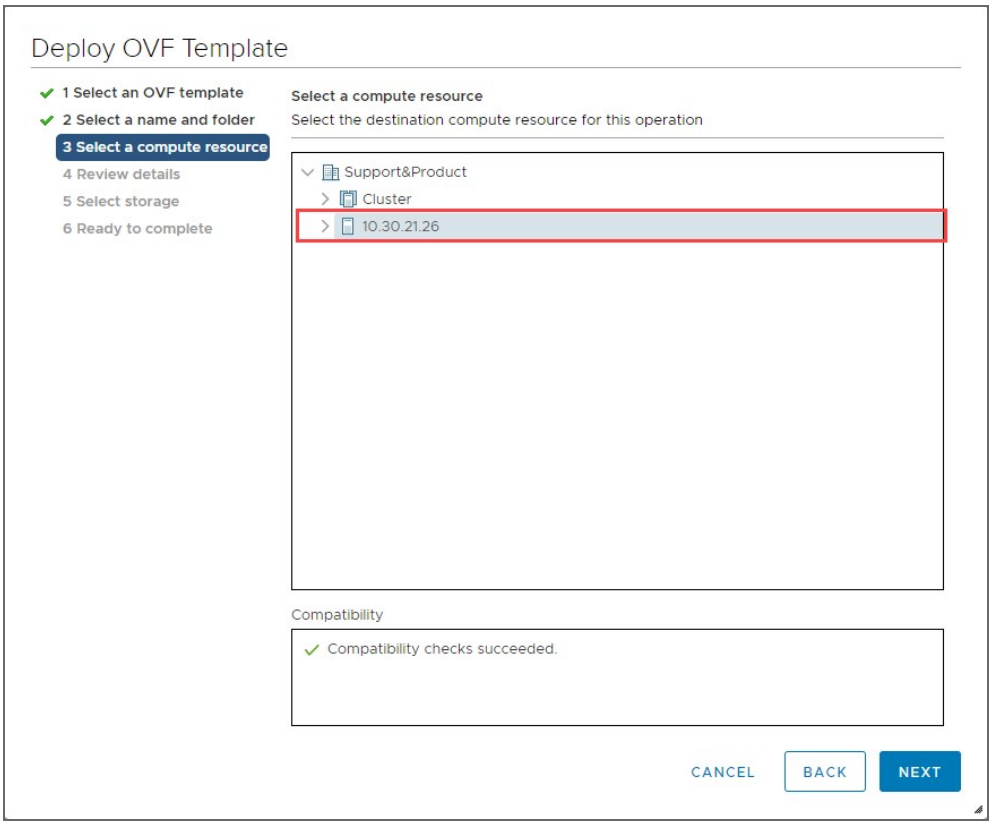

7. On the **Review details** page, review the template details and click **Next**.

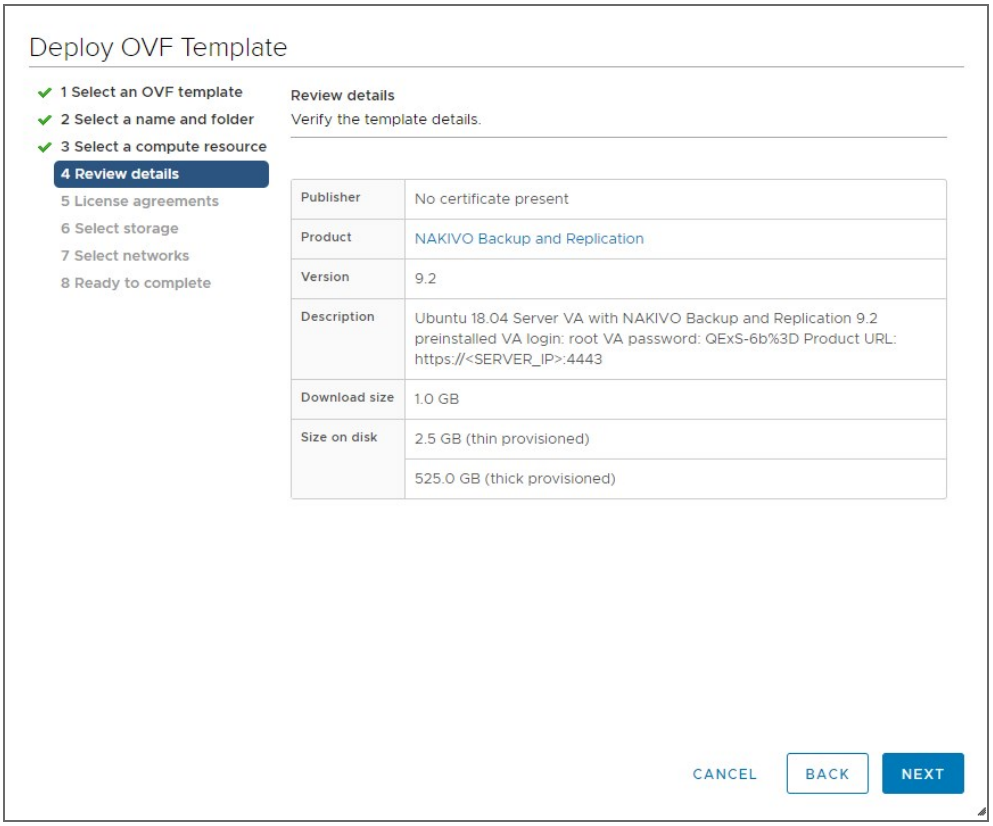

8. On the **License agreements** page, read the end-user license agreement (EULA). If you agree to its terms, select **I accept all license agreements** and then click **Next**.

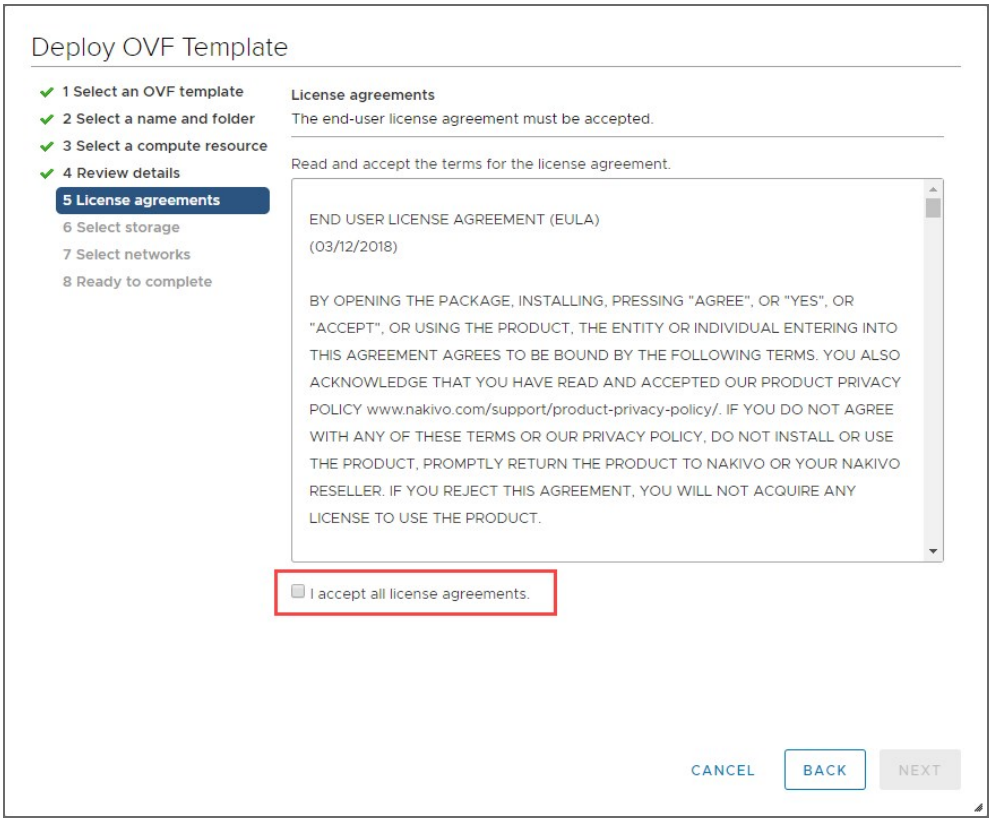

9. On the **Select storage** page, select a datastore in which you would like to keep the Virtual Appliance disk, virtual disk format (*Thin Provisioning* is recommended), VM storage policy and click **Next**.

   Important

   If you use thick provisioning instead of thin provisioning, keep in mind that NAKIVO Backup & Replication can take up to 0,5 TB of data. Check to see if it is 0,5 TB by default for all cases.

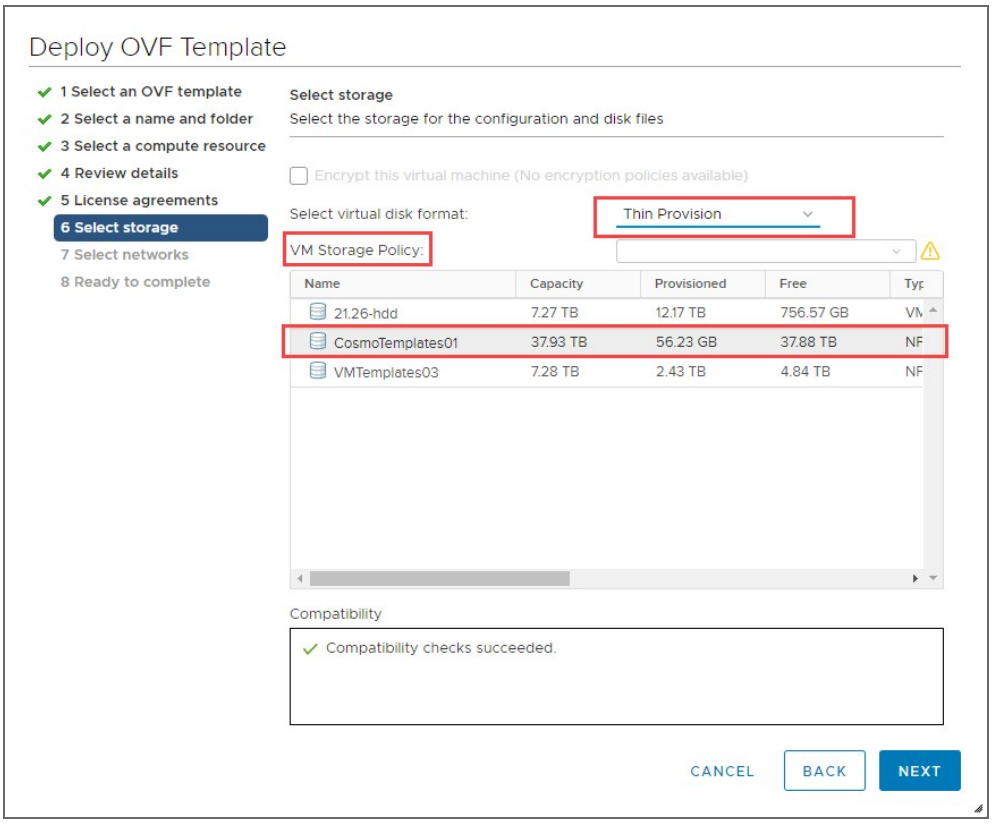

10. On the **Select networks** page, select a network to which the Virtual Appliance will be connected. Opting for a network with DHCP and Internet access is recommended. Click **Next**.

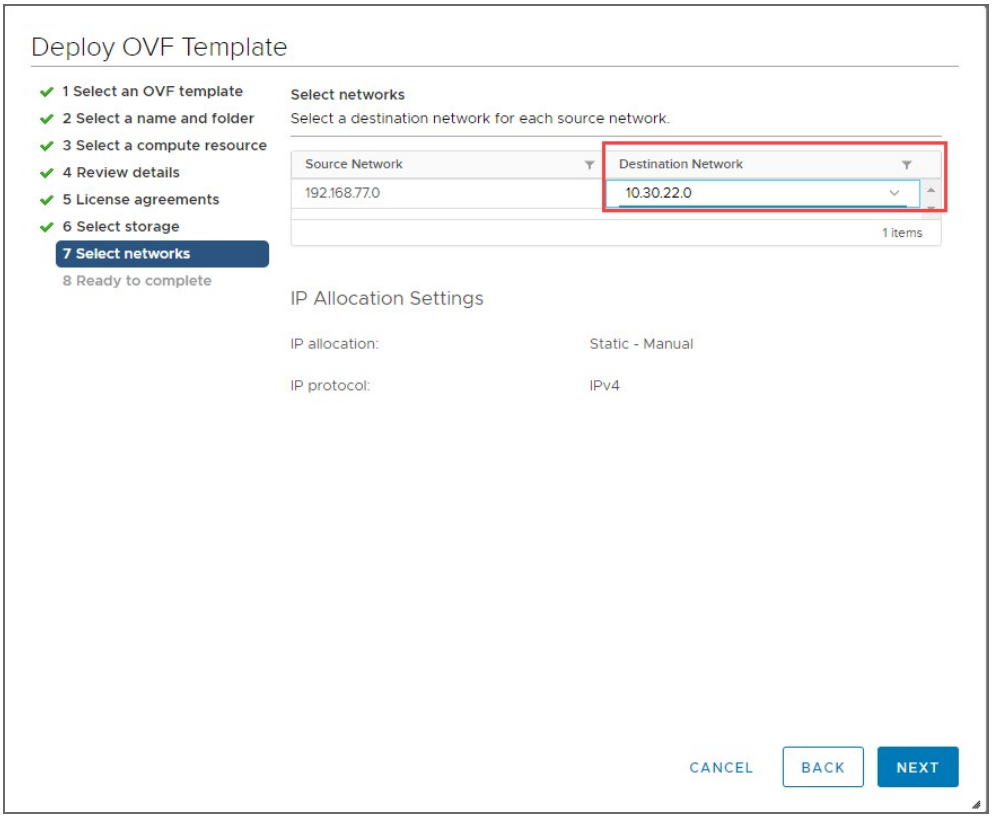

11. On the **Ready to complete** page, review the summary of the setups you have configured and click **Finish** to complete deployment.

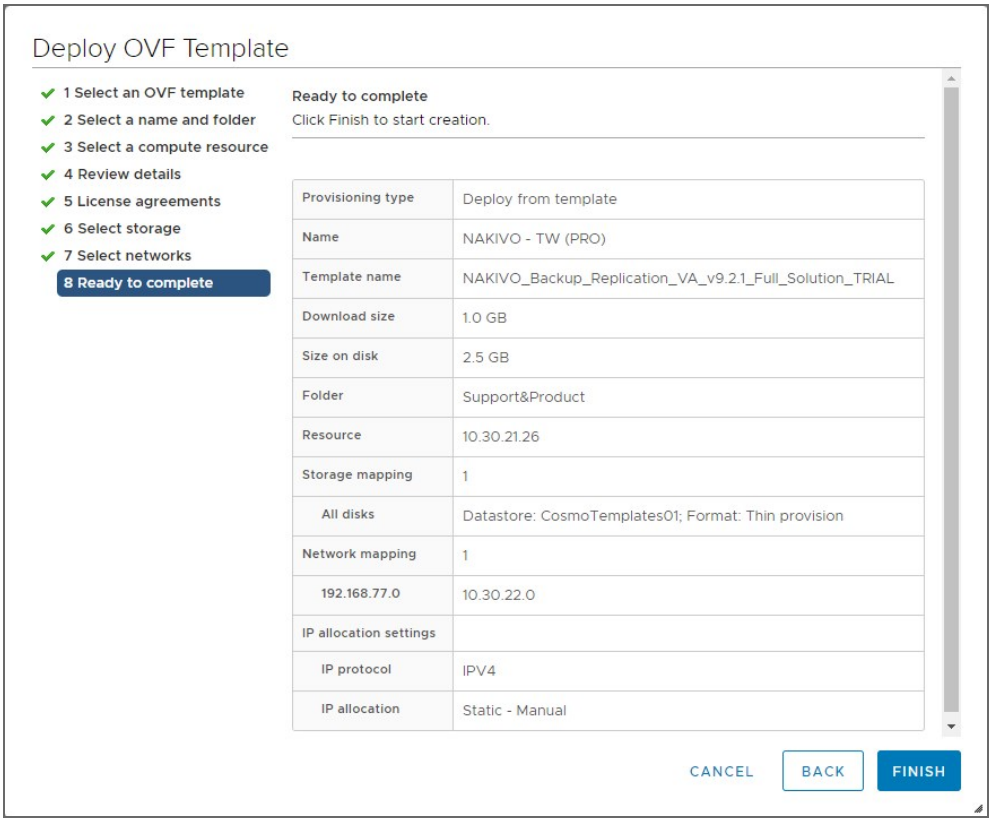

After the Virtual Appliance is deployed, you may need to configure it.

Important

If you plan to expose the Virtual Appliance to the Internet, change the default credentials and set up a login and password for the Web interface

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 OV Template.**

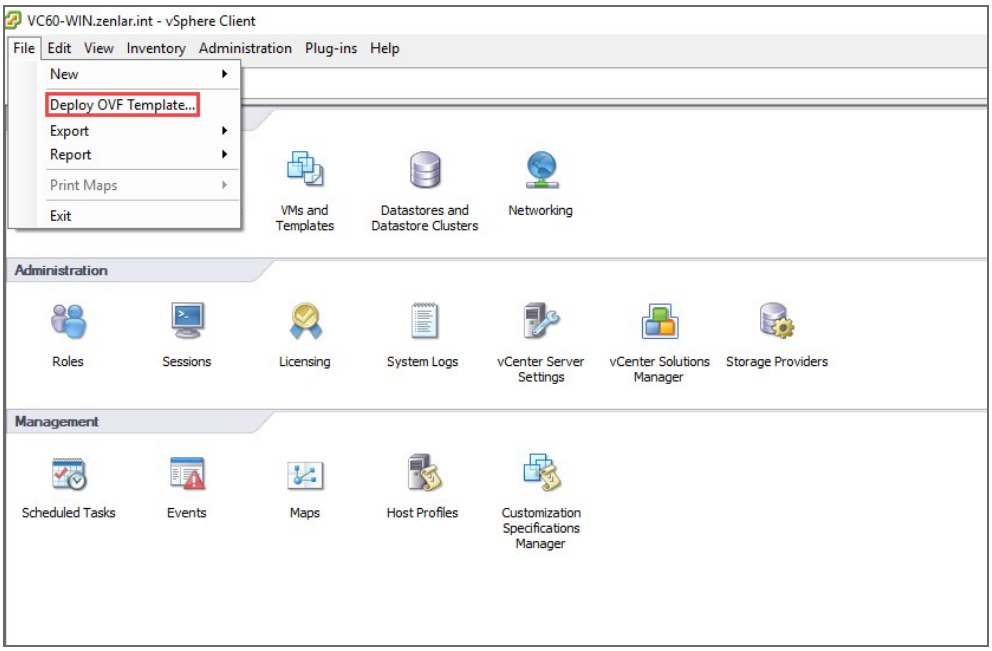

3. On the **Source** page of the **Deploy OVF Template** wizard, select and locate the file with the template. Click **Next**.
4. On the **OVF Template Details** page, review the template details and click **Next**.

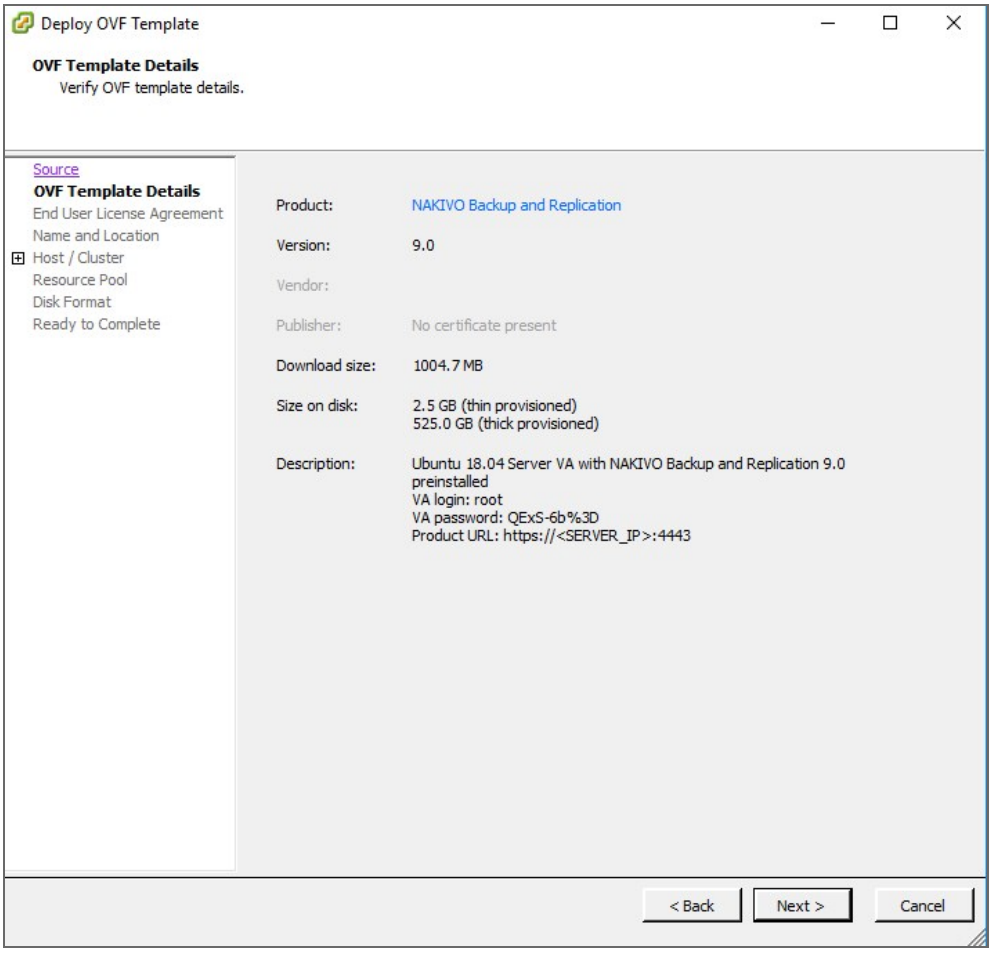

5. On the **End User License Agreement** page, read the license agreement. If you agree to its terms, click **Accept** and then click **Next.**

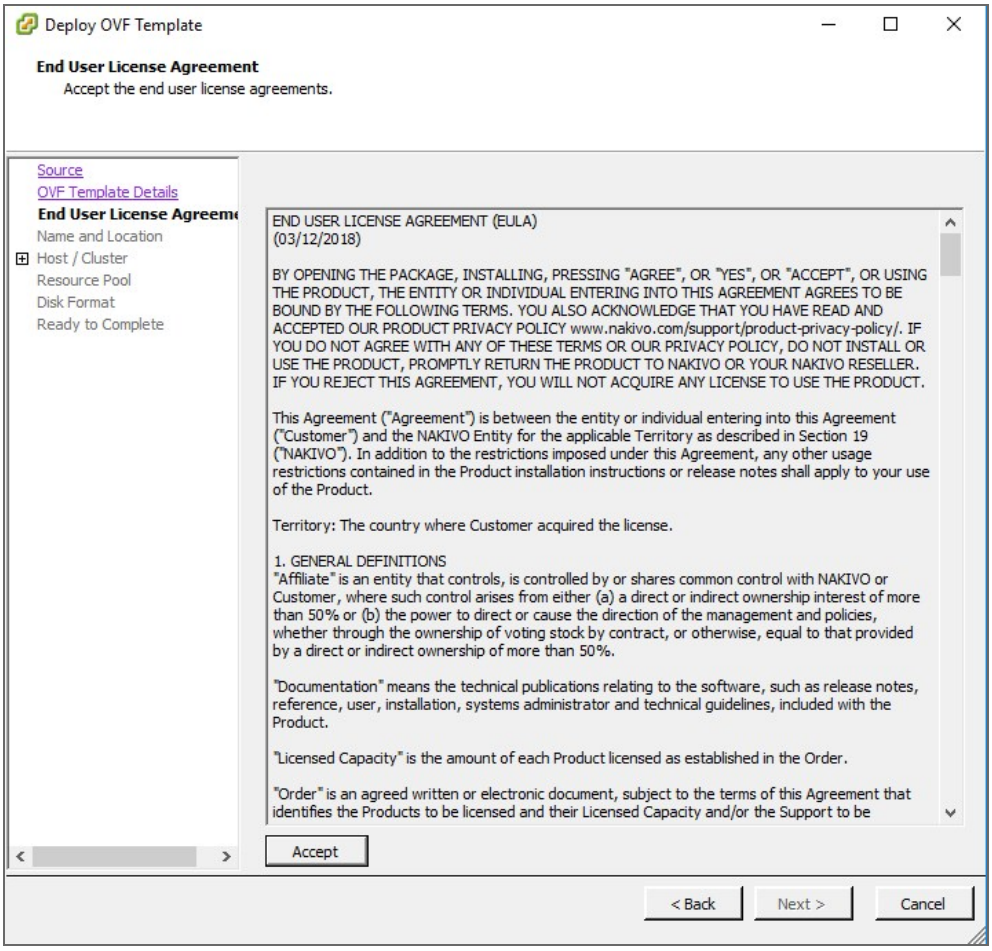

6. On the **Name and Location** page, specify a name and location for the deployed VA and click **Next**.

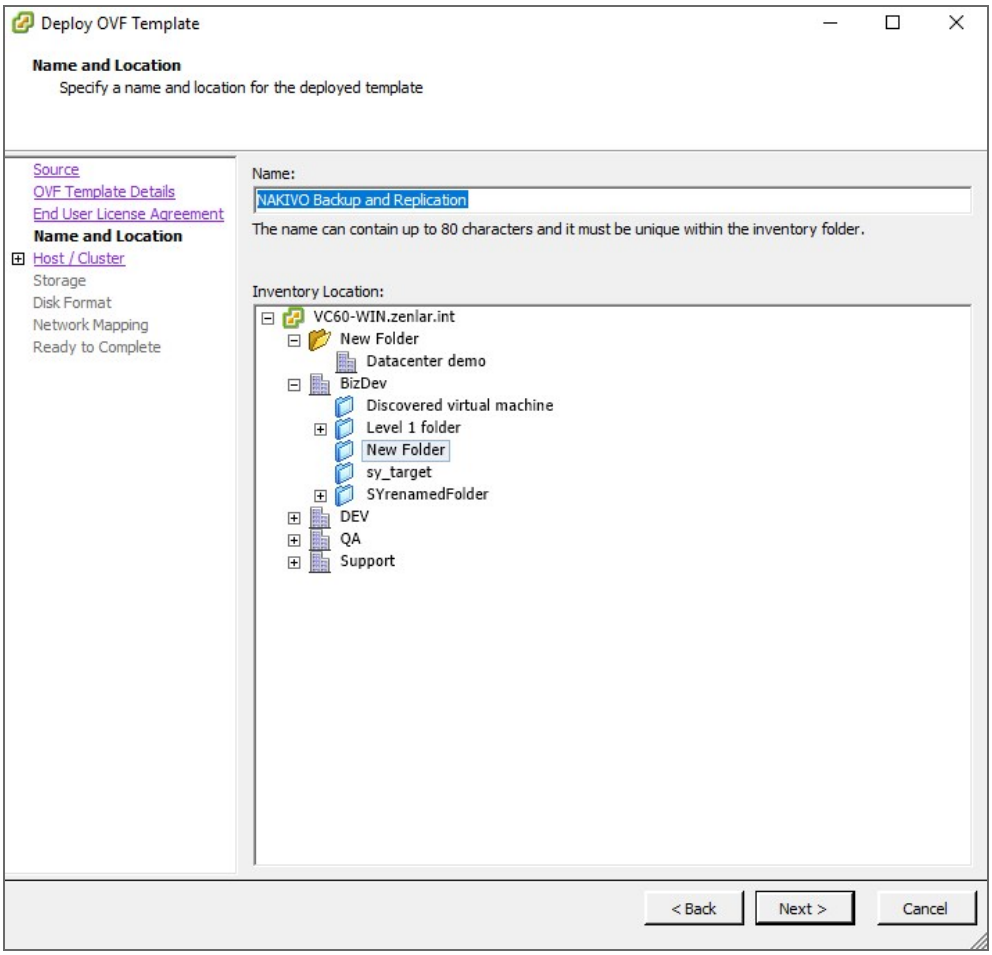

7. On the **Host/Cluster** page, select the host or cluster on which you wish to run the deployed template and click **Next**.

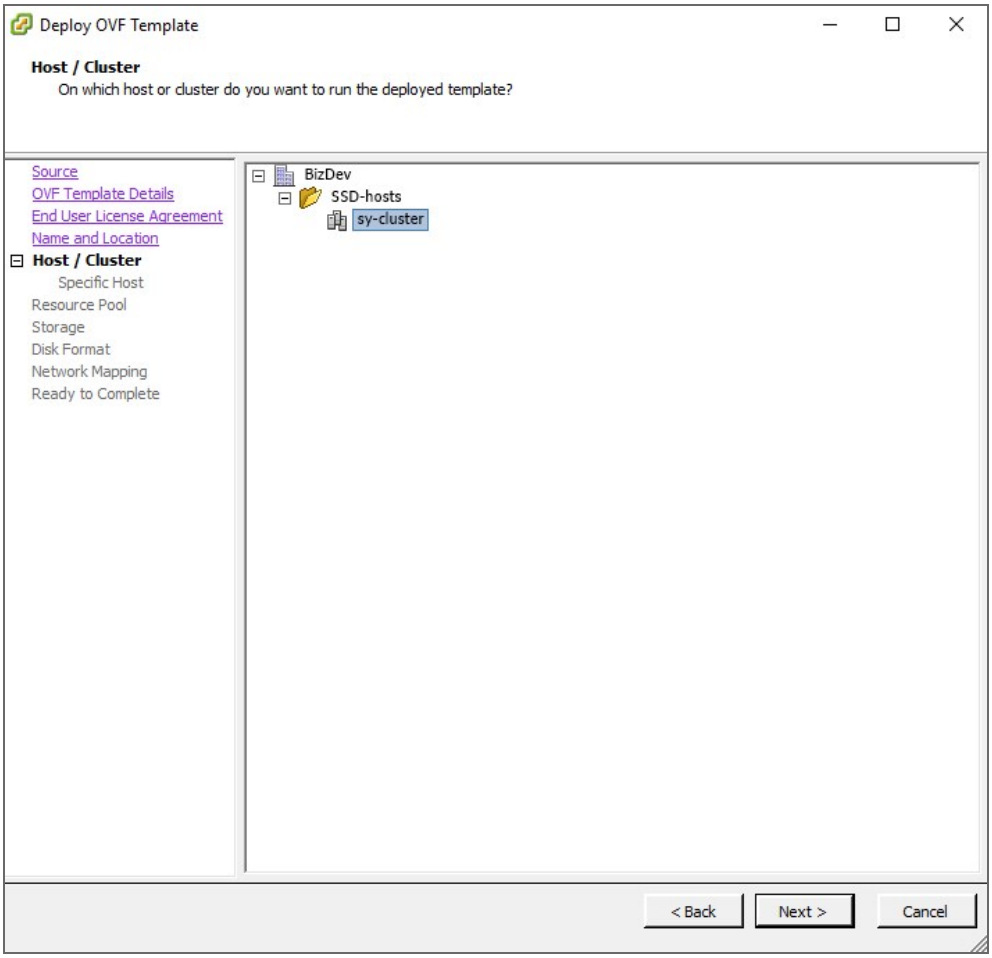

8. On the **Storage** page, select a datastore where you would like to keep the VA disk and click **Next**.

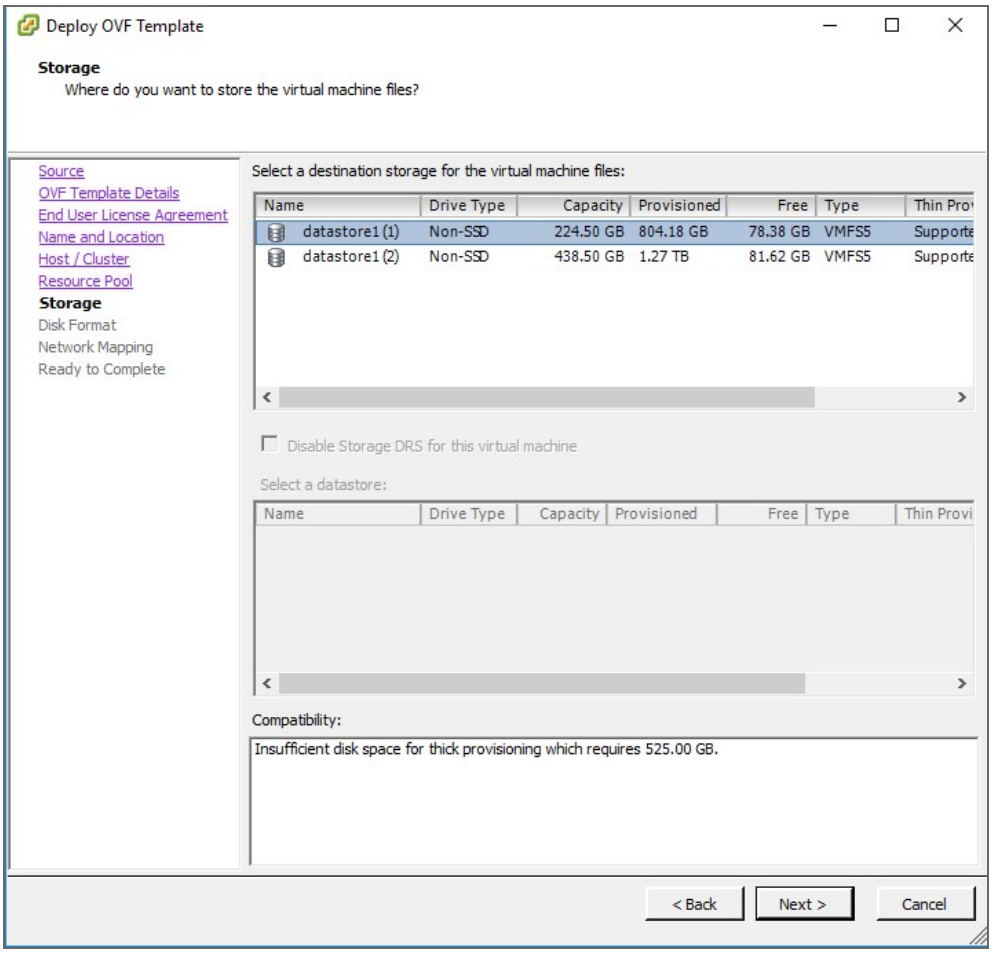

9. On the **Disk Format** page, select a virtual disk format (**Thin Provision** is recommended) and click **Next**.

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.

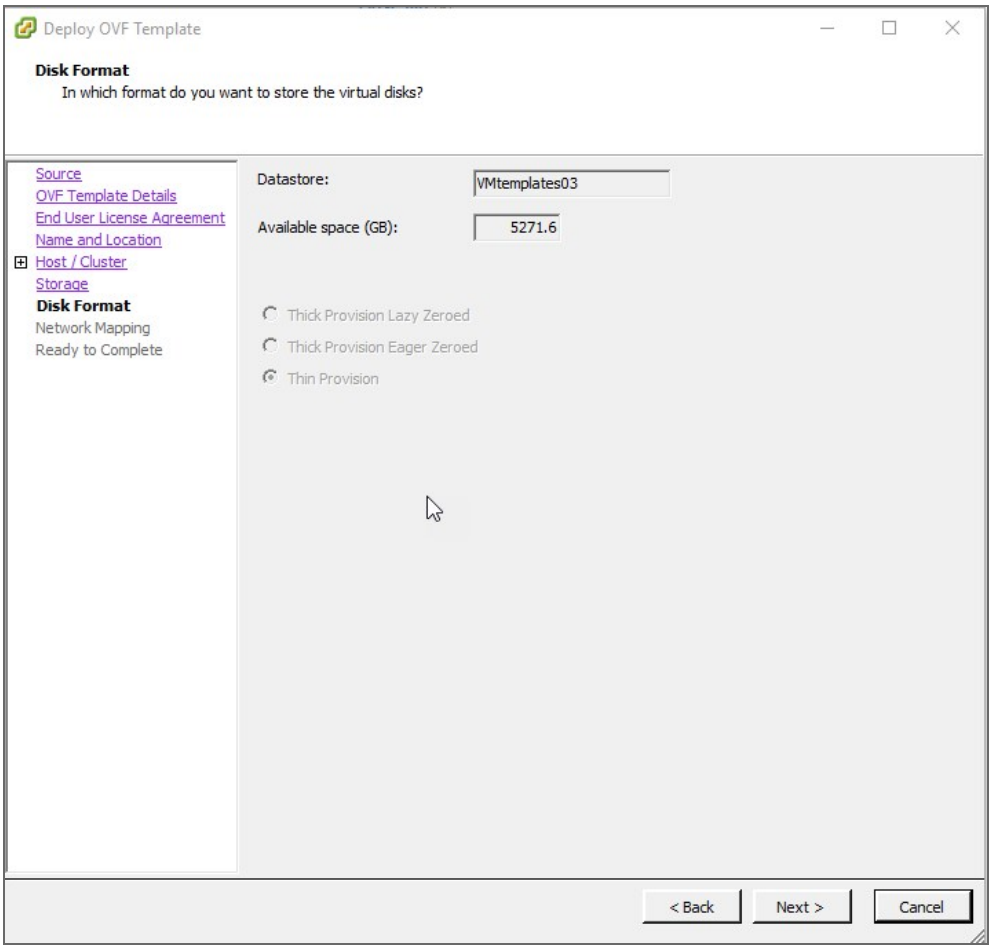

10. On the **Network Mapping** page, select a network to which the VA will be connected. It is recommended that you choose a network with DHCP and Internet access. Click **Next**.

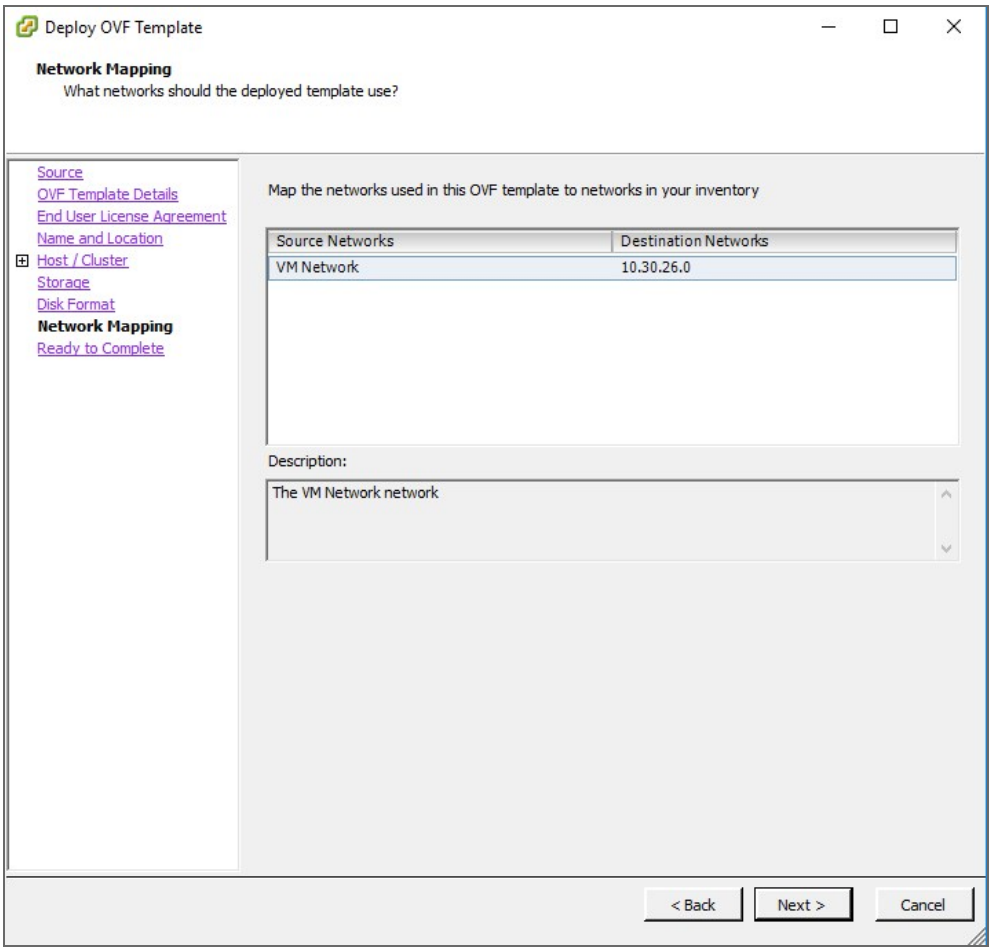

11. On the **Ready to Complete** page, review the summary of the options you have configured and select the **Power on after deployment** option.

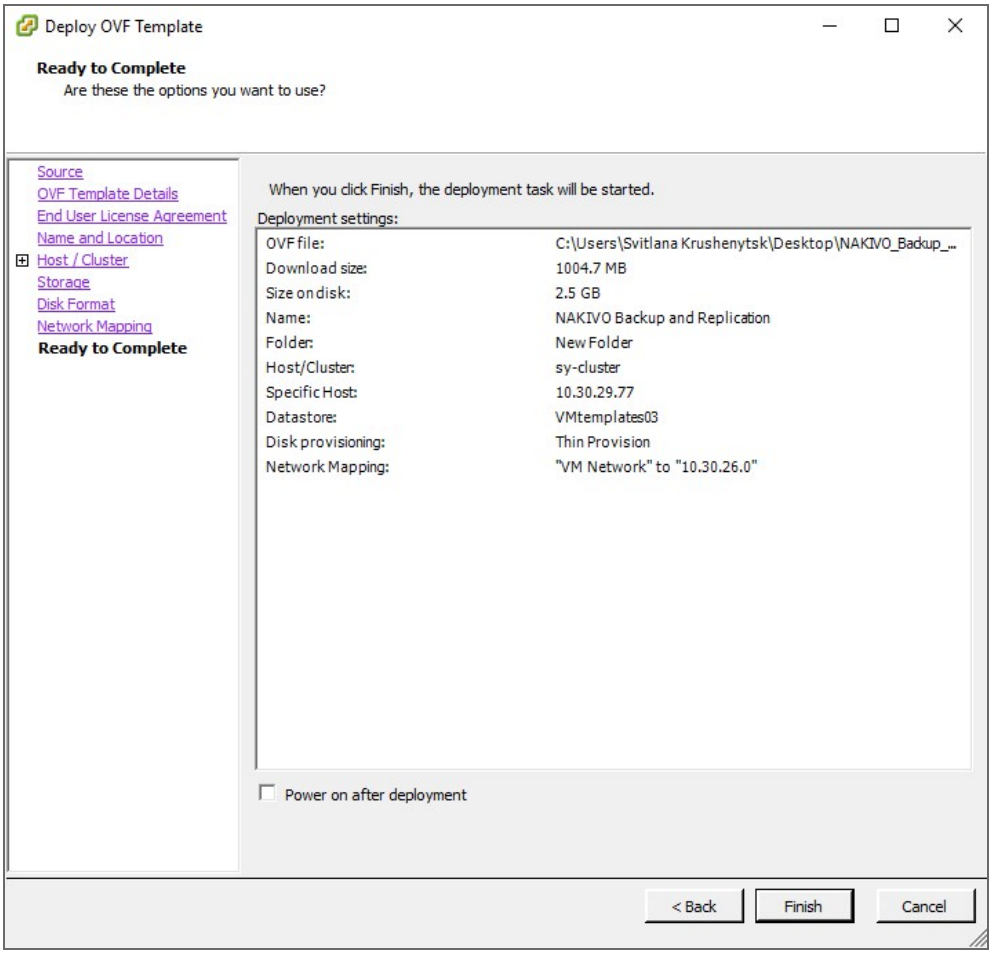

12. Click **Finish** to complete the deployment.
13. After the Virtual Appliance is deployed, configure it if necessary.

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`

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.

To enable Backup Immutability for **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
  - Turns on execshield protection
  - 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
- Prevents IP Spoofing
- Installs and configures fail2ban

Web Interface Login

Open the following URL to access the product's web interface: `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**.

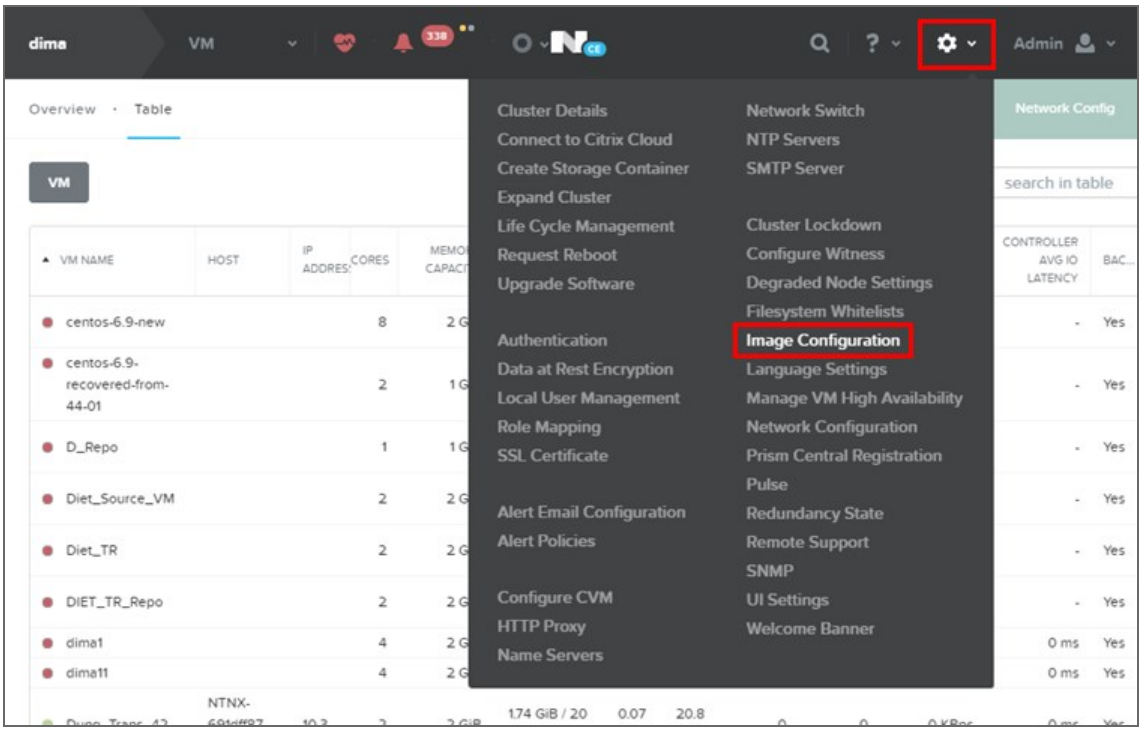

4. In the **Image Configuration** dialog, click **Upload Image**.

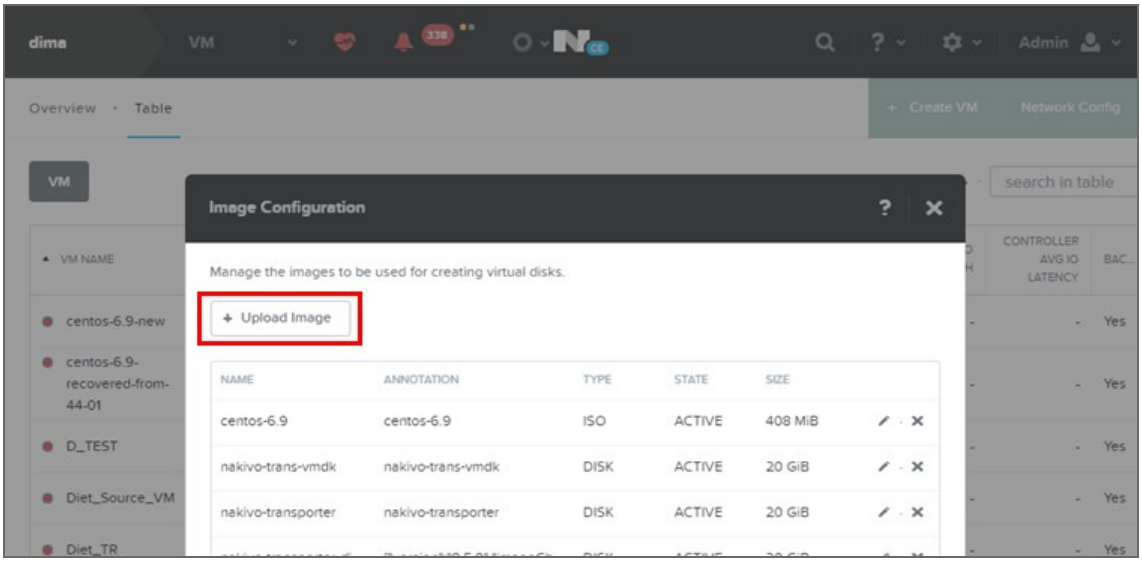

5. In the **Create Image** dialog, fill in the following options:
   - **Name**: Enter a name for the new image.
   - **Image Type**: From the drop-down list, select **DISK**.
   - **Storage Container**: Select the storage container you wish to use from the drop-down list. The list includes all storage 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.

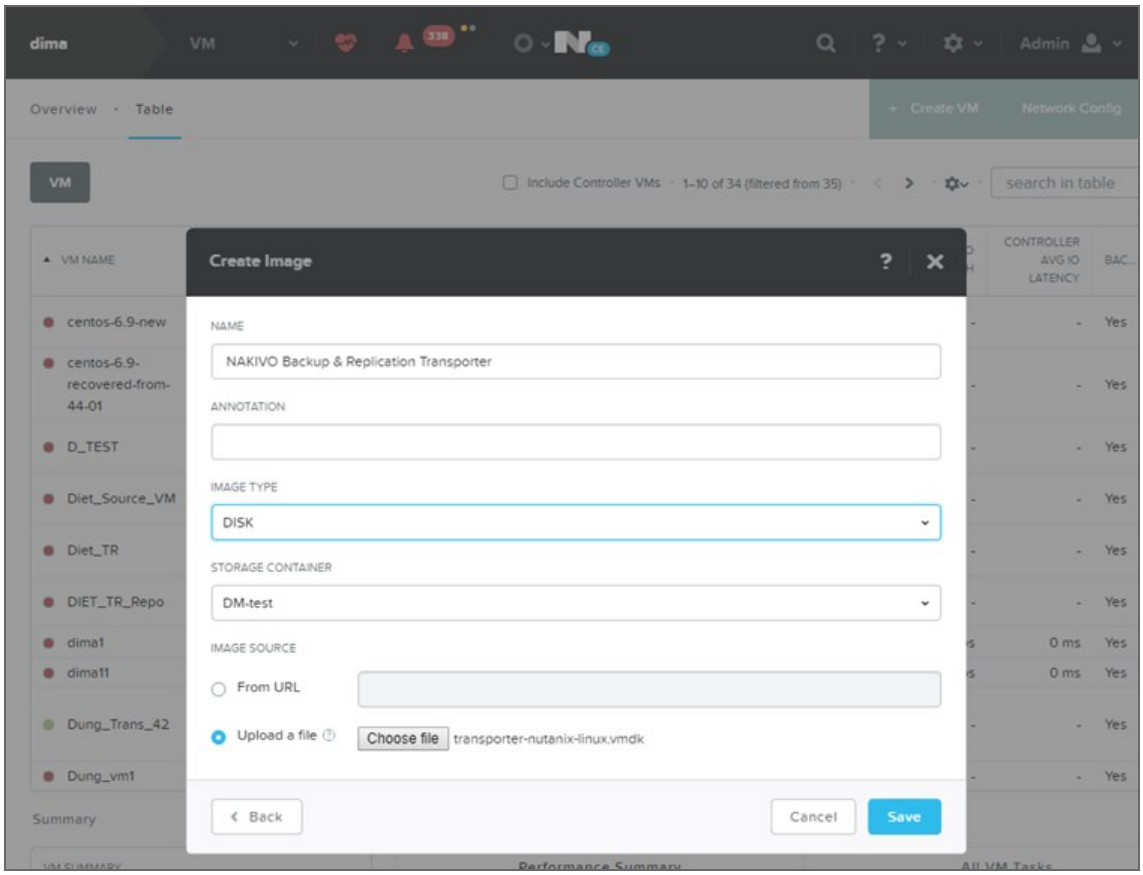

6. When all fields are correct, click the **Save** button.
   After the file uploading completes, the **Create Image** window closes and the **Image Configuration** window reappears with the new image present in the list.

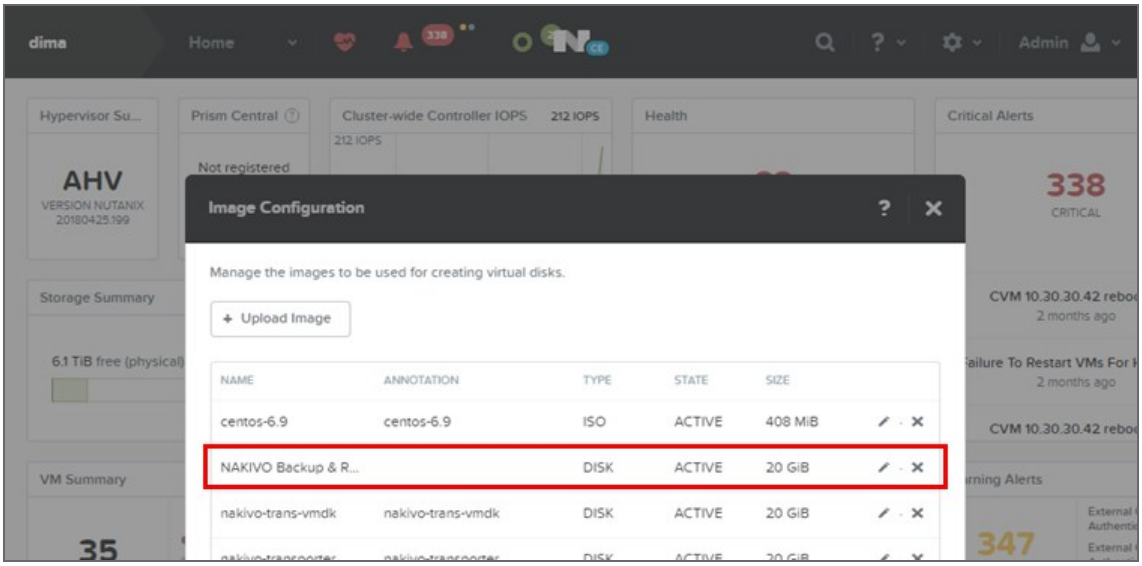

Note

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**.

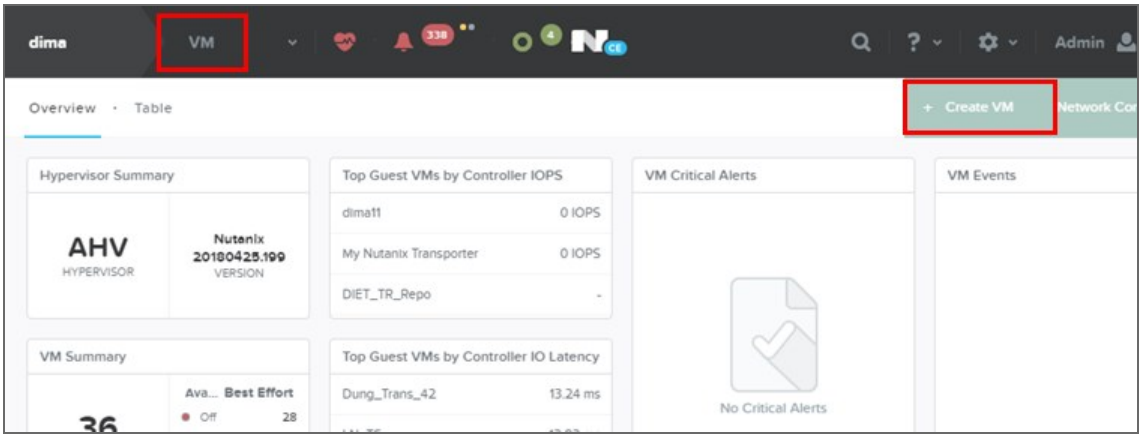

8. In the **Create VM** dialog, fill in the following options:
   - **Name**: Enter a name for the VM.
   - **vCPU(s)**: Enter the number of virtual CPUs to allocate to this VM (minimum 1).
   - **Number of Cores per vCPU**: Enter the number of cores assigned to each virtual CPU (minimum 2).
   - **Memory**: 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.

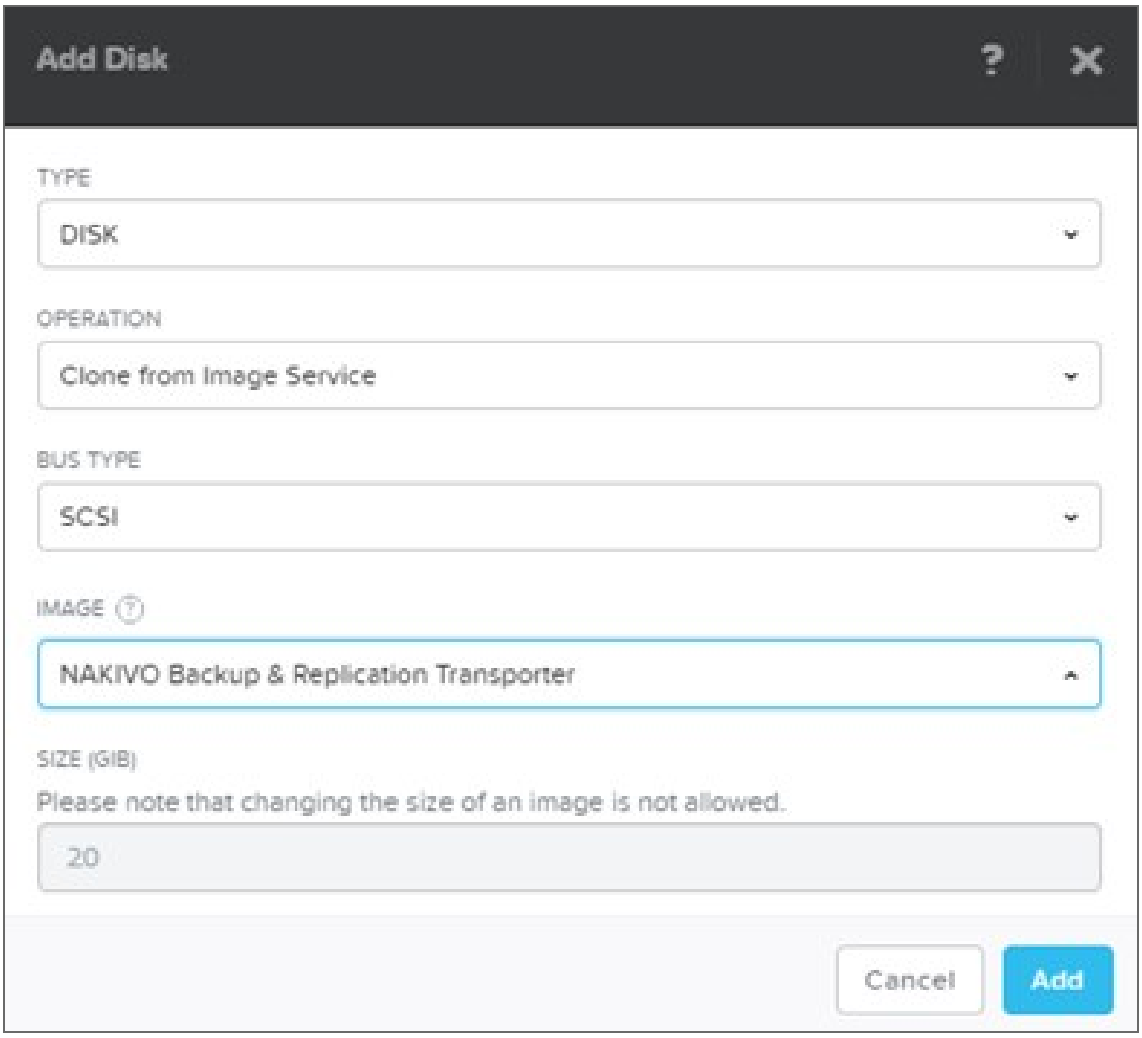

- In the **Network Adapters (NIC)** section, click **Add New NIC** and select an available VLAN from the list.

9. Click **Save**.

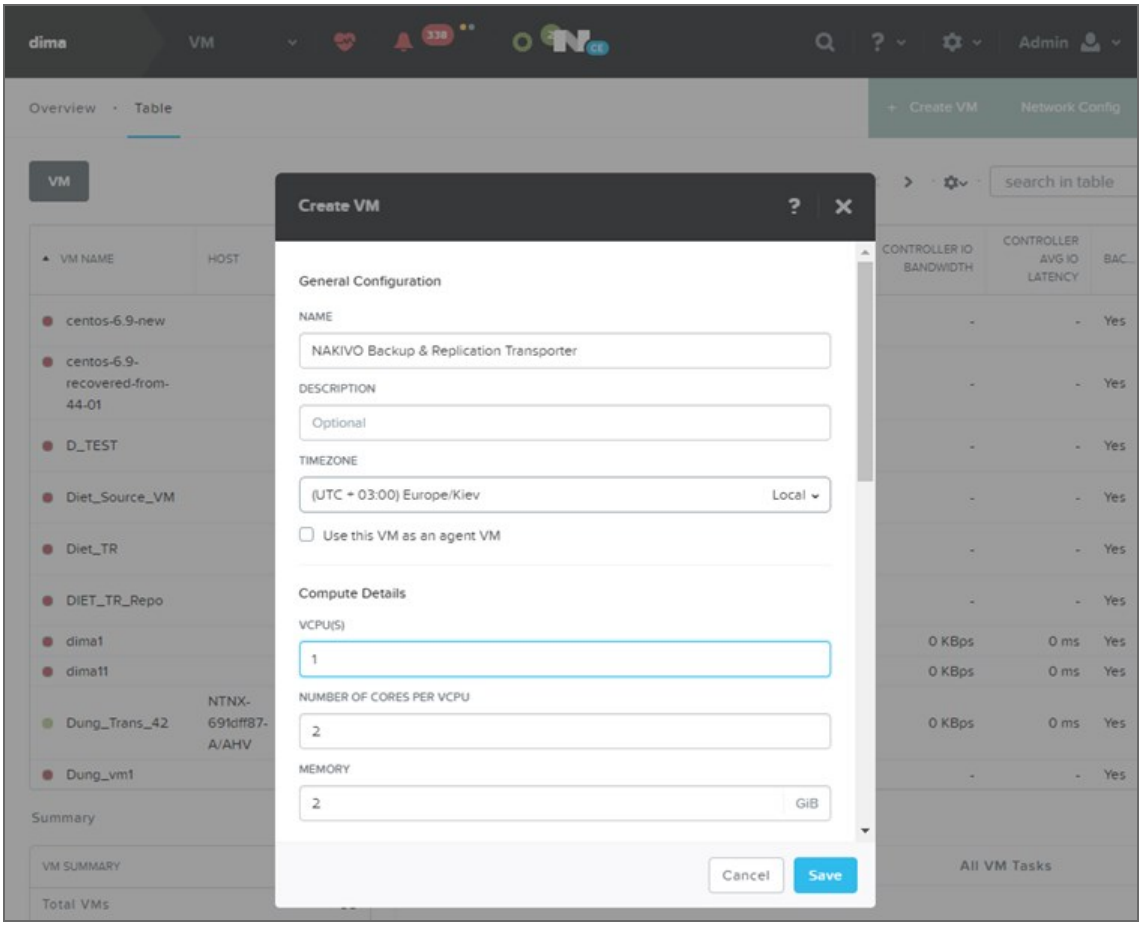

10. Wait until the process of VM creation is complete and locate your newly-created VM on the list.
11. Select your VM and click **Power On**.

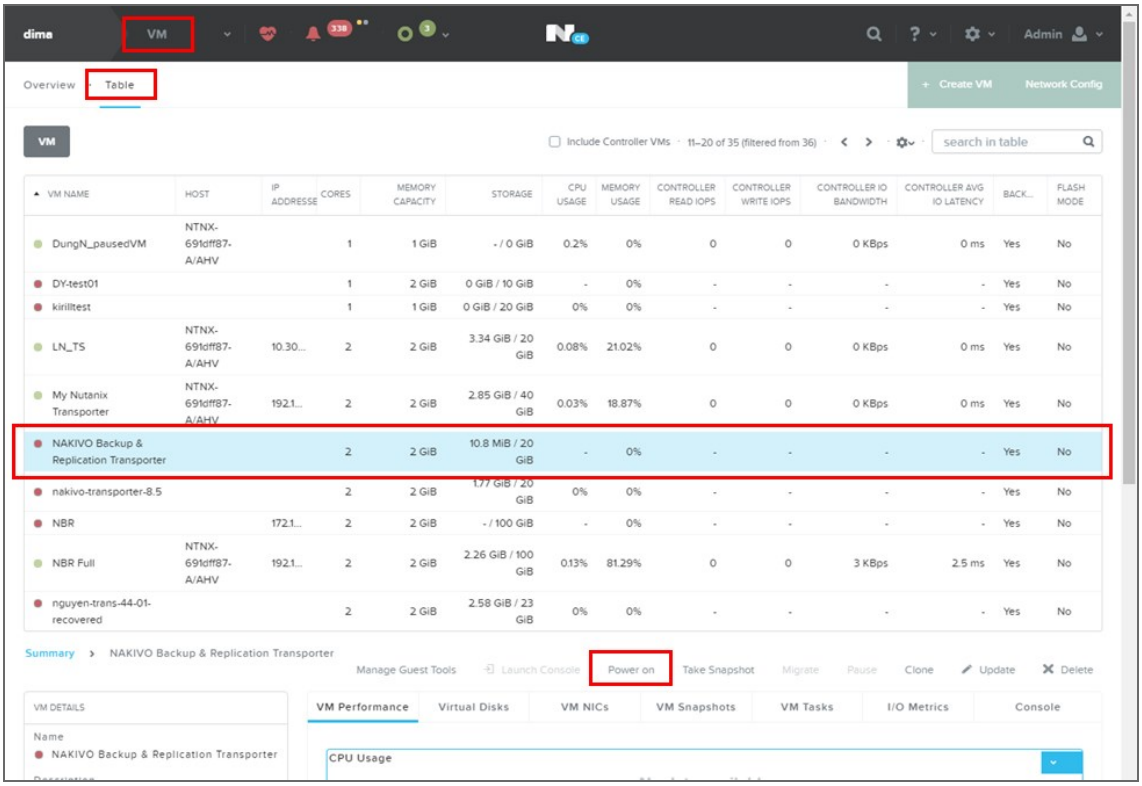

12. After the Virtual Appliance is deployed and powered on, you may need to configure it.

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.

Virtual Appliance OS, Credentials, and Security

The appliance runs Ubuntu 18.04, 64-bit. Use the following credentials to log in to the appliance:

- **Username**: root
- **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.

To enable Backup Immutability for **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
    - Turns on execshield protection
    - 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
- Prevents IP Spoofing
- Installs and configures fail2ban

Web Interface Login

Open the following URL to access the product's web interface: https://Appliance_VM_IP:4443.

Refer to "Getting Started" on page 234 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 fill out our download form, you will get a link to the AWS marketplace page where you can download the AMI.

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 from the Internet.
3. **Security Group**: Either use the "All Traffic" rule or create a set of rules listed below:

Type	Port Range	Source	Description
SSH	22	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 Transporters
Custom TCP	9448-10000	0.0.0.0/0	Enables access to a remote Transporters
All ICMP	0-65535	0.0.0.0/0	Enables access to a remote Transporter

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 the access to the private key file.
   Refer to "Getting Started" on page 234 to understand better how to continue working with NAKIVO Backup & Replication.

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.
   - 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.
     - 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`

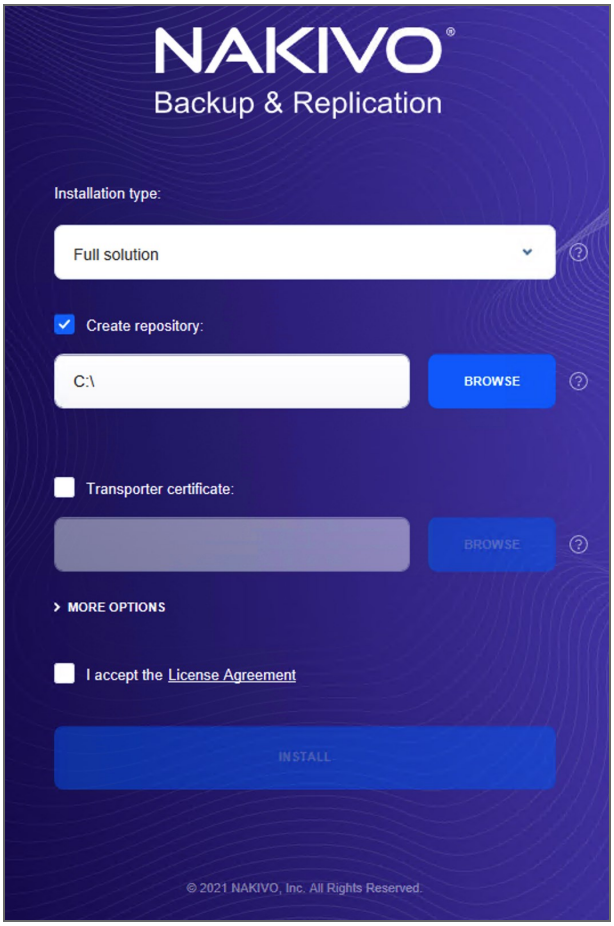

2. Click **MORE OPTIONS** to set up more installation options:
   - **Installation path**: The location where NAKIVO Backup & Replication will be installed. If you want to change the default path to NAKIVO Backup & Replication, click **Browse** and select a new location.
   - **Director Web UI port**: 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.

- **Send daily support bundles during evaluation**: When this option is selected, NAKIVO Backup & Replication automatically creates, encrypts, and uploads support bundles once a day to a NAKIVO support server during the evaluation period. The NAKIVO support team may use this information to improve the product experience and may be able to identify and resolve product issues faster.

3. **I accept the License Agreement**: Select this option to confirm that you have read and agreed to the License Agreement.
4. Click **Install**.

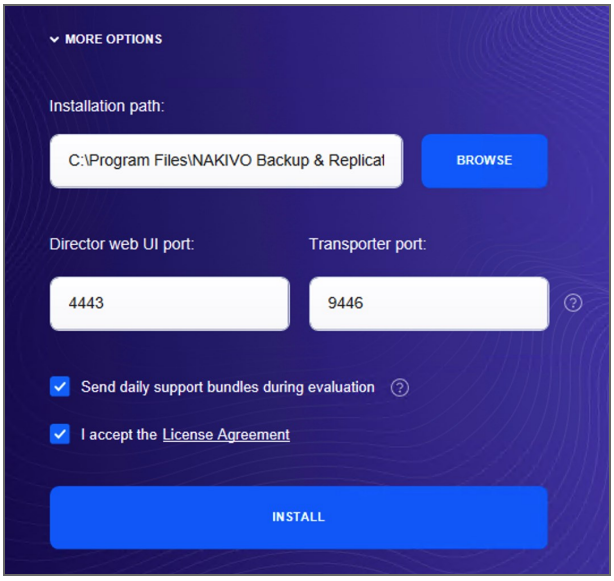

5. Click **Finish** to complete the installation process or **Finish & Open** to complete installation and start using NAKIVO Backup & Replication.

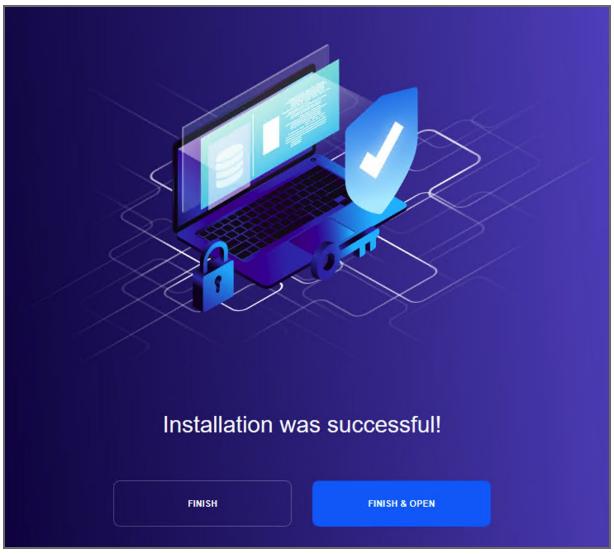

6. To prevent unauthorized access to the product, create your user account. Fore more details, refer to "Logging in to NAKIVO Backup & Replication" on page 235 .

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.

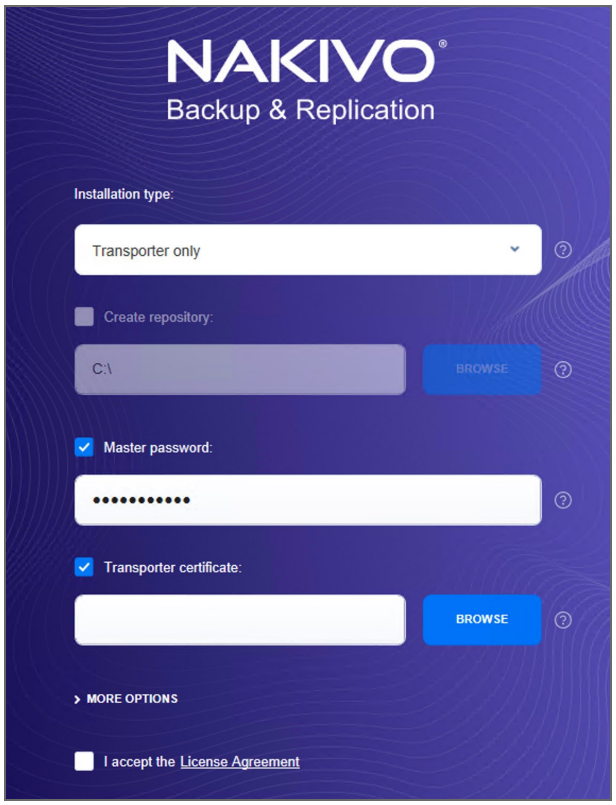

3. Optionally, you can select the **Master password** checkbox and enter the password that will be used to generate a pre-shared key and secure the Transporter.

   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.

4. 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.
   - 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
  ```

5. 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.
   - **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.
6. **I accept the License Agreement**: Select this option to confirm that you have read and agreed to the License Agreement.
7. Click **Install**.

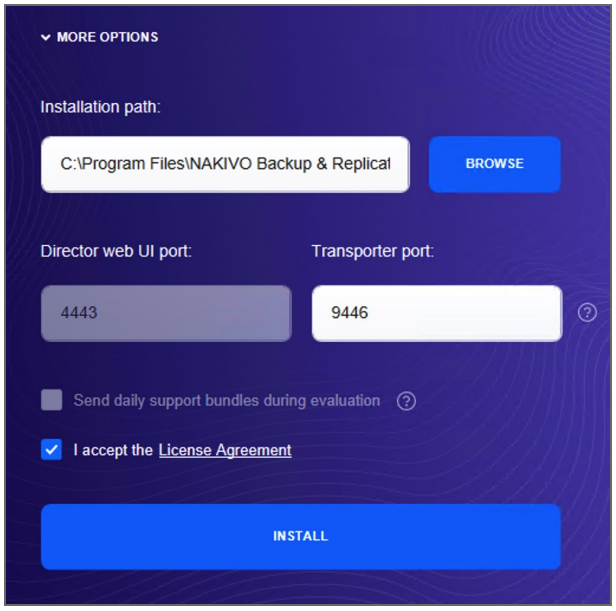

8. When the installation is complete the **Transporter installation was successful** notification appears.

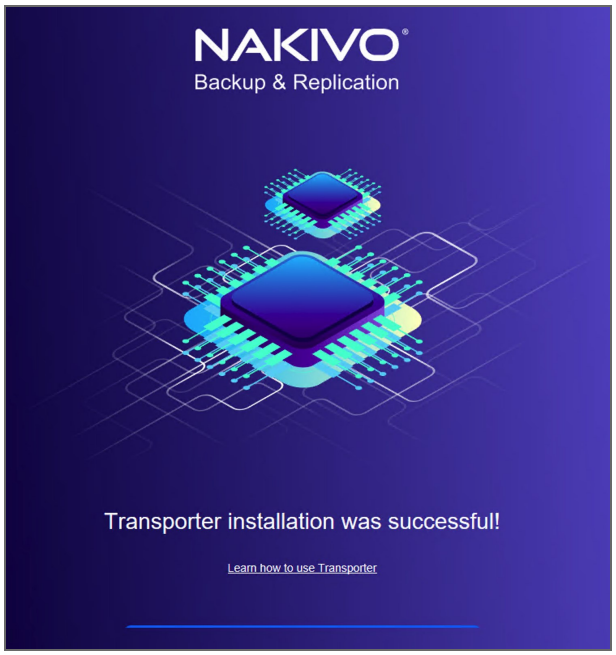

9. Add the Transporter to NAKIVO Backup & Replication.

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. Choose **Multi tenant solution** from the **Installation type** drop-down list.
2. 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.
   - 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
```

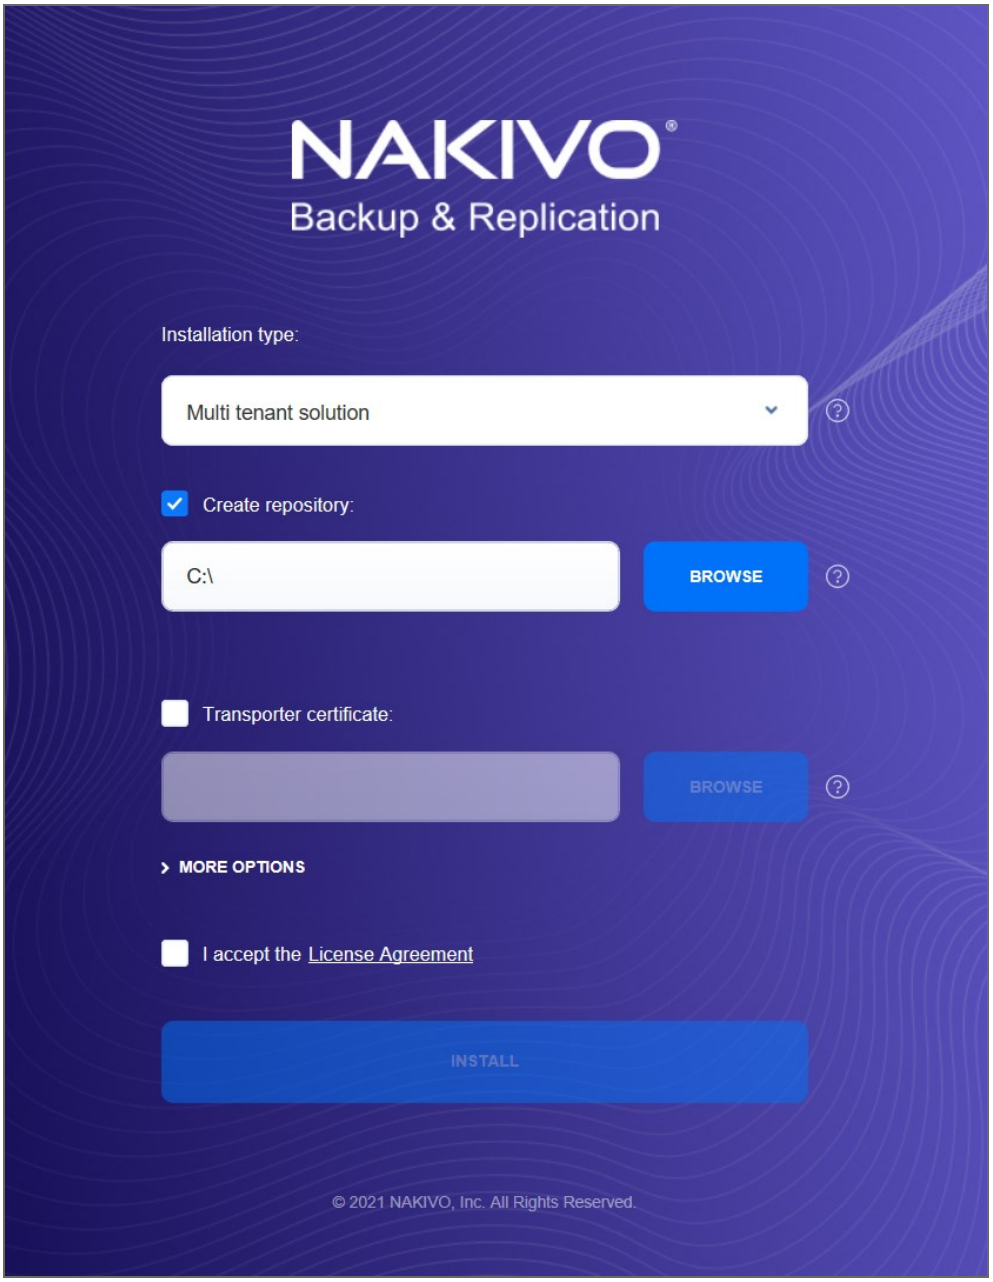

3. Click **MORE OPTIONS** to set up more installation options:
   - **Installation path**: The location where NAKIVO Backup & Replication will be installed. If you want to change the default path to the product, click **Browse** and select a new location.
   - **Director Web UI port**: The default port 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.

- **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.

4. **I accept the License Agreement**: Select this option to confirm that you have read and agreed to the License Agreement.
5. Click **Install**.

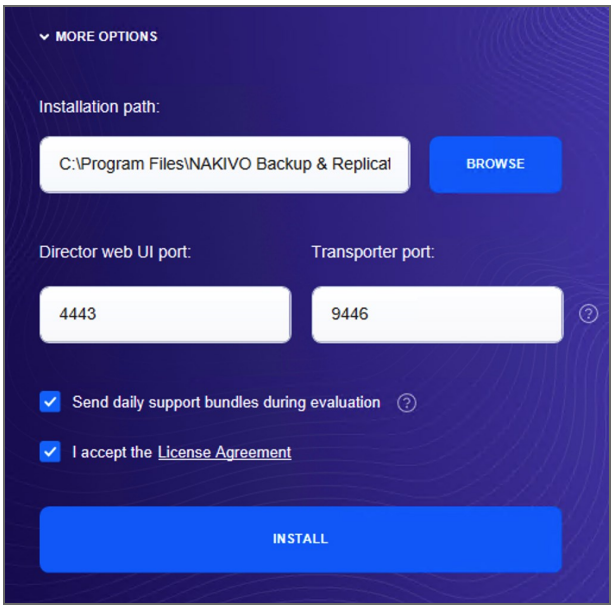

6. Click **Finish** to complete the installation process or **Finish & Open** to complete installation and start using NAKIVO Backup & Replication.

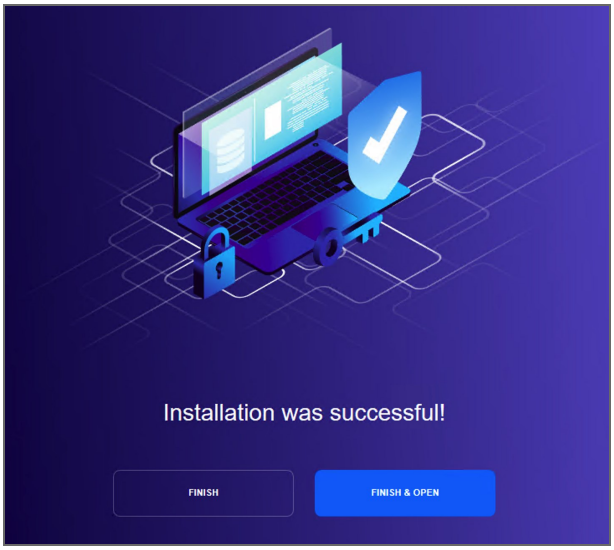

Note

The onboard backup repository for the Master Tenant is automatically created after the installation.

7. Create an account by completing the form. Fore details, refer to "Logging in to NAKIVO Backup & Replication" on page 235 .

Credentials are not required to log in as Master Admin after installation. However, the default credentials are required to log into the product after the first tenant is created. To log in as Master Admin, specify "admin" as the username and leave the password field empty. You can change credentials in the product configuration.

Silent Installation on Windows

You can install NAKIVO Backup & Replication in silent mode via a command line by running the following command: **installer.exe -f --eula-accept** This installs all product components (Director, Transporter, and Backup Repository), and you will be able to use all product features after installation.
The following arguments are available:

Argument	Description
-h	Display the list of available arguments without starting the installation.
--eula-accept, -ea	Indicates that you have read and agree to the End User License Agreement.
-f	Performs the silent installation of the full solution (Director and Transporter).
-t	Performs the silent installation of Transporter only.
-m	Performs the silent installation of the full solution in multi-tenant mode.
-u	Performs the silent update of the installed product components.
--release-notes, -n	Indicates the user has read the release notes for the new release during an update.
-sii	Performs the silent install or update ignoring the single installer instance check.
--ignore-pre-install-action-failures, -ipiaf	All pre-install action failures are ignored.

--cert	Allows to set up a custom Transporter certificate.
--master-pass (short version: **-b**)	Allows to set up a custom master password for the Transporter.

Installing on Linux

- Linux Installation Prerequisites
- Silent Installation on Linux
- Installing Full Solution on Linux
- Installing Transporter on Linux
  - Transporter Installation Prerequisites
  - Transporter Installation
- Installing Full Solution in Multi-Tenant Mode on Linux

Linux Installation Prerequisites

In order to install and use NAKIVO Backup & Replication on a Linux OS, make sure the following 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
-h, -help, help	Display the list of available arguments without starting the installation.
--eula-accept, -ea	Indicates that you have read and agree to the End User License Agreement.
-f	Shall perform the silent installation of the full solution (Director and Transporter).

Argument	Description
-t	Shall perform the silent installation of Transporter only.
-m	Shall perform the silent installation of the full solution in multi-tenant mode.
-u	Shall perform the silent update of the installed product components.
-e	Shall install Transporter on Amazon EC2, or update Transporter installed on Amazon EC2. Refer to Updating on Amazon EC2 for details.
-a	Shall enable uploading support bundles to support team server (Call Home). Refer to System Settings for details.
-y	Shall accept limitations silently.
-i <install_path>	Shall install to the specified installation path.
-d <director_port>	Shall provide a custom Director port.
-p <transporter_port>	Shall provide a custom Transporter port.
-r <port1>-<port2>	Shall provide a custom transporter data ports range.
-C	Shall suppress creating the repository.
-c <repo_path>	Shall create the repository. The **<repo_path>** parameter is optional.
--rt <repo_type>	Shall create a repository of the specified type. The **<repo_type>** parameter may accept the following values: **1** – "Forever incremental with deduplication"; **2** – "Forever incremental without deduplication"; **3** – "Incremental with full backups (deduplication devices)".
--rc <compress_level>	Shall specify the repository compression level. The parameter may accept the following values: Disabled; Fast; Medium; Best. Refer to Creating Backup Repositories for details.
--pnp-cleanup	Shall clean up the database of the device manager for the Linux kernel.
--cert	Allows to set up a custom Transporter certificate.
--**master-pass** (short version: **-b**)	Allows to set up a custom master password for the Transporter.

Installing Full Solution on Linux

Follow the steps below to install all components of NAKIVO Backup & Replication (both Director and Transporter) on a Linux OS:

1. Upload the installer file to the machine on which you want to install NAKIVO Backup & Replication in the *binary transfer mode*. For example:
   - Upload the installer from a Windows-based machine.
   - Upload the product from a Linux-based machine: run the following command: `wget 'server_ip/shared/NAKIVO_Backup_&_Replication_TRIAL.sh'`
2. Log in to the Linux machine and allow the execution of the installer file.
   For example: `chmod +x NAKIVO_Backup_&_Replication_TRIAL.sh`
3. Execute the installer file with root privileges.
   For example: `sudo ./NAKIVO_Backup_&_Replication_TRIAL.sh`
4. Review the license agreement (press **Space** to go to the next page of the agreement). If you agree to the terms of the license agreement, press "Y" and then press **Enter**.
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 running the command on the machine housing the Transporter. Follow these steps:
     - Enter the following command: `bhsvc -b P@ssword123`
     - 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 234 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**.

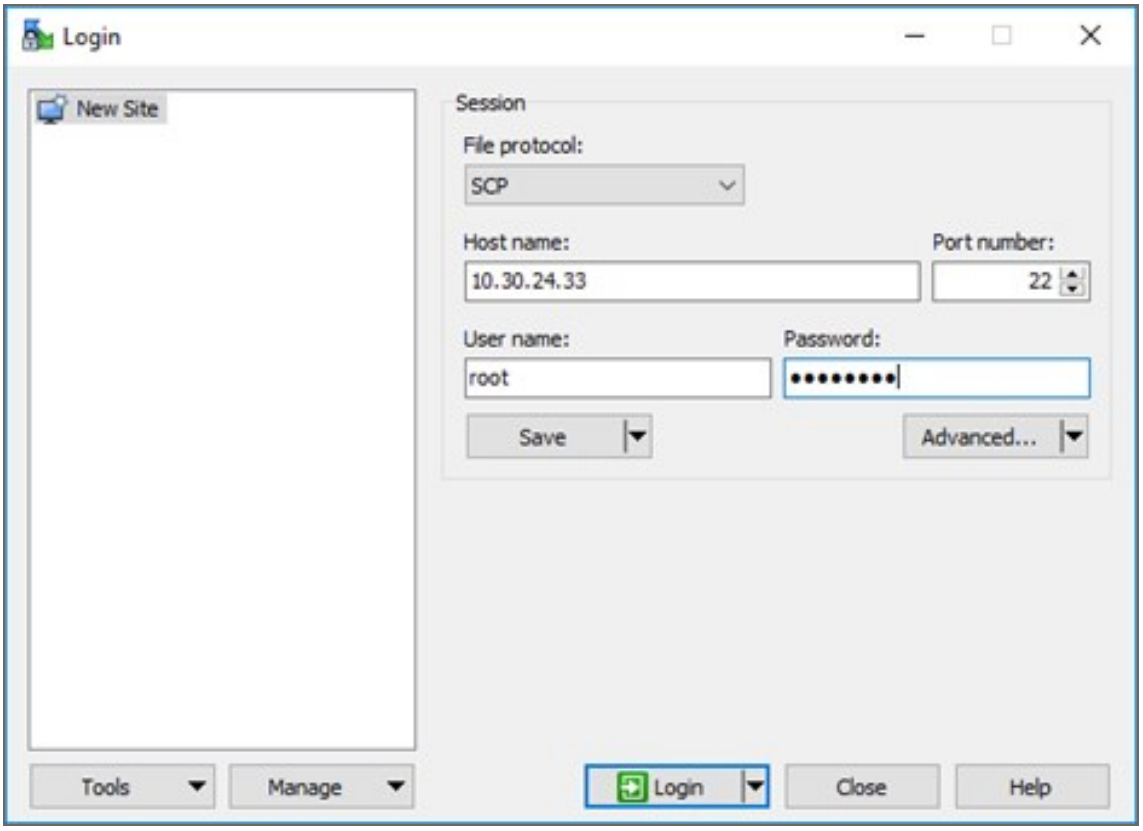

6. Click **Yes** in the dialog box that opens.
7. In the left pane, find the folder that contains the Linux installer, in the right pane, go up to the root folder.
8. Drag and drop the installer from left to the right pane.
9. Choose **Binary** from the **Transfer settings** drop-down list in the Copy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

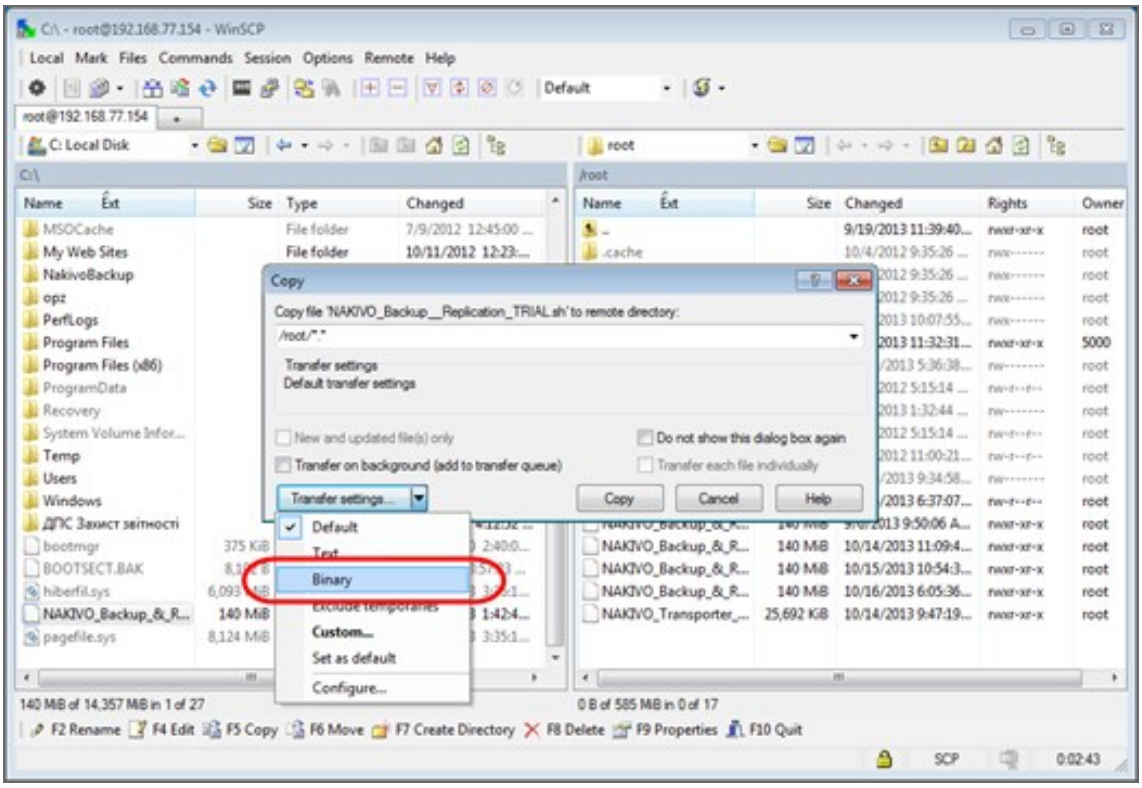

10. Click **Copy**.

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 page 175
- "Installing on Synology NAS Manually" on page 177

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 353.

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.

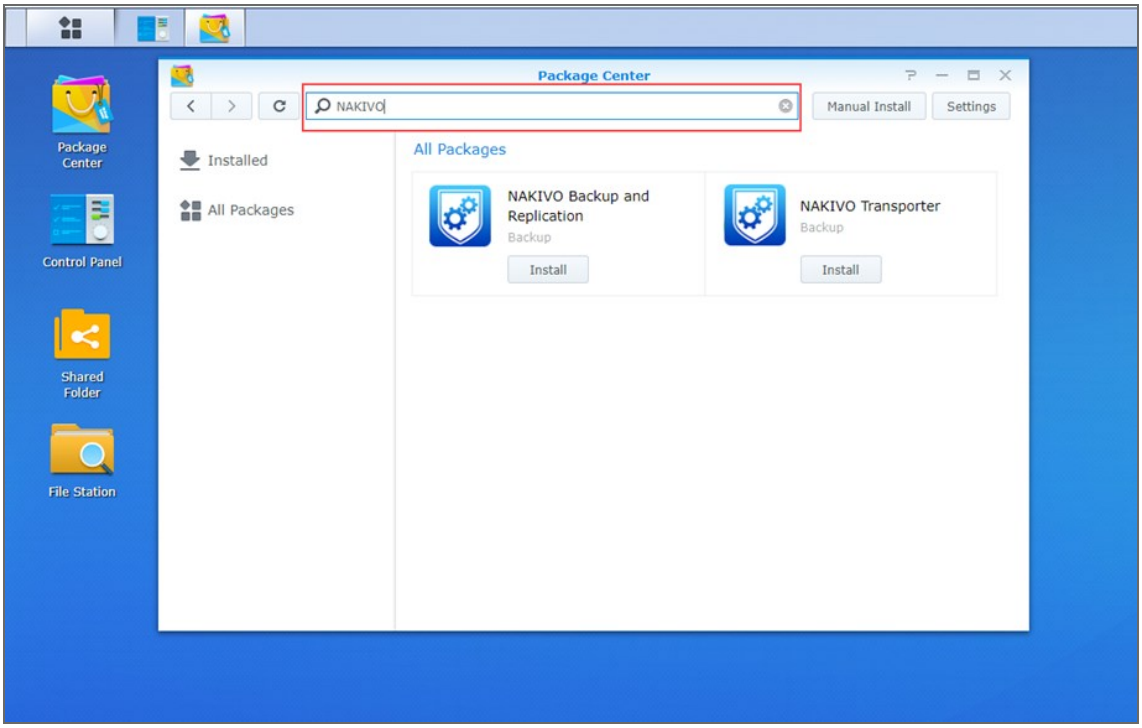

3. Click **Install** on one of the following:
   - **NAKIVO Backup and Replication** to install all product components.
   - **NAKIVO Transporter**to install a Transporter only.

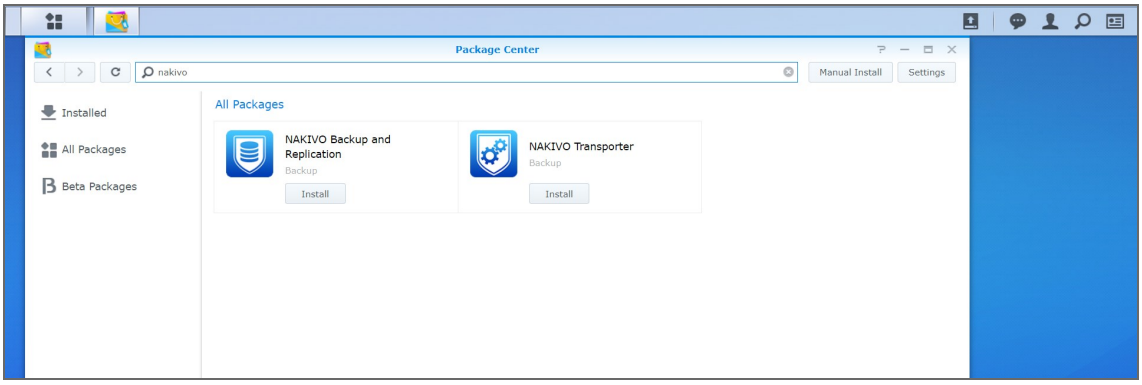

4. Select the **I accept the terms of the license agreement** checkbox and click **Next**.

5. In the **Confirm settings** dialog box, click **Apply**.

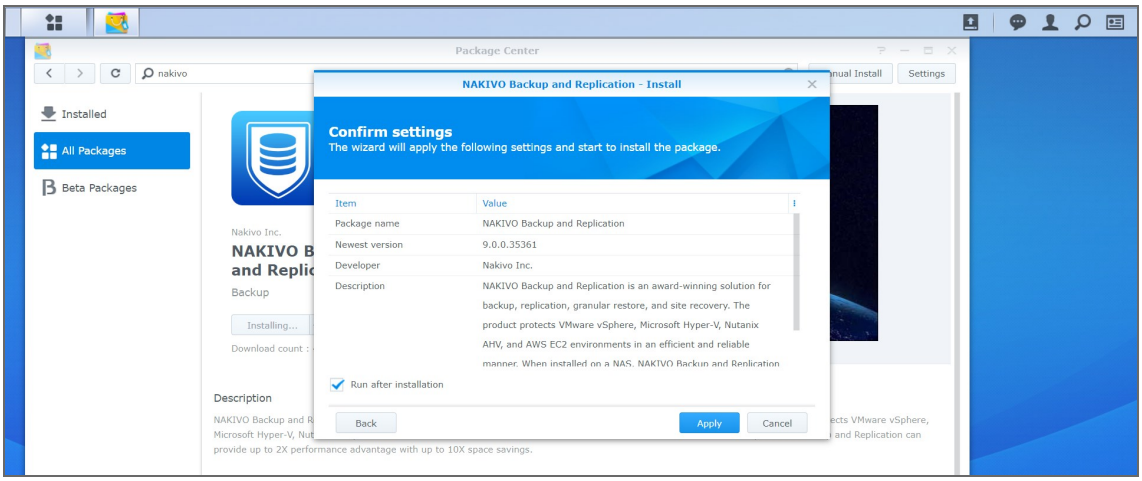

Refer to "Getting Started" on page 234 to better understand how to continue working with NAKIVO Backup & Replication.

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.

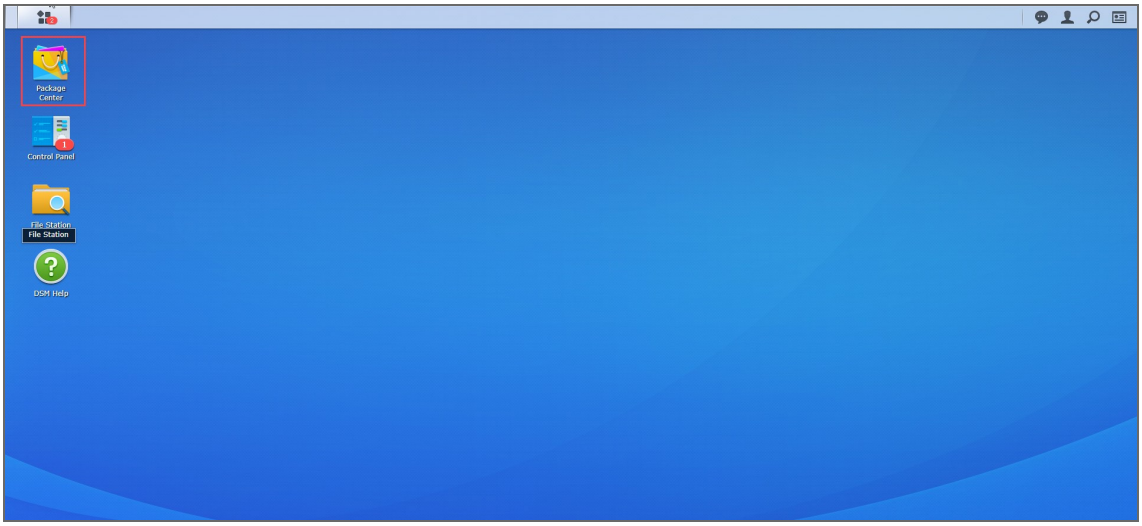

3. Click **Manual Install.**

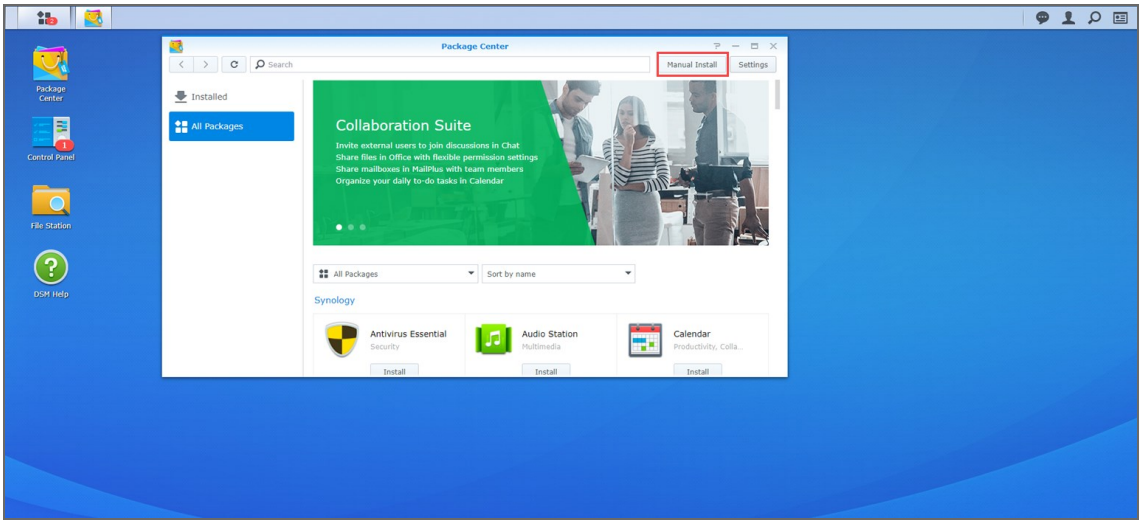

4. Click **Browse**, navigate to the Synology NAS package that you have downloaded, select it, and click **Open**.

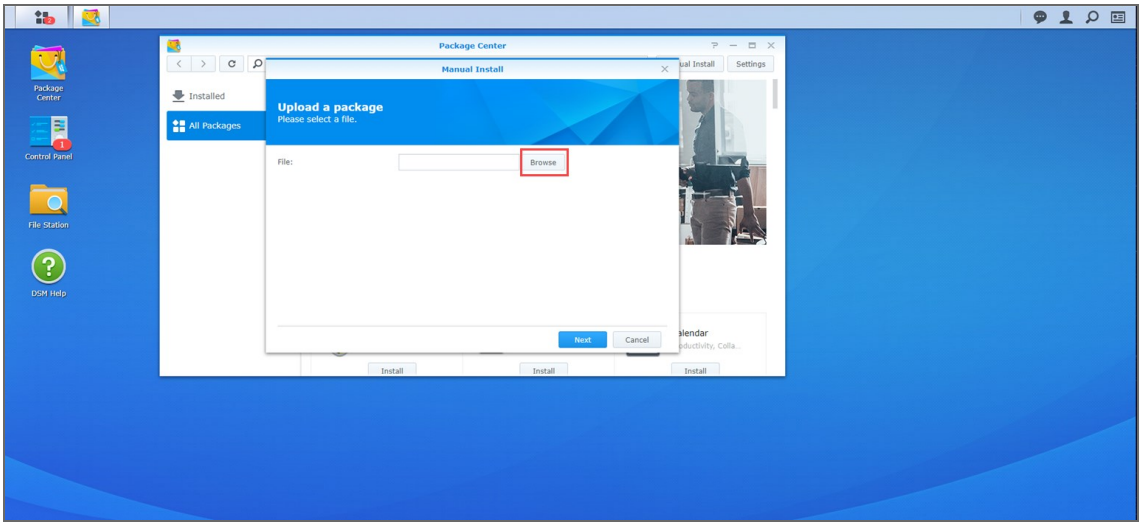

5. Click **Yes** to proceed.

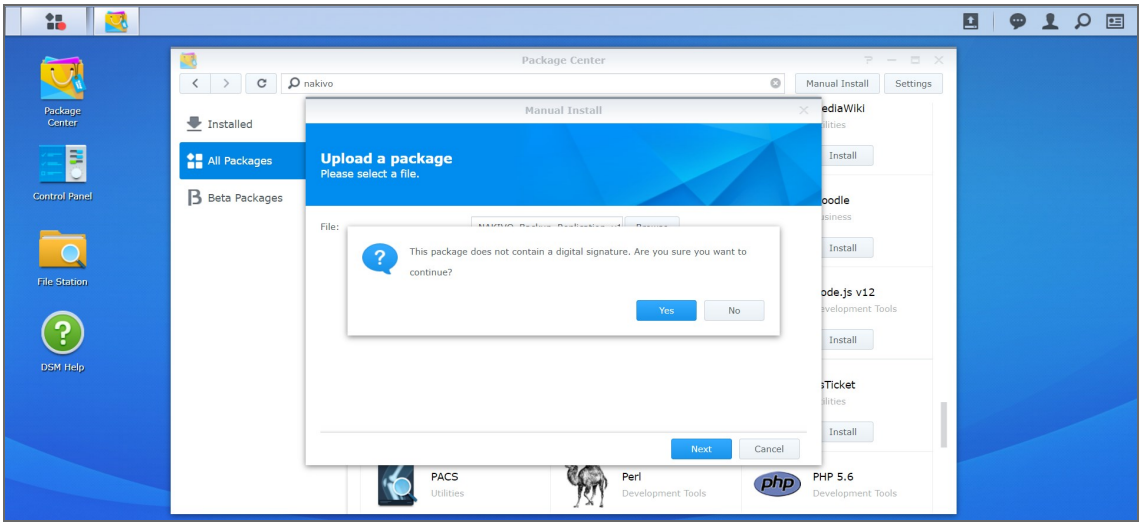

6. After reading through the License Agreement, check **I accept the terms of the license agreement** and click **Next**.

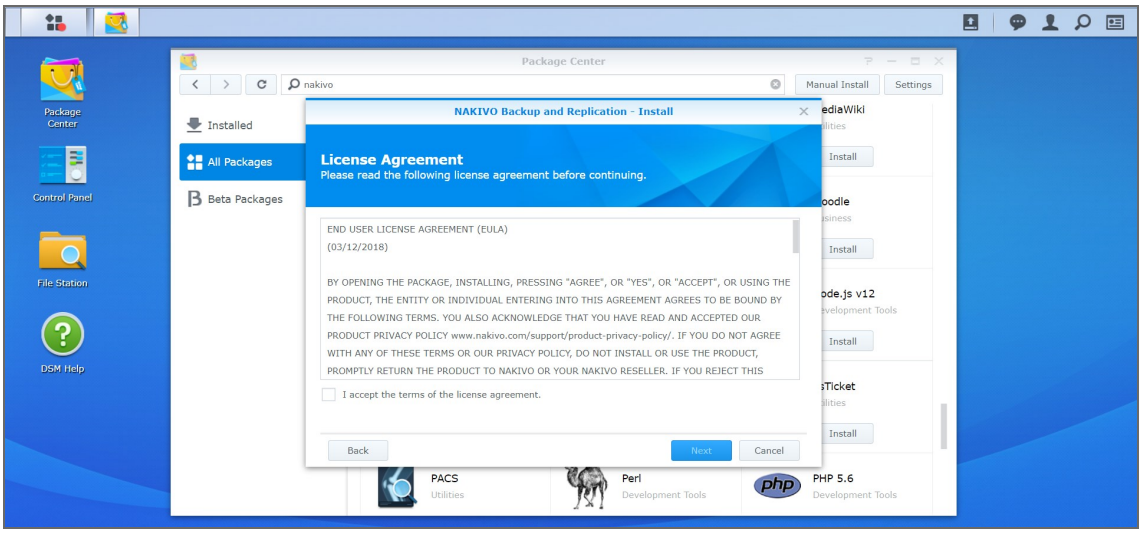

7. Optionally check **Run after installation** to start NAKIVO Backup & Replication immediately after the install process is finished. Click **Apply.**

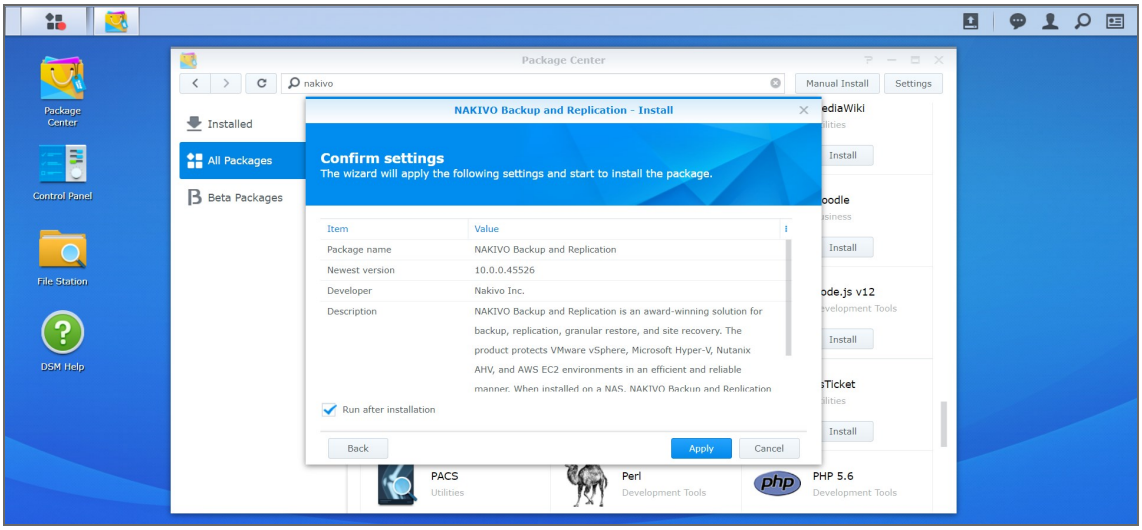

8. Now NAKIVO Backup & Replication is installed on your NAS. To open the NAKIVO Backup & Replication Web interface, go to the following address in your web browser: `https://NAS_IP_address:4443`, or click the NAKIVO Backup & Replication icon in the main menu of the NAS.

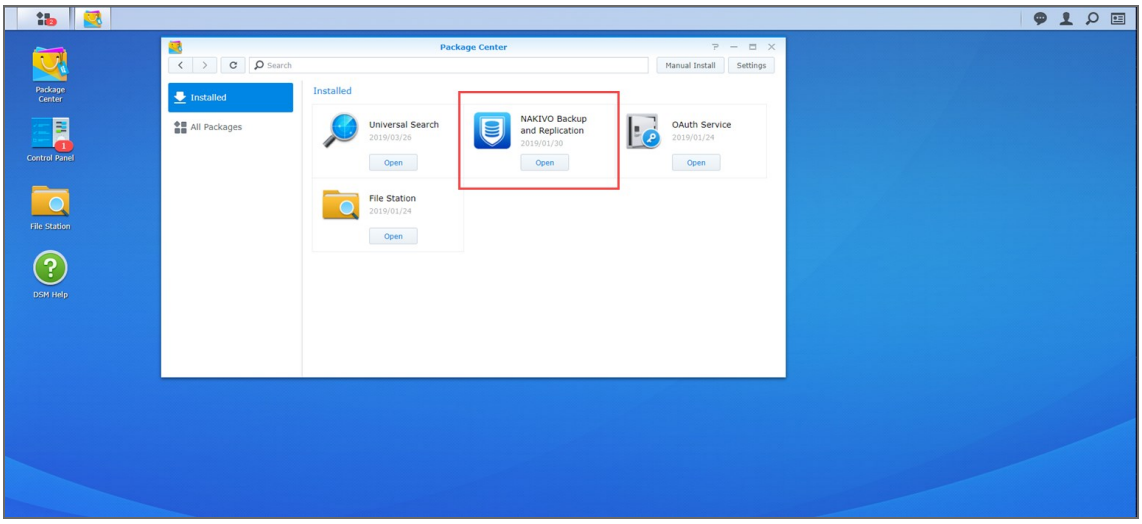

Refer to "Getting Started" on page 234 to better understand how to continue working with NAKIVO Backup & Replication.

Installing on QNAP NAS

You can install a QNAP package with either all NAKIVO Backup & Replication components (Director, Transporter, Backup Repository) or a Transporter only.

NAKIVO Backup & Replication can be installed directly on a supported QNAP NAS to create your own, high-performance backup appliance. With this appliance, all VM data protection components are unified in a single system that is quick to deploy and easy to manage, while also not consuming your environment's valuable resources. Moreover, you are getting an all-in-one backup hardware, backup software, backup storage, and data deduplication in a single box. This results in a zero VMware footprint, less power and cooling, less required maintenance, time, money, and – most of all – higher VM backup performance.

You can install NAKIVO Backup & Replication either via QNAP store or manually.

- "Installing on QNAP NAS via QNAP Store" on page 182
- "Installing on QNAP NAS Manually" on page 184

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 353.

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.

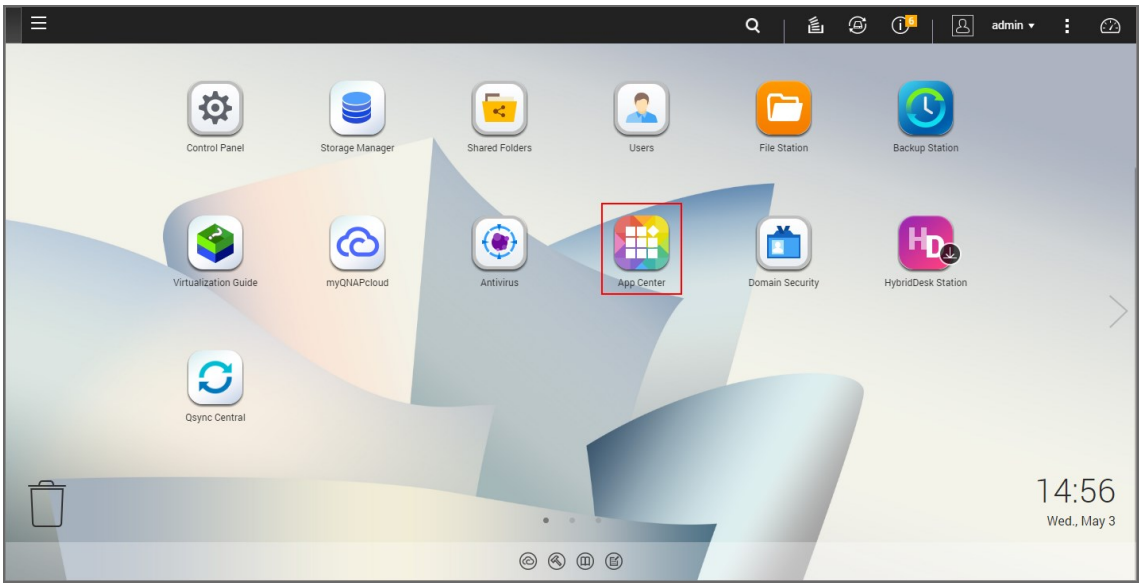

2. Go to **App Center**.
3. Select the **Backup/Sync** category and locate NAKIVO Backup & Replication. Alternatively, you can use the search bar at the top of the App Center window. Click on the magnifying glass icon and enter 'Nakivo'.

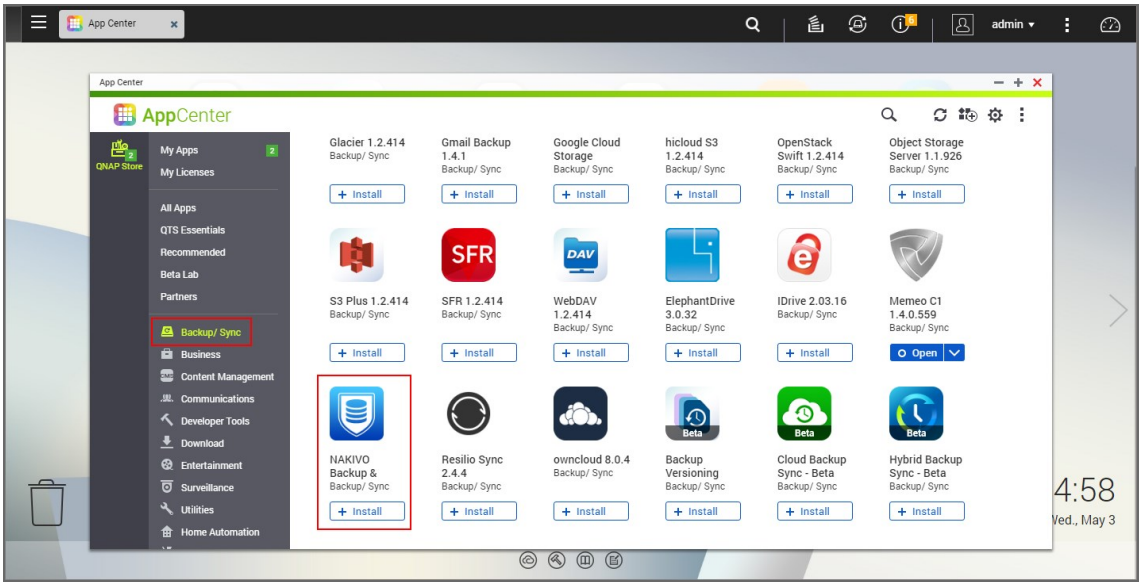

4. Click **Install**.
5. Wait till the installation is completed.

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`.
Refer to "Getting Started" on page 234 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.

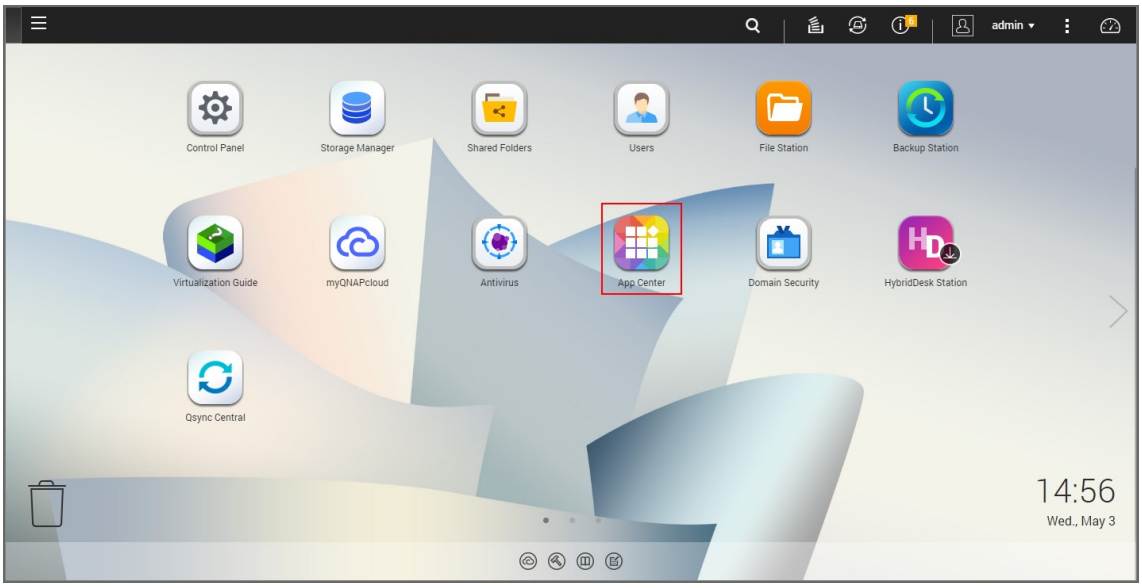

2. Go to **App Center**.
3. Click the **Install Manually** icon.

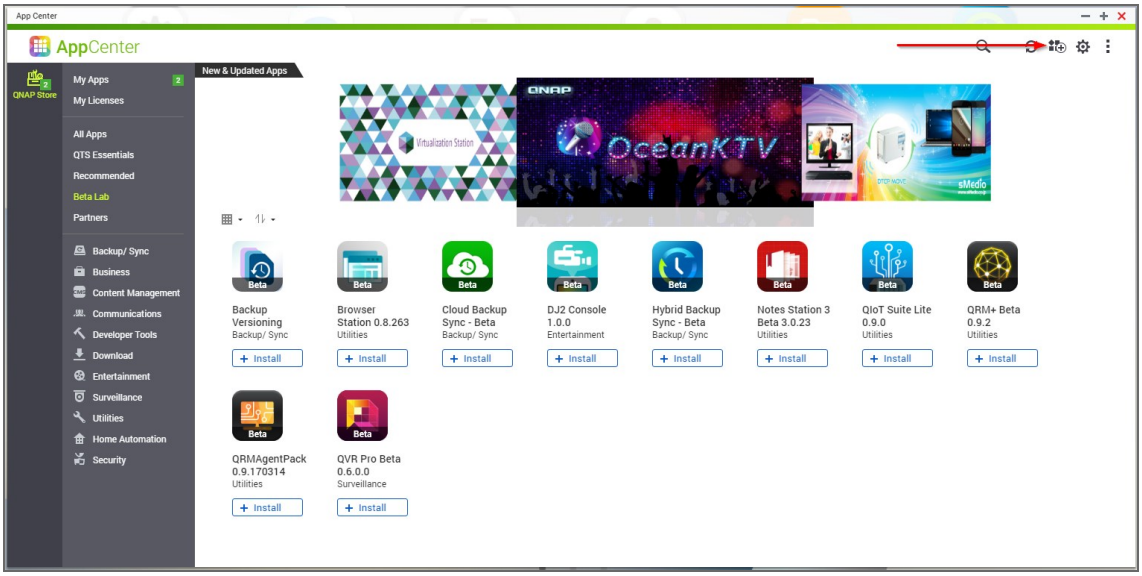

4. Click **Browse** in the window that appears and locate the installer (.qpkg file) on your computer.

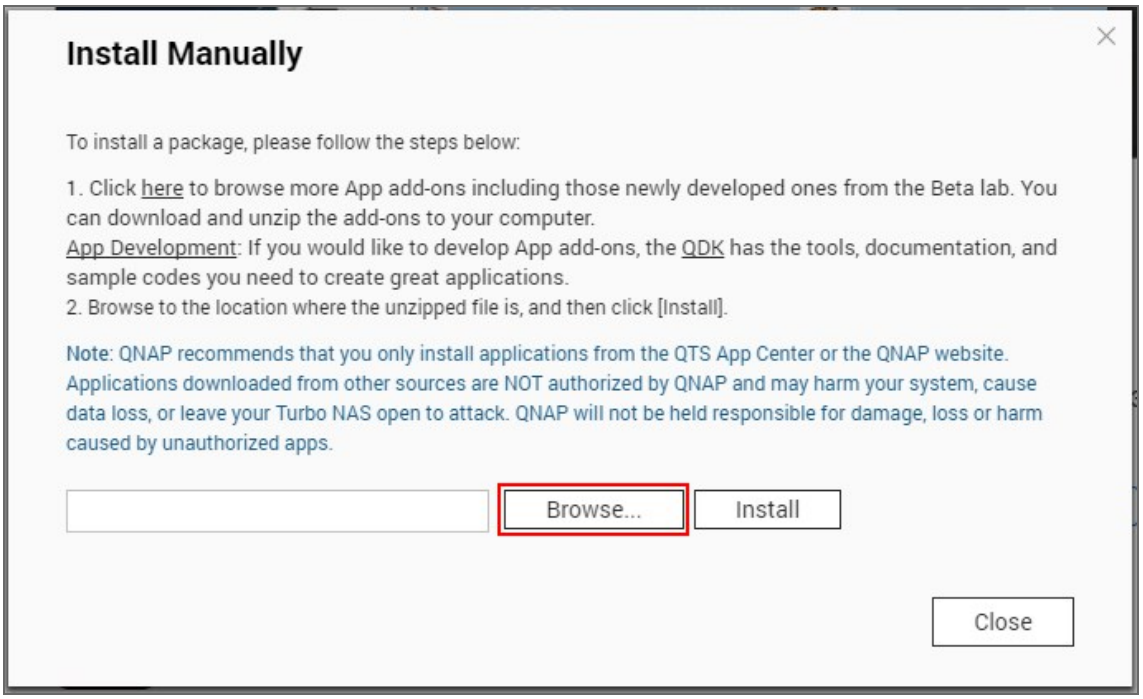

5. Click **Install**.
6. Wait until the installation is complete.

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`.

Refer to "Getting Started" on page 234 to better understand how to continue working with NAKIVO Backup & Replication.

Installing on ASUSTOR NAS

You can install an ASUSTOR package with either all NAKIVO Backup & Replication components (Director, Transporter, Backup Repository) or a Transporter only.

NAKIVO Backup & Replication can be installed directly on a supported ASUSTOR NAS to create your own, high-performance backup appliance. With the appliance, all VM data protection components are unified in a single system that is fast to deploy and easy to manage, while also not consuming your environment's valuable resources. Moreover, you are getting an all-in-one backup hardware, backup software, backup storage, and data deduplication in a single box.

- "Installing on ASUSTOR NAS via App Central" on page 187
- "Installing on ASUSTOR NAS Manually" on page 189

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 353.

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**.

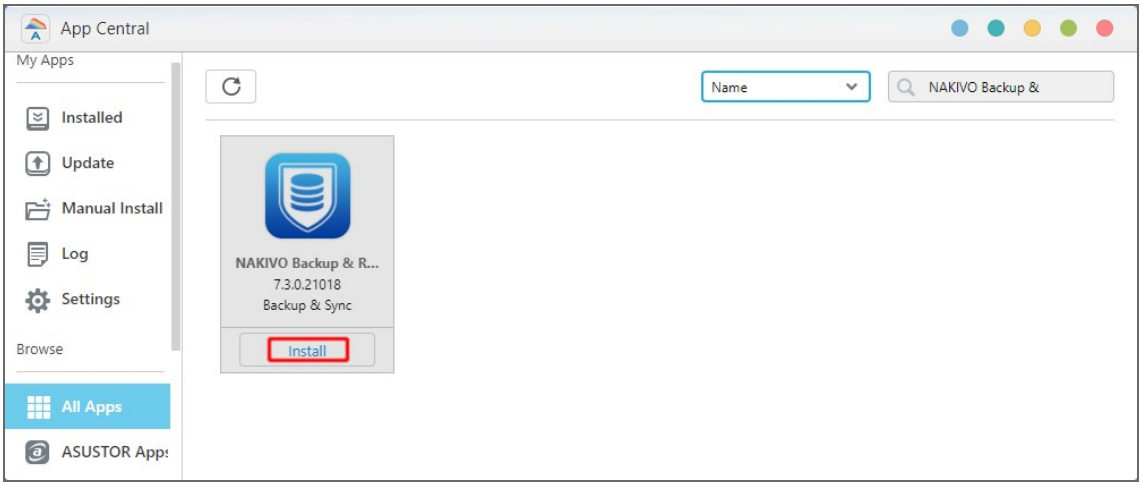

6. In the **About This App** dialog box that opens, select **Enable port forwarding for NAKIVO Backup & Replication** and then click **Install**.

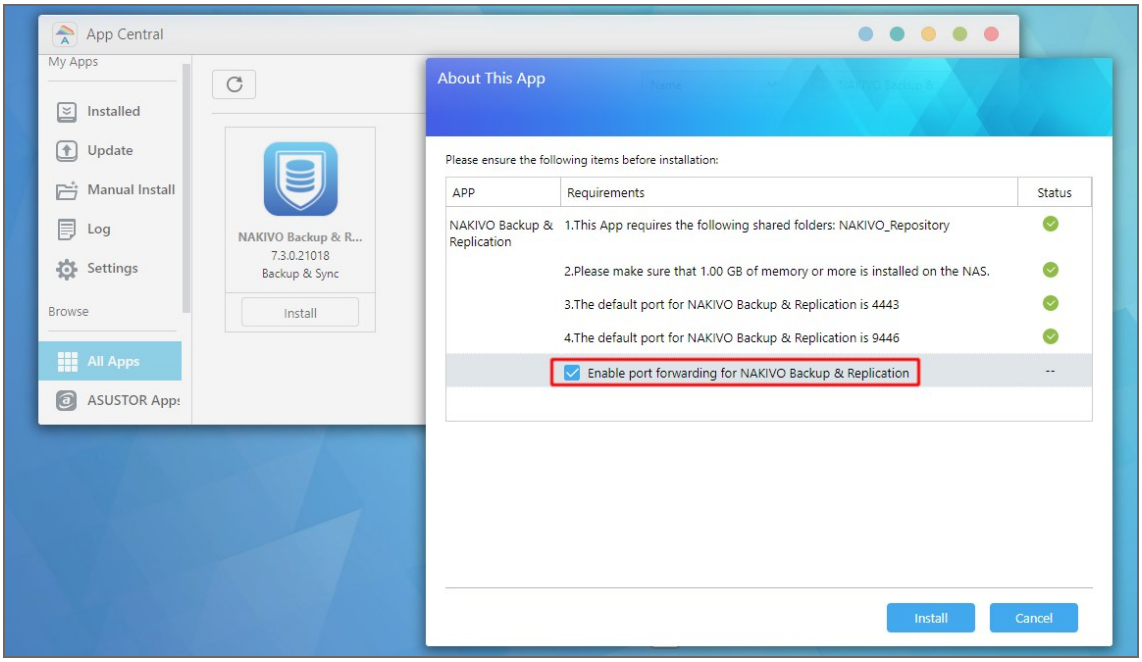

7. Wait until the installation is complete.

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`.

Refer to "Getting Started" on page 234 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**.

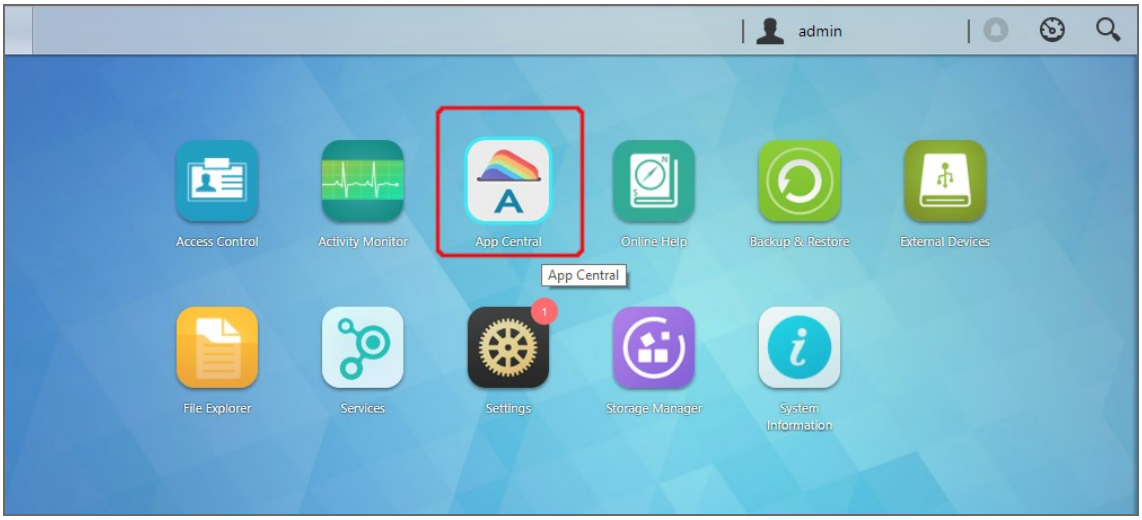

3. Click **Manual Install**.

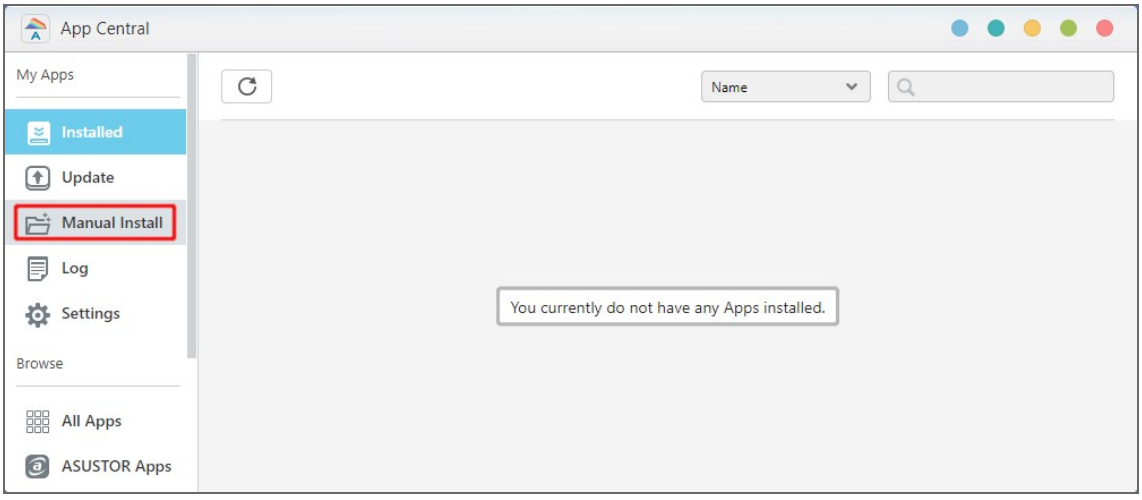

4. Click **Browse**. In the dialog box that opens, locate the installer (.apk file) on your computer.
5. Click **Upload**.
6. In the **About This App** dialog box that opens, check **Enable port forwarding for NAKIVO Backup & Replication**.

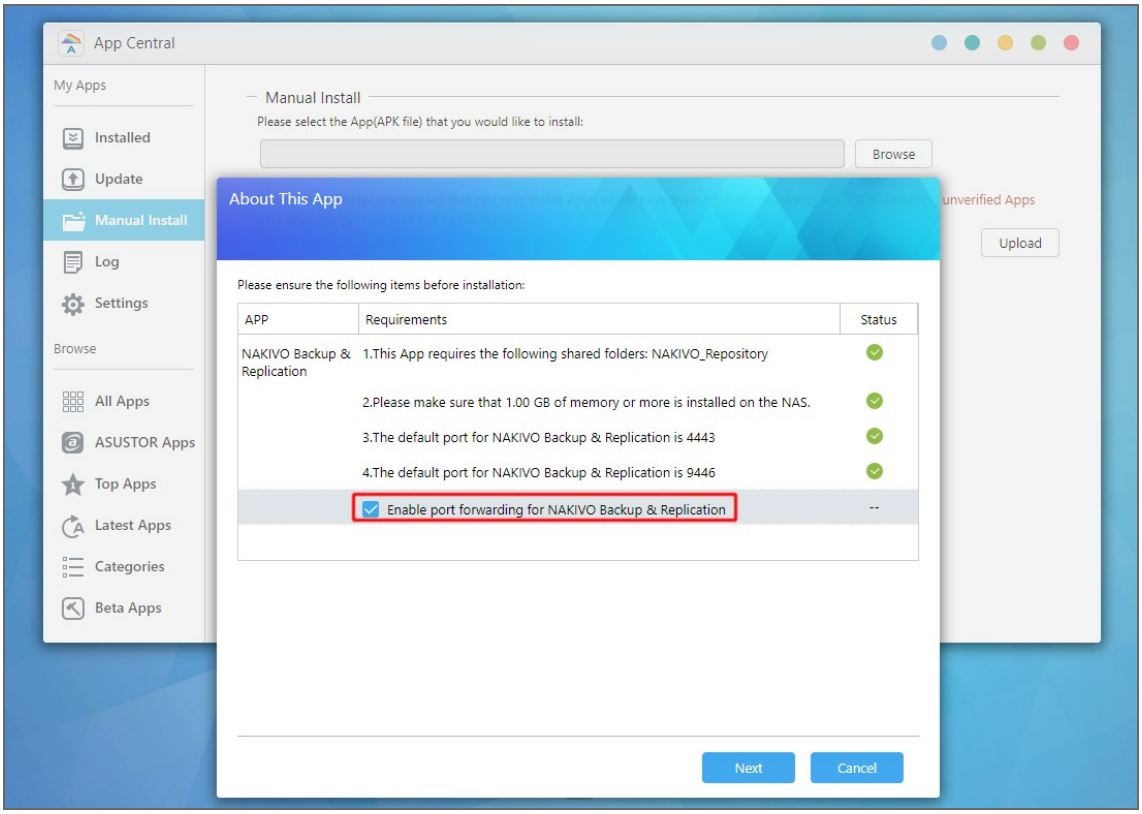

7. Click **Next**.
8. In the warning dialog box that opens, select **I understand the risks associated with installing unverified apps**.
9. Click **Install**.
10. Wait until the installation is complete.

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`.

Refer to "Getting Started" on page 234 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 353.

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.

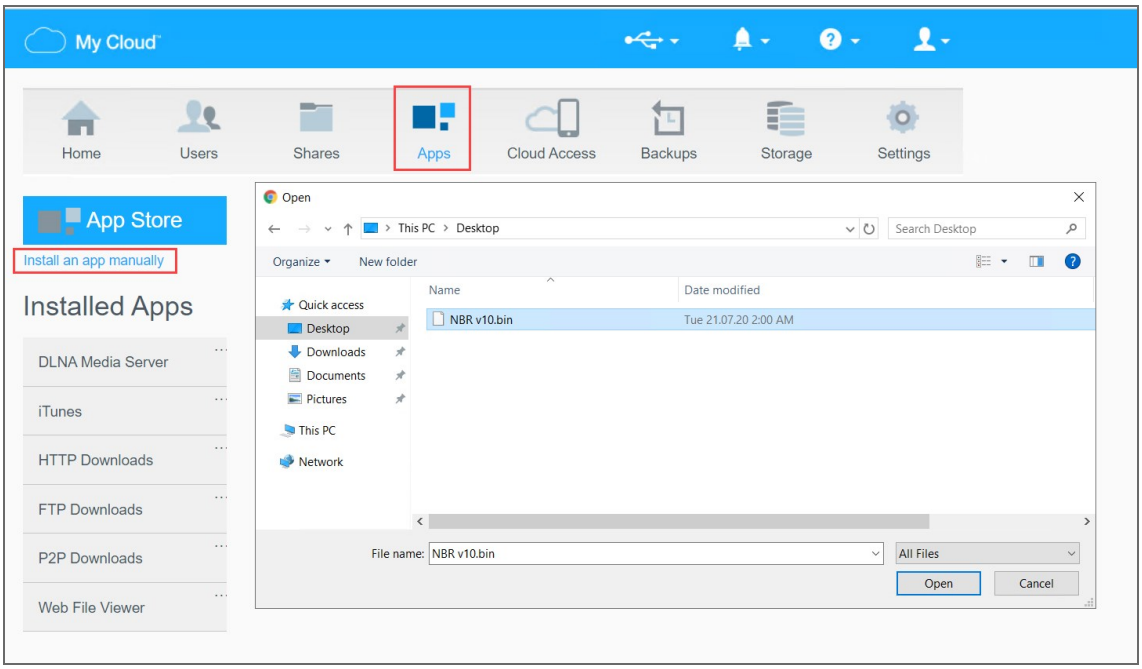

3. In the **File Upload** dialog, navigate to your copy of NAKIVO Backup & Replication installer and click **Open**. The installation progress bar opens.
4. When the installation finishes successfully, a dialog box opens with a message informing you about it. Click **OK** to close the dialog box.

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 234 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” on page 194
- “Installing on NETGEAR ReadyNAS Manually” on page 195

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 353.

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.

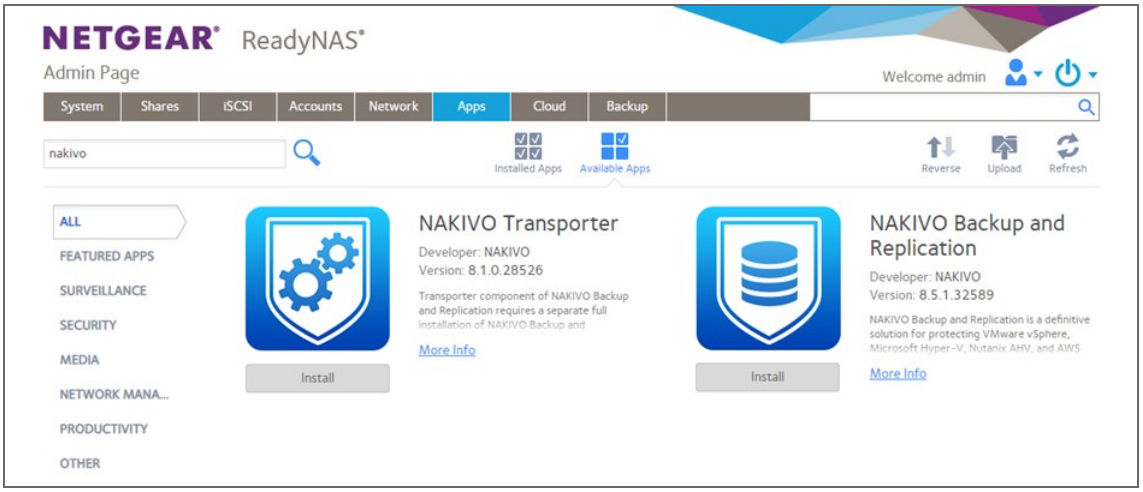

Note

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 234 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**.

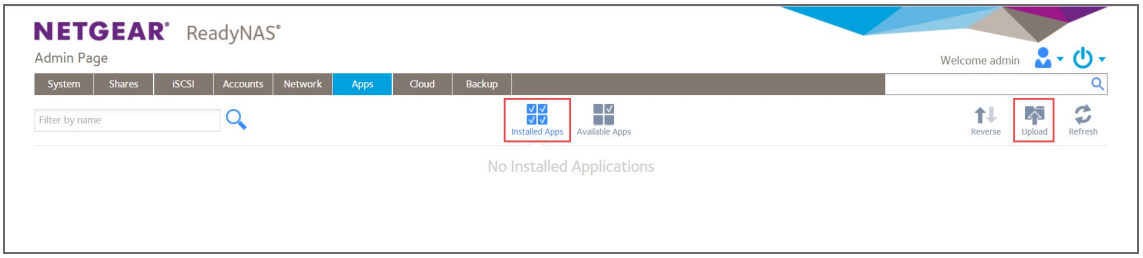

3. The **Install Application** dialog box opens. Click **Browse**.

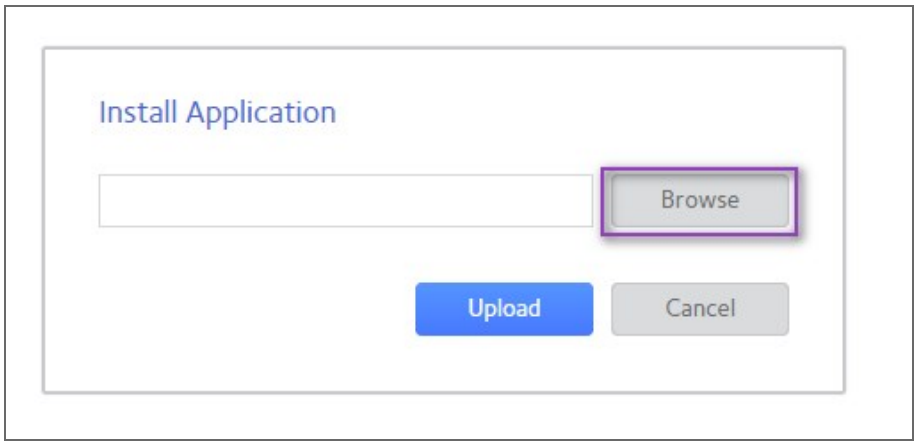

4. In the dialog box that opens, locate the downloaded installer (`.deb` file) and then click **Upload**.
5. Wait until the installation has been completed.

   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 234 to understand better how to continue working with NAKIVO Backup & Replication.

Installing on FreeNAS

Make sure the following prerequisites are met:

1. You have access to the FreeNAS system.
2. Your FreeNAS system meets system requirements for installing NAKIVO Backup & Replication.
3. The `iocage` jail/container manager is installed on your FreeNAS system. Refer to the `iocage` README page for a description.
4. A storage pool is created on your FreeNAS 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 system:

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

     Note

     If a utility for downloading files like `wget` or `curl` is missing on your FreeNAS 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 system.
   - For the full NAKIVO Backup & Replication installation on a FreeNAS v11.3:
     ```
     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 v11.3:
     ```
     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 353 for details.

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.

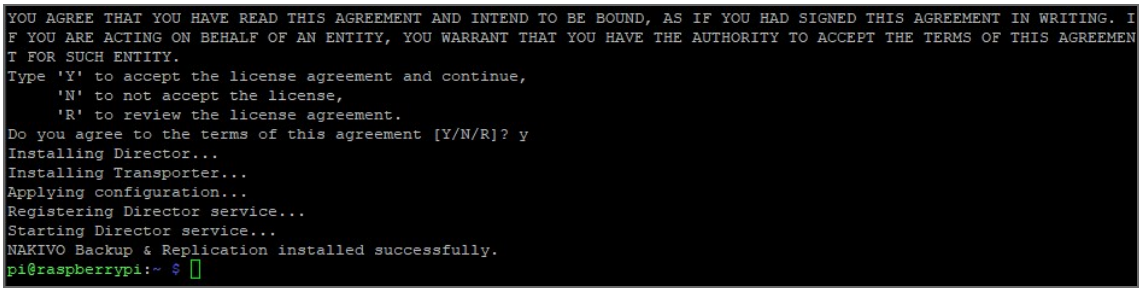

Refer to "Getting Started" on page 234 to better understand how to continue working with NAKIVO Backup & Replication.

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 99 .
- For the installation procedure, refer to "Installing on Generic ARM-based Device" on page 197 .

Refer to "Getting Started" on page 234 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:

- "Auto Updating NAKIVO Backup and Replication" on page 200
- "Updating Virtual Appliance" on page 204
- "Updating on Windows" on page 209
- "Updating on Linux" on page 211
- "Updating on Synology NAS" on page 212
- "Updating on Western Digital NAS" on page 214
- "Updating on Amazon EC2" on page 215
- "Updating on QNAP NAS" on page 221
- "Updating on ASUSTOR NAS" on page 224
- "Updating on NETGEAR ReadyNAS" on page 226
- "Updating on FreeNAS" on page 228
- "Updating on Generic ARM-based Device" on page 229

Auto Updating NAKIVO Backup and Replication

- Download & Update Option
- Download Option
- Forgetting Update

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. Once the update becomes available, the **Update available** notification appears in the main menu of the product. You can choose to either download and update the product immediately or download the update and run it at a later time.

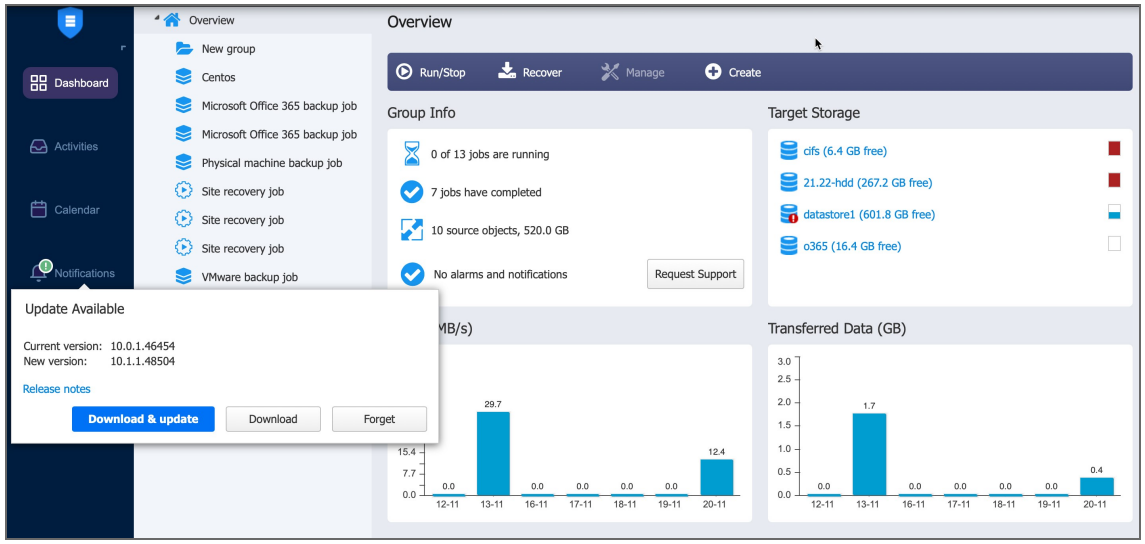

Note

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 in a single click, do the following:

1. Click the **Update Available** button.
2. Optionally, click **Release Notes** to see features and improvements implemented in the new product version.
3. In the **Update Available** dialog box, click **Download & update**.

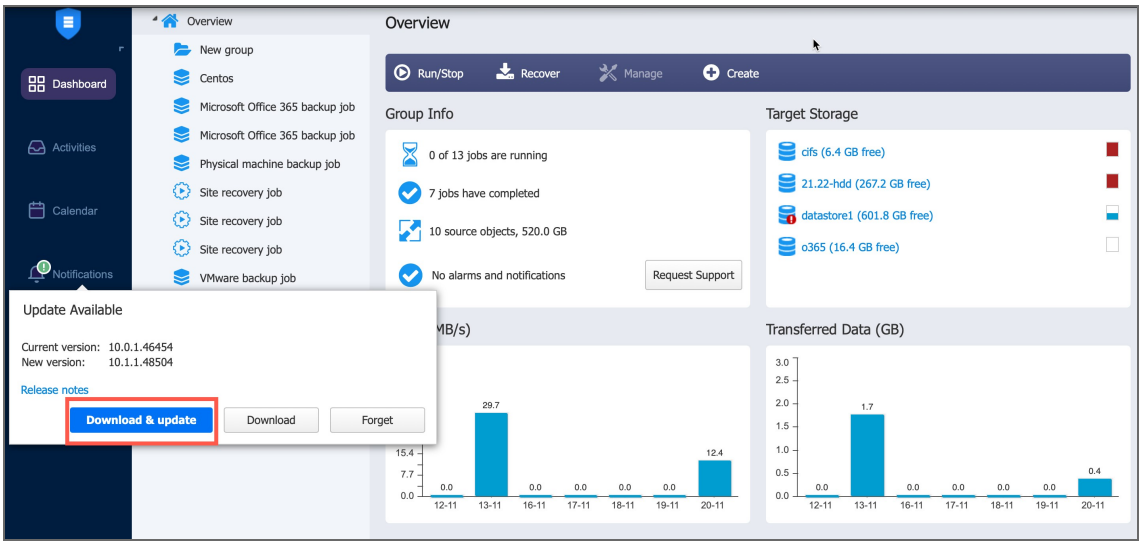

4. In the **Update Available** dialog box, click **Proceed** to confirm stopping all current activities and start downloading the update. When the download is complete, the product updating process will begin.

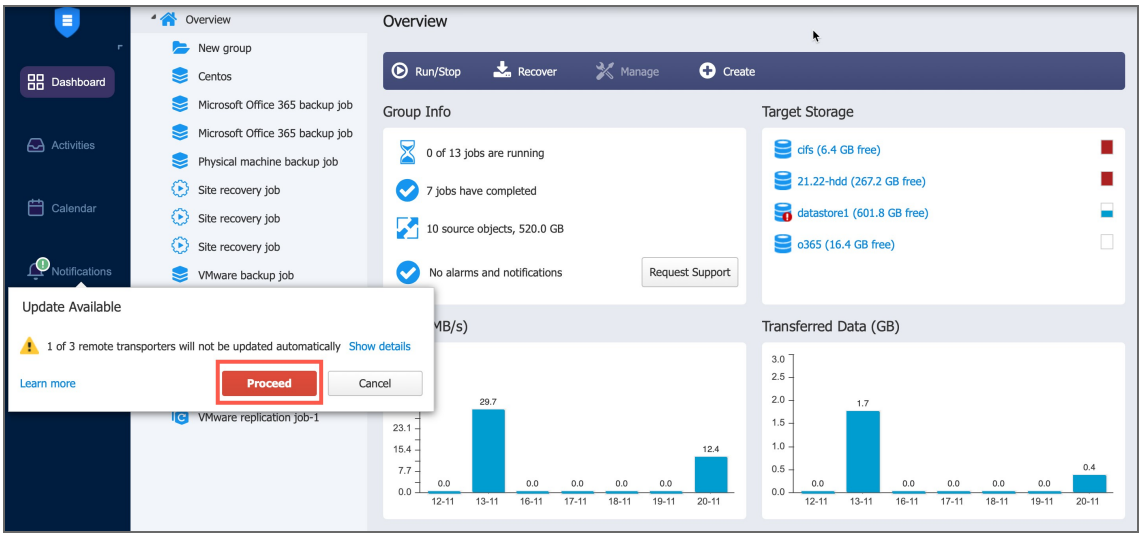

The product will download the update to the Director first. When the Director is updated, the update will be downloaded to the Transporters that in turn will be updated simultaneously. If some Transporters are not updated, you can update them outside the product. Refer to the corresponding articles for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

5. If the full solution is updated successfully, the **Update Complete** dialog box opens. Click **Close** to dismiss the dialog box.

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. Click the **Update Available** button in the product interface.
2. In the **Update Available** dialog box, click **Download** to start downloading the update.

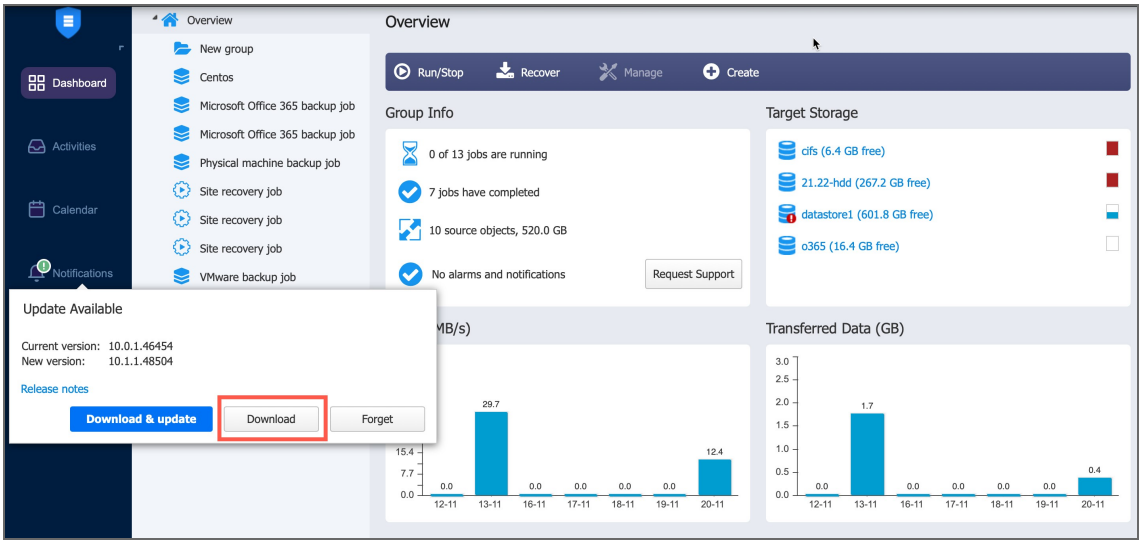

3. When downloading is complete, the **Update Ready** dialog box appears in **Notifications**.
4. Do either of the following:
   - Click **Update Now** if you want to start the updating process. In the confirmation dialog box that opens, click **Proceed**. Updating the product will stop all current activities, including running jobs, recovery jobs, repository maintenance, etc.

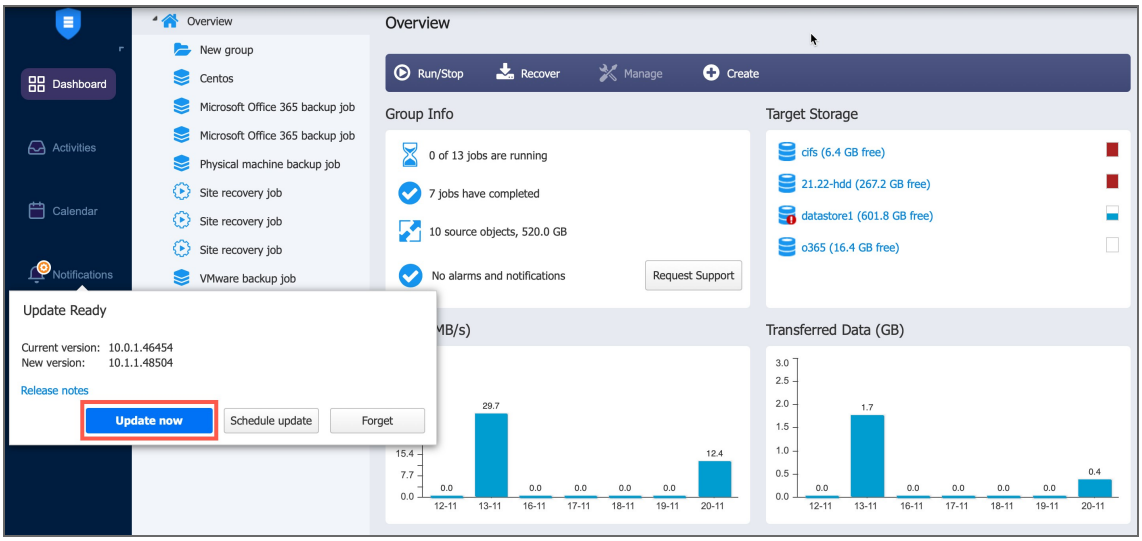

- Click **Schedule Update** to update the solution on a schedule:
  1. In the dialog box that opens, pick a day and time for updating. Click **Apply**.

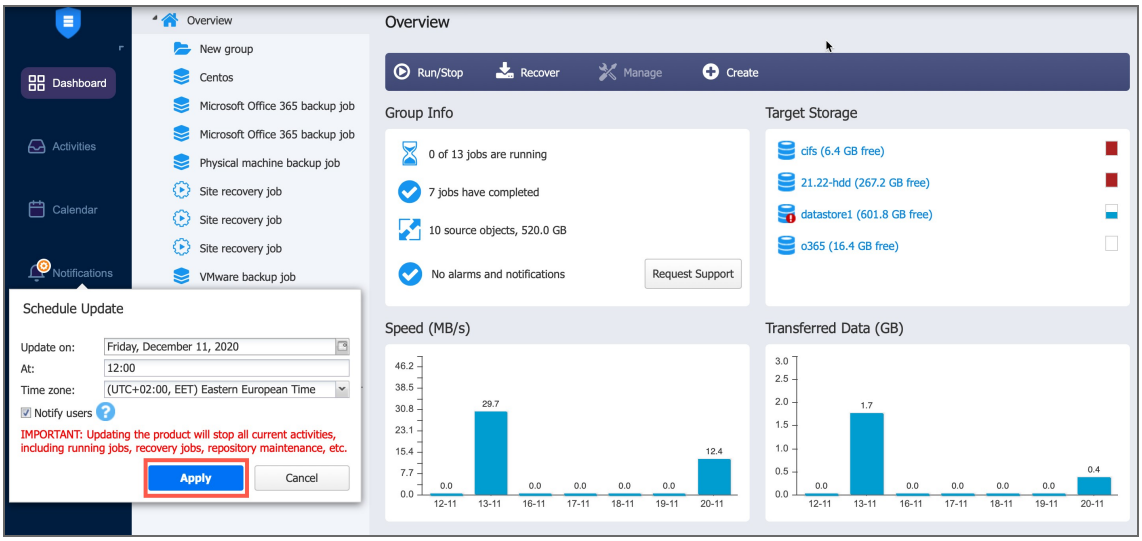

  2. On a working day before the scheduled update, you will see the notification in the product menu with the **Update Reminder** dialog box. Do any of the following:
     a. Click **Reschedule** if you want to reschedule the update and pick a different time.
     b. Click **Cancel update** to cancel updating of the full solution.

        Note

        A notification about the update will also be sent to your email if email settings are configured.

Forgetting Update

In the **Update Available** dialog box, you can click **Forget** to dismiss all notifications. If you select this option, you will not receive a notification regarding a product update until the next update is available.

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.

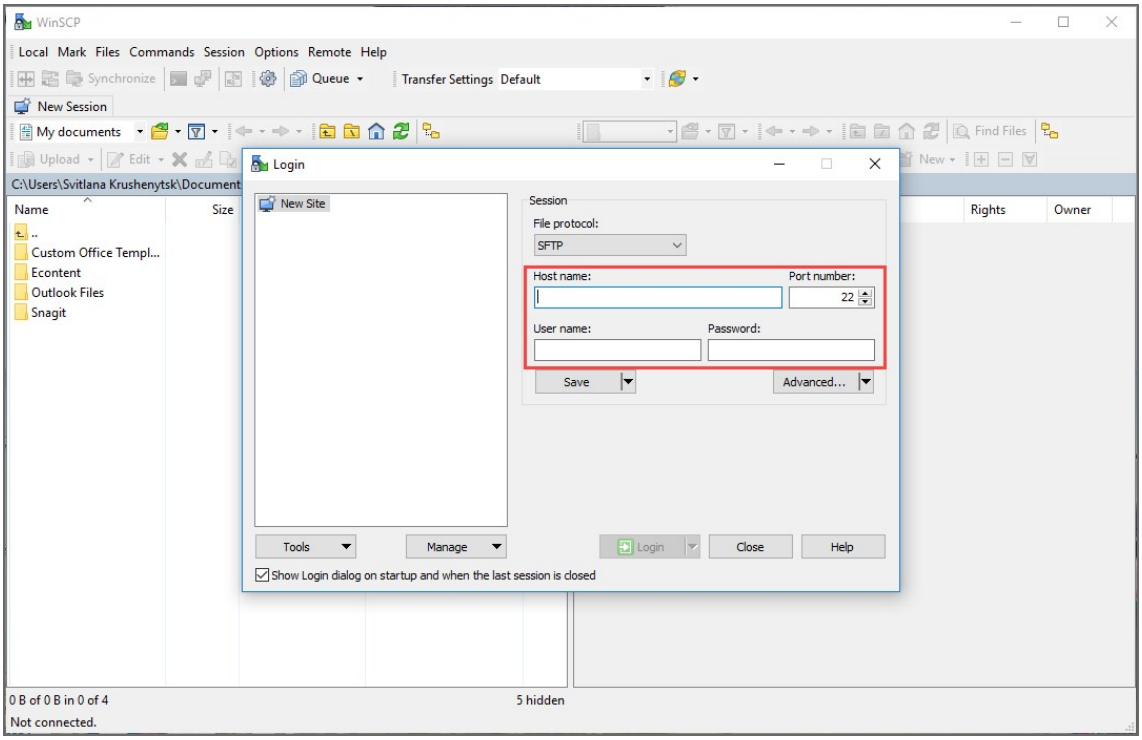

2. Download the latest *VA and Linux updater* from www.nakivo.com/resources/download/update/.

3. Change the directory to `/opt/nakivo/updates` and locate the updater.

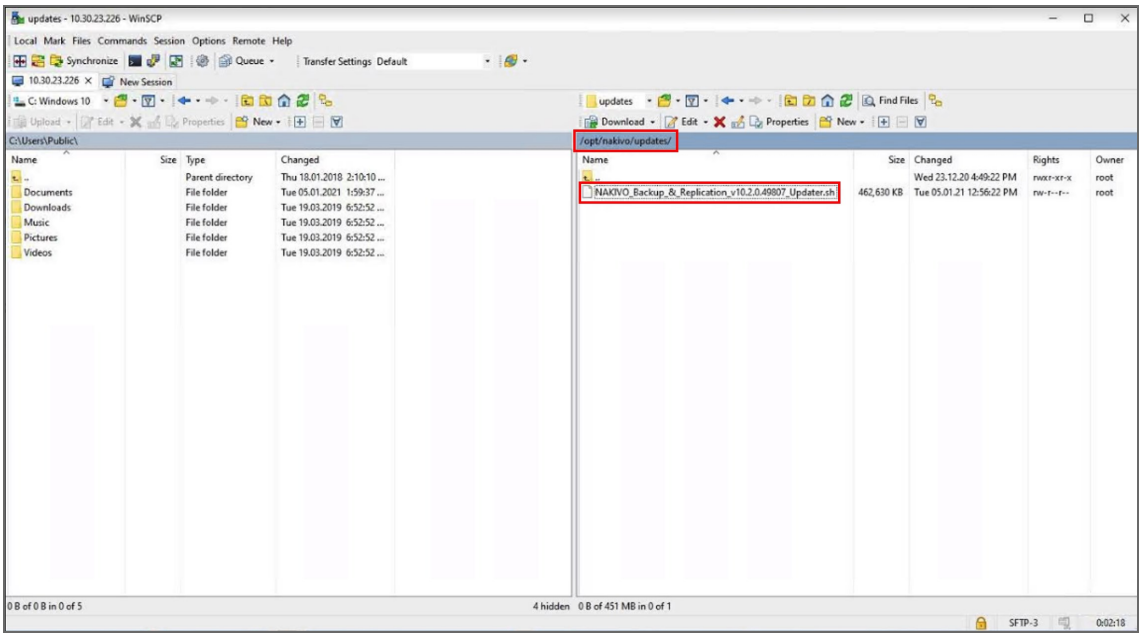

4. Log out from the SSH client.

5. Log in to your vSphere client, navigate to your VA and click **Launch Web Console**.

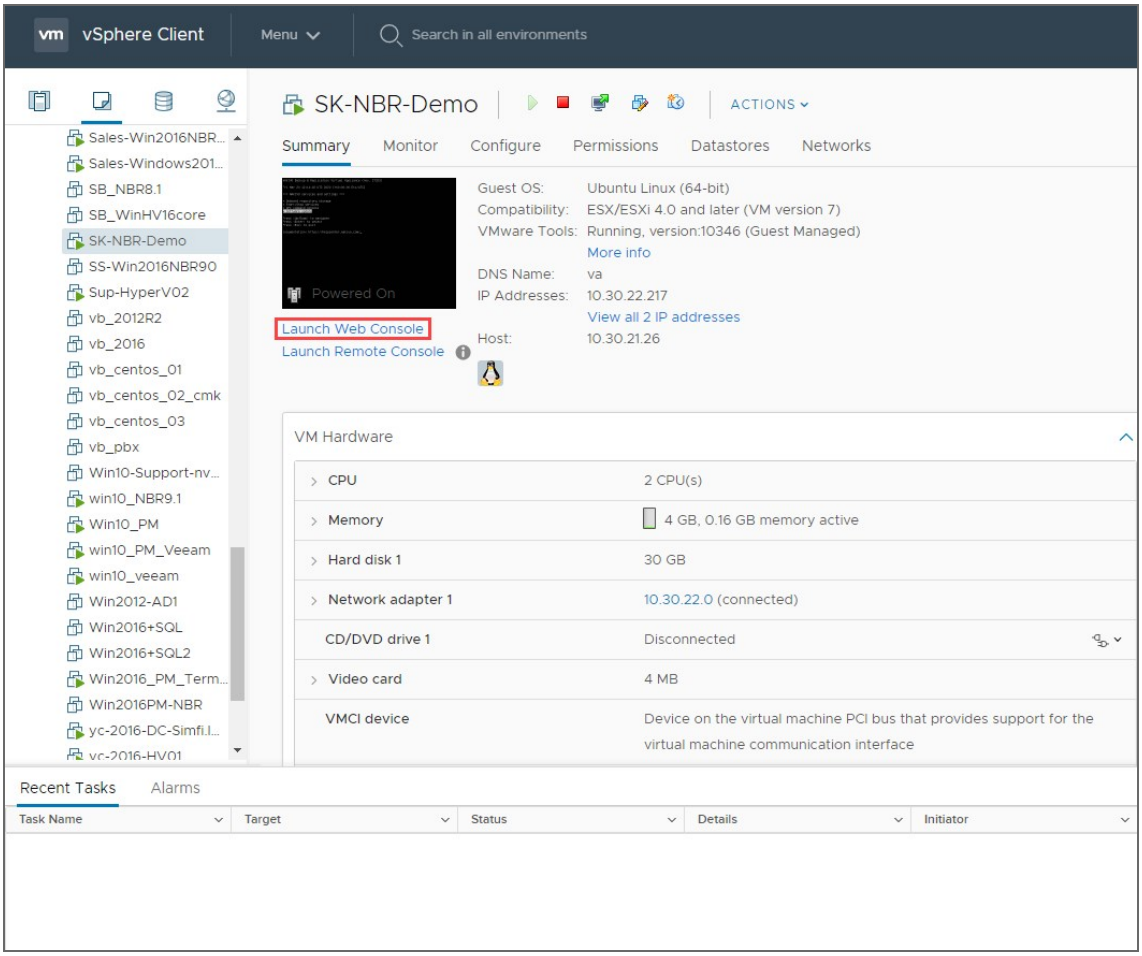

6. Do one of the following depending on the NAKIVO Backup & Replication version you use:
   - For the product Version 8.1 and above:
     1. In the VA menu, select **Manage NAKIVO services** and press **Enter**.

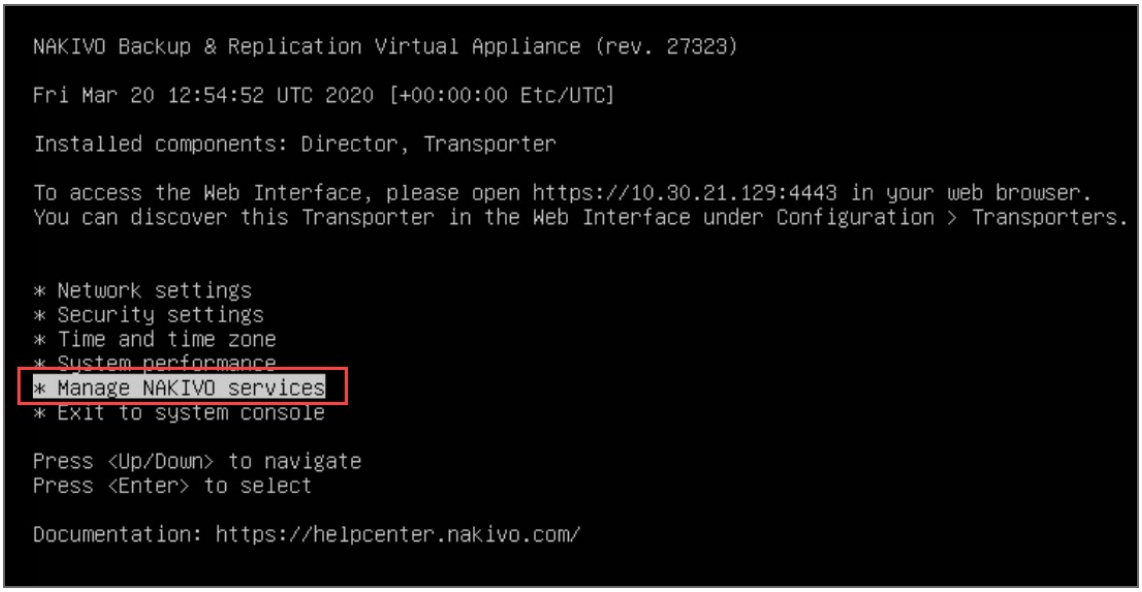

     2. In the menu that opens, select **Software update** and press **Enter**.

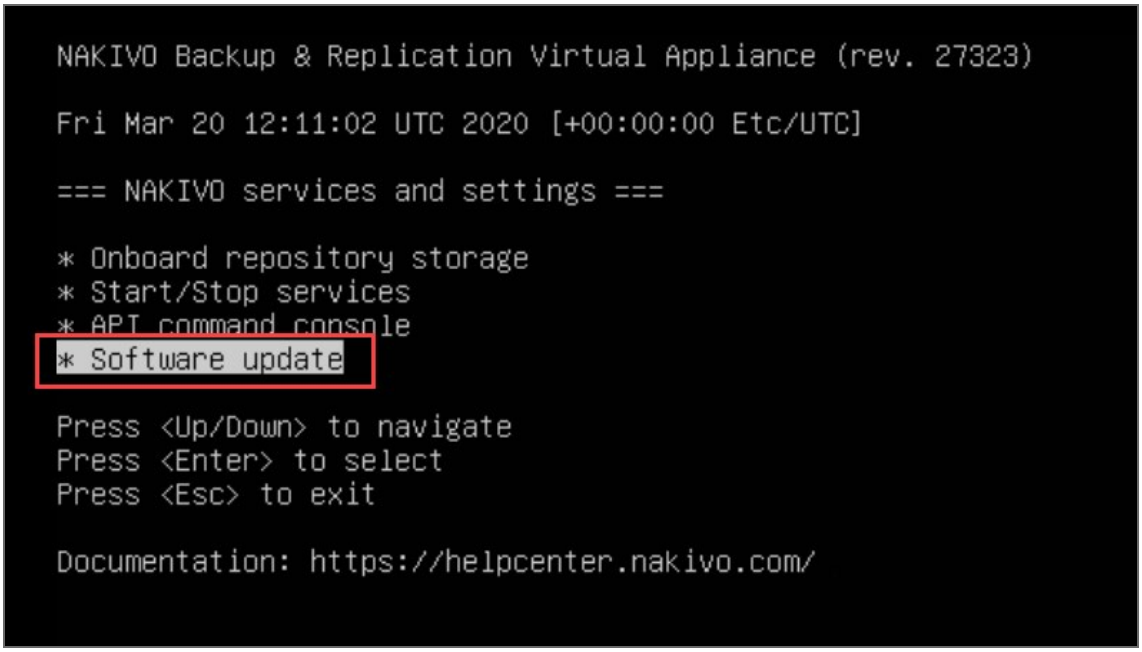

3. Select the updater that you have downloaded and press **Enter**.

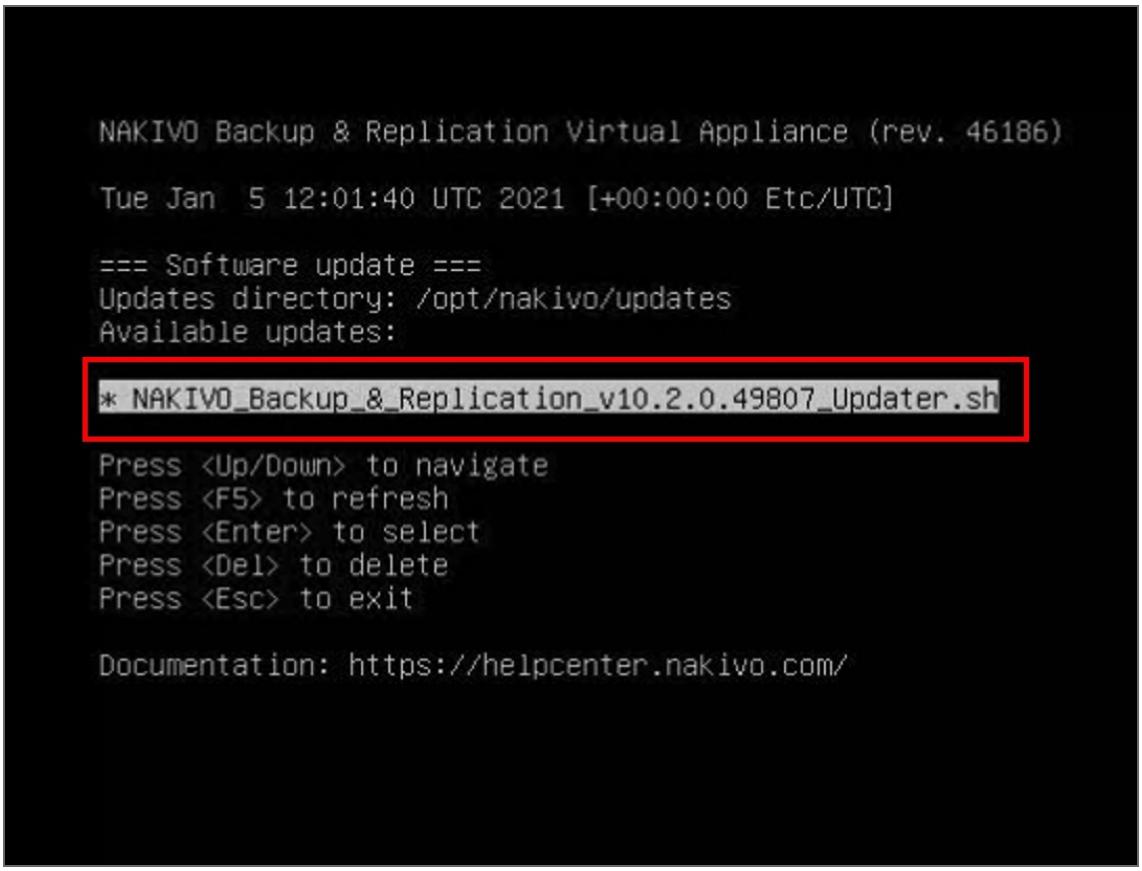

4. Review the End User License Agreement. Press **Space** to go to the bottom of it. If you agree to the terms of the agreement, type **Y** and then press **Enter** to begin the update process.

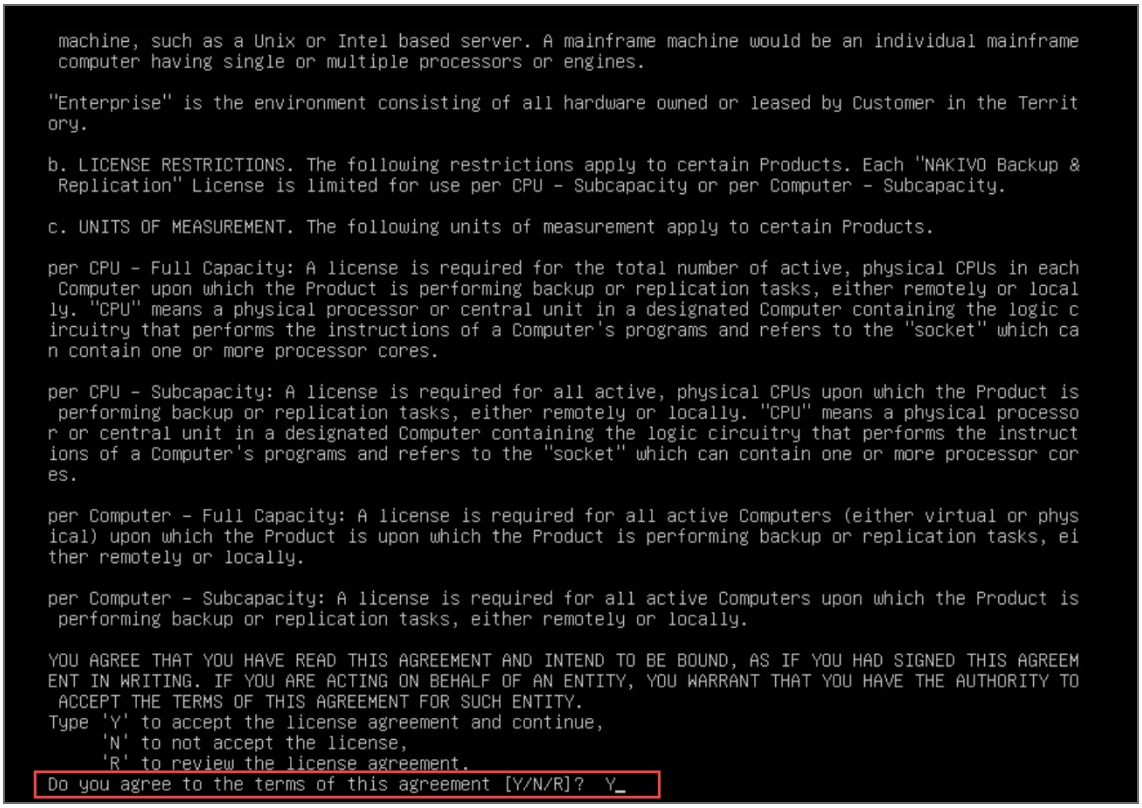

- For earlier product versions:
    1. In the VA menu, select **Software update** and press **Enter.**
    2. Select the updater that you have downloaded and press **Enter.**
    3. Review the End User License Agreement. Press **Space** to go to the bottom of it. If you agree to the terms of the agreement, type **Y** and then press **Enter** to begin the update process.

7. When the update process is complete, a message will appear to inform you about it. Exit the VA console.
8. Update all machines on which you have deployed an additional Transporter.

Note

Updating your VA with versions prior to the previous major version (for example, updating VA version 6.1 to version 9.0) is prohibited. Please update your VA to the next major version first.

Updating on Windows

If auto-update within the NAKIVO Backup & Replication interface is not supported, follow the steps below to update the product manually:

1. Download the latest Windows updater from 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 sameTransporter 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/.
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 88.

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**.

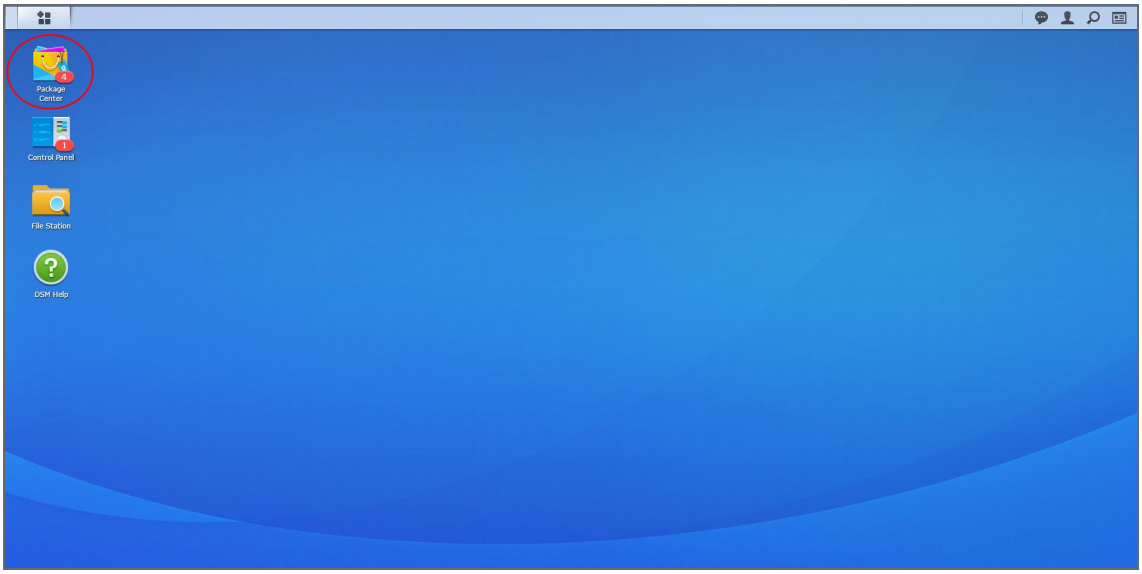

3. Go to the **Installed** section.
4. If there is a new version of NAKIVO Backup & Replication available, you will see an **Update** button.

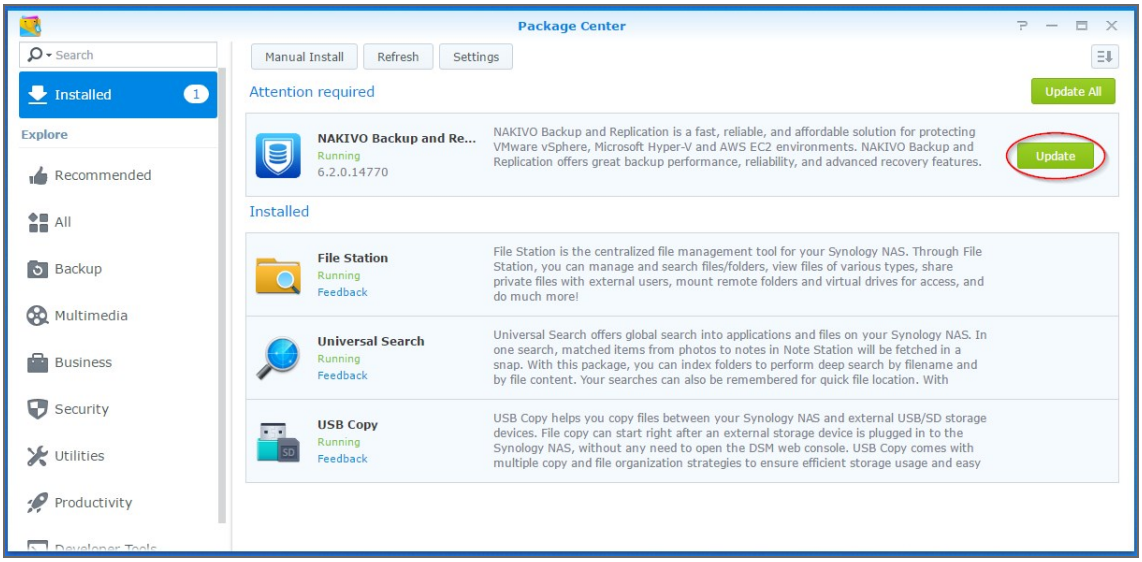

5. Click **Update**.

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/.
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**.

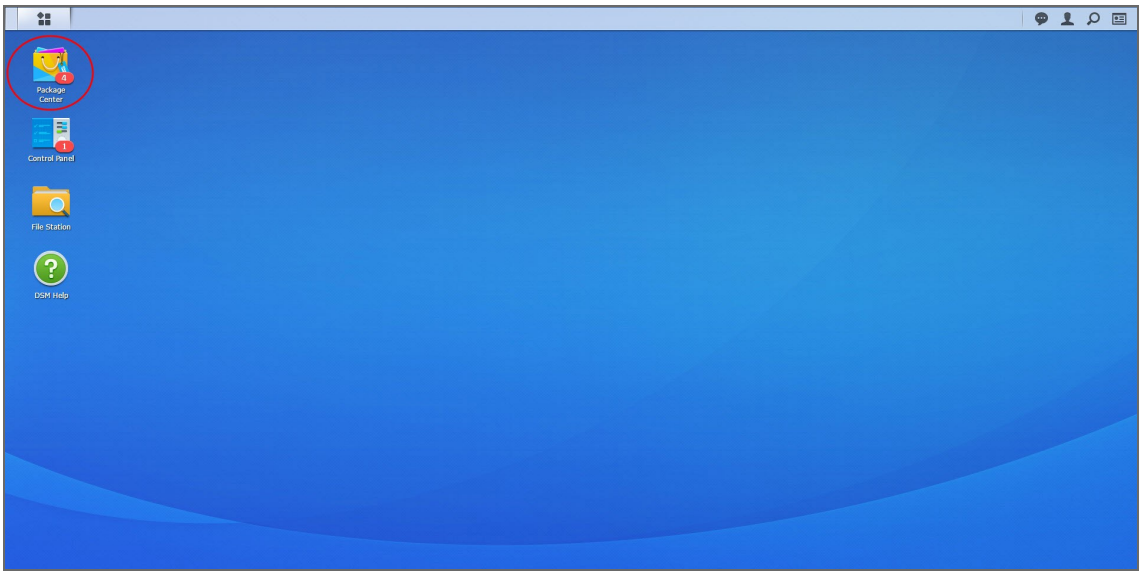

4. Click **Manual Install**.

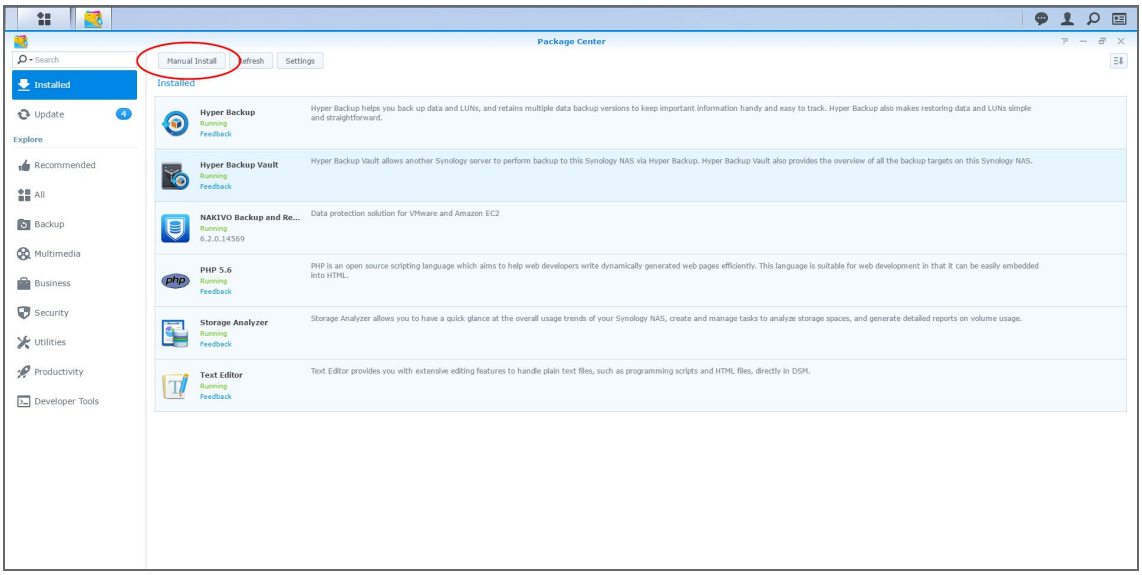

5. Click **Browse**, navigate to the Synology NAS package that you have downloaded, select it, and click **Open**.
6. Click **Next**. the package is uploaded to your NAS.
7. Click **Apply**.
8. Run an appropriate updater on all machines on which you have also installed a Transporter.

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.

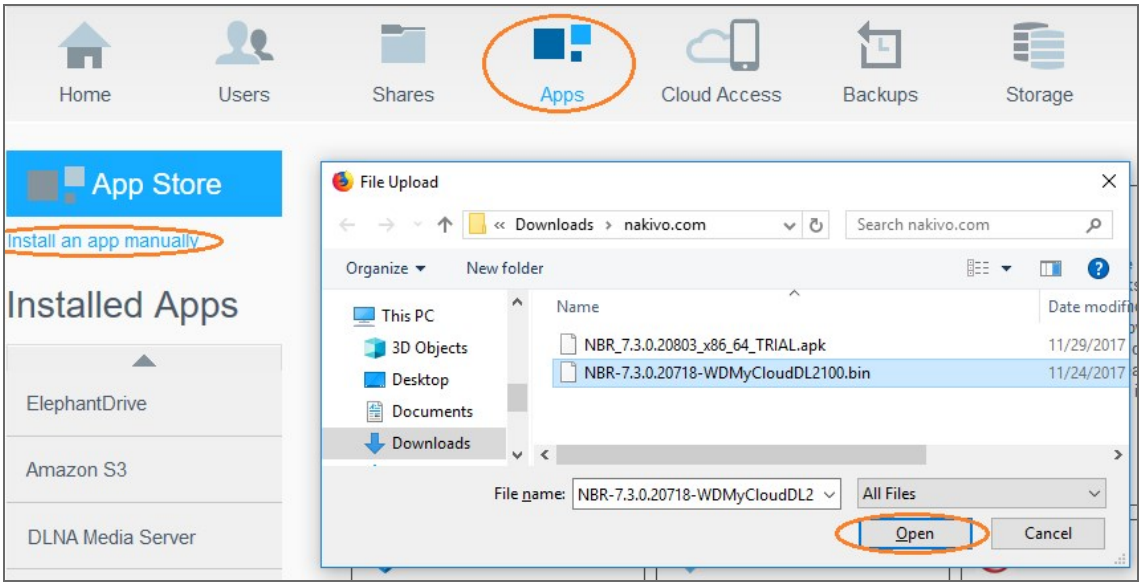

4. Once the update has successfully finished, a dialog box opens with a message including said information. Click **OK** to close the dialog box.

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 165 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.

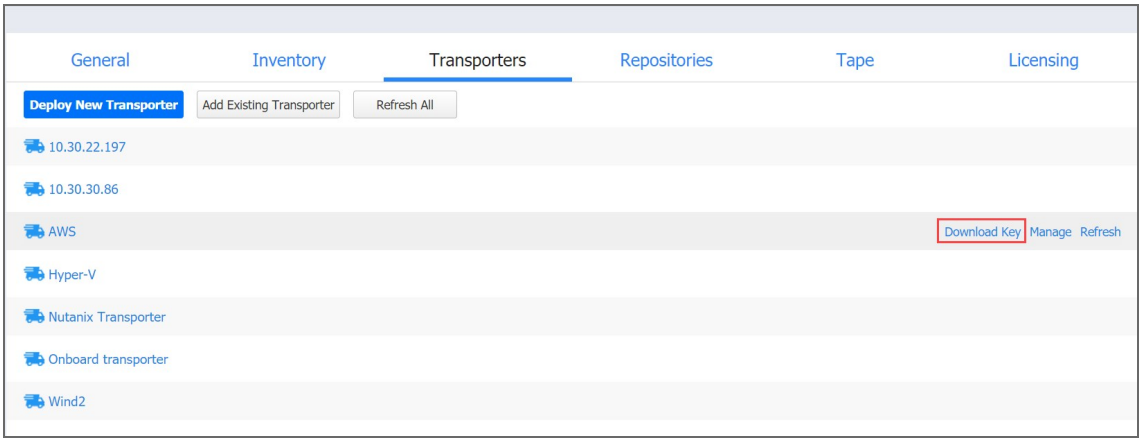

4. Click on the Transporter to view its details. Copy or remember the IP-address/hostname of the Amazon instance.
5. Unzip the folder with the key.
6. Convert the key using PuTTYgen:

1. In PuTTYgen menu, go to *Conversions > Import.*

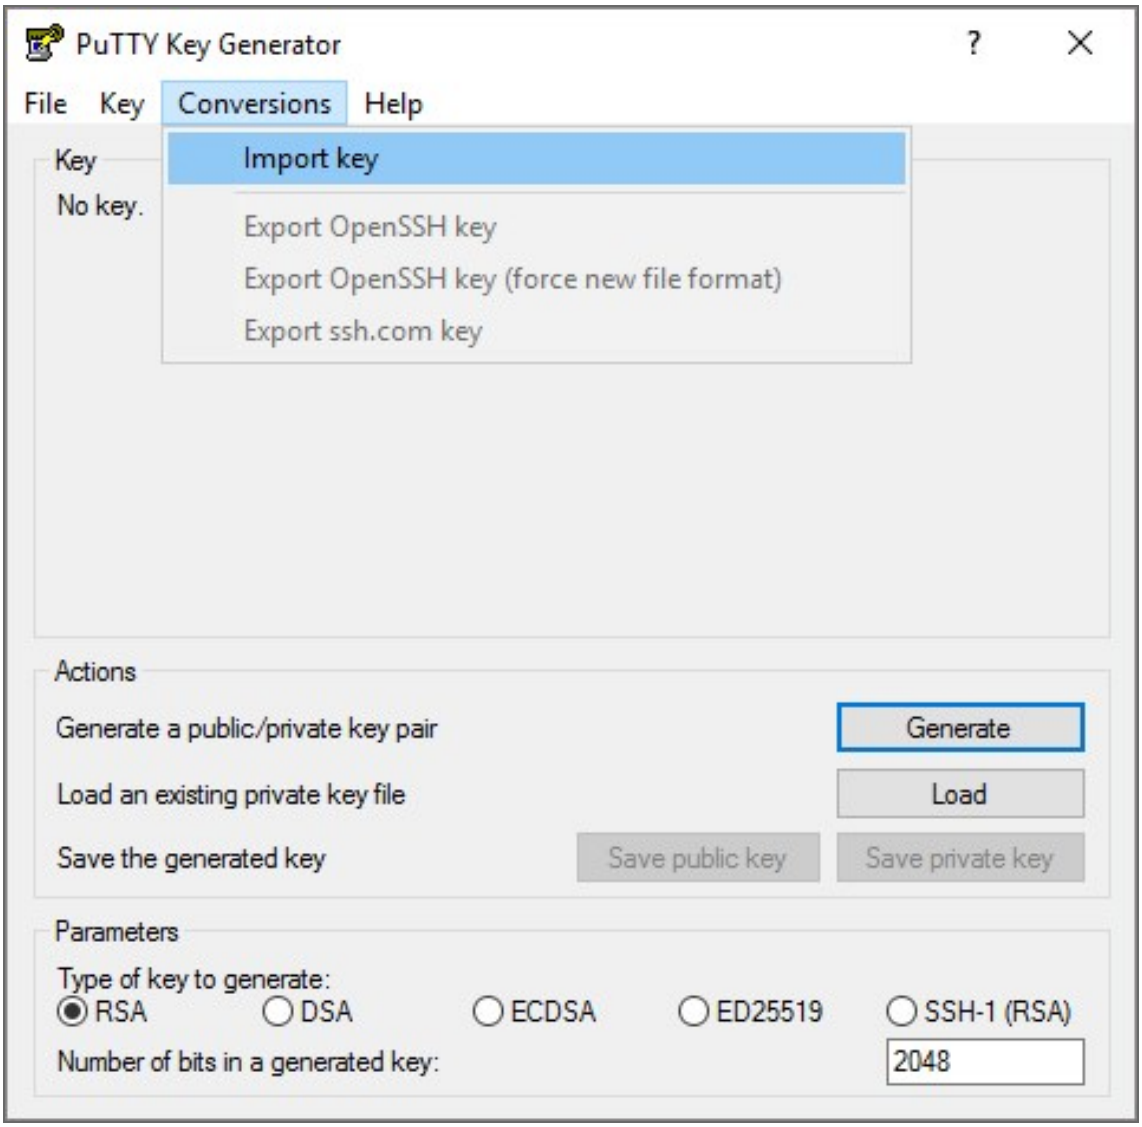

2. Locate the `SSH_key.pem` you just downloaded and unzipped. If you don't see it in the **Open**... dialogue box, change the file type to **All files**.

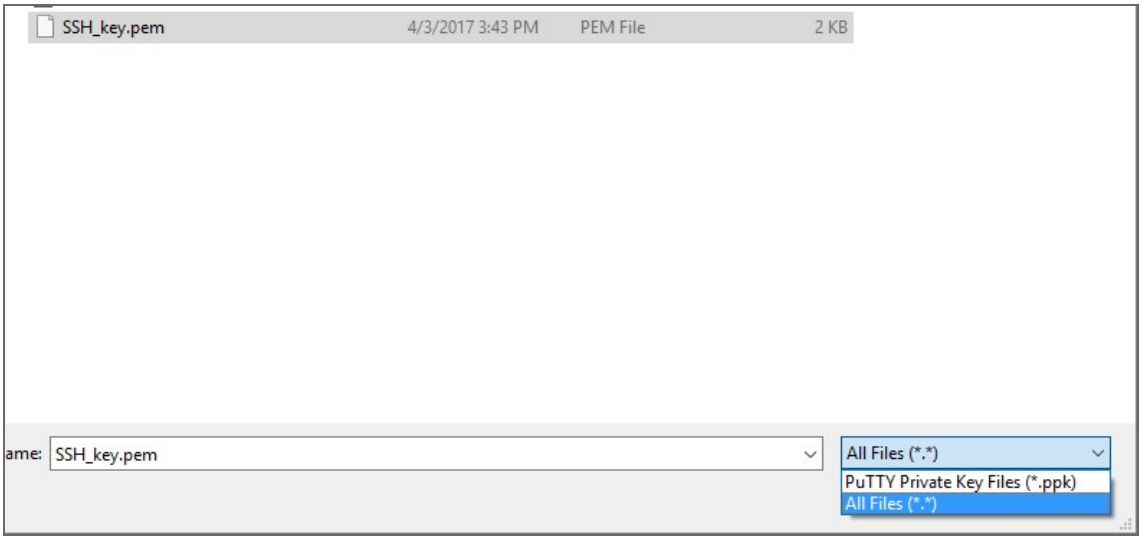

3. Click on **Save private key**. If PuTTYgen asks you to save the key without a passphrase, click **Yes**.

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.

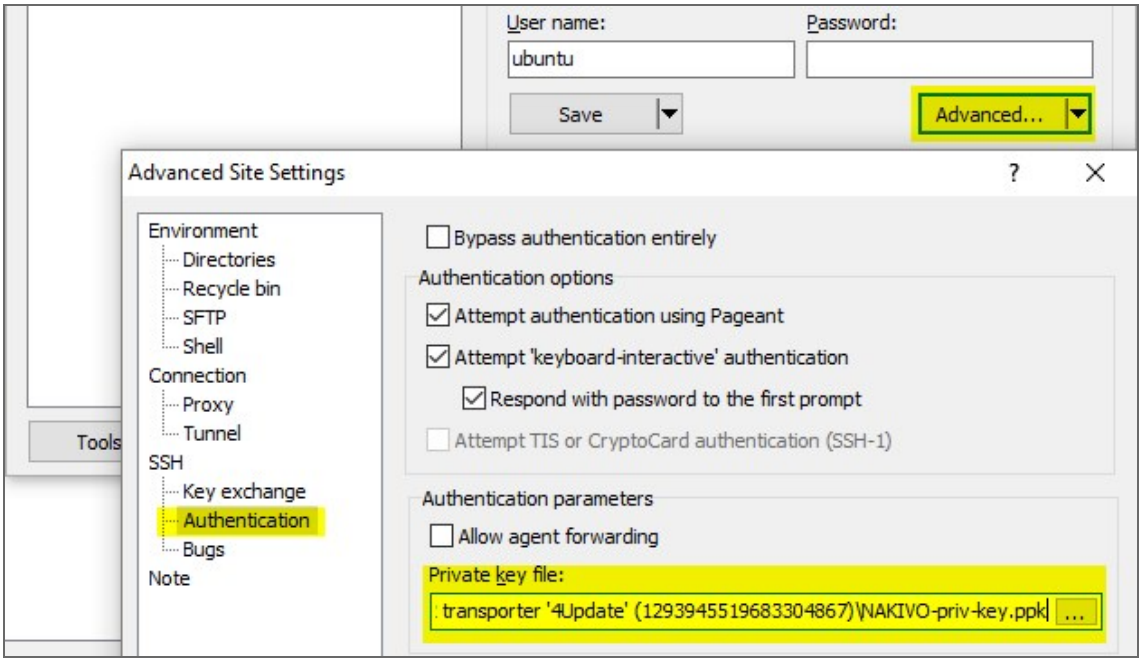

3. Click **OK**.

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.

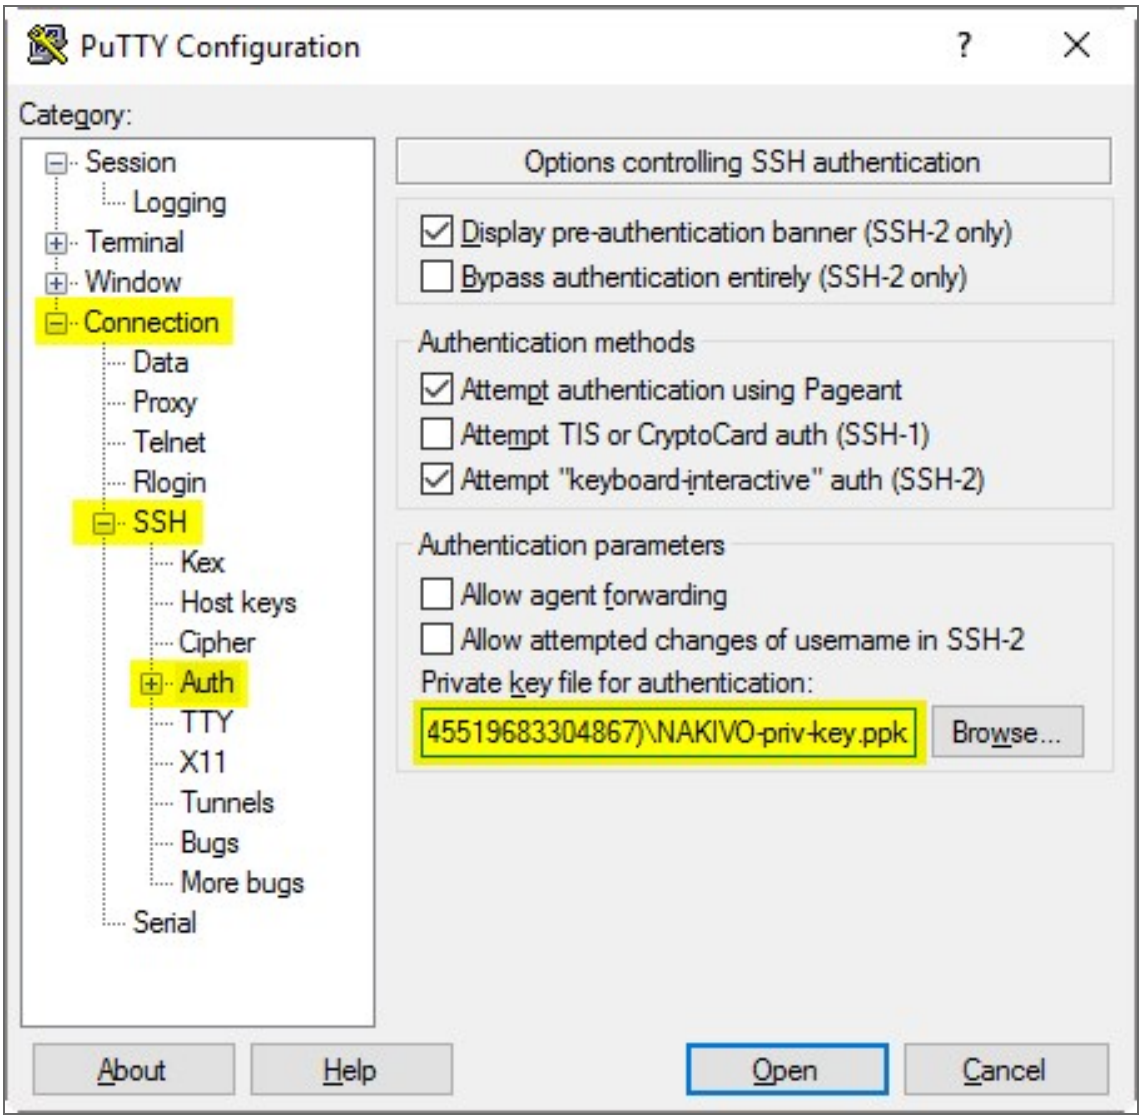

10. Click **Open**.
11. In the command line prompt that opens: log in to the Amazon EC2 instance:
    1. For **login**, enter `ubuntu`
    2. For **password**, leave a blank line.

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.

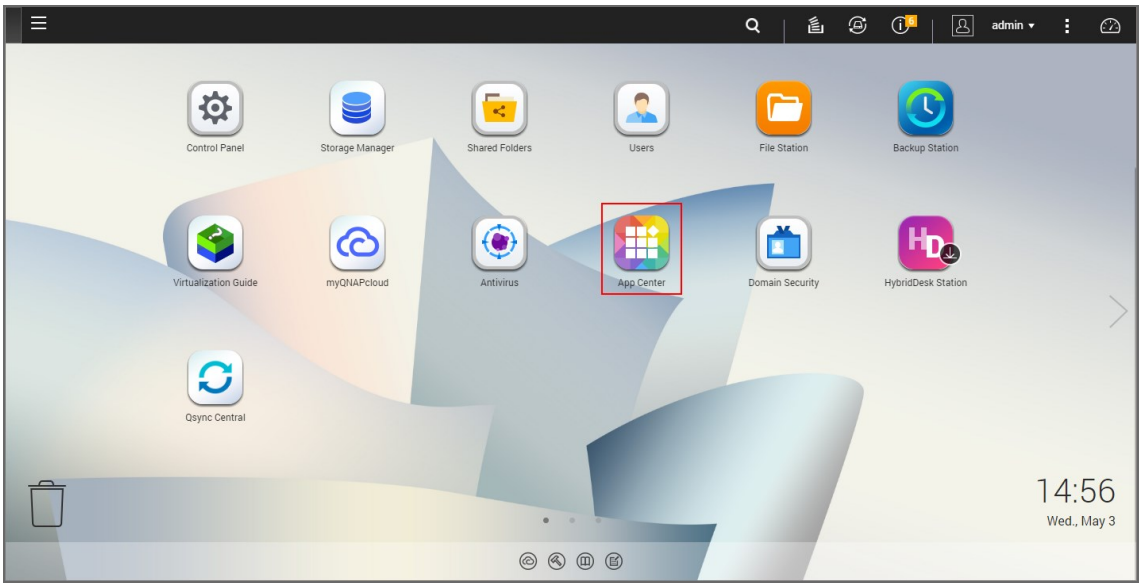

2. Go to **App Center**.
3. Select the *Backup/Sync* category and find NAKIVO Backup & Replication. Alternatively, use the search box at the top of the App Center window: click on the magnifier icon and enter "Nakivo".
4. If the new version of NAKIVO Backup & Replication is available in the QNAP App Center, you will see a green **Update** button.

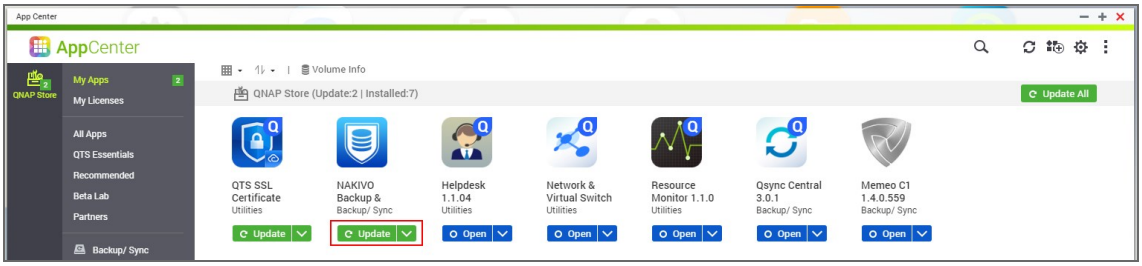

5. Click the **Update button** and wait till update finishes.

Updating Manually

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

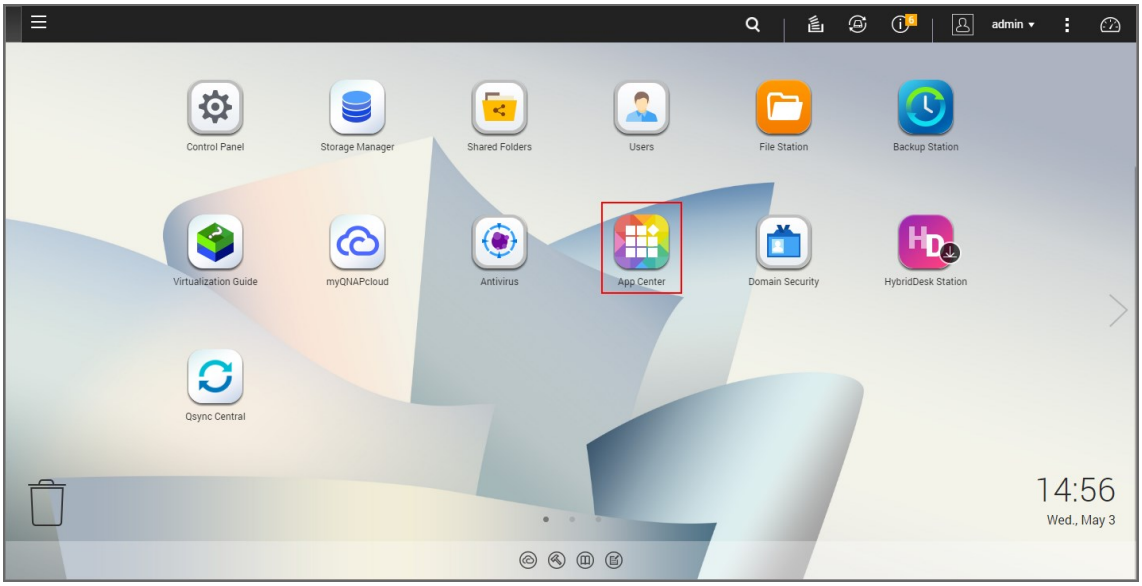

3. Go to **App Center**.
4. Click the **Install Manually** icon.

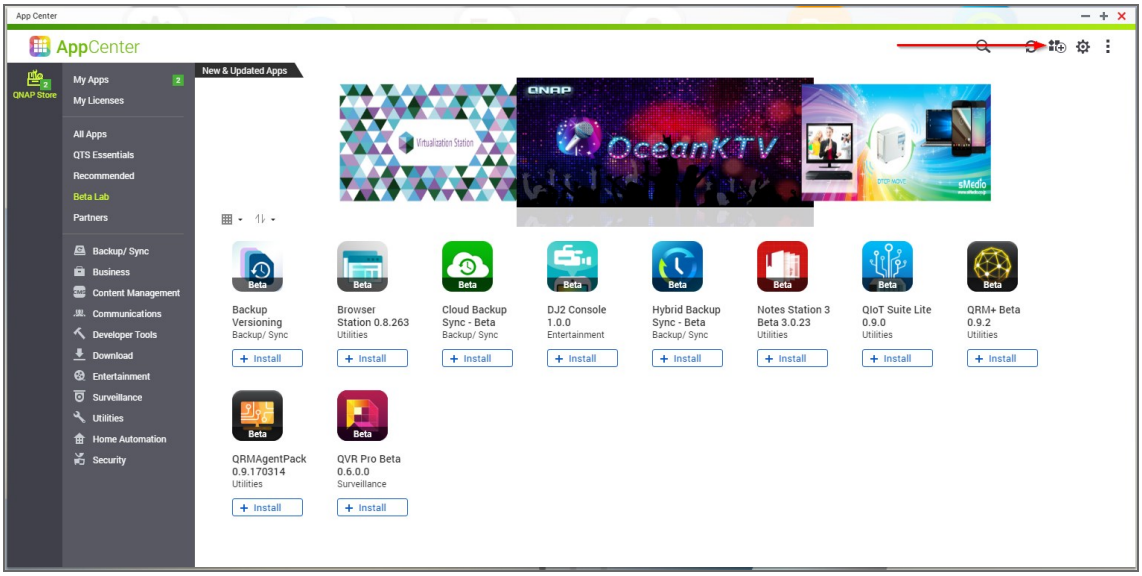

5. Click **Browse**. In the window appears, locate the installer (.qpkg file) on your computer.

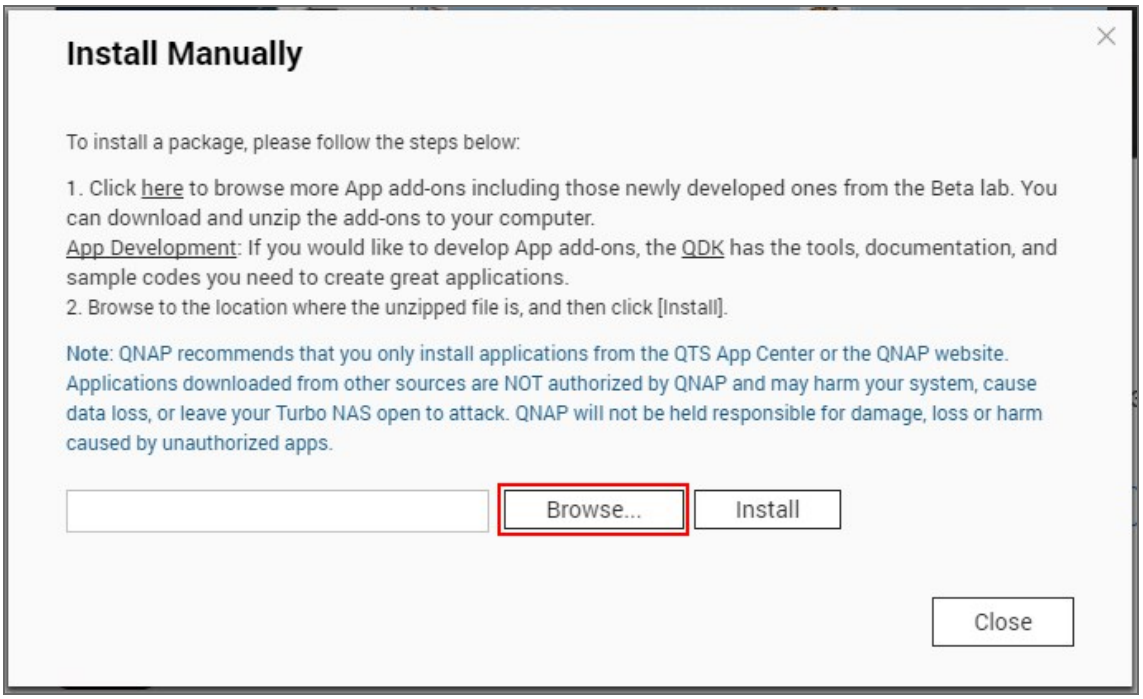

6. Click **Install**.
7. Wait until the update process is finished.

Updating on ASUSTOR NAS

- Updating on ASUSTOR NAS Manually
- Updating on ASUSTOR NAS via App Central

Updating on ASUSTOR NAS Manually

Prior to updating NAKIVO Backup & Replication on ASUSTOR NAS manually, make sure the following requirements are met:

- You have access to the ASUSTOR NAS.
- NAKIVO Backup & Replication installer is available for your ASUSTOR NAS.

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.

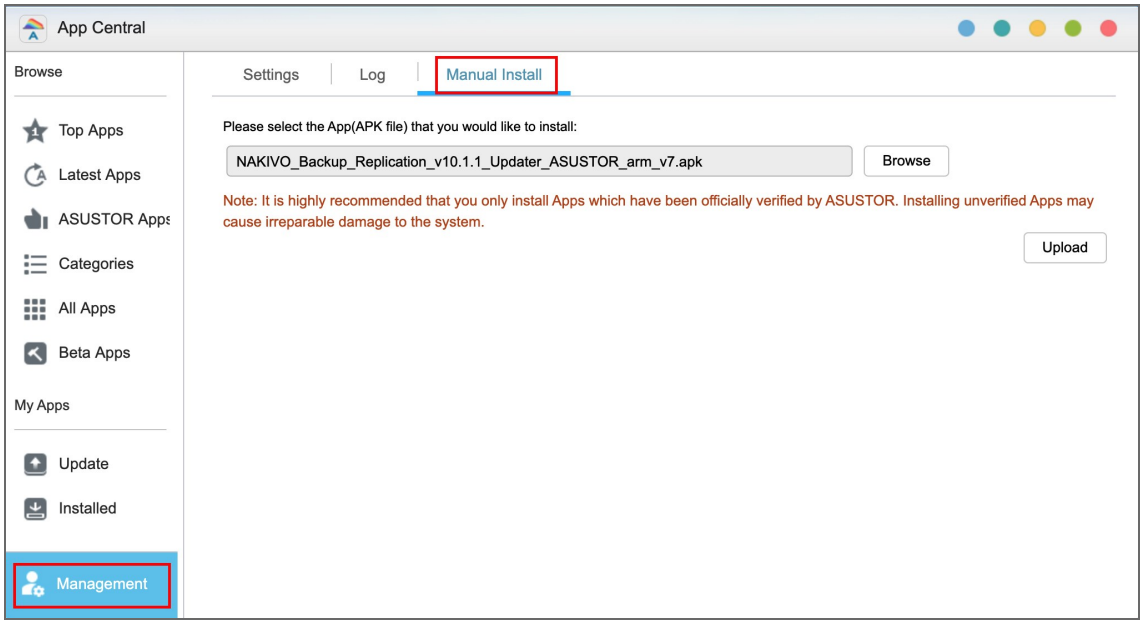

Updating on ASUSTOR NAS via App Central

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.

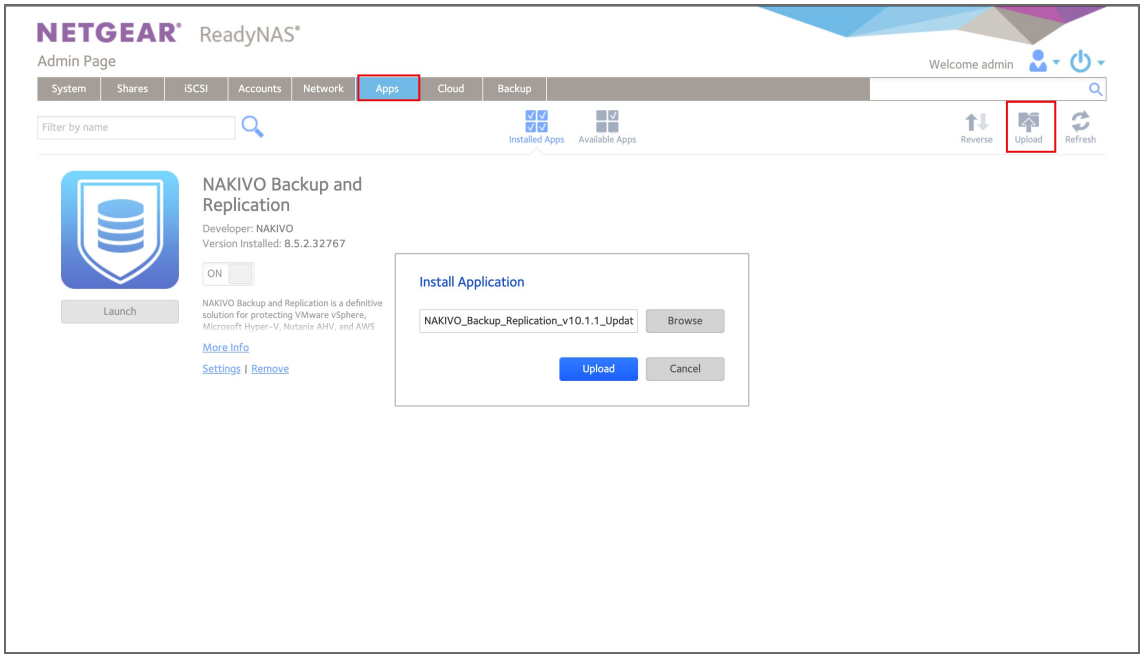

Updating on NETGEAR ReadyNAS via Available Apps

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 Linux/VA 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_Updater.sh
   ```
4. Change the updater file permission with the `chmod` command:
   ```
   chmod +x NAKIVO_Backup_&_Replication_vX.X.X_Updater.sh
   ```
5. Run the updater in silent mode:
   ```
   ./NAKIVO_Backup_&_Replication_vX.X.X_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/.
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” on page 235
- “First Steps with NAKIVO Backup & Replication” on page 241
- “Web Interface Components” on page 244
- “Managing Jobs and Activities” on page 249

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.

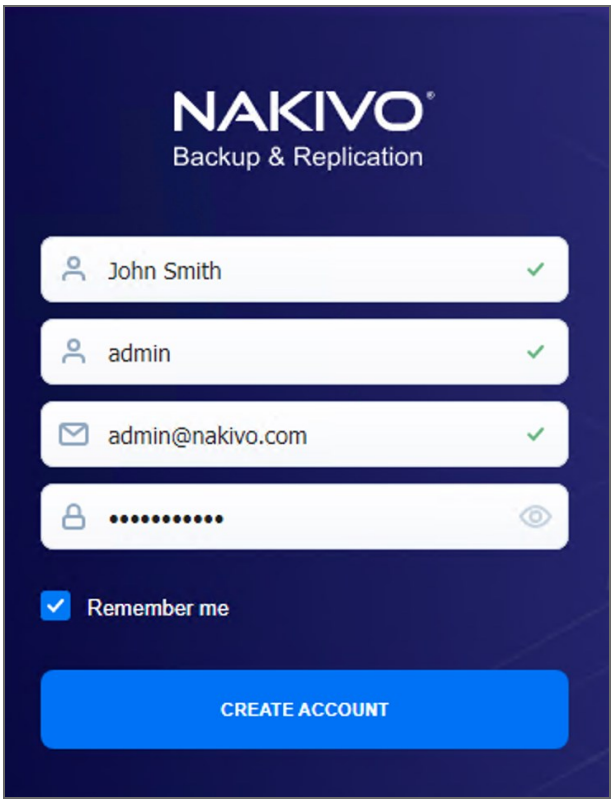

NAKIVO Backup & Replication opens in your browser displaying the configuration wizard. Refer to First Steps with NAKIVO Backup & Replication to learn how to start using NAKIVO Backup & Replication.

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.

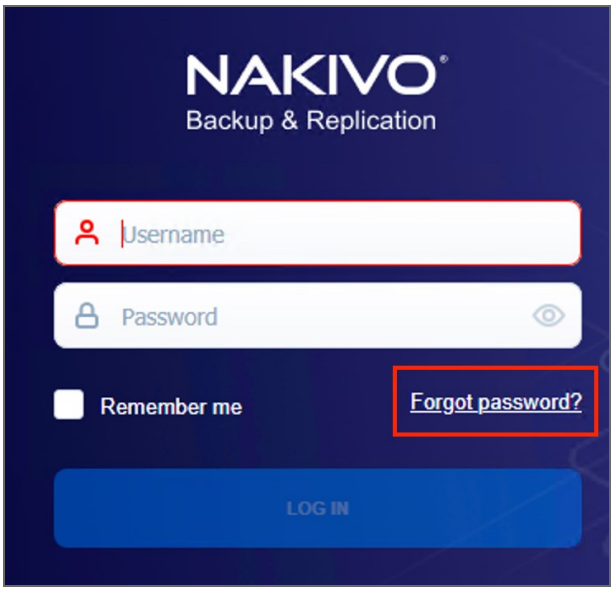

3. Do one of the following:
   - If you have set up email settings in NAKIVO Backup & Replication, enter your email address and click **Done**.

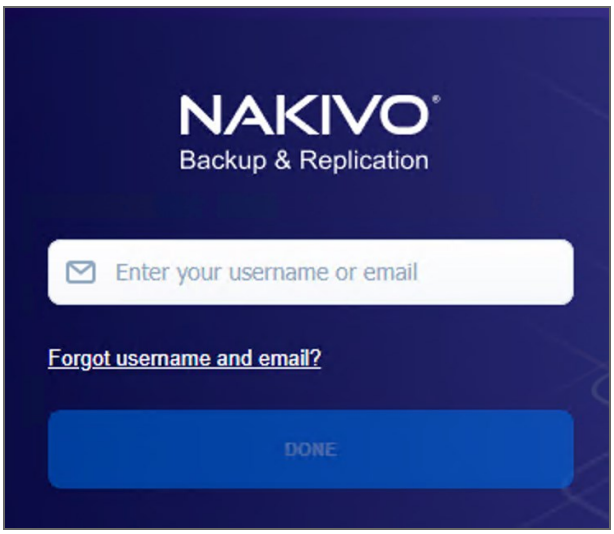

A temporary password, which is a security string, is sent to your inbox. Enter this password the next time you log in to your NAKIVO Backup & Replication instance. Once you are logged in, it's recommended that you change the temporary password for your user account. To change the temporary password:

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**.

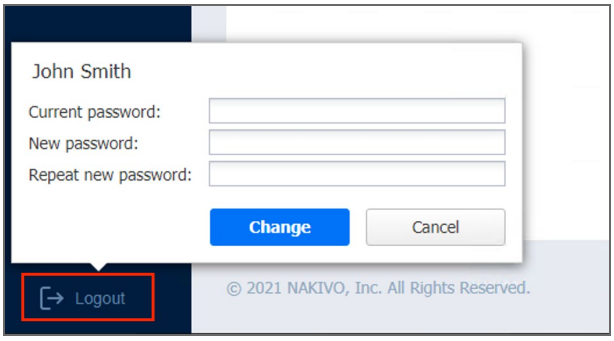

You can also change your temporary password in **Settings**>**General**>**Users and Roles**

- If you have not set up email settings in NAKIVO Backup & Replication:

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

b. Go to the product installation folder and locate the “forgot_password.txt” file.

Important

For security reasons, only a root user (Linux) or a member of the Administrators group (Windows) is allowed to access the installation folder and the "forgot_password.txt" file.

c. Paste the security string from the file in the appropriate field.

d. Click **Done**.

Notes

- If you are using a Virtual Appliance (VA), go to the VA console, then go to the command line and enter: `cat /opt/nakivo/director/forgot_password.txt` The security string will be displayed on the screen. You can copy and paste it into the web interface.
- 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 options below:
    - VMware vCenter or ESXi host
    - VMware Cloud Director
    - Microsoft Hyper-V host or cluster
    - Physical machine(s)
    - Microsoft 365 account
    - Nutanix AHV cluster
    - AWS/Wasabi account
    - HPE 3PAR storage device
    - Oracle database

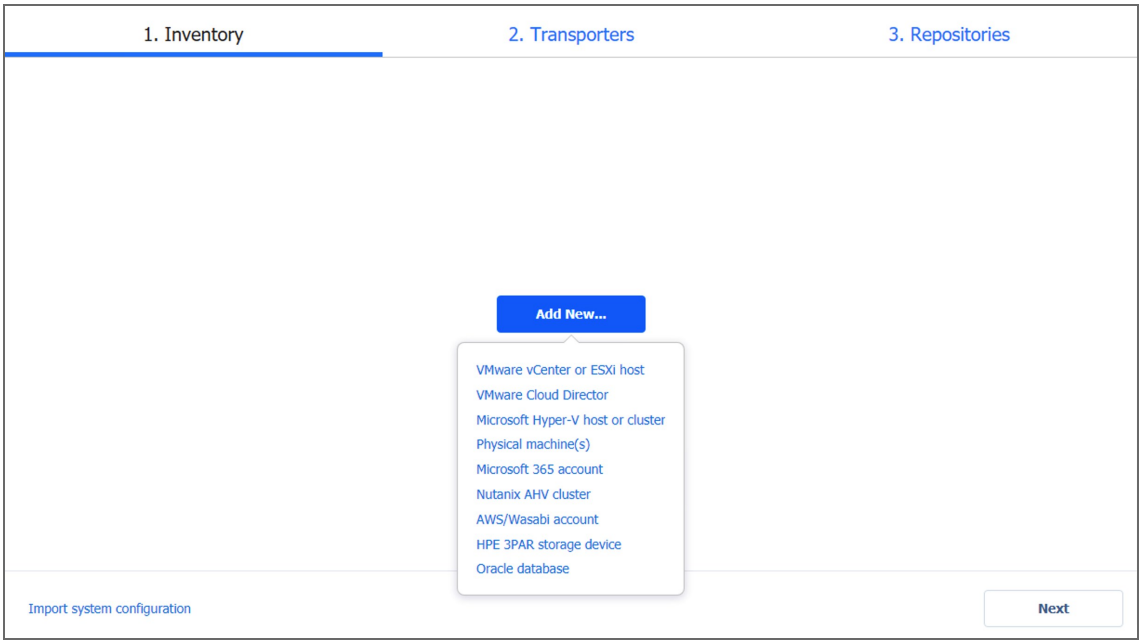

3. Proceed with adding items as described in the Inventory article.
4. On the **Transporters** page of the wizard, you will find information about the Transporter component of the NAKIVO Backup & Replication.
5. To deploy a new Transporter or add an existing one, click **Got it** and proceed as described in the Transporters article.
6. To move to the next page of the wizard, click **Next**.

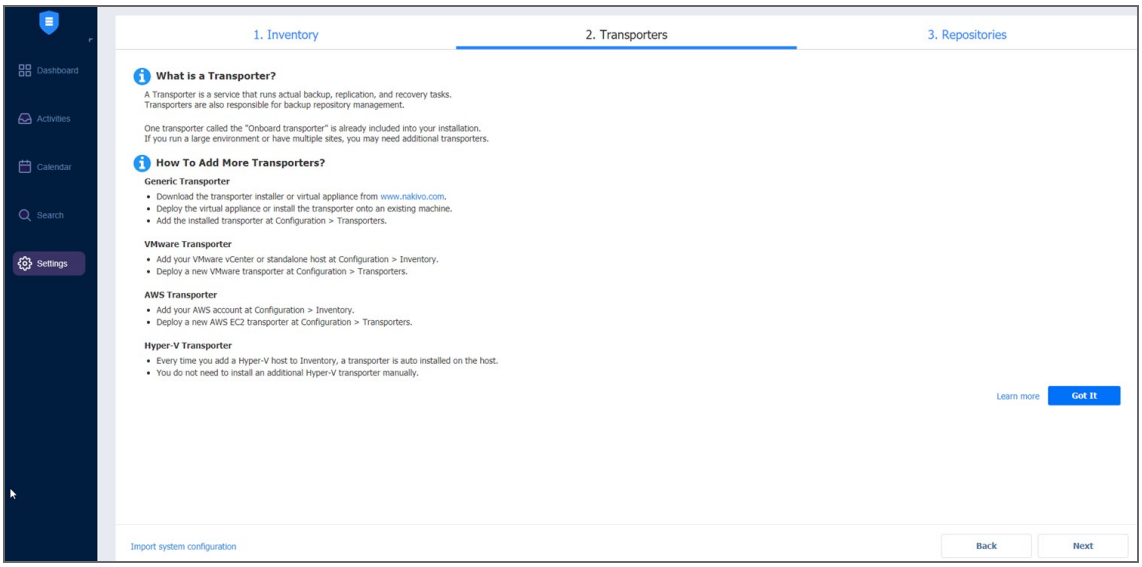

7. On the **Repositories** page of the wizard, you can add a local or a remote Backup Repository to your application by clicking **Add Backup Repository**.

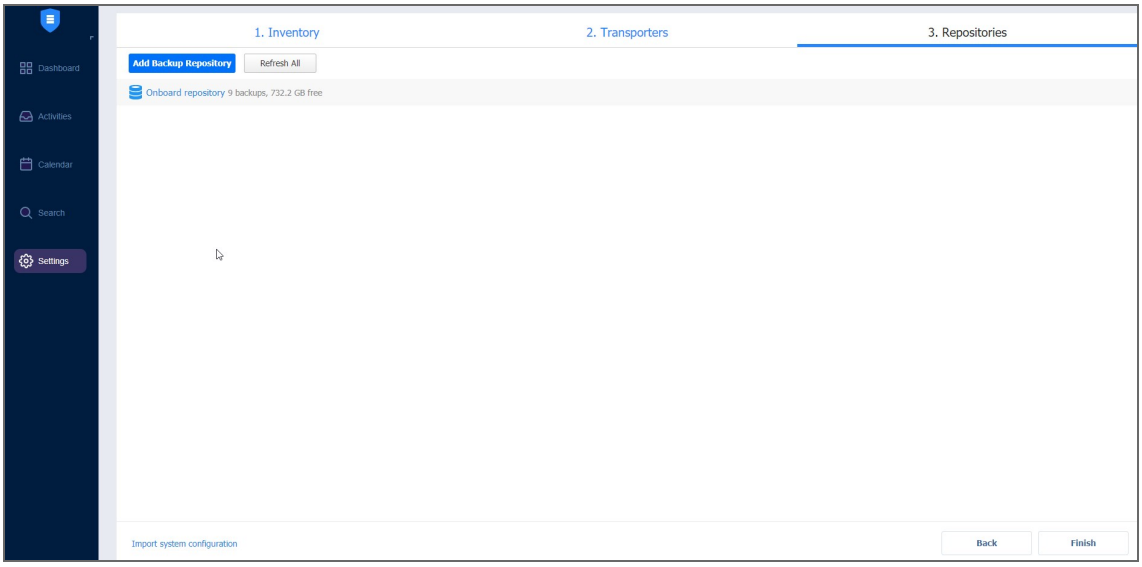

8. Click **Finish**.
9. The **Dashboard** of the application opens. Proceed with creating your backup and replication jobs.

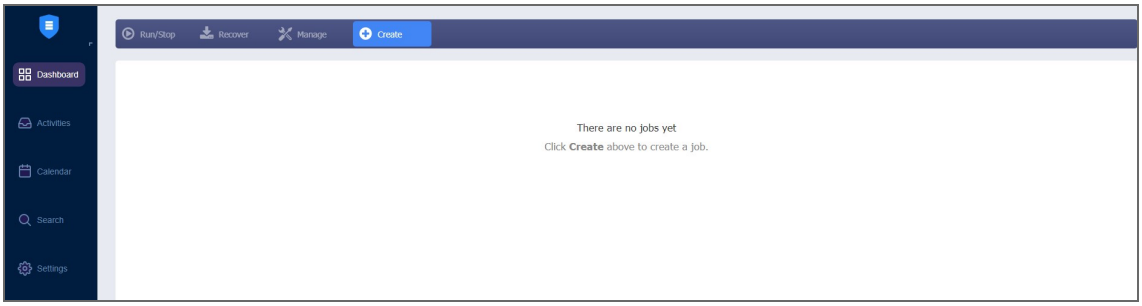

If your current license type is **Free** and the **Trial** license has not yet been applied to the current deployment of NAKIVO Backup & Replication, a dialog box appears. Using this dialog box, you can contact the sales team to change your license type or try the full functionality of the solution for 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
- 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.

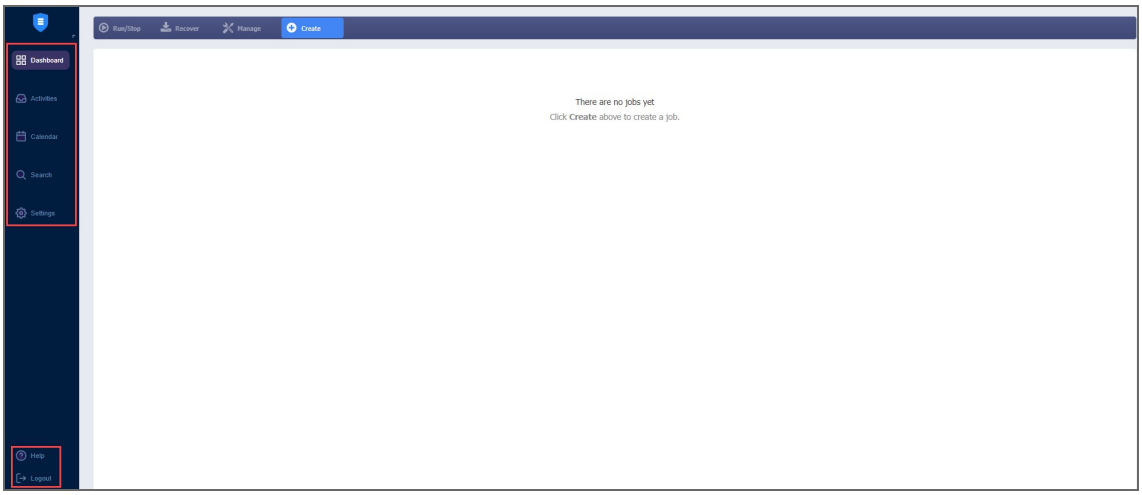

Dashboard

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

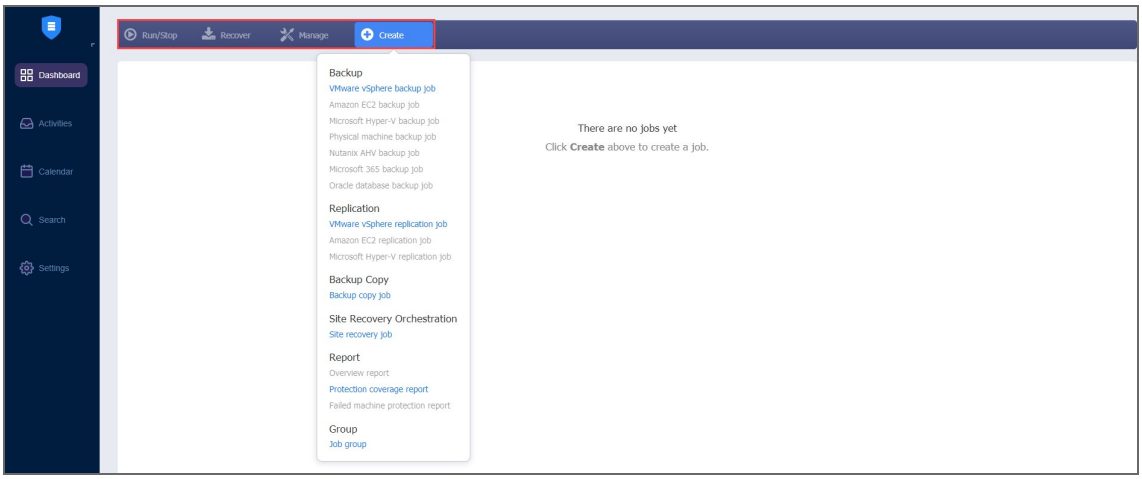

Activities

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: Tape scan, Tape erase, reading Tape
- Other

For further details and information, refer to “Managing Activities” on page 271 .

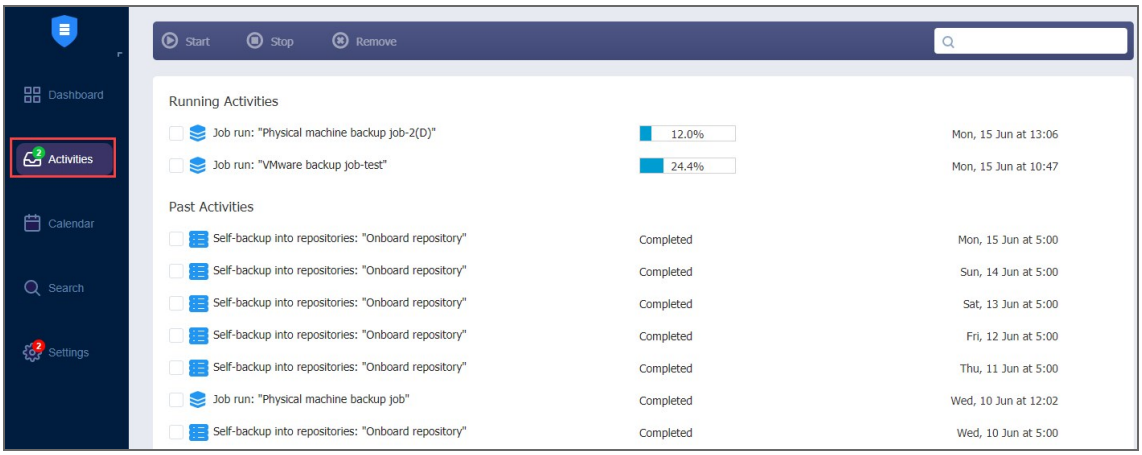

Calendar

The **Calendar** allows you to schedule jobs and to view the history of all job runs right from the calendar. For more information, refer to “Using Calendar” on page 275.

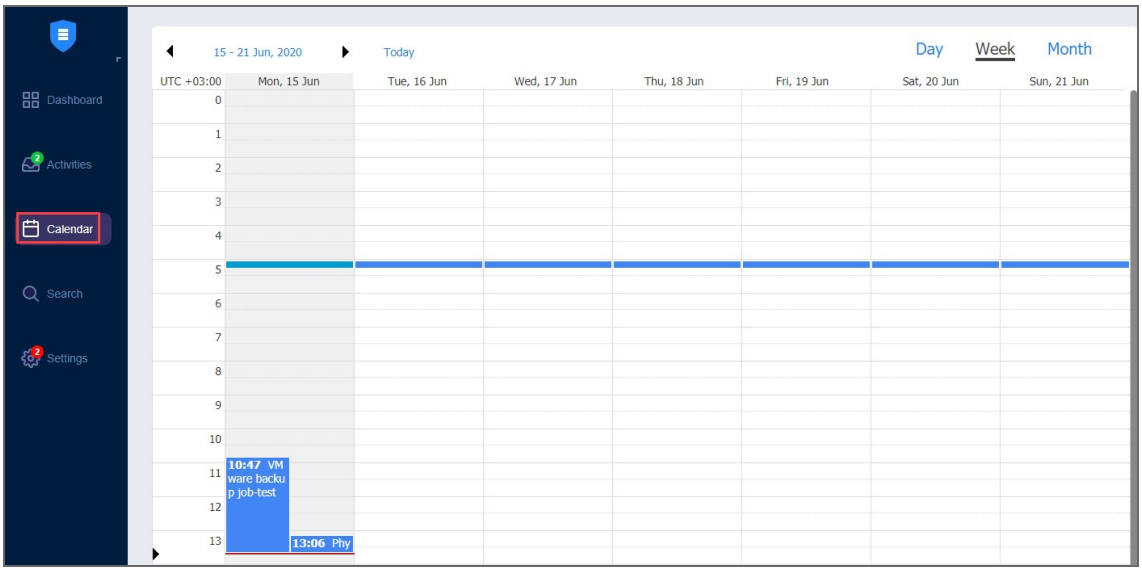

Search

The **Search** page allows you to search for items within the entire N AKIVO Backup & Replication – in the inventory, Transporters, Backup Repositories, tape devices, jobs, backups, replicas, and other. For more details, refer to “Using Global Search” on page 276.

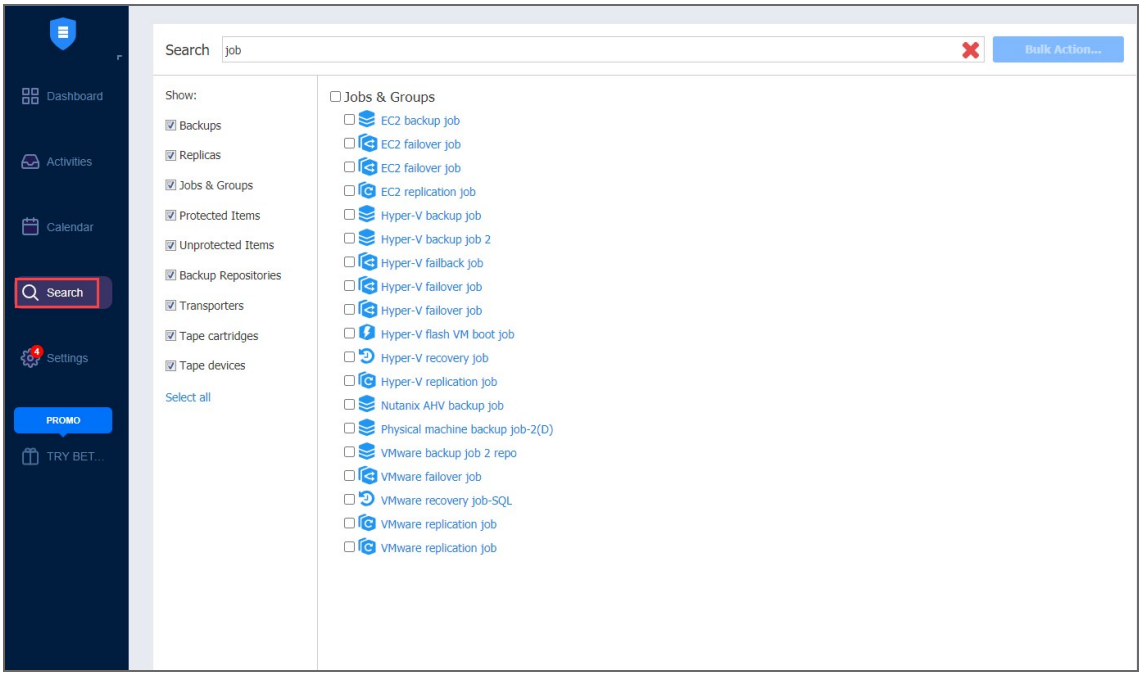

Settings

On the **Settings** page, you can configure NAKIVO Backup & Replication general settings, inventory, Transporters, Backup Repositories, Tape, and licensing. Refer to "Settings" on page 279 for more details.

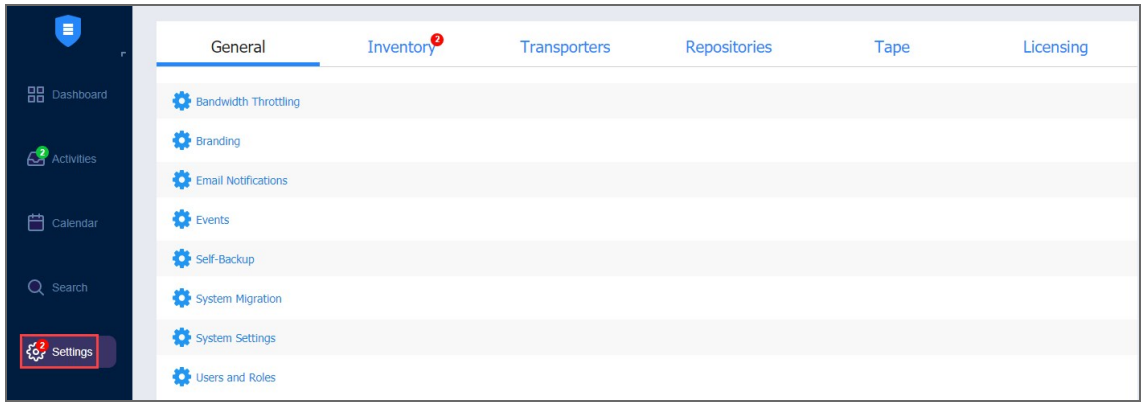

Help Menu

Use the **Help** menu to request technical support and access the NAKIVO online help center. If you are evaluating NAKIVO Backup & Replication, you can use the **How to Buy** section of the **Help** menu to view pricing, upgrade your Free license to Trial for 15 days with the **Try full functionality** option, and request a live demo or a quote.

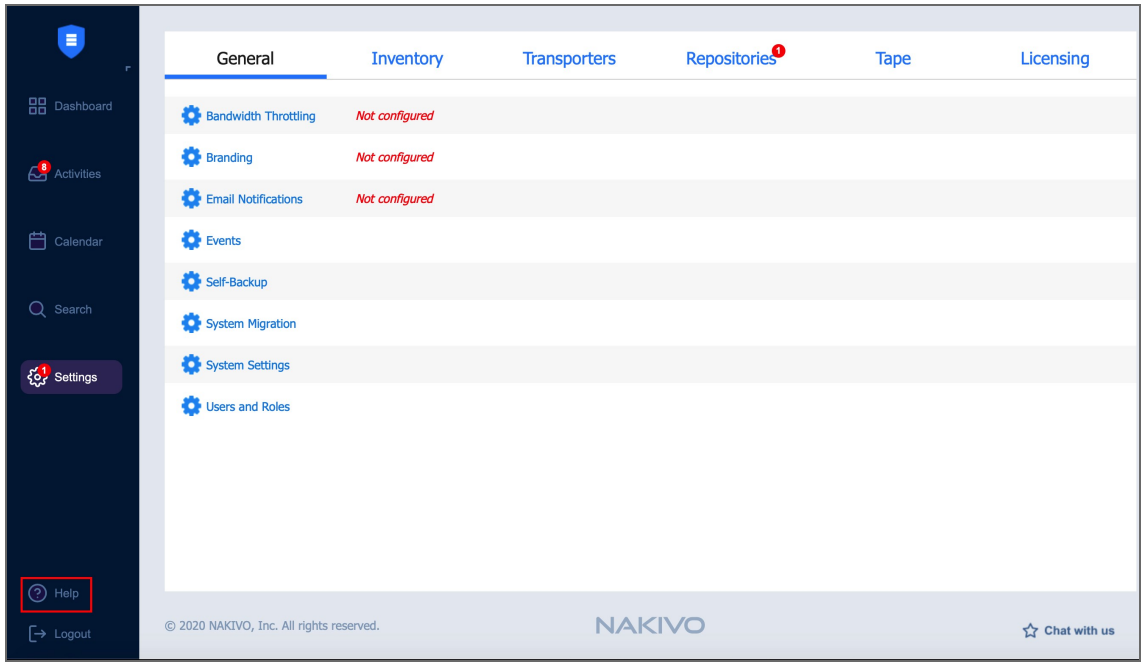

Online Chat Dialog

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 297 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 page 250
- “Managing Jobs” on page 253
- “Managing Job Policies” on page 262
- “Managing Policy Rules” on page 265
- “Jobs Alarms and Notifications” on page 269
- “Managing Activities” on page 271
- “Using Calendar” on page 275
- “Using Global Search” on page 276

Running Jobs on Demand

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

- Starting Jobs
- Stopping 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 VM/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**: The job will run for previously failed VMs/backups/physical machines/ items only.
    b. If the type of the Backup Repository for your job is Incremental with full backups, 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. In the dialog that opens, define the scope of your job:

3. Click the **Run** button to confirm your operation.

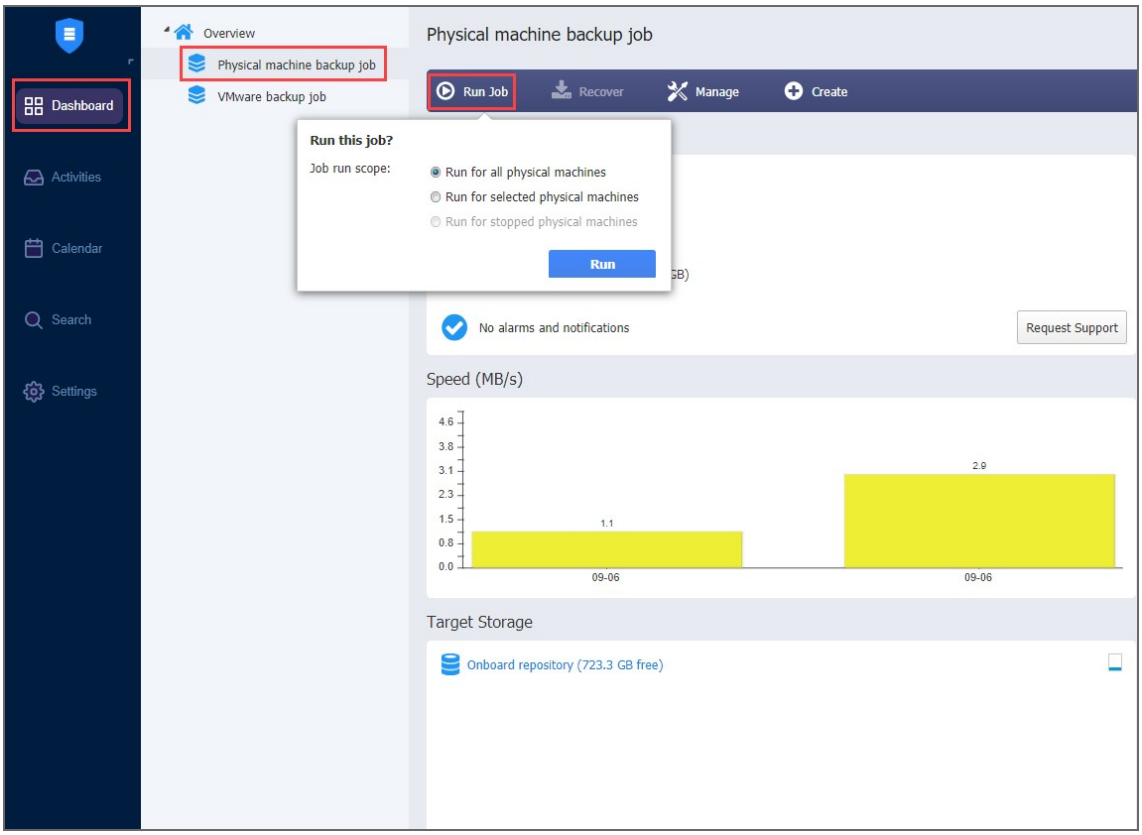

The product will close the dialog box and start running your job.

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.

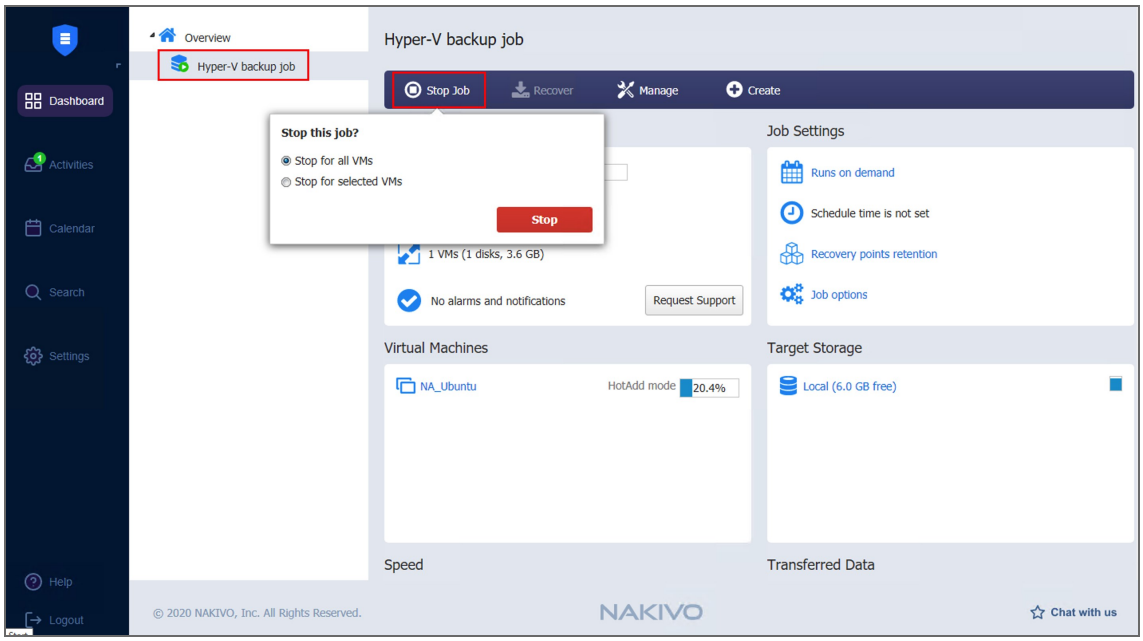

The product will close the dialog box and stop your job.

Managing 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**.

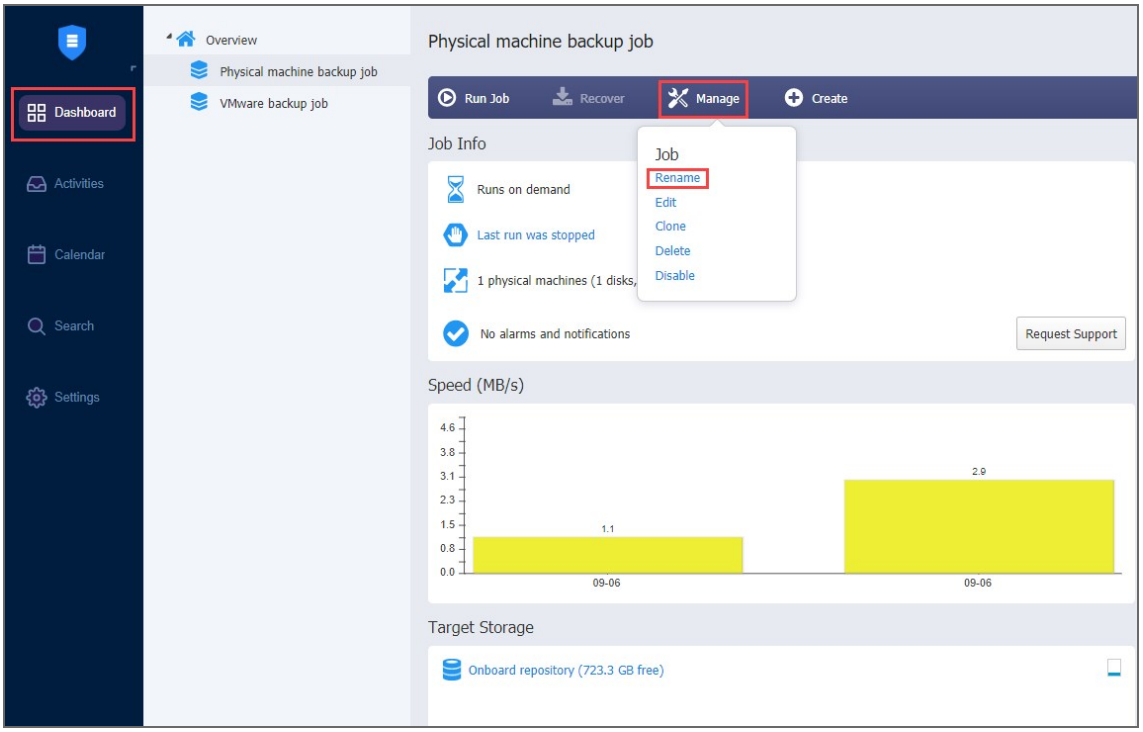

Note

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**.

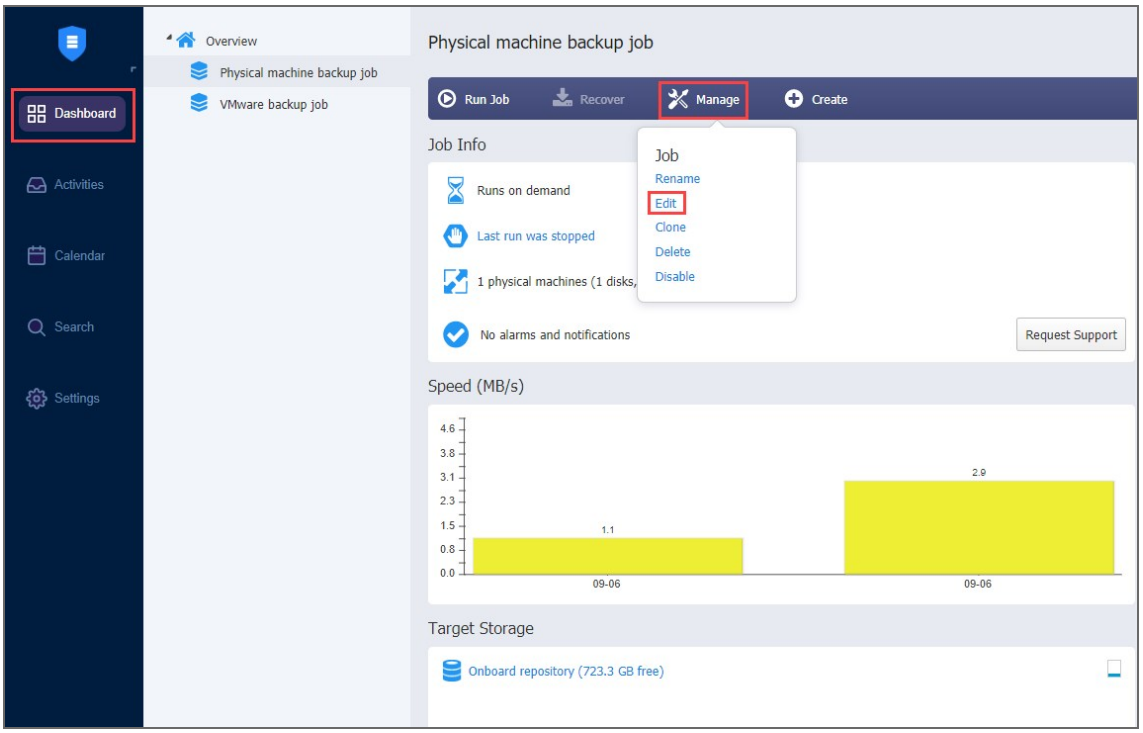

4. In the **Edit** wizard, click the necessary page to open it for editing.
5. Make the required changes and then click **Save** or **Save & Run**.

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**.

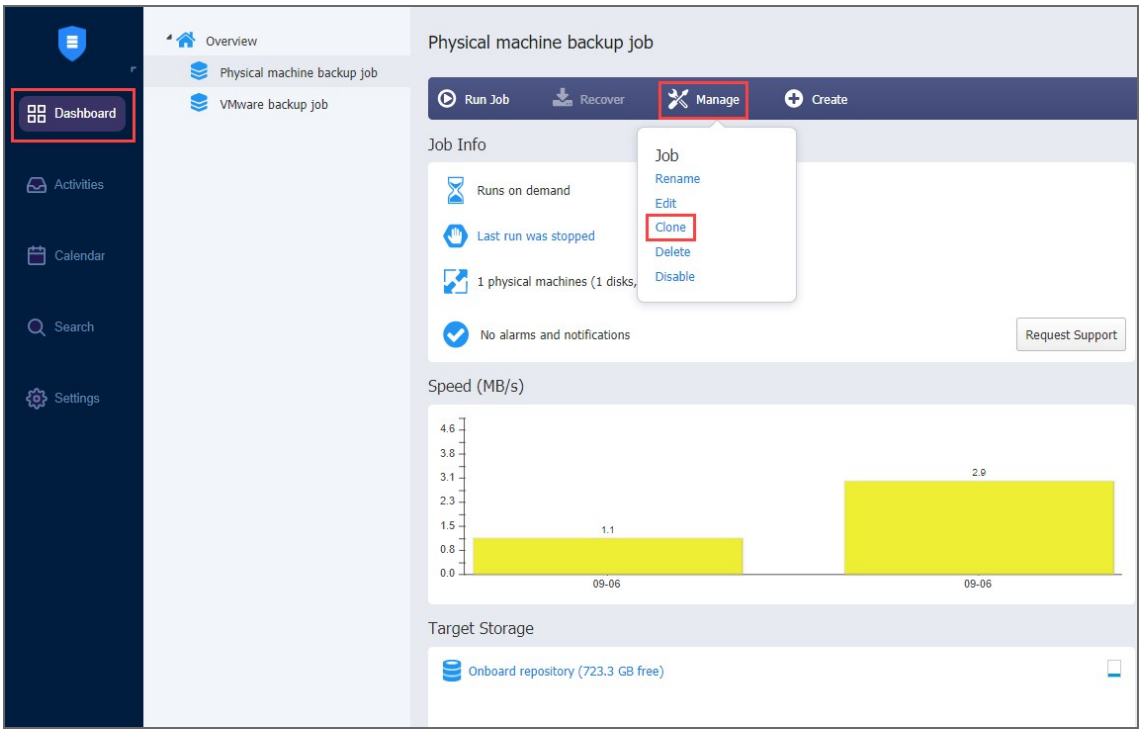

Note

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.

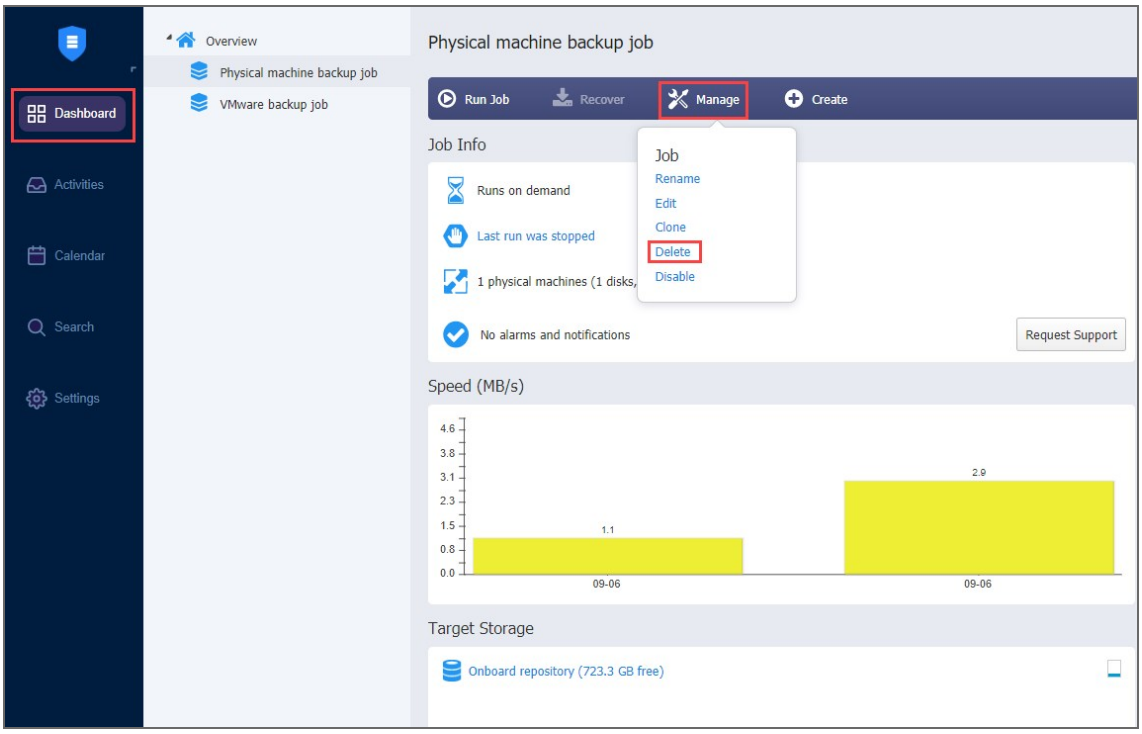

Disabling and Enabling Jobs

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**.

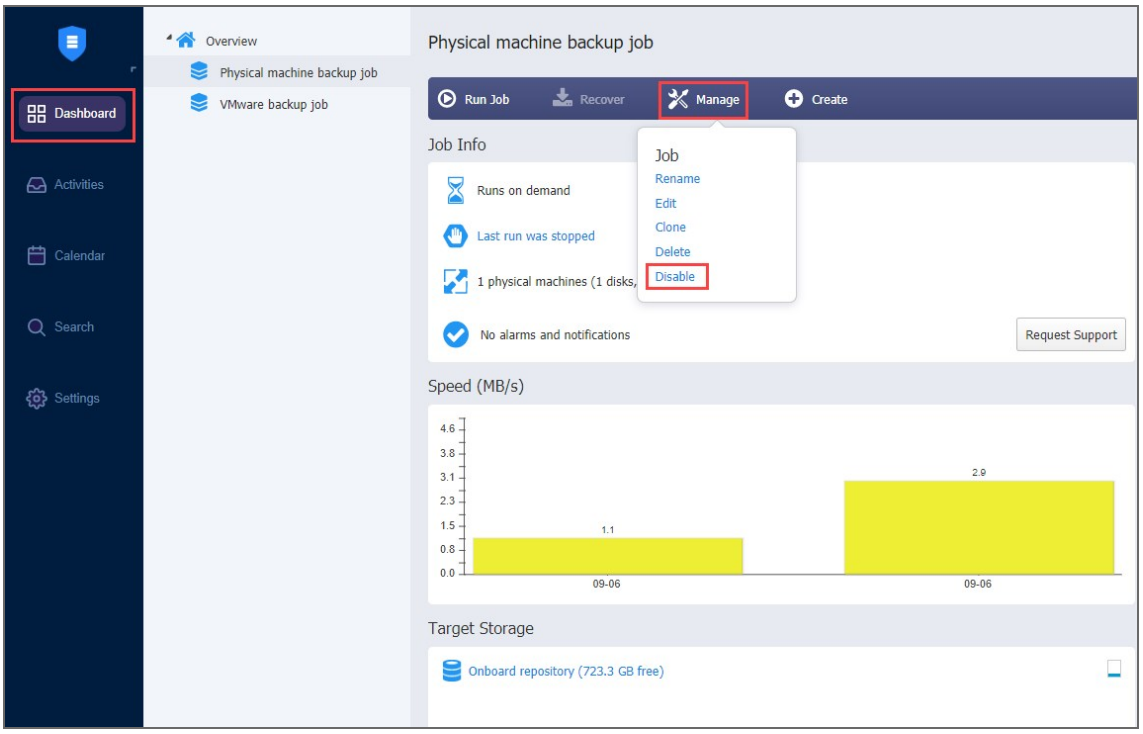

To enable a job, select **Enable** from the **Manage** menu.

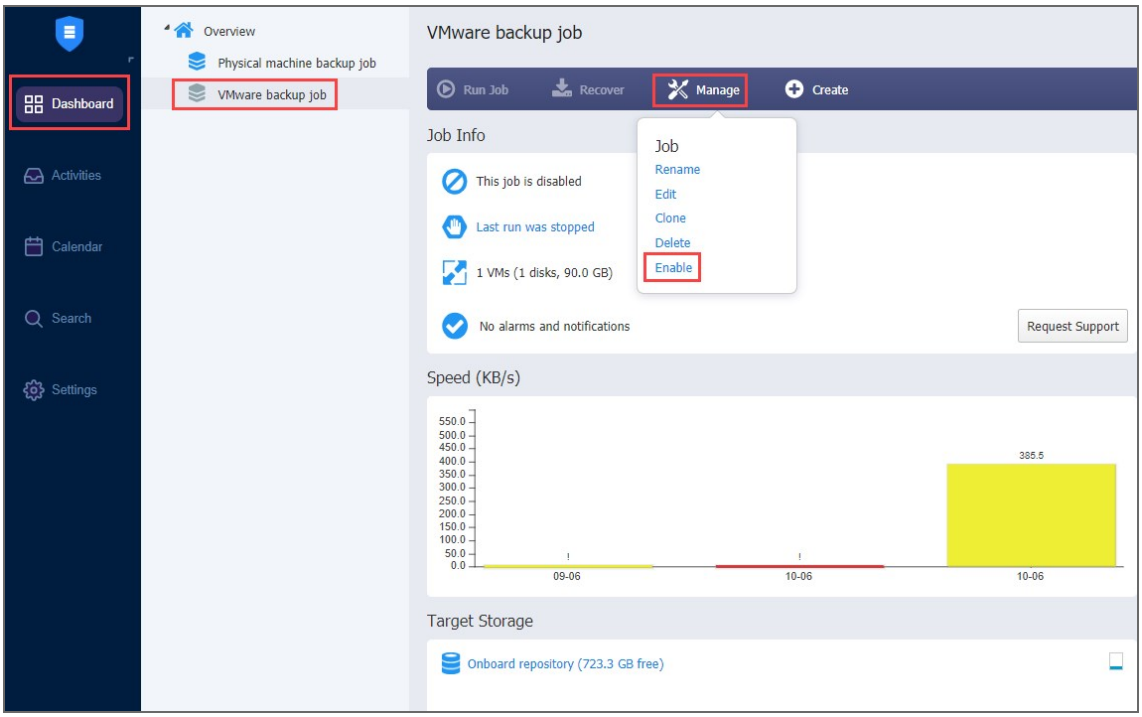

Note

You can also disable/enable jobs by right-clicking on a job and selecting **Disable/Enable** from the **Manage Job** menu.

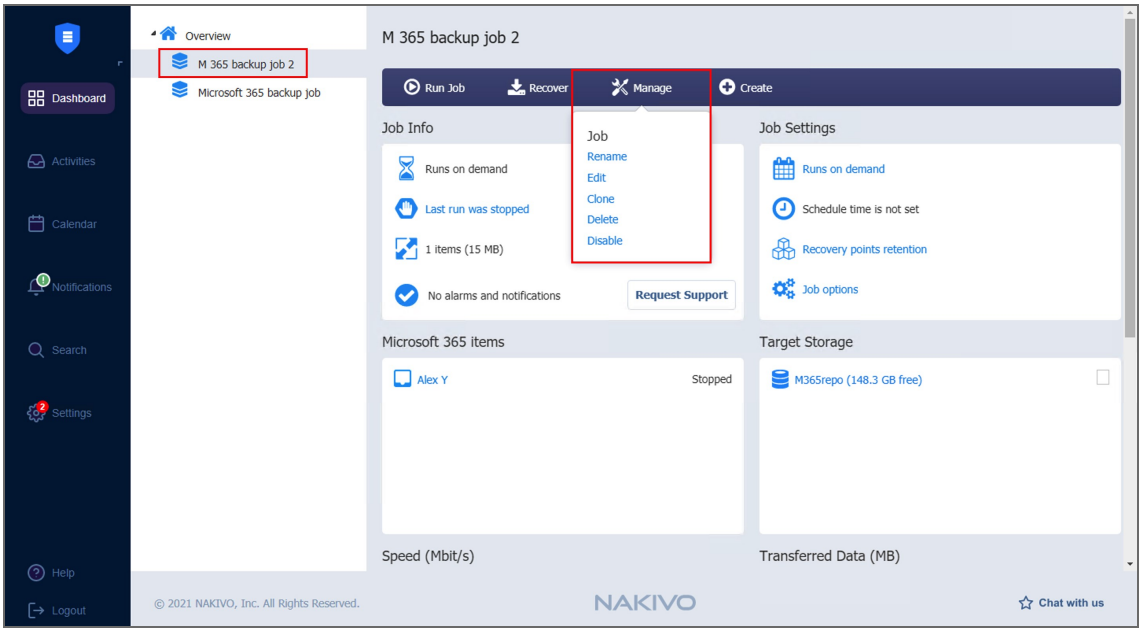

Grouping Jobs

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**.

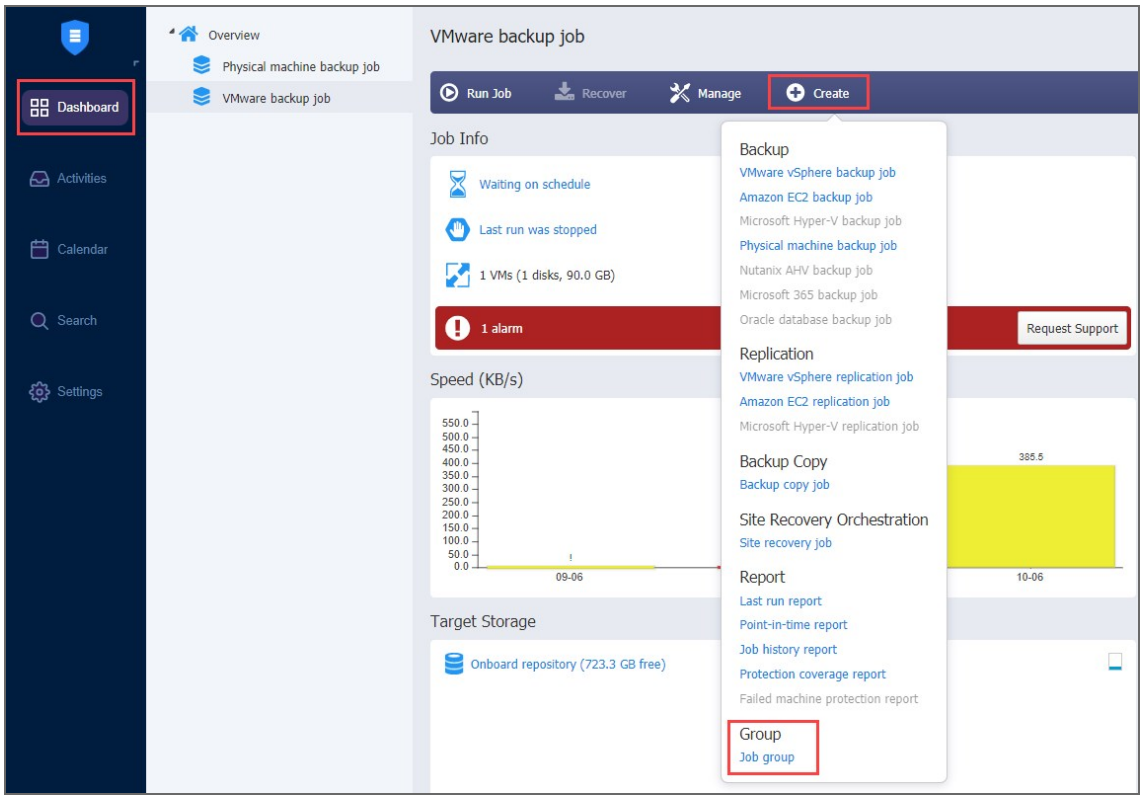

The following actions are available to manage groups:

- 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.
- **Protection coverage report**: Contains information about all VMs and instances protected by backup and/or 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 machine 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 machine protection report** and click **Create**.

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

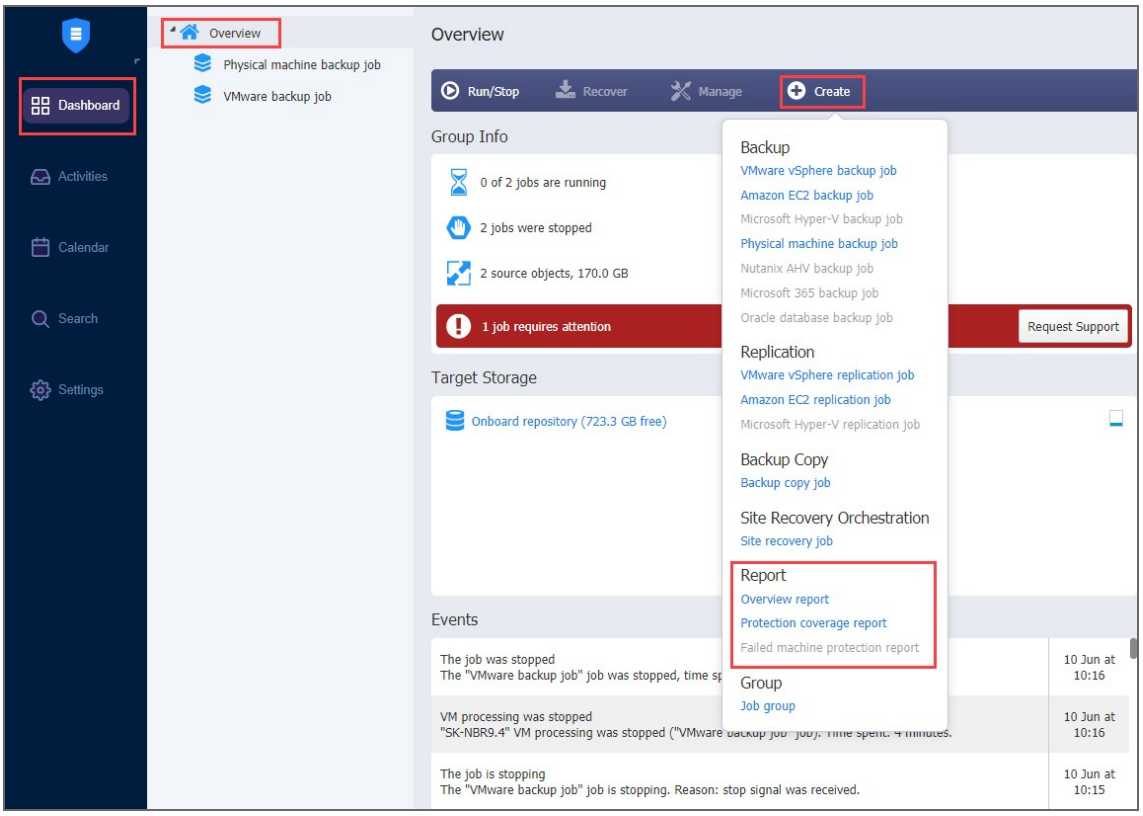

To generate reports from for an individual job, do the following:

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 popup that appears 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 in the popup that appears and click **Create**.
   - **Protection coverage report**:Contains information about all VMs and instances protected by backup and/or replication jobs, as well as about all unprotected VMs and instances.

- **Failed machine protection report**: Contains information about job objects processing of which were failed during the last job run. Only backup and replication jobs 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.

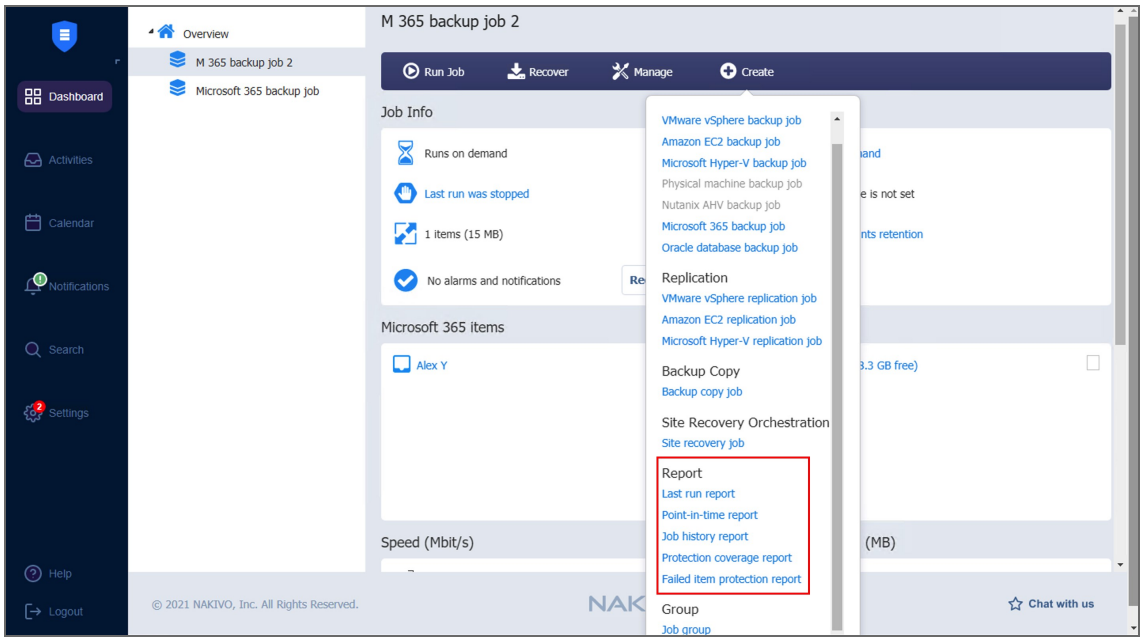

Managing Job Policies

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.

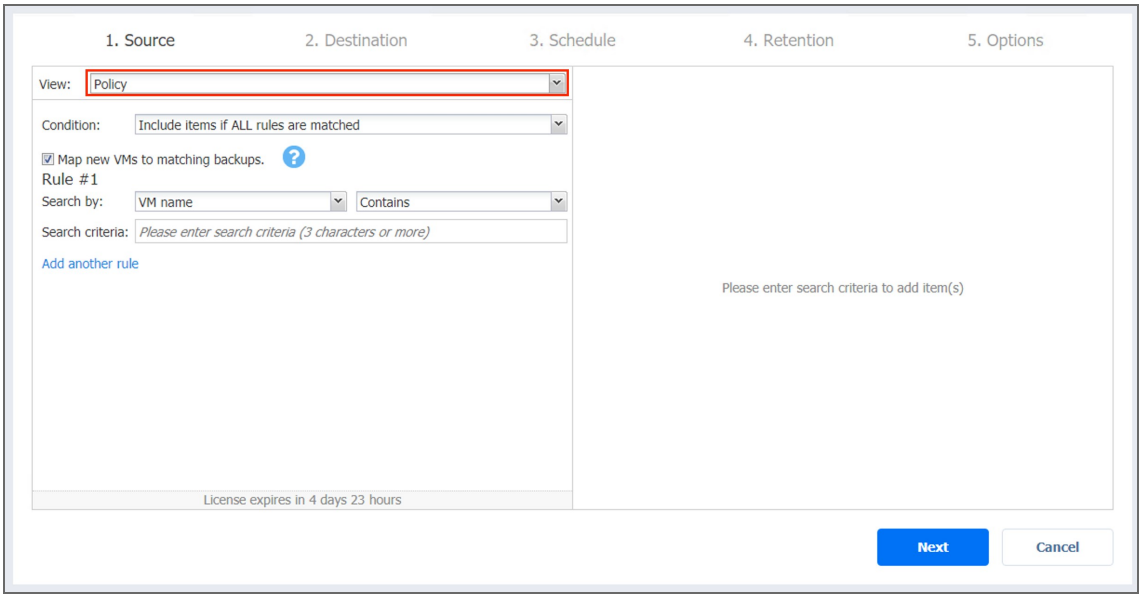

Every job policy contains at least one rule. Refer to "Managing Policy Rules" on page 265 for details.

Job policies are available for the following job types:

- Backup jobs
- Replication jobs
- Backup copy jobs
- Failover jobs
- Failback jobs
- Several actions of the Site Recovery job

Learn how to save, edit, and remove job policies in these sections:

- Saving Job Policy
- Editing Job Policy
- Removing Job Policy

Saving Job Policy

Follow the steps below to save a policy rule:

1. Make sure your job is opened in the **Policy** view.
2. Choose either of the following **Condition** for your job policy:

- **Include items if ALL rules are matched**: If selected, the logical AND will be applied to the set of policy rules.
- **Include items if ANY rule is matched**: If selected, the logical OR will be applied to the set of policy rules.

3. **Map new VMs/instances/machines to matching backups**: If the checkbox is selected, NAKIVO Backup & Replication maps new workloads, added to the job as compliant to the configured policy rules, to matching backups within the specified destination. This option is only available for VMware/Hyper-V/Amazon EC2/Physical machine backup jobs.
4. Provide the necessary policy rules. Refer to "Managing Policy Rules" on page 265 for details. Make sure that at least one item matches the available set of policy rules.
5. Save your job.

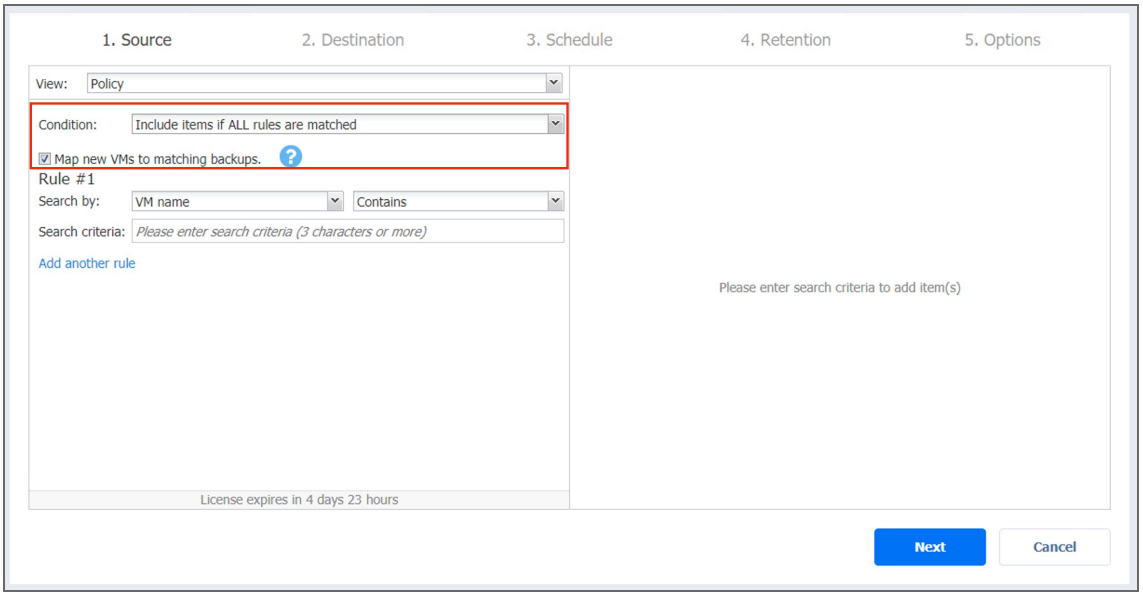

Editing Job Policy

Follow the steps below to edit a job policy:

1. Make sure your job is opened in the **Policy** view.
2. Change the necessary parameters of your job policy:
    1. **Condition**.
    2. Add, edit or delete policy rules. Refer to "Managing Policy Rules" on page 265 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.

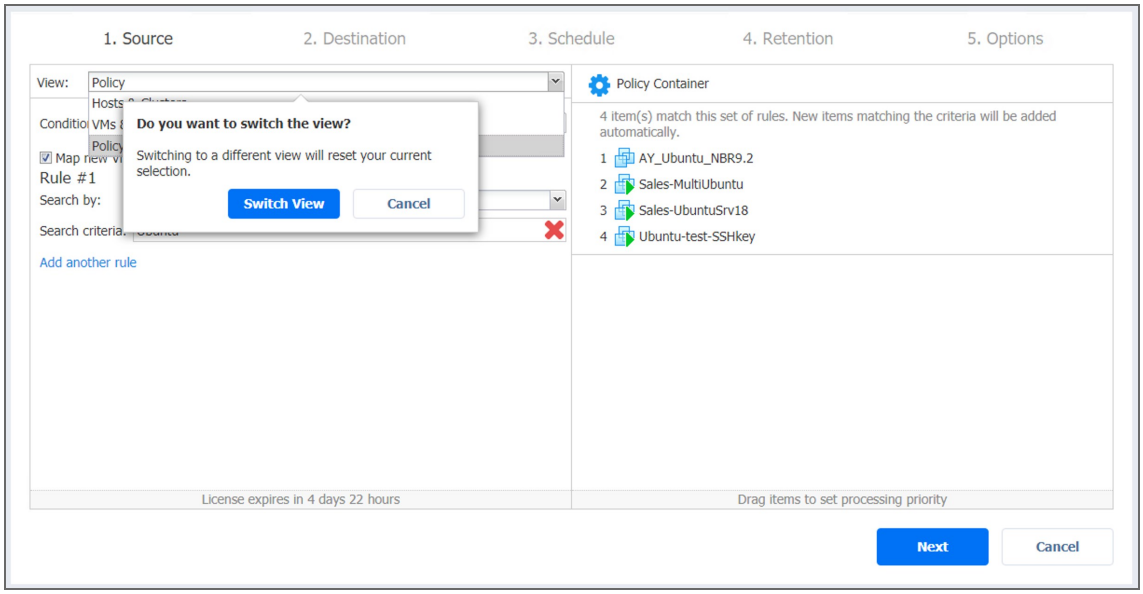

Managing Policy Rules

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 name / Instance name / Backup name**: The rule is to be applied based on the name of the object.
   - **VM tag / Instance tag**: The rule is to be applied based on the tag of the object.
   - **VM location / Instance location**: The rule is to be applied based on the location of the object.
   - **Name of VM datastore / VM Path**: The rule is to be applied based on the name of the VM datastore or the VM path.
   - **Name of VM network / Name of Subnet**: The rule is to be applied based on the name of the VM network or instance network.
   - **Size of VM / Size of instance**: The rule is to be applied based on the size of the object.
   - **Amount of VM/Instance RAM**: The rule is to be applied based on the amount of VM/Instance RAM.
   - **Number of VM CPU sockets / Number of VM processors / Number of Instance virtual CPUs**: The rule is to be applied based on the number of VM CPU sockets, VM processors, or Instance virtual CPUs correspondingly.

     Note

     The criteria are available as follows:
     - **Number of VM CPU sockets**: This is for your VMware backup and replication jobs.
     - **Number of VM processors**: This is for your Hyper-V backup and replication jobs.
     - **Number of instance virtual CPUs**: This is for your Amazon EC2 backup and replication jobs.
   - **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.

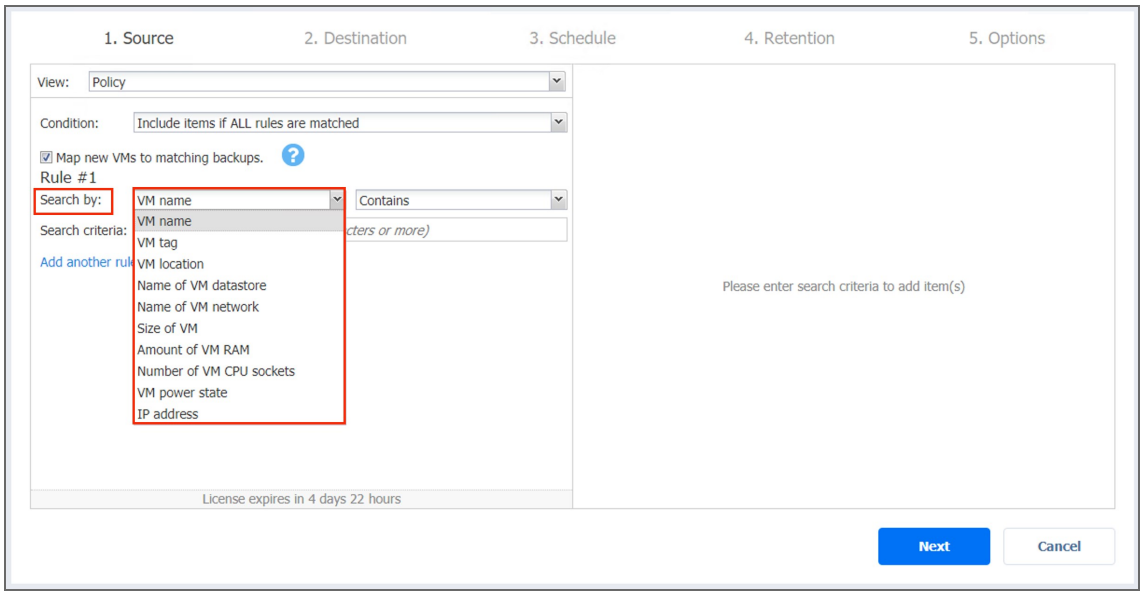

2. Search parameter: You can choose either of the following:
   - For **VM name / Instance name / Backup name, Name of VM network / Name of subnet, Name of VM datastore / VM Path, VM tag / Instance tag, and IP Address**:
     - **Contains**
     - **Does not Contain**
     - **Equals** (always applied to the VM tag)
     - **Does not equal**
     - **Starts with**
     - **Ends with**
   - For **Amount of VM/Instance RAM, Number of VM CPU sockets / Number of VM processors / Number of Instance virtual CPUs, and Size of VM / Size of Instance**, you can choose any of the following search parameters:
     - **Is more than**
     - **Is less than**
     - **Equals**
     - **Does not equal**
   - For **VM Power State / Instance power state and VM Location / Instance location**:
     - **Is**
     - **Is not**

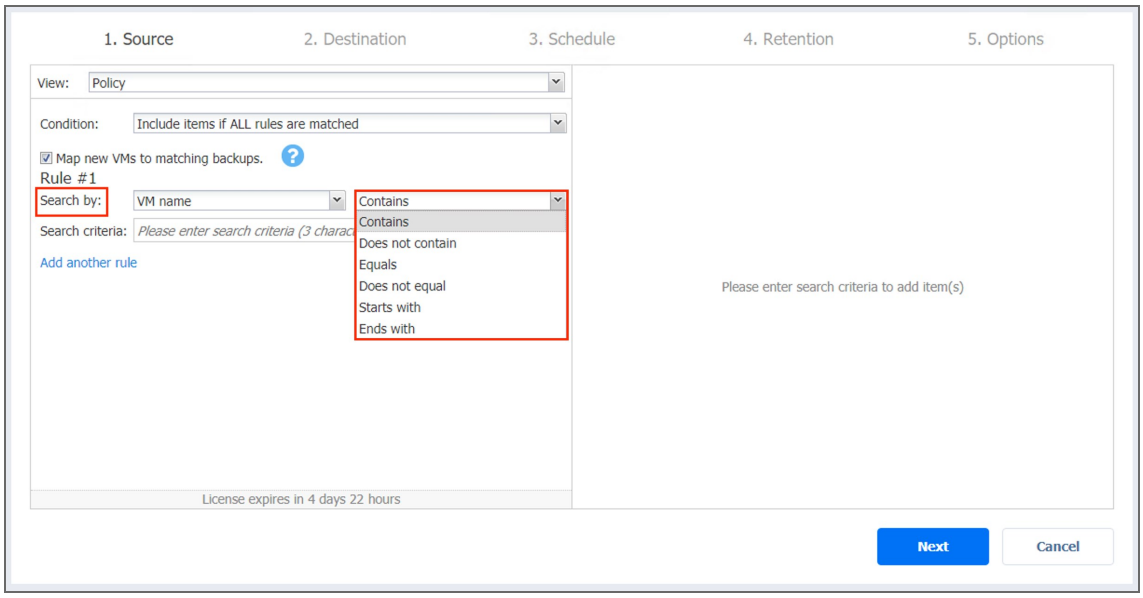

3. **Search criteria**: A text string or a numeric value to be used by the policy rule.

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 another rule**.
3. The wizard displays a new policy rule, *Rule #N*. Provide the necessary parameters of your policy rule. Make sure that at least one item matches the available set of policy rules.

4. Click **Next** when all parameters are set.

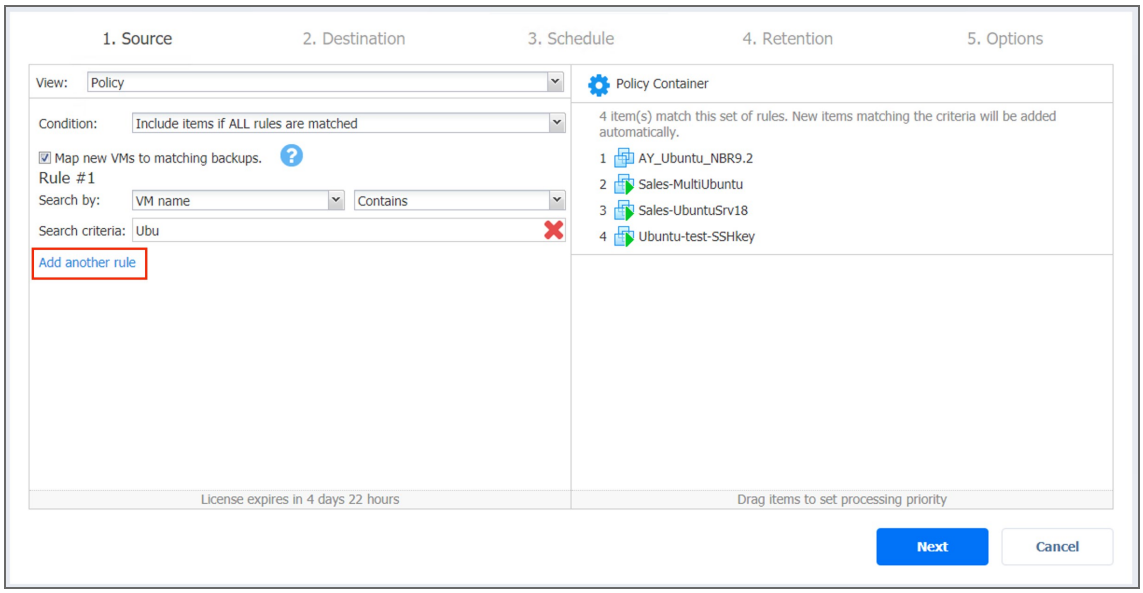

Removing Policy Rule

Follow the steps below to remove a policy rule:

1. Make sure your job is opened in the **Policy** view.
2. Locate your policy rule in the left pane of the view. If necessary, use the scroll bar.
3. Click **Remove**.
4. Click **Next** when all parameters are set.

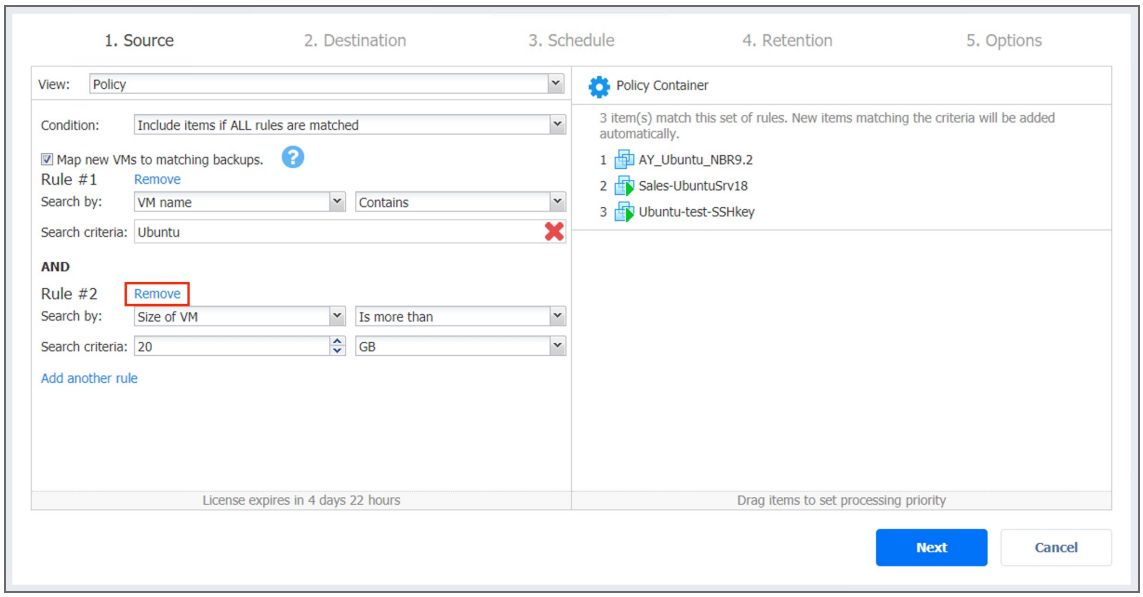

Note

You cannot remove all policy rules. A job policy must have at least one rule.

Jobs 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 box in the Job info widget.

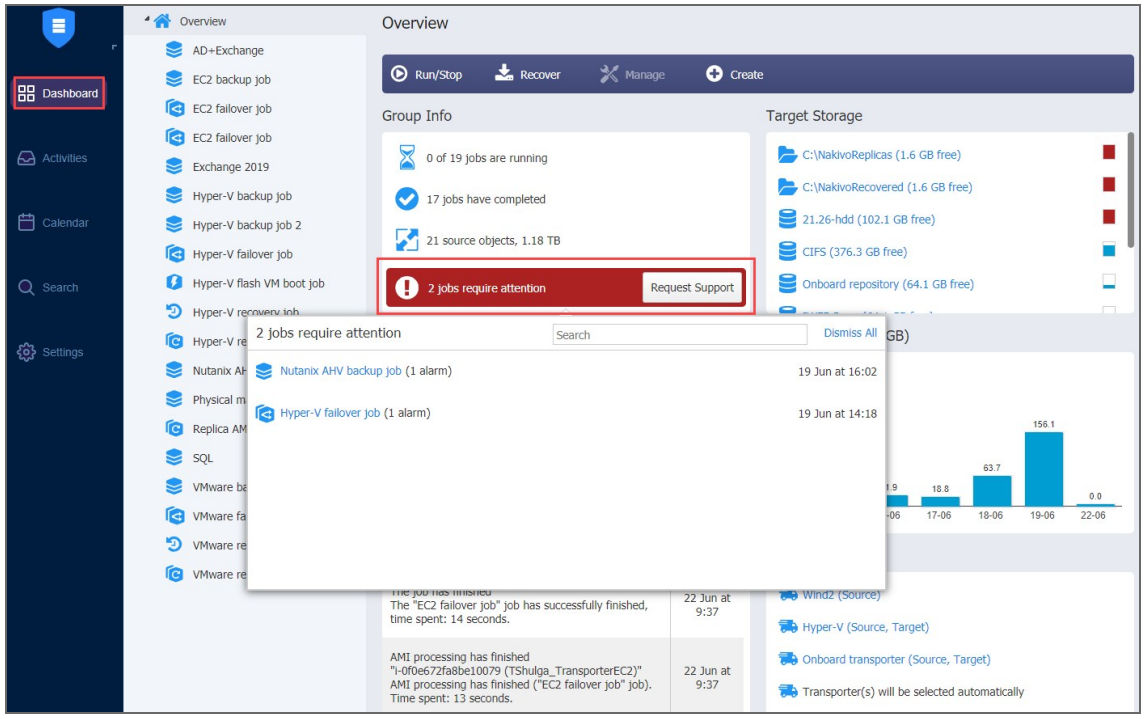

Dismissing Alarms and Notifications

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**.

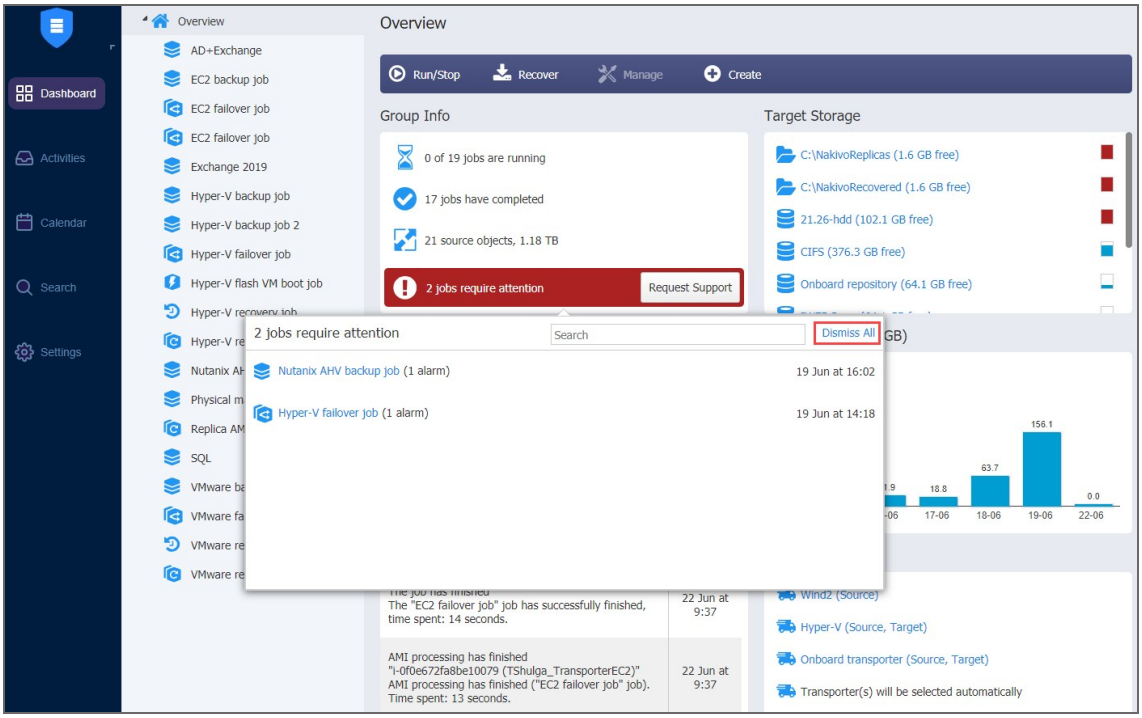
Overview
Dashboard
Activities
Calendar
Search
Settings
Overview
AD+Exchange
EC2 backup job
EC2 failover job
EC2 failover job
Exchange 2019
Hyper-V backup job
Hyper-V backup job 2
Hyper-V failover job
Hyper-V flash VM boot job
Hyper-V recovery job
Hyper-V re
Nutanix AH
Physical m
Replica AM
SQL
VMware ba
VMware fa
VMware re
VMware re
Overview
Run/Stop
Recover
Manage
Create
Group Info
0 of 19 jobs are running
17 jobs have completed
21 source objects, 1.18 TB
2 jobs require attention
Request Support
Target Storage
C:\NakivoReplicas (1.6 GB free)
C:\NakivoRecovered (1.6 GB free)
21.26-hdd (102.1 GB free)
CIFS (376.3 GB free)
Onboard repository (64.1 GB free)
2 jobs require attention
Search
Dismiss All
Nutanix AHV backup job (1 alarm)
19 Jun at 16:02
Hyper-V failover job (1 alarm)
19 Jun at 14:18
GB)
156.1
63.7
18.8
1.9
0.0
17-06
18-06
19-06
22-06
The "EC2 failover job" job has successfully finished, time spent: 14 seconds.
22 Jun at 9:37
AMI processing has finished "i-0f0e672fa8be10079 (TShulga_TransporterEC2)" AMI processing has finished ("EC2 failover job" job). Time spent: 13 seconds.
22 Jun at 9:37
Win2 (Source)
Hyper-V (Source, Target)
Onboard transporter (Source, Target)
Transporter(s) will be selected automatically

Managing Activities

The **Activities** page displays current and past tasks performed by NAKIVO Backup & Replication. From this dashboard, the following actions can be done:

- Viewing Activities
- Searching for Activities
- Viewing Activity Details
- Stopping Running Activities
- Running Activities Again
- Removing Activities

Past activities are stored for the number of days specified in the **Store job history for the last X days** setting in the General tab.

Viewing Activities

The **Activities** dashboard allows viewing all your current and past activities in the application.

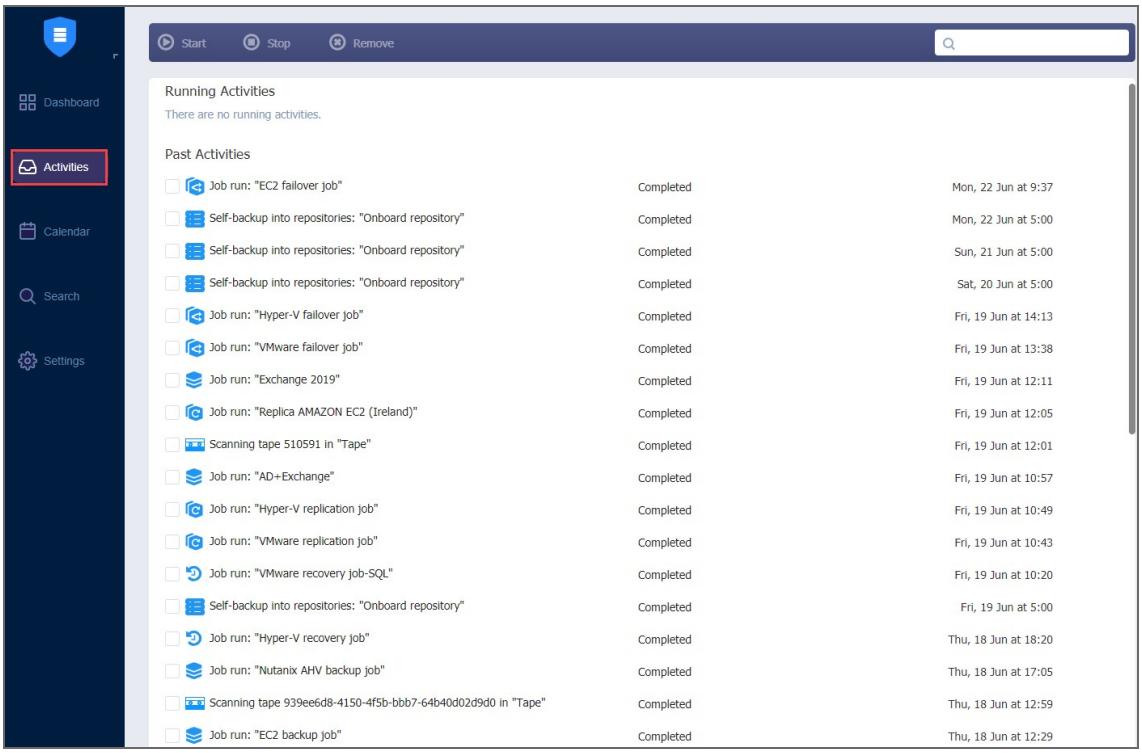

Searching for Activities

Find activity by typing in part of its name in the **Search** field.

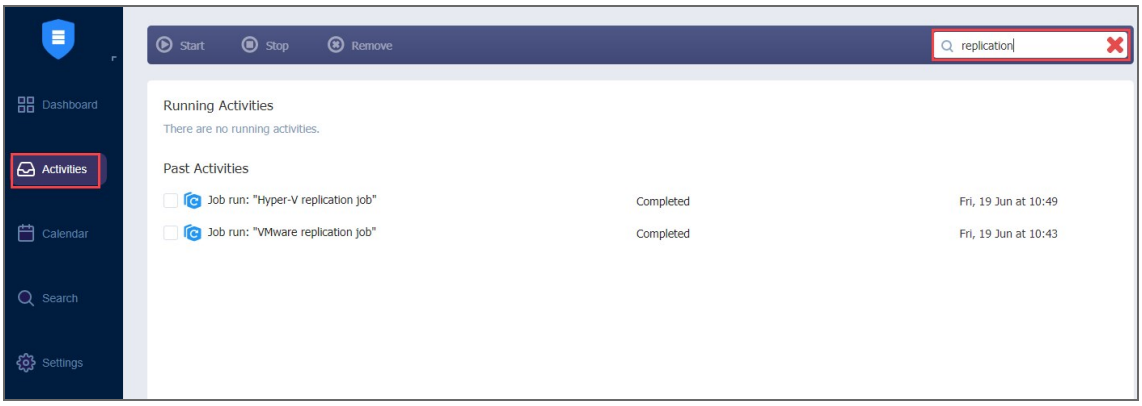

Viewing Activity Details

View the details of an activity by selecting an activity name.

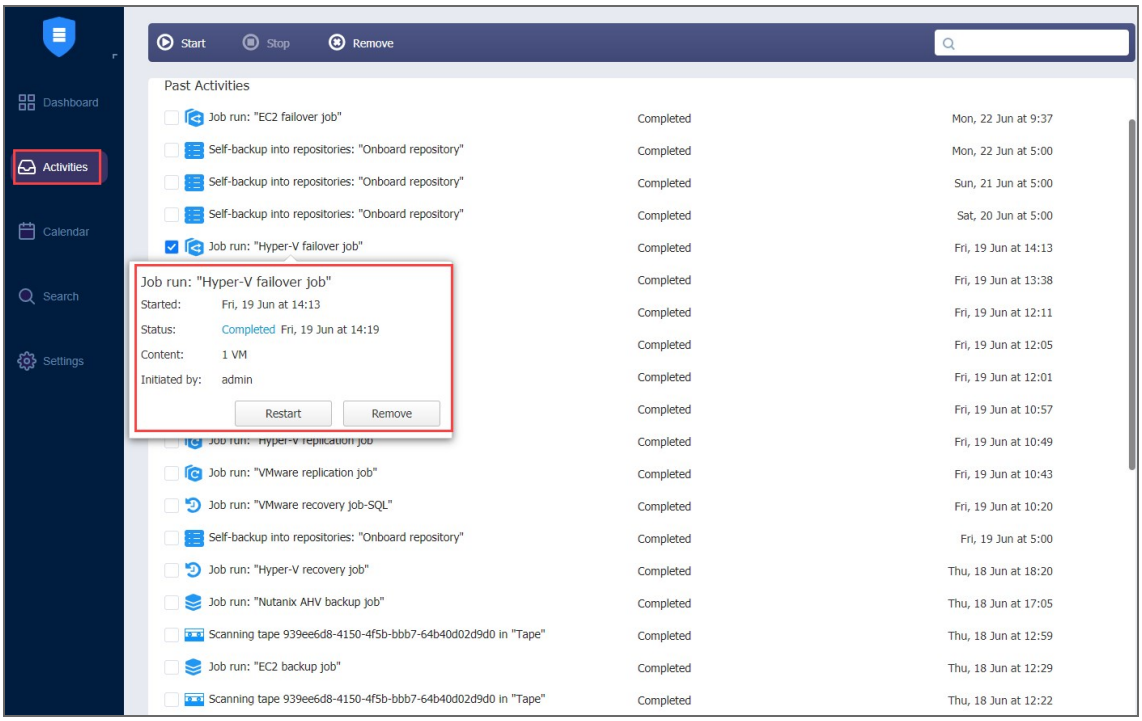

Stopping Running Activities

To stop a running activity, select the checkbox next to the activity and click **Stop** in the toolbar at the top. Optionally, you can stop multiple activities by selecting the checkboxes next to them or stop all running activities by clicking **Select/Deselect all**.

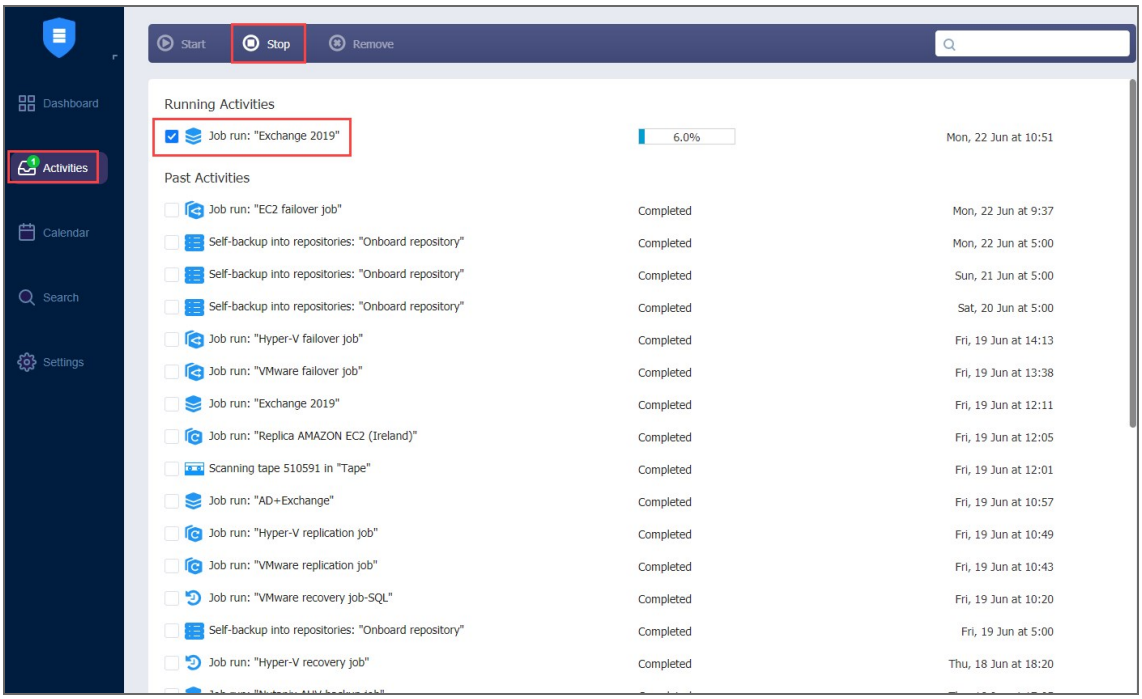

Running Activities Again

To run an activity again (if possible), select the checkbox next to the activity and click **Start** in the toolbar at the top. Optionally, you can run multiple activities by selecting the checkboxes next to them or run all stopped activities by clicking **Select/Deselect all**.

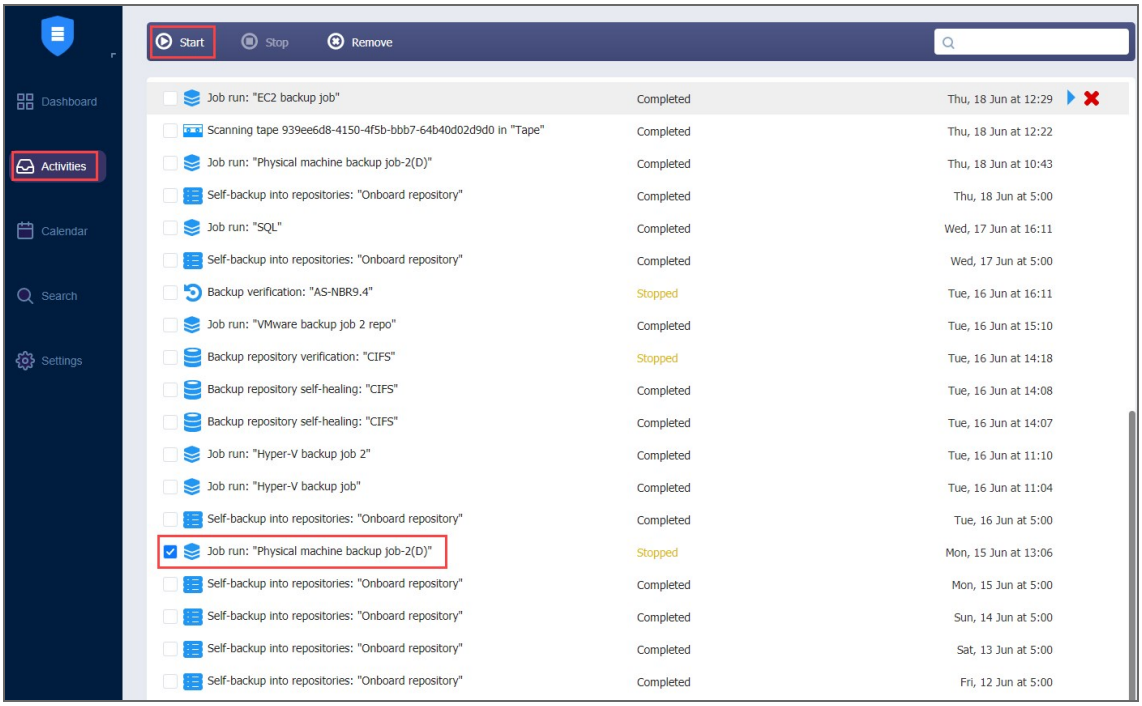

Removing Activities

Remove an activity from the list by selecting the checkbox next to the activity and clicking **Remove** in the toolbar.

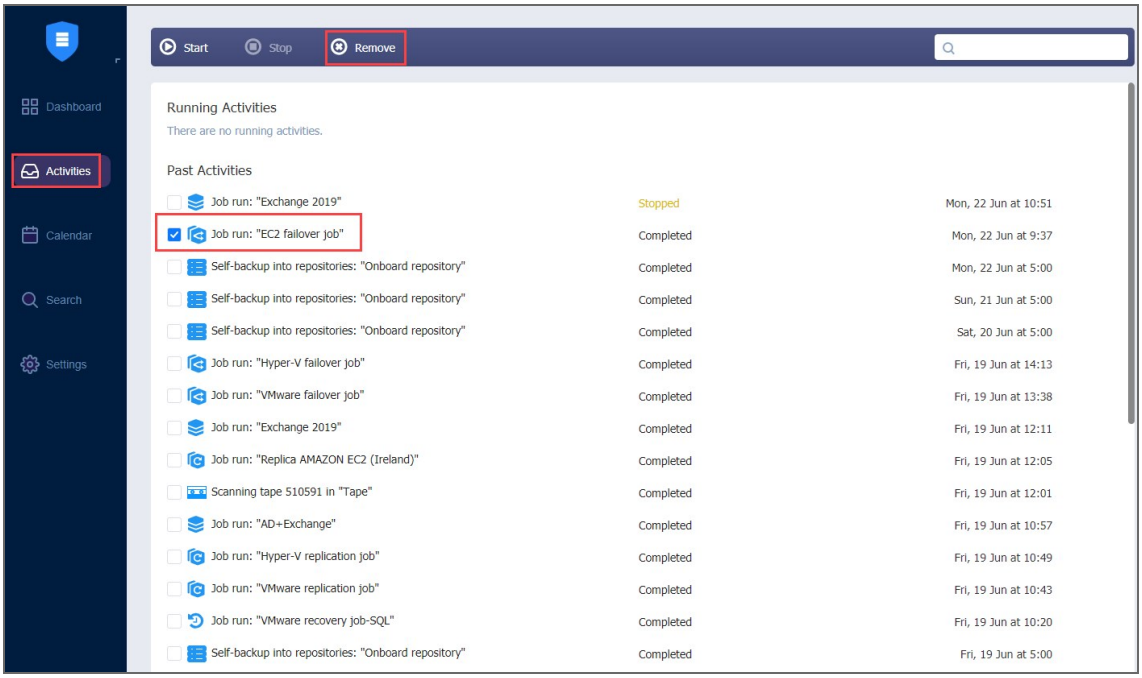

Using Calendar

The **Calendar** allows you to schedule and view the history of past job runs.

- Creating Jobs with Calendar
- Editing Jobs with Calendar

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 overriding the schedule.
- Edit a job.
- Clone a job
- Delete the job. If the job is repeated, this action will affect all job runs.
- Disable/Enable a job. If the job is repeated, this action will affect all job runs.
- Open the Dashboard.
- Create a report.

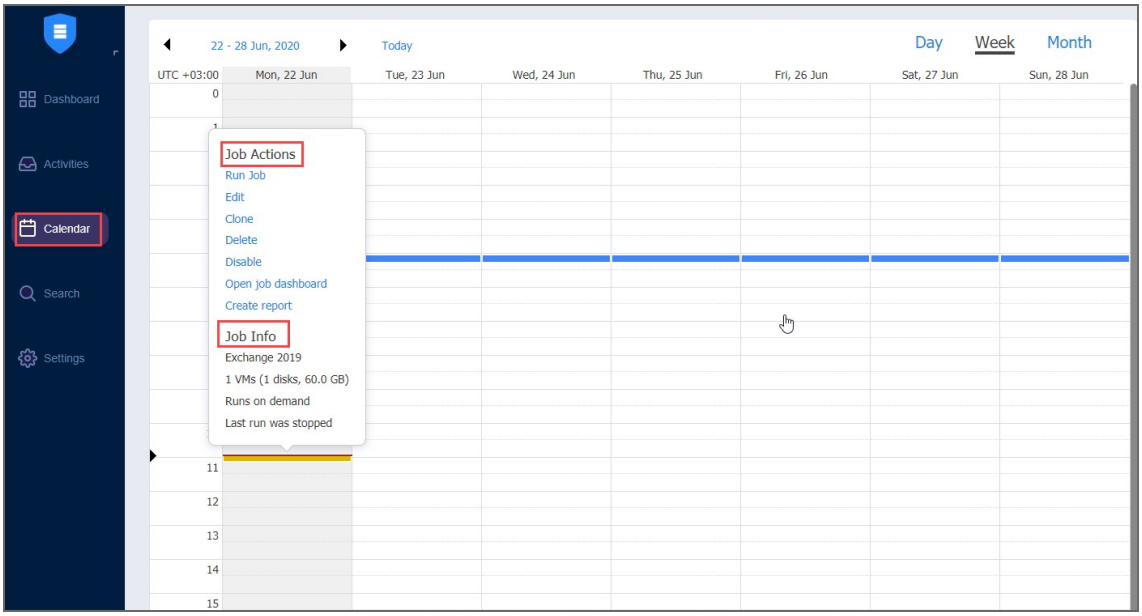

Using Global Search

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-Tenancy” on page 79
- “Multi-Tenant Mode” on page 769

Opening Global Search

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

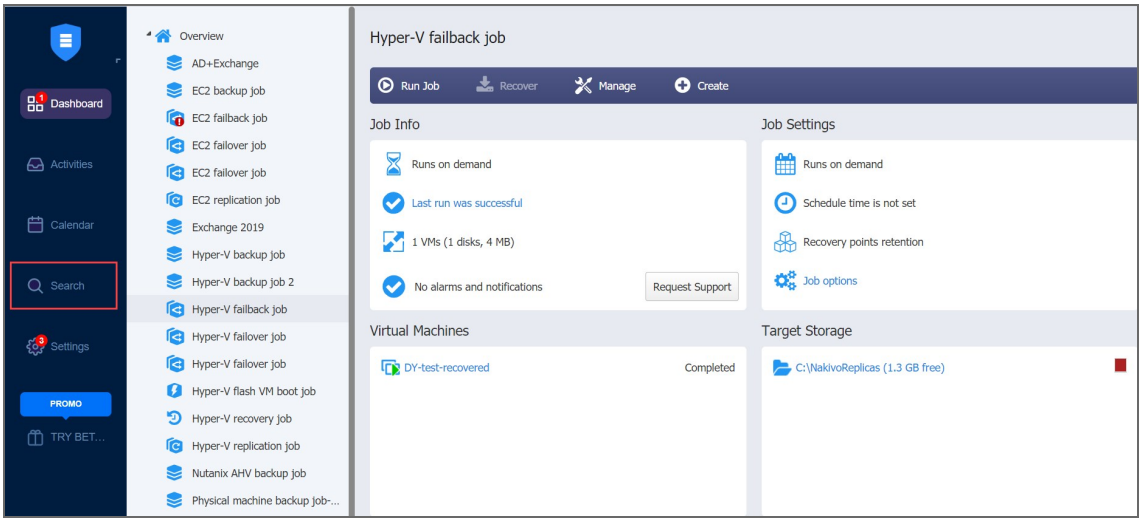

Running Global Search

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.

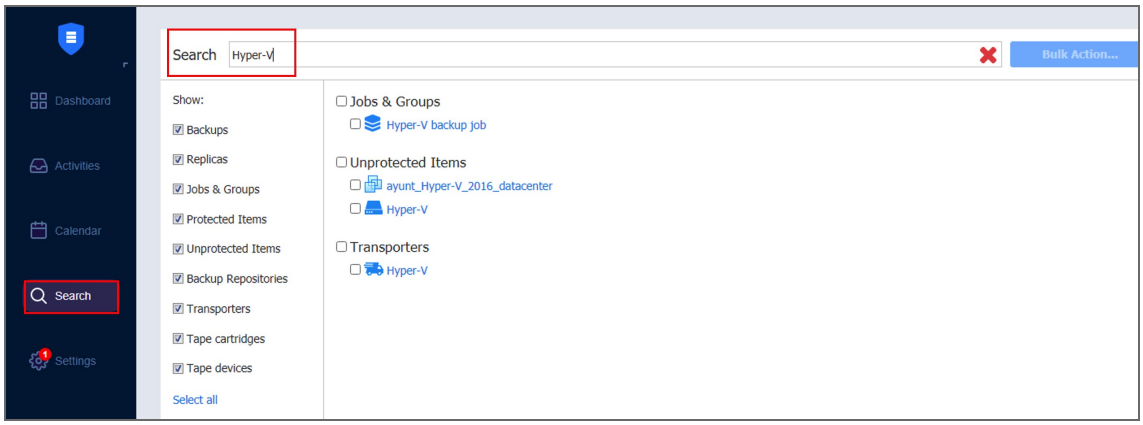

Please note the following:

- 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.

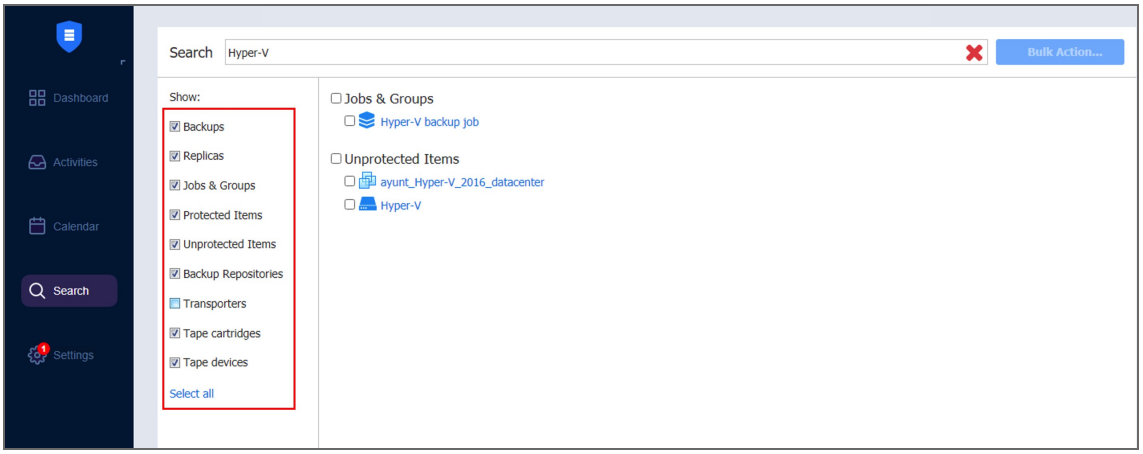

To get back to the default filtering settings, click **Select all** below the categories list.

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**.

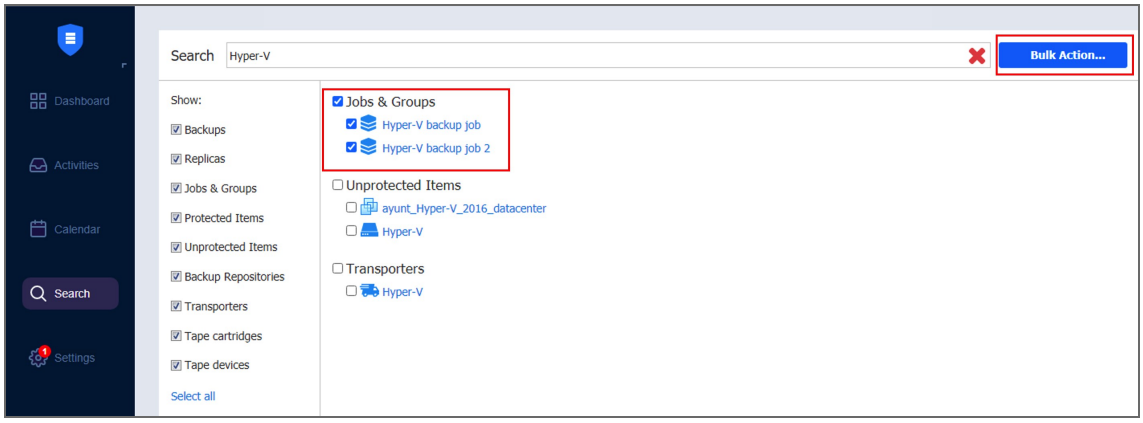

A dialog opens with the list of actions applicable to the selected items. To proceed with the necessary action, click the corresponding item in the list of actions.

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.

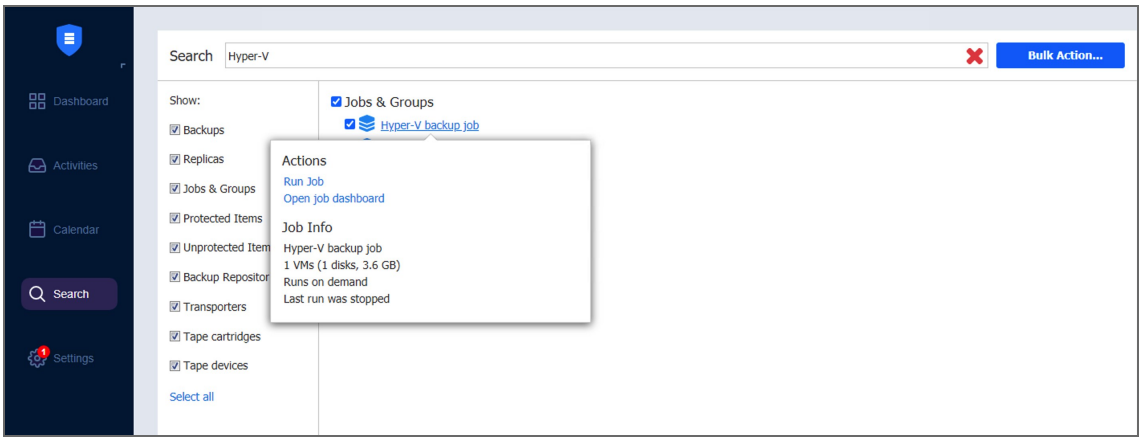

A dialog opens displaying object info, along with the list of typical actions applicable to the object.

Settings

This section covers the following topics:

- "General" on page 280
- "Inventory" on page 329
- "Transporters" on page 344
- "Backup Repositories" on page 369
- "Tape" on page 432
- "Replacing License" on page 481
- "Expert Mode" on page 482
- "Multi-Tenant Mode Configuration" on page 463
- "Support Bundles" on page 472
- "Built-in Support Chat" on page 474

General

This section contains the following topics:

- "Bandwidth Throttling" on page 281
- "Branding Configuration" on page 284
- "Email Notifications" on page 287
- "Configuring Events" on page 286
- "Self-Backup Configuration" on page 291
- "System Migration" on page 295
- "System Settings" on page 297
- "Users and Roles" on page 327

Bandwidth Throttling

With bandwidth throttling settings, you can configure how LAN/WAN bandwidth is consumed by NAKIVO Backup & Replication jobs. Bandwidth consumption 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**.

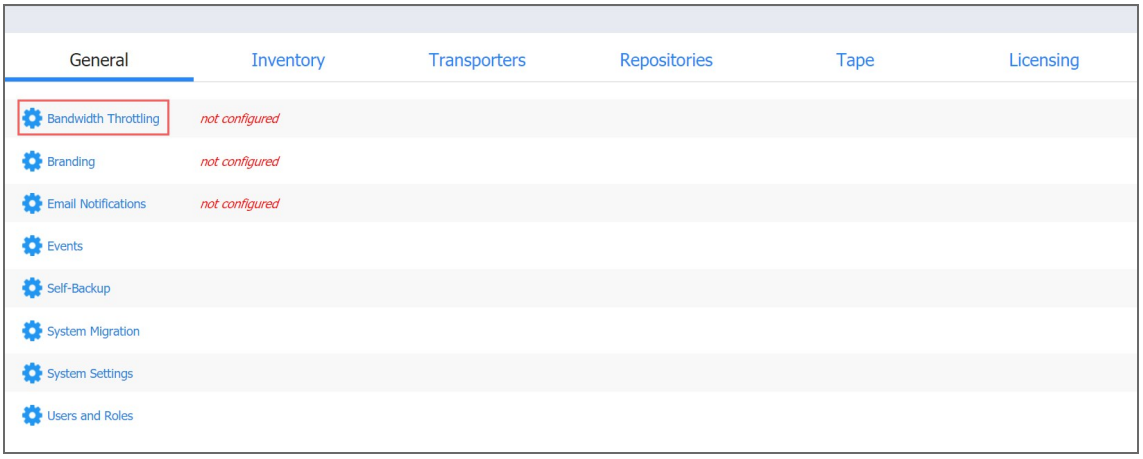

The *Bandwidth throttling* section opens.

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 **Create New Rule**.
2. The New Bandwidth Rule dialog opens. Proceed as follows:

a. Choose a type for your bandwidth rule:
   - **Global**: The rule will be applied to the whole application.
   - **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. **Name**: Enter a name for your bandwidth rule.

c. **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.

d. **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.

e. **Show/Hide jobs**: This option is available for the **Per job** rule type only. Select jobs for which the rule will be active.

3. Click **Save**.

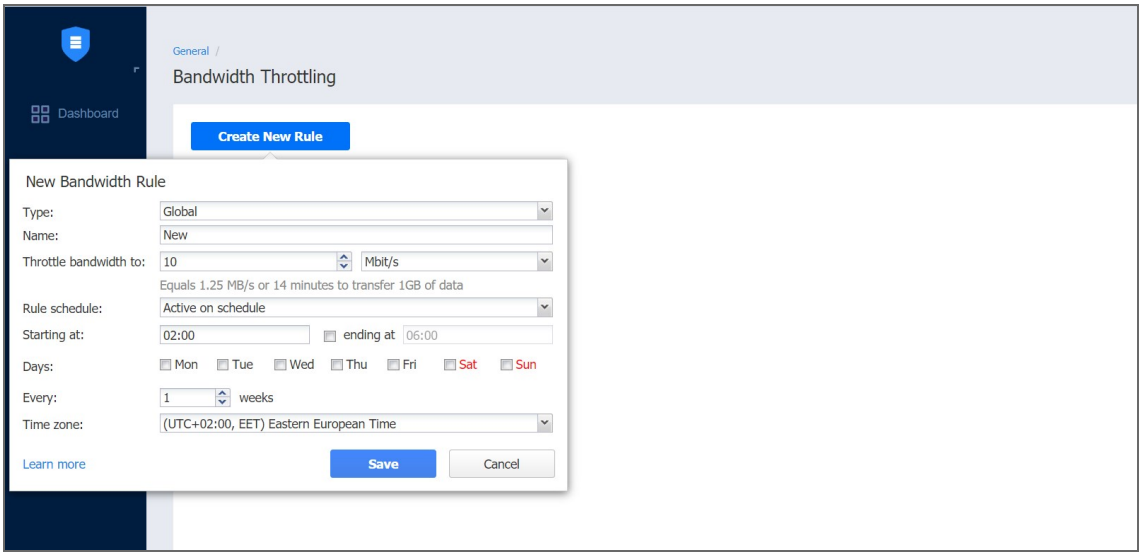

Managing Bandwidth Rules

You can 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.

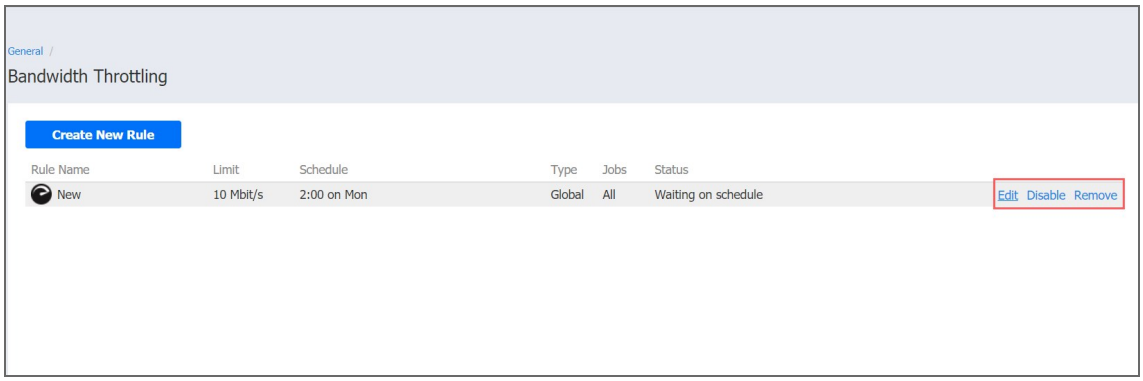

Note

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 Hyper-V Backup Jobs” on page 528
- “Creating Backup Copy Jobs” on page 552
- “Creating Hyper-V Replication Jobs” on page 497
- “Hyper-V VM Recovery” on page 665

Branding Configuration

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**.

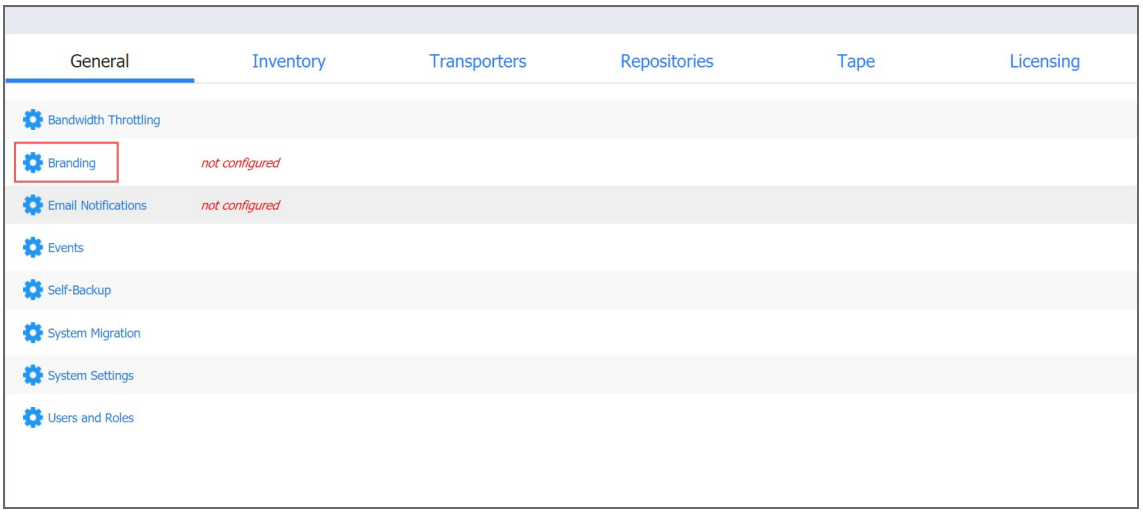

3. Change the following, as appropriate:
   - Product title
   - Company name
   - Website URL
   - Contact email
   - Support email
   - Contact phone
   - Bookmark icon
   - Global logo
   - Page background

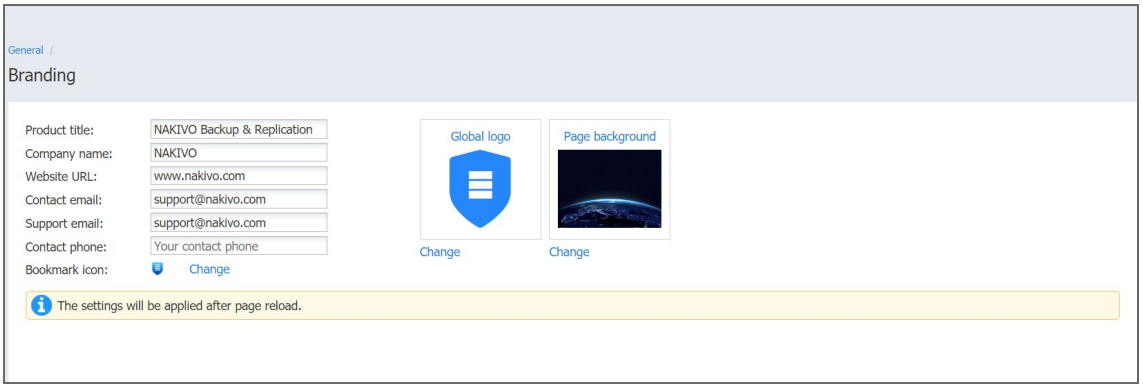

4. Click **Apply**.

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	40x40
Page background	.jpeg	1920x1440
Bookmark icon	.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 **Show only warnings and alarms**. If selected, the events list displays only warnings and alarms.
5. Optionally, you can select **Filter by date** and enter the beginning and ending dates for events filtering. This allows you to limit the events list within a specific time period.
6. If required, navigate between pages by using the **Page** control.

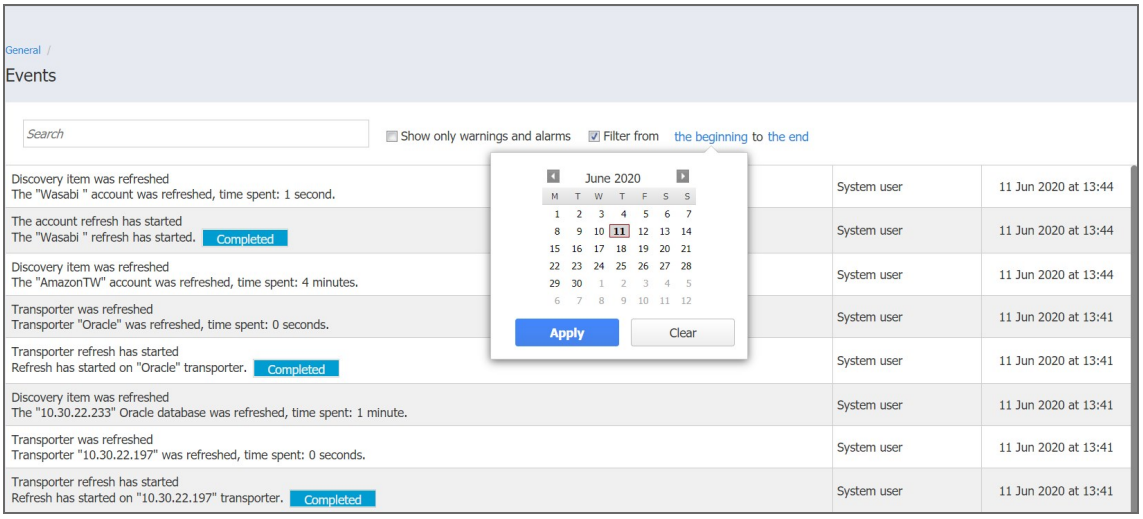

Email Notifications

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

- Email Settings
- 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 **Email notifications** to configure email settings, email notifications, and automatic reports section on the page that opens.

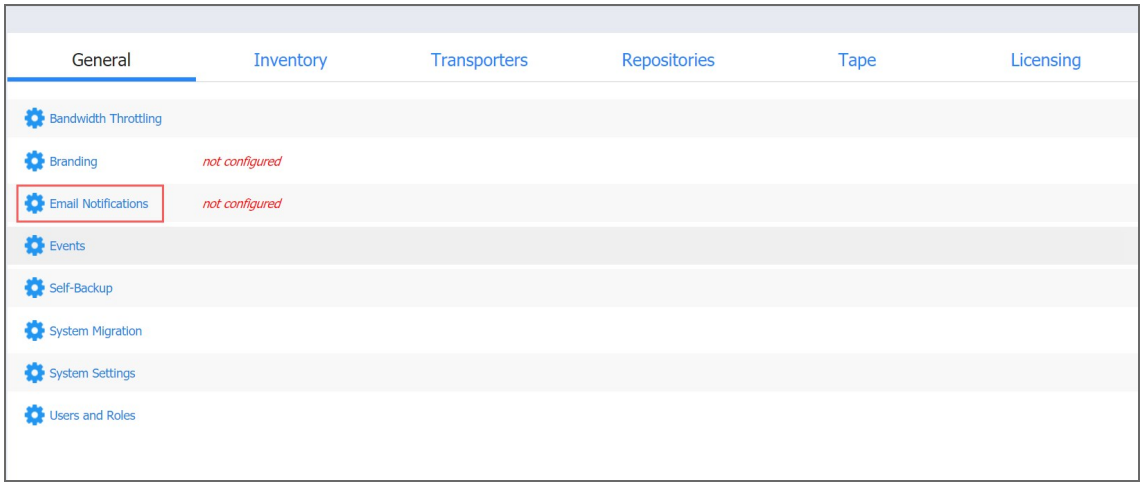

Email Settings

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.
   - **Encrypted connection**: Select the type of encryption:

- **Never**: 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.

2. 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.



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.
- **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).
- **Limit email notification frequency**: This option provides you with the ability to set up a notification email frequency and hourly limit. If notification emails exceed the hourly limit, all new notifications will be delivered the next hour. If deselected, notification emails will be sent every 5 minutes with no hourly limit.
- **Email notification recipients**: Specify the recipients who will be receiving alarm and warning notifications (if

enabled).

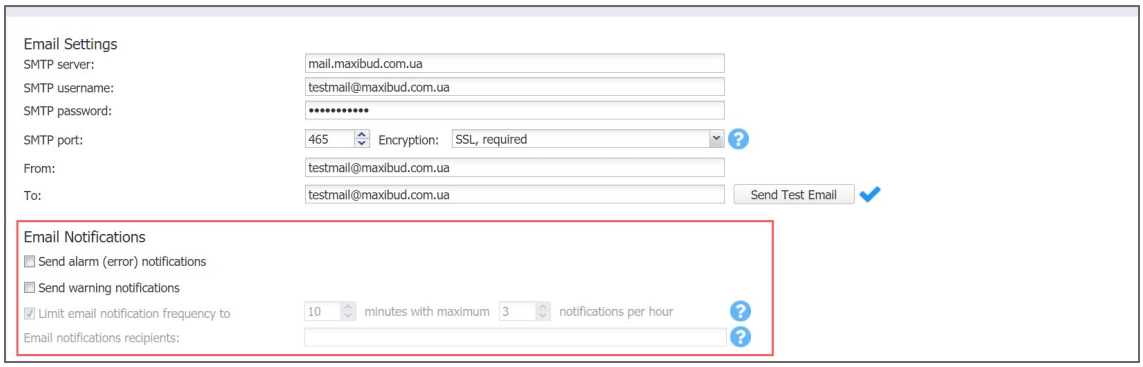

Automatic Reports

To set automatic reports, fill out the fields in the *Automatic Reports* section:

- **Attach PDF copy to all automatic reports**: If selected, a PDF copy of the report will be attached to each automatic job report and to the Overview report. Note that this increases the size of email messages.
- **Send job reports on each job completion**: 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.
- **Send Overview report on schedule to**: 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.
- **Send Protection Coverage report on schedule to**: 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.

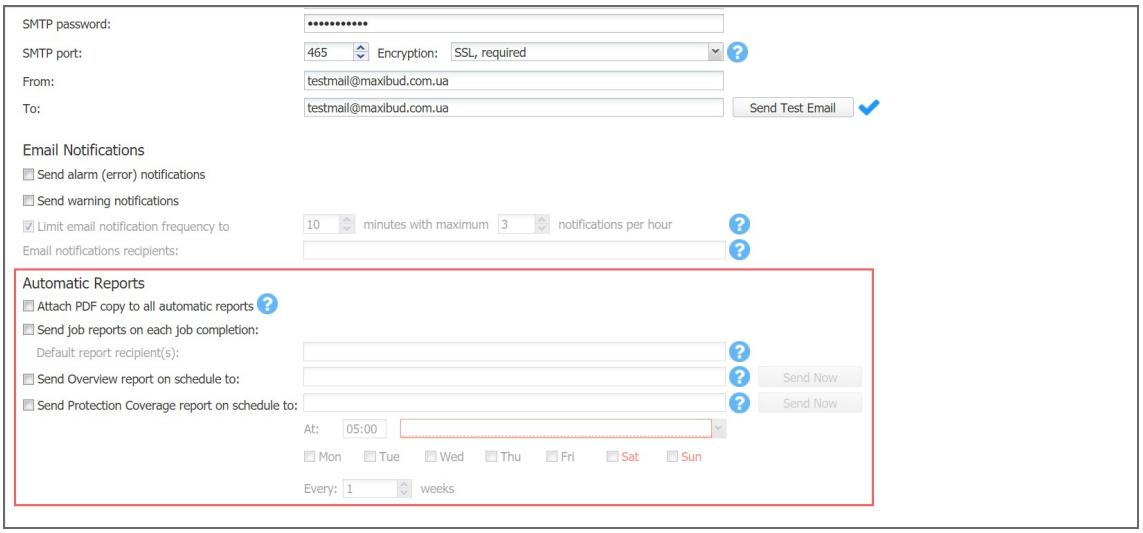

Click **Apply** when all settings are configured.

Self-Backup Configuration

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" below
- "Setting Up Self-Backup Destination" below
- "Self-Backup Schedule" on page 293
- "Self-Backup Retention" on page 293
- "Recovering from Self-Backup" on page 294

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**.

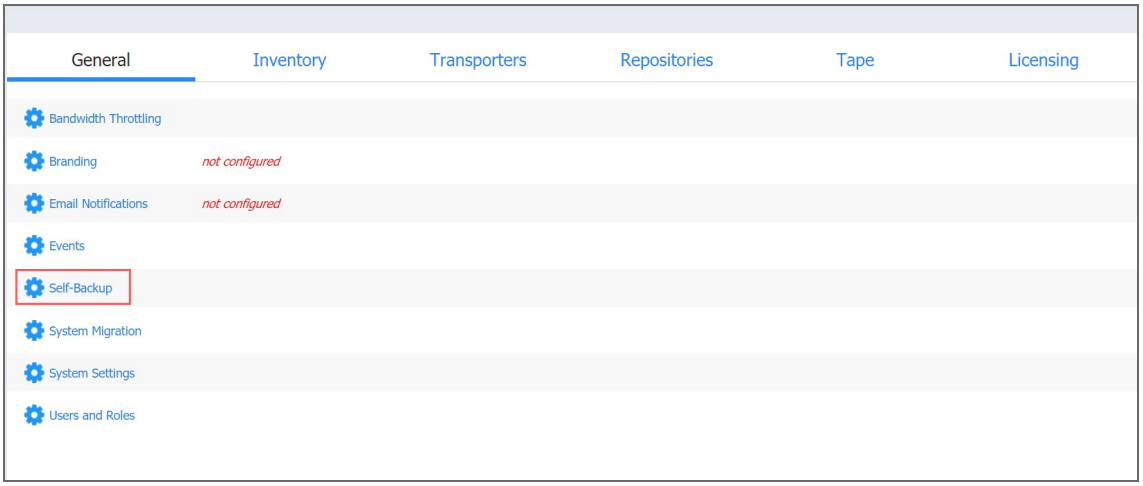

Setting Up Self-Backup Destination

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.

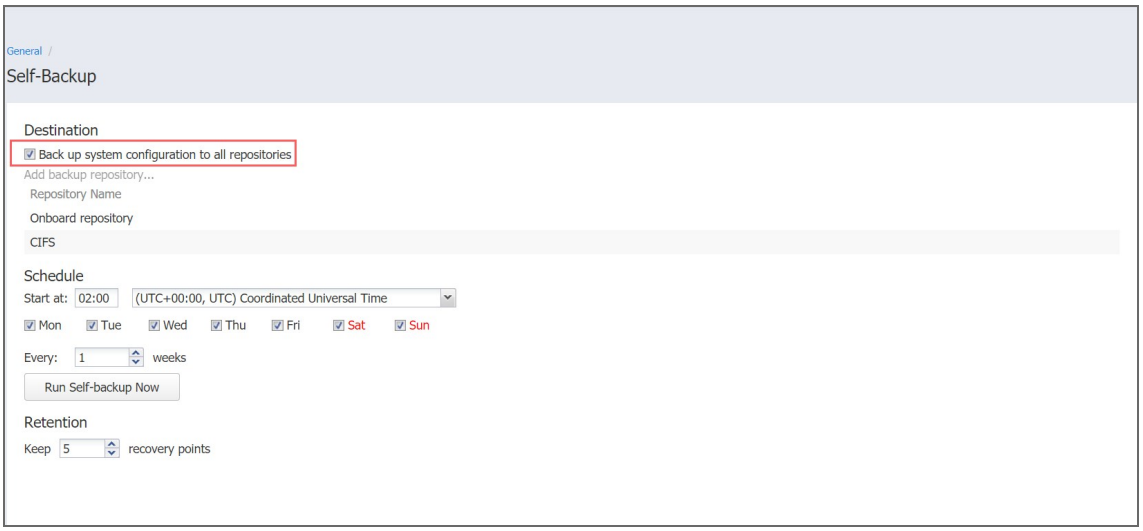

2. If necessary, remove a Backup Repository from the list of repositories for self-backup:
    1. Hover the pointer over the header of the Backup Repository and then click **Remove**.
    2. In the dialog that opens, choose either of the following:
        - **Remove repository and keep self- backups**: Removes the selected Backup Repository from the list and keeps the self-backups.
        - **Remove repository and self-backups**: Removes both the selected Backup Repository and self-backups.
    3. Click **Remove** to confirm your operation.

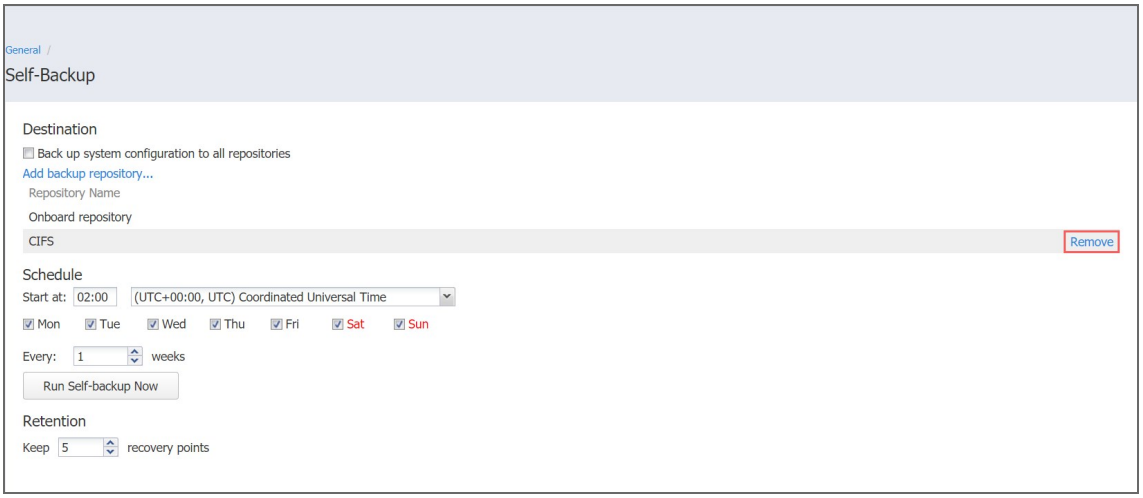

3. If necessary, add a Backup Repository to the list:

a. Click **Add backup repository** to add repositories to the list of repositories for system backing up.

b. In the **Backup Repositories** dialog that opens, select the necessary repositories and close the dialog.

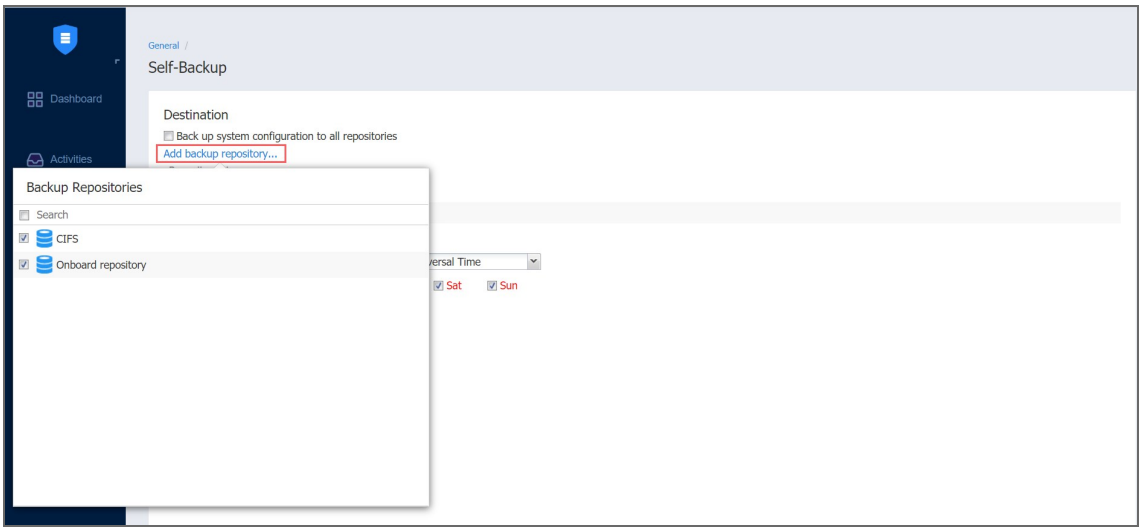

4. When ready with configuring the self-backup destination, click **Apply**.

Self-Backup Schedule

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

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

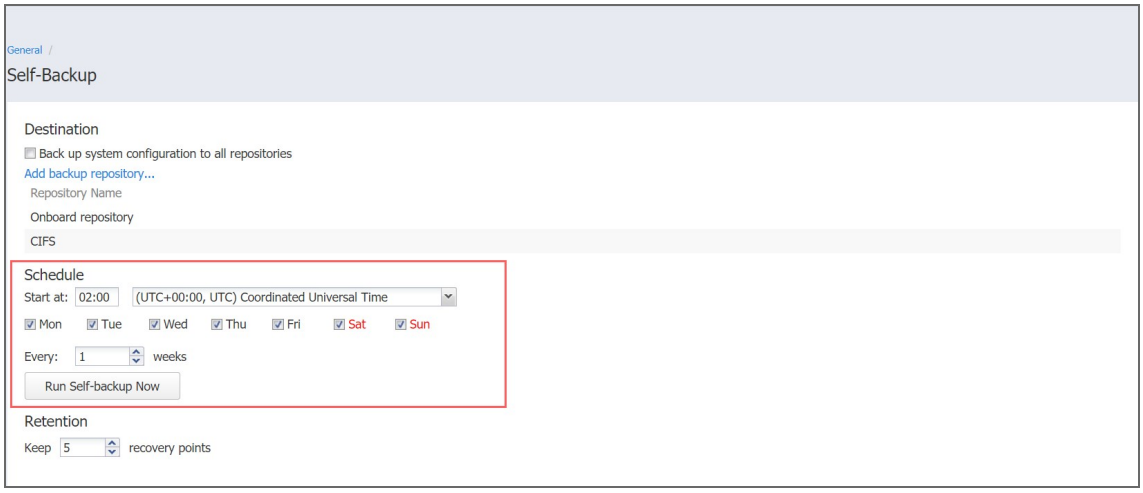

Self-Backup Retention

In the **Retention** 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.

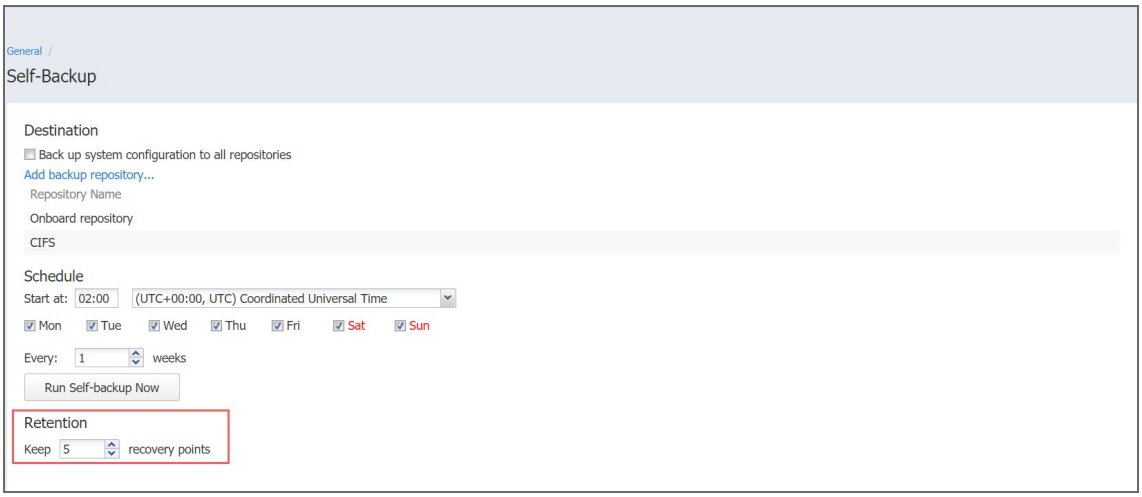

Recovering from Self-Backup

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.

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. Click **System Migration** in the **General** tab.
3. Click **Export system configuration**.
4. In the dialog window that appears, click **Export**.

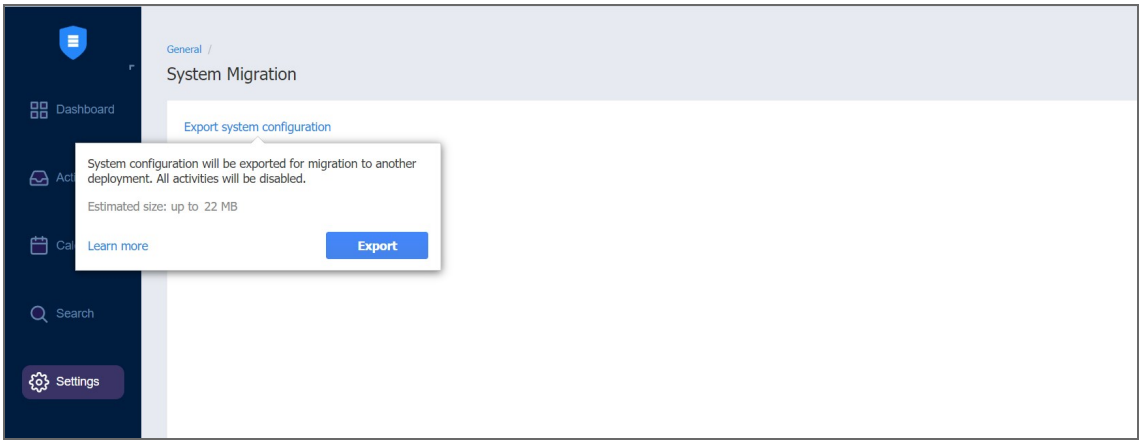

5. Click **Proceed** to confirm the operation.

   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. Click **System Migration** in the **General** tab.
3. Click **Import system configuration**.
4. In the dialog window that appears, locate the system configuration bundle using the **Browse** button.

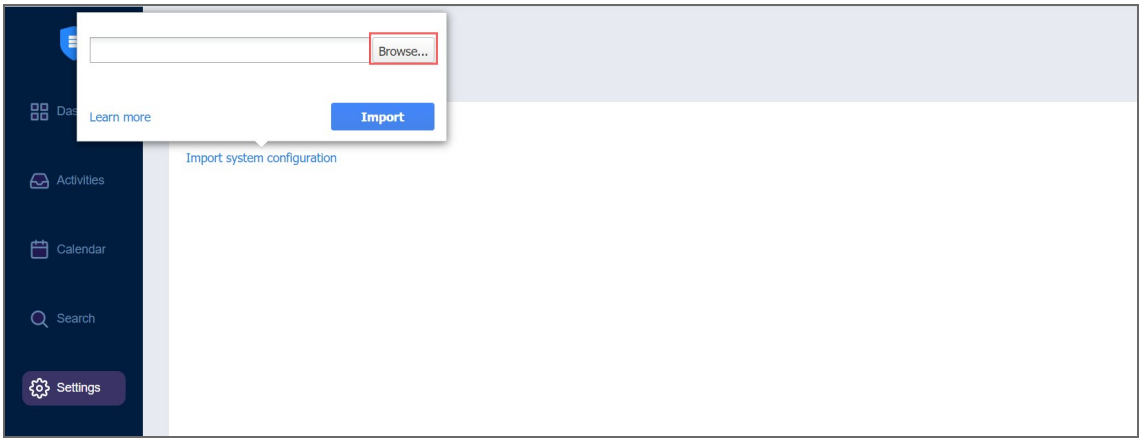

5. Click **Import**.
6. Click **Proceed** to confirm the operation.

   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.

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 *System* section:
        - **Store system events for the last x days**: Events older than the specified number of days (can be from 10 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.
        - **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.
        - **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.
        - **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.
        - **Display special offers**: When this option is enabled, the NAKIVO special offers toolbar appears in the NAKIVO Backup & Replication interface.
        - **Continue to update even if self-back fails**: When this option is selected, updates proceed even if self-backup cannot be performed.
        - 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.
    - In the *Tape Options* section:
        - **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.

- **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.
- **Auto refresh tapes every**: Select how often the contents of the tapes are refreshed in minutes or hours. Deselect if refreshing is not required.

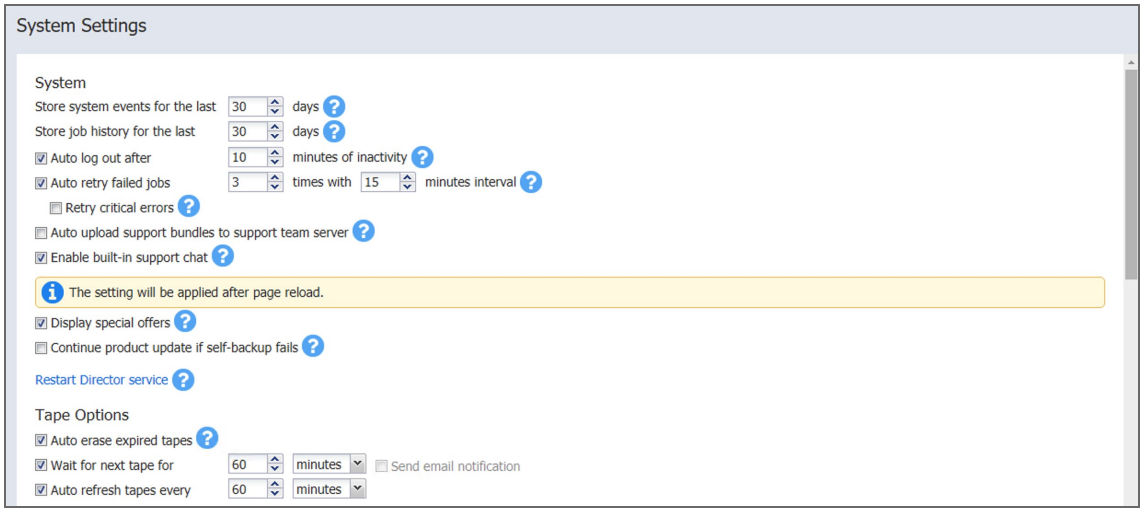

- In the *Processing Options* section:
  - **Auto remove deleted VMs and instances from jobs**: This option applies to a protected container (such as a VMware cluster or EC2 region). When this option is selected, if NAKIVO Backup & Replication discovers (during the inventory refresh) that a VM(s) and/or 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 machine 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.
  - **Skip swap files and partitions during processing**: When this option is selected, swap files and partitions are skipped during backup and replication to reduce backup size.
  - **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 *Integration Options* section:
  - **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" on page 763.
- In the *Auto Refresh* section:
  - **Auto refresh inventories 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.

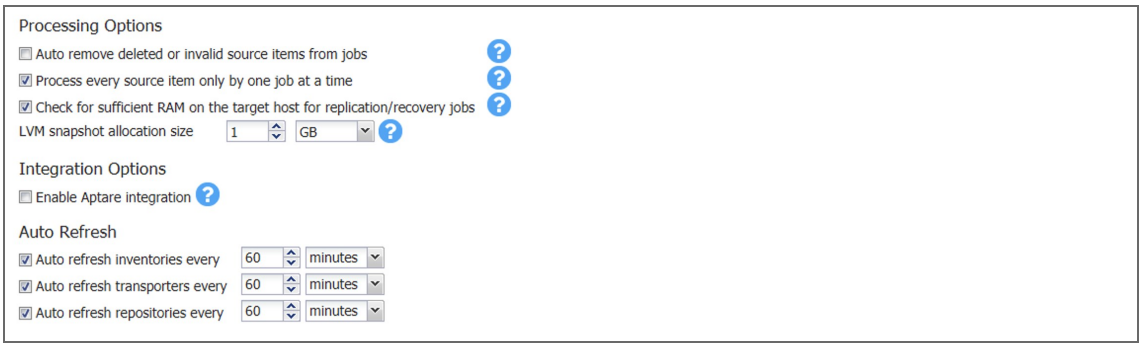

- In the *Regional options* section, set:
  - **Clock format**
  - **Short date format**
  - **Long date format**
  - **First day of week**
  - **Decimal symbol**
  - **Default time zone**

Note

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

- In the *Security section*you can enable the followin:
  - **Accept all transporter certificates by default**: When this option is selected, all Transporter certificates are accepted by default. Enabling this option is not recommended due to the security risks involved.

- **Enforce usage of pre-shared keys for all transporters**: When this option is selected, all Transporters start requiring the pre-shared keys to function and become inaccessible without them.

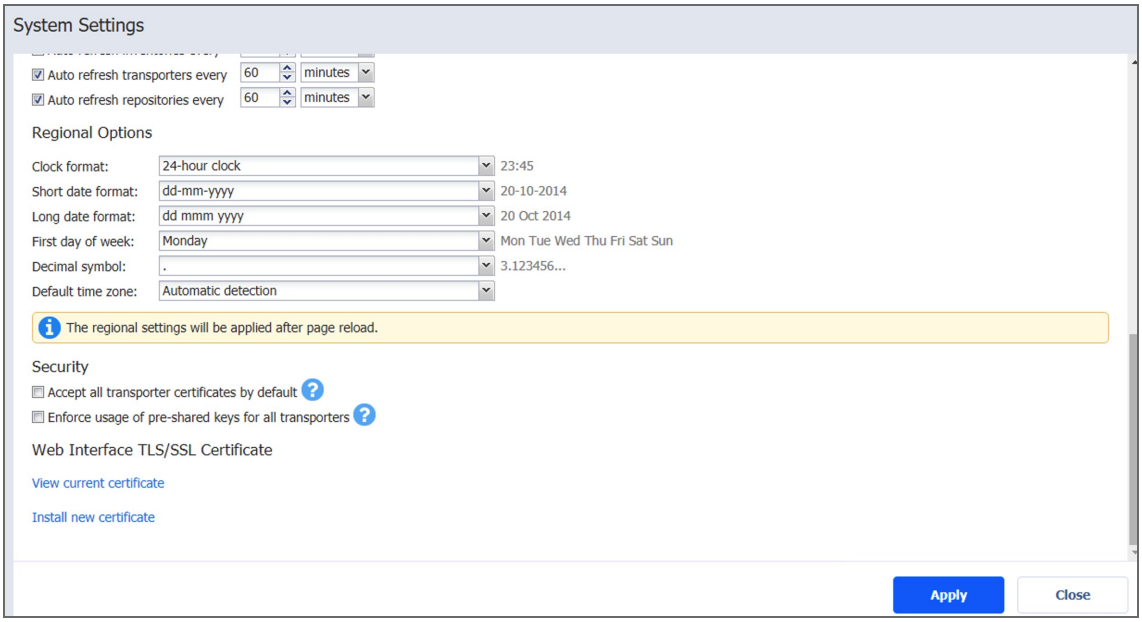

- In the *Web Interface TLS/SSL Certificate* section, you can either:
  - **View current certificate**: A dialog containing the current certificate information opens.
  - **Install new certificate**: A dialog opens allowing you to install a new TLS/SSL certificate for the NAKIVO Backup & Replication 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 one of the following formats: *.pem, *.crt, *.cer, *.p7b, *.p7s.

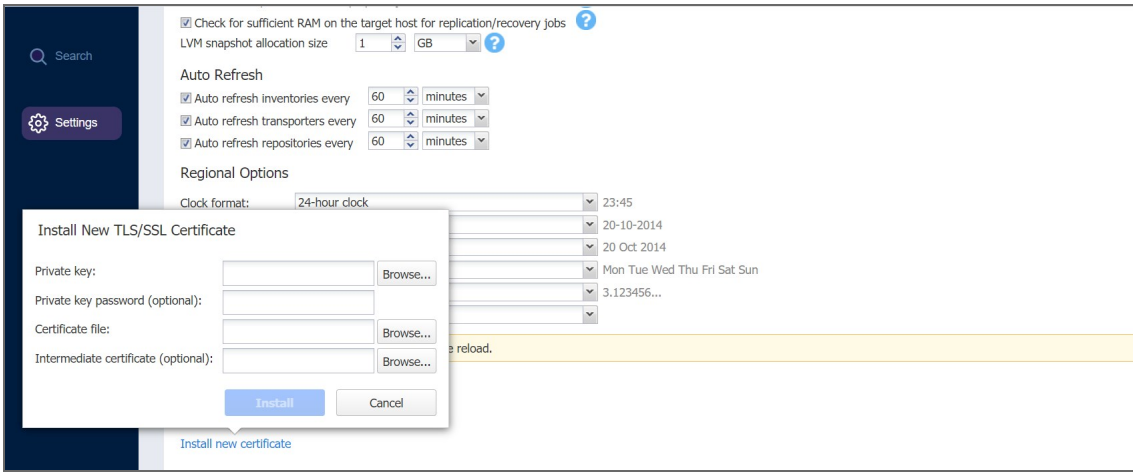

4. Click **Install**.

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.

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 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**.

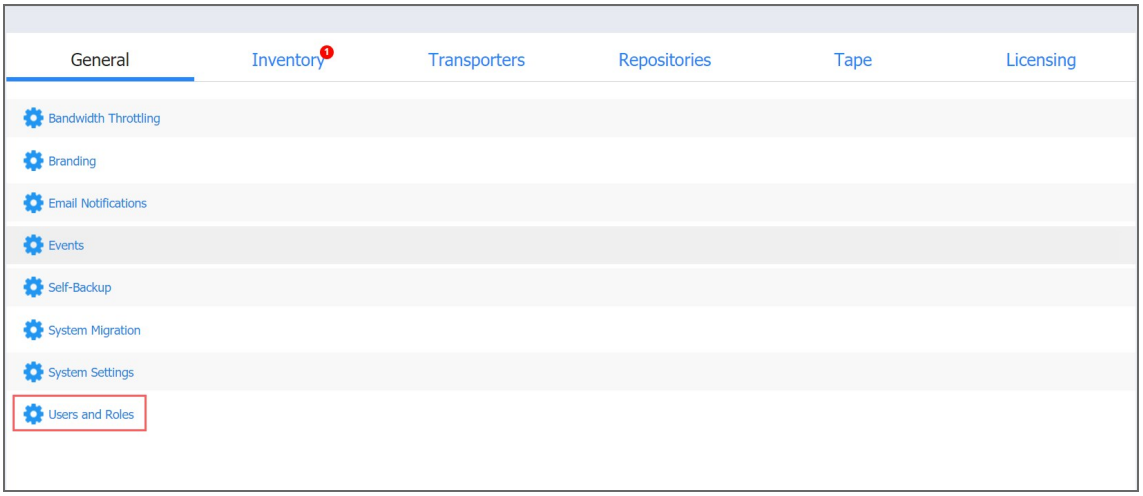

Navigating Users View

To see the list of all local users, select the **Users** view from the drop-down list in the upper-left 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**, or **Group** by clicking on the respective name of the column.
- 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 **Name**, **User name**, **Email**, **Group**, **Role**, and **Status**.
- Add a new local user by clicking **Add User**.
- 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. This can also be done in bulk by checking the box in the upper left pane to select all users and clicking **Bulk Action**.

  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 AD Groups View

To see the list of all Active Directory groups, select the **AD Groups** view from the drop-down list in the upper-left 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 **Name**, **Users**, **2FA**, or **Role** by clicking on the respective name of the column.
- 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 **Name**, **Role**, and **Status**.
- Add a new AD group by clicking **Add AD group**.
- Integrate Active Directory account by clicking **AD Integration**.
- Edit, delete, disable, enable Two-factor authentication, and assign a new role to the AD group individually. This can also be done in bulk by checking the box in the upper left pane to select all groups and clicking **Bulk Action**.

  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 local use 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 304
- “Managing Local Users” on page 311
- “Managing User Roles” on page 319
- Configuring Two-Factor Authentication

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" on page 305
- "Assigning Role to Active Directory User" on page 306
- "Configuring Active Directory Integration" on page 307
- "Deleting Active Directory User" on page 308
- "Disabling Active Directory User" on page 309
- "Editing Active Directory User" on page 310

Adding Active Directory User

After configuring AD integration in the **Active Directory Configuration** wizard, you can proceed with adding AD user(s). 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.
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. In the **Role** list, choose a user role to be assigned to the users.

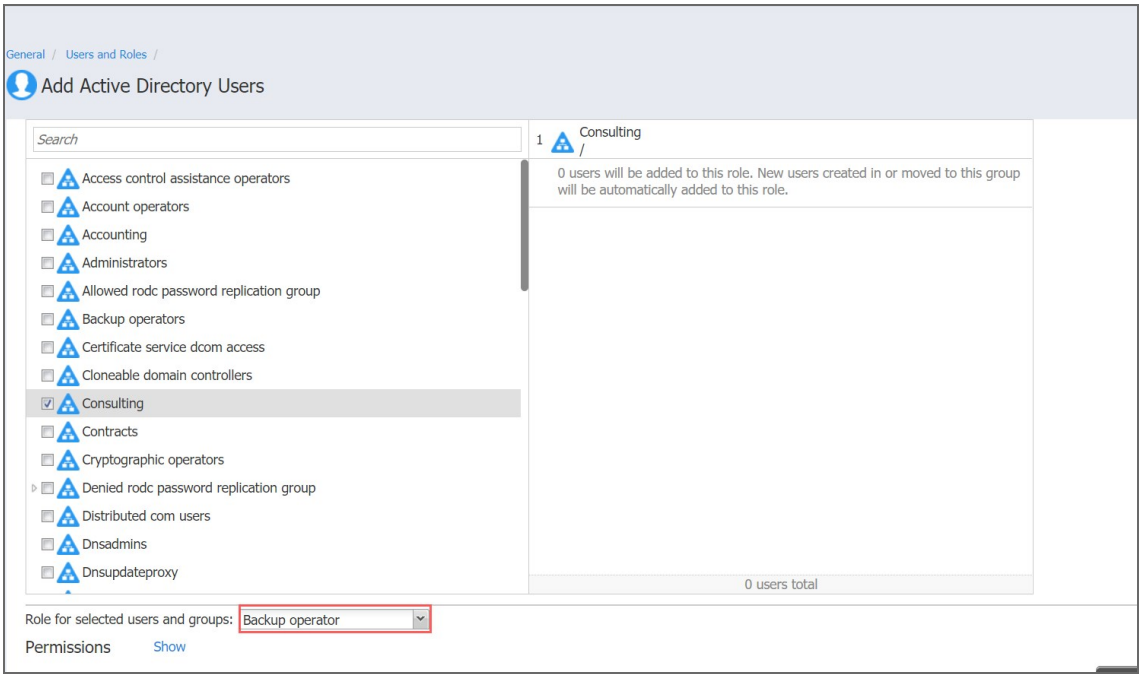

6. To unhide permissions granted to the users based on the role assigned, click the **Show** button beside the **Permissions** label. To hide the user permissions again, click **Hide**.
7. In the lower right corner of the page, click **Add**. Active Directory users appear in the NAKIVO Backup & Replication 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 & Roles**.
2. The **Users and Roles** page opens in the **Users** view. Hover over the Active Directory user, and then click **Manage** in the rightmost column of the row.
3. In the menu that opens, click **Assign role**.
4. In the dialog box that opens, select a new user role from the **Role** list and then click **Save**.

The Active Directory user appears in the list of users with the assigned role.

Configuring Active Directory Integration

To configure Active Directory integration, follow these steps:

1. Go to **Settings > General > Users & Roles**.
2. The **Users & Roles** page opens in the **Users** view. 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.

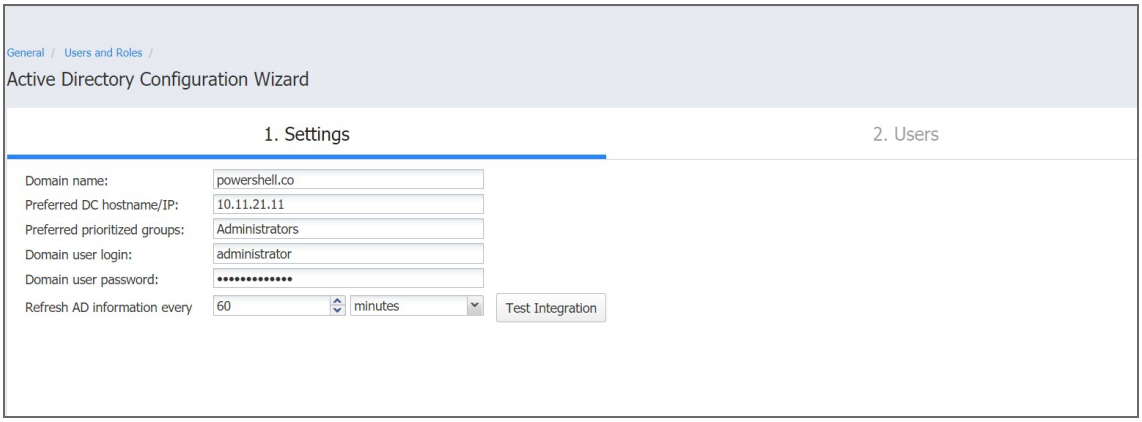

   h. If Active Directory integration is tested successfully, a checkmark appears beside the **Test Integration** button. Then click **Next** to go to the next page of the wizard. If you fail to connect to the AD domain, refer to the Knowledge Base article for possible causes.
   i. On the **Users** page 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 & 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 **Manage** 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 & 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 **Manage** 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 & 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 **Manage** 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 **Description** box, edit the user description.
    b. In the **Role** list, edit the user role.
    c. To unhide permissions granted to the user based on the role assigned, click the **Show** button beside the **Permissions** label. To hide the user permissions again, click **Hide**.
    d. 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" on page 312
- "Assigning Role to Local User" on page 314
- "Deleting Local User" on page 315
- "Disabling Local User" on page 316
- "Editing Local User" on page 317

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 in the **Users** view. Click **Add User**.
3. In the menu that opens, click **Local User**.

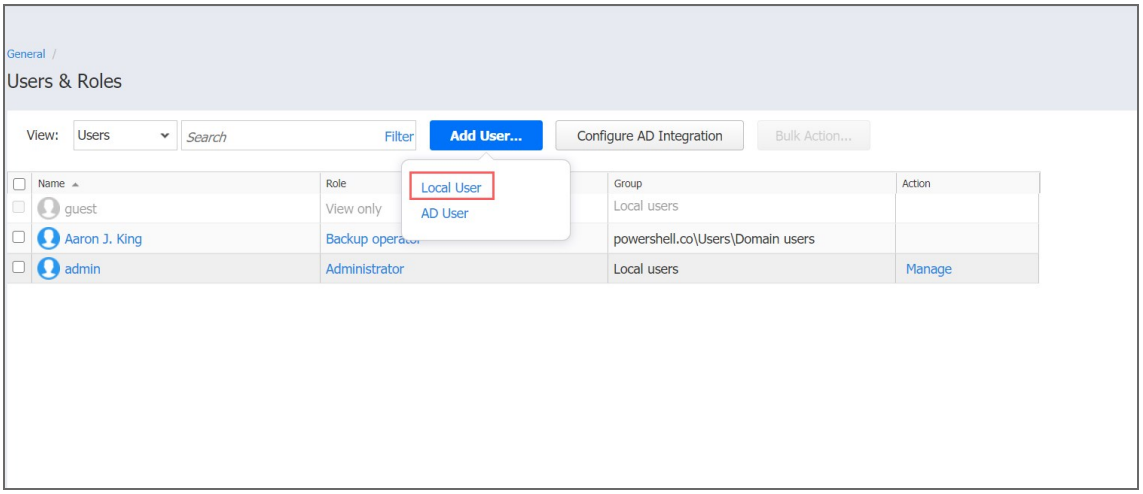

4. The **Add Local User** page opens. Proceed as follows:
   a. In the **Username** box, enter the user name.
   b. In the **Name** box, enter the user's real name.
   c. In the **Password** box, enter the user password. To generate a password automatically and send it to the user, select **Generate password and send by**.
   d. In the **Repeat password** box, re-enter the user password.
   e. In the **Email** box, enter the user's email address.
   f. In the **Description** box, optionally enter a user description.
   g. In the **Access level** dropdown list, select an access level for the new user (for multi-tenant solutions only).
   h. In the **Role** dropdown list, select a user role. Refer to "Managing User Roles" on page 319 for more details about user roles.
   i. To unhide permissions granted to the user based on the role assigned, click the **Show** button beside the **Permissions** label. To hide the user permissions again, click **Hide**.
   j. To proceed with creating another user after creating the current one, select **Create another user**.

k. In the lower right corner of the page, click **Add**.

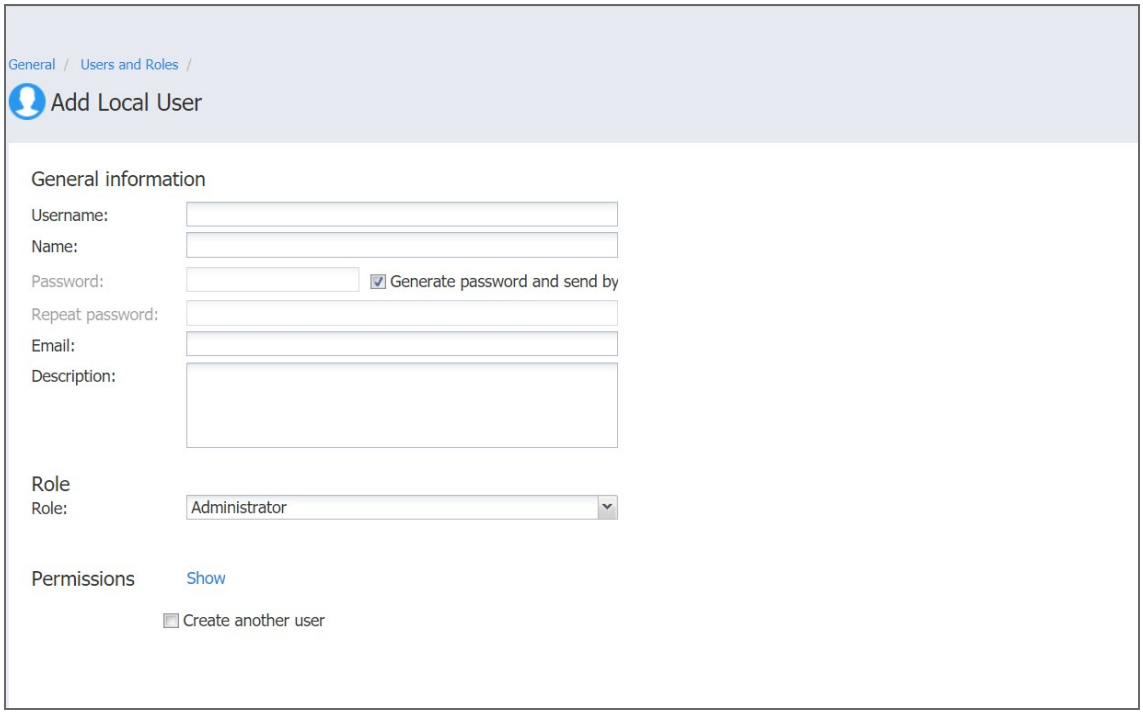

The local user appears in the list of users.

Assigning Role to Local User

Follow the steps below to assign a role to a local user:

1. Go to **Settings > General > Users & Roles**.
2. The **Users and Roles** page opens in the **Users** view. Hover over the local user, and then click **Manage** 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** drop-down list and then click **Save**.

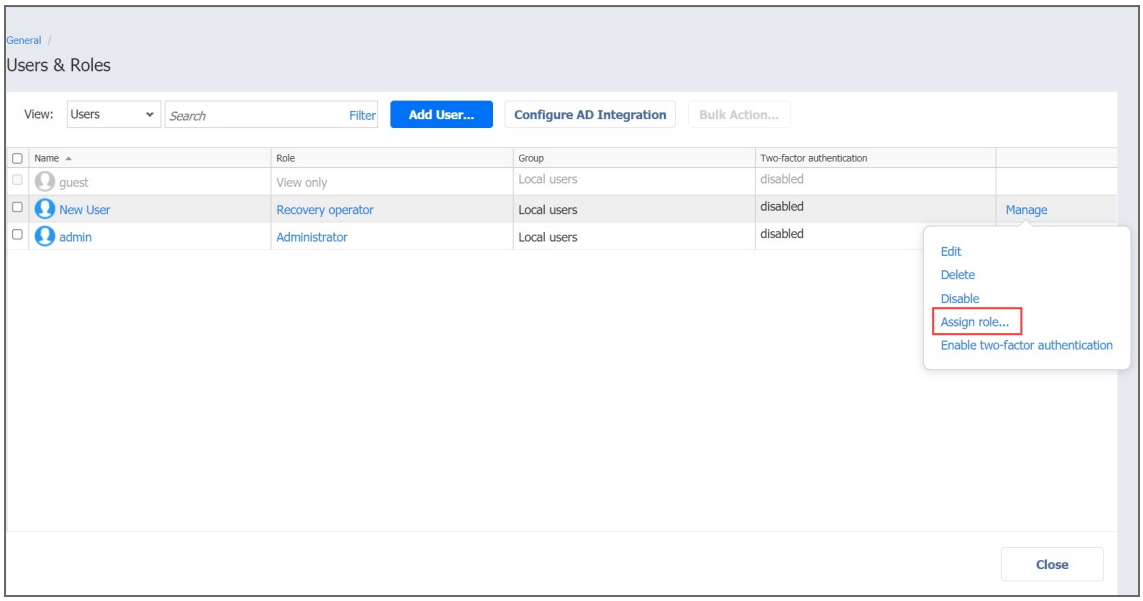

The local user appears in the list of users with the assigned role.

Deleting Local User

Follow the steps below to delete a local user:

1. Go to **Settings > General > Users & Roles**.
2. The **Users and Roles** page opens in the **Users** view. Hover over the local user you wish to be deleted, and then click **Manage** 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 delete the local user.

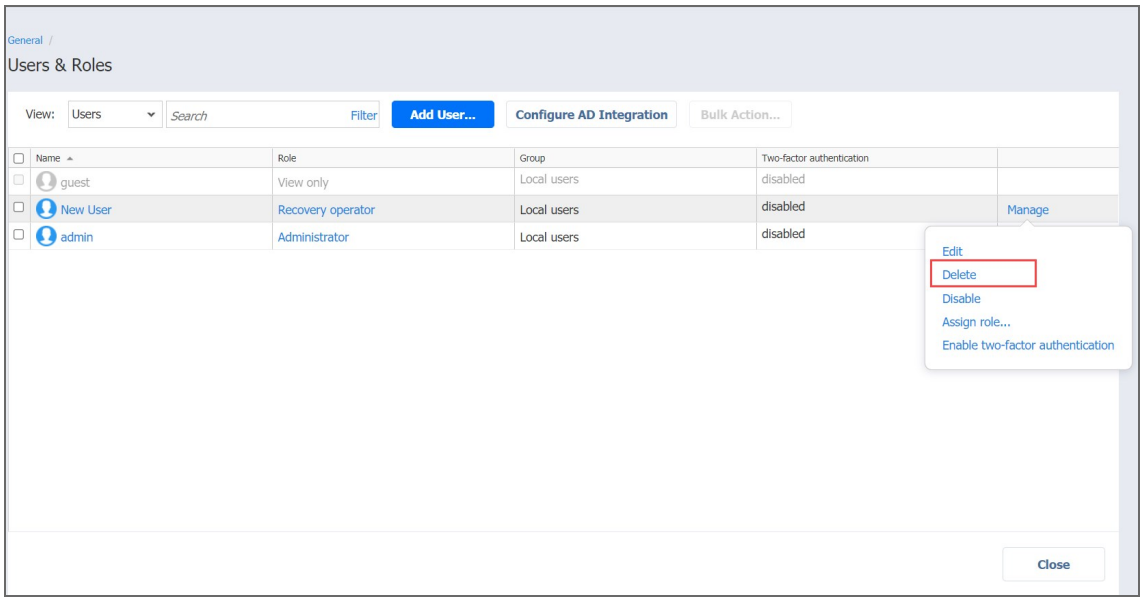

The deleted user disappears from the list of users.

Disabling Local User

Follow the steps below to disable a local user:

1. Go to **Settings > General > Users & Roles**.
2. The **Users and Roles** page opens in the **Users** view. Hover over the local user you wish to be disabled, and then click **Manage** 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 disable the local user.

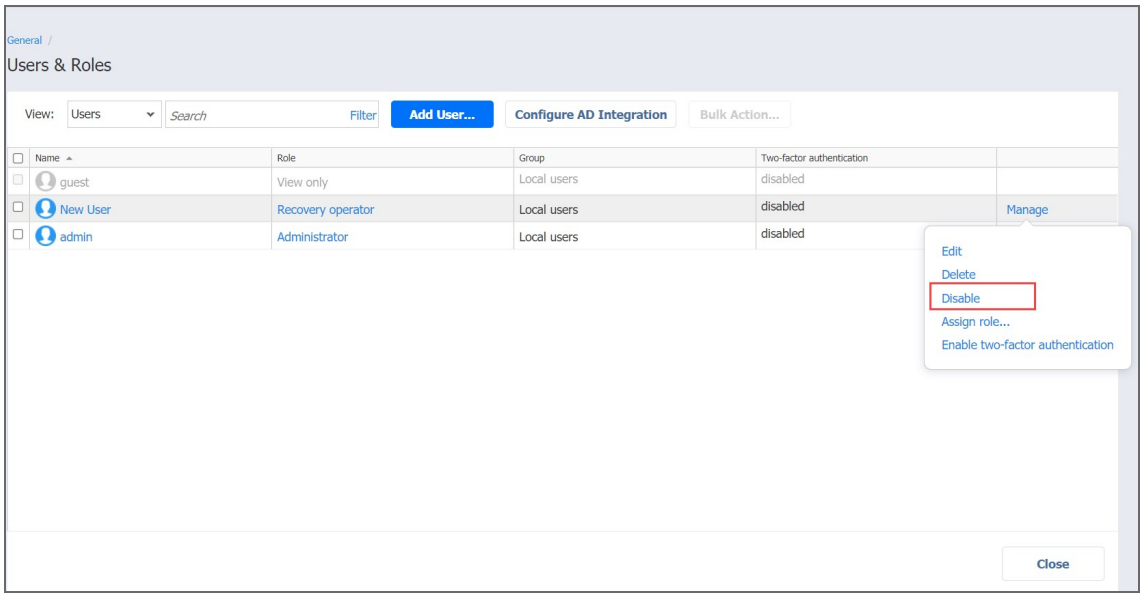

The disabled user appears dimmed in the list of local users.

Editing Local User

Please follow the steps below to edit a local user:

1. Go to **Settings > General > Users & 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.
   b. Hover over the local user, click **Manage** in the rightmost column of the row and then click **Edit**.

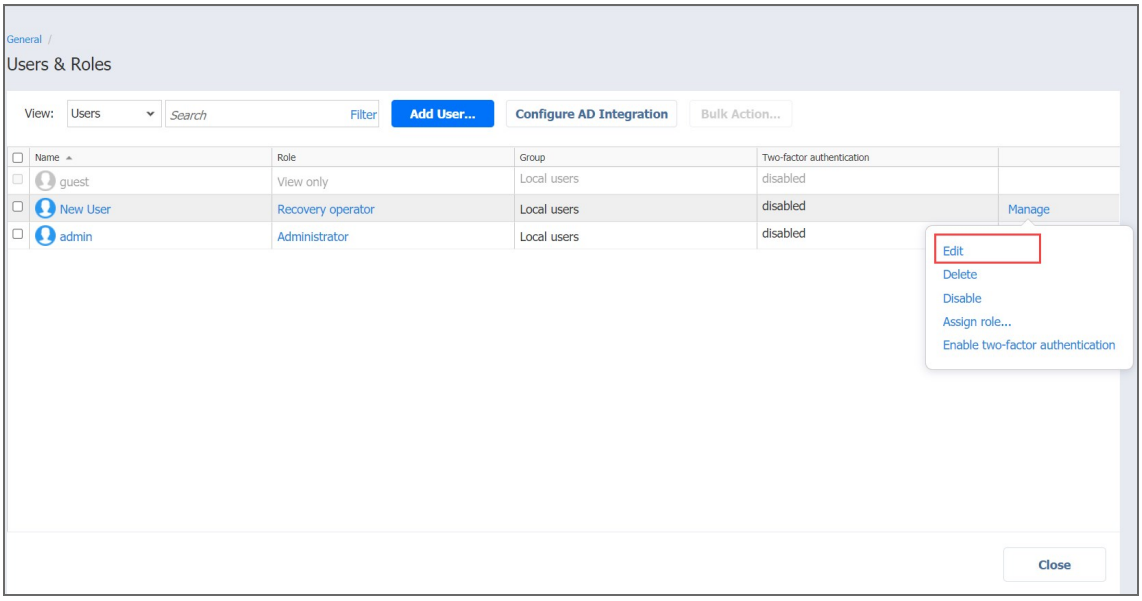

3. The **Edit User** page opens. Edit the local user properties if needed:
   a. In the **Name** box, edit the user name.
   b. In the **Password** box, edit the user password.
   c. If you edited the user password, re-enter the user password in the **Repeat password** box.
   d. In the **Email** box, 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** list, edit the user role.
   h. To unhide permissions granted to the user based on the role assigned, click the **Show** button beside the **Permissions** label. To hide the user permissions again, click **Hide**.

i. Click **Save** to save your modifications to the local user.

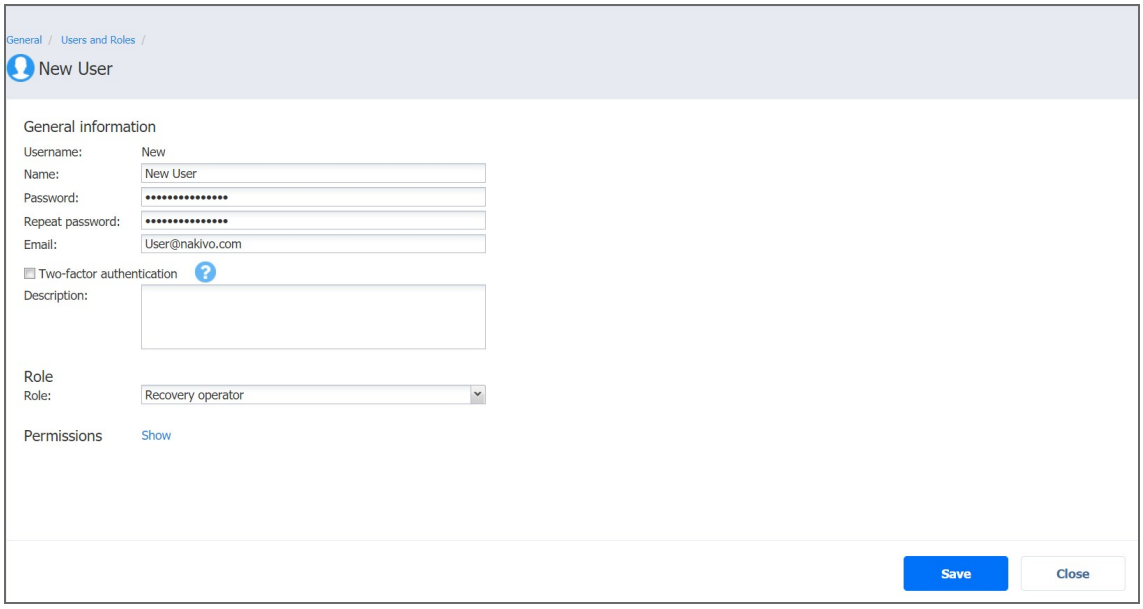

Managing User Roles

A user role with full access to the **User management** permission is assigned to your user profile to manage user roles. You cannot edit or delete the user role that is assigned to your user profile. The following topics describe how to manage roles of NAKIVO Backup &Replication users in detail:

- "Overview of User Roles" on page 325
- "Adding User Role" on page 320
- "Editing User Role" on page 323
- "Cloning User Role" on page 321
- "Deleting User Role" on page 322

Adding User Role

Follow the steps below to add a user role:

1. Go to **Settings > General > Users & Roles**.
2. On the **Users & Roles** page, switch to the **Roles** view.
3. Click **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.
    d. To unhide permissions to be granted to the role, click the **Show** button beside the **Permissions** label.
    e. A list of permissions opens. Specify necessary permissions for the user role.
    f. Click **Add** in the lower right corner of the page.

The user role appears in the list of roles.

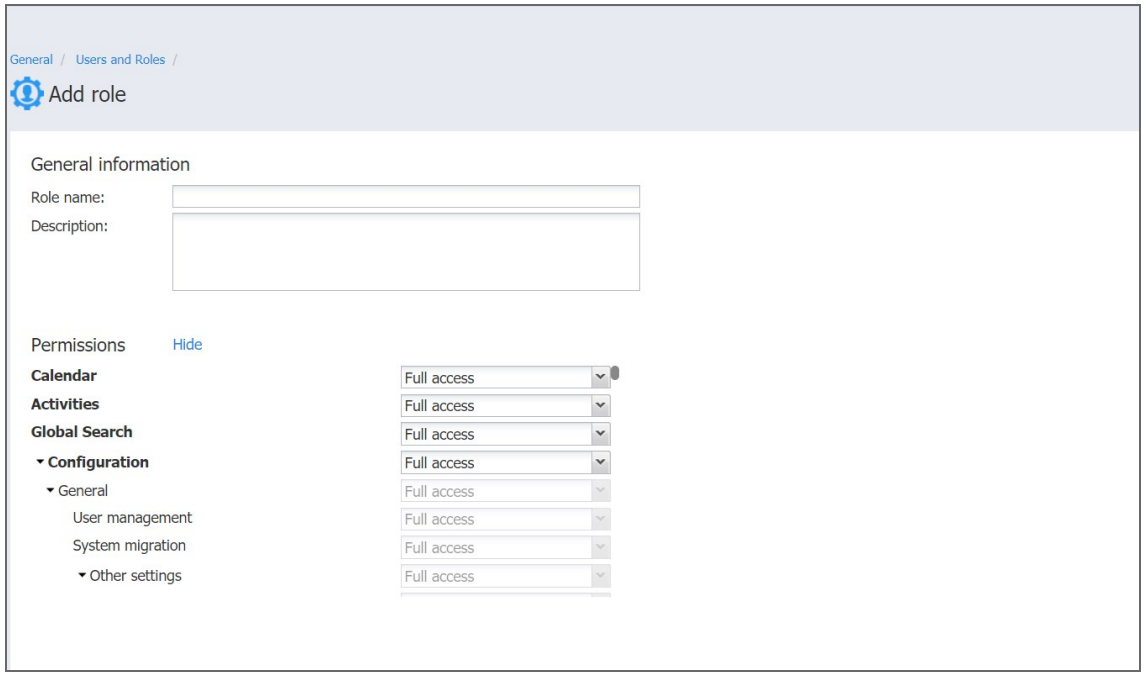

Cloning User Role

Follow the steps below to clone a user role:

1. Go to **Settings > General > Users & Roles**.
2. On the **Users & Roles** page, switch to the **Roles** view.
3. Hover over the user role, click **Manage** in the rightmost column of the row and then click **Clone**.
4. A dialog opens asking you to enter the name of the new user role. Enter the name of the new user role and click **Save**.

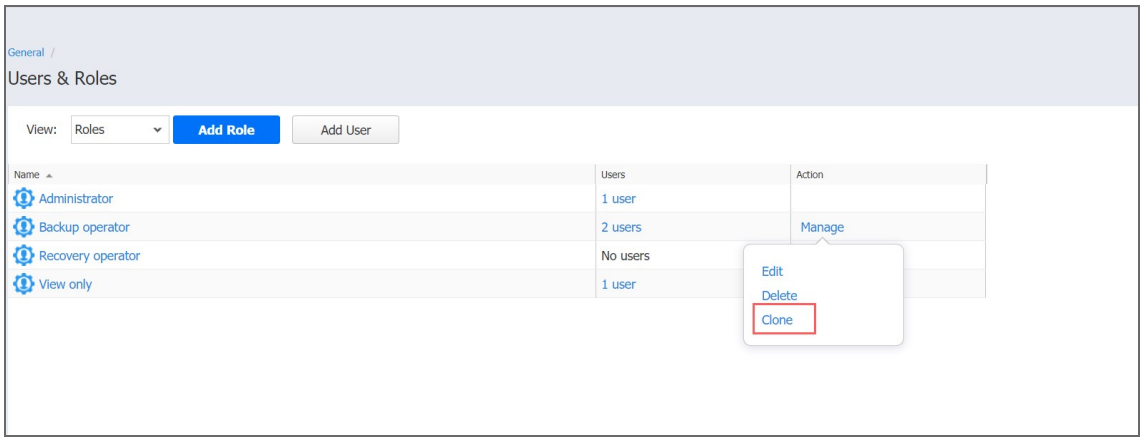

The new user role appears in the list of roles.

Deleting User Role

Follow the steps below to delete a user role:

1. Go to **Settings > General > Users & Roles**.
2. On the **Users & Roles** page, switch to the **Roles** view.
3. Hover over the user role, click **Manage** in the rightmost column of the row and then click **Delete.**
4. In the dialog box that opens, click **Delete** to confirm deleting the local user.

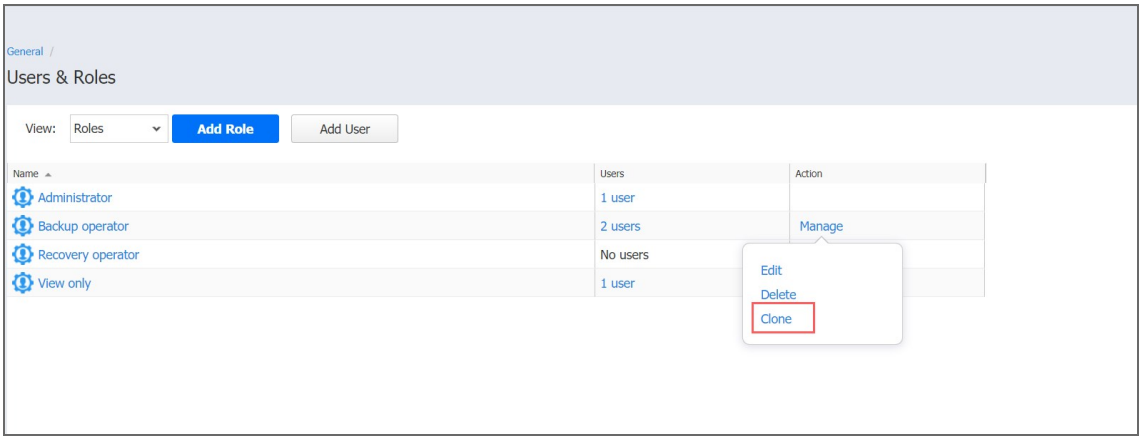

The user role disappears from the list of roles.

Editing User Role

Follow the steps below to edit a user role:

1. Go to **Settings > General > Users & Roles**.
2. On the **Users & Roles** page, switch to the **Roles** view.
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 **Manage** in the rightmost column of the row.
   c. Click **Edit**.

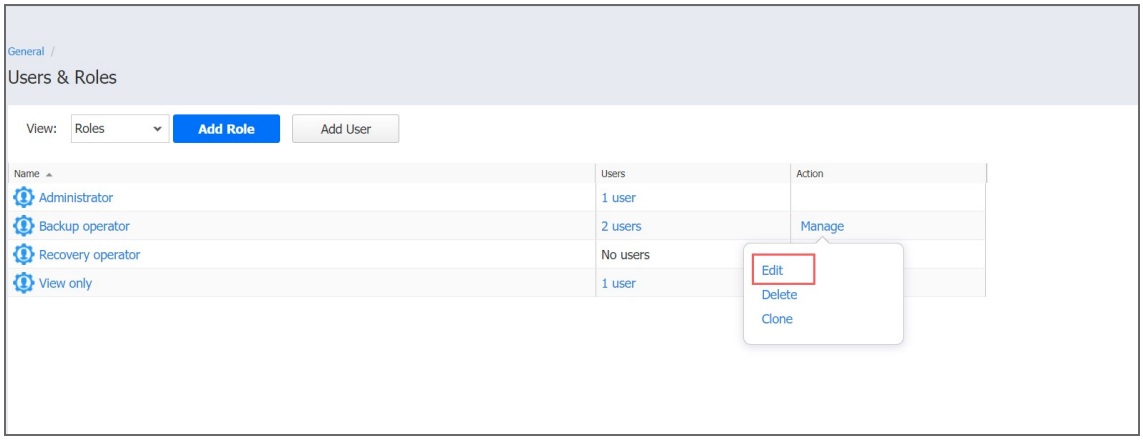

4. The **Edit User Role** page opens. Edit the user role properties if needed:
   a. In the **Role name** box, edit the user role name.
   b. If you are working with a multi-tenant environment, you can change the access level for this role by choosing another tenant, master tenant, or all tenants in the **Access level** list.
   c. In the **Description** box, edit the user description.
   d. You can view the **Number of users** assigned with this role and click **view all**to see their full list on a new page.
   e. To unhide permissions to be granted to the role, click the **Show** button beside the **Permissions** label.
   f. A list of permissions opens. Edit necessary permissions for the user role.

g. When ready to save the user role, click **Save** in the lower right corner of the page.

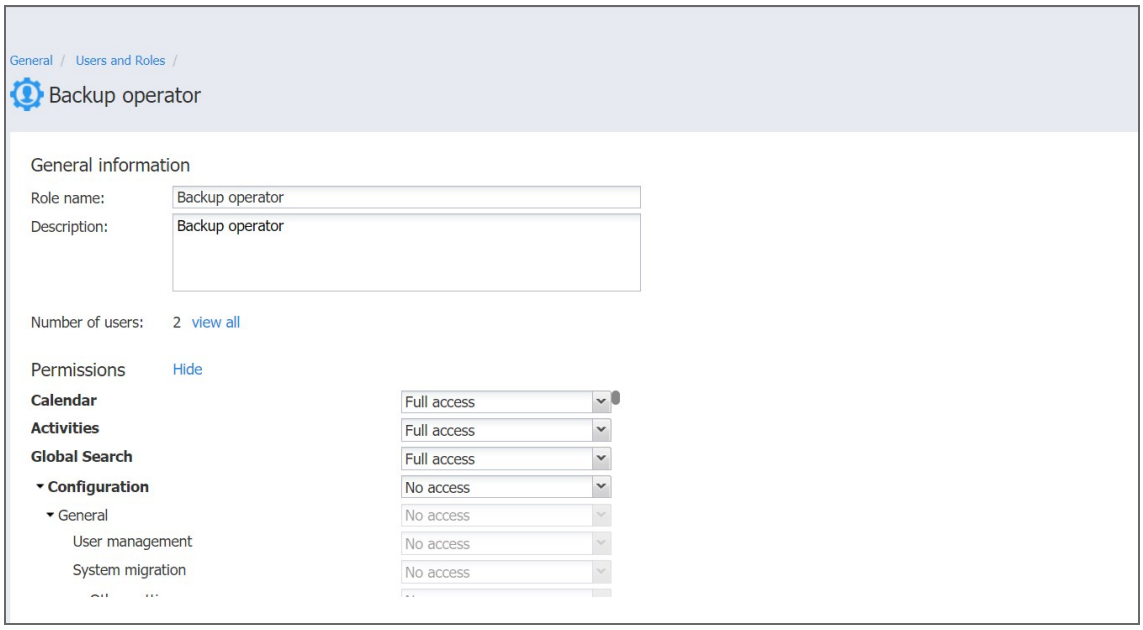

Overview of User Roles

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.

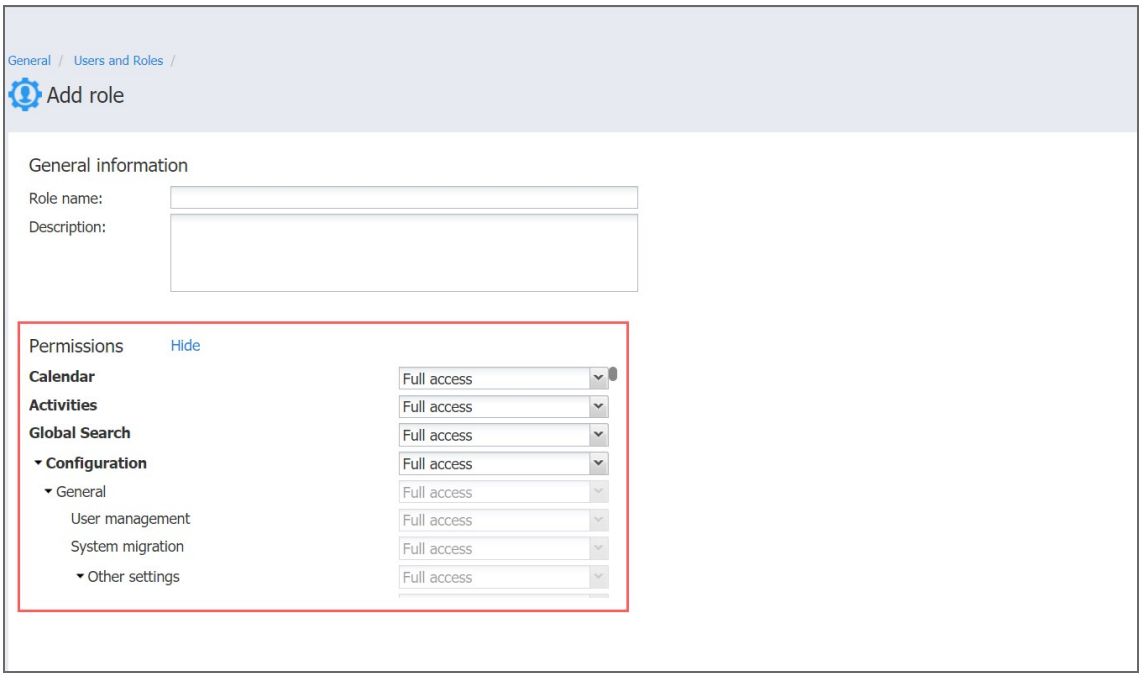

Access Levels

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.

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 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**.

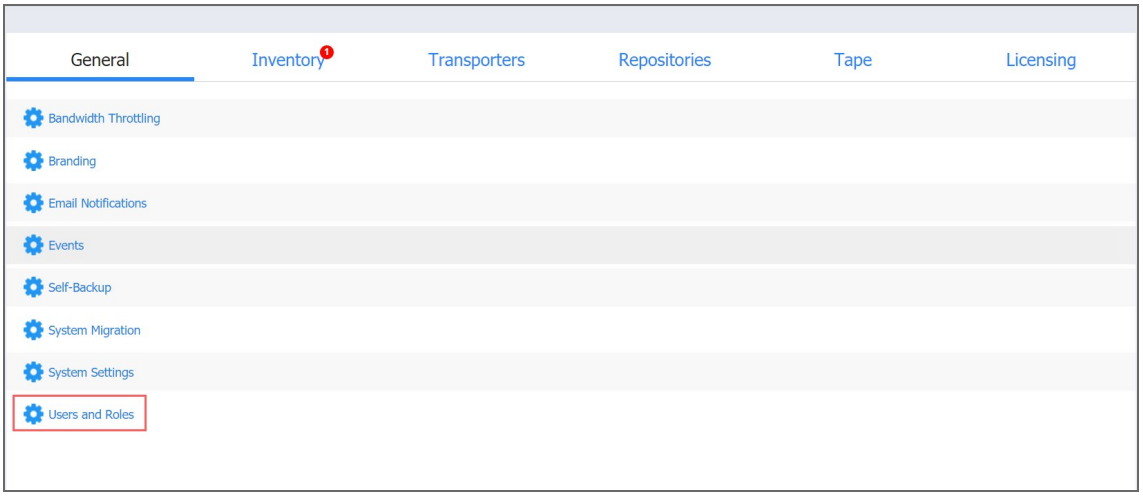

Navigating Users View

To see the list of all local users, select the **Users** view from the drop-down list in the upper-left 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**, or **Group** by clicking on the respective name of the column.
- 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 **Name**, **User name**, **Email**, **Group**, **Role**, and **Status**.
- Add a new local user by clicking **Add User**.

- 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. This can also be done in bulk by checking the box in the upper left pane to select all users and clicking **Bulk Action**.

  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 AD Groups View

To see the list of all Active Directory groups, select the **AD Groups** view from the drop-down list in the upper-left 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 **Name**, **Users**, **2FA**, or **Role** by clicking on the respective name of the column.
- 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 **Name**, **Role**, and **Status**.
- Add a new AD group by clicking **Add AD group**.
- Integrate Active Directory account by clicking **AD Integration**.
- Edit, delete, disable, enable Two-factor authentication, and assign a new role to the AD group individually. This can also be done in bulk by checking the box in the upper left pane to select all groups and clicking **Bulk Action**.

  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 local use 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 304
- “Managing Local Users” on page 311
- “Managing User Roles” on page 319
- Configuring Two-Factor Authentication

Inventory

Prior to creating backup, replication, or recovery jobs, you need to add your virtual/cloud infrastructure to the product's Inventory. The discovered infrastructure is added to the internal product database, which is refreshed every 1 hour by default. Refer to the following sections to learn more:

- "Adding Hyper-V Servers" on page 332
- "Managing Credentials" on page 341
- "Managing Inventory" on page 336

Adding Amazon EC2 Accounts

Add an Amazon EC2 account to NAKIVO Backup & Replication as described in the sections below.

- “Creating AWS Access Key ID and Secret Access Key” below
- “Adding an Amazon EC2 Account to Inventory” below

Creating AWS Access Key ID and Secret Access Key

Prior to adding your AWS account to inventory, you need to create and retrieve an AWS Access Key ID and Secret Access Key. They are used by NAKIVO Backup & Replication to sign programmatic requests 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, do the following:

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 user name (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 the only copy. Keep it confidential in order to protect your account.

Adding an Amazon EC2 Account to Inventory

To add an Amazon EC2 account to the product, follow the steps below:

1. Click **Settings**.
2. Go to the **Inventory** tab and click **Add New AWS/Wasabi Account**.

3. Fill out the fields in the dialog that appears and click **Add**.

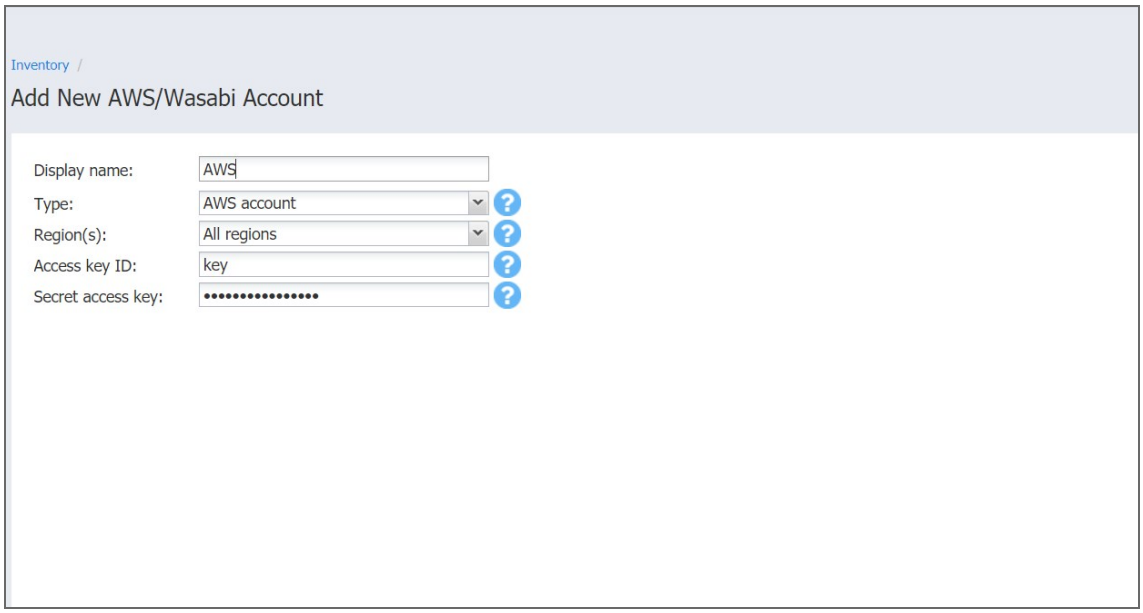

Adding Hyper-V Servers

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

1. In the main menu, click **Settings**.

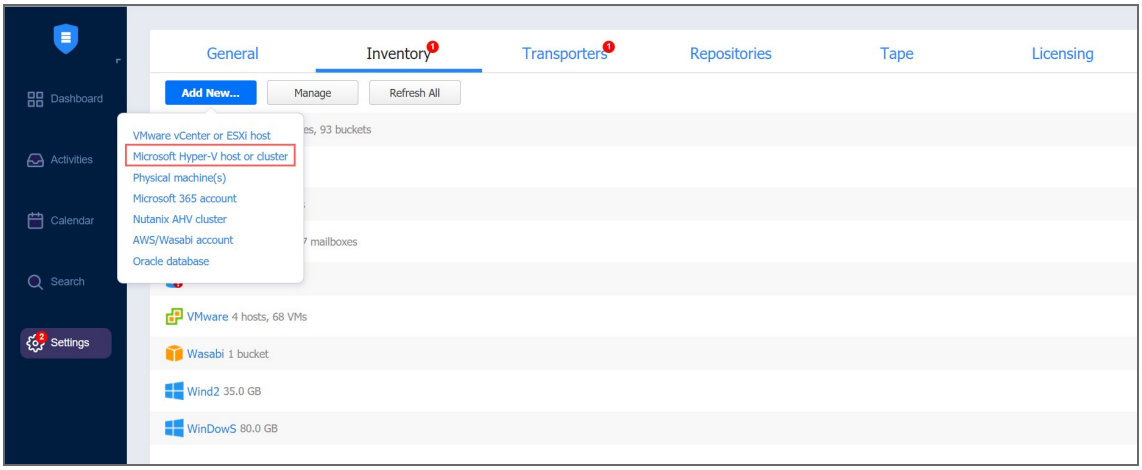

2. Go to the **Inventory** tab and click **Add New**.
3. Click **Microsoft Hyper-V host or cluster** in the dialog that opens.
4. Fill out the fields as follows:
   - **Display name:** Enter a name for the host that will be displayed in the inventory
   - **Type:** Select whether you want to add a standalone 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 Hyper-V server.
   - **Username and Password:** Specify the credentials of the 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.

5. Click **Add**.

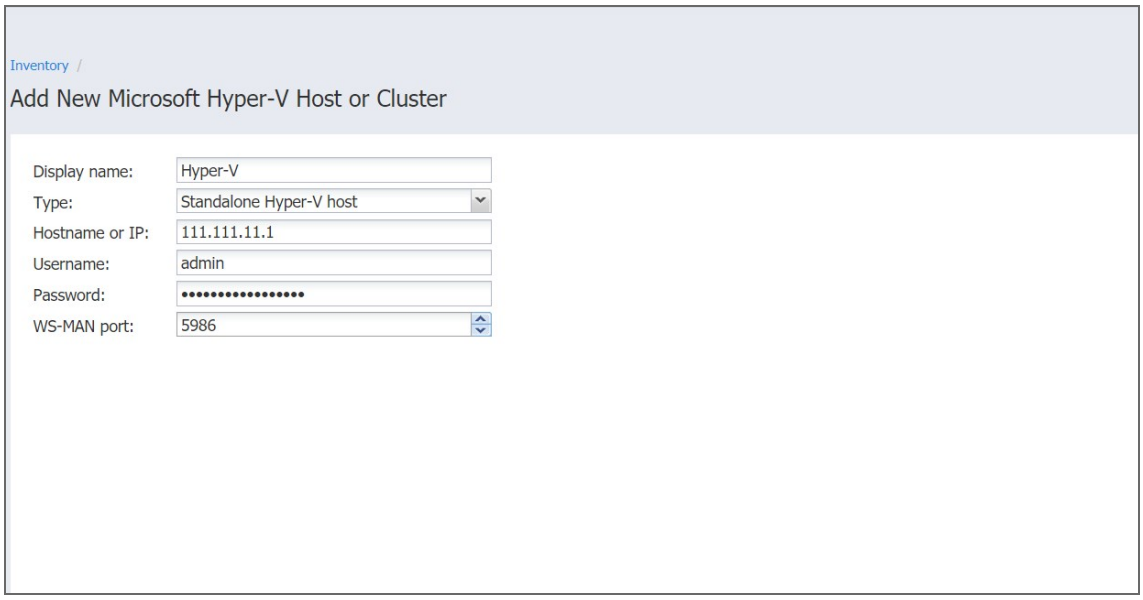

After the process has completed successfully, you can exit **Settings** and create jobs with the newly discovered VMs.

Notes

- When adding a new 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 Hyper-V host to the inventory of an additional NAKIVO Backup & Replication instance, Transporter installation fails.

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 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 programmatic requests 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, do the following:

1. If you don't have a Wasabi account, create a new one at wasabi.com/sign-up/
2. Log in to your Wasabi account.
3. Navigate to the main menu and click **Access Keys**.
4. Click **Create New Access Key**.

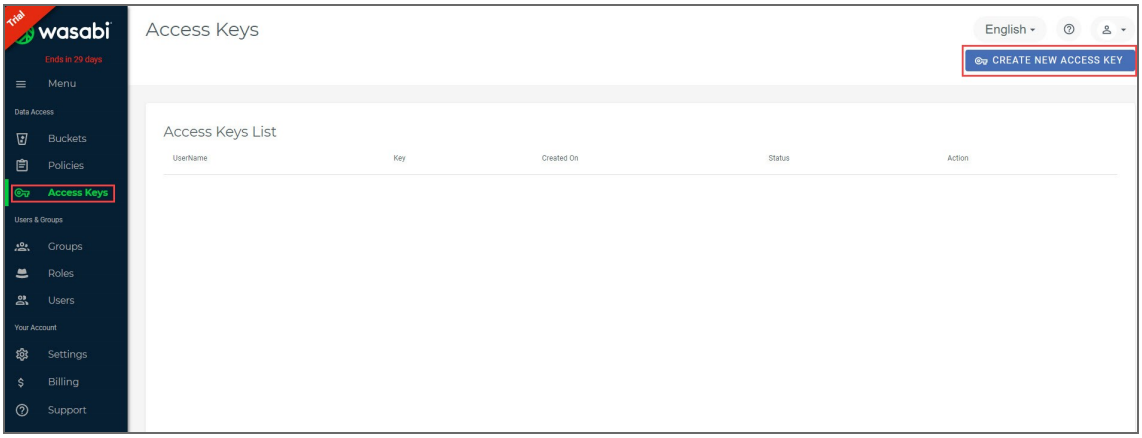

5. In the dialog box that opens, select one of the following:
   - **Root user** and click **Create**.

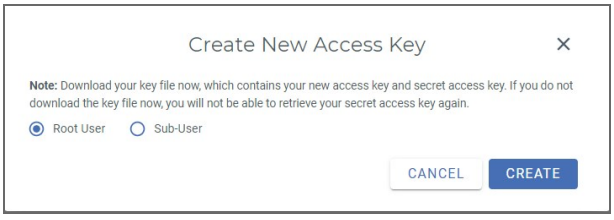

   - **Sub-User.** Select a sub-user from the **Assign to a user** drop-down list and click **Create**. Note that the original user access key of the selected user will be changed.

Note

To use the Sub-User option, you need to have at least one user to be created within your Wasabi account. For details, refer to Creating a User and How do I set up Wasabi for user access separation?

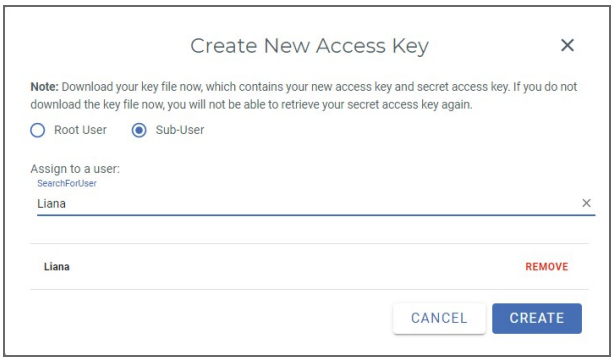

6. Click **Download CSV** and save the file with generated keys in a secure location. Keep the Access key confidential in order to protect your account.

Adding a Wasabi Account to Inventory

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

1. Click **Settings** in the main menu of NAKIVO Backup & Replication.
2. Go to the **Inventory** tab and click **Add New**.
3. In the dialog that opens, click **AWS/Wasabi Account**.
4. In the dialog box that opens:
   - Enter the name in the **Display name box**.
   - Select a Wasabi account from the **Type** drop-down list.
   - Select a Wasabi region from the **Region(s)** drop-down list.
   - Enter the Access Key ID and Secret Access Key of a root user or a sub-user in the corresponding fields.
5. Click **Add**.

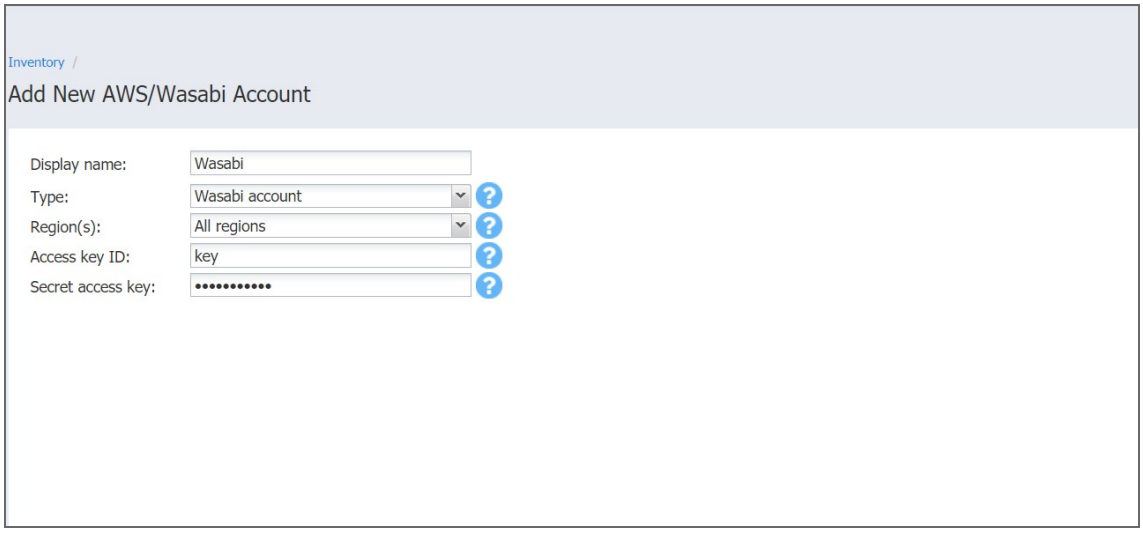

Managing Inventory

Refer to the following topics:

- “Refreshing Inventory” on page 338
- “Editing Inventory Items” on page 337
- “Removing Items from Inventory” on page 340

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. Click the item you want to edit.
4. In the title of the item, click **Manage** and then click **Edit**.

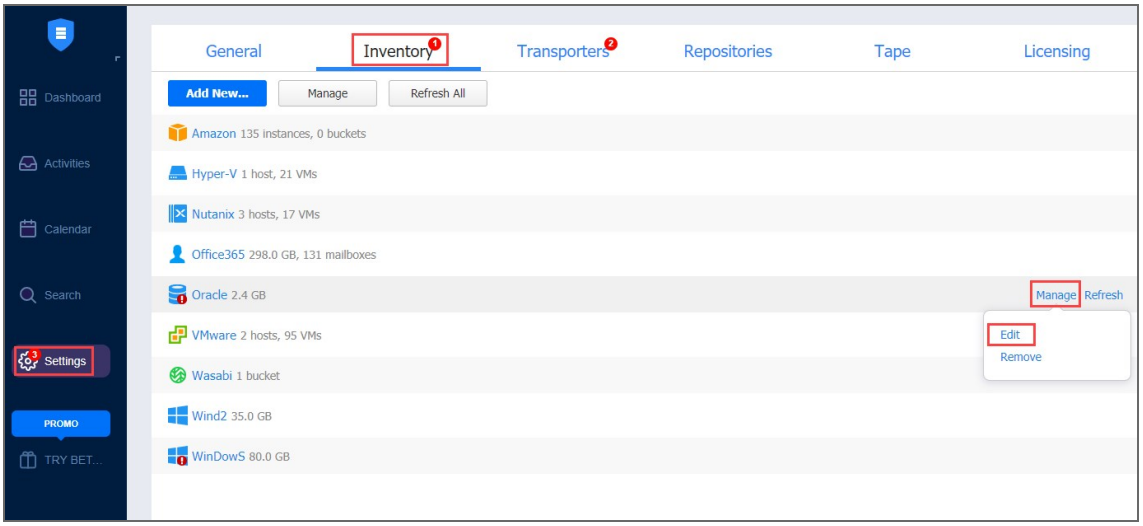

5. Update the appropriate fields and click **Apply**.

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 **Inventory** 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 setting are applied instantly and do not need to be saved.

Manually Refreshing All Inventory

To update 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 a Discovered Item

To update a single discovered 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 update.
3. In the title of the item, click **Refresh**.

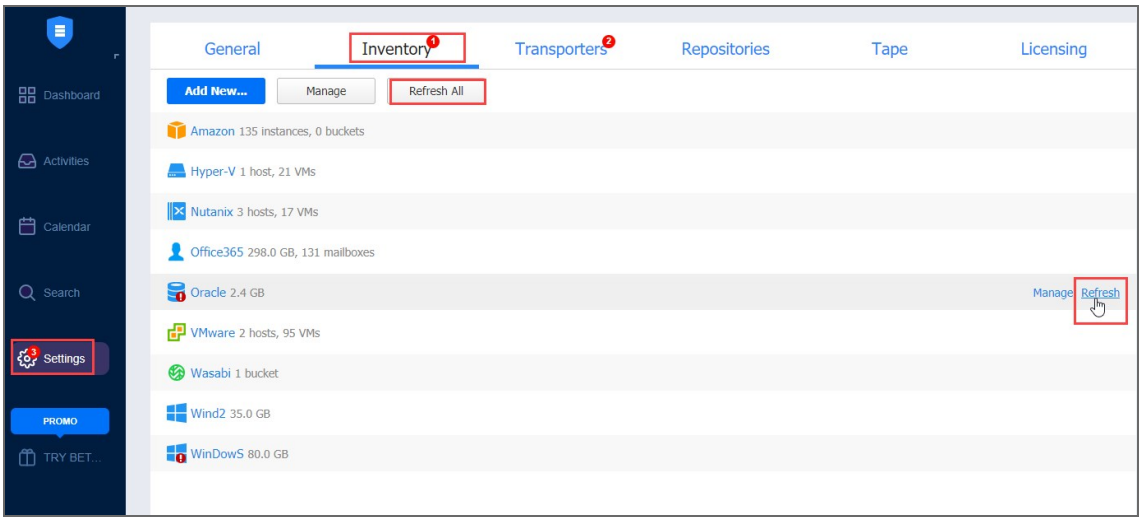
Dashboard
Activities
Calendar
Search
Settings
PROMO
TRY BET...
General
Inventory
Transporters
Repositories
Tape
Licensing
Add New...
Manage
Refresh All
Amazon 135 instances, 0 buckets
Hyper-V 1 host, 21 VMs
Nutanix 3 hosts, 17 VMs
Office365 298.0 GB, 131 mailboxes
Oracle 2.4 GB
Manage
Refresh
VMware 2 hosts, 95 VMs
Wasabi 1 bucket
Wind2 35.0 GB
WinDowS 80.0 GB

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. Click the item that you wish to remove from inventory.
3. In the item title, click **Manage**, and then click **Remove**.
4. In the dialog that opens, click **Remove**.

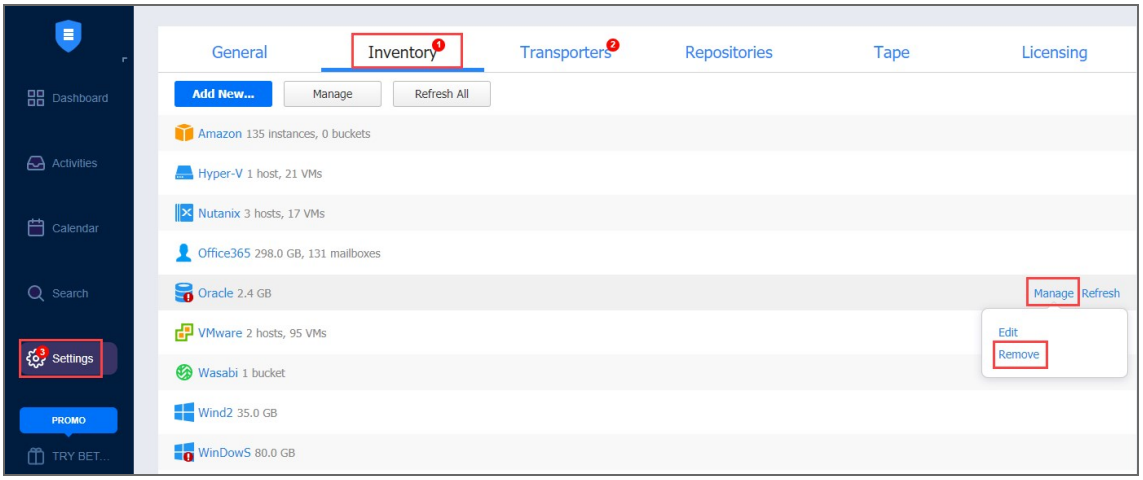

Managing Credentials

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

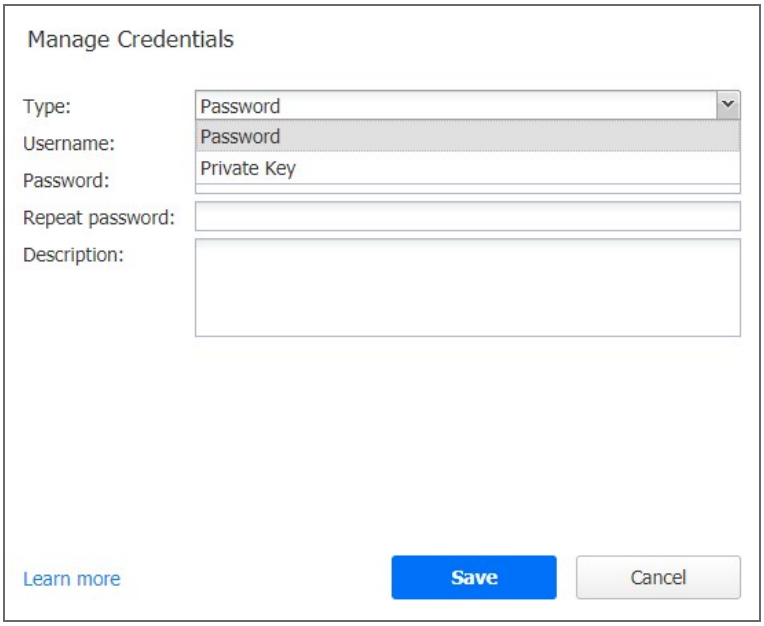

6. Then do the following:
   - **Type**: Select the type of credentials:
     - To add a username and password, fill out the **Username**, **Password**, and **Description** fields and click **Save**.
     - To add a private key to an Amazon EC2 instance or a Linux physical machine, do the following:
       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.**

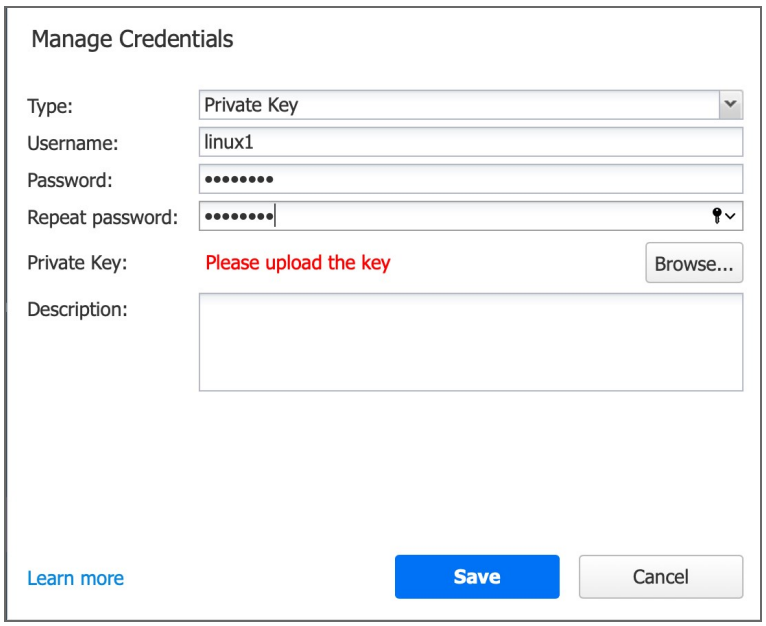

You can now assign the credentials while creating jobs.

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.**

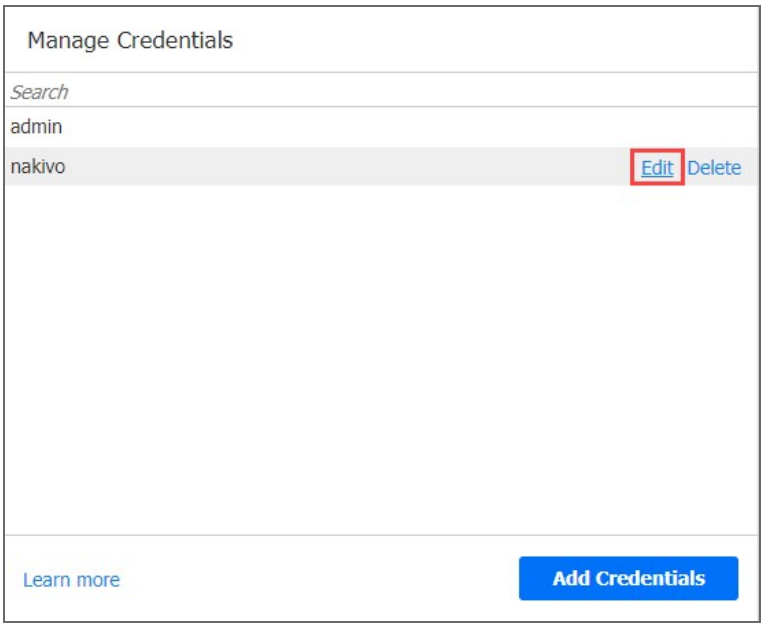

5. Make any required changes, and then click **Save**.

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.**

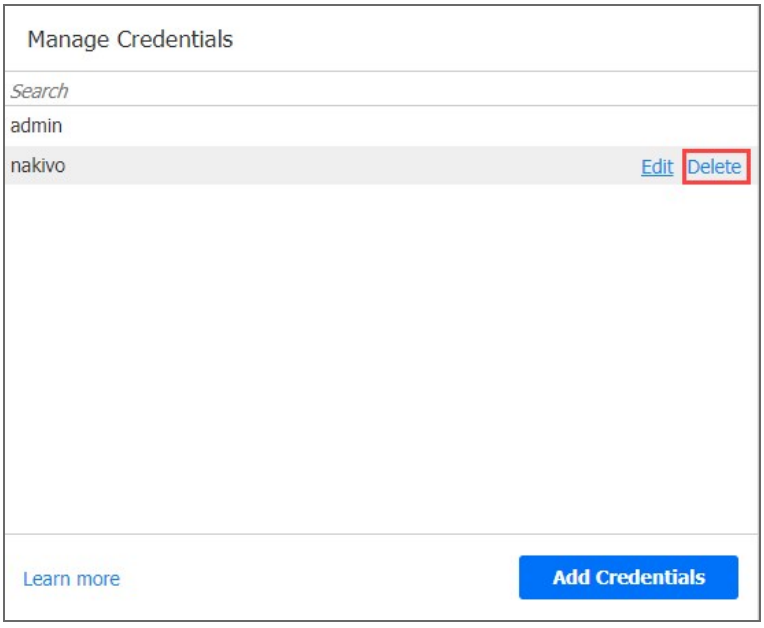

5. Click **Delete** in the confirmation dialog that opens.

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" on page 353
- "Managing Transporters" on page 362

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 **Settings** > **Transporters** and then click **Deploy New Transporter**.
2. In the dialog that opens, click **VMware appliance**.

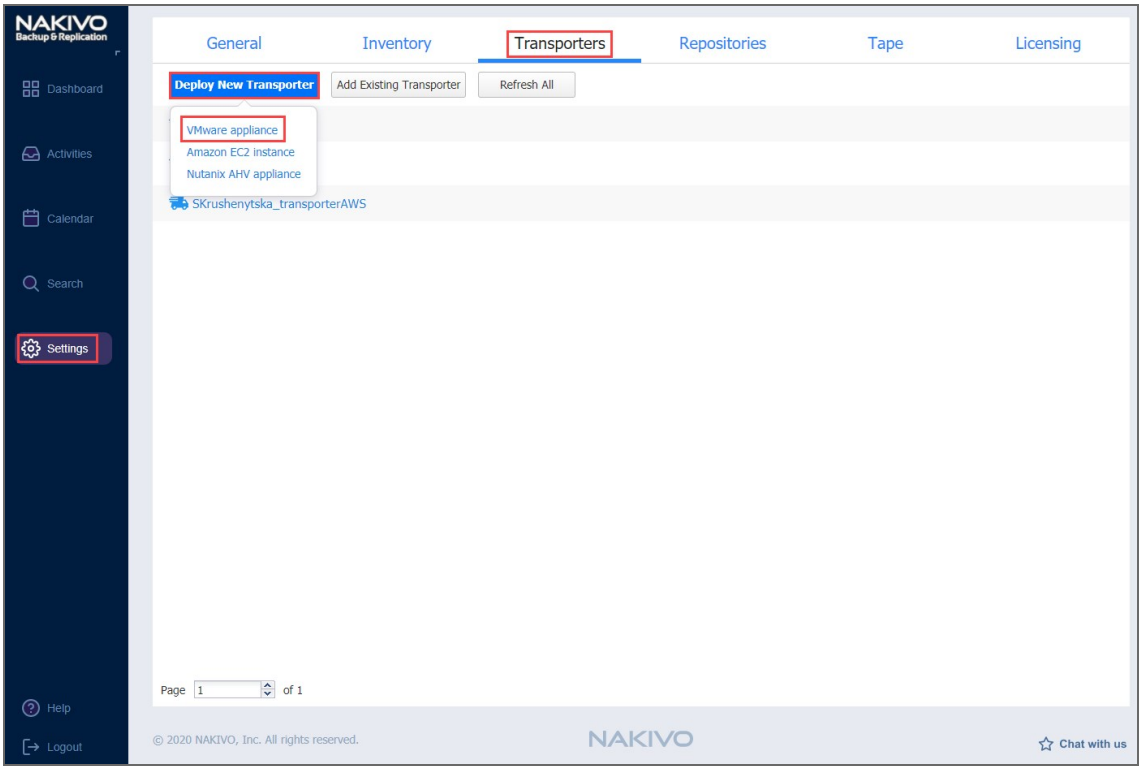

3. In the **Deploy New Transporter - VMware Appliance** dialog that opens, proceed as follows:
   - **Transporter name**: Enter a name for your Transporter.
   - **Host or cluster**: Select a target host or cluster.
   - **Datastore**: Select a target datastore.
   - **Virtual network**: Select a target 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.

4. Click **Deploy** to confirm deploying the Transporter.

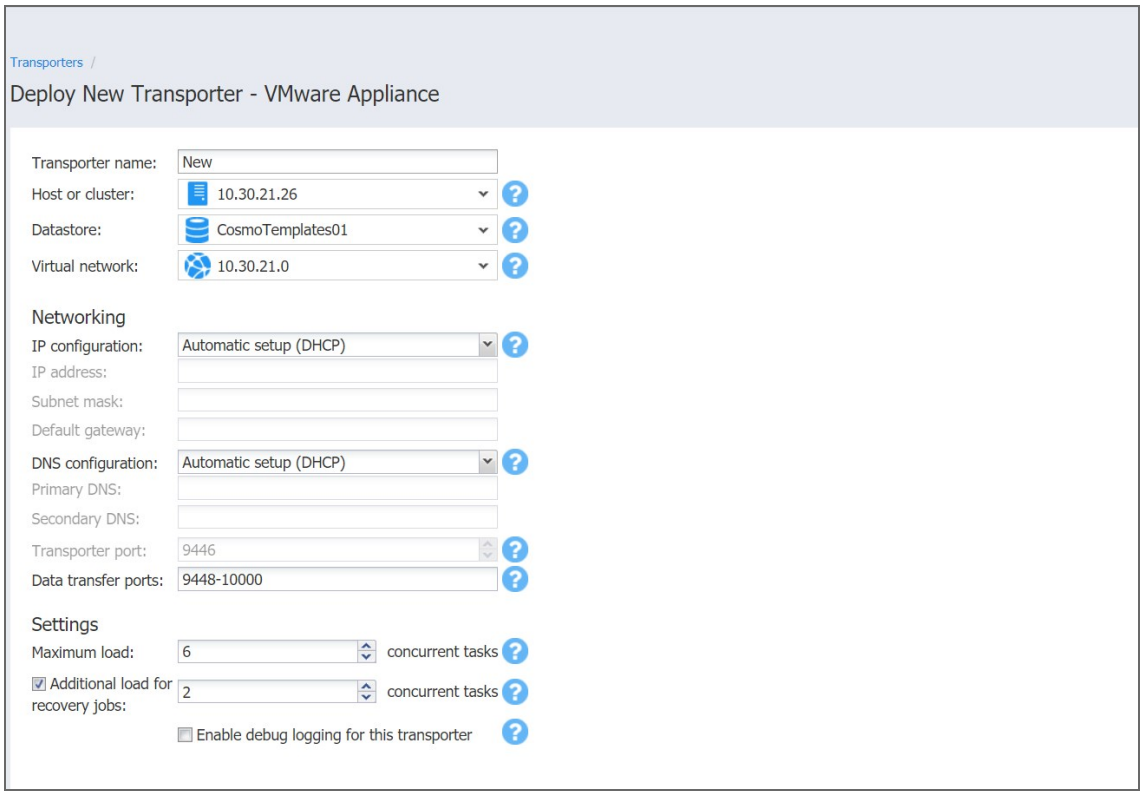

When deployment of the new Transporter finishes successfully, a message appears informing you about it.

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 **Settings > Transporters** tab.
2. On the **Transporters** tab, click **Deploy New Transporter** and select **Nutanix AHV Appliance** from the drop-down list.

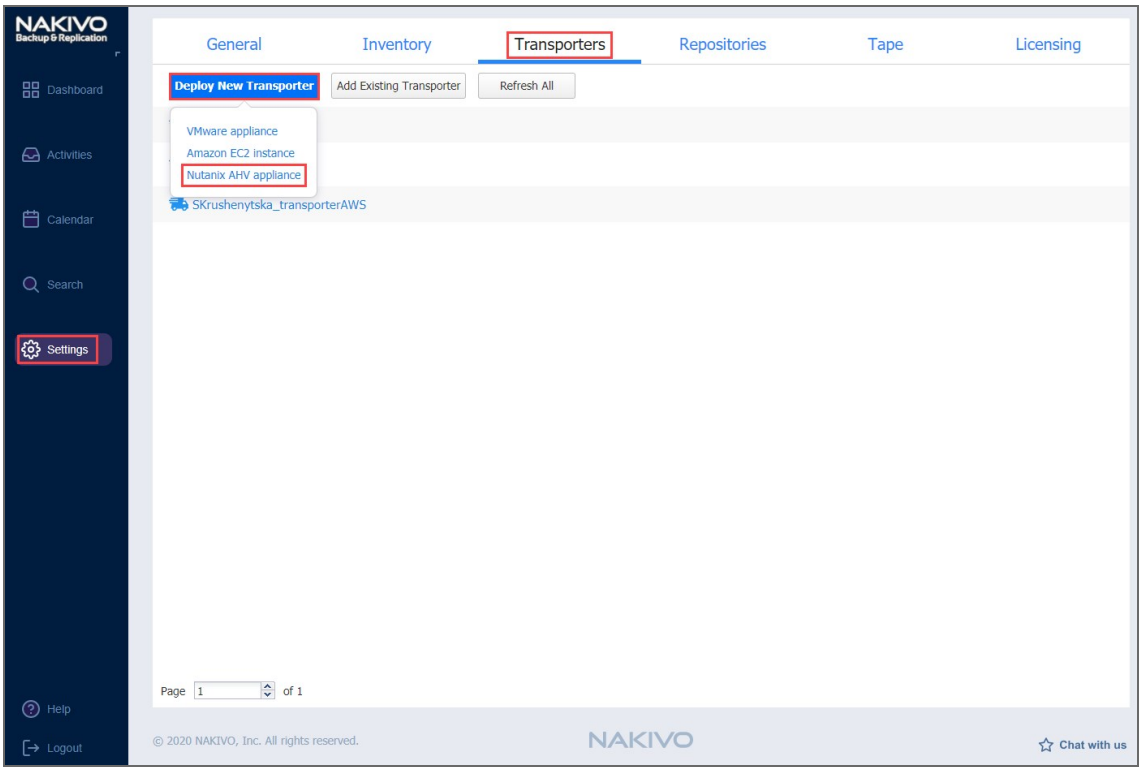

3. In the **Deploy New Transporter - Nutanix AHV Appliance** dialog, specify the following options:
   - **Transporter name**: Enter a name for the new Transporter.
   - **Cluster**: Select a cluster where the transporter VM will run.
   - **Storage container**: Select a storage container where the transporter VM will be located.
   - **Virtual network**: Select a virtual network where the transporter VM will be connected.

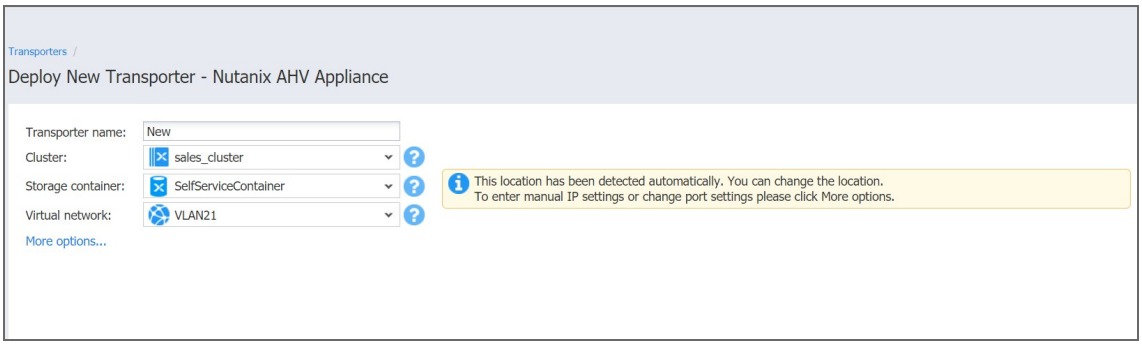

4. Click **Deploy** if you want to go with the automatically selected networking options and default Transporter load configuration.
5. Alternatively, click **More options** if you wish to manually set the following options:
   - **IP configuration**: Can be either **Automatic setup (DHCP)** or **Manual setup**. With manual setup selected, specify an **IP address**, **Subnet mask** and **Default gateway**.
   - **DNS configuration**: Can be either **Automatic setup (DHCP)** or **Manual setup**. With manual setup selected, specify **Primary** and **Secondary DNS**.
   - **Transporter port**: Enter a communication port for your Transporter.
   - **Data transfer ports**: Enter a port range that will be used by your Transporter for actual data transfer.
   - **Maximum load**: Specify the maximum number of tasks that the Transporter should process simultaneously. A task, for example, is a backup or replication of a single VM disk, or one granular recovery session.
   - **Additional load for recovery jobs**: If selected, the specified quantity of tasks will be added to set the maximum Transporter load to be used for recovery jobs exclusively. This allows for running the specified amount of concurrent recovery jobs along with other types of jobs without the need to wait for their completion.
   - **Enable debug logging for this transporter**: If needed, enable debug level logging for the current transporter. Using this option on a permanent basis is not recommended.

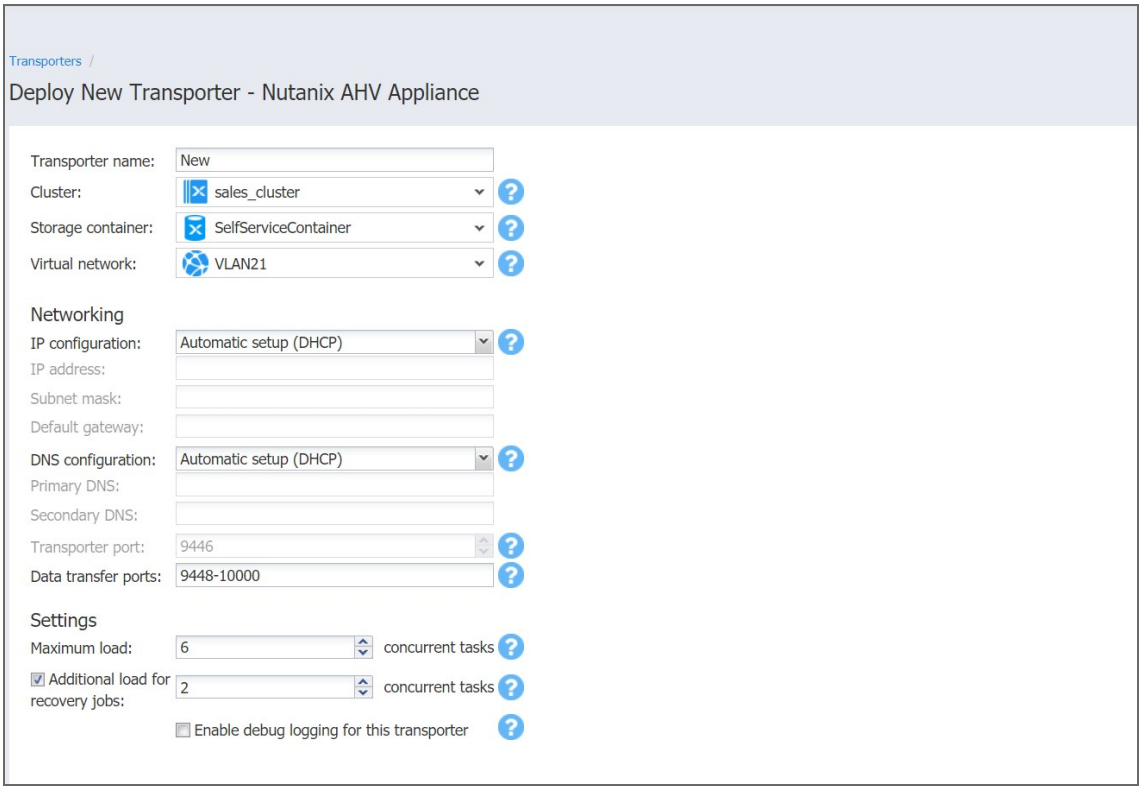

6. Click **Deploy**. The deployment process starts. Successfully deployed Transporter is displayed in the **Transporters** tab.

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 dashboard and then go to the **Transporters** tab.
2. Click **Deploy New Transporter** and then click **Amazon EC2 instance**.

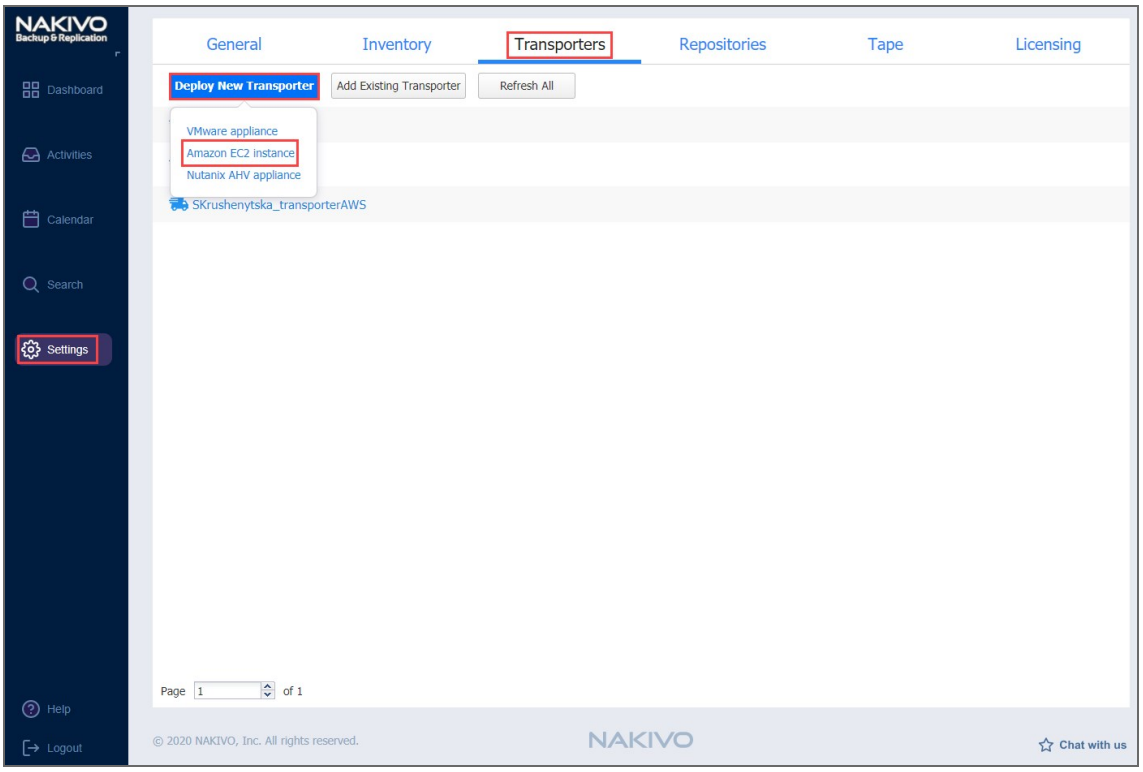

3. The **Amazon EC2 instance** dialog opens. Fill out the fields as described below:
   - **Transporter name**: Enter a name for the Transporter.
   - **Region**: Select an Amazon EC2 region where you wish to deploy the Transporter. This will enable you to create a backup repository in the region as 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** and do the following:
     - 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 only 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.

Note

Refer to "Amazon EC2 Concepts" on page 17 for the definitions of Amazon EC2-related terms.

4. Click **Deploy**.

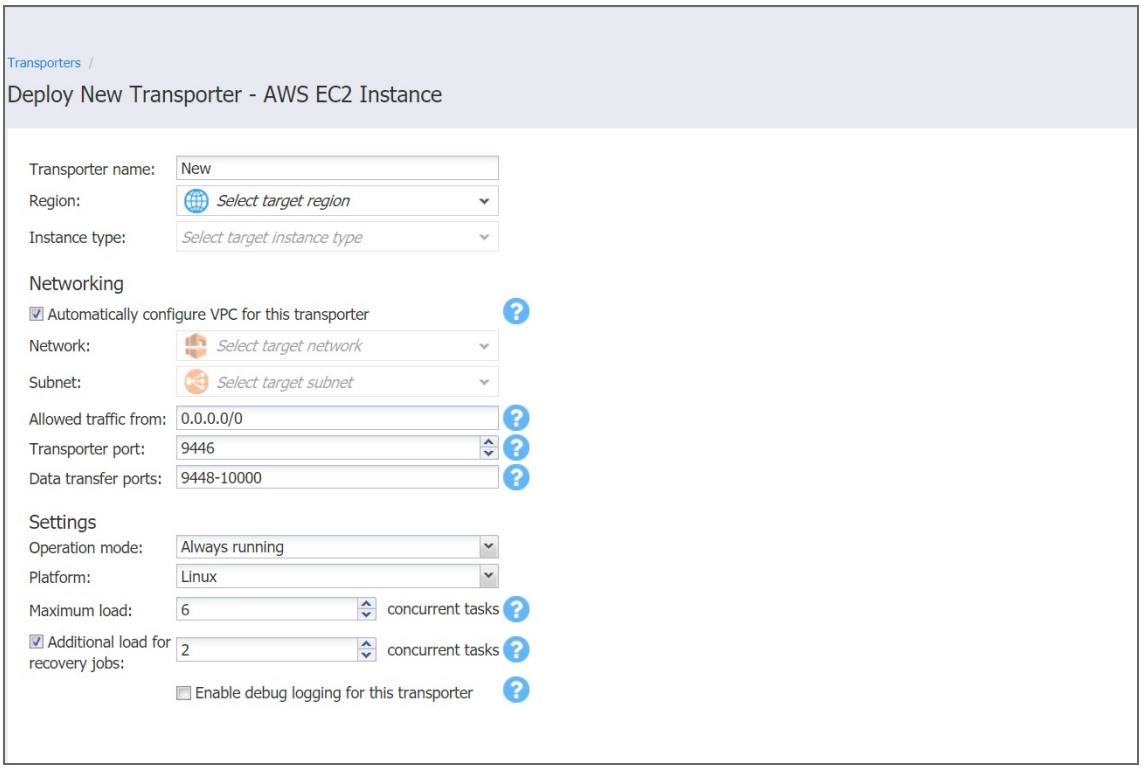

Important

- After you have deployed a Transporter in Amazon EC2, you need to download the Transporter Keys. The Transporter Keys are 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 Keys 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.

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

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

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 **Installed service** in the dialog that opens.

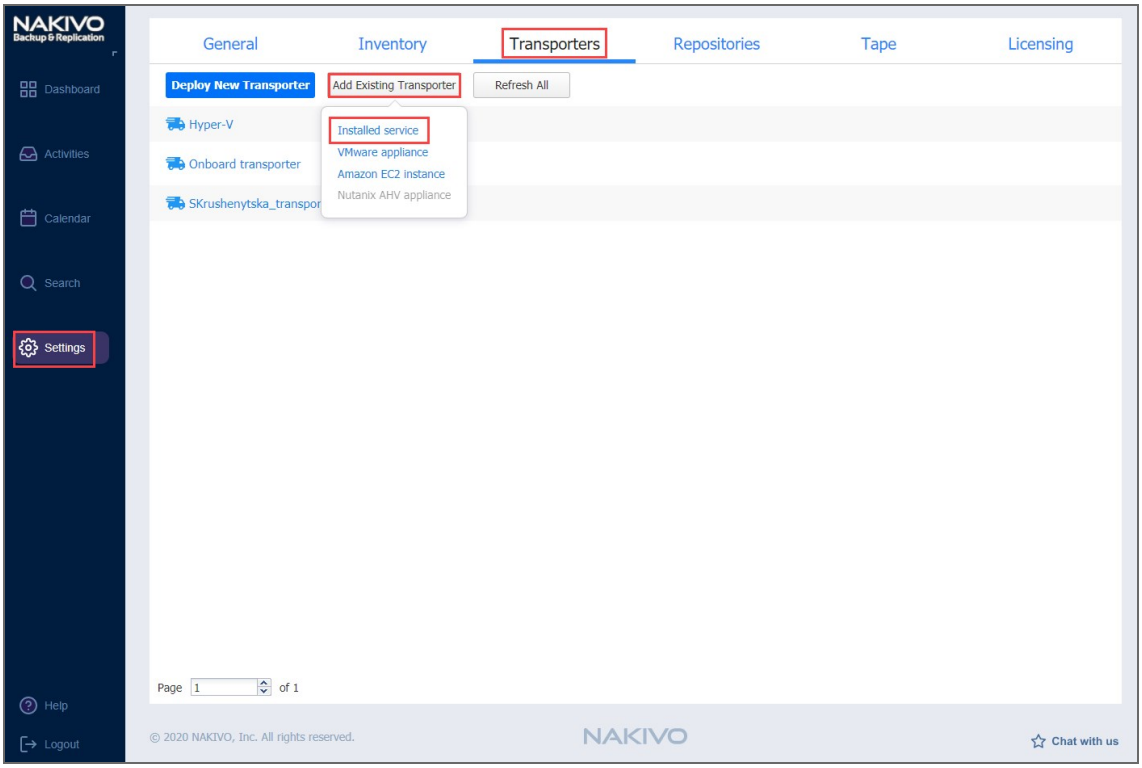

3. The **Add Existing Transporter - Installed Service** dialog opens. In the **Hostname or IP** box, enter the IP address or a hostname of the machine on which the Transporter is installed.

   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** and fill out 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 Acceptance** 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:
    - The 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 be externally available 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**.

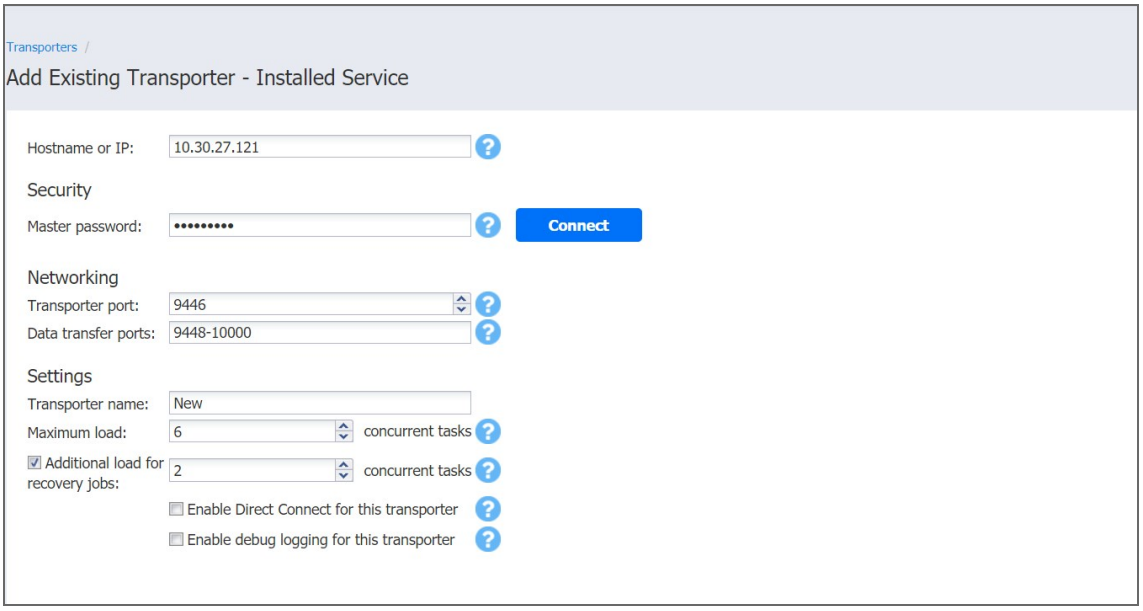

The Transporter is added to the product and is ready to be used for backup, replication, and recovery.

VMware Appliance

Please 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 appliance** in the dialog that opens.

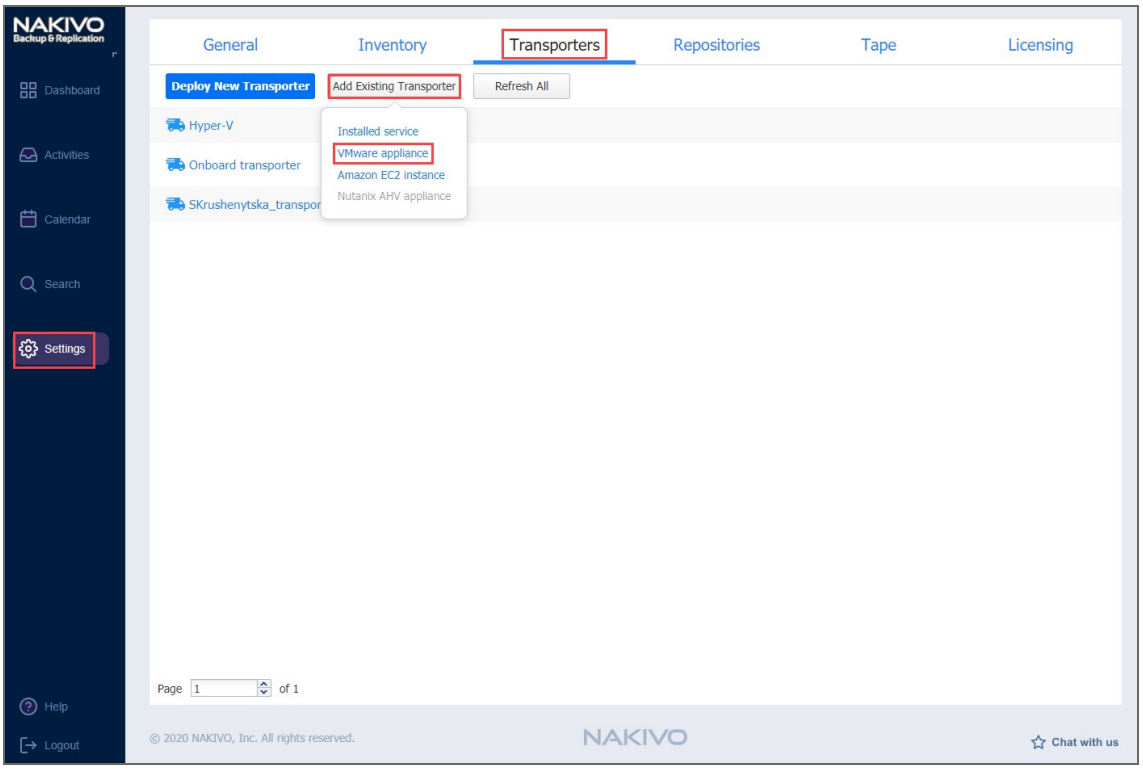

3. The `VMware Appliance` dialog opens. Fill out the fields as described below:
    - In the **Host or cluster** box, enter the IP address or name of the host or the cluster where the corresponding virtual machine is deployed.
    - In the **Virtual machine** box, specify the virtual machine on which the Transporter is installed.
    - In the **OS Username** and **OS Password** fields, specify credentials for accessing the virtual machine.
    - Click **More options** and fill out 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.

4. Click **Add**.

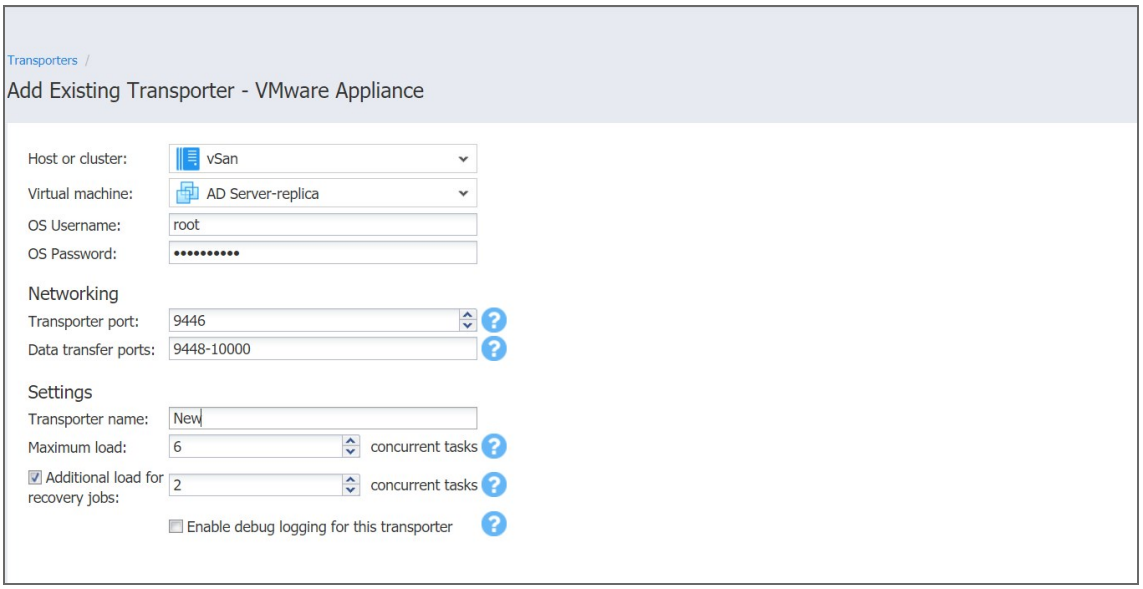

The Transporter is added to the product and is ready to be used for backup, replication, and recovery.

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 dashboard and go to the **Transporters** tab.
2. Click **Add Existing Transporter** and then click **Amazon EC2 instance** in the popup that opens.

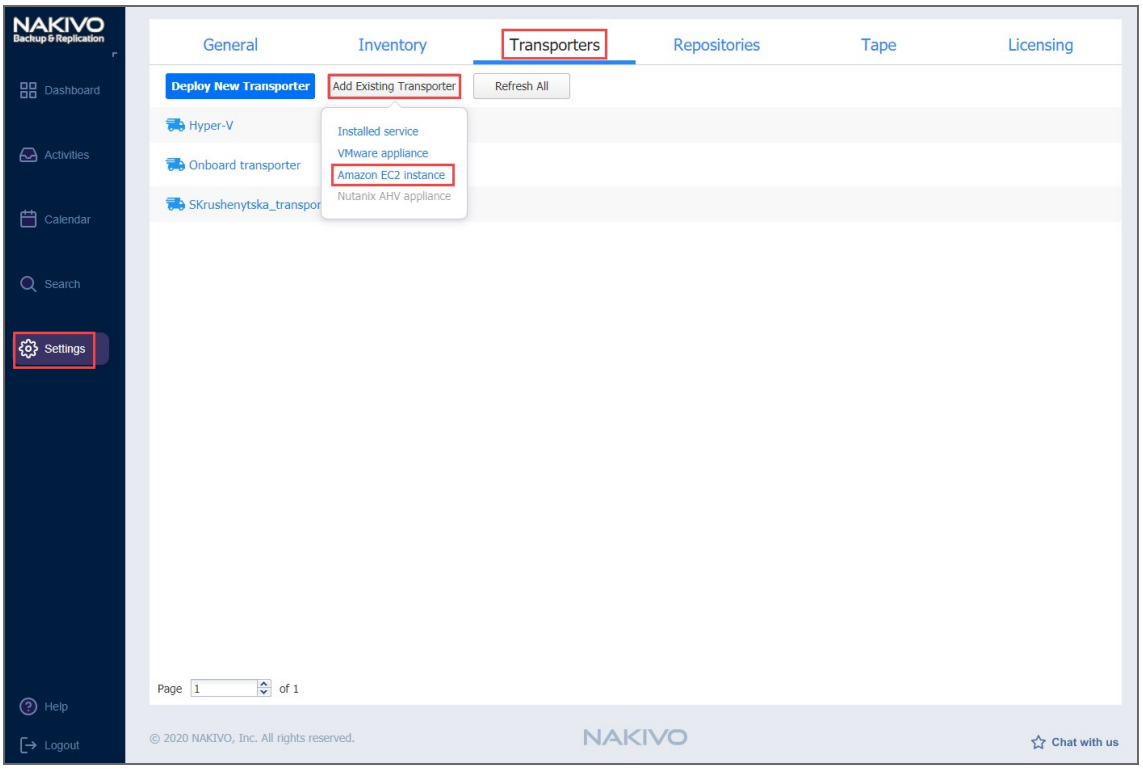

3. The **Amazon EC2 Instance** dialog opens. Fill out the fields as described below:
    - **AWS account**: Choose an appropriate Amazon AWS Account from the list of Amazon AWS Accounts added to the Inventory.
    - **Region**: Choose a region in which an AWS EC2 instance with the Transporter is deployed.
    - **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** and fill out 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 either 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**.

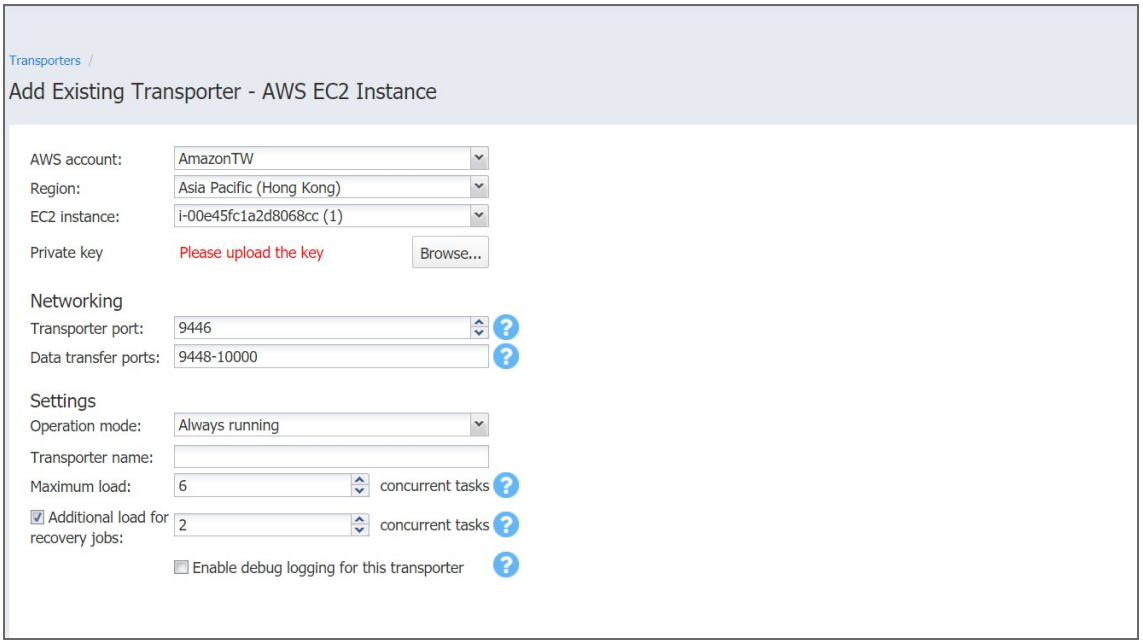

The Transporter is added to the product and is ready to be used for backup, replication, and recovery tasks.

Nutanix AHV Appliance

Please 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 dashboard and go to the **Transporters** tab.
2. Click **Add Existing Transporter** and then select **Nutanix AHV Appliance**.

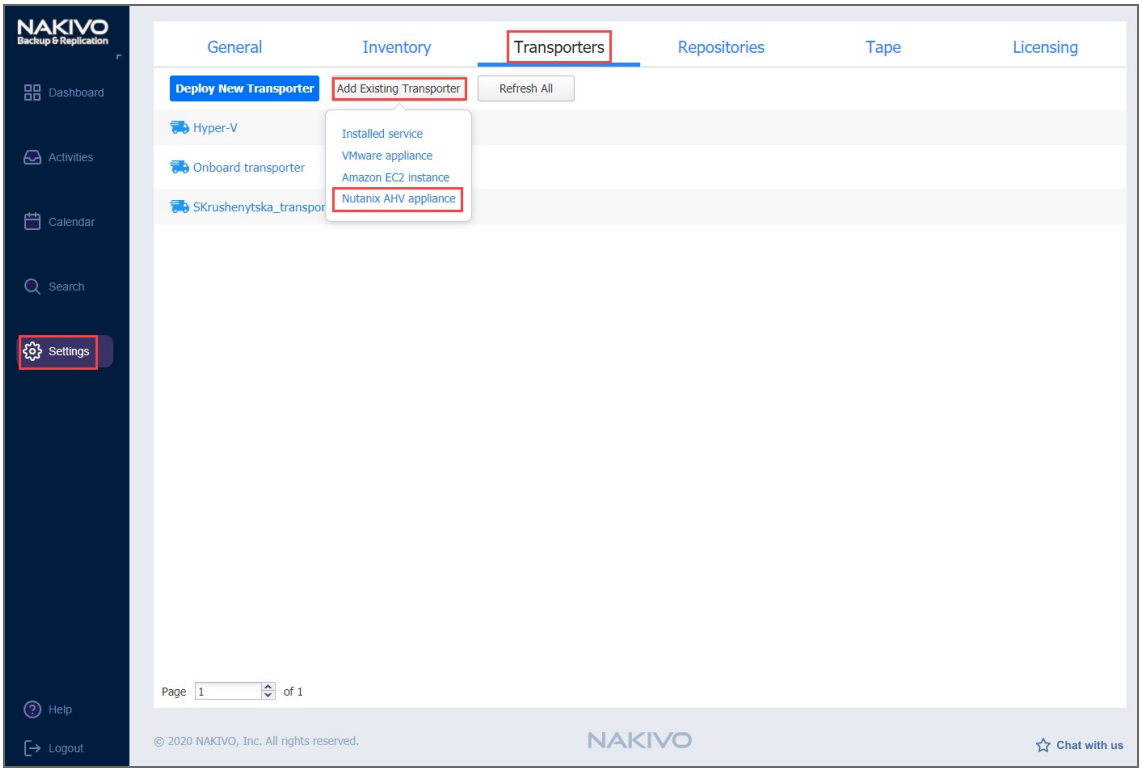

3. In the **Add Existing Transporter - Nutanix AHV Appliance** dialog, enter the following options:
   - **Cluster**: Select a cluster where the corresponding virtual machine is deployed.
   - **Virtual machine**: Select the virtual machine on which the Transporter is installed.
   - **Username/Password**: Enter credentials for accessing the virtual machine where the Transporter is installed.
   - **Transporter port**: Enter the port number that will be used to connect to the Transporter.
   - **Data transfer ports**: Enter 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.
   - **Transporter name**: Enter a display name for the Transporter.
   - **Maximum load**: Select 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.

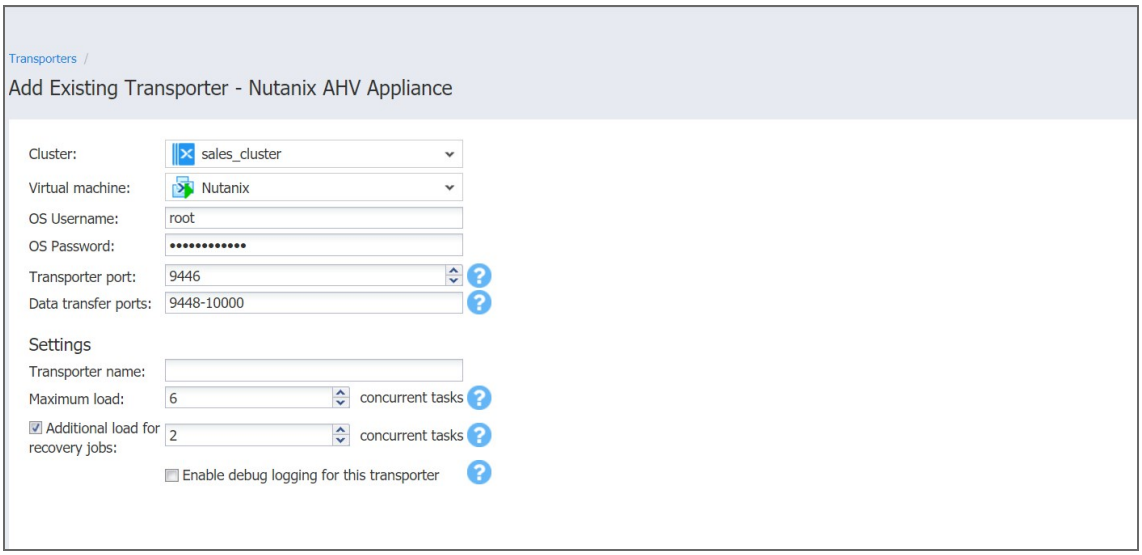

4. Click **Add**. The Transporter is added to the product and can be used for backup, replication, and recovery.

Managing Transporters

Refer to the following topics:

- “Editing Transporters” on page 364
- “Refreshing Transporter Details” on page 366
- “Downloading Transporter's Credentials” on page 363
- “Removing (Deleting) Transporters” on page 368

Downloading Transporter's Credentials

If you want to re-import an Amazon EC2 or VMware Transporter into another installation of the NAKIVO Backup & Replication, you need to have the Transporter's credentials. To obtain the credentials, click the **Download key** link next to the corresponding Transporter.

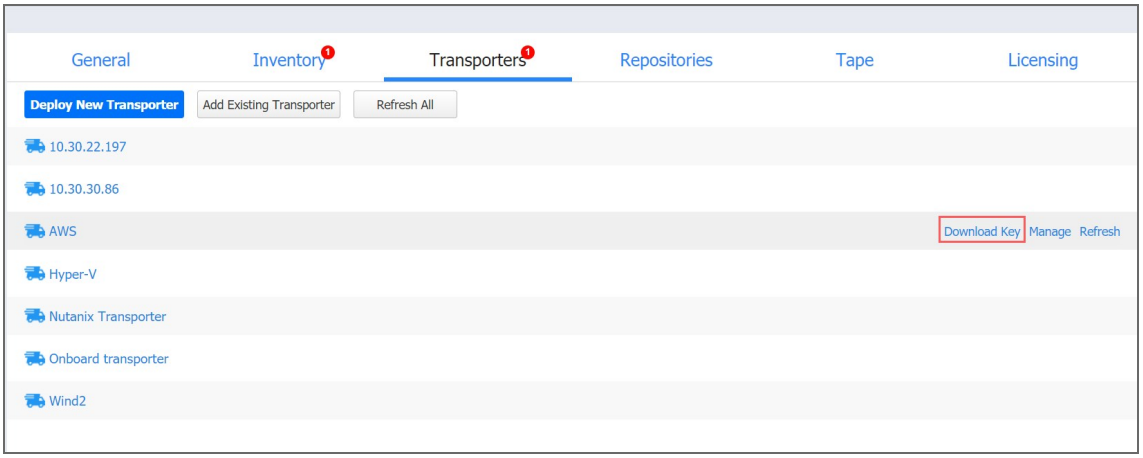

Clicking the link downloads the ZIP file containing the Transporter's credentials.

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 click on the Transporter you would like to edit.
3. In the Transporter title, click **Manage** and then click **Edit**.

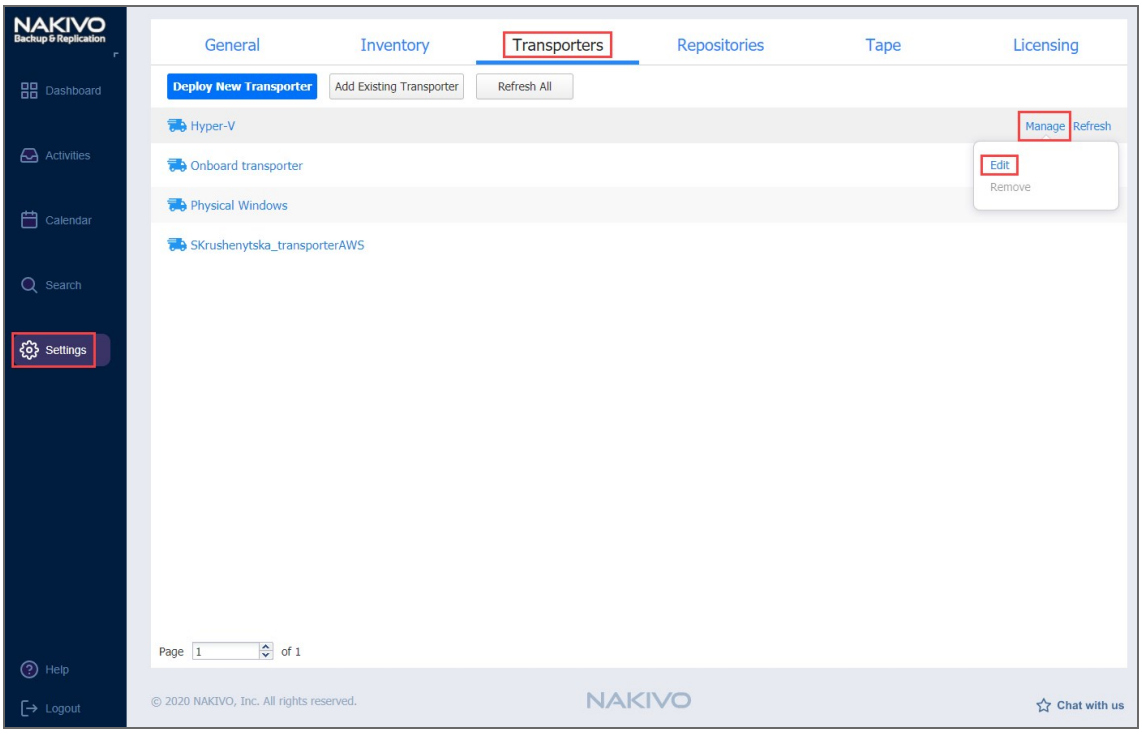

4. A dialog opens for editing the Transporter settings. Edit the settings as required:
    - **Hostname or IP**: Here you can edit the IP address or hostname of the machine on which the Transporter is installed.
        - In the *Networking* section:
            - **Transporter port**: Enter a communication 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 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 **Appy** to save your changes.

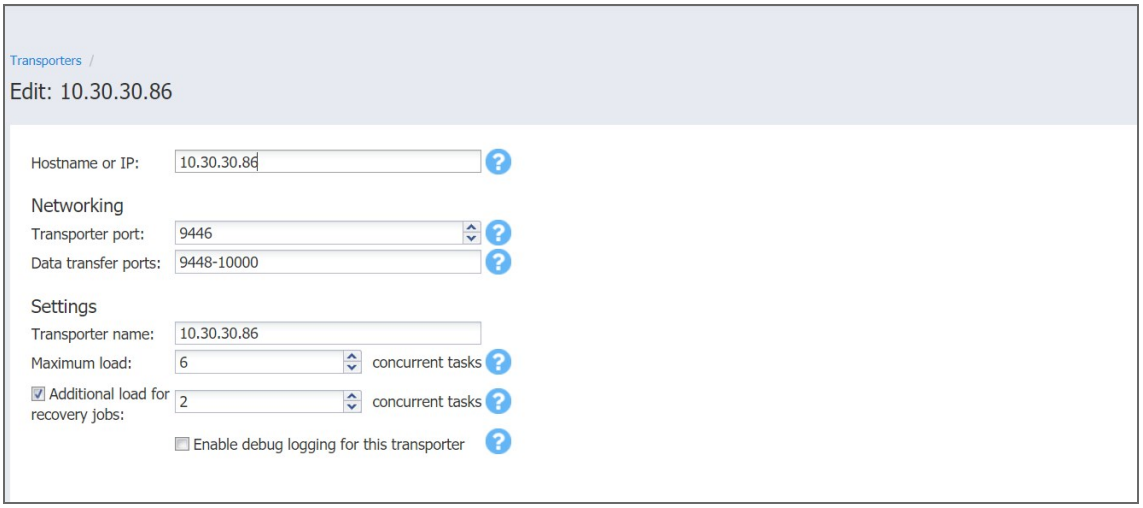

The changes you have made are applied to the Transporter.

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 update 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.**

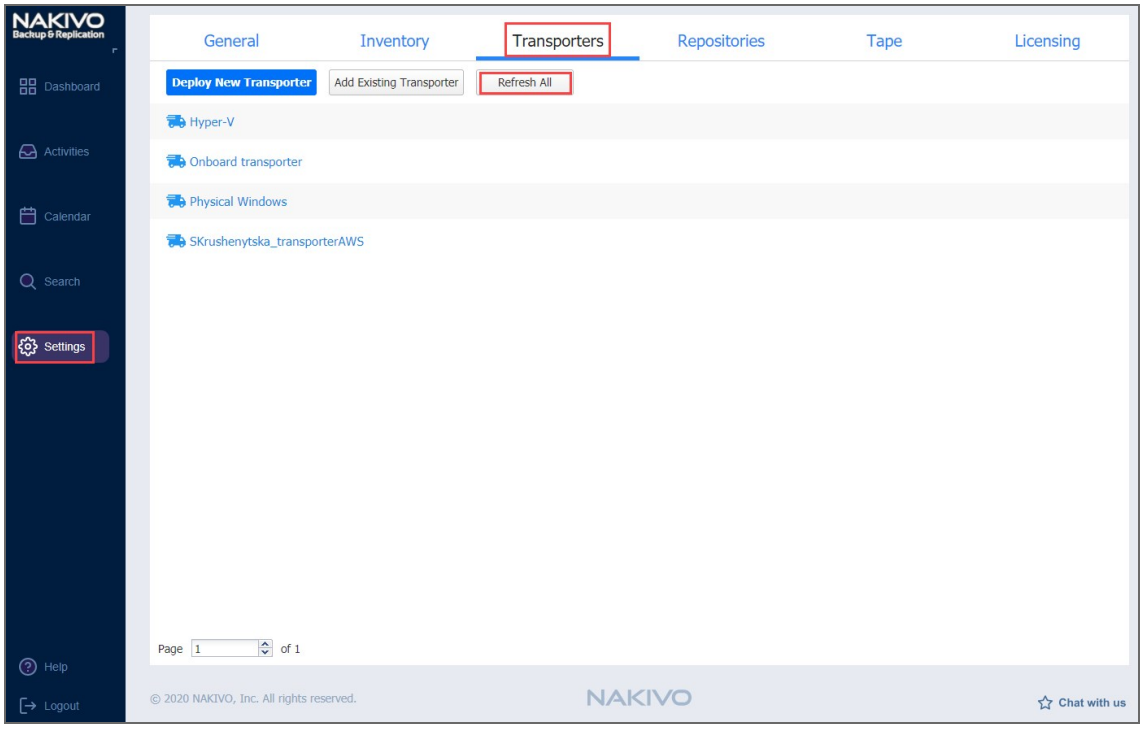

The update of all Transporters starts.

Manually Refreshing a Single Transporter

To update a single Transporter, follow the steps below:

1. Click **Settings** in the left pane of the product.
2. Go to the **Transporters** tab.
3. Select the Transporter you would like to update.

4. In the title of the Transporter, click **Refresh.**

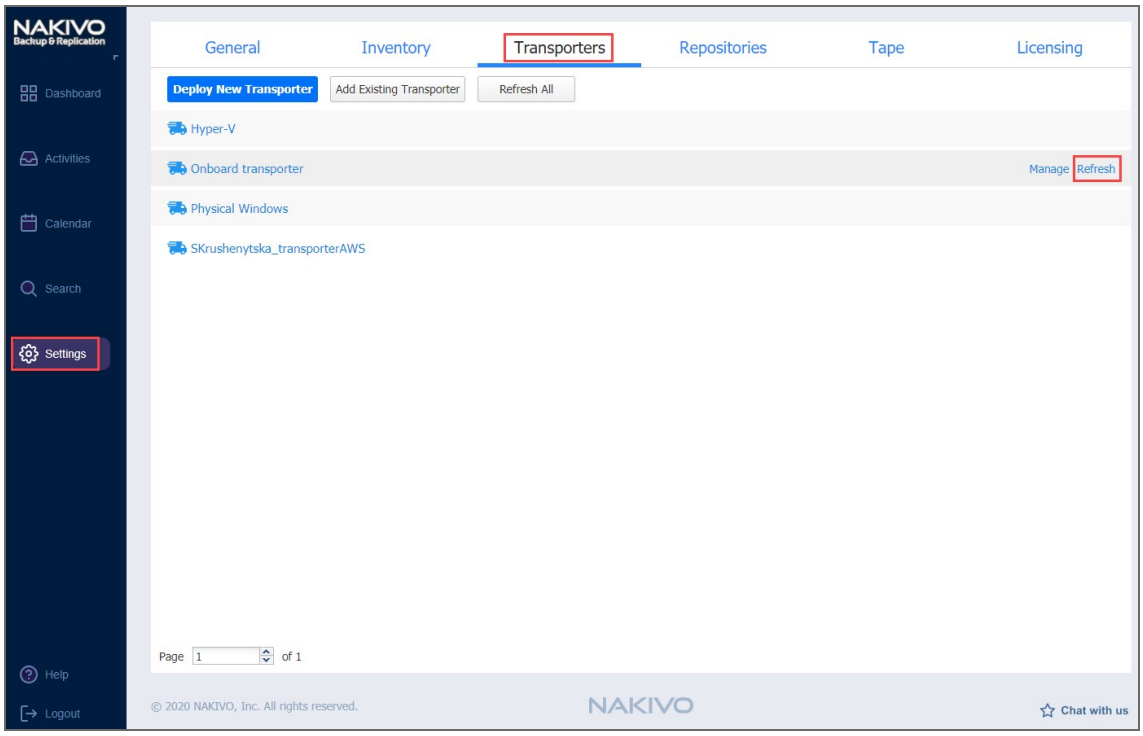

The Transporter refresh starts.

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. Select the Transporter you would like to remove.
4. In the Transporter title, click **Manage** and then click **Remove.**

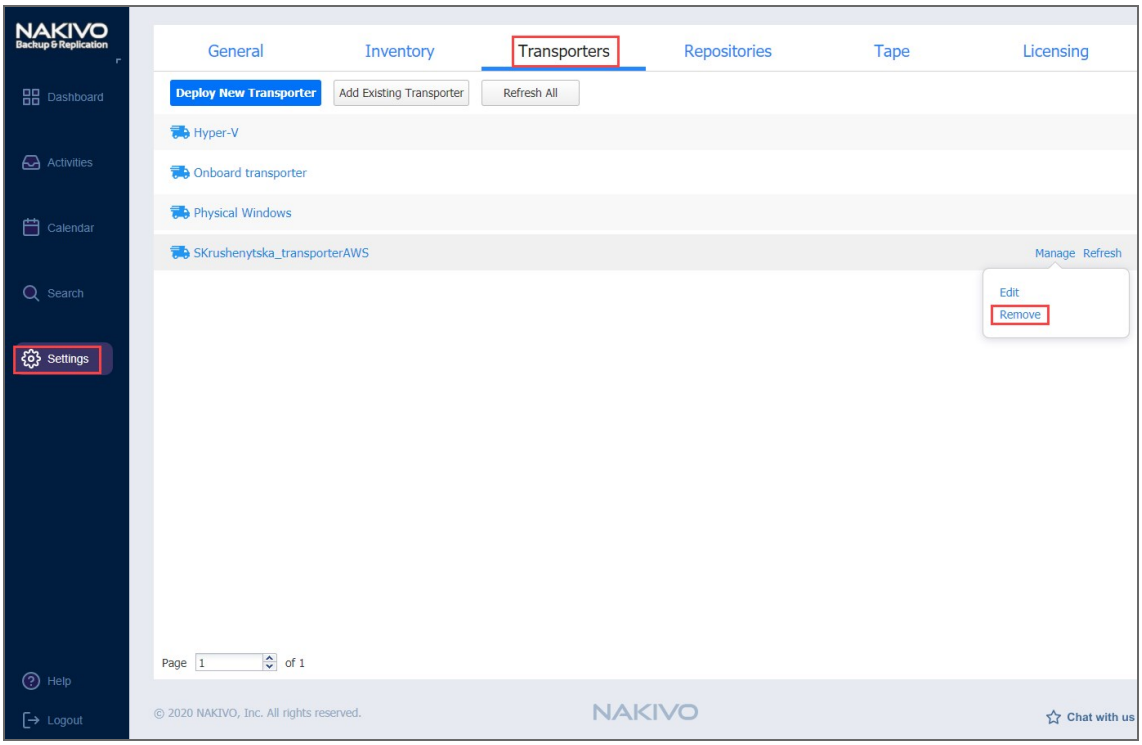

5. Click **Remove** in the message that appears.

Important

The following Transporters cannot be removed:

- The Onboard Transporter (which is installed with the "Director" on page 86 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 92.

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

- “Creating Backup Repositories” on page 373
- “Adding Existing Backup Repositories” on page 370
- “Viewing Backup Repository Details” on page 408
- “Managing Backup Repositories” on page 412

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.

Important

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. Go to the main menu and click **Settings**.
2. Go to the **Repositories** tab and click **Add Backup Repository**.
3. Click **Add existing backup repository** in the dialog that opens.

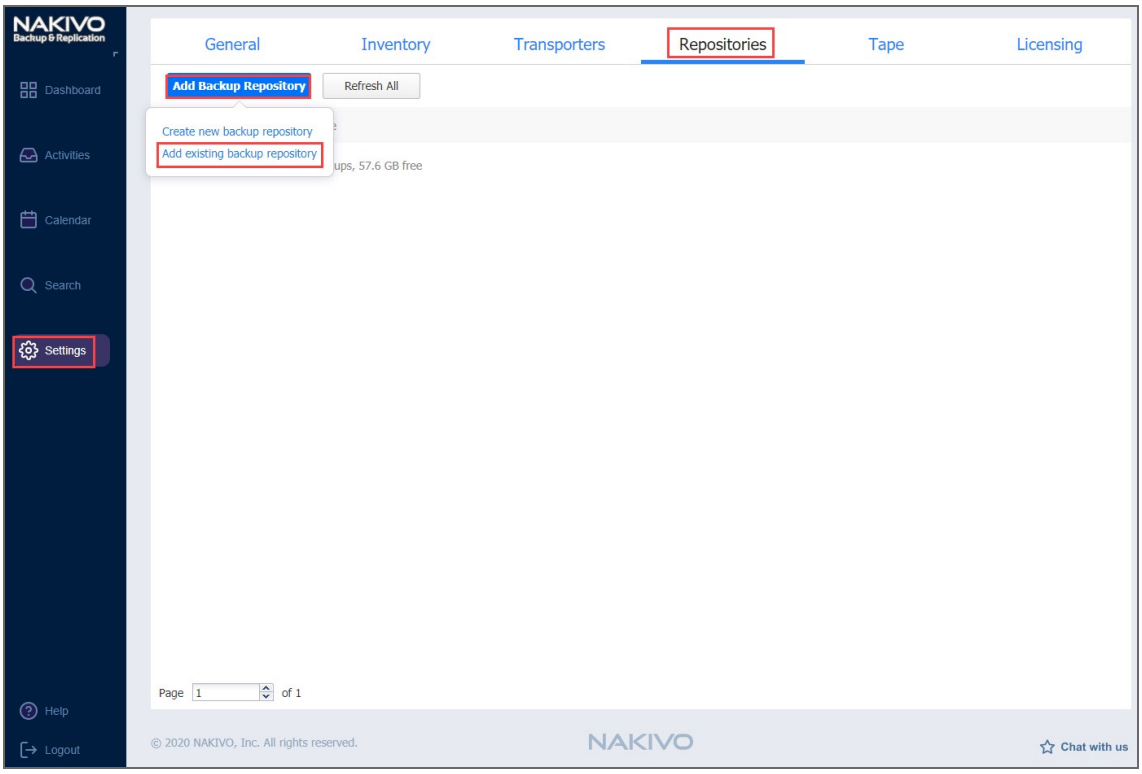

4. The **Add Existing Backup Repository** wizard opens. On the **Type** page of the wizard, select one of the following Backup Repository types:
   - **Local Folder**
   - **CIFS Share**
   - **NFS Share**
   - **Amazon EC2**
   - **Amazon S3**

- **Wasabi**
- **Deduplication Appliance**

5. On the **Name & Location** page of the wizard, fill out all the necessary fields the way it's 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.

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 can be lost after the solution update.

1. In the NAKIVO Backup & Replication, navigate to **Settings**.
2. Go to the **Repositories**tab and click **Add Backup Repository**.
3. Click **Create new backup repository**.

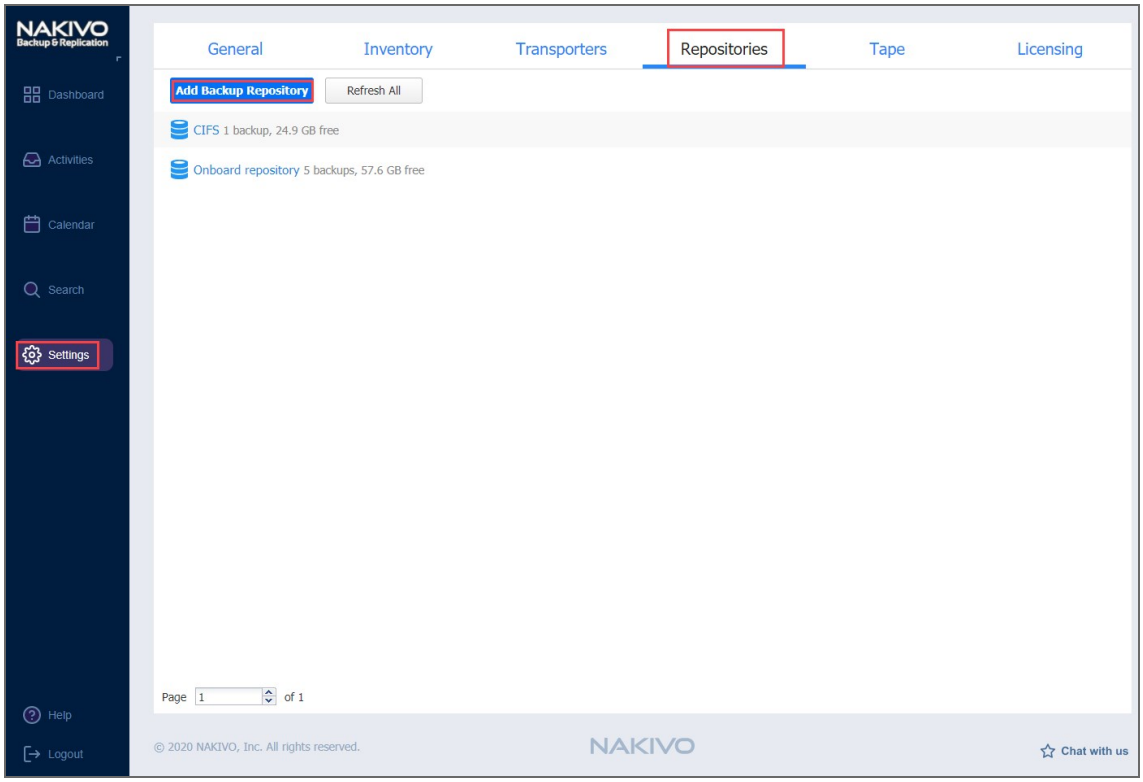

Choose one of the locations for storing your backups by completing the **Create Backup Repository** wizard as it's described in the sections below:

- "Local Backup Repository" on page 374
- "Backup Repository on CIFS Share" on page 384
- "Backup Repository on NFS Share" on page 379
- "Backup Repository on Deduplication Appliance" on page 401

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.

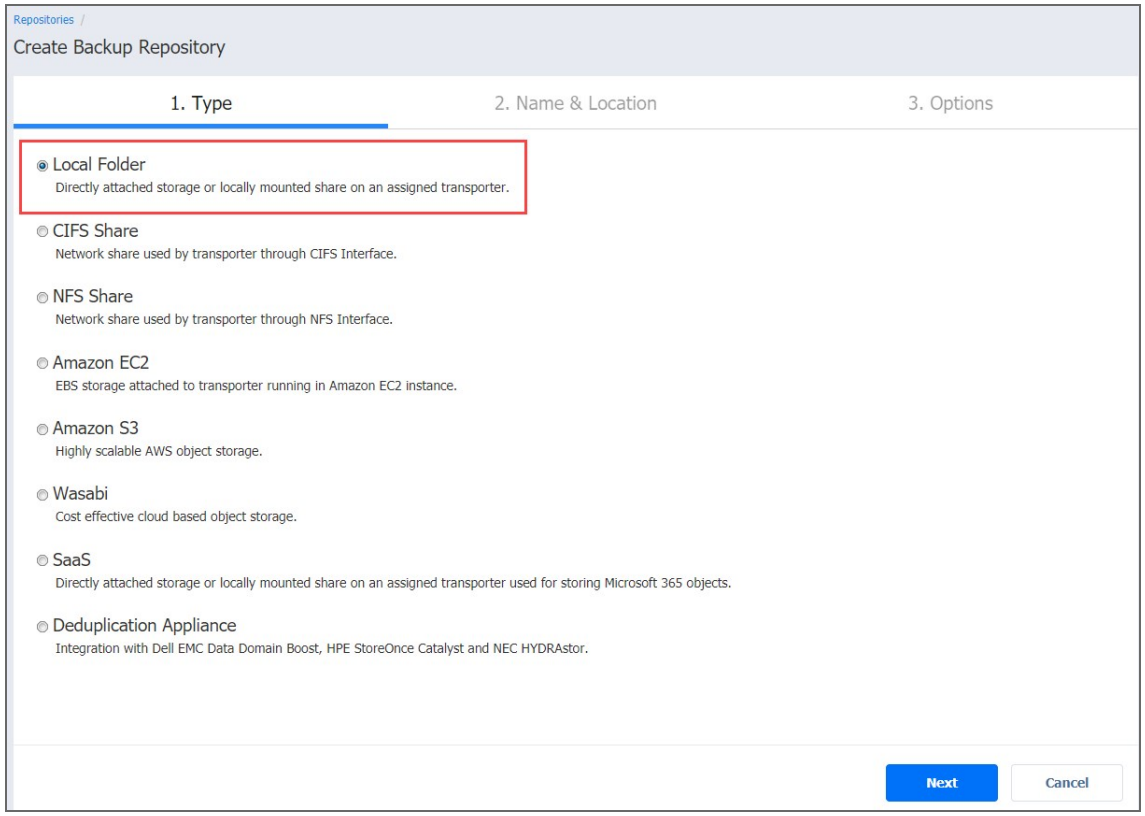

Create Backup Repository: Name and Location

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.

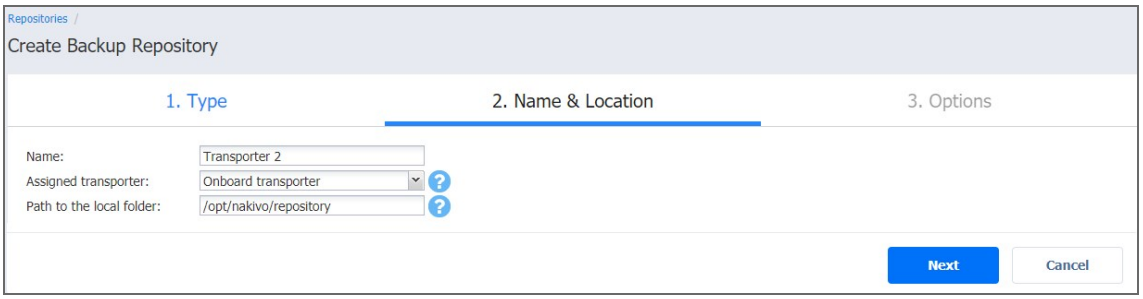

Create Backup Repository: Options

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.

- **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.

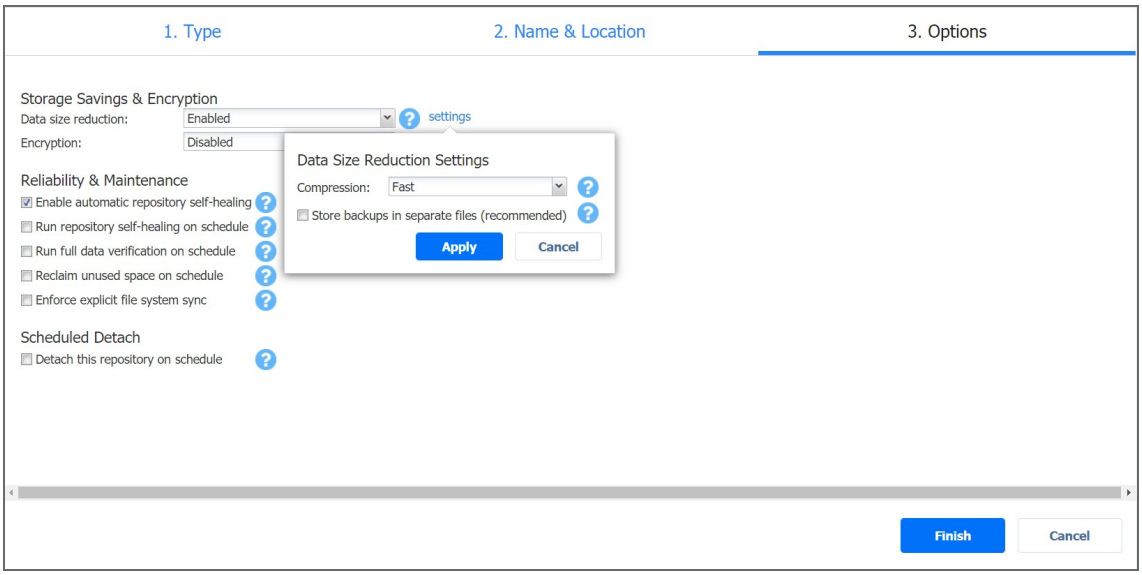
1. Type
2. Name & Location
3. Options
Storage Savings & Encryption
Data size reduction:
Enabled
settings
Encryption:
Disabled
Reliability & Maintenance
Enable automatic repository self-healing
Run repository self-healing on schedule
Run full data verification on schedule
Reclaim unused space on schedule
Enforce explicit file system sync
Scheduled Detach
Detach this repository on schedule
Data Size Reduction Settings
Compression:
Fast
Store backups in separate files (recommended)
Apply
Cancel
Finish
Cancel

Backup Repository on NFS Share

Choose this option if you wish to create a Backup Repository on an NFS share. Before creating a Backup Repository on an NFS share, make sure that all the necessary prerequisites are met:

- The folder where you would like to create the Backup Repository exists on the share.
- The share can be accessed from the machine on which the Assigned Transporter is installed.
- You are using credentials with read and write permissions to the share.

To create a repository on an NFS share, proceed as described in the following sections:

- Create Backup Repository: Type
- Create Backup Repository: Name and Location
- Create Backup Repository: Options

Create Backup Repository: Type

On the **Type** page of the **Create Backup Repository** wizard, select **NFS Share** and click **Next** to move to the next page of the wizard.

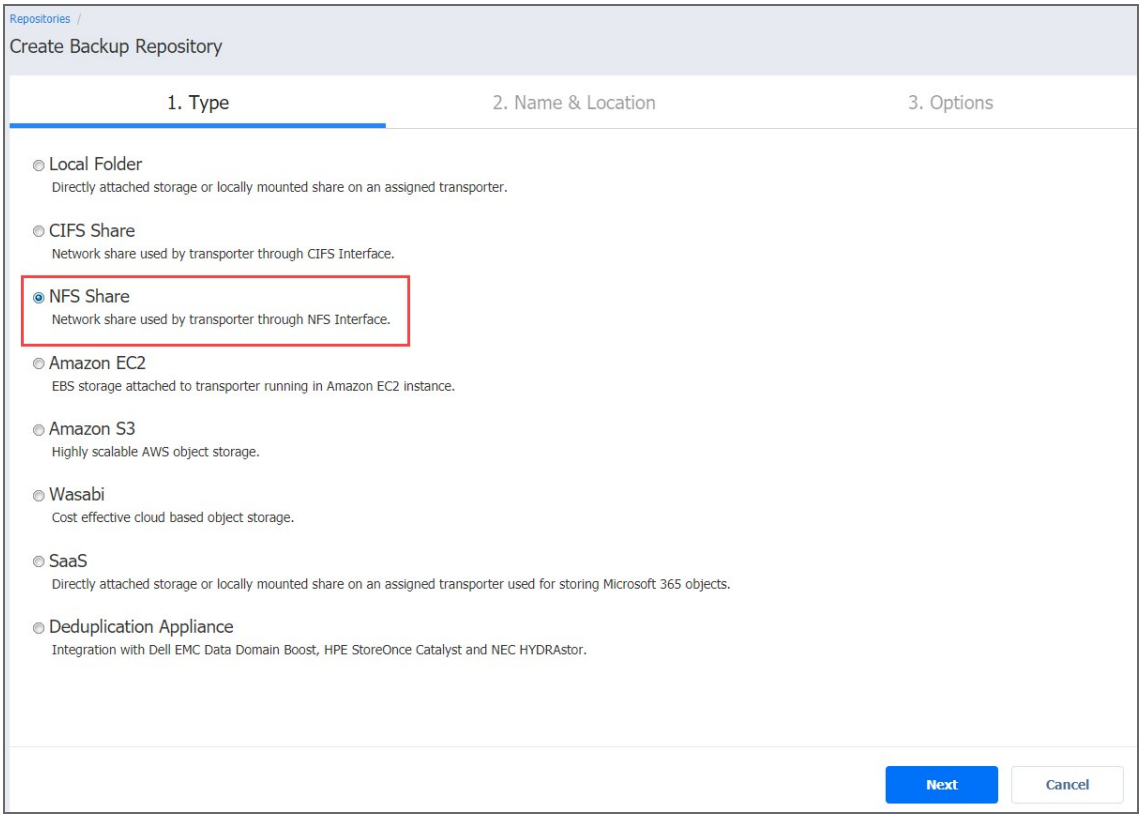

Create Backup Repository: Name and Location

On the **Name & Location** page of the wizard, do the following

1. Enter the name of the Backup Repository in the **Name** box.
2. Select the Transporter from the **Assigned transporter** drop-down list.
3. Enter the path to the NFS share.

   Examples

   QNAP share path: `10.30.30.109:/ayunt_nfs`

   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.

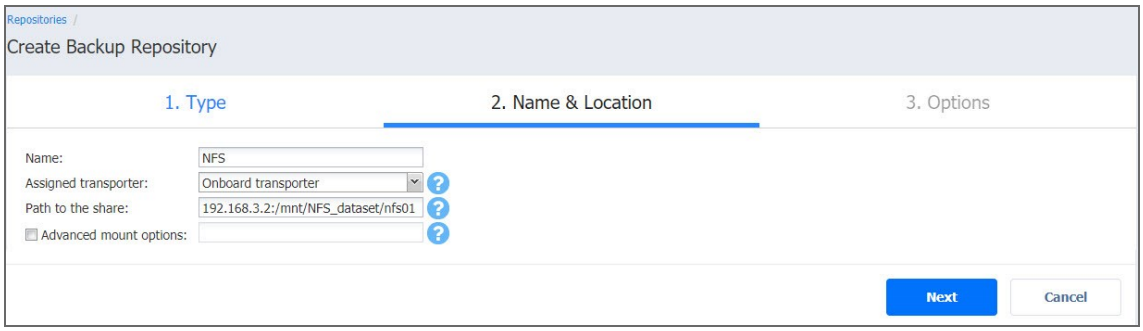

Create Backup Repository: Options

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.

- **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.

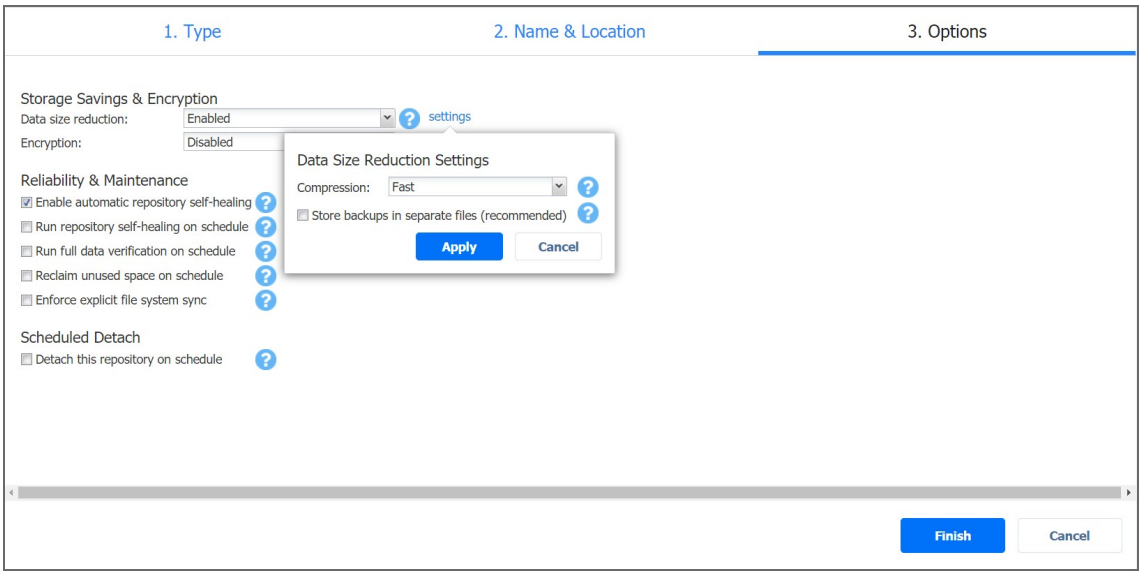

Backup Repository on CIFS Share

Choose this option if you want to create a Backup Repository on a Windows CIFS share. Before creating a Backup Repository on a CIFS share, make sure that all the necessary prerequisites are met:

- The folder where you would like to create the Backup Repository exists on the share.
- The share can be accessed from the machine on which the Assigned Transporter is installed.
- You are using credentials with read and write permissions to the share.
- The share is compatible with Version 2 or later of the SMB protocol.

To create a Backup Repository on a Windows CIFS share, proceed as described in the following sections:

- Create Backup Repository: Type
- Create Backup Repository: Name and Location
- Create Backup Repository: Options

Create Backup Repository: Type

On the **Type** page of the **Create Backup Repository** wizard, select **CIFS Share** and click **Next** to move to the next page of the wizard.

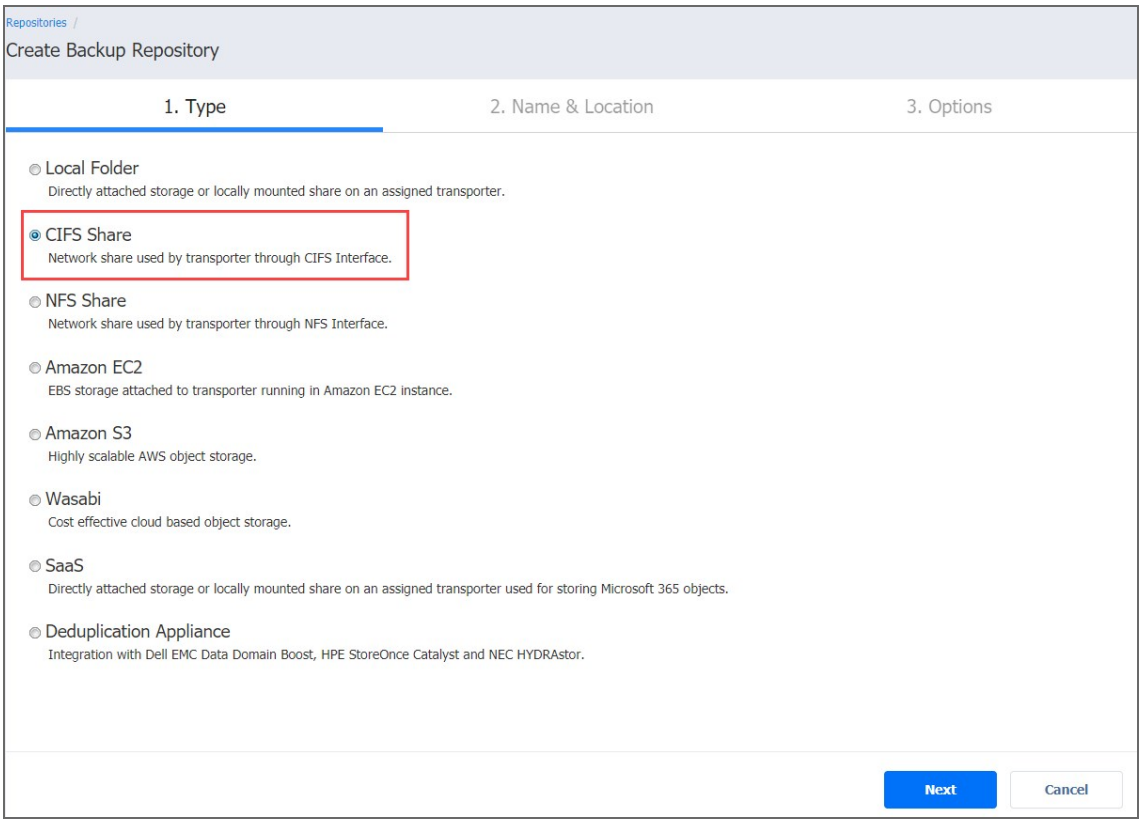

Create Backup Repository: Name and Location

On the **Name & Location** page of the wizard, do the following

1. Enter the name of the Backup Repository in the **Name** box.
2. Select the Transporter from the Assigned transporter drop-down list.
3. Enter the path to the CIFS share.

   Example

   Synology share path: `\\10.30.30.61\ayunt_cifs1`
4. Provide username and password in the appropriate boxes.

   Note

   If you're using domain credentials to access the share, enter your domain username via the following format: domain\username.
5. Select **Advanced mount options** if needed. Refer to the mount `man` pages for a detailed description of CIFS share mount options.
6. Click **Next** to go to the next page of the wizard.

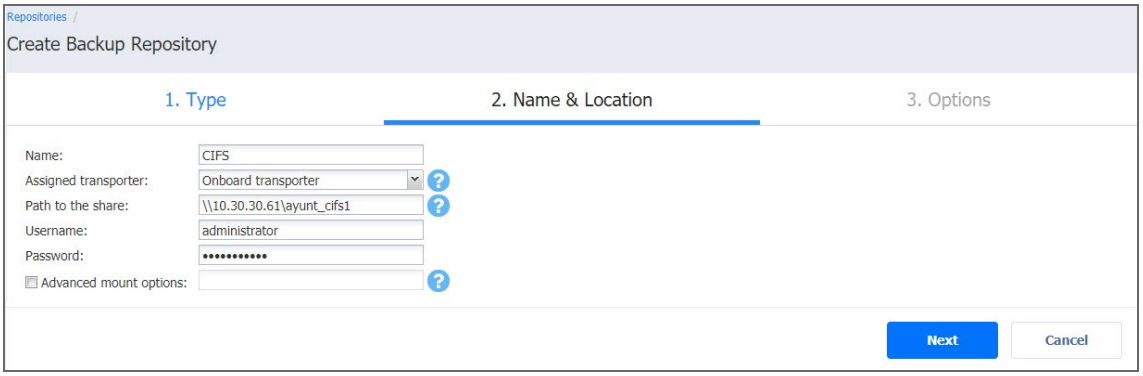

Create Backup Repository: Options

On the **Options** page, do the following:

1. Set up *Storage Savings & Encryption* options:
   - **Data size reduction**: If this option is enabled, NAKIVO Backup & Replication enables the use of data size reduction for this repository to save disk space. Note that this may put additional load on the CPU. Disabling data size reduction is required if the target is a deduplication storage appliance. Click settings to configure the settings. A popup window appears. Set the following:
     - **Compression**: Select a compression level that will be used to reduce the data size in the Backup Repository. Note that higher compression levels consume considerably more CPU and may slow down VM backup speed. The following options are available:
       - **Disabled:** The data in the Backup Repository will not be compressed.
       - **Fast:** Lowest compression level.
       - **Medium:** Medium compression level.
       - **Best:** Maximum compression level.

Note

This option cannot be configured after creating the Backup Repository.

- **Store backups in separate files**: Select this option to enable this backup repository to store data of every machine in separate backup files. Enabling this option is highly recommended to ensure higher reliability and performance.

- **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.

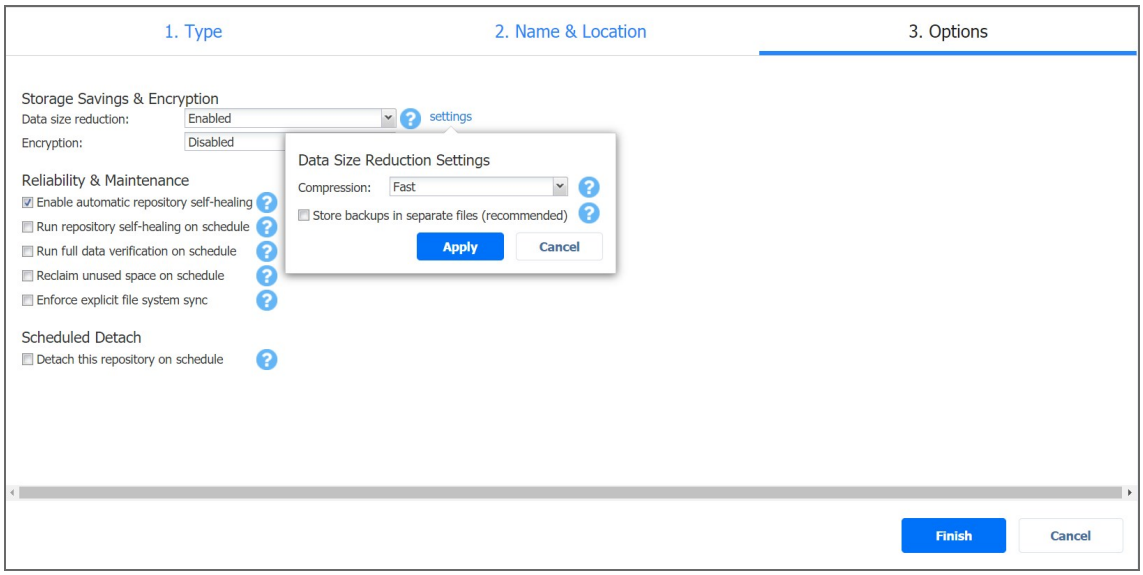
1. Type
2. Name & Location
3. Options
Storage Savings & Encryption
Data size reduction:
Enabled
settings
Encryption:
Disabled
Reliability & Maintenance
Enable automatic repository self-healing
Run repository self-healing on schedule
Run full data verification on schedule
Reclaim unused space on schedule
Enforce explicit file system sync
Scheduled Detach
Detach this repository on schedule
Data Size Reduction Settings
Compression:
Fast
Store backups in separate files (recommended)
Apply
Cancel
Finish
Cancel

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 Transported 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.

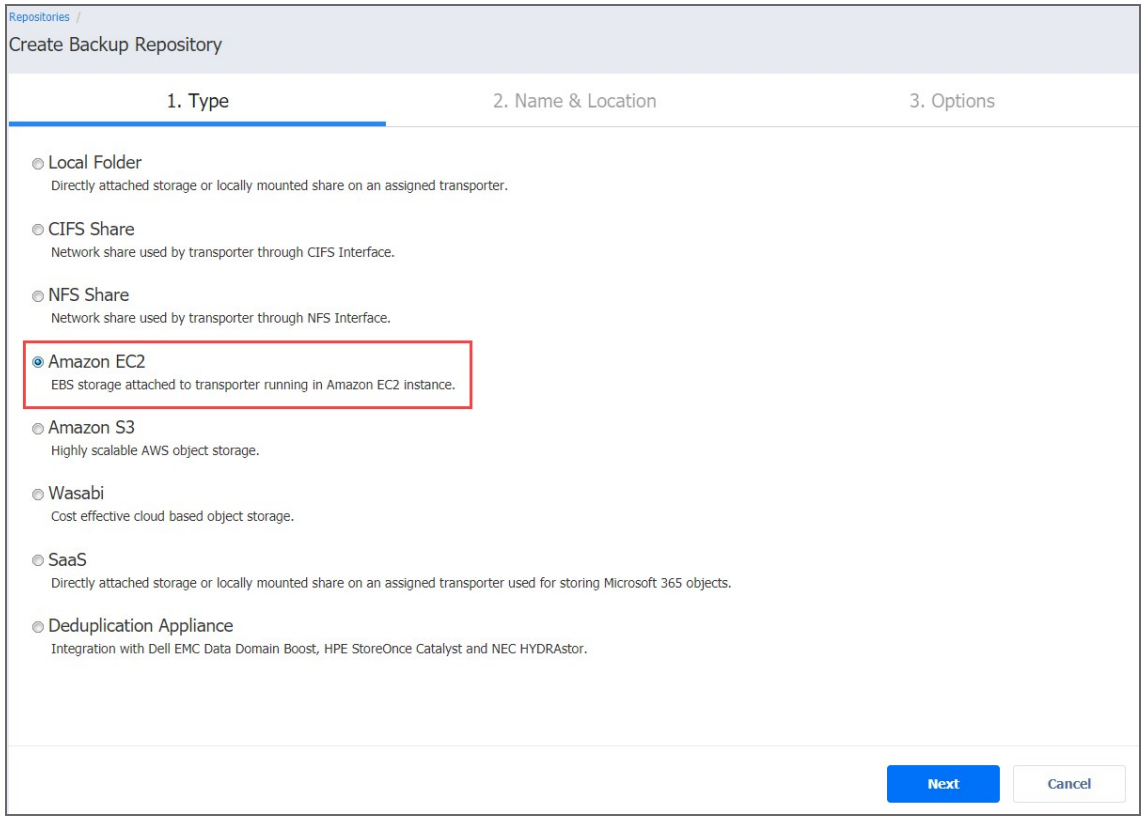

Create Backup Repository: Name and Location

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.

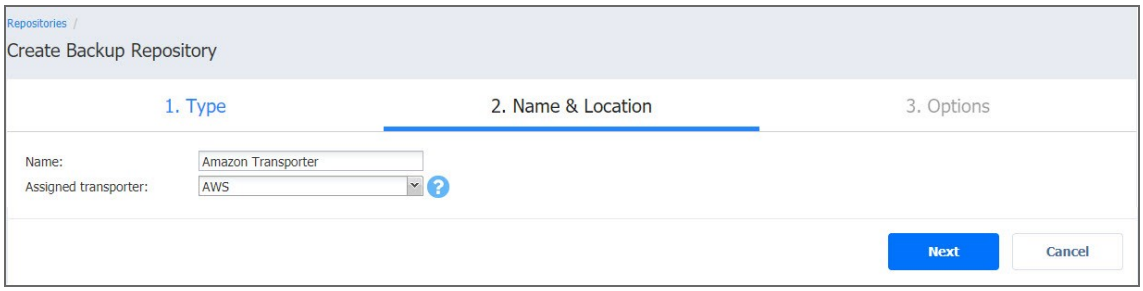

Create Backup Repository: Options

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)
  - 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.
- **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.

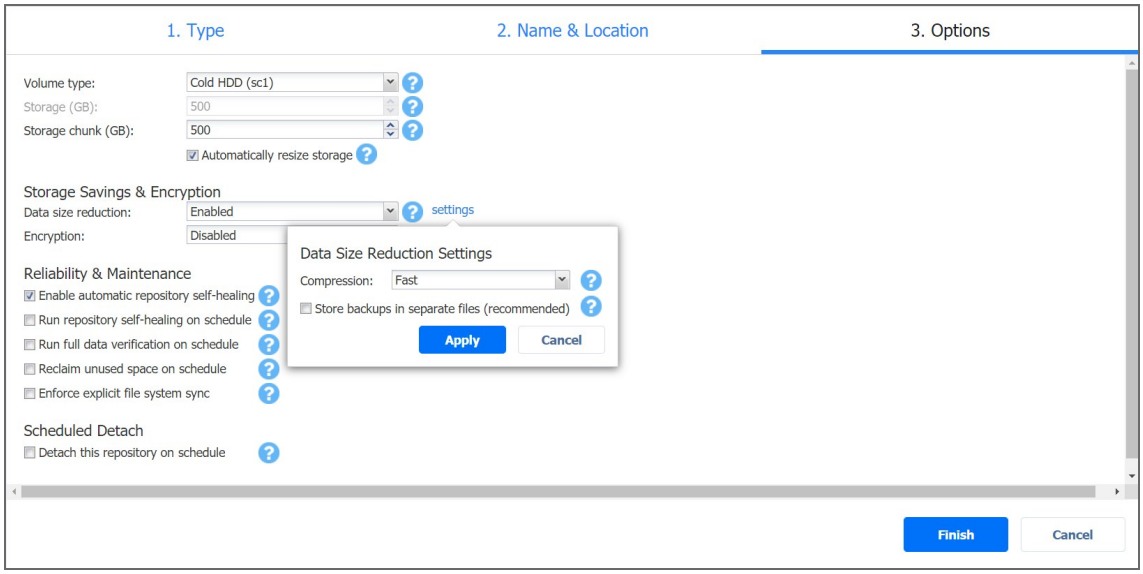

Backup Repository in Amazon S3

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.

Important

- You will be charged for Amazon S3 storage/traffic according to AWS tariffs.
- Forever incremental backups are not supported by this location.

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.

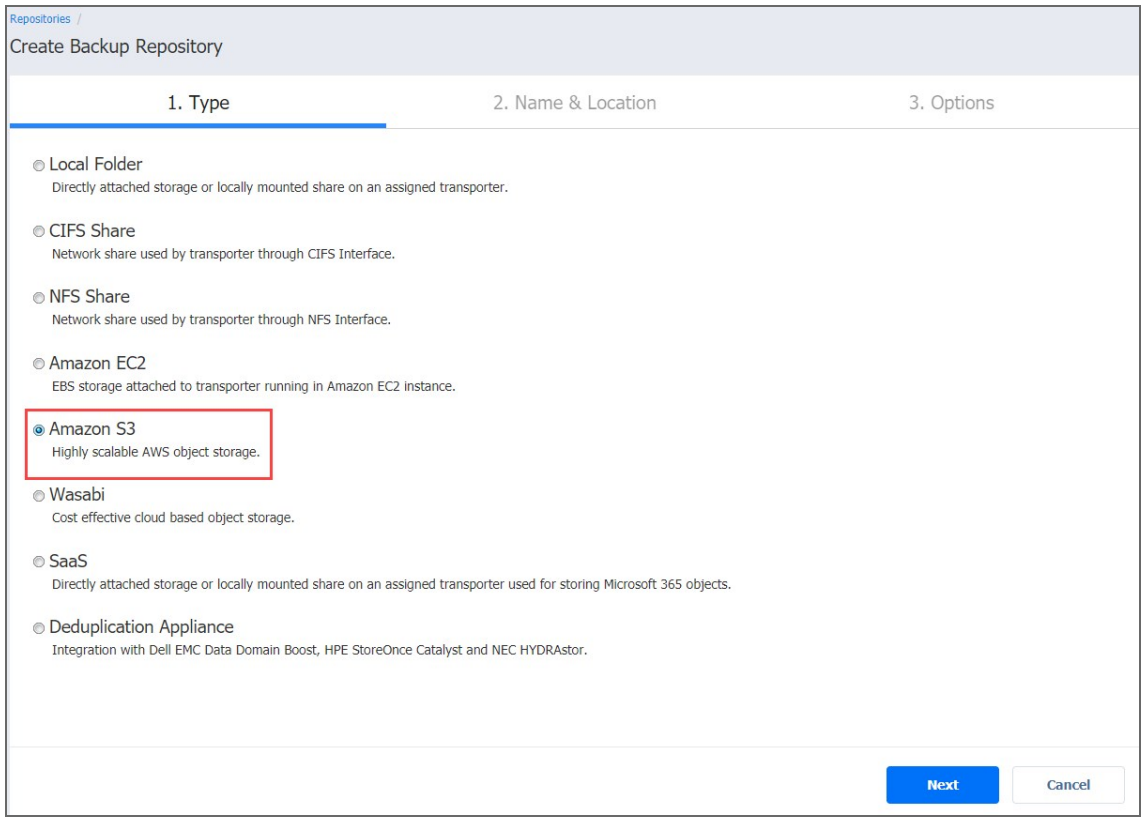

Create Backup Repository: Name and Location

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.

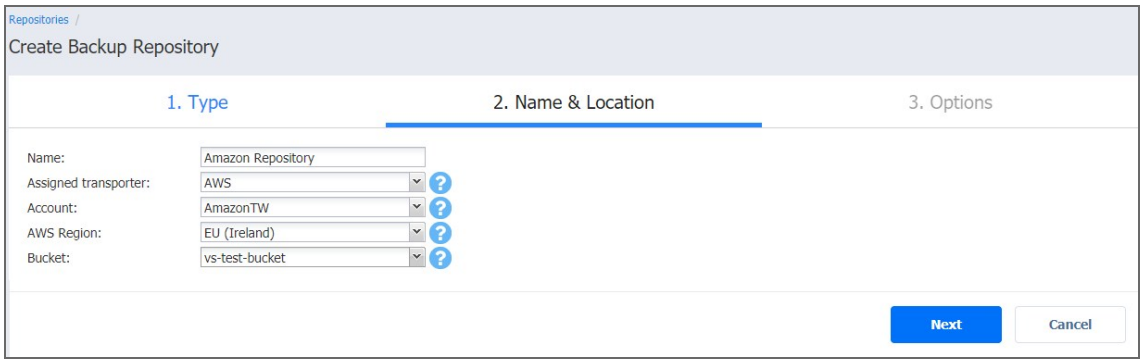

Create Backup Repository: Options

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.

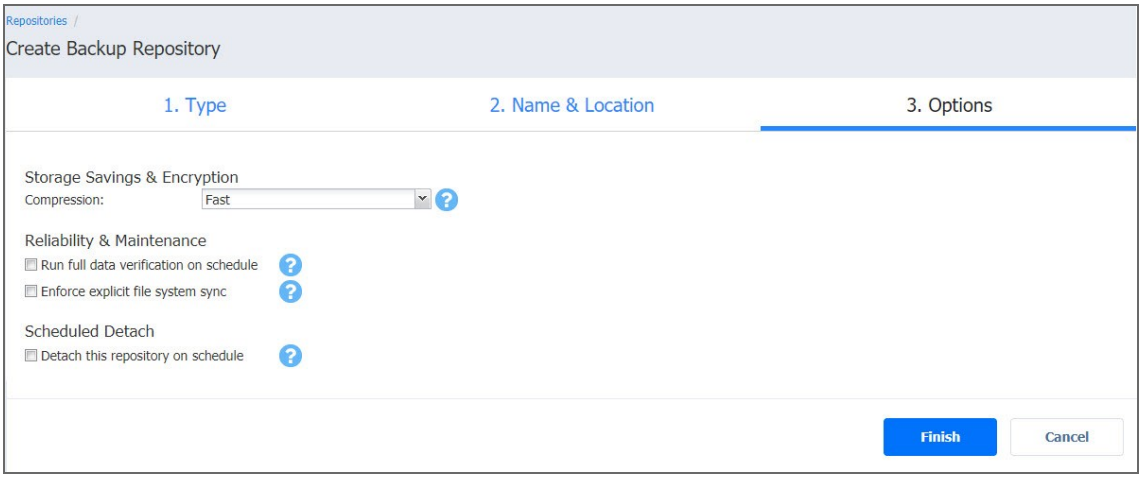
Repositories /
Create Backup Repository
1. Type
2. Name & Location
3. Options
Storage Savings & Encryption
Compression:
Fast
Reliability & Maintenance
Run full data verification on schedule
Enforce explicit file system sync
Scheduled Detach
Detach this repository on schedule
Finish
Cancel

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.

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.

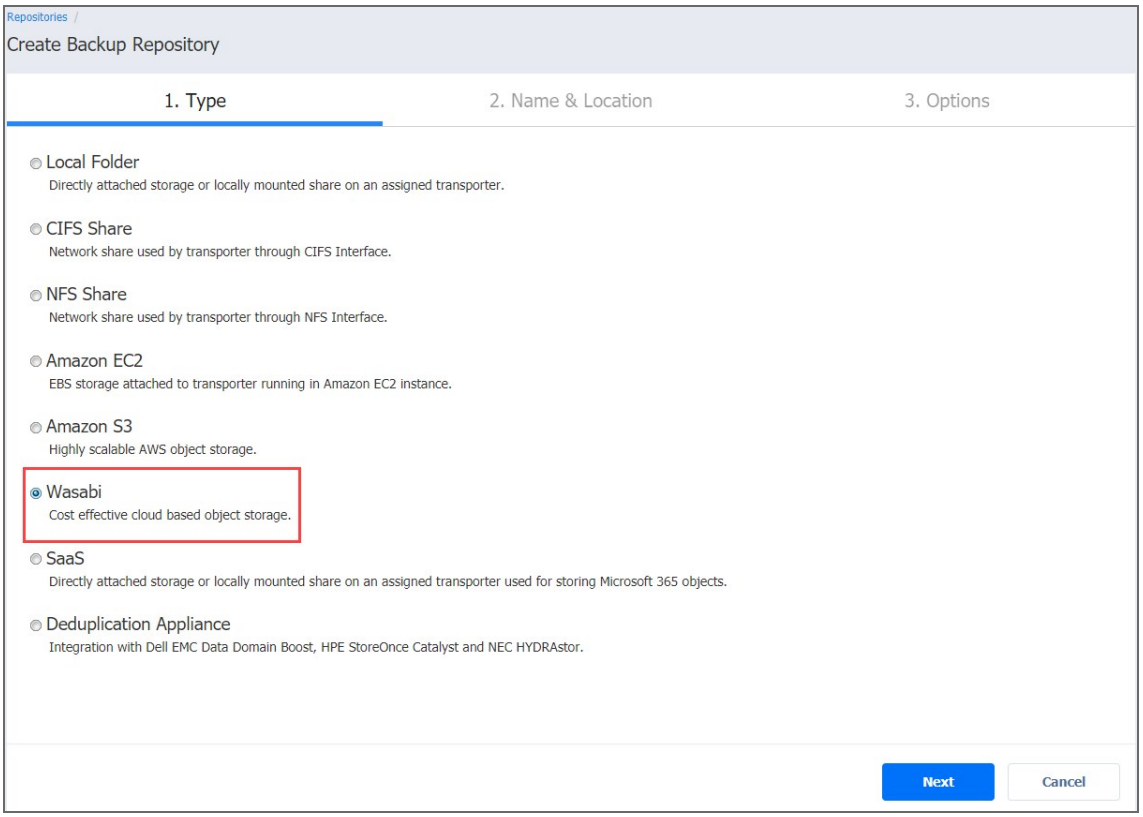

Create Backup Repository: Name and Location

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.

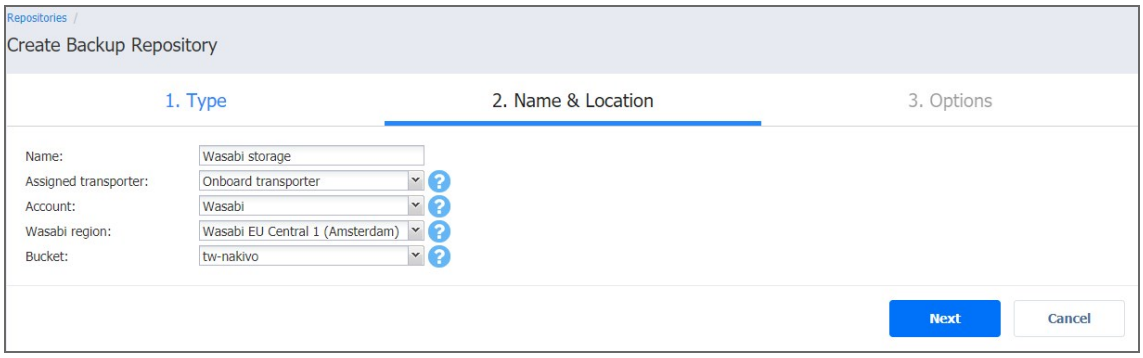

Create Backup Repository: Options

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.

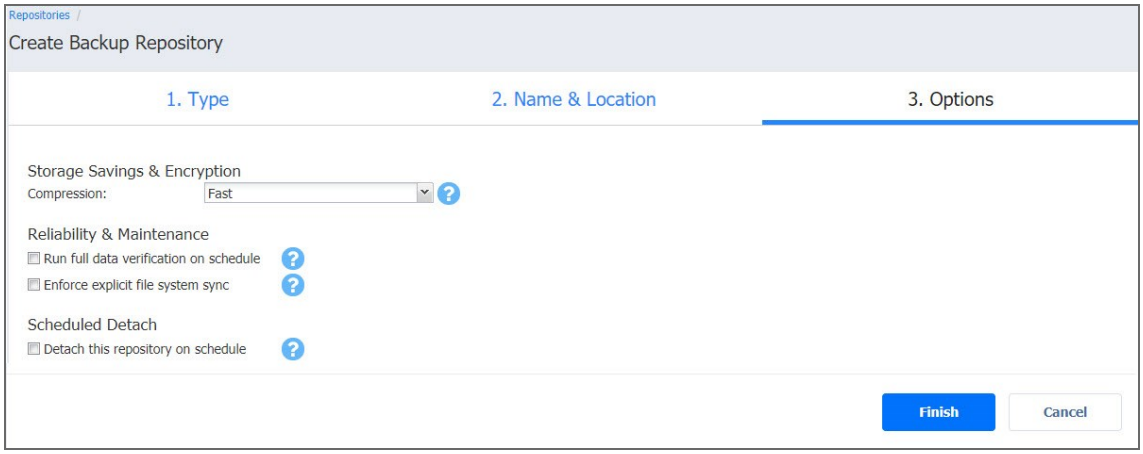

Backup Repository on Deduplication Appliance

NAKIVO Backup & Replication allows you to use advanced deduplication appliances for data protection.

Notes

- Before creating a Backup Repository on a Dell EMC DD, you need to install BoostFS Plugin and create a storage unit on the data domain backup appliance. Refer to Integrating with EMC DD Boost for details.
- Before creating a Backup Repository on an NEC HYDRAstor, you need to configure the NEC HYDRAstor and the machine on which NAKIVO Transporter is installed. Refer to Integrating with NEC HYDRAstor for details.
- To create a Backup Repository on other deduplication appliances, refer to "Backup Repository on NFS Share" on page 379.

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 96 to see the list of supported advanced deduplication appliances.

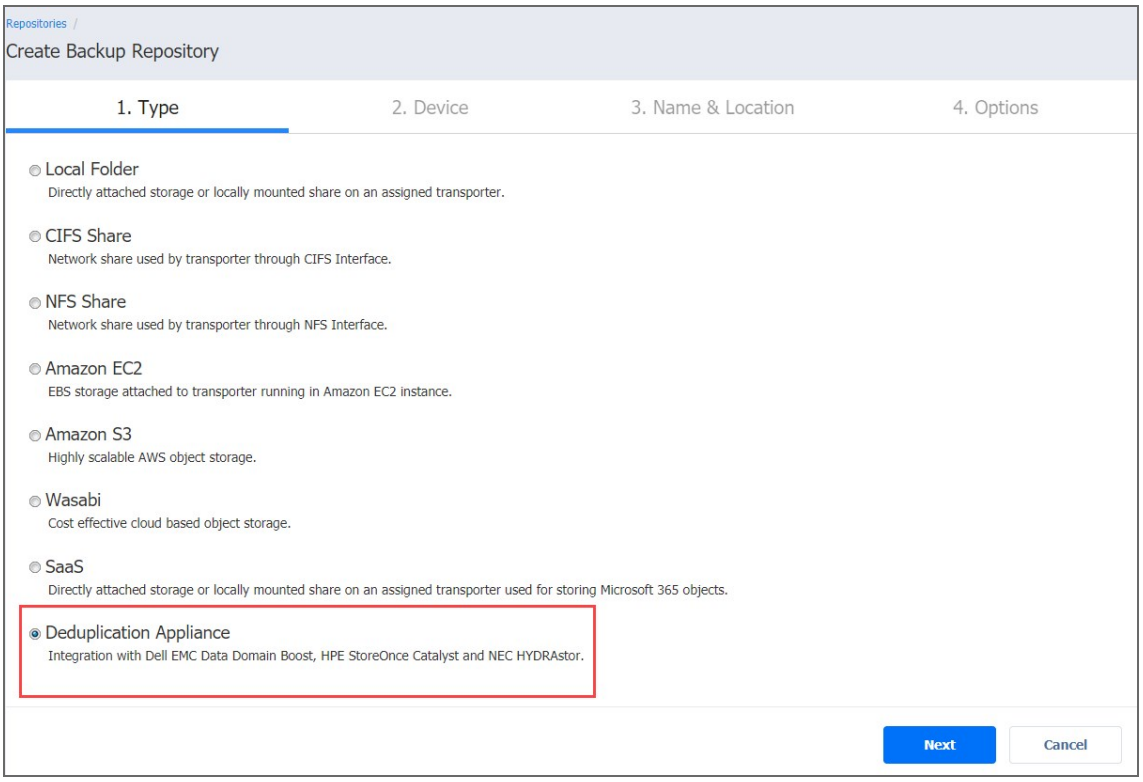

Create Backup Repository: Device

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

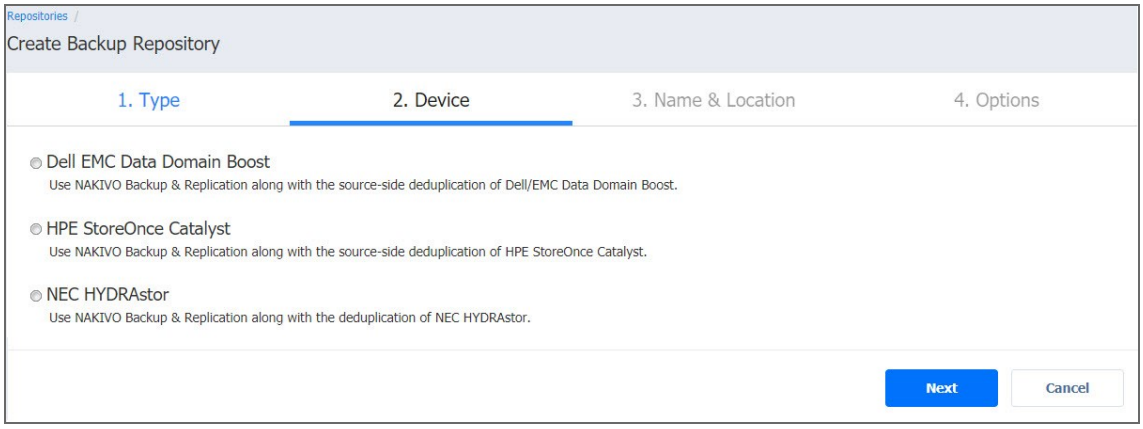

Create Backup Repository: Name and Location

On the **Name & Location** page, specify the following:

1. **Name**: Enter a name for the Backup Repository.
2. **Assigned transporter**: Choose a Transporter that will manage (that is, write data to and read data from) this Backup Repository.
3. Depending on the deduplication appliance, provide the following information:
   - **Dell EMC Data Domain Boost**
     1. **Name**: Enter the name of your Backup Repository.
     2. **Assigned transporter**: Select the assigned Transporter.
     3. **Path to the share**: Enter the path to the share folder in the following format: `<backup_appliance>:/<storage_unit>`. Refer to Creating a NAKIVO Backup & Replication Backup Repository on EMC Data Domain Backup Appliance for details.

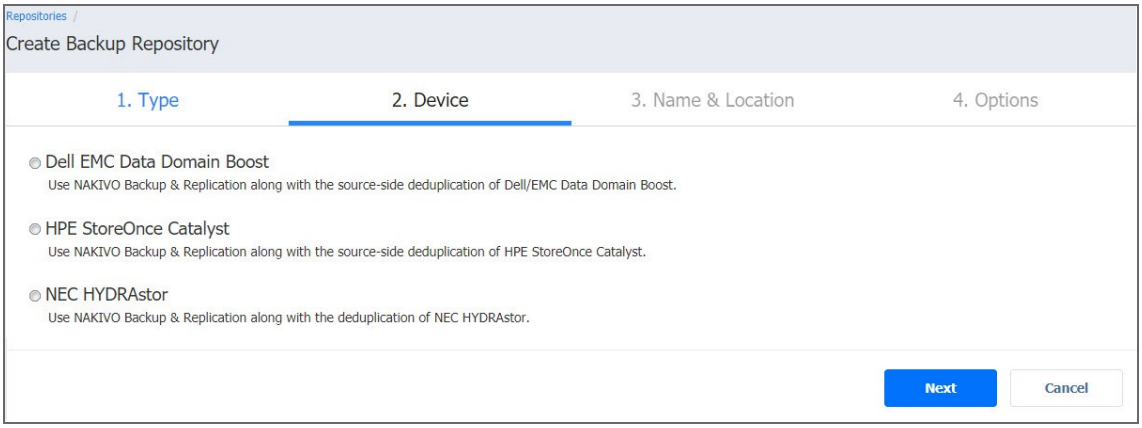

   - **HPE StoreOnce Catalyst**
     1. **Name**: Enter the name of your Backup Repository.
     2. **Assigned transporter**: Select the assigned Transporter.
     3. **Connection type**: Select one of the connection types to be used to access the Backup Repository:
        - **IP address**
        - **Fibre Channel**
     4. Depending on the connection type, do the following:
        - **Server name** (if IP address connection type is selected): Enter the server name or IP address of the HPE StoreOnce Catalyst.
        - **COFC identifier** (if Fibre Channel connection type is selected): Enter the COFC identifier. You can find your COFC identifier by going to **Catalyst Settings>Fibre Channel** in the **StoreOnce Management Console**.
     5. **Catalyst store name**: Enter the Catalyst store name.
     6. **Username**: Provide the username to the Catalyst store.

7. **Password**: Provide the password to the Catalyst store.

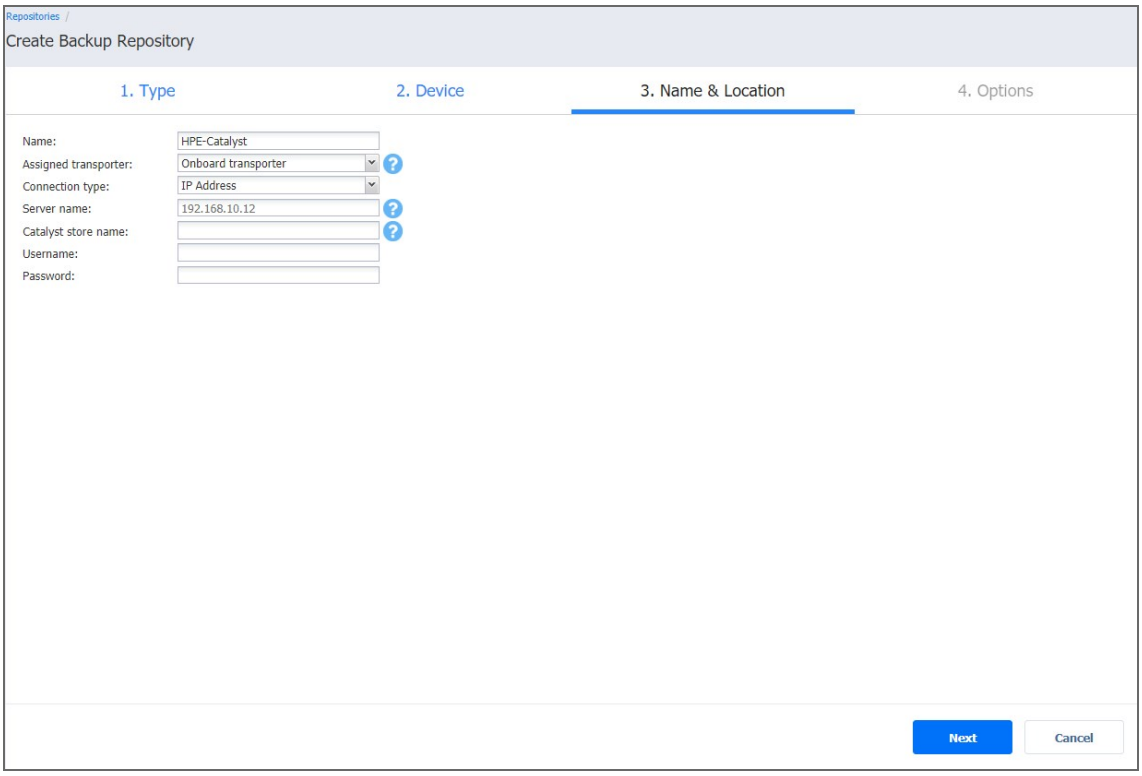

- **NEC HYDRAstor**
  1. **Name**: Enter the name of your Backup Repository.
  2. **Assigned transporter**: Select the assigned Transporter.
  3. **Path to the mount point**: Enter the path to the mount point in the following format: `/opt/nakivo/repository/hsva`.

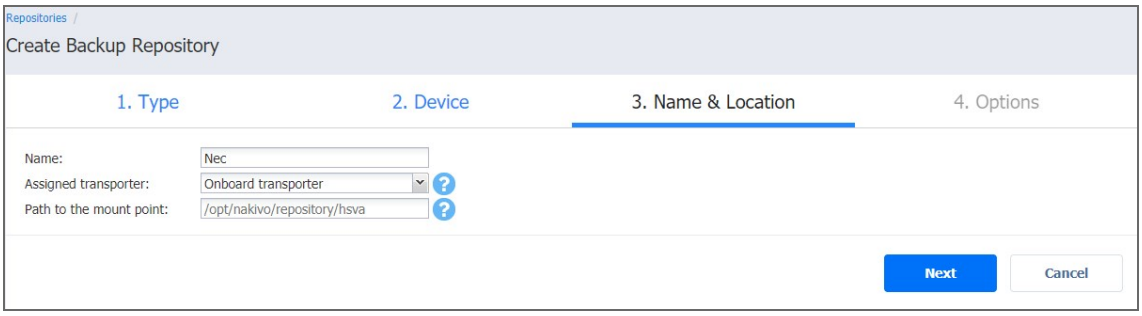

4. Click **Next** to go to the next page of the wizard.

Create Backup Repository: Options

On the **Options** page, do the following:

1. Set up *Storage Savings & Encryption* options:
    - **Data size reduction**: If this option is enabled, NAKIVO Backup & Replication enables the use of data size reduction for this repository to save disk space. Note that this may put additional load on the CPU. Disabling data size reduction is required if the target is a deduplication storage appliance. Click settings to configure the settings. A popup window appears. Set the following:
        - **Compression**: Select a compression level that will be used to reduce the data size in the Backup Repository. Note that higher compression levels consume considerably more CPU and may slow down VM backup speed. The following options are available:
            - **Disabled:** The data in the Backup Repository will not be compressed.
            - **Fast:** Lowest compression level.
            - **Medium:** Medium compression level.
            - **Best:** Maximum compression level.

            Note

            This option cannot be configured after creating the Backup Repository.

        - **Store backups in separate files**: Select this option to enable this backup repository to store data of every machine in separate backup files. Enabling this option is highly recommended to ensure higher reliability and performance.
    - **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.

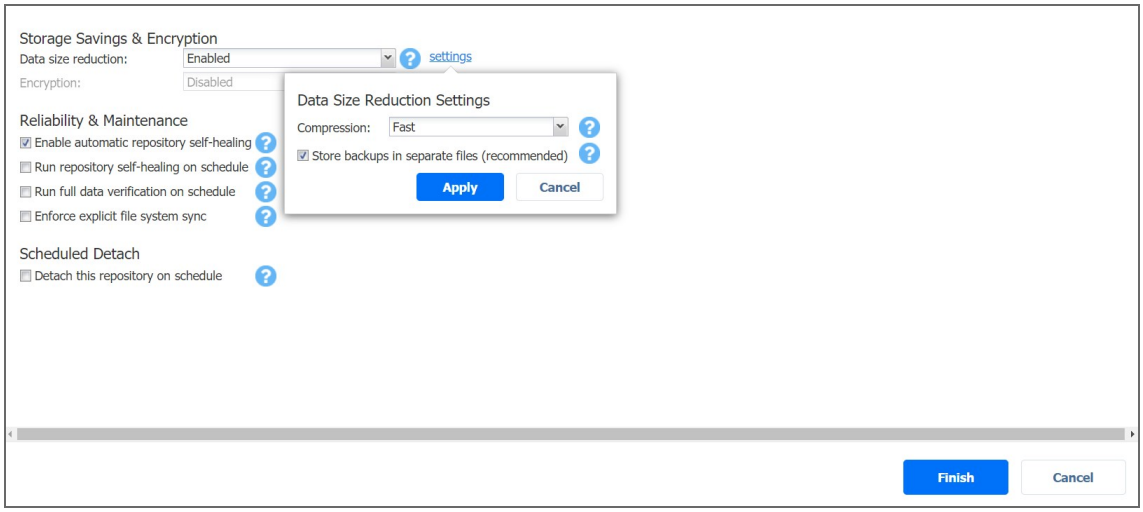

Viewing Backup Repository Details

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:

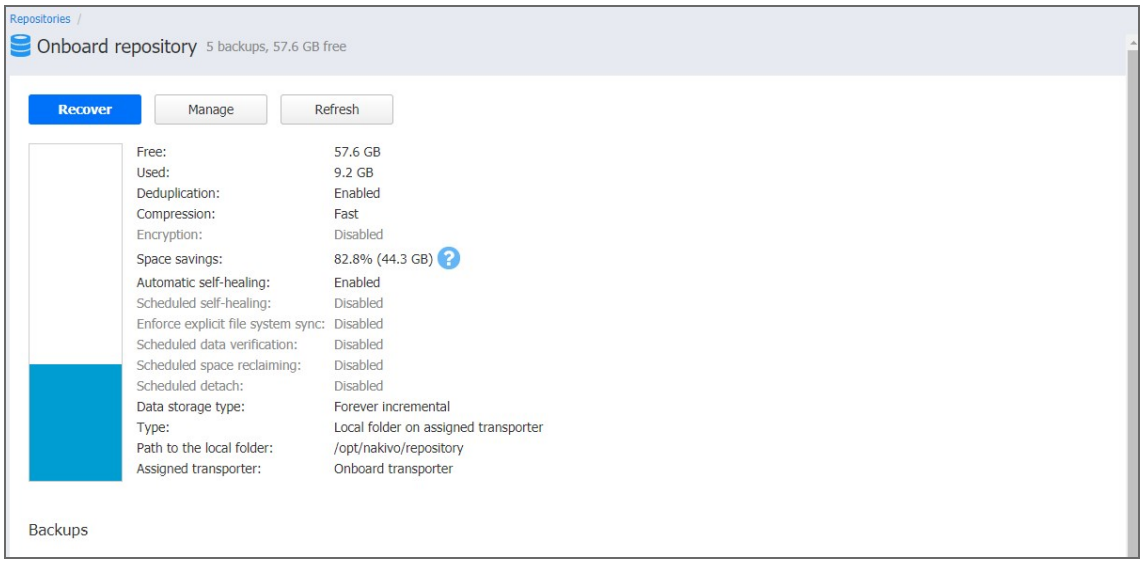

- **Free**: The amount of free space currently available for the Backup Repository.
- **Used**: The amount of space that the Backup Repository occupies on a disk. The amount of space that can be reclaimed is displayed in parentheses.
- **Deduplication**: The status of deduplication in the current Backup 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.

- **Data storage type**: The type of Backup Repository, which can be one of the following:
  - **Forever incremental**: After the initial full backup, all subsequent backups will transfer and store only changed data (increments) to the backup repository.
  - **Incremental with full backups**: After the initial full backup, backup jobs will store changed data (increments) in separate files and will periodically create full backups as specified in the job settings. Backup deduplication is not available when this option is selected.
- **Type**: The location of the Backup Repository, which can be one of the following:
  - **Local folder on the 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.

Hover over the name of the 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.
- **Delete**: Select this option to delete the backup from the repository.

Click on the 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.
- **Job name**: Name of the job.

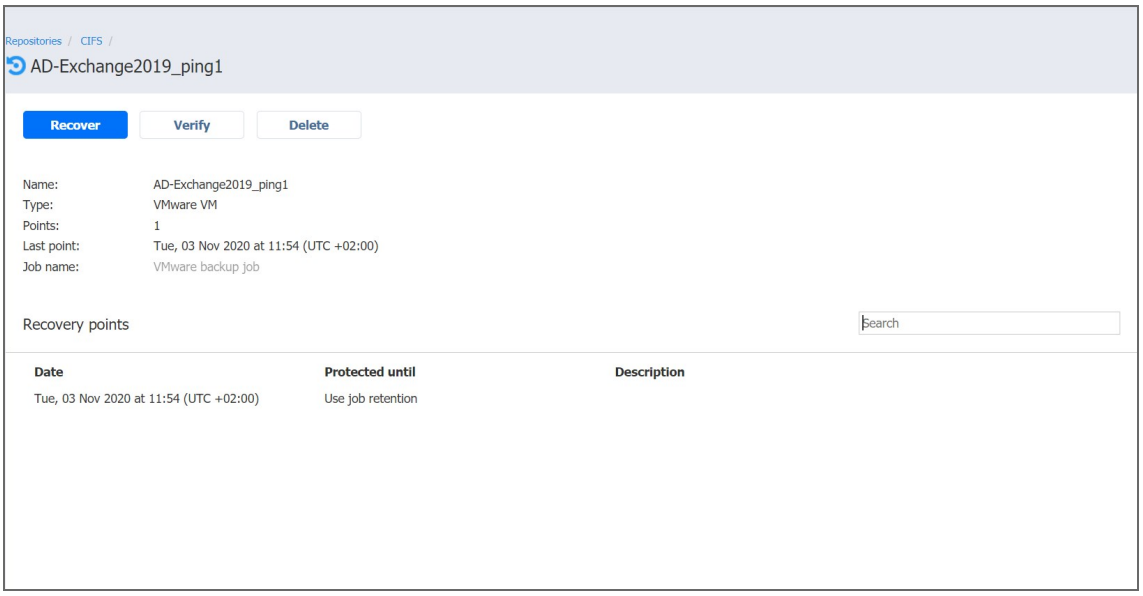

Viewing Recovery Point Details

You can view the details of the recovery point in the lower part of the screen. 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.
- **Protected until**: The date until which the recovery point remains protected.
- **Description**: Description of the recovery point.

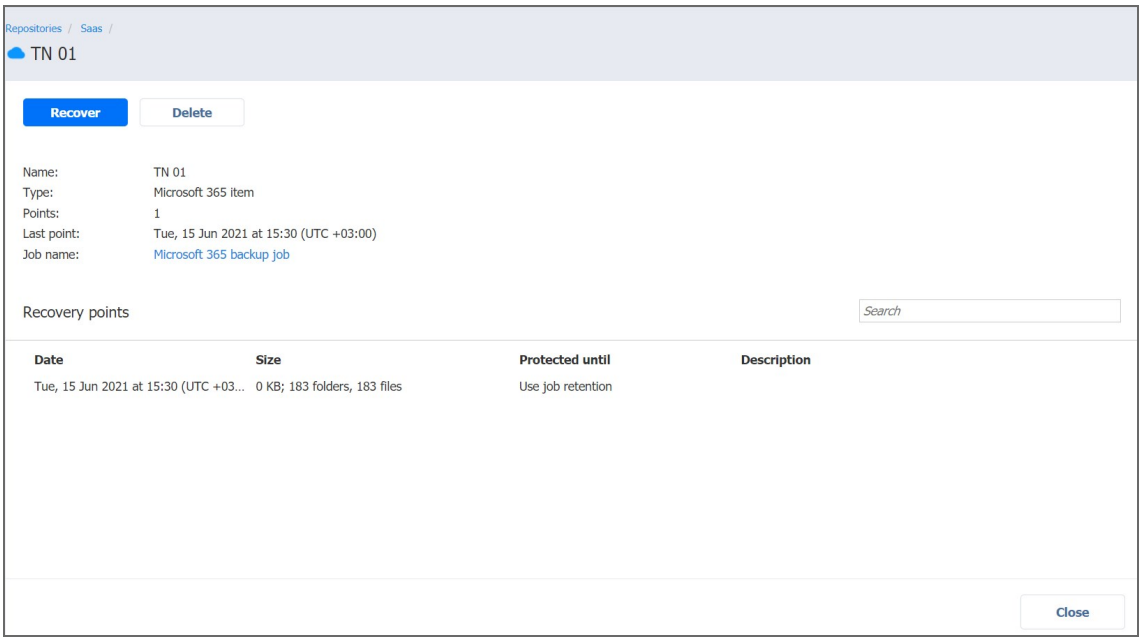

Note

Size and **Type** are displayed only if the Backup Repository has **Incremental with full backup data storage** type selected.

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 the **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.

Managing Backup Repositories

Refer to the following topics:

- "Attaching Backup Repositories" on page 413
- "Detaching Backup Repositories" on page 414
- "Editing Backup Repositories" on page 415
- "How to Copy Backup Repository to Tape" on page 416
- "Reclaiming Backup Repository Space" on page 417
- "Refreshing Backup Repositories" on page 419
- "Removing and Deleting Backup Repositories" on page 421
- "Repairing Backup Repository" on page 423
- "Running Backup Repository Self-Healing" on page 426
- "Running Block-Level Backup Verification" on page 428

Attaching Backup Repositories

If you have detached a Backup Repository, you can reattach it to the product 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 click a Backup Repository.
3. In the Backup Repository title, click **Manage** and then click **Attach.**

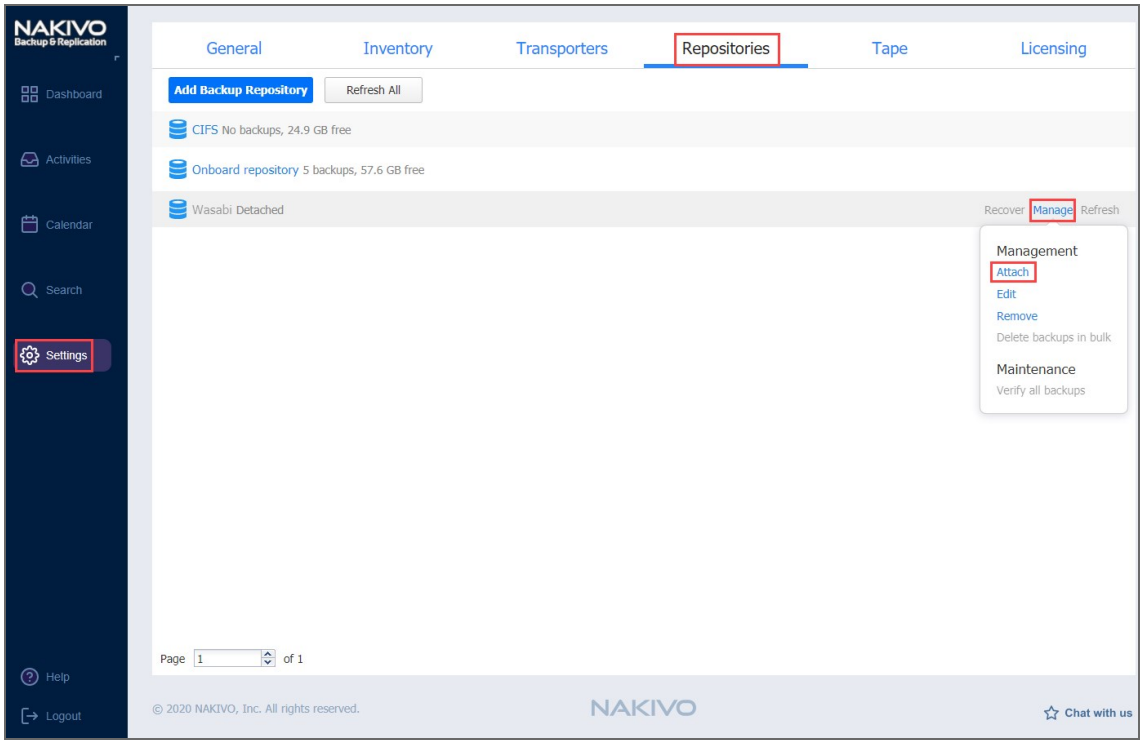

The Backup Repository is reattached to NAKIVO Backup & Replication. You can now back up to the attached Backup Repository.

Detaching Backup Repositories

Detaching a Backup Repository saves the Backup Repository data and metadata in a consistent state and then stops the product's interaction with the Backup Repository (such as read and write of data and metadata, and so on). You may want to detach a Backup Repository in order to move it to a different location or to put the storage with the Backup Repository on maintenance.

Important

Since the product stops working with detached backup repositories, jobs that back up VMs to a detached Backup Repository will fail.

To detach a Backup Repository, follow the steps below:

1. Go to the main menu and click **Settings**.
2. Go to the **Repositories** tab and choose a Backup Repository.
3. In the Backup Repository title, click **Manage** and then click **Detach.**

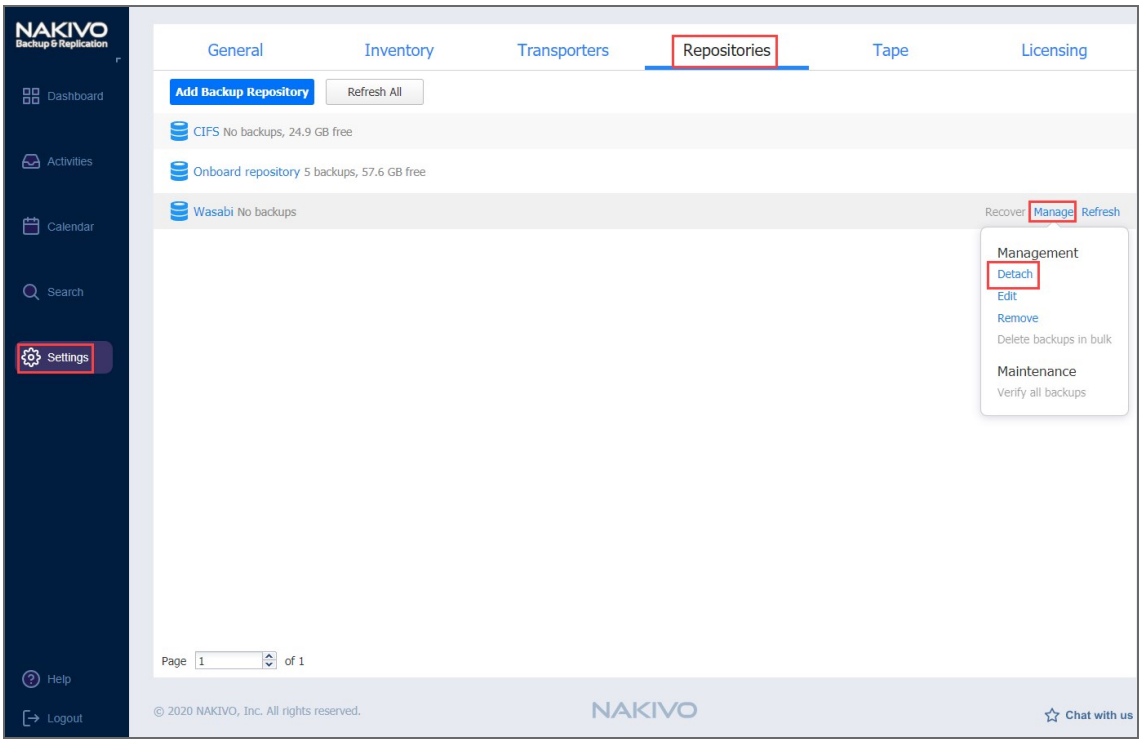

Note

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. 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 Backup Repository title, click **Manage** and then click **Edit**.

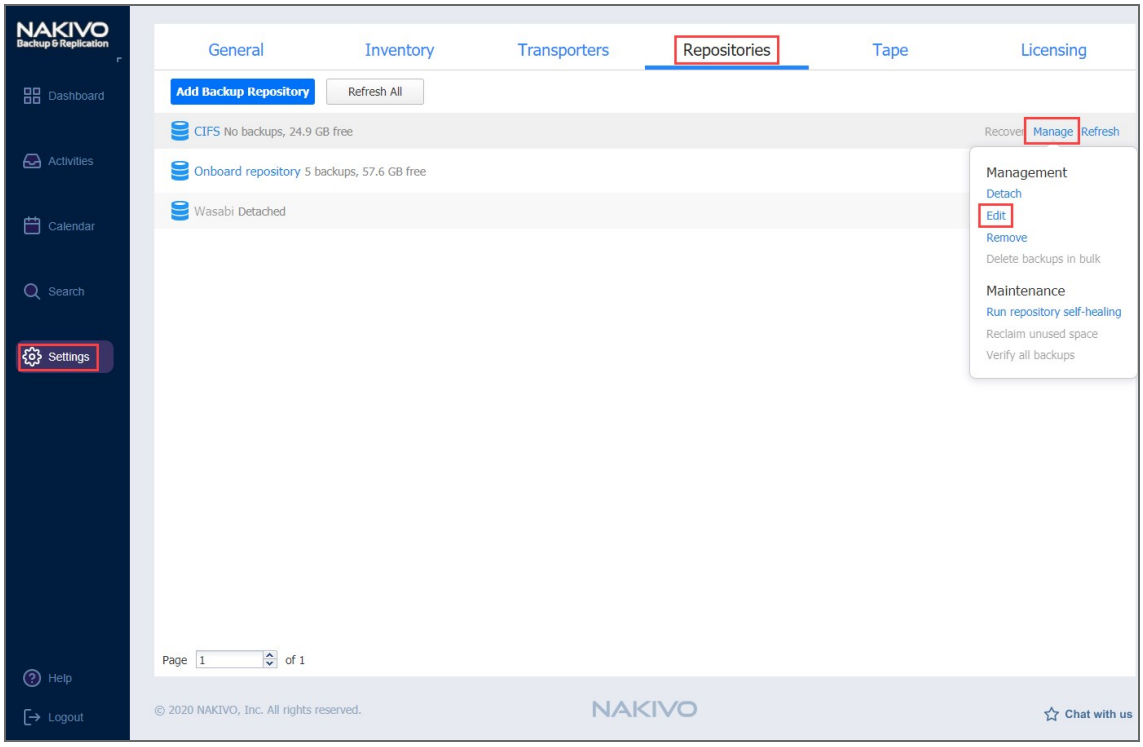

Note

A Backup Repository cannot be edited if there is a job that backs up to this Backup Repository is concurrently 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

   The Onboard Backup Repository is forever incremental by default. If you want to store backups that include incremental and full backup files, you should create a new backup repository and configure it as incremental with full backups. This can be done 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” below
- “Stopping the Space Reclaiming Process” on the next page

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 information about Backup Repositories every hour. During the refreshing process, the product collects all required information about Backup Repositories (such as the amount of free space, number of backups and recovery points, and so on).

Only one Backup Repository is refreshed at a time. Therefore, if you have more than one Backup Repository, all others will remain in 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. Go to the main menu of NAKIVO Backup & Replication and click **Settings**.
2. Go to the **Repositories** tab.
3. Click **Refresh All.**

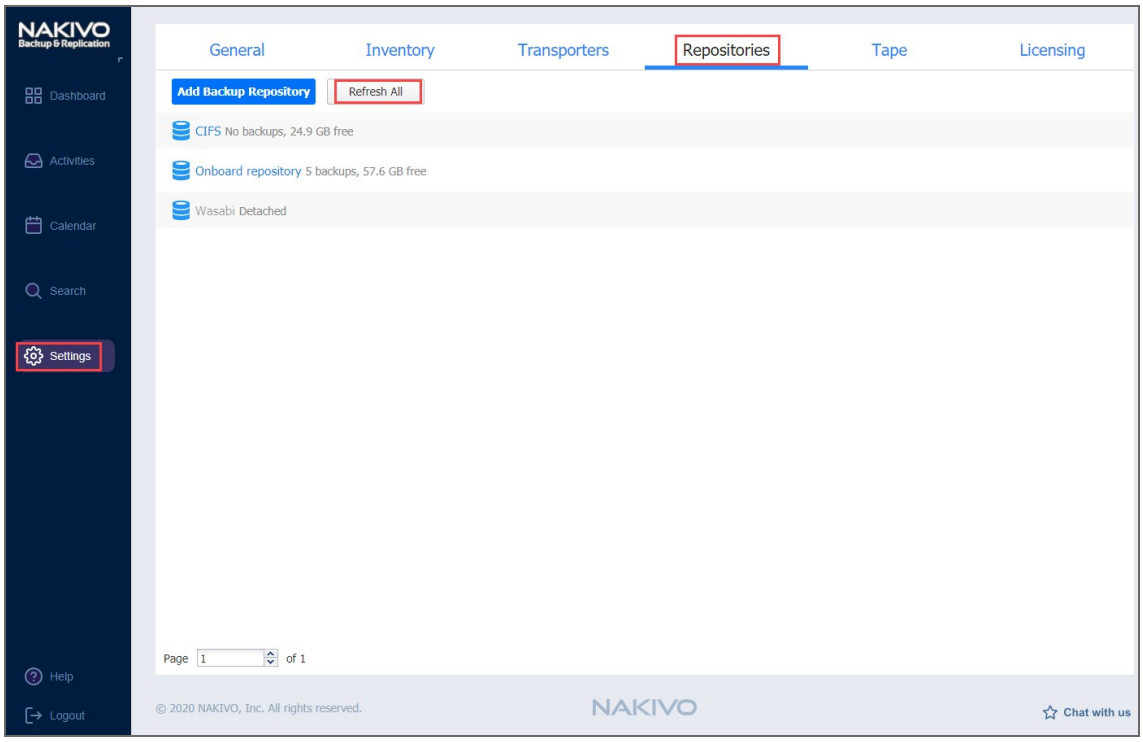

The Backup Repositories refresh starts.

Refreshing a Single Backup Repository

To update a single Backup Repository, follow the steps below:

1. Go to the main menu of NAKIVO Backup & Replication and click **Settings**.
2. Go to the **Repositories** tab.
3. Click the Backup Repository that you wish to update.
4. In the title of the Backup Repository, click **Refresh.**

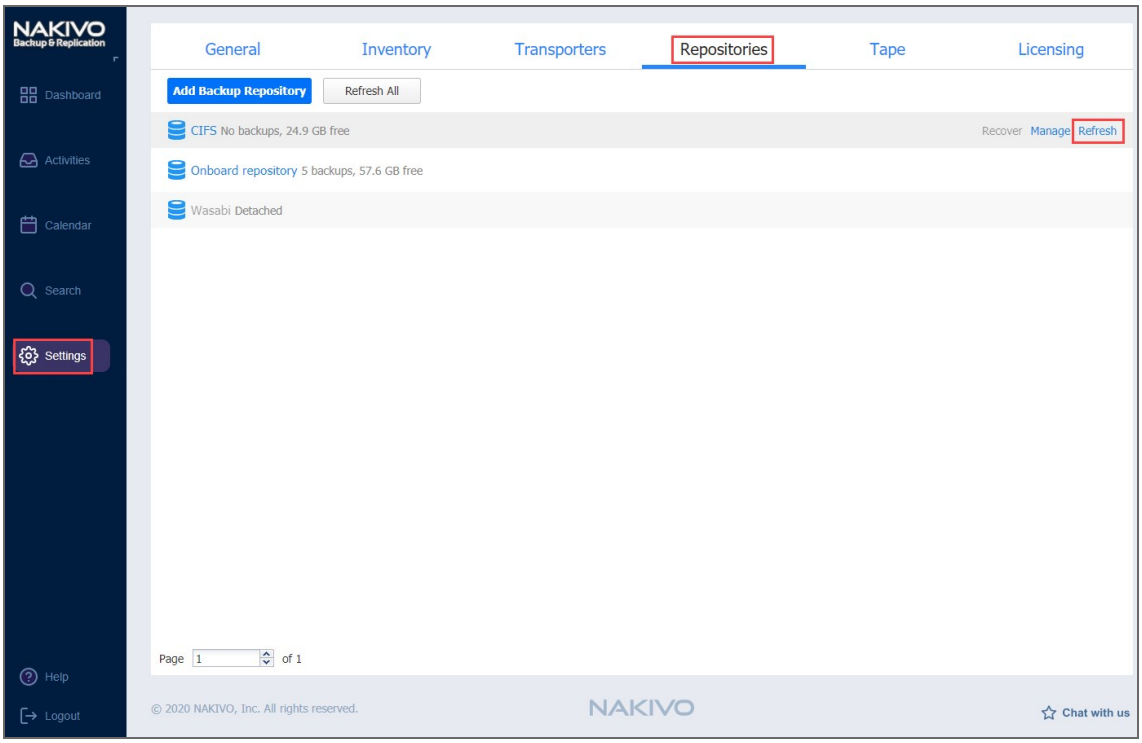

The Backup Repository refresh starts.

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.

Important

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, you first need to delete (or edit) the corresponding jobs so no items are backed up to the Backup Repository that is being removed.

To permanently delete or remove a Backup Repository from the product, 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. In the Backup Repository title, click **Manage** and then click **Remove.**

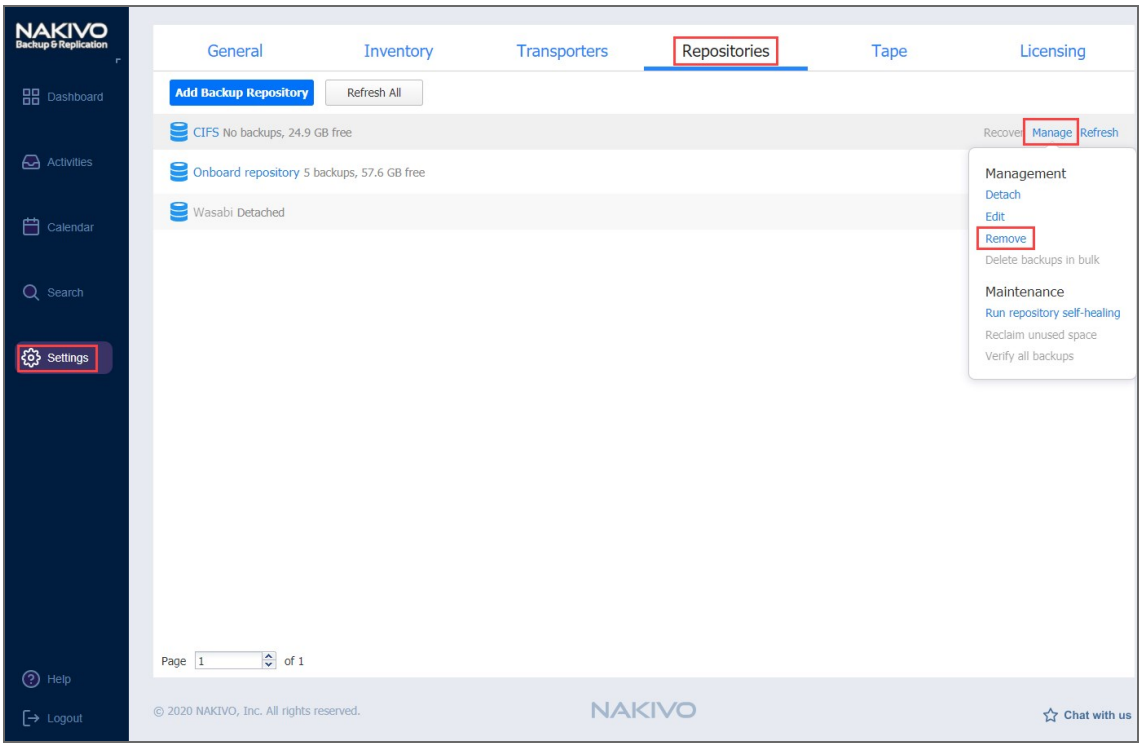

5. Do the following when the confirmation message appears:
   - To remove the Backup Repository from NAKIVO Backup & Replication and keep the Backup Repository on a disk, click the **Remove Repository and Keep Backups** button.

     Note

     You can import the removed Backup Repository back to the same or to new product installation.

- To permanently delete the Backup Repository and all its data, click the **Remove Repository and Delete Backups** button.

  Important

  This operation will permanently delete the Backup Repository and all VM backups.

Repairing Backup Repository

In case the backup with immutability enabled or the Backup Repository itself is corrupted, it is possible to initiate the repair process. During this process, NAKIVO Backup & Replication tries to revert the Backup Repository or a specific backup to its uncorrupted state.

Refer to the following topics:

- "Running Repair Process for the Backup Repository" below
- "Running Repair Process for the Specific Backup Object" on the next page

Running Repair Process for the Backup Repository

To run repair for the Backup Repository in **Local Folder** or **Amazon S3** *Type* of 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:

   - 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.

3. Check the desired options. Choose among the following:
   - **Overwrite repository metadata:** If checked, the metadata file is overwritten even if it's present and valid. If the metadata file is not present, the new file is then created regardless of whether this option is checked or not.
   - **Overwrite backup objects:** If checked, the locked backup objects are overwritten with the immutable data during the repair process.

- **Verify backup objects:** If checked, NAKIVO Backup & Replication runs verification on the backup object after the repair process has completed. In case **Verify backup objects** is not selected, NAKIVO Backup & Replication runs automatic self-healing after the repair process is finished.

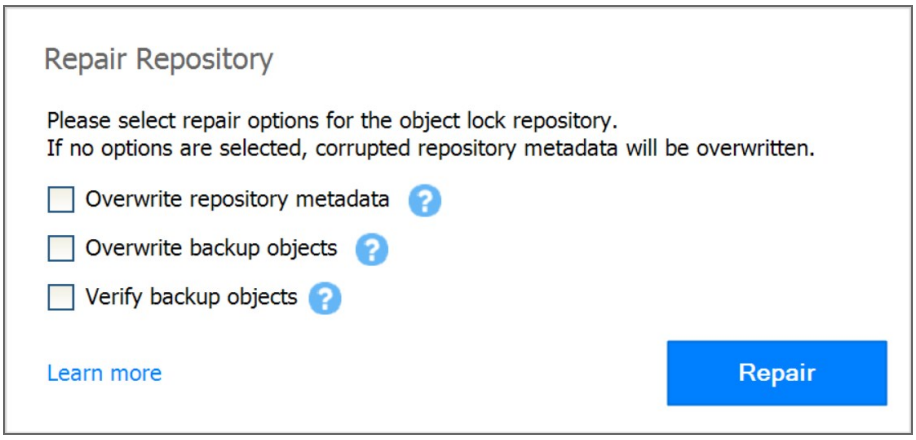

4. Click **Repair** to begin the repair process.

Running Repair Process for the Specific Backup Object

To run repair for a specific backup object located in a **Local Folder** or **Amazon S3** *Type* Backup Repository, do the following:

1. Go to **Settings > Repositories** and can click on the name of the Backup Repository. Hover over the name of the backup and click **Repair** to start the repair process. Alternatively, you can click on the name of the backup and then click **Repair**.

   Note

   You can also perform the **Repair** process for a backup object when all files except immutable files were manually deleted from the Backup Repository. The **Repair** option is only available in the following cases:

   - The Backup Repository is inaccessible, was created in Amazon S3, and has Object Lock enabled.
   - The local Backup Repository is inaccessible and meets the conditions specified in the feature requirements section.

2. Optionally, check the **Verify backup object** option. If you select this option, NAKIVO Backup & Replication runs verification on the backup object after the repair process has completed. In case **Verify backup object** is not selected, NAKIVO Backup & Replication runs automatic self-healing after the repair process is finished.



3. Click **Repair** to begin the repair process.

Running Backup Repository Self-Healing

The self-healing process verifies Backup Repository integrity and automatically repairs errors wherever possible. Namely, the process performs the following tasks:

- Verifies that the data blocks of each recovery point are present in the Backup Repository.
- Cleans up “in progress” blocks of data from failed/crashed backup job runs that did not have a proper cleanup.
- Verifies and repairs Backup Repository metadata so that it correctly describes available data.
- Restores the consistent state of the Backup Repository to enable subsequent backup jobs.

Before the self-healing process begins, the Backup Repository is detached from the product to keep data in a consistent state. Jobs that back up VMs to such Backup Repository will fail while the self-healing process is in progress.

Refer to the following topics to learn more:

- “Starting the Self-Healing Process” below
- “Stopping the Self-Healing Process” on the next page

Starting the Self-Healing Process

To run the Backup Repository self-healing, follow the steps below:

1. 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 **Run repository self-healing.**

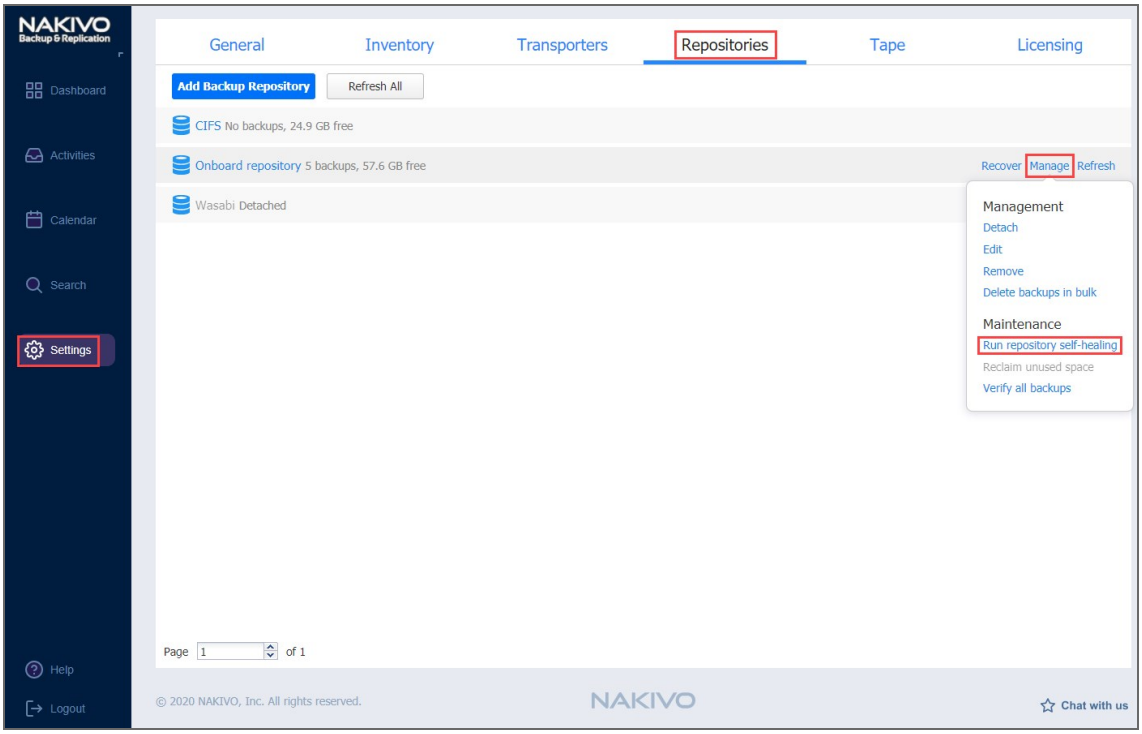

4. In the dialog box that appears, click **Start**.

   The self-healing process is started.

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. 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 repository self-healing.**

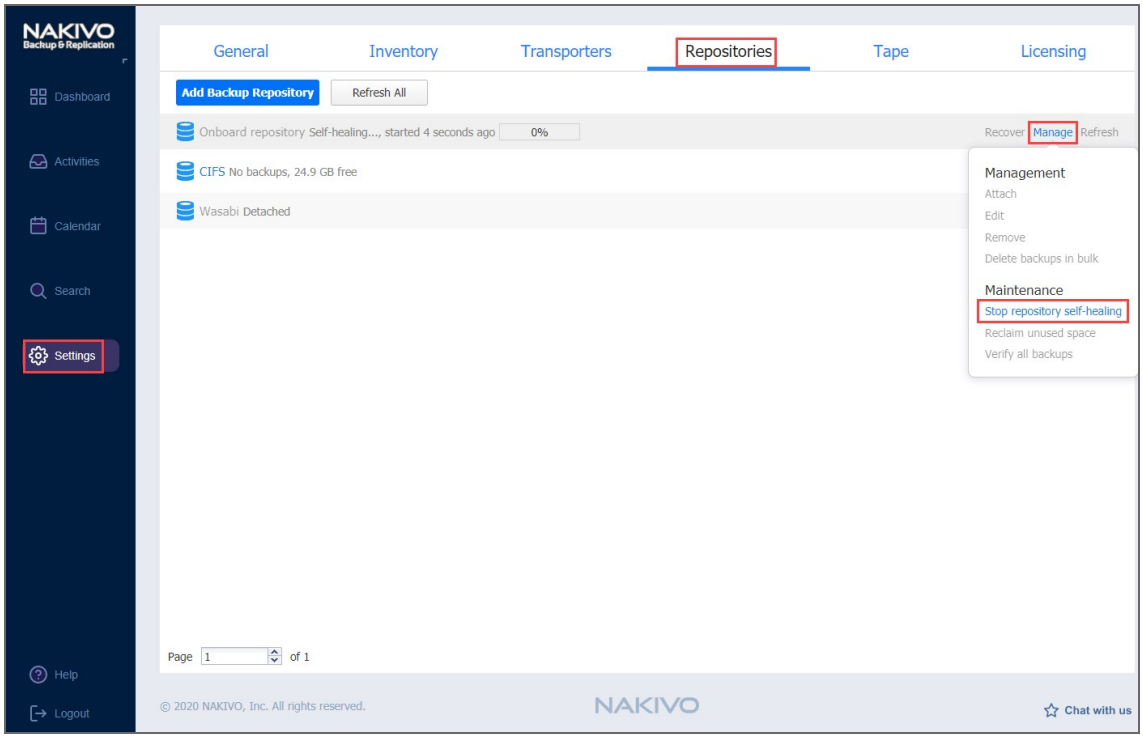

The self-healing process is stopped.

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" below
  - "Verifying a Single Backup" on the next page
- "Stopping the Backup Verification Process" on page 430
  - "Stopping Backup Verification for a Backup Repository" on page 430
  - "Stopping Backup Verification for a Single Backup" on page 430

Verifying Backups

Important

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. 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 **Verify all backups.**
   The backup verification process cannot be started if a job that backs up to this Backup Repository is running.
4. In the dialog box that opens, click **Start**. The backup verification process is started.

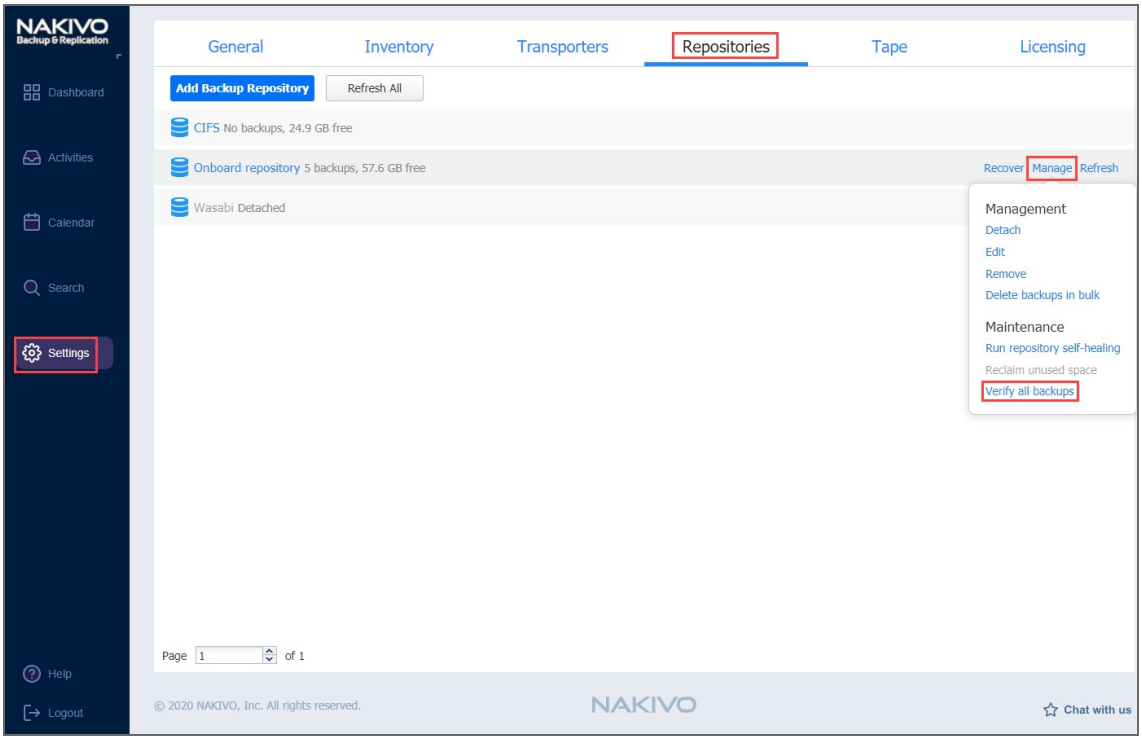

Verifying a Single Backup

To verify a single VM backup in a repository, follow the steps below:

1. In the main menu of NAKIVO Backup & Replication, click **Settings**.
2. Go to the **Repositories** tab and click a Backup Repository to expand it.
3. In the title of a VM backup, click **Verify**.

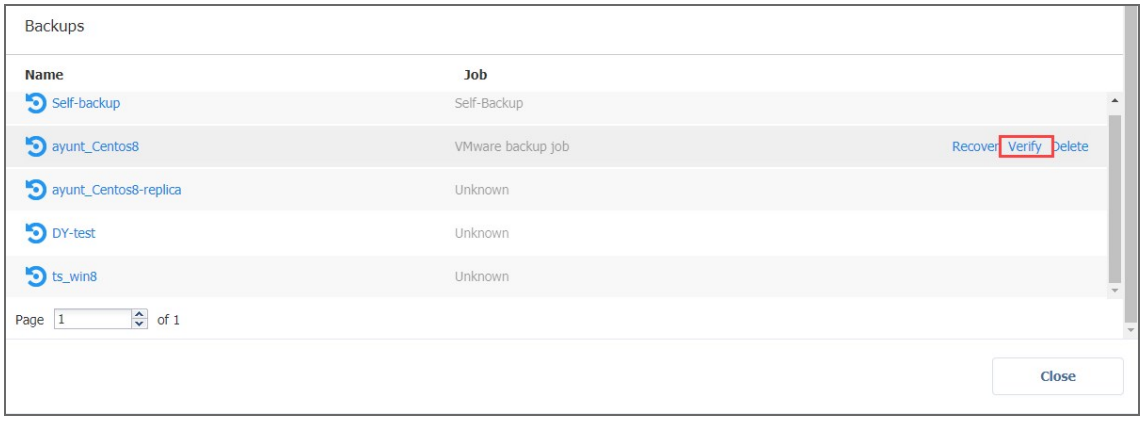

Stopping the Backup Verification Process

You can stop the backup verification process at any time (for example, to run a recovery job, move your Backup Repository to a new location, or put your backup storage on maintenance).

Stopping Backup Verification for a Backup Repository

To stop the backup verification process, 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 backup verification.**

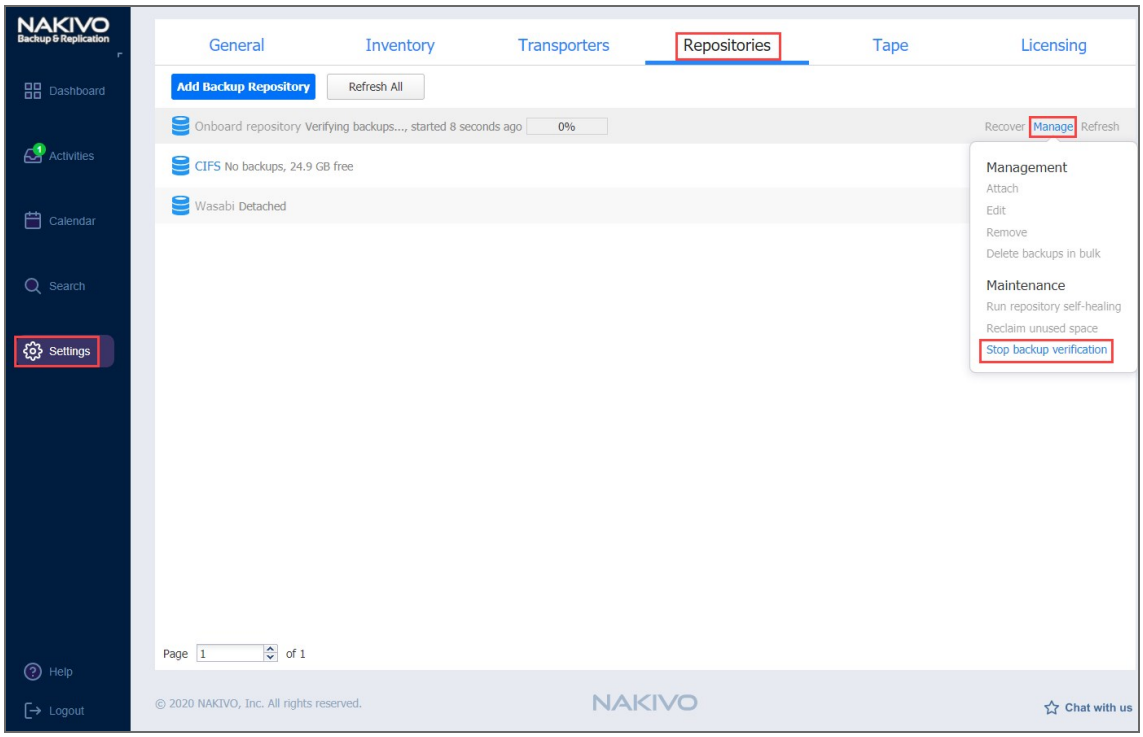

The backup verification process is stopped.

Stopping Backup Verification for a Single Backup

To stop the backup verification 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 click a Backup Repository to expand it.

3. In the title of a VM backup, click **Stop verifying**.

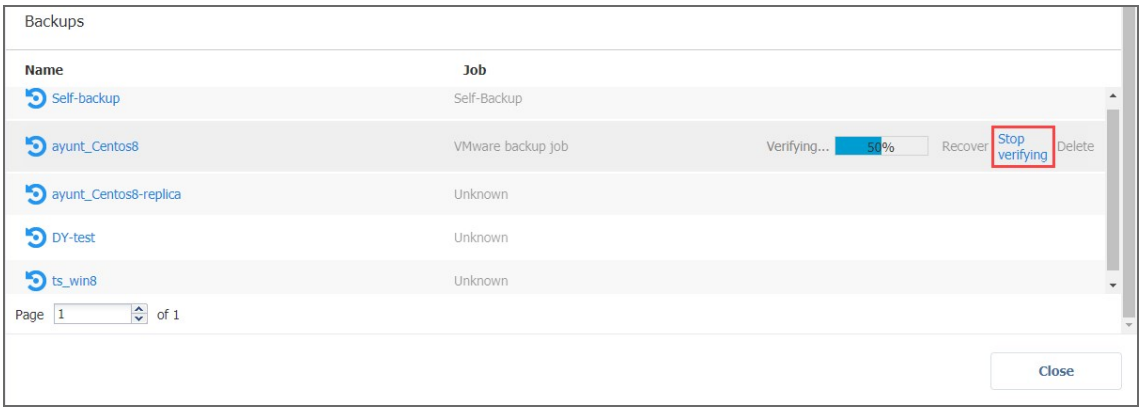

For near-instant backup verification, refer to the "VM Verification" on page 47 feature.

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.

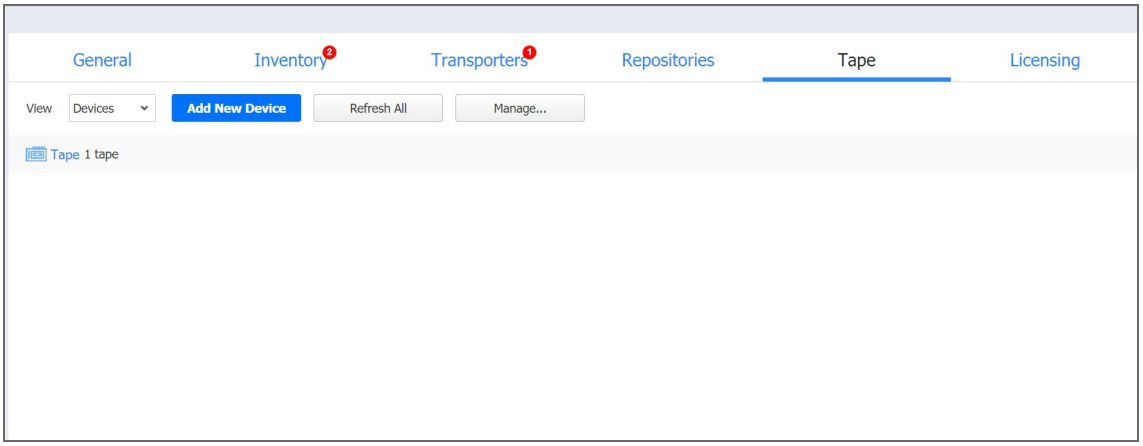

The default view of the **Tape** page is set to **Devices**, and once you add your tape devices, you will be able to view and manage them here. Also, by selecting different views from the **View** drop-down list you will be able to work with tapes and backups.

On the **Tape** page, you can perform the following operations:

- “Adding Robotic Tape Libraries or VTLs” on page 433
- “Adding Standalone Tape Drives” on page 438
- “Managing Backups” on page 441
- “Managing Locations” on page 444
- “Managing Media Pools” on page 446
- “Managing Tape Devices” on page 458
- “Managing Tape Cartridges” on page 448

Adding Robotic Tape Libraries or VTLs

Important

Please observe the following prerequisites before adding Robotic Tape Libraries or VTLs 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 VTL to NAKIVO Backup & Replication includes the following steps:

- Launching Wizard
- Selecting Transporter
- Selecting Changers
- Selecting Drives
- Selecting Options
- Managing Added Tape Library

Launching Wizard

To add a Robotic tape library or VTL to the system:

Note

Before adding a Robotic tape library or VTL, make sure that the on-premises VM or Amazon EC2 instance meets the necessary feature requirements.

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**.

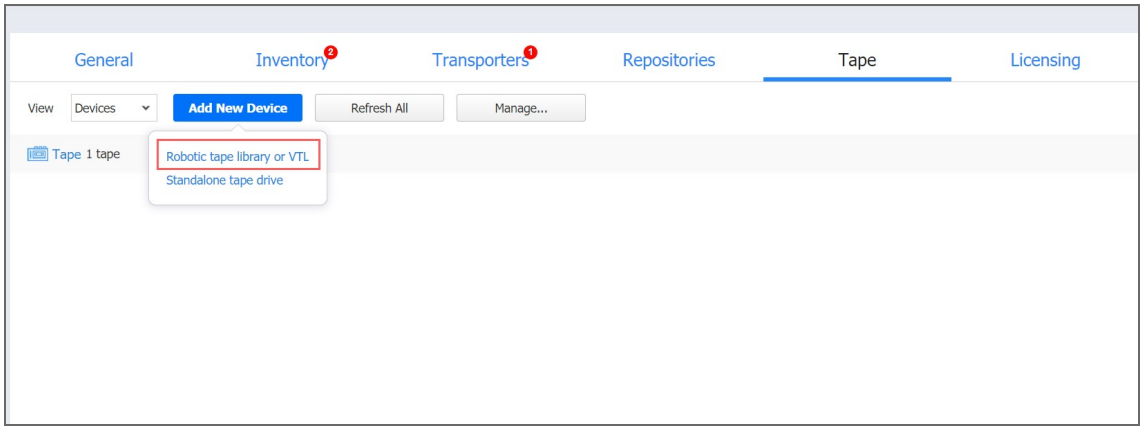

The **Add New Robotic Tape Library or Virtual Tape Library** wizard opens. Follow the steps below to add a new device.

Selecting Transporter

Important

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.

Refer to "Deploying Transporter as VMware Appliance" on page 345 and "Adding Installed Transporters" on page 353 for more information on Transporter deployment.

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.

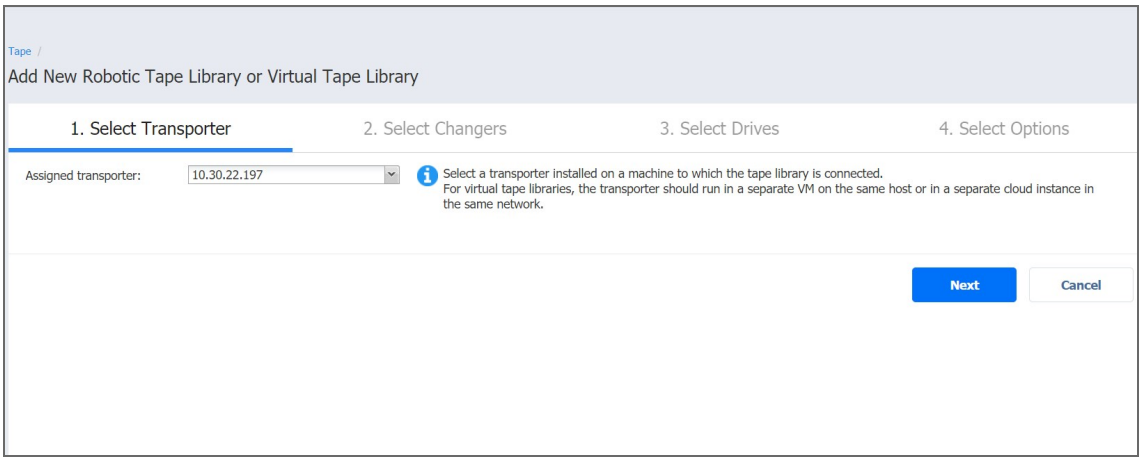

2. Click **Next**.

Selecting Changers

The **Select Changers** page displays the list of media changers on the selected Transporter.

Note

If no media changers were found on the specified transporter, make sure the devices are connected, powered on, and the appropriate drivers are installed.

Select one media charger from the list. Media changers already being used in another discovered tape library are disabled.

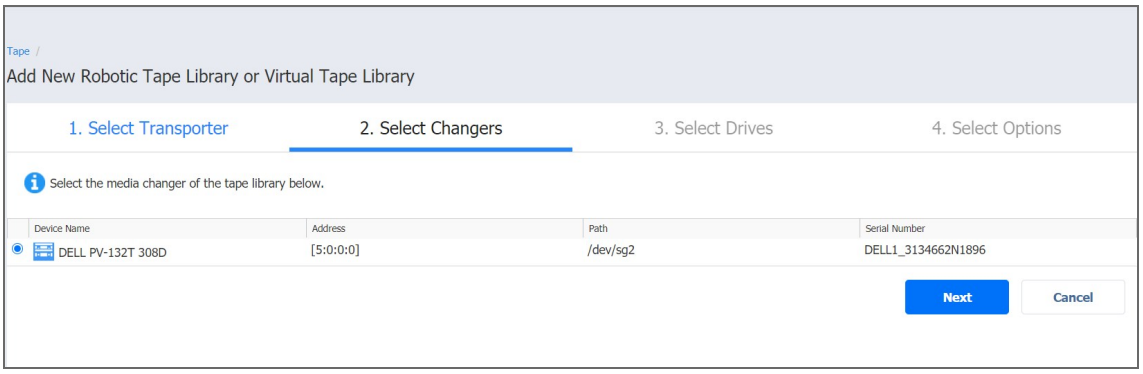

The following information is displayed for each media changer to facilitate the selection:

- **Device name**: Indicates device's vendor and model, separated by space
- **Address**: Indicates the hardware address including the bus and node numbers
- **Path**: Indicates location in the operating system
- **Serial number**: Indicates the serial number of the device

Selecting Drives

On the **Select Drives page**, you can select tape drives from the tape library and specify the actual drive number for each drive. Drives already used in another discovered tape device are disabled and cannot be selected. The table provides the following information:

- **Device name**: Indicates the device's vendor and model.
- **Address**: Displays the hardware address including the bus and node numbers.
- **Path**: Shows the location in the operating system.
- **Serial number**: Shows the serial number of the drive.
- **Drive Number**: Indicates the drive number and allows changing it. Changing the drive number may be required to address situations, where iSCSI targets are assigned incorrectly to the mounted drives.

  Note

  If more than one drive is selected, such drives should use the same host\buses.

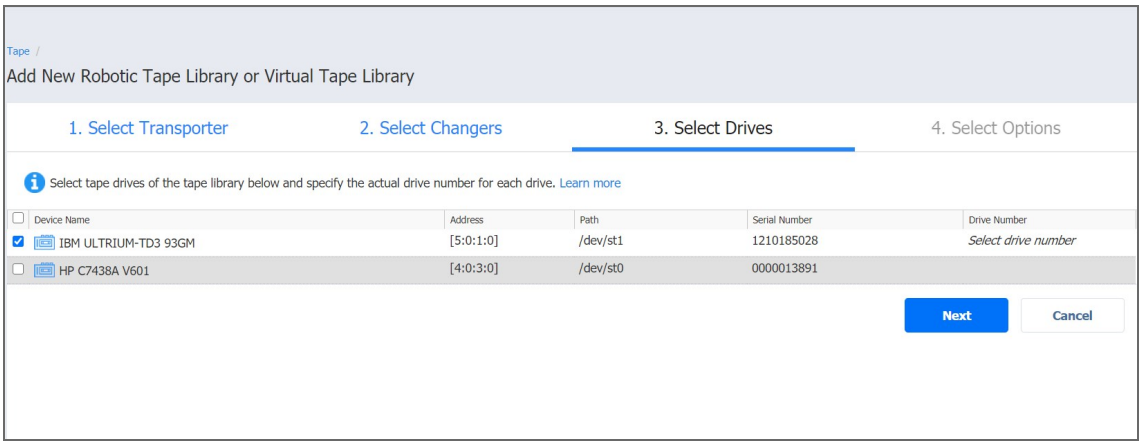

Click **Next** to proceed to the next page.

Selecting Options

The last step of adding a tape library is selecting its options.

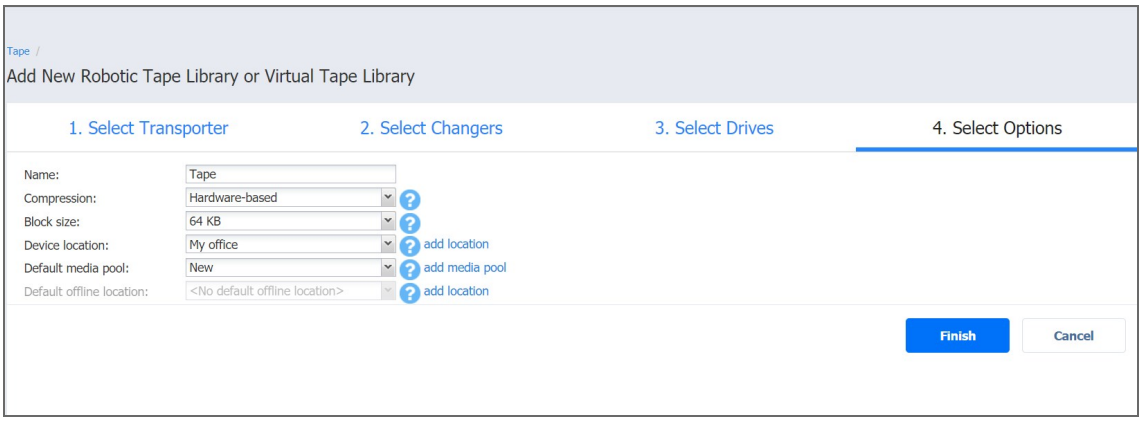

1. Specify the following parameters:

- **Name**: Enter the name for the tape library
- **Compression**: Select a compression level of the tape device:
  - 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
  - 128 KB
  - 256 KB (default)
  - 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 444
- **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 446
- **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.



Adding Standalone Tape Drives

Important

Vendor drivers should be installed on tape devices prior to adding them to NAKIVO Backup & Replication inventory.

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

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**.

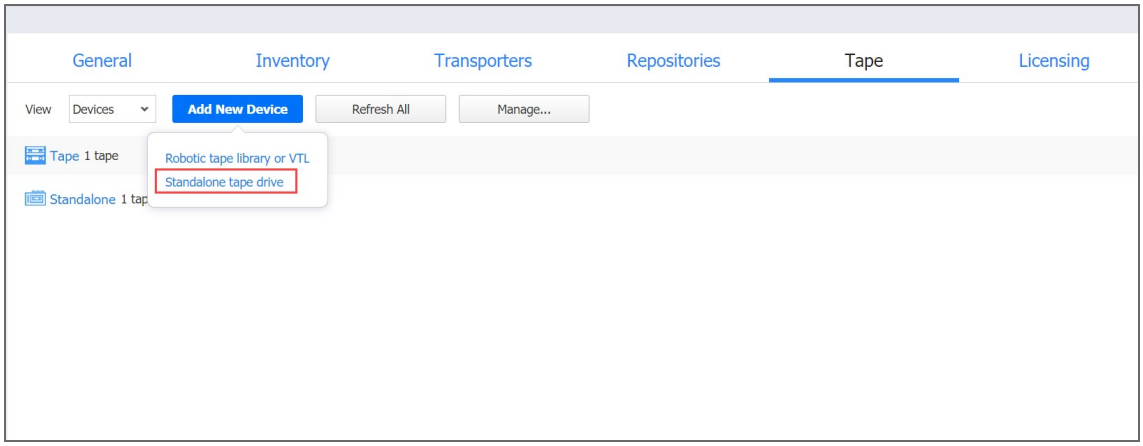

The **Add New Standalone Tape Drive** wizard opens. Follow the steps below to add a new tape drive.

Selecting Transporter

Important

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.

Refer to "Deploying Transporter as VMware Appliance" on page 345 and "Adding Installed Transporters" on page 353 for more information on Transporter deployment.

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.

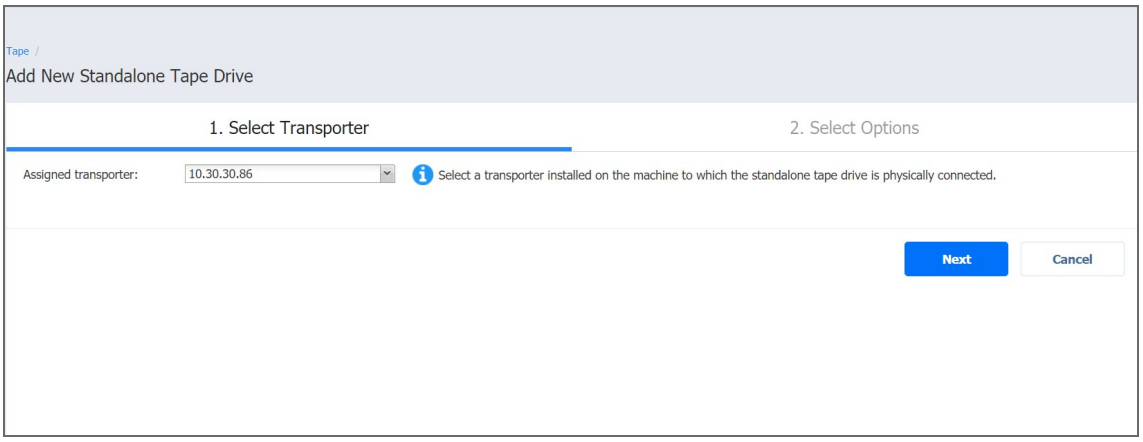

2. Click **Next**.

Selecting Options

The last step of adding a tape drive is selecting its options.

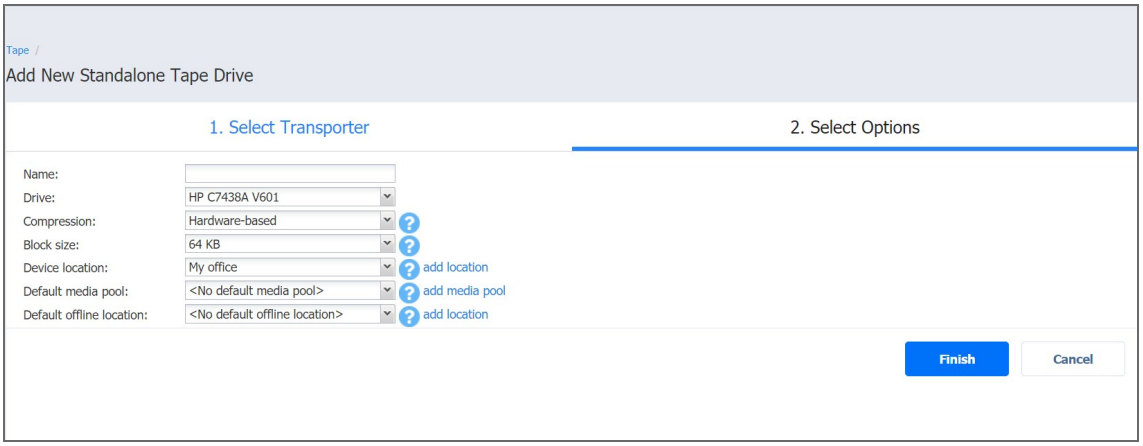

1. Specify the following parameters:

- **Name**: Enter the name of the tape library
- **Drive**: Select one of the standalone tape drives on the assigned transporter
- **Compression**: Select a compression level of the tape device:
  - 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
  - 128 KB

    - 256 KB (default)
    - 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 444
- **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 446
- **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.

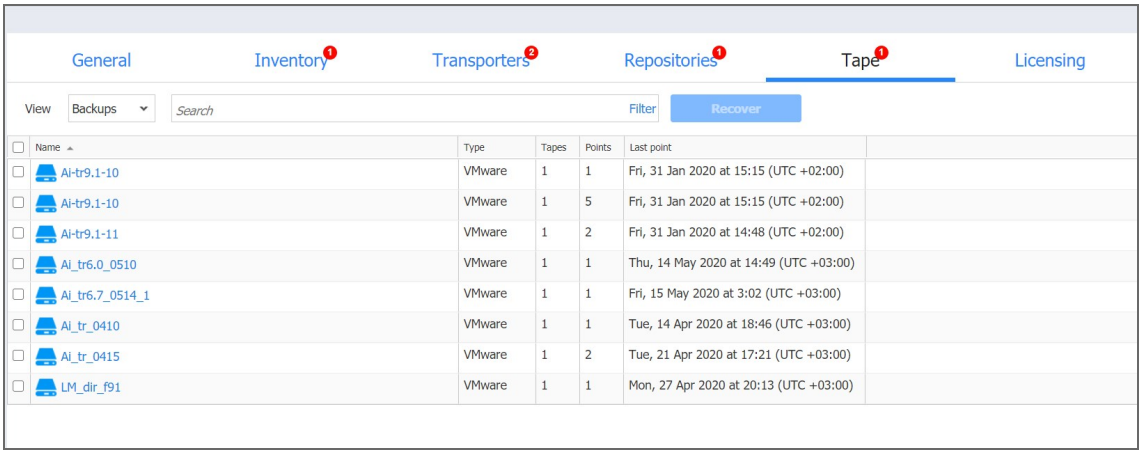

From the **Backups** view, you can search for backups, recover from backups, and view backups' details.

- 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.

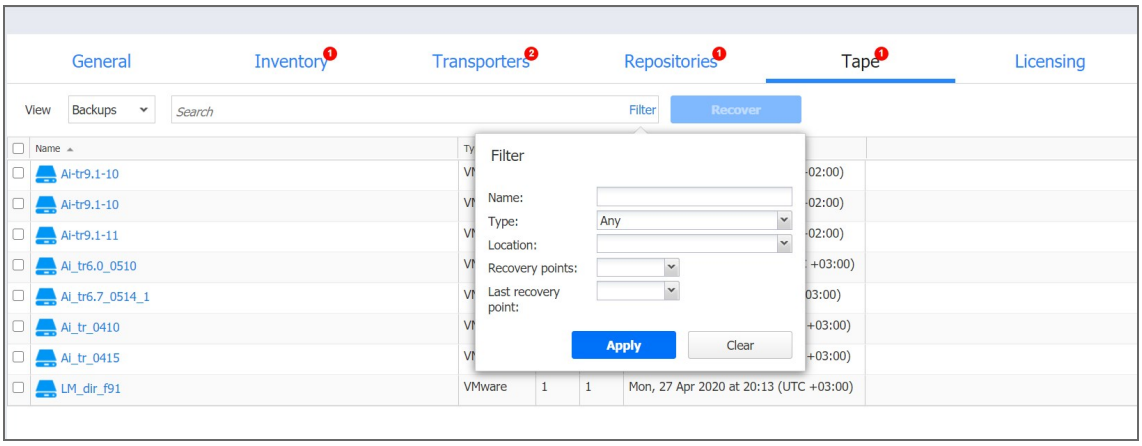

You can apply the following filtering criteria:

- **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/later 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.
- **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.

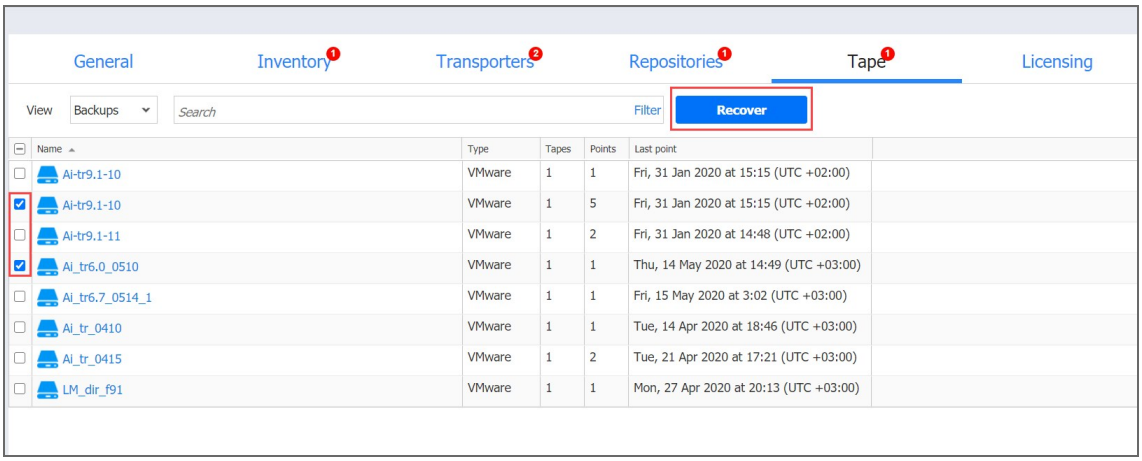

The **New Recovery Job Wizard** opens with the specified backups and their latest recovery points selected.

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 **Devices** view selected, click **Manage** and select **Locations**. The **Location Management** dialog opens.

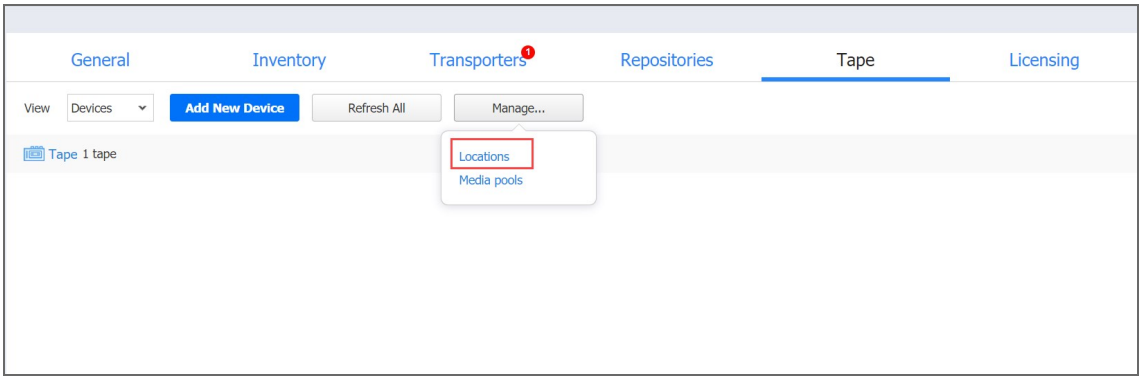

3. Click **Add New Location**.
4. In the **Add New Location** dialog box, specify a name for the device location and provide its description (optionally).

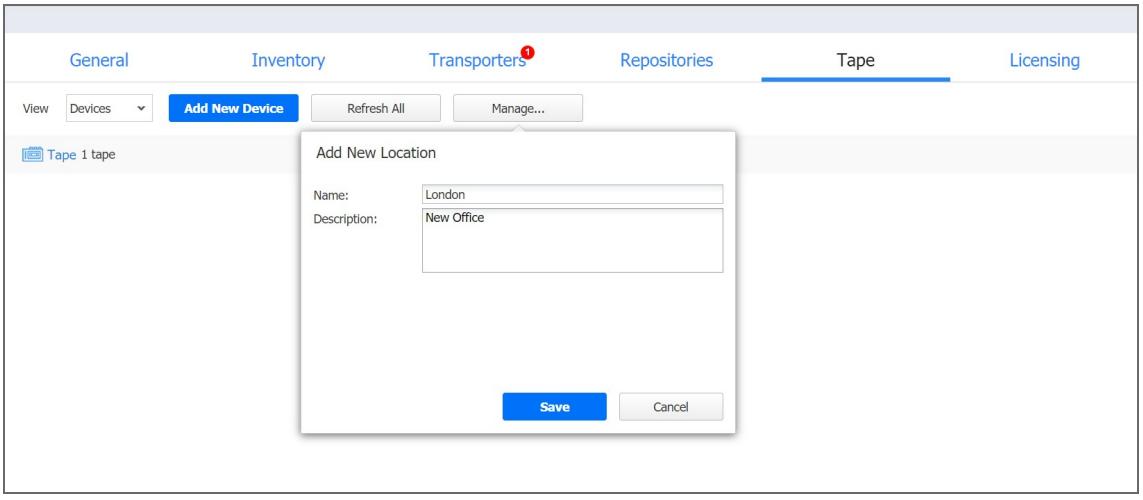

5. Click **Save**. The new device location is added to the list.

Managing Device Locations

From the **Location Management** screen page, 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.

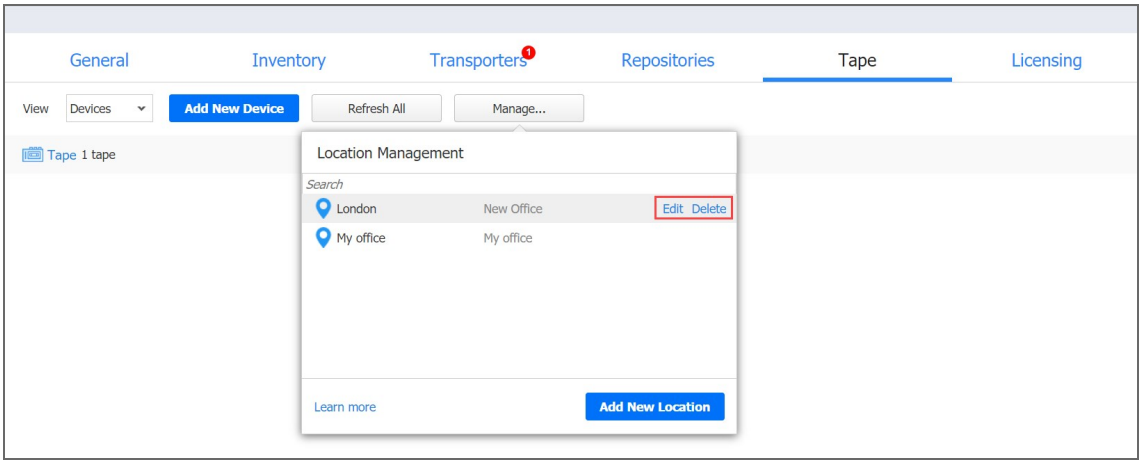

Managing Media Pools

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** view selected, click **Manage** and select **Media Pools**. The **Media Pool Management** dialog box opens.

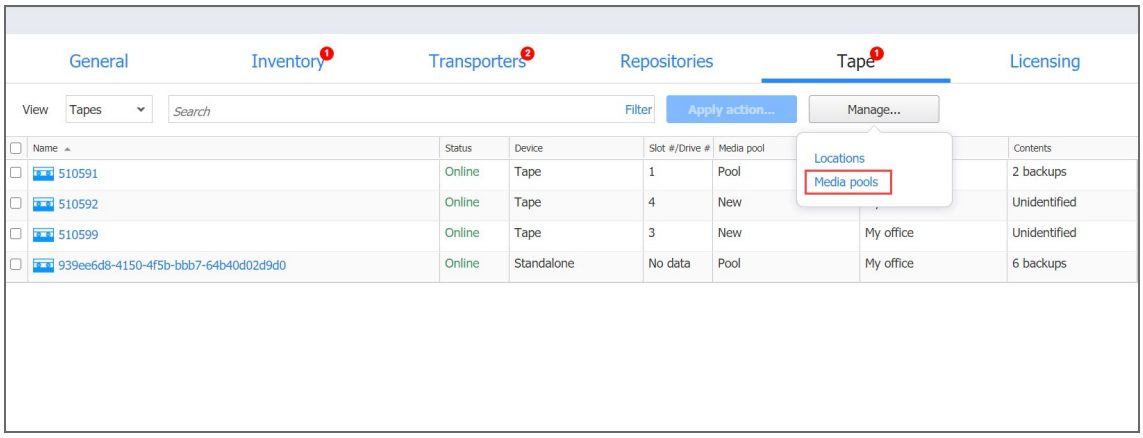

3. Click **Create Media Pool**.
4. In the **Create Media Pool** dialog box, specify the name for the Media Pool and provide its description (optionally).
5. From the **Move Offline Tapes To** drop-down list, select a device location to determine which location is automatically set for all offline tapes 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.

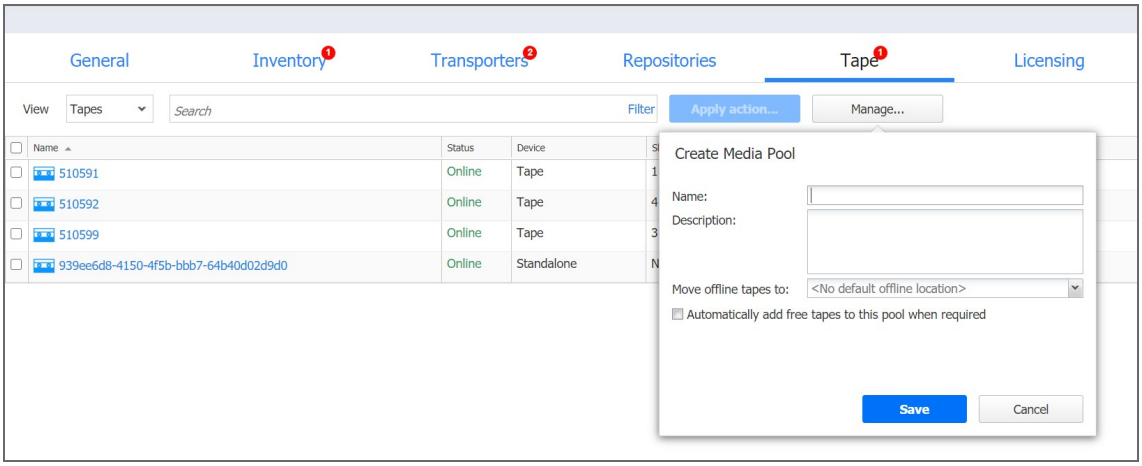

7. Click **Save**.The new Media Pool is created.

Managing Media Pools

From the **Media Pool Management** page, 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.

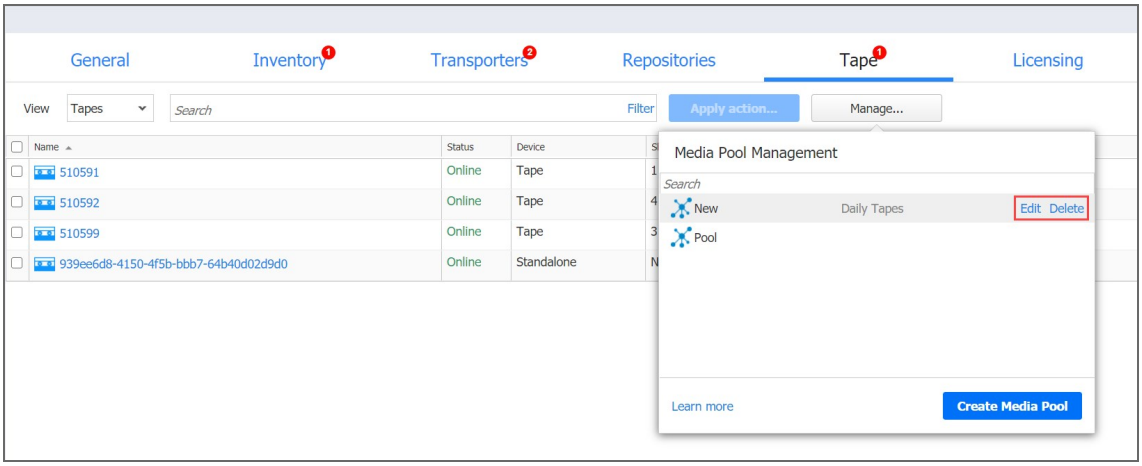

Managing Tape Cartridges

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.

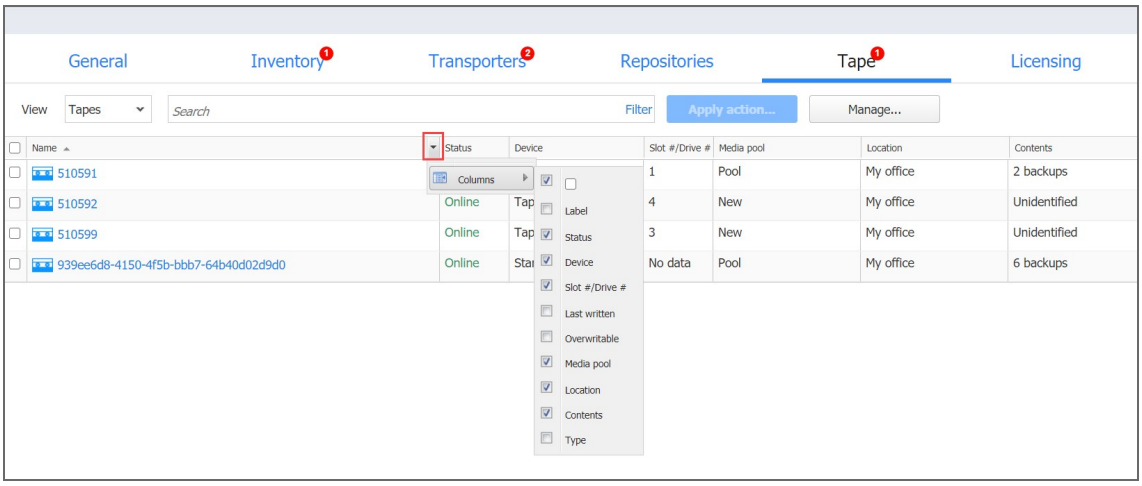

Searching for Tape Cartridges

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.

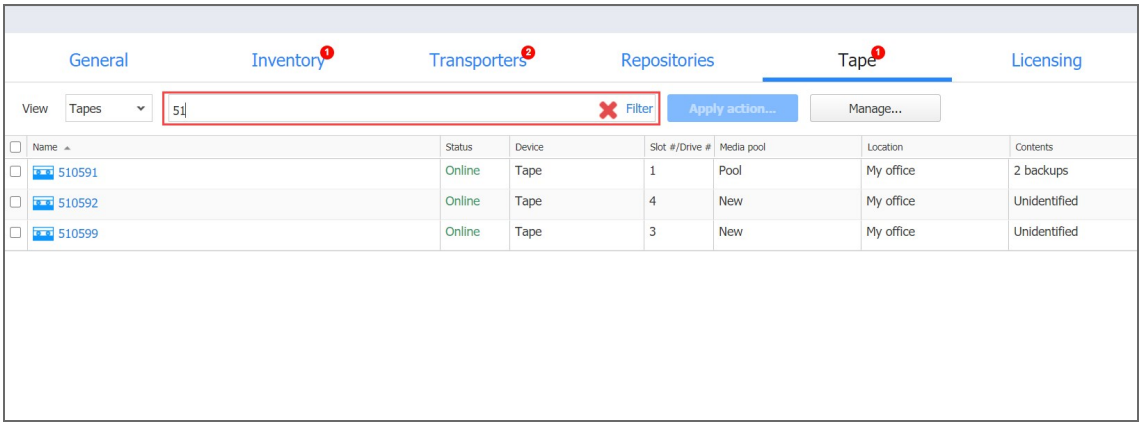

Clicking the **Clear** button in the search field will clear the query and the table will display all tape cartridges.

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.

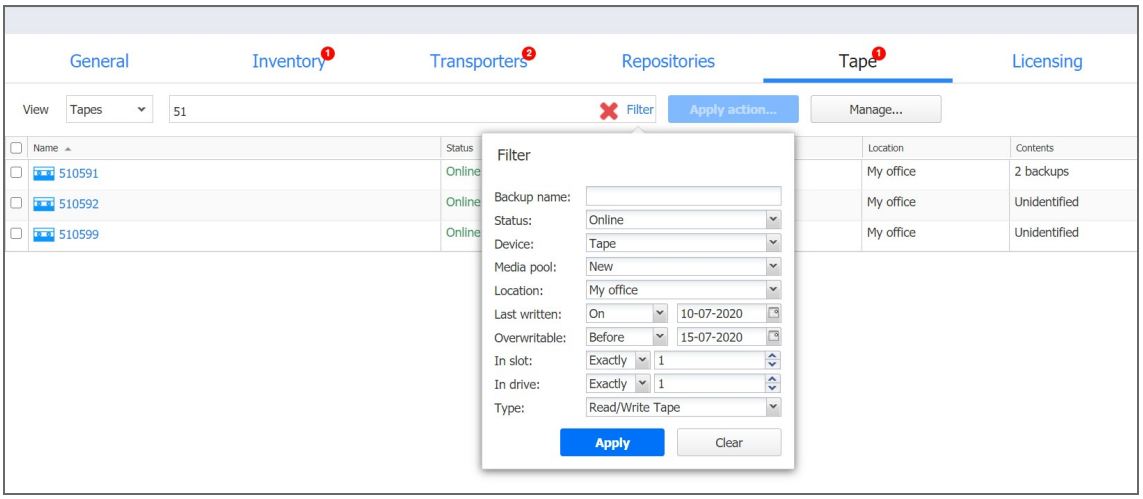

You can apply the following filtering criteria:

- **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

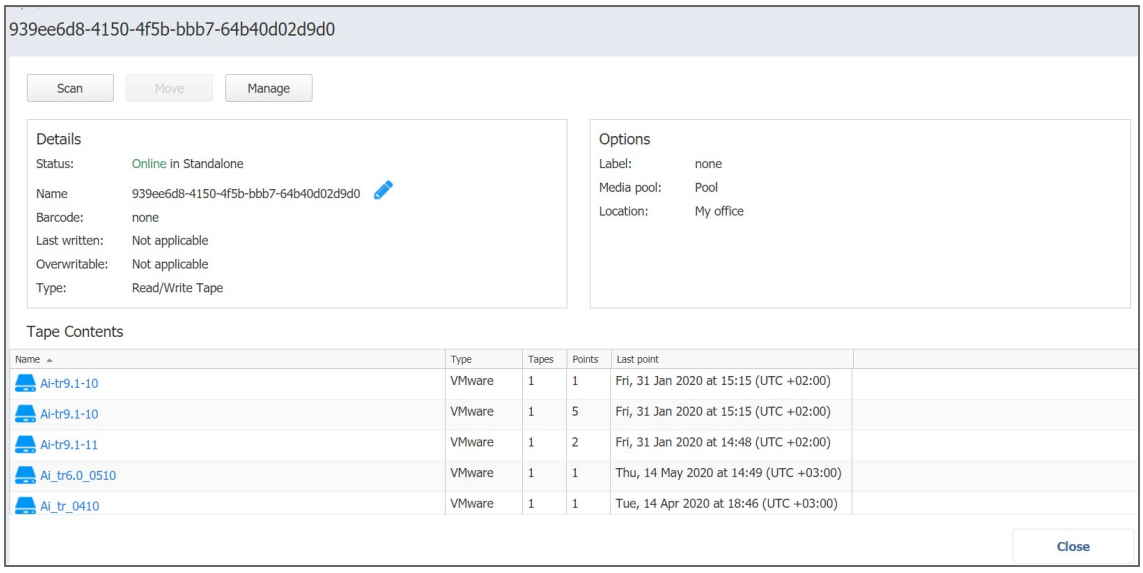

Action Buttons

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 456.

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.

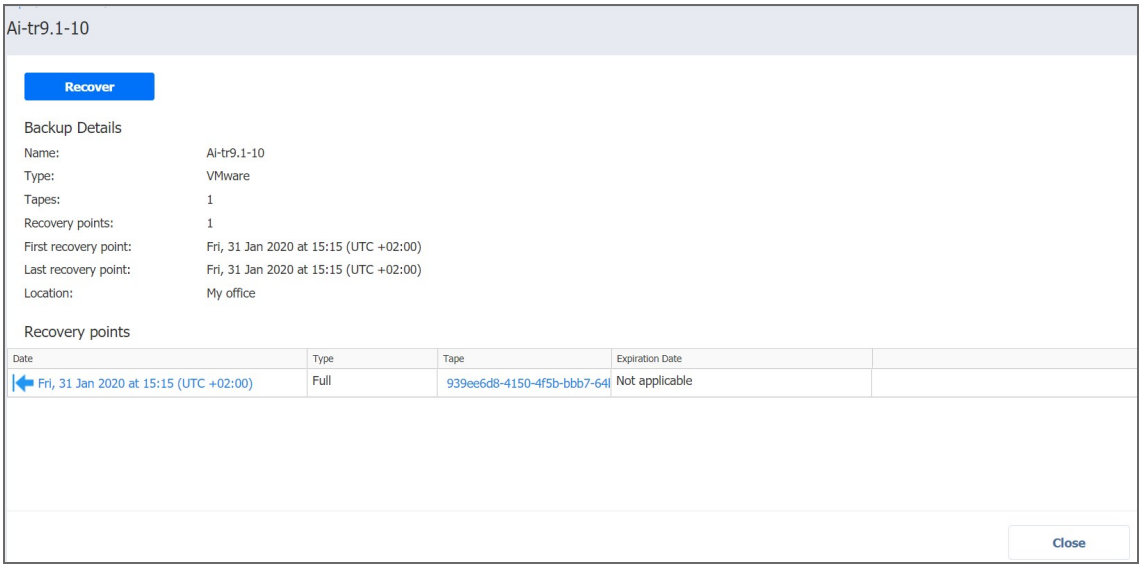

The **Backup Details** section provides the following information about the backup:

- **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 678 .

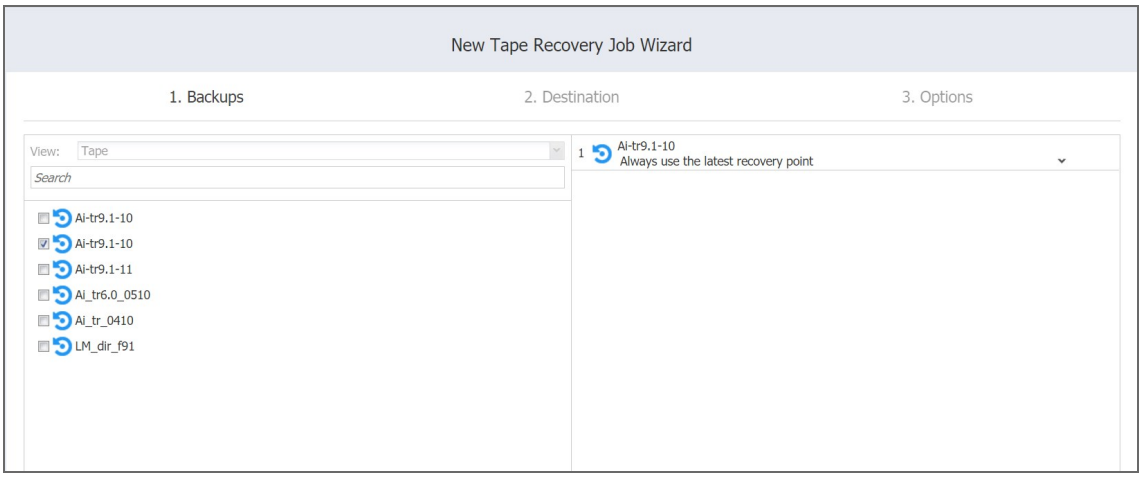

Bulk Tape Cartridge Management

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.

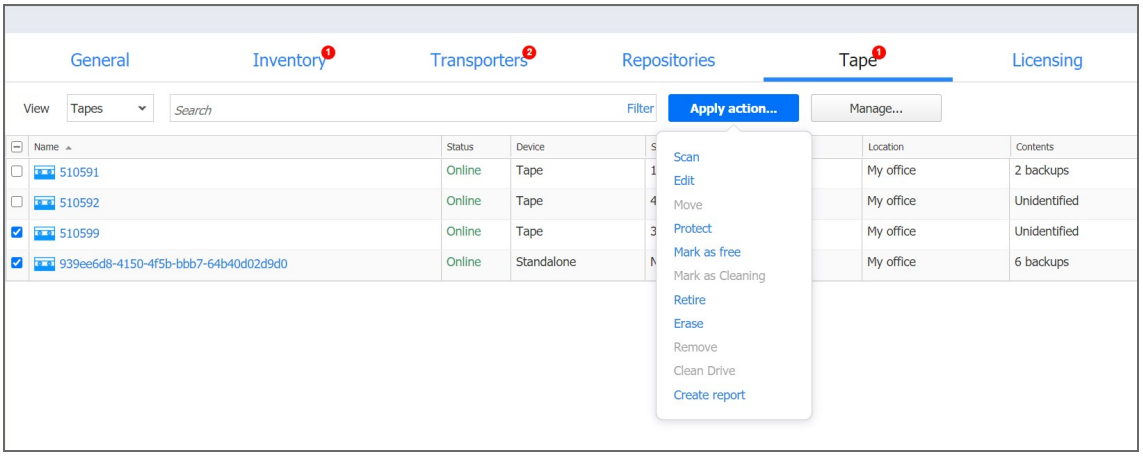

The **Create report** action is unique to the bulk tape cartridge management and is used to generate reports about selected tape cartridges. The report is created as a PDF file and is stored locally on your computer.

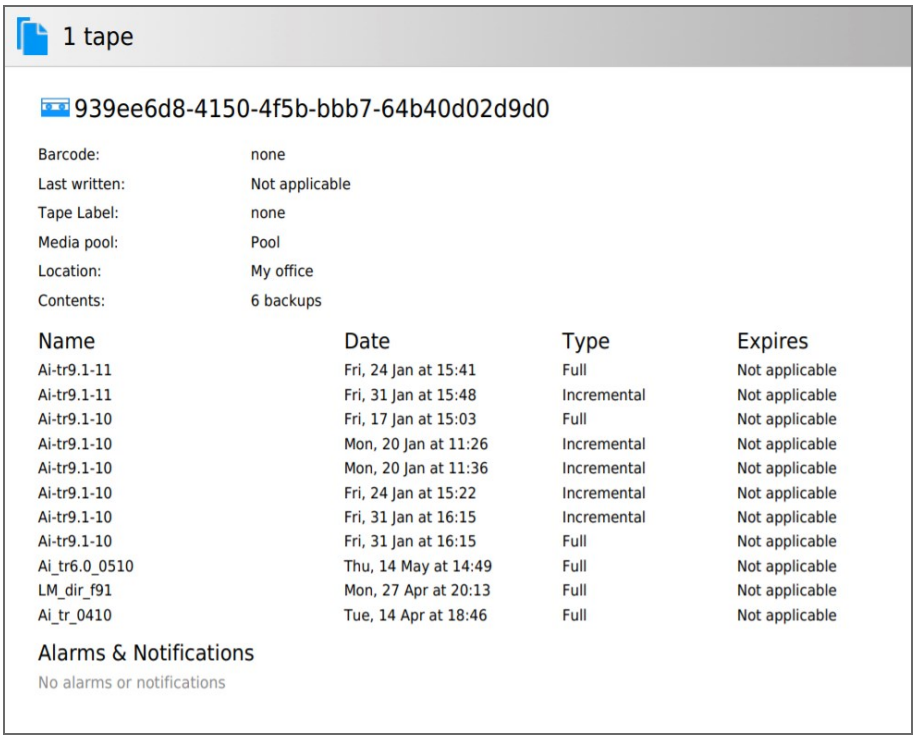
1 tape
939ee6d8-4150-4f5b-bbb7-64b40d02d9d0
Barcode: none
Last written: Not applicable
Tape Label: none
Media pool: Pool
Location: My office
Contents: 6 backups
Name Date Type Expires
Ai-tr9.1-11 Fri, 24 Jan at 15:41 Full Not applicable
Ai-tr9.1-11 Fri, 31 Jan at 15:48 Incremental Not applicable
Ai-tr9.1-10 Fri, 17 Jan at 15:03 Full Not applicable
Ai-tr9.1-10 Mon, 20 Jan at 11:26 Incremental Not applicable
Ai-tr9.1-10 Mon, 20 Jan at 11:36 Incremental Not applicable
Ai-tr9.1-10 Fri, 24 Jan at 15:22 Incremental Not applicable
Ai-tr9.1-10 Fri, 31 Jan at 16:15 Incremental Not applicable
Ai-tr9.1-10 Fri, 31 Jan at 16:15 Full Not applicable
Ai_tr6.0_0510 Thu, 14 May at 14:49 Full Not applicable
LM_dir_f91 Mon, 27 Apr at 20:13 Full Not applicable
Ai_tr_0410 Tue, 14 Apr at 18:46 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.

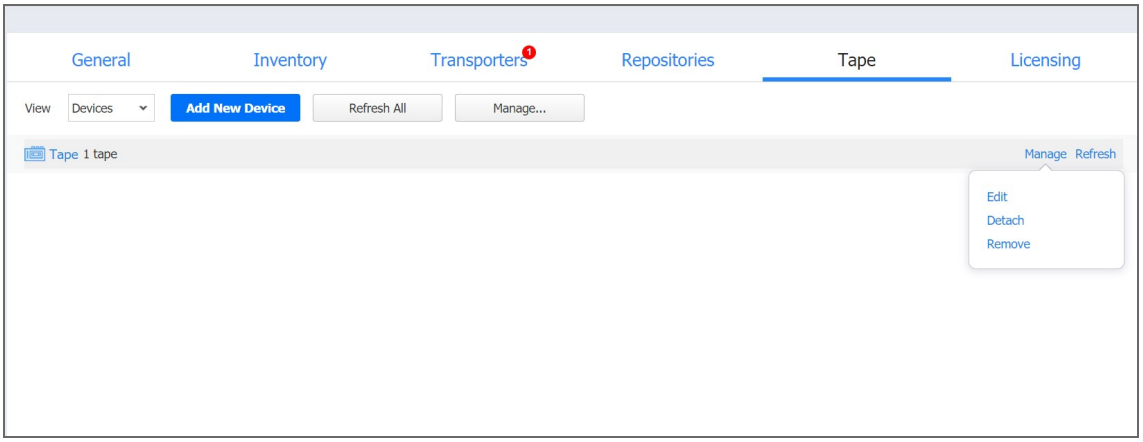

Clicking the name of the tape device opens the device's details window where you can manage it and view the

device's detailed information.

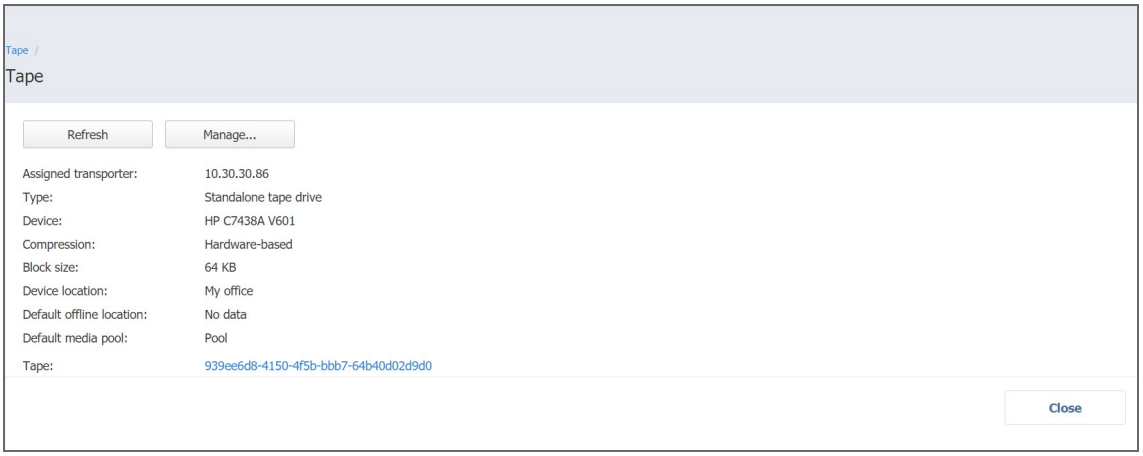

Virtual Appliance Configuration

This section covers the following topics:

- "Configuring Network Settings of Virtual Appliance" on page 460
- "Increasing Backup Repository Size on Virtual Appliance" on page 461
- "Removing the Disk with Backup Repository from Virtual Appliance" on page 462

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" on page 464
- "Configuring Branding Settings in Multi-Tenant Mode" on page 465
- "Configuring Email Notifications in Multi-Tenant Mode" on page 467
- "Configuring Email Settings in Multi-Tenant Mode" on page 468
- "Configuring System Settings in Multi-Tenant Mode" on page 469
- "Exporting and Importing Configuration in Multi-Tenant Mode" on page 471

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 **Configuration** in the upper right corner of the product.
3. Go to the **General** tab and click **Branding**.

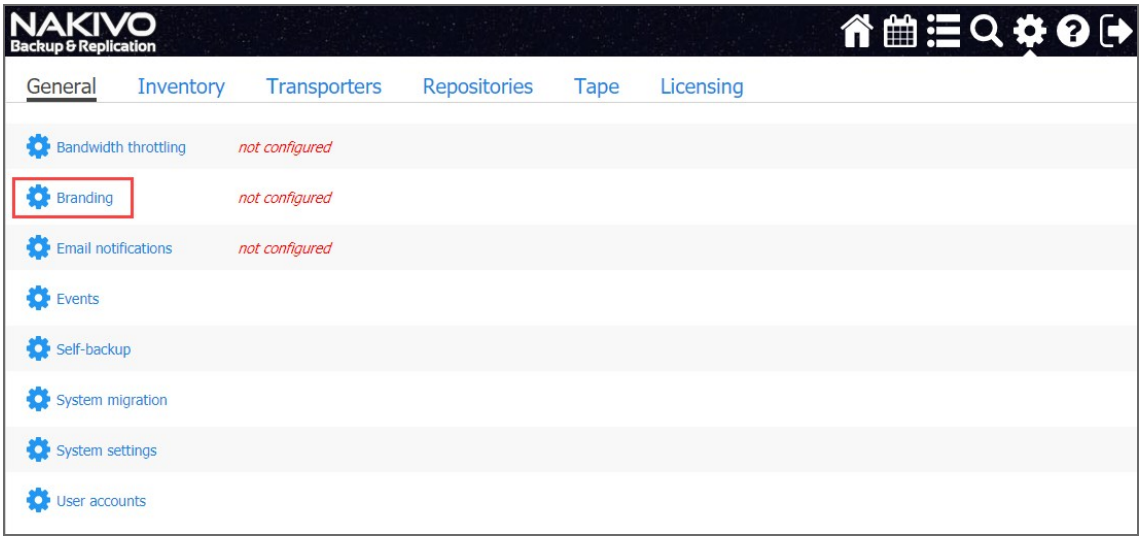

4. Do the following:
   - To change the product title, company name, website URL, contact email, support email, and contact phone, type a new value in the appropriate field
   - To change the product logo, background, and default tenant logo, click **Change** under the appropriate box, select a new image, and click **Open**.

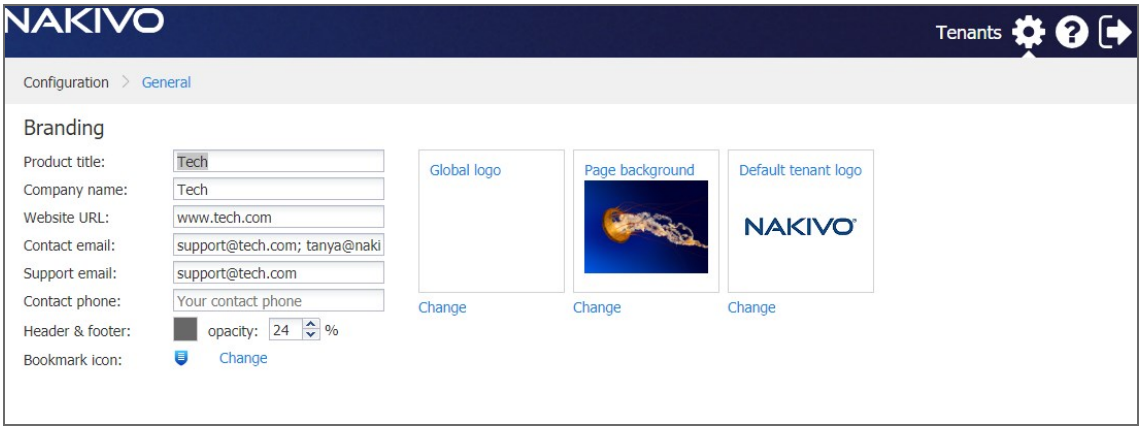

5. Click **Apply**.

NOTE: During upload, the logo and bookmark icon images are resized internally while preserving the aspect ratio. The background image is used as it is. To get the best image quality, follow the recommendations below:

Image	Best format	Best resolution
Global logo	.png	40x40
Page background	.jpeg	1920x1440
Bookmark icon	.png	16x16
Default tent logo	.png	120x95

Configuring Email Notifications in Multi-Tenant Mode

NAKIVO Backup & Replication can send notifications and reports over email. To configure the email notifications, follow the steps below:

1. Make sure you have configured your email settings.
2. Log in to NAKIVO Backup & Replication as a Master Admin.
3. Click **Configuration** in the upper right corner 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 **Configuration** in the upper right corner of the product.
3. Go to the **General** tab and click **Email notifications**.
4. In the **Email Settings** section, enter data in the boxes, and click **Send Test Email** to verify the settings are correct.

After the email settings are configured, you can configure the product email notifications.

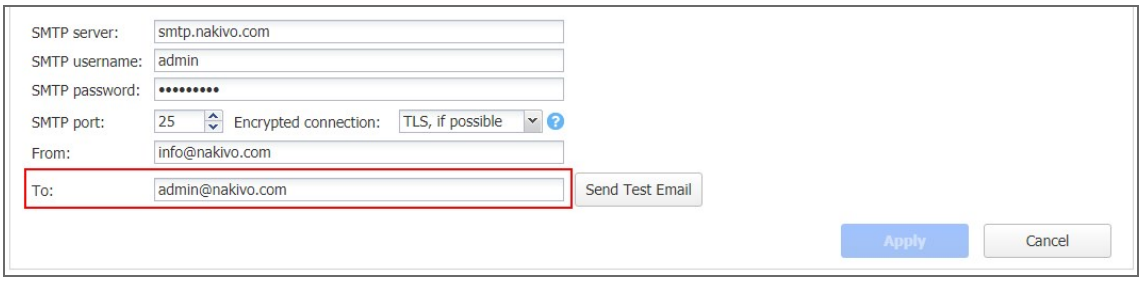

Configuring System Settings in Multi-Tenant Mode

To configure the system settings, follow the steps below:

1. Log in to NAKIVO Backup & Replication as a Master Admin.
2. Click **Configuration** in the upper right corner 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 **Configuration** of 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 **Configuration** of 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

To create a support bundle, follow the steps below:

Important

Before creating a support bundle, make sure Email settings are configured.

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.

10. Optionally, click **Download** to save the support bundle on your machine.

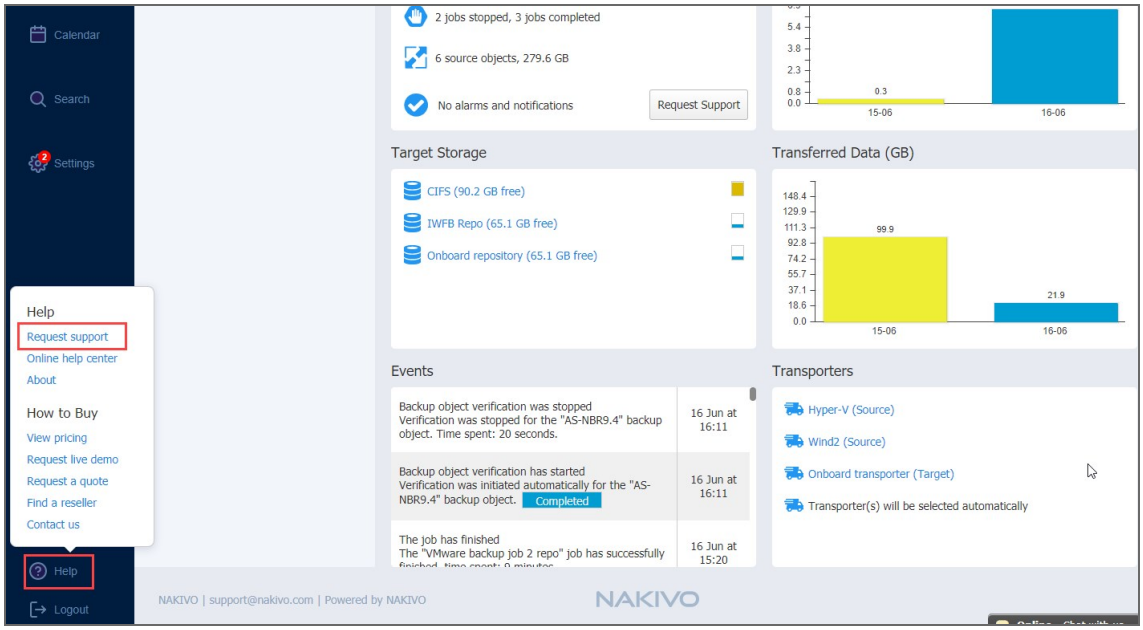

Sending Support Bundles Manually

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:

1. Open the [Upload Files to NAKIVO Support]page.
2. In the *Files* section, click **Browse** and select up to three files. You can select more than three files by clicking **Add Another File**.

   Note

   You can upload any files relevant to your issue: logs, file dumps, or the support bundles that you have manually downloaded from the product's UI.
3. Enter your email address in the **Contact email** field.
4. You can also enter the ID of your support ticket in the **Ticket ID** field if you have one opened.
5. Optionally, enter a description in the **Description** field.
6. Click **Upload** when you're done uploading the file(s).

   Note

   Wait for a successful upload notification before closing the page.

Built-in Support Chat

You have the possibility to contact a NAKIVO representative via chat in the NAKIVO Backup & Replication interface.

- 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.

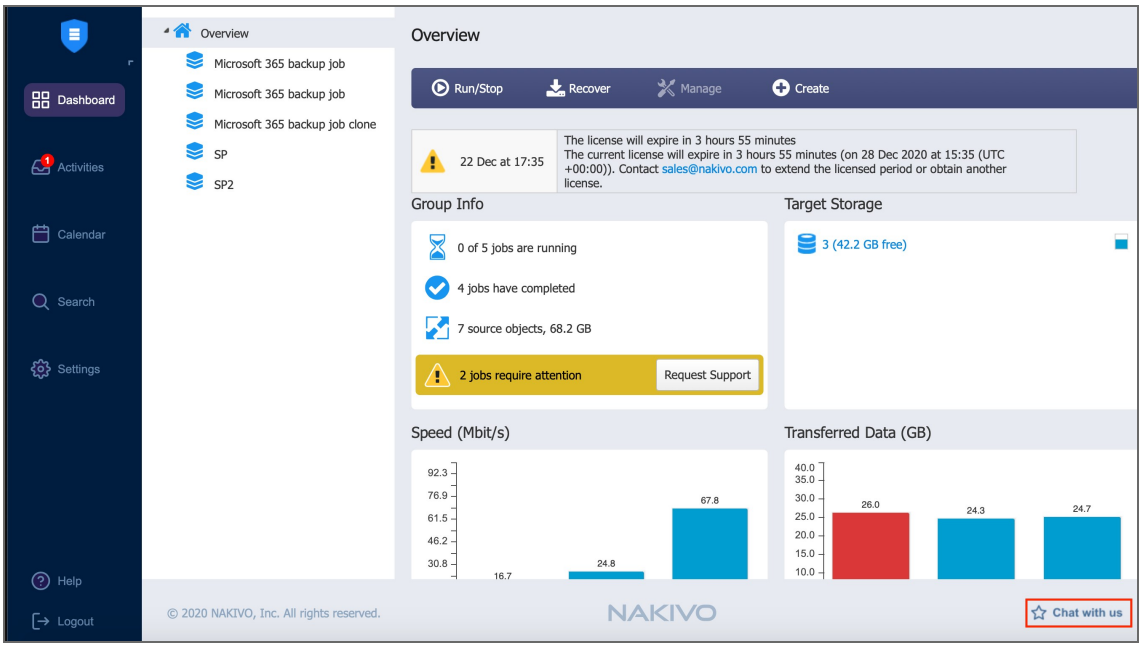

2. The **NAKIVO Support** dialog box opens. Introduce yourself by providing the following information:
   a. In the upper box of the dialog box, enter your name.
   b. In the box below, enter your email address.
3. Choose a department from the list of available departments.

4. Enter your message text and click **Start Chatting**.

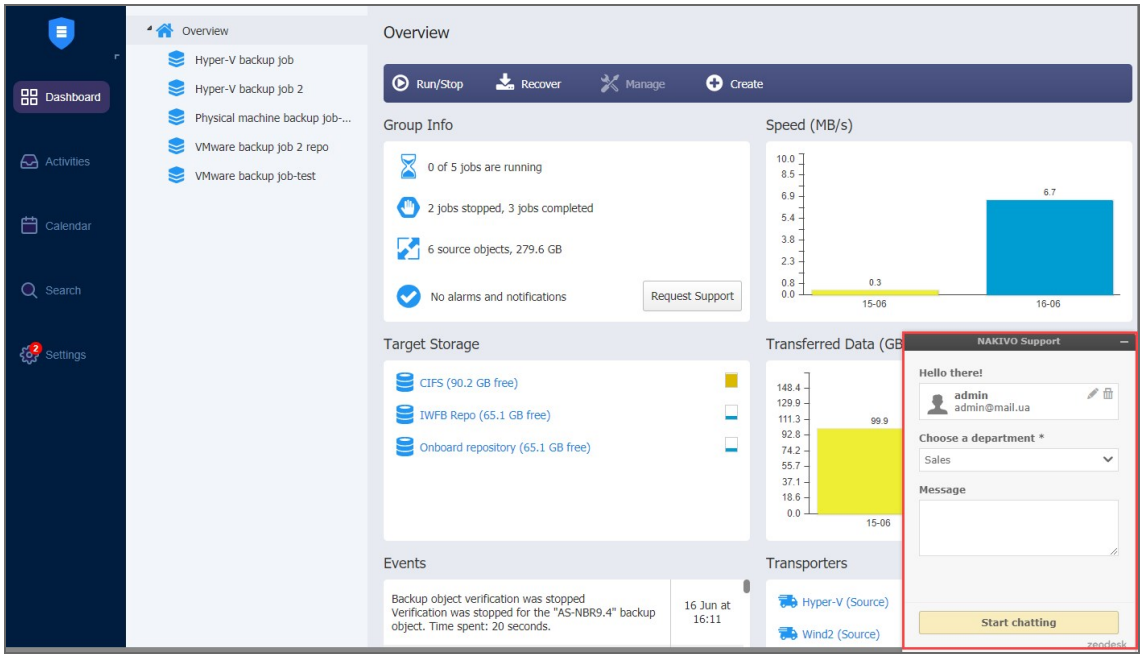

5. Your message is sent to a NAKIVO representative and will be processed as soon as possible. If needed, click the **Send Another** button to proceed with sending another chat message.

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**.

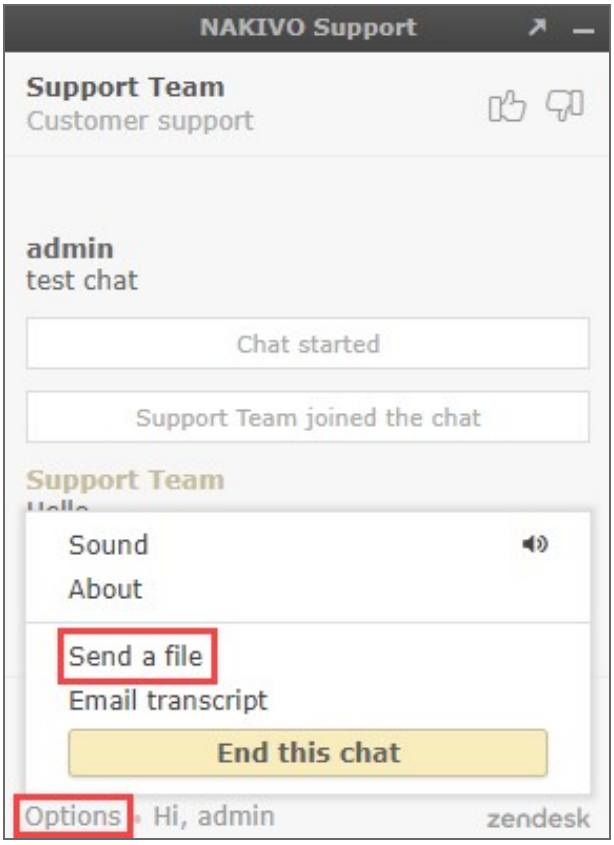

Note

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**.

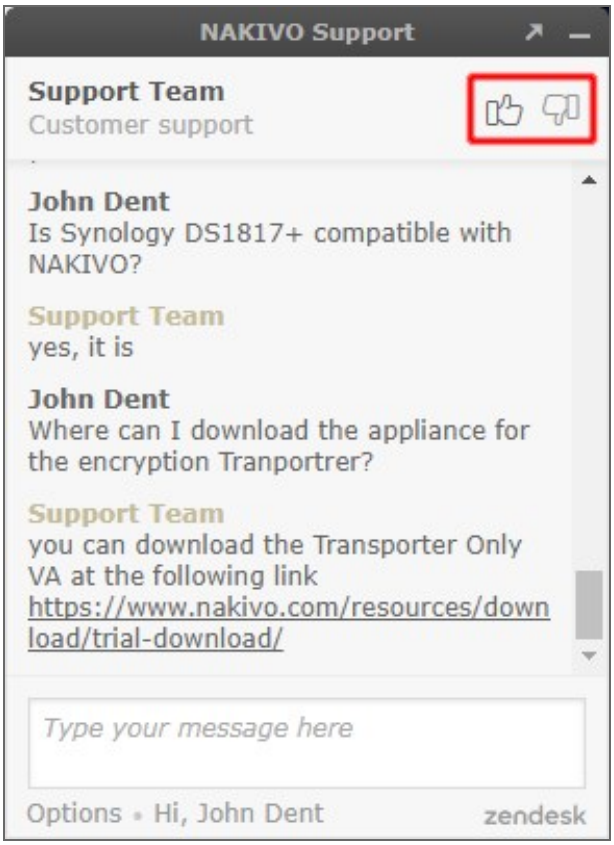

Sending Email Transcript of Built-in Support Chat

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.

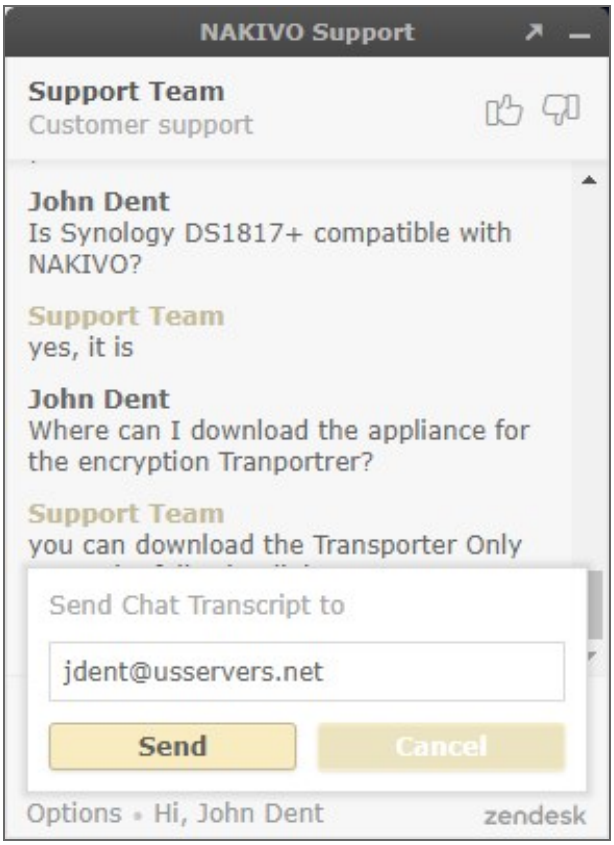

Disabling/Enabling Sound Notifications

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.

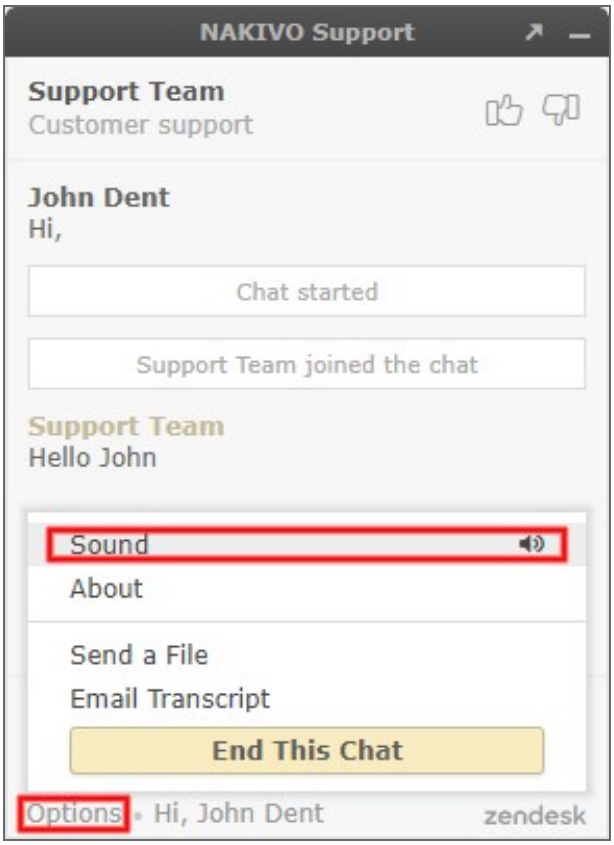

Disabling Built-in Support Chat

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.

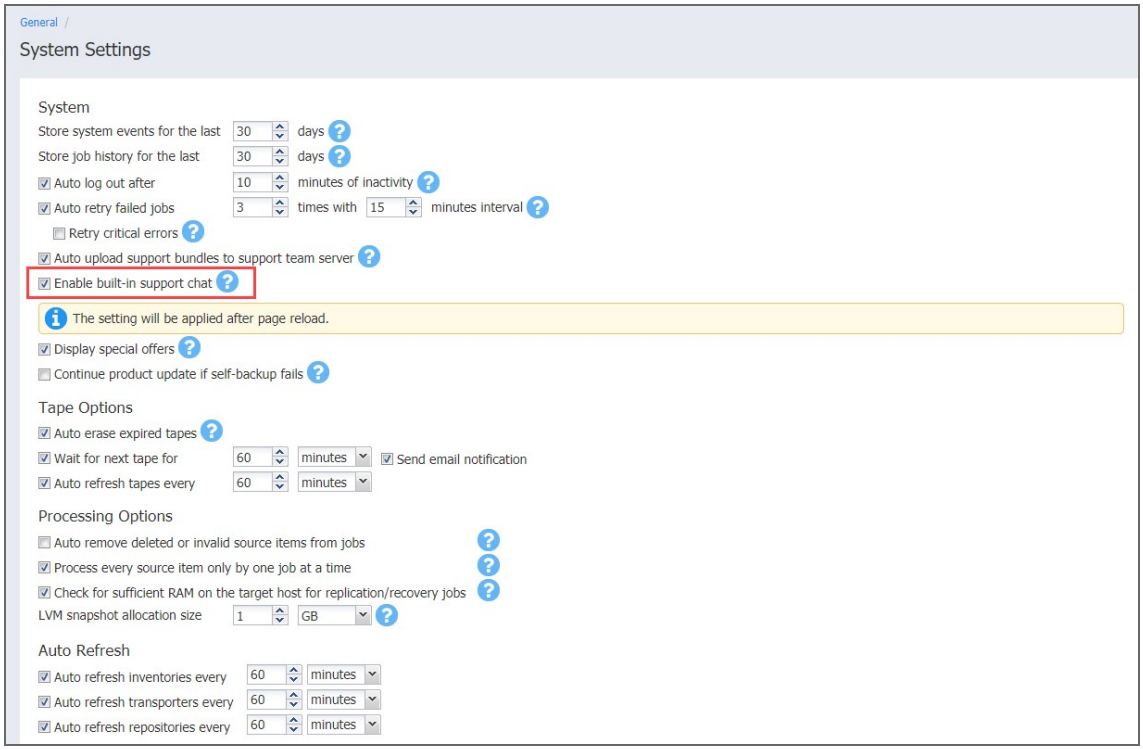

Note

When disabled, the Built-in Support Chat will not be available in all tenants of the NAKIVO Backup & Replication instance in multi-tenant mode.

Replacing License

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**.

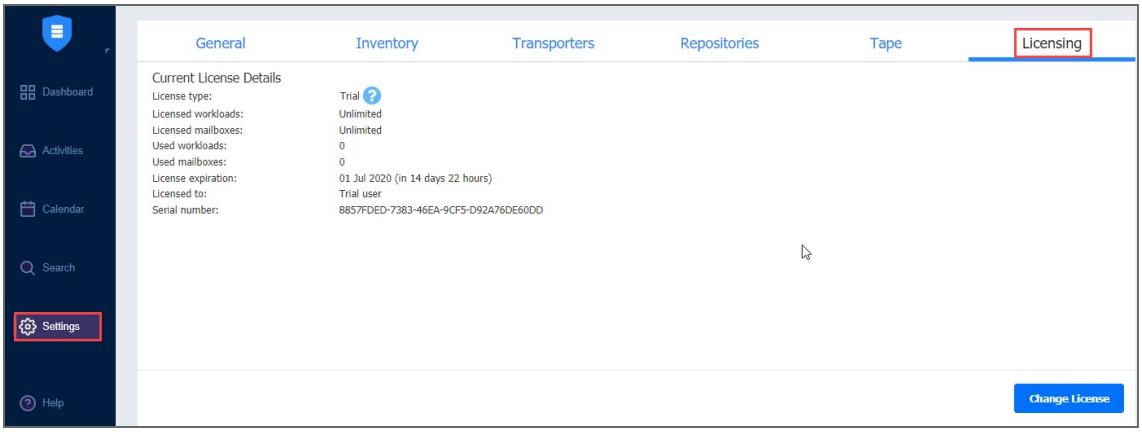

3. Locate and open the license file in the window that appears.

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.

Expert Mode

For advanced NAKIVO Backup & Replication configuration, you can enable the Expert mode.

To do this, take the following steps:

1. Log in to your NAKIVO Backup & Replication instance.
2. Add the word “expert” to the URL parameters of the **Settings** page.

   Examples:

   https://localhost:4443/c/configuration?expert or

   https://localhost:4443/c/configuration?action=&targetId=&backUrl=&wizard=false&expert
3. Click the **Expert** tab.

Configuring Settings

To configure advanced product settings, make the necessary changes in the following parameters:

Parameters	Description	Possible Values
system.email.smtp.localhost.mode:	Specifies how to determine the name of the localhost that is used in the SMTP HELO or EHLO commands.	• Default • Use DNS • Provide custom hostname
system.email.smtp.localhost.name:	Specifies the name of the localhost that is used in the SMTP HELO or EHLO commands. This setting is valid for custom hostname resolution mode only.	
system.email.notifications.skip.event.list:	List of event names to skip when creating an email digest. Use space or "," or ";" as separators. The event names can be found in events.log.	error60 error61 other
system.vmware.esxi.ssh.port:	For VMware only. Specifies the SSH port to connect to ESXi (global setting).	• Default value: 22 • Minimum value: 1

		• Maximum value: 65535
system.vmware.skip.outdated.tools.checking:	For VMware only. When enabled, the system does not check VMware Tools outdated status when creating quiescing snapshot.	• Checked • Unchecked (default)
system.vmware.skip.tag.discovery:	VMware only. When enabled, the system does not discover VMware Tags. This is applied to all tenants.	• Checked • Unchecked (default)
system.vmware.job.run.hotadd.target.datastore.min.free.space.percent:	For VMware only. Specifies the minimum free space for the target datastore during replication job run with HotAdd enabled. If the free space goes below this value, an alarm is generated.	• Default value: 63 • Minimum value: 1 • Maximum value: 100
http.max.upload.size:	Specifies the max upload size for file upload operations, **bytes** (global setting). If multiple files are uploaded, this is the total size. Use -1 for unlimited. Example: 200MB: 200000000	• Default value: 1073741824 • Minimum value: 1 • Maximum value: 9999999999
system.auth.use.lockout:	Enables or disables the login lockout feature. When enabled, the offending IP address is not allowed to login after several failed attempts.	• Checked • Unchecked (default)
system.auth.max.login.attempt.count:	Specifies the maximum number of failed login attempts to trigger the login lockout feature for the offending IP.	• Default value: 5 • Minimum value: 1 • Maximum value: 9999

system.auth.lockout.timeout:	Specifies the timeout (minutes) for the login lockout feature. The offending IP is allowed to login again after the timeout expires.	• Default value: 15 • Minimum value: 1 • Maximum value: 9999
system.auth.login.history.period:	Specifies the period (minutes) to calculate the maximum number of failed login attempts for the login lockout feature.	• Default value: 5 • Minimum value: 1 • Maximum value: 9999
system.auth.ad.integration.follow.referrals:	Defines LDAP/Active Directory behavior for referrals. When set to follow, all referrals are resolved (can be slow); otherwise they are ignored. What are the implications of the ignore option? * If you only have one domain, there should be no effects. * If you have multiple domains joined in a forest, then any cross-domain memberships will not be resolved. More info: https://docs.oracle.com/javase/jndi/tutorial/ldap/referral/jndi.html	• follow (default) • ignore
system.job.pool.queue.length:	Specifies the length of the job queue. A job is placed in a queue before execution. Requires restart.	• Default value: 200 • Minimum value: 10 • Maximum value: 9999

system.job.pool.thread.min:	Specifies the minimum thread pool size for jobs. A job requires 1 thread from the job pool to start running. Requires restart.	• Default value: 30 • Minimum value: 10 • Maximum value: 9999
system.job.pool.thread.max:	Specifies the maximum thread pool size for jobs. A job requires 1 thread from the job pool to start running. When the pool thread limit is reached, the job is placed in the job queue. Requires restart. If using Linux and systemd, please add the following to the service startup script: TasksMax=infinity	• Default value: 200 • Minimum value: 10 • Maximum value: 9999
system.job.resolve.host.hostname.on.transporter:	If set, sends the source and/or target host hostname as is to Transporter. Transporter will resolve the hostname to the IP address(es) and check if it is reachable. This is done during the Transporter to host checks on a job run. 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).	• Checked • Unchecked (default)
system.job.resolve.transporter.hostname.on.transporter:	If set, sends the source and/or target Transporter hostname as is to Transporter. Transporter will resolve the hostname to the IP address(es) and check if it is reachable. This is done during Transporter to Transporter checks on job run.	• Checked • Unchecked (default)

	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).	
system.job.bandwidth.throttling.source:	If set, applies bandwidth throttling for data reading from source.	• Checked (default) • Unchecked
system.job.bandwidth.throttling.target:	If set, applies bandwidth throttling for data writing to target.	• Checked (default) • Unchecked
system.job.bandwidth.throttling.network:	If set, applies bandwidth throttling for data transfer between source and target.	• Checked (default) • Unchecked
system.task.pool.queue.length:	Specifies the length of the task queue. A task is placed in the queue before execution. Requires restart.	• Default value: 200 • Minimum value: 10 • Maximum value: 9999
system.task.pool.thread.min:	Specifies the minimum thread pool size for tasks. A task requires 1 thread from the task pool to start running. Task example: repository refresh, Transporter refresh, support bundle creation. Requires restart.	• Default value: 30 • Minimum value: 10 • Maximum value: 9999

system.task.pool.thread.max:	Specifies the maximum thread pool size for tasks. A task requires 1 thread from the task pool to start running. When the pool thread limit is reached, the task is placed in the task queue. Task example: repository refresh, Transporter refresh, support bundle creation. Requires restart.	• Default value: 200 • Minimum value: 10 • Maximum value: 9999
system.repository.min.free.space.byte:	Specifies the minimum free space (bytes) for the repository. If the free space goes below this value, an alarm is generated.	• Default value: 5368709120 • Minimum value: 1024 • Maximum value: 1099511627776
system.repository.min.free.space.percent:	Specifies the minimum free space (percent) for the Backup repository. If the free space goes below this value, an alarm is generated.	• Default value: 5 • Minimum value: 1 • Maximum value: 99
system.repository.maintenance.interrupt.timeout.seconds:	Specifies the timeout (seconds) to wait for repository maintenance stop during job run.	• Default value: 300 • Minimum value: 1 • Maximum value: 86400

system.repository.refresh.backup.size.calculation:	Specifies the backup size calculation on the repository refreshing. True: Always calculates backup size. False: Skips backup size calculation and only calculates backup size with necessary backups.	• Checked (default) • Unchecked
system.product.skip.update.server.ssl.certificate.verification:	The product update check process requires the remote server certificate to be trusted. This parameter disables such check. It can be useful when secure (SSL/TLS) connections are being intercepted by third-party software.A product restart is required to apply.	• Checked • Unchecked (default)
system.debug.mode.enabled:	The debug mode prints more information into the logs, including some sensitive one (hardware UUIDs, MAC addresses, etc). The passwords are not printed unless they are present in raw communication dumps (e.g., SOAP/XML/JSON).	• Checked • Unchecked (default)
system.debug.mode.log.passwords:	When debug mode is enabled, also log passwords. This can be a security risk.	• Checked • Unchecked (default)
system.debug.mode.log.api.requests:	When debug mode is enabled, also log product API requests/responses. The data is logged as is and will contain plaintext passwords. This can be a security risk.	• Checked • Unchecked (default)
system.hyperv.optimize.queries:	Hyper-V only. Instructs to use a faster query method to read VM and host information. This will speed up the refresh process in large environments.	• Checked (default) • Unchecked

system.hyperv.discovery.host.thread.count:	Hyper-V only. Sets the max parallel threads to run when refreshing cluster hosts during discovery. Each cluster host can be refreshed separately. This will speed up the refresh process in large environments.	• Default value: 20 • Minimum value: 1 • Maximum value: 20
system.hyperv.discovery.vm.thread.count:	Hyper-V only. Sets the max parallel threads to run when refreshing host VMs during discovery. When increasing the setting value, make sure to test its impact on host CPU usage during refresh. This will speed up the refresh process in large environments.	• Default value: 2 • Minimum value: 1 • Maximum value: 10
system.database.scheduled.backup.path:	Specifies the target path for database backups. The tenant databases will be stored in subfolders, if present. The path can be local or absolute. The folder will be created automatically if it does not exist.	• Text field • Optional • Default value: "userdata"
system.database.scheduled.backup.max.count:	Specifies the maximum number of files for periodic database backups. The number is applied separately to each tenant database. The master and tenants product databases are backed up each day.	• Default value: 5 • Minimum value: 0 • Maximum value: 365
system.logging.max.index:	Specifies the maximum index of log files. This works globally for all log files. Set 0 to use default value (configured in log4j.xml).	• Default value: 0 • Minimum value: 0 • Maximum value: 999

system.product.min.free.space.byte:	Specifies the minimum free space (bytes) for the product installation folder. If the free space goes below this value, an alarm is generated.	• Default value: 2147483648 • Minimum value: 10485760 • Maximum value: 10737418240
system.product.free.memory.threshold:	Specifies the minimum ratio for JVM free memory. If the free JVM memory goes below this value, an alarm is generated.	• Default value: 0.1 • Minimum value: 0.01 • Maximum value: 0.9
system.nutanix.discovery.vm.thread.count:	Nutanix AHV only. Sets the max parallel threads to run when refreshing host VMs during discovery. When increasing the setting value, make sure to test its impact on host CPU usage during refresh. This will speed up the refresh process in large environments.	• Default value: 2 • Minimum value: 1 • Maximum value: 10
system.plugin.flr.operation.timeout.seconds:	Specifies the timeout (seconds) to wait for plugin session FLR/OLR. This is a low-level setting that is only sent to Transporter and used during iSCSI interaction.	• Default value: 900 • Minimum value: 1 • Maximum value: 86400
system.physical.skip.os.checking:	Physical Windows host discovery only. When enabled, the system will not check the supported OS version.	• Checked • Unchecked (default)

system.transporter.agent.injection.skip.vc.redist:	When enabled, the system will not automatically install VC redistributable during Transporter/agent injection.	• Checked • Unchecked (default)
system.transporter.load.max.time.created.state.hours:	Specifies the timeout (hours) to wait for getting Transporter load request. Default is 3 hours.	• Default value: 3 • Minimum value: 1 • Maximum value: 72
system.transporter.modern.min.heap.size.megabyte:	Megabytes. 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. **Note**: Transporter restart is required to apply the setting.	• Default value: 512 • Minimum value: 256 • Maximum value: 65536
system.transporter.modern.max.heap.size.megabyte:	Megabytes. 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. **Notes**: -Xmx does not limit the total amount of memory that the JVM can use. Transporter restart is required to apply the setting.	• Default value: 3072 • Minimum value: 256 • Maximum value: 65536

system.transporter.modern.thread.stack.size.kilobyte:	Kilobytes. -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. **Note**: Transporter restart is required to apply the setting.	• Default value: 512 • Minimum value: 64 • Maximum value: 2048
system.transporter.modern.job.handler.max.thread.count:	Specifies the job thread count for modern Transporter. **Notes:** • 1 job thread equals ~200MB of memory, consider changing the related setting. • Transporter restart is required to apply the setting.	• Default value: 10 • Minimum value: 1 • Maximum value: 128
system.transporter.modern.service.handler.max.thread.count:	Specifies the service thread count for modern Transporter. **Note**: Transporter restart is required to apply the setting.	• Default value: 10 • Minimum value: 1 • Maximum value: 128
system.transporter.jvm.ram.requirement:	Bytes. For NASes only. Specifies the minimal ram required on NASes to create a SaaS repository.	• Default value: 0 • Minimum value: 0 • Maximum value: 10995 11627776
system.deleted.users.groups.remove.frequency:	Specifies the scheduled time for removing unnecessary deleted users, groups (in second).	• Default value: 86400 • Minimum value: 300

		• Maximum value: 1.797693134 8623157e+3 08
system.inventory.allow.duplicated:	Microsoft 365 and physical machines only. When enabled, the system allows duplicated discovery items.	• Unchecked (default) • Checked
system.inventory.optimize.discovery.time:	Microsoft 365 (SharePoint Online) only. When enabled, the system skips some attributes to optimize the discovery time.	• Unchecked (default) • Checked
system.0365.suppress.throttling.event:	Suppress throttling warning.	• Unchecked (default) • Checked
system.product.register.disable.periodic.data.collection:	When enabled, the product will not send data bundles every 30 days.	• Unchecked (default) • Checked
system.exchange.enable.direct.recovery:	When enabled, you can recover Exchange items without using a recovery server. For example, you can download items to the browser or forward them to a certain email. To do this, select **Download items or forward via email** on the **Destination** page of the job wizard and then select the appropriate recovery type on the **Options** page. Note that Google limits the total size of attachments within a message to 25 MB. Forwarding messages containing attachments that exceed this limit will fail.	• Unchecked (default) • Checked

system.repository.skip.periodic.refresh.on.transporter.busy.with.job:	When enabled and any Transporter repository is locked by a running job, the product skips periodic refresh for this Transporter repository.	• Unchecked (default) • Checked
system.job.skip.manual.transporter.data.path.validation	If set, transporter data path validation will be skipped for manually configured transporters. This setting is useful when ICMP (ping) cannot be used.	• Unchecked (default) • Checked
system.metadata.disable.ec2.instance.id.update	Disables EC2 instance ID detection on product startup. The detection is done via a HTTP request to http://169.254.169.254/latest/meta-data/instance-id This is required for proper product functioning in the AWS cloud.	• Unchecked (default) • Checked
system.transporter.allow.new	Allows using newer Transporter versions.	• Unchecked (default) • Checked
system.job.ict.skip.new.disk	If set, new disks added to the source item will not be added to the job automatically.	• Unchecked (default) • Checked

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**

Replication

With NAKIVO Backup & Replication, you can perform replication of virtual machines. Replication creates and maintains an identical copy of the source VM at the target location.

Refer to the following topics for more information:

- "Creating Hyper-V Replication Jobs" on page 497
- "Staging (Seeding) VM Replication" on page 526

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**.

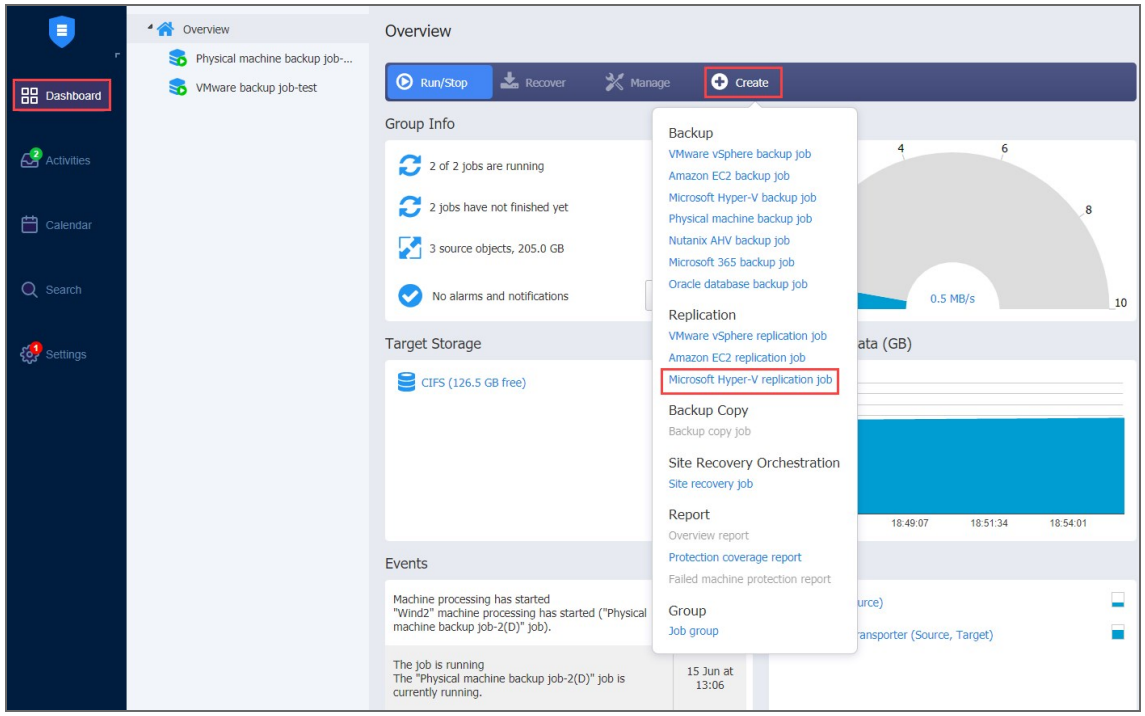

The **New Replication Job Wizard for Microsoft Hyper-V** opens. Complete the wizard to create a replication job.

- “Replication Job Wizard for Hyper-V: Source” on page 498
- “Replication Job Wizard for Hyper-V: Destination” on page 502
- “Replication Job Wizard for Hyper-V: Networks” on page 507
- “Replication Job Wizard for Hyper-V: Re-IP” on page 509
- “Replication Job Wizard for Hyper-V: Schedule” on page 511
- “Replication Job Wizard for Hyper-V: Retention” on page 515
- “Replication Job Wizard for Hyper-V: Options” on page 516

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 red “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.

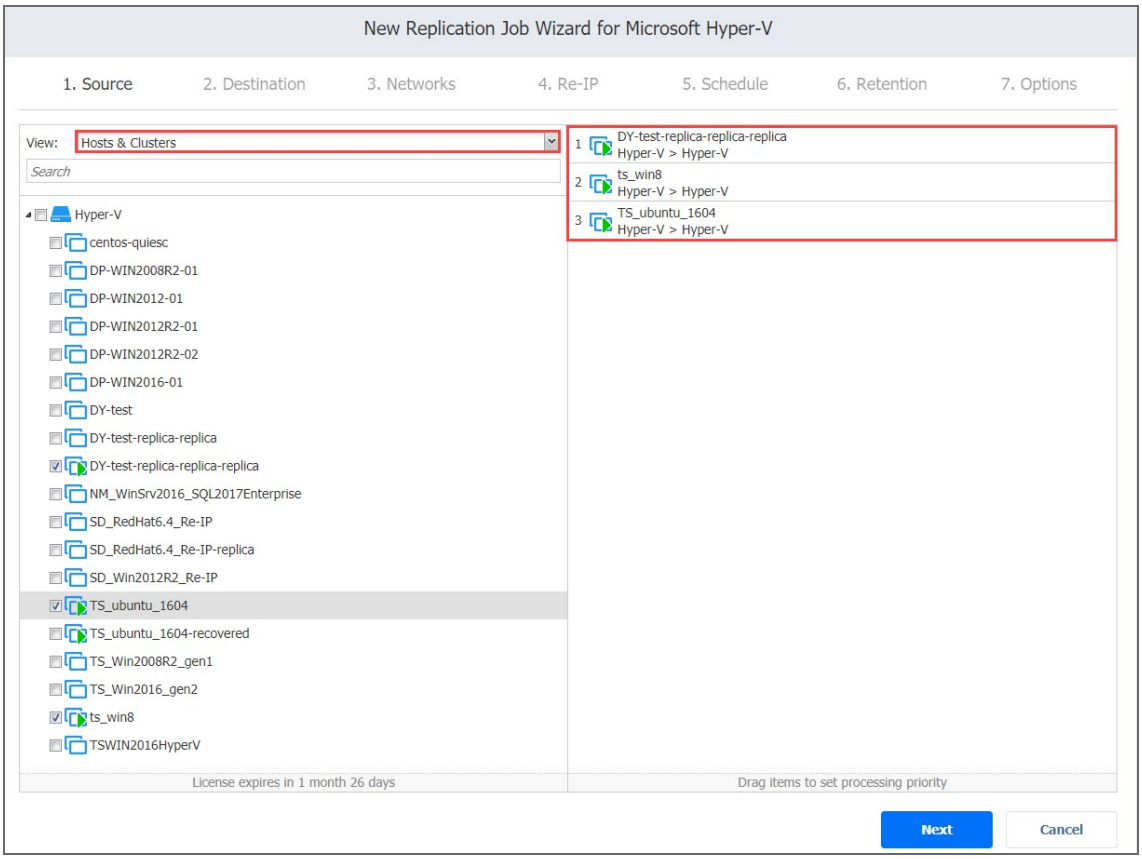

Jobs and Groups

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.

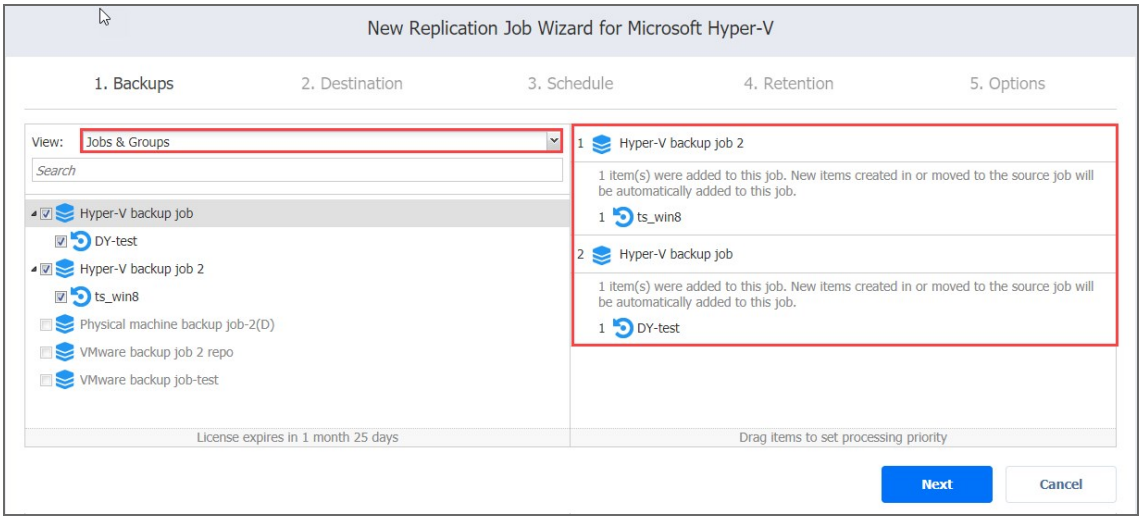

Backup Repositories

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.

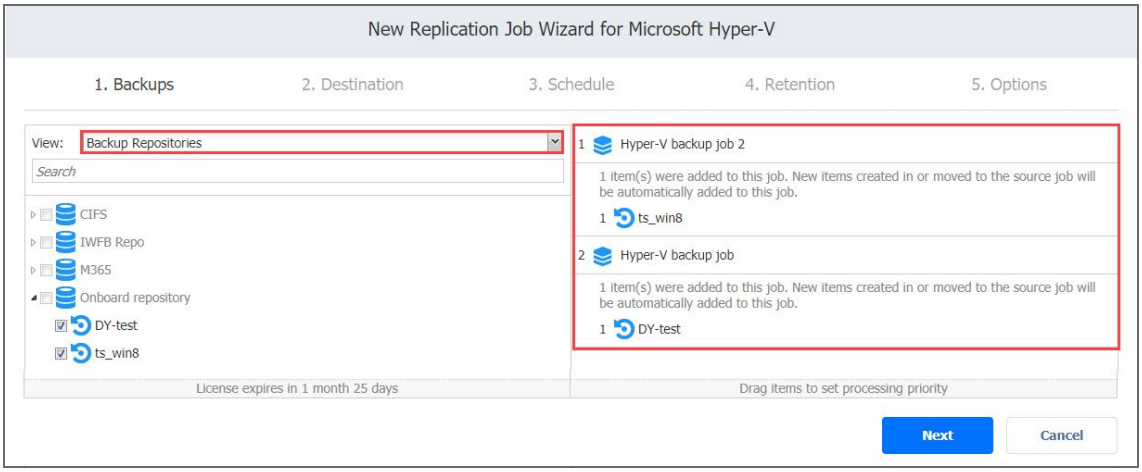

Policy

When the **Policy** view is selected, job policies can be used. Refer to "Managing Job Policies" on page 262 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 265 for details.

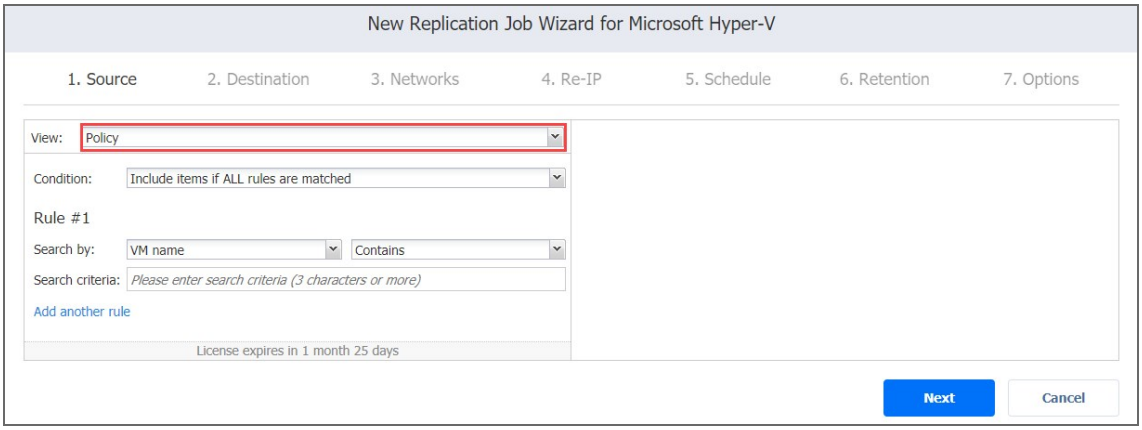

Notes

- 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 95 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 Hyper-V Servers" on page 332 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 create a Replication job from the 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 page 498.

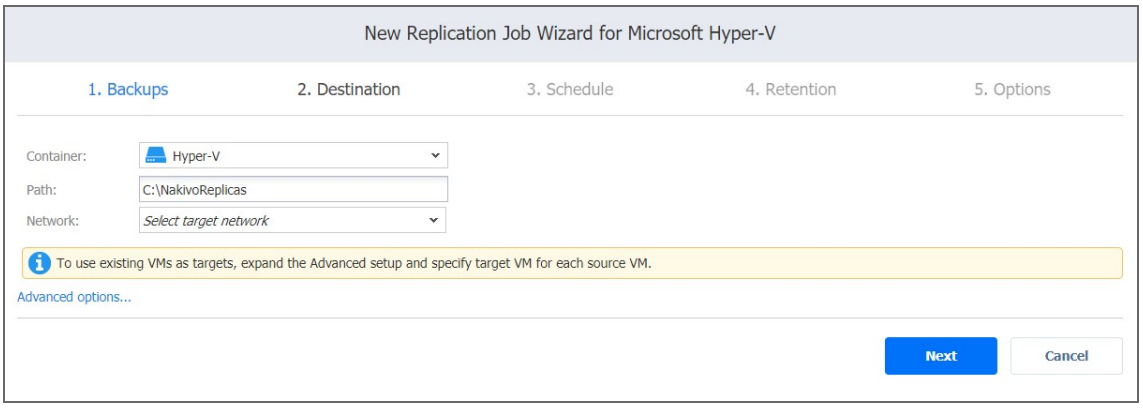

Setting the Default Destination for Replicas

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 options** 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 the backup job on the **Source** page, you can choose a **Default Network**.
4. Optionally, you can also choose a **Default VM folder.**

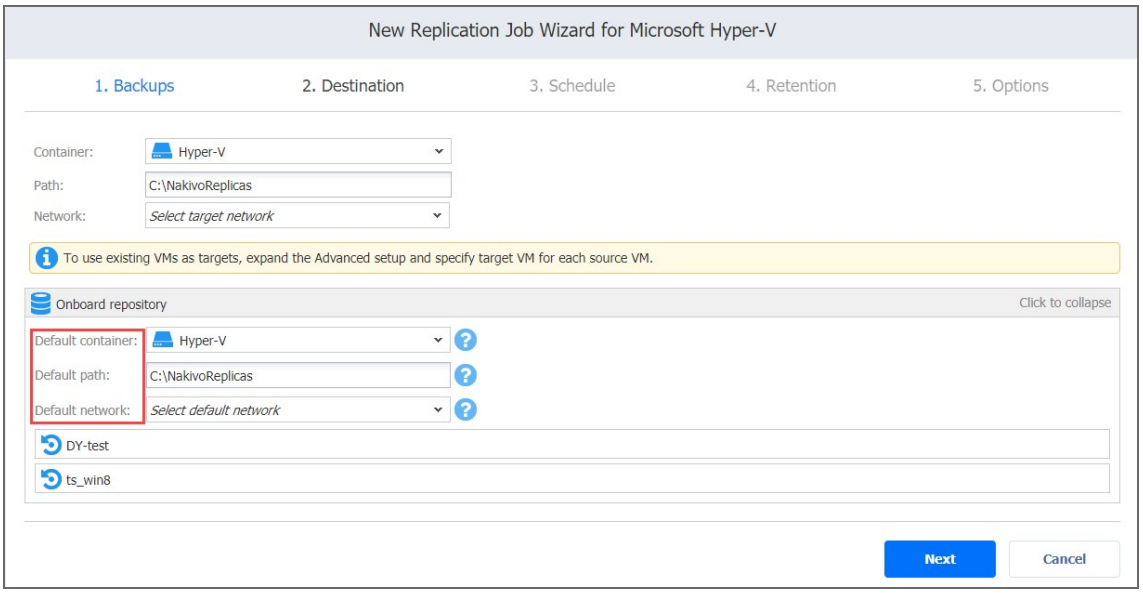

Setting Different Options for Replicas

To specify different replication options for VMs, follow the steps below:

1. Click **Advanced options**.
2. Set a target server, path, and network for each VM.

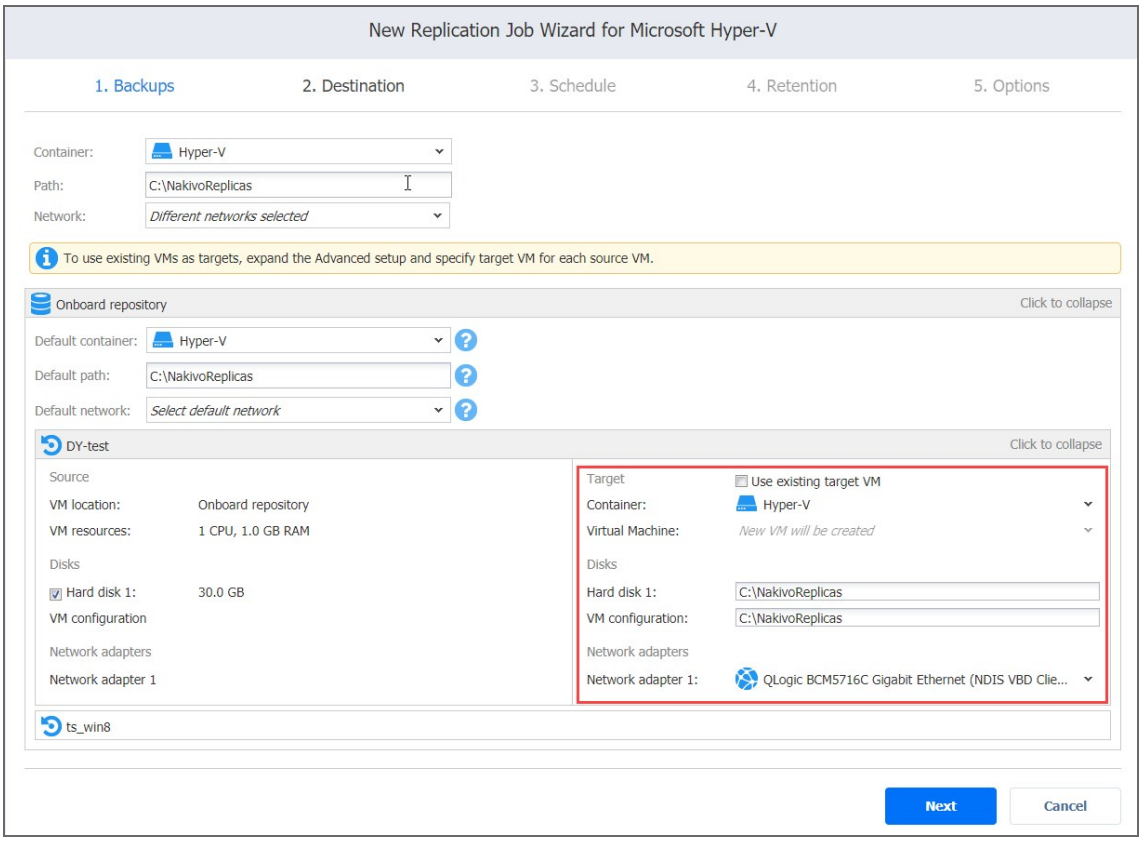

Mapping Source VMs to Existing Replicas

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 options**.
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.

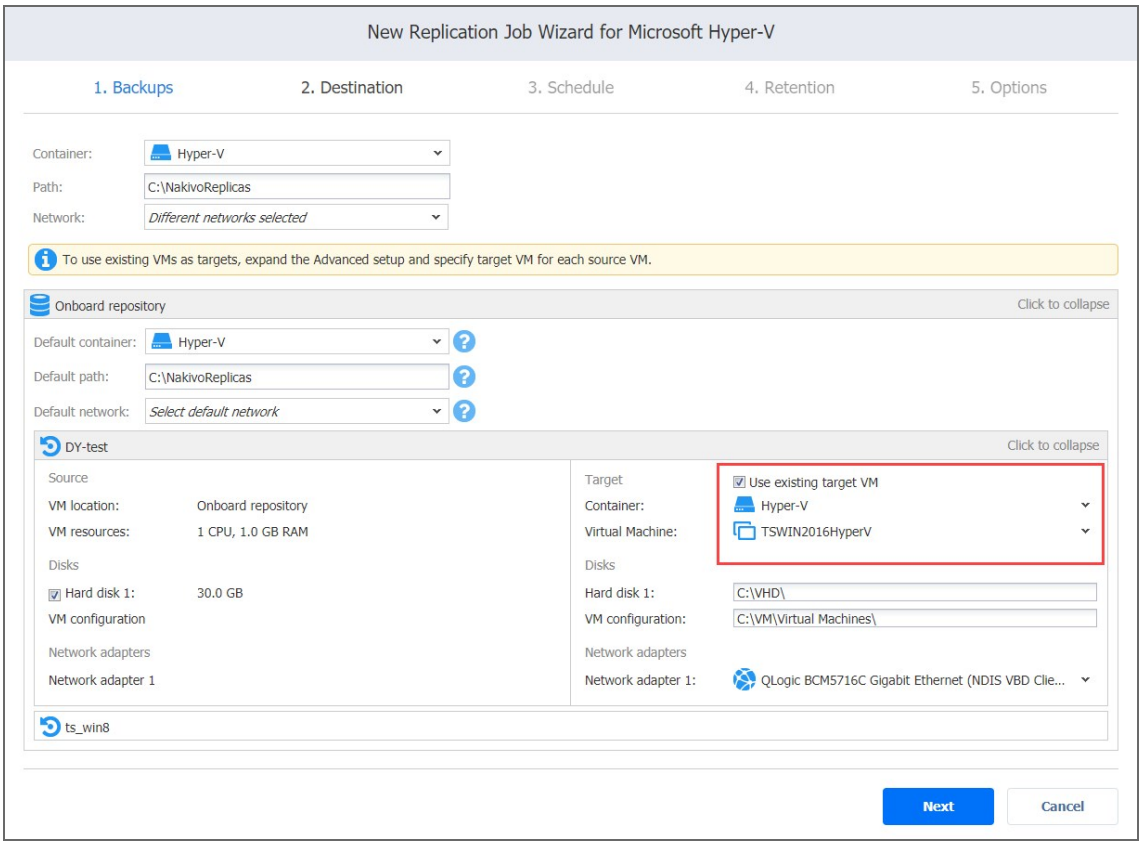

When you run the job, the product analyzes the target VM you have selected, determines how it is different from the source VM, and transfers only the differential data. VM replication mapping can be a time-consuming process that can be equal to the time required to create a full VM replication. After 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.

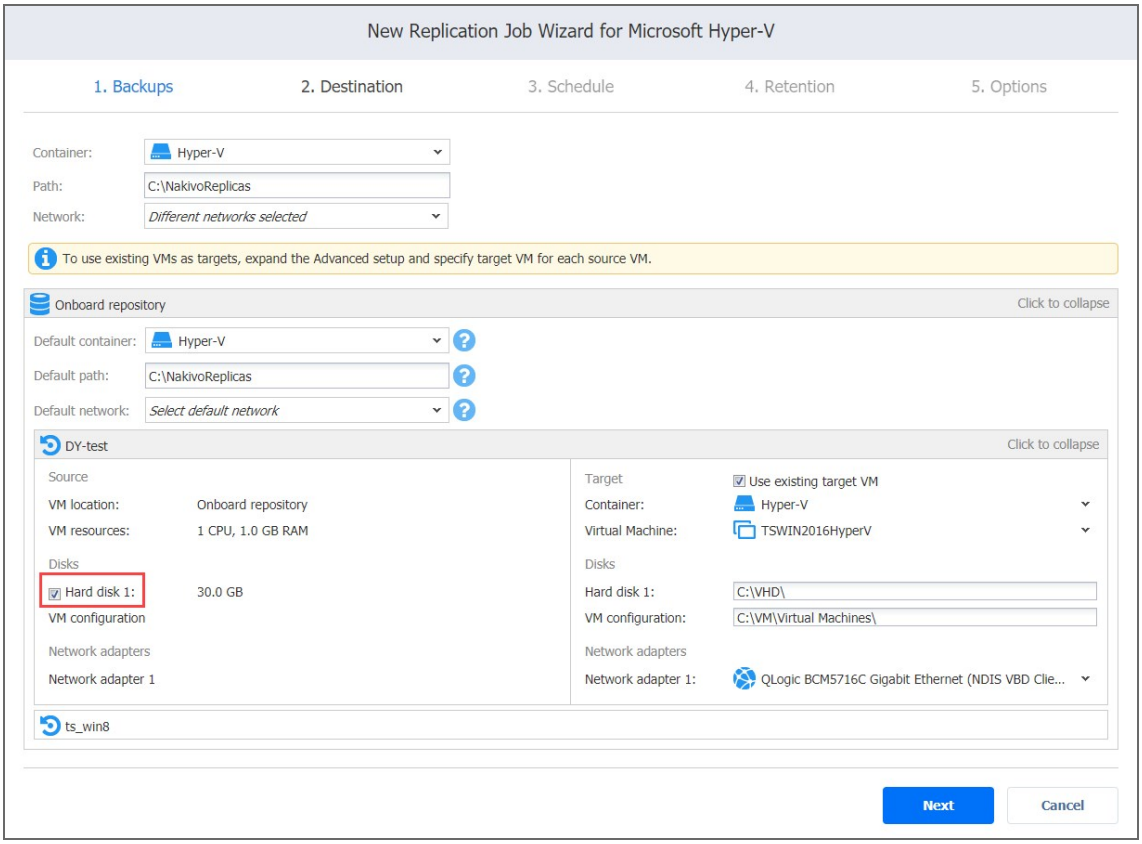
New Replication Job Wizard for Microsoft Hyper-V
1. Backups
2. Destination
3. Schedule
4. Retention
5. Options
Container:
Hyper-V
Path:
C:\NakivoReplicas
Network:
Different networks selected
To use existing VMs as targets, expand the Advanced setup and specify target VM for each source VM.
Onboard repository
Click to collapse
Default container:
Hyper-V
Default path:
C:\NakivoReplicas
Default network:
Select default network
DY-test
Click to collapse
Source
VM location:
Onboard repository
VM resources:
1 CPU, 1.0 GB RAM
Disks
Hard disk 1:
30.0 GB
VM configuration
Network adapters
Network adapter 1
Target
Use existing target VM
Container:
Hyper-V
Virtual Machine:
TSWIN2016HyperV
Disks
Hard disk 1:
C:\VHD\
VM configuration:
C:\VM\Virtual Machines\
Network adapters
Network adapter 1:
QLogic BCM5716C Gigabit Ethernet (NDIS VBD Clie...
ts_win8
Next
Cancel

Replication Job Wizard for Hyper-V: Networks

Note

The **Networks** page is skipped if you have chosen an existing backup as the target 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**.

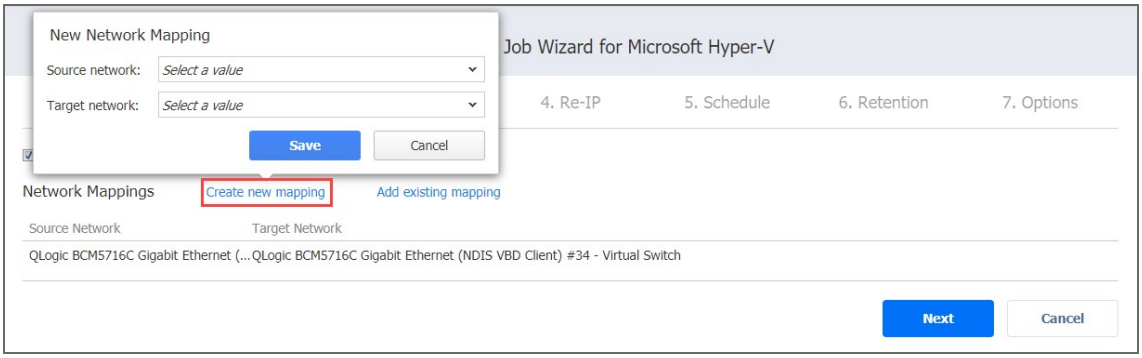

   - Add an existing mapping:
     a. Click **Add existing mapping**.
     b. The **Network Mappings** dialog box opens. Choose one or more appropriate network

mappings and close the dialog box.

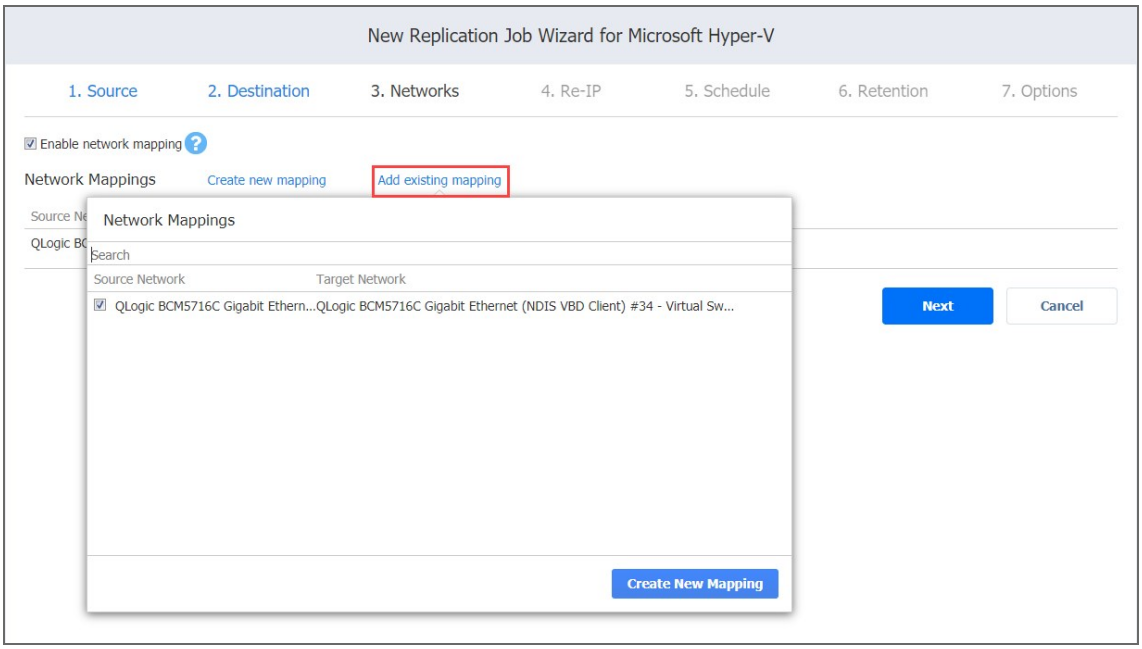

- Edit an existing mapping:
  a. Hover the pointer over the required item in the **Network Mappings** list and click the **Edit** button to the right of the item.
  b. The **Edit Network Mapping** dialog box opens. Choose the appropriate item from the **Target network** drop-down list and click **Save**.

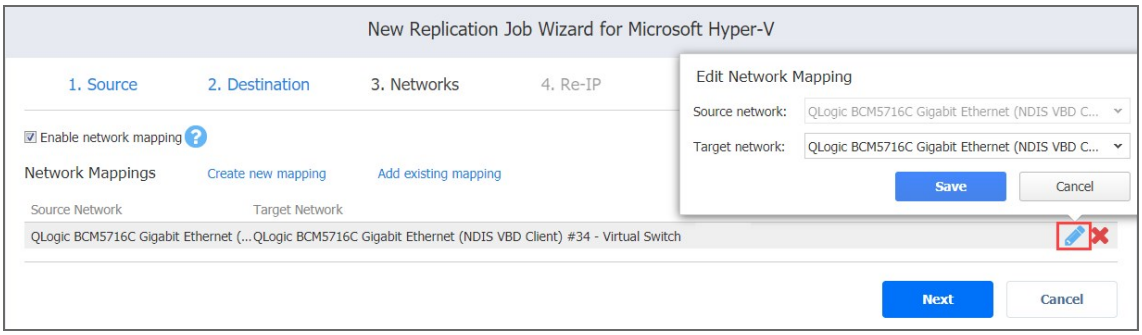

- Delete an existing mapping: Hover the pointer over the required item in the **Network Mappings** list and click the delete icon to the right of the item.

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 700 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**.

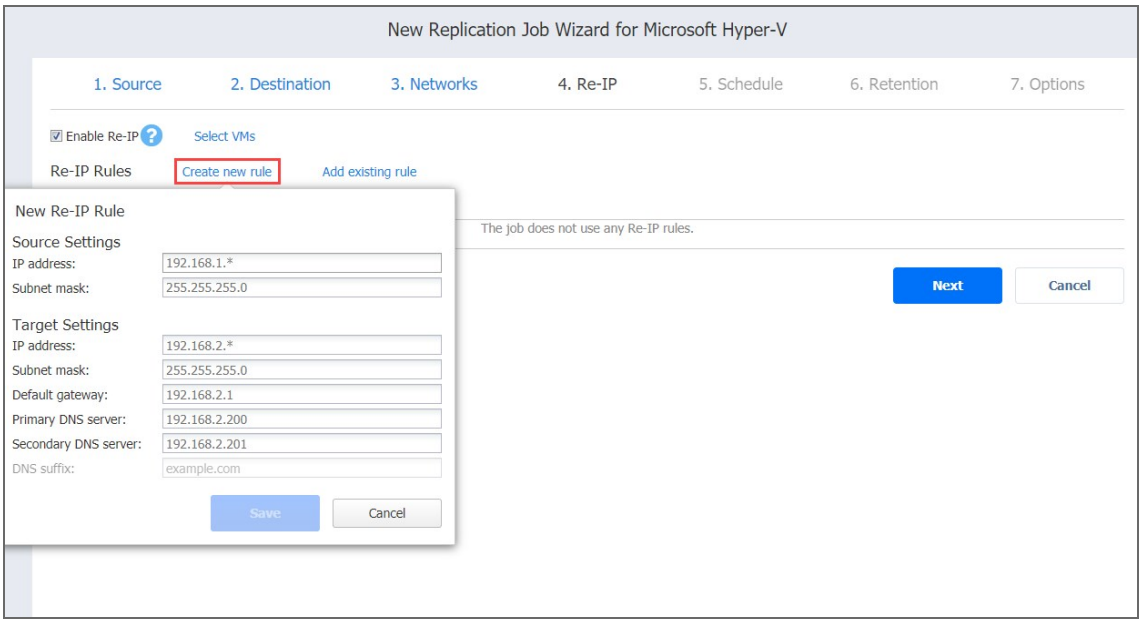

Note

You can use wildcards for IP addresses. Refer to the "Failover Job Wizard for Microsoft Hyper-V: Re-IP" on page 700 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.

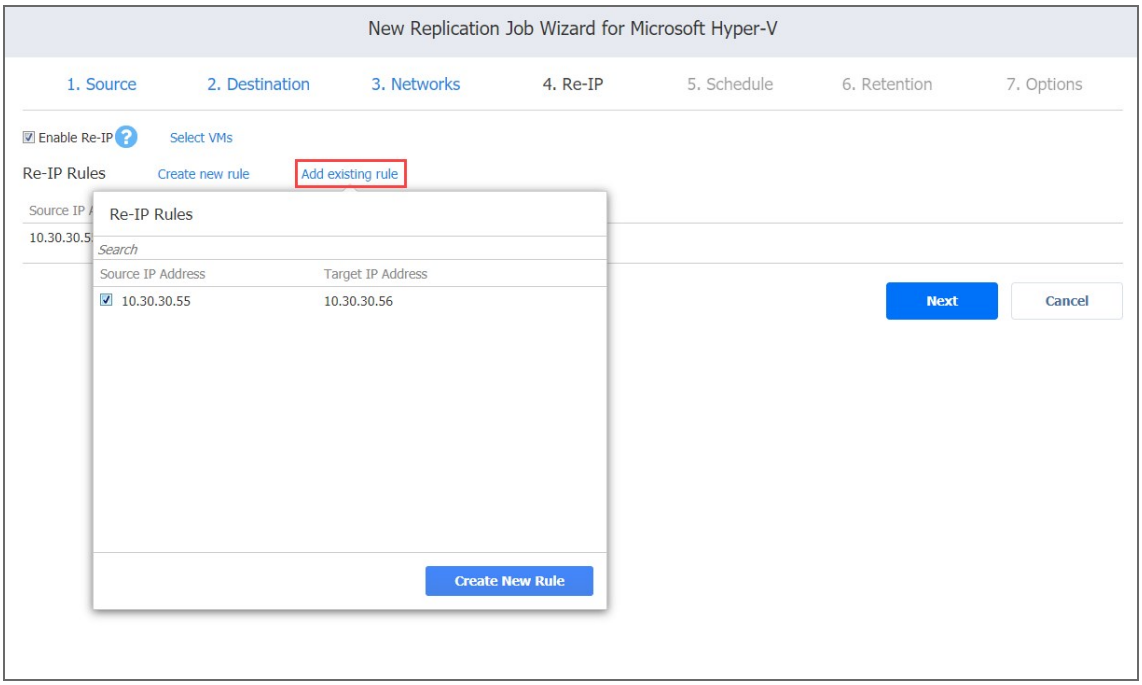

- Edit an existing Re-IP rule:

  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**.

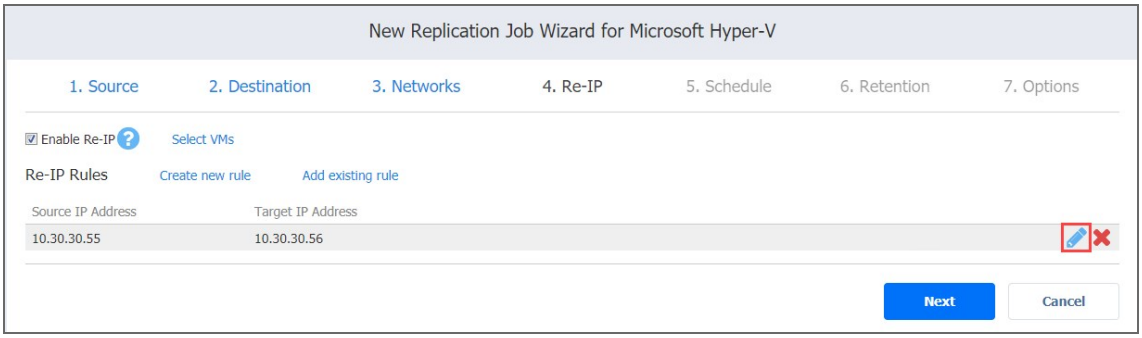

- Delete an existing mapping: Hover the pointer over the required item in the **Re-IP Rules** list and click the delete icon to the right of the item.

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.

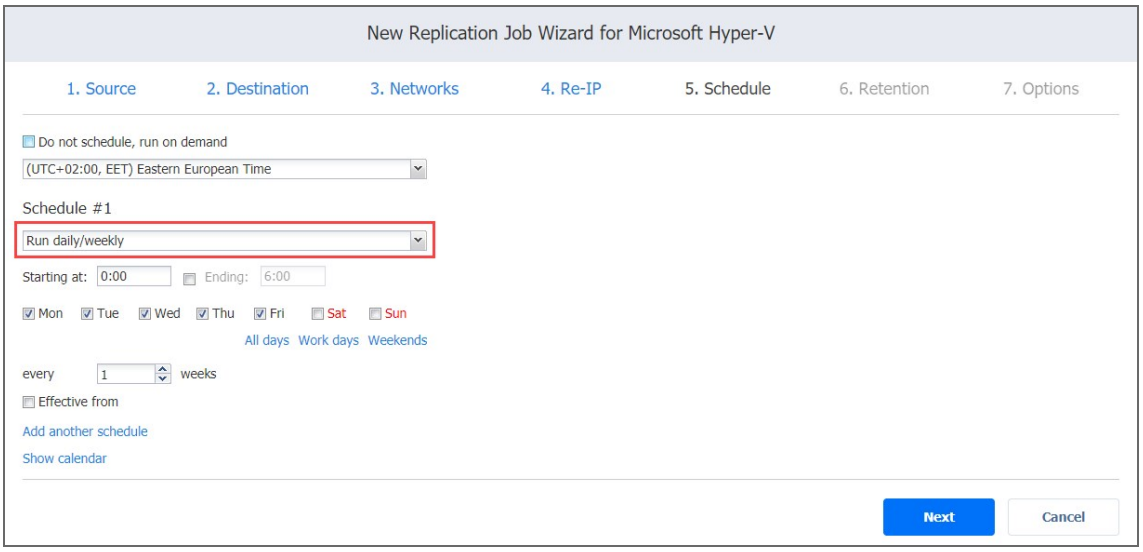

Monthly or Yearly Replication

To run the job monthly or yearly, choose **Run monthly/yearly** from the schedule drop-down list and do the following:

- Specify the job start schedule in the appropriate boxes.
- Specify the day and month when the job should be started in the **Run every** boxes.
- Specify the time when the job should be started in the **Starting at** box.
- Specify the end time for the job in the **Ending** box. If the job has not completed by the time specified, the job will be stopped.
- If necessary, select the **Effective from** checkbox and pick the date when the schedule comes into effect.

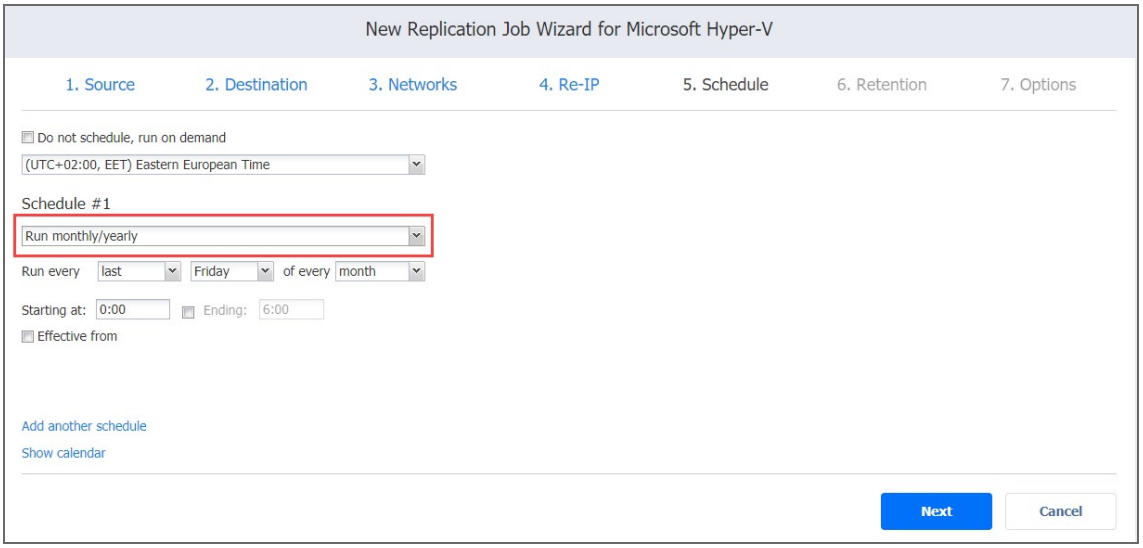

Periodic VM Replication

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.

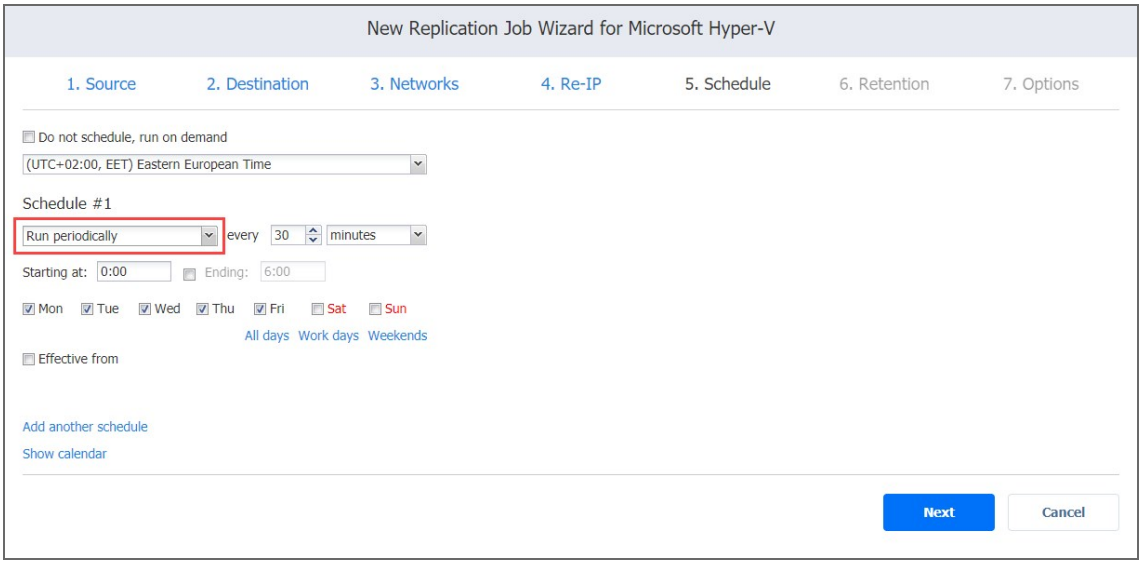

Chained Job

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.

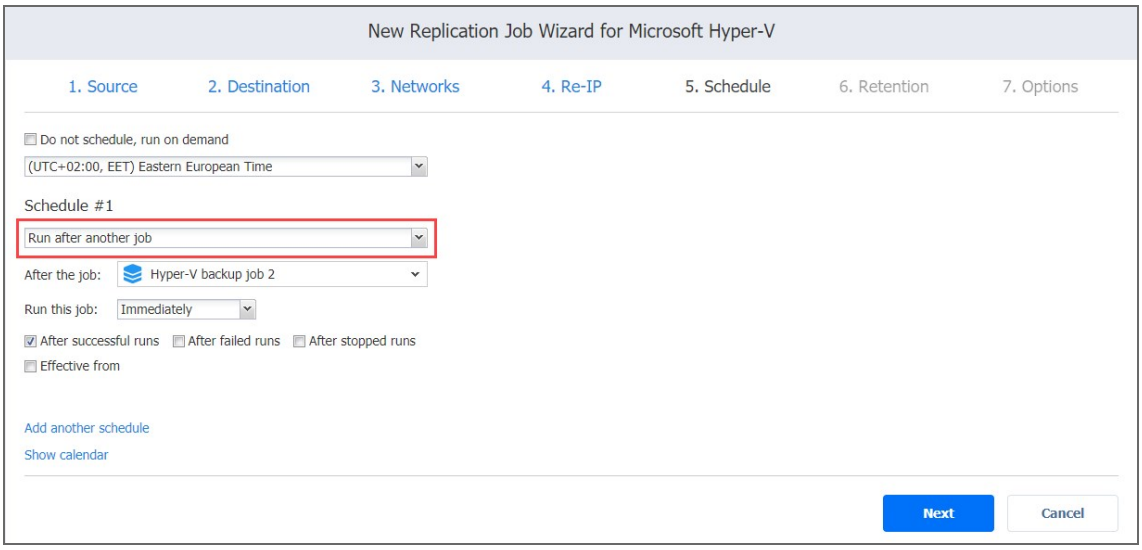

Additional Schedule

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.

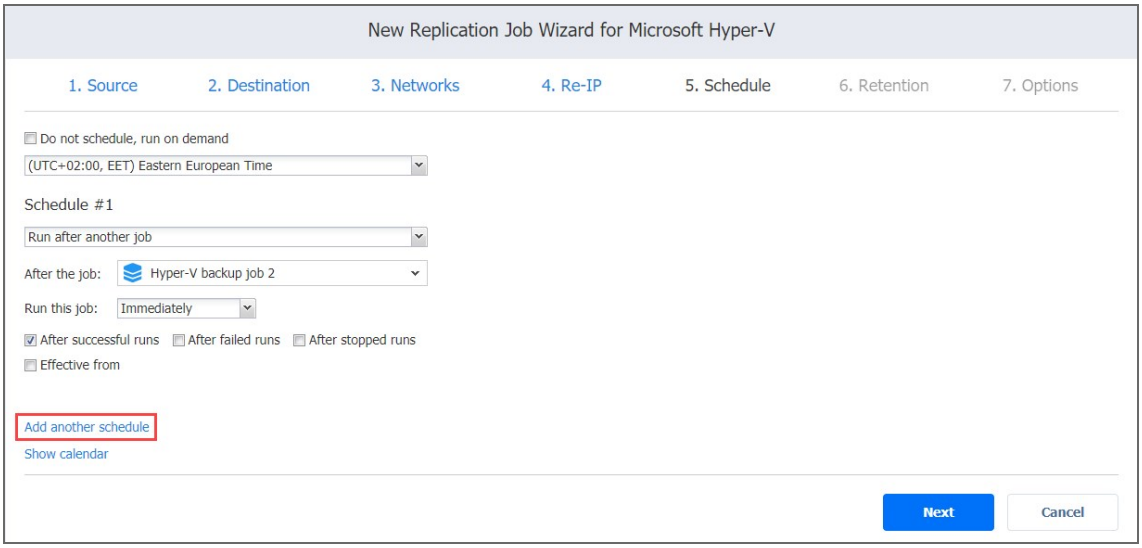

Replication Job Wizard for Hyper-V: Retention

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.



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 520
- "Pre and Post Actions" on page 521
- "Data Transfer" on page 523
- "Completing the New Replication Job Wizard for Microsoft Hyper-V" on page 525

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, skipping swap files, partitions and unused blocks.

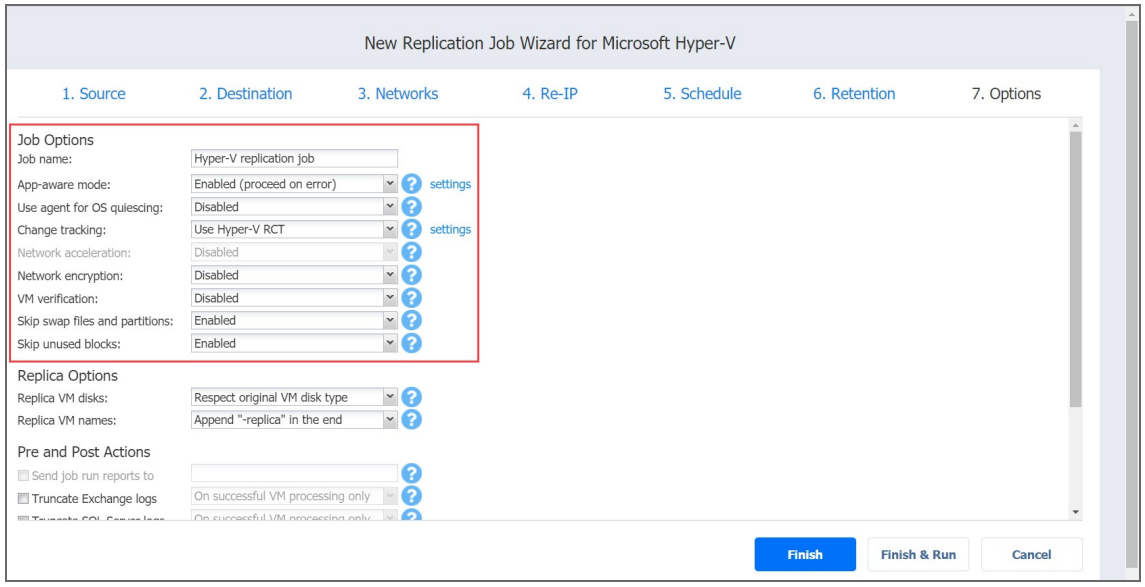

Job Name

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.

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 "Email Notifications" on page 287.) 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.

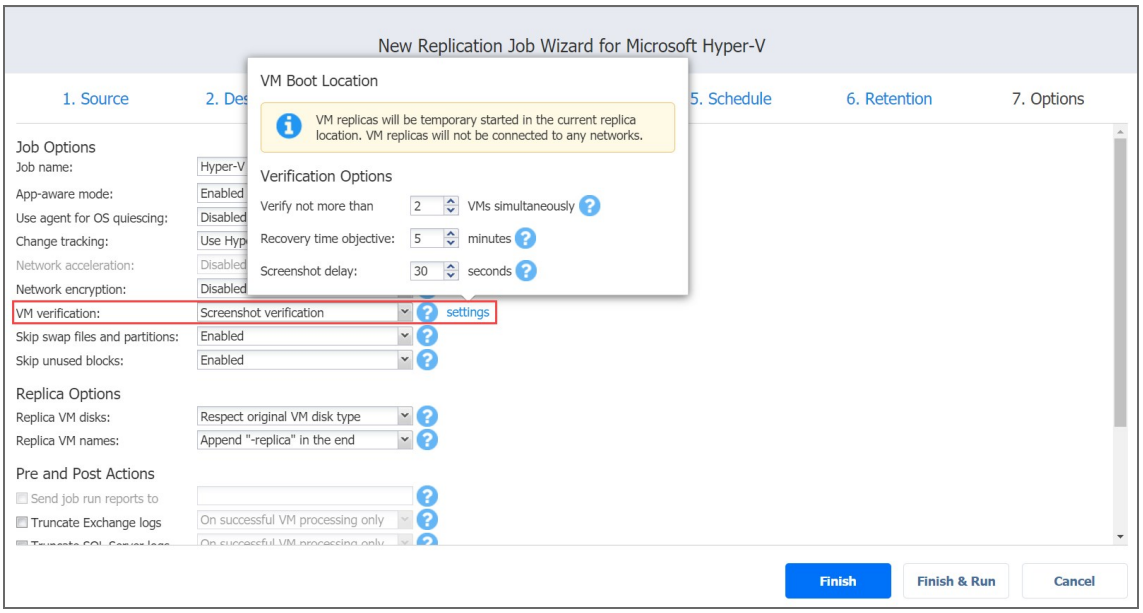

After choosing **Boot verification**, set verification options in the dialog box that opens:

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.

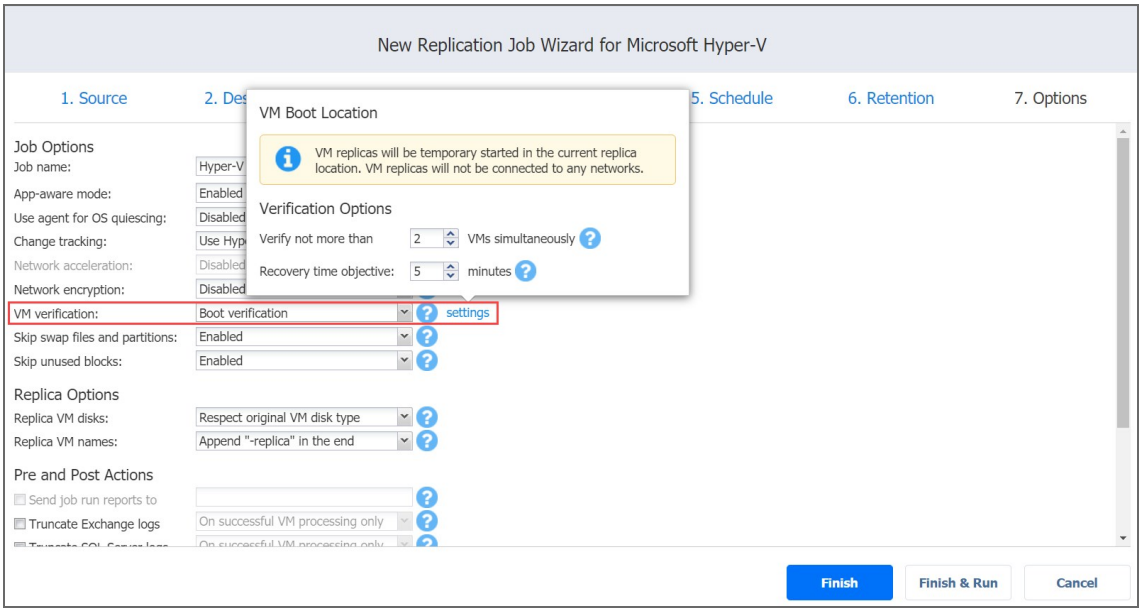

Skip Swap Files and Partitions

With this option enabled, NAKIVO Backup & Replication automatically skips swap files and partitions during the backup process.

Note

The feature is not available for replication from backup jobs.

Skip Unused Blocks

With this option enabled, NAKIVO Backup & Replication automatically skips 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.

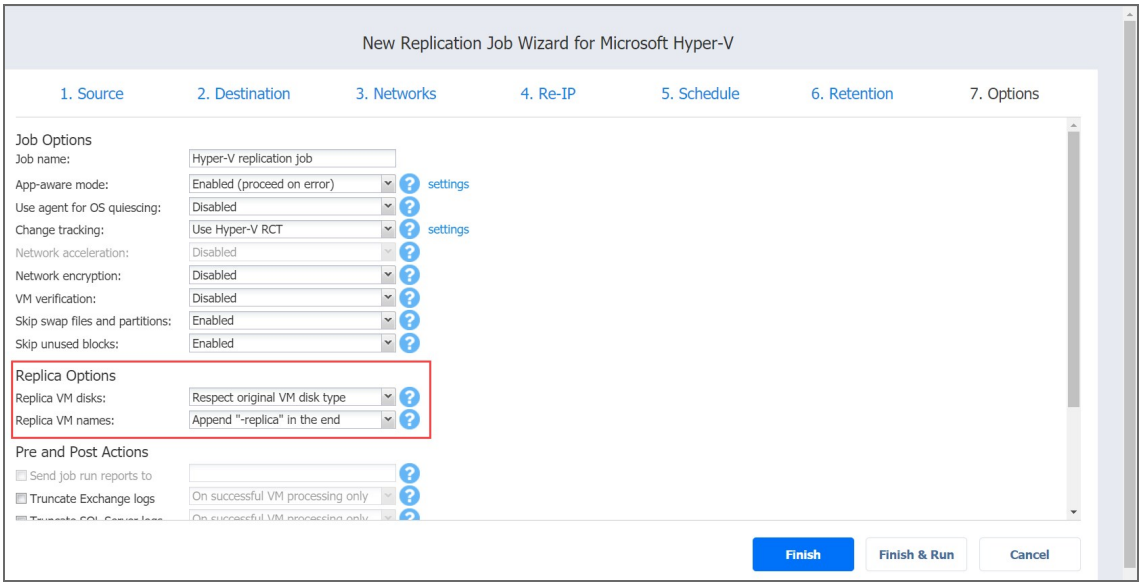

Replica VM Disks

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.

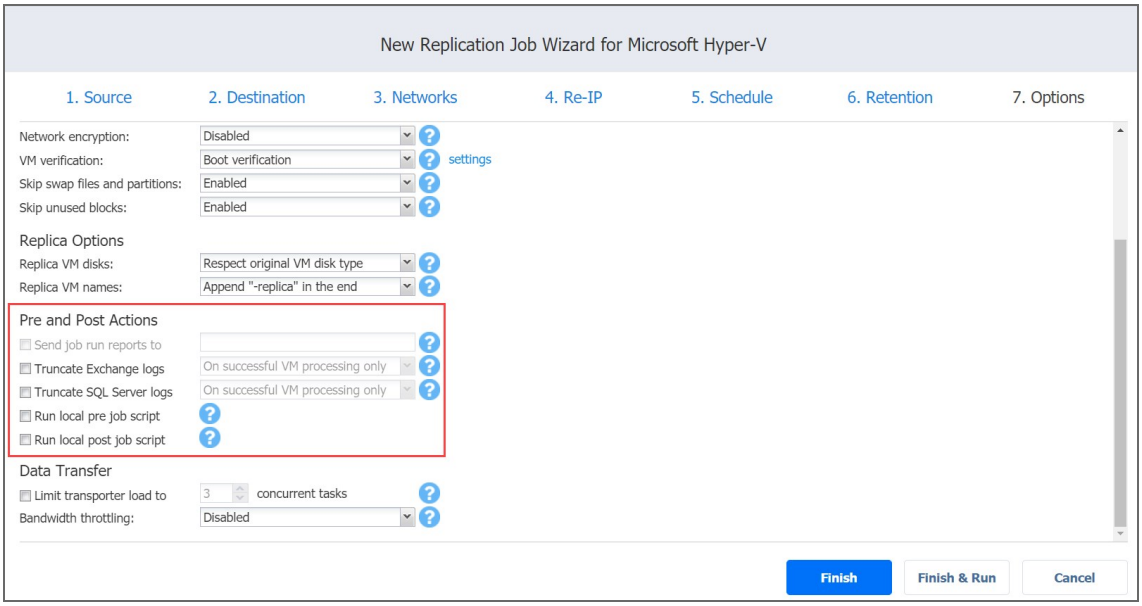

Data Transfer

In the *Data Transfer* section on the **Options** page, you can set the transporter load and configure bandwidth throttling.

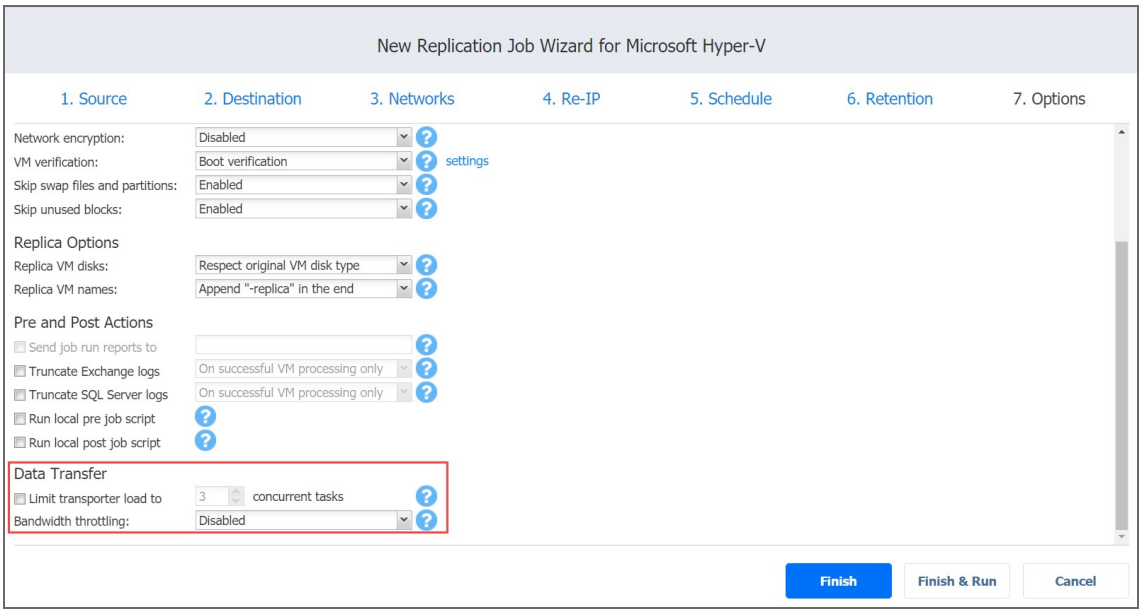

Transporter Load

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 281 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 281 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.

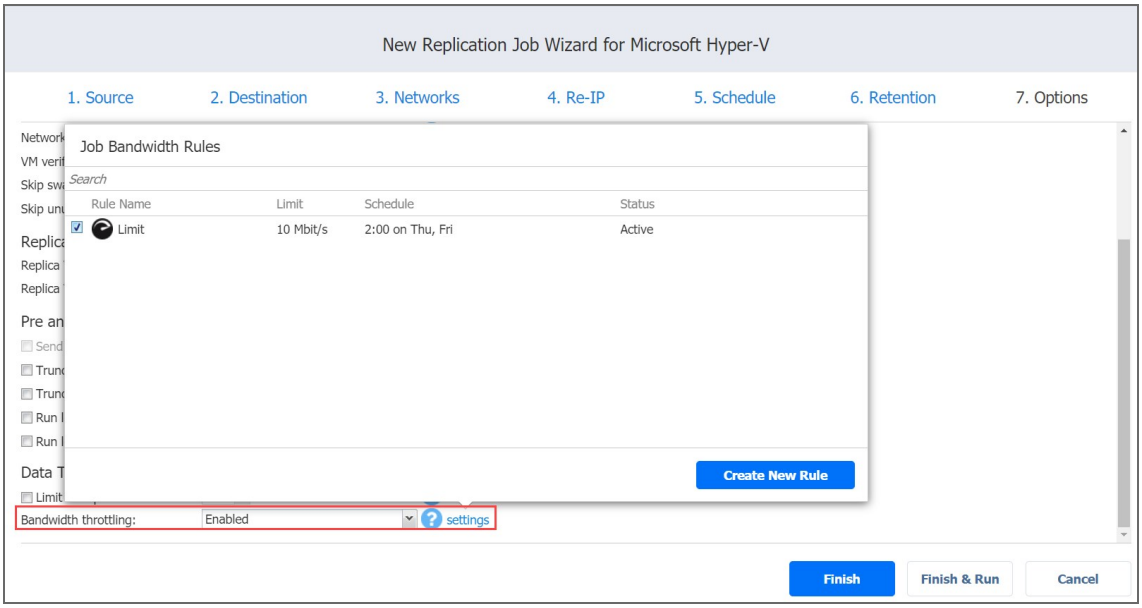

Completing the New Replication Job Wizard for Microsoft Hyper-V

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 250 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.

Backup

This section contains the following topics:

- "Creating Hyper-V Backup Jobs" on page 528
- "Creating Backup Copy Jobs" on page 552
- "Backing Up to Tape" on page 574
- "Staging (Seeding) Initial Backup" on page 596
- "Deleting Backups" on page 591

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**.

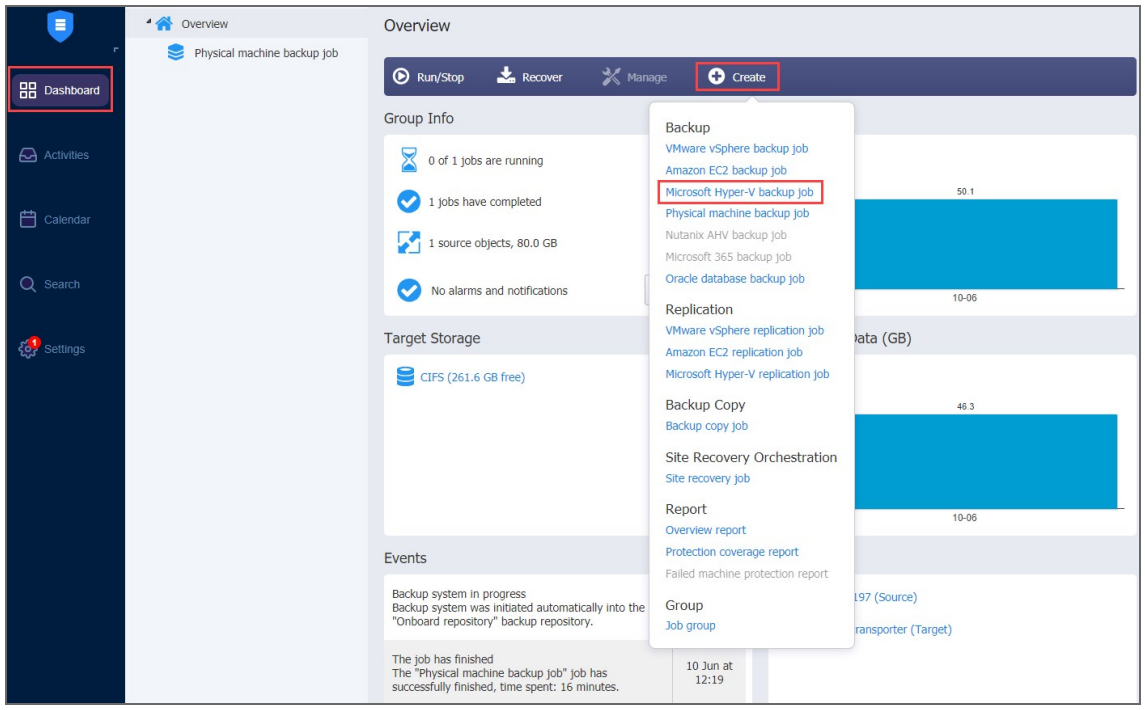

The **New Backup Job Wizard for Microsoft Hyper-V** opens. Complete the wizard as described in the sections below:

- “Backup Job Wizard for Hyper-V: Source” on page 529
- “Backup Job Wizard for Hyper-V: Destination” on page 533
- “Backup Job Wizard for Hyper-V: Schedule” on page 536
- “Backup Job Wizard for Hyper-V: Retention” on page 540
- “Backup Job Wizard for Hyper-V: Options” on page 542

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:

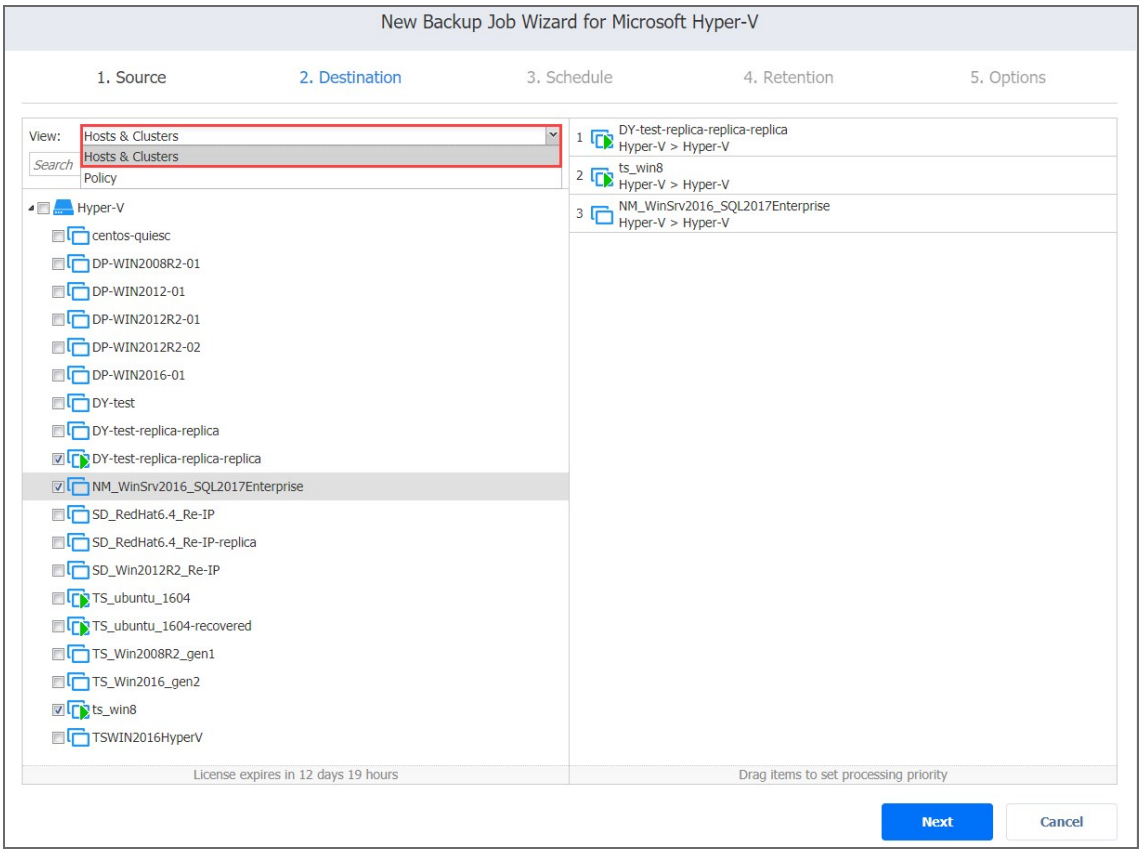

     a. Optionally, filter the inventory tree by entering a string to the **Search** box. You can enter a part or the whole item name.
     b. Select Hyper-V items by selecting the checkbox next to the item.
     c. The selected items appear in the right pane of the page. If necessary, reorder the selected 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.

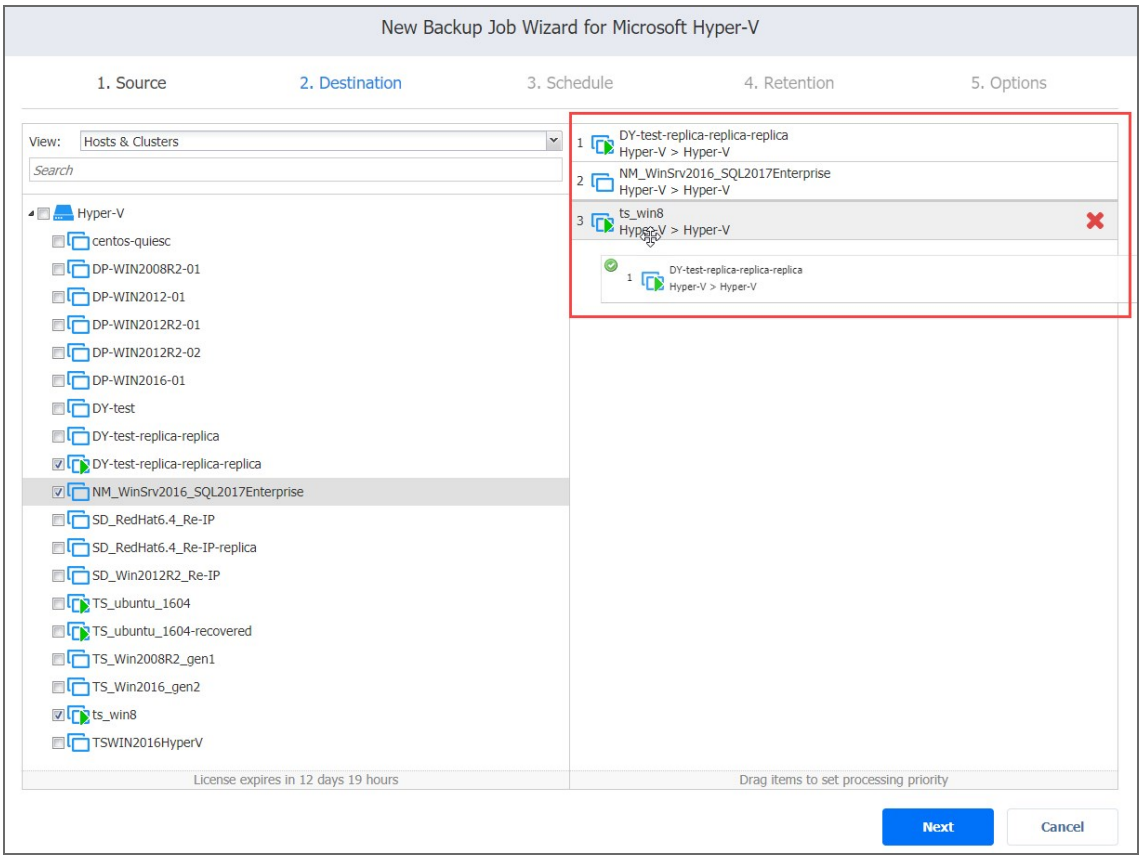

d. Review the list of selected Hyper-V items. If needed, remove a selected VM or a container from the backup job in either of the following ways:

- 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.

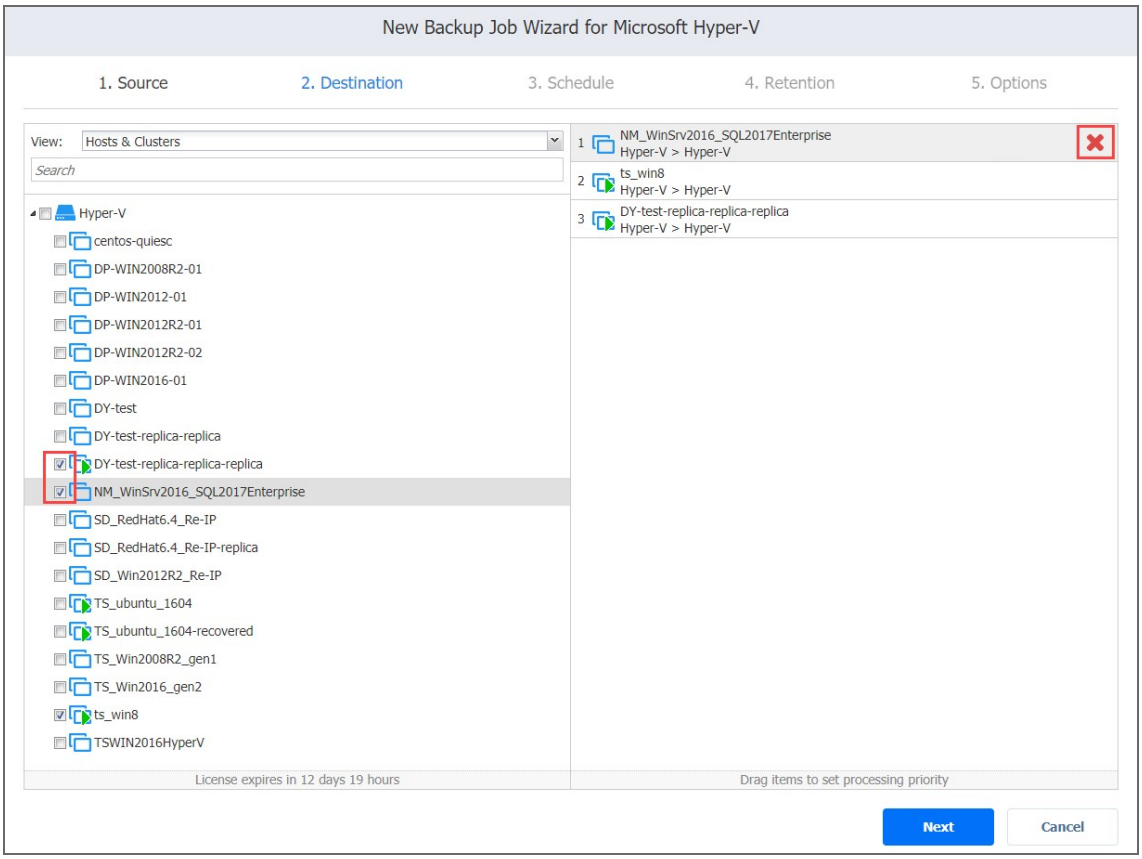

- **Policy**: If selected, this allows you to use job policies; for details, refer to "Managing Job Policies" on page 262. Please follow the steps below:
  a. If items are selected in the alternate view, a dialog opens warning you that switching to the **Policy** view will reset your current selection. 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 265 .

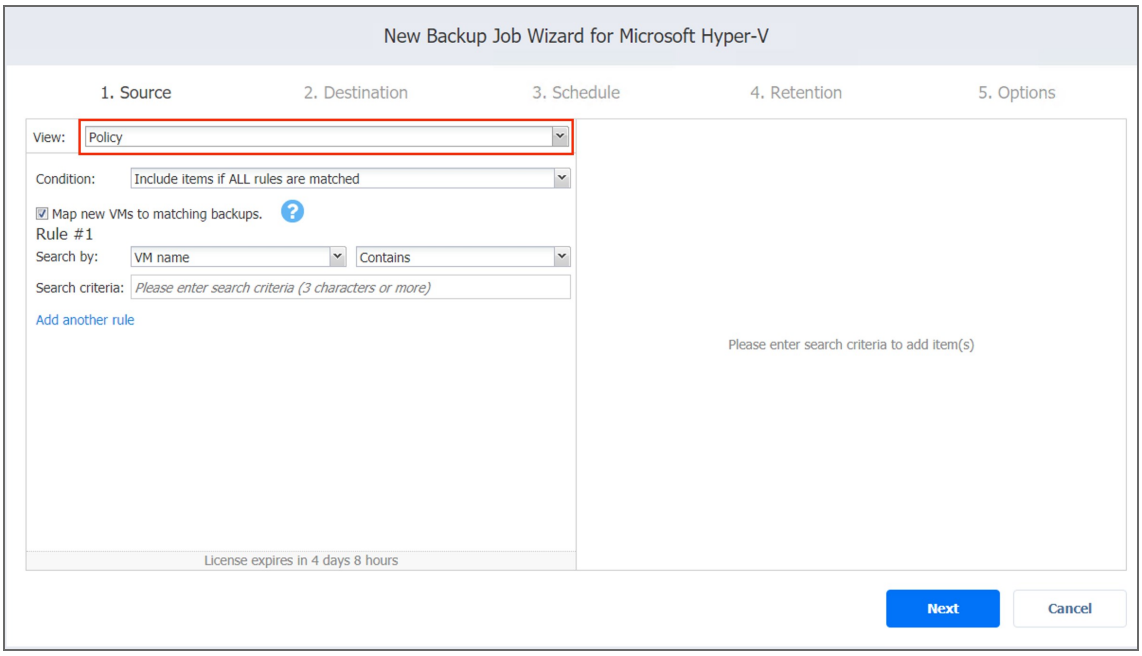

2. Click **Next** to confirm that you wish to add the selected VMs to the backup job.

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.
- 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.

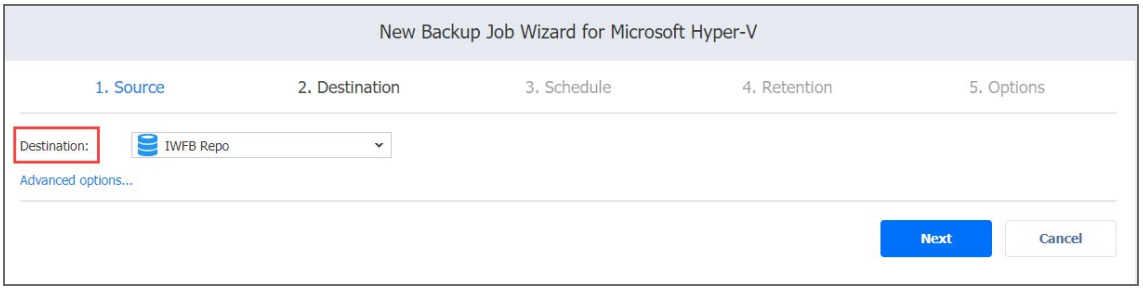

Setting Different Backup Repositories for VMs

To back up the selected VMs to different Backup Repositories, follow the steps below:

1. Click **Advanced options**.
2. In the VM boxes, choose a Backup Repository for each VM individually.

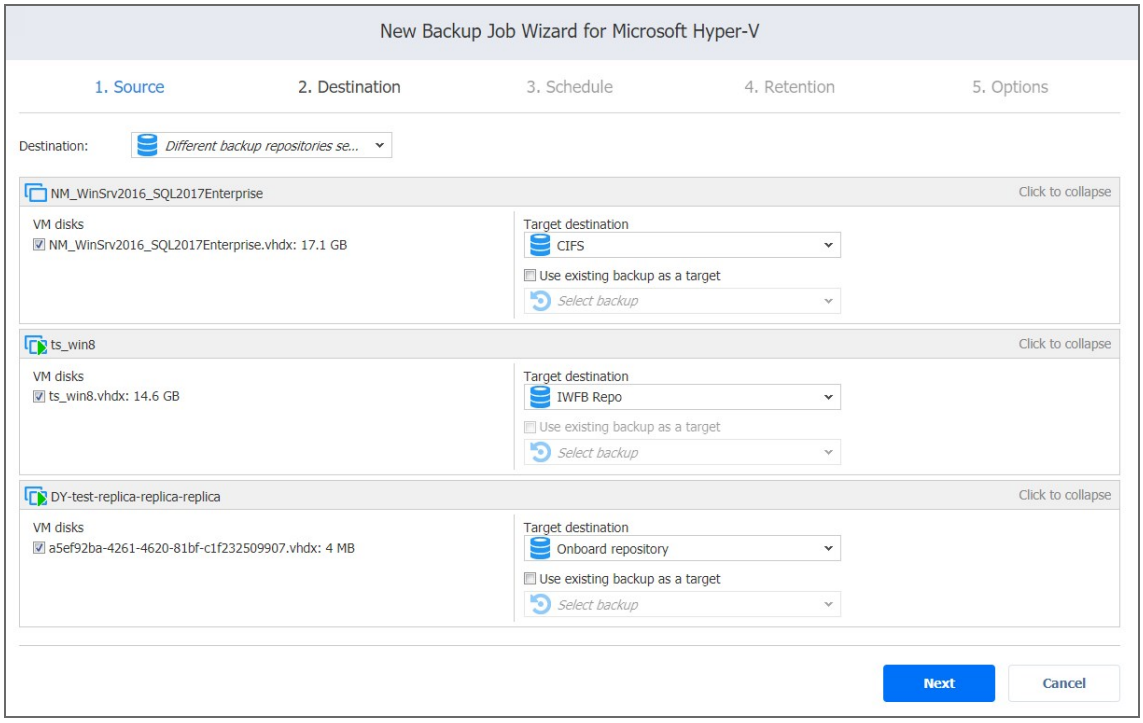

Mapping Source VMs to Existing Backups

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 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 options**.
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.

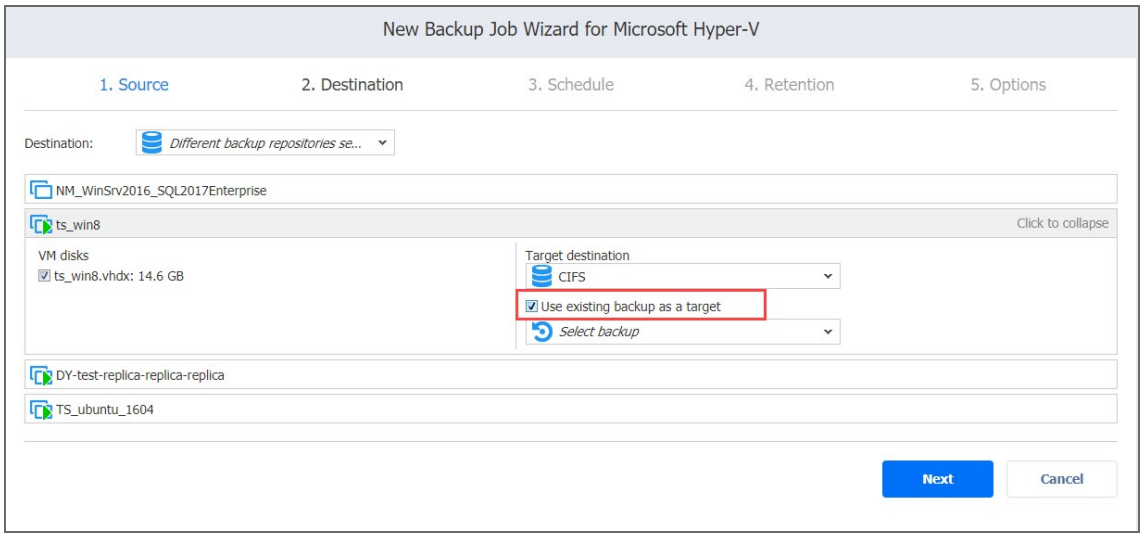

When you run the job, the product will analyze the target VM you have selected, determine how it is different from the source VM, and transfer only the differential data. VM backup mapping can be a time-consuming process that can be equal to the time required to create a full VM backup. After the job 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 options**.
2. Cancel the selection of the VM disks that you would not like to back up.

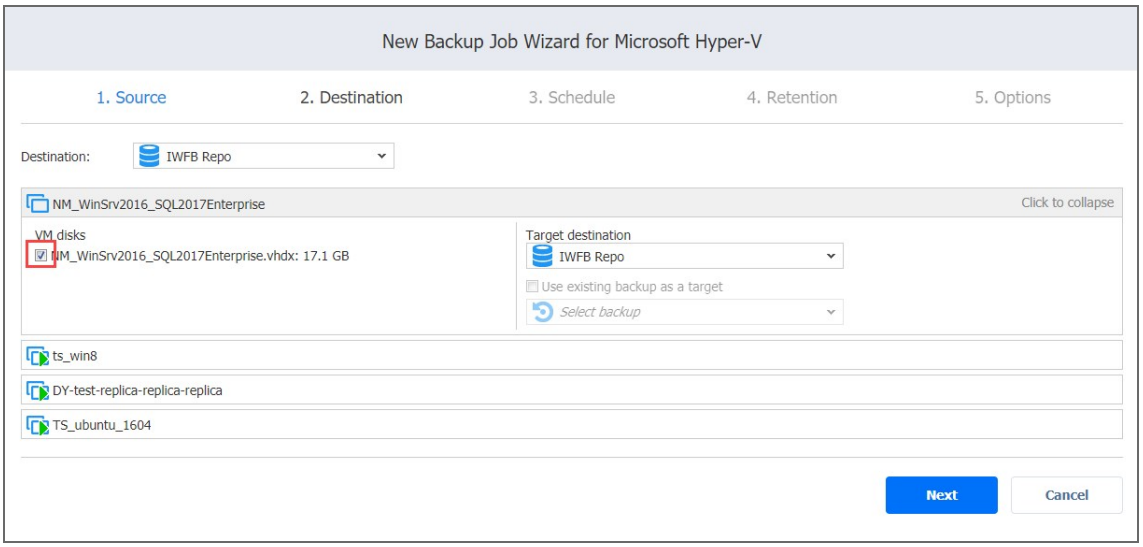

Backup Job Wizard for Hyper-V: Schedule

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.

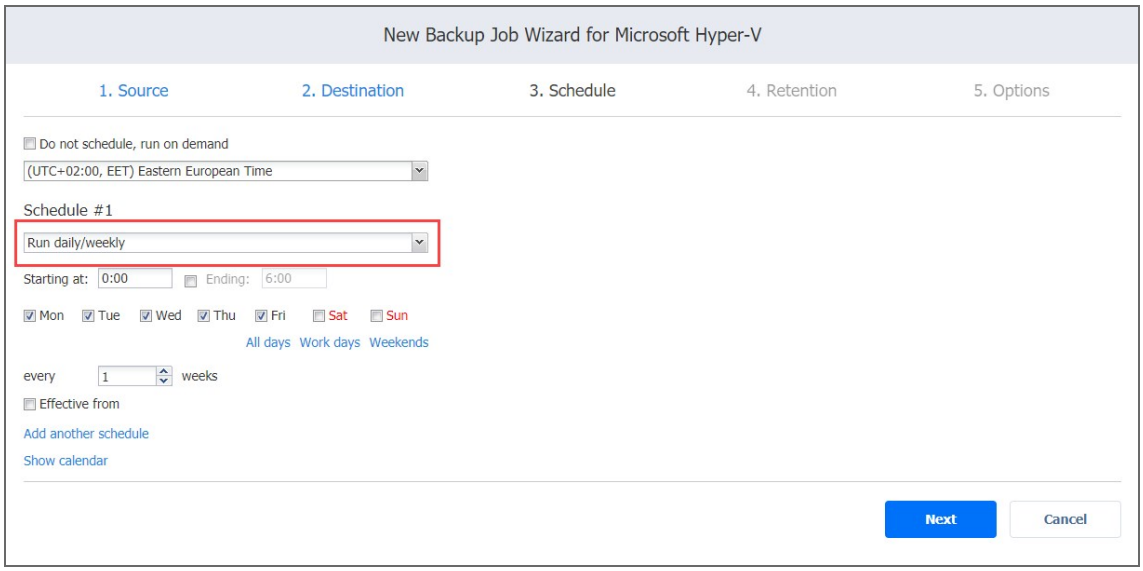

Monthly or Yearly Backup

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.

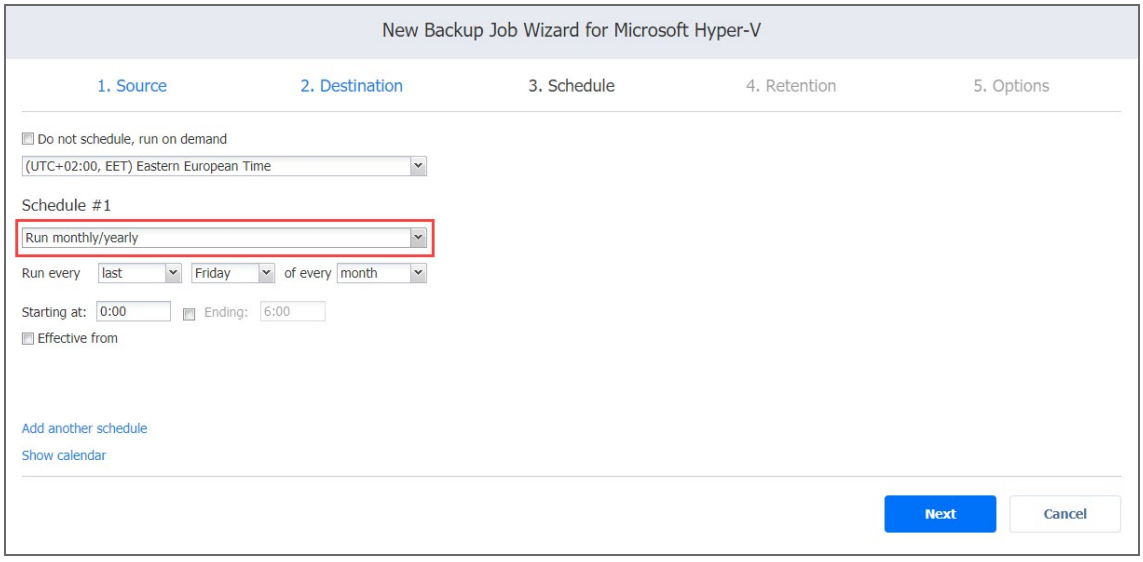

Periodic Backup

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.

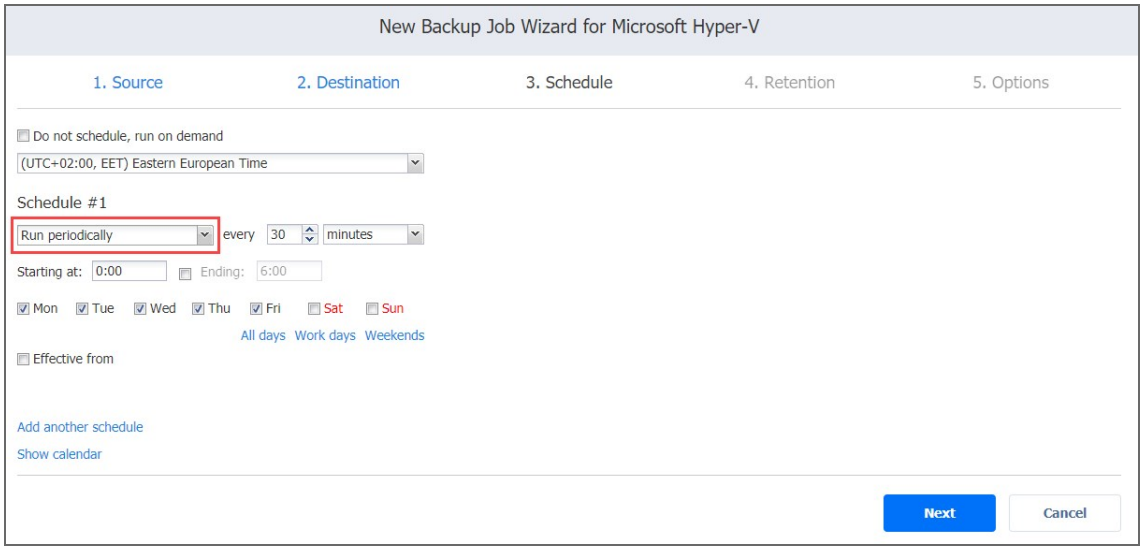

Chained Job

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.

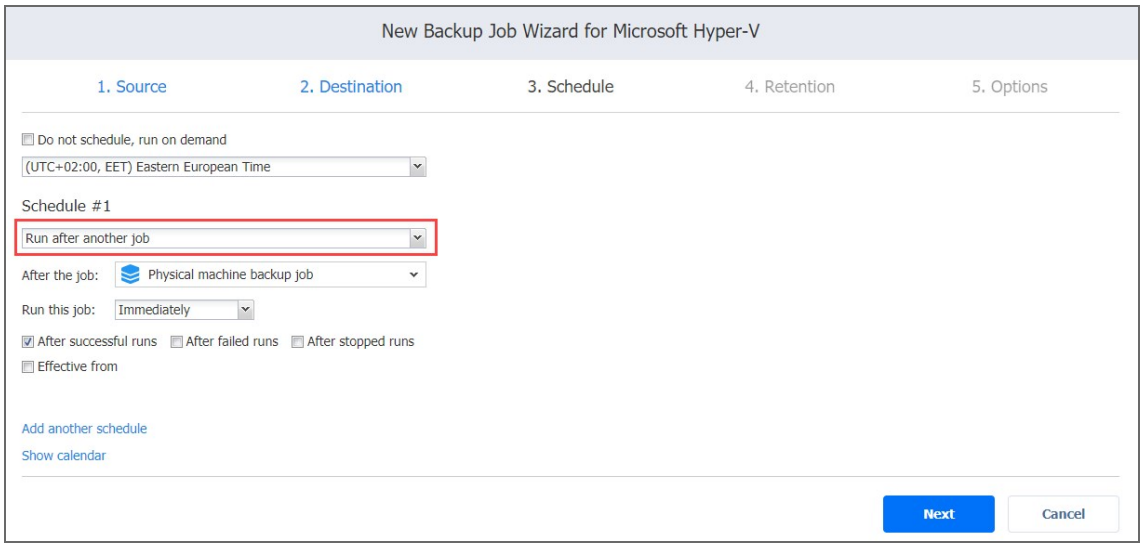

Additional Schedule

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.

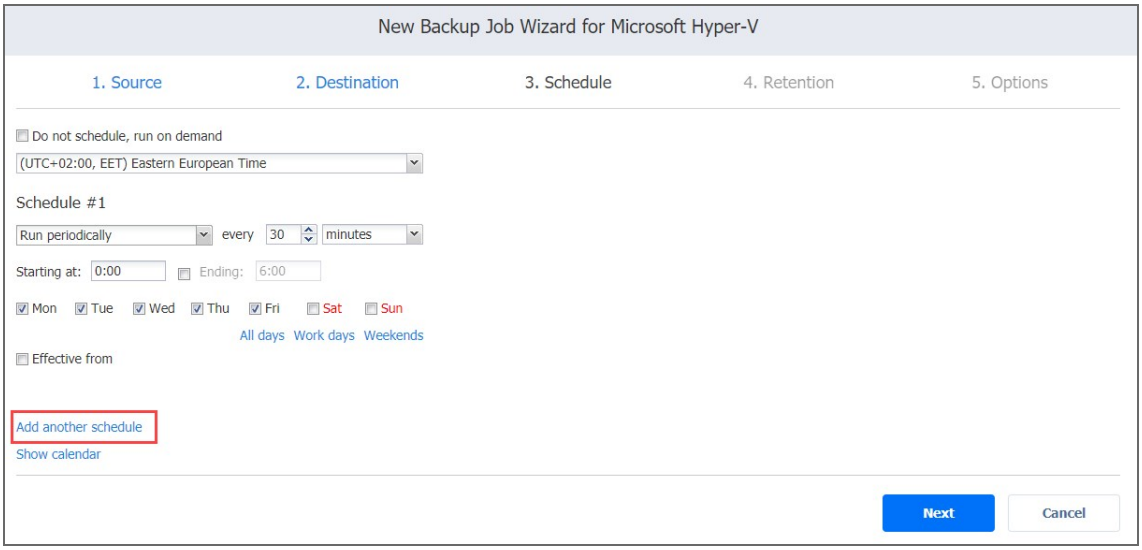

Backup Job Wizard for Hyper-V: Retention

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. Immutable recovery points cannot be overwritten, deleted, or changed by anyone, not even the root user, until the specified period has expired.

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.

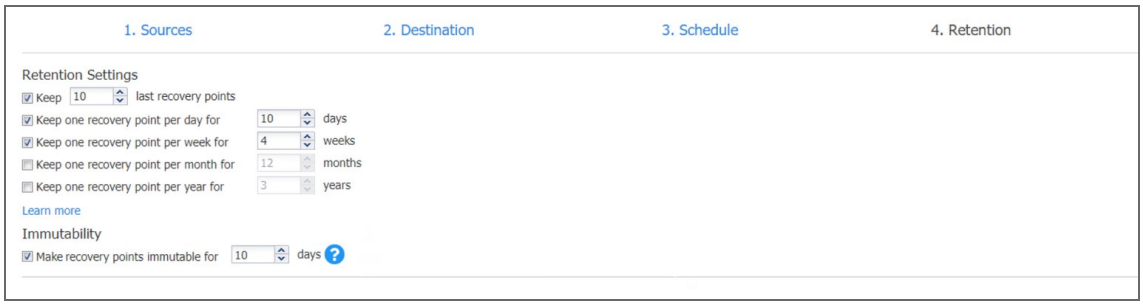

For more details and an example of job retention settings, refer to the Keeping Recovery Points article in the Knowledge Base.

Backup Job Wizard for 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.

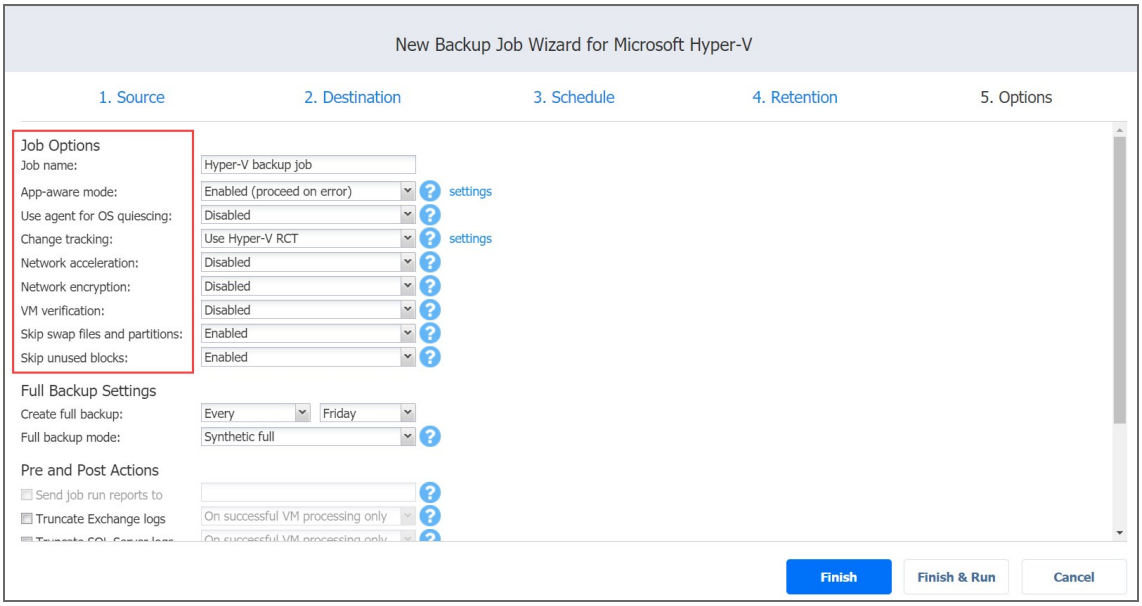

Job Name

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.

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 47.
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 "Email Notifications" on page 287) 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.

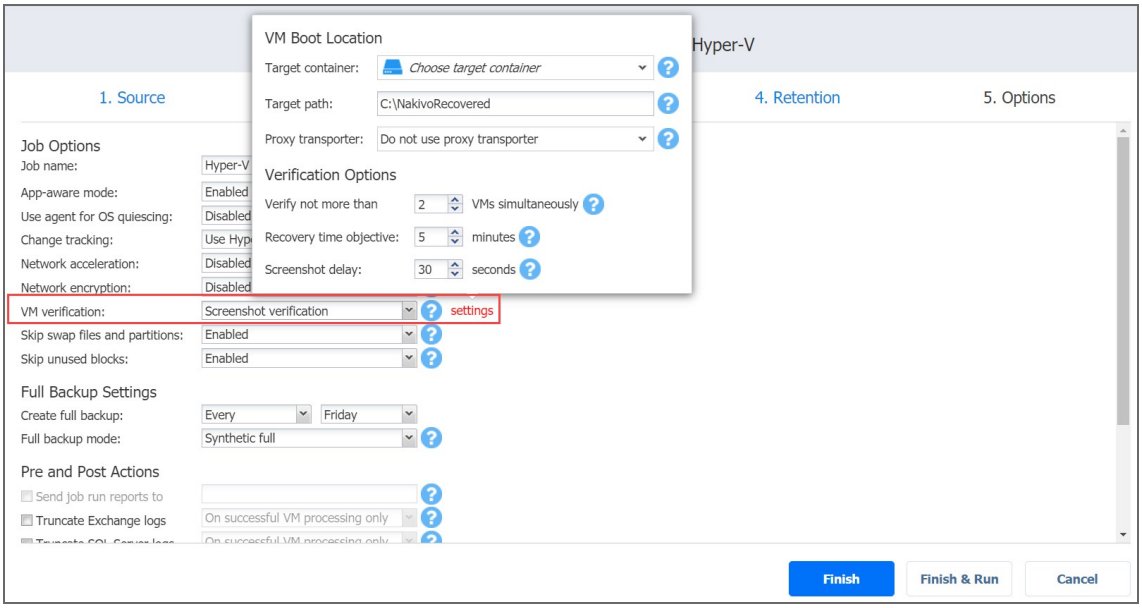

After selecting **Boot verification**, do the following in the dialog box that opens:

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.

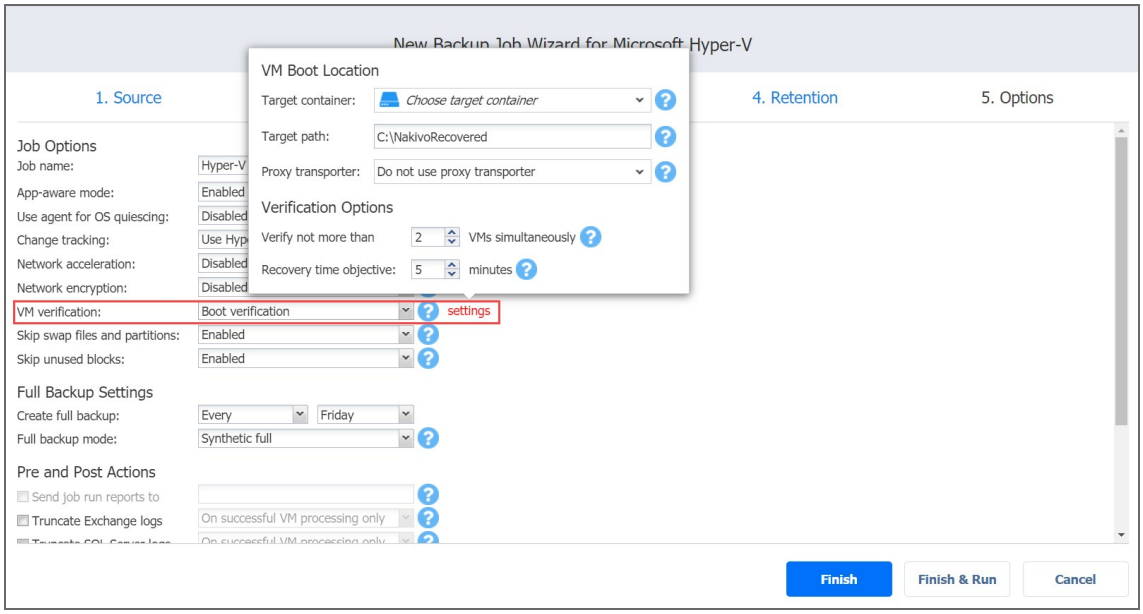

Skip Swap Files and Partitions

With this option enabled, NAKIVO Backup & Replication automatically skips swap files and partitions during the backup process.

Skip Unused Blocks

With this option enabled, NAKIVO Backup & Replication automatically skips 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**, you can specify the following options:

- **Create full backup**: Specify how often full backups should be created.
- **Full backup mode**: Specify how a 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. The benefits of this approach are as follows:
    - 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.

    Note

On a full active backup, the entire source disks will be read even if Hyper-V RCT is selected.

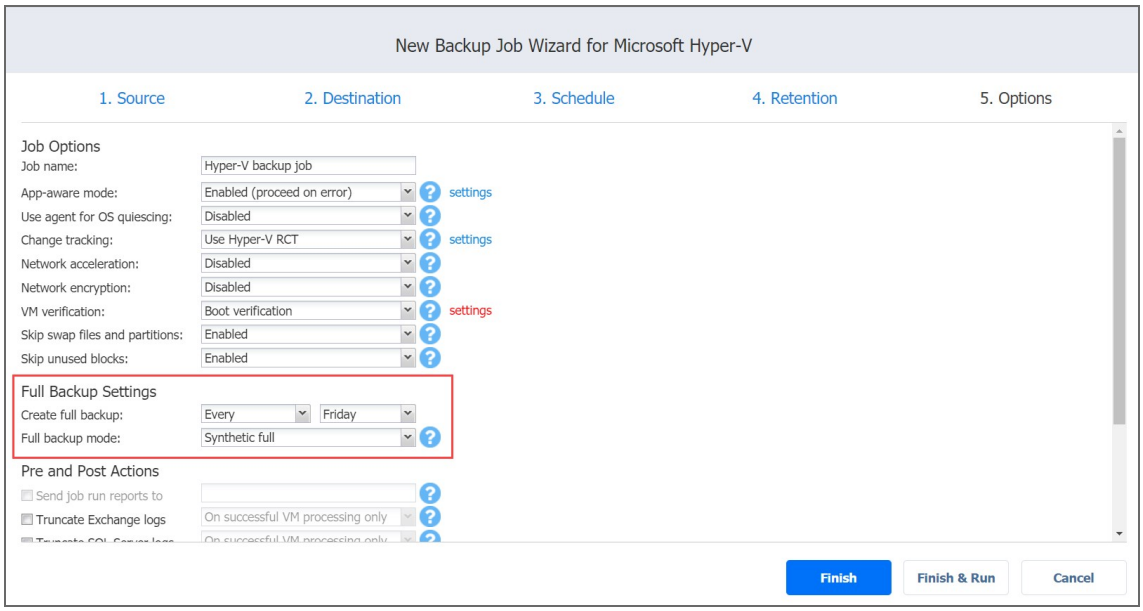

Pre and Post Job Actions

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.

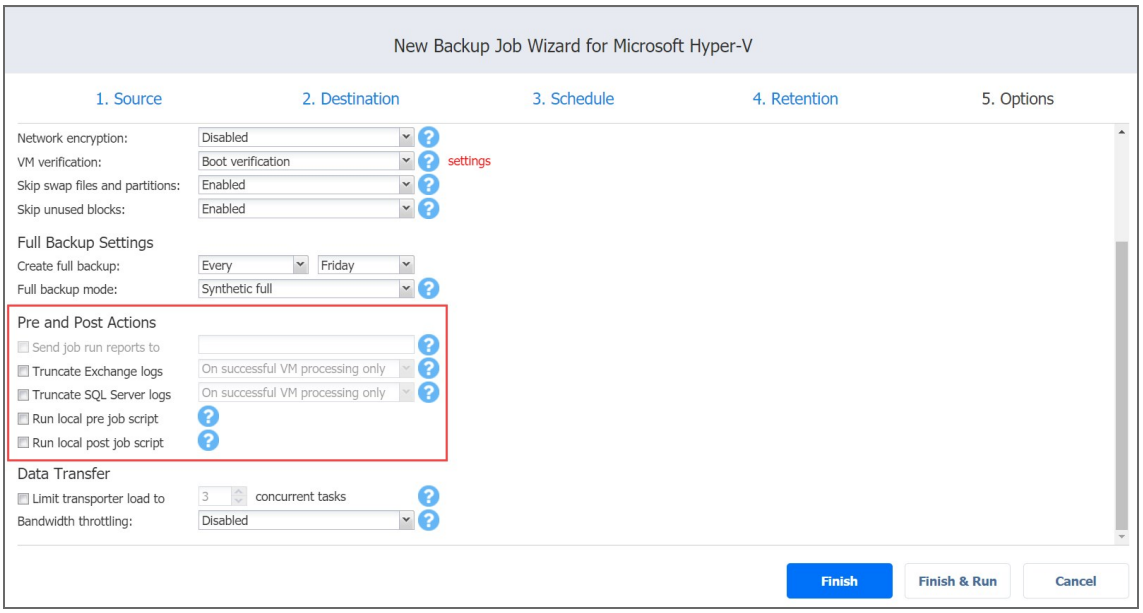

Data Transfer

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 281 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 281 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.

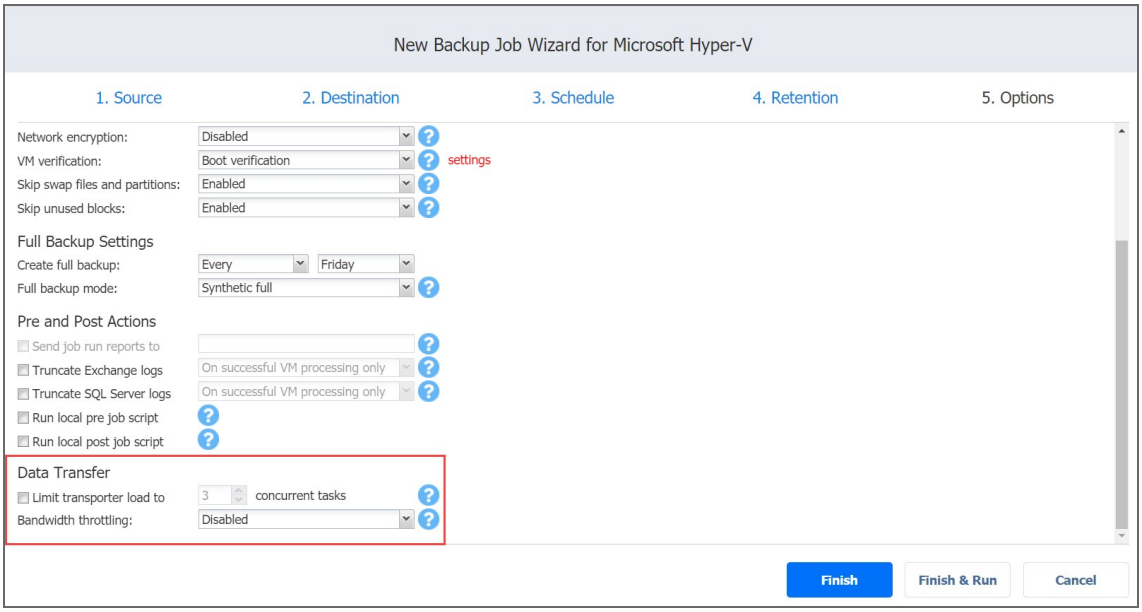

Completing New Backup Job Wizard for Microsoft Hyper-V

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 250 for details.

Creating Backup Copy Jobs

To create a backup copy job, click **Create** on the **Dashboard**, and then click **Backup copy job**.

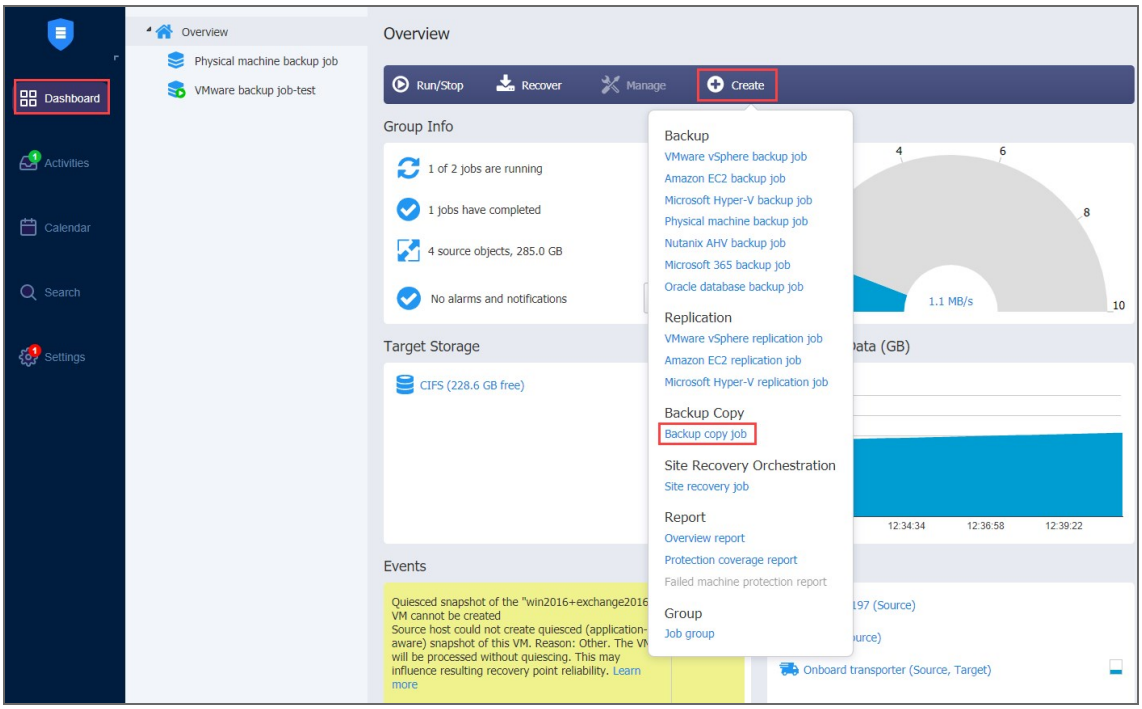

The **New Backup Copy Job Wizard** opens. Complete the wizard as described in the sections below:

- “Backup Copy Job Wizard: Backups” on page 553
- “Backup Copy Job Wizard: Destination” on page 556
- “Backup Copy Job Wizard: Schedule” on page 559
- “Backup Copy Job Wizard: Retention” on page 563
- “Backup Copy Job Wizard: Options” on page 565

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.
3. The selected items appear in the right pane of the page. If necessary, reorder the selected items by dragging them to a new position. By doing so, you can specify what items you wish to back up first.
4. Review the list of the selected items. If needed, remove a selected backup from the backup copy job in either of the following ways:
    - Cancel the selection of the item(s) in the left pane. This will remove the item(s) 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(s) in the left pane.

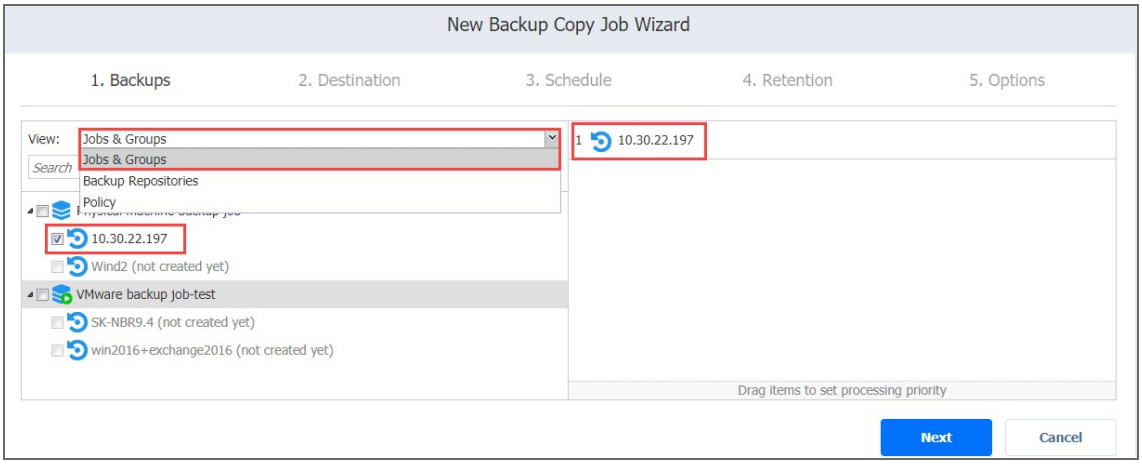

Creating Backup Copies Using Backup Repositories

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.

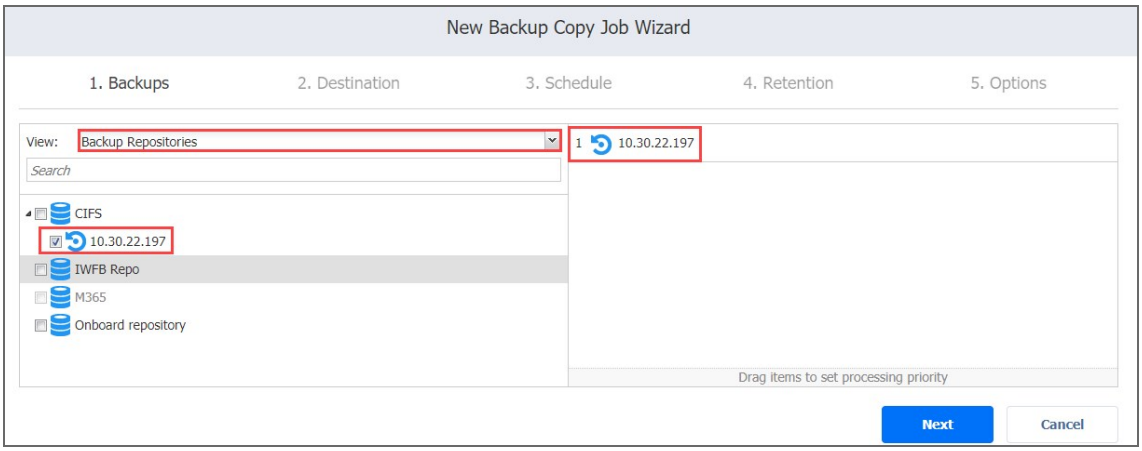

Creating Backup Copies Using Policies

When the **Policy** view is selected, it allows you to use job policies; refer to "Managing Job Policies" on page 262 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 265 for details.

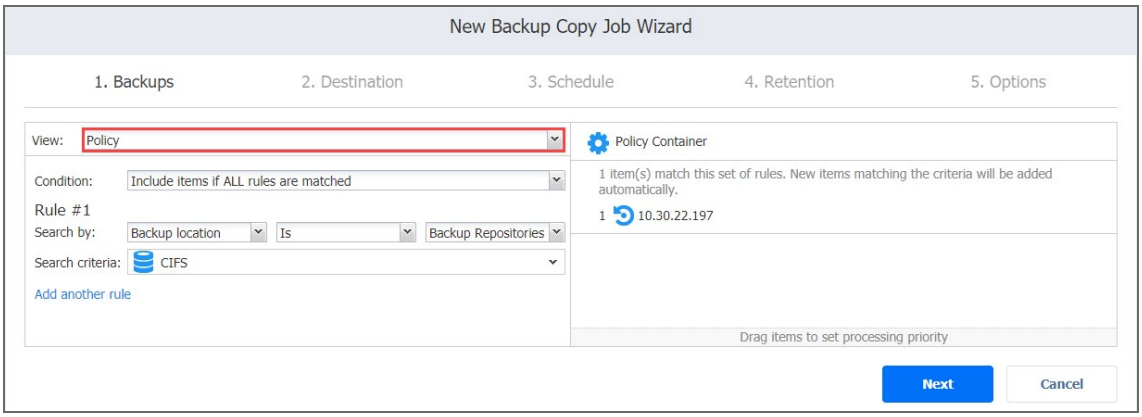

Click **Next** to confirm that you wish to add selected items to the backup copy job. The wizard will display the next page.

Notes

- When you add a container – a group, job, or a 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.

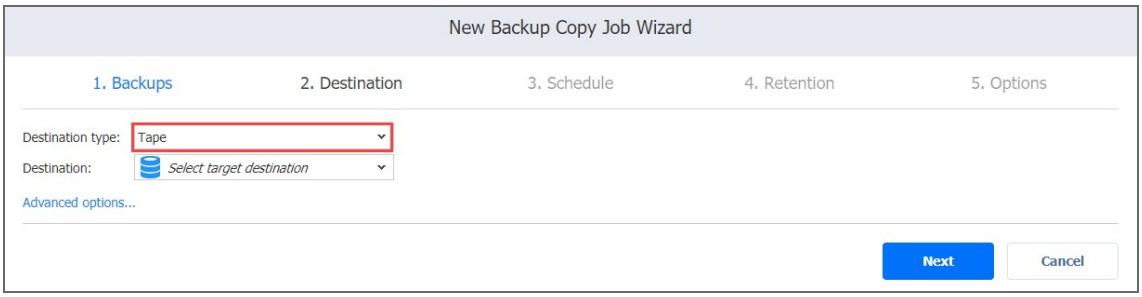

Selecting a Target Backup Repository

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.

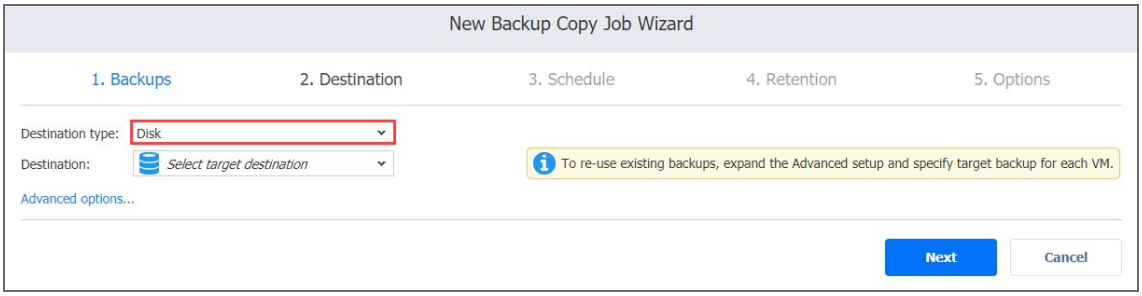

- To copy backups to different Backup Repositories, follow the steps below:

a. Click **Advanced options**.

b. For each backup, select a target Backup Repository.

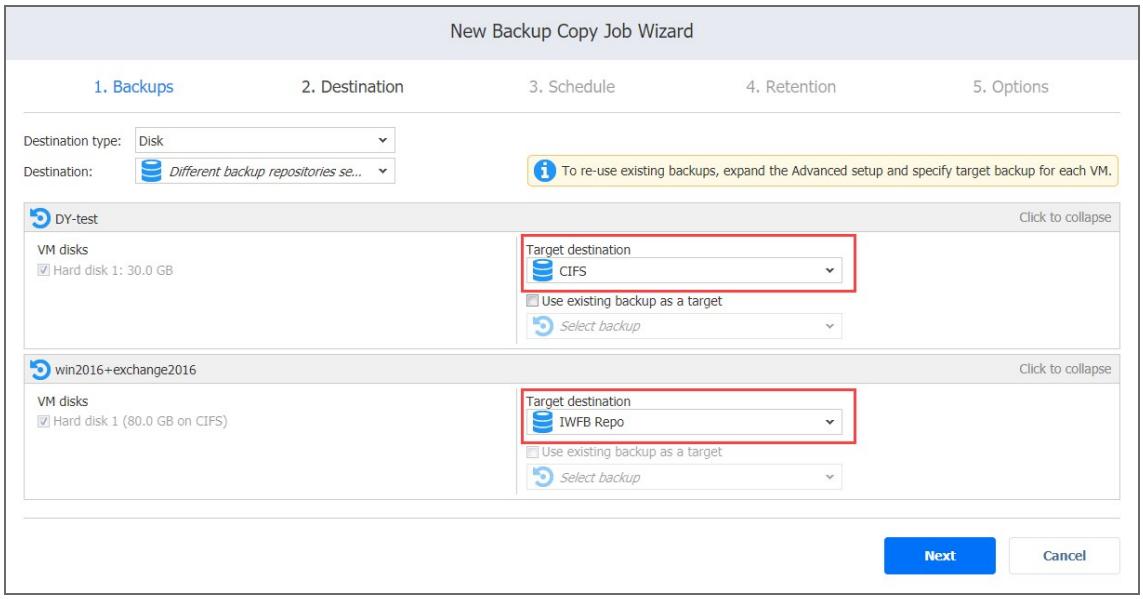

Mapping Source Backups to Existing Backups

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 options**.
2. From the **Backup repository** drop-down list, choose a Backup Repository that contains a copy of the source backup.
3. Select the **Use existing backup as a target** option and select the existing backup copy from the drop-

down list.

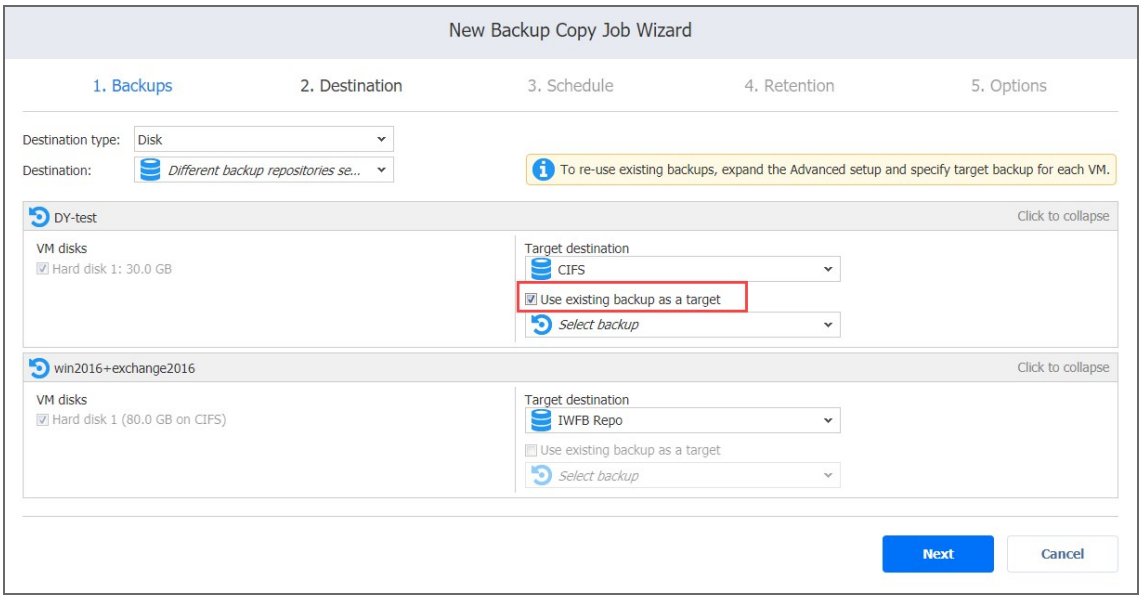

When running the job, the product analyzes the existing backup copy you have selected, determines how it is different from the source backup, and transfers only the differential data.

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.

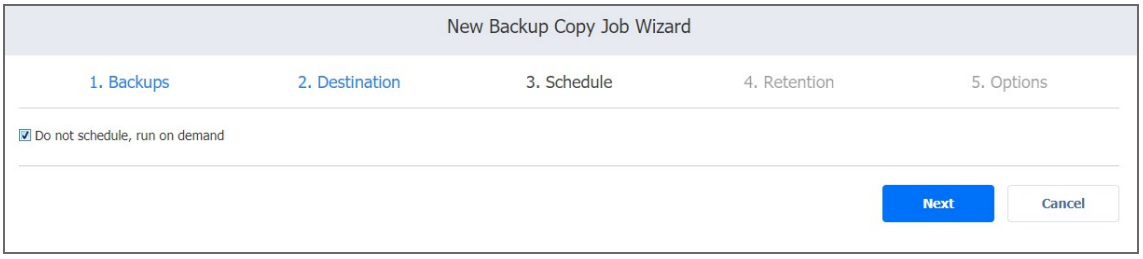

Daily or Weekly Backup

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.

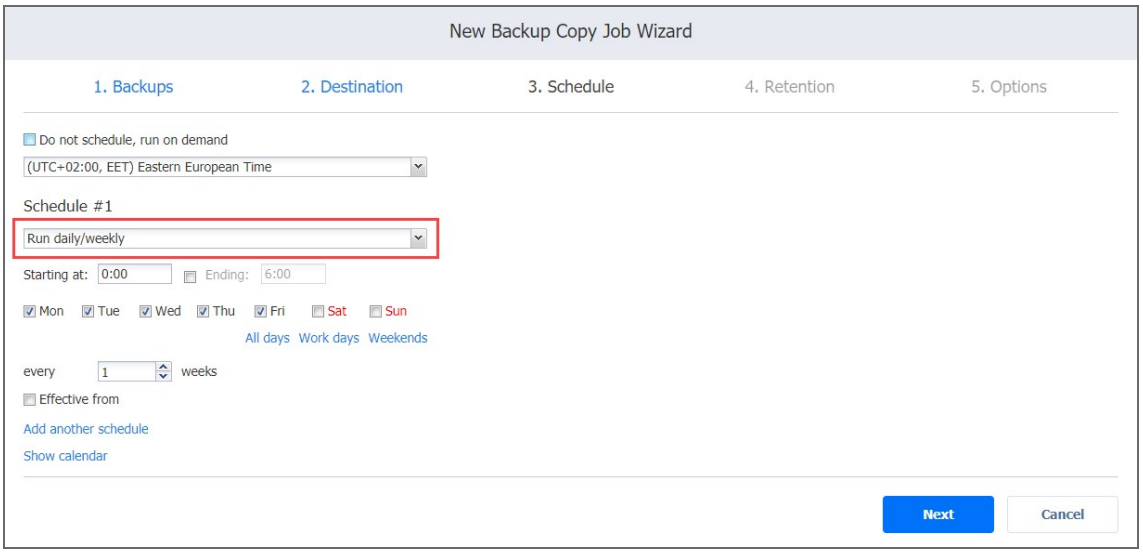

Monthly or Yearly Backup

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.

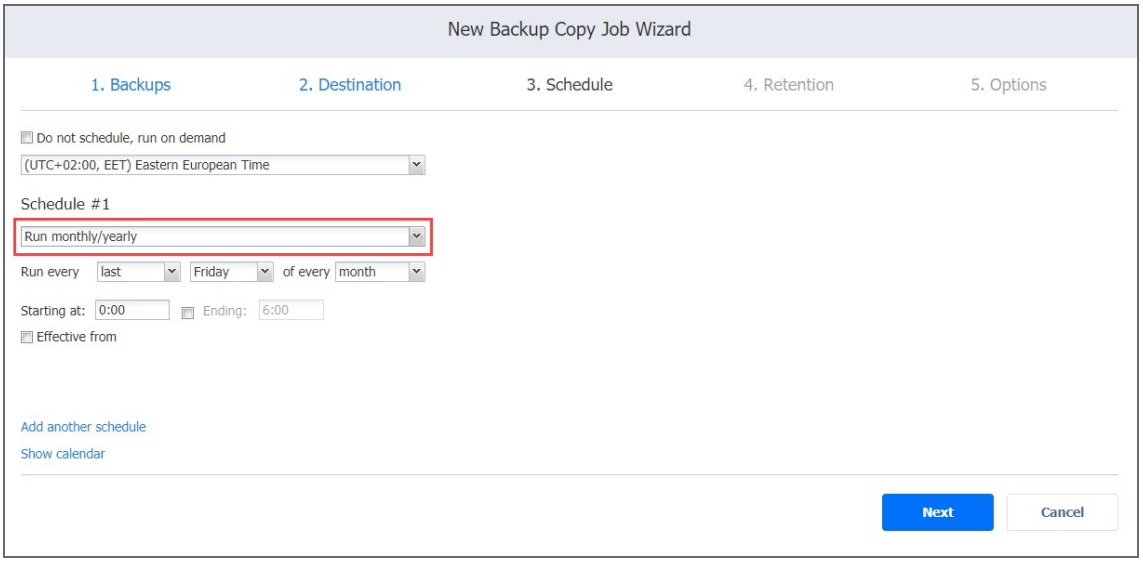

Periodic Backup

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.

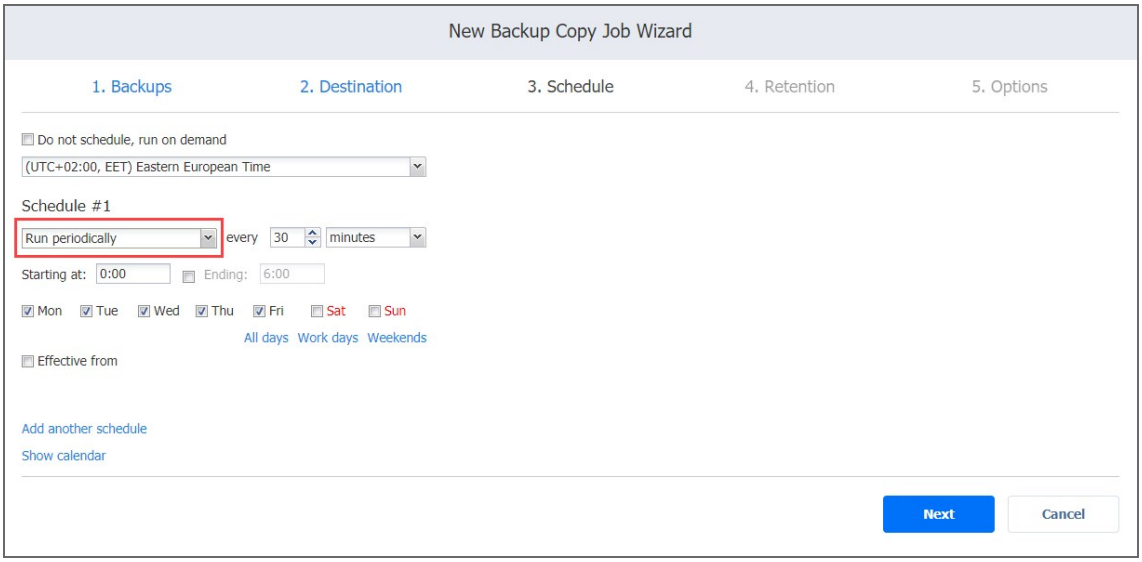

Chained Job

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.

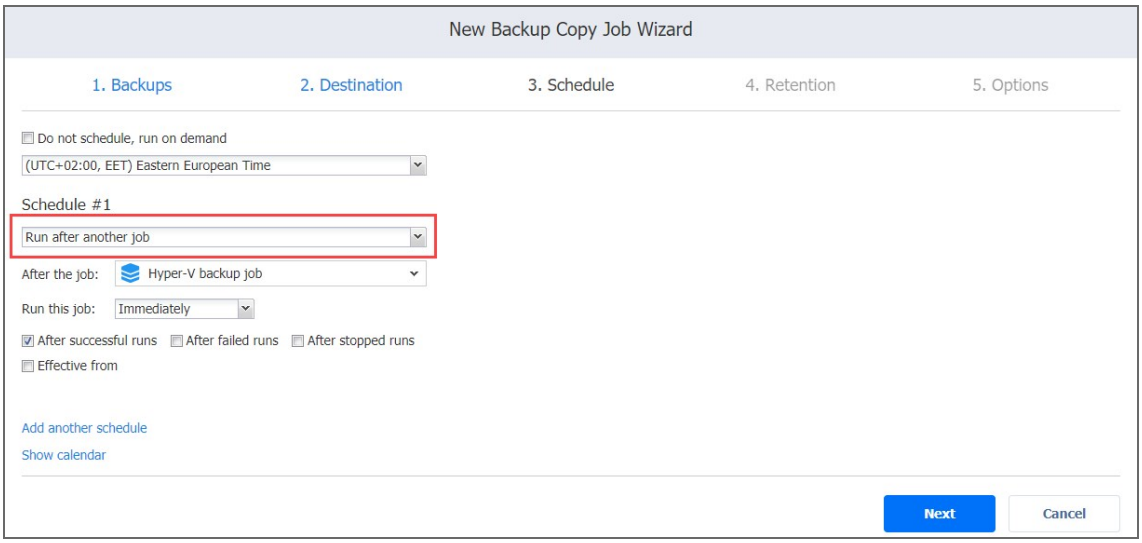

Adding Additional Schedule

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.

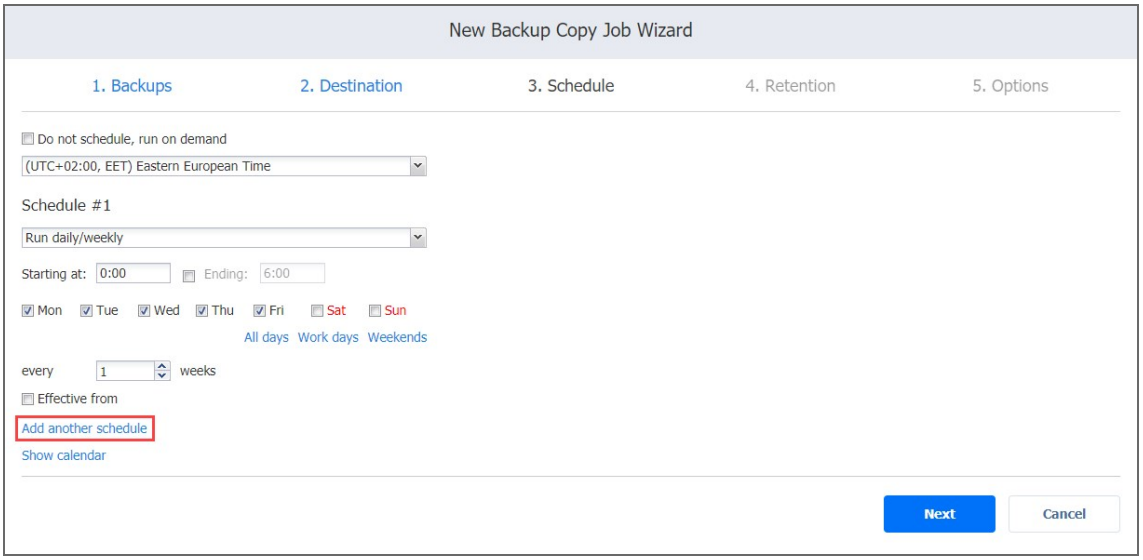

Backup Copy Job Wizard: Retention

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. Immutable recovery points cannot be overwritten, deleted, or changed by anyone, not even the root user, until the specified period has expired.

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.

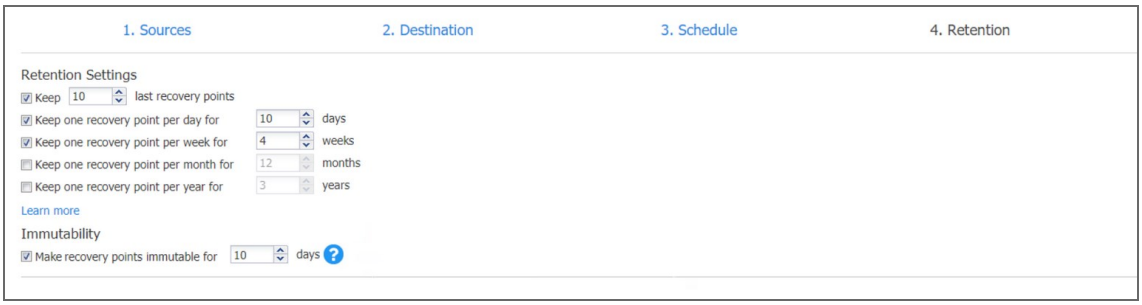

For more details and an example of job retention settings, refer to the Keeping Recovery Points article in the Knowledge Base.

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
- 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.

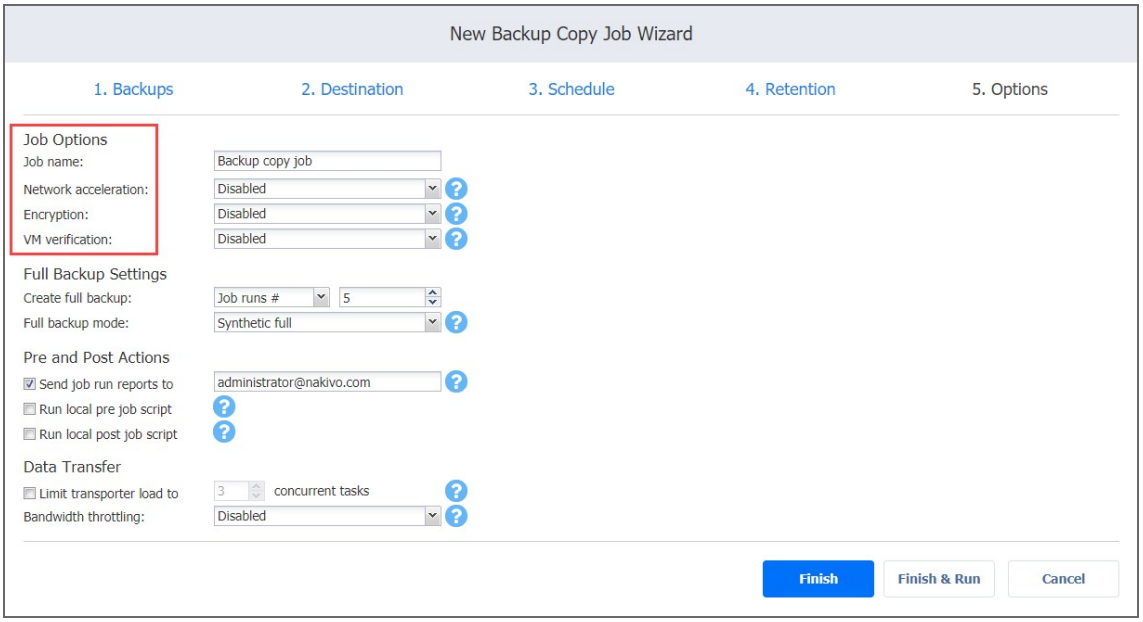

Job Name

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 47 .

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.

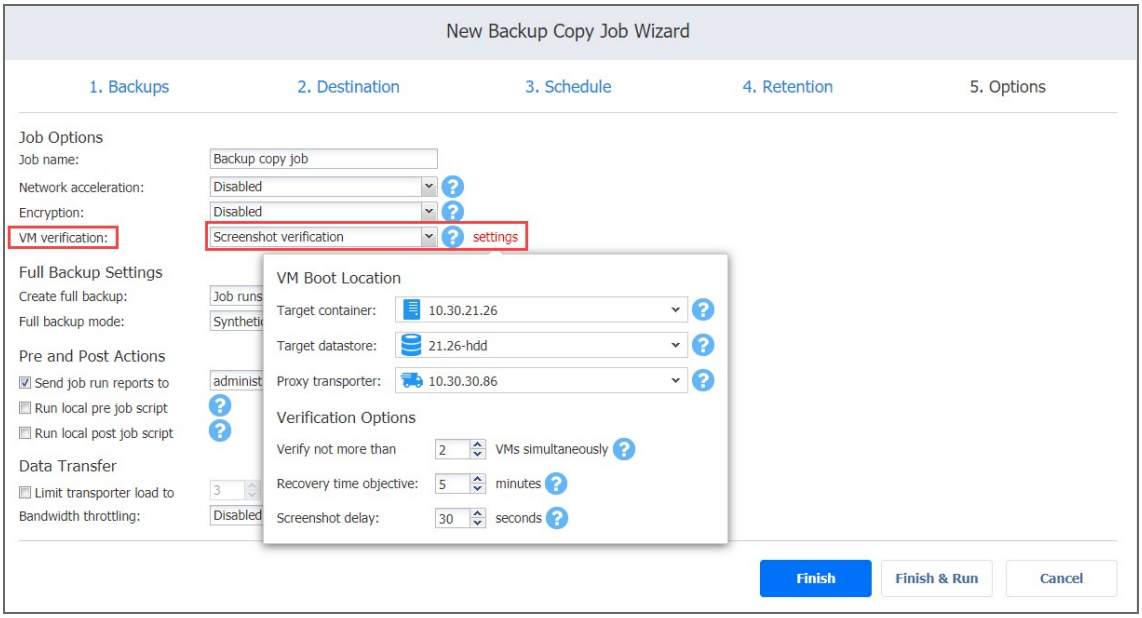

After selecting **Boot verification**, do the following in the dialog box that opens:

1. Provide the location of the VMs to be booted as described for the **Screenshot verification** option.
2. Set verification options:
    - **Verify not more than x VMs simultaneously**: Specify the maximum number of VMs that can be started on the Target Container simultaneously.
    - **Recovery time objective**: Specify the amount of time allocated for verification of each VM backup. If a VM OS does not start within the specified amount of time, verification will be

considered failed.

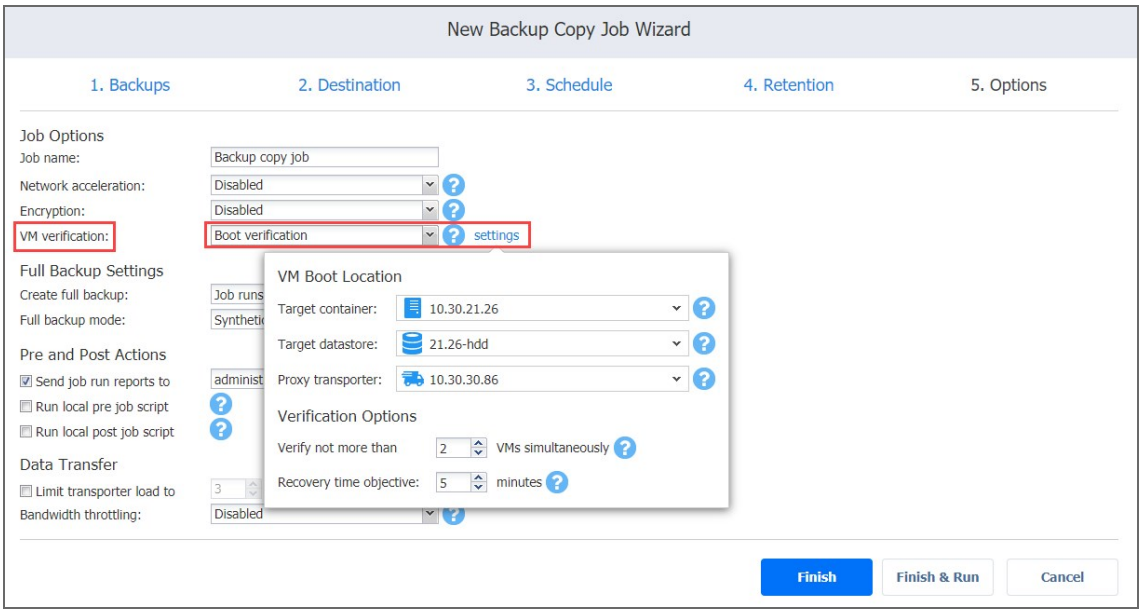

Full Backup

If the type of the Backup Repository that you selected on the Destination page is set to **Incremental with full backups**, 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 has changed since the last backup) and will then transform the available data into a full backup file. The benefits of this approach are as follows:
    - 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.

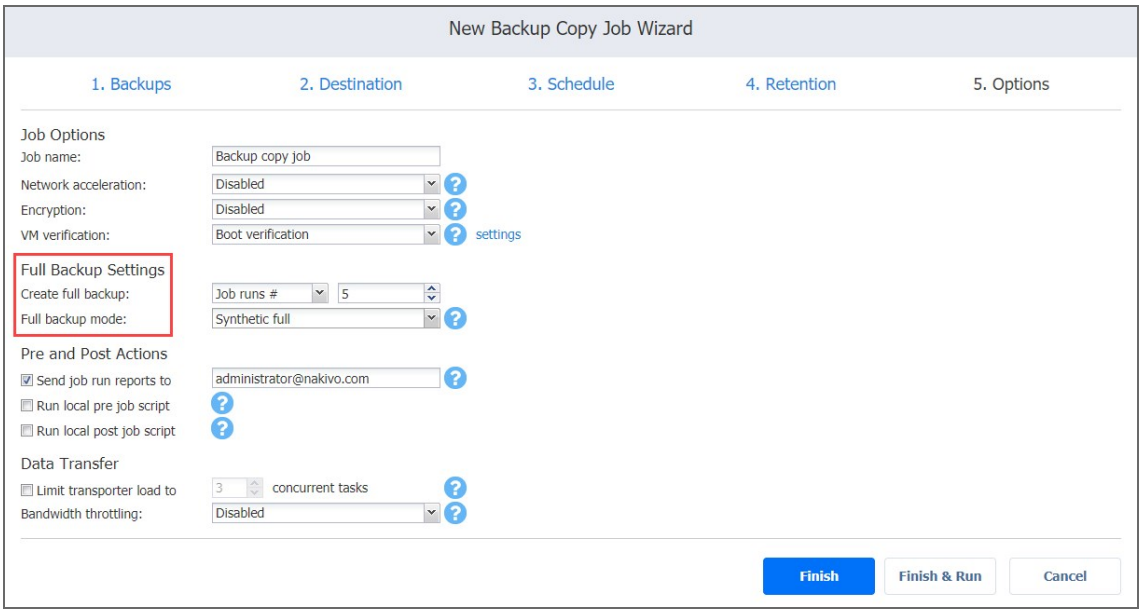

Pre and Post Actions

NAKIVO Backup & Replication allows you to set up certain actions before a backup copy job begins and after it has completed. You can choose to send job run reports to the email provided and run local pre and post job scripts.

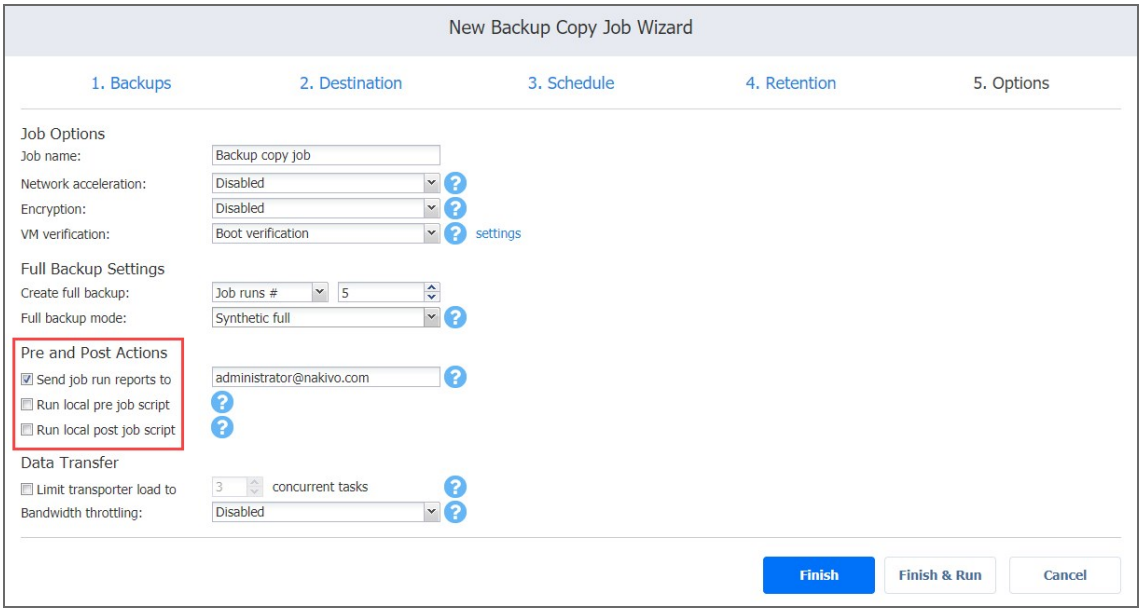

Email Notifications

NAKIVO Backup & Replication can send email notifications on job completion status to specified recipients. This feature complements global notifications and provides you with the ability to configure notifications on a per-job level.

To enable this option, configure your Email settings.

To send email notifications, do the following:

1. In the *Pre and Post Actions* section, select the **Send job run reports to** option.
2. Specify one or more email addresses in the text box. Separate multiple email addresses with a semicolon.

Pre Job Script

To run a script before the product begins copying backups, do the following:

1. Place a script file on the machine where the Director is installed.
2. In the *Pre and Post Actions* section, select the **Run local pre job script** option and click the **settings** link. Specify the following parameters in the dialog box that opens:

- **Script path**: Specify a local path to the script on the machine where the Director is installed. Script interpreter should be specified.
  Example (Windows): *cmd.exe /c D:\script.bat*

  Example (Linux): *bash /root/script.sh*

- **Job behavior**: Choose either of the following job behaviors in relation to script completion:
  - **Wait for the script to finish**: If this option is selected, the backup copy will not be started until the script is completed.
  - **Do not wait for the script to finish**: If this option is selected, the product will run the script and will start copying backups at the same time.
- **Error handling**: Choose either of the following job behaviors in relation to script failure:
  - **Continue the job on script failure**: If this option is selected, the job will perform backup copy even if the script has failed.
  - **Fail the job on script failure**: If this option is selected and the script fails, the job will be failed and the backup copy will not be performed.

Post Job Script

To run a script after the product has finished copying all backups, do the following:

1. Place a script file on the machine on which the Director is installed.
2. In the *Pre and Post Actions* section, select the **Run local post job script** option and click the **settings** link. Specify the following parameters in the dialog box that opens:

- **Script path**: Specify a local path to the script on the machine on which the Director is installed. Script interpreter should be specified.
  Example (Windows): *cmd.exe /c D:\script.bat*
  Example (Linux): *bash /root/script.sh*
- **Job behavior**: Choose either of the following job behaviors in relation to script completion:
  - **Wait for the script to finish**: If this option is selected, the job will be in the "running" state until the script is completed.
  - **Do not wait for the script to finish**: If this option is selected, the job will be completed even if the script execution is still in progress.
- **Error handling**: Choose either of the following job behaviors in relation to script failure.
  - **Continue the job on script failure**: If this option is selected, script failure will not influence the status of the job.
  - **Fail the job on script failure**: If this option is selected and the script has failed, the job status will be set to "failed" even if VM backup has been successful.

Notes

- Pre- and post-job scripts can be executed only on the machine on which the Director is installed.
- When Integration Services are used on Hyper-V 2016 and above, custom pre/post scripts are unavailable for Windows VMs.

Data Transfer

In the *Data Transfe*r section of the **Options** page, you can specify a Transporter load and configure bandwidth throttling.

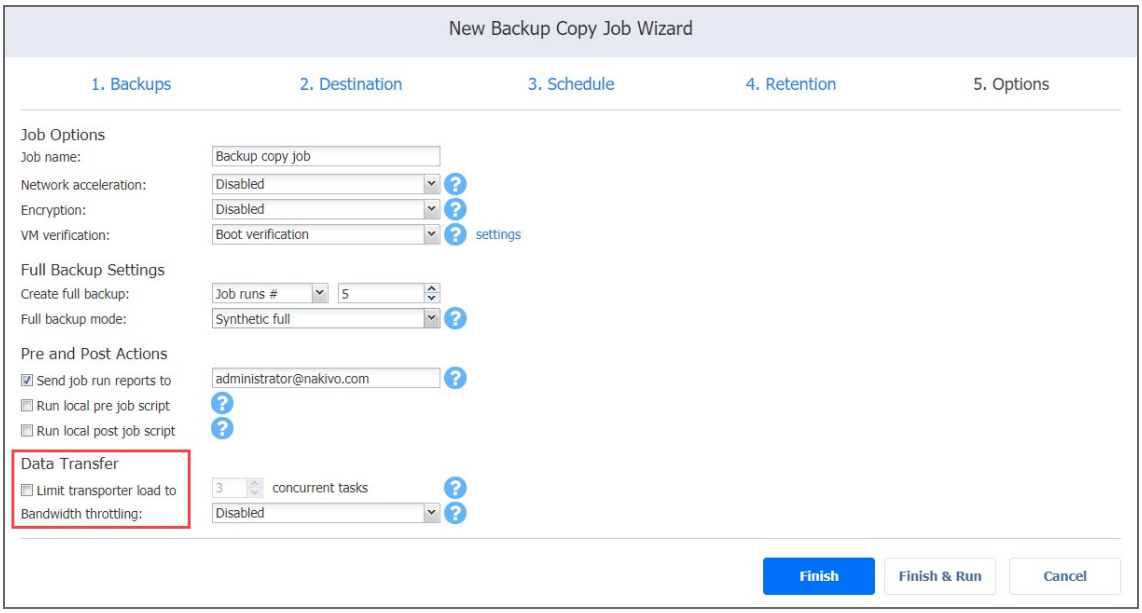

Transporter Load

You can limit the maximum number of Transporter tasks used by the job. By default, it is set to 3 concurrent tasks.

To change the default number of tasks, do the following:

1. In the *Data Transfer* section, select the **Limit transporter load to checkbox**.
2. Specify the number of concurrent tasks in the corresponding box.

Bandwidth Throttling

Follow the steps below to regulate the speed of data transfer over the network for your backup copy job:

1. For the **Bandwidth throttling** option, choose **Enabled**.

   Note

   If bandwidth throttling is disabled for the current job, global bandwidth rules may still apply to your job.
2. Click the **settings** link that becomes available.
3. The **Job Bandwidth Rules** dialog box opens displaying you the list of available rules. You have the following options:
   - Create a new bandwidth rule for your backup copy job:
     a. Click the **Create New Rule** button.
     b. The **New Bandwidth Rule** dialog box opens. Refer to "Bandwidth Throttling" on page 281 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.

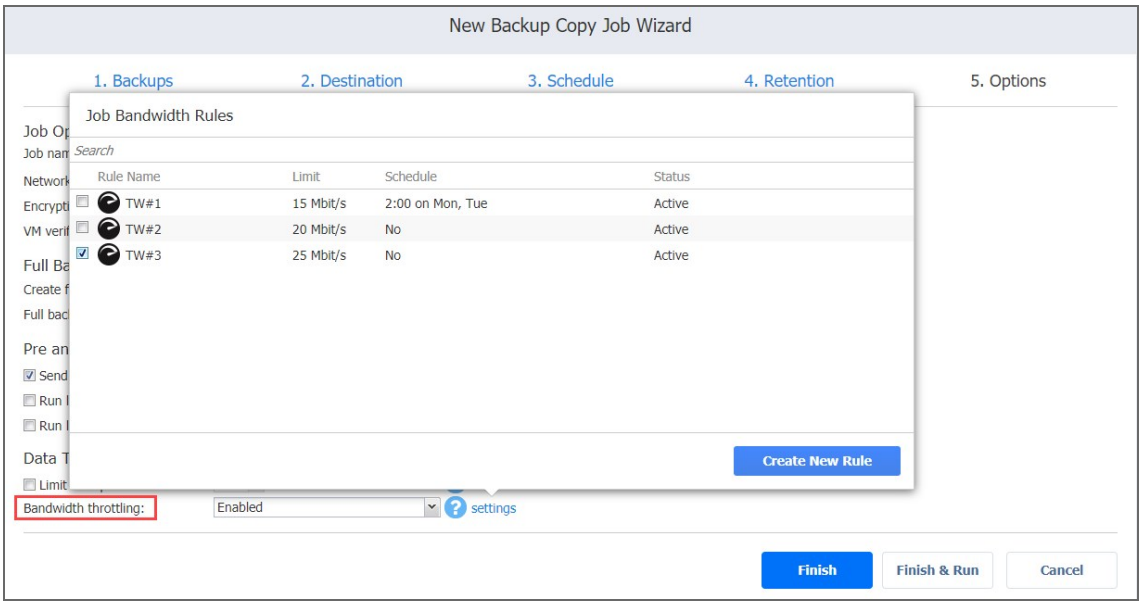

Completing the New Backup Copy Job Wizard

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 250 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**.

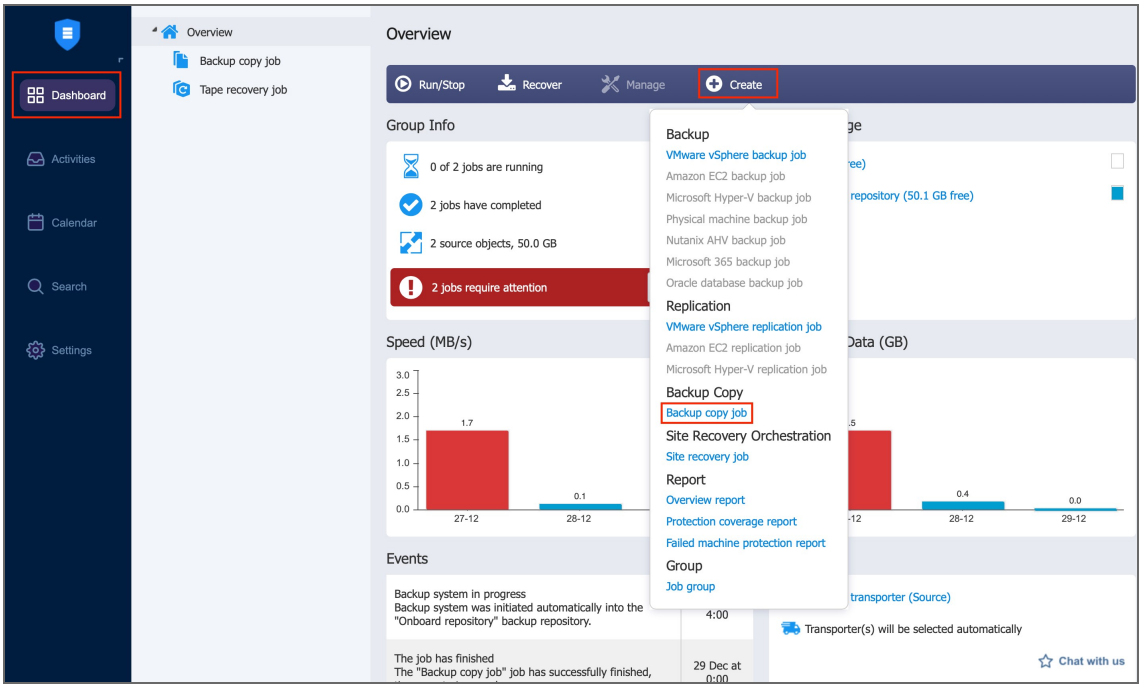

The **New Backup Copy Job Wizard** opens. Complete the wizard as described in the sections below:

- “Tape Backup Wizard: Backups” on page 575
- “Tape Backup Wizard: Destination” on page 578
- “Tape Backup Wizard: Schedule” on page 579
- “Tape Backup Wizard: Retention” on page 583
- “Tape Backup Wizard: Options” on page 584

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.

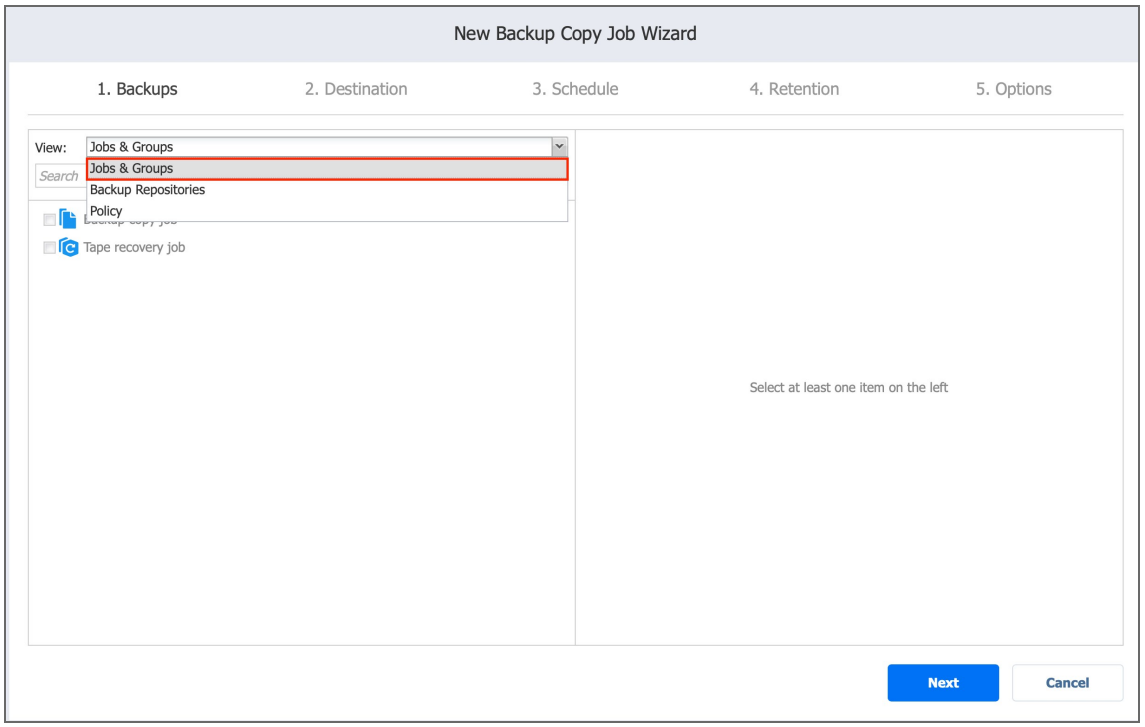

- **Backup Repositories**: If chosen, the inventory tree shows available Backup Repositories along with the backups in them. Proceed as it is described for the **Jobs & Groups** view above.

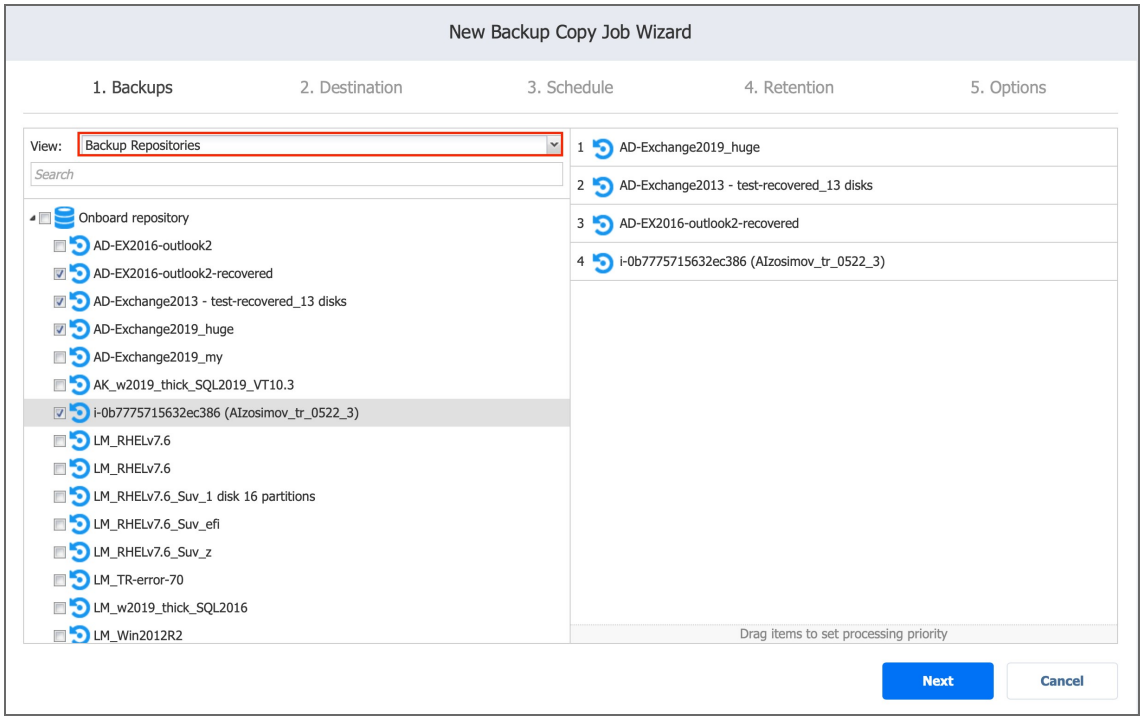

- **Policy**: If selected, this allows you to use job policies; refer to "Managing Job Policies" on page 262 for details. Please follow the steps below:
    a. If items were selected in alternate views, a dialog box opens warning you that switching to the **Policy** view will reset your current selection. 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 265 for details.

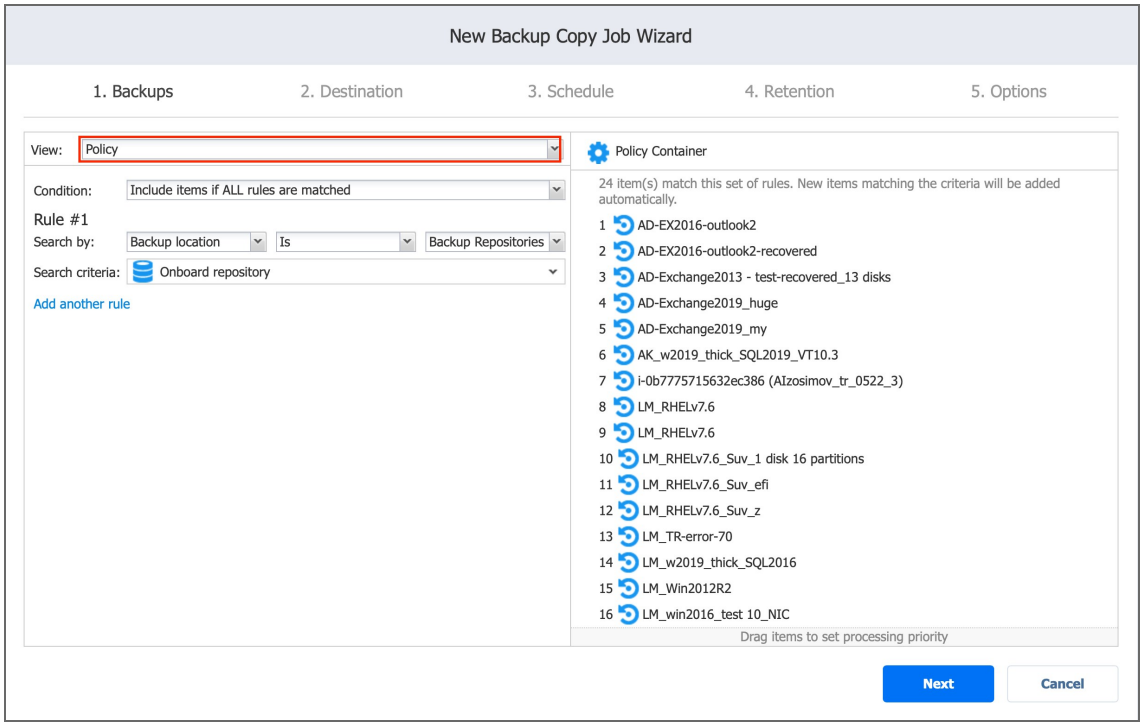

2. Click **Next** to confirm adding selected items to the backup copy job.

The wizard will display the next page.

Notes

1. If you add a container – a group, job, or a 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.

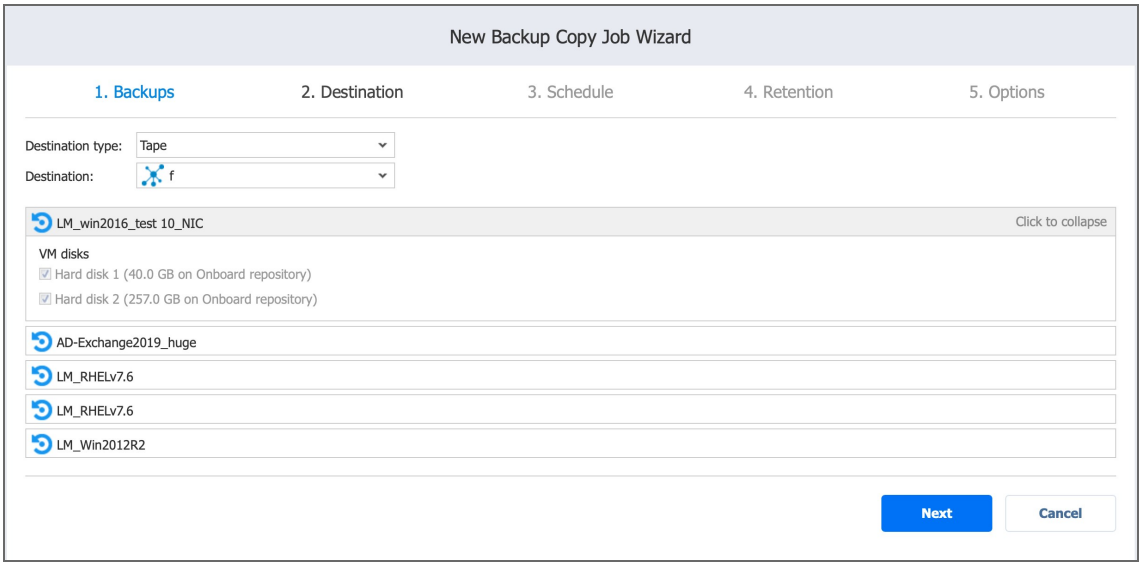

4. Click **Next** to proceed to the next page.

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.

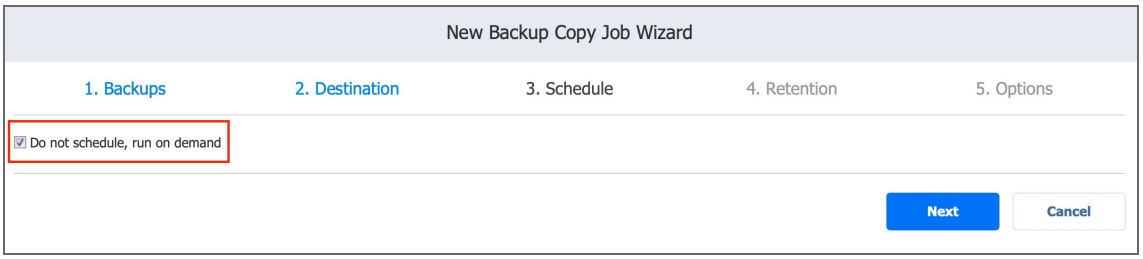

Daily or Weekly Backup

To run the job once a day, choose **Run daily/weekly** from the schedule drop-down list and do the following:

- Specify the time when the job should be started in the **Starting** box.
- Specify the end time for the job in the **Ending** box. If the job has not completed by the time specified, the job will be stopped.
- Choose a time zone that should be used for the job start and end times from the time zone drop-down list.
- Select the days of the week during which the job will be started.
- If necessary, select the **Effective from** checkbox and pick the date when the schedule comes into effect.

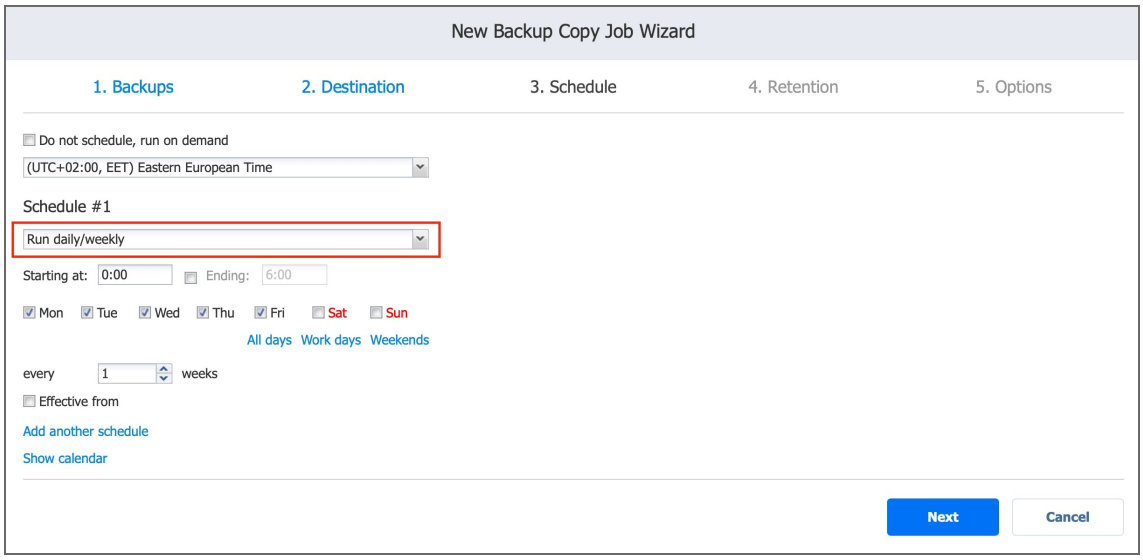

Monthly or Yearly Backup

To run the job monthly or yearly, choose **Monthly/yearly** from the schedule drop-down list and do the following:

- Specify the job start schedule in the appropriate boxes.
- Specify the day and month when the job should be started in the **Run every** boxes.
- Specify the time when the job should be started in the **Starting** box.
- Specify the end time for the job in the **Ending** box. If the job has not completed by the time specified, the job will be stopped.
- If necessary, select the **Effective from** checkbox and pick the date when the schedule comes into effect.

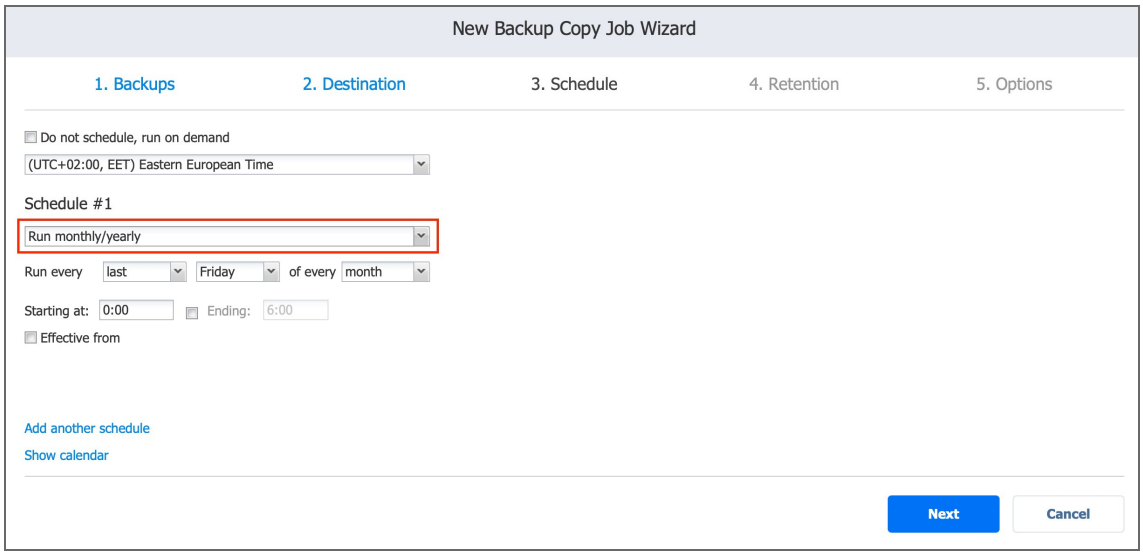

Periodic Backup

To run the job multiple times per day, choose **Run periodically** from the schedule drop-down list and then choose a time period from the appropriate boxes:

- Specify the time when the job should be started in the **Starting** box.
- Specify the end time for the job in the **Ending** box. If the job has not completed by the time specified, the job will be stopped.
- Select the days of the week during which the job will be started.
- If necessary, select the **Effective from** checkbox and pick the date when the schedule comes into effect.

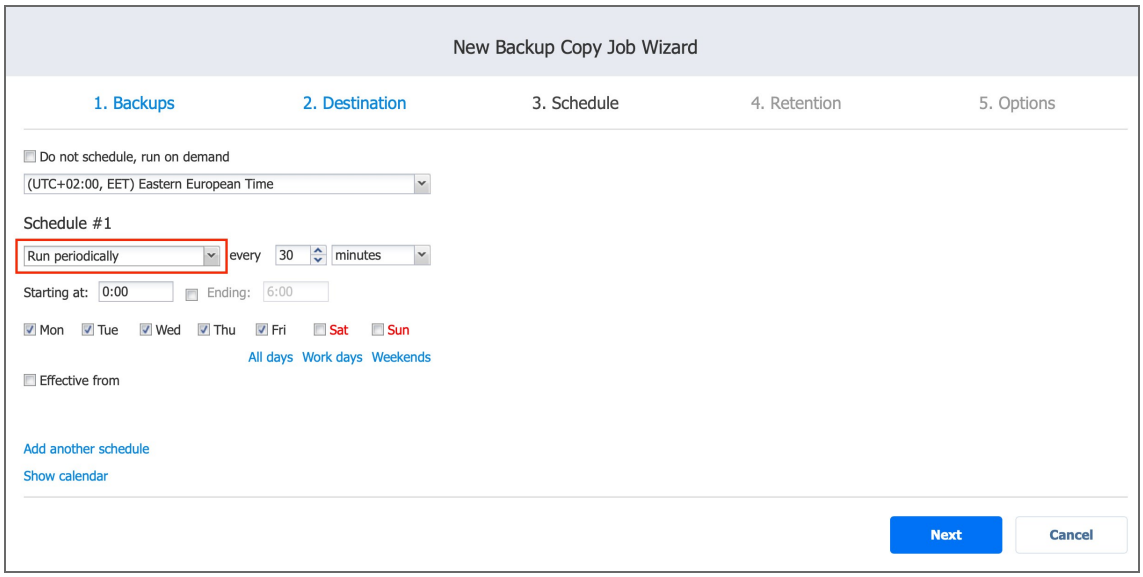

Chained Job

To run the job after a previous one has completed, choose **Run after another job** from the schedule drop-down list and set the options as follows:

- **After the job**: Select a job after which the current job will be started.
- **Run this job**: Choose whether to run the current job immediately after the previous one has completed or within a delay.
- **After successful runs**: If selected, the job will run if the previous one has completed successfully.
- **After failed runs**: If selected, the job will run if the previous one has failed.
- **After stopped runs**: If selected, the job will run if the previous one has been stopped.

- **Effective from**: If selected, the schedule will come into effect on the date picked.

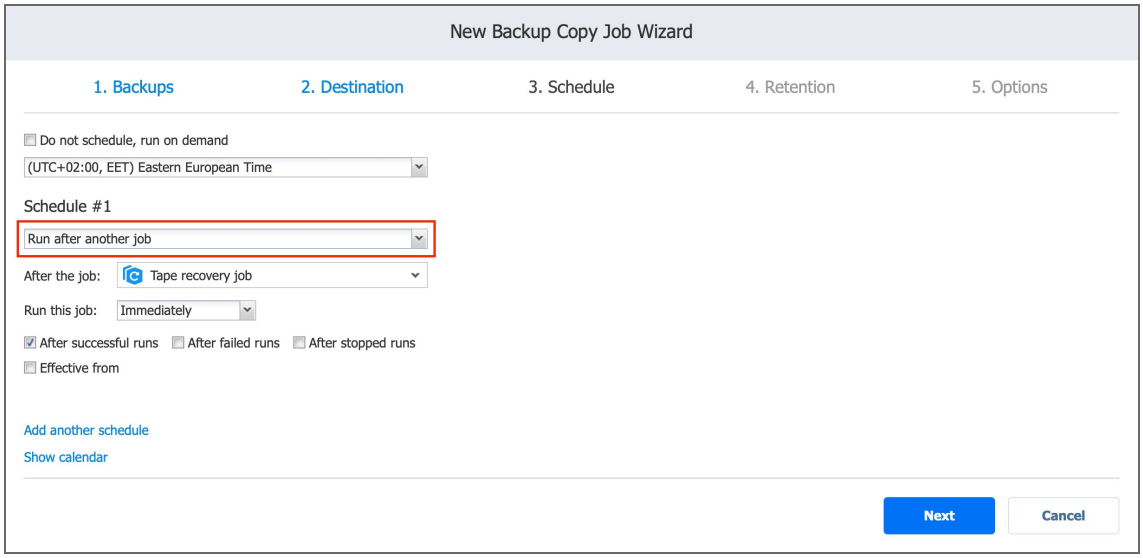

Adding Another Schedule

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.

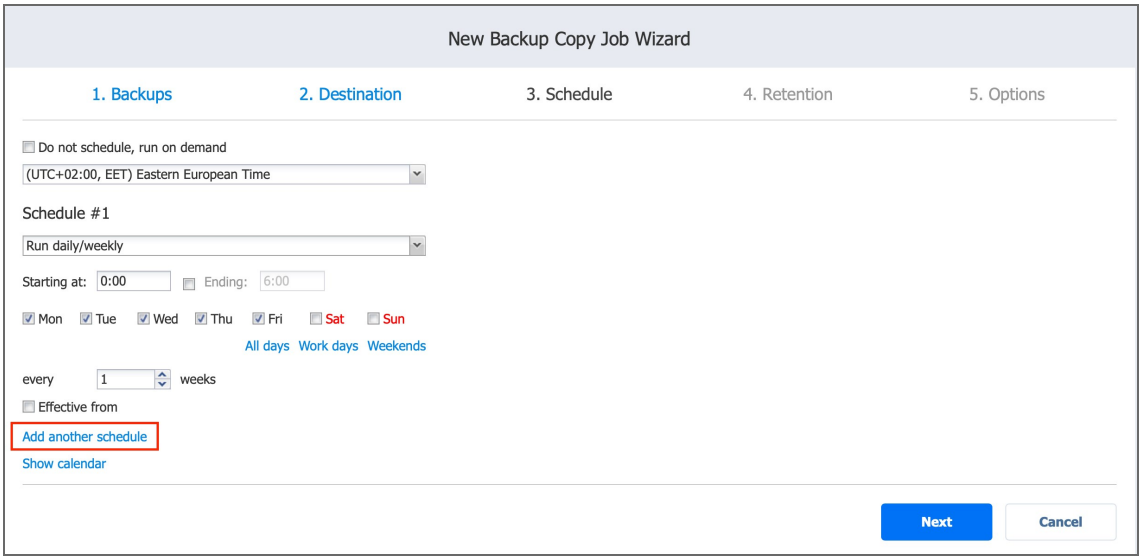

Tape Backup Wizard: Retention

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.

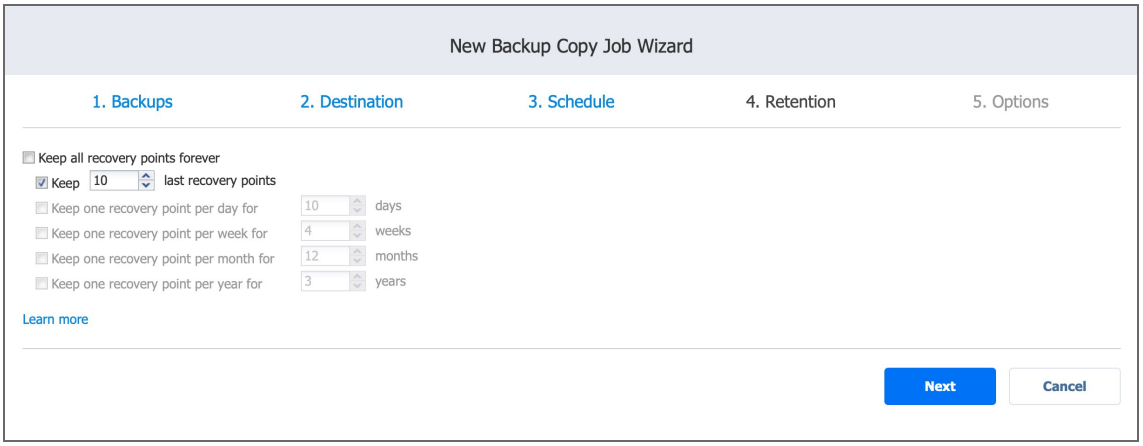

For an example of job retention settings and further explanation, refer to the Keeping Recovery Points Knowledge Base article.

Tape Backup Wizard: Options

On the **Options** page of the wizard, you can name the job and enable/disable network acceleration and encryption. What's more, 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.

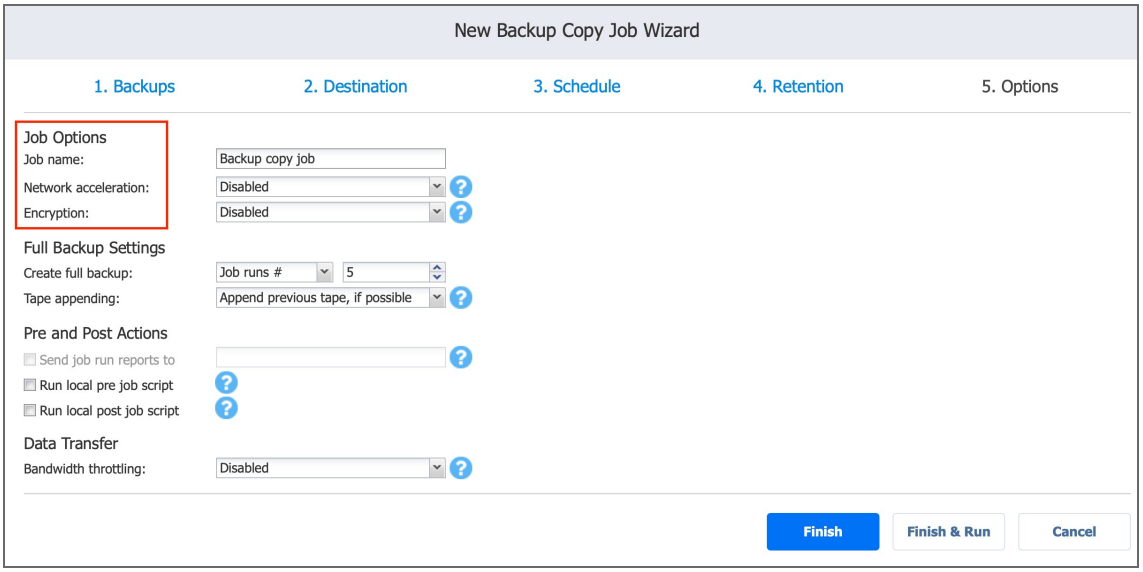

Job Name

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.

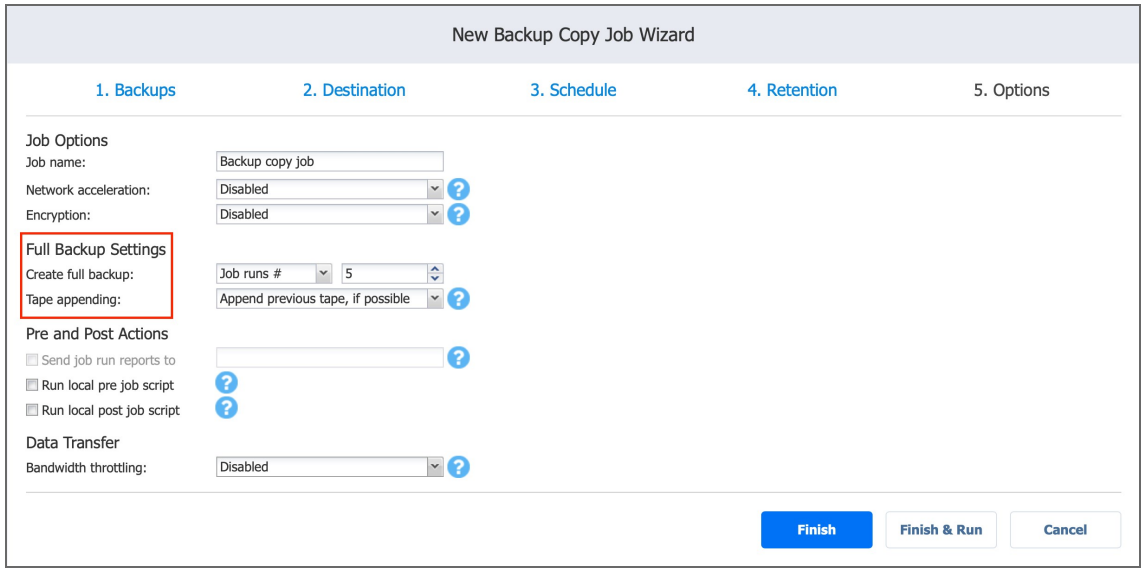

Creating Full Backup

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

Note

Only **Always** and **Job runs #** options are always available for selection. The rest of the options' availability depends on the scheduling settings specified on the Schedule page.

Tape Appending

The **Tape appending** feature allows you to set up the rules for data appending. The following options are available:

- **Append previous tape if possible**:
  - The job run appends data to the last tape cartridge.
  - If the tape cartridge that was last written during the last job run is not available in the device or is full, the job starts with an empty cartridge:
    - All job objects within the job run are written to the selected tape cartridge one by one.
    - If the selected tape cartridge runs out of space, the next empty tape cartridge is selected and the process repeats until all job objects are written.
- **Start full backup with an empty tape**:
  - In case the backup modes of all job objects within the job run are defined as **full**:
    - The job run starts with an empty tape cartridge.
    - All job objects within the job run are written to the selected tape cartridge one by one.
    - If the selected tape cartridge runs out of space, the next empty tape cartridge is selected and the process repeats until all job objects are written.
  - In case the backup modes of all job objects within the job run are defined as **incremental**:
    - The job run appends data to the last tape cartridge.
    - All job objects within the job run are written to the selected tape cartridge one by one.

        - If the selected tape cartridge runs out of space, the next empty tape cartridge is selected and the process repeats until all job objects are written.
    - In case the backup modes of all job objects within the job run are defined as a **mix of full and incremental** modes:
        - The job run appends data to the last tape cartridge.
        - All job objects within the job run are written to the selected tape cartridge one by one.
        - If the selected tape cartridge runs out of space, the next empty tape cartridge is selected and the process repeats until all job objects are written.
- **Always start with an empty tape**:
    - The job starts with an empty tape cartridge.
    - All job objects within the job run are written to the selected tape cartridge one by one.
    - In case the selected tape cartridge runs out of space, the next empty tape cartridge is selected and the process repeats until all job objects are written.

Pre and Post Job Actions

NAKIVO Backup & Replication provides you with the ability to enable certain actions before a tape backup job begins and after it has completed. You can choose to send job run reports to the email provided and run local pre and post job scripts.

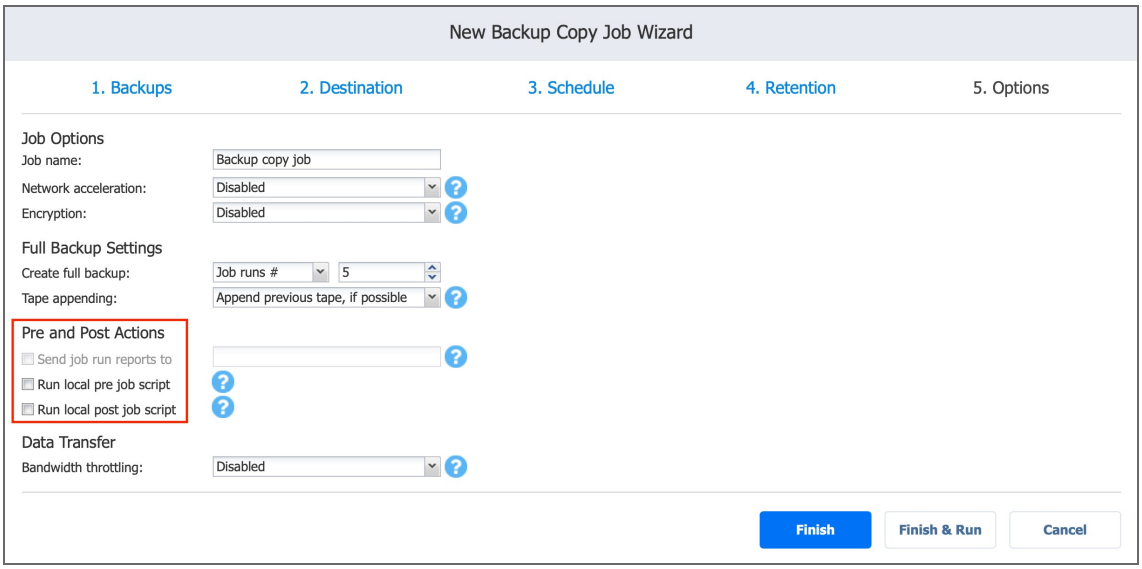

Email Notifications

NAKIVO Backup & Replication can send email notifications on job completion status to specified recipients. This feature complements global notifications and provides you with the ability to configure notifications on a per-job level.

To enable this option, configure your Email settings.

To send email notifications, do the following:

- In the *Pre- and Post- Actions* section, select the **Send job run reports to** option and specify one or more email addresses in the text box. Separate multiple email addresses with a semicolon.

Pre Job Script

To run a script before the product begins copying backups, do the following:

1. Place a script file on the machine on which the Director is installed.
2. In the *Pre and Post Actions* section, select the **Run local pre job script** option and click the **settings** link. Specify the following parameters in the dialog that opens:

- **Script path**: Specify a local path to the script on the machine on which the Director is installed. Script interpreter should be specified.
  Example (Windows): *cmd.exe /c D:\script.bat*
  Example (Linux): *bash /root/script.sh*
- **Job behavior**: Choose either of the following job behaviors in relation to script completion:
  - **Wait for the script to finish**: If this option is selected, tape backup will not be started until the script is completed.
  - **Do not wait for the script to finish**: If this option is selected, the product will run the script and will start copying backups at the same time.
- **Error handling**: Choose either of the following job behaviors in relation to script failure:
  - **Continue the job on script failure**: If this option is selected, the job will perform tape backup even if the script has failed.
  - **Fail the job on script failure**: If this option is selected and the script fails, the job will be failed and tape backup will not be performed.

Post Job Script

To run a script after the product has finished copying all backups, do the following:

1. Place a script file on the machine on which the Director is installed.
2. In the *Pre and Post Actions* section, select the **Run local post job script** option and click the **settings** link. Specify the following parameters in the dialog box that opens:

- **Script path**: Specify a local path to the script on the machine on which the Director is installed. Script interpreter should be specified.

  Example (Windows): *cmd.exe /c D:\script.bat*

  Example (Linux): *bash /root/script.sh*
- **Job behavior**: Choose either of the following job behaviors in relation to script completion:
  - **Wait for the script to finish**: If this option is selected, the job will be in the "running" state until the script is completed.
  - **Do not wait for the script to finish**: f 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.

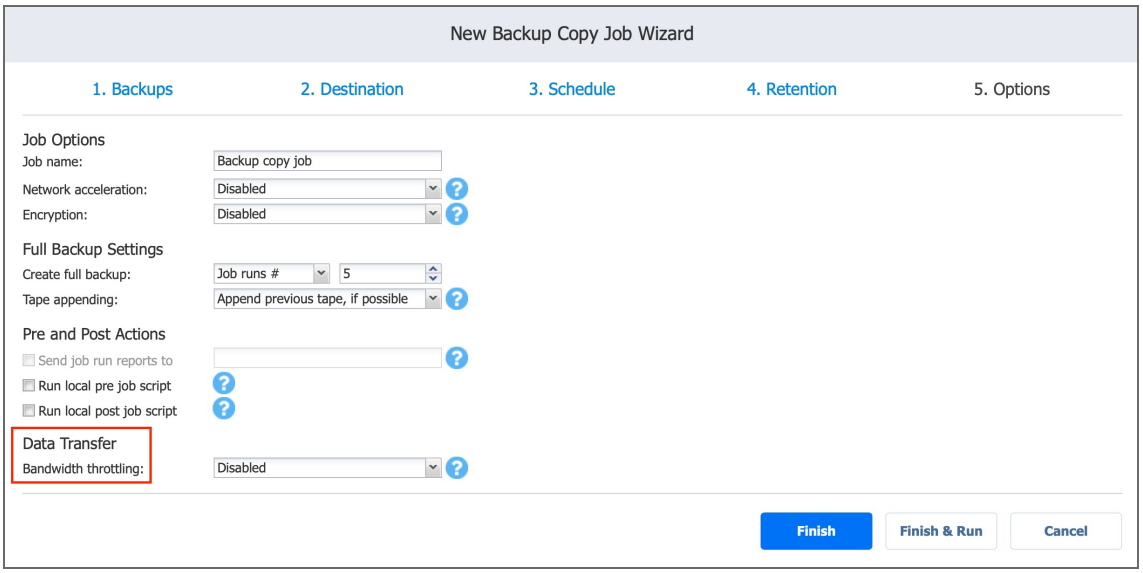

Bandwidth Throttling

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 281 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 281 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.

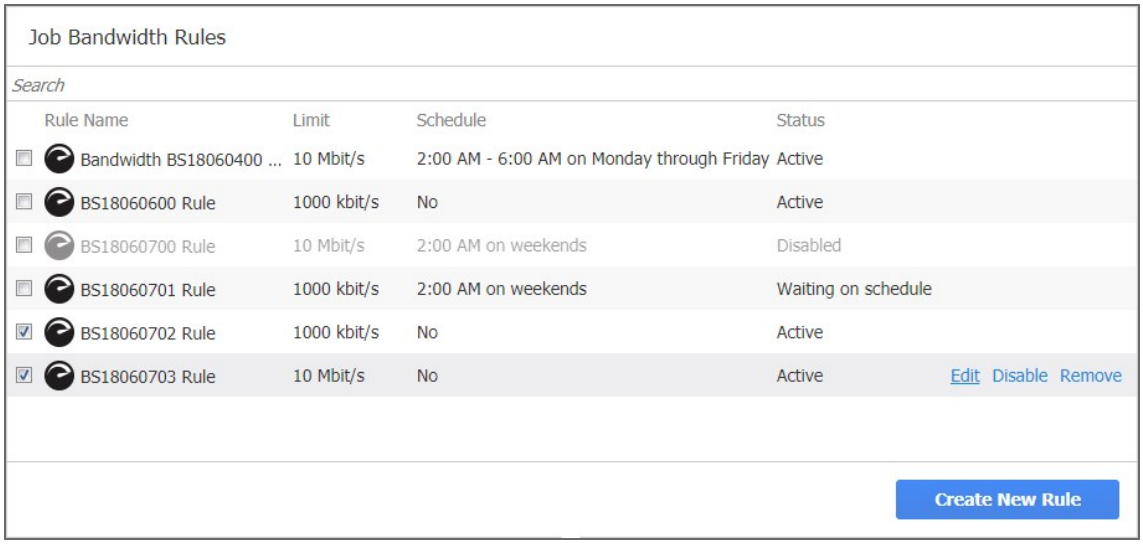

Completing Tape Backup Wizard

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 250 for details.

Deleting Backups

With NAKIVO Backup & Replication, you can permanently delete a backup (including all of its recovery points) if it is available in a Backup Repository. Also, you can delete specific recovery points of a backup without affecting any of the others.

Note

You can only delete a backup if you have deleted the corresponding backup job or edited the backup job so that it does not include the backup's source VM or a 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. Click **Settings** in the main menu of the product.
2. Go to the **Repositories** tab and click a Backup Repository.
3. Click the backup you want to delete.
4. In the backup page, click **Delete**.
5. Click **Delete** in the dialog box that opens.

Note

For a forever-incremental Backup Repository, the space occupied by the backup will be marked as "free" and reused by new data blocks on the next job runs, though the actual size of the Backup Repository may not change. To free up the space occupied by the backup, you can reclaim free space.

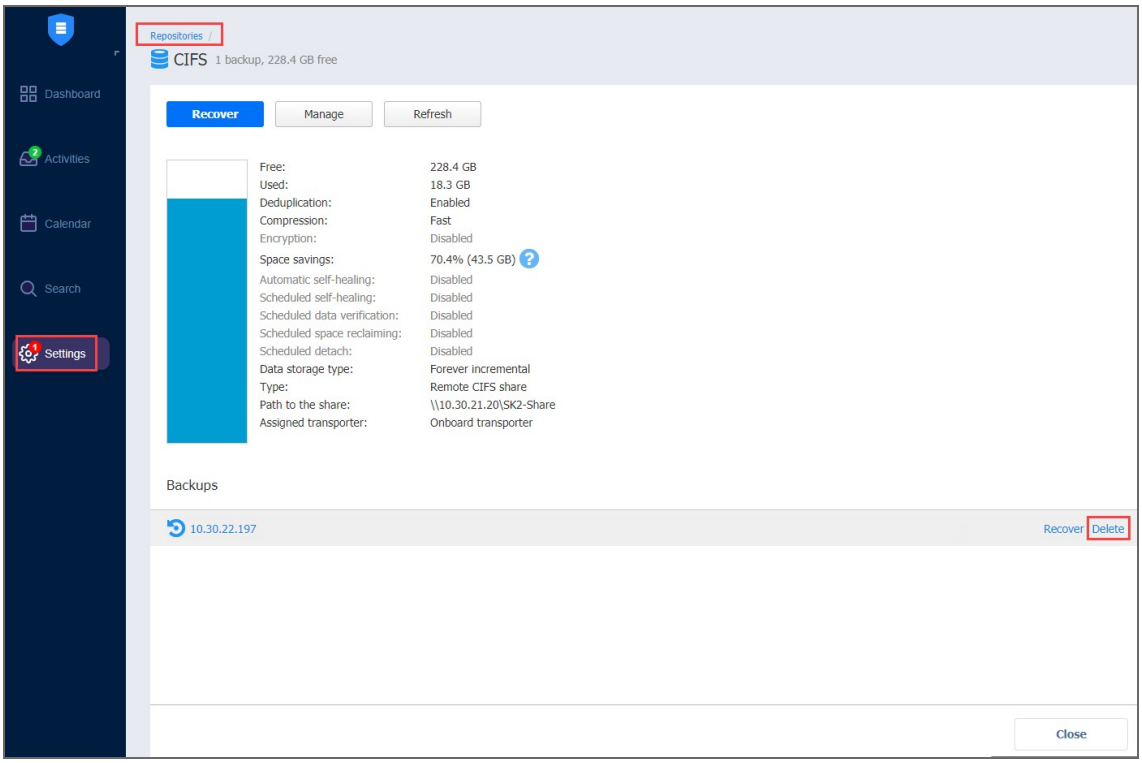

Deleting Backups in Bulk

To permanently delete a number of backups that match particular criteria, follow the steps below:

1. Click **Settings** in the main menu.
2. Go to the **Repositories** tab and hover over a Backup Repository.
3. Click **Manage** and then click **Delete backups in bulk**.

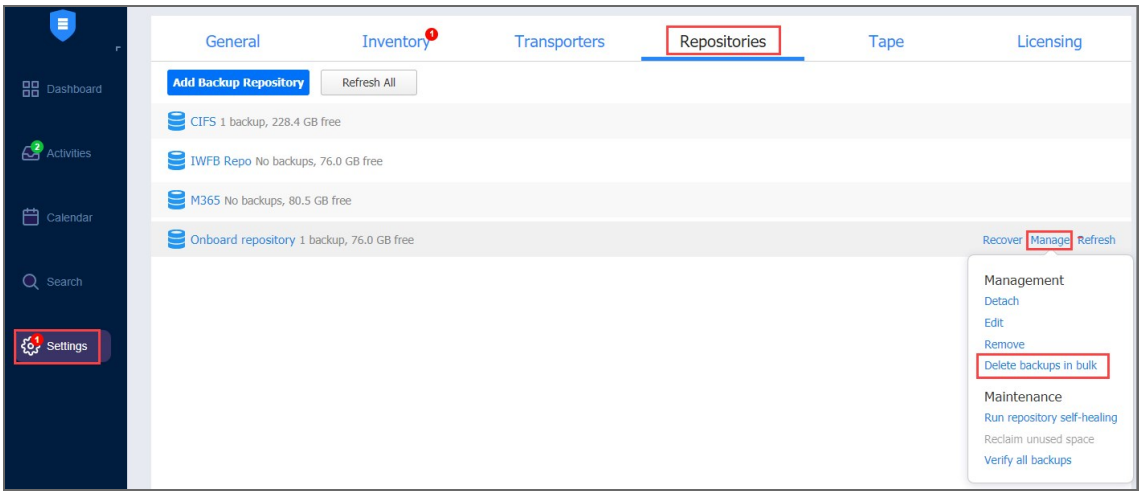

The **Bulk Delete Backups** dialog box opens.

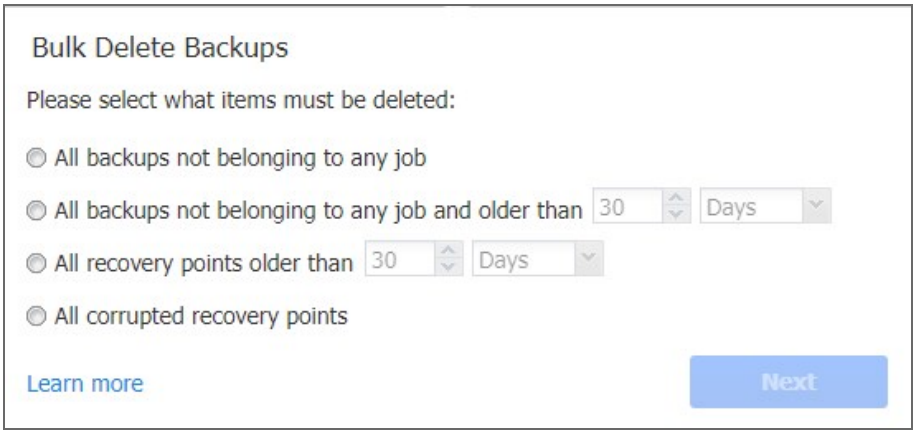

4. Select one of the available options:
   - All backups not belonging to any job
   - All backups not belonging to any job and older than X <time_units>, where X is an integer and <time_units> is either Days, Weeks, or Months

     The dialog shows the number of backups to be deleted.
5. Click **Next**.
6. A list of backups/recovery points to be deleted opens. Click **Delete**.

Note

For a forever-incremental Backup Repository, the space occupied by the backup will be marked as “free” and reused by new data blocks on the next job runs, though the actual size of the Backup Repository may not change. To free up space occupied by the backup, you can reclaim free space.

Deleting Recovery Points

You can delete either a single recovery point, all corrupted recovery points, or all recovery points older than the specified number of days.

Notes

1. If a backup is used by a VM/Amazon EC2/physical machine backup job, you cannot delete its last recovery point.
2. A recovery point cannot be deleted while the source VM/physical machine backup job is running.

Deleting a Single Recovery Point

1. Click **Settings** in the main menu.
2. Go to the **Repositories** tab.
3. Click a Backup Repository.
4. Click the backup where you want to delete a recovery point.
5. Hover over the recovery point you want to delete.
6. Click **Delete**.

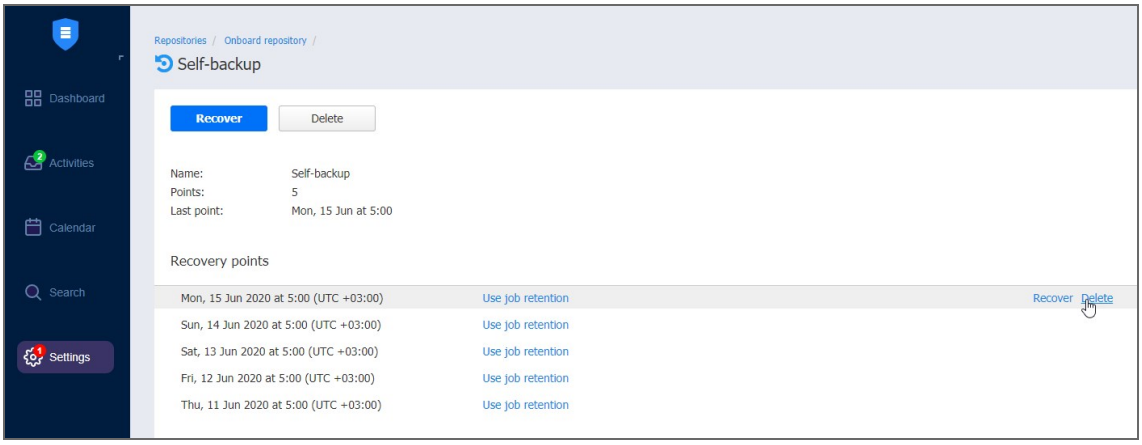

7. Click **Delete Recovery Point** in the message box that opens.

Note

For forever-incremental Backup Repository, the space occupied by the recovery point will be marked as "free" and reused by new data blocks on the next job runs, though the actual size of the Backup Repository may not change. To free up space occupied by the recovery point, you can reclaim free space.

Bulk Recovery Points Deletion

1. Click **Settings** in the main menu of the product.
2. Go to the **Repositories** tab.
3. Click a Backup Repository.
4. Click **Manage**.
5. Click **Delete backups in bulk**.
6. The **Bulk Delete Backups dialog** box 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 will be deleted.

      Note

      The following deletion exclusions are applicable:

        - **For forever-incremental repositories**: if all recovery points of a backup match the deletion criteria, the latest recovery point – no matter whether it is corrupted or not – will not be deleted.
        - **For incremental-with-full backups repositories**:
            - Recovery points that are older than the end of the time interval, but have dependent recovery points that are newer than the beginning of the time interval, will not be deleted.
            - If all recovery points of a backup match the deletion criteria, the latest full recovery point – no matter whether it is corrupted or not – will not be deleted.
    - **All corrupted recovery points**: If selected, all recovery points that are corrupted will be deleted. For incremental-with-full Backup Repositories, this will delete all recovery points that are dependent on corrupted recovery points. The following deletion exclusions are applicable:
        - **For forever-incremental repositories**: if a backup is used by a backup job and all its recovery points are corrupted, the latest recovery point will not be deleted.
        - **For incremental-with-full-backups repositories**: if all recovery points of a backup are corrupted or depend on a corrupted recovery point, and match the deletion criteria, the latest full recovery point will not be deleted.

          Note

          This option is unavailable for Microsoft 365 backups.
7. The **Bulk Delete Recovery Points** dialog box opens displaying the list of recovery points to be deleted. Click **Delete** to confirm the deletion of recovery points.

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.

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 647
- "Granular Recovery" on page 598
- "Planning Disaster Recovery" on page 694

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 VMware, Microsoft Hyper-V, and AWS virtual environments. You can also recover Microsoft Exchange emails, and Microsoft Active Directory and Microsoft SQL Server objects directly from compressed and deduplicated backups.

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

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 23 for mode 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 page 600
- “File Recovery Wizard: Backup” on page 602
- “File Recovery Wizard: Recovery Server” on page 603
- “File Recovery Wizard: Files” on page 605
- “File Recovery Wizard: Options” on page 608

Opening File Recovery Wizard

You can start the recovery process either 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 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**.

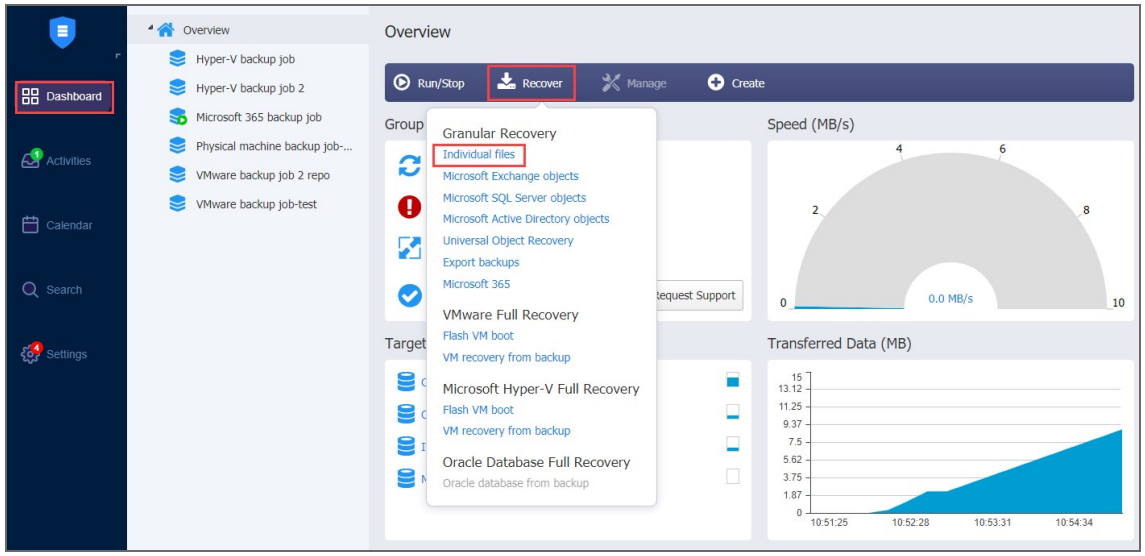

Starting File Recovery from Backup Repository

To start file recovery from a Backup Repository, do the following:

1. Go to the main menu of NAKIVO Backup & Replication 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 **Individual Files**. The **File Recovery Wizard** opens.

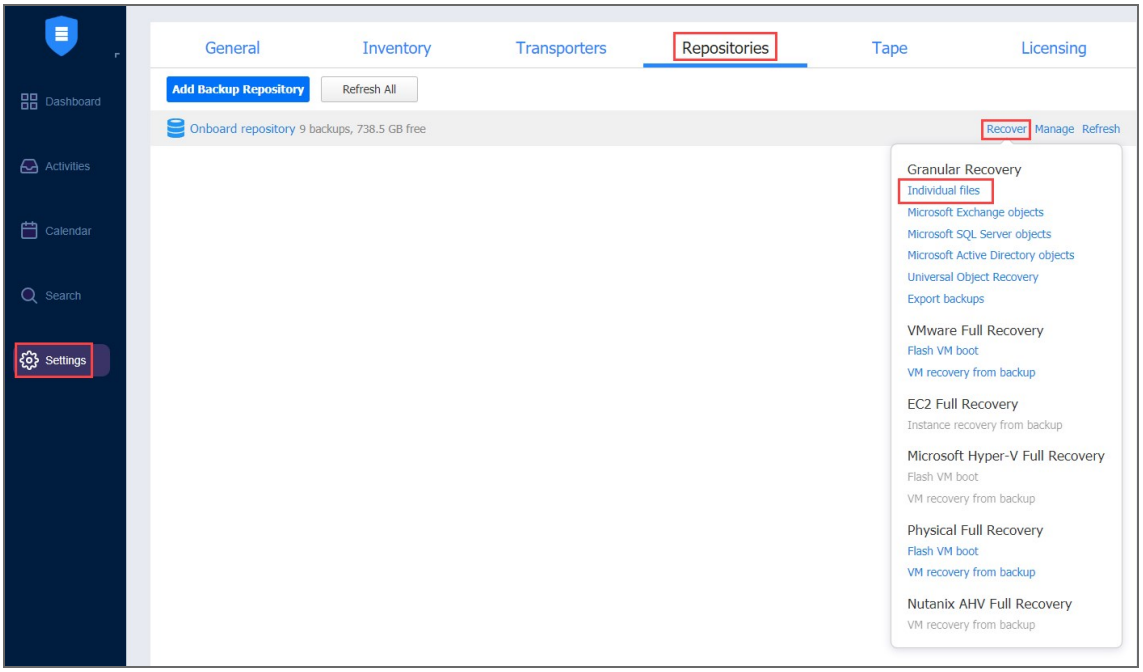

File Recovery Wizard: Backup

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.

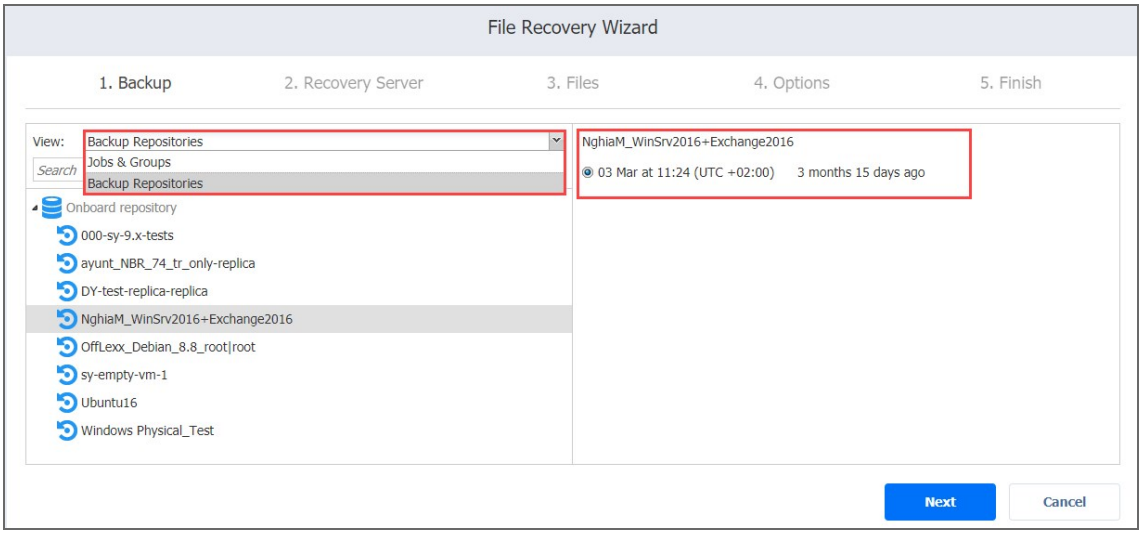

Click **Next** to go to the next page of the wizard.

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 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. **Username**: Enter a username with administrative privileges for the recovery server entered above. Refer to "Feature Requirements" on page 120 for a full list of requirements for recovering files to server.
   e. **Password**: Enter a user password.
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**.

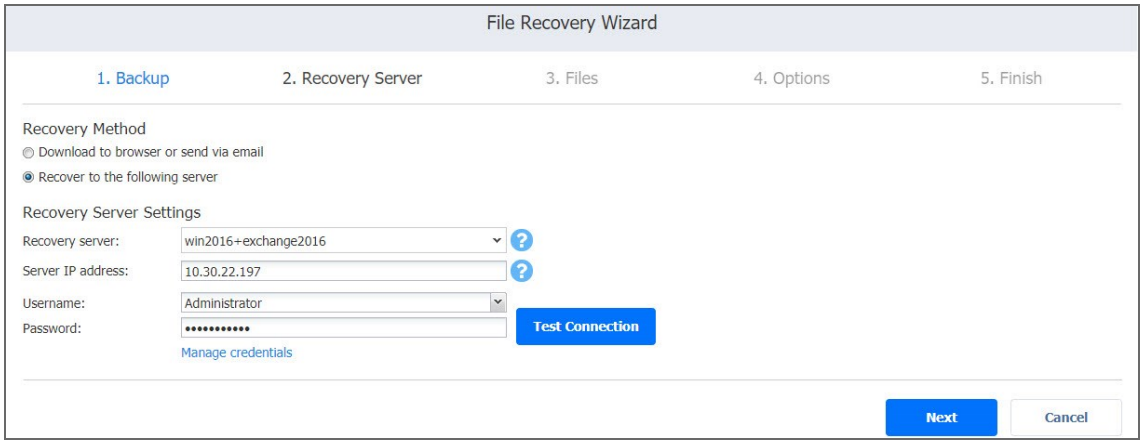

After NAKIVO Backup & Replication prepares a recovery point, the next page of the wizard opens.

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**.

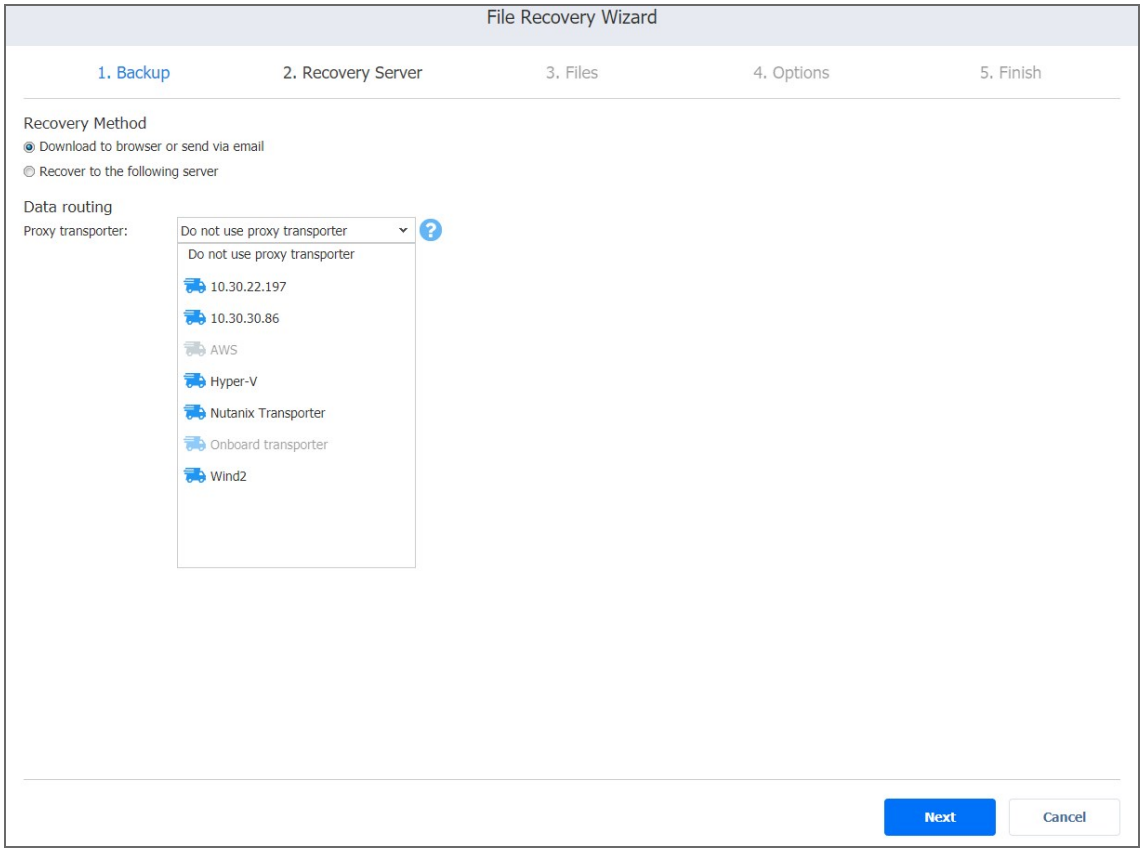

Note

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.

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**.

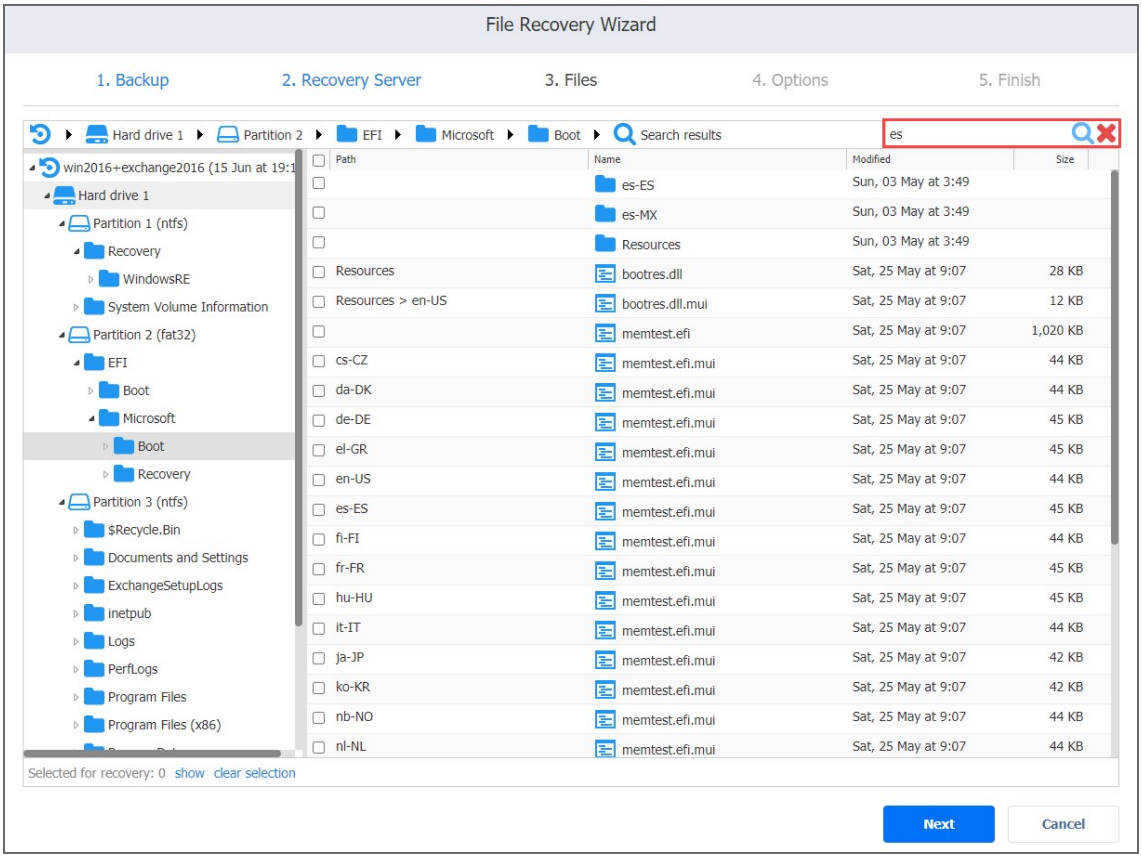

Notes

- 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:

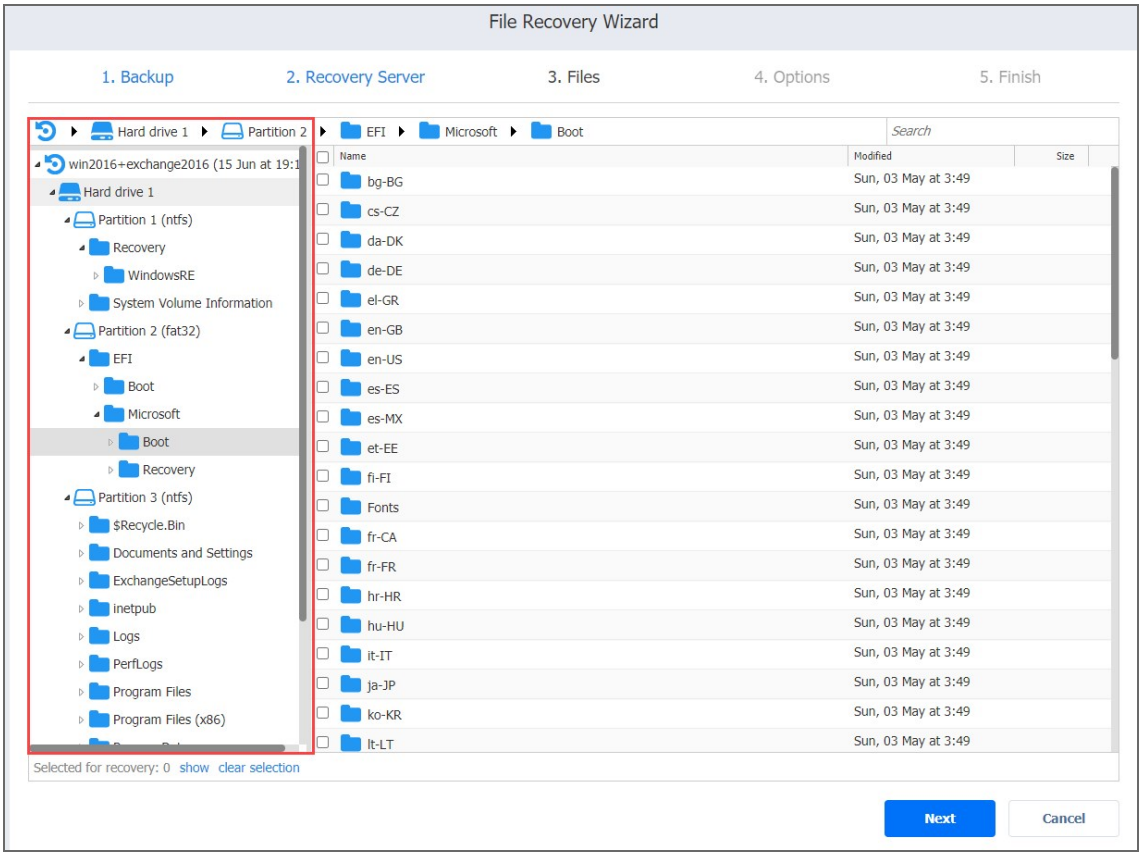

If a VM backup contains Linux LVM volumes or Windows dynamic disks, the navigation pane will display these logical groups in addition to all hard drives available in the VM backup. If a hard drive does not contain any partitions and servers as a part of a Linux LVM volume or a Windows dynamic disk, this 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.

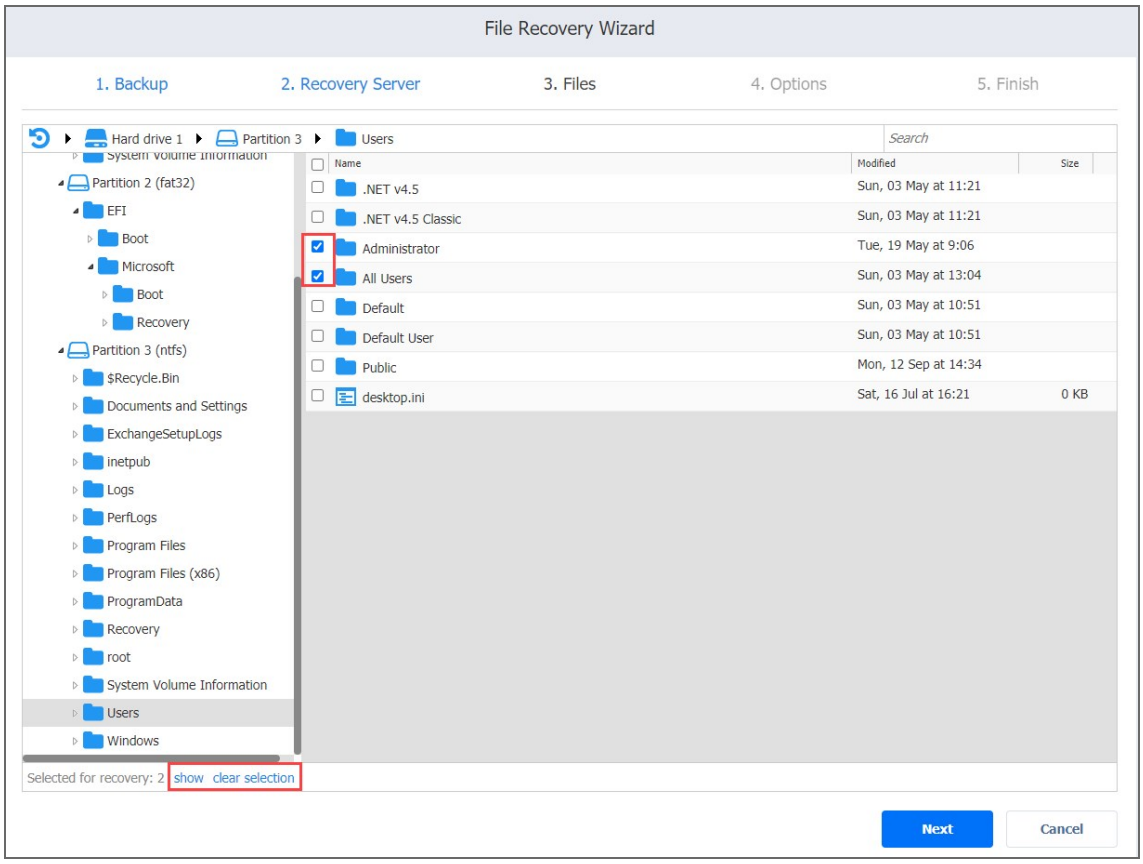

Click **Next** to go to the next page of the wizard.

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.

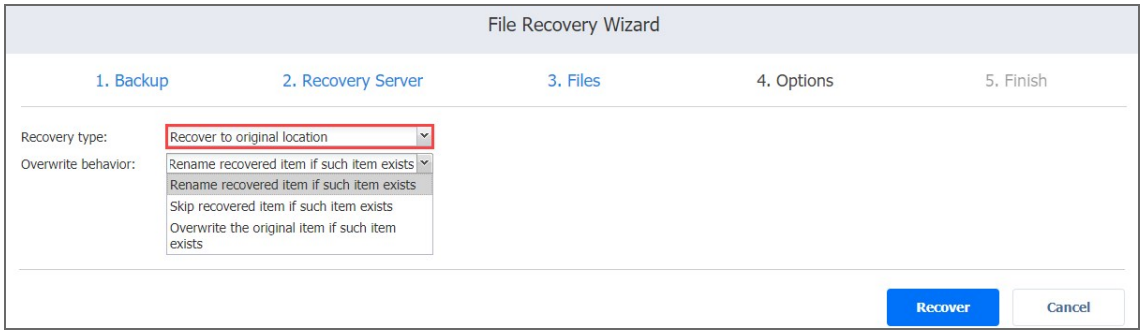

Recovering Files to a Custom Location

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**.

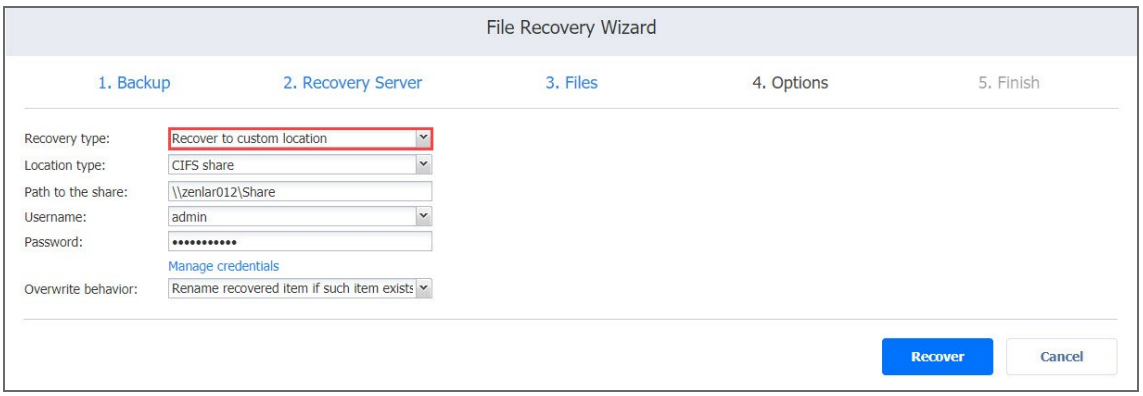

Downloading Files to Browser or Sending Files via Email

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**.

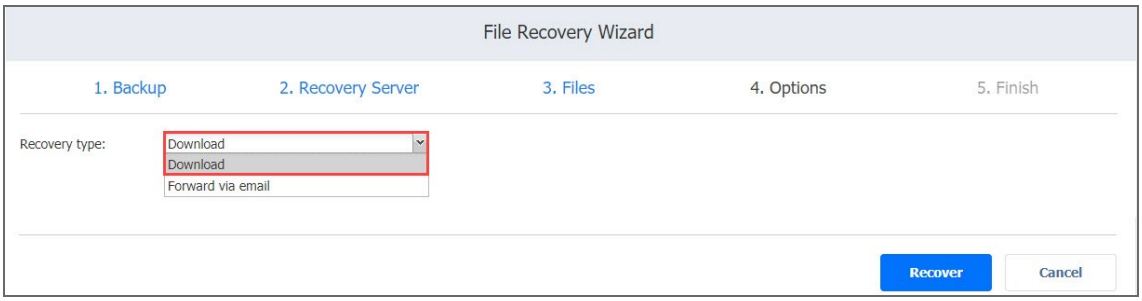

When the download has finished successfully, the archive with the recovered items appears in the browser downloads folder.

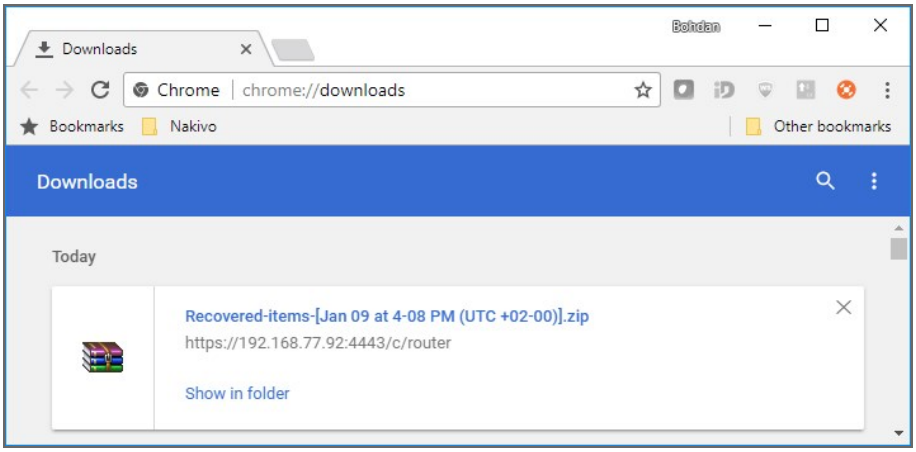

Forwarding Files via Email

Note

To use this recovery type, your Email settings must be properly configured in the NAKIVO Backup & Replication Configuration. Refer to "Email Notifications" on page 287 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**.

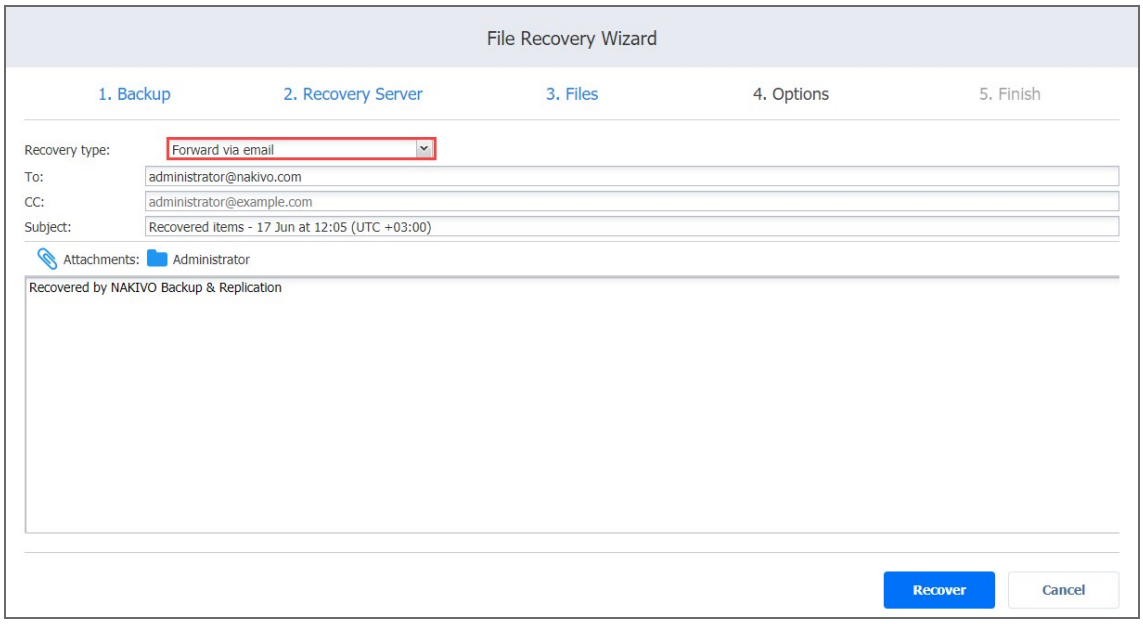

Object Recovery for Microsoft Exchange

The object recovery feature in NAKIVO Backup & Replication allows you to browse, search, and recover Microsoft Exchange emails directly from compressed and deduplicated backups. Recovery can also be performed back to the source or any other location including CIFS share. The Object Recovery for Microsoft Exchange feature is agentless, works right out of the box, and does not require creating a special lab or running a special backup type.

Refer to the following topics for more information:

- "Starting Object Recovery for Microsoft Exchange" on page 613
- "Object Recovery Wizard for Microsoft Exchange: Backup" on page 615
- "Object Recovery Wizard for Microsoft Exchange: Recovery Method" on page 616
- "Object Recovery Wizard for Microsoft Exchange: Objects" on page 617
- "Object Recovery Wizard for Microsoft Exchange: Options" on page 619

Starting Object Recovery for Microsoft Exchange

You can start the recovery process either 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).

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**.

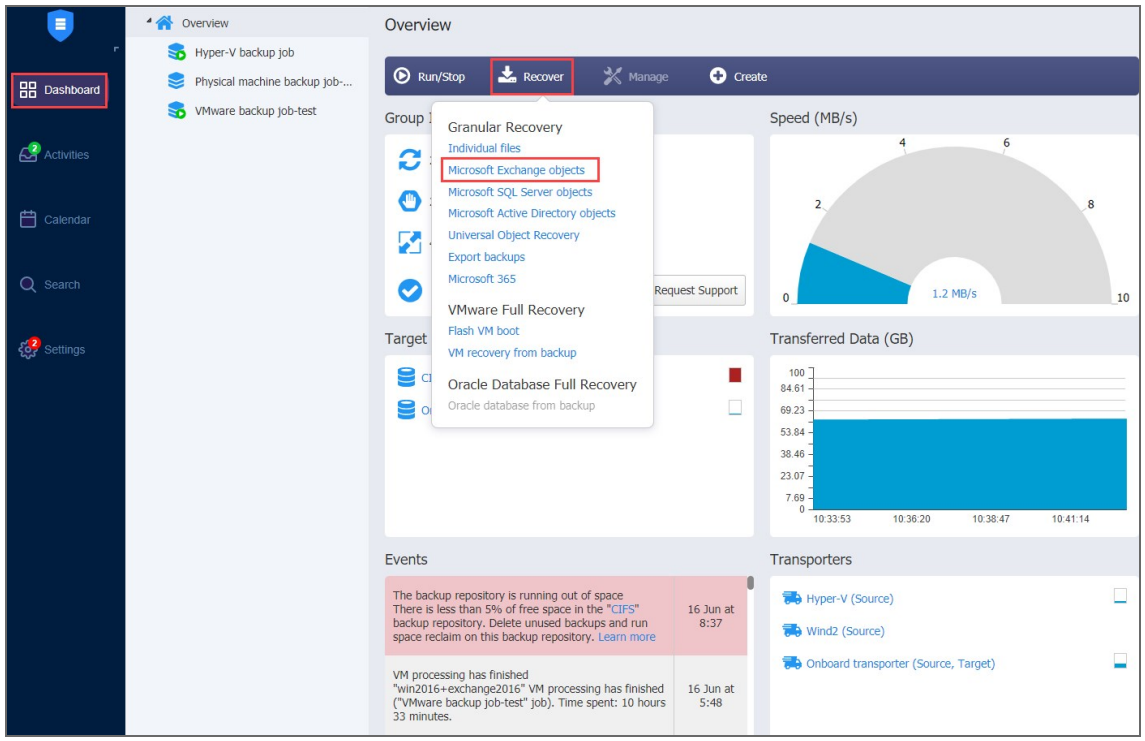

Starting Object Recovery for Microsoft Exchange from Backup Repository

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**.

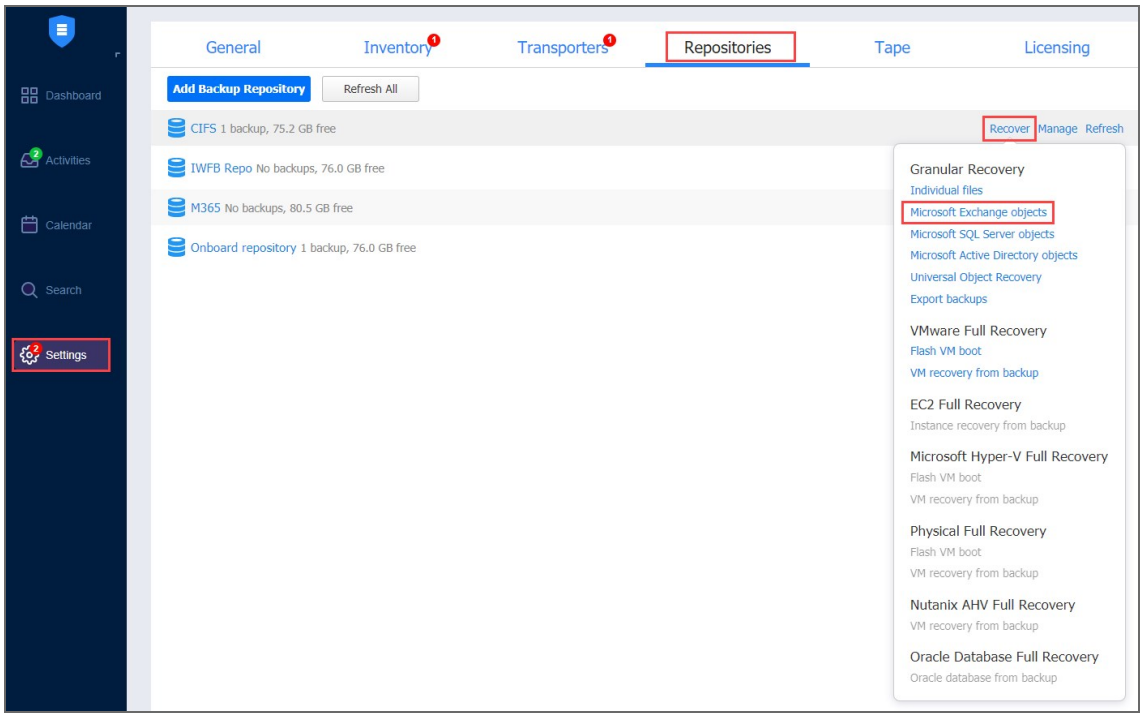

The **Object Recovery Wizard for Microsoft Exchange** opens.

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.

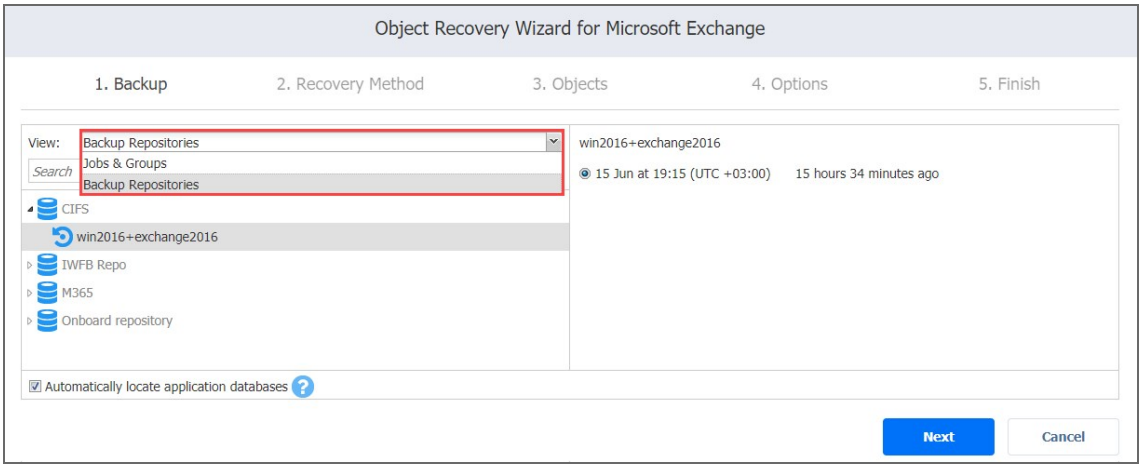

By default, NAKIVO Backup & Replication automatically searches the selected recovery point for Microsoft Exchange databases (files with .edb extension) from which application objects can be recovered. This process can take a few minutes. If you want to manually specify the location of the database file, 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.
- **Username**: Enter the username to log in to the VM.
- **Password**: Enter the password to log in to the VM.
- **Test Connection**: Click this button to verify the credentials. You won't be able to proceed until after the connection has been successfully established.
- Click **Next** to proceed to the next step.

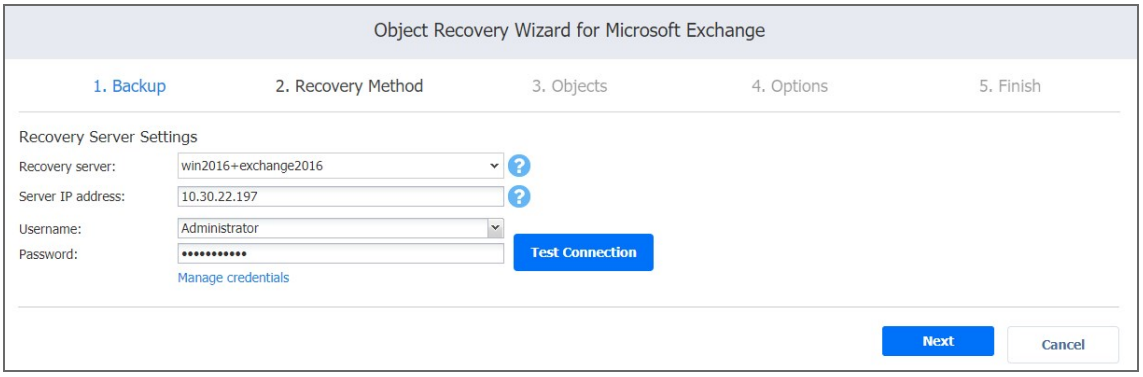

Info

To download items to a browser or forward them via email, enable the **system.exchange.enable.direct.recovery** setting in the Expert tab. Note that contacts and calendar items will not be recoverable with this enabled setting.

Object Recovery Wizard for Microsoft Exchange: Objects

On the **Objects** page of the wizard, select Microsoft Exchange objects for recovery. Proceed as described in the following sections:

- Searching for Microsoft Exchange Objects
- Browsing Microsoft Exchange Objects
- Viewing Microsoft Exchange Objects
- Selecting Microsoft Exchange Objects to Recover

Searching for Microsoft Exchange Objects

NAKIVO Backup & Replication allows you to search for emails. The search functionality, however, has the following limitations:

- The product can search for emails only by email subject or email body
- If text formatting (such as “bold text”) is applied to a keyword that is searched for, the search may not find the keyword due to formatting conversion issues.
- The product does not create or maintain an index of the Exchange database contents. The search is performed on the fly and can take a long time to complete.

To speed up the search, perform the search within a particular folder, rather than in a mailbox.

To search for an email by its subject or body, type a word in the **Search** field and press **Enter**.

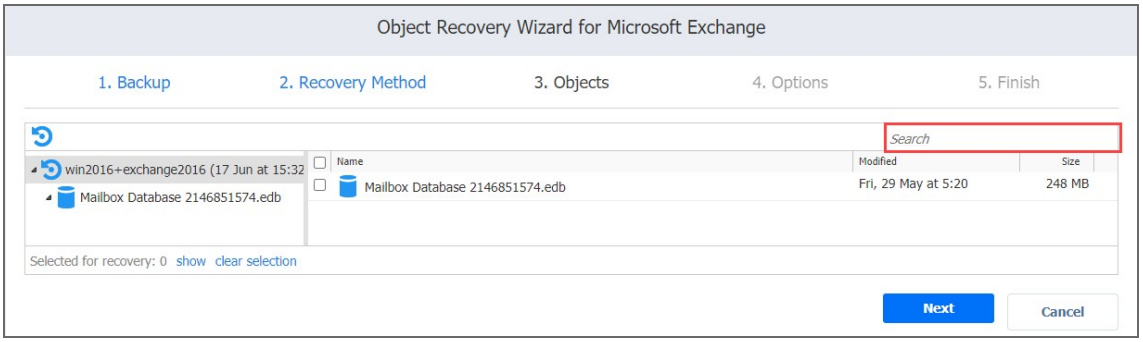

The search is performed starting from the point selected in the left (Navigation) pane. For example, if you have selected Mailbox Database > John Smith, the search will be performed only inside the John Smith mailbox.

Browsing Microsoft Exchange Objects

NAKIVO Backup & Replication scans the selected recovery point for Microsoft Exchange databases (files with “.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.

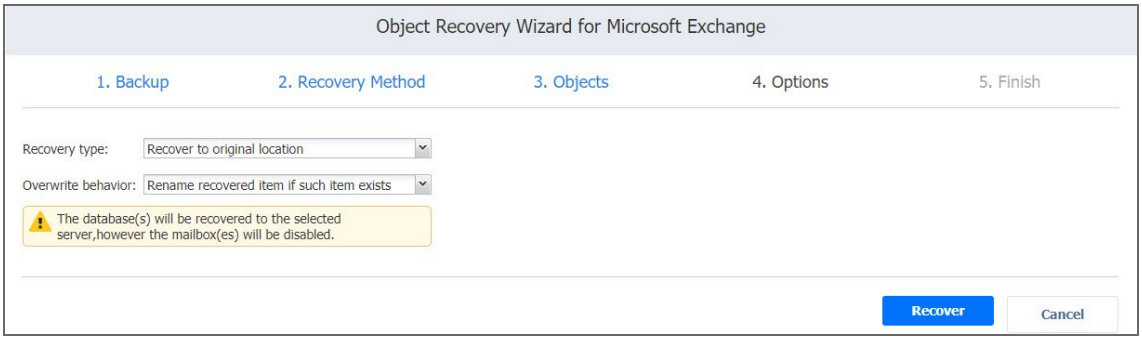

Recovering to a Custom Location

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.

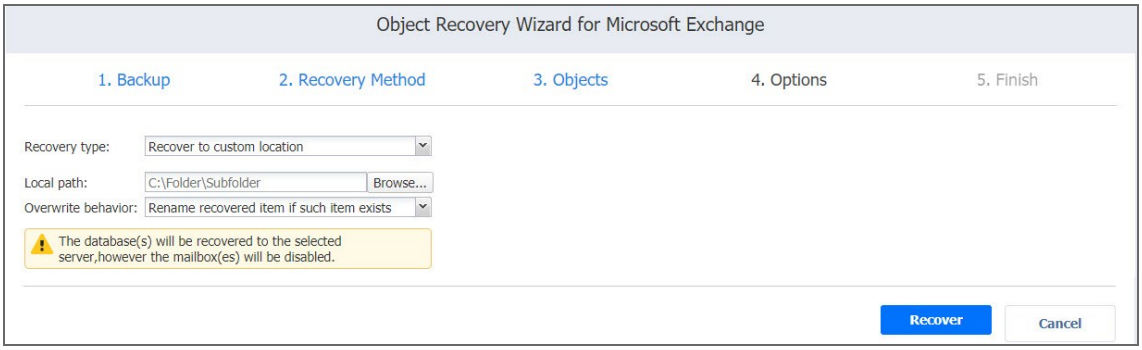

Exporting to a Custom Location

In the **Recovery type** drop-down list, you can choose **Export** to export Microsoft Exchange objects to a CIFS share. In this case, you'll need to provide the path to the share and 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**

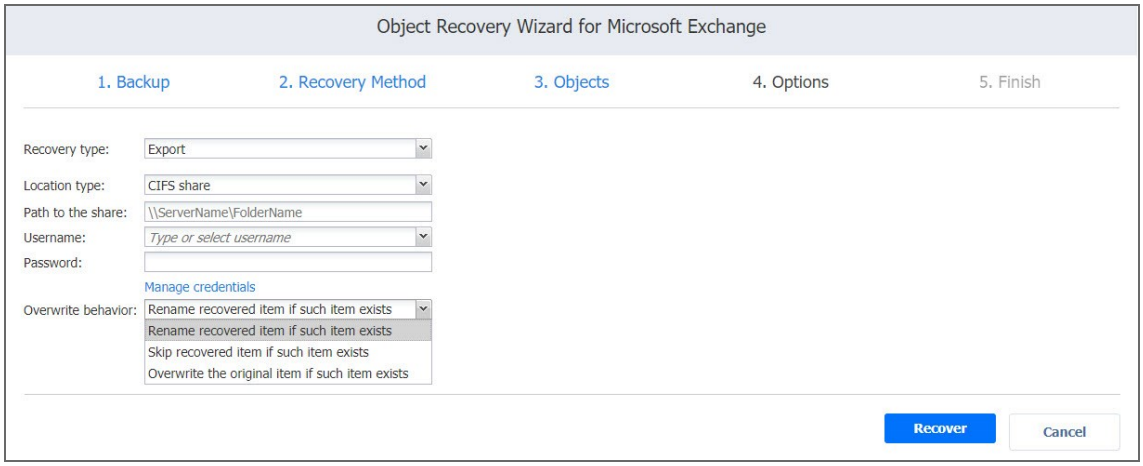

Click **Recover** to proceed with the recovery process. The **Finish** page is displayed. You cannot return to the previous pages of the wizard at this point, however, you can check the progress of the job execution by clicking the **Activities** link.

Object Recovery for Microsoft Active Directory

The instant object recovery feature allows you to browse, search and recover Microsoft Active Directory objects directly from compressed and deduplicated backups. This feature is agentless, works right out of the box, and does not require you to create a special lab or run a special type of backup. Microsoft Active Directory objects can be recovered in .ldif format and then be imported back to the Active Directory Server. Refer to the following topics for more information:

- "Starting Object Recovery for Microsoft Active Directory" on page 622
- "Object Recovery Wizard for Microsoft AD Server: Backup" on page 624
- "Object Recovery Wizard for Microsoft AD Server: Recovery Server" on page 625
- "Object Recovery Wizard for Microsoft AD Server: Objects" on page 626
- "Object Recovery Wizard for Microsoft AD Server: Options" on page 629

Starting Object Recovery for Microsoft Active Directory

You can start the recovery process either from the **Dashboard**, 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**.

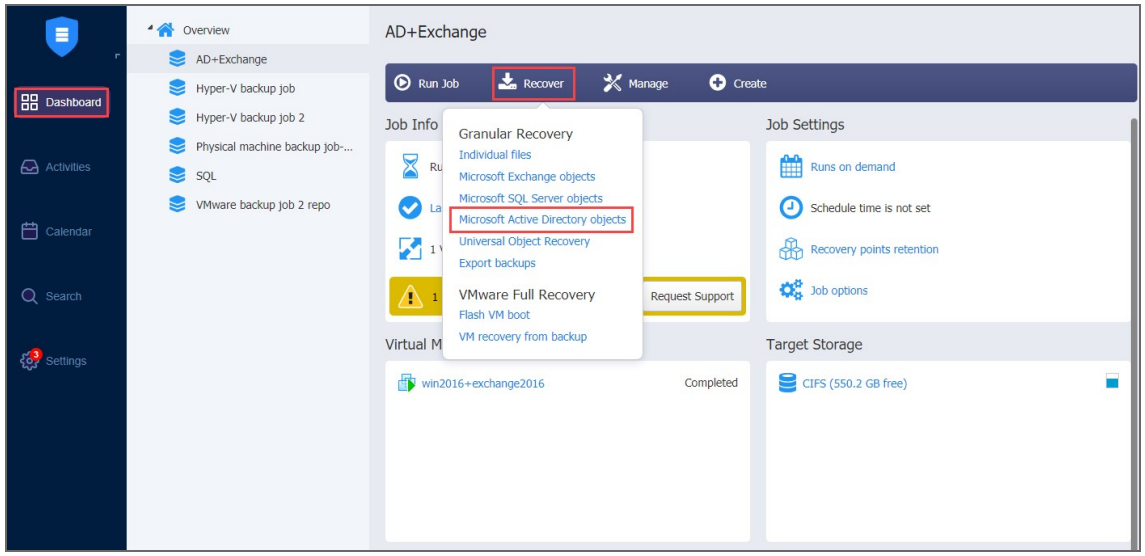

Starting Active Directory Object Recovery from a Backup Repository

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**.

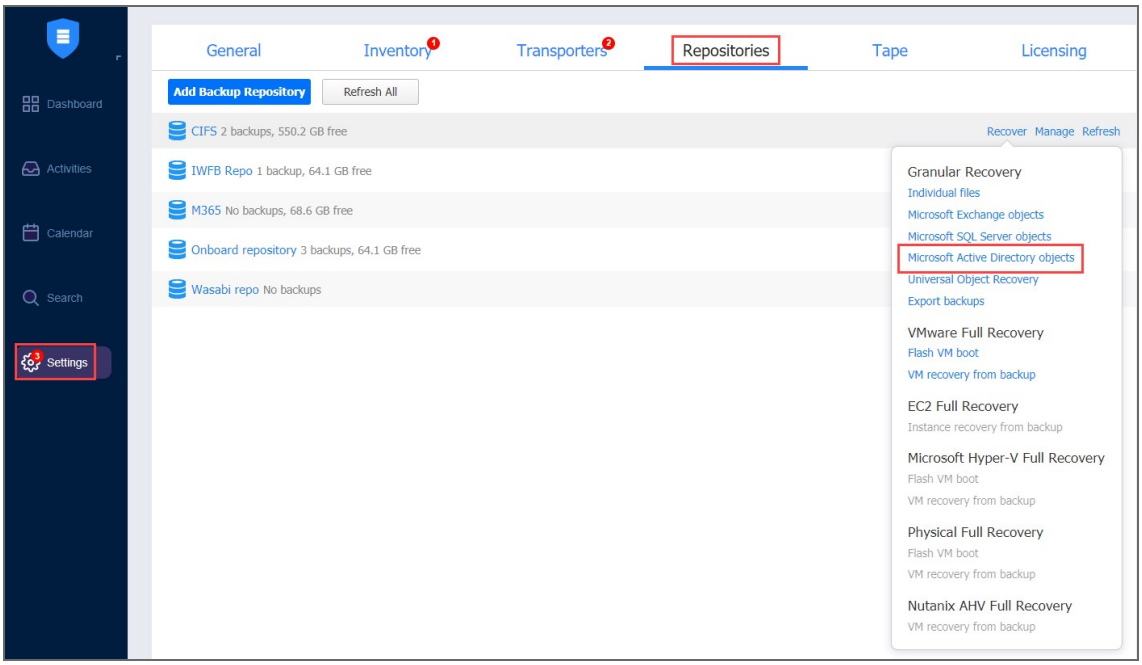

The **Object Recovery Job Wizard for Microsoft AD Server** opens.

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.

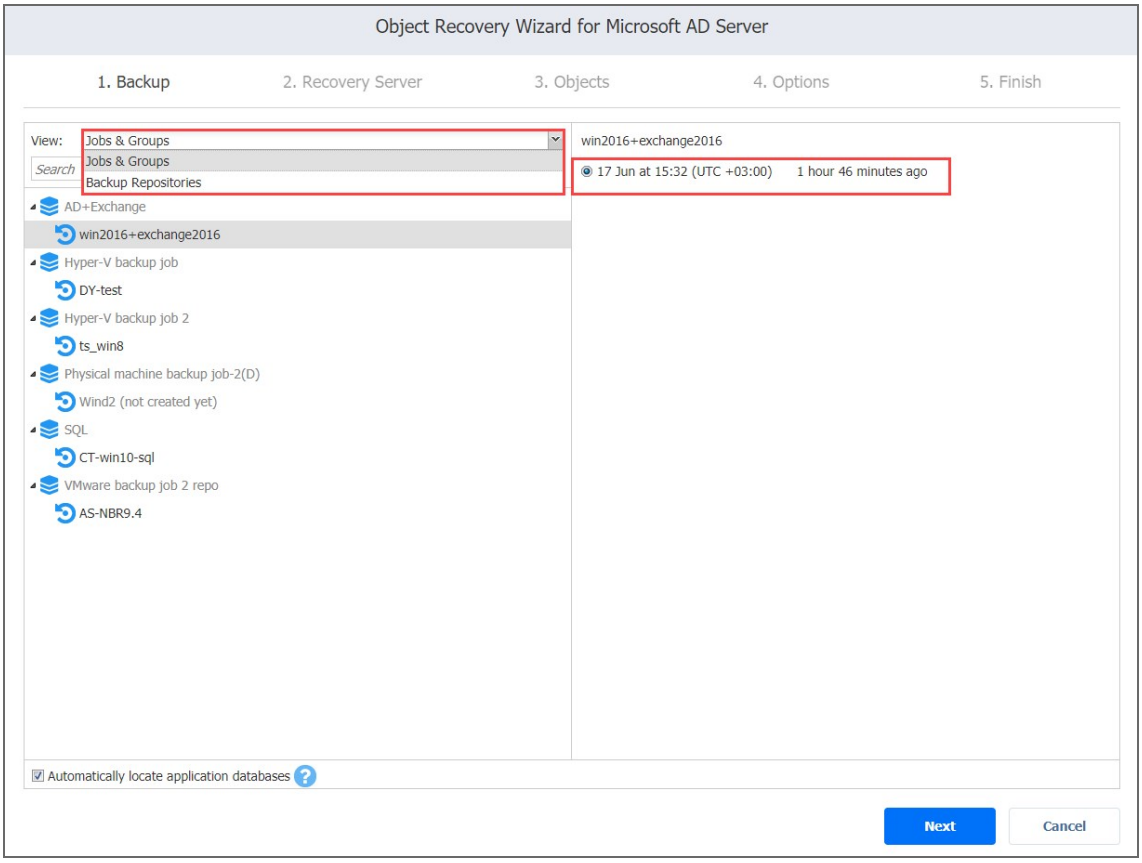

By default, NAKIVO Backup & Replication automatically searches the selected recovery point for the Microsoft Active Directory database from which application objects can be recovered. This process can take a few minutes. If you want to manually specify the location of the database file, deselect the **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.
- In the **Username** box, enter the name of a user with administrative privileges for the recovery server.
- In the **Password** box, enter a user password.
- Click the **Test Connection** button to test your credentials for the specified recovery server. If your credentials are correct, a checkmark appears to the right of the button.
- **Create snapshot before recovery:** When selected, a snapshot of the VM will be taken if recovery fails, and the VM will be reverted to this snapshot.
- Click **Next** to go to the next page of the wizard.

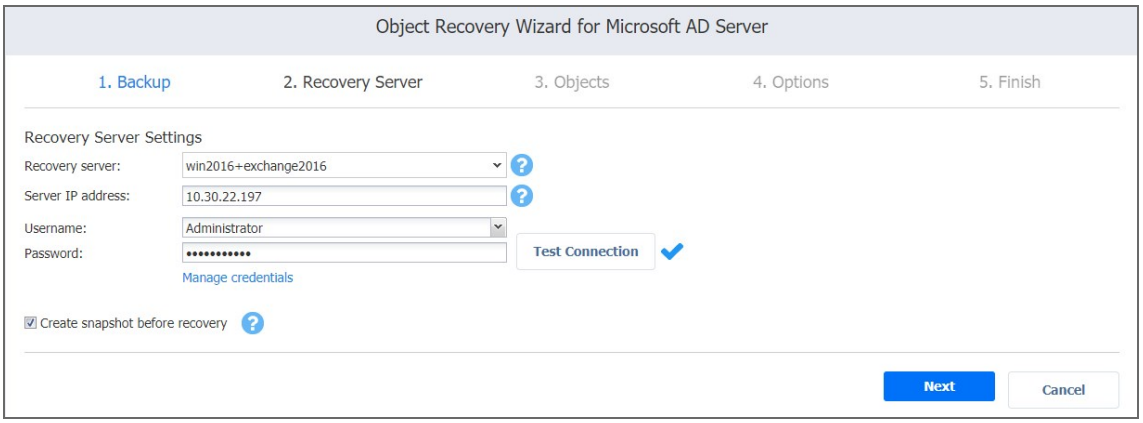

Object Recovery Wizard for Microsoft AD Server: Objects

On the **Objects** page of the wizard, select Active Directory objects you want to recover.

- Searching for Active Directory Objects
- Browsing Active Directory Objects
- Viewing Active Directory Objects
- Selecting Active Directory Objects to Recover

Searching for Active Directory Objects

NAKIVO Backup & Replication allows you to search Active Directory objects by name. To find an object by its name, enter a word in the **Search** box and press **Enter**

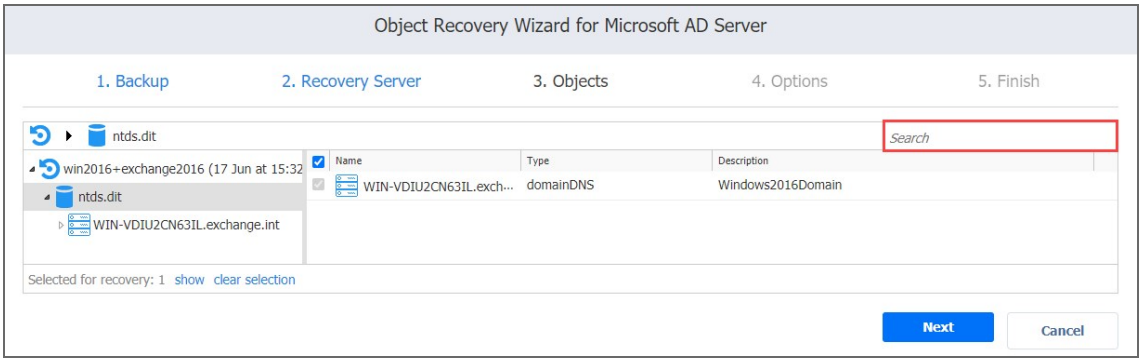

The search is performed starting from the point selected in the left (navigation) pane. For example, if you have selected the **Users** group, the search will only be performed inside the **Users** group.

Browsing Active Directory Objects

NAKIVO Backup & Replication scans the selected recovery point for Active Directory databases (files with ".edb" extension) and displays the list of identified databases in the left (Navigation) pane. To browse Microsoft Active Directory objects, simply expand the appropriate database in the left pane. You can also browse the tree by using the scroll bar.

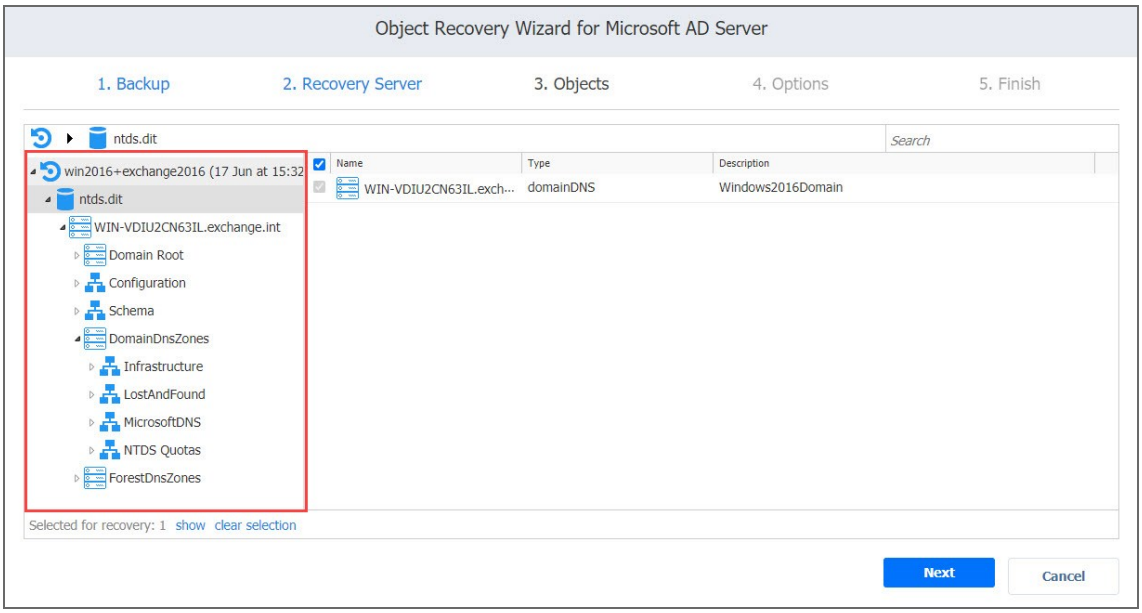

Viewing Active Directory Objects

To view a Microsoft Active Directory object, click the object. The object contents will be displayed. Use the close buttons to close the item.

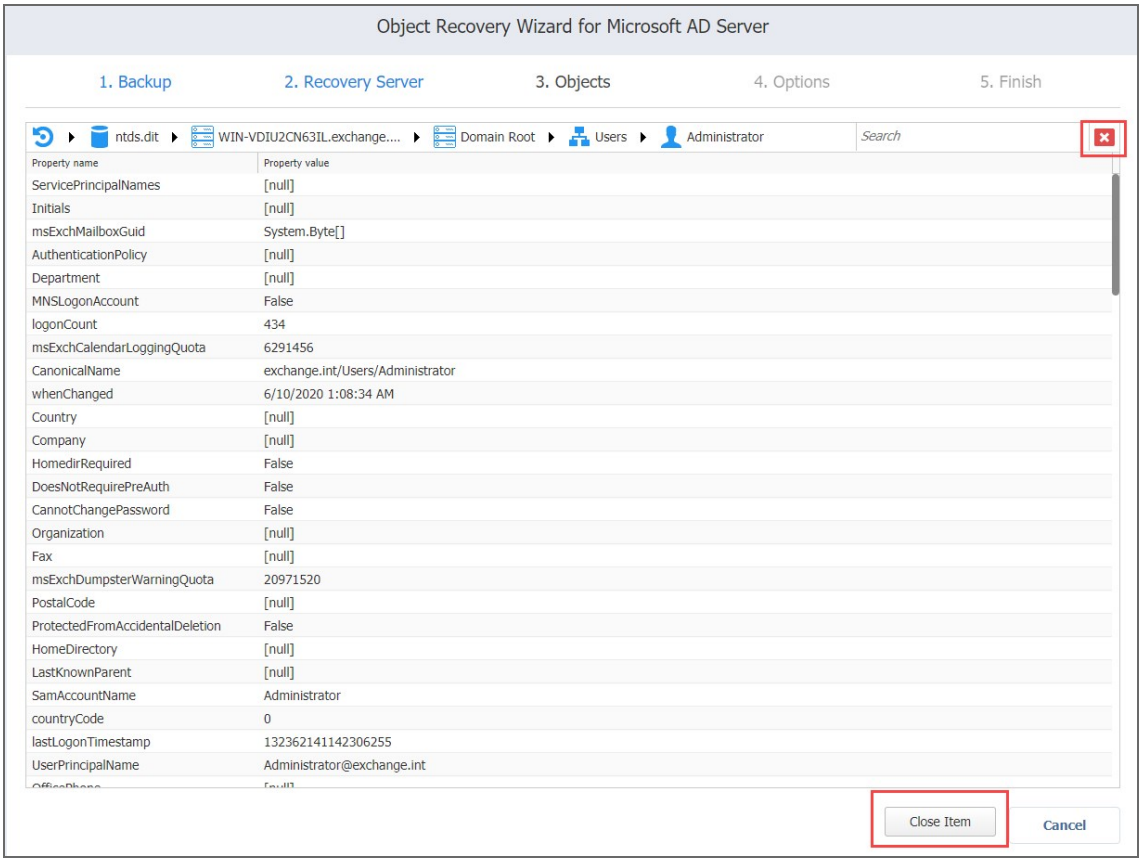

Selecting Active Directory Objects to Recover

In the **Contents** pane to the right, select a checkbox next to the items you want to recover. The number of items selected for recovery is displayed at the bottom of the wizard page. You can also:

- Click **show** to view the list of all items selected for recovery.
- Click **clear selection** to clear the list of items selected for recovery.
- Click **hide** to hide the list of items selected for recovery.

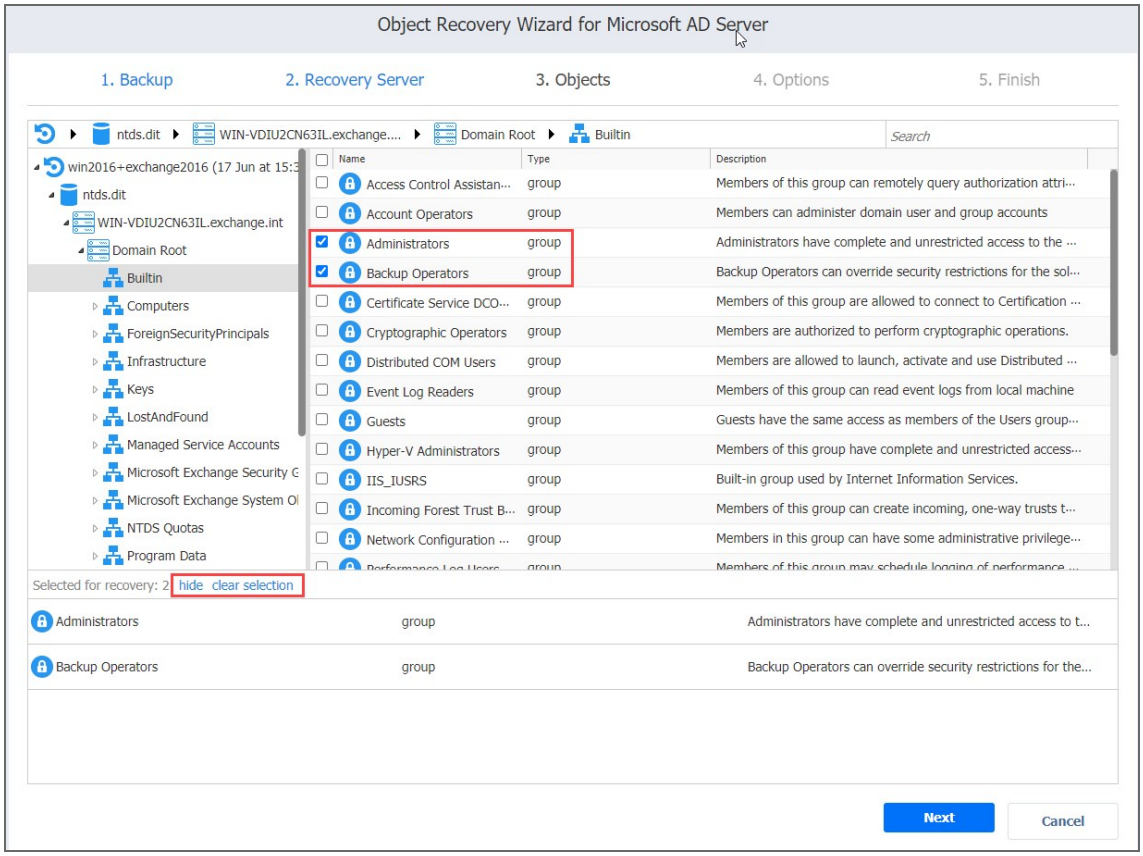

When ready with selecting Microsoft Active Directory objects for recovery, click **Next** to go to the next page of the wizard

Object Recovery Wizard for Microsoft AD Server: Options

On the **Options** page of the wizard, you can set up the following options for your object recovery job:

- Recovering Objects to the Original Location
- Exporting Active Directory Objects

Recovering Objects to the Original Location

Follow the steps below to recover objects of your Microsoft Active Directory server to the original location:

1. In the **Recovery type**list, select **Recover to original location**.
2. If you have selected multiple objects or container(s) that include one or more “user” objects, the **Recover of user object** list becomes available. Select either of the following options:
   - **User will be disabled**: If this option is selected, NAKIVO Backup & Replication disables all recovered “user” objects and the corresponding user accounts are disabled after importing these objects to Active Directory.
   - **User must change password at next log on**: If this option is selected, NAKIVO Backup & Replication generates a new password for each recovered “user” object. The passwords.txt file is added to the ZIP archive with recovered objects, and it contains the new passwords. After importing the “user” objects to Active Directory, corresponding users are forced to change the password on the next log on.

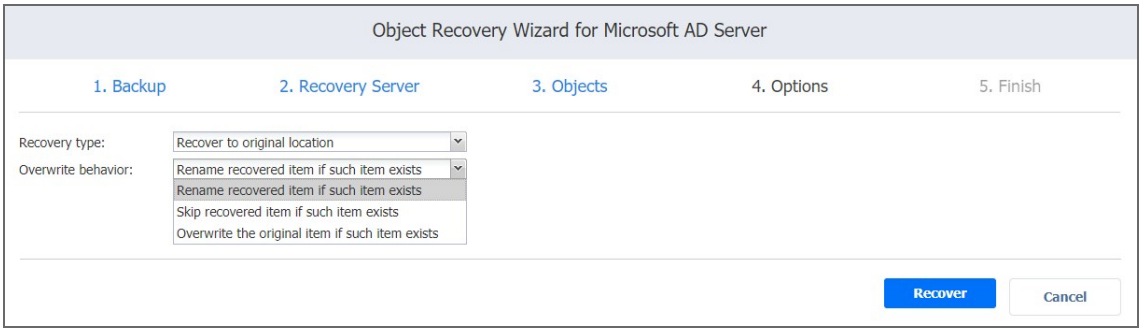

3. In the **Overwrite behavior** list, select what you wish to do if the recovered item conflicts with an existing one:
   - **Rename recovered item if such item exists**
   - **Skip recovered item if such item exists**
   - **Overwrite the original item if such item exists**
4. If you have chosen to recover a full database (`ntds.dit` file) on the **Recovery Server** pageput a check-mark in the**Stop Microsoft Active Directory instance before recovery** checkbox to stop the instance before the recovery process begins.This option is recommended for the safe recovery of Active Dir-

ectory objects.

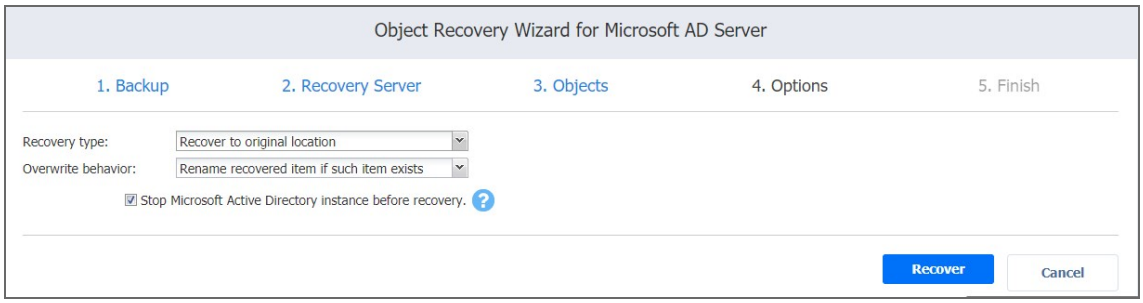

5. Click **Recover**.

Notes

- Some attributes may be skipped for the selected object(s) depending on the Active Directory system settings.
- In case the recovery process fails, theVM will be reverted to the snapshot taken on the **Recovery Server** page of the wizard

Exporting Active Directory Objects

Follow the steps below to export recovered objects of your Microsoft Active Directory server to a custom location:

1. In the **Recovery type** list, select **Export**. A number of options become available for setting up a custom location.
2. In the **Export location** list, select the appropriate location type:
    - **Local folder on Recovery Server:** If this option is selected, you will have to enter the path to a local folder on the recovery server in the **Local path** field.

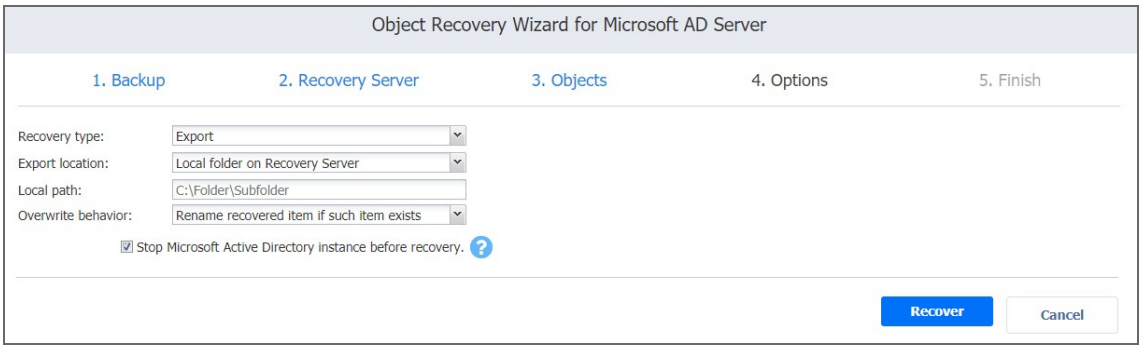

    - **CIFS share**: If this option is selected, enter the following values:
        1. **Path to the share**
        2. **Username**

3. **Password**

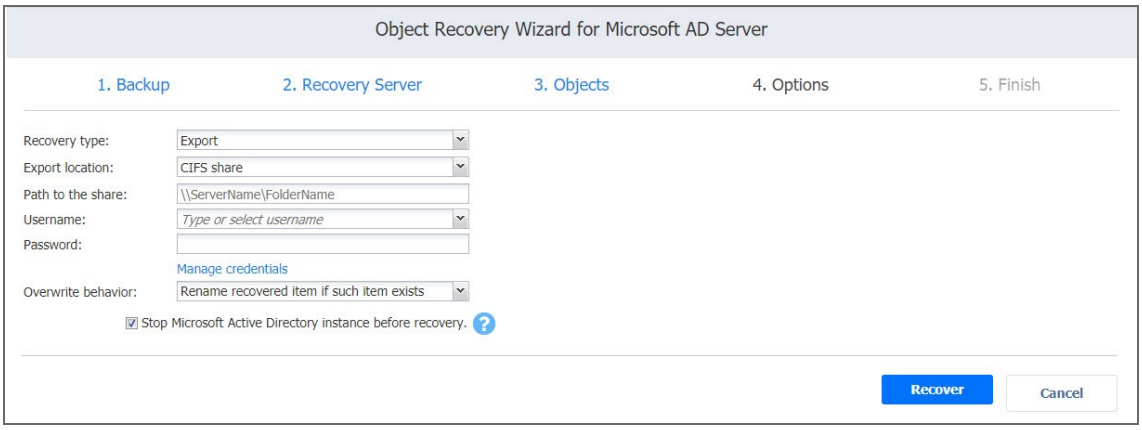

3. In the **Overwrite behavior** list, select what needs to be done if the recovered item conflicts with an existing item. Refer to the section above for an explanation.
4. Click **Recover**.

The **Finish** page of the wizard opens informing you that Microsoft Active Directory object recovery has started. To view the object recovery progress, 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 page 634
- "Object Recovery Wizard for Microsoft SQL Server: Backup" on page 636
- "Object Recovery Wizard for Microsoft SQL Server: Recovery Server" on page 637
- "Object Recovery Wizard for Microsoft SQL Server: Objects" on page 638
- " Object Recovery Wizard for Microsoft SQL Server: Options" on page 639

Starting Object Recovery for Microsoft SQL Server

You can start the recovery process either from the **Dashboard** 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**.

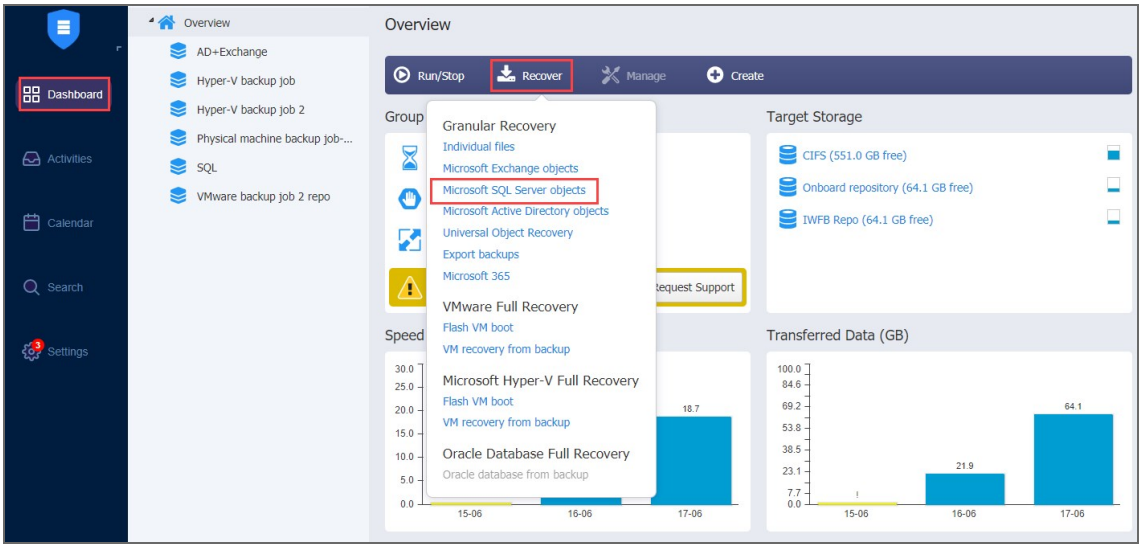

Starting SQL Server Object Recovery from Backup Repository

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.

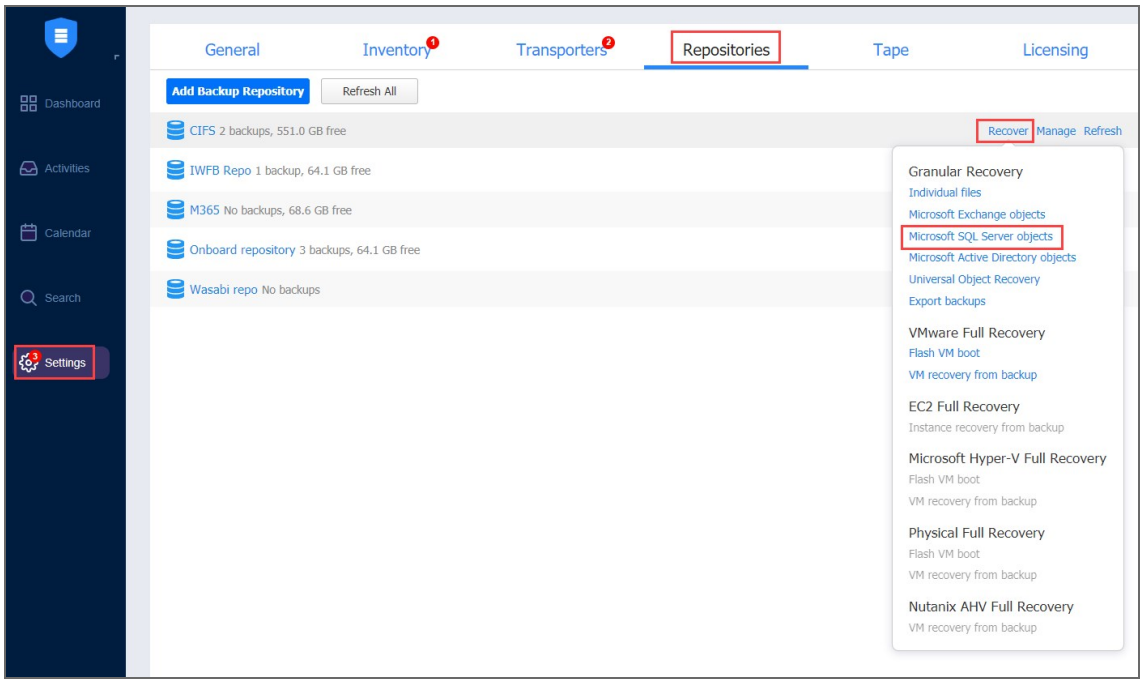

The **New Object Recovery Wizard for Microsoft SQL Server** opens.

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.

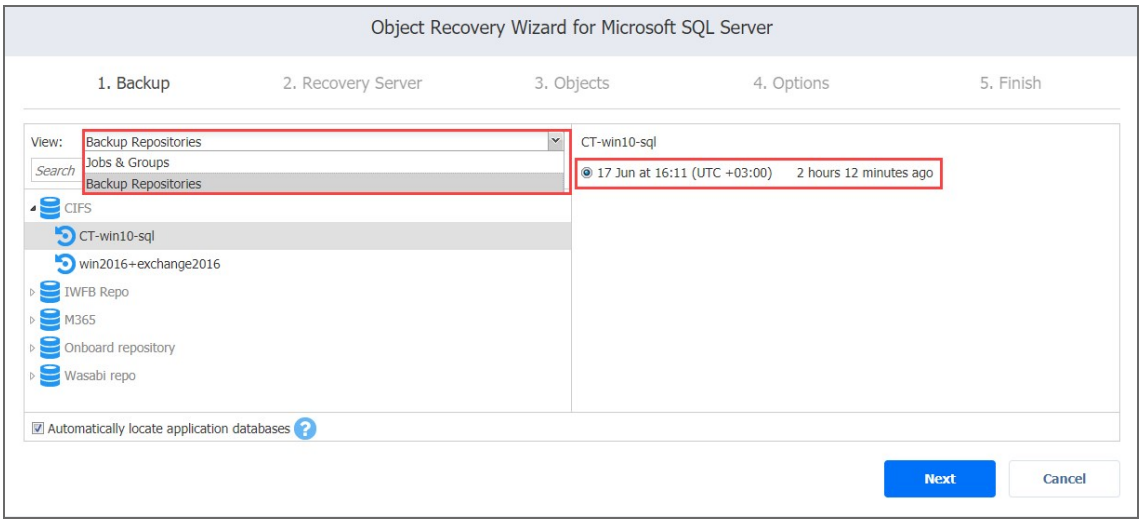

3. Click **Next**.

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. **Username**: Enter a user name with administrative privileges for the recovery server entered above. Refer to "Requirements for Microsoft SQL Server Object Recovery" on page 125 for a full list of requirements for recovering files to server.
   e. **Password**: Enter a user password.
   f. **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**.

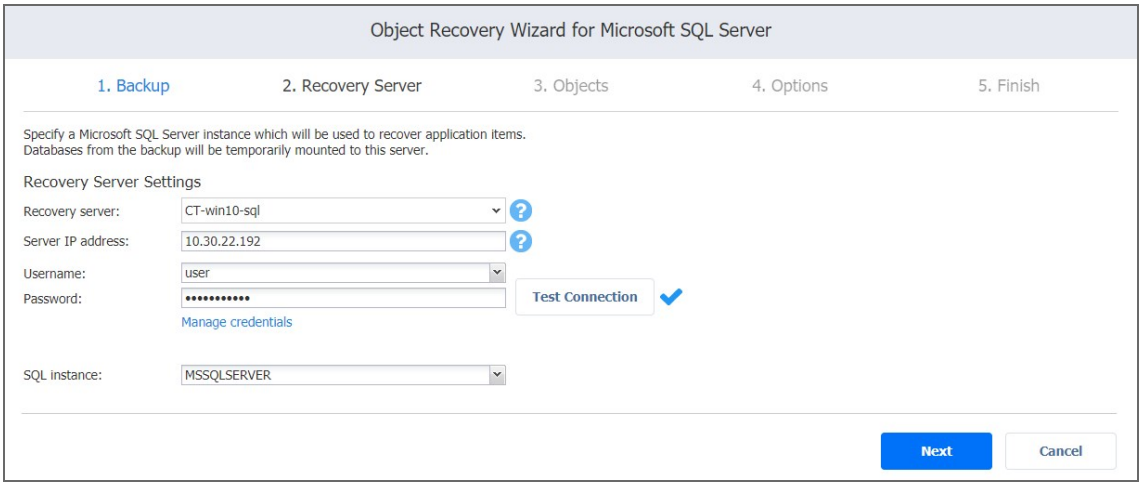

Object Recovery Wizard for Microsoft SQL Server: Objects

On the **Objects** page of the wizard, select objects for recovery. You can select either entire databases or individual objects for recovery.

1. Select the database in the left pane.
2. Select the objects in the right pane. If you want to restore an entire database, select all objects in this pane.
3. Alternatively, you can look for objects using the **Search** bar.
4. When you are done, click **Next**.

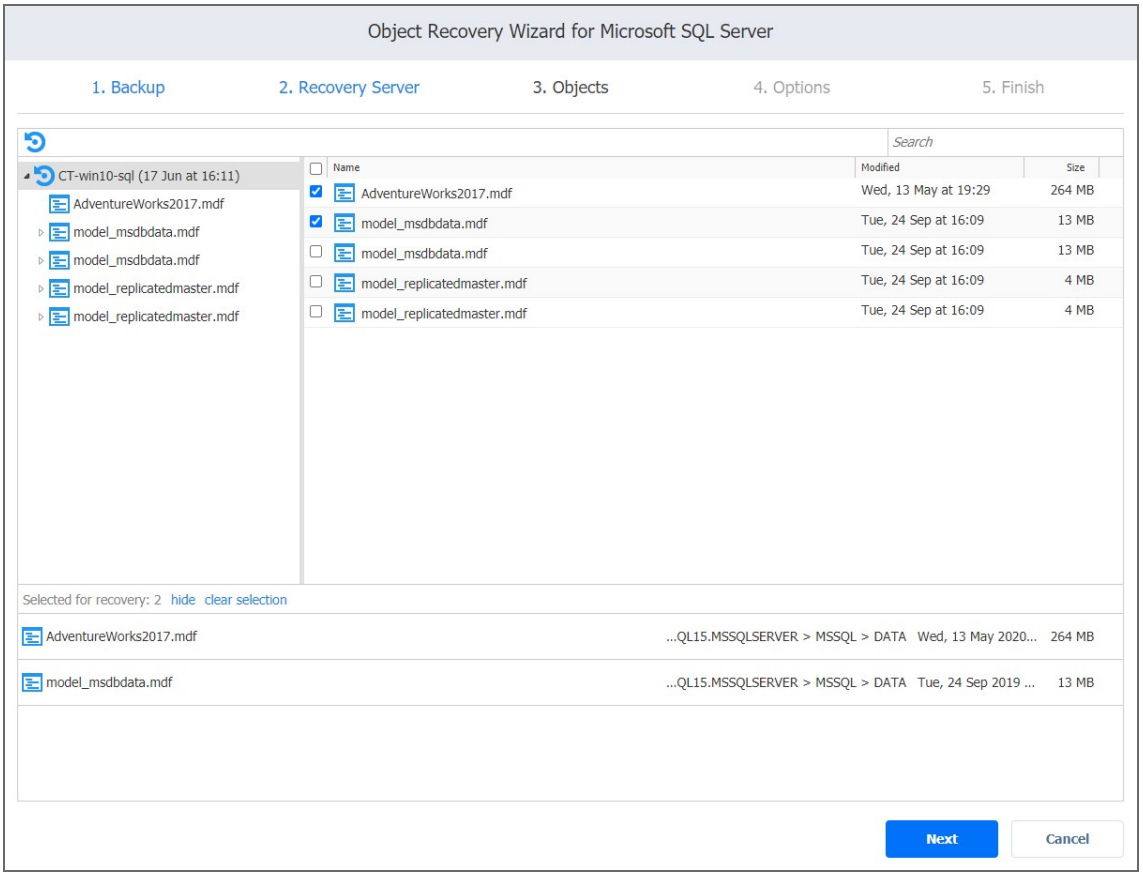

Object Recovery Wizard for Microsoft SQL Server: Options

On the **Options** page of the wizard, set the options for the recovery job.

- Recovery Scope
- Recovery Settings
- Overwrite Behavior

Recovery Scope

Set the recovery scope by selecting either **Recover schema and data** or Recover only schema.

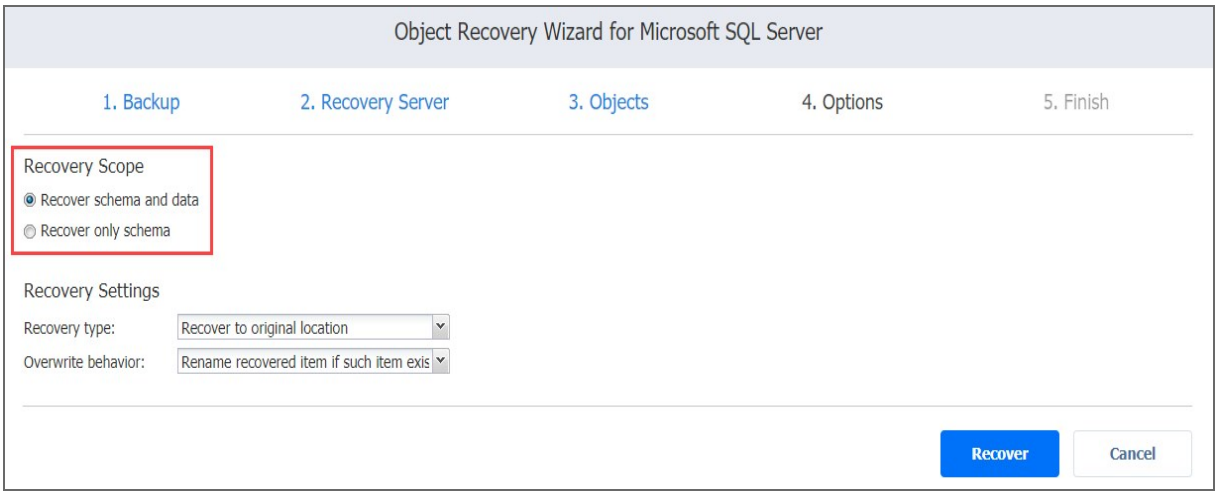

Recovery Settings

Set up the recovery type and overwrite behavior.

Recovery Type

- **Recover to original location**: Recover objects to the same server and SQL instance where they were originally located.
- **Recover to custom location**: Recover objects to a different instance.
- **Export** : Export objects as files to a specified location.
  - **SQL instance**: Select the target SQL instance.
  - **Target database**: Select the target database of the selected instance.
  - **Local folder on Recovery Server**: Specify a path to save objects.
  - **CIFS share**: Specify a remote CIFS (Windows) file share and your credentials for it (or select them

from the Manage credentials list).

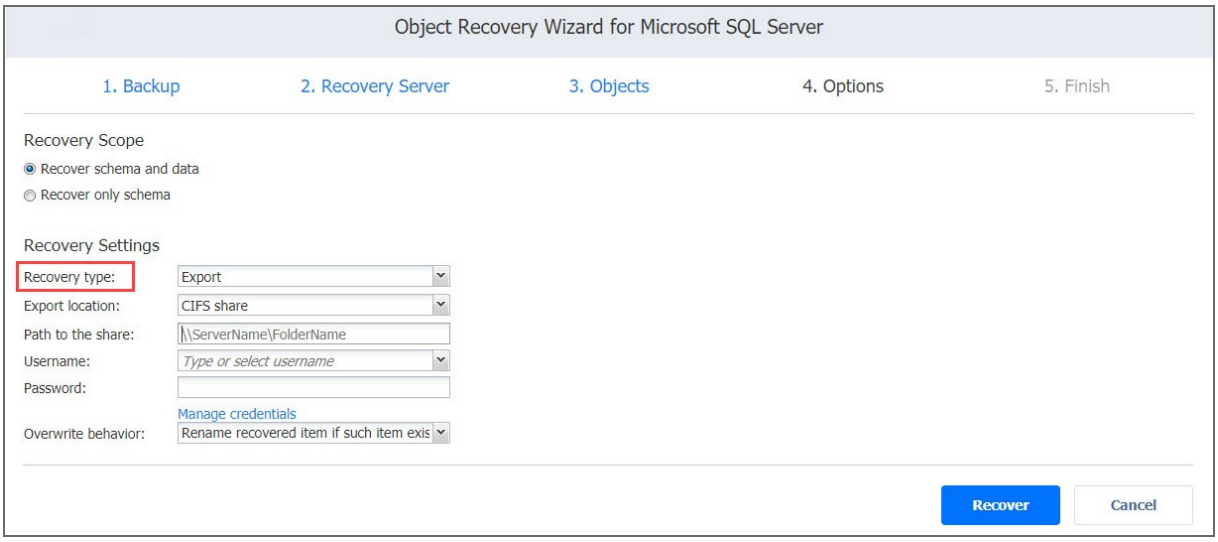

If you are using a domain name, enter it in the following format: domain\username

Overwrite Behavior

Select what to do if the recovered item conflicts with an existing one:

- **Rename recovered item if such item exists**
- **Skip recovered item if such item exists**
- **Overwrite the original item if such item exists**

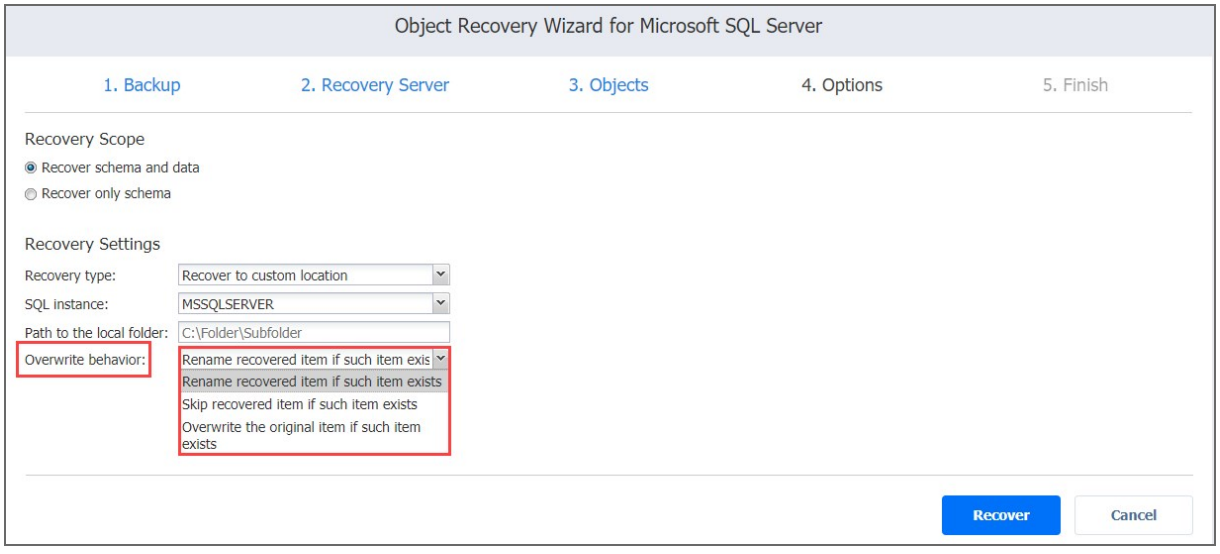

Click **Recover** to start the object recovery process. The **Finish** page opens.

Performing Universal Object Recovery

With Universal Object Recovery you can choose a disk from a VM recovery point and mount it to a target machine. This will allow you to recover backup data located on the mounted disk. Before creating a Universal Object Recovery job, make sure the System Requirements for recovering files to a server are met.

Please refer to the following topics for creating a Universal Object Recovery job:

- "Opening Universal Object Recovery Wizard" on page 642
- "Universal Object Recovery Wizard: Backup" on page 644
- "Universal Object Recovery Wizard: Disks" on page 645
- "Universal Object Recovery Wizard: Options" on page 646

Opening Universal Object Recovery Wizard

You can start the recovery process either from the **Dashboard** 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**.

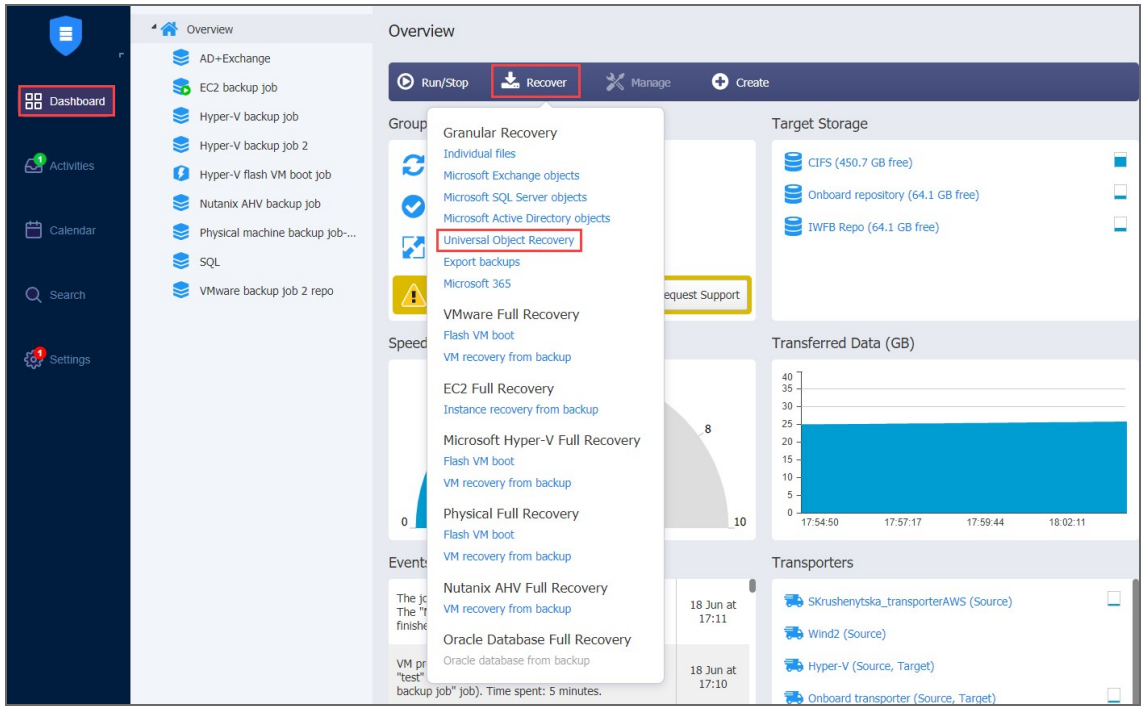

Starting Universal Object Recovery from a Backup Repository

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**.

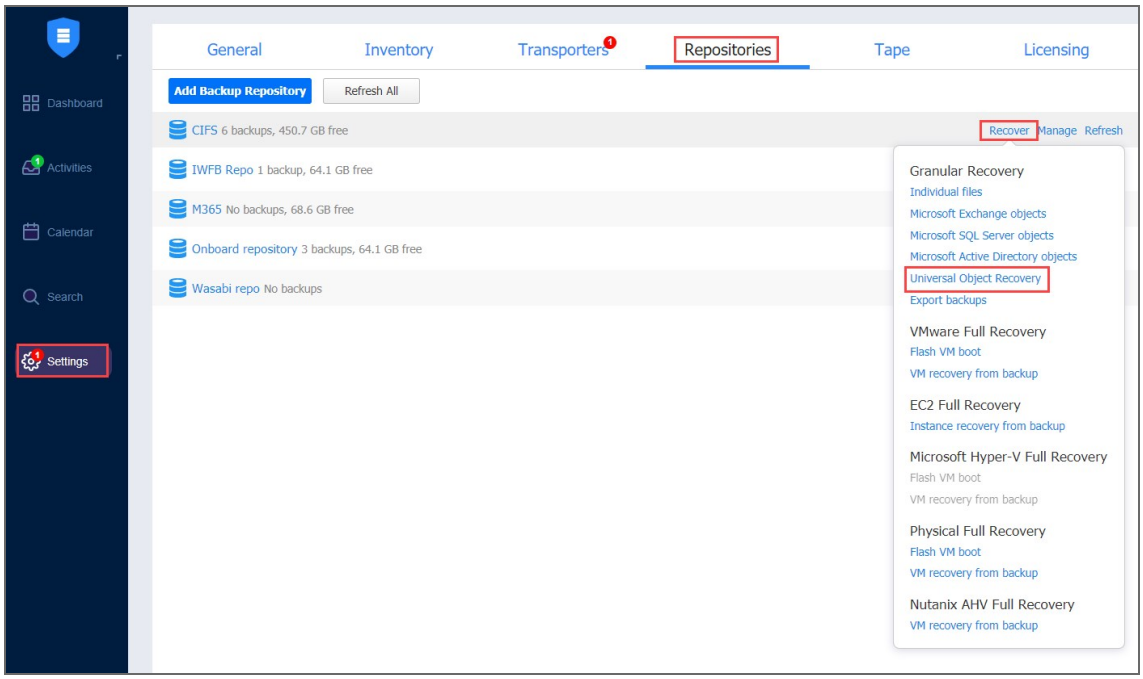

The new **Universal Recovery Job Wizard** opens.

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.

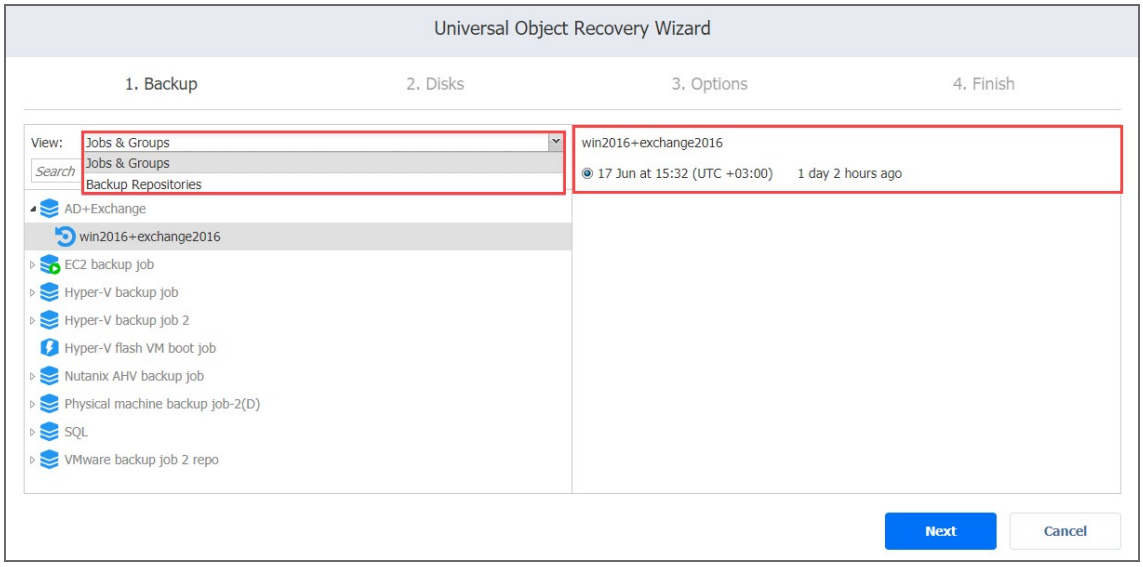

Universal Object Recovery Wizard: Disks

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.
   - **Username**: Enter a username with administrative privileges for the recovery server specified above.
   - **Password**: Enter a user password.
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.

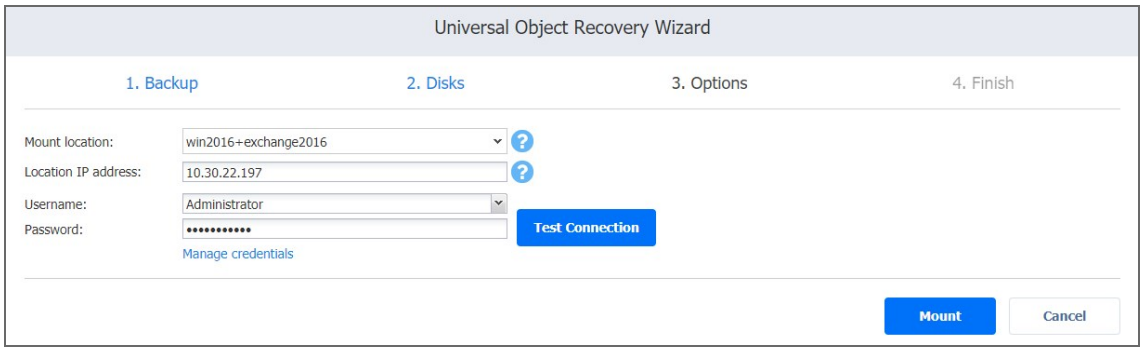

   The Universal Object Recovery is started and the **Finish** page of the wizard opens.
4. Click the **Activities** link to go to the **Activities** page if you want to view the progress of the Universal Object Recovery.
5. Click **Close** to close the Universal Object Recovery Wizard. Upon successful 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.

For more details, refer to the corresponding articles below:

- "Performing Flash VM Boot Recovery" on page 648
- "Hyper-V VM Recovery" on page 665
- "Recovery From Tape" on page 678
- "Performing Cross-Platform Recovery" on page 687

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.

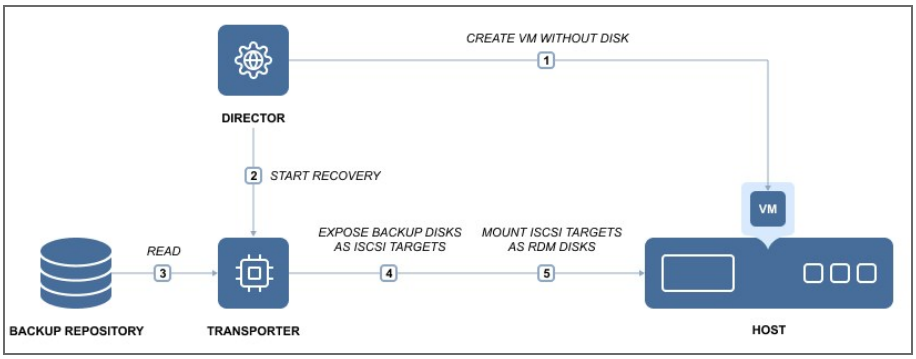

When the VMware VM is created, NAKIVO Backup & Replication takes a snapshot of the VM: this way all changes that occur to the VM are temporarily stored in the snapshot and discarded when you stop the job. When the Hyper-V VM is created, the application temporarily stores the changes to the VM in a disk-based write cache in the Backup Repository; changes are discarded when the job is stopped. For more information, refer to the following topics:

- “Creating Hyper-V Flash VM Boot Jobs” on page 649
- “Migrating Recovered VMs Using Flash Boot” on page 663

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**.

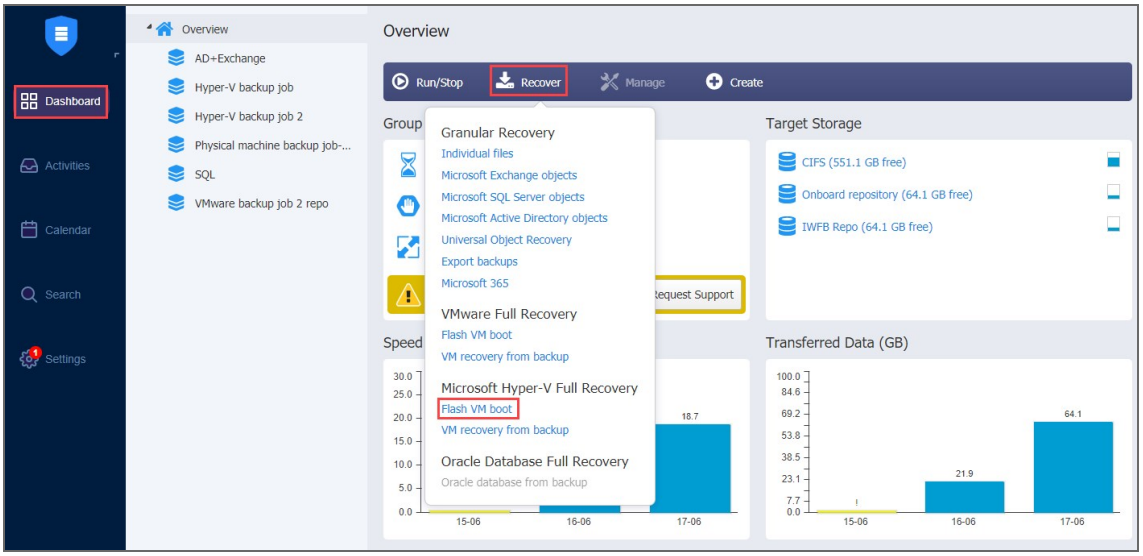

- Open the Flash VM Boot wizard from the **Repositories** tab by following the steps below:
  1. Click **Settings** in the main menu of the product.
  2. Go to the **Repositories** tab.
  3. Hover over the name of the Backup Repository, click **Recover** and then click **Flash VM Boot**.

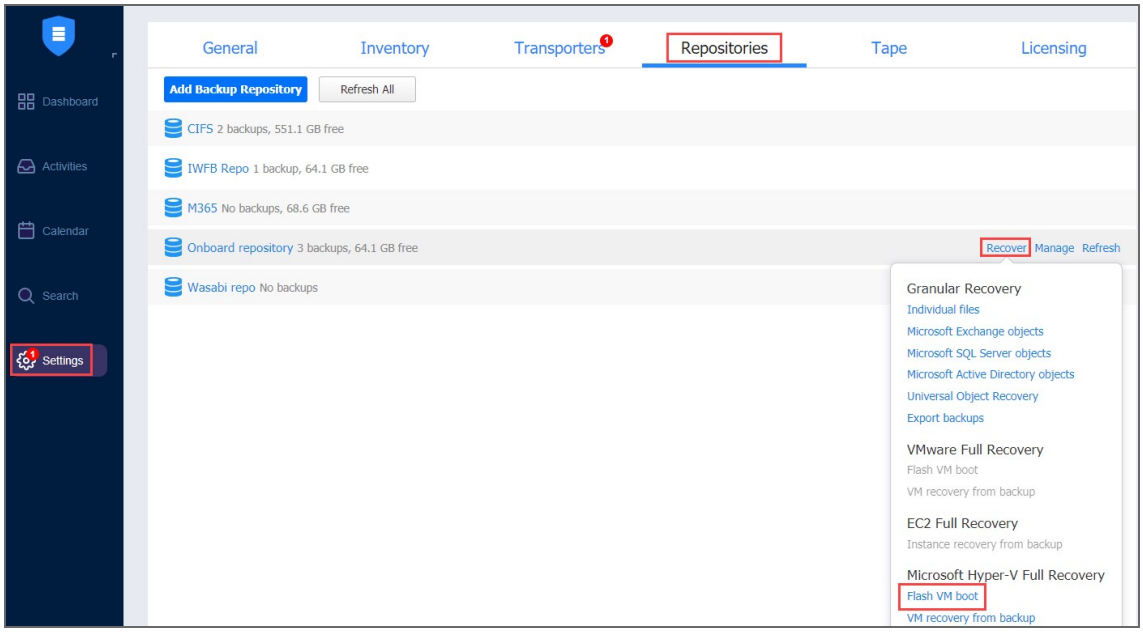

The **New Flash Boot Job Wizard** opens.

- "Hyper-V Flash Boot Job Wizard: Backups" on page 651
- "Hyper-V Flash Boot Job Wizard: Destination" on page 652

- “Hyper-V Flash Boot Job Wizard: Schedule” on page 654
- “Hyper-V Flash Boot Job Wizard: Options” on page 658

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

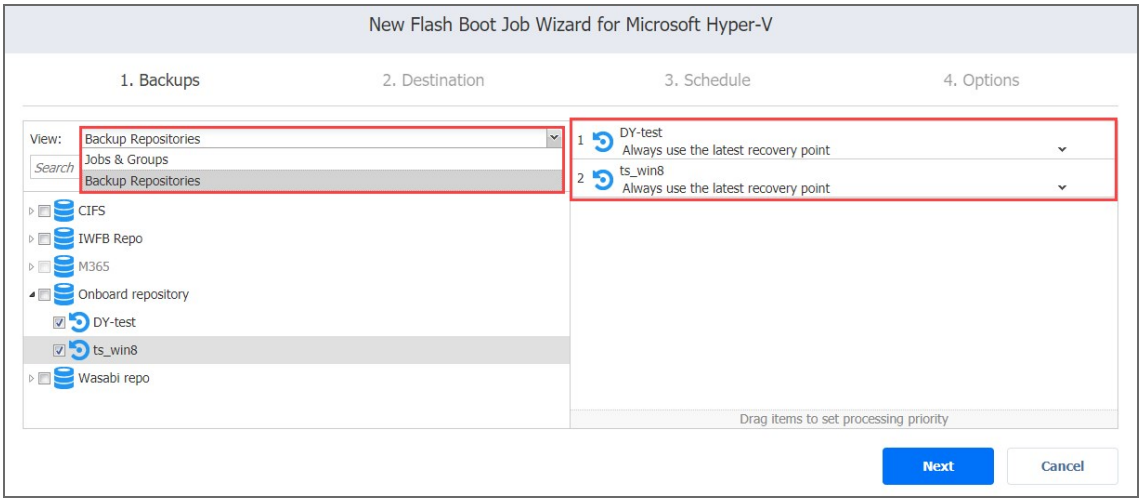

2. Select a recovery point for each backup in the right pane.
3. Click **Next** to go to the next page.

Hyper-V Flash Boot Job Wizard: Destination

On the **Destination** page, select a destination for the recovered VMs.

Important

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.

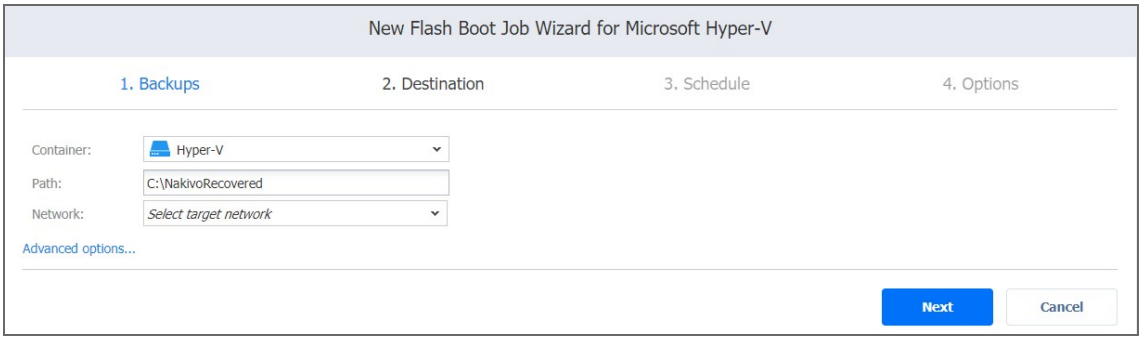

Note

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 options**.
2. Choose a target container and network adapter for each individual VM.

3. Click **Next**.

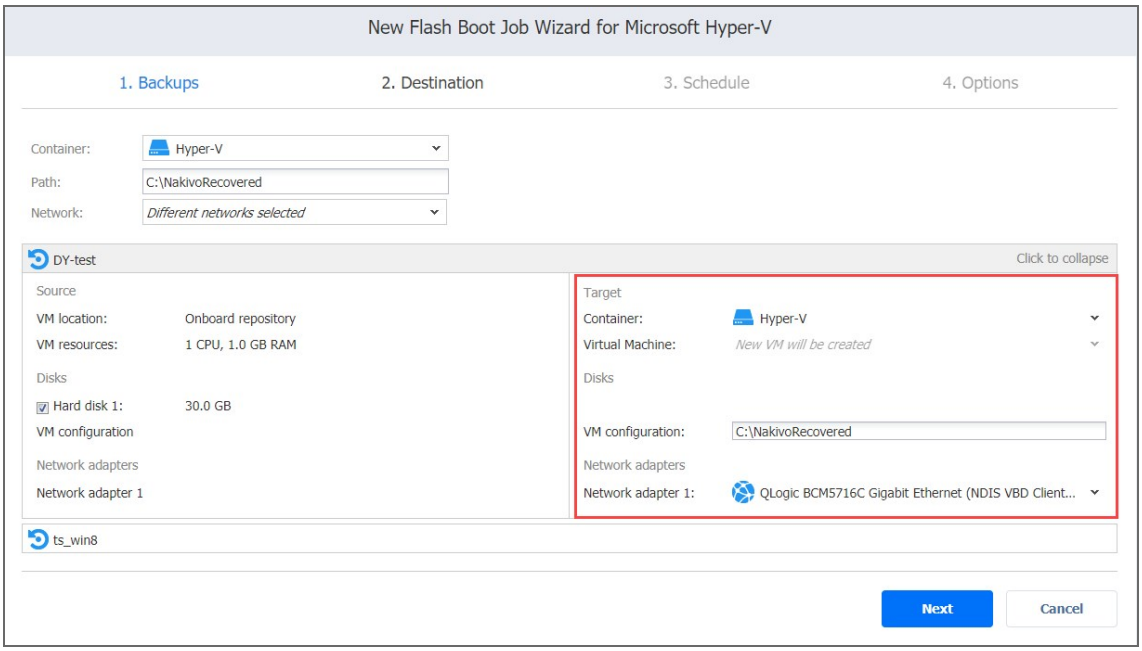

Note

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**.

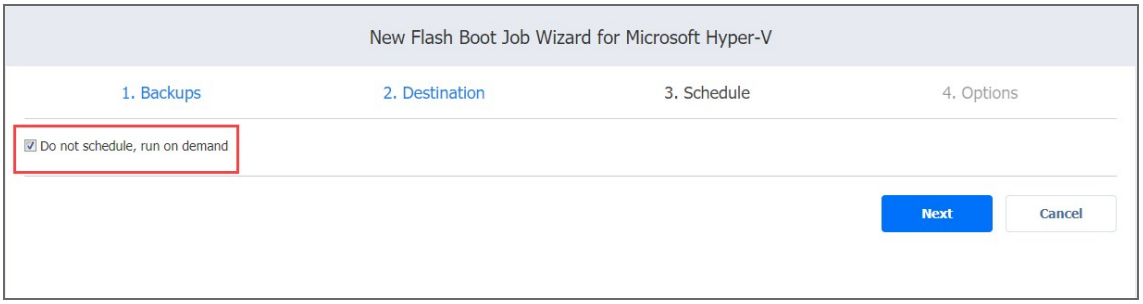

Daily Job Execution

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**.

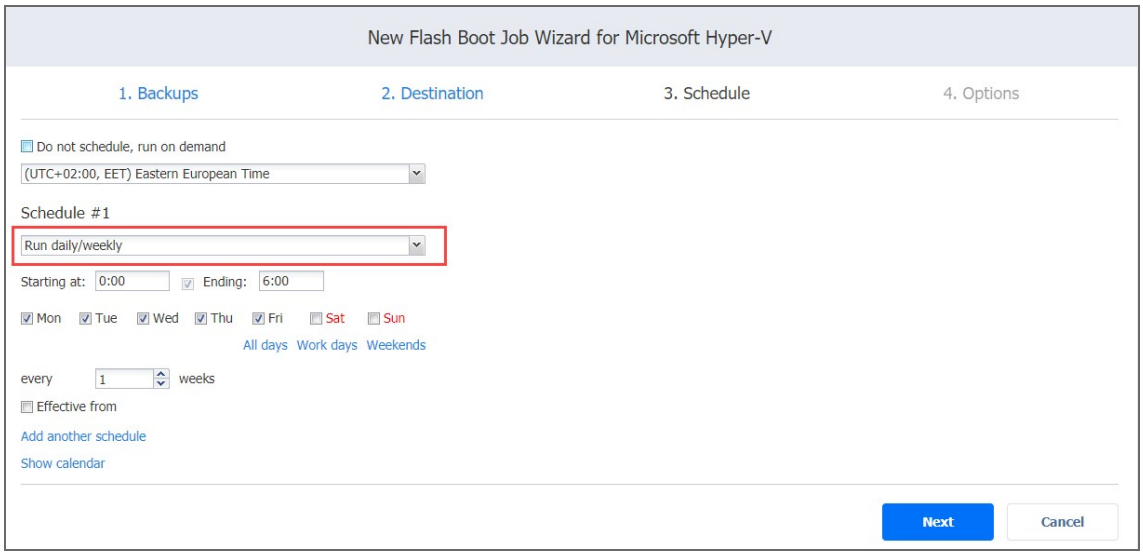

Monthly or Yearly Job Execution

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**.

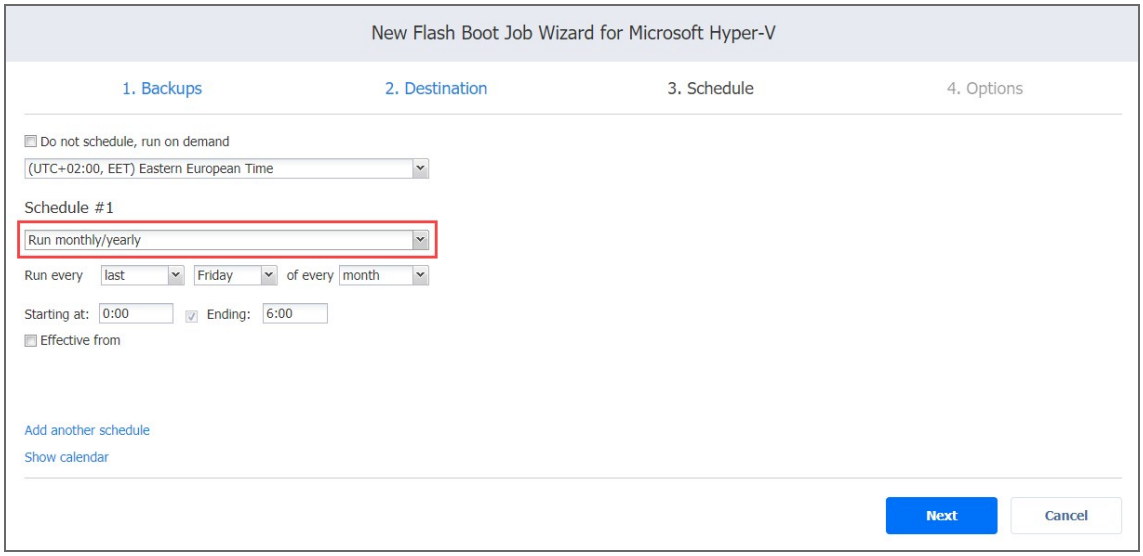

Periodic Job Execution

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**.

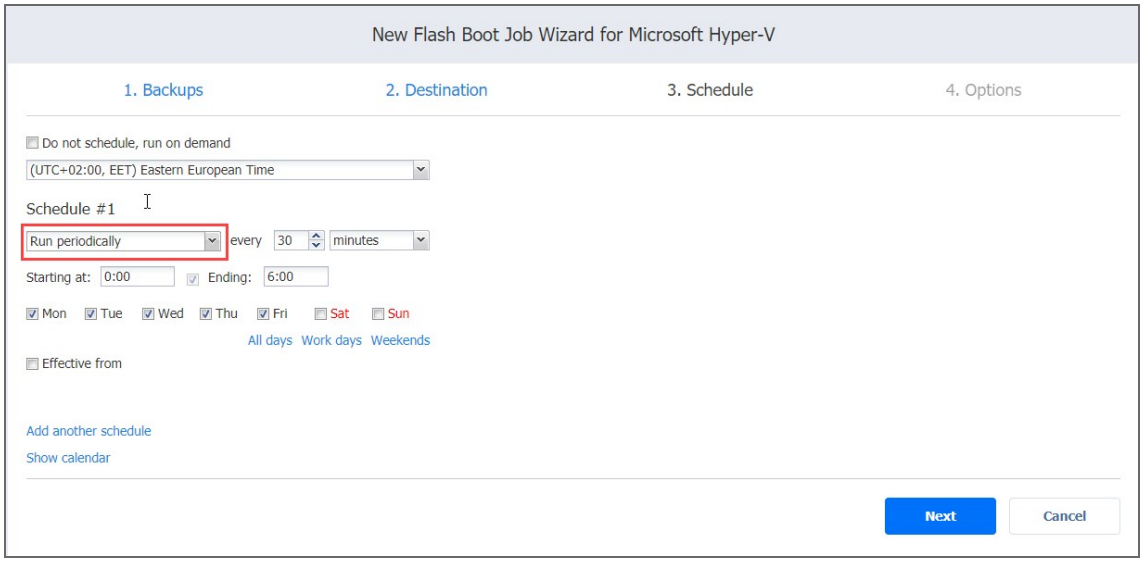

Chained Job

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**.

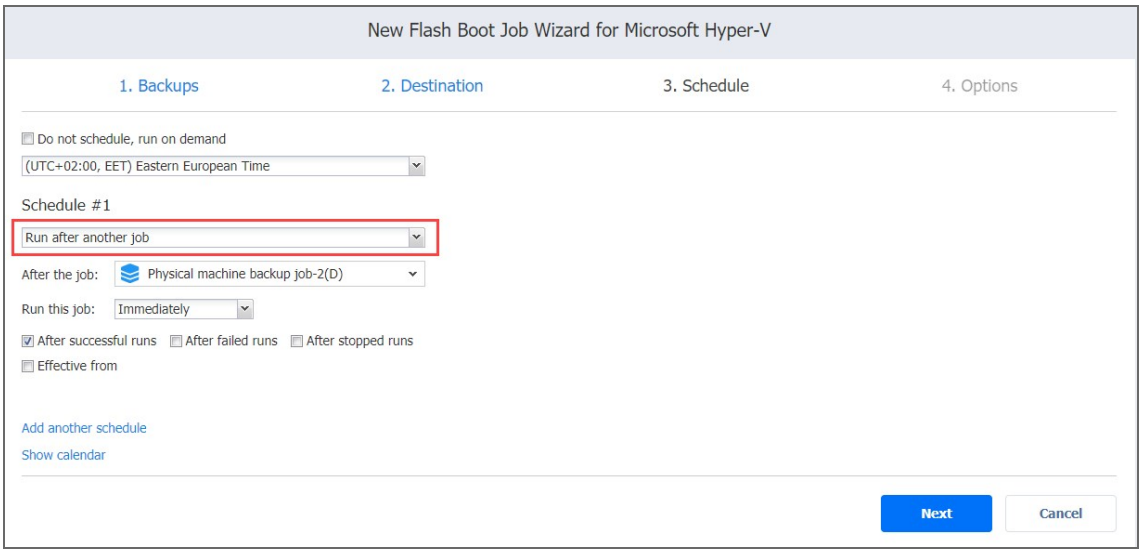

Additional Schedule

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 47.

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 "Email Notifications" on page 287) 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.

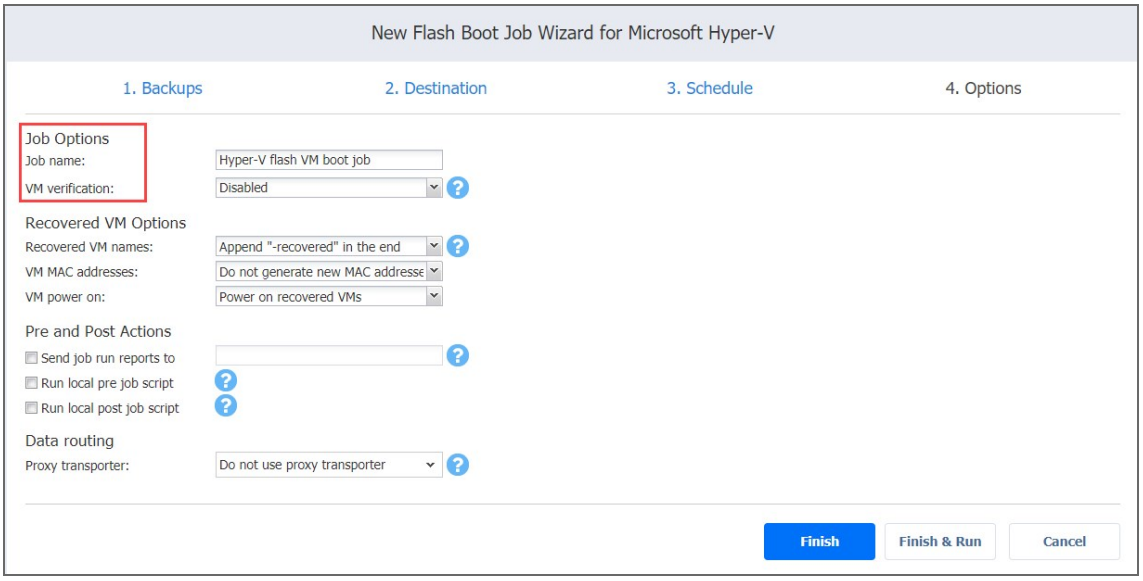

Recovered VM Options

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**

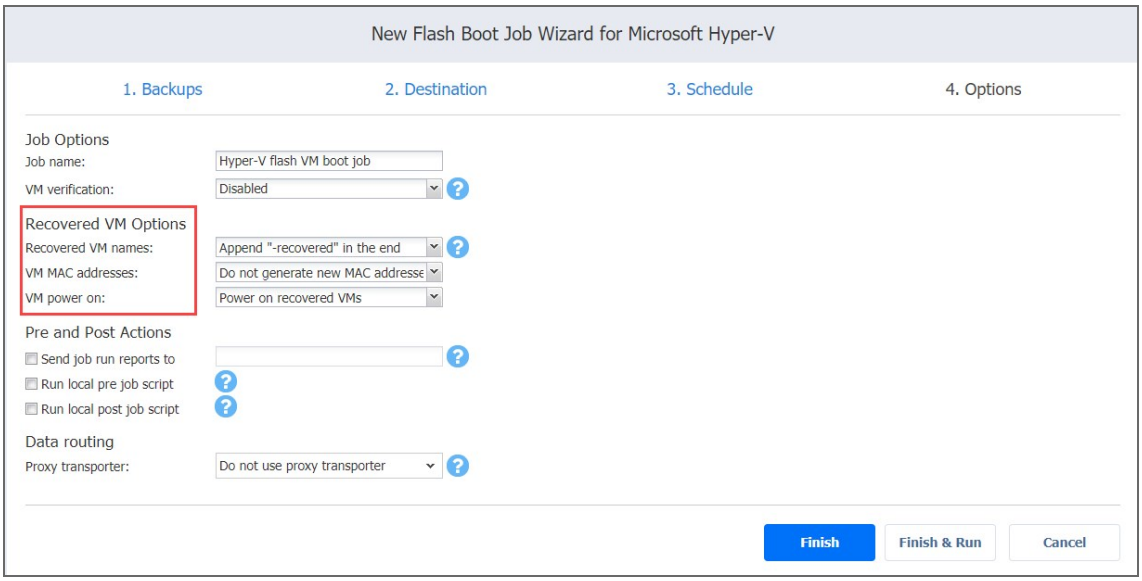

Pre and Post Actions

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 70 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.

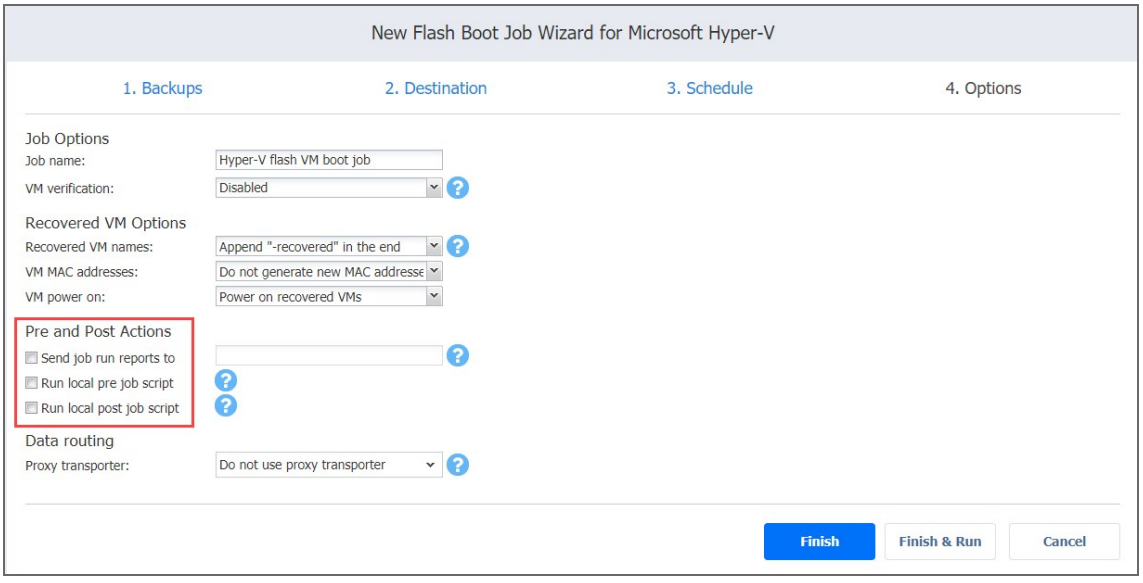

Data Routing

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.

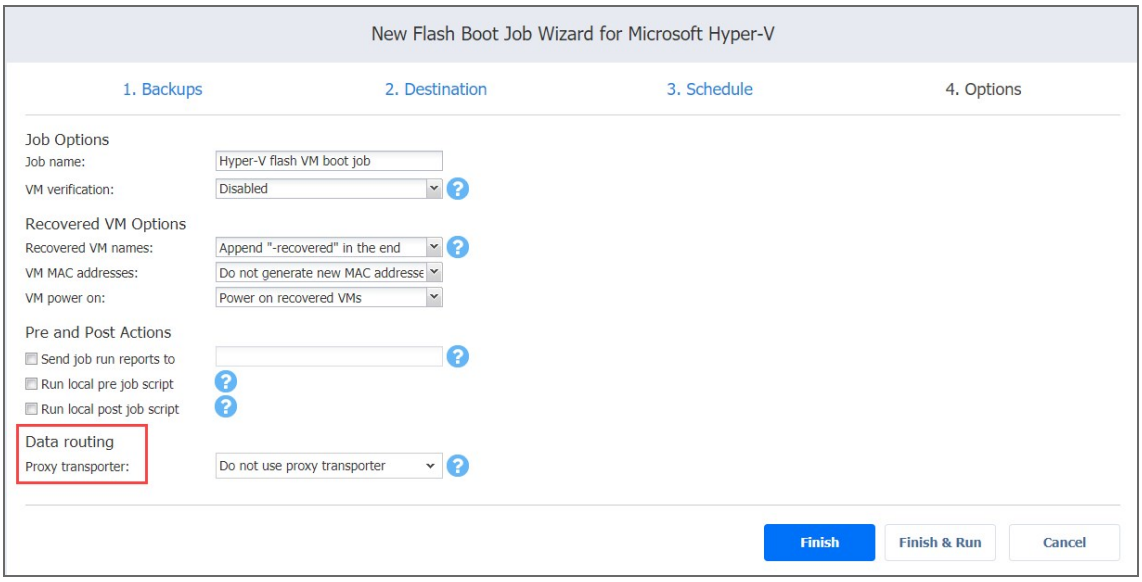

Completing the New Flash VM Boot Job Wizard

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 250 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 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.

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" on page 666
- "Recovery Job Wizard for Hyper-V: Backups" on page 668
- "Recovery Job Wizard for Hyper-V: Destination" on page 669
- "Recovery Job Wizard for Hyper-V: Options" on page 672

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.**

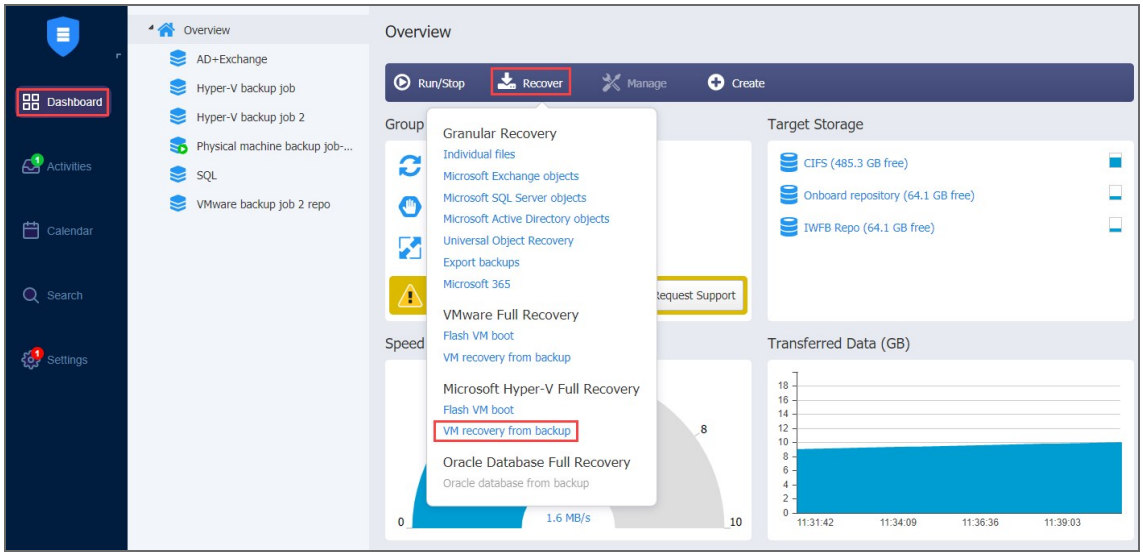

- Open the **New Recovery Job Wizard** from the **Repositories** tab following the steps below:
  1. Go to the **Repositories** tab.
  2. Click **Settings** in the main menu of NAKIVO Backup & Replication.
  3. Hover over the Backup Repository containing the necessary backup, click **Recover** and then click **VM recovery from backup.**

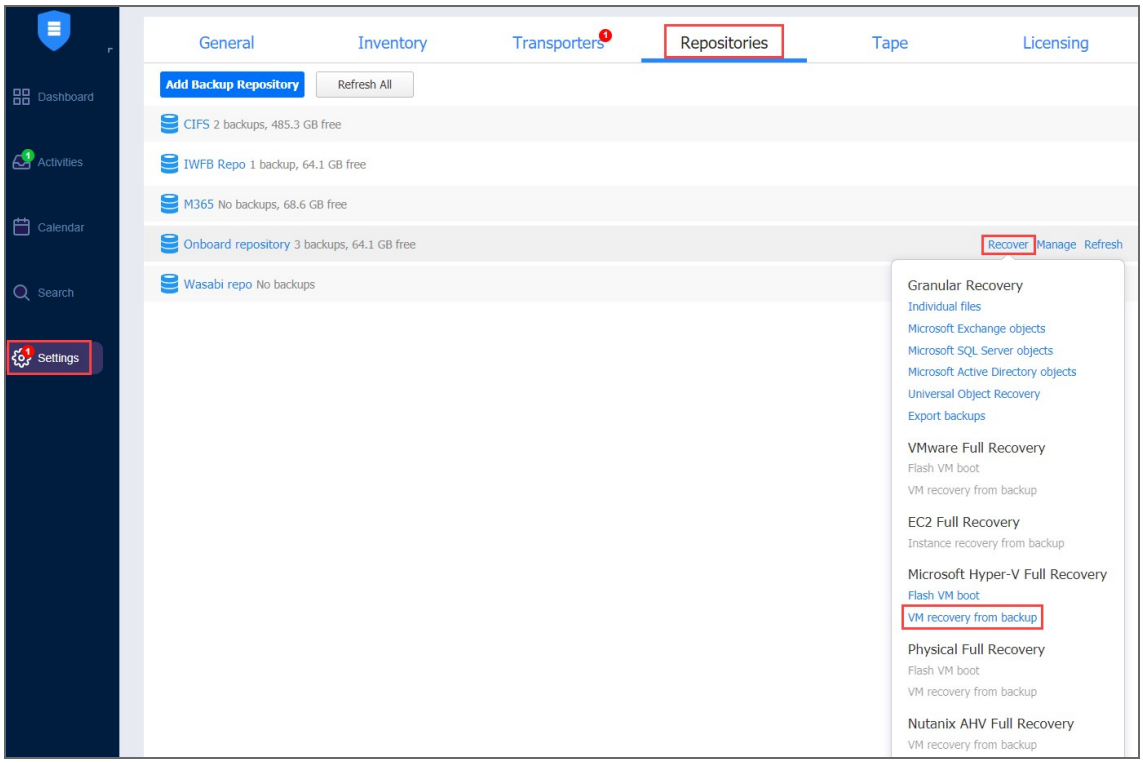

The **New Recovery Job Wizard for Hyper-V** opens.

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.

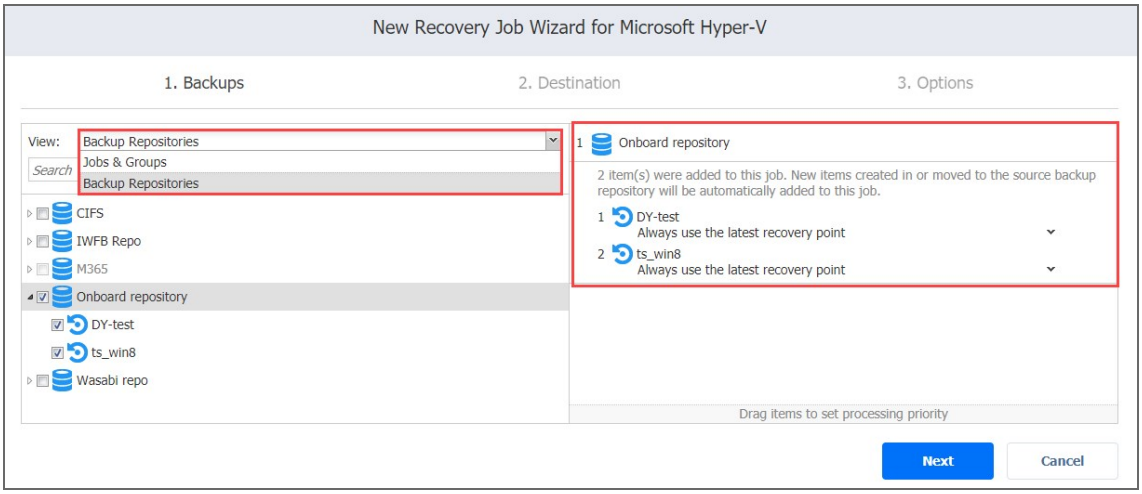

3. Click **Next** to go to the next step.

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.

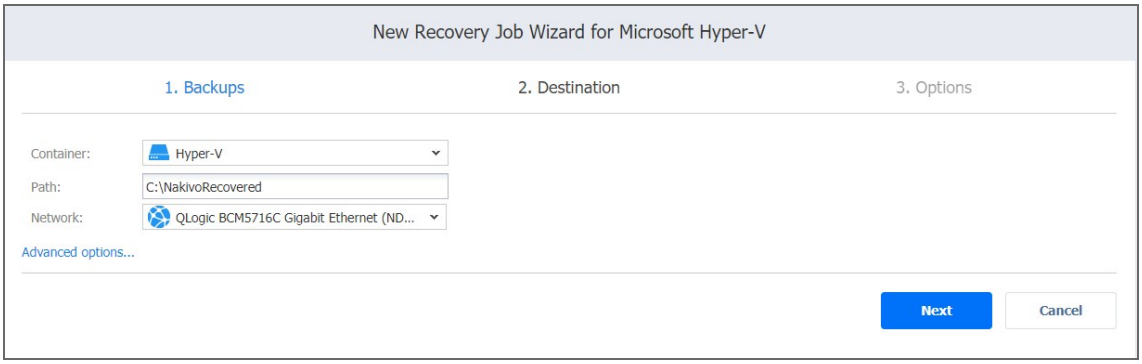

Note

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, folder, or a resource pool.
2. Choose a **Default container**.
3. If you have chosen the backup job on the **Source** page, you can choose a **Default Network**.
4. Optionally, you can also choose a **Default VM folder.**

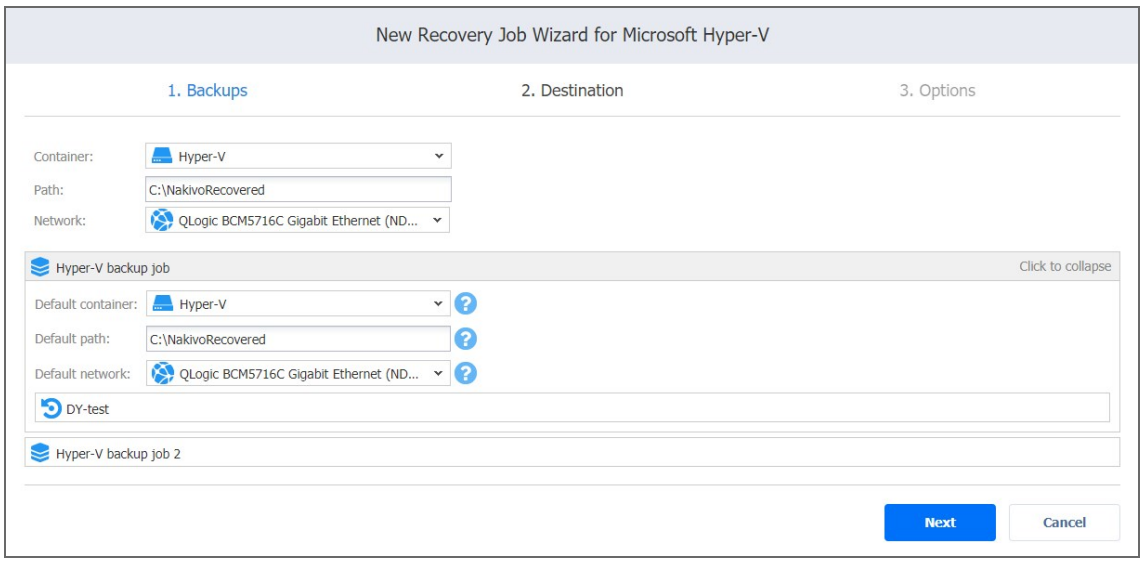

Setting Different Options for Recovered VMs

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.

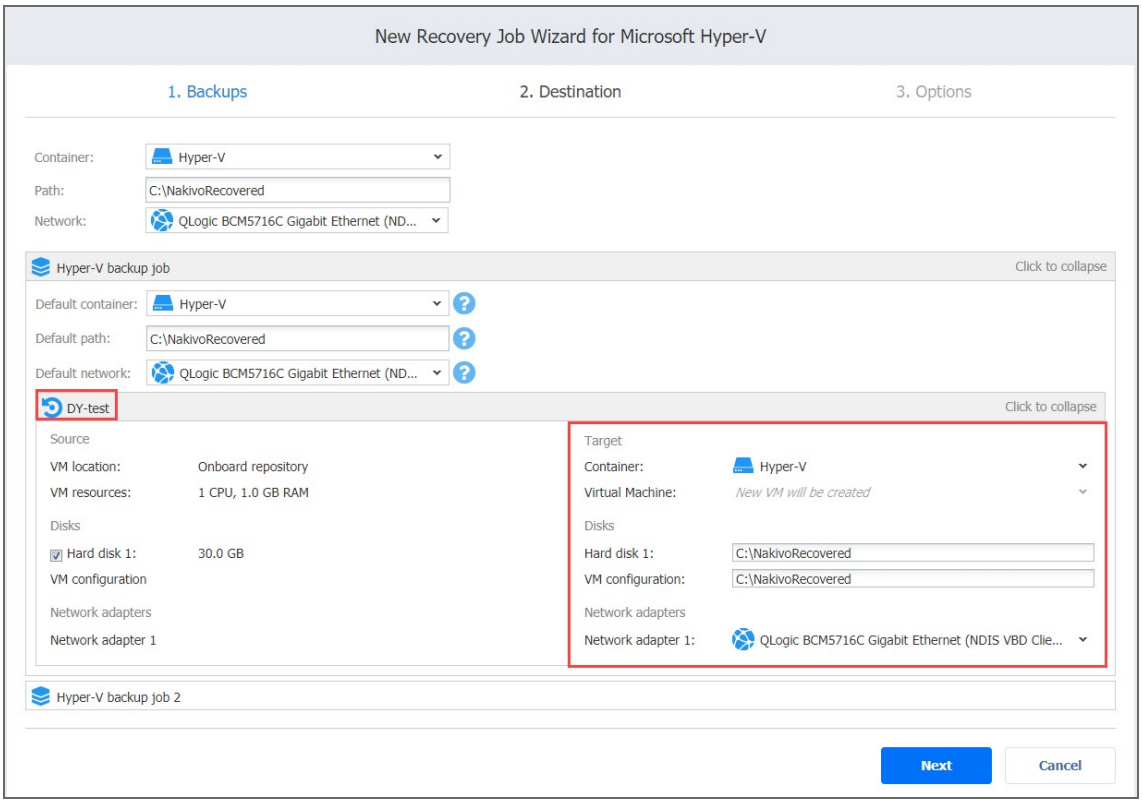

Recovery Job Wizard for Hyper-V: Options

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 674
  - "Setting Up a Pre Job Script" on page 674
  - "Setting Up a Post Job Script" on page 675
  - "Email Notifications" on page 675
- "Data Transfer" on page 676
  - "Transporter Load" on page 676
  - "Bandwidth Throttling" on page 676
- "Completing the New Recovery Job Wizard for Microsoft Hyper-V" on page 677

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.

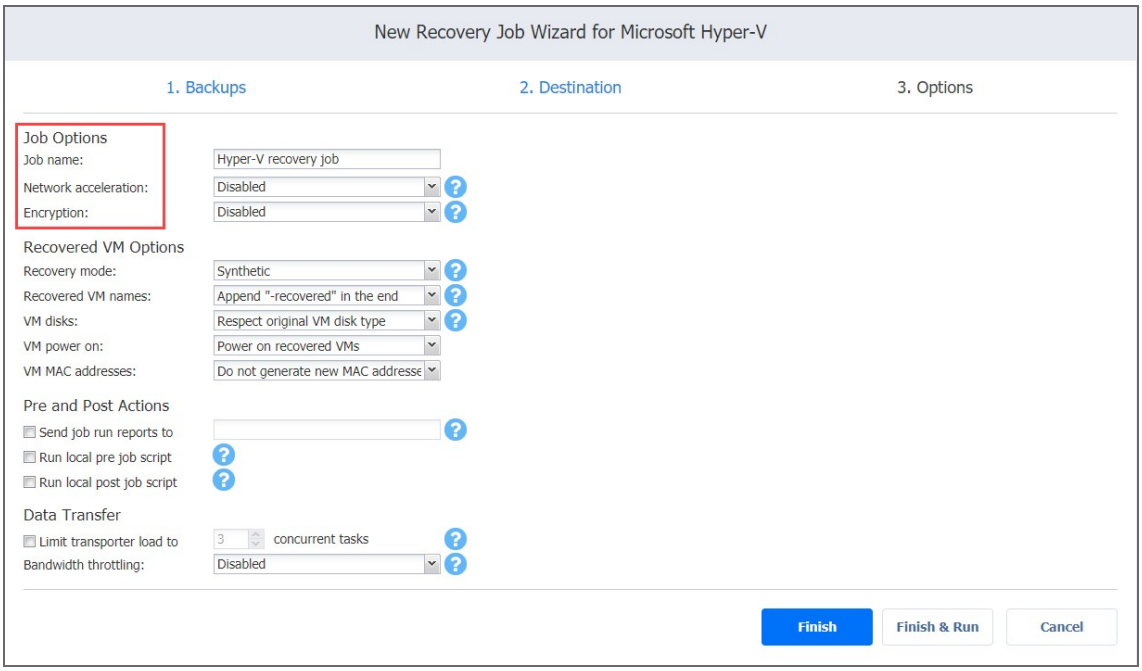

Recovery Options

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 “-recoverd” 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.

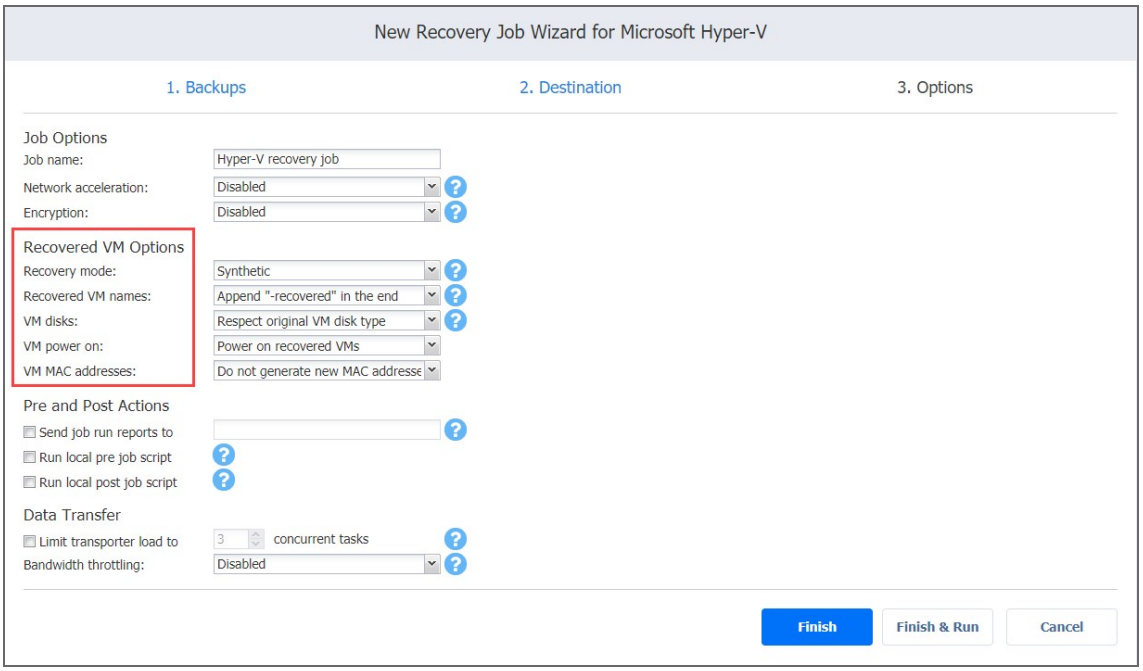

Pre and Post Job Actions

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.

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.

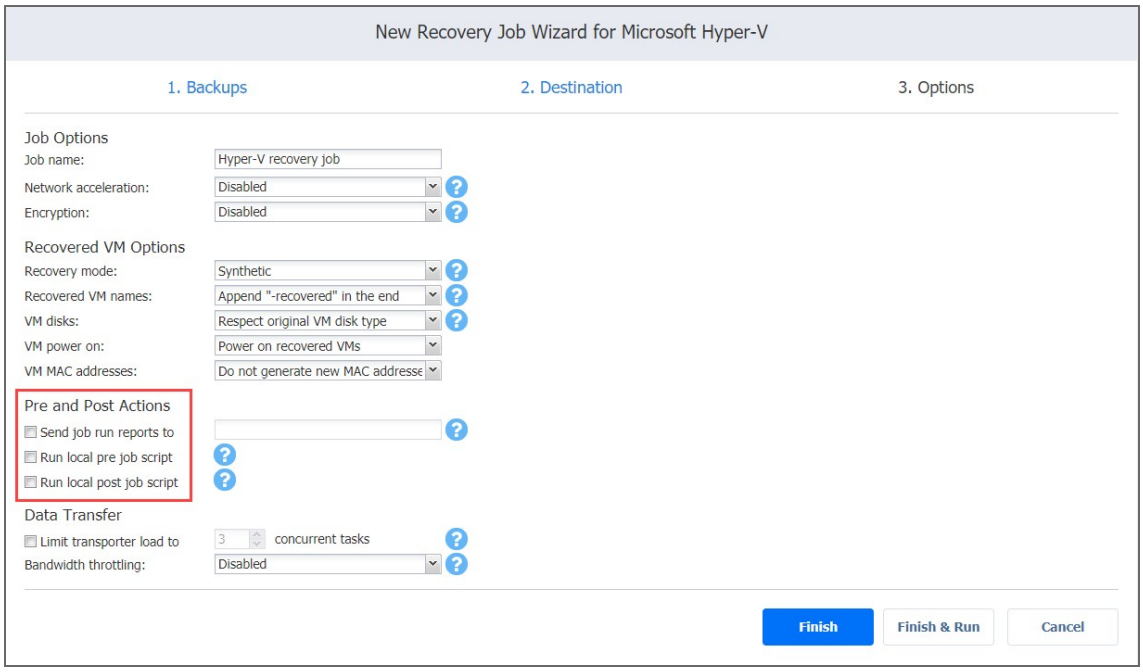

Data Transfer

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 281 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 281 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.

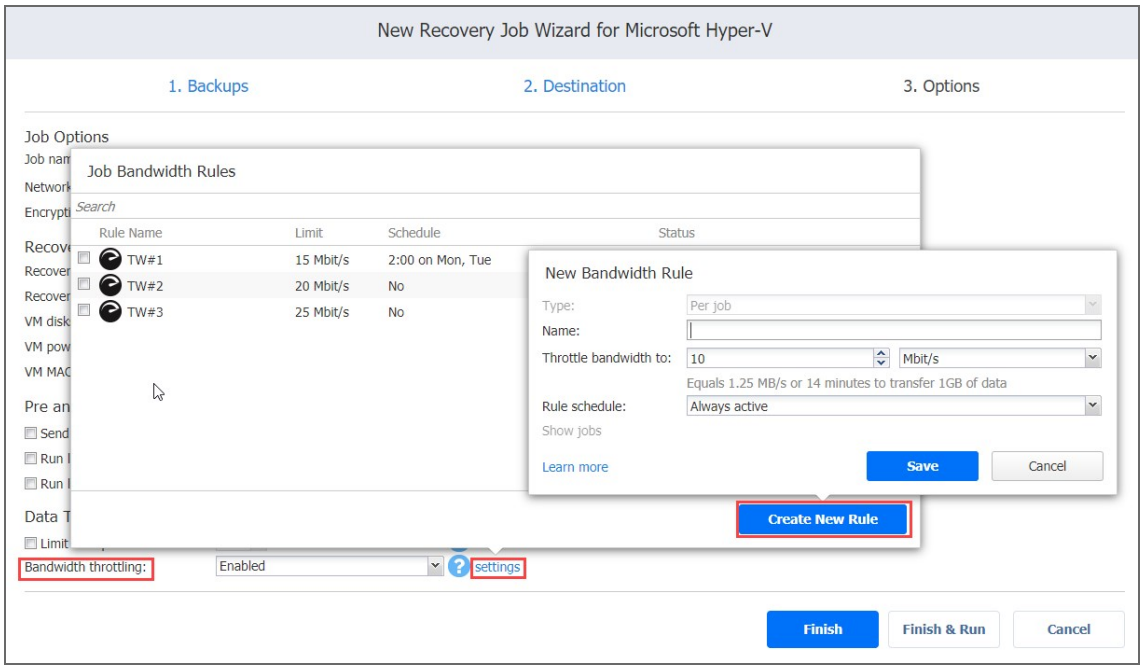

Completing the New Recovery Job Wizard for Microsoft Hyper-V

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 250 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.

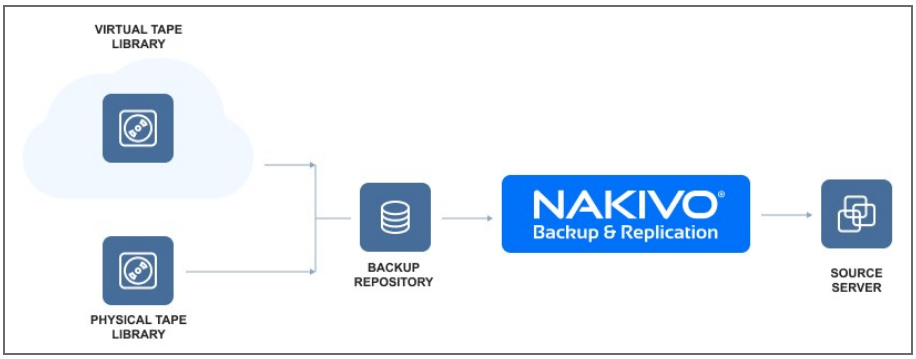

Refer to the following topics for more information:

- “Starting Recovery from Tape Wizard” on page 679
- “Recovery from Tape Wizard: Backups” on page 680
- “Recovery from Tape Wizard: Destination” on page 681
- “Recovery from Tape Wizard: Options” on page 682

Starting Recovery from Tape Wizard

To launch the recovery from the Tape wizard, do the following:

1. Go to **Inventory**, click the **Tape** tab and select **Backups** from the **View** drop-down list.

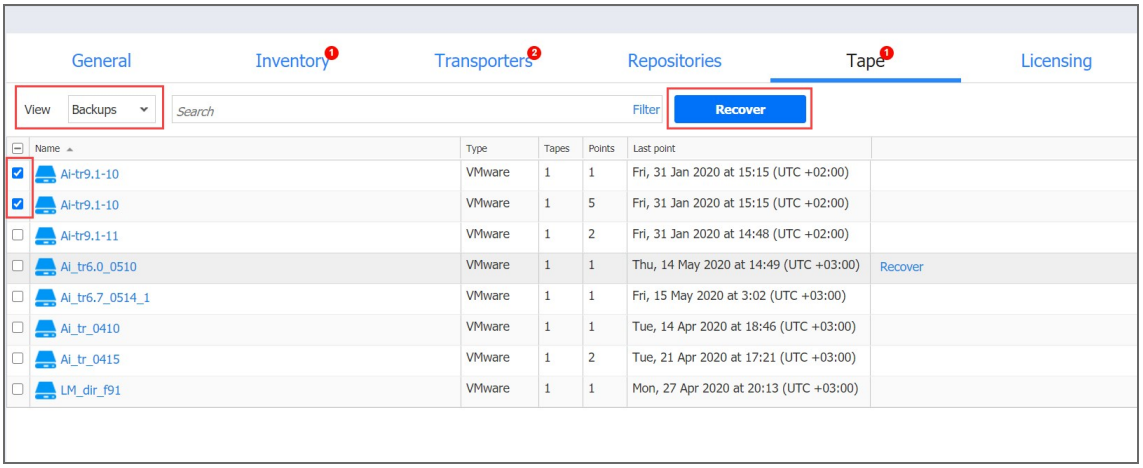

2. In the **Backups** table, do one of the following:
   - Select the checkbox next to one or several backups that you want to recover and click the **Recover** button. This opens the Recovery wizard with specified backups and their latest recovery point selected.

     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 448 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.

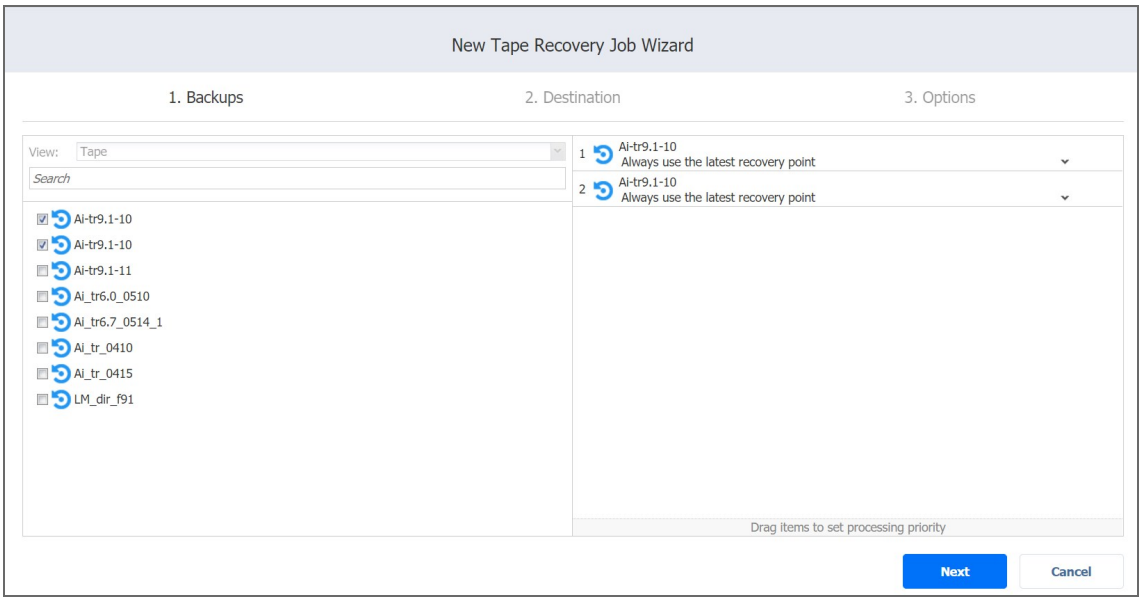

After you are done, click **Next**.

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**.



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 the next page
  - "Setting Up a Post Job Script" on page 684
  - "Email Notifications" on page 684
- "Data Transfer" on page 685
  - "Bandwidth Throttling" on page 685

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 62.
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.

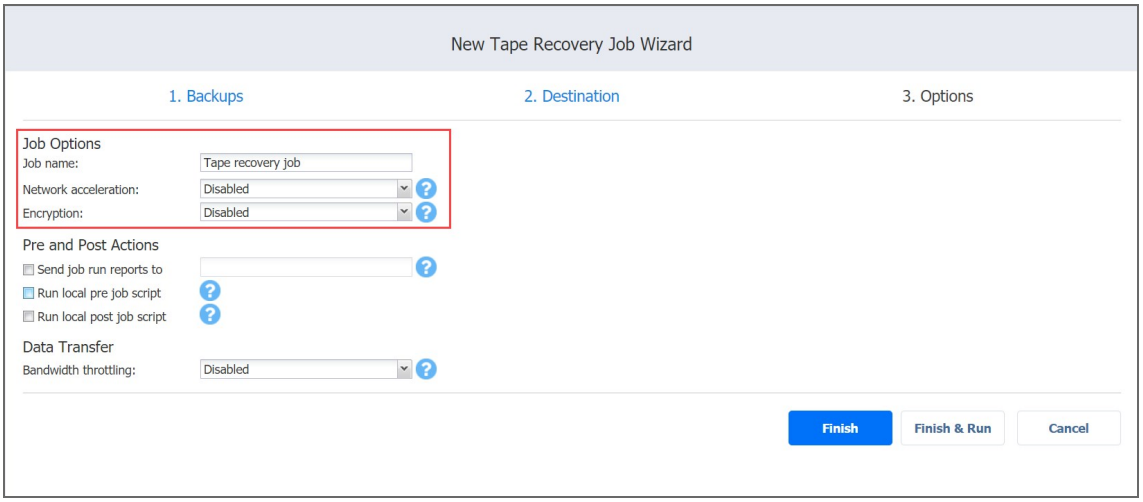

Pre and Post Actions

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 70 for details.

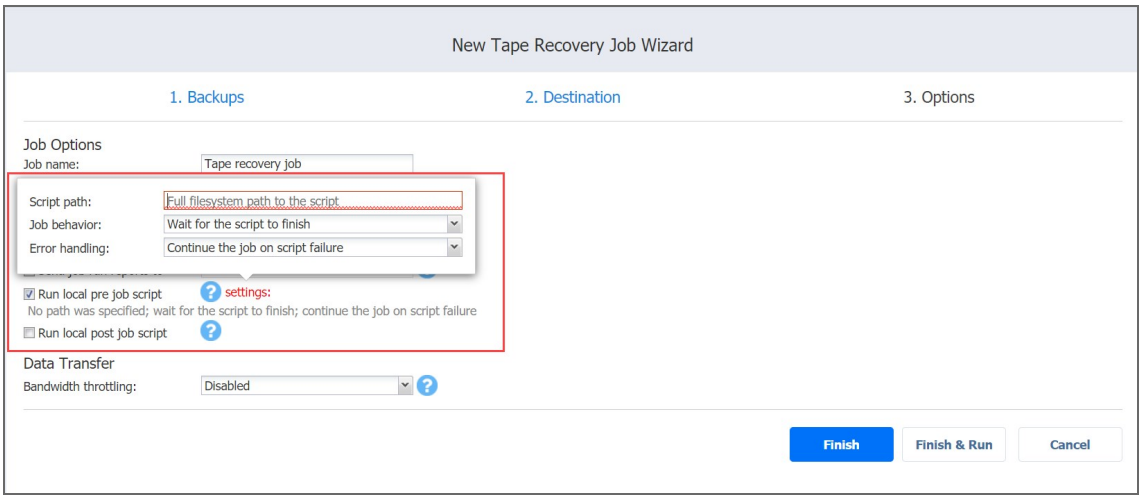

Setting Up a Pre-Job Script

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 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.

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 "Email Notifications" on page 287 for details.



Data Transfer

In the **Data Transfer** section of the **Options** page, you can set or configure bandwidth throttling rules.

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 281 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 281 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.

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 120 for the supported scenarios for cross-platform recovery.

NAKIVO Backup & Replication does not run VM preparation when exporting the backups into a specific format. If you plan to import the VM into a different platform and VM preparation is required, prepare your VM in advance.

This section includes the following topics:

- “Opening Backup Export Wizard” on page 688
- “Backup Export Wizard: Backups” on page 690
- “Backup Export Wizard: Disks” on page 691
- “Backup Export Wizard: Options” on page 692
- “Backup Export Wizard: Finish” on page 693

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**.

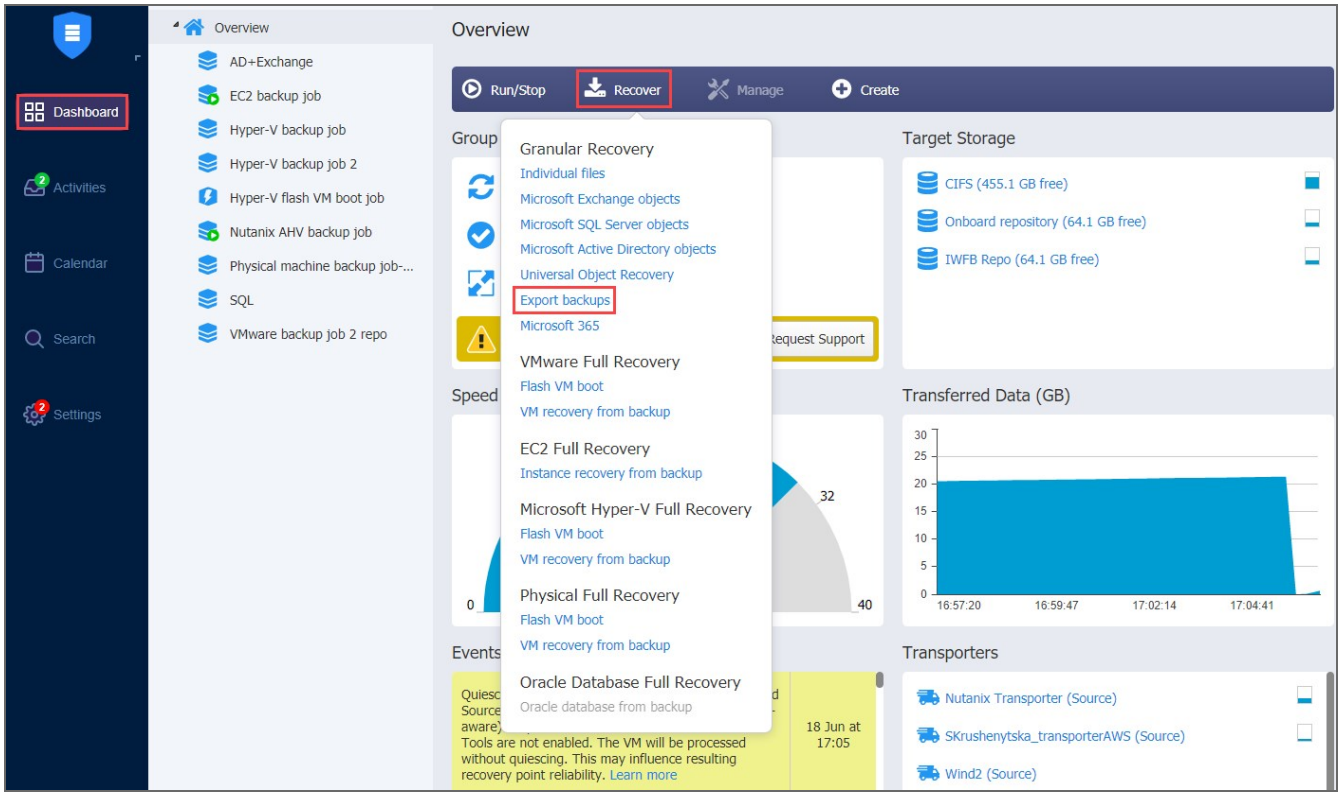

- On the *Settings* page:
    1. Click the **Repositories** tab.
    2. In the list of repositories, click a repository to expand it.
    3. In the list of backups, hover over a backup and then click **Recover**.

4. In the menu that opens, click **Export Backups**.

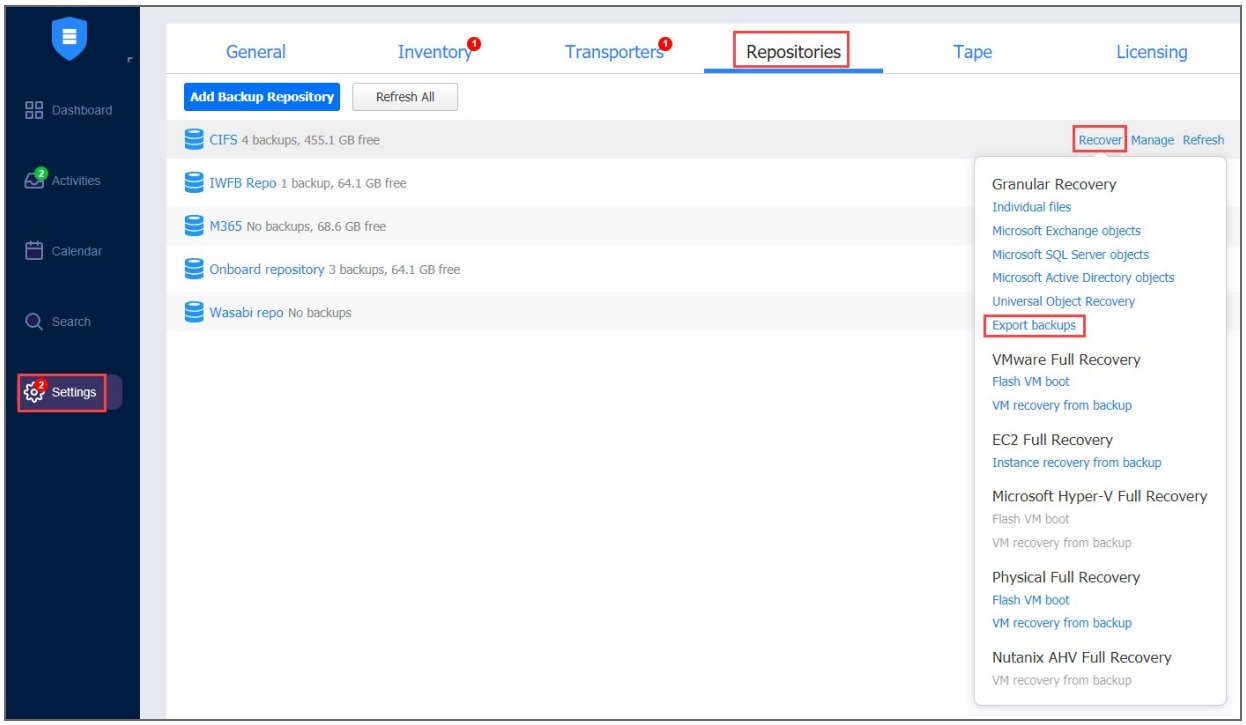

The **Backup Export Wizard** opens.

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.

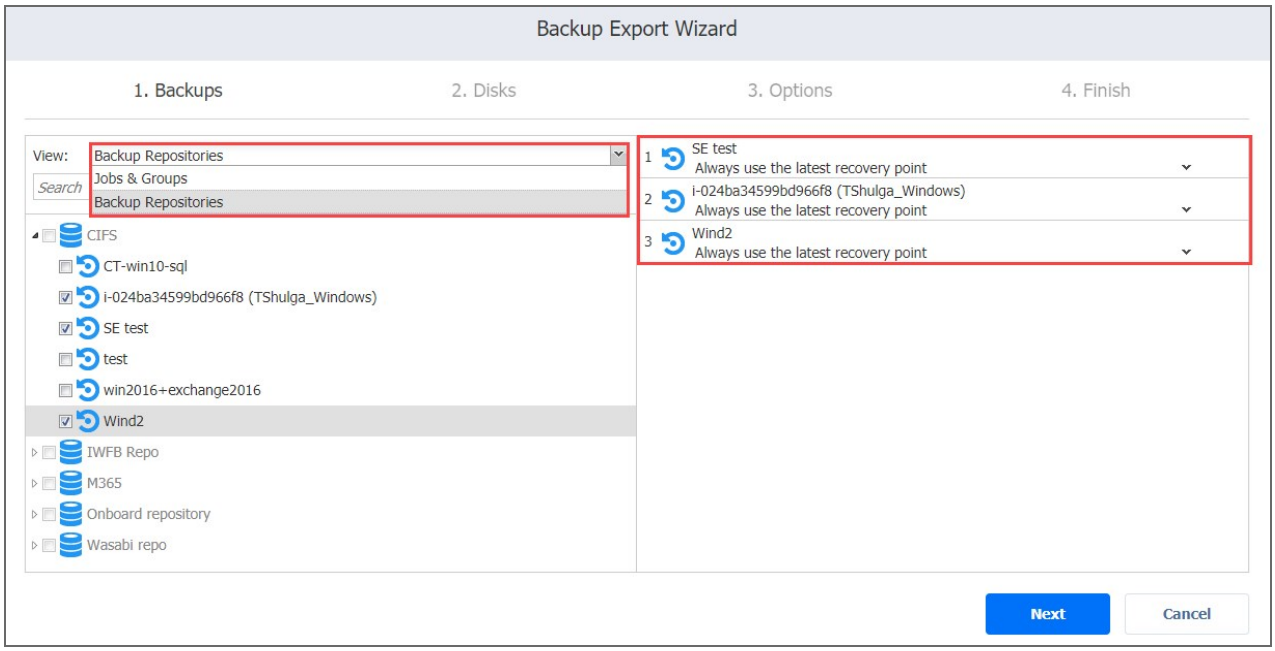

Backup Export Wizard: Disks

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.

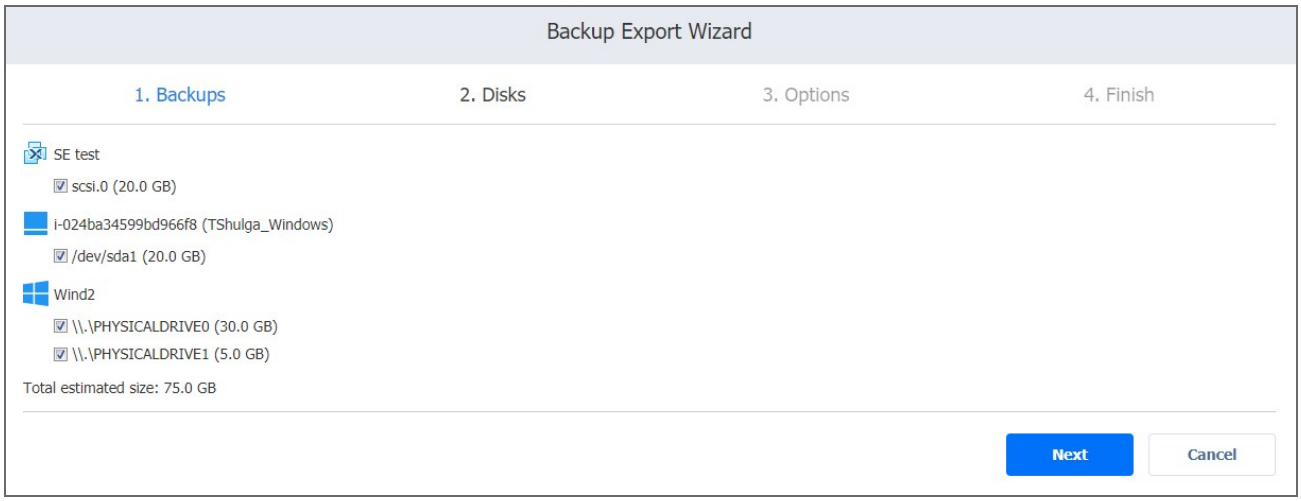

Backup Export Wizard: Options

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
        b. Click **Test Connection** to check your credentials for the specified share.
    - **NFS share**: With this option selected, proceed as follows:
        a. Enter Path to the share.
        b. Click **Test Connection** to check the connection to the specified share.
3. Click **Export** to go to the next page of the wizard.

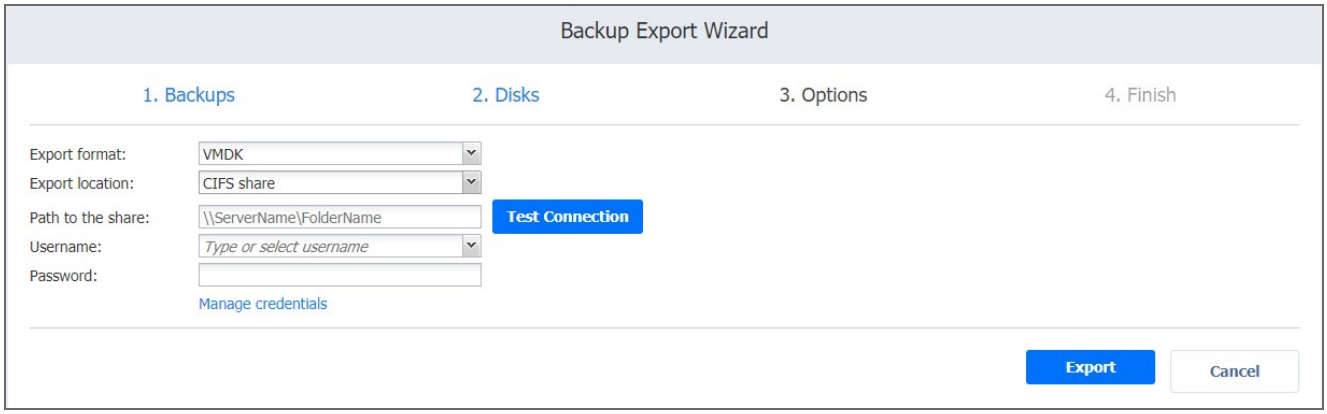

Backup Export Wizard: Finish

The **Finish** page of the wizard informs you that your backup export has started. To view the status of your backup export, go to **Activities**.

To view the backup export progress, go to **Settings** > **General** > **Events**.

To close the **Backup Export Wizard**, click **Close**.

Planning Disaster Recovery

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 AWS 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 Microsoft Hyper-V" on page 695
- "Replica Failback for Microsoft Hyper-V" on page 704
- "Site Recovery Job" on page 714

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" on page 696
- "Failover Job Wizard for Microsoft Hyper-V: Source" on page 697
- "Failover Job Wizard for Microsoft Hyper-V: Networks" on page 698
- "Failover Job Wizard for Microsoft Hyper-V: Re-IP" on page 700
- "Failover Job Wizard for Microsoft Hyper-V: Options" on page 702

Starting Microsoft Hyper-V Failover to Replica

To start the procedure of switching a Hyper-V system workload to a backup VM, click Recover in the NAKIVO Backup & Replication Dashboard, and then click VM replica failover.

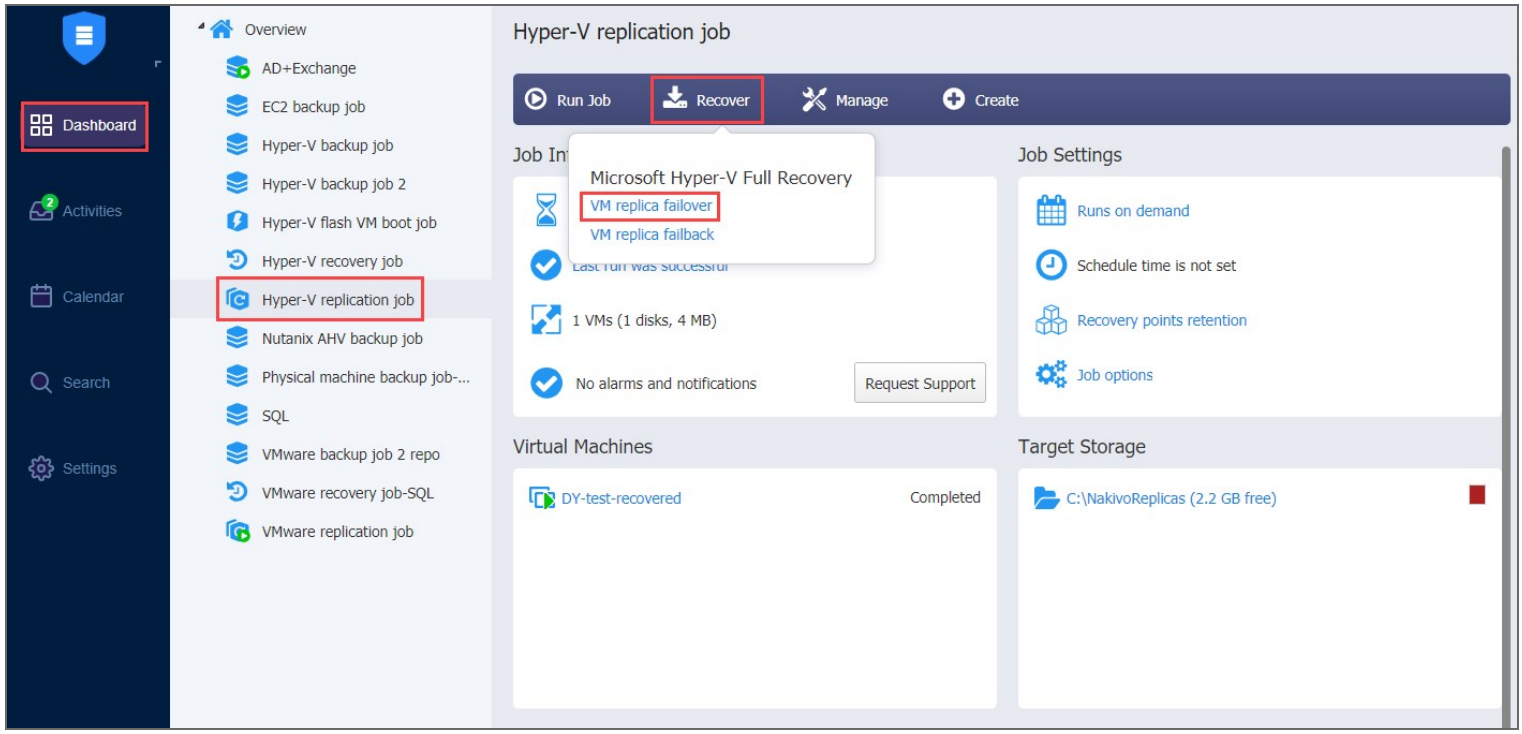

The **New Failover Job Wizard for Microsoft Hyper-V** opens.

Failover Job Wizard for Microsoft Hyper-V: Source

On **Source** page of the wizard, do the following:

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 262 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 265 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 next page.

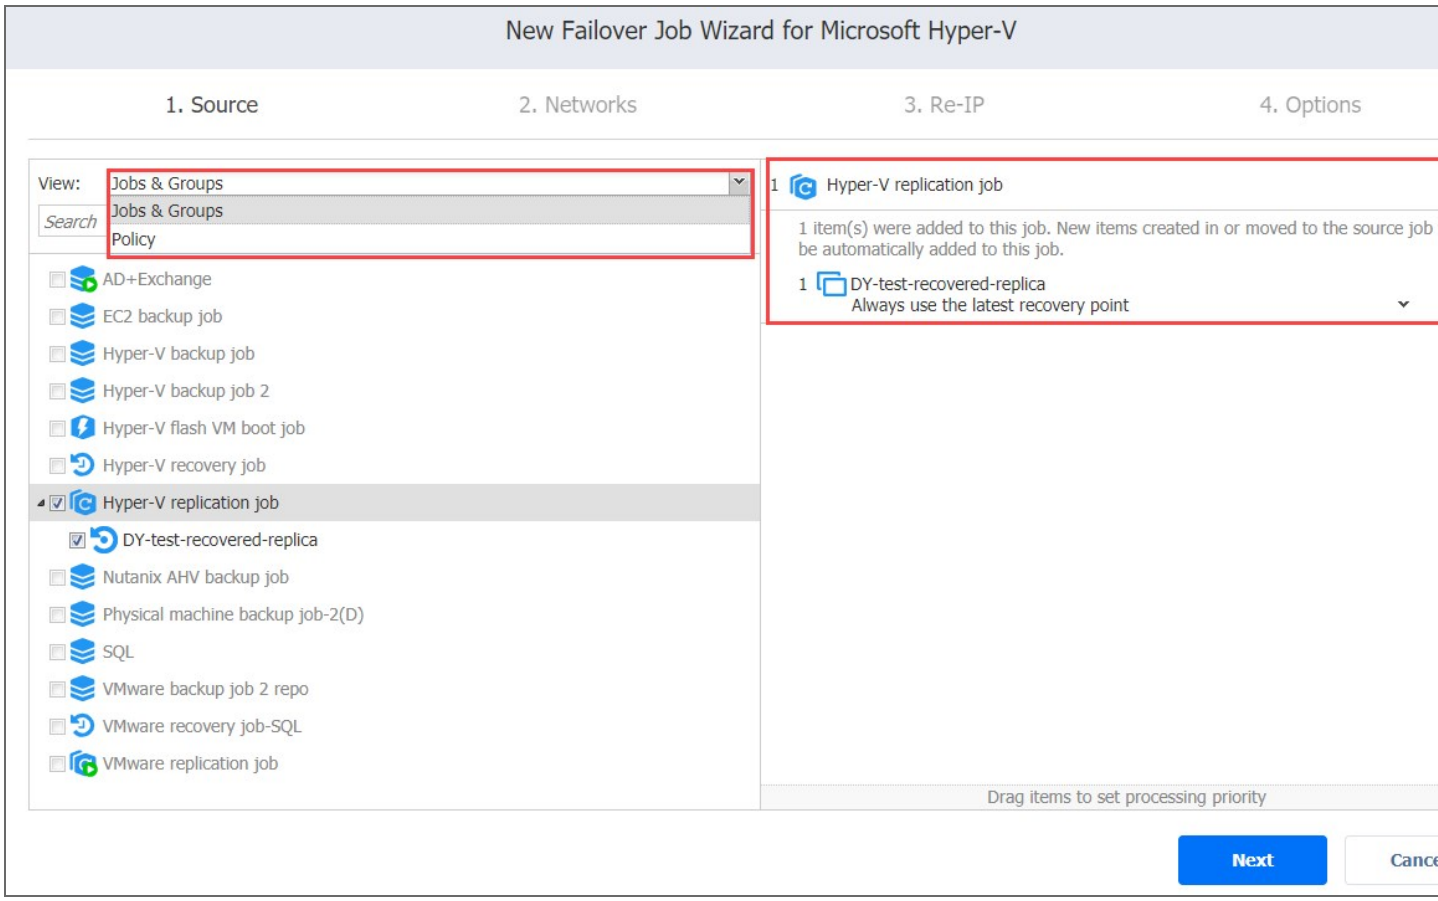

Failover Job Wizard for Microsoft Hyper-V: Networks

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**.

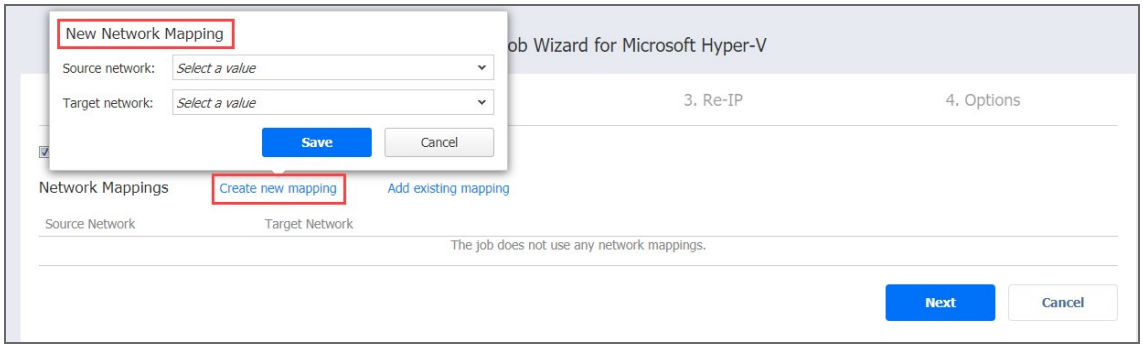

    - Add an existing network mapping:
        a. Click **Add exisitng mapping**.
        b. The **Network Mappings** dialog box opens. Choose an appropriate network mapping and close the dialog box.

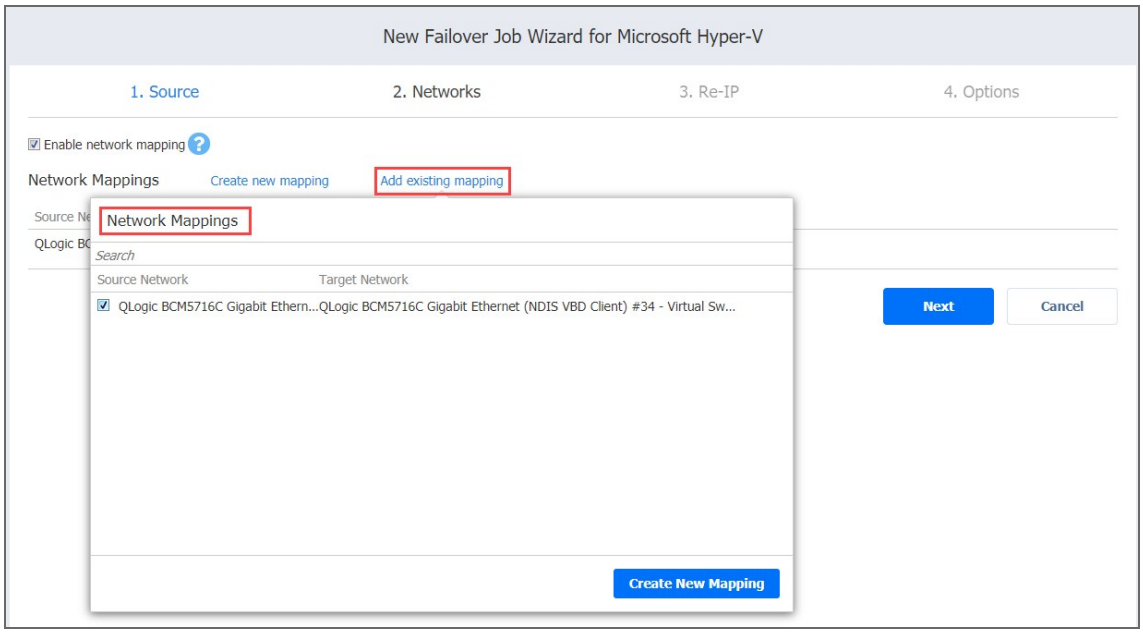

- Edit an existing network mapping:
    a. Hover over the required item in the **Network Mappings** list and click **Edit** to the right of the item.
    b. The **Edit Network Mapping** dialog box opens. Choose the required item from the **Target network** drop-down list and click **Save**.

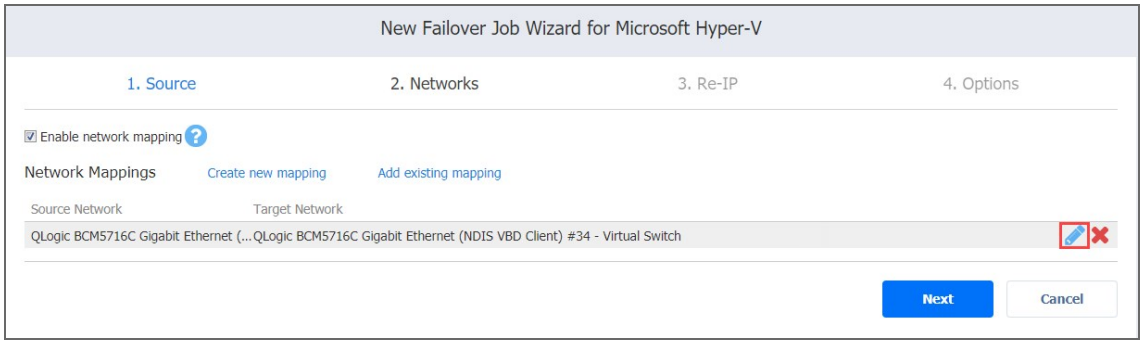

- Delete an existing mapping: Hover over the required item in the **Network Mappings** list and click the "X" icon to the right of the item.
- To leave the list of existing network mappings intact, fo to the next page of the wizard.

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**.

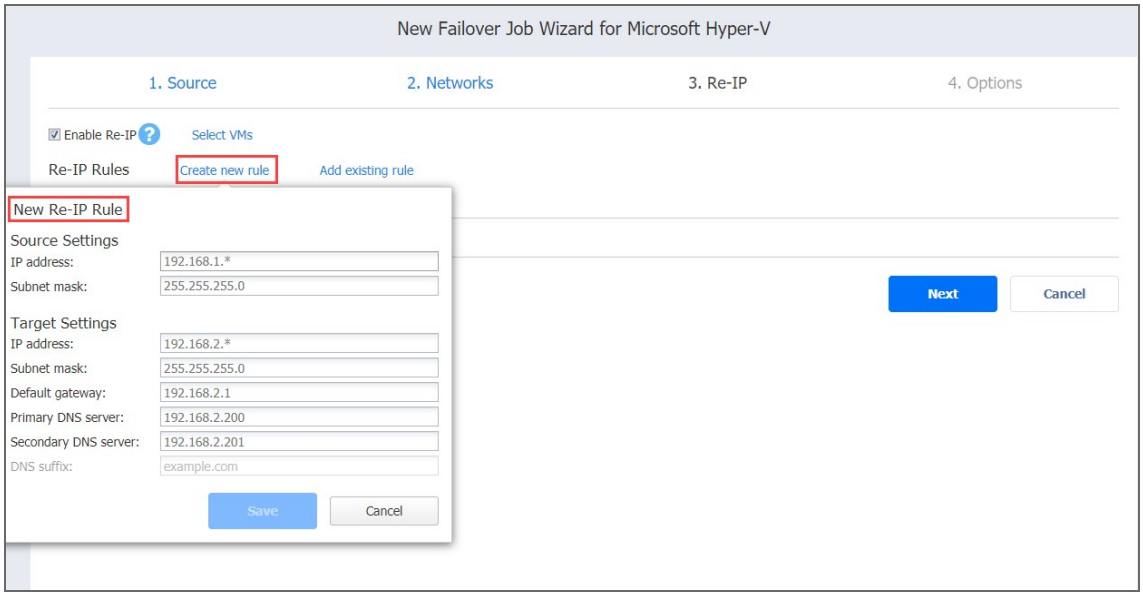

Note

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.

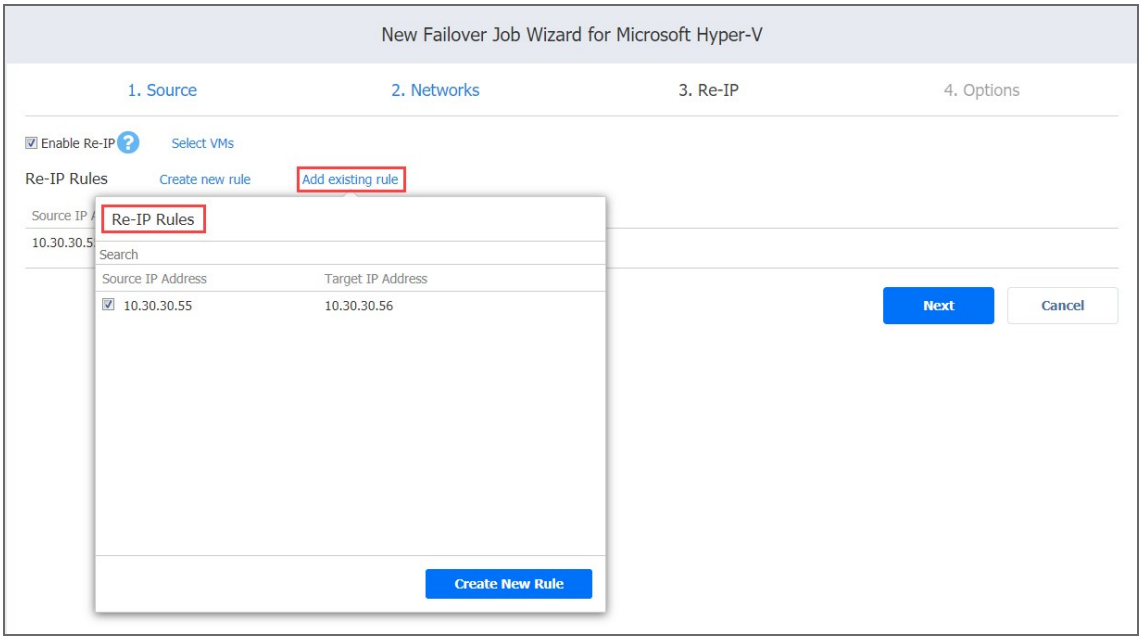

- Edit an existing re-IP rule:
  a. Hover over the required item in the **Re-IP Rules** list and click the **Edit** button to the right of the item.
  b. The **Edit Re-IP Rule** dialog box opens. Edit the required properties of the re-IP rule and click **Save**.
- Delete an existing rule: Hover over the required item in the **Re-IP Rules** list and click the **X** icon to the right of the item.

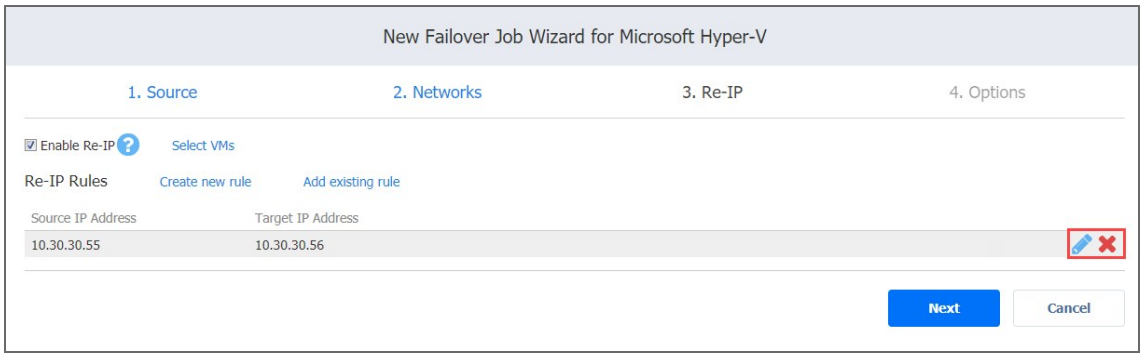

To leave the list of existing re-IP rules intact, go to the next page of the wizard.

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.

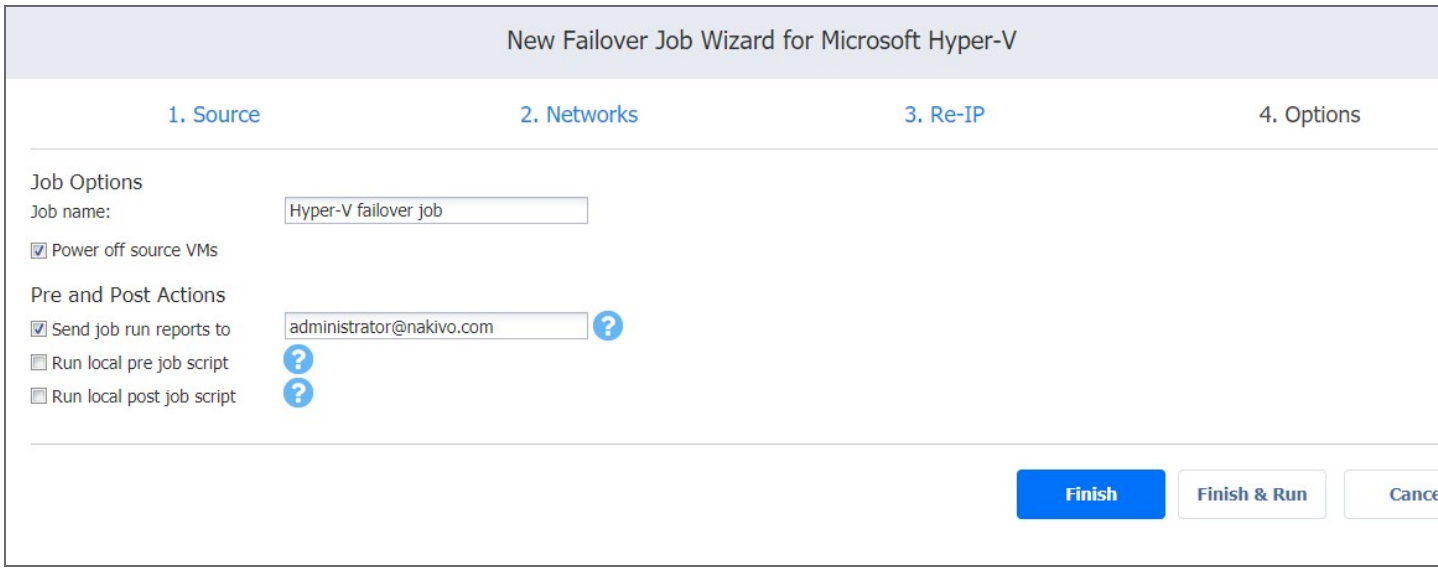

Completing the New Failover Job Wizard for Microsoft Hyper-V

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.

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" on page 705
- "Failback Job Wizard for Microsoft Hyper-V: Source" on page 706
- "Failback Job Wizard for Microsoft Hyper-V: Destination" on page 707
- "Failback Job Wizard for Microsoft Hyper-V: Networks" on page 709
- "Failback Job Wizard for Microsoft Hyper-V: Re-IP" on page 710
- "Failback Job Wizard for Microsoft Hyper-V: Options" on page 712

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 695 for details.

To start the failback procedure, click **Recover** in the **Dashboard** of NAKIVO Backup & Replication and then click **VM replica failback**.

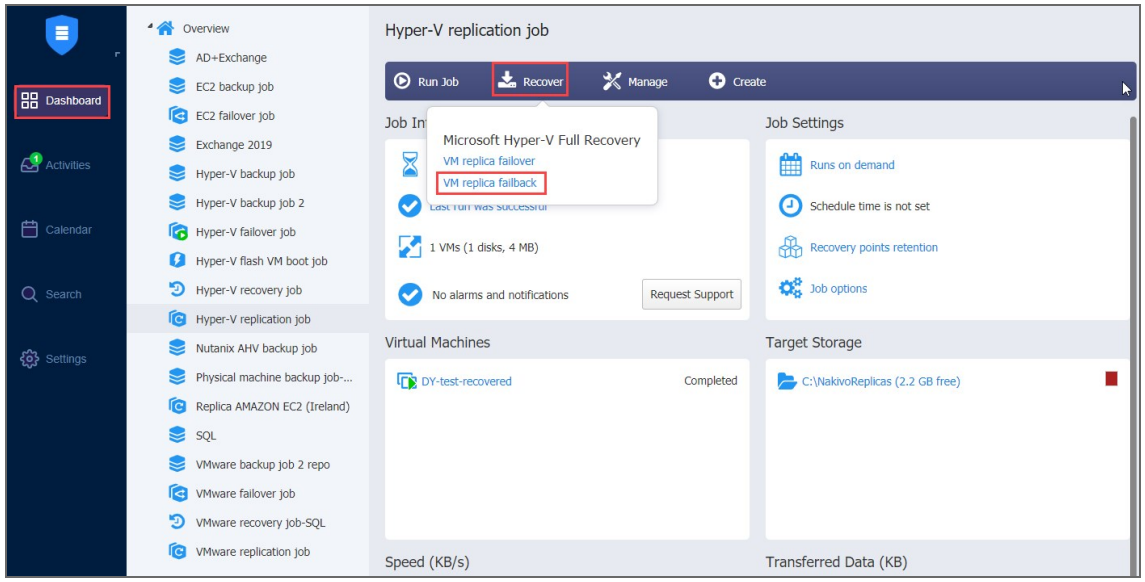

Failback Job Wizard for Microsoft Hyper-V: Source

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 262. 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 265.
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.

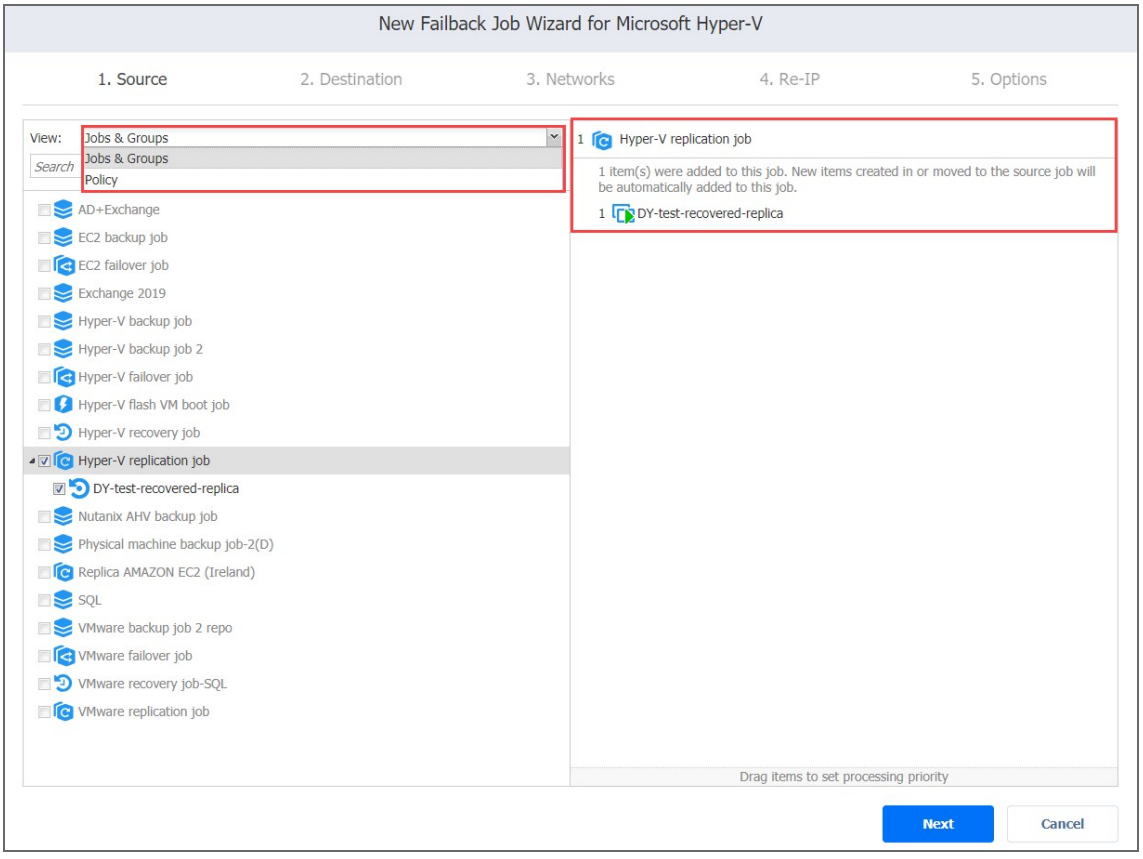

Failback Job Wizard for Microsoft Hyper-V: Destination

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.

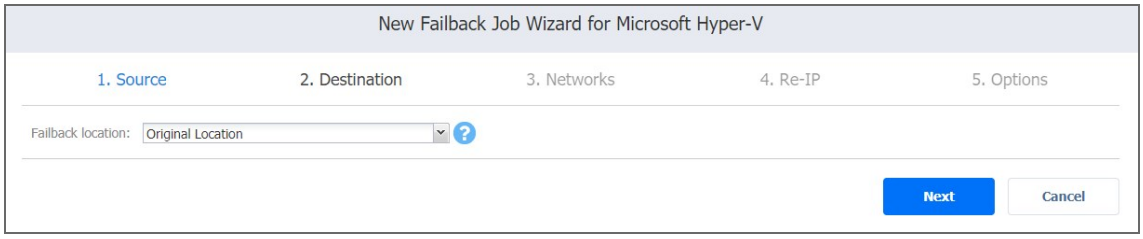

Creating New Hyper-V Failback Location

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.

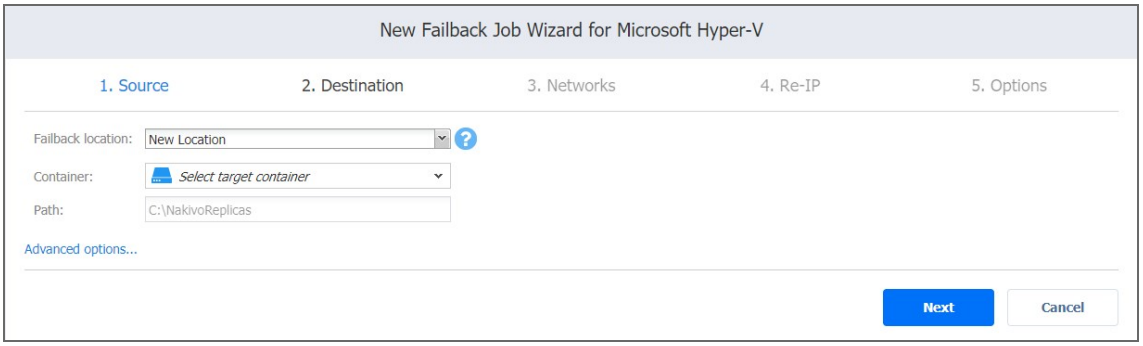

Creating a Different Hyper-V Failback Location for All Replicas

To create a new failover location for every replica you've added to the job:

1. Click **Advanced options.**
2. Select a target container and target disks.
3. Click **Next** to go to the next page.

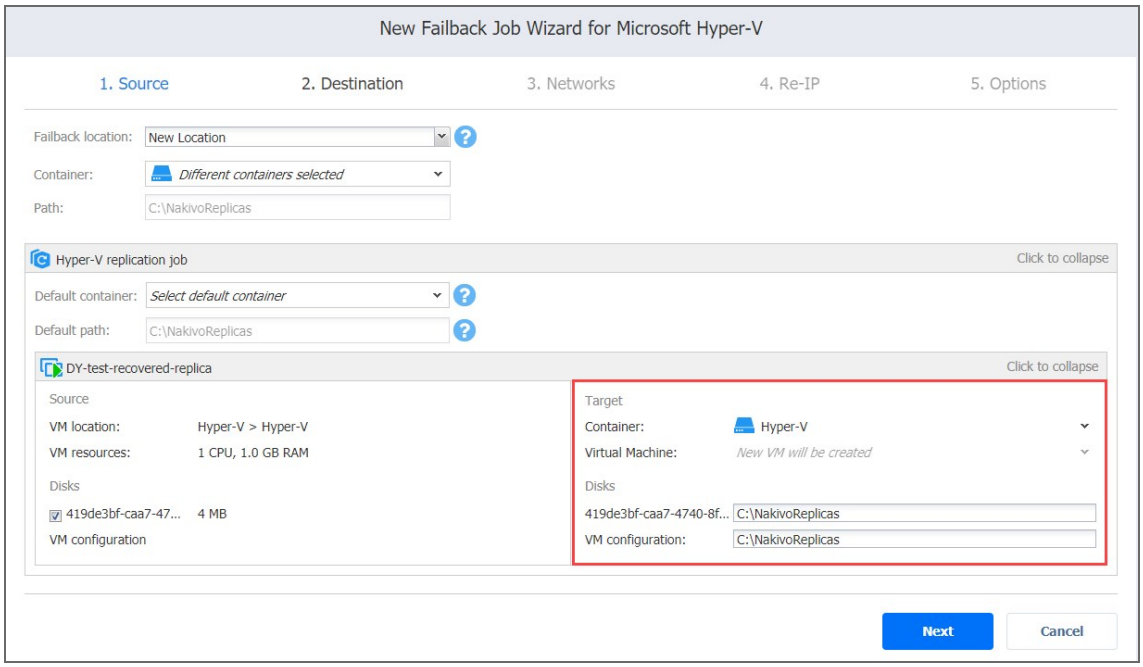

Failback Job Wizard for Microsoft Hyper-V: Networks

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.

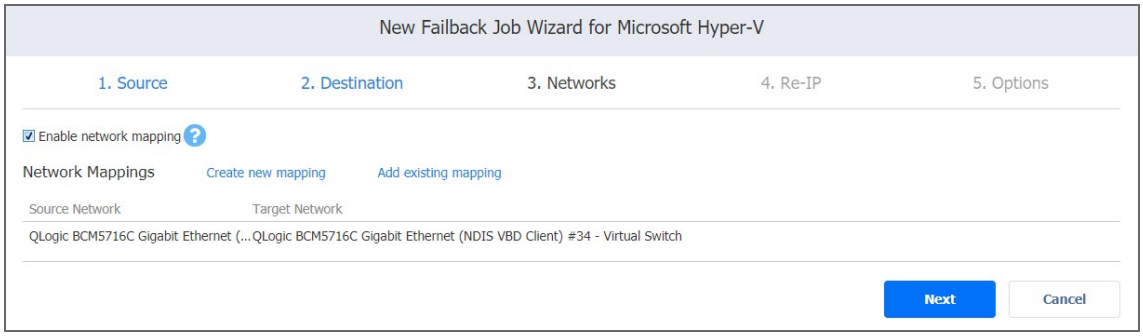

Failback Job Wizard for Microsoft Hyper-V: Re-IP

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.

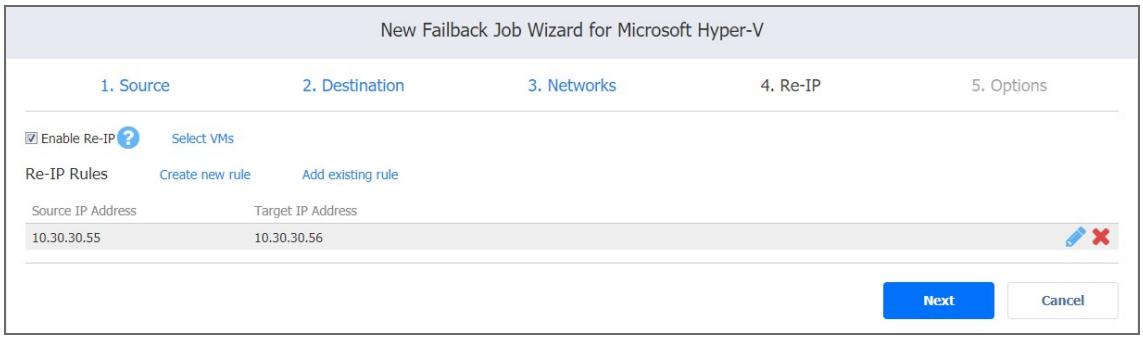

Failback Job Wizard for Microsoft Hyper-V: Options

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.

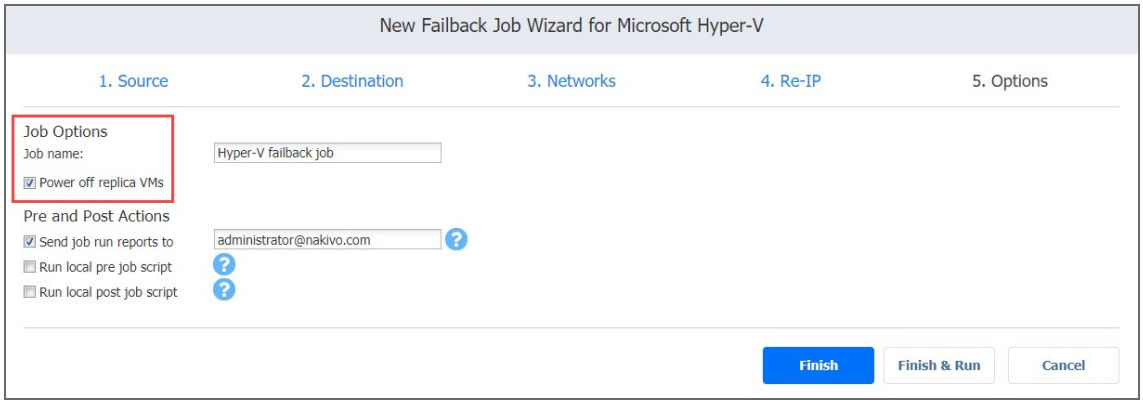

Pre and Post Actions

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.

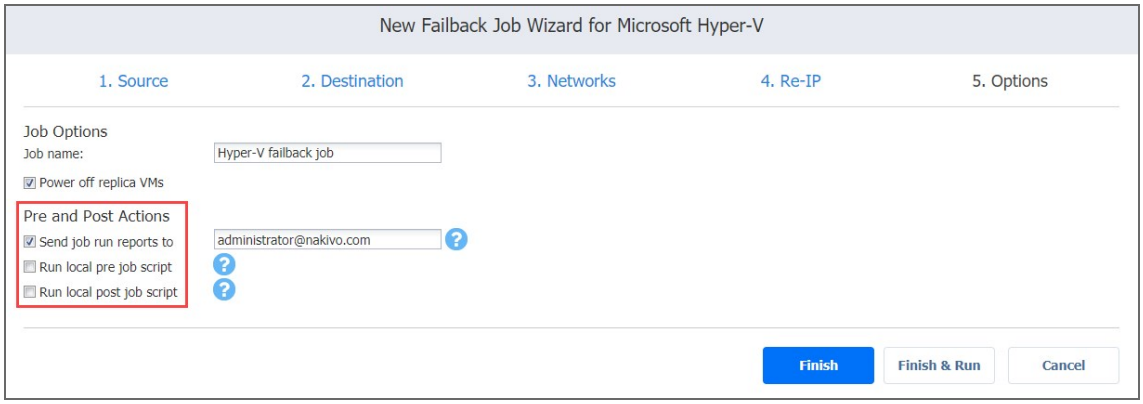

Note

If you click **Finish & Run**, you will have to define the scope of your job. Please refer to "Running Jobs on Demand" on page 250 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” on page 715
- “Running Site Recovery Job” on page 746

Creating Site Recovery Job

Creating a Site Recovery Job is done with a wizard and includes the following steps:

- "Starting Site Recovery Job Wizard" on page 716
- "Site Recovery Job Wizard: Actions" on page 717
- "Site Recovery Job Wizard: Networks" on page 736
- "Site Recovery Job Wizard: Re-IP" on page 738
- "Site Recovery Job Wizard: Test Schedule" on page 741
- "Site Recovery Job Wizard: Options" on page 745

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**.

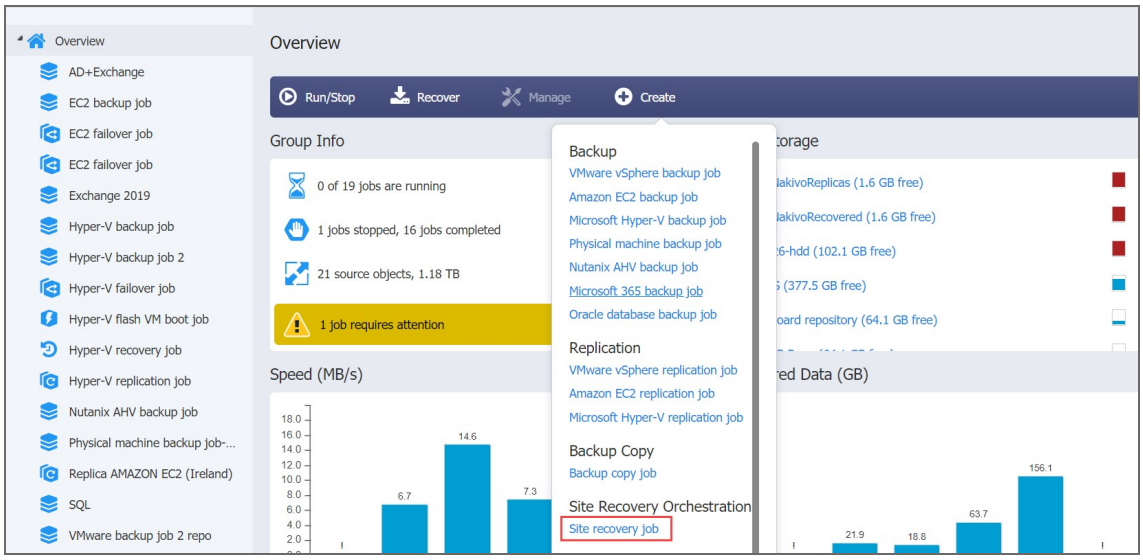

The **New Site Recovery Job Wizard** opens.

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 Hyper-V VMs Action" on page 720
- "Failover Hyper-V VMs Action" on page 720
- "Start VMs / Instances Action" on page 726
- "Stop VMs / Instances Action" on page 727
- "Run / Stop Jobs Action" on page 729
- "Run Script Action" on page 730
- "Attach / Detach Repository Action" on page 731
- "Send Email Action" on page 732
- "Wait Action" on page 733
- "Check Condition Action" on page 734

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.=

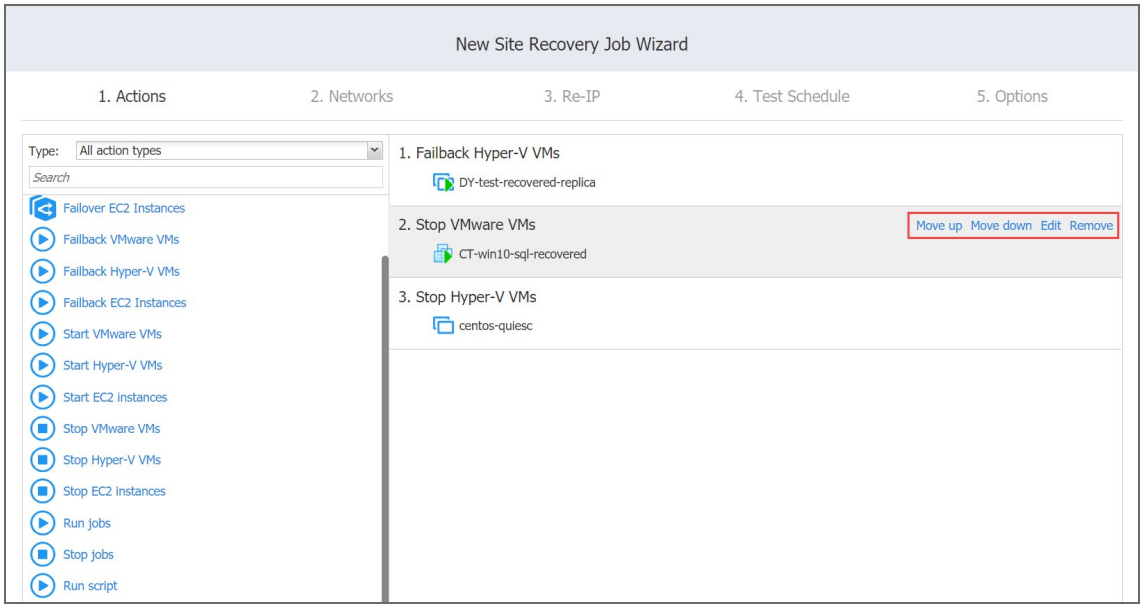

Options Common to Most Actions

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.
- **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 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 262 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 265 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.

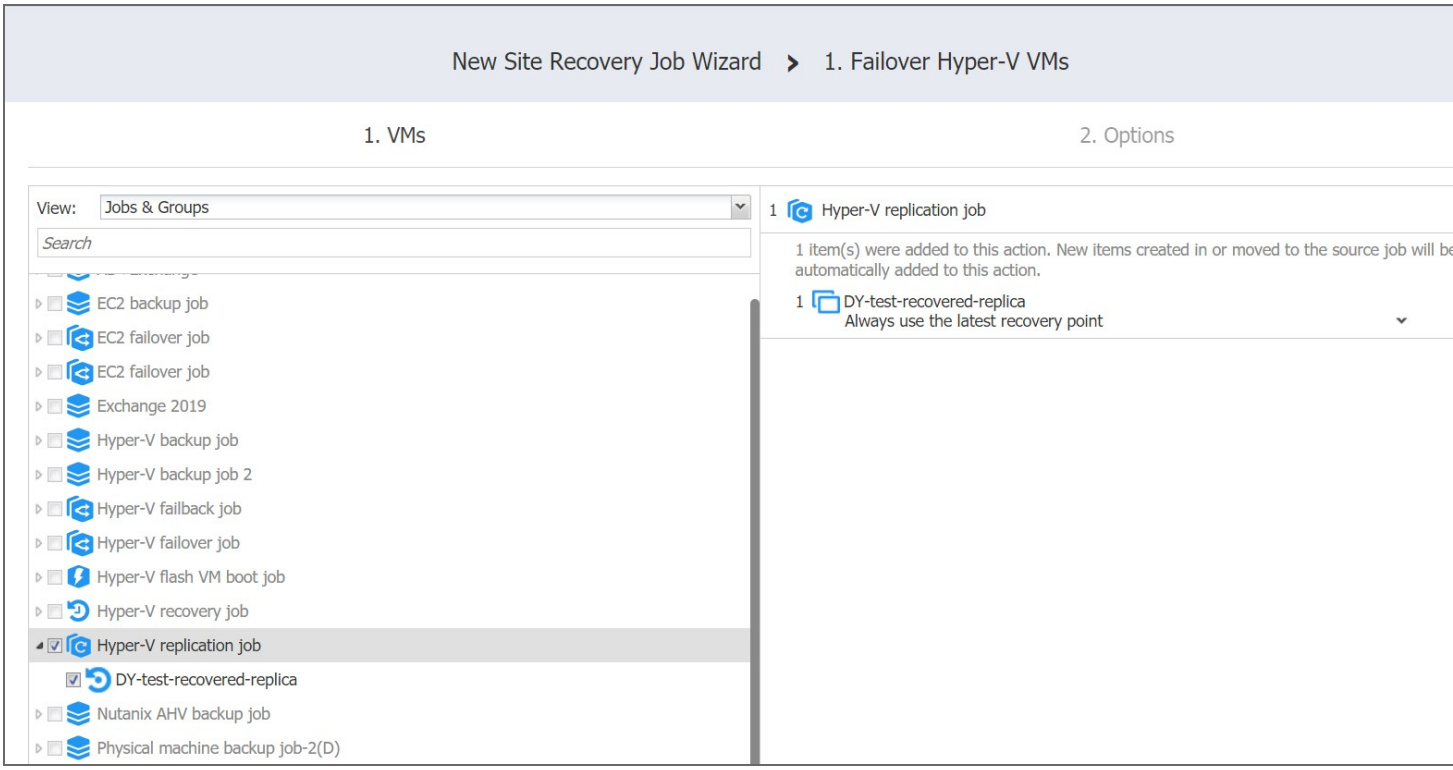

Failover Hyper-V VMs: Options

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 717 for details.
3. Click**Save**.

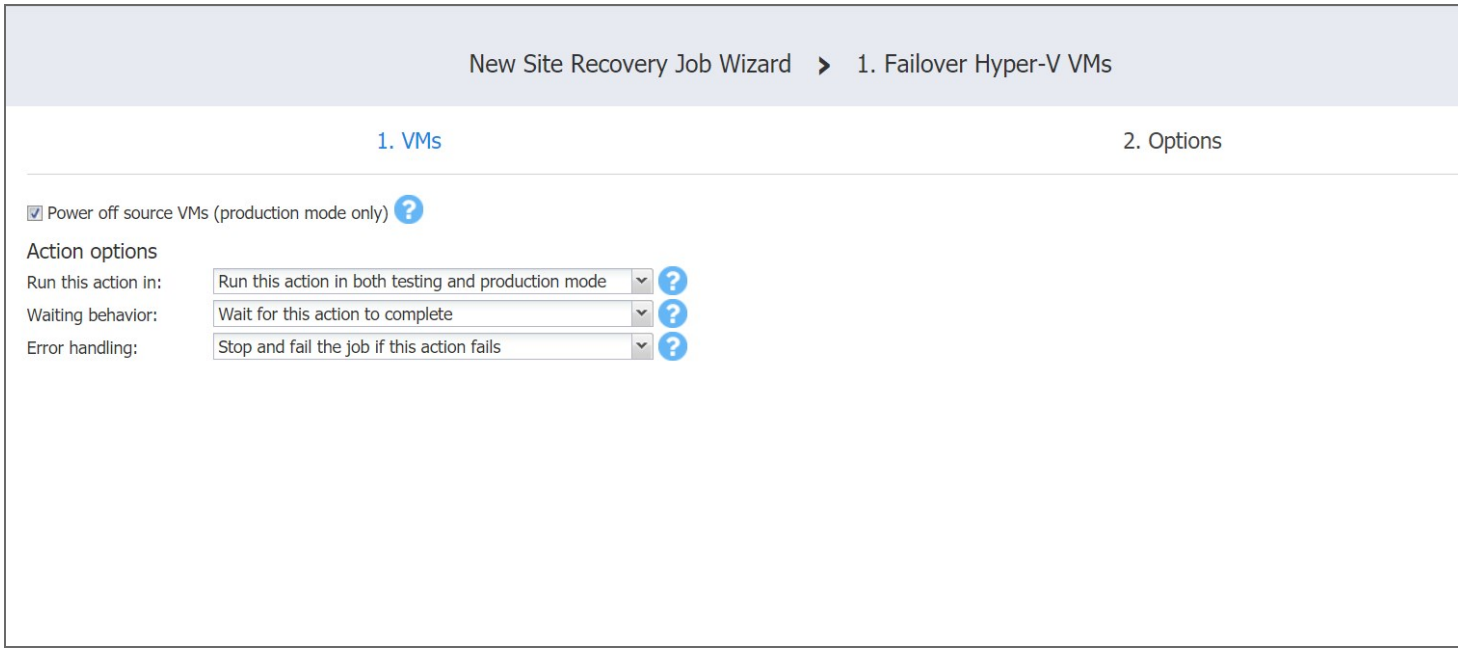

The **Options** page closes and your **Failover Hyper-V VMs**action is added to the Site Recovery Job.

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 262 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 265 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.

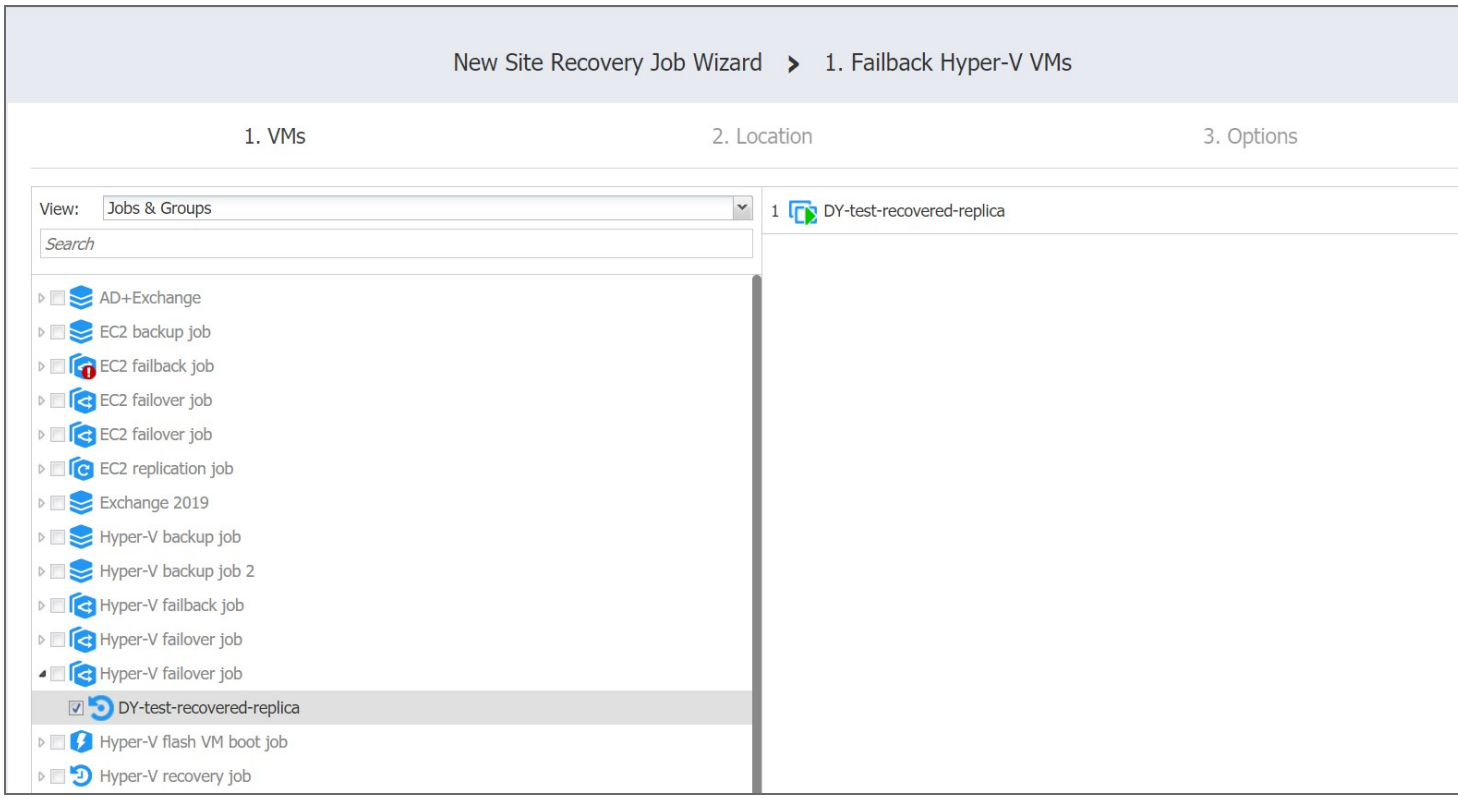

Refer to the following subsections on providing a location for your **Failback Hyper-V VMs** action:

- "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.

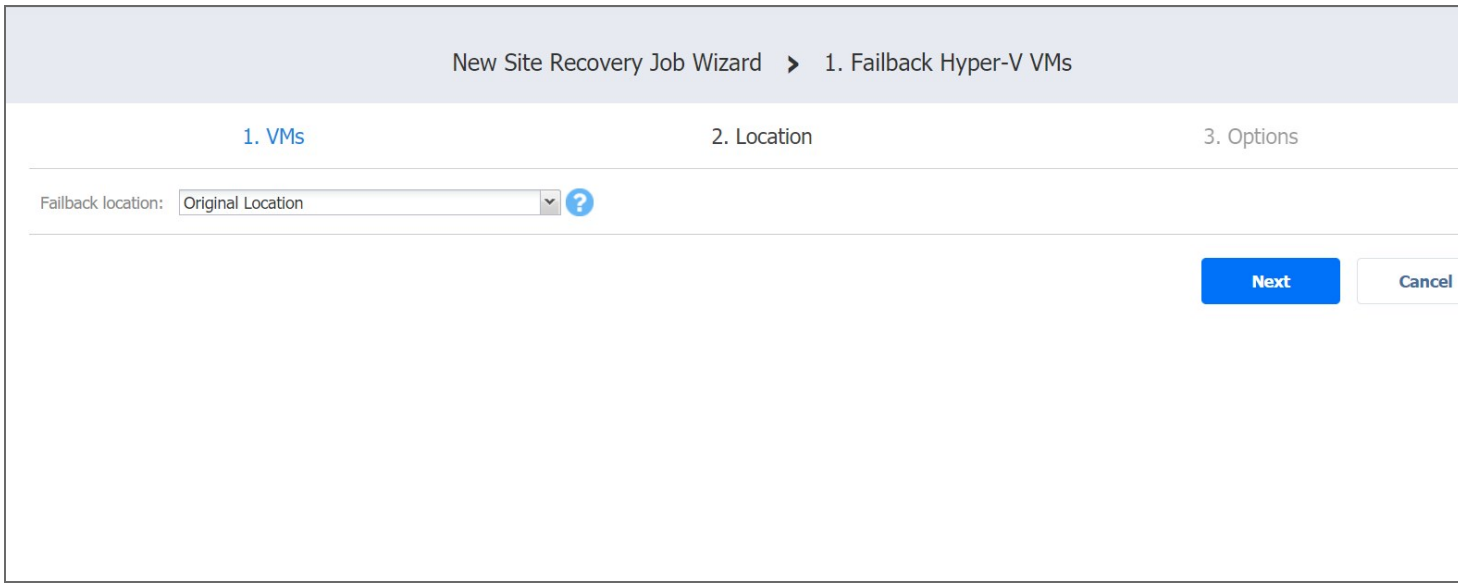

Creating New Hyper-V Failback Location

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.

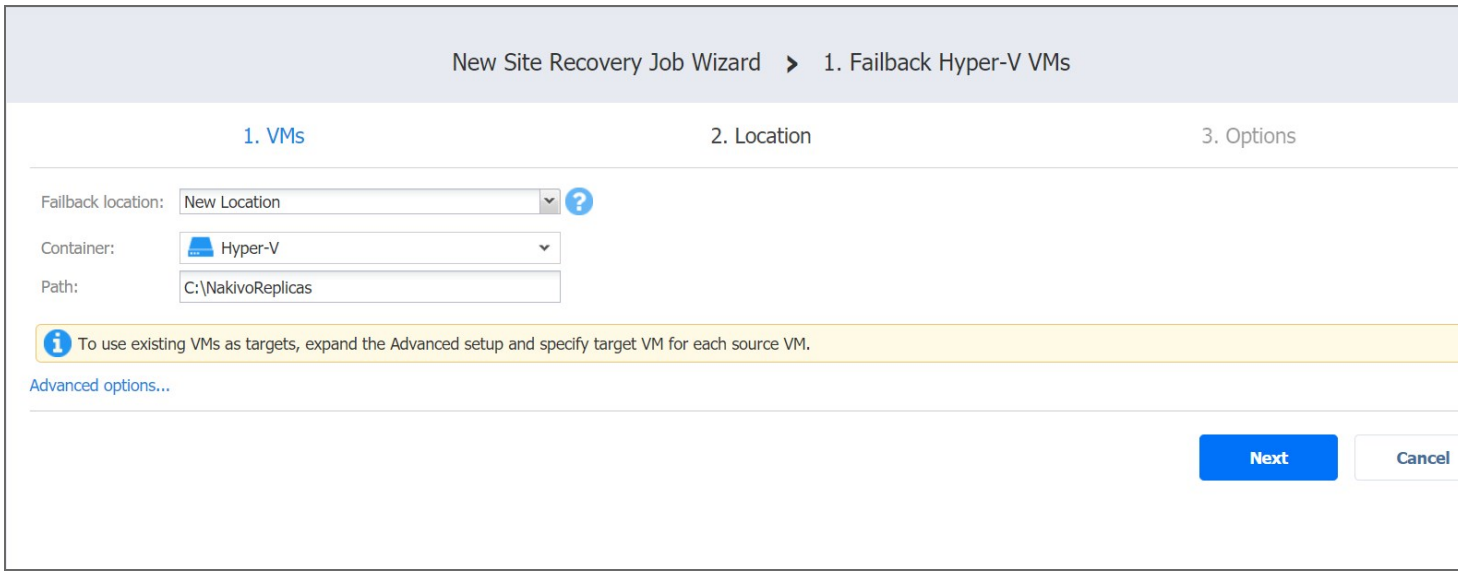

Reusing Existing Hyper-V Failback Location

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 717 for details.

3. Click **Save**.

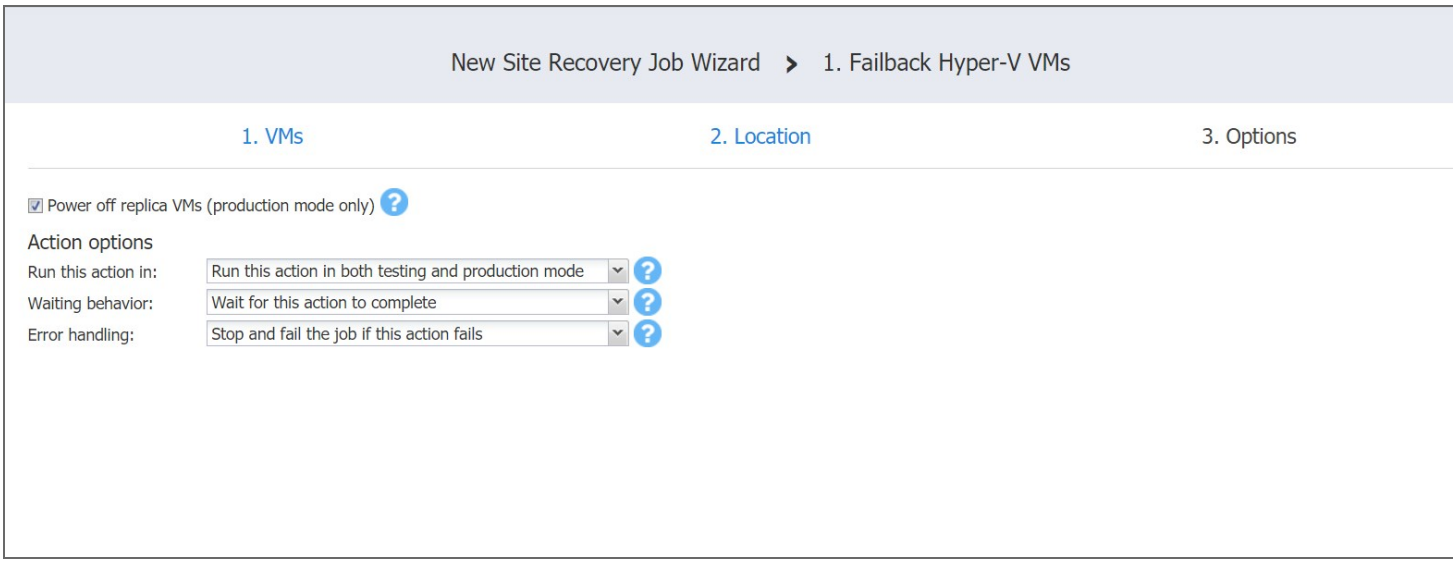

The **Options** page closes, and your Failback Hyper-V VMs action is added to the Site Recovery job.

Start VMs / Instances Action

The topic covers the following actions of a Site Recovery Job:

- Start VMware VMs
- Start 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.

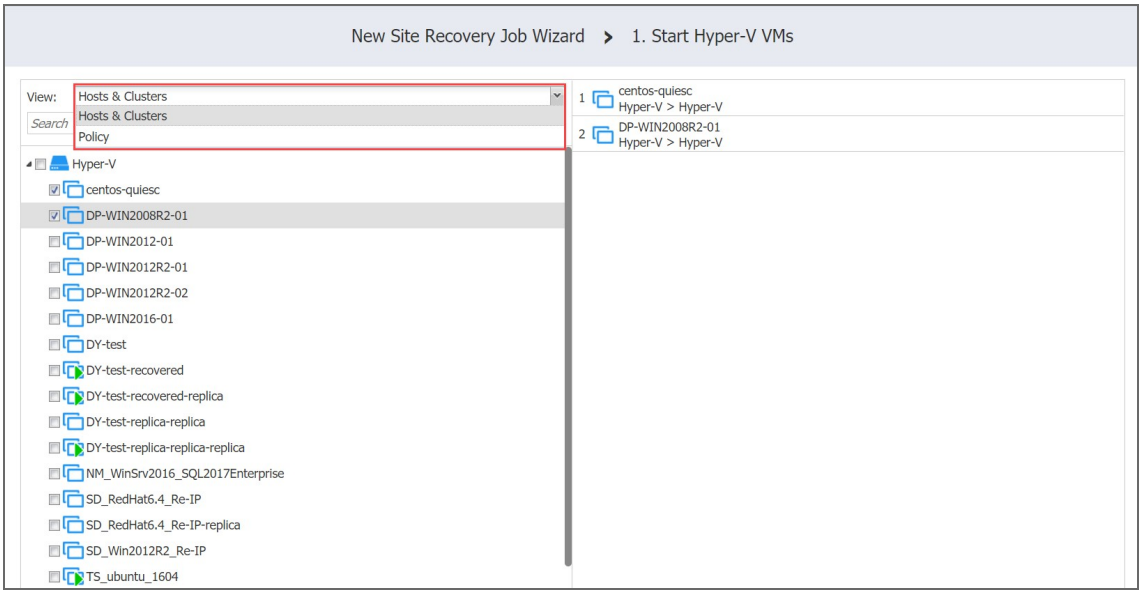

2. Optionally, you can filter the inventory tree by entering a string into the **Search** box. You can enter a part of or the entire name of the item.
3. Select at least one VMware/Hyper-V VM or EC2 instance in the inventory tree.
4. Set the action options. Refer to "Site Recovery Job Wizard: Actions" on page 717 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 262 for details.

     Note

     Switching to an alternative view resets your selection in the current view.

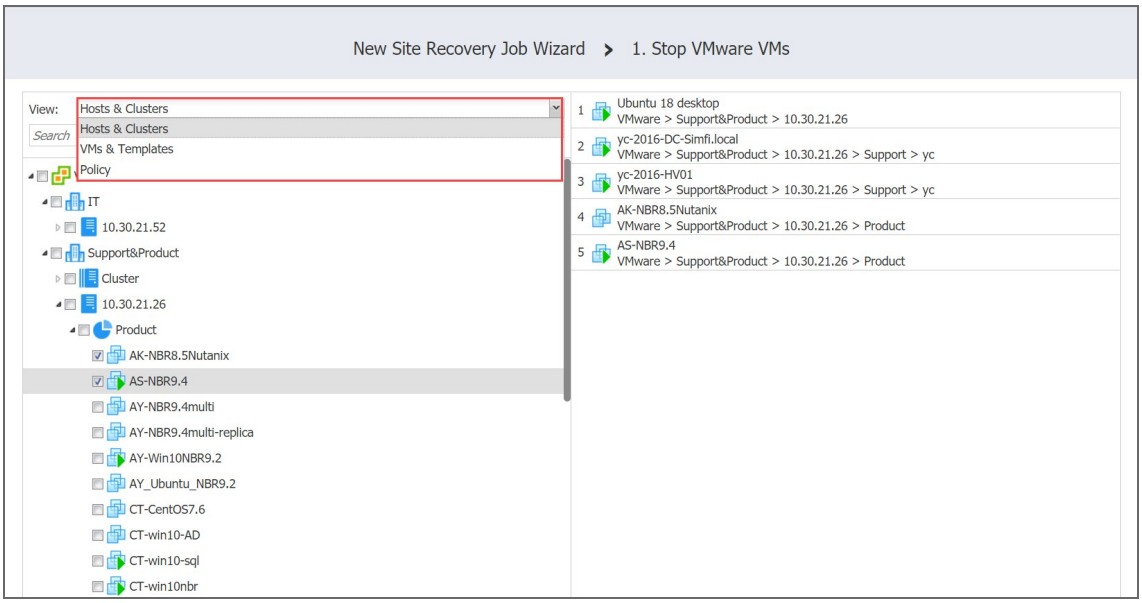

2. Optionally, you can filter the inventory tree by entering a string into the **Search** box. You can enter a part of the entire name of the item.
3. Select at least one VMware/Hyper-V VM or EC2 instance in the inventory tree.
4. Set action options. Refer to "Site Recovery Job Wizard: Actions" on page 717 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 717 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 717 for details.
3. Click **Save**.

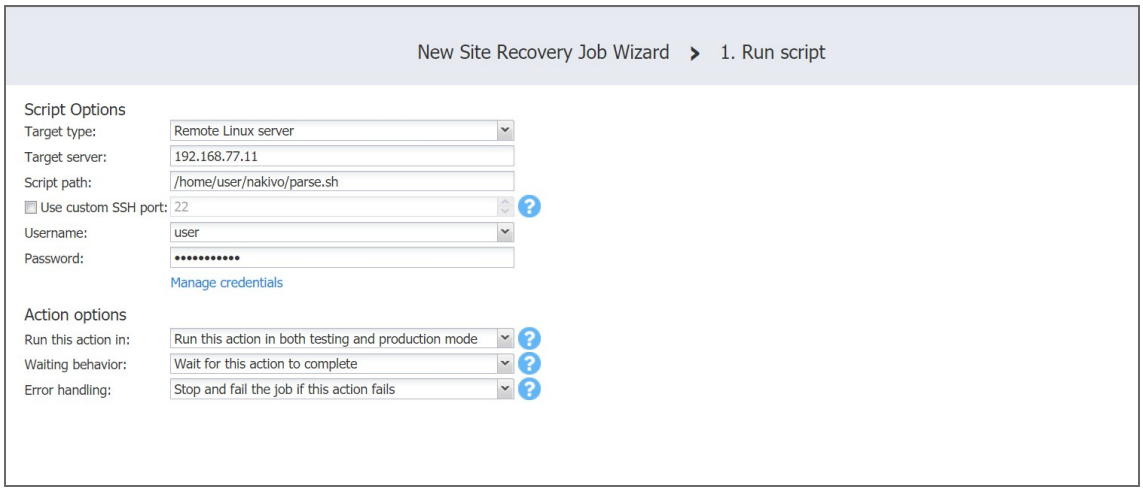

The page of the wizard closes and your **Run Script** action is added to the Site Recovery Job.

Attach / Detach Repository Action

The topic covers the following actions of a Site Recovery Job:

- Attach Repository
- Detach Repository

On the **Attach / Detach Repository** page that opens from the **Actions** page of the Site Recovery Job Wizard, do the following:

1. Select a repository in the list of repositories.
2. Specify action options. Refer to Options Common to Most Actions for details.
3. Click **Save**.

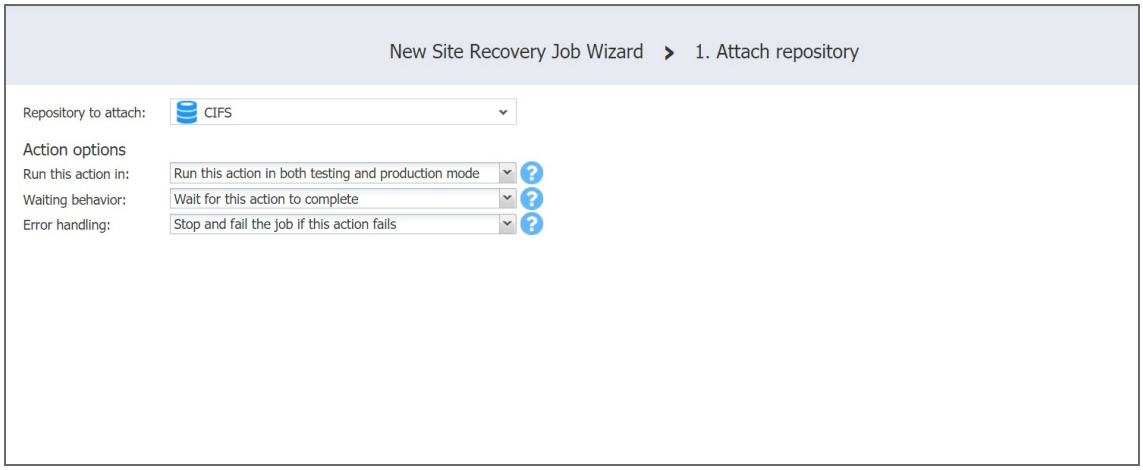

The page of the wizard closes and your **Attach / Detach Repository** action is added to the site recovery job.

Send Email Action

On the **Send Email** page of the wizard, do the following to add the action to your site recovery job:

1. In the **To** box, enter a valid email address for the recipient.
2. In the **Cc** box, optionally you can enter an email address of the carbon copy recipient.
3. In the **Subject** box, optionally you can enter a subject of the mail.
4. Optionally, you can add attachments to your mail with the **Browse** button.
5. Enter your message text in the email body box.
6. Set up action options. Refer to "Site Recovery Job Wizard: Actions" on page 717 for details.
7. Click **Save**.

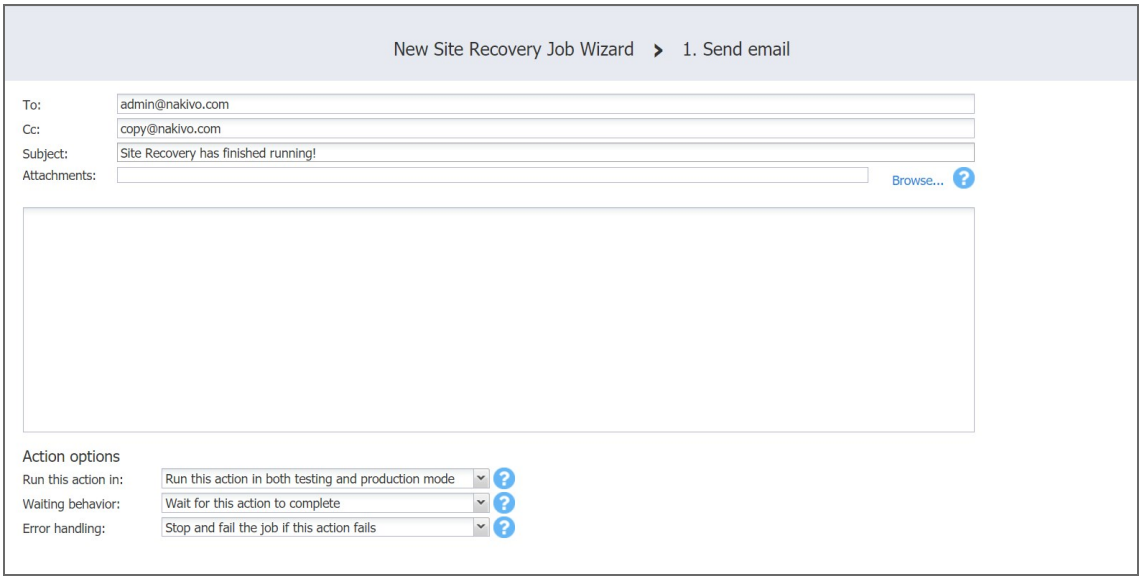

The page of the wizard closes, and your **Send Email** action will be added to the site recovery job.

Wait Action

To add a wait action to your Site Recovery job, on the **Wait** page of the wizard, do the following:

1. Enter the time to wait in minutes or hours.
2. Set the action options. Refer to "Site Recovery Job Wizard: Actions" on page 717 for details.
3. Click **Save**.

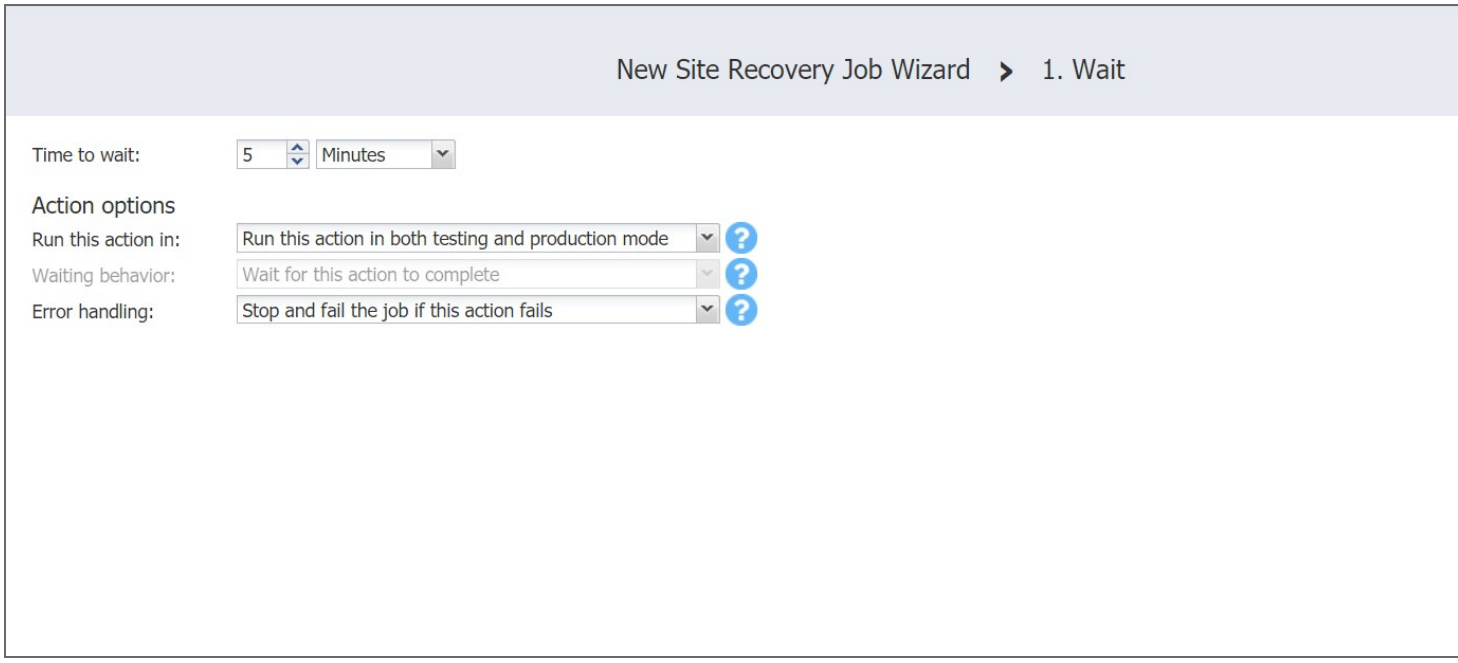

The page of the wizard closes and your **Wait** action is added to the Site Recovery job.

Check Condition Action

On the *Check Condition* page of the wizard, do the following to add the action to your site recovery job:

1. Choose a condition type from the list:
    - **Resource exists**: This condition checks whether the specified resource exists. With this option, you have to set the following:
        a. Choose a resource type from the list:
            - **VMware VM**
            - **Hyper-V VM**
            - **EC2 Instance**
        b. Define your identification method using two lists:
            i. In the first list, choose either **Name** of **ID**.
            ii. In the second list, choose one of the following:
                - **Equals**
                - **Contains**
                - **Starts with**
                - **Ends with**
        c. Enter your condition criterion in the **Search string** box.
    - **Resource is running**: This condition checks whether the specified resource is running. With this option, you have to set the options as described for the **Resource exists** option above.
    - **IP/hostname is reachable**: This condition checks whether the specified IP/hostname is reachable. With this option, you have to enter the following:
        a. Choose the source endpoint among the following:
            - **This server (Director)**: With this option, the condition checks whether the Director can reach the **IP/hostname** specified in the box below.
            - **Remote transporter**: With this option, the condition checks whether one of your remote Transporters can reach the **IP/hostname** specified in the box below.
        b. In the **IP/hostname** box, enter an IP address or a host name of the resource to be checked for reachability.
2. In the **Action if True** section of the page, choose an action to be taken if the condition criterion is satisfied:
    - **Continue site recovery job**: Your Site Recovery job will be continued.
    - **Stop and fail site recovery job**: Your Site Recovery job will be stopped as failed.
    - **End site recovery job successfully**: Your Site Recovery job will be ended as successful.
    - **Go to another site recovery Job action**: Another action of your Site Recovery job will be initiated. If you choose this action type, a new box will open to allow you to choose the necessary action.

3. In the **Action if False** section of the page, choose an action to be taken if the condition criterion is not satisfied. Available options are similar to those described in the **Action if True** section above.
4. Click **Save**.

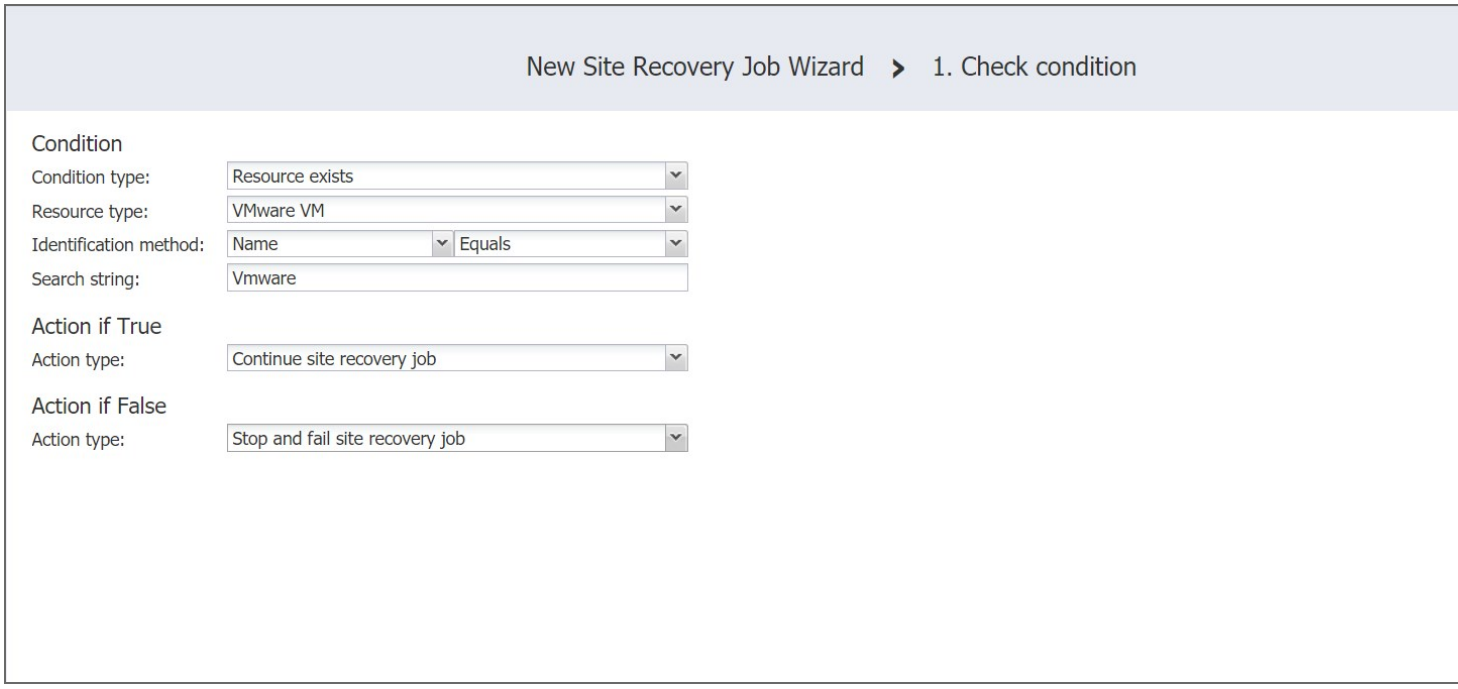

The page of the wizard closes and your **Check Condition** action is added to the Site Recovery job.

Site Recovery Job Wizard: Networks

On the **Networks** page of the **Site Recovery Job Wizard** you can map source VM virtual networks to appropriate target virtual networks and test networks.

Please proceed as follows:

1. Select **Enable network mapping**.

   Note

   A failover or a failback action needs to be on the actions list to allow enabling network mapping for your site recovery job.

2. The *Network Mapping* section opens. You have the following options:

   - Create a new network mapping:

     a. Click **Create new mapping**.

     b. The **New Network Mapping** dialog opens. Choose a source network, a target network and a test network, and click **Save**.

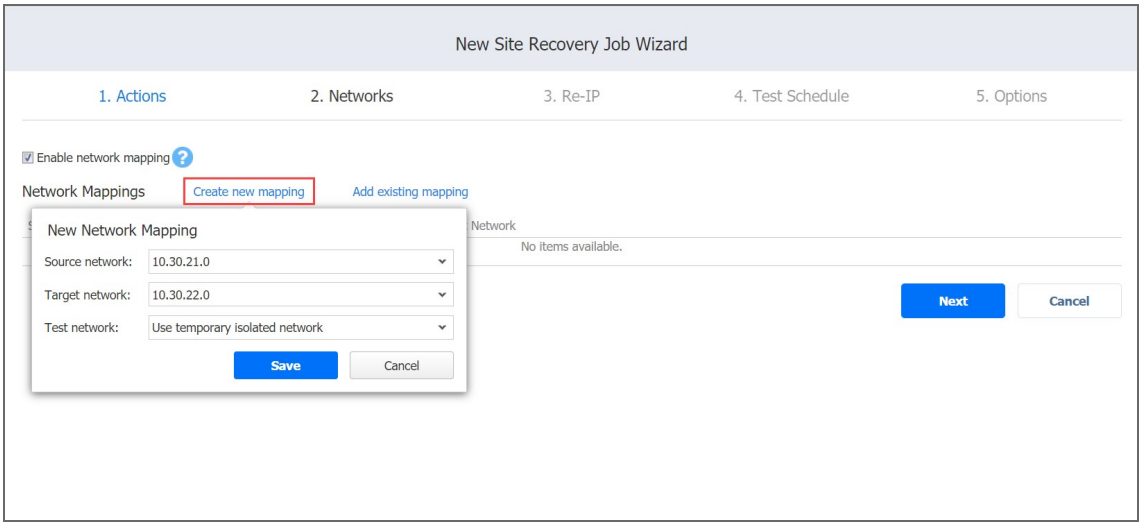

   - Add an existing network mapping:

     a. Click **Add existing mapping**.

     b. The **Network Mappings** dialog opens. Choose an appropriate network mapping and close

the dialog box.

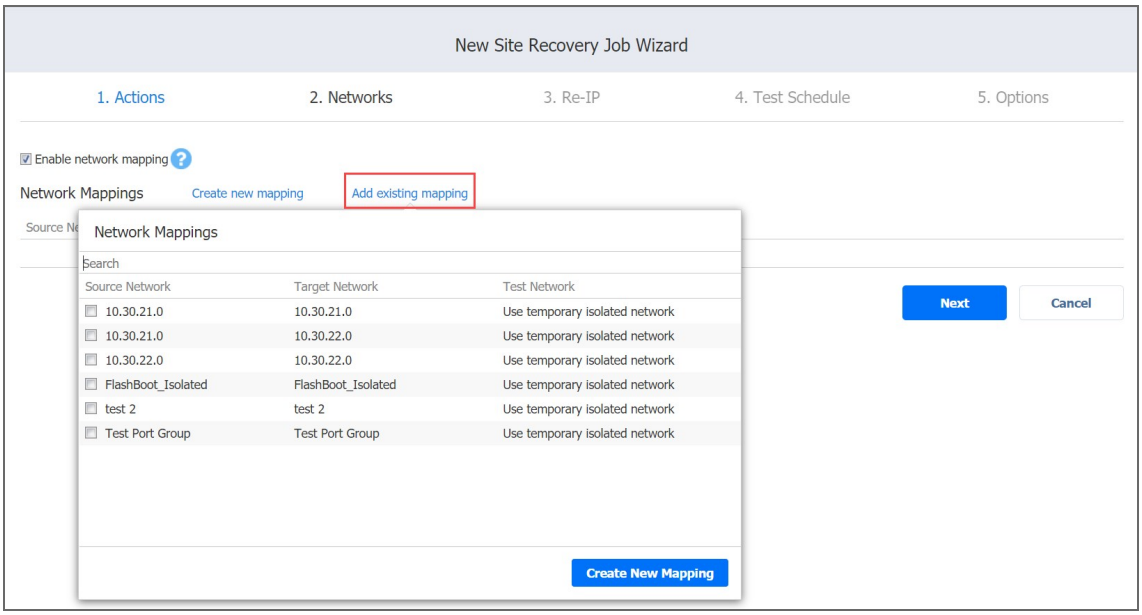

- Edit an existing network mapping:
  a. Hover the pointer over the necessary item in the **Network Mappings** list and then click the **Edit** button to the right of the item.
  b. The **Edit Network Mapping** dialog box opens. Choose appropriate items from the **Target network** and the **Test network** lists and then click **Save**.

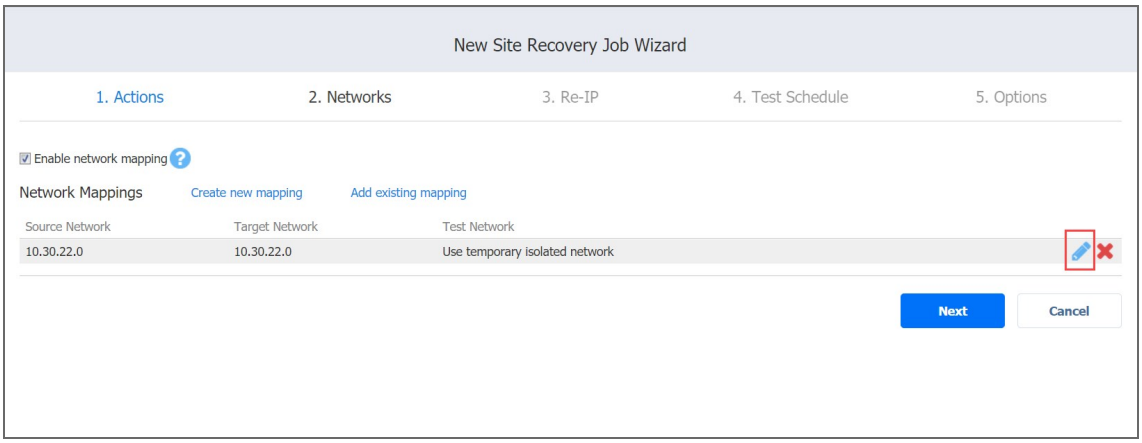

- Delete an existing mapping: hover the pointer over the necessary item in the **Network Mappings** list and then click the **Delete** icon to the right of the item.
- To leave the list of existing network mappings intact, go to the next page of the wizard.

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

Site Recovery Job Wizard: Re-IP

In the **Re-IP** page of the **Site Recovery Job Wizard** you can map a source VM IP address to a specific target IP address.

Please proceed as follows:

1. Select **Enable Re-IP**.

   Note

   A failover or a failback action needs to be in the actions list to allow enabling Re-IP for your site recovery job.

2. The *Re-IP Rules* section opens. Click the **Select VMs** link.
3. The **Re-IP** dialog opens. In the list of your source VMs, select at least one, select the credentials to be used for each VM, and close the dialog.

   Note

   Re-IP rules will be applied only to VMs that have a static IPv4 address configured.

4. You have the following options:
   - Create a new rule:
     a. Click **Create new rule**.
     b. The **New Re-IP Rule** dialog opens. Enter source and target settings for the Re-IP rule and click **Save**.

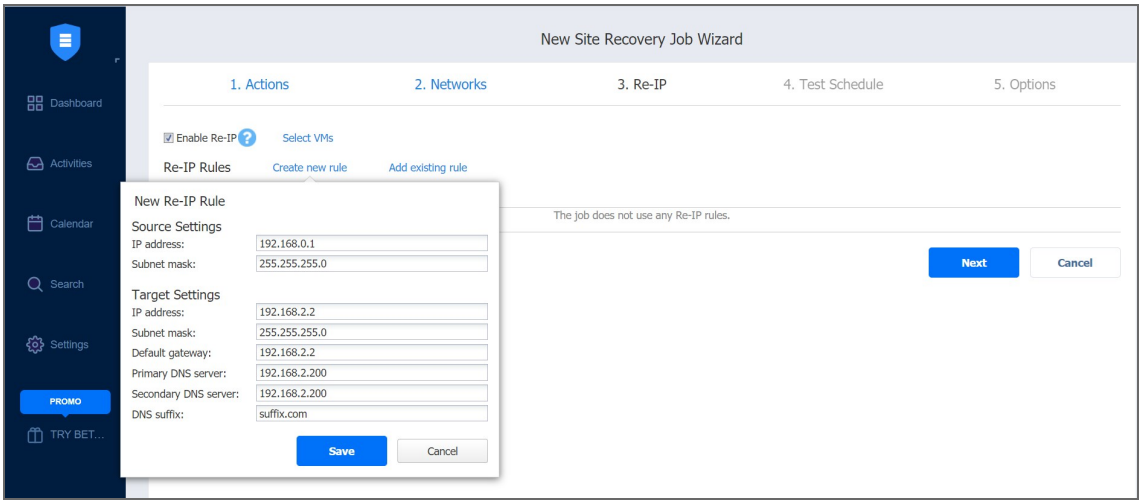

Note

You can use wildcards for IP addresses. For example, when the 192.168.1.* -> 10.30.30.* Re-IP rule is available, it will change the source VM IP address like 192.168.1.50 to the 10.30.30.50 IP address, for your site recovery job.

When there are several Re-IP rules applicable to your source VM, the application will define the most suitable one and apply it to the source VM IP address.

- Add an existing rule:
  a. Click **Add existing rule**.
  b. The **Re-IP Rules** dialog box opens. Choose an appropriate Re-IP rule and close the dialog.

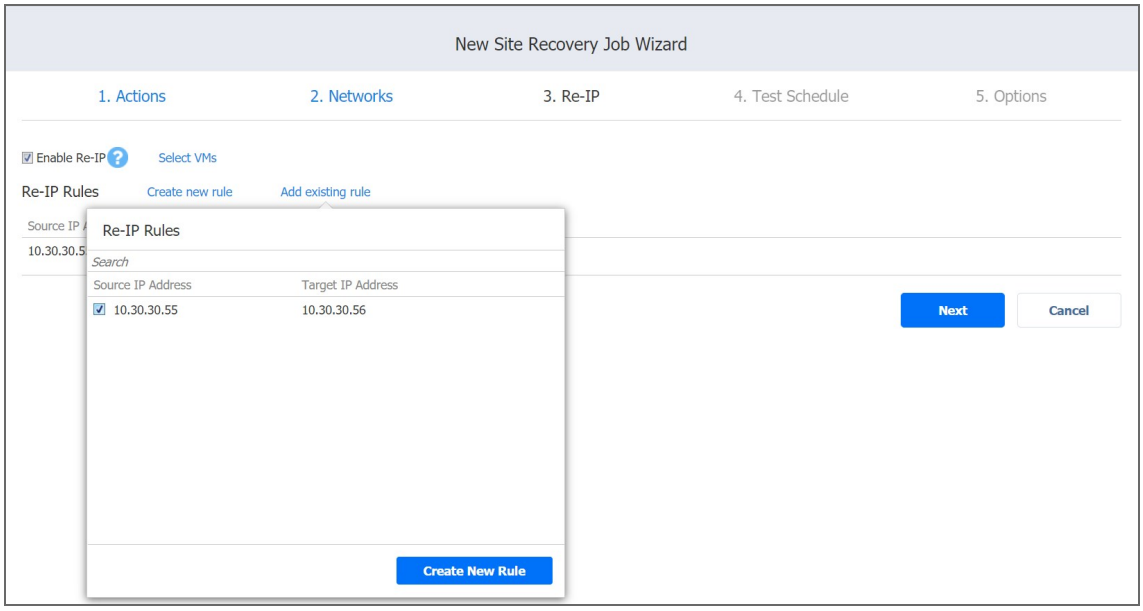

- Edit an existing Re-IP rule:
  a. Hover the pointer over the necessary item in the **Re-IP Rules** list and then click the **Edit** button to the right of the item.
  b. The **Edit Re-IP Rule** dialog opens. Edit the necessary properties of the Re-IP rule and then click **Save**.

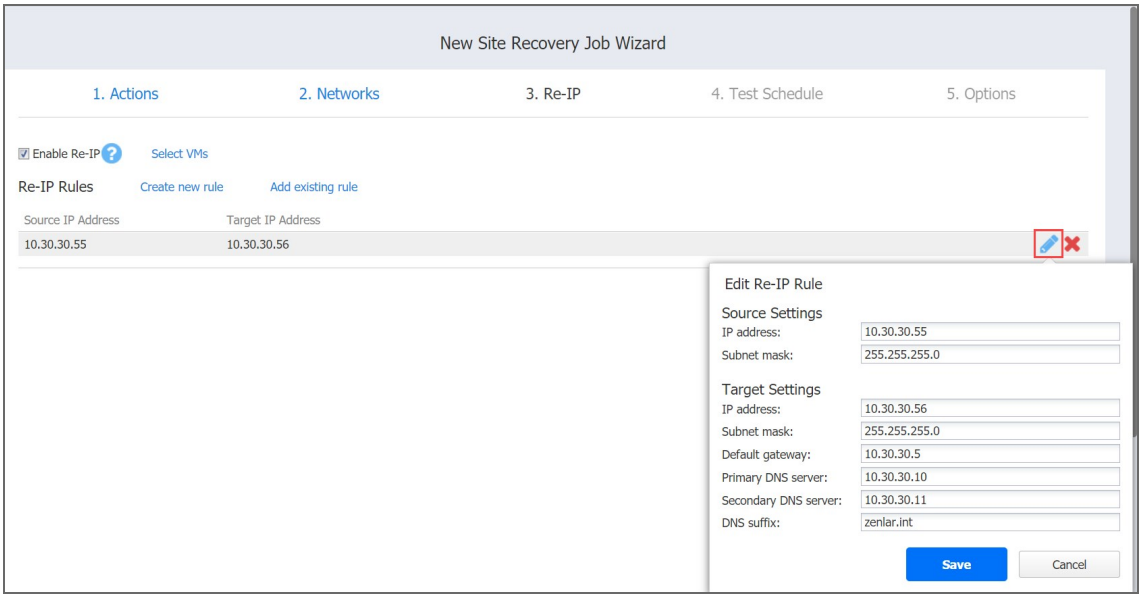

- Delete an existing mapping: hover the pointer over the necessary item in the **Re-IP Rules** list and then click the **Delete** icon to the right of the item.



- To leave the list of existing Re-IP rules intact, go to the next page of the wizard.

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

Site Recovery Job Wizard: Test Schedule

On the **Test Schedule** page of the **Site Recovery Job Wizard** you can schedule testing your site recovery job. Please refer to the following sections for details:

- "Disabling Site Recovery Job Test Schedule" below
- "Daily Site Recovery Job Testing" below
- "Monthly or Yearly Site Recovery Job Testing" on the next page
- "Periodic Site Recovery Job Testing" on page 743
- "Chained Site Recovery Job" on page 743
- "Additional Schedule" on page 744

Disabling Site Recovery Job Test Schedule

If you only want to start the site recovery job manually (without any test schedule), select the **Do not schedule, test on demand** check box.



Then click **Next** to go to the *Options* page of the *Site Recovery Job Wizard*.

Daily Site Recovery Job Testing

To test your site recovery job once a day, do the following:

- Choose a time zone that should be used for the site recovery job start and end times from the time zone list.
- Choose **Run daily/weekly** from the **Schedule #1** list.
- Specify the time when the site recovery job should be started in the **Starting at** box.
- Specify the end time for the site recovery job in the **Ending** box. If the site recovery job has not completed by the time specified, the site recovery job will be stopped.
- Select the days of the week during which the site recovery job will be started.
- To specify a date when the job test schedule comes into effect, click **Effective from,** click **date** and then pick a date in the calendar that opens.

- Click **Next** to go to the **Options** page of the **Site Recovery Job Wizard**.

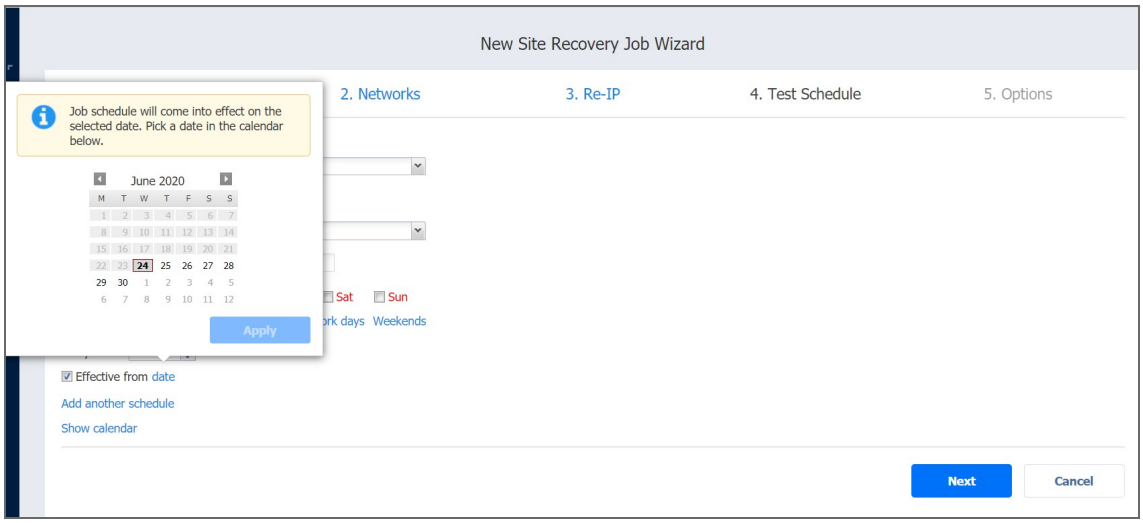

Monthly or Yearly Site Recovery Job Testing

To test your site recovery job monthly or yearly, do the following:

- Choose **Monthly/yearly** from the schedule list.
- Choose a time zone that should be used for the job start and end times, in the list of available time zones.
- Specify a site recovery job start schedule in the appropriate boxes.
- Specify the time when the site recovery job should be started, in the **Starting at** box.
- Specify the end time for the site recovery job, in the **Ending** box. If the site recovery job has not completed by the time specified, the site recovery job will be stopped.
- Select the days of the week during which the job will be started.
- To specify a date when the job test schedule comes into effect, click **Effective from,** click **Date** and then pick a date in the calendar that opens.
- Click **Next** to go to the **Options** page of the **Site Recovery Job Wizard**.

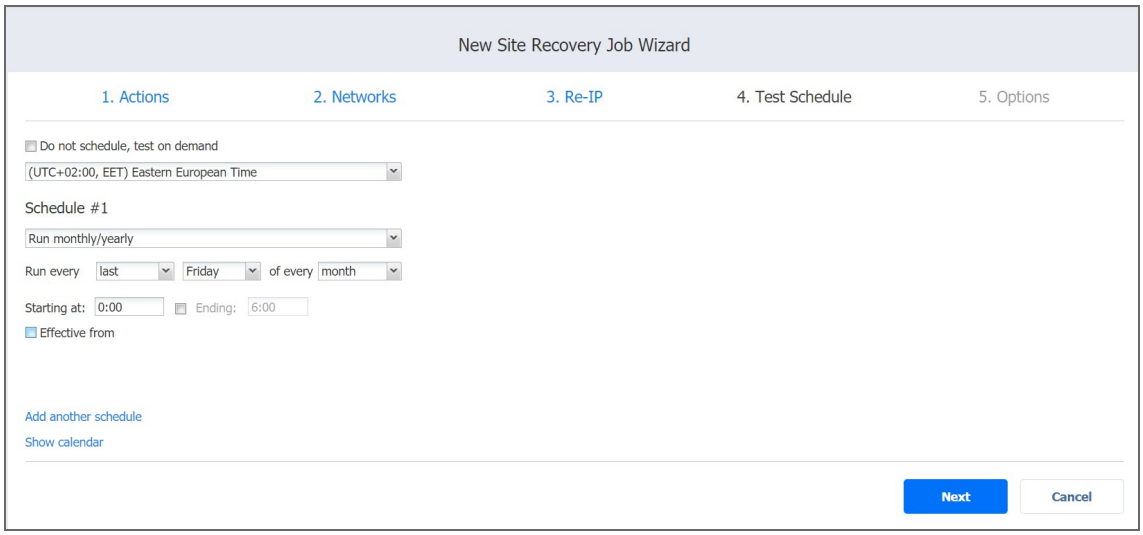

Periodic Site Recovery Job Testing

To test your site recovery job multiple times per day, do the following:

- Choose a time zone that should be used for the site recovery job start and end times from the list of time zones.
- Choose **Run periodically** from the **Schedule #1** list and then choose a time period from the appropriate boxes.
- Specify the time when the site recovery job should be started in the **Starting at** box.
- Specify the end time for the job in the **Ending** box. If the site recovery job has not completed by the time specified, the site recovery job will be stopped.
- To specify a date when the job test schedule comes into effect, click **Effective from,** click **Date** and then pick a date in the calendar that opens.
- Click **Next** to go to the **Options** page of the **Site Recovery Job Wizard**.

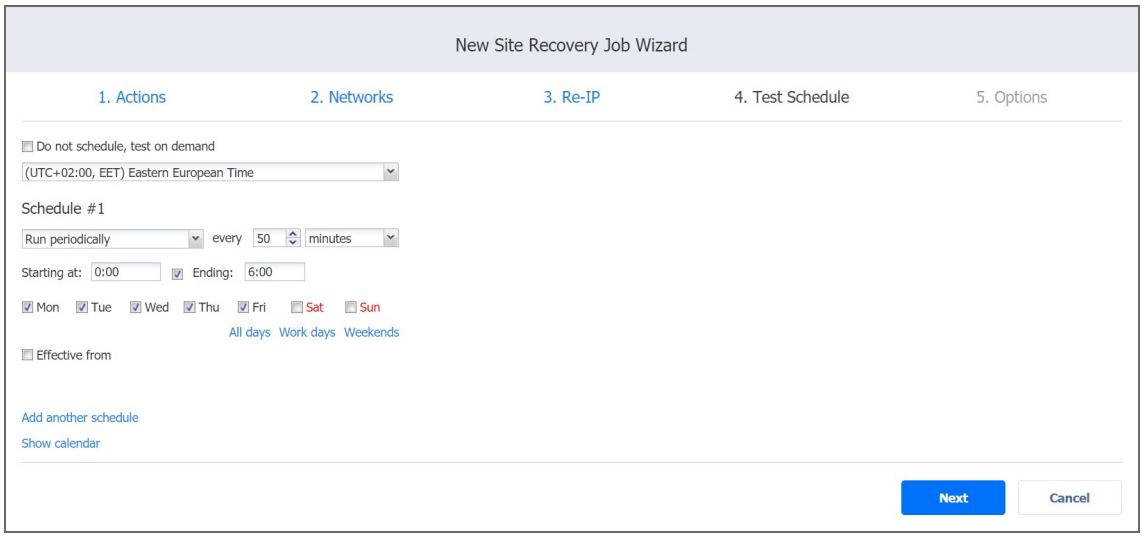

Chained Site Recovery Job

To run the site recovery job after a previous job has completed, do the following:

1. Choose **Run after another job** from the **Schedule #1** list.
2. Set the options as follows:
    - **After the job**: select a job after which the current site recovery job will be started.
    - **Run this job**: Choose whether to run the current site recovery job immediately after the previous job has completed, or specify a delay.
    - **After successful runs**: If selected, the site recovery job will run if the previous job has completed successfully.
    - **After failed runs**: If selected, the site recovery job will run if the previous job has failed.
    - **After stopped runs**: If selected, the site recovery job will run if the previous job has been stopped.
3. To specify a date when the job test schedule comes into effect, click **Effective from**, click **date** and then pick a date in the calendar that opens.

4. Click **Next** to go to the *Options* page of the **Site Recovery Job Wizard**.

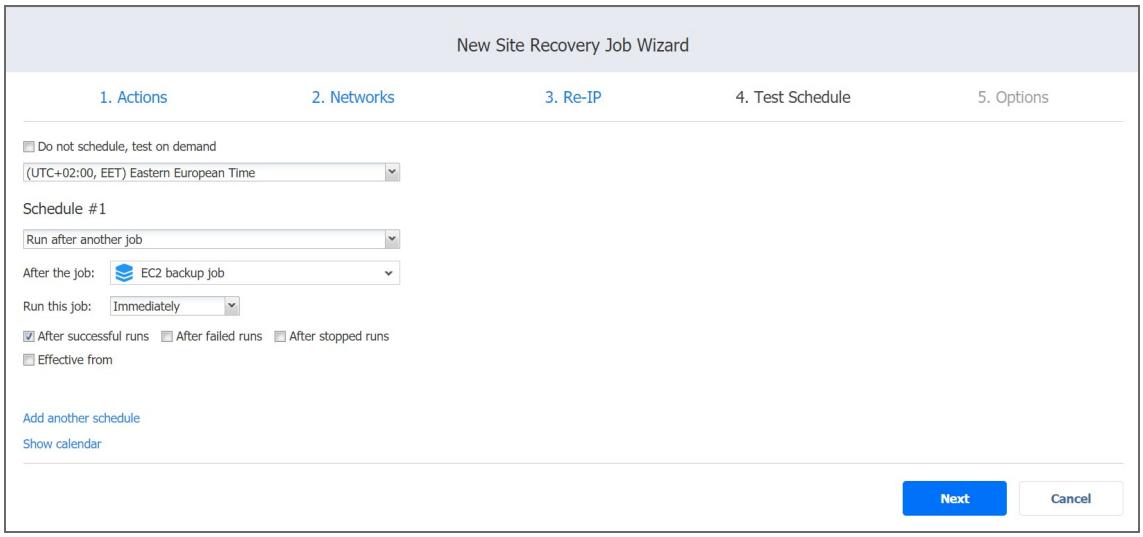

Additional Schedule

To add more than one schedule to your site recovery job, do the following:

1. Click **Add another schedule**.
2. The *Schedule #2* section opens. Proceed with instructions provided in the sections above.
3. When ready with adding an additional schedule to your site recovery job, click **Next**.

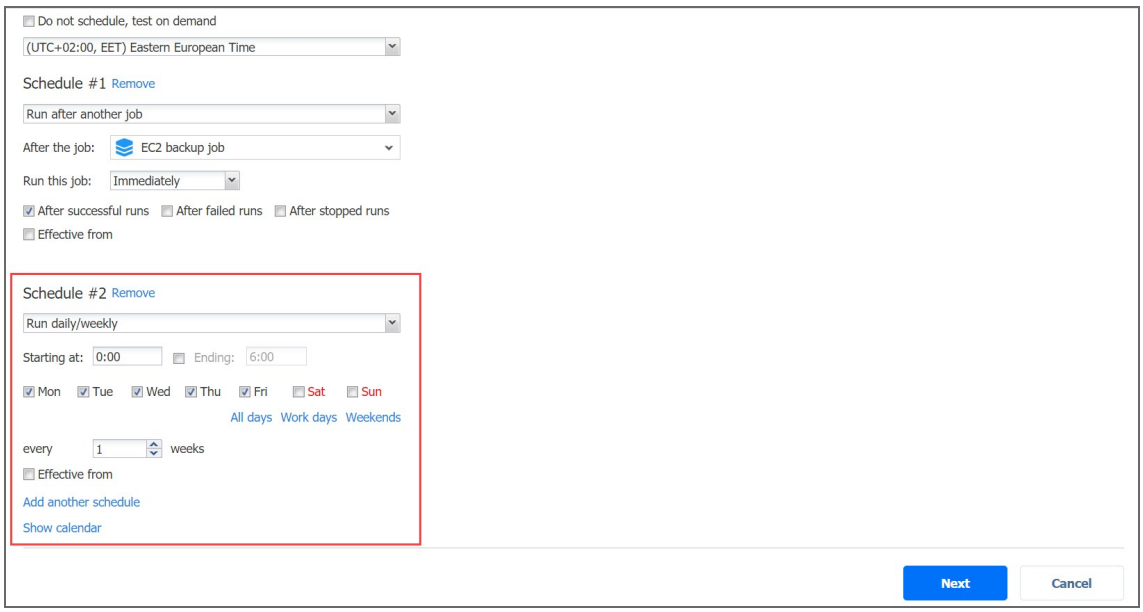

Site Recovery Job Wizard: Options

On the **Options** page of the **Site Recovery Job Wizard** you can specify 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.



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**.

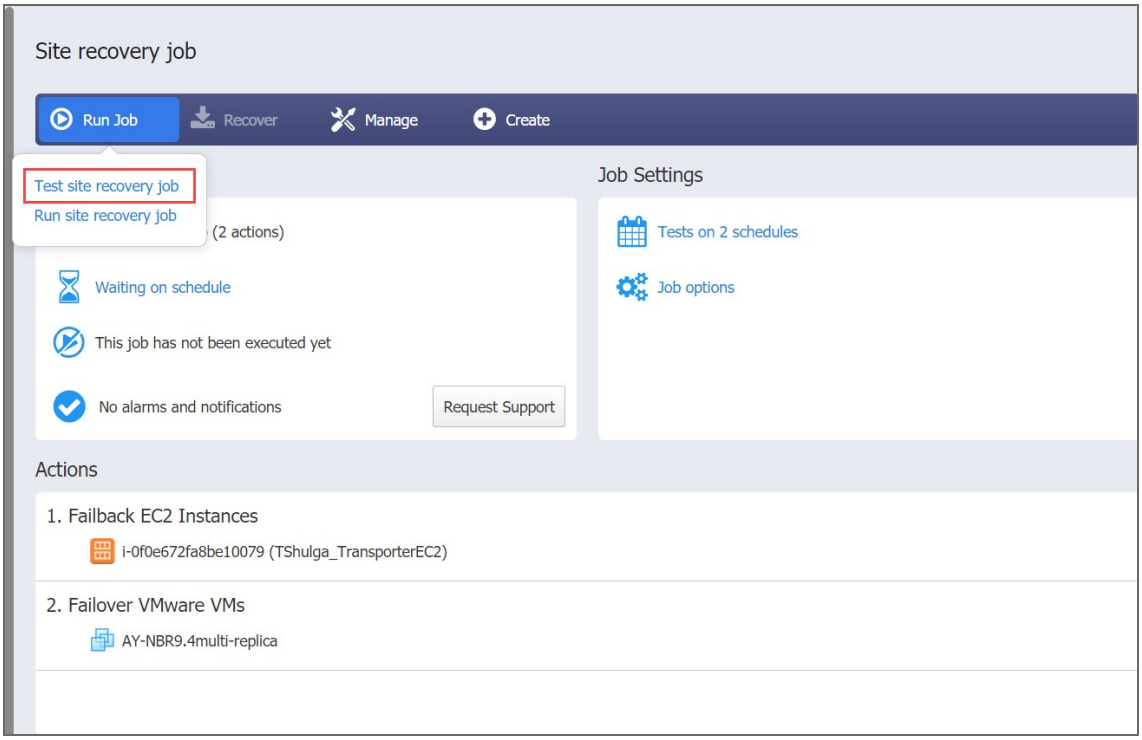

3. The **Recovery time objective** dialog opens. Here you can:

- 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.

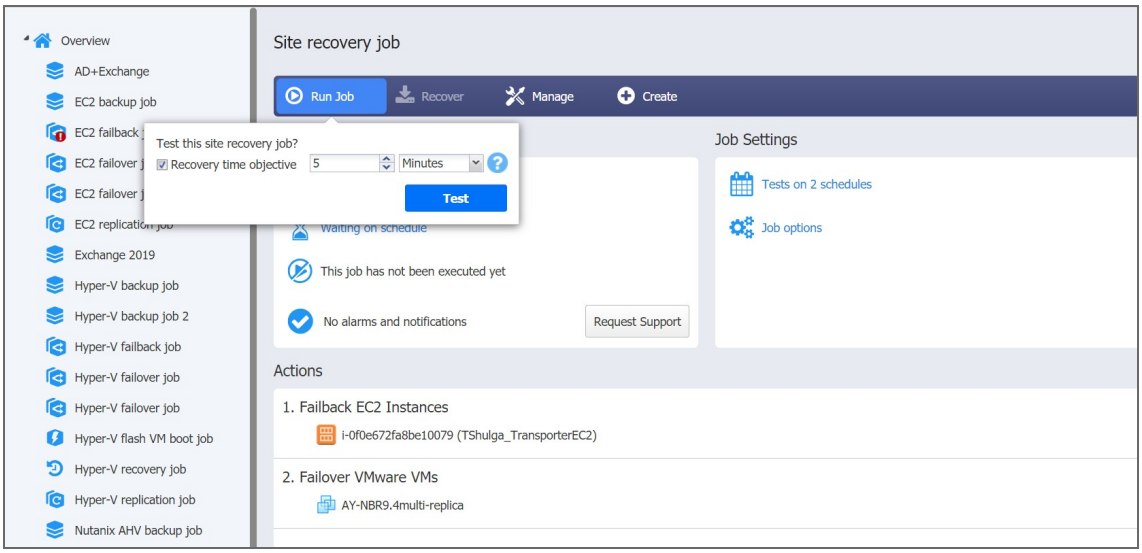

4. Click **Test** when ready. The Site Recovery Job starts running in the test mode.

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**.

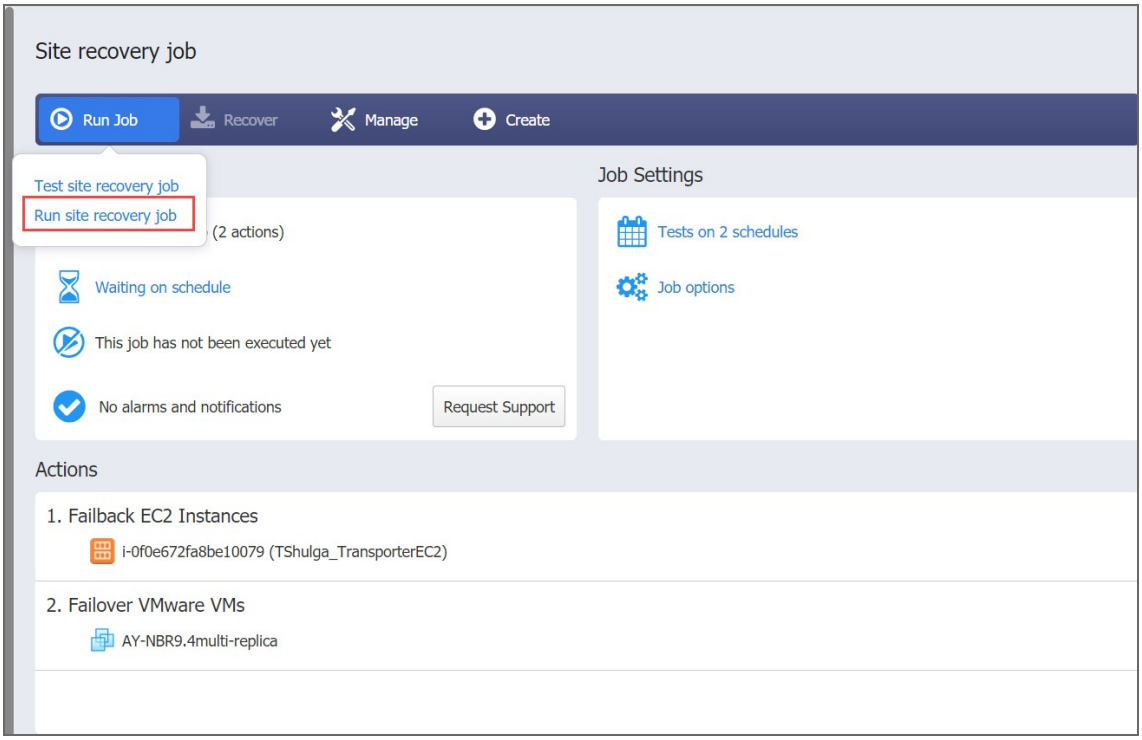

3. The **Failover type** dialog opens. Choose either of the following failover types:
   - **Planned failover**: The application will sync replica data with the source VM before switching workloads to the replica.
   - **Emergency failover**: The application will switch workloads from the source VM to the replica immediately.

   Note

   The **Failover type** option is only available for Site Recovery Jobs containing a Failover action.

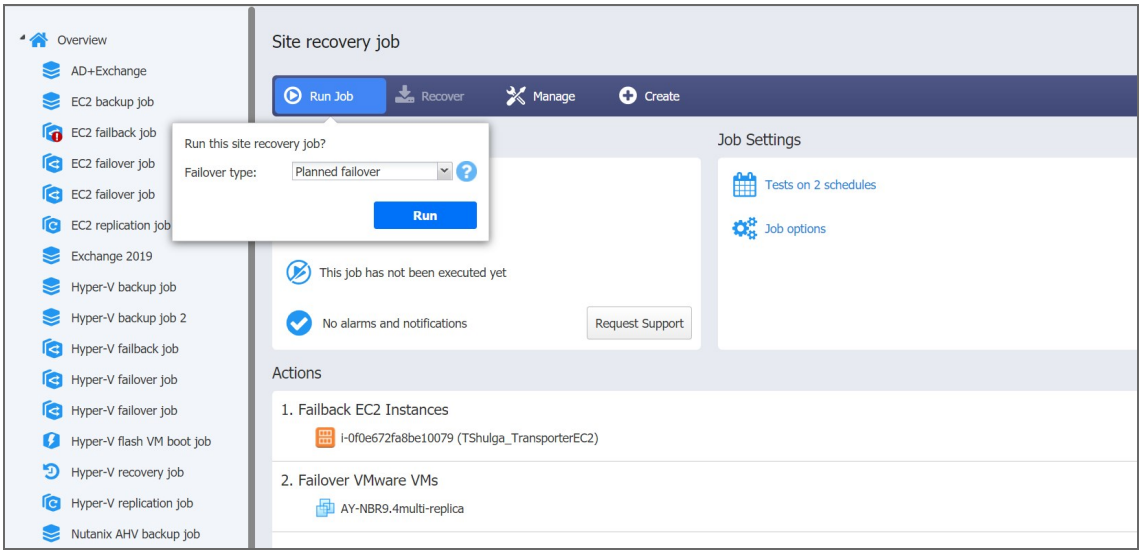

4. Click **Run**. The Site Recovery Job starts running in the production mode.

Integration and Automation

This section contains the following topics:

- "Command Line Interface" on page 751
- "Automation with HTTP API" on page 762
- "Aptare IT Analytics Integration" on page 763

Command Line Interface

This section covers the following topics:

- "Using Command Line Interface" on page 752
- "Available Commands" on page 754
- "Exit Codes" on page 761

Using Command Line Interface

- "Operation Modes of Command Line Interface" below
- "Using Command Line Interface Locally" below
- "Using Command Line Interface Remotely" below
- "Using Command Line Interface in Multi-Tenant Mode" on the next page

NAKIVO Backup & Replication allows you running actions from the product's command line interface (CLI). In case credentials are configured for the product, running an action via CLI requires providing administrator credentials as arguments, namely, `--username [login] --password [password]`, where `[login]` is the administrator user name and `[password]` is the administrator password.

Operation Modes of Command Line Interface

You can run CLI in either of the following modes:

- *Interactive mode.* This allows you to use a single login for a session. When opened in the interactive mode, CLI allows you executing commands without dashes.
  To open CLI in the interactive mode, enter `cli.bat --interactive --username [login] --password [password]` and press **Enter**. To exit the CLI interactive mode, enter **Ctrl-C**.
- *Non-interactive mode.* This requires entering your credentials for each command. You will have to enter dashes before commands. For example: `cli.bat --username [login] --password [password] --inventory-list`

Using Command Line Interface Locally

To use CLI on the machine where NAKIVO Backup & Replication Director is installed, follow the steps below:

1. Run the CLI executable:
   - If NAKIVO Backup & Replication is installed on a Windows OS, run the `cli.bat` file located in the `bin` folder inside the product installation folder ("`C:\Program Files\NAKIVO Backup & Replication`" by default).
   - If NAKIVO Backup & Replication is installed on a Linux OS, run the `cli.sh` file located in the `director/bin` folder inside the product installation folder (`/opt/nakivo/` by default).
2. Run available commands.

Using Command Line Interface Remotely

To use CLI from a remote machine, follow the steps below:

1. Copy the CLI executable and `.jar` files to the machine from where you plan to use the CLI:
    - If NAKIVO Backup & Replication is installed on a Windows OS, copy the `cli.bat` and `cli.jar` files located in the `bin` folder inside the product installation folder ("`C:\Program Files\NAKIVO Backup & Replication`" by default).
    - If NAKIVO Backup & Replication is installed on a Linux OS, copy the `cli.sh` and `cli.jar` files located in the `director/bin` folder inside the product installation folder (`/opt/nakivo/` by default).
2. On the machine from where you plan to use the CLI, configure the `PATH` system variable as described at http://java.com/en/download/help/path.xml
3. Run commands using the following format: `<command> <host> <port> <username> <password>`

Example

To get a list of jobs of the product which is installed on the machine with the `192.168.10.10` IP address, uses the `4443` port number for the **Director Web HTTPS port**, and has "`admin`" as login and password for the product's web UI, run the following command: `--job-list --host 192.168.10.10 --port 4443 --username admin --password admin`

Using Command Line Interface in Multi-Tenant Mode

Triggering an action inside a tenant in the multi-tenant mode via command line interface requires providing a tenant ID as an argument:

```
cli.bat --repository-detach [repo_id] --username [login] --password [password] --tenant [tenant-id]
```

Available Commands

You can run CLI commands either in interactive or non-interactive mode. Refer to the *Operation Modes of Command Line Interface* subsection of the "Using Command Line Interface" on page 752 topic.

Use either long or short form of the commands*.

Command	Long form	Short form	Output
Help			
General help	cli.bat --help	cli.bat -h	• Command name • Description
Job Management			
List all jobs	cli.bat --job-list	cli.bat -jl	• Job ID • Job name • Current job status • Job last run result
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] • The command with no arguments creates the job report and saves it to the current directory.	cli.bat -jp [job_id] • The command with no arguments creates the job report and saves it to the current directory.	

Command	Long form	Short form	Output
	• To save the report to other directory: cli.bat --job-report [job_id] --save-to [dir_path] • To send the report to default email(s): cli.bat. --job-report [job_id] --send-by-email • To send the report to other email: cli.bat. --job-report [job_id] --send-by-email [email_ address]	• To save the report to other directory: cli.bat -jp [job_id] -f [dir_path] • To send the report to default email(s): cli.bat. -jp [job_id] -eml • To send the report to other email: cli.bat. -jp [job_id] -eml [email_ address]	
Return information about a job	cli.bat --job-info [job_id]	cli.bat -ji [job_id]	• Job ID • Job name • Current job status • Job last run result
Inventory			
List all inventory items	cli.bat --inventory-list	cli.bat -il	• Item ID • Item IP/host name • Item type (host/vCenter) • Item children count (X hosts, E VMs) • Item current state • Item current status
Update all inventory items	cli.bat --inventory-update	cli.bat -iu	
Update an inventory item	cli.bat --inventory-update [item_id]	cli.bat -iu [item_id]	
Return information about an inventory item	cli.bat --inventory-info [item_id]	cli.bat -ii [item_id]	• Item ID • Item IP/host name

Command	Long form	Short form	Output
			• Item type (host/vCenter) • Item children count (X hosts, E VMs) • Item current state • Item current status
Transporters			
List all transporters	cli.bat --transporter-list	cli.bat -trl	• Transporter ID • Transporter IP/host name • Transporter current load • Transporter maximum load • Transporter current state • Transporter current status
Update all transporters	cli.bat --transporter-update	cli.bat -tru	
Update a transporter	cli.bat --transporter-update [transporter_id]	cli.bat -tru [transporter_id]	
Return information about a transporter	cli.bat --transporter-info [transporter_id]	cli.bat -tri [transporter_id]	• Transporter ID • Transporter IP/host name • Transporter current load • Transporter maximum load • Transporter current state • Transporter current status

Command	Long form	Short form	Output
Repositories			
List all repositories	cli.bat --repository-list	cli.bat -rl	• Repository ID • Repository name • Assigned transporter • Backup count • Free space • Attached or detached • Consistent or inconsistent • Repository current state • Repository current status
Update all repositories	cli.bat --repository-update	cli.bat -ru	
Update a repository	cli.bat --repository-update [repo_id]	cli.bat -ru [repo_id]	
Detach a repository	cli.bat --repository-detach [repo_id]	cli.bat -rd [repo_id]	
Attach a repository	cli.bat --repository-attach [repo_id]	cli.bat -ra [repo_id]	
Start repository maintenance	cli.bat --repository-maintenance [repo_id] [parameter] Parameters: • --selfheal • --verify • --spacereclaim	cli.bat -rm [repo_id] [parameter] Parameters: • --selfheal • --verify • --spacereclaim	
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]	• Repository ID • Repository name

Command	Long form	Short form	Output
			• Assigned transporter • Backup count and free space • Attached or detached • Consistent or inconsistent • Repository current state • Repository current status
Support			
Generate a support bundle	cli.bat --bundle-create • The command with no parameters will create a support bundle and save it in the current directory. • To save the bundle to other directory: cli.bat --bundle-create --save-to [dir_path] • To send the bundle to support over email: cli.bat --bundle-create --send-to-support • To send the bundle to other email: cli.bat --bundle-create --send-by-email [email_ address]	cli.bat -bc • The command with no parameters will create a support bundle and save it in the current directory. • To save the bundle to other directory: cli.bat -bc -f [dir_path] • To send the bundle to support over email: cli.bat -bc -sup • To send the bundle to other email: cli.bat -bc -eml [email_ address]	
Licensing			
Get the current license information	cli.bat --license-info	cli.bat -li	

Command	Long form	Short form	Output
Replace the current license with a new license file	cli.bat --license-replace [file_path]	cli.bat -lin [file_path]	
Multi Tenancy			
List all tenants	cli.bat --tenant-list	cli.bat -tl	• Tenant ID • Tenant name • Allocated items type and count • Tenant status • Enabled or disabled
Disable a tenant	cli.bat --tenant-disable [tenant_id]	cli.bat -td [tenant_id]	
Enable a tenant	cli.bat --tenant-enable [tenant_id]	cli.bat -te [tenant_id]	
Return information about a tenant	cli.bat --tenant-info [tenant_id]	cli.bat -ti [tenant_id]	• Tenant ID • Tenant Account ID • Tenant name • Allocated items type and count • Tenant status • Enabled or disabled
Create a support bundle for master admin level	• Generate the support bundle for master level only: cli.bat --bundle-create • Generate the support bundle with all tenants logs: cli.bat --bundle-create --include-tenants	• Generate the support bundle for master level only: cli.bat -bc • Generate the support bundle with all tenants logs: cli.bat -bc -ite	
Miscellaneous			
Get the CLI version	cli.bat --version	-	

Command	Long form	Short form	Output
	The command returns the CLI version which is equal to the full version of NAKIVO Backup & Replication.		
Run a command in the debug mode	cli.bat --repository-info [repo_id] --debug This is an option that can be added to any other CLI command. With the debug mode turned on, the commands will return the full error text.	cli.bat -ri [repo_id] --debug	

*Examples are given for Windows OS.

Exit Codes

NAKIVO Backup & Replication CLI provides the following exit codes:

- **0:** Normal
- **1:** Unknown command
- **2:** Cannot login
- **3:** Command failed
- **4:** Local failure
- **5:** No arguments

Automation with HTTP API

HTTP API allows you to run common NAKIVO Backup & Replication commands outside of the product web interface.

The API is JSON-RPC based. For detailed request and response syntax, refer to API Reference.

Aptare IT Analytics Integration

APTARE IT Analytics is a storage resource management platform for integrating storage and backup solutions. The integration with NAKIVO Backup & Replication is based on an APTARE data collector that sends storage component information to the system's platform. The steps for integrating NAKIVO Backup & Replication with APTARE IT Analytics are as follows:

1. On the machine where NAKIVO Backup & Replication is deployed, do the following:
   a. Install APTARE StorageConsole Data Collector with NAKIVO connector.
   b. When the installation has been successfully completed, make sure that the APTARE Agent service is running.
2. Open your NAKIVO Backup & Replication instance and run your backup jobs.
3. Log in to the APTARE portal.
4. Go to the **ADMIN** tab and take the following steps:

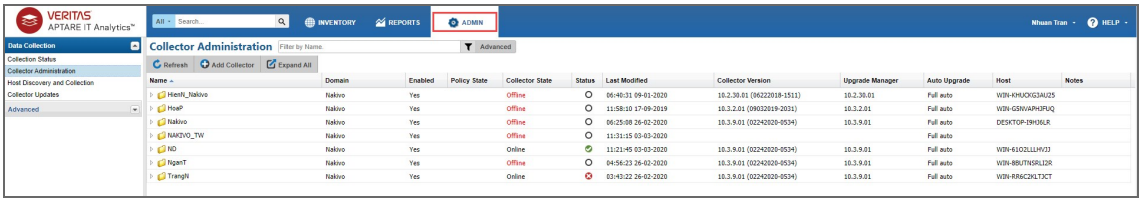

   a. Add a Collector. For details, refer to the Managing and Monitoring Data Collection subsection of the APTARE IT Analytics User Guide.

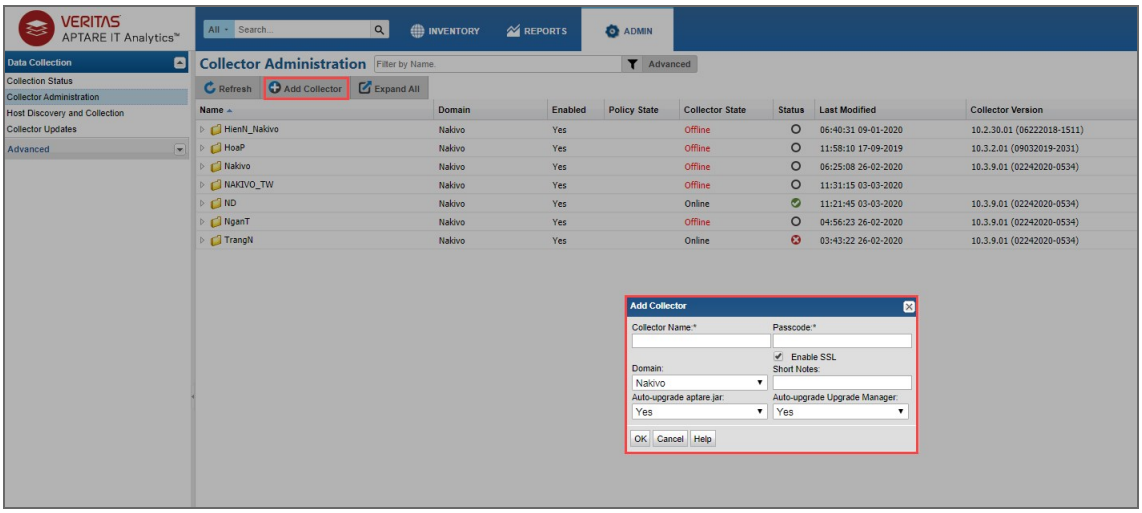

b. Add a NAKIVO Backup & Replication data protection policy with a connection to your NAKIVO Backup & Replication instance. For details, refer to the Pre-Installation Setup for Generic Backup subsection of the APTARE IT Analytics User Guide.

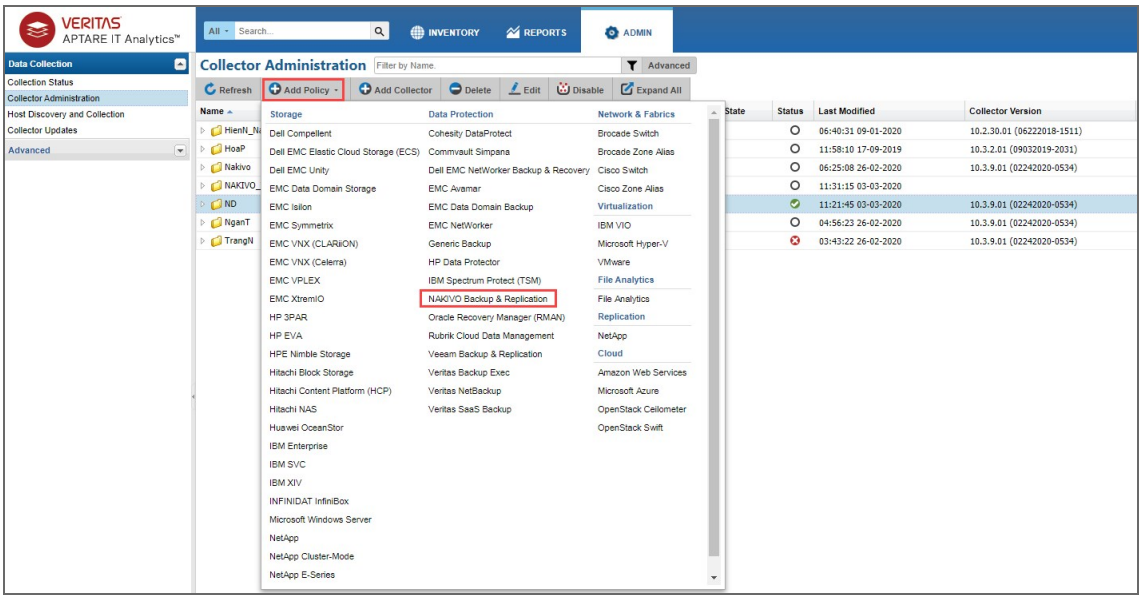

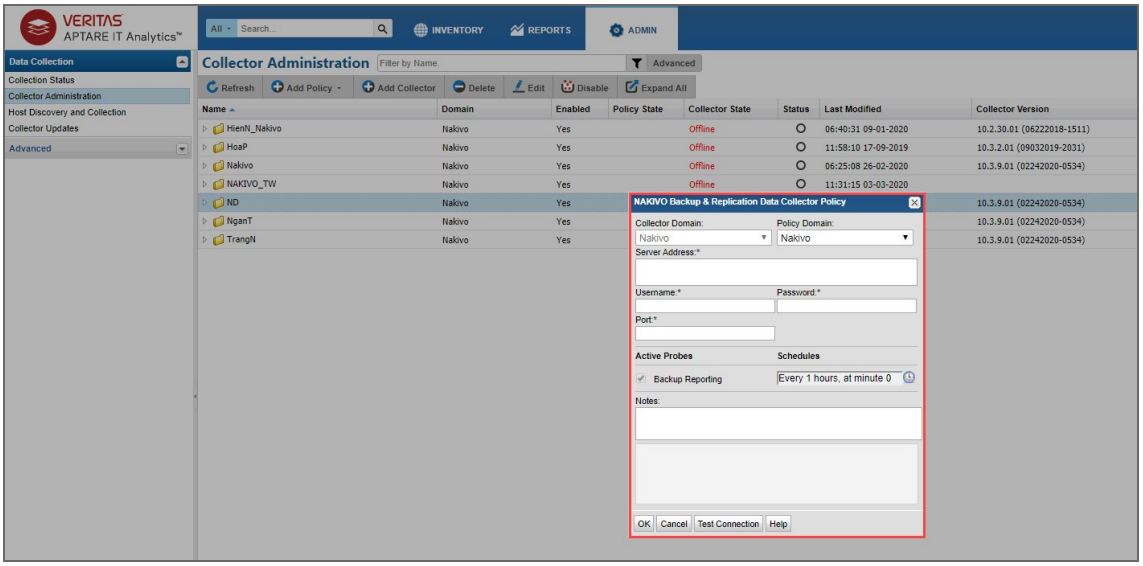

c. Run your policy.

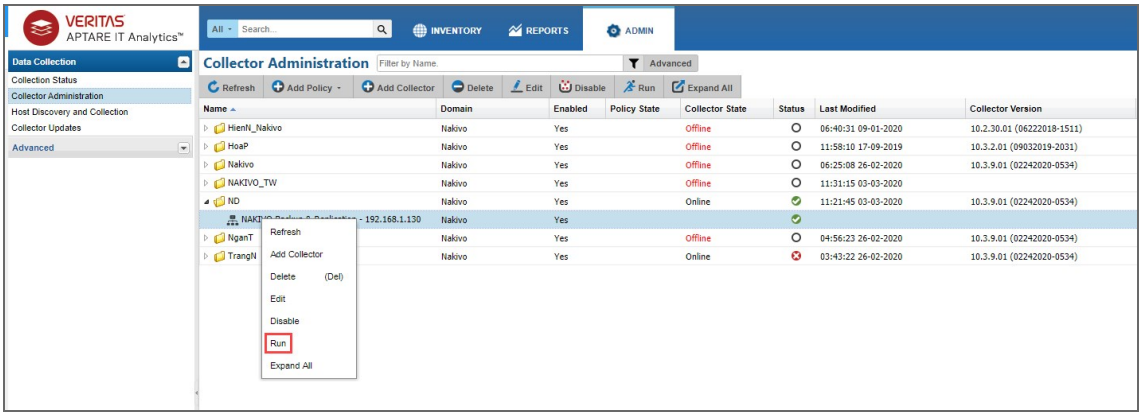

5. Go to the **REPORTS** tab in the APTARE portal and take the following actions:

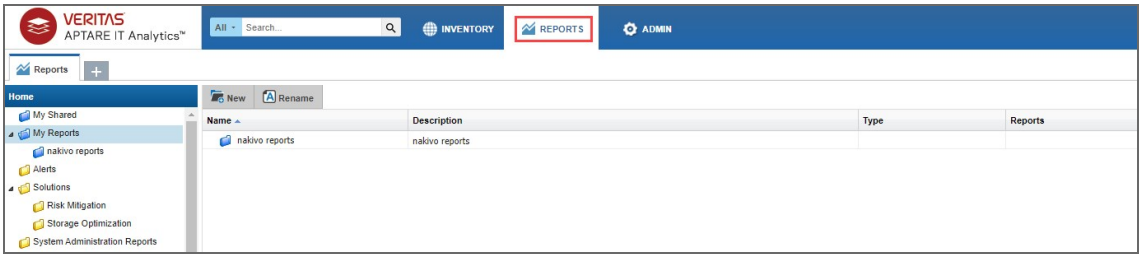

a. Create and configure the report for your backup job the following way:

i. Right-click on your report folder and select **New SQL Template**.

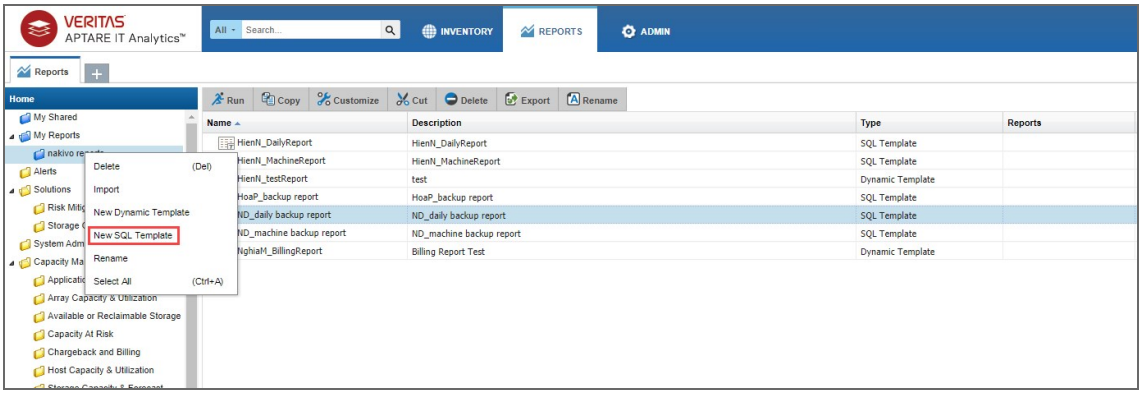

ii. Select the template designer that will be used to gather user input for the report.

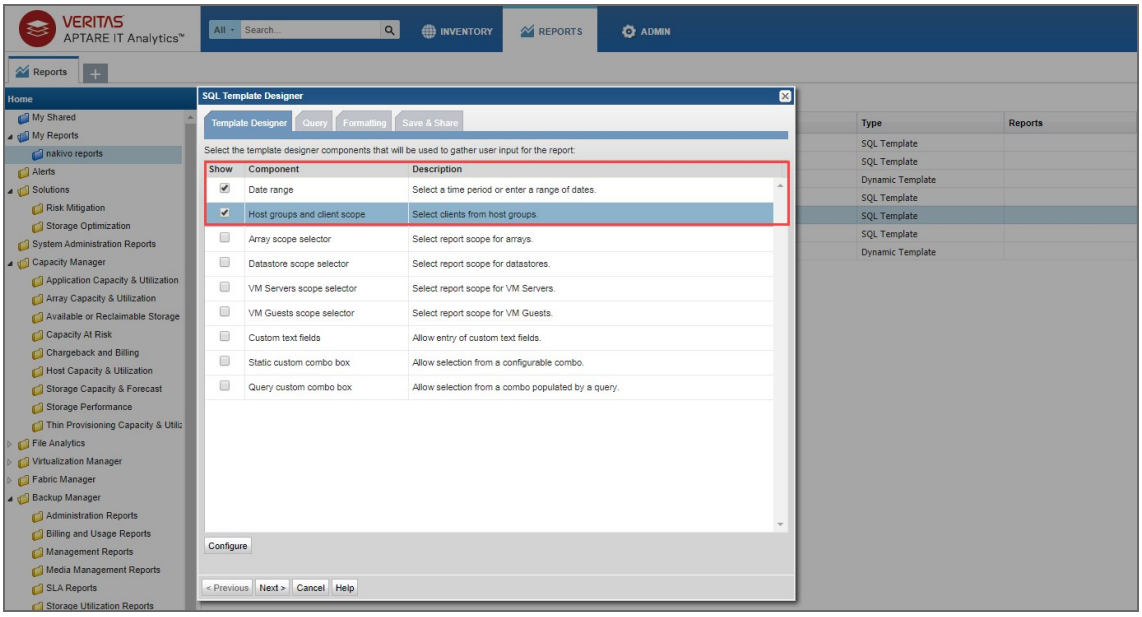

iii. Build an SQL query for your machine backup report or daily backup report, use either:

Daily Backup Report

```
select * from SDK_V_NKVO_NBR_DAILYBACKUP where (creation_date
between ${startDate} and ${endDate}) and (server_instance_id in
(select server_instance_id from apt_v_server_instance where server_
id in (${hosts}))) order by report_id desc
```

Machine Backup Report

```
select * from SDK_V_NKVO_NBR_MACHINEBACKUP where (creation_
date between ${startDate} and ${endDate}) and (server_
instance_id in (select server_instance_id from apt_v_server_
instance where server_id in (${hosts}))) order by report_id
desc
```

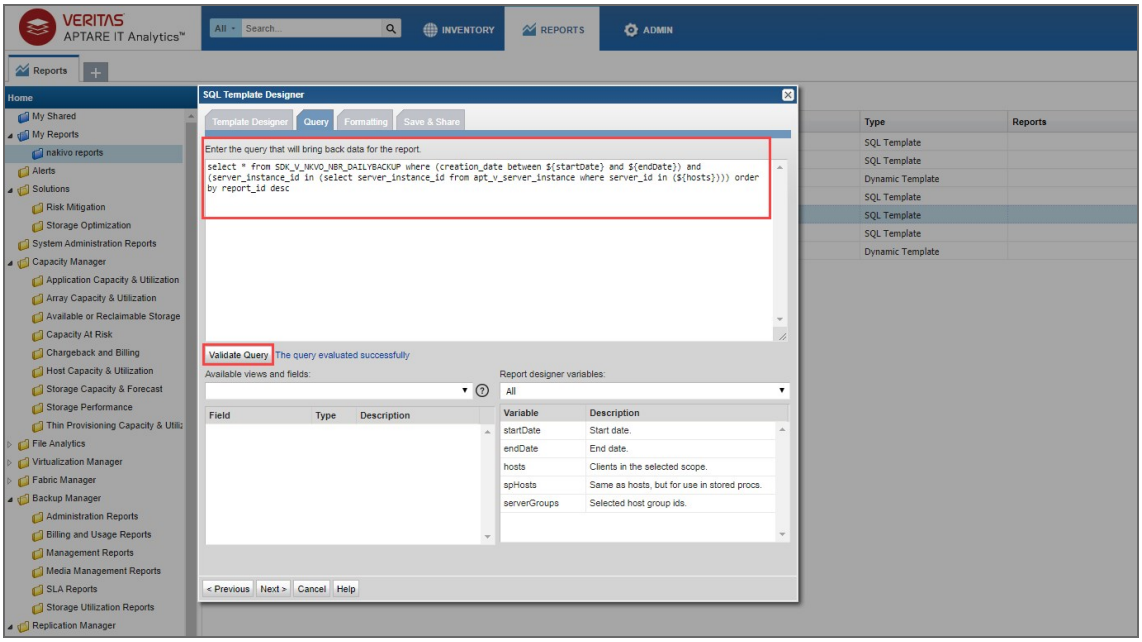

iv. Change formatting options if necessary.

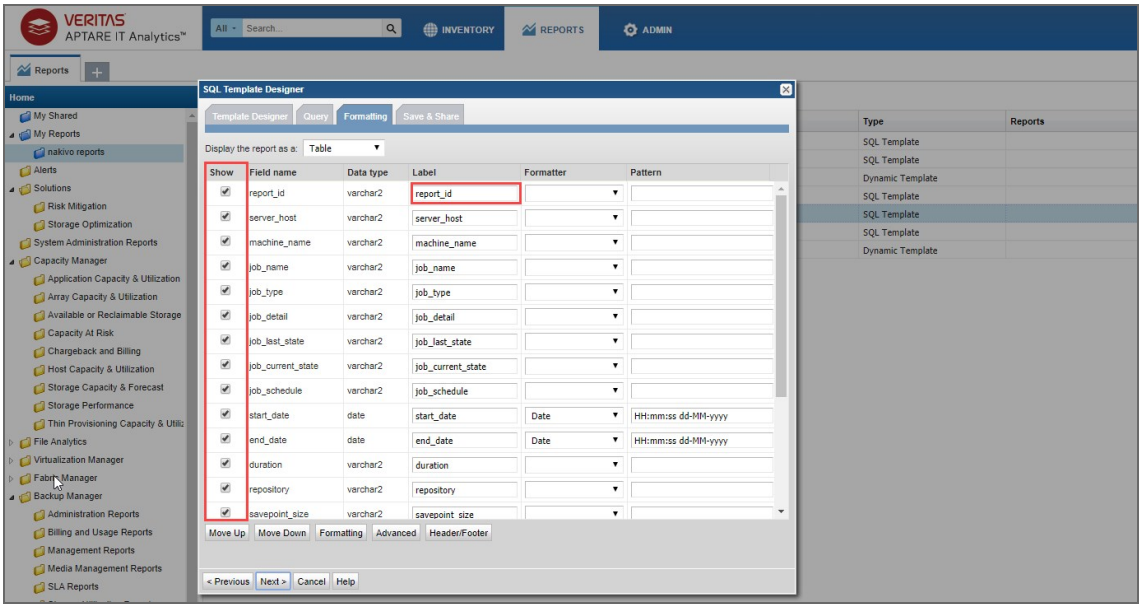

v. Provide a name and description of the report and select users to share it with. Click **Finish**.

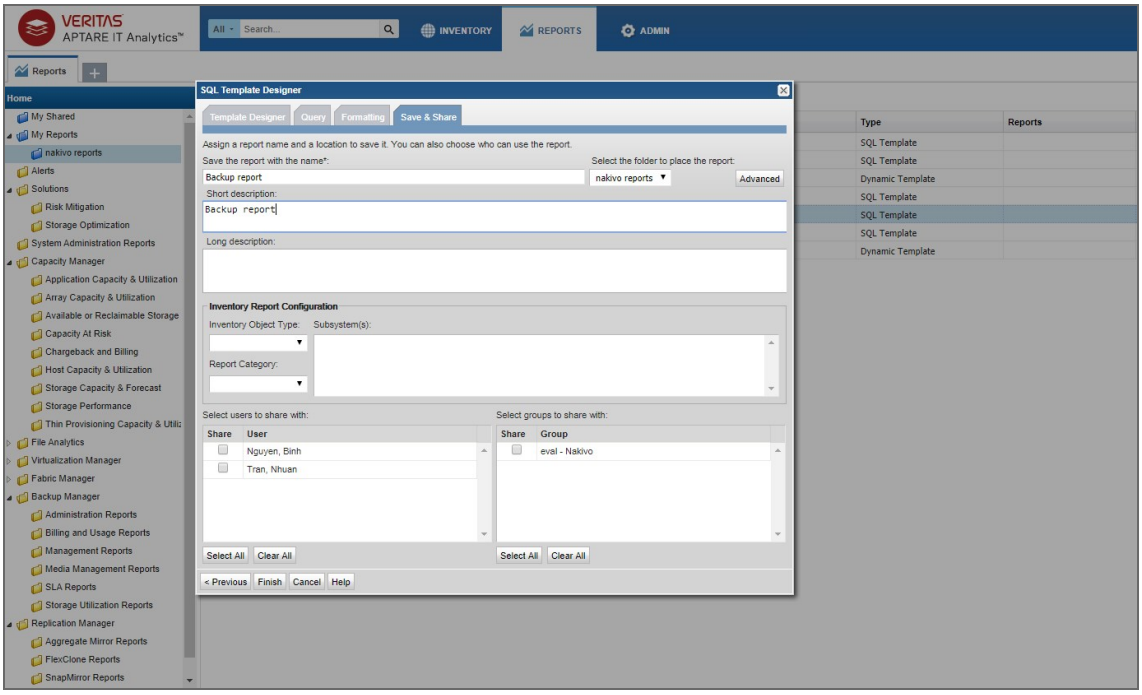

b. Double-click on the report. In the dialog box that opens, enter the necessary time and report scope. Click **Generate** to generate your report. For details, refer to the Generating and Maintaining Reports subsection of the APTARE IT Analytics User Guide.

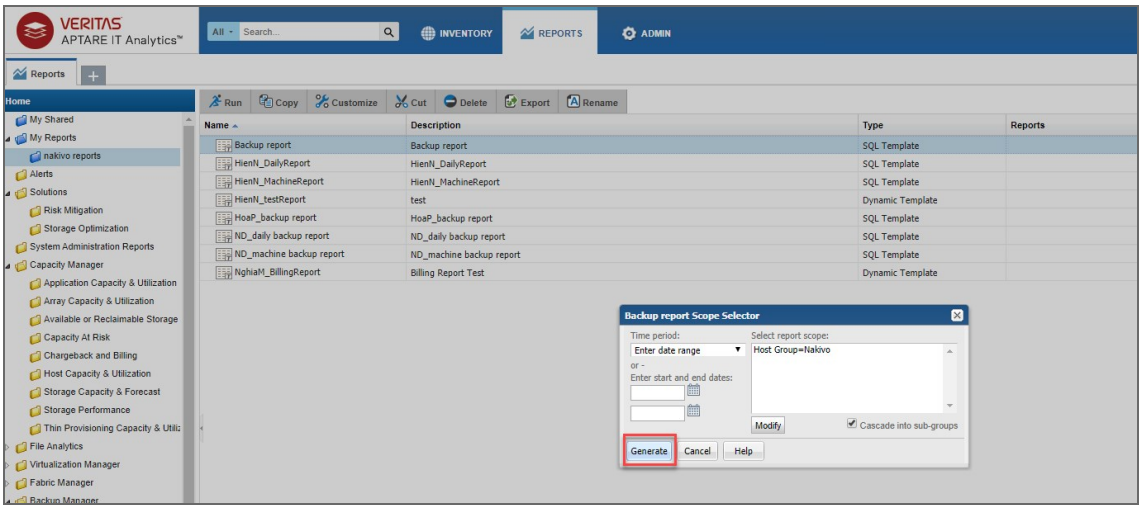

To know more about APTARE IT Analytics, refer to the APTARE IT Analytics User Guide.

Multi-Tenant Mode

This section covers the following topics:

- "Tenant Creation" on page 770
- "Tenant Configuration" on page 777
- "Tenant Management" on page 778
- "Granting Self-Service Access" on page 789

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**.

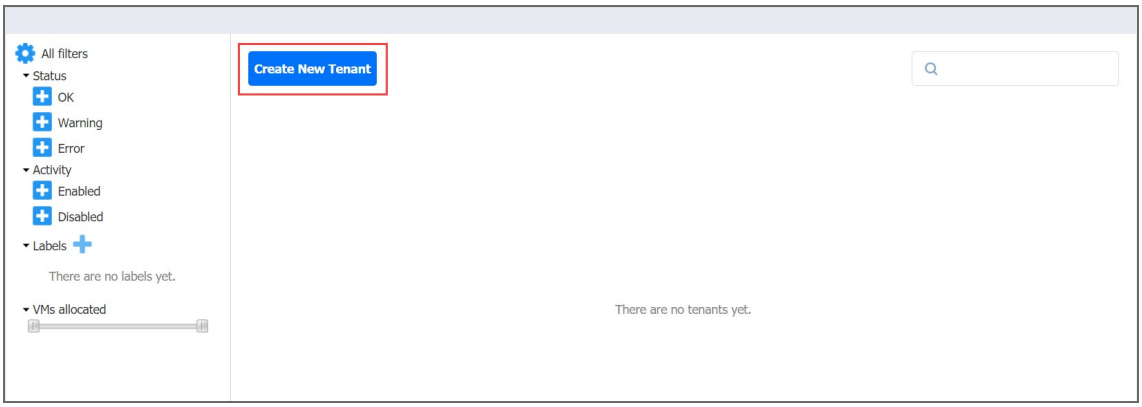

3. Complete the wizard as described in the topics below to finish the tenant creation process:

- “Tenant Creation Wizard: Tenant” on page 771
- “Tenant Creation Wizard: Inventory” on page 773
- “Tenant Creation Wizard: Transporters” on page 774
- “Tenant Creation Wizard: Repositories” on page 775
- “Tenant Creation Wizard: Users” on page 776

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.

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.
    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.

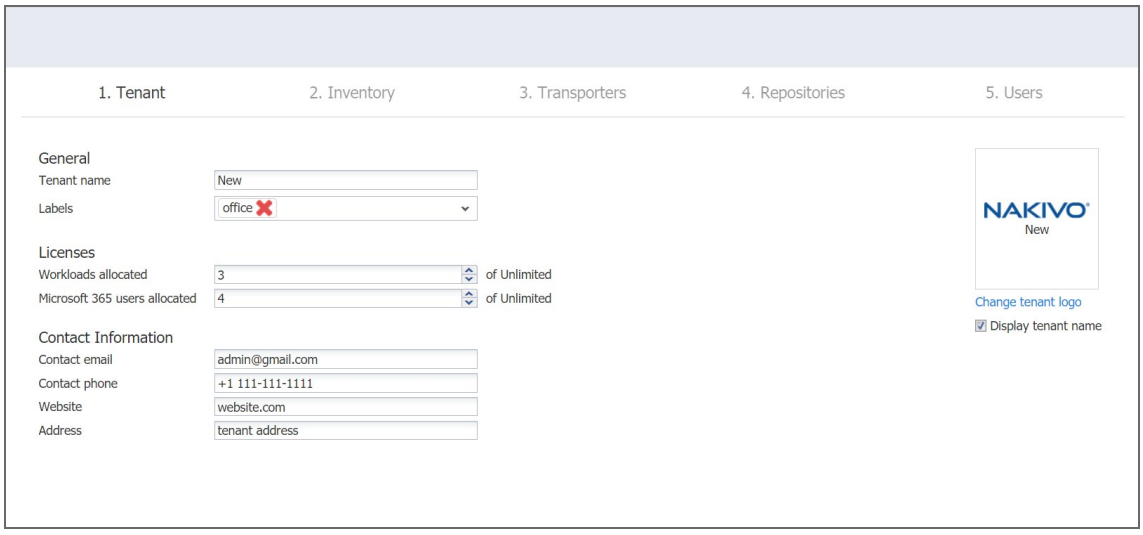
1. Tenant
2. Inventory
3. Transporters
4. Repositories
5. Users
General
Tenant name
New
Labels
office
Licenses
Workloads allocated
3
of Unlimited
Microsoft 365 users allocated
4
of Unlimited
Contact Information
Contact email
admin@gmail.com
Contact phone
+1 111-111-1111
Website
website.com
Address
tenant address
NAKIVO
New
Change tenant logo
Display tenant name

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.

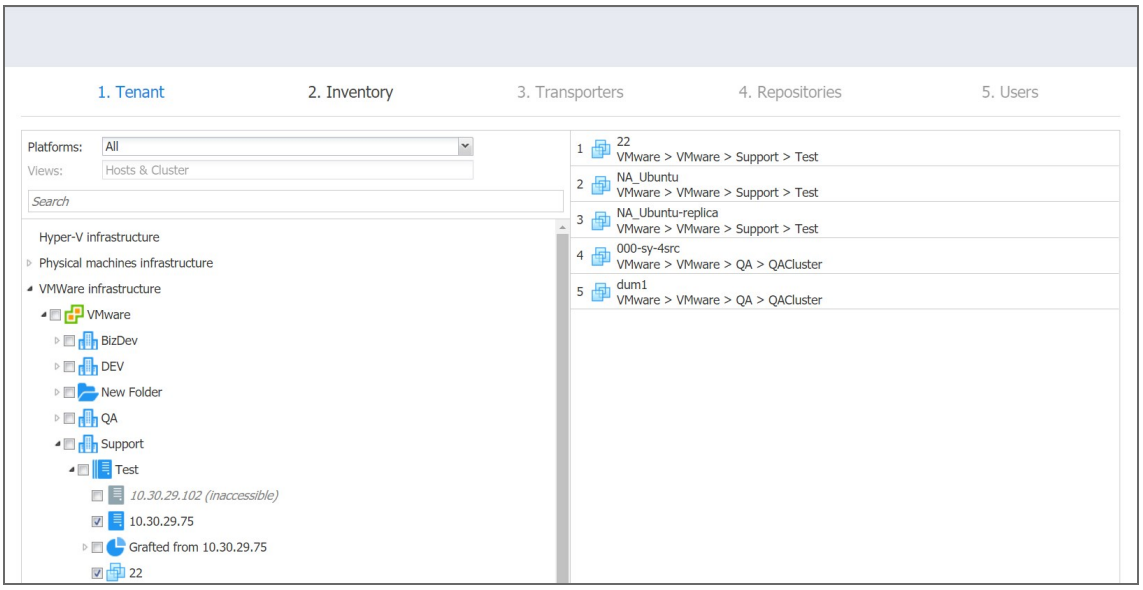

4. Click **Next** to proceed to the **Transporters** page.

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.

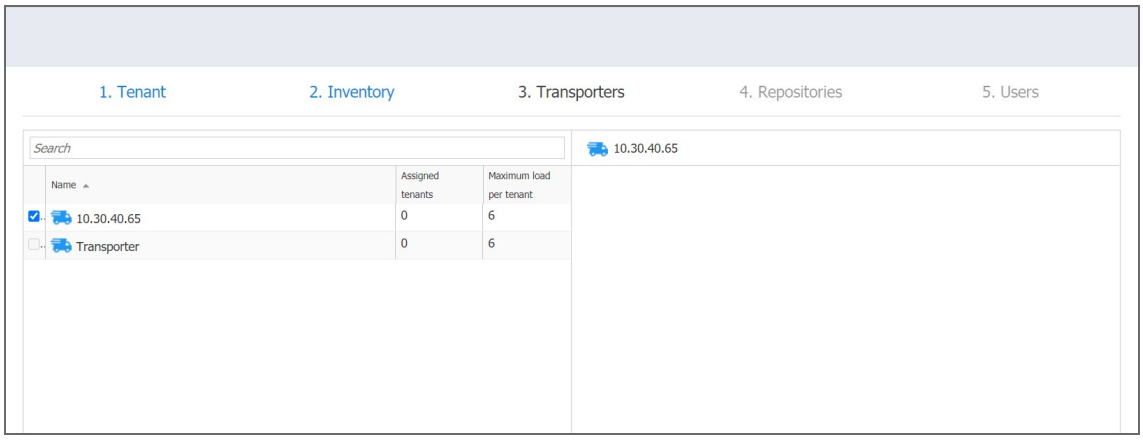

Tenant Creation Wizard: Repositories

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.

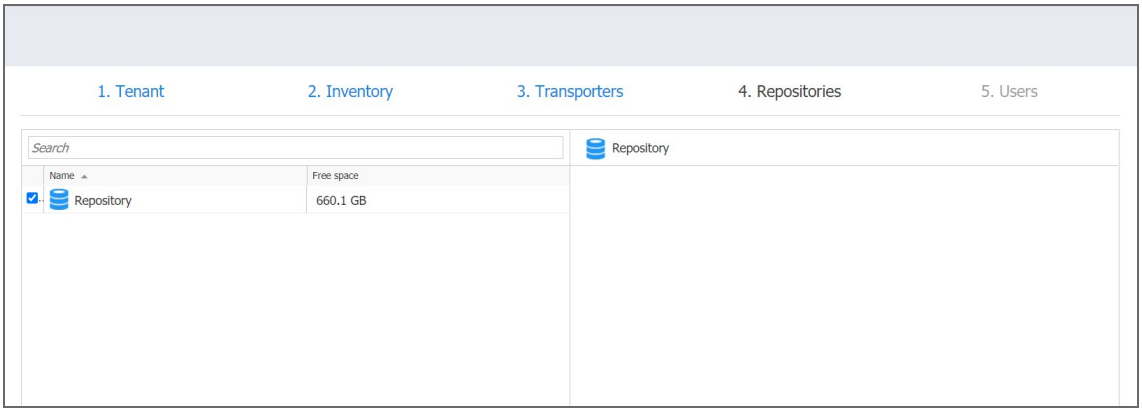

3. Click **Next** to proceed to the next page of the wizard.

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.

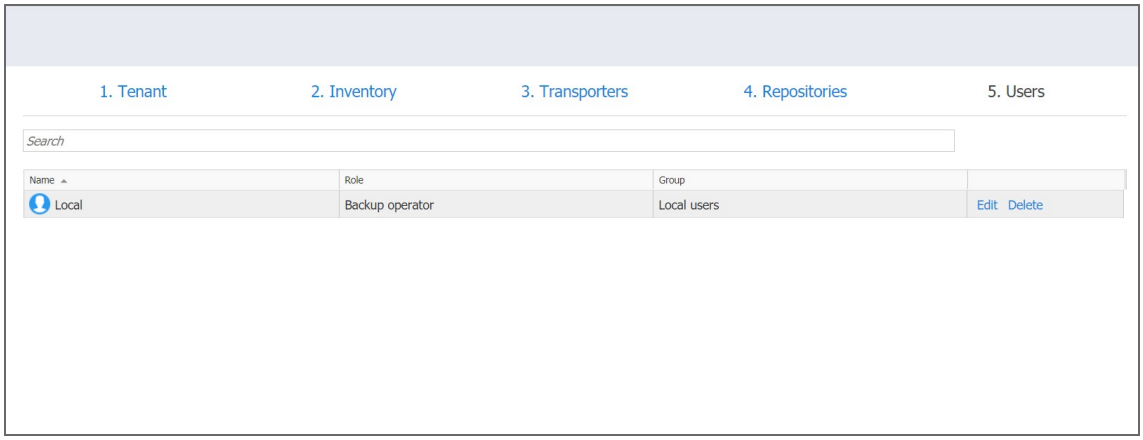

Tenant Configuration

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 241 for a description of the initial configuration wizard.

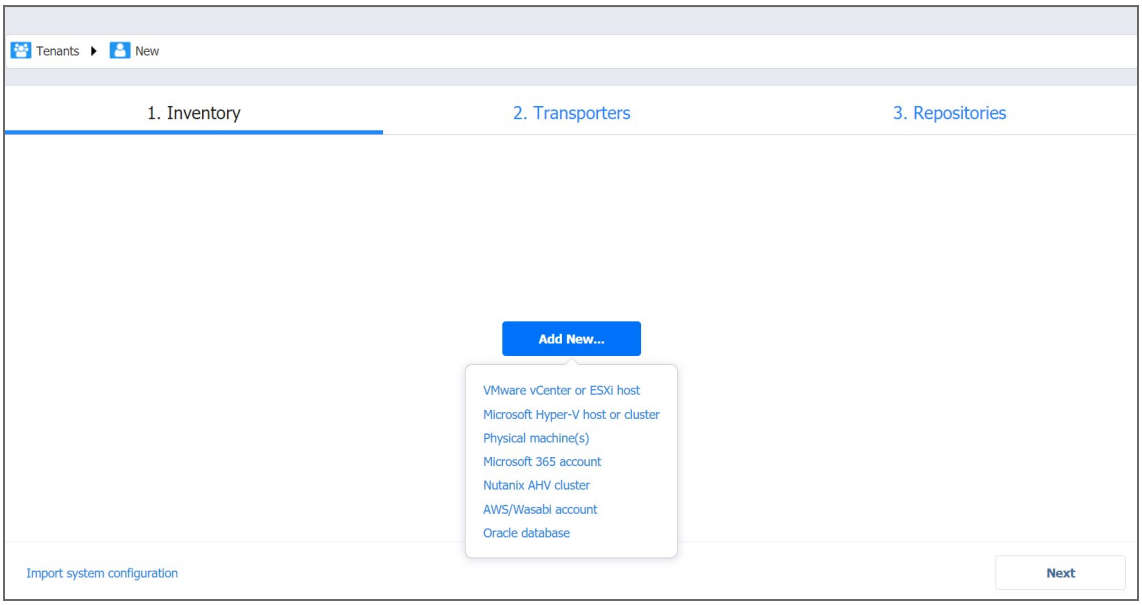

Tenant Management

This section covers the following topics:

- "Using Filters" on page 779
- "Using Labels" on page 781
- "Viewing Tenant Information" on page 784
- "Opening Tenant Dashboard" on page 785
- "Disabling Tenants" on page 786
- "Editing Tenants" on page 787
- "Deleting Tenants" on page 788

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.

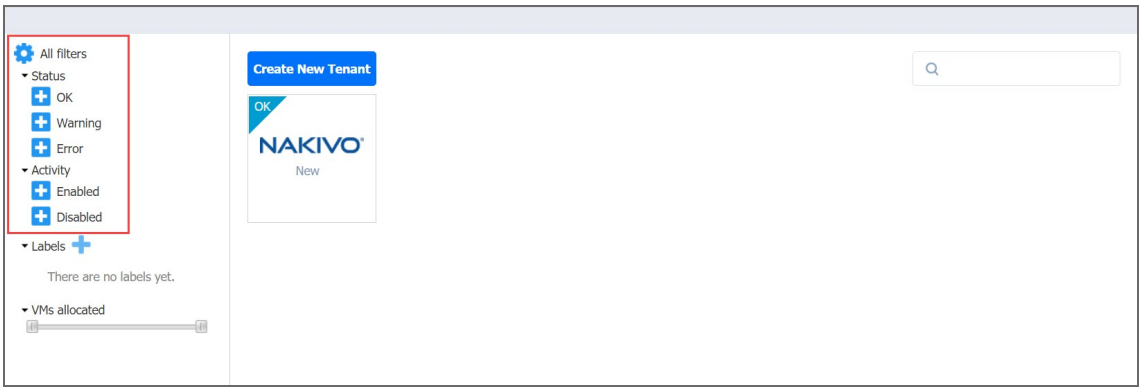

The filters that are currently applied are displayed under the **Active Filters**.

Dismissing Filters

To dismiss a filter, click the filter name under **Active filters**.

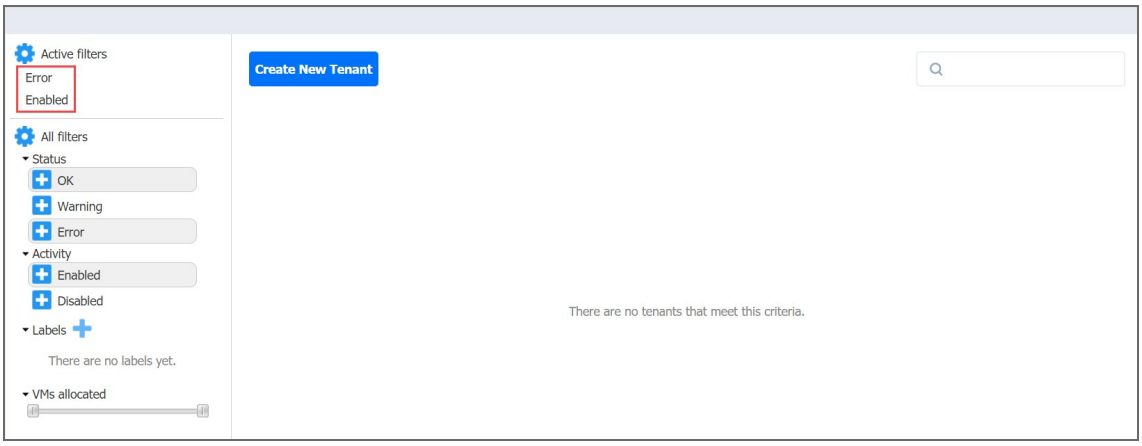
Active filters
Error
Enabled
All filters
Status
OK
Warning
Error
Activity
Enabled
Disabled
Labels
There are no labels yet.
VMs allocated
Create New Tenant
There are no tenants that meet this criteria.

Using Labels

- About Labels
- Creating Labels
- Assigning Labels to Tenants
- Editing Label Names
- Deleting Label

About Labels

With NAKIVO Backup & Replication, you can create custom labels and assign them to tenants. Assigning a label to a tenant allows you to quickly sort existing tenants into different categories, such as location, SLA level, etc.

Creating Labels

To create a new label, click the **Plus** icon next to **Labels** and enter a name for the new label, and press the **Enter** key.

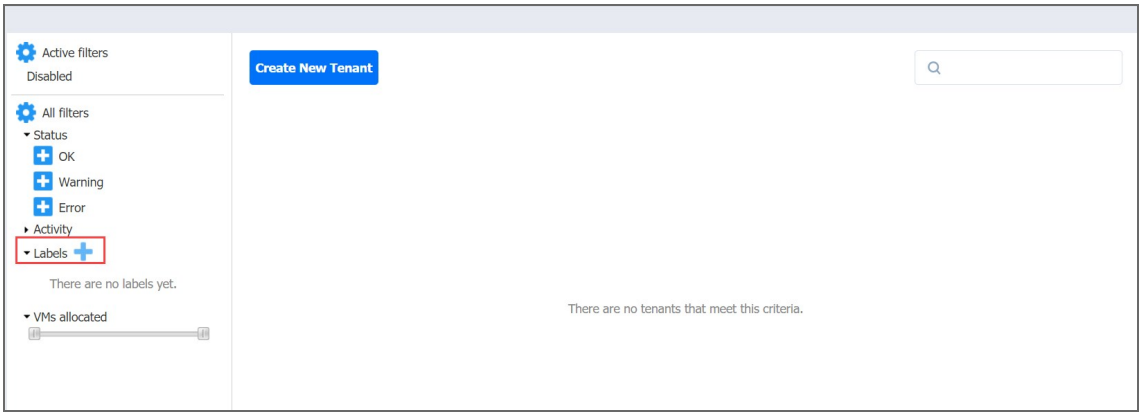

You can also create a new label when creating a new tenant.

Assigning Labels to Tenants

You can assign a label to a tenant either during the tenant creation or by editing the tenant.

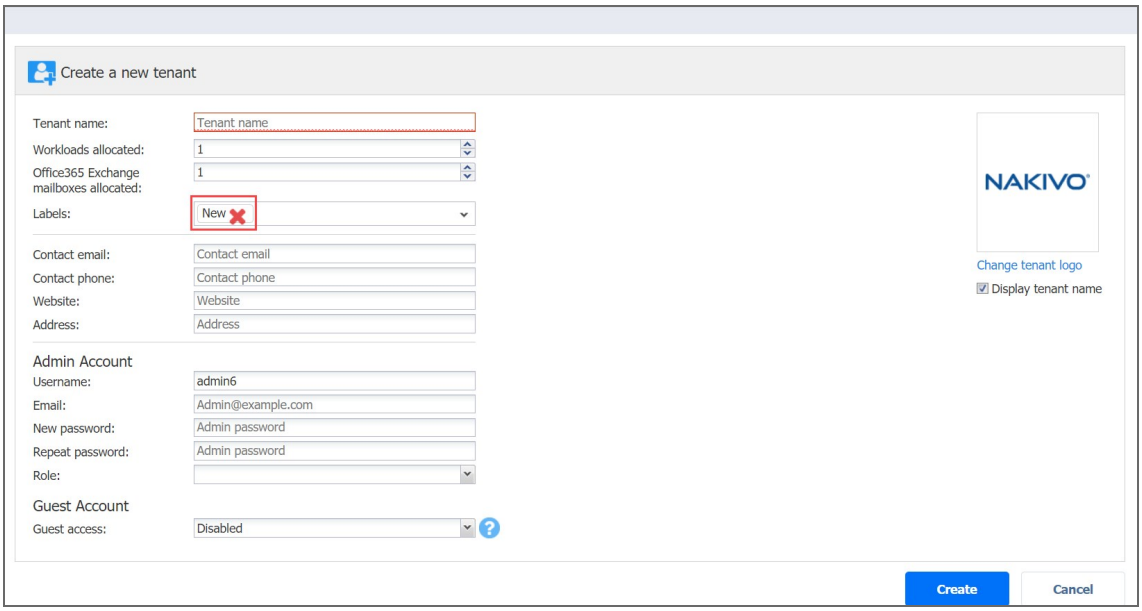

Editing Label Names

To change a label name, do the following:

1. Hover over the label.
2. Click the **Edit** icon.

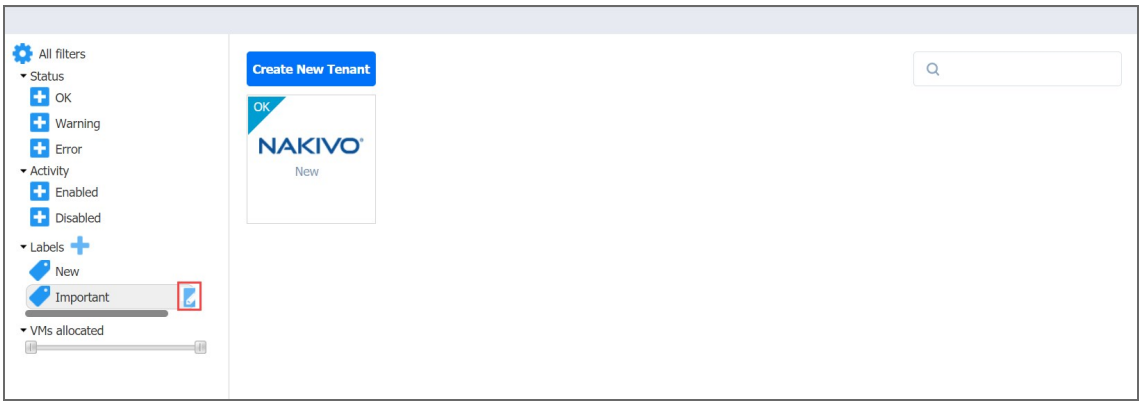

3. Enter the new label name and press the **Enter** key.

Deleting Labels

To permanently delete a label, do the following:

1. Hover the mouse pointer over a label.
2. Click the **Delete** icon.

3. In the dialog box that opens, click **Delete** to confirm that you wish to permanently delete the label



Viewing Tenant Information

To view tenant information, hover over the tenant and click on the **Info** button.

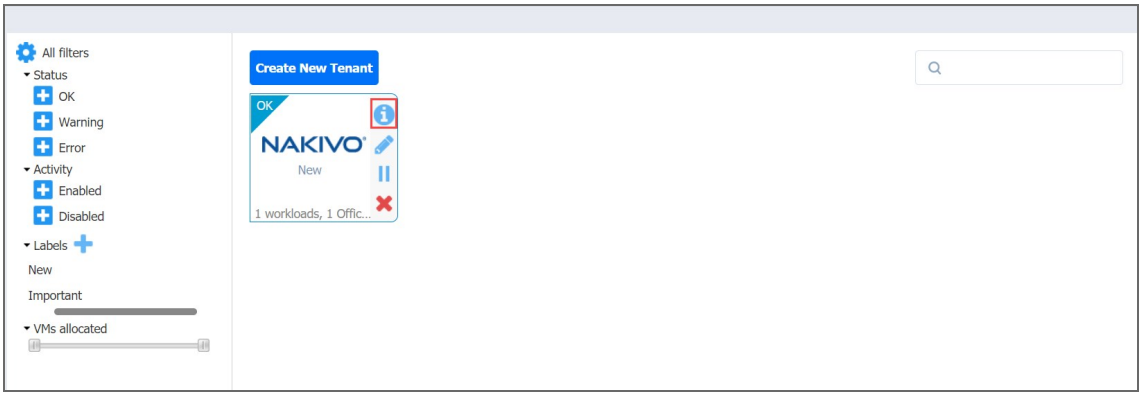

The tenant information is displayed.

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.

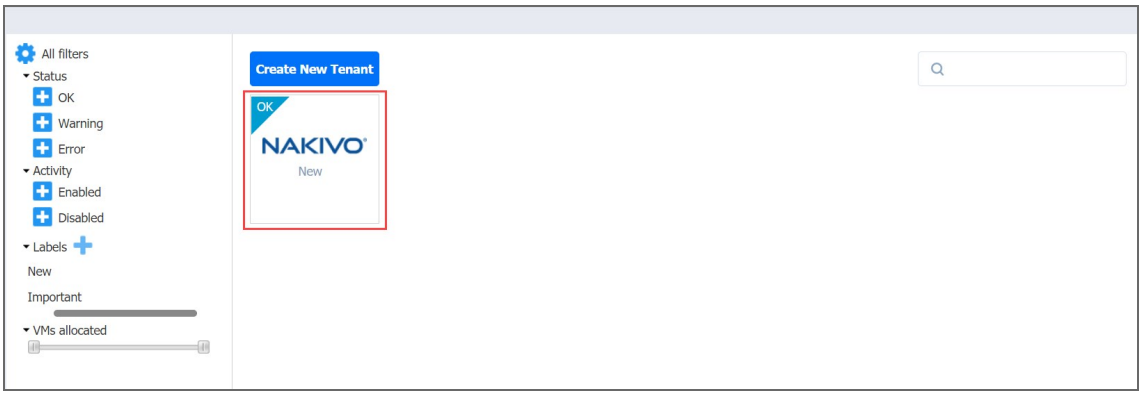

Returning Master Admin Dashboard

To return to the **Master Admin** dashboard, click **Tenants** in the navigation bar.

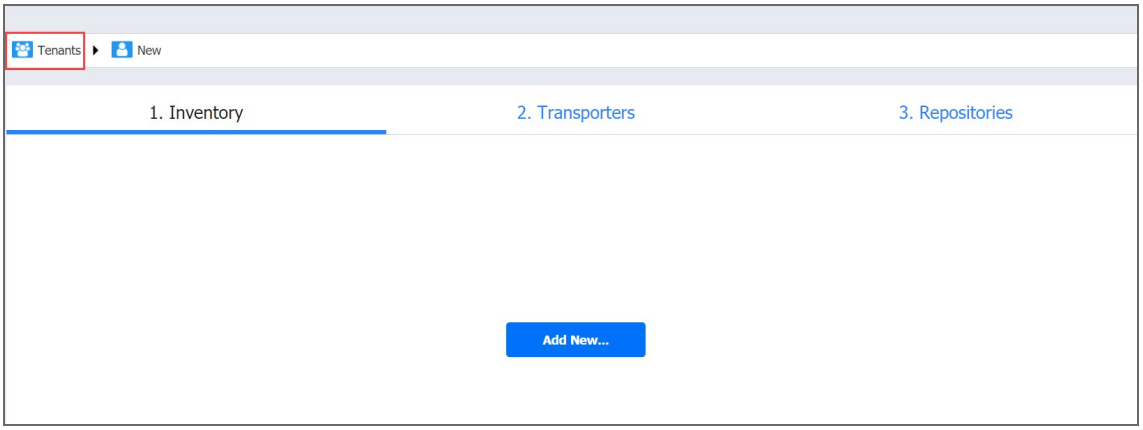

Disabling Tenants

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.

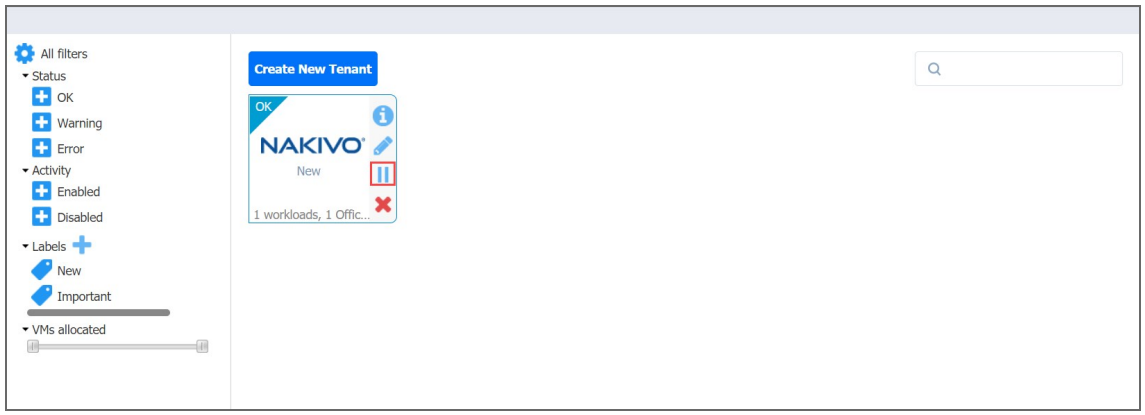

Editing Tenants

To edit a tenant, do the following:

1. Hover over a tenant box and click the **Edit** icon.

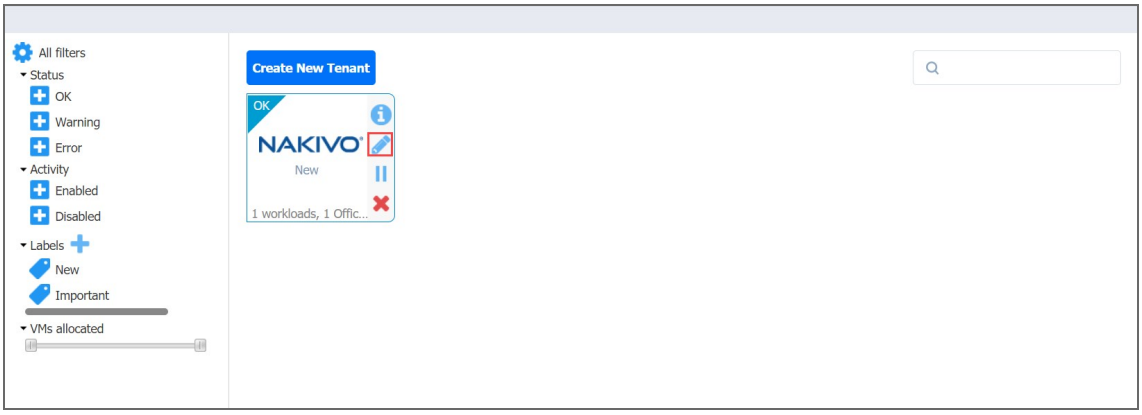

2. In the **Edit** dialog that opens, make the required changes and click **Save**.

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