NAKIVO Backup & Replication v11.2

User Guide for Microsoft 365

Table of Contents

NAKIVO Backup & Replication Overview **5**

Microsoft 365 Backup 6

Microsoft 365 Object Recovery 7

Amazon EC2 Concepts 7

External Product Database Support 9

Recovery Point Retention 9

Backup Encryption 12

Backup Copy To Tape Recovery 15

Backup Immutability 16

Self-Backup Feature 17

Two-Factor Authentication 19

Advanced Bandwidth Throttling 19

Jobs and Concurrent Tasks 22

Network Acceleration 23

Pre and Post Job Scripts 24

BaaS 25

NAKIVO Licensing Policy 32

Deployment **47**

Architecture 48

System Requirements 58

Deployment Scenarios 111

Installing NAKIVO Backup & Replication 125

Updating NAKIVO Backup & Replication 187

Uninstalling NAKIVO Backup & Replication 218

Getting Started **221**

Logging in to NAKIVO Backup & Replication 222

First Steps with NAKIVO Backup & Replication 229

Web Interface Components 232

Managing Jobs and Activities 241

Settings **283**

General 284

Inventory 403

Nodes 430

Backup Repositories 460

Federated Repositories 487

Expert Mode 512

Maintenance Mode 538

Virtual Appliance Configuration 540

Multi-Tenant Mode Configuration 564

Support Bundles 579

Built-in Support Chat 581

Backup **589**

Creating Microsoft 365 Backup Jobs 589

Deleting Backups 614

Object Recovery for Microsoft 365 **623**

Starting Object Recovery for Microsoft 365 624

Microsoft 365 Object Recovery Wizard: Backup 626

Microsoft 365 Object Recovery Wizard: Recovery Account 629
Microsoft 365 Object Recovery Wizard: Objects 630
Microsoft 365 Object Recovery Wizard: Options 634

Backup Copy to Tape Recovery 641

Starting Backup Copy to Tape Recovery Wizard 642
Backup Copy to Tape Recovery Job Wizard: Backups 644
Backup Copy to Tape Recovery Job Wizard: Destination 646
Backup Copy to Tape Recovery Job Wizard: Options 647

Multi-Tenant Mode 652

Creating a Local Tenant 652
Creating a Remote Tenant 660
Using the MSP Console 663
Using the MSP Dashboard 676
Granting Self-Service Access 687

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.

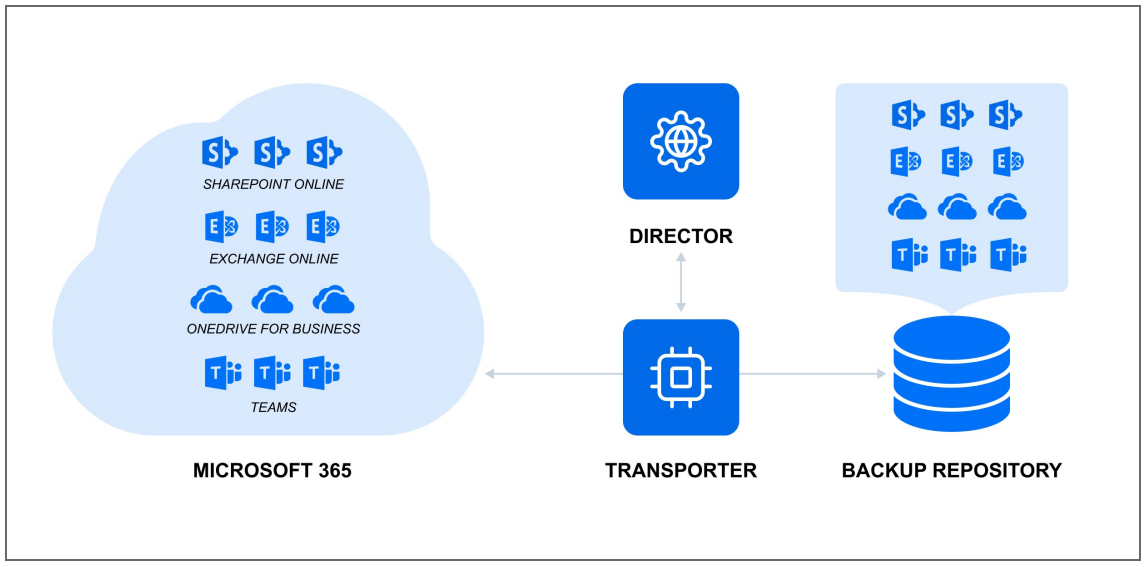
Microsoft 365 Backup

With NAKIVO Backup & Replication, you can back up and recover data from Microsoft Exchange Online, OneDrive for Business, SharePoint Online, and Microsoft Teams.

You can back up items such as user and group mailboxes, drives, sites, teams, and more directly to cloud storage, local storage, CIFS or NFS shares, or deduplication appliances. You can also send copies of Microsoft 365 backups to these destinations as well as to tape for air-gapped storage. When recovery is required, data can be recovered to Microsoft 365 accounts, including mailboxes, emails, contacts, calendar events, files, folders, sites, document libraries, lists, channels, tabs, posts, and other items.

A common misconception among Microsoft 365 users is that data stored in the cloud is automatically protected and does not need to be backed up. Under Microsoft's shared responsibility model, Microsoft is responsible for service availability and for the physical and virtual infrastructure. However, protecting data beyond basic geo-redundancy—which is not equivalent to backup—is the responsibility of the customer.

To address this responsibility, NAKIVO Backup & Replication provides Backup for Microsoft 365 to ensure that application data remains accessible and recoverable at all times. By backing up Microsoft 365 data, you ensure that if data loss occurs, required items can be recovered easily to their original location or to a custom destination.



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

- Add your Microsoft 365 accounts to the inventory
- Create a Microsoft 365 job

Microsoft 365 Object Recovery

NAKIVO Backup & Replication allows you to instantly browse, search, and recover Microsoft 365 objects directly from compressed backups. Microsoft 365 objects can be restored to their original location or to a specified Exchange Online mailbox, OneDrive for Business instance, or SharePoint Online site.
The feature simplifies, automates, and accelerates the data recovery process and is available out of the box in NAKIVO Backup & Replication.
For more details, refer to “Object Recovery for Microsoft 365” on page 623.

Amazon EC2 Concepts

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

Instance

An *Amazon EC2 Instance* is a virtual server in Amazon's Elastic Compute Cloud (EC2). Amazon EC2 offers a variety of instance types so you can select the CPU, memory, storage, and networking capacity that meets your requirements.

EBS Volume

An *Amazon EBS Volume* is a virtual disk that can be attached to any Amazon EC2 Instance within the same Availability Zone. Amazon EBS volumes persist independently of the instance lifecycle. This means that deleting an Amazon EC2 Instance does not remove the EBS volumes previously attached to it.

Region

An *Amazon EC2 Region* is a geographic area where Amazon EC2 Instances are hosted. Multiple Regions are available so you can run your Amazon EC2 Instances in areas that best meet your performance, compliance, or availability needs. Each Region operates independently and is 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 Zones, while still being connected to the others in the same Region through low-latency network links.

VPC

A *virtual private cloud* (VPC) is a virtual network in Amazon EC2. It is dedicated to your AWS account and logically isolated from other networks in the AWS cloud. Similar to a traditional network, you can configure your VPC by selecting IP ranges, creating subnets, configuring routing, gateways, and security settings. After configuration, you can connect your Amazon EC2 Instances to the VPC.

Subnet

A *subnet* is a range of IP addresses in a VPC. You can assign Amazon EC2 Instances to any subnet you choose. Public subnets provide Internet access, while private subnets do not.

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 can associate one or more security groups with it. Each security group contains rules that allow specific traffic. Changes you make to a security group are applied automatically to all associated instances. When Amazon EC2 evaluates whether to allow traffic to reach an instance, it considers all rules from every associated security group.

Key Pair

Amazon EC2 uses *key pairs* to encrypt and decrypt login information. A key pair consists of a public key, which encrypts passwords, and a private key, which decrypts them. When creating a new Amazon EC2 Instance, you can generate a new key pair or assign an existing one. To log in, you must provide the corresponding private key.

- Linux instances do not use passwords; you log in using SSH with the private key.
- Windows instances use the key pair to decrypt the Administrator password before connecting via RDP.

Elastic Network Adapter

Elastic Network Adapter (ENA) is a specialized network interface that enables Enhanced Networking on Amazon EC2 Instances. ENA is designed for high throughput, high packet-per-second performance, and consistently low latency. Depending on the instance type, ENA can provide up to 20 Gbit/s of network bandwidth. For more information, refer to the corresponding article on the AWS website.

External Product Database Support

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

- Database Options
- Troubleshooting External Database Connection Issues

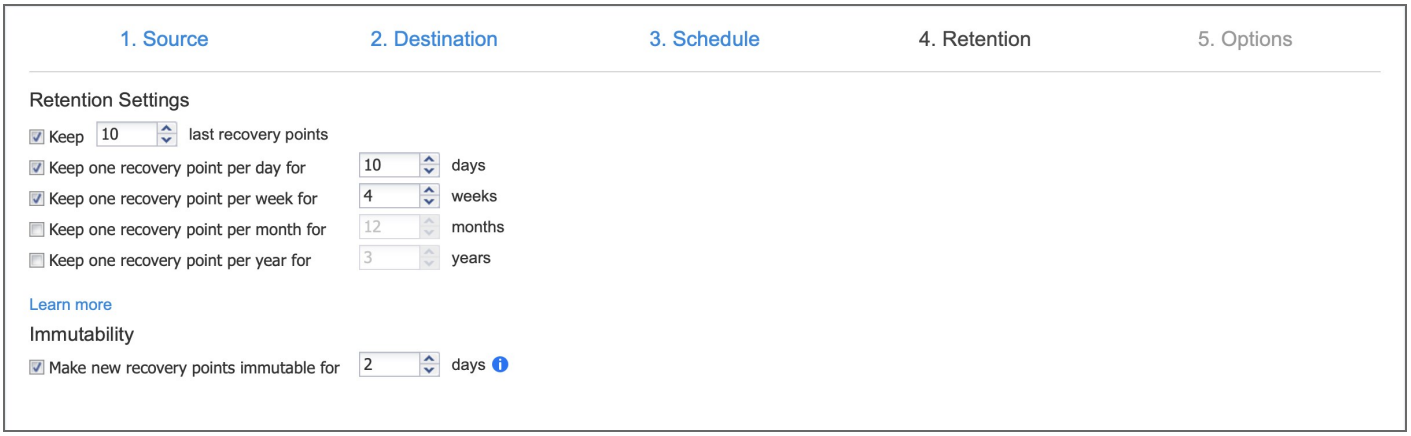
Recovery Point Retention

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

Legacy Retention Approach

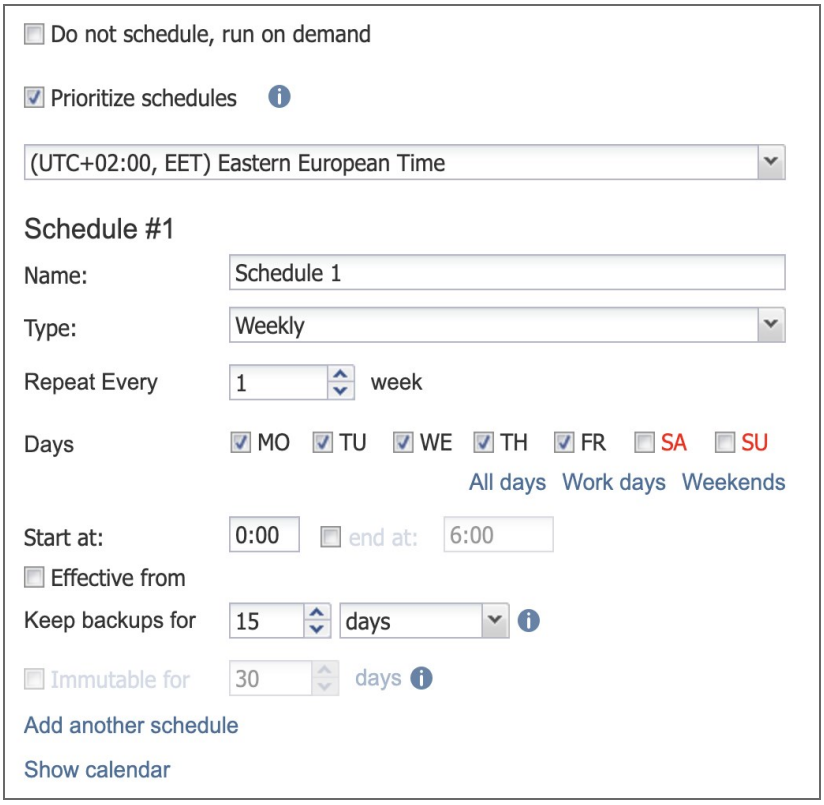
With the legacy retention method, NAKIVO Backup & Replication offers Grand-Father-Son (GFS) retention. This method allows you to save storage space while retaining the recovery points for any period that you need with the following options:

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



Schedule Retention Approach

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



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

Notes

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

- To view the expiration dates and other details of recovery points created with this approach, refer to "Viewing Backup Repository Details" on page 478.
- To learn more about how expiration dates are assigned to recovery points, refer to this article in the Knowledge Base.

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

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

The retention rule applied to the recovery points stored in the backup repository of the **Incremental with full** type is as follows:

- For existing jobs associated with the corresponding backup object (the job currently uses the backup object):
  - The product deletes a full chain of recovery points including the full recovery point if the next full recovery point exists or all recovery points in the chain (including full and all dependent increments until the next full) are expired.
  - The product deletes all incremental recovery points in the chain except for the full recovery point if all incremental recovery points in the chain (excluding the full) are expired.

    Exception:

    The expert setting **system.repository.remove.backups.recoverypoints** is enabled.
- For non-existing jobs associated with the corresponding backup object (the backup object is currently not used by any job):
  - The product deletes a full chain of recovery points including the full recovery point in the following cases:
    - All recovery points in the chain (including the full recovery point and all dependent increments until the next full recovery point) are expired.
    - All recovery points of the backup object are expired and removed, the backup object is removed.

- The product deletes all incremental recovery points in the chain except the full recovery point if all incremental recovery points in the chain (excluding the full recovery point) are expired.

Backup Encryption

NAKIVO Backup & Replication lets you configure encryption to protect backup data, recover encrypted backups, and manage encryption passwords.

How Backup Encryption Works

To encrypt backup data, do the following:

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

During operation:

- The solution automatically generates a password hash based on the user password.
- The cryptographic salt used for hash creation is stored in the recovery point metadata.
- The password hash is used to generate a single-use encryption key that encrypts both backup data and File System Indexing (FSI) data. (FSI must be enabled in the job.)

Note

Data compression is performed before encryption to preserve efficiency. Encryption does not reduce or interfere with compression results.

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

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

- “Backup Job Wizard for Microsoft 365: Options” on page 605
- “Microsoft 365 Object Recovery Wizard: Backup” on page 626

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

Trust Zones

Trust Zones are secure areas or system components where encrypted data is processed, transmitted, or stored. These zones preserve data integrity and confidentiality throughout backup, recovery, and backup copying operations.

- **Backup Jobs**: The trust zone includes the source transporter, which reads data and prepares it for secure transmission.
- **Recovery Jobs**: The trust zone includes the target transporter, which securely writes restored data to the target location.
- **Backup Copy Jobs**: The trust zone may be the source transporter or the target transporter, depending on job configuration.
  To define trust zones:
- Select **Options** > **Backup Encryption** > **Enable on source**, to encrypt the backup data with the source transporter. In this case, the source transporter works in the trusted zone and the encrypted backup can be sent to the target transporter in the untrusted zone.
- Select **Options** > **Backup Encryption** > **Enable on target**, to encrypt the backup data with the target transporter. In this case, the source transporter works in the untrusted zone and the backup encryption can be done with the target transporter in the trusted zone.

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

Reverse Connection

In environments with restrictive firewall policies, a reverse connection can be used. Instead of the source transporter initiating the connection, the target transporter in the untrusted zone connects back to the source using a proxy host. This helps maintain security by ensuring that only validated, encrypted requests are processed.

Restoring from Encrypted Backups

NAKIVO Backup & Replication stores all information required for encryption and decryption, so you do not need to enter the password each time you restore data from encrypted backups. Even if a password is forgotten, the **Backup Encryption** feature enables recovery of encrypted backup data.

To recover encrypted backup data, do the following:

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

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

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

2. Proceed to the next step of the wizard.

Important

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

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

- "Microsoft 365 Object Recovery Wizard: Backup" on page 626

How Encryption/Decryption of System Configuration Works

- To safely encrypt your system configuration bundle, initiate exporting system configuration in the **Settings > General > System Settings > Configuration** tab, set a password, and proceed to exporting.
- To recover from the system configuration bundle, initiate importing system configuration in the **Settings > General > System Settings > Configuration** tab, provide the password to decrypt and import the configuration, and proceed to importing.

Refer to "System Migration" on page 343 for more information.

Password Management

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

Notes

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

Refer to "Managing Passwords" on page 307 for more details.

Backup Copy To Tape Recovery

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

The following backup types are supported:

- File Share Backup
- Microsoft 365 Backup

Note

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

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

Refer to "Starting Backup Copy to Tape Recovery Wizard" on page 642 for more information.

Backup Immutability

When creating a backup job and selecting the **Amazon EC2**, **Amazon S3**, **generic S3-compatible storage**, **Wasabi**, **Azure Blob Storage**, **Backblaze B2 Cloud Storage**, **Dell EMC DataDomain**, **NEC HYDRAstor**, or **Local Folder** type of Backup Repository, NAKIVO Backup & Replication allows you to make the recovery points in these repositories immutable. With immutability enabled, the recovery points are stored using the write-once-read-many (WORM) model. This provides an additional security layer by preventing ransomware encryption as well as accidental deletion or modification of recovery points.

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

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

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

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

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

- "Deploying VMware Virtual Appliance" on page 126
- "Deploying Nutanix AHV Virtual Appliance" on page 132
- "Deploying Amazon Machine Image in Amazon EC2" on page 139

Self-Backup Feature

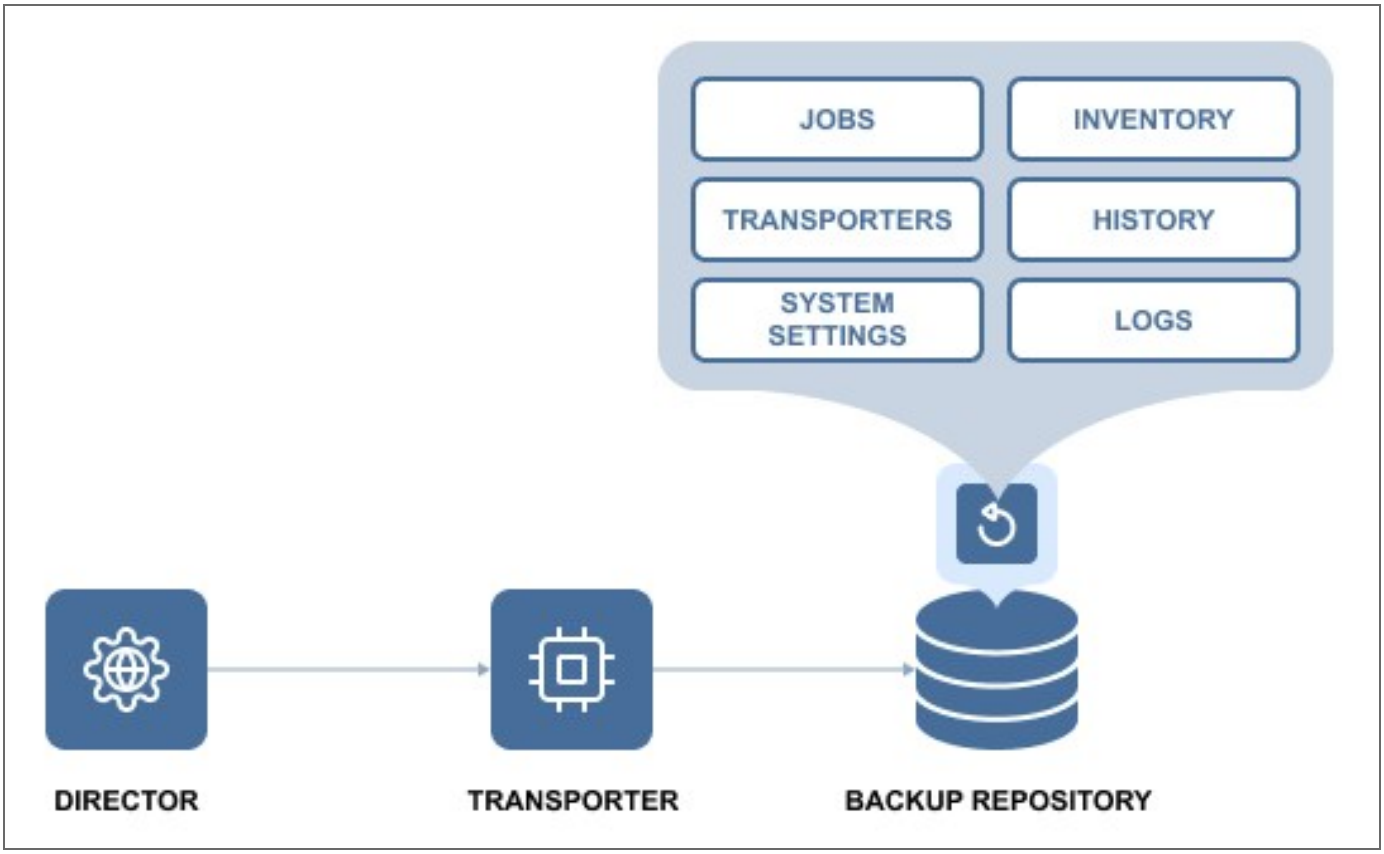
The Self-Backup feature provides automated protection for everything you configure in NAKIVO Backup & Replication. A comprehensive data protection solution must back up not only VMs, but also its own configuration. There are several scenarios in which this becomes critical. For example, the VM running NAKIVO Backup & Replication may become corrupted, infected by malware, or accidentally deleted. Regardless of the cause, restoring the product as quickly as possible is essential.

Although a new instance of NAKIVO Backup & Replication can be deployed in less than one minute, you would still need to restore the product configuration, including jobs, inventory, and backup history. To simplify this process, NAKIVO Backup & Replication automatically backs up the entire configuration. This includes all jobs, inventory data, information about connected Transporters, Backup Repositories, and other system settings.

The Self-Backup feature is enabled by default. NAKIVO Backup & Replication creates self-backups on a daily basis and sends them to the first five available Backup Repositories. By default, each self-backup is retained for five days. If needed, you can customize the backup targets, schedule, and retention policy.

If unintended changes are made to the product configuration, you can easily roll back to a previous system state using a self-backup. Migrating the configuration to a new product instance is also straightforward. You simply deploy a new instance of NAKIVO Backup & Replication, import a Backup Repository that contains a self-backup, and select the desired recovery point. The complete product configuration, along with all settings, is restored.

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



For information on the Self-Backup configuration, refer to “Self-Backup” on page 324.

Two-Factor Authentication

NAKIVO Backup & Replication allows you to add an additional layer of security by enabling two-factor authentication (2FA). When 2FA is enabled, an extra verification step is added to the user login process to help prevent unauthorized access to the solution and the organization's backup data.

With two-factor authentication enabled, user authentication requires entering a verification code generated in one of the following ways:

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

You can find more information in the following articles:

- "Logging in to NAKIVO Backup & Replication" on page 222
- "Configuring Two-Factor Authentication" on page 380

Advanced Bandwidth Throttling

NAKIVO Backup & Replication is designed to transfer data at the maximum available speeds to complete VM backup, replication, and recovery jobs as efficiently as possible. However, running data protection jobs during business hours can overload your LAN or WAN networks. This may impact the performance of business applications and degrade user experience—for example, causing delays in email delivery or slow website loading times. To address this challenge, NAKIVO Backup & Replication provides the flexible **Advanced Bandwidth Throttling** feature. With this feature, you can set limits on your data protection jobs to ensure they never consume more bandwidth than you can allocate.

The Advanced Bandwidth Throttling feature allows you to create global rules that control the data transfer speeds used by your backup processes. These rules can be applied to multiple jobs and scheduled according to your operational requirements. For instance, you can define a global rule that restricts backup jobs to 50 MB/s during business hours while leaving bandwidth usage unrestricted on Sundays. If you require more precise control, you can also configure bandwidth throttling rules on a per-job basis. Individual limits override global rules, eliminating the need to adjust global settings for each job.

The Advanced Bandwidth Throttling feature of NAKIVO Backup & Replication is an effective way to optimize backup operations and manage network traffic. By applying both global and individual limits to data transfer speeds, the feature helps ensure that your business applications maintain consistent performance—even when bandwidth is limited. With these bandwidth rules, the solution can restrict LAN or WAN usage by NAKIVO Backup & Replication jobs to a specific, predefined amount.

For more information, refer to the following sections:

- About Bandwidth Rules

About Bandwidth Rules

A bandwidth rule defines the amount of bandwidth that can be used by one job, several jobs, or all applicable jobs. When a job with multiple VMs runs under an active bandwidth rule, the available bandwidth is automatically distributed among the job's tasks.

Bandwidth rules apply to the following NAKIVO Backup & Replication jobs:

- Backup Jobs for:
  - VMware
  - Hyper-V
  - Nutanix AHV
  - Amazon EC2
  - Physical Machines
  - Proxmox VE
  - VMware Cloud Director
- Backup Copy Jobs (Disk to Disk, Disk to Tape)
- Full VM Recovery Jobs for:
  - VMware
  - Hyper-V
  - Nutanix AHV
  - Amazon EC2
  - Physical Machines (P2V)
  - Proxmox VE
  - VMware Cloud Director
- Instant VM Recovery Job for:
  - VMware
  - Hyper-V
  - Physical Machines (P2V)
  - Proxmox VE
- Recovery from tape
- Replication Jobs for:
  - VMware
  - Hyper-V
  - Proxmox VE
- Replica Failback/Failover Jobs for:
  - VMware
  - Hyper-V

Bandwidth rules may be always active, active on schedule, or disabled. Refer to "Bandwidth Throttling" on page 285 for more details.

A bandwidth rule can be:

- **Global Rule** – a bandwidth rule applied to all supported Jobs (see the supported jobs list above).
- **Per Job Rule** – a bandwidth rule associated with and applied to specific supported Job(s).

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.

Up to 100 bandwidth rules can be created and enforced at the level of a Single-tenant product or a Tenant of a Multi-tenant product. These rules are applied at the Transporter level, stored by the Director, and enforced when processing of a job object begins, provided that the job object falls within the scope of the active rules. Each rule can be enabled or disabled individually.

When enabled, a rule limits the bandwidth of all Job Object Data Transfers (JODTs) that fall under that rule.

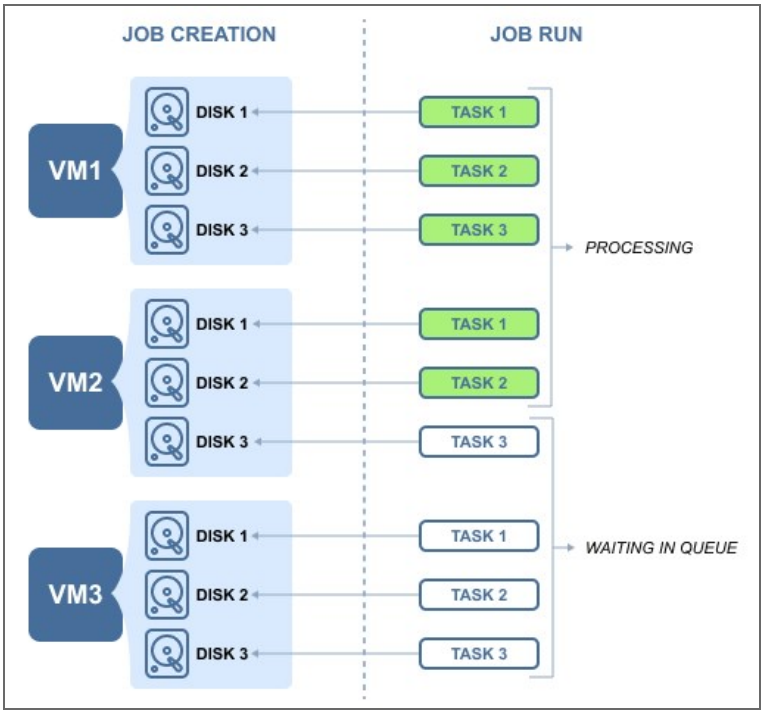
> **Note**
>
> Keep in mind that if the job starts when the Bandwidth throttling rule is active, it will stay at that speed even if the rule time is over.

Jobs and Concurrent Tasks

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

- Backup jobs
- Replication jobs
- Recovery Jobs

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



Each job object may consist of one or more machine disks, Oracle databases, Exchange Online mailboxes, OneDrive for Business instances or SharePoint Online sites that have to be processed within a job run. Data processing that is related to a specific VM disk or service constitutes a single task, in the scope of the corresponding job. Such tasks are processed by a Transporter. For the sake of managing the load over the infrastructure, any Transporter is configured to process a limited number of concurrent tasks. When a task is processed, the Transporter starts processing another task if available. A task can be one disk, file or recovery session, Oracle database, Exchange Online mailbox, OneDrive instance, or a SharePoint Online site. By default, NAKIVO Backup & Replication is set to process 6 concurrent tasks per one Transporter. Refer to “Editing Nodes” on page 448 to learn how to change the Transporter maximum load.

Network Acceleration

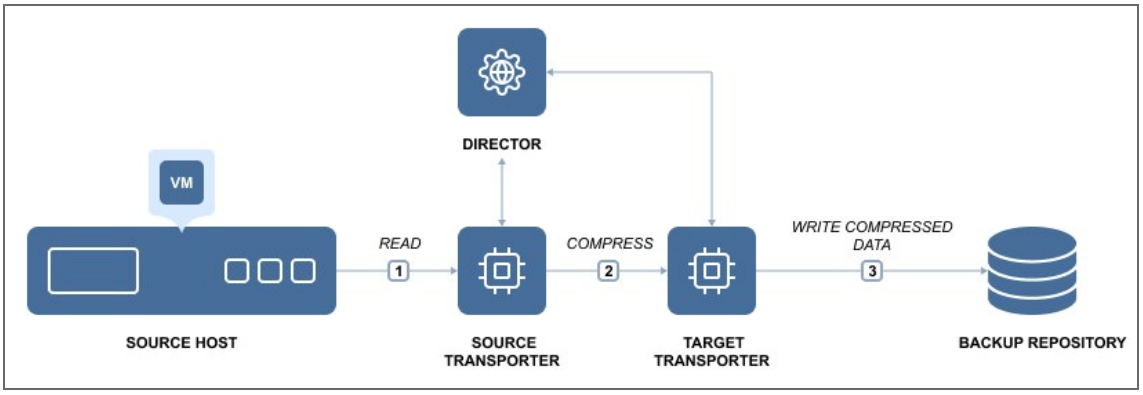
Whether you run VM backup and replication jobs during business hours or send VM backups and replicas offsite over the Internet, efficient use of network bandwidth is critical. NAKIVO Backup & Replication provides the Network Acceleration feature to speed up VM backup and replication jobs, shorten backup windows, and reduce network load simultaneously.

When network acceleration is enabled, NAKIVO Backup & Replication can increase VM backup, replication, and recovery performance by up to two times in WAN environments and busy LAN networks.

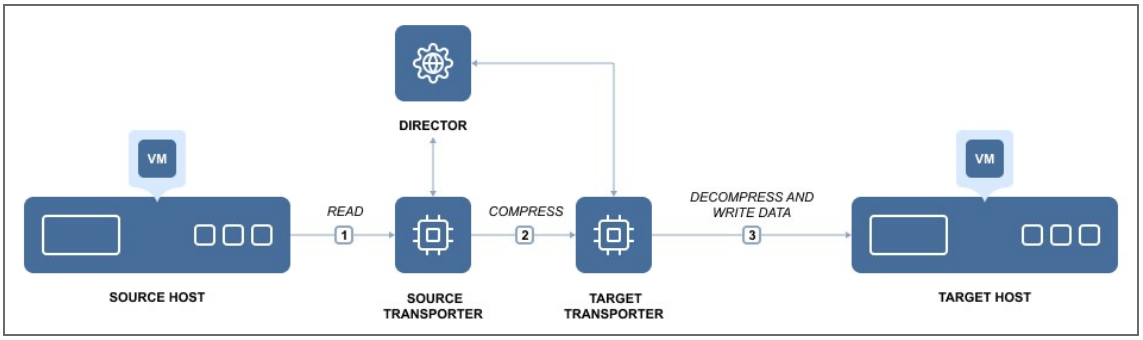
Network acceleration is achieved by the use of two instances of Transporter. Transporter is the product component that performs all data protection and recovery tasks, such as backup, replication, recovery, encryption, and so on. To simplify deployment and configuration, one instance of Transporter is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



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 wake servers, establish connections, mount volumes, start and stop services, send commands to 3rd-party reporting, monitoring and automation tools, and etc.

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 26
- "License Delegation" on page 27
- "MSP Console" on page 29
- "Multi-Tenancy" on page 28
- "Self-Service" on page 31

Branding

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

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

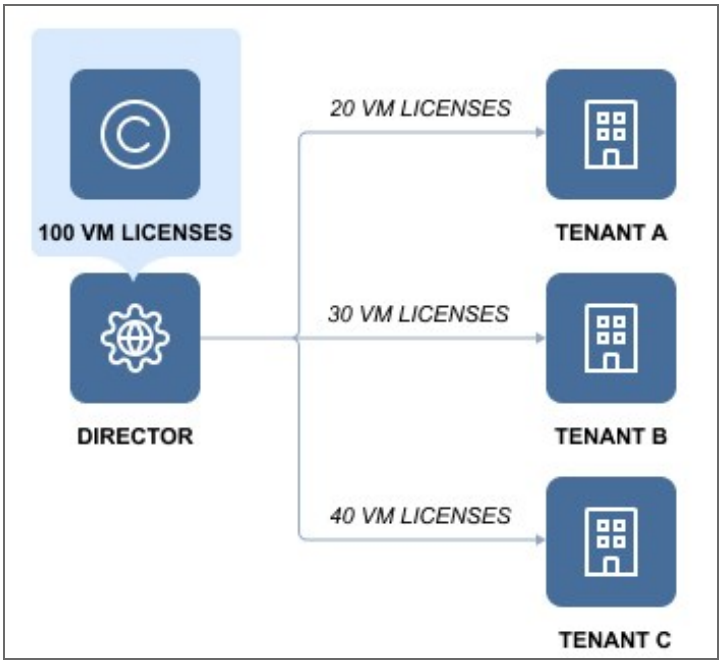
For information on branding configuration, refer to “Branding” on page 288.

License Delegation

In Multi-tenant mode, NAKIVO Backup & Replication allows you to create multiple isolated tenants within a single product instance. These tenants can represent branch offices/departments in enterprise environments or clients in Cloud Provider environments. Because tenants are isolated and require controlled access to licensing resources, NAKIVO Backup & Replication includes the License Delegation feature.

In Multi-tenant mode, a Master Admin (tenant manager) installs a single multi-socket license in the product and then assigns a specific number of licenses to each tenant. This approach provides centralized license management while ensuring that each tenant can use only the licenses allocated to it.

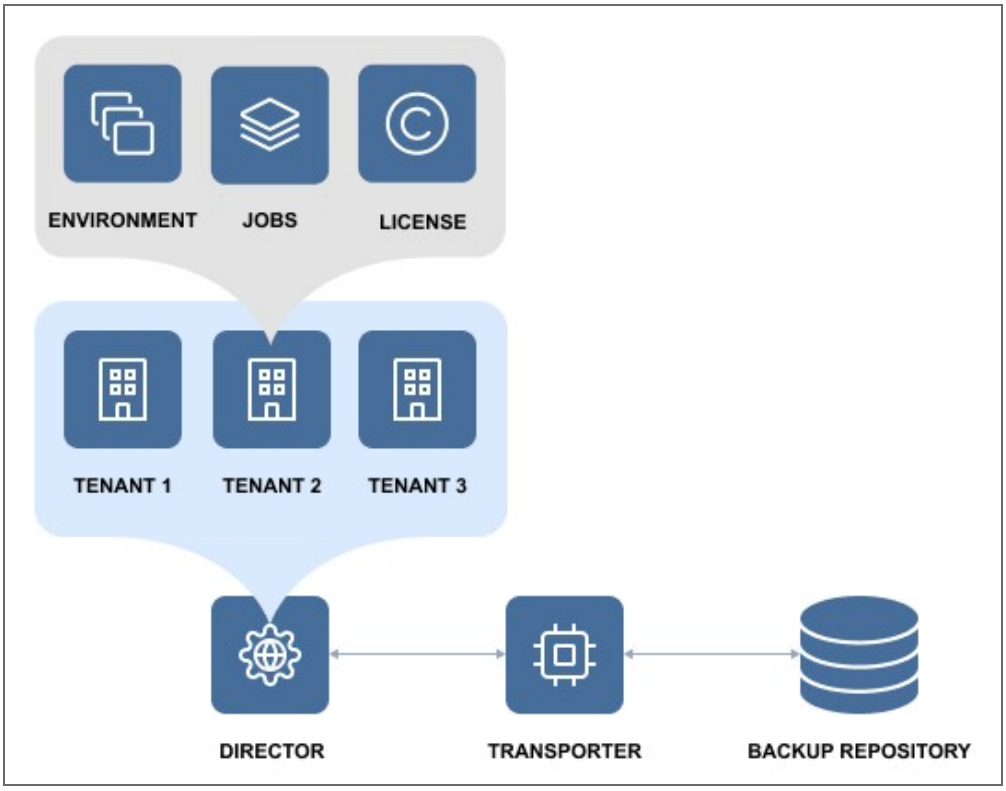
For example, a Master Admin can install a 20-socket license in the Multi-tenant mode of NAKIVO Backup & Replication and assign 3 licenses to Tenant A, 2 licenses to Tenant B, and 4 licenses to Tenant C, leaving the remaining 11 licenses unassigned and available for future use.



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 100 isolated tenants within a single copy of the product. Tenants can represent business units, branch offices, departments, customers, and any other entities.



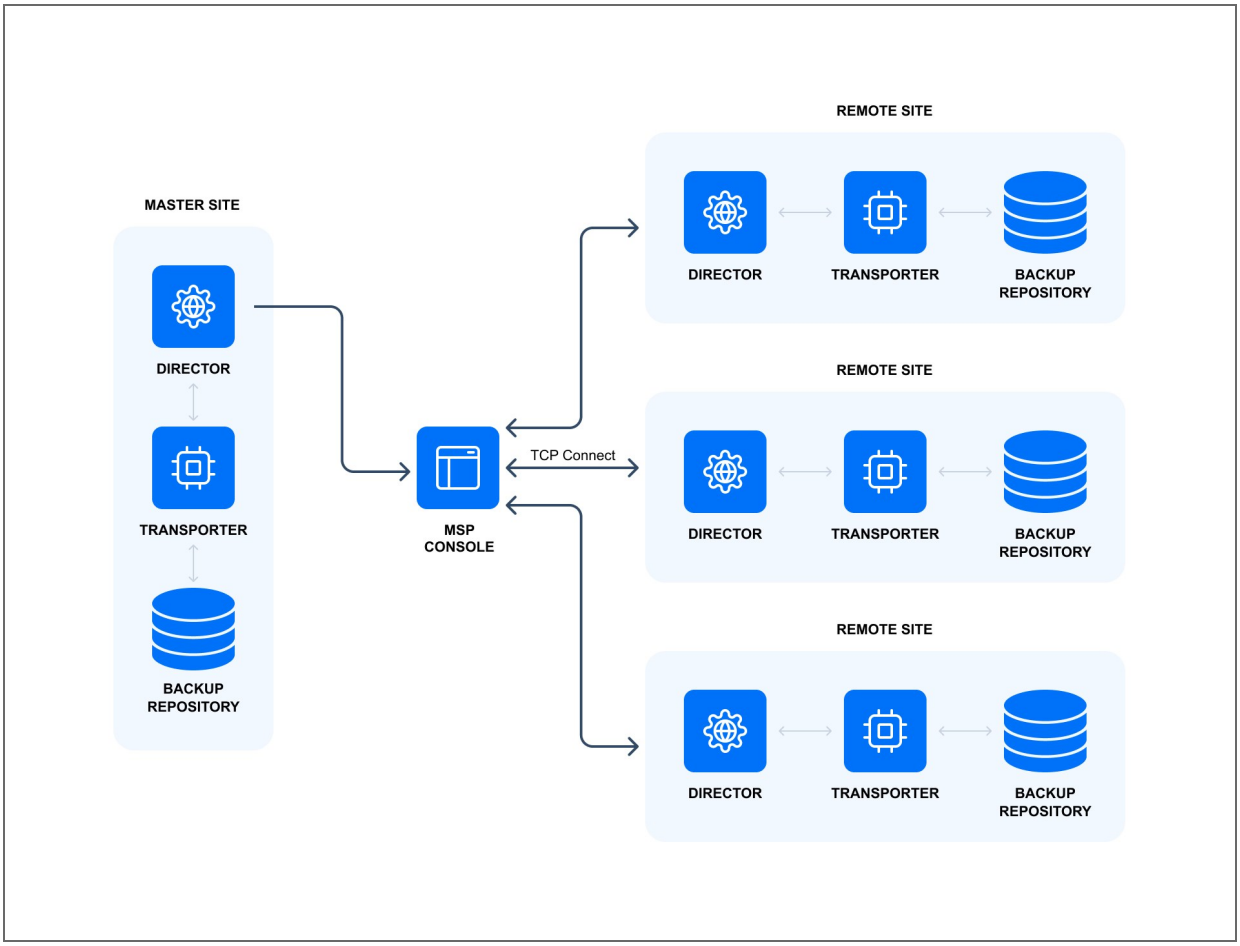
In Multi-tenant mode, each tenant can access their own environment through a self-service portal, and perform all data protection and recovery tasks. At the same time, tenants are isolated from each other and cannot access the environment and jobs of other tenants.

With Multi-tenancy, you can:

- Deliver Backup-as-a-Service, Replication-as-a-Service, and Disaster-Recovery-as-a-Service 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.

MSP Console

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



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

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

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

See the topics below for more information:

- "MSP Architecture" on page 314
- "Using the MSP Console" on page 663

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

- “Creating New Tenants” on page 665
- “Opening MSP Console” on page 664

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

- “MSP Tab in Single-Tenant Mode” on page 314
- “Adding an MSP” on page 317
- “Managing an MSP Connection” on page 319

Self-Service

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

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

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

NAKIVO Licensing Policy

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

- Licensing for NAKIVO Backup & Replication
  - Perpetual Licenses
  - Per-Workload Subscription Licenses
  - License Units
  - NAS File Share Backup Rules
  - License Expiration
- Licensing for Backup for Microsoft 365
- "Licensing for NAS File Share Backup" on page 43
- Real-Time Replication for VMware Licensing
- IT Monitoring Licensing
- MSP Licensing
- Frequently Asked Questions
- Additional Resources

Licensing for NAKIVO Backup & Replication

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

Perpetual Licenses

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

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

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

- It is possible to purchase a license for a single bundle of 5 physical servers in case you also purchase a license for a bundle of 10 physical workstations along with it.
- For NAS backup, the solution is licensed per terabyte (see NAS File Share Backup Rules for more details).
- For Oracle Database, the solution is licensed per database (available with the Enterprise Plus edition only).

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

Notes

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

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

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

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

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

Per-Workload Subscription Licenses

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

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

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

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

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

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

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

License Units

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

Units for Perpetual Licenses

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

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

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

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

Workloads in Per-Workload Subscription License

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

NAS File Share Backup Rules

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

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

    - If the same file share and/or its contents are protected by multiple jobs, the source data is still summed.
    - If a job run reaches or exceeds the licensed data size, the job will become disabled. It will not be possible to create new file share backup jobs, and the current job cannot be re-enabled until it is edited to exclude a sufficient amount of backup data.
- Adding shares to inventory does **not** consume licenses.
- File share backup metadata does **not** contribute to licensed file share size.
- If both per-TB NAS capacity and workloads are present in the license, File Share backups consume the per-TB NAS capacity in the license first. Once the licensed NAS capacity is fully consumed, the system automatically consumes workloads in the license using a conversion rate of 1 workload = 0.5 TB of source file share data.

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

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

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

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

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

License Expiration

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

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

Limited mode is disabled upon changing to the active license.

Licensing for Backup for Microsoft 365

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

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

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

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

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

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

Note

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

The following Student SKUs are supported:

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

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

SharePoint Online Backup Licensing Rules

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

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

Licensing for NAS File Share Backup

File share backups are licensed based on the size of the source file share data. File share backup licensing can be applied **Per workload** or **Per terabyte (TB)**. The minimum licensed amount for File share backups is:

- 1 workload, which covers 500 GB of source file share data (for workload-based licenses)
- 1 TB of source file share data (for per-TB licenses)

Licenses are consumed in increments of 500 GB.

If both per-TB NAS capacity and workload licenses are available:

- File share backups consume the per-TB NAS capacity first.
- After the NAS capacity is fully consumed or exceeded, the system automatically uses workload licenses, where 1 workload covers 500 GB of source file share data.

Notes

- There is no free licensed capacity for File share backups.
- Creating File share backup job does not consume licenses. License consumption starts when file share data is backed up by a job.
- File share backup metadata does not count toward licensed file share size.

Real-Time Replication for VMware vSphere Licensing

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

Perpetual Licenses for VMware Real-Time Replication

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

Note

Valid support is required to receive product updates.

Subscription Licenses for VMware Real-Time Replication

- Licensed per VM
- Annual subscriptions (1 to 5 years) that are billed upfront

- Include 24/7 Support for the licensed period

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

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

IT Monitoring Licensing

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

Perpetual Licenses for IT Monitoring

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

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

Subscription Licenses for IT Monitoring

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

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

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

MSP Licensing

MSP licensing in NAKIVO Backup & Replication is designed specifically for Managed Service Providers that deliver backup and disaster recovery services to multiple customers from a single, centrally managed environment.

NAKIVO offers three MSP editions:

- MSP Pro
- MSP Enterprise
- MSP Enterprise Plus

Each edition defines supported platforms and feature availability. Functionally, MSP editions mirror regular licenses of the same level, with additional MSP-only capabilities.

MSP licenses explicitly enable multi-tenancy:

- One master license is installed on the MSP instance
- Licensed units are allocated to tenants by the MSP administrator
- Tenants are logically isolated but centrally managed
- License usage is tracked at both:
  - Master level
  - Tenant level

MSP licenses are billed on a monthly subscription basis. An MSP may optionally enter into a one-year commitment agreement with NAKIVO. In this case, the partner commits to purchasing a minimum number of licenses each month for the duration of the agreement and receives more favorable pricing compared to a standard monthly subscription.

Notes

- MSP licenses are available only to MSP partners, not to regular end users. Check here to get an MSP quote.
- The MSP Licence allows the Direct Connect and MSP Console features.
- MSP licenses can be allocated on a per-workload basis.

Frequently Asked Questions

Q: What is a socket?

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

Q: Does adding another Transporter require an additional license?

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

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

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

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

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

Additional Resources

NAKIVO Pricing & Editions

NAKIVO Customer Support Policy

NAKIVO Customer Support Agreement

End-User License Agreement

Deployment

This section contains the following topics :

- “Architecture” on page 48
- “System Requirements” on page 58
- “Deployment Scenarios” on page 111
- “Installing NAKIVO Backup & Replication” on page 125
- “Updating NAKIVO Backup & Replication” on page 187
- “Uninstalling NAKIVO Backup & Replication” on page 218

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 49
- “Transporter” on page 51
- “Backup Repository” on page 55

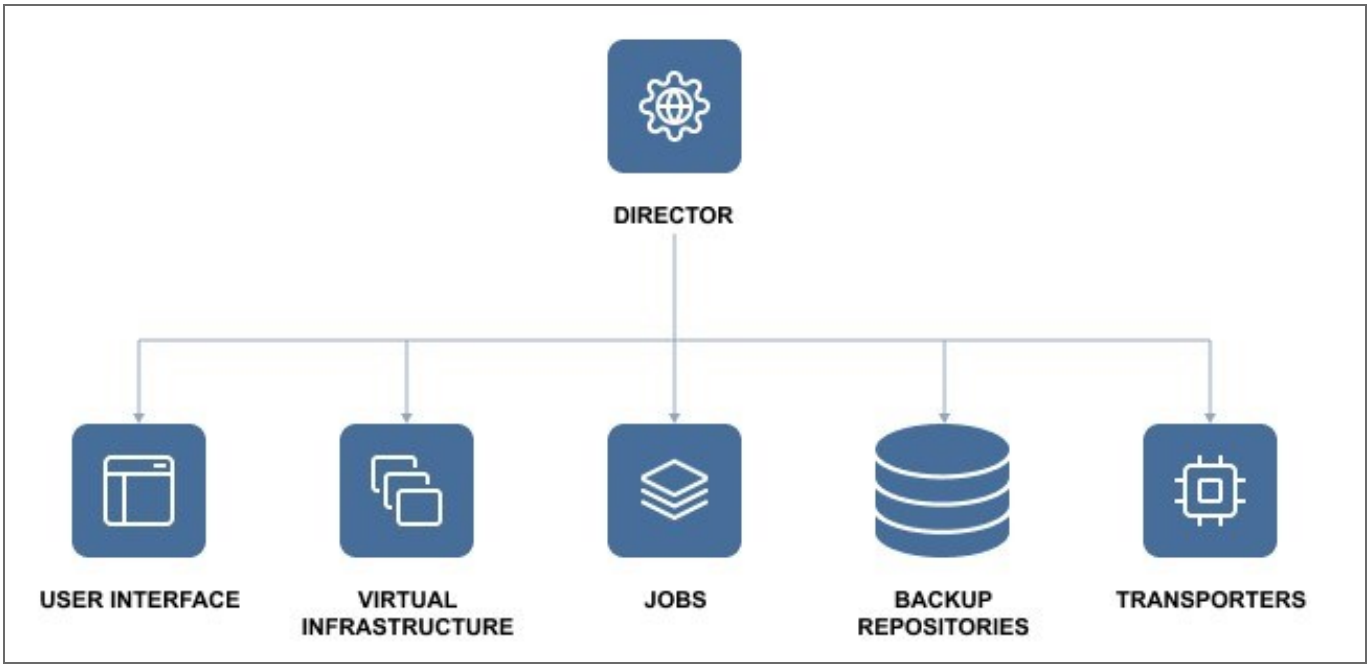
All components can be installed on a single machine or can be distributed across multiple machines and geographical locations.

Director

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

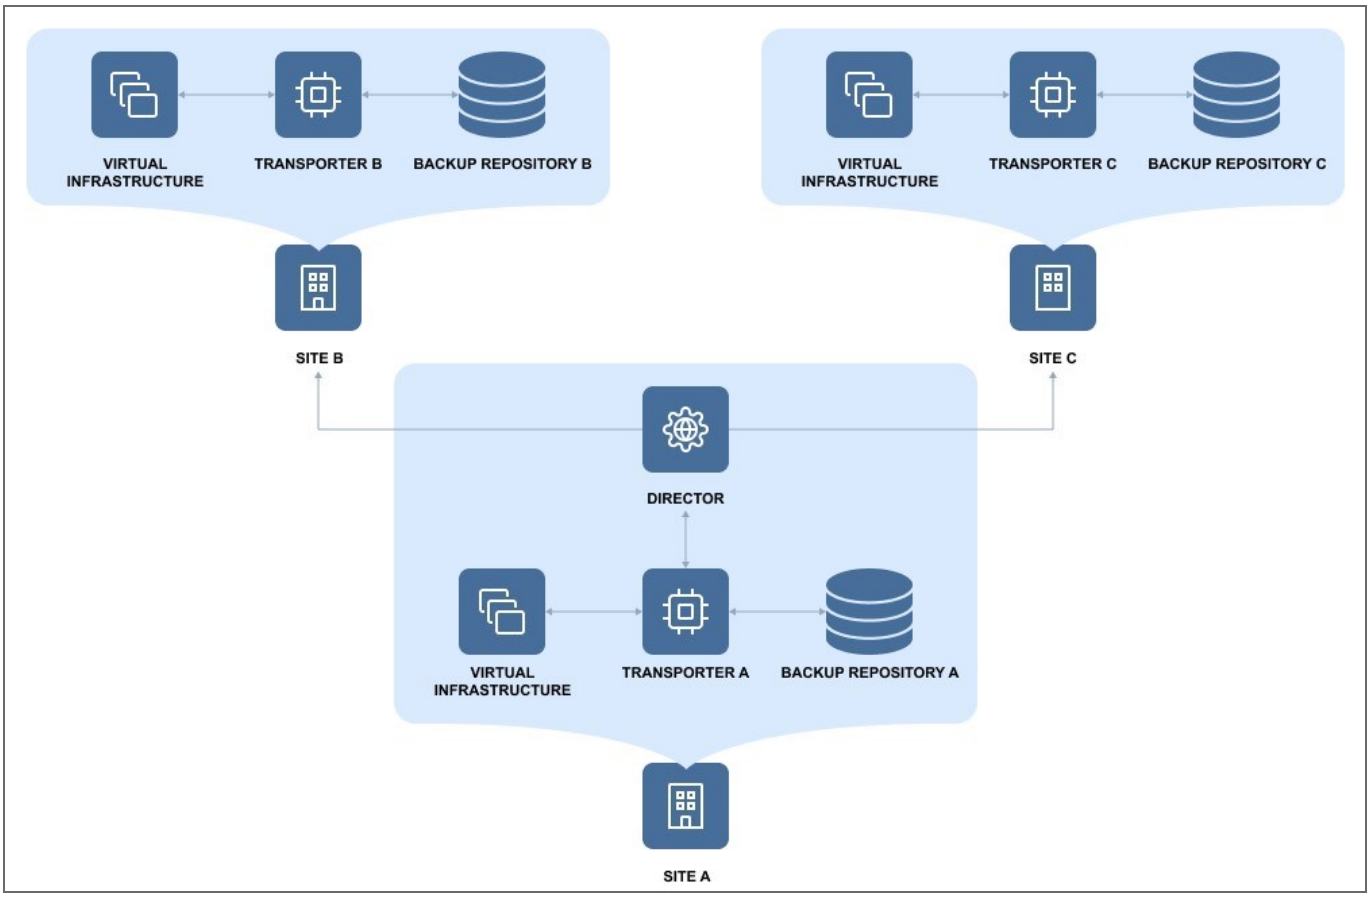
What is Director?

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



How Many Directors Should be Deployed

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

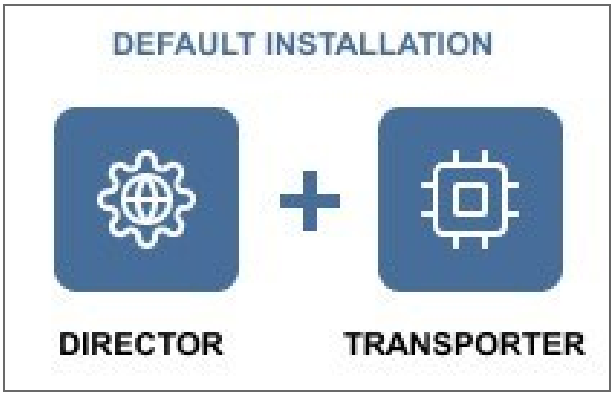

VIRTUAL INFRASTRUCTURE
TRANSPORTER B
BACKUP REPOSITORY B
VIRTUAL INFRASTRUCTURE
TRANSPORTER C
BACKUP REPOSITORY C
SITE B
SITE C
DIRECTOR
VIRTUAL INFRASTRUCTURE
TRANSPORTER A
BACKUP REPOSITORY A
SITE A

Transporter

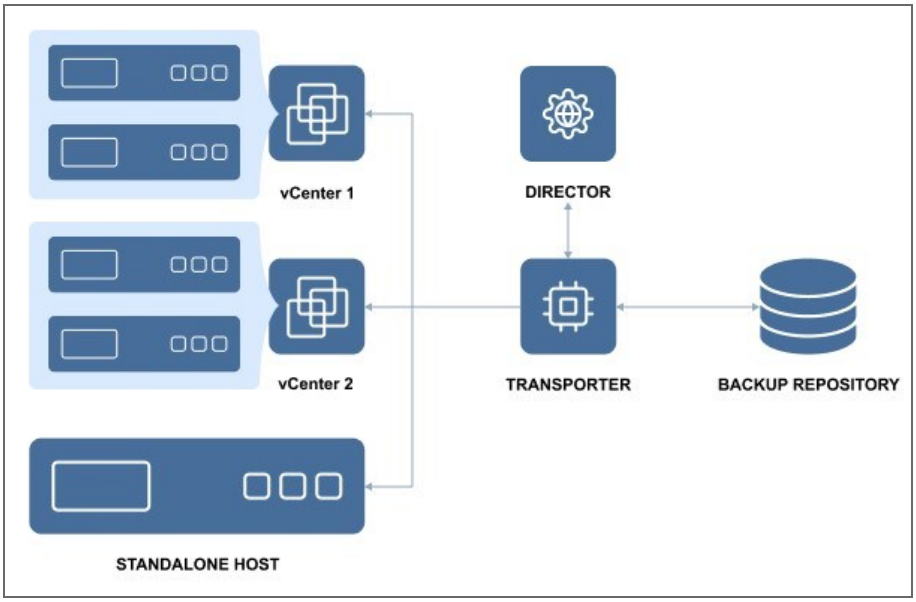
- What is a Transporter?
- How many Transporters Should be Deployed?
- How Transporters are Selected for Jobs
- Transporter Security

What is a Transporter?

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



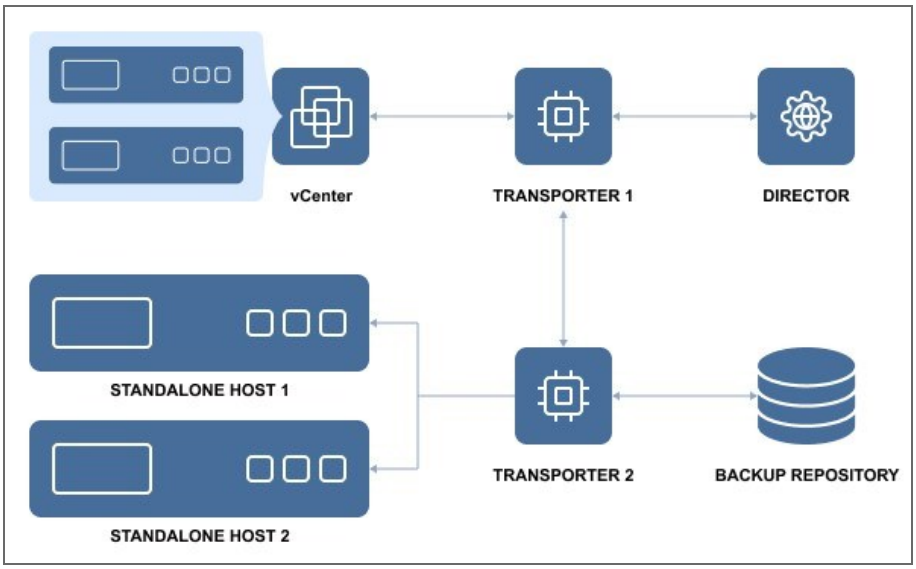
A single **Transporter** can back up, replicate, and recover multiple VMs and cloud instances.



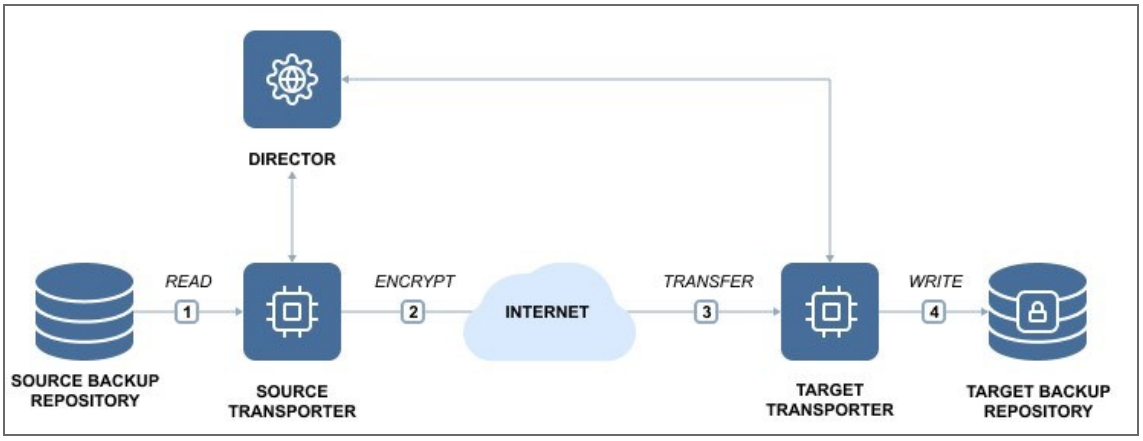
One Transporter can simultaneously process multiple source disks (6 by default) during backup, replication, and recovery. If jobs contain more disks than the **Transporter** is set to process simultaneously, the disks will be put in a queue and will be processed once the **Transporter** frees up.

How Many Transporters Should be Deployed?

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



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

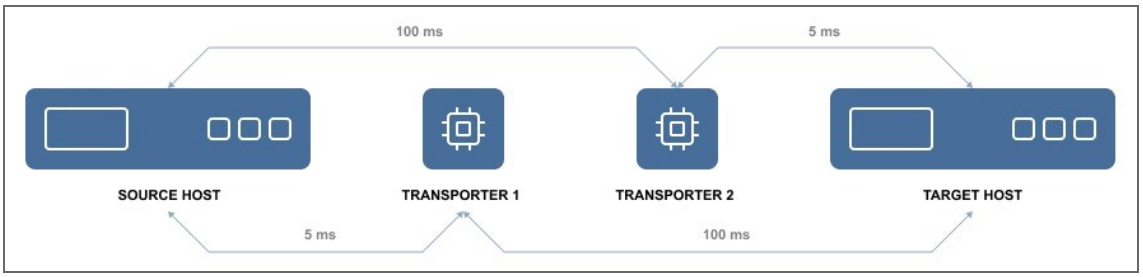


If you plan to transfer data over WAN without a VPN connection from your source site to the target site, make sure the source and target **Transporters** are added to the product using external IP addresses or DNS names that can be properly resolved in WAN, so that the two **Transporters** can connect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

How Transporters are Selected for Jobs

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

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



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

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

Transporter Security

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

This option is available for the following supported target platforms:

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

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

Backup Repository

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

What is a Backup Repository?

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

Important

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

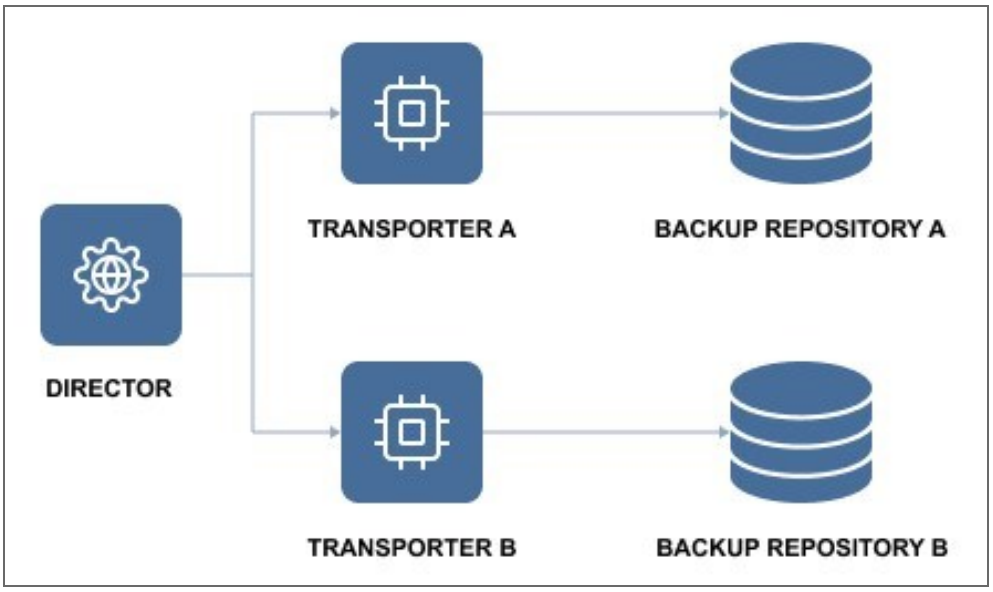
By default, a **Backup Repository** is created when the full solution (both **Director** and **Transporter**) is installed. The default **Backup Repository** is named “Onboard repository”.

How Much Data Can Be Stored in a Backup Repository?

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

How is a Backup Repository Managed?

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



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

Federated Repository

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

What Is a Federated Repository?

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

Important

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

How Does a Federated Repository Work?

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

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

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

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

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

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

How Many Members Can Be Added to a Federated Repository?

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

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

Note

You cannot simultaneously select:

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

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

System Requirements

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

- "Supported Platforms" on page 59
- "Deployment Requirements" on page 62

Supported Platforms

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

- VMware vSphere v5.5 - 9.0.1.0 (Free ESXi is supported for 8.0U3e, 8.0U3f, 8.0U3g)
- VMware Cloud Foundation (VCF)
  This platform is supported through its individual VMware components. For details on how specific VMware components correspond to particular VCF versions, refer to the VMware Knowledge Base article.
- Nutanix AHV v10.0 - Nutanix AHV v10.3
- Microsoft Windows 8 Professional
- Microsoft Windows 10 Home/Pro/Enterprise (20H2 / 21H1 / 21H2 / 22H2)
- Microsoft Windows 11 Home/Pro/Enterprise (21H2 / 22H2 / 23H2 / 24H2)
- Microsoft 365 (Exchange Online, OneDrive for Business, SharePoint Online, Teams)

Notes

- To learn about the limitations of NAKIVO Backup & Replication related to supported platforms, refer to the Platform Limitations section of the latest Release Notes.
- To add a supported platform to NAKIVO Backup & Replication, make sure that your system has been updated with the latest patch and all the necessary requirements are met.
- The support for sub-versions that are not stated in the user guide can be clarified with the support team.
- NAKIVO Backup & Replication does not support backup and recovery of ACL data via NFSv3.
- For ReFS and NTFS volumes containing files with Reparse Point (L) attributes, enable the Windows Data Deduplication feature on the Transporter host that mounts backup data. If this feature is missing, such files may not be recovered or included in the output archive. For more information refer to File-Level Recovery Fails on Windows.

Find the necessary requirements below:

- Microsoft 365 Requirements

Microsoft 365 Requirements

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

- Exchange Online must be purchased as a part of the Microsoft 365 plan or a standalone service.
- The required API permissions must be provided to NAKIVO Backup & Replication. Refer to Required API Permissions for Microsoft 365.

- For automatic enabling of the "Site Collection Administrator" role, a Microsoft 365 account user must have the "SharePoint admin" or "Global admin" role. This is required for backing up and recovery of SharePoint Online sites.
- Microsoft 365 account must be accessible over the network.
- Conditional Access must be disabled in Azure portal. Alternatively, the account must be added as "excluded" from the Conditional Access.
- When recovering from an incremental with full repository, the */nakivo/transporter/java/tmp* folder on the installed transporter must have more space than the size of 3 largest items being recovered, otherwise recovery may fail.
- To backup and restore Microsoft 365 services, a Transporter should be installed on one of the following operating systems:

Windows

- Windows Server 2025 (x64)
- Windows Server 2022 (21H2) (x64)
- Windows Server 20H2 (x64)
- Windows Server 2019 Standard (x64)
- Windows Server 2016 Standard (x64)
- Windows Server 2012 Standard (x64)
- Windows Server 2012 R2 Standard (x64)
- Microsoft Windows 11 (x64):
    - Home, Pro, and Enterprise editions
    - Versions: 21H2, 22H2, 23H2, and 24H2
- Microsoft Windows 10 (x64):
    - Home, Pro, and Enterprise
    - Versions: 20H2, 21H1, 21H2, and 22H2

Linux

- Debian 13 (64-bit)
- Debian 12.1 (64-bit)
- Debian 11.7 (64-bit)
- Debian 11.6 (64-bit)
- Debian 11.5 (64-bit)
- Debian 11.4 (64-bit)
- Debian 11.3 (64-bit)
- Debian 11.2 (64-bit)
- Debian 11.1 (64-bit)
- Debian 11.0 (64-bit)

- Debian 10.13 (64-bit)
- Debian 10.12 (64-bit)
- Debian 10.11 (64-bit)
- Debian 10.10 (64-bit)
- Debian 10.9 (64-bit)
- Debian 10.8 (64-bit)
- Debian 10.7 (64-bit)
- Debian 10.6 (64-bit)
- Debian 10.5 (64-bit)
- Debian 10.4 (64-bit)
- Debian 10.3 (64-bit)
- Debian 10.2 (64-bit)
- Debian 10.1 (64-bit)
- Ubuntu 26.04 Server (x64)
- Ubuntu 25.04 Server (x64)
- Ubuntu 24.10 Server (x64)
- Ubuntu 24.04 Server (x64)
- Ubuntu 22.04.4 Server (x64)
- Ubuntu 22.04 Server (x64)
- Ubuntu 20.04 Server (x64)
- Ubuntu 18.04 Server (x64)
- Ubuntu 16.04 Server (x64)
- SUSE Linux Enterprise Server 12 SP1 – 15 SP6 (x64)
- Red Hat Enterprise Linux v7.0 – v10.0 (x64)
- CentOS v7.0 – v8.5 (x64)
- CentOS Stream 8 – 10 (x64)

NAS

Refer to "Deployment Requirements" on page 62 to see the list of supported NAS devices.

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

Deployment Requirements

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

- Hardware
  - VM or Physical Machine
  - Network Attached Storage
  - Scalability and UI Performance
- Operating Systems
- Networking Requirements
  - Required TCP Ports
  - “External Resources” on page 72
  - 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 RAM
    - additional 150 MB RAM 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 RAM
    - additional 150 MB RAM for each concurrent Java Transporter task
- **Free space**: 5 GB

Note

ARMv7 CPU is not supported.

Network Attached Storage

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

Director and Onboard Transporter:

- **CPU**: x86-64, 2 cores
- **RAM**: 1 GB + 250 MB for each concurrent task
  - For SaaS Backup Repository-related activities:
    - additional 4 GB RAM
    - additional 150 MB RAM 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:
    - additional 4 GB RAM
    - additional 150 MB for each concurrent Java Transporter task
- **Free space**: 5 GB

Note

- To ensure smoother operation, at least 8 GB of RAM is recommended, particularly if you expect to run concurrent backup or recovery tasks.
- Onboard Transporters installed on NAS devices with ARM CPU do not support VMware infrastructures. Refer to Transporter Does Not Support VMware vSphere for a solution.
- ARMv7 CPU is not supported.

Supported NAS Devices

- **Synology**: For a full list of supported models, refer to “Supported Synology NAS Devices” on page 87
- **QNAP**: For a full list of supported models, refer to “Supported QNAP NAS Devices” on page 79
- **ASUSTOR**: For a full list of supported models, refer to “Supported ASUSTOR NAS Devices” on page 74
- **NETGEAR**: For a full list of supported. For a full list of supporter models, refer to “Supported NETGEAR NAS Devices” on page 76.
- **Western Digital**: For a full list of supported models, refer to “Supported Western Digital NAS Devices” on page 91.

Scalability and UI Performance

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

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

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

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

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

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

The maximum number of sessions per Transporter is 25. Note that exceeding these limits may result in unexpected errors and performance issues.

Operating Systems

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

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

Windows

- Microsoft Windows Server 2025 (x64)
- Microsoft Windows Server 2022 (21H2) (x64)
- Microsoft Windows Server 20H2 (x64)
- Microsoft Windows Server 2019 Standard (x64)
- Microsoft Windows Server 2016 Standard (x64)
- Microsoft Windows Server 2012/2012 R2 (x64)

- Microsoft Windows 11 (x64)
    - 21H2
    - 22H2
    - 23H2
- Microsoft Windows 10 Enterprise (x64)
- Microsoft Windows 10 Home (x64)
- Microsoft Windows 10 Professional (x64)
- Microsoft Windows 8 (x64)

Linux

- Rocky Linux 9.5 (64-bit)
- Rocky Linux 9.4 (64-bit)
- Rocky Linux 9.3 (64-bit)
- Rocky Linux 9.2 (64-bit)
- Rocky Linux 9.1 (64-bit)
- Rocky Linux 9.0 (64-bit)
- Rocky Linux 8.9 (64-bit)
- Rocky Linux 8.8 (64-bit)
- Rocky Linux 8.7 (64-bit)
- Rocky Linux 8.6 (64-bit)
- Rocky Linux 8.5 (64-bit)
- Rocky Linux 8.4 (64-bit)
- Rocky Linux 8.3 (64-bit)
- Debian 12.10 (64-bit)
- Debian 12.9 (64-bit)
- Debian 12.8 (64-bit)
- Debian 12.7 (64-bit)
- Debian 12.6 (64-bit)
- Debian 12.5 (64-bit)
- Debian 12.4 (64-bit)
- Debian 12.3 (64-bit)
- Debian 12.2 (64-bit)
- Debian 12.1 (64-bit)
- Debian 11.7 (64-bit)
- Debian 11.6 (64-bit)
- Debian 11.5 (64-bit)
- Debian 11.4 (64-bit)

- Debian 11.3 (64-bit)
- Debian 11.2 (64-bit)
- Debian 11.1 (64-bit)
- Debian 11.0 (64-bit)
- Debian 10.13 (64-bit)
- Debian 10.12 (64-bit)
- Debian 10.11 (64-bit)
- Debian 10.10 (64-bit)
- Debian 10.9 (64-bit)
- Debian 10.8 (64-bit)
- Debian 10.7 (64-bit)
- Debian 10.6 (64-bit)
- Debian 10.5 (64-bit)
- Debian 10.4 (64-bit)
- Debian 10.3 (64-bit)
- Debian 10.2 (64-bit)
- Debian 10.1 (64-bit)
- Oracle Linux 9.5 (64-bit)
- Oracle Linux 9.4 (64-bit)
- Oracle Linux 9.3 (64-bit)
- Oracle Linux 9.2 (64-bit)
- Oracle Linux 9.1 (64-bit)
- Oracle Linux 9.0 (64-bit)
- Oracle Linux 8.9 (64-bit)
- Oracle Linux 8.8 (64-bit)
- Oracle Linux 8.7 (64-bit)
- Oracle Linux 8.6 (64-bit)
- Oracle Linux 8.5 (64-bit)
- Oracle Linux 8.4 (64-bit)
- Oracle Linux 8.3 (64-bit)
- Oracle Linux 8.2 (64-bit)
- Oracle Linux 8.1 (64-bit)
- Oracle Linux 8.0 (64-bit)
- Oracle Linux 7.9 (64-bit)
- Oracle Linux 7.8 (64-bit)
- Oracle Linux 7.7 (64-bit)

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

- Red Hat Enterprise Linux 7.5 (x64)
- Red Hat Enterprise Linux 7.4 (x64)
- CentOS Stream 10 (x64)
- CentOS Stream 9 (x64)
- CentOS Stream 8 (x64)
- CentOS Linux 8.5 (x64)
- CentOS Linux 8.4 (x64)
- CentOS Linux 8.3 (x64)
- CentOS Linux 8.2 (x64)
- CentOS Linux 8.1 (x64)
- CentOS Linux 8.0 (x64)
- CentOS Linux 7.9 (x64)
- CentOS Linux 7.8 (x64)
- CentOS Linux 7.7 (x64)
- CentOS Linux 7.6 (x64)
- CentOS Linux 7.5 (x64)
- CentOS Linux 7.4 (x64)
- CentOS Linux 7.3 (x64)
- CentOS Linux 7.2 (x64)
- CentOS Linux 7.1 (x64)
- CentOS Linux 7.0 (x64)
- AlmaLinux 9.5 (64-bit)
- AlmaLinux 9.4 (64-bit)
- AlmaLinux 9.3 (64-bit)
- AlmaLinux 9.2 (64-bit)
- AlmaLinux 9.1 (64-bit)
- AlmaLinux 9.0 (64-bit)
- AlmaLinux 8.7 (64-bit)

NAS

- ASUSTOR ADM v3.5 - v5.1.1.RCI1
- Netgear ReadyNAS OS v6.9 - v6.10.9
- Synology DSM v6.0 - v7.3.2
- QNAP QTS v4.3 - v5.2.8.3359
- QNAP QuTS Hero h4.5.3 - h5.2.0
- QNAP QuTScloud v4.5.1 - c5.1.0

- WD MyCloud v5
- TrueNAS CORE v13.0-U6.2 - v13.3-U1

Supported Operating System Localizations

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

- English
- Italian
- German
- French
- Spanish

Networking Requirements

Required TCP Ports

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

TCP Port (Default)	Where	Description
NAKIVO Backup & Replication		
80 8080	Director	Used to redirect to the secure TCP port (HTTPS) 4443.
4443	Director	Used to access the Director web UI. Must be opened on the Director machine.
9446	Transporter	Used by Director and Transporters to communicate with the Transporter. Must be opened on the Transporter machine.
9448 - 10000	Transporter	Used by Director and Transporters for cross-Transporter data transfer. Must be opened on the Transporter machine.
VMware		

443	vCenter Server, ESXi host	Used by Director and Transporters to access VMware infrastructure. Must be opened on vCenter Servers and ESXi hosts.
902	ESXi host	Used by Transporters to access VMware infrastructure and by Journal Service (Real-time Replication) to access the journal disks via LAN transport mode. Must be opened on ESXi hosts.
3260	Transporter, ESXi host	Used by Proxy Transporters to access VMware infrastructure during a Flash boot. Must be opened on the Transporter machine and the ESXi host used as the target for a Flash boot.
Hyper-V		
137 - 139	Hyper-V hosts	Used by Director to upload files and install configuration service. Must be opened on Hyper-V servers.
445	Hyper-V hosts	Used by Director to upload files and install configuration service.
5986 (opens automatically)	Hyper-V hosts	Used by Transporter to add a host to inventory and establish a connection with it.
9445 (opens automatically)	Hyper-V hosts	Used by Director to upload files and install configuration service. Must be opened on Hyper-V host if NAKIVO Backup & Replication is installed on a host and this host is added to inventory simultaneously.
9446 (opens automatically)	Hyper-V hosts	Used by Director and Transporters to communicate with the Transporter. Used by Transporters for cross-Transporter data transfer. Must be opened on the Transporter machine.

9448 -10000 (opens automatically)	Hyper-V hosts	Used by Transporters for cross-Transporter data transfer. Must be opened on the Transporter machine.
Physical machine (Windows)		
445	Windows machine	Used by Director to upload files and install configuration service via SMB.
9446 (opens automatically)	Windows machine	Used to create the Transporter installed by default.
9448 -10000 (opens automatically)	Windows machine	Used by Transporters for cross-Transporter data transfer. Must be opened on the Transporter machine.
Physical machine (Linux)		
22	Linux machine	Used by Director to access a Linux physical machine via SSH.
9446	Linux machine	Used to create the Transporter installed by default.
9448 -10000	Linux machine	Used by Transporters for cross-Transporter data transfer. Must be opened on the Transporter machine.
Real-time replication		
33088	ESXi host	Used by the I/O Filter (source daemon installed on ESXi host) for receiving requests (or commands) from the Director.
33089	Transporter, ESXi host	Used by the Journal Service (installed on the target Transporter) for receiving requests (or commands) from the Director and the I/O Filter. (Outgoing ports on the I/O Filter).

33090	Transporter, ESXi host	Used by the Journal Service (installed on the target Transporter) for receiving changed data from the I/O Filter. (Outgoing ports on the I/O Filter).
33091	Transporter, ESXi host	Used by the Journal Service (installed on the target Transporter) for receiving heartbeat requests from the I/O Filter. (Outgoing ports on the I/O Filter).
Nutanix AHV		
9440	Nutanix REST APIs	Used to manage traffic to and from Prism Web Console and SSH console for nCLI. Prism is the central management interface for Nutanix clusters that offers an end-to-end management solution for virtualized datacenter environments that streamlines and automates common workflows.
Proxmox VE		
8006	Proxmox VE hosts	Proxmox VE API
22	Proxmox VE hosts	Used by Director to manage (install/uninstall/upgrade) Transporter and check temp folder's existence via SSH.
9446	Proxmox VE hosts	Used by Director and Transporters to communicate with the Transporter.
9448-10000	Proxmox VE hosts	Used by Transporters for cross-Transporter data transfer.

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

External Resources

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

https://www.nakivo.com/nbr_update_check-/ga/ https://receptor.nakivo.com:443/ https://collector.nakivo.com:443/ https://upload.nakivo.com/	Used for updating, checking, and supporting bundle upload.
https://nbr-va.s3.us-west-2.amazonaws.com/	Used for downloading VA.
https://d111xps0uy3x00.cloudfront.net/	Used for downloading installers.

Network Conditions

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

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

Web Browsers

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

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

Supported ASUSTOR NAS Devices

NAKIVO Backup & Replication supports the following ASUSTOR NAS devices :

Director and Onboard Transporter

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

- AS-606T
- AS-608T
- AS6508T
- AS6510T
- AS7110T
- AS6602T
- AS6604T
- AS7116RDX
- AS7112RDX
- AS1102T
- AS1104T
- AS3302T
- AS3304T
- AS6504RD
- AS6504RS
- AS6512RD
- AS5402T
- AS5404T
- FS6706T
- FS6712X
- AS1102T Lite
- AS1104T Gen 2
- AS3304T Gen 2
- AS6702T
- AS6704T
- AS3302T v2
- AS3304T v2
- AS5004U
- AS6706T
- AS7212RDX
- AS7216RDX
- AS6702T v2
- AS6704T v2
- AS6706T v2
- AS7224RDX

For minimum hardware requirements, refer to “Network Attached Storage” on page 63.

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
- RN3130
- RN31342E
- RN3138
- RN3220
- RR2312
- RR3312
- RN4220
- RR4312X
- RR4312S
- RR4360X

- RR4360S

For minimum hardware requirements, refer to “Network Attached Storage” on page 63

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-432XU
- TS-432XU-RP
- TS-451
- TS-451+
- TS-451A
- IS-400 Pro
- IS-453S
- TBS-453A
- TBS-453DX
- TS-128A
- TS-253 Pro
- TS-253A
- TS-253B
- TS-228A
- TS-451U
- TS-453 mini
- TS-453 Pro
- TS-453A
- TS-453B
- TS-453Be
- TS-453Bmini
- TS-453BT3
- TS-453BU
- TS-453BU-RP
- TS-453U
- TS-453U-RP
- TS-463U

- TS-463U-RP
- TS-463XU
- TS-463XU-RP
- TS-473
- TS-473A
- TS-563
- TS-653 Pro
- TS-653A
- TS-653B
- TS-653B
- TS-673
- TS-673A
- TS-677
- TS-832X
- TS-832XU
- TS-832XU-RP
- TS-853 Pro
- TS-853A
- TS-853BU
- TS-853BU-RP
- TS-853U
- TS-853U-RP
- TS-863U
- TS-863U-RP
- TS-863XU
- TS-863XU-RP
- TS-873
- TS-873U
- TS-873U-RP
- TS-877
- TS-877XU
- TS-877XU-RP
- TS-883XU
- TS-883XU-RP
- TS-932X
- TS-963X
- TS-977XU
- TS-977XU-RP
- TS-983XU
- TS-983XU-RP
- TS-1232XU
- TS-1232XU-RP
- TS-1253BU

- TS-1253BU-RP
- TS-1253U
- TS-1253U-RP
- TS-1263U-RP
- TS-1263U
- TS-1263XU
- TS-1263XU-RP
- TS-1273U
- TS-1273U-RP
- TS-1277
- TS-1277XU-RP
- TS-1283XU-RP
- TS-1635AX
- TS-1673U
- TS-1673U-RP
- TS-1677X
- TS-1677XU-RP
- TS-1683XU-RP
- TS-1685
- TS-2477XU-RP
- TS-2483XU-RP
- TVS-463
- TVS-471
- TVS-472XT
- TVS-473e
- TVS-473
- TVS-663
- TVS-671
- TVS-672XT
- TVS-673
- TVS-673e
- TVS-682
- TVS-682T
- TVS-863
- TVS-863+
- TVS-871
- TVS-871T
- TVS-871U-RP
- TVS-872XT
- TVS-872XU
- TVS-872XU-RP
- TVS-873e
- TVS-873

- TVS-882
- TVS-882T
- TVS-882ST2
- TVS-882BR
- TVS-882BRT3
- TVS-882ST3
- TVS-951X
- TVS-972XU
- TVS-972XU-RP
- TVS-1271U-RP
- TVS-1272XU-RP
- TVS-1282
- TVS-1282T
- TVS-1282T3
- TVS-1582TU
- TVS-1672XU-RP
- TVS-2472XU-RP
- SS-EC1279U-SAS-RP
- SS-EC1879U-SAS-RP
- SS-EC2479U-SAS-RP
- TDS-16489U
- TES-3085U
- TES-1885U
- TS-EC880U
- TS-EC880U R2
- TS-EC1280U
- TS-EC1280U R2
- TS-EC1680U
- TS-EC1680U R2
- TS-EC2480U
- TS-EC2480U R2
- TVS-EC880
- TVS-EC1080
- TVS-EC1080+
- TVS-EC1280U-SAS-RP
- TVS-EC1580MU-SAS-RP
- TVS-EC1680U-SAS-RP
- TVS-EC1680U-SAS-RP R2
- TVS-EC2480U-SAS-RP
- TVS-EC2480U-SAS-RP R2
- TVS-EC2480U-SAS-RP R2
- TVS-EC1580MU-SAS-RP R2
- TVS-EC1280U-SAS-RP R2

- TDS-16489U-SE1-R2
- TDS-16489U-SE2-R2
- TDS-16489U-SF2-R2
- TDS-16489U-SF3-R2
- TS-2888X-W2195-512G
- TS-2888X-W2195-256G
- TS-2888X-W2195-128G
- TS-2888X-W2175-512G
- TS-2888X-W2175-256G
- TS-2888X-W2175-128G
- TS-2888X-W2145-512G
- TS-2888X-W2145-256G
- TS-2888X-W2145-128G
- TS-2888X-W2133-64G
- TS-2888X-W2123-32G
- ES2486dc
- TS-1886XU-RP
- TS-230
- TS-251C
- TS-251D
- TS-253D
- TS-451DeU
- TS-453D
- TS-653B
- TS-653D
- TS-h1277XU-RP
- TS-h1283XU-RP
- TS-h977XU-RP
- TVS-472XT-PT
- TVS-672N
- TVS-872N
- TVS-EC2480U-SAS-RP-R2
- TS-h686-D1602
- TS-h886-D1622
- TS-873AU
- TS-873AU-RP
- TS-1273AU-RP
- TS-1673AU-RP
- TS-932PX
- GM-1001
- TS-432PXU
- TS-432PXU-RP
- TS-832PXU

- TS-832PXU-RP
- TS-1232PXU-RP
- TS-451D2
- TS-h2490FU-7232P-64G
- TS-h2490FU-7302P-128G
- TS-h1886XU-RP
- TS-h1683XU-RP
- TS-h2483XU-RP
- TVS-h1288X
- TVS-h1688X
- TS-h973AX-8G
- TS-h973AX-32G
- TS-832PX
- TS-h3088XU-RP-W1270-64G
- TS-h3088XU-RP-W1250-32G
- TS-453DU-4G
- TS-873A
- TS-EC879U-RP
- TS-EC879U-RP
- TS-h987XU-RP
- TS-h3077AFU
- TS-h1277AXU-RP
- TS-h1677AXU-RP
- TVS-h674T
- TVS-h874T
- TBS-574TX
- TS-AI642
- TS-855X
- TS-879U-RP
- TS-262
- TS-462
- TS-h1887XU-RP
- TS-1655
- TS-855eU
- TS-473A
- TBS-h574TX
- TDS-h2489FU-4309Y
- TDS-h2489FU-4314
- TS-130
- TS-233
- TS-253E
- TS-264
- TS-364

- TS-410E
- TS-433
- TS-453E
- TS-464
- TS-464eU
- TS-464U-4G
- TS-464U-RP-4G
- TS-473A-8G
- TS-664
- TS-864eU
- TS-864eU-RP
- TS-h1090FU
- TS-h1290FX
- TS-h1677XU-RP
- TS-h2477XU-RP
- TS-hx87XU-RP
- TS-i410X
- TVS-672X-i3-8G
- TVS-675
- TVS-872X-i3-8G
- TVS-h474
- TVS-h674
- TVS-h874
- TS-216G
- TS-432X
- TS-632X
- TS-433eU
- TS-h2287XU-RP
- TS-h765eU
- TS-h1277AFX
- TS-1264U-RP
- TDS-h2489FU
- TS-873AeU-RP-4G
- TS-435XeU-4G
- ES2486dc
- ES1686dc R2
- TS-h2287XU-RP
- TS-h2490FU
- TS-433-4G
- TS-h3087XU-RP
- TVS-AIh1688ATX
- TS-h2477AXU-RP
- TS-h1077AFU

- TS-h1655XeU-RP
- TS-855eU-RP

For minimum hardware requirements, refer to “Network Attached Storage” on page 63.

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+
- RS815RP+
- RS815+
- RS814RP+

- RS814+
- RS812RP+
- RS812+
- RS810RP+
- RS810+
- RC18015xs+
- DS3617xs
- DS3615xs
- DS3612xs
- DS3611xs
- DS3018xs
- DS2415+
- DS2413+
- DS2411+
- DS1819+
- DS1817+
- DS1815+
- DS1813+
- DS1812+
- DS1618+
- DS1517+
- DS1515+
- DS1513+
- DS1512+
- DS1511+
- DS918+
- DS916+
- DS718+
- DS716+II
- DS716+
- DS713+
- DS712+
- DS710+
- DS418
- DS418play
- DS418j
- DS416play
- DS415+
- DS412+
- DS411+II
- DS411+
- DS218+
- DS218

- DS218play
- DS216+II
- DS216+
- DS118
- DS1019+
- DS2419+
- DS420+
- DS420j
- DS620slim
- DS720+
- DS920+
- DS925+
- FS3400
- FS3600
- FS6400
- RS819
- RS820+
- RS820RP+
- SA3200D
- SA3400
- SA3600
- DS1520+
- DS1621+
- DS1621xs+
- DS1821+
- RS1221+
- RS1221RP+
- RS2421+
- RS2421RP+
- RS4021xs+
- RS2423+
- RS2423RP+
- DS124
- DS224+
- DS223j
- DS723+
- DS273+
- DS223
- DS1823xs+
- DS1821Plus
- DS423+
- DS423
- SA3400D

- SA3610
- SA3410
- DS120j
- DS1522
- DS1522+
- DS220j
- DS2422+
- DS3621xs+
- DS3622xs+
- DS923+
- FS2500
- FS3410
- RS2821RP+
- RS3621RPxs
- RS3621xs+
- RS422+
- RS822+
- RS822RP+
- DS1525+
- DS1825+
- DS725+
- DS425+
- RS2825RP+
- DS225+

Transporter Only

- DS220j

Important

For minimum hardware requirements, refer to “Network Attached Storage” on page 63.

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 63.

Feature Requirements

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

- Auto-Update
- Backup Encryption
- Backup Malware Scan
- External Database
- Federated Repository
- Granular User Notifications
- Merge Jobs
- MSP Console
- Object Recovery and Log Truncation for Microsoft Exchange
- Object Recovery and Log Truncation for Microsoft SQL Server
- Object Recovery for Microsoft Active Directory
- Proxmox VE Support
- Transporter Load Balancing
- VM Limitation for Multi-Tenancy

Auto-Update

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

- Linux
- Windows

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

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

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

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

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

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

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

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

> **Note**
> To maintain maximum security on your Linux OS, always keep it up-to-date; however, if necessary, you may disable auto-updates by editing the `/etc/apt/apt.conf.d/20auto-upgrades` file.
> For more details, refer to https://linuxconfig.org/disable-automatic-updates-on-ubuntu-22-04-jammy-jellyfish-linux.

Backup Encryption

With the **Backup encryption** feature, you can customize your protection preferences by encrypting backup data created by the product, managing passwords, and performing recovery from encrypted backups.
To use the **Backup encryption** feature, make sure you meet the following requirements:
Supported Data Types:

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

Supported Job Types:

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

Supported Destinations:

The following repository types are supported for storing encrypted backups:

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

> **Note**
>
> SaaS cannot be used for storing encrypted backups.

- Federated repositories
- Tape
    - Supported for the data copied to tape media
    - Not supported if hardware encryption is enabled on a target tape device

Supported Recovery Jobs:

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

Supported Recovery Sessions:

The following recovery sessions are supported to recover encrypted backups:

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

Supported Key Management Service:

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

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

> **Note**
>
> The list of required AWS permissions for creating the (AWS) KMS Keys:
>
> - *kms:ListKeys*
> - *kms:ListAliases*
> - *kms:CreateKey*
> - *kms:CreateAlias*
> - *kms:PutKeyPolicy*
> - *kms:TagResource*
> - *tag:GetResources*
> - *iam:CreateServiceLinkedRole*
> - *iam:ListUsers*
> - *iam:ListRoles*
> - *iam:GetUser*
>
> Refer to Creating keys - AWS Key Management Service for more details.

General Requirements:

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

Supported Platforms:

Encryption across all platforms is supported.

Immutability:

Immutability can be applied to encrypted backups.

Backup Malware Scan

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

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

The scan server must adhere to the following requirements:

- A supported version of antivirus software must be installed on the scan server.
- iSCSI must be available and active on the scan server.
- SMB must be enabled on the scan servers using Windows OS.
- SSH must be enabled on the scan servers using Linux OS.
- TCP port *9445* must not be blocked.
- SElinux service must be disabled on the scan servers using Linux OS.
- The NTFS-3G package must be enabled on the scan servers using Linux OS.
- The LVM2 package must be enabled and updated to the latest version on the scan servers using Linux OS.
- Special permissions must be configured for NAKIVO Backup & Replication recovery service.

> **Note**
> Adding more than 500 scan servers is not recommended as it may lead to unexpected problems and performance issues.

The following antivirus software is supported:

- **Windows OS**:
  - Microsoft Windows Defender
    - Antimalware Client Version: 4.10.14393 or higher
    - Engine Version: 1.1.12805 or higher
    - Must support the command line: *Scan -ScanType %type% -File %path% -DisableRemediation -BootSectorScan*
    - See more details here
  - ESET NOD32 Antivirus:
    - Version 14.2.24 or higher
    - Must support the command line: *%path% /clean-mode=None /no-symlink*
    - See more details here
  - Kaspersky Internet Security:
    - Version 2015 (15.0) or higher
    - Must support the command line: *scan %path% -i0*
    - See more details here
  - Sophos Intercept X:
    - Version 2.0.20 or higher
    - Must support the command line: *scan --noui --expand_archives %path%*
    - See more details here
- **Linux OS**:
  - Clam AV:
    - Version clamav-0.99.0 or higher
    - Must support the command line: *clamscan --infected --recursive %path%*
    - See more details here
  - Sophos Protection:
    - Version 1.1.8 or higher
    - Must support the command line: *avscanner %path% --scan-archives*
    - See more details here

Notes

- Sophos Protection can be used only for Linux server recovery.
- Sophos AV software is limited to the specific Linux operating systems.
- Antivirus software on Windows OS cannot scan Linux OS backups.
- Scanning process may not detect malware if the antivirus software has the run-time (real-time) protection enabled.

Important

NAKIVO Backup & Replication consistently mounts the recovery point and provides full access to its data for malware scanning. The scanning process, file analysis, and detection logic are fully performed by the antivirus engine installed on the Scan Server. Scan behavior and results may vary depending on how the antivirus engine reads and processes the mounted data under current system conditions.

Any limitations related to the antivirus engine, such as skipped files, inconsistent scan results, or filename/path reporting issues (including Unicode handling), originate from the antivirus engine and are not controlled by NAKIVO Backup & Replication.

External Database

The following external databases are supported:

- PostgreSQL v10-17x

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

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

Notes

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

Federated Repository

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

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

Notes

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

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

Notes

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

Supported Backup Location Policy

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

Supported Functionality

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

Supported Federated Repository Member Types

- Local Folder
- CIFS Share
- NFS Share

Supported Repository Attributes

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

File System Indexing

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

File System Indexing indexes only files and folders located on disks that are included in the backup job, except VM disks using a controller other than iSCSI.

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

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

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

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

- VMware Tools or Hyper-V Integration Services must be running on VMware or Hyper-V VMs correspondingly.
- *iSCSI* must be available.
- *SMB* must be enabled on Windows OS.
- *SSH* must be enabled on Linux OS.
- TCP port `9445` must not be blocked.
- *SElinux* service must be disabled on Linux OS.
- Special permissions must be configured for Linux OS.
- Special permissions must be configured for Windows OS.

- To run **File System Indexing** on RHEL 7.9 and RHEL 7.4, the **Netstat** (network statistics) command line tool must be installed on the source VM.
- The default port for SSH is 22.

Granular User Notifications

The feature supports the following jobs:

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

Merge Jobs

The feature supports the following types of jobs:

- Backup
- Backup copy
- Replication

Job merging can be performed in the following cases:

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

Job merging cannot be performed in the following cases:

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

MSP Console

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

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

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

Notes

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

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 Server SE
- Microsoft Exchange 2019
- Microsoft Exchange 2016
- Microsoft Exchange 2013

Permissions

The following requirements should be met for log truncation:

- Selected users should have permissions to "Log on as a batch job".
- Active Directory Module For Windows PowerShell must be installed.

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

Services and Settings

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

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

Object Recovery and Log Truncation for Microsoft SQL Server

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

General Requirements

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

Supported Versions of Microsoft SQL Server

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

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

Supported Editions

NAKIVO Backup & Replication supports the following editions of Microsoft SQL Server:

- Enterprise
- Standard
- Web
- Developer
- Express edition

Permissions

- A user logging in to Microsoft SQL Server must have a `sysadmin` role.
- The user running Microsoft SQL Service should have permissions to "Log on as a batch job".

Services and Settings

- NAKIVO Backup & Replication requires PowerShell v2 or later.
- VMware Tools or Hyper-V Integration Services must be running on VMware or Hyper-V VMs correspondingly.
- `sqlcmd` utility must be installed on the machine running Microsoft SQL server.
- Ports *137-139* must be opened for `cifs`.
- The SMB v1 or higher version of SMB protocol should be enabled.

Requirements for Microsoft SQL Server Object Recovery

- The user running Microsoft SQL service must have executive permissions to the `Data` folder and all other folders in which the databases are located.
- If "Rename recovered item if such item exists" option is selected during the recovery, NAKIVO Backup & Replication skips keys, constraints, indexes, and statistical properties when recovering a table to an original location.
- If "Overwrite the original item if such item exists" option is chosen, all the above properties are preserved. Tables that contain a foreign key cannot be recovered with this option.
- Full administrative permissions are required.
- Default administrative shares must be enabled.
- The "File server" role must be enabled.
- Ports *445* and 9*445* must be opened on the instance.

Requirements for Microsoft SQL Server Log Truncation

- VMware VM must be running on VMware ESXi 5.0 and later.
- System databases are skipped during the log truncation.
- Databases with the "Simple" recovery model are skipped during the log truncation.

- A database must be in the "online" state.
- The SMB v1 or higher version of SMB protocol should be enabled.

Object Recovery for Microsoft Active Directory

Supported Versions

NAKIVO Backup & Replication supports the following versions of Microsoft Active Directory for objects recovery:

- Windows Server 2022 (21H2)
- Windows Server 20H2 (20H2)
- Windows Server 2019
- Windows Server 2016

Requirements for Object Recovery for Microsoft Active Directory

- The ISCI Initiator service must be running on the recovery server.
- The vc_redist.x86.exe (v.2015) file must be installed on the recovery server.
- Active Directory Web Services must be running.
- Port *5000* must not be blocked by other services and must be opened in the firewall of AD.
- Active Directory Module For Windows PowerShell must be installed.

Proxmox VE Support

Proxmox Virtual Environments (VE) are supported on all supported browsers for NAKIVO Backup & Replication and all supported operating systems.

To successfully discover Proxmox VE items (clusters and standalone hosts (nodes)), add them to the **Inventory**, back up Proxmox VE VMs, create backup copies, and perform full and file level recovery, make sure you meet the following requirements:

- **Proxmox VE** are supported for the following Proxmox VE versions:
  - **Proxmox VE 8.x**. For more details, refer to Proxmox VE versions table.

Notes

- The Proxmox VE (host/cluster) of the version lower than the oldest supported version cannot be discovered.
- The Proxmox VE (host/cluster) of the version higher than the newest supported version can be discovered, but confirmation of the operation is required to proceed.
- The product can discover all Proxmox VE storage types (See Proxmox VE Storage Types), except the following:
  - PBS
  - ESXi

- A transporter must be available inside the Proxmox VE host for running backup and recovery jobs.

Notes

- For the Proxmox VE cluster, the transporter is injected into each cluster host.
- Transporters injected into Proxmox VE hosts are supported for auto-update.

Supported Functionality

- Backup of Proxmox VE VMs to Incremental with Full and Forever Incremental repositories
  - Immutability
  - Change Block Tracking (CBT)
    - Proprietary method
    - Native change tracking
  - Application awareness

Notes

App-aware mode on Proxmox VE VM requires the following to function:

- QEMU Guest Agent must be enabled and running on the VM. You can test the communication qm agent by running the following command. If the QEMU Guest Agent runs correctly in the VM, it will return without an error message.

```
qm agent <vmid> ping
```

- Virtio-serial driver must be installed on the Windows based VM. For Linux based VMs, you only need to run the install command.
  - For Debian/Ubuntu: *apt-get install qemu-guest-agent*
- For FreeBSD-based VMs, create the script files in the source VM usr/sbin/ folder:
  - "pre-freeze-script"
  - "post-thaw-script"

  Make sure they have sufficient permissions.

    - Backup encryption
    - VM verification
    - Log truncation for Proxmox VE backups
- Backup copy of Proxmox VE VM can be saved to the following types of storage:
    - Repositories/stream repositories with creating immutable copies for Proxmox VE VM backups
    - Data storage types:
        - Forever Incremental
        - Incremental with periodic full
    - Repository types:
        - Local Folder
        - CIFS Share
        - NFS Share
        - Cloud & S3-Compatible Storages
            - Amazon S3
            - Wasabi
            - Backblaze B2 Cloud Storage
            - Microsoft Azure Blob Storage
            - Amazon EC2
            - Generic S3-compatible Storage

    - Deduplication Appliances
        - Dell EMC Data Domain Boost
        - HPE StoreOnce Catalyst
        - NEC HYDRAstor
- Backup copy to tape
- Replication of Proxmox VE VMs
- Flash VM Boot from Proxmox VM backups
- Full recovery of Proxmox VE VM backups
- File recovery from Proxmox VE VM backups
- Direct recovery of Proxmox VE VM backups from tape
- Object recovery from Proxmox VE VM backups
- Universal object recovery for Proxmox VE VM backups
- Backup export for Proxmox VE VM backups
- Granular notifications for Proxmox VE backup jobs
- Backup, recovery, and replication is also supported for Proxmox VMs located on a Ceph erasure-coded cluster as hyper-converged storage.

Feature Limitations

Native Change Tracking

- If the VM image is stored in QEMU format, the dirty bit map (CBT) is persistent, and VM restart does not affect CBT (we can continue conducting incremental backups).
- If it is stored in RAW format, a non-persistent dirty bit map is used, and VM restart leads to a dirty bitmap reset and the next backup should be full.

Backups can be performed on disks mapped to ISCSI storage LUNs but they cannot be used as target storage for recovery due to the following reasons:

- When an ISCSI LUN is mapped to multiple disks or VMs, the recovery process writes data to that LUN, which affects all other VMs on that LUN.
- There is a requirement for LUNs to be specified for every recovery disk, which is not supported by the current UI.
- NAKIVO Backup & Replication uses iSCSI storage for Flash VM Boot jobs. For this reason, VM disks cannot be mapped to storages and can only be either skipped or recovered.

Proxmox VE Limitations

- Proxmox VE VM network devices cannot be connected directly to a physical network; they must be connected to a bridge.
- Freeze and thaw hooks in the QEMU guest agent are only supported on Linux VMs. Custom quiescing scripts for Proxmox VE are not supported on Windows VMs.
- Proxmox VE storages must contain **Disk Image** in their **Content** field to store VM disks.

- To successfully perform backups, the number of free nbd devices on a Proxmox VE host must be more than or equal to the number of VM disks to back up.
- For successful Proxmox VE VM replication, the target storage must support snapshot, and source disk format will be converted to “qcow2”.
- The Proxmox's log truncation function uses "QEMU Agent" to execute the script with "local system" account (NT AUTHORITY\SYSTEM):
    - For SQL Server, it requires using an account having "sysadmin" role.
    - If the account does not have "sysadmin" role by default, you cannot execute some SQL commands (Eg: BACKUP DATABASE).
    - It is also not possible to switch to an account having "sysadmin" role when executing a command.
    - For the log truncation to work, do the following:
        - Re-config "QEMU Agent" service to run using an account having SQL's "sysadmin" role, and restart the service.
        - Re-config SQL's "NT AUTHORITY\SYSTEM" account: Assign "sysadmin" role.
- Do not manually unlock or reboot a Proxmox VM from the Proxmox UI while a backup job is running. Doing so may cause temporary snapshot files to remain after the successful job run.
- A target network must always be selected for all VMs in a recovery or replication job. Otherwise, the NIC will not be created on the target VM.

Transporter Load Balancing

The **Transporter load balancing** feature ensures the even distribution of tasks among available transporters while processing backup/replication/recovery jobs.

The following platforms are supported:

- VMware vSphere
- Nutanix AHV
- Amazon EC2

Requirements

To use the **Transporter load balancing** feature, ensure that the product is set to determine the transporter automatically for at least one VM/instance. To verify, go to the backup job **Options** tab and check that the **Data Transfer** > **Transporters** option is set to **Automatic selection**.

Feature Limitations

- Direct Connect (proxy transporter)
    - For items that were added using the Direct Connect feature, the product uses the transporter initially used to add an item.

VM Limitation for Multi-Tenancy

The feature is only available if a license with Socket limit mode is installed.

The following hypervisors are supported:

- VMware vSphere
- VMware Cloud Director
- Microsoft Hyper-V
- Nutanix AHV

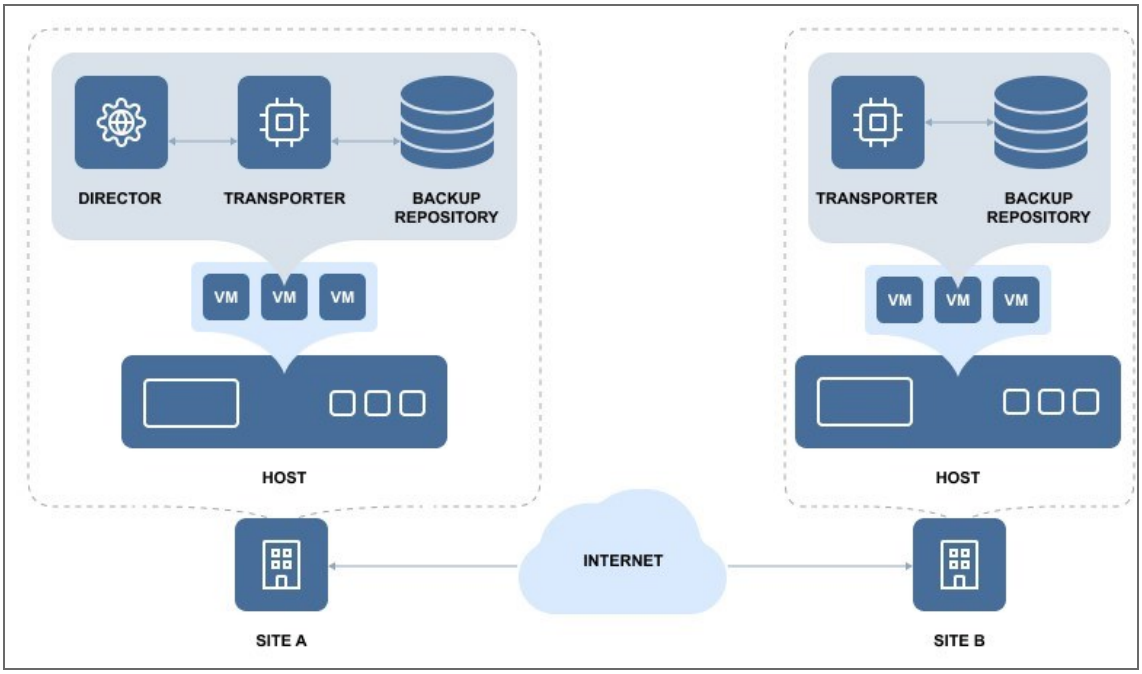
Deployment Scenarios

NAKIVO Backup & Replication is a modular solution that can be fully installed on a single machine to protect small and mid-sized environments, as well as scale out horizontally and support large distributed environments. Refer to the sections below to learn more about the product deployment scenarios.

- “Distributed Deployment” on page 112
- “Multi-Tenant Deployment” on page 113
- “Single Site Deployment” on page 124

Distributed Deployment

If you have multiple sites and need to back up and/or replicate over WAN, install the Director and Transporter on one site, and at least one Transporter on all other sites.



Note

Make sure the required ports are open on the appropriate endpoints. The full list of required ports can be found in Deployment Requirements.

Multi-Tenant Deployment

Installation of a multi-tenant solution of NAKIVO Backup & Replication allows you to create multiple isolated tenants within a single product deployment and manage them from a single pane of glass. In the Multi-Tenant mode, tenants can access the self-service portal to offload backup, and recovery tasks from the service provider.

Backup from Remote Site to Master Site

- Deployment Scenario
- Deployment Steps
- Connections

Deployment Scenario

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

> **Example**
> A service provider needs to back up customers' VMs to the service provider's datacenter so that the customers don't see each other's backups and can recover their own files and emails through a self-service interface.

Deployment Steps

To deploy the above scenario, perform the following steps:

1. Install the Director in multi-tenant mode at the master site.
2. Install at least one Transporter at the master site.
3. Install at least one Transporter at each remote site.
4. For each tenant, prepare a separate folder at the master site for creating separate Backup Repositories.

Connections

The implementation of the above scenario requires that the following connections be available:

Connection	Description
A	Connection from the machine on which the Director is installed to vCenter servers and ESXi hosts that run source VMs at remote sites. The port used for communication with vCenter servers and ESXi hosts (443 by default) is open in firewalls.
B	Connection from the machine on which the Director is installed to machines at remote sites on which Transporters are installed. The port used for communication with the Transporters (9446 by default) is open in firewalls.
C	Connection from the machine on which the Director is installed to vCenter servers and ESXi hosts at the master site where VM replicas will be created. The port used for communication with vCenter servers and ESXi hosts (443 by default) is open in firewalls.
D	Connection from the machine on which the Director is installed to the machine at the master site where the Transporter is installed. The port used for communication with the Transporters (9446 by default) is open in firewalls.

E	Connection from the machine at the Master site where the Transporter is installed to ESXi hosts at the master site where VM replicas will be created.
F	Connection from the machine at the Master site where the Transporter is installed to machines at remote sites where Transporters are installed. The ports used for data transfer between a pair of Transporters are open in firewalls.
G	At remote sites, connections from machines on which Transporters are installed to vCenter servers and ESXi hosts running source VMs.

Note

For security purposes, a VPN connection should be established between the master site and remote sites.

Replication from Remote Site to Master Site

- Deployment Scenario
- Deployment Steps
- Connections

Deployment Scenario

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

> **Example**
> A service provider wants to introduce Replication-as-a-Service to customers and replicate their VMs to the service provider's datacenter.

Deployment Steps

To deploy the above scenario, perform the following steps:

1. Install the Director in multi-tenant mode at the master site.
2. Install at least one Transporter at the master site.
3. Install at least one Transporter at each remote site.
4. For each tenant, prepare a separate ESXi host that will serve as a replication target.

Connections

The implementation of the above scenario requires that the following connections be available:

Connection	Description
A	Connection from the machine on which the Director is installed to vCenter servers and ESXi hosts that run source VMs at remote sites. The port used for communication with vCenter servers and ESXi hosts (443 by default) is open in firewalls.
B	Connection from the machine on which the Director is installed to machines at remote sites on which Transporters are installed. The port used for communication with the Transporters (9446 by default) is open in firewalls.
C	Connection from the machine on which the Director is installed to vCenter servers and ESXi hosts at the master site where VM replicas will be created. The port used for communication with vCenter servers and ESXi hosts (443 by default) is open in firewalls.
D	Connection from the machine on which the Director is installed to the machine at the master site where the Transporter is installed. The port used for communication with the Transporters (9446 by default) is open in firewalls.
E	Connection from the machine at the master site where the Transporter is installed to ESXi hosts at the master site where VM replicas will be created.

F	Connection from the machine at the master site where the Transporter is installed to machines at remote sites where Transporters are installed. The ports used for data transfer between a pair of Transporters are open in firewalls.
G	At remote sites, connections from machines on which Transporters are installed to vCenter servers and ESXi hosts running source VMs.

Note

For security purposes, a VPN connection should be established between the master site and remote sites.

Local Backup at Remote Site

- Deployment Scenario
- Deployment Steps
- Connections

Deployment Scenario

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

> **Example**
> An Enterprise has two branch offices running VMware virtual infrastructure. The IT manager located at the headquarters is responsible for the Enterprise data protection and needs to back up VMs locally at their branch offices to ensure fast operational recovery. Employees of the branch offices should have access to their VM backups and be able to recover their files and emails.

Deployment Steps

To deploy the above scenario, perform the following steps:

1. Install the Director in multi-tenant mode at the master site.
2. Install at least one Transporter at each remote site.
3. For each tenant, prepare a separate folder at a remote site for creating a Backup Repository.

Connections

The implementation of the above scenario requires that the following connections be available:

Connection	Description
A	Connection from the machine on which the Director is installed to vCenter servers and ESXi hosts that run source VMs at remote sites. The port used for communication with vCenter servers and ESXi hosts (443 by default) is open in firewalls.
B	Connection from the machine on which the Director is installed to machines at remote sites where the Transporters are installed. The port used for communication with the Transporters (9446 by default) is open in firewalls.
C	Connection from the machines on which the Transporters are installed at remote sites to vCenter servers and ESXi hosts running source VMs.

> **Note**
> For security purposes, a VPN connection should be established between the master site and remote sites.

Local Replication at Remote Site

- Deployment Scenario
- Deployment Steps
- Connections

Deployment Scenario

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

> **Example**
> An Enterprise has two branch offices running VMware virtual infrastructure. The IT manager located at the headquarters is responsible for the Enterprise data protection and needs to replicate business critical VMs locally at the branch offices for high availability.

Deployment Steps

To deploy the above scenario, perform the following steps:

1. Install the Director in multi-tenant mode at the master site.
2. Install at least one Transporter at each remote site.
3. For each tenant, prepare a separate folder at the remote site for creating a Backup Repository.

Connections

The implementation of the above scenario requires that the following connections be available:

Connection	Description
A	Connection from the machine on which the Director is installed to vCenter servers and ESXi hosts that run source VMs at remote sites. The port used for communication with vCenter servers and ESXi hosts (443 by default) is open in firewalls.
B	Connection from the machine on which the Director is installed to machines at remote sites where Transporters are installed. The port used for communication with the Transporters (9446 by default) is open in firewalls.
C	Connection from the machine on which the Director is installed to vCenter servers and ESXi hosts where VM replicas will be created at remote sites. The port used for communication with vCenter servers and ESXi hosts (443 by default) is open in firewalls.
D	At remote sites, connections from machines where Transporters are installed to vCenter servers and ESXi hosts running source VMs.
E	At remote sites, connections from machines where Transporters are installed to vCenter servers and ESXi hosts where VM replicas will be created.

Note

For security purposes, a VPN connection should be established between the master site and remote sites.

Backup at Master Site

- Deployment Scenario
- Deployment Steps
- Connections

Deployment Scenario

In this scenario, tenant VMs are running at the master site and the backing up of tenant VMs is also performed at the master site.

> **Example**
> A service provider runs VMs of customer A and customer B in the service provider's datacenter. The Service Provider seeks to offer Backup-as-a-Service to both customers. The customers should be able to recover their files and emails through a self-service interface without being able to see each other's backups.

Deployment Steps

To deploy the above scenario, perform the following steps:

1. Install the Director in multi-tenant mode at the master site.
2. Install at least one Transporter at the master site.
3. For each tenant, prepare a separate folder at the master site for creating a Backup Repository.

Connections

The implementation of the above scenario requires that the following connections be available:

Connection	Description
A	Connection from the machine on which the Director is installed to the machine on which the Transporter is installed. The port used for communication with the Transporters (9446 by default) is open in firewalls.
B	Connection from the machine on which the Director is installed to vCenter servers and ESXi hosts that run source VMs. The port used for communication with vCenter servers and ESXi hosts (443 by default) is open in firewalls.
C	Connection from the machine on which the Transporter is installed to vCenter servers and ESXi hosts running source VMs.
D	Connection from the machine on which the Transporter is installed to the folders where tenant Backup Repositories will be created.

Note

For security purposes, a VPN connection should be established between the master site and remote sites.

Replication at Master Site

- Deployment Scenario
- Deployment Steps
- Connections

Deployment Scenario

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

> **Example**
> A service provider runs customers' VMs in the service provider's datacenter. To ensure high availability of tenant VMs, the service provider seeks to replicate customer VMs to a different server.

Deployment Steps

To deploy the above scenario, perform the following steps:

1. Install the Director in multi-tenant mode at the master site.
2. Install at least one Transporter at the master site.
3. For each tenant, prepare a separate ESXi host that will serve as a replication target.

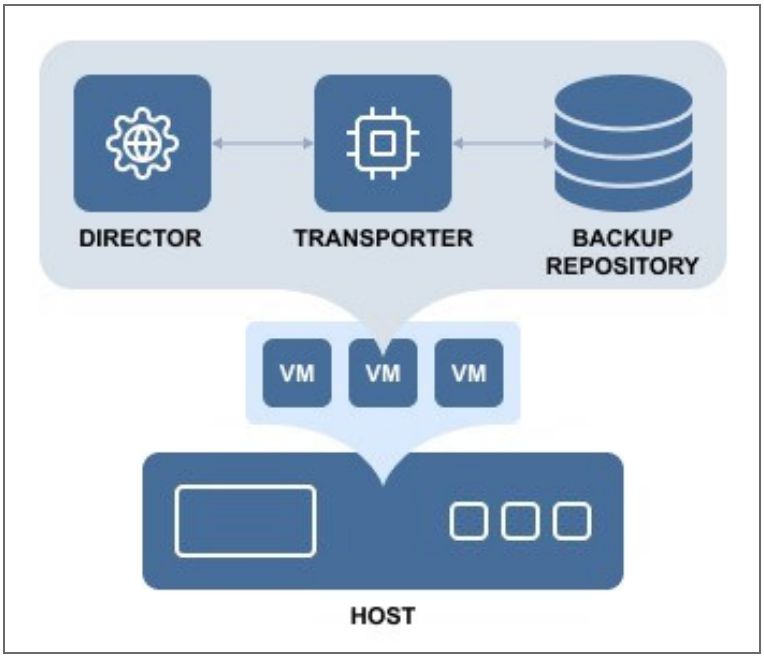
Connections

The implementation of the above scenario requires that the following connections be available:

Connection	Description
A	Connection from the machine on which the Director is installed to vCenter servers and ESXi hosts that run source VMs. The port used for communication with vCenter servers and ESXi hosts (443 by default) is open in firewalls.
B	Connection from the machine on which the Director is installed to the machine on which the Transporter is installed. The port used for communication with the Transporters (9446 by default) is open in firewalls.
C	Connection from the machine on which the Director is installed to vCenter servers and ESXi hosts where VM replicas will be created. The port used for communication with vCenter servers and ESXi hosts (443 by default) is open in firewalls.
D	Connection from the machine on which the Transporter is installed to vCenter servers and ESXi hosts running source VMs.
E	Connection from the machine on which the Transporter is installed to vCenter servers and ESXi hosts where VM replicas will be created.

Single Site Deployment

For a single site deployment, it is often sufficient to install both the Director and Transporter on a single VM/physical machine within your infrastructure.



This deployment provides you with the ability to back up, replicate, and recover multiple VMs from multiple source hosts.

Installing NAKIVO Backup & Replication

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

- "Deploying VMware Virtual Appliance" on page 126
- "Deploying Nutanix AHV Virtual Appliance" on page 132
- "Deploying Amazon Machine Image in Amazon EC2" on page 139
- "Installing on Windows" on page 141
- "Installing on Linux" on page 152
- "Installing on Synology NAS" on page 162
- "Installing on QNAP NAS" on page 169
- "Installing on Western Digital NAS" on page 174
- "Installing on ASUSTOR NAS" on page 176
- "Installing on NETGEAR ReadyNAS" on page 181

Note

Before installing your NAKIVO Backup & Replication solution, make sure that your system has been updated with the latest patch and all the necessary requirements are met.

Deploying VMware Virtual Appliance

- Deploying Virtual Appliance with vSphere Web Client
- Virtual Appliance OS, Credentials, and Security
- Web Interface Login

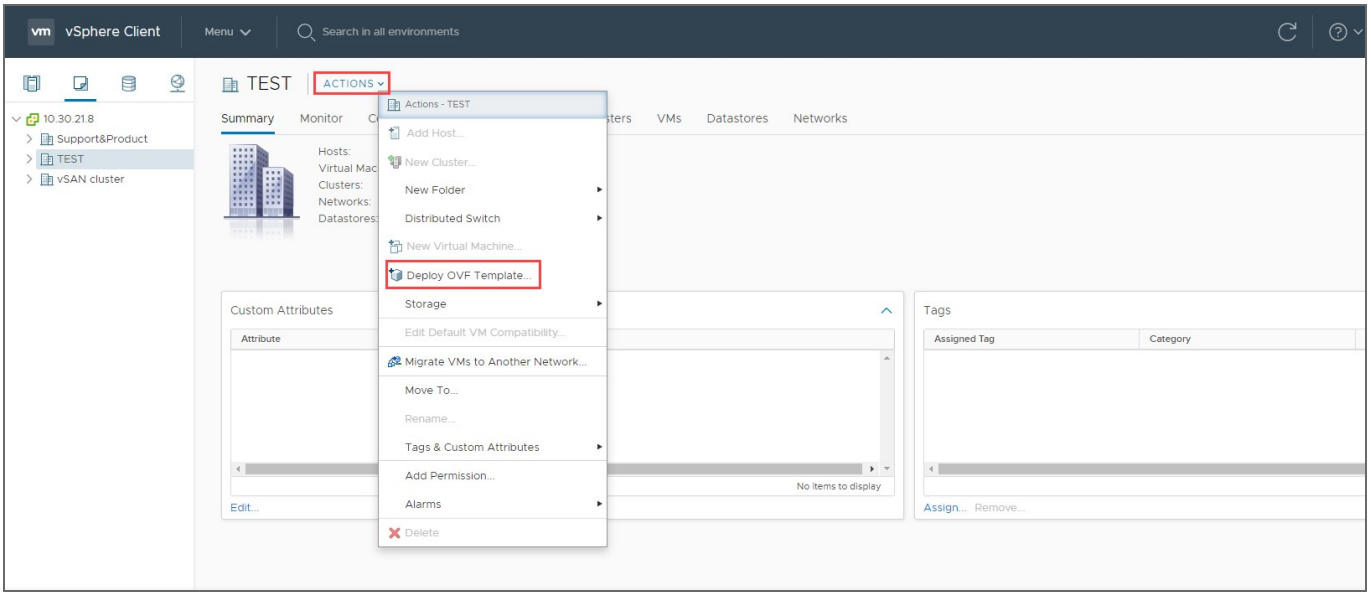
NAKIVO Backup & Replication offers the following VA deployment options:

- Full Solution
- Full Solution without **Backup Repository**
- **Transporter**-only
- **Transporter** with **Backup Repository**
- Multi-tenant **Director**

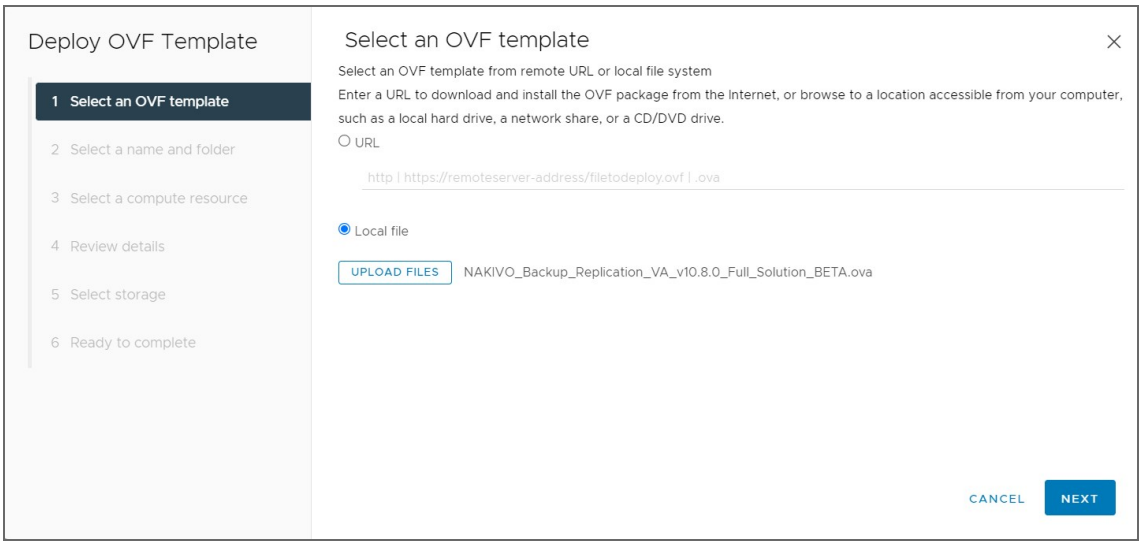
The Virtual Appliance (VA) has two disks: the first (30 GB) contains a Linux OS with NAKIVO Backup & Replication, and the second (500 GB) is used as a **Backup Repositor**y. If you deploy the Virtual Appliance disks using the **Thin Provision** option, then the disks will not reserve space on your datastore and will only consume space when actual data (such as your backups) is written to disks.

Deploying Virtual Appliance with vSphere Web Client

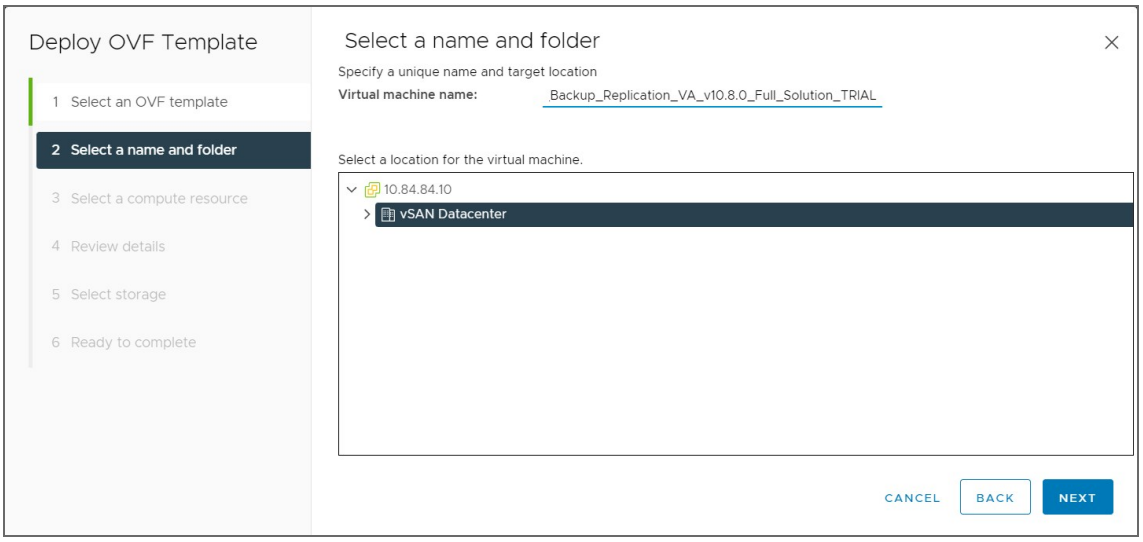
1. Download NAKIVO Backup & Replication VA.
2. Log in to your vSphere vCenter with the vSphere Web Client.
3. Select **Deploy OVF Template** from the **Actions** menu. Note that the Client Integration Plug-in must be installed to enable OVF functionality.



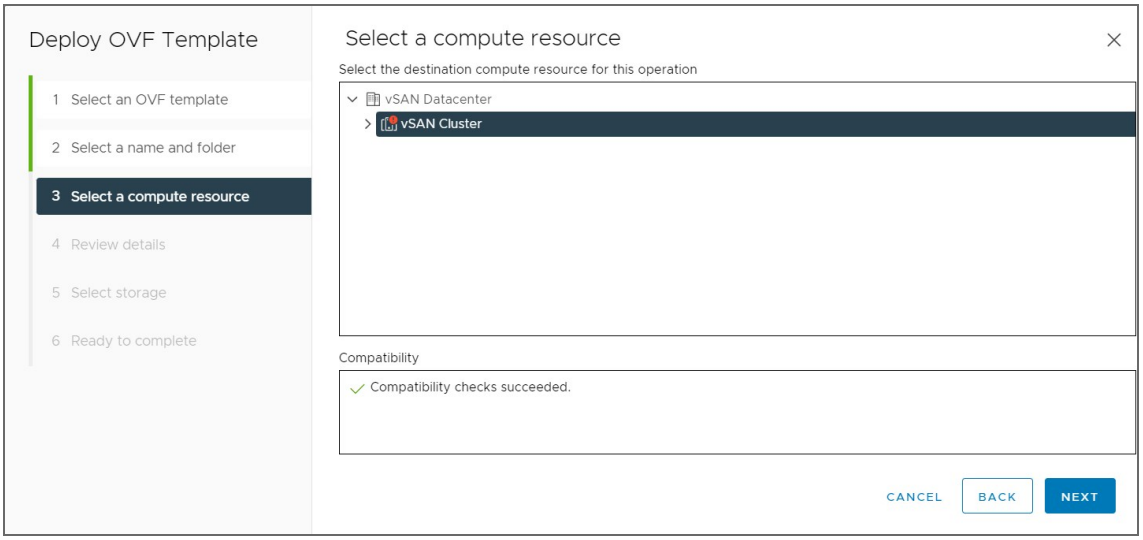
4. On the **Select an OVF template** page of the **Deploy OVF Template** wizard, select **Local file** and upload the VA file (.ova) you've downloaded. Click **Next**.



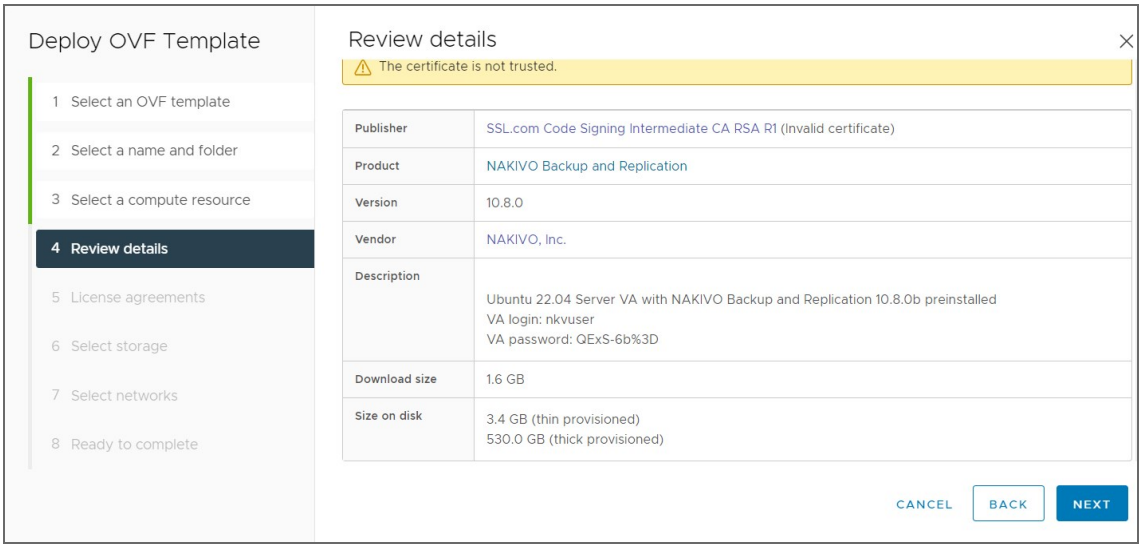
5. On the **Select a name and folder** page, specify a unique name and target location for the Virtual Appliance. Click **Next**.



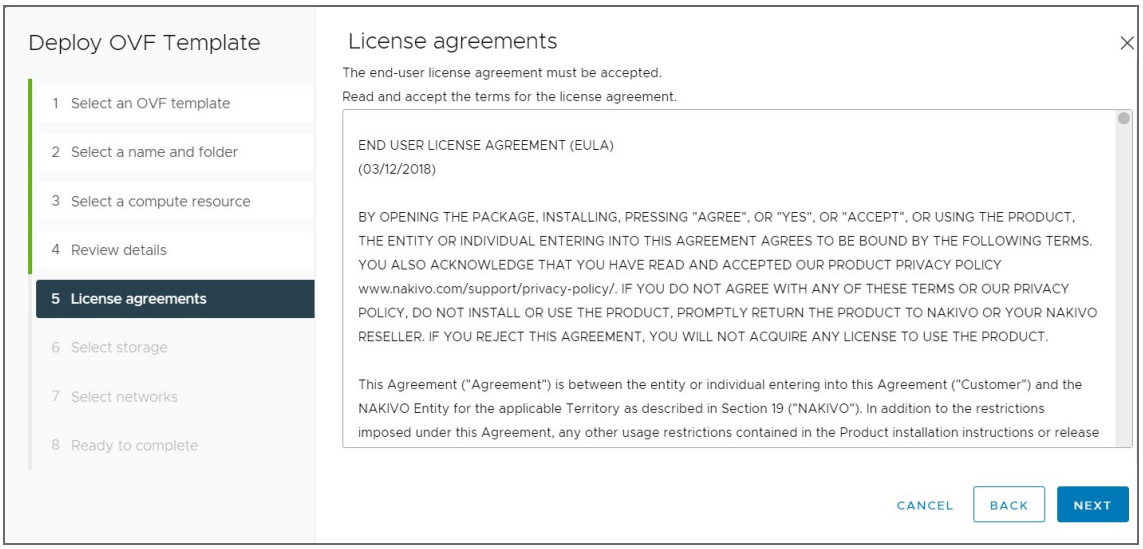
6. On the **Select a computer resource** page, select the resource pool within which you would like to deploy the Virtual Appliance and click **Next**.



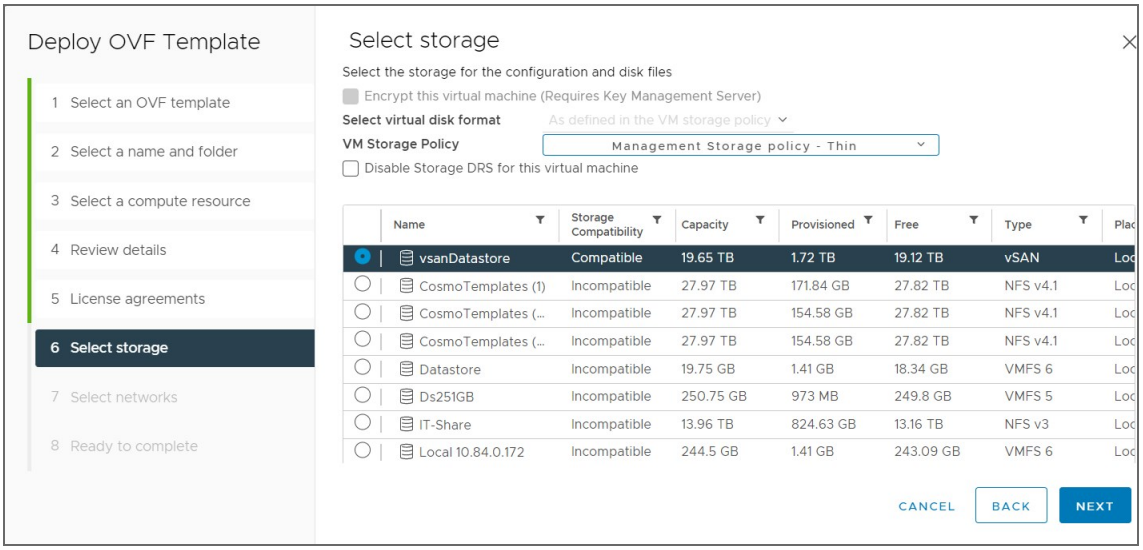
7. On the **Review details** page, review the template details and click **Next**.



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



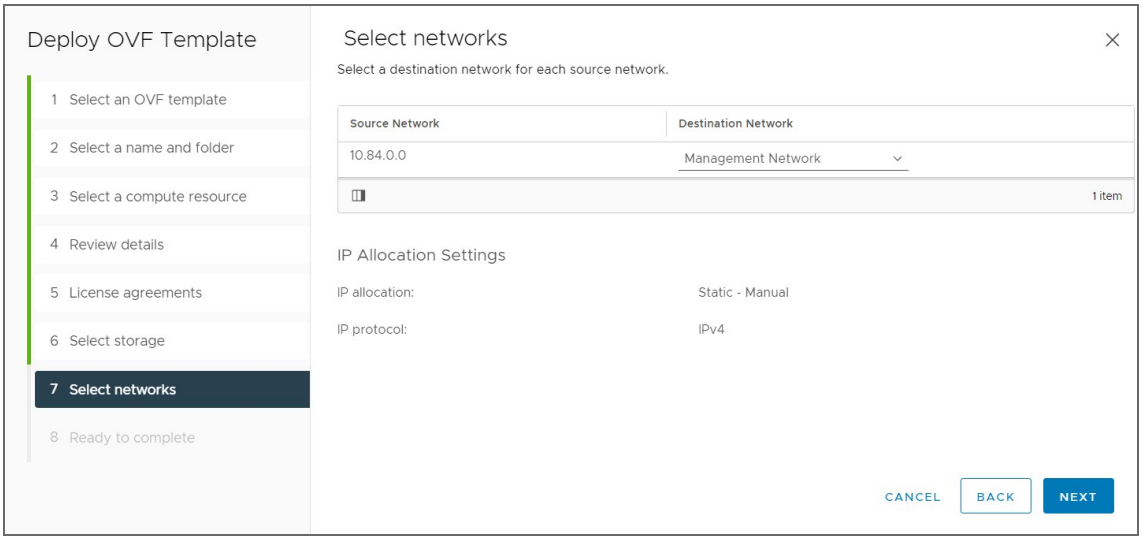
9. On the **Select storage** page, select a datastore in which you would like to keep the Virtual Appliance disk, virtual disk format (*Thin Provisioning* is recommended), VM storage policy and click **Next**.



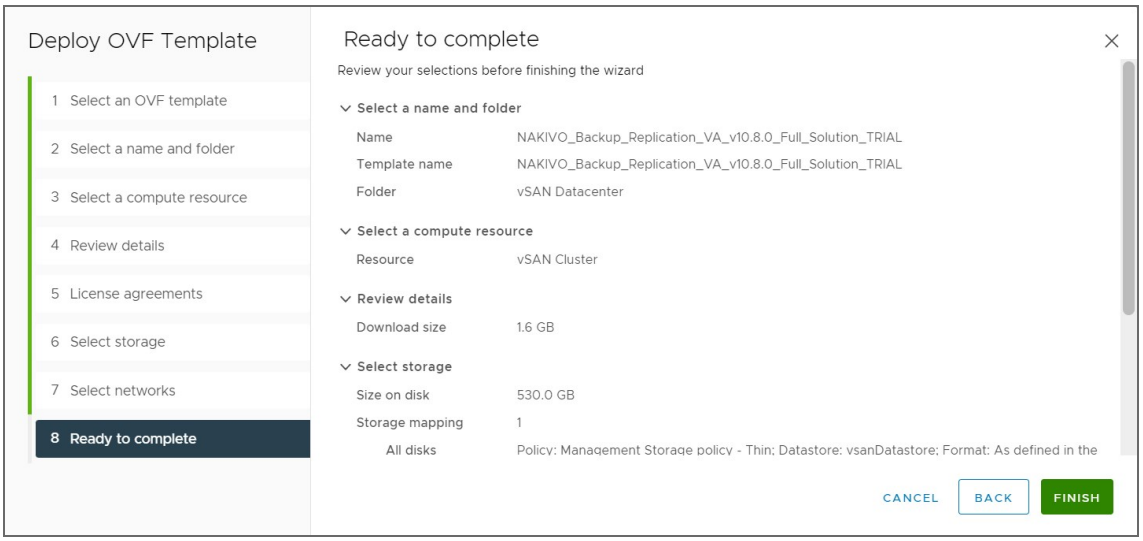
Important

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

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



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



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

Important

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

Virtual Appliance OS, Credentials, and Security

The appliance runs Ubuntu 24.04, 64-bit. Use the following credentials to log in to the appliance:

- **Username**: nkvuser
- **Password**: QExS-6b%3D

The following vSphere versions are supported for the VA deployment:

- 6.5
- 6.7
- 7.0
- 8.0
- 9.0

Note

This limitation on supported versions only applies to when deploying the VA using VMware desktop client. When a VA is deployed via NAKIVO Backup & Replication UI (usually a Transporter), no problems should occur.

For the versions of NAKIVO Backup & Replication older than 7.2, the password is `root`.

Important

- If you plan to expose the Virtual Appliance to the Internet, change the default VA credentials and set up a login and password for the Web interface.
- It is recommended to run an update on all packages in your Virtual Appliance at least once a month.

Web Interface Login

Open your browser and go to the VA's web interface. Refer to the Getting Started section to better understand how to continue working with NAKIVO Backup & Replication.

Deploying Nutanix AHV Virtual Appliance

- Deploying Nutanix AHV Virtual Appliance
- Virtual Appliance OS, Credentials, and Security
- Web Interface Login

Deploying Nutanix AHV Virtual Appliance

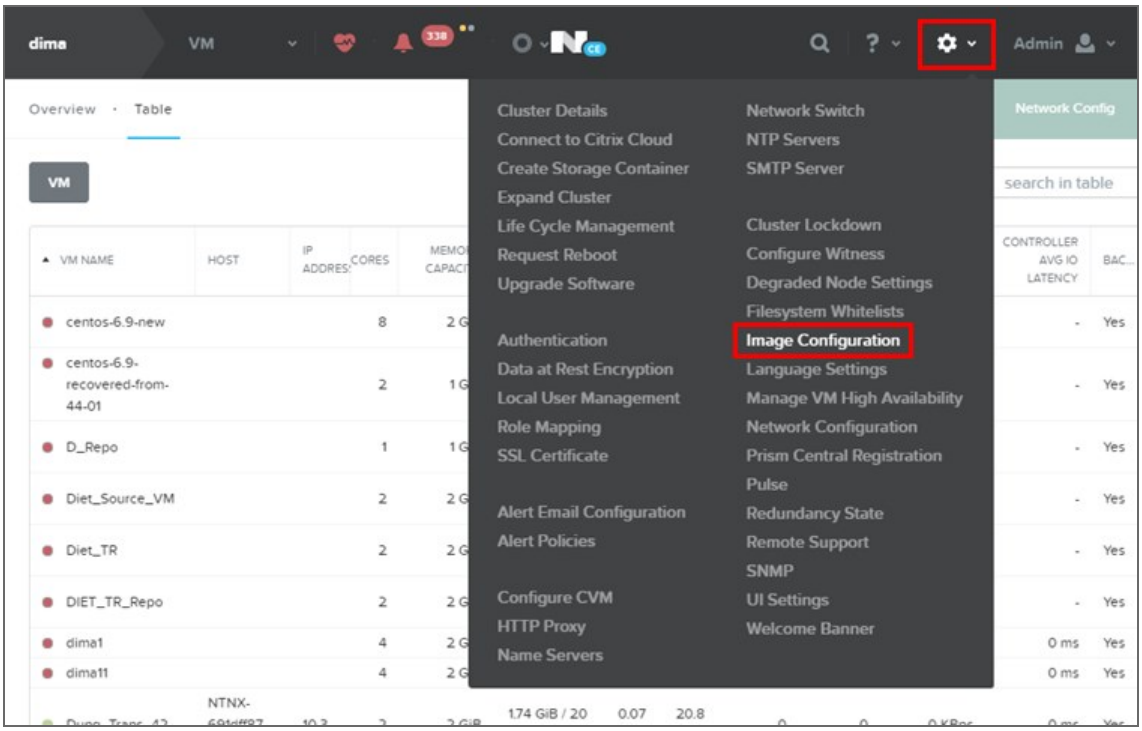
The NAKIVO Backup & Replication instance must be deployed in a Nutanix AHV cluster in order to enable backup and recovery functions.

NAKIVO Backup & Replication offers the following solutions:

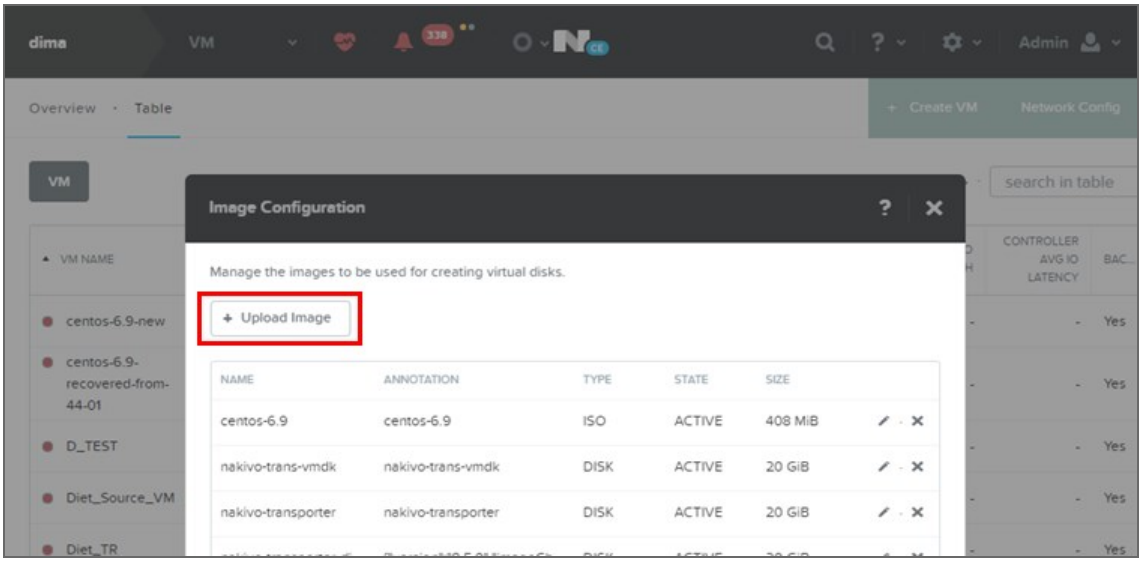
- Full Solution (Single Tenant) - requires a 100 GB thin provisioned disk
- **Transporter**-only - requires a 20 GB thin provisioned disk

To deploy a virtual appliance via the Nutanix Prism application, follow the steps below:

1. Download the `.VMDK` file with a full or transporter-only image from the Nakivo website and store it locally.
2. Log in to the Prism console.
3. From the **Configurations** menu, select **Image Configuration**.

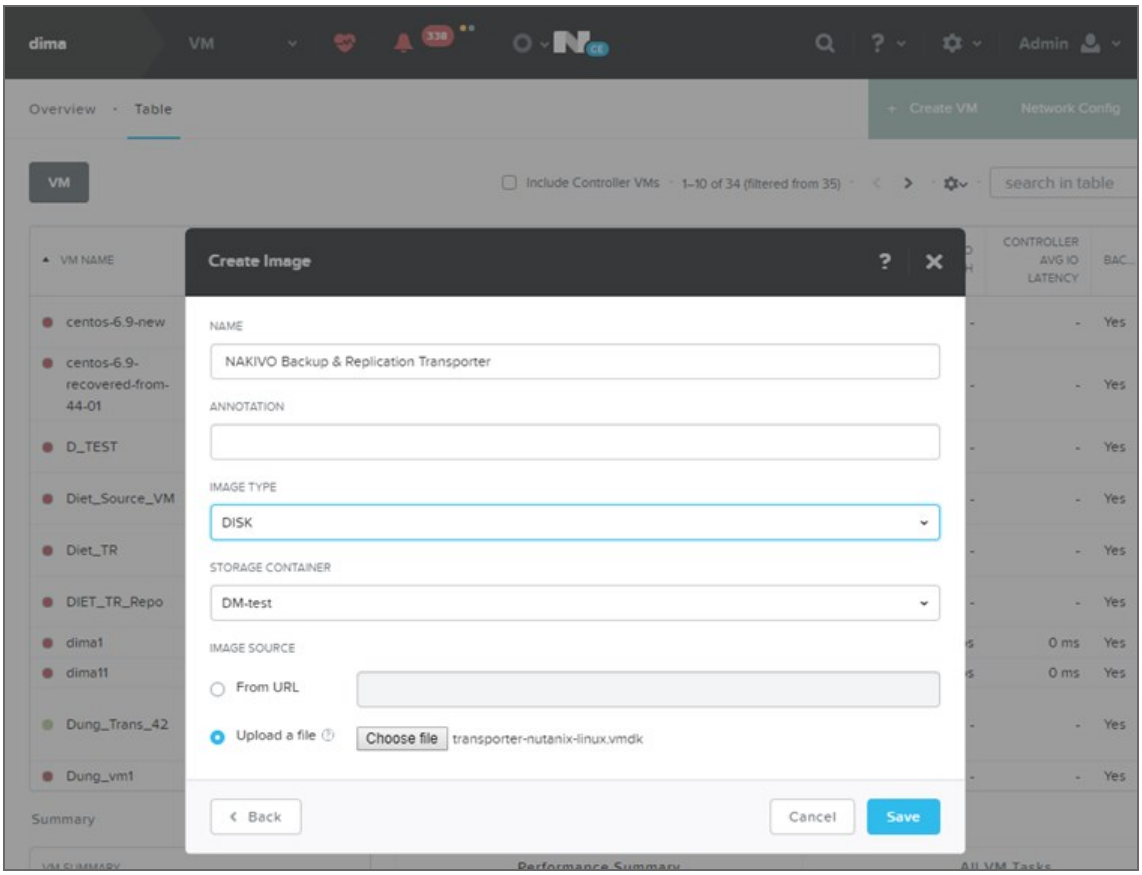


4. In the **Image Configuration** dialog, click **Upload Image**.

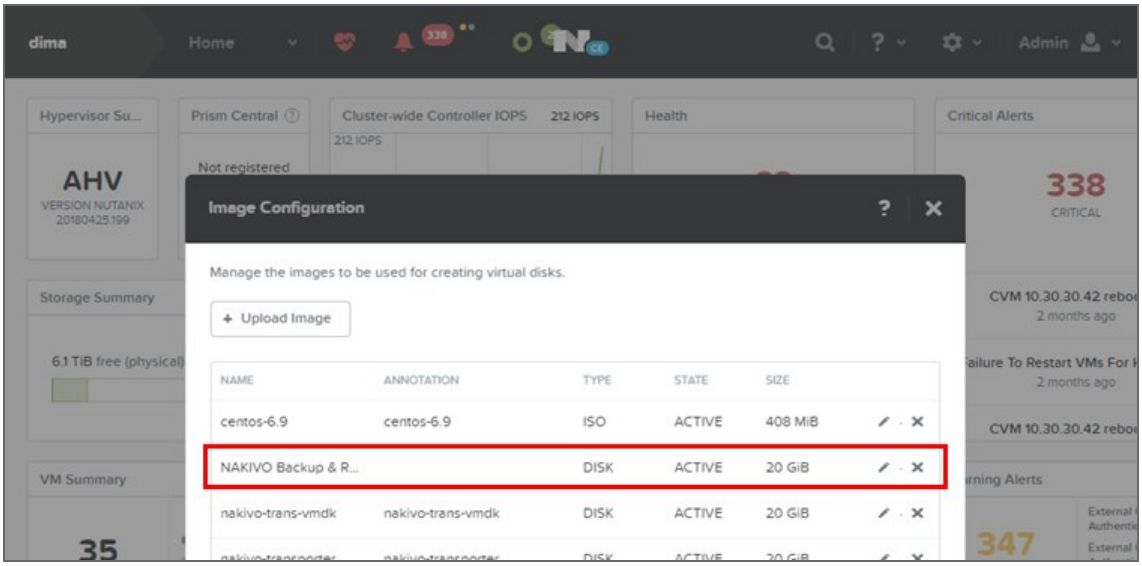


5. In the **Create Image** dialog, fill in the following options:
    - **Name**: Enter a name for the new image.
    - **Image Type**: From the drop-down list, select **DISK**.
    - **Storage Container**: Select the storage container you wish to use from the drop-down list. The list includes all storage containers created for this cluster. If there are no storage containers currently available, a **Create Storage Container** link is displayed.
    - **Image Source**: Click the **Upload a file** radio button to upload a file from your workstation. Click

the **Choose File** button and then select the file to upload from the file search window.



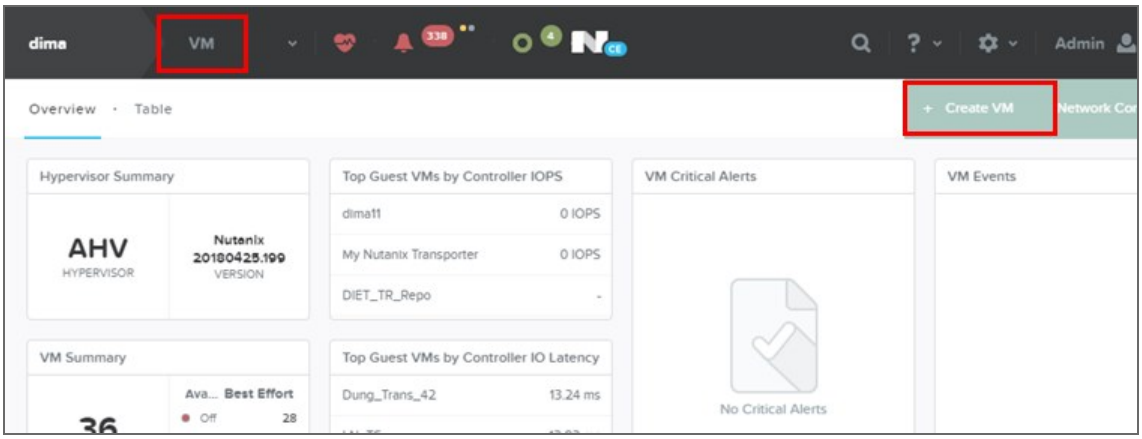
6. When all fields are correct, click the **Save** button.
   After the file uploading completes, the **Create Image** window closes and the **Image Configuration** window reappears with the new image present in the list.



Note

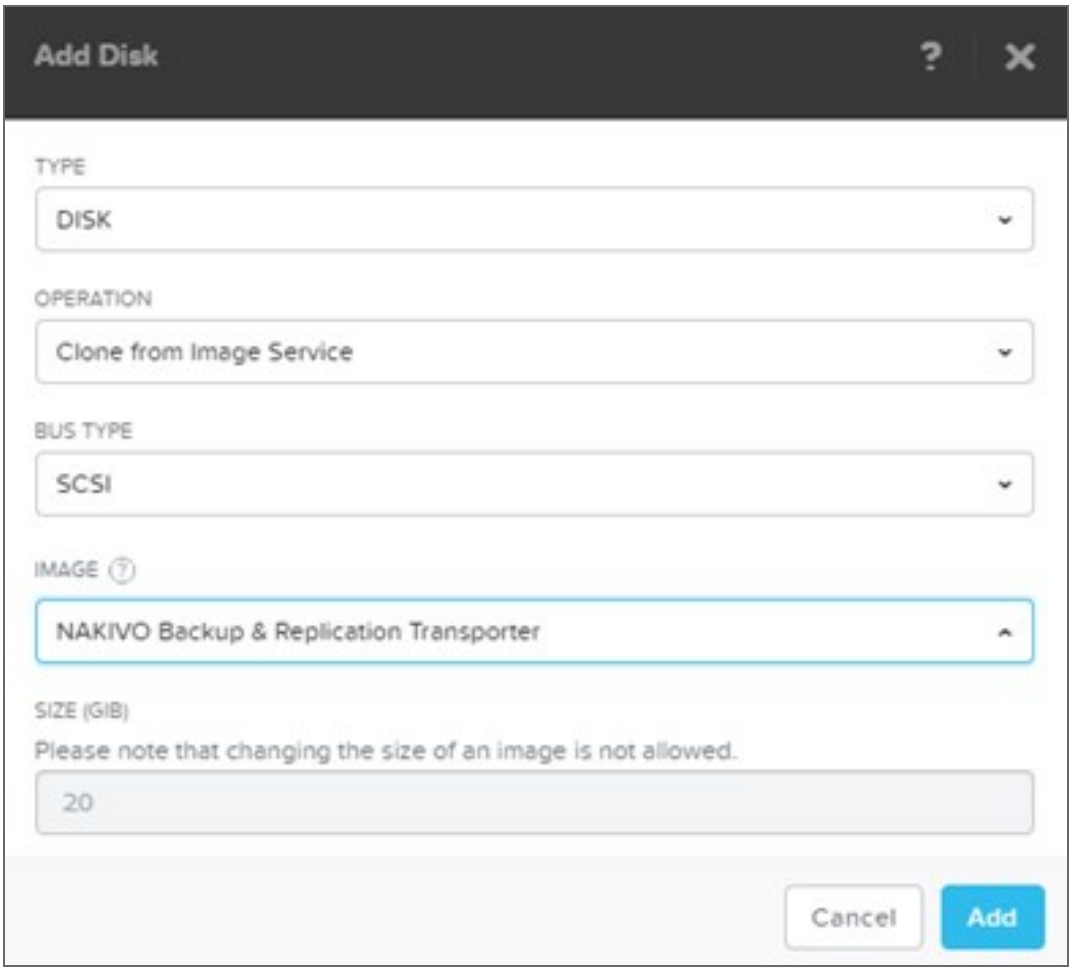
Make sure the status of the disk is **Active** before proceeding to the next step.

7. Close the **Image Configuration** window, go to the **VM** view and click **Create VM**.



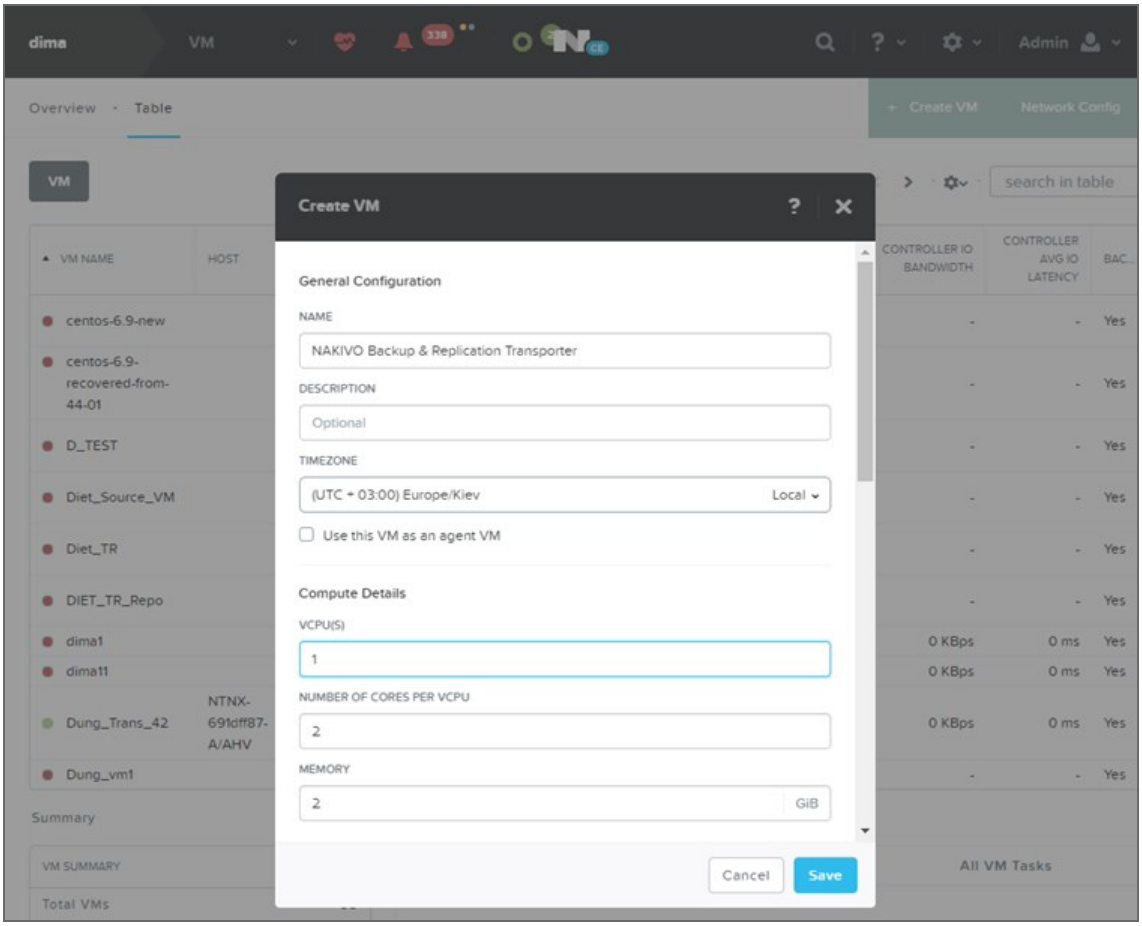
8. In the **Create VM** dialog, fill in the following options:
   - **Name**: Enter a name for the VM.
   - **vCPU(s)**: Enter the number of virtual CPUs to allocate to this VM (minimum 1).
   - **Number of Cores per vCPU**: Enter the number of cores assigned to each virtual CPU (minimum 2).
   - **Memory**: Enter the amount of memory (in GBs) to allocate to this VM (minimum 4 GB + 250 MB for each concurrent job for full solution/minimum 2 GB + 250 MB for each concurrent job **Transporter**-only solution).
   - In the **Disk** section, click **Add New Disk**, and specify the following settings in the **Add Disk** dialog:
     a. **Type**: Select **Disk**.
     b. **Operation**: Select **Clone from Image Service**.
     c. **Bus Type**: Select **SCSI**.

d. **Image**: Select your uploaded image from the list.

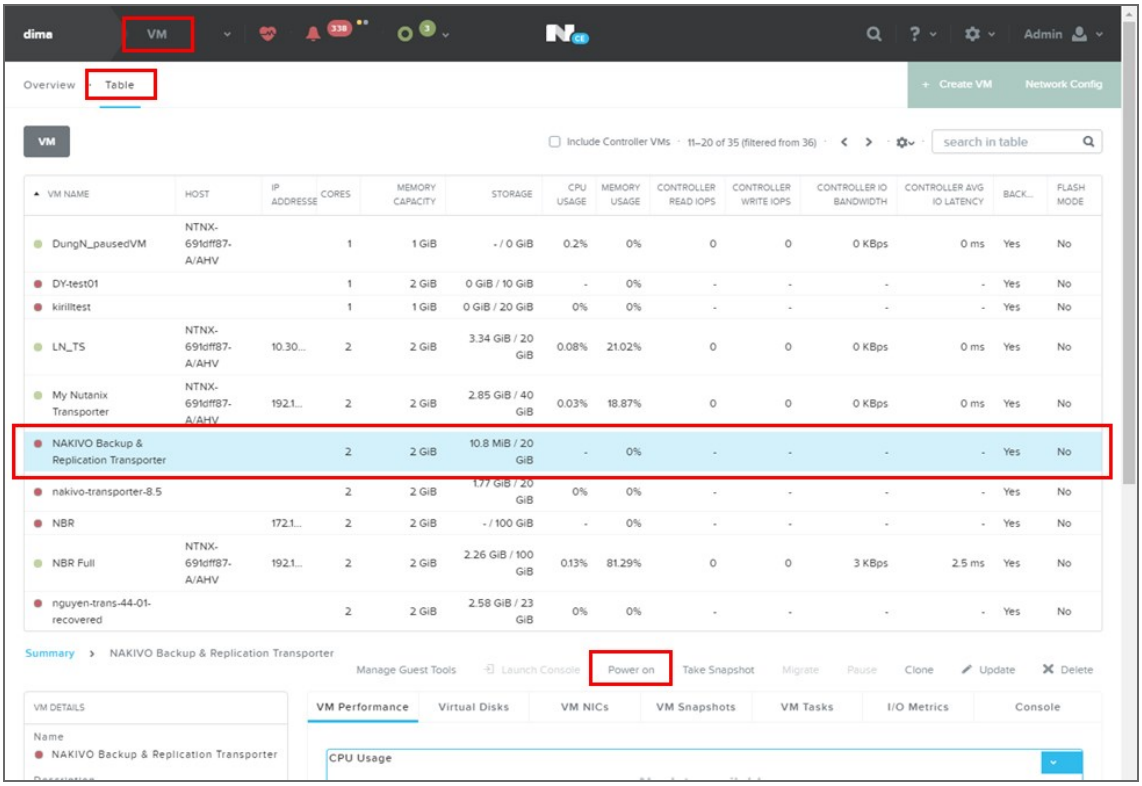


- In the **Network Adapters (NIC)** section, click **Add New NIC** and select an available VLAN from the list.

9. Click **Save**.



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



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

Virtual Appliance OS, Credentials, and Security

The appliance runs Ubuntu 24.04, 64-bit. Use the following credentials to log in to the appliance:

- **Username**: nkvuser
- **Password**: QExS-6b%3D

For the versions of NAKIVO Backup & Replication older than 7.2, the password is `root`.

Important

- If you plan to expose the Virtual Appliance to the Internet, change the default VA credentials and set up a login and password for the Web interface.
- It is recommended to run an update on all packages in your Virtual Appliance at least once a month.

Web Interface Login

Open your browser and go to the VA web interface. Refer to "Getting Started" on page 221 to better understand how to continue working with NAKIVO Backup & Replication.

Deploying Amazon Machine Image in Amazon EC2

You can deploy NAKIVO Backup & Replication as a pre-configured Amazon Machine Image (AMI) in Amazon EC2. After you complete the download form, you get a link to the AWS marketplace page where you can download the AMI.

Configuring AMI Parameters

Configure the following AMI parameters:

1. **Instance Type**: More powerful instances can process tasks faster and run more tasks simultaneously. The minimum requirement for NAKIVO Backup & Replication is the t2.micro instance type; the t2 medium instance type is recommended.
2. **Instance Details**: Assign a public IP to the instance if you wish to access the instance over the Internet.
   - **Security Group**: Use the "All Traffic" rule or create a set of rules listed below:

Type	Port Range	Source	Description
SSH	2221	0.0.0.0/0	Enables remote SSH access to the instance
Custom TCP	80	0.0.0.0/0	Enables access to the web interface
Custom TCP	443	0.0.0.0/0	Required for local **Transporter** import
Custom TCP	902	0.0.0.0/0	Required for local **Transporter** import
Custom TCP	4443	0.0.0.0/0	Enables access to the web interface
Custom TCP	9446	0.0.0.0/0	Enables access to a remote **Transporter**
Custom TCP	9448-10000	0.0.0.0/0	Enables access to a remote **Transporter**
All ICMP	0-65535	0.0.0.0/0	Enables access to a remote **Transporter**

Note
Older AMIs may still use SSH Port *22* instead of *2221*.

4. **Key pair**: Select an existing key pair or create a new key pair for your instance. If you select an existing key pair, make sure you have access to the private key file.

Note
The AMI deliverable uses Ubuntu 22.04 OS and a standalone EC2 instance with a **Director** and **Transporter**. Instead of the default system user **ubuntu**, the AMI uses the username **nkvuser**.

Refer to “Getting Started” on page 221 to better understand 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.

1. Set the installation options as follows:
   - **Installation type**: Leave the **Full solution** option selected to install the key product components (Director and Transporter)
   - **Create repository**: Leave the checkbox selected to create a **Backup Repository** on the machine on which NAKIVO Backup & Replication is installed.
   - Optionally, click **Browse** and select a folder to change the default location of the **Backup Repository**.



2. Click **MORE OPTIONS** to set up more installation options:

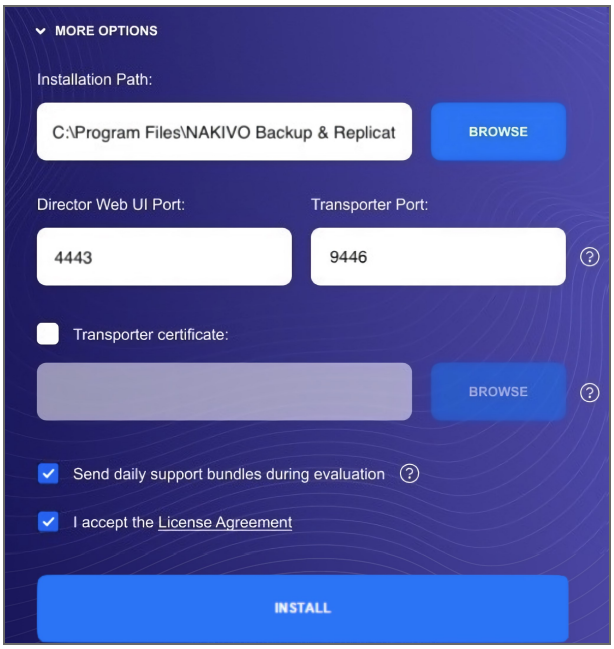
- **Installation path**: The location where NAKIVO Backup & Replication will be installed. If you want to change the default path to NAKIVO Backup & Replication, click **Browse** and select a new location.
- **Director Web UI port**: The default port that will be used to connect to the Web UI of NAKIVO Backup & Replication. Make sure that the port you specify is open in your firewall.
- **Transporter port**: The default port that will be used by the **Director** to communicate with the Onboard **Transporter**. Make sure that the port you specify is open in your firewall.
- **Transporter certificate**: This allows you to use a CA Certificate. Enter the path to the folder containing the certificate file in the field.

Note

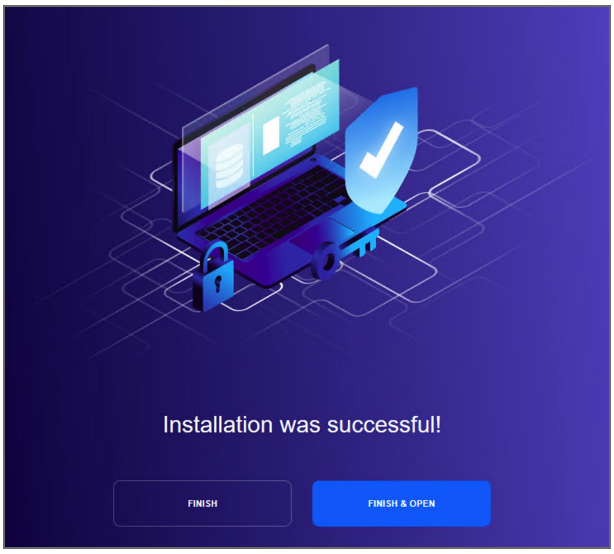
- When the checkbox is not selected, NAKIVO Backup & Replication automatically installs a self-signed certificate.
- It is possible to install a CA-signed certificate for the **Transporter** by conducting silent installation using the command-line arguments passed to the installer:
  - Use the following command for Windows OS:
    `installer.exe --cert C:\certificate.pem --eula-accept`
    The short option for the Windows OS command is the following:
    `installer.exe -ct C:\certificate.pem -ea`
  - Use the following command for Linux OS:
    `installer.sh --cert /tmp/certificate.pem --eula-accept`

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

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



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



6. To prevent unauthorized access to the product, create your user account. Fore more details, refer to “Logging in to NAKIVO Backup & Replication” on page 222 .

Installing Transporter-Only on Windows

If you have already installed the full solution (both **Director** and **Transporter**) and wish to deploy an additional **Transporter**, run the NAKIVO Backup & Replication installer for Windows and follow the steps below:

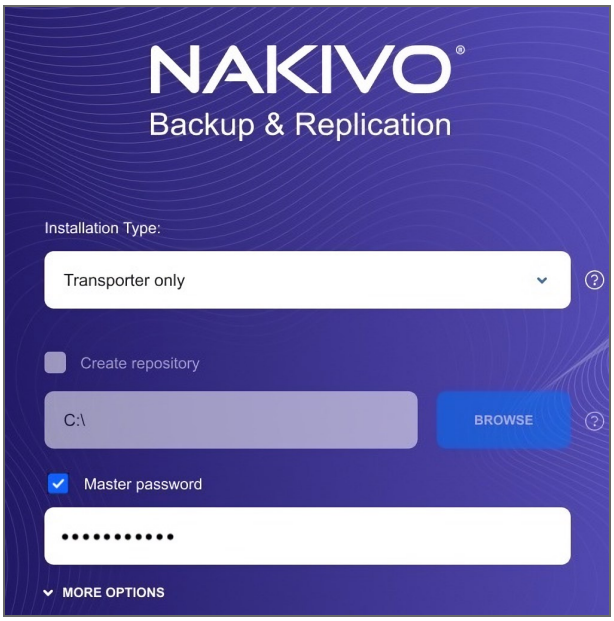
Transporter Installation Prerequisites

Prior to installing the **Transporter**, make sure the following prerequisites are met:

- Make sure the machine on which you plan to install the **Transporter** has a connection to the relevant items below:
  - The machine on which the **Director** is installed.
  - VMware/Hyper-V/Nutanix AHV servers on which you plan to back up or replicate VMs (provided that you plan to retrieve VM data using the Transporter you are about to install)
  - Machines on which you have installed other **Transporters** (provided that you plan to set up data transfer between an existing Transporter and the one you are about to install)
  - **Backup Repositor**y (provided that you plan to assign the **Transporter** you are about to install to a **Backup Repository**)
  - VMware/Hyper-V/Nutanix AHV servers which you plan to use as a destination for replicated VMs (provided that you plan to write data to the target servers and datastores using the Transporter you are about to install)
- For VMware/Hyper-V/Nutanix AHV servers discovered with DNS names, make sure those DNS names can be resolved on the machine on which to install the Transporter.

Transporter Installation Process

1. Run the NAKIVO Backup & Replication installer.
2. Choose **Transporter only** from the **Installation type** drop-down list.



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

Note

- The master password must adhere to the following requirements:
  - Length between 8 and 511 characters
  - At least one uppercase (A-Z) and one lowercase letter (a-z)
  - At least one digit (0-9)
  - At least one special character (e.g., !@#$%^&*()_+-=[]{}|;:'",.<>?/)
  - Standard English characters and spaces only (no emojis or diacritics)

  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.
- When installing NAKIVO Backup & Replication components using console or silent installation modes on Windows, password input is subject to Windows command-line parsing rules:
  - Passwords containing spaces must be enclosed in double quotation marks. For example: -b "Abcd 1ljk".
  - Special characters such as !, &, |, (, and) must be escaped using the caret (^). For example: -b password1^!A.
  - Passwords containing both spaces and special characters must be enclosed in double quotation marks, with special characters escaped. For example: -b "Abcd 1^!jk".

4. Click **MORE OPTIONS** and set up the following:
   - **Installation path**: The location where the **Transporter** will be installed. If you want to change the default path to the **Transporter** installation folder, click **Browse** and select a new location.
   - **Transporter port**: The default port that will be used by the **Director** to communicate with the Onboard **Transporter**. Make sure that the port you specify is open in your firewall.

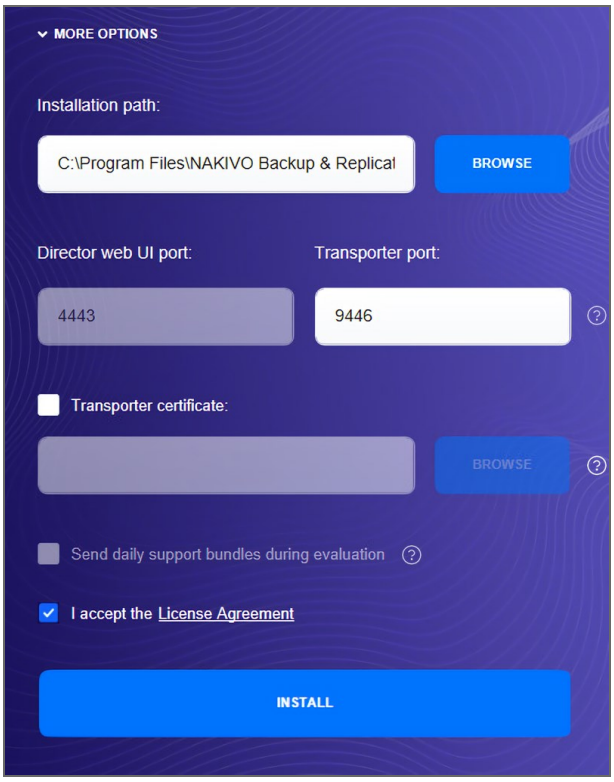
- **Transporter certificate**: This allows you to use a CA Certificate. Enter the path to the folder containing the certificate file in the field.

  Note

  - When the checkbox is not selected, NAKIVO Backup & Replication automatically installs a self-signed certificate.
  - It is possible to set up a master password and CA-signed certificate for the **Transporter** by conducting silent installation using the command-line arguments passed to the installer:
    - Use the following command for Windows OS:
      ```
      installer.exe --cert C:\certificate.pem --master-pass P@ssword123 --eula-accept
      ```
      The short option for the Windows OS command is the following:
      ```
      installer.exe -ct C:\certificate.pem -b P@ssword123 -ea
      ```
    - Use the following command for Linux OS:
      ```
      installer.sh --cert /tmp/certificate.pem -b P@ssword123 --eula-accept
      ```

- **Send daily support bundles during evaluation**: If this option is selected, NAKIVO Backup & Replication will automatically create, encrypt, and upload support bundles once a day to a NAKIVO support server during the evaluation period. NAKIVO Support team may use this information to improve the product experience and may be able to identify and resolve product issues faster.

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



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

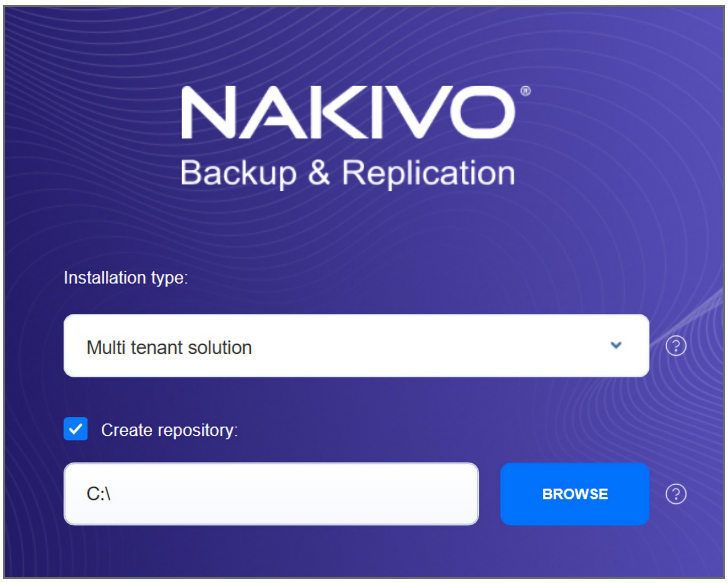


8. Add the Transporter to NAKIVO Backup & Replication.

Installing Full Solution in Multi-Tenant Mode on Windows

To install the full solution in multi-tenant mode on a Windows OS, run the NAKIVO Backup & Replication installer for Windows and follow the steps below:

1. Set the installation options as follows:
    - **Installation type**: Select the **Multi tenant solution** option from the **Installation type** drop-down list.
    - **Create repository**: Leave the checkbox selected to create a **Backup Repository** on the machine on which NAKIVO Backup & Replication is installed.
        - Optionally, click **Browse** and select a folder to change the default location of the **Backup Repository**.



2. Click **MORE OPTIONS** to set up more installation options:
    - **Installation path**: The location where NAKIVO Backup & Replication will be installed. If you want to change the default path to the product, click **Browse** and select a new location.
    - **Director Web UI port**: The default port that will be used to connect to the Web UI of NAKIVO Backup & Replication. Make sure that the port you specify is open in your firewall.
    - **Transporter port**: The default port that will be used by the **Director** to communicate with the Onboard **Transporter**. Make sure that the port you specify is open in your firewall.

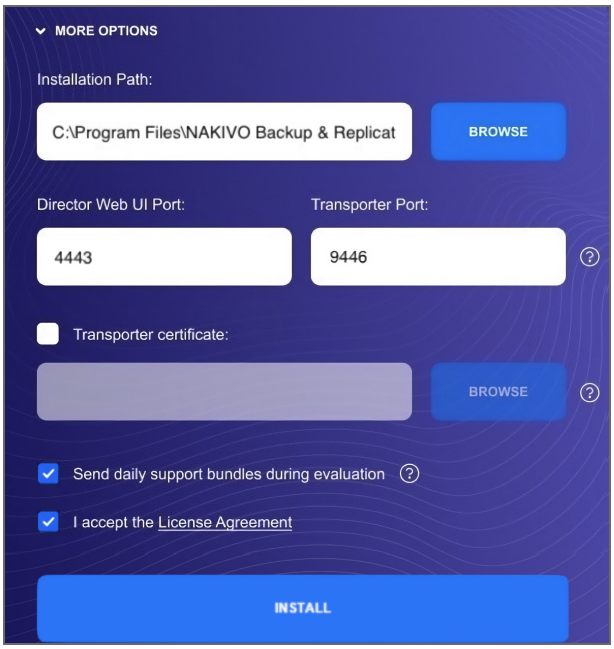
- **Transporter certificate**: This allows you to use a CA Certificate. Enter the path to the folder containing the certificate file in the field.

  Note

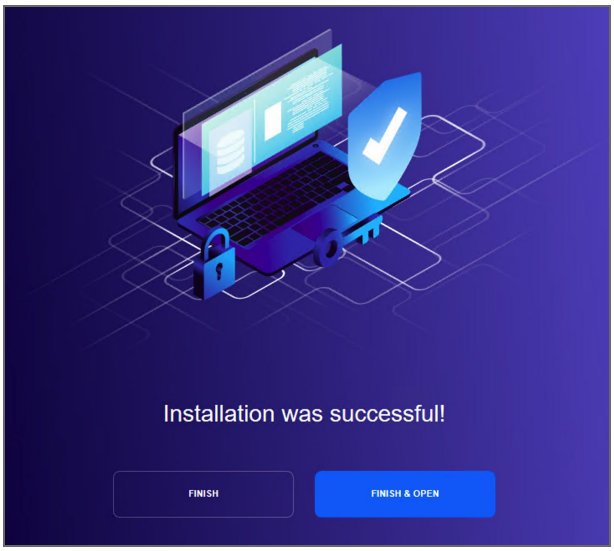
  - When the checkbox is not selected, NAKIVO Backup & Replication automatically installs a self-signed certificate.
  - It is possible to install a CA-signed certificate for the **Transporter** by conducting silent installation using the command-line arguments passed to the installer:
    - Use the following command for Windows OS:
      `installer.exe --cert C:\certificate.pem --eula-accept`
      The short option for the Windows OS command is the following:
      `installer.exe -ct C:\certificate.pem -ea`
    - Use the following command for Linux OS:
      `installer.sh --cert /tmp/certificate.pem --eula-accept`

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

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



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



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

6. Create an account by completing the form. For details, refer to “Logging in to NAKIVO Backup & Replication” on page 222.

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. When installing NAKIVO Backup & Replication components using console or silent installation modes on Windows, password input is subject to Windows command-line parsing rules: • Passwords containing spaces must be enclosed in double quotation marks. For example: -b "Abcd 1ljk". • Special characters such as !, &, |, (, and) must be escaped using the caret (^). For example: -b password1^!A. • Passwords containing both spaces and special characters must be enclosed in double quotation marks, with special characters escaped. For example: -b "Abcd 1^!jk".

Installing on Linux

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

Linux Installation Prerequisites

In order to install and use NAKIVO Backup & Replication on a Linux OS, make sure the following packages are installed:

- cifs-utils
- iscsi-initiator-utils
- ntfs-3g
- tar

Silent Installation on Linux

You can install NAKIVO Backup & Replication in silent mode via a command line. To install the full solution, simply run the following command: `installer.sh -f --eula-accept` This will install all product components (Director, Transporter, Backup Repository) and you will be able to use all product features after installation.

The following arguments are available:

Argument	Description
-h, -help, help	Display the list of available arguments without starting the installation.
--eula-accept, -ea	Indicates that you have read and agree to the End User License Agreement.
-f	Shall perform the silent installation of the full solution (**Director** and **Transporter**).
-t	Shall perform the silent installation of **Transporter** only.
-m	Shall perform the silent installation of the full solution in multi-tenant mode.
-u	Shall perform the silent update of the installed product components.

Argument	Description
-e	Shall install **Transporter** on Amazon EC2, or update **Transporter** installed on Amazon EC2. Refer to Updating on Amazon EC2 for details.
-a	Shall enable uploading support bundles to support team server (Call Home). Refer to System Settings for details.
-y	Shall accept limitations silently.
-i <install_path>	Shall install to the specified installation path.
-d <director_port>	Shall provide a custom **Director** port.
-p <transporter_port>	Shall provide a custom **Transporter** port.
-r <port1>-<port2>	Shall provide a custom transporter data ports range.
-C	Shall suppress creating the repository.
-c <repo_path>	Shall create the repository. The **<repo_path>** parameter is optional.
--rt <repo_type>	Shall create a repository of the specified type. The **<repo_type>** parameter may accept the following values: **1** – "Forever incremental with deduplication"; **2** – "Forever incremental without deduplication"; **3** – "Incremental with full backups (deduplication devices)".
--rc <compress_level>	Shall specify the repository compression level. The parameter may accept the following values: Disabled; Fast; Medium; Best.
--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**. Passwords can be passed as literal strings by enclosing them in single quotation marks to prevent shell expansion of special characters. For example: -b 'Pass! Word'

Installing Full Solution on Linux

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

1. Upload the installer file to the machine on which you want to install NAKIVO Backup & Replication in the *binary transfer mode*. For example:

- Upload the installer from a Windows-based machine.
- Upload the product from a Linux-based machine: run the following command: `wget 'server_ip/shared/NAKIVO_Backup_&_Replication_TRIAL.sh'`

2. Log in to the Linux machine and allow the execution of the installer file.
   For example: `chmod +x NAKIVO_Backup_&_Replication_TRIAL.sh`
3. Execute the installer file with root privileges.
   For example: `sudo ./NAKIVO_Backup_&_Replication_TRIAL.sh`
4. Review the license agreement (press **Space** to go to the next page of the agreement). If you agree to the terms of the license agreement, press "Y" and then press **Enter**.

   Note
   On some Linux consoles you need to press "Q".

5. Type "S" to install the full solution and press **Enter**.
6. Optionally, you can install CA **Transporter** certificate. Enter the path to the folder containing the certificate file and press **Enter**.

   Notes
   - If no path to the CA certificate was provided, NAKIVO Backup & Replication automatically installs a self-signed certificate.
   - It is possible to set up CA-signed certificate for the **Transporter** by conducting silent installation using the command-line arguments passed to the installer. Use the following command:
     `installer.sh --cert /tmp/certificate.pem --eula-accept`

7. Specify the installation path for the product: Press **Enter** to accept the default installation path "/opt/nakivo" or enter a custom path and press **Enter**.
8. Specify the **Director** HTTPS port (which will be used to access the Web UI of NAKIVO Backup & Replication): Press **Enter** to accept the default port *4443* or enter a custom port number and press **Enter**. Make sure the port you specify is open in your firewall.
9. Specify whether to allow the product to automatically send support bundles to a NAKIVO server during the evaluation period (Call Home). If this option is enabled, NAKIVO Backup & Replication will automatically create, encrypt, and upload support bundles once a day to a NAKIVO support server during the evaluation period. NAKIVO Support team may use this information to improve the product experience and will be able to identify and resolve product issues faster.
10. Specify the Transporter port (which will be used to connect to the **Transporter** that is installed by default with the **Director**): Press **Enter** to accept the default port "9446" or enter a custom port number (1 to 65535) and press **Enter**. Make sure the port you specify is open in your firewall.

11. Specify a range of port numbers (from *1* to *65535*) that will be used to transfer data by the Onboard **Transporter** (default are *9448-10000*). The range you specify should contain at least 100 ports. Make sure that the ports you specify are open in your firewall.
12. Specify a path to the default Backup Repository: Press **Enter** to accept the default path "/opt/nakivo/repository" or enter a custom path and press **Enter** to begin the installation process.

After the installation is complete, you can log in to NAKIVO Backup & Replication by opening the product's web interface in your browser (use the machine's IP address or DNS name with the appropriate HTTPS port). By default, login name and password are not required to access NAKIVO Backup & Replication. To prevent unauthorized access to the product, you can set up credentials in **Configuration**.

Installing Transporter on Linux

If you have already installed the full solution (both **Director** and **Transporter**) and want to deploy an additional **Transporter**, run the NAKIVO Backup & Replication installer for Windows and follow the steps below:

Transporter Installation Prerequisites

Prior to installing a **Transporter**, make sure the following prerequisites are met:

1. Make sure the machine on which you plan to install the **Transporter** has a connection to the relevant items below:
    - The machine on which the **Director** is installed
    - VMware/Hyper-V servers on which you plan to back up or replicate VMs (if you plan to retrieve VM data using the **Transporter** you are about to install)
    - Machines on which you have installed other **Transporters** (if you plan to set up data transfer between an existing **Transporter** and the one you are about to install)
    - Backup repository (if you plan to assign the **Transporter** you are about to install to a **Backup Repository**)
    - VMware/Hyper-V servers which you plan to use as a destination for replicated VMs (if you plan to write data to the target servers and datastores using the **Transporter** you are about to install)
2. If you have discovered VMware/Hyper-V servers using DNS names, make sure those DNS names can be resolved on the machine on which you plan to install the **Transporter**.

Transporter Installation

1. Upload the installer file to the machine on which you wish to install NAKIVO Backup & Replication in the *binary transfer mode*. For example:
    - Upload the installer from a Windows-based machine.
    - Upload the product from a Linux-based machine: run the following command: `wget 'server_ip/shared/NAKIVO_Backup_&_Replication_TRIAL.sh'`

2. Allow the execution of the installer file. For example: `chmod +x NAKIVO_Backup_&_Replication_TRIAL.sh`
3. Execute the installer file with root privileges. For example:`sudo ./NAKIVO_Backup_&_Replication_TRIAL.sh`
4. Review the license agreement (press **Space** to go to the next page of the agreement). If you agree to the terms of the license agreement, press "Y" and then press **Enter**.
5. Type "T" to install only the **Transporter** and press **Enter**.

> **Note**
> Alternatively, you can use the **-t** argument to install the **Transporter** silently:
> `sudo ./NAKIVO_Backup_&_Replication_TRIAL.sh -t`

6. Optionally, enter the master password that will be used to generate a pre-shared key and secure the **Transporter** and then press **Enter**.

> **Notes**
> - The master password must adhere to the following requirements:
>   - Length between 8 and 511 characters
>   - At least one uppercase (A-Z) and one lowercase letter (a-z)
>   - At least one digit (0-9)
>   - At least one special character (e.g., !@#$%^&*()_+-=[]{}|;:'",.<>?/)
>   - Standard English characters and spaces only (no emojis or diacritics)
> - Passwords can be passed as literal strings by enclosing them in single quotation marks to prevent shell expansion of special characters. For example: -b 'Pass! Word'
> - The master password can be set and re-set manually by following these steps:
>   1. Switch to root using the following command:
>   `sudo -i`
>   2. Stop the **Transporter** service.
>   3. Go to the **Transporter** folder with the following command:
>   `cd /opt/nakivo/transporter`
>   4. Run the following command to set the master password:
>   `./bhsvc -b P@ssword123`
>   5. Restart the Transporter service.

7. Specify the installation path for the product: Press **Enter** to accept the default installation path "/opt/nakivo" or enter a custom path and press **Enter**.

8. Optionally, you can install CA Transporter certificate. Enter the path to the folder containing the certificate file and press **Enter**.

Notes

- If no path to the CA certificate was provided, NAKIVO Backup & Replication automatically installs a self-signed certificate.
- It is possible to set up a master password and a CA-signed certificate for the **Transporter** by conducting silent installation using the command-line arguments passed to the installer. Use the following command:
  `installer.sh --cert /tmp/certificate.pem -b P@ssword123 --eula-accept`

9. Specify the **Transporter** port (used to connect to the **Transporter**): Press **Enter** to accept the default port *9446* or enter a custom port number and press **Enter** to begin the installation process. Make sure the port you specify is open in your firewall.

After the installation is complete, add the **Transporter** to NAKIVO Backup & Replication.

Installing Full Solution in Multi-Tenant Mode on Linux

Follow the steps below to install the full solution in multi-tenant mode on a Linux OS:

1. Upload the installer file to the machine on which you want to install NAKIVO Backup & Replication in the *binary transfer mode*. For example:
   - Upload the installer from a Windows-based machine.
   - Upload the product from a Linux-based machine: run the following command: `wget 'server_ip/shared/NAKIVO_Backup_&_Replication_TRIAL.sh'`
2. Log in to the Linux machine and allow the execution of the installer file.
   For example: `chmod +x NAKIVO_Backup_&_Replication_TRIAL.sh`
3. Execute the installer file with root privileges.
   For example: `sudo ./NAKIVO_Backup_&_Replication_TRIAL.sh`
4. Review the license agreement (press **Space** to go to the next page of the agreement). If you agree to the terms of the license agreement, press "Y" and then press **Enter**.
5. Type “M” to install the **Director** in Multi-tenant mode and press **Enter**.

   Note
   Alternatively, you can use the **-m** argument to install the solution in multi-tenant mode silently:
   `sudo ./NAKIVO_Backup_&_Replication_TRIAL.sh -m`

6. Optionally, you can install CA **Transporter** certificate. Enter the path to the folder containing the certificate file and press **Enter**.

Notes

- If no path to the CA certificate was provided, NAKIVO Backup & Replication automatically installs a self-signed certificate.
- It is possible to set up CA-signed certificate for the **Transporter** by conducting silent installation using the command-line arguments passed to the installer. Use the following command:
  `installer.sh --cert /tmp/certificate.pem --eula-accept`

7. Specify the installation path for the product: Press **Enter** to accept the default installation path "/opt/nakivo" or enter a custom path and press **Enter**.
8. Specify the Director HTTPS port (which will be used to access the Web UI of NAKIVO Backup & Replication): Press **Enter** to accept the default port *4443* or enter a custom port number and press **Enter**. Make sure the port you specify is open in your firewall.
9. Specify whether to allow the product to automatically send support bundles to a NAKIVO server during the evaluation period. If this option is enabled, NAKIVO Backup & Replication will automatically create, encrypt, and upload support bundles once a day to a NAKIVO support server during the evaluation period. NAKIVO Support team may use this information to improve the product experience and will be able to identify and resolve product issues faster.
10. Specify the **Transporter** port (which will be used to connect to the **Transporter** that is installed by default with the **Director**): Press **Enter** to accept the default port "9446" or enter a custom port number (*1* to *65535*) and press **Enter**. Make sure the port you specify is open in your firewall.
11. Specify a range of port numbers (from 1 to 65535) that will be used to transfer data by the Onboard **Transporter** (default are *9448-10000*). The range you specify should contain at least 100 ports. Make sure that the ports you specify are open in your firewall.
12. The onboard backup repository for the Master Tenant is automatically created after the installation.
13. Specify a path to the default backup repository: Press **Enter** to accept the default path `/opt/nakivo/repository` or enter a custom path and press **Enter** to begin the installation process.

Note

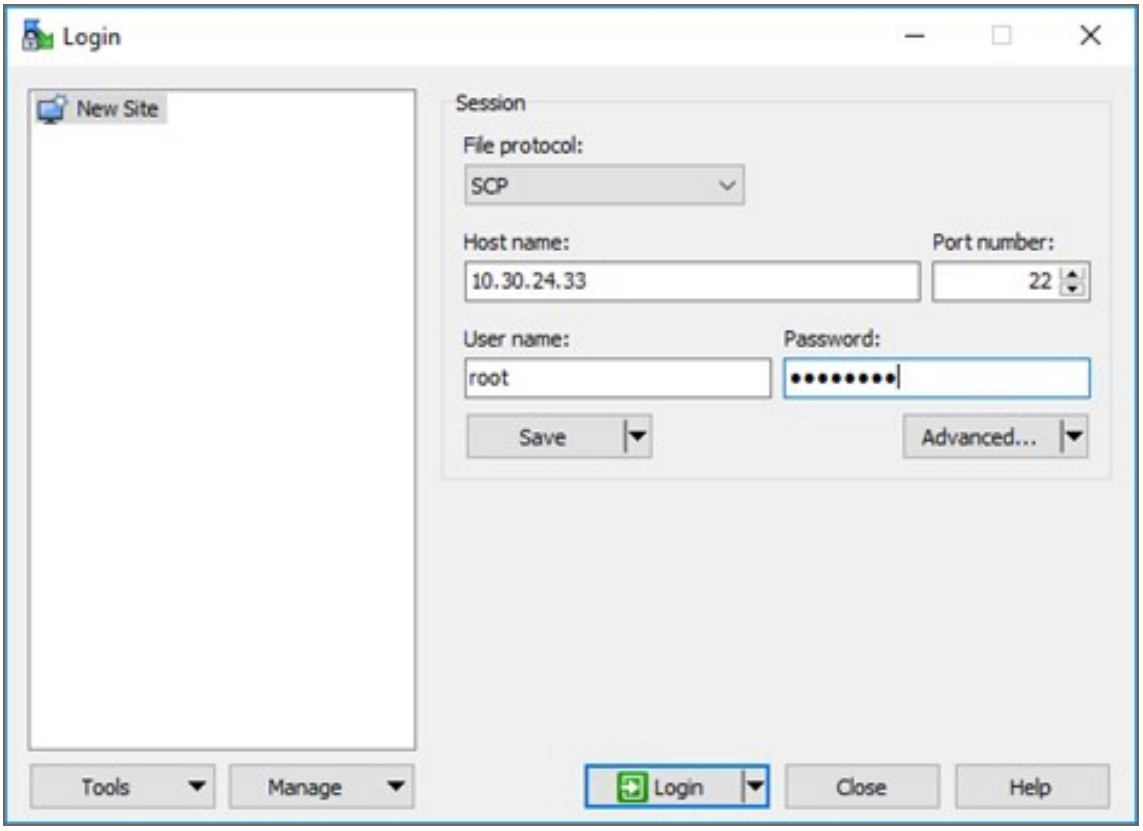
The onboard backup repository for the Master Tenant is automatically created after the installation.

After the installation is complete, you can log in to NAKIVO Backup & Replication by opening its web interface in your browser. Use the machine's IP address or DNS name together with the Director HTTPS port. Refer to "Getting Started" on page 221 to know how to continue working with NAKIVO Backup & Replication.

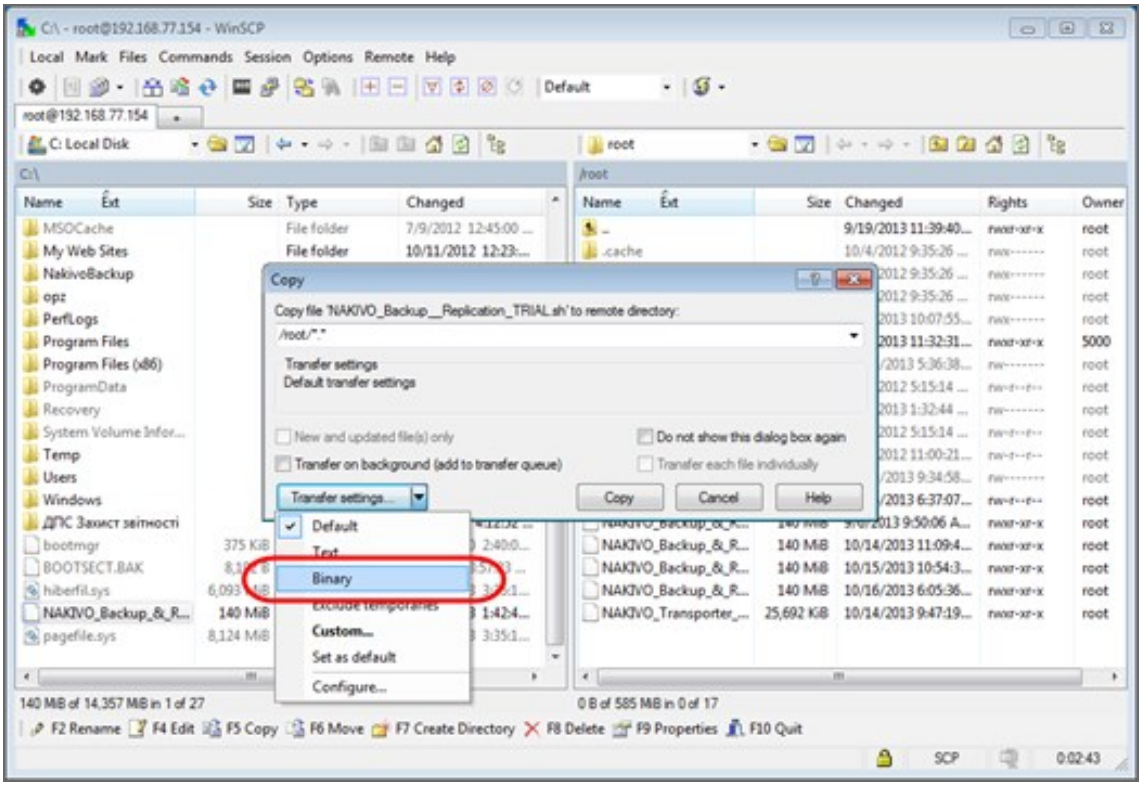
Uploading Installer from Windows Machine to Linux Machine

To upload the installer from a Windows-based machine, follow the steps below:

1. Download the free WinSCP client from http://winscp.net, install, and run it.
2. Choose **SCP** from the **File protocol** list.
3. Specify the IP address or the hostname of the Linux machine on which you would like to install the product in the **Host name** field.
4. Specify the username and password to the Linux machine in the appropriate boxes.
5. Leave other options as is and click **Login**.



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


C:\ - root@192.168.77.154 - WinSCP
Local Mark Files Commands Session Options Remote Help
Default
root@192.168.77.154
C: Local Disk
root
C:\
/root
Name Ext Size Type Changed
Name Ext Size Changed Rights Owner
MSOCache File folder 7/9/2012 12:45:00 ...
My Web Sites File folder 10/11/2012 12:23:...
NakivoBackup
opt
PerfLogs
Program Files
Program Files (x86)
ProgramData
Recovery
System Volume Infor...
Temp
Users
Windows
ДПС Захист звітності
bootmgr 375 KiB
BOOTSECT.BAK
hiberfil.sys
NAKIVO_Backup_&_R... 140 MiB
pagefile.sys 8,124 MiB
.cache
NAKIVO_Backup_&_R... 140 MiB 10/14/2013 11:09:4...
NAKIVO_Backup_&_R... 140 MiB 10/15/2013 10:54:3...
NAKIVO_Backup_&_R... 140 MiB 10/16/2013 6:05:36...
NAKIVO_Transporter_... 25,692 KiB 10/14/2013 9:47:19...
Copy
Copy file 'NAKIVO_Backup__Replication_TRIAL.sh' to remote directory:
/root/*.*
Transfer settings
Default transfer settings
New and updated file(s) only
Do not show this dialog box again
Transfer on background (add to transfer queue)
Transfer each file individually
Transfer settings...
Copy
Cancel
Help
Default
Text
Binary
Custom...
Set as default
Configure...
140 MiB of 14,357 MiB in 1 of 27
0 B of 585 MiB in 0 of 17
F2 Rename F4 Edit F5 Copy F6 Move F7 Create Directory F8 Delete F9 Properties F10 Quit
SCP
0:02:43

9. Click **Copy**.

Installing on Synology NAS

NAKIVO Backup & Replication can be installed directly on a supported Synology NAS to create your own, high-performance backup appliance. With the appliance, all VM data protection components are unified in a single system that is fast to deploy and easy to manage, while also not consuming your environment's valuable resources. Moreover, you are getting an all-in-one backup hardware, backup software, backup storage, and data deduplication in a single box. This results in a zero VMware footprint, less power and cooling, less required maintenance, time, money, and – most of all – higher VM backup performance.
You can install a Synology package with either all NAKIVO Backup & Replication components (**Director**, **Transporter**, **Backup Repository**) or a **Transporter** only. The product can be installed via Package Center or manually. For more details, refer to the corresponding topics below:

- "Installing on Synology NAS via Package Center" on page 163
- "Installing on Synology NAS Manually" on page 165

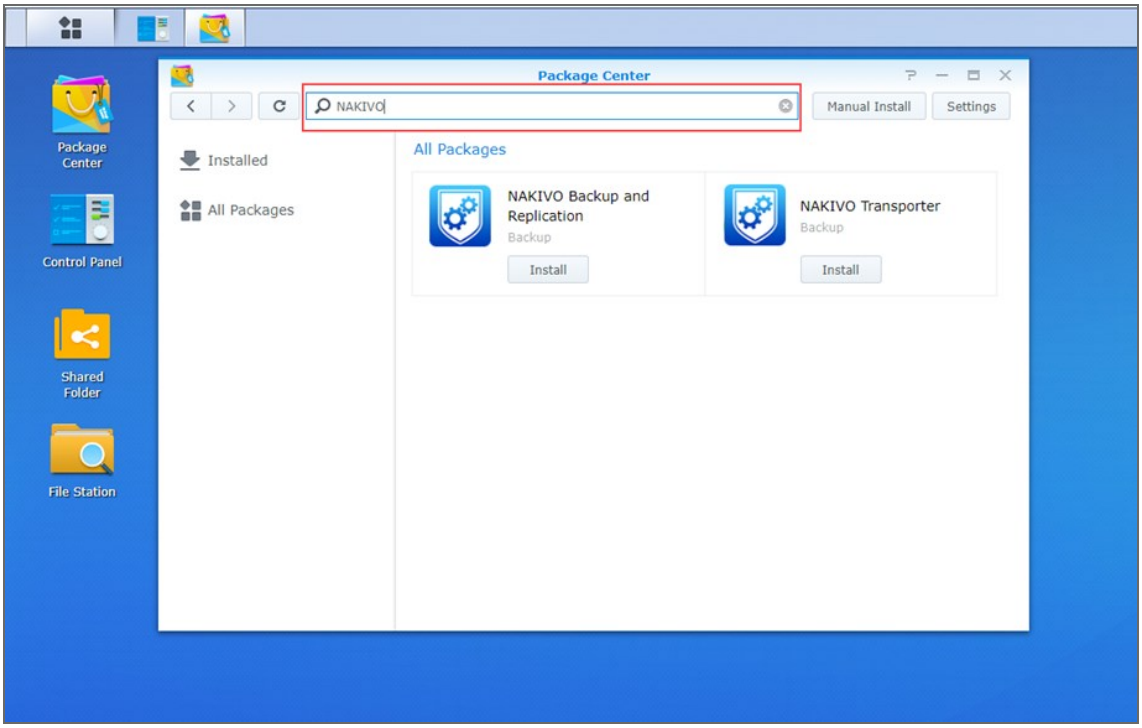
Note

A pre-shared key is not created during **Transporter**-only installation. When adding this Transporter to NAKIVO Backup & Replication, filling out the master password field is not required. The master password can be manually set and reset later. For details, refer to "Installed Service" on page 431.

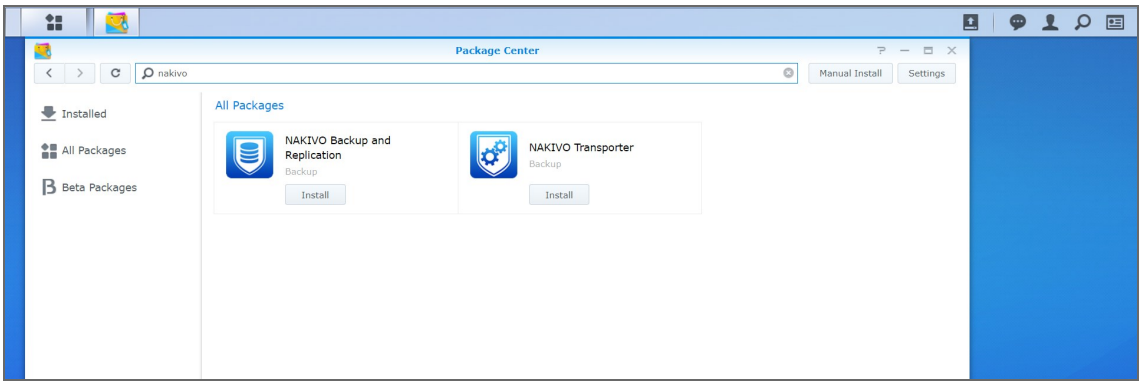
Installing on Synology NAS via Package Center

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.

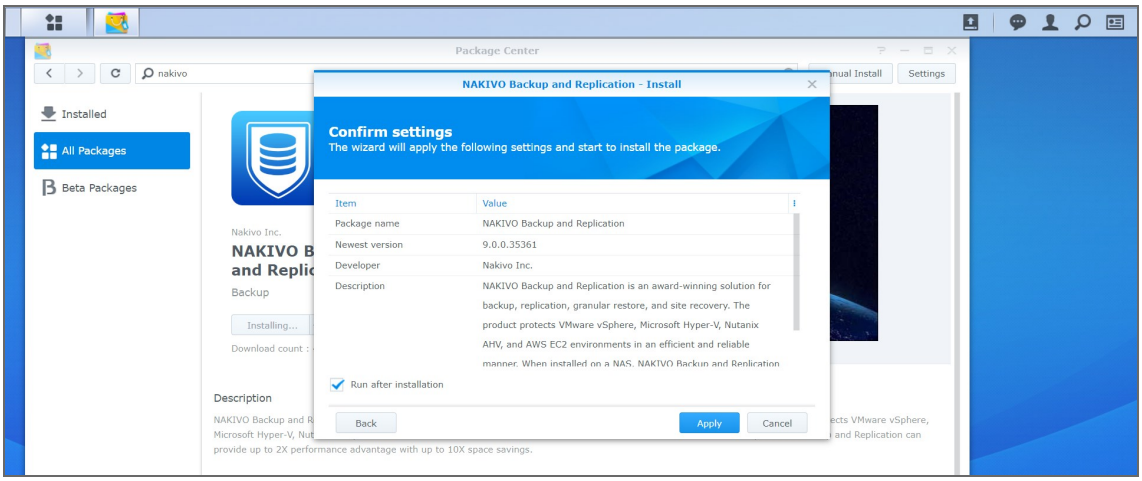


3. Click **Install** on one of the following:
   - **NAKIVO Backup and Replication** to install all product components.
   - **NAKIVO Transporter** to install a **Transporter** only.



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

Installing on Synology NAS Manually

If for any reason installation of NAKIVO Backup & Replication via Package Center is not available for your Synology NAS, you can install it manually.

The following packages are available for manual installation:

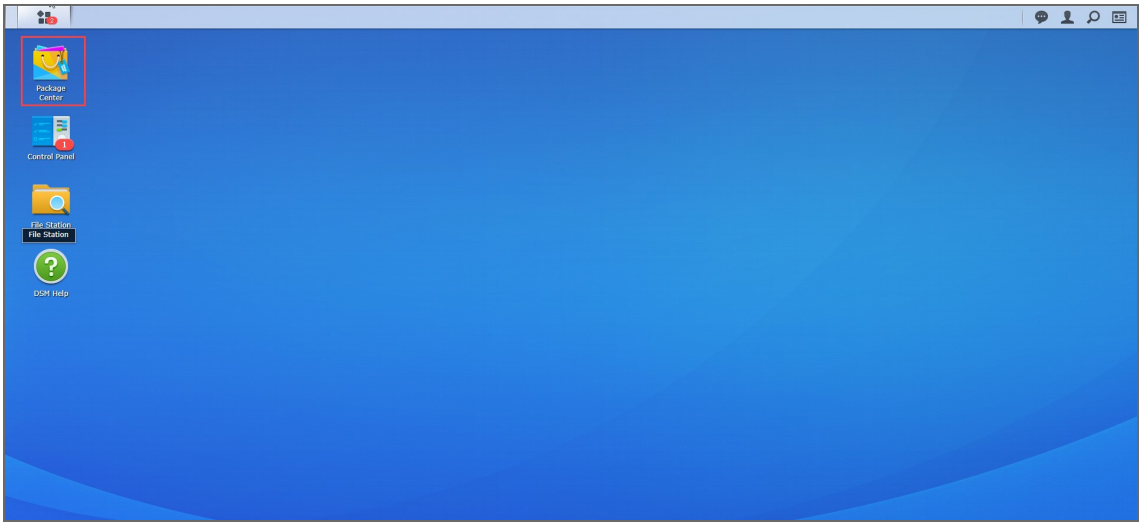
- Synology package
- Synology **Transporter** package
- Synology ARM package
- Synology ARM **Transporter** package

To manually install NAKIVO Backup & Replication on a Synology NAS, do the following:

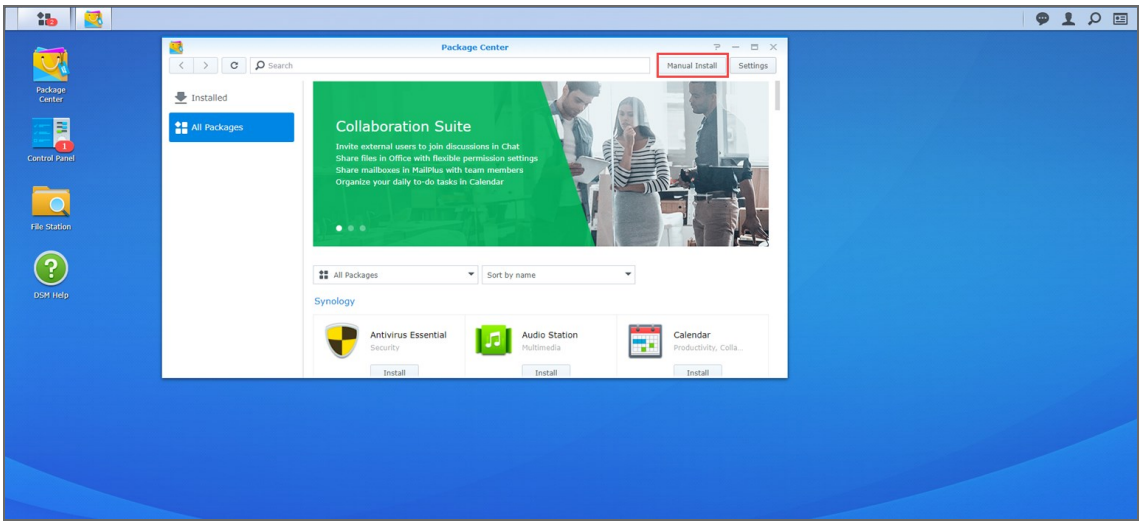
1. Download a Synology NAS package.

> **Note**
>
> Installing the NAKIVO Backup & Replication instance on your NAS-based backup appliance requires accurate knowledge of the hardware system information.
>
> Refer to the following page to learn how to identify your Synology NAS device CPU model: How to Identify the CPU Architecture of a Synology NAS Processor.

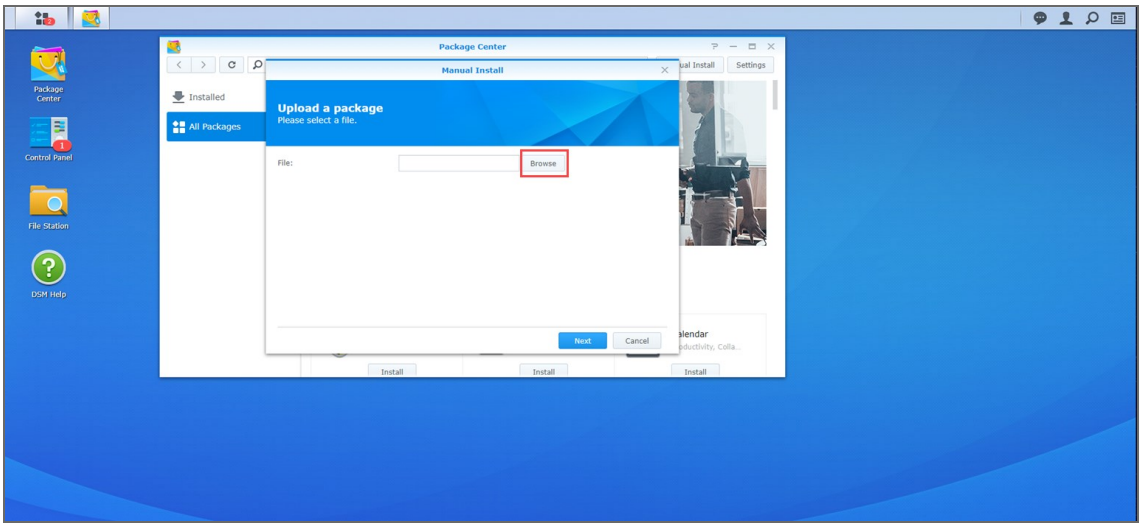
2. Log in to your Synology account and open the **Package Center** in the management interface.



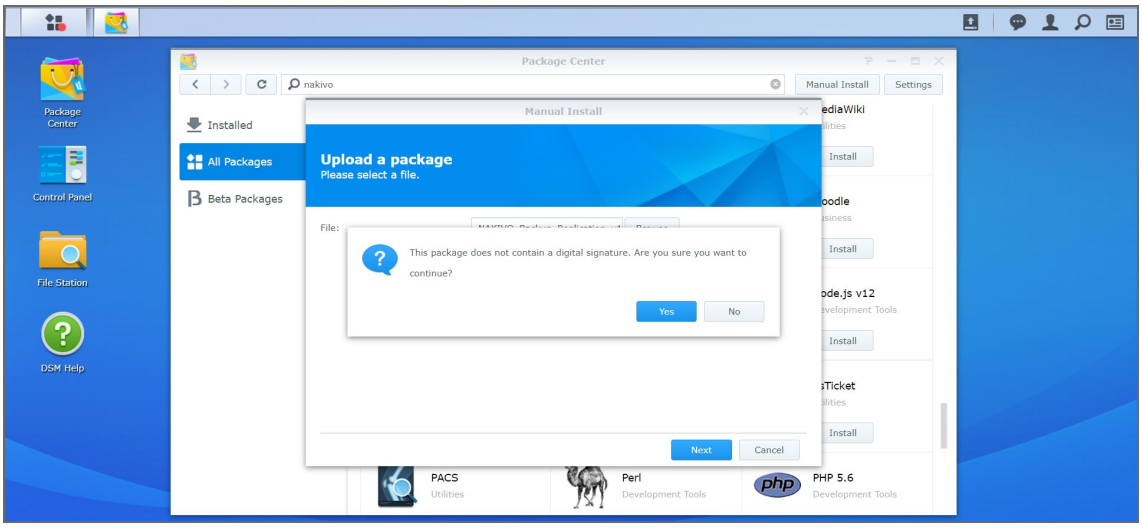
3. Click **Manual Install.**



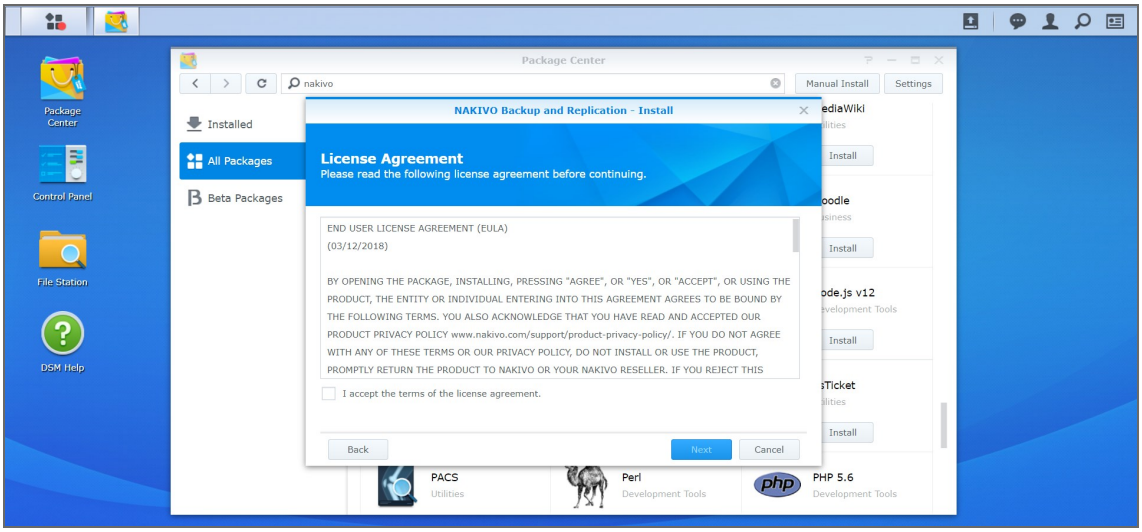
4. Click **Browse**, navigate to the Synology NAS package that you have downloaded, select it, and click **Open**.



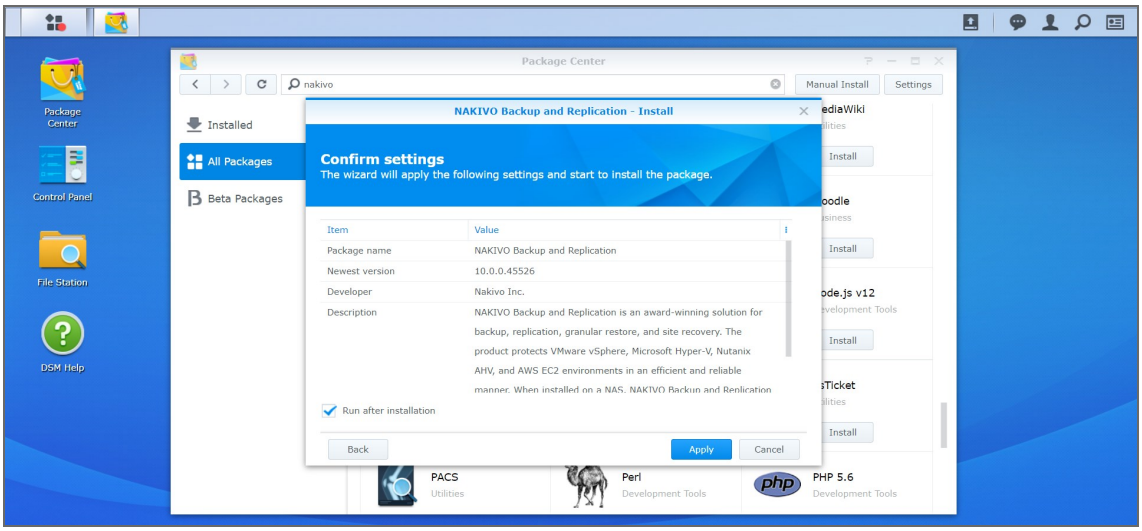
5. Click **Yes** to proceed.



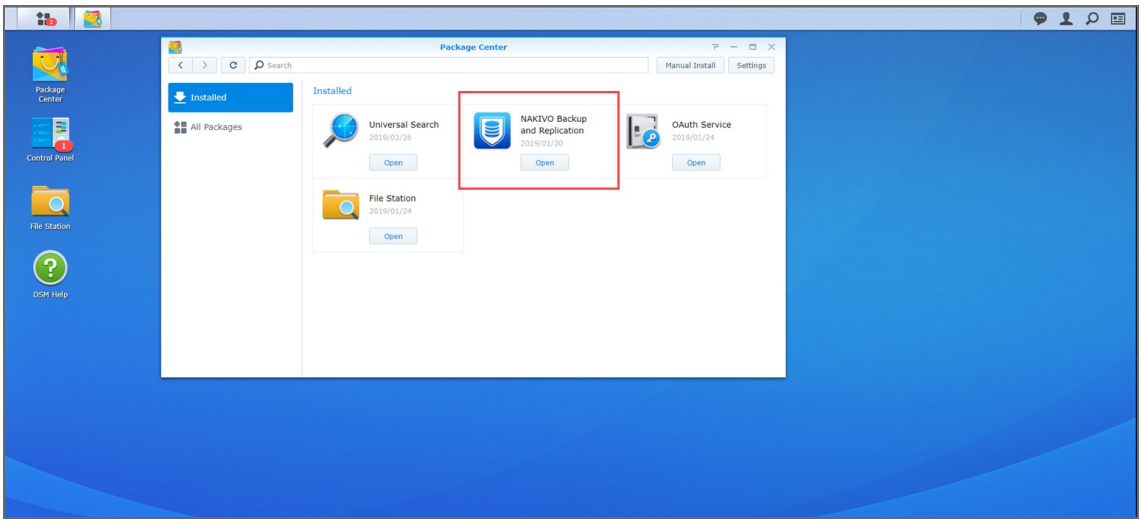
6. After reading through the License Agreement, check **I accept the terms of the license agreement** and click **Next**.



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



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



Refer to “Getting Started” on page 221 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 170
- "Installing on QNAP NAS Manually" on page 172

Note

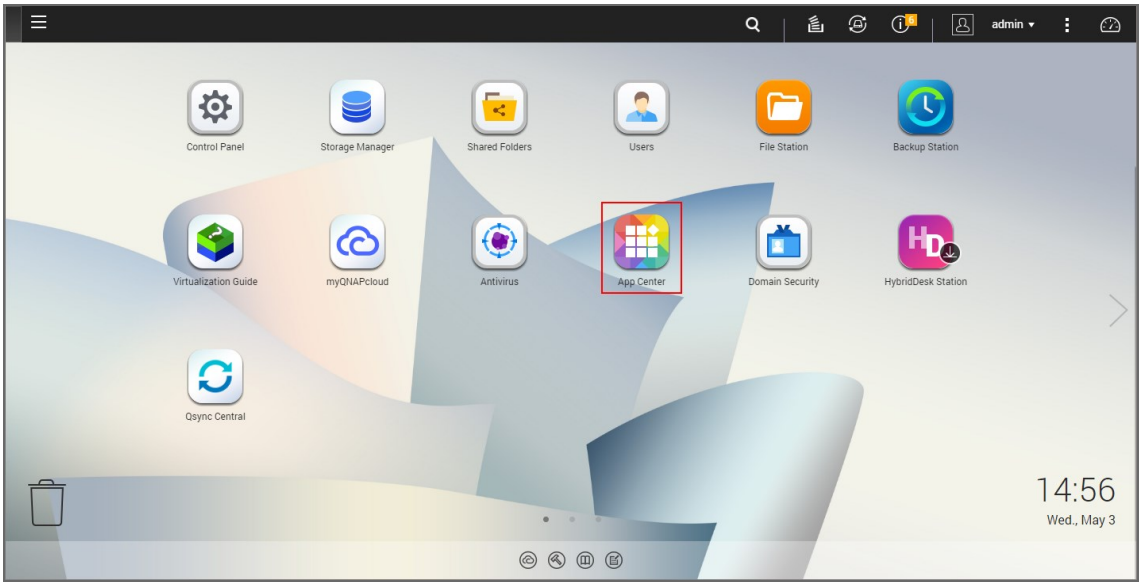
A pre-shared key is not created during **Transporter**-only installation. When adding this **Transporter** to NAKIVO Backup & Replication, filling out the master password field is not required. The master password can be manually set and reset later. For details, refer to "Installed Service" on page 431.

Installing on QNAP NAS via QNAP Store

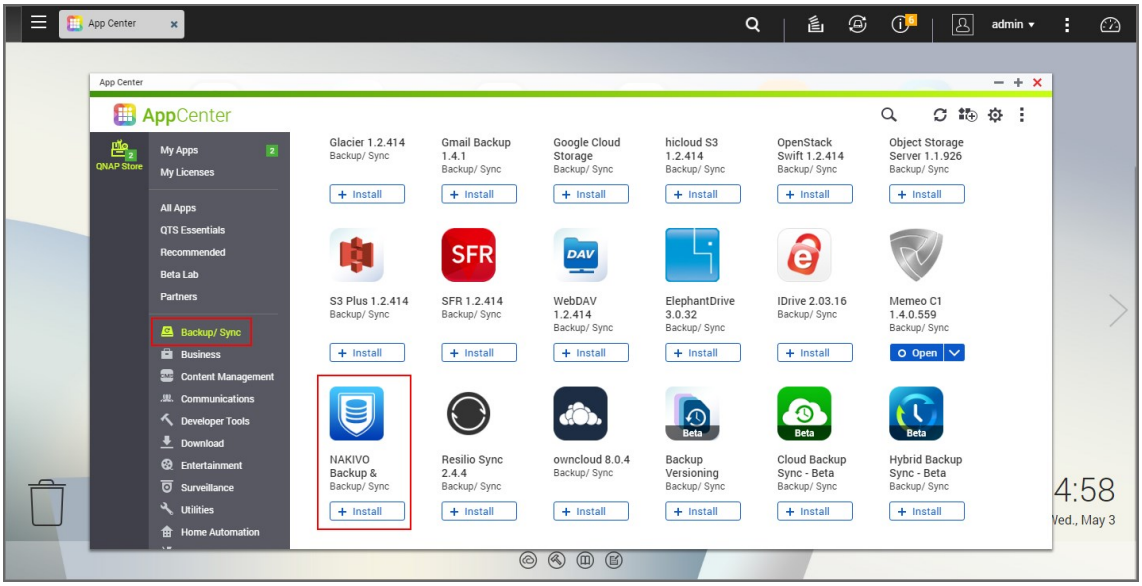
Check to see if your NAS model is supported before you begin installing NAKIVO Backup & Replication on a QNAP NAS.

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 221 to know how to continue working with NAKIVO Backup & Replication.

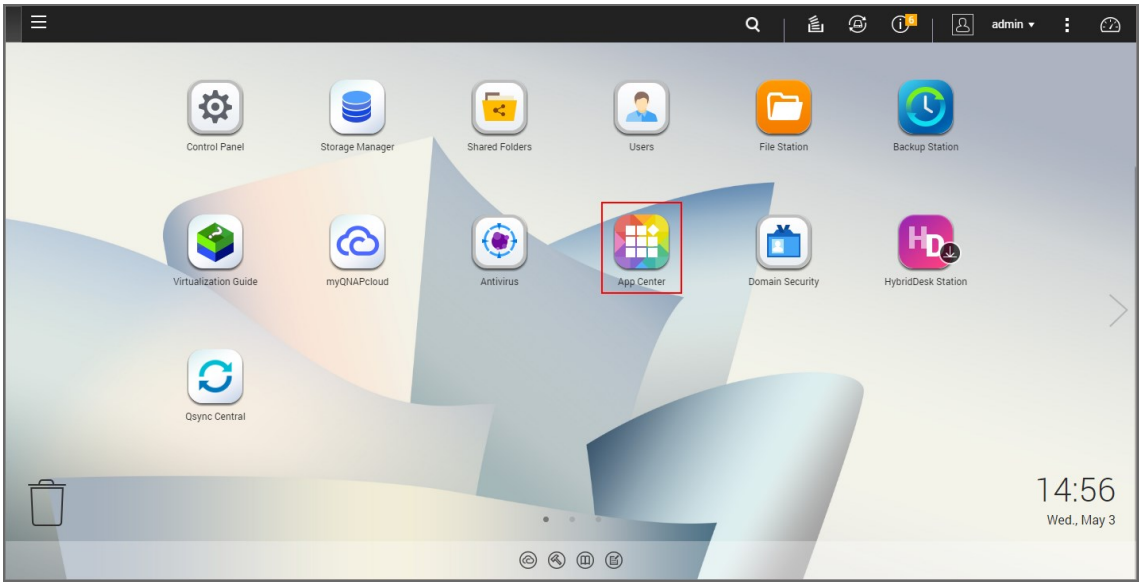
Installing on QNAP NAS Manually

Before you begin installing NAKIVO Backup & Replication on a NAS, make sure your NAS model is supported and you have downloaded the installer (`.qpkg` file) for QNAP NAS.

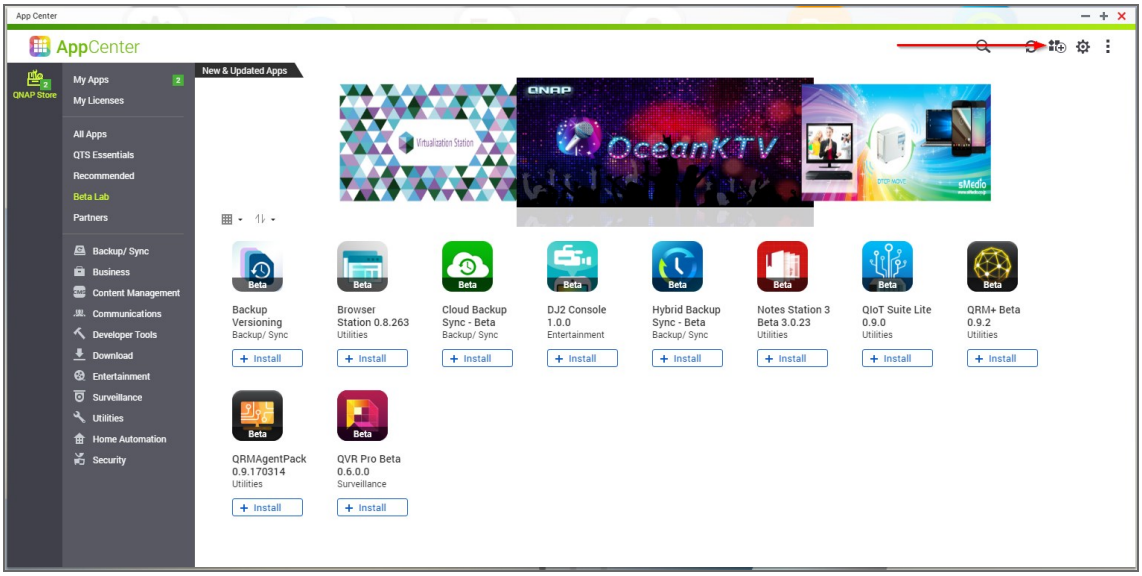
> **Note**
>
> Installing updates of NAKIVO Backup & Replication on your NAS-based backup appliance requires accurate knowledge of the hardware system information.
>
> Refer to the following page to learn how to identify your QNAP device CPU model: How to Identify the CPU Architecture of a QNAP NAS Processor.

To install NAKIVO Backup & Replication on a NAS:

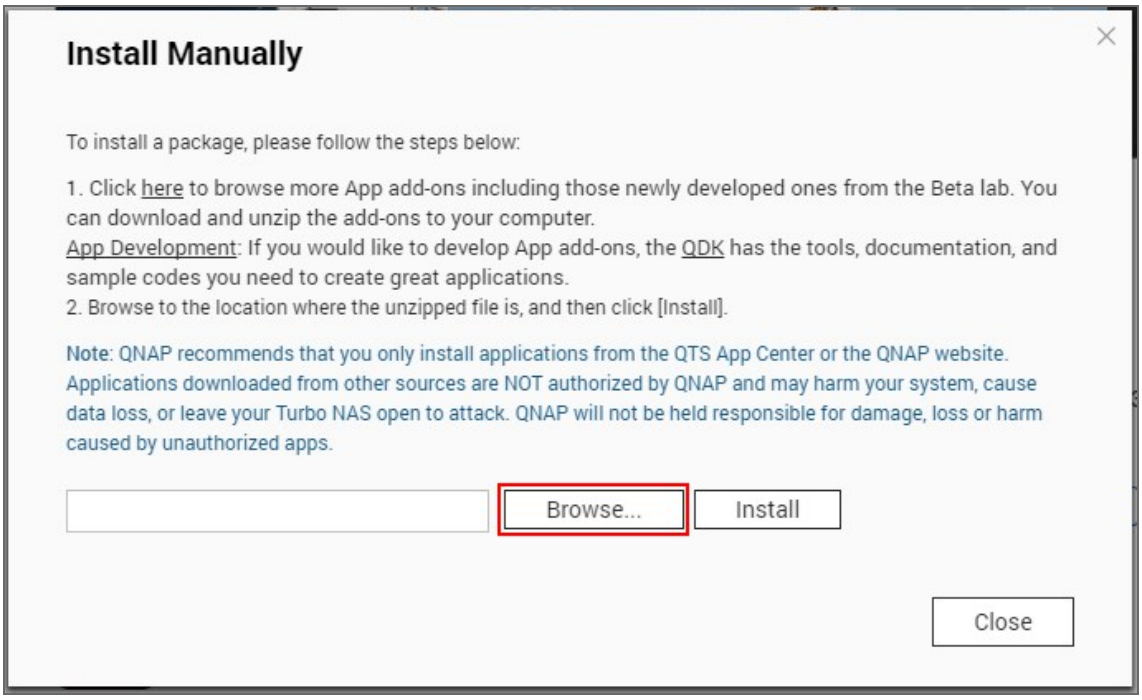
1. Open the QNAP Desktop in your browser by entering the IP address of your QNAP NAS.



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



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



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

By default, NAKIVO Backup & Replication interface is available at the IP address of your QNAP NAS on the port *4443*: `https://<IP_address_of_QNAP_NAS>:4443`.

Refer to "Getting Started" on page 221 to better understand how to continue working with NAKIVO Backup & Replication.

Installing on Western Digital NAS

You can install a Western Digital MyCloud package with either all NAKIVO Backup & Replication components (Director, Transporter, Backup Repository) or a **Transporter** only. The following packages are available:

- Western Digital MyCloud DL2100 package
- Western Digital MyCloud DL2100 **Transporter** package
- Western Digital MyCloud DL4100 package
- Western Digital MyCloud DL4100 **Transporter** package
- Western Digital MyCloud PR2100 package
- Western Digital MyCloud PR 2100 **Transporter** package
- Western Digital MyCloud PR 4100 package
- Western Digital MyCloud PR 4100 **Transporter** package

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

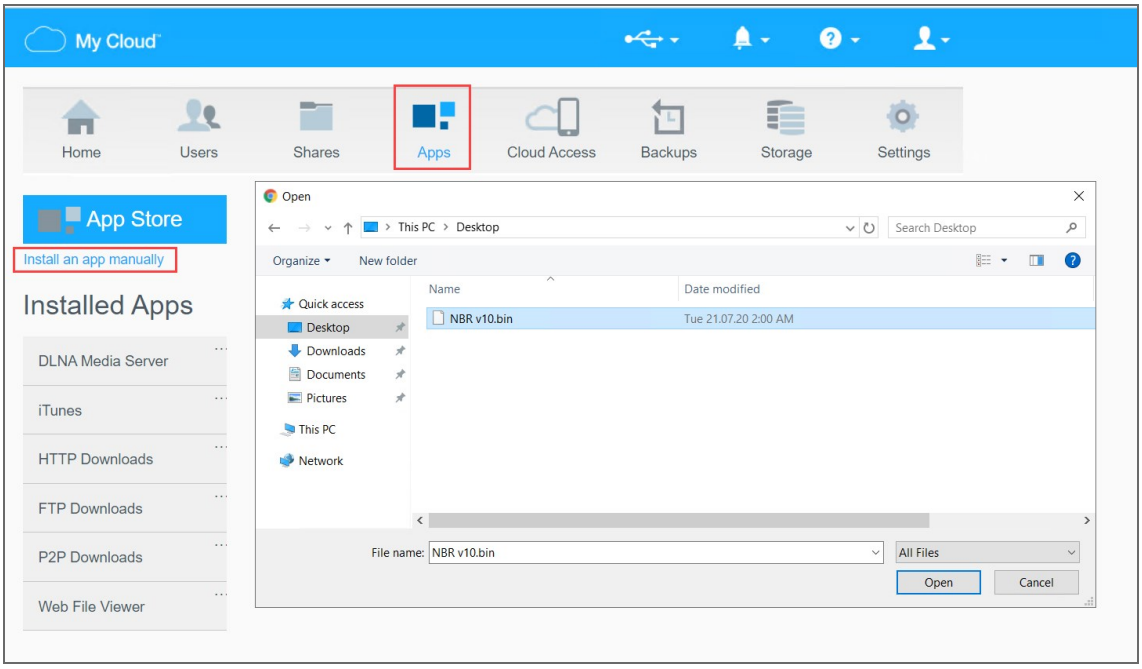
> **Note**
> A pre-shared key is not created during **Transporter**-only installation. When adding this **Transporter** to NAKIVO Backup & Replication, filling out the master password field is not required. The master password can be manually set and reset later. For details, refer to "Installed Service" on page 431.

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 221 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 177
- "Installing on ASUSTOR NAS Manually" on page 179

Note

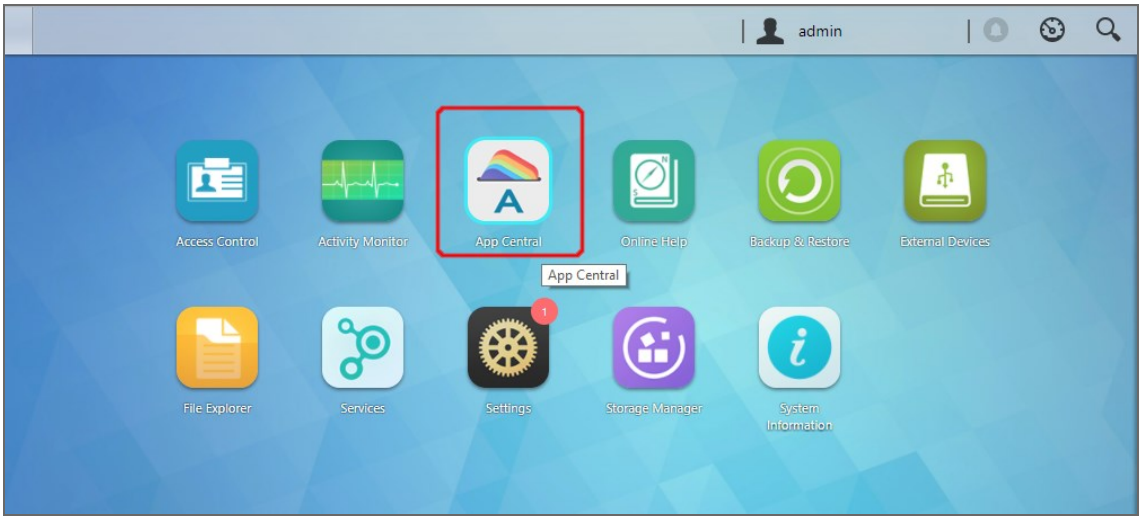
A pre-shared key is not created during Transporter-only installation. When adding this Transporter to NAKIVO Backup & Replication, filling out the master password field is not required. The master password can be manually set and reset later. For details, refer to "Installed Service" on page 431.

Installing on ASUSTOR NAS via App Central

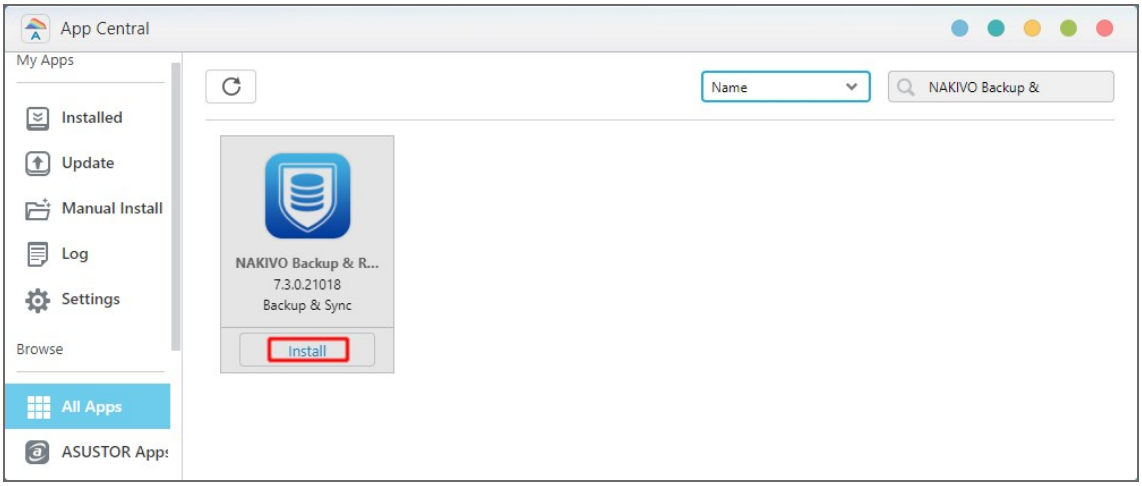
Before you begin installing NAKIVO Backup & Replication on a NAS make sure your NAS model is supported.

To install NAKIVO Backup & Replication on ASUSTOR NAS via App Central:

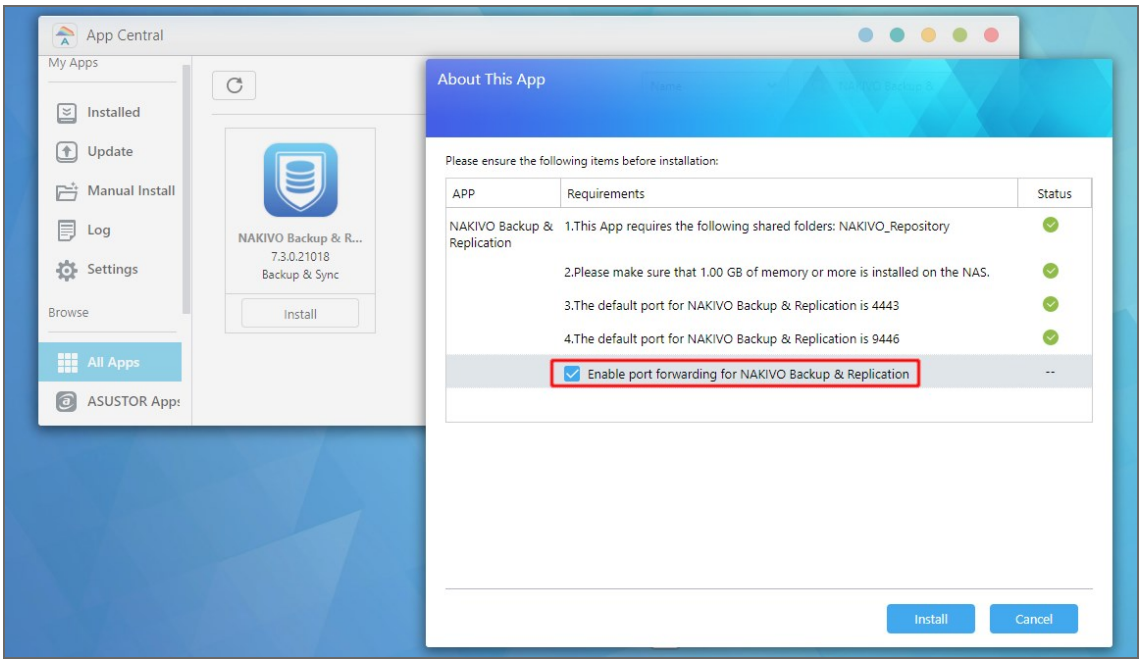
1. Open the ASUSTOR Desktop in your browser by entering the IP address of your ASUSTOR NAS.



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

Installing on ASUSTOR NAS Manually

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

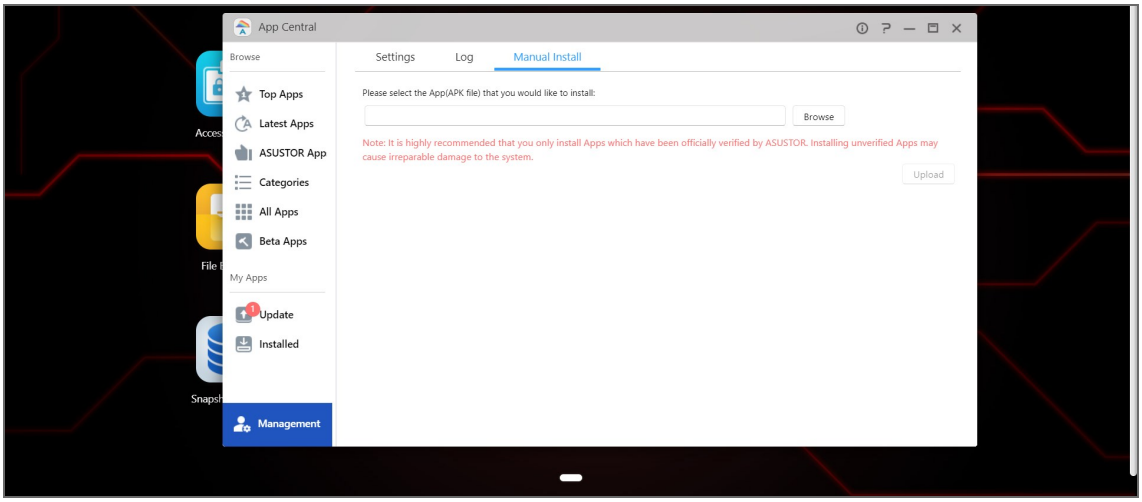
> **Note**
>
> Installing the NAKIVO Backup & Replication instance on your NAS-based backup appliance requires accurate knowledge of the hardware system information.
>
> Refer to the following page to learn how to identify your ASUSTOR device CPU model: How to Identify the CPU Architecture of an ASUSTOR NAS Processor.

To manually install NAKIVO Backup & Replication on ASUSTOR NAS:

1. Open the ASUSTOR Desktop in your browser by entering the IP address of your ASUSTOR NAS.
2. Go to **App Central**.

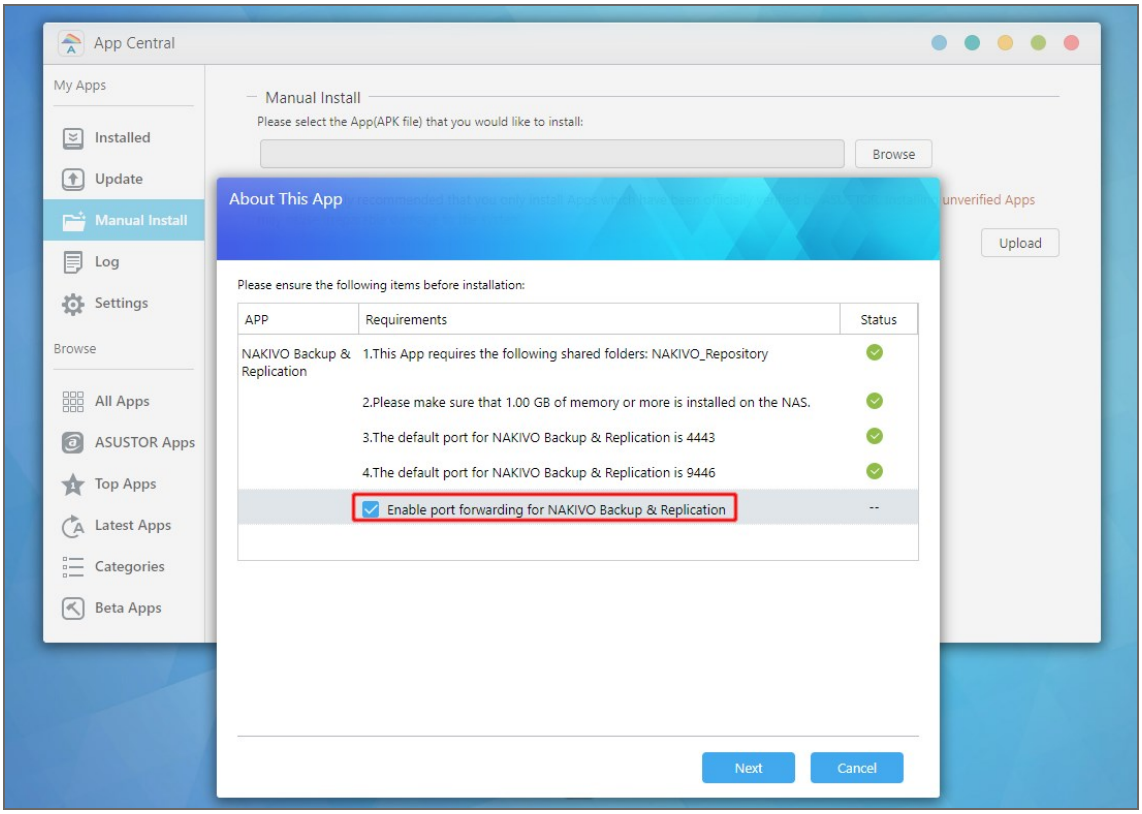


3. Click **Management** > **Manual Install**.



4. Click **Browse**. In the dialog box that opens, locate the installer (.apk file) on your computer.

5. Click **Upload**.
6. In the **About This App** dialog box that opens, check **Enable port forwarding for NAKIVO Backup & Replication**.



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

By default, the NAKIVO Backup & Replication interface is available at the IP address of your ASUSTOR NAS on the port 4443: `https://<IP_address_of_ASUSTOR_NAS>:4443`.

Refer to "Getting Started" on page 221 to understand better how to continue working with NAKIVO Backup & Replication.

Installing on NETGEAR ReadyNAS

You can install the NETGEAR package that includes all NAKIVO Backup & Replication components (Director, Transporter, Backup Repository) or the NETGEAR Transporter package.

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

- “Installing on NETGEAR ReadyNAS via Available Apps” on page 182
- “Installing on NETGEAR ReadyNAS Manually” on page 183

Note

A pre-shared key is not created during **Transporter**-only installation. When adding this **Transporter** to NAKIVO Backup & Replication, filling out the master password field is not required. The master password can be manually set and reset later. For details, refer to “Installed Service” on page 431.

Installing on NETGEAR ReadyNAS via Available Apps

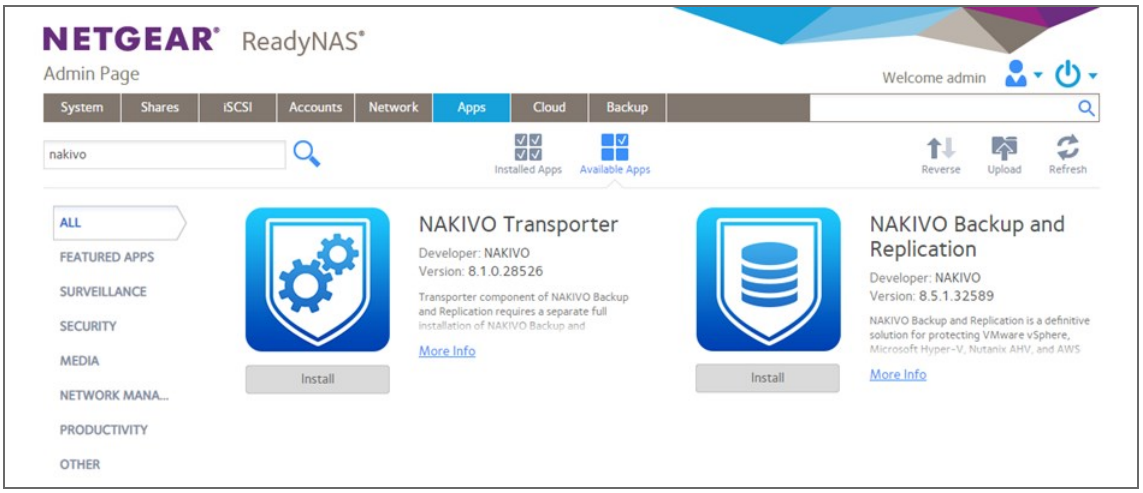
Important

This method is not possible for ReadyNAS OS 6.10.10. To use this method for installing NAKIVO Backup & Replication, you need to use or downgrade to ReadyNAS OS 6.10.9.

Before you begin installing NAKIVO Backup & Replication or NAKIVO Transporter on a NETGEAR ReadyNAS device, please check if your NETGEAR ReadyNAS model is supported.

To install NAKIVO Backup & Replication or NAKIVO **Transporter**, take the following steps:

1. Open the NETGEAR ReadyNAS **Admin Page** in your browser by entering the IP address of your NAS.
2. Go to **Apps** -> **Available Apps**.
3. Find **NAKIVO Backup & Replication** or **NAKIVO Transporter** in the list of available applications. Alternatively, you can enter `NAKIVO` to the filtering box in the upper left corner of the **Admin Page**.
4. Click the **Install** button below the corresponding item.



Note

Make sure that only one instance of the NAKIVO solution - either Full Product or **Transporter**-only - is installed on the device concurrently. Having both products installed at once may lead to incorrect operation.

5. Wait until the installation is completed.

By default, the **NAKIVO Backup & Replication** interface is available at the IP address of your NETGEAR ReadyNAS on the port 4443: `https://<IP_address_of_ NETGEAR_ReadyNAS>:4443`.

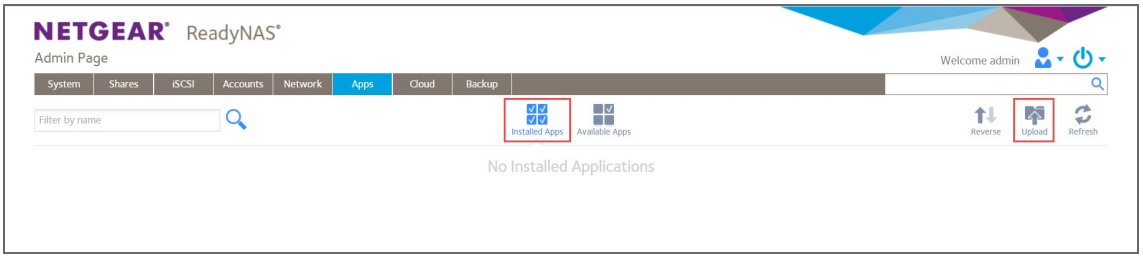
Refer to "Getting Started" on page 221 to know how to continue working with NAKIVO Backup & Replication.

Installing on NETGEAR ReadyNAS Manually

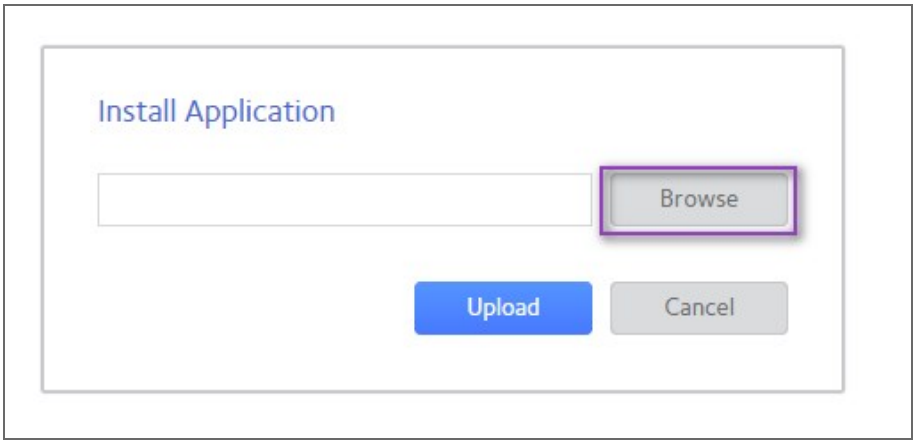
Before you begin installing NAKIVO Backup & Replication or NAKIVO Transporter on a NETGEAR ReadyNAS device, make sure your NAS model is supported and you have downloaded a relevant installer (`.deb` file) for NETGEAR ReadyNAS.

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

Installing on Generic ARM-Based Devices

NAKIVO Backup & Replication can be deployed on ARMv7/ARMv8 computers by downloading and running an appropriate script within a Linux-based OS supported by NAKIVO Backup & Replication.

1. Download a package suitable for your setup from the downloads page.

   Note
   Installing the NAKIVO Backup & Replication instance on your ARM-based backup appliance requires accurate knowledge of the hardware system information.
   Refer to the following page to learn how to identify your ARM-based device CPU model: How to Identify the CPU Architecture of a Generic ARM-Based Device.

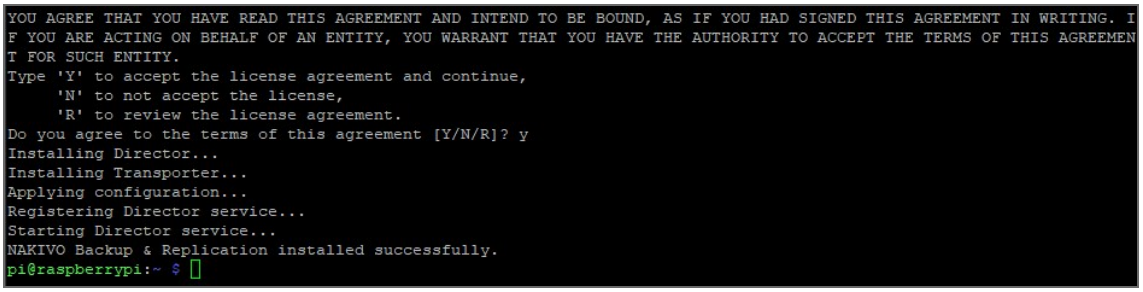
2. Upload the installer file to the machine on which you wish to install NAKIVO Backup & Replication in the binary transfer mode. For example:
   1. Upload the installer from a Windows-based machine
   2. Upload the product from a Linux-based machine: run the following command: wget 'server_ip/shared/NAKIVO Backup_&_Replication v8.5.0.30224 Installer-NAS-ARM-TRIAL.sh'
3. Log in to the Linux machine and allow for the execution of the installer file.

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

4. Execute the installer file with root privileges.

   Example
   sudo ./NAKIVO Backup_&_Replication v8.5.0.30224 Installer-NAS-ARM-TRIAL.sh

5. Accept the License Agreement by typing [Y] and hit **Enter**. You can review the license agreement by typing [R]. Rejecting [N] the license agreement will terminate the installation process and the product will not be installed.
6. The system will notify you when the installation is successfully completed.



Refer to "Getting Started" on page 221 to better understand how to continue working with NAKIVO Backup & Replication.

Installing on TrueNAS

Make sure the following prerequisites are met:

- You have access to the TrueNAS system.
- Your TrueNAS system meets system requirements for installing NAKIVO Backup & Replication.
- The `iocage` jail/container manager is installed on your TrueNAS system. Refer to the `iocage` README page for a description.
- A storage pool is created on your TrueNAS system. Make sure the pool has enough storage for all NAKIVO Backup & Replication functionality. Refer to TrueNAS User Guide for more details on creating storage pools.
    - To create local Repository on TrueNAS local storage outside the jail, make sure you use the following mount points:
      Source: */mnt/share/*
      Destination: */mnt/test/iocage/jails/nbr/root/usr/repo*

    Notes
    - The *repo* is an empty folder created by the user.
    - The path for adding a repository on NAKIVO Backup & Replication is */usr/repo*.
- Make sure to check limitations for TrueNAS here.

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

1. Log in to the TrueNAS system via SSH.
2. Go to the `tmp` folder: `cd /tmp`
3. Download the necessary `json` file:
    - for the full NAKIVO Backup & Replication installation on a TrueNAS v13:
      `wget` https://github.com/NAKIVO/iocage-plugin-nbr/raw/master/nbr.json
    - for the NAKIVO Backup & Replication Transporter installation on a TrueNAS v13:
      `wget` https://github.com/NAKIVO/iocage-plugin-nbr-transporter/raw/master/nbr-transporter.json

    Notes
    If a utility for downloading files like `wget` or `curl` is missing on your TrueNAS system, you can first download the necessary file to your local machine and then upload it to TrueNAS with a third-party tool like WinSCP or FileZilla.

4. Install NAKIVO Backup & Replication with the `iocage` jail/container manager:

Note

Make sure that the jail IP address is not the IP address of your TrueNAS system.

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

Note

Run the following command to select the desired release if it was not found:
`iocage fetch --server ftp-archive.freebsd.org --root-dir /pub/FreeBSD-Archive/old-releases/amd64/`

5. For the NAKIVO Backup & Replication **Transporter** installation, add the **Transporter** to the **Director**. Refer to "Adding Existing Nodes" on page 431 for details.

Updating NAKIVO Backup & Replication

NAKIVO Backup & Replication automatically checks for updates once each day. If an update is available, a notification is displayed in the product web interface. Click the notification link to view information about the update.

Starting from v8.5, a full solution of the NAKIVO Backup & Replication installed on Windows or Linux can be updated automatically. Should you find that product auto updating is not supported or there are some network issues, you can update the product manually. For more details, refer to the corresponding articles below.

Notes

- It is highly recommended that you update your product from any version within its support period. Check NAKIVO Backup & Replication Support Lifecycle Policy for more details.
- No new job runs can be started during the update.
- Updating to a major (e.g. 10.8 to 11.0), minor (e.g. 10.0 to 10.5), or sub-minor version (e.g. 10.0.0 to 10.0.1) will not reset the current license. Check NAKIVO Licensing Policy for more details about product licensing.
- Updating the product running the NAKIVO Backup & Replication version more than N-1 requires an interim update. For example, to update the product from v9.3 (N-2) to v11.0 (N), it first must be updated to v10.x, where N is the current major version.

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

To update your copy of the product to a newer version, you need to download an appropriate updater and run it on:

- Each machine on which you have additionally installed the **Transporter**.
- The machine on which the **Director** is installed.

Refer to the following topics for more information:

- "NAKIVO Backup & Replication Support Lifecycle Policy" on the next page
- "Updating Virtual Appliance" on page 192
- "Updating on Windows" on page 197
- "Updating on Linux" on page 199
- "Updating on Synology NAS" on page 201
- "Updating on Western Digital NAS" on page 204
- "Updating on Amazon EC2" on page 205
- "Updating on QNAP NAS" on page 211

- “Updating on ASUSTOR NAS” on page 214
- “Updating on NETGEAR ReadyNAS” on page 216
- “Software Update” on page 329

NAKIVO Backup & Replication Support Lifecycle Policy

The level of support for various versions of NAKIVO Backup & Replication varies based on their lifecycle phase. The most recent product versions receive full support, which includes updates, hotfixes, and patches. Conversely, older versions may be covered with limited support.

The periods of the NAKIVO Backup & Replication Support Lifecycle are as follows:

- **Hotfix period:** 1 year after release. Within this period, NAKIVO can release hotfixes/patches to fix different product issues.
- **Support period**: 3 years after release. No support will be offered when the support period expires.

Important

It is highly recommended that you update your product from any version within its support period.

Please refer to the table below for the list of NAKIVO Backup & Replication product versions and their lifecycle periods.

Version	Release date	End of Fix	Security Patches	End of Support
11.2	10 March 2026	10 March 2027	10 March 2028	10 March 2029
11.1	6 October 2025	6 October 2026	6 October 2027	6 October 2028
11.0.4	1 September 2025	1 September 2026	1 September 2027	1 September 2028
11.0.3	22 April 2025	22 April 2026	22 April 2027	22 April 2028
11.0.2	17 March 2025	17 March 2026	17 March 2027	17 March 2028
11.0.1	13 January 2025	13 January 2026	13 January 2027	13 January 2028
11.0	4 November 2024	4 November 2025	4 November 2026	4 November 2027
10.11.3	17 September 2024	17 September 2025	17 September 2026	17 September 2027
10.11.2	17 June 2024	17 June 2025	17 June 2026	17 June 2027

10.11.1	17 April 2024	17 April 2025	17 April 2026	17 April 2027
10.11	29 January 2024	29 January 2025	29 January 2026	29 January 2027
10.10.1	6 November 2023	6 November 2024	6 November 2025	6 November 2026
10.10	9 October 2023	9 October 2024	9 October 2025	9 October 2026
10.9	4 July 2023	4 July 2024	4 July 2025	4 July 2026
10.8	23 January 2023	23 January 2024	23 January 2025	23 January 2026
10.7.2	7 November 2022	7 November 2023	7 November 2024	7 November 2025
10.7.1	12 October 2022	12 October 2023	12 October 2024	12 October 2025
10.7	31 August 2022	31 August 2023	31 August 2024	31 August 2025
10.6.1	9 June 2022	9 June 2023	9 June 2024	9 June 2025
10.6	3 May 2022	3 May 2023	3 May 2024	3 May 2025
10.5.1	22 December 2021	22 December 2022	22 December 2023	22 December 2024
10.5	16 November 2021	16 November 2022	16 November 2023	16 November 2024
10.4.1	26 October 2021	26 October 2022	26 October 2023	26 October 2024
10.4	2 August 2021	2 August 2022	2 August 2023	2 August 2024
10.3	4 May 2021	4 May 2022	4 May 2023	4 May 2024
10.2	25 January 2021	25 January 2022	25 January 2023	25 January 2024
10.1.1	30 October 2020	30 October 2021	30 October 2022	30 October 2023
10.1	21 September 2020	21 September 2021	21 September 2022	21 September 2023

10.0.1	26 August 2020	26 August 2021	26 August 2022	26 August 2023
10.0	23 July 2020	23 July 2021	23 July 2022	23 July 2023
9.4	3 June 2020	3 June 2021	3 June 2022	3 June 2023

Important

- All earlier versions of the product not mentioned here are not supported.
- Updating the product (Director and Transporter) to a new version is prohibited in case the current version is two or more major versions older than the new version. For example: you cannot upgrade your product from 9.x to 11.x.
- Updating any version of the product to any Beta version of the product is prohibited.
- Updating any Beta version to any version of the product is prohibited

Product Lifecycle Terminology and Milestones

Versioning Scheme

NAKIVO classifies releases as major or minor:

- A major release contains significant new features and/or structural changes. Changes may also include deprecating existing features and breaking backward compatibility.
- A minor release may include some new functionalities and changes that enhance performance and usability. These releases preserve backward compatibility within the related major release.

For information about changes in different releases, refer to the Release Notes.

Release Lifecycle

The lifecycle of major and minor NAKIVO Backup & Replication releases consists of 3 phases:

- **Phase 1 - Early Access / Beta**: Releases in this phase are meant for testing and feedback purposes and should not be used in production environments.
- **Phase 2 - General Availability**: Releases in this phase are fully supported and can be used in production environments.

Notes

- These releases may receive minor fixes within the first year from the date of general availability. During first year after GA, fixes are provided to the current version. After 1 year passed, for the next 2 years fixes are provided to the version which is on 1st year.
- Support in this case means that NAKIVO continues to deliver technical support and updates for the duration of that period.
- For information on technical support services available to you if you have a valid NAKIVO Backup & Replication subscription and/or perpetual license, refer to the NAKIVO Technical Support Policy.

- **Phase 3 - End of Support**: End of support for a minor or major release means that NAKIVO will no longer provide technical support or issue updates, patches and fixes for a specific version.

A version reaches end of support after 3 years from the date of general availability. End of support for a version applies to all solution editions and license types, including NFR licenses.

If your version is nearing the end of support, we recommend that you plan an update to a newer NAKIVO Backup & Replication version to avoid support gaps and maintain optimal data protection. You can update your instance in the product UI or download an updater from the NAKIVO website.

Notes

- During the general availability period, if you have a valid perpetual and/or subscription license, you can update your NAKIVO Backup & Replication version to any newer supported version.
- NAKIVO does not provide patches or fixes for third-party technology and services that reach end of life within the 3-year support period.

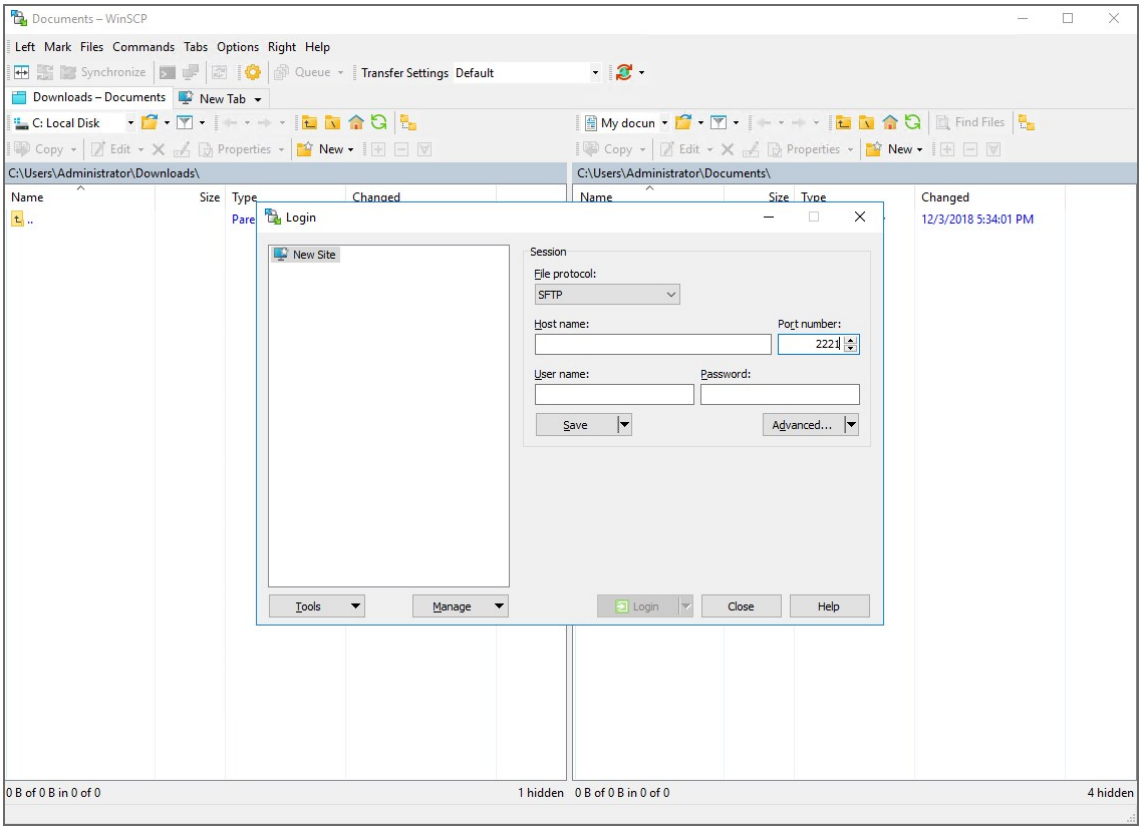
Updating Virtual Appliance

Prior to updating your virtual appliance (VA):

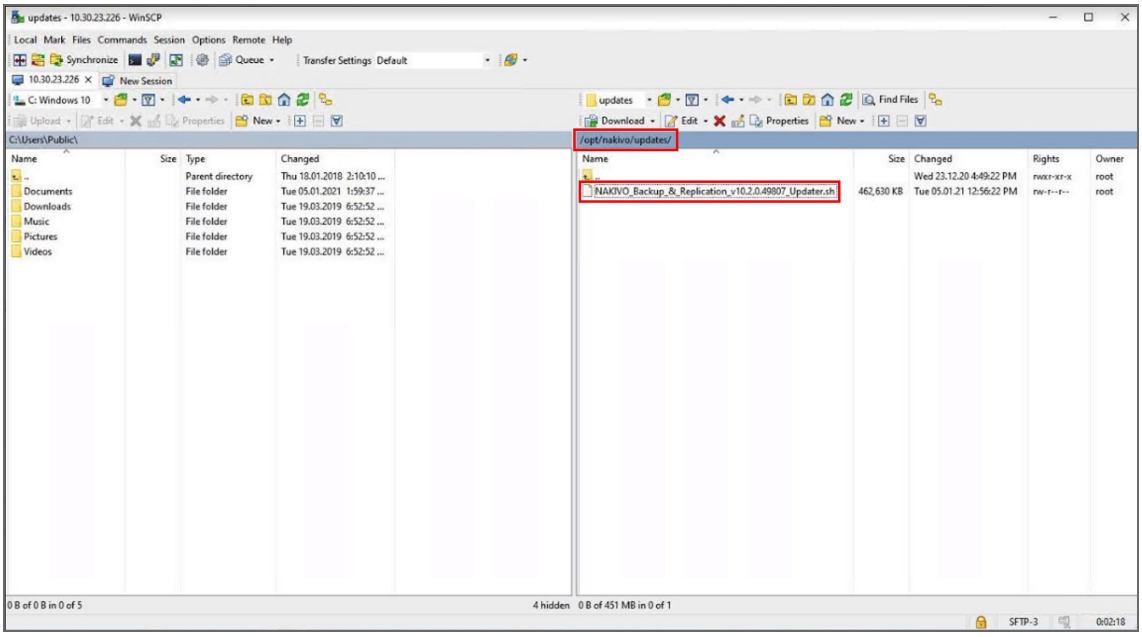
1. Make sure that no jobs or repository maintenance tasks are running in the product.
2. Create a snapshot of the VA to revert to the previous version in case any failure occurs.

Follow the steps below to update your VA:

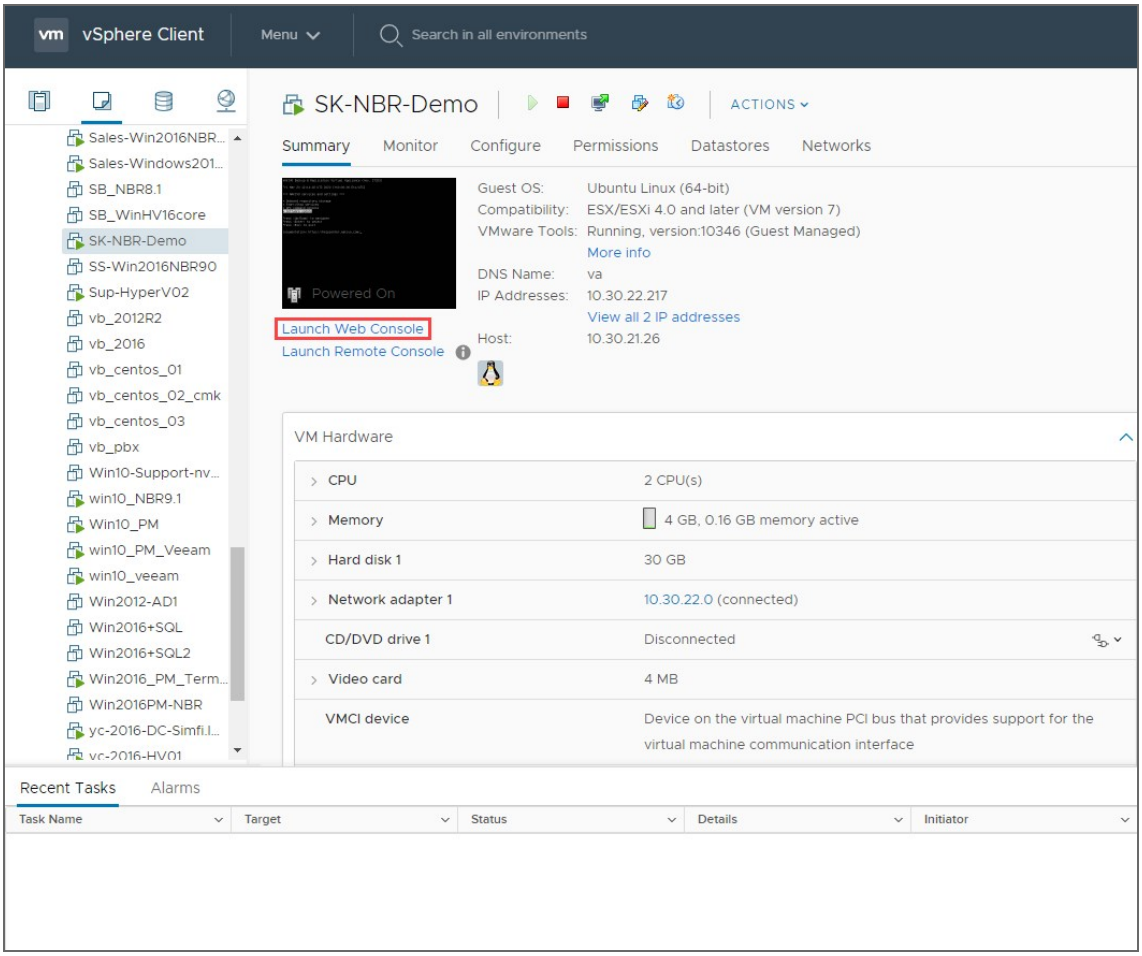
1. Using SSH client, log in to the VA that needs to be updated.



2. Download the latest *VA and Linux updater* from here.
3. Change the directory to `/opt/nakivo/updates` and locate the updater.

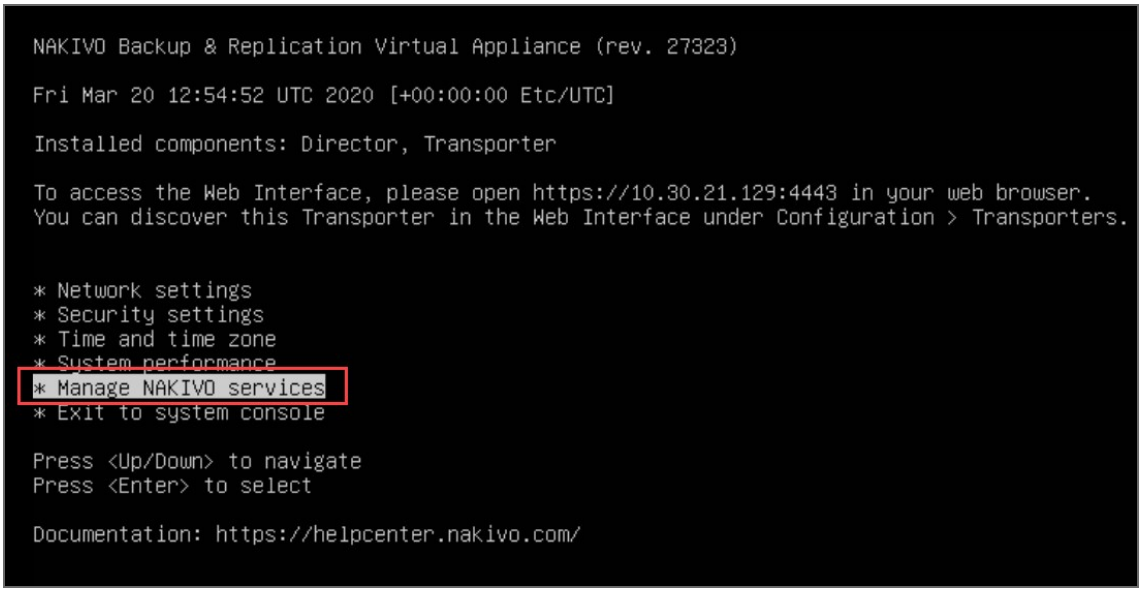


4. Log out from the SSH client.
5. Log in to your vSphere client, navigate to your VA and click **Launch Web Console**.

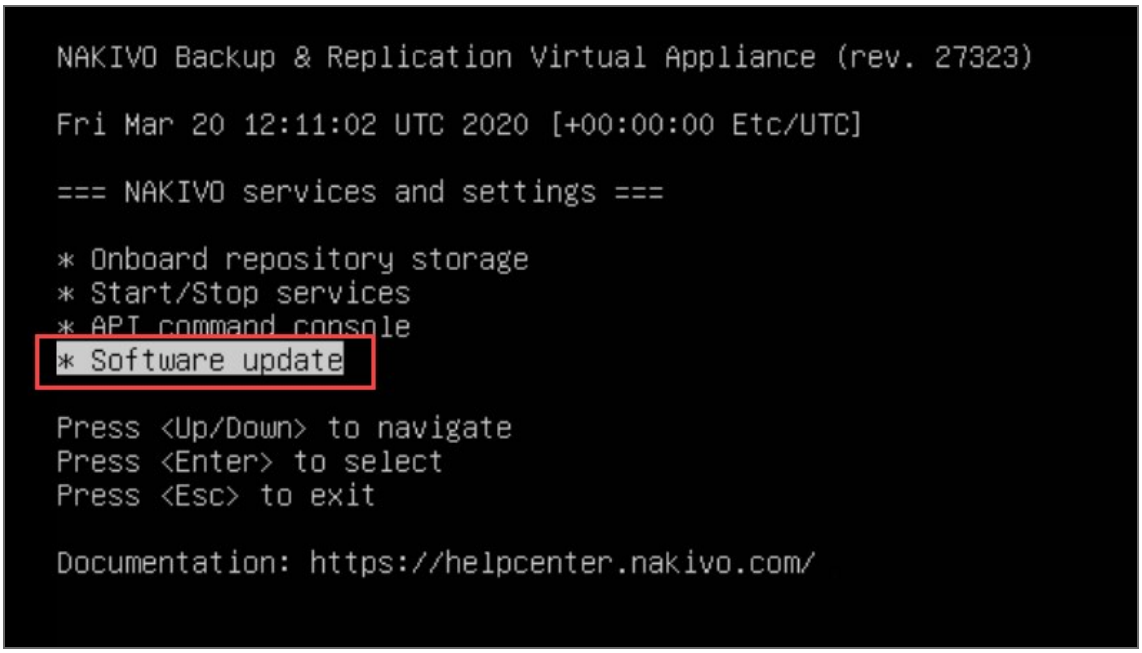


6. Do one of the following depending on the NAKIVO Backup & Replication version you use:

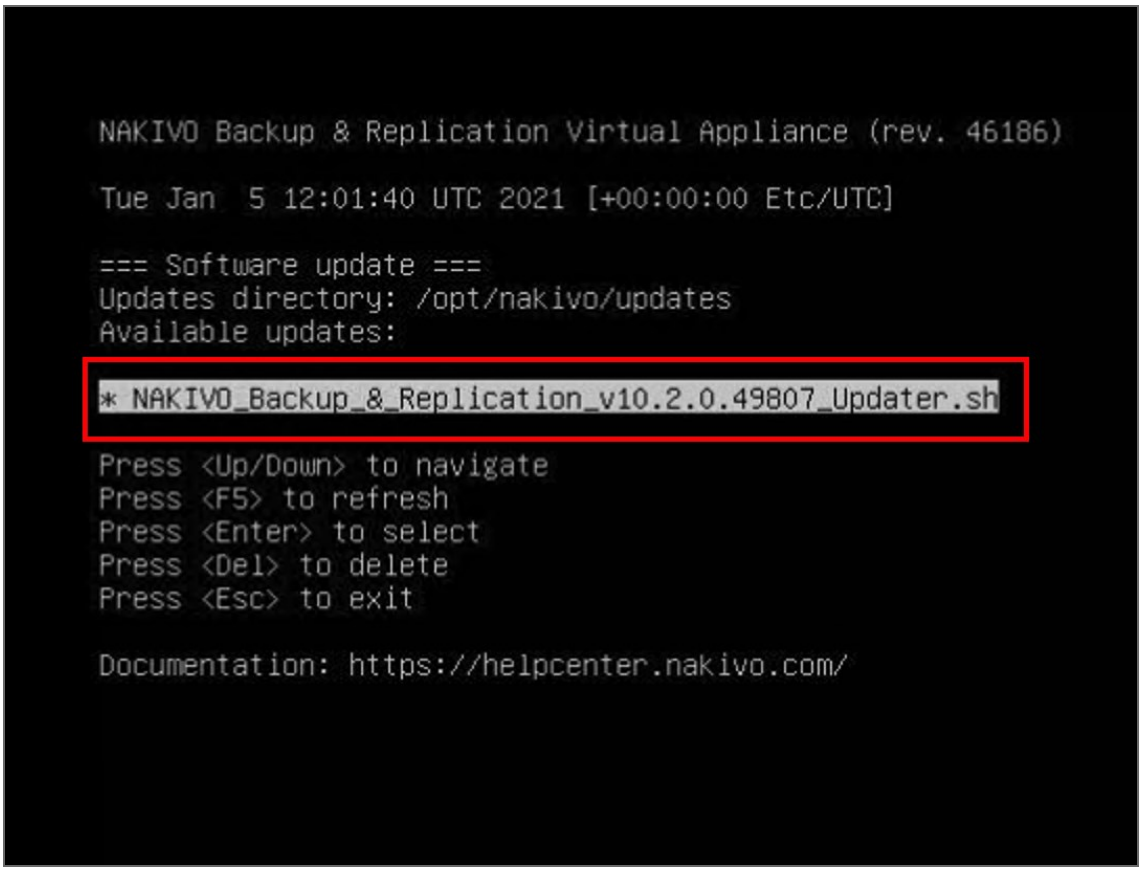
- For the product Version 8.1 and above:
  1. In the VA menu, select **Manage NAKIVO services** and press **Enter**.



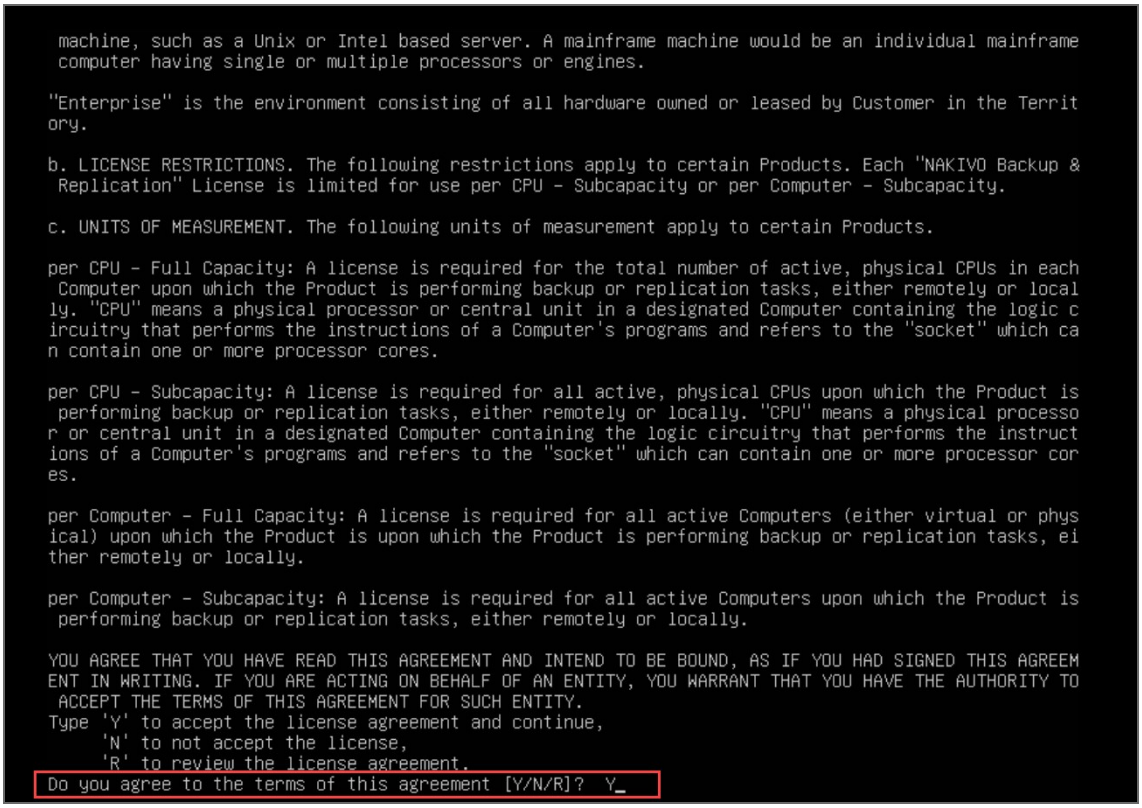
  2. In the menu that opens, select **Software update** and press **Enter**.



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



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



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

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

Note

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

Updating on Windows

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

1. Download the latest Windows updater from here.
2. Make sure that no jobs or repository maintenance tasks are running in the product.
   If NAKIVO Backup & Replication is installed on a VM, create a snapshot of the VM before updating the product.
3. Run the updater on the machine on which the Director is installed, and also on all machines on which you have additionally deployed a **Transporter**.
4. Optionally, you can select the **Master password** checkbox and enter the password that will be used to generate a pre-shared key and secure the **Transporter**. This option is available only for the **Transporter**-only update.

Notes

- The master password must adhere to the following requirements:
  - Length between 8 and 511 characters
  - At least one uppercase (A-Z) and one lowercase letter (a-z)
  - At least one digit (0-9)
  - At least one special character (e.g., !@#$%^&*()_+-=[]{}|;:'",.<>?/)
  - Standard English characters and spaces only (no emojis or diacritics)
- 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.
- When updating NAKIVO Backup & Replication components using console or silent installation modes on Windows, password input is subject to Windows command-line parsing rules:
  - Passwords containing spaces must be enclosed in double quotation marks. For example: -b "Abcd 1ljk".
  - Special characters such as !, &, |, (, and) must be escaped using the caret (^). For example: -b password1^!A.
  - Passwords containing both spaces and special characters must be enclosed in double quotation marks, with special characters escaped. For example: -b "Abcd 1^!jk".

5. Optionally, you can select the **Transporter certificate** checkbox. This allows you to use a CA Certificate. Enter the path to the folder containing the certificate file in the field.

   Notes

   - When the checkbox is not selected, NAKIVO Backup & Replication automatically installs a self-signed certificate.
   - If the **Transporter Certificate** checkbox is not selected, a warning window appears prompting you to install it. Click **Continue** to proceed.

6. Click **Update**.
7. When the update is complete, click **Finish**.
8. If you have entered the new master password on step 4, do the following:
   a. Go to **Settings > Transporters** and click on the **Transporter** you have changed the master password for.
   b. Select **Edit**.
   c. Enter the new master password and click **Connect**.
   d. The **Certificate Acceptance** dialog box appears. Verify the certificate details, and click **Accept**.
   e. Click **Apply** to save the changes.
   f. Click on the same **Transporter** once again and select **Refresh** to refresh the **Transporter**.

Updating on Linux

If updating on a Linux OS within the NAKIVO Backup & Replication interface is not supported, follow the steps below to update the product manually:

1. Download the latest Linux/VA updater from here.
2. Upload the updater to the machine on which the **Director** is installed.

   Important
   Make sure you are using the *binary transfer mode* when uploading the updater to the machine with a Linux OS. For example:

   - Upload the installer from a Windows-based machine
   - Upload the product from a Linux-based machine: run the following command: `wget 'server_ip/shared/NAKIVO_Backup_Replication_vX.X.X_Updater.sh'`
3. Log in to the Linux machine and allow the execution of the updater file. For example: `chmod +x NAKIVO_Backup_Replication_vX.X.X_Updater.sh`
4. Make sure that no jobs or repository maintenance tasks are running in the product.
   If NAKIVO Backup & Replication is installed on a VM, create a snapshot of the VM prior to updating the product.
5. Run the updater file with root privileges. For example: `sudo ./NAKIVO_Backup_Replication_vX.X.X_Updater.sh`
6. Review the license agreement (press **Space** to go to the next page of the agreement). If you agree to the terms of the license agreement, press “Y” and then press **Enter**.
7. Enter the “Y” key and then press **Enter** to confirm that you wish to stop the services and begin the update process.
8. Update all machines on which you have additionally deployed a “Transporter” on page 51.

Note

In some rare cases, CIFS/NFS shares may remain mounted on the OS level even after detaching a repository in the NAKIVO web UI.

To avoid risks of damaging CIFS/NFS-based repositories during the update, follow the steps below:

1. Make sure there is no activity in NAKIVO.
2. Stop NAKIVO transporter services assigned to the attached backup repositories.
3. Check if any active mounts are left.
4. Run the `umount` command.
5. Run the NAKIVO update.

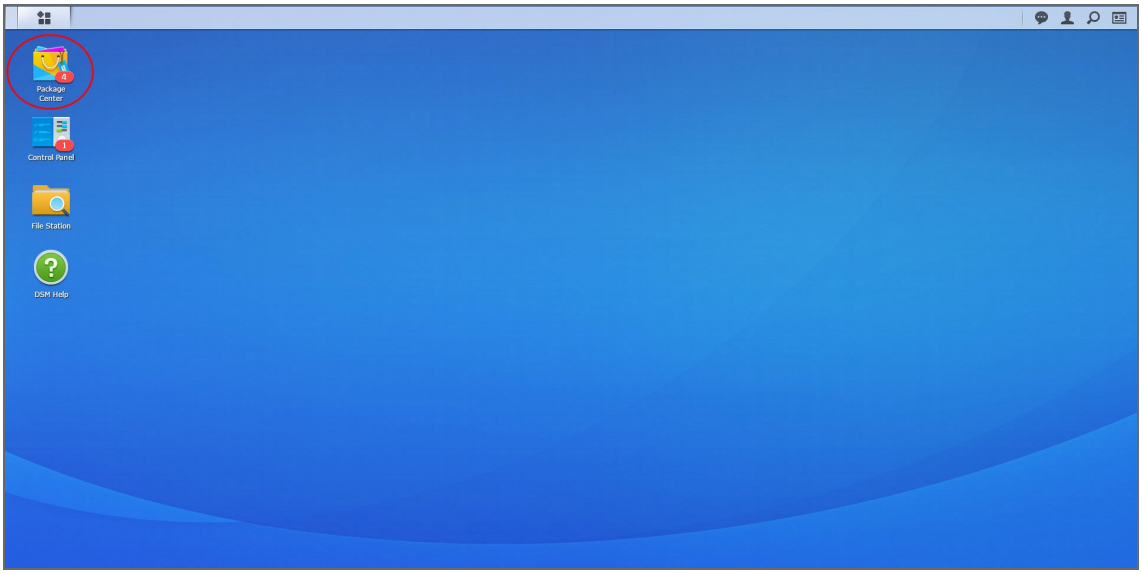
The same flow is recommended for M365 repositories.

Updating on Synology NAS

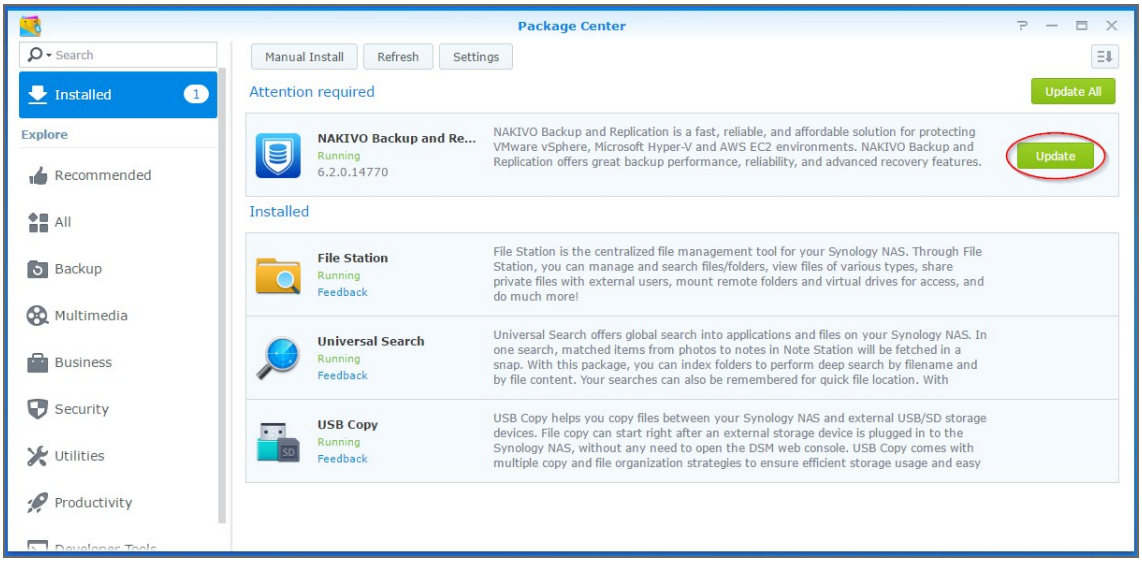
- Updating via Synology Package Center
- Updating Manually

Updating via Synology Package Center

1. Make sure that no jobs or repository maintenance tasks are running in the product.
2. In the Synology NAS management interface, open the **Package Center**.



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



5. Click **Update**.

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

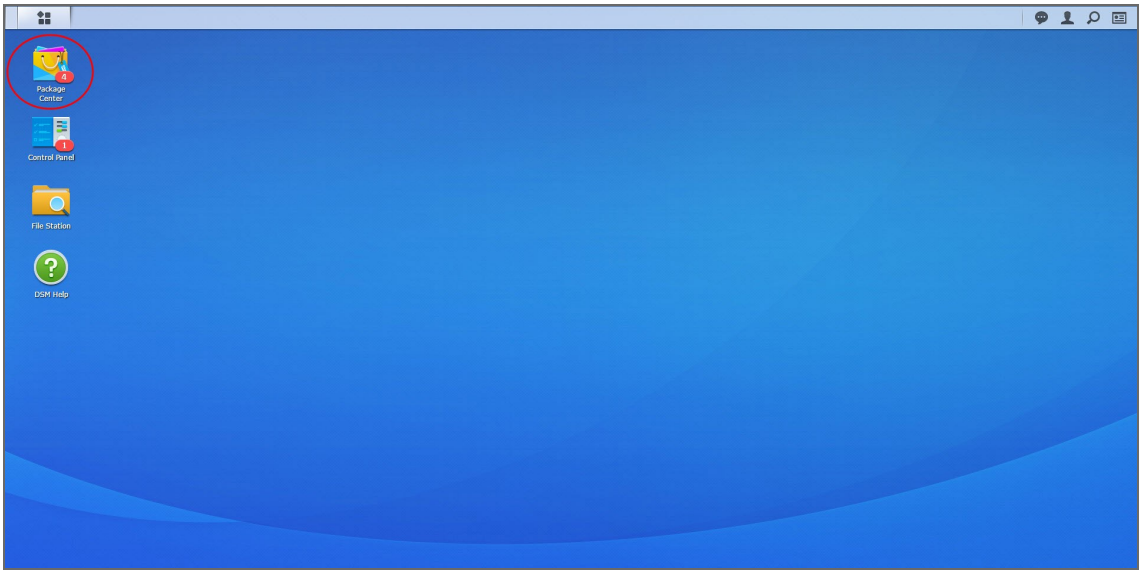
> **Note**
> If the latest version of NAKIVO Backup & Replication is not available in the **Synology Package Center**, you may update manually by following the instructions in this Knowledge Base article.

Updating Manually

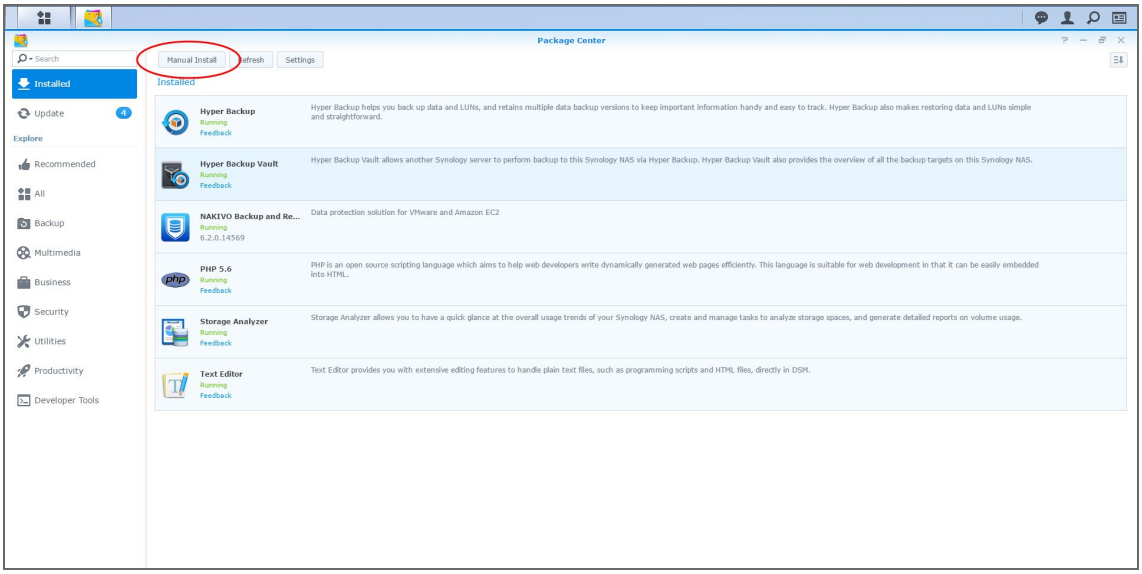
1. Download the latest Synology NAS updater from here.

> **Note**
> Installing updates of NAKIVO Backup & Replication on your NAS-based backup appliance requires accurate knowledge of the hardware system information.
> Refer to the following page to learn how to identify your Synology device CPU model: How to Identify the CPU Architecture of a Synology NAS Processor.

2. Make sure that no jobs or repository maintenance tasks are running in the product.
3. In the Synology NAS management interface, open the **Package Center**.



4. Click **Manual Install**.



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

Now, NAKIVO Backup & Replication has been updated.

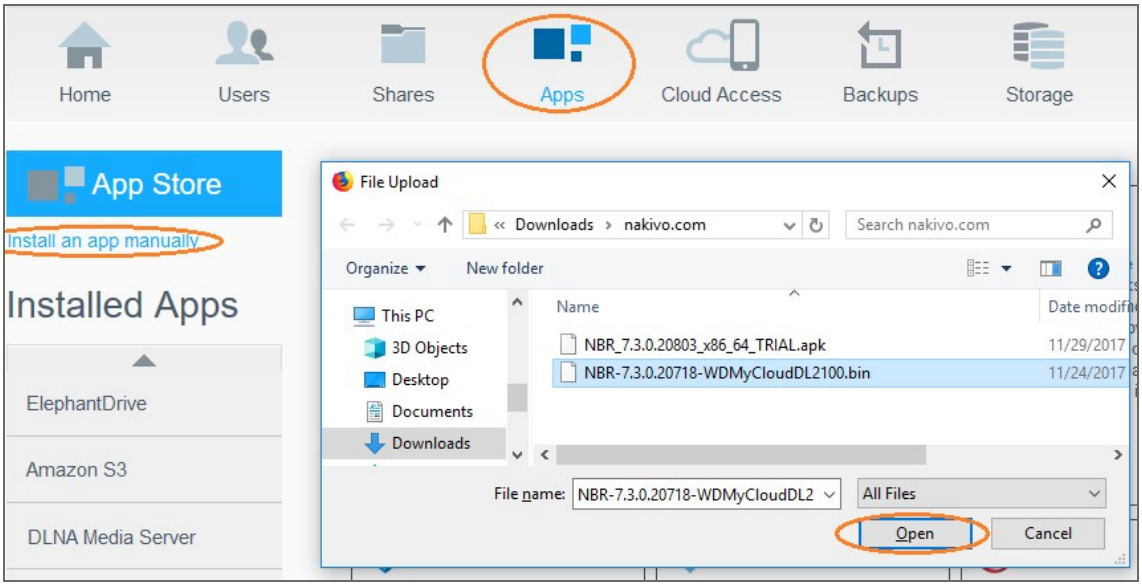
Updating on Western Digital NAS

Prior to updating NAKIVO Backup & Replication on Western Digital MyCloud NAS, make sure the following requirements have been met:

- You have access to the Western Digital NAS MyCloud Dashboard.
- NAKIVO Backup & Replication installer is available for your Western Digital NAS.

Please follow the steps below to update NAKIVO Backup & Replication on a Western Digital MyCloud NAS device:

1. Make sure that no jobs or repository maintenance tasks are running in the product.
2. In the **My Cloud** Dashboard, click **Apps**. The list of installed NAS applications opens on the left side of the page.
3. Above the list of NAS installed applications, click **Install an app manually**. The **File Upload** dialog opens.
4. In the **File Upload** dialog, navigate to your copy of the NAKIVO Backup & Replication installer for Western Digital NAS and click **Open**. The update progress bar opens.



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

Updating on Amazon EC2

The main installation of NAKIVO Backup & Replication (**Director** and **Transporter**) must be updated the way it is done on Linux.

Notes

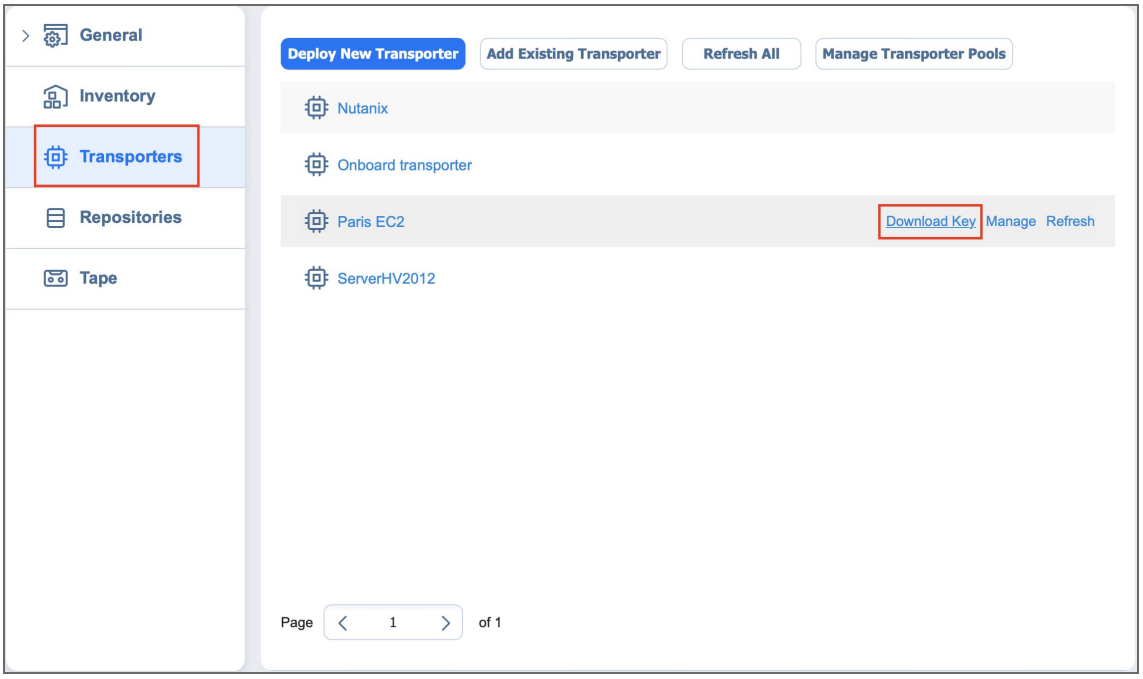
- You have to apply the **-e** argument for executing the installer, in order to avoid changing the Amazon EC2 **Transporter** with the regular Linux **Transporter**. Refer to "Installing on Linux" on page 152 for a description of the available arguments.
- Only the main installation of NAKIVO Backup & Replication needs to be updated manually. **Transporters** installed on Amazon EC2 instances are updated automatically.

Connecting to Amazon EC2 from Windows

You can use the following free tools to connect to your Amazon EC2 instance:

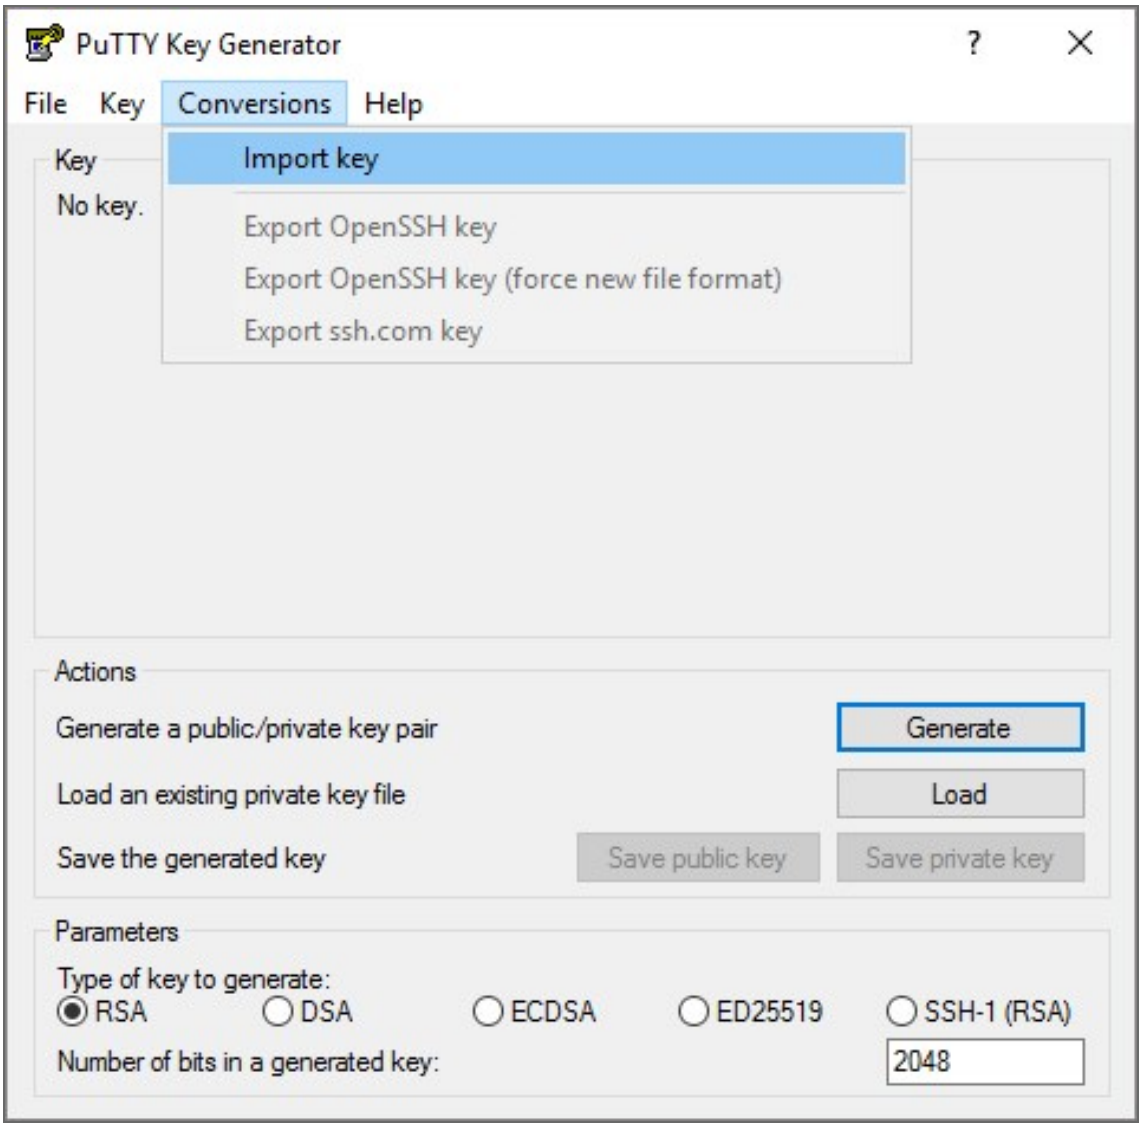
- WinSCP to upload the installer file.
- PuTTYgen tool to convert the private key.
- PuTTY tool to connect to an Amazon instance securely.

1. Log in to NAKIVO Backup & Replication.
2. Go to **Settings** > **Transporters**.
3. Download the keys of your Amazon instance.

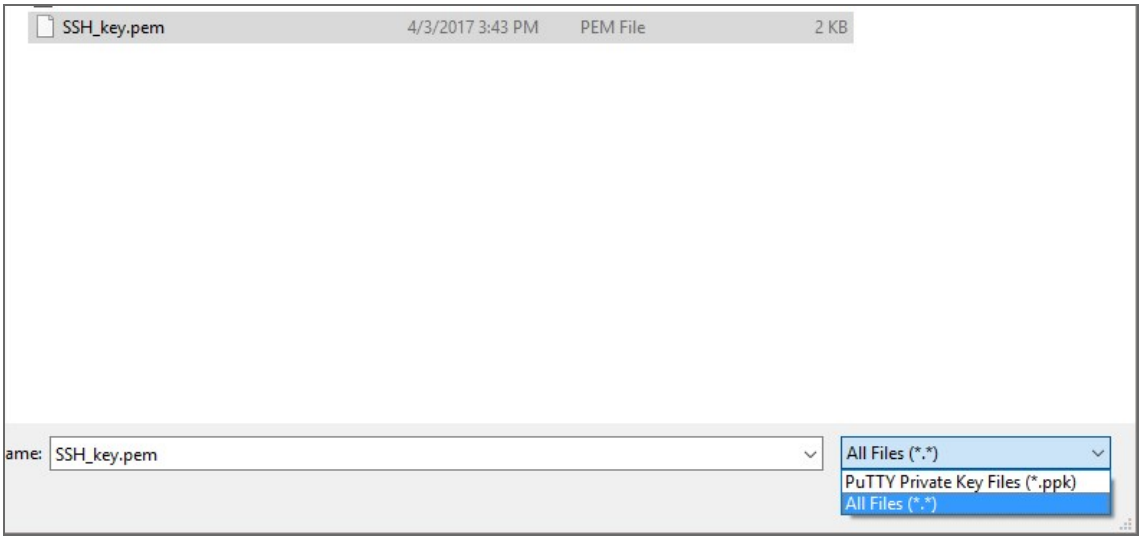


4. Click on the **Transporter** to view its details. Copy or remember the IP-address/hostname of the Amazon instance.
5. Unzip the folder with the key.
6. Convert the key using PuTTYgen:

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

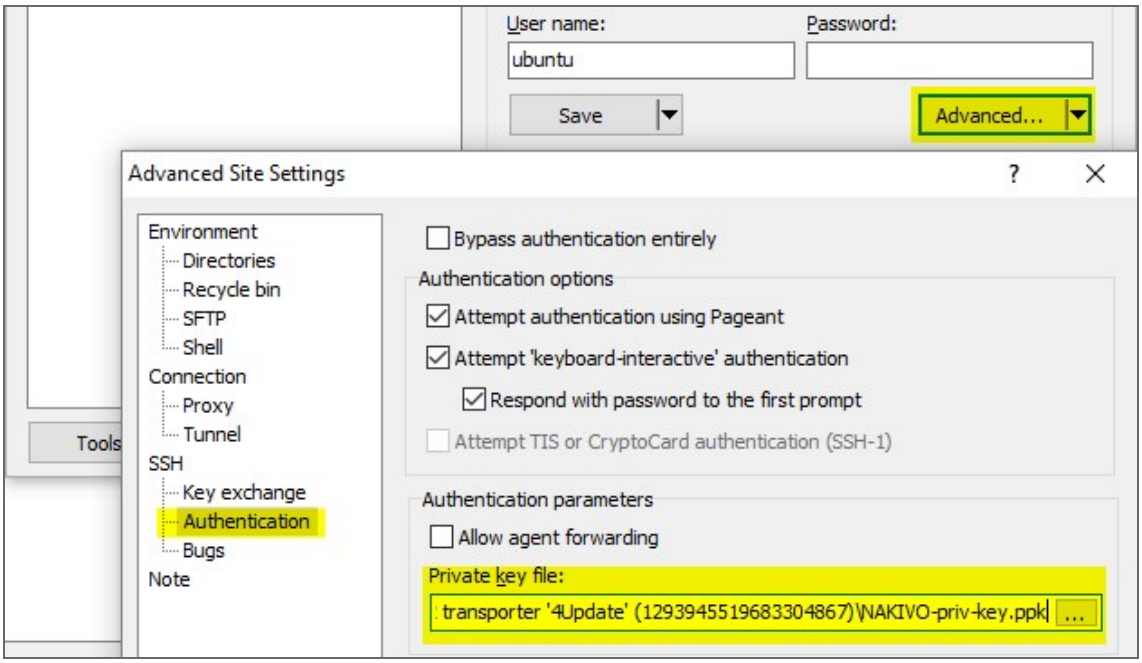


2. Locate the `SSH_key.pem` you just downloaded and unzipped. If you don't see it in the **Open**... dialogue box, change the file type to **All files**.



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

7. Open WinSCP.
8. Create a new session:
    a. Add the hostname or IP address of your Amazon instance you received on step 4 into the **Host Name** box.
    b. In the **Username** box, enter `nkvuser`.
    c. Leave the **Password** box empty.
    d. Add the private key to WinSCP:
        1. Click the **Advanced...** button.
        2. The **Advanced Site Settings** dialog box opens. Go to *SSH > Authentication > Private key file:* and select the key file you generated on step 6.



3. Click **OK**.

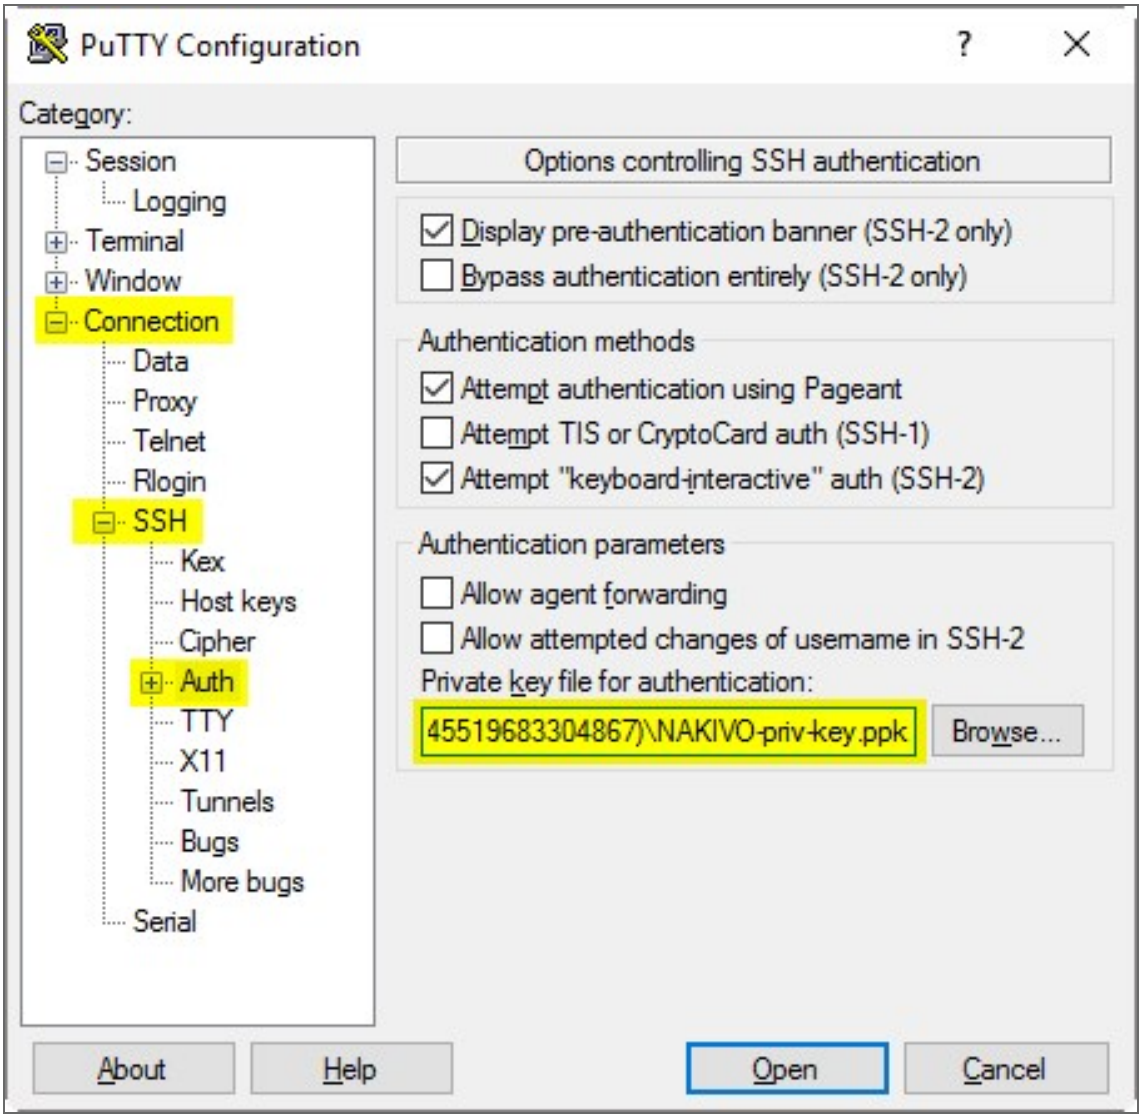
e. Click **Login**.

f. Upload the updater file.

g. Open PuTTY.

h. Enter the IP-address or hostname of the Amazon EC2 instance.

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



j. Click **Open**.

k. In the command line prompt that opens: log in to the Amazon EC2 instance:

   1. For **login**, enter `nkvuser`.
   2. For **password**, leave a blank line.

9. Update NAKIVO Backup & Replication following the instructions.

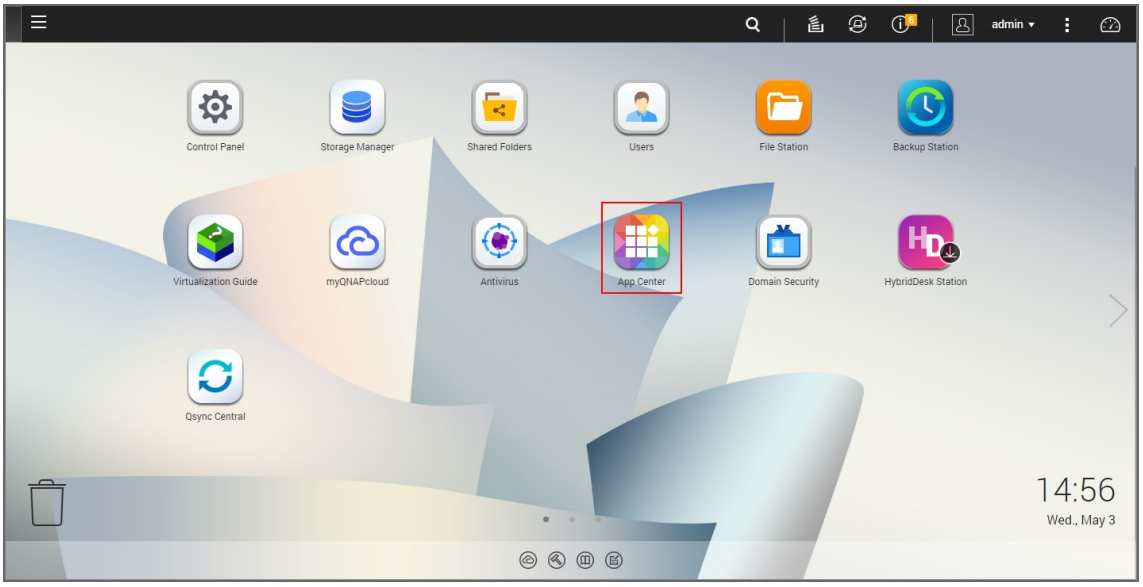
Updating on QNAP NAS

You can update NAKIVO Backup & Replication via QNAP AppCenter or manually. Refer to the following subtopics for details:

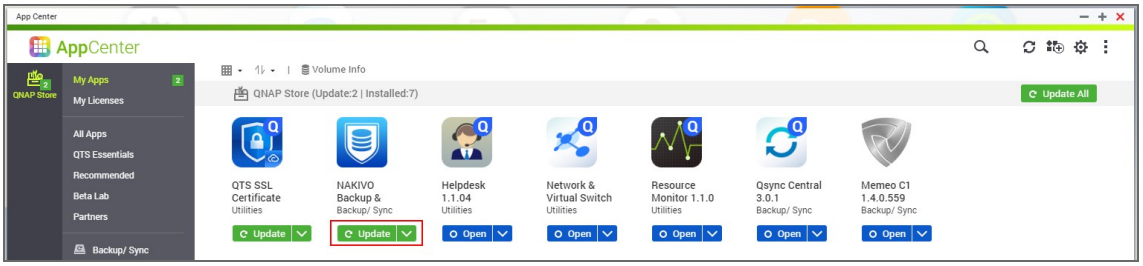
- Updating via QNAP AppCenter
- Updating Manually

Updating via QNAP AppCenter

1. Open the QNAP Desktop in your browser by entering the IP address of your QNAP NAS.
2. Make sure that no jobs or repository maintenance tasks are running in the product.
3. Go to **App Center**.



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



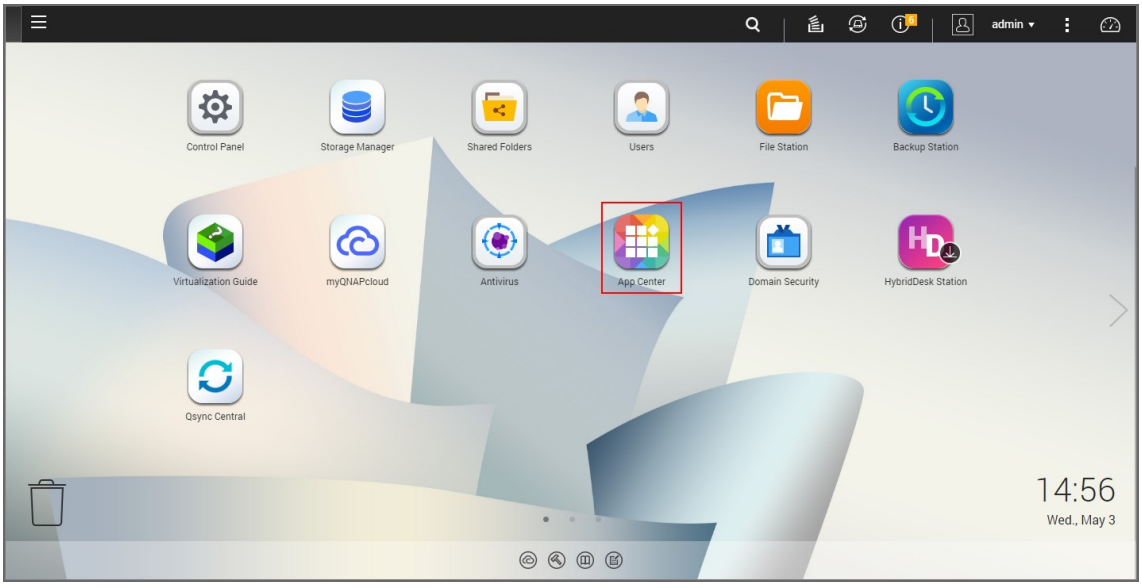
6. Click the **Update button** and wait till update finishes.

Updating Manually

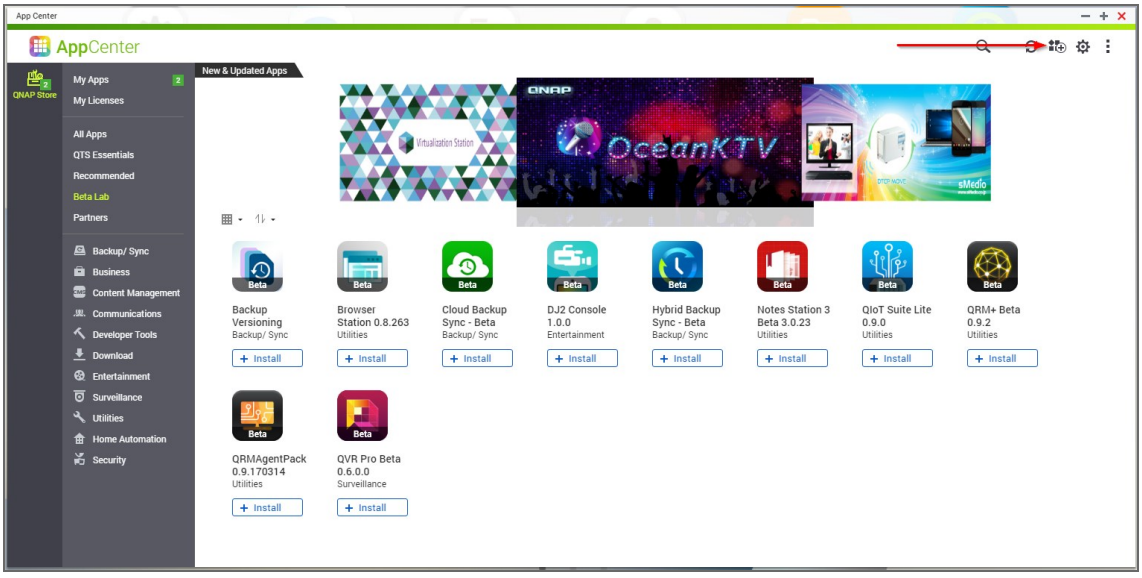
1. Download the update package from here.

> **Note**
>
> Installing updates of NAKIVO Backup & Replication on your NAS-based backup appliance requires accurate knowledge of the hardware system information.
>
> Refer to the following page to learn how to identify your QNAP device CPU model: How to Identify the CPU Architecture of a QNAP NAS Processor.

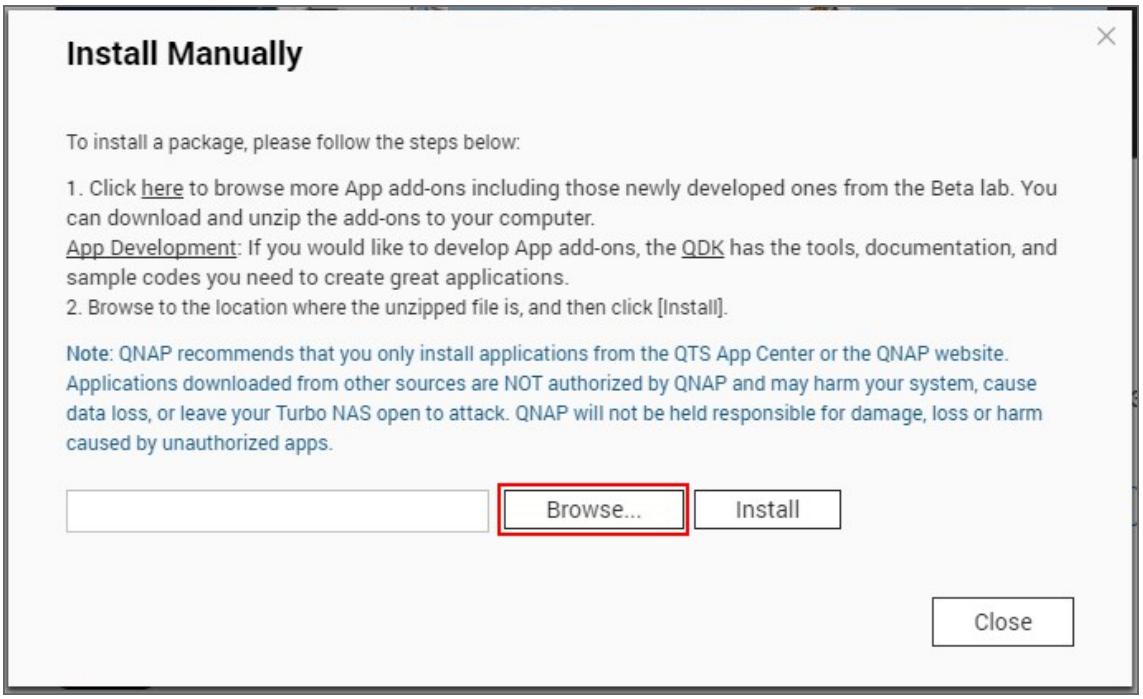
2. Open the QNAP Desktop in your browser by entering the IP address of your QNAP NAS.
3. Go to **App Center**.



4. Click the **Install Manually** icon.



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



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

Updating on ASUSTOR NAS

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

Updating on ASUSTOR NAS Manually

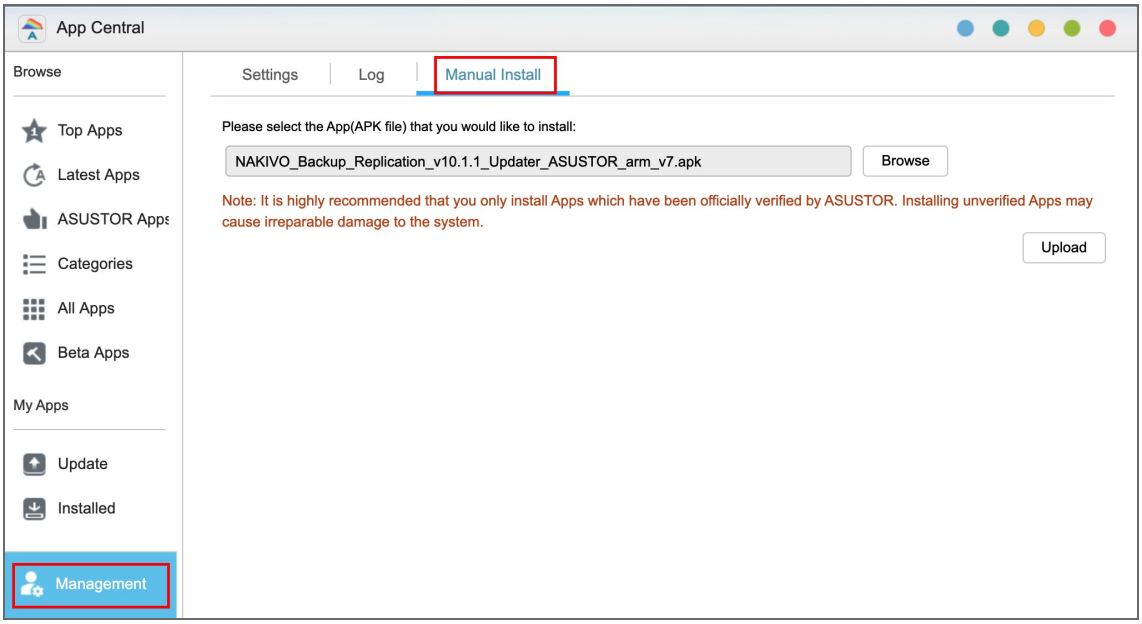
Prior to updating NAKIVO Backup & Replication on ASUSTOR NAS manually, make sure the following requirements are met:

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

> **Note**
>
> Installing updates of NAKIVO Backup & Replication on your NAS-based backup appliance requires accurate knowledge of the hardware system information.
>
> Refer to the following page to learn how to identify your ASUSTOR device CPU model: How to Identify the CPU Architecture of an ASUSTOR NAS Processor.

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

1. Make sure that no jobs or repository maintenance tasks are running in the product.
2. Open the **App Central** from the ASUSTOR NAS **Desktop**.
3. Click **Management** in the bottom left corner and click **Manual Install**.
4. The **Manual Install** pane opens to the right of the **App Central**. Click **Browse**.
5. The **Open** dialog box opens. Locate your copy of NAKIVO Backup & Replication installer for ASUSTOR NAS and click the **Open** button.
6. The **Open** dialog closes, and the **Upload** button becomes enabled. Click the **Upload** button.
7. When the upload finishes, the **About This App** dialog opens. If you are sure the requirements are met, click the **Next** button.
8. The **About This App** dialog opens a message asking you to review the summary of the NAKIVO Backup & Replication update. Select the checkbox **I understand the risks associated with installing unverified Apps** and click **Install**.
9. The **About This App** dialog closes, and the **Installed** pane of the **App Central** opens.
10. Wait until the update of NAKIVO Backup & Replication is complete.



Updating on ASUSTOR NAS via App Central

Follow the steps below to update NAKIVO Backup & Replication on ASUSTOR NAS via App Central:

1. Open the **App Central** from the ASUSTOR NAS **Desktop**.
2. In the **Browse** menu to the left, click **All Apps**. The list of applications available in **the App Central** opens in the right pane.
3. In the search box in the upper right corner of the pane, enter "Nakivo". Installations of the NAKIVO Backup & Replication application that are available at App Central are now displayed.
4. Click the **Update** button below the required NAKIVO Backup & Replication application to start uploading the update.
5. When the update is uploaded successfully, the **About This App** dialog opens. Click the **Update** button if you are sure that all the requirements are met.
6. The **About This App** dialog closes, and the **Installed** pane of the **App Central** opens. Wait until the update of the NAKIVO Backup & Replication is completed.

Updating on NETGEAR ReadyNAS

- Updating on NETGEAR ReadyNAS Manually
- Updating on NETGEAR ReadyNAS via Available Apps

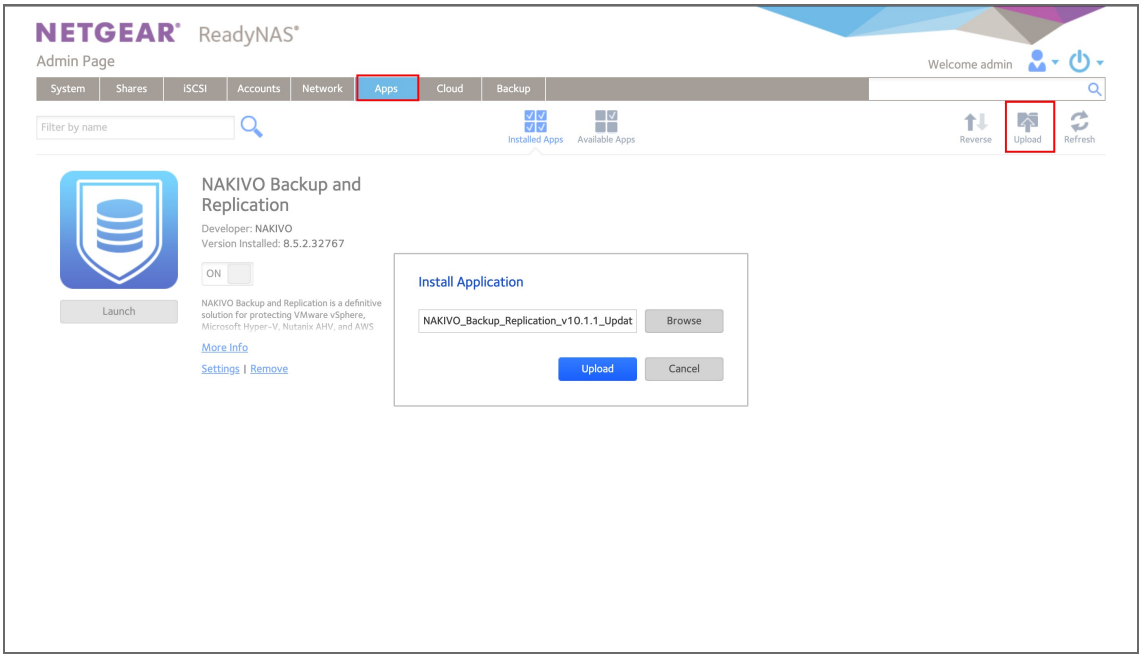
Updating on NETGEAR ReadyNAS Manually

Prior to updating NAKIVO Backup & Replication on NETGEAR ReadyNAS manually, make sure the following requirements have been met:

- You have access to the NETGEAR ReadyNAS.
- NAKIVO Backup & Replication update is available for your NETGEAR ReadyNAS.

Follow the steps below to update NAKIVO Backup & Replication on NETGEAR ReadyNAS manually:

1. Make sure that no jobs or repository maintenance tasks are running in the product.
2. Open the NETGEAR ReadyNAS **Admin Page** in your browser by entering the IP address of your NAS.
3. Go to **Apps** and click **Upload**.
4. The **Install Application** dialog box opens. Click **Browse**.
5. In the dialog box that opens, locate the downloaded installer (.deb file) and then click **Upload**.
6. Wait until the update is completed.



Updating on NETGEAR ReadyNAS via Available Apps

Important

This method is not possible for ReadyNAS OS 6.10.10. To use this method of updating NAKIVO Backup & Replication, you need to use or downgrade to ReadyNAS OS 6.10.9.

Prior to updating NAKIVO Backup & Replication on NETGEAR ReadyNAS via Available Apps, make sure that you have access to NETGEAR ReadyNAS.

Follow the steps below to update NAKIVO Backup & Replication on NETGEAR ReadyNAS via Available Apps:

1. Open the NETGEAR ReadyNAS **Admin Page** in your browser by entering the IP address of your NAS.
2. Go to **Apps** > **Available Apps**.
3. Find **NAKIVO Backup & Replication** in the list of available applications. Alternatively, enter NAKIVO to the filtering box in the upper left corner of the **Admin Page**.
4. If a new version of NAKIVO Backup & Replication is available in the NETGEAR **Available Apps**, the **Update** button will be available below the application item. Click the **Update** button.
5. Wait until the update is complete.

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 Windows

To uninstall NAKIVO Backup & Replication, run the uninstaller:

1. Go to **Start** -> **Control Panel** and run **Programs and Features**.
2. Select **NAKIVO Backup & Replication** and click **Uninstall**.
3. In the **NAKIVO Backup & Replication Uninstallation** wizard, click **Uninstall**.
4. Click **Close** when the uninstallation process is completed.

Uninstalling on Linux or Generic ARM-based NAS

Refer to the sections below to learn how to uninstall NAKIVO Backup & Replication on a Linux OS or a generic ARM-based NAS.

Uninstalling Director and Onboard Transporter on Linux or Generic ARM-based NAS

To uninstall the **Director** and Onboard **Transporter**, which is installed with the **Director** by default, follow the steps below:

1. Run the "uninstall" script which is located in the **Director** folder inside the product installation folder. If the product is installed in the default location, run: /opt/nakivo/director/uninstall
2. Enter "U" and then press **Enter** to confirm uninstalling the application.

Uninstalling Transporter on Linux or Generic ARM-based NAS

To uninstall the Transporter, follow the steps below:

1. Run the "uninstall" script which is located in the transporter folder inside the product installation folder. If the product is installed in the default location, run: /opt/nakivo/transporter/uninstall

2. Enter "U" and then press **Enter** to confirm uninstalling the application.

Uninstalling on Synology NAS

Follow the steps below to uninstall NAKIVO Backup & Replication on a Synology NAS:

1. In the Synology NAS management interface, open the **Package Center**.
2. Click NAKIVO Backup & Replication.
3. Choose **Uninstall** from the **Actions** list.
4. Click **OK** in the message box that opens to confirm that you wish to uninstall the application.

When the uninstallation process is completed, NAKIVO Backup & Replication will be removed from the list of installed applications.

Uninstalling on Western Digital NAS

Follow the steps below to uninstall NAKIVO Backup & Replication on a Western Digital NAS:

1. Open the NAS My Cloud Dashboard and click **Apps**.
2. In the **Installed Apps** list, select NAKIVO Backup & Replication.
3. The NAKIVO Backup & Replication item opens to the right of the installed applications list. Click the **Uninstall** button.
4. The **Uninstall NAKIVO Backup and Replication** dialog opens. Click **OK** to confirm that you wish to uninstall the application and delete all application data and settings.
5. The **Updating** progress bar opens. Wait until the uninstallation completes.

When the uninstallation process is completed, NAKIVO Backup & Replication will be removed from the list of installed applications.

Uninstalling on QNAP NAS

Follow the steps below to uninstall NAKIVO Backup & Replication on a QNAP NAS:

1. Open the QNAP NAS Desktop and click **App Center**.
2. The **App Center** dialog opens. In the **My Apps** list, locate the NAKIVO Backup & Replication application and open the list of applicable actions by clicking the drop-down button.
3. In the list of applicable actions, click **Remove**.
4. In the dialog that opens, click **OK** to confirm removing the application and application-relevant user data.
5. Wait until the uninstallation is complete.

When the uninstallation process is completed, NAKIVO Backup & Replication will be removed from the list of installed applications.

Uninstalling on ASUSTOR NAS

Follow the steps below to uninstall NAKIVO Backup & Replication on a QNAP NAS:

1. Open the ASUSTOR NAS Desktop and click **App Central**.
2. In the list of installed applications, locate NAKIVO Backup & Replication, select it and then click the **Remove** button.
3. In the dialog that opens, click **OK** to confirm that you wish to remove the application.
4. The **Removing** progress bar opens. Wait until the uninstallation is completed.

When the uninstallation process is completed, NAKIVO Backup & Replication will be removed from the list of installed applications.

Uninstalling on NETGEAR ReadyNAS

Follow the steps below to uninstall NAKIVO Backup & Replication on NETGEAR ReadyNAS:

1. Open the NETGEAR ReadyNAS **Admin Page** and go to **Apps** > **Installed Apps**.
2. Locate **NAKIVO Backup & Replication** in the list of available applications. Alternatively, enter NAKIVO to the filtering box in the upper left corner of the **Admin Page**.
3. Click the **Remove** button below the application item.
4. The **Confirm Deletion** dialog box opens. Click **Yes** to confirm that you wish to uninstall NAKIVO Backup & Replication on NETGEAR ReadyNAS.
5. Wait until the uninstallation is completed.

When the uninstallation process is completed, NAKIVO Backup & Replication will be removed from the list of installed applications.

Terminating on Amazon EC2

Follow the steps below to terminate NAKIVO Backup & Replication that is launched as an Amazon EC2 instance:

1. Open AWS Management Console and go to **EC2 Dashboard**.
2. In the **Instances** menu, click **Instances**.
3. In the list of instances, locate the necessary NAKIVO Backup & Replication instance and select it.
4. In the **Actions** menu, go to **Instance State** and click **Terminate**.
5. In the **Terminate Instances** dialog, click **Yes, Terminate** to confirm that you wish to terminate your instance of NAKIVO Backup & Replication.
6. Wait until the instance is terminated.

In about 60 minutes, the terminated NAKIVO Backup & Replication instance will be removed from the list of Amazon EC2 instances.

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 222
- “First Steps with NAKIVO Backup & Replication” on page 229
- “Web Interface Components” on page 232
- “Managing Jobs and Activities” on page 241

Logging in to NAKIVO Backup & Replication

- Getting to the Login Page
- Creating a User Account
- Resetting Password
- Default Password in Amazon EC2
- Passing Verification

Getting to the Login Page

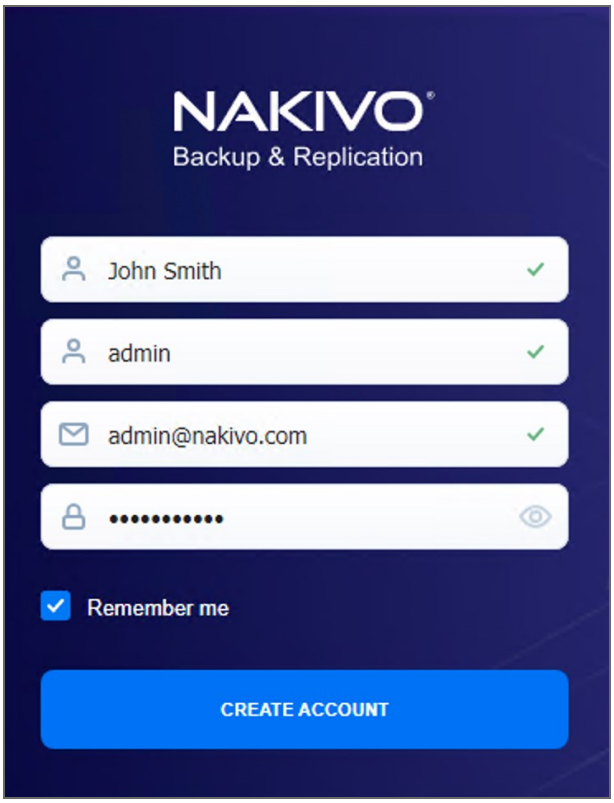
To open the NAKIVO Backup & Replication login page, enter the machine's IP address or DNS name in your web browser, followed by the HTTPS port (default port is 4443).

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

Creating a User Account

When you open the NAKIVO Backup & Replication login page for the first time, you are prompted to create a new user account. This user account is the admin account to be used to access your instance of NAKIVO Backup & Replication. Fill out the fields in the form:

1. **Name**: Provide your real name.
2. **Username**: Enter an admin username to log in to NAKIVO Backup & Replication.
3. **Email**: Provide an email.
4. **Password**: Enter a password.
5. Optionally, you can select **Remember me** to save your credentials.
6. Click **CREATE ACCOUNT**.



Note
If NAKIVO Backup & Replication is deployed in an Amazon EC2 instance, you will first be prompted to enter the Amazon EC2 instance ID.

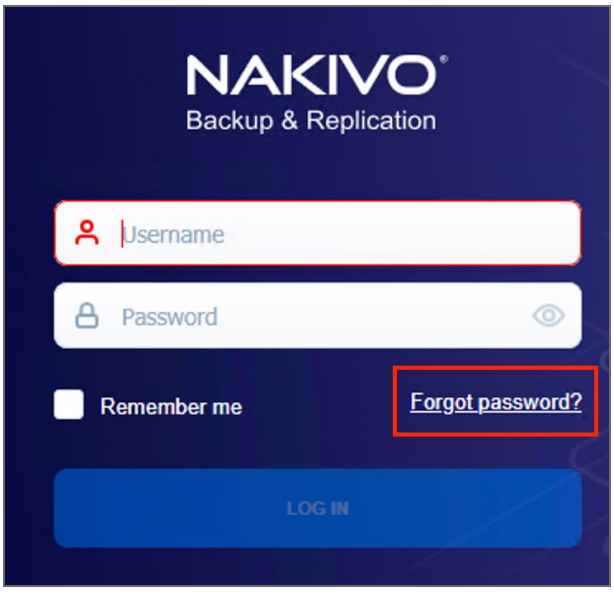
NAKIVO Backup & Replication opens in your browser displaying the configuration wizard. Refer to First Steps with NAKIVO Backup & Replication to learn how to start using NAKIVO Backup & Replication.

To log out, click **Logout** in the bottom left corner.

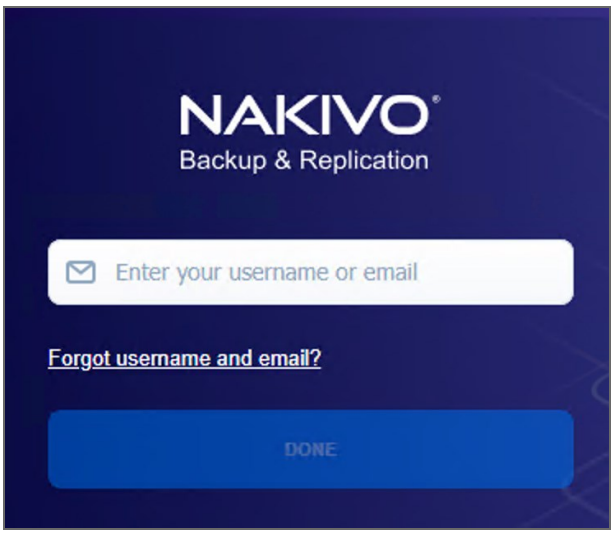
Resetting Password

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

1. Go to NAKIVO Backup & Replication login page.
2. Click the **Forgot password** link.



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

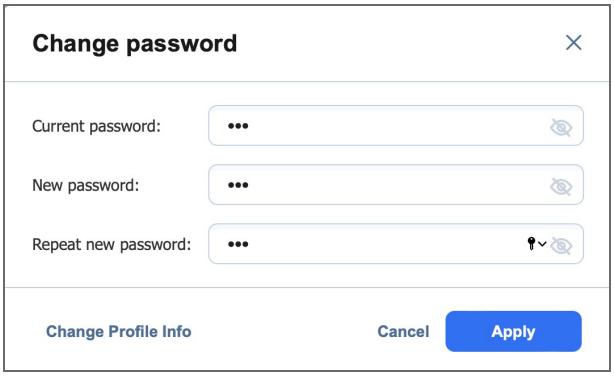


A temporary password, which is a security string, is sent to your inbox. Enter this password the next time you log in to your NAKIVO Backup & Replication instance. Once you are logged in, it's recommended that you change the temporary password for your user account. To change the temporary password:

a. Click **Logout** in the bottom left corner.
b. Select **Profile**.
c. Click **Change password**.
d. In the dialog box that opens, fill out the following fields:
   - **Current password**: Enter the temporary password that you received to your inbox.
   - **New password**: Enter a new password.

- **Repeat new password**: Enter the new password again.

e. Click **Change**.



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

- If you have not set up email settings in NAKIVO Backup & Replication:
  a. Enter your username and click **Done**.
  b. Go to the product installation folder and locate the “forgot_password.txt” file.

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

  c. Paste the security string from the file in the appropriate field.
  d. Click **Done**.

Notes

- If you are using a Virtual Appliance (VA), go to the VA console, then go to the command line and enter: `cat /opt/nakivo/director/forgot_password.txt` The security string will be displayed on the screen. You can copy and paste it into the web interface.
- For Windows OS, you can find forgot_password.txt file in the following default installation folder: *C:\Program Files\NAKIVO Backup & Replication*. Please note that if you changed the installation path, the file will be located in the folder you specified.
- If you are using a NAS, open an SSH connection to your device and read the forgot_password.txt file in the following folders:
  - For ASUSTOR NAS: `/usr/local/AppCentral/NBR`
  - For NETGEAR NAS: `/apps/nbr`
  - For QNAP NAS: `/share/CACHEDEV1_DATA/.qpkg/NBR`
  - For Synology NAS: `/volume1/@appstore/NBR`
  - For Western Digital NAS: `/mnt/HD/HD_a2/Nas_Prog/NBR`
- To learn how to open an `SSH` connection to your NAS device and read text files, refer to the NAS vendor documentation.

Default Password in Amazon EC2

If you have deployed NAKIVO Backup & Replication as an Amazon machine image in Amazon EC2, use the following default credentials to log in:

- **Username**: admin
- **Password**: The password is the ID of the NAKIVO Backup & Replication instance in Amazon EC2.

Passing Verification

If two-factor authentication was configured, verification needs to be passed after entering the credentials to access your NAKIVO Backup & Replication instance. This can be done in one of the following ways:

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

If Two-factor authentication was enabled but never configured, it must be configured now. Do the following:

1. Click **Continue**.
2. Optionally, click on the **change your email** link to enter the new email address for the user. Select **Continue** to proceed.

3. Enter the verification code that was sent to the specified email and click **Continue**. Optionally, click **Resend email** in case you did not receive it.
4. Optionally, enter the alternative email address that can be used in case the primary one becomes unavailable, and select **Continue**. Alternatively, select **Skip this step**.
5. If you have entered the alternative email address for the previous step, enter the verification code that was sent to the specified email, and click **Continue**. Optionally, click **Resend email** in case you did not receive it.
6. Follow instructions on screen to download and install Google Authenticator, and click **Continue**.
7. Add your NAKIVO Backup & Replication user account to Google Authenticator. Use one of the following methods:
    - Select **Scan QR Code** option and scan the QR code in the popup window.
    - Select **Enter a Code** option and follow the instructions to enter the shown code into the Google Authenticator app.
8. Enter the 6-digit verification code from Google Authenticator into the field. Note that the verification code is time-based. Click **Continue** to proceed.
9. A pairing key is displayed which can be used to add multiple devices to your account.

    Important
    It is highly recommended that you save the pairing key or write it down.

    You have the following options:
    - Optionally, click on the **Copy the key** link to copy your key and save it for future use.
    - Optionally, click on the **Download pairing information** link to download and save instructions on how to use the pairing key.
    - Click **Continue** when you're done.
10. Four backup codes are displayed on the next page. These one-time codes can be used to log in when you are unable to provide a verification code. Click on the **Save as PDF** link to download and save these codes in PDF format or write them down. Click **Continue**.
11. Enter one of the backup codes in the next popup window to confirm that you have saved them, and click **Finish**.

Google Authenticator Verification

If you have selected the **Google Authenticator** verification method on the **Managing Two-Factor Authentication** page, do one of the following:

- Enter the verification code from Google Authenticator into the field, and click **Proceed**.
- Enter one of the one-time backup codes.
- Click **More verification options** to use email verification.

Email Verification

If you have selected the **Email** verification method on the **Managing Two-Factor Authentication** page, do one of the following:

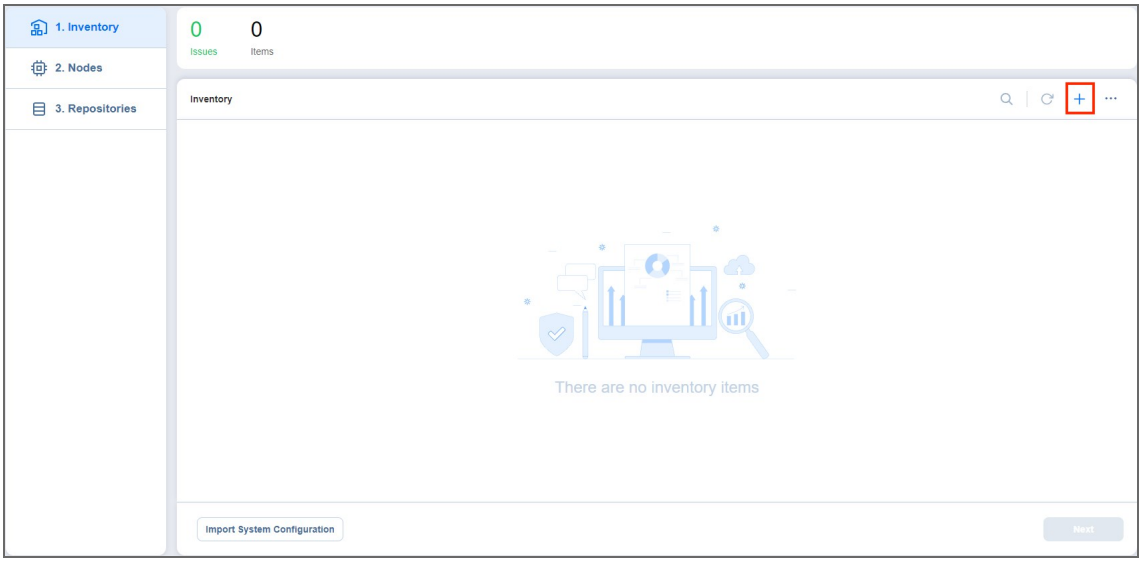
- Select one of the email addresses verified previously, and click **SEND VERIFICATION CODE**. Then click **OK**.
- Enter one of the one-time backup codes.
- Alternatively, click **More verification options** to choose a different email for verification.

First Steps with NAKIVO Backup & Replication

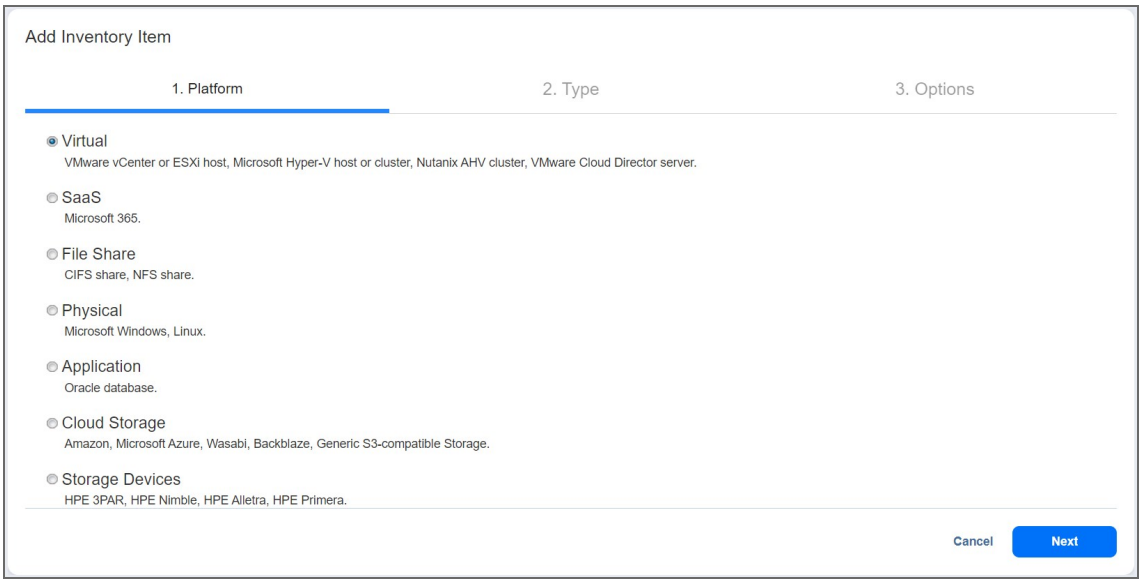
When you log in to NAKIVO Backup & Replication for the first time, the initial configuration wizard opens.

Proceed as follows:

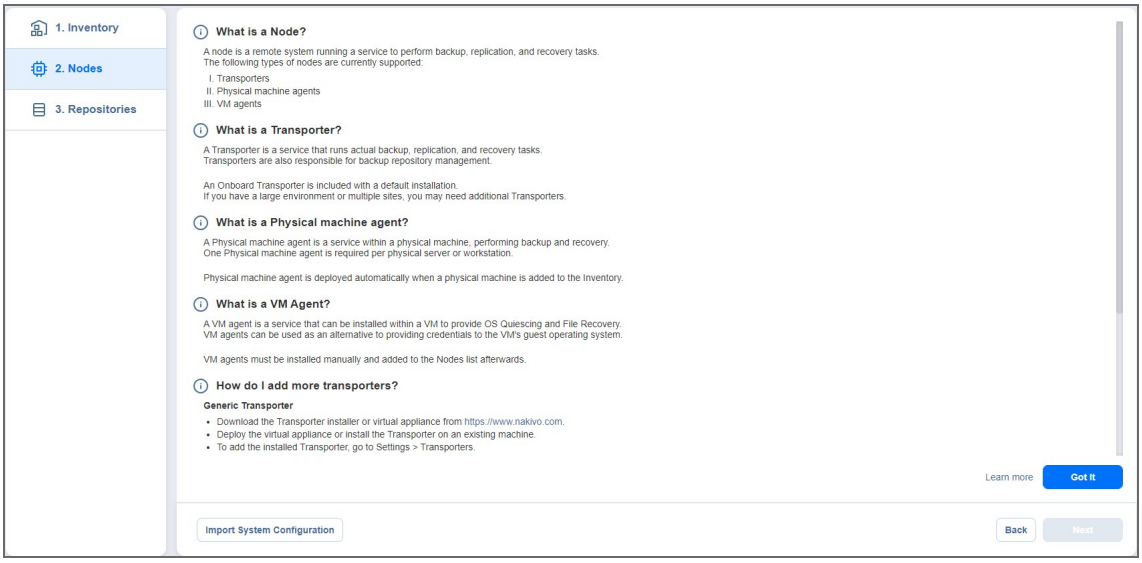
1. On the **Inventory** page of the wizard, click **Add New**.



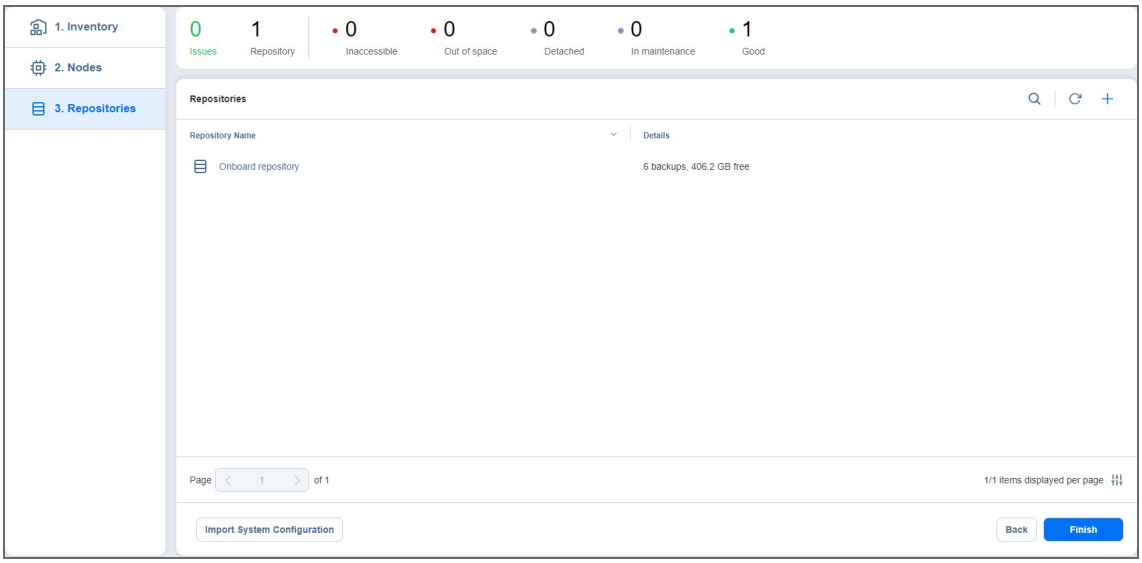
2. Select one of the given options:
    - Virtual
    - SaaS
    - File Share
    - Physical
    - Application
    - Cloud Storage
    - Storage Devices



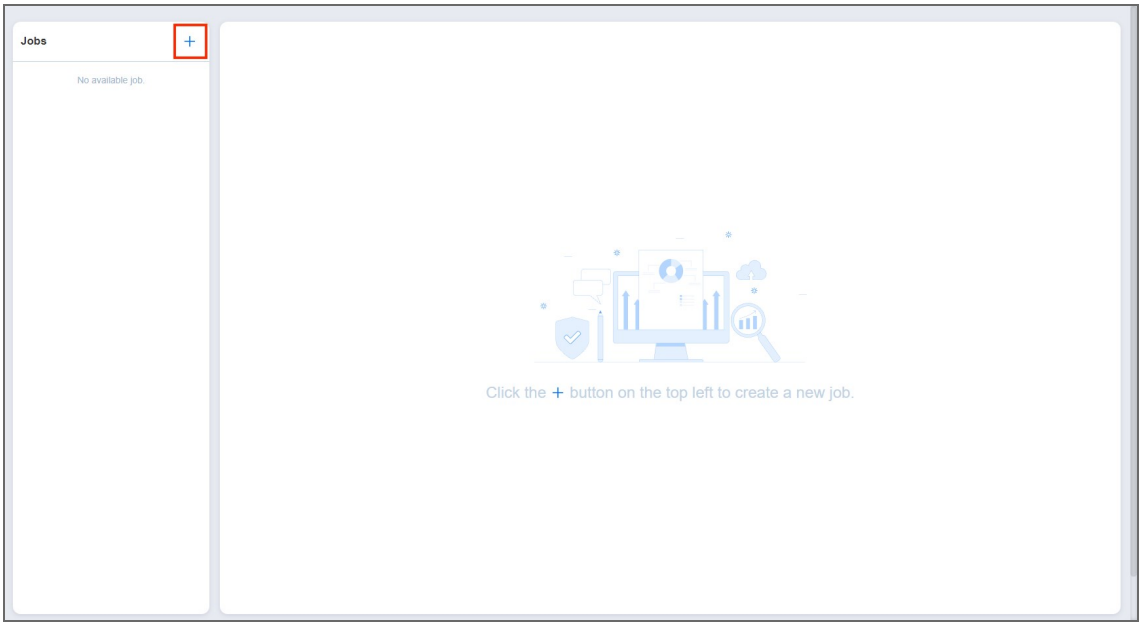
3. Proceed with adding items as described in the Inventory article.
4. On the **Nodes** page of the wizard, you will find information about the **Transporter** component of the NAKIVO Backup & Replication.



5. To deploy a new **Transporter** or add an existing one, click **Got it** and proceed as described in the **Transporters** article.
6. To move to the next page of the wizard, click **Next**.
7. On the **Repositories** page of the wizard, you can add a local or a remote **Backup Repository** to your application by clicking **Add Backup Repository.**



8. Click **Finish**.
9. The **Jobs** menu of the application opens. Proceed by creating backup jobs.



If your current license type is **Free** and the **Trial** license has not yet been applied to the current deployment of NAKIVO Backup & Replication, a dialog box appears. Using this dialog box, you can contact the sales team to change your license type or try the full functionality of the solution for 15 days. If you do not want to upgrade your license type right away, you can do it at any time in the Help menu.

Note

If you switch the license type to **Trial**, the product will automatically go back to using your **Free** license after expiration.

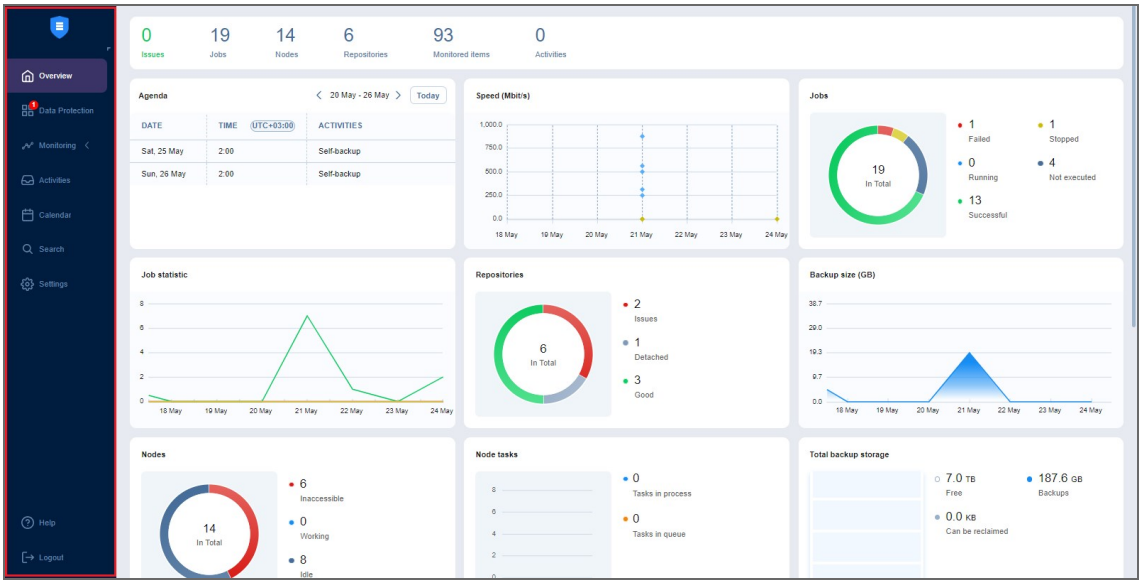
Web Interface Components

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

- Main Menu
- Overview
- Data Protection
- Monitoring
- Activities
- Calendar
- Search
- Settings
- Help Menu
- Online Chat Dialog
- Language Selection
- Special Offers Toolbar
- Tenants Dashboard

Main Menu

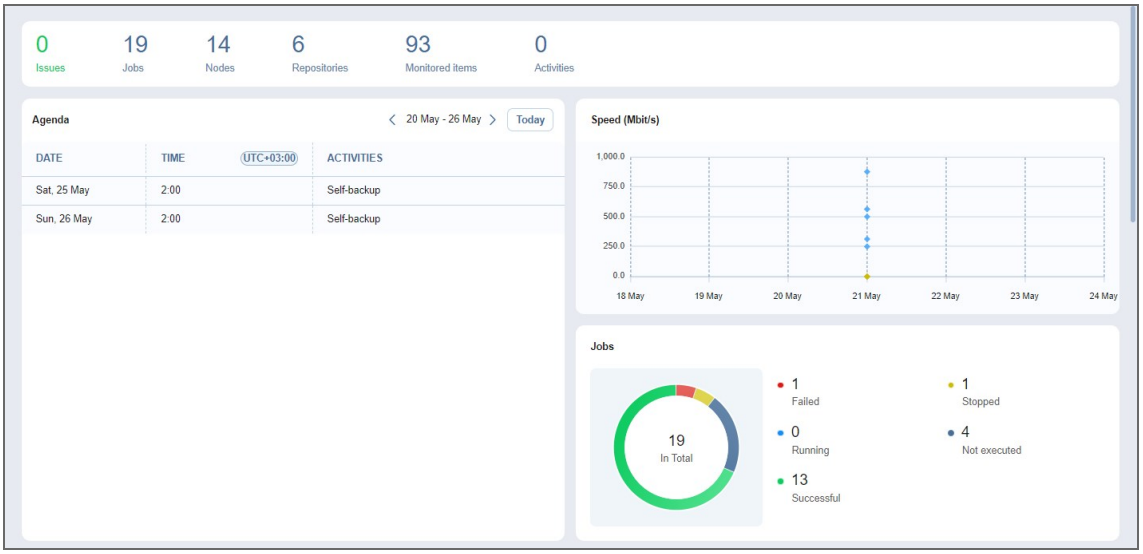
The main menu of NAKIVO Backup & Replication is located on the left side of the product interface. It provides access to the overview dashboard, jobs, activities, calendar, global search, and product settings. It also contains the **Help** menu and **Log Out** button.



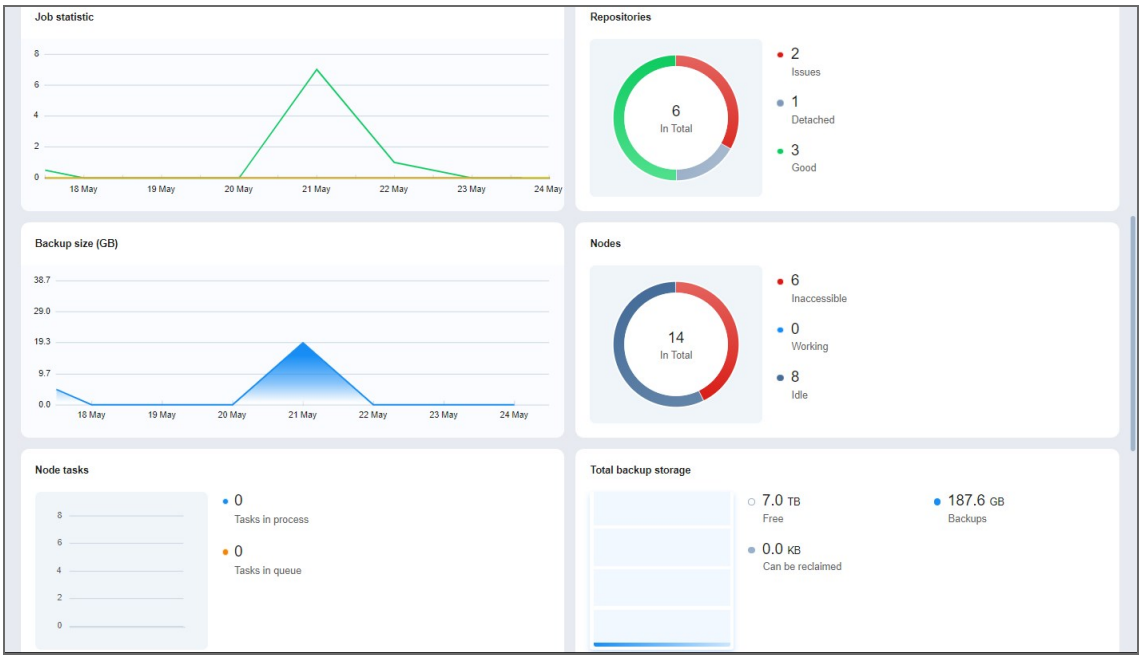
Overview

The **Overview** page displays the key statistics for your instance of NAKIVO Backup & Replication. The information is displayed in the following widgets:

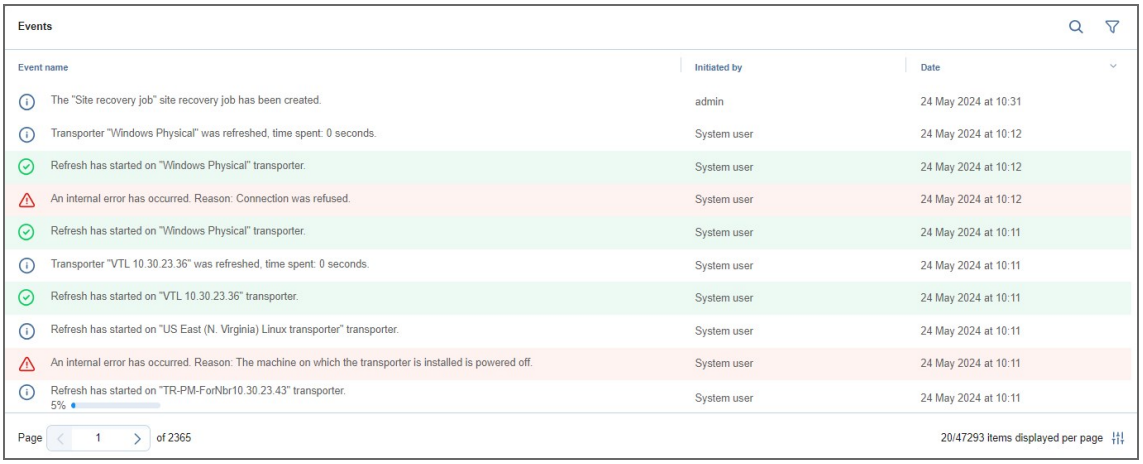
- **Summary bar**: Lists the total number of issues (errors and alarms), jobs, transporters, repositories, monitored items, and running activities.
- **Agenda**: Lists running and scheduled activities for a given week. By default, this widget displays the current week.
- **Speed**: Displays the speed at which raw data has been transferred during successful job runs in the previous seven days.
- **Jobs**: Displays the total number of available jobs and their respective last run statuses.



- **Job statistic**: Shows a graph of the number of successful, stopped, and failed jobs for each day in the previous seven days.
- **Repositories**: Displays the total number of available repositories and their statuses.
- **Backup size**: Displays the total size of backups created for each day in the previous seven days. Note that backups in forever incremental (**Store backups in separate files option** is not selected) Backup Repositories are considered 0KB, and thus are not reflected in the **Backup size** graph.
- **Transporters**: Displays the total number of available transporters that have been added or deployed successfully and their statuses.
- **Transporter tasks**: Displays the total number of tasks being processed or waiting to be processed by all transporters.
- **Total backup storage**: Displays the total amount of storage of all available repositories.



- **Events**: Lists all events, including errors, warnings, and general status information, sorted by date by default. Includes search and filter functions that simplify finding specific events by name, type, or date range.

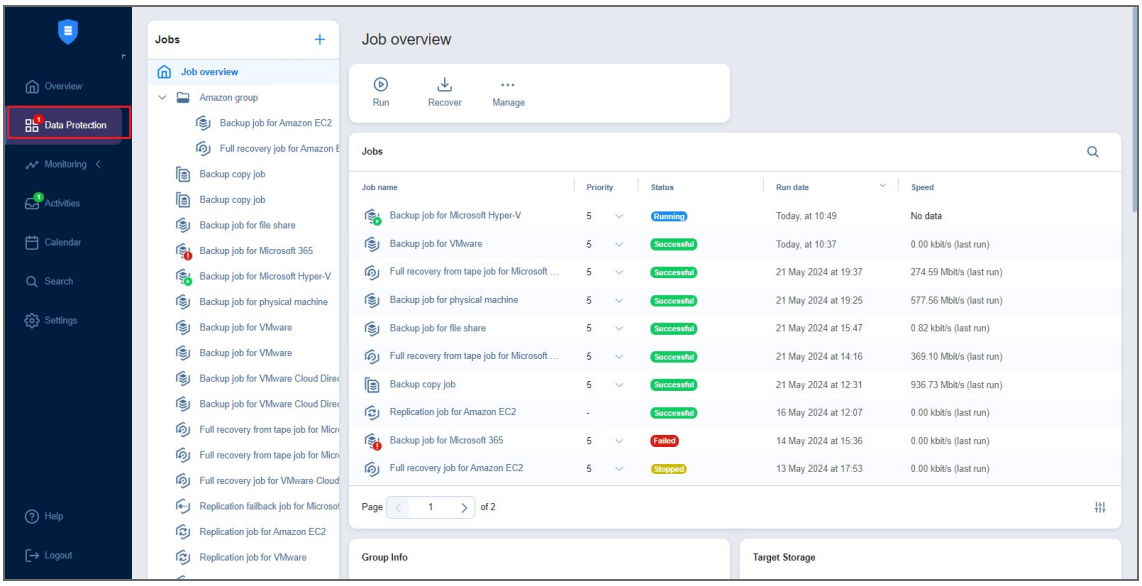


Data Protection

Using the **Data Protection** page, you can:

- View, run, and stop jobs on demand or on schedule
- Recover files, objects, VMs, and sites
- Manage jobs
- Create backup, backup copy, replication, recovery, and flash boot jobs
- Create job reports

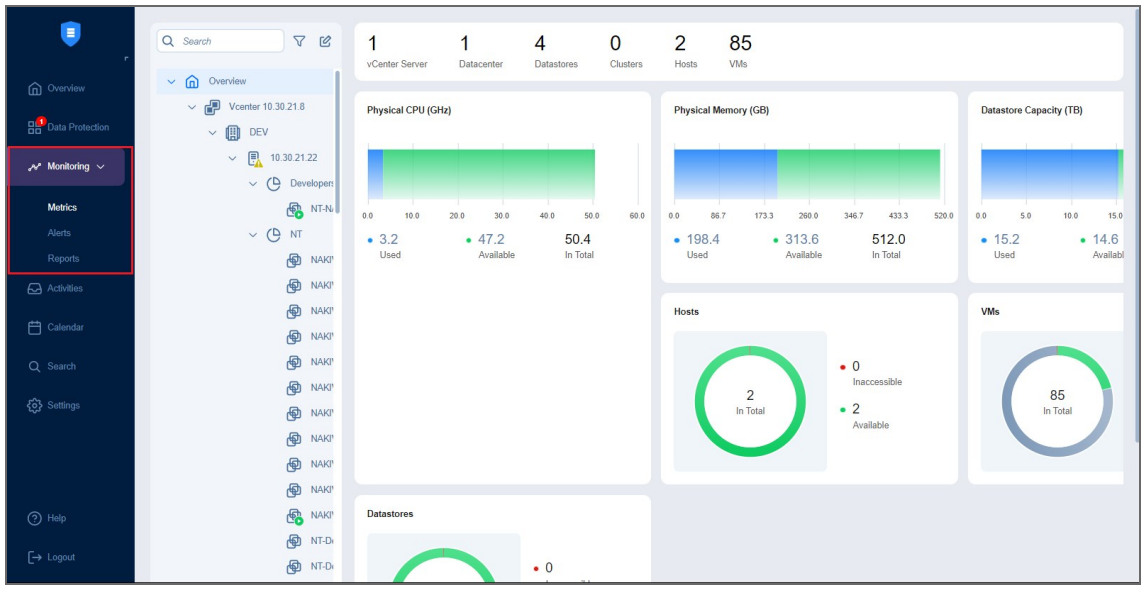
- Create and manage job groups



Monitoring

On the **Monitoring** page, you can check the following metrics (current and historical):

- For VMware VMs: CPU usage, memory usage, and disk usage
- For VMware hosts: CPU usage and memory usage
- For VMware datastores: Disk usage

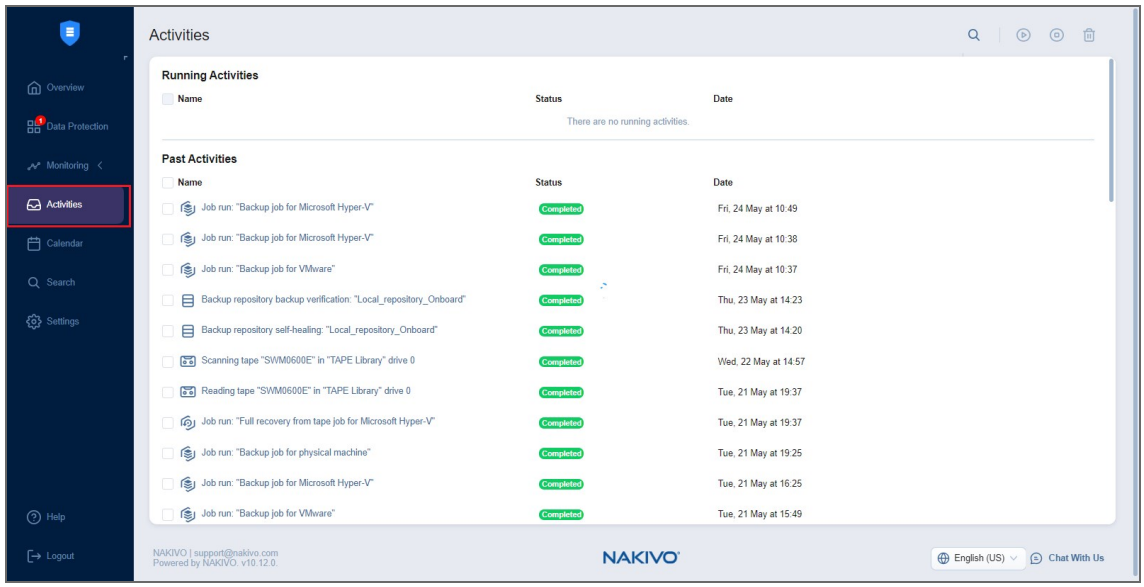


Activities

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

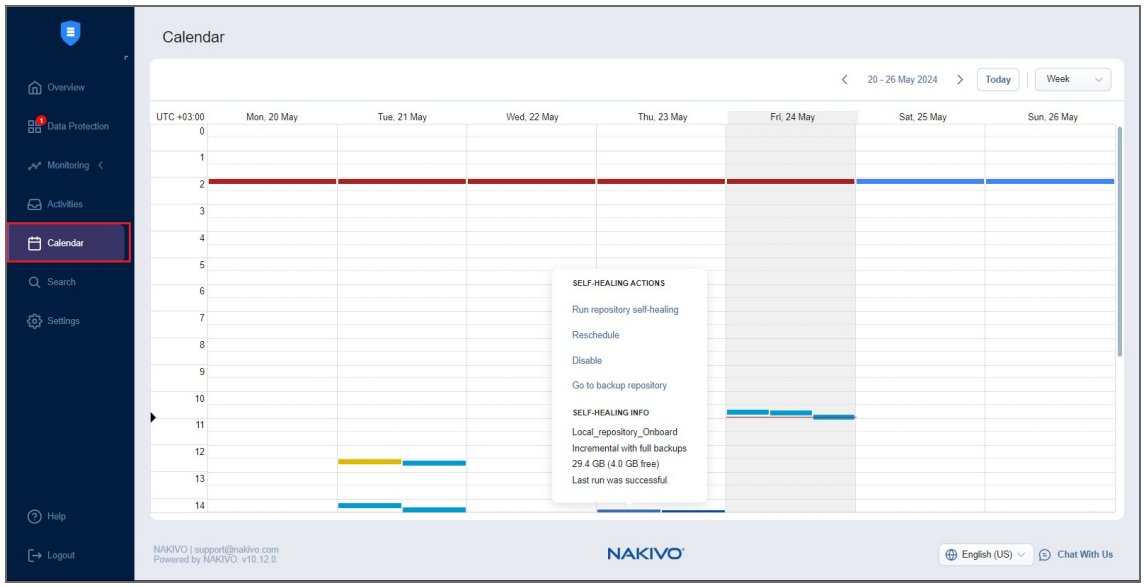
- Job run
- Repository Self-Backup
- Other

For further details and information, refer to “Managing Activities” on page 270.



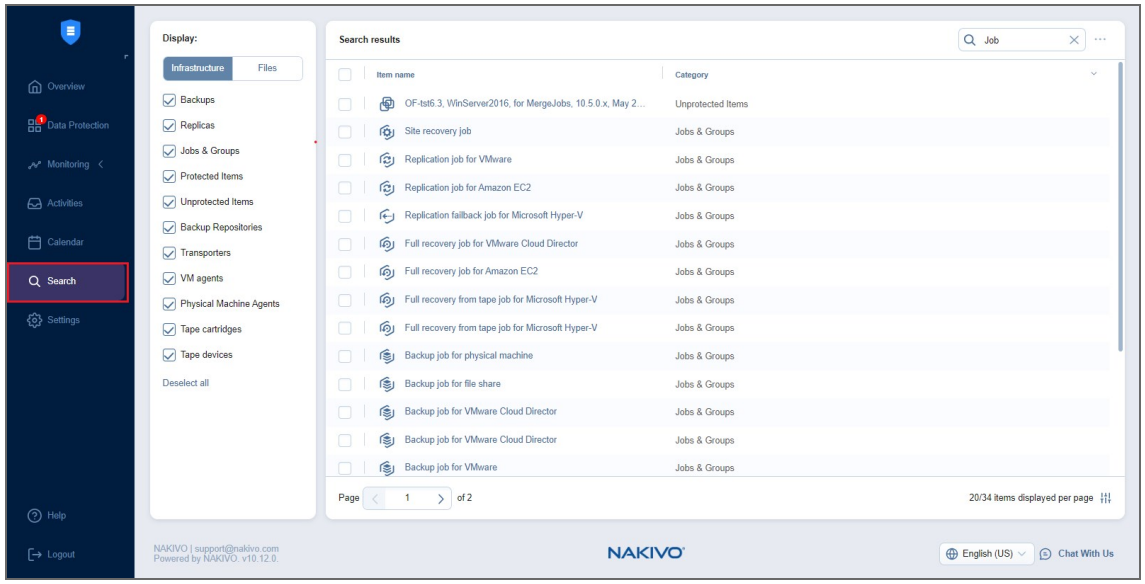
Calendar

The **Calendar** allows you to schedule jobs and view the history of all job runs in one organized space. For more information, refer to “Using Calendar” on page 274.



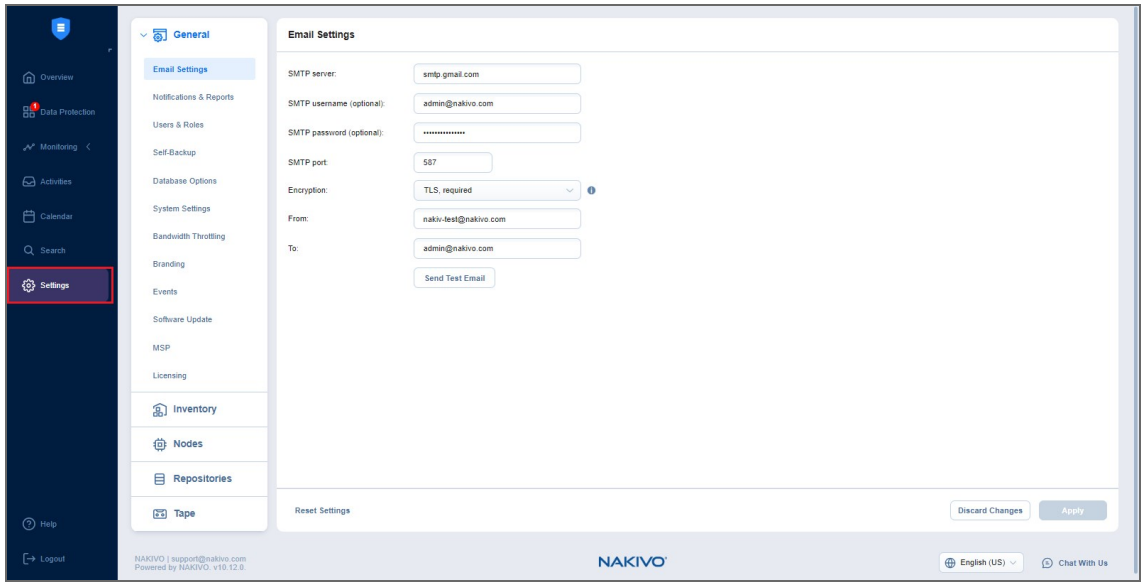
Search

The **Search** page allows you to search for items within the entire NAKIVO Backup & Replication instance–the Inventory, Transporters, Repositories, tape devices, jobs, backups, replicas, and more. For more details, refer to “Using Global Search” on page 275.



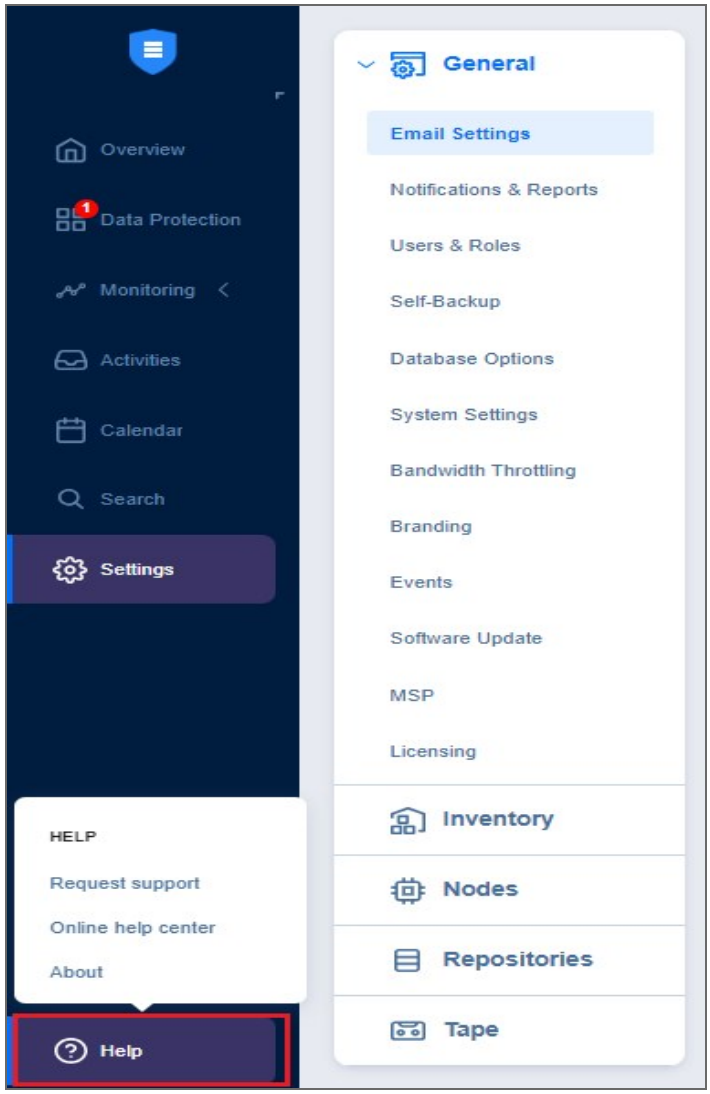
Settings

On the **Settings** page, you can configure NAKIVO Backup & Replication General, Inventory, Transporters, Repositories, and settings. Refer to “Settings” on page 283 for more detailed information.



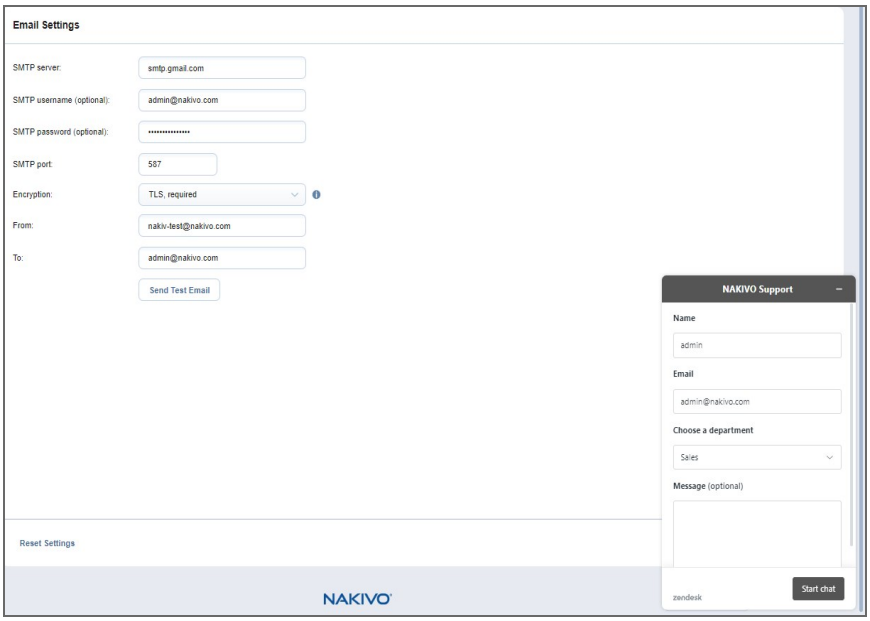
Help Menu

Use the **Help** menu to request technical support and access the NAKIVO online help center or see information about your product. If you are evaluating NAKIVO Backup & Replication, you may also check the the **How to Buy** section of the **Help** menu to view pricing, request a live demo or quote, find a reseller, or contact Sales. If you are using a Free license, you may also upgrade to a Trial license for 15 days with the **Try full functionality** option.



Online Chat Dialog

The **NAKIVO Support** online chat is located in the right bottom corner of the application. It enables you to quickly request help from a sales or technical support representative.



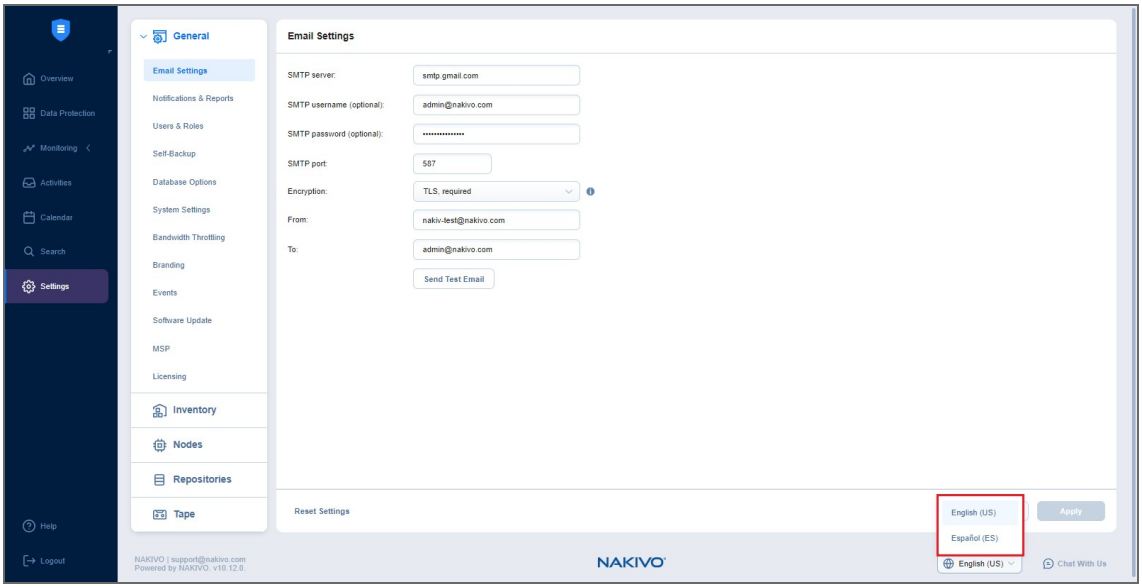
Language Selection

NAKIVO Backup & Replication supports the following languages:

- Spanish
- French
- German
- Italian
- Polish
- Chinese (simplified)

To select the language for your NAKIVO Backup & Replication instance UI, do the following:

1. Log in to your NAKIVO Backup & Replication instance.
2. In the lower right corner of the NAKIVO Backup & Replication interface, select the needed language from the dropdown menu.



3. The current page reloads automatically with the UI and content switched to the selected language.

Notes

- The language switching menu is available to users with the **Administrator** role only; the language change is applied globally (across all users).
- In Multi-tenant deployment, it is available at the Master Admin level only; the language change is applied globally (across all local tenants/users).
- Only users with the **Administrator** role can change the language in a remote tenant.
- You can also change language in the product installer UI, in silent installation interface, and using the command line interface

Special Offers Toolbar

This element of the interface is located to the left of the NAKIVO Backup & Replication dashboard. The toolbar contains special offers. If you click the button, a dialog opens displaying information about a specific offer. If needed, the **Special Offers** toolbar can be disabled. Refer to "System Settings" on page 332 for details.

Tenants Dashboard

If you use NAKIVO Backup & Replication in a multi-tenant mode, the **Tenants** dashboard allows you to create, manage, and configure tenants.

Managing Jobs and Activities

Using NAKIVO Backup & Replication interface, you can manage jobs and tasks. This section covers the following topics:

- "Jobs Dashboard" below
- "Running Jobs on Demand" on page 250
- "Managing Jobs" on page 257
- "Job Alarms and Notifications" on page 268
- "Managing Activities" on page 270
- "Using Calendar" on page 274
- "Using Global Search" on page 275

Jobs Dashboard

The **Data Protection** (**Jobs**) dashboard is a detailed interface where you can create and manage jobs, as well as get an overview of job details. For a detailed explanation of each component in the Jobs dashboard, see the sections below.

- Group/Job Overview Dashboard
  - Action Bar
  - Summary Bar
  - Jobs Table
  - Group Info
  - Overview Panes
- Job Dashboard
  - Action Bar
  - Summary Bar
  - Job Info
  - Job Settings
  - Job Objects
  - Overview Panes

Group/Job Overview Dashboard

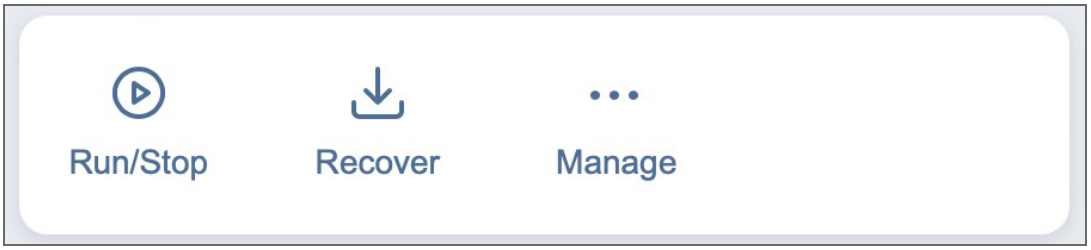
The **Job overview** and job group views offer an overview of multiple jobs. See the sections below for more information.

Action Bar

The group and **Job overview** action bars contain the following three job actions:

- **Run/Stop**: Opens the Run/Stop Jobs dialog box
- **Recover**: Brings up a list of recovery options for the selected group of jobs
- **Manage**: Brings up the options to **Rename**, **Delete**, or **Disable** a job group, as well as change the destination for all backup jobs in the group

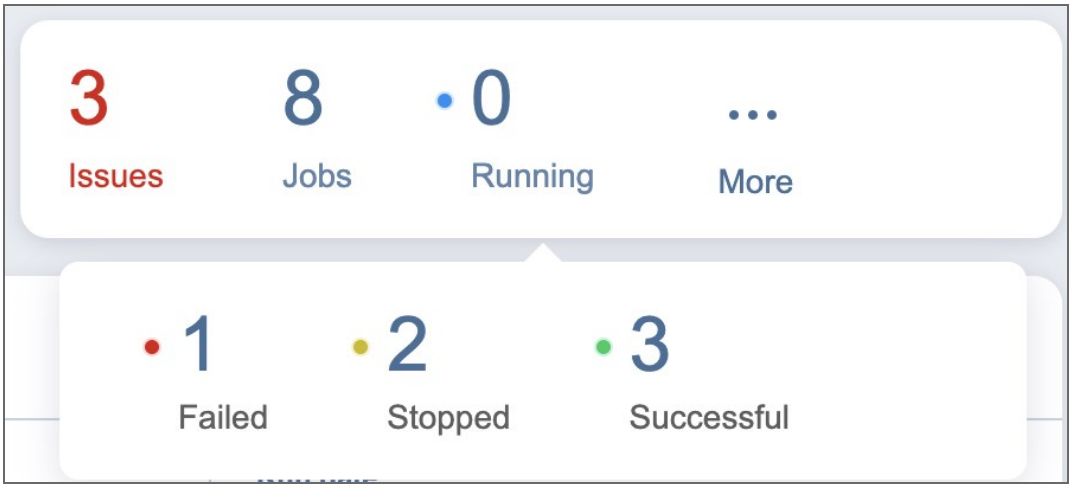
For more information on using the action bar, see "Running Jobs on Demand" on page 250.



Summary Bar

The summary bar displays information about the jobs in a given group. The data displayed is as follows:

- **Issues**: Total number of issues/alarms for the group of jobs. When clicked, this displays the Alarms & Notifications dialog box.
- **Jobs**: Total number of jobs in the group.
- **Running**: The number of running jobs in the group.
- **Failed**: The number of failed jobs in the group.
- **Stopped**: The number of stopped jobs in the group.
- **Successful**: The number of successful jobs in the group.



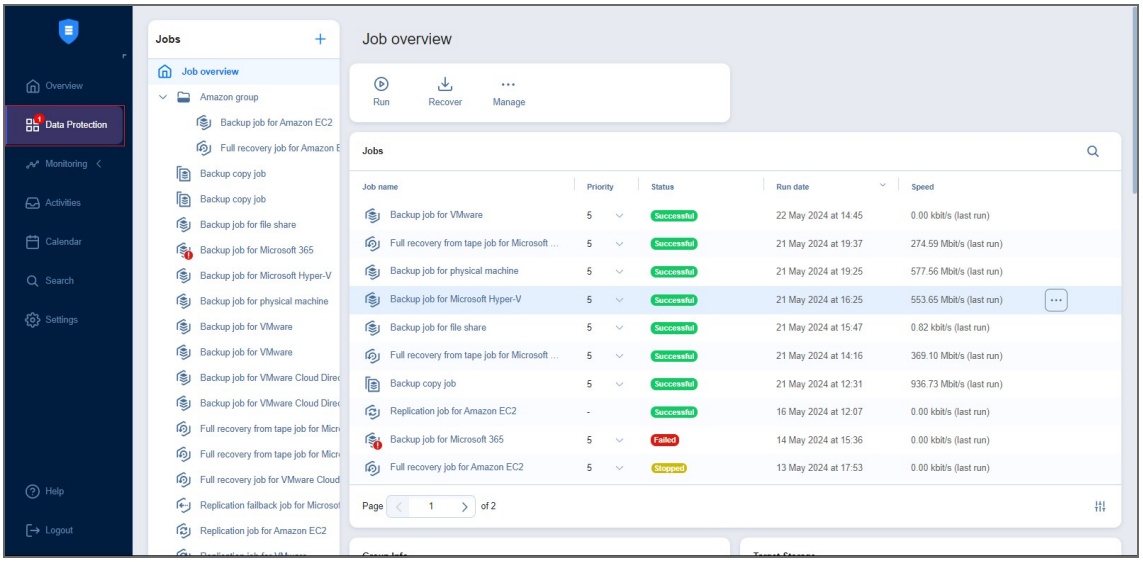
Jobs Table

The **Jobs** table shows a list of jobs and the information about each job in the following columns:

- **Job name**: The name of a given job in the group.
- **Priority**: The priority level of a given job in the group. Click the arrow button to the right of this parameter to change the priority level of a job.

> **Note**
> This column is only available in the Enterprise, Enterprise Essentials, Enterprise Plus, MSP Enterprise, and MSP Enterprise Plus editions.

- **Status**: The status of a given job:
  - **Successful**: The last job run was successfully completed.
  - **Failed**: The last job run failed.
  - **Running**: The job is currently running.
  - **Stopped**: The last job run was stopped.
  - **Not executed yet**: The job has not been executed yet.
- **Run date**: The date of a given job's last run.
- **Speed**: If the job is currently running, displays the current job run speed. If the job is not currently running, displays the speed of the last job run.



To customize the sorting of the **Jobs** table, click the head of the column you wish to sort by. To change the order of the columns, drag and drop a column to the needed position. You may also search for a job by clicking the **Search** button at the top of the table. To manage a job in the table, hover over a job and click the ellipsis **Manage** button on the right side.

Group Info

This pane displays current information about the jobs in the selected group. The information includes:

- Currently running jobs, displayed as a ratio to the total number of jobs
- The status of recent jobs; **Completed**, **Failed**, or **Stopped**

- The number of source objects and their respective total size (if applicable)



Overview Panes

There are several other panes that give an overview for the chosen job group. These panes are as follows:

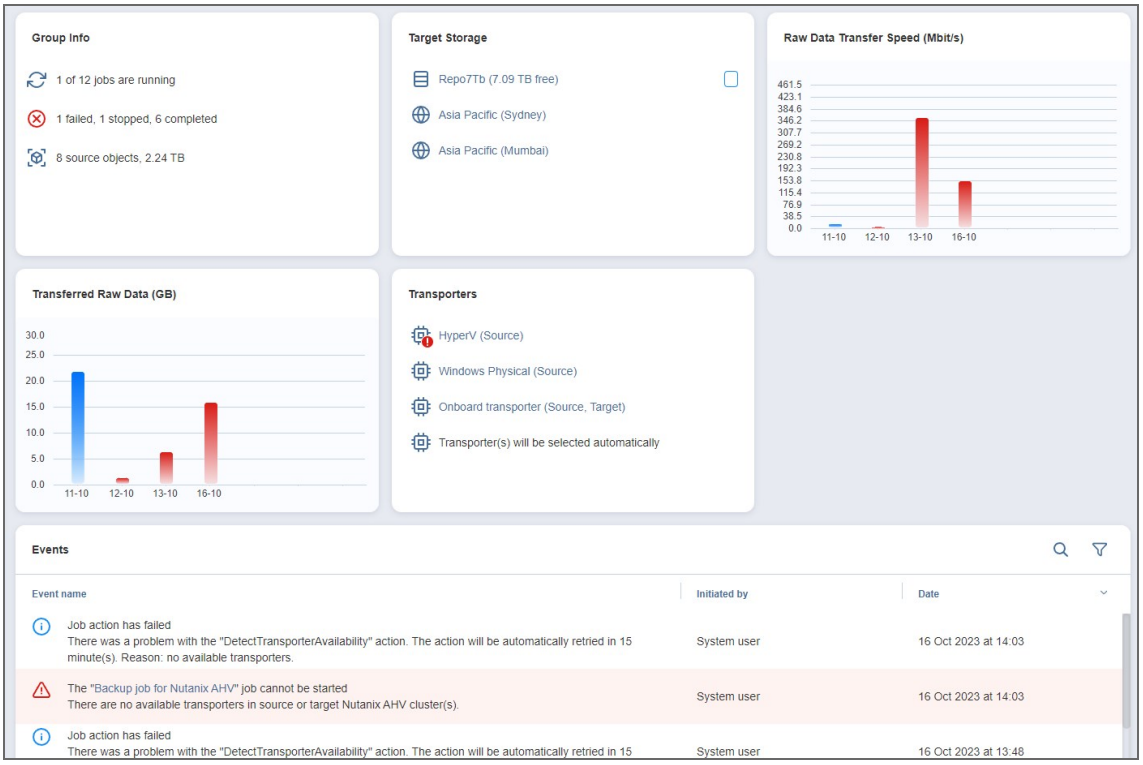
- **Target Storage**: The target storage(s) of the jobs in the chosen group. To open the popup with additional information about it (**Data storage type**, **Free space**, **Used space**, **Transporter**, **Path**, **Compression**, **Deduplication**, **Space savings** fields), click this target storage.

  > **Note**
  > If a federated repository is selected as a destination, only **Data storage type** (*Incremental with full backups*), **Free space**, **Used space** fields are displayed.

- **Raw Data Transfer Speed (Mbit/s)**: The raw data transfer speed for previous job runs (if no job in the group is currently running) or current job run (if a job is currently running). If a job run includes multiple backup objects, the aggregated data transfer speed of all backup objects is displayed.
- **Transferred Raw Data (GB)**: The amount of raw transferred data before compression/deduplication for the current job run or past job run(s).

  > **Note**
  > The **Raw Data Transfer Speed** and **Transferred Raw Data** widgets show:
  > - Raw (uncompressed) data size for disk-based backups.
  > - Actual transferred (compressed) data size for file-based backups (File-level backup and File share backup jobs).

- **Transporters**: Table of the Transporters used by the group of jobs
- **Events**: Table of alarms/notifications for the given group of jobs



Job Dashboard

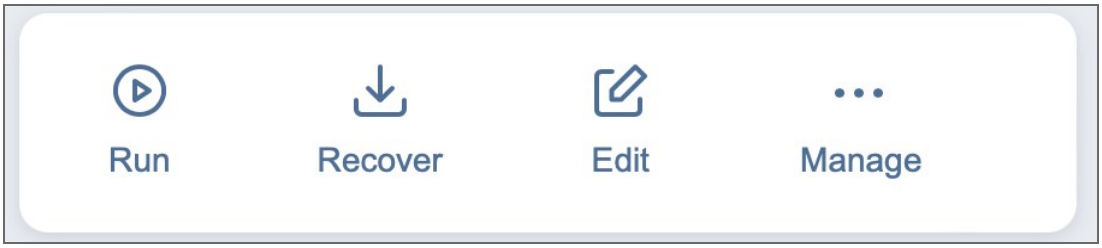
When selecting a specific job in the **Jobs** menu, the following information is displayed.

Action Bar

The job action bar contains the following four actions:

- **Run**: Opens the job run dialog box
- **Recover**: Brings up a list of recovery options for the given job (backup and replication jobs only)
- **Edit**: Opens the job edit wizard
- **Manage**: Brings up the options to **Clone**, **Merge**, **Rename**, **Create report** for, **Enable**/**Disable**, or **Delete** the job.

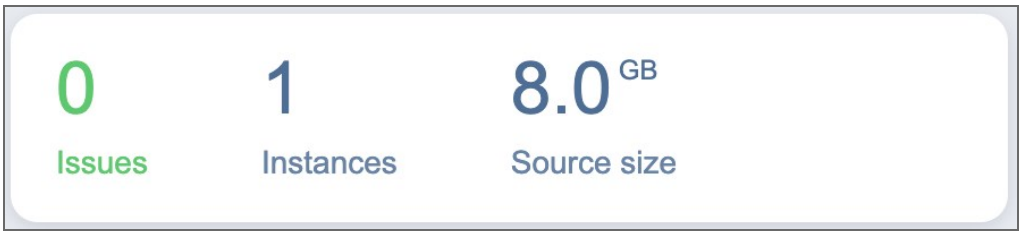
For more information on using the action bar, see “Running Jobs on Demand” on page 250.



Summary Bar

The summary bar displays information about a job. The data displayed is as follows:

- **Issues**: Total number of issues/alarms for the job. When clicked, the Alarms & Notifications dialog box is displayed.
- **Objects**: Total number of objects covered by the job
- **Source size**: Total size of the objects covered by the job

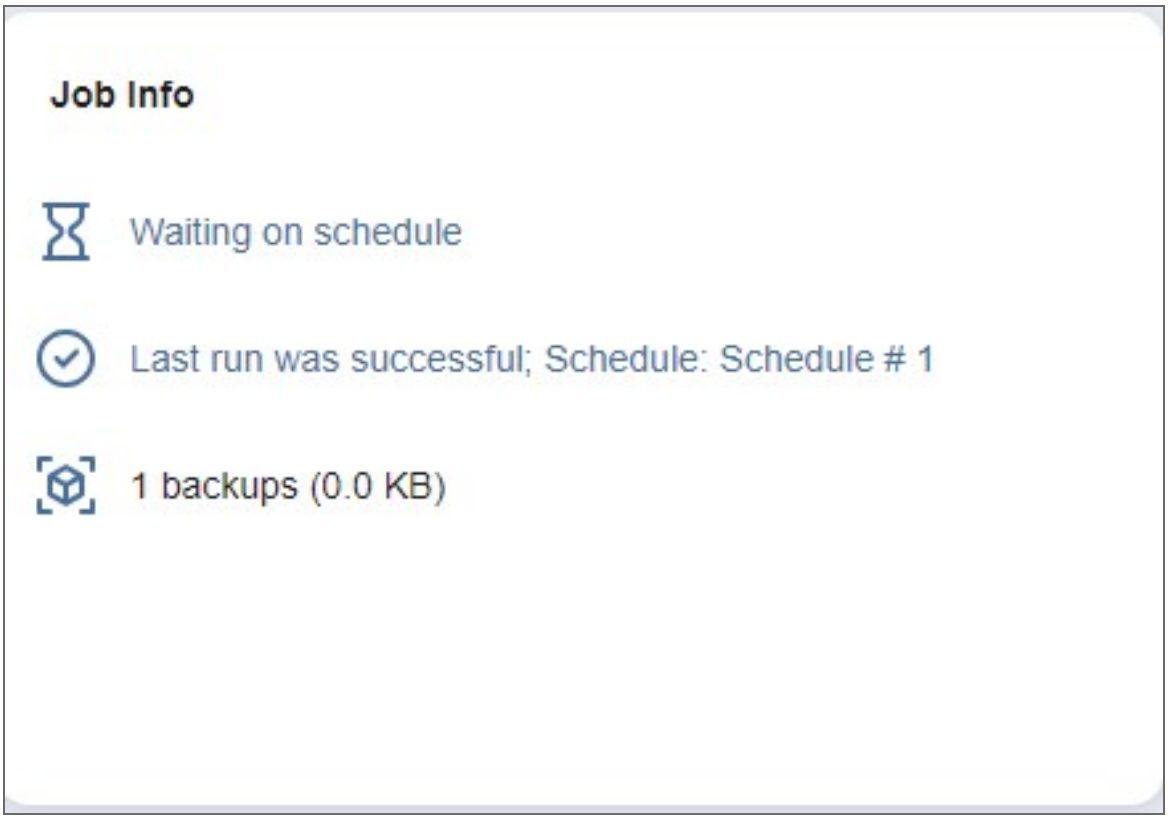


Job Info

This pane displays current information about the job. The information includes:

- The running schedule for this job. NAKIVO Backup & Replication displays the name of the schedule associated with the last job run with the following exceptions:
  - If the schedule name is not set, the product displays a number of the schedule.
  - If the job was run manually and in the manual run no specific schedule settings were used, the “Manual run” is displayed instead of the schedule name.
- The status of this job; **Successful**, **Failed**, **Running**, **Stopped**, or **Not executed yet**

- The number of source objects and their respective total size (if applicable)



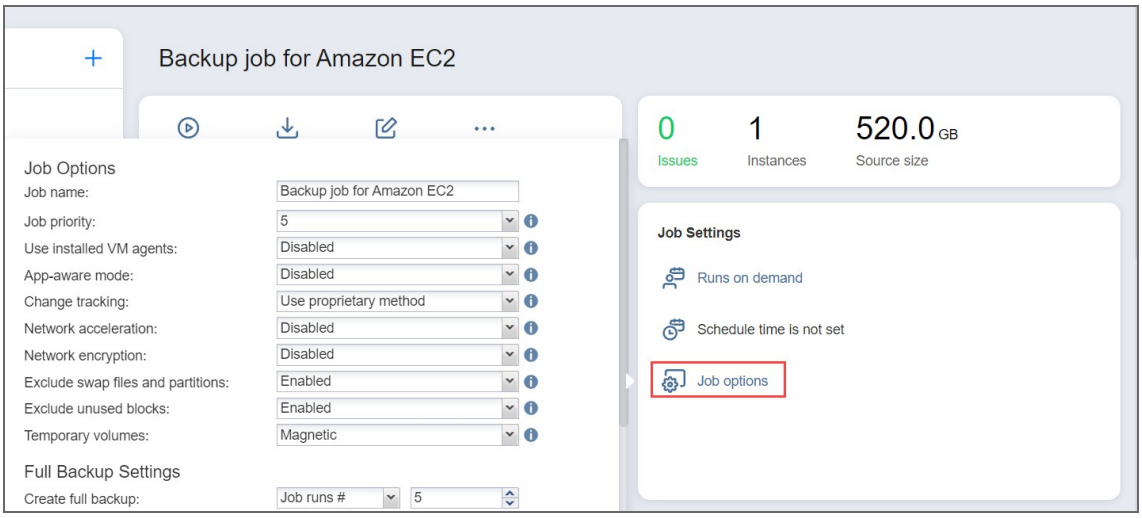
Job Settings

This pane allows you to view and edit certain options for a job. The settings displayed are as follows:

- The running schedule for this job (if applicable)
- Scheduled run time(s) for this job (if applicable)
- **Recovery points retention**: Clicking this opens a dialog box with **Retention Settings** for this job (if applicable)

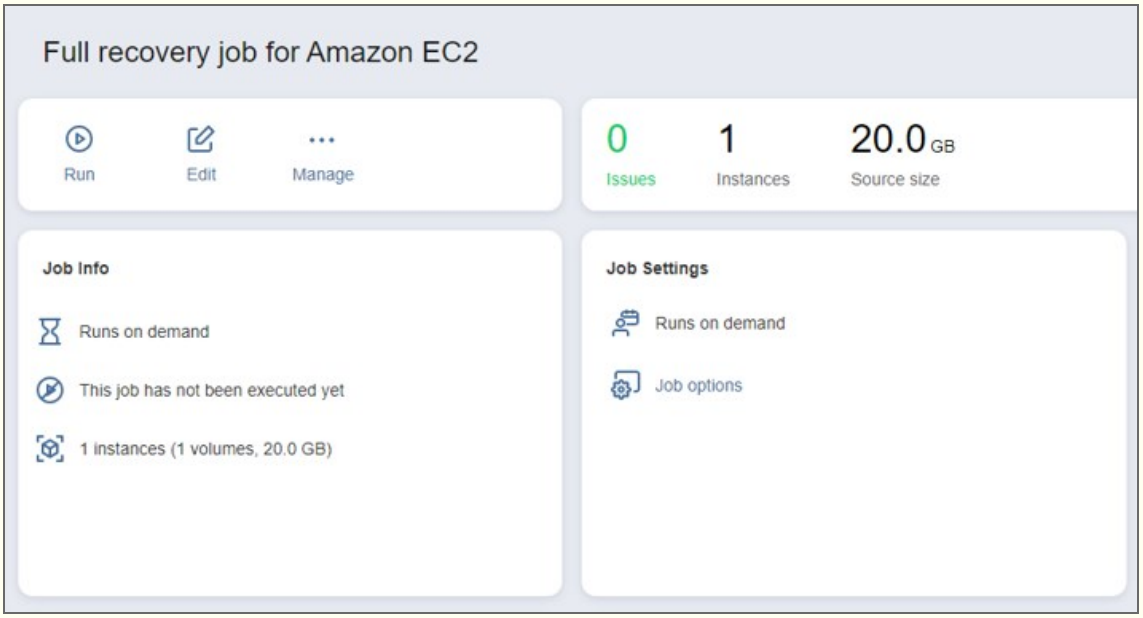
> **Note**
>
> This option is only available in the following cases:
>
> - Your version of NAKIVO Backup & Replication is older than v10.8.
> - You have updated NAKIVO Backup & Replication from a version older than v10.8 to v10.8 or newer and have not enabled the new scheduler for the respective job.
> - You enabled legacy retention in the Expert tab.

- **Job options**: Clicking this opens a dialog box with **Job Options** for this job.



Note

The **Job Settings** pane for NAKIVO Backup & Replication recovery jobs only displays the settings relevant to this particular job:

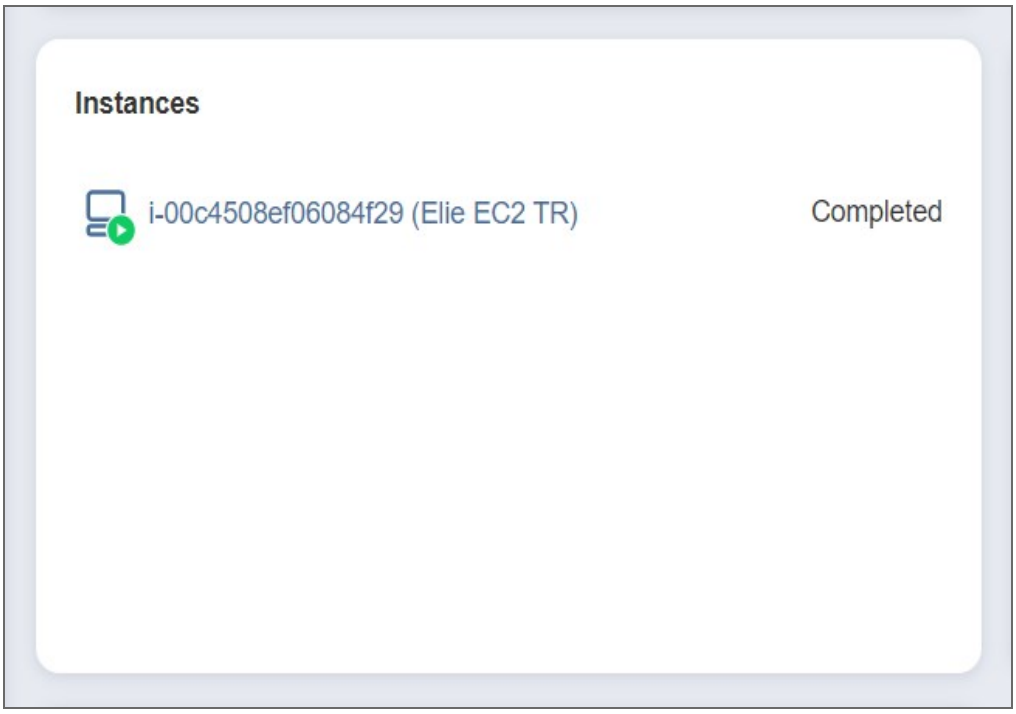


Job Objects

This pane displays a list of backup/replication/recovery objects based on the respective object type. The objects can be one of the following:

- Virtual Machines
- Instances
- Backups
- Physical Machines

- Databases
- Microsoft 365 Items
- File Share Items



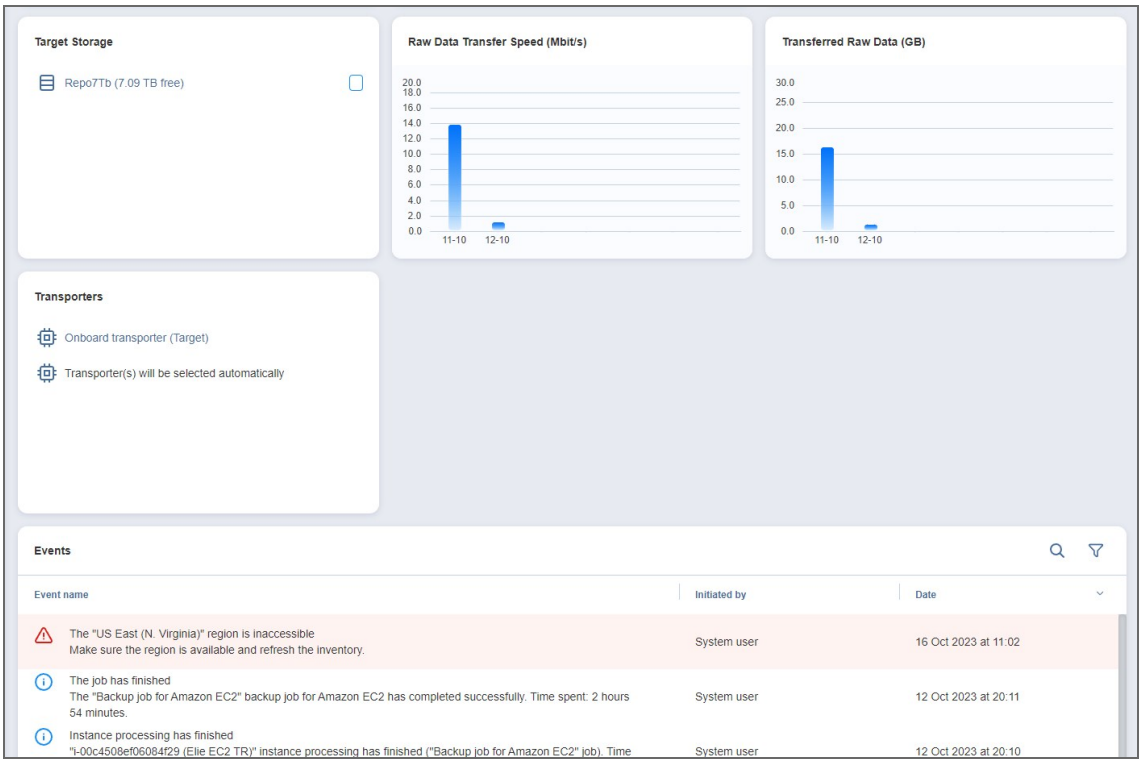
Overview Panes

Several other panes give an overview of the chosen job. These panes are as follows:

- **Target Storage**: The target storage of the chosen job
- **Raw Data Transfer Speed (Mbit/s)**: The raw data transfer speed for the current job run or previous job runs if the job is not currently running. If a job run includes multiple backup objects, the aggregated data transfer speed of all backup objects is displayed.
- **Transferred Raw Data (GB)**: The amount of raw transferred data before compression/deduplication for a current job run or past job run(s).

> **Note**
> The **Raw Data Transfer Speed** and **Transferred Raw Data** widgets show:
> - Raw (uncompressed) data size for disk-based backups.
> - Actual transferred (compressed) data size for file-based backups (File-level backup and File share backup jobs).

- **Transporters**: Table of the Transporters used by the job
- **Events**: Table of alarms/notifications for the given job



Running Jobs on Demand

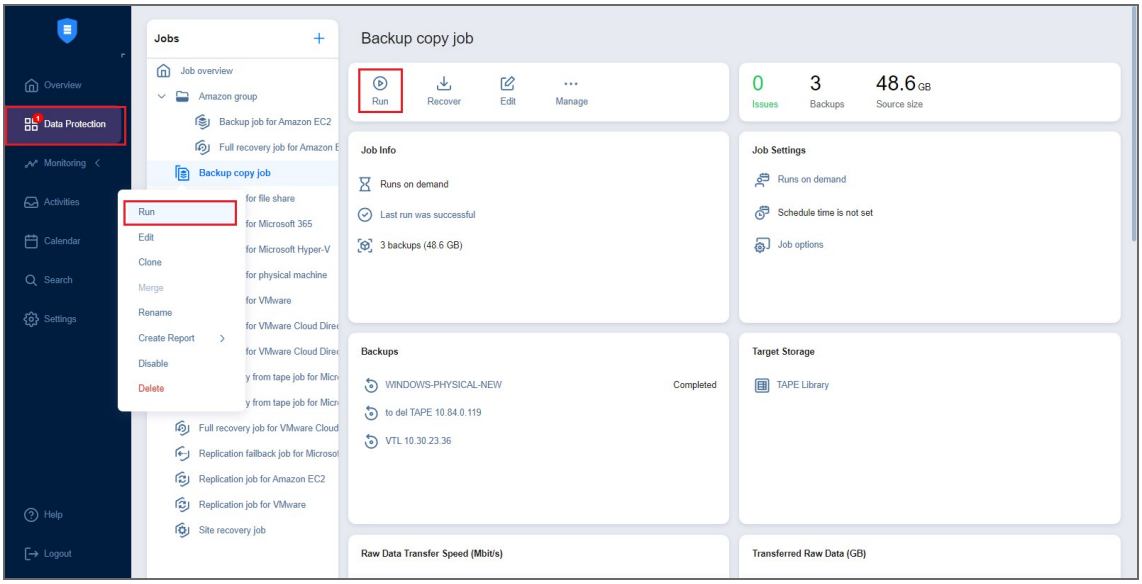
Use the **Data Protection** menu to start and stop jobs on demand.

- Starting Jobs
- Stopping Jobs
- Managing Grouped Jobs

Starting Jobs

To start a job, follow the steps below:

1. Go to the **Data Protection** menu, select the job from the list of jobs, and click **Run**. Alternatively, right-click a job to bring up the action menu and click **Run**.

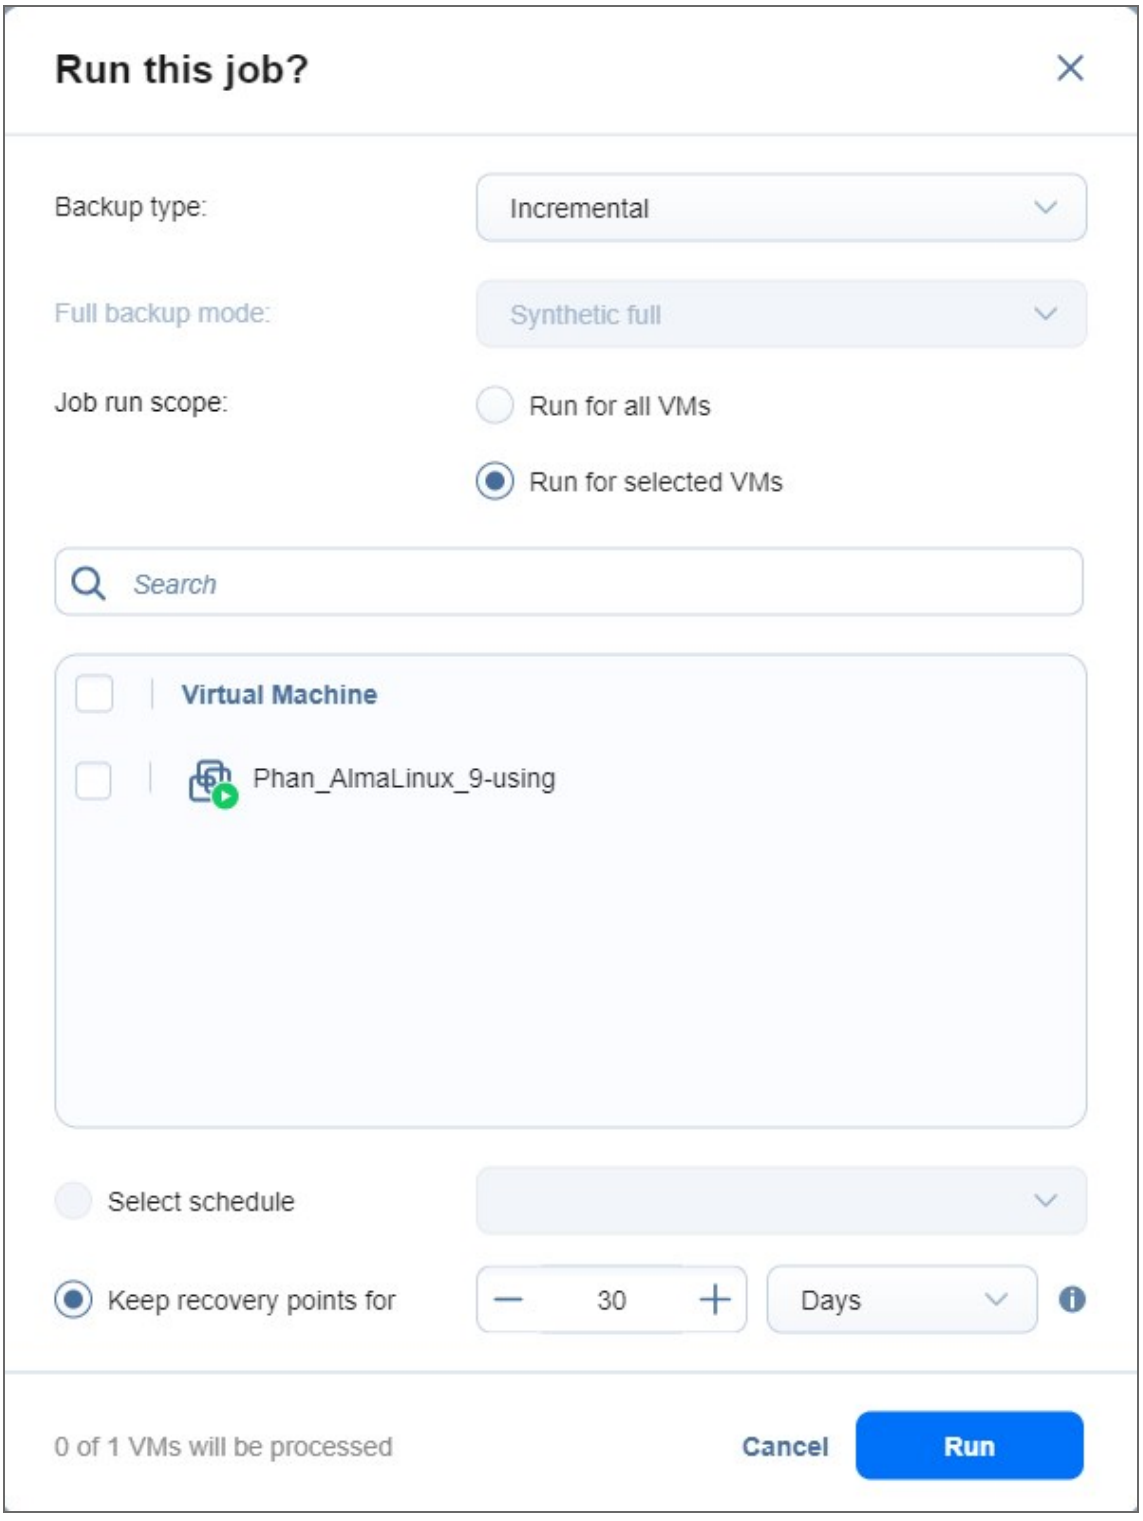


2. Choose one of the following options:
   - **Run for all VMs/VM templates/backups/physical machines/instances/databases/sites/accounts/items**: The job runs for all job objects.
   - **Run for selected VMs/VM templates/backups/physical machines/instances/databases/sites/accounts/items**: The job runs for the job objects that you select.
   - **Run for failed VMs/VM templates/backups/physical machines/instances/databases/sites/accounts/items**: If applicable, the job runs for previously failed job objects only.
   - **Run for stopped VMs/VM templates/backups/physical machines/instances/databases/sites/accounts/items**: If applicable, the job runs for previously job stopped objects only.
3. If backups in the Backup Repository selected for a job are stored in separate files, you have to choose between the following backup types:
   - **Incremental**: The job creates an incremental backup.
   - **Full**: The job creates a full backup. When you choose this option, choose one of the full backup modes:
     - **Synthetic full**: The application first creates an incremental backup—that is, transfers only the data that changed since the last backup—and then creates a new full backup using the last full backup and the chain of subsequent incremental recovery points.
     - **Active full**: The application reads all source machine data and transfers it to the backup repository to create a full backup.

4. For backup and backup copy jobs, you can use preconfigured retention settings by selecting **Use job retention** (legacy retention approach) or **Select schedule** (schedule retention approach), or specify custom retention settings for a manual job run by selecting **Keep recovery points for**.
    - **Use job retention**: Select this option to use the preconfigured legacy retention settings for a job run. If a previous run for this job was stopped or failed, the settings used for that run are selected by default.
    - **Select schedule**: Select this option to choose a preconfigured schedule and its retention settings for this job. If a previous run for this job was stopped or failed, the settings used for that run are selected by default. Recovery points created by a manual job run using this option are automatically assigned expiration dates.
    - **Keep recovery points for**: The recovery points created by this job run are kept for the specified period of time and then expire. The expired recovery points are removed during the following job run.

    > **Note**
    > If a job does not support retention or has the **Do not schedule, run on demand** option selected, only the **Keep recovery points for** option will be available.

5. Click the **Run** button to confirm your selection.

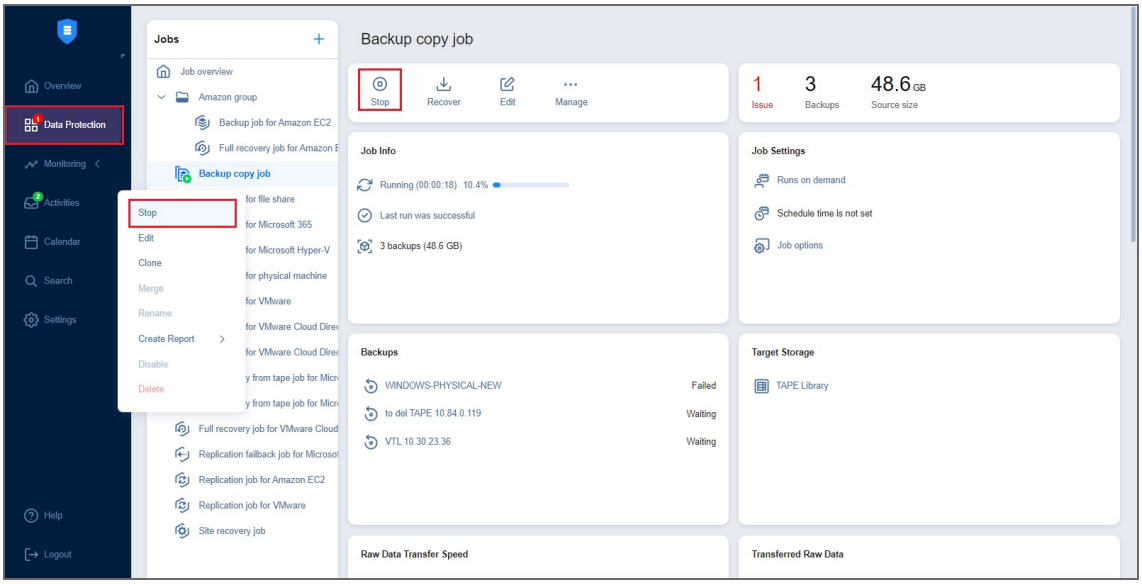


The product will close the dialog box and start running your job.

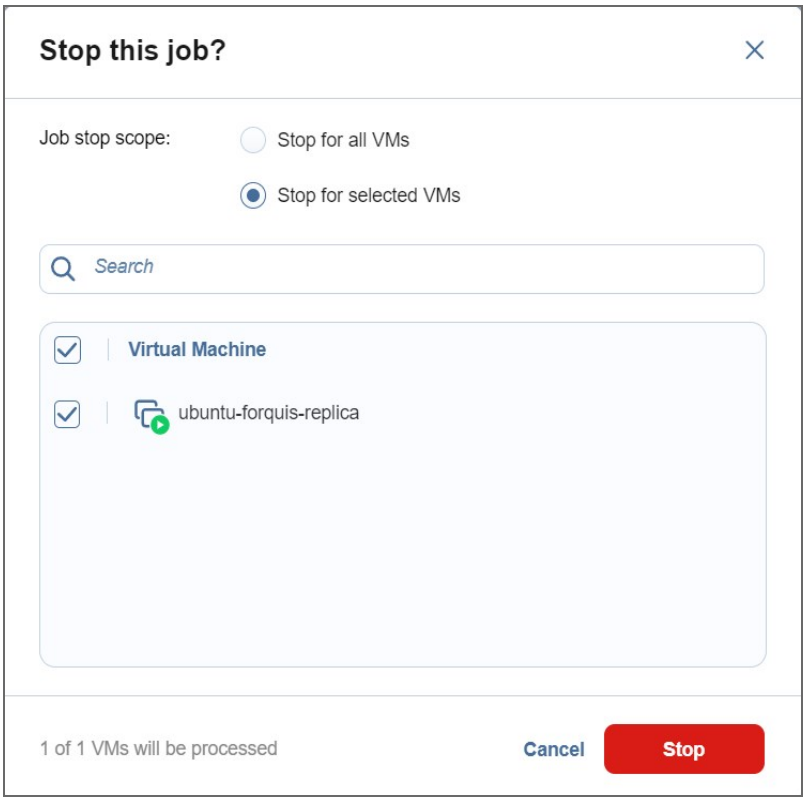
Stopping Jobs

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

1. Go to the **Data Protection** menu, select the job from the list of jobs, and click **Stop** . Alternatively, right-click a job to bring up the action menu and click **Stop**.



2. In the dialog box that opens, choose one of the following:
   - **Stop for all VMs/VM templates/backups/physical machines/instances/databases/sites/accounts/items**: The job stops for all job objects.
   - **Stop for selected VMs/VM templates/backups/physical machines/instances/databases/sites/accounts/items**: The job stops for the job objects that you select.
3. Click the **Stop** button in the dialog box to confirm your selection.

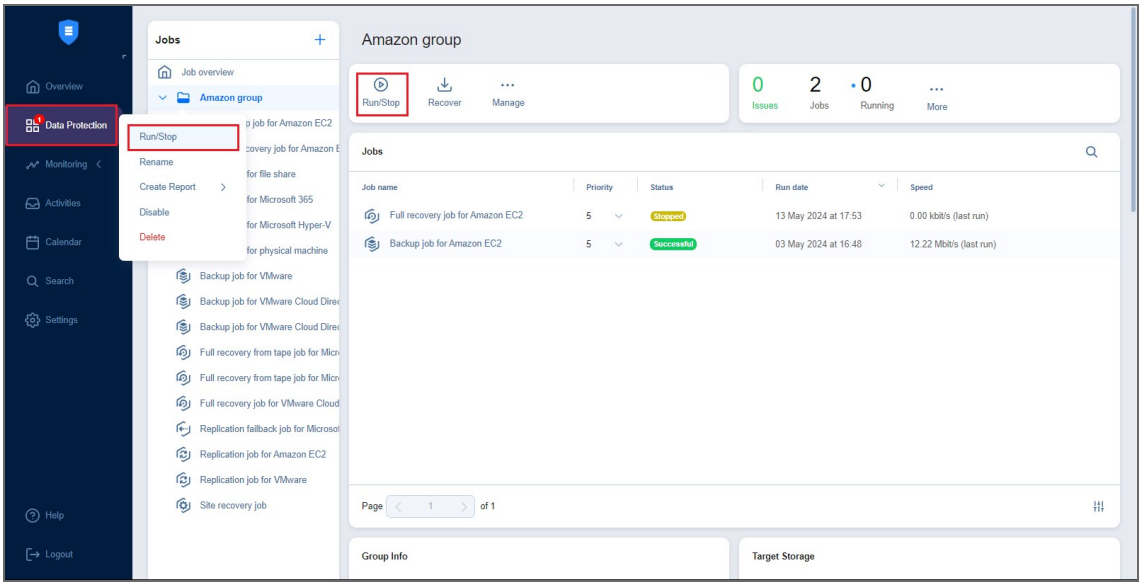


The application closes the dialog box and stops your job.

Managing Grouped Jobs

To efficiently start or stop jobs in bulk (run all failed jobs, for example), follow these steps:

1. From the **Data Protection** menu, select the needed job group and click **Run/Stop**. To manage all jobs and groups at once, select **Overview** and click **Run/Stop**. Alternatively, right-click on the needed job group to bring up the action menu and click **Run/Stop**.



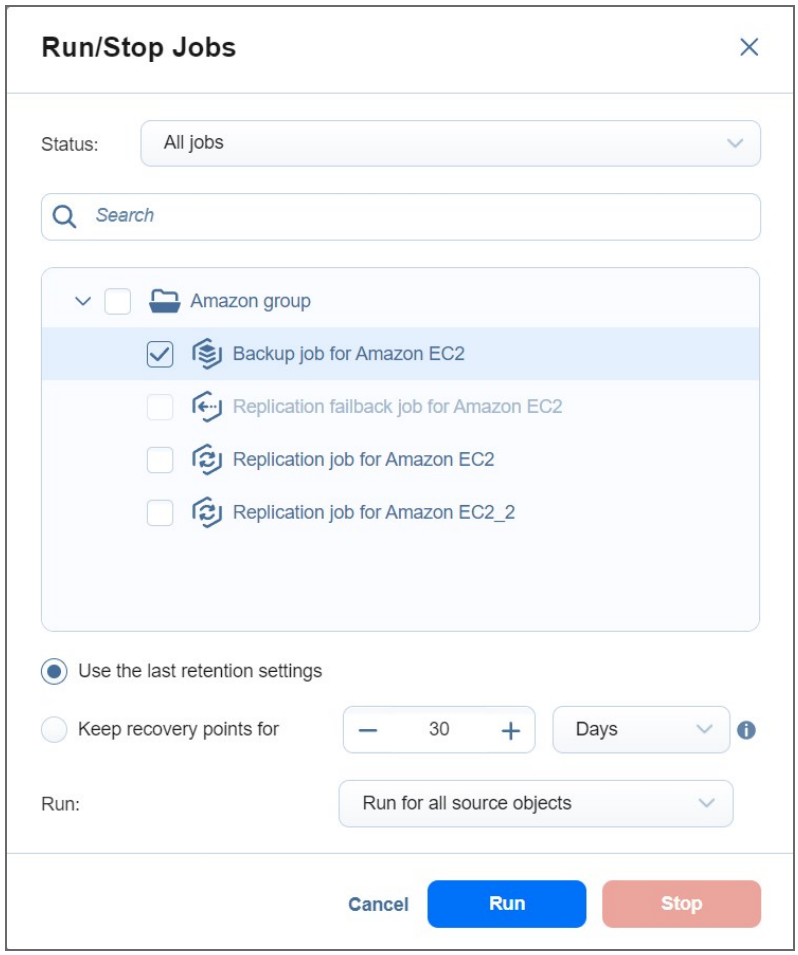
2. In the drop-down **Status** menu, select one of the following:

- **All jobs:** Displays all jobs in the group
- **Failed jobs:** Displays all failed jobs in the group
- **Stopped jobs:** Displays all stopped jobs in the group

3. Select the jobs you want to run/stop.
   a. When running backup or backup copy jobs, specify the retention settings with one of the following options:
      - **Use the last retention settings**: Select this option to use the retention settings from the last job run for the manual job run.
      - **Keep recovery points for**: The recovery points created by this job run are kept for the specified period of time and then expire. The expired recovery points are removed during the following job run.

      > **Note**
      > If the group of jobs contains at least one job that isn't a backup/backup copy job, does not support retention, or has had its retention settings changed since the previous run, only the **Keep recovery points for** option will be available.

   b. In the lowest drop-down menu, specify (if applicable) whether you want the operation to run for failed source objects, stopped source objects, or all source objects.
4. Click the **Run** or **Stop** button to confirm your selection.



Managing Jobs

Using the **Data Protection** menu, you can easily manage jobs. Use the **Manage** menu to rename, edit, merge, delete, and enable/disable jobs.

- Viewing Job Run Details
- Renaming Jobs
- Editing Jobs
- Cloning Jobs
- Merging Jobs
- Deleting Jobs
- Disabling and Enabling Jobs
- Grouping Jobs
- Creating Groups
- Creating Job Reports

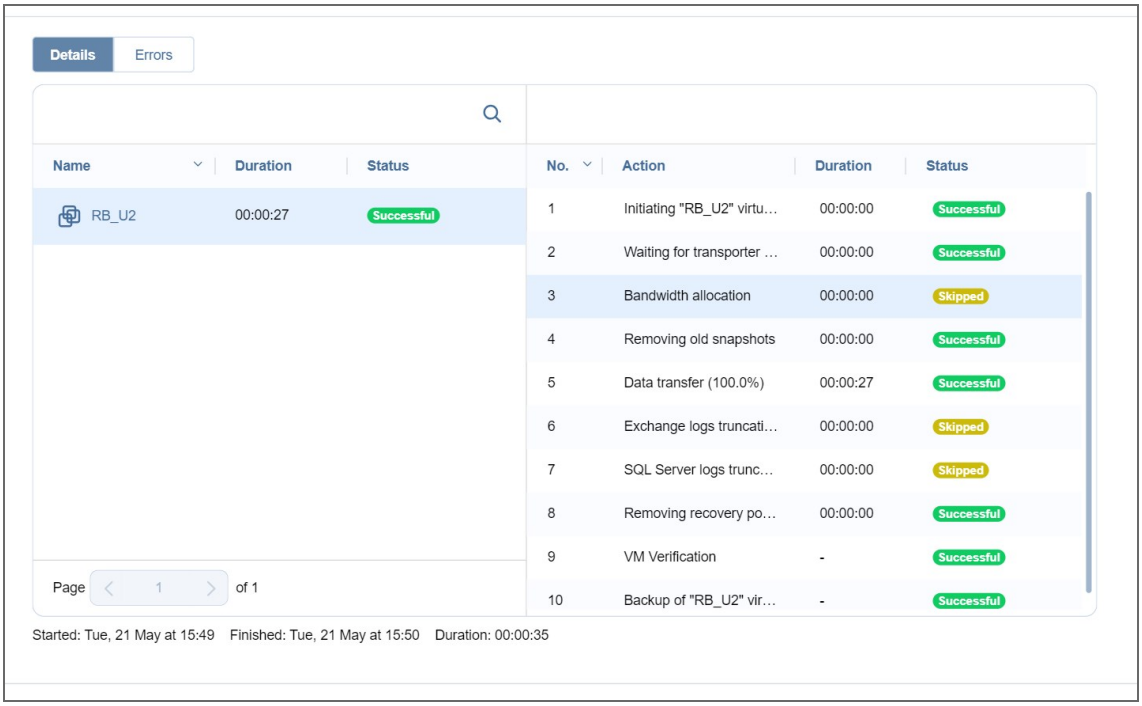
Viewing Job Run Details

You can view job run details in the **Data Protection** or **Activities** dashboard.

- **Jobs dashboard**: In the **Jobs** pane on the right, click the job status to view details.
- **Activities dashboard**: Click on the job name and then click on the **details** link.

The **Job Run Details** popup window opens displaying the following details:

- On the **Details** tab:
  - **Name**: The column displays the name of the job object.
  - **Duration**: The column displays the real-time duration for job processing. For completed jobs, the column shows the total duration for job processing.
  - **Status**: The column displays the last run status of the job.



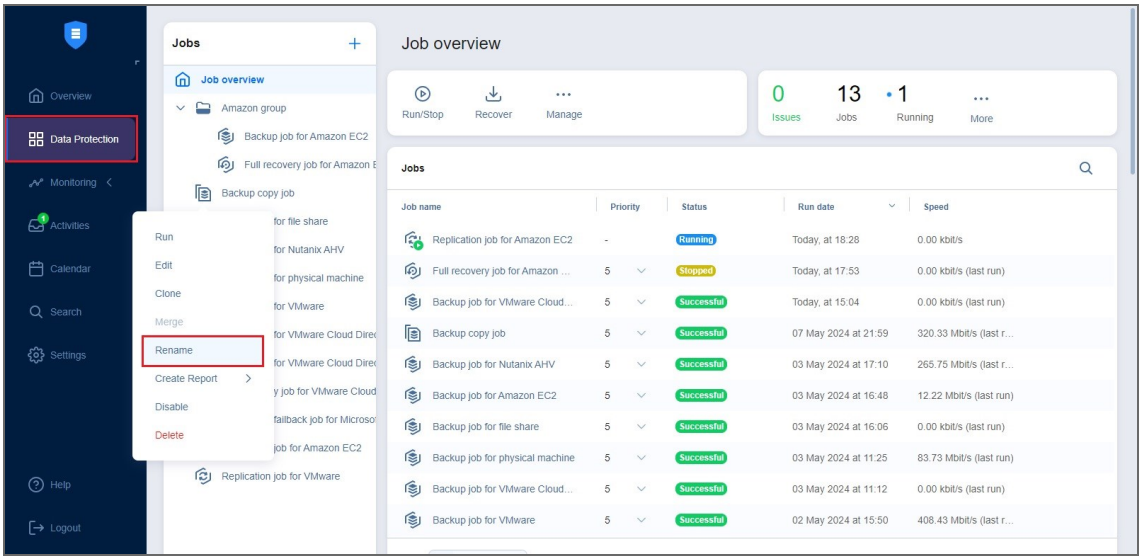
Notes

- You can click on any of the columns in the window to sort the list alphabetically or numerically. Additionally, you can use the search field for the specific job object.
- On the right side, you can find the list of the job runs for this job and their durations. Clicking each specific run displays the detailed list of actions that were performed during the job run, their status, and the amount of time each of these actions took.
- If the job was never retried, the right pane shows the actions that NAKIVO Backup & Replication took when running this job.

- On the **Errors** tab, you can view all the alarms associated with this job.

Renaming Jobs

1. From the list of jobs, right-click on the job you want to rename.
2. Click **Rename**.
3. In the dialog box that opens, specify the new name for the job and click **Rename**.



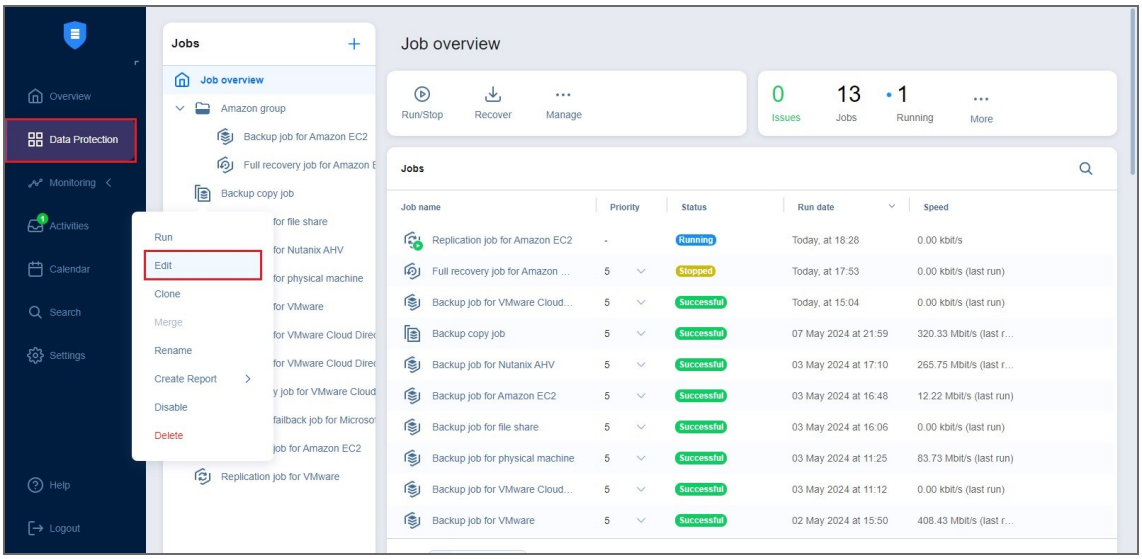
> **Note**
> You can also rename jobs by selecting a job and clicking **Manage** > **Rename**.

Editing Jobs

To edit a job, follow the steps below:

1. Right-click on the job you want to edit from the list of jobs.
2. Click **Edit**.
3. In the **Edit** wizard, click the needed page to open it for editing.

4. Make the required changes and click **Save** or **Save & Run**.



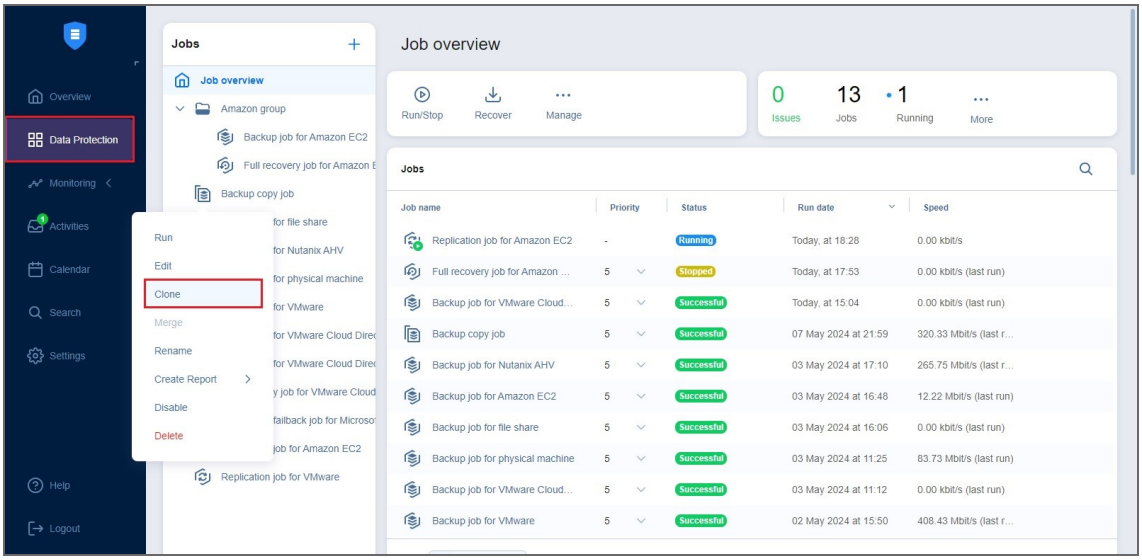
Notes

- You can edit the job while it is running, but the changes will be applied only when the job run has completed.
- You can also edit jobs by selecting a job and clicking **Manage** > **Edit**.

Cloning Jobs

To clone a job, follow the steps below:

1. Right-click on the job you want to clone from the list of jobs.
2. Click **Clone**.



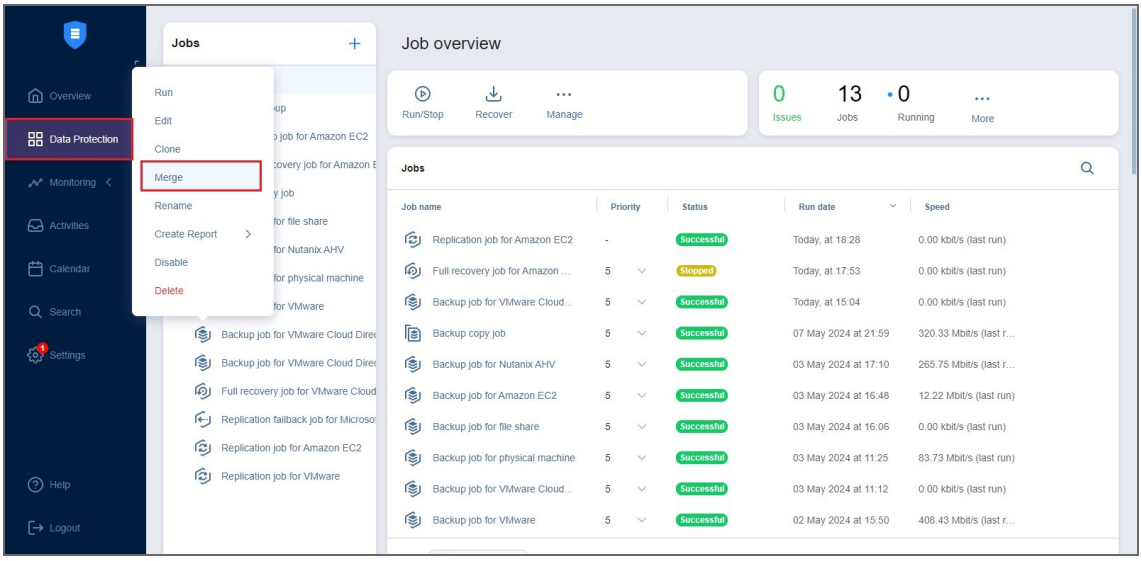
Note

You can also clone jobs by selecting a job and clicking **Manage** > **Clone**.

Merging Jobs

NAKIVO Backup & Replication allows you to merge jobs of the same type. Before doing this, make sure to check feature requirements and how the feature works. To merge the jobs, do the following:

1. From the list of jobs, right-click on the source job you want to merge.
2. Click **Merge**.
3. Choose the target job for the merge and click **Apply**.
4. After the merge is finished, click **Close** to close the popup.



Notes

The **Merge** button may be unavailable in the following cases:

- The selected job does not meet the feature requirements.
- The selected job is currently running.
- There are no target jobs available to merge the selected job with.

Deleting Jobs

To delete a job follow the steps below:

1. Right-click on the job you want to delete from the list of jobs.
2. Click **Delete**.
3. In the dialog box that opens, select one of the following:
   - **Delete job and keep backups**
   - **Delete job and backups**

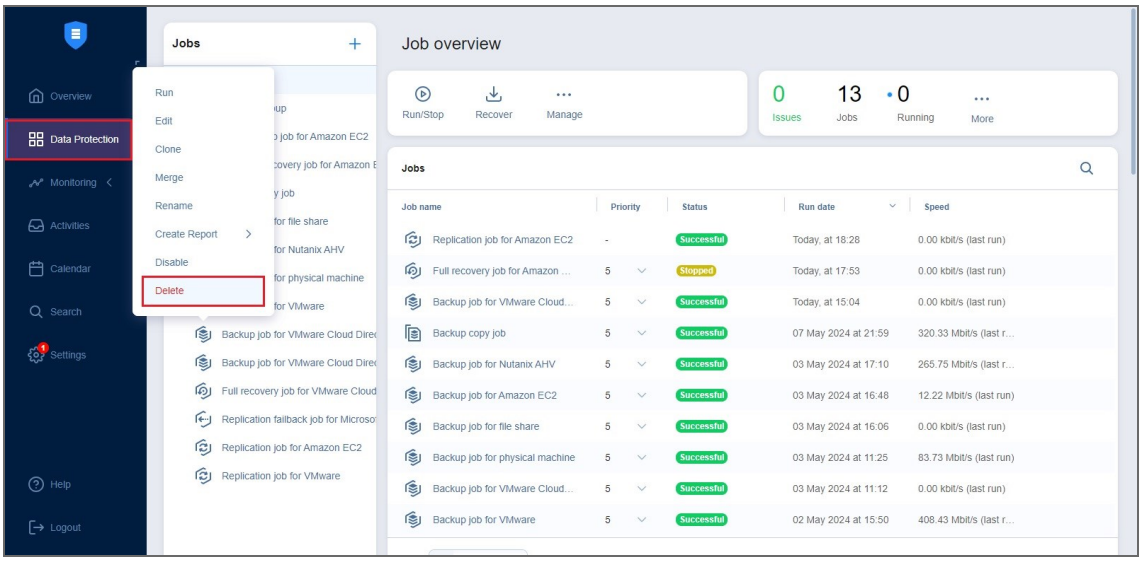
Note

If the job has target objects, the options in the dialogue box allow you to keep or delete the job and either of the following:

- backups
- replicas
- recovered VMs
- recovered instances
- replicas and their journals

Note that replicas, recovered VMs, and recovered instances are deleted in the powered off state.

4. Click **Delete**



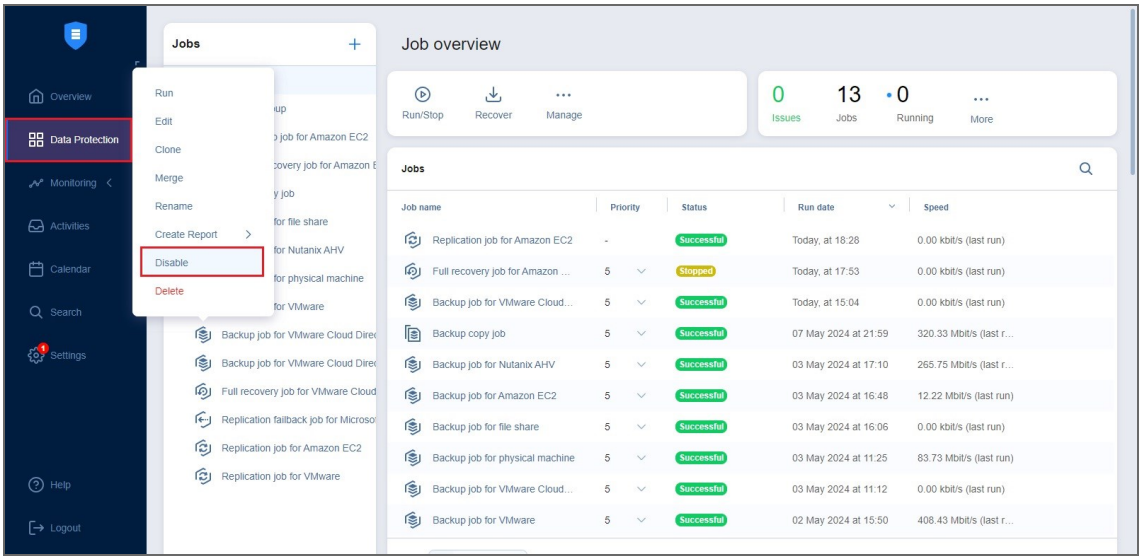
Notes

- You can also delete jobs by selecting a job and clicking **Manage** > **Delete**.
- Backups can also be deleted from Backup Repositories.

Disabling and Enabling Jobs

NAKIVO Backup & Replication provides you with the ability to disable jobs. A disabled job does not run based on the schedule and cannot be run on demand.

1. From the list of jobs, right-click on the job you want to disable.
2. Click **Disable**.



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

> **Note**
> You can also manage jobs by selecting a job and selecting the desired action from the **Manage** menu.

Grouping Jobs

Groups are folders which allow you to:

- Logically arrange jobs (to represent organizations, locations, services, etc.).
- Perform bulk actions with all or selected jobs in a group.

Creating Groups

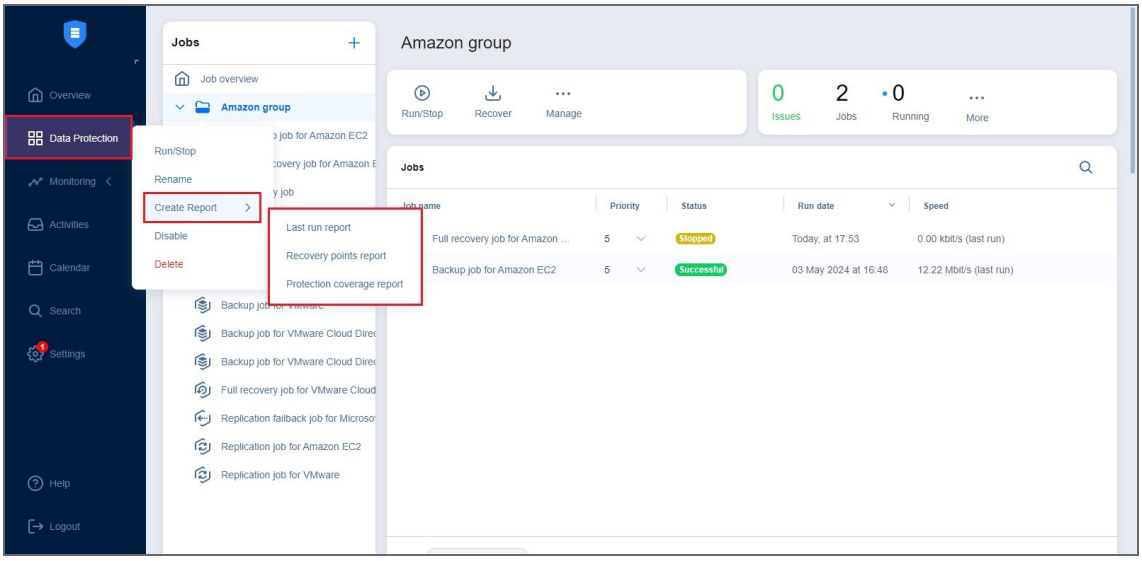
To create a group, follow the steps below:

1. In the **Data Protection** menu, click **Create** and then click **Job group**.
2. Type in the group name in the dialog box that opens and click **Create**.

The following actions are available to manage groups:

- To add a job to a group, simply drag the job into the group.
- To remove a job from the group, drag the job outside the group.
- To delete a group, right-click the group and choose **Delete** from the shortcut menu that opens. Confirm the group deletion when prompted to do so. Note that when deleting a group, the jobs in the group are not deleted. The jobs are moved to the parent group (or to *Overview*).
- To rename a group, double-click the group and enter a new name.
- To enable or disable all jobs inside a group, click the **Enable/Disable** switch.
- To run jobs available in a group, click **Run/Stop** and then click **Run** Jobs. In the dialog box that opens, select the jobs you want to run and click **Run Jobs**.

- To stop running the jobs available in a group, click **Run/Stop** and then click **Stop Jobs**. In the dialog box that opens, select the jobs you want to stop and click **Stop Jobs**.
- To create a report, click **Create Report**. In the dialog box that opens, select one of the following reports from the **Create Report** menu:
  - **Last run report**: Provides data on the last run of the job. Choose either the PDF or CSV format and click **Create**.
  - **Recovery points report**: Contains information regarding the sizes of recovery points for backups/replicas for the chosen job or jobs.
  - **Protection coverage report**: Contains information about all VMs and instances protected by backup/replication jobs, as well as about all unprotected VMs and instances.



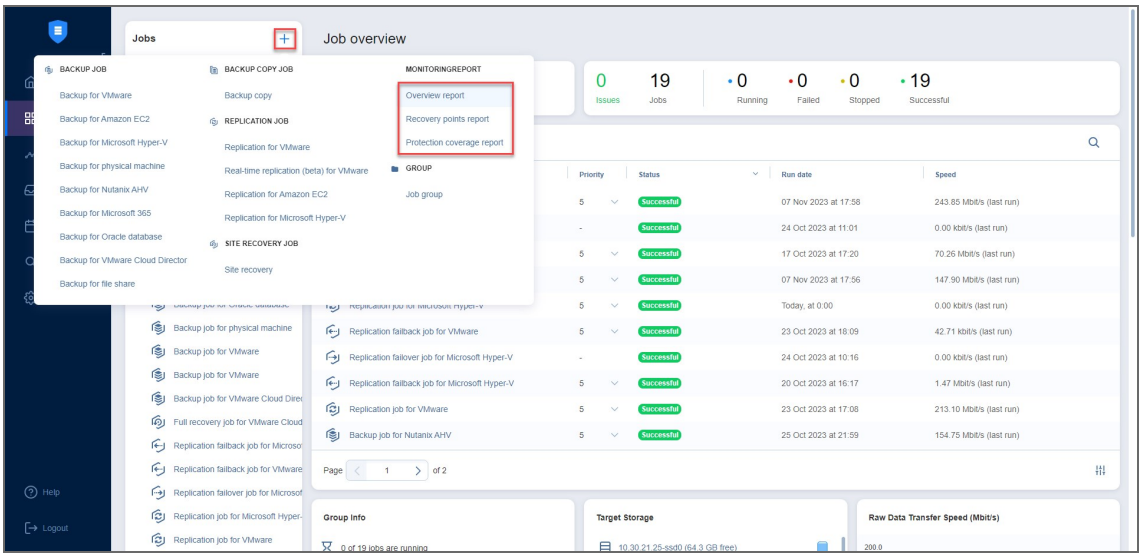
- Choose either the PDF or CSV format and click **Create**.

Creating Job Reports

To create a general report for all your jobs:

1. Select **Overview** in the **Data Protection** menu.
2. Click **Create Report**.
3. Choose one of the following reports:
   - **Overview report**: Contains information about the status and errors of all jobs.
   - **Recovery points report**: Contains information about the recovery points of backups/replicas for the chosen job or jobs created within a specific period.
   - **Protection coverage report**: Contains information about all VMs and instances protected by backup/replication jobs, as well as about all unprotected VMs and instances. Choose either the PDF or CSV format for the **Protection coverage report** and click **Create**.

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



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

1. Go to the list of jobs.
2. Select the job that you need to generate a report for and right-click it or click **Create**.
3. Select one of the following reports from the **Create Report** menu:
   - **Last run report**: Provides data on the last run of the job.
   - **Job history report**: Provides data on job runs that occurred during a specified time period. To generate a report, pick a start date on the left and finish date on the right side of the popup and click **Create**.

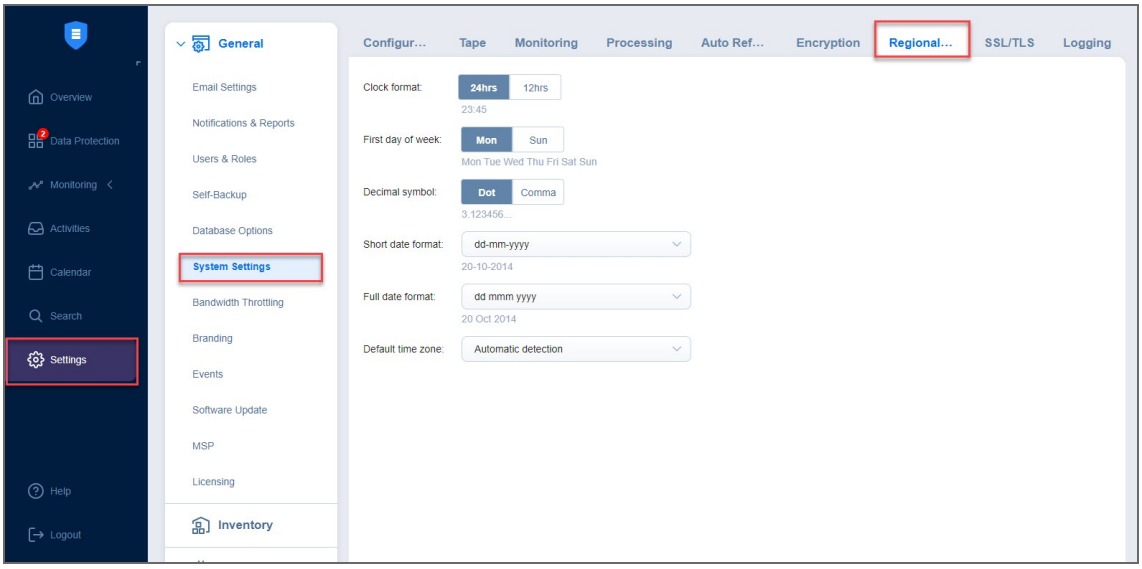
     Job history reports feature the **Calendar View** hyperlink. Clicking on it opens the high-level calendar view page with the summary of the backup status of every backup item within a single job.

     The report data is organized in separate blocks of weeks each consisting of two columns:
     - **Item Name** - displays the job object(s) name as ordered in the corresponding job
     - **Day of the week (Date)** - displays the status for processing each job object

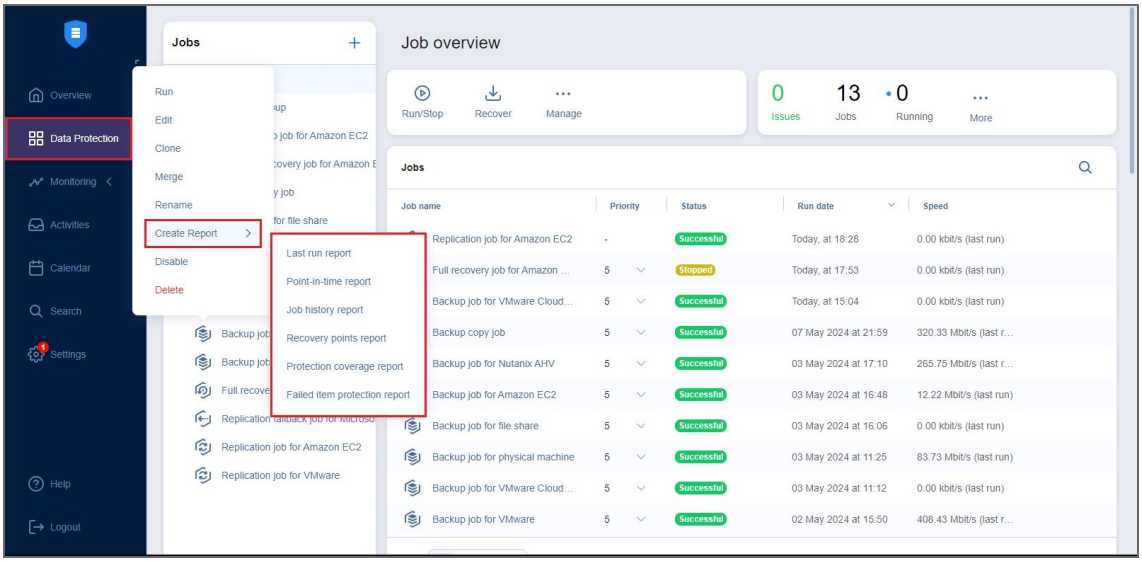
       Possible statuses:
       - **Successful** – displayed for successful object processing
       - **Failed** – displayed for failed object processing
       - **Skipped** – displayed for skipped object processing
       - **Stopped** – displayed for stopped object processing
     - The status field is empty if no object processing was scheduled on a particular day.
     - The columns representing weekdays that are not included in the reported range are empty.
     - Weekdays are displayed when there is at least one run on the week.

Notes

- The enhanced **Calendar View** report data is displayed from old to new for all reports where the **Job Runs** block is shown except for the **Monitoring** area and the **Point-in-time report**.
- To customize the first day of the week, go to the **Settings > General tab > System Settings > Regional Format** tab and select the needed day in the **First day of week** selector.



- **Point-in-time report**: Provides data on a particular job run. To generate a report, choose a date in the popup and click **Create.**
- **Recovery points report**: Contains data of recovery points created within specific period.
- **Protection coverage report**: Contains information about all VMs and instances protected by backup/replication jobs, as well as about all unprotected VMs and instances.
- **Failed item protection report**: Contains information about job objects for which processing failed during the last job run. Only backup and replication jobs are included.
- **Site recovery job report**: Contains a summary of the site recovery job, including the result of passing the **Recovery time objective** value, information about all actions performed, and all registered alarms and notifications.


Overview
Data Protection
Monitoring
Activities
Calendar
Search
Settings
Help
Logout
Jobs
Run
Edit
Clone
Merge
Rename
Create Report
Disable
Delete
Last run report
Point-in-time report
Job history report
Recovery points report
Protection coverage report
Failed item protection report
Replication job for Amazon EC2
Replication job for VMware
Job overview
Run/Stop
Recover
Manage
0
Issues
13
Jobs
0
Running
More
Jobs
Job name
Priority
Status
Run date
Speed
Replication job for Amazon EC2
Today, at 18:28
0.00 kbit/s (last run)
Full recovery job for Amazon ...
Stopped
Today, at 17:53
0.00 kbit/s (last run)
Backup job for VMware Cloud...
Today, at 15:04
0.00 kbit/s (last run)
Backup copy job
07 May 2024 at 21:59
320.33 Mbit/s (last r...
Backup job for Nutanix AHV
03 May 2024 at 17:10
265.75 Mbit/s (last r...
Backup job for Amazon EC2
03 May 2024 at 16:48
12.22 Mbit/s (last run)
Backup job for file share
03 May 2024 at 16:06
0.00 kbit/s (last run)
Backup job for physical machine
03 May 2024 at 11:25
83.73 Mbit/s (last run)
Backup job for VMware Cloud...
03 May 2024 at 11:12
0.00 kbit/s (last run)
Backup job for VMware
02 May 2024 at 15:50
408.43 Mbit/s (last r...

Job Alarms and Notifications

NAKIVO Backup & Replication displays:

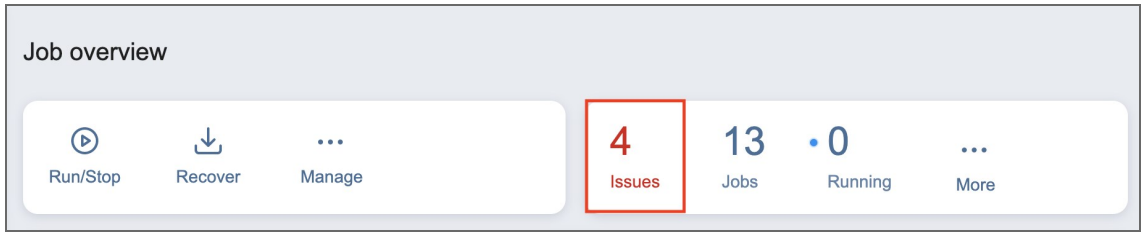
- **Alarms:** Job failures
- **Notifications**: Infrastructure changes and minor errors that do not lead to processing failure

For details, refer to the following sections:

- Viewing Alarms and Notifications
- Dismissing Alarms and Notifications

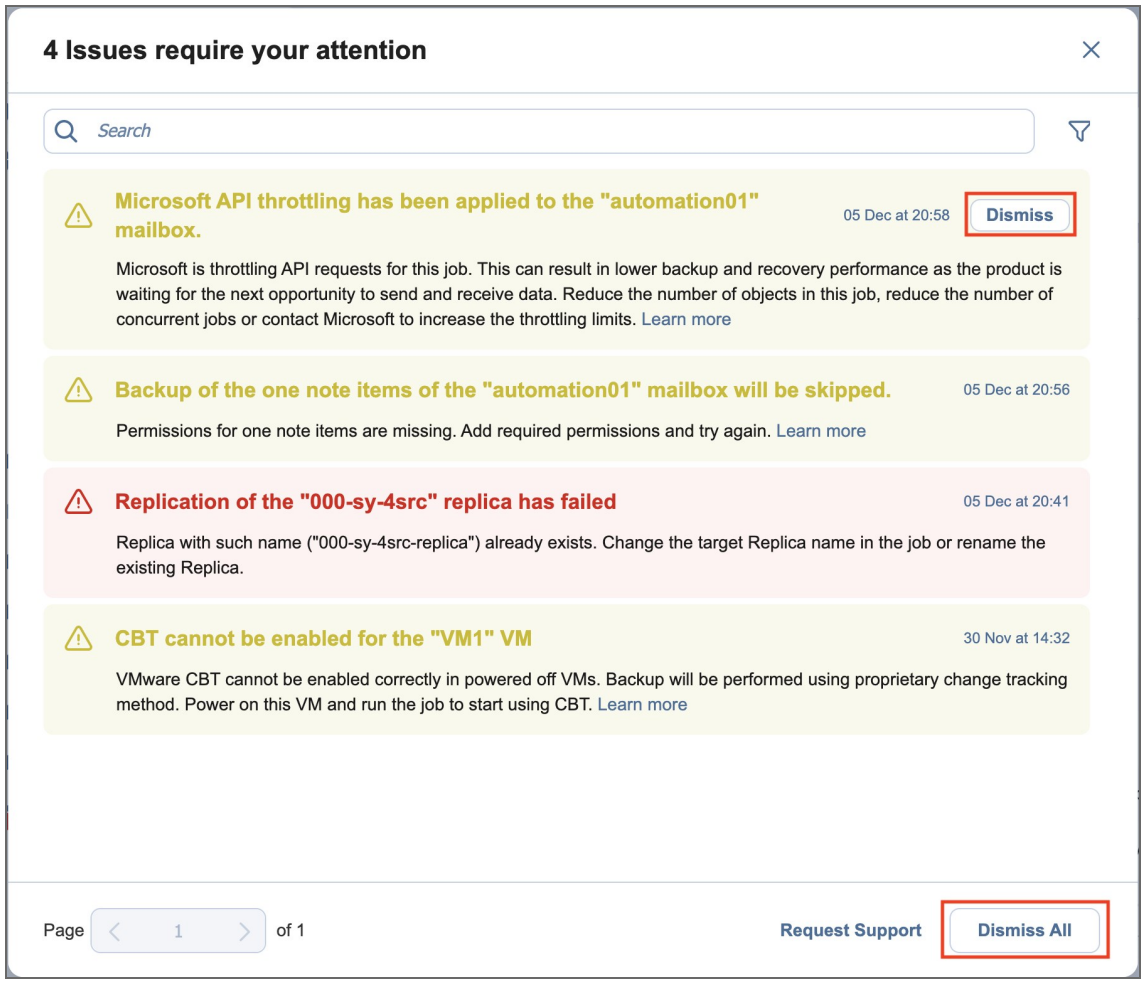
Viewing Alarms and Notifications

To view alarms and notifications, click the red **Issues** number in the **Summary bar**.



Dismissing Alarms and Notifications

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


4 Issues require your attention
Search
Microsoft API throttling has been applied to the "automation01" mailbox.
05 Dec at 20:58
Dismiss
Microsoft is throttling API requests for this job. This can result in lower backup and recovery performance as the product is waiting for the next opportunity to send and receive data. Reduce the number of objects in this job, reduce the number of concurrent jobs or contact Microsoft to increase the throttling limits. Learn more
Backup of the one note items of the "automation01" mailbox will be skipped.
05 Dec at 20:56
Permissions for one note items are missing. Add required permissions and try again. Learn more
Replication of the "000-sy-4src" replica has failed
05 Dec at 20:41
Replica with such name ("000-sy-4src-replica") already exists. Change the target Replica name in the job or rename the existing Replica.
CBT cannot be enabled for the "VM1" VM
30 Nov at 14:32
VMware CBT cannot be enabled correctly in powered off VMs. Backup will be performed using proprietary change tracking method. Power on this VM and run the job to start using CBT. Learn more
Page
1
of 1
Request Support
Dismiss All

Managing Activities

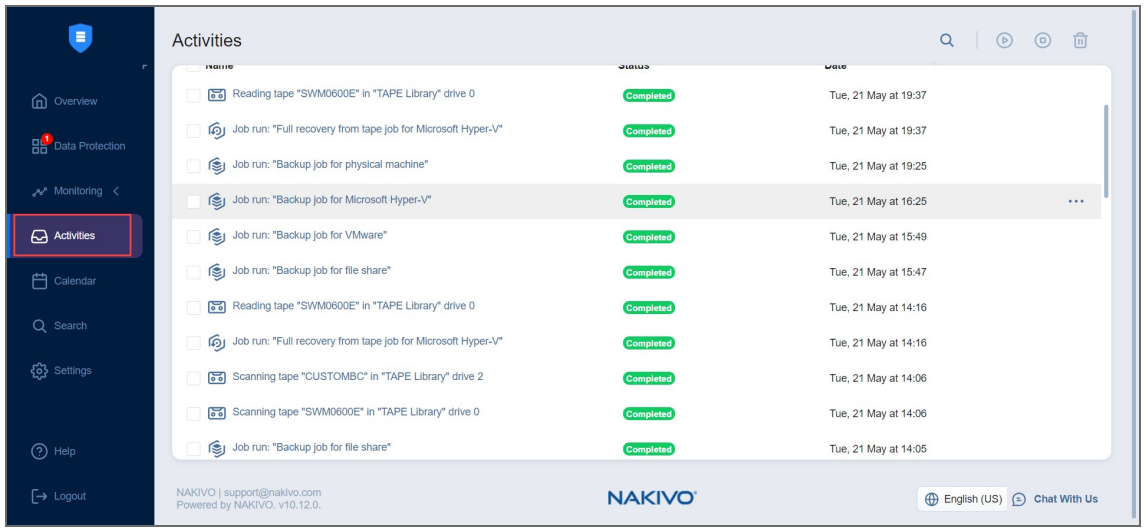
The **Activities** page displays current and past tasks performed by NAKIVO Backup & Replication. From this dashboard, the following actions can be done:

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

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

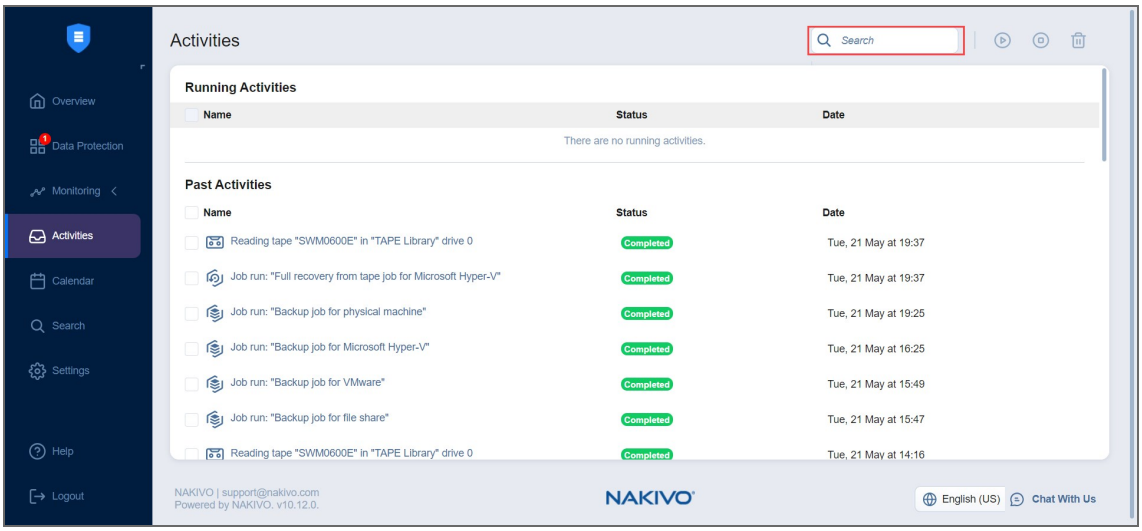
Viewing Activities

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



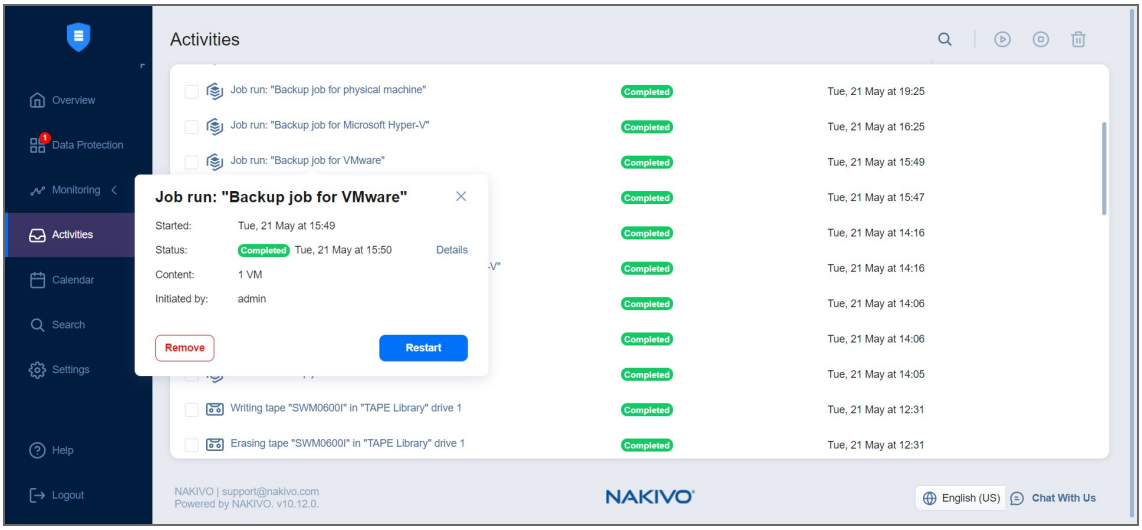
Searching for Activities

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



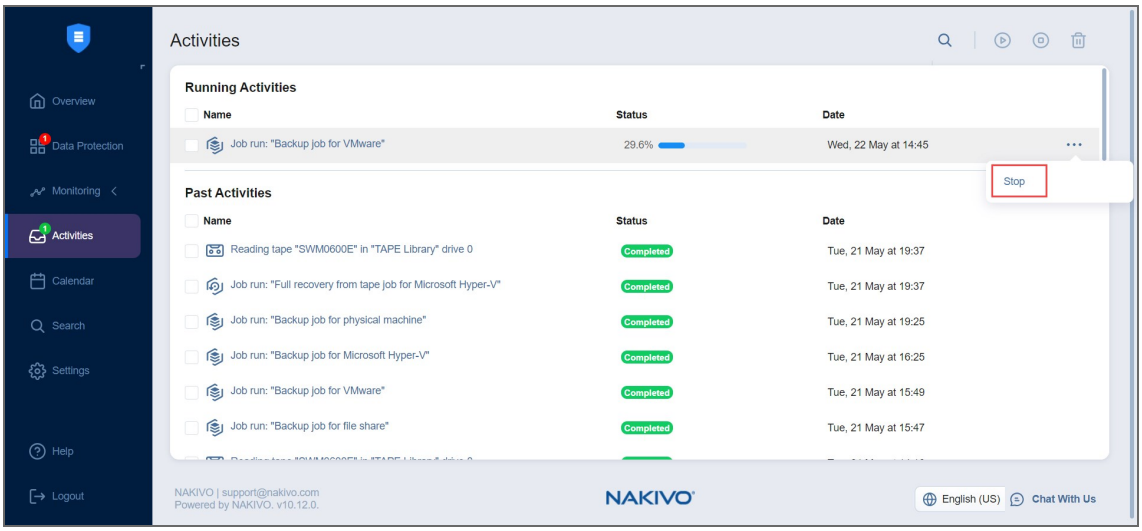
Viewing Activity Details

View the details of an activity by selecting an activity name. Note that you can view job run details by clicking the **details** link.



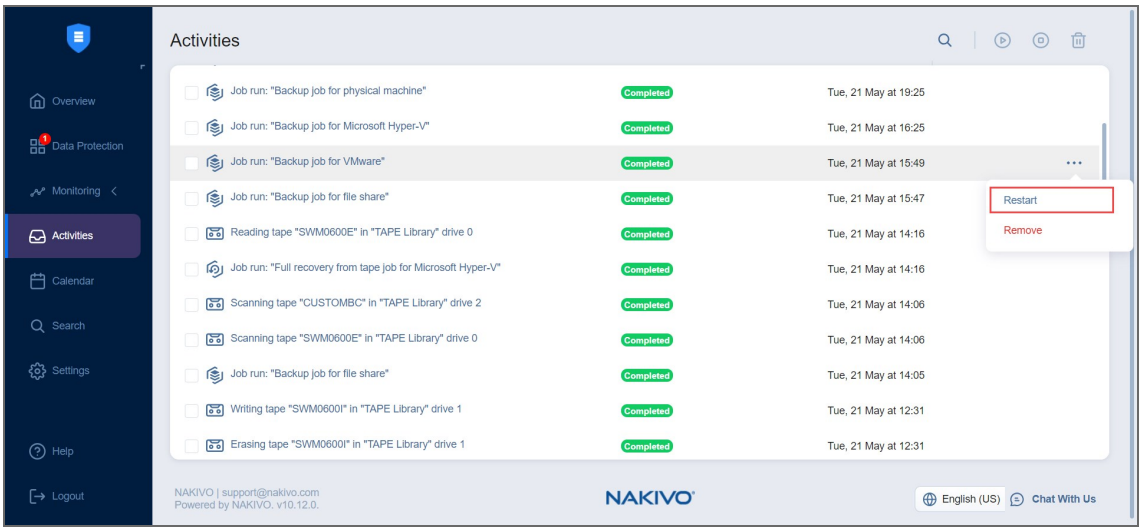
Stopping Running Activities

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



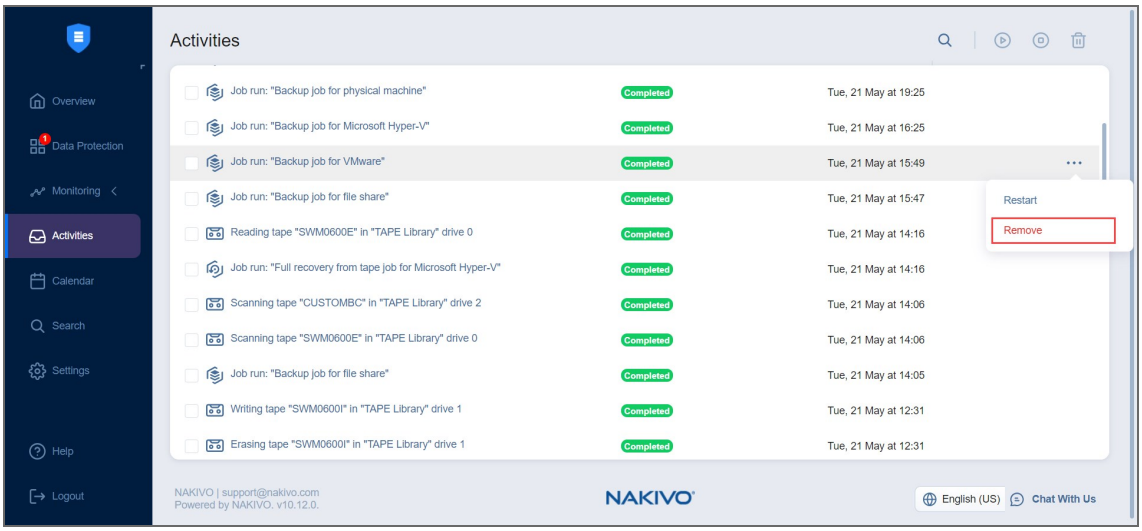
Running Activities Again

To run activities again (if possible), tick the checkbox next to each desired activity and click **Start** in the toolbar above. To run all activities again at once, tick the **Select/Deselect all** checkbox at the top and click **Start**. You can also run a single activity by clicking the **Start** icon that appears when you hover over a specific activity.



Removing Activities

To remove activities from the list, select the checkbox next to an activity, then click "**...**" and select **Remove**. To remove all activities from the list at once, tick the checkbox next to the **Name** and click **Remove** icon..


Activities
Overview
Data Protection
Monitoring
Activities
Calendar
Search
Settings
Help
Logout
Job run: "Backup job for physical machine" Completed Tue, 21 May at 19:25
Job run: "Backup job for Microsoft Hyper-V" Completed Tue, 21 May at 16:25
Job run: "Backup job for VMware" Completed Tue, 21 May at 15:49
Job run: "Backup job for file share" Completed Tue, 21 May at 15:47
Reading tape "SWM0600E" in "TAPE Library" drive 0 Completed Tue, 21 May at 14:16
Job run: "Full recovery from tape job for Microsoft Hyper-V" Completed Tue, 21 May at 14:16
Scanning tape "CUSTOMBC" in "TAPE Library" drive 2 Completed Tue, 21 May at 14:06
Scanning tape "SWM0600E" in "TAPE Library" drive 0 Completed Tue, 21 May at 14:06
Job run: "Backup job for file share" Completed Tue, 21 May at 14:05
Writing tape "SWM0600I" in "TAPE Library" drive 1 Completed Tue, 21 May at 12:31
Erasing tape "SWM0600I" in "TAPE Library" drive 1 Completed Tue, 21 May at 12:31
Restart
Remove
NAKIVO | support@nakivo.com
Powered by NAKIVO. v10.12.0.
NAKIVO
English (US)
Chat With Us

Using Calendar

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

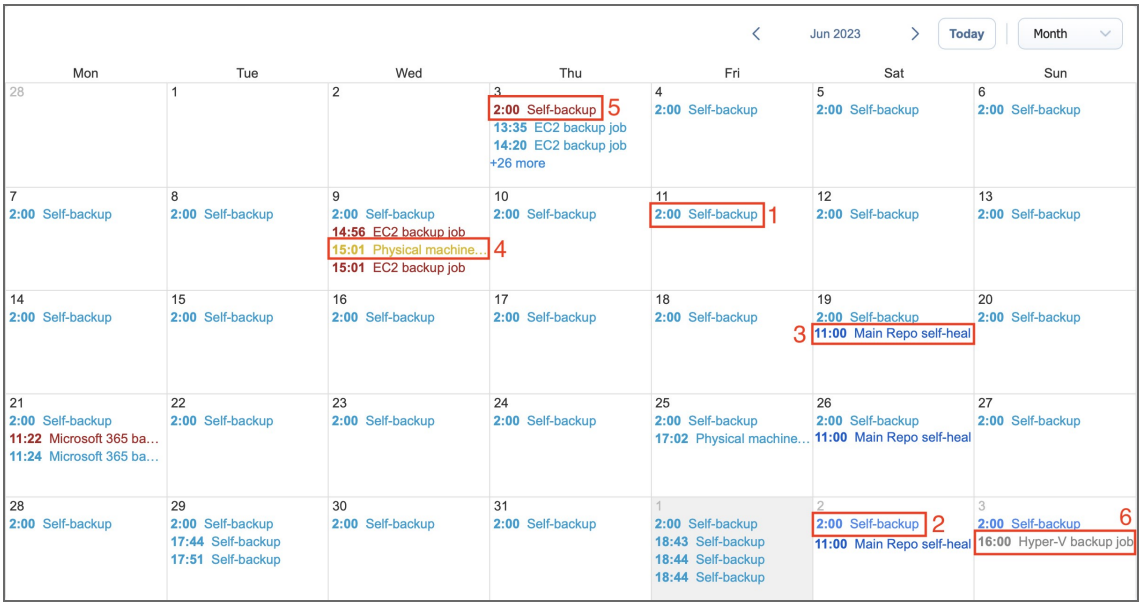
- Understanding Calendar Formatting
- Creating Jobs with Calendar
- Editing Jobs with Calendar

Understanding Calendar Formatting

Jobs in the **Calendar** view are formatted by start/end time and color coded by status. The color coding format is as follows:

1. Successful job runs are marked in teal.
2. Future scheduled job runs and currently running jobs are marked in sky blue.
3. Repository maintenance jobs (such as scheduled self-healing) are marked in navy blue.
4. Stopped job runs are marked in yellow.
5. Failed job runs are marked in red.
6. Job runs belonging to disabled jobs are marked in gray.

> **Note**
> Job runs that complete later than their start date are marked in the Calendar for the appropriate number of days. In **Month** view, such jobs are also marked with background fill. For example, a job that started on a Monday and finished on a Wednesday will be marked in one continuous solid light blue line across three days.



Creating Jobs with Calendar

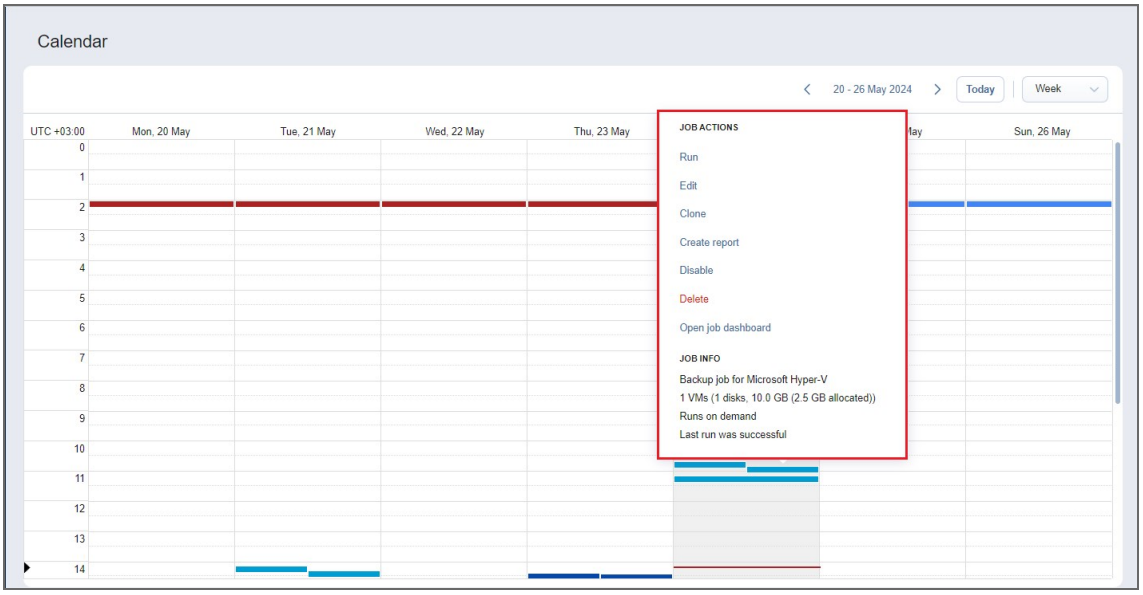
To create a job:

1. Click on the date and time when you'd like to run the job
2. Select the type of job you need.
3. On the **Schedule** page of the wizard, the time you've selected in the **Calendar** will be selected.

Editing Jobs with Calendar

If you click on the job title on the **Calendar** dashboard, the **Job Actions** and **Job info** menus appear. Using them, you can:

- Run a job on demand.
- Edit a job.
- Clone a job.
- Delete a job. If the job repeats on schedule, this action will affect all job runs.
- Disable/Enable a job. If the job repeats on schedule, this action will affect all job runs.
- Open the **Job Dashboard**.
- Create a report.
- View the job info.



Using Global Search

Using the **Global Search** dashboard, search for items within the entire inventory of NAKIVO Backup & Replication, Transporters, Backup Repositories, jobs, backups, and replicas.

- Opening Global Search
- Running Global Search

- Filtering Search Results
- Applying Bulk Action
- “Searching Indexed Files” on page 280
- Viewing Object Info
- Viewing Group Info

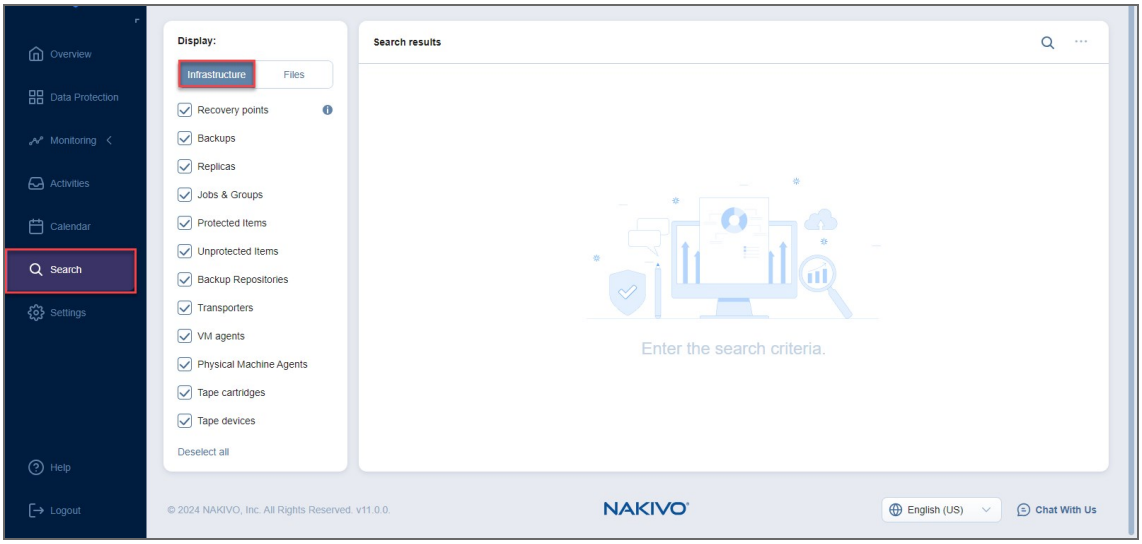
Note

When the multi-tenant mode is enabled, **Global Search** operates within a specific tenant. For more information about multi-tenancy in NAKIVO Backup & Replication, please consult with the following resources:

- “Multi-Tenancy” on page 28
- “Multi-Tenant Mode” on page 652

Opening Global Search

To open **Global Search**, click the **Search** icon in the main toolbar of the application. The **Global Search** dashboard opens with the **Infrastructure** tab selected by default.



Running Global Search

Click the **Search** button and enter your search string into the search box. The search results are immediately returned in the form of a list.

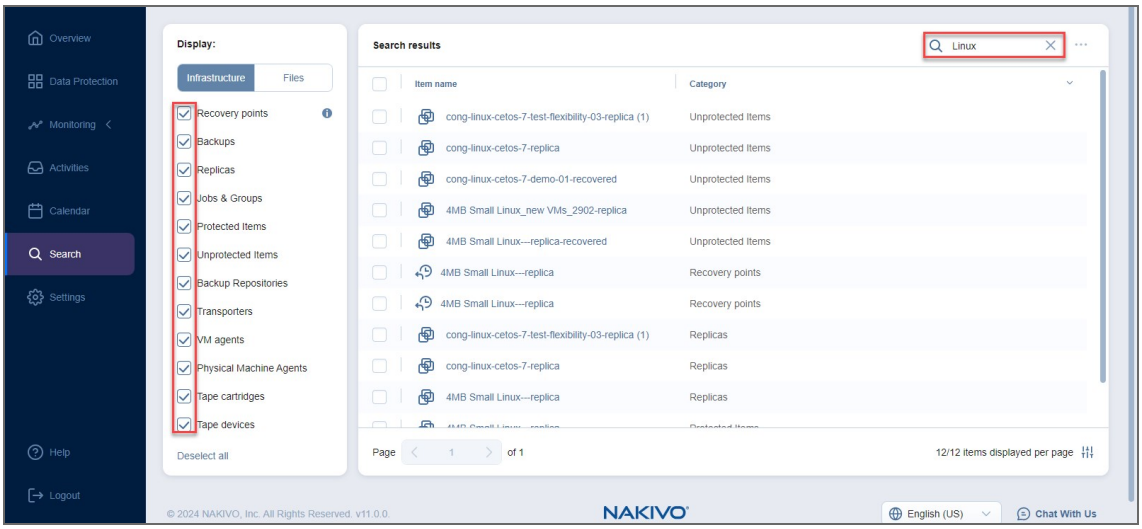
By default, your search results are non-filtered. This means that the search is applied to all categories of NAKIVO Backup & Replication objects that are enabled by default.

Note

To fine-tune search results, you can use the following wildcards:

- "**?**" representing a single character.

  "*****" representing zero or more characters.



Please note the following:

- Search is case insensitive.
- Search results are grouped by categories and the items within the same category are sorted alphabetically.
- Sorting by the Item name column is applied if multiple items of the same category are returned.

Configuring Table Columns

You can show/hide columns or modify the number of items per page in the Search Results table. In the lower right corner, click the controls icon. In the dialog window that opens, select/deselect checkboxes and click **Apply**. The following options are available:

- Item name
- Category
- Path
- Recovery Point Date

Note

You cannot modify the **Item name** and **Category** columns in this dialog window.

Click **Cancel** to exit or **Reset Settings** to reset to default settings.

Filtering Search Results

To narrow your search results, you can deselect some of the categories listed in the Infrastructure tab:

- Recovery points
- Backups
- Replicas
- Jobs & Groups
- Protected Items
- Unprotected Items
- Backup Repositories
- Transporters
- VM agents
- Physical Machine Agents
- Tape cartridges
- Tape devices

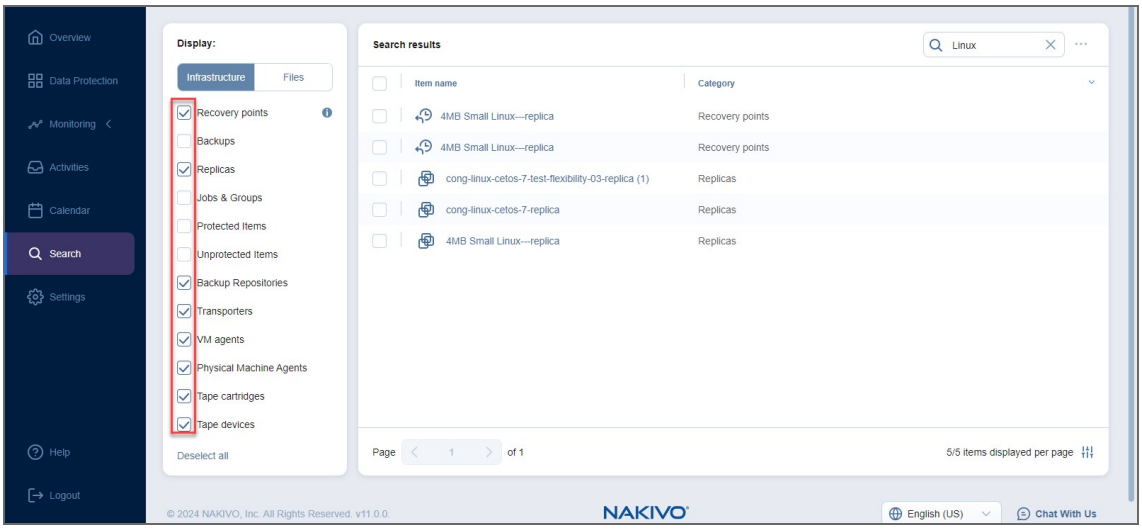
Notes

- When the search matches the recovery point name, all chains of recovery points of the source backup object to which this recovery point belongs are returned.

  Exception

  Self-backup recovery points are not returned in the search result.
- If you have added a federated repository to the product, **Global Search** displays the federated repository name in the search result instead of its members.
- The following fields are not displayed in the federated repository popup:
  - Backup repository path
  - Backup repository compression
  - Transporter
  - Deduplication
  - Space savings

The filtered search results are displayed immediately as the search results list.



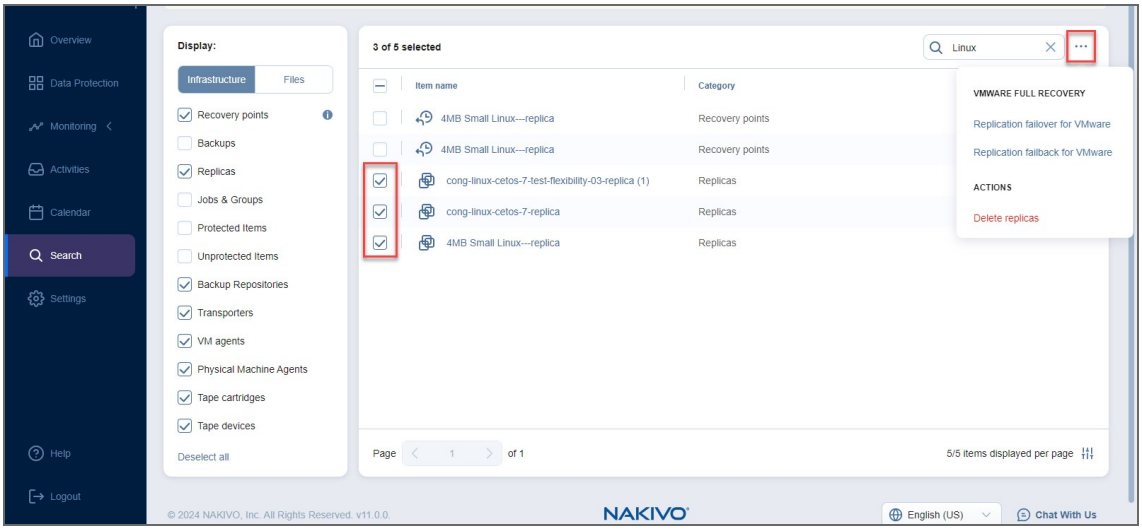
To get back to the default filtering settings, first click **Deselect all** and then **Select all** below the categories list.

Applying Bulk Action

With NAKIVO Backup & Replication Global Search, you can apply a bulk action to objects belonging to the same category and of the same type.

Proceed as follows to apply a bulk action:

1. In the search result list, select similar objects.
2. The **Bulk Action** button becomes active in the upper right corner of the dialog. Click **Bulk Action**.



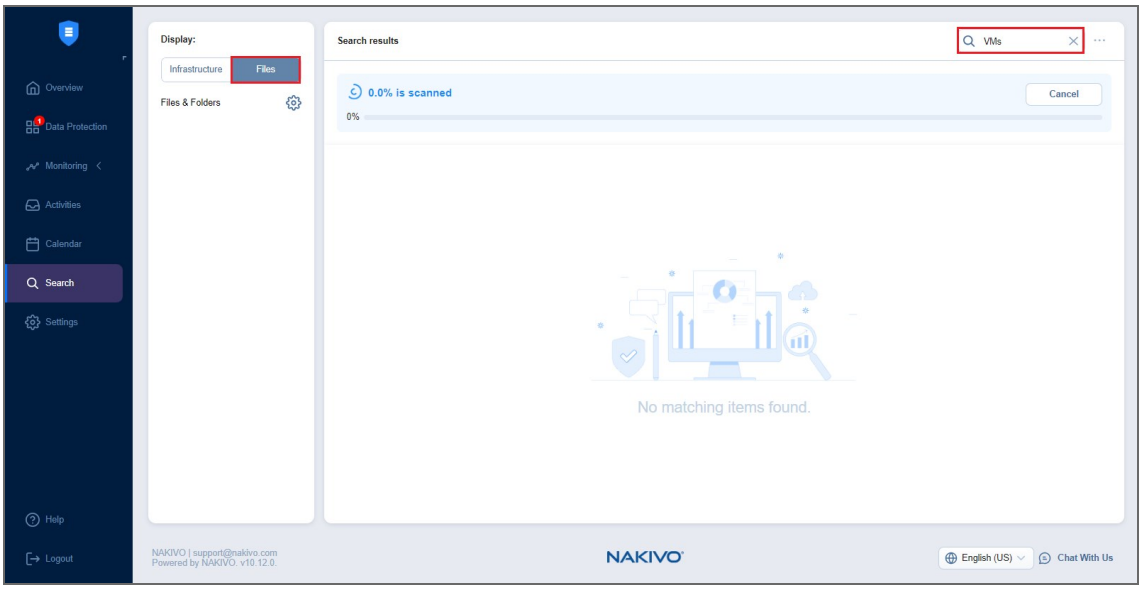
3. A dialog opens with the list of actions applicable to the selected items. To proceed with the necessary action, click the corresponding item in the list of actions.

Notes

- Bulk actions do not apply to NAKIVO Backup & Replication dissimilar objects.
- Bulk actions do not apply to items with mixed destinations (disk and tape).

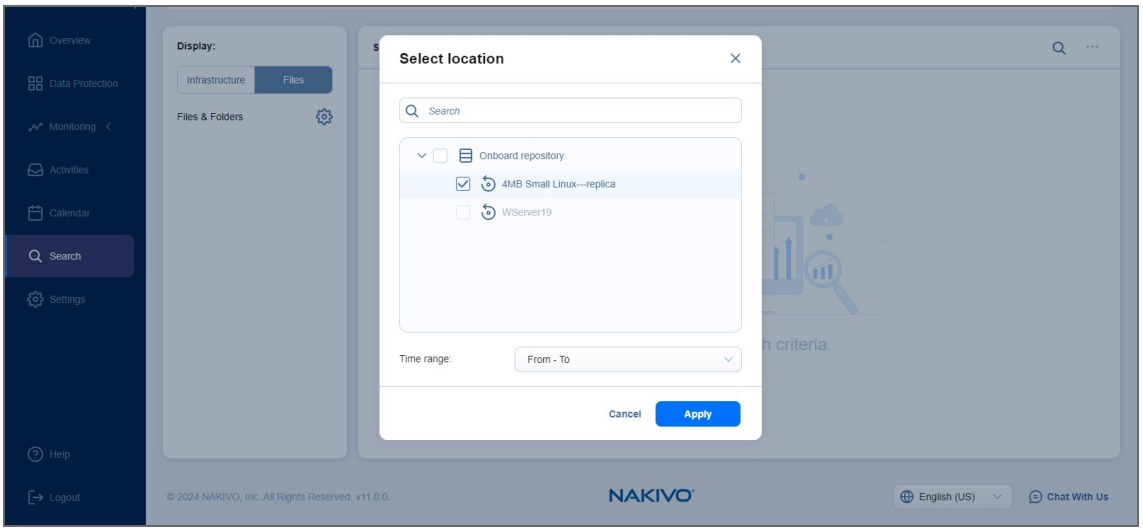
Searching Indexed Files

To search for a file or folder, switch to the **Files** tab. Click the **Search** button and enter your search string into the search box.

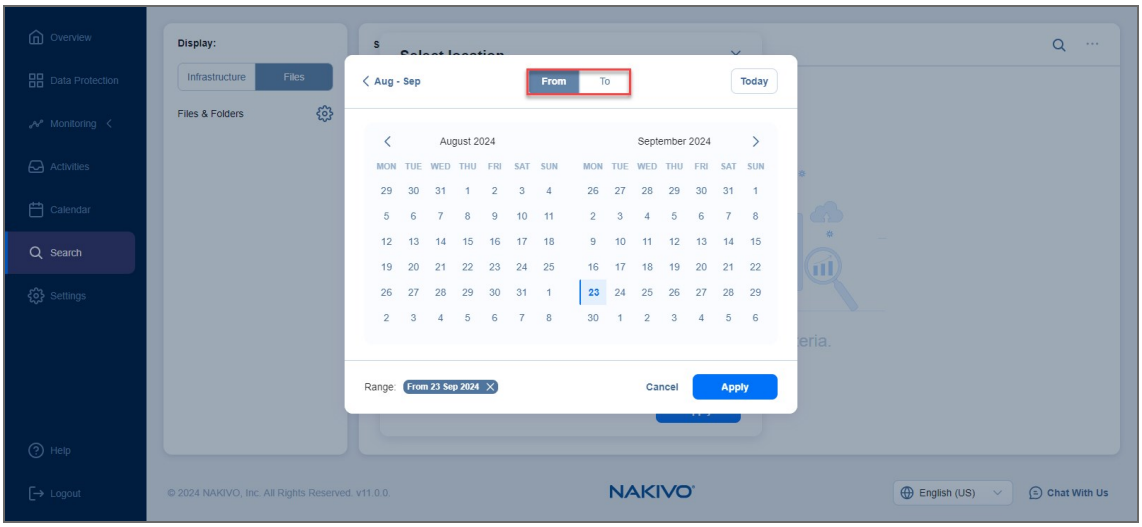


The **Search results** pane displays a progress bar (displayed when scanning is enabled only) and the results of the search.

You can also narrow your search results in the **Files** tab. To do that, click the gear icon next to **Files & Folders** to open the **Select location** popup.



Open the date picker dropdown and set the needed search range by navigating between the **From** and **To** panes.



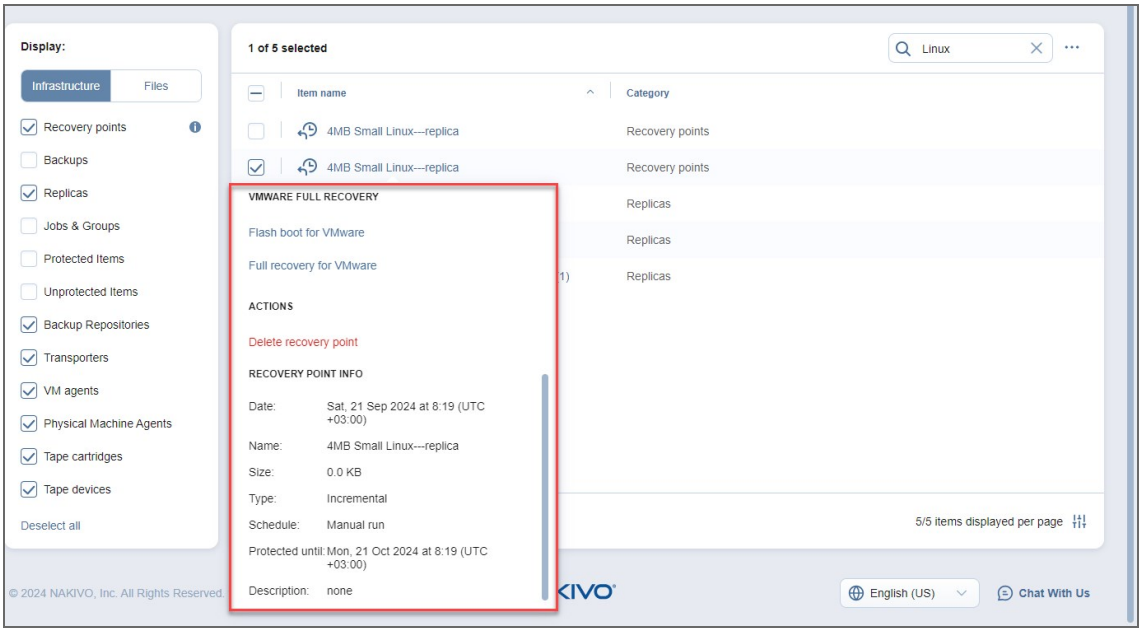
Click **Apply** to filter the search results.

The index of the following items included in the backup job is compressed and recorded into the **Director** file storage to be returned as **Global Search** results:

- File name
- File extension
- File path
- Linked recovery point
- Disk label on OS level
    - Windows OSs: disk name is a drive letter of volume disk (e.g. C, D, E, F, etc.)
    - Linux OSs: disk name is a block device name of disk (e.g: /dev/sda1, /dev/sdb2, etc.).

Viewing Object Info

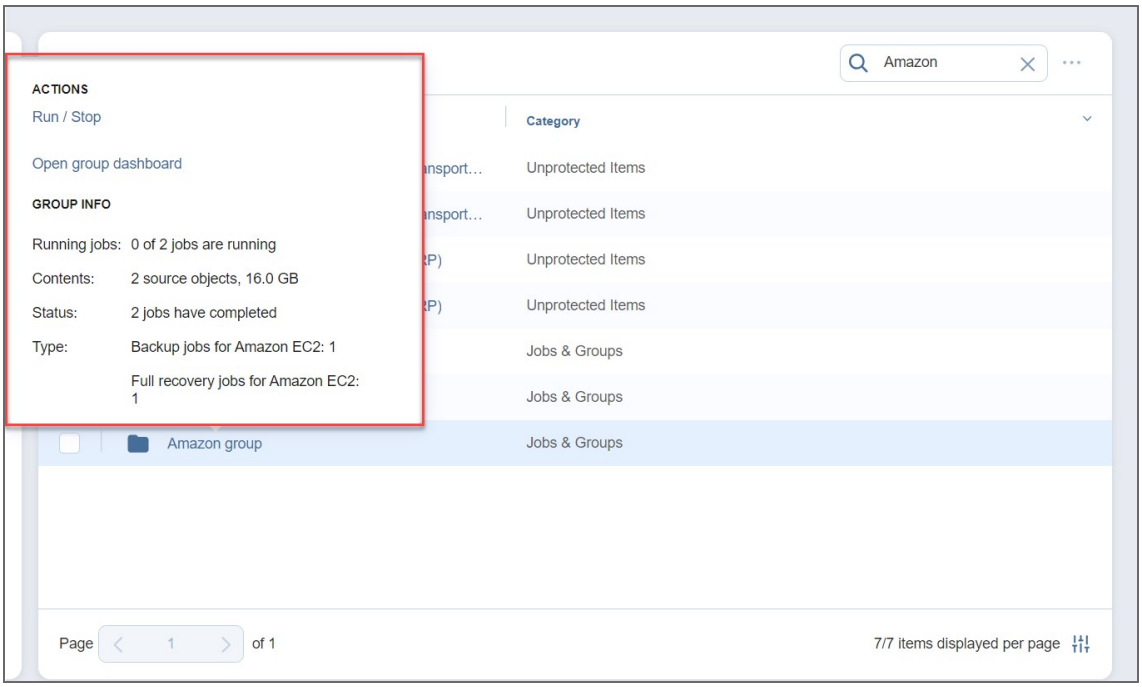
To view info on a specific object available in the search result, select the object.



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

Viewing Group Info

To view information about a specific group available in the search result, select the group.



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

Settings

This section covers the following topics:

- “General” on page 284
- “Inventory” on page 403
- “Nodes” on page 430
- “Backup Repositories” on page 460
- “Federated Repositories” on page 487
- “Expert Mode” on page 512
- “Maintenance Mode” on page 538
- “Multi-Tenant Mode Configuration” on page 564
- “Support Bundles” on page 579
- “Built-in Support Chat” on page 581

General

This section contains the following topics:

- "Bandwidth Throttling" on page 285
- "Branding" on page 288
- "Configuring Events" on page 290
- "Database Options" on page 291
- "Email Settings" on page 296
- "Licensing" on page 300
- "Managing Backup Encryption" on page 303
- "MSP Tab in Single-Tenant Mode" on page 314
- "Notifications & Reports" on page 321
- "Self-Backup" on page 324
- "Software Update" on page 329
- "System Settings" on page 332
- "Users and Roles" on page 346

Bandwidth Throttling

With bandwidth throttling settings, you can control the throughput of the data processing by setting specific limits for all or for separate jobs. Bandwidth throttling is managed with bandwidth rules. When a bandwidth rule is applied to your job, the speed of data transfer from source to target will not exceed the specified limit.

To use bandwidth throttling, make sure the following requirements are met:

- At least one bandwidth rule is created.
- At least one source and target endpoint exist.
  - Hypervisor host, **Transporter**, or **Backup Repository** act as a source or target endpoint.
  - The current job object data transfer falls into the bandwidth rule.
- At least one **JODT** exists.

> **Note**
>
> **Source Endpoint** - an endpoint that sends data.
>
> **Target Endpoint** - an endpoint that receives data.
>
> **Job Object Data Transfer** (**JODT**) is a step of a single job object processing which transfers data of the job object from the source endpoint to the target endpoint.

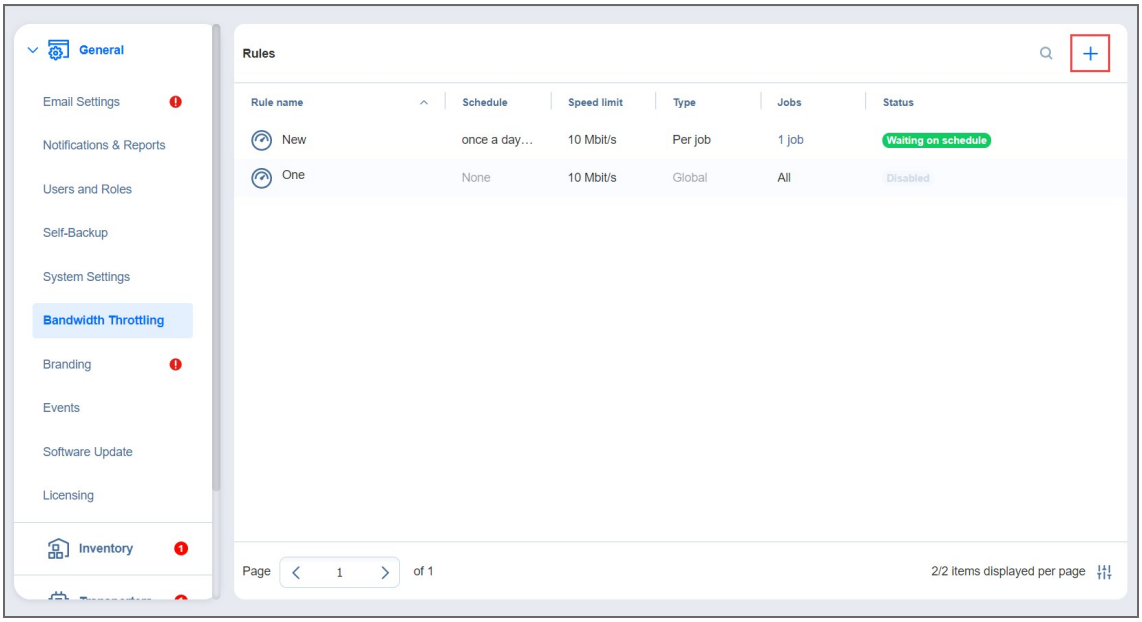
This topic contains the following instructions:

- Accessing Bandwidth Throttling Settings
- Creating Bandwidth Rules
- Managing Bandwidth Rules

Accessing Bandwidth Throttling Settings

To access bandwidth throttling settings, follow the steps below:

1. Click **Settings** in the left pane of the application to open the **Settings** dashboard.
2. In the **General** tab of the **Settings** dashboard, click **Bandwidth Throttling**. The *Bandwidth Throttling* section opens.



Creating Bandwidth Rules

Please follow the steps below to create a bandwidth rule:

1. In the *Bandwidth Throttling* section of the **General** tab of **Settings**, click the "**+**" icon.
2. The **New Bandwidth Rule** wizard opens. Proceed as follows:
    a. Choose a type for your bandwidth rule and click **Next**:
        - **GLOBAL**: The rule will be applied to all applicable jobs.
        - **PER-JOB**: The rule will be applied to the selected jobs.

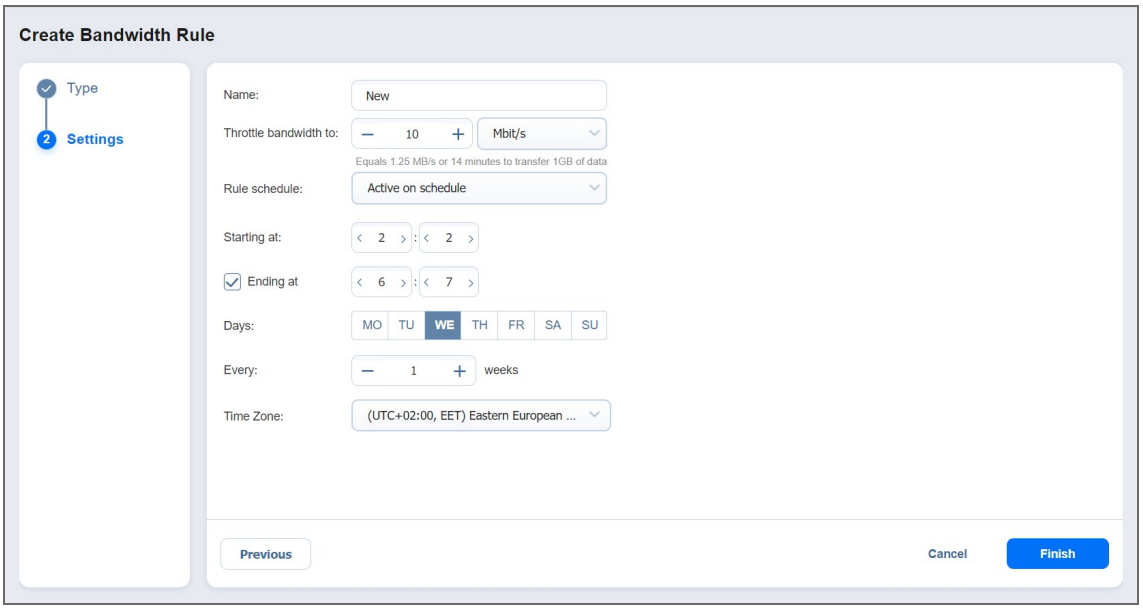
> **Note**
> When applied to specific jobs, **Per-Job** bandwidth rules have higher priority over **Global** bandwidth rules.

    b. **Job**: Choose a job to apply to the **Per-Job** bandwidth rule. Click **Next**.
    c. **Settings**: Configure the following settings:
        a. **Name**: Enter a name for your bandwidth rule.
        b. **Throttle bandwidth to**: Enter the value of the bandwidth limit; and choose the measurement unit: Mbit/s or Gbit/s.

> **Notes**
> - For your convenience, a description is available below the value you've entered, explaining what the value means.
> - In some cases, the actual data transfer speed may exceed the limit you set by up to 0.3 MByte/s or 2.4 Mbit/s.

c. **Rule schedule**: Choose either of the following:
   - **Always active**: The rule will always be active.
   - **Active on schedule**: The rule will be active on schedule. When chosen, the following options are available:
     a. **Starting at** and **ending at**: Enter the time, in hours and minutes, when the rule will be active.
     b. **Days**: Select weekdays for which the rule will be active.
     c. **Time Zone**: Choose a time zone of your rule.
   - **Disabled**: The rule will be disabled.

3. Click **Finish**.



Refer to "Advanced Bandwidth Throttling" on page 19 for more details about bandwidth rules.

> **Note**
> When set, the bandwidth rule will only limit the bandwidth itself and will not affect the number of VMs processed at the same time.

Managing Bandwidth Rules

You can search for the specific rule by clicking the **magnifying glass** icon in the upper-right part of the screen and entering the name in the search box.

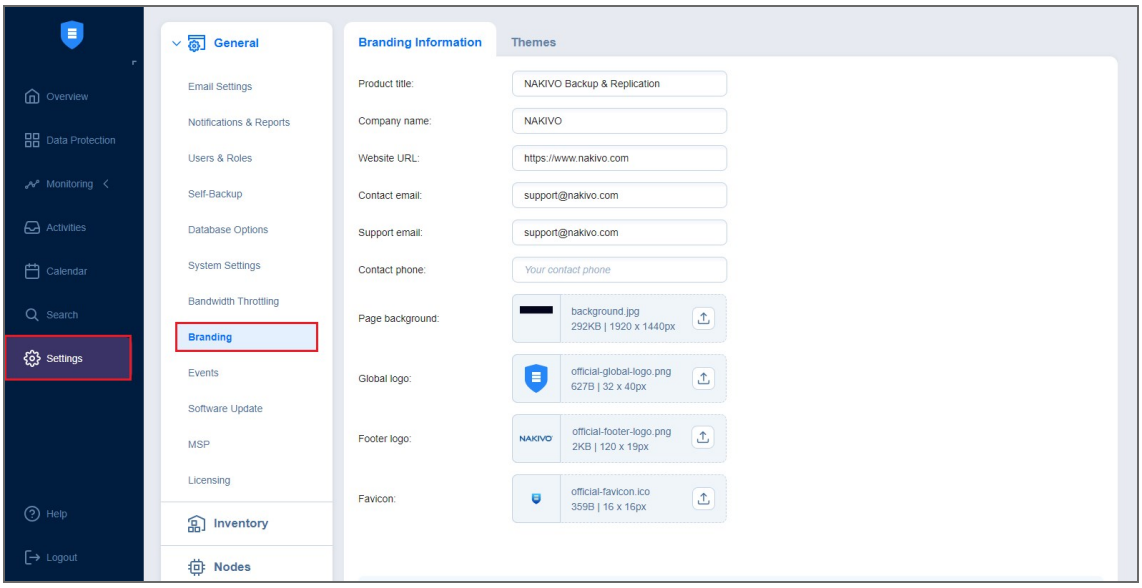
Click on the **ellipsis** to the right of the rule's name to manage bandwidth rules with the following commands:

- **Edit**: The **Edit Bandwidth Rule** dialog opens where you can modify your rule.
- **Disable/Enable**: When applied, the command will disable/enable the rule.
- **Delete**: When applied, a dialog will open asking you to confirm the operation. Click **Delete** to confirm that you wish to delete your rule.

Branding

You can change the product branding settings such as product name, logo, background, and so on. To configure these product settings, follow the steps below:

1. Click **Settings** in the left pane of the product.
2. Go to the **General** tab and click **Branding**.



3. Change the following, as appropriate:
   - **Product title**
   - **Company name**
   - **Website URL**
   - **Contact email**
   - **Support email**
   - **Contact phone**
   - **Page background**
   - **Global logo**
   - **Footer logo**
   - **Favicon**
4. On the **Themes** tab, you can configure the colors of your NAKIVO Backup & Replication instance.
5. After making the necessary changes, click **Apply**. Alternatively, click **Discard Changes** to discard any changes you have made.
6. Optionally, click **Reset Settings** to return all the settings to their default values.

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

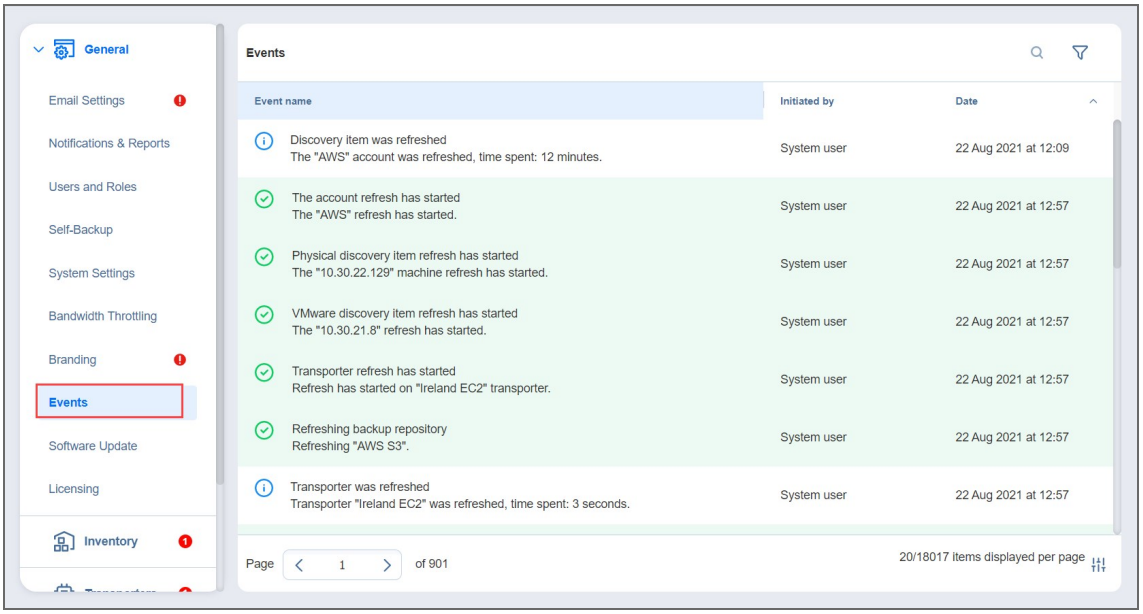
Image	Best format	Best resolution
Global logo	.png	32x40
Footer logo	.png	32x40
Favicon	.png	16x16

Configuring Events

NAKIVO Backup & Replication can store and display system events. By default, events are stored for 60 days; you can change the time period in **Settings**.

To view events, follow the steps below:

1. Click **Settings** in the left pane of the product.
2. Open the **General** tab and click **Events**. The **Events** page opens, displaying the NAKIVO Backup & Replication system events.



3. Optionally, you can enter a search string to the **Search** box. This allows you to see events related only to NAKIVO Backup & Replication items – **Transporters**, repositories, jobs, backups, and replicas,– contained in your search string.
4. Optionally, you can select filter the events by the following parameters:
    - **Initiated by**: Select one of the users of the product in the dropbox
    - **Event type**: Choose among the following event types:
        - Info
        - Warning
        - Error
        - Debug
    - **Date**: After selecting this parameter, choose the start and end dates. This allows you to limit the events list within a specific time period.

Database Options

- System Requirements for PostgreSQL
- Installing PostgreSQL
- Migrating Database

NAKIVO Backup & Replication allows you to migrate the internal H2 database used by the NAKIVO Backup & Replication Director to an external database. To do that, take the following steps:

Important

- If you migrate the internal H2 database to an external database, you will not be able to switch back to the internal database or an external database of the same type later.
- For multi-tenant mode, only the Master Admin can perform database migration. The functionality is not available for the local tenant. The remote tenant can still perform database migration as described below if they log in as a single tenant into NAKIVO Backup & Replication.
- The migration occurs for all local tenants at the same time. If the migration fails for one of the tenants, the product reverts to the previous database type automatically.
- Insufficient resources may cause PostgreSQL to crash during operations.
- PostgreSQL can be installed either on the same server as the Director or on a separate server. If you choose to install it on the Director's server, ensure sufficient resources are available to run both services.

System Requirements for PostgreSQL

1. Make sure that the PostgreSQL server has at least:
   - 4 CPUs and 8GB RAM for Single-Tenant mode
   - 8 CPUs and 36GB RAM for Multi-Tenant mode
2. Set the following parameters in the **postgresql.conf** file:

- for Single-Tenant mode:

```
max_connections = 500
shared_buffers = 256MB
work_mem = 10MB
```

- for Multi-Tenant mode for around 100 tenants including 50 tenants with monitoring data:

```
max_connections = 3000
shared_buffers = 256MB
work_mem = 10MB
```

- for Multi-Tenant mode for around 100 tenants with monitoring data:

```
max_connections = 5000
shared_buffers = 256MB
work_mem = 6MB
```

See also "Feature Requirements" on page 91.

Installing PostgreSQL

Installing PostgreSQL on Windows

1. Download the PostgreSQL installer from the official website.
2. Run the installer and follow the instructions.
3. Edit the **postgresql.conf** to set the required system resources configuration.
4. Restart the PostgreSQL service to apply changes.

Installing PostgreSQL on Linux

PostgreSQL is available integrated with the package management on most Linux platforms.

1. Install the PostgrSQL

   ```
   apt install postgresql
   ```

2. Configure PostgreSQL by editing the **postgresql.conf** to set the required system resources configuration.
3. Restart the PostgreSQL service to apply changes.

   ```
   systemctl restart postgresql
   ```

Note

When the PostgreSQL is installed, ensure that the database user has the following privileges:

- Create Database
- Create user with password
- Grant privileges on database to the created user
- Alter the created user

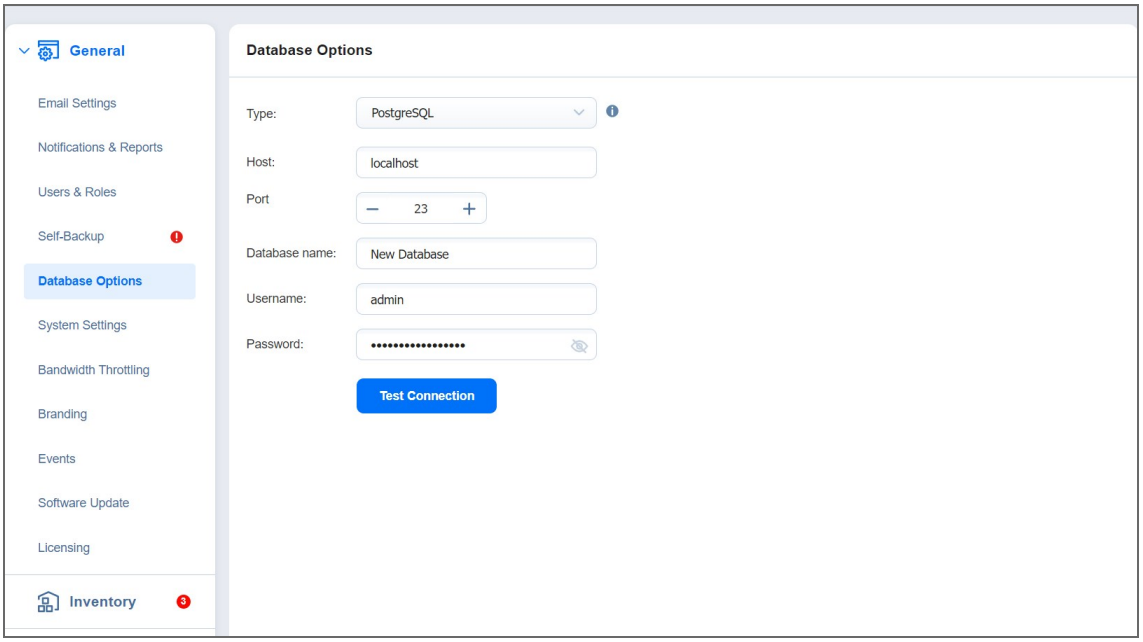
Migrating Database

1. Go to **Settings > General > Database Options**.
2. Select the external database from the list of supported platforms in the **Type** drop-down list. Note that the internal database is selected by default.
3. In the **Host** field, enter the hostname or IP of the server housing the database.

4. In the **Port** field, enter the relevant port number.
5. Enter the name for your database in the **Database name** field.
6. Enter **Username** and **Password** in the corresponding fields.
7. Click **Test Connection**.
8. If the test is successful, click **Apply Settings**:
    - If the database does not exist, a dialog box appears asking if you would like to create one and proceed with the migration. Click **Migrate**.
    - If the database belongs to the current NAKIVO Backup & Replication installation, a dialog box appears asking if you would like to update the settings of the existing database. Click **Update** to proceed.
    - If the database already exists and is compatible with the current NAKIVO Backup & Replication installation, a dialog box appears asking if you would like to use it, cleanup all its records and replace the contents of the database with the new data. Click **Proceed**.

Important

Check ALL your repositories after migration and make sure they are still attached.



Notes

- If you have the Self-Backup feature enabled, the self-backup process starts before the database switch and runs again after the switch is completed.
- Self-backup of an external database is possible only with a single-tenant instance of the solution.
- If the external database is installed on another VM or is using an IP address instead of *localhost*, take the following steps before migration:
  1. Open the *pg_hba.conf* file located in the external database installation folder.
  2. Change **IPv4 local connections** settings *from 127.0.0.1/32* to *0.0.0.0/0*.
  3. Save changes.
  4. Restart external database service.
- If the connection between PostgreSQL and NAKIVO Backup & Replication cannot be established, add the following string to the *pg_hba.conf* file:

  host DATABASE USER ADDRESS METHOD [OPTIONS]

  host all all 0.0.0.0/0 md5

  or

  host DATABASE USER ADDRESS METHOD [OPTIONS]

  host all all 0.0.0.0/0 scram-sha-256

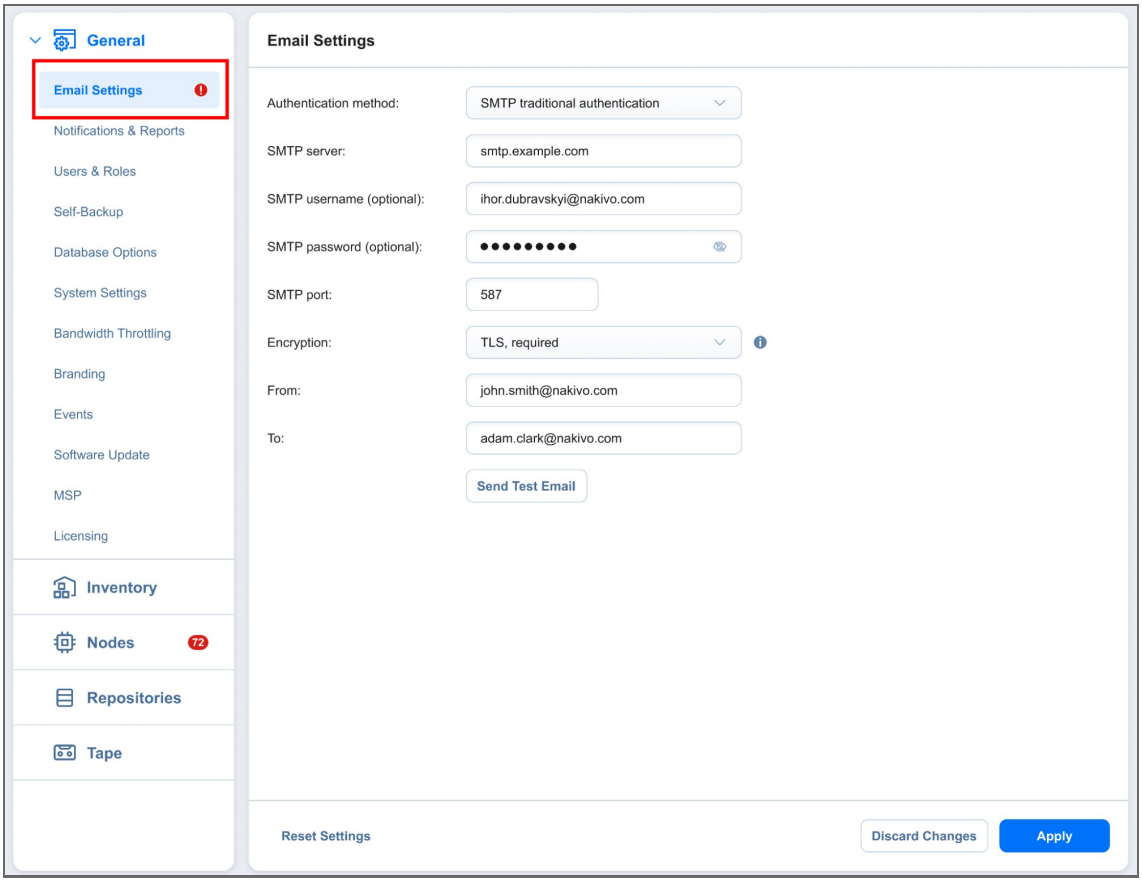
  Note that method (md5/scram-sha-256) may be different for some versions of PostgreSQL. Check the respective method for your version of PostgreSQL before applying the changes.
- If Master Tenant connects to existing database that already houses the data from previous migrations, such database is automatically mapped to the tenants during the new migration using the database UUID.
- It is not possible to recover from a self-backup and system migration in the following cases:
  - The NAKIVO Backup & Replication installation uses the H2 database while the self-backup contains data from an external database.
  - The NAKIVO Backup & Replication installation uses an external database while the self-backup contains data from the H2 database.
- It is not possible to edit the external database **Host**, **Port**, and **Database name** after a successful migration.
- If the internal database is used, the product checks the performance capability adequacy of this database to the current product workload:

- This check is performed every 10 days by default.
- If the total number of protected workloads for single tenant or per tenant for Multi-Tenant mode exceeds the limit of 100 VMs/instances/physical machines/oracle databases, the product displays the notification with recommendation to switch to the external database.

Email Settings

On this page, you can configure your email settings. Do this by following the steps below:

- Log in to NAKIVO Backup & Replication.
- Click **Settings** in the left pane of the product.
- Go to the **General** tab.
- Click **Email Settings** to configure email settings on the page that opens.

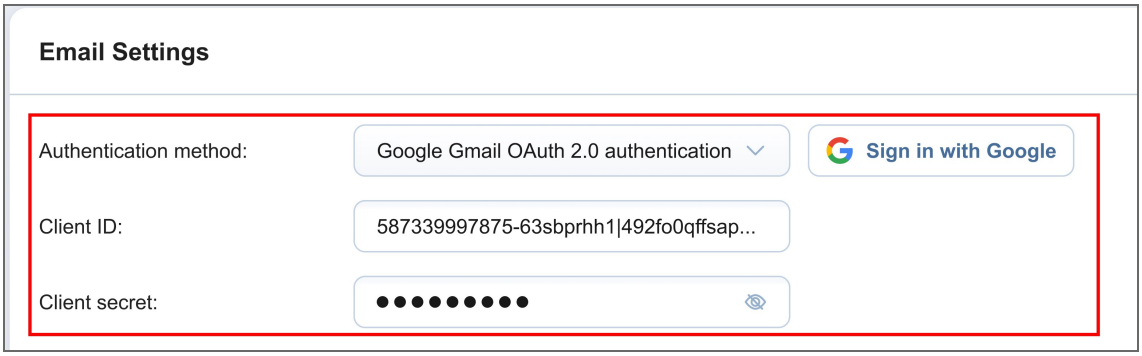


1. To set email settings, fill out the fields in the Email settings section:
    - **Authentication method**: Select the authentication method from the dropdown.
        - **SMTP traditional authentication**: Select the traditional email authentication to connect your mail server, then configure the following parameters. See also Using Email with Two-Factor Authentication.

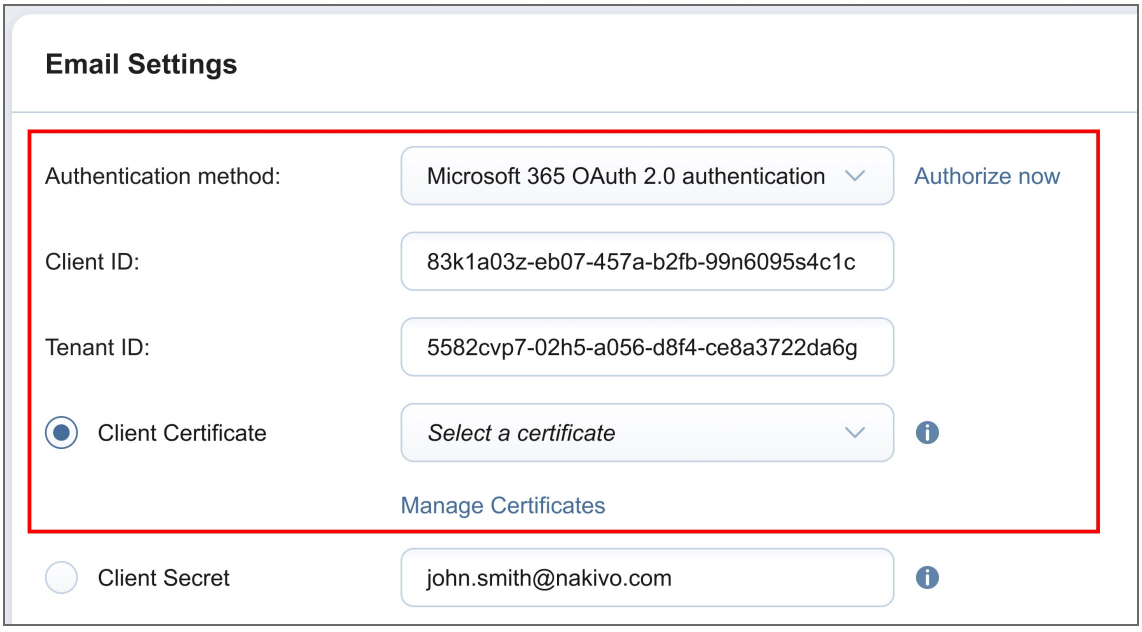
> **Important**
> If you use an email with two-factor authentication, grant access permissions to NAKIVO Backup & Replication via your account security settings and generate a unique password. When configuring email setting of the product, enter this password in the **SMTP password** box.

- **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.

- **Google Gmail OAuth 2.0 authentication**: Select the Google Gmail OAuth to connect your mail server, then configure the following parameters. For more information, refer to "Using Email with Gmail OAuth Authentication" on page 388.



  - **Client ID**: Enter the unique identifier generated in your Google Cloud Console when you create OAuth 2.0 credentials.
  - **Client Secret**: Enter the confidential key paired with your **Client ID**. Click the **Sign in with Google** to authorize NAKIVO Backup & Replication to access your Gmail account through OAuth 2.0 authentication. Ensure that the credentials are correct and the account is authenticated successfully.
- **Microsoft 365 OAuth authentication**: Select the **Microsoft 365 OAuth** to connect your mail server and configure the following parameters:



- **Client ID**: Enter the unique identifier generated with your Microsoft 365 account when you create OAuth 2.0 credentials.
- **Tenant ID**: Enter your Tenant ID (your Tenant ID is displayed in the **Tenant information** section in **Microsoft Entra Admin center**).
- **Client Certificate**: Choose a saved certificate from the drop-down list to use an existing certificate. If you don't have any saved certificates or want to generate a new one, click **Manage certificates**. In the resulting pop-up, you can edit and delete existing certificates, generate new certificates, or import existing certificates.
- **Client Secret**: Select and enter the Client Secret, generated on the Microsoft Azure portal.

  Click **Authorize now** to authorize NAKIVO Backup & Replication to access your Microsoft 365 account through OAuth 2.0 authentication. Ensure that the credentials are correct and the account is authenticated successfully. For more information, refer to "Using Email with Microsoft OAuth Authentication" on page 393.

Notes

- When Google Gmail OAuth 2.0 authentication or Microsoft 365 OAuth 2.0 authentication is selected as the Authentication method, the Encription field is hiden.
- When Google Gmail OAuth 2.0 authentication or Microsoft 365 OAuth 2.0 authentication is selected as the Authentication method, the From field is disabled and and automatically populated with the email address used during authentication.

- **Encryption**: Select the type of encryption:
  - **None**: Always use a plaintext connection. Not recommended.
  - **TLS, if possible**: Start with plaintext, then use `STARTTLS` to switch to secure connection if supported by the server.
  - **TLS, required**: Start with plaintext, then use `STARTTLS` to switch to secure connection; drop the connection if not supported by the server.
  - **SSL, required**: Use the SSL-encrypted connection.
  - **From**: Specify the sender email address
  - **To**: Specify the receiver email address

2. Click **Send Test Email** to verify that the settings are correct.
3. Click **Apply** to save the settings.
4. Alternatively, click **Discard Changes** to discard any changes you have made to the email configuration.

5. Optionally, click **Reset Settings** to return all the settings to their default values.

Licensing

The Licensing page displays information about your current license and allows you to manage licensing options for NAKIVO Backup & Replication.

Checking License Details

To check your license details, follow these steps:

1. Go to the main menu of NAKIVO Backup & Replication and click **Settings** > **General**.
2. Go to the **Licensing** tab to see license details.

In the **License Information** section, you can find detailed license information, including:

- **Type**: Type of the license
- **Edition**: Edition of the license
- **Serial number**: Serial number of the license
- **License expiration date**: Date when the license expires



In the **Perpetual licensing section**, you can see the following information:

- Number of licensed and used CPU sockets
- Number of licensed and used VMs
- Number of licensed and used physical servers
- Number of licensed and used physical workstations
- Number of licensed and used EC2 instances
- Number of licensed and used Oracle databases

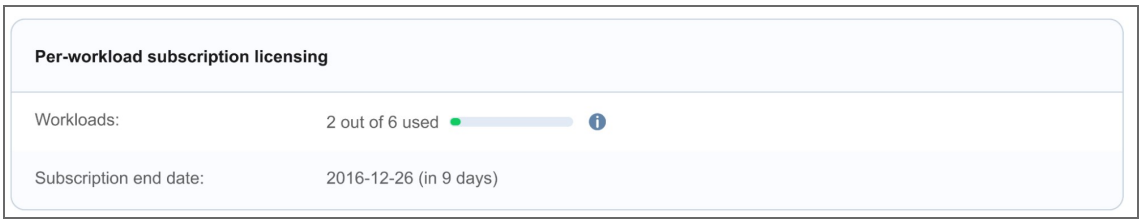


Note

This section is displayed when a perpetual license is applied to the product.

In the **Per-workload subscription licensing** section, you can see the following information:

- Number of licensed and used workloads
- Subscription end date



Note

This section is displayed when a subscription license is applied to the product.

In the **Microsoft 365 subscription licensing** section, you can see the following information:

- Number of licensed and used Microsoft 365 users
- Subscription end date

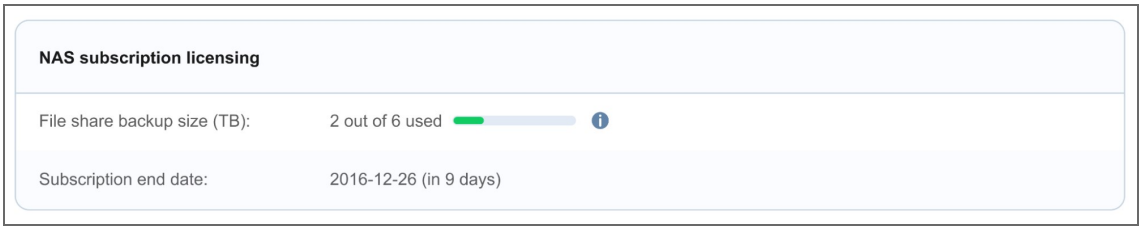


Note

In Multi-Tenant mode, Microsoft 365 users displays the total number of Microsoft 365 users used across all tenants for the Master tenant, and the total number of Microsoft 365 users used by a specific tenant.

In the **NAS subscription licensing** section, you can see the following information:

- File share backup size in TB
- Subscription end date



Note

In Multi-Tenant mode, the File share backup size in TB displays:

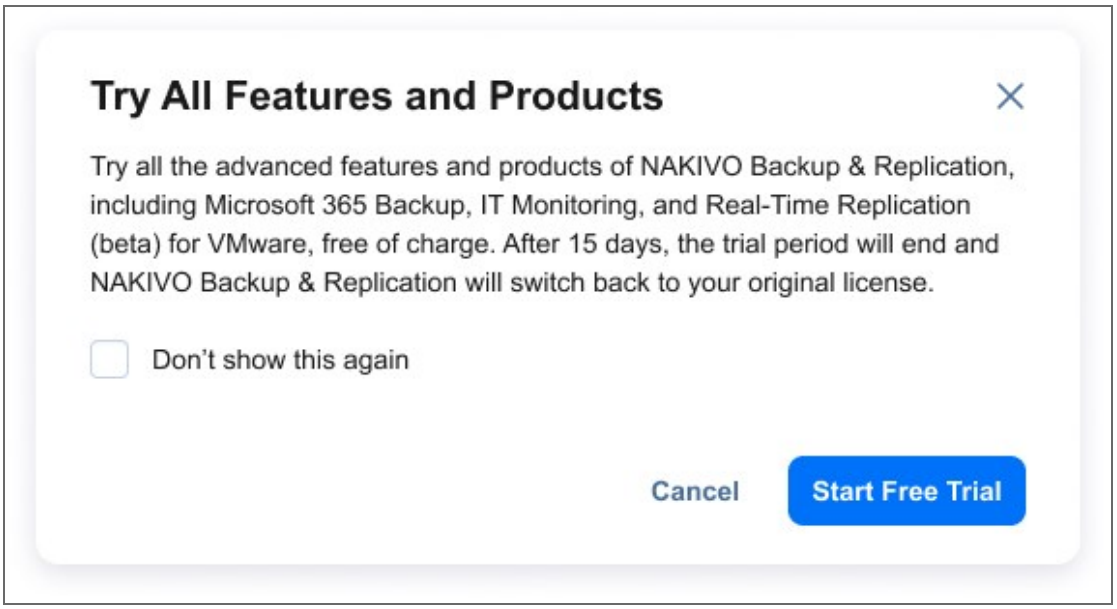
- The total size of licensed file share objects (per 0.5 TB) used across all tenants for the Master tenant.
- The size of licensed file share source objects (per 0.5 TB) allocated to the Local tenant.

If you are logged in as a tenant in multi-tenant mode, the following information is displayed in the **Obtain more licenses** section:

- **Email address** of the master tenant
- **Contact phone** of the master tenant
- **Company website** of the master tenant

Enabling Trial Period

Optionally, you can click the **Try All Features and Products** button to enable all Enterprise Plus license features for 15 days. The trial period can be activated only once per product version. After this time period ends, NAKIVO Backup & Replication automatically switches back to your original license.



Notes

The button is not displayed in the following cases:

- You are using NAKIVO Backup & Replication as a tenant in Multi-tenant mode.
- You are using one of the following license editions:
  - Free
  - Trial
  - Beta
  - Promo
  - Enterprise Plus
  - MSP Enterprise Plus

Changing the License

To change your license, follow the steps below:

1. Go to the main menu of NAKIVO Backup & Replication and click **Settings**.
2. Go to the **Licensing** tab and click **Change License**.
3. In the window that appears, locate and select the new license file.
4. Click **Apply** to apply the new license.

Upgrading from Free License

If your license type is **Free** and the **Trial** license has not yet been applied to you deployment of NAKIVO Backup & Replication, you can try the full functionality of the solution for 15 days. To do that:

1. Open the **Help Menu**.
2. Select the **Try full functionality** option. A new popup window appears.
3. Click **Start Free Trial.**

Note

Once the **Trial** license expires, the product automatically switches back to the **Free** license.

Managing Backup Encryption

With NAKIVO Backup & Replication, you can personalize your protection settings through encryption of backup data, password management, and restoration from encrypted backups.

By using the **Backup Encryption** feature, you can encrypt the following:

1. Product workloads' backups
2. System configuration stored as a self-backup and used in the scope of system migration
3. File System Indexing (FSI) data stored on the Director

Note

Encryption does not alter the quality or effectiveness of the compression process, as compression reduces the data size before encryption is applied.

The product offers an efficient password management dialog box, so that you do not need to enter your password every time you restore encrypted data.

AWS KMS is highly recommended as an additional layer of protection and to restore the password hash in case you forget it.

Refer to the following sections to learn more about managing backup encryption:

- “Enabling Backup Encryption” below
- “Managing Passwords” on page 307
- “Enabling KMS” on page 311

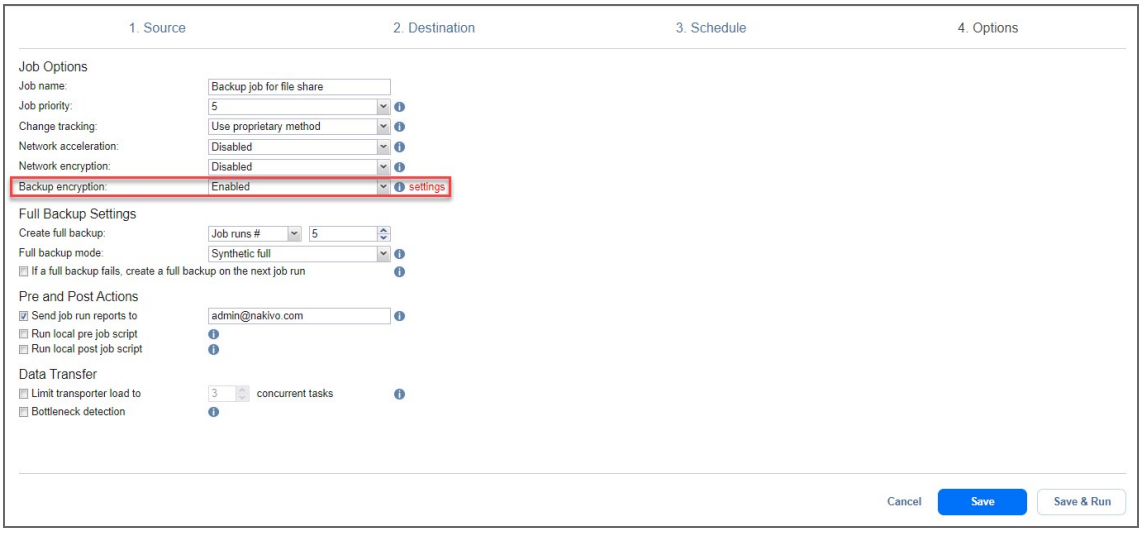
Enabling Backup Encryption

To set up backup encryption, do the following:

1. In the **Options** step of the corresponding backup/backup copy job wizard, select **Enabled** from the **Backup encryption** drop-down list.
2. After the **Backup encryption** mode is enabled, the **settings** link becomes available.

Notes

- If not configured, the settings link is highlighted in red.
- For backup copy jobs, you can select from two options: **Enabled on source** and **Enabled on target** to enable **Backup encryption**. For more details, refer to Backup Encryption.

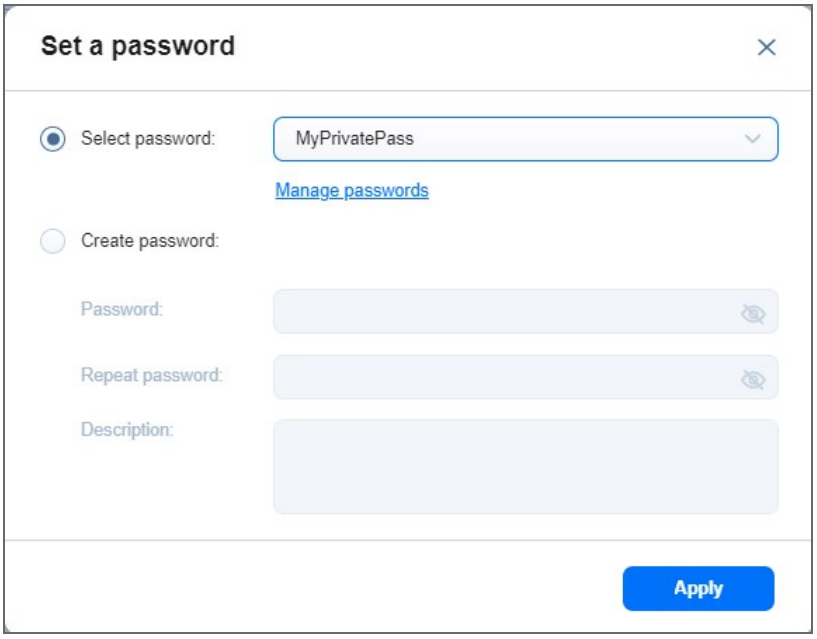


3. Click the **settings** link to open the **Set a Password** dialog box.

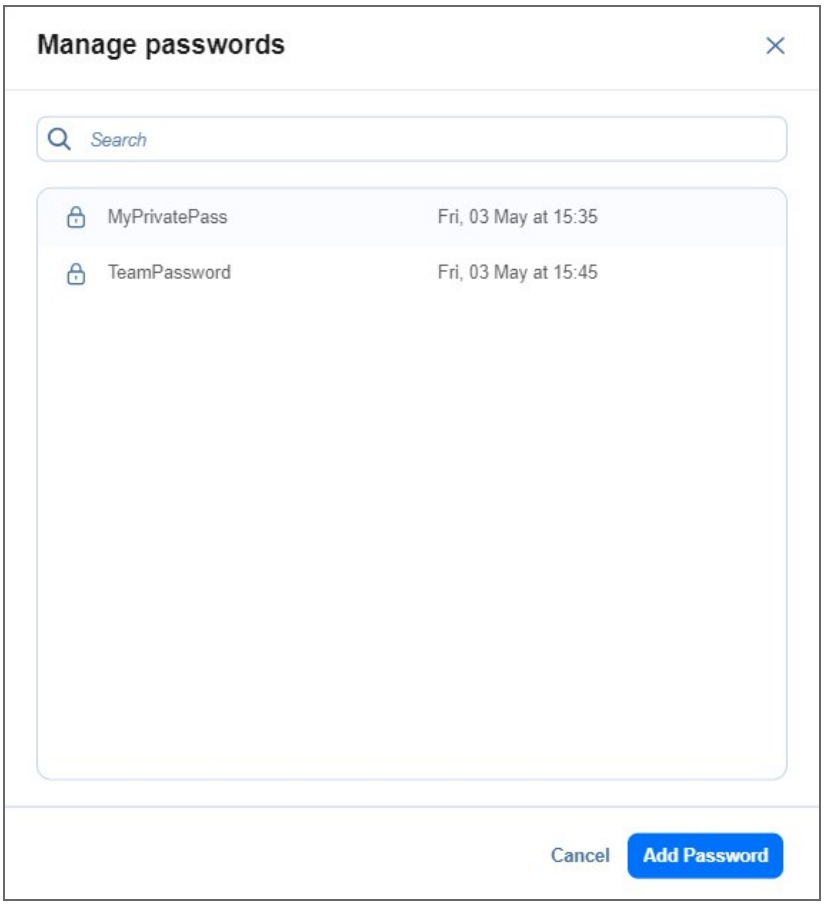
Notes

- If the **Key Management Service (KMS)** is not enabled, a warning message is displayed. To enable KMS, click the **Encryption Tab** link to go to the **Settings > General > System Settings > Encryption** tab. After enabling KMS, you can proceed with setting a password.
- It's recommended that you enable the **(AWS) Key Management Service**. If **AWS** is enabled, all backup encryption passwords encrypted with the Key Management Service cryptographic key are available for recovery in case of product re-installation. For more information, refer to "Enabling KMS" on page 311.

4. In the dialog box that appears, select the needed password or create a new one. Refer to "Setting Password" on page 307 for more details.



5. Optionally, you can click the **Manage passwords** link to manage the existing or add a new password to the list of passwords.



6. Click **Apply** to proceed.
   The product automatically generates the password hash based on the user password
   The cryptographic salt used for hash creation is saved in the recovery point metadata.
   The password hash is used to generate a single-use encryption key to encrypt the backup and FSI data (if the **FSI** option is enabled in the job).
   - If **KMS** is enabled, the password hash is encrypted with the AWS KMS key and saved in the recovery point metadata.
   - If **KMS** is not enabled, a dialog box opens, warning you that if you lose the password, it will be impossible to decrypt your data, and this data will be lost forever.
7. Click the **Proceed** button to go to the next step of the wizard.
   - To enable KMS, click the **Encryption Tab** link to go to the **Settings > General > System Settings > Encryption** tab.
   - Click the **Cancel** or **X** button to close the dialog box without applying any changes.

Notes

- Single chain of incremental recovery points has to have consistent encryption settings including encryption password.
- Changing encryption settings in a job (including changing encryption password) results in creating an active full backup and starting a new chain or recovery points.

Managing Passwords

With NAKIVO Backup & Replication, you can create and then manage your passwords for encrypting/decrypting backups and FSI data stored on the Director or exporting/importing the system configuration stored as self-backups.

Setting Password

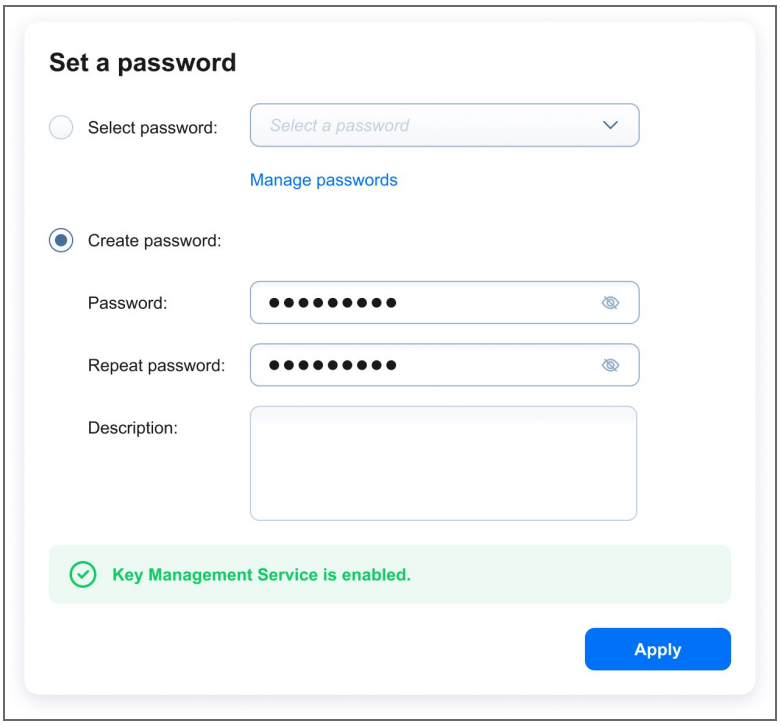
To set a password, proceed as follows:

1. In the **Set a Password** dialog box, choose the **Select password** option and select the needed password from the list of passwords to apply it to the job. Close the dialog box.

   Note

   The **Select Password** option is disabled if no passwords are available.

2. Optionally, you can click the **Manage passwords** link to manage the existing or add a new password to the list of passwords.
3. To create a new password, choose the **Create password** option. In the fields below, enter the password, repeat it, and add a description (required).



Note

By default, for the source-side backup encryption, the encryption key is applied to data in the trusted zone from the target side. If the Allow sending encryption key to source side. Not Recommended checkbox is selected, the encryption key will be sent to the non-trusted zone to arrange communication from the trusted to a non-trusted zone.

4. Click **Apply** to close the dialog box and apply the set password to the job.
   Click the **Cancel** or **X** button to close the dialog box without applying any changes.
   - If the **Key Management Service (KMS)** is enabled, the password hash is encrypted with the AWS KMS key and saved in the recovery point metadata.
   - If the **KMS** is not enabled, a warning message is displayed.
   - Click the **Proceed** button to dismiss the warning and apply the set password to the job and close the dialog box.
   - To enable KMS, click the **Encryption Tab** link to go to the **Settings > General > System Settings > Encryption** tab.

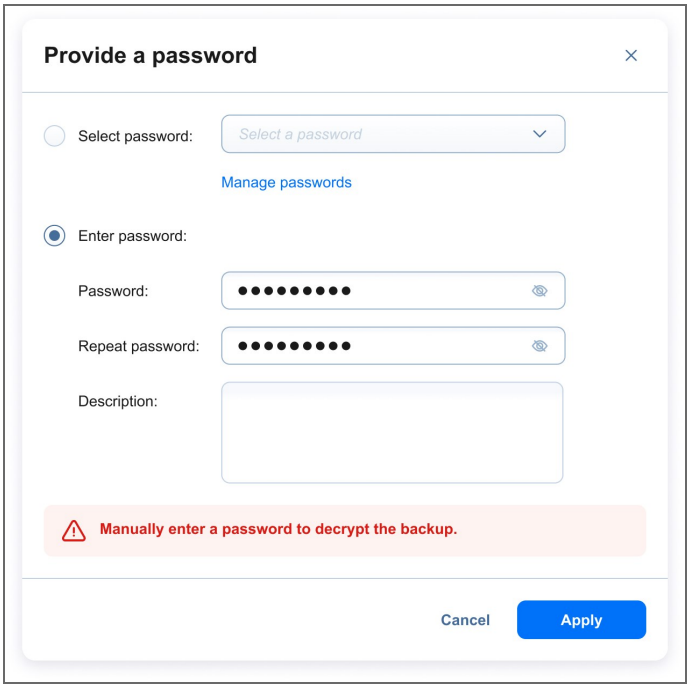
Note

It's recommended that you enable the **(AWS) Key Management Service**. If **KMS** is enabled, all backup encryption passwords encrypted with the Key Management Service cryptographic key are available for recovery in case of product re-installation. For more information, refer to "Enabling KMS" on page 311.

Providing Password

To provide the password, do the following:

1. In the **Provide a Password** dialog box, select the needed password from the list of passwords or enter it in the provided fields.
2. Optionally, you can click the **Manage passwords** link to manage the existing or add a new password to the list of passwords.
3. To enter a password, choose the **Enter password** option. In the fields below, enter the password, repeat it, and add a description (required).
4. Click **Apply** to proceed.
   Click the **Cancel** or **X** button to close the dialog box without applying any changes.



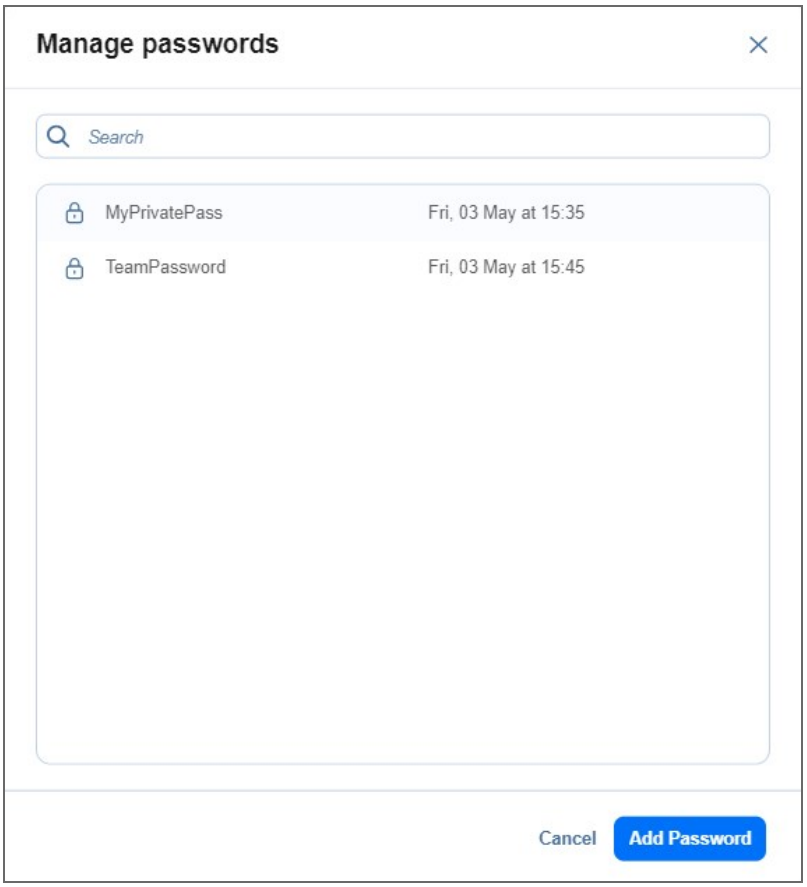
Managing Passwords

In the **Manage passwords** dialog box, you can do the following:

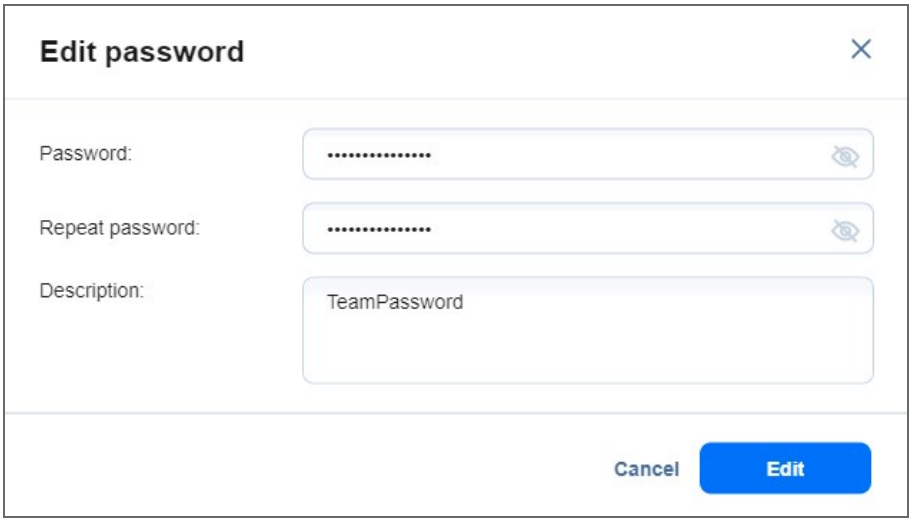
- View the list of the saved passwords in the current tenant with their descriptions and the timestamps of their last updates.

  Note
  If the description or timestamp does not fit it is truncated. To display full password information, hover over it.

- Search for passwords by entering the description (fully or partially) into the **Search** bar.
- Manage the existing passwords by hovering over the needed item and selecting from the list of typical actions applicable to it:



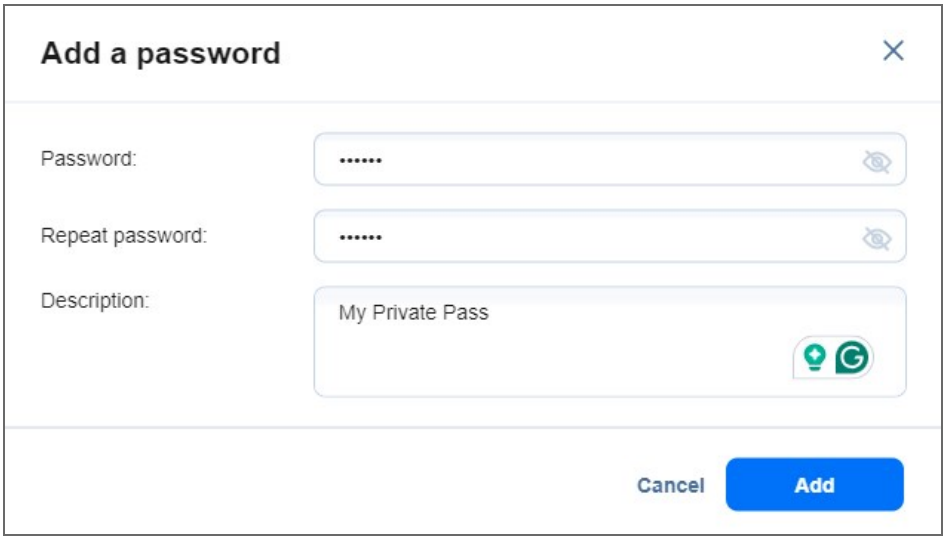
- Select **Edit** to open the **Edit password** dialog box, where you can edit the existing password.
    - Click the **Edit** button to confirm changes or the **Cancel** or **X** button to close the dialog box without applying any changes.



- Select **Delete** to display a confirmation dialog box, where you can confirm deleting the existing password. Click **Cancel** to dismiss the dialog box without applying any changes.

> **Note**
> The confirmation dialog box is displayed if the rule is not used by any jobs.

- Add a new password by clicking the **Add Password** button. In the dialog box that appears, enter the password, repeat it, and provide its description (optional), then click **Add**.



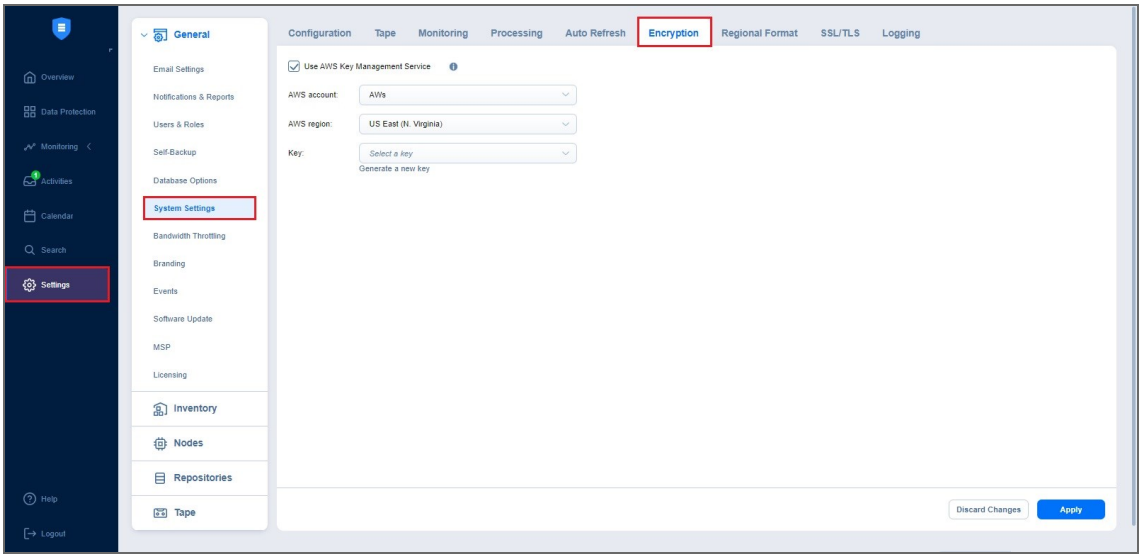
- Click the **Cancel** or **X** button to close the dialog box without applying any changes.

Enabling KMS

For encrypting the password hash, the AWS Key Management Service is used.

To enable the **(AWS) Key Management Service** proceed as follows:

1. Open the **Settings > General > System Settings > Encryption** tab.
2. Check the **Use AWS Key Management Service** checkbox (disabled by default).
3. Specify the AWS account (the AWS account should be discovered first) by selecting from the **AWS account** dropdown (the option is disabled if no AWS accounts have been discovered or if the discovered AWS account has insufficient permissions to use KMS). For the list of required AWS permissions for creating the (AWS) KMS Keys refer to the Feature Requirements section.
4. Specify the AWS region by selecting from the **AWS region** dropdown (the option is disabled if no AWS account has been discovered).



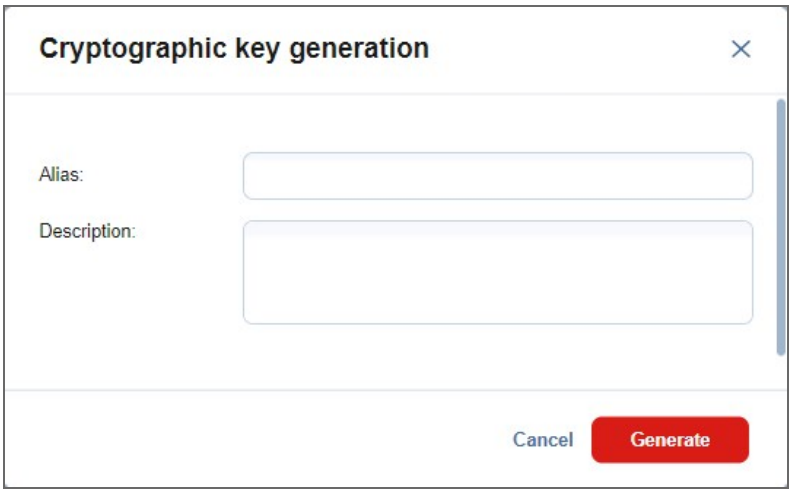
5. From the **Key** dropdown, select a key from the list of existing symmetric cryptographic keys available to the specified AWS account (the option is disabled if no keys are available).
6. Click **Apply** to apply the changes.
7. If encryption is set for a job, the password hash is generated based on the provided password. The password hash gets encrypted with the KMS cryptographic key with *base64* and is saved in the recovery point metadata.

Notes

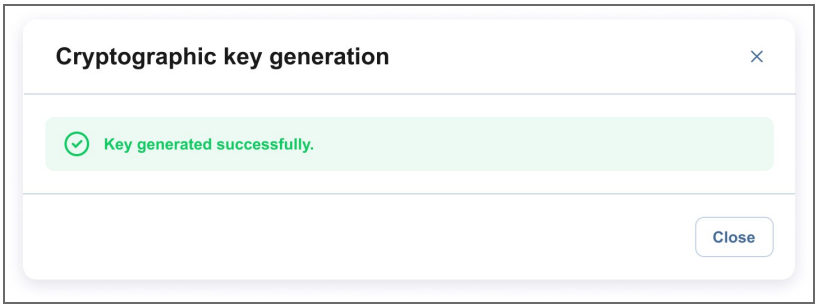
- If the AWS Key Management Service is enabled, the password hashes are automatically be restored in case of product reinstallation. Otherwise, passwords need to be provided manually in case the password hashes are not available.
- The AWS Key Management Service is not applied to self-backup and system configuration encryption.

Generating New KMS Cryptographic Key

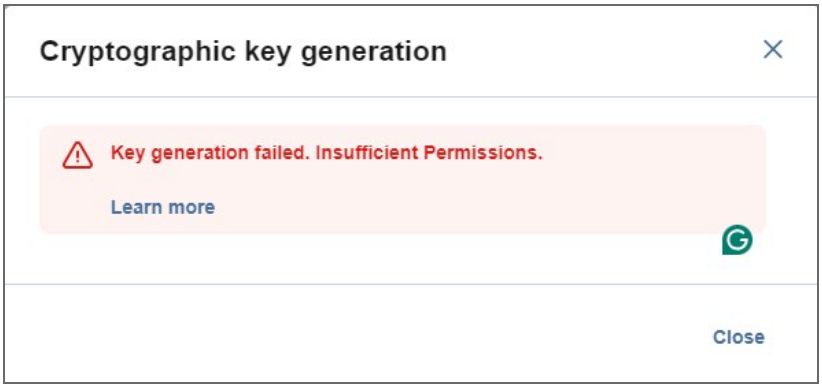
1. Optionally, click **Generate a new key** to open the **Cryptographic key generation** dialog box.



2. In the dialog box that appears, enter the alias and its description (optional) and click **Generate** to initiate generation of a new symmetric cryptographic key in the specified account. For more details, refer to the AWS::KMS::Alias page. Click the **Cancel** or **X** button to close the dialog box without applying any changes.
3. If the key was generated successfully, the following dialog appears:



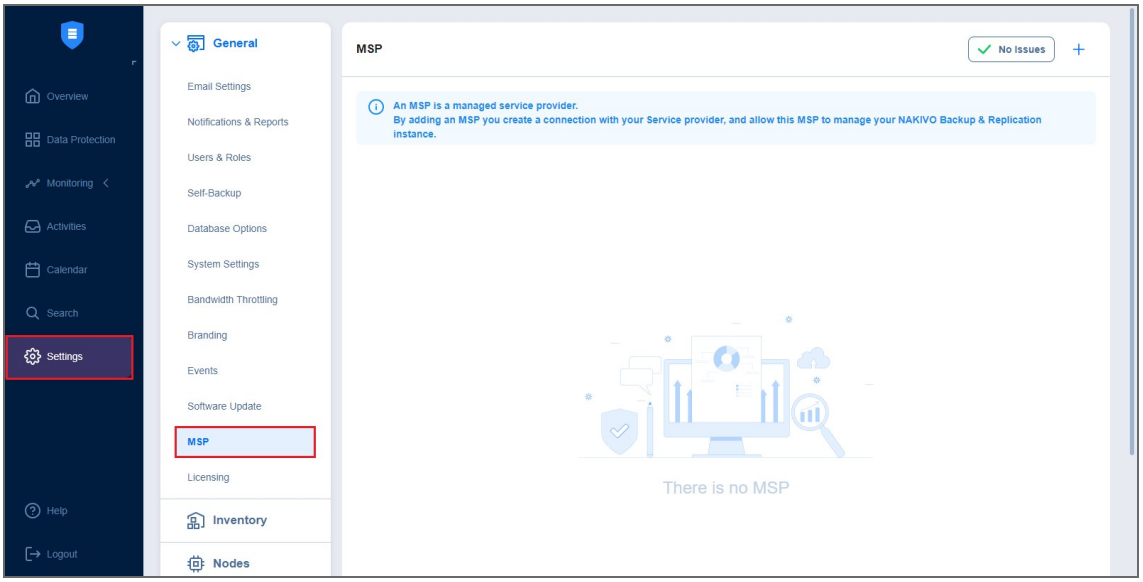
4. If the key generation fails, a dialog box with the reason for the failure appears:



5. Click the **Close** or **X** button to proceed with key generation.

MSP Tab in Single-Tenant Mode

To allow an MSP to manage your standalone instance of NAKIVO Backup & Replication, you need to link your instance to the MSP's instance of the solution. To do so, you should first add and configure the MSP details on the **MSP** page in **Settings**. These details can be provided by the MSP after Remote Tenant Configuration. Establishing a connection to an MSP as a remote tenant allows the MSP to monitor and manage your instance of NAKIVO Backup & Replication.



As a remote tenant, you can manage an established connection to an MSP on the **MSP** page.

See the topics below for more information:

- "MSP Architecture" below
- "Adding an MSP" on page 317
- "Managing an MSP Connection" on page 319

MSP Architecture

The MSP infrastructure in NAKIVO Backup & Replication consists of the following components:

- **MSP Console**: The centralized control hub that enables service providers to efficiently manage the entire cloud backup and replication processes. It is available to users with an *MSP license*, *Beta instance*, *Promo license*, or *Trial license*.
- **MSP Dashboard**: The Master Tenant Dashboard that is available for a multi-tenancy product with a non-MSP license.
- **Director**: This core solution component orchestrates and coordinates the other components of the NAKIVO Backup & Replication infrastructure for MSPs, ensuring a cohesive and streamlined operation.
- **Transporter**: Transporters are responsible for efficiently moving data between different components and optimizing the backup and replication processes.

- **Backup Repository**: These are storage locations within the NAKIVO Backup & Replication cloud infrastructure. Backup repositories are designed to store the backups of local tenant workloads. As primary and secondary backup storage, backup repositories can be used to implement the recommended 3-2-1 backup strategy.
- **Master Site and Remote Site(s)**: These terms encompass the main NAKIVO Backup & Replication infrastructure components (that is, **Director**, **Transporter(s)**, and **Backup Repositories**) in an MSP setting. The Master site (on the MSP side) serves as the central point for managing multiple Remote sites (client sites).
- **Remote Tenant**: This is a tenant with a standalone NAKIVO Backup & Replication instance installed in single-tenant mode in their infrastructure and added to the multi-tenant installation of the MSP for monitoring purposes by the Master tenant (the MSP).
- **Local Tenant**: This is a tenant created and managed by the Master tenant (MSP) to provide backup and disaster recovery as a service to clients. To establish a connection with client remote resources, the MSP should enable Direct Connect.

Notes

- Only users with an *MSP license*, *Beta instance*, *Promo license*, or *Trial license* can additionally access the MSP Console and the Licensing > Tenants tab. This allows them to efficiently oversee all independent instances of NAKIVO Backup & Replication associated with a managed service provider (MSP) as well as local tenants from a unified interface, eliminating the need to navigate through individual tenants. Users with multi-tenant instances of NAKIVO Backup & Replication without these license types only have access to the MSP Dashboard.
- Multiple remote tenants, each tied to a separate client environment, can connect to a single instance of the MSP Director.
- Only local tenants consume MSP workloads. Since remote tenants are licensed independently, they have their own workloads.
- The Master tenant (MSP) operates each remote tenant independently from other tenants, managing and facilitating a specific tenant's backup, replication, and recovery operations.
- The **Direct Connect** feature allows the NAKIVO Backup & Replication instance installed at the MSP site to access the resources at the remote site of a client via a single port connection (the port should be exposed on the client's infrastructure) without the need to establish a VPN connection.

The MSP architecture in NAKIVO Backup & Replication empowers the MSP Director to oversee and synchronize with each tenant (client). It ensures the establishment of secure and isolated communication channels, facilitating both data exchange and administration within the distinct data protection environment of each tenant.

To establish these isolated communication channels, the MSP initiates the creation of a remote tenant account for standalone users. After a remote tenant at the MSP Director is created, the client with a standalone instance of NAKIVO Backup & Replication can use the provided credentials to connect to the MSP's NAKIVO Backup & Replication instance. Refer to "Creating a Remote Tenant" on page 660 for more details.

Once the setup is done on both sides, the MSP can manage the Remote tenant like a regular user instance on the remote tenant side. This way, the MSP (Master tenant) manages the remote tenants (client sites) with a single tenant from the NAKIVO Backup & Replication instance installed on another Director. For more information, refer to "Managing an MSP Connection" on page 319.

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

Running an MSP Connection

The general workflow for running a connection between the remote tenant instance and the MSP's instance of NAKIVO Backup & Replication is as follows:

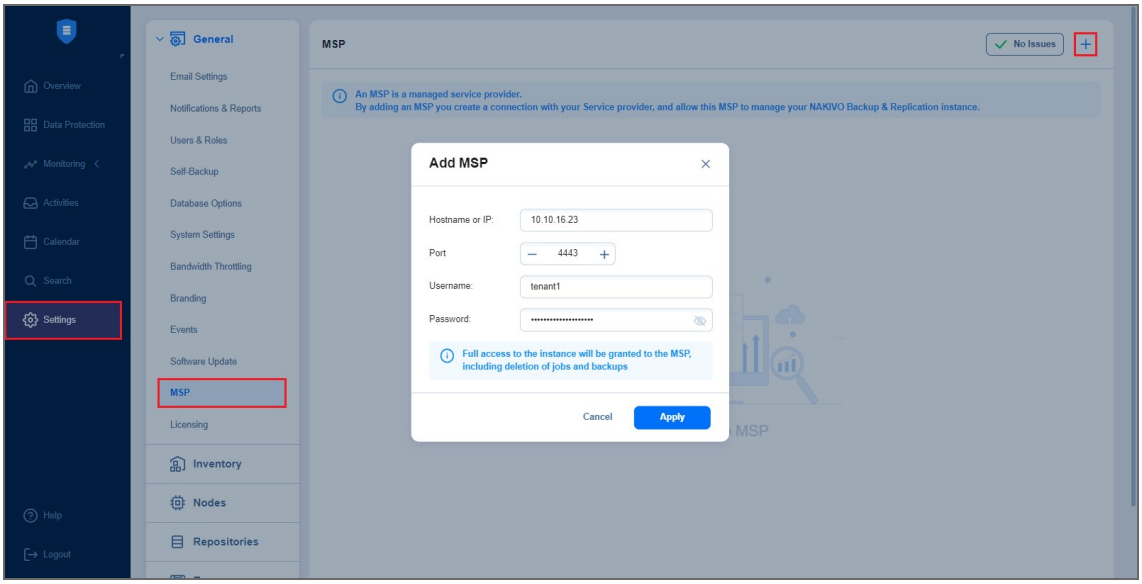
- At startup, each remote tenant connects to the MSP Director to establish and maintain a single TCP connection (using port *4443*) for various types of traffic, including:
  - MSP Dashboard traffic (between MSP and remote tenant): the established connection on port *6702 TCP* for bidirectional communication between the MSP Dashboard and the storage transporter on the client side.
  - Remote tenant traffic: All remote tenant-related activities, including drill-through into data, remote tenant actions, and other communication between the MSP and the storage transporter.
- In the course of initial connection establishment, the MSP informs the remote tenant about the listening port used.
- MSP opens the port in the firewall.
- Each remote tenant then maintains a dedicated and separate TCP connection with the MSP Director (using port *6702 TCP*) for bidirectional communication.
- The MSP Director, in turn, uses a separate listening port (*6702 TC*P by default) for communication with a remote tenant's instance (while *4443 TCP* is used for initial MSP - remote tenant connection setup/checks only).
- These connections serve as the conduits for all data transfers, administrative commands, monitoring, and management actions related to the specific remote tenant.

- Both ports, *6702 TCP* and *4443 TCP*, use TLS for securing the communication channels between the MSP Director and each remote tenant. TLS ensures encryption, authentication, and data integrity for the traffic exchanged between these components, maintaining data security for each tenant's connection.
- The ports used by the MSP Director (*6702 TCP*) and the initial connection setup port (*4443 TCP*) can be configured or customized by the MSP using Expert Settings to suit specific deployment requirements or security protocols.
- If the MSP changes the listening port used, the connection may be interrupted. For more information on the required TCP ports, see the **MSP Console** section in Feature Requirements.
- During a product update (on the Master tenant or Remote tenant side) the connection is suspended. After the update is completed, the remote tenant continuously tries to reestablish the connection with the Master tenant. In turn, the Master tenant side starts listening to the connection attempts from the Remote tenant side until the connection is reestablished.
- In case of a network connectivity issue, the connection between the MSP and the remote tenant can be lost. Once the issue is resolved, the Remote tenant side initiates infinite attempts to reestablish the connection with the Master tenant side. In turn, the Master tenant side starts listening to the connection attempts from the remote tenant side until the connection is reestablished. Alternatively, the connection can be reestablished manually.

Adding an MSP

To add an MSP to which you would like to link your NAKIVO Backup & Replication instance, you need the Master tenant's hostname or IP address, port number, and the remote tenant credentials they have generated for you. Once you obtain the above information, follow the steps below:

1. Go to **Settings** and click the **+** button in the top right corner.
2. Fill in the **Hostname or IP**, **Port**, **Username**, and **Password** fields based on the information provided by your Master tenant (MSP).

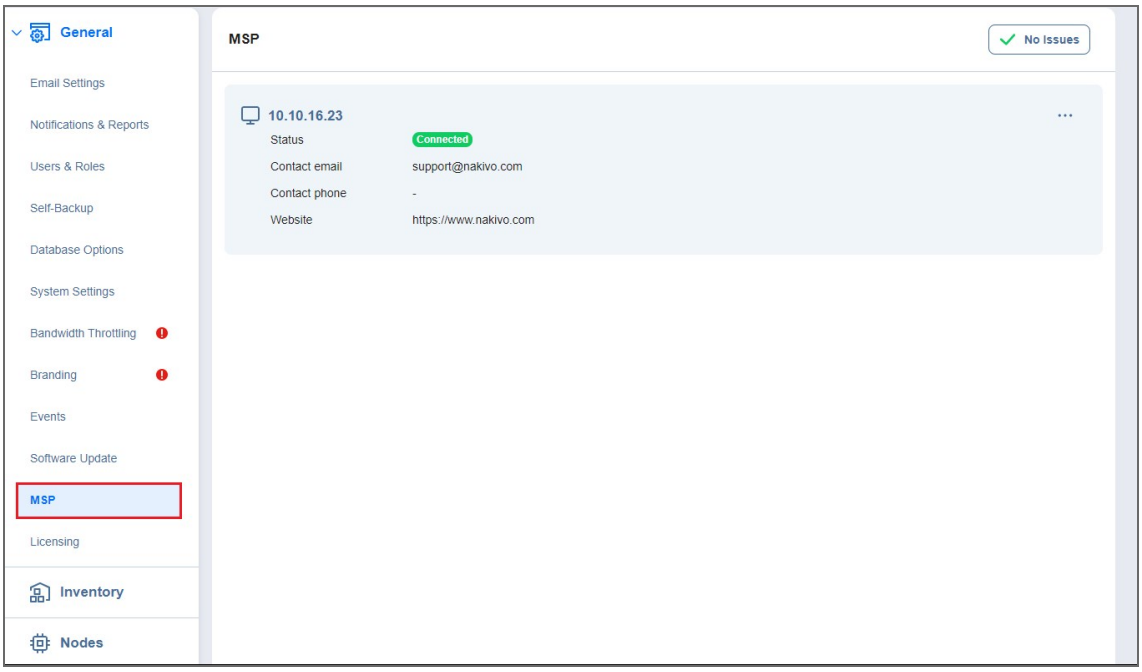


3. Click **Apply**. The following screen will display the MSP's certificate details.

Notes

- Your version of NAKIVO Backup & Replication must be the same as the MSP's version. Otherwise, you cannot connect to MSP's instance of NAKIVO Backup & Replication.
- The MSP uses a separate listening port for communication with a remote tenant's instance (port *6702* is used by default). If the MSP changes the listening port used, the connection may be interrupted. For more information on required TCP ports, see the **MSP Console** section in "Feature Requirements" on page 91.

4. Read through the MSP's certificate details and click **Accept**.
5. The added MSP should now appear in your **MSP** menu.



Once the MSP is added, you have successfully established a connection with the Master tenant as a remote tenant. This allows the Master tenant to access your instance of NAKIVO Backup & Replication as long as the connection is active.

Managing an MSP Connection

Several options are available for managing an established connection to an MSP. See the sections below for more information on managing yyour connection as a Remote tenant.

- Viewing MSP Details
- Disconnecting/Reconnecting an MSP
- Deleting an MSP

Viewing MSP Details

The MSP block in the **MSP** menu contains the following information:

- **Hostname or IP**: The hostname or IP address of the Master tenant.
- **Connection status**: Current status of your connection to the MSP.
  - **Connected**: Your NAKIVO Backup & Replication is connected to the MSP
  - **Disconnected**: Your NAKIVO Backup & Replication is disconnected from the MSP
  - **Connecting**: Your NAKIVO Backup & Replication is actively trying to establish a connection to the MSP
- **Contact information**: The Master tenant's email address, phone number, and website, updated automatically.



Disconnecting/Reconnecting an MSP

To disconnect your instance of NAKIVO Backup & Replication from an MSP, click the ellipsis **Manage** button in the top right corner of the MSP block. In the popup, click **Disconnect**. This will suspend the connection to the MSP until resumed.



To reconnect to the MSP, simply click the ellipsis **Manage** button in the top right corner of the MSP block and click **Connect** in the popup. You will not be asked to provide the same details you did when first connecting to the MSP unless the MSP has changed their certificate or your remote tenant credentials.

Deleting an MSP

To delete an MSP connection, click the ellipsis **Manage** button in the top right corner of the MSP block. In the popup, click **Delete** and confirm the action. This action will erase all tenant data from the MSP's side and vice versa.



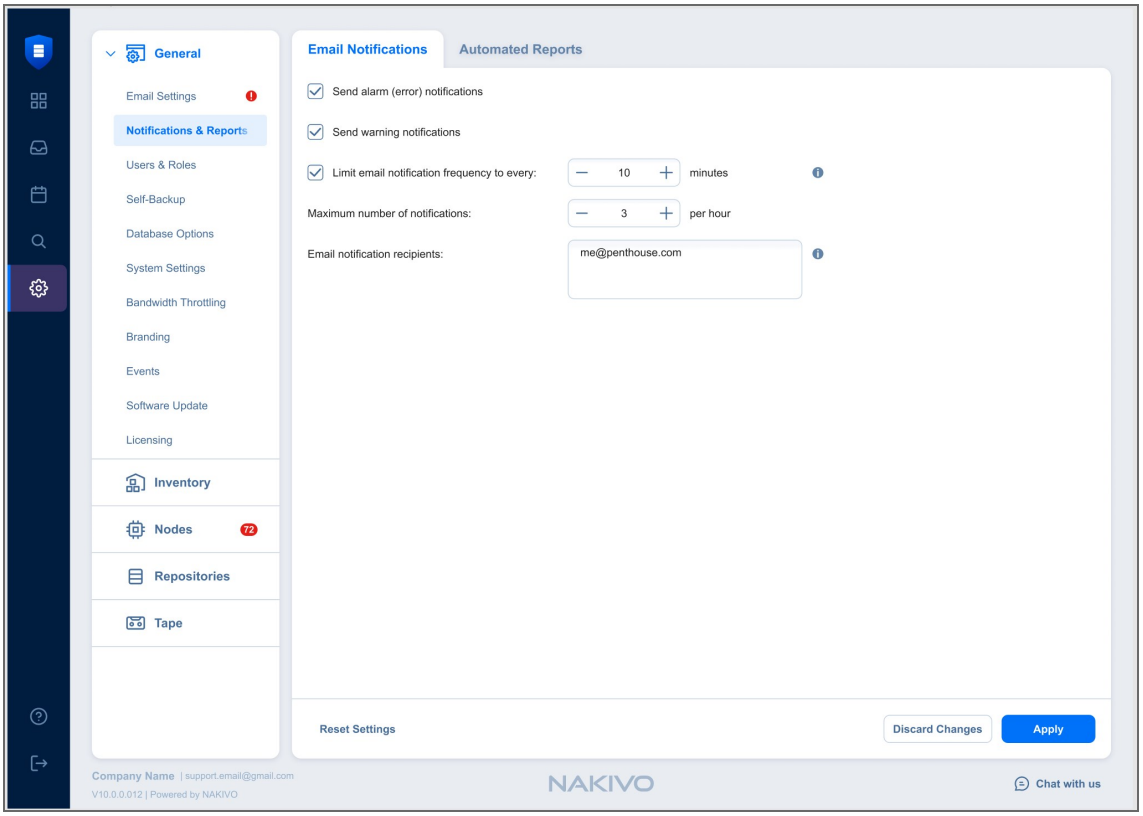
Notifications & Reports

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

- Email Notifications
- Automatic Reports

To receive automatic notifications, configure email settings by following the steps below:

1. Log in to NAKIVO Backup & Replication.
2. Click **Settings** in the left pane of the product.
3. Go to the **General** tab.
4. Click **Notifications & Reports** to configure notifications and automatic reports section on the page that opens.
5. Click **Apply** to save the settings after you're done.
6. Alternatively, click **Discard Changes** to discard any changes you have made to the email configuration.
7. Optionally, click **Reset Settings** to return all the settings to their default values.



Note

To configure email notifications and automatic reports, you must first configure email settings.

Email Notifications

To set Email notifications, fill out the fields in the *Email notifications* section:

- **Send alarm (error) notifications**: If this option is selected, NAKIVO Backup & Replication will send email notifications to the specified recipients in case an error (for example, a job failure) occurs in the product. For users in Multi-Tenant Mode, these notifications also identify the relevant tenant and the instance where the error occurred.
- **Send warning notifications**: If this option is selected, NAKIVO Backup & Replication will send email notifications to the specified recipients in case the product generates a warning message (for example, lost connection to a host or **Backup Repository**). For users in Multi-Tenant Mode, these notifications also identify the relevant tenant and the instance that generated the warning.
- **Limit email notification frequency**: This option allows you to set up an email notification frequency in minutes. If deselected, notification emails will be sent every 5 minutes with no hourly limit.
- **Maximum number of notifications**: Use this option to change the limit of email notifications receivable per hour. If this limit is reached, any additional notifications will be delivered the following hour.
- **Email notification recipients**: Specify the recipients who will be receiving alarm and warning notifications (if enabled).

Automatic Reports

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

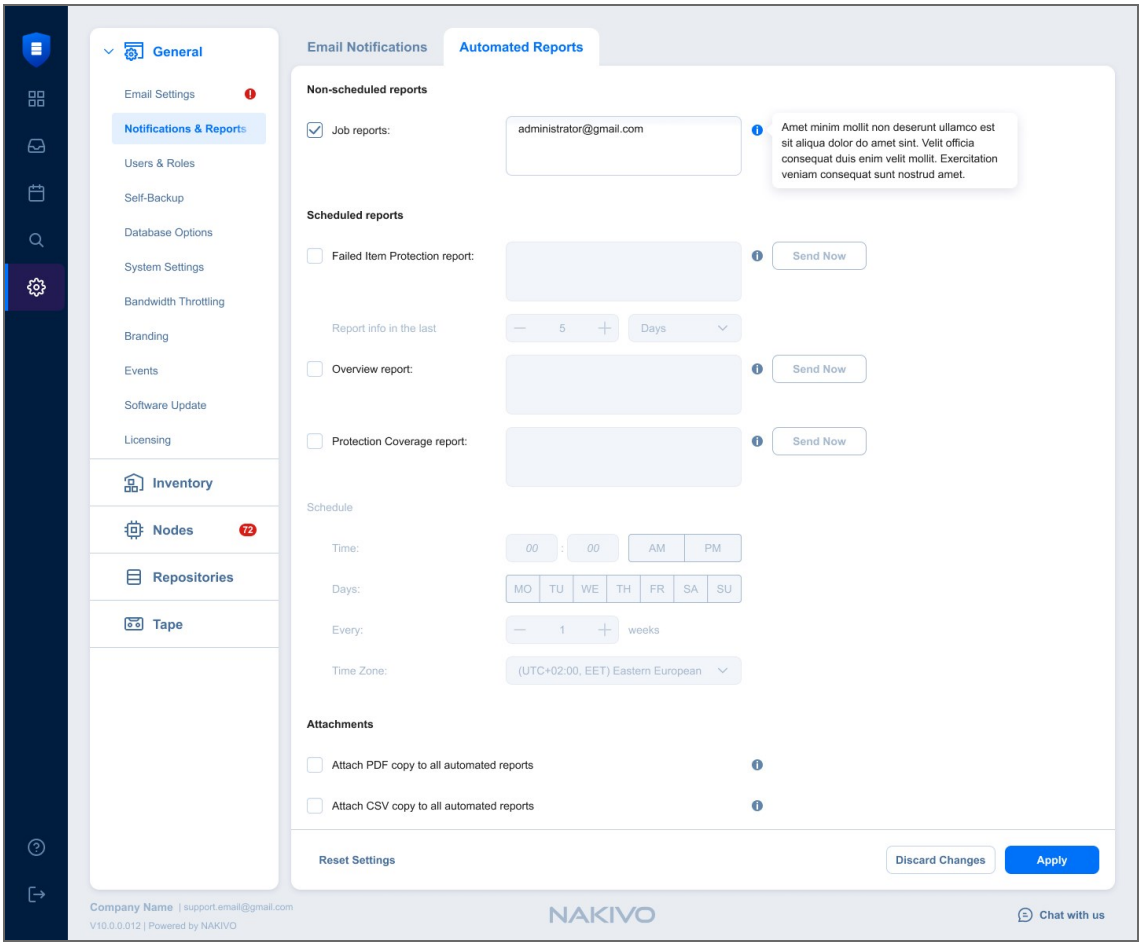
- **Job reports**: If this option is selected, NAKIVO Backup & Replication will send an HTML report after the completion of every job (regardless of the job success or failure) to email addresses specified in the text field. Use a semi-colon to separate multiple email addresses.
- **Failed Item Protection report**: Contains information about all items which had failed to be protected by backup and/or replication jobs, and the error message. Additionally, configure **Report info in the last** option by entering the number of days you want to get the report for.
- **Overview report**: If this option is selected, NAKIVO Backup & Replication will generate the Overview report (which includes information about all jobs and groups in the product) on the date and time specified in the scheduler and will send the report to the recipients specified in the text field. Use a semi-colon to separate multiple email addresses.
- **Protection Coverage**: If this option is selected, NAKIVO Backup & Replication will generate the Protection Coverage report. This includes information about all VMs & instances protected by backup and/or replication jobs as well as the information about all unprotected VMs & instances. The report will be sent to the recipients specified in the text field on the date and time specified in the scheduler. Use a semi-colon to separate multiple email addresses.
- **Schedule**: Configure the schedule at which you want to get the reports.

- **Attach PDF copy to all automated reports**: Select this option to get the additional attached copy of the report in the PDF format.
- **Attach CSV copy to all automated reports**: Select this option to get the additional attached copy of the report in the CSV format.

Note

NAKIVO Backup & Replication supports the following special characters in reports:

- US special characters
- Characters in the following languages:
  - Vietnamese
  - Japanese
  - Korean
  - Chinese
  - Arabic



Self-Backup

The self-backup feature allows you to automatically protect configuration settings of your NAKIVO Backup & Replication instance. For more information, refer to “Self-Backup Feature” on page 17.

Notes

- Self-backup is not supported for the multi-tenant configuration.
- A federated repository cannot be used for self-backup.

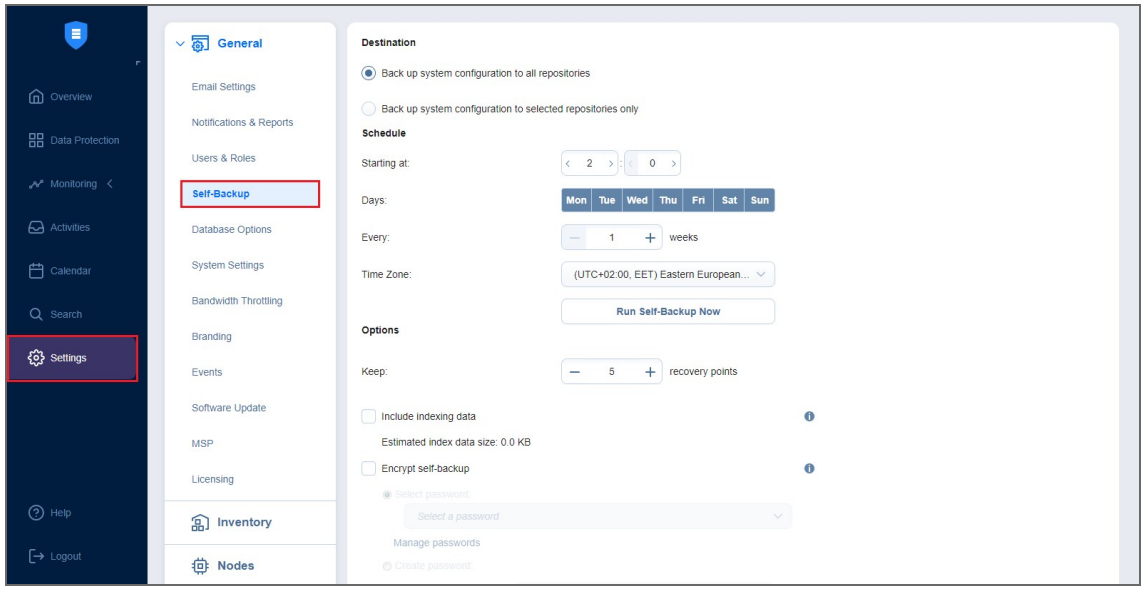
To configure self-backup options, proceed as described in the following sections:

- Accessing Self-Backup Options
- Setting Up Self-Backup Destination
- Self-Backup Schedule
- Self-Backup Options
- Self-Backup Encryption
- Recovering from Self-Backup

Accessing Self-Backup Options

To access self-backup options, follow the steps below:

1. Click **Settings** in the left pane of NAKIVO Backup & Replication.
2. Go to the **General** tab and click **Self-Backup**.
3. After making the necessary changes, click **Apply**. Alternatively, click **Discard Changes** to discard any changes you have made.



Setting Up Self-Backup Destination

To configure a self-backup destination, follow the steps below:

1. Select **Back up system configuration to all repositories** to enable all repositories in the list of repositories where system configuration is backed up. If deselected, you can remove specific repositories from the list.
2. Alternatively, select **Back up system configuration to selected repositories only** and select specific repositories you wish to use for self-backup.
3. If necessary, add a **Backup Repository** to the list:
    - Click the "**+**" icon to add repositories to the list of repositories for system backing up.
    - In the **Add Backup Repositories** dialog that opens, select the necessary repositories and close the dialog.

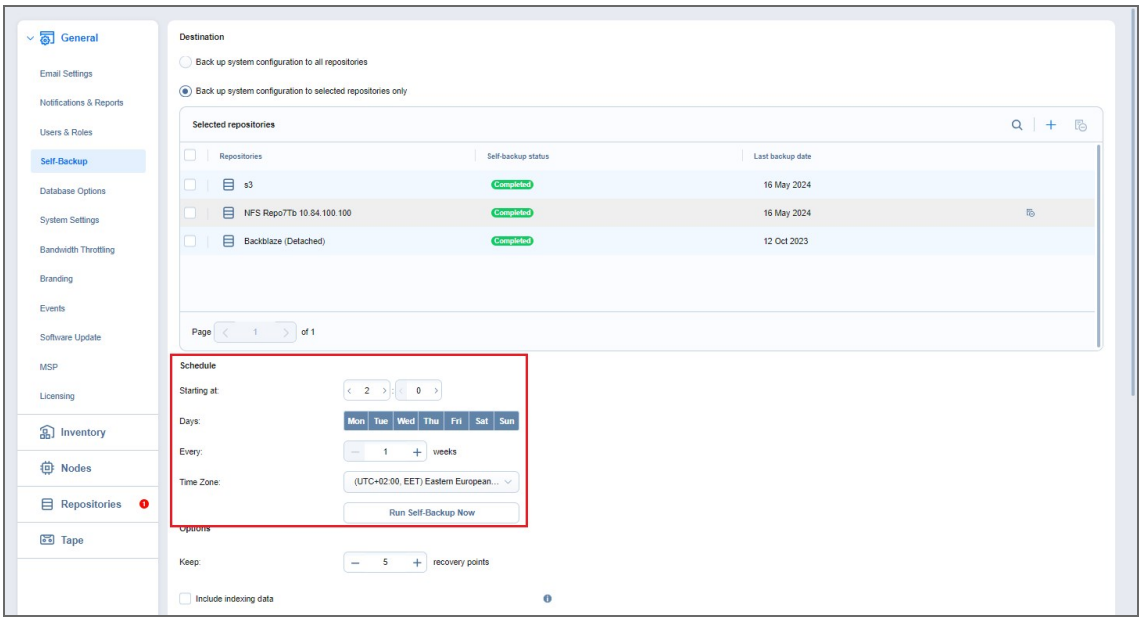
Notes

- Federated repositories are disabled on the list of repositories used as a target for self-backup.
- Federated repository members are not displayed on the list of repositories available for self-backup.

Self-Backup Schedule

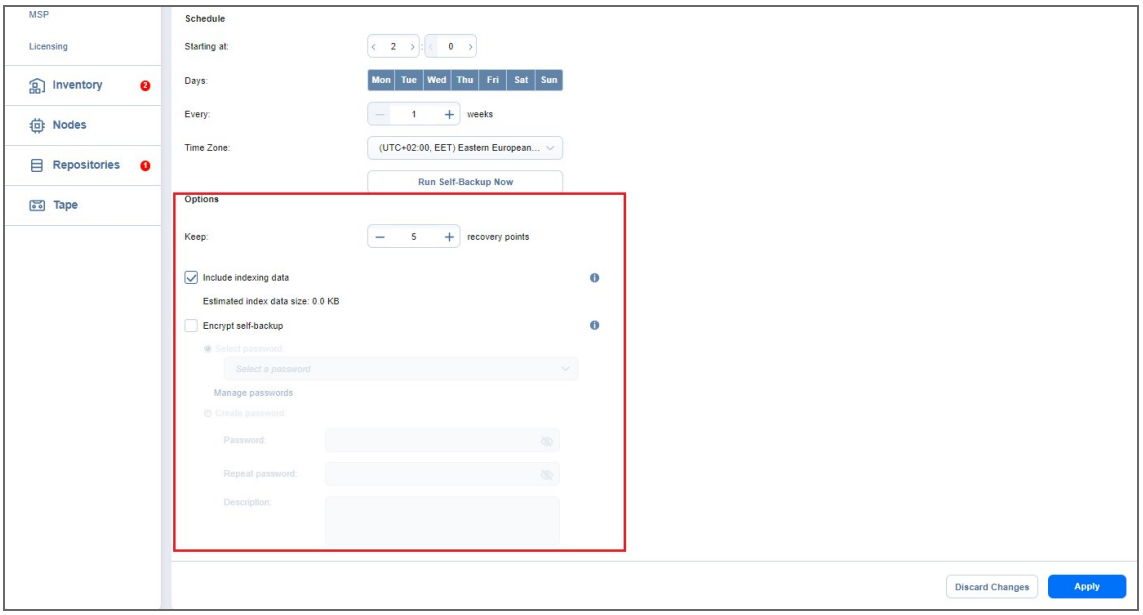
To configure the self-backup schedule, follow the steps below:

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



Self-Backup Options

In the **Options** section of the self-backup settings, you can enter a number of recovery points to be kept for the self-backup. To apply your settings, click the **Apply** button.



Self-Backup Encryption

A system configuration stored as a self-backup can be encrypted to safeguard data from loss, corruption, or unauthorized access to data.

Select **Encrypt self-backup** to encrypt your backup for additional security. Afterwards, select the password or create a new one in the corresponding fields. See Managing Passwords for more details.

Note

AWS Key Management Service is not applied to self-backup encryption.

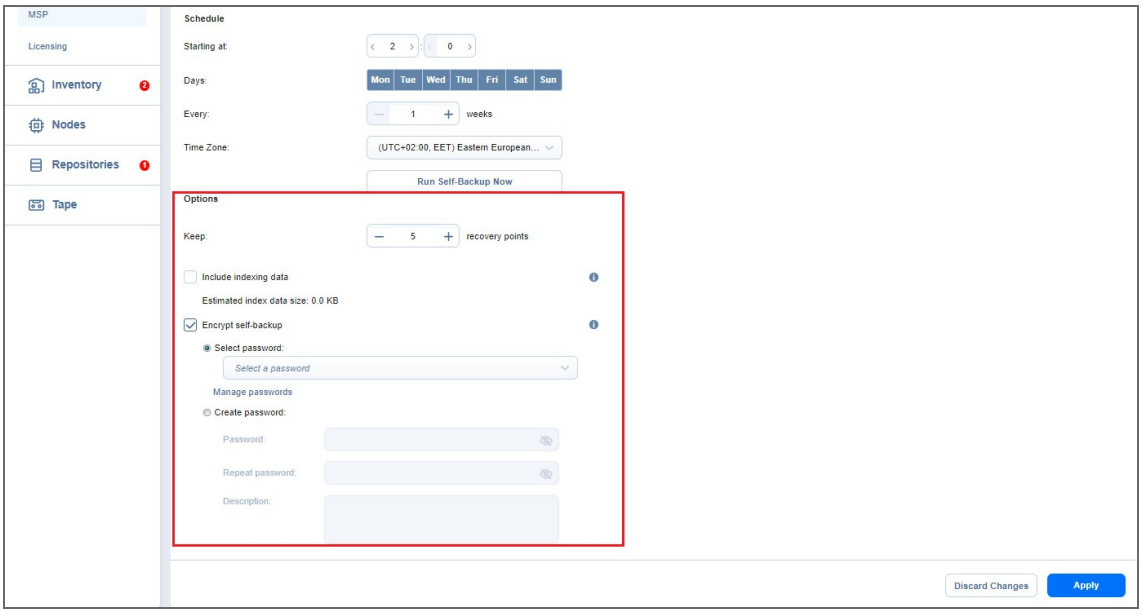
Set up self-backup encryption as follows:

1. Select the **Encrypt self-backup** checkbox (disabled by default) to store self-backups in an encrypted form.
2. Select the **Select password** option and click on the field to display a list of saved passwords with descriptions.
3. Select the needed password.

Note

The **Select Password** option is disabled if no passwords are available.

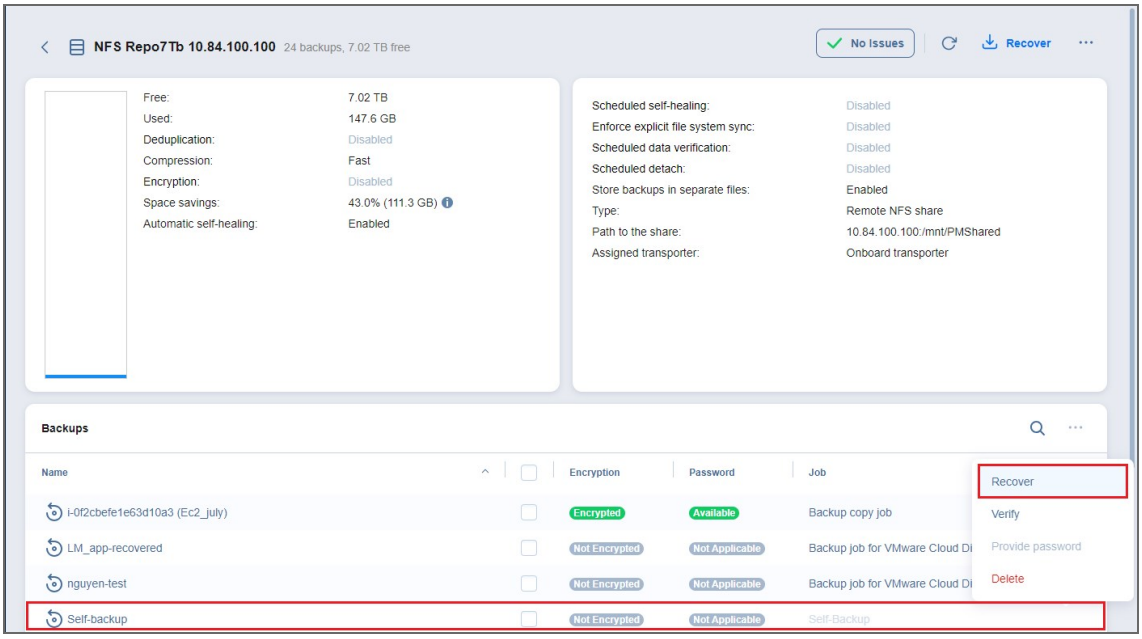
4. To create a new password, select the **Create password** option.
5. Enter the new password and its description and repeat it.
6. If a newly created password is applied, the **Select password** option becomes selected and the new password is pre-selected in the **Select password** dropdown and is added to the list of saved passwords.
7. Click the **Cancel** or **X** button to close the dialog box without applying any changes.
8. Optionally, you can click the **Manage passwords** link to manage the existing or add a new password to the list of passwords. For more information, refer to Managing Backup Encryption.



Recovering from Self-Backup

To recover the configuration of NAKIVO Backup & Replication from a self-backup stored in a **Backup Repository**, do the following:

1. Go to **Settings** > **Repositories**.
2. Select one of the repositories that contain a self-backup.
3. Select the self-backup from the **Backups** list and click **Recover**.



4. Click **Restore.**
5. Wait for the system configuration to be restored. When the self-backup recovery process is completed, a message announcing success appears.

> **Note**
>
> After a system configuration backup, the encrypted recovery points passwords that were created and saved in the password manager become unavailable and providing them manually is required.

Software Update

- Download & Update Option
- Download Option

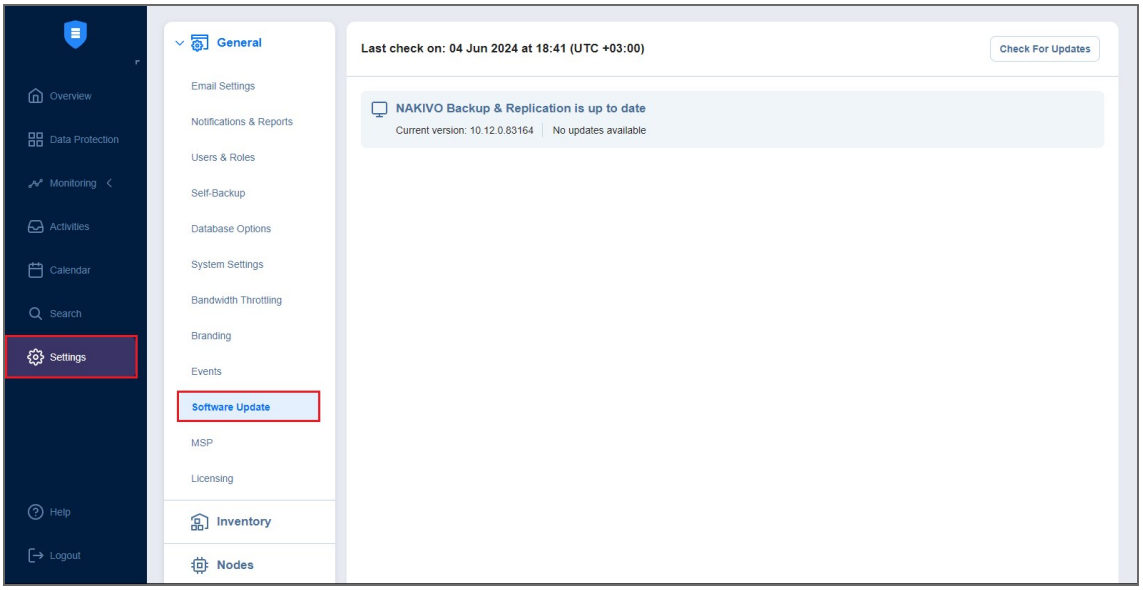
When the full solution of NAKIVO Backup & Replication (that is, the Director and the Transporter) is installed on a Windows or Linux machine, you can download product updates from the **Software Update** tab in the web interface. This feature automatically updates your NAKIVO Backup & Replication instance, the Onboard Transporter, and any other nodes that support auto-update.

For a list of supported nodes and requirements for the auto-update feature, see the **Auto-Update** section in "Feature Requirements" on page 91.

> **Note**
>
> In case **I/O Filter** was installed on the Cluster before the software update and Auto Refresh Inventory is disabled, manually refresh the **Inventory** after an software update to check that **I/O Filter Daemons** are displayed in the **Nodes** tab, and ensure Real-Time Replication jobs run successfully.

To check if an update is available, do the following:

1. Click **Settings** in the left pane of the product.
2. Go to the **General** tab.
3. Go to the **Software Update** page.
4. Click **Check for Updates** if needed.



Note

If you are using a multi-tenant solution, only master-tenant users with the appropriate permissions can see and manage software updates.

Download & Update Option

To download and install the update, do the following:

1. Optionally, click **Release Notes** to see features and improvements implemented in the new product version.
2. Select the **I have read the Release Notes** checkbox.
3. Click **Download & update**.
4. Click **Update Now**.

Before downloading the update, the product performs a self-backup and stops all current activities including running jobs, recovery jobs, repository maintenance, etc. When the download is complete, the product updating process begins. The product downloads the update to the **Director** first. When the **Director** is updated, the update is automatically uploaded to the **Transporters** which are then updated simultaneously. If some **Transporters** are not updated, you can update them manually. Refer to the corresponding articles for details.

Notes

- For a list of supported **Transporters**, see the **Auto-Update** section in "Feature Requirements" on page 91.
- Only 20 **Transporters** can be updated simultaneously. All other **Transporters** will be sent to a queue and updated once the previous update is completed.

Download Option

If you wish to postpone an update or schedule it, take the following steps to download the update without installing it:

1. Optionally, click **Release Notes** to see features and improvements implemented in the new product version.
2. Select the **I have read the Release Notes** checkbox.
3. Click **Download**.
4. After the download is completed, do one of the following:
    - Click **Update Now** if you want to start the updating process. Updating the product will stop all current activities, including running jobs, recovery jobs, repository maintenance, etc.
    - Click **Schedule Update** to update the solution at a specific time:

1. In the dialog box that opens, pick a day and time for updating. Click **Apply**.
2. On a working day before the scheduled update, you will see a notification in the product menu with the **Update Reminder** dialog box. By hovering over this notification, you can:
    a. Click **Reschedule** if you want to reschedule the update and pick a different time.
    b. Click **Cancel update** to cancel updating the full solution.

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

System Settings

To configure the system settings, follow the steps below:

1. Click **Settings** in the main menu on the left.
2. Go to the **General** tab and click **System settings.**
3. Set the System Settings options as shown in the following sections:
    - Setting the Configuration Options
    - Configuring Tape Settings
    - Configuring Monitoring Settings
    - Configuring Processing Settings
    - Configuring Auto Refresh Settings
    - Configuring Encryption Settings
    - Configuring Regional Format Settings
    - Configuring SSL/TLS Settings

    4. After making the necessary changes, click **Apply**. Alternatively, click **Discard Changes** to discard any changes you have made.

Notes

- NAKIVO Backup & Replication supports Certificates with the RSA algorithm only.
- In the *Web Interface TLS/SSL Certificate* section, you can see a notification about imminent TLS/SSL Certificate expiration in 30 days and onwards. If your certificate has expired, you will be asked to install a valid certificate.

Setting the Configuration Options

With the **Configuration** tab, you can define how long system events and job histories are stored, automatically retry failed jobs, manage user sessions and inactivity timeouts, and decide how the system responds to critical and non-critical errors. Set the following options:

- **Store system events for the last x days**: Events older than the specified number of days (can be from 5 to 365) are deleted.
- **Store job history for the last x days**: The history of the jobs older than the specified number of days (can be from 5 to 90) is deleted.

    Note

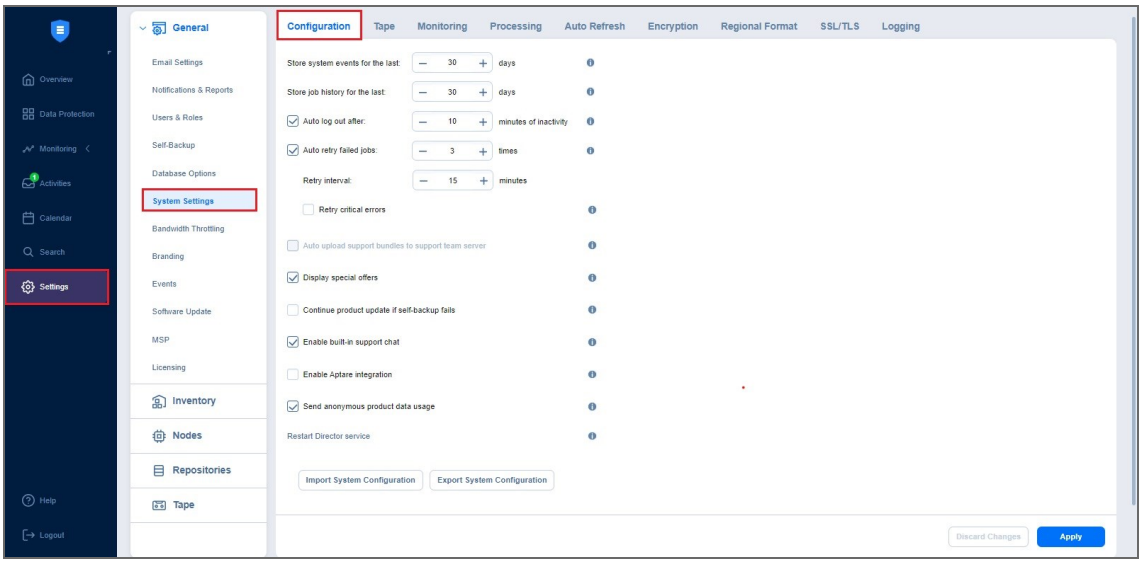
    This option is not displayed for the Master tenant in Multi-tenancy mode.

- **Auto log out after x minutes of inactivity**: When this option is selected, the current user is automatically logged out of NAKIVO Backup & Replication after the specified period of inactivity.

- **Auto retry failed jobs x times with y minutes interval**: When this option is selected, failed jobs are automatically retried the specified number of times (from 2 to 10) and with the specified time interval (from 1 to 60). Jobs with failed backup, replication, and recovery remain in the "running" state until all retries have either succeeded or failed.
- **Retry critical errors**: When this option is selected, NAKIVO Backup & Replication tries to automatically rerun jobs with critical and non-critical errors a specified number of times.

> **Notes**
>
> - The term **critical error** refers to persistent errors that are unlikely to change without any additional intervention, that is, hardware failure.
> - The term **non-critical error** refers to non-persistent errors that are likely to change without any additional intervention, that is, unstable network connection.

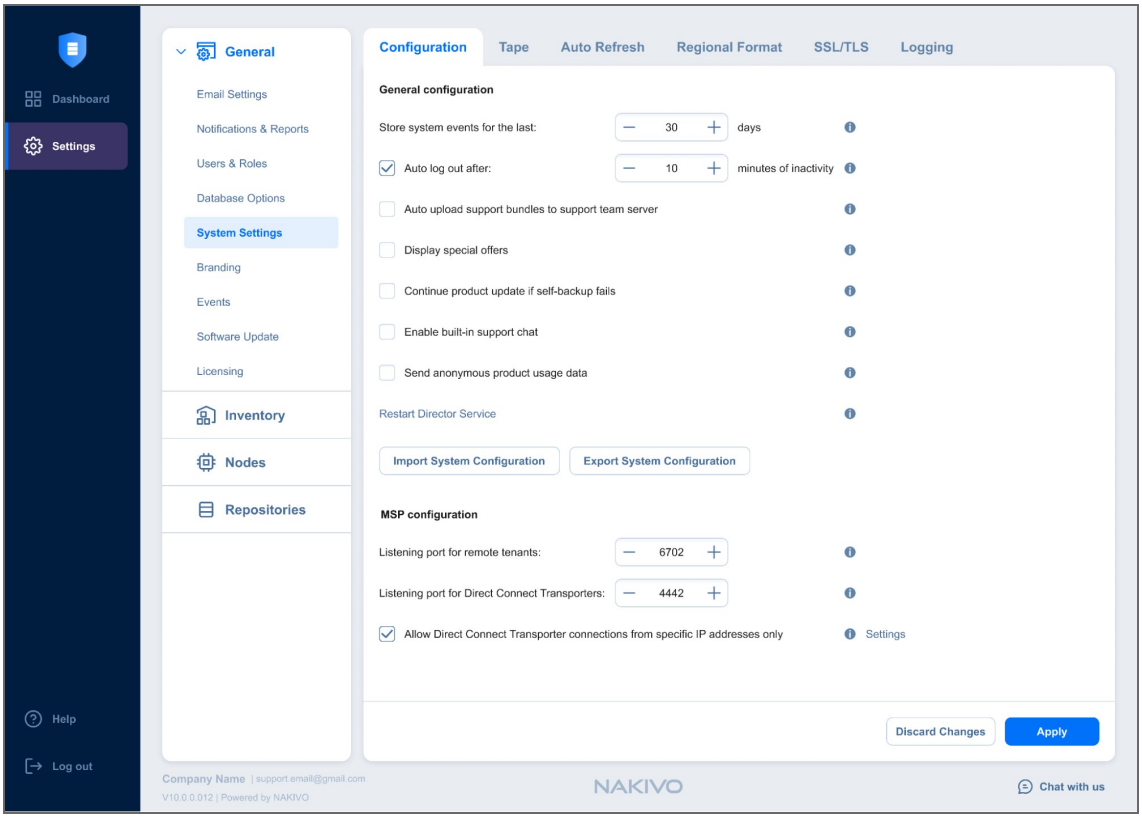
- **Auto upload support bundles to support team server**: When this option is selected, NAKIVO Backup & Replication automatically creates, encrypts, and uploads support bundles once a day to a NAKIVO support server. The NAKIVO Support Team may use this information to improve the product experience and to identify and resolve product issues faster.
- **Display special offers**: When this option is enabled, the NAKIVO special offers toolbar appears in the NAKIVO Backup & Replication interface.
- **Continue product update if self-backup fails**: When this option is selected, updates proceed even if self-backup cannot be performed.
- **Enable built-in support chat**: When this option is selected, you can contact a NAKIVO representative via chat in the NAKIVO Backup & Replication interface. When selected in the multi-tenant mode, the built-in support chat is available to all tenants of the NAKIVO Backup & Replication instance.
- **Enable Aptare Integration**: Select this option to integrate the APTARE storage resource management platform with NAKIVO Backup & Replication.
- **Send anonymous product usage data**: Enable this option to send anonymous product usage data to NAKIVO for efficient product development and enhancement. Note that no personal data is collected.
- Click the **Use New Scheduler** link to enable the use of a new scheduler that merges the retention and schedule steps. The scheduler allows you to set backup retention settings per schedule and get expiration dates for recovery points.
- You can click **Restart Director service** to stop all current activities and restart the **Director**. After clicking the link, a confirmation window appears. Click **Reboot** to confirm the restart.
- **Import System Configuration**: Find more information on the topic here.
- **Export System Configuration**: Find more information on the topic here.


Overview
Data Protection
Monitoring
Activities
Calendar
Search
Settings
Help
Logout
General
Email Settings
Notifications & Reports
Users & Roles
Self-Backup
Database Options
System Settings
Bandwidth Throttling
Branding
Events
Software Update
MSP
Licensing
Inventory
Nodes
Repositories
Tape
Configuration
Tape
Monitoring
Processing
Auto Refresh
Encryption
Regional Format
SSL/TLS
Logging
Store system events for the last: 30 days
Store job history for the last: 30 days
Auto log out after: 10 minutes of inactivity
Auto retry failed jobs: 3 times
Retry interval: 15 minutes
Retry critical errors
Auto upload support bundles to support team server
Display special offers
Continue product update if self-backup fails
Enable built-in support chat
Enable Aptare integration
Send anonymous product data usage
Restart Director service
Import System Configuration
Export System Configuration
Discard Changes
Apply

- You can also configure the following options in **Multi-tenant mode**:
  - **Listening port for remote tenants**:
    1. Specify the listening port for remote tenant connections in a multi-tenant environment (6702 is the default port). You can change this option only in the master tenant.
    2. Restart the **Director** to apply the changes.
  - **Listening port for Direct Connect Transporters**:
    1. Specify the listening port for a **Direct Connect Transporter**. The allowed range is 1 – 65,535 (4442 is the default port).

       > **Note**
       > The port must be different from the **MSP Director** port configured in the installer. You can change this option only in the master tenant.

    2. Restart the **Director** to apply the changes. For more information, see MSP Direct Connect.
  - **Allow Direct Connect Transporter connections from specific IP addresses only**: Select the checkbox to allow only whitelisted Direct Connect Transporters to connect to the MSP Director.
  - **Settings**: Click the link to open the Direct Connect Transporter IP whitelist dialog box.



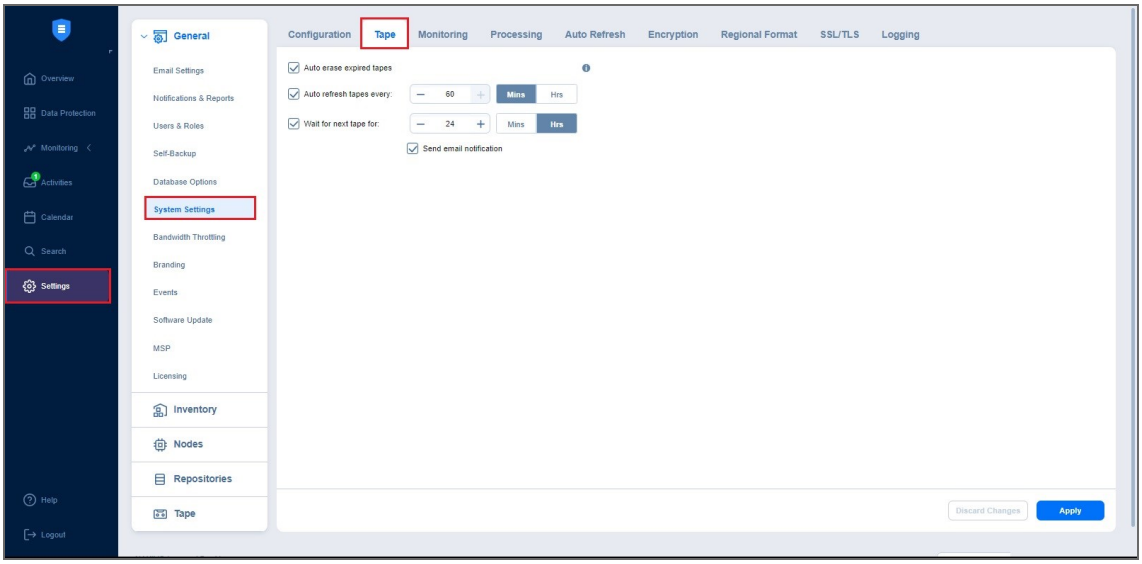
Configuring Tape Settings

In this section, you can control how NAKIVO Backup & Replication manages tape media throughout its lifecycle. Using these settings, you can automate the erasure of expired tapes, schedule how often tape contents are refreshed, and define how the system behaves when the required tape is not available. Set the following options:

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

> **Important**
>
> When this option is selected, the following prerequisites must be met for a cartridge to be erased:
>
> - All recovery points within the tape cartridge have expired.
> - There are no dependent recovery points on other tape cartridges.
> - The product keeps at least one full chain of recovery points.

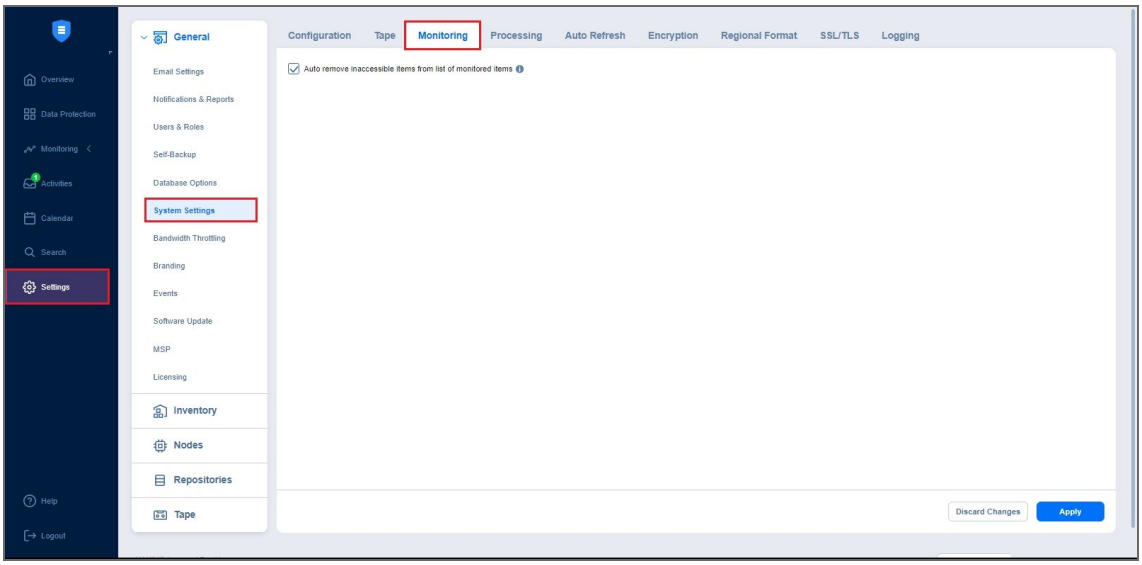
- **Auto refresh tapes every**: Select how often the contents of the tapes are refreshed in minutes or hours. Deselect if refreshing is not required.
- **Wait for next tape for**: Specify how long the system should wait for the next tape if there is no appropriate amount. Select the **Send email notification** checkbox to receive email notifications.



Configuring Monitoring Settings

In the **Monitoring** tab, you can control how NAKIVO Backup & Replication maintains and cleans up the list of monitored items. Set the following options:

- **Auto remove inaccessible items from list of monitored items**: When this option is selected, all inaccessible items are removed automatically from the list of monitored items.



.

Configuring Processing Settings

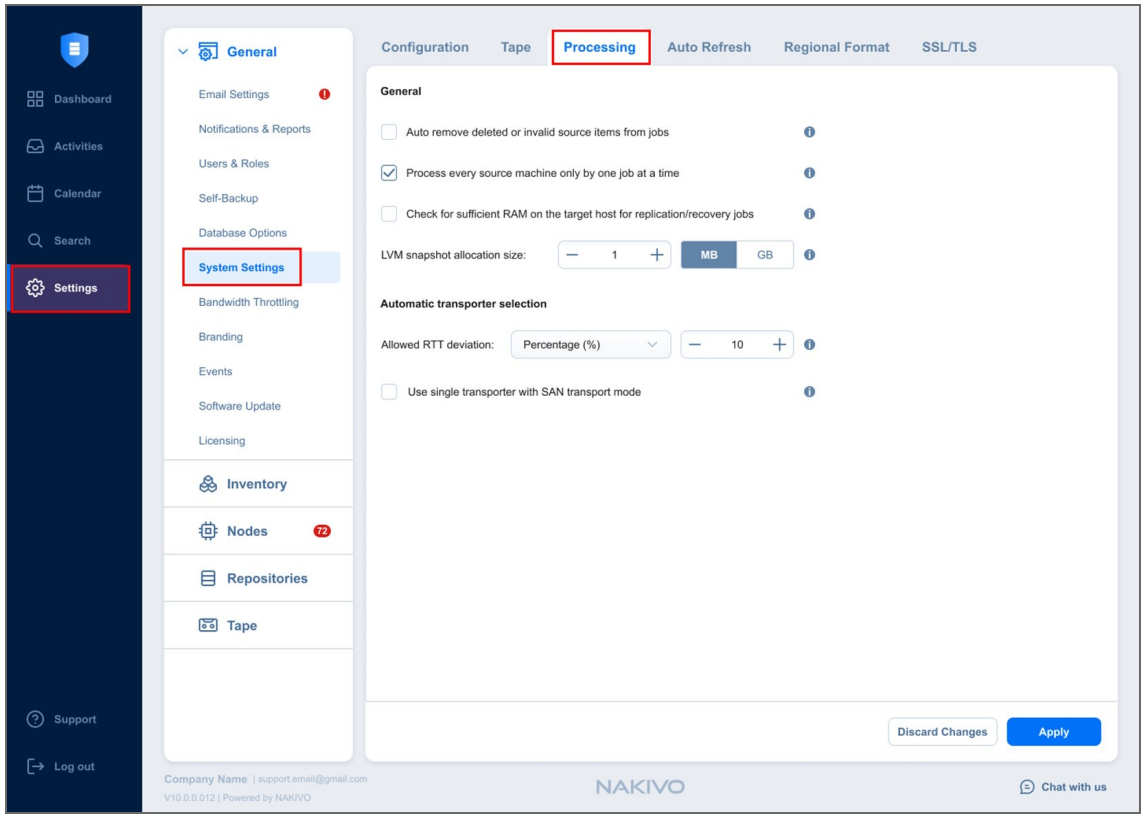
In this section, you can fine-tune how NAKIVO Backup & Replication processes source data during backup, replication, and recovery operations. Set the following options:

- **General**:
  - **Auto remove deleted or invalid source items from jobs**: This option applies to a protected container (such as a VMware cluster or EC2 region). When this option is selected, if during an inventory refresh, NAKIVO Backup & Replication discovers that at least one VM and/or EC2 instance is no longer available in the protected container, these VMs and EC2 instances are automatically removed from all jobs.
  - **Process every source item only by one job at a time**: When this option is selected, all machines in backup and replication jobs are processed by one job at a time only. Running jobs and respective source objects are not affected after changing this setting. This option is always enabled for physical servers and real-time replication jobs.
  - **Check for sufficient RAM on the target host for replication/recovery jobs**: When this option is deselected, NAKIVO Backup & Replication does not check whether the amount of RAM on the target host is sufficient for replication and recovery jobs.
  - **LVM snapshot allocation size**: This option allows you to set an LVM allocation snapshot size for a Linux physical server backup. The default size is 1 GB. The maximum size is 1000 GB.
- **Automatic transporter selection**: This section allows you to set the target transporter as a source transporter when finding the best data transfer path for a job object. To use a single transporter as a target and as a source, check the Use single transporter with SAN transport mode checkbox.

Notes

- This setting can be applied to the jobs with the Transporters option set to Automatic selection (Options step, DataTransfer section).
- This setting can be applied to jobs using the SAN transport mode.

- **Allowed RTT deviation**: This option allows you to define the acceptable range for selecting data paths during job execution. You can specify the deviation as a percentage (%) or a threshold (ms). This setting helps you optimize job performance and evenly distribute the workload across transporters. When a transporter is selected automatically, the data path with the best RTT is used during the job. If load balancing is needed during data transfer, data paths with RTT values within the specified range are chosen to ensure smooth operation.
- **Use a single transporter with SAN transport mode**: Select this checkbox to limit data transfer to a single transporter that is working as the target and source transporter when using SAN transport mode with automatic transporter selection enabled.

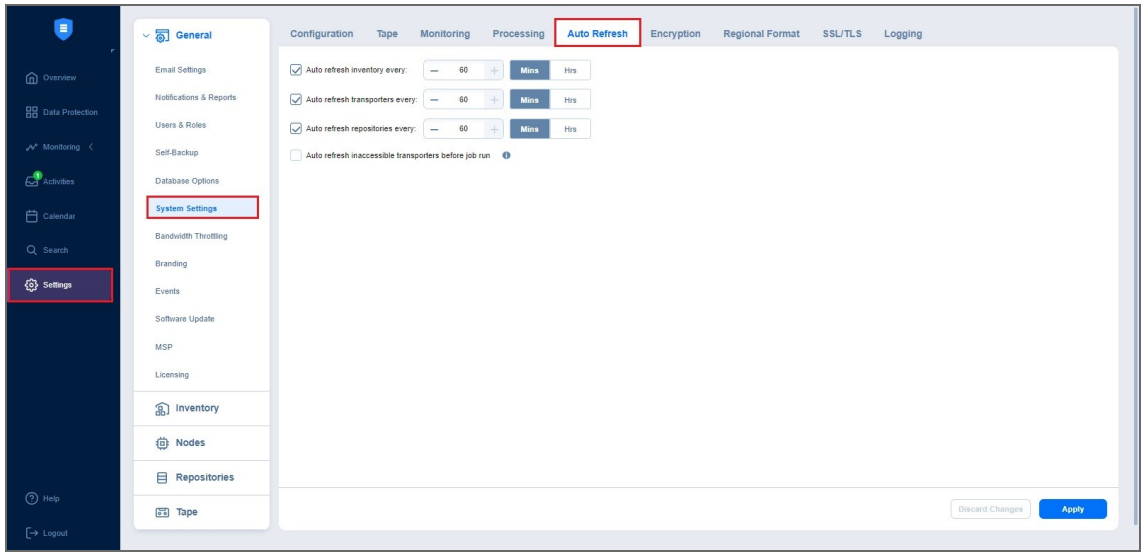


Configuring Auto Refresh Settings

In this section, you can control how often NAKIVO Backup & Replication automatically updates system inventories and infrastructure status. These options help keep system information accurate, reduce job failures caused by outdated data, and minimize the need for manual refresh operations.

Set the following options:

- **Auto refresh inventory every X minutes**: Specify how often you want your inventories to be refreshed.
- **Auto refresh transporters every X minutes**: Specify how often you want your transporters to be refreshed.
- **Auto refresh repositories every X minutes**: Specify how often you want your repositories to be refreshed.
- **Auto refresh inaccessible transporters before job run**: If this option is enabled, the product refreshes all inaccessible transporters before a job is processed.



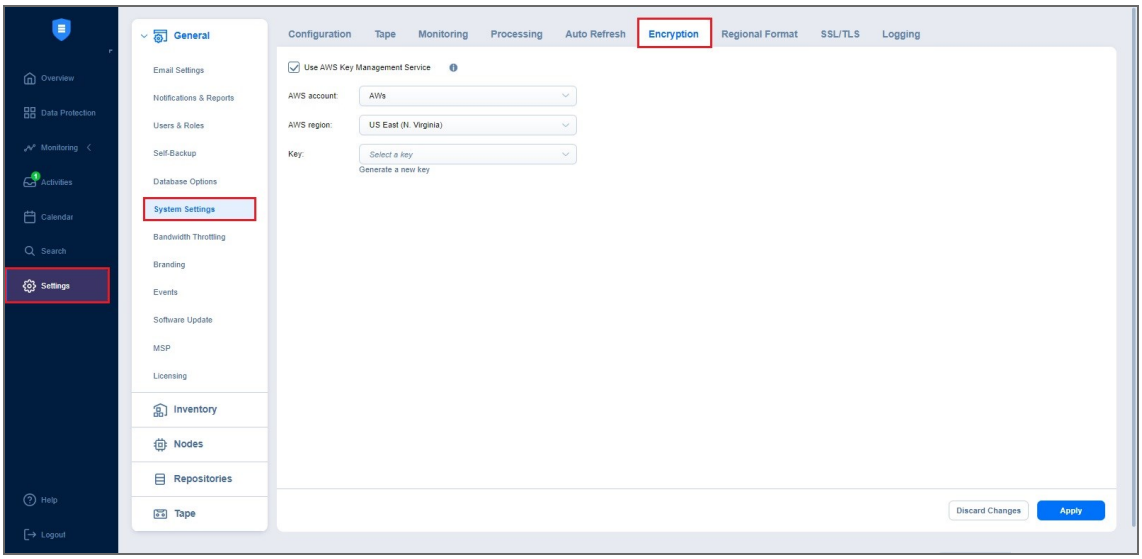
Configuring Encryption Settings

In this section, you can control how encryption keys and passwords are managed for protected data. These options allow you to strengthen security, centralize key management, and protect sensitive encryption data while respecting tenant isolation in multi-tenant deployments.

Set the following options:

- **Use (AWS) Key Management Service**: If enabled, all backup encryption passwords encrypted with the Key Management Service cryptographic key are available for recovery in case of product re-installation. You can select an existing symmetric cryptographic key or create a new one.

For more details, refer to Managing Backup Encryption.



Notes

- The **Encryption** tab is displayed at the Tenant level and in single tenant mode and is hidden at the Master level.
- The passwords are not propagated from one tenant to another.
- In case encryption is set for a job, the password hash is generated based on the configured password.
- The password hash gets encrypted with the KMS cryptographic key and is saved in the recovery point metadata.
- The AWS Key Management Service is not applied to self-backup and system configuration encryption. For more details, refer to Self-Backup Encryption and System Migration.

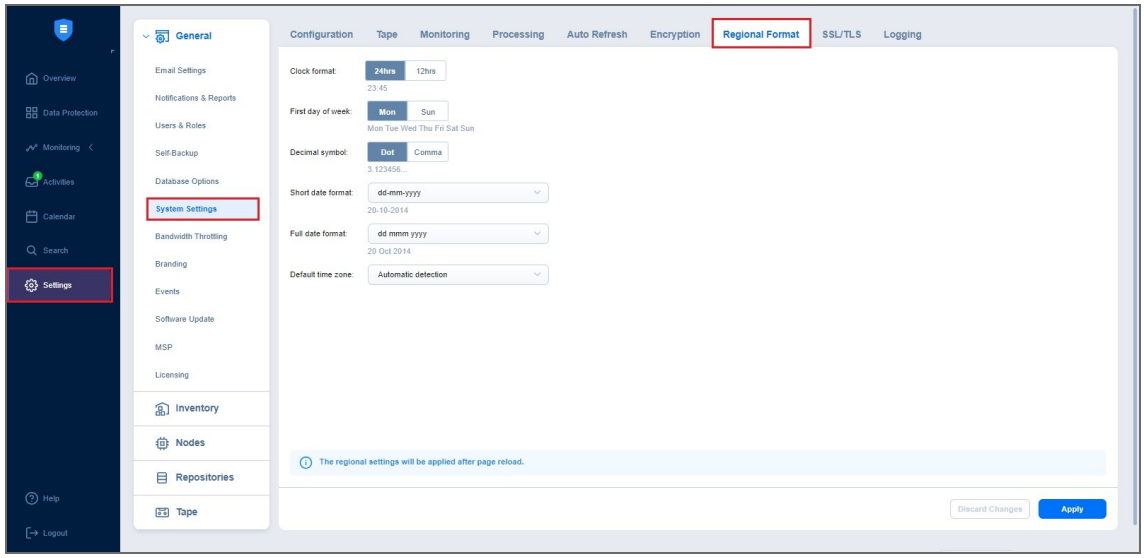
Configuring Regional Format Settings

In this section, you can define how date, time, and numeric values are displayed across the NAKIVO Backup & Replication interface. These options help you align the interface with local and regional conventions and ensure consistent interpretation of time-based data.

Set the following options:

- **Clock format**
- **First day of week**
- **Decimal symbol**
- **Short date format**

- **Full date format**
- **Default time zone**



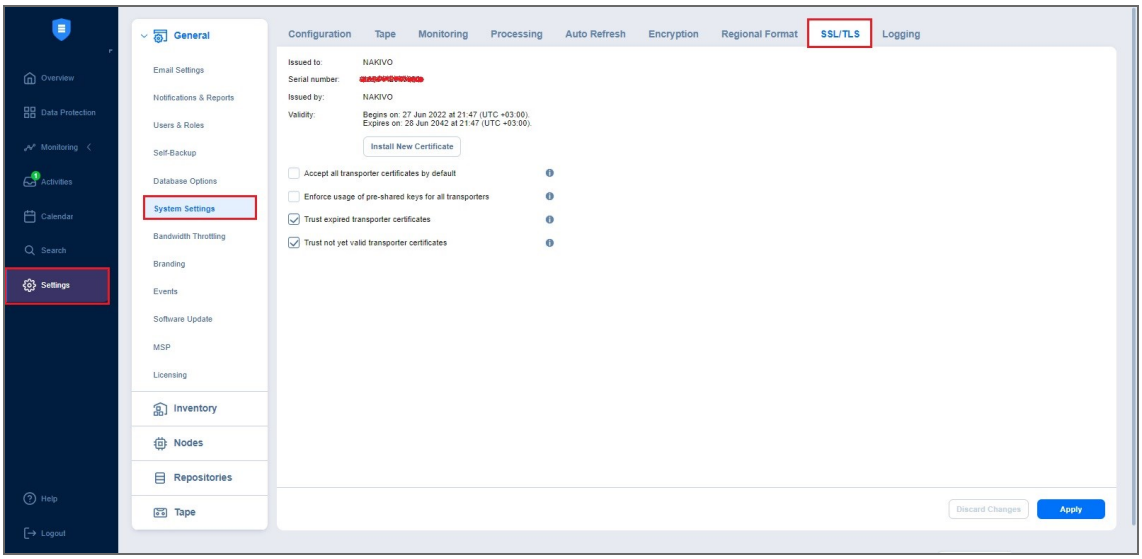
Configuring SSL/TLS Settings

With these settings, you can install and manage TLS/SSL certificates for the web interface, define how transporter certificates are validated, enforce the use of pre-shared keys, and control trust behavior for self-signed or expired certificates.

Set the following options:

- **Install new certificate**: A dialog opens allowing you to install a new TLS/SSL certificate for the NAKIVO Backup & Replication web interface. Certificates are generated either internally or through certification authorities. Proceed as follows to install a new certificate:
  - Click **Browse** and navigate to the location of either of the following certificate file types:
    - **Private key**: A file in the *.key format.
    - **Private key password (optional)**: A password for your private key.
    - **Certificate file**: A file in the *.pem, *.crt, *.cer, *.p7b, or *.p7s format.
    - **Intermediate certificate (optional)**: A file in one of the following formats: *.pem, *.crt, *.cer, *.p7b, *.p7s.
- **Accept all transporter certificates by default**: Select this option to automatically accept all transporter certificates. After selecting the option, click **Continue** in the warning popup window that appears to confirm the selection.
- **Enforce usage of pre-shared keys for all transporters**: Selecting this option makes sure that transport function only when pre-shared key is installed.
- **Trust expired self-signed transporter certificates**: Selecting this option makes the solution trust

the expired self-signed transporter certificates.



Notes

- Repositories assigned to the inaccessible transporters are refreshed as well during the auto refresh.
- The **Auto refresh inaccessible transporters before job run** option is supported only for the following platforms:
  - VMware vSphere
  - AWS
  - Cloud Director
  - Nutanix AHV
- If any time zone other than **(UTC+00:00, UTC) Coordinated Universal Time** is chosen, daylight savings times are honored.

System Migration

NAKIVO Backup & Replication provides you with the ability to migrate all your settings (including inventory, jobs, credentials, transporter settings, and so on) to a new instance (copy) of the product.

Important

System configuration export and import are designed for migration purposes only, and not to serve as a system configuration backup. After you have exported system configuration from an old instance of the product, do not run jobs in that old instance. Doing so results in failed jobs in the new instance after the migration. All jobs have to be recreated, and full initial job run is required.

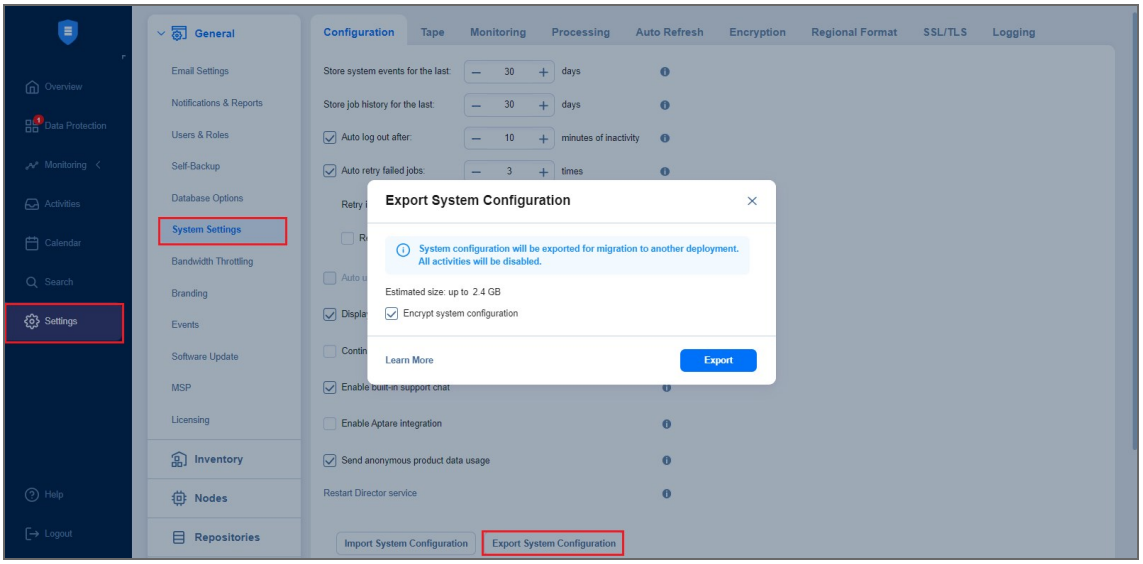
See the topics below for more information:

- Exporting System Configuration
- Importing System Configuration

Exporting System Configuration

To export system configuration from the old deployment, follow the steps below:

1. Click **Settings** in the left pane of the product.
2. Select **System Settings** tab in the **General** section.
3. On the **Configuration** tab, click **Export System Configuration**.
4. In the dialog window that appears, click **Export**.



Notes

- Select the **Include indexing data** checkbox to include the indexing data into the package.
- Selecting the **Include indexing data** checkbox displays the **Estimated index data size** label and includes the size of index data in the displayed data size.
- Optionally, select the **Encrypt system configuration** checkbox.

5. Click **Export** to confirm the operation.
6. If the **Encrypt system configuration** checkbox is selected, clicking the **Export** button opens the **Set a password** dialog box. For more details, refer to Setting Password.

Notes

- All activities in the old instance (such as jobs and recovery sessions) are automatically stopped and all jobs are disabled.
- AWS Key Management Service is not applied to system configuration encryption.

7. Wait until the export is completed and download the export bundle.
8. Do not run jobs in the old instance.

Importing System Configuration

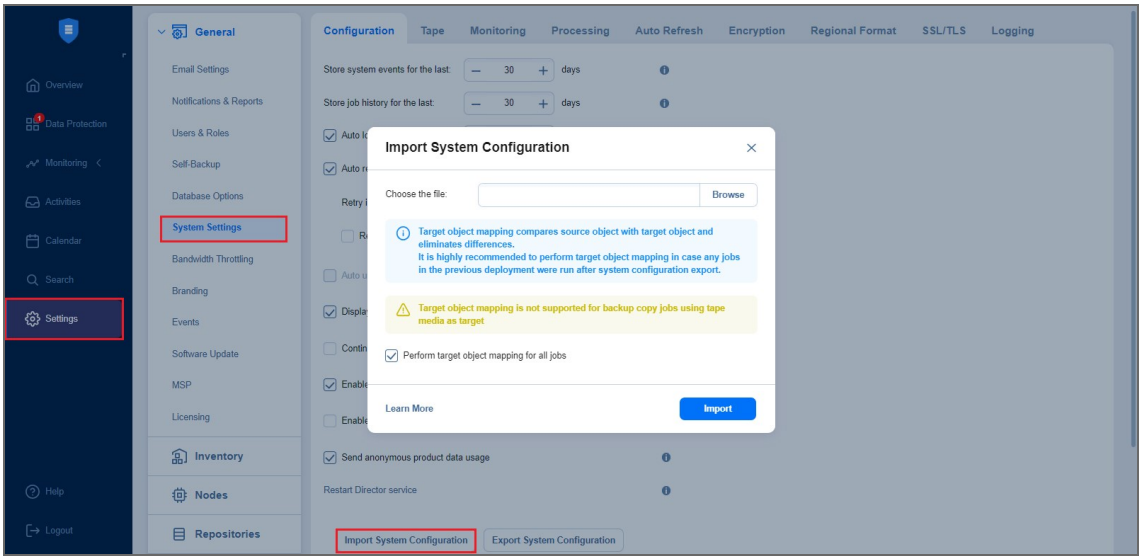
To import system configuration into a new instance of the product, follow the steps below:

1. Click **Settings** in the left pane of the product.
2. Select **System Settings** tab in the **General** section.
3. On the **Configuration** tab, click **Import System Configuration**.
4. In the dialog box that appears, locate the system configuration bundle using the **Browse** button.
5. Click **Import**.

Note

If the configuration bundle is encrypted, clicking on the **Import** button opens the **Provide a password** dialog box. In the dialog box that appears, select the needed password or create a new one. For more information, refer to Managing Passwords.

6. Click the **Cancel** or **X** button to close the dialog box without applying any changes.



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

Important

- If there is any existing data in the new instance, it will be overwritten with the import operation.
- Target object mapping is skipped for backup copy jobs using tape media as target.
- If a physical configuration of your source deployment differs from a target deployment, a **Backup Repository** may become inaccessible after the bundle import is completed.

8. Wait until the import is completed, and close the dialog window.

Notes

- **Backup Repositories** are not migrated by the system configuration export and import. If you have a local **Backup Repository** on the old instance of the product, you may want to move it to the new location. After moving the **Backup Repository**, you may need to edit **Backup Repository** settings in the new instance so that the new settings refer to the actual **Backup Repository** location.
- In case a custom TLS/SSL certificate of the Web server was used in the old instance, a manual service restart is required in the new deployment.

Users and Roles

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

- Managing Users and Roles
- Navigating Users View
- Navigating Roles View
- "Navigating AD Domains View" on page 350
- Navigating AD Groups View

Managing Users and Roles

Managing users and roles can be done by following these steps:

1. Log in to NAKIVO Backup & Replication.
2. Click **Settings** (cog icon) in the left pane of the product.
3. Go to the **General** tab and click **Users & Roles**.

Navigating Users View

To see the list of all users, select the **Users** view in the upper pane. On this page of the solution you can do the following:

- See the list of all users added to NAKIVO Backup & Replication.
- Sort the list by **User Name**, **Role**, **Tenant**, **Domain**, **Access level**, **Group** or **2 Factor-authentication** by clicking on the respective name of the column.
- Search for users by entering the name of the user (fully or partially) into the **Search** bar.

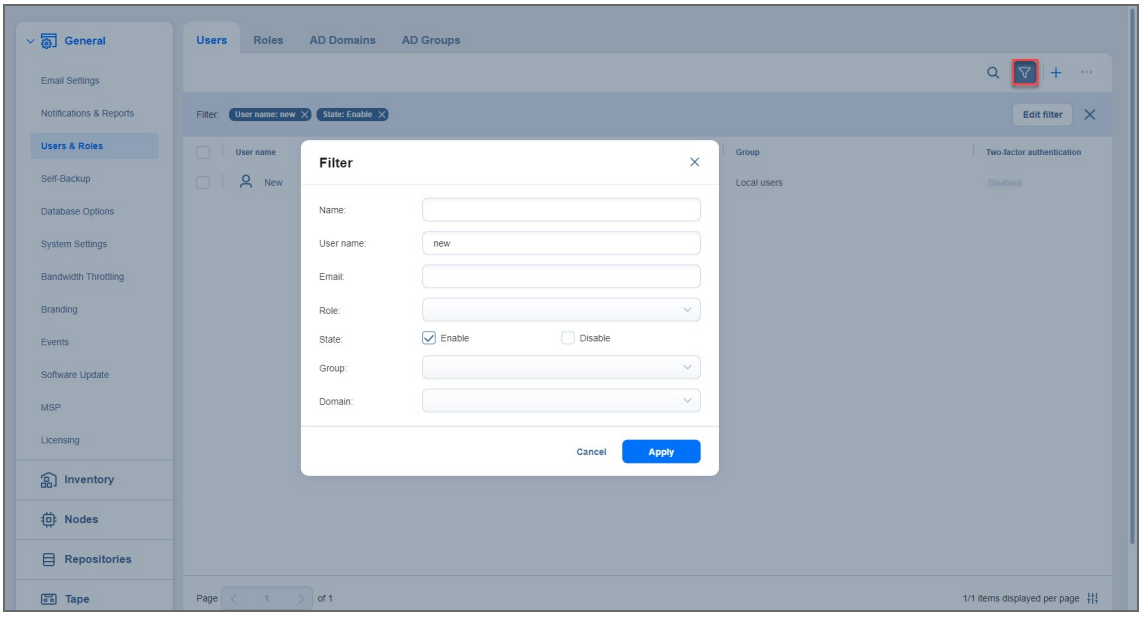
  Notes

  - The search phrase must be contained in the user registration name, username, or email address.
  - The **Search** option allows you to perform search on the following columns:
    - **User name**
    - **Role**
    - **2 Factor-authentication**
    - **Access level**
    - **Group**

- Filter the list of users by selecting the **Filter** option.

Note

Clicking **Filter** opens a new window that allows you to filter the list of local users according to **User name**, **Role**, **State**, **Group**, and **Domain**.



- Add a new local user by clicking the **+** icon.
- Delete, disable, enable two-factor authentication, and assign a new role to the local user individually. These actions, except **Edit**, can also be done in bulk by checking the box in the upper left pane to select all users and clicking the **...** (ellipsis) button.

Notes

- When selecting all local users to apply a bulk action, NAKIVO Backup & Replication selects only those users that are displayed on the screen.
- When a user is deleted, the name of the user becomes *Deleted* in places where this user is referenced to.
- You can create another user with the same username as the deleted user.

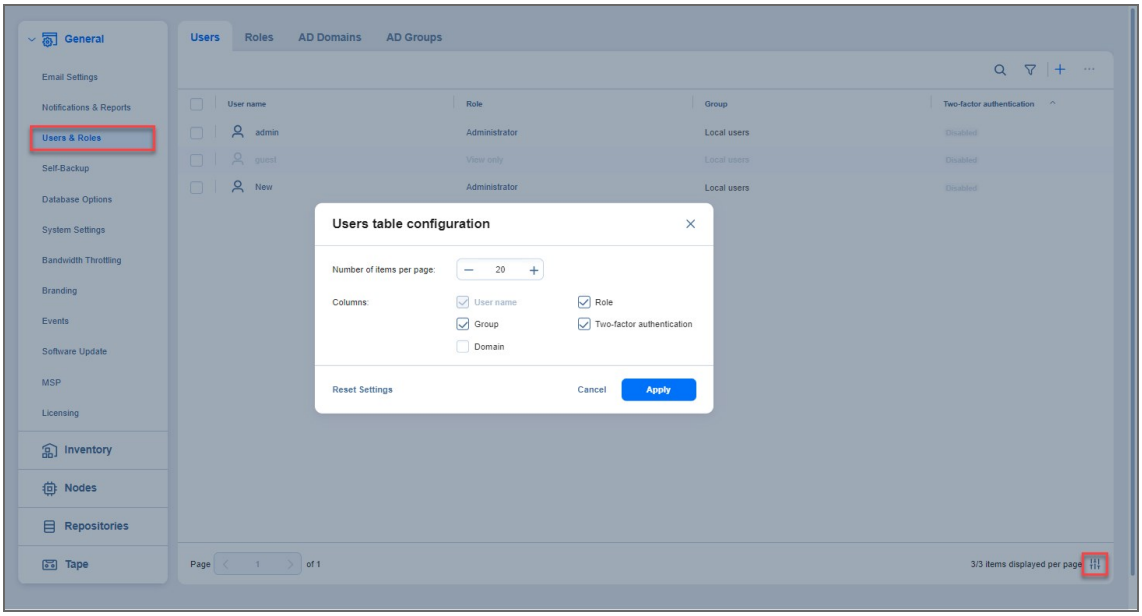
- Edit the role assigned to the local use by clicking on the name of the role in the respective column.

Configuring Users Table

You can show/hide columns or modify the number of items per page in the **Users** table. In the lower right corner, click the controls icon. In the dialog window that opens, select/deselect checkboxes and click **Apply**. The following options are available:

- User name
- Group
- Domain

- Role
- Two-factor authentication



Note

You cannot modify the **User name** column in this dialog window. Refer to "Editing Local User" on page 368 to learn more about editing users.

Navigating Roles View

To see the list of all local users, select the **Roles** view in the upper pane. On this page of the solution you can do the following:

- See the list of all user roles added to NAKIVO Backup & Replication.
- Sort the list by **Role name**, **Access level**, or **Number of users** by clicking on the respective name of the column.

  Note

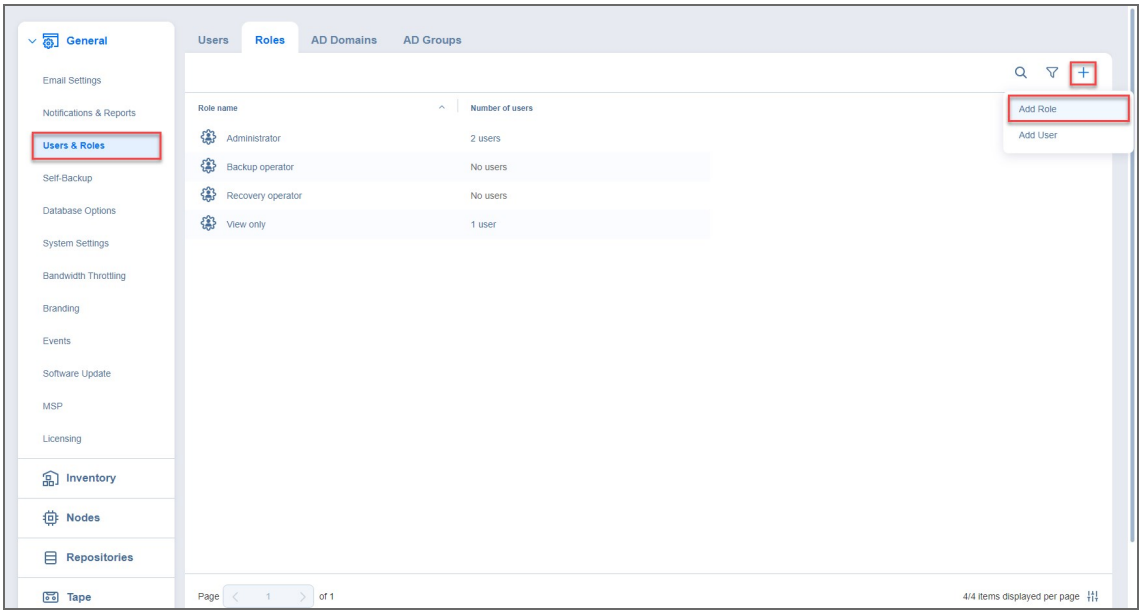
  The **Access level** column is displayed only for the **Master tenant** in Multi-tenant mode. It displays the access level that the role has.

- Filter the list of users by entering the name of the user fully or partially into the **Search** bar or by selecting the **Filter** option.

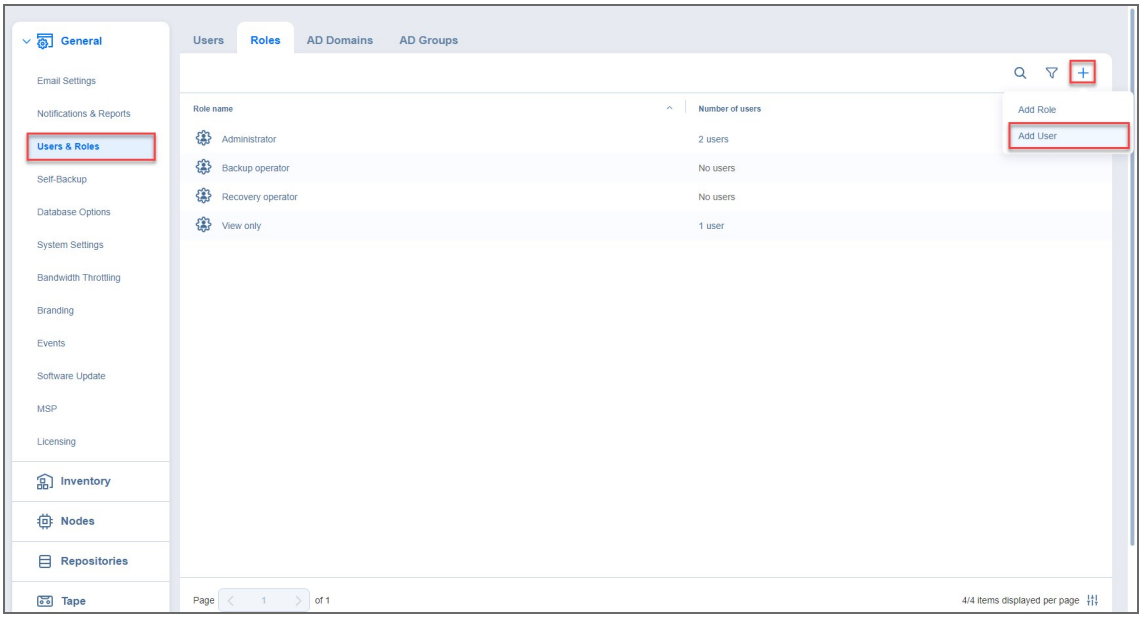
  Note

  Clicking **Filter** opens a new window that allows you to filter the list of local users according to **Role name** or **Number of users**.

- Add a new role by clicking the **+** icon and selecting **Add Role**.



- Add a new local user by clicking the **+** icon and selecting **Add User**.

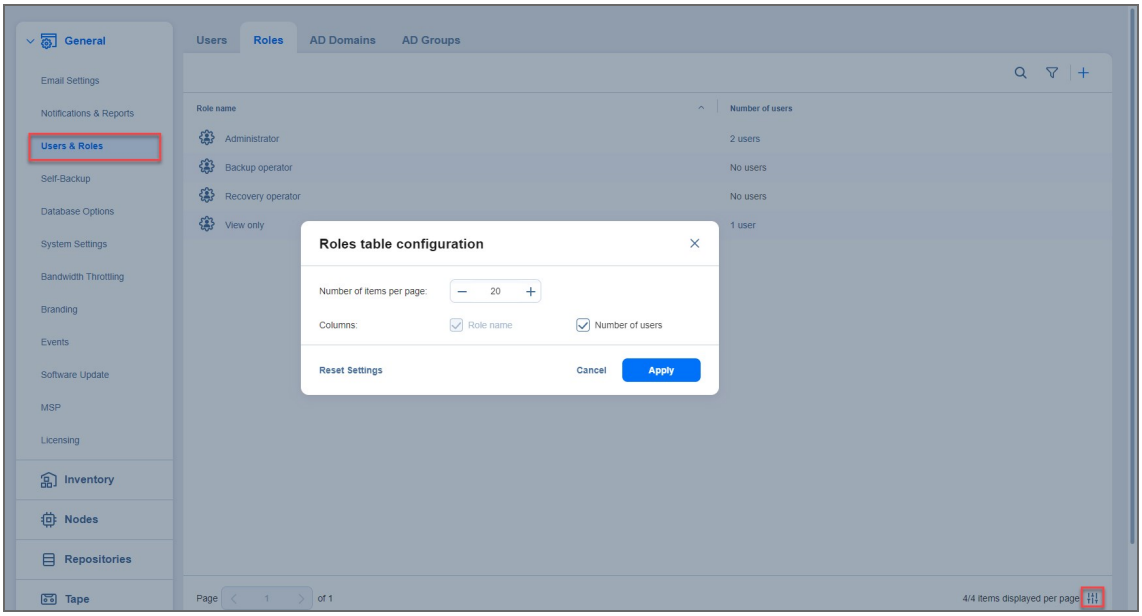


- Edit, delete, or clone the user roles individually. These actions, except **Edit**, can also be done in bulk by checking the box in the upper left pane to select all users and clicking the **...** (ellipsis) button. When selecting all local users to apply a bulk action, NAKIVO Backup & Replication selects only those users that are displayed on the screen.
- **Edit**, **Delete**, **Clone** the role by clicking the ellipses to the right of the role's name.

Configuring Roles Table

You can show/hide columns or modify the number of items per page in the **Roles** table. In the lower right corner, click the controls icon. In the dialog window that opens, select/deselect checkboxes and click **Apply**. The following options are available:

- Role name
- Number of users



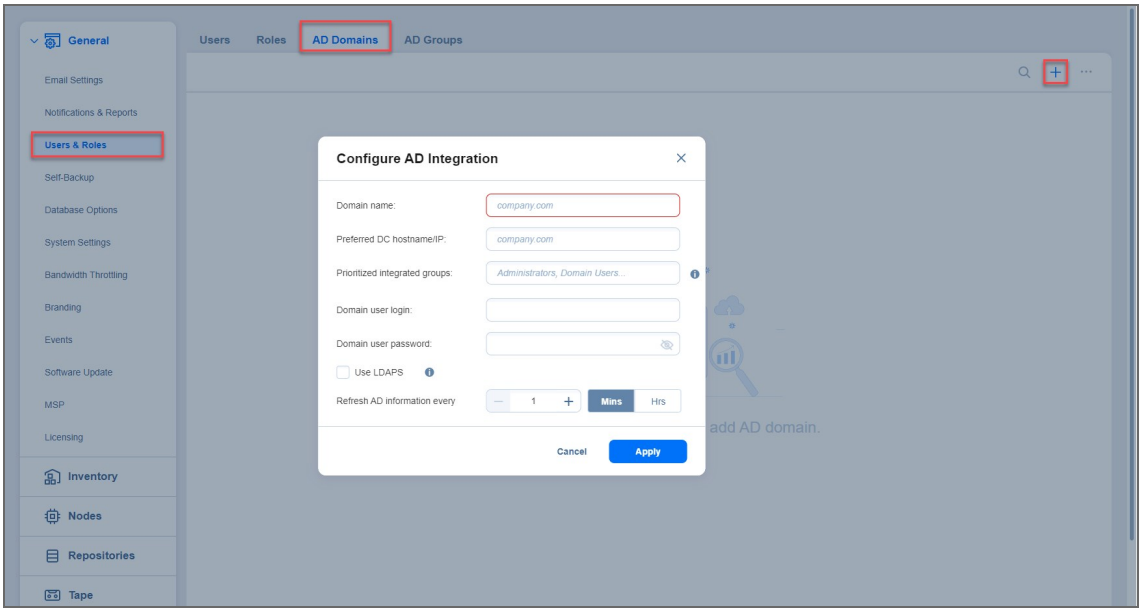
Note

You cannot modify the **Role name** column in this dialog window. Refer to "Editing User Role" on page 375 to learn more about editing alert user roles.

Navigating AD Domains View

To see the list of all Active Directory domains, select the **AD Domains** view in the upper pane. On this page of the solution you can do the following:

- See the list of all AD domains added to NAKIVO Backup & Replication.
- Sort the list by **Domain name**, **DC/Hostname IP**, **Prioritized integrated groups**, **Domain user**, **Refresh AD** or **Status** by clicking on the respective name of the column.
- Search the list of AD domains by entering the name of the user fully or partially into the **Search** bar.
- Add a new AD domain by clicking the **+** icon.



Note

You can add and then manage two or more AD domains and add AD groups from multiple domains. Refer to Configuring Active Directory Integration for more details.

- Edit, delete, disable/enable AD domains. These actions, except **Edit**, can also be done in bulk by checking the box in the upper left pane and clicking the **...** (ellipsis) button.

Note

The user associated with the disabled AD domain cannot log in to the solution.

Navigating AD Groups View

To see the list of all Active Directory groups, select the **AD Groups** view in the upper pane.

On this page of the solution you can do the following:

- See the list of all AD groups added to NAKIVO Backup & Replication.
- Sort the list by **Group name**, **Logged in users**, **Access level**, or **Role** by clicking on the respective name of the column.

  Note

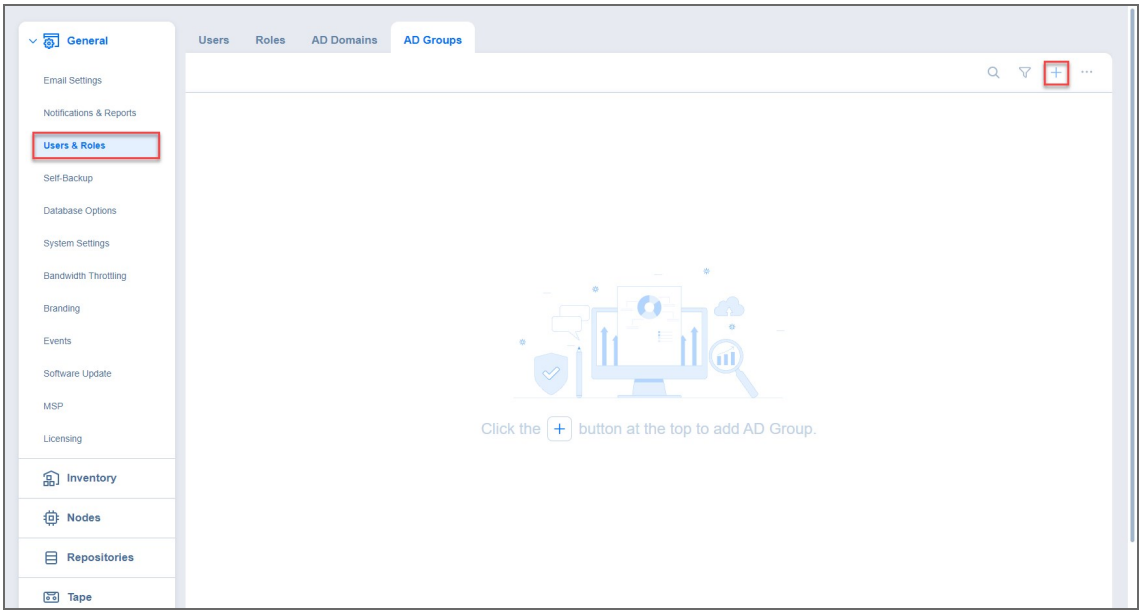
  The **Access level** column is displayed only for **Master tenant** in Multi-tenant mode. It displays the access level assigned to the AD group.

- Filter the list of users by entering the name of the user (fully or partially) into the **Search** bar or by selecting the **Filter** option.

Note

Clicking **Filter** opens a new window that allows you to filter the list of local users according to **Group name**, **Role**, **Number of users**, and **Status**.

- Add a new AD group by clicking the **+** icon.



Notes

- You can add AD groups from multiple domains.
- The AD groups list is disabled if the domain is disabled.

- Edit, delete, disable, enable two-factor authentication, and assign a new role to the local user individually. These actions, except **Edit**, can also be done in bulk by checking the box in the upper left pane to select all users and clicking the **...** (ellipsis) button.

Note

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

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

For details, refer to the following sections:

- “Managing Active Directory Users” on page 353
- “Managing Local Users” on page 361
- “Managing User Roles” on page 370
- “Configuring Two-Factor Authentication” on page 380
- “Managing Application Passwords” on page 378

Managing Active Directory Users

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

- "Adding Active Directory User" on page 354
- "Assigning Role to Active Directory User" on page 357
- "Configuring Active Directory Integration" on page 358
- "Disabling Active Directory User" on page 360
- "Editing Active Directory User" on page 361

Adding Active Directory User

After configuring AD integration in the **Active Directory Configuration** wizard, you can proceed with adding AD user(s). Alternatively, switch to **AD Groups** tab and then click on the "**+**" icon.

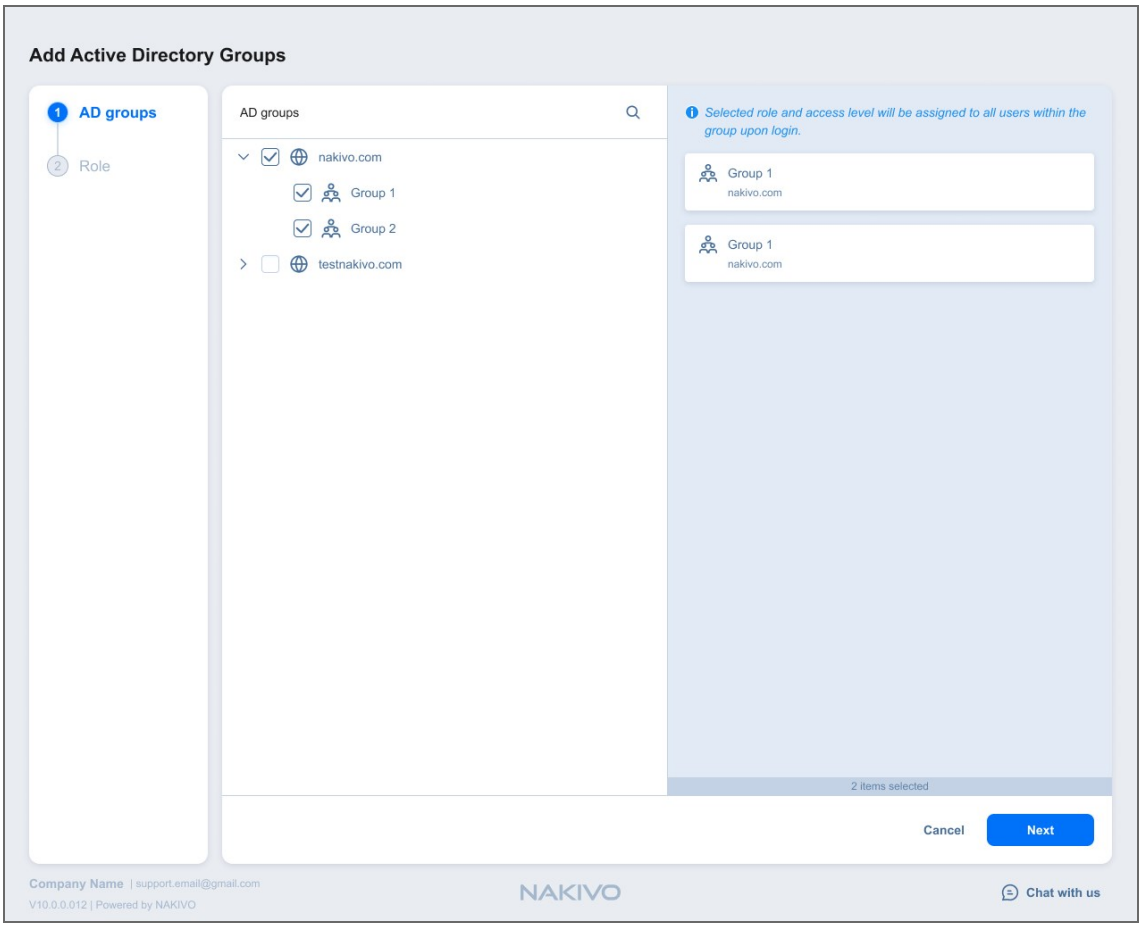
Proceed as follows:

1. Optionally, you can filter the tree of Active Directory users by entering a string to the **Search** box. You can enter a section or the whole name of the item.
2. Select **Active Directory users and groups** by placing a checkmark to their left.
3. The selected items appear in the right pane of the page. If necessary, reorder the selected items by dragging them to a new position. By doing so, you can specify to add the most important users and groups first.

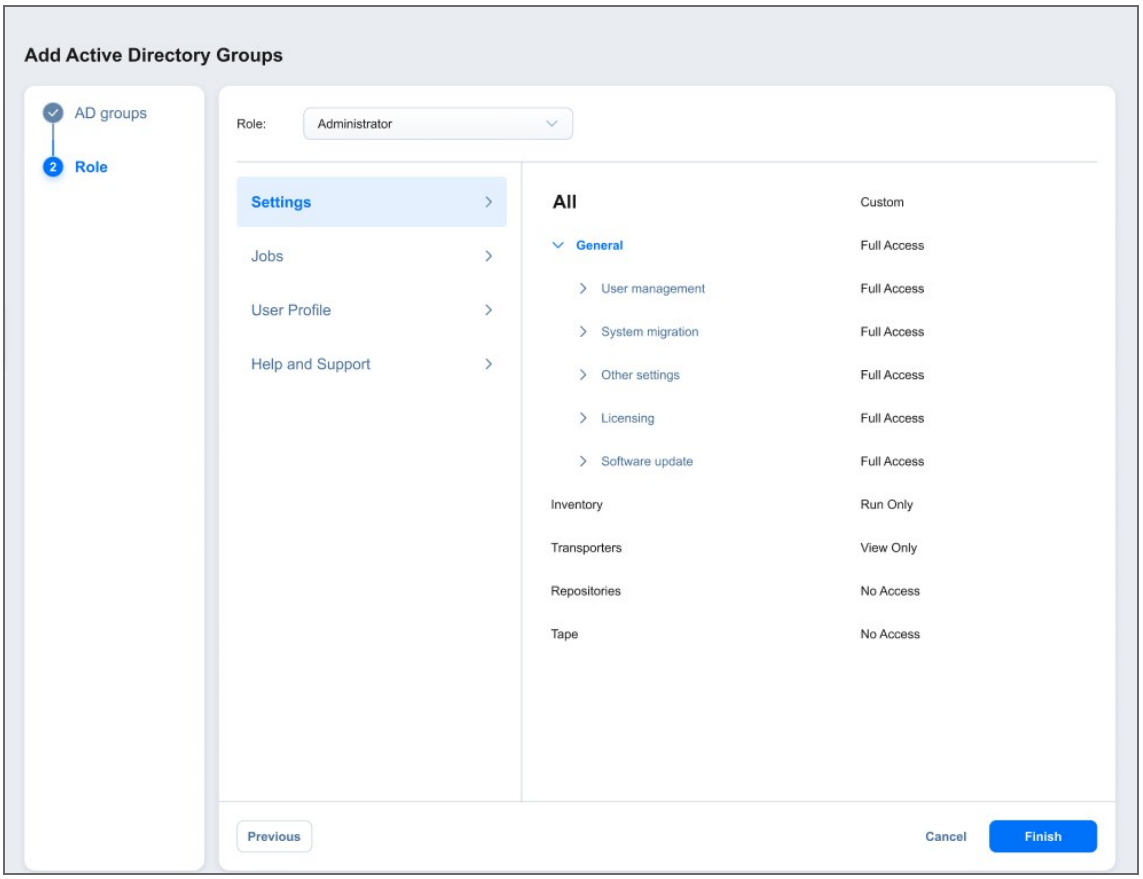
   > **Note**
   > Only logged in users that belong to the group can be added.

4. Review the list of selected items. If necessary, remove a selected user or group from the list in either of the following ways:
   - Deselect the item in the left pane. This will remove the item from the right pane.
   - In the right pane, hover the pointer over the item you wish to remove and click the **Remove**

button. This will deselect the item in the left pane.



5. Click **Next** to proceed to the **Role** tab.
6. On the **Role** tab, choose a user role to be assigned to the users.



7. In the lower right corner of the page, click **Finish**. The selected AD groups with selected access level and role appear in the **AD Groups** tab.

Assigning Role to Active Directory User

Follow the steps below to assign a role to an Active Directory user:

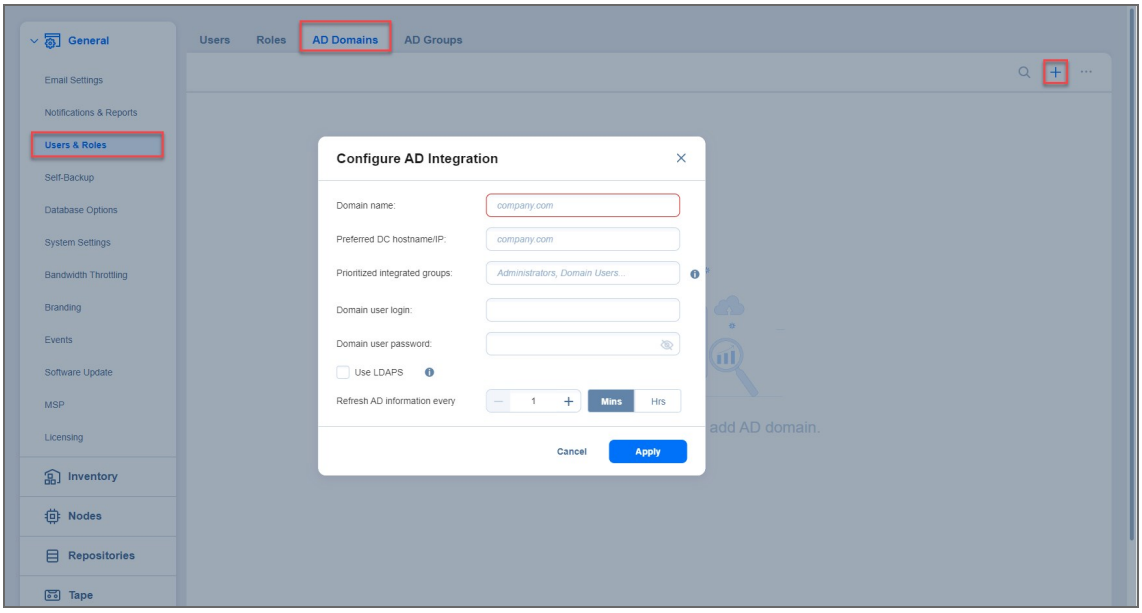
1. Go to **Settings > General > Users & Roles**.
2. The **Users & Roles** page opens in the **Users** view. Hover over the Active Directory user, and then click the **…** (ellipsis) button in the rightmost column of the row.
3. In the menu that opens, click **Assign role**.
4. In the dialog box that opens, select a new user role from the **Role** list and then click **Save**.

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

Configuring Active Directory Integration

To configure Active Directory integration, follow these steps:

1. Go to **Settings > General > Users & Roles**.
2. Select the **AD Domains** tab and then click the "**+**" icon.



3. The **Active Directory Configuration Wizard** opens. Proceed as follows:
   a. In the **Domain name** box, enter the domain name.
   b. In the **Preferred DC hostname/IP** box, enter the name of the preferred domain controller or its IP address.
   c. Optionally, you can enter the name of the preferred Active Directory groups in the **Prioritized integrated groups** box.

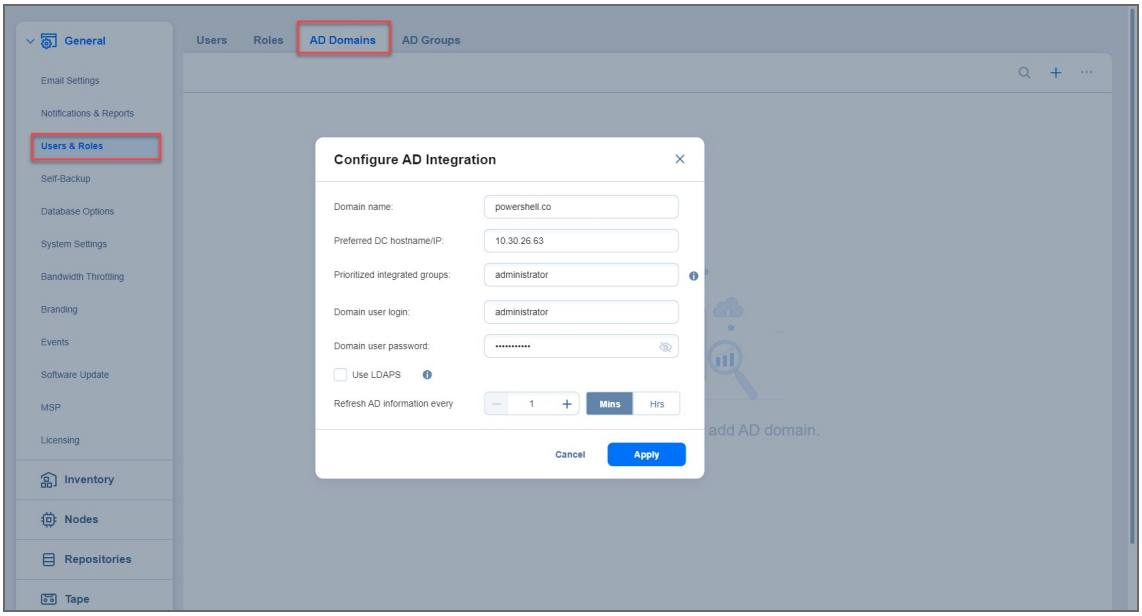
> **Note**
> If a user is a member of two or more Active Directory groups, enter the prioritized group's name in this field.

   d. In the **Domain user login** box, enter the username that will be applied when integrating Active Directory.
   e. In the **Domain user password** box, enter the user password that will be applied when integrating Active Directory.
   f. Optionally, enable **Use LDAPS** option. If checked, port *636* is used for LDAP (Lightweight Directory Access Protocol) over SSL.
   g. **Refresh AD information every**: Specify a periodicity of refreshing Active Directory information.

h. In case Active Directory integration was successfully completed before, you can optionally click **Remove AD Integration** to cancel the AD integration.

> **Note**
>
> The **Remove AD Integration** option is disabled if AD integration is not configured.

i. Click **Apply** after you're done.



j. On the **Users** page of the wizard, proceed with adding an Active Directory user.

Notes

- You can add up to 10 AD domains and up to 5.000 AD groups to your solution. Refer to "Navigating AD Domains View" on page 350 and "Navigating AD Groups View" on page 351 for more details about how to work with AD domains and groups.
- A new domain must be assigned a unique name.
- A new AD domain will be added to the list of Active Directory domains and displayed for the corresponding user on the **Users** tab.
- After a new AD domain is added, new AD groups that belong to this domain are created.
- After the existing AD domain is edited/changed, the AD groups that belong to this domain are edited/changed accordingly, for example:
  - If the AD groups are no longer associated with any AD domain, they are deleted.
  - After the items belonging to the AD groups marked as deleted are removed from the product (for example, by another user), these AD groups are deleted.

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

Disabling Active Directory User

Follow the steps below to disable an Active Directory user:

1. Go to **Settings > General > Users & Roles**.
2. The **Users & Roles** page opens in the **Users** view. Hover over the Active Directory user you want to disable, and then click the **…** (ellipsis) button in the rightmost column of the row.
3. In the menu that opens, click **Disable**.
4. In the dialog box that opens, click **Disable** to confirm that you want to disable the Active Directory user.

The Active Directory user appears dimmed in the list of users.

> **Note**
> You cannot enable the AD user if its parent AD group is disabled.

Editing Active Directory User

Follow the steps below to edit an Active Directory user:

1. Go to **Settings > General > Users & Roles**.

   The **Users & Roles** page opens in the **Users** view. In the list of users, do either of the following:

   a. Locate the Active Directory user and click its name.

   b. Hover over the Active Directory user, click the **…** (ellipsis) button in the rightmost column of the row.

   c. Click **Edit**.

2. The **Edit Active Directory User** page opens. Edit the Active Directory user properties if necessary:

   a. In the **General** tab, edit the user.

   b. In the **Role** tab, edit the user role.

   c. Click **Save** to save your modifications to the Active Directory user.

Managing Local Users

With NAKIVO Backup & Replication, you can freely add, edit, disable, delete local users, or assign a role to them. For details, refer to the topics below:

- "Adding Local User" on page 362
- "Assigning Role to Local User" on page 365
- "Deleting Local User" on page 366
- "Disabling Local User" on page 367
- "Editing Local User" on page 368

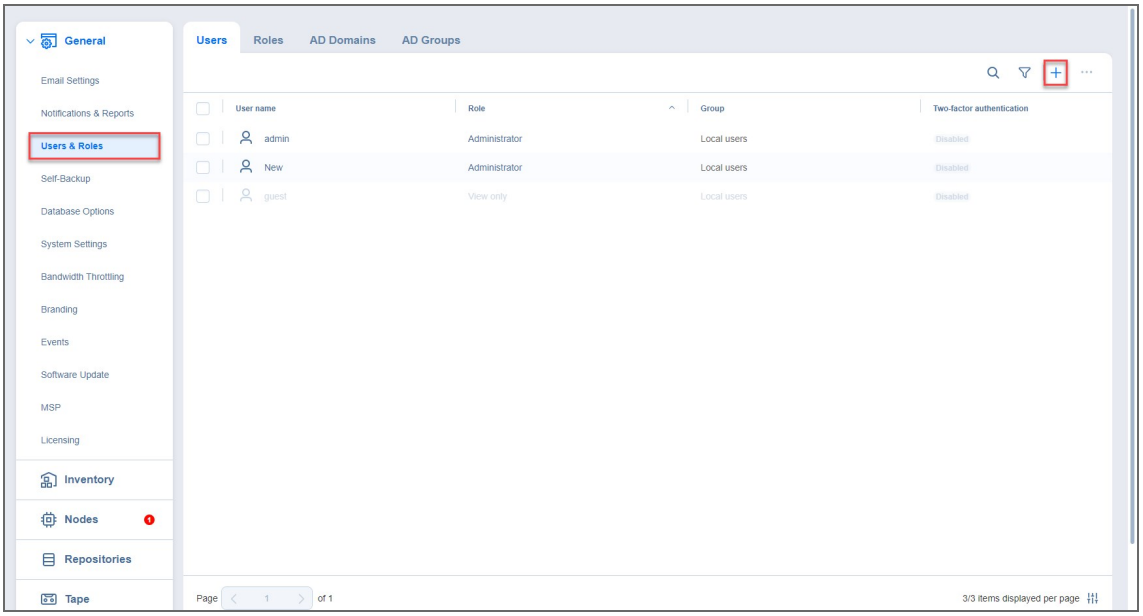
The application has the following built-in local users:

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

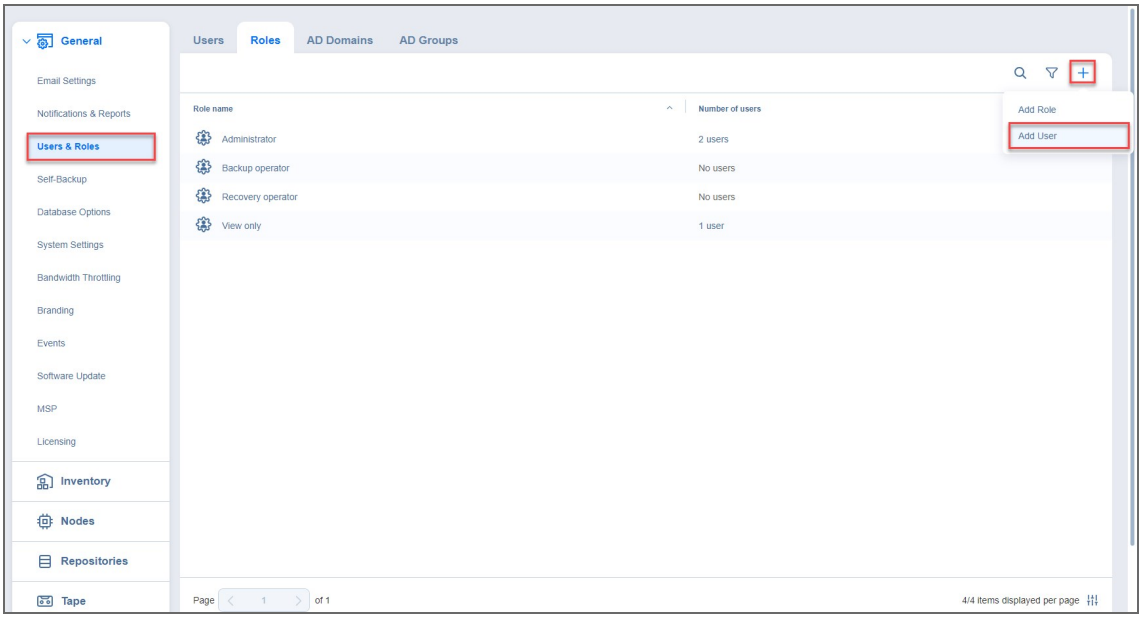
Adding Local User

Follow the steps below to add a local user:

1. Go to **Settings > General > Users & Roles**
2. The **Users & Roles** page opens on the **Users** tab.
3. Click the **+** icon.

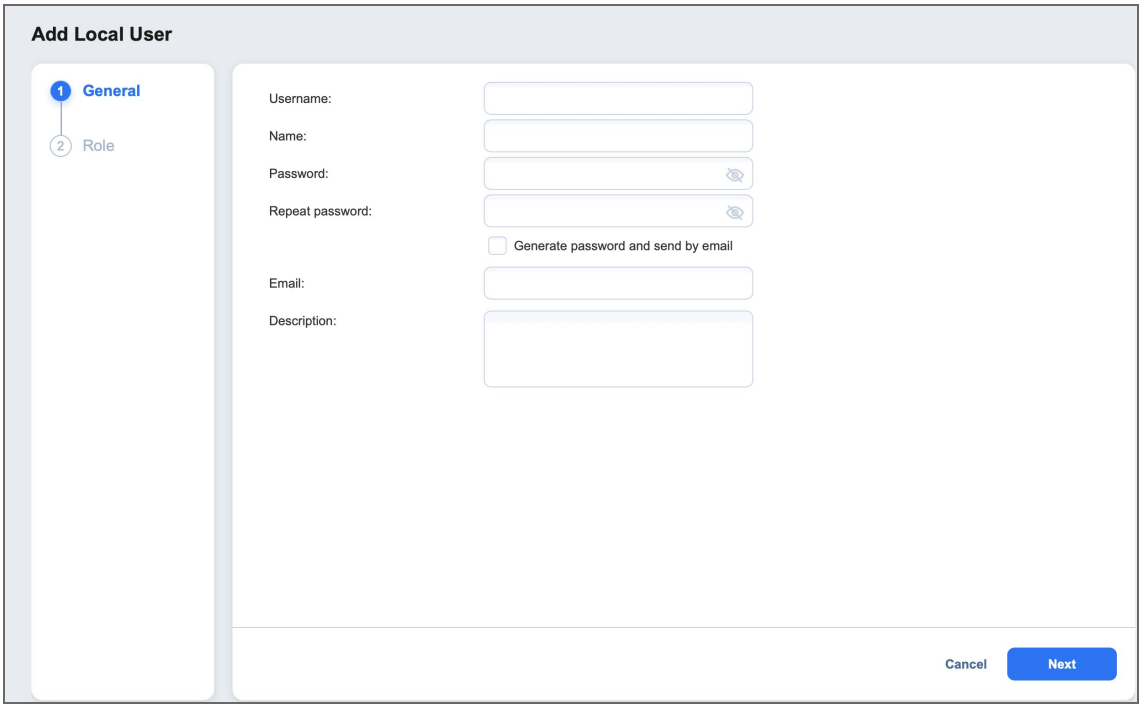


Alternatively, you can switch to the **Roles** tab, click on the "+" icon and select **Add User**.

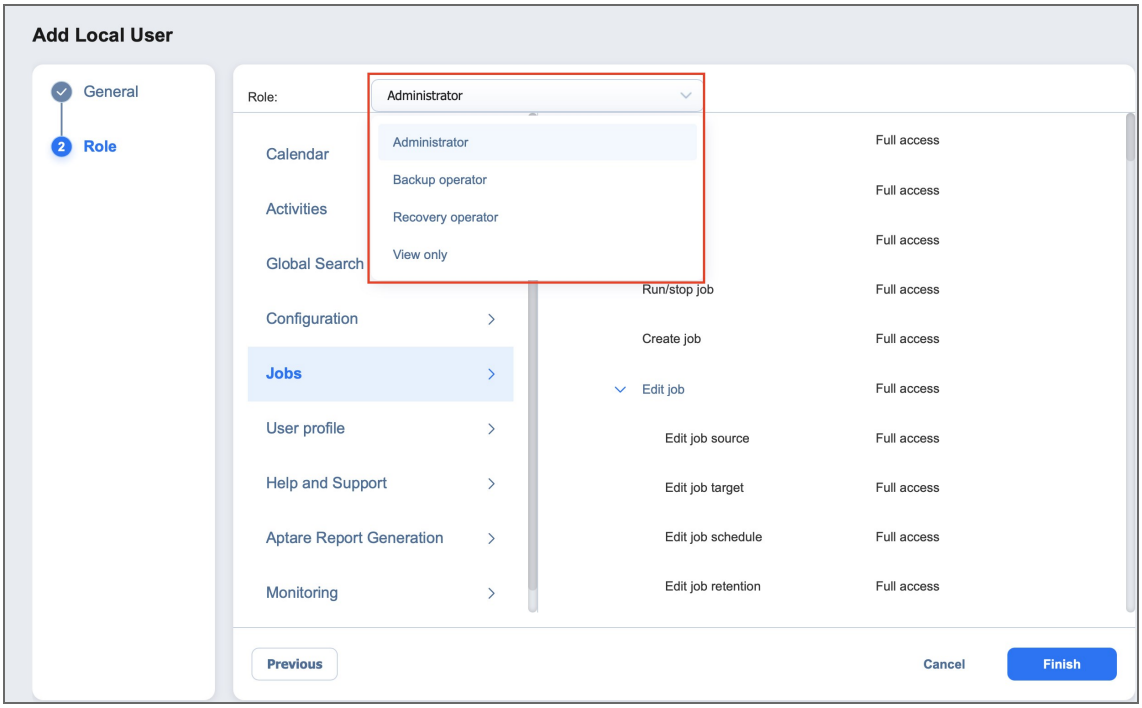


4. The **Add Local User** page opens. Proceed as follows:
   a. In the **Username** box, enter the user name.
   b. In the **Name** box, enter the user's real name.

c. In the **Password** box, enter the user password. To generate a password automatically and send it to the user, select **Generate password and send by email**.
d. In the **Repeat password** box, re-enter the user password.
e. In the **Email** box, enter the user's email address.
f. In the **Description** box, optionally enter a user description.



g. Click **Next** to proceed to the **Role** Tab.
h. In the **Access level** dropdown list, select an access level for the new user (for multi-tenant solutions only).
i. In the **Role** dropdown list, select a user role. Refer to "Managing User Roles" on page 370 for more details about user roles.
j. In the lower right corner of the page, click **Finish**. The local user will appear in the list of users.


Add Local User
General
2 Role
Role:
Administrator
Administrator
Backup operator
Recovery operator
View only
Calendar
Activities
Global Search
Configuration
Jobs
User profile
Help and Support
Aptare Report Generation
Monitoring
Full access
Full access
Full access
Run/stop job
Full access
Create job
Full access
Edit job
Full access
Edit job source
Full access
Edit job target
Full access
Edit job schedule
Full access
Edit job retention
Full access
Previous
Cancel
Finish

Assigning Role to Local User

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

1. Go to **Settings > General > Users & Roles**. The **Users & Roles** page opens in the **Users** view.
2. Hover over the local user, and then click the **…** (ellipsis) button in the rightmost cell of the row.

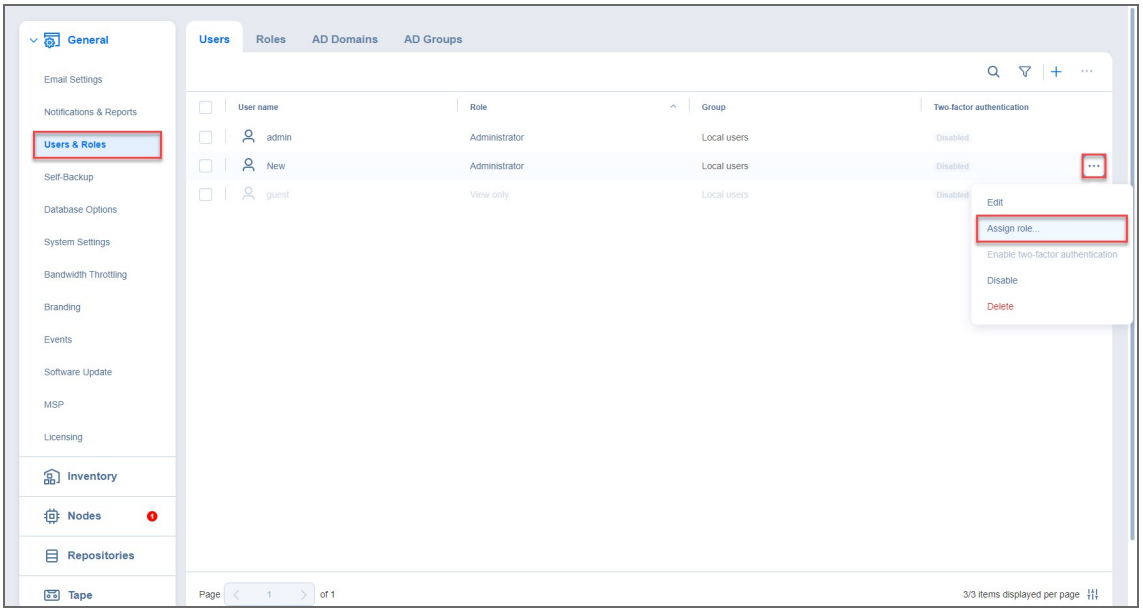
> **Note**
> To assign a role to multiple users, select them using the checkboxes, and then click the ellipsis button for bulk edit.

3. In the resulting menu, click **Assign role**.

> **Important**
> - The **Assign role** button is disabled if:
> - At least one Direct Connect user is selected
> - There are disabled users in the selected users
> - The default administrator is among the selected users

   There are local tenant users and remote tenant users among the selected ones

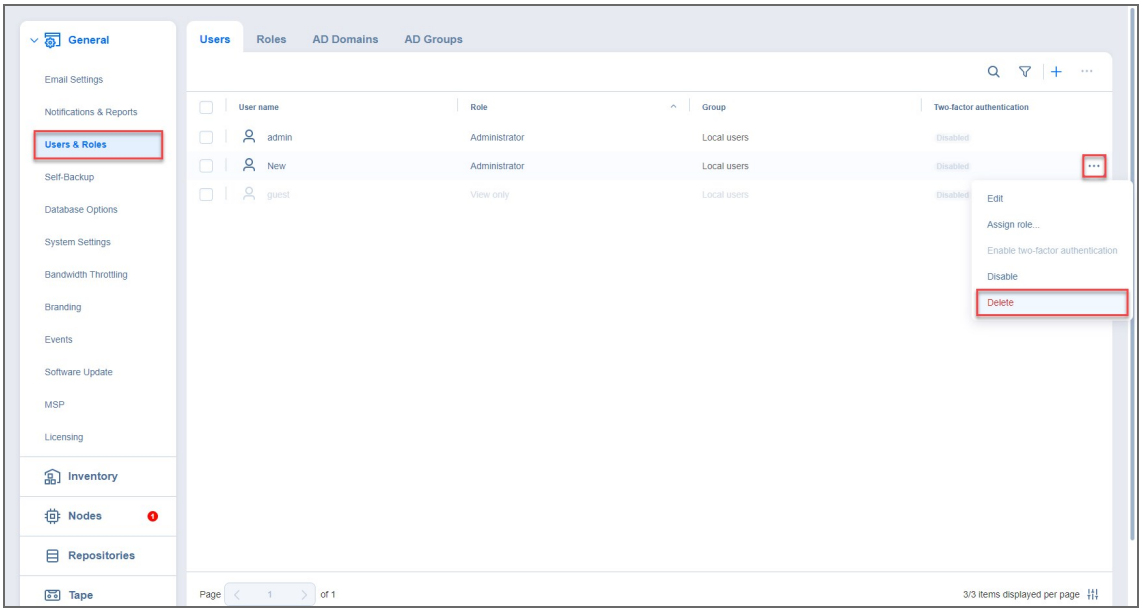
4. In the dialog box that opens, select a new user role from the **Role** drop-down list and then click **Save**.
   The local user will appear in the list of users with the assigned role.



Deleting Local User

Follow the steps below to delete a local user:

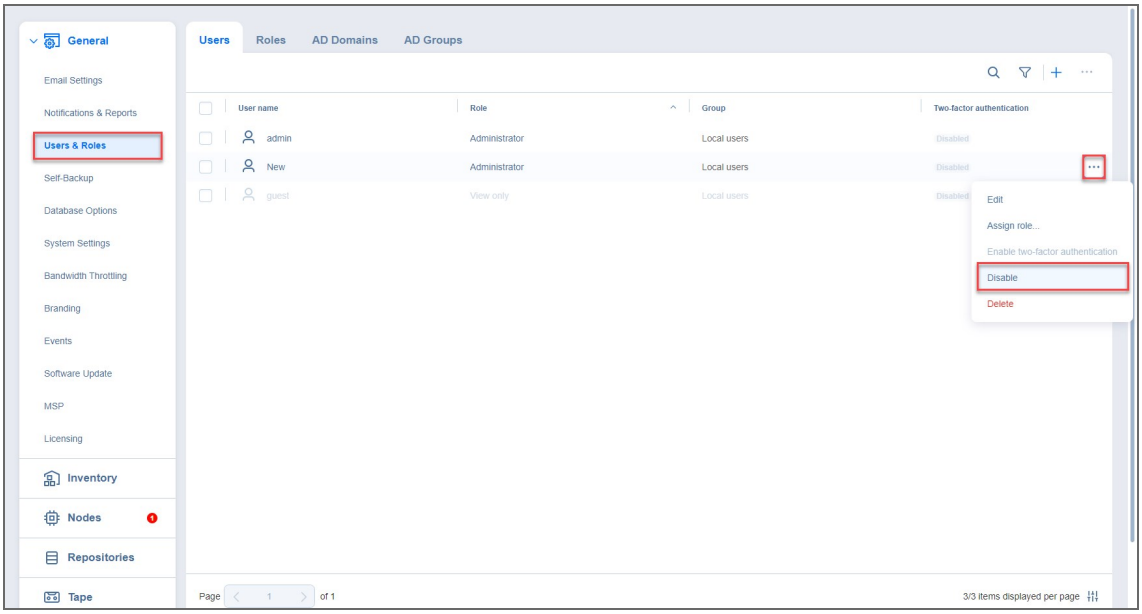
1. Go to **Settings > General > Users & Roles**.
2. The **Users & Roles** page opens in the **Users** view. Hover over the local user you wish to be deleted, and then click the **...** (ellipsis) icon in the rightmost cell of the row.
3. In the resulting menu, click **Delete**.
4. In the dialog box that opens, click **Delete** again to delete the local user. The user will disappear from the list of users.



Disabling Local User

Follow the steps below to disable a local user:

1. Go to **Settings > General > Users & Roles**.
2. The **Users & Roles** page opens in the **Users** view. Hover over the local user you wish to be disabled, and then click the ... (ellipsis) button in the rightmost cell of the row.
3. In the resulting menu, click **Disable**.
4. In the dialog box that opens, click **Disable** again to disable the local user. The user will appear dimmed in the list of local users.

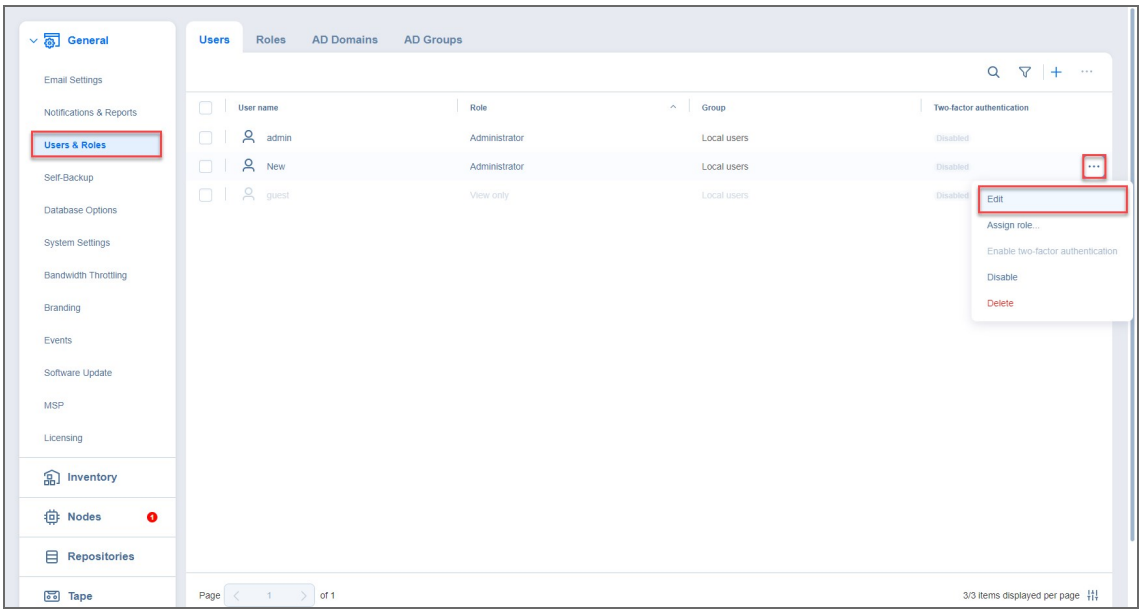


The disabled user cannot log in to the product.

Editing Local User

Please follow the steps below to edit a local user:

1. Go to **Settings > General > Users & Roles**.
2. The **Users & Roles** page opens in the **Users** view. In the list of users, do either of the following:
    a. Locate the local user that you want to edit and click on the user name.
    b. Hover over the local user and click the **...** (ellipsis) button in the rightmost cell of the row. In the resulting menu, click **Edit**.



3. The **Edit User** page opens. Edit the local user properties if needed:
    a. In the **Username** box, edit the user name.
    b. In the **Name** box, edit the user's real name.
    c. In the **Password** box, edit the user password.
    d. If you edited the user password, re-enter the user password in the **Repeat password** box.
    e. In the **Email** box, edit the user's email address.
    f. Optionally, enable **Two-factor authentication**.

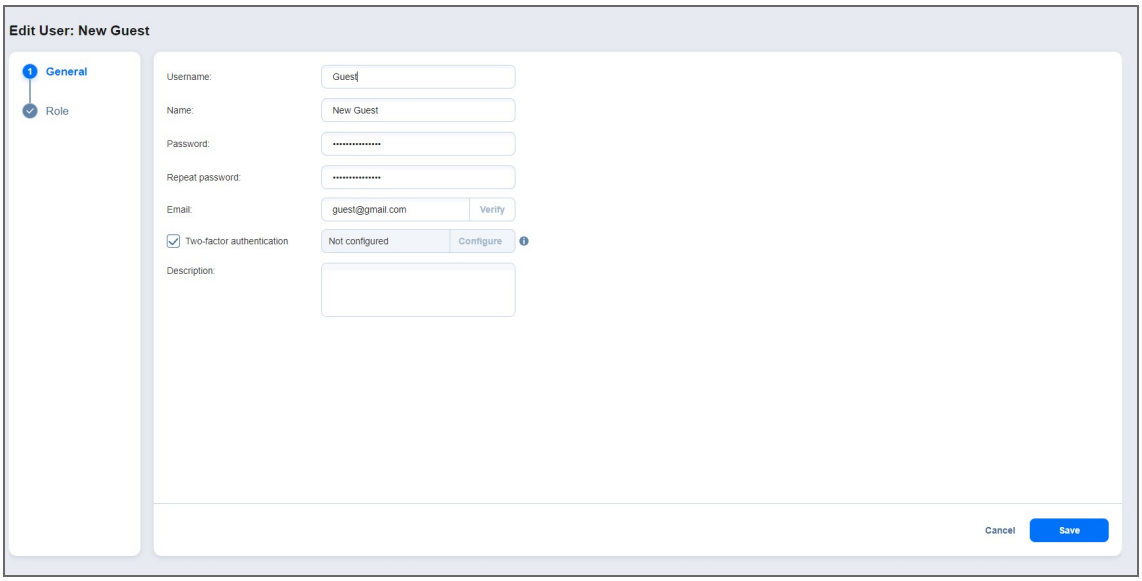
Note

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

g. In the **Description** box, edit the user description.

h. In the **Role** tab, edit the user's role.

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



Managing User Roles

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

- “Overview of User Roles” on page 371
- “Adding User Role” on page 373
- “Editing User Role” on page 375
- “Cloning User Role” on page 377
- “Deleting User Role” on page 378

Overview of User Roles

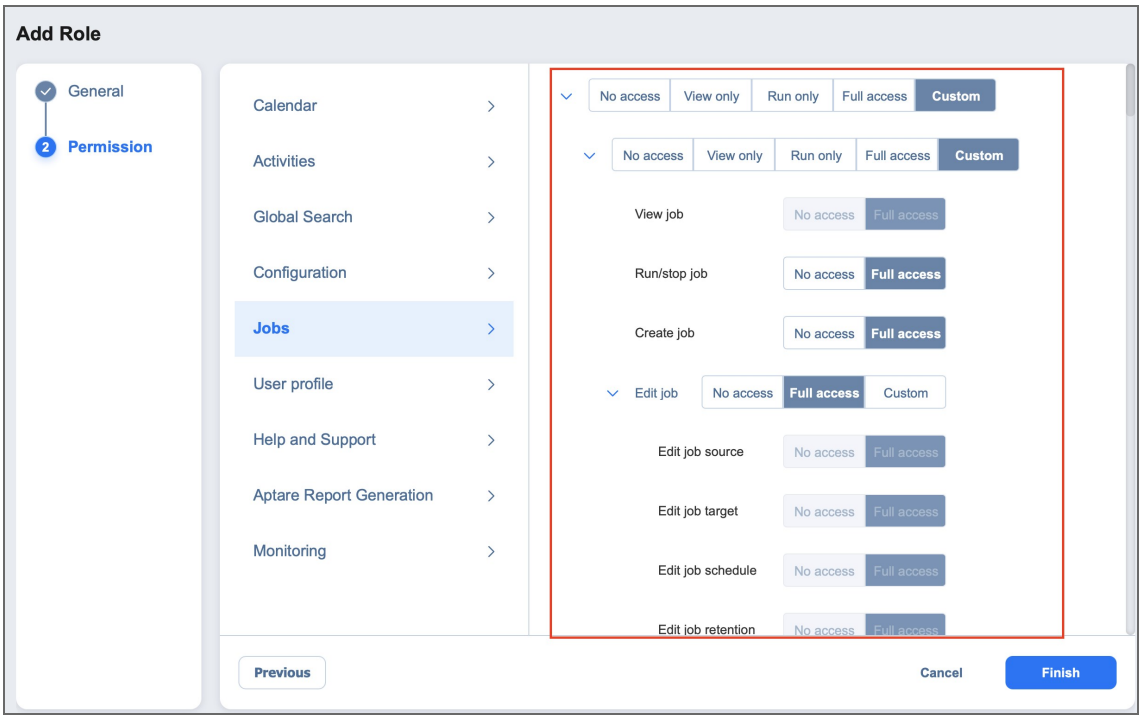
NAKIVO Backup & Replication allows you to assign roles and grant specific permissions to users of the product.

- User Roles
- Access Levels
- Built-in User Roles

User Roles

A user role consists of a set of permissions that can be granted to a NAKIVO Backup & Replication user. Available permissions are grouped by the following product objects:

- **Calendar**: Contains permissions for accessing the **Calendar** dashboard.
- **Activities**: Contains permissions for accessing the **Activities** dashboard.
- **Global Search**: Contains permissions for accessing **Global Search**.
- **Configuration**: Contains a series of permissions for accessing configuration of NAKIVO Backup & Replication.
- **Jobs**: Contains a series of permissions for managing jobs.
- **User profile**: Contains a series of permissions for managing user profile.
- **Help and Support**: Contains a series of permissions for accessing email support, online help center, chat support, and system information.
- **Aptare Report Generation**: Contains permissions for managing Aptare report generation.
- **Monitoring**: Contains permissions for managing the **Monitoring** feature.



Access Levels

There are the following access levels that can be set up for particular permission:

- **No access**: The user cannot view, edit, and run the commands, neither from the graphical interface nor from the command line.
- **View only**: The user can view the commands in the graphical interface but cannot edit or run them; using the command line, the user can only run the commands that do not change NAKIVO Backup & Replication objects.
- **Run only**: The user can only view and run commands, both from the graphical interface and the command line.
- **Full access**: The user can view, edit, and run the commands, both from the graphical interface and the command line.
- **Custom**: A custom set of permissions is configured for a product object.

Built-In User Roles

The product offers you a number of built-in user roles:

- Backup operator
- Recovery operator
- Self-service administrator
- Self-service user
- View only

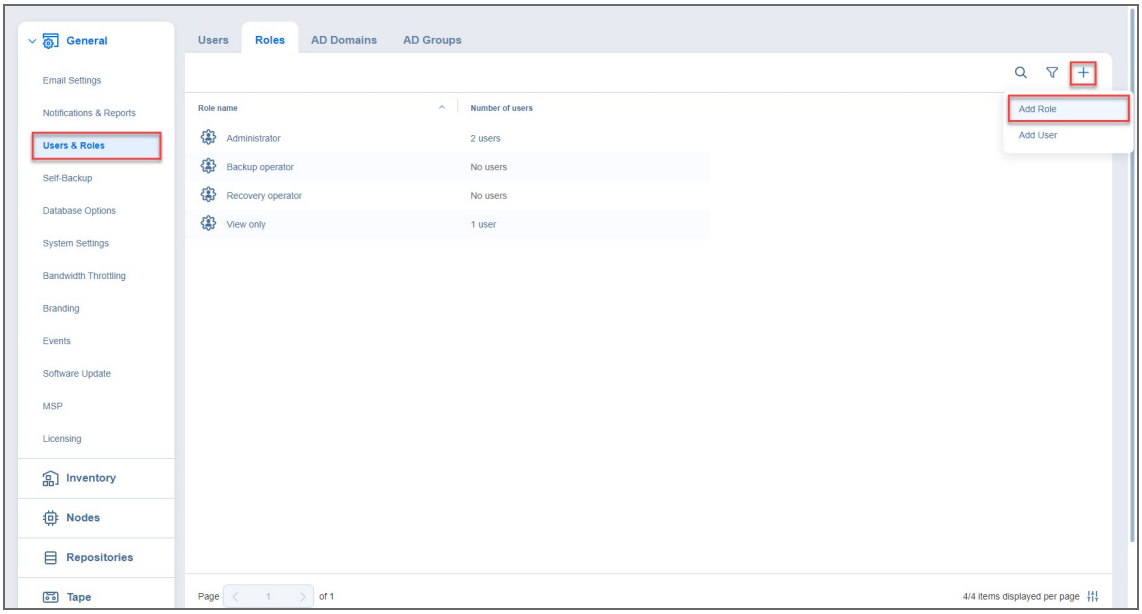
> **Note**
> For editions other than Enterprise Plus, the roles **Backup operator**, **Recovery operator**, and custom roles are not available.

Built-in user roles can be used for performing typical user management tasks. If you need an extra level of security, you can add a new user role or take a built-in user role as a starting point by cloning it.
The user profile can only have a single role assigned.

Adding User Role

Follow the steps below to add a user role:

1. Go to **Settings > General > Users & Roles**.
2. On the **Users & Roles** page, switch to the **Roles** tab.
3. Click the **+** icon and then select **Add Role**.

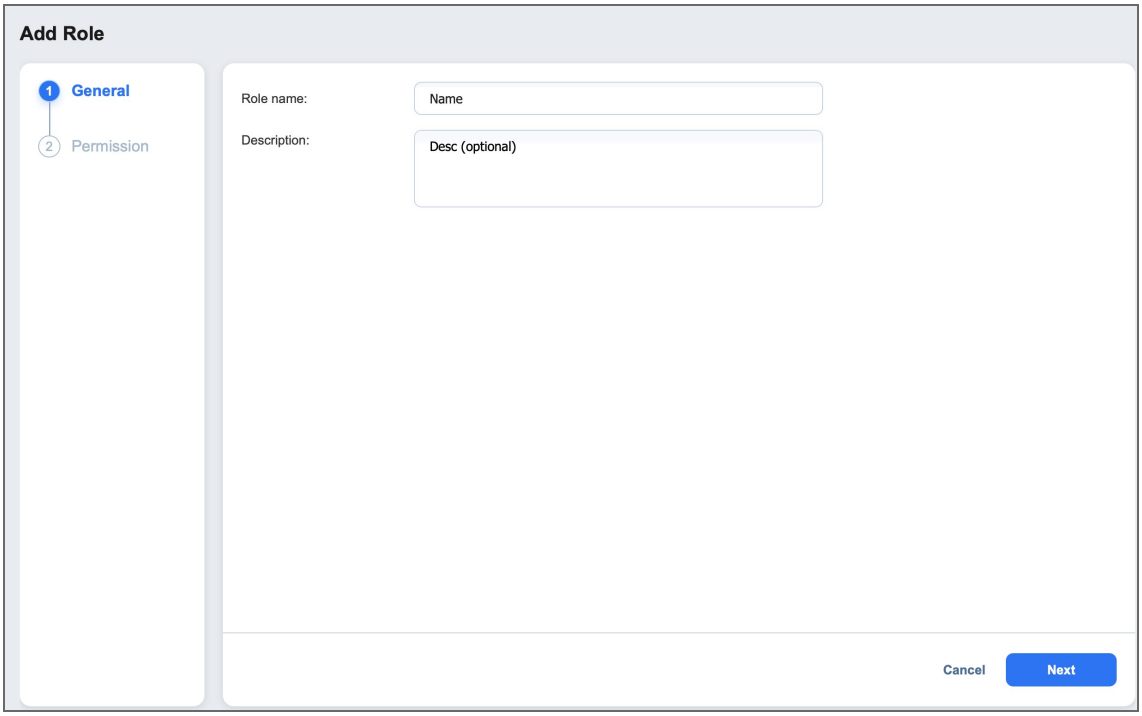


4. The **Add Role** page opens. Proceed as follows:
   a. In the **Role name** box, enter the role name.

   > **Note**
   > In the Multi-tenant mode, every local tenant must have a unique role name.

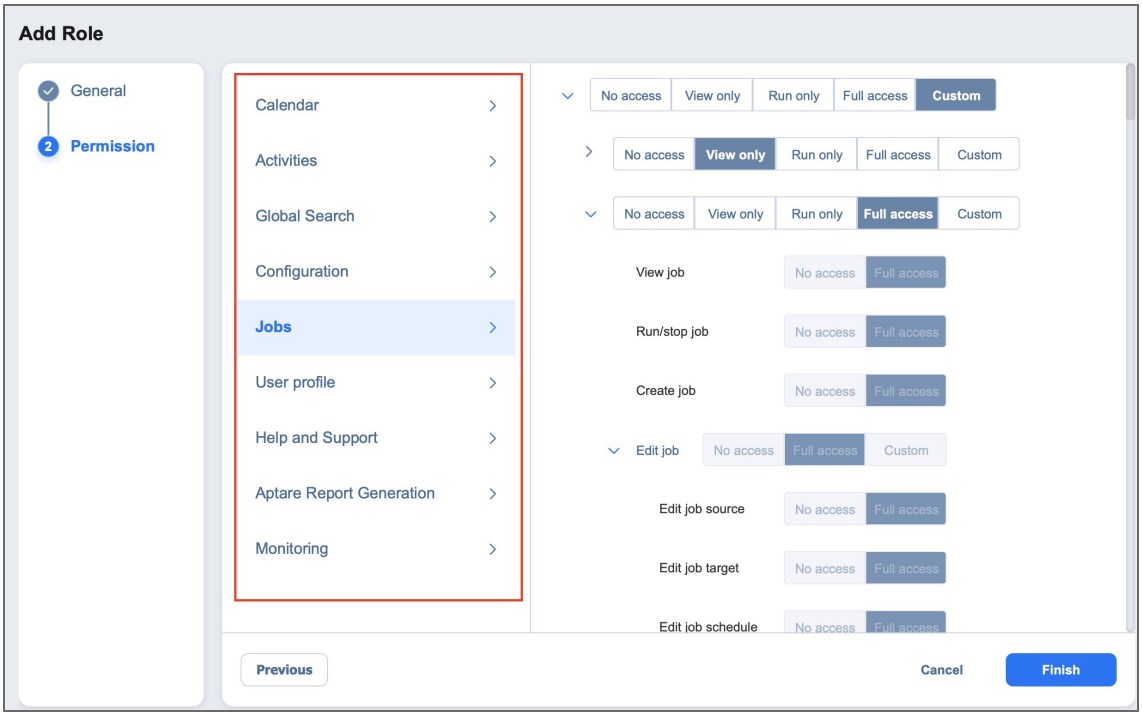
   b. If you are working with a multi-tenant environment, choose either a tenant, master tenant, or all tenants, from the **Access level** list.
   c. In the **Description** box, optionally enter a user description.

   > **Note**
   > For editions other than Enterprise Plus, the roles **Backup operator**, **Recovery operator**, and custom roles are not available.



d. Click **Next** to proceed to the **Permission** tab.

e. A list of permissions opens. Specify necessary permissions for the user role.



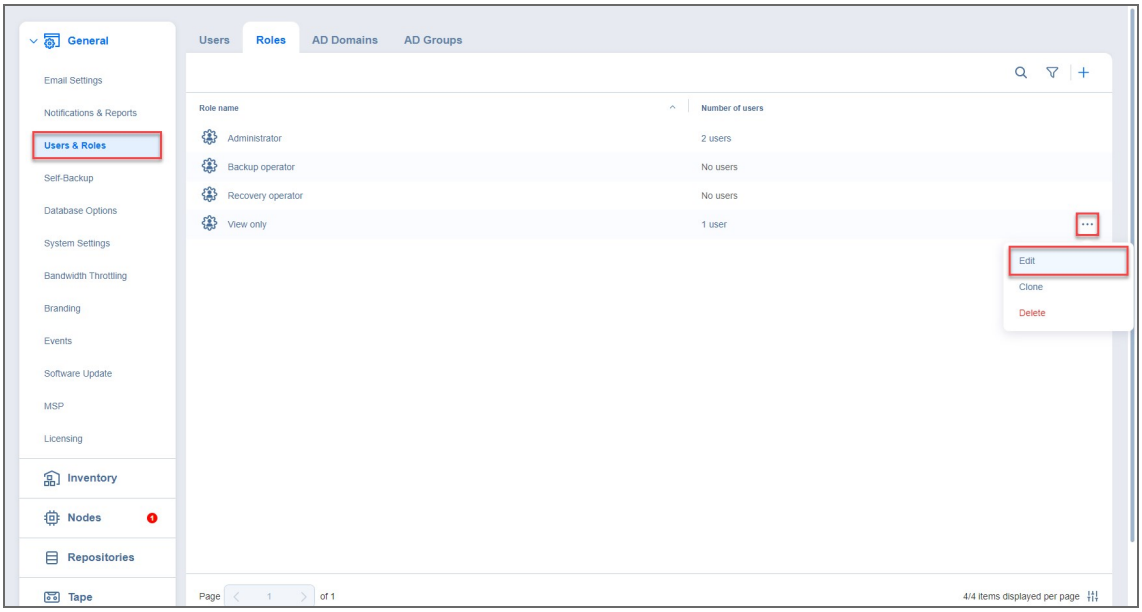
f. Click **Finish** in the lower right corner of the page.

The user role appears in the list of roles.

Editing User Role

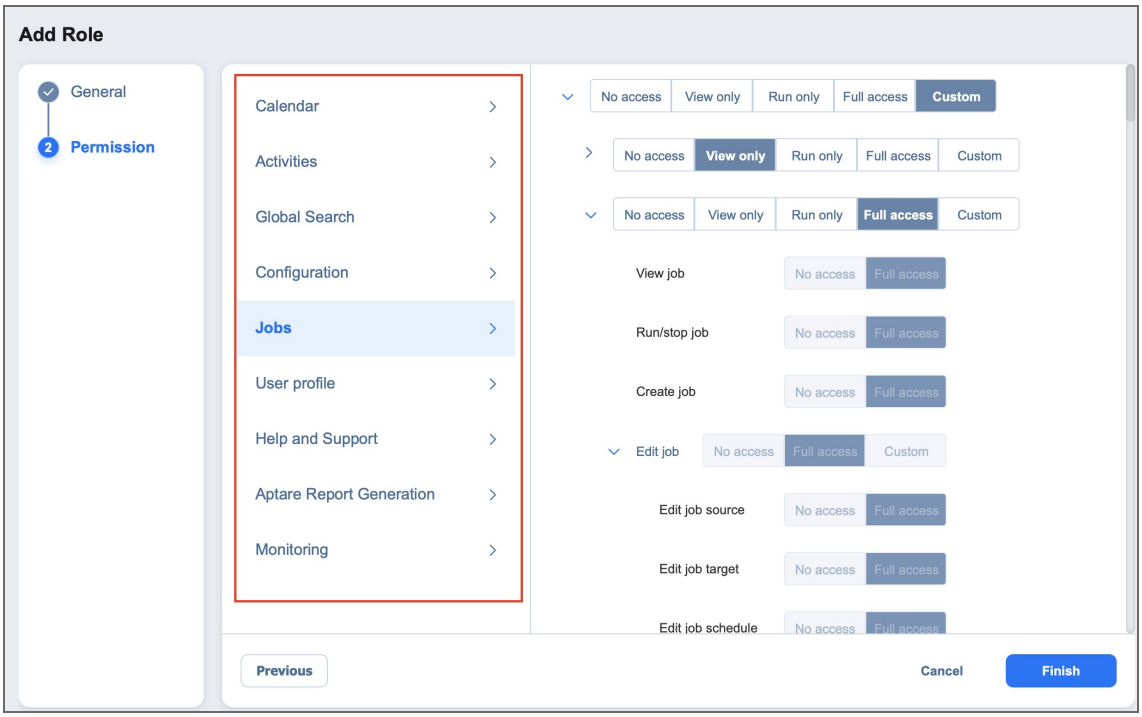
Follow the steps below to edit a user role:

1. Go to **Settings > General > Users & Roles**.
2. On the **Users & Roles** page, switch to the **Roles** tab.
3. In the list of roles, do either of the following:
   a. Locate the user role and click on it.
   b. Hover over the user role, click the **...** (ellipsis) button in the rightmost column of the row, and click **Edit**.



4. The **Edit Role** page opens. Edit the user role properties if needed:
   a. In the **Role name** box, edit the user role name.
   b. If you are working with a multi-tenant environment, you can change the access level for this role by choosing another tenant, master tenant, or all tenants in the **Access level** list.
   c. In the **Description** box, edit the user description.
   d. You can view the **Number of users** with this role, as well as view a full list by clicking the ***x* users** button.
   e. In the **Permissions** tab, you can edit all necessary permissions for the user role.

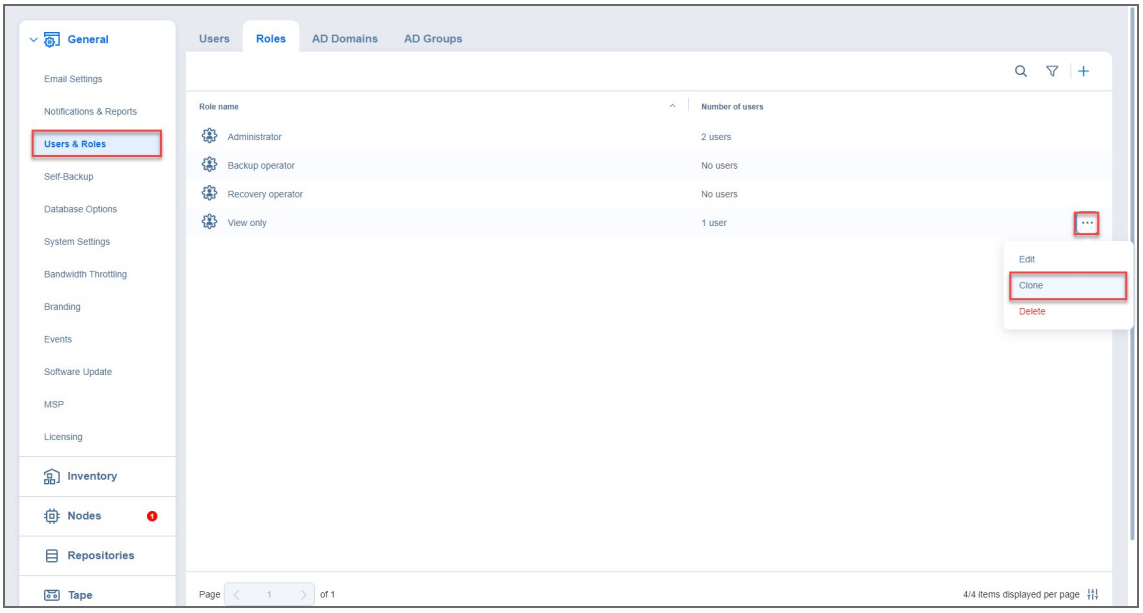
f. When finished, click **Save** in the lower right corner of the page.



Cloning User Role

Follow the steps below to clone a user role:

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

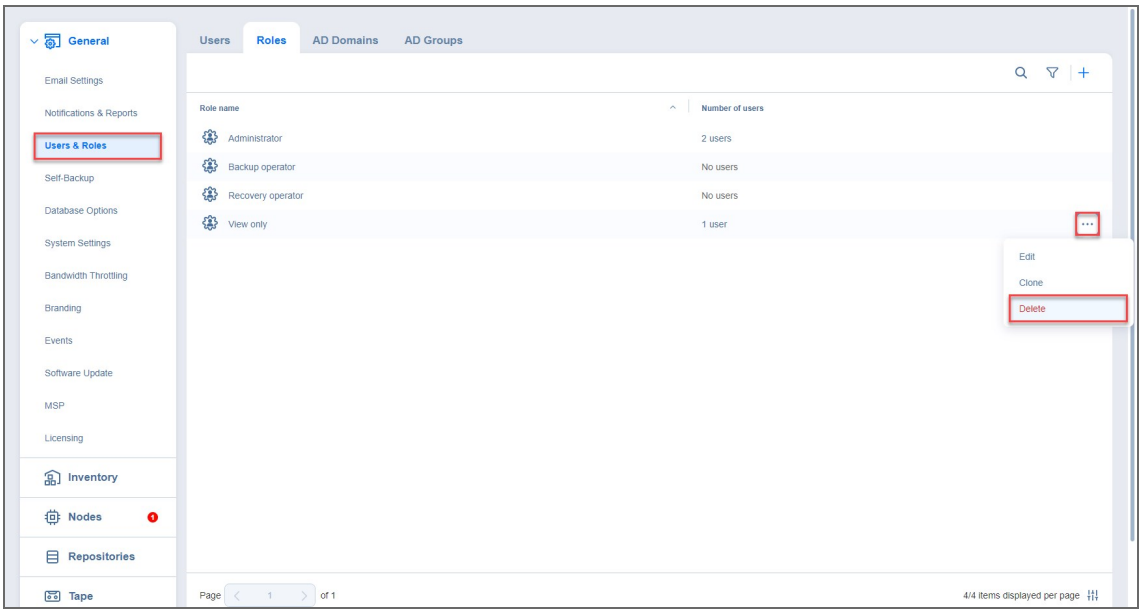


The new user role appears in the list of roles.

Deleting User Role

Follow the steps below to delete a user role:

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



The user role disappears from the list of roles.

Managing Application Passwords

An application password is a unique, 16-character credential that can be generated for a user with the Direct Connect role.

Generate App Password

To generate an application password for the Direct Connect user:

1. Select the Direct Connect user in the list.
2. Click the … (ellipsis) and select **Generate app password**. The application password is generated for the user.

Notes

- If a password has already been generated, **Generate app password** is disabled.
- The user can have only one application password, which remains unaffected by changes to their regular login credentials.
- You can download the existing user application password as a text file.

Download App Password

To download the application password, proceed as follows:

1. Select the Direct Connect user from the list.
2. Click the … (ellipsis) and select **Download app password**. The application password is downloaded as a text file in the browser.

Note

You can generate and download application passwords only for users with the **Direct Connect** role.

Delete App Password

To delete the application password that was assigned to a user, perform the following steps:

1. Select a user with an application password from the list.
2. Click the … (ellipsis) and select **Delete app password**. NAKIVO Backup & Replication deletes the application password for the user.

Configuring Two-Factor Authentication

NAKIVO Backup & Replication allows you to add an additional layer of security with two-factor authentication (2FA). For details, refer to the topics below:

- Enabling Two-Factor Authentication
- Managing Two-Factor Authentication
- Setting Up Google Authenticator

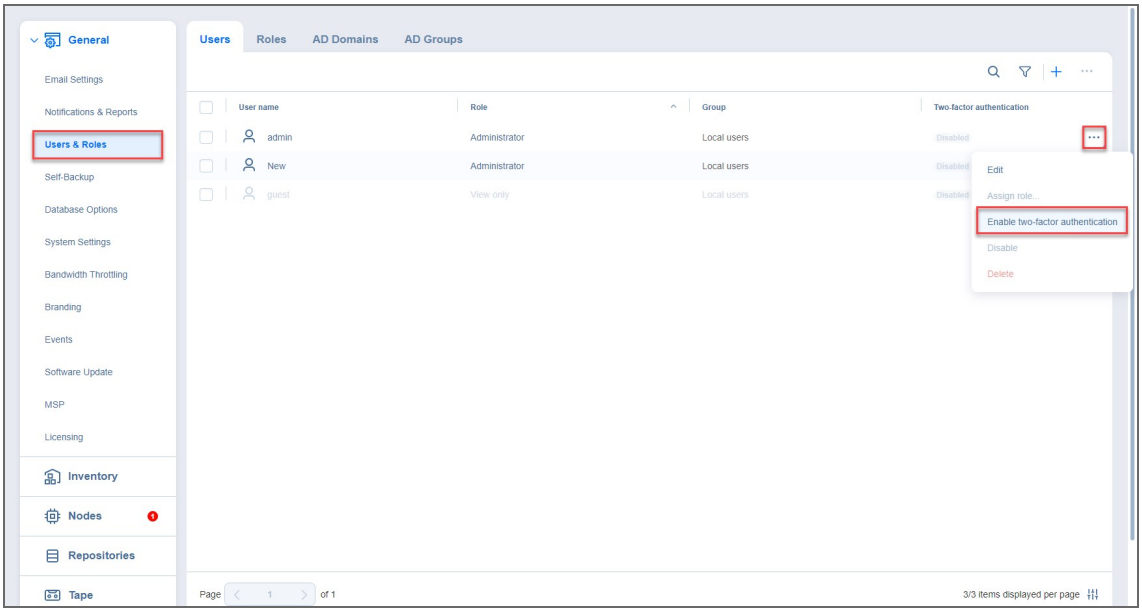
Enabling Two-Factor Authentication

Two-factor authentication can be enabled in either of the following pages:

- On the **Editing local user** page, select the **Two-factor authentication** checkbox.

> **Notes**
>
> - The two-factor authentication is disabled for:
>   - Direct Connect users
>   - Users without **User management** permission
> - Users without **User management** permission cannot enable Two-factor authentication.
> - Users without **Administrator** role or **Configuration** permission can only configure Two-factor authentication on the login screen of NAKIVO Backup & Replication.
> - It is possible to enable Two-factor authentication only after configuring Email Notifications.

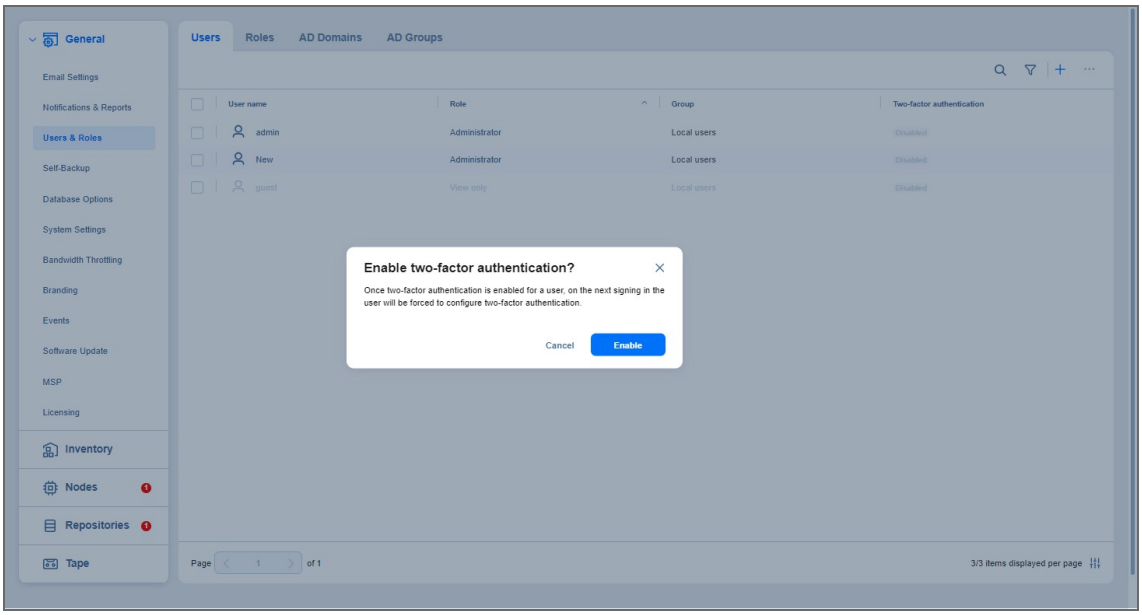
- On the **Users** view, hover over user's name and select **Manage > Enable two-factor authentication**.



> **Note**
>
> - Once two-factor authentication is enabled but not configured for a user, on the next signing in the user will be forced to configure it.

Proceed with configuring two-factor authentication:

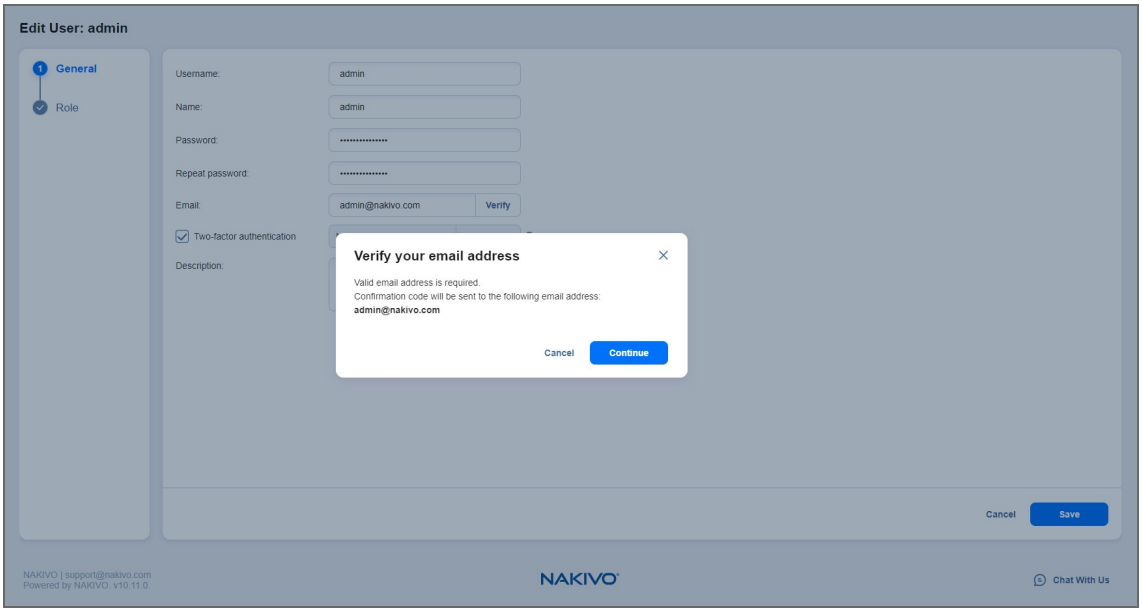
1. Click **Enable** in the dialogue window that appears.



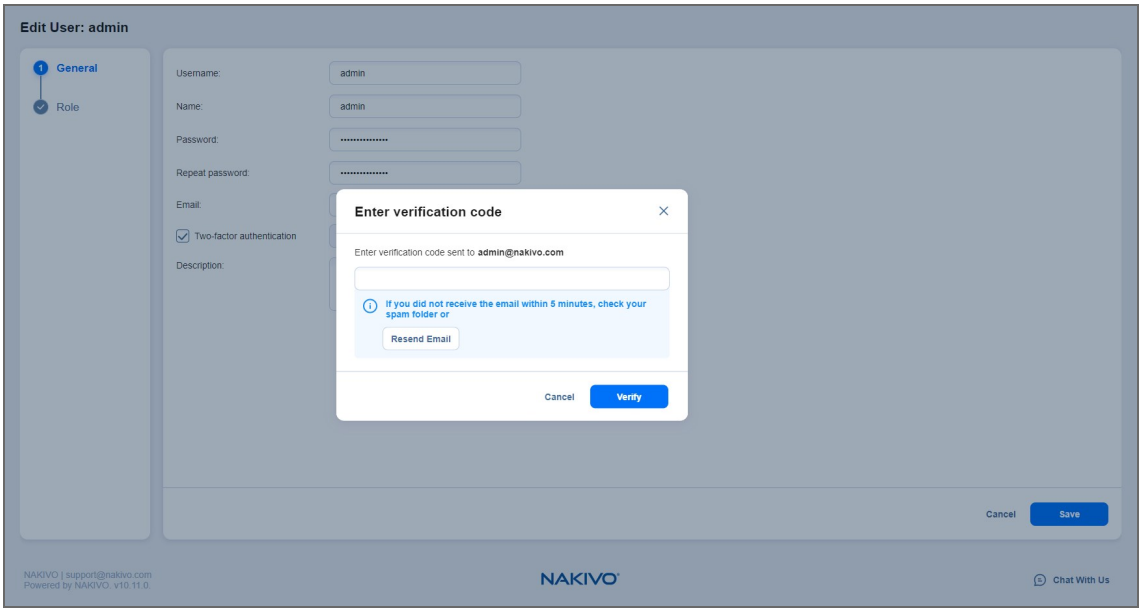
> **Note**
>
> - The two-factor authentication feature is disabled for MSP users drilled down into an ST instance (remote tenant).

2. Click **Continue** in the **Verify your email address** popup that appears.

> **Note**
>
> - Optionally, you can enter the new email address for the user.
> - Also, you can select **Configure** to proceed with 2FA configuration.



3. Enter the verification code that was sent to the specified email address, and click **Verify**.



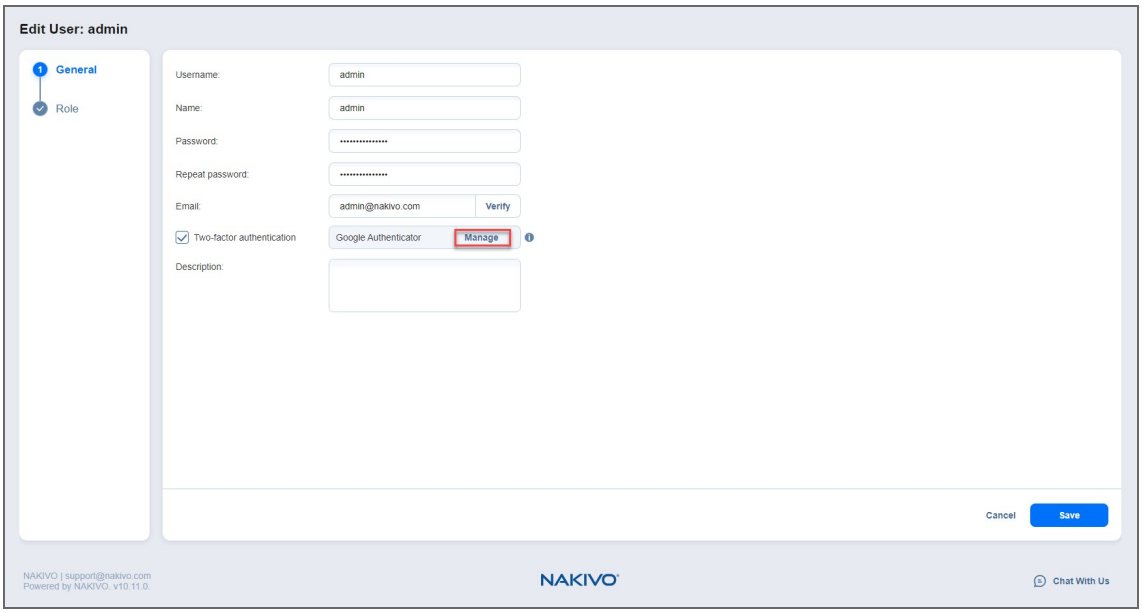
4. Optionally, enter the alternative email address that can be used in case the primary one becomes unavailable, and select **Continue**. Alternatively, select **Skip** to skip this step.
5. If you have entered the alternative email address during the previous step, enter the verification code that was sent to the specified email, and click **Continue** to proceed with Google Authenticator configuration. Alternatively, when configuring 2FA on the **Editing local user** page, select **Cancel** on the **Get Google Authenticator** popup to set up Google Authenticator later.

> **Note**
> When configuring 2FA on the login screen, clicking **Cancel** returns you to the main login screen.

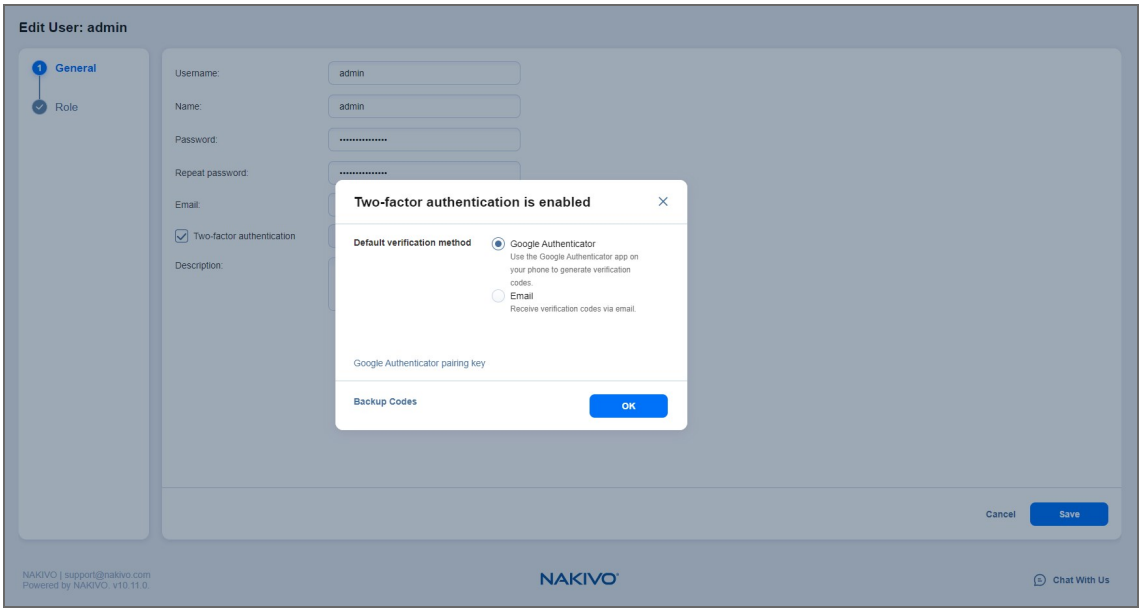
Managing Two-Factor Authentication

You can manage two-factor authentication in the following way:

1. Click the **Manage** link to the right of **Two-factor authentication** checkbox.



2. Choose one of the following verification methods:
    - **Google Authenticator**: Choose this option to use the Google Authenticator app to generate verification codes. Optionally, click on the **Google Authenticator pairing key** link to see your pairing key or on the **Backup codes** link to view your backup codes.
    - **Email**: Choose this option to receive verification codes via email. Optionally, you can view and change your primary email and add an alternative email by clicking the **add** link. Here you can also view your backup codes by clicking the **Backup codes** link.

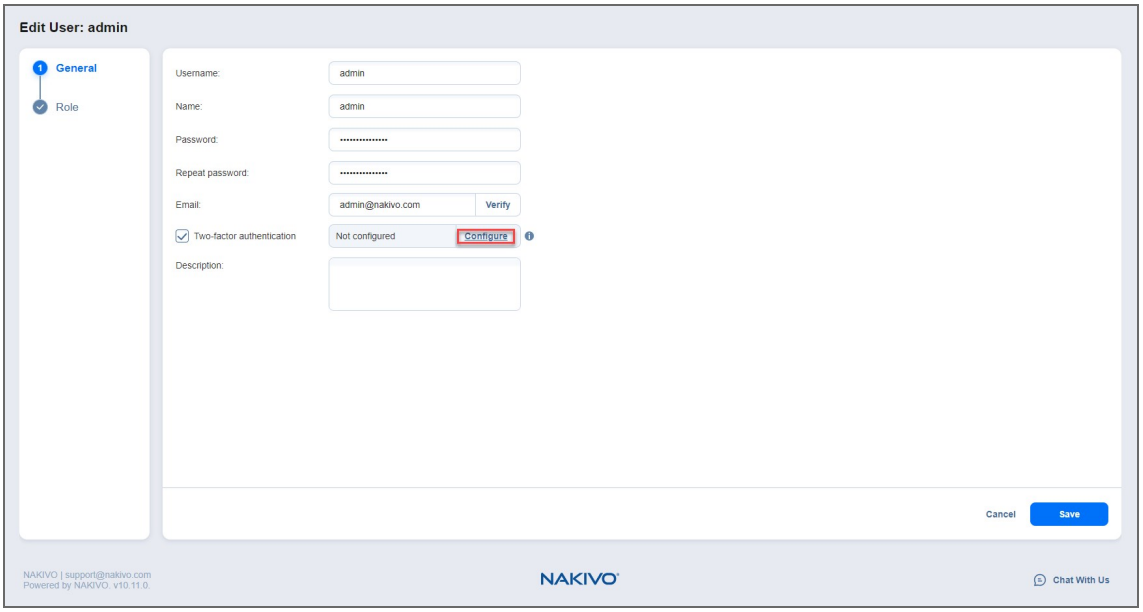


3. Click **OK** when you're done.

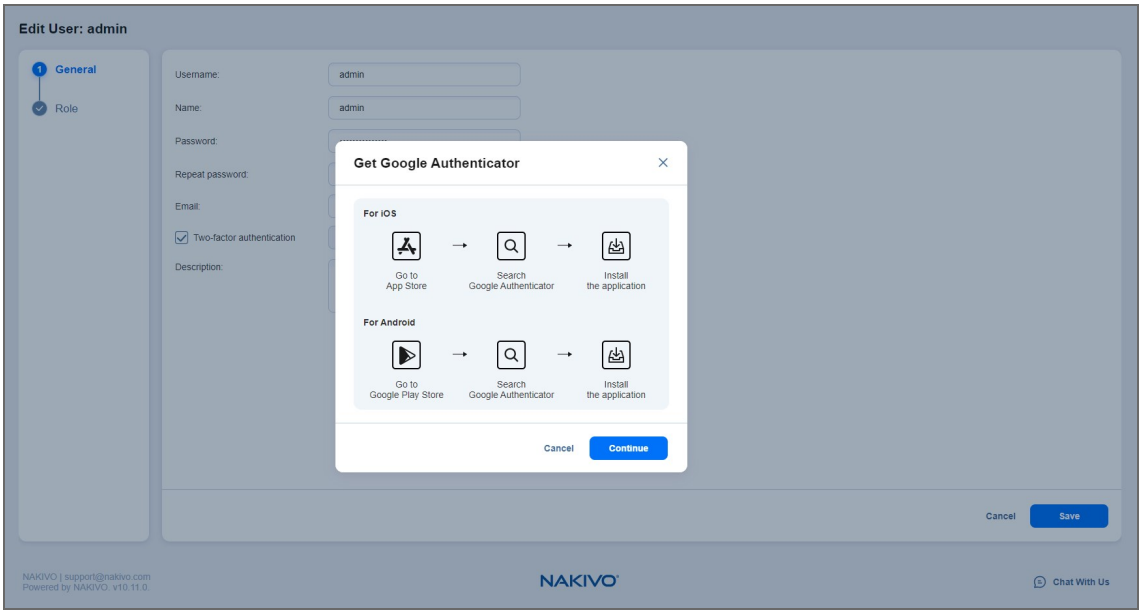
Setting Up Google Authenticator

NAKIVO Backup & Replication uses Google Authenticator for two-factor authentication. To set up Google Authenticator, do the following:

1. Optionally, if you selected **Cancel** on the **Get Google Authenticator** popup, click the **Configure** link to the right of the **Two-factor authentication** checkbox if you are configuring.



2. Follow the instructions in the popup window to download and install Google Authenticator.



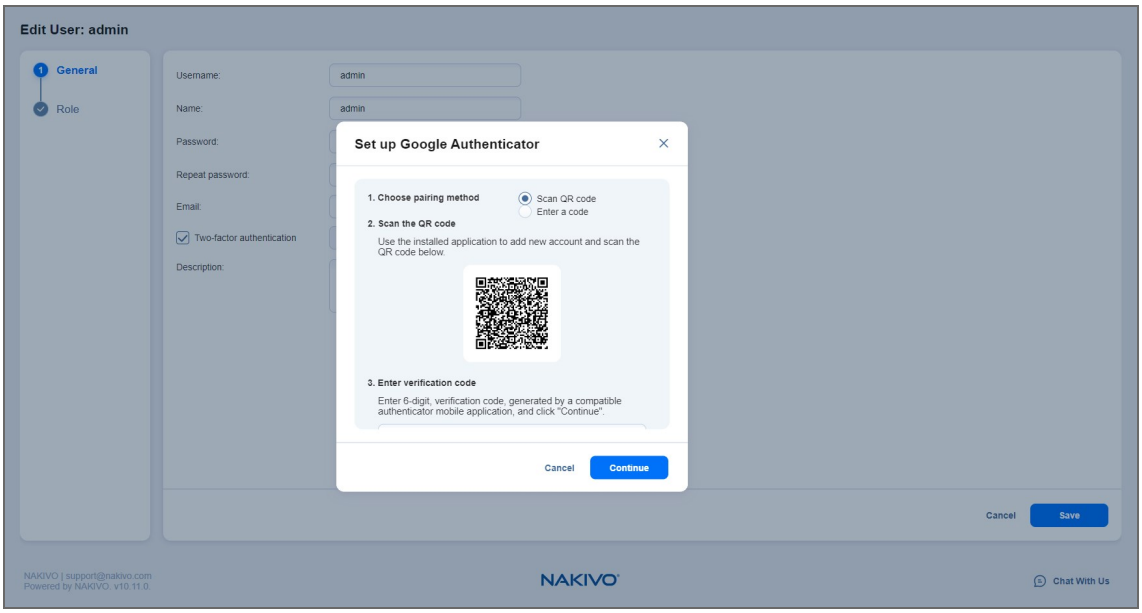
3. Add your NAKIVO Backup & Replication user account to Google Authenticator. Use one of the following methods:

- Select **Scan QR Code** option, and scan the QR code in the popup window.
- Select **Enter a Code** option, and follow the instructions in the popup window to enter the shown code into the Google Authenticator app.
- A popup window appears containing the pairing key, which can be used for adding multiple devices to your account.

> **Important**
> It is highly recommended that you save the pairing key or write it down.

You have the following options:

- Optionally, click on the **copy the Key** link to copy your key and save it for future use.
- Optionally, click on the **download pairing information** link to download and save instructions on how to use the pairing key.
- Click **Continue** when you're done.



4. The **Backup codes** popup window with four backup codes appears. These one-time codes can be used to log in when you are unable to provide a verification code. Click on the **download as PDF** link to download and save these codes in PDF format or write them down. Additionally, you can click the **generate backup codes** link to generate new codes. Click **Continue**.
5. Enter one of the backup codes in the next popup window to confirm that you have saved them, and click **Finish**.

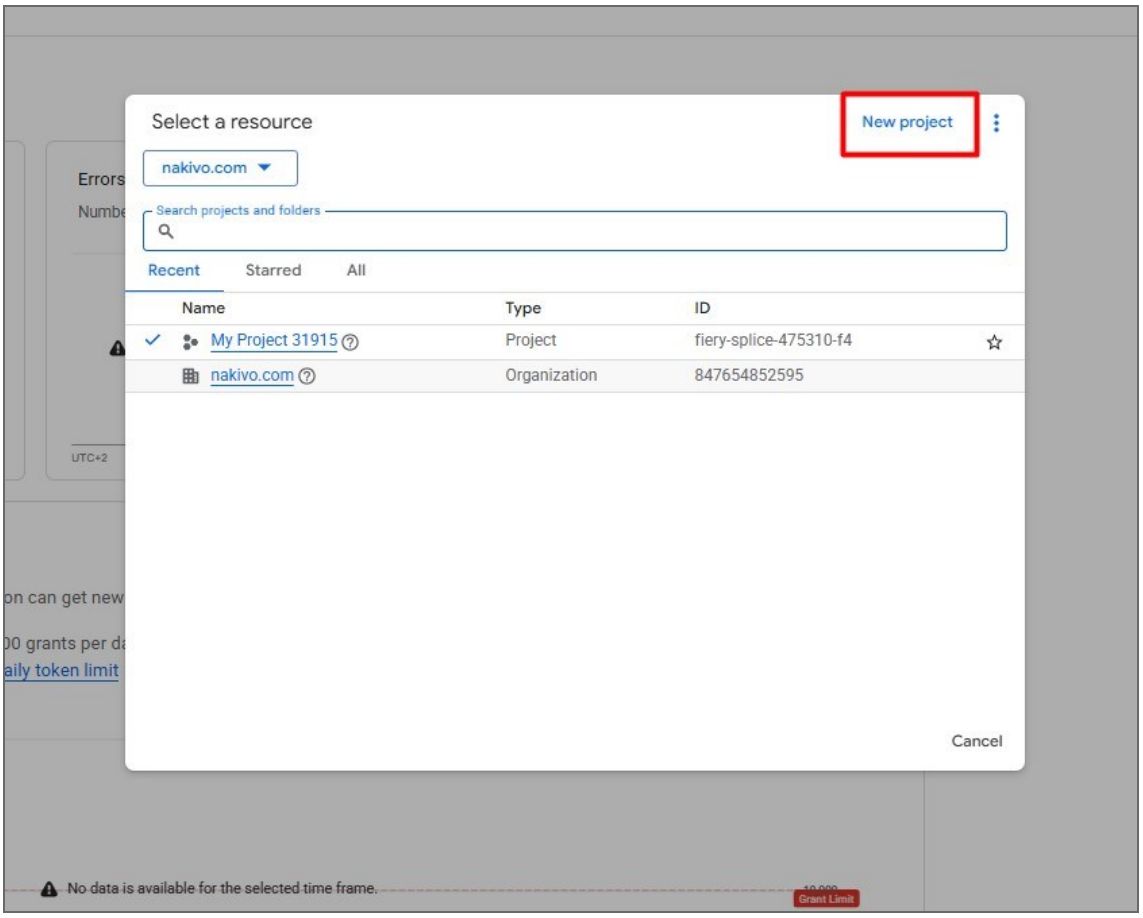
> **Notes**
> - The backup code used in this step remains valid for one more use.
> - The **Manage** link replaces the **Configure** link after this step has been completed.

Using Email with Gmail OAuth Authentication

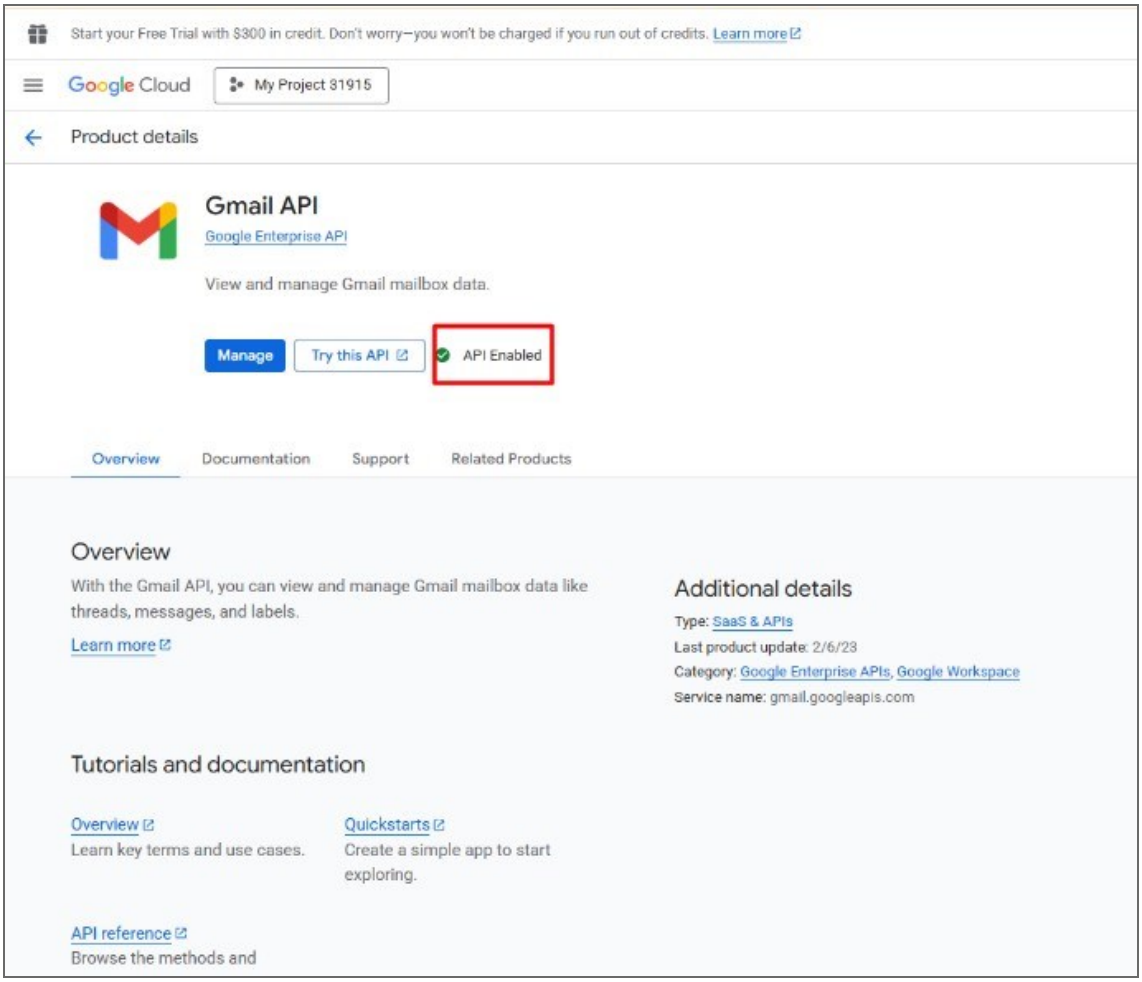
To configure Gmail OAuth 2.0 authentication for NAKIVO Backup & Replication, follow the steps below.

In Google Cloud Console

1. Open Google Cloud Console and create a new project for OAuth configuration. For more information, refer to Google Documentation.



2. Navigate to **APIs & Services** > **API Library**.
3. Find **Google Workspace** and connect the **Gmail API**: select the icon and click **Enable**.



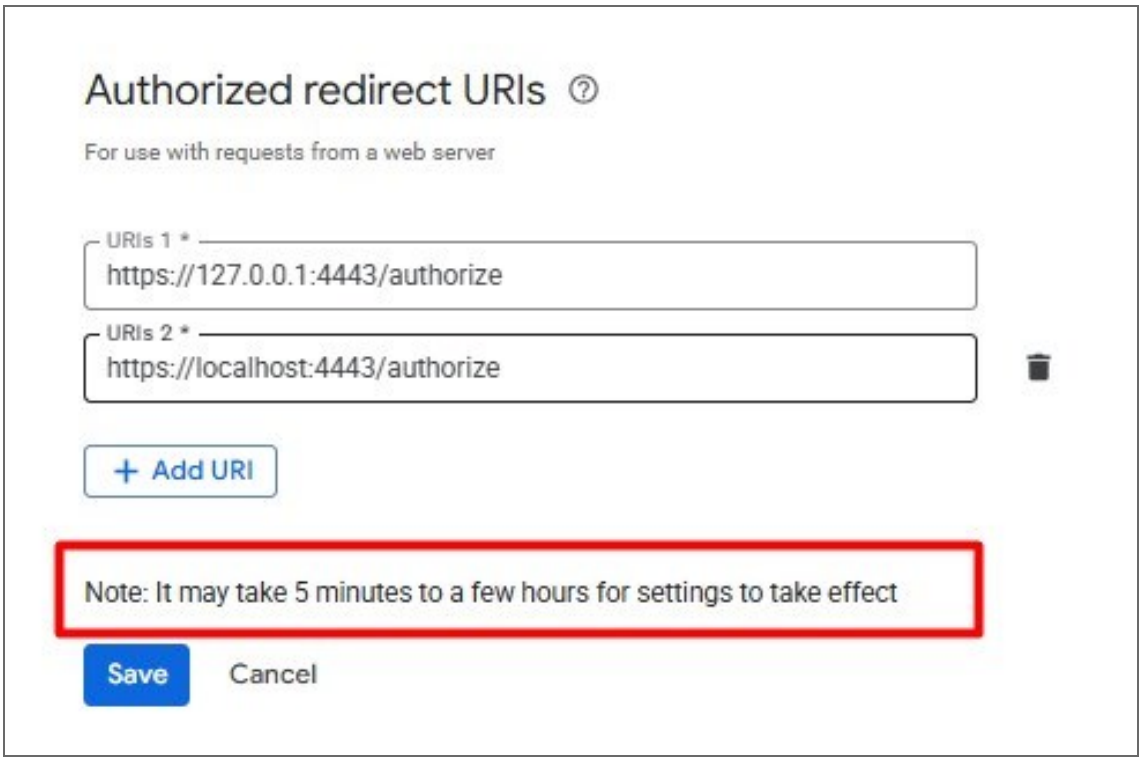
4. Go to the **APIs & Services** > **OAuth consent screen** > **Data Access** tab and add the required Gmail API scopes. See OAuth 2.0 Scopes for Gmail for a list of available permissions.
5. From the **Clients** tab, create a client and set the name and Application type to **Web application**.
6. In the **Authorized redirect URIs** field, enter the following values (when you are logged into NAKIVO Backup & Replication on the VM where it is installed):

```
https://127.0.0.1:4443/authorize
https://localhost:4443/authorize
```

Note

Avoid using private IP addresses (e.g., *https://10.84.0.70:4443/*), as they are not supported.

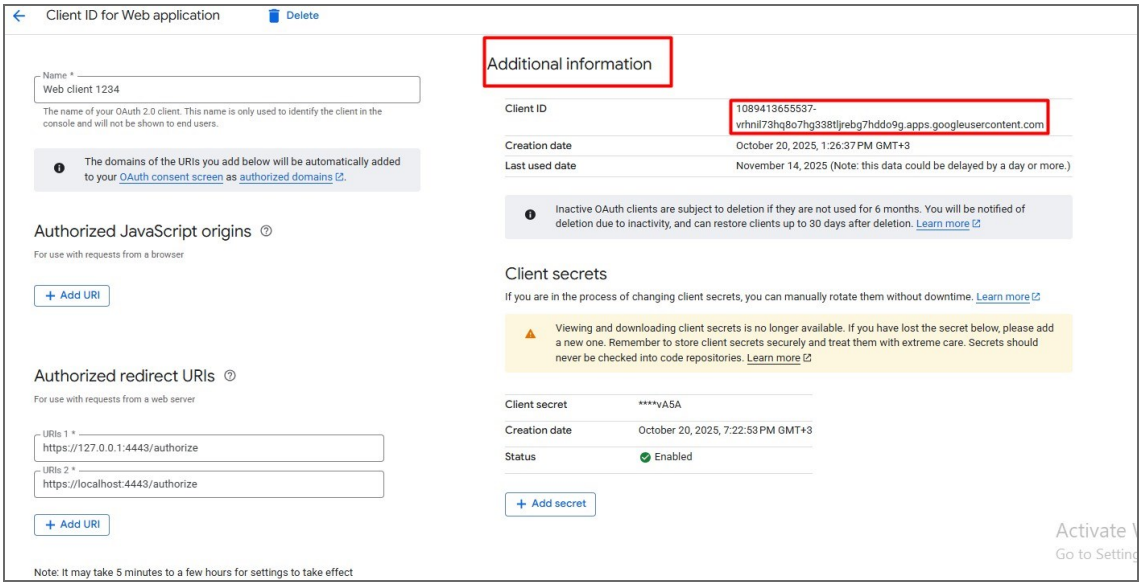
7. Click **Save** to keep the changes. After saving your OAuth configuration, allow several minutes for the changes to take effect.



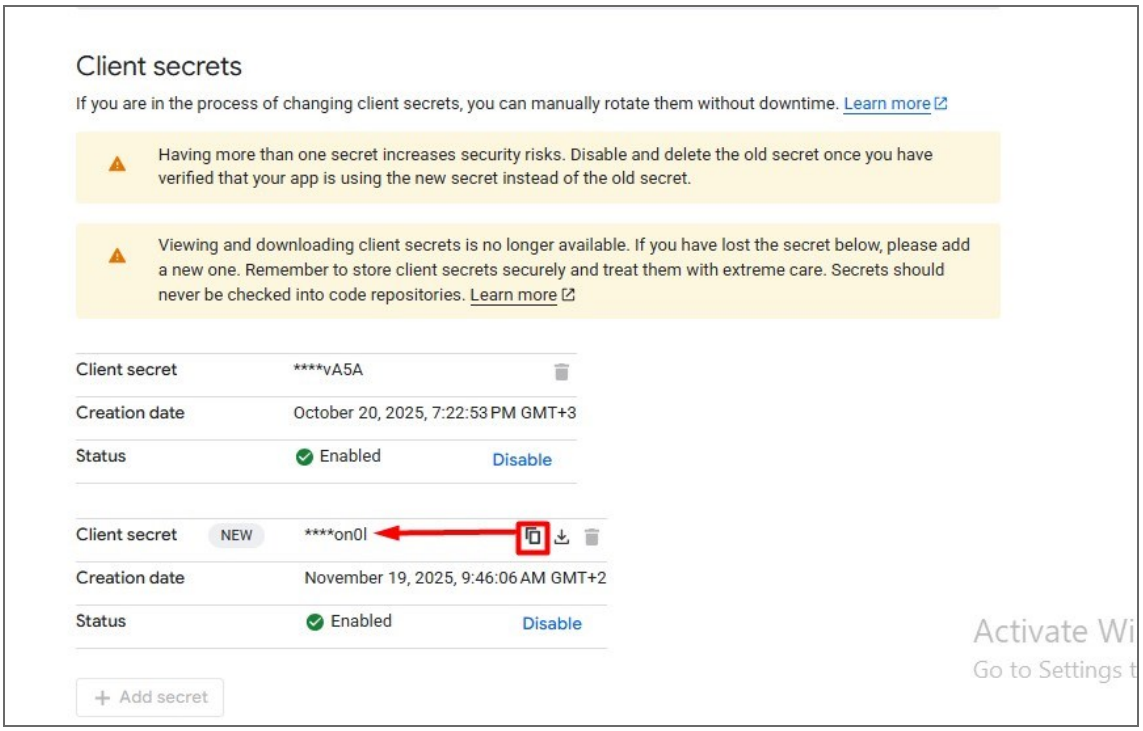
Important

Before starting the authorization process, log in to your Gmail account in the browser. Keep the session active to ensure successful authentication.

8. From the **Additional Information** section, get your **Client ID**.



9. From the Client secrets section, copy your **Client secret**.



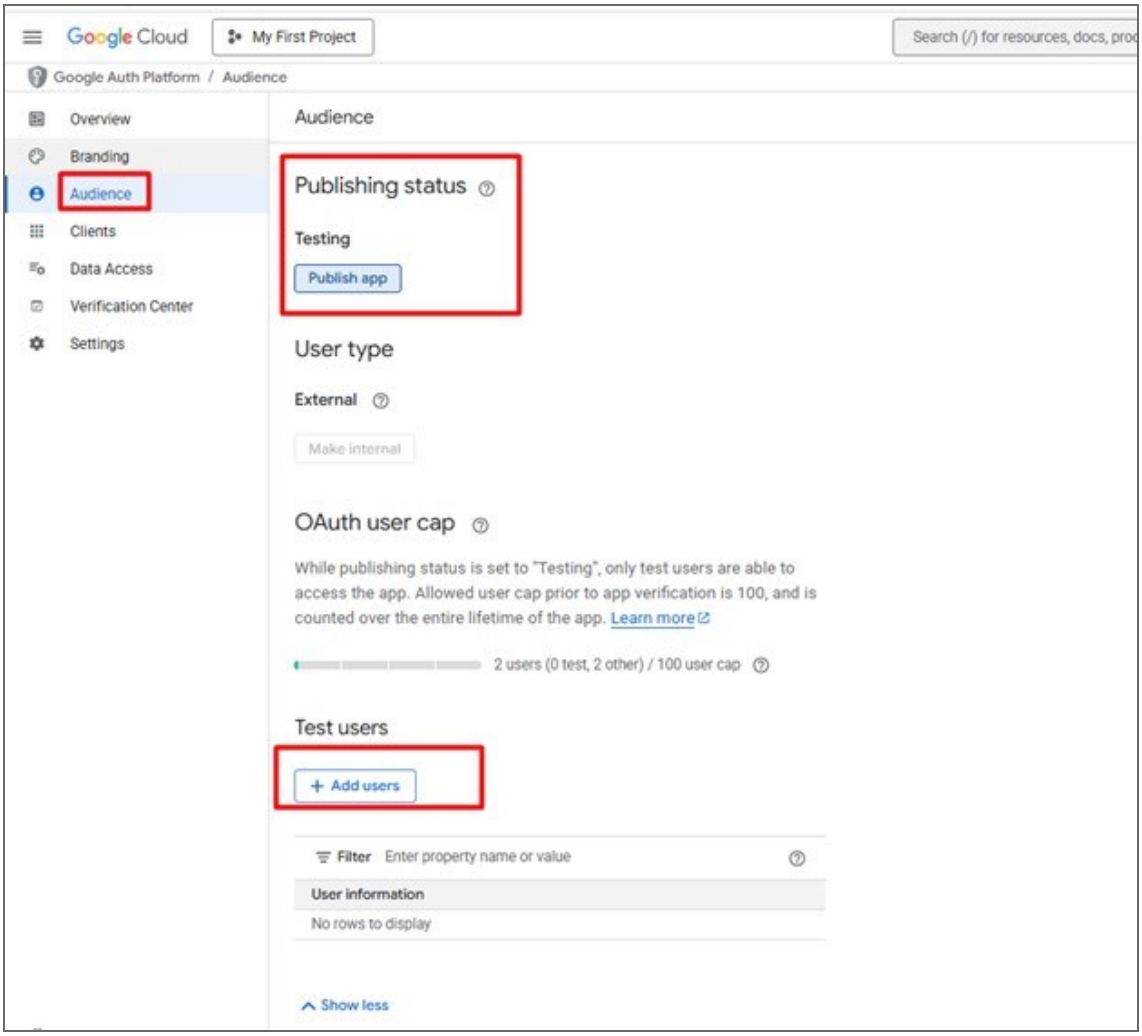
Note

The **Client secret** is temporary. You cannot copy the **Client secret** when it is expired. In this case create a new one.

For New Gmail Accounts

If you are setting up OAuth for a new Gmail account:

1. Follow the same steps as for the Corporate email above.
2. Open the **Audience** tab and add the email address that will be used in **Test users**.
3. Change the **Publishing status** from **Testing** to **In production**.

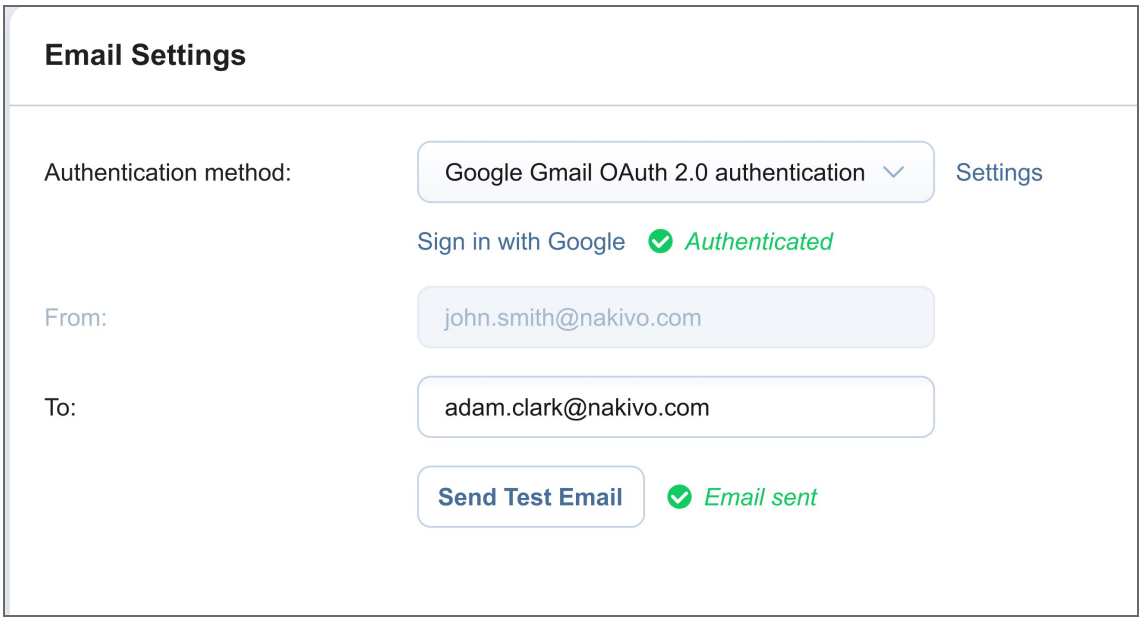


In NAKIVO Backup & Replication

1. Navigate to **Settings** > **Email Settings**.
2. Under **Authentication method**, select **Google Gmail OAuth 2.0 authentication**.
3. Enter the following information:
   - **Client ID**: Enter the value obtained from Google Cloud Console.
   - **Client Secret**: Enter the corresponding secret key.
4. Click **Sign in with Google**.
5. In the web browser window, provide the Google account credentials to connect to NAKIVO Backup & Replication. If the authentication is successful, the 'Authenticated' notice appears.
6. Enter the email address in the **To** field.
7. Click **Send Test Email** to confirm that the configuration works.

Notes

- Ensure the Gmail API is enabled and the correct redirect URIs are configured.
- If authorization fails, verify credentials and permissions in Google Cloud Console.
- Ensure you are signed in to your Google account during setup.

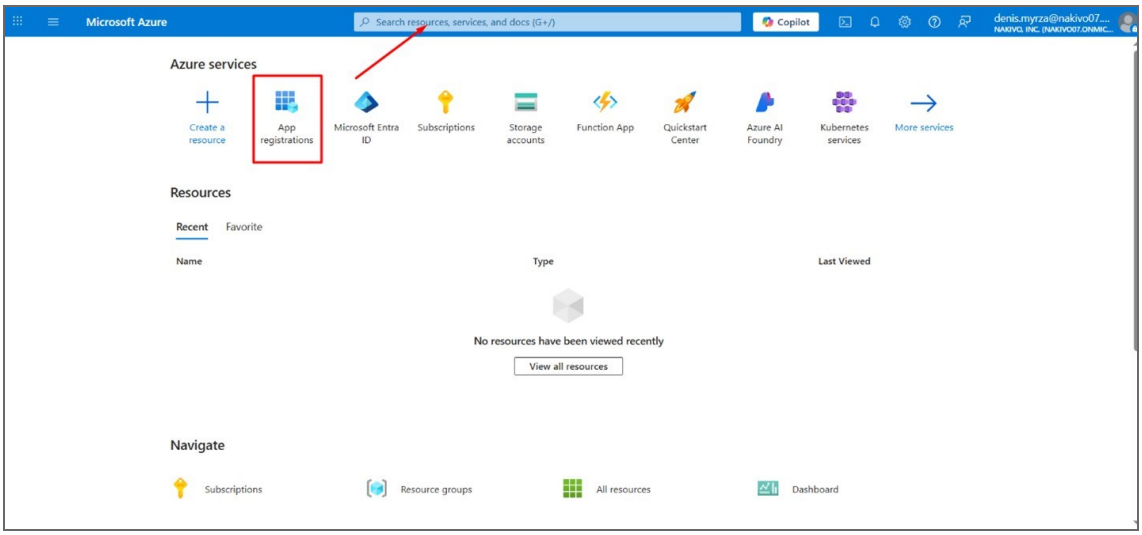


Using Email with Microsoft OAuth Authentication

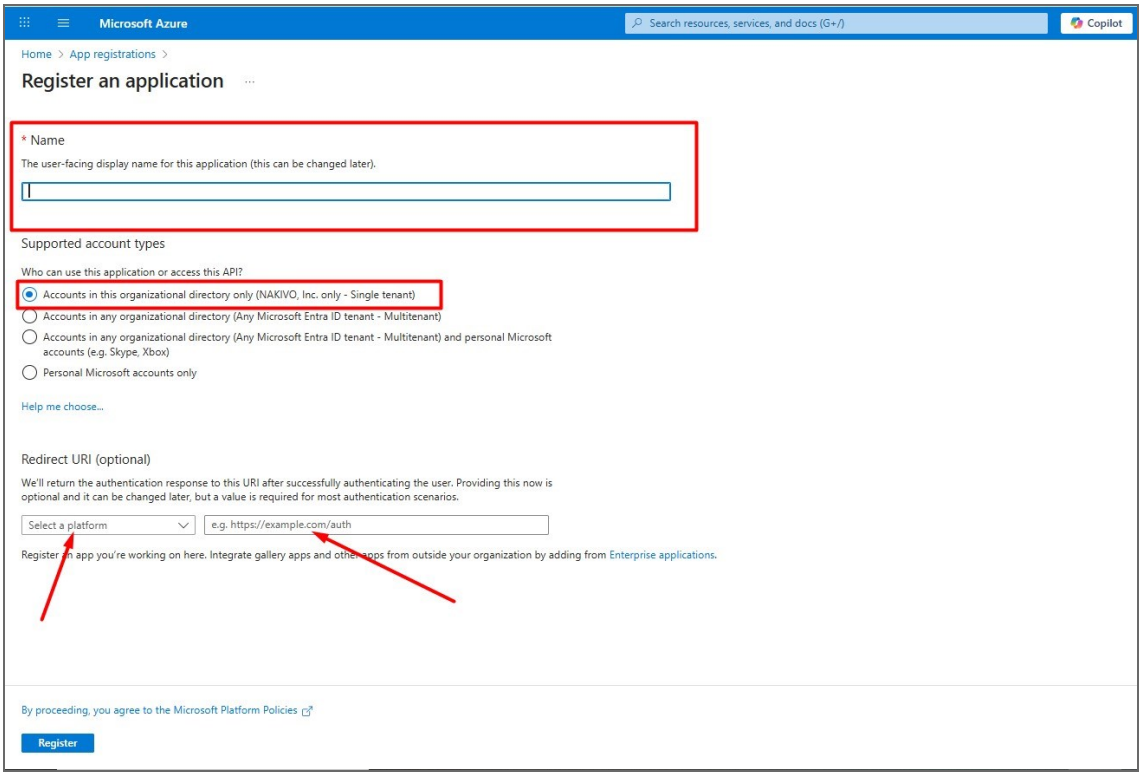
This section explains how to configure Microsoft 365 so that NAKIVO Backup & Replication can send emails using OAuth 2.0 authentication. Follow the steps below before adding your Microsoft 365 credentials in the Email Settings.

Register an Application in Microsoft Azure

1. Sign in to your Microsoft Azure account and open the **App registrations** tab.
2. Click **New registration** to open the registration form.



3. Enter a name for the application.
4. Under **Supported account types**, select **Accounts in this organizational directory only (NAKIVO, Inc. only - Single tenant)**.



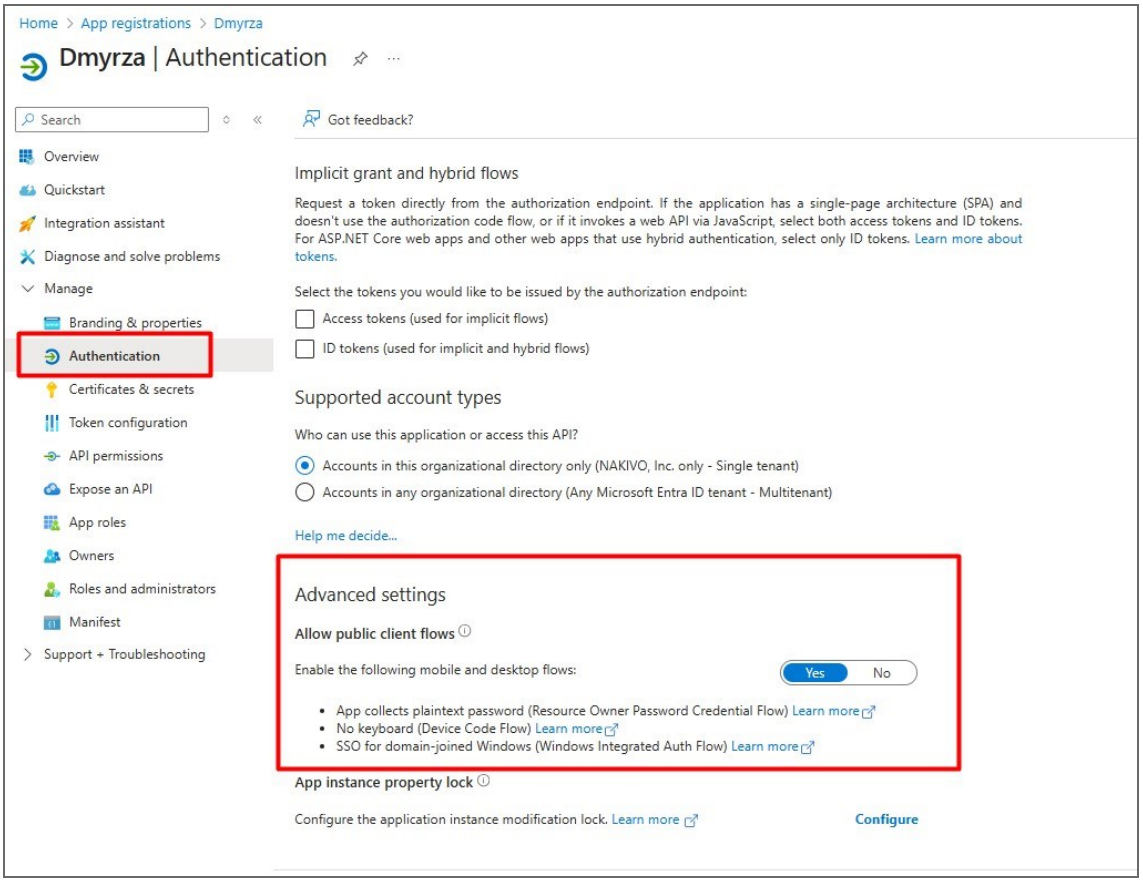
5. In **Redirect URI (Optional)**, select the **Web platform** and enter the appropriate URL based on how you access NAKIVO Backup & Replication:

- https://localhost:4443/authorize – if accessing NAKIVO Backup & Replication locally inside the VM where it is installed.
- https://127.0.0.1:4443/authorize – alternative local address.
- https://<Director-IP>:4443/authorize – if accessing NAKIVO Backup & Replication remotely.

These URLs represent where Azure will send the authorization code after login. If you intend to authenticate using a client certificate, upload your certificate to Azure > App Registration > Certificates & secrets > Upload certificate before authorizing NBR.

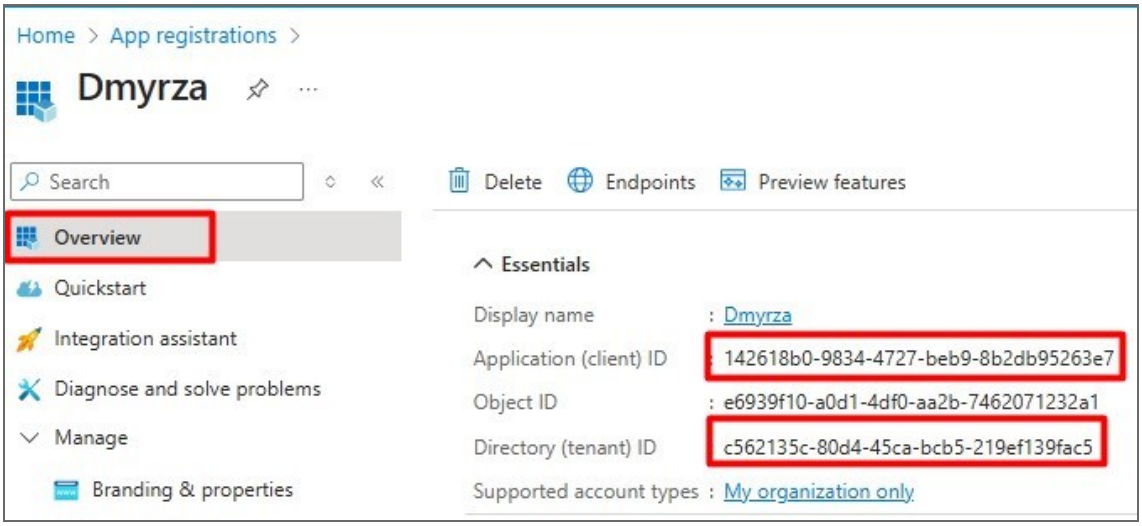
Configure Authentication Settings

1. Go to **Authentication** > **Advanced settings**.
2. Click **Yes** to enable the flows.



3. Go to **Overview** > **Essentials** and copy the following credentials to configure the **Email Settings** in NAKIVO Backup & Replication:

- **Client ID**
- **Tenant ID**

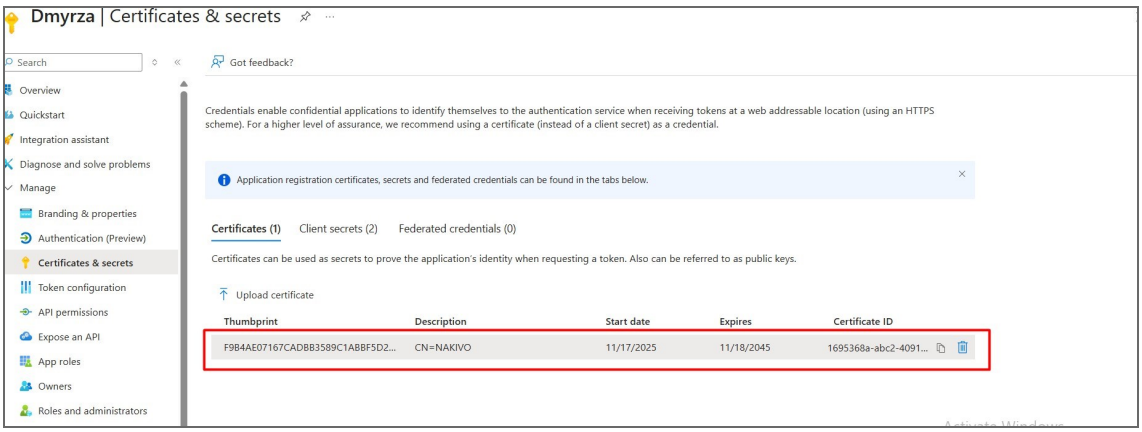


Authentication Settings

NAKIVO Backup & Replication supports two authentication methods. Please select one of the following:

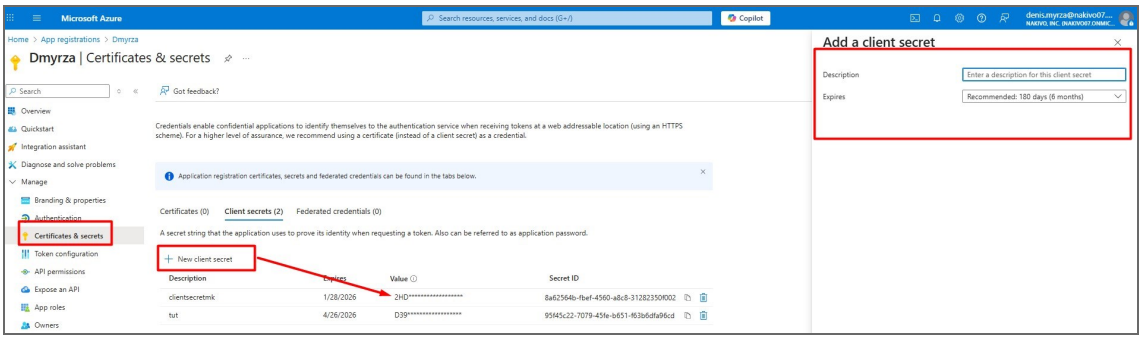
- **Client Certificate (Recommended)**: This option is selected by default. To ensure successful authorization, make sure to use the correct certificate format for each system:
  - **Microsoft Azure**: Upload the public key (typically .pem or .cer) to the Certificates & secrets section of your App Registration.
  - **NAKIVO Backup & Replication**: Upload the corresponding private key (usually a matching .pfx file) to the Client Certificate field.

> **Note**
>
> To download or import certificates in PFX/PEM format from Azure Key Vault, ensure you have the appropriate permissions (Get, Import, and Download). Your Azure administrator may need to assign these permissions.



- **Azure Client Secret**: If you prefer not to use a certificate, or if a valid Client Secret already exists, you can switch to this method (it will disable the Client Certificate fields). To generate a new secret:

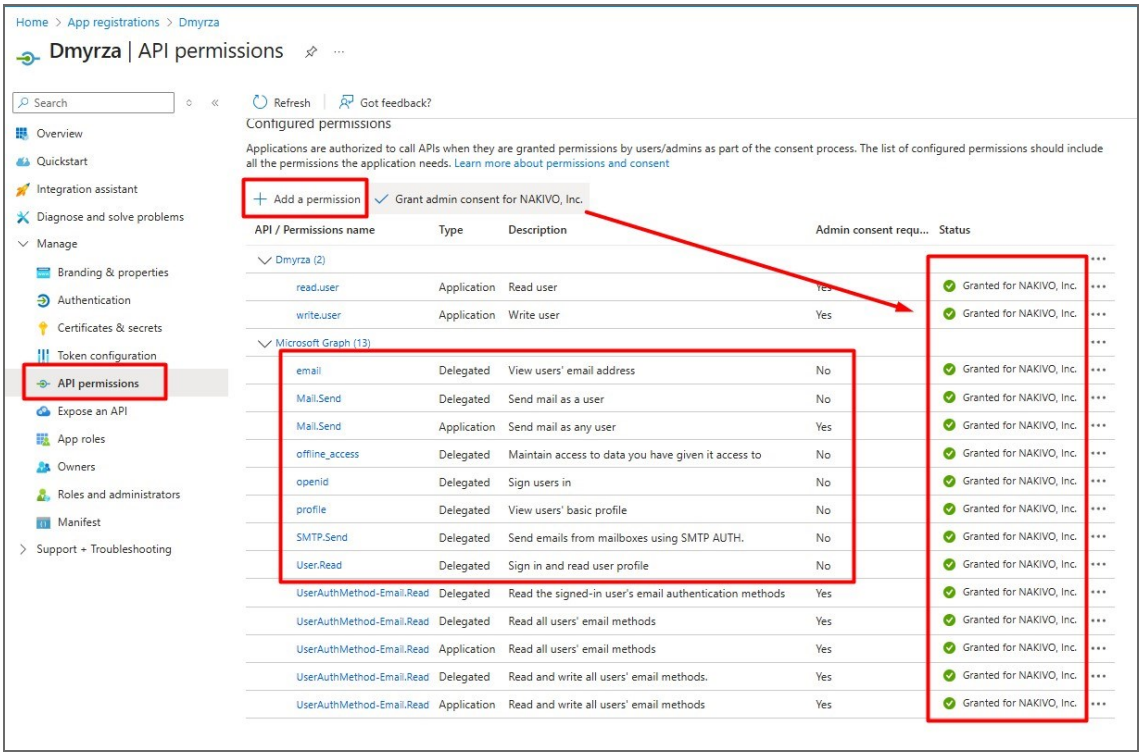
1. In the Azure Portal, go to **Certificates & secrets** and click **New client secret**.
2. Specify a description and expiration period, then click **Add**.
3. Copy the generated value immediately and paste it into NAKIVO Backup & Replication.



Grant Required API Permissions

1. Open the API permissions tab.
2. Add the following Microsoft Graph permissions:

Permission	Type	Description
Mail.Send	Delegated	Send mail as a user
Mail.Send	Application	Send mail as any user
offline_access	Delegated	Maintain access to data you have given it access to
openid	Delegated	Sign users in
profile	Delegated	View users' basic profile
SMTP.Send	Delegated	Send emails from mailboxes using SMTP AUTH.
User.Read	Delegated	Sign in and read user profile



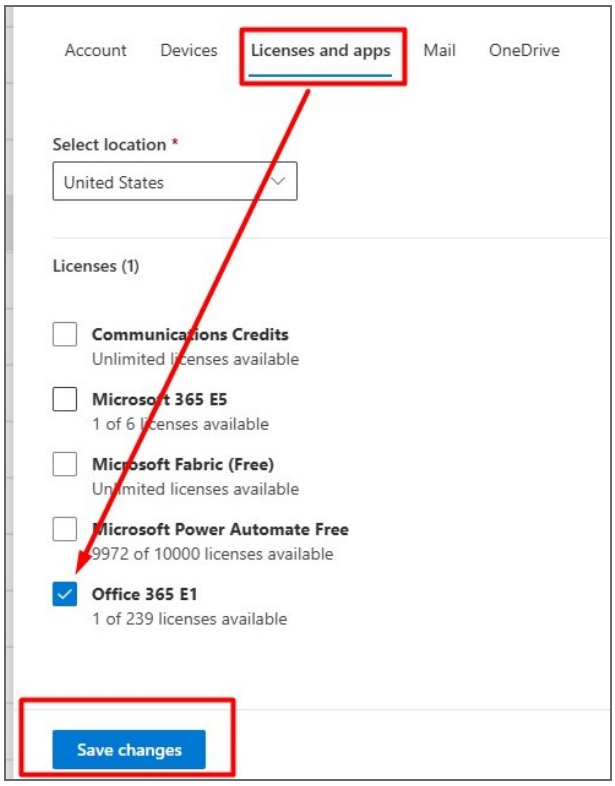
3. Click **Grant admin consent** to activate the permissions. For the full list of permission descriptions, refer to Microsoft Graph permissions reference.

Verify User Licensing

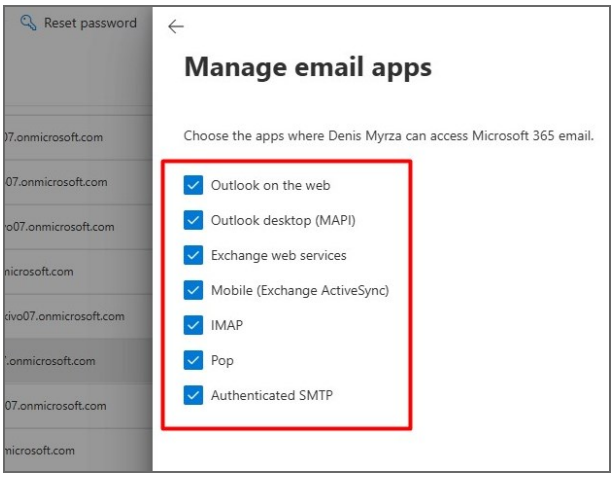
To authenticate SMTP via OAuth, the user account must have a valid Exchange-enabled license.

1. Open Microsoft Admin page.
2. Go to **Users** > **Active users**.
3. Select your user and verify the **License** status.
4. If the account is **Unlicensed**, assign **Office 365 E1** or any other plan that includes **Exchange Online**.

> **Note**
> This is required to connect Exchange and the Authenticated SMTP function.



5. After licensing, reload the page.
6. Open **Mail** > **Email apps** and click **Manage email apps**.
7. Ensure that **Authenticated SMTP** is enabled for the user. For more information, refer to Enable or disable authenticated client SMTP submission in Exchange Online.



In NAKIVO Backup & Replication

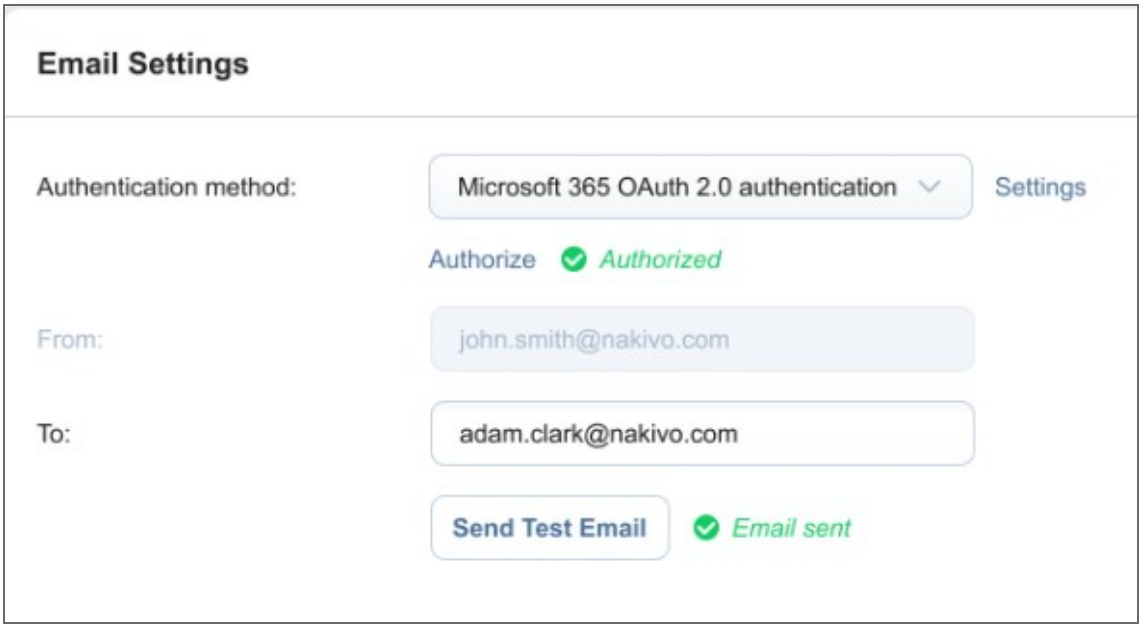
1. Navigate to **Settings** > **Email Settings**.
2. Under **Authentication method**, select **Microsoft OAuth 2.0 authentication**.
3. Enter the following information:

- **Client ID**: Enter the unique identifier generated with your Microsoft 365 account when you create OAuth 2.0 credentials.
- **Tenant ID**: Enter your Tenant ID (your Tenant ID is displayed in the Tenant information section in Microsoft Entra Admin center).
- **Client Certificate**: Choose a saved certificate from the drop-down list to use an existing certificate. If you don't have any saved certificates or want to generate a new one, click **Manage certificates** and create a new certificate. Make sure that the downloaded certificate file has the .cer format.
- **Client Secret**: Select and enter the **Client Secret**, generated on the Microsoft Azure portal.

4. Add the certificate to **Microsoft Azure** > **App Registration** > **Certificates & secrets** > **Upload certificate**.
5. In Nakivo Backup & Replication, click **Authorize now** to get access to your Microsoft 365 account through OAuth 2.0 authentication.
6. Enter the email address in the **To** field.
7. Click **Send Test Email** to confirm that the configuration works.

Note

If OAuth settings are incomplete or Azure permissions are missing, sending a test email from NAKIVO Backup & Replication may fail. Complete all Microsoft 365 configuration steps before attempting authentication.

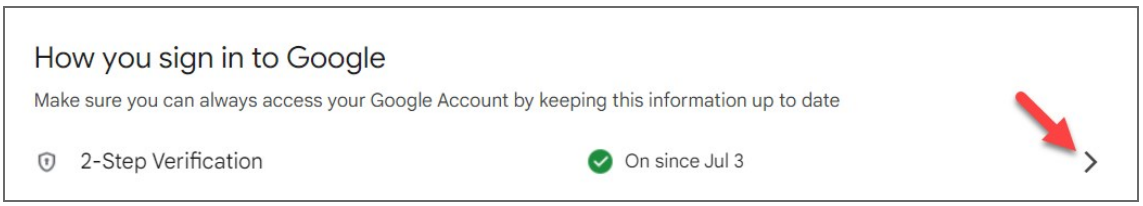


Using Email with Two Factor Authentication

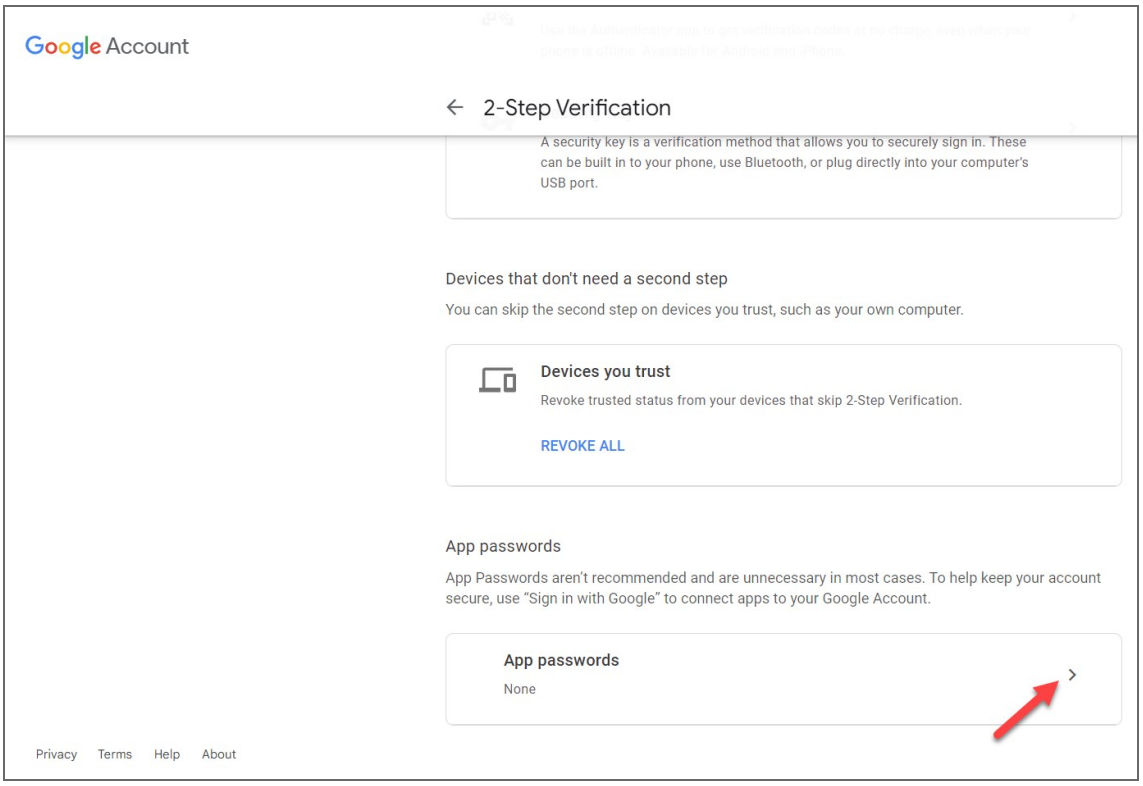
Many email providers require that the third-party apps or devices use two-factor authentication to sign into their accounts. Therefore, for correct email notification setup, you must grant access permissions to NAKIVO Backup & Replication via your account security settings and generate a unique **SMTP password**.

As an example, you can follow the instructions for setting up a Gmail account notifications below:

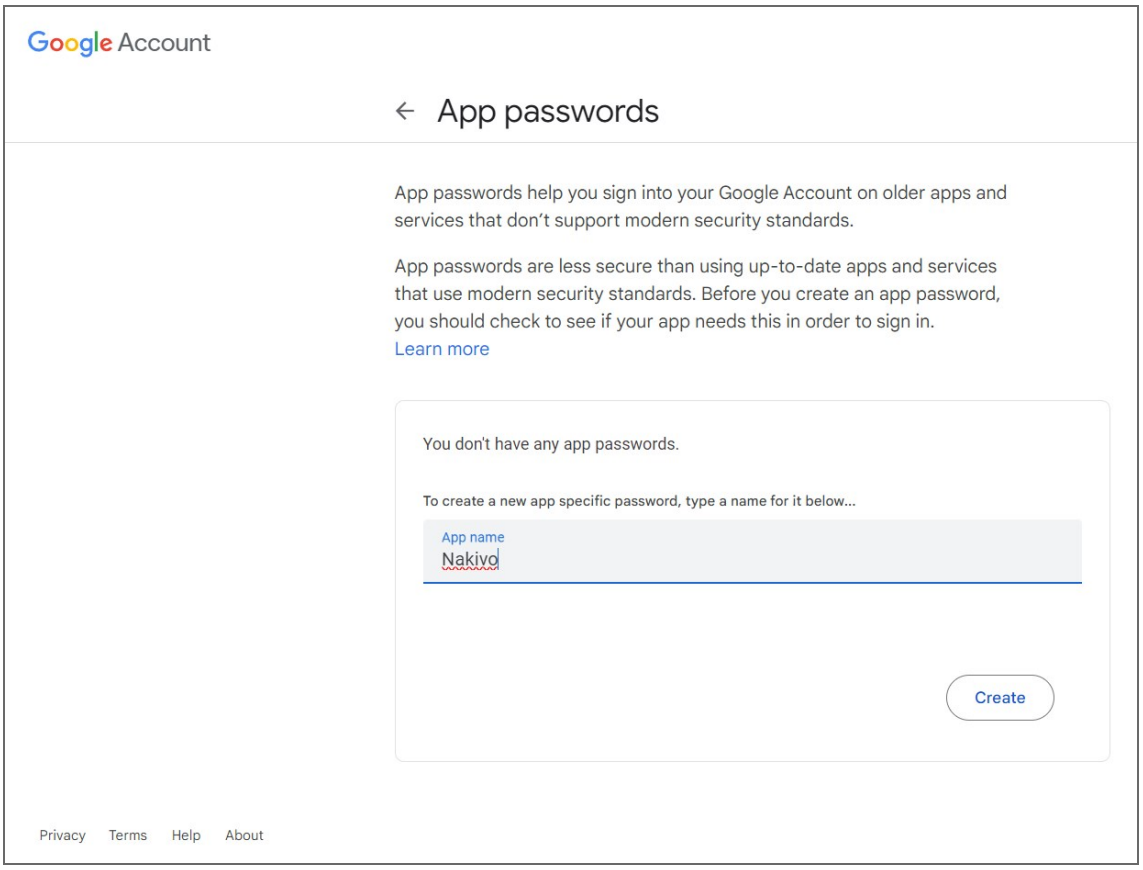
1. Log into your Google account and allow 2-step verification (you will be asked to complete a second step to verify it's you when you sign in):



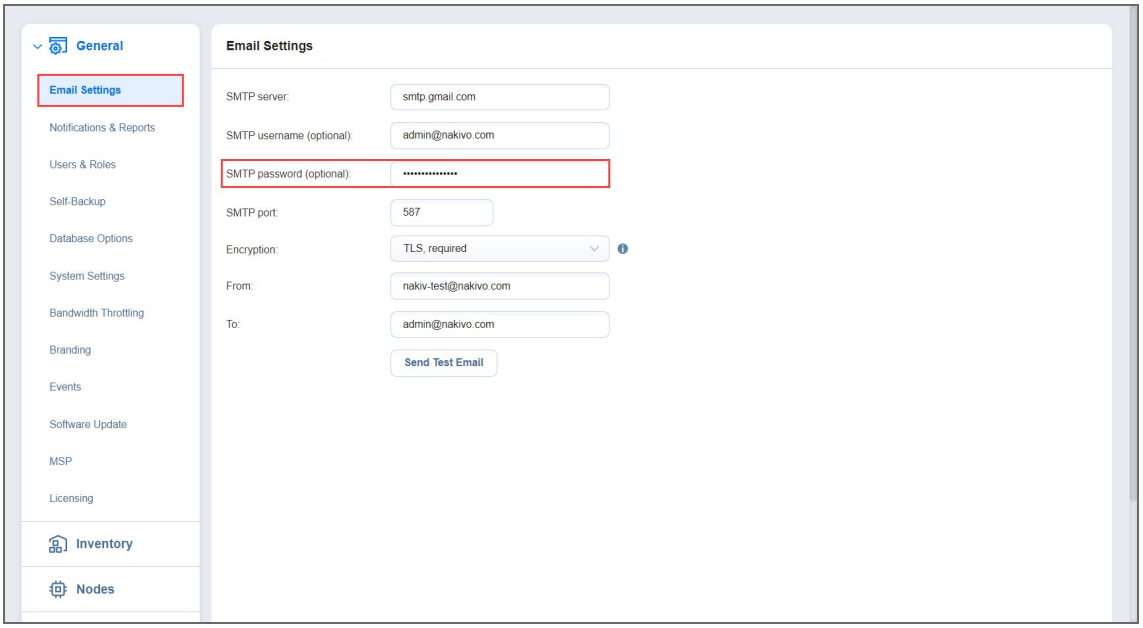
2. Scroll down to the *App passwords* section and click the arrow button to proceed with the password.



3. Type your app name and click **Create** to generate the password.



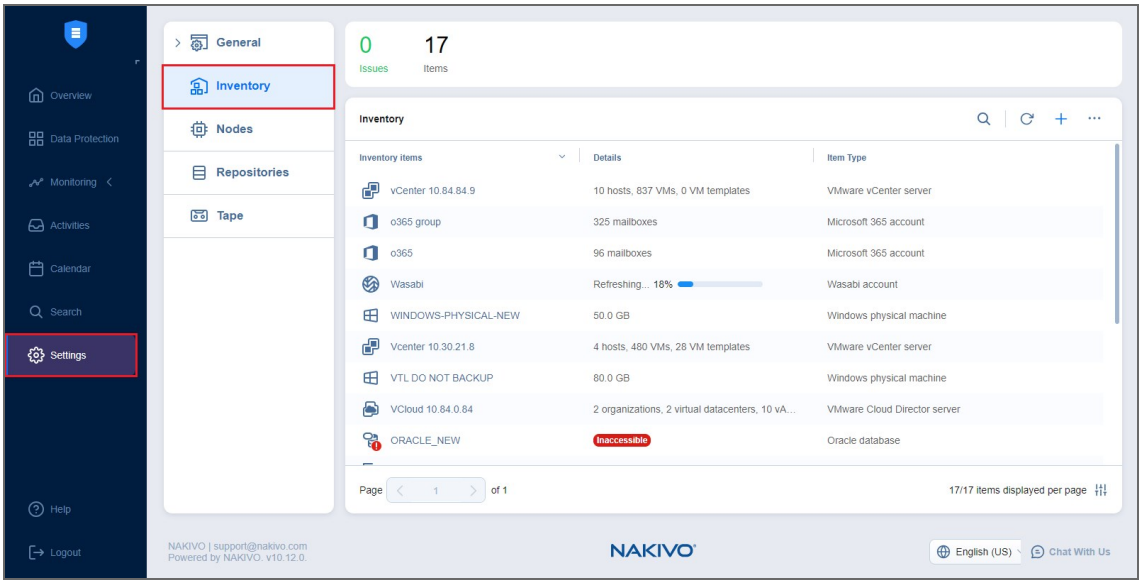
4. Now, copy the generated password and paste it into the **SMTP password** field in the **Settings > General > Email Settings** tab.



Inventory

Prior to creating backup, replication, or recovery jobs, you need to add your virtual/cloud/physical infrastructure, Microsoft 365 account, Oracle database, or supported storage device to the product's **Inventory**. The discovered item is added to the internal product database, which is refreshed every 1 hour by default. The **Inventory** tab contains a **Summary** bar, which offers an overview of all **Inventory** items. The data displayed is as follows:

- **Issues**: Total number of issues/alarms related to **Inventory** items
- **Items**: Total number of items in the **Inventory**



Refer to the following sections to learn more about adding and managing **Inventory** items:

- "Adding Microsoft 365 Accounts" on page 404
- "Managing Inventory" on page 416
- "Managing Credentials" on page 425

Adding Microsoft 365 Accounts

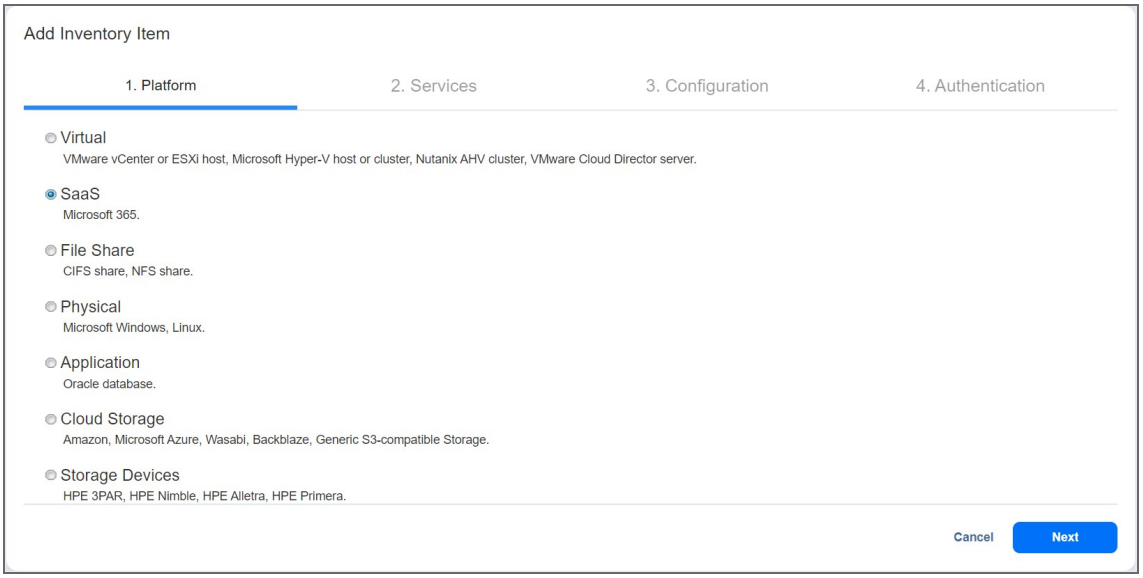
Before you start backing up items from your organization's Microsoft 365 account, you need to add the Microsoft 365 account to Inventory in NAKIVO Backup & Replication.

- Adding Microsoft 365 Account to Inventory
- Activating Access to Microsoft Graph Teams Export APIs
- Obtaining Microsoft 365 Credentials

Adding Microsoft 365 Account to Inventory

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

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



5. On the **Services** page, provide the following information:
   - **Display name**: Enter a name for the account.
   - **Services**: Select the Microsoft 365 services that you want to back up:
     - **Exchange Online Users**
     - **OneDrive for Business**

       Note

       OneDrives without added email addresses or assigned licenses are ignored.

     - **Exchange Online Groups**

- **SharePoint Online**

  Note

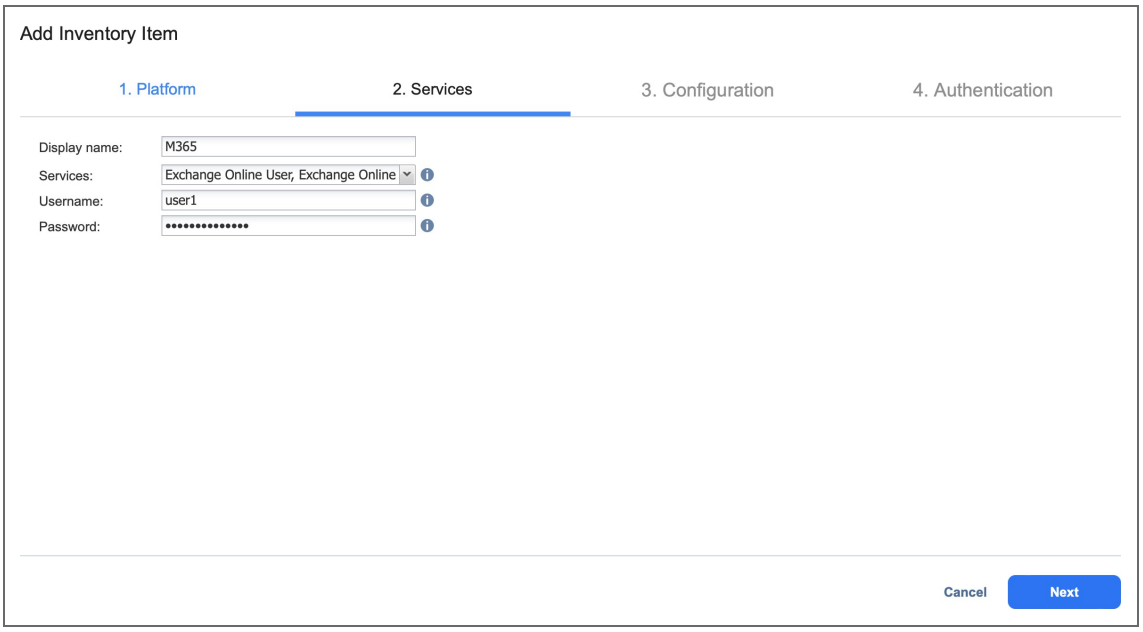
  Optionally, select the **Use Modern Authentication with enabled MFA account** checkbox if MFA is enabled for your Microsoft 365 accounts.

- **Teams**

  Notes

  - Optionally, you can enable the **Use Teams Export APIs to back up Teams chat messages** to use Teams Export APIs instead of MS Graph APIs to back up Teams chat messages. Backup performance may be improved with this option. As of August 25, 2025, Microsoft has confirmed that Teams Export APIs are free to use and no longer incur additional charges.
  - For NAKIVO Backup & Replication to successfully discover **Teams**, the following conditions must be met:
    - Microsoft Teams Exploratory experience license or higher (access to Microsoft Teams API is required).
    - The administrator account must have the Global Administrator role or Team Administrator role assigned.
    - If you prefer not to use a Global Administrator role:
      1. Register the application and manually assign all necessary permissions. For more information, refer to Required API Permissions for Microsoft 365.
      2. In the **Add Inventory Item** wizard select **Use an existing AD Application** and deselect **Automatically grant required permissions**, and proceed with the setup.

6. Click **Next** to proceed.

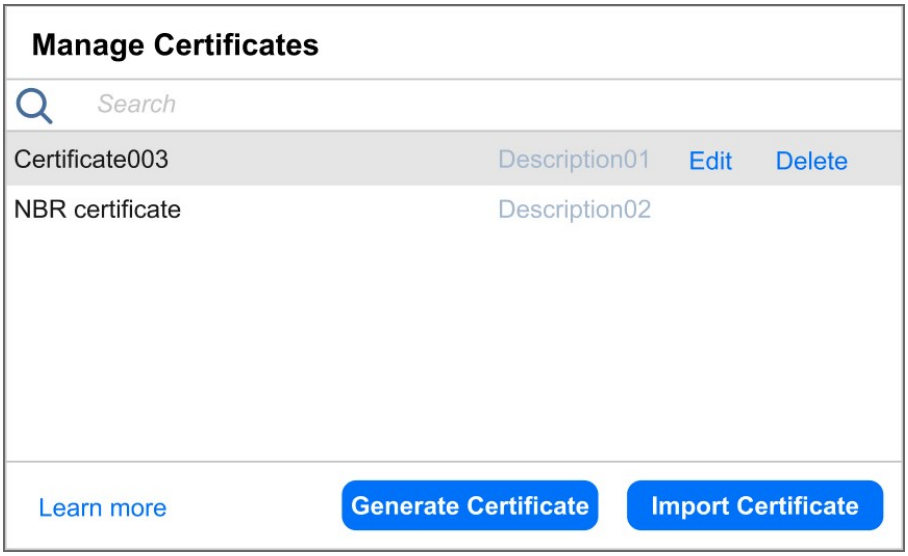


7. On the **Configuration** page, choose one of the following:
   - **Automatically register a new Azure AD application**: When this option is selected, NAKIVO Backup & Replication automatically registers the Microsoft 365 account as a new application along with a new self-signed certificate in Azure Active Directory and grants the required permissions after authentication.
   - **Use an existing Azure AD application**: When this option is selected, you can manually add an existing Microsoft 365 application already registered in Azure Active Directory.

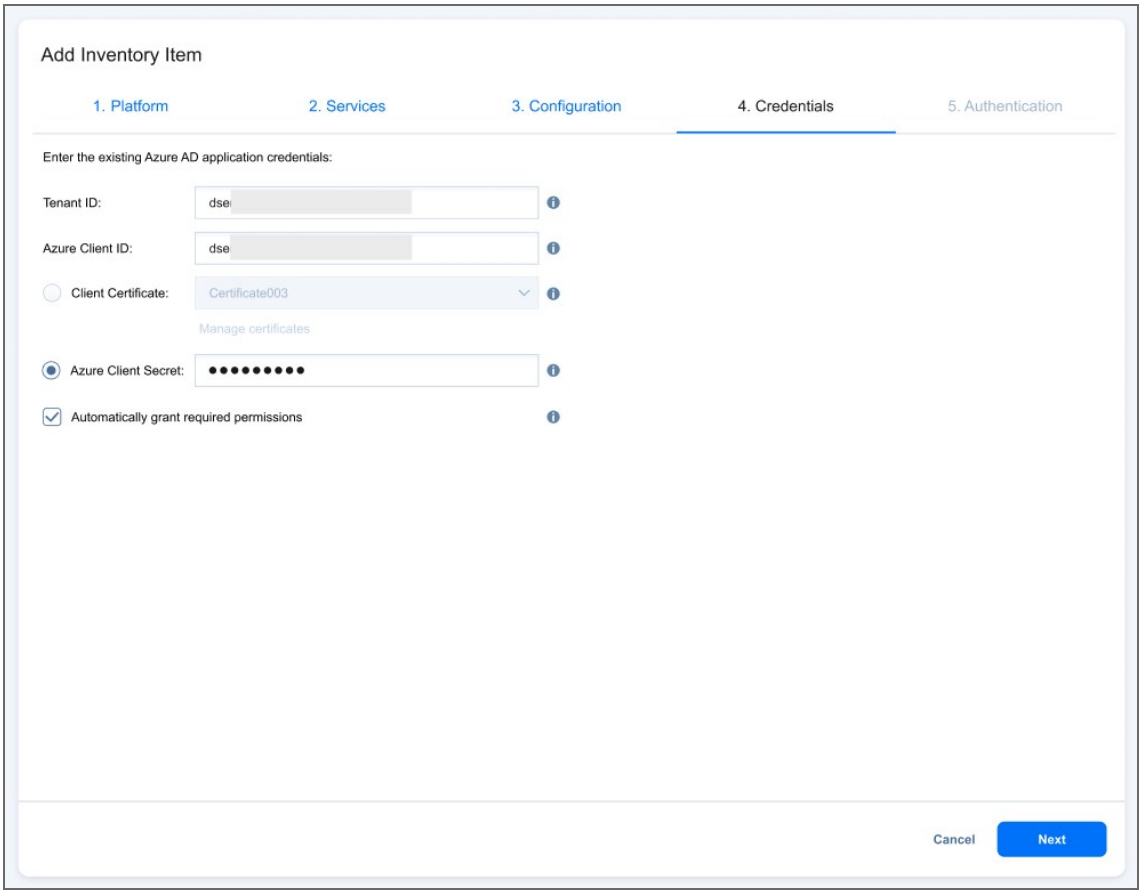
   > **Note**
   > In order to add a Microsoft 365 account to NAKIVO Backup & Replication, the account has to be registered in Azure Active Directory.

8. Click **Next** to proceed. If you selected **Automatically register a new Azure AD application**, skip to step 12.
9. The **Credentials** page opens if you have chosen **Use an existing Azure AD application** option on the **Configuration** page. Provide the following information:
   - **Tenant ID**: Enter your Azure Tenant ID created when registering your Microsoft 365 account in the Azure Portal.
   - **Azure Client ID**: Enter your Azure Client ID created when registering your Microsoft 365 account in the Azure Portal.
10. Choose one of the following authentication methods:

- **Client Certificate**: Choose a saved certificate from the drop-down list to use an existing certificate. If you don't have any saved certificates or want to generate a new one, click the **Manage certificates** button. In the resulting pop-up, you can edit and delete existing certificates, generate new certificates, or import existing certificates.
    - To generate a certificate, click the **Generate Certificate** button. In the **Generate New Self-Signed Certificate** pop-up, enter a display name and (optionally) a description for the certificate. Clicking **Save** will generate a new self-signed certificate and add it to the list of saved certificates.
    - To import an existing certificate, click the **Import Certificate** button. In the **Import Certificate** pop-up, enter a display name for the certificate and click the **Browse…** button to upload a certificate in the required .pfx format. Then, enter the certificate password if required, and (optionally) add a description for the certificate. Clicking **Save** will add the certificate to the list of saved certificates.



- **Azure Client Secret**: Enter your Azure Client Secret obtained from the Azure Portal. For more information on obtaining the Azure credentials, refer to the Obtaining Microsoft 365 Credentials below. Note that this option is disabled if **Use Modern Authentication with enabled MFA account** is enabled on the **Services** page; you must use **Client Certificate** in this case.



11. Enter the administrator **Username** and **Password** required for **SharePoint Online**. This is only required if you have provided the Azure Client Secret.
12. Depending on the authentication method you selected, you can optionally select **Automatically grant required permissions** (for Azure Client Secret authentication) or **Automatically grant required permissions and register certificate** (for Client Certificate authentication). This allows NAKIVO Backup & Replication to automatically check the required API permissions for the selected services of your organization's Microsoft 365 account and to add them if they are missing. If you do not select this option, you have to manually grant the required permissions. If this option is selected for Client Certificate authentication, NAKIVO Backup & Replication will also automatically register the selected certificate with the specified application, if it is not already registered. If you do not select this option, you have to manually register the certificate with the application. If you have selected either option, click **Next**. Alternatively, click **Finish** to complete adding the item to the **Inventory**.
13. If you have selected **Automatically grant required permissions** or **Automatically grant required permissions and register certificate** on the **Credentials** page or **Automatically register a new Azure AD** application on the **Configuration** page, the **Authentication** page opens. Do the following:
    - If you selected **Automatically register a new Azure AD application** on the **Configuration** page, enter a name for the application in the **Application name** field.

- If you selected **Automatically register a new Azure AD application** on the **Configuration** page, add a client certificate as described in step 10.
- The **Code** field shows the code that needs to be entered at https://microsoft.com/devicelogin.
- Clicking the link will open the page in a new tab and you will be required to log in to the Microsoft Azure Cross-Platform Command-Line Interface (xplat-cli) using the provided code.
- If the provided credentials are correct, the account will be authenticated and you will be able to continue with the discovery process.

Notes

- While using a Global Administrator account is recommended for automatic permissions setup, you can use a non-Global Admin account if you manually assign the required Microsoft 365 permissions listed in Required API Permissions for Microsoft 365.
- If **SharePoint Online** was selected on the **Services** page, the SharePoint Administrator role must be assigned to the account in Microsoft 365 admin center for the authentication process to be successful.
- It is possible to bypass the **Authentication** step if the IP address or hostname of your NAKIVO Backup & Replication installation location is added as a trusted location in Azure Active Directory. For more information, refer to this page.
- If **Automatically register a new Azure AD application** is selected in the **Configuration** step, then the credentials to the newly registered application will be downloaded to the browser after successful authentication.
- When generating a new certificate, the new self-signed certificate will be downloaded to the browser in .pfx format after clicking **Save** in the **Generate New Self-Signed Certificate** pop-up.
- If NAKIVO Backup & Replication is updated from a version that did not include support for Microsoft 365 Groups to a version that does, it is possible for the group mailboxes and group sites to be discovered automatically for existing Microsoft 365 accounts:
  - If you have selected **Automatically grant required permissions** on the **Credentials** page, the group mailboxes and group sites will be added automatically with all of the required permissions granted.
  - If you have not selected **Automatically grant required permissions** on the **Credentials** page, the group mailboxes and group sites will remain undiscovered by the solution.
- If NAKIVO Backup & Replication is updated from a version that did not include support for Microsoft Teams to a version that does, Teams will not be discovered automatically. In order to discover Teams, you will need to edit the appropriate Microsoft 365 account in Inventory and add **Teams** as a service.

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

Activating Access to Microsoft Graph Teams Export APIs

To back up Teams chat messages using Teams Export APIs, the following requirements must be met:

- You have an active Azure subscription.
- You have an active app registration in Azure.
- You have created a resource group as described here.
- Make sure Azure CLI is installed on the running machine.
- Run the command line and enter the required information to enable an application to use metered APIs and services in Microsoft Graph and become associated with an Azure subscription.

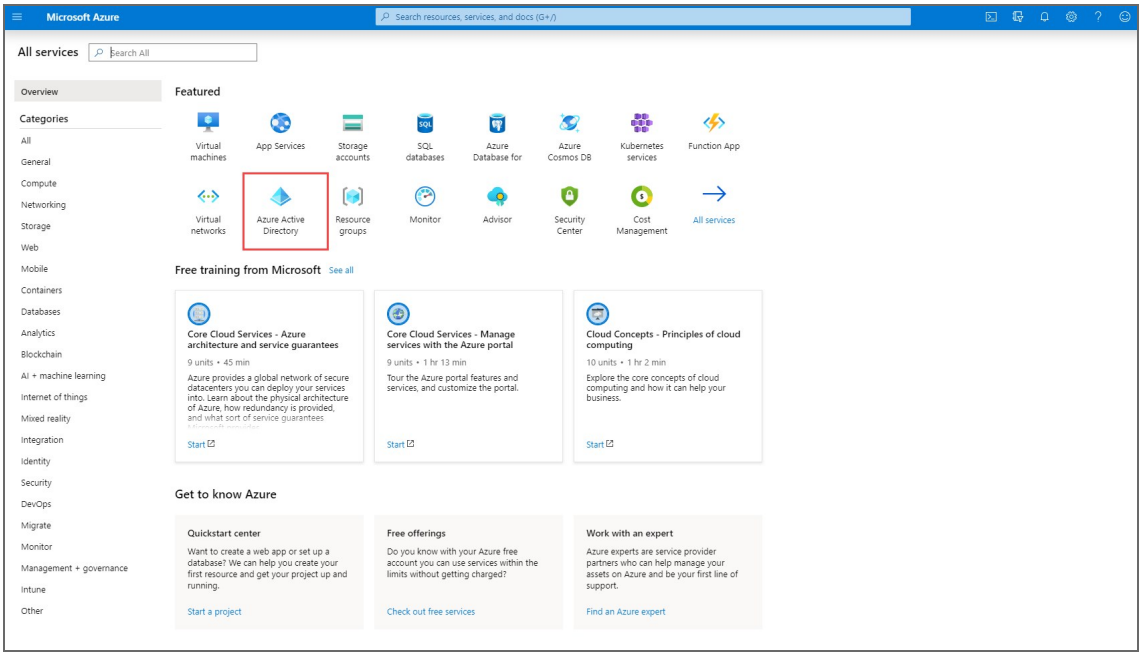
Notes

- As of August 25, 2025, Microsoft has confirmed that Teams Export APIs are free to use and no longer incur additional costs.
- The feature is not supported for Microsoft 365 US Government, Germany, and China.

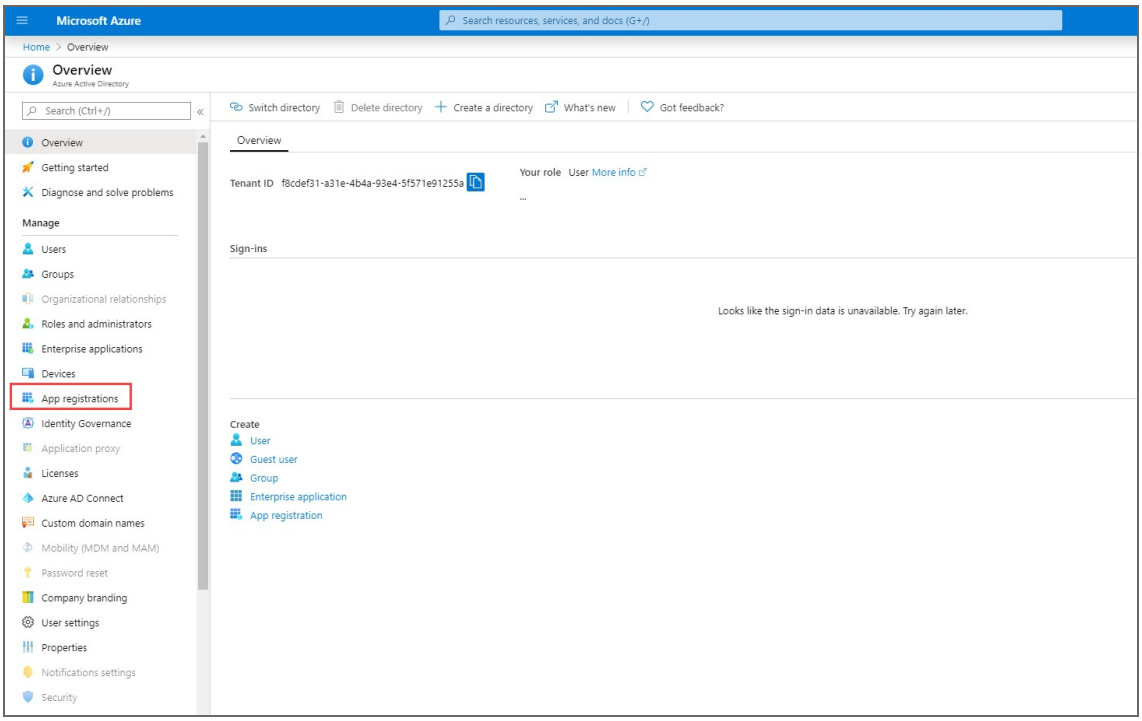
Obtaining Microsoft 365 Credentials

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

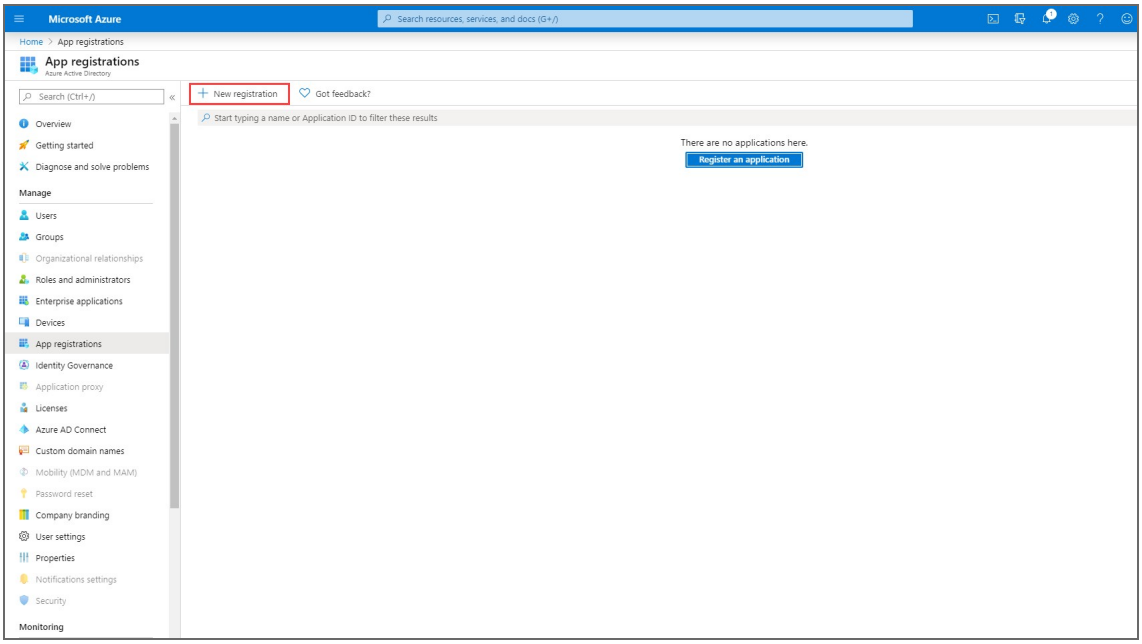
1. Open the Azure Portal by going to portal.azure.com
2. Sign in to Microsoft Azure with your Microsoft 365 account credentials.
3. Select **Azure Active Directory** from the **Dashboard** or from the Portal Menu.



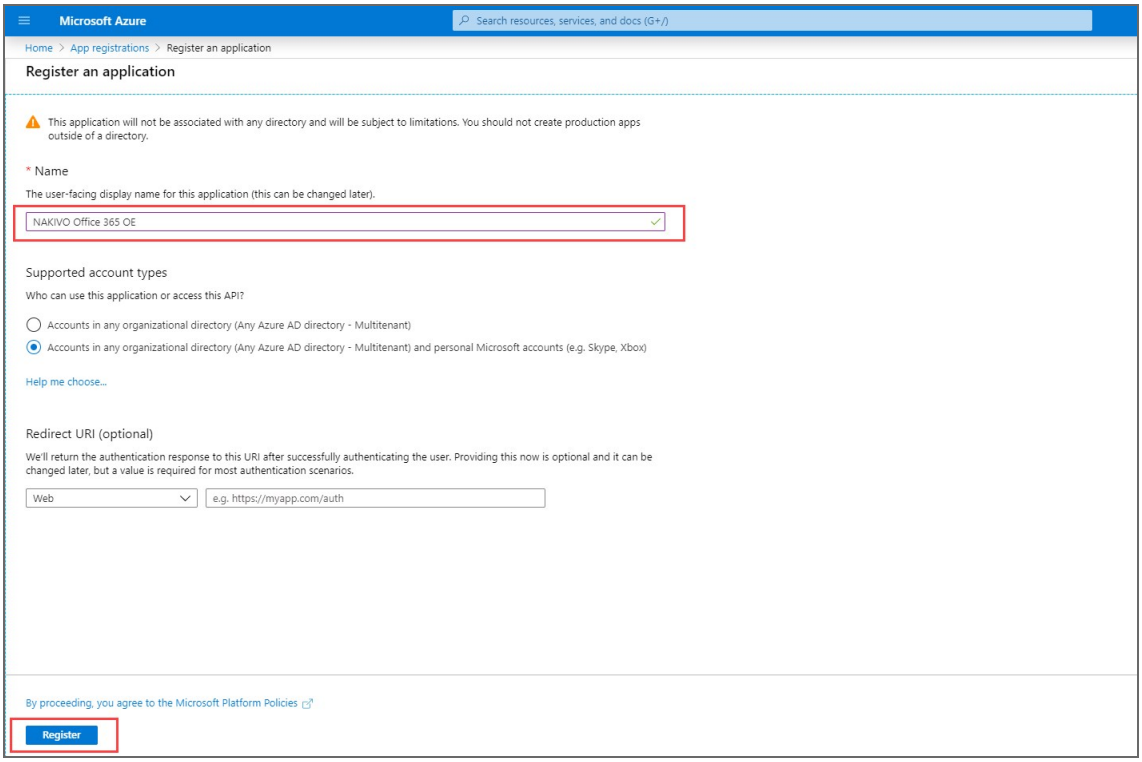
4. In the left menu, click **App registrations**.



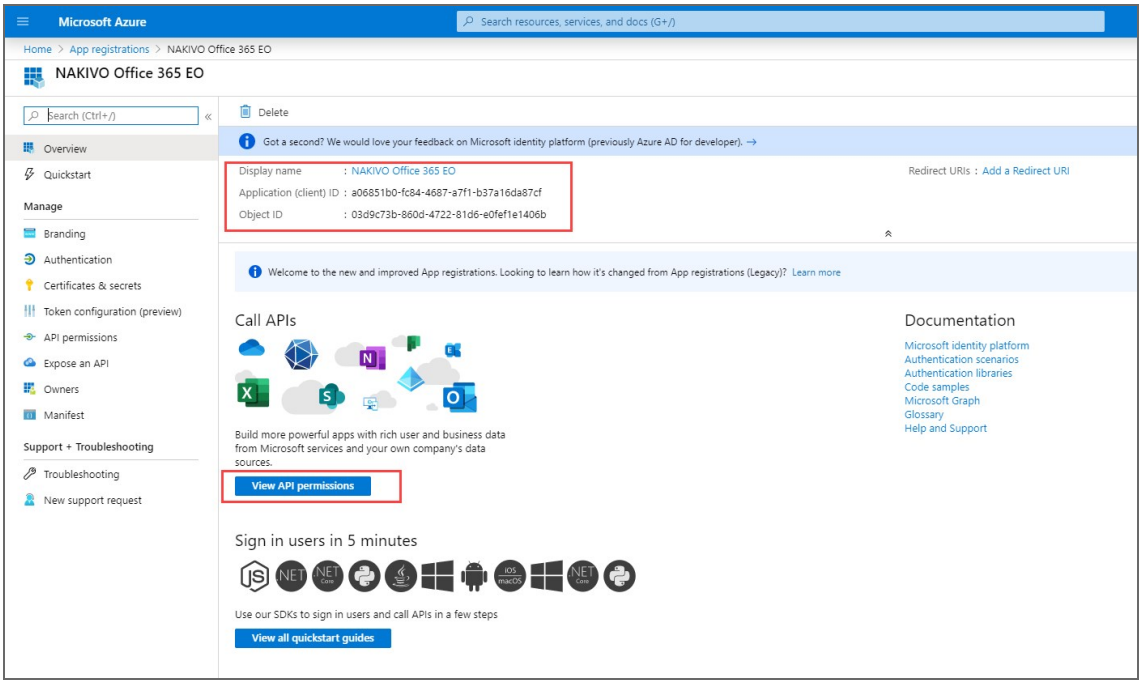
5. Click **New registration** on the **App registrations** page.



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



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

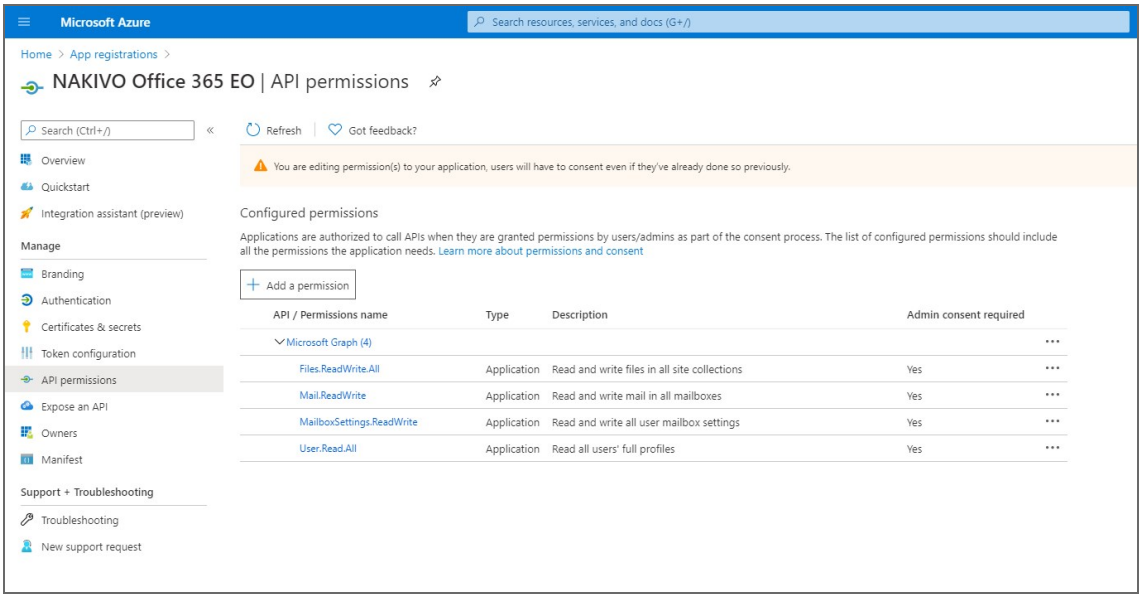


8. Click **Microsoft Graph**.
9. Click the **Application permissions** tab.
10. Provide the necessary API Permissions. Refer to Required API Permissions from Microsoft 365 for details.

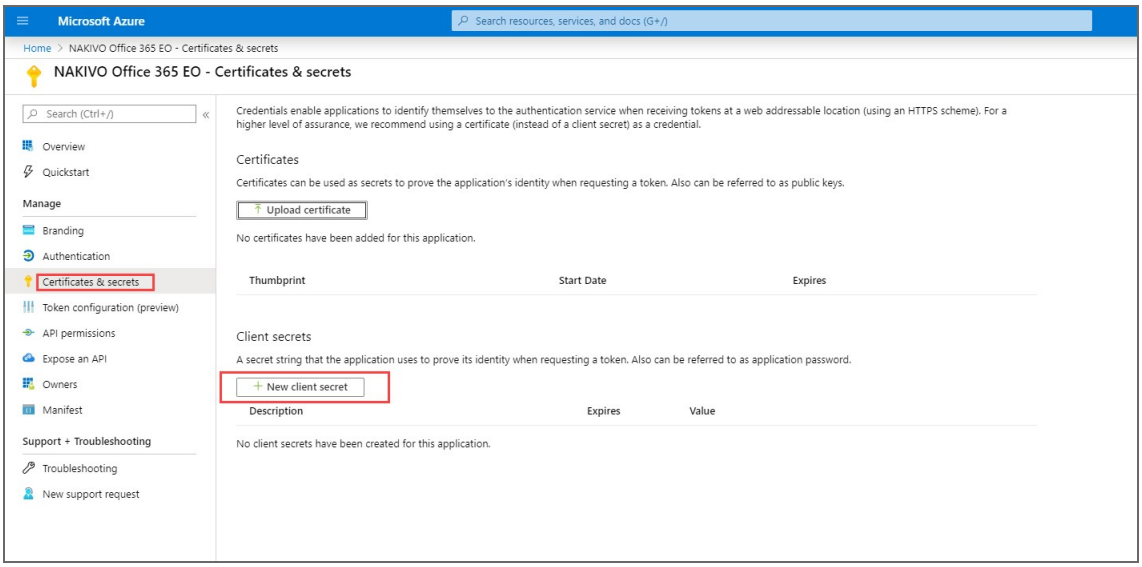
Notes

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

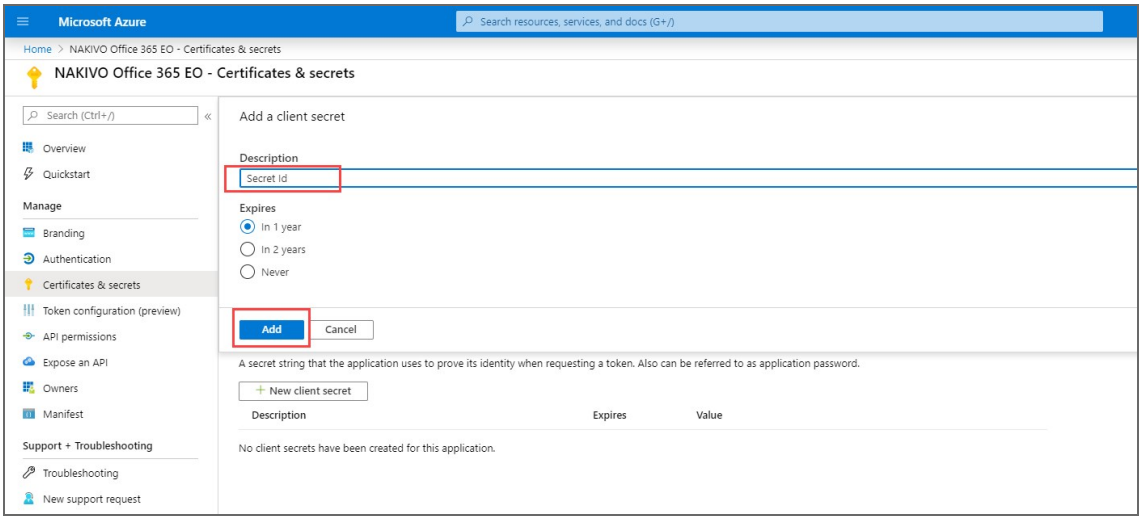
11. Click **Update Permissions**.



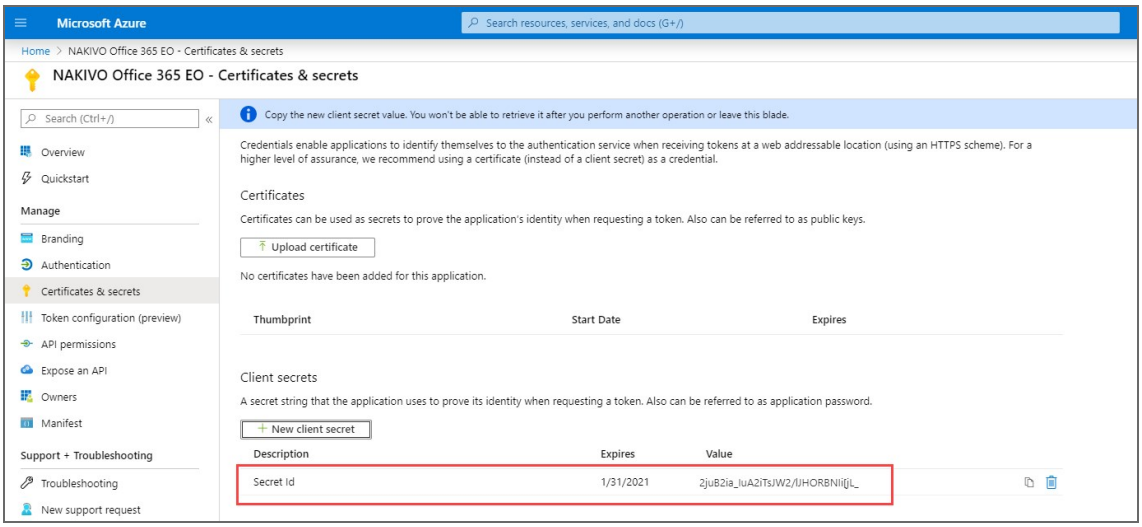
12. Click **Certificates & secrets > New client secret** to create a new client secret for your app.



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



The new **Client secret** is generated.



Managing Inventory

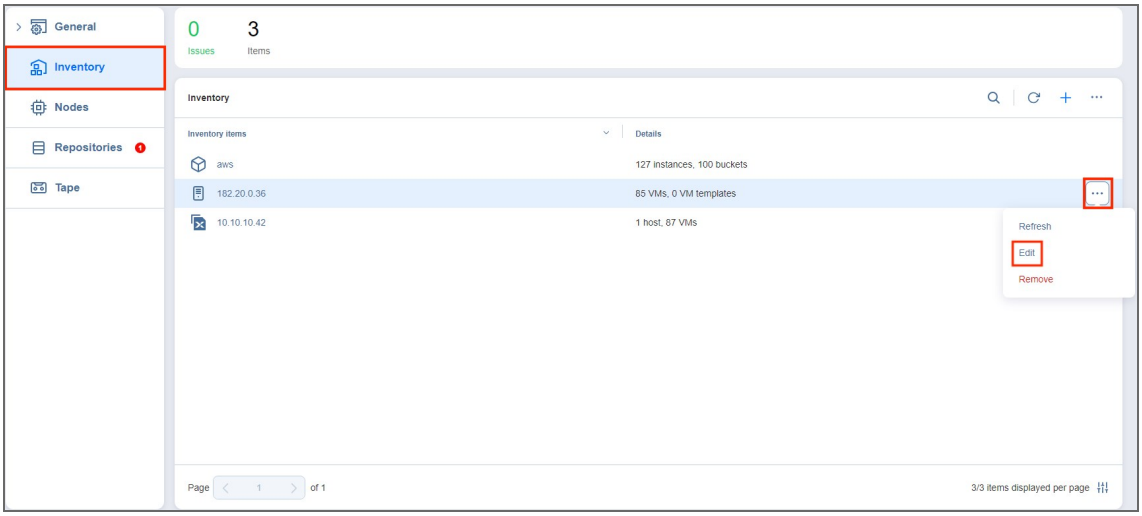
Refer to the following topics:

- “Editing Inventory Items” on page 417
- “Refreshing Inventory” on page 418
- “Removing Items from Inventory” on page 420

Editing Inventory Items

If the credentials of an inventory item are no longer correct, the connection to the **Inventory** item will be lost. To re-establish a connection, update the required fields in the product by following the steps below:

1. Click **Settings** in the left pane of the product.
2. Go to the **Inventory** tab.
3. Hover over the item you would like to edit.
4. Click ••• on the right side and then click **Edit**.



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

Refreshing Inventory

NAKIVO Backup & Replication keeps the information about the discovered infrastructure in its internal database, which is refreshed every 1 hour by default. During the **Inventory** refresh, the product collects all required information about your virtual infrastructure, such as a list of hosts and VMs, their power state, and so on.

Only one item can be refreshed at a time. If you have added multiple items to the inventory, they will remain in the queue until they are able to be refreshed. Refer to the sections below to learn how to refresh the discovered infrastructure.

- Changing Inventory Refresh Frequency
- Manually Refreshing All Inventory
- Manually Refreshing a Discovered Item

Changing Inventory Refresh Frequency

1. Click **Settings** in the left pane of the product.
2. Go to the **System setting > Auto refresh** tab.
3. Do either of the following:

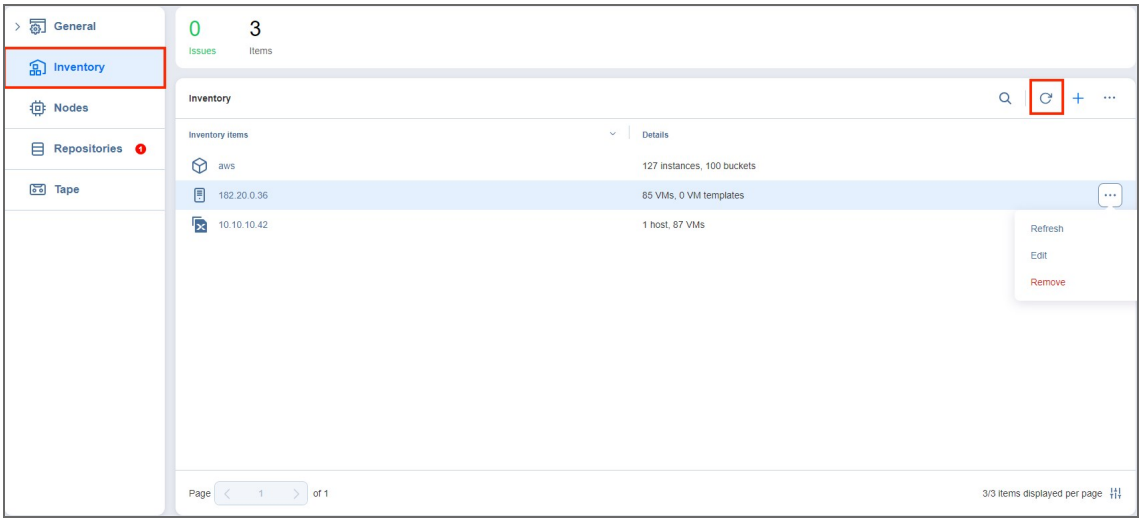
- To prevent the product from automatically refreshing the **Inventory**, deselect the **Refresh inventory every X [time period]** checkbox.
- To change the inventory refresh frequency, enter a new value in the **Refresh inventory every X [time period]** field (from 1 to 60 minutes or from 1 to 24 hours).

> **Note**
> New settings are applied instantly and do not need to be saved.

Manually Refreshing the Entire Inventory

To refresh all **Inventory** items, follow the steps below:

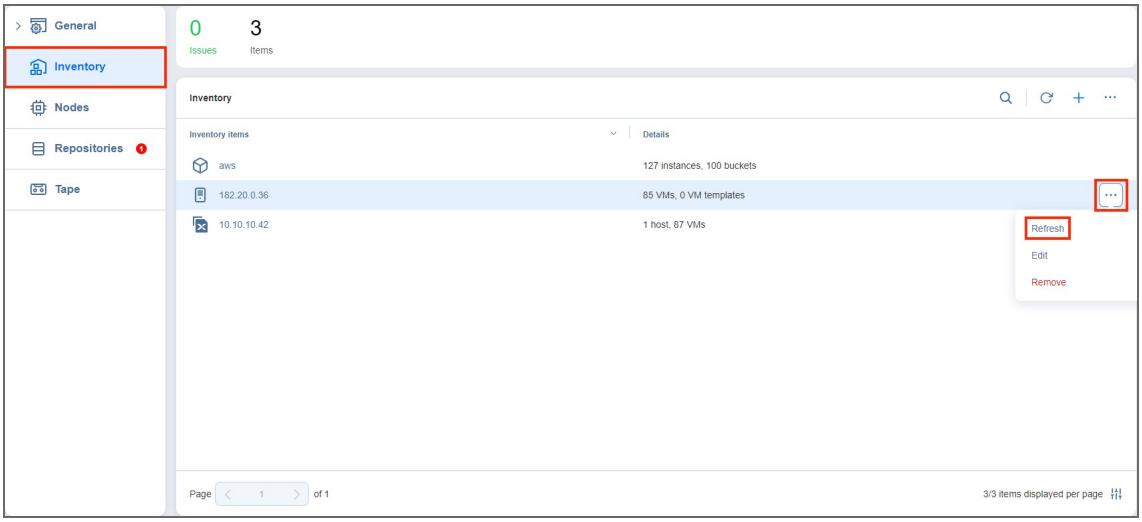
1. Click **Settings** in the left pane of the product and go to the **Inventory** tab.
2. Click the **Refresh All** button.



Manually Refreshing an Inventory Item

To refresh a single **Inventory** item, follow the steps below:

1. Click **Settings** in the left pane of the product and go to the **Inventory** tab.
2. Click the ••• button next to the item that you would like to refresh.
3. Click **Refresh**.

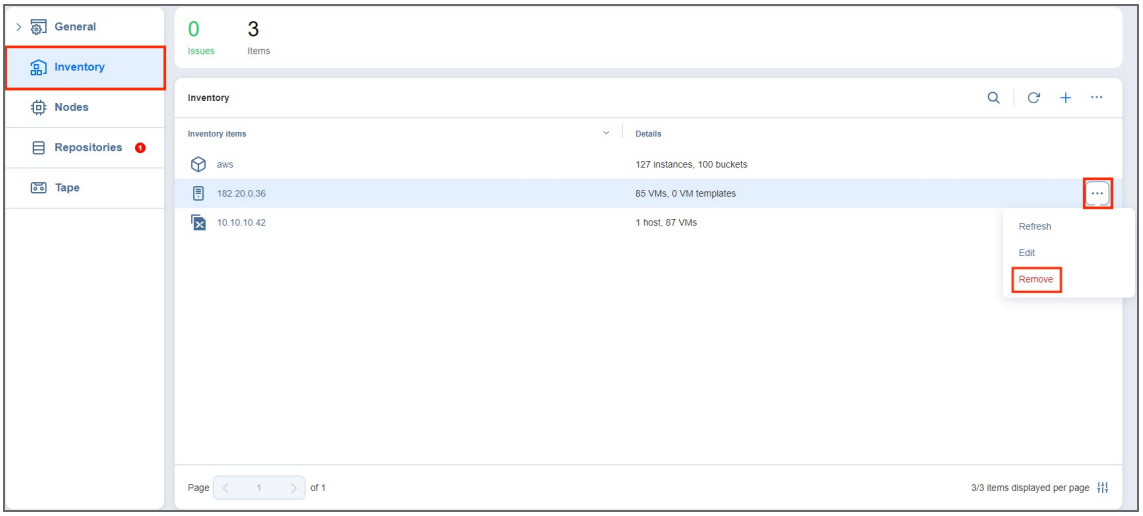


Removing Items from Inventory

You cannot remove an inventory item if there is at least one backup or replication job that uses the item or its children. In order to remove such items from the **Inventory**, you first need to delete (or edit) the corresponding jobs so no VMs/Instances are backed up or replicated on the host/server/account being removed.

To remove an item from the **Inventory**, follow the steps below:

1. Click **Settings** in the left pane of the product and go to the **Inventory** tab.
2. Hover over the item that you would like to remove from the **Inventory**.
3. Click **Manage** on the right side and click **Remove**.



Managing Certificates

NAKIVO Backup & Replication provides you with the ability to create and store certificates for Microsoft 365 authentication. Refer to the following topics:

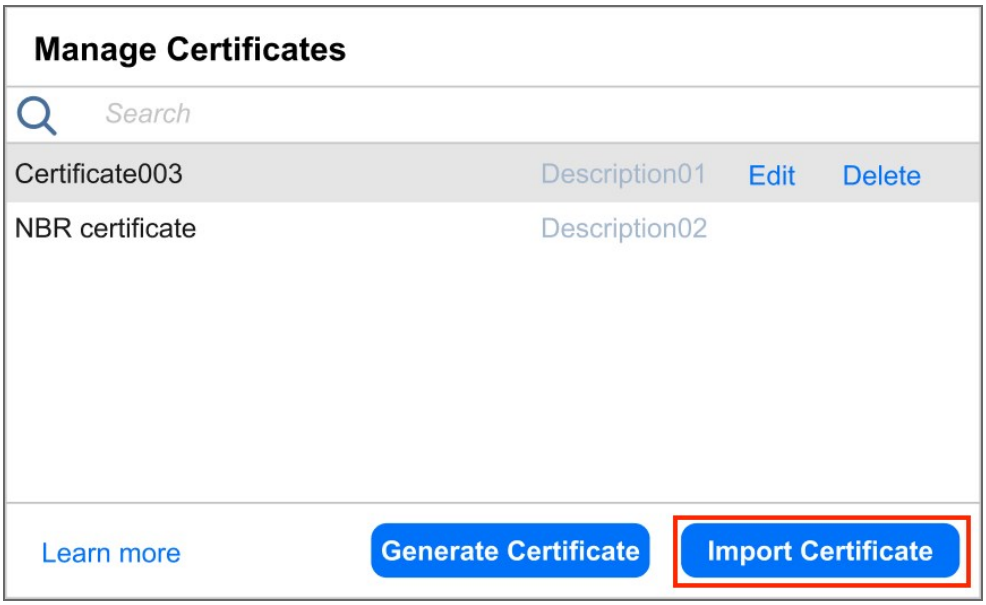
- Importing a Certificate
- Generating a New Certificate
- Editing Certificates
- Deleting Certificates

Importing a Certificate

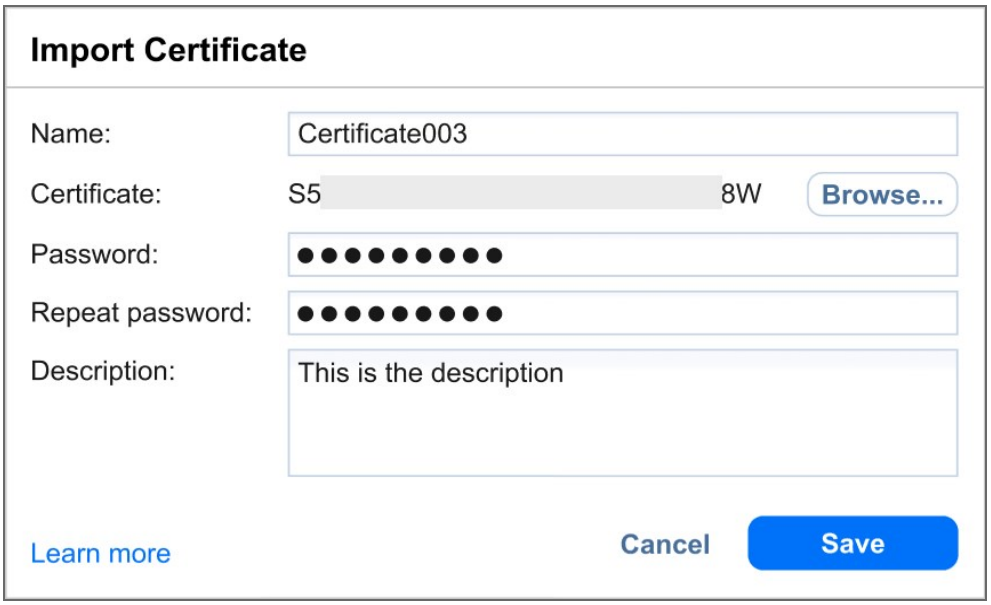
To import an existing certificate, do the following:

1. Click **Settings** in the left pane of the product.
2. Go to the **Inventory** tab.
3. Click **Manage**.
4. In the dialog that opens, click **Certificates**.

5. In the **Manage Certificates** dialog that opens, click **Import Certificate**.



6. In the resulting **Import Certificate** pop-up, enter a display name for the certificate in the **Name** field.
7. In the **Certificate** row, click **Browse…** and upload a certificate in the required `.pfx` format.
8. Enter the certificate password if needed in the **Password** and **Repeat password** fields.
9. Optionally, add a description for the certificate in the **Description** field.
10. Click **Save**.

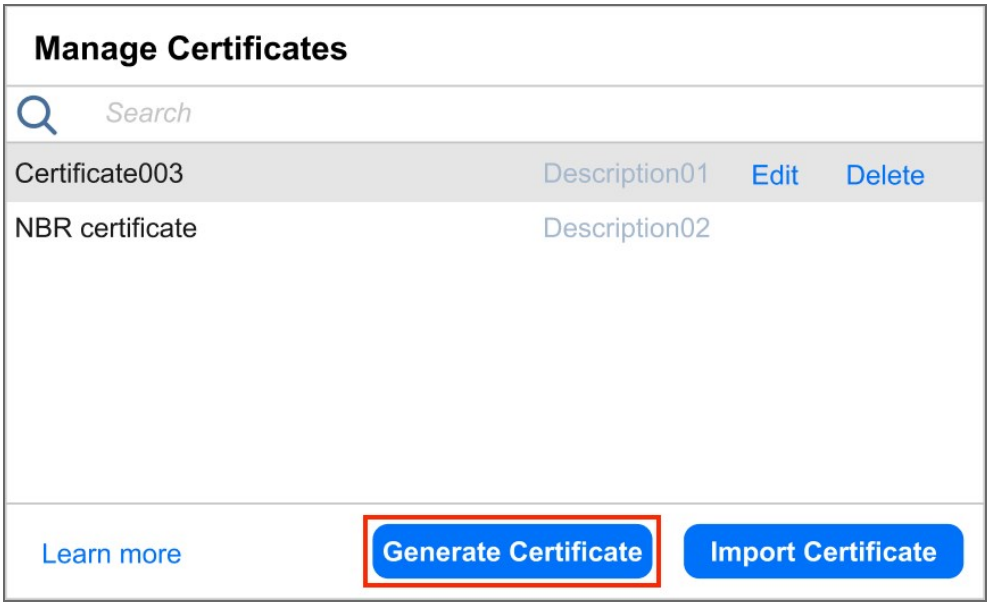


NAKIVO Backup & Replication will check if the certificate is valid. If it is, the certificate will become available for Microsoft 365 authentication.

Generating a New Certificate

To generate a new self-signed certificate, do the following:

1. Click **Settings** in the left pane of the product.
2. Go to the **Inventory** tab.
3. Click **Manage**.
4. In the dialog that opens, click **Certificates**.
5. In the **Manage Certificates** dialog that opens, click **Generate Certificate**.



6. In the resulting **Generate New Self-Signed Certificate** pop-up, enter a display name for the certificate in the **Name** field.
7. Optionally, add a description for the certificate in the **Description** field.
8. Click **Save** to generate a new self-signed certificate and add it to the list of saved certificates. NAKIVO Backup & Replication will automatically download the new certificate to the browser.

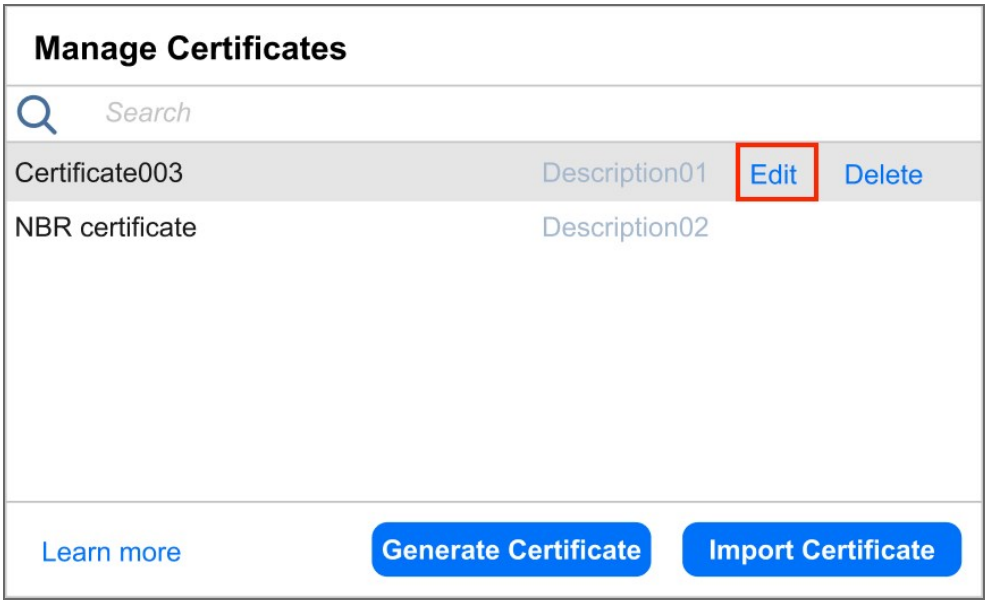


You can now use this certificate for Microsoft 365 authentication.

Editing Certificates

To edit a certificate, do the following:

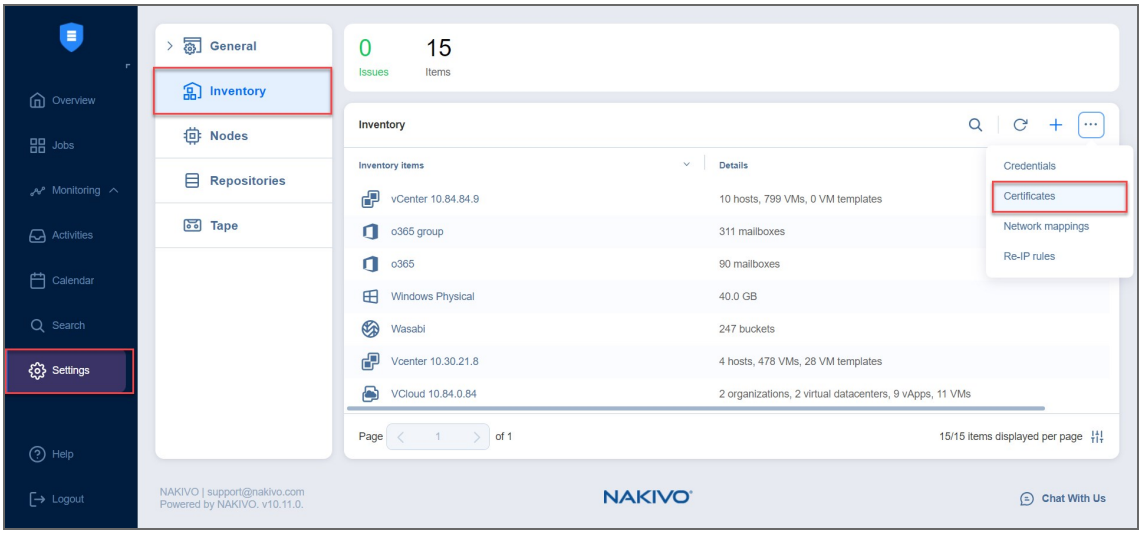
1. Click **Settings** in the left pane of the product.
2. Go to the **Inventory** tab.
3. Click **Manage**.
4. In the dialog that opens, click **Certificates**.
5. In the **Manage Certificates** dialog box that opens, hover over the certificate you want to edit and click **Edit**.
6. Edit any necessary fields and click **Save** when done.



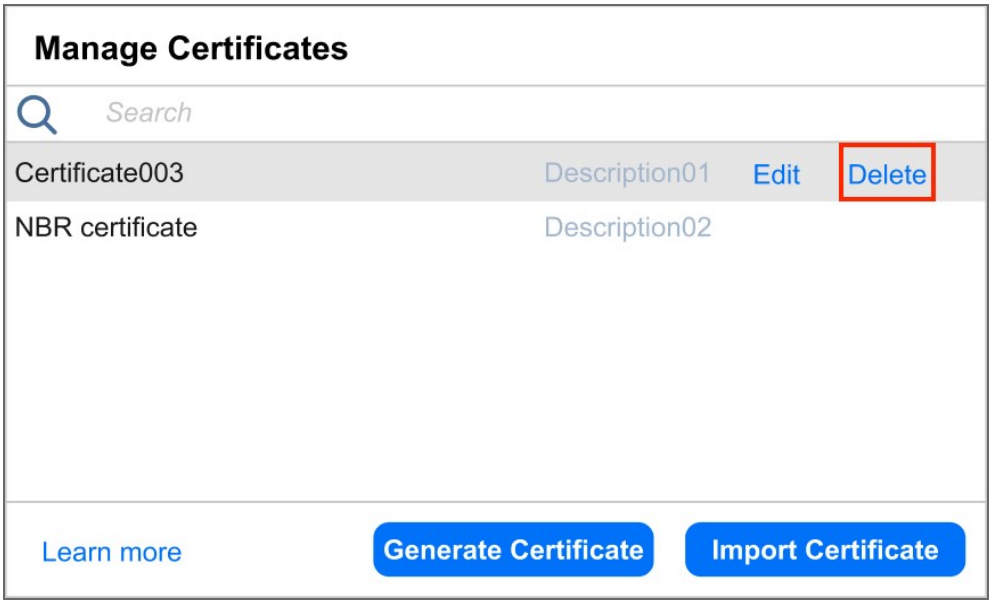
Deleting Certificates

To delete a certificate, do the following:

1. Click **Settings** in the left pane of the product.
2. Go to the **Inventory** tab.
3. Click **Manage**.
4. In the dialog that opens, click **Certificates**.



5. In the **Manage Certificates** dialog box that opens, hover over the certificate you want to delete and click **Delete**.
6. In the confirmation pop-up that opens, click **Delete** to confirm deletion.



Managing Credentials

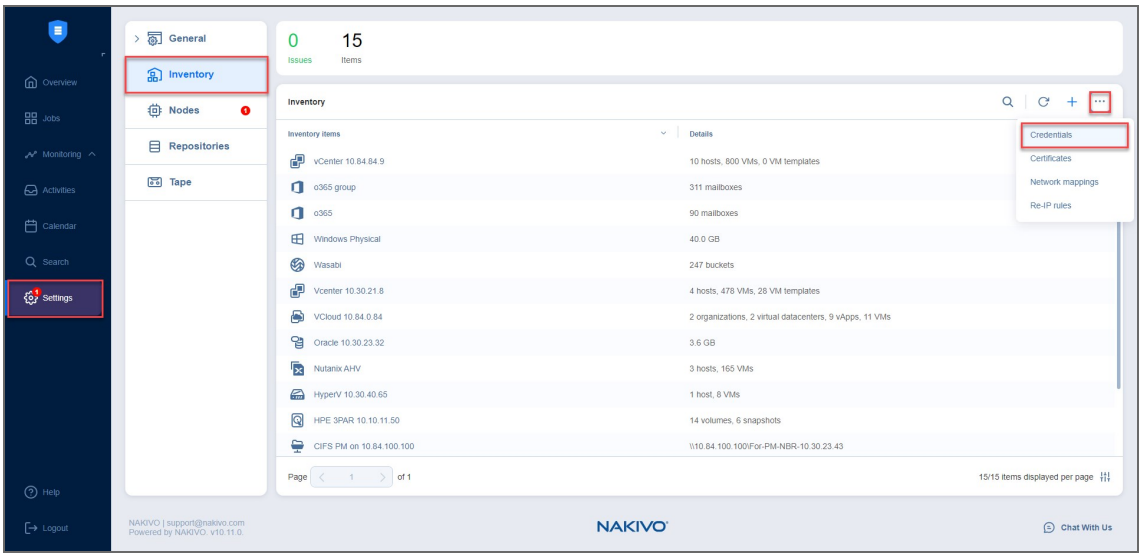
NAKIVO Backup & Replication provides you with the ability to store your OS login and password, Amazon EC2 instance private keys, and SSH keys to your Linux machines. Refer to the following topics:

- Adding Credentials
- Editing Credentials
- Deleting Credentials

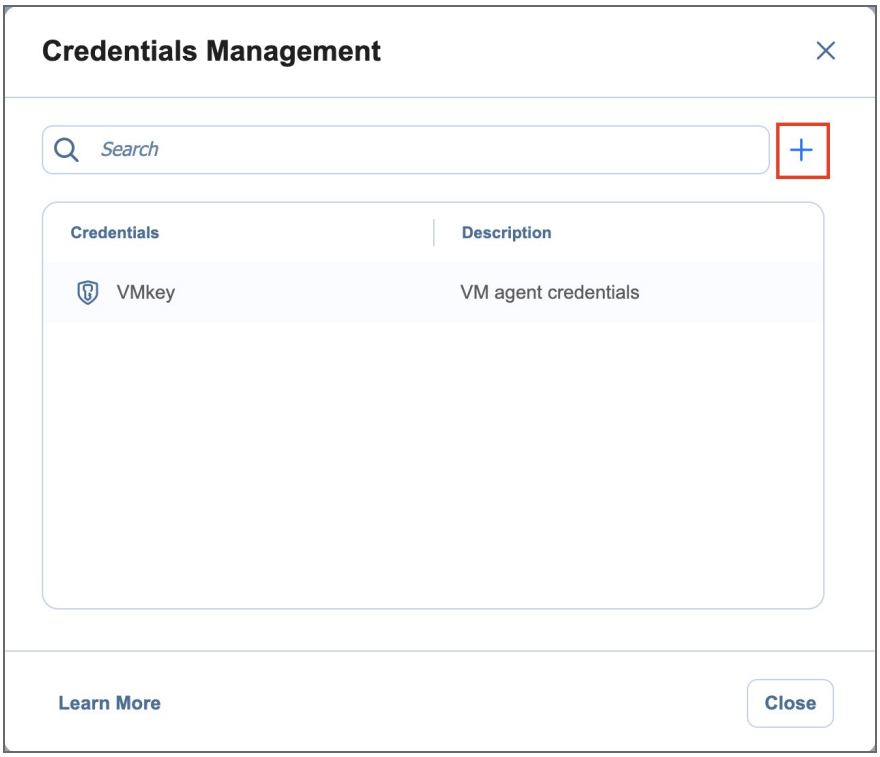
Adding Credentials

To add new credentials, do the following:

1. Click **Settings** in the left pane of the product.
2. Go to the **Inventory** tab.
3. Click **Manage**.
4. In the dialog box that opens, click **Credentials**.



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



6. Then, do the following:
   - **Type**: Select the type of credentials:
     - To set up a basic username and password, fill out the **Username**, **Password**, and (optionally) **Description** fields and click **Save**.
     - To set up a master password, select **Master password** from the drop-down list and fill out the **Name**, **Password**, and (optionally) **Description** fields and click Save.
     - To add a private key to an Amazon EC2 instance or a Linux physical machine, do the following:
       a. **Private key**: Select **Private Key** from the **Type** menu.
       b. **Username**: Enter a username for the private key.
       c. **Password**: Create a password for the private key.
       d. **Repeat password**: Repeat password.

          > **Note**
          > If you generated your key with a passphrase, you have to enter this passphrase into the **Password** and **Repeat password** boxes.

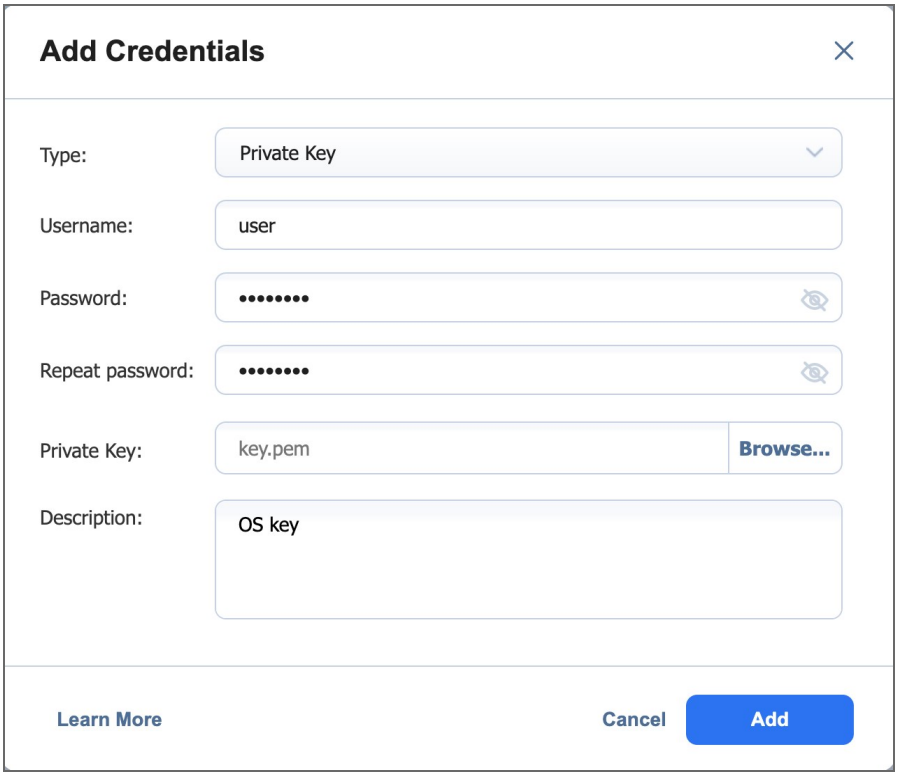
       e. Locate and select the private key.

Notes

- Supported key formats: RSA, *DSA*
- By default, newer versions of *ssh-keygen* generate keys with the unsupported -----*BEGIN OPENSSH PRIVATE KEY*----- format. To generate a key with the -----*BEGIN RSA PRIVATE KEY*----- format, include *-m PEM* in your *ssh-keygen* command.
- Supported file extensions: no extension, `.pem`, `.key`, `.cer`, `.der`, `.txt`

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

g. Click **Save**.

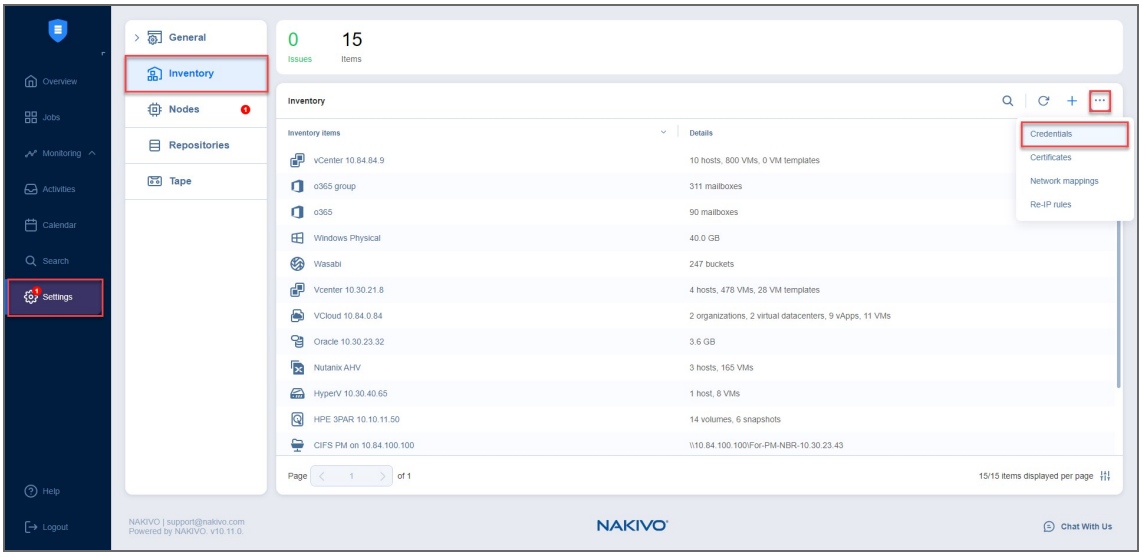


You can now assign the credentials while creating jobs or setting up VM agents.

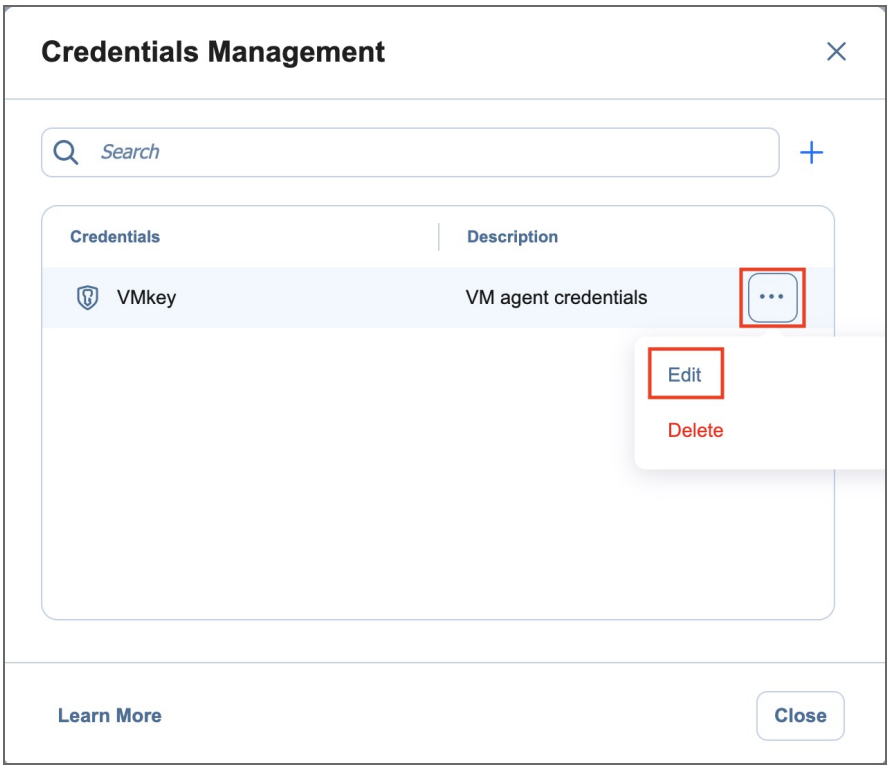
Editing Credentials

To edit credentials, do the following:

1. Click **Settings** in the left pane of the product.
2. Go to the **Inventory** tab.
3. Click **Manage**.
4. In the dialog that opens, click **Credentials**.



5. Hover over the record that you would like to edit and click **Manage > Edit**.

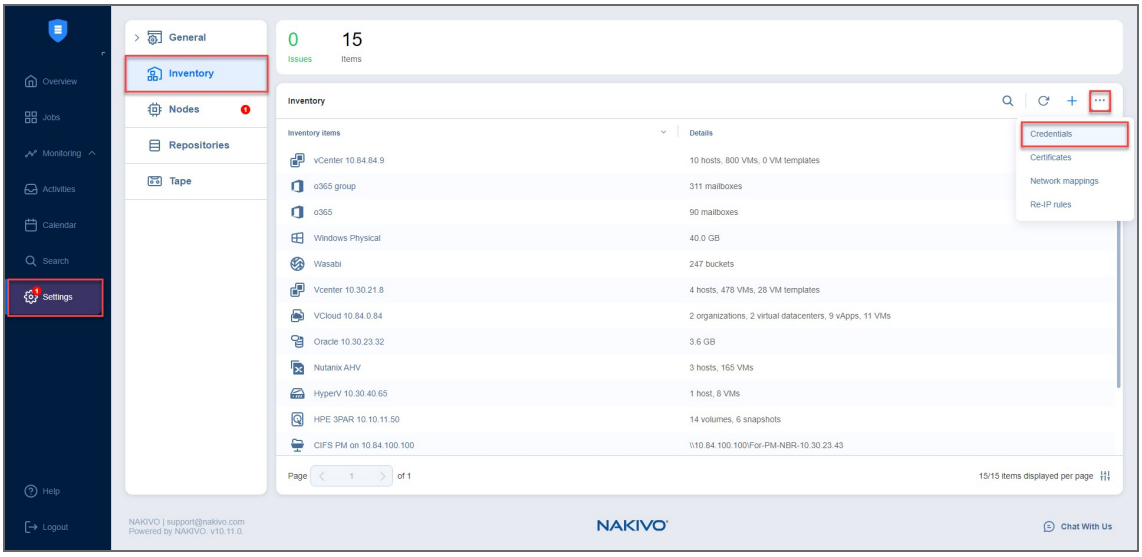


6. Make any required changes, and then click **Save**.

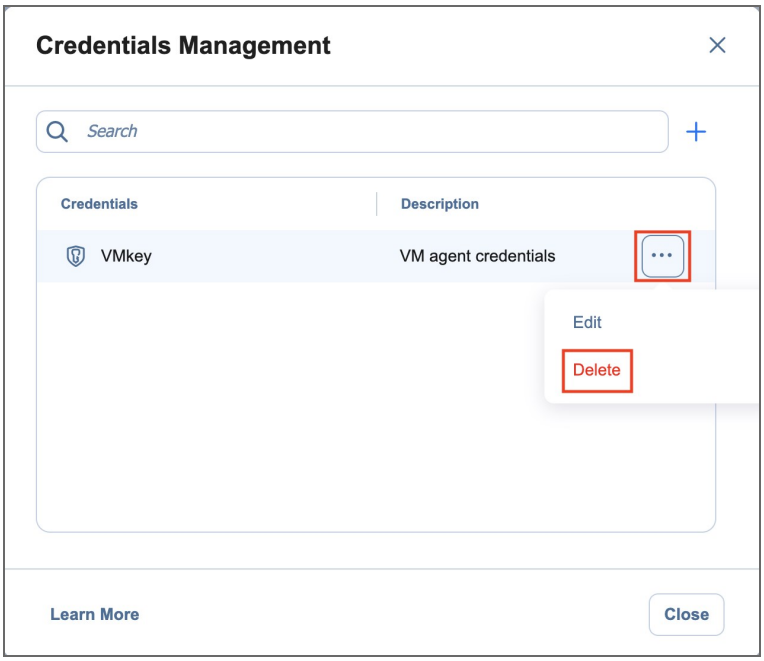
Deleting Credentials

To delete credentials, do the following:

1. Click **Settings** in the left pane of the product.
2. Go to the **Inventory** tab.
3. Click **Manage**.
4. In the dialog that opens, click **Credentials**.



5. Hover over the record that you would like to delete and click **Manage > Delete**.

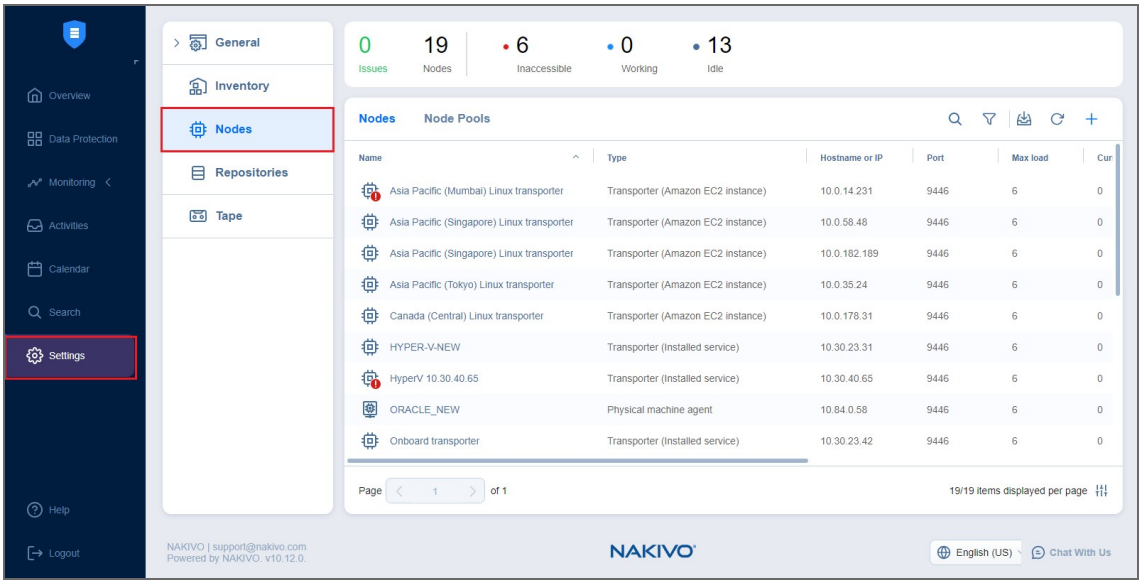


6. Click **Delete** in the confirmation dialog box that opens.

Nodes

Nodes are an essential component of NAKIVO Backup & Replication. They include **Transporters**, VM Agents, and Physical Machine Agents. The **Transporter**, for example, performs backup, replication, and recovery, as well as data compression, deduplication, and encryption. The **Nodes** tab contains a **Summary** bar, which offers an overview of all nodes. The data displayed is as follows:

- **Issues**: Total number of issues/alarms related to nodes
- **Nodes**: Total number of nodes
- **Inaccessible**: Total number of inaccessible nodes
- **Working**: Total number of working nodes
- **Idle**: Total number of idle nodes



To learn how to add nodes and manage them, refer to the topics below:

- “Configuring Nodes” below
- “Managing Nodes” on page 448

Configuring Nodes

Refer to the following topics:

- “Adding Existing Nodes” on the next page
- “Deploying Transporter as Nutanix AHV Appliance” on page 440
- “Deploying Transporter as VMware Appliance” on page 441
- “Deploying Transporters in Amazon EC2” on page 444

Adding Existing Nodes

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

Important

Before adding the existing **Transporter** to your NAKIVO Backup & Replication, make sure that this **Transporter** is not used by any other **Director** as it may lead to unforeseen errors.

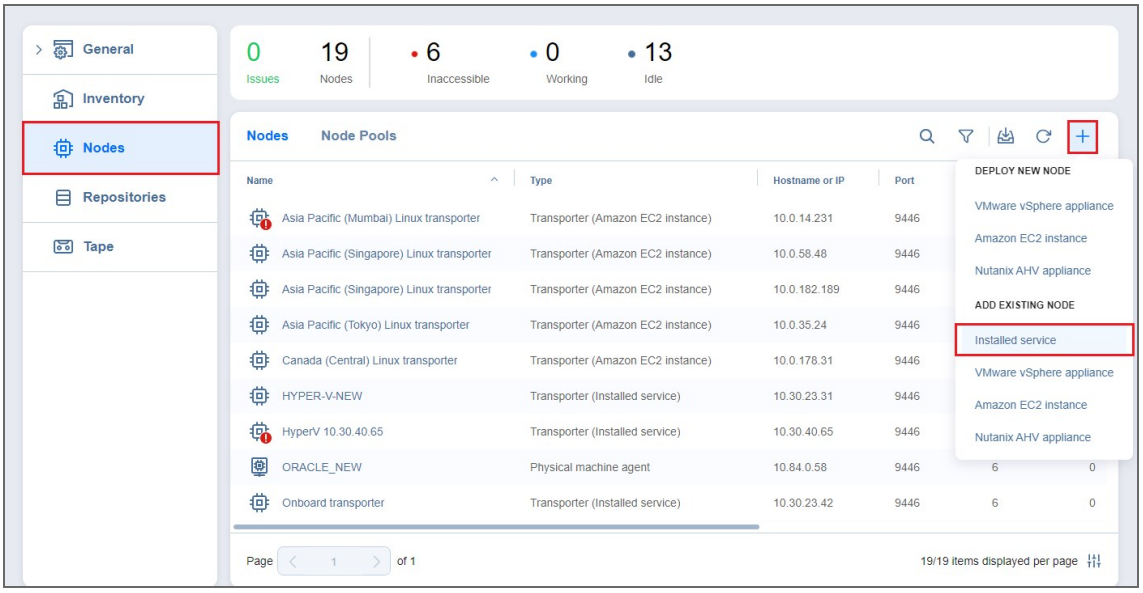
Refer to the following topics:

- Installed Service
- VMware Appliance
- Amazon EC2 Instance
- Nutanix AHV Appliance

Installed Service

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

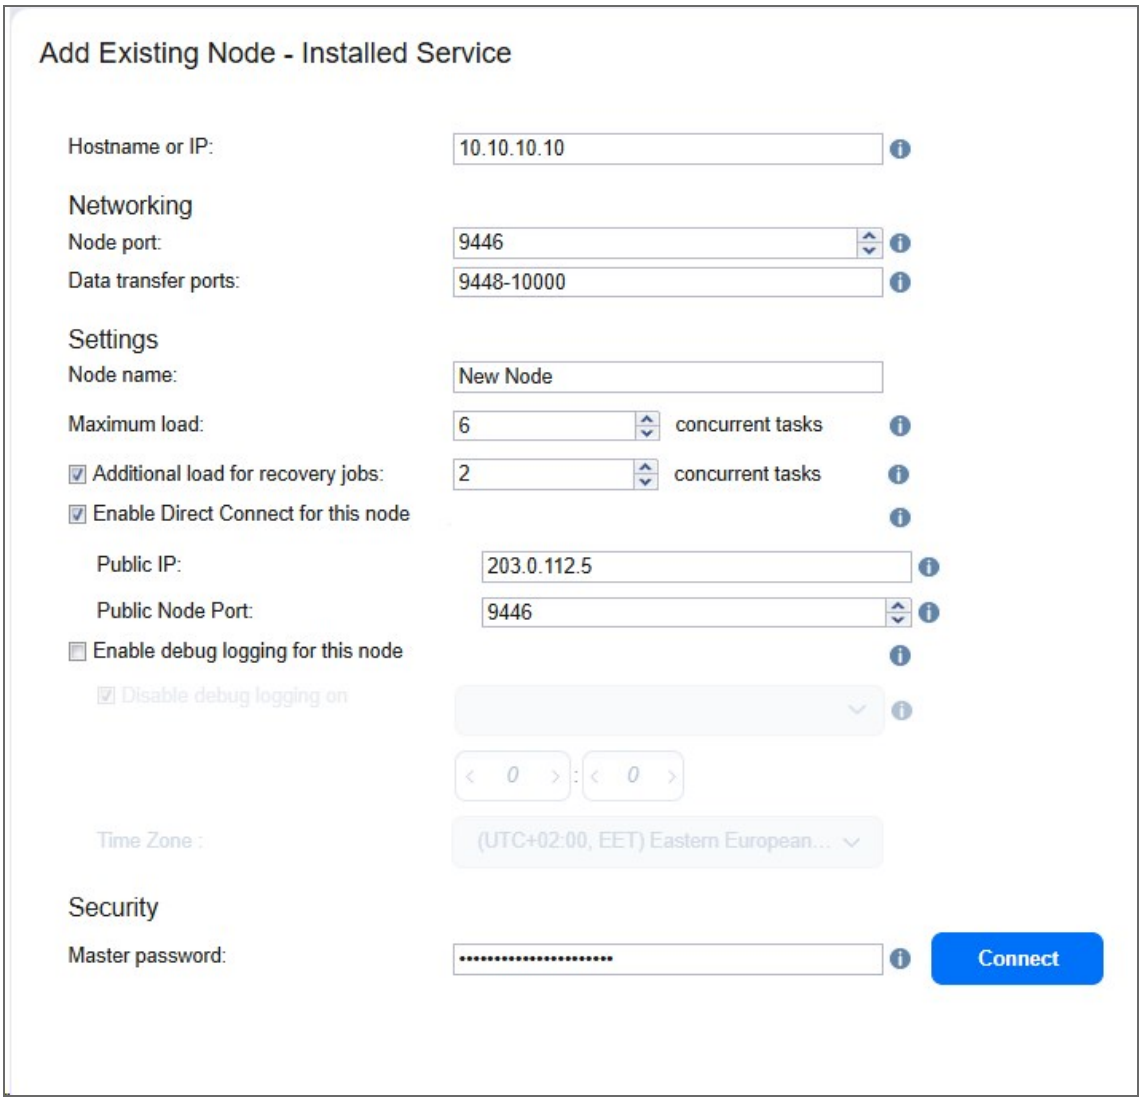
1. Click **Settings** in the left pane of the product, go to the **Nodes** tab, and click **+**.
2. In the **Add Existing Node** section, click **Installed service**.



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

Note

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



3. Click **More options...** to reveal and edit the following fields:
    - In the *Networking* section:
        - **Node port**: Specify the port number that will be used to connect to the node.
        - **Data transfer ports**: Specify a range of port numbers (from 1 to 65535) that will be used to transfer data. The range you specify should contain at least 100 ports. Make sure that the ports you specify are open in your firewall.
    - In the *Settings* section:
        - **Node name**: Specify a display name for the node.
        - **Maximum load**: Specify the maximum number of tasks that the node should process simultaneously. A task, for example, is a backup or replication of a single VM disk, or one granular recovery session.

- **Additional load for recovery jobs**: If selected, the specified amount of tasks will be added to set maximum node load to be used for recovery jobs exclusively. This allows running the specified amount of concurrent recovery jobs along with other types of jobs without the need to wait for their completion.
- **Enable Direct Connect for this node**: When this option is enabled, you can access remote resources via a single port connection without establishing a VPN connection. For more information refer to Enabling Direct Connect.
    - **Public IP**: The IP address of external **MSP Transporter** that is reachable from the internet and used for communication with the **Direct Connect Transporter**.
    - **Public Node Port**: The port of external **MSP Transporter** that is reachable from the internet and used for communication with the **Direct Connect Transporter**. Ensure that the entered combination of the **Public Node Port** and **Public IP** address is not used by other transporters.

  The following conditions must be met at the remote infrastructure to enable this feature:
    - A NAKIVO **Transporter** or Agent must be installed.
    - A master password must be set for security reasons. A pre-shared key is generated based on the entered master password.
    - The node port on the local machine must be exposed to external availability via the Internet.
    - The **Public IP** and **Public Node Port** fields are available only with the license that supports MSP Direct Connect.
- **Enable debug logging for this node**: If needed, enable debug level logging for the current node. It is not recommended to use this option on a permanent basis.

- In the *Security* section:
    - **Master Password**: Optionally, you can set a password to secure the connection. The set password must match the one configured on the **Transporter** or Agent. Note that setting a master password is required when the **Enable Direct Connect for this node** option is enabled. Proceed as follows:
        a. After entering the password, click **Connect**.
        b. The **Certificate Details** dialog box appears. Verify the certificate details, and click **Accept**.

Notes

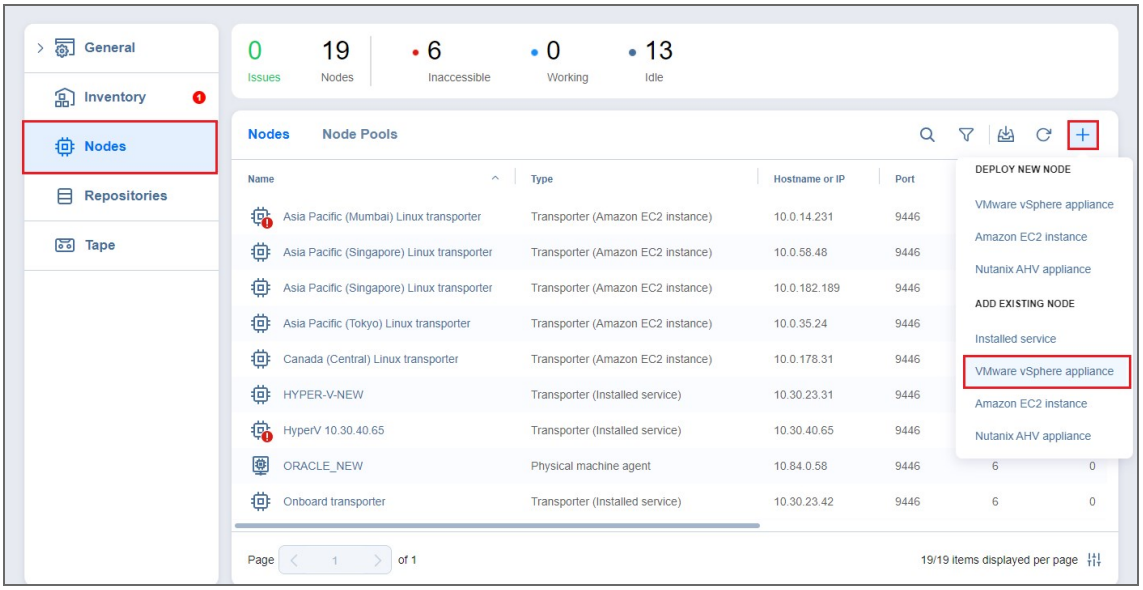
- The master password must adhere to the following requirements:
  - Minimal length - 5 characters.
  - Maximum length - 50 characters.
- The master password can be set and re-set manually by running the command on the machine housing the **Transporter** or Agent. Follow these steps:
  - Enter the following command `bhsvc -b password`, replacing "password" with your master password.
  - Restart the **Transporter** or Agent.

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

VMware Appliance

Follow the steps below to add a **Transporter** that is deployed as a VMware appliance:

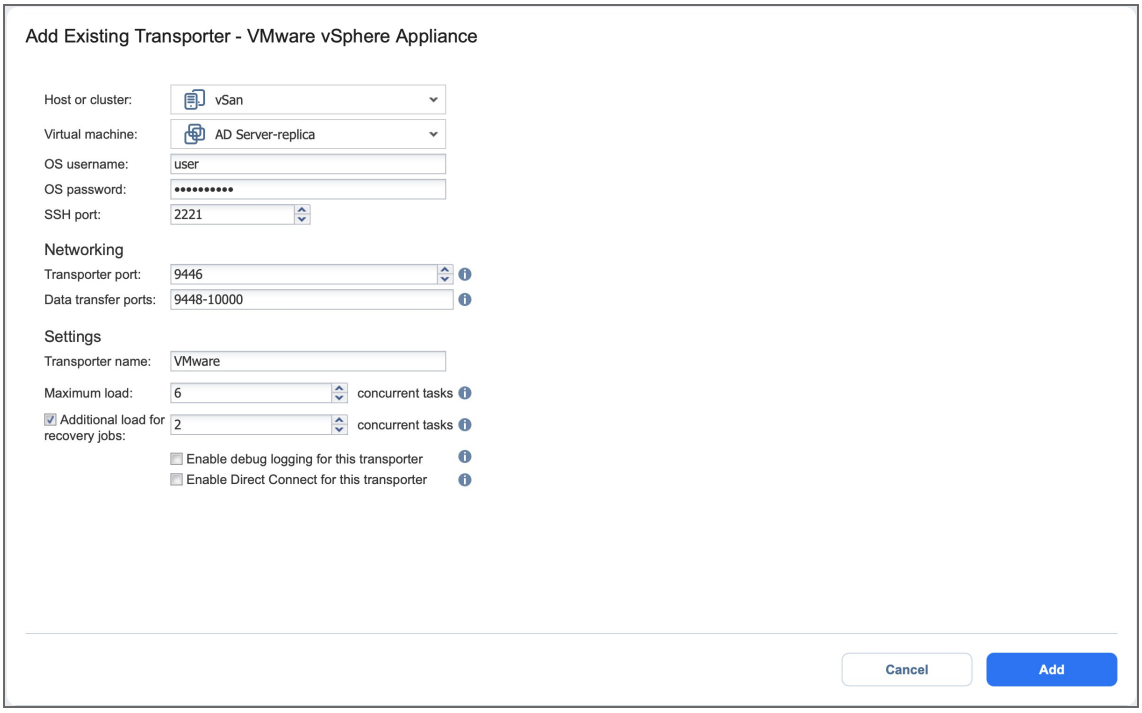
1. Click **Settings** in the left pane of the product, go to the **Nodes** tab, and click **+**.
2. In the **Add Existing Node** section, click **VMware vSphere appliance**.



3. The **Add Existing Transporter - VMware vSphere Appliance** dialog opens. Fill out the fields as described below:
   - In the **Host or cluster** box, specify the location of the host or cluster where the corresponding virtual machine is deployed.
   - In the **Virtual machine** box, specify the virtual machine on which the **Transporter** is installed.

- In the **OS Username** and **OS Password** fields, specify credentials for accessing the virtual machine.
- In the **SSH port** box, enter the SSH port if needed.
- Click **More options** to reveal and edit the following fields:
    - In the *Networking* section:
        - **Transporter port**: Specify the port number that will be used to connect to the Transporter.
        - **Data transfer ports**: Specify a range of port numbers (from 1 to 65535) that will be used to transfer data. The range you specify should contain at least 100 ports. Make sure that the ports you specify are open in your firewall.
    - In the *Settings* section:
        - **Transporter name**: Specify a display name for the **Transporter**.
        - **Maximum load**: Specify the maximum number of tasks that the **Transporter** should process simultaneously. A task, for example, is a backup or replication of a single VM disk, or one granular recovery session.
        - **Additional load for recovery jobs**: Selecting this option reserves the **Transporter**'s resources exclusively for recovery jobs. This allows you to run recovery jobs concurrently with other types of jobs without the need to wait for their completion. The **Transporter** resources will be reserved according to the specified number.
        - **Enable debug logging for this transporter**: If needed, enable debug level logging for the current transporter. It is not recommended to use this option on a permanent basis.
        - **Enable Direct Connect for this transporter**: When this option is enabled, you can access remote resources via a single port connection without establishing a VPN connection. The following conditions must be met at the remote infrastructure to enable this feature:
            - A NAKIVO Transporter must be installed.
            - A master password must be set for security reasons. A pre-shared key is generated based on the entered master password.
            - The **Transporter** port on the local machine must be exposed to external availability via the Internet.

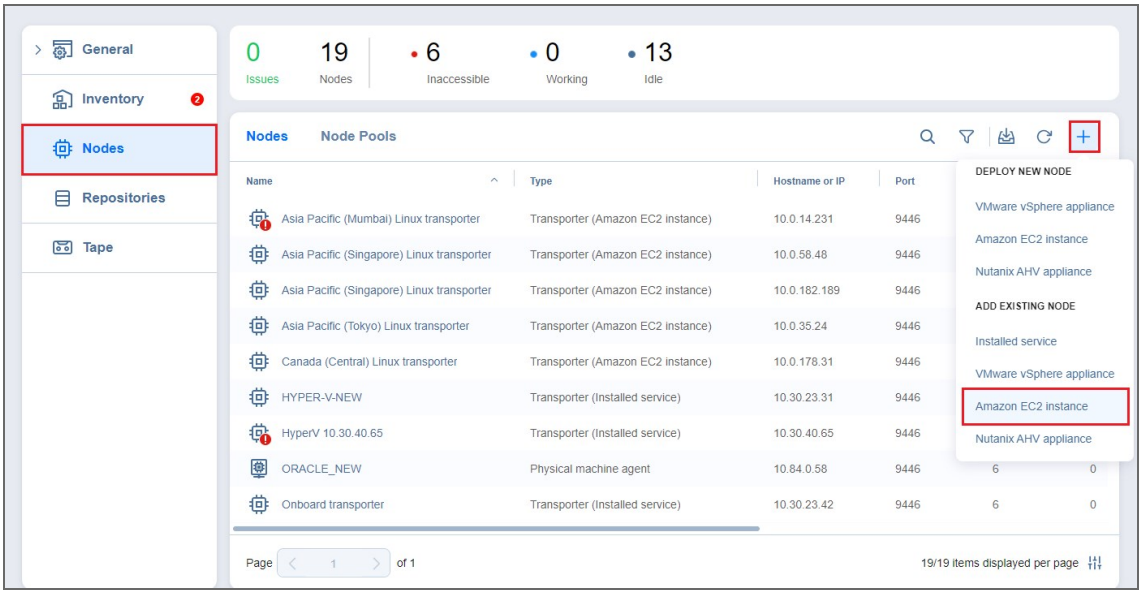
4. Click **Add**. The **Transporter** is added to the product and can be used for backup, replication, and recovery jobs.



Amazon EC2 Instance

If you have already deployed a **Transporter** in Amazon EC2 and now wish to re-import the **Transporter** in a new instance of NAKIVO Backup & Replication, do the following:

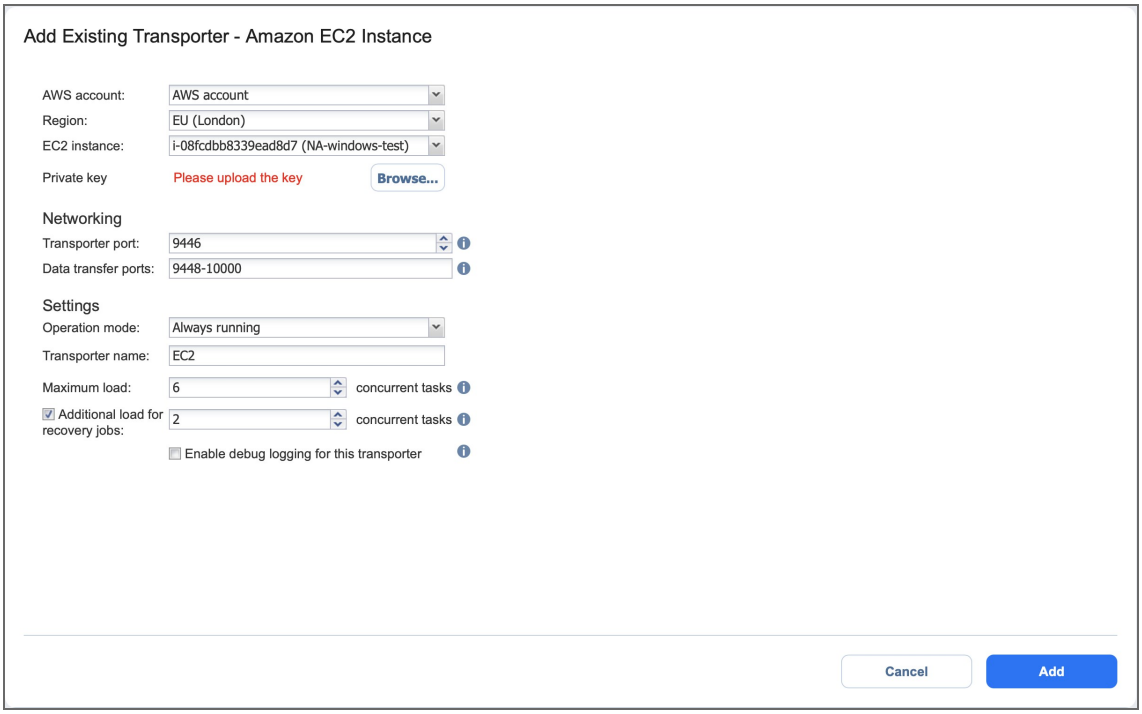
1. Click **Settings** in the left pane of the product, go to the **Nodes** tab, and click **+**.
2. In the **Add Existing Node** section, click **Amazon EC2 instance**.



3. The **Add Existing Transporter - Amazon EC2 Instance** dialog opens. Fill out the fields as described below:

- **AWS account**: Choose an appropriate Amazon AWS Account from the list of Amazon AWS Accounts added to the Inventory.
- **Region**: Choose a region in which an AWS EC2 instance with the **Transporter** is deployed.
- **EC2 Instance**: Select the Amazon EC2 Instance with the **Transporter** that you wish to add to the product.
- **Private key**: Click the **Browse** button to locate and upload the Private key for the **Transporter** Instance that was created when you deployed the **Transporter** in the cloud.
- Click **More options** to reveal and edit the following fields:
    - In the *Networking* section:
        - **Transporter port**: Specify the port number that will be used to connect to the **Transporter**.
        - **Data transfer ports**: Specify a range of port numbers (from 1 to 65535) that will be used to transfer data. The range you specify should contain at least 100 ports. Make sure that the ports you specify are open in your firewall.
    - In the *Settings* section:
        - **Operation mode**: Choose one of the following **Transporter** operation modes:
            - **Always running**
            - **Running while required**
        - **Transporter name**: Specify a display name for the **Transporter**.
        - **Maximum load**: Specify the maximum number of tasks that the **Transporter** should process simultaneously. A task, for example, is a backup or replication of a single VM disk, or one granular recovery session.
        - **Additional load for recovery jobs**: If selected, the specified amount of tasks will be added to set maximum transporter load to be used for recovery jobs exclusively.
        - **Enable debug logging for this Transporter**: If needed, enable debug level logging for the current transporter. It is not recommended that you use this option on a permanent basis.

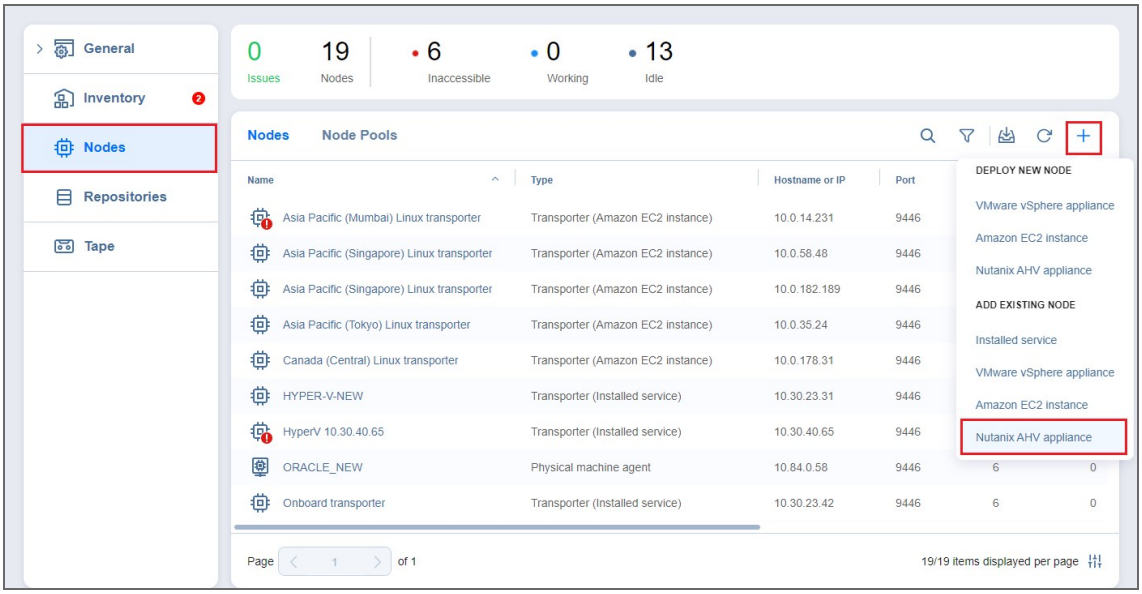
4. Click **Add**. The **Transporter** is added to the product and can be used for backup, replication, and recovery jobs.



Nutanix AHV Appliance

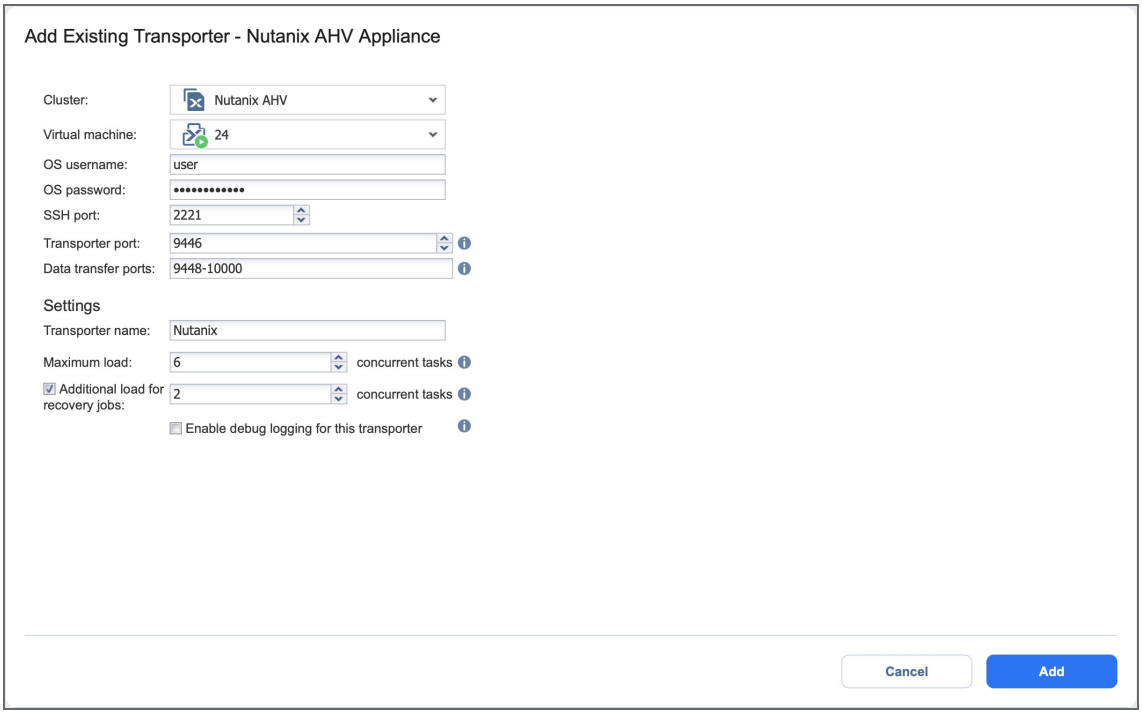
Follow the steps below to add a **Transporter** that is deployed as a Nutanix AHV appliance:

1. Click **Settings** in the left pane of the product, go to the **Nodes** tab, and click **+**.
2. In the **Add Existing Node** section, select **Nutanix AHV appliance**.



3. In the **Add Existing Transporter - Nutanix AHV Appliance** dialog, enter the following options:
   - In the **Cluster** box, select the cluster where the corresponding virtual machine is deployed.
   - In the **Virtual machine** box, specify the virtual machine on which the **Transporter** is installed.

- In the **OS Username** and **OS Password** fields, specify credentials for accessing the virtual machine.
- In the **SSH port** box, enter the SSH port if needed.
- Click **More options** to reveal and edit the following fields:
  - In the *Networking* section:
    - **Transporter port**: Specify the port number that will be used to connect to the **Transporter**.
    - **Data transfer ports**: Specify a range of port numbers (from 1 to 65535) that will be used to transfer data. The range you specify should contain at least 100 ports. Make sure that the ports you specify are open in your firewall.
  - In the *Settings* section:
    - **Transporter name**: Specify a display name for the **Transporter**.
    - **Maximum load**: Specify the maximum number of tasks that the **Transporter** should process simultaneously. A task, for example, is a backup or replication of a single VM disk, or one granular recovery session.
    - **Additional load for recovery jobs**: If selected, the specified amount of tasks will be added to set maximum transporter load to be used for recovery jobs exclusively.
    - **Enable debug logging for this Transporter**: If needed, enable debug level logging for the current transporter. It is not recommended that you use this option on a permanent basis.



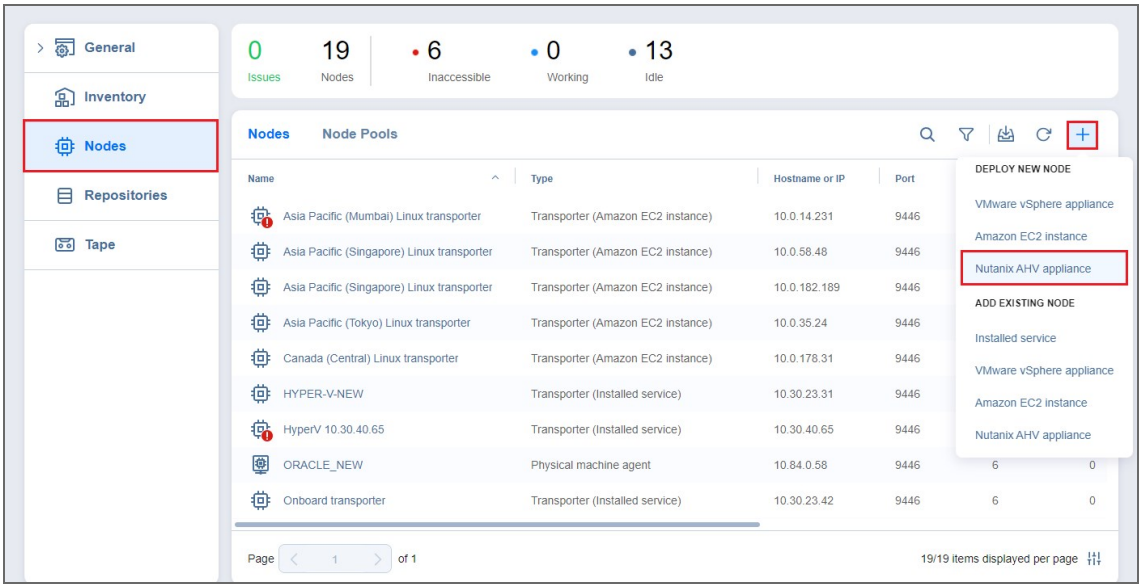
4. Click **Add**. The **Transporter** is added to the product and can be used for backup, replication, and recovery jobs.

Deploying Transporter as Nutanix AHV Appliance

To enable NAKIVO Backup & Replication to create and run jobs within a Nutanix AHV cluster, a dedicated **Transporter** must be deployed as a Nutanix appliance in that cluster.

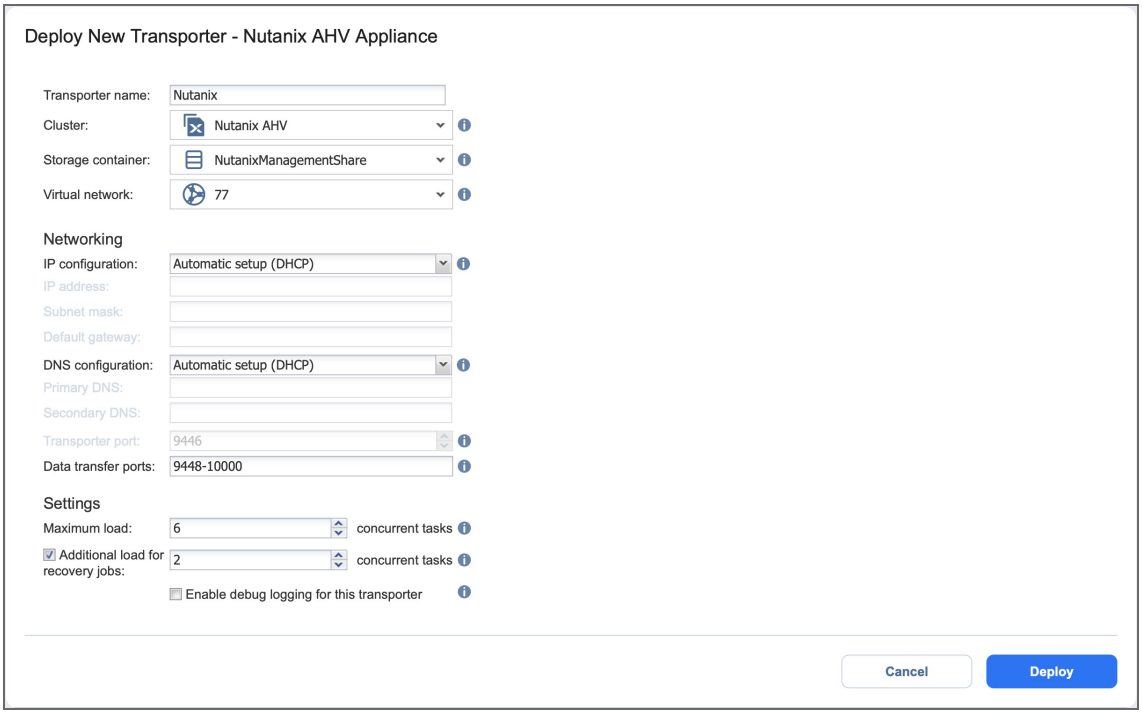
Please follow the steps below to add a transporter as a Nutanix appliance:

1. Click **Settings** in the left pane of the product, go to the **Nodes** tab, and click **+** .
2. In the **Deploy New Node** section, click **Nutanix AHV appliance**.



3. In the **Deploy New Transporter - Nutanix AHV Appliance** dialog, specify the following options:
   - **Transporter name**: Enter a name for the new **Transporter**.
   - **Cluster**: Select a cluster where the transporter VM will run.
   - **Storage container**: Select a storage container where the transporter VM will be located.
   - **Virtual network**: Select a virtual network where the transporter VM will be connected.
4. Click **Deploy** to proceed with the automatically selected networking options and default **Transporter** load configuration.
5. Alternatively, click **More options** if you wish to manually set the following options:
   - **IP configuration**: Can be either **Automatic setup (DHCP)** or **Manual setup**. With manual setup selected, specify an **IP address**, **Subnet mask** and **Default gateway**.
   - **DNS configuration**: Can be either **Automatic setup (DHCP)** or **Manual setup**. With manual setup selected, specify **Primary** and **Secondary DNS**.
   - **Transporter port**: Enter a communication port for your **Transporter**.
   - **Data transfer ports**: Enter a port range that will be used by your **Transporter** for actual data transfer.

- **Maximum load**: Specify the maximum number of tasks that the **Transporter** should process simultaneously. A task, for example, is a backup or replication of a single VM disk, or one granular recovery session.
- **Additional load for recovery jobs**: If selected, the specified quantity of tasks will be added to set the maximum **Transporter** load to be used for recovery jobs exclusively. This allows for running the specified amount of concurrent recovery jobs along with other types of jobs without the need to wait for their completion.
- **Enable debug logging for this transporter**: If needed, enable debug level logging for the current transporter. Using this option on a permanent basis is not recommended.



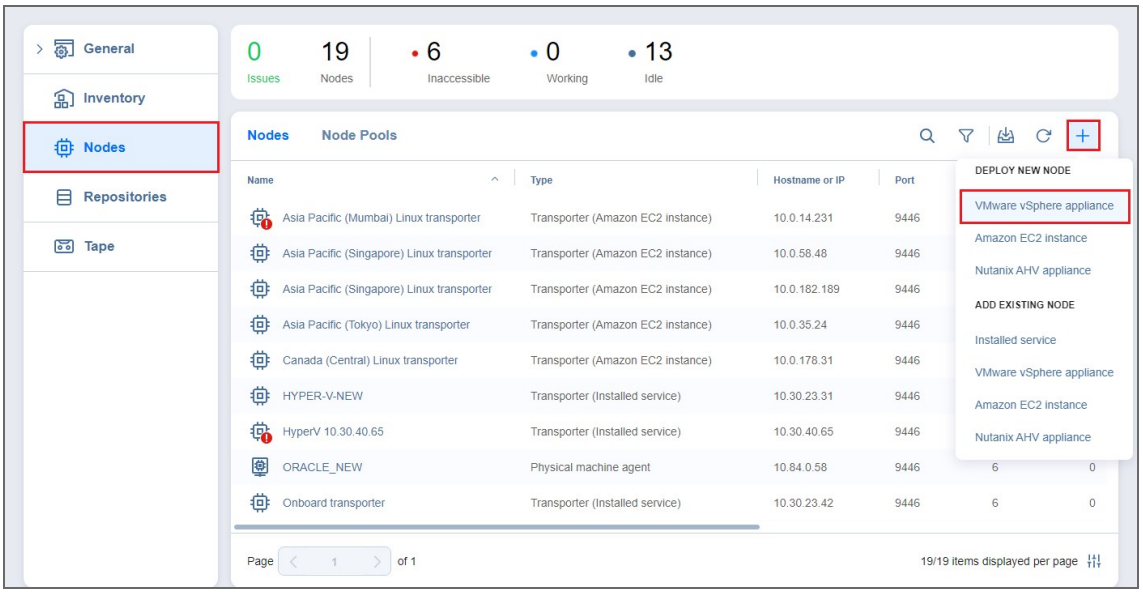
6. Click **Deploy** to begin the deployment process. Successfully deployed transporters are displayed in the **Transporters** tab.

Deploying Transporter as VMware Appliance

> **Note**
> If your instance of NAKIVO Backup & Replication is installed on ARM-based NAS, an external **Transporter** needs to be deployed to work with VMware vCenters and ESXi hosts. This is because certain features are not supported by ARM-based NASes.

Please follow the steps below to deploy a Transporter that supports VMware vCenter:

1. Click **Settings** in the left pane of the product, go to the **Nodes** tab, and click **+** .
2. In the **Deploy New Node** section, click **VMware vSphere appliance**.



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

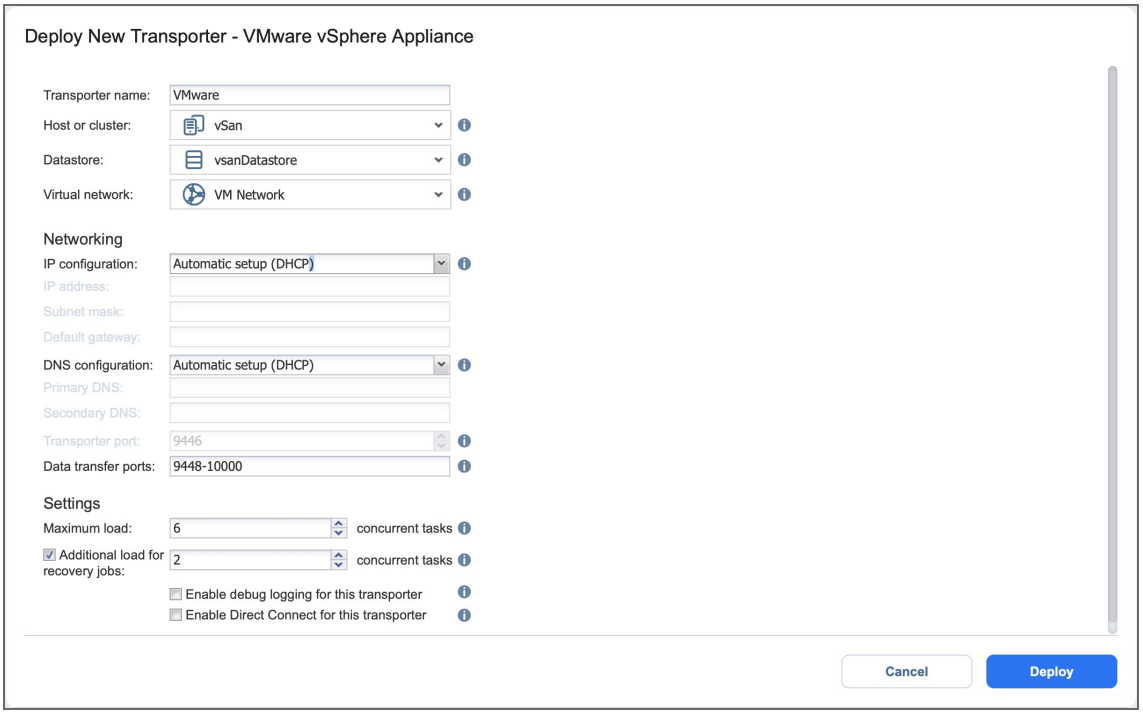
    > **Note**
    > An internet connection is required to deploy a new Transporter as a VMware appliance on the target host or cluster.

    - If necessary, access the advanced options for your **Transporter** by clicking **More options** and then entering data for the following parameters:
        - In the *Networking* section:
            - **IP configuration**: It can be either **Automatic setup (DHCP)**, or **Manual setup**.
            - **IP address**: If you have chosen **Manual setup** for the IP configuration, enter a **Transporter** IP address.
            - **Subnet mask**: If you have chosen **Manual setup** for the **IP configuration**, enter a subnet mask.
            - **Default gateway**: If you have chosen **Manual setup** for the **IP configuration**, enter a default gateway.
            - **DNS configuration**: It can be either **Automatic setup (DHCP)**, or **Manual setup**.
            - **Primary DNS**: If you have chosen **Manual setup** for the **DNS configuration**, enter a primary DNS server IP address.

- **Secondary DNS**: If you have chosen **Manual setup** for the **DNS configuration**, enter a secondary DNS server IP address.
- **Transporter port**: Enter a communication port for your transporter.
- **Data transfer ports**: Enter a port range that will be used by your transporter for actual data transfer.

- In the *Settings* section:
  - **Maximum load**: A number of tasks concurrently processed by the **Transporter**.
  - **Additional load for recovery jobs**: If selected, the specified amount of tasks will be added to set maximum transporter load to be used for recovery jobs exclusively. This allows for running the specified amount of concurrent recovery jobs along with other types of jobs without the need to wait for their completion.
  - **Enable debug logging for this transporter**: When selected, it enables debug level logging for the Transporter. It is not recommended to have this option selected on a permanent basis.
  - **Enable Direct Connect for this transporter**: When this option is enabled, you can access remote resources via a single port connection without establishing a VPN connection. The following conditions must be met at the remote infrastructure to enable this feature:
    - A NAKIVO **Transporter** must be installed.
    - A master password must be set for security reasons. A pre-shared key is generated based on the entered master password.
    - The **Transporter** port on the local machine must be exposed to external availability via the Internet.

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



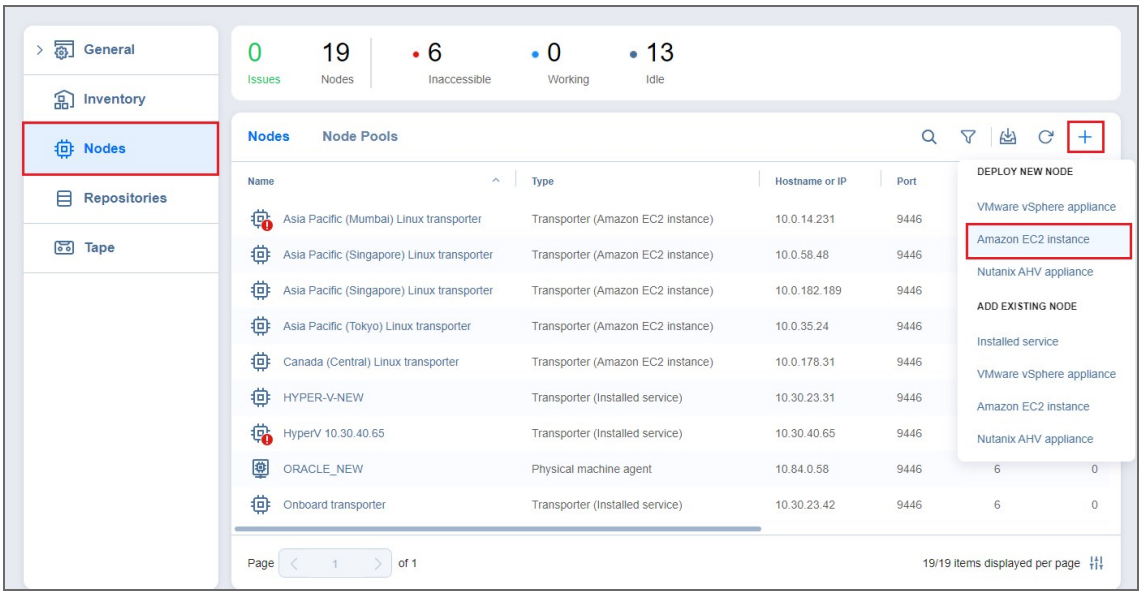
Deploying Transporters in Amazon EC2

You need to deploy a **Transporter** in Amazon EC2 to enable the following features:

- Backing up VMware VMs and/or Amazon EC2 Instances to a backup repository located in Amazon EC2.
- Backing up Amazon EC2 Instances in a particular Amazon EC2 Region.

NAKIVO Backup & Replication automates deploying a **Transporter** in Amazon EC2. To deploy a **Transporter** in Amazon EC2 within the product interface, follow the steps below:

1. Click **Settings** in the left pane of the product, go to the **Nodes** tab, and click **+** .
2. In the **Deploy New Node** section, click **Amazon EC2 instance**.



3. The **Deploy New Transporter - Amazon EC2 Instance** dialog opens. Fill out the fields as described below:
    - **Transporter name**: Enter a name for the **Transporter**.
    - **Region**: Select an Amazon EC2 region where you wish to deploy the **Transporter**. This will enable you to create a backup repository in the region as well as back up Amazon EC2 Instances available in the region.
    - **Instance type**: Choose a type of Amazon EC2 Instance (for example, "t2.medium") that will be used to deploy the **Transporter**. Note that more powerful instances may be able to process data faster, but will cost more to run on Amazon EC2.

      Note
      ARM-based instances cannot be selected if you have chosen **Windows** for the **Platform** option.

    - Click **More options** to reveal and edit the following options:
        - In the *Networking* section:
            - **Automatically configure VPC for this transporter**: If selected, a new VPC with a single public subnet will be created and used to deploy this transporter. If you want to deploy the **Transporter** into a different VPC and subnet, deselect this option.
                - **Network**: Select a network to which the Amazon EC2 instance with the **Transporter** will be connected.
                - **Subnet**: Select a subnet for the Amazon EC2 Instance with the **Transporter**.
            - **Allowed traffic from**: Enter the IP addresses of the machines that can connect to the Amazon EC2 instance with the **Transporter**. Access from other IP addresses will be restricted.

              Important
              By default, the Amazon EC2 security group is not restricted; that is, the **Transporter** can be accessed by and receive tasks from any machine. For security purposes, restrict traffic to trusted IP addresses.

            - **Transporter Port**: Specify the port number that will be used to connect to the **Transporter**.

- **Data transfer ports**: Specify a range of port numbers (from 1 to 65535) that will be used to transfer data. The range you specify should contain at least 100 ports. Make sure that the ports you specify are open in your firewall.

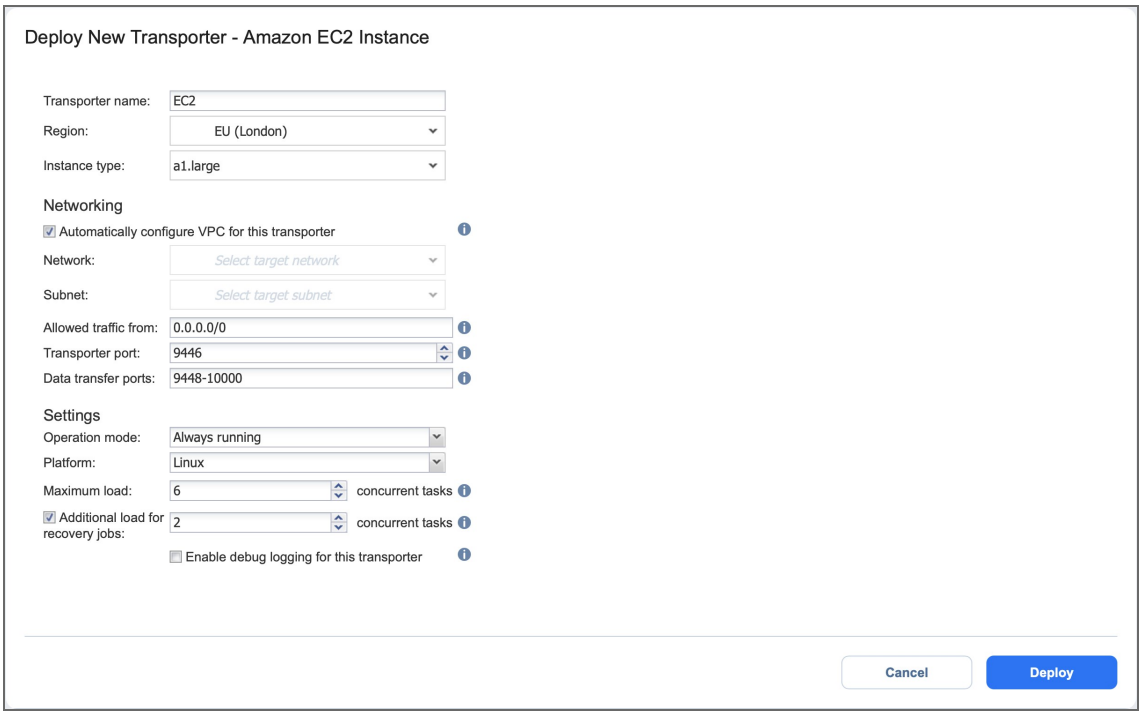
- In the *Settings* section:
  - **Operation mode**: If you select the **Running while required** option, the Amazon EC2 Instance with the **Transporter** will be powered on only when the **Transporter** is required to run a backup, replication, and recovery tasks.
  - **Platform**: Choose an OS for the instance where the **Transporter** will be deployed.

    Note
    Windows OS is not supported for ARM-based instances.

  - **Maximum load**: Specify the maximum number of tasks that the **Transporter** should process simultaneously. An example of a task is processing a single VM disk or a single file recovery session.
  - **Additional load for recovery jobs**: If selected, the specified amount of tasks will be added to set the maximum **Transporter** load to be used for recovery jobs exclusively. This allows for running the specified quantity of concurrent recovery jobs along with other types of jobs without the need to wait for their completion.
  - **Enable debug logging for this Transporter**: Enables debug level logging for the current **Transporter**. Since this feature slows down **Transporter** performance, it is recommended that you enable debug logging only for the investigation of support issues.

Note
Refer to "Amazon EC2 Concepts" on page 7 for the definitions of Amazon EC2-related terms.

4. Click **Deploy**.



Note

- After deploying a **Transporter** in Amazon EC2, you need to download the Transporter Key. A Transporter Key is used by NAKIVO Backup & Replication to access and manage the **Transporter** in Amazon EC2. If you lose the current instance of NAKIVO Backup & Replication and install a new copy of the product, you will need to provide the Transporter Key to access the **Transporter**.
- You may be additionally charged for using a 3rd-party resource. Please refer to the 3rd-party resource provider documentation for details.

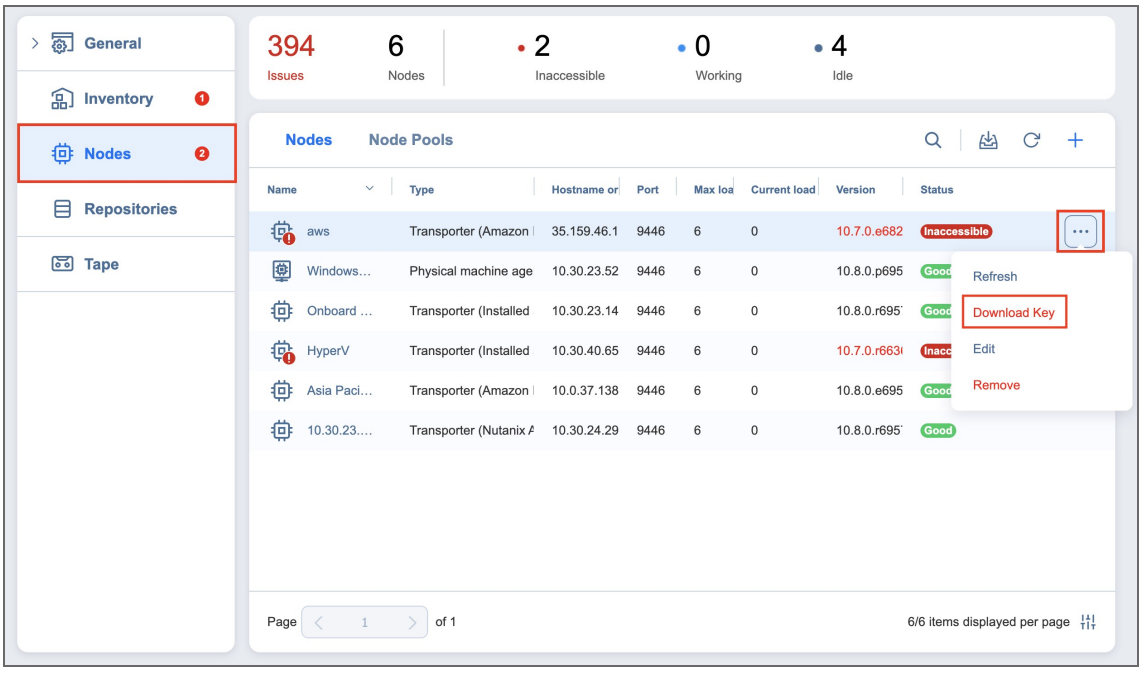
Managing Nodes

Refer to the following topics:

- “Editing Nodes” below
- “Downloading Transporter Credentials” below
- “Managing Node Pools” on page 456
- “Refreshing Node Details” on page 457
- “Removing (Deleting) Nodes” on page 458

Downloading Transporter Credentials

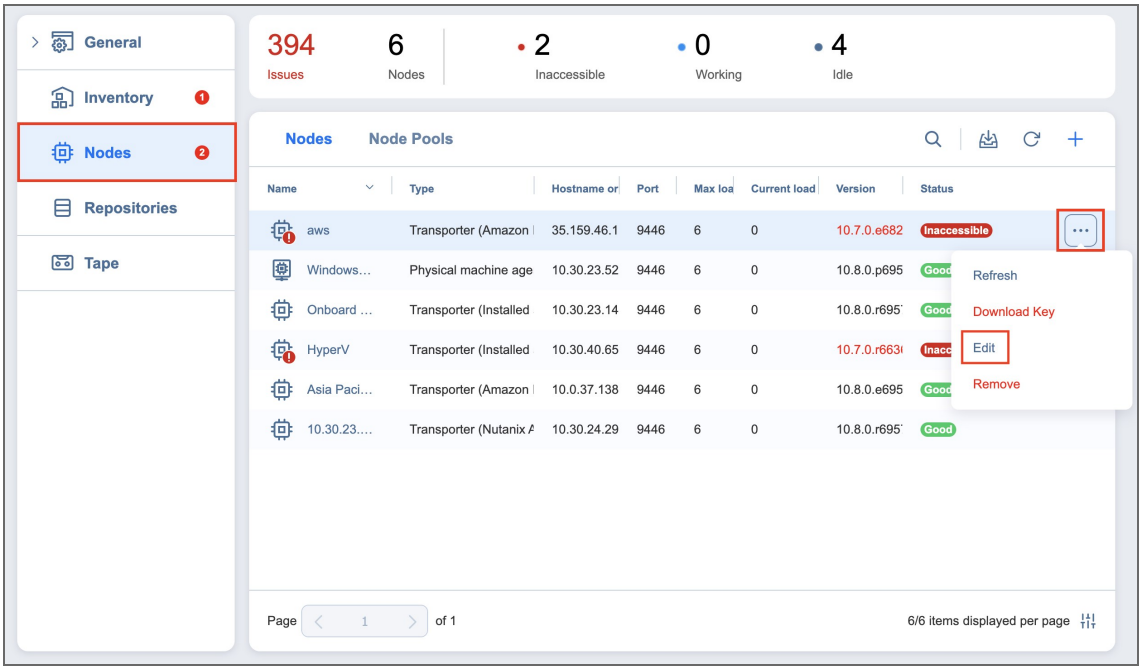
If you would like to import an Amazon EC2, Nutanix AHV, or VMware **Transporter** into another installation of NAKIVO Backup & Replication, you need to download the **Transporter**'s credentials. To obtain the credentials, hover over the desired **Transporter** and click the ellipsis **Manage** button on the right side. In the dialog box, click **Download Key**. This begins the download of a ZIP file containing the **Transporter**'s credentials.



Editing Nodes

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

1. Click **Settings** in the left pane of the product.
2. Go to the **Nodes** tab and hover over the node you would like to edit.
3. On the right side, click the ellipsis **Manage** button and then click **Edit**.



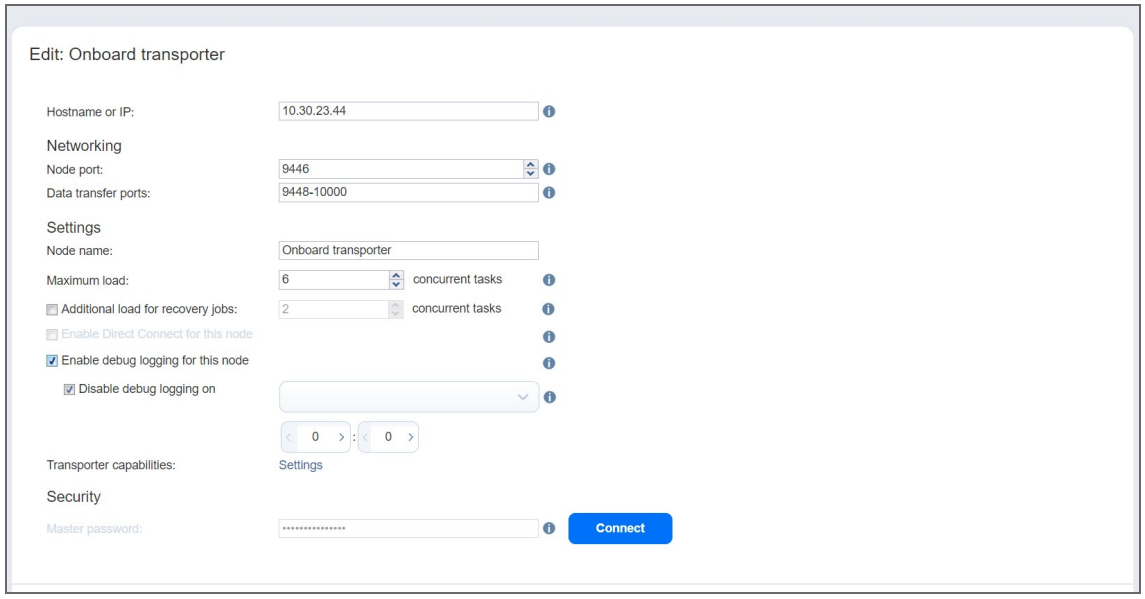
4. A dialog opens for editing the node's settings. Edit the settings as required:
    - **Hostname or IP**: Here you can edit the IP address or hostname of the machine on which the node is installed. Not applicable to Nutanix AHV Appliances, VMware vSphere appliances, or Amazon EC2 instances.
        - In the *Networking* section:
            - If editing a Nutanix AHV Appliance, VMware vSphere appliance, or Amazon EC2 instance, you can edit the following options:
                - **OS username**: Enter the username used to access the virtual machine.
                - **OS password**: Enter the password for the username entered previously (not applicable to EC2 instances).
                - **SSH port**: Enter the SSH port if needed.
            - If editing other node types:
                - **Node port**: Enter a communication port for your node.
                - **Data transfer ports**: Enter a port range that will be used by your node for actual data transfer.
        - In the *Settings* section:
            - **Node name**: Edit the name of your node.
            - **Maximum load**: Edit the number of tasks concurrently processed by the node.
            - **Additional load for recovery jobs**: If selected, the specified amount of tasks will be added to the set maximum node load to be used for recovery jobs exclusively. This allows for running the specified amount of concurrent recovery jobs along with other types of jobs without the need to wait for their completion.

- **Enable Direct Connect for this node**: When this option is enabled, you can access remote resources via a single port connection without establishing a VPN connection. The following conditions must be met at the remote infrastructure to enable this feature:
  - A NAKIVO **Transporter** or Agent must be installed.
  - A master password must be set for security reasons. A pre-shared key is generated based on the entered master password.
  - The node port on the local machine must be exposed to external availability via the Internet.
- **Enable debug logging for this node**: Enable/disable debug level logging for the node. Having this option enabled on a permanent basis is not recommended.
- **Transporter capabilities**: Click the **settings** button next to this option to modify the additional capabilities of the selected **Transporter**. Click **Apply** to save your changes.

Notes

- **Backup repository management** cannot be disabled if the **Transporter** is assigned to a repository.
- **VMware vSphere support** cannot be disabled if the **Transporter** is selected in a VMware backup, replication, or recovery job.
- **Transporter** capabilities is not supported for physical agent.

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



Enabling Direct Connect

To start using this feature, you can choose different approaches. For feature requirements, refer to Direct Connect Requirements.

Direct Connect

To start using this feature, take the following steps:

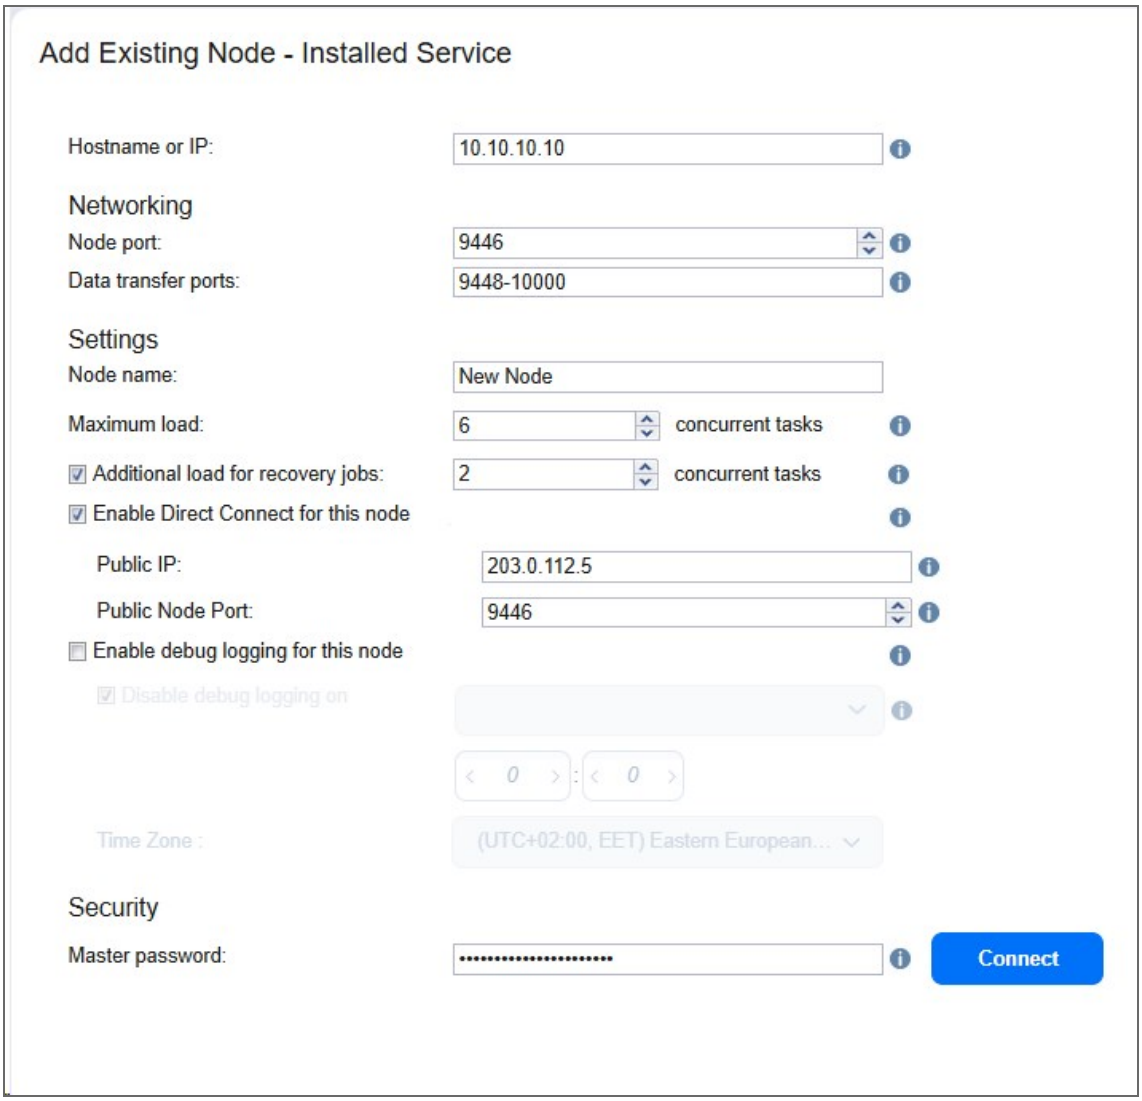
1. Download NAKIVO Backup & Replication installer or a VA with **Transporter**.
    1. Install Transporter at the remote environment and set the master password during the installation. Alternatively, deploy Transporter VA in the remote environment. The master password for the VA can be set after the deployment.

    Notes

    - For Windows, the installer UI and silent installations are supported.
    - For Linux, interactive and silent installations are supported.
    - For NAS, see the Supported NAS devices.
    - Master Password is required for Direct Connect Transporter installation.

2. If you didn't set up the master password during the installer setup or if you have deployed Transporter VA, you need to generate a pre-shared key. This can be done at any time after the installation. Execute the following command in **Transporter** installation folder:
   bhsvc -b "password"
3. Restart the Transporter service after generating the password to apply the change.
4. Expose the Transporter port on the local machine to be externally available via the Internet.
5. Add the **Transporter** to NAKIVO Backup & Replication by entering the **Transporter** details.
6. Select **Enable Direct Connect** for this transporter.
7. Enter **Public IP** and **Public Node Port** (IP and Port of the external **MSP Transporter** that is reachable from the internet and used for communication with the **Direct Connect Transporter**). Ensure that the entered combination of the **Public Node Port** and **Public IP** address is not used by other transporters.
8. Enter master password.
9. Click **Connect**.
10. View Transporter certificate information.
11. Accept received certificate.
12. Confirm adding the **Transporter**.

The client's **Direct Connect Transporter** is ready to be used for data protection activities.



MSP Direct Connect

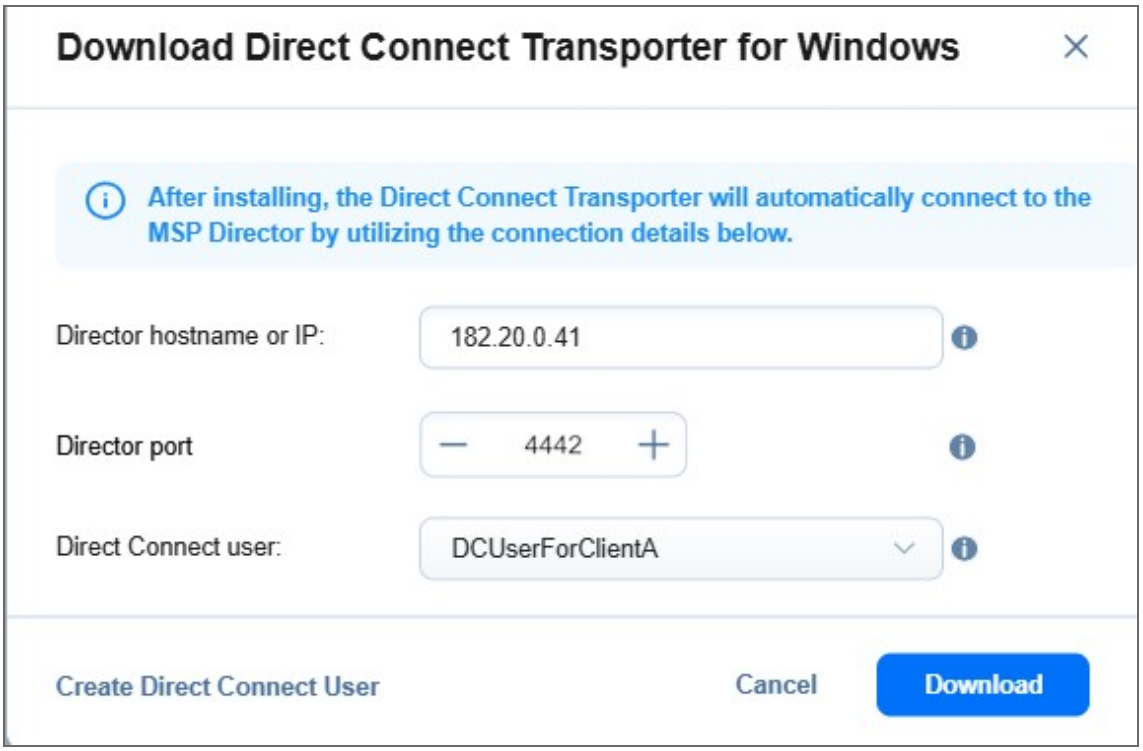
To enable MSP Direct Connect, the **Direct Connect Transporter** must be installed on the client's Windows/Linux operating system.

Installation and setting up the **Direct Connect Transporter** is performed according to the following flow:

1. Go to **Settings** > **Nodes**, click **Download**, and select **Direct Connect Transporter for Windows** or **Direct Connect Transporter for Linux** from the drop-down.

> **Note**
> The Direct Connect Transporter for Windows / Direct Connect Transporter for Linux options are disabled if your current license does not support the MSP Direct Connect feature.

2. In the dialog box, set the installation options as follows:



- **Director hostname or IP** field: Enter the IP address or hostname of the machine on which the **MSP Director** is installed. This should be an external address that is reachable from the Internet.
- **Director port** field: Enter the **MSP Director port** used for communication with the **Direct Connect Transporter**. By default, port 4442 is filled in automatically, so make sure the port is open and not being used by other applications. This setting is global and can only be changed inside the master tenant **Settings** > **System Settings** > **Configuration** tab.
  You can specify an external port that is reachable from the Internet instead of default value, for example:
  - Public port 10005 for TCP Port 4442 of MSP communication.
- **Direct Connect user** dropdown: Click to select the user with the **Direct Connect** role. By default, the newest created user with the **Direct Connect** role is assigned.
  Notes
  - This user is used to generate an application password required for authentication when the **Direct Connect Transporter** from the client site connects to the **MSP Director**. This option is disabled if no users with the **Direct Connect** role exist in the tenant.
  - To create a new user with the **Direct Connect** role, click the **Create Direct Connect user** link, and the **Add Local User** wizard that opens proceeds as described in Adding Local Users.
  - When the wizard is opened this way, the default value of the **Role** dropdown is set to **Direct Connect** and can't be changed.

3. Click **Download** to start downloading the installer.

> **Notes**
>
> - The **Download** button is disabled if the required dropdown/fields are not filled in or contain invalid values.
> - Optionally, click the **Cancel** or **X** button to close the dialog box without applying any changes.

4. After the Direct Connect Transporter installer is downloaded, run the installer at the client's site and follow the steps below:

> **Notes**
>
> - For Windows, the installer UI and silent installations are supported.
> - For Linux, interactive and silent installations are supported.
> - Master Password is required for Direct Connect Transporter installation.

5. After the **Direct Connect Transporter** is installed successfully at the client's site, the **MSP Director** automatically adds the **Direct Connect Transporter** to the **Nodes** tab of the local tenant from which it was downloaded.

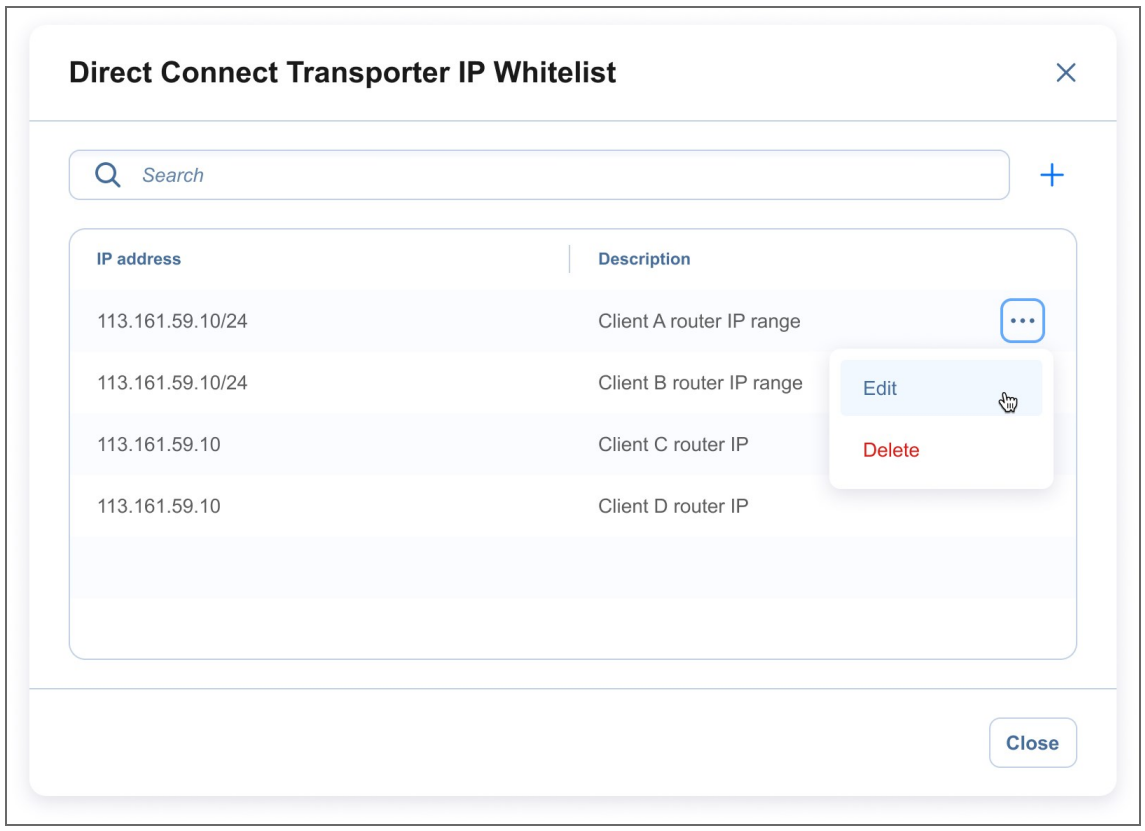
> **Note**
>
> If the installer is downloaded from the **Nodes** tab from the master tenant UI, the **Direct Connect Transporter** will be added there.

6. The added **Direct Connect Transporter** is displayed with the **Pending** status in the **Nodes** dashboard of the local tenant instance or master tenant, depending on where the installer was downloaded from.
7. Hover the Pending **Direct Connect Transporter**, click **Accept**, and enter the configured Master password to accept the **Direct Connect Transporter**. The client's Direct Connect Transporter is ready to be used for data protection activities.

Notes

- To back up items from the client's site to the MSP Repository, you must ensure at the MSP site, the **MSP Transporter** handling the backup must have Direct Connect enabled and be assigned to the MSP repository.
- You cannot add the **Direct Connect Transporter** manually.
- After the **Direct Connect Transporter** is installed, a config file is generated in the **Direct Connect Transporter** installation folder (**activedc.conf**). The **activedc.conf** file contains the following details about the connection to the **MSP Director**:
  - MSP Director hostname or IP
  - MSP Director port
  - Direct Connect user application password
  - MSP Director fingerprint
- The following permissions are required to edit the **activedc.conf** file on Linux:
  - Owner: bhsvc
  - Group: bhsvc
  - Permissions: 644
- If the **Direct Connect Transporter** does not appear in **MSP Director** after installation, check or modify this configuration file.

IP Whitelist Management



Only the **Direct Connect Transporters** from the specified IP addresses can connect to the **MSP Director** when enabled. If no whitelist is configured, any **Direct Connect Transporter** can connect to the **MSP Director** as long as its config file contains matching values (MSP Director IP, port, application password).
To do it, follow the steps below:

1. Go to the **General** > **System Settings** > **Configuration** tab.
2. Select the **Allow Direct Connect Transporter connections from specific IP addresses only** checkbox. The **Settings** link appears.
3. Click the **Settings** link.
4. In the **Direct Connect Transporter IP Whitelist** dialog box, click the **(+)** button to add the IP address to the whitelist.
5. Enter the required data and click **Add**. The IP address is added to the whitelist.

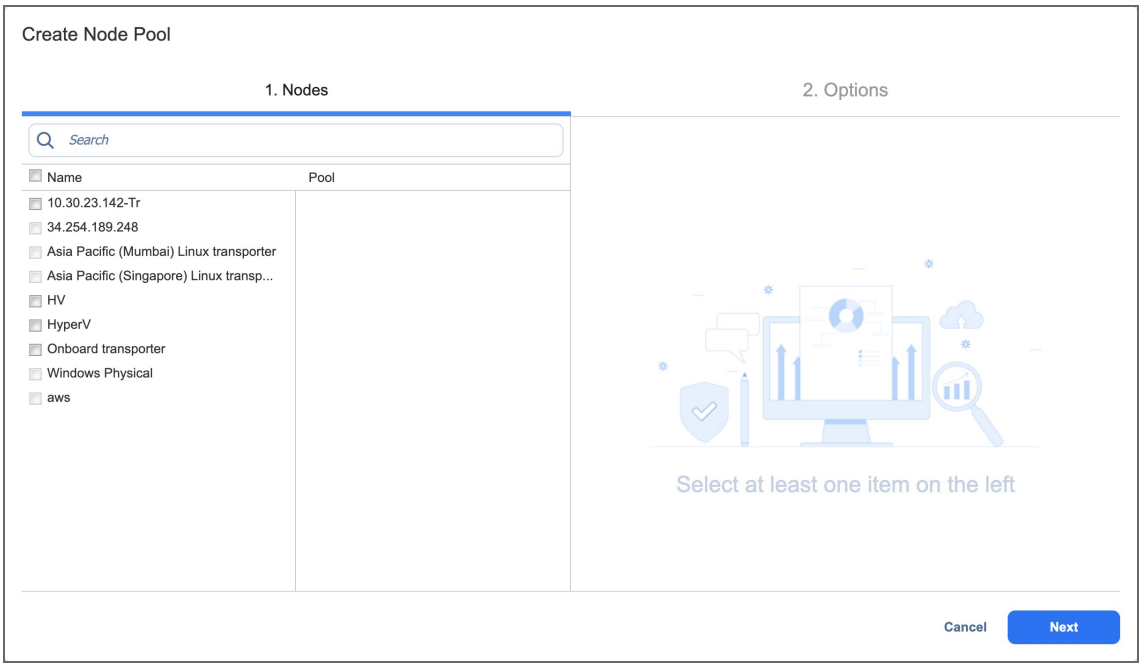
Notes

- If the IP address was on the whitelist and you removed it, the Direct Connect Transporter from that IP address remains connected to the MSP Director until you manually remove it from the Nodes tab.
- The whitelist can be configured at either the master tenant or the local tenant level, depending on where the Direct Connect Transporter will be added.

Managing Node Pools

NAKIVO Backup & Replication allows you to group nodes into pools to optimize backup, replication, and recovery jobs. To create a node pool, do the following :

1. Navigate to **Settings**.
2. Click the **Nodes** tab.
3. Open the **Node Pools** tab, then click the plus **Add** button.
4. Complete the **Create Node Pool** wizard and click **Finish**.



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

Refreshing Node Details

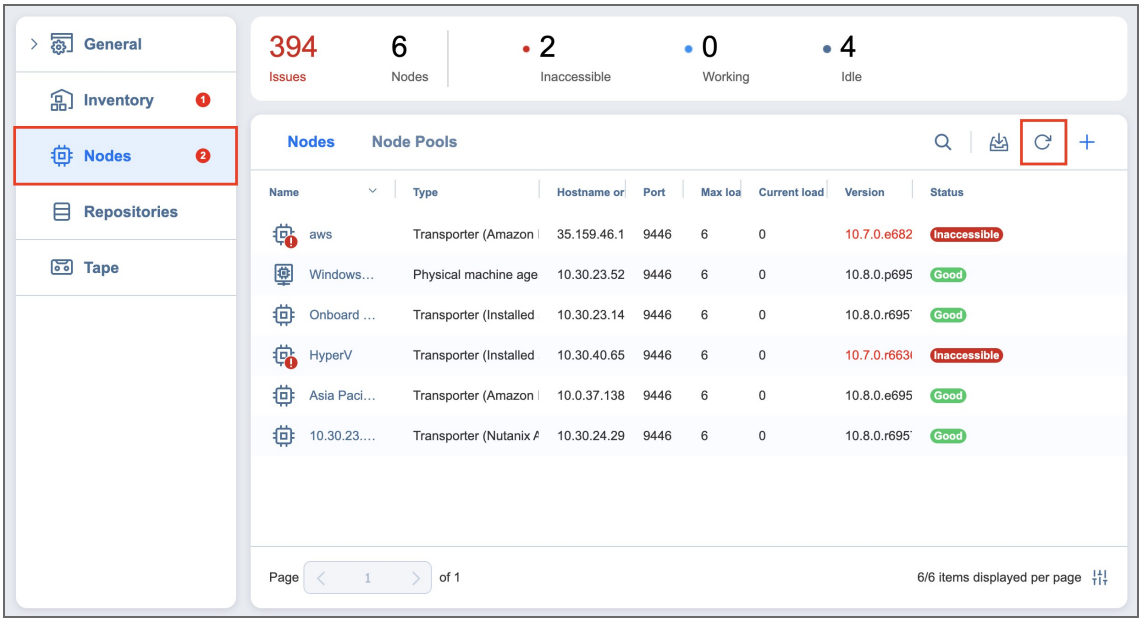
By default, NAKIVO Backup & Replication refreshes the information about **Transporters** every hour. During the refreshing process, the product collects all the required information about all transporters. Only one **Transporter** can be refreshed at a time. If you have more than one **Transporter**,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

- Manually Refreshing All Nodes
- Manually Refreshing a Single Node

Manually Refreshing All Nodes

To refresh all nodes, follow the steps below:

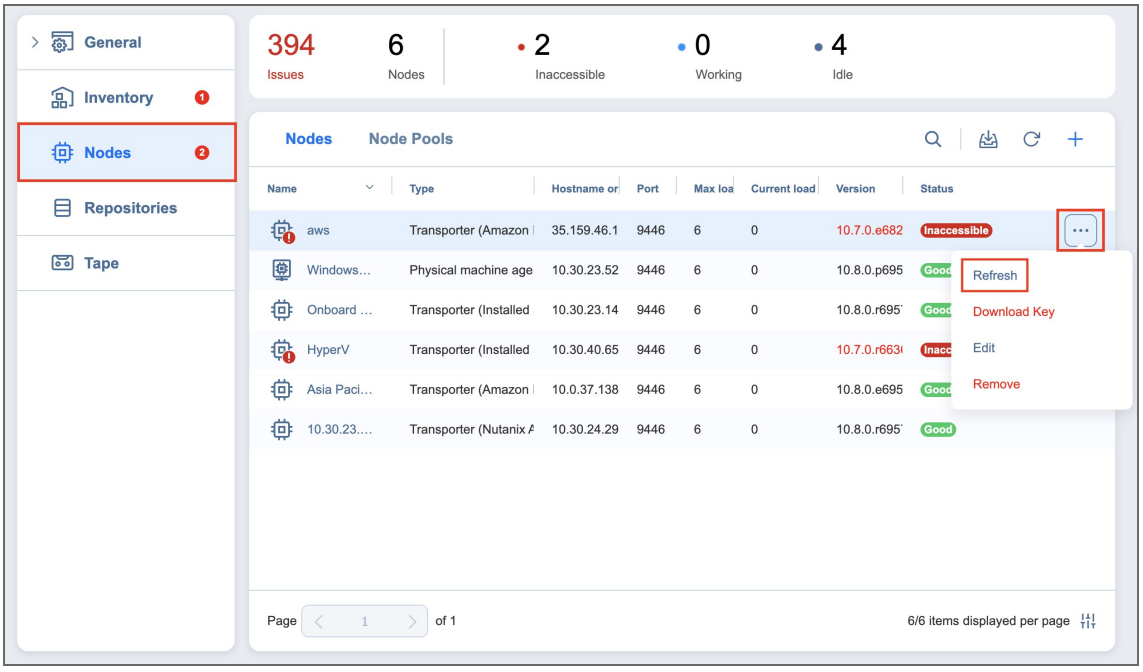
1. Click **Settings** in the left pane of the product and go to the **Nodes** tab.
2. Click the **Refresh** button above the **Nodes** table.



Manually Refreshing a Single Node

To refresh a single node, follow the steps below:

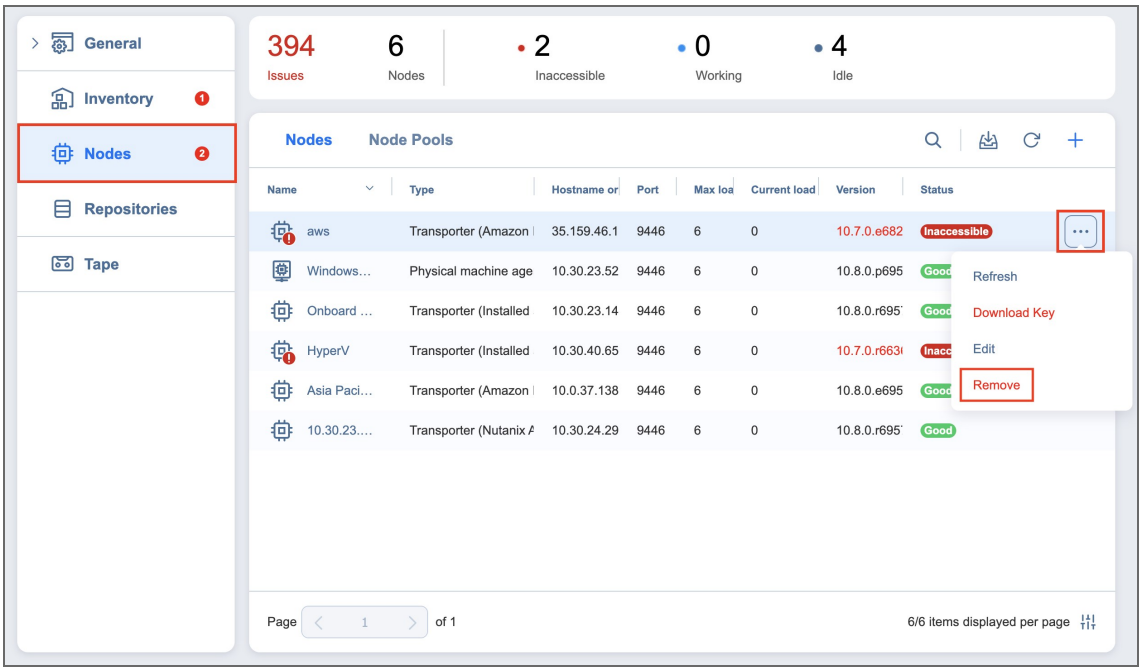
1. Click **Settings** in the left pane of the product.
2. Go to the **Nodes** tab.
3. Hover over the node you would like to refresh and click the ellipsis **Manage** button.
4. Click **Refresh**.



Removing (Deleting) Nodes

To remove a Transporter from NAKIVO Backup & Replication, follow the steps below:

1. Click **Settings** in the left pane of the product.
2. Go to the **Nodes** tab.
3. Hover over the node you would like to remove.
4. On the right side, click the ellipsis **Manage** button and then click **Remove**.



Note

The following nodes cannot be removed:

- The Onboard **Transporter** (which is installed with the **"Director" on page 49** by default)
- Nodes manually assigned to a job
- **Transporters** assigned to backup repositories

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 55.

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

- “Adding Existing Backup Repositories” on page 461
- “Creating Backup Repositories” on page 463
- “Managing Backup Repositories” on page 467
- “Viewing Backup Repository Details” on page 478

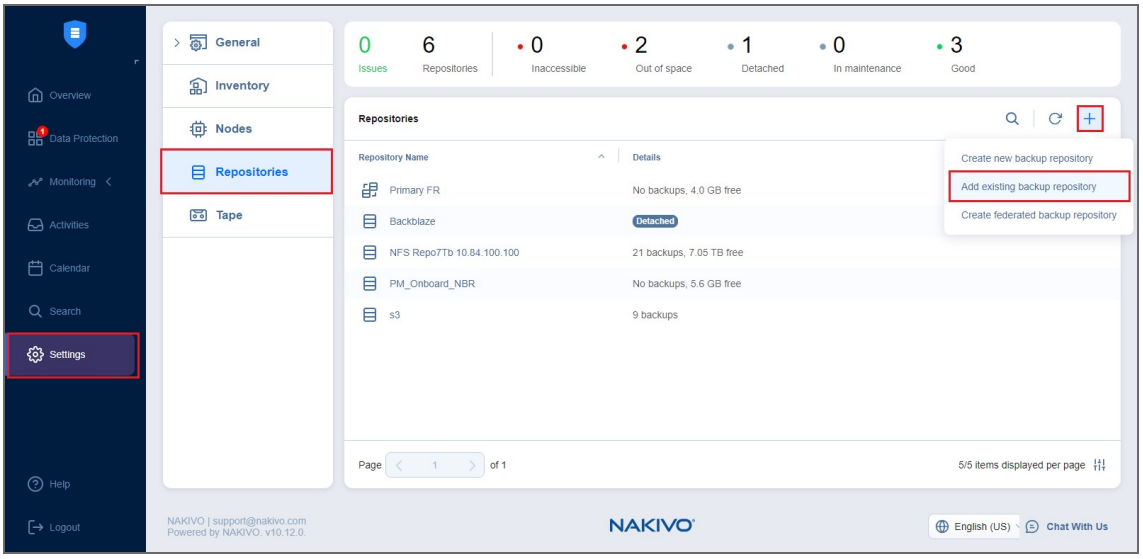
Adding Existing Backup Repositories

NAKIVO Backup & Replication allows you to add an existing **Backup Repository** to a new copy of the product.

> **Note**
> During the import process, NAKIVO Backup & Replication searches for the *NakivoBackup* folder in the specified location. If your **Backup Repository** is located in *E:\backup\NakivoBackup*, you should specify the following path: *E:\backup*

To import an existing **Backup Repository**, do the following:

1. In the main menu, click **Settings**.
2. Go to the **Repositories** tab and click **+**.
3. Click **Add existing backup repository** in the dialog box that opens.



4. The **Add Existing Backup Repository** wizard opens. On the **Type** page of the wizard, select the following **Backup Repository** :
   - **SaaS**
5. When you select **Cloud**, the **Vendor** page opens. Select the cloud storage vendor from the following options:
   - Amazon S3
   - Microsoft Azure
   - Wasabi
   - Backblaze
   - Amazon EC2
   - Generic S3-compatible storage

6. When you select **Deduplication Appliance**, the **Device** page opens. Select the device from the following options:
    - Dell EMC Data Domain Boost
    - HPE StoreOnce Catalyst
    - NEC HYDRAstor
7. On the **Name & Location** page of the wizard, fill out all the necessary fields as described in the article for the corresponding **Backup Repository** type.
8. On the **Options** page of the wizard, the following options can be available for configuration:
    - 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.
    - **Detach this repository on schedule**: Select this option if you want to detach and then reattach the **Backup Repository** based on a schedule. Detaching a **Backup Repository** saves its data and metadata in a consistent state and stops the product interaction with the **Backup Repository** (so that it can be copied or moved). You can use this feature, for example, for the disk-to-disk-to-tape (D2D2T) data protection approach: backups are stored on a disk for fast operational recovery and copied to a tape (while the repository is detached) for archiving and long-term storage.
    - **Delete and re-create the repository on attach**: When this option is selected, all data in the **Backup Repository** is erased prior to attaching it to the product. As a result, jobs that write to this **Backup Repository** create full VM backups. You can use this option, for example, to create full daily, weekly, or monthly VM backups and write them to tape or removable media.
9. Click **Finish**. The **Backup Repository** is imported to the list.

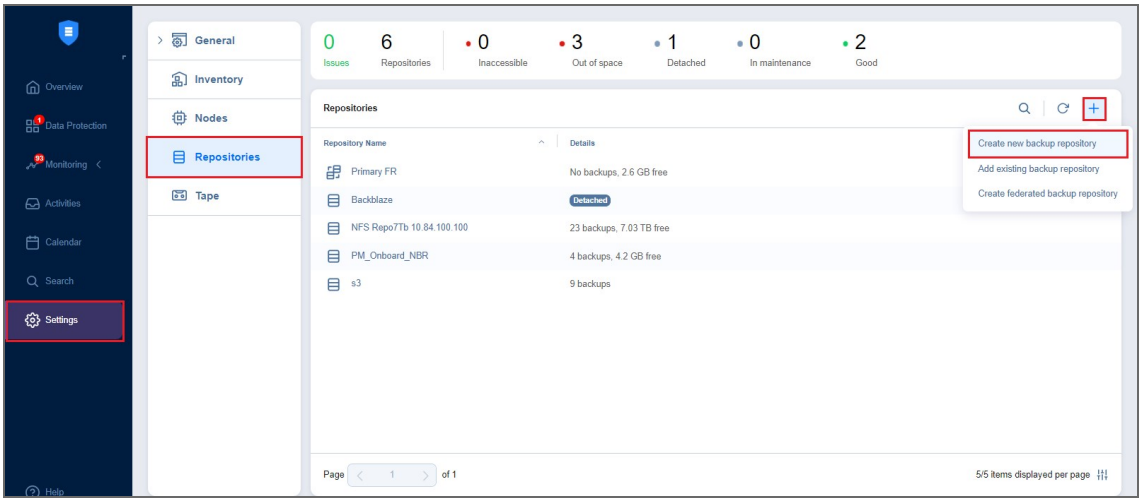
Creating Backup Repositories

NAKIVO Backup & Replication allows you to create additional **Backup Repositorie**s for storing backups. You can use a local folder, NFS share, CIFS share, public cloud, or deduplication appliance as a **Backup Repository** location. To create a new **Backup Repository**, follow the steps below.

Important

Do not create **Backup Repositorie**s inside NAKIVO Backup & Replication installation folders. The data inside **Director** and **Transporter** folders may be lost after a solution update.

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



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

- “SaaS Backup Repository” on page 464

SaaS Backup Repository

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

Note

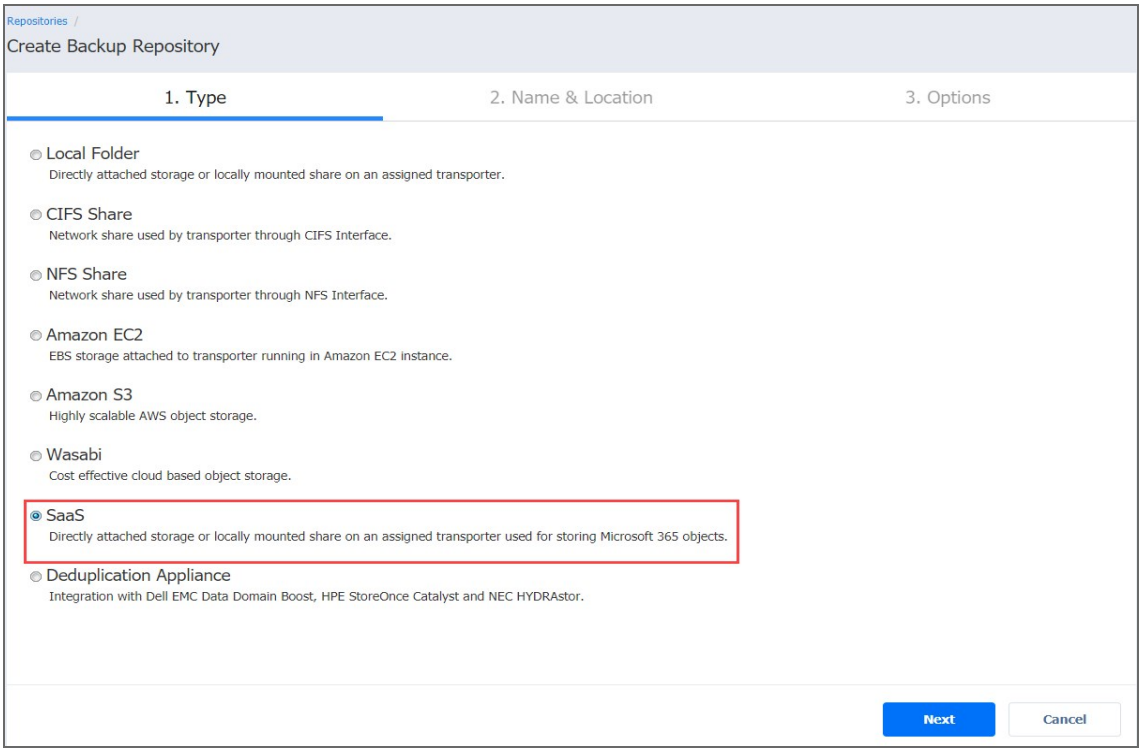
- For SaaS repositories, only directly attached storage or locally mounted shares are supported as a location.
- Before creating a SaaS repository, provide read and write permissions to the local folder where the repository will be located.
- For Office 365 SaaS repositories, manual removal of backup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

To create a SaaS **Backup Repository**, proceed as described in the following sections:

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

Create Backup Repository: Type

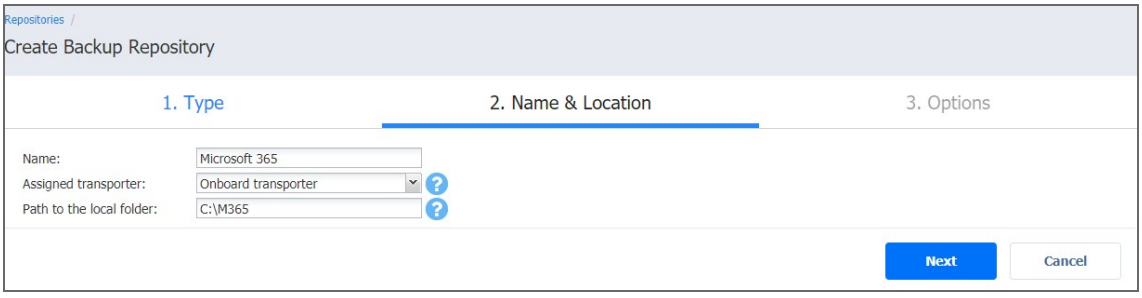
On the **Type** page of the **Create Backup Repository** wizard, select **SaaS** and click **Next** to go to the next page of the wizard.



Create Backup Repository: Name and Location

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

1. Enter the name of the **Backup Repository** in the **Name** box.
2. Select the **Transporter** from the **Assigned transporter** drop-down list.
3. Enter a path to the local folder in the corresponding box.
4. Click **Next** to go to the next page of the wizard.



Create Backup Repository: Options

On the **Options** page:

1. Schedule detaching of the **Backup Repository**
    - **Detach this repository on schedule:** Select this option if you want to detach and then attach the **Backup Repository** on a schedule. Detaching a **Backup Repository** saves its data and metadata in a consistent state and then stops the product's interaction with the **Backup Repository** (so it can be copied or moved). You can use this feature, for example, for the disk-to-disk-to-tape (D2D2T) data protection approach, in which backups are stored on a disk for fast operational recovery, and copied to a tape (while the repository is detached) for archiving and long-term storage.
    - **Delete and re-create the repository on attach:** If this option is selected, all data in the **Backup Repository** will be erased prior to attaching it to the product. As a result, jobs that write to this Backup Repository will create full VM backups. You can use this option, for example, to create full daily, weekly, or monthly VM backups and write them to tape or removable media.
2. Click **Finish** to finish creating the **Backup Repository**.



Managing Backup Repositories

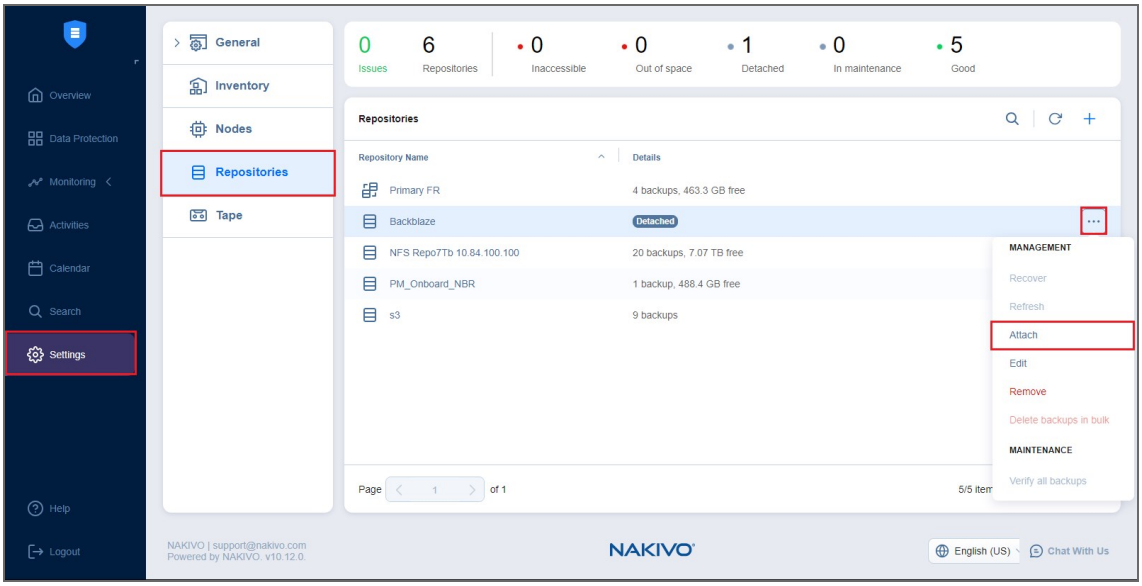
Refer to the following topics:

- “Attaching Backup Repositories” on page 468
- “Deleting Backups in Bulk” on page 468
- “Detaching Backup Repositories” on page 471
- “Editing Backup Repositories” on page 472
- “Managing Backup Encryption” on page 303
- “Refreshing Backup Repositories” on page 475
- “Removing and Deleting Backup Repositories” on page 477

Attaching Backup Repositories

If you have detached a **Backup Repository**, you can reattach it to the product by following the steps below:

1. From the main menu of NAKIVO Backup & Replication, click **Settings**.
2. Go to the **Repositories** tab and hover over a **Backup Repository**.
3. On the right side, click ••• and then click **Attach.**

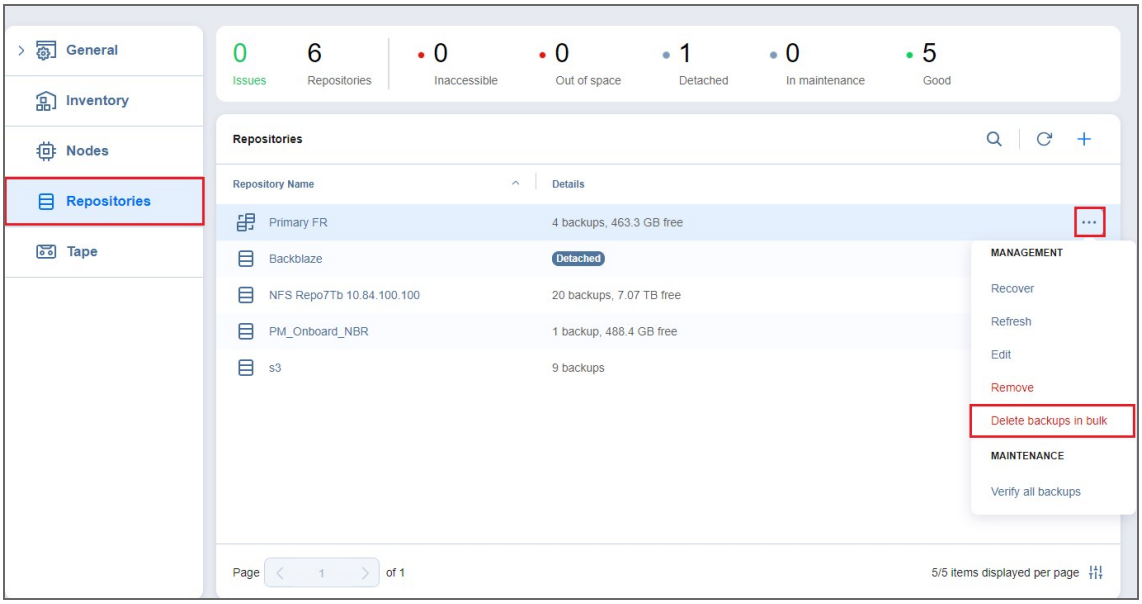


The **Backup Repository** is reattached to NAKIVO Backup & Replication. You can now back up to the attached **Backup Repository**.

Deleting Backups in Bulk

To permanently delete several backups that match specific criteria, follow the steps below:

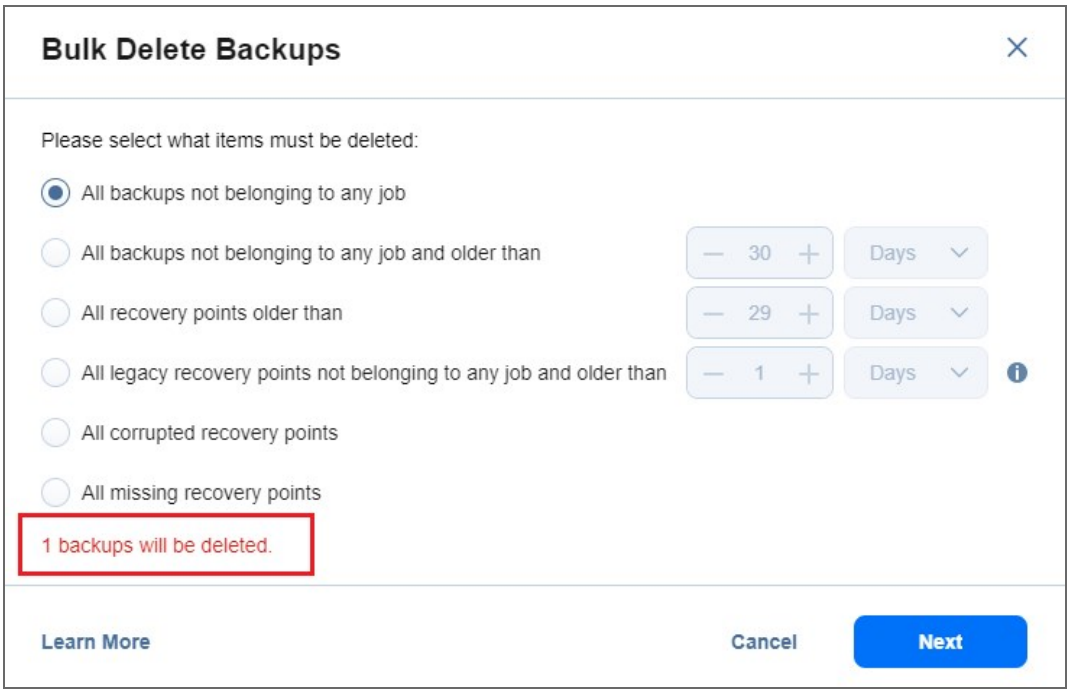
1. In the main menu of NAKIVO Backup & Replication, click **Settings**.
2. Go to the **Repositories** tab.
3. Hover over a repository.
4. On the right side, click the ellipsis **Manage** button and then click **Delete backups in bulk**.



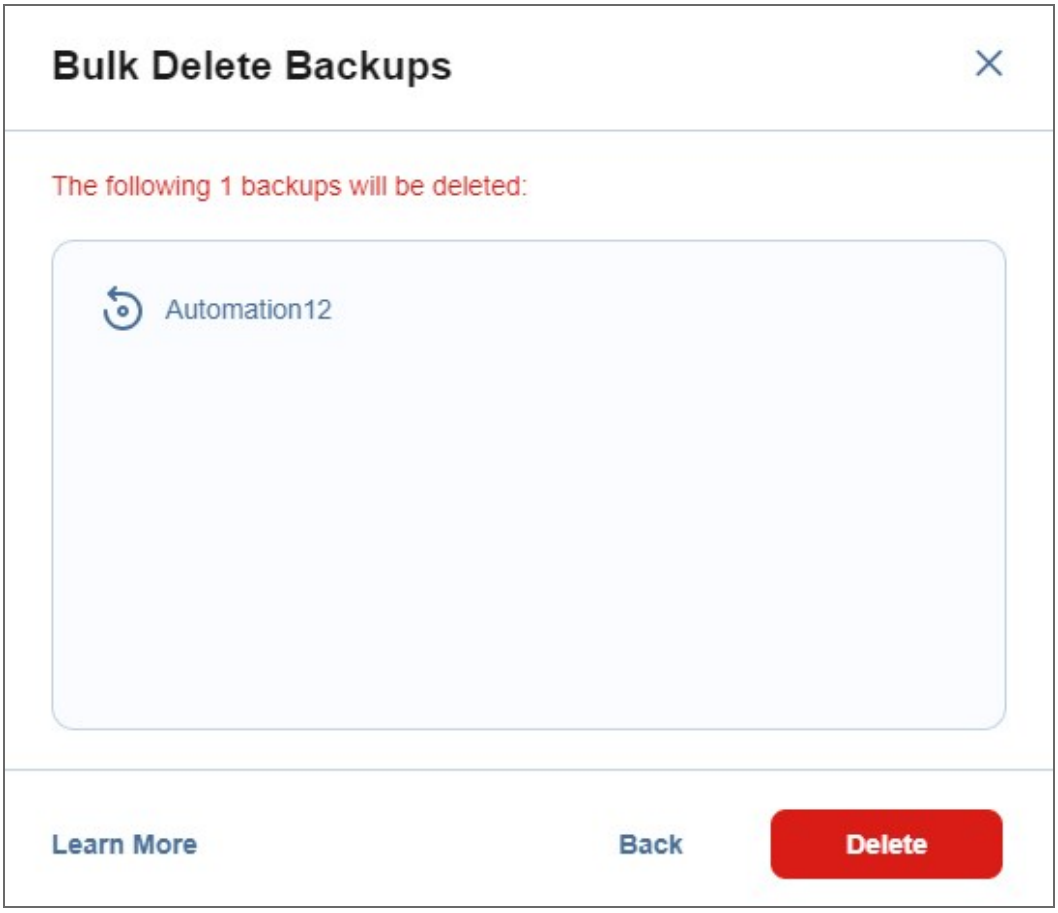
5. In the **Bulk Delete Backups** dialog box that opens, select one of the available options:

- **All backups not belonging to any job**
- **All backups not belonging to any job and older than** X <time_units>, where X is an integer and <time_units> is either days, weeks, or months
- **All recovery points older than** X <time_units>, where X is an integer and <time_units> is either days, weeks, or months
- **All legacy recovery points not belonging to any job and older than** X <time_units>, where X is an integer and <time_units> is either days, weeks, or months
- **All corrupted recovery points**
- **All missing recovery points**

6. The dialog shows the number of backups to be deleted.



7. Click **Next**.
8. The **Bulk Delete Backups** dialog box opens displaying the list of backups to be deleted.
9. Click **Delete** to confirm the deletion of backups.



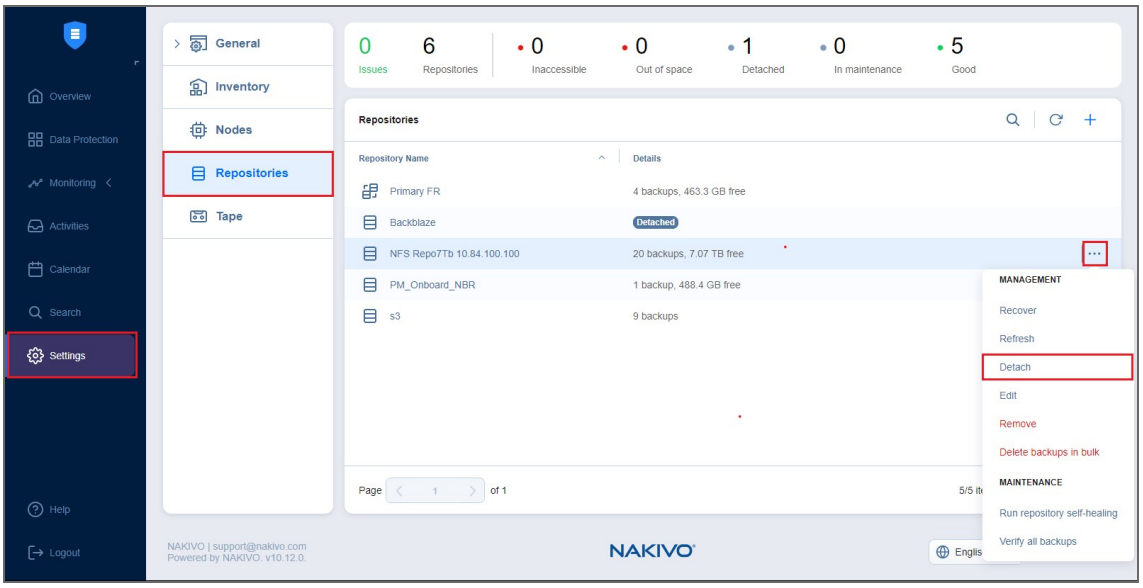
Detaching Backup Repositories

Detaching a **Backup Repository** saves its data and metadata in a consistent state and stops the product's interaction with the repository (e.g. reading and writing of data or metadata). You may want to detach a **Backup Repository** in order to move it to a different location or to put the associated storage in maintenance.

> **Note**
> As the product does not interact with detached repositories, jobs with detached **Backup Repositories** as target storage will fail.

To detach a **Backup Repository**, follow the steps below:

1. From the main menu, click **Settings**.
2. Go to the **Repositories** tab and hover over a **Backup Repository**.
3. On the right side, click ••• and then click **Detach**.



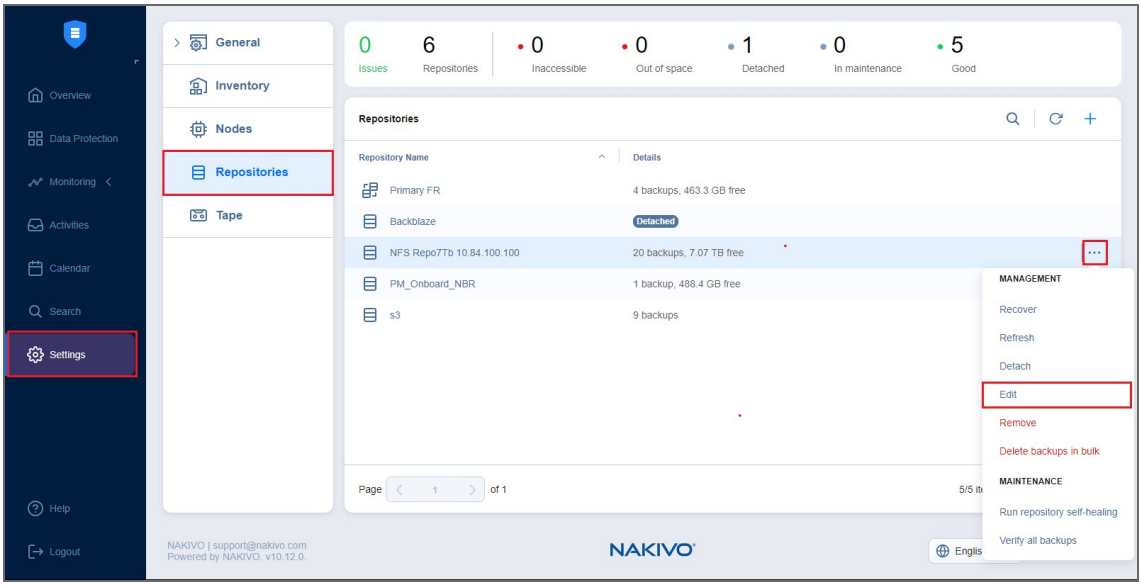
> **Note**
> A **Backup Repository** cannot be detached if a job that backs up to this **Backup Repository** is running.

The **Backup Repository** is detached from the product. You can reattach the **Backup Repository** to NAKIVO Backup & Replication when needed.

Editing Backup Repositories

To modify the settings of an existing **Backup Repository**, follow the steps below:

1. From the main menu of NAKIVO Backup & Replication, click **Settings**.
2. Go to the **Repositories** tab and hover over a **Backup Repository**.
3. On the right side, click ••• and then click **Edit**.



> **Note**
>
> A **Backup Repository** cannot be edited while a job that backs up to this **Backup Repository** is running.

4. Update the fields as necessary.
5. Click **Apply**. Changes you have made are applied and the **Backup Repository** update starts.

Reclaiming Backup Repository Space

When a backup or recovery point is deleted in a **Backup Repository**, the space occupied by that backup or recovery point is marked as “free” and can be reused by new data blocks on the next job runs. However, the actual size of the **Backup Repository** may not change. The size of a **Backup Repository** can be reduced by rearranging the data blocks so there are no “free” ones occupying storage space. The amount of space that can be freed up is displayed in parentheses after the amount of used space. This is applicable if the repository type is **Forever-incremental**. Otherwise, if the repository type is **Incremental with full backups**, space reclaiming is not required. It is enough to delete the backups or recovery points to free up space and continue backing up to the repository.

For the incremental with full backup **Backup Repository** type, it is technically impossible to remove recovery points if there is no full backup after them. Make a full backup before deleting older recovery points.

Reclaiming free space can take the same amount of time as copying the entire **Backup Repository** to the storage where it is located (that is, if your repository size is 500 GB, reclaiming free space can take the same amount of time as copying 500GB of data to the storage where the **Backup Repository** is located).

Refer to the following topic to learn how to start and stop the reclaiming process:

- Starting the Space Reclaiming Process
- Stopping the Space Reclaiming Process

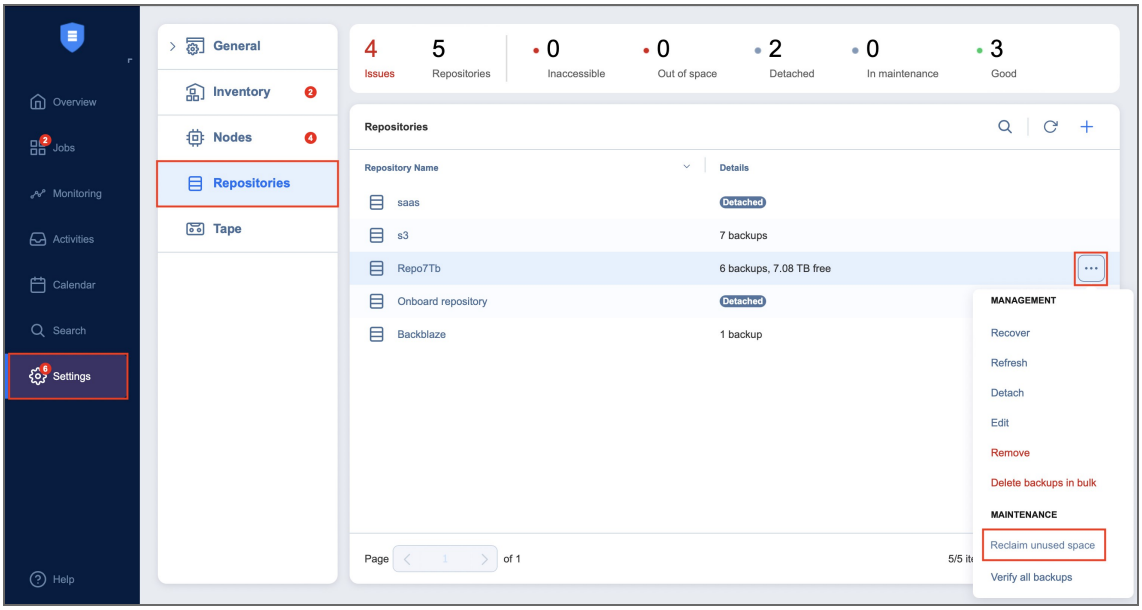
Starting the Space Reclaiming Process

Important

Space reclaim requires at least 500 MB of free space on the repository storage in order to start.

To reclaim free space, follow the steps below:

1. Go to the main menu of NAKIVO Backup & Replication and click **Settings**.
2. Go to the **Repositories** tab and choose a **Backup Repository**.
3. In the title of the **Backup Repository**, click ••• and then click **Reclaim unused space.**



The space reclaiming process cannot be started if a job that backs up to this **Backup Repository** is concurrently running.

4. In the dialog box that opens, leave the **Interrupt space reclaim task if backup or recovery is started** option selected to pause the space reclaiming process when a backup or recovery is started. The space reclaiming process will be resumed once the backup or recovery job is completed. If you deselect the option, backup jobs will fail and recovery jobs will not start until the space reclaim process is completed.
5. Click **Start**. The process of rearranging data blocks is started, and progress is displayed in the title of the **Backup Repository**.

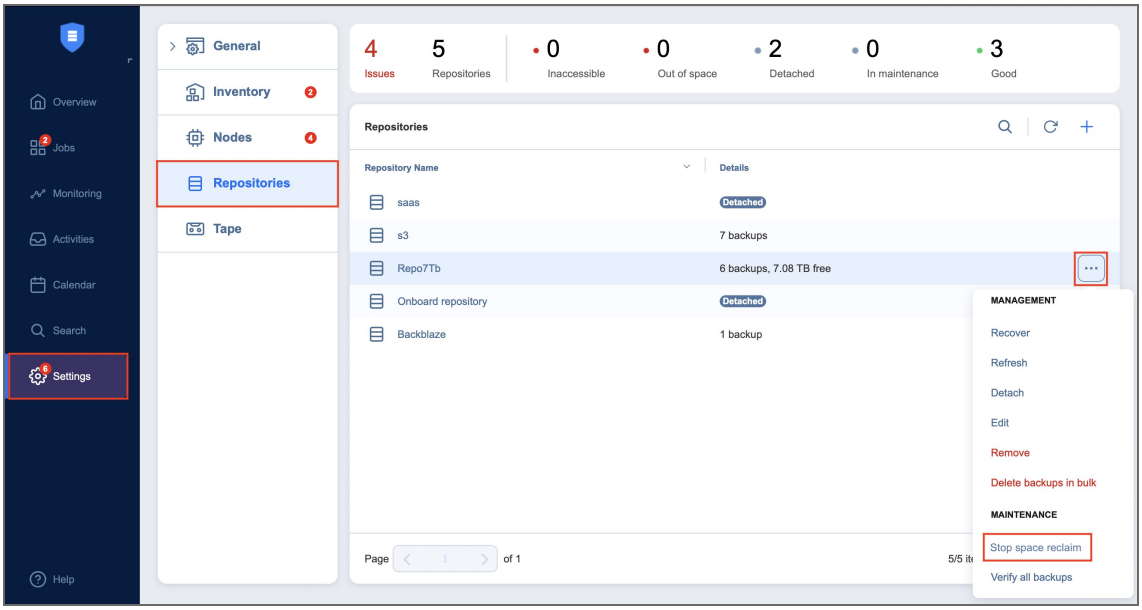
Stopping the Space Reclaiming Process

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

Before the space reclaiming process begins, the **Backup Repository** is detached from the product to keep data in a consistent state. The space reclaiming process stops if job that backs up VMs to such a **Backup Repository** is started and resumes after it is finished.

To stop the space reclaim process, follow the steps below:

1. Go to the main menu of NAKIVO Backup & Replication and click **Settings**.
2. Go to the **Repositories** tab and choose a **Backup Repository**.
3. In the title of the **Backup Repository**, click ••• and then click **Stop space reclaim**.



Refreshing Backup Repositories

By default, NAKIVO Backup & Replication refreshes **Backup Repository** information hourly. During the refreshing process, the product collects all required information about **Backup Repositories**, such as the amount of free space, number of backups, and number of recovery points.

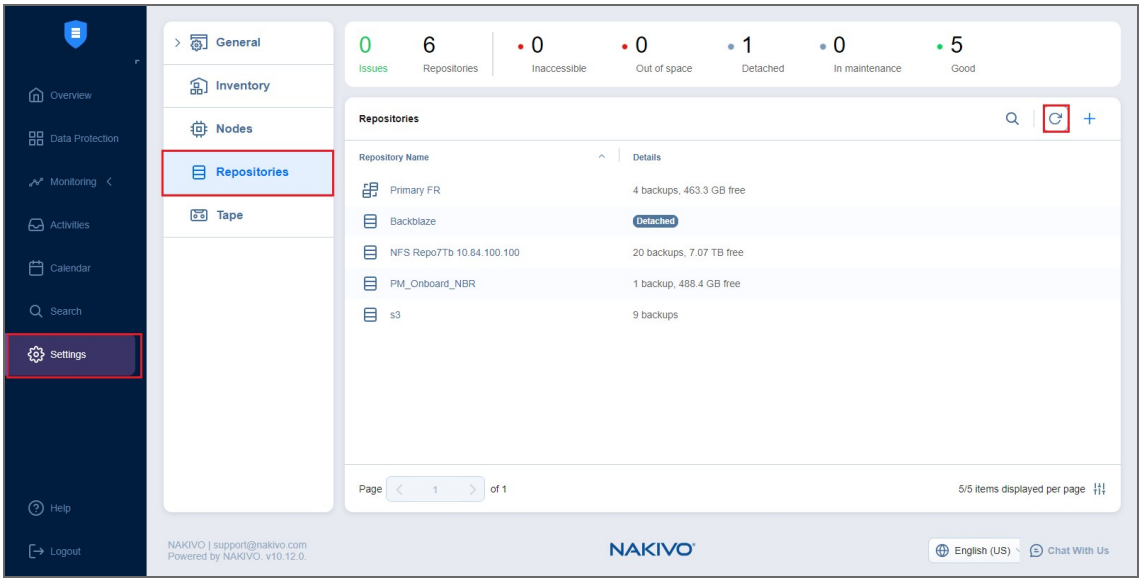
Only one **Backup Repository** is refreshed at a time. Therefore, if you attempt to refresh multiple **Backup Repositories**, all but one will be added to a queue.

- Refreshing All Backup Repositories
- Refreshing a Single Backup Repository

Refreshing All Backup Repositories

To refresh all backup repositories, follow the steps below:

1. From the main menu of NAKIVO Backup & Replication, click **Settings**.
2. Go to the **Repositories** tab.
3. Click the **Refresh All** button.



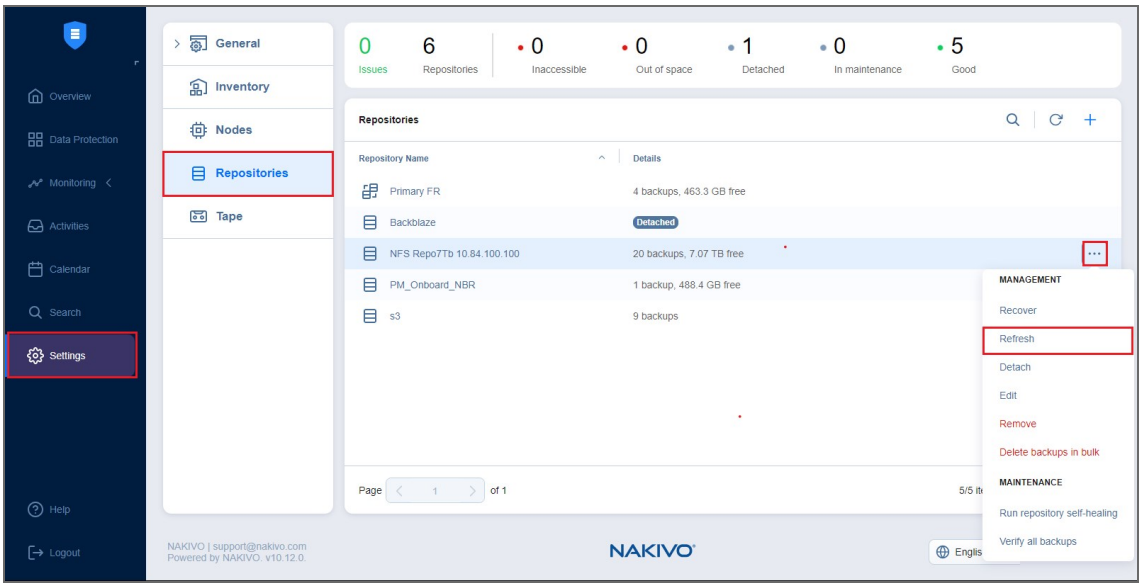
The **Backup Repository** refresh process begins.

Refreshing a Single Backup Repository

To refresh a single **Backup Repository**, follow the steps below:

1. From the main menu of NAKIVO Backup & Replication, click **Settings**.
2. Go to the **Repositories** tab.
3. Hover over the **Backup Repository** that you wish to refresh and click •••.

4. Click **Refresh**.



The **Backup Repository** refresh begins.

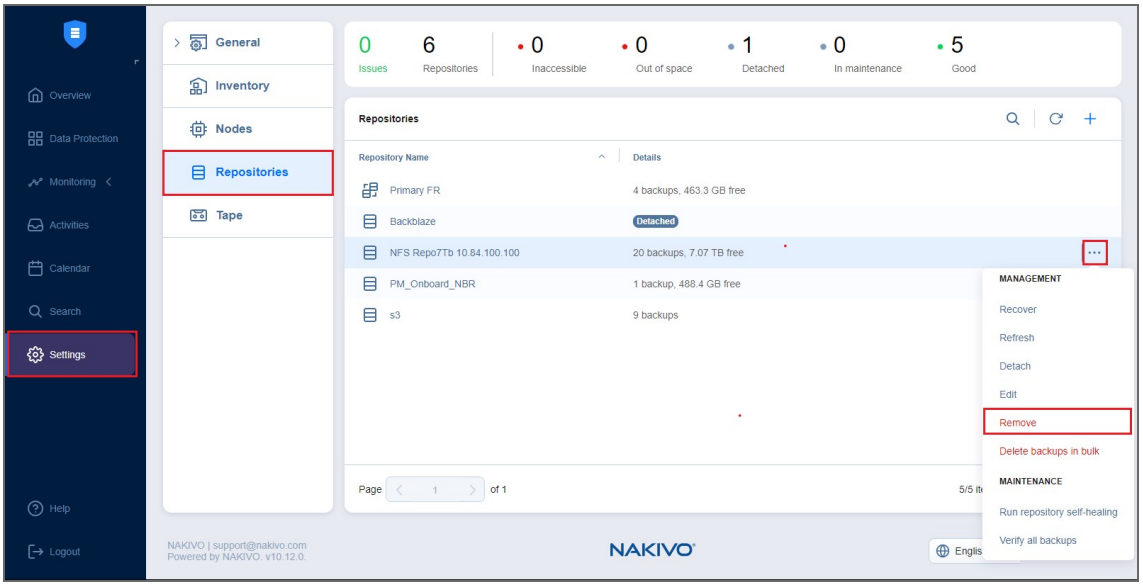
Removing and Deleting Backup Repositories

In NAKIVO Backup & Replication, you can either permanently delete a **Backup Repository** and all of its data or remove only the **Backup Repository** from the product while maintaining all of its data. After removing a **Backup Repository** you will be able to import it into the same or a new instance of the product.

> **Note**
> You will not be able to remove a **Backup Repository** if there is a job that backs up to this **Backup Repository**. To remove such a **Backup Repository**, delete (or edit) the corresponding jobs so no items are backed up to the aforementioned repository.

To permanently delete or remove a **Backup Repository** from the product, follow the steps below:

1. From the main menu of NAKIVO Backup & Replication, click **Settings**.
2. Go to the **Repositories** tab.
3. Hover over a **Backup Repository.**
4. On the right side, click **Manage** and then click **Remove.**



5. Do the following when the confirmation message appears:
   - To remove the **Backup Repository** from NAKIVO Backup & Replication and keep the **Backup Repository** on a disk, select **Remove repository and keep backups**.

     > **Note**
     > You can import the removed **Backup Repository** back to the same instance or to a new installation.

   - To permanently delete the **Backup Repository** and all its data, select **Remove repository and delete backups**.

Note

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

Viewing Backup Repository Details

To view the instance **Backup Repositories** details, follow the steps below:

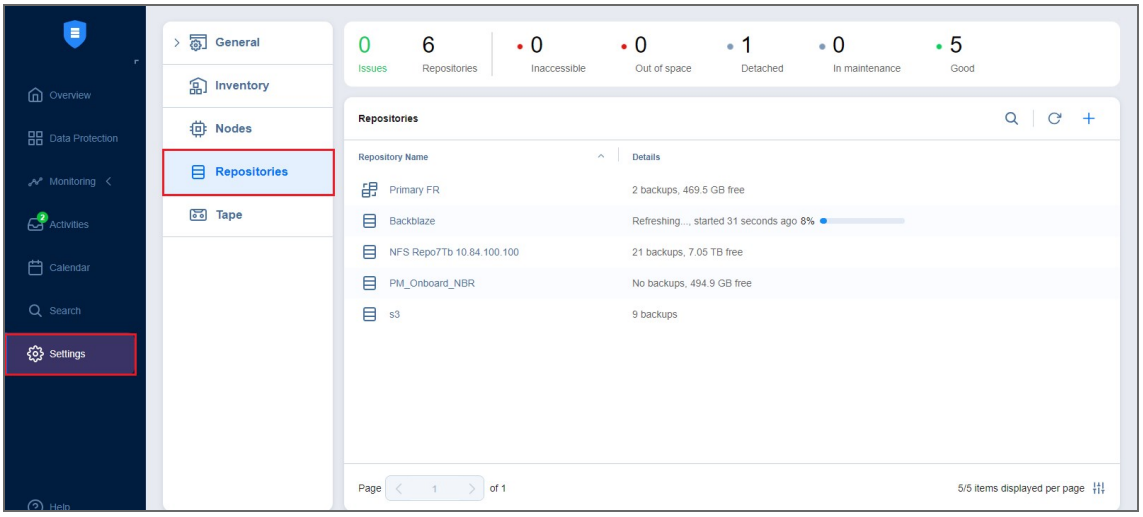
1. Go to the main menu of NAKIVO Backup & Replication and click **Settings**.
2. Go to the **Repositories** tab.

The **Repositories** tab contains a **Summary** bar, which offers an overview of all backup repositories both federated and standalone (not used as members of federated repositories). The data displayed is as follows:

- **Issues**: Total number of issues/alarms related to repositories
- **Repositories**: Total number of repositories
- **Inaccessible**: Total number of inaccessible repositories
- **Out of Space**: Total number of repositories that are out of storage space
- **Detached**: Total number of detached repositories
- **In Maintenance**: Total number of repositories in maintenance
- **Good**: Total number of usable repositories

To see information about specific standalone repositories, backups, and recovery points, see the sections below.

To see information about federated repositories and their members as well as the backups and recovery points assigned to them, refer to Viewing Federated Repository.



In this view, you can also search for a specific repository by entering its name (or part of its name) into the **Search** box.

In the **Repositories** table below, you can see information about specific repositories, their backups and recovery points.

The table data is displayed as follows:

- **Repository Name**: The name of the repository
- **Details**: The information about the repository:
  - Number of backups
  - Amount of free space in the backup repository
  - Repository status, which can be detached or inaccessible
  - The progress bar displayed when any of in progress statuses is in action (for example, self-healing or verifying backups)

To manage the **Backup Repository**, hover over it, and on the right side, click the ellipsis **Manage** button, and then select the necessary option in the drop-down list. For more details, refer to "Managing Backup Repositories" on page 467.

Refer to the following sections to learn more about viewing Backup Repository details:

- "Viewing Standalone Backup Repository Details" below
- "Viewing Backup Details" on page 481
- "Viewing Recovery Point Details" on page 484

Viewing Standalone Backup Repository Details

To view the details for a standalone **Backup Repository**, follow the steps below:

1. Go to the main menu of NAKIVO Backup & Replication and click **Settings**.
2. Go to the **Repositories** tab.

> **Note**
>
> All repositories added to the product are listed in the **Repositories** table in the following order:
>
> - Federated repositories (for more details, refer to Federated Repositories)
> - Standalone repositories (not used as members of federated repositories)

3. Click a standalone **Backup Repository**  to open it in a new window.

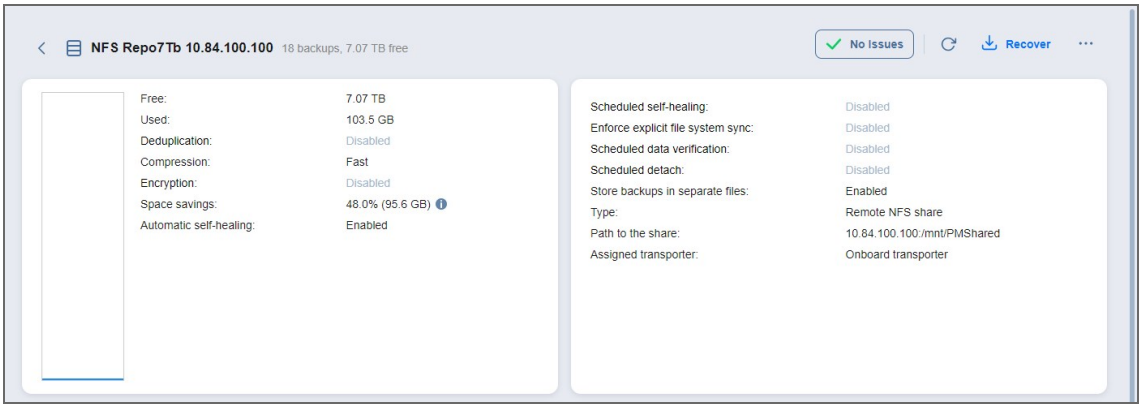
The following data is displayed:

- **Free**: The amount of free space available for the **Backup Repository**

> **Note**
>
> If the repositories are placed on the same same disk/share/etc., the free space metric can be incorrect.

- **Used**: The amount of space that the **Backup Repository** occupies on a disk. The amount of space that can be reclaimed is displayed in parentheses.

- **Deduplication**: The status of deduplication in the **Backup Repository**
- **Compression**: The compression level specified for the **Backup Repository**
- **Encryption**: The status of encryption in the **Backup Repository**
- **Space savings**: The estimated percentage and amount of space saved by compression and deduplication. For example, if 200 GB of data were backed up and the size of the backup was reduced to 50 GB, the ratio is calculated as 75%.
- **Automatic self-healing**: The state of the automatic self-healing option for the **Backup Repository**
- **Scheduled self-healing**: The state of the scheduled self-healing option for the **Backup Repository**
- **Enforce explicit file system sync**: The state of the enforce explicit file system sync option for the **Backup Repository**
- **Scheduled data verification**: The state of the scheduled data verification option for the **Backup Repository**
- **Scheduled space reclaiming**: The state of the scheduled space reclaiming option for the **Backup Repository**
- **Scheduled detach**: The state of the scheduled detach option for the **Backup Repository**
- **Store backups in separate files**: The behavior of the **Backup Repository** on backup data storage
- **Type**: The location of the **Backup Repository**, which can be one of the following:
    - **Local folder on assigned Transporter**
    - **Remote CIFS Share**
    - **Remote NFS Share**
    - **Amazon EC2**
    - **Microsoft 365**
    - **Microsoft Azure Blob Storage**
    - **Amazon S3**
    - **Generic S3-Compatible Storage**
    - **Wasabi**
    - **Backblaze B2 Cloud Storage**
    - **Deduplication Appliance**
- **Path to the folder**: The path to the **Backup Repository** folder
- **Assigned transporter**: The **Transporter** that manages the **Backup Repository** (that is, the **Transporter** that reads data from and writes data to the **Backup Repository**)
- **Backups**: List of available backups in the **Backup Repository**



Note

Standalone repositories used as federated repository members are not displayed in the **Repositories** table.

To view federated repository members, select the **Federated repository** to open the Federated repository details dashboard. In the **Storage** table, click the federated repository member name to open the **Backup Repository details** screen.

Viewing Backup Details

Below, in the **Backups** table, you can view the details of the backups stored in the selected Backup Repository.

Large numbers of backups are split into pages. To find a specific backup, you can scroll through the pages manually or search for a specific backup by entering its name (or part of its name) into the **Search** box. The table dynamically changes to display the search results matching your query. Clicking the **Clear** button in the search box clears the query, and the table displays all backups.

The **Backups** table provides the following detailed information about each backup:

- **Name**: The name of the backup
- **Encryption**: The state of encryption of the backup object, which can be one of the following:
  - **Not encrypted**: None of the recovery points of the backup object are encrypted.
  - **Encrypted**: All the recovery points of the backup object are encrypted.
  - **Partially encrypted**: Some recovery points of the backup object are encrypted.
- **Password**: The state of the password of the backup object, which can be one of the following:
  - **Available**: All or some recovery points of the backup object are encrypted, and the password hash is available for all encrypted recovery points.

> **Note**
> In case of a self-backup, the status is displayed if all or some recovery points of the backup object are encrypted and the product database contains details of encryption for all encrypted recovery points.

- **Not available**: All or some recovery points of the backup object are encrypted, and the password hash is not available for all encrypted recovery points.

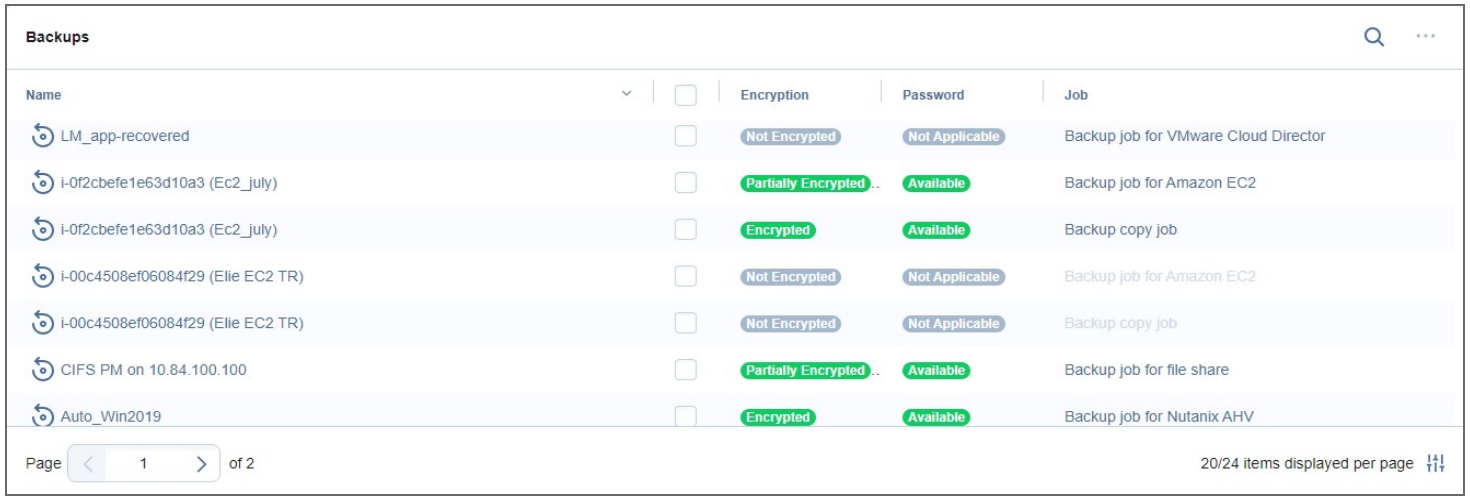
> **Note**
> In case of a self-backup, the status is displayed if all or some recovery points of the backup object are encrypted, and the product database does not contain the encryption details for all encrypted recovery points.

- **Partially available**: All or some recovery points of the backup object are encrypted, and the password hash is not available for some encrypted recovery points.

> **Note**
> In case of a self-backup, the status is displayed if all or some recovery points of the backup object are encrypted, and the product database does not contain the encryption details for some encrypted recovery points.

- **Not applicable**: The backup object does not contain encrypted recovery points.

- **Job**: The job type that created this backup
- **Size**: The total size of the backup



You can also initiate the recovery, verification, deletion, or password management process from here. Hover over the name of a backup and click the ellipsis **Manage** button on the right side to select one of the following options:

- **Recover**: Select this option to proceed with recovery.
- **Verify**: Select this option to verify the backup.
- **Provide password**: Select this option to manually provide the password if the currently selected backup is encrypted (some or all recovery points of the backup are encrypted) and the corresponding password hash is not available.

  Do the following:

  1. In the dialog box that appears, select the needed password or create a new one.

     Optionally, you can click the **Manage passwords** link to manage the existing or add a new password to the list of passwords.
  2. Click **Apply** to proceed.

  Notes

  - Optionally, you can select the checkboxes next to the backup names you want to manage, click the ellipsis **Manage** button, and then select **Provide password** in the drop-down list.
  - The **Provide password** option is disabled if the selected backups are not encrypted (none of the recovery points are encrypted).
  - The **Provide password** option is disabled if the selected backups are encrypted (some or all recovery points are encrypted) and the password hash is available for all encrypted recovery points.
  - It's recommended that you enable the **(AWS) Key Management Service** in the **Settings > General > System Settings > Encryption** tab. If **AWS** is enabled, all backup encryption passwords are encrypted with the Key Management Service cryptographic key to be available for recovery in case of product re-installation. For more information, refer to Enabling KMS.
- **Delete**: Select this option to delete the backup from the repository. Refer to "Deleting Backups" on page 614 for more details.

Clicking a backup name opens the **Backup Details** page where you can view the backup information and see all recovery points available for this backup.

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

- **Name**: The name of the backup item
- **Type**: The type of the backup
- **Tapes**: The number of tape cartridges that the backup occupies
- **Points**: The number of recovery points within the backup
- **Last point**: The date of the most recent recovery point of the backup
- **Size**: The total size of the backup

- **Job name**: The name of the job associated with the backup



Viewing Recovery Point Details

You can view the details of a recovery point in the lower part of the screen. To find a recovery point for a specific date, you can use the **Search** bar on the right. The following information is displayed:

- **Date**: The date when the recovery point was created
- **Size**: The size of the recovery point
- **Type**: Type of backup used to create the recovery point
- **Schedule**: If applicable, the schedule that was used to create the recovery point
- **Encryption**: The state of encryption of the backup object, which can be one of the following:
  - **Not encrypted**: None of the recovery points of the backup object are encrypted.
  - **Encrypted**: All the recovery points of the backup object are encrypted.
- **Password**: The state of the backup object password, which can be one of the following:
  - **Available**:The recovery point is encrypted and the corresponding password hash required for its decryption is available.

    Note

    In case of a self-backup, the status is displayed if the recovery point is encrypted and the product database contains encryption details for this recovery point.

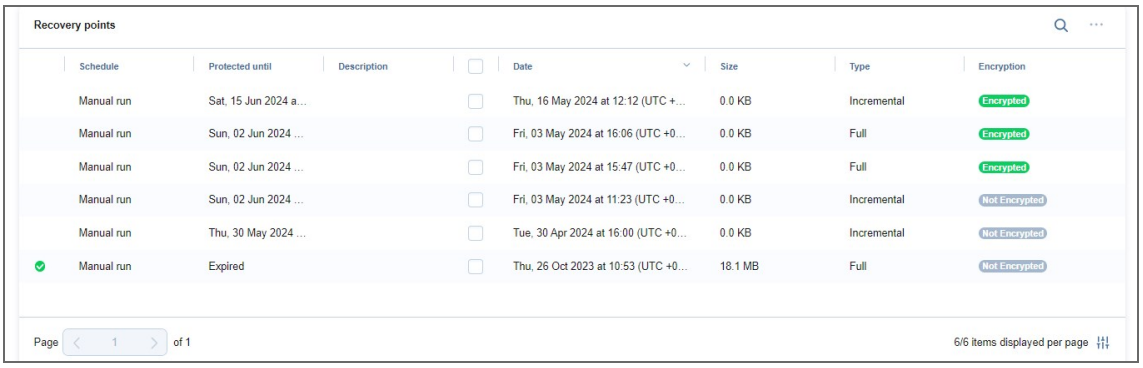
  - **Not available**: The recovery point is encrypted, and the corresponding password hash required for its decryption is not available.

Notes

- In case of a self-backup, the status is displayed if the recovery point is encrypted and the product database does not contain encryption details for this recovery point.
- In case of a self-backup, the status is displayed if you are upgrading your product from the version that does not support backup encryption to the version that supports it (for example, from 10.11 to 11.0).
- After a system configuration backup, the encrypted recovery points passwords that were created and saved in the password manager become unavailable and providing them manually is required.

    - **Not applicable**: The recovery point is not encrypted.

- **Immutable until**: If applicable, the date when the recovery point immutability expires
- **Protected until**: The date until which the recovery point is retained, displayed only for recovery points belonging to jobs that use the schedule retention approach
- **Description**: The description of the recovery point if one was provided



Notes

- The **Size**, **Type**, and **Immutable until** details are displayed only if the **Store backups in separate files** option (under **Storage Savings & Encryption**) is selected when creating or editing a Backup Repository.
- For recovery points belonging to jobs using legacy retention settings, **Use job retention** is displayed under **Protected until** instead.

Date, **Type**, and **Description** can also be viewed when selecting recovery points in the **Recovery Job Wizard**. Hover over the name of the recovery point and click the ellipsis **Manage** button on the right side to select one of the following options:

- **Recover**: Select this option to proceed with recovery.
- **Edit**: Select this option to edit the recovery point. Do the following:

- Optionally, you can add a **Description** to your recovery point.
- Choose the date until which the recovery point should be kept. The following options are available:
  - **Use job retention**: Choose this option to use the retention settings configured in the job for this recovery point.
  - **Keep forever**: Choose this option to keep this recovery point forever.
  - **Protect until**: Choose this option to keep this recovery point until a specific date. After selecting this option, choose the date in the calendar pop-up.

- **Provide password**: Select this option to manually provide the password if the currently selected recovery point is encrypted and the corresponding password hash is not available.

  Do the following:
  - In the dialog box that appears, select the needed password or create a new one.
  - Optionally, you can click the **Manage passwords** link to manage the existing or add a new password to the list of passwords.
  - Click **Apply** to proceed

  Notes
  - Optionally, you can select the checkboxes next to the recovery points you want to manage, click the ellipsis **Manage** button, and then select **Provide password** in the drop-down list.
  - The **Provide password** option is disabled if the selected recovery points are not encrypted (none of the recovery points are encrypted).
  - The **Provide password** option is disabled if the selected recovery points are encrypted (some or all recovery points are encrypted) and the password hash is available for all encrypted recovery points.
  - The **Provide password** option is disabled for self-backup recovery points created from an older version that does not support backup encryption.
  - It's recommended that you enable the **(AWS) Key Management Service** in the **Settings > General > System Settings > Encryption** tab. If **AWS** is enabled, all backup encryption passwords are encrypted with the Key Management Service cryptographic key to be available for recovery in case of product re-installation. For more information, refer to Enabling KMS.

- **Delete**: Select this option to delete the recovery point from the repository.

Note

To learn about the limitation on the chain of incremental recovery points, refer to the Knowledge Base article.

Federated Repositories

A **Federated repository** is a backup repository that consists of one or more existing backup repositories (members) with identical attributes. A **Federated repository** can be scaled horizontally to consume more backup repositories. Thus, backup jobs can continue to run even if one of the members is unavailable or has insufficient storage space. For detailed information, refer to Federated Repository.

This section covers federated repository-related topics such as creation, management, etc. of **Federated repositories** and contains the following articles:

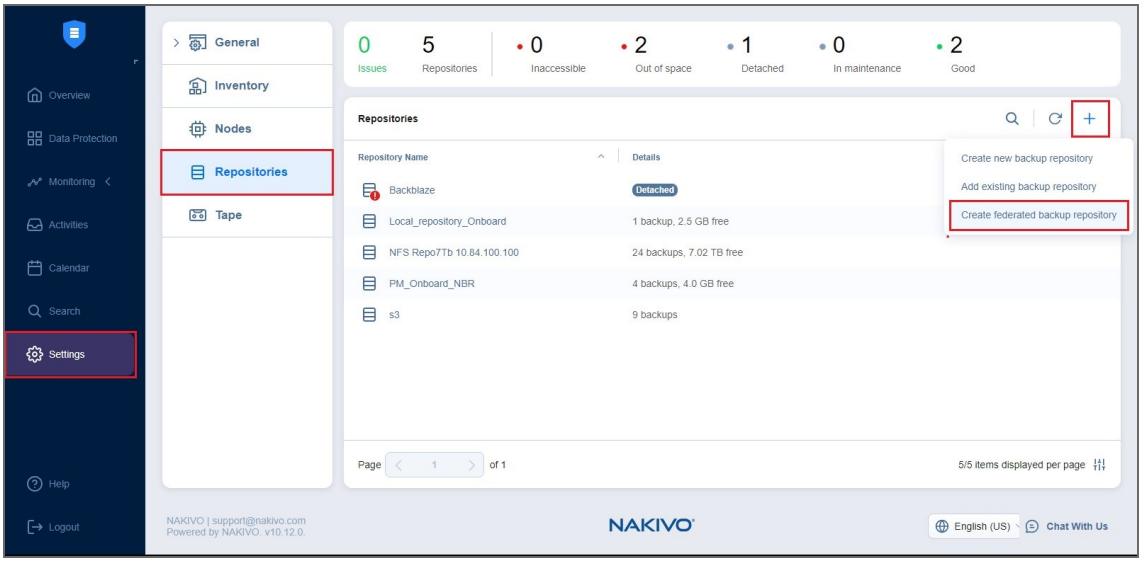
- Creating Federated Repositories
- Viewing Federated Repository Details
- Managing Federated Repositories

Creating Federated Repositories

NAKIVO Backup & Replication allows you to create federated repositories for storing backups.

To create a new federated repository, follow the steps below.

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



Important

- The **Federated repository** feature supports backing up data to **Incremental with full backups** repositories only.
- The **Create federated backup repository** option is disabled if no backup repositories of the supported type have been added to the product.

4. Proceed by completing the **Create Federated Backup Repository** wizard as described in the following sections:
    - Creating Federated Repository: Members
    - Creating Federated Repository: Options

Creating Federated Backup Repository: Members

On the **Members** page of the **Create Federated Backup Repository** wizard, use the table of available standalone repositories to choose and define one or more existing backup repositories (members) to be added to a federated repository.

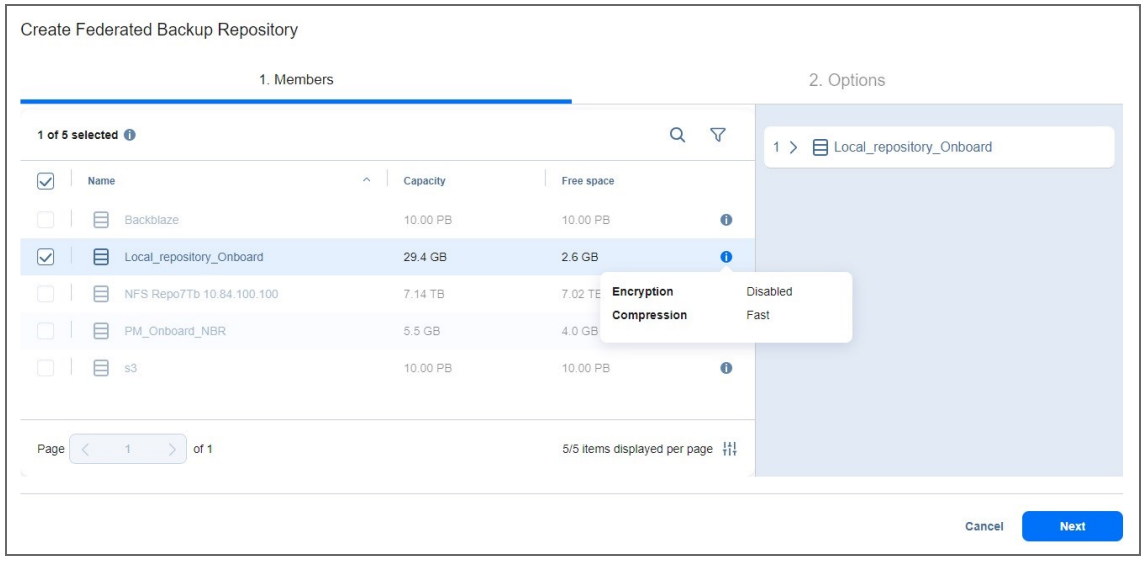
The table displays all available repositories with the following columns:

- **Repository name**: The name of the repository
- **Capacity**: The repository capacity in GB (gigabytes)
- **Free space**: The repository free space in GB (gigabytes)

Click the **Info** icon next to the repository to open a popup displaying the following information about the following supported federated repository member attributes:

Encryption (*Enabled* or *Disabled*)

Compression (*Disabled*, *Fast*, *Medium*, or *Best*)



Notes

- Only **Incremental with full backups** repositories are available for selection.
- Repositories that are already used as members of other federated repository(ies), repositories (and their members) containing a self-backup or selected as a self-backup storage, and those assigned to tenants are disabled.
- You cannot simultaneously select:
  - the members that support immutability and contain at least one immutable object or are associated with job(s) configured to create an immutable recovery point(s) and
  - the members that do not support immutability
- Selecting either type disables the member(s) of the other type.
- Federated repository members can use different transporters.

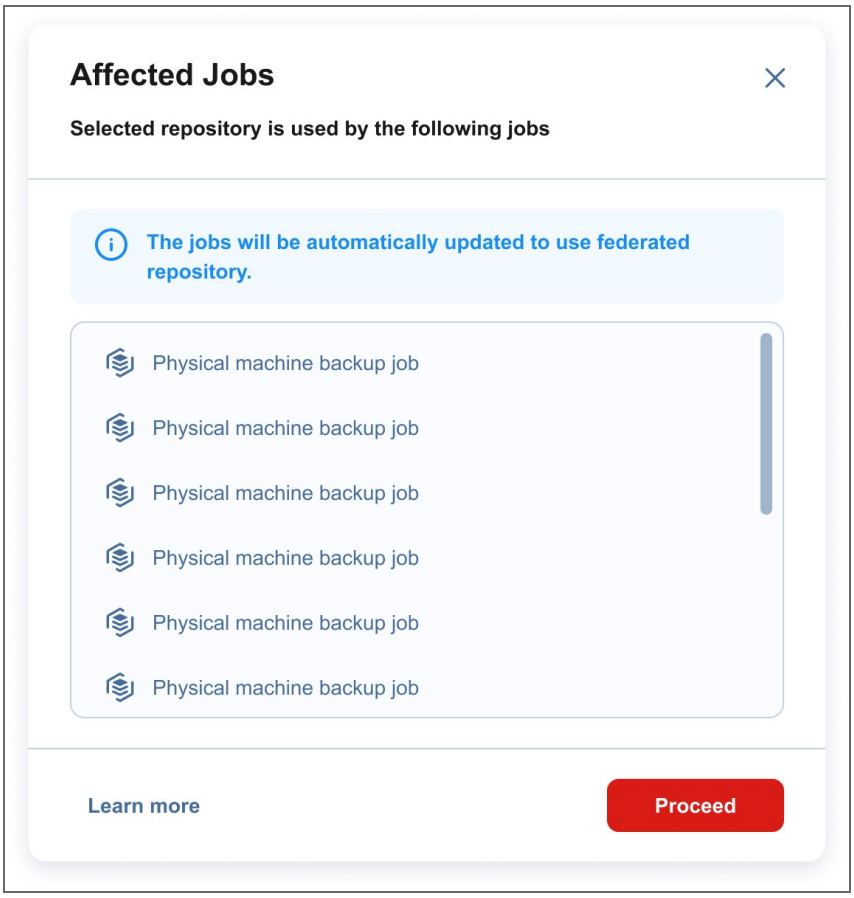
To choose and/or define one or more existing backup repositories (members) to be added to your federated repository, follow the steps below:

1. Select the needed repository item(s) by placing a checkmark to the left of the item(s).
2. Optionally, click the **Select All** checkbox to select all the enabled repositories listed.
3. The selected item(s) appear in the right pane of the page. You can remove a selected repository member in one of the following ways:
   a. Deselect the item in the left pane. This removes the repository member from the right pane; OR
   b. In the right pane, hover the pointer over the item you want to remove and click the "**X**" icon. This deselects the repository member in the left pane.

Notes

- If a federated repository is selected as the destination, the product automatically assigns an available federated repository member during the job run.
- Optionally, you can search members by clicking the **Search** button. This displays the search field where you can enter the partial or full name of the member.
- Click on the **Filter** button to filter the repository tree items by **Name** or **Type** (*Local folde*r, *NFS share*, or *CIFS share*).
- Federated repository members can be later added/removed to/from the federated backup repository or managed/refreshed individually.

4. If the selected member is used as a destination for existing jobs, the **Affected Jobs** dialogue is displayed with the list of jobs using the selected repository as a destination or source. Click **Proceed** to lock the selected repository before adding it as a new member and update the jobs after the federated repository is created.



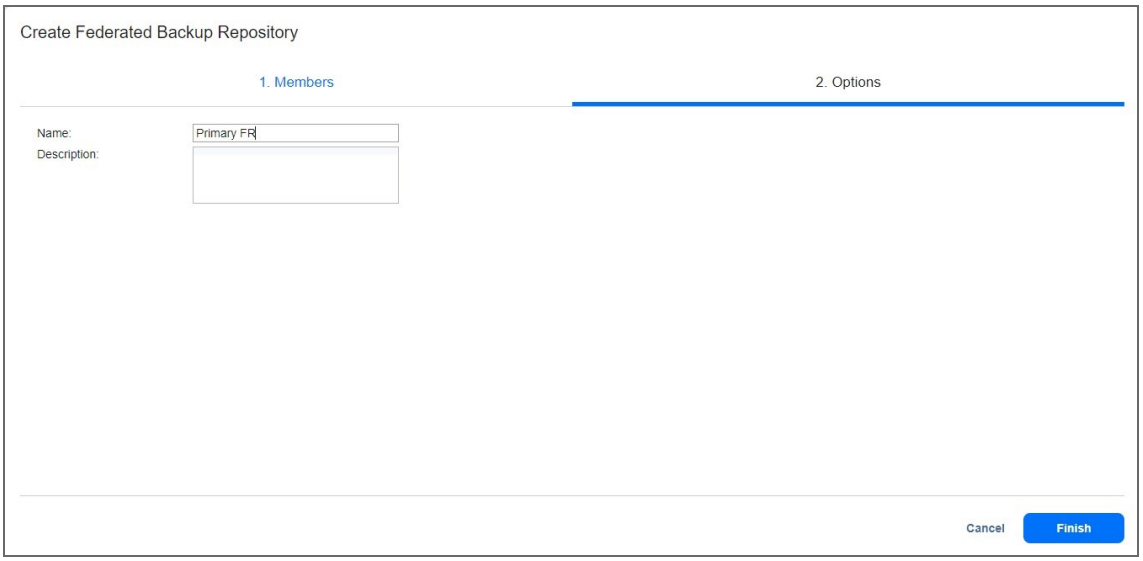
Notes

- After a backup repository is added as a member to a federated repository, it cannot be selected as a target for a new backup/backup copy job.
- After a federated repository is saved, the jobs using the repositories now selected to be part of the federated repository are automatically reconfigured to use this federated repository as a target.

5. Click **Next** to move to the next page of the wizard.

Creating Federated Backup Repository: Options

1. On the **Options** page, provide the name and description (optional) for the federated backup repository.
2. Click **Finish** to confirm saving the federated repository.



You can now view the federated repository by going to the **Settings > Repositories** tab.

Note

After a federated repository is saved, the jobs using the repositories now selected to be part of the federated repository are automatically reconfigured to use this federated repository as a target.

Managing Federated Repositories

For more detailed information about managing federated repositories, refer to the following topics:

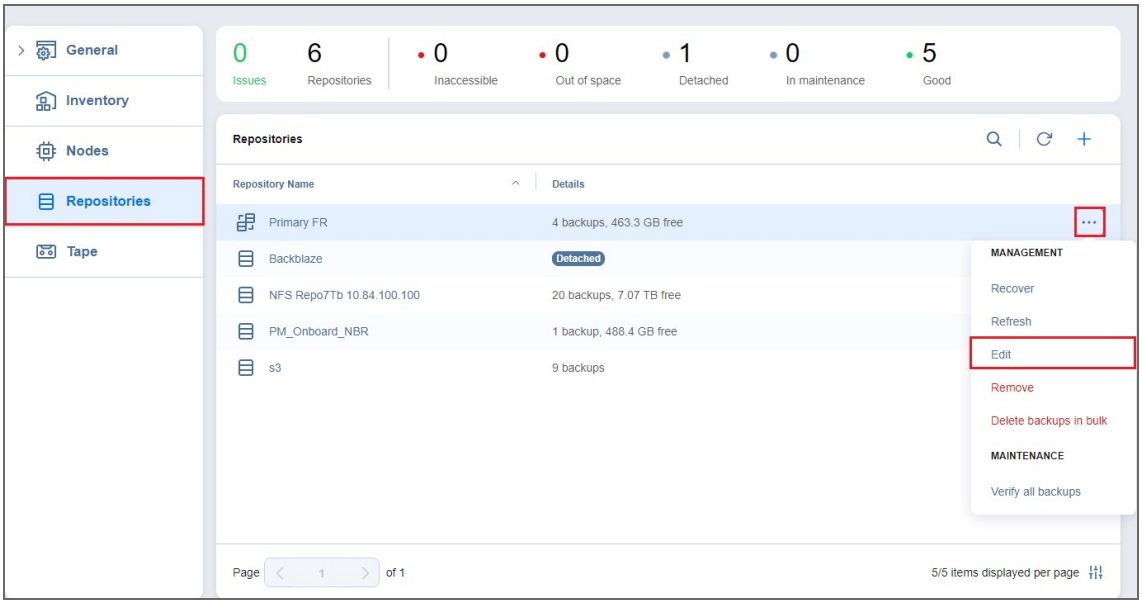
- Refreshing Federated Repositories
- Editing Federated Repositories
- Removing Federated Repositories
- Verifying All Backups
- Managing Federated Repository Members
- Migrating Backups between Federated Repository Members

Editing Federated Repositories

To edit the settings of an existing federated repository, follow the steps below:

In the main menu of NAKIVO Backup & Replication, click **Settings**.

1. Go to the **Repositories** tab.
2. Hover over a **Federated repository**.
3. On the right side, click **Manage**.
4. Click **Edit**.



Update the fields as necessary.

In the **Edit** mode, you can deselect federated repository members, remove them from the federated repository, or add new members.

5. Click **Apply**. Changes you have made are applied and the federated repository update starts.

Note

A federated repository cannot be edited while a backup job with this federated repository as target is running.

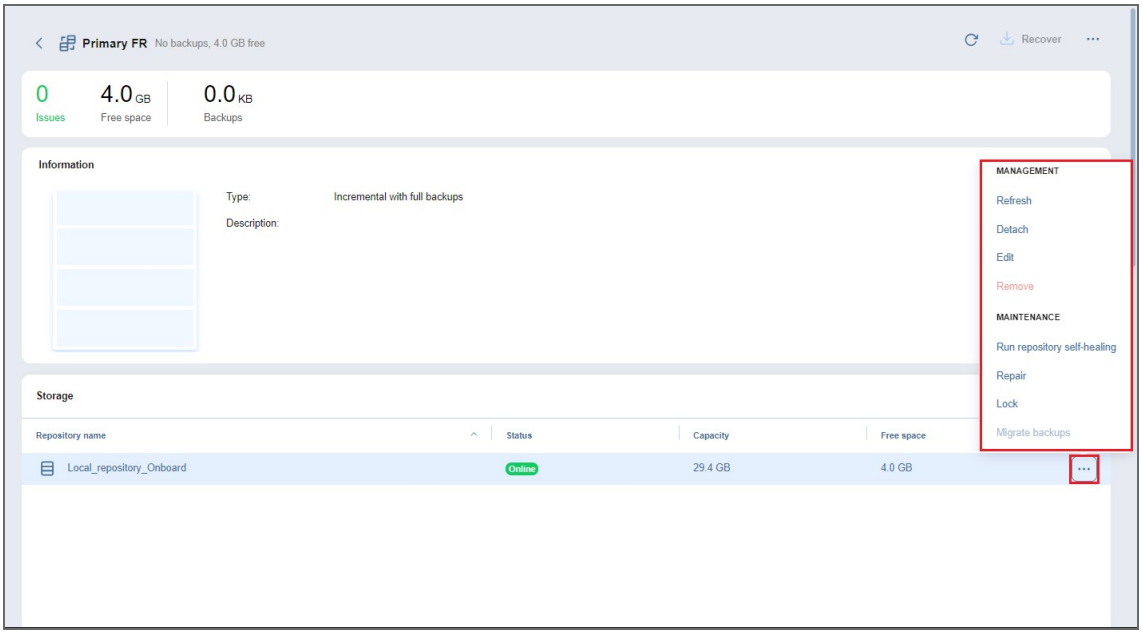
Managing Federated Repository Members

If needed, you can manage the standalone repositories used as members of federated repositories, add new members to, or remove them from the federated repository.

Managing Members in a Federated Repository

To manage/maintain members of a federated repository, follow the steps below:

1. In the main menu of NAKIVO Backup & Replication, click **Settings**.
2. Go to the **Repositories** tab.
3. Click a **Federated repository** to open the **Storage** table.
4. In the **Storage** table, hover over the needed federated repository member.
5. On the right side, click the ellipsis **Manage** button.



6. Select the needed option from the dropdown.

The following action buttons are available:

- **Refresh**: Opens the **Refresh repository** dialog box
- **Detach/Attach**: Opens the **Detach/Attach** dialog box

> **Note**
> If the federated repository member is currently detached, the **Attach** option is available. Select it to reattach this member.

- **Edit**: Brings up the option to edit the repository member
- **Remove**: Brings up the option to remove the repository member
- **Run repository self-healing**: Starts the repository self-healing process
- **Repair**: Opens the Repair Repository dialog box
- **Lock/Unlock**: Opens the Lock the member dialog box. Select it to suspend saving new recovery points to this repository member

> **Note**
> If the federated repository member is currently locked, the **Unlock** option is available. When you select **Unlock**, the **Unlock the member?** dialog box opens. Select it to resume saving new recovery points to this member.

- **Migrate Backups**: Brings up the option of migrating backups between federated repository members.

Notes

- Actions not applicable to a federated repository member are disabled.
- The option to migrate backups is disabled if the member is not locked or if the federated repository runs out of sufficient space for storing migrated data.

Adding Members to Federated Repository

To add members to a federated repository, follow the steps below:

1. In the main menu of NAKIVO Backup & Replication, click **Settings**.
2. Go to the **Repositories** tab.
3. Hover over a **Federated repository**.
4. On the right side, click **Manage**.
5. Click **Edit**.
6. Add a repository member from the list of available standalone repositories.

Notes

- After the changes to the federated repository are saved, the jobs using the repositories and now selected to be part of the federated repository are automatically reconfigured to use this federated repository as a target.
- Closing the federated member edit wizard opens the parent federated repository details screen.

7. Save the changes.

Deselecting Federated Repository Members

To deselect members in a federated repository, follow the steps below:

1. In the main menu of NAKIVO Backup & Replication, click **Settings**.
2. Go to the **Repositories** tab.
3. Hover over a **Federated repository**.
4. On the right side, click **Manage**.
5. Click **Edit**.
6. In the **Storage** table, deselect a federated repository member.

Note

A federated repository member cannot be deselected if it contains at least one backup object referenced by existing jobs.

Removing Members from Federated Repository

To remove members from a federated repository, follow the steps below:

1. In the main menu of NAKIVO Backup & Replication, click **Settings**.
2. Go to the **Repositories** tab.
3. Hover over a **Federated repository**.
4. On the right side, click **Manage**.
5. Click **Edit**.
6. In the **Storage** table, select a federated repository member.
7. Click **Manage**, then click **Remove**.

Notes

- The federated repository member can only be removed at the federated repository level.
- Removing a member that is not referenced by any existing jobs but contains recovery points from a backup object chain distributed across several members may break these backup chain(s). Consider migrating the backup objects before removing this member.

Migrating Backups between Federated Repository Members

Backups from one federated repository member can be moved to another available qualified federated repository member(s).

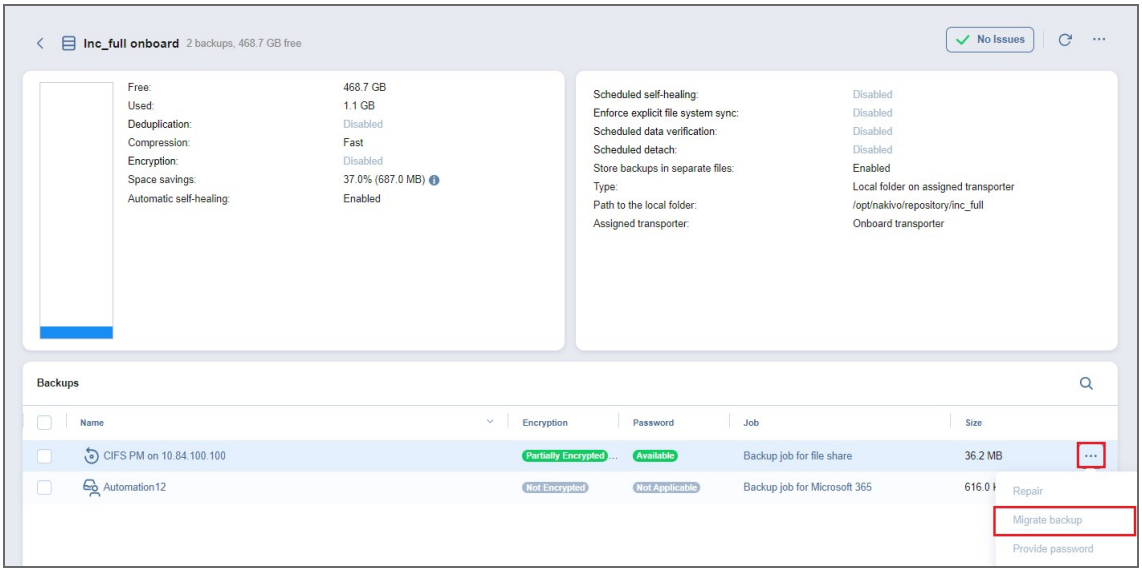
When considering backup migration, you can select one of two options:

- Migration of selected backups
- Migration of all backups

Migration of Selected Backups

To migrate a selected backup from one federated repository member to another available federated repository member(s), follow the steps below:

1. Go to the main menu of NAKIVO Backup & Replication and click **Settings**.
2. Go to the **Repositories** tab.
3. Select the **Federated repository** to open the Federated repository details dashboard.
4. In the **Storage** table, select the federated repository member name to open the **Backup Repository details** screen.
5. In the **Backups** table, hover over a backup name and click the ellipsis **Manage** button.
6. Click **Migrate backup**.



Notes

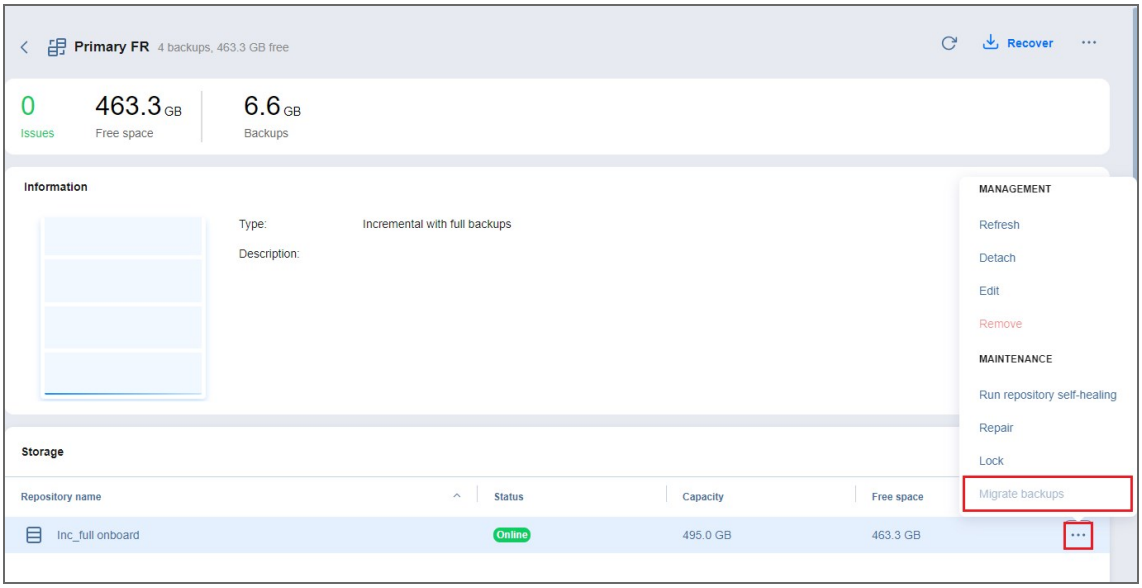
- The option of migrating backups is disabled if the member is not locked or if the federated repository runs out of sufficient space to store migrated data.
- The option is disabled if there are no other members with corresponding backup location policy.
- If a recovery point could not be cleaned up in the course of backup migration and it duplicates another recovery point, it is marked as inaccessible.
- Transporters assigned to both the source and the target federated repository members participate in the backup migration process.

Migration of All Backups

To migrate all the backups saved on one federated repository member to another federated repository member(s), follow the steps below:

1. Go to the main menu of NAKIVO Backup & Replication and click **Settings**.
2. Go to the **Repositories** tab.
3. Select the **Federated repository** to open the Federated repository details dashboard.
4. In the **Storage** table, hover over a federated repository member name.
5. On the right side, click the ellipsis **Manage** button and then click **Lock the member**. Now, saving new recovery points to this repository member is suspended.
6. Click **Migrate backups**.

The process of migrating backups between federated repository members begins.



Notes

- The progress based on the size of migrated backups can be monitored in the Action Bar.
- In the course of backup migration, the product redistributes backups stored on the selected federated repository member to available qualified federated repository members. Then the product removes the successfully redistributed backups from the selected source federated repository member. However, the immutable backups are not removed.
- The following functionality is enabled by default (if applicable) in the course of backup migration:
  - Network acceleration
  - Network encryption
- If the cleanup of some recovery points/backup objects fail (Clean up is skipped for immutable backups.), these remaining recovery points are not considered part of the backup objects associated with existing jobs.
- Duplicates of existing recovery points are marked as inaccessible, and the corresponding notification is displayed.
- If backup migration fails, the product re-tries the operation.
- If migration has been stopped and some backups have not been migrated yet, they remain in the source member.
- In case of a backup migration failure, information is logged and a dialog with the alarms list is displayed in the Activities dashboard.

Refreshing Federated Repositories

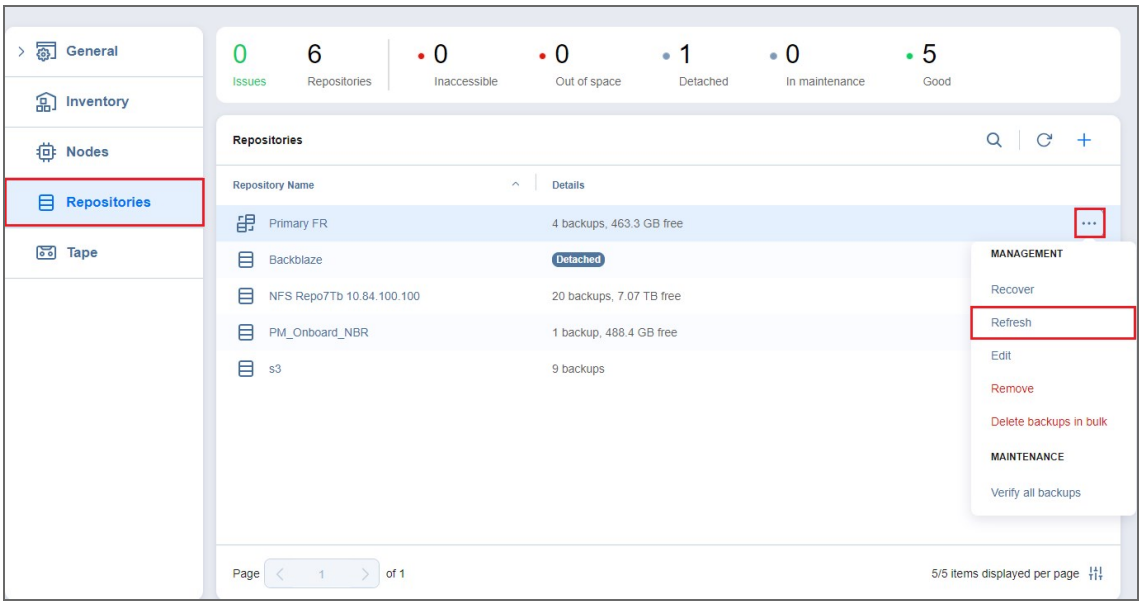
By default, NAKIVO Backup & Replication refreshes federated repository information hourly. During the refreshing process, the product collects all required information about federated repositories, such as the amount of available free space, number of backups, and number of recovery points.

Only one federated repository is refreshed at a time. Therefore, if you attempt to refresh multiple federated repositories, while one is being refreshed, all other federated repositories are queued.

To refresh a federated repository, follow the steps below:

In the main menu of NAKIVO Backup & Replication, click **Settings**.

1. Go to the **Repositories** tab.
2. Hover over a **Federated repository**.
3. On the right side, click **Manage**.
4. Click **Refresh**.



The federated repository refresh process begins.

Removing Federated Repositories

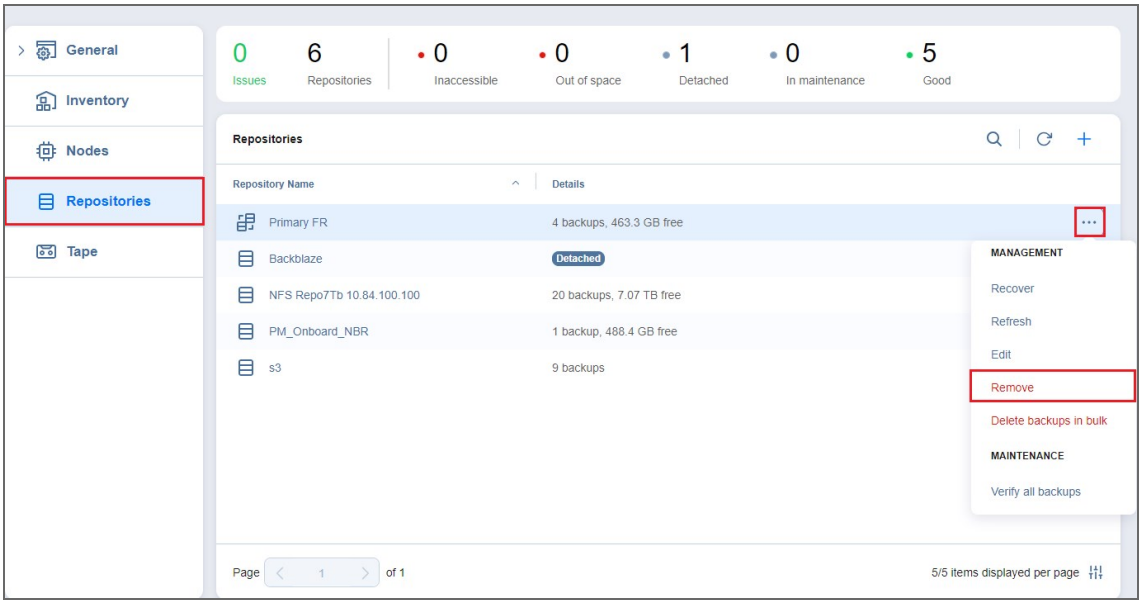
When you remove a federated repository, federated repository members and their contents are not removed and are treated as standalone backup repositories. However, the backup chains may be broken.

Note

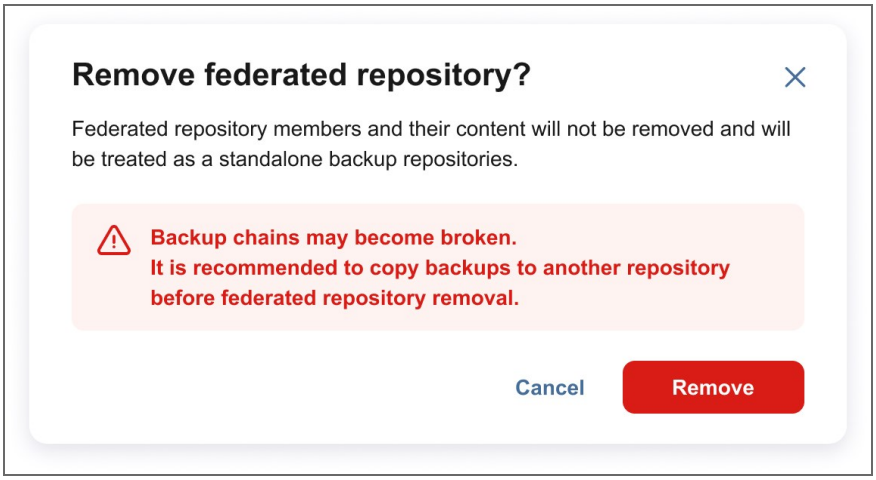
- A federated repository cannot be removed while a backup job using this federated repository as a target is running or if it contains at least one backup object referenced by existing job(s). To remove such a federated repository, delete (or edit) the corresponding job so that no items are backed up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

To remove a federated repository from the product, follow the steps below:

1. In the main menu of NAKIVO Backup & Replication, click **Settings**.
2. Go to the **Repositories** tab.
3. Hover over a **Federated repository**.
4. On the right side, click **Manage**.



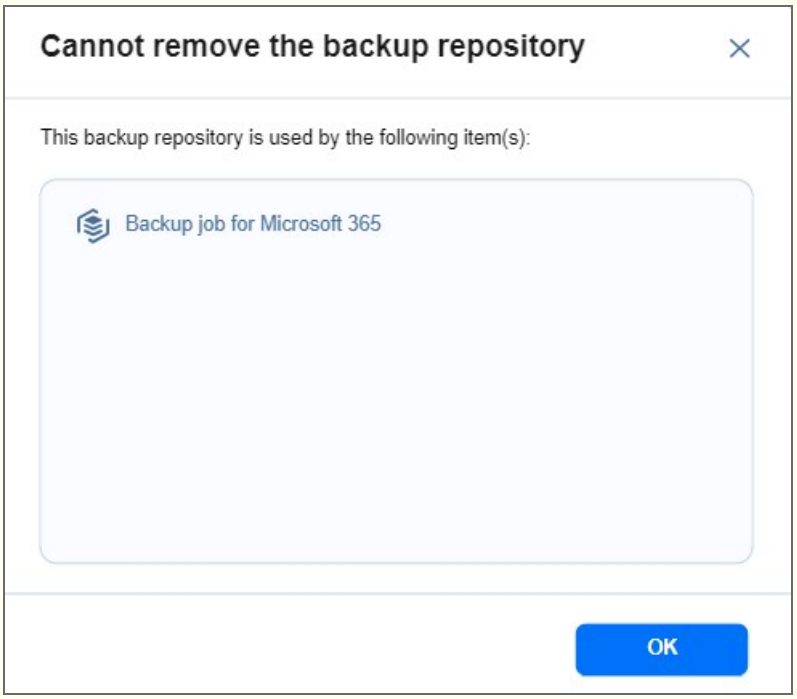
5. Click **Remove**.
6. Click **Remove** to confirm the action when the following confirmation message appears:



Make sure you copy the backups it contains to another repository.

Note

If a federated repository is used by one or more objects/processes or contains at least one backup object referenced by existing job(s), the following dialog box is displayed:



Important

You cannot remove the federated repository if it is used by one or more objects/processes.

Verifying All Backups

Backup verification is a process of checking the integrity of a backup by verifying its recovery points.

The **Federated repository** feature introduces enhanced backup integrity verification:

- Performs comprehensive check of all recovery points within the repository
- Validates metadata completeness
- Ensures data integrity by comparing with checksums

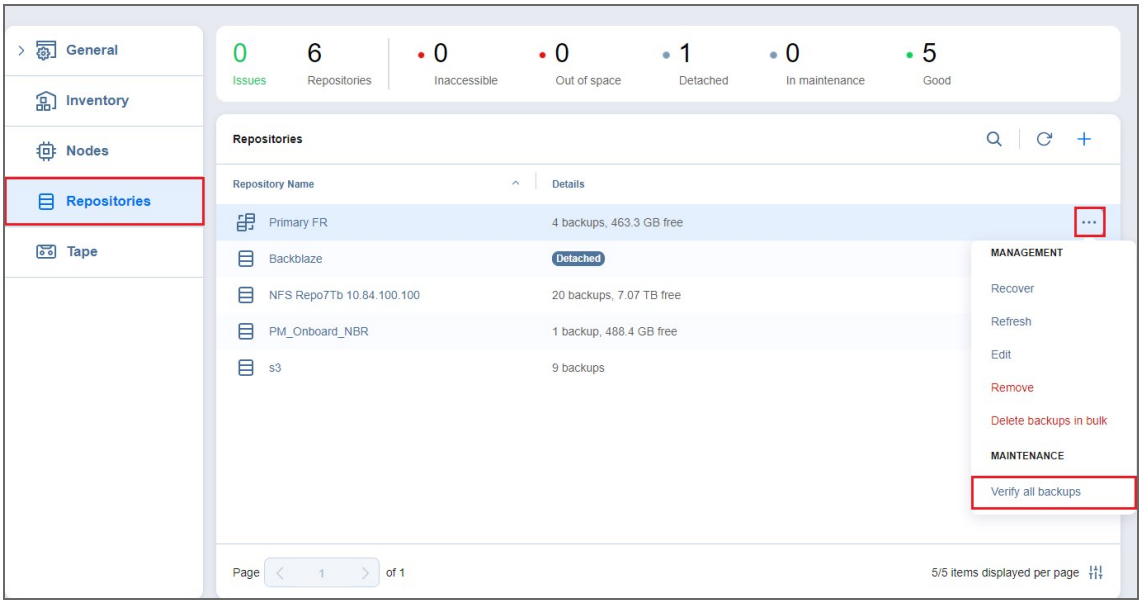
For specific recovery points, the following verification tasks will be done:

- Read timeline file:
  - Retrieves recovery point metadata from the timeline file
- Metadata verification:
  - Validates the presence and currency of metadata files specified in the timeline file
  - Compares timestamps for accuracy
- Incremental verification:
  - Searches for the parent FULL recovery point if the current recovery point is incremental
  - Validates the presence and correctness of parent recovery points, raw files, and metadata

- Raw file and metadata presence check:
  - Ensures the existence and correctness of raw files and their backup metadata (headers and records).

To start backup verification, follow the steps below:

1. In the main menu of NAKIVO Backup & Replication, click **Settings**.
2. Go to the **Repositories** tab.
3. Hover over a **Federated repository**.
4. On the right side, click the ellipsis **Manage** button and then click **Verify all backups**.



> **Note**
>
> The backup verification process cannot be started if a backup job using this Federated repository as a target is running.

5. To stop verification, hover over a federated repository, click the ellipsis **Manage** button and then click **Stop backup verification**.

Viewing Federated Repository

To see information about specific federated backup repositories, their members, backups, and recovery points, check out the sections below.

Viewing Federated Repository Details

Viewing Federated Repository Member Details

Viewing Backup Details

To manage federated repositories, in the **Repositories** table, hover over the name of a federated repository and click the ellipsis **Manage** button on the right side. For more detailed information, refer to "Managing Federated Repositories" on page 491.

Viewing Backup Details

To view the federated repository backup details, follow the steps below:

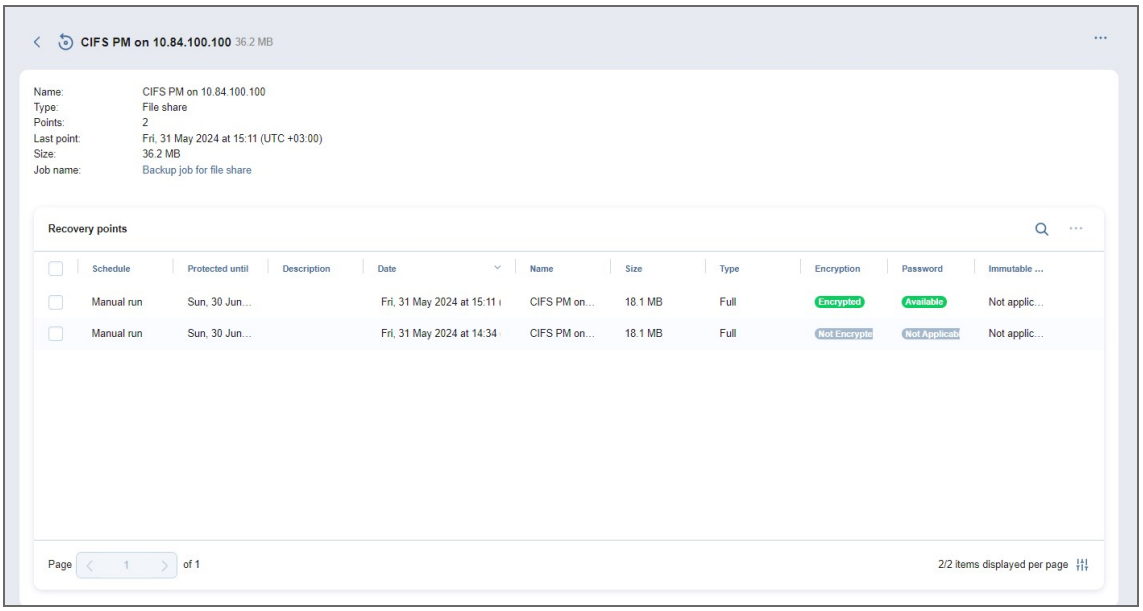
1. Go to the main menu of NAKIVO Backup & Replication and click **Settings**.
2. Go to the **Repositories** tab.
3. Select the **Federated repository**  to open the Federated repository details dashboard.
4. Scroll down to the **Backups** table and click the backup name to open the backup details screen in a new view.

The following information is displayed:

- **Name**: The name of the backup item
- **Type**: The type of job
- **Points**: The number of recovery points available
- **Last point**: The date of the latest recovery point
- **Size**: The total size of the backup
- **Job name**: The name of the job

Notes

- Backup details only contain the list of recovery points available in the selected backup object.
- Large numbers of backups are separated into pages to reduce clutter. To find a specific backup, you can scroll through the pages manually or simply look it up using the **Search** bar.



You can view the details of a recovery point in the lower part of the screen. To find a recovery point for a specific date, you can use the **Search** bar on the right. The following information is displayed:

- **Date**: The date when the recovery point was created
- **Size**: The size of the recovery point
- **Type**: Type of backup used to create the recovery point
- **Schedule**: If applicable, the schedule that was used to create the recovery point
- **Encryption**: The state of encryption of the backup object, which can be one of the following:
  - **Not encrypted**: None of the recovery points of the backup object are encrypted.
  - **Encrypted**: All the recovery points of the backup object are encrypted.
- **Password**: The state of the backup object password, which can be one of the following:
  - **Available**:The recovery point is encrypted and the corresponding password hash required for its decryption is available.
  - **Not available**: The recovery point is encrypted, and the corresponding password hash required for its decryption is not available.
  - **Not applicable**: The recovery point is not encrypted.
- **Immutable until**: If applicable, the date when the recovery point immutability expires
- **Protected until**: The date until which the recovery point is retained, displayed only for recovery points belonging to jobs that use the schedule retention approach
- **Description**: The description of the recovery point if one was provided

Notes

- For recovery points belonging to jobs using legacy retention settings, **Use job retention** is displayed under **Protected until** instead.
- **Date**, **Type**, and **Description** can also be viewed when selecting recovery points in the **Recovery Job Wizard**.

Optionally, hover over the name of the recovery point and click the ellipsis **Manage** button on the right side to select one of the following options:

- **Recover**: Select this option to proceed with recovery.
- **Edit**: Select this option to edit the recovery point. Do the following:
  - Optionally, you can add a **Description** to your recovery point.
  - Choose the date until which the recovery point should be kept. The following options are available:
    - **Use job retention**: Choose this option to use the retention settings configured in the job for this recovery point.
    - **Keep forever**: Choose this option to keep this recovery point forever.
    - **Protect until**: Choose this option to keep this recovery point until a specific date. After selecting this option, choose the date in the calendar pop-up.
- **Delete**: Select this option to delete the recovery point from the repository. Refer to "Deleting Recovery Points" on page 618 for more details.

Viewing Federated Repository Details

To view the federated repository details, follow the steps below:

1. Go to the main menu of NAKIVO Backup & Replication and click **Settings**.
2. Open the **Repositories** tab.

The **Repositories** tab contains a **Summary** bar, which offers an overview of all backup repositories both federated and standalone (not used as members of federated repositories). The data displayed is as follows:

- **Issues**: Total number of issues/alarms related to repositories
- **Repositories**: Total number of repositories
- **Inaccessible**: Total number of inaccessible repositories
- **Out of Space**: Total number of repositories that are out of storage space
- **Detached**: Total number of detached repositories
- **In Maintenance**: Total number of repositories in maintenance
- **Good**: Total number of usable repositories

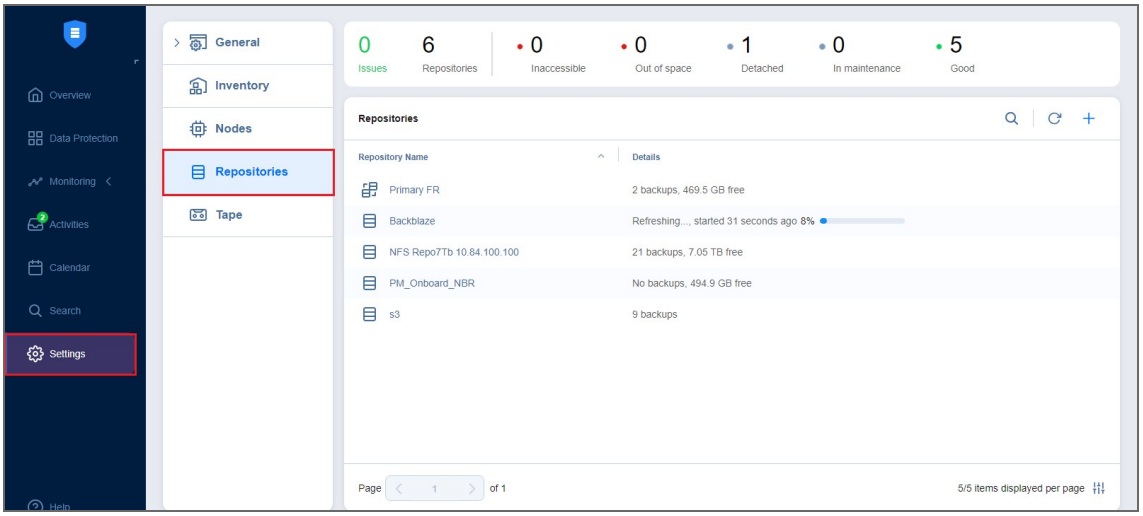
Note

In the **Repositories** table below, all repositories added to the product are listed according to the following priorities:

- Federated repositories
- Standalone repositories (those not used as members of federated repositories)

The data in the **Repositories** table is displayed as follows:

- **Repository name**: The name of the selected repository and its icon
- **Details**: The number of repository backups, amount of free space, and repository status (displayed for detached and inaccessible repositories only)



Notes

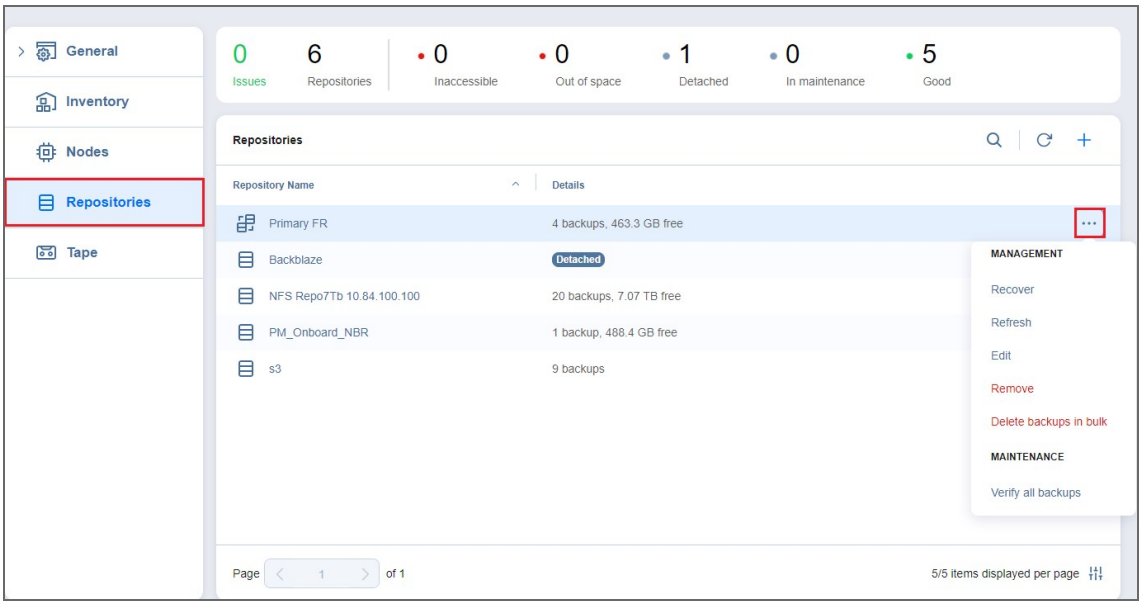
- If a repository status is displayed, it replaces the **Number of backups** and **Amount of free space data**.
- If a repository is both **Inaccessible** and **Detached**, the **Detached** status is displayed.
- If a repository is both **Out of space** and **Detached**, the **Detached** status is displayed.
- If some federated repository members are detached, other members are inaccessible, and there are no accessible attached members, the **Inaccessible** status is displayed for the federated repository.
- If any of the *in progress* statuses is in action (for example, *Self-healing*, *Verifying backups*), the progress bar is displayed.
- The progress bar is not displayed for **detached** and **inaccessible** repositories.

3. In the **Repositories** table, select the federated repository  to open the Federated repository details dashboard.

Notes

- Standalone repositories used as federated repository members are not displayed in the **Repositories** table. To view them, in the **Federated repository details** dashboard, proceed to the **Storage** table and click the federated repository member name to open the **Backup Repository details** screen.
- Optionally, you can search for standalone repositories by clicking the **Search** button. This displays the search field where you can enter the partial or full name of the repository.
- Click the **Refresh** button to refresh the list of items.
- Click the **+** button to create a new or add existing repository to your instance.
- To sort the table by **Repository name** click the head of the corresponding column.

Optionally, in the **Repositories** table, hover over the name of a federated repository and click the ellipsis **Manage** button on the right side to select one of the available options.



For more detailed information, refer to Managing Federated Repositories.

Federated Repository Details Dashboard

The **Federated repository details** dashboard offers a detailed overview of the selected federated repository details.

For a detailed explanation of each component in the **Federated repository details** dashboard, see the sections below.

- "Action Bar" on the next page
- "Summary Bar" on the next page
- "Information Pane" on the next page
- "Storage Table" on page 508

- “Backups Table” on the next page
- “Events Table” on page 509

Action Bar

The **Federated repository** action bar contains the generic icon, the title of the federated repository, total amount of free space in the federated repository, and the following three action buttons:

Refresh: Opens the Refresh repository dialog box

Recover: Brings up a list of recovery options for the selected federated repository

Manage: Brings up the options to edit, remove, delete backups in bulk, or verify all backups.

For more information, see Managing Federated Repositories.

Summary Bar

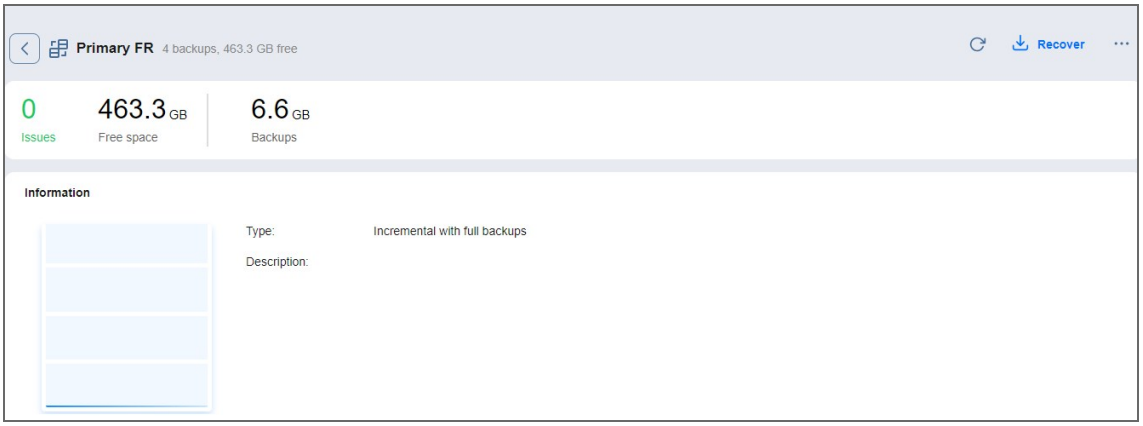
The **Federated backup repository** tab contains a **Summary** bar, which offers an overview of the selected federated repository. The data displayed is as follows:

- **Issues**: Total number of issues associated with the selected federated repository.
- **Free space**: Total amount of free space in the federated repository in GB (Gigabytes).
- **Backups**: Total amount of used space for backups in the federated repository in GB (Gigabytes).

Information Pane

This pane displays the following information about the selected repository:

- **Bar chart**: represents the amount of free/used space.
- **Type**
- **Description** (optional)



Important

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

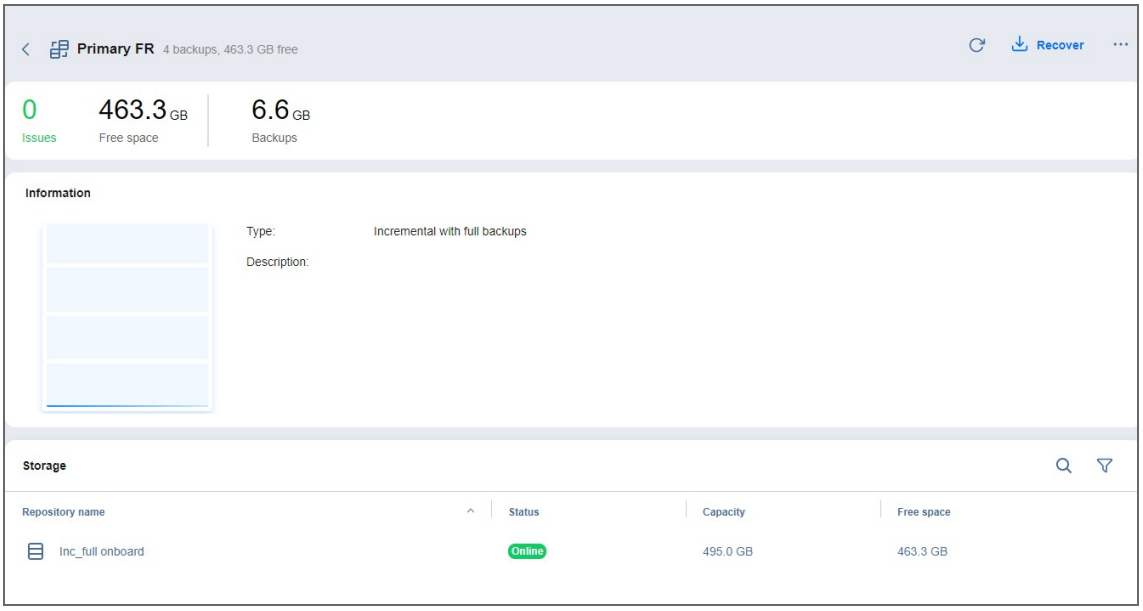
Storage Table

The **Storage** table displays a list of federated repository members and the information about each of them in the following columns:

- **Repository name**: The name of the repository.
- **Status**: The repository status, such as *Online* (displayed when the repository is accessible, attached, consistent and not under maintenance), *Detached*, *Locked* (replaced by *Self-healing…* or *Verifying backups…* if those are the cases), *Inaccessible*, or *Inconsistent*).
- **Capacity**: The repository capacity in GB (Gigabytes).
- **Free space**: The repository free space in GB (Gigabytes).

Notes

- Optionally, you can search for members by clicking the **Search** button. This displays the search field where you can enter the partial or full name of the member.
- Click on the **Filter** button to filter the storage tree items by **Name**, **Status** (*Online*, *Detached*, *Locked*, *Inaccessible*, or *Inconsistent*) .



To manage federated repository members, in the **Storage** table, hover over the name of a federated repository member and click the ellipsis **Manage** button on the right side. For more detailed information about the available options, refer to Managing Federated Repository Members.

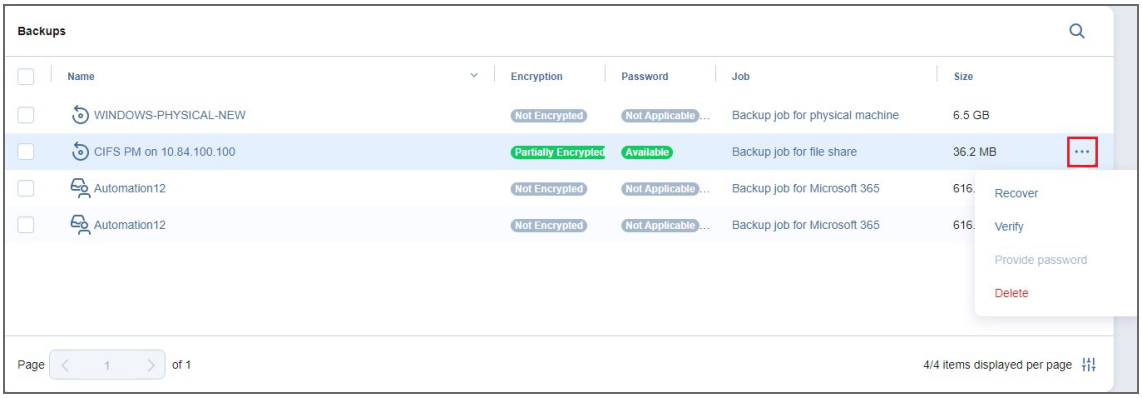
Click the federated repository member name to open the **Backup repository details** screen.

For more detailed information, refer to Viewing Federated Repository Details.

Backups Table

You can view the details of the backups stored in the federated repository. The following information is displayed in the **Backups** table:

- **Backup name**: Name of the backup
- **Encryption**: The state of encryption of the backup object, which can be one of the following:
  - **Not encrypted**: None of the recovery points of the backup object are encrypted.
  - **Encrypted**: All recovery points of the backup object are encrypted.
  - **Partially encrypted**: Some recovery points of the backup object are encrypted.
- **Password**: The state of the password of the backup object, which can be one of the following:
  - **Available**: All or some recovery points of the backup object are encrypted, and the password hash is available for all encrypted recovery points.
  - **Not available**: All or some recovery points of the backup object are encrypted, and the password hash is not available for all encrypted recovery points.
  - **Partially available**: All or some recovery points of the backup object are encrypted, and the password hash is not available for some encrypted recovery points.
  - **Not applicable**: The backup object does not contain encrypted recovery points.
- **Job**: The job type that created this backup
- **Size**: The total size of the backup



Notes

- You can use backups stored in a federated repository to create replicas.
- Recovery point management is only available on the federated repository level.

Click the backup name to open the backup details screen in a new view.

Note

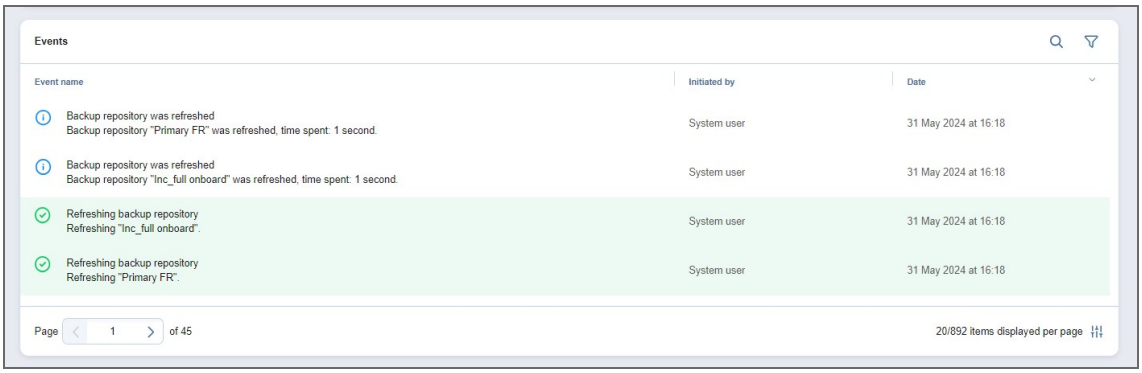
You can also migrate all the backup(s) from one federated repository member to the other available federated repository member(s). Refer to Migration of All Backups to learn more.

Events Table

The **Events** table displays a list of events associated with this federated repository in the following columns:

- **Event name**: The name of a given event in the group
- **Initiated by**: The initiator of the event

- **Date**: The date and time of the event



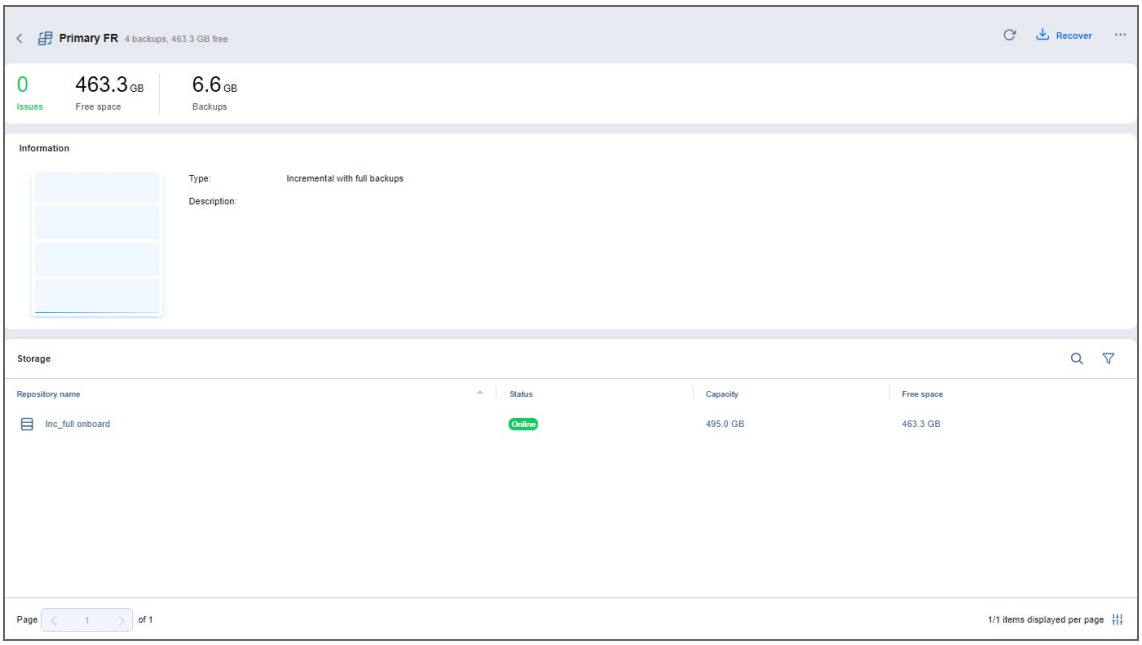
Viewing Federated Repository Member Details

Important

Standalone repositories used as federated repository members are not displayed in the **Repositories** table.

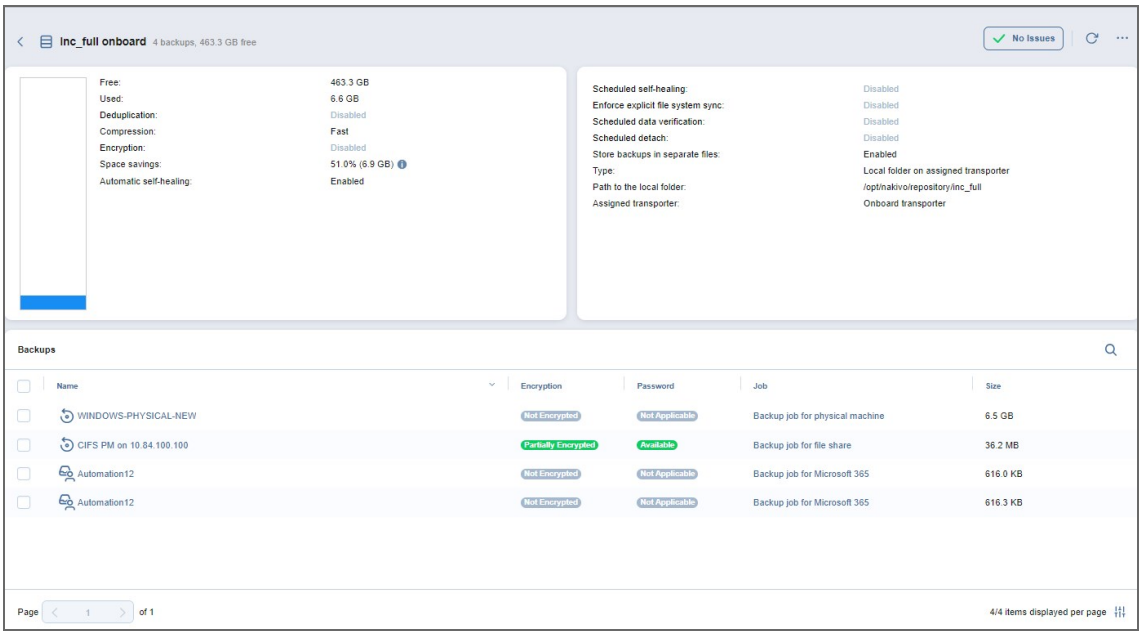
To view the federated repository member details, follow the steps below:

1. Go to the main menu of NAKIVO Backup & Replication and click **Settings**.
2. Go to the **Repositories** tab.
3. Select the **Federated repository** to open the **Federated repository details** dashboard.

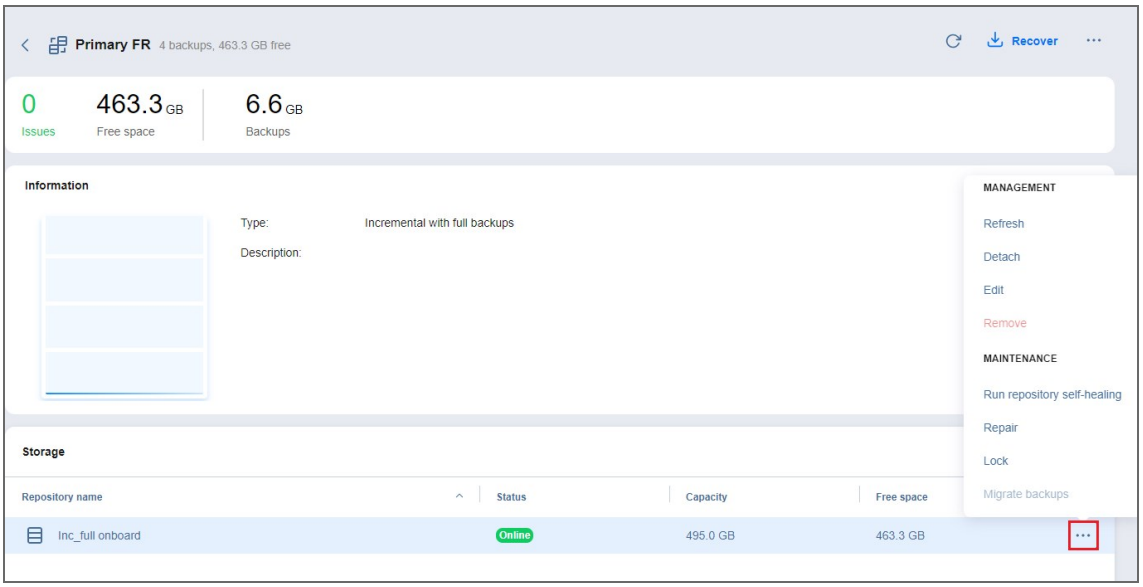


4. In the **Storage** table, click the federated repository member name to open the dashboard with the

federated repository member details.



To manage federated repository members, in the **Storage** table, hover over the name of a federated repository member and click the ellipsis **Manage** button on the right side.



For more detailed information, refer to Managing Members in a Federated Repository.

Notes

- Backup details only contain the list of recovery points available in the selected member.
- No actions are available at this level.

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.smtp.tls.version	Specifies the TLS version to use for SMTP server communication when TLS is configured in the Email Settings.	• Default • TLS10 • TLS11 • TLS12 • TLS13

system.email.notifications.skip.event.list	List of event names to skip when creating an email digest. Use space or "," or ";" as separators. The event names can be found in events.log.	Event names (example: error60)
system.vmware.esxi.ssh.port	For VMware only. Specifies the SSH port to connect to ESXi (global setting).	• Default value: 22 • 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.	• Unchecked (default) • Checked
system.vmware.skip.tag.discovery	VMware only. When enabled, the system does not discover VMware Tags. This is applied to all tenants.	• Unchecked (default) • Checked
system.debug.mode.log.vmware.api.incoming.requests	VMware only. When enabled, the incoming message will be printed for VIJAVA API received response. The option only works if **system.debug.mode.enabled** is checked.	• Unchecked (default) • Checked
system.debug.mode.log.vmware.api.outgoing.requests	VMware only. When enabled, the outgoing message will be printed for VIJAVA API sent request. The option only works if **system.debug.mode.enabled** is checked.	• Unchecked (default) • Checked

http.max.upload.size	Specifies the max upload size for file upload operations, **bytes** (global setting). If multiple files are uploaded, this is the total size. Use -1 for unlimited. Example: 200MB: 200000000	• Default value: 10737 41824 • Minimum value: 1 • Maximum value: 99999 9999999
system.auth.use.lockout	Enables or disables the login lockout feature. When enabled, the offending IP address is not allowed to login after several failed attempts.	• Unchecked (default) • Checked
system.auth.max.login.attempt.count	Specifies the maximum number of failed login attempts to trigger the login lockout feature for the offending IP.	• Default value: 5 • 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.auth.ad.integration.connect.timeout	Specifies the timeout (miliseconds) for connecting LDAP/Active Directory.	• Default value: 2000 • Minimum value: 2000 • Maximum value: 100000
system.auth.ad.integration.read.timeout	Specifies the timeout (miliseconds) for reading LDAP/Active Directory operations.	• Default value: 10000 • Minimum value: 10000 • Maximum value: 100000
system.auth.max.login.2fa.attempt.count	Specifies the attempts for the login lockout feature. The offending IP is allowed to login again after the timeout expires.	• Default value: 5 • Minimum value: 1 • Maximum value: 9999

system.auth.lockout.2fa.timeout	Specifies the timeout (minutes) for the login lockout feature. The offending IP is allowed to login again after the timeout expires.	• Default value: 5 • Minimum value: 1 • Maximum value: 9999
system.job.block.size	Select block size for processing data. **Notes** • Deduplication can only be efficient with recovery points using the same block size. • Once the value is changed, the existing backup jobs, previously using a different block size, will produce a full backup on the next run. • Mapping to a backup with a different block size will be skipped	• 4 MB (default) • 2 MB • 1 MB • 512 KB • 256 KB • 128 KB • 64 KB • 32 KB • 16 KB • 8 KB • 4 KB
system.job.map.new.source.item.scope	The scope to search for the existing backup when adding a new source item to the job.	• Default location (default) • Default transporter's locations • All locations
system.job.pool.queue.length	Specifies the length of the job queue. A job is placed in a queue before execution. Requires restart.	• Default value: 200 • 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).	• Unchecked (default) • Checked

system.job.resolve.transporter.hostname.on.transporter	If set, sends the source and/or target Transporter hostname as is to Transporter. Transporter will resolve the hostname to the IP address(es) and check if it is reachable. This is done during Transporter to Transporter checks on job run. The default behavior is to do the resolution locally, get hostnames for all resolved IP addresses, and then send them to Transporter. This can be a problem in complex network topologies (VPN, etc).	• Unchecked (default) • Checked
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.job.ict.skip.new.disk	If set, new disks added to the source item will not be added to the job automatically.	• Checked (default) • Unchecked
system.job.replica.vm.suffix	The default suffix to append to replica VMs. This setting is global and can only be changed inside the master tenant.	Can be between 1 and 20 characters ("-replica" by default)

system.job.recovered.vm.suffix	The default suffix to append to recovered/flash-booted VMs. This setting is global and can only be changed inside the master tenant.	Can be between 1 and 20 characters ("-recovered" by default)
system.job.skip.manual.transporter.data.path.validation	If set, transporter data path validation will be skipped for manually configured transporters.	• Unchecked (default) • Checked
system.metadata.disable.ec2.instance.id.update	Disables EC2 instance ID detection on product startup. The detection is performed through an HTTP request to the AWS instance metadata service. This process is required for proper product functioning in the AWS cloud.	• Unchecked (default) • Checked
system.task.pool.queue.length	Specifies the length of the task queue. A task is placed in the queue before execution. Requires restart.	• Default value: 200 • 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.ec2.min.free.space.resize.percent	In case the free space is less than the set percentage of the total current storage, one minimum chunk will be added to the storage.	• Default value: 10 • Minimum value: 1 • Maximum value: 100
system.repository.ec2.max.free.space.resize.percent	In case the free space is more than the set percentage of the total current storage, one minimum chunk will be removed from the storage.	• Default value: 15 • Minimum value: 1 • Maximum value: 100
system.repository.maintenance.interrupt.timeout.seconds	Specifies the timeout (seconds) to wait for repository maintenance stop during job run.	• Default value: 300 • 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.repository.refresh.timeout.seconds	Specifies the timeout (seconds) to wait for repository refresh.	• Default value: 600 • Minimum value: 1 • Maximum value: 86400
system.repository.remove.backups.used.by.job	The setting allows to remove backup objects associated with existing jobs, and remove the last RP of a backup object in case such RP is due to be removed according to the retention policy. If enabled, removal of the aforesaid objects can be done manually or automatically, in accordance with the configured retention policy.	• Unchecked (default) • Checked
system.product.skip.update.server.ssl.certificate.verification	The product update check process requires the remote server certificate to be trusted. This parameter disables such check. It can be useful when secure (SSL/TLS) connections are being intercepted by third-party software.A product restart is required to apply.	• Unchecked (default) • Checked

system.debug.mode.enabled	The debug mode prints more information into the logs, including some sensitive one (hardware UUIDs, MAC addresses, etc). The passwords are not printed unless they are present in raw communication dumps (e.g., SOAP/XML/JSON).	• Unchecked (default) • Checked
system.debug.mode.log.passwords	When debug mode is enabled, also log passwords. This can be a security risk.	• Unchecked (default) • Checked
system.debug.mode.log.api.requests	When debug mode is enabled, also log product API requests/responses. The data is logged as is and will contain plaintext passwords. This can be a security risk.	• Unchecked (default) • Checked
system.hyperv.optimize.queries	Hyper-V only. Instructs to use a faster query method to read VM and host information. This will speed up the refresh process in large environments.	• Checked (default) • Unchecked
system.hyperv.discovery.host.thread.count	Hyper-V only. Sets the max parallel threads to run when refreshing cluster hosts during discovery. Each cluster host can be refreshed separately. This will speed up the refresh process in large environments.	• Default value: 2 • 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.	
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.aws.discovery.region.thread.count	AWS only. Sets the max parallel threads to run when refreshing the AWS Regions during discovery. When increasing the setting value, make sure to test its influence on host CPU usage during refresh. This will speed up the refresh process in large environments.	• Default value: 4 • Minimum value: 1 • Maximum value: 10
system.aws.discovery.other.thread.count	AWS only. Sets the max parallel threads to run when refreshing other AWS entities inside the Region during discovery. When increasing the setting value, make sure to test its influence on host CPU usage during refresh. This will speed up the refresh process in large environments.	• Default value: 4 • 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 5 hours.	• Default value: 5 • 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: 4294967296 • Minimum value: 0 • Maximum value: 10995 11627776

system.transporter.modern.thread.pool.size	Specifies the session factory thread pool size for modern Transporter. **Note**: Transporter restart is required to apply the setting.	• Default value: 1000 • Minimum value: 100 • Maximum value: 1000
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: 1.7976931348623157e+308
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.o365.suppress.throttling.event	Suppress throttling warning.	• Unchecked (default) • Checked
system.event.skip.creating.event.list	List of event/alarm/notification names to skip when creating an event. The event is still logged and handled. Use space or , or ; as separators. The names can be found in events.log.	Event names (example: error60)

system.events.use.windows.event.integration	Use Windows Event log integration. Some product events will also be created in the Application log. This setting is global and can only be changed inside the master tenant.	• 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.olr.dsamain.mount.port	TCP port where DSAMAIN mounts NTDS.dit (AD database) for.	• Default value: 5000 • Minimum value: 1 • Maximum value: 65535
system.product.register.disable.periodic.data.collection	When enabled, the product will not send data bundles every 30 days.	• 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.pql.custom.file.name	PQL file name in the userdata folder. Empty by default. If empty, the file will be downloaded from web. Example: pql.json.	• Default value: empty • Minimum value: 1 • Maximum value: 2048
system.pql.cache.ttl.hours	Time to keep PQL file cache, in hours. Use 0 to disable the cache.	• Default value: 8 • Minimum value: 0 • Maximum value: 72
system.transporter.allow.new	Allows using newer Transporter versions.	• Unchecked (default) • Checked
system.transporter.allow.old	Allows using older (outdated) Transporter versions.	• Unchecked (default) • Checked
system.transporter.modern.idle.timeout	Specifies the timeout (milliseconds) for modern Transporter IDLE. If you set it to 0, it will be an unlimited timeout, meaning the transporter can only be stopped manually. **Note**: Transporter restart is required to apply the setting.	• Default value: 3600000 • Minimum value: 0 • Maximum value: 86400000
system.volatile.object.processing.type	**Default**: try to remove the volatile objects periodically until their time to live (fixed) is reached. **Alternative**: fine-tune the settings. See the other system.volatile.object variables. The setting is global and can be changed inside the master tenant only.	• Default (default option) • Alternative

system.volatile.object.retry.count	Alternative processing type only. The maximum number of retries for volatile objects removal. 0 means no retries, so only one removal attempt will happen. The setting is global and can be changed inside the master tenant only.	• Default value: 7 • Minimum value: 0 • Maximum value: 256
system.volatile.object.retry.interval	Alternative processing type only. Minutes. The desired delay between each removal retry. The real delay depends on the queue and on the exponential retry factor (configurable). The setting is global and can be changed inside the master tenant only.	• Default value: 60 • Minimum value: 5 • Maximum value: 14400
system.volatile.object.exponential.retry.interval.factor	Alternative processing type only. The ratio to use when calculating the delay time for the next retry. The next delay equals interval * (factor^retry). Example: the interval is 60 minutes, the factor is 2. The first retry will happen in +60 minutes, the second in +240 minutes , ... The setting is global and can be changed inside the master tenant only.	• Default value: 2 • Minimum value: 1 • Maximum value: 10
system.visual.notification.service.disable	Disables the visual notification service. This can speed up the UI when the database contains many event entries. This setting is global and can be changed inside the master tenant only.	• Unchecked (default) • Checked
system.msp.console.listening.port	TCP port used by the MSP product for listening to remote tenants. Port 6702 is used by default.	• Default: 6702 • Minimum value: 1 • Maximum value: 65535

system.events.use.windows.event.integration	Use Windows Event log integration. Some product events will also be created in the Application log. The setting is global and can be changed inside the master tenant only.	• Unchecked (default) • Checked
system.transporter.load.path.cost.variation.percent	Percent. Specifies the allowed data path cost variation. During the job run, automatic transporter selection may happen. The first step is to choose the top N (by cost in milliseconds) data paths. The second step is to choose the best data path based on the lowest transporter load. For example, 20 paths were found based on round trip time between source and target host. If the setting is 10%, the best path cost is 2 (N=2), then only paths with costs in range 2 ... 2.2 (2+0.2) will be chosen on the first step. The setting is global and can be changed inside the master tenant only.	• Default value: 10 • Minimum value: 1 • Maximum value: 10000
system.vmware.discovery.vm.detect.ipaddress.by.dns.skip	The setting is applicable only to VMware vSphere infrastructure. If enabled, the detection of VM IP address via DNS resolution will be skipped. Note: Detection of VM IP address via DNS resolution is applied in case VMware Tools are not installed on the VM.	• Checked (default) • Unchecked

system.job.default.retention.approach	Schedule-retention fusion: New backup and backup copy jobs will use the new schedule-retention step. Legacy: New backup and backup copy jobs will use the legacy schedule and retention steps.	• Schedule-retention fusion • Legacy
system.job.run.skip.infrastructure.refresh	If set, infrastructure refresh will be skipped during job run.	• Unchecked (default) • Checked
system.job.ict.skip.new.source.item	If set, new source items will not be added to the job automatically.	• Unchecked (default) • Checked
system.transporter.min.connect.timeout.seconds	Specifies the minimum connect timeout (seconds) for transporter connection. Set 0 to use default value in the product (90 seconds).	• Default value: 0 • Minimum value: 0 • Maximum value: 86400
system.transporter.min.write.timeout.seconds	Specifies the minimum write timeout (seconds) for transporter command. Set 0 to use default value in the product (60 seconds).	• Default value: 0 • Minimum value: 0 • Maximum value: 86400
system.transporter.min.idle.timeout.seconds	Specifies the minimum read idle timeout (seconds) for transporter command. Set 0 to use default value in the product (300 seconds).	• Default value: 0 • Minimum value: 0 • Maximum value: 86400

system.events.use.syslog.integration	Use Syslog log integration. The product events will be copied to the remote log. The setting is global and can be changed inside the master tenant only	• Unchecked (default) • Checked
system.events.syslog.host	Specifies the Syslog server hostname. The setting is global and can be changed inside the master tenant only.	• IP/hostname field (empty by default)
system.events.syslog.port	Specifies the Syslog server port. syslog-udp usually uses port 514, syslog-tls usually uses port 6514. The setting is global and can be changed inside the master tenant only.	• Default value: 514 • Minimum value: 1 • Maximum value: 65535
system.events.syslog.transport	Specifies the transport for sending Syslog messages. The setting is global and can be changed inside the master tenant only.	• UDP (default) • TCP • TLS
system.events.syslog.connect.timeout	Milliseconds. Specifies the connect timeout when connecting to the Syslog server, TCP/TLS case. The setting is global and can be changed inside the master tenant only.	• Default value: 1000 • Minimum value: 10 • Maximum value: 180000

system.events.syslog.message.facility	Specifies the facility to be used in Syslog messages. Facility code Keyword Description 0 kern Kernel messages 1 user User-level messages 2 mail Mail system 3 daemon System daemons 4 auth Security/authentication messages 5 syslog Messages generated internally by syslogd 6 lpr Line printer subsystem 7 news Network news subsystem 8 uucp UUCP subsystem 9 cron Cron subsystem 10 authpriv Security/authentication messages 11 ftp FTP daemon 12 ntp NTP subsystem 13 security Log audit 14 console Log alert 15 solaris-cron Scheduling daemon 16–23 local0 – local7 Locally used facilities The setting is global and can be changed inside the master tenant only.	• Default value: 1 • Minimum value: 0 • Maximum value: 23
system.events.syslog.message.app.name	Specifies a custom source application name to be used in Syslog messages. Leave blank for default. The setting is global and can be changed inside the master tenant only.	• Empty (default)

system.events.syslog.message.hostname	Specifies a custom source host name to be used in Syslog messages. Some Syslog cloud services may use this field to transmit a secret key. Leave blank for default. The setting is global and can be changed inside the master tenant only.	• Empty (default)
system.events.syslog.message.format	Specifies the message format to be used in Syslog messages. RFC 5424 has more fields and is used by rsyslog by default. The setting is global and can be changed inside the master tenant only.	• RFC_3164 • RFC_5424 (default) • RFC_5424
system.events.syslog.message.style	Specifies the message style to be used for Syslog messages. The setting is global and can be changed inside the master tenant only.	• TEXT • JSON (default)
system.inventory.o365.skip.throttled.items	Microsoft 365 only. When enabled, the system will skip throttled items.	• Unchecked (default) • Checked
system.job.run.test.treat.immutable.days.as.minutes	Test only. Treat 'Immutable for X days' job schedule option as 'Immutable for X minutes'. The setting is global and can be changed inside the master tenant only.	• Unchecked (default) • Checked
system.job.run.test.treat.keep.backup.days.as.minutes	Test only. Treat 'Keep backups for X days' job schedule option as 'Keep backups for X minutes'. The setting is global and can be changed inside the master tenant only.	• Unchecked (default) • Checked
system.debug.config	Predefined debug logging configuration. Default values will be used if input invalid JSON. Set to empty to reset default values. Should be edited by Dev only to debug specific feature.	• Format value : JSON

ui.hidden.features	The list of features to hide in the UI. Reload the UI to apply. Available options: • Repositories - S3 • Repositories - SaaS • Repositories - deduplication device • Repositories - federated • Backup job for AWS EC2 • Backup job for M365 • Backup job for Oracle Database • Backup job for Vmware vCloud Director • RtR job for VMware • Replication job for AWS EC2 • Nodes - AWS EC2 • Nodes - VMware RtR • Inventory - SaaS • Inventory - Application • Inventory - Cloud • Inventory - Storage • Inventory - VMware vCloud Director • Configuration - System - Encryption • Configuration - System - Aptare • Configuration - MSP	• Unchecked (default) • Checked

system.repository.remove.backups.recoverypoints	The setting allows to remove backup objects associated with existing jobs, and remove all recovery points including the last recovery point of a backup object in case such recovery points are due to be removed according to the retention policy. If enabled, removal of the aforesaid objects can be done manually or automatically, in accordance with the configured retention policy.	• Unchecked (default) • 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.
- **Clean up job history**: By clicking the link, you can immediately apply the configured **Store job history for the last** setting.

In the text box, you can see the report on the actions.

Example 1

Request 1: sending (Remove all events)...

Request 1: success=true (Remove all events).

Example 2

Request 1: sending (Forget all passwords (except users))...

Request 1: success=true (Forget all passwords (except users)).

Example 3

Request 1: sending (Clean up job history)...

Request 1: success=true (Clean up job history).

Packages

By clicking the **Packages** tab, you can see the following information:

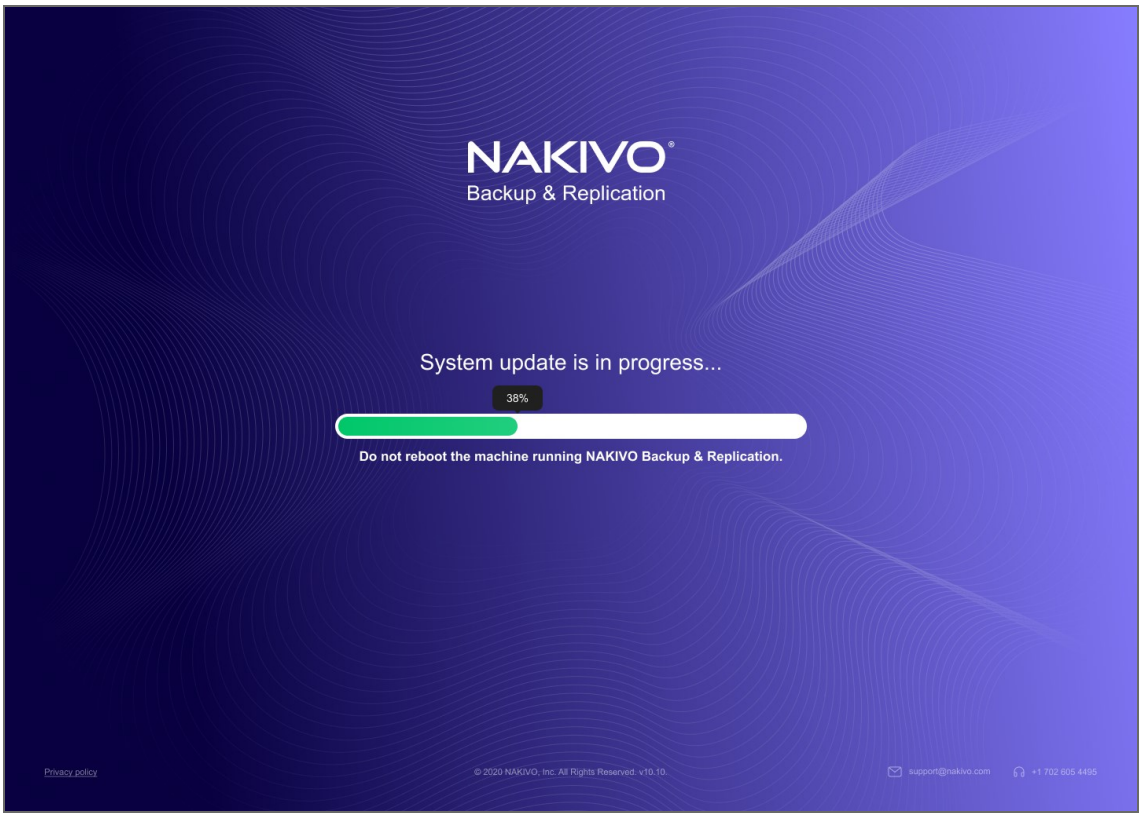
- **Base local path:** packages. Location of packages in product installation directory
- List of **Existing packages**
- List of **Supported packages**

Maintenance Mode

Your NAKIVO Backup & Replication instance can enter the Maintenance mode in the following cases:

- The product is updating
- External database migration is in progress
- The Director is being rebooted
- The Director is running out of space.

After the product enters the Maintenance mode, all the activities are stopped, and the maintenance page is displayed with the corresponding information about the maintenance activity. Below is an example of a maintenance page for system update.



Note

The informational text displayed on the maintenance page also depends on the following:

- Maintenance activities
- The instance deployment type (for example, if you are using a standalone solution or have been added as a tenant to the multi-tenant solution of NAKIVO Backup & Replication)
- The User type (Master Admin, local or remote tenant, etc.)

Director is Running out of Space

The product checks the free disk space every 30 seconds.

If less than 1 GB of free disk space is available:

- The instance user is disabled.
- NAKIVO Backup & Replication enters Maintenance mode.
- All activities stop, and the maintenance page is displayed.

To check the current free space, click **Check Now**.



If more than 2 GB of free disk space is available:

- The product exits Maintenance mode and resumes normal operation.
- All Backup Repositories and Tape Devices detached during maintenance mode are reattached.
- All Jobs that were disabled when entering maintenance mode are enabled.

Virtual Appliance Configuration

This section covers the following topics:

- “Virtual Appliance Interface” below
- “Configuring Network Settings of Virtual Appliance” on page 561
- “Increasing Backup Repository Size on Virtual Appliance” on page 561
- “Removing Disk with Backup Repository from Virtual Appliance” on page 562

Virtual Appliance Interface

The NAKIVO Backup & Replication Virtual Appliance Interface allows users to configure certain features of the NAKIVO Backup & Replication VM Appliances run in text mode.

It is displayed automatically when the VM console is launched.

It does not require installing any additional software.

Notes

- The VA Interface can also be made available by logging in remotely (for example, via SSH).
- If needed, the VA Appliance Interface can be installed or updated/disabled.
- Refer to Updating Virtual Appliance to learn more about how you can update your virtual appliance (VA).

The NAKIVO Backup & Replication Virtual Appliance Interface allows the user to configure the below features of the VM Appliance. Refer to the following topics for more information:

- "Network Settings" on page 543
- "Security Settings" on page 547
- "Time and Time Zone" on page 550
- "System Performance" on page 553
- "Manage NAKIVO Backup & Replication Services" on page 555
- "Exit to System Console" on page 560

Screen Modes and Navigation

All screens of the NAKIVO Backup & Replication Virtual Appliance Interface are available in three basic modes and follow similar navigation rules:

List mode

In this mode, a list of available items is displayed with the first list item highlighted by default. You can navigate between list items and choose the needed item by using the *<Up/Down>* arrow keys.
To open the screen that corresponds to the highlighted item, press *<Enter>*.
Pressing the *<Esc>* key opens the higher-level screen (if any).
Pressing the *<F5>* key refreshes the list.
Press the *<F10>* key to save and exit.

View mode

In this screen mode, particular VA information is displayed.
Pressing the appropriate key (screen specified) performs the required action.
Pressing the *<Esc>* key opens the higher-level screen (if any).

Wizard mode

In this mode, the user is guided by text instructions through action steps. You can enter data or choose options (screen specified).

Pressing the *<Esc>* key will open the confirmation screen. If the action is confirmed, the higher-level screen is opened.

Main Screen

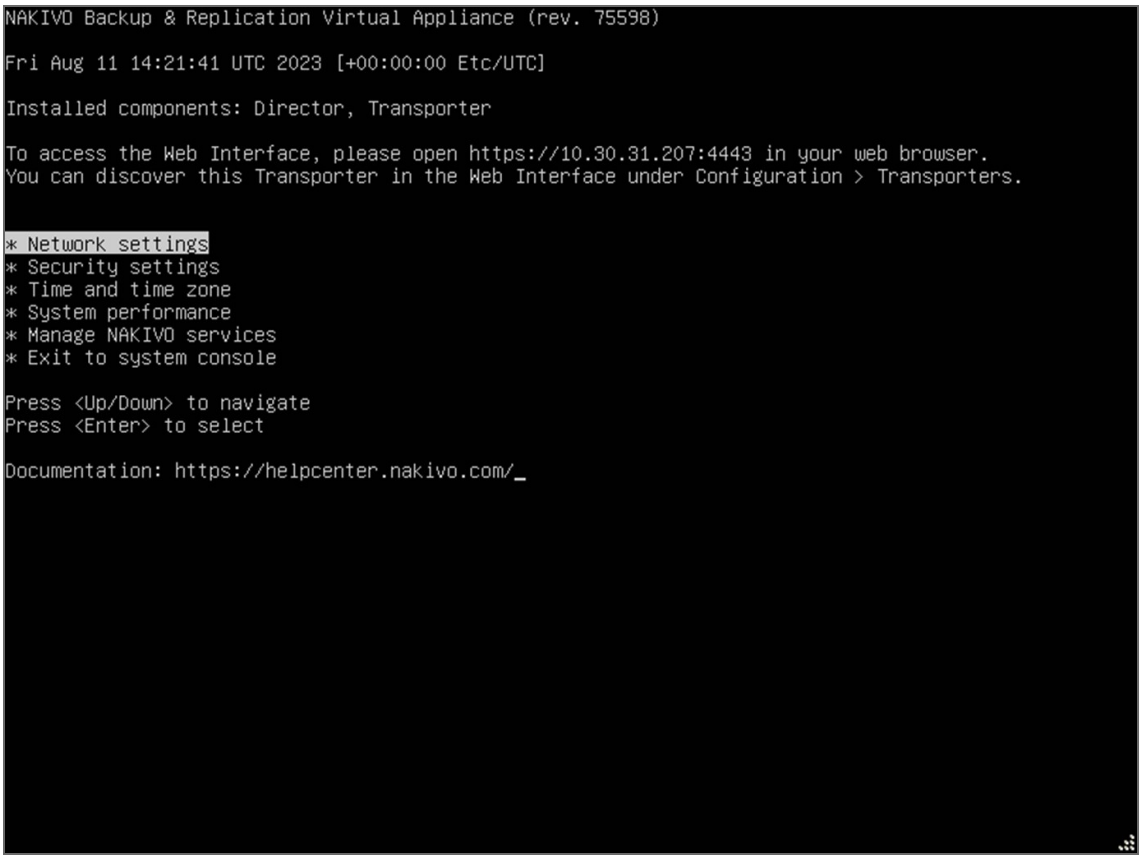
The main screen opens in the **List** mode displaying the information on the NAKIVO Backup & Replication VM Appliance interface revision version, its build number and date.

> **Important**
> The **NAKIVO Backup & Replication command console** option is displayed only if the **Director** component is present.

The installed components are listed below: "Director", "Transporter", or "Director, Transporter" (depending on the installed components).

> **Note**
> If the machine has no configured IP addresses, the following message is displayed:
> *Networking is not configured. Please open the Network settings and configure the networking.*

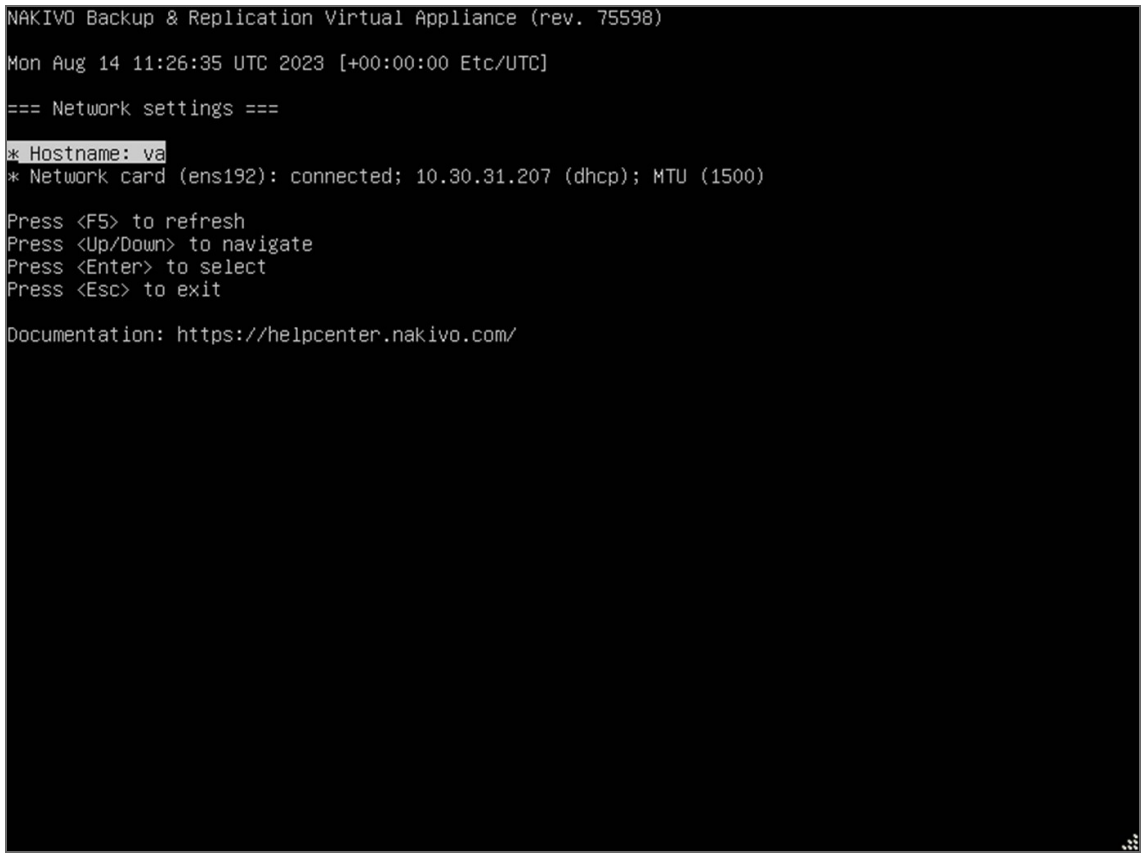
If the machine has at least one IP address configured and the **Director** component is present, the NAKIVO Backup & Replication command console option becomes available and the following message is displayed:
To access the Web Interface, please open http://<current_ip> in your web browser.
If there are multiple IP addresses, the first configured IP address is displayed.



The main screen is a starting point from where access to the VM Appliance configuration settings is granted. By default, the first list item, **Network settings**, is highlighted but you are free to navigate the list and select any other option to proceed with configuring your VM.

Network Settings

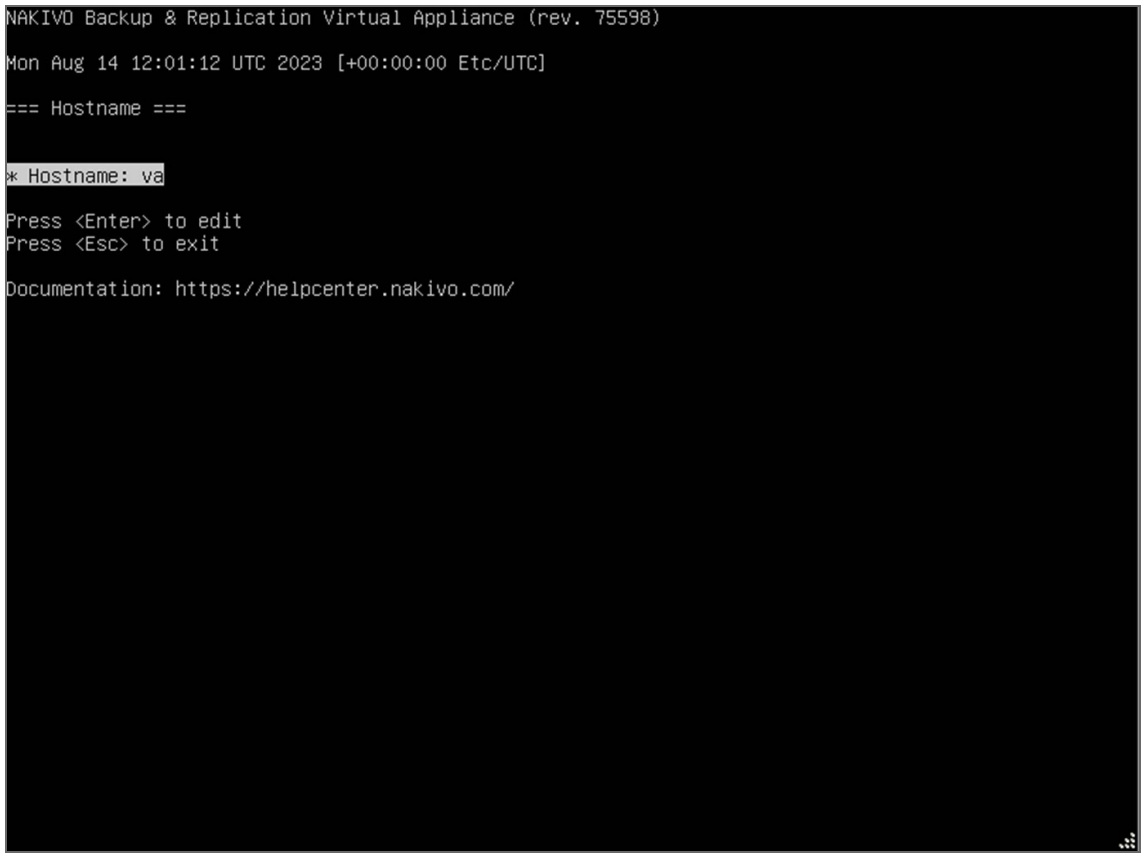
In this screen menu, you can select between two options, Hostname and Network card.



Hostname

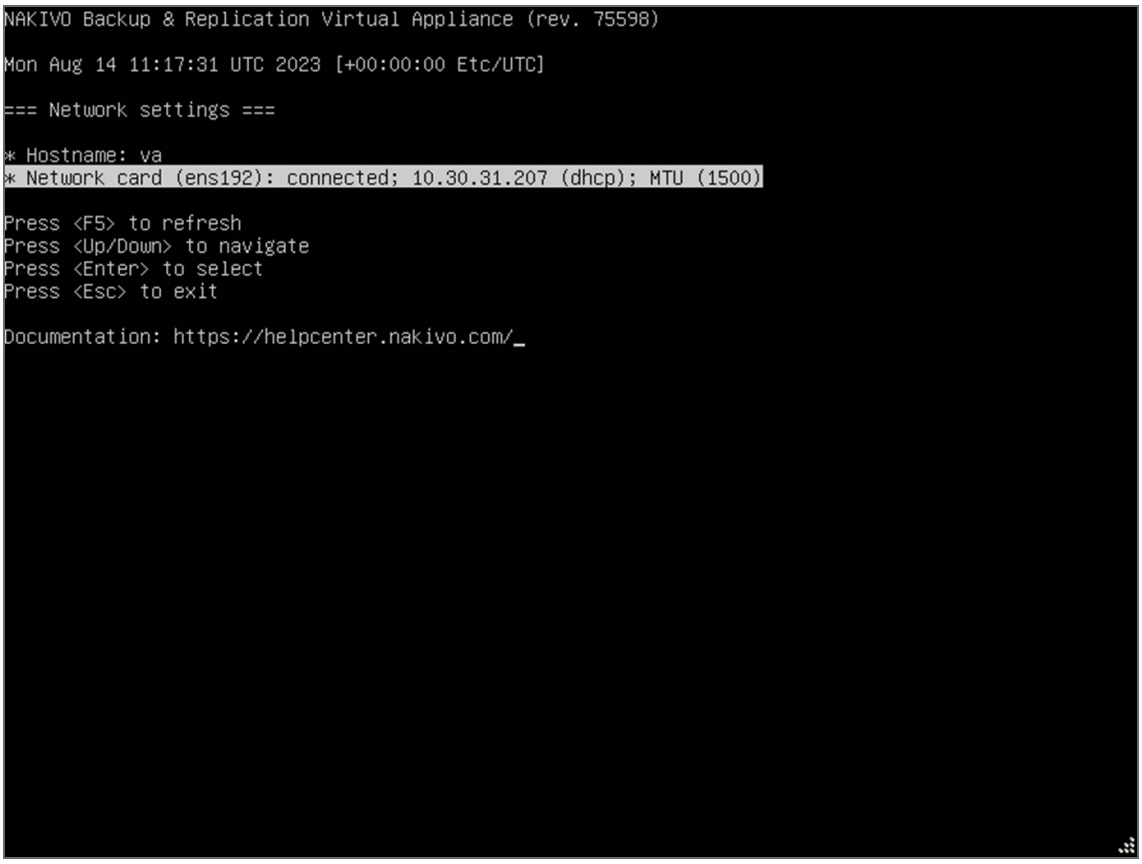
Select to view or edit the VM Appliance hostname and domain.

To edit the VM Appliance hostname, select it by pressing *<Enter>* and then follow the instructions on the screen.



Network card

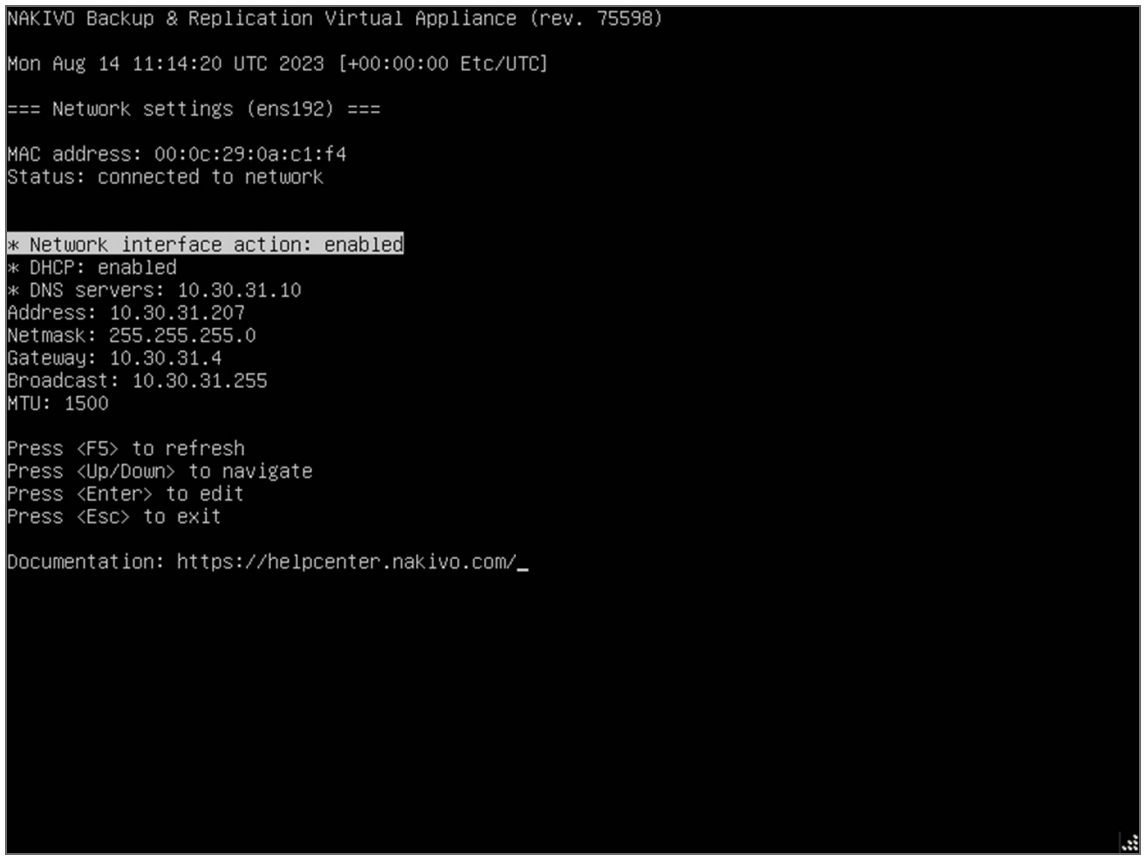
Select to view or edit the settings of the network adapters of the VM Appliance.



To open the network card menu data, select the needed item by pressing <*Enter*>.

Three options become available for editing:

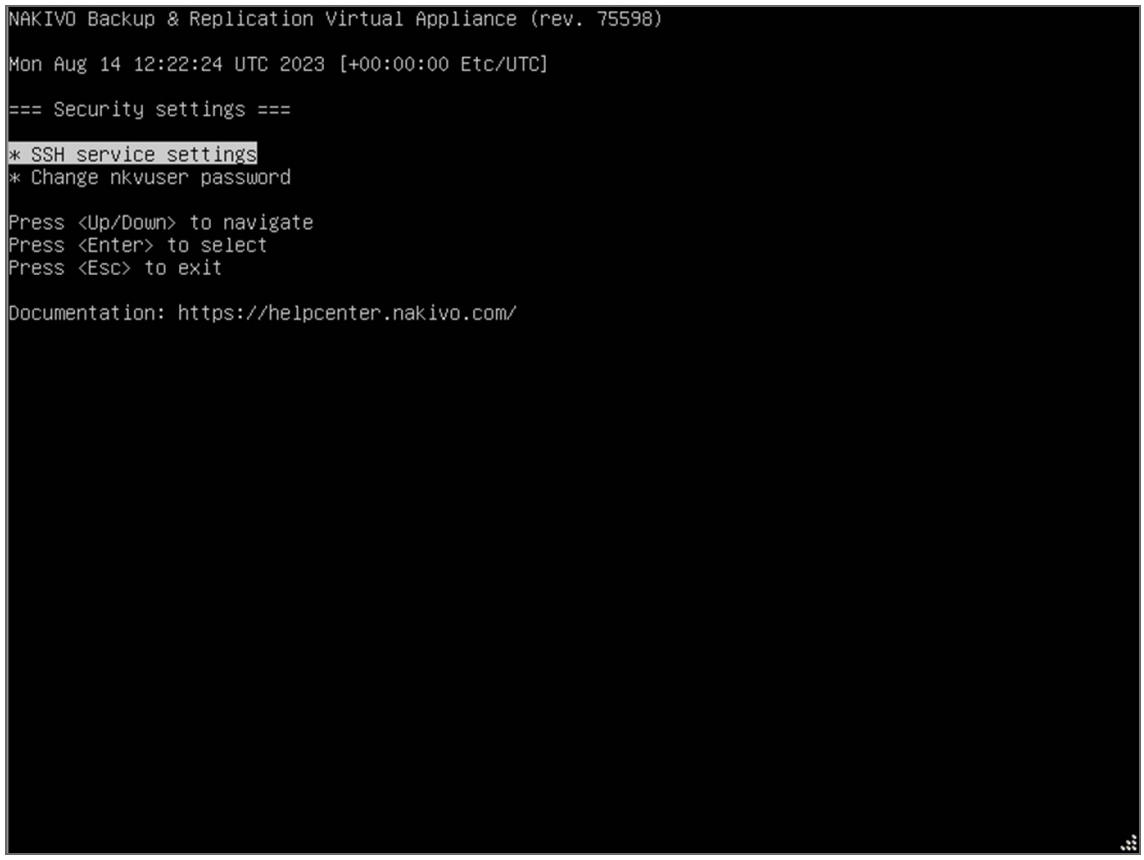
- **Network interface actions** (*enabled* or *disabled*).
- **DHCP** (*enabled* or *disabled*).
- **DNS servers**. Multiple DNS servers are supported but only valid values are accepted after entering and submitting.



To edit the VM Appliance network card settings, select it by pressing <Enter> and then follow the instructions on the screen.

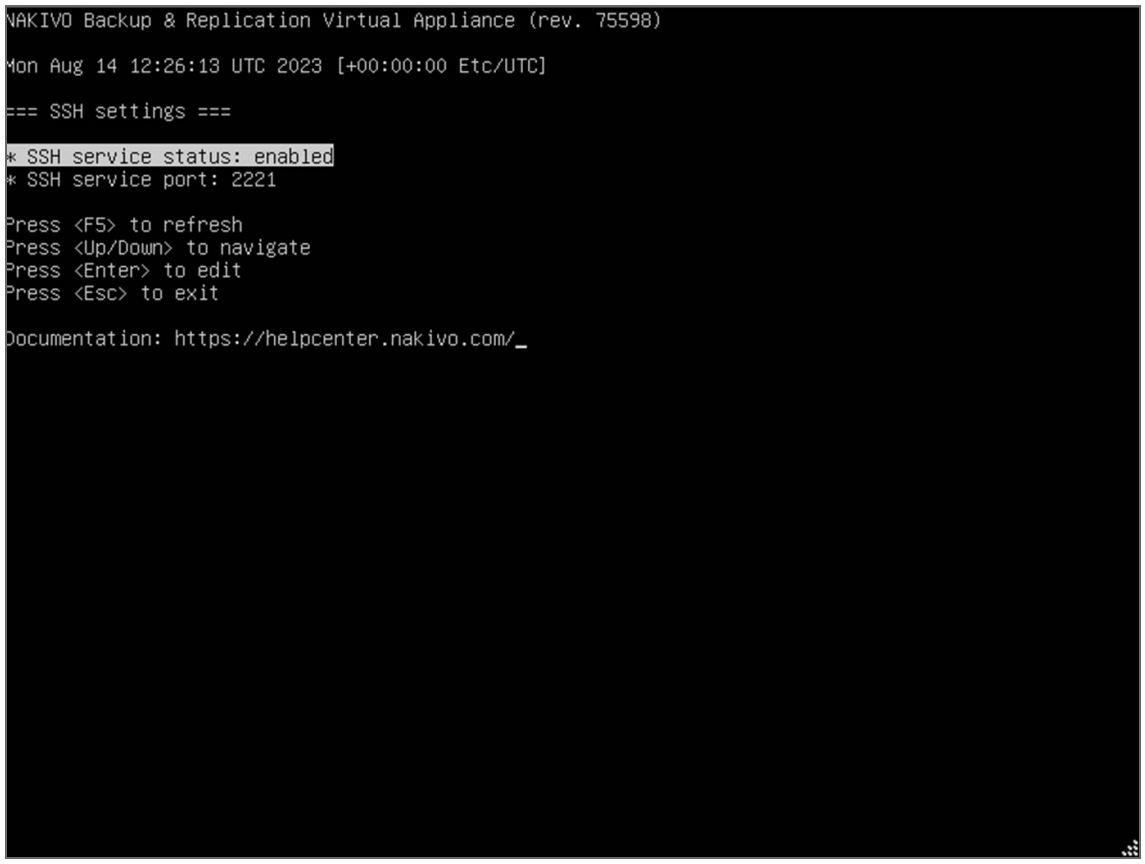
Security Settings

In this screen menu, you can select between two options, SSH service settings and Change [username] password.



SSH service settings

Select to view or edit the SSH service status and SSH service port by following the instructions on the screen.

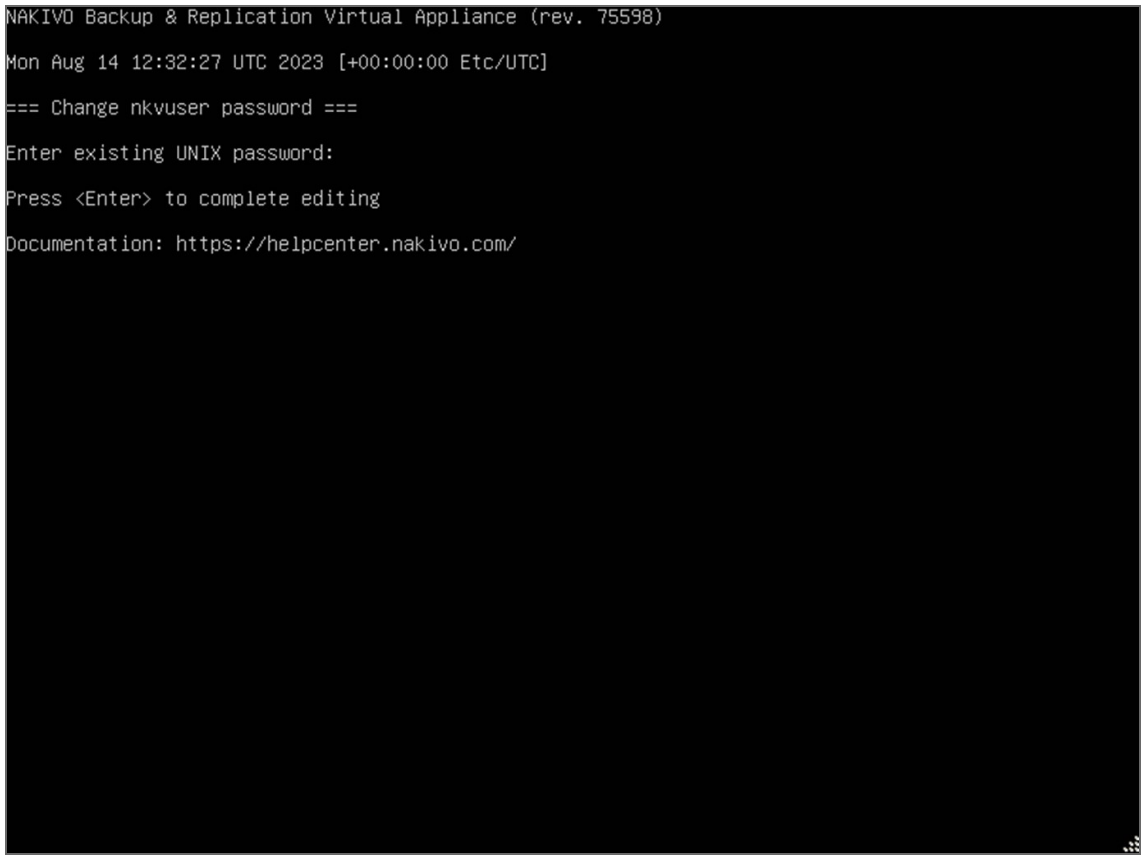


Change [username] password

Select to view or change the [username] password.

To apply the password change, entering the existing UNIX password is required.

If access to password change is authorized, enter the new password, repeat it, and then press *<Enter>* to confirm.

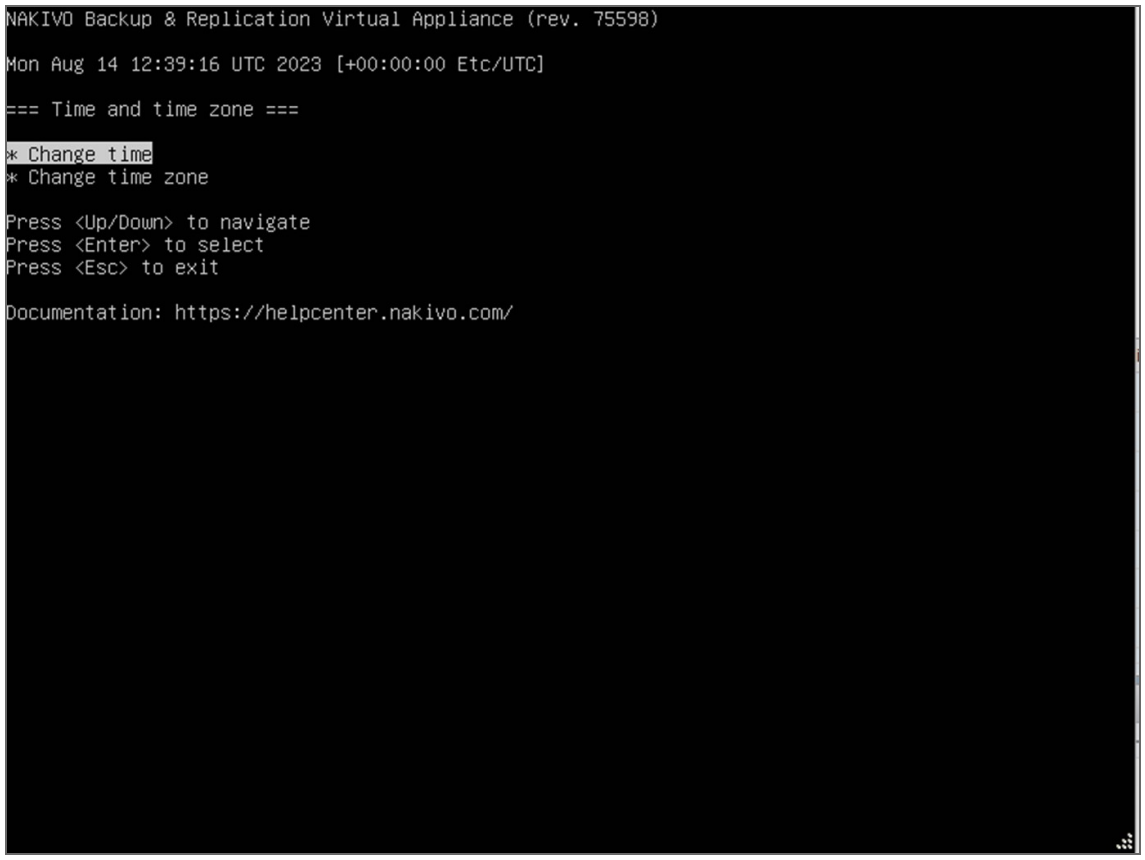


Note

To exit this screen, press *<Esc>*.

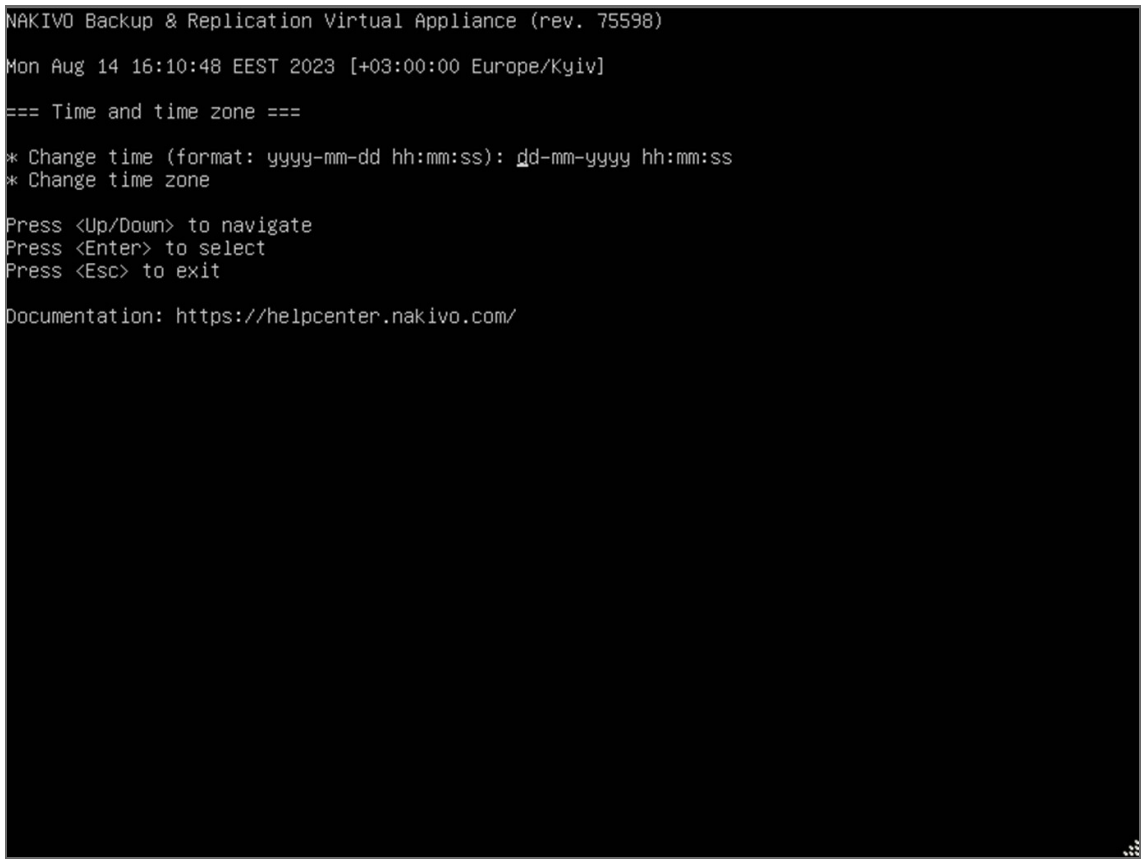
Time and Time Zone

In this screen menu, you can select between two options, Change time and Change time zone.



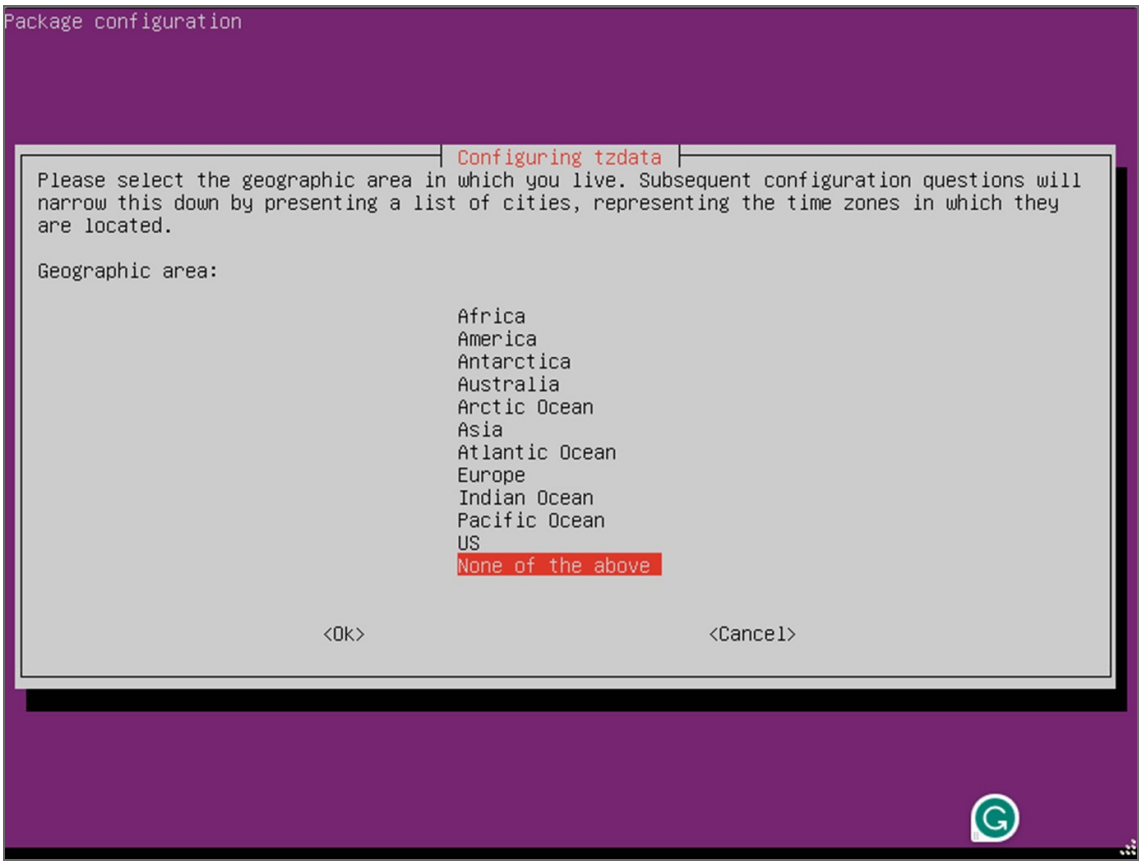
Change time

Select to view or change the VA time by following the instructions on the screen.



Change time zone

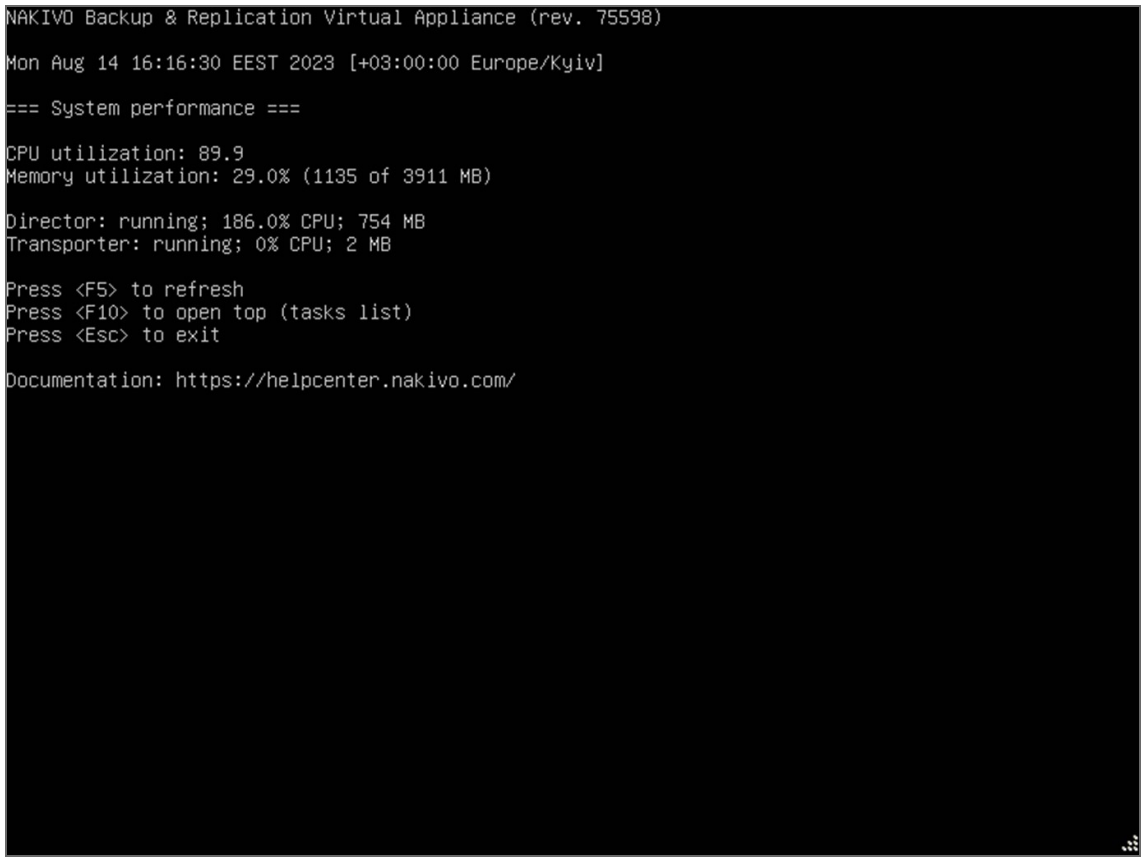
Select to view or change the VA time zone configuration by following the instructions on the screen.



System Performance

In this screen, you can view the current system consumption and the consumption of the following system resources:

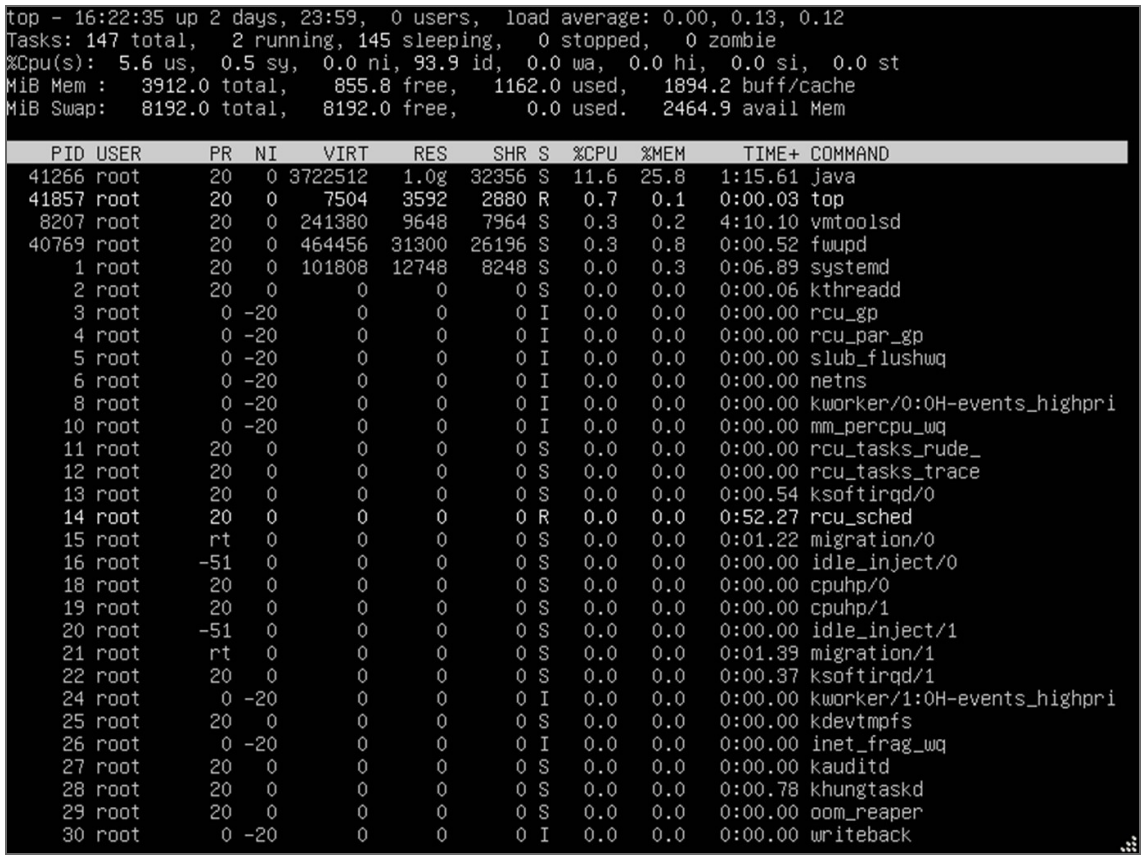
- CPU utilization by **Director** and **Transporter**
- Memory utilization by **Director** and **Transporter**



Notes

- The "Director" item is displayed only if the **Director** component is present.
- The "Transporter" item is displayed only if the **Transporter** component is present.
- The "Transporter" item accounts for all **Transporte**r processes.

To open the top tasks list, press *<F10>*.



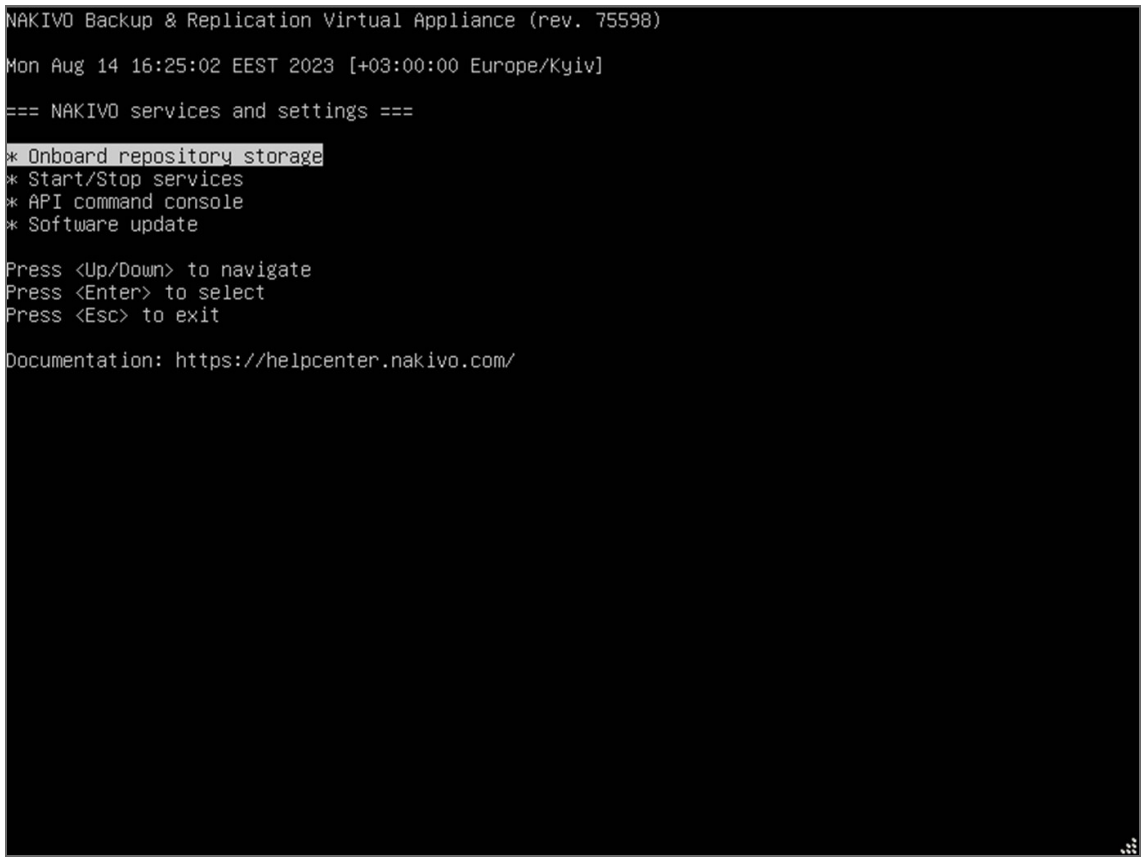
Notes

- Press *<F5>* to refresh.
- To exit this screen, press *<Esc>*.

Manage NAKIVO Backup & Replication Services

Click on **Manage NAKIVO services** opens the NAKIVO services and settings menu where you can select between four options:

- “Onboard repository storage” on the next page
- “Start/Stop services” on page 558
- “Manage NAKIVO Backup & Replication Services” above
- “Software update” on page 559



Onboard repository storage

Select to view information about the VM Appliance backup repository disk(s) kept on a logical volume (spread across multiple physical volumes).

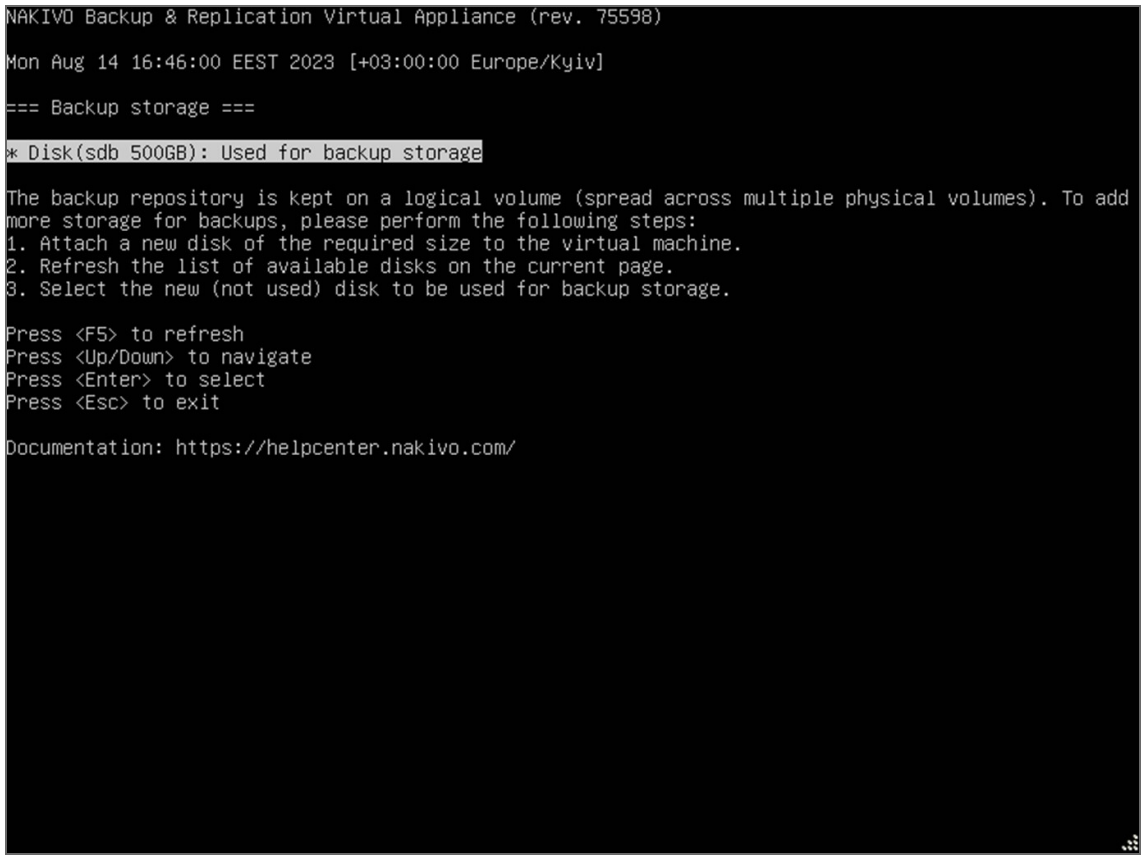
In this screen, you can also attach a new disk to your VA and then configure it for further use as backup storage by following the instructions on the screen.

To add more storage for backups, perform the following steps:

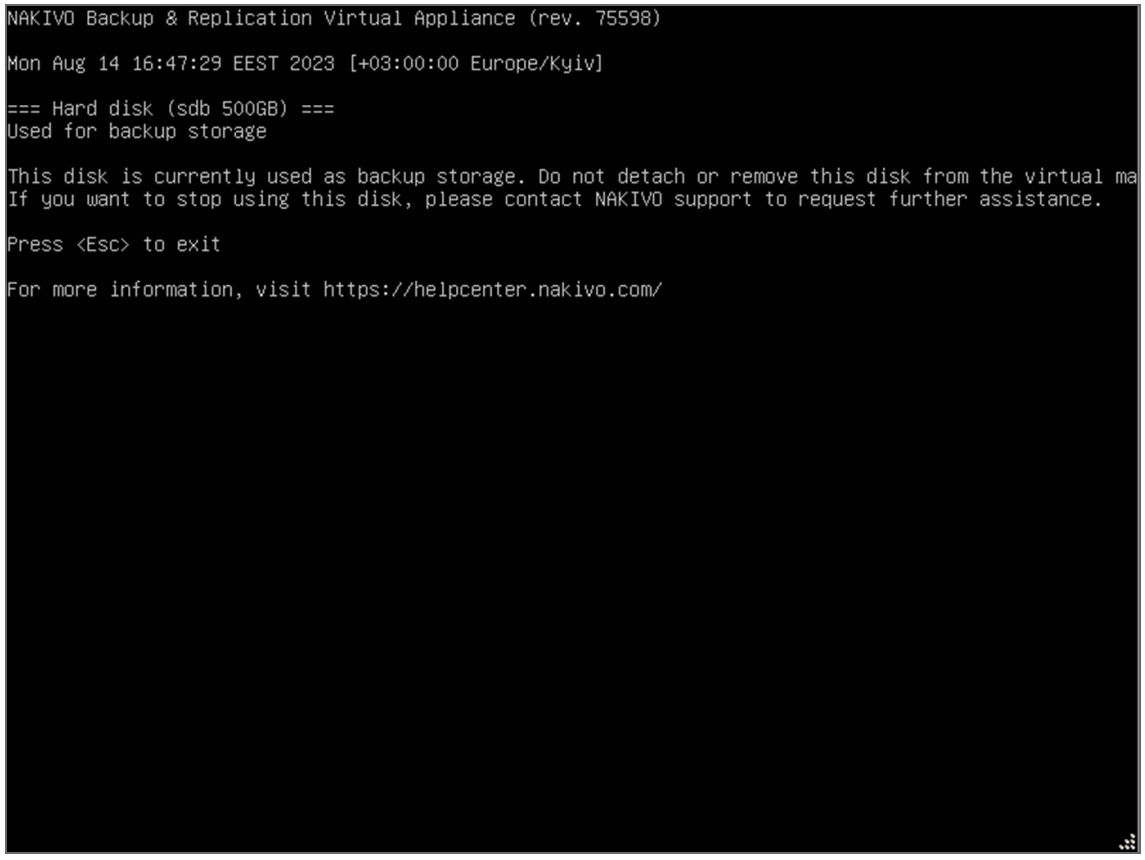
1. Attach a new disk of the required size to the virtual machine.
2. Refresh the list of available disks on the current page.
3. Select the new (not used) disk to be used for backup storage.

Notes

- The screen is not available if the **Transporter** component is not installed.
- Creation of a logical structure for backup storage is supported.



For example, if the **Disk** option is selected, the screen with the hard disk details opens.



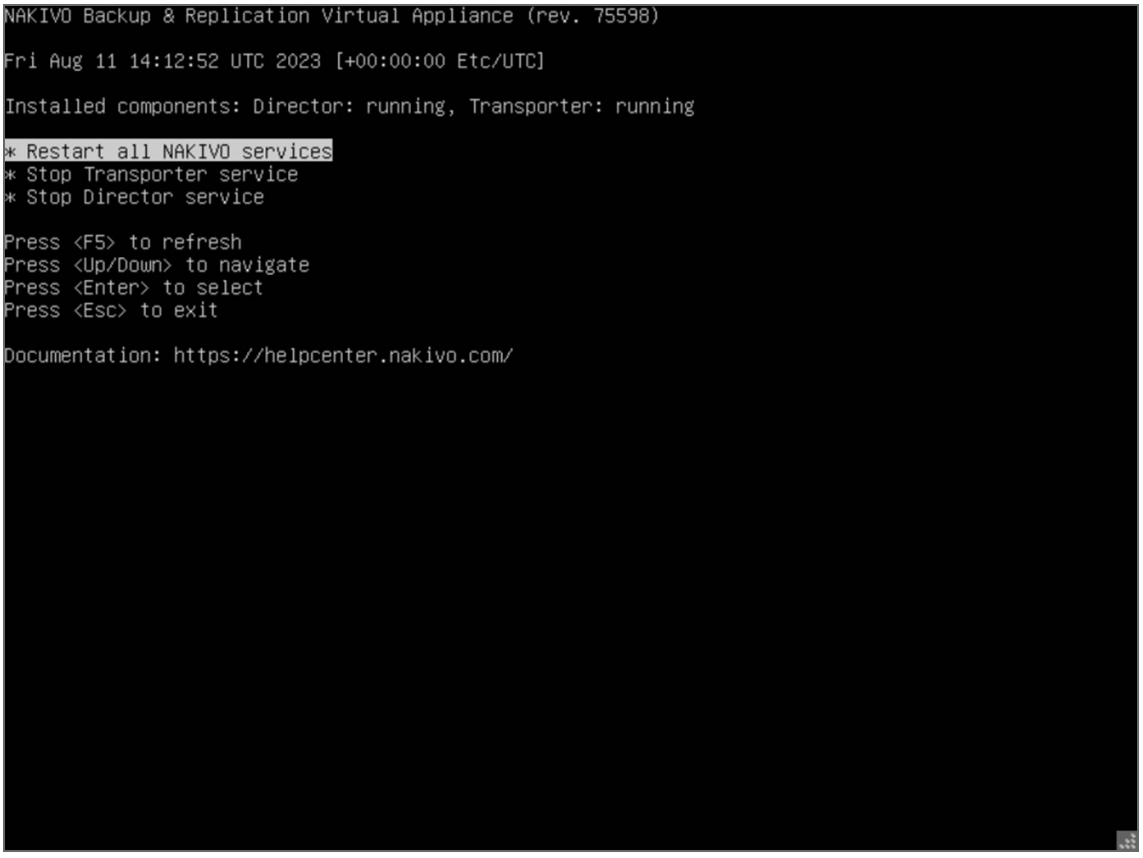
Important

- Do not detach or remove the disk currently used for backup storage from the virtual machine.
- If you want to stop using this disk, please contact NAKIVO support to request further assistance.
- If you select to use the disk that is not currently used for backup storage, it will be formatted and appended to the existing logical volume.
- If there is no logical volume structure yet, it will be created automatically.

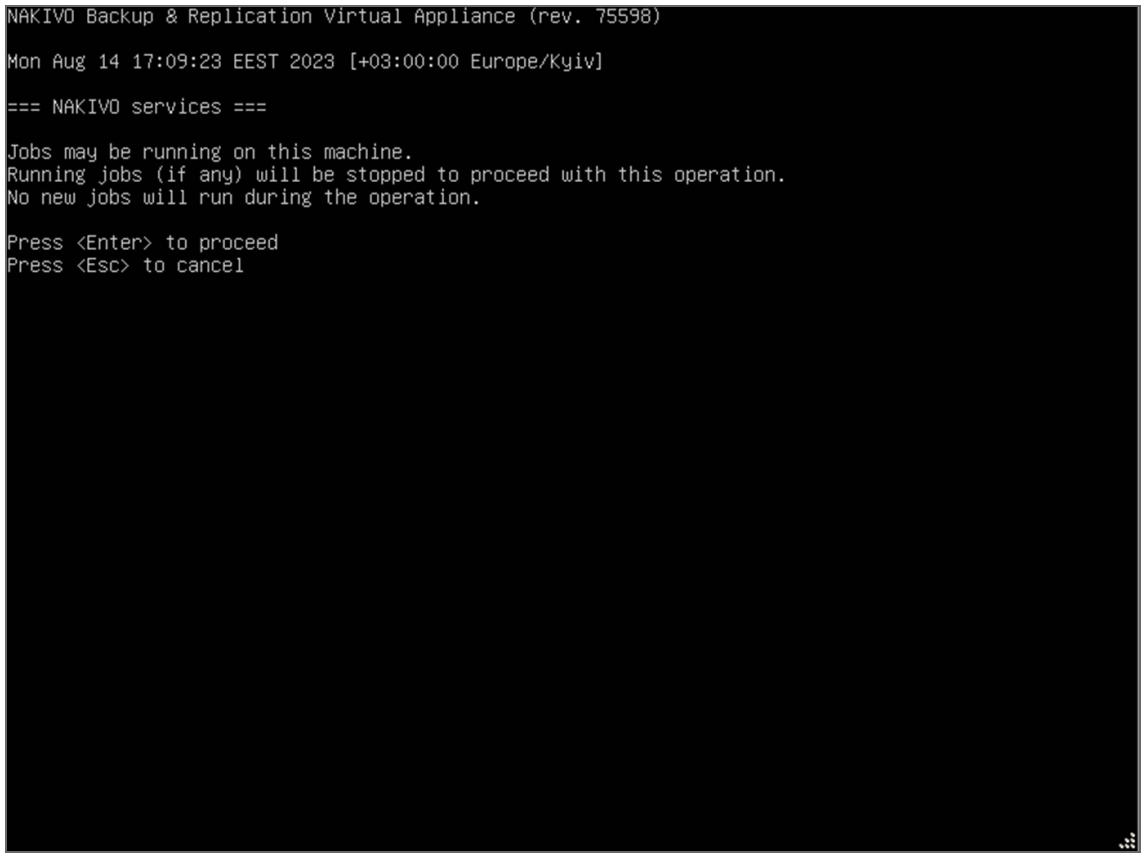
Start/Stop services

Select to manage the NAKIVO Backup & Replication Virtual Appliance services. The menu offers to select between three options:

1. Restart all NAKIVO services
2. Stop Transporter service
3. Stop Director service



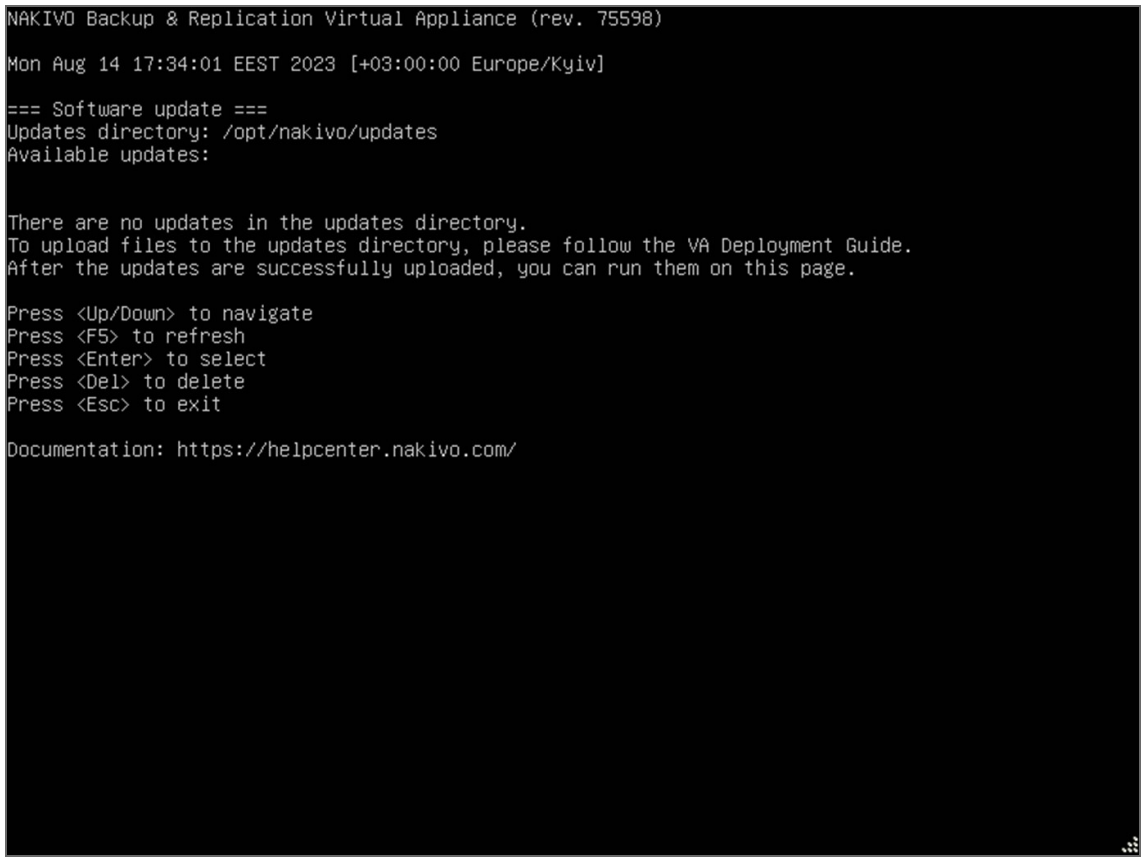
Note that if you decide to select any option, you are asked to confirm your selection.



Software update

Select to check for available software updates in a certain directory and manage the available updates (if any).

If you decide to start the software update, upload files to the updates directory by following the VA Deployment Guide (available at www.nakivo.com/documentation.htm).



After the updates are uploaded, you can apply them on the current page.

Exit to System Console

Select to close the VM Appliance Interface and exit to the system console.



Configuring Network Settings of Virtual Appliance

To configure networking on the Virtual Appliance (VA), follow the steps below:

1. Open the VA console.
2. On the main menu, select the **Network Settings** option and press **Enter**.
3. Do either of the following:
   - To change the Virtual Appliance hostname, select the **Hostname** option, press **Enter**, enter a new hostname, and press **Enter** again.
   - To configure a network card, select it and press **Enter**. Press **Enter** to switch between DHCP and manual network settings. If you set the **DHCP** option to **disabled**, you can manually set up network settings by selecting an option, pressing **Enter**, entering a new value, and pressing **Enter** again. Press **F10** to save your changes and exit.

   For more details, check the "Network Settings" on page 543 page.

Increasing Backup Repository Size on Virtual Appliance

A **Backup Repository** on a Virtual Appliance (VA) is located in a logical volume (that can spread across multiple physical volumes). To extend the **Backup Repository** size on the VA, you need to add a new disk to the VA and then use the VA console to extend the **Backup Repository** to the new disk.

The **Backup Repository** size on the VA cannot be increased by extending existing VA disks.

The backup repository size on the VA cannot be increased by extending existing VA disks. To increase the size of the backup repository on the Virtual Appliance, follow the steps below:

1. Attach a new disk to the VA.
2. Open the VA console in your hypervisor's client.
3. Run the following commands in the VA console depending on the NAKIVO Backup & Replication version you use:
    - For the product Version 8.1 and higher:
        a. Select **Manage NAKIVO services** in the main menu and press **Enter**.
        b. Select **Onboard repository storage** and press **Enter**.
    - For earlier product versions, select **Backup storage** in the main menu and press **Enter**.
4. Refresh the list of disks by pressing **F5**.
5. Select the disk that you have created and press **Enter**.
6. Press **Enter** again to confirm the procedure. The disk is formatted and added to the **Backup Repository** on the VA.

For more details, check the "Manage NAKIVO Backup & Replication Services" on page 555 page.

Removing Disk with Backup Repository from Virtual Appliance

The Virtual Appliance (VA) comes with a 500 GB disk on which a **Backup Repository** is created. If you have deployed the Virtual Appliance disks using the **Thin Provision** option, then the disk does not consume 500 GB of space on your datastore – only the space occupied by VM backups is consumed.

If you still would like to delete the 500GB disk after you have deployed the Virtual Appliance, follow the steps below:

1. Log in to NAKIVO Backup & Replication.
2. Go to the **Configuration** > **Repositories** tab.
3. Click **Onboard repository**
4. Click **Manage** and choose **Remove** from the menu.
5. In the message that opens, click the **Remove Repository and Delete Backups** button.
6. Click **Remove** to confirm that you wish to remove the **Backup Repository**.
7. Open the vSphere client and launch the console of the VA.
8. In the Virtual Appliance interface, select the **Exit to system console** option and press **Enter**.
9. Enter a login and password (default are `root/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 565
- "Configuring Branding Settings in Multi-Tenant Mode" on page 566
- "Configuring Email Notifications in Multi-Tenant Mode" on page 568
- "Configuring Email Settings in Multi-Tenant Mode" on page 569
- "Configuring System Settings in Multi-Tenant Mode" on page 570
- "Exporting and Importing Configuration in Multi-Tenant Mode" on page 572
- "Viewing Tenant Licensing Details in Multi-Tenant Mode" on page 573

Changing Login and Password in Multi-Tenant Mode

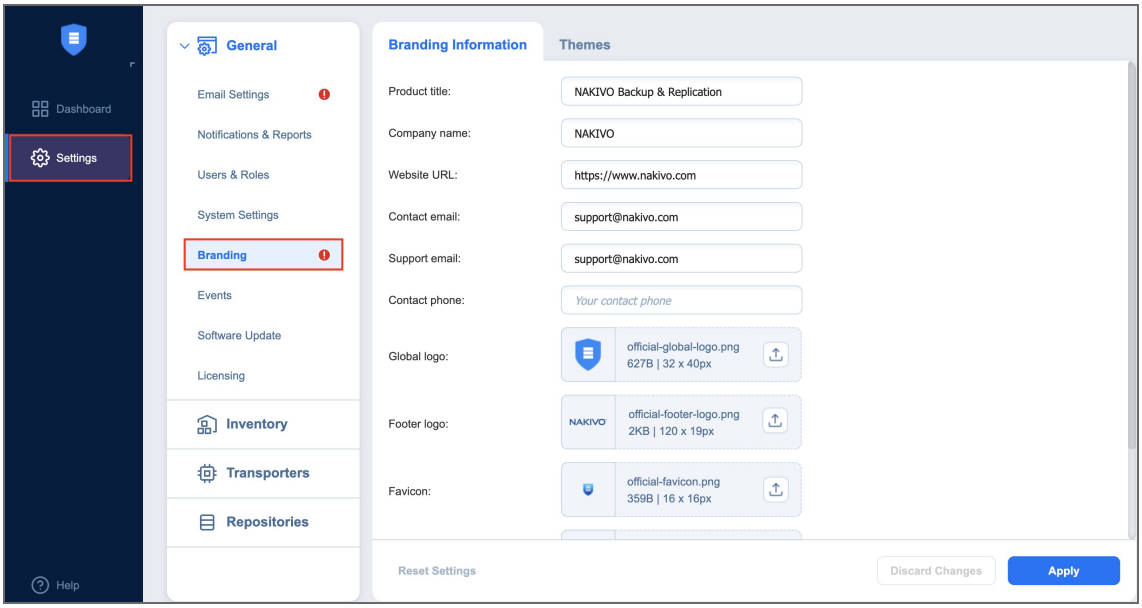
To change the login and password of the Master Admin, follow the steps below:

1. Log in to NAKIVO Backup & Replication as a Master Admin.
2. Click **Configuration** in the upper right corner of the product.
3. Go to the **General** tab and click **Users and Roles**.
4. In the list of users that opens, click the Master Admin user.
5. For the Master Admin, enter data in the **Login, Password**, **Confirm Password**, and **Admin email** boxes and click **Apply.**

Configuring Branding Settings in Multi-Tenant Mode

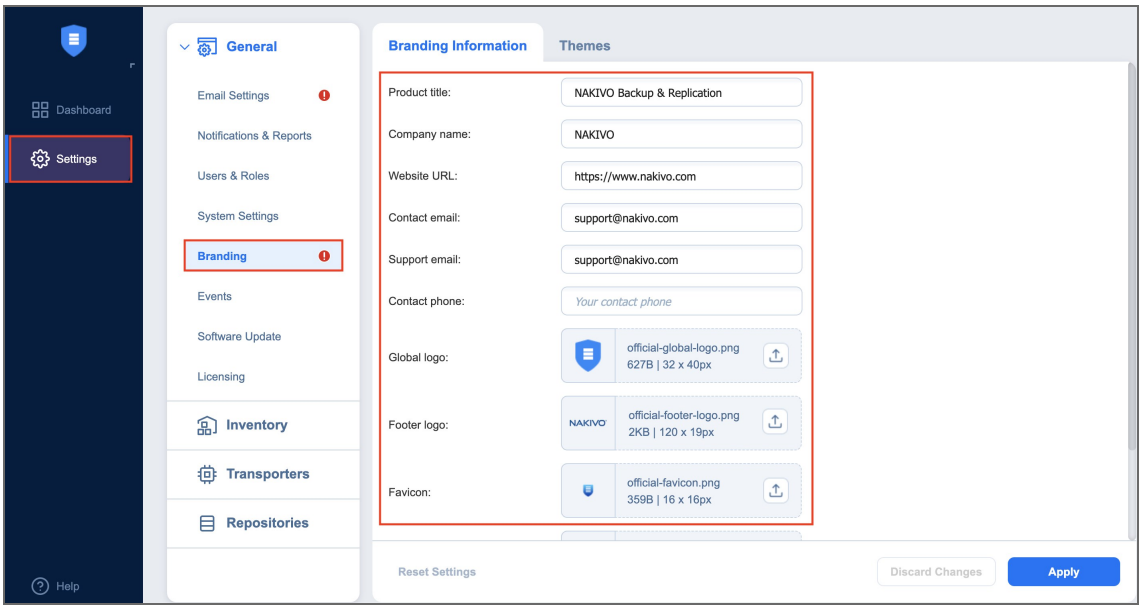
In the multi-tenant mode, you can change the product branding settings such as product name, logo, background, and so on. To configure the system settings, follow the steps below:

1. Log in to NAKIVO Backup & Replication as a Master Admin.
2. Click **Settings** in the left pane of the product.
3. Go to the **General** tab and click **Branding**.



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

appropriate box, select a new image, and click **Open**.



5. Click **Apply**.

Note

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

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

Configuring Email Notifications in Multi-Tenant Mode

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

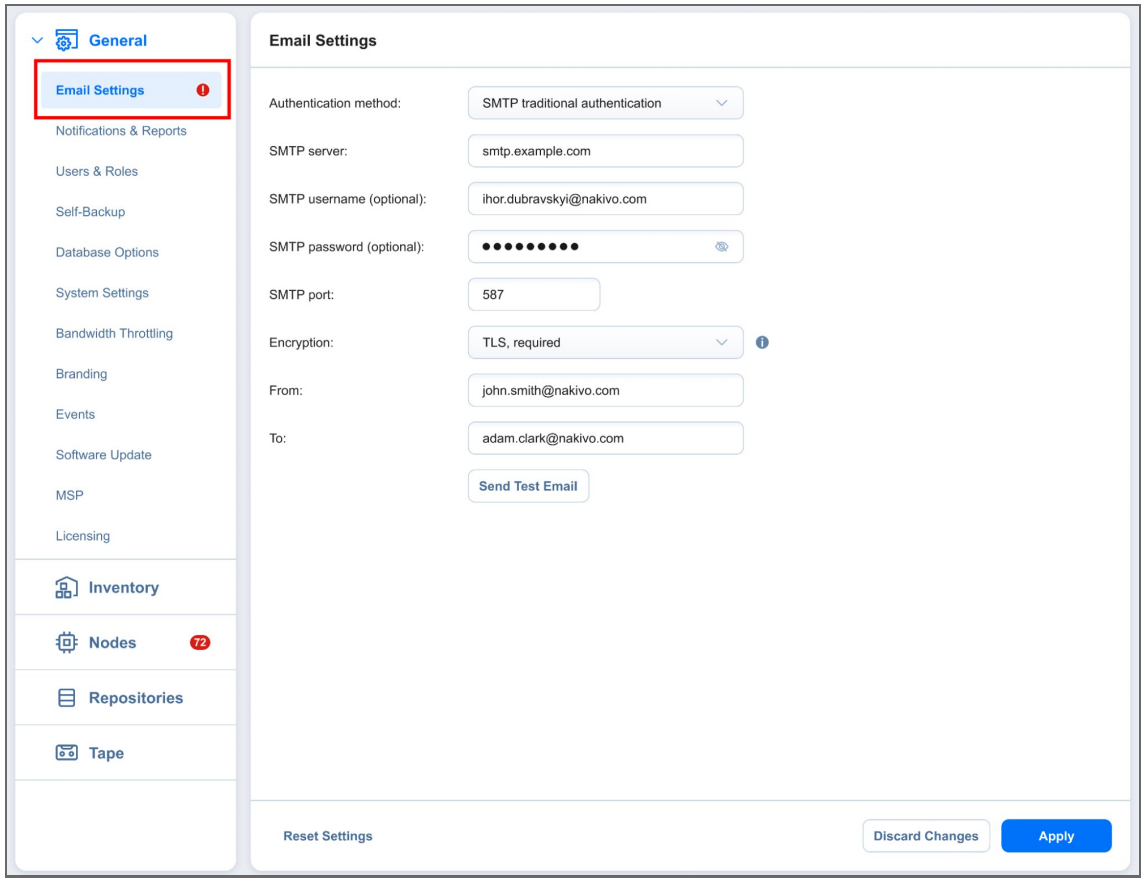
1. Make sure you have configured your email settings.
2. Log in to NAKIVO Backup & Replication as a Master Admin.
3. Click **Settings** in the left pane of the product and go to the **General** tab.
4. Click **Email settings**.
5. In the **Email Notifications** section, select the options as appropriate:
    a. **Send alarm (error) notifications**: If selected, this will send notifications about a job, repository, infrastructure, connection, and other failures to email addresses specified in the text field. Use a semicolon to separate multiple email addresses.
    b. **Send warning notifications**: If selected, this will send warning notifications on non-critical events, such as infrastructure change, to email addresses specified in the text field. Use a semicolon to separate multiple email addresses.
    c. **Limit email notification frequency to**: Set a limit to how often email notifications are sent.
6. In the **Automatic Reports** section, select or deselect the following automatic reports options:
    - **Attach PDF copy to automatic reports**: Specify whether you wish to include a copy of the PDF report with notifications.
    - **Send tenant Overview reports on schedule to**: If this option is selected, NAKIVO Backup & Replication will generate an Overview report (which includes information about all jobs and groups in the product) on the date and time specified in the scheduler and will send the report to the recipients specified in the text field. Use a semicolon to separate multiple email addresses.
    - **Send tenant Protection Coverage reports on schedule to**: If this option is selected, NAKIVO Backup & Replication will generate the Protection Coverage report (which includes information about all VMs & instances protected by backup and/or replication jobs as well as the information about all unprotected VMs & instances) on the date and time specified in the scheduler and will send the report to the recipients specified in the text field. Use a semicolon to separate multiple email addresses.
    - Click **Apply**.

Configuring Email Settings in Multi-Tenant Mode

Configure email settings so that NAKIVO Backup & Replication can send email notifications as well as reports over email. If email settings are not configured, tenants will not be able to configure email notifications for their jobs. To configure email settings**,** follow the steps below:

1. Log in to NAKIVO Backup & Replication as a Master Admin.
2. Click **Settings** in the left pane of the product.
3. Go to the **General** tab and click **Email notifications**.
4. In the **Email Settings** section, enter data in the boxes, and click **Send Test Email** to verify the settings are correct.

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



Configuring System Settings in Multi-Tenant Mode

To configure the system settings, follow the steps below:

1. Log in to NAKIVO Backup & Replication as a Master Admin.
2. Click **Settings** in the left pane of the product.
3. Go to the **General** tab and click **System settings**.
4. Select or deselect the following options:
    - **Store system events for the last X days**: This option specifies the time period (from 10 to 365 days) during which the application events will be kept. Older events are automatically deleted.
    - **Auto log out after X minutes of inactivity**: If this option is selected, the current user will be automatically logged out of the product after the specified period of inactivity.
    - **Auto upload support bundles to support team server**: If this option is enabled, NAKIVO Backup & Replication will automatically create, encrypt, and upload support bundles once a day to a NAKIVO support server during the evaluation period. The NAKIVO Support team may use this information to improve the product experience and will be able to identify and resolve product issues faster.
    - **Enable built-in support chat:** If selected, this will allow you to chat with the NAKIVO support team.
    - **Display special offers**: If selected, this will show a toolbar with special offers in the GUI.
    - **Continue product update if self-backup fails**: If selected, product update will proceed even if automatic self-backup cannot be performed.
    - **Tape options**: These present you with setting options for tape devices:
        - **Auto erase expired tapes**: If selected, expired tape cartridges will be erased automatically.
        - **Wait for next tape for**: Specify for how long the system needs to wait for the next tape cartridge if there is no appropriate one. Select the **Send email notification** checkbox to allow you to receive email notifications.
        - **Auto refresh tapes every**: Select how often the contents of tape cartridges are to be refreshed in minutes or hours. Deselect if no refreshing is required.
    - **Regional options**: Set the clock format, short date format, long date format, first day of the week, decimal symbol, and default time zone in the corresponding fields.
- In the **Web Interface TLS/SSL Certificate** section, you can either**:**
    - **View current certificate**: A dialog containing the current certificate information opens.
    - **Install new certificate**: A dialog opens, allowing you to install a new TLS/SSL certificate for the NAKIVO Backup & Replication web interface. Certificates are generated either internally or through certification authorities. Proceed as follows to install a new certificate:

- Click **Browse** and navigate to the location of either of the following certificate file types:
  - **Private key**: A file in the *.key format.
  - **Private key password (optional)**: A password for your private key.
  - **Certificate file**: A file in the *.pem, *.crt, *.cer, *.p7b, or *.p7s format.
  - **Intermediate certificate (optional)**: A file in the *.pem, *.crt, *.cer, *.p7b, *.p7s format.
- Click **Install**.

Note

In the Web Interface TLS/SSL Certificate section, you can see a notification about imminent TLS/SSL certificate expiration in 30 days and onwards. If your certificate has expired, you will be asked to install a valid certificate.

Exporting and Importing Configuration in Multi-Tenant Mode

System configuration export and import are recommended for easy migration to new product deployment. System configuration, such as jobs, user credentials, inventory items, Transporter and Backup Repository settings, is all exported into a single export bundle.

The export bundle can be applied to a new deployment.

To export system configuration from the old deployment, follow the steps below:

1. Open **Settings** in the old deployment.
2. Go to the **General** tab and click **System migration**.
3. Click **Export system configuration**.
4. In the dialog box that opens, click **Export**.
5. Click **Proceed** to confirm the operation.

> **Note**
> All activities in the old deployment (such as jobs and recovery sessions) will be automatically stopped and disabled.

6. Wait until the export is completed, and download the export bundle.

To import system configuration into the new deployment, follow the steps below:

1. Open **Settings** in the new deployment.
2. Go to the **General** tab and click **System migration**.
3. Click **Import system configuration**.
4. In the dialog window that appears, locate the export bundle using the **Browse** button.
5. Click **Import**.
6. Click **Proceed** to confirm the operation.

> **Note**
> If there is any existing data in the new deployment, it will be overwritten with the import operation.

7. Wait until the import is completed, and close the dialog box.

Notes

1. Data contained in backup repositories is not migrated to the new location automatically. If you are using a locally attached Backup Repository, the physical data must be copied or moved to the new location manually.
   After moving the files you may need to edit the Backup Repository settings in the new deployment so that the new settings refer to the actual Backup Repository location.
2. If a custom TLS/SSL certificate of the Web server was used in the old deployment, a manual service restart will be required in the new deployment.

Viewing Tenant Licensing Details in Multi-Tenant Mode

Important

This tab is only displayed for users with an *MSP license*, *Beta instance*, *Promo license*, or *Trial license*.

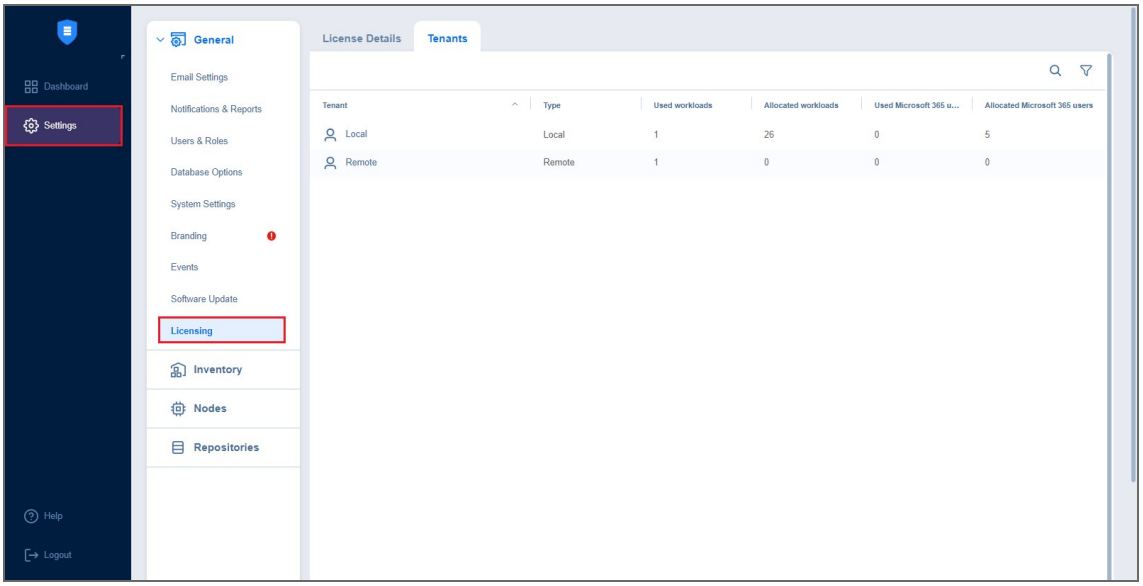
To check the license details of all local and remote tenants connected to a managed service provider (MSP) through the **MSP Console**, follow these steps:

1. Go to the main menu of NAKIVO Backup & Replication and click **Settings** > **General**.
2. Go to the **Licensing** > **Tenants** tab.

   Note

   To display the **MSP Console**, you must have the multi-tenancy mode installed.

3. The tab opens the page with the **Tenants** table that allows viewing the licensing data of all existing tenants.

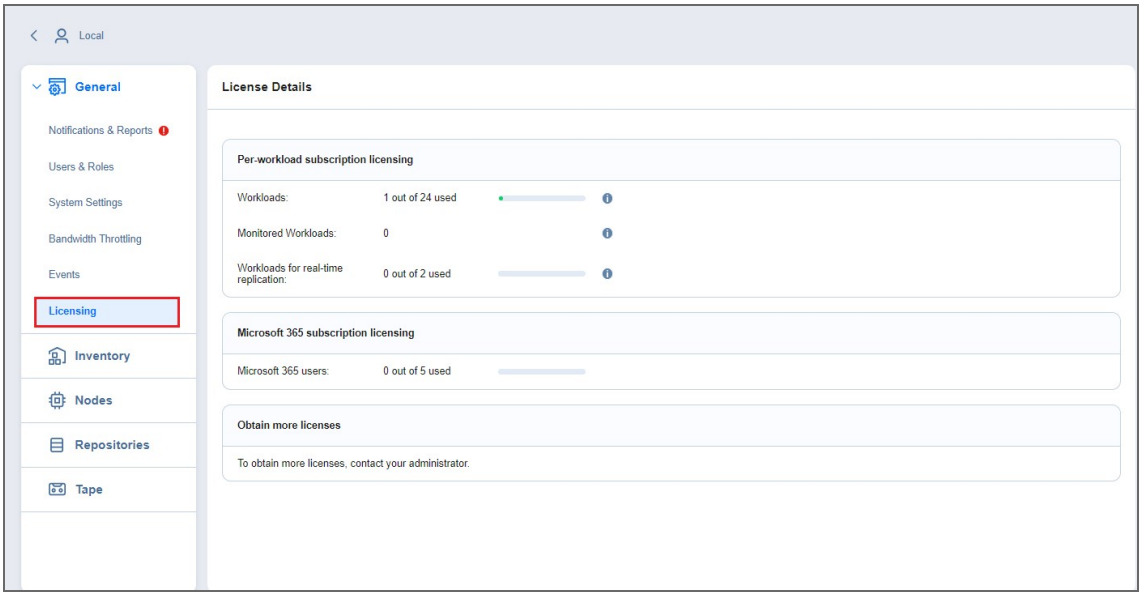


The table has the following columns:

- **Tenant**: The name of the tenant.

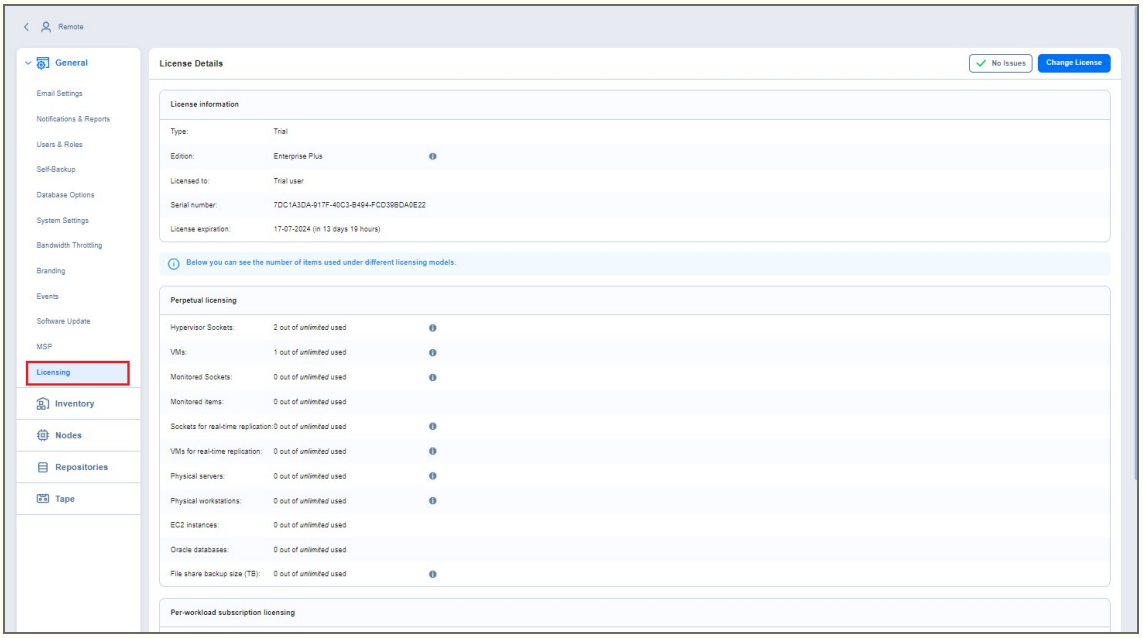
> **Notes:**
>
> - Clicking the name opens the **Licensing** tab of the corresponding tenant.
> - *Disabled*, *Disconnected*, and *Inaccessible* tenants are disabled.
> - Hovering over a tenant's data point (for example, tenant's name, type, workloads, etc.) displays the text with the time since the tenant's data was last updated.

- **Type**: The tenant's connection type – **Local** or **Remote**.
- **Used workloads**: The number of a tenant's used workloads.



Note

For the tenants with a perpetual license, the **Perpetual** link is displayed instead; clicking the link opens the **Tenant licensing details** popup displaying the **Perpetual licensing block** in the **Licensing** tab of the corresponding tenant.



- **Allocated workloads**:
  - For remote tenants, displays the number of workloads available to the tenant.
  - For local tenants, displays the number of workloads allocated to the tenant.

> **Note:**
> For the tenants with a perpetual license, the **Perpetual** link is displayed instead. Clicking the link opens the **Tenant licensing details** popup displaying the **Perpetual licensing block** in the **Licensing** tab of the corresponding tenant.

- **Used Microsoft 365 users**: The number of Microsoft 365 users used by the tenant.
- **Allocated Microsoft 365 users**:
  - For local tenants, displays the number of Microsoft 365 users allocated to the tenant.
  - For remote tenants, displays the number of Microsoft 365 users available to the tenant.

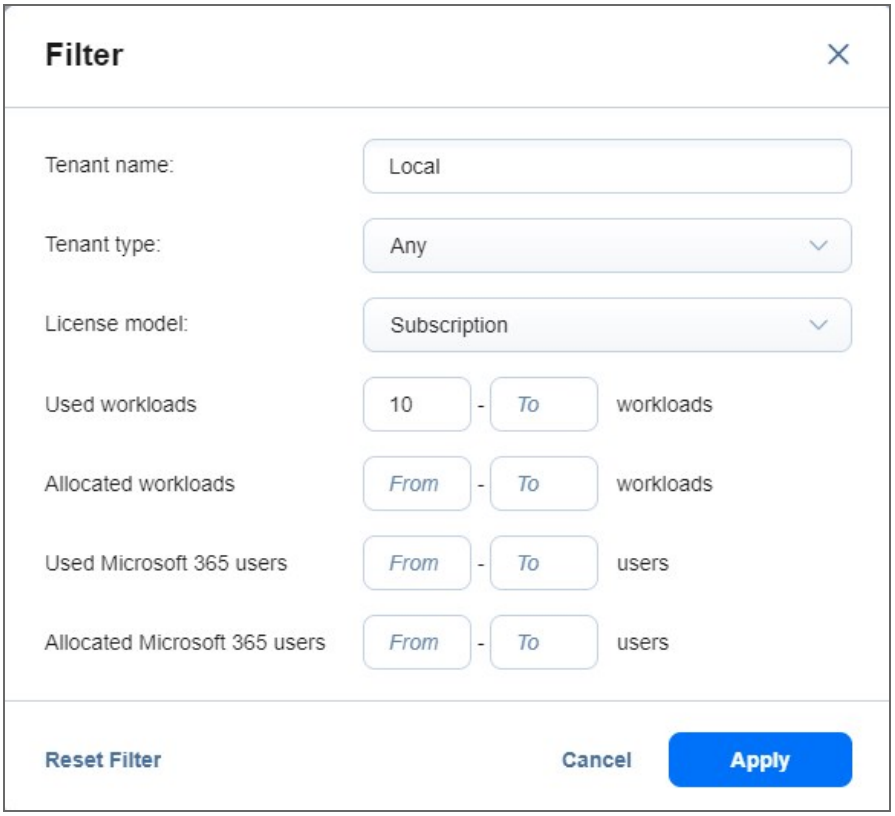
Using Filtering

You can search or filter tenants from the **Tenants** table.

Use the **Search** field to browse or search for a specific tenant. Search can be performed on the **Name** column only.

To access filtering options, click the **Filter** icon in the top right corner. In the dialog box that opens, you can select one or more filtering criteria. The following filtering options are available:

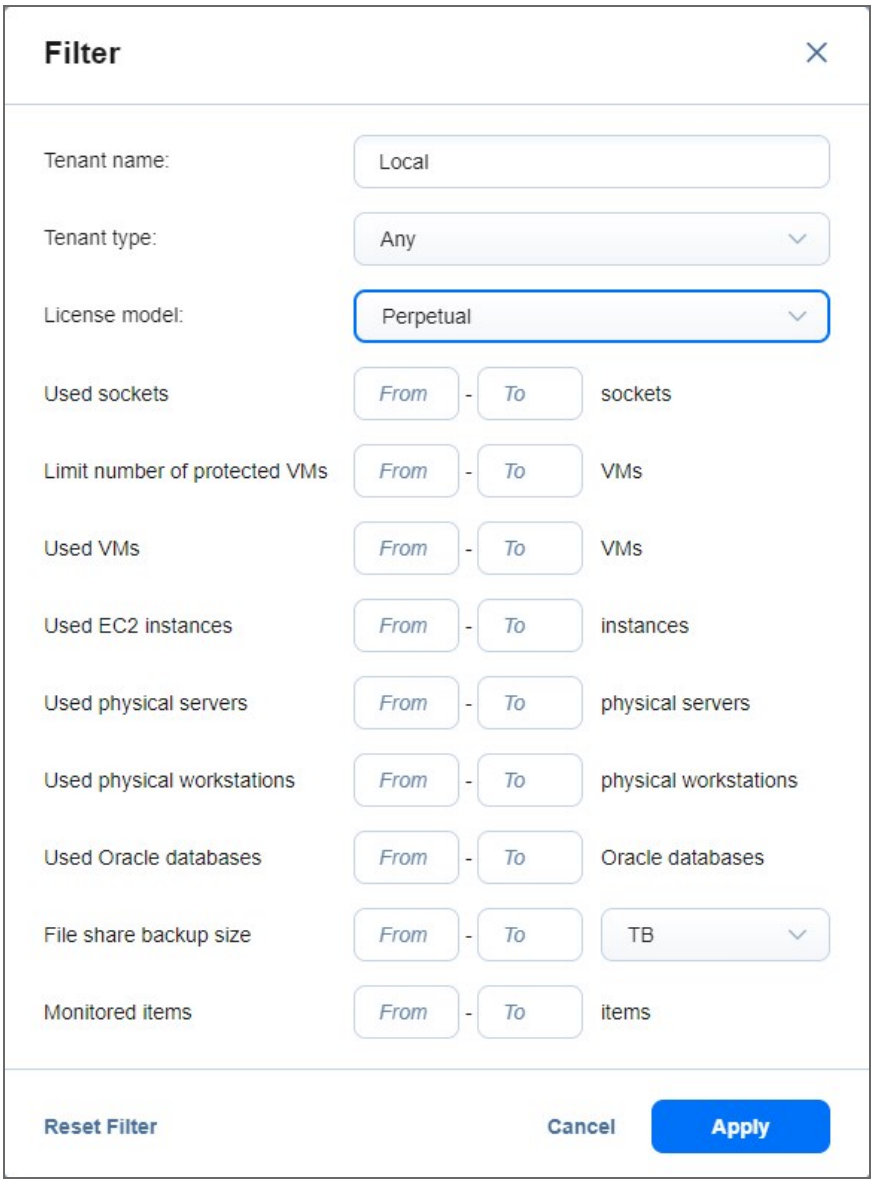
- **Tenant name**: Allows you to filter by tenant name.
- **Tenant type**: Allows you to filter by the following options:
  - **Any** (default): Both **Local** and **Remote** tenant types
  - **Local**
  - **Remote**
- **License model**: Allows you to filter by the following options:
  - **Any** (default): If selected, no extra fields are displayed.
  - **Subscription**: If selected, the following range input fields are displayed:
    - **Used workloads**
    - **Allocated workloads**
    - **Used Microsoft 365 users**
    - **Allocated Microsoft 365 users**



> **Note**
> The **Allocated workloads** and **Allocated Microsoft 365** users fields are not displayed if **Remote** is selected in the **Type** dropdown.

  - **Perpetual**: If selected, the following range input fields are displayed:
    - **Used sockets**
    - **Limit number of protected VMs**

- **Used VMs**
- **Used EC2 instances**
- **Used physical servers**
- **Used physical workstations**
- **Used Oracle databases**
- **File share backup size** with two dropdown options: **TB** (default) and **GB**
- **Monitored items**



Note

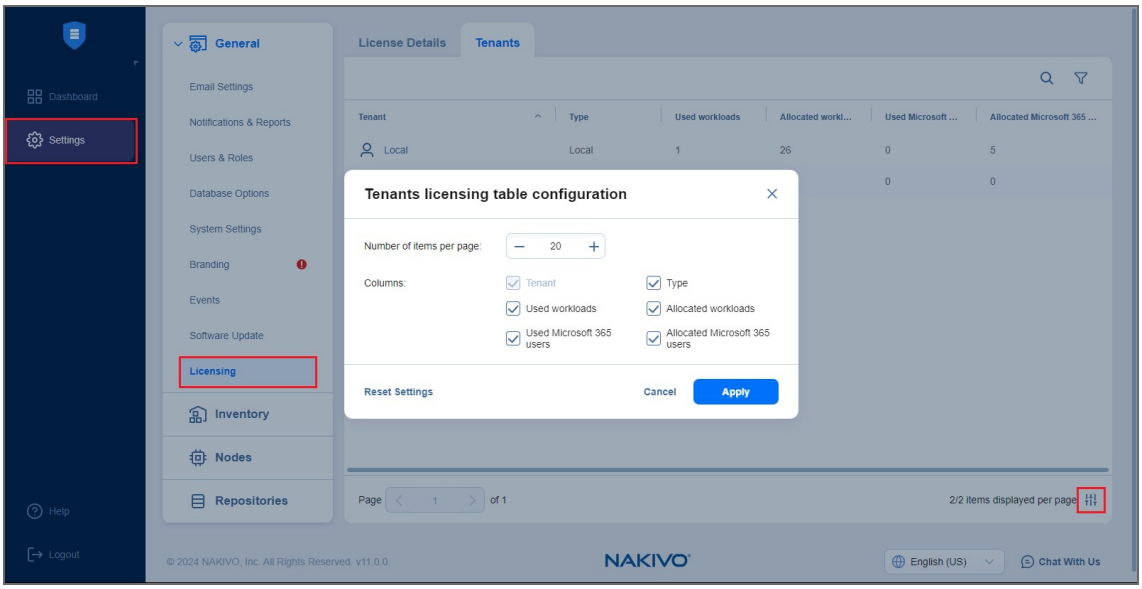
All range inputs can accept numeric values from *0* to *9999*.

Click **Apply** to close the dialog box and apply the filtering.

Click **Reset Filter** to reset the fields in the **Filter** dialog box to the default.

Click the **Cancel** or **X** button to close the dialog box without applying any changes.

Optionally, you can show/hide columns or modify the number of items per page in the **Tenants licensing table configuration** table. In the lower right corner, click the controls icon. In the dialog window that opens, select/deselect checkboxes and click **Apply**.



Click **Reset Settings** to reset the configuration settings to the default.

Click the **Cancel** or **X** button to close the dialog box without applying any changes.

Support Bundles

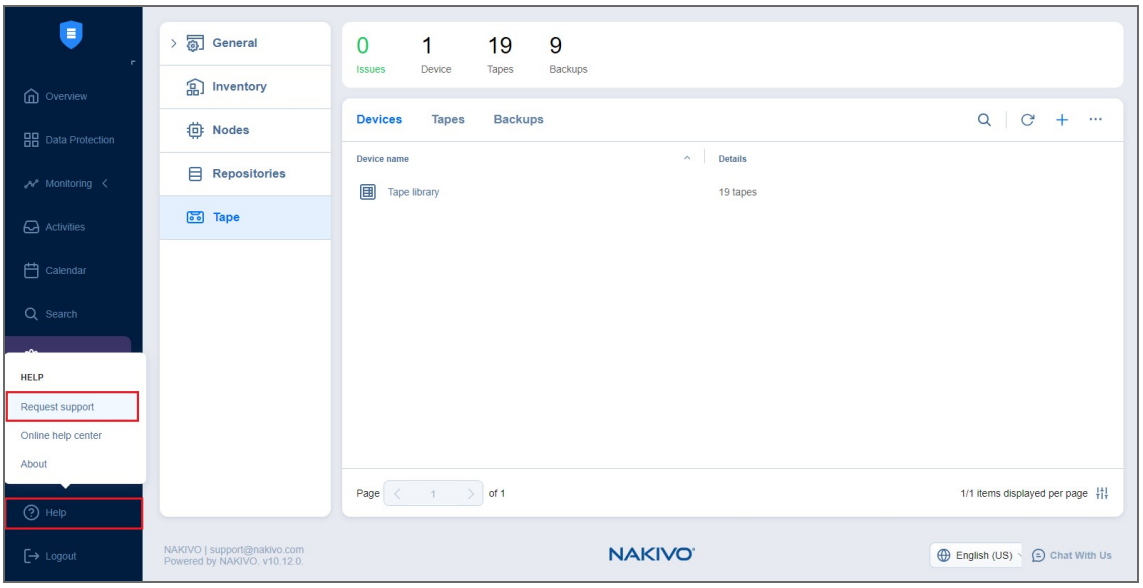
NAKIVO Backup & Replication provides you with the ability to create support bundles – a zipped collection of the product logs and system information. Sending a support bundle to the NAKIVO Support Team allows them to quickly identify the root cause of issues and suggest a proper solution.

- Creating Support Bundles
- Sending Support Bundles

Creating Support Bundles

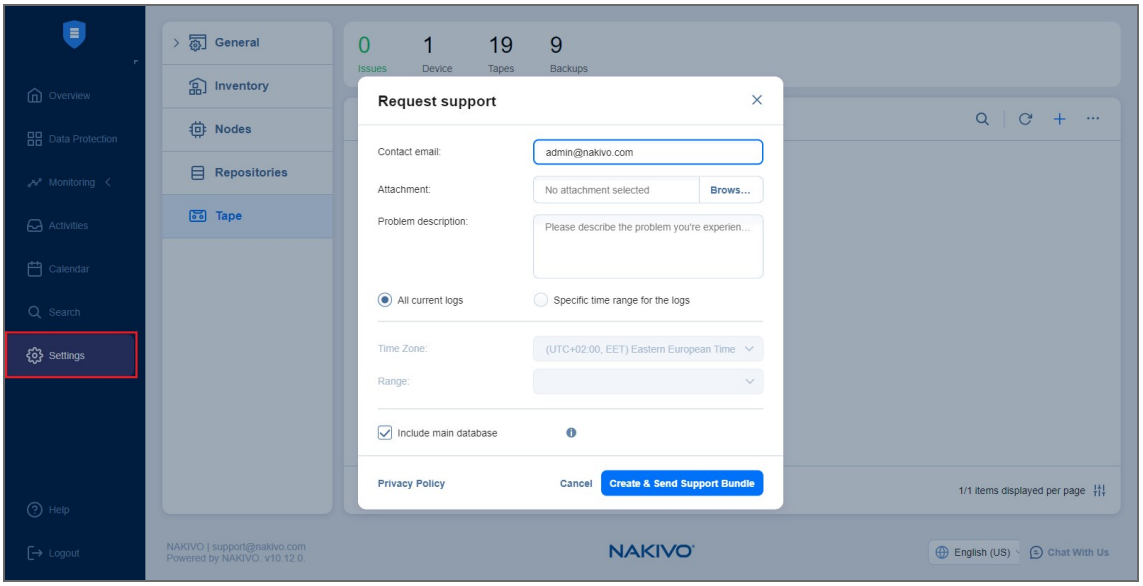
Before creating a support bundle, make sure Email settings are configured. To create a support bundle, follow these steps:

1. Click the **"?" (Help)** icon in the lower-left corner of the web UI.
2. Select and click **Request support**. The dialog box will appear.



3. Enter a description of your problem in the **Please describe the problem you're experiencing** box.
4. Enter your email address in the **Contact email** box.
5. If necessary, upload an attachment by clicking **Browse**.
6. Select **All current logs** if you wish to include all current log files to the support bundle.
7. Select **Specific time range for the logs** if you want to include the log files within the specific time range.
8. To specify the time range, click the **Range** dropdown and use the **From** and **To** selectors to set the time range; click **Apply**.
9. Optionally, click **Include main database** (recommended).

10. Click **Create & Send Support Bundle** to send the support bundle to NAKIVO Support Team. You will receive an answer from the NAKIVO Support Team within one business day.



Sending Support Bundles Manually

Some support bundles may become overly large in size. This can occur due to large log files or file dumps. In such cases, it is recommended to upload these files manually.

To do this, follow these steps:

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

> **Note**
> You can upload any files relevant to your issue: logs, file dumps, or the support bundles that you have manually downloaded from the product's UI.

3. Enter your email address in the **Contact email** field.
4. You can also enter the ID of your support ticket in the **Ticket ID** field if you have one opened.
5. Optionally, enter a description in the **Description** field.
6. Click **Upload** when you're done uploading the file(s).

> **Note**
> Wait for a successful upload notification before closing the page.

Built-in Support Chat

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

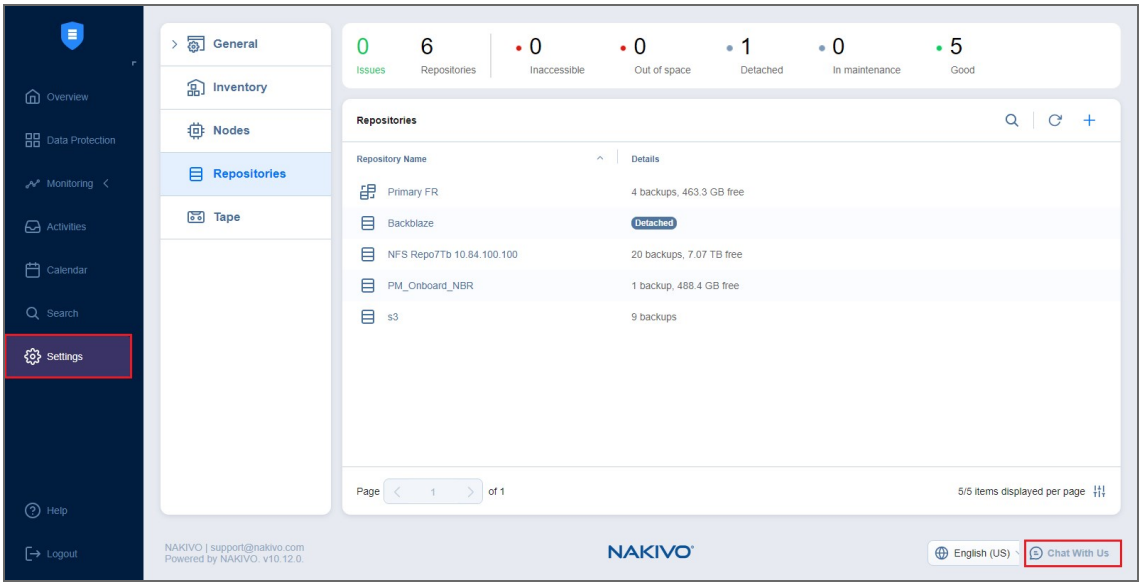
The built-in **Nakivo Support** chat is an interface where you can get live professional support. For a detailed explanation of each component, see the sections below:

- Opening Built-in Support Chat
- Sending Files in Built-in Support Chat
- Sending Email Transcript of Built-in Support Chat
- Editing Contact Details
- Disabling/Enabling Sound Notifications
- Disabling Built-in Support Chat
- Sending Feedback to Built-in Support Chat

Opening Built-in Support Chat

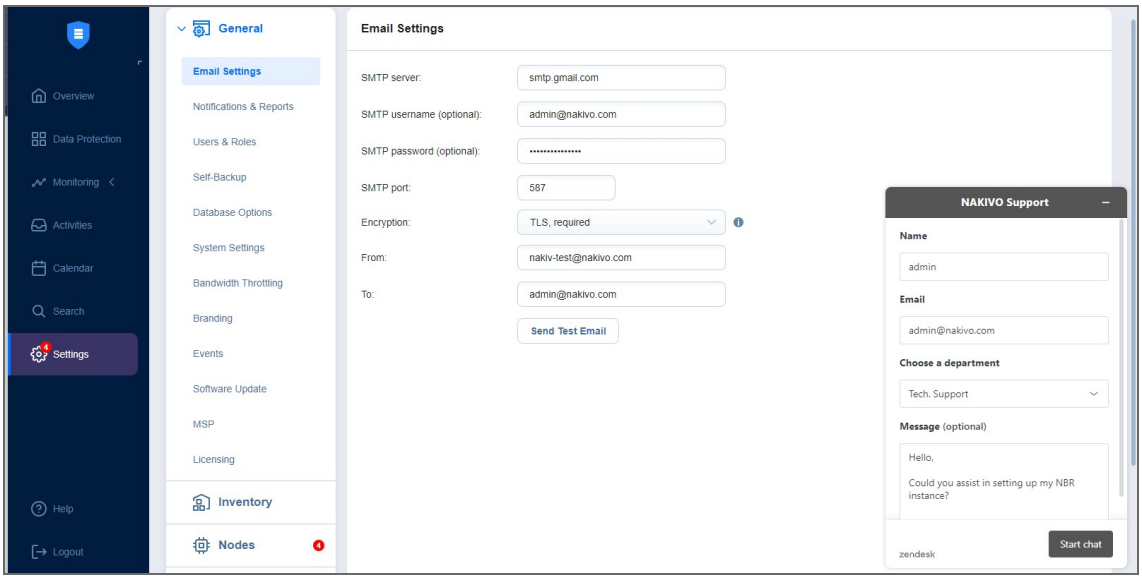
To open the built-in **Live Support Chat**, follow the steps below:

1. In the lower right corner of the NAKIVO Backup & Replication interface, click the **Chat With Us** button.

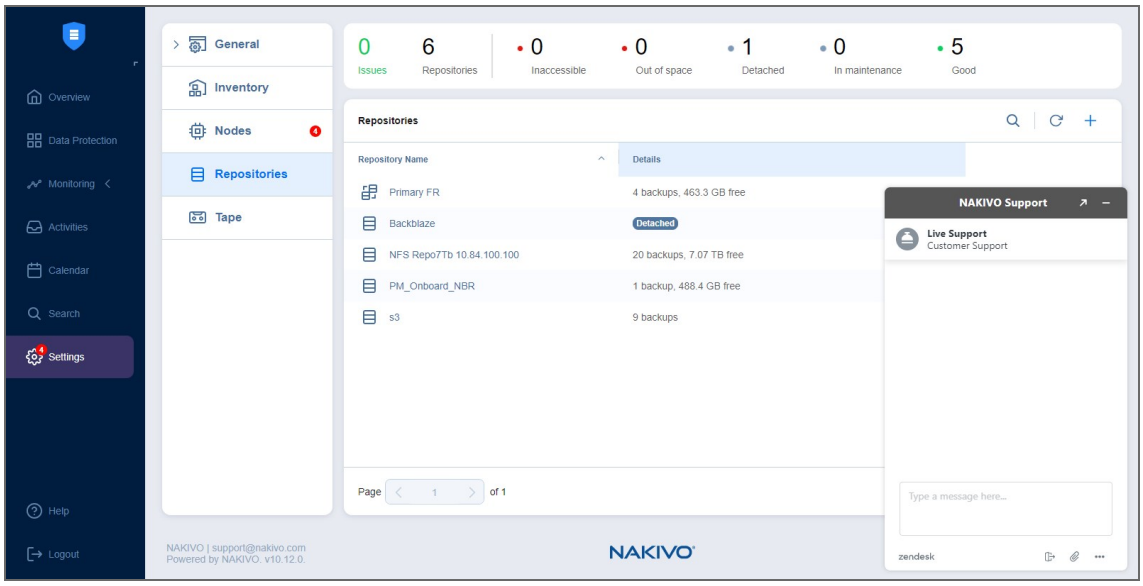


2. The **NAKIVO Support** dialog box opens. Introduce yourself by providing the following information:
   a. In the upper box of the dialog box, enter your name.
   b. In the box below, enter your email address.
3. Choose a department from the list of available departments.

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

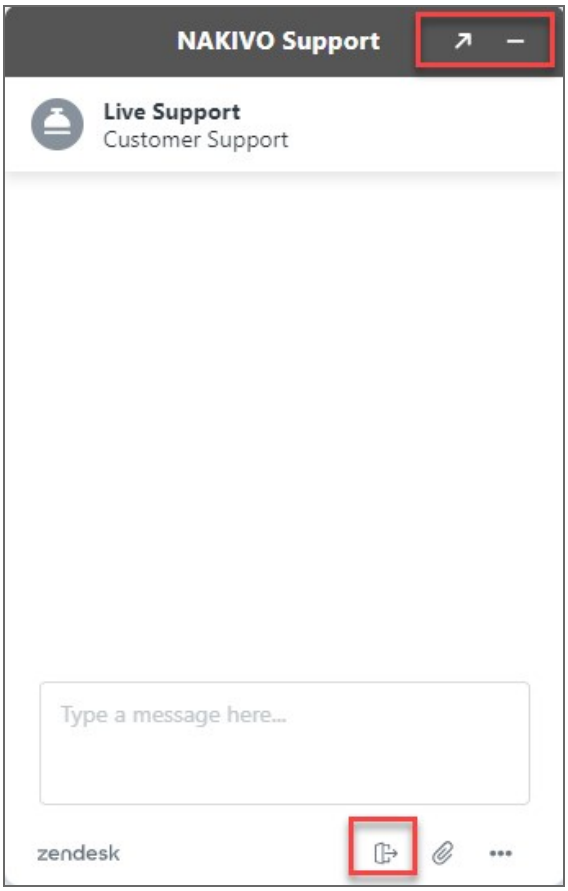


5. Your message is sent to a NAKIVO representative and the built-in **Live Support Chat** dialog box opens.



6. Optionally, you can open the **Web Widget Live Chat** by clicking the **Up right arrow** ↗ button in the chat header.

7. To leave the chat, click the minus (-)button in the chat header or click the **Exit** button in the lower right corner of the built-in **Live Support Chat** dialog box.

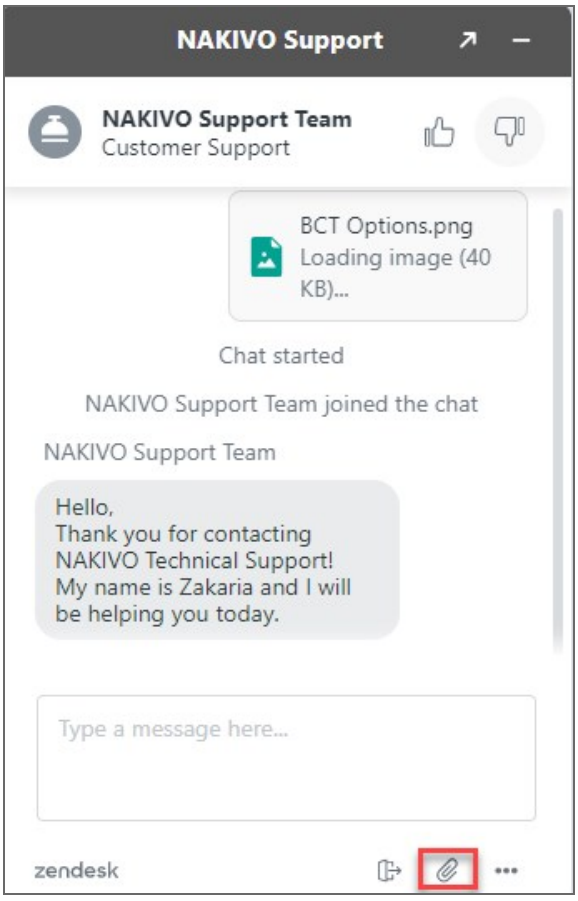


Sending Files in Built-in Support Chat

Please use either of the following ways to send your files in the built-in **Live Support Chat**:

- Drag and drop: open **Windows File Explorer**, select necessary files, and then drag them and drop to the chat dialog.
- The built-in **Live Support Chat** interface:
  1. In the lower right corner of the built-in **Live Support Chat** dialog box, click the **Attach** button.
  2. In the window that opens, select the needed file to attach to the chat.
  3. If needed, you can provide additional details in the message field.

4. Your message is sent to a NAKIVO representative and will be processed as soon as possible.



Note

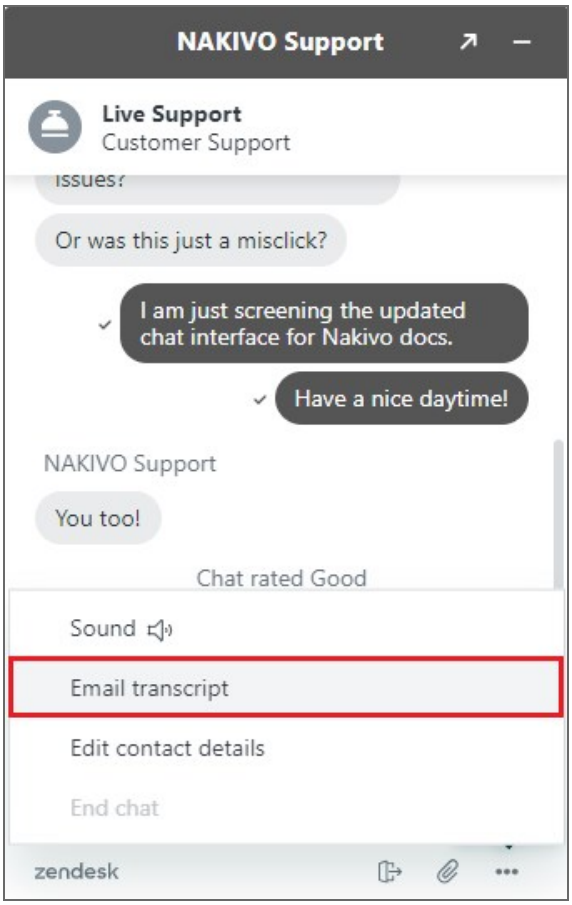
The following file formats are allowed: `.pdf, .png, .jpeg, .gif, .txt`. The maximum file size is 20 MB.

Sending Email Transcript of Built-in Support Chat

Follow the steps below to send the transcript of your built-in **Live Support Chat** session:

1. In the lower right corner of the built-in **Live Support Chat** dialog, click the **...** (ellipsis) button.
2. In the dialog box that opens, click **Email transcript**.
3. In the dialog box that opens, make sure the email address of the recipient is correct, and then click **Send**.

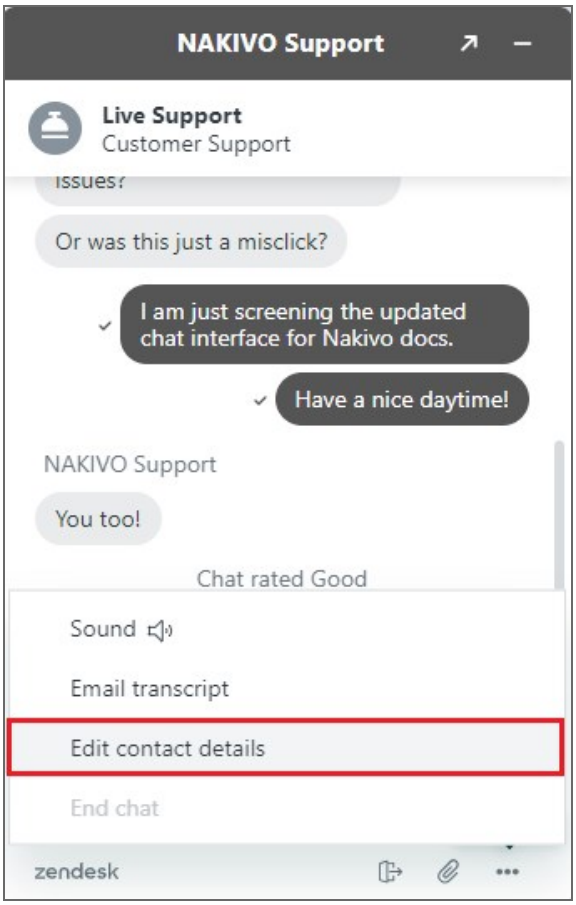
Your built-in **Live Support Chat** transcript will be sent to the specified email recipient.



Editing Contact Details

Follow the steps below to edit the contact details of your built-in **Live Support Chat** session:

1. In the lower right corner of the built-in **Live Support Chat** dialog, click the **...** (ellipsis) button.
2. In the dialog box that opens, click **Edit contact details**.
3. In the dialog box that opens, edit the contact name and the email address, and then click **Save**.



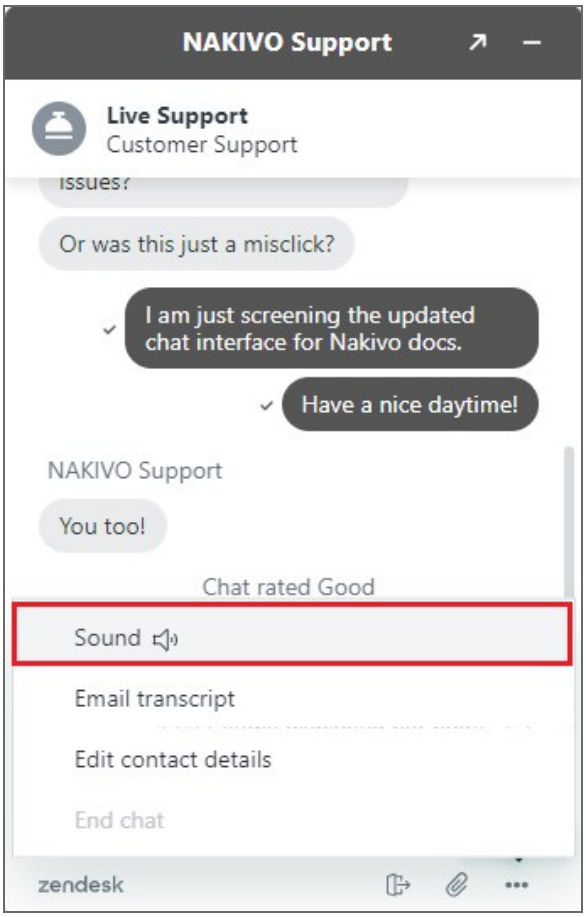
Disabling/Enabling Sound Notifications

By default, sound notifications are enabled for the built-in **Live Support Chat**.

Do the following to disable sound notifications:

1. In the lower right corner of the built-in **Live Support Chat** dialog, click the **...** (ellipsis) button.
2. In the dialog box that opens, click **Sound**.
3. Close the options dialog.

Sound notifications will be disabled for the built-in **Live Support Chat**.

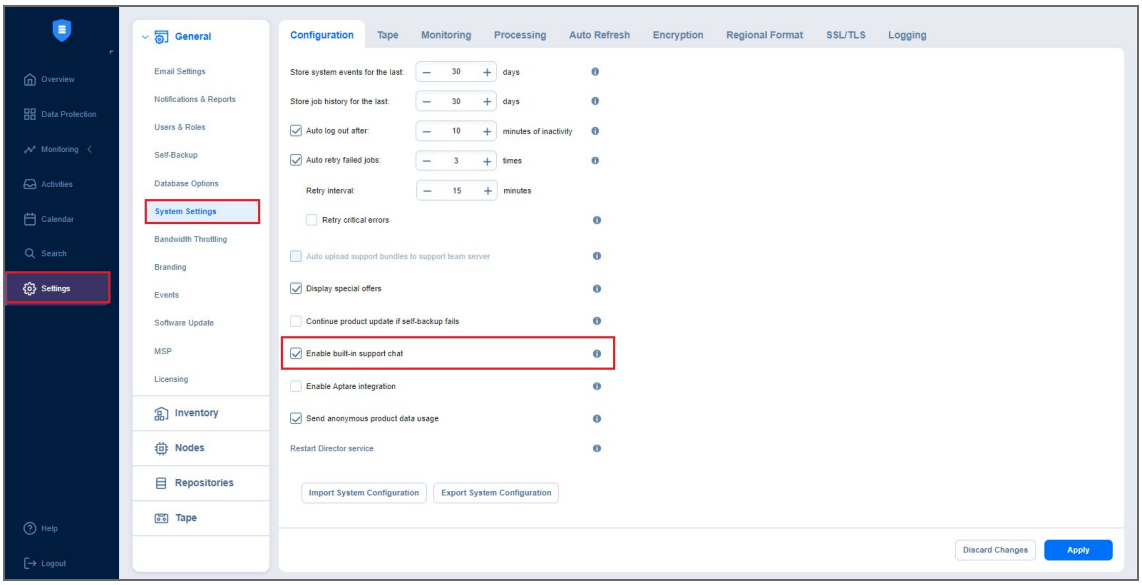


Disabling Built-in Support Chat

By default, the built-in support chat is enabled in your instance of NAKIVO Backup & Replication.

Do the following to disable the built-in **Live Support Chat**:

1. Go to **Settings** > **General** > **System Settings**.
2. Deselect the **Enable built-in support chat** checkbox.
3. Click the **Apply** button.



Note

When disabled, the **Built-in Support Chat** is not available in all tenants of the NAKIVO Backup & Replication instance in the Multi-tenant mode.

Sending Feedback to Built-in Support Chat

You have the possibility of sending feedback to built-in **Live Support Chat**: in the upper right corner of the dialog, click **Good** or **Bad**, as you deem appropriate.

If appropriate, leave a comment for the NAKIVO Support Team: click **Leave a comment** and in the text box that opens, enter your comment about the chat service. Then click **Send**.

Backup

This section contains the following topics:

- "Creating Microsoft 365 Backup Jobs" below
- "Deleting Backups" on page 614

Creating Microsoft 365 Backup Jobs

With NAKIVO Backup & Replication, you can back up an entire Microsoft 365 account or individual mailboxes, drives, sites, and Teams within that account. The product also allows you to send Microsoft 365 backups to tape for long-term data archiving and then recover these backup copies to a target backup repository. For details, refer to Backup Copy to Tape Recovery.

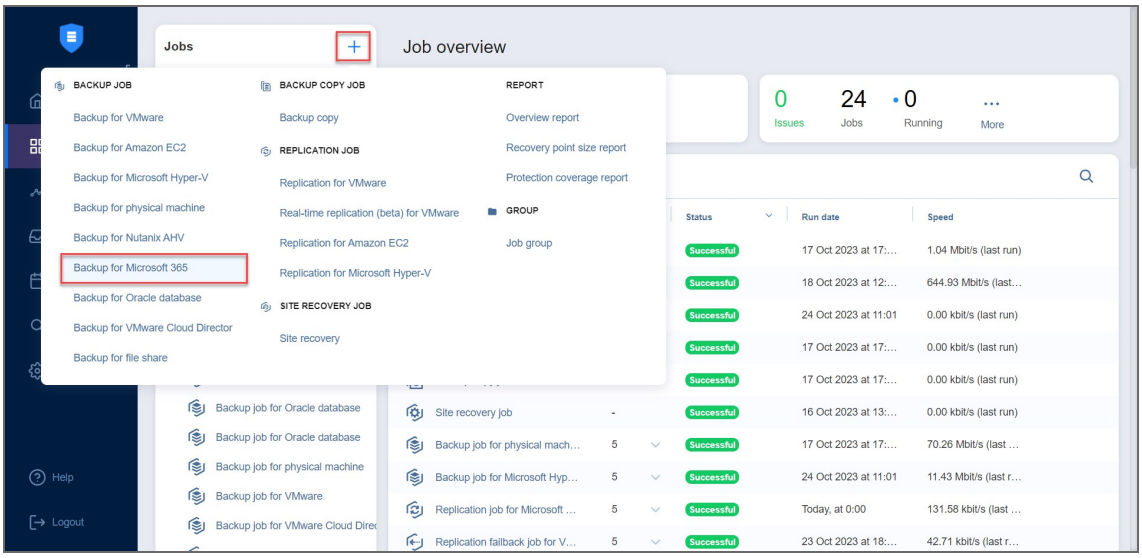
When creating a backup job for your Microsoft 365 account(s), you can specify which items to back up, where to store the backups, how often the backup job is run, and other backup options.

To create a backup job, click the plus **+** button in the **Data Protection** menu, and then click **Backup for Microsoft 365**.

Notes

- Before creating a Microsoft 365 backup job, you must add a SaaS Backup Repository or Incremental with Full repository.

  Refer to Required API Permissions for Microsoft 365 to see the list of required permissions for backing up Microsoft 365 objects.
- Refer to Microsoft 365 Platform Limitations of the latest Release Notes for the full list of Microsoft 365 items that are not supported.
- For mailbox backups, the size displayed in the Transferred Raw Data widget shows the amount of data that was actually transferred during backup. This size may differ from the total mailbox size displayed in the Job Info widget. Refer to "Backup Job Wizard for Microsoft 365: Source" on the next page for a list of items that are skipped during the backup process.



The **New Backup Job Wizard for Microsoft 365** opens. Complete the wizard as described in the sections below:

- “Backup Job Wizard for Microsoft 365: Source” below
- “Backup Job Wizard for Microsoft 365: Destination” on page 592
- “Backup Job Wizard for Microsoft 365: Schedule” on page 594
- “Backup Job Wizard for Microsoft 365: Retention” on page 604
- “Backup Job Wizard for Microsoft 365: Options” on page 605

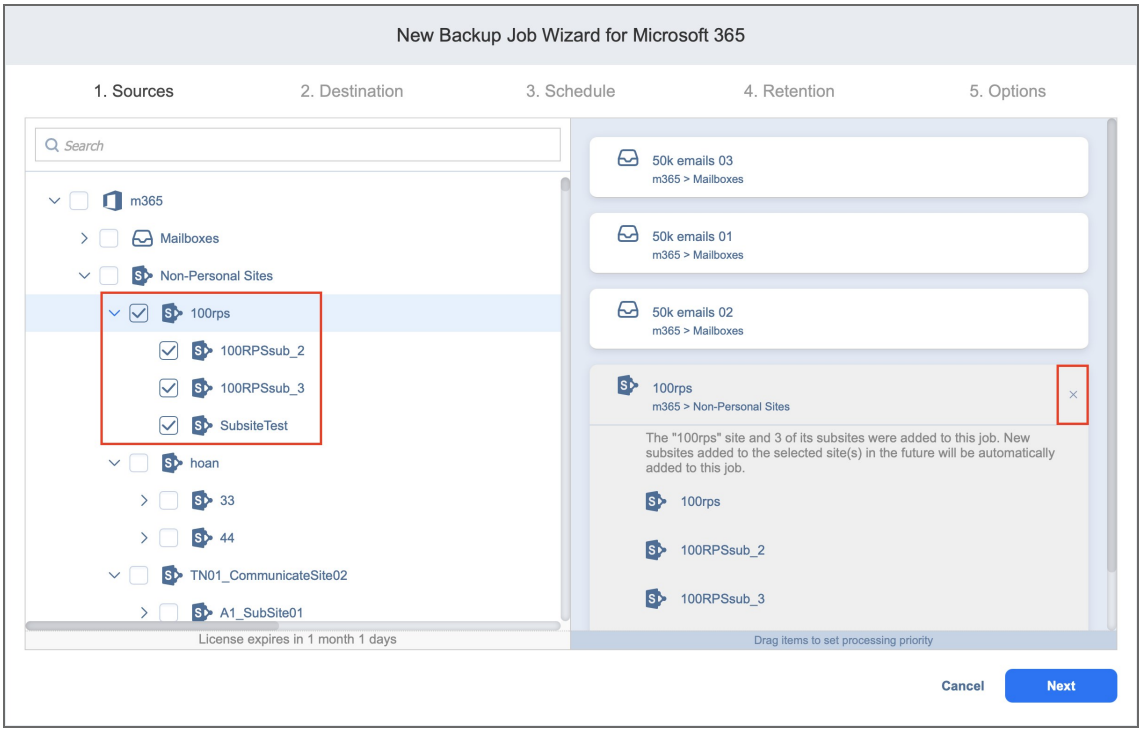
Backup Job Wizard for Microsoft 365: Source

On the **Source** page of the wizard, you can add Microsoft 365 account(s), user, shared, and group mailboxes, user drives, personal, communication, and team sites, and public, private, and organization-wide teams to your backup job. Proceed as follows:

1. In the left pane of the page, select the items you want to back up.
2. To quickly find an item, use the **Search** functionality. You can enter a part of or the full name of the item. Selected items are displayed in the right pane. You can remove items from the pane if necessary or change the priority of the item by dragging and dropping it in the required position. The priority determines the order in which the item will be processed during the job run.
3. Click **Next** to confirm the selection and go to the next page of the wizard.

Notes

- If you select a Microsoft 365 account, all mailboxes, SharePoint sites, OneDrives, and Teams contained in that account are added to the backup job.
- If you select a group mailbox, only the group mailbox data is backed up. This does not include the data of the group members, such as user mailbox, user OneDrive, or personal site.
- If you select a group site, only the group site data is backed up. This does not include the data of the group members, such as user mailbox, user OneDrive, or personal site.
- The following mailbox items are **not** backed up:
    - Outbox folders
    - Calendar event messages
    - Event request messages
    - ReadOnly folders in calendars
- Backing up a Team will **not** back up the connected team members' mailboxes, OneDrives, or personal sites. Teams backups include the following:
    - Team metadata, including team settings, team members roles & permissions, channels, tabs, applications, etc.
    - Team document library
    - Messages from posts
- Only the following types of channel tabs are included in a Teams backup:
    - Website
    - Word
    - Excel
    - PowerPoint
    - Visio
    - PDF
    - Document Library
    - SharePoint
    - SharePoint Pages
    - SharePoint Lists



Backup Job Wizard for Microsoft 365: Destination

On the **Destination** page of the wizard, you can specify the storage location for the backup. You can specify a single location for all items in the backup or select different Backup Repositories for different items.

You can select repositories in the following locations to back up your Microsoft 365 items:

- Amazon S3
- Generic-S3 compatible storage
- Wasabi
- Backblaze B2
- Azure Blob
- Local Folder
- CIFS Share
- NFS Share
- Amazon EC2
- Deduplication Appliance

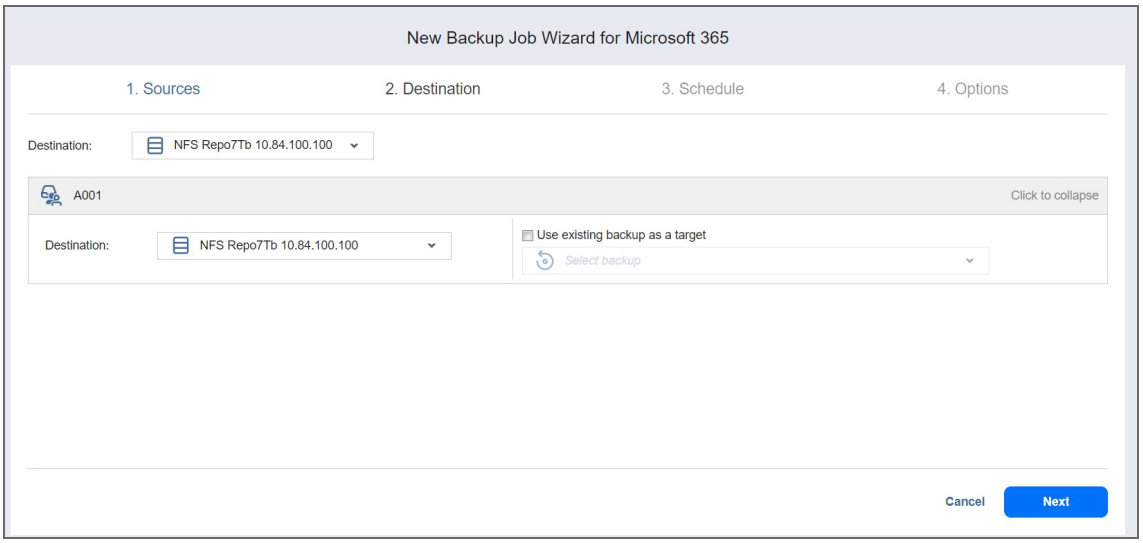
Notes

- The size of source items is not displayed in the dropdown menu.
- Only incremental with full repository types or the SaaS repository type can be selected as target locations.
- Stream repository types can be mixed the SaaS repository type within the same job.

- Setting a Single Backup Repository for All Items
- Setting a Different Backup Repository for Each Item
- Mapping Source Microsoft 365 Items to Existing Backups

Setting a Single Backup Repository for All Items

To back up all items selected on the **Sources** page to a single Backup Repository, choose a Backup Repository from the **Destination** drop-down list.



Both federated repositories and standalone repositories (those not used as members of federated repositories) can be selected as a destination.

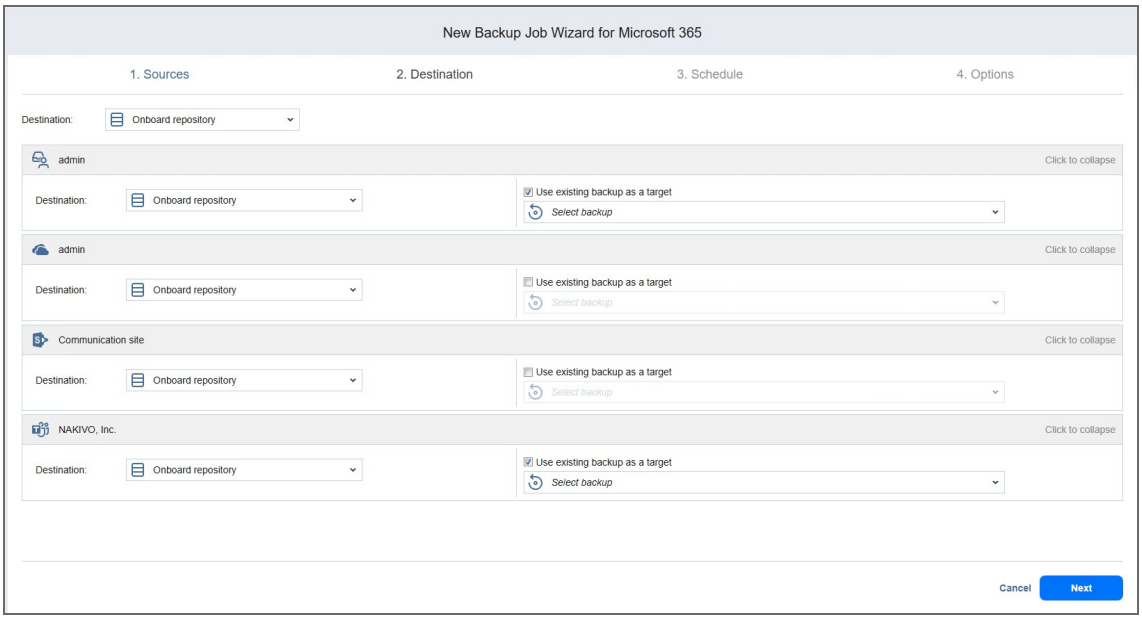
If a federated repository is selected as the destination, the product automatically selects an available qualified federated repository member during the job run.

Setting a Different Backup Repository for Each Item

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

1. Click **Advanced setup** and do one of the following:
   - If you have selected a Microsoft 365 account(s) on the **Source** page, the account block is displayed.
     a. Click the Microsoft 365 name to expand it and view all of the mailboxes, drives, sites, and teams.
     b. In the **Default Destination** drop-down list, select the Backup Repository for storing the backups of all services within the Microsoft account(s).
   - If you need a specific mailbox(es), drive, site, or team to be stored in a different Backup Repository, click the name of the service and select a different location from the **Destination** drop-down list.

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



Mapping Source Microsoft 365 Items to Existing Backups

If you have previously backed up Microsoft 365 items and then lost the backup job due to accidental job deletion or a need to recreate jobs in a new copy of the product, you can map source Microsoft 365 items to existing backups in order to avoid running full backups again.

To map source Microsoft 365 items to existing backups, follow the steps below:

1. Click **Advanced setup**.
2. From the **Backup repository** drop-down list, choose a Backup Repository that contains an existing Microsoft 365 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 Microsoft 365 items you have selected, determine how it is different from the source items, and transfer only the differential data.

Backup mapping can be a time-consuming process that can be equal to the time required to create a full backup.

After the job is completed, a new recovery point will be created and existing recovery points will not be changed or overwritten.

Backup Job Wizard for Microsoft 365: Schedule

On the **Schedule** page of the wizard, select to run the backup job manually or schedule the job to run on a regular basis.

Proceed as described in the sections below:

- Switching to the improved retention approach
- Creating New Schedules
  - Weekly
  - Monthly
  - Yearly
  - Periodical
  - After another job
- Creating Legacy Schedules
  - Daily or Weekly Backup
  - Monthly or Yearly Backup
  - Periodic Backup
  - Chained Job

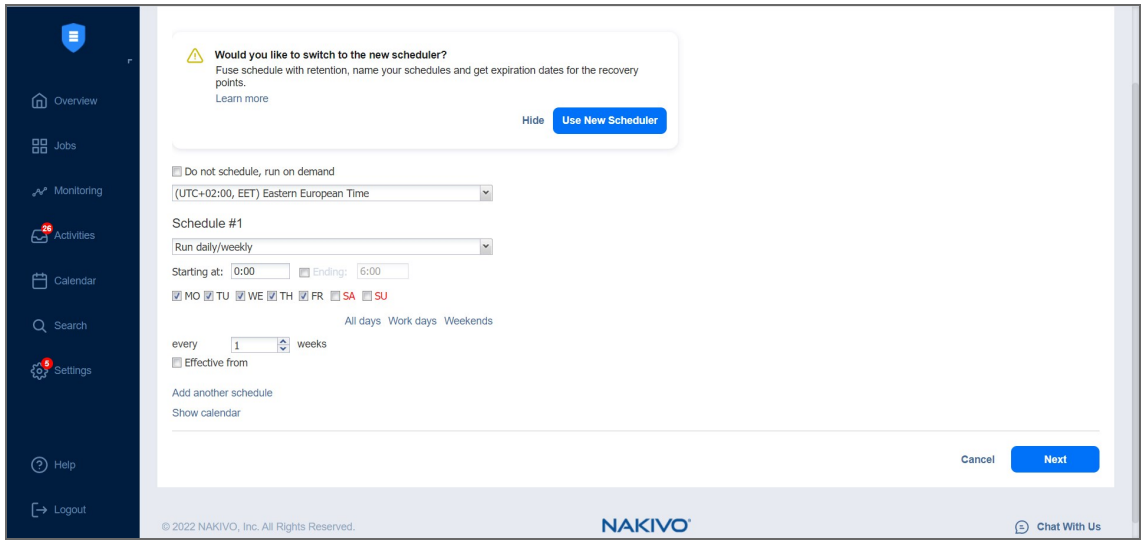
Switching to Improved Retention Approach

NAKIVO Backup & Replication offers two approaches to retention and scheduling: the legacy or the improved approach. To learn more about how the legacy and improved approaches work, go here. If you create a new job or edit the existing one that uses the legacy approach, a popup appears offering that you to switch to the improved retention approach in the following cases:

- You have updated your instance of the product to v10.8 or later from an older version.
- You have imported a configuration to an instance of NAKIVO Backup & Replication v10.8 or later from an older version.

Note

If you install NAKIVO Backup & Replication v10.8 or higher, the new approach is enabled by default.



After the popup appears, do one of the following actions:

- If you do not want to switch to the new scheduler, click **Hide** to close the popup. You can later click **Use New Scheduler** on the **Schedule** page to proceed with the change if you change your mind.
- Alternatively, click **Use New Scheduler** in the popup. Next, choose one of the following options:
  - **MIGRATE SETTINGS**: When you select this option, the existing schedules are automatically converted to new schedules and the existing retention settings are mapped to the new schedules.
  - **CREATE NEW SCHEDULES**: When you select this option, you can create new schedules using the existing retention settings. Old schedules will be deleted.
  - **CONFIGURE SETTINGS ANEW**: Select this option to reset all existing schedules and retention settings and configure them from scratch.

Notes

- After switching to the new scheduler, the legacy schedule and retention settings are displayed on the right side of the page.
- After switching to the new scheduler, reverting to the legacy schedule and retention settings is impossible.
- You can learn how expiration dates are assigned to recovery points after migration here.

Creating New Schedules

Before creating a new schedule, you can optionally enable the following settings:

- **Do not schedule, run on demand**: Enable this option if you want to start the job manually.
- **Prioritize schedules**: When this option is selected, NAKIVO Backup & Replication starts treating schedules based on their priority. The **Yearly** schedule will have higher priority than the **Monthly** schedule, etc. In case 2 or more schedules overlap, the schedules with lower priority will be skipped.
- Optionally, when creating any type of schedule, click **Show Calendar** to show the calendar or **Hide Calendar** to hide it.

When **Amazon EC2**, **Amazon S3**, **Wasabi**, **Azure Blob Storage**, **Backblaze B2 Cloud Storage**, or **Local Folder** is selected as the Backup Repository Type and the only backup destination, you can make recovery points in these repositories immutable during schedule creation. With immutability enabled, the recovery points are immutable and stored using the write-once-read-many (WORM) model. With **Amazon EC2**, **Amazon S3**, **Wasabi**, **Azure Blob Storage**, or **Backblaze B2 Cloud Storage** types of Backup Repository, immutable recovery points cannot be overwritten, deleted, or changed by the root user until the specified period has expired. For Local Folder type of Backup Repository, the root user can still clear immutability.

Notes

For the Immutability section to be available, the following conditions must be met:

- **Amazon EC2**, **Amazon S3**, **Wasabi**, **Azure Blob Storage**, **Backblaze B2 Cloud Storage**, or **Local Folder** must be selected for Backup Repository Type on the Destination page of the wizard.
- When **Amazon EC2**, **Amazon S3**, **Wasabi**, **Azure Blob Storage**, or **Backblaze B2 Cloud Storage** is selected as the Backup Repository Type, the bucket or blob container with the repository must have Object Lock or version-level immutability enabled respectively as well as Versioning.
- For **Local Folder** type of Backup Repository, see feature requirements.

When creating the schedules, you can create schedules of the following types:

Weekly

You can configure the following options for this schedule type:

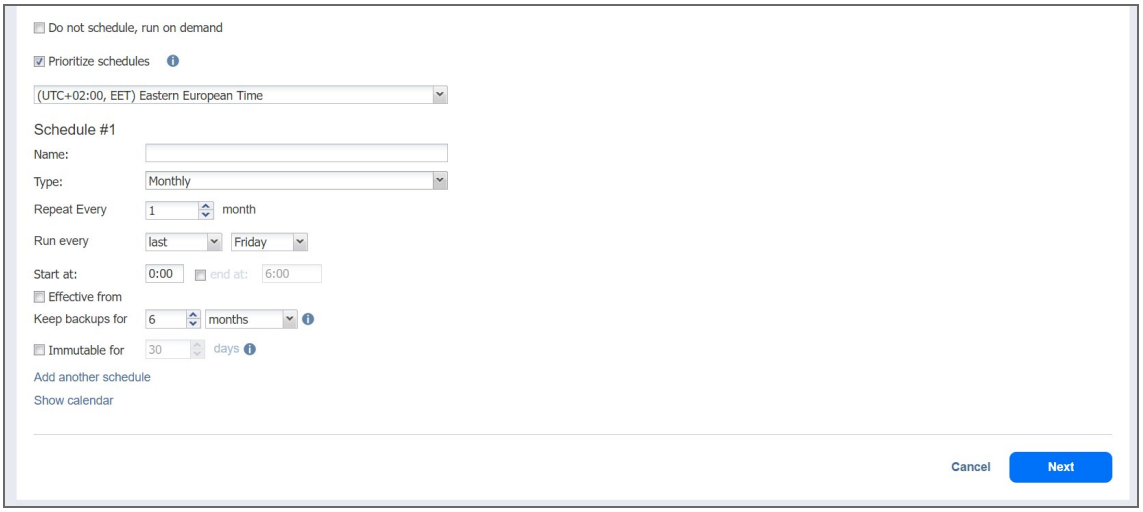
- **Schedule title**: Designates the number of a schedule.
- **Name**: Enter the name of your schedule.
- **Repeat every X weeks**: Indicates how often the schedule is repeated.
- **Days**: Select specific days when the schedule executes the job.
- **Start at**: Specify the time when the job should start.
- **End at**: Specify the time when the job should end.
- Optionally, select the **Effective from** checkbox and choose the date when the schedule should come into effect.
- **Keep backups for**: Specify how many days, months, or years NAKIVO Backup & Replication should retain the backups.
- Optionally, click **Add another schedule** if you want to add more than one schedule.



Monthly

You can configure the following options for this schedule type:

- **Schedule title**: Designates the number of a schedule.
- **Name**: Enter the name of your schedule.
- **Repeat every X months**: Indicates how often the schedule is repeated.
- **Run every**: Select specific days of the month when NAKIVO Backup & Replication executes the job.
- **Start at**: Specify the time when the job should start.
- **End at**: Specify the time when the job should end.
- Optionally, select the **Effective from** checkbox and choose the date when the schedule should come into effect.
- **Keep backups for**: Specify how many days, months, or years NAKIVO Backup & Replication should keep the backups.
- Optionally, click **Add another schedule** if you want to add more than one schedule.

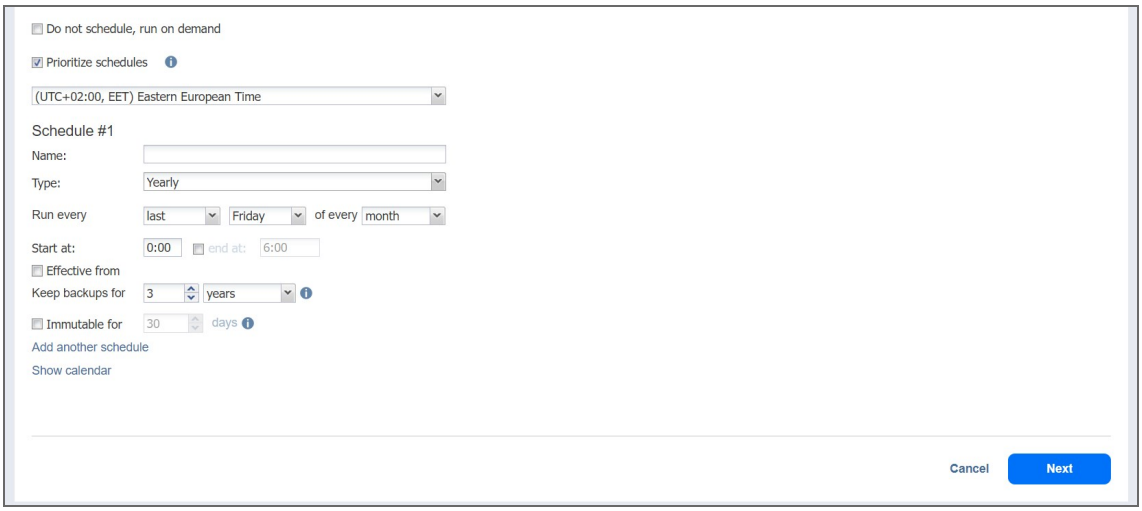


Yearly

You can configure the following options for this schedule type:

- **Schedule title**: Designates the number of a schedule.
- **Name**: Enter the name of your schedule.
- **Run every**: Select specific days of the specific month when NAKIVO Backup & Replication executes the job.
- Optionally, select the **Effective from** checkbox and choose the date when the schedule should come into effect.
- **Keep backups for**: Specify how many days, months, or years NAKIVO Backup & Replication should keep the backups.

- Optionally, click **Add another schedule** if you want to add more than one schedule.



Periodical

You can configure the following options for this schedule type:

- **Schedule title**: Designates the number of a schedule.
- **Name**: Enter the name of your schedule.
- **Run every**: Select the period measured in minutes, hours, or days when NAKIVO Backup & Replication executes the job.
- **Start at**: Specify the time when the job should start.
- **End at**: Specify the time when the job should end.
- **Days**: Select specific days when the schedule executes the job.
- Optionally, select the **Effective from** checkbox and choose the date when the schedule should come into effect.
- **Keep backups for**: Specify how many days, months, or years NAKIVO Backup & Replication should keep the backups.
- Optionally, click **Add another schedule** if you want to add more than one schedule.
- Optionally, when creating any type of schedule, click **Show Calendar** to show the calendar or **Hide**

Calendar to hide it.



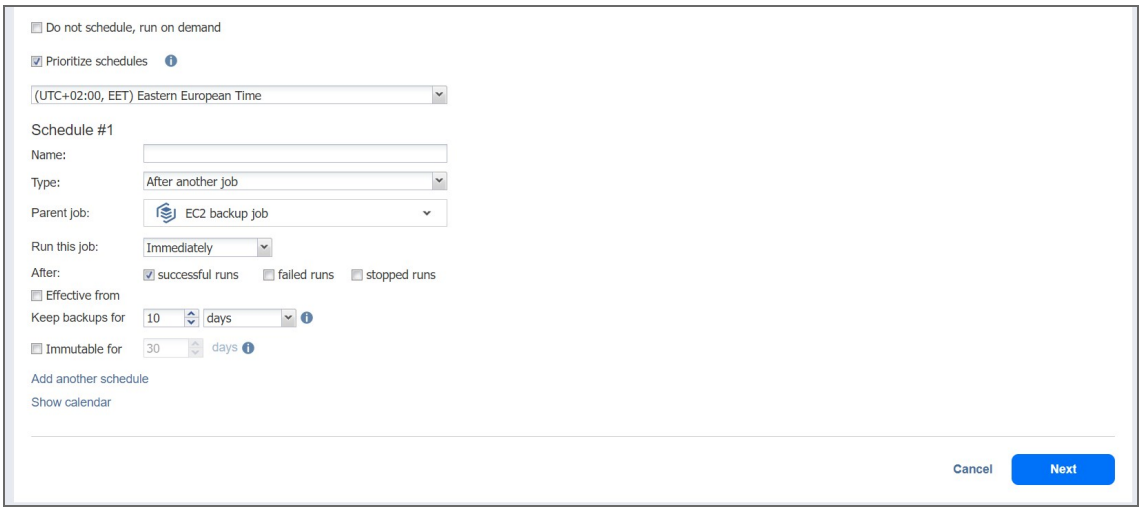
After Another Job

You can configure the following options for this schedule type:

> **Note**
> This option is disabled if there are no other jobs.

- **Schedule title**: Designates the number of a schedule.
- **Name**: Enter the name of your schedule.
- **Parent job**: Select the job after which this job starts running.
- **Run this job**: Select one of the following options:
  - **Immediately**: The schedule starts right after the parent job is completed.
  - **Delayed**: The schedule starts after the specified number of **minutes** or **hours** following parent job completion.
- Optionally, select the **Effective from** checkbox and choose the date when the schedule should come into effect.
- **Keep backups for**: Specify how many days, months, or years NAKIVO Backup & Replication should keep the backups.
- Optionally, click **Add another schedule** if you want to add more than one schedule.

- Optionally, click **Show Calendar** to show the calendar or **Hide Calendar** to hide it.



Creating Legacy Schedules

Before creating a new schedule, you can optionally enable the following settings:

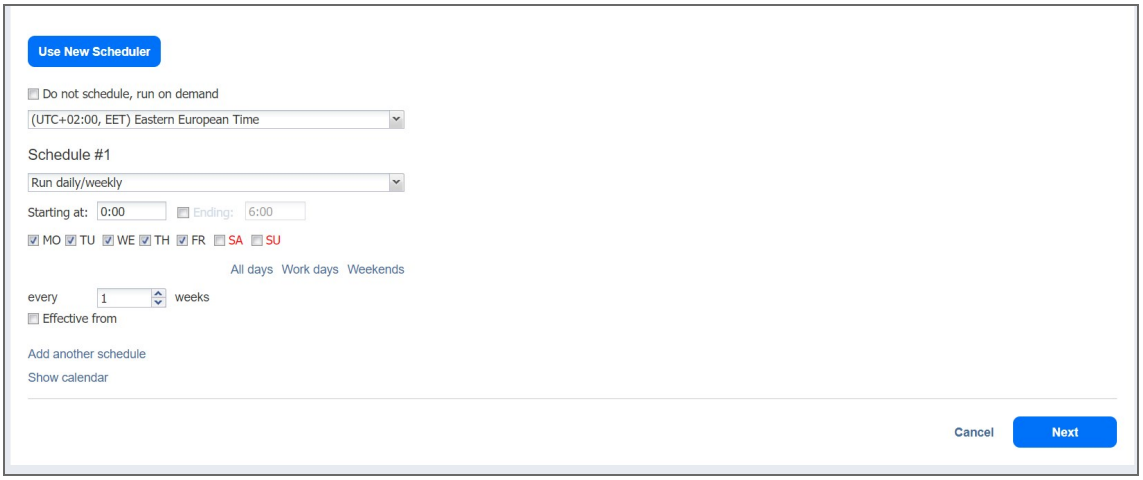
- Click **Use New Scheduler** to switch to the Improved retention approach.
- **Do not schedule, run on demand**: Enable this option if you want to start the job manually.
- Optionally, when creating any type of schedule, click **Show Calendar** to show the calendar or **Hide Calendar** to hide it.
- Optionally, click **Add another schedule** if you want to add more than one schedule.

Daily or Weekly Backup

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

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

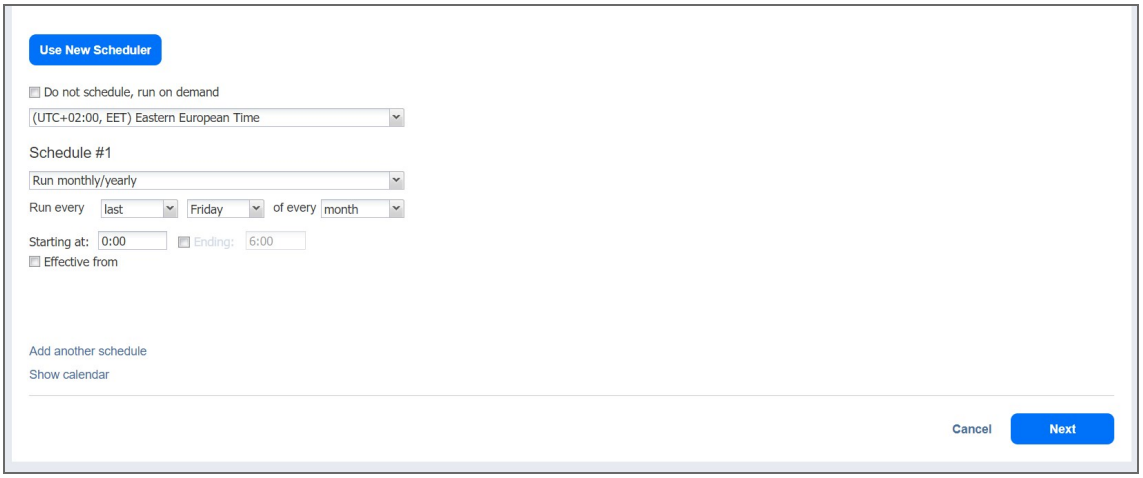
- If necessary, select the **Effective from** checkbox and pick the date when the schedule comes into effect.



Monthly or Yearly Backup

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

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



Periodic Backup

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

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



Chained Job

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

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



Backup Job Wizard for Microsoft 365: Retention

Important
This page is not displayed if the new scheduler is enabled.

After each job run, NAKIVO Backup & Replication creates a recovery point for each item in the Backup Repository. A recovery point represents the backed up Microsoft 365 account or mailboxes as of a particular moment in time and allows you to recover individual objects or the entire account from the Backup Repository. You can specify how many recovery points should be preserved in the Backup Repository using the Grandfather-Father-Son (GFS) backup rotation scheme.

When **Amazon EC2**, **Amazon S3**, **generic S3-compatible storage**, **Wasabi**, **Azure Blob Storage**, **Backblaze B2 Cloud Storage**, **Dell EMC DataDomain**, or **Local Folder** is selected as the Backup Repository Type for the only backup destination, you can make recovery points in these repositories immutable. With immutability enabled, the recovery points are immutable and stored using the *write-once-read-many* (WORM) model. In case of **Amazon EC2**, **Amazon S3**, **generic S3-compatible storage**, **Wasabi**, **Azure Blob Storage**, **Dell EMC DataDomain**, or **Backblaze B2 Cloud Storage** types of Backup Repository, immutable recovery points cannot be overwritten, deleted, or changed by the root user, until the specified period has expired. For **Local Folder** type of Backup Repository, the root user can still clear immutability.

Retention Settings

Here you can set the retention settings for the backup job. Set the following options:

- **Keep x last recovery points**: Retains the specified number of last recovery points for each item in the job.
- **Keep one recovery point per day for x days**: Retains one last recovery point per day for the specified number of days.
- **Keep one recovery point per week for x weeks**: Retains the last available backup of every week for the specified number of weeks.
- **Keep one recovery point per month for x months**: Retains the last available backup of every month for the specified number of months.
- **Keep one recovery point per year for x years**: Retains the last available backup of every year for the specified number of years.

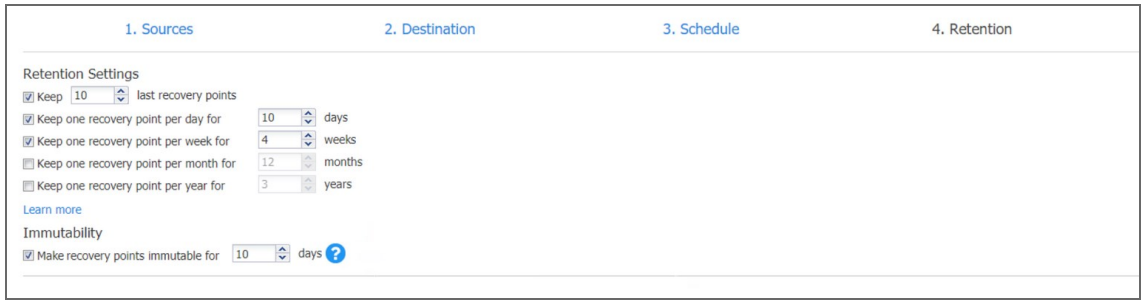
Immutability

In this section, you can configure the **Make recovery points immutable for x days** option. The recovery points remain immutable for the specified number of days.

Note

For the *Immutability* section to be available, the following conditions must be met:

- **Amazon EC2**, **Amazon S3**, **generic S3-compatible storage**, **Wasabi**, **Azure Blob Storage**, **Backblaze B2 Cloud Storage**, **Dell EMC DataDomain**, or **Local Folder** must be selected for Backup Repository Type on the Destination page of the wizard.
- If **Amazon EC2**, **Amazon S3**, **generic S3-compatible storage**, **Wasabi**, **Azure Blob Storage**, **Dell EMC DataDomain**, or **Backblaze B2 Cloud Storage** is selected as the Backup Repository type, Object Lock or version-level immutability support and Versioning must be enabled bucket or blob container respectively where your Backup Repository is located.
- For **Local Folder** type of Backup Repository, see feature requirements.



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

Backup Job Wizard for Microsoft 365: Options

On the **Options** page, you can specify the job's name, define pre- and post-job actions, and limit the Transporter load. Proceed as described in these sections:

- Job Options
- Mailbox Processing
- Full Backup Settings
- Pre and Post Job Actions
  - Email Notifications
  - Pre Job Script
  - Post Job Script
- Data Transfer
  - Transporter Load
  - Bottleneck Detection
- Completing the New Backup Job Wizard for Microsoft 365

Job Options

In this section, you can specify a name and a priority level for the backup job, change tracking, encryption, and other options. Proceed as described below:

- Enter a name for the job to the **Job name** field.
  - Select a **Job priority** level between 1 and 5, with 1 being the highest priority. Jobs with higher priority levels are prioritized by Transporters during job processing.

    Note
    This option is only available in the Enterprise, Enterprise Essentials, Enterprise Plus, MSP Enterprise, and MSP Enterprise Plus editions.

- Configure **Change tracking**. Do the following:
  - Choose from the following options in the dropdown menu:
    - **Use Microsoft change tracking**: With this option selected, after the initial job run, NAKIVO Backup & Replication uses Microsoft APIs to process exclusively the changed data from the last job run. This allows for a significant increase in processing speed.
    - **Use proprietary method**: With this option selected, NAKIVO Backup & Replication performs an incremental backup using the NAKIVO proprietary change tracking technology.
    - **No change tracking (always full)**: When change tracking is disabled, the full data set is transferred on every job run.
  - Optionally, click **settings** to select one of the following job behaviors in relation to the **Microsoft change tracking** error:
    - **switch to proprietary method**: If Microsoft change tracking fails to provide information on changed data for an object and this option is selected, NAKIVO Backup & Replication performs an incremental backup of the object using the NAKIVO proprietary change tracking technology.
    - **fail object processing**: If Microsoft change tracking fails to provide information on changed data for an object and this option is selected, NAKIVO Backup & Replication does not process the object and the job is failed.

Notes

- The **Change tracking** option is only displayed if at least one team is selected in the Sources step and its account has disabled Teams Export APIs.
- If the **Use Microsoft change tracking** option is selected, the product uses delta queries in Microsoft Graph API to process only the changed data from the last run for Teams chat messages.
- The following items are not supported and the proprietary method is used by default:
  - Contact lists
  - Mailboxes and its content
  - Public folder items
  - SharePoint Online sites and lists
  - OneNotes.
- OneDrive files (including group mailbox files) and document libraries use delta queries by default and no configuration in the product is necessary.
- The first backup run after switching to **Use Microsoft change tracking** may take longer to process, as it re-scans all source data using Microsoft change tracking.

- If the type of the Backup Repository that you've selected on the **Destination** page of the wizard is set to **Incremental with full backups** (**Store backups in separate files option** is selected), you can specify the following options:
  - **Network Acceleration**: When enabled, NAKIVO Backup & Replication uses compression and traffic reduction techniques to speed up data transfer. Select this option if you plan to back up over WAN or slow LAN links.

    Note

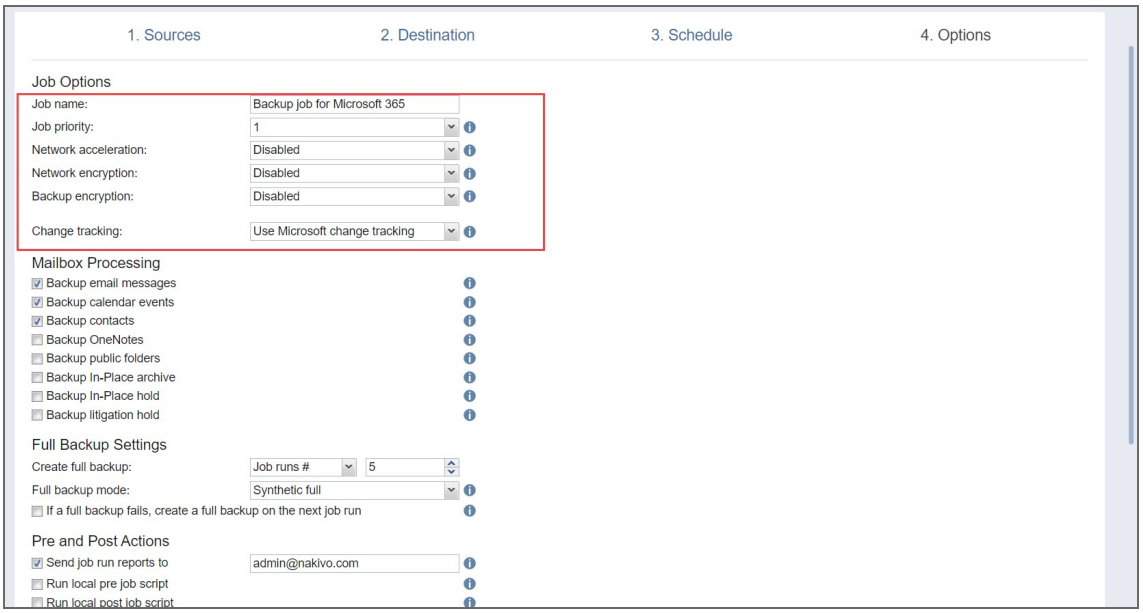
    The **Network acceleration** option is not available if the **Backup encryption** option is enabled.

  - **Network Encryption**: When enabled, VM data is protected with AES 256 encryption while traveling over the network.
    - Data encryption increases the backup time and CPU load on machines running Transporters. Select this option if you back up over WAN without a VPN connection.

- You need at least one **Transporter** at source and target sites to enable network encryption.

- Configure **Backup encryption** to protect your backup data with AES 256 block cipher encryption with a 256-bit key length. You can protect the backup file by creating a new password or selecting an existing one. For more information, refer to Enabling Backup Encryption.

> **Note**
>
> It's recommended that you enable the **(AWS) Key Management Service**. If **KMS** is enabled, all backup encryption passwords encrypted with the Key Management Service cryptographic key are available for recovery in case of product re-installation. For more information, refer to "Enabling KMS" on page 311.

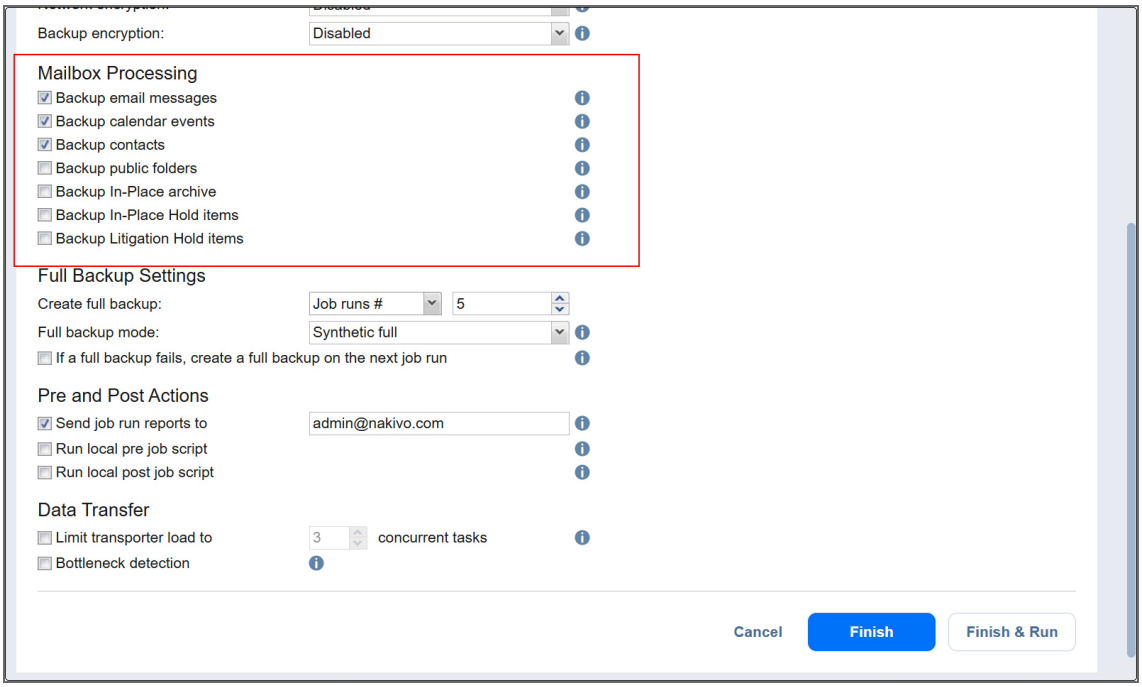


> **Notes**
>
> - If **Backup encryption** is enabled, the created recovery points are encrypted.
> - The **Backup encryption** option is not displayed for a backup job where forever incremental repositories are selected as the only target repositories.
> - The **Backup encryption** feature is not supported for Microsoft 365 backup jobs where SaaS repository is the only selected target repository type or one of repositories. Backup encryption is supported for Microsoft 365 backup jobs where only stream repositories are selected as target repositories.

Mailbox Processing

If you selected a mailbox on the **Source** page of the wizard, select at least one of the following mailbox items for backup by checking the respective option(s):

- Email messages
- Calendar events
- Contacts
- OneNotes
- Public folders
- In-Place archive
- In-Place hold items
- Litigation hold items
- Group mailbox files



Notes

- Selecting **Back up Group mailboxes files** may cause some files to be duplicated if you are also backing up the connected Group Site.
- Selecting **Back up public folders** may cause the folders to be duplicated if you back up other mailboxes with the same folders. Only public folders located in the `IPM_SUBTREE` folder can be backed up. Other public folders are skipped.
- If NAKIVO Backup & Replication is updated from a version that did not include support for Microsoft 365 Groups to a version that does, the **Back up Group mailboxes files** option is not selected automatically.
- If other Microsoft 365 items were selected in addition to Exchange Online Group mailboxes , the **Backup Group mailboxes files** option can still be selected and applies only to group mailboxes.

Full Backup Settings

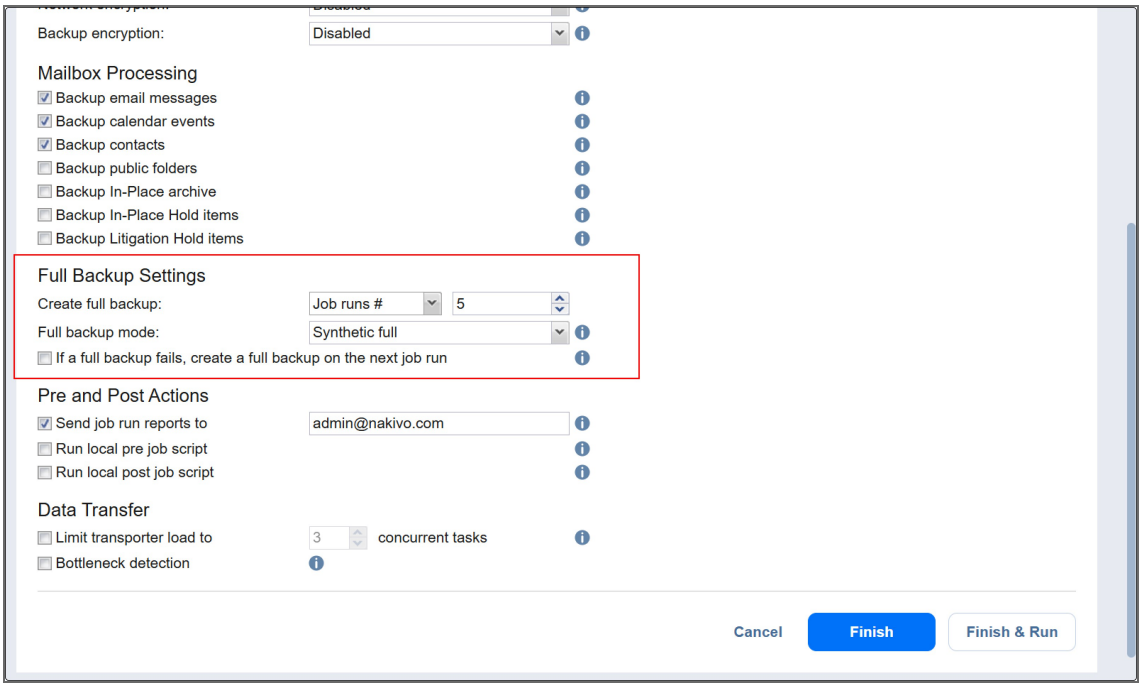
If the type of the **Backup Repository** that you've selected on the **Destination** page of the wizard is set to **Incremental with full backups** (**Store backups in separate files** option is selected), you can specify the following options:

- **Create full backup**: Specify how often full backups should be created.
- **Full backup mode**: Specify how the full backup should be created. You can choose one of the following options:
  - **Synthetic Full**: When this option is selected, NAKIVO Backup & Replication first performs an incremental backup (that is, transfers only the data that has changed since the last backup) and then transforms the available data into a full backup file. The benefits of this approach are:
    - The Synthetic Full backup is usually faster than the Active Full backup.
    - The load on the network is lower as less data is transferred.
    - The load on the source datastores running your production VMs is lower.
  - **Active Full**: When this option is selected, NAKIVO Backup & Replication reads all VM data from the source datastore and transfers it to the Backup Repository.
- **If a full backup fails, create a full backup on the next job run**: With this option selected, the next job run creates a full backup if the current job run fails to do so.

Notes

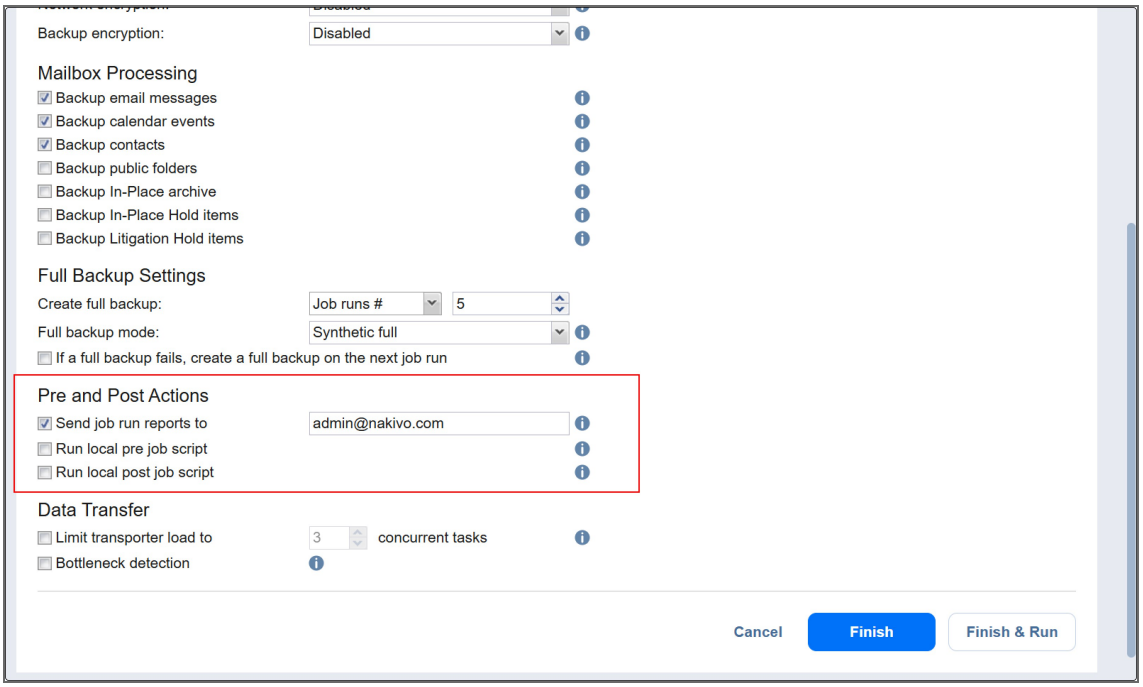
- A recovery point of a Microsoft 365 backup stored in an Incremental with full repository may have a bigger size compared to the same recovery point stored in a SaaS repository because each recovery point in Incremental with full repositories has an individual embedded database versus SaaS repositories, which only has one database for all recovery points.
- Compression rates for Microsoft 365 items stored in Incremental with full repositories may vary depending on the items that are backed up; no minimum compression rate can be defined.
- Synthetic full backups may have larger sizes than Active full backups because with Active full backups, empty clean metadata databases are created. As for Synthetic Full backup, metadata database from the previous successful savepoint is taken, and it's content is modified (add delta data in it. etc).
- During a backup run, the database file is processed and consumes some data. The

size of this file may be displayed in transferred data size.



Pre and Post Job Actions

With NAKIVO Backup & Replication, you can enable certain actions before a backup job begins and after it is completed. You can choose to send job run reports and run local pre- and post-job scripts. For more information, refer to "Pre and Post Job Scripts" on page 24.



Email Notifications

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

> **Note**
> To enable this option, configure your email settings. For details, refer to "Email Settings" on page 296.

To send email notifications, select **Send job run reports to** and specify one or more email addresses in the text box. Use the semi-colon character (;) to separate multiple email addresses.

Pre Job Script

To run a script before the product begins backing up your items, do the following:

1. Place a script file on the machine on which the Director is installed.
2. Select the **Run local pre job script** option.
3. Specify the following parameters in the dialog box that opens:
    - **Script path**: Specify a local path to the script on the machine on which the Director is installed. The script interpreter should be specified.
      Example (Windows): *cmd.exe /c D:\script.bat*
      Example (Linux): *bash /root/script.sh*
    - **Job behavior**: Choose one of the following job behaviors in relation to script completion:
        - **Wait for the script to finish**: When this option is selected, the backup is started only after the script is completed.
        - **Do not wait for the script to finish**: When this option is selected, the product runs the script and starts backing up your items at the same time.
    - **Error handling**: Choose one of the following job behaviors in relation to script failure:
        - **Continue the job on script failure**: When this option is selected, the product performs the backup job even if the script has failed.
        - **Fail the job on script failure**: When this option is selected and the script fails, the job fails and the backup is not performed.

Post Job Script

To run a script after the product has finished backing up all items, do the following:

1. Place a script file on the machine on which the Director is installed.
2. Select the **Run local post job script** option.
3. Specify the following parameters in the dialog box that opens:

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

Note
Pre- and post-job scripts can be executed only on the machine on which the Director is installed.

Data Transfer

In the *Data Transfer* section, you can configure the following options.

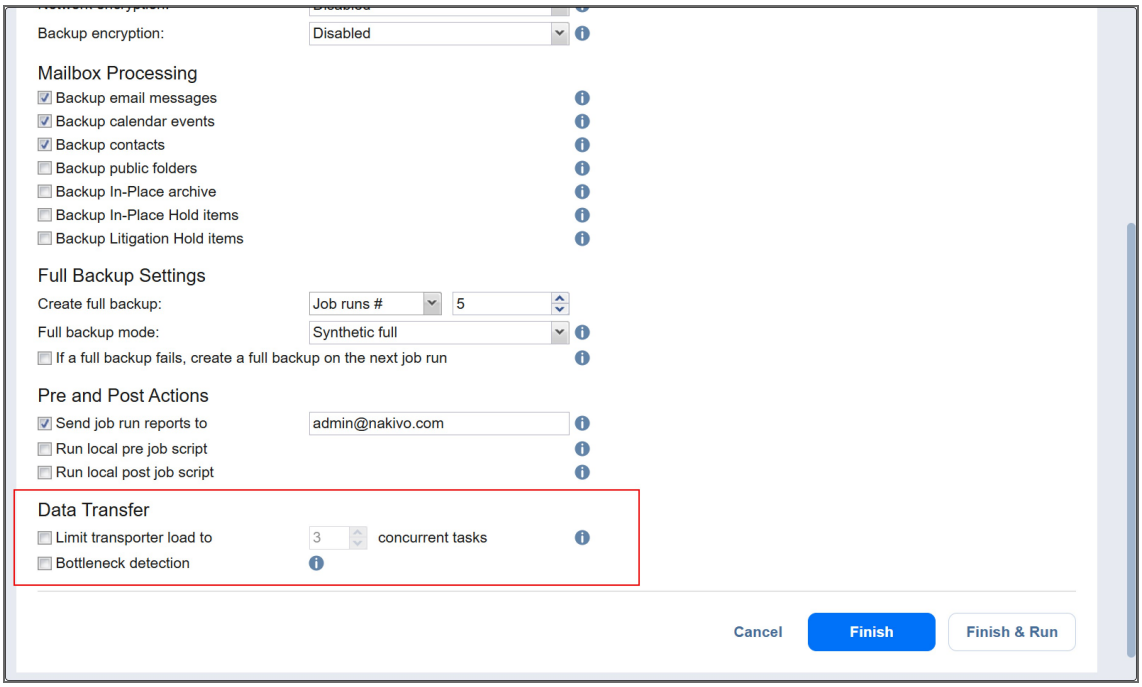
Transporter Load

You can limit the maximum number of Transporter tasks used by a job. By default, the number is set to 3 concurrent tasks. For a Microsoft 365 backup job, one task is equal to processing one mailbox or one OneDrive instance.
To change the default number of tasks, do the following:

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

Bottleneck Detection

When the **Bottleneck detection** option is enabled, additional information is collected and recorded in NAKIVO Backup & Replication logs in the course of data transfer for the purpose of bottleneck detection. Check this option to enable the **Bottleneck detection** capability of the **Transporters** engaged in the job.



Completing the New Backup Job Wizard for Microsoft 365

Click **Finish** or **Finish & Run** to complete the job creation.

> **Note**
>
> If you click **Finish & Run**, you will have to define the scope of your job.

Deleting Backups

With NAKIVO Backup & Replication, you can permanently delete a backup with all of its recovery points if this backup is available in a Backup Repository. You can also delete specific recovery points in a backup without affecting any of the other recovery points. The option to delete a specific recovery point can be used if you get an alert about corrupted recovery points in a backup.

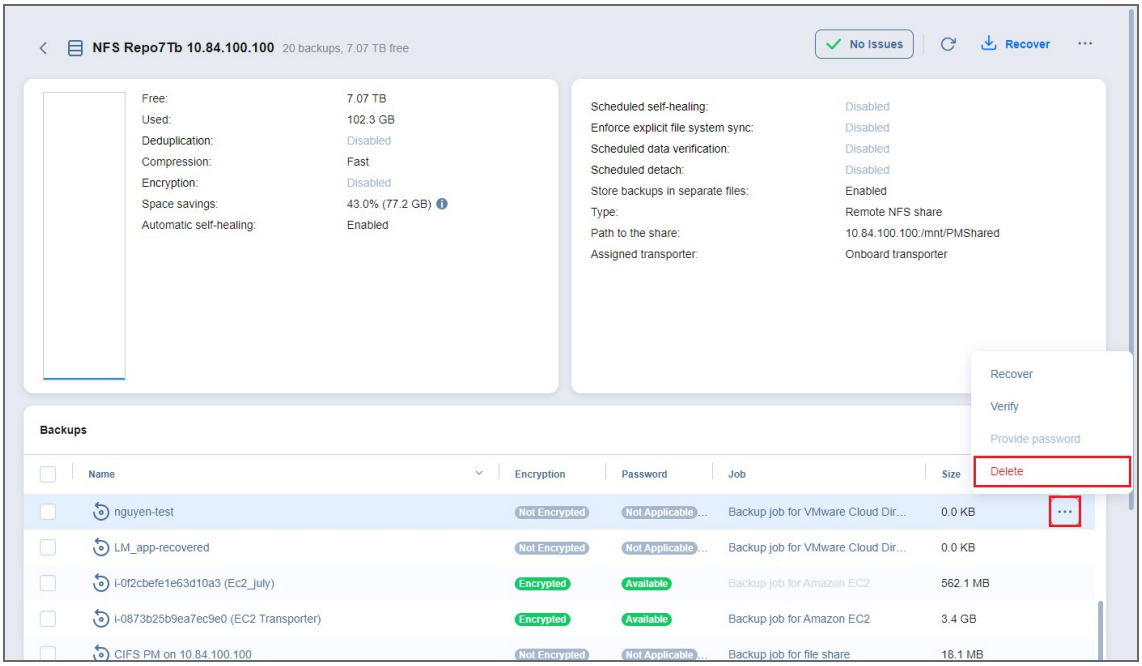
Refer to one of the following sections:

- Deleting a Single Backup
- Deleting Backups in Bulk
- Deleting Recovery Points
    - Deleting a Single Recovery Point
    - Bulk Recovery Points Deletion

Deleting a Single Backup

To delete a backup permanently, follow the steps below:

1. In the main menu, click **Settings**.
2. Go to the **Repositories** tab and click the Backup Repository you need.
3. Hover over the backup you want to delete, and on the right side, click **Delete**.



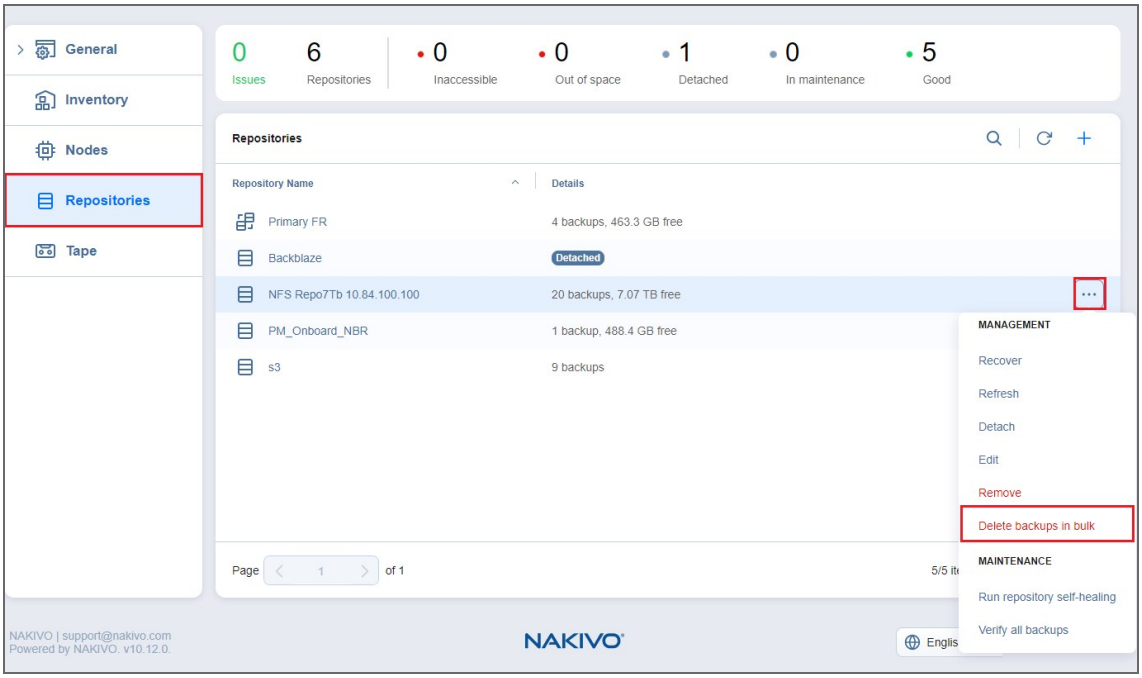
4. Click **Delete** in the dialog box that opens.

For SaaS Backup Repositories, manually removing backup data may not return space to the operating system correctly.

Deleting Backups in Bulk

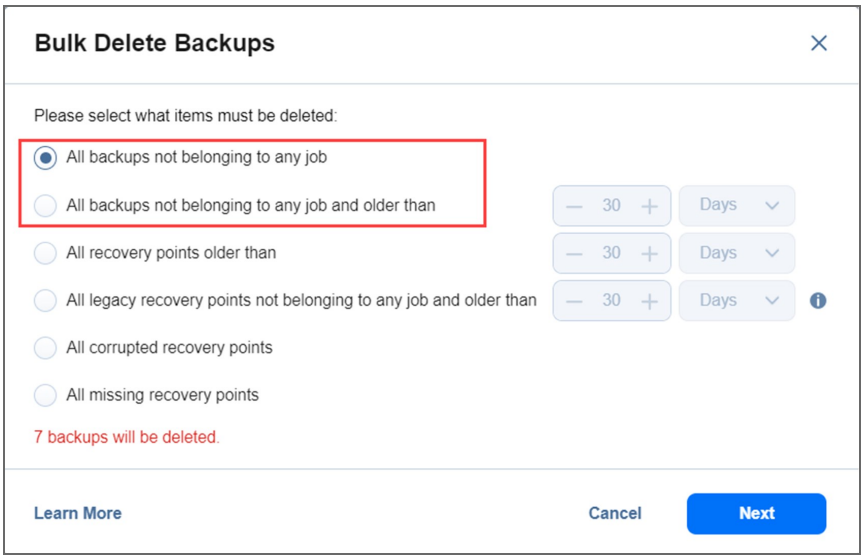
To permanently delete several backups that match specific criteria, follow the steps below:

1. In the main menu, click **Settings**.
2. Go to the **Repositories** tab and hover over the Backup Repository you need.
3. Click the ellipsis **Manage** button and then click **Delete backups in bulk**.

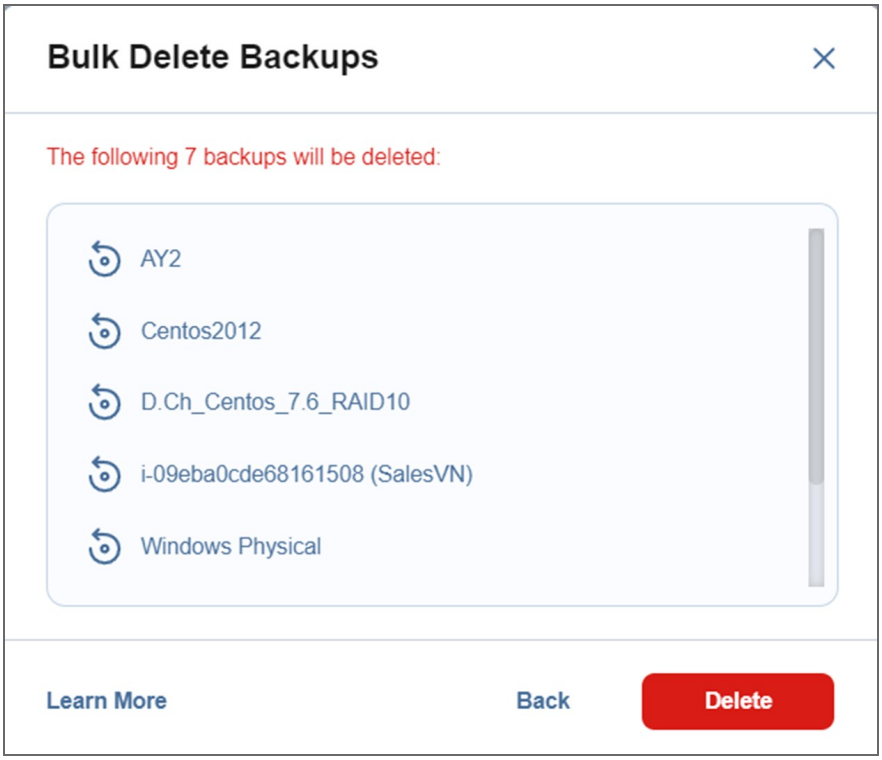


4. In the **Bulk Delete Backup** dialog box that opens, select one of the available options:
   - **All backups not belonging to any job**
   - **All backups not belonging to any job and older than** X <time_units>, where X is an integer and <time_units> is either days, weeks, or months
   - **All recovery points older than** X <time_units>, where X is an integer and <time_units> is either days, weeks, or months
   - **All legacy recovery points not belonging to any job and older than** X <time_units>, where X is an integer and <time_units> is either days, weeks, or months
   - **All corrupted recovery points**
   - **All missing recovery points**

   The dialog also shows the number of backups to be deleted.



5. Click **Next**.
6. The **Bulk Delete Backups** dialog box opens displaying the list of backups to be deleted. Click **Delete** to confirm the deletion.



For a **Forever-incremental** Backup Repository (that is, when the **Store backups in separate files** option is not selected), the space that was occupied by the deleted backup is marked "free" and reused by new data blocks on subsequent job runs. However, the actual size of the Backup Repository may not change. To free up the space that was occupied by the deleted backup, you can reclaim the free space.

Deleting Recovery Points

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

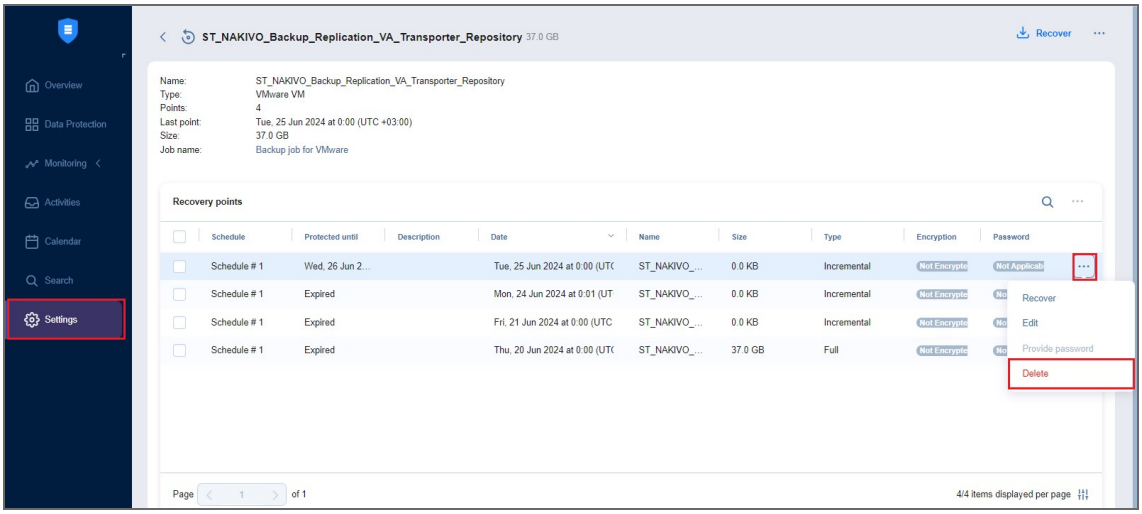
The **Delete** action is disabled:

- For VMware Cloud Director VMs underlying vApps
- For the only remaining and/or uncorrupted recovery point, except in the cases when the recovery point is set to be removed automatically according to the configured retention policy. In this case, removal of the last recovery point removes the corresponding backup object and data pertaining thereto.
- For the recovery point if there is at least one other recovery point depending on the current recovery point
- For inaccessible backup objects
- For pending removal recovery points
- For inconsistent repositories
- For the recovery point that is currently in use.

Deleting a Single Recovery Point

To delete a single recovery point in response to a corruption alert or for functional requirements, do the following:

1. In the main menu, click **Settings**.
2. Go to the **Repositories** tab and click the Backup Repository you need.
3. Click the backup with the recovery point you want to delete.
4. Hover over the recovery point that you want to delete, and on the right side, click **Delete**.

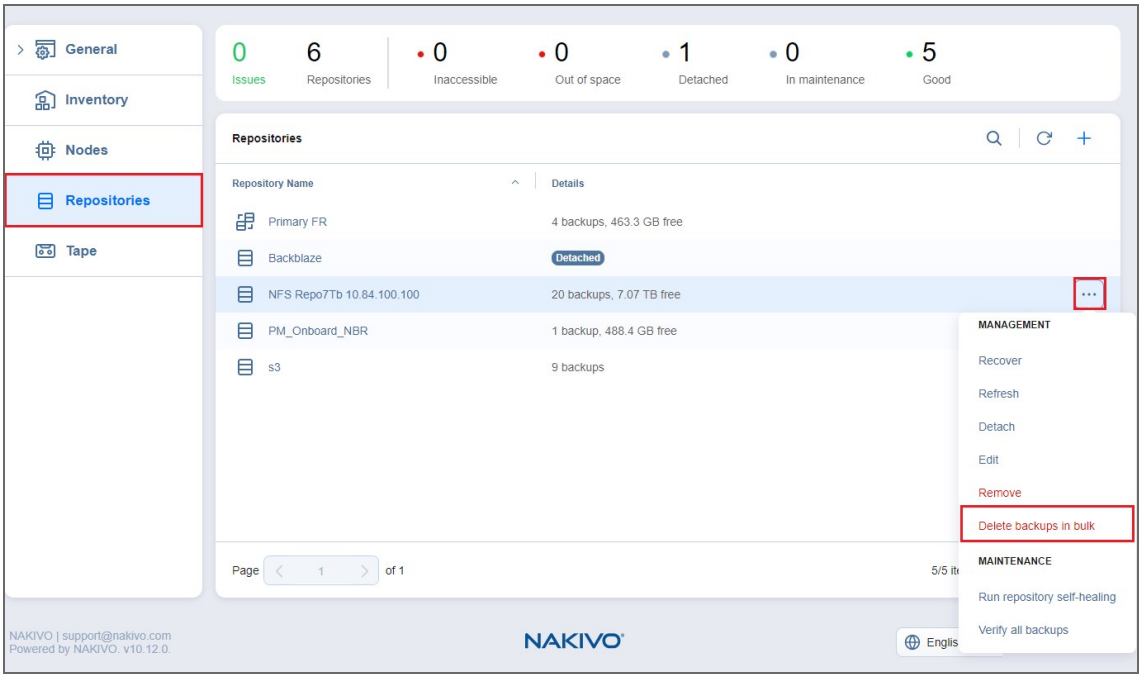


5. Click **Delete** in the dialog box that opens.

For SaaS Backup Repositories, manually removing backup data may not return space to the operating system correctly.

Bulk Recovery Points Deletion

1. In the main menu, click **Settings**.
2. Go to the **Repositories** tab and click the Backup Repository you need.
3. Click **Manage** and then click **Delete backups in bulk**.



4. In the **Bulk Delete Backups** dialog box that opens, select criteria for recovery points to be deleted:

- **All recovery points older than** X <time_units>, where X is an integer and <time_units> is either days, weeks, or months. When selected, the recovery points that are older than the specified time interval are deleted.

  - For **Forever-incremental** repositories (that is, when the **Store backups in separate files** option is not selected): If all recovery points of a backup match the deletion criteria, the latest recovery point – whether corrupted or not – is not deleted.
  - For **Incremental-with-full-backups** repositories (that is, when the **Store backups in separate files** option is selected):
    - Recovery points that are older than the end of the time interval that have dependent recovery points that are newer than the beginning of the time interval are not deleted.

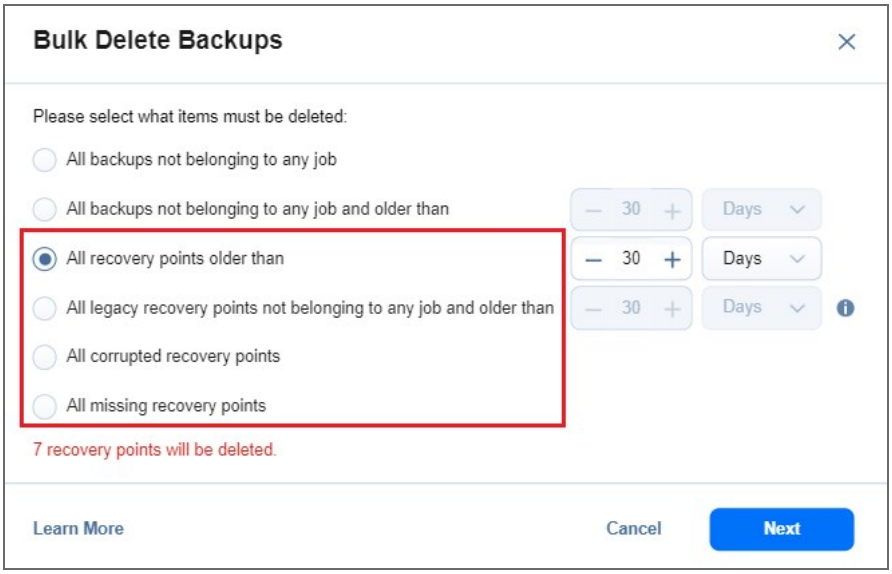
- **All corrupted recovery points**: When this option is selected, all recovery points that are corrupted are deleted. Recovery point selection criteria include the following:
  - For **Forever-incremental** repositories (that is, when the **Store backups in separate files** option is not selected), if a backup is used by a backup job and all its recovery points are corrupted, the latest recovery point is not deleted.
  - For **Incremental-with-full-backups** repositories (that is, when the **Store backups in separate files** option is selected), this option also deletes all recovery points that are dependent on corrupted recovery points. If all recovery points in a backup are corrupted or depend on a corrupted recovery point and match the deletion criteria, the latest full recovery point is not deleted.

    Note
    This option is not available for Microsoft 365 backups.

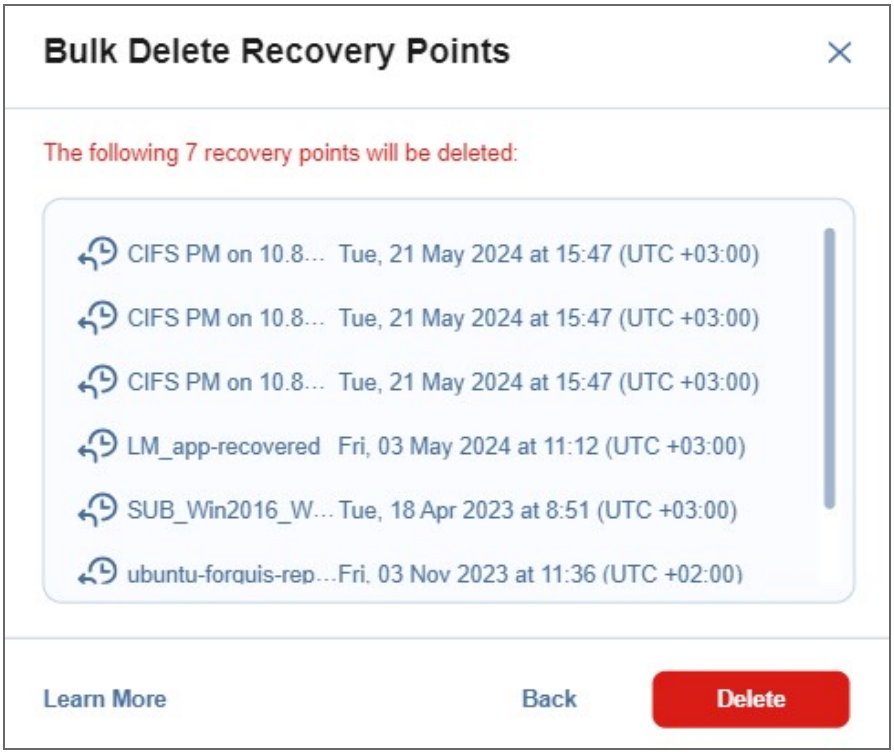
- **All missing recovery points**: When selected, all missing recovery points are deleted. Recovery point selection criteria include the following:
  - For **Forever-incremental** repositories (that is, when the **Store backups in separate files** option is not selected), this option deletes all missing recovery points. If all recovery points in a backup are missing, the latest recovery point is not deleted.
  - For **Incremental-with-full-backups** repositories (that is, when the **Store backups in separate files** option is selected), this option deletes all missing recovery points and any recovery points that are dependent on them. If all recovery points in a backup are missing or depend on missing recovery points, the latest full recovery point is not deleted.

> **Note**
> The recovery points of a missing backup object are not accounted as missing recovery points.

The dialog box shows the number of recovery points to be deleted.



5. The **Bulk Delete Recovery Points** dialog box opens displaying the list of recovery points to be deleted. Click **Delete** to confirm deleting the recovery points.



Note

For SaaS Backup Repositories, manual removal of backup data may not return space to the operating system correctly.

Object Recovery for Microsoft 365

The Object Recovery for Microsoft 365 feature enables browsing, searching, and recovering mailboxes, OneDrives, sites, and Teams directly from backups. This feature is agentless, works right out of the box, and does not require creating a special lab or running a special backup type.

Important

- Refer to Microsoft 365 Platform Limitations of the latest Release Notes for the full list of Microsoft 365 items that are not supported.
- When recovering from an incremental with full repository, the */nakivo/transporter/java/tmp* folder on the installed transporter must have more space than the size of 3 largest items being recovered, otherwise recovery may fail.
- Refer to Required API Permissions for Microsoft 365 to see the list of required permissions for recovering Microsoft 365 objects.
- Object Recovery for Microsoft 365 is not supported for backups created from 4K native (4Kn) disks.

With NAKIVO Backup & Replication, you can recover Microsoft 365 backup copies from tape devices to a target backup repository. For more details, refer to Backup Copy to Tape Recovery.

Refer to the following topics for more information:

- "Starting Object Recovery for Microsoft 365" on page 624
- "Microsoft 365 Object Recovery Wizard: Backup" on page 626
- "Microsoft 365 Object Recovery Wizard: Recovery Account" on page 629
- "Microsoft 365 Object Recovery Wizard: Objects" on page 630
- "Microsoft 365 Object Recovery Wizard: Options" on page 634
- "Starting Backup Copy to Tape Recovery Wizard" on page 642

Starting Object Recovery for Microsoft 365

You can start the recovery process either from the **Data Protection** menu, by using the Search function or from the **Repositories** page in Settings (for example, if you no longer have a backup job, but still have the backup). Refer to the following sections for more details:

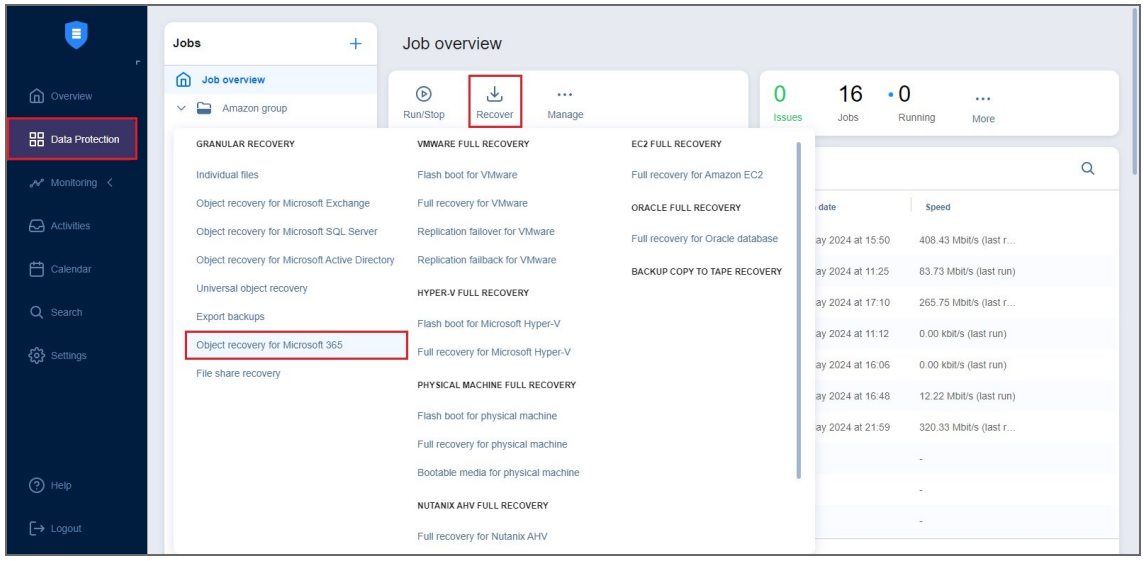
- Starting Object Recovery for Microsoft 365 from Data Protection Menu
- Starting Object Recovery for Microsoft 365 from a Backup Repository

Note

Object-level recovery may fail if the file or folder name, or the full path length exceeds the maximum length supported by the target operating system.

Starting Object Recovery for Microsoft 365 from Data Protection Menu

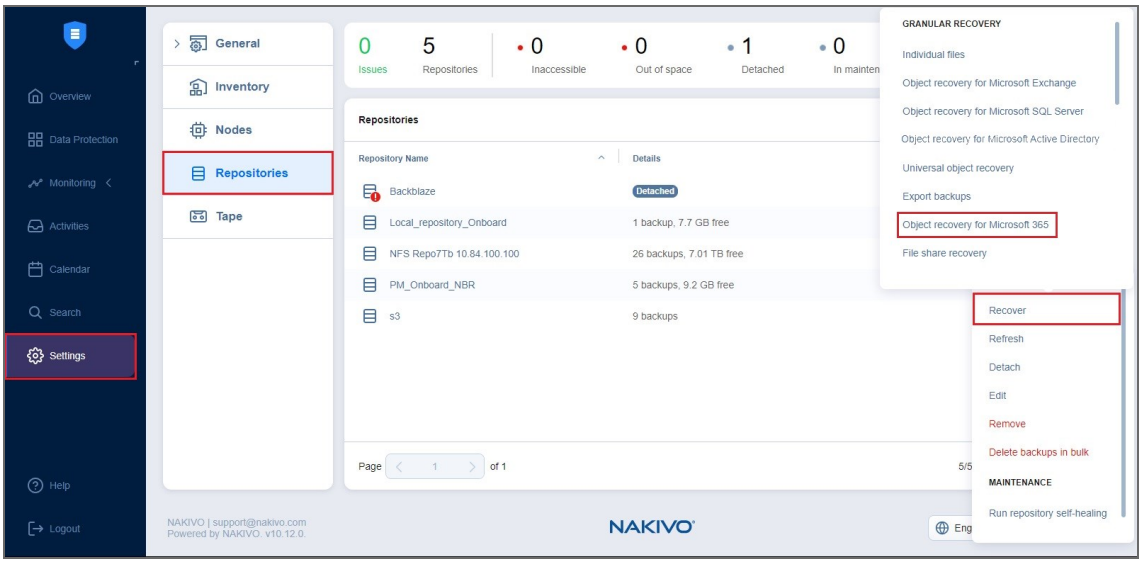
To start object recovery for Microsoft 365 from the **Data Protection** menu, click **Recover** and then click **Object recovery for Microsoft 365**.



Starting Object Recovery for Microsoft 365 from a Backup Repository

To start object recovery for Microsoft 365 from a **Backup Repository**, do the following:

1. Go to the main menu of NAKIVO Backup & Replication and click **Settings**.
2. Go to the **Repositories** tab and hover the cursor over the **Backup Repository** containing the required backup.
3. Click the ellipsis **Manage** button, click **Recover**, and select **Object recovery for Microsoft 365**.



The **Object Recovery Wizard for Microsoft 365** opens.

Microsoft 365 Object Recovery Wizard: Backup

On the **Backup** page of the **Object Recovery Wizard for Microsoft 365** proceed as follows:

1. Select a backed up mailbox, OneDrive, site, or Team using either a **Backup Repositories** or **Jobs & Groups** view in the left pane.

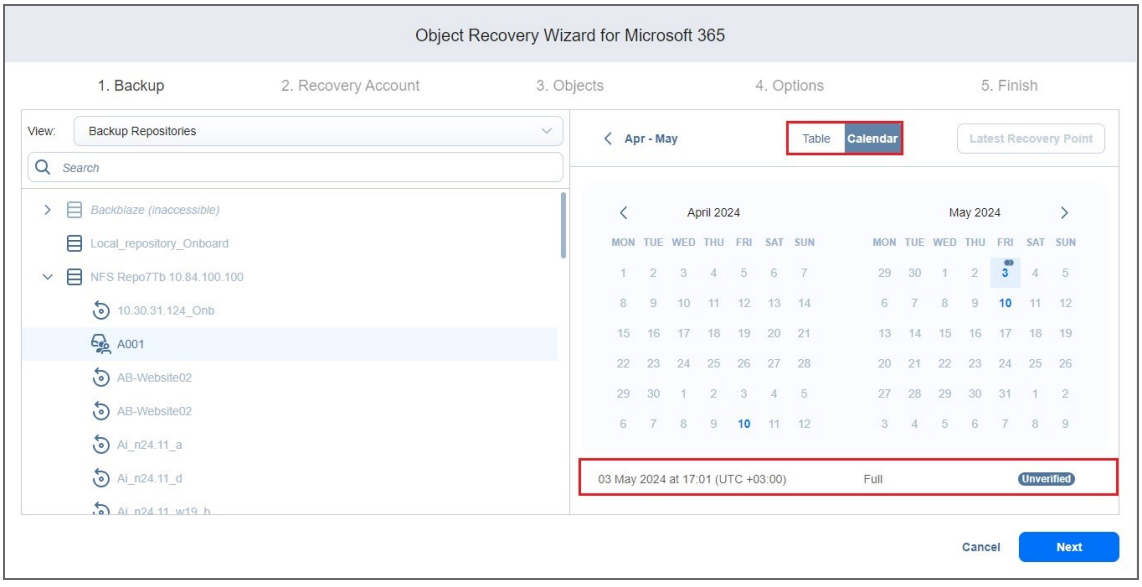
Note, that only one backup can be selected at a time.

Selecting an item in the tree displays all of its corresponding recovery points to the right pane.

Notes

- You can search for a specific backup, job, group, or repository (depending on the selected view) by entering its name into the **Search** field.
- Selecting a single backup object opens the **Calendar** view of the date picker if this view was the last user-selected view. Selecting a single backup object opens the **Table** view if this was the last view the user selected.
- You can select a backup object with the corrupted, inaccessible, or pending removal recovery points, but you cannot select the backup objects of the detached repository.
- It is not recommended to have more than one recovery session per backup object at the same time; other sessions may fail while the first one is running.

2. The latest recovery point is selected by default.



Notes

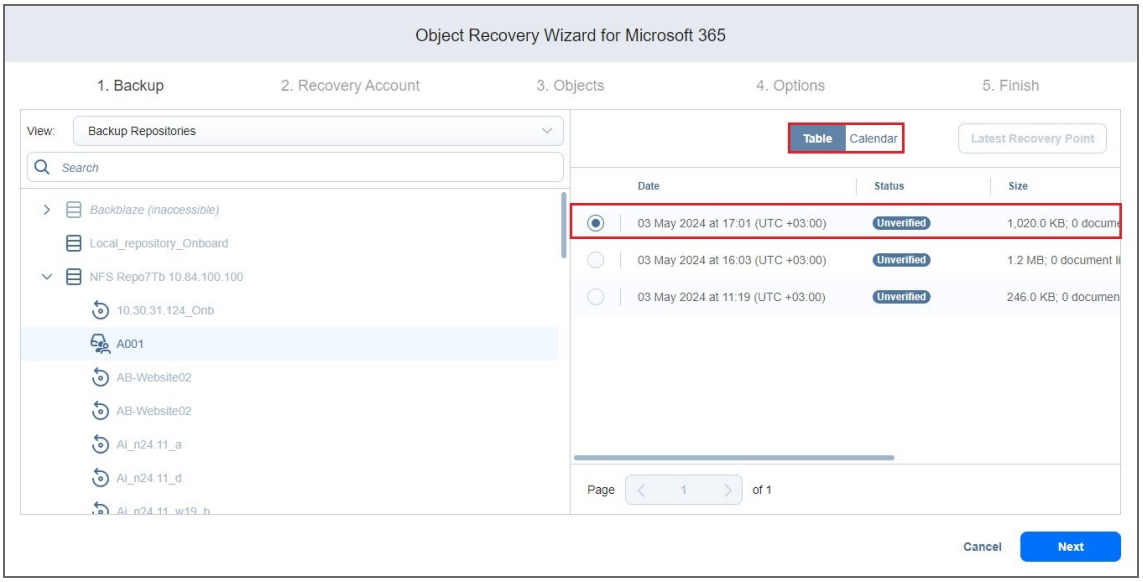
- The selected date is highlighted.
- If a recovery point is selected in the **Calendar** or **Table** view, this recovery point is selected by default the next time you open the **Calendar** or **Table** view.
- The selected view, either **Calendar** or **Table**, is saved on a per-user basis.

3. If necessary, toggle between the **Calendar** and **Table** views:
   - In the **Calendar** view, select a date or click **Latest Recovery Point** to select the date with the latest recovery point.

Notes

- The **Latest Recovery Point** option is disabled if:
  - The latest recovery point is currently selected.
  - All the recovery points of the backup object are corrupted, inaccessible, or pending removal.
- A small dot is displayed on top of the date if there is at least one recovery point created on that date.
- The selected recovery point has a preview displayed under the calendar showing the recovery point status:
- **green** – verified recovery points
- **red** – inaccessible/corrupted/pending removal recovery points
- **dark grey** – unverified recovery points
- Corrupted, inaccessible, or pending removal recovery points cannot be recovered.

   - In the **Table** view, select a date or click **Latest Recovery Point** to select the date with the latest recovery point.



Note

Clicking the **Latest Recovery Point** option selects the date with the latest recovery point that is not corrupted, inaccessible, or pending removal.

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

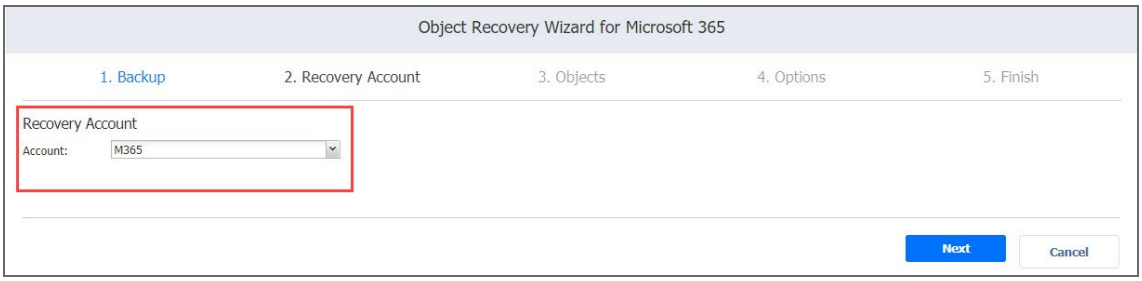
Note

You cannot proceed to the next step if there is at least one selected recovery point to which a password has not been provided.

Microsoft 365 Object Recovery Wizard: Recovery Account

On the **Recovery Account** page, specify the Microsoft 365 account to which you want to recover your items. You can choose to recover to a specific Online Exchange mailbox, OneDrive for Business, or a SharePoint site later in the wizard. Select the required account from the **Account** drop-down list. The list contains all Microsoft 365 accounts added to the inventory. Refer to "Adding Microsoft 365 Accounts" on page 404 for details.

> **Note**
> Recovering to a different Microsoft 365 account may be blocked in case the user has insufficient permissions.



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

Microsoft 365 Object Recovery Wizard: Objects

On the **Objects** page, the following items can be selected for recovery:

- Mailboxes including In-Place archive mailboxes, Public folders, In-Place hold and litigation hold items, contacts, contact lists, calendar events, and emails.
- Drives, folders, and files.
- Sites, subsites, document libraries, lists, list items, individual and public files and folders.
- Teams, channels, tabs, posts, and files.

Notes

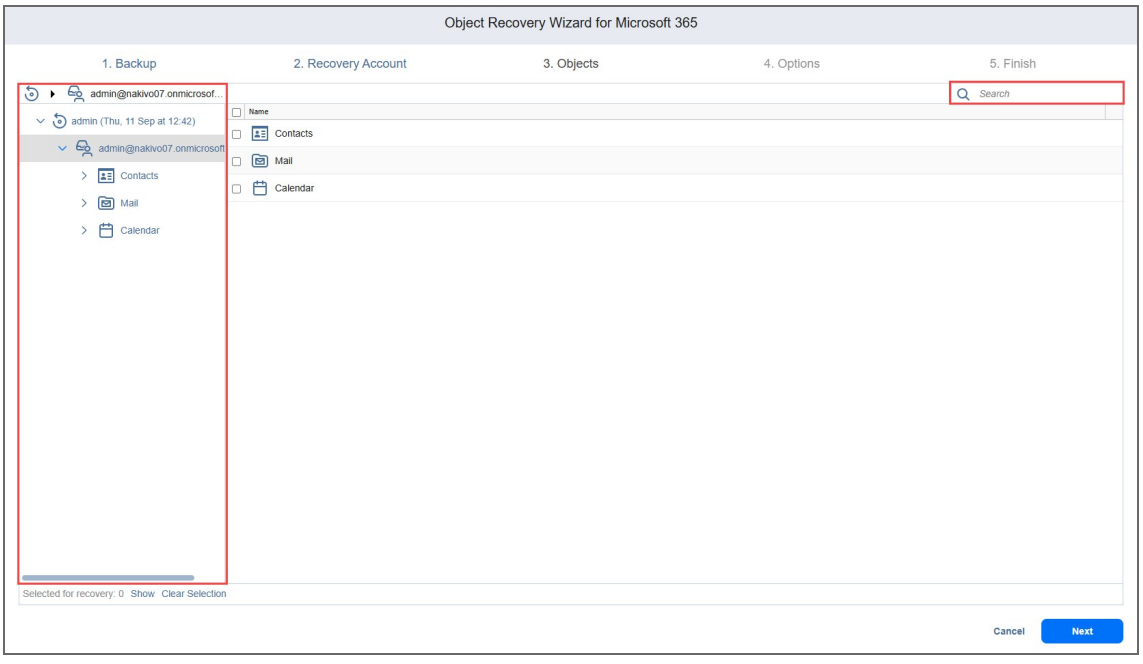
- Linked contacts are displayed as separate contact items.
- The lookup data of SharePoint Online lists displayed in the recovery wizard is inaccurate, and this data can't be recovered.
- Read-only calendars and calendars added from the directory can't be recovered.
- The following Teams tabs are not supported for recovery: Planner, OneNote, Stream, Forms, Power BI, Flow, Azure DevOps.
- If you are only recovering a tab to a different team and the tab contains a link to a document library, the recovery will fail with the error message: "Cannot recover a tab to a different team if the tab links to a document library."
- If recovering a team or channel to a different team and the team or channel includes a tab that links to a document library, the tab will be skipped for recovery and the recovery will be considered successful.
- If recovering files/folders that belong to the Team SharePoint document library (and do not belong to a specific channel) to a different channel, the files/folders will be recovered to the root "General" channel folder.
- If recovering files (regardless of whether they're linked to a tab or message) to a different location, the files will be recovered to the target channel following the selected overwrite behavior (skip/overwrite/rename). No additional folders will be created.
- Messages that were backed up but have the "This message has been deleted" tag cannot be recovered.
- When recovering a public folder, the target mailbox email address must be included in the `IPM_SUBTREE` list of users who can access the public folder. The user must have the "Owner" permission level for the public folder with the following permissions selected (refer to Adding Microsoft 365 Public Folder Permissions for more information):
  - Create items
  - Read items
  - Create subfolders
  - Edit own
  - Edit all
  - Folder owner
  - Folder contact
  - Folder visible
  - Delete own
  - Delete all

- Refer to the Microsoft 365 Platform Limitations of the latest Release Notes for the full list of Microsoft 365 items that are not supported.
- If Teams chat messages are the only items selected in the **Objects** step, a full snapshot of chat messages is restored instead of using overwrite behavior. Content of selected chat messages is exported to an HTML file and uploaded to a new tab in the target channel.

The **Objects** page contains:

- **Navigation**: Use the Navigation pane on the left to locate the items you need to restore. Selecting the container in the Navigation pane loads the container contents in the right pane. If a container has sub-folders, they are also displayed in the right pane.
- **Search**: The search box allows you to search for the objects that you want to recover. The search is performed within:
  - Mailboxes, folders, contacts, calendar items, files, and emails;
  - Drives, folders, and individual files;
  - Sites, document libraries, folders, files, lists, and list items.
  - Teams, channels, tabs, posts, and files.

  Enter the word or a part of the word into the search box and hit Enter on the keyboard. All items matching the criteria will be displayed in the *Contents* pane on the right.

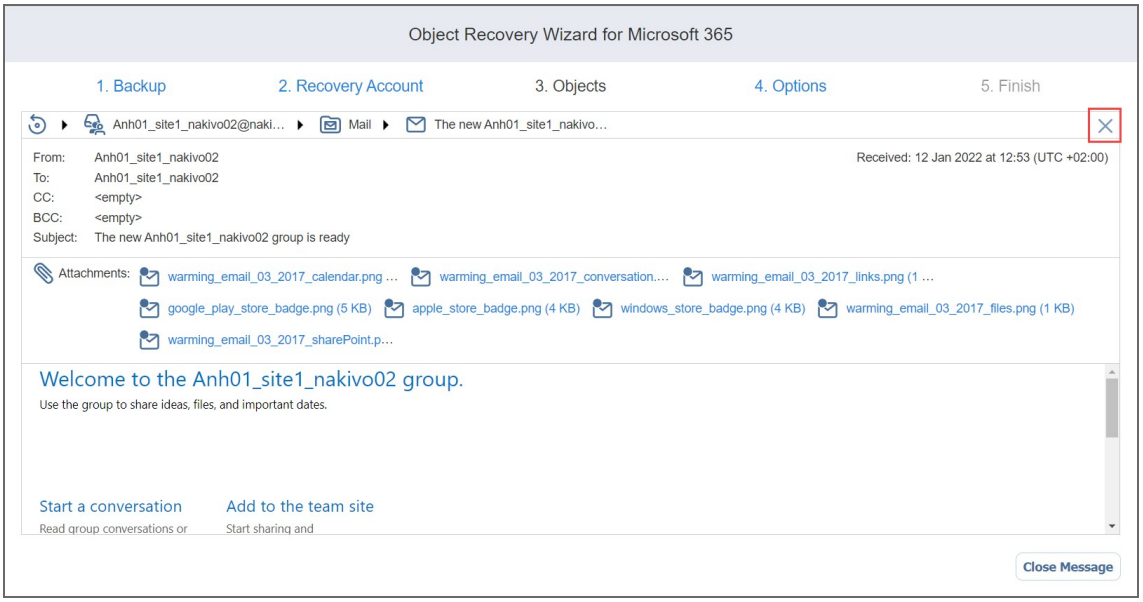


When recovering emails, contacts, or calendar events, you can also preview the contents of these items by clicking the item name in the right pane.

Note

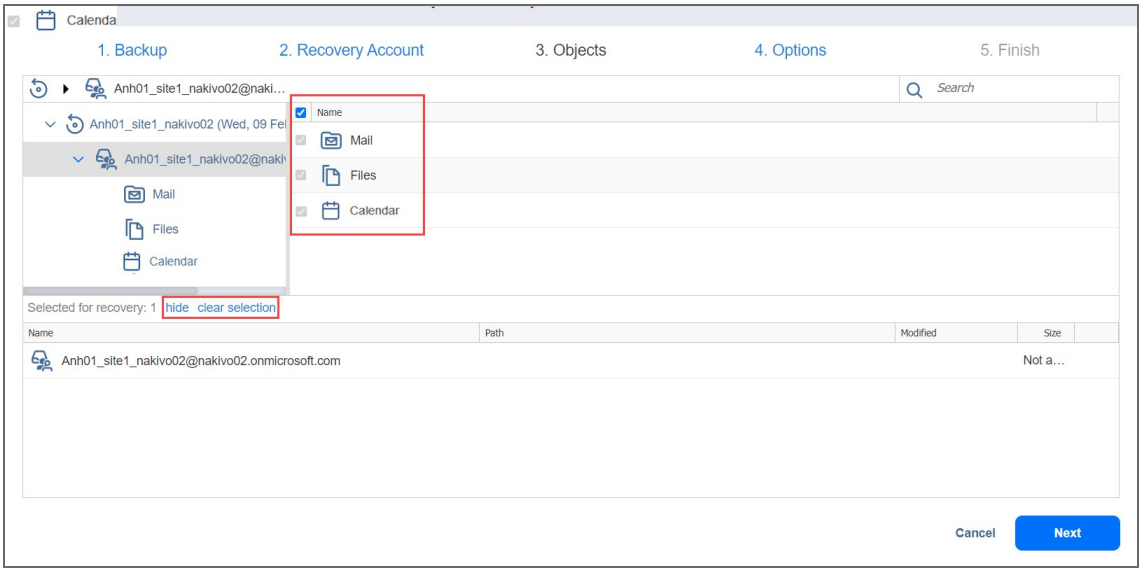
Emails may be blocked from reading in case the user has insufficient permissions.

To close the email message, contact, or calendar event preview, click the **Close** button at the bottom or click the **X** button above the item body.



After locating the items you need to recover, select the checkboxes next to their names. You can select different object types for recovery at the same time. The number of items selected for recovery is displayed at the bottom of the wizard page. You can also:

- Click **show** to view the list of all items selected for recovery.
- Click **clear selection** to clear the list of items selected for recovery.
- Click **hide** to hide the list of items selected for recovery.



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

Microsoft 365 Object Recovery Wizard: Options

On the **Options** page, specify the location to which the items will be restored and overwrite behavior for folders. Proceed as described in the sections below:

- Recovering to the Original Location
- Recovering to Mailbox or to Group
- Recovering to OneDrive
- Recovering to Site
- Recovering to Team

Recovering to the Original Location

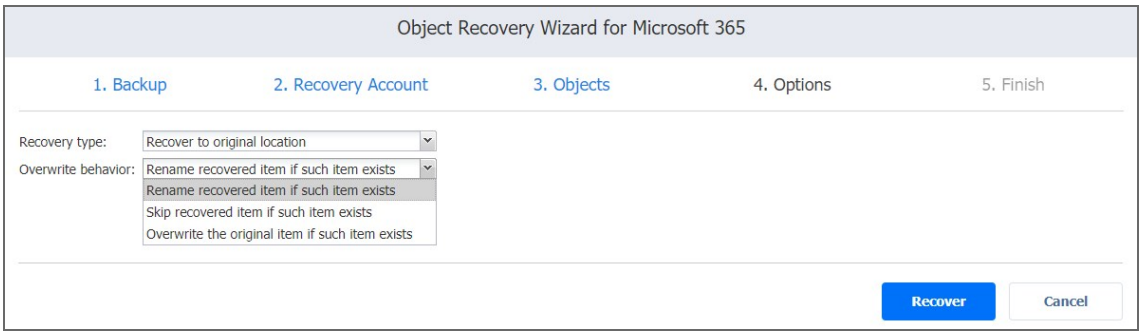
To recover Microsoft 365 items to the original location, take the following steps:

Notes

- Group mailboxes can only be recovered to the original location or other group mailbox types.
- User mailboxes can be recovered to the original location or other group mailbox types and cannot be recovered to the group mailbox type.
- Personal site non-personal site items can be recovered to group site type and vice versa.

1. Select **Recover to original location** from the **Recovery type** drop-down list. All selected items will be recovered to their original locations within the Microsoft 365 account. This option is disabled if a mailbox, OneDrive, or site no longer exists (deleted from the product, for example) or if you have selected a recovery account on the **Recovery Account** page of the wizard that is different from the original account.
2. Choose the naming convention for the recovered items by selecting one of the following **Overwrite behavior** options:
    - **Rename recovered item if such item exists**
    - **Skip recovered item if such item exists**
    - **Overwrite the original item if such item exists**

3. Click **Recover**.



Recovering to Mailbox or to Group

If you want to recover items to a group or to a specific mailbox, take the following steps:

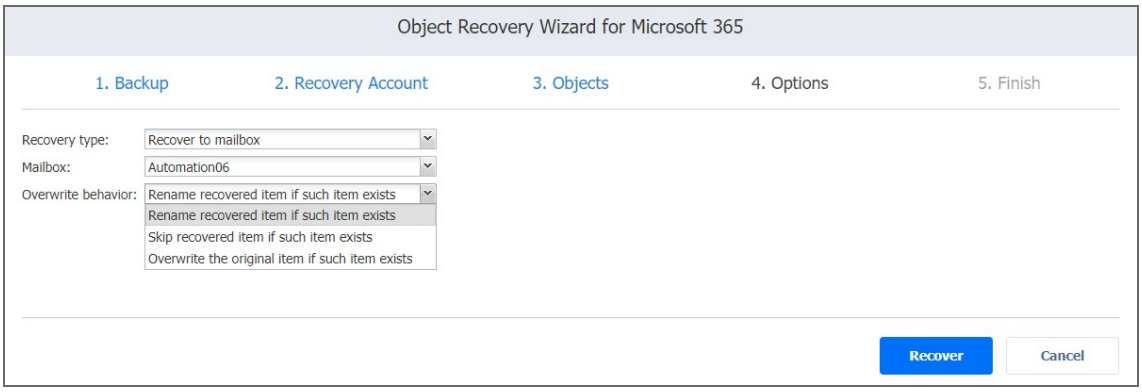
> **Note**
> Recovering to a different mailbox or group mailbox may be blocked in case the user has insufficient permissions.

1. Select **Recover to mailbox** from the **Recovery type** drop-down list. Note that **Recover to group** is displayed if group mailbox items were selected.
2. Select the required mailbox from the **Mailbox** drop-down list. You may also search for the required mailbox by typing its name or part of its name into the search bar.

> **Notes**
> - **Group** is displayed if group mailbox items were selected.
> - The following may occur if recovering Group messages to a different Group mailbox:
>   - The original sender of the message may be changed to the name of the user used for account discovery.
>   - The original receiver of the message may be changed to the target user group mailbox name.
>   - The timestamp of the message may be changed.

3. Choose the overwrite behavior for the recovered items by selecting one of the following **Overwrite behavior** options:
   - **Rename recovered item if such item exists**
   - **Skip recovered item if such item exists**
   - **Overwrite the original item if such item exists**

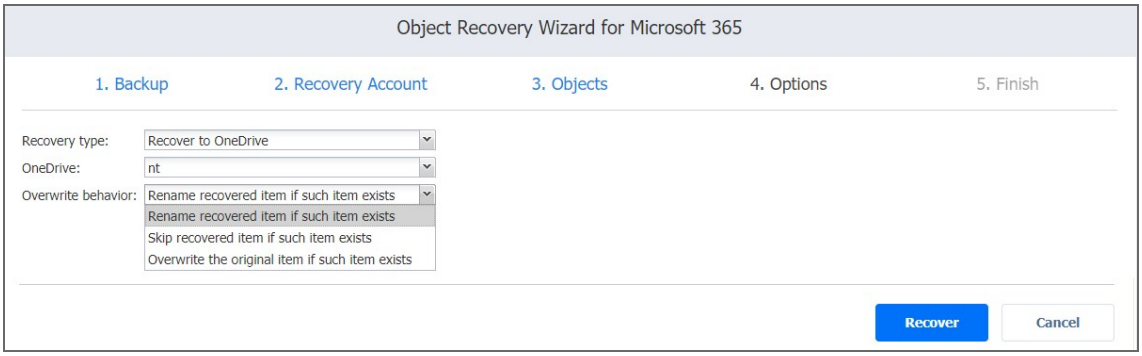
4. Click **Recover**.



Recovering to OneDrive

To recover OneDrive items to a specific OneDrive account, take the following steps:

1. Select **Recover to OneDrive** from the **Recovery type** drop-down list.
2. Select the required OneDrive from the **OneDrive** drop-down list. You may also search for the required OneDrive by typing its name or part of its name into the search bar.
3. Choose the overwrite behavior for the recovered items by selecting one of the following **Overwrite behavior** options:
    - **Rename recovered item if such item exists**
    - **Skip recovered item if such item exists**
    - **Overwrite the original item if such item exists**
4. Click **Recover**.



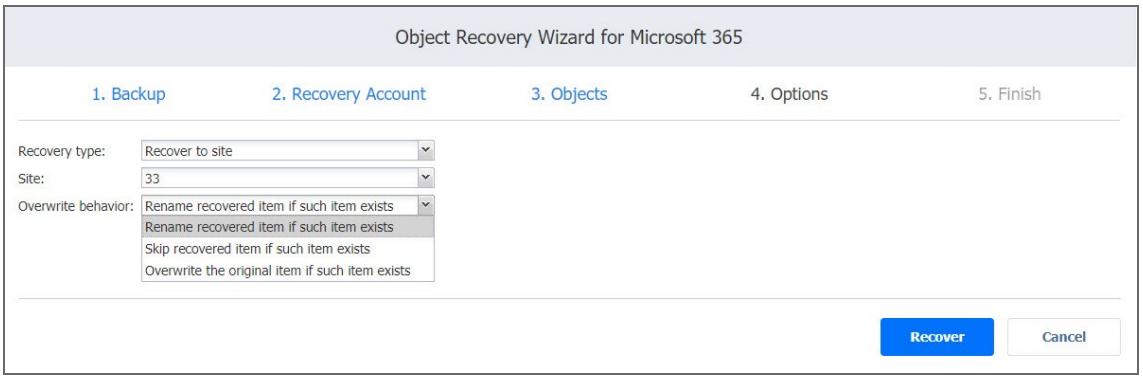
Recovering to Site

To recover SharePoint objects to a specific site, take the following steps:

1. Select **Recover to site** from the **Recovery type** drop-down list.

> **Note**
> If **Recover to site** is selected, the source site, including the subsites, is recovered to the target site as a subsite.

2. Select the required site from the **Site** drop-down list. You may also search for the required site by typing its name or part of its name into the search bar.
3. Choose the overwrite behavior for the recovered items by selecting one of the following **Overwrite behavior** options:
   - **Rename recovered item if such item exists**
   - **Skip recovered item if such item exists**
   - **Overwrite the original item if such item exists**
4. Click **Recover**.



When the recovery process is completed, the **Finish** page is displayed. You cannot return to the previous pages at this point; however, you can check the progress of the job by clicking the **Activities** link.

Recovering to Team

If an **entire team** was selected in the **Objects** step, do the following:

1. Select **Recover to team** from the **Recovery type** drop-down list.
2. Select the required team from the **Team** drop-down list. You may also search for the required team by typing its name or part of its name into the search bar.
3. Choose the overwrite behavior for the recovered items by selecting one of the following **Overwrite behavior** options:
   - **Rename recovered item if such item exists**
   - **Skip recovered item if such item exists**
   - **Overwrite the original item if such item exists**
4. To revert your team's settings (including member permissions, @mentions, fun stuff, etc.) to those in the chosen backup, check the **Restore team settings** option.

Note

This option restores the following settings:

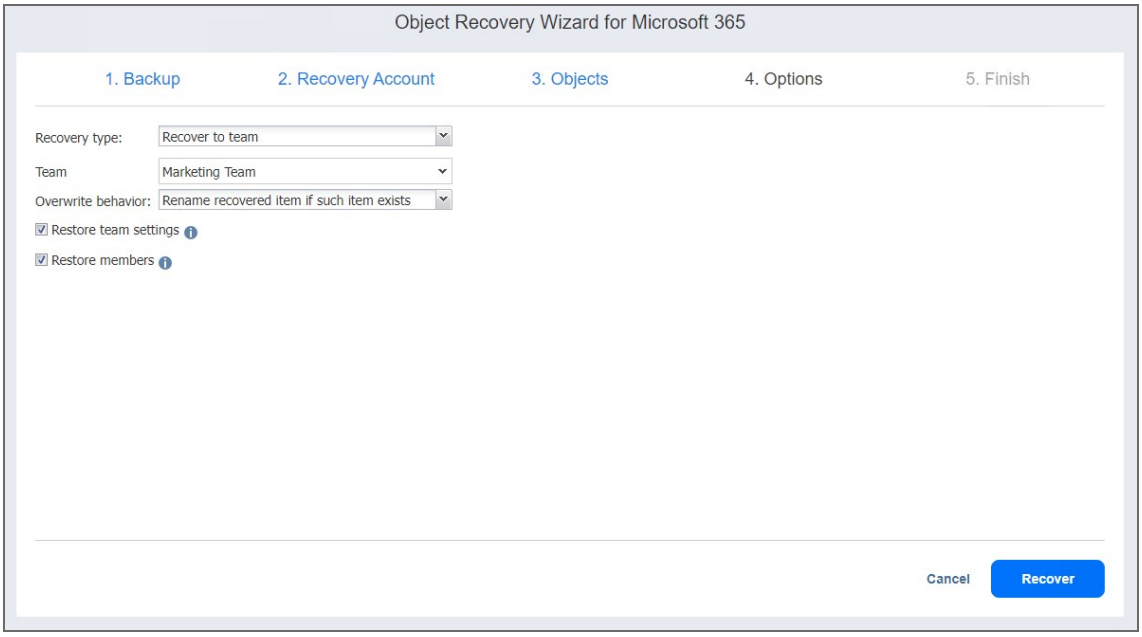
- Description
- isMembershipLimitedToOwners
- funSettings (allowGiphy, giphyContentRating, allowStickersAndMemes, allowCustomMemes)
- guestSettings (allowCreateUpdateChannels, allowDeleteChannels)
- memberSettings (allowCreateUpdateChannels, allowCreatePrivateChannels, allowDeleteChannels, allowAddRemoveApps, allowCreateUpdateRemoveTabs, allowCreateUpdateRemoveConnectors)
- messagingSettings (allowUserEditMessages, allowUserDeleteMessages, allowOwnerDeleteMessages, allowTeamMentions, allowChannelMentions)
- specialization (none, education_standard, healthcare_standard, etc)
- visibility (private, public, hidden_membership)
- discoverySettings (showInTeamsSearchAndSuggestions)
- team photo (may not always be restored due to an internal error on Microsoft's side, same as for Groups)

5. To revert your list of team members (including tags related to team members) to those in the chosen backup, check the **Restore members** option.

Note

If a user does not belong to the target account, that user is not added as a member to the target team. Recovery of member permissions for this user is skipped.

6. Click **Recover**.



If **team items** (channels, tabs, posts, etc.) were chosen in the **Objects** step, do the following:

1. Select **Recover to team** from the **Recovery type** drop-down list.
2. Select the required team from the **Team** drop-down list. You may also search for the required team by typing its name or part of its name into the search bar.
3. Select the required channel from the **Channel** drop-down list.
4. Choose the overwrite behavior for the recovered items by selecting one of the following **Overwrite behavior** options:
    - **Rename recovered item if such item exists**
    - **Skip recovered item if such item exists**
    - **Overwrite the original item if such item exists**
5. Click **Recover**.

Note

- Depending on the selected overwrite behavior, the following channel settings are recovered if a Team or channel is selected for recovery:
  - Name
  - Description
  - membershipType: standard, private, shared
  - isFavoriteByDefault
  - channel moderation settings (allowNewMessageFromBots, allowNewMessageFromConnectors, replyRestriction (everyone, authorAndModerators), userNewMessageRestriction (everyone, everyoneExceptGuests, moderators))
- If chat messages (posts) are selected for recovery, original chat messages are not overwritten or modified. Instead, an HTML file containing the selected messages is created and added as a new tab in the target channel. This preserves the original content and authorship and ensures compatibility with Microsoft API limitations.

If multiple channels are selected for recovery to a different team:

- Posts and tabs in selected channels are recovered to the target channel.
- Files in selected channels are recovered to newly created folders in the target channel.

Backup Copy to Tape Recovery

The Backup Copy to Tape Recovery feature allows recovering backups from tape devices to a target backup repository for file share backups of physical machines and Microsoft 365 backups.

With NAKIVO Backup & Replication, you can back up/send backup copies of selected folders or send immutable recovery points to repositories, then recover from these repositories and send the backup copies to tape.

The following repository destinations are supported:

- Onboard repository
- NFS and SMB share
- Local folder
- Amazon S3
- Amazon EC2
- Wasabi
- Backblaze
- S3-compatible storage
- Azure Blob
- HPE StoreOnce
- Dell EMC Data Domain
- NEC HYDRAstor

Notes

- The feature supports recovering only file share and Microsoft 365 backups from tape to Backup Repository.
- 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.

This section includes the following topics:

- "Starting Backup Copy to Tape Recovery Wizard" on the next page
- "Backup Copy to Tape Recovery Job Wizard: Backups" on page 644
- "Backup Copy to Tape Recovery Job Wizard: Destination" on page 646
- "Backup Copy to Tape Recovery Job Wizard: Options" on page 647

Starting Backup Copy to Tape Recovery Wizard

You can start the recovery process either from the **Data Protection** menu, by using the **Search** function, or from the **Repositories** tab in **Settings** (for example, if you no longer have a backup job, but still have the backup). Refer to the following sections for more details:

- Starting Backup Copy to Tape Recovery from Data Protection Menu
- Starting Backup Copy to Tape Recovery from Tape Tab

Notes

- The following backup types are supported in the Backup Copy to Tape Recovery Job Wizard:
  - File share backup
  - Microsoft 365 backups
- File share backup and Microsoft 365 backup options are disabled in the following recovery job wizards when recovering backups from tape:
  - New Recovery Job Wizard for VMware vSphere
  - New Recovery Job Wizard for Amazon EC2
  - New Recovery Job Wizard for Microsoft Hyper-V
  - New Recovery Job Wizard for Nutanix AHV
  - New Recovery Job Wizard for Physical Machine

Starting Backup Copy to Tape Recovery from Data Protection Menu

To start backup copy to tape recovery from the **Data Protection** menu, click **Recover**, and in the **Backup Copy to Tape Recovery** section, click **Recover backup to repository** to open the **New Backup Copy to Tape Recovery Job Wizard**.



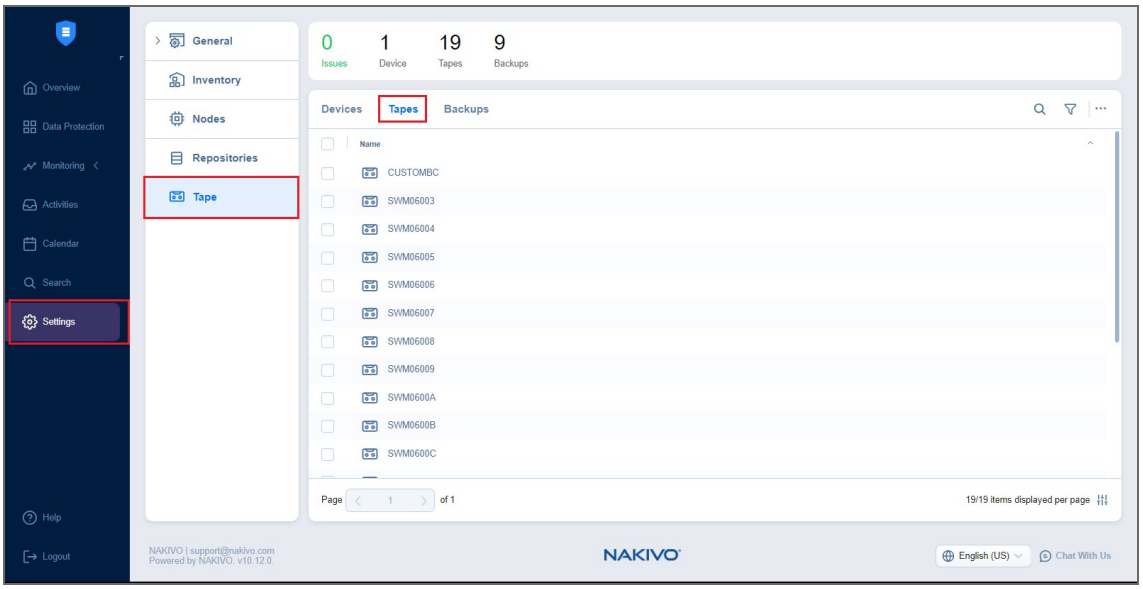
Notes

- The **Backup Copy to Tape Recovery** section is displayed only for backup copy jobs to tape whose sources are file share and Microsoft 365 backups.
- If no backup repositories are added to the product, the **Recover backup to repository** option is disabled.

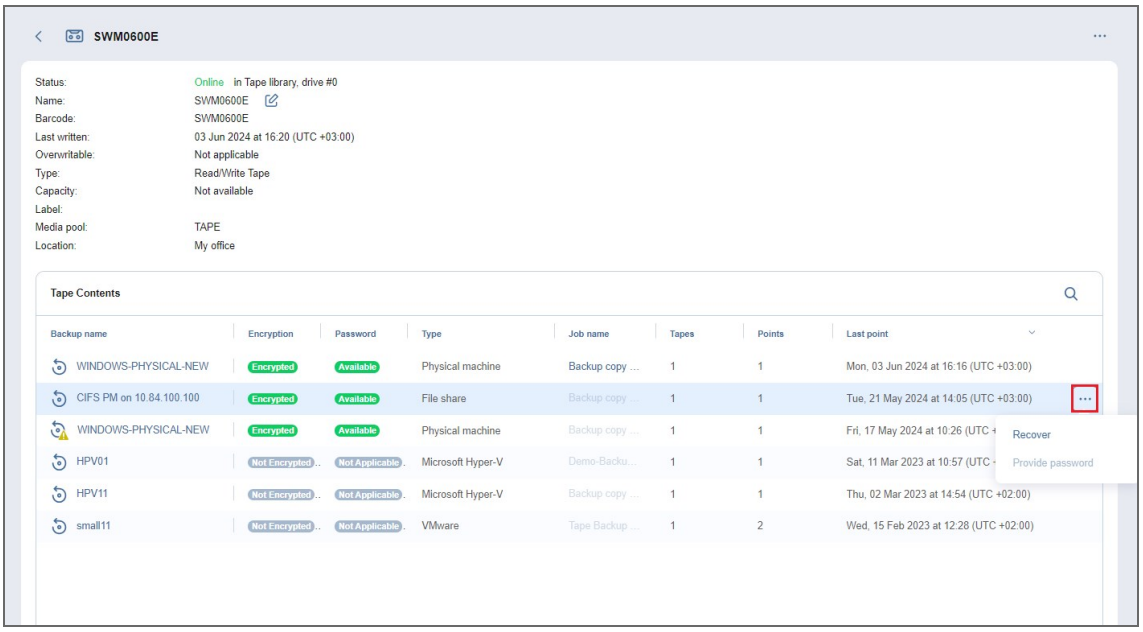
Starting Backup Copy to Tape Recovery from Tape Tab

To start backup copy to tape recovery from the **Tape** tab, do the following:

1. In the main menu of NAKIVO Backup & Replication, click **Settings**.
2. Go to the **Tape** tab and search for the tape cartridge containing the needed backup.



3. Click the cartridge name to open the tape cartridge management page.



4. In the **Tape Contents** table, select the file share or Microsoft 365 backup.
5. Hover over the needed recovery point and click the ellipsis **Manage** button.
6. In the dialog box that opens, click **Recover** to open the **New Backup Copy to Tape Recovery Job Wizard**.

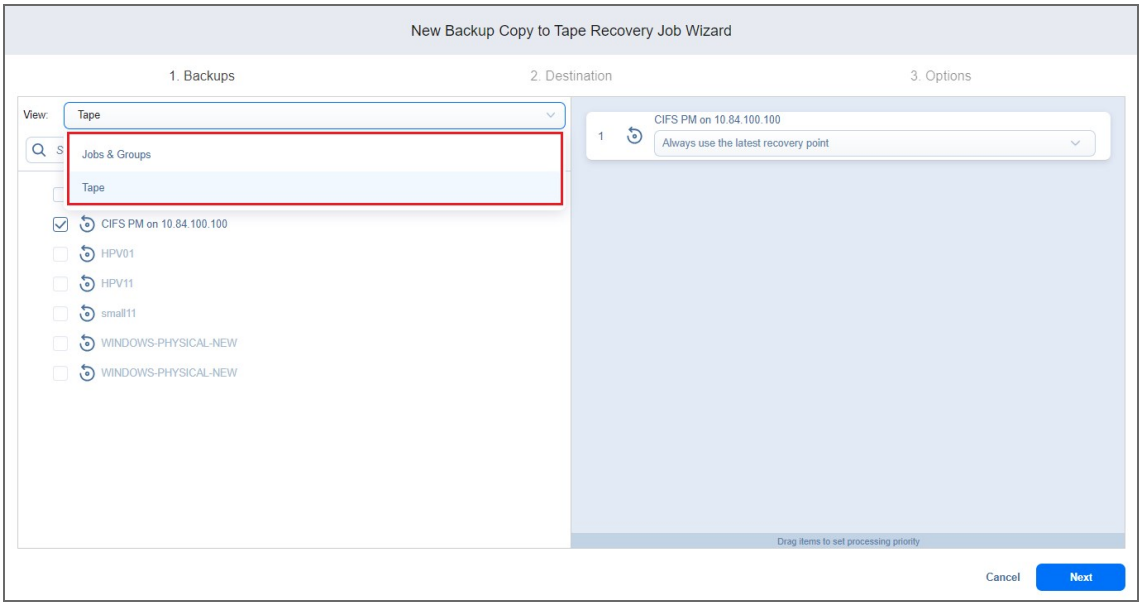
Backup Copy to Tape Recovery Job Wizard: Backups

On the **Backups** page of the wizard, proceed as follows:

1. Select one of the following views:
   - **Jobs & Groups**: In this view, do the following:
     a. Select one or more backups of the appropriate hypervisor type.
     b. Select one or more backup or backup copy jobs that contain backups of the appropriate hypervisor type.

> **Notes**
>
> - Backup and backup copy jobs and backups of unsupported hypervisor types are disabled.
> - Backup and backup copy jobs using the backup repository as a destination are disabled.
> - Replication, recovery, flash VM boot jobs, and their children are disabled.
> - If the wizard is initiated from the **Group** dashboard and all jobs in the group use tape as a destination, the wizard opens in the **Recover from tape** mode.
> - The item selected in the navigation tree during the recovery initiation is expanded by default.
> - If java transporter (or java transporter extension) is uninstalled from repository's transporter, items may not load in recovery wizard.

   - **Tape**: In this view, select one or more tape backups available for recovery in the left pane.
2. You can search for a specific backup, job, or group (depending on the selected view) by entering its name into the **Search** field.



3. Click **Next** to go to the next step.

Backup Copy to Tape Recovery Job Wizard: Destination

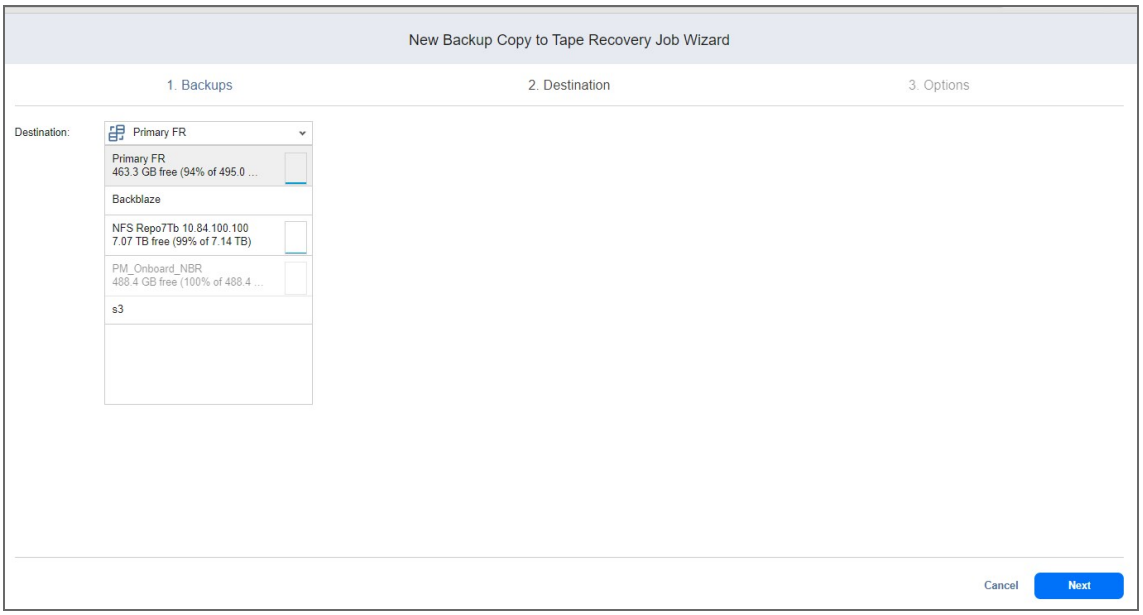
On the **Destination** page of the wizard, select a location for the recovered backup copies.

Proceed as described in the section below.

Selecting a Target Backup Repository

Backup Copy to Tape Recovery jobs can copy backups from one Backup Repository to another. Select a target Backup Repository as described below:

- To copy all the backups that 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.



> **Note**
>
> SaaS backup repositories are not available for selection at this step.

- To copy backups to different Backup Repositories, follow the steps below:
  1. Click **Advanced setup…**.
  2. For each backup, select a target Backup Repository.
  3. Select **Use existing backup as a target** option and select the existing backup copy from the drop-down list.

Backup Copy to Tape Recovery Job Wizard: Options

On the **Options** page, set up recovery job options as described in the sections below:

- Job Options
- Pre and Post Actions
  - Email Notifications
  - Pre Job Script
  - Post Job Script
- Data Transfer
  - Bandwidth Throttling
- Completing the New Backup Copy to Tape Recovery Job Wizard

Job Options

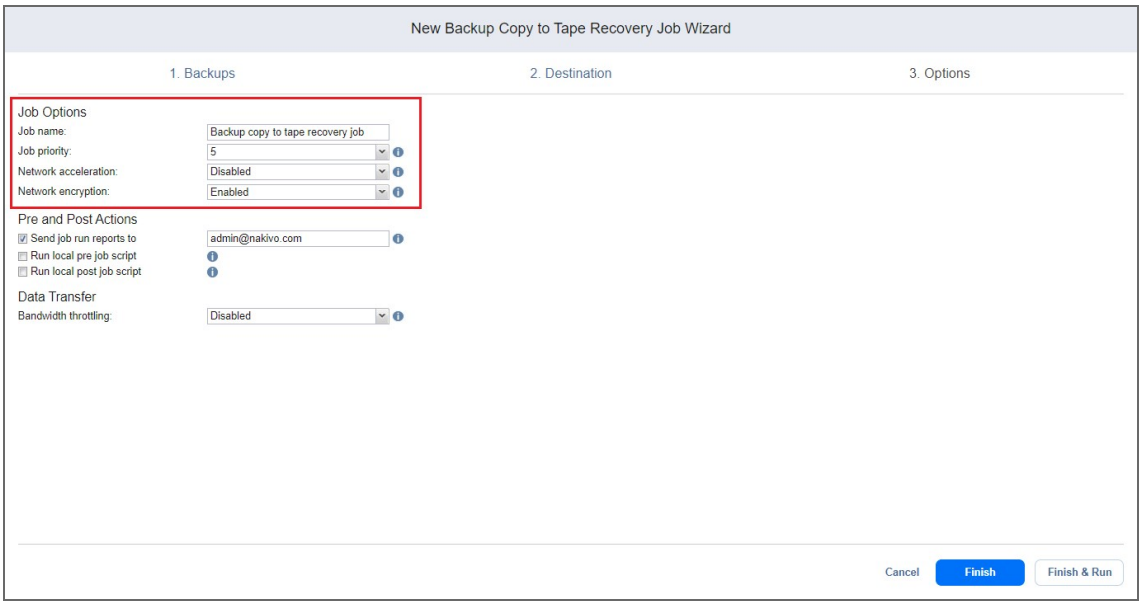
Specify the general options as follows:

1. **Job name:** Specify a name for the recovery job.
2. **Job priority**: Select a job priority level between 1 and 5, with 1 being the highest priority. Jobs with higher priority levels are prioritized by Transporters during job processing.

   > **Note**
   > This option is only available in the Enterprise, Enterprise Essentials, Enterprise Plus, MSP Enterprise, and MSP Enterprise Plus editions.

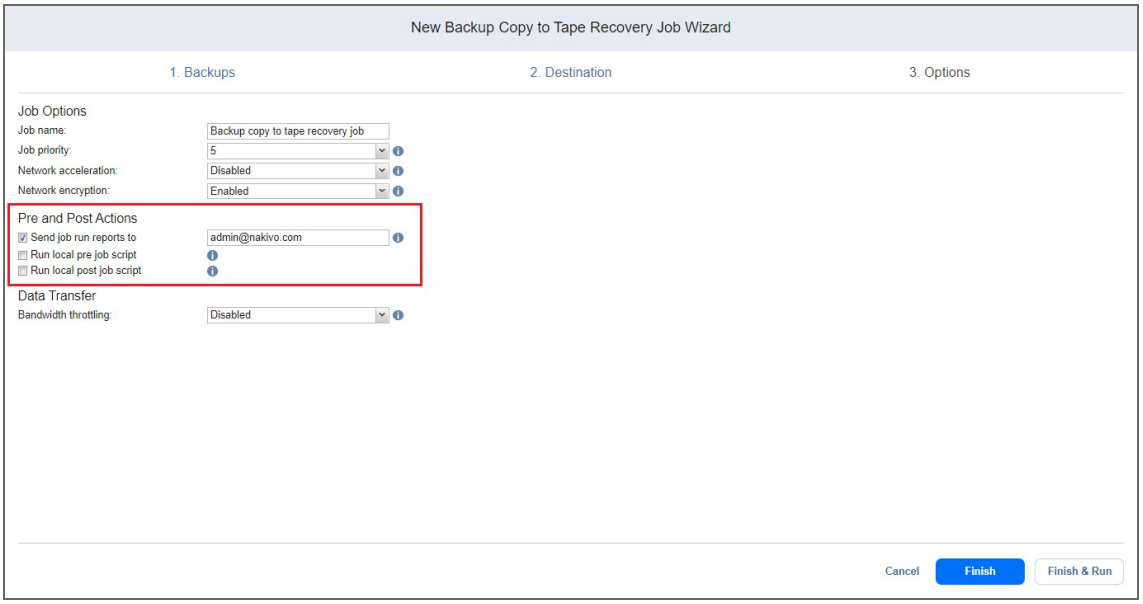
3. **Network acceleration**: With network acceleration enabled, NAKIVO Backup & Replication uses compression and traffic reduction techniques to speed up data transfer. Enable this option if you plan to recover VMs over WAN or slow LAN links.
4. **Network encryption**: With network encryption enabled, VM data is protected with AES 256 encryption while traveling over the network. Data encryption increases the backup time and CPU load on machines running Transporters. Select this option when recovering over WAN without a VPN

connection.



Pre and Post Actions

NAKIVO Backup & Replication allows you to run a script before recovery begins (a pre-job script) and after the recovery of all backup copies in the job has been completed (a post-job script). The scripts can only be executed on the machine on which the Director is installed.



Setting a Pre-Job Script

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

1. Place a script file on the machine on which the **Director** is installed.
2. Select the **Run local pre job script** option and click the **settings** link.

3. Specify the following options in the dialog box that appears:
    - **Script path**: Specify a local path to the script on the machine where the **Director** is installed. A script interpreter should be specified.
      Example (Windows): *cmd.exe /c D:\script.bat*
      Example (Linux): *bash/root/script.sh*
    - **Job behavior**: Select one of the following job behaviors in relation to script completion:
        - **Wait for the script to finish**: With this option selected, the 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 the backup at the same time.
    - **Error handling**: Select one of the following job behaviors in relation to script failure:
        - **Continue the job on script failure**: With this option selected, the job performs a backup even if the script fails.
        - **Fail the job on script failure**: With this option selected, if the script fails, the job fails, and the backup is not performed.

Setting a Post-Job Script

To run a script after the product has finished backing up, do the following:

1. Place a script file on the machine on which the **Director** is installed.
2. Select the **Run local post job script** option and click the **settings** link.
3. Specify the following options in the dialog box that appears:
    - **Script path**: Specify a local path to the script on the machine on which the **Director** is installed. A script interpreter should be specified.
      Example (Windows): *cmd.exe /c D:\script.bat*
      Example (Linux): *bash/root/script.sh*
    - **Job behavior**: Select one of the following job behaviors in relation to script completion:
        - **Wait for the script to finish**: With this option selected, the job remains in the “running” state until the script is executed.
        - **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**: Select one of the following job behaviors in relation to script failure:
        - **Continue the job on script failure**: With this option selected, script failure does not influence the status of the job.
        - **Fail the job on script failure**: With this option selected, if the script fails, the job status is set to “failed” even if the backup is successful.

Email Notifications

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

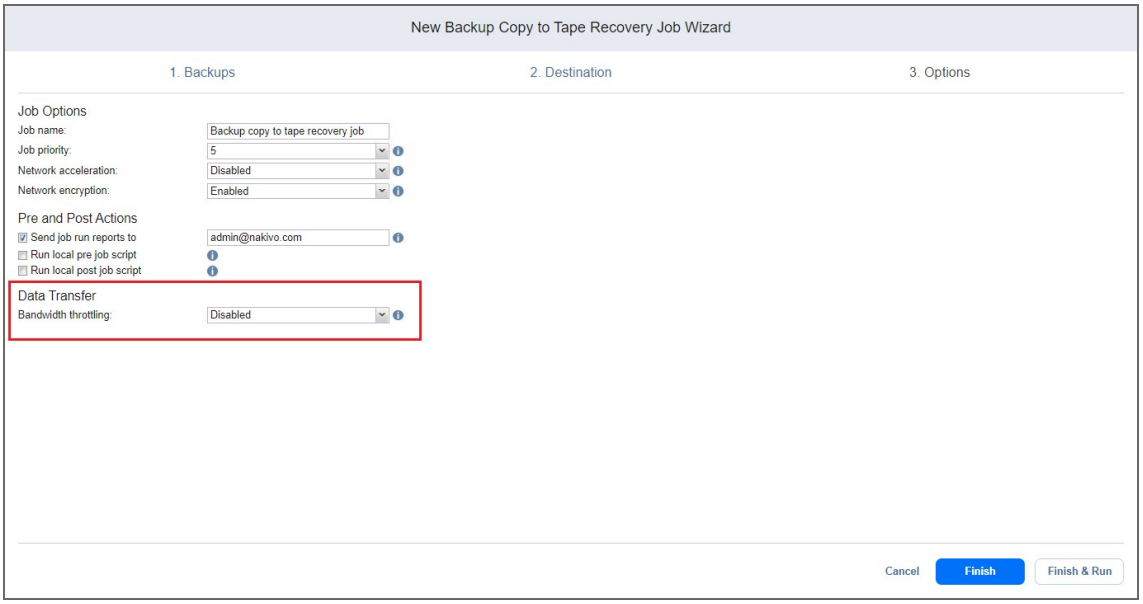
To enable this option, make sure that Email settings are configured.

To send email notifications, do the following:

1. In the **Pre and Post Actions** section, select **Send job run reports to**.
2. Specify one or more email addresses in the text box. Use semicolons to separate multiple email addresses.

Data Transfer

In this section, choose a transport mode for writing data, select which Transporter to use for reading data from the source backup, and set bandwidth throttling.



Bandwidth Throttling

Follow the steps below to set the speed of data transfer over the network for your recovery job:

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

> **Note**
>
> If bandwidth throttling is disabled for the current job, global bandwidth rules may still apply to your job. Refer to Bandwidth Throttling for details.

2. Click the **Settings** link that becomes available.

3. The **Job Bandwidth Rules** dialog box opens, displaying the list of available rules. You have the following options:
    - Create a new bandwidth rule for your recovery job:
        1. Click the **Create New Rule** button.
        2. The **New Bandwidth Rule** dialog box opens. Refer to Bandwidth Throttling for details on creating a bandwidth rule.
        3. Click **Save**.
    - **Activate an existing bandwidth rule for your job**: Select the checkbox to the left of the required bandwidth rule. To deactivate a bandwidth rule for your job, deselect the corresponding checkbox.
    - **Edit a bandwidth rule**: Click the **Edit** link for a bandwidth rule and modify it in the **Edit Bandwidth Rule** dialog box that opens.
    - **Disable a bandwidth rule**: Click the **Disable** link. The bandwidth rule is disabled for all jobs.
    - **Remove a bandwidth rule**: Click the **Remove** link and then click **Delete** to confirm your operation.

Completing the New Backup Copy to Tape Recovery Job Wizard

Click **Finish** or **Finish & Run** to complete job creation.

> **Note**
> If you click **Finish & Run**, you will have to define the scope of your job. Please refer to Running Jobs on Demand for details.

Multi-Tenant Mode

This section covers the following topics:

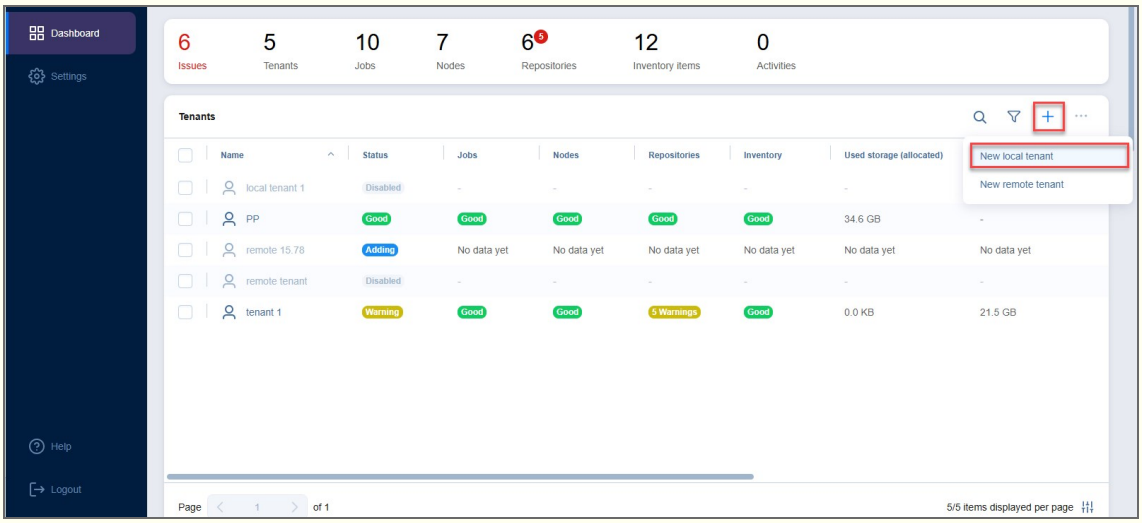
- "Creating a Local Tenant" below
- "Creating a Remote Tenant" on page 660
- "Using the MSP Console" on page 663
- "Using the MSP Dashboard" on page 676
- "Granting Self-Service Access" on page 687

Creating a Local Tenant

This section covers the topics describing the local tenant creation process in NAKIVO Backup & Replication. The data protection resources (Inventory items, Backup Repositories, and Nodes) of a local tenant account can only be added and edited by the Master tenant.

Notes

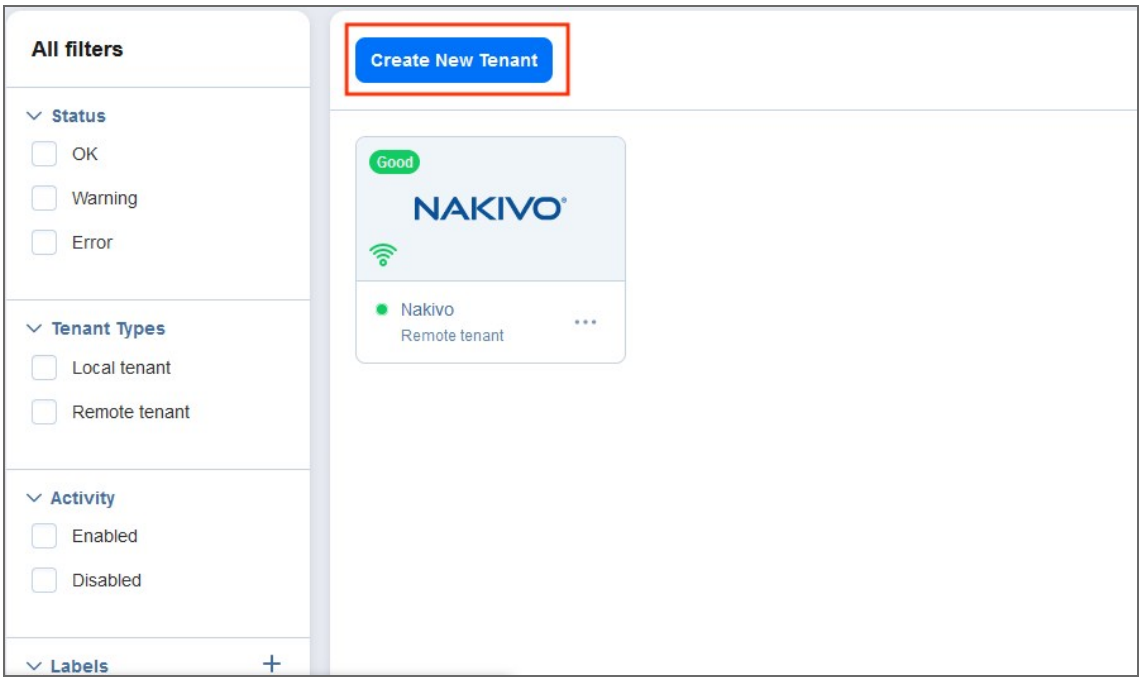
- You can create new local tenants by using the **MSP Console**. To learn how to do it, check here.



- The **MSP Console** is only available for users with an *MSP license*, *Beta instance*, *Promo license*, or *Trial license*. Users with multi-tenant instances of NAKIVO Backup & Replication without these license types have access to the MSP Dashboard.

To create a new local tenant, follow the steps below:

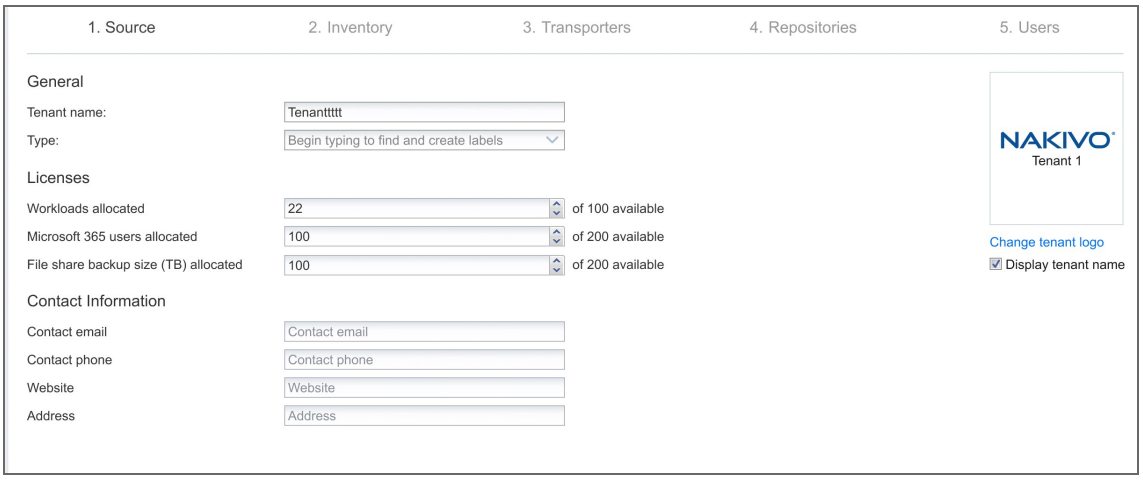
1. Log in to NAKIVO Backup & Replication as a Master Admin.
2. Click **Create New Tenant**.



3. In the popup, select **New local tenant**.
4. Complete the wizard as described in the topics below to finish the tenant creation process.
    - “Local Tenant Creation Wizard: Tenant” below
    - “Local Tenant Creation Wizard: Inventory” on page 655
    - “Local Tenant Creation Wizard: Nodes” on page 656
    - “Local Tenant Creation Wizard: Repositories” on page 657
    - “Local Tenant Creation Wizard: Users” on page 658
    - “Local Tenant Configuration” on page 659

Local Tenant Creation Wizard: Tenant

On this page of the wizard, you can provide a name for the local tenant, assign licenses to the local tenant, and enter or edit contact information for the local tenant. Additionally, the master tenant can enable VM Limitation for the new local tenant. When this option is enabled, the tenant cannot exceed the number of allocated VMs for the purpose of backup and replication. Tenants can see the number of allocated and used VMs in the licensing tab and in the job creation wizard.



Proceed as follows:

1. To add a tenant logo, click **Change tenant logo**, navigate to a new image, select it, and click **Open**. The uploaded image is resized and displayed on the right side of the page.
2. In the **Tenant name** field, enter a name for the local tenant. By default, the tenant name is displayed under the tenant logo. If you do not want the tenant name to be displayed, deselect the **Display tenant name** checkbox.
3. Optionally, in the **Labels** field, select the tags you want to assign to the tenant. Additionally, you can enter the name of the new label in the field and click **Create new label** to create and add it to the **Labels** field automatically.
4. In case the Trial or Subscription license is installed, do the following:
    a. In the **Workloads allocated** field, enter the number of workloads you want to assign to the local tenant.
    b. In the **Microsoft 365 users allocated** field, enter the number of Microsoft 365 users you want to assign to the local tenant.
5. In case a Perpetual license is installed, do the following:
    a. In the **Sockets allocated** field, enter the number of sockets you want to assign to the local tenant.
        i. Optionally, enable the **Limit number of protected VMs** option.

> **Note**
> In case the option is not available, make sure that the feature requirements are met.

        ii. Enter the number of protected VMs for the tenant.

> **Note**
> Even with VM limitation enabled, licenses are counted on a per-socket basis.

b. In the **Physical servers allocated** field, enter the number of physical server licenses you want to assign to the local tenant.
c. In the **Physical workstations allocated** field, enter the number of physical workstation licenses you want to assign to the local tenant.
d. In the **Microsoft 365 users** allocated field, enter the number of Microsoft 365 users you want to assign to the local tenant.
e. In the **File share backup size (TB) allocated** field, enter the amount of file share data (in terabytes) allocated for backup to assign to the local tenant. You can enter whole numbers, decimal values in increments of 0.5, or 0.
f. In the **Oracle databases** allocated field, enter the number of Oracle Database licenses you want to assign to the local tenant.

6. Optionally, in the **Contact email** field, enter the email address of the local tenant.
7. Optionally, in the **Contact phone** field, enter the phone number of the local tenant.
8. Optionally, in the **Website field**, enter the website URL of the local tenant.
9. Optionally, in the **Address** field, enter the address of the local tenant.
10. Click **Next** to proceed to the **Inventory** page or **Save** to save the changes.

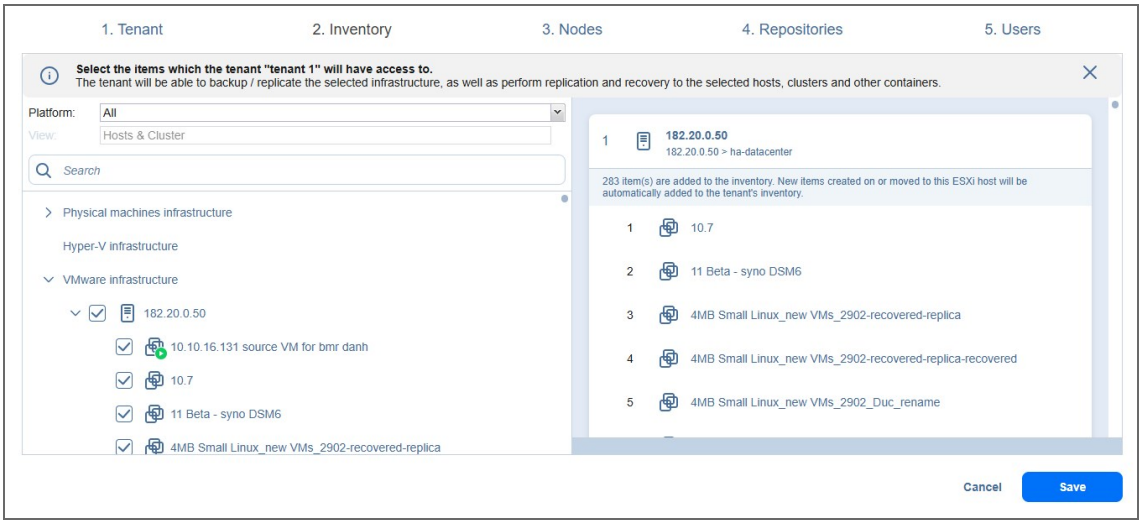
Local Tenant Creation Wizard: Inventory

On this page, you can assign inventory items to the local tenant or edit the items assigned to this tenant. Proceed as follows:

1. Choose the platform to display the items added to the inventory. **All** is selected by default.

> **Note**
> Items that are assigned to other tenants are visible, but cannot be selected.

2. Optionally, you can filter the Inventory tree by entering a string into the **Search** box. You can enter either a part or the entire name of the item.
3. Select the items you want to assign to the local tenant. The selected items appear in the right pane.



4. Click **Next** to proceed to the **Transporters** page or **Save** to save the changes.

Local Tenant Creation Wizard: Nodes

On this page of the wizard, you can assign/edit the **Transporters** that the local tenant will be able to use for backup, recovery, and replication jobs. Proceed as follows:

1. In the **Search** field, you can enter the name or part of the name of the **Transporter** to find the specific ones you need.

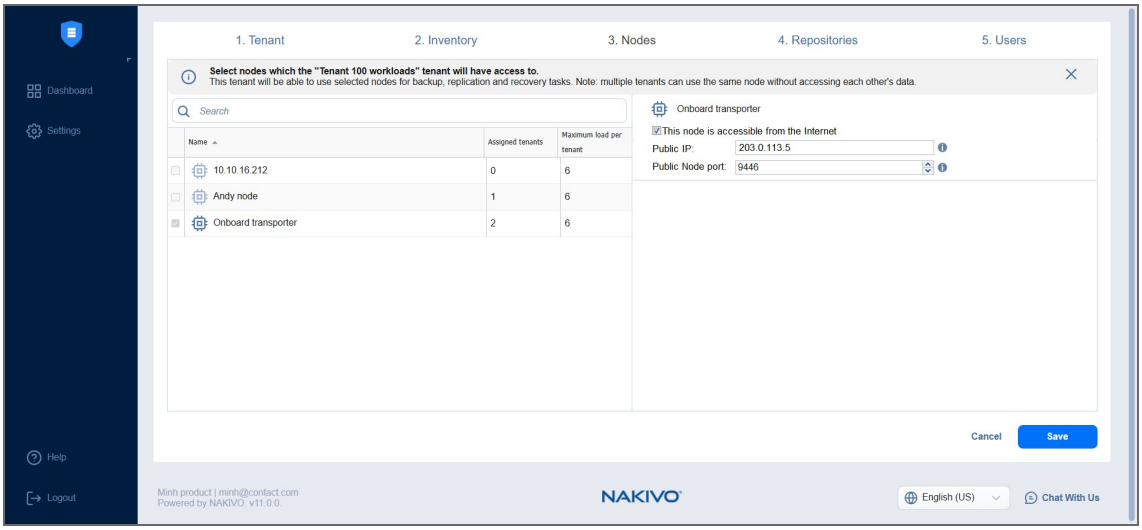
> **Notes**
>
> - When you assign an Inventory item with a dependent **Transporter** to the local tenant on the Inventory page of the wizard, that Transporter is selected automatically and cannot be deselected. If an Inventory item with a dependent **Transporter** was not assigned to the local tenant, that **Transporter** cannot be selected on this page.
> - The **Transporter** deployed in the virtual appliance cannot be assigned to multiple tenants.

2. On the left pane of the screen, you can select the **Transporters** to be assigned to the tenant. The following information is available:
   - **Name**: Name of the **Transporter**.
   - **Assigned tenants**: The number of tenants assigned to the **Transporter**. Multiple tenants can use the same **Transporter** without accessing each others data.
   - **Maximum load per tenant**: The maximum number of tasks that the **Transporter** can perform at the same time per each assigned tenant.
3. The selected **Transporters** appear in the right pane of the screen, where you can edit the following **Transporter** details:

- **This node is accessible from the Internet**: Select this checkbox to enable **Public IP** and **Public Node** port fields and establish **Direct Connect** communication via specified IP address and port accessible from the internet.
- **Public IP**: The external **MSP Transporter** IP address reachable from the internet and used for communication with the **Direct Connect Transporter**.
- **Public Node port**: The external **MSP Transporter** port reachable from the internet and used for communication with the **Direct Connect Transporter**. Ensure that the entered combination of the port and IP address is not used by other transporters.

> **Note**
> These fields are displayed for each **MSP Transporter** with enabled **Direct Connect** option.

4. Click **Next** to proceed to the **Repositories** page or **Save** to save the changes.



Local Tenant Creation Wizard: Repositories

On this page of the wizard, you can assign/edit Backup Repositories that the local tenant will be able to use for backup, recovery, and replication jobs. Note that a single repository cannot be used by multiple tenants. Proceed as follows:

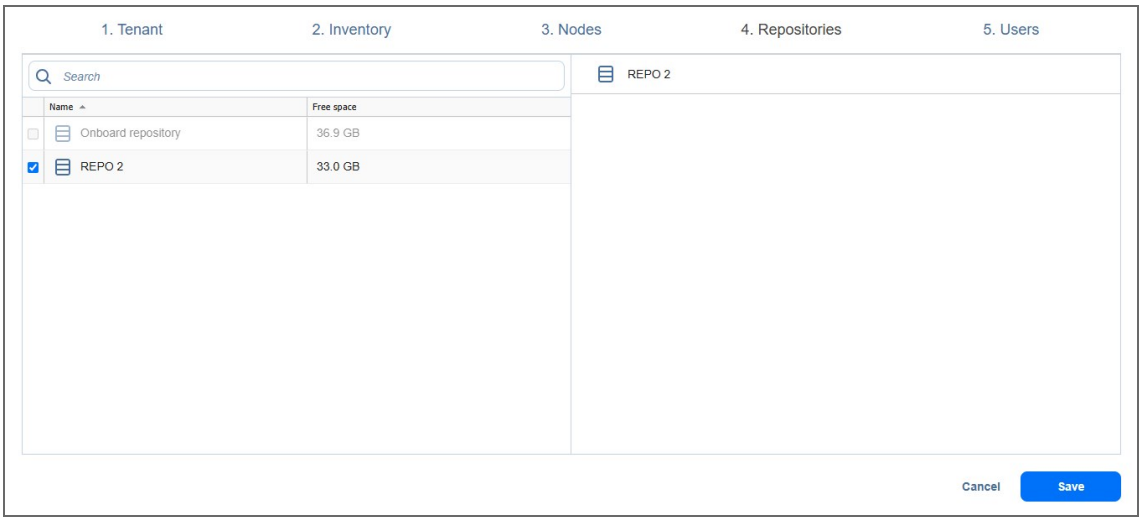
1. In the **Search** field, you can enter either a part or the entire name of the Backup Repository to find the specific ones you need.

Notes

- If the dependent Transporter was not chosen on the **Nodes** page of the wizard, the Backup Repositories assigned to this Transporter would not be available for selection.
- A federated repository can be assigned to a tenant. A federated repository member cannot be assigned.

2. On the left pane of the screen, you can select the Backup Repositories to be assigned to the local tenant. The following information is available
   - **Name**: Name of the Backup Repository.
   - **Free space**: The amount of free space available on the Backup Repository.

   The selected Backup Repositories appear in the right pane.



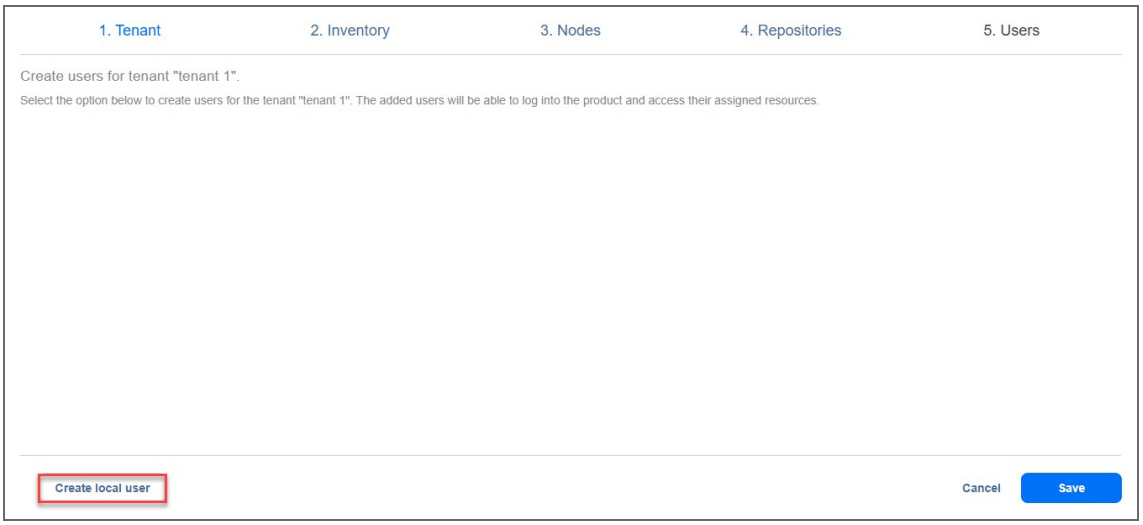
3. Click **Next** to proceed to the next page of the wizard or **Save** to save the changes.

Local Tenant Creation Wizard: Users

On this page of the wizard, you can create/edit local users or import Active Directory users for the tenant. The added users can use the product and have access to the allocated resources. Do the following:

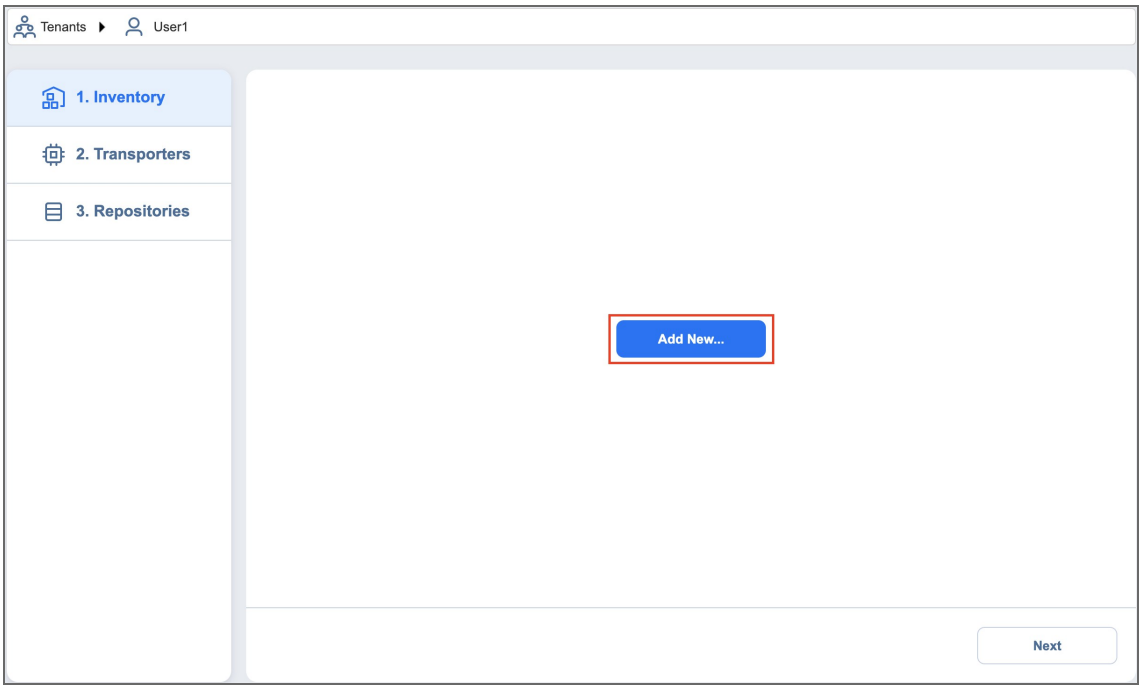
1. In the lower-left pane of the screen, click **Create local user** to create a new local user for the tenant.
2. If you have successfully configured AD integration, you can click **Add AD user** to import AD user for the tenant.
3. Once you're done, click **Finish** to complete the Local Tenant Creation Wizard or **Save** to save the

changes.



Local Tenant Configuration

After creating a new tenant, click the tenant to open the initial Tenant Configuration Wizard which will guide you through the tenant setup process. Refer to “First Steps with NAKIVO Backup & Replication” on page 229 for a description of the initial configuration wizard.

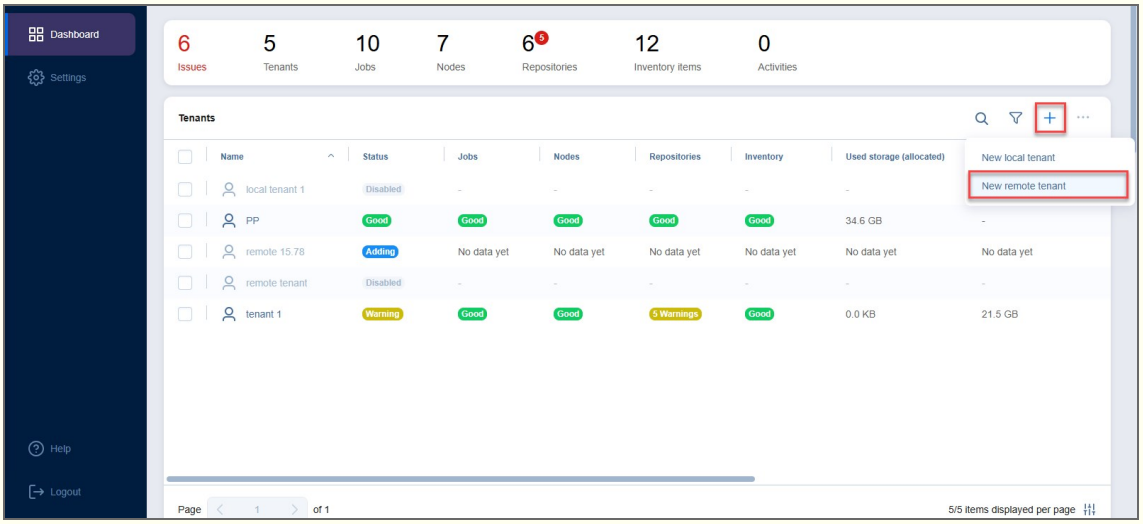


Creating a Remote Tenant

This section covers the topics describing the remote tenant creation process. Creating and configuring a Remote tenant allows a Master tenant to monitor a standalone instance of NAKIVO Backup & Replication. The remote tenant account retains the ability to manage the resources in the client's data protection infrastructure.

Notes

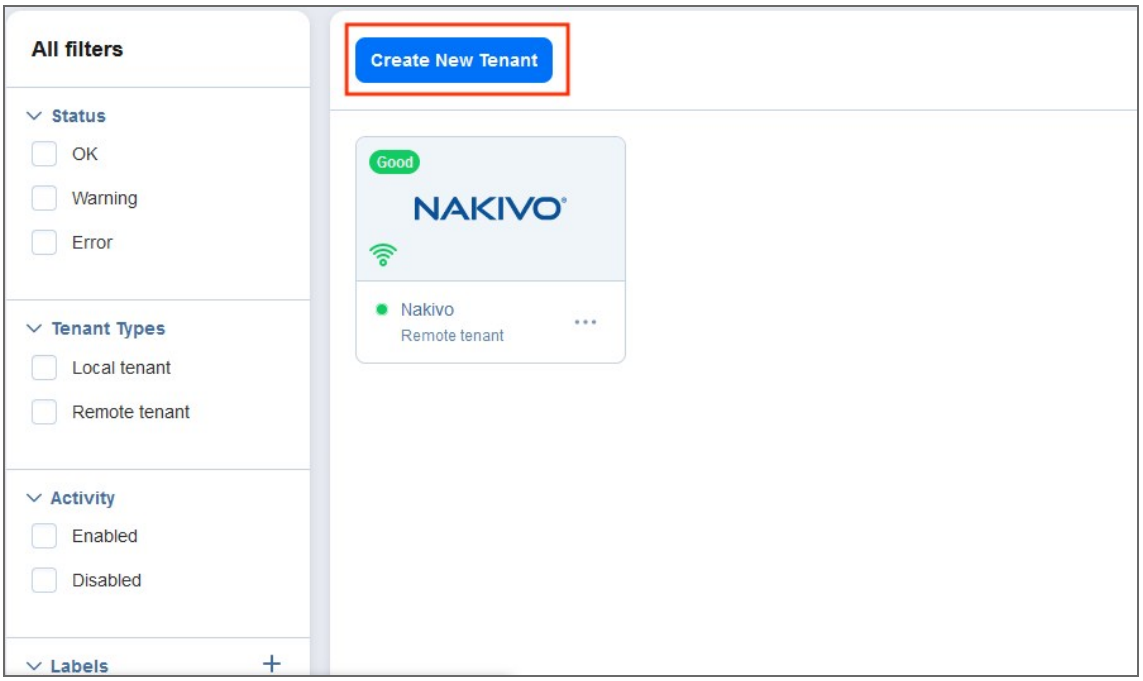
- You can create new remote tenants by using the **MSP Console**. To learn how to do it, refer to Creating New Tenants.



- The **MSP Console** is only available for users with an *MSP license*, *Beta instance*, *Promo license*, or *Trial license*. Users with multi-tenant instances of NAKIVO Backup & Replication without these license types have access to the MSP Dashboard.

To create a new remote tenant, follow the steps below:

1. Log in to NAKIVO Backup & Replication as a Master Admin.
2. Click **Create New Tenant**.



3. In the popup, select **New remote tenant**.
4. Complete the remote tenant creation process as described in the topics below:
    - “Remote Tenant Creation Wizard: Tenant” below
    - “Remote Tenant Creation Wizard: User” on the next page
    - “Remote Tenant Configuration” on page 663

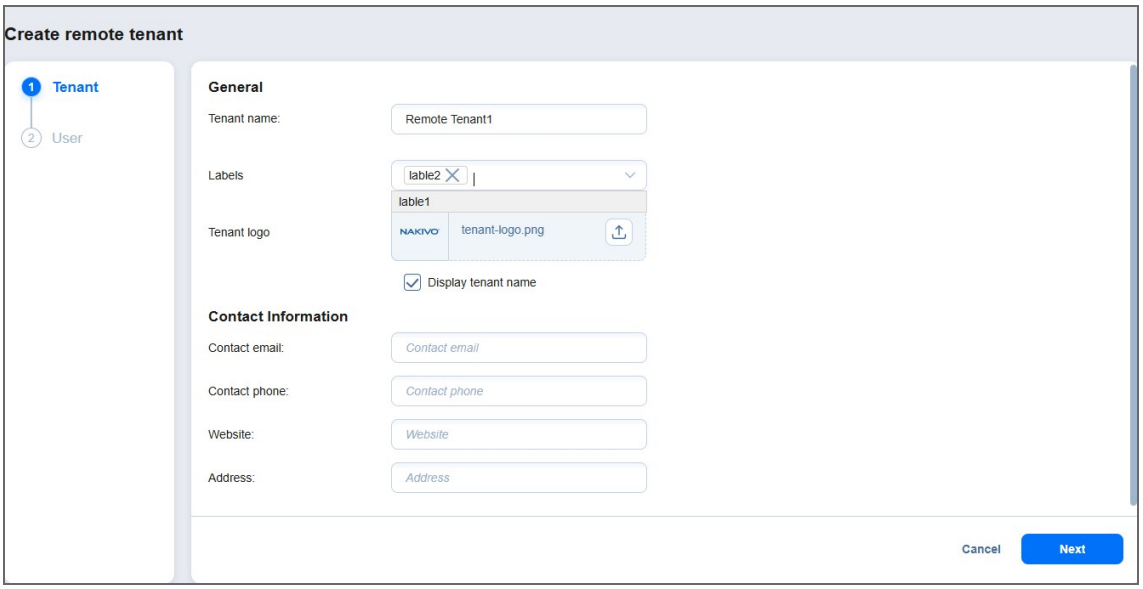
Remote Tenant Creation Wizard: Tenant

Complete the **Tenant** section of the **Remote Tenant Creation** wizard by configuring the following fields:

1. **Tenant name**: Specify the name of the remote tenant.
2. **Labels**: Optionally, you can create a new tag or assign existing tags to the remote tenant using the drop-down menu.
3. **Tenant logo**: Upload a logo to be displayed for the remote tenant in the Multi-tenancy Dashboard. The photo is automatically resized and a preview is generated.
4. **Display tenant name**: Enable this option if you want the tenant name to be displayed in the **Master Tenant Dashboard**.

Optionally, add contact information for the remote tenant by filling in the following fields:

- **Contact email**
- **Contact phone**
- **Website**
- **Address**

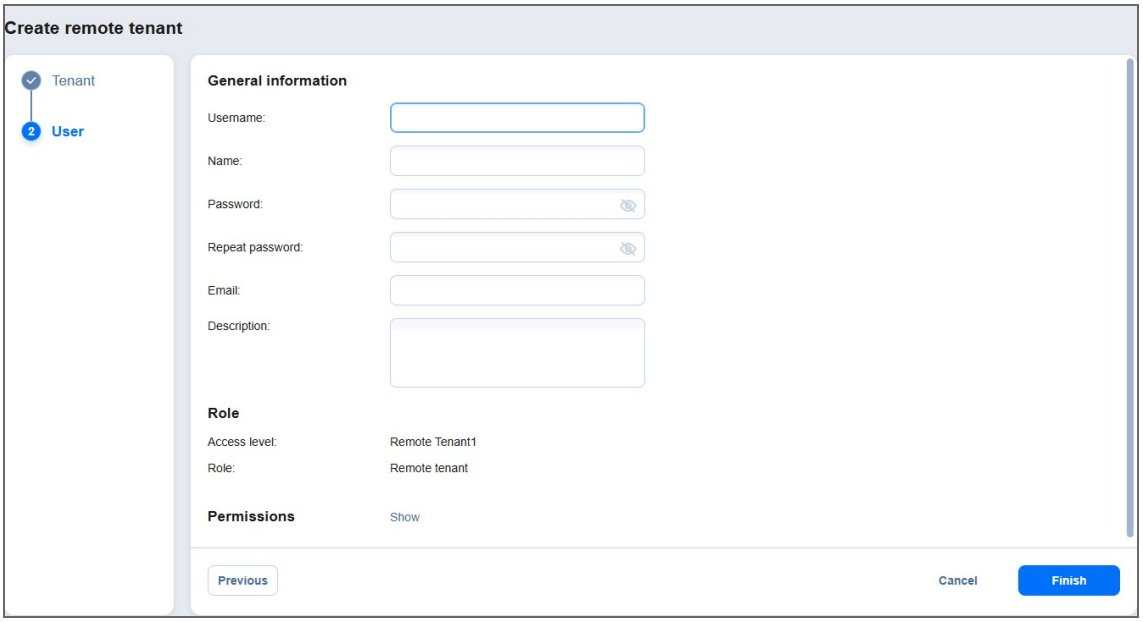


When you're done, click **Next** to move to the next page of the wizard.

Remote Tenant Creation Wizard: User

Complete the **User** section of the **Remote Tenant Creation** wizard by configuring the following fields:

1. **Username**: Specify a username for the remote tenant **User**.
2. **Name**: Specify the remote tenant display name.
3. **Password**: Create a password for this user and repeat it in the **Repeat Password** field below.
4. **Email**: Enter the user's email address.
5. **Description**: Optionally, you can add a description for this **User**.



The **Remote tenant** role and its permissions are added to the **User** automatically. Click **Finish** to confirm the creation of the remote tenant.

Remote Tenant Configuration

When a remote tenant is created, it is automatically added to the list of tenants in the **Master Tenant Dashboard**. To connect a remote tenant to your multi-tenant installation of NAKIVO Backup & Replication, follow the steps below:

1. Provide the remote tenant with the credentials created for the remote tenant user, as well as your hostname/IP address and **Director** port number (*4443* by default). In addition, you will need to open a separate listening port for communication with the remote tenant's instance (port *6702* is used by default). For more information on the required TCP ports, see the **MSP Console** section in "Feature Requirements" on page 91.
2. The remote tenant must go to **Settings** > **MSP** in their own instance of NAKIVO Backup & Replication and add the MSP using the above information, and then click **Add**.
3. A popup with certificate details appears. The remote tenant should click **Apply** to add the MSP to the **MSP** tab.
4. In your **Master Tenant Dashboard**, the remote tenant should now have a green **Connected** icon on the tenant card. Clicking on the remote tenant's name allows you to drill down and monitor their instance.

For more information on tenant-side **MSP Console** configuration, refer to "Adding an MSP" on page 317.

Using the MSP Console

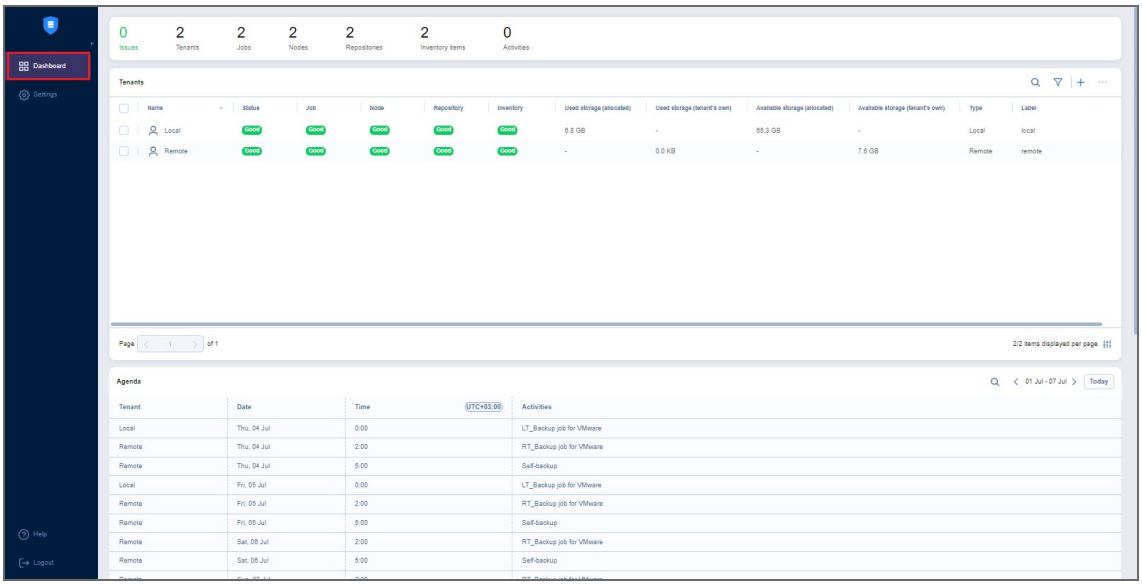
The **MSP Console** feature allows users to connect their standalone NAKIVO Backup & Replication instance to a managed service provider (MSP) and vice versa. Conversely, using the **MSP Console**, MSPs can create, add, and connect to local and remote tenants' environments to monitor a remote tenant's instance of NAKIVO Backup & Replication once a connection has been established on both sides.

Important

The **MSP Console** is only available for users with an *MSP license*, *Beta instance*, *Promo license*, or *Trial license*. Users with multi-tenant instances of NAKIVO Backup & Replication without these license types have access to the MSP Dashboard.

The **MSP Console** displays the key statistics for each local and remote tenant in the MSP's instance of NAKIVO Backup & Replication.

Using the **MSP Console**, MSPs can save time on tenants' management activities monitoring their statuses without drilling down into each tenant once a connection has been established on both sides.



See the topics below for more information:

- “MSP Architecture” on page 314
- “Opening MSP Console” below
- “Creating New Tenants” on the next page
- “Viewing Tenants Details” on page 666
- “Managing Tenants Details” on page 671
- “Using Search / Filter Options” on page 674

Opening MSP Console

As a managed service provider (MSP), you can use the **MSP Console** to monitor and manage all standalone instances of NAKIVO Backup & Replication connected to your instance of the solution and the data protection activities of the tenants you create.

> **Note**
> To use the MSP Console feature as an MSP (Master tenant), you must have NAKIVO Backup & Replication installed in the multi-tenant mode.

To establish a connection with a standalone instance of the solution, the MSP must first create a remote tenant account for the client and share the credentials with the client. The remote tenant must then use these to connect their instance to the MSP’s instance of NAKIVO Backup & Replication.

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

Important

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

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

Once the standalone NAKIVO Backup & Replication instance is added to the **MSP Console**, it appears as a remote tenant on the Tenants table and the Licensing > Tenants tab.

Note

Only users with an *MSP license*, *Beta instance*, *Promo license*, or *Trial license* can access the MSP Console and the Licensing > Tenants tab. This allows them to efficiently oversee all independent instances of NAKIVO Backup & Replication associated with a managed service provider (MSP) as well as local tenants from a unified interface, eliminating the need to navigate through individual tenants.

Users with multi-tenant instances of NAKIVO Backup & Replication without these license types only have access to the MSP Dashboard.

Creating New Tenants

The MSP Console allows MSPs to create, add, and connect to Local and Remote tenant environments.

1. To create a new tenant, in the top right corner of the Tenants table, click **+** (plus) > **Create New Tenant**.
2. In the dialog box that opens, select one of the options: **New local tenant** or **New remote tenant**, then proceed to create a remote or local tenant.



3. Complete the wizard as described in the topics below to finish the tenant creation process.

- "Local Tenant Creation Wizard: Tenant" on page 653
- "Local Tenant Creation Wizard: Inventory" on page 655
- "Local Tenant Creation Wizard: Nodes" on page 656
- "Local Tenant Creation Wizard: Repositories" on page 657
- "Local Tenant Creation Wizard: Users" on page 658
- "Local Tenant Configuration" on page 659

Notes

- You cannot assign any resources to a remote tenant.
- All the resources (inventory items, transporters, repositories) in the added single tenant NAKIVO Backup & Replication instance are treated as local tenant resources.
- Added standalone NAKIVO Backup & Replication instances are visible in the Tenants Dashboard as remote tenants. To learn how to create new remote tenants by using the **Tenants Dashboard**, refer to "Creating a Remote Tenant" on page 660.

Viewing Tenants Details

In the MSP Console, tenant data is displayed in the following widgets:

- Summary bar: Displays the number of tenant-related issues (errors and notifications), jobs, transporters, repositories, monitored items, and running activities. This widget includes search and filter functions that simplify finding specific issues by tenant, category, type, or date range.
- Tenants table: Allows MSPs to view the real-time status of all existing tenants, tenant-related jobs, nodes, repositories, inventory items, allocated resources, etc. You can also create new local and remote tenants, enable/disable existing tenants, and edit or delete them.
- Agenda widget: Allows MSPs to view information about the running and scheduled activities of all available tenants. By default, this widget displays the current week.
- Events widget: Lists all events, including errors and/or warnings of all available tenants, and is sorted by date by default.

Notes

- The Master tenant collects data from tenants every 1 minute and stores this data in the master tenant's database.
- The **Agenda** and **Events** widgets do not collect data from *Disabled*, *Inaccessible*, and *Disconnected* tenants.

See the sections below for more information on monitoring and managing tenants.

- "Summary Bar" on the next page
  - "Alarms" on the next page

- "Tenants Table" on the next page
- "Agenda" on page 670
- "Events" on page 670

Summary Bar



The **Summary bar** displays the following data collected from tenants: total number of issues (errors and notifications), tenants, jobs, nodes, repositories, monitored items, and running activities.
The data displayed is as follows:

- **Issues**: Total number of alarms/notifications related to all the available tenants (including the Master tenant). Clicking the number displays the Alarms dialog box.

  > **Note**
  > Alarms and notifications of *Disabled*, *Inaccessible*, and *Disconnected* tenants are not displayed, except for the reasons why the tenants are inaccessible/disconnected.

- **Tenants**: Total number of existing tenants (excluding Master tenant).
- **Jobs**: Total number of jobs inside tenants.
- **Nodes**: Total number of nodes inside tenants (including Master tenant). The nodes include allocated transporters, independently added/deployed transporters inside the tenant, and VMA and PMA.
- **Repositories**: Total number of tenants' repositories (including Master tenant). The repositories include both allocated and independently added/created repositories.
- **Inventory items**: Total number of tenants' inventory (including Master tenant). The inventory includes both allocated inventory and independently added inventory.
- **Activities**: Total number of running activities inside tenants (excluding Master tenant).

> **Notes**
> - Data is not collected for *Disabled*, *Inaccessible*, and *Disconnected* tenants.
> - The small number inside the red circle displays the total number of alarms and notifications related to the corresponding data section: nodes, repositories, or inventory items of all tenants (including the Master tenant).

Alarms

The **Alarms** dialog box displays all the alarms and notifications relating to existing tenants.
Here, you can browse or search for a specific issue in the **Search** field.

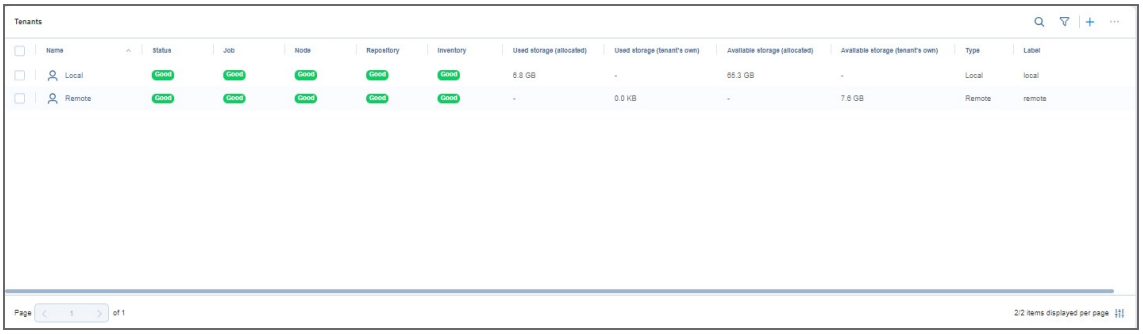
Optionally, click the **Filter** button to display issues based on their characteristics. The following filters are available:

- **Tenant**: Select to filter the alarms by tenant name. Multiple selections are supported.
- **Category:** Click to filter by selecting the following options inside the dropdown:
  - **Job**
  - **Inventory**
  - **Node**
  - **Repository**
  - **Tape**
  - **Monitoring**
  - **Users & Roles**
  - **Auto Update**
  - **Licensing**
- **Type**: Select to filter the alarms and notifications by type. The following types are available:
  - **Warning**
  - **Error**
  - **Dismissible**
  - **Non-dismissible**
- **Date**: Click to filter by date.

Once the filtering options are set, click **Apply** to start filtering.

Click the **Cancel** or **X** button to close the dialog box without applying any changes.

Tenants Table



The **Tenants** table gives MSPs an overview of all local and remote tenants:

- **Name**: The icon (**active** or **disabled**) and the name of a tenant. Clicking the tenant's name opens the Overview dashboard of the corresponding tenant.

> **Note**
>
> *Disabled*, *Inaccessible*, and *Disconnected* tenants are not clickable.

- **Status**: The status of a tenant:
  - **Adding**: A remote tenant has been added to the MSP console but the first data collection has not been completed yet.
  - **Good**: The tenant is functional.
  - **Warning**: The tenant has notifications only.
  - **Error**: The tenant has alarms.
  - **Inaccessible**: The tenant is not functional.
  - **Disconnected**: The remote tenant is disconnected from the MSP. You cannot drill down to a disconnected tenant.
  - **Disabled**: The tenant is disabled. Disabled tenants still store all the data but cannot log in and perform any activities.

> **Note**
> If a tenant has multiple statuses, the priorities of displaying these statuses are: *Disabled* > *Inaccessible/Disconnected* > *Good/Warning/Error*.

- **Job**: The status of a tenant's jobs. Clicking the status indicator opens the **Job Overview Dashboard** of the corresponding tenant. *Adding*, *Disabled*, *Inaccessible*, and *Disconnected* tenants are not clickable.
- **Node**: The status of the tenant's nodes. Clicking the status indicator opens the **Nodes** tab of the corresponding tenant. *Adding*, *Disabled*, *Inaccessible*, and *Disconnected* tenants are not clickable.
- **Repository**: The status of a tenant's repositories. Clicking the status indicator opens the **Repositories** tab of the corresponding tenant . *Adding*, *Disabled*, *Inaccessible*, and *Disconnected* tenants are not clickable.
- **Inventory**: The status of the tenant's inventory. Clicking the status indicator opens the **Inventory** tab of the corresponding tenant . *Adding*, *Disabled*, *Inaccessible*, and *Disconnected* tenants are not clickable.
- **Used storage (allocated)**: The amount of allocated storage by the Master tenant that has been used up by the tenant. For *Adding* tenants, **No data yet** is displayed.
- **Used storage (tenant's own)**: The amount of the tenant's own storage that has been used by the tenant. For *Adding* tenants, **No data yet** is displayed.
- **Available storage (allocated)**: The amount of allocated storage by the Master tenant that is available to the tenant. For *Adding* tenants, **No data yet** is displayed.
- **Available storage (tenant's own)**: The amount of the tenant's own storage that is currently available. For *Adding* tenants, **No data yet** is displayed.

> **Note**
> If the repositories are placed on the same same disk/share/etc., the free space metric can be incorrect.

- **Label**: The tenant's labels.
- **Type**: The tenant's connection type – **Local** or **Remote**.

Agenda

The **Agenda** widget allows MSPs to view information about all the running and scheduled activities of tenants (except for *Disabled*, *Inaccessible*, and *Disconnected* tenants).

Clicking the tenant name, date, time, or activity name opens the **Calendar dashboard** of the corresponding tenant.

The **Agenda** widget allows you to search for **Tenant** name on the currently displayed week.



Events

The **Events** widget gives an overview of all **Error** and **Warning** events for all tenants (except for *Disabled*, *Inaccessible*, and *Disconnected* tenants) and provides information about each event in the following columns:

- **Tenant**: Displays information about all tenants and their events. Clicking the tenant name opens the **Events** page of the corresponding tenant.
- **Event name**: Displays the icon and the message of the event.
- **Category**: Displays one of the following tenant event categories:
  - **Job**
  - **Inventory**
  - **Node**
  - **Repository**
  - **Tape**
  - **Monitoring**
  - **User & Roles**

    - **Auto Update**
    - **Licensing**
- **Initiated by**: Displays the name of the tenant that initiated the event.
- **Date**: Displays the date and time of the event.

You can browse or search for a specific event in the **Search** field. Search can be performed on the **Tenant** column only.

You can sort the events by clicking the respective name of the column.

To access filtering options, click the **Filter** icon in the top right corner. In the dialog box that opens, you can select one or several filtering criteria.

The following filtering options are available (multiple selections are supported):

- **Initiated by**: Allows you to filter by tenant name.
- **Event type**: Allows you to filter by the following options:
    - **Info**
    - **Error**
    - **Warning**
    - **Debug**
- **Date**: Click to select the date range.

Optionally, you can show/hide columns or modify the number of items per page in the **Events table configuration** table. In the lower right corner, click the controls icon. In the dialog box that opens, select/deselect checkboxes and click **Apply**.

Click **Reset Settings** to reset the configuration settings to the default.

Click the **Cancel** or **X** button to close the dialog box without applying any changes.

Managing Tenants Details

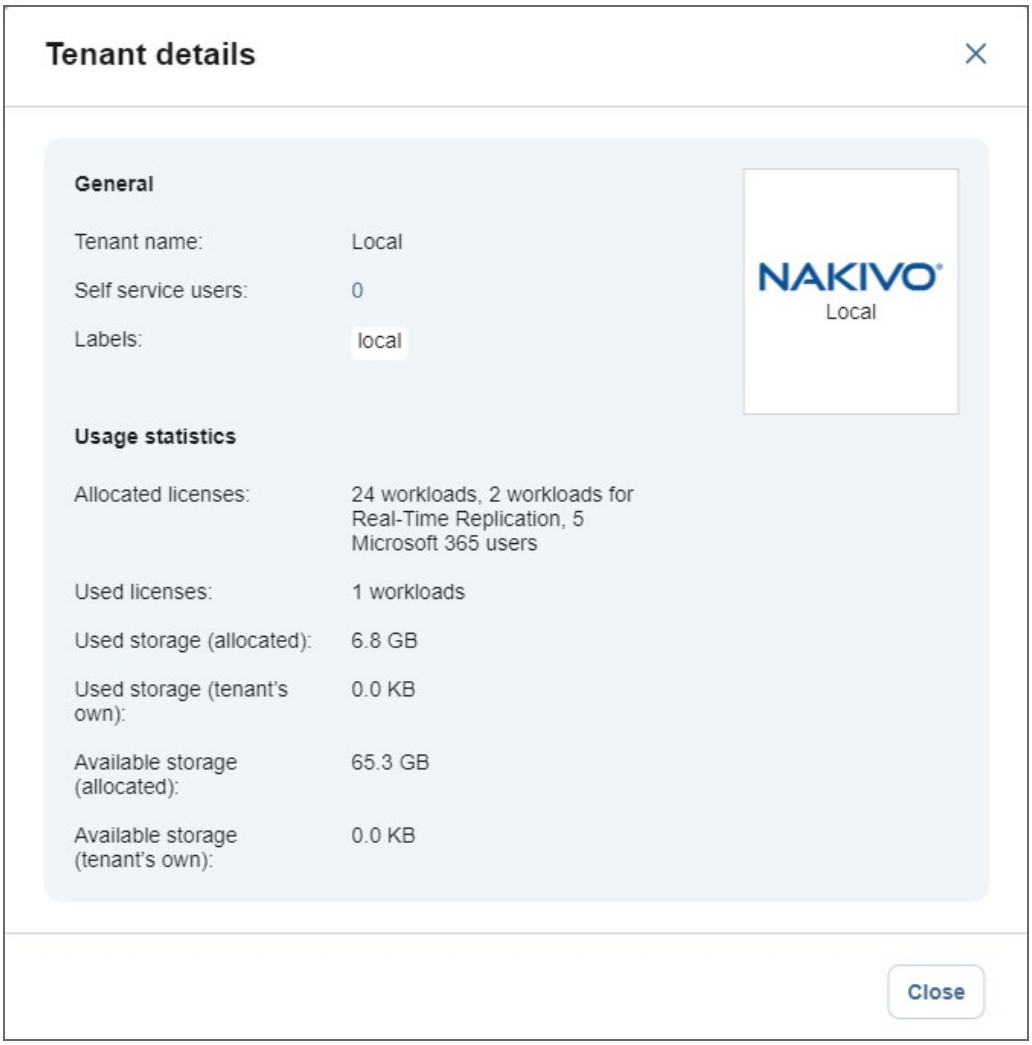
In the Tenants table, you can edit, disable, enable, or delete tenants.

1. To open the **Tenant details** popup window, hover over the corresponding row and click the ellipsis **Manage** button.

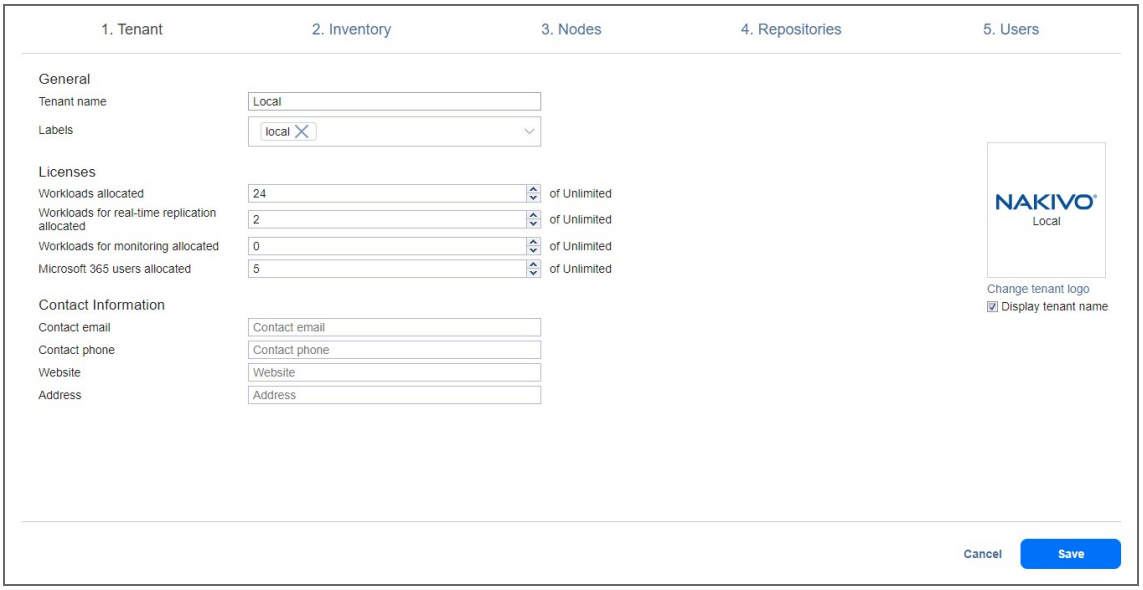


2. Then select the corresponding item in the list of actions:

- **View details**: Click to open the **Tenant details** popup.



- **Edit**: Click to open the tenant editor.



- **Enable/Disable**: Click to enable/disable the selected tenants.

> **Note**
>
> The option is not displayed if the tenant is disabled/enabled.

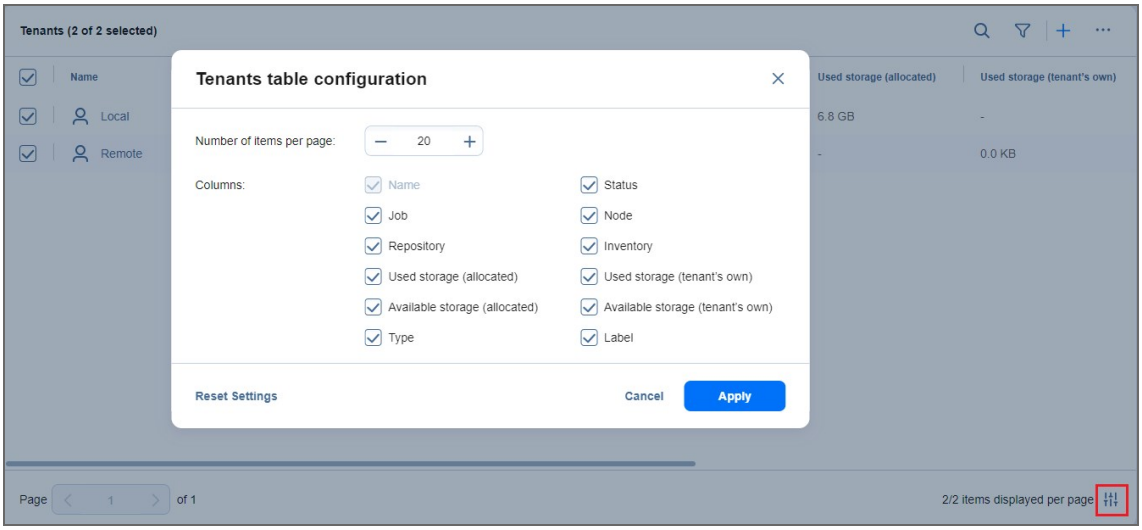
- **Delete**: Click to delete the selected tenants. You can select to delete the tenants and keep the tenant data or completely delete the tenant and all related data.

3. In the confirmation dialog box that opens, confirm or cancel the selected action.

These actions, except **Edit**, can also be done in bulk by selecting at least 2 items in the table or checking the box in the upper left pane to select all users and clicking the **...** (ellipsis) button.



Optionally, you can show/hide columns or modify the number of items per page in the **Tenants table configuration** table. In the lower right corner, click the controls icon. In the dialog box that opens, select/deselect checkboxes and click **Apply**.



Click **Reset Settings** to reset the configuration settings to the default.

Click the **Cancel** or **X** button to close the dialog box without applying any changes.

Using Search / Filter Options

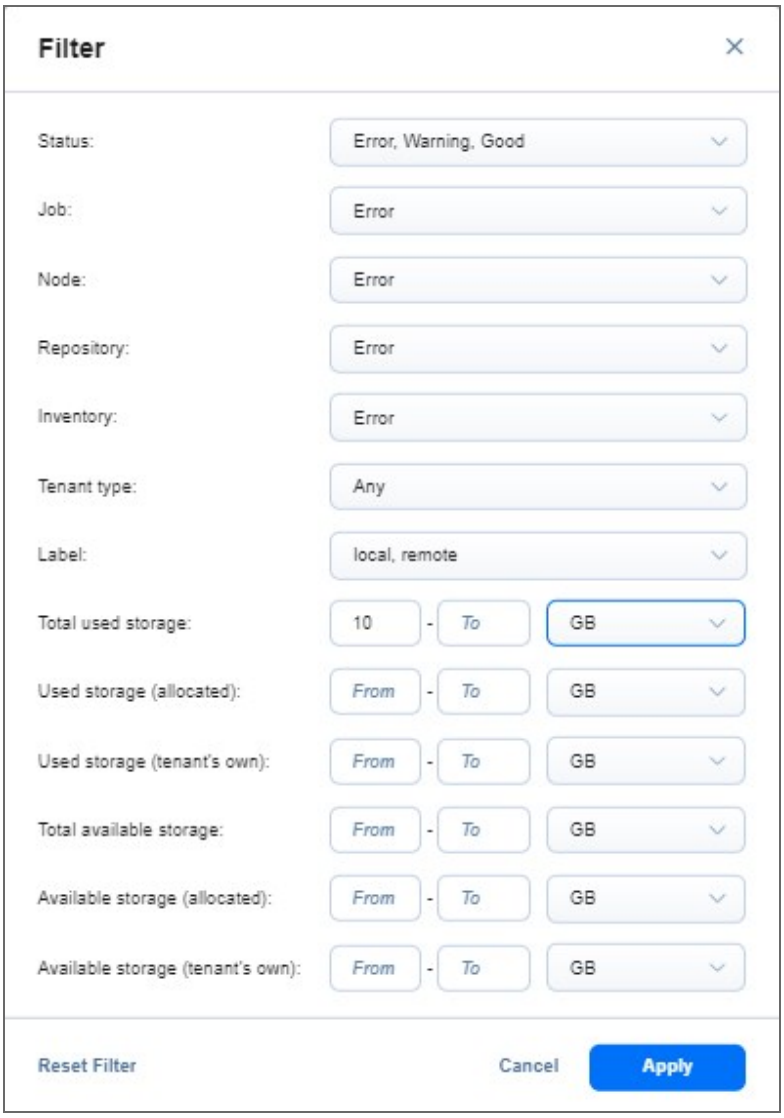
You can search or filter tenants from the Tenants table.

Use the **Search** field to browse or search for a specific tenant. Search can be performed on the **Name** column only.

To access filtering options, click the **Filter** icon in the top right corner. In the dialog box that opens, you can select one or several filtering criteria.

Click **Apply** to filter the search results.

Click the **Cancel** or **X** button to close the dialog box without applying any changes.



The following filtering options are available:

- **Status**: Allows you to filter by the following tenant status options:
  - **Good**
  - **Warning**
  - **Error**

    - **Inaccessible**
    - **Disconnected**
    - **Disabled**
- **Job**: Allows you to filter by the following job options:
    - **Good**
    - **Warning**
    - **Error**
- **Node**: Allows you to filter by the following node options:
    - **Good**
    - **Warning**
    - **Error**
- **Repository**: Allows you to filter by the following repository options:
    - **Good**
    - **Warning**
    - **Error**
- **Inventory**: Allows you to filter by the following inventory options:
    - **Good**
    - **Warning**
    - **Error**
- **Tenant type**: Allows you to filter by the following tenant type options:
    - **Any** (default)
    - **Local**
    - **Remote**
- **Label**: Allows you to filter by the following label options:
    - **Good**
    - **Warning**
    - **Error**
- **Total used storage**: From the dropdown, select to filter the items displayed by:
    - **MB**
    - **GB** (default)
    - **TB**
- **Used storage (allocated)**: From the dropdown, select to filter the items displayed by:
    - **MB**
    - **GB** (default)
    - **TB**
- **Used storage (tenant's own)**: From the dropdown, select to filter the items displayed by:

    - **MB**
    - **GB** (default)
    - **TB**
- **Total available storage**: From the dropdown, select to filter the items displayed by:
    - **MB**
    - **GB** (default)
    - **TB**
- **Available storage (allocated)**: From the dropdown, select to filter the items displayed by:
    - **MB**
    - **GB** (default)
    - **TB**
- **Available storage (tenant's own)**: From the dropdown, select to filter the items displayed by:
    - **MB**
    - **GB** (default)
    - **TB**

Using the MSP Dashboard

Tenant management for a multi-tenancy product with a non-MSP license is available from the MSP (Master Tenant) Dashboard.

This section covers the following topics:

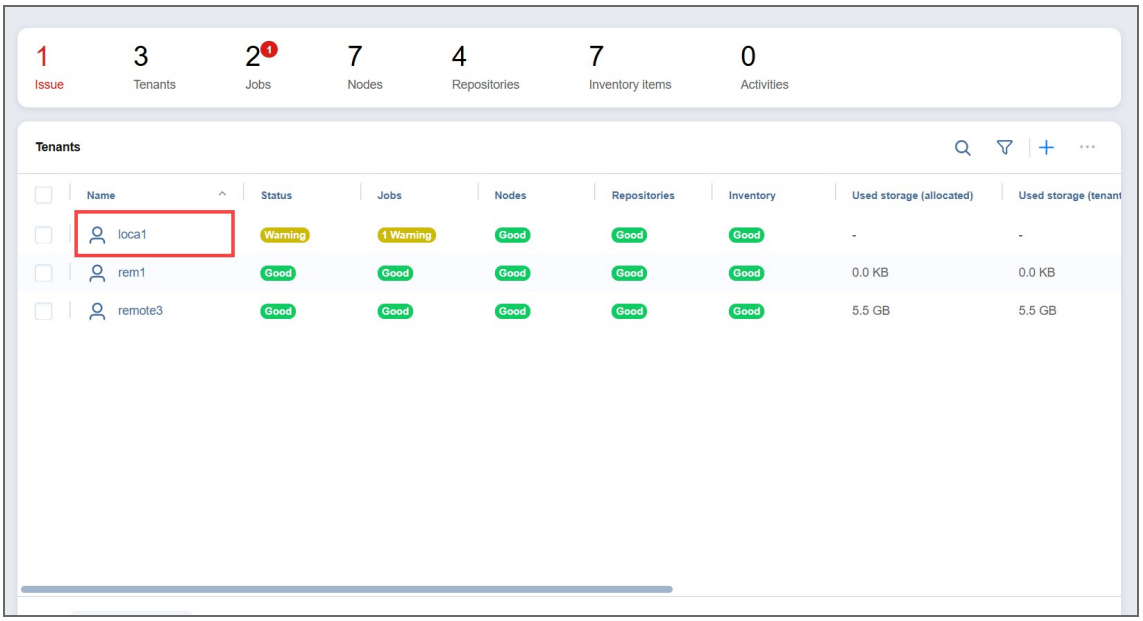
- "Opening Tenant Dashboard" on the next page
- "Viewing Tenant Information" on page 678
- "Using Filters" on page 680
- "Using Labels" on page 681
- "Editing Tenants" on page 684
- "Disabling Tenants" on page 685
- "Deleting Tenants" on page 686

Notes

- Users with an *MSP license*, *Beta instance*, *Promo license*, or *Trial license* can additionally access the MSP Console and the Licensing > Tenants tab. This allows them to efficiently oversee all independent instances of NAKIVO Backup & Replication associated with a managed service provider (MSP) as well as local tenants from a unified interface, eliminating the need to navigate through individual tenants.
- With the MSP Console feature, the standalone user is referred to as a remote tenant. Unlike local tenants in the traditional Multi-Tenancy workflow, remote tenants retain the ability to manage their resources in their data protection infrastructure.

Opening Tenant Dashboard

To work with a tenant's instance, you should open the tenant's **Dashboard**. For local tenants, this allows you to configure the tenant, create jobs and groups, and perform recovery. For remote tenants, this allows you to monitor their instance of NAKIVO Backup & Replication. To open a tenant's **Dashboard**, simply click the tenant.

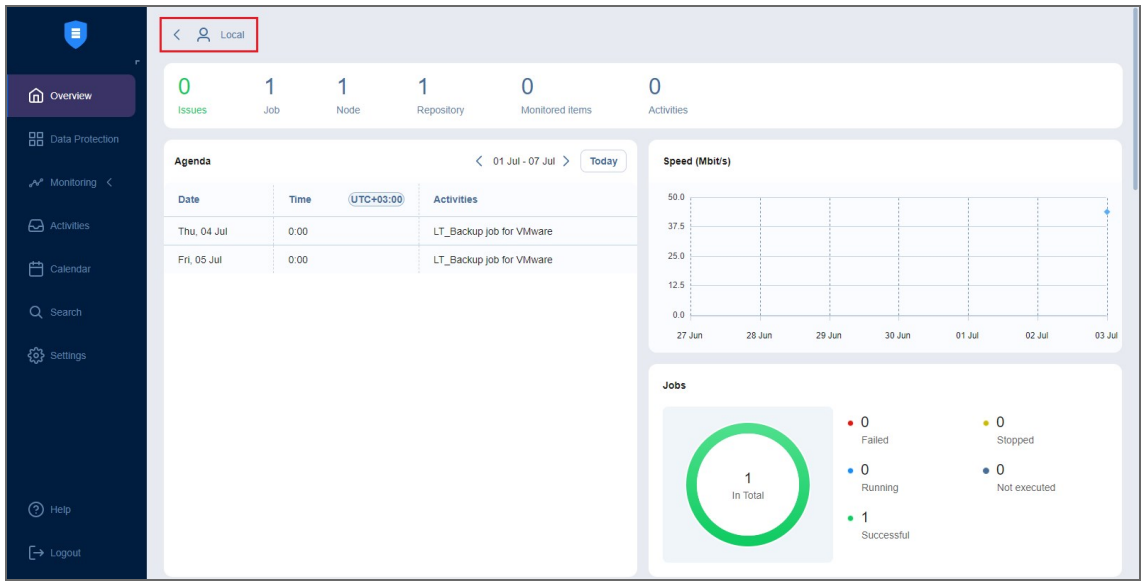


Note

Users with an *MSP license*, *Beta instance*, *Promo license*, or *Trial license* can additionally access the MSP Console and the Licensing > Tenants tab. This allows them to efficiently oversee all independent instances of NAKIVO Backup & Replication associated with a managed service provider (MSP) as well as local tenants from a unified interface, eliminating the need to navigate through individual tenants.

Returning to Master Admin Dashboard

To return to the **Master Tenant Dashboard** from a local tenant's instance, click **Tenants** in the navigation bar. To return from a remote tenant's instance, click the arrow to the left of the tenant name.

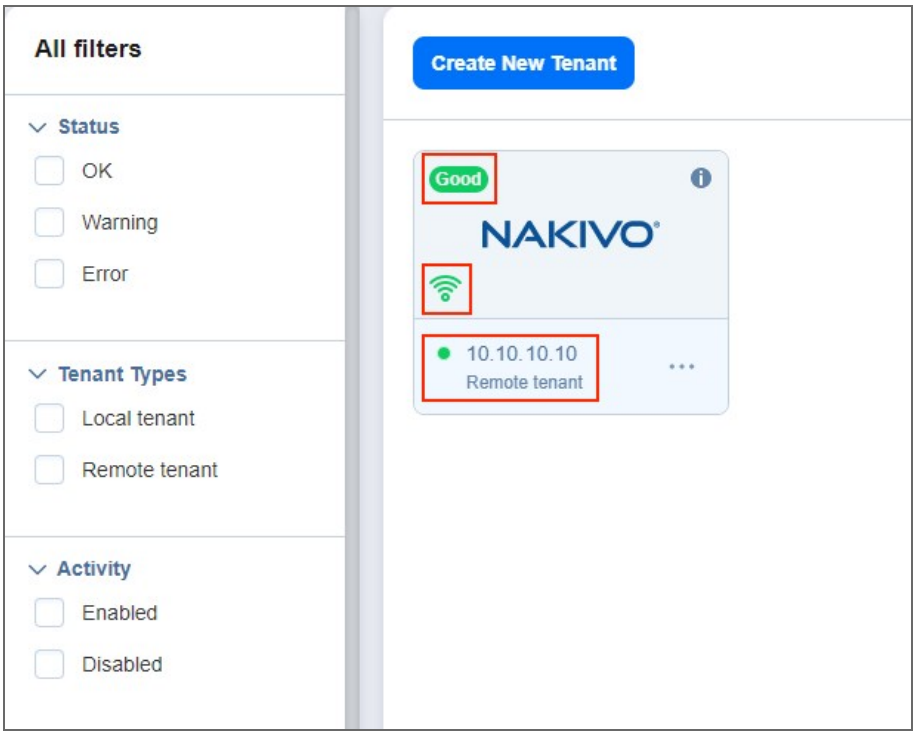


Viewing Tenant Information

On the **Master Tenant Dashboard**, you can view information about each tenant's instance of NAKIVO Backup & Replication. The information readily displayed on a tenant's card is as follows:

- **Tenant status**: The color and content of this indicator gives an overview of the tenant instance's alarms and notifications status. The given number reflects the number of alarms and/or notifications present at the remote tenant. A green **Good** indicator means there are no outstanding alarms and notifications. Other colors represent the following:
  - **Yellow**: There are outstanding notifications.
  - **Red**: There are outstanding alarms.
  - **Grey**: The tenant is disabled.

- **Connection** (remote tenants only): A green signal icon on a remote tenant's card indicates that a connection has been established between the remote tenant and Master Tenant instances; that is to say, a green signal icon will appear on the remote tenant's card once they have successfully added the MSP. A red signal icon means the connection could not be established or has been interrupted.
- **Accessibility**: A green circle icon next to the tenant's name indicates that the tenant is currently accessible by the Master Tenant.
- **Tenant name and type**: Lastly, the tenant card indicates the name and type of a given tenant.

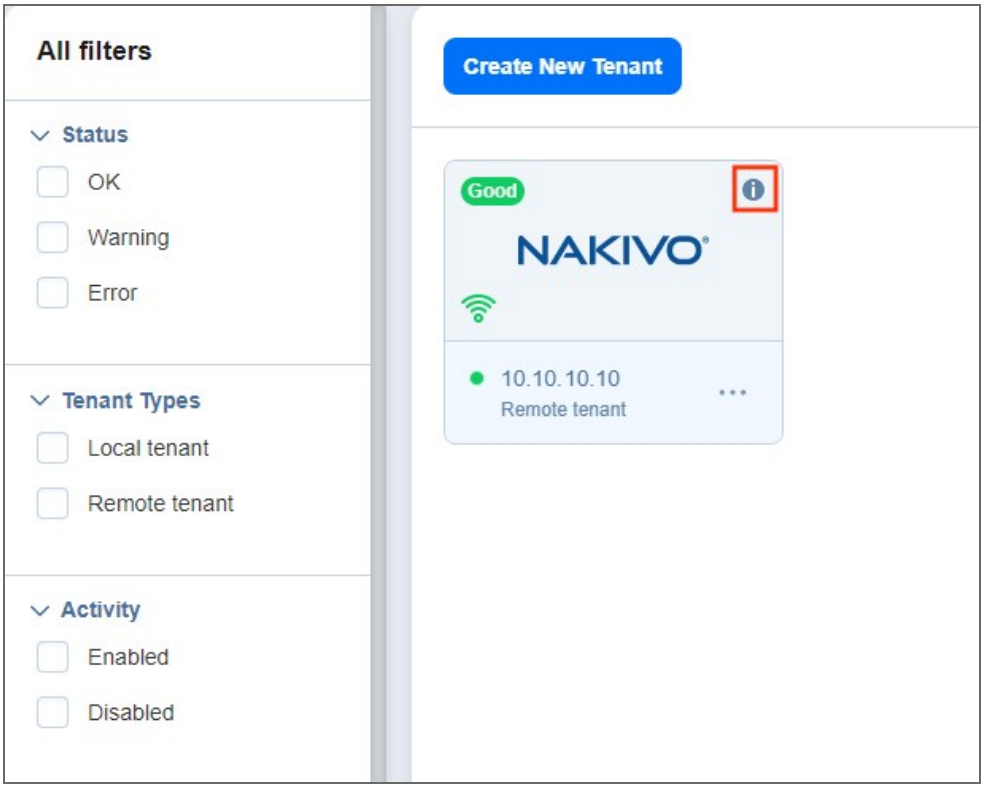


Note

Users with an *MSP license*, *Beta instance*, *Promo license*, or *Trial license* can additionally access the MSP Console and the Licensing > Tenants tab. This allows them to efficiently oversee all independent instances of NAKIVO Backup & Replication associated with a managed service provider (MSP) as well as local tenants from a unified interface, eliminating the need to navigate through individual tenants.

To learn more, refer to "Using the MSP Console" on page 663.

For more tenant information, hover over the tenant card and click on the **Info** button. A pop-up window opens with general tenant details and usage statistics.



Using Filters

- About Filters
- Applying Filters

About Filters

The **Master Tenant Dashboard** has 5 filter categories, which allow you to quickly display tenants based on their characteristics. The following filters are available:

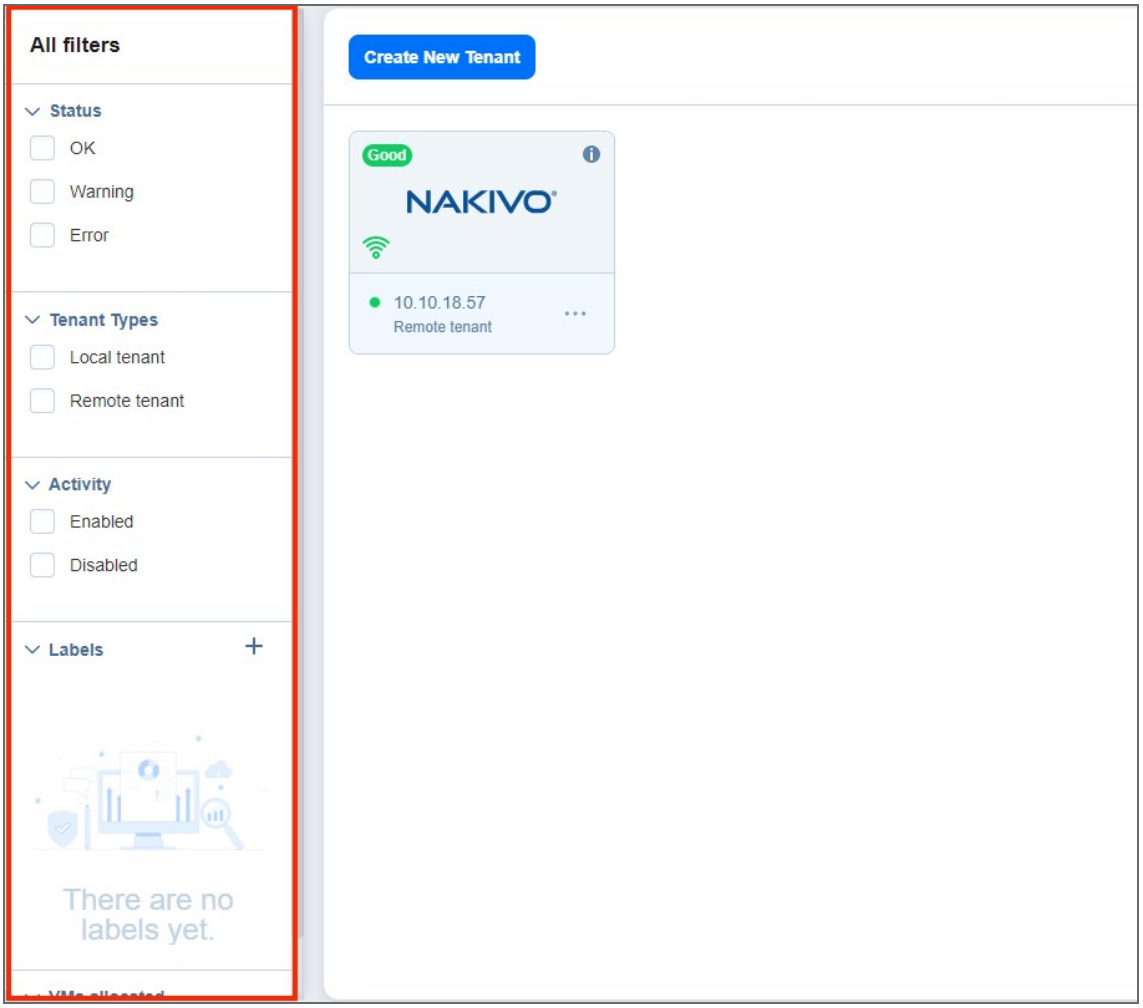
- **Status**:
  - **OK**: Displays tenants that have no errors and notifications
  - **Warning**: Displays only tenants that have notifications
  - **Error**: Displays only tenants that have errors
- **Type**:
  - **Local**: Displays only local tenants
  - **Remote**: Displays only remote tenants
- **Activity**:
  - **Enabled**: Displays only enabled tenants
  - **Disabled**: Displays only disabled tenants
- **Labels**: Filters tenants by labels assigned to them
- **VMs Allocated**: Filters tenants by the number of VMs allocated to them

Note

Users with an *MSP license*, *Beta instance*, *Promo license*, or *Trial license* can additionally access the MSP Console and the Licensing > Tenants tab. This allows them to efficiently oversee all independent instances of NAKIVO Backup & Replication associated with a managed service provider (MSP) as well as local tenants from a unified interface, eliminating the need to navigate through individual tenants.

Applying Filters

To apply a filter, check the box to the left of a filter name.



To dismiss a filter, simply uncheck the box to the left of the name of an active filter.

Using Labels

- About Labels
- Creating Labels
- Assigning Labels to Tenants

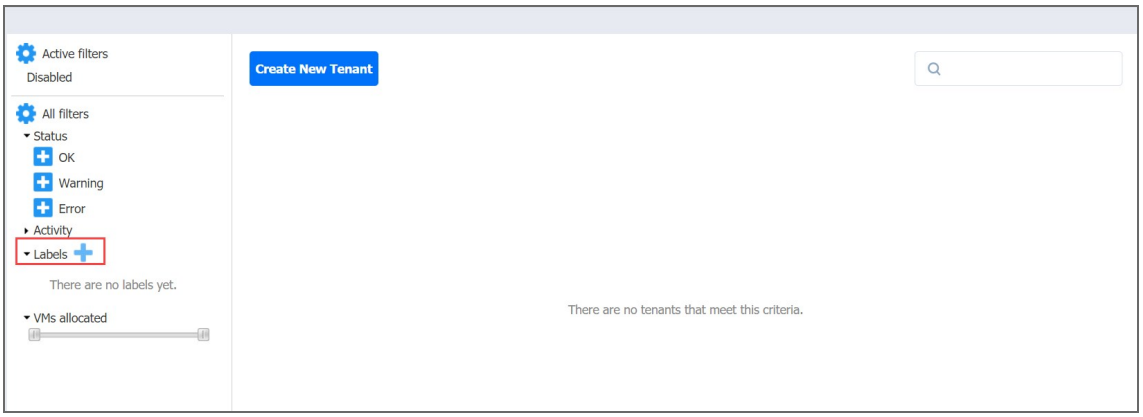
- Editing Label Names
- Deleting Label

About Labels

With NAKIVO Backup & Replication, you can create custom labels and assign them to tenants. Assigning a label to a tenant allows you to quickly sort existing tenants into different categories, such as location, SLA level, etc.

Creating Labels

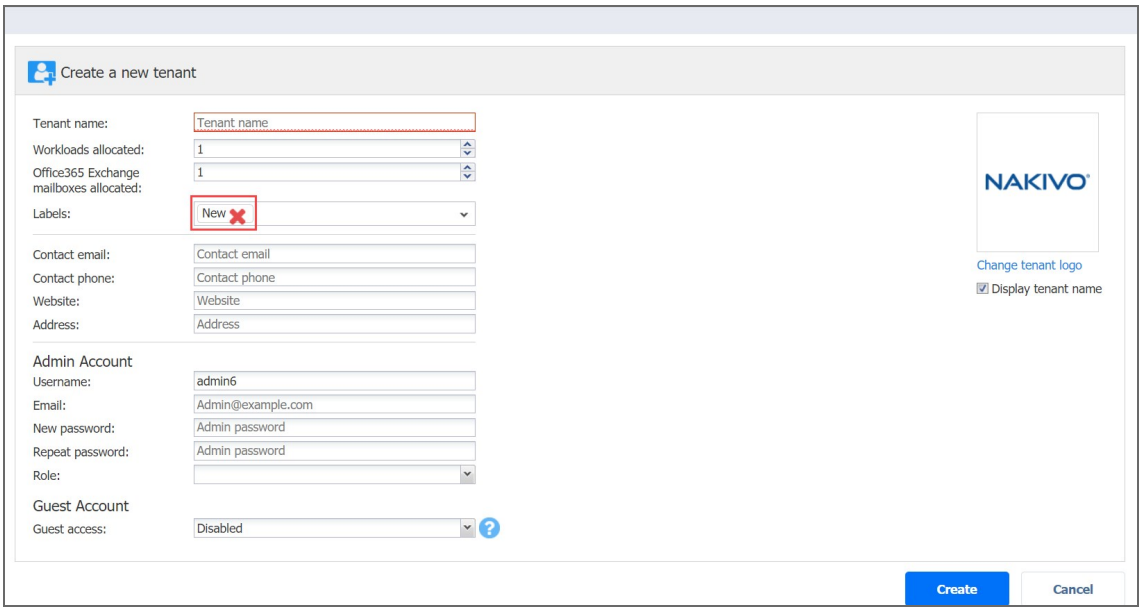
To create a new label, click the **Plus** icon next to **Labels** and enter a name for the new label, and press the **Enter** key.



You can also create a new label when creating a new tenant.

Assigning Labels to Tenants

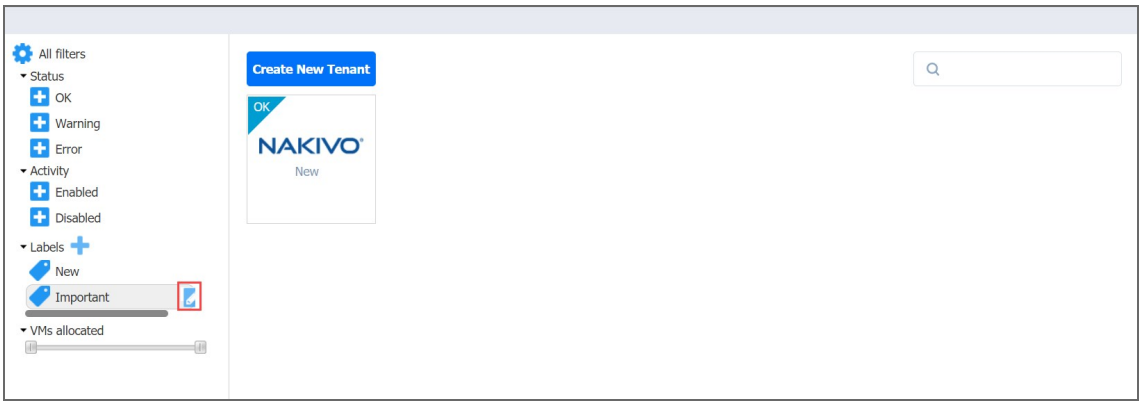
You can assign a label to a tenant either during the tenant creation or by editing the tenant.



Editing Label Names

To change a label name, do the following:

1. Hover over the label.
2. Click the **Edit** icon.



3. Enter the new label name and press the **Enter** key.

Deleting Labels

To permanently delete a label, do the following:

1. Hover the mouse pointer over a label.
2. Click the **Delete** icon.

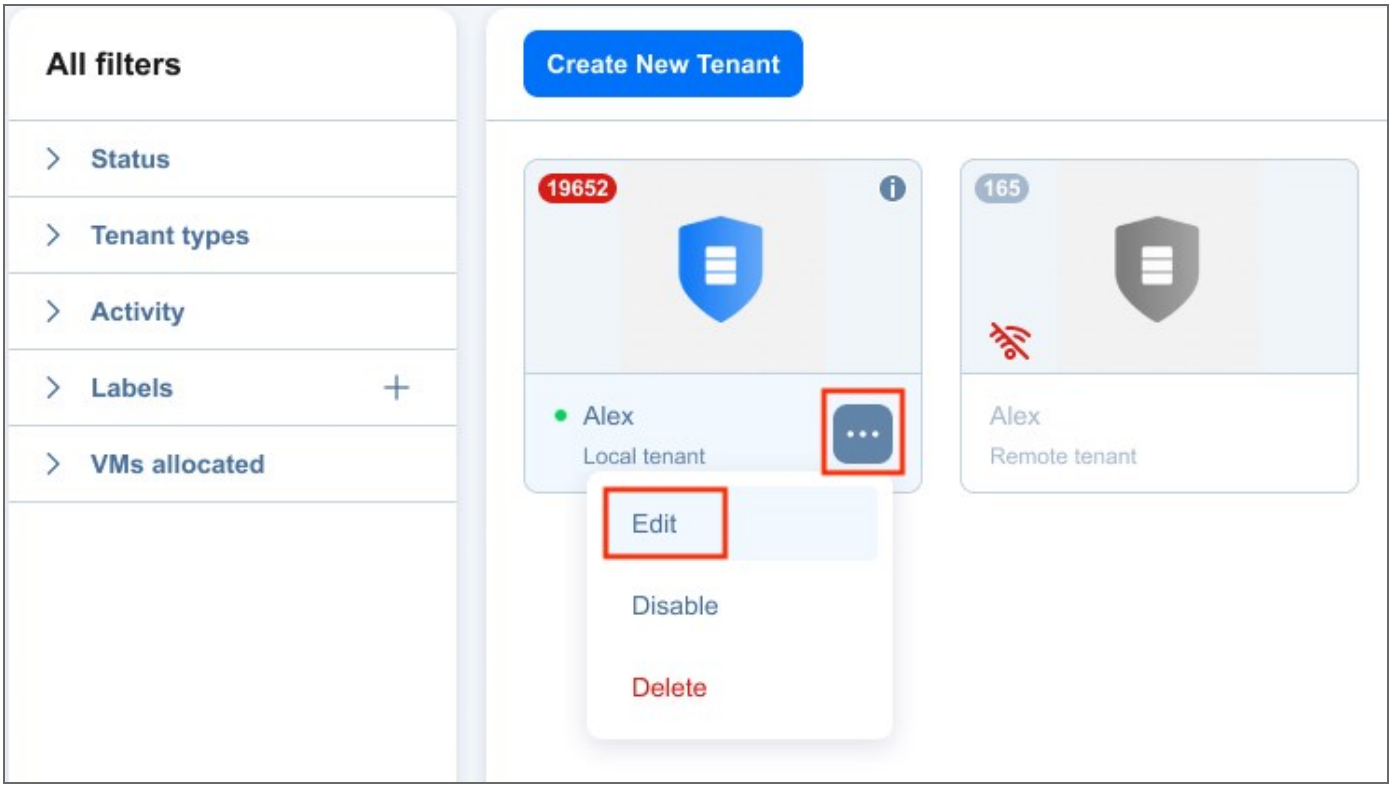
3. In the dialog box that opens, click **Delete** to confirm that you wish to permanently delete the label



Editing Tenants

To edit a tenant, do the following:

1. Hover over the tenant card and click the ellipsis **Manage** button, then click **Edit**.



2. In the **Edit** dialog box that opens, make the required changes and click **Save**.

Note

Users with an *MSP license*, *Beta instance*, *Promo license*, or *Trial license* can additionally access the MSP Console and the Licensing > Tenants tab. This allows them to efficiently oversee all independent instances of NAKIVO Backup & Replication associated with a managed service provider (MSP) as well as local tenants from a unified interface, eliminating the need to navigate through individual tenants.

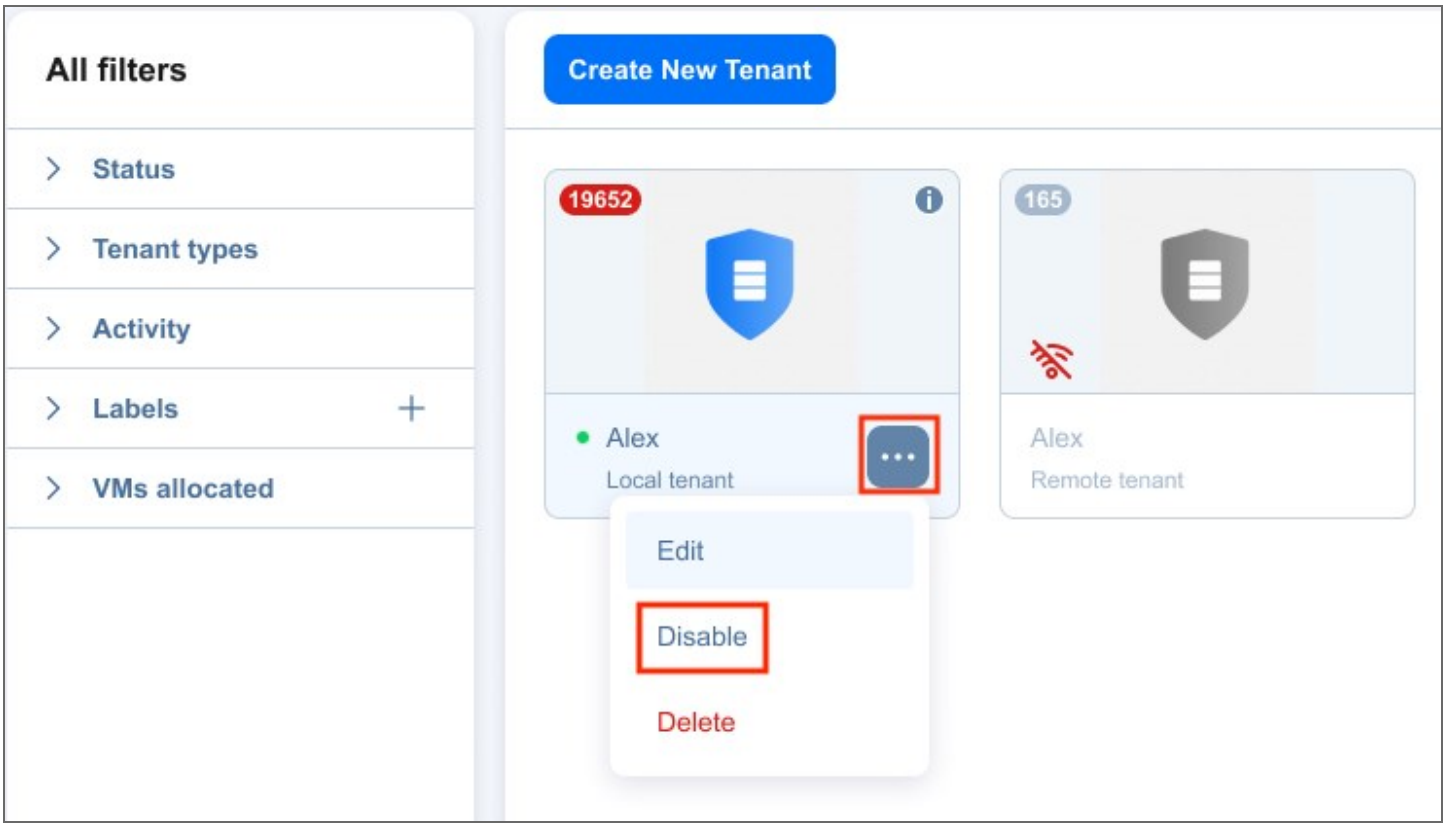
To learn more, refer to "Using the MSP Console" on page 663.

Disabling Tenants

In multi-tenant mode, Master Admin can disable a tenant to temporarily stop delivering backup, replication, and recovery services for that tenant. After disabling a tenant:

- Tenant admin and tenant guest will not be able to log in to the self-service interface. A message saying that the service has been disabled will be displayed after login attempts.
- Existing jobs will not be run on schedule.
- All currently running jobs will be allowed to complete.

To disable a tenant, hover over the tenant card and click the ellipsis **Manage** button, then click **Disable**.



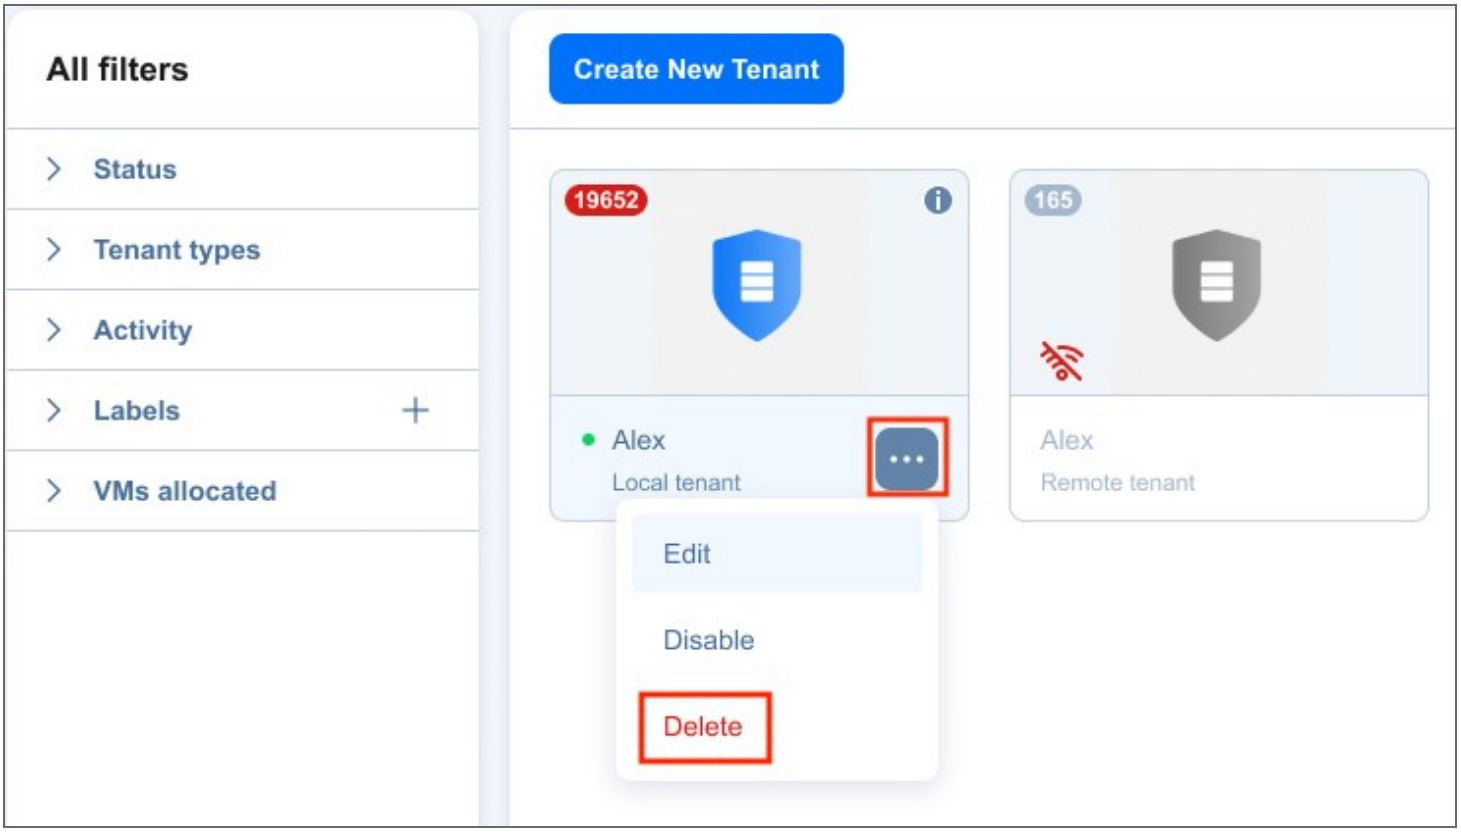
Note

Users with an *MSP license*, *Beta instance*, *Promo license*, or *Trial license* can additionally access the MSP Console and the Licensing > Tenants tab. This allows them to efficiently oversee all independent instances of NAKIVO Backup & Replication associated with a managed service provider (MSP) as well as local tenants from a unified interface, eliminating the need to navigate through individual tenants.

To learn more, refer to "Using the MSP Console" on page 663.

Deleting Tenants

To permanently delete a tenant from the product, hover over the tenant card and click the ellipsis **Manage** button, then click **Delete**.



The tenant will be permanently deleted from NAKIVO Backup & Replication.

Tenant Transporters are not uninstalled and the Tenant Backup Repositories are not removed.

Note

Users with an *MSP license*, *Beta instance*, *Promo license*, or *Trial license* can additionally access the MSP Console and the Licensing > Tenants tab. This allows them to efficiently oversee all independent instances of NAKIVO Backup & Replication associated with a managed service provider (MSP) as well as local tenants from a unified interface, eliminating the need to navigate through individual tenants.

To learn more, refer to "Using the MSP Console" on page 663.

Granting Self-Service Access

In the Multi-tenant mode, you can provide local tenants with access to their dashboards. By default, a tenant admin account is automatically created when you create a new local tenant. The tenant admin has full control over the product features inside the tenant dashboard (such as the ability to edit and update tenant inventory, transporters, and backup repositories, and create and manage jobs and groups). For each local tenant, one guest account can also be created. The tenant guest has limited permissions inside the tenant instance and can only generate job and group reports by default. To provide a local tenant with access to the self-service interface, send the following information to the tenant:

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