N2WS Backup & Recovery (CPM)

User Guide

V3.0.0b

Contents

1 Introduction to N2WS Backup & Recovery (CPM)

N2WS Backup & Recovery (CPM), known as N2WS, is an enterprise-class backup, recovery and disaster recovery solution for the Amazon Web Services (AWS). Designed from the ground up to support AWS, N2WS uses cloud native technologies (e.g. EBS snapshots) to provide unmatched backup and, more importantly, restore capabilities in AWS.
N2WS is sold as a service. When you register to use the service, you get permission to launch a virtual Amazon Machine Image (AMI) of an EC2 instance. Once you launch the instance, and after a short configuration process, you can start backing up your data using N2WS.
Using N2WS, you can create backup policies and schedules. Backup policies define what you want to back up (i.e. Backup Targets) as well as other parameters, such as:

- Frequency of backups
- Number of backup generations to maintain
- Whether to copy the backup data to other AWS regions, etc.
- Whether to back up a resource immediately

Backup targets can be of several different types, for example:

- EC2 instances (including some or all of the instance's EBS volumes)
- Independent EBS volumes (regardless of whether they are attached and to which instance)
- Amazon Relational Database Service (RDS) databases
- RDS Aurora clusters, except for Aurora Serverless
- Redshift clusters
- DynamoDB tables
- Elastic File System (EFS)

In addition to backup targets, you also define backup parameters, such as:

- In Windows achieving application consistency using Microsoft Volume Shadow Copy Service (VSS)
- Running backup scripts
- Number of retries in case of a failure

Schedules are used to define how you want to time the backups. You can define the following:

- A start and end time for the schedule, including time zone of data
- Backup frequency, e.g. every 15 minutes, every 4 hours, every day, etc.
- Days of the week to run the policy
- Special times to disable the policy

A policy can have one or more schedules associated with it. A schedule can be associated with one or more policies. As soon as you have an active policy defined with a schedule, backups will start automatically.

1.1 Purchasing N2WS on the AWS Marketplace

N2WS is available in several different editions which support different usage tiers of the solution (e.g. number of protected instances, number of AWS accounts supported, etc.) The

price for using the N2WS software is a fixed monthly price which varies between the different N2WS editions.
To see the different features for each edition, along with pricing and details, go to the N2W Software Web site. Once you subscribe to one of N2WS' editions, you can launch a N2WS Server instance and begin protecting your AWS environment. Only one N2WS Server per subscription will actually perform backup. If you run additional instances, they will only perform recovery operations (section 1.3.3).

1.1.1 Moving between N2WS Editions

If you are already subscribed and using one N2WS edition and want to move to another that better fits your needs, you need to perform the following steps:

Note:	Before proceeding, it is highly recommended to create a snapshot of your CPM data volume. You can delete that snapshot once your new N2WS Server is up and running. The data volume is typically named **N2WS – Data Volume**.

1. Terminate your existing N2WS instance. It is recommended that you do so while no backup is running.
2. Unsubscribe from your current N2WS edition. It is important since you will continue to be billed for that edition if you don't cancel your subscription. You will only be able to unsubscribe if you don't have any running instances of your old edition. You manage your subscriptions on the AWS Marketplace site in the Your Software page.
3. Subscribe to the new N2WS Edition and launch an instance. You need to launch the instance in the same Availability Zone (AZ) as the old one. If you want to launch your new N2WS Server in a different zone or region, you will need to create a snapshot of the data volume and either create the volume in another zone or copy the snapshot to another region and create the volume there.
4. During configuration, choose **Use Existing Data Volume** and select the existing data volume.
5. Once configuration completes, continue to work with your existing configuration with the new N2WS edition.

1.1.2 Downgrading

If you moved to a lower N2WS edition, you may find yourself in a situation where you exceed the resources your new edition allows. For example, you used N2WS Advanced Edition and you moved to N2WS Standard Edition, which allows fewer instances. N2WS will detect such a situation as a compliance issue, will cease to perform backups, display a message, and issue an alert detailing the problem.

To fix the problem:

- Move back to a N2WS edition that fits your current configuration, or
- Remove the excessive resources, e.g. remove users, AWS accounts or instances from policies.

Once the resources are back in line with the current edition, N2WS will automatically resume normal operations.

1.2 N2WS Architecture

The N2WS Server is a Linux based virtual appliance. It uses AWS APIs to access your AWS account. It allows managing snapshots of EBS volumes, RDS instances and clusters, Redshift clusters, and DynamoDB tables. Except in cases where the user chooses to install our Thin Backup Agent for Windows Servers, N2WS does not directly access your instances. Access is performed by the agent, or by a script that the user provides, which performs application quiescence.

N2WS consists of three parts, all of which reside on the N2WS virtual server:

- A database that holds your backup related metadata.
- A Web/Management server that manages metadata.
- A backup server that actually performs the backup operations. These components reside in the N2WS server.

The N2WS architecture is shown below. N2WS Server is an EC2 instance inside the cloud, but it also connects to the AWS infrastructure to manage the backup of other instances. N2WS does not need to communicate or interfere in any way with the operation of other instances. The only case where the N2WS server communicates directly with, and has software installed on, an instance, is when backing up Windows Servers for customers who want to use Microsoft VSS for application quiescing. If you wish to have VSS or script support for application quiescence, you will need to install the N2WS Thin Backup Agent. The agent will get its configuration from the N2WS server, using the HTTPS protocol.

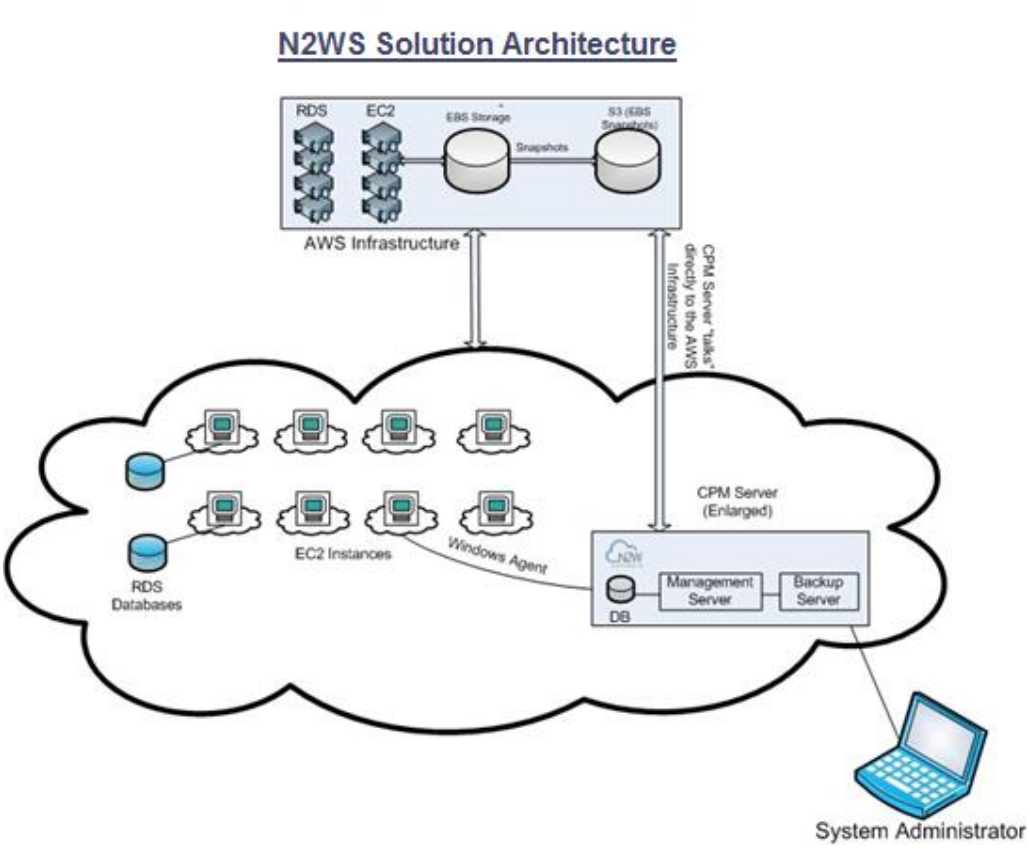

1.3 N2WS Server Instance

The N2WS instance is an EBS-based instance with two EBS volumes. One is the root device, and the other is the CPM data volume. All persistent data and configuration information reside on the data volume. From N2WS' perspective, the root device is dispensable. You can always terminate your N2WS instance and launch a new one, then using a short configuration process continue working with your existing data volume.

1.3.1 Root Volume

Although you have access to the N2WS Server instance via SSH, N2W Software expects the N2WS Server instance will be used as a virtual appliance. N2W Software expects you not to change the OS and not to start running additional products or services on the instance. If you do so and it affects N2WS, N2W Software will not be able to provide you with support. Our first requirement will be for you to launch a clean N2WS server.

Note:	Remember that all your changes in the OS will be wiped out as soon as you upgrade to a new release of N2WS, which will come in the form of a new image (AMI). If you need to install software to use with backup scripts (e.g. Oracle client) or you need to install a Linux OS security update, you can. N2W Software recommends that you consult N2W Software support before doing so.

1.3.2 Backing up the N2WS Server

N2WS server runs on an EBS-based instance. This means that you can stop and start it whenever you like. But if you create an image (AMI) of it and launch a new one with the system and data volume, you will find that the new server will not be fully functional. It will load and will allow you to perform recovery, but it will not continue performing backup as this is not the supported way to back up N2WS servers. What you need to do, is to back up only the data volume, and to launch a fresh N2WS server and connect it to a recovered data volume. See section 11.4.3.

1.3.3 N2WS Server with HTTP Proxy

N2WS needs connectivity to AWS endpoints to be able to use AWS APIs. This requires Internet connectivity. If you need N2WS to connect to the Internet via an HTTP Proxy, that is fully supported. During configuration you will be able to enable proxy use and enter all the required details and credentials: proxy address, port, user and password. User and password are optional and can be left empty if the proxy server does not require authentication. Once you configure proxy settings at the configuration stage, they will also be set for use in the main application.

The proxy setting can be modified at any time in the **Proxy** tab of N2WS Server Settings > **General Settings**. Select or clear **Enable Proxy**. If enabled, enter the requested proxy information.

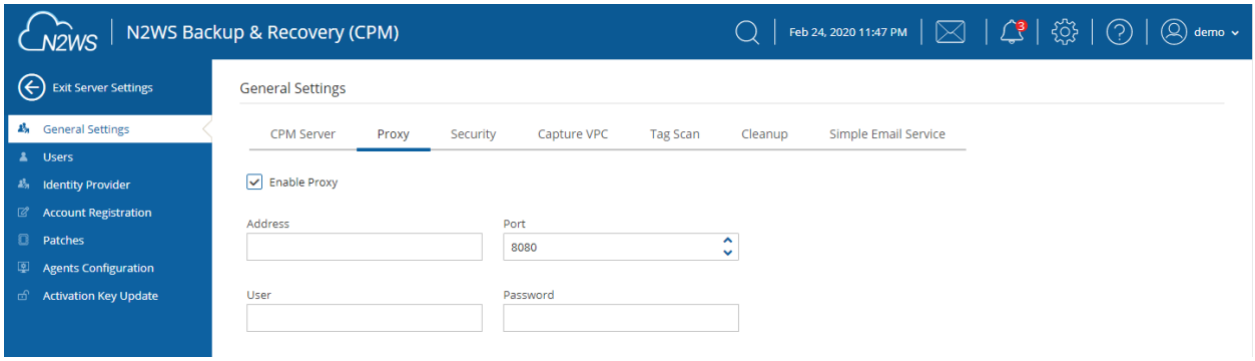

1.3.4 Multiple N2WS Servers

If you are trying to launch multiple N2WS servers of the same edition in the same account, you will find that from the second one on, no backup will be performed. Each such server will assume it is a temporary server for recovery purposes and will allow only recovery. Typically, one N2WS server should be enough to back up your entire EC2 environment. If you need more resources, you should upgrade to a higher edition of N2WS. If you do need to use more than one N2WS server in your account, contact N2W Software support.

1.3.5 Upgrading the N2WS Server Instance

At certain times, you may need to terminate the current N2WS Server instance and start a fresh one. The typical scenario is upgrading to a newer N2WS version or update the N2WS edition.

To upgrade/restart the N2WS Server Instance:

1. Terminate the old CPM instance, preferably while no backup, DR, or Cleanup is being performed.
2. Launch a new N2WS Server instance in the same region and AZ as the old one. You can launch the instance using the Your Marketplace Software page on the AWS web site.
3. To determine the AZ of the new instance, launch the instance using the EC2 console rather than using the 1-click option.
4. Wait until the old CPM instance is in the **terminated** state. When the new instance is in the **running** state, connect to it with a browser using HTTPS.
5. Approve the exception to the SSL certificate.
6. Choose **Use Existing Data Volume** in step #4, "Data Volume and Proxy".
7. Select your old data volume from the **Existing CPM Data Volume** list in step #5, "Server Configuration".
8. Select **Configure System** in step #6, "Register Your Account". CPM will automatically resume operations.
9. If you are using backup scripts that utilize SSH, you may need to login to the N2WS Server once and run the scripts manually so that the use of the private key will be approved.

1.4 N2WS Technology

As part of the cloud ecosystem, N2WS relies on web technology. The management interface through which you manage backup and recovery operations is web-based. The APIs which N2WS uses to communicate with AWS, are web-based. All communication with the N2WS server is performed using the HTTPS protocol, which means it is all encrypted. This is important, since sensitive data will be communicated to/from the N2WS server, for example, AWS credentials, N2WS credentials, object IDs of your AWS objects (instances, volumes, databases, images, snapshot IDs etc.).

1.5 Browser Support

Most interactions with the N2WS server are performed via a web browser.

- Since N2WS uses modern web technologies, you will need your browser to be enabled for Java Script.
- N2WS supports Microsoft Edge, Mozilla Firefox, and Google Chrome.
- Other browsers are not supported.

1.6 Viewing Tutorial and Free Installation

If you want to view a getting-started tutorial, or to try the fully-functional N2WS free for 30 days, go to https://n2ws.com/support/video-tutorials/getting-started.
Follow the instructions in the 'Getting Started with N2WS Backup & Recovery for AWS' video.

Note:	It is not necessary to reinstall N2WS after purchasing a license.

1.7 Customized Free Trial

It is now possible to have a free trial of N2WS with the usage limitations customized for your specific AWS infrastructure. Contact N2W Software sales at info@n2ws.com to start your customized free trial. The N2W Software sales team may provide a reference code for your customized installation.

2 Configuring N2WS

Important: **BEFORE** upgrading to version 3.0 from versions 2.4-2.7, Copy to S3 customers must review section 2.3.2 (Step 4) about special conditions for data recovery.

The N2WS management console is accessed via a web browser over HTTPS.

- When a new N2WS Server is launched, the server will automatically generate a new self-signed SSL certificate. This certificate will be used for the web application in the configuration step.
- If no other SSL certificate is uploaded to the N2WS Server, the same certificate will be used also for the main N2WS application.
- Every N2WS Server will get its own certificate.
- Since the certificate is not signed by an external Certificate Authority, you will need to approve an exception in your browser to start using N2WS.

When configuring the N2WS server, define the following settings:

- AWS Credentials for the N2WS root user.
- Time zone for the server.
- Whether to create a new CPM data volume, or attach an existing one from a previous N2WS server.
- Whether to create an additional N2WS server from an existing data volume during Force Recovery Mode.
- Proxy settings. Configure proxy settings in case the N2WS server needs to connect to the Internet via a proxy. These settings will also apply to the main application.
  The port the web server will listen on. The default is 443. See section 1.3.3.
- Whether to upload an SSL certificate and a private key for the N2WS server to use. If you provide a certificate, you will also need to provide a key, which must not be protected by a passphrase.
- Register the AWS account with N2W Software. This is mandatory only for free trials but is recommended for all users. It will allow N2W Software to provide quicker and enhanced support. Registration information is not shared.

For the configuration process to work, as well as for normal N2WS operations, N2WS needs to have outbound connectivity to the Internet, for the HTTPS protocol. Assuming the N2WS server was launched in a VPC, it needs to have:

- A public IP, or
- An Elastic IP attached to it, or
- Connectivity via a NAT setup, Internet Gateway, or HTTP proxy.

If an access issue occurs, verify that the:

- Instance has Internet connectivity.
- DNS is configured properly.
- Security groups allow outbound connections for port 443 (HTTPS) or other (if you chose to use a different port).

Following are the configuration steps:

1. Approve the end-user license agreement.
2. Define the root user name, email, and password.
3. Define the time zone of the N2WS Server and usage of data volumes.
4. Fill in the rest of the information needed to complete the configuration process.

2.1 Instance ID

To initially be identified as the owner of this instance, you are required to type or paste the N2WS server instance ID. This is just a security precaution.

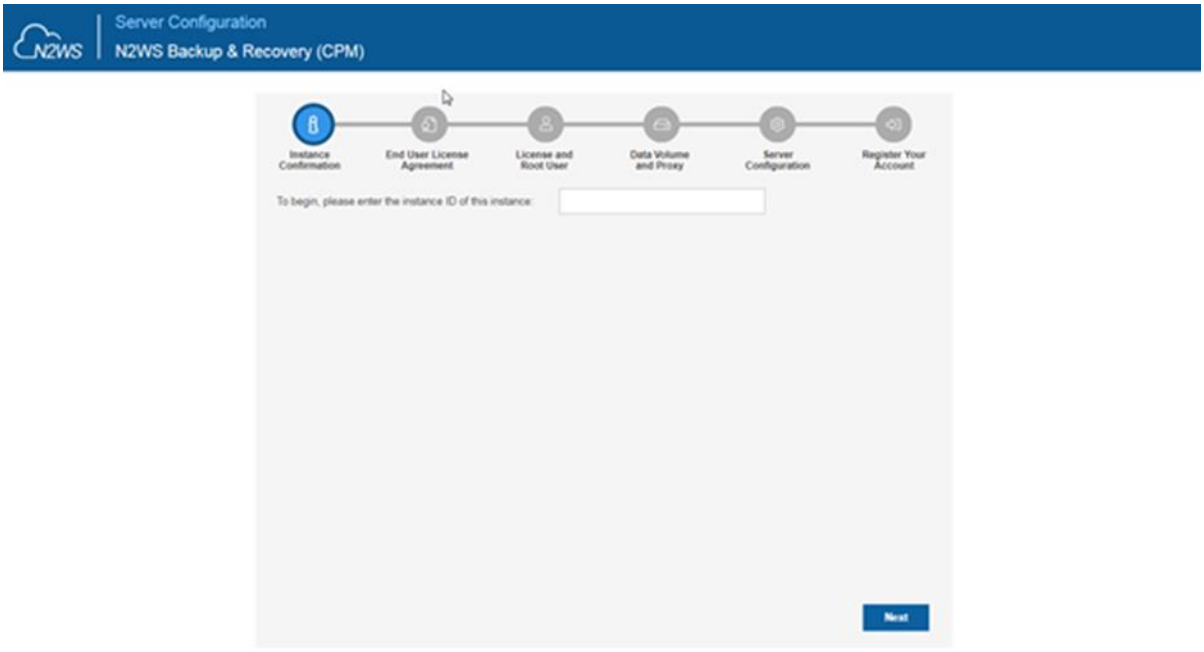

In the next step of the configuration process, you will also be required to approve the end-user license agreement.

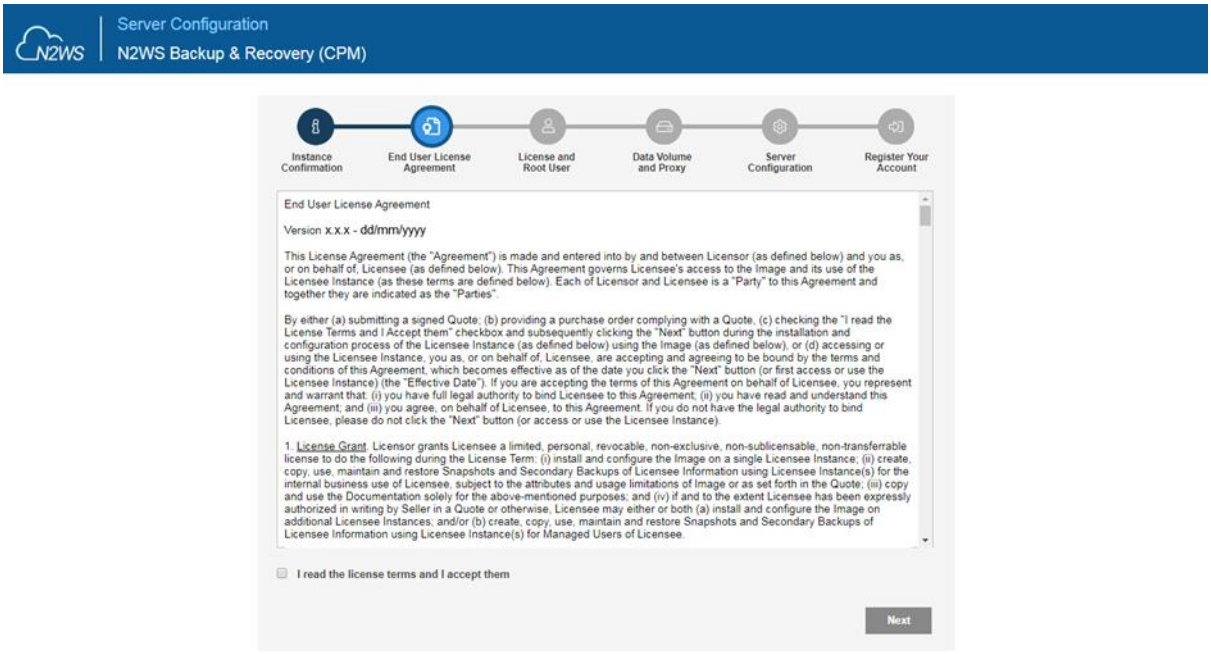

2.2 License Agreement and Root User

The **License** field is presented. Select **I'm starting a free trial** for a free trial. Otherwise, select the appropriate license option in the list, such as Bring Your Own License (BYOL) Edition. Alternatively, if your organization purchased a license directly from N2W Software, additional instructions are shown.

The AWS root user (IAM User) is no longer allowed to control the operation of the N2WS server. A user with the Authentication credentials for **N2WS Instance IAM Role** is the only user allowed to install N2WS, log on to the system server and operate it. As shown below, you need to define the root user name, email, and password. This is the third step in the configuration process. The email may be used when defining Amazon Simple Notification Service (SNS) based alerts. Once created, choose to automatically add this email to the SNS topic recipients.

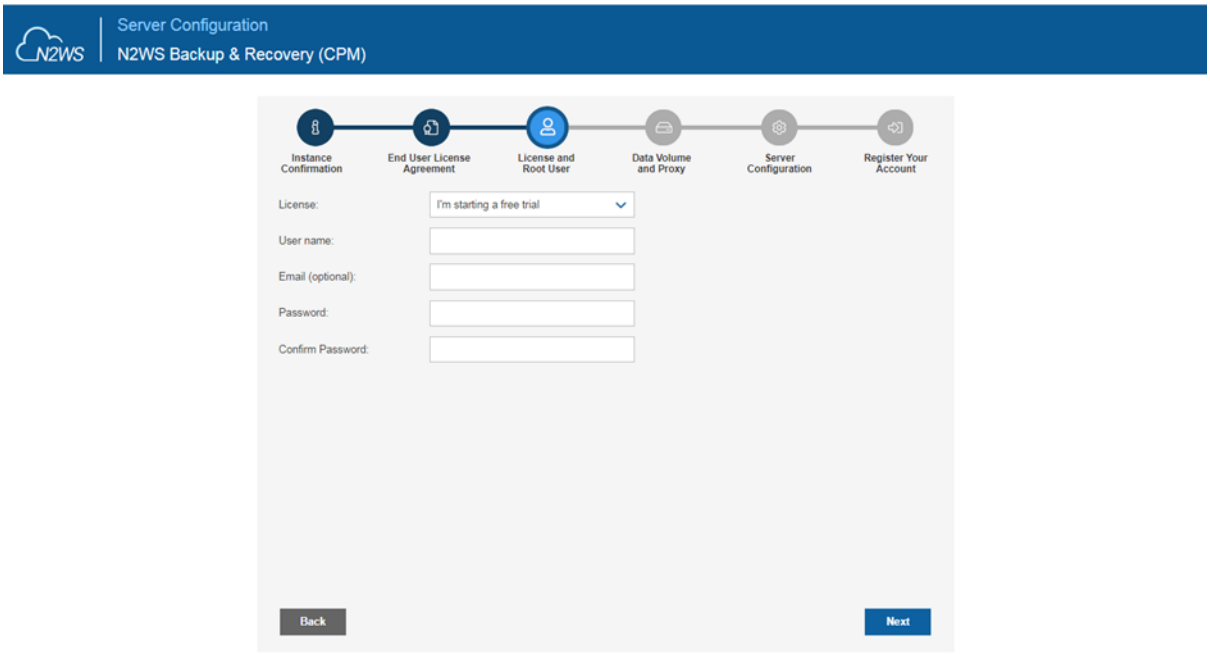

Note:	Passwords: N2W Software does not enforce any password policy, however, it is recommended that you use passwords that are difficult to guess and that are changed from time to time.

2.3 Defining Time Zone, Data Volume, Force Recovery Mode, Web Proxy

In the fourth step of the configuration process, you can:

- Set the time zone of the N2WS Server.
- If using a paid license, choose whether to create a new data volume, or use an existing one. Your AWS credentials will be used for the data volume setup process.
- Create an additional N2WS server in recovery mode only, by choosing an existing data volume and set **Force Recovery Mode.**
- Configure proxy settings for the N2WS server. See section 2.3.3.

As you will see in section 4.1.3, all scheduling of backup is performed according to the local time of the N2WS Server. You will see all time fields displayed by local time; however, all time fields are stored in the N2WS database in UTC. This means that if you wish to change the time zone later, all scheduling will still work as before.

As you can see below, the choice of new or existing data volume is made here. Actual configuration of the volume will be accomplished at the next step.

AWS credentials are required to create a new Elastic Block Storage (EBS) data volume if needed and to attach the volume to the N2WS Server instance.

- If you are using AWS Identity and Access Management (IAM) credentials that have limited permissions, these credentials need to have permissions to view EBS volumes in your account, to create new EBS volumes, and to attach volumes to instances. See section 16.3. These credentials are kept for file-level recovery later on and are used only for these purposes.

- If you assigned an IAM Role to the N2WS Server instance, and this role includes the needed permissions, select **Use Instance's IAM Role** and then you will not be required to enter credentials.

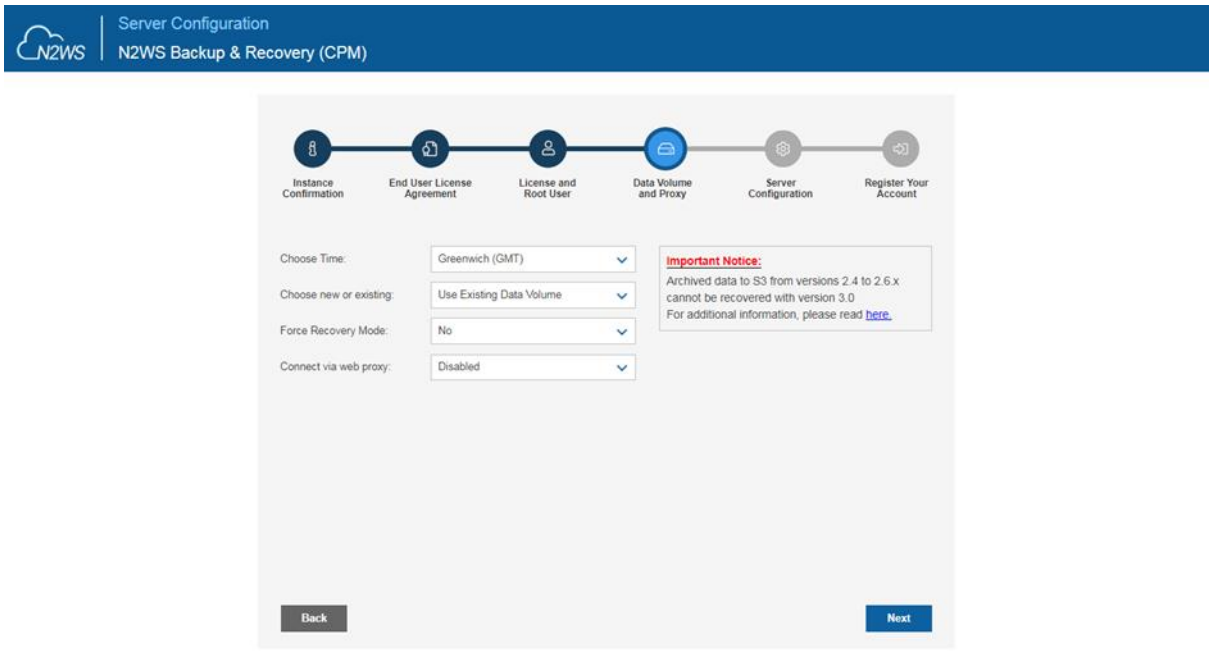

2.3.1 New Data Volume

When creating a new data volume, the only thing you need to define is the capacity of the created volume. You also have the option to encrypt the volume, as described in section 2.4.1. The volume is going to contain the database of N2WS's data, plus any backup scripts or special configuration you choose to create for the backup of your servers. The backup itself is stored by AWS, so normally the data volume will not contain a large amount of data.
The default size of the data volume is 5 GiB.

- This is large enough to manage roughly 50 instances, and about 3 times as many EBS volumes.
- If your environment is larger than 50 instances, increase the volume at about the ratio of 1 GiB per 10 backed-up instances.

The new volume will be automatically created in the same AZ as the N2WS instance It will be named **N2WS Data Volume**. During the configuration process, the volume will be created and attached to the instance. The N2WS database will be created on it.

2.3.2 Existing Data Volume

Important notice for Copy to S3 customers BEFORE upgrading to version 3.0:

- All data previously archived to S3, using versions 2.4-2.6.x, **cannot** be recovered using version 3.0.
- To allow recovery of such data in the future, create an AMI of your current N2WS instance **BEFORE upgrading** to version 3.0.
- To do this, follow all the steps outlined in the version 3.0 Upgrade Instructions **BEFORE continuing** your upgrade.

- For additional information, see Release Notes.

The Existing data volume option is used if:

- You have already run N2WS and terminated the old N2WS server, but now wish to continue where you stopped.
- You are upgrading to new N2WS releases.
- You are changing some of the configuration details.
- You want to configure an additional N2WS server in recovery mode only. See section 2.3.4.

The select box for choosing the volumes will show all available EBS volumes in the same AZ as the N2WS Server instance. When choosing the volumes, consider the following:

- It is important to create the instance in the AZ your volume was created in the first place.
- Another option is to create a snapshot from the original volume, and then create a volume from it in the AZ you require.

Note: Although CPM data volumes typically have a special name, it is not a requirement. If you choose a volume that was not created by a N2WS server for an existing data volume, the application will *not* work.

2.3.3 Proxy Settings

If the N2WS server needs an HTTP proxy to connect to the Internet, define the proxy address, port, user, and password. The proxy settings will be kept as the default for the main application.

In the N2WS UI, proxy settings are made in the **Proxy** tab of **Server Settings** > **General Settings**.

Note: Make sure to enable SSH connections (port 22) through your proxy.

2.3.4 Force Recovery Mode

You can configure an additional N2WS server, in recovery mode only, by choosing an existing data volume:

- In step 4, choose to use an existing volume and in the **Force Recovery Mode**, select **Yes**.
- In step 5, in the **Existing CPM Data Volume** list, select the volume that holds your backup records.

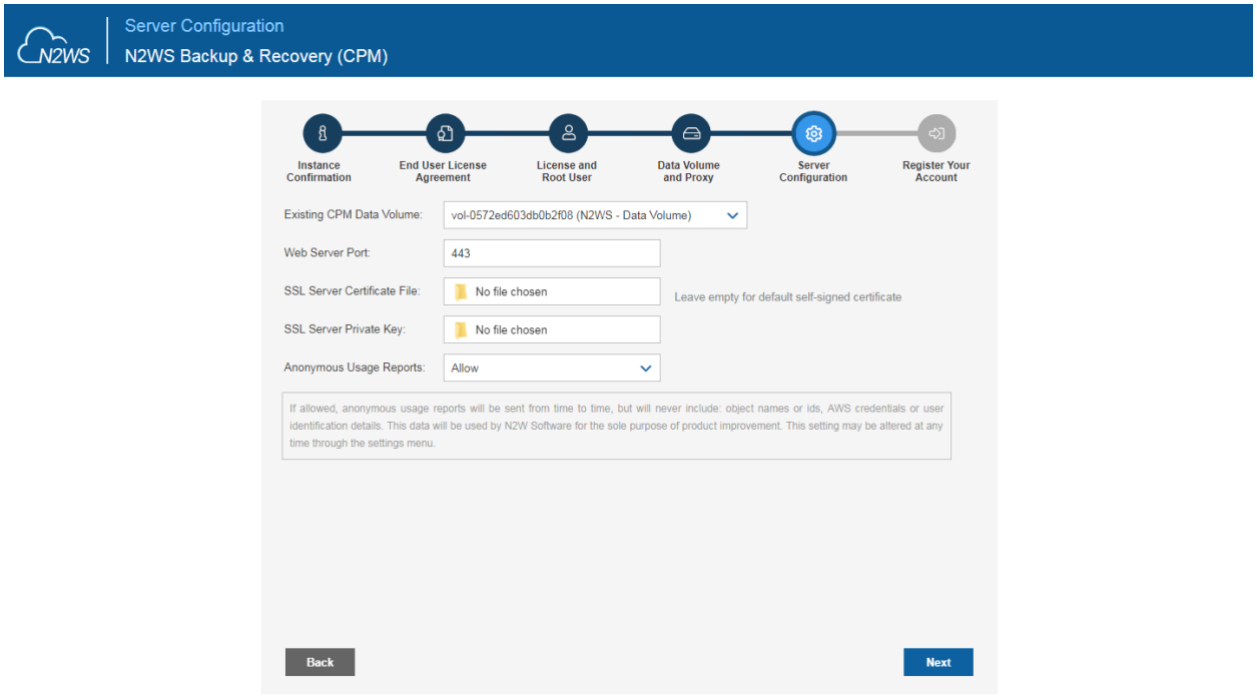

Note: The N2WS server configured for recovery mode will NOT:

- Perform backups.
- Perform Data Lifecycle Management operations.
- Perform Recovery Scenario.
- Have Resource Control management.
- Perform any scheduled operations.

2.4 Complete Remaining Fields in N2WS Configuration

In the fifth step, you will fill in the rest of the information needed for the configuration of the data volume for the N2WS Server.

If you chose to create a new volume, you can choose the volume capacity, type, and whether to encrypt.

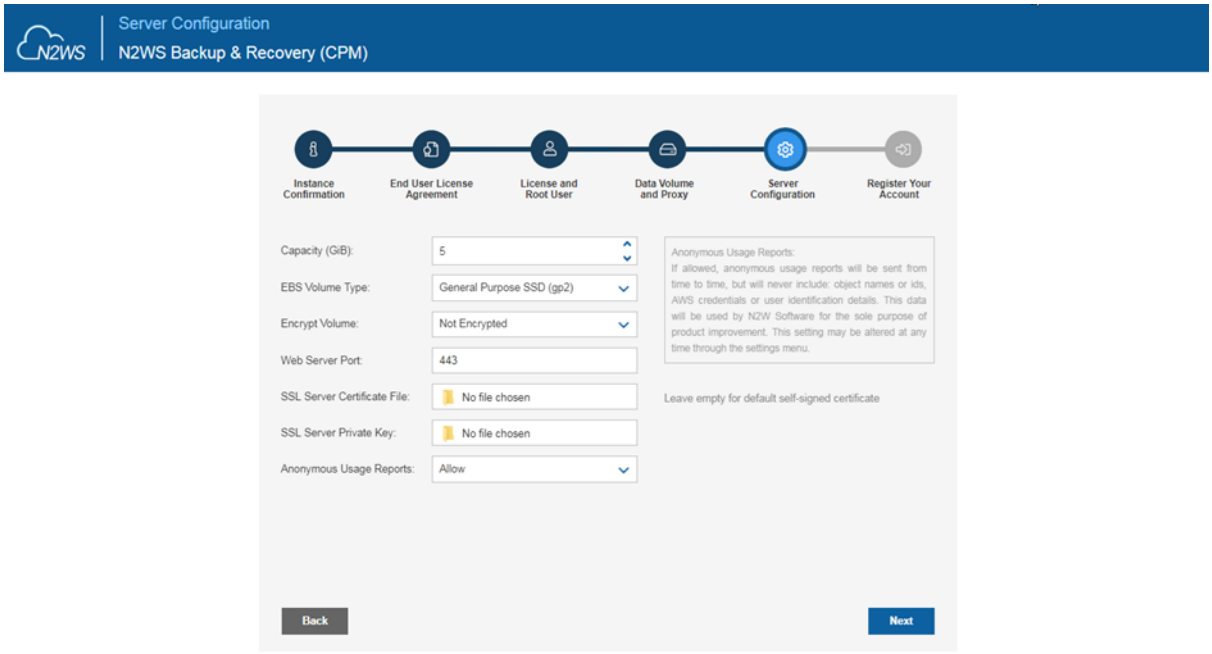

If you chose to use an existing volume, you will see a drop-down volume selection box instead of the volume capacity field:

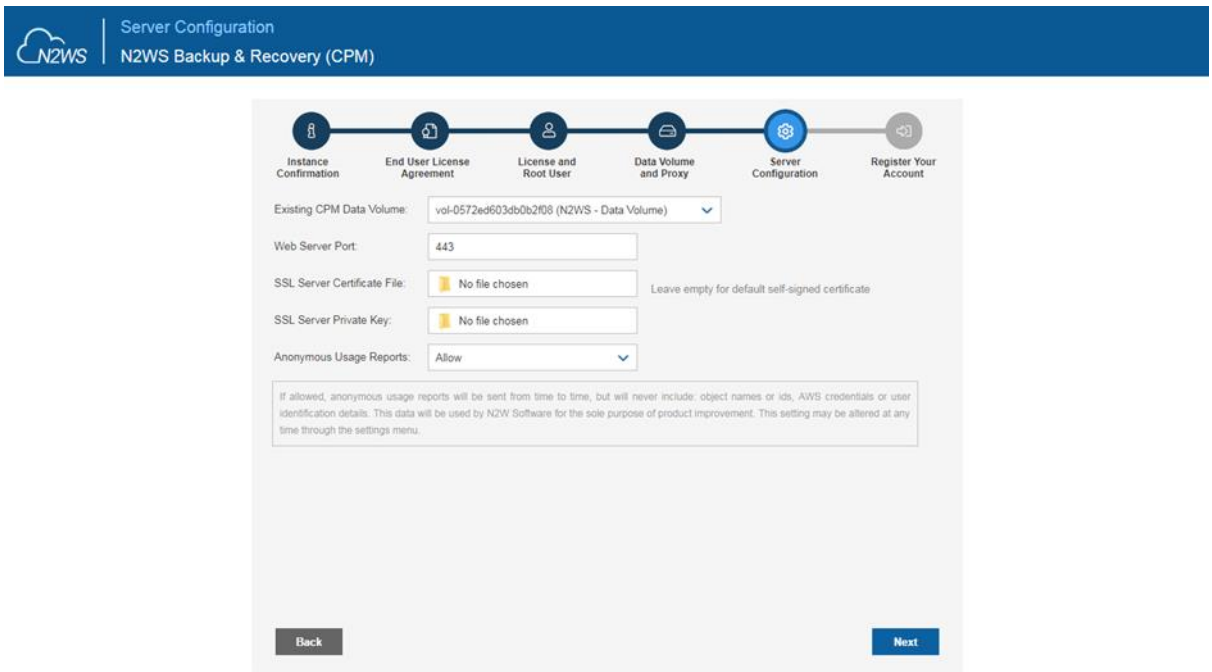

2.4.1 Encrypting a New Data Volume

If you choose a new data volume, you have an option to encrypt CPM user data. You also have the option to encrypt a new data volume if using the silent configuration mode. See section 2.8.

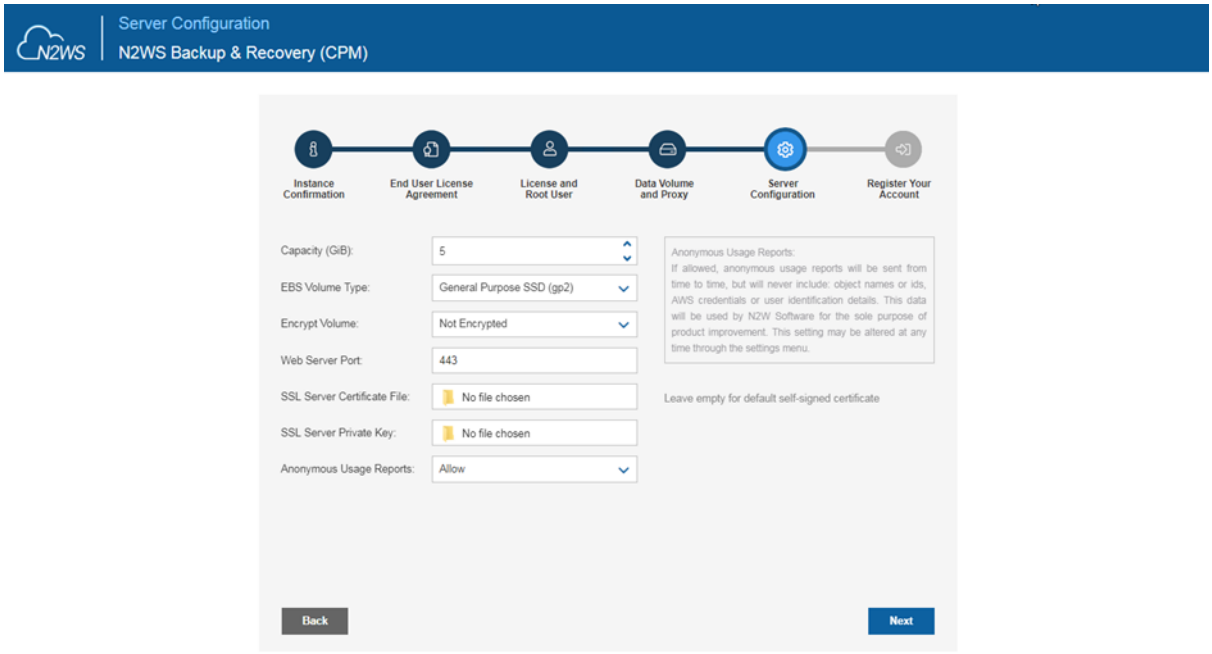

Select **Encrypted** in the **Encrypt Volume** drop-down list and choose a key in the **Encryption Key** list. You have the option to use a custom ARN.

2.4.2 Web Server Settings

Port 443 is the default port for the HTTPS protocol, which is used by the N2WS manager. If you wish, you can configure a different port for the web server. But, keep in mind that the specified port will need to be open in the instance's security groups for the management console to work, and for any Thin Backup Agents that will need to access it.
The final detail you can configure is an SSL certificate and private key.

- If you leave them empty, the main application will continue to use the self-signed certificate that was used so far.
- If you choose to upload a new certificate, you need to upload a private key as well. The key cannot be protected by a passphrase, or the application will not work.

Warning: If a corrupted SSL certificate is installed, it will prevent the N2WS server from starting.

2.4.3 Anonymous Reports Setting

Leaving the Anonymous Usage Reports value as **Allow** permits N2WS to send anonymous usage data to N2W Software. This data does not contain any identifying information:

- No AWS account numbers or credentials.
- No AWS objects or IDs like instances or volumes.
- No N2WS names of objects names, such as, policy and schedule.

It contains only details like:

- How many policies run on a N2WS server
- How many instances per policy
- How many volumes

- What the scheduling is, etc.

You can change this setting at any time using the enable/disable anonymous usage reports link at the bottom of N2WS's main page.

2.5 Registering and Finalizing the Configuration

After filling in the details in the last step, you are prompted to register. This is mandatory for free trials and optional for paid products.

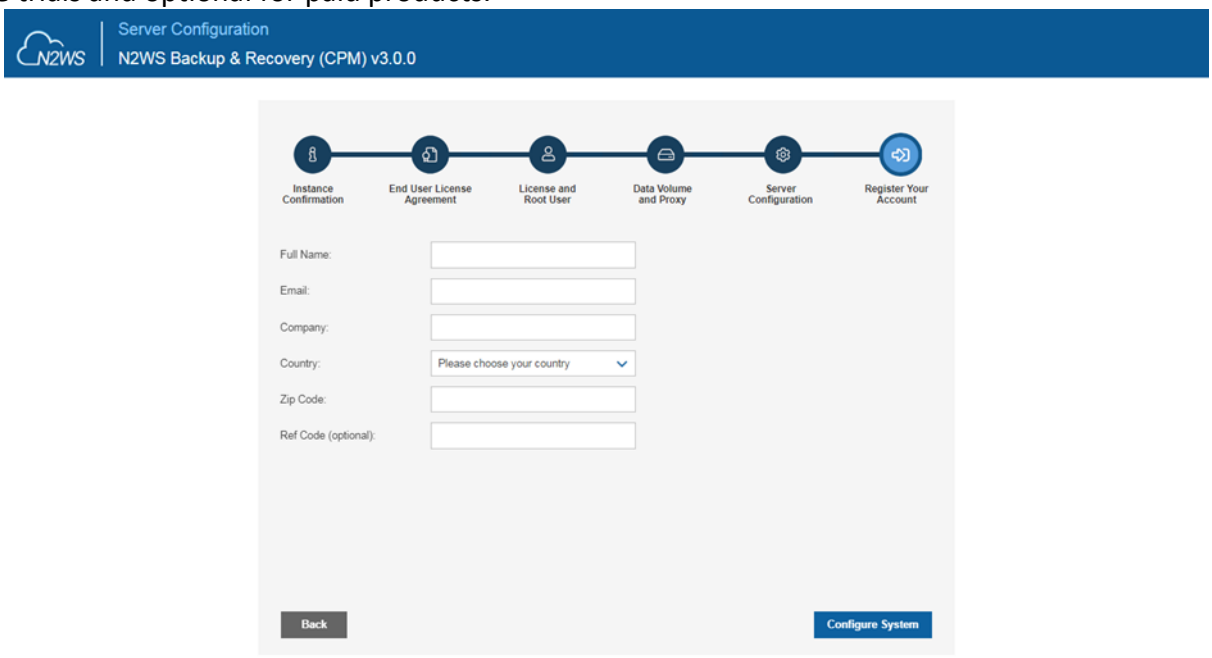

Select **Configure System** to finalize the configuration. The configuration will take between 30 seconds and 3 minutes for new volumes, and usually less for attaching existing volumes. After the configuration is complete, a 'Configuration Successful – Starting Server …' message appears. It will take a few seconds until you are redirected to the login screen of the N2WS application.

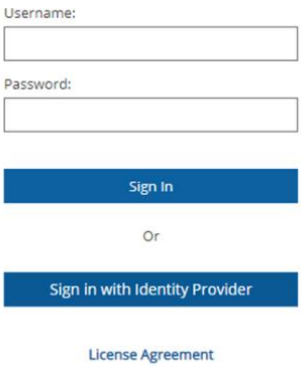

If you are not redirected, refresh the browser manually. If you are still not redirected, reboot the N2WS server via AWS Management Console, and it will come back up, configured and running.

2.6 Configuration Troubleshooting

Most inputs you have in the configuration steps are validated when you select **Next**. You will get an informative message indicating what went wrong.
A less obvious problem you may encounter is if you reach the third step and get the existing volume select box with only one value in it: **No Volumes found**. This can arise for two reasons:

- If you chose to use an existing volume and there are no available EBS volumes in the N2WS Server's AZ, you will get this response. In this case, you probably did not have your existing data volume in the same AZ.

  To correct this:

  - Terminate and relaunch the N2WS server instance in the correct zone and start over the configuration process, or
- Take a snapshot of the data volume, and create a volume from it in the zone the server is in.
- If there is a problem with the credentials you typed in, the "No Instances found" message may appear, even if you chose to create a new data volume. This usually happens if you are using invalid credentials, or if you mistyped them.

  To fix, go back and enter the credentials correctly.

In rare cases, you may encounter a more difficult error after you configured the server. In this case, you will usually get a clear message regarding the nature of the problem. This type of problem can occur for several reasons:

- If there is a connectivity problem between the instance and the Internet (low probability).
- If the AWS credentials you entered are correct, but lack the permissions to do what is needed, particularly if they were created using IAM.
- If you chose an incorrect port, e.g. the SSH port which is already in use.
- If you specified an invalid SSL certificate and/or private key file.

In case you cannot discover the problem, try again. If it persists, contact N2W Software support (support@n2ws.com).
If the error occurred after completing the last configuration stage, it is recommended that you:

1. Terminate the N2WS server instance.
2. Delete the new data volume (if one was already created).
3. Try again with a fresh instance.

2.7 Modifying the Configuration of a N2WS Server

If you need to change the configuration of your N2WS server after it has already been created, you may need to:

- Change the time zone
- Reset the N2WS root user password
- Change SSL credentials
- Change the HTTPS port

The process to make these changes is to terminate the current N2WS server instance and create a new one. After you terminate the N2WS server, the data volume becomes available. Configure the server as needed and connect to the old (existing) data volume.

Note: Remember to launch the new server in the same AZ.

For the N2WS root user, you may change the email or the password. The username of the root user cannot be changed. If, during the configuration process, you type a different username than the original, N2WS will assume you forgot the root username. In that case, the username will not change, and a file named `/tmp/username_reminder` will be created on the N2WS server. It will contain the username. You can connect to N2WS server using SSH to view this file. See section 7.1.

2.8 Configuring N2WS in Silent Mode

There is an option to configure N2WS using a special "user data" script. The **user data** script is a configuration in `ini` file format, stating the configuration of the new N2WS instance.
Create the **user data** file with `CPMCONFIG` in the first line, `[SERVER]` in the second line, followed by the configuration details.
N2WS assumes that the N2WS instance has an IAM role that is used for the configuration process, so no credentials are required.
Following is an example of the whole script:

```
CPMCONFIG

[SERVER]

user=<username for the N2WS user>

password=<password>

volume_option=<new or existing>

volume_size=<in GB, used only for the new volume option>

volume_id=<Volume ID for the data volume, used only in the existing
volume option>

volume_type=<set your storage performance and cost.
The default is "gp2". It can be set to "io1" or "gp2">

snapshot_id=<snapshot ID to create the data volume from, used only with
the existing volume option, and only if volume_id is not present>

encryption_key=<encrypt user-data volume by setting the ARN of the
KMS key. used only for the new volume option>

time_zone=<set N2WS server's local time.
The default timezone is GMT. Available time zones are listed in
Appendix C – Timezone Choices>

allow_anonymous_reports=<send anonymous usage data to N2W Software.
The default is "False">

force_recovery_mode=<allow additional N2WS server to perform recovery
operations only. The default is "False". If it set to "True" - it
requires volume_option=existing>
```

Additionally, if you need the N2WS server to connect to the Internet via an HTTP proxy, add a `proxy` section:

```
[PROXY]
proxy_server=<address of the proxy server>
proxy_port=<proxy port>
proxy_user=<user to authenticate, if needed>
proxy_password=<password to authenticate, if needed>
```

The snapshot option does not exist in the UI. It can be used for automation of a Disaster Recovery (DR) server recovery. Additionally, if you state a volume ID from another AZ, N2WS will attempt to create a snapshot of that volume and migrate it to the AZ of the new N2WS server. This option can be used in a high availability setup.

Note:	You are not required to select to approve the license terms when using the silent configuration option, since you already approved the terms when subscribing to the product on AWS Marketplace.

3 Start Using N2WS

N2WS opens to the Dashboard – an overview of recent backups, recoveries, alerts, resources, and costs.

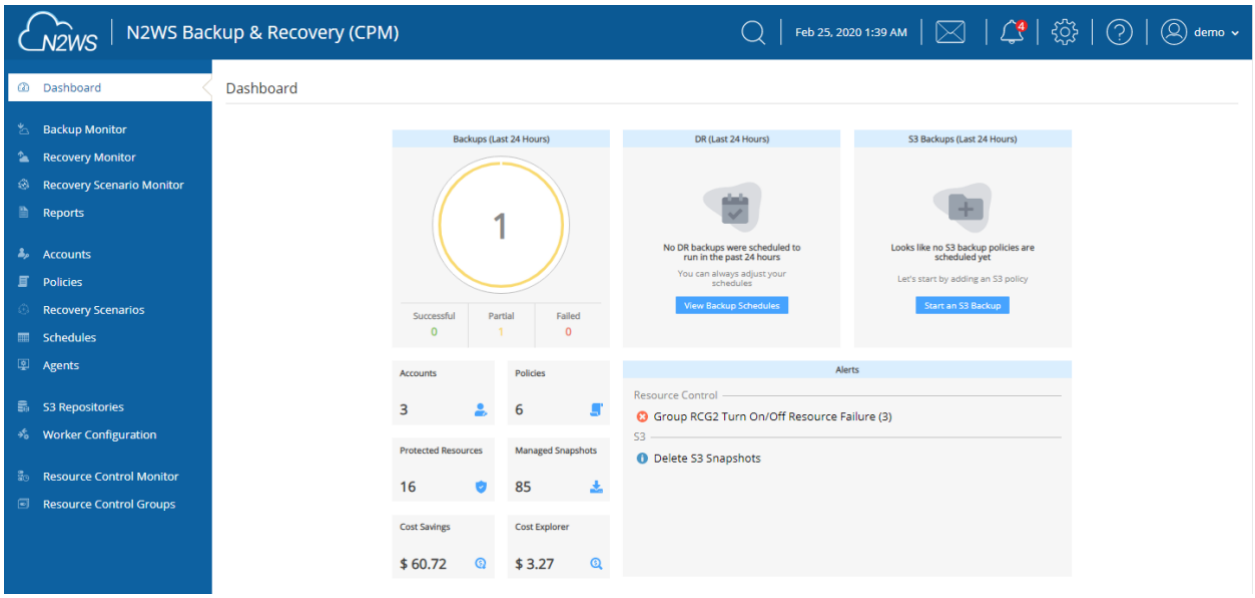

Depending on your device resolution, the **Alerts** tile may not appear in the Dashboard.

Regardless of resolution, all Alerts are available by selecting  in the toolbar.

3.1 Associating an AWS Account

To associate an AWS account for File Level Recovery, you will need to either:

- Enter AWS credentials consisting of an access key and a secret key, or
- Use an IAM role, either on the N2WS server instance or cross-account roles.

There are two steps to associating a N2WS account with an AWS account:

1. To manage your users and roles and obtain AWS credentials, go to the IAM console at https://console.aws.amazon.com/iam/home?#users
   a. Follow the directions to either add a new account or view an existing account.
   b. Capture the AWS credentials.
2. To associate the AWS account with a N2WS account, go to N2WS:
   a. In the left panel, select the **Accounts** tab and then select + **New**.
   b. Complete the fields, entering the required information for the **Account Type** and **Authentication** method.

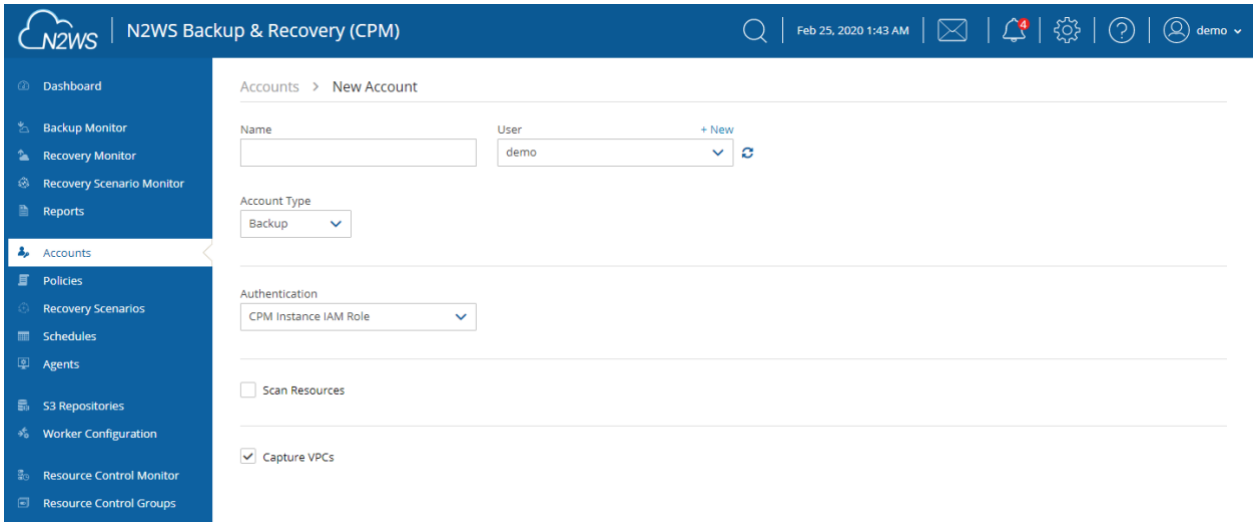

3.1.1 Account Type

If you are using the Advanced or Enterprise Edition or a free trial, you will need to choose an account type.

- The **Backup** account is used to perform backups and recoveries and is the default 3333
- **DR Account** is used to copy snapshots to as part of cross-account functionality.
  If this is a DR Account, you choose whether this account is allowed to delete snapshots. If the account not allowed to delete snapshots when cleaning up, the outdated backups will be tagged. Not allowing N2WS to delete snapshots of this account implies that the presented IAM credentials do not have the permission to delete snapshots.

3.1.2 Authentication

N2WS Supports three methods of authentication:

- **IAM User** - Authentication using IAM credentials, access and secret keys.

  Warning: Using IAM User credentials is not recommended as they are less secure than using IAM roles.

- **CPM Instance IAM Role** – If an IAM role was assigned to the N2WS server at launch time or later, you can use that IAM role to manage backups in the same AWS account the N2WS server is in.

  Note: Only the root/admin N2WS user is allowed to use the IAM role.

- **Assume Role** – This type of authentication requires another AWS account already configured in N2WS. If you want to use one account to access another, you can define a cross-account role in the target account and allow access from the first one. The operation of using one account to take a role and accessing another account is called **assume role**.

To allow account authentication using Assume Role in N2WS:

1. In the **Authentication** box, choose **Assume Role**.
2. In the **Account Number** box, type the 12-digit account number, with no hyphens, of the target account.

3. In the **Role to Assume** box, type the role name, not the full Amazon Resource Name (ARN) of the role. N2WS cannot automatically determine what the role name is, since it is defined at the target account, which N2WS has no access to yet.
4. The **External ID** box is optional unless the cross-account role was created with the **3rd party** option.
5. In the **Assuming Account** list, choose the account that will assume the role of the target account.

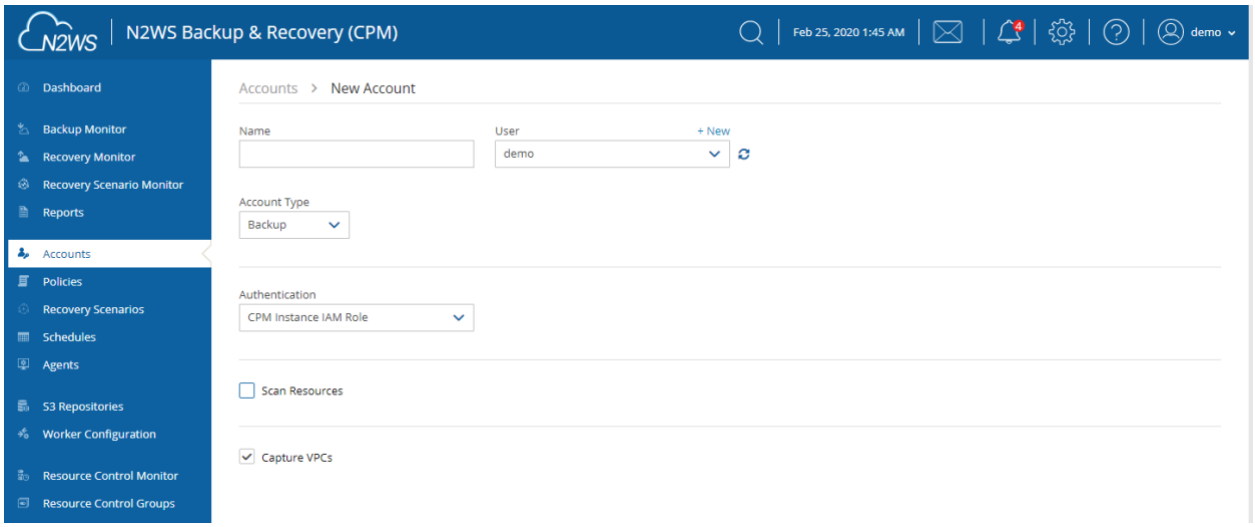

If you are the root user or independent user and have managed users defined, an additional selection list will appear enabling you to select the user.

6. Select **Scan Resources** to include the current account in tag scans performed by the system. Once **Scan Resources** is **Enabled**:
   - In the **Scan Regions** list, select the regions to scan. To select all regions, select the check box at the top of the list. To filter regions, start typing in the search box.

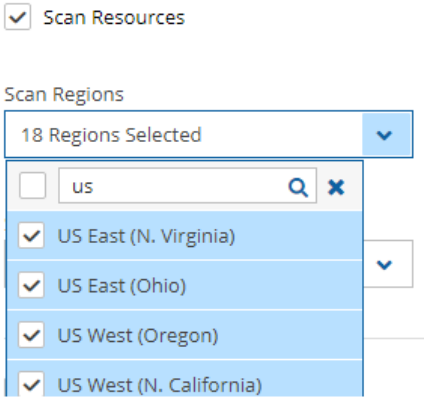

   - The **Scan Resource Types** list, select the types of resource to scan. Select the top check box for all, or use the search box to filter types.

Note: Scanning only specific resource types can reduce tag scan processing time and is useful when user permissions are limited to certain resource types.

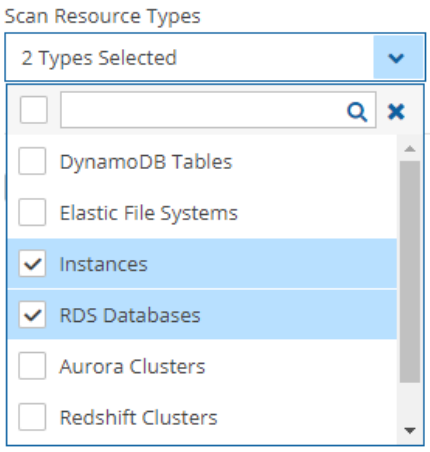

7. The **Capture VPCs** option defaults to enabled. Clear **Capture VPCs** to disable for this account. See section 23.

Note: You can add as many AWS accounts as your N2WS edition permits.

3.2 N2WS Support

For support issues, contact N2W Software support. To collect and download support logs, select in the toolbar and then select **Download Logs**.
In the Download Support Logs dialog box, select the relevant logs and time frame, and then select **Download Logs**.

- **Check AWS permissions** – Against the required permissions for AWS services and resources.
- **Collect S3 Worker Logs** – When S3 support is needed.
- **Collect System Logs** – For comprehensive system debugging.
- **Collect Backup Logs from Last** - Select Day, Week, or Month in the list.

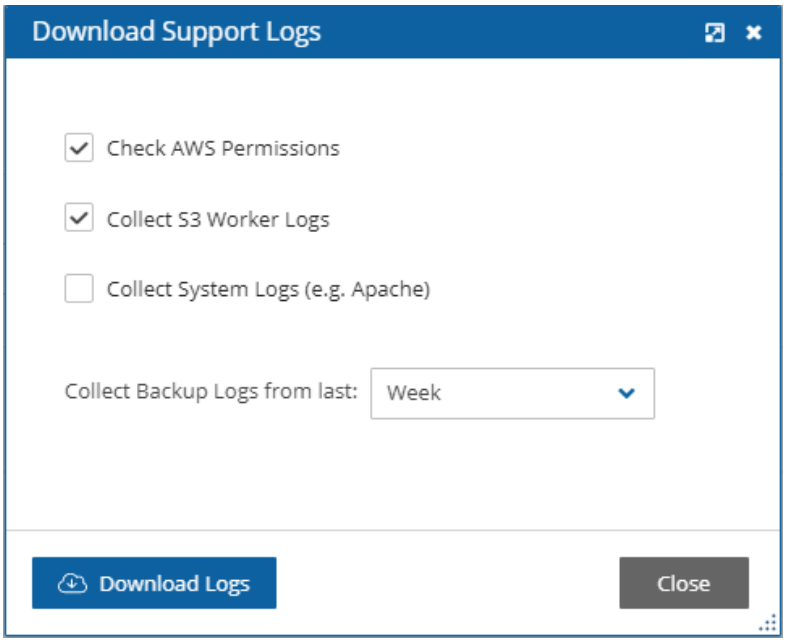

4 Defining Backup Policies

The backbone of the N2WS solution is the backup policy. A backup policy defines everything about a logical group of backed-up objects. A policy defines:

- What will be backed up - **Backup Targets.**
- How many generations of backup data to keep.
- When to back up – **Schedules.**
- Whether to use backup scripts.
- Whether VSS is activated (Windows Servers 2008, 2012, 2016, and 2019 only).
- Whether backup is performed via a backup agent (Windows only).
- The retry policy in case of failure.
- DR settings for the policy.
- Lifecycle events: Copy to S3, Archive to Glacier

The following sections explain the stages for defining a policy and its schedule. Schedule definition is addressed first as it one of the attributes of a policy and Scheduled Reports.

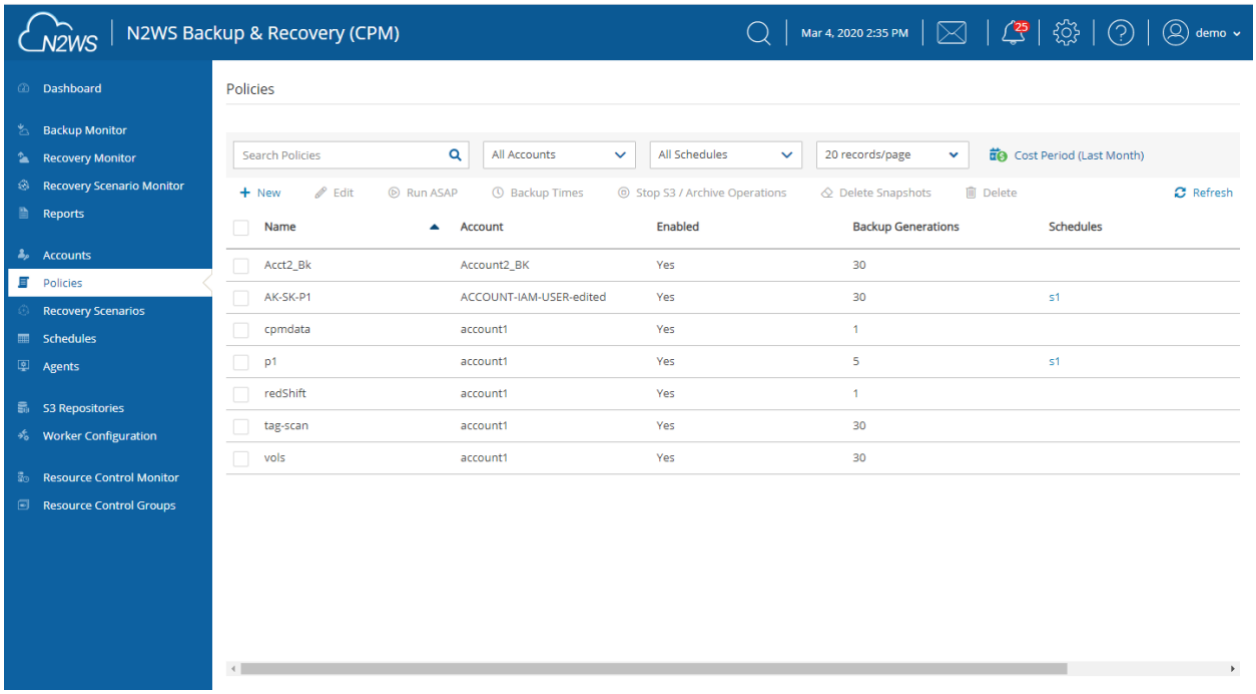

4.1 Schedules

Schedules are the objects defining **when** to perform a backup

- Schedules are defined separately from policies and Scheduled Reports.
- One or more schedules can be assigned to both policies and Scheduled Reports.

Schedules can be managed in the **Schedules** tab found in the left panel. They can be added also during the creation of a new Policy.

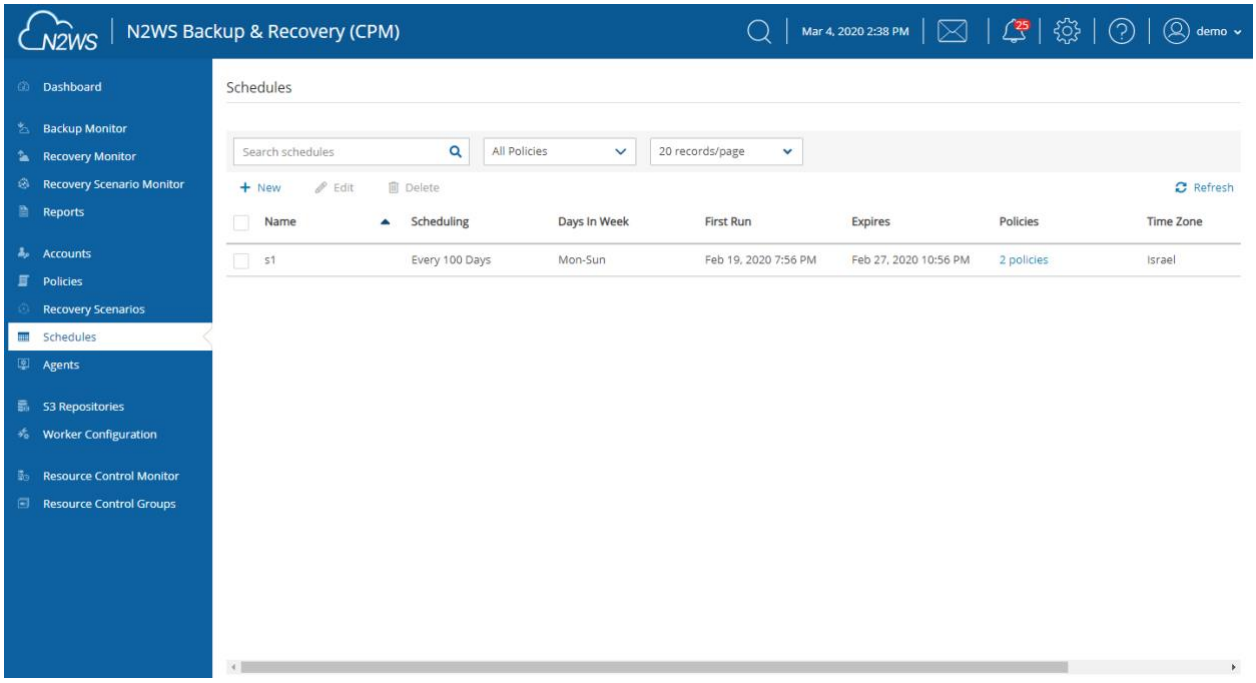

Or, added when creating a new scheduled report in the **Scheduled Reports** tab of the **Reports** tab in the left panel.

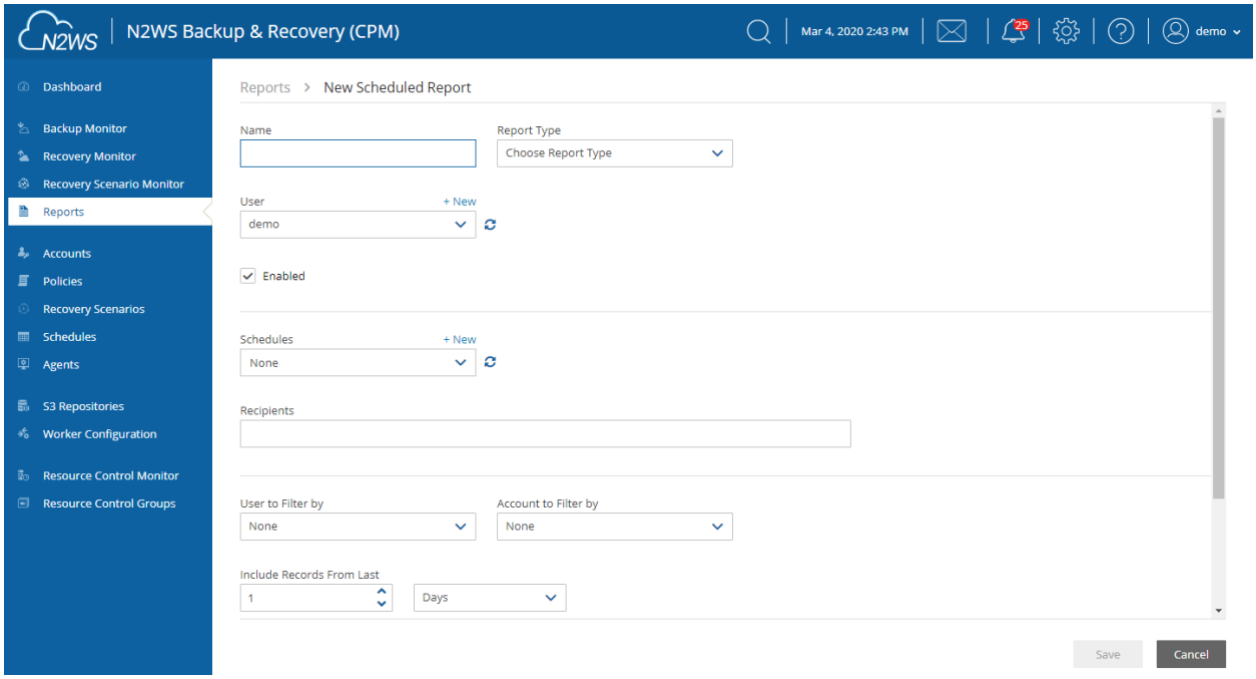

Note: Both interfaces include all defined schedules and the same definition options.

You can define schedules to:

- Run for the first time at a date and time in the future.
- Run forever or have a specific date and time to stop.
- Repeat every ‘n’ minutes, hours, days, weeks, months.

- Selectively enable for certain minutes, hours, and day of the week, but not for weeks and months.
- Repeat every day of the week, or only run on certain days.
- Exclude running a report or policy during certain time ranges within the scheduled times.

For the root/admin user, if you have created additional managed users, you can select the user to whom the schedule belongs.

Important: For weekly or monthly backups and report generation, the start time will determine the day of week of the schedule and *not* the days of week check boxes.

4.1.1 Defining Schedules

The same schedule can be used in all of the following:

- Policy backup operations.
- Recovery Scenarios for policies with schedules.
- Generating **Scheduled Reports**.

Note: All start times are derived from the **First Run** time.

A schedule can be added either from the **Schedules** tab or as part of adding or editing a **Scheduled Report** in the **Reports** tab:

1. In the **Schedules** tab, select + **New**:
2. In the **Reports** tab:
   - To create a new Scheduled Report, select + **New**. In the **New Scheduled Report** page, select + **New** above the **Schedules** list.
   - To add a new schedule for an existing report, select the report, select **Edit** and then select + **New** above the **Schedules** list.

To define a schedule:

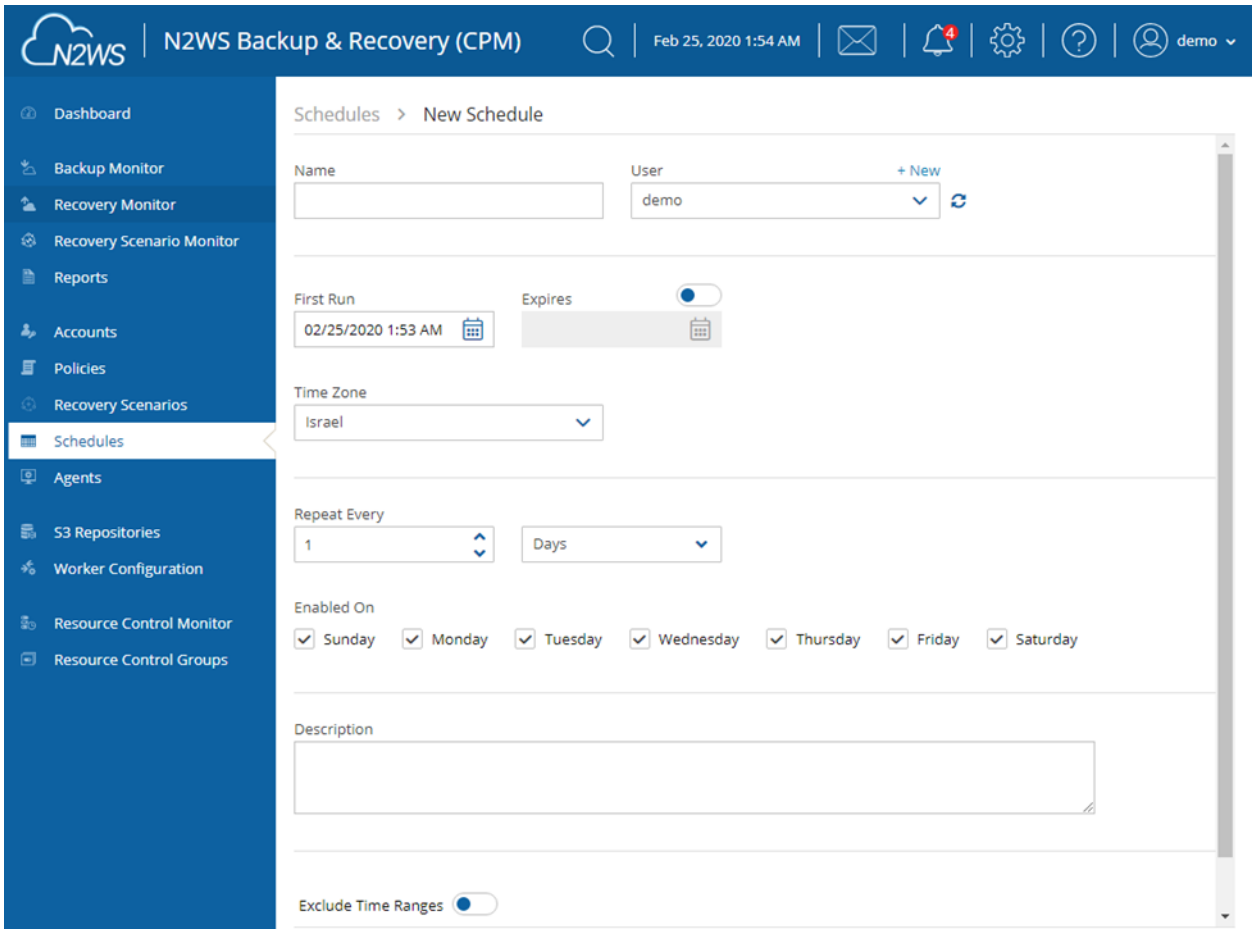

1. Type a name for the schedule and an optional description.
2. In the **First Run** box, if the First Run is other than immediately, select **Calendar** to choose the date and time to first run this schedule.

   Note: The time set in **First Run** becomes the regular start time for the defined schedule.

   - If you want a daily backup to run at 10:00 AM, set **Repeats Every** to one day and the start time to 10:00 AM.
   - If you want an hourly backup to run at 17 minutes after the hour, set **Repeats Every** to one hour and **First Run** to `nn`:17, where `nn` is the hour of the **First Run**.
3. If this schedule expires, turn on the **Expires** toggle and select the date and time the schedule ends. By default, schedules never expire.
4. In the **Repeats Every** list, select the frequency of the backups for this schedule. The possible units are months, weeks, days, hours, and minutes.
5. In the **Enabled On** section, select the day-of-week check boxes on which to run the schedule.
6. To exclude time ranges within the defined schedule, turn on the **Exclude Time Ranges** toggle. See section 4.1.4.
7. Select **Save**.

4.1.2 Overriding Schedules

It is possible to override existing schedules and run backups and Scheduled Reports immediately:

- Ad hoc backups are initiated by the **Run ASAP** command in the **Policies** tab. See section 4.2.8.
- Ad hoc generation of Scheduled Reports are initiated by the **Run Now** command in the **Reports** page. See section 17.9.3.

4.1.3 Scheduling and Time Zones

When you configure a N2WS server, its time zone is set. See section 2.3. In the N2WS management application, all time values are in the time zone of the N2WS server. Schedule times, however, may be set and executed according to a specific time zone. A policy's **Backup Times** will be according to the time zone.

Important: Even when you are backing up instances that are in different time zones, the backup time that is shown on the monitor and Backup Log is always according to the N2WS server's local time.

In the N2WS database, times are saved in UTC time zone (Greenwich). So, if, at a later stage, you start a new N2WS server instance, configure it to a different time zone, and use the same CPM data volume as before, it will still perform backup at the same times as before.

4.1.4 Disabled Times

While or after defining a schedule, you can set specific times when the schedule should not start a backup or generate a Scheduled Report. For example, you want a backup or report to run every hour, but not on Tuesdays between 01:00 PM and 3:00 PM. You can define that on Tuesdays, between these hours, the schedule is disabled.
You can define a disabled time where the finish time is earlier than the start time. The meaning of disabling the schedule on **Monday** between 17:00 and 8:00 is that it will be disabled every Monday at 17:00 until the next day at 8:00. The meaning of disabling the schedule for **All** days between 18:00 and 6:00 will be that every day the schedule will be disabled after 18:00 until 6:00.
Be careful not to create contradictions within a schedule's definition:

- It is possible to define a schedule that will never start backups or generate a report.
- You can define a weekly schedule which runs on Mondays, and then deselect Monday from the week days.
- It is also possible to create different "disabled times", which would effectively mean that the schedule is always disabled.

4.1.4.1 Defining Disabled Times

For each schedule, you can define as many excluded time ranges as you need.

To define disabled times:

1. In **New Schedule** screen, turn on the **Exclude Time Ranges** toggle.

2. In the **Day** list, select a day of the week to exclude from the schedule.
3. In the **Start Time** and **End Time** lists, select the start and end time of the exclusion.

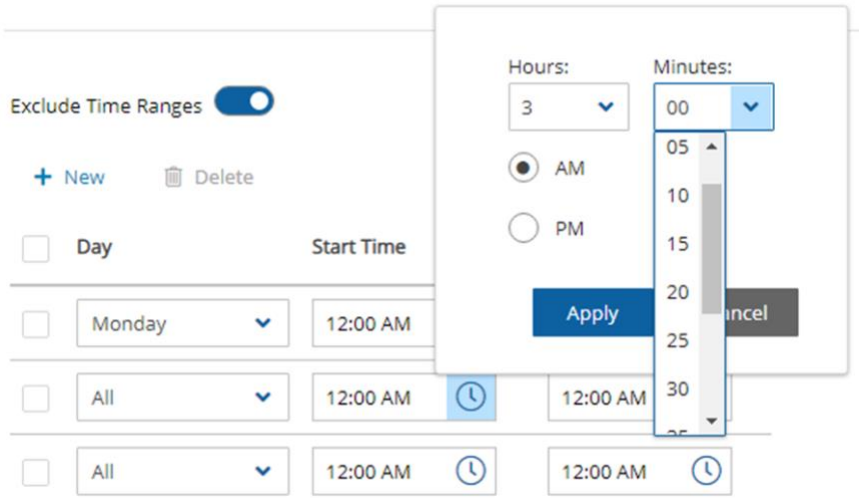

4. Select **Apply** after each definition. To add another time range, select **+ New**.
5. Select the check boxes of the excluded time ranges to enable, and then select **Save**.

4.2 Policies

Policies are the main objects defining backups. A policy defines:

- What to back up
- How to back it up
- When to perform the backup by assigning schedules to the policy

In the new UI for v3.0.0, policy definitions are spread among the following tabs:

- **Policy Details** – Basic policy details: name, user, account, enabled, schedules, auto target removal, backup generations, and description
- **Backup Targets** – Choose and configure backup targets
- **More Options** – Backup options, such as activate Linux backup scripts, define successful backup, retry parameters, and number of failures to trigger alert
- **DR** – Enable DR
- **Lifecycle Management (Snapshot /S3 / Glacier)** – Configure Lifecycle operations, including number of backup generations, copy to S3 and archive to Glacier.

For the `cpmdata` policy, the relevant tabs are **Policy Details** and **DR**.

4.2.1 Creating a New Policy

Note:	As of v2.7.0, the `cpmdata` policy is no longer using scripts as the default. Users can enable application consistent scripts by selecting **Application Consistent** for the `cpmdata` policy in the **Policy Details** tab.

To define a new policy:

1. In the left panel, select **Policies** and then select **+ New**. The **Policy Details** tab opens.

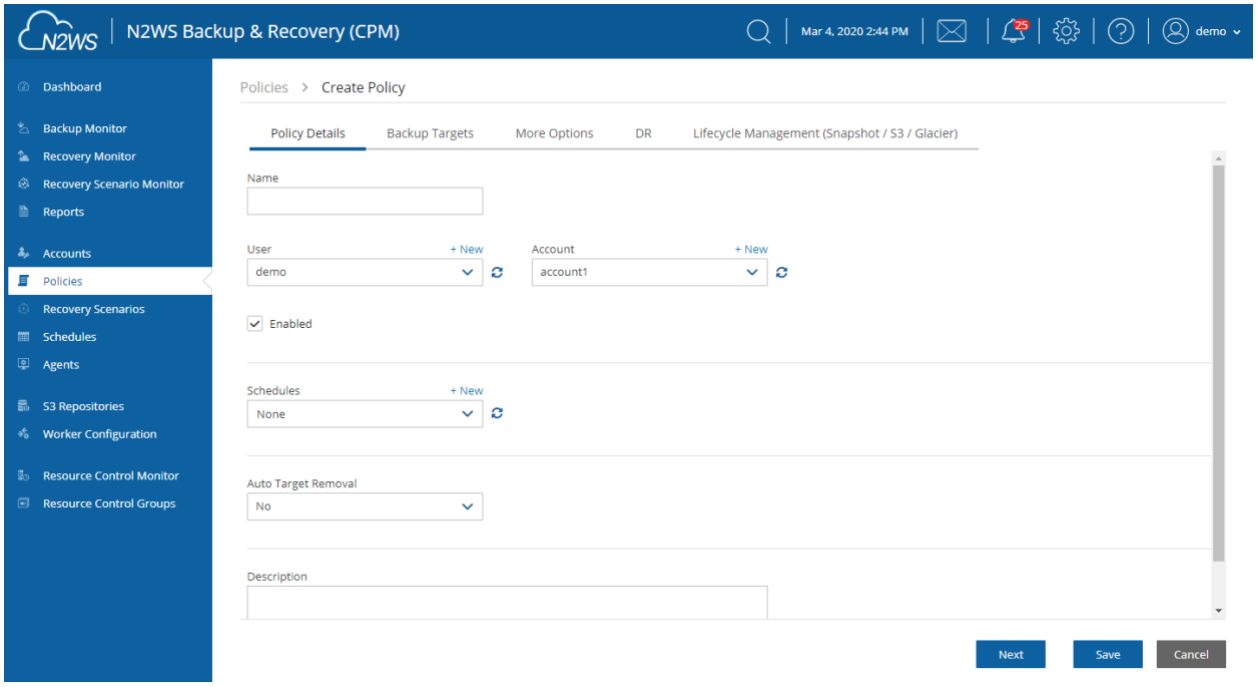

2. In the **Name** box, type a name for the policy.
3. For the root/admin user, if you have created additional managed users, select the policy owner in the **User** box.
4. If you have more than one account, in the **Account** list, select the account that the policy is associated with. The account cannot be modified after the policy has been created.

   Note: In order to avoid a Policy creation error, if the **Account** does not exist yet, click + **New** to create the policy and then return to the Policy creation**.**

5. In the **Auto Target Removal** list, specify whether to automatically remove resources that no longer exist. If you enable this removal, if an instance is terminated, or an EBS volume deleted, the next backup will detect that and remove it from the policy. Choose **yes and alert** if you want the backup log to include a warning about such a removal.
6. To use application consistent scripts as the default, in the `cmpdata` policy, select **Application Consistent**.

   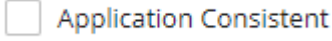

7. In the **Description** box, optionally type a description.
8. To add Backup Targets now, select the **Backup Targets** tab. See section 4.2.2. The policy will be saved when you save the Backup Targets definition.
9. To add Backup Targets later, in the **Policy Details** tab, select **Save**. The new policy is included in the list of policies in the **Policies** screen. To add Backup Targets, select the new policy, select **Edit** and then select the **Backup Targets** tab.

4.2.2 Adding Backup Targets

Backup targets define what a policy is going to back up. You define backup targets by selecting the **Backup Targets** tab of the **Create Policy** screen.

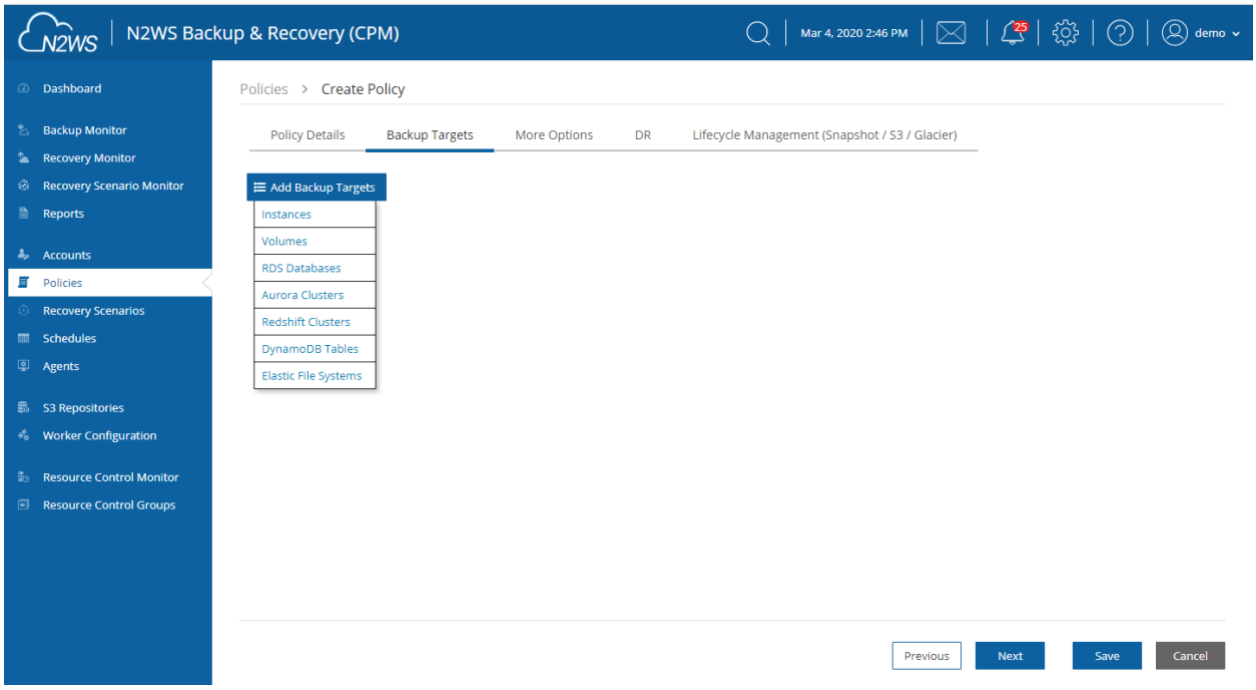

Following are the types of backup targets:

- **Instances** – This is the most common type. You can choose as many instances as you wish for a policy up to your license limit.
  For each instance, allowed under your license, define:
  - Whether to back up all its attached volumes, some, or none.
  - Whether to take snapshots (default for Linux), take snapshots with one initial AMI (default for Windows), or just create AMIs.

  See section 4.2.3.
- **EBS Volumes** – If you wish to back up volumes, not depending on the instance they are attached to, you can choose volumes directly. This can be useful for backing up volumes that may be detached part of the time or moved around between instances (e.g. cluster volumes).
- **RDS Databases** – You can use N2WS to back up RDS databases using snapshots. There are advantages with using the automatic backup AWS offers. However, if you need to use snapshots to back up RDS, or if you need to back up databases in sync with instances, this option may be useful.
- **Aurora Clusters** – Even though Aurora is part of the RDS service, Aurora is defined in clusters rather than in instances. Use this type of backup target for your Aurora clusters.
  - Aurora cluster storage is calculated in increments of 10 GiB with the respect to the license. For example, if you have over 10 GiB of data but less than 20 GiB, your data is computed as 20 GiB.
  - Keep in mind that clusters can grow dynamically and may reach the storage limits of your license. If the total storage is approaching your license limit, N2WS will issue a warning.
- **Redshift Clusters** – You can use N2WS to back up Redshift clusters. Similar to RDS, there is an automatic backup function available, but using snapshots can give an extra layer of protection.

- **DynamoDB Tables** – You can use N2WS to back up DynamoDB tables. The recommended best practice is a backup limit of 10 DynamoDB tables per policy.
    - When defining your backup targets, keep in mind that DynamoDB table storage is calculated in increments of 10 GiB with the respect to the license. For example, if you have over 10 GiB of data but less than 20 GiB, your data is computed as 20 GiB.
    - Tables can grow dynamically and may reach the storage limits of your license. If the total storage is approaching your license limit, N2WS will issue a warning.
- **Elastic File Systems (EFS)** – You can use N2WS to back up and restore your EFS snapshot data to AWS using AWS Backup service. Configuration options include backup vault, IAM role, transition to cold storage, and expiration.

In the **Add Backup Targets** drop-down menu, select the relevant resource type. A window containing a list of the available resources for the selected type and region opens. To expand the window, select **Expand** in the upper right corner. Select **Expand** again to collapse the window.

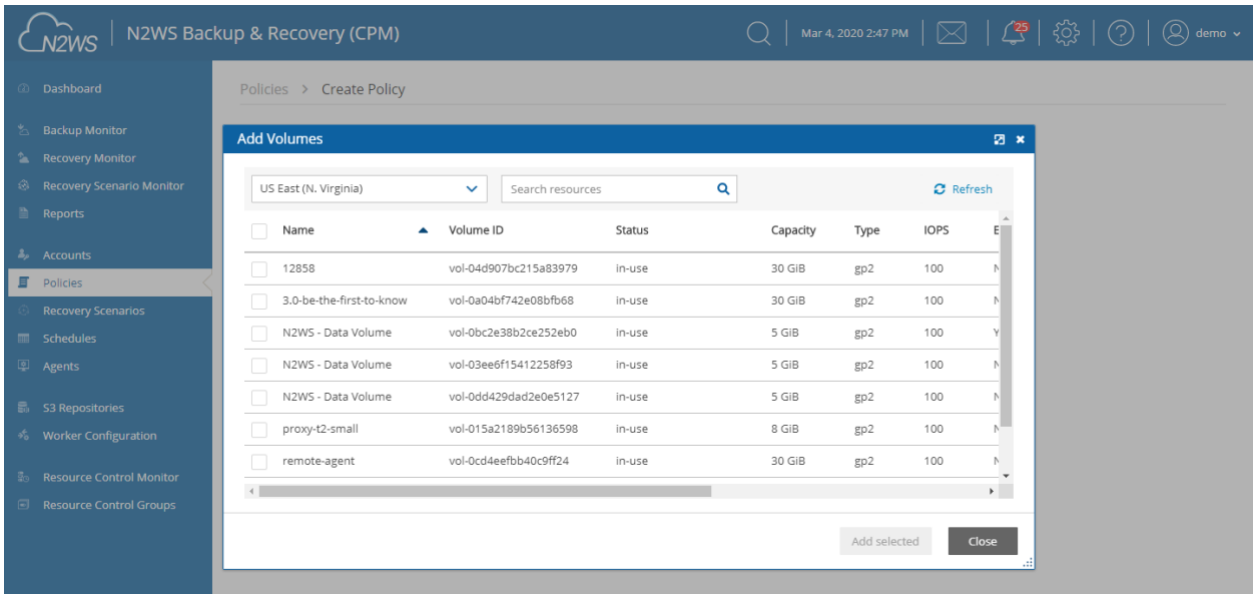

When adding backup targets of the resource type to the policy:

- You will see all the backup targets of the requested type that reside in the current region, except the ones already in the policy.
- You can select targets in another region by choosing from the region drop-down list.
- If there are many targets, you have the ability to:
    - Filter by typing part of a resource name in the search box and select **Search**. To clear a search, select **x**.
    - Sort by a column by selecting its heading. Select it again to change the sort direction.
    - Scroll between pages.
- For each backup target, the **Policies** column shows the policy, or number of policies, the target is already in. Select the link to see which policies it is in.

To add a backup target to the policy:

1. Select a region in the drop-down list. The resources in the selected region are shown.
2. To filter the resources in the region, enter all or part of a resource name in the Search resources box and select **Search** .
3. Select the check box of the desired target resources.
4. Select **Add Selected**. The selected objects are added to the policy's backup target list.
5. Repeat as many times as needed, from multiple regions if relevant.
6. Select **Close** when finished. The selected targets are listed in the **Backup Targets** tab.

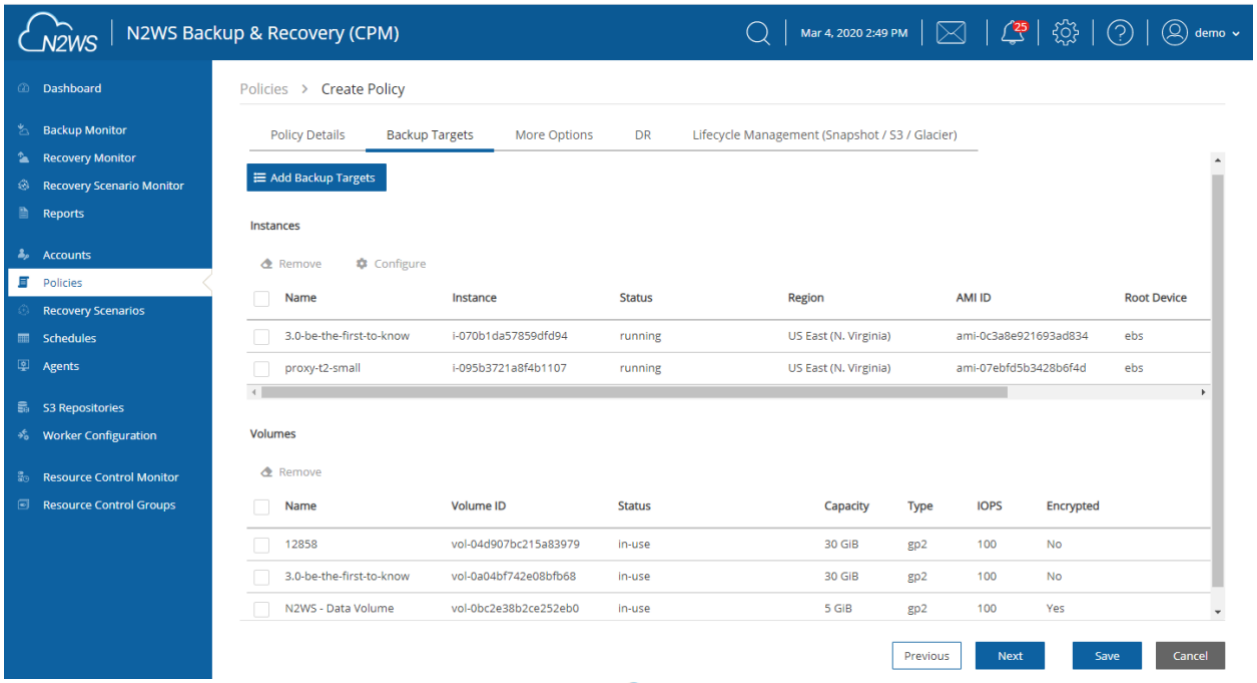

7. For Instances, Volumes, and EFS targets, select **Configure** and complete the backup options.
8. In the **Backup Targets** screen, select **Save** to save the listed selections to the policy.

4.2.3 Instance Configuration

Note: With N2WS you can now create multi-volume snapshots, which are point-in-time snapshots for all EBS volumes attached to a single EC2 instance. After it's created, a multi-volume snapshot behaves like any other snapshot. You can perform all operations, such as restore, delete, and copy across Regions and Accounts. After creating your snapshots, they appear in your EC2 console created at the exact point-in-time.

In the case of EC2 instances, you can set options for any instance.
By default, Copy to S3 is performed incrementally. In order to ensure the correctness of your data, you can force the copy of the full data for a single iteration to your S3 Repository. While defining the **Backup Targets** for a policy with Copy to S3, select **Force a single full copy to S3.**

To configure an instance:

1. Select a policy, select the **Backup Targets** tab, and then select an instance.
2. Select **Configure**. The Policy Instance and Volume Configuration screen opens.

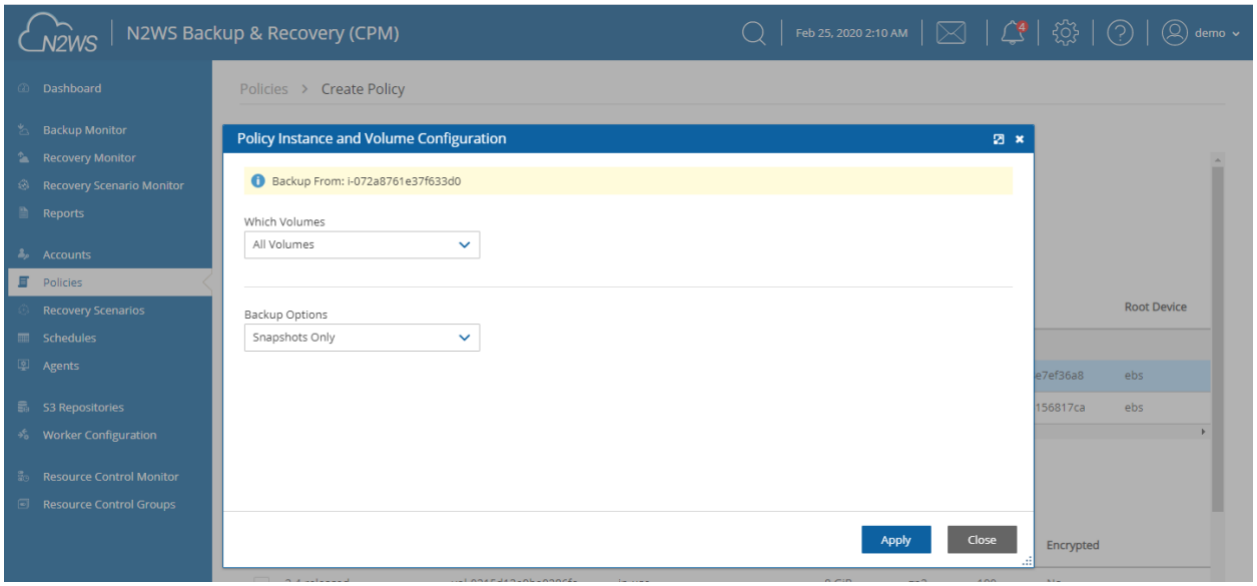

3. In the **Which Volumes** list, choose whether to back up all the volumes attached to this instance, or include or exclude some of them. By default, N2WS will back up all the attached storage of the instance, including volumes that are added over time.
4. If **All Volumes** was not selected, in the volumes table, clear or select the volume check boxes.

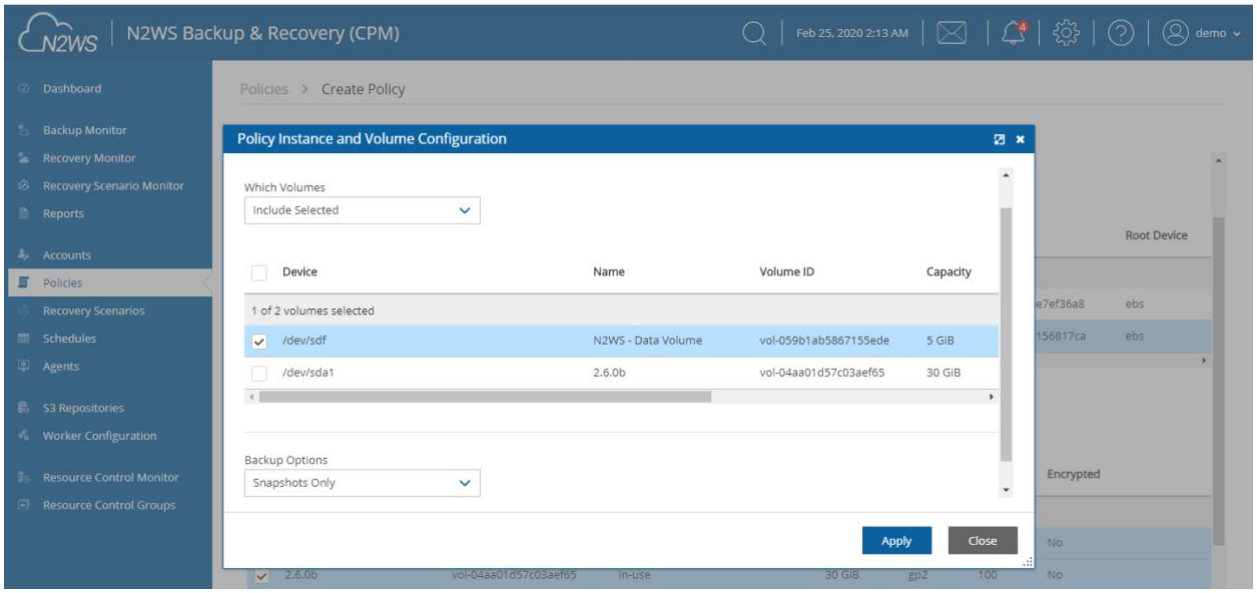

5. In the **Backup Options** list, select one of the following:
   - **Snapshots Only -** Default for Linux-based instances
   - **Snapshots with Initial AMI -** Take an initial AMI and then snapshots- the default for Windows-based instances
   - **AMIs Only** - Just schedule AMI creation. If a reboot is required after the backup, select **Reboot**. See section 4.2.3.1.

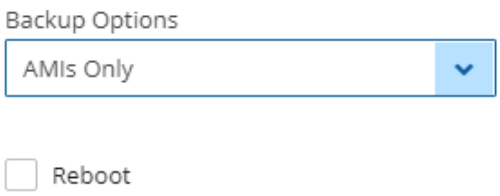

6. When Copy to S3 is enabled for the policy, to have a full copy of the data copied to your S3 Repository, select **Force a single full copy to S3.**
7. Select **Apply**.
8. Continue to add instances from other regions, and select **Close** when finished with the resource type. Control returns to the **Backup Targets** tab list.
9. At the bottom of the **Backup Targets** tab list, select **Save** to update the policy for all listed targets.

4.2.3.1 AMI Creation

If you choose to just create AMIs:

- N2WS will create AMIs for that instance instead of taking direct snapshots. App-aware backup per agent does not apply for AMI creation.
- You can choose whether to reboot the instance during AMI creation or not to reboot, which leaves a risk of a data corruption. As opposed to AMI creation in the EC2 console, the default in N2WS is no reboot.

Note:	Try not to schedule AMI creations of an instance in one policy, while another policy running at the same time backs up the same instance using snapshots. This will

cause the AMI creation to fail. N2WS will overcome this issue by scheduling a retry, which will usually succeed. However, it is recommended that you avoid such scheduling conflicts.

4.2.3.2 Initial/Single AMI

Single or Initial AMIs are meant to be used mainly for Windows instance recovery.

- N2WS will keep a single AMI for each instance with this setting. A single AMI will contain *only* the root device (boot disk).
- N2WS will rotate single AMIs from time to time. It will create a new AMI and delete the old one. N2WS aims to optimize cost by not leaving very old snapshots in your AWS account.
- By default, N2WS will rotate single AMIs every 90 days. That value can be configured in in the **Server Settings** > **General Settings** > **Cleanup** tab to any number of days, or to 0, if you prefer no rotation at all.

4.2.3.3 Limitations with AMI creation:

AMIs will be copied across region for DR, but they will not be copied across accounts.

Important: If you use cross-account backup, be aware that if you need to recover the instance at the remote account, you need to make sure you have an AMI ready in that account.

4.2.4 More Options

To see more policy options, select the **More Options** tab for a policy in the **Policies** tab. The options consist of:

- Activating Linux backup scripts and defining script timeout and output
- Defining backup success, retries, and failures for alerting

Backup scripts refers to those running on the N2WS server. See section 7.2.

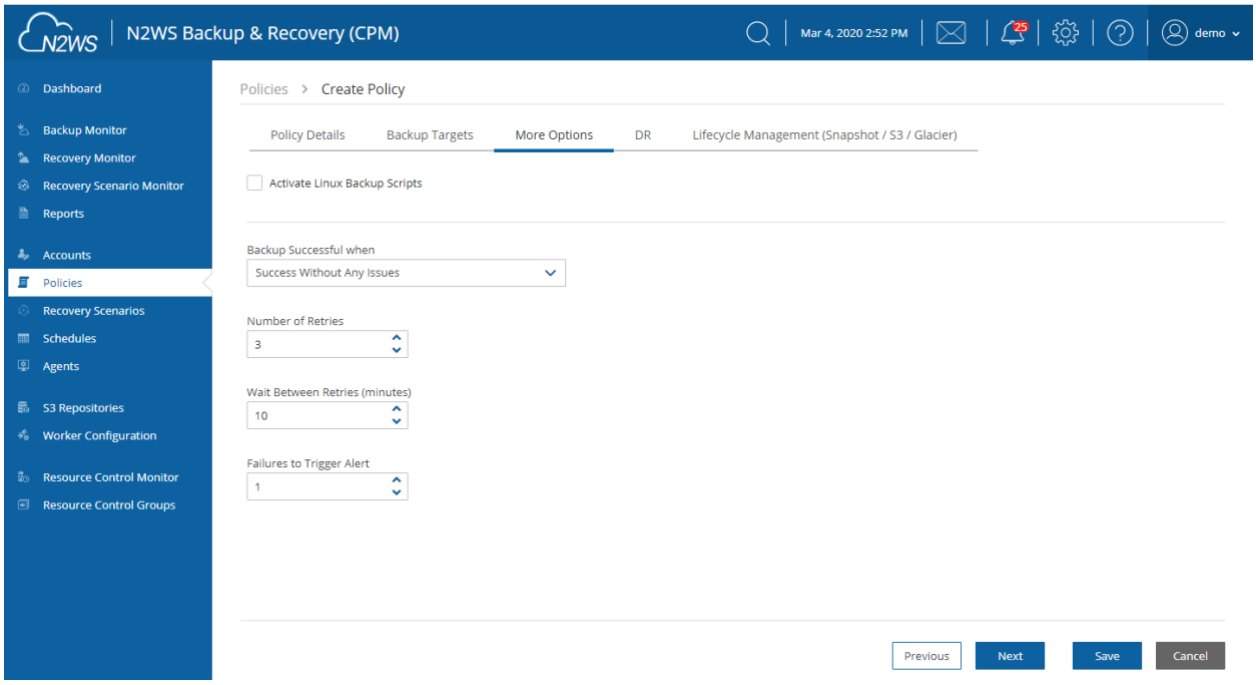

- **Activate Linux Backup Scripts** – This option is turned off by default. If you select **Activate Linux Backup Scripts**, the following options appear:
  - **Scripts Timeout** – Timeout (in seconds) to let each script run. When a backup script runs, after the timeout period, it will be killed, and a failure will be assumed. The default is 30 seconds.
  - **Collect Scripts Output** – N2WS can collect the output of backup scripts to the standard error (stderr). This may be useful for debugging. It can also be used by a script to log the operations it is performing and write useful information. This output is captured, saved in the N2WS database, and can be viewed from the Recovery Panel screen. To turn this option on, choose Collect. The default option is Ignore.

    Note: The output of a script is typically a few lines. However, if it gets really big (MBs), it can affect the performance of N2WS. If it gets even larger, it can cause crashes in N2WS processes. To avoid the risk of too much data going to `stderr`, redirect the output elsewhere.
- **Backup Successful when** - This indicates whether a backup needs its scripts/VSS to complete, in order to be considered a valid backup. This has a double effect:
  - For retries, a successful backup will not result in a retry;
  - For the automatic retention management process, a backup which is considered successful is counted as a valid generation of data.
    The possible values are:
    - **Success Without Any Issues** – If scripts and/or VSS are defined for this policy, the backup will be considered successful only if everything succeeds. If backup scripts or VSS fails and all snapshots succeed, the backup is not considered successful. You can still recover from it, but it will cause a retry (if any are defined), and the automatic retention management process will not count it as a valid generation of data. This is the stricter option and is also the default.

- **Snapshot Succeeded with Possible VSS or Script Failure** – This is the less strict option and can be useful if scripts or VSS fail often, which can happen in a complex environment. Choosing this option accepts the assumption that most applications will recover correctly even from a crash-consistent backup.

- Retry information - The last three options deal with what to do when a backup does not succeed:
  - **Number of Retries** – The maximal number of retries that can be run for each failed backup. The default is three. After the retries, the backup will run again at the next scheduled time.
  - **Wait Between Retries (minutes)** – Determines how much time N2WS will wait after a failure before retrying. Backup scripts and VSS may sometimes fail or timeout because the system is busy. In this case, it makes sense to substantially extend the waiting time until the next retry when the system may be more responsive.
  - **Failures to Trigger Alert** – The minimum number of failures to trigger an alert.

4.2.5 Enabling Disaster Recovery

To enable Disaster Recovery for the policy, select the **DR** tab for a policy in the **Policies** tab screen.

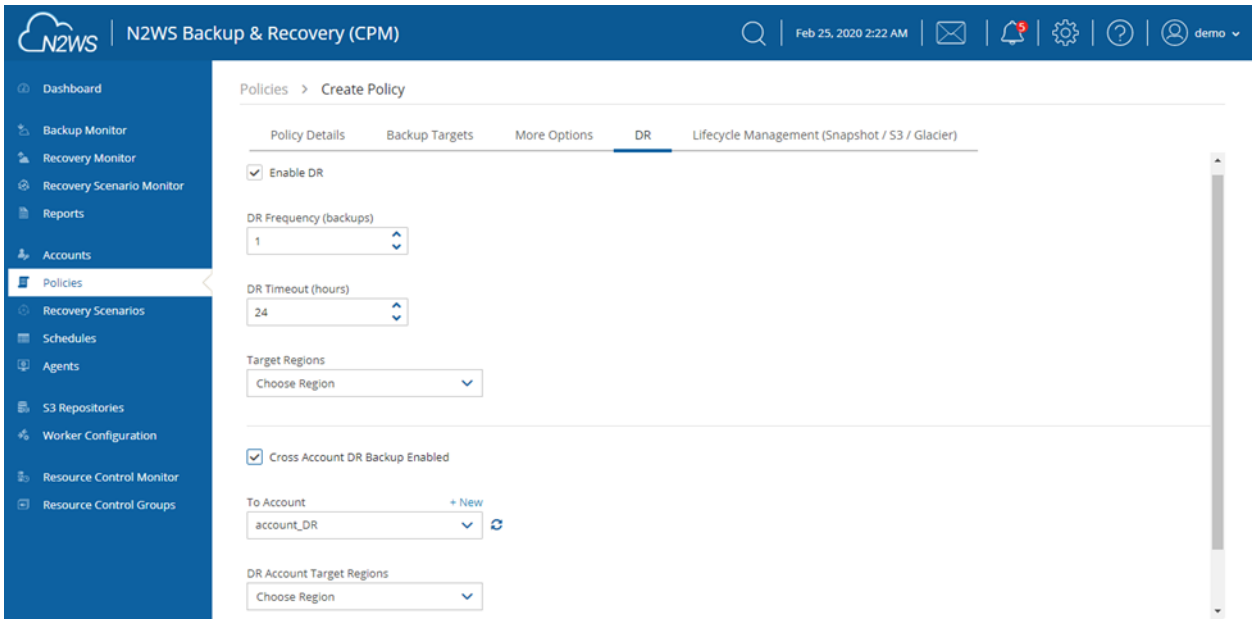

1. Select **Enable DR**.
2. Select the **DR Frequency** for backups, **DR Backup Timeout**, and **Target Regions**.
3. To enable **Cross Account DR Backups**, select the check box, and:
   a. Choose the **To Account** and **DR Account Target Regions** in the lists.

   Note: If the DR **To Account** does not exist yet, click **New** to create the account and then return to the policy creation.

   b. To **Keep Original Snapshots**, select the check box.

4.2.6 Managing Lifecycle

The **Lifecycle Management** tab for a policy allows you to do the following:

- Set retention policy for snapshots.
- Backup to S3 – Enable, set copy and retention policies, and choose storage options. See section 21.3.
- Archive to Glacier - Enable, set copy and retention policies, and choose the Storage Class. See section 21.4.3.

Note: **Archive to Glacier** requires that the **Backup to S3** option is enabled first.

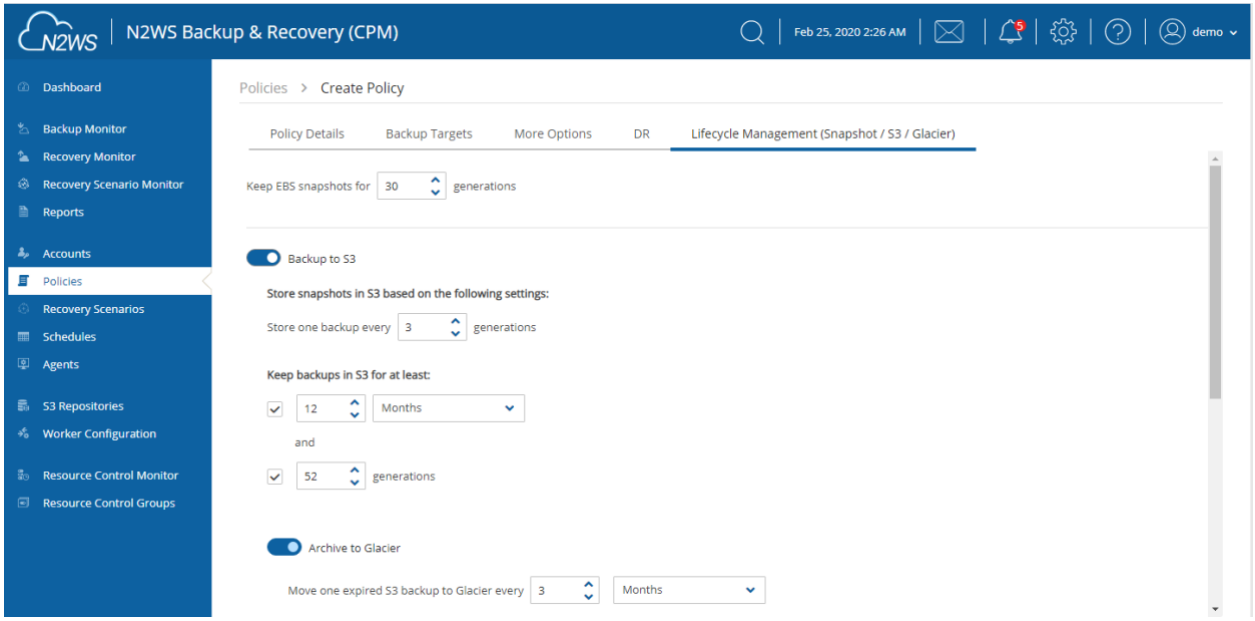

4.2.7 Viewing Policy Backup Times

In the **Policies** tab, select a policy and then select **Backup Times** to open the Planned Backups window, which is ordered by Time Zone. You can change the **Time Period** from the default 'Next Day' to a number of days, weeks, or the next month.

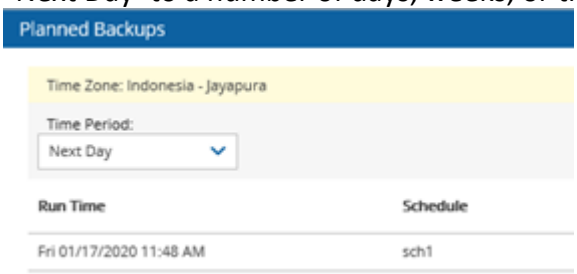

Backup Times are relevant to the schedules of the selected policy.

4.2.8 Running an Ad Hoc Backup

An ad hoc backup will execute the selected Policy and create backups of all its targets.

Note: An ad hoc backup counts as another generation if it completes successfully.

To run a backup immediately:

1. In the left panel, select the **Policies** tab and then select a policy in the list.
2. To start the backup, in the Policies tab, select **Run ASAP.**
3. To follow the progress of the backup, select the **Open Backup Monitor** link in the 'Backup started' message Backup started (Open Backup Monitor) at the top right corner, or select the **Backup Monitor** tab. To update the list, select **Refresh**.

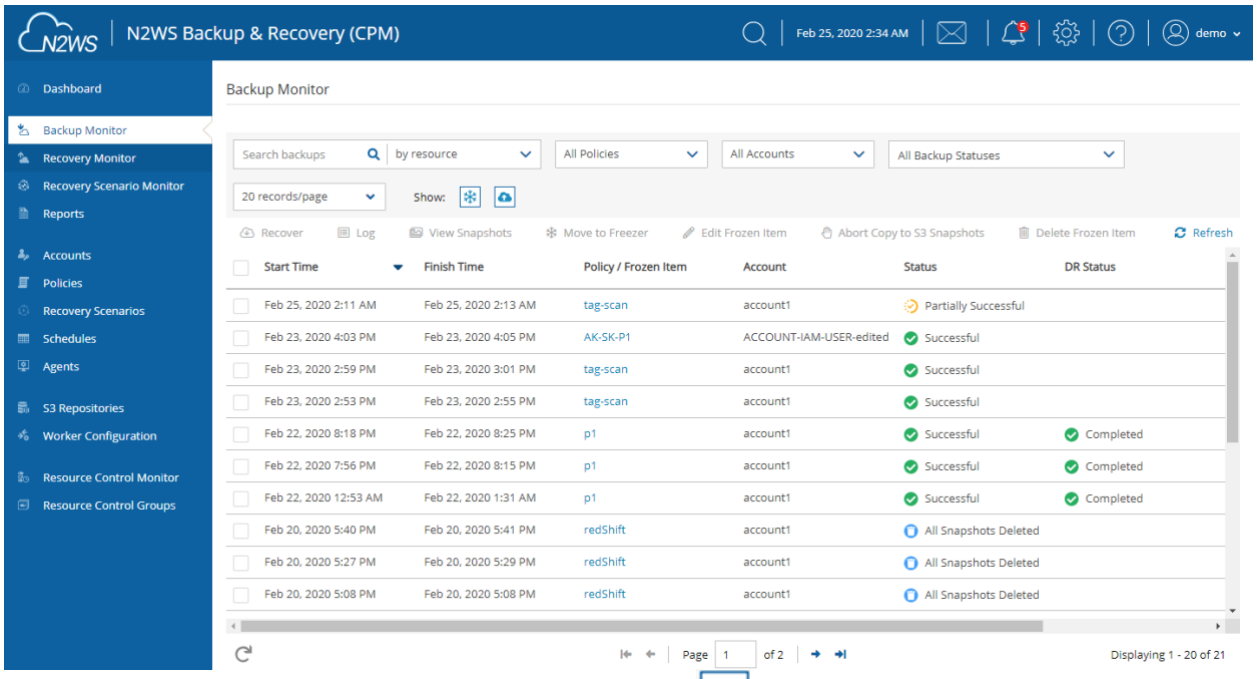

   a. You can switch between showing backup records *in* the Freezer and backup records *not* in the Freezer. In the **Show** area on the far right of the filters line:
      - To show only the backup records that were moved to the Freezer, select **Backup** once. Select **Backup** again to show all backup records.
      - To show only the backup records that were *not* moved to the Freezer, select **Freezer** once. Select **Freezer** again to show all backup records.
4. To view or download backup details, select **Log**. Select **Refresh** as needed.

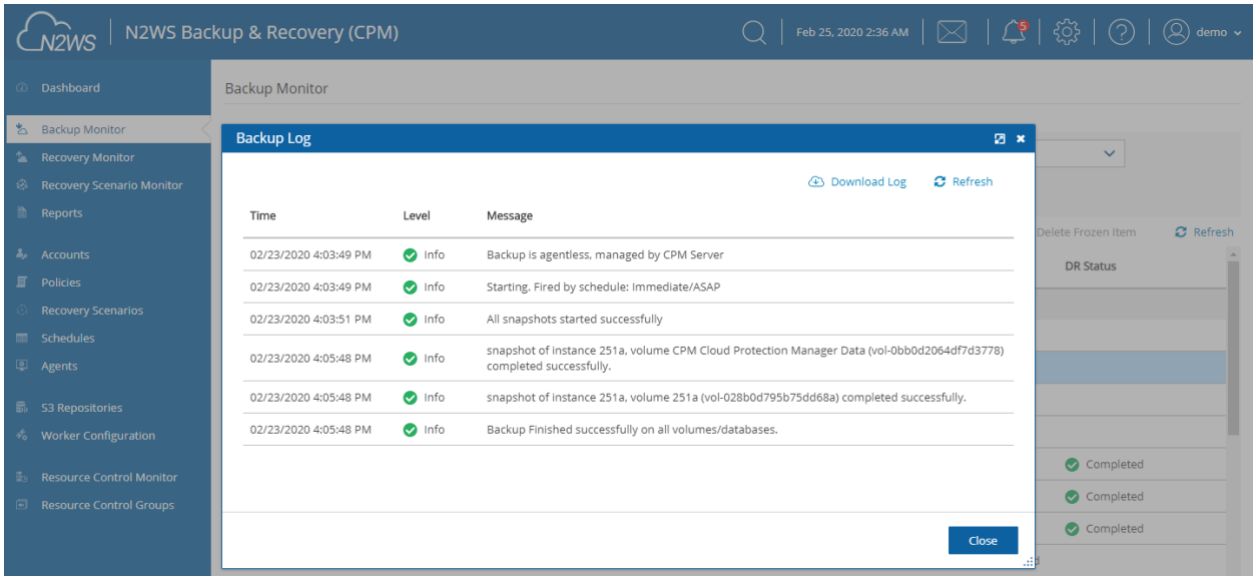

5. To delete a particular snapshot in AWS or to delete all AWS snapshots after a successful run, select **View Snapshots**. Select one or more backups and then select **Delete**.

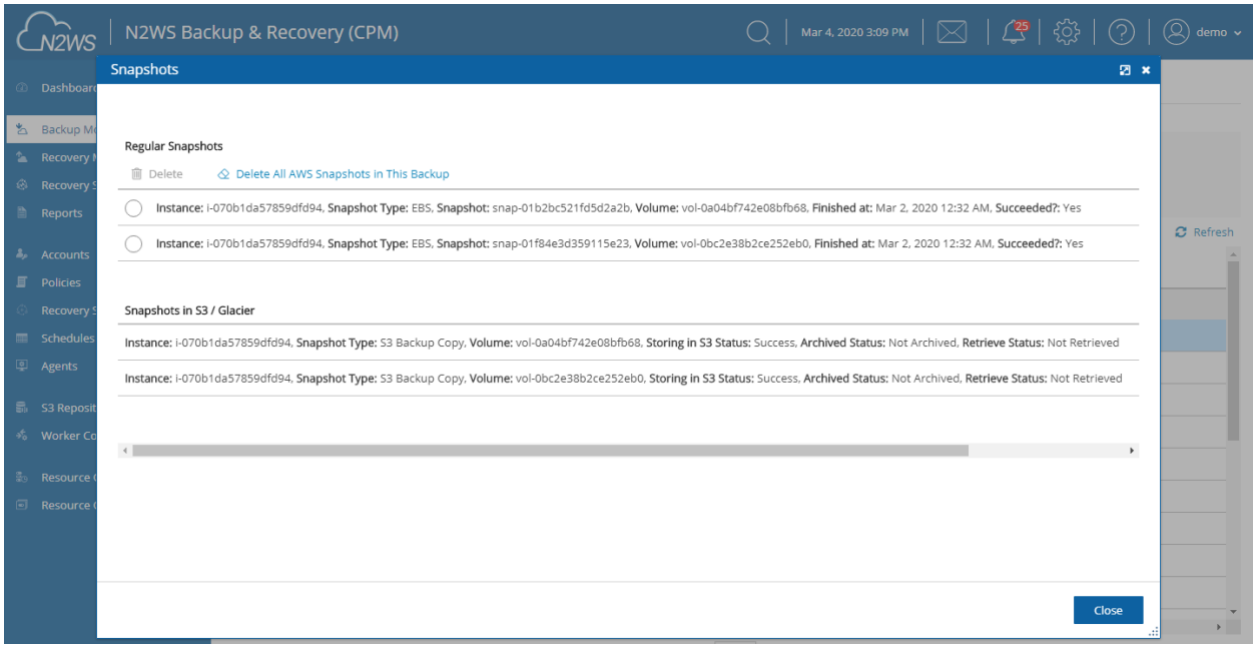

5 Consistent Backup

This guide explains a few key concepts to help you use N2WS correctly.

5.1 Crash-Consistent Backup

By default, snapshots taken using N2WS are Crash-consistent. When you back up an EC2 instance at a certain time, and later want to restore this instance from backup, it will start the same as a physical machine booting after a power outage. The file system and any other applications using EBS volumes were not prepared or even aware that a backup was taking place, so they may have been in the middle of an operation or transaction.
Being in the middle of a transaction implies that this backup will not be consistent, but actually this is not the case. Most modern applications that deal with important business data are built for robustness. A modern database, be it MySQL, Oracle or SQL Server, has transaction logs. Transaction logs are kept separately from the data itself, and you can always play the logs to get to a specific consistent point in time. A database can start after a crash and use transaction logs to get to the most recent consistent state. NTFS in Windows and EXT3 in Linux have implemented journaling, which is not unlike transaction logs in databases.

5.2 Application-Consistent Backup

During application-consistent backups, any application may be informed about the backup progress. The application can then prepare, freeze and thaw **in minimal required time** to perform operations to make sure the actual data on disk is consistent before the backup starts., making minimal changes during backup time (**backup mode**) and returning to full scale operation as soon as possible.
There is also one more function that application-consistent backups perform especially for databases. Databases keep transaction logs which occasionally need to be deleted to recover storage space. This operation is called **log truncation**. When can transaction logs be deleted without impairing the robustness of the database? Probably after you make sure you have a successful backup of the database. In many cases, it is up to the backup software to notify the database it can truncate its transaction logs.

5.3 N2WS and a Point in Time

When taking snapshots, the **point in time** is the exact time that the snapshot started. The content of the snapshot reflects the exact state of the disk at that point in time, regardless of how long it took to complete the snapshot.

5.4 Summary or What Type of Backup to Choose

The type of backup to choose depends on your needs and limitations. Every approach has its pros and cons:

5.4.1 Crash-Consistent

Pros:

- Does not require writing any scripts.
- Does not require installing agents in Windows Servers.
- Does not affect the operation and performance of your instances and applications.
- Fastest.

Cons:

- Does not guarantee consistent state of your applications.
- Does not guarantee exact point in time across multiple volumes/disks.
- No way to automatically truncate database transaction logs after backup.

5.4.2 Application-Consistent

Pros:

- Prepares the application for backup and therefore achieves a consistent state.
- Can ensure one exact point in time across multiple volumes/disks.
- Can truncate database transaction logs automatically.

Cons:

- May require writing and maintaining backup scripts.
- Requires installing a N2WS Thin Backup Agent for Windows Servers.
- May slightly affect the performance of your application, especially for the freezing/flushing phase.

6 Windows Instances Backup

From the point of view of the AWS infrastructure, there is not much difference between backing up Linux/Unix instances or Windows instances. You can still run snapshots on EBS volumes. However, there is one substantial difference regarding recovering instances:

- In Unix/Linux instances, you can back up system volumes (root devices), and later launch instances based on the snapshot of the system volume. You can create an image (AMI) based on the system volume snapshot and launch instances.
- This option is currently not available for Windows Servers. Although you can take snapshots of the system volume of a Windows Server, you cannot create a launchable image (AMI) from that snapshot.
  Because of this limitation, N2WS needs an AMI to start a recovery of a Windows instance. N2WS will still make sure all the volumes, including the root device (OS volume) will be from the point-in-time of the recovered backup. By default, N2WS will create an initial AMI when you start backing up a Windows instance. That AMI will be used as the default when recovering this instance.

6.1 Configuring N2WS Thin Backup Agent

If crash-consistent backup is sufficient for your needs, you do not need to install any agent. However, to use VSS or run backup scripts, you will need to install N2WS Thin Backup Agent. The N2WS Thin Backup Agent is used for Windows instances that need to perform application quiescence using VSS or backup scripts.

- The agent communicates with the N2WS Server using the HTTPS protocol.
- No sensitive information passes between the backup agent and the N2WS Server.

Any Windows instance in a policy can have a backup agent associated with it.

6.1.1 Associating an Agent with a Policy

After adding your Windows instance in the backup targets page (section 4.2.2), the next step is to configure its agent by associating it with a `cpmdata` policy.

To associate an agent with the cpmdata policy:

1. In the **Policies** tab, select the 'cpmdata' policy and then select **Edit.**
2. In the **Policy Details** tab, select **Application-Consistent**.

By default, VSS quiescence will be activated for this policy.

Note:	In case the agent represents a Windows 2003 instance, VSS will fail every time. You need to turn off this option and use only backup scripts. If you have a Windows 2003 instance and you do not need scripts, there is no use installing an agent, so just perform backups without one.

6.1.2 Downloading and Distributing the Agents Configuration

1. You can download the installation package of the agent from the **Agents** tab in the left panel. Select **Download Thin Backup Agent** to download a standard Windows package (`CPMAgentService.msi`) to the Downloads folder.
2. After downloading the agent installation package, select **Server Settings** and then select **Agents Configuration**.
3. Enter the configuration syntax as described in Appendix B (page 222), and select **Publish**.

6.1.3 Installing the Agent

The agent can be installed on any Windows 2003, 2008, 2012, 2016, or 2019 instance, 32 or 64-bit. After accepting the license agreement, the Setup screen opens.

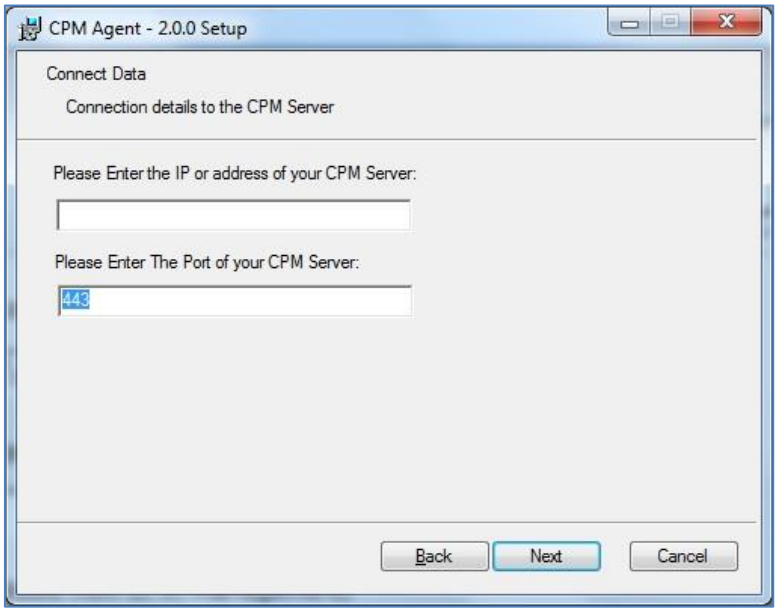

The required fields are:

- The address of the N2WS server that is reachable from this instance.
- The default port is 443 for HTTPS communication. Change it if you are using a custom port.

After finishing the installation, the N2WS agent will be an automatic service in your Windows system.

Important: After the Agent is installed and configured and a policy with target that points at it is configured and enabled, the users must wait to see it registered in the remote **Agents** screen in the N2WS. It may take a few minutes until the N2WS connects.

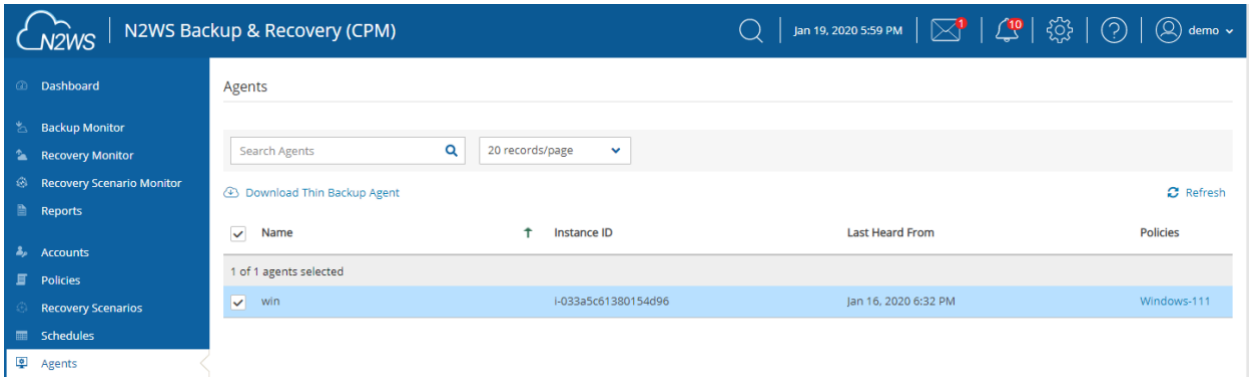

6.1.4 Changing Agent Configuration

To change the configuration of the backup agent after installation, edit the backup agent configuration file.

To change the agent configuration file:

1. Before proceeding, it is recommended that you make a copy of the `cpmagent.cfg` file, which resides in the N2WS Agent installation folder.
2. If the address or port of the N2WS Server had changed, edit the agent configuration file manually. Make the change after the equation sign.
3. After making the changes, restart the **N2WS Agent Service**, in the Windows Service Manager console.

   As an alternative, you could uninstall and reinstall the agent.

6.1.5 Using the Agent with an HTTP Proxy

N2WS supports configurations where the backup agent for a Windows instance can reach the CPM server only through a proxy.

To configure the agent with an HTTP proxy:

1. See section 6.1.4 about editing `cpmagent.cfg`, and:
2. Add the following lines under the general section:
   ```
   proxy_address=<dns name or ip address of the proxy server>
   proxy_port=<port for the proxy (https)>
   ```
3. If your proxy server requires authentication, add the following two lines as well:
   ```
   proxy_user=<proxy user name>
   proxy_password=<proxy password>
   ```
4. Restart the N2WS Agent service from the Windows Service Manager.

6.2 Using VSS

VSS, or Volume Shadow Copy Service, is a backup infrastructure for Windows Servers. It is beyond the scope of this guide to explain how VSS works. You can read more at http://technet.microsoft.com/en-us/library/cc785914%28v=WS.10%29.aspx. However, it is important to state that VSS is the standard for Windows application quiescence, and all recent releases of many of the major applications that run on Windows use it, including Microsoft

Exchange, SQL Server, and SharePoint. It is also used by Windows versions of products not developed by Microsoft, like Oracle.
N2WS supports VSS for backup on Windows Servers 2008, 2012, 2016, and 2019 *only*. Trying to run VSS on other Windows OSs will always fail. VSS is turned on by default for every Windows agent. For unsupported OSs, you will need to disable it yourself. This can be done in the instance configuration screen, see section 6.1.1.

Any application that wishes to be **backup aware** has a component called **VSS Writer.** Every vendor who would like to support copying the actual backup data (making shadow copies) provides a component called a **VSS Provider**. The operating system comes with a **System Provider**, which knows how to make shadow copies to the local volumes. Storage hardware vendors have specialized **Hardware Providers** that know how to create shadow copies using their own hardware snapshot technology. Components that initiate an actual backup are called **VSS Requestors**.
When a requestor requests a shadow copy, the writers flush and freeze their applications. At the point of time of the shadow copy, all the applications and the file systems are frozen. They all get thawed after the copy is started (copy-on-write mechanisms keep the point in time consistent, similar to EBS snapshots). When the backup is complete, the writers get notified that they have a consistent backup for the point in time of the shadow copy. For example, Microsoft Exchange automatically truncates its transaction logs when it gets notified that a backup is complete.

6.2.1 N2WS' Use of VSS

The N2WS Agent performs under the role of a **VSS Requestor** to request the VSS **System Provider** to perform shadow copies. The process is:

- When N2WS initiates a backup, it **requests** the N2WS Backup Agent to invoke a backup of all relevant volumes. The agent then requests the VSS System Provider to start the shadow copy.
- VSS only creates differential copies, which means that in order for the N2WS to fully backup each volume, a few extra MBs are needed for the backup to complete. The amount of MBs depends on the size of the volume and the amount of data written since last backup. Once the backup is complete, the N2WS agent will request the VSS Provider to delete the shadow copies. The N2WS Agent will notify all relevant VSS writers that the backup is complete, only after making sure all the EBS snapshots are completed successfully.

You can see the process depicted below.

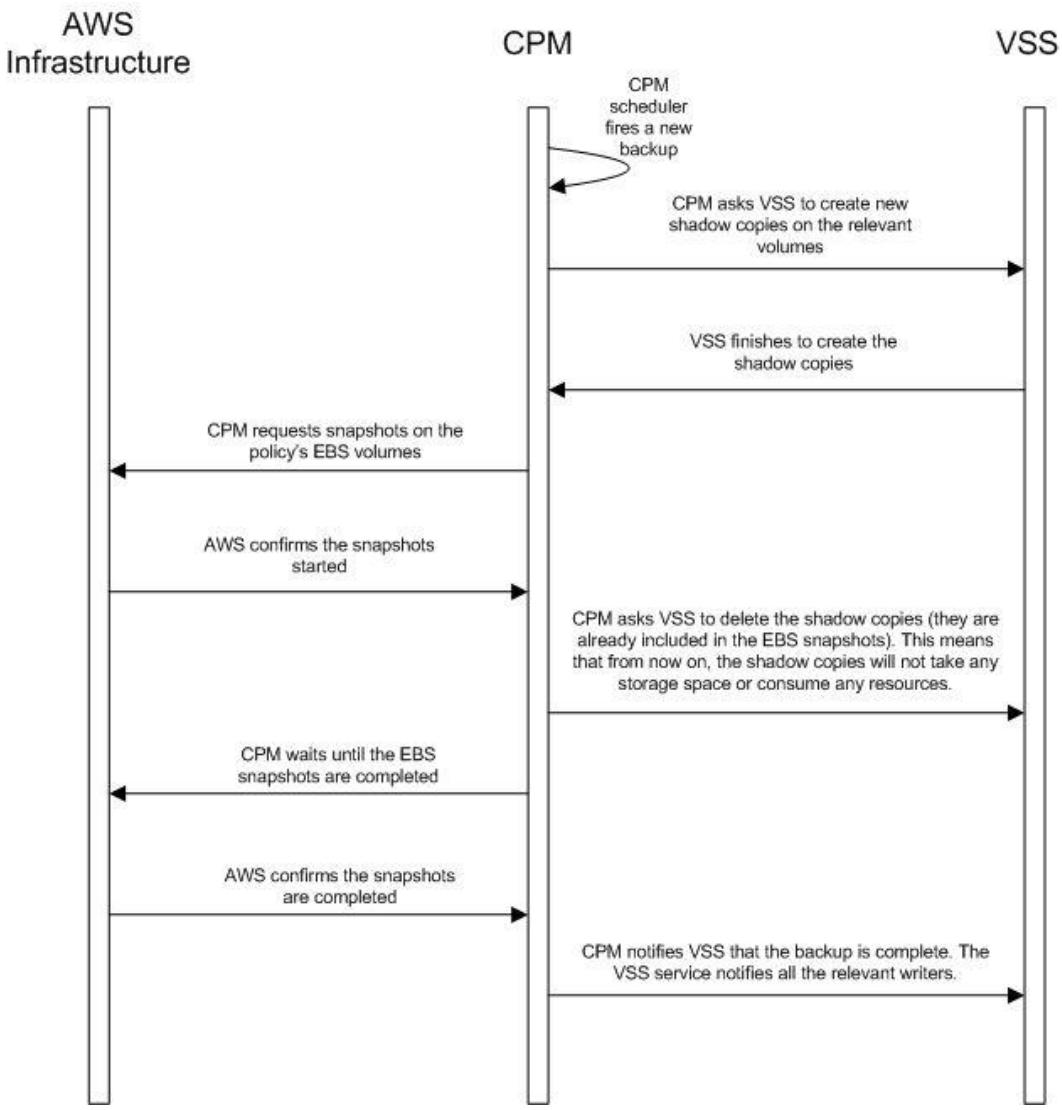

6.2.2 Configuring VSS

By default, VSS is enabled when a N2WS Thin Backup Agent is associated with an instance in a policy. In many cases, you will not need to do anything. By default, VSS will take shadow copies of all the volumes. However, you may want to exclude some volumes. For example, since the system volume (typically C:\) cannot be used to recover the instance in a regular scenario, you may want to exclude it from the backup.

To make shadow copies of only some of the volumes:

1. In the Instance and Volume configuration screen, change the value of **Volumes for shadow copies**.
2. Type drive letters followed by a colon, and separate volumes with a comma, e.g. **D:, E:, F:.**

6.2.3 Excluding and Verifying VSS Writers

Writer exclusions and inclusions are configured using a text file, not the UI.
You may wish to exclude **VSS Writers** from the backup process in cases where the writer is:

- Failing the backup.

- Consuming too many resources.
- Not essential for the backup's consistency.

To exclude VSS writers:

In the subfolder `scripts` under the installation folder of the Thin Backup Agent (on the backed-up instance), create a text file named `vss_exclude_writers_<policy name>.txt`. with the following structure:

- Each line will contain a writer ID (including the curly braces)
- If you write in one of the lines `all`, all writers will be excluded. This can be handy sometimes for testing purposes.

In some cases, you want to make sure that certain writers are included (verified) in the shadow copy, and if not, fail the operation.

To verify writers:

In the subfolder `scripts` under the installation folder of the Thin Backup Agent (on the backed-up instance), create a text file named `vss_verify_writers_<policy name>.txt` with the following structure:

- Each line will contain a writer ID (including the curly braces).
- `all` is not an option.

An example for a line in either of the files is:

```
{4dc3bdd4-ab48-4d07-adb0-3bee2926fd7f}
```

6.2.4 Troubleshooting VSS Issues

When a VSS-enabled policy runs, you will see its record in the backup monitor tab of N2WS's main screen.

- If it finished with no issues, the status of the record will be **Backup Successful**.
- If there were issues with VSS, the status will be **Backup Partially Successful**.

To troubleshoot:

- To view the errors that VSS encountered, look in the backup log.
- To view the output of the exact VSS error, select **Recover**.
- To view the VSS Disk Shadow log, select its link in the recovery panel. There is a link for each of the agents in the policy, with the instance ID stated.
- In most cases, VSS will work out of the box with no issues. There can be a failure from time to time in stressed system conditions.
- If the writers do not answer to the **freeze** request fast enough, the process times out and fails. Often, the retry will succeed.
- When VSS is constantly failing, it is usually a result of problems with one of the writers. This could be due to some misconfiguration in your Windows system.
- In most cases the problem is out of the scope of N2WS. The best way to debug such an issue is to test VSS independently. You can run the Diskshadow utility from a command line window and use it to try and create a shadow copy. Any issue you have with VSS using N2WS should also occur here.
- To learn how to use the Diskshadow utility, see: https://docs.microsoft.com/en-us/windows-server/administration/windows-commands/diskshadow

- You may see failures in backup because VSS times out or is having issues. You will see that the backup has status **Backup Partially Successful**. Most times you will not notice it since N2WS will retry the backup and the retry will succeed.
- If the problem repeats frequently, check that your Windows Server is working properly. You can check the application log in Window's Event Log. If you see VSS errors reported frequently, contact N2W Software support.

6.2.5 VSS Recovery

Recovering instances using N2WS is covered in section 10. When recovering a Windows Server that was backed up with VSS, you need to revert back to the shadow copies in the recovery volumes to get the consistent state of the data.

To revert back to shadow copies after VSS recovery:

1. Connect to the newly recovered instance.
2. Stop the services of your application, e.g. SQL Server, Exchange, SharePoint, etc.
3. Open an administrator command line console and type `diskshadow`.
4. In the recovery panel screen, select the **VSS DiskShadow Data** link to find the IDs of the shadow copies made for the required backup.
5. Type `revert {shadow id}` for each of the volumes you are recovering, except for the system volume (C: drive). After finishing, the volumes are in a consistent state.
6. Turn the services on and resume work.

If you wish to recover a system disk, it cannot be reverted to the shadow copy using this method. The system volume should not contain actual application data as it is not a recommended configuration, and, therefore, you should be able to skip this revert operation. However, you can expose the system disk from the shadow and inspect its contents.

To expose the system disk from the shadow:

1. In the Diskshadow utility, type: `expose {shadow id} volletter:`
2. After finishing, remember to unexpose the disk.
3. To avoid unnecessary resource consumption, delete the shadow: (`delete shadow {shadow id}`).

Reverting to a shadow copy for a system volume

If you have a strict requirement to recover the consistent shadow copy for the system volume as well, do the following:

1. Before reverting for other volumes, stop the instance; wait until it is in **stopped** state.
2. Using the AWS Console, detach the EBS volume of the C: drive from the instance and attach it to another Windows instance as an "additional disk".
3. Using the Windows Disk Management utility, make sure the disk is online and exposed with a drive letter.
4. Go back to the process in the previous section (VSS Recovery), and revert to the snapshot of drive C which will now have a different drive letter. Since it is no longer a system volume, it is possible to do so.

5. Detach the volume from the second Windows instance, reattach to the original instance using the original device, which is typically `/dev/sda1`, and turn the recovered instance back on.

Note: Shadow copy data is stored by default in the volume that is being shadowed. However, in some cases it is stored on another volume. In order for you to be able to recover, you need to make sure you also have the volume the shadow copy is on included in the backup and the recovery operation.

Important: If you revert a volume that contains another volume's shadow data, the reversion will take the volume to a state where it no longer contains the second volume's backup data, as the second volume would need to be reverted before the first volume can be reverted. If you accidentally restore the shadow copies in the wrong order, just delete the recovered instance and its data volumes and begin the recovery operation again from N2WS, taking care to revert the shadow copies in the correct order.

6.3 Using Backup Scripts on Windows

Besides VSS, there is also the option to run backup scripts to achieve backup consistency. It is also possible to add backup scripts in addition to VSS.

- You enable backup scripts in the Instance and Volume Configuration screen of the instance in the policy.
- As opposed to Linux, Windows backup scripts run directly on the agent. All the scripts are located in the subfolder `scripts` under the installation folder of N2WS Thin Backup Agent.
- If the N2WS user that owns the policy is not the root user, the scripts will be under another subfolder with the user name (e.g. `…\scripts\cpm_user1`).
- All scripts are named with a prefix plus the name of the policy.
- There are 3 types of events. If scripts are used, a script must be provided for each of these events. If all of the scripts are not defined, N2WS will treat the missing script as a failing script.
    - Before the VSS backup - BEFORE_<policy name>.<ext>
    - After the VSS backup started - AFTER_<policy name>.<ext>
    - After the VSS backup has completed - COMPLETE_<policy name>.<ext>
- Scripts can have any extension as long as they are executable. They can be batch scripts, VBS scripts, Power Shell, or even binary executables. However, N2WS cannot run PowerShell scripts directly as Windows scripts.
- Scripts are launched by N2WS Thin Backup Agent, so their process is owned by the user that runs the agent service. By default, this is the local system account. However, if you need to run it under a different user you can use the service manager to change the logged-on user to a different one. For example, you might want to run it with a user who has administrative rights in a domain.
- All scripts must be set with exit code `0`.

6.3.1 Before Script

The `before_<policy name>.<ext>` runs before backup begins. Typically, this script is used to move applications to backup mode. The **before** script leaves the system in a **frozen** state. This state will stay for a very short while, until the snapshots of the policy start, which is when the **after** script is started.

6.3.2 After Script

The `after_<policy name>.<ext>` script runs after all the snapshots of the policy start. It runs shortly after the **before** script, generally less than 2-3 seconds. This script releases anything that may have been frozen or locked by the **before** script.
This script accepts the success status of the **before** script. If the **before** script succeeded, the argument will be `1`. If it failed, crashed, or timed out, the argument will be `0`.

Note: This is the opposite of the exit status. Think of it as an argument that is true when the **before** script succeeded.

6.3.3 Complete Script

The `complete_<policy name>.<ext>` script runs after all snapshots are completed. Usually the script runs quickly since snapshots are incremental. This script can perform clean-up after the backup is complete and is typically used for transaction log truncation.
The script accepts one argument. If the entire backup was successful and all the previous scripts were successful, it will be `1`. If any issues or failures happened, it will be `0`. If this argument is `1`, truncate logs.

Important: When you enable backup scripts, N2WS assumes you implemented all three scripts. Any missing script will be interpreted as an error and be reflected in the backup status. Sometimes the “complete” script is often not needed. In this case, write a script that just exits with the code `0`, and the policy will not experience errors.

6.3.4 Capturing the Output of Backup Scripts

You can have the output of backup scripts collected and saved in the N2WS Server. See section 4.2.4.

7 Linux/Unix Instances Backup

Making application-consistent backup of Linux instances does not require any agent installation. Since the N2WS server is Linux based, backup scripts will run on it. Typically, such a script would use SSH to connect to the backed-up instance and perform application quiescence. However, this can also be accomplished using custom client software.
There is no parallel to VSS in Linux, so the only method available is by running backup scripts.

7.1 Connecting to the N2WS Server

In order to create, test, and install backup scripts, you will need to connect to the N2WS server using SSH with `cpmuser`. The only way to authenticate `cpmuser` is by using the private key from the key pair you used when you launched the N2WS server instance. As long as your key is not compromised, no unauthorized person will be able to connect to the N2WS server.
With `cpmuser`, you will be able to copy (using secure copy), create, and test your scripts. `cpmuser` is the same user N2WS will use to run the scripts. If you need to edit your scripts on the N2WS Server, you can use the Vim or nano editors. Nano is simpler to use.

7.2 Backup scripts

Backup scripts should be placed in the path `/cpmdata/scripts`. If the policy belongs to a N2WS user other than the root user, scripts will be located in a subfolder named like the user (e.g. `/cpmdata/scripts/cpm_user1`). This path resides on the CPM data volume, and will remain there even if you terminate the N2WS server instance and wish to launch a new one. Backup scripts will remain on the data volume, together with all other configuration data. As `cpmuser`, you have read, write, and execute permissions in this folder.

- All scripts should exit with the code `0` when they succeed and `1` (or another non-zero code) when they fail.
- All scripts have a base name (detailed for each script in the coming sections) and may have any addition after the base name. The delimiter between the base part of the name and the file extension is a period (.). For example:
  `before_policy1.v11.5.bash`
  where '`before_policy1`' is the base name, '`v11.5`' is the optional additional part of the name, and '`bash`' is the file extension.
- Scripts can be written in any programming language: shell scripts, Perl, Python, or even binary executables.
- You only have to make sure the scripts can be executed and have the correct permissions.

Warning: Having more than one script with the same base name can result in unexpected behavior. N2WS does not guarantee which script it will run, and even to run the same script every backup.

There are three scripts for each policy:

- Before

- After
- Complete

7.2.1 Before Script

The `before_<policy name>[.optional_addition].<ext>` script runs before backup begins. Typically, this script is used to move applications to backup mode. The **before** script typically leaves the system in a frozen state for a short time until the snapshots of the policy are fired. One option is to issue a `freeze` command to a file system like `xfs`.

7.2.2 After Script

The `after_<policy name>[.optional_addition].<ext>` script runs after all the snapshots of the policy fire. It runs within a few seconds after the **before** script. This script releases anything that may have been frozen or locked by the **before** script. This script accepts the success status of the **before** script. If the **before** script succeeded, the argument will be `1`. If it failed, crashed, or timed out, the argument will be `0`.

Note:	This is the opposite of the exit status. Think of this as an argument that is true when the **before** script succeeds.

7.2.3 Complete Script

The `complete_<policy name>[.optional addition].<ext>` script runs after all snapshots are completed. Usually, it runs quickly since snapshots are incremental. This script can perform clean-up after the backup is complete and is typically used for transaction logs truncation. The script accepts one argument. If the entire backup was successful and all the previous scripts were successful, it will be 1. If any issues or failures happened, it will be `0`. If this argument is `1`, truncate logs.

7.2.4 Capturing the Output of Backup Scripts

You can have the output of backup scripts collected and saved in the N2WS Server, see section 4.2.4.

7.2.5 Troubleshooting and Debugging Backup Scripts

You can use the output collected by N2WS to debug backup scripts. However, the recommended way is to run them independently of N2WS, on the N2WS Server machine using SSH. You can then view their outputs and fix what is needed. Once the scripts work correctly, you can start using them with N2WS. Assuming these scripts are using SSH, during the first execution you will need to approve the SSH key by answering `yes` at the command line prompt. If you terminate your N2WS Server and start a new one, you will need to run the scripts again from the command line and approve the SSH key.

7.2.6 Example Backup Scripts

Following is an example of a set of backup scripts that use SSH to connect to another instance and freeze a MySQL Database:

- The **before** script will flush and freeze the database.
- The **after** script will release it.
- The **complete** script will truncate binary logs older than the backup.

Note: These scripts are presented as an example *without* warranties. Test and make sure scripts work in your environment as expected before using them in your production environment.

The scripts are written in `bash`:

before_MyPolicy.bash

```
#!/bin/bash
ssh -i /cpmdata/scripts/mysshkey.pem sshuser@ec2-host_name.compute-1.amazonaws.com "mysql -u root -p<MySQL root password> -e 'flush tables with read lock; flush logs;'"
if [ $? -gt 0 ]; then
  echo "Failed running mysql freeze" 1>&2
  exit 1
else
  echo "mysql freeze succeeded" 1>&2
fi
```

This script connects to another instance using SSH, and then runs a MySQL command. Another approach would be to use a MySQL client on the N2WS Server, and then the SSH connection will not be necessary.

After that script is executed, the N2WS server will start the snapshots, and then call the next script:

after_MyPolicy.bash

```
#!/bin/bash
if [ $1 -eq 0 ]; then
  echo "There was an issue running first script" 1>&2
fi
ssh -i /cpmdata/scripts/mysshkey.pem sshuser@ec2-host_name.compute-1.amazonaws.com "date +'%F %H:%M:%S' >  sql_backup_time; mysql -u root -p<MySQL root password> -e 'unlock tables;'"
if [ $? -gt 0 ]; then
  echo "Failed running mysql unfreeze" 1>&2
  exit 1
else
  echo "mysql unfreeze succeeded" 1>&2
fi
```

This script checks the status in the first argument and then does two things:

- First, it saves an exact timestamp of the of the current backup of the frozen database to a file,
- Then, it releases the lock on the MySQL table.

After that, when all snapshots succeed, N2WS runs the **complete** script:

complete_MyPolicy.bash

```
#!/bin/bash
if [ $1 -eq 1 ]; then
  cat /cpmdata/scripts/complete_sql_inner |ssh -i
/cpmdata/scripts/mysshkey.pem sshuser@ec2-host_name.compute-
1.amazonaws.com "cat > /tmp/complete_ssh; chmod 755 /tmp/complete_ssh;
/tmp/complete_ssh"
  if [ $? -gt 0 ]; then
    echo "Failed running mysql truncate logs" 1>&2
    exit 1
  else
    echo "mysql truncate logs succeeded" 1>&2
  fi
else
  echo "There was an issue during backup - not truncating logs" 1>&2
fi
```

It calls an inner script, `complete_sql_inner`:

```
butime=`<sql_backup_time`
mysql -u root -p<MySQL root password> -e 'PURGE BINARY LOGS BEFORE
"'"$butime"'"'
```

These two scripts purge the binary logs only if the **complete** script receives `1` as the argument. They purge logs earlier than the time in the timestamp files.

7.2.7 Scripts and SSH Access in a Multi-user Environment

If your N2WS configuration requires multiple users, which are separated from each other, you may wish to allow users to access N2WS using SSH to create and debug backup scripts:

- Create additional Linux users in the N2WS instance and allowing each user access to their own scripts folder only.
- `cpmuser` will need to be able to access and execute the scripts of all users. This can be achieved by assigning the user `cpmuser` as the group of all user subfolders and scripts. Then, if given **executable** permissions for the group, `cpmuser` will be able to access and execute all scripts.

8 Using Elastic File System (EFS) with N2WS

Configuring EFS on N2WS allows you to determine backup:

- Schedule and frequency
- Retention
- Lifecycle policy, including moving backups to cold storage, defining expiration options, and deleting them at end of life.

With AWS Backup, you pay only for the amount of backup storage you use and the amount of backup data you restore in the month. There is no minimum fee and there are no set-up charges.

Important: EFS Backup and Restore is performed by AWS Backup Service.

When adding an EFS target for the first time in a region, you must create the default backup vault in AWS. Go to the AWS Backup console and choose **Backup vaults**.

For more information regarding the AWS Backup Service, refer to https://docs.aws.amazon.com/efs/latest/ug/awsbackup.html

Notes: Before continuing, consider the following:

- Currently, AWS Backup service doesn't support cross-account DR for EFS resources.
- Check AWS for regions that are available for EFS backup on the AWS Backup service.
- AWS Backup is not available for EFS in the following regions: Asia Pacific (Hong Kong), Europe (Stockholm), South America (Sao Paulo), and Middle East (Bahrain).
- Backup transitions and expirations are performed automatically according to the configured lifecycle.
- A default or custom IAM role must exist in AWS to create and manage backups on behalf of N2WS. The IAM identity contains the backup and restore policies allowing operations on EFS. If a default was not automatically created, or you prefer to use a custom IAM role, see section 8.2.

8.1 Configuring EFS

1. In the AWS Console, create the EFS in one of the available regions listed in section 8.
2. In N2WS, in the **Backup Targets** tab of a Policy, select **Elastic File Systems** in the **Add Backup Targets** menu.
3. In the **Add Elastic File System** screen list, select one or more EFS targets and then select **Add Selected**.
4. In the **Backup Targets** tab, select an EFS target and then select **Configure**.
5. Configure the EFS backup and restore options and select **Apply**.

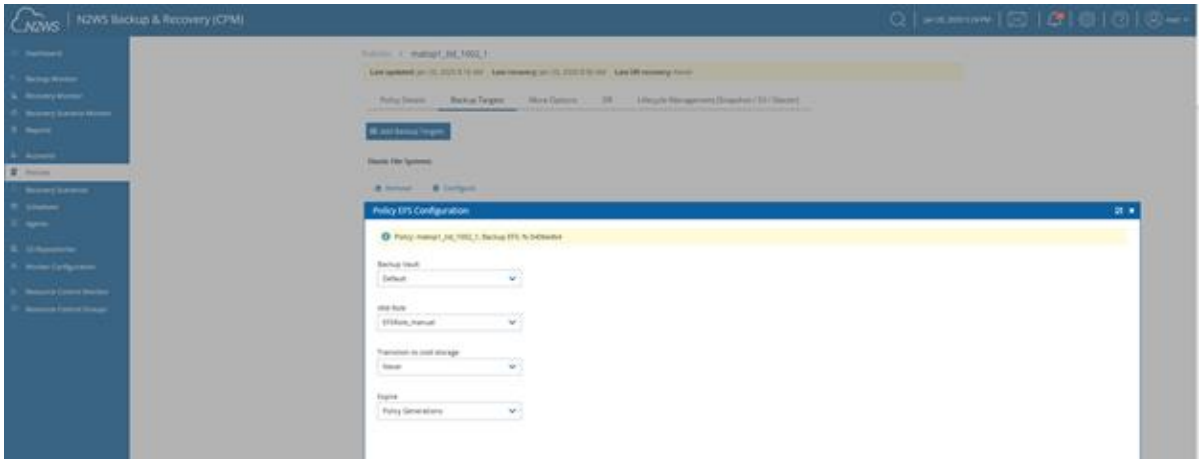

- **Backup Vault** – A logical backup container for your recovery points (your EFS snapshots) that allows you to organize your backups.

  Note: Default Backup vaults are created in AWS: **AWS Backup > Backup vaults**.

- **IAM Role** – An IAM identity that has specific permissions for all supported AWS backup services. The following AWS backup permissions should be attached to your IAM role:
  - **AWSBackupServiceRolePolicyForBackup** - Create backups on your behalf across AWS services.
  - **AWSBackupServiceRolePolicyForRestores** - Perform restores on your behalf across AWS services.

  If a default IAM role was not automatically created by AWS, or you require a custom IAM role, see section 8.2. Selecting the preferred IAM role is only required during the EFS policy configuration.
- **Transition to cold storage**– Select the transition lifecycle of a recovery point (your EFS snapshots). The default is **Never**.
- **Expire** – When does a protected resource expire. The default is **Policy Generations**.

  Note: Moving a backup to the Freezer will set **Expire** to **Never**.

6. When finished, select **Apply**.
7. Select **Save** in the Backup Targets screen to save to the configuration to the policy.

8.2 Creating IAM Roles in AWS

A default or custom IAM role is necessary for AWS to perform EFS operations on behalf of N2WS.

To create a default IAM Role:

1. Go to the AWS Backup Service:
   https://us-east-1.console.aws.amazon.com/backup/
2. Select **Create an on-demand backup**.
3. For **Resource type**, select **EBS**.
4. For **Volume ID**, select **any EBS volume to backup**.
5. Select **Default IAM Role**.
6. Select **Create on-demand backup**. Ignore the error provided by AWS.

7. Verify that the following role was created on AWS IAM Service:

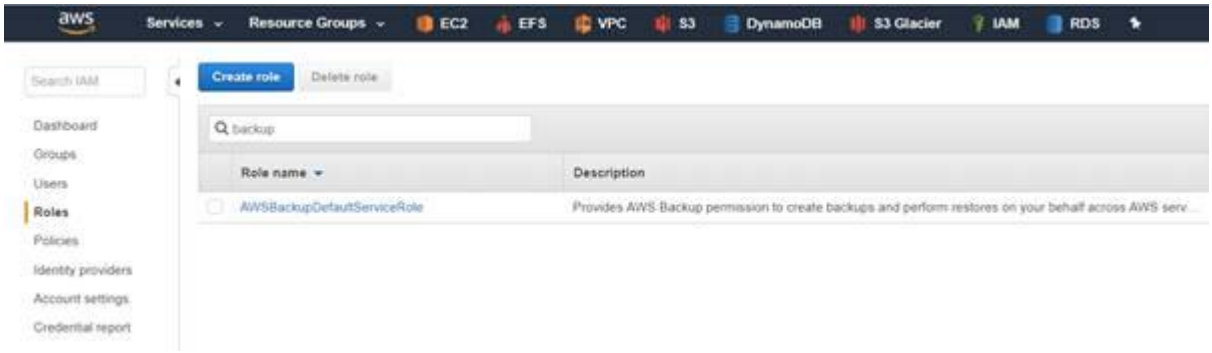

To create a custom IAM Role:

1. Go to AWS IAM Service:
   https://console.aws.amazon.com/iam/home#/roles
2. Select **Create role**.
3. Select **AWS Backup** and then select **Next: Permissions**.
4. Search for **BackupService**.
5. Select the following AWS managed policies:

   AWSBackupServiceRolePolicyForBackup

   AWSBackupServiceRolePolicyForRestores
6. Select **Next: Tags** and then select **Next: Review.**
7. Enter a **Role name** and select **Create role.**

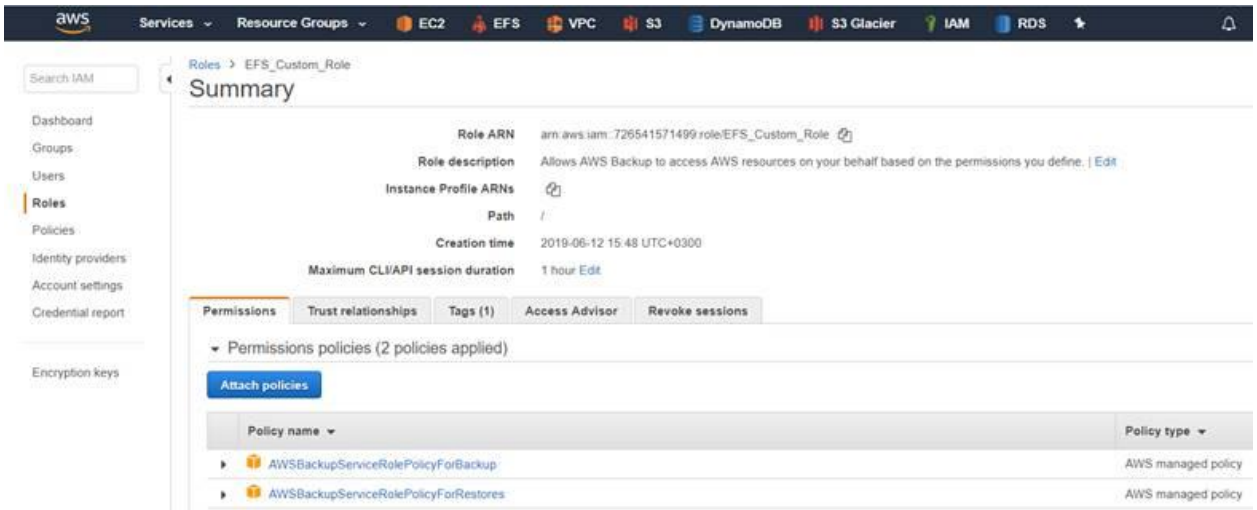

8.3 Backup Options for EFS Instances

EFS can be configured by creating the **cpm backup** tag with the following options. In this case, N2WS will override the EFS configuration with the tag values:

Key	Value
vault	Vault. Example: Default
role_arn	ARN of role. Example: arn:aws:iam::040885004714:role/service-role/AWSBackupDefaultServiceRole

Key	Value
cold_opt	Lifecycle transition: N – Never D – Days W – Weeks M – Months Y - Years
cold_opt_val	Integer for D, W, M, Y only
exp_opt	When does resource expire: P – Policy Generations N – Never D – Days W- Weeks M – Months Y - Years
exp_opt_val	Integer for D, W, M, Y only

Example:

```
cpm backup my_policy+vault=Default+exp_opt=D+exp_opt_val=1
```

CPM will backup EFS to the default vault, and set its expiration date to 1 day.

Note: The max length for the **cpm backup** value is limited to 256 characters.

9 Additional Backup Topics

9.1 N2WS in a VPC Environment

The N2WS Server runs in a VPC, except in old environments utilizing EC2 Classic. For N2WS to work correctly, it will need outbound connectivity to the Internet. To use AWS endpoints, see AWS Regions and Endpoints.

- You will need to provide such connectivity using one of the following methods:
  - Attaching an elastic IP,
  - Using a dynamic public IP, which is not recommended unless there is a dynamic DNS in place,
  - Enabling a NAT configuration, or
  - Using a proxy
- You will need to access it using HTTPS to manage it and possibly SSH as well, so some *inward* access will need to be enabled.
- If you will run Linux backup scripts on it, it will also need network access to the backed-up instances.
- If N2WS backup agents will need to connect, they will need access to it (HTTPS) as well.
- If backup scripts are enabled for a Linux backed-up instance, it will need to be able to get an *inbound* connection from the N2WS Server.
- If a Thin Backup Agent is used in a Windows backed-up instance, the agent will need *outbound* connectivity to the N2WS Server.

9.2 Backup when an Instance is Stopped

N2WS continues to back up instances even if they are stopped. This may have important implications:

- If the policy has backup scripts and they try to connect to the instance, they will fail, and the backup will have **Backup Partially Successful** status.
- If the policy has no backup scripts and VSS is not configured, or if the policy's options indicate that **Backup Partially Successful** is considered successful (section 4.2.2), backup can continue running, and automatic retention will delete older backups. Every new backup will be considered a valid backup generation.
- Snapshots will soon take no storage space since there will be no changes in the volumes, and EBS snapshots are incremental.
- Assuming the instance was shut down in an orderly manner and did not crash, backups will be consistent by definition.

Note: It is recommended that if you are aware of an instance that will be stopped for a while, you disable the policy by selecting its name and changing **status** to **disabled**.

Another way to proceed is to make sure the policy is not entirely successful when the instance is stopped by using backup scripts, and to keep the default stricter option that treats script failure as a policy failure. This will make sure that the older generations of the policy, before it was stopped, will not be deleted.

Important: If you disable a policy, you need to be aware that this policy will not perform backup until it is enabled again. If you disable it when an instance is stopped, make sure you enable it again when you need the backup to resume.

9.3 The Freezer

Backups belonging to a policy eventually get deleted. Every policy has its number of generations, and the retention management process automatically deletes older backups. To keep a backup indefinitely and make sure it is not deleted, move it to the Freezer. There can be several reasons to freeze a backup:

- An important backup of an instance you already recovered from so you will be able to recover the same instance again if needed.
- A backup of interest, such as the first backup after a major change in the system or after an important update.
- You want to delete a policy and only keep one or two backups for future needs.
- Elements in the freezer will not be deleted by the automatic **Cleanup** process.

To move a backup to the Freezer:

Note: Once a backup is moved to the freezer, you will *not* be able to move it back.

1. In the left panel, select the **Backup Monitor** tab.
2. Select the backup and then select **Move to Freezer**.
3. Type a unique name and an optional description for identification and as keywords for searching and filtering later.

After a backup is in the Freezer:

- Frozen backups are identified by the frozen icon in the **Lifecycle Status** column of the **Backup Monitor** tab.
- It will only be deleted if you do so explicitly. Use **Delete Frozen Item**.
- It will still remain even if you delete the whole policy, frozen backups from the policy will still remain.
- It is recovered the same way as from a regular backup.
- You can search and filter frozen backups using as keywords the name or description. To change the name or description, select **Edit Frozen Item**.

While in the **Backup Monitor,** you can switch between showing backup records *in* the Freezer and backup records *not* in the Freezer. In the **Show** area on the far right of the filters line:

- To show only the backup records that were moved to the Freezer, select **Backup** once. Select **Backup** again to show all backup records.
- To show only the backup records that were *not* moved to the Freezer, select **Freezer** once. Select **Freezer** again to show all backup records.

9.4 Running Automatic Cleanup

Automatic Cleanup allows you to manage the frequency of the cleanup process and the:

- Number of days to keep backup record, even if the backup is deleted.
- Number of days after which to rotate single AMIs.

Note:	Keeping backups for long periods of time can cause the N2WS database to grow and therefore affect the size you need to allocate for the CPM data volume. N2W Software estimates that every GiB will accommodate the backup of 10 instances. N2W Software estimates that 10 instances are correct when every record is kept for around 30 days. If you want to keep records for 90 days, triple the estimate, i.e. for 10 instances make the estimate 3 GiB, for 20 make the estimate 6 GiB, etc.

To manage the number of generations saved:

1. In the toolbar, select  **Server Settings**.
2. In the **General Settings** tab, select **Cleanup**.
3. In the **Cleanup Interval** list, select the number of hours between cleanup runs. Select **Cleanup Now** to start a cleanup immediately.
4. In each list, select the number of days to:
   - Rotate Single AMIs
   - Keep Deleted Records
   - Keep User Audit logs
   - Keep Resource Control Records

Note:	The number of days is counted since the backup was created and not deleted. If you want to make sure every backup record is saved for 90 days after creation, even if it was already deleted, select 90.

If **Explore** sessions are running (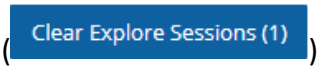

), you can select **Clear Explore Sessions** to terminate all sessions.
The S3 Cleanup runs independently according to the retention period configured for the policy in the backup copy settings. See section 21.1.

9.5 Backing up Independent Volumes

Backing up independent volumes in a policy is performed regardless of the volumes attachment state. A volume can be attached to any instance or not attached at all, and the policy will still back it up. Backup scripts can determine which instance is the active node of a cluster and perform application quiescence through it.

9.6 Excluding Volumes from Backup

Note:	If you enable the **Exclude volumes** option in the **Tag Scan** tab of the **General Settings:**

- The **Exclude volumes** option overrides the exclusion of volumes performed through the UI.
- Tagged instances are not included in the **Exclude volumes** option and are excluded from backup *only* when tagged with **'#exclude'** for the policy.

Following are the ways to exclude volumes from backup:

- Enabling the **Exclude volumes** option in **General Settings**:
  - In the toolbar, select **Server Settings** > **General Settings**.
  - In the **Tag Scan** tab, select **Exclude volumes** and then select **Scan Now**.

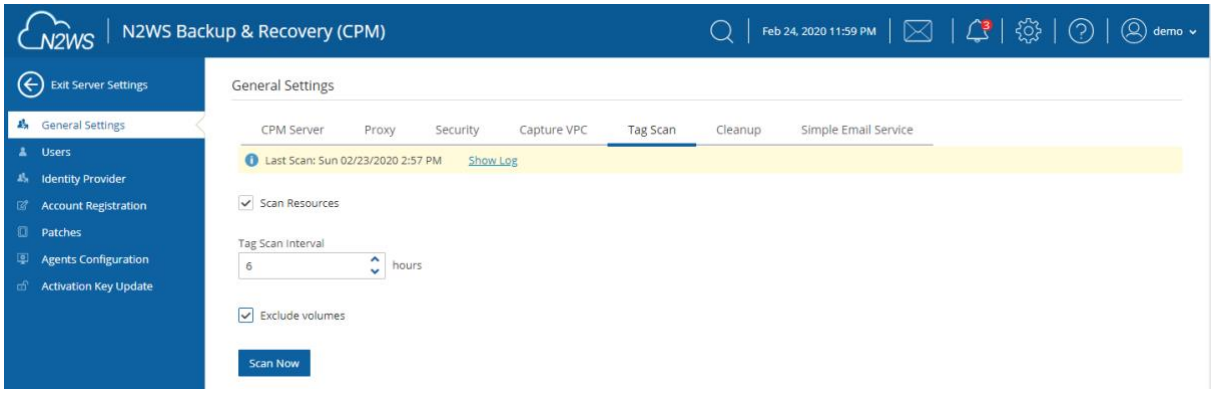

- Excluding a volume from a policy configuration in the UI. See section 4.2.3
- Disabling a scheduled backup time. See section 4.1.4.
- Using an '#exclude' tag for the policy. See section 14.1.6.

9.7 Regions Disabled by Default

In order to perform certain actions on Asia Pacific (Hong Kong) and Middle East (Bahrain) AWS regions, managing Session Token Services (STS) is required, as Session Tokens from the global endpoint (https://sts.amazonaws.com) are only valid in AWS Regions that are enabled by default.
For AWS Regions not enabled by default, users have to configure their AWS Account settings.

To configure AWS Account settings to enable Session Tokens for all regions:

5. Go to your AWS console and sign in: https://console.aws.amazon.com/iam
6. In the navigation pane, select Account settings.
7. In the 'Security Token Service (STS)' section, select Change Global endpoint.
8. In the Change region compatibility of session tokens for global endpoint dialog box, select Valid in all AWS Regions.

Note: Session tokens that are valid in all AWS regions are larger. If you store session tokens, these larger tokens might affect your system.

For more information on how to manage your STS, see https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_temp_enable-regions.html

10 Performing Recovery

N2WS offers several options for data recovery. Since all N2WS backup is based on AWS's snapshot technology, N2WS can offer rapid recovery of instances, volumes, and databases. In the **Backup Monitor**, when you select **Recover** for a certain backup, you are directed to the **Backup Monitor Recover** screen. You can Search by Resource using the resource ID or name.

Note: You can also use the toolbar search box to quickly find backups for recovery. Enter the target resource's AWS ID or name in the search box and select the search icon.

For backups with multiple resource types as targets, the Recover screen will have a separate tab for each type. Select a backup. The **Recover** screen opens.

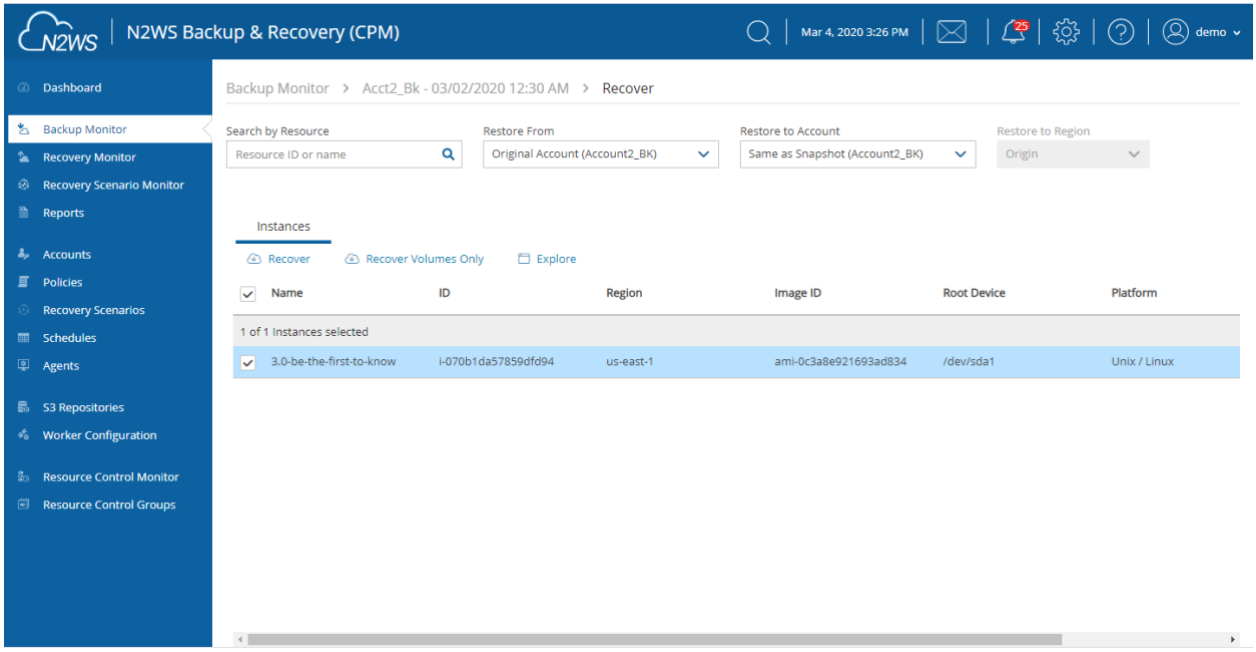

Depending on the specifics of the backup, the recover screen includes:

- A search box for locating a resource by ID or name.
- Tabs for recovering the backed-up instances, independent volumes, databases, etc.
- Outputs of any backup scripts and VSS if it exists. These reference outputs may be important during a recovery operation.
- If this backup includes DR to another region, there will be a **Restore to Region** drop-down menu to choose in which region to perform the recovery.
- If you have cross-account functionality enabled for your N2WS license, there are two other drop-down menus:

- **Restore to Account** list where you can choose to restore the resources to another account.
- If you defined cross-account DR for this policy, you will have the **Restore from Account** list for choosing from which account to perform recovery.

Note: All the choices about regions and accounts you make in the recover screen apply to all the recovery operations that you initiate from this screen.

Choose the backups to recover and then select the **Recover** resource type button.

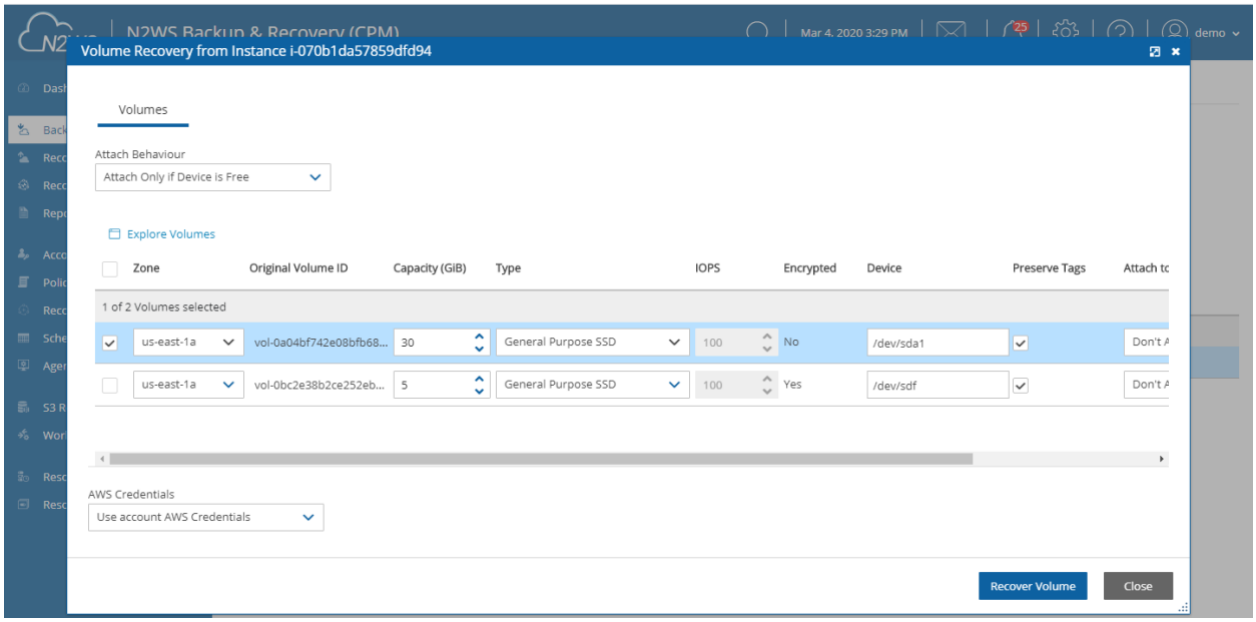

Recommendation: N2W Software strongly recommends that you perform recovery drills occasionally to make sure your recovery scenarios work. It is not recommended to try it for the first time when your servers are down. Each policy on the policy screen shows the last time recovery was performed on it. Use the last recovery time data to track recovery drills.

10.1 Recovery AWS credentials

All recovery screens have a drop-down list at the bottom labelled **AWS Credentials.** By default, the account AWS credentials used for backup will be used for recovery operations also. Depending on the backup, you can select **Provide Alternate AWS Credentials** and fill in different credentials for recovery. This can be useful if you want to use IAM-created backup credentials that do not have permissions for recovery. See section 16.3. When using custom credentials, N2WS verifies that these credentials actually belong to the recovery account.

To use custom credentials:

1. Select **Provide Alternate AWS Credentials** in the list. The custom credential boxes appear.
2. In the **AWS Access Key** box, enter your access key.
3. In the **AWS Secret Key** box, enter your secret key.

10.2 Instance Recovery

With Instance recovery, you can recover a complete instance with its data for purposes, such as:

- An instance crashed or is corrupted and you need to create a new one
- Creating an instance in a different AZ
- Creating an instance in a different region. See section 11.5.1.
- Creating an instance from a frozen image

When you recover an instance, by default, you recover it with its configuration, tags, and data, as they were at the time of the backup. However, you can change these elements:

- Instance type
- Placement
- Architecture
- User data, etc.

You can also choose how to recover the system itself:

- For Linux EBS-based instances: if you have a snapshot of the boot device, you will, by default, use this snapshot to create the boot device of the new instance. You can, however, choose to create the new instance from its original image or a different one.
- For instance-store-based: you will only have the image option. This means you cannot use the snapshot of the instance's root device to launch a new instance.
- For EBS-based Windows Servers: there is a limitation in AWS, prohibiting launching a new instance from a snapshot, as opposed to from an AMI.
  N2WS knows how to overcome this limitation. You can recover an instance from a snapshot, but you also need an AMI for the recovery process. By default, N2WS will create an initial AMI for any Windows instance it backs up and use that AMI for the recovery process. Usually, you do not need to change anything to recover a Windows instance.
- Your data EBS volumes will be recovered by default to create a similar instance as the source. However, you can choose:

- To recover some or none of the volumes.
- To enlarge volume capacity, change their device name, or IOPS value.
- To preserve tags related to the instance and/or data volumes, or not.

The instance recovery screen has tabs for **Basic Options, Volumes,** and **Advanced Options**. At the bottom of each screen, there is an option to change **AWS Credentials**.

10.2.1 Basic Options

The **Basic Options** tab is divided into the general section and the Networking section:

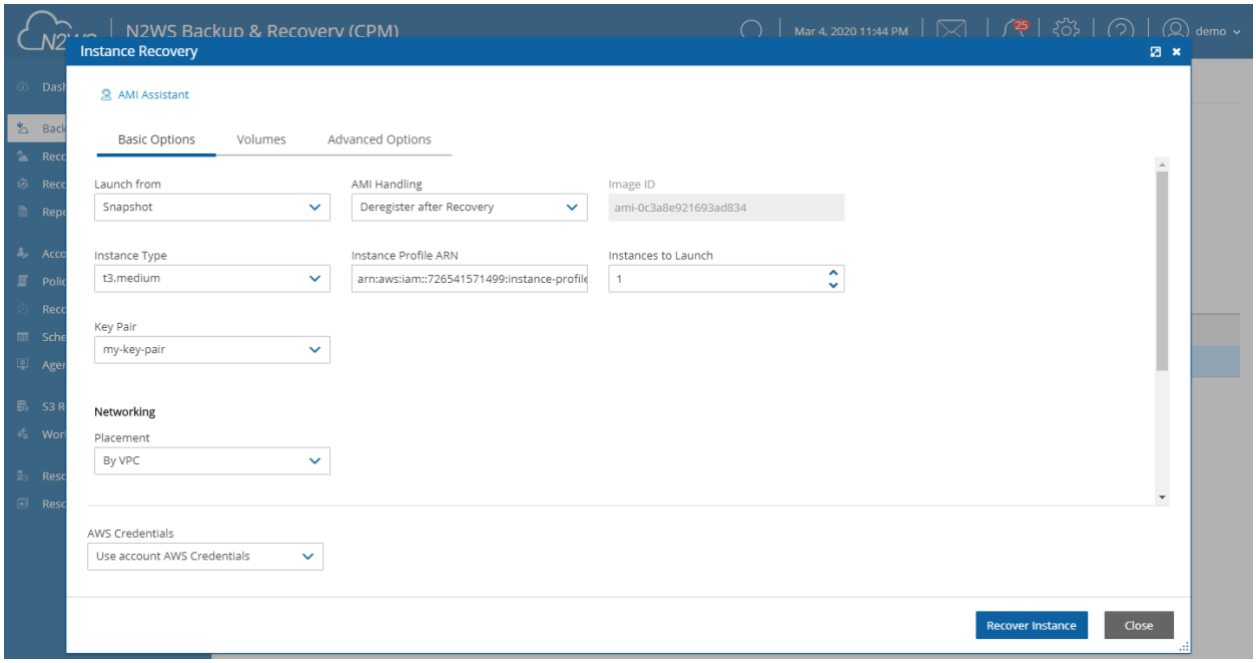

- **AMI Assistant** – Select to view the details of the AMI used to launch your instance and find similar AMIs.
- **Launch From** – Whether to launch the boot device (**Image**) from an existing image or a snapshot. The **Snapshot** option is available only if this is an EBS-based instance, and a snapshot of the boot device is available in this backup.
- **AMI Handling** – This option is relevant only if **Launch From** is set to **snapshot**.
- **Image ID** – This is only relevant if **Launch From** is set to **image** or if you are recovering a Windows instance. By default, this will contain the initial AMI that N2WS created, or if it does not exist, the original AMI ID from which the backed-up instance was launched. You can type or paste a different AMI ID here, but you cannot search AMIs from within N2WS. You can search for it with the AWS Management Console.
- If this instance is launched from a snapshot, a new AMI image will be registered and defined as follows:
  - **De-Register after Recovery** – This is the default. The image will only be used for this recovery operation and will be automatically de-registered at the end. This option will not leave any images behind after the recovery is complete.
  - **Leave Registered after Recovery** – The new created image will be left after recovery. This option is useful if you want to hold on to this image to create future instances. The snapshots the image is based on will not be deleted by the automatic retention process. However, if you want to keep this image and use it in the future, move the whole backup to the Freezer. See section 9.3.
  - **Create AMI without Recovery** – This option creates and keeps the image but does not launch an instance from it. This is useful if you want to launch the instance/s from outside N2WS. If you wish to keep using this image, move the backup to the Freezer.
- **Instance Type** – Choose the instance type of the new instance/s. The instance type of the backed-up instance is the default.

- If you choose an instance type that is incompatible with the image or placement method, the recovery operation will fail.

Note: Since not all instance types are available in all AWS regions, recovery of an unsupported instance type in a certain region might fail. Where the instance type is not yet supported by AWS, we recommend that you configure a supported instance type.

- **Instance Profile ARN** – The ARN of the instance role (IAM Role) for the instance. To find the ARN, select the Role name in IAM Management Console and then select the **Summary** tab. The default will be the instance role of the backed-up instance if it had one.
- **Instances to Launch** – Specifies how many instances to launch from the image. The default is one, which is the sensible choice for production servers. However, in a clustered environment you may want to launch more than one. It is not guaranteed that all the requested instances will launch. Check the message at the end of the recovery operation to see how many instances were launched, and their IDs.
- **Key Pair** – The key pair you want to launch the instance with. The default is the key that the backed-up instance was created with. You can choose a different one from the list. Keys are typically needed to connect to the instance using SSH (Linux).

Note: Keys cannot be used to decrypt the Windows password of a restored instance.

10.2.1.1 Networking Section

The main purpose of the **Networking** section is to define what will be the placement of the instance. By default, it will be the same placement as the backed-up instance. An instance can be placed using three methods which are not all necessarily available.

- **By VPC** – Default placement if you have VPC subnets defined in your account.
- **By Availability Zone** – This is the most basic type and the only one which is always available. You can choose in which AZ to launch the instance. Additional options are:
  - You can choose a different AZ from the backed-up instance.
  - By default, if the backed-up instance was not in a VPC, it will have the same zone as the backed-up instance. Choose a different AZ from the list.
- **By Placement Group** – If you have placement groups defined, this option is available. This is an instance type that can be placed in a placement group. See AWS documentation for details.

If you chose **By VPC** in **Placement**, the following fields are available:

- **VPC** –You can choose the VPC the instance is to be recovered to. By default, it will contain the VPC of the original instance.
- **Clone VPC** - Option to recover to a clone of the selected VPC environment. Control switches to the Account's Clone VPC screen. Choose the date of the source VPC capture for the clone and an optional new destination name. See section 10.2.5. After the cloning process is completed, the name of the newly cloned VPC will appear in the VPC box.
- **VPC Subnet** – This will hold all the subnets in the currently selected VPC.

- **Security Group** – Choose security groups to be applied with the new instance. This is a multiple-choice field. By default, the security groups of the backed-up instance will be chosen.

  Note: Security groups for VPC instances are different than groups of non-VPC instances. This field has a filter to help you find the security group that you need.

- **VPC Assign IP** – If the backed-up instance was in a VPC subnet, the default value will be the IP assigned to the original instance.
  - If the assigned IP is still taken, it can fail the recovery operation. You can type a different IP here. When you begin recovery, N2WS will verify the IP belongs to the chosen subnet.
  - If this field is empty, an IP address from the subnet will be automatically allocated for the new instance.

If you chose **By Availability Zone** in **Placement**:

- **Availability Zone** - By default, if the backed-up instance was not in a VPC, it will have the same zone as the backed-up instance. However, you can choose a different one from the list.
- **Security Group** - Choose security groups to be applied with the new instance. This is a multiple-choice field. By default, the security groups of the backed-up instance will be chosen.

If you chose **By Placement Group**:

- **Placement Group** - Choose the placement group from the list.

For all Placement options, the following boxes are also available:

- **Additional NICs** - If you want to add additional NICs.
- **AWS Credentials** - You can choose to use different AWS credentials for the recovery operation.

10.2.2 Volumes

Select the **Volumes** tab to choose which volumes to recover and how.

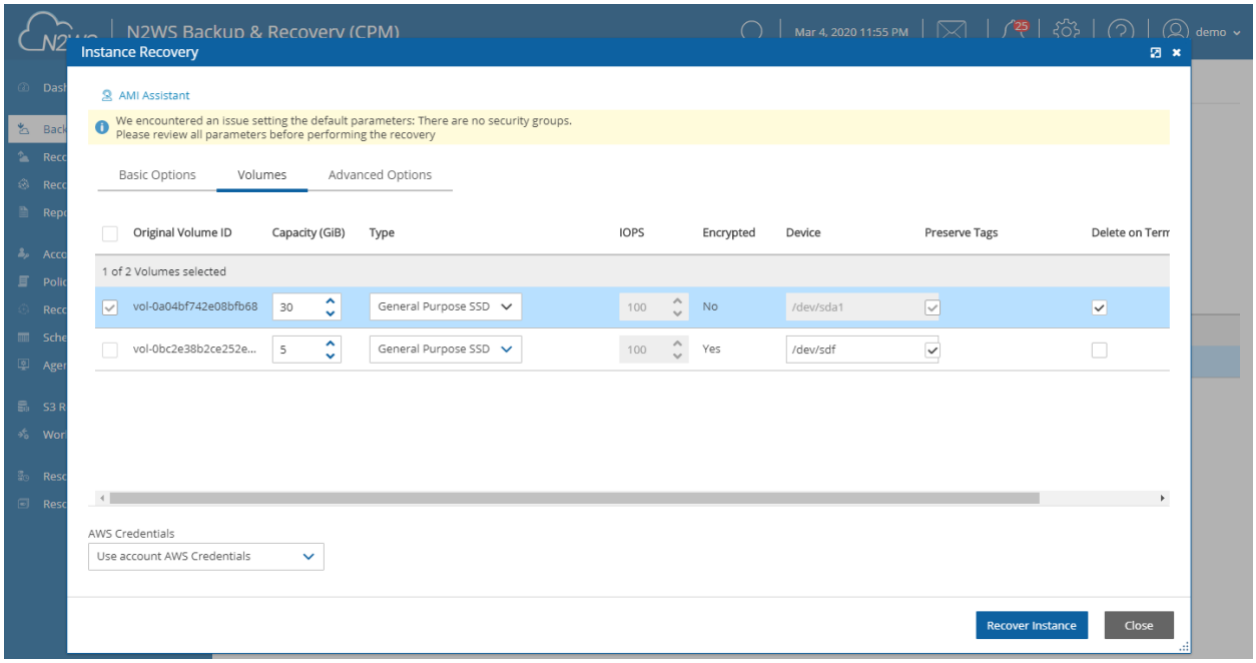

All data volumes in the policy except the boot device are listed here. Their default configuration is the same as it was in the backed-up instance at the time of the backup.
Select a volume to include it in the recovery. You can make adjustments to the volumes, as follows:

- Enlarge capacity of the volume.
- Change the device and device type.
- Change IOPS.
- Exclude any tags associated with the volume, such as its name.
- By default, tags associated with the volume, such as name, are preserved. Clear **Preserve Tags** to exclude all tags.
- By default, the volumes are not deleted on termination of instances recovered from a snapshot. Select **Delete on Termination** For instances recovered from a snapshot, delete the volume on termination of the instance ().

10.2.3 Advanced Options

Note:	It is possible to recover to a different account and region by recovering to a clone of an original VPC environment. See the **Clone VPC** option below.

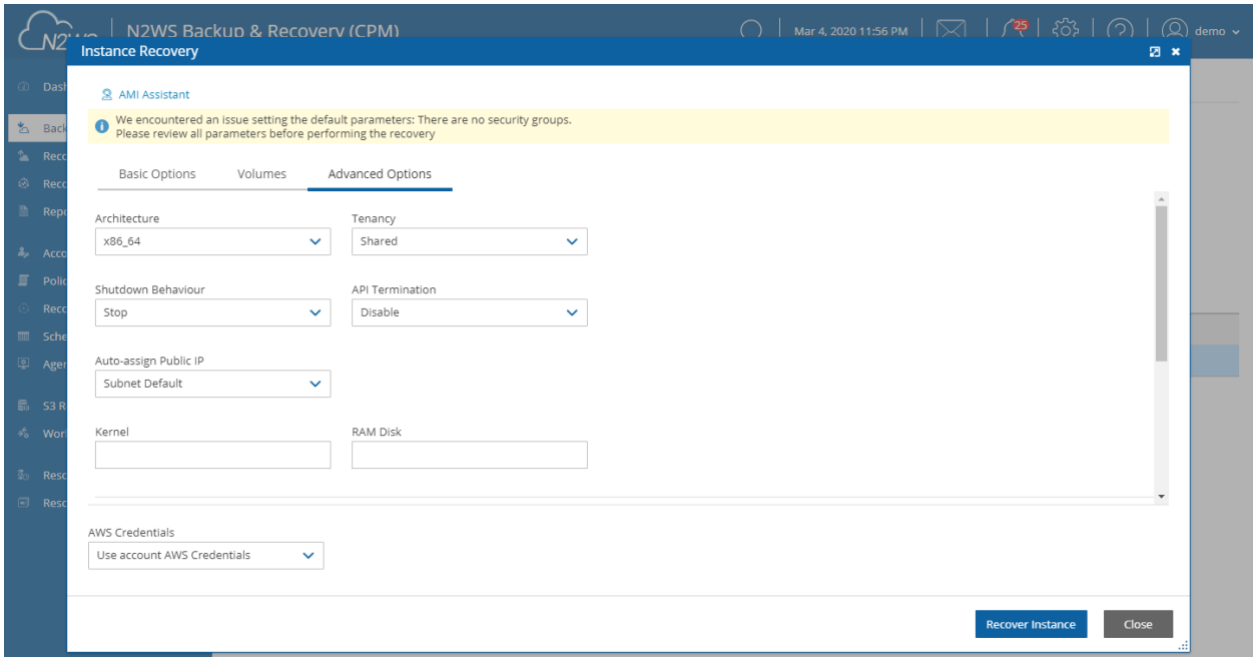

Advanced options include the following:

- **Architecture** – The default will be the architecture of the backed-up instance. Options are:
  - i386 – which is X86 – 32-bit
  - x86_64 – which is X86 – 64-bit

  Note: Changing the architecture may result in an error if the image is incompatible with the requested architecture. For example, if your image is a native 64-bit image and you choose **i386**, the recovery operation will fail.
- **Tenancy** – Choose the tenancy option for this instance.
- **Shutdown Behaviour** – The value of the original instance is the default. If the recovered instance is instance-store-based, this option is not used. The choices are:
  - **stop** – If the instance is shut down, it will not be terminated and will just move to **stopped** state.
  - **terminate** – If the instance is shut down, it will also be terminated.
- **API Termination** – Whether terminating the new instance by API is enabled or not. The backed-up instance value is the default.
- **Auto-assign Public IP** - Whether to assign a public IP to the new instance. This is for public subnets. By default, it will behave as the subnet defines.
- **Kernel** – Will hold the Kernel ID of the backed-up instance. You can type or paste a different one. However, you cannot search for a kernel ID from within N2WS. Change this option only if you know exactly which kernel you need. Choosing the wrong one will result in a failure.
- **RAM Disk** - Will hold the RAM Disk ID of the backed-up instance. You can type or paste a different one. However, you cannot search for a RAM Disk ID from within N2WS. Change this option only if you know exactly which RAM Disk you need. Choosing the wrong one will result in a failure.
- **Preserve Tags** – Whether to associate the same tags, such as the volume name, to the recovered volume. The default is yes.

- **Allow Monitoring** – Select if monitoring should be allowed for the new instance. The value in the backed-up instance is the default.
- **ENA** – Select to support Extended Network Adaptor.
- **EBS Optimized** –Select to launch an EBS Optimized instance. The value from the backed-up instance is the default.
- **Enable User Data** – Whether to use user data for this instance launch. If selected, the **User Data** box opens. Enter the text. The text of the user data. Special encoding or using a file as the source is not currently supported from within N2WS.

Advanced Options include different additional choices depending on whether **Placement** is **By VPC**, **By Availability Zone** or **By Placement Group**.

To complete the recovery operation, select **Recover Instance** and then confirm. If there are errors that N2WS detects in your choices, you will return to the recover instance screen with error messages. Otherwise, you will be redirected back to the recovery panel screen, and a message will be displayed regarding the success or failure of the operation.

10.2.4 AMI Assistant

The AMI Assistant is a feature that lets you view the details of the AMI used to launch your instance, as well as find similar AMIs. N2WS will record the details of the AMI when you start backing up the instance. If the AMI is deleted sometime after the instance started backing up, N2WS will remember the details of the original AMI.

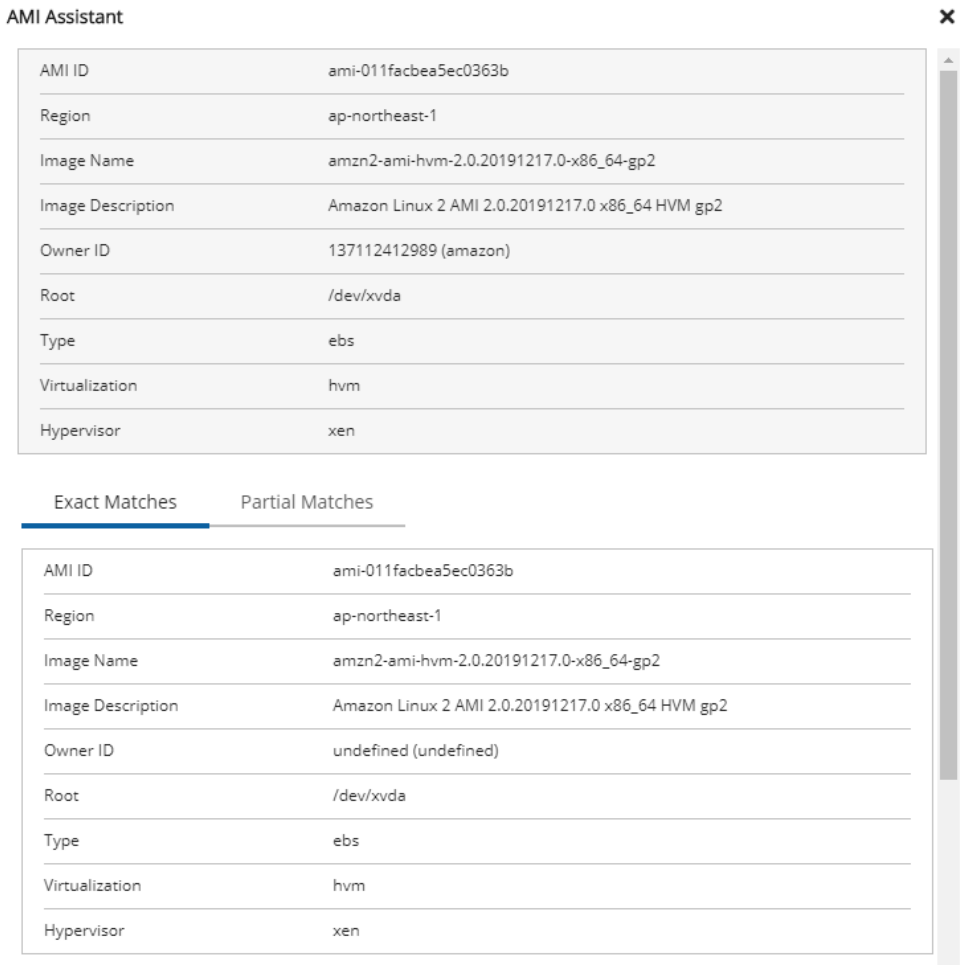

After selecting **AMI Assistant** in the instance recovery screen, you will see these details:

- AMI ID
- Region
- Image Name
- Image Description
- Owner ID
- Root – Device
- Type
- Virtualization
- Hypervisor

To find AMIs with properties that are exactly like the original, select the **Exact Matches** tab.
If the **Exact Matches** search does not find matches, select the **Partial Matches** tab which will search for AMIs similar to the original.
AMI Assistant searches can be useful in the following scenarios:

- You want to recover an instance by launching it from an image, but the original AMI is no longer available.
- You want to recover an instance by launching it from an image, but you want to find a newer version of the image. The fuzzy search will help you.
- You are using DR (section 11) and you need to recover the instance in a different region. You may want to find the matching AMI in the target region to use it to launch the

instance, or you may need its kernel ID or ram disk ID to launch the instance from a snapshot.

10.2.5 Recovering to a Cloned VPC

When you select **Clone VPC** in the **Basic Options** tab, the **Clone VPC** screen opens.

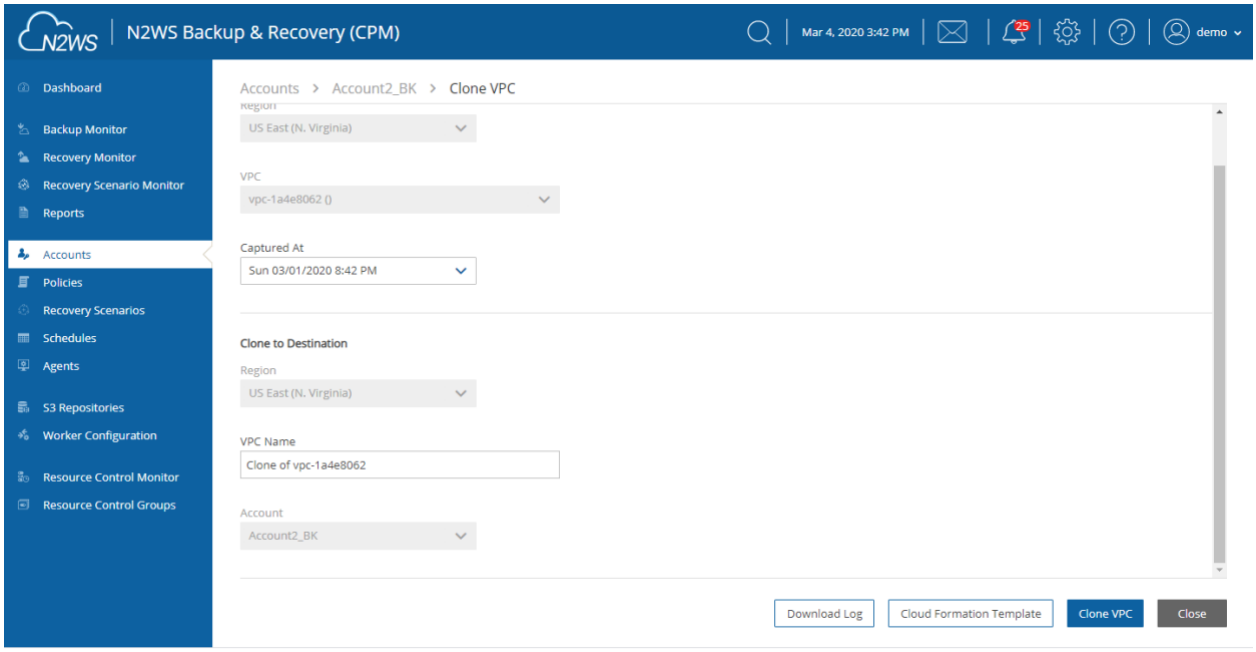

N2WS will have pre-set the following fields according to the selections made in the **Advanced Options** section:

- **Capture Source**:
  - Region and VPC
  - Captured at date and time – You can select a different date and time to clone in the drop-down list of captures.
- **Clone to Destination**:
  - Region and Account
  - VPC Name – You can change the suggested name for the new VPC.

Note: As part of the cloning process, N2WS uses CloudFormation. If the CloudFormation template is over 50 kB, the **Cloud Formation Template** button appears: It requires an existing S3 bucket for uploading. See section 23.5.

When finished, select **Clone VPC**.
If you changed the suggested **VPC Name**, it will appear in the **VPC** box.
To view the results of the clone VPC action, select the **Download Log** button.

10.3 Volume Recovery

Volume recovery means creating EBS volumes out of snapshots. In N2WS, you can recover volumes that were part of an instance's backup or recover EBS volumes that were added to a policy as an independent volume. The recovery process is basically the same.

To recover volumes belonging to an instance:

1. In the left panel, select the **Backup Monitor**.
2. In the **Backup Monitor** tab, select a backup and then select **Recover**.

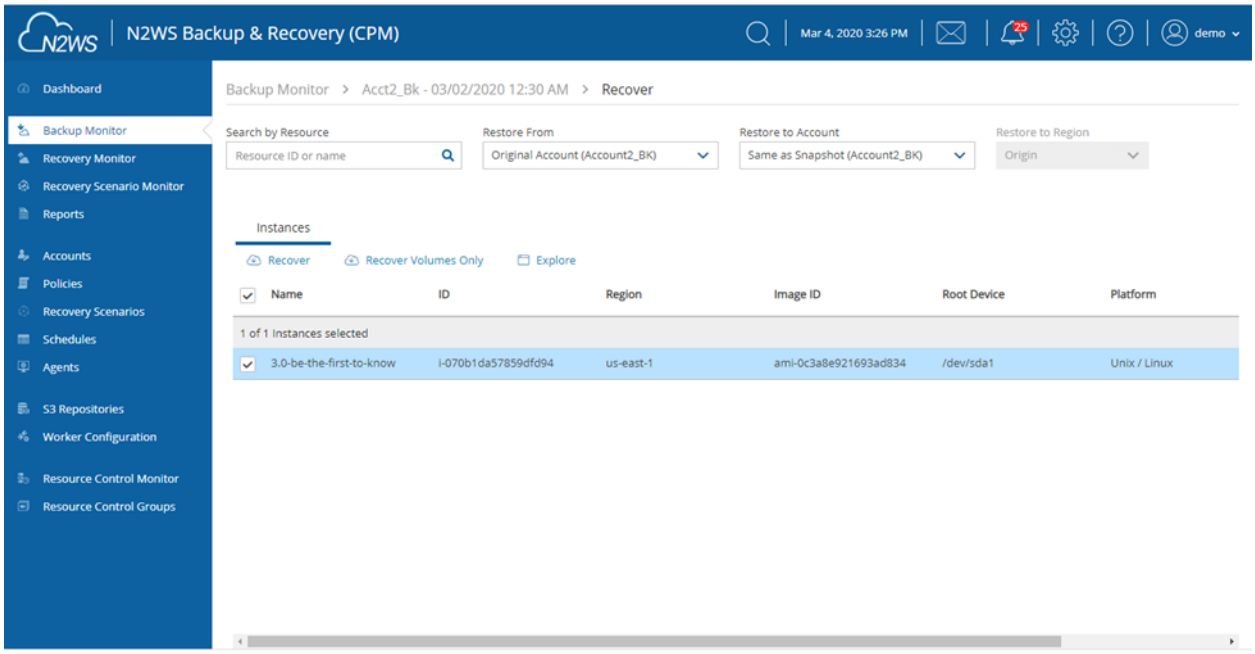

3. Select an instance and then select **Recover Volumes Only**.

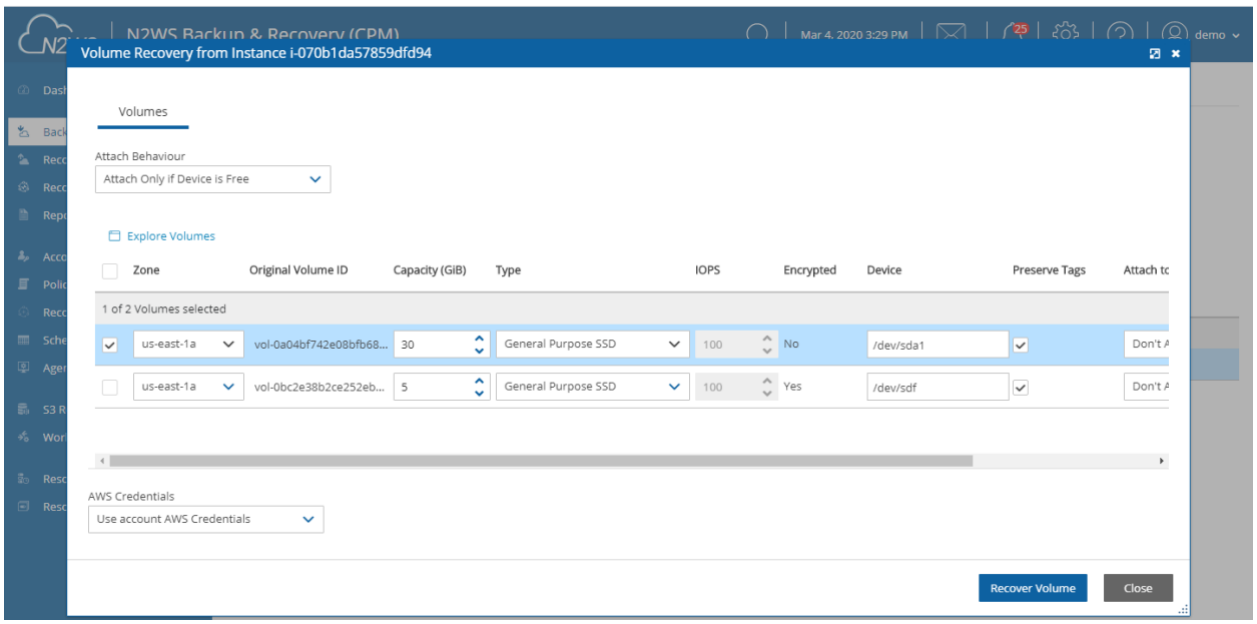

4. Change the fields as needed:
   - **Attach Behaviour** – This applies to all the volumes you are recovering, if you choose to attach them to an instance:

- Attach **only if Device is Free** – If the requested device is already taken in the target instance, the attach operation will fail. You will get a message saying the new volume was created but was not attached.
- **Switch Attached Volumes** – This option will work only if the target instance is in **stopped** state. If the instance is running, you will get an error message. N2WS will not try to forcefully detach volumes from a running instance, since this can cause systems to crash.
- **Switch Attached Volumes and Delete Old Ones** – This option will work only on stopped instances. This option will also delete the old volumes that are detached from the instance.

Important: If you choose **Switch Attached Volumes and Delete Old Ones**, make sure you do not need the old volumes. N2WS will delete them after detaching them from the target instance.

- **Recover** – Enabled by default. Clear **Recover** if you do not want that volume recovered.
- **Zone** – AZ. The default is the original zone of the backed-up volume.
- **Original Volume ID** – ID of original volume.
- **Capacity** – Enlarge the capacity of a volume. You cannot make it smaller than the size of the original volume, which is the default.
- **Type** – Type of the EBS volume.
- **IOPS** – Number of IOPS. This field is used only if the type of volume you chose is **Provisioned IOPS SSD**. The default will be the setting from the original volume. Values for IOPS should be at least 100, and the volume size needs to be at least 1/10 that number in GiBs. For example, if you want to create a 100 IOPS volume, its size needs to be at least 10 GiB. If you will not abide to this rule, the recovery operation will fail.
- **Encrypted** – Whether device is encrypted.
- **Device** – Which device it will be attached as. This is only used if you choose to automatically attach the recovered volume to an instance. If the device is not free or not correct, the attach operation will fail.
- **Preserve Tags** – Whether to associate the same tags, such as the volume name, to the recovered volume. The default is yes.
- **Attach to Instance** – Whether to attach the newly recovered volume to an instance. Start typing in the list to initiate a filter. The list holds instances that are in the same AZ as the volume. Changing the **Zone** will refresh the content of this list.
- **AWS Credentials** - As with other recovery screens, you can choose to use different AWS credentials for the recovery operation.

5. After selecting **Recover Volumes** and confirming, if there was a logical error in a field that N2WS detected, you will be returned to the screen with an error notification.
6. To follow the progress of the recovery, select the **Open Recovery Monitor** link in the 'Recovery started' message 

at the top right corner, or select the **Recovery Monitor** tab.

To recover independent volumes:

1. Select a backup and then select the **Independent Volumes** tab above the table.

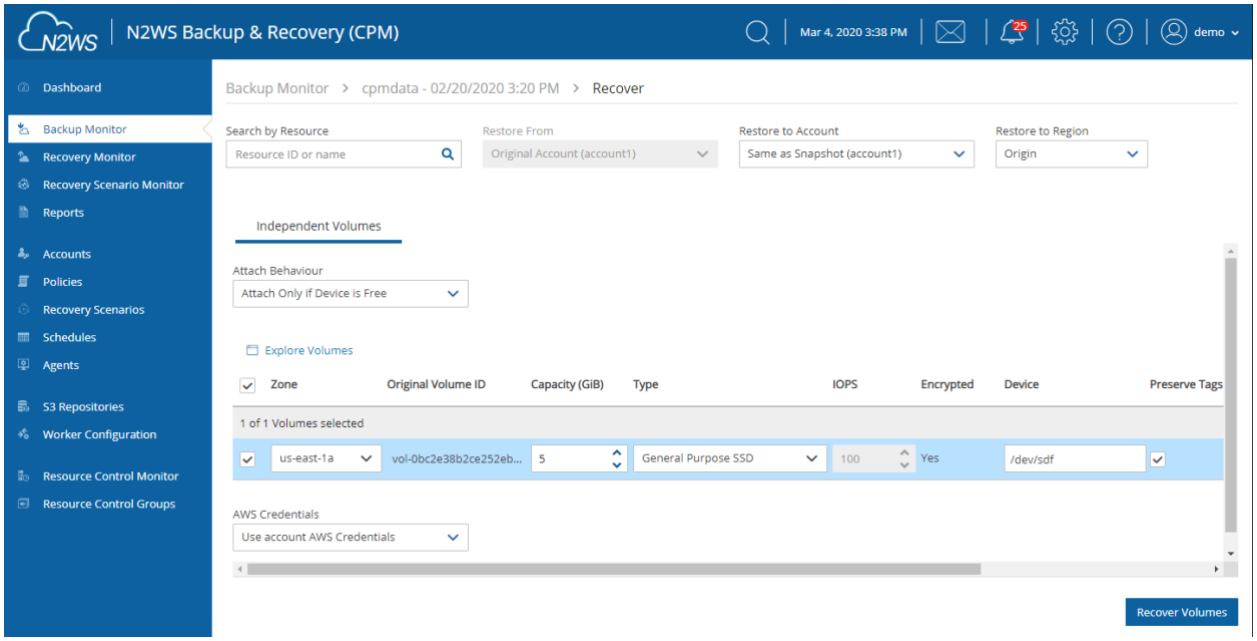

2. A screen similar to recover instance volumes opens. See above.

10.4 RDS Database Recovery

When a backup includes snapshots of RDS databases, selecting **Recover** to bring you to the RDS Databases tab. You will see a list of all RDS databases in the current backup. You can change the following options:

- **Recover** – Clear **Recover** to *not* recover the current database.
- **Zone** – The AZ of the database. By default, it will be the zone of the backed-up database, but this can be changed. Currently, recovering a database into a VPC subnet is not supported by N2WS. You can recover from the snapshot using AWS Management Console.
- **DB Instance ID** – The default is the ID of the original database. If the original database still exists, the recovery operation will fail. To recover a new database, type a new ID.
- **DB Instance Class** – The default is the original class, but you can choose another.
- **Storage Type** – Type of storage.
- **IOPS** - Number of IOPS. This field is used only if the type of volume you chose is **Provisioned IOPS SSD**. The default will be the setting from the original volume. Values for IOPS should be at least 100, and the volume size needs to be at least 1/10 that number in GiBs.
- **Port** –The default is the port of the original backed-up database, but you can choose another.
- **Multi AZ** – Whether to launch the database in a multi AZ configuration or not. The default is the value from the original backed-up database.
- **Subnet Group** – Whether to launch the database in a VPC subnet or not, and to which subnet group. The default will be the value from the original backed-up database. You can recover a database from outside a VPC to a VPC subnet group, but the other way around is not supported and will return an error.

- **Publicly Access**. – Whether the database will be publicly accessible or not. The default is the access from the original backed-up database.
- **AWS Credentials** - As in other types of recovery, you can choose to use different AWS credentials and enter your keys.

10.5 Aurora Cluster Recovery

Aurora recovery is similar to RDS recovery, with a few important differences.

- Aurora introduces the concept of clusters to RDS. You no longer launch and manage a DB instance, but rather a DB cluster that contains DB instances.
- An Aurora cluster may be created in a single AZ deployment, and the cluster will contain one instance.
- Or, as in production deployments, the cluster will be created in a multi-AZ deployment, and the cluster will have reader and writer DB instances.
- When recovering an Aurora cluster, N2WS will recover the DB cluster and then will create the DB instances for it.

After selecting a backup with Aurora Clusters, select **Recover**. The **Aurora Clusters Recover** screen opens. In this screen all Aurora clusters that were backed up are listed. You can change the following options:

- **Recover** – Clear to not recover the current Aurora cluster.
- **RDS Cluster ID** – The default will be the ID of the original cluster. If the original cluster still exists, the recovery operation will fail, unless you change the ID.
- **RDS Instance ID** – The default will the ID of the original instance. If the original instance still exists, the recovery operation will fail.
  Type a new ID to recover a new database. N2WS will use this instance ID for the writer, and in the case of multi-AZ, it will create the reader with this name with `_reader` added at the end.
- **RDS Cluster Snapshot ID** – Displays the snapshot ID.
- **Instance Type** – The type or class of the DB instances.
- **Port** – The port of the database. The default is the port of the original backed-up database.
- **Zone** – The AZ of the cluster in case of single AZ. If using a subnet group, leave as is.
- **Subnet Group** – Whether to launch the cluster in a VPC subnet or not, and to which subnet group. The default is the value from the original backed-up cluster.
- **Publicly Access** – Whether the cluster will be publicly accessible or not. The default is the access from the original backed-up instance.

Select **Recover Aurora Clusters** when finished.

10.6 Redshift Cluster Recovery

When a backup to recover includes snapshots of Redshift clusters, the **Redshift Clusters** tab opens. All Redshift clusters in the current backup are listed. You can change the following options:

- **Recover** – Clear **Recover** to not recover the current cluster.

- **Zone** – The AZ of the cluster. By default, it will be the zone of the backed-up cluster, but this can be changed.
  Currently, recovering a cluster into a VPC subnet is not supported by N2WS. You can always recover from the snapshot using AWS Management Console.
- **Cluster ID** – The default will the ID of the original cluster. If the original cluster still exists, the recovery operation will fail. To recover a new cluster, type a new ID.
- **Cluster Snapshot ID**– Displays the snapshot ID.
- **Node Type** and **Nodes** – For information only. Changing these fields is not supported by AWS.
- **Port** – The port of the cluster. The default is the port of the original backed-up cluster.
- **Subnet Group** – Whether to launch the cluster in a VPC subnet or not, and to which subnet group. The default will be the value from the original backed-up cluster. You can recover a cluster from outside a VPC to a VPC subnet group, but the other way around is not supported.
- **AWS Credentials** - You can choose to use different AWS credentials and enter your keys.

10.7 DynamoDB Table Recovery

When a backup to recover includes DynamoDB Table backups, the **DynamoDB Tables** tab opens.

Note:	If you reach the limit of the number of tables that can be recovered at one time, you will need to wait until they have completed before starting the recovery of additional tables.

All DynamoDB tables in the current backup are listed. You can change the following options:

- **Recover** – Clear **Recover** to not recover a table.
- **Region** – The Region where the table will be recovered, which is the same region as the backup.
- **Table Name** – The default will the Name of the original table. However, if the original table still exists, the recovery operation will fail. To recover to a new table, type a new Name.
- **Backup Name** – Displays the name of the backup.
- **AWS Credentials** - You can choose to use different AWS credentials and enter your keys.

During backup, N2WS retains the DynamoDB tags at the table level and the Time To Live (TTL) metadata and enables these attributes on recovery.
During the recovery process, a confirmation message appears with a reminder to recreate the following settings on the restored DynamoDB tables *MANUALLY*: Auto Scaling policies, IAM policies, CloudWatch metrics and alarms.

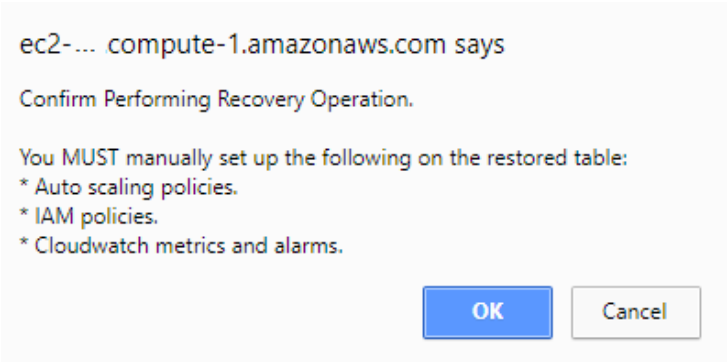

10.8 EFS Recovery

When a backup includes EFS backups, the **Recover EFS** tab is available.

Note: The AWS role "**AWSBackupDefaultServiceRole**" is required for recovery.

1. In the **Backup Monitor** screen, select an EFS backup and then select **Recover**.
2. In the Target EFS list, select the target to restore to:
   - **New** - Recover to a separate EFS
   - **Original** - Recover to the same EFS
3. Select **Recover Volumes**.

To view the progress of the recovery:

1. In N2WS, select the **Recovery Monitor** tab.
2. To view details of the recovery process, select the recovery record and select **Log**. Select **Refresh** as needed.

Note: Regular recoveries to original and new EFSs are supported. For DR, only recovery to a new EFS is supported.

11 Disaster Recovery (DR)

N2WS' DR (Disaster Recovery) solution allows you to recover your data and servers in case of a disaster. DR will help you recover your data for whatever reason your system was taken out of service. N2WS flexibility allows users to copy their backup snapshots to multiple AWS regions as well as to various AWS accounts, combining cross-account and cross-region options.
What does that mean in a cloud environment like EC2? Every EC2 region is divided into AZs which use separate infrastructure (power, networking, etc.). Because N2WS uses EBS snapshots you will be able to recover your EC2 servers to other AZs. N2WS' DR is based on AWS's ability to copy EBS snapshots between regions and allows you the extended ability to recover instances and EBS volumes in other regions. You may need this ability if there is a full-scale outage in a whole region. But it can also be used to migrate instances and data between regions and is not limited to DR. If you use N2WS to take RDS snapshots, those snapshots will also be copied and will be available in other regions.

- DynamoDB Tables - DR for DynamoDB tables is currently not supported by AWS.
- Redshift Clusters - Currently N2WS does not support DR of Redshift clusters. If you enable DR on a policy containing Redshift clusters, they will be ignored at the DR stage. You can enable copying Redshift snapshots between regions automatically by enabling cross-region snapshots using the EC2 console.

11.1 Configuring DR

After defining a policy, select the **DR** tab.

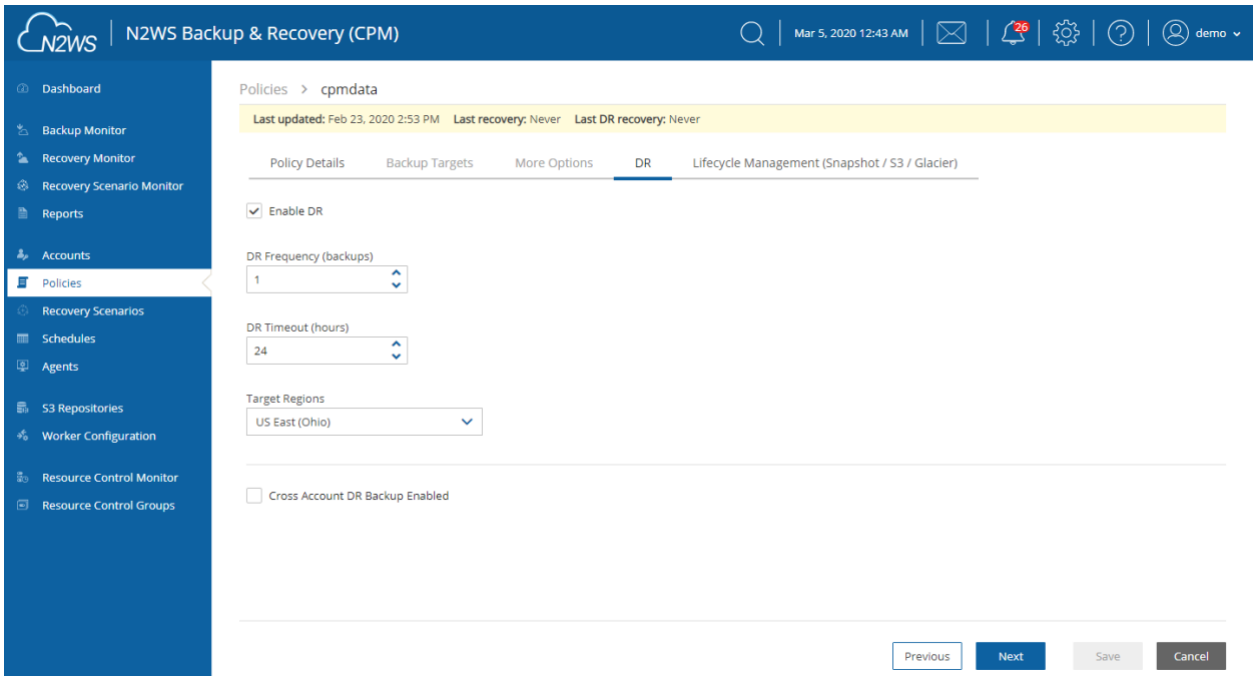

In the DR Options screen, configure the following and then select **Save**.

- **Enable DR** – Select to display additional fields.

- **DR Frequency (backups)** – Frequency of performing DR in terms of backups. On each backup, the default is to copy snapshots of all supported backups to other regions. To reduce costs, you may want to reduce the frequency. See section 11.4 below for considerations in planning DR.
- **DR Timeout (hours)** – How long N2WS waits for the DR process on the policy to complete. DR copies data between regions over a WAN (Wide Area Network) which can take a long time. N2WS will wait on the copy processes to make sure they are completed successfully. If the entire DR process is not completed in a certain timeframe, N2WS assumes the process is hanging, and will declare it as failed. Twenty-four hours is the default and should be enough time for a few 1 TiB EBS volumes to copy. Depending on the snapshot, however, you may want to increase or decrease the time.
- **Target Regions** – List of regions of region or regions that you want to copy the snapshots of the policy to.

To configure Cross-Account backup, see section 12.1.

11.2 About the DR Process

Things to know about the DR process:

- N2WS' DR process runs in the background.
- It starts when the backup process is finished. N2WS determines then if DR should run and kicks off the process. In the **Backup Monitor,** you will see the 'In Progress' status.

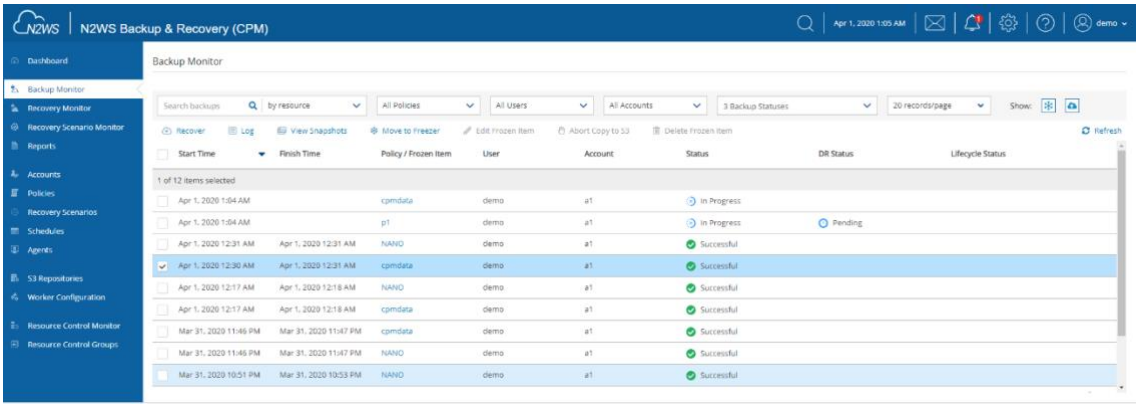

- N2WS will wait until all copy operations are completed successfully before declaring the DR status as **Completed** as the actual copying of snapshots can take time.
- As opposed to the backup process that allows only one backup of a policy to run at one time, DR processes are completely independent. This means that if you have an hourly backup and it runs DR each time, if DR takes more than an hour to complete, the DR of the next backup will begin before the first one has completed.
- Although N2WS can handle many DR processes in parallel, AWS limits the number of copy operations that can run in parallel in any given region to avoid congestion. See section 11.4.2.
- N2WS will keep all information of the original snapshots and the copied snapshots and will know how to recover instances and volumes in all relevant regions.
- The automatic retention process that deletes old snapshots will also clean up the old snapshots in other regions. When a regular backup is outside the retention window and its snapshots are deleted, so are the DR snapshots that were copied to other regions.

11.3 DR and Mixed-region Policies

N2WS supports backup objects from multiple regions in one policy. In most cases, it would probably not be the best practice, but sometimes it is useful. When you choose a target region for DR, DR will copy all the backup objects from the policy which are not already in this region to that region. For example, if you back up an instance in Virginia and an instance in North California, and you choose N. California as a target region, only the snapshots of the Virginia regions will be copied to California. So, you can potentially implement a mutual DR policy: choose Virginia and N. California as target regions and the Virginia instance will be copied to N. California and vice versa. This can come in handy if there is a problem or an outage in one of these regions. You can always recover the instance in the other region.

11.3.1 DR of EFS

Disaster Recovery is supported for EFS to a new EFS *only*. DR to an original EFS is not supported.

- If the source and target regions are the same for DR, no action is required as the target region will default to the source.
- If the target region is different than the source, the target region must have a backup vault with the same name as the source and must be specified using a tag before the DR begins.
  - For the "Default" vault, if this is the initial time copying a snapshot to the DR region, go to the AWS Backup console and activate the vault by clicking **Backup vaults**.
  - For a non-default custom vault, a vault with the same name needs to be created in the DR region. For example, if the source region's vault name is "Test", the DR region also must include a vault with the name "Test".

To set a custom vault name for cross-region EFS DR:

Before the DR, add a tag to the resource with the key '**cpm_dr_backup_vault**' and the value of the custom backup vault ARN:

Key='**cpm_dr_backup_vault**:REGION', Value ='BACKUP_VAULT_ARN'

Add a key for each target region that is different from the source.

Note: Cross-account DR for EFS is not currently supported.

11.4 Planning your DR Solution

11.4.1 Time and Financial Considerations

There are some fundamental differences between local backup and DR to other regions. It is important to understand the differences and their implications when planning your DR solution. The differences between storing EBS snapshots locally and copying them to other regions are:

- Copying between regions is transferring data over a WAN. It means that it will be much slower than moving data locally. A data transfer from the U.S to Australia or Japan will take considerably more time than a local copy.
- AWS will charge you for the data transfer between regions. This can affect your AWS costs, and the prices are different depending on the source region of the transfer. For

example, in March 2013, transferring data out of U.S regions will cost 0.02 USD/GiB and can climb up to 0.16 USD/GiB out of the South America region.

As an extreme example: You have an instance with 4 x TiB EBS volumes attached to it. The volumes are 75% full. There is an average of 3% daily change in data for all the volumes. This brings the total size of the daily snapshots to around 100 GiB. Locally you take 4 backups a day. In terms of cost and time, it will not make much of a difference if you take one backup a day or four, which is true also for copying snapshots, since that operation is incremental as well. Now you want a DR solution for this instance. Copying it every time will copy around 100 GiB a day. You need to calculate the price of transferring 100 GiB a day and storing them at the remote region on top of the local region.

11.4.2 Timing your DR processes

You want to define your recovery objectives both in local backup and DR according to your business needs. However, you do have to take costs and feasibility into consideration. In many cases it is ok to say: For local recovery I want frequent backup, four times a day, but for DR recovery it is enough for me to have a daily copy of my data. Or, maybe it is enough to have DR every two days. There are two ways to define such a policy using N2WS:

- In the definition of your policy, select the frequency in **DR Frequency (backups).** If the policy runs four times a day, configure DR to run once every four backups. The DR status of all the rest will be **Skipped**.
- Or, define a special policy for the DR process. If you have a **sqlserver1** policy, define another one and name it something like **sqlserver1_dr**. Define all targets and options the same as the first policy, but choose a schedule relevant for DR. Then define DR for the second policy. Locally it will not add any significant cost since it is all incremental, but you will get DR only once a day.

11.4.3 Performing DR on the N2WS Server (The cpmdata Policy)

To perform DR recovery, you will need your N2WS server up and running. If the original server is alive, then you can perform recovery on it across regions. You want to prepare for the case where the N2WS server itself is down. You may want to copy your N2WS database across regions as well. Generally, it is not a bad idea to place your N2WS server in a different region than your other production data. N2WS has no problem working across regions and even if you want to perform recovery because of a malfunction in only one of the AZs in your region, if the N2WS server happens to be in that zone, it will not be available.

To make it easy and safe to back up the N2WS server database, there is a special policy named `cpmdata`. Although N2WS supports managing multiple AWS accounts, the only account that can back up the N2WS server is the one that owns it, i.e. the account used to create it. Define a new policy and name it `cpmdata` (case insensitive), and it will automatically create a policy that backs up the CPM data volume.

Note:	Application consistency is disabled by default for the `cpmdata` policy. When enabled, the CPM will run application consistent scripts. See section 4.2.1.

Not all options are available with the `cpmdata` policy, but you can control Scheduling, Number of generations, and DR settings.

When setting these options, remember that at the time of recovery you will need the most recent copy of this database, since older ones may point to snapshots that no longer exist and not have newer ones yet. Even if you want to recover an instance from a week ago, you should always use the latest backup of the `cpmdata` policy.

11.5 DR Recovery

DR recovery is similar to regular recovery with a few differences:

- When you select **Recover** for a backup that includes DR (DR is in **Completed** state), you get the same Recovery Panel screen with the addition of a drop-down list.
- The DR Region default is **Origin**, which will recover all the objects from the original backup. It will perform the same recovery as a policy with no DR.
- When choosing one of the target regions, it will display the objects and will recover them at the selected region.

Recommendation: N2W Software strongly recommends that you perform recovery drills occasionally to be sure your recovery scenario works. It is not recommended to try it for the first time when your servers are down. Each policy on the policy screen shows the last time recovery was performed on it. Use the last recovery time data to track recovery drills.

11.5.1 DR Instance Recovery

Volume recovery is the same in any region. With instance recovery there are a few things that need considering. An EC2 instance is typically related to other EC2 objects:

- Image ID (AMI)
- Key Pair
- Security Groups
- Kernel ID
- Ram disk ID

These objects exist in the region of the original instance, but they do not mean anything in the target region. In order to launch the instance successfully, you will need to replace these original objects with ones from the target region:

- **Image ID (AMI)** - If you intend to recover the instance from a root device snapshot, you will not need a new image ID. If not (as in all cases with Windows and instance store-based instances), you will need to type a new image ID. If you use AMIs you prepared, you should also prepare them at your target regions and make their IDs handy when you need to recover. If needed, AMI Assistant can help you find a matching image. See section 10.2.4.
- **Key Pair** - You should have a key pair created with AWS Management Console ready so you will not need to create it when you perform a recovery.
- **Security Groups** - In a regular recovery, N2WS will remember the security groups of the original instance and use them as default. In DR recovery, N2WS cannot choose for you. You need to choose at least one, or the instance recovery screen will display an error. Security groups are objects you own, and you can easily create them in AWS

Management Console. You should have them ready so you will not need to create them when you perform recovery.

- **Kernel ID** - Linux instances need a kernel ID. If you are launching the instance from an image, you can leave this field empty, N2WS will use the kernel ID specified in the AMI. If you are recovering the instance from a root device snapshot, you need to find a matching kernel ID in the target region. If you do not do so, a default kernel will be used, and although the recovery operation will succeed and the instance will show as running in AWS Management Console, it will most likely not work. AMI Assistant can help you find a matching image in the target region. See section 10.2.4. When you find such an AMI, copy and paste its kernel ID from the AMI Assistant window.
- **RAMDisk ID** - Many instances do not need a RAM disk at all and this field can be left empty. If you need it, you can use AMI Assistant the same way you do for Kernel ID. If you're not sure, use the AMI Assistant or start a local recovery and see if there is a value in the RAMDisk ID field.

11.5.2 Setting a Tag with an AMI ID for Cross-Account DR Recovery

The CPM can add a tag with an AMI ID to a resource during backup. The tag will hold the AMI ID that is expected to be present on the AWS account in case of recovery to a different AWS account.

Example of tag format that will be used only on the region/account combination specified:

Key = '**cpm_dr_recover_ami:**REGION:ACCOUNT'; Value = 'ami-XXXXX'

In this case, the region and account are optional.

Example of tag format for a tag that will be used on any region/account combination:

Key = '**cpm_dr_recover_ami';** Value = 'ami-XXXXX'

When this tag is found and there is no other proper option for instance recovery, the CPM will then use this AMI if the recovery region and account fits.

Note: The AMI with an ID that is provided in the value of the tag is expected to be available on the other AWS account while recovering the instance.

11.5.3 DR of Encrypted Volumes, AMIs and RDS Instances

N2WS supports DR of encrypted EBS volumes. If you are using KMS keys for encryption:

- N2WS will seek a KMS key in the target region, which has the same alias.
- The AWS ID of the DR account should be added to the 'Other AWS accounts' section on a Backup account.

To configure your cross-region DR:

Create a matching-alias key in the source and in the remote region for N2WS to use automatically in the DR copy process.

- If a matching key is not found in the target region, the DR process will fail.
- If the key uses the default encryption, then it will be copied to the other region with the default encryption key as well.

- N2WS supports copy of AMIs with encrypted volumes with the same logic it uses for volumes.
- N2WS supports cross-region DR of encrypted RDS databases, except from the Asia Pacific (Hong Kong) region.

To add the AWS ID of the DR Account to the 'Other AWS accounts' section of KMS on a Backup account:

1. Log on to your Backup AWS account and navigate to the KMS console.
2. Select your Customer managed keys.
3. Go to the 'Other AWS accounts' section.
4. Select Add other AWS accounts.
5. In the box, enter the AWS account ID of the DR account.

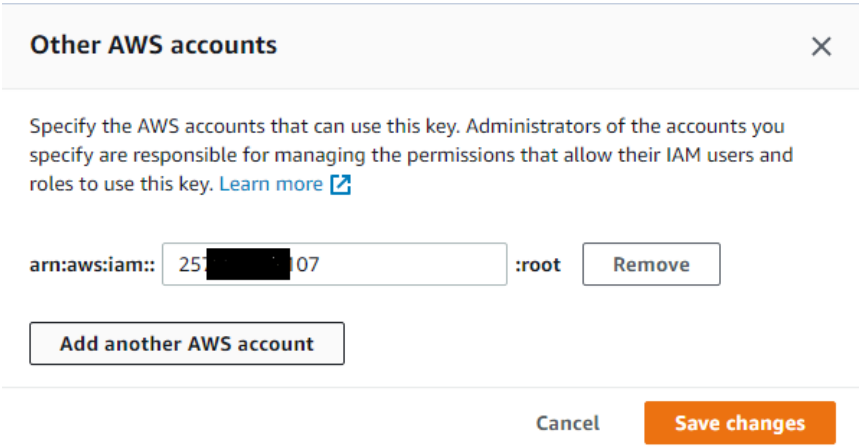

Note:	To let N2WS see keys and aliases, add these two permissions to the IAM policy that N2WS is using: **kms:ListKeys, kms:ListAliases.**

To recover an EFS with a non-default KMS, in AWS configure the KMS as follows:
In the **Key user** field, select **Add** and choose "**AWSBackupDefaultServiceRole**".

For information about support for custom DR encryption keys for different regions and accounts using the '**cpm_dr_encryption_key**' tag, see https://support.n2ws.com/portal/kb/articles/cpm-supports-custom-encryption-keys-for-dr.

11.5.4 A Complete Disaster Recovery Scenario

Let's assume a real disaster recovery scenario: The region of your operation is completely down. It means that you do not have your instances or EBS volumes, and you do not have your N2WS Server, as it is down with all the rest of your instances. Here is Disaster Recovery step by step:

1. In the AWS Management Console:
   a. Find the latest snapshot of your `cpmdata` policy by filtering snapshots with the string `cpmdata`. N2WS always adds the policy name to any snapshot's description.
   b. Sort by **Started** in descending order and it will be the first one on the list.

c. Create a volume from this snapshot by right-selecting it and choosing **Create Volume from Snapshot**. You can give the new volume a name so it will be easy to find later.

2. Launch a new N2WS Server at the target region. You can use the Your Software page to launch the AWS Marketplace AMI. Wait until you see the instance in **running** state.
3. As with regular configuration of a N2WS server:
   a. Connect to the newly created instance using HTTPS.
   b. Approve the SSL certificate exception.
   Assuming the original instance still exists, N2WS will come up in **recovery** mode, which means that the new server will perform recovery and not backup.
   c. If you are running the BYOL edition and need an activation key, most likely you do not have a valid key at the time, and you do not want to wait until you can acquire one from N2W Software.
   You can quickly register at N2WS Free Edition. In step 2 of the registration, use your own username and type any password. In step 3, choose the volume you just created for the CPM data volume. Afterwards, complete the configuration.
4. With a working N2WS server, you can perform any recovery you need at the target (current) region:
   a. Select the backup you want to recover.
   b. Select Recover.
   c. Choose the target region from the drop-down list.

Note:	If your new server allows backup (it can happen if you registered to a different edition or if the original one is not accessible), it can start to perform backups. If that is not what you want, it is best to disable all policies before you start the recovery process.

   d. You can recover all the backed-up objects that are available in the region.

11.6 DR Monitoring and Troubleshooting

DR is a straightforward process. If DR fails, it probably means that either a copy operation failed, which is not common, or that the process timed-out. You can track DR's progress in the **Recovery Monitor** **Log** screen where every stage and operation during DR is recorded:

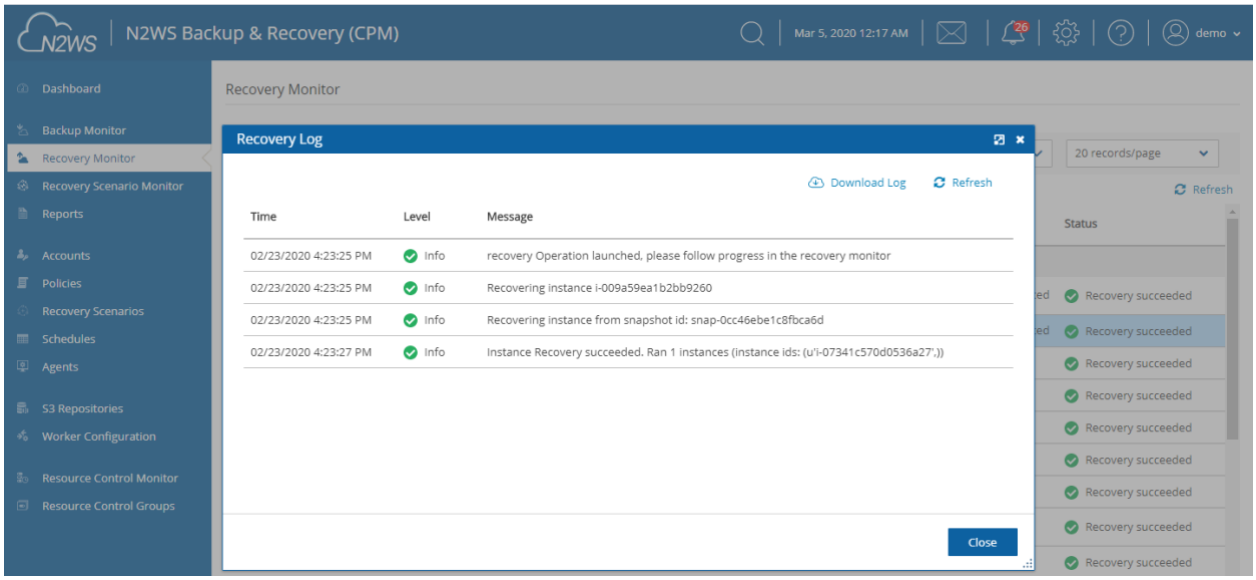

You can also view DR snapshot IDs and statuses in the **View Snapshots** screen of the backup:

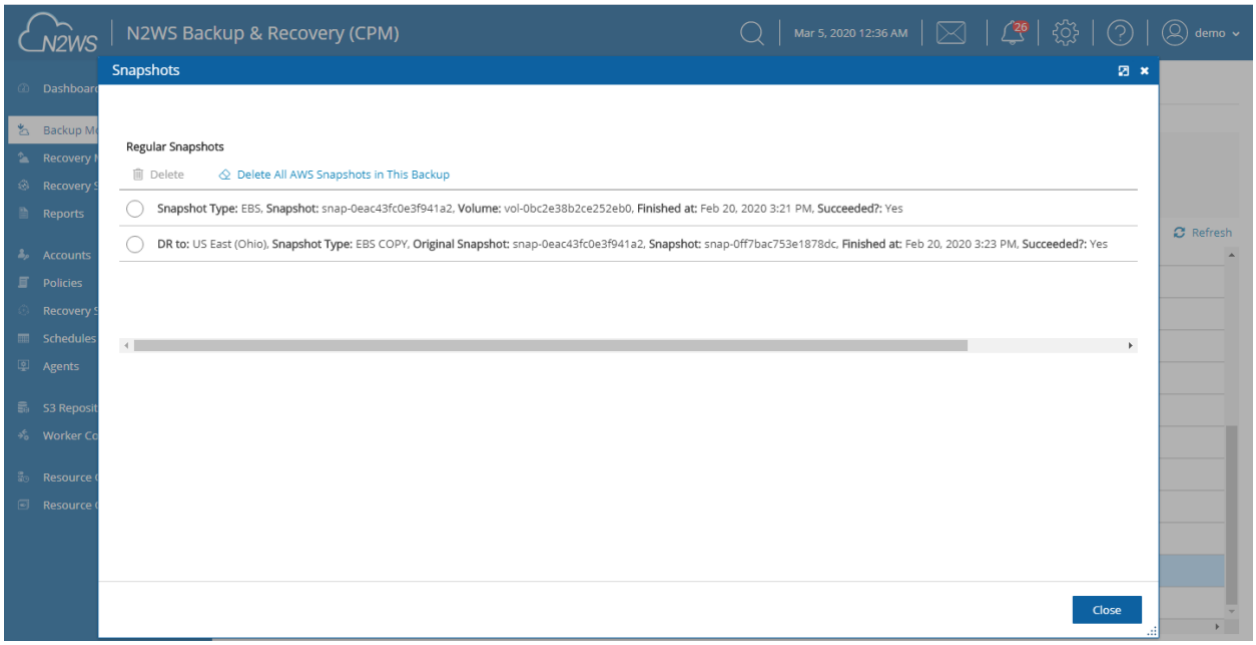

Every DR snapshot is displayed with region information and the IDs of both the original and the copied snapshots. In the **Snapshots** list, you can choose to **Delete All AWS Snapshots in This Backup**.

If DR fails, you will not be able to use DR recovery. However, some of the snapshots may exist and be recoverable. You can see them in the snapshots screen and, if needed, you can recover from them manually.

If DR keeps failing because of timeouts, you may need to increase the timeout value for the relevant policy. The default of 24 hours should be enough, but there may be a case with a very large amount of data, that may take longer.

Reminder: You can only copy a limited number of snapshots to a given region at one time. Currently the number is 5. If the limit is reached, N2WS will wait for copy operations to finish before it continues with more of them which can affect the time it takes to complete the DR process.

12 Cross-Account DR, Backup and Recovery

Available only in Advanced and Enterprise Editions, N2WS' cross-account functionality allows you to automatically copy snapshots between AWS accounts as part of the DR module. With cross-region DR, you can copy snapshots between regions as well as between accounts and any combination of both. In addition, you can recover resources (e.g. EC2 instances) to a different AWS account even if you did not copy the snapshots to that account. This cross-account functionality is important for security reasons.
The ability to copy snapshots between regions can prove crucial if your AWS credentials have been compromised and there is a risk of deletion of your production data as well as your snapshot data. N2WS utilizes the **snapshot share** option in AWS to enable copying them across accounts. Cross-account functionality is currently supported only for EC2 instances, EBS volumes and RDS instances, including Aurora.
Cross-account functionality is enabled for encrypted EBS volumes and instances with encrypted EBS volumes and RDS databases.

Note: Cross-region DR is not supported for RDS databases in the Asia Pacific (Hong Kong) region.

- Users will need to share the encrypted key used for the encryption of the volumes or RDS instance to the other account as N2WS will not do it.
- In addition, N2WS expects to find a key in the target account with the same alias as the original key (or just uses the default key).

For information on sharing encryption keys between different accounts, see https://support.n2ws.com/portal/kb/articles/cpm-supports-custom-encryption-keys-for-dr

If a matching encryption key is not found with an alias or with custom tags, the behavior of the backup depends on the setting in the **Encryption Key Detection** list in the **Security** tab of the **General Settings** screen:

- **Use Default Key** – If the encryption key is not matched, the default encryption key is used.
- **Strict** – DR encryption key must match, either with an alias or a custom tag.
- **Use Default Key & Alert** – Use the default key and send an alert.

N2WS can support a DR scheme where a special AWS account is used only for snapshot data. This account's credentials are not shared with anyone and used only to copy snapshots to. The IAM credentials used in N2WS can have limited permissions that do not allow snapshot deletion.
N2WS will tag outdated snapshots instead of actually deleting them, allowing an authorized user to delete them separately using the EC2 console or a script. The tag '**cpm_deleted**' will have a value of '**CPM deleted this snapshot (**<time_of_deletion>**)**'. Also, you may choose to keep the snapshots only in the vault account and not in their original account. This will allow you to save storage costs and utilize the cross-recovery capability to recover resources from the vault account back to the original one.

12.1 Configuring Cross-Account Backup

Once you have created a DR Account with the **Account Type** DR, you can configure cross-account DR from the **DR** tab of a policy

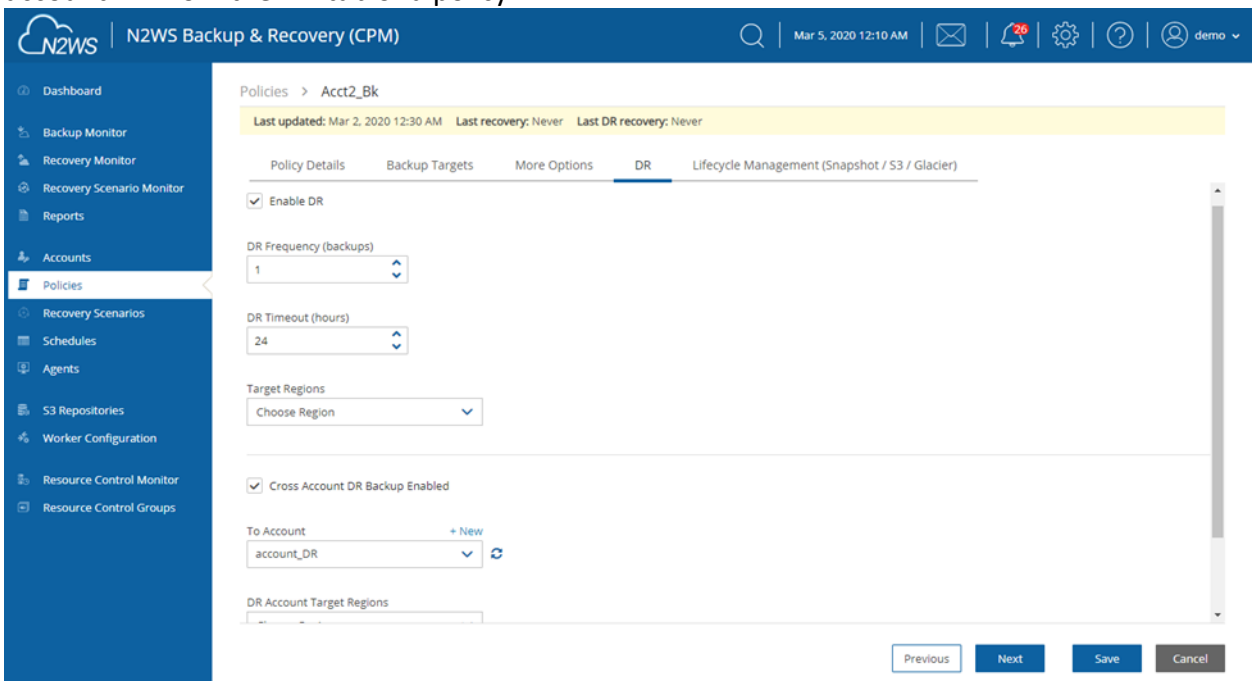

Cross-account fields will be available only if your N2WS is licensed for cross-account functionality. See the pricing and registration page in our website to see which N2WS editions include cross-account backup & recovery.

Once you select **Cross-Account DR Backup Enabled**, other fields become visible:

- **To Account** – Which account to copy the snapshots to. This account needs to have been defined as a **DR Account** in the **Accounts** screen.
- **DR Account Target Regions** – Which region or regions you want to copy the snapshots of the policy to. To include all of the Target Regions selected for backup, select **Original** in the list. Select additional regions as needed.
- **Keep Original Snapshots** – When enabled, the original snapshot from the source region will be kept. If disabled, once the snapshot is copied to the DR account, it will be deleted from the source region.

12.2 Cross-Account DR and Clean-Up

N2WS performs clean-up on backup policies and deletes backups and snapshots that are out of the retention window, according to the policy's definition. By default, N2WS will clean up snapshots copied to other accounts as well. However, if you do not wish for N2WS to clean up, because you want to provide IAM credentials that are limited and cannot delete data, you have that option. If you defined the DR Account with **Allow Deleting Snapshots** set as False, N2WS will not try to delete snapshots in the DR Account. It will rather flag a snapshot for subsequent deletion by adding a tag to the snapshot called **cpm_deleted**. The tag value will contain the time when the snapshot was flagged for deletion by N2WS.

When using this option, occasionally make sure that these snapshots are actually deleted. You can either run a script on a schedule, with proper permissions, or make it delete all snapshots with the tag **cpm_deleted**. Or, using the EC2 console, filter snapshots by the tag name and delete them.

12.3 Cross-Account with Cross-Region

If you configure the backup policy to copy snapshots across accounts as well as across regions, be aware of how the increased number of copies might affect your AWS costs.

12.4 Cross-Account Recovery

If you have cross-account functionality enabled in your N2WS license, and even if you actually configured N2WS to copy snapshots between accounts, you can recover across accounts. This is already mentioned in the recovery section (section 10). You need to choose which account to recover the resource (EC2 instance, EBS volume or RDS database) to.

Note:	Only account type **DR Account** may be the target of a cross-account recovery.

When copying snapshots between accounts and not keeping the original snapshots, you will also have the option to restore the instance/volume to the original account. N2WS will utilize the AWS **share snapshot** option to enable recovering resources across accounts.

Note:	There is an AWS limitation for restoring encrypted RDS snapshots from a DR AWS account. Directly restoring a cross-account DR copy of encrypted RDS snapshots is not supported. As a workaround, you can either restore directly to the DR AWS account, or the snapshot data can be copied back to the original AWS account, and then the restore can work as intended from there.

13 File-level Recovery

N2WS supports file-level recovery. N2WS does backup on the volume and instance level and specializes in instant recovery of volumes and complete instances. However, in many cases a user may want to access specific files and folders rather than recovering an entire volume.
In previous versions of N2WS, you could recover a volume, attach it to an instance, mount it and then access the data from within that instance. After completing the restore, assuming the volume is no longer needed, the user needed to unmount, detach and delete the volume.
N2WS now automates this entire process.

Note: The root volume of instances purchased from Amazon Marketplace, such as instances with a product code, cannot be partially recovered by file-level recovery. The data volumes of such instances, if they exist, will be recovered.

In the **Backup Monitor**, select an instance or volume backup and then select **Recover**. In the Recover screen, select **Explore** in the **Instances** or **Independent Volumes** tab. When selecting **Recover Volumes Only**, **Explore** is available in the **Volume Recovery** screen.

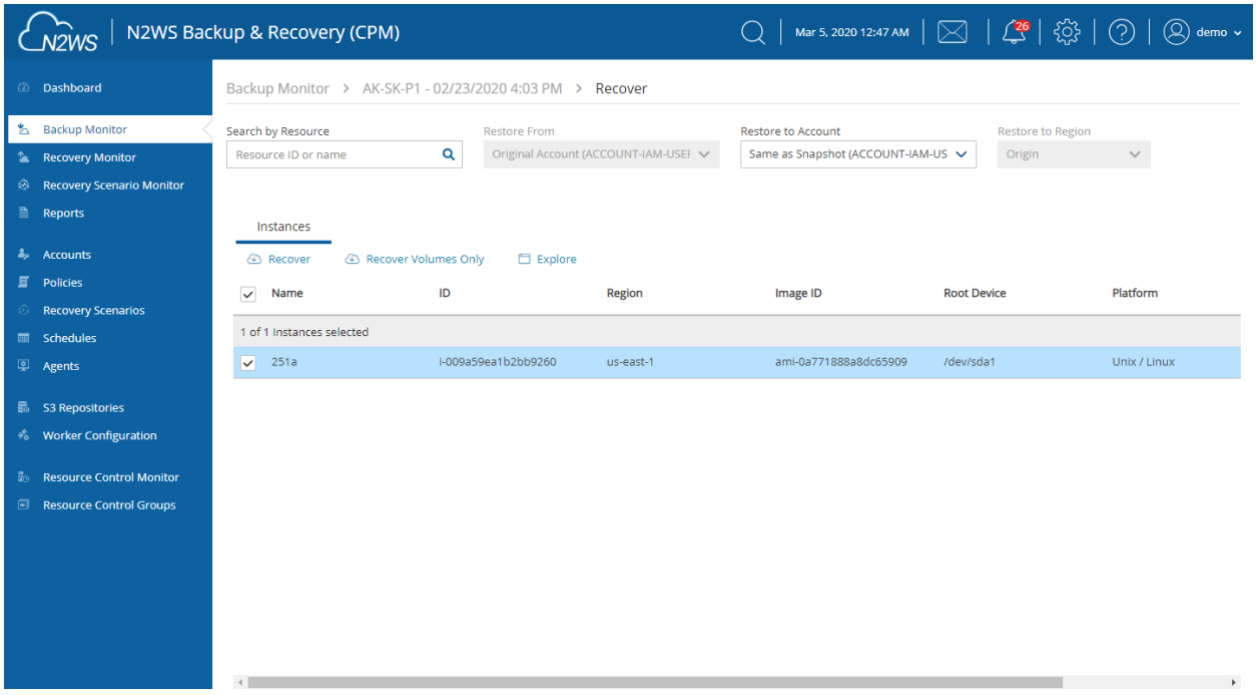

N2WS will open a new browser tab showing a **File Explorer-like** view of the entire instance or a specific volume. You will be able to browse, search for files, and download files and folders.

Note: Files in an **Explore** volume may actually be soft links (symbolic links) to other files. Trying to access this type of file may result in an error. However, the file is accessible via its real path. For example, `root/folder2/file2` is a soft link to `root/folder1/file1`, where `/root/folder1/file1` is the real path.

On the right side of any file or folder, there is a green Download arrow for downloading the file or folder. Folders are downloaded as uncompressed zip files.

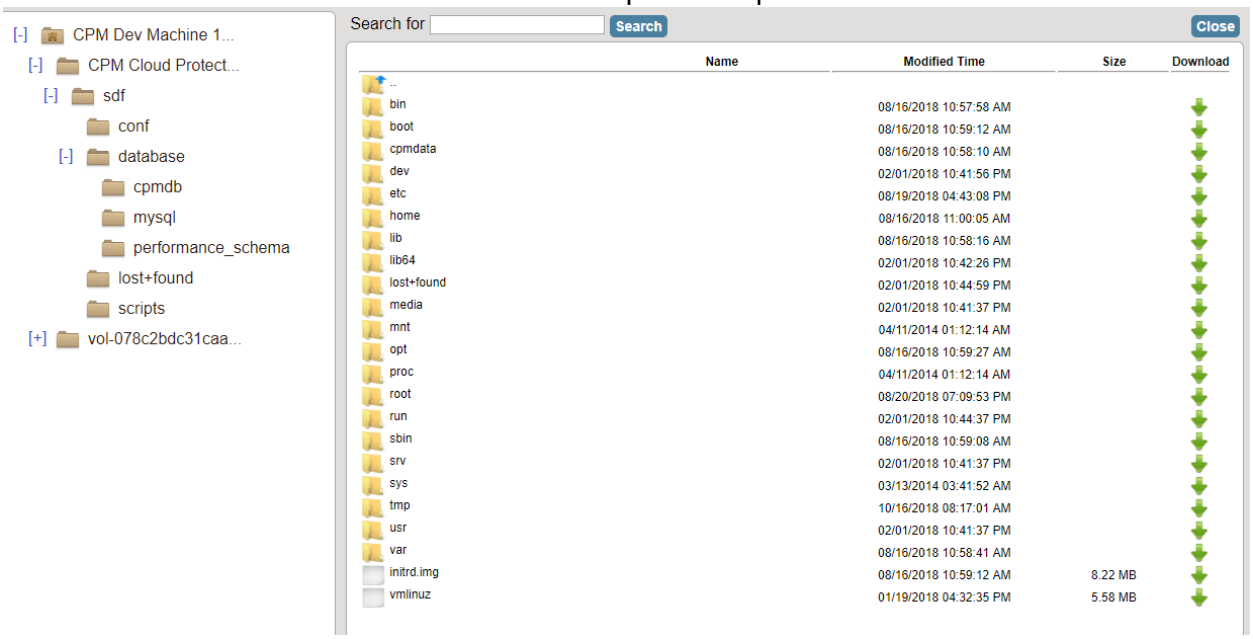

To perform these operations, N2WS needs to be able to use AWS credentials belonging to the N2WS server instance account, with sufficient permissions to create and attach volumes. By default, N2WS will use the same credentials used to initially configure the instance, but they can be modified using the **General Settings** screen.
File-level recovery requires N2WS to recover volumes in the background and attach them to a 'worker' launched for the operation. The worker will be launched in the same account and region as the snapshots being explored, using a pre-defined worker configuration. See section 22 to configure a 'worker' instance in the region that the snapshots exist.

Note:	The worker will communicate with the N2WS server over both HTTPS and SSH. Verify that your configuration allows such communication.

There are a few limitations:

- File-level recovery is supported only for file system types Ext2, Ext3, Ext4, NTFS, XFS, Btrfs.

Note:	If several XFS volumes have the same UUID, they cannot be mounted.

- **Explore** works only on simple volumes and Logical Volume Management (LVM). LVM is supported with file-level restore on Linux, as well as for Windows dynamic disks. Additionally, disks defined with Microsoft Storage Spaces are not supported.
- In order to **Explore** snapshots taken in a different region than where the N2WS server is, it is required to configure a 'worker' instance in the region that the snapshots exist. See section 22.

After you complete the recovery operation, select **Close** for all the resources to be cleaned-up and to save costs. Even if you just close the tab, N2WS will detect the redundant resources and clean them up, but it is recommended that you use **Close**.

14 Tag-based Backup Management

Cloud and specifically AWS, is an environment based largely on automation. Since all the functionality is available via an API, scripts can be used to deploy and manage applications, servers and complete environments. There are very popular tools available to help with configuring and deploying these environments, like Chef and Puppet.
N2WS allows configuring backup using automation tools by utilizing AWS tags. By tagging a resource (EC2 instance, EBS volume, EFS, DynamoDB, RDS instance, Aurora Cluster or Redshift cluster), N2WS can be notified what to do with this resource, and there is no need to use the UI. Since tagging is a basic functionality of AWS, it can be easily performed via the API and scripts.
To tag Aurora clusters, tag one of the cluster's DB instances and N2WS will pick it up and back-up the entire cluster.

Note: For information on using tags with Resource Control, see section 15.5.

14.1 The "cpm backup" Tag

To automate backup management for a resource, you can add a tag to that resource named **cpm backup** (lower case with a space). N2WS will identify this tag and parse its content. In this tag you will be able to specify whether to:

- Ignore the resource and remove it from all backup policies.
- Add the resource to a policy or list of policies.
- Create a new policy, based on an existing one (template), and then add the resource to it.

Note: The policy name on the 'cpm backup' tag is case sensitive and should be aligned with the policy name create on CPM.

If an AWS resource has 2 AWS tags with the same tag name, differing only by the case of the letters (upper, lower), then N2WS will back up just one tag. The tag name will be in the format of the first tag N2WS scans, and the tag value *may* be from the second tag. Check that tag names are in the same case.

Following is a summary table of all **cpm backup** tag values:

Purpose	cpm backup Tag Value	Examples
Add resource to existing backup policy. See 14.1.1.	*policy1*	*policy1 policy2 policy3*
Create policy from a template. See 14.1.2.	*new_policy1*:*existing_policy1*	

Purpose	cpm backup Tag Value		Examples
Set backup options for EC2 instances. See 14.1.3.	**only-snaps** (create AMIs without reboot) **initial-ami** **only-amis** **only-amis-reboot** (create AMIs with reboot) **app-aware** (Windows instance backup agent is same as snapshot and AMI options) **app-aware-vss** (Enable application consistent with VSS) **app-aware-script** (Enable application consistent without VSS)		*policy1***#only-snaps** *new_policy:existing_policy***#only-amis** *policy1***#initial-ami#app-aware**
Set backup options for EFS instances. N2WS will override EFS configuration with tag values. See 14.1.4.	**key**	**value**	*policy1***+vault=**Default**+exp_opt=D+exp_opt_val**=1
	vault **role_arn** **cold_opt** **cold_opt_val** **exp_opt** **exp_opt_val**	Default (example) ARN of role Lifecycle transition: N, D, W, M, Y Integer for D,W,M,Y only When resource expires: P (Policy Gen), N, D, W, M, Y Integer for D, W, M, Y only	
Remove resource from all policies. See 14.1.5.	**no-backup**		
Exclude volumes from backup. See 14.1.6.	*policy1***#exclude** Note**:** Tagged instances are excluded from the **Exclude volumes** option in **General Settings** for **Tag Scan.** Tagged instances are only excluded with the **'#exclude'** tag.		*policy*1**#exclude** *policy2***#exclude**

14.1.1 Adding to a Policy or Policies

To add a resource (e.g. an EC2 instance) to an existing backup policy, all you need to do is to create the tag for this resource and specify the policy name (e.g. **policy1**):

tag key: **cpm backup**, tag value: **policy1**

To add the resource to multiple policies all you need to do is to add a list of policy names, separated by spaces:

policy1 policy2 policy3

14.1.2 Creating a Policy from a Template

To create a new policy and to add the resource to it, add a new policy name with a name of an existing policy which will serve as a template (separated by semicolon):

tag value: **new_policy1:existing_policy1**

You can also add multiple policy name pairs to create additional policies or create a policy (or policies) and to add the resource to an existing policy or policies.

When a new policy is created out of a template, it will take the following properties from it:

- Number of generations
- Schedules
- DR configuration
- Script/agent configuration
- Retry configuration

It will not inherit any backup targets, so you can use a real working policy as a template or an empty one.

For Script definitions:

If backup scripts are defined for the template policy, the new one will keep that definition but will not initially have any actual scripts. You are responsible to create those scripts. Since the N2WS server is accessible via SSH you can automate script creation. In any case, since scripts are required, the backups will have a failure status and will send alerts, so you will not forget about the need to create new scripts.

For Windows instances with a backup agent configured:

If that was the configuration of the original policy, the new instance (assuming it is a Windows instance) will also be assigned as the policy agent. However, since it does not have an authentication key, and since the agent needs to be installed and configured on the instance, the backups will have a failure status. Setting the new authentication key and installing the agent needs to be made manually.

Auto Target Removal for the new policy will always be set to **yes and alert**, regardless of the setting of the template policy. The basic assumption is that a policy created by a tag will automatically remove resources which do not exist anymore, which is the equivalent as if their tag was deleted.

14.1.3 Setting Backup Options for EC2 Instances

When adding an instance to a policy, or creating a new policy from template, you may make a few decisions about the instance:

- To create snapshots only for this instance.
- To create snapshots with an initial AMI.
- To schedule AMI creation only.

If this option is not set, N2WS will assume the default:

- Snapshots only for Linux.
- Snapshots with initial AMI for Windows instances by adding a backup option after the policy name. The backup option can be one of the following values:
  - **only-snaps**
  - **initial-ami**
  - **only-amis**
  - **only-amis-reboot**

  For example, with existing policy: `policy1#only-snaps`.
  Or, for a new policy based on template and setting AMI creation:
  `my_new_policy:existing_policy#only-amis`

Note: The **only-amis** option will create AMIs without rebooting them. The option **only-amis-reboot** will create AMIs with reboot.

For a Windows instance, you can also define backup with **app-aware**, i.e. a backup agent. It is used the same as the snapshots and AMI options.

- When adding the **app-aware** option, the agent is set to the default: VSS is enabled and backup scripts are disabled.
  - **app-aware-vss** - Enable application consistent with VSS.
  - **app-aware-script** - Enable application consistent without VSS.
- Additional configurations need to be made manually, and not with the tag.

You can also combine the backup options: `policy1#initial-ami#app-aware`.

14.1.4 Setting Backup Options for EFS Instances

EFS can be configured by creating the **cpm backup** tag with the following values. In this case, N2WS will override the EFS configuration with the tag values:

Key	Value
vault	Vault. Example: Default
role_arn	ARN of role. Example: arn:aws:iam::040885004714:role/service-role/AWSBackupDefaultServiceRole
cold_opt	Lifecycle transition: N – Never D – Days W – Weeks M – Months Y - Years
cold_opt_value	Integer for D, W, M, Y only
exp_opt	When does resource expire: P – Policy Generations N – Never D – Days W- Weeks M – Months Y - Years
exp_opt_val	Integer for D, W, M, Y only

Example:

```
cpm backup my_policy+vault=Default+exp_opt=D+exp_opt_val=1
```

N2WS will back up EFS to the default vault, and set its expiration date to 1 day.

Note: The max length for the **cpm backup** value is limited to 256 characters.

14.1.5 Tagging a Resource to be Removed from All Policies

By creating the **cpm backup** tag with the value **no-backup** (lower case), you can tell N2WS to ignore the resource and remove this resource from all policies. Also, see section 14.1.

14.1.6 Excluding Volumes from Backup

N2WS can exclude a volume from an instance which is backed up on policy using the "**cpm backup**" tag with **'#exclude'** added to the end of the policy name value.

- Add a tag to an instance that you want to back up:
  Key = **cpm backup**; Value = policy_name1 policy_name2

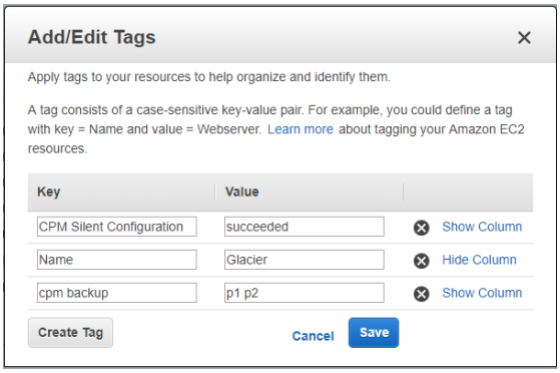

- Add a tag to volumes that you would like to exclude from being backed up: Key = **cpm backup**; Value = policy_name1**#exclude** policy_name2**#exclude**

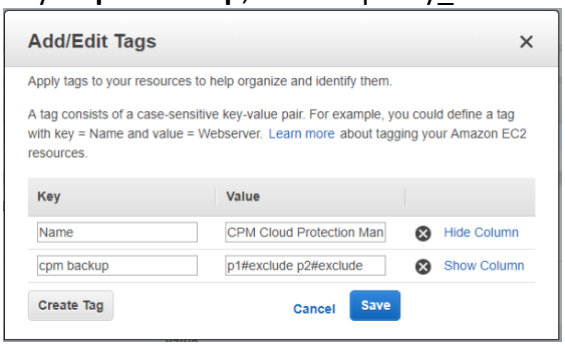

For example, if instance1 has 3 volumes and has a '**cpm backup**' tag with the value 'policy1', adding the '**cpm backup**' tag with value 'policy1**#exclude**' to a volume will remove it from the policy.

The instance with the excluded volume(s) will be added automatically as a backup target to the policy, after running Scan Tag.

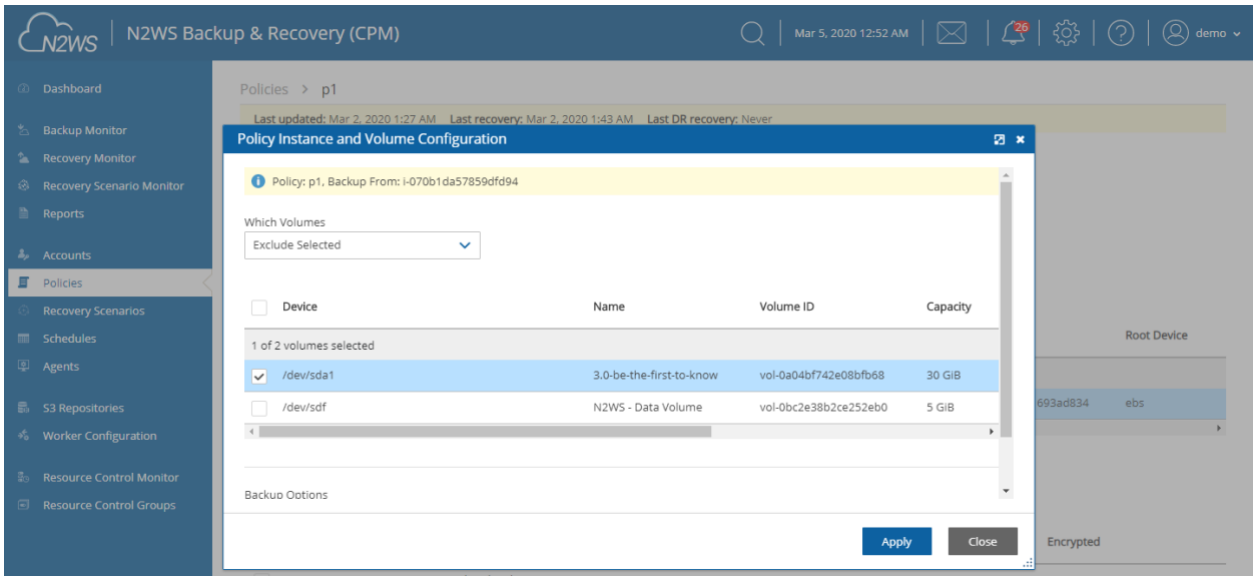

Note: Tagged instances are not included in the **Exclude volumes** option in the **Tag Scan** tab of **General Settings** and are excluded from backup only when tagged with '**#exclude'** for the policy.

14.2 Tag Scanning

Tag scanning can only be controlled by the admin/root user. When the scan is running, it will do so for all the users in the system but will only scan AWS accounts that have **Scan Resources** enabled. This setting is disabled by default. N2WS will automatically scan resources in all AWS regions.

1. In the **General Settings** tab, select the **Tag Scan** tab.
2. Select **Scan Resources**.
3. In the **Tag Scan interval** list, set the interval in hours for automatic scans.
4. To override the exclusion of volumes specified in the UI and to exclude instances tagged with '**#exclude**' for the policy, select **Exclude volumes**. See section 9.6.
5. Select **Save**.
6. To initiate a tag scan immediately, select **Scan Now**.

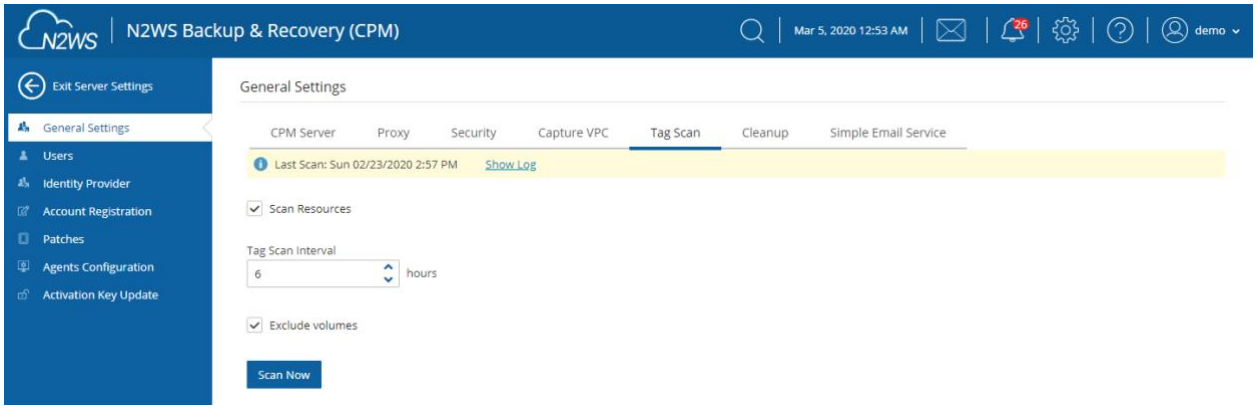

7. To view the Last Scan, select **Show Log**.

Note: Even if scanning is disabled, selecting **Scan Now** will initiate a scan.

If you do want automated scans to run, keep scanning enabled and set the interval in hours between scans using the **General Settings** screen. You will also need to enable **Scan Resources** for the relevant AWS accounts.

14.3 Pitfalls and Troubleshooting

The following topics should help guide you when developing tags.

14.3.1 Pitfalls

There are potential issues you should try to avoid when managing your backup via tags:

- The first is not to create contradictions between the tags content and manual configuration. If you tag a resource and it is added to a policy, and later you remove it from the policy manually, it may come back at the next tag scan. N2WS tries to warn you from such mistakes.
- Policy name changes can also affect tag scanning. If you rename a policy, the policy name in the tag can be wrong. When renaming a policy, correct any relevant tag values.
- When you open a policy that was created by a tag scan to edit it, you will see a message at the top of the dialog window: "* This policy was automatically added by tag scan".

Note:	Even if all the backup targets are removed, N2WS will not delete any policy on its own, since deletion of a policy will also delete all its data. If you have a daily summary configured (section 17.5), policies without backup targets will be listed.

- If the same AWS account is added as multiple accounts in N2WS, the same tags can be scanned multiple times, and the behavior can become unpredictable. N2W Software generally discourages this practice. It is better to define an account once, and then allow delegates (section 18.4) access to it. If you added the same AWS account multiple times (even for different users), make sure only one of the accounts in N2WS has **Scan Resources** enabled.

14.3.2 Troubleshooting

Sometimes you need to understand what happened during a tag scan, especially if the tag scan did not behave as expected, such as a policy was not created. In the **General Settings** screen, you can view the log of the last tag scan and see what happened during this scan, as well as any other problems, such as a problem parsing the tag value, that were encountered. Also, if the daily summary is enabled, new scan results from the last day will be listed in the summary.

Ensure tag format is correct:

Tips for ensuring correct tag formats are:

- When listing multiple policy names, make sure they are separated by spaces.
- When creating new policy, verify using a colon ':' and not a semi-colon ';'. The syntax is `new_policy1:existing_policy1`.
- Use a valid name for the new policy or it will not be created. An error message will be added to scan log.
- Use correct names for existing/template policies.
- Resource scanning order is NOT defined, so use policy names as existing/template only if you are sure that it exists in N2WS – defined manually or scanned previously.

15 Resource Control

- Resource Control allows users to stop and start Instances and RDS Databases for each Account during the course of a week. It also allows users to stop the resources at a designated time in the future.

  Note: RDS Aurora Clusters are *not* supported by Resource Control.

- A Group is the controlling entity for the stopping and starting of selected resources. Resource Control allows for stopping on one day of a week and starting on another day of the same week. Once an Off/On schedule is configured for a Group, N2WS will automatically stop and start the selected resource targets.
- Resources that are eligible and enabled for hibernation in AWS will be hibernated regardless of whether their current operation is On or Off if their controlling Resource Control Group is enabled for hibernation. Hibernated instances are restarted by an On operation.
- See AWS hibernation prerequisites in the User Guide for Linux Instances.
- For enabling hibernation in N2WS, see the Hibernation description in section 15.1.
- The stopping and starting of targets identified for each Group is independent of the backup schedule for an Account's policy.
- It is possible to turn off operations for a long period of time even though the Group was never turned on.
- Ad hoc Off and On operations are available in addition to the Resource Control schedule.
- Off/On operations are not allowed for Groups with a **Status** of 'disabled'.

Recommendation: N2WS recommends that you not execute a stop or start operation on critical servers.

Following are Resource Control tabs in the left panel of the N2WS user interface:

- **Resource Control Monitor** – Lists the current operational status of Groups under Resource Control. The **Log** lists the details of the most recent operation for a Group.

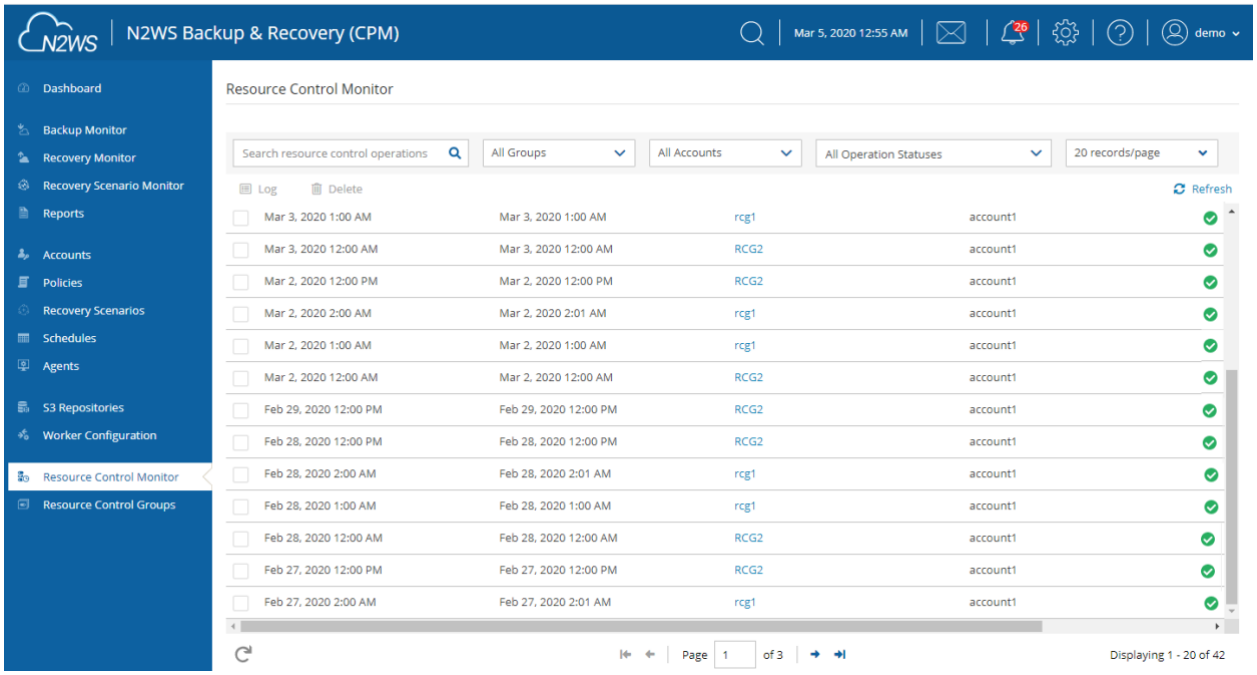

- **Resource Control Groups** – Use the **Groups** tab to add and configure a Group: the account, the days and off/on times, which Resource Targets are subject to the Group control, and other features. You can also delete a group and activate ▷ **Turn On Now** / ☐ **Turn Off Now** controls.

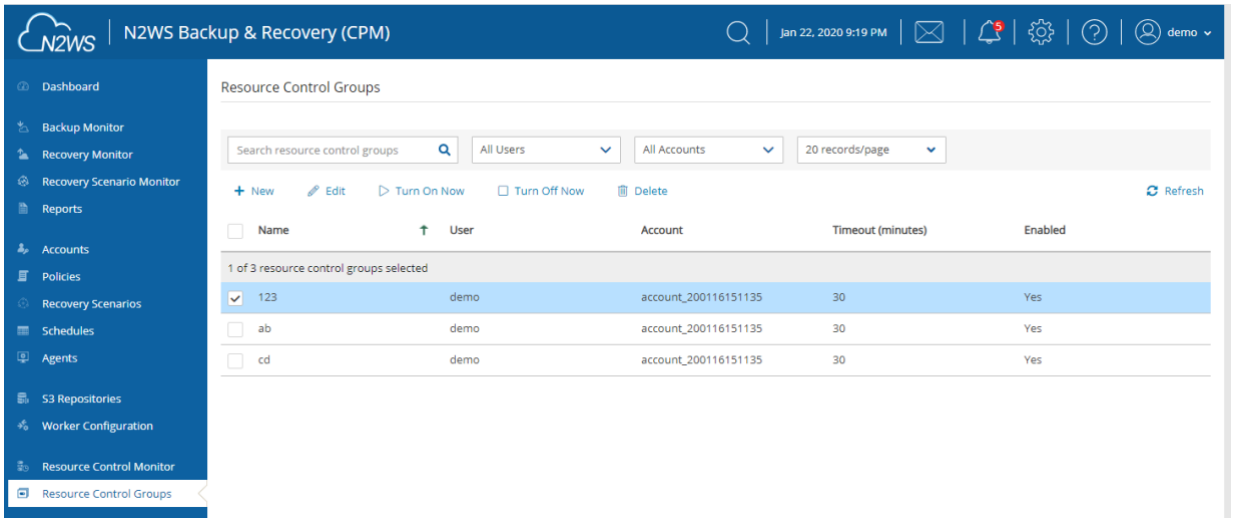

After configuring a group, you can add resources in the **Operation Targets** tab. See section 15.2.

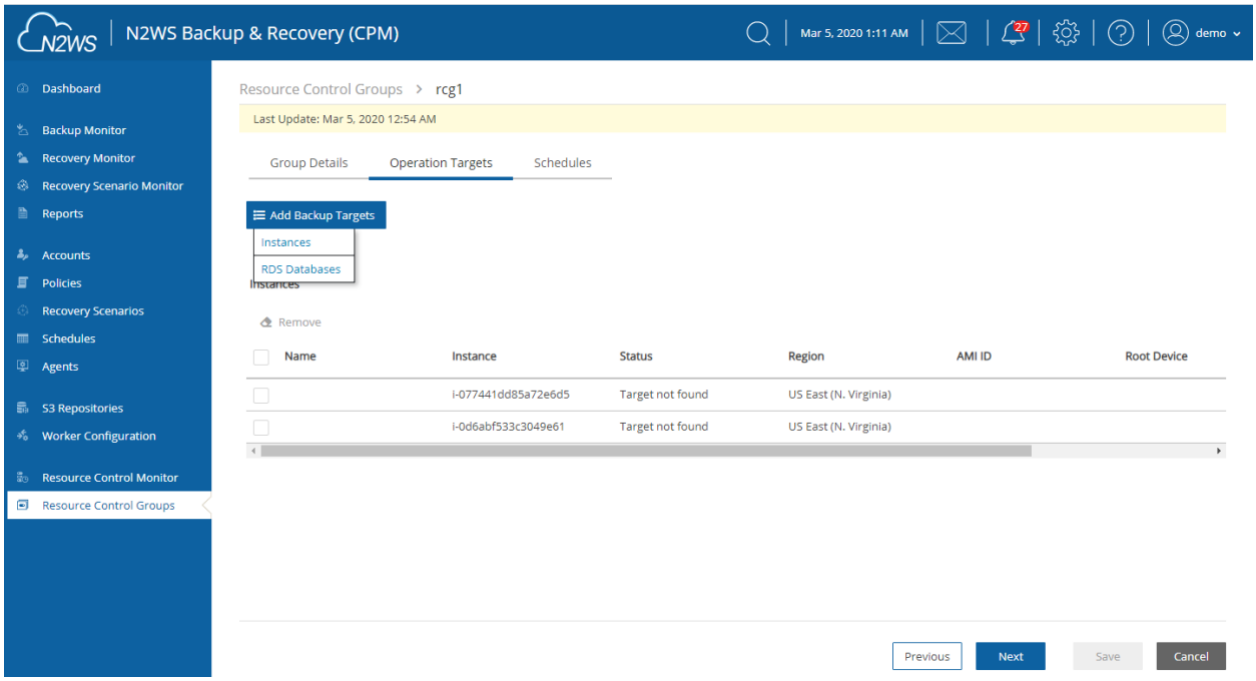

15.1 Adding a Resource Control Group

In the **Resource Control Groups** tab, select **+ New** and complete the **Group Details** screen fields:

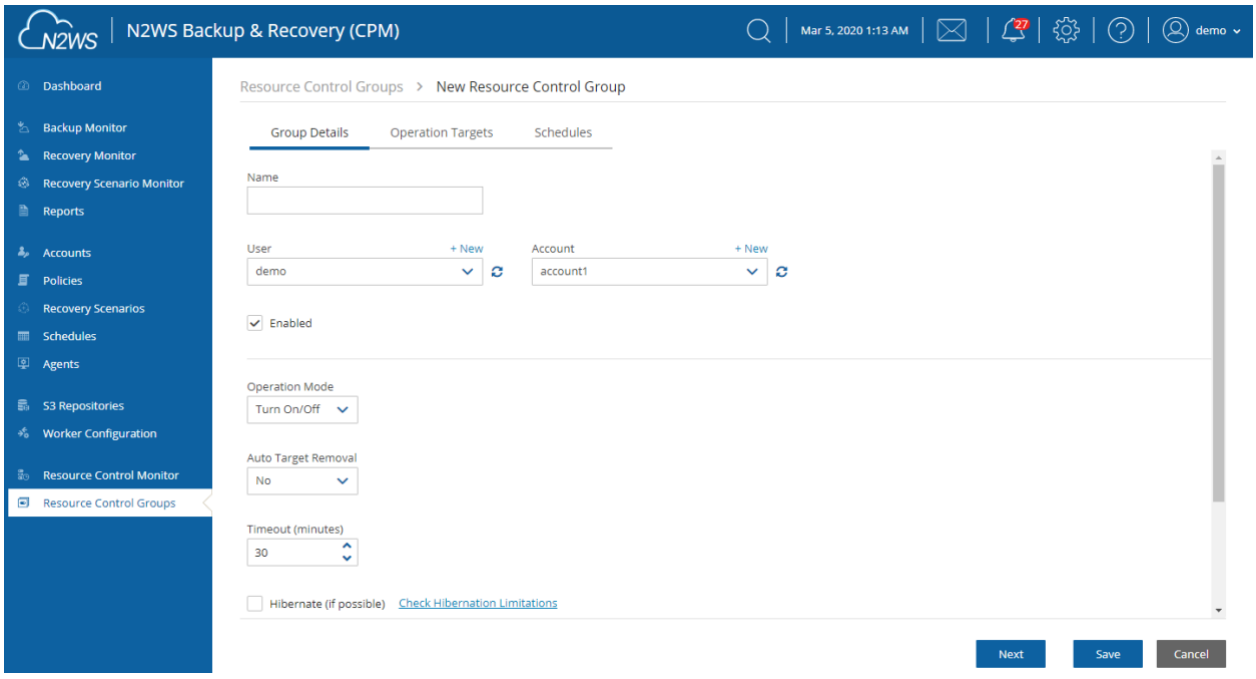

- **Name** –Only alphanumeric characters and the underscore allowed (no spaces).

  Note: A Group may belong to *only* one Account.

- **Account** – Owner of Group. Users are configured for a maximum number of Resource Control entities. See section 18.

- **Enabled** – Whether the Group is enabled to run.

  Note: Off/On operations are not allowed for Groups that are Disabled.

- **Operation Mode** – Two options for controlling operation:
  - **Turn On/Off** – Turn Group on and off according to schedule.
  - **Turn Off Only** – Turn off for an undefined long period of time without having to ever had the Group turned on.
- **Auto Target Removal** – Whether a target resource is automatically removed from the Group if the resource no longer exists in AWS.
- **Timeout (in minutes)** - How long will operation wait in minutes until finished. Default is 30 minutes. Failure from exceeding the timeout does not necessarily mean that the operation of stopping or starting the resource has failed. The Log will show run status for each resource.
- **Hibernation (if possible)** – Whether eligible instances will be hibernated. If enabled, only instances within the Group's target resources that are eligible for hibernation by AWS will be hibernated. See Note on limitations below.
  If an enabled Group contains mixed types of instances, only some of which are eligible for hibernation, then the Off operation will 'hibernate' only the eligible instances, while the remaining instances will 'stop'.

  Note: Select the "Check Hibernation Limitations" link to view current AWS limitations on hibernating instances. During instance creation in AWS, hibernation would have been enabled and encryption configured. If the resource is eligible and the Group is enabled, instances that are 'stopped' move to 'hibernation' state.

- **Description** – Optional description of the Resource Control Group function.

After adding a Group, select **Next** at the bottom of the screen or select the **Operation Targets** tab and configure the **Operation Targets** (section 15.2) and the **Off/On Times** (section 15.3).

15.2 Adding Resource Targets to a Group

Instances and RDS Databases may be added to the Group.

Note: A Resource Target (Instance or RDS Database) may belong to *only* one Group.

- Eligible resources within a Group enabled for hibernation that has been stopped have a **Status** of 'stopped-hibernation'.
- The **Status** column shows whether a target is 'running' or 'stopped'.

Select the **Operation Targets** tab. In the **Add Backup Targets** menu, select a resource type to add to the Group.

Note: It is important to not configure a critical server as part of a Group.

1. If you selected **Instances**, the **Add Instances** screen opens. The following instance types are omitted from **Add Instances** and not allowed to be part of a Group**:**
   - CPM

- Instance-Store type
- Worker - See section 22.

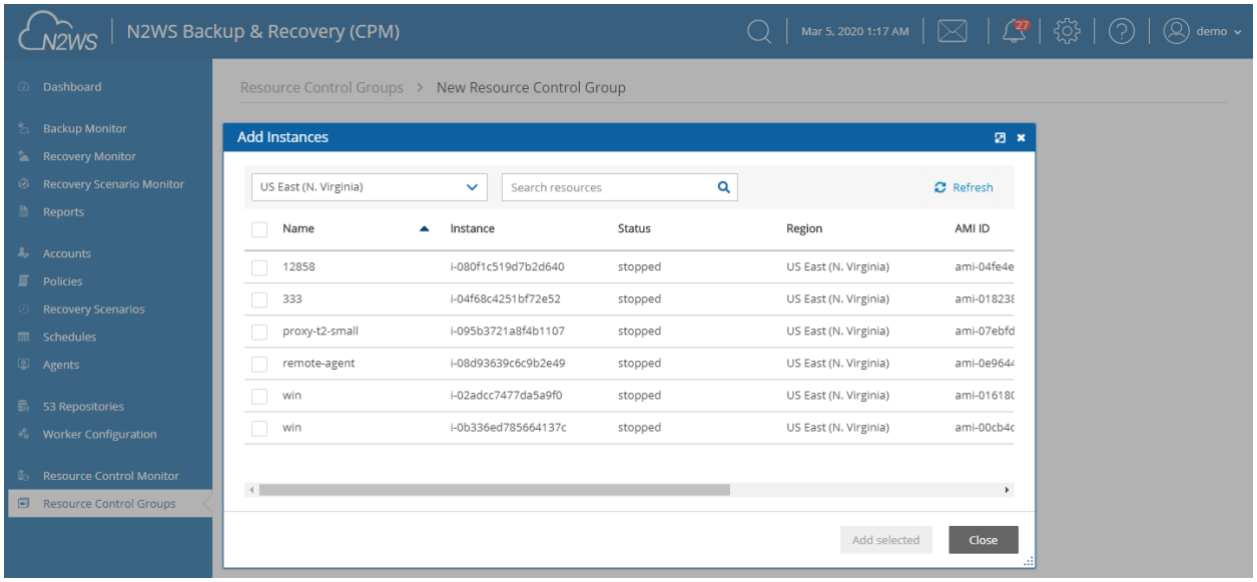

2. If you selected RDS Databases, the following screen opens:

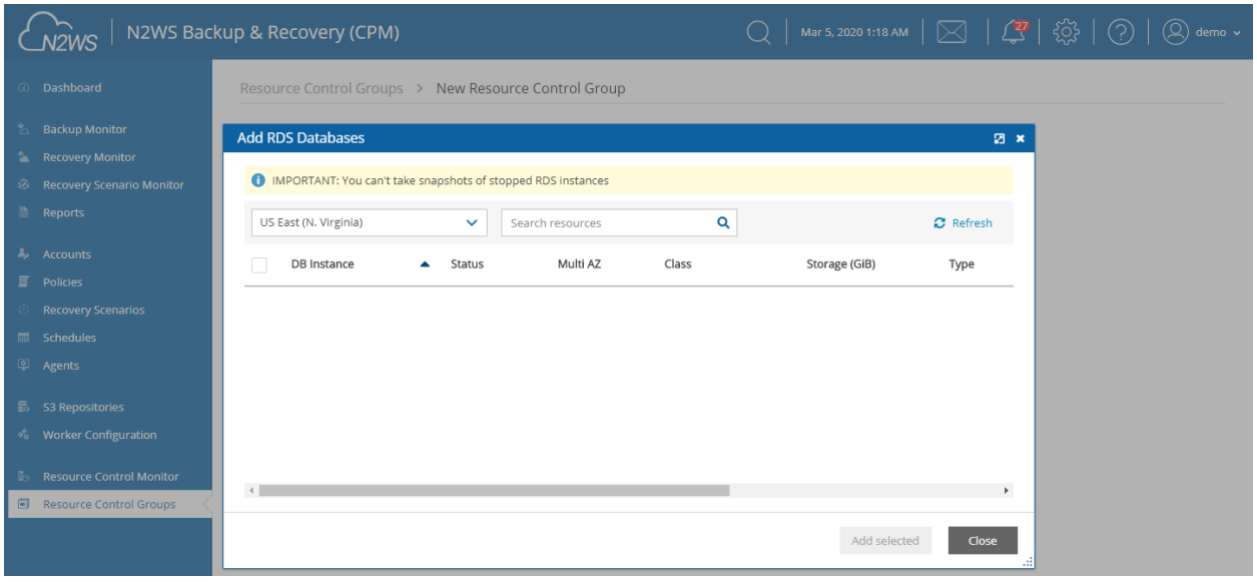

Note: If an RDS database is stopped, a regularly scheduled backup will fail.

3. Check the **Status** column to determine whether a resource is eligible for adding to the Group.
4. Select one or more resources, and then select **Add Selected**. Selected resources are removed from the table.
5. Continue until you are finished and select **Close** to return to the **Operations Targets** screen.
6. Select **Save** to save the **Operation Targets** selections.

15.3 Configuring Off/On Scheduler

- Scheduling overlapping off and on time ranges is invalid. For example:
- A resource is turned off at 20:00 on Wednesday and turned on at 23:00 the same day.

- Then, an attempt to schedule the same resource to be turned off on Wednesday at 9:00 and turned off at 22:00 on Wednesday will result in an invalid input error.

1. Select **Next** to advance to the **Schedules** tab for the group.
2. Select **+ New** to open a default time range row ready for your changes.

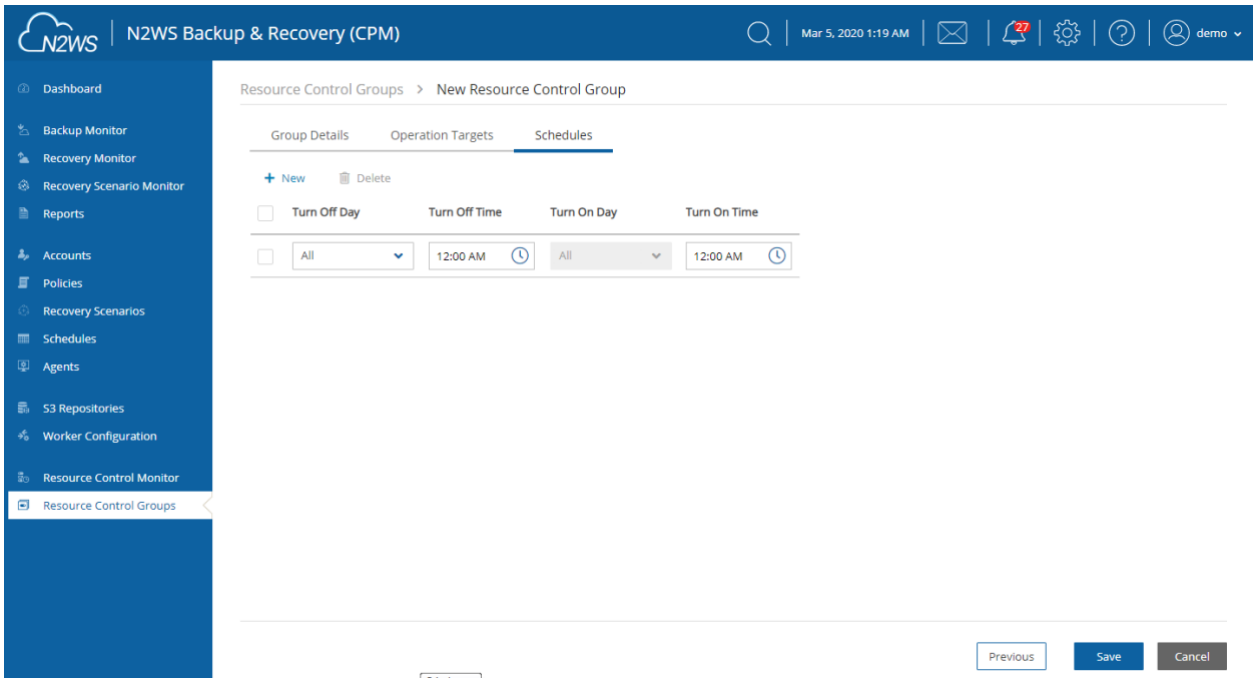

3. In the time range row, select the **Turn Off Day** and **Time** and the **Turn On Day** and **Time** values from the drop-down lists, choosing **AM** or **PM** as required. After each time selection, select **Apply**.

Note: There must be 60 minutes between each operation in order for them to work.

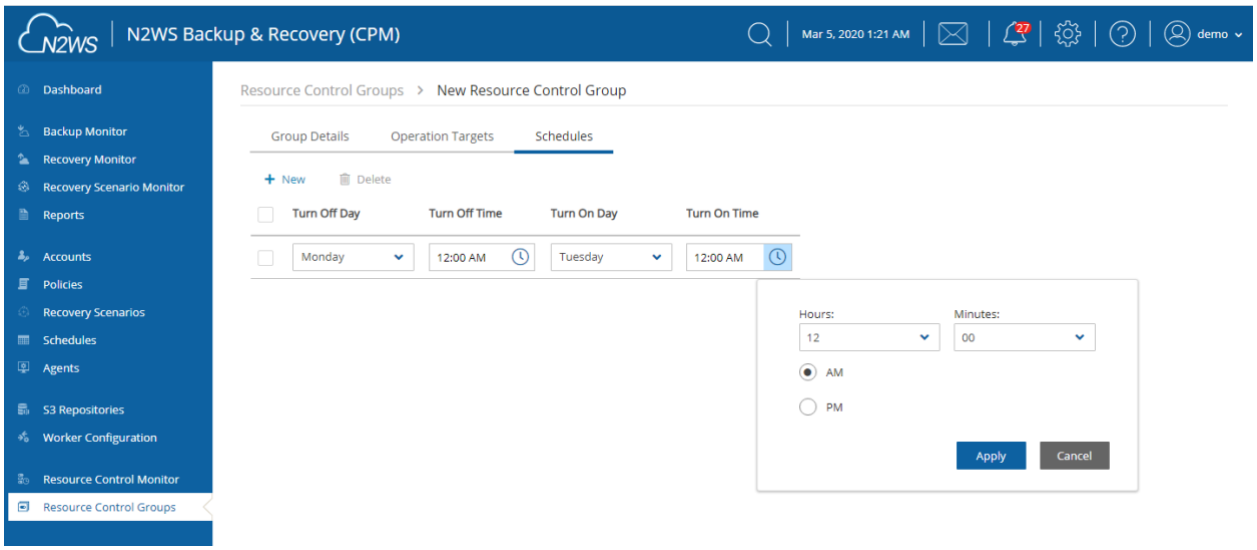

4. Select **+ New** to open another time range row.
5. When finished creating the time ranges, select the required time range rows, and then select **Save**.

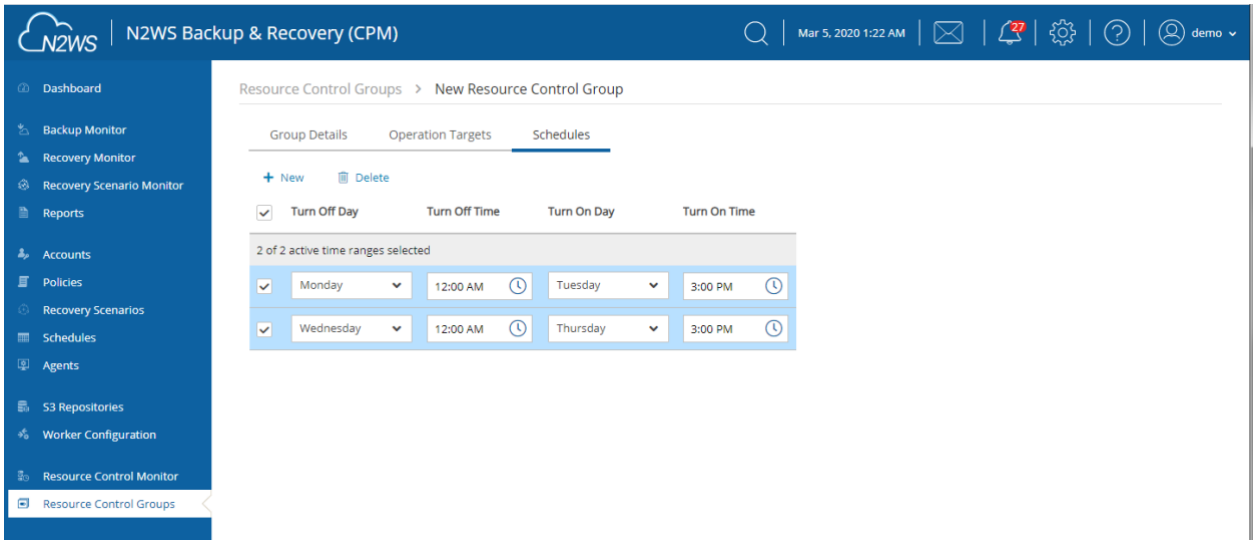

15.4 Overriding a Resource Control Schedule

After creating the Group, you can initiate a stop or start action outside of the scheduled times by selecting the **Turn On Now** or **Turn OFF Now** in the **Resource Control Groups** tab.

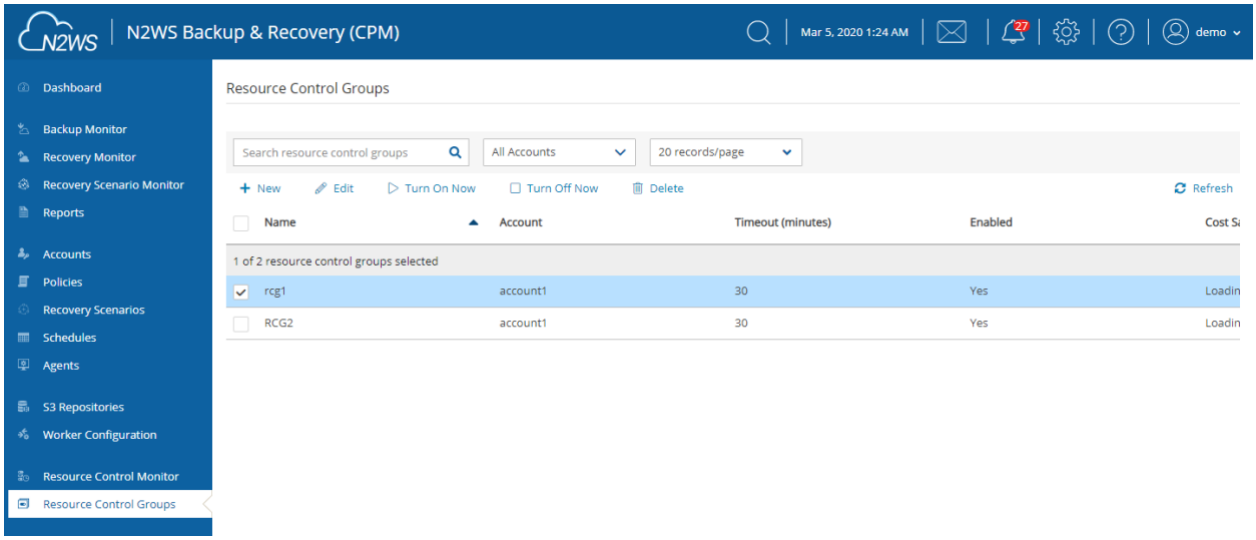

15.5 Using Scan Tags with Resource Control

Scan tags for Resource Control can be used to:

- Create a new Group based on an existing Group's configuration.
- Add a resource to a Group.
- Remove a tagged or untagged resource from a Group.

The tag format is:

Key: **cpm_resource_control**

with one of the following values:

- Value: <group name> or <group name>:<based on group>

- If the value in <group name> equals 'g1', the resource will be added to the g1 group.

- The template <group name>:<based on group> means, in the case of g1:g2:
  If g1 exists, add the resource to g1.
  Otherwise, create a new group g1 based on group g2, and add the resource to it.
  - Value: **no-resource-control** - Remove the resource instance or RDS database from the Group whether it is tagged or not.
  - Value: <no value> - Remove the tagged resource instance or RDS database from the Group.

15.6 Resource Control Reporting

Resource Control provides individual logs of off and on operations and a summary report of all operations.
The individual log contains timestamps for each step within the operation, from firing to completion, and is downloadable as a CSV file. To view individual logs, in the **Resource Control Monitor** tab, select a group and then select **Log**.

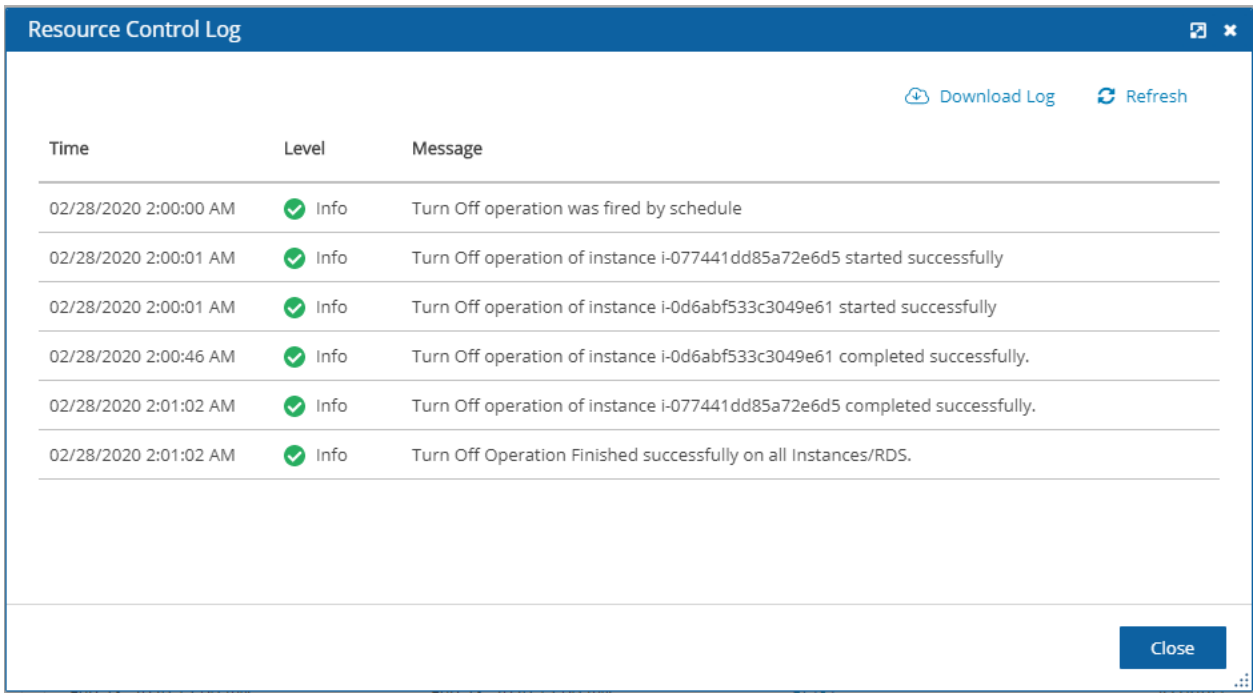

To download the individual log, select **Download Log**.

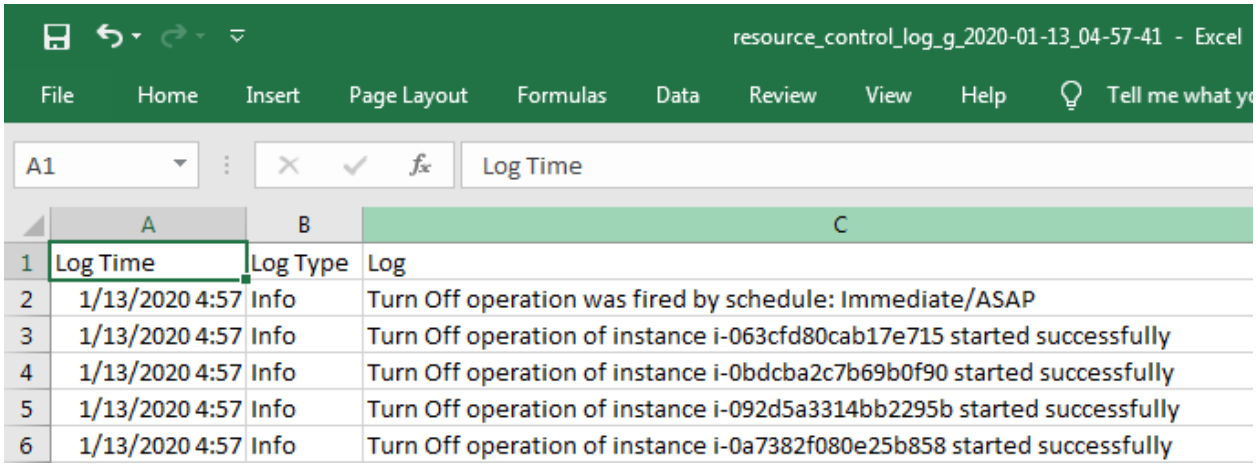

	A	B	C
1	Log Time	Log Type	Log
2	1/13/2020 4:57	Info	Turn Off operation was fired by schedule: Immediate/ASAP
3	1/13/2020 4:57	Info	Turn Off operation of instance i-063cfd80cab17e715 started successfully
4	1/13/2020 4:57	Info	Turn Off operation of instance i-0bdcba2c7b69b0f90 started successfully
5	1/13/2020 4:57	Info	Turn Off operation of instance i-092d5a3314bb2295b started successfully
6	1/13/2020 4:57	Info	Turn Off operation of instance i-0a7382f080e25b858 started successfully

To generate the summary log:

1. Select the **Reports** tab in the left panel.
2. Select the **Immediate Report Generation** tab and then select **Resource Control Operations** in the **Report Type** list.
3. Complete the filter and time range boxes.
4. Select **Generate Report**. The report is automatically downloaded as a CSV file

The Resource Control Operations Report contains information for all saved operations for all accounts. For each operation it contains:

- **Resource Control Operation ID** – A sequential number for each operation.
- **User** – User generating the report.
- **Account** – The N2WS owner of the Resource Control Group.
- **AWS Account Number** – The AWS account number of the owner of the resources.
- **Resource Control Group** – The N2WS Resource Control Group name.
- **Status** – Operation status.
- **Start Time** – Start date and time.
- **End Time** – End date and time.
- **Marked for Deletion** – Whether the resource is marked for deletion.

16 Security Concerns and Best Practices

Security is one of the main issues and barriers in decisions regarding moving business applications and data to the cloud. The basic question is whether the cloud is as secure as keeping your critical applications and data in your own data center. There is probably no one simple answer to this question, as it depends on many factors.
Prominent cloud service providers like Amazon Web Services, are investing a huge amount of resources so people and organizations can answer 'yes' to the question in the previous paragraph. AWS has introduced many features to enhance the security of its cloud. Examples are elaborate authentication and authorization schemes, secure APIs, security groups, IAM, Virtual Private Cloud (VPC), and more.
N2WS strives to be as secure as the cloud it is in. It has many features that provide you with a secure solution.

16.1 N2WS Server

N2WS Server's security features are:

- Since you are the one who launches the N2WS server instance, it belongs to your AWS account. It is protected by security groups you control and define. It can also run in a VPC.
- All the metadata N2WS stores, is stored in an EBS volume belonging to your AWS account. It can only be created, deleted, attached, or detached from within your account.
- You can only communicate with the N2WS server using HTTPS or SSH, both secure protocols, which means that all communication to and from N2WS is encrypted. Also, when connecting to AWS endpoints, N2WS will verify that the SSL server-side certificates are valid.
- Every N2WS has a unique self-signed SSL certificate. It is also possible to use your own SSL certificate.
- AWS account secret keys are saved in an encrypted format in N2WS' database.
- N2WS supports using different AWS credentials for backup and recovery.
- N2WS Server supports IAM Roles. If the N2WS Server instance is assigned an adequate IAM role at launch time, you can use cross-account IAM roles to "assume" roles from the main IAM role of the N2WS instance account to all of the other AWS accounts you manage and not type AWS credentials at all.
- To manage N2WS, you need to authenticate using a username and password.
- N2WS allows creating multiple users to separately manage the backup of different AWS accounts, except in the Free Edition.

16.2 Best Security Practices for N2WS

Implementing all or some of the following best practices depends on your company's needs and regulations. Some of the practices may make the day-to-day work with N2WS a bit cumbersome, so it is your decision whether to implement them or not.

16.2.1 Avoid using AWS Credentials

By using the N2WS Server instance IAM role and cross-account IAM role, you can manage multiple AWS accounts without using AWS credentials (access and secret keys) at all. This is the most secure way to manage multiple AWS accounts and the one recommended by AWS.

16.2.2 Credentials Rotation

Assuming you have to use AWS credentials, you should follow AWS practices. It is recommended that you rotate account credentials from time to time. See http://docs.amazonwebservices.com/AWSSecurityCredentials/1.0/AboutAWSCredentials.html#CredentialRotation
After changing credentials in AWS, you need to update them in N2WS. Select on the account name in the **Accounts** management screen and modify the access and secret keys.

16.2.3 Passwords

Create a strong password for the N2WS server and make sure no one can access it. Change passwords from time to time. N2WS does not enforce any password rules. It is the user's responsibility to create strong passwords.

16.2.4 Security Groups

Since N2WS server is an instance in your account, you can define and configure its security groups. Even though N2WS is a secure product, you can block access from unauthorized addresses:

- You need HTTPS access (original 443 port or your customized port) from:
  - Any machine which will need to open the management application
  - Machines that have N2WS Thin Backup Agent installed on them. See section 6.1.
- You will also need to allow SSH access to create and maintain backup scripts.
- Blocking anyone else will make N2WS server invisible to the world and therefore completely bullet-proof.

Note: The only problem with this approach is that any time you will try to add new backup agents or connect to the management console or SSH from a different IP, you will need to change the settings of the security groups.

16.3 Using IAM

N2WS keeps your AWS credentials safe. However, it is preferable to use IAM roles and not use credentials at all. Additionally, N2WS will not accept root user credentials. To minimize risk, try:

- To provide credentials that are potentially less dangerous if they are compromised, or
- To set IAM roles, which will save you the need of typing in credentials at all.

You can create IAM users/roles and use them in N2WS to:

1. Create a user/role using IAM.
2. Attach a user policy to it.
3. Use the policy generator to give the user custom permissions.

Warning: Using IAM User credentials is not recommended as they are less secure than using IAM roles.

An IAM role can also be used in the N2WS Server (for the account the N2WS Server was launched in) and for instances running N2WS Agent to perform the configuration stage as well as normal operations by combining some of the policies. You can attach more than one IAM policy to any IAM user or role.
The permissions the IAM policy must have depend on what you want to policy to do. For more information about IAM, see IAM documentation: http://aws.amazon.com/documentation/iam/

16.3.1 N2WS Server Configuration Process

AWS credentials in the N2WS configuration process are only used for configuring the new server. However, if you want to use IAM credentials for the N2WS configuration process, or to use the IAM role associated with the N2WS Server instance, its IAM policy should enable N2WS to:

- View volumes instances, tags and security groups
- Create EBS volumes
- Attach EBS volumes to instances
- Create tags

Generally, if you want to use IAM role with the N2WS Server instance, you will need the following policy and the policies for N2WS Server's normal operations, as described in the next section.

Minimal IAM Policy for N2WS Configuration:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "ec2:AttachVolume",
        "ec2:AuthorizeSecurityGroupEgress",
        "ec2:AuthorizeSecurityGroupIngress",
        "ec2:CreateTags",
        "ec2:CreateVolume",
        "ec2:DescribeAvailabilityZones",
        "ec2:DescribeInstanceAttribute",
        "ec2:DescribeInstanceStatus",
        "ec2:DescribeInstances",
        "ec2:DescribeSecurityGroups",
        "ec2:DescribeTags",
        "ec2:DescribeVolumeAttribute",
        "ec2:DescribeVolumeStatus",
        "ec2:DescribeVolumes"
      ],
      "Sid": "Stmt1374233119000",
      "Resource": [
        "*"
      ],
      "Effect": "Allow"
    }
  ]
}
```

16.3.2 N2WS Server IAM Settings

You can use the N2WS Server's IAM role to manage backups of the same AWS account. If you manage multiple AWS accounts, you will still either need to create cross-account roles or enter the credentials for other accounts. If you want to use an IAM user for an account managed by N2WS Server (or the IAM role), you need to decide whether you want to support backup only or recovery as well. There is a substantial difference:

- For backup you only need to manipulate snapshots.
- For recovery you will need to create volumes, create instances and create RDS databases. Plus, you will need to attach and detach volumes and even delete volumes. If your credentials fall into the wrong hands, recovery credentials can be more harmful.
- If you use a backup-only IAM user or role, then you will need to enter ad hoc credentials when you perform a recovery operation.
- Generally, if you want to use the IAM role with the N2WS Server instance, you will need a certain policy, or policies, for N2WS Server's normal operations. For details, see the N2W Software Knowledge Base article on minimal IAM policies at https://support.n2ws.com/portal/kb/articles/what-are-the-required-minimal-aws-permissions-roles-for-cpm-operation

Warning: Using IAM User credentials is not recommended as they are less secure than using IAM roles.

You can check on the permissions required for AWS services and resources, such as backup, RDS, and DynamoDB, and compare them to the policies which cover the requirements. In the **Accounts** tab, select an account and then select **Check AWS Permissions**. To expand a line, select its down arrow ▾.

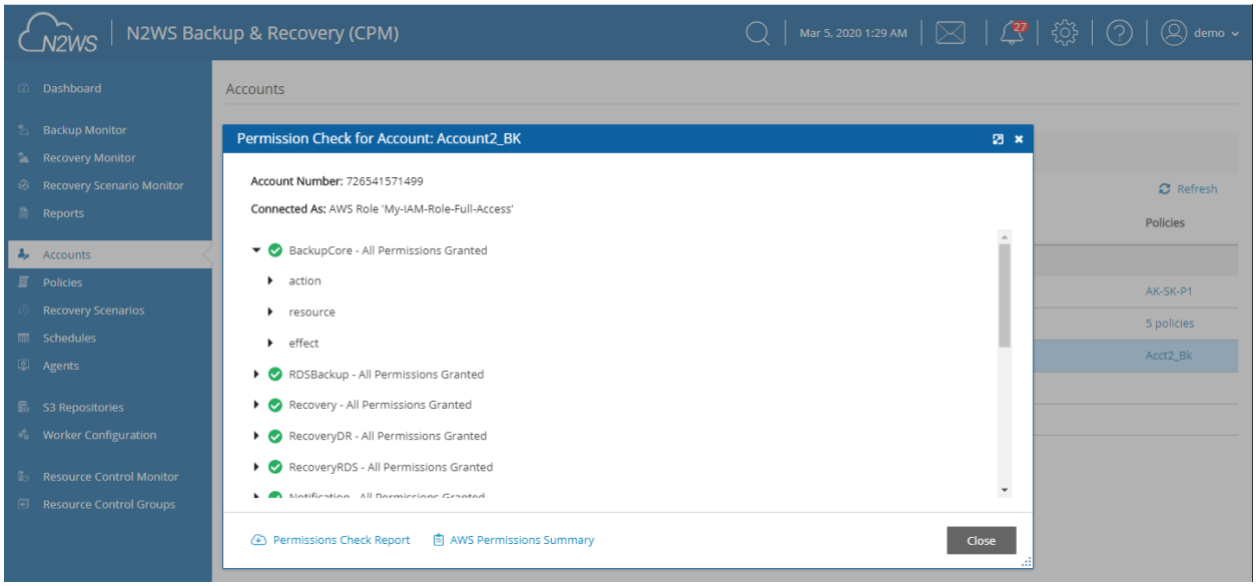

- To download a CVS report, select **Permissions Check Report**.
- To download a JSON file, select **AWS Permissions Summary**.

16.3.3 Configure N2WS' IAM Role with CloudFormation

CloudFormation is an AWS service that allows you to treat a collection of AWS resources as one logical unit. CloudFormation provides a common language for you to describe and provision all the infrastructure resources in your cloud environment, across all regions and accounts in an automated and secure manner.
The IAM role will automatically contain the required permissions for N2WS operations.
See section 20.

17 Alerts, Announcements, Notifications and Reporting

N2WS manages the backup operations of your EC2 servers. In order to notify you when something is wrong and to integrate with your other cloud operations, N2WS allows sending alerts, notifications and even raw reporting data. And when something is not wrong, N2WS can send you an announcement of interest, such as a new feature or money-saving promotion.
So, if you have a network operations center (NOC), are using external monitoring tools or just want an email to be sent to the system administrator whenever a failure occurs, N2WS has an answer for that.

17.1 Alerts

Alerts are notifications about issues in your N2WS backup solution. Whenever a policy fails, in backup or DR, an alert is issued so you will know this policy is not functioning properly. If there are current alerts, the Alerts icon in the toolbar has a number to show you how many there are. Select **Alerts** to open the Alerts list.

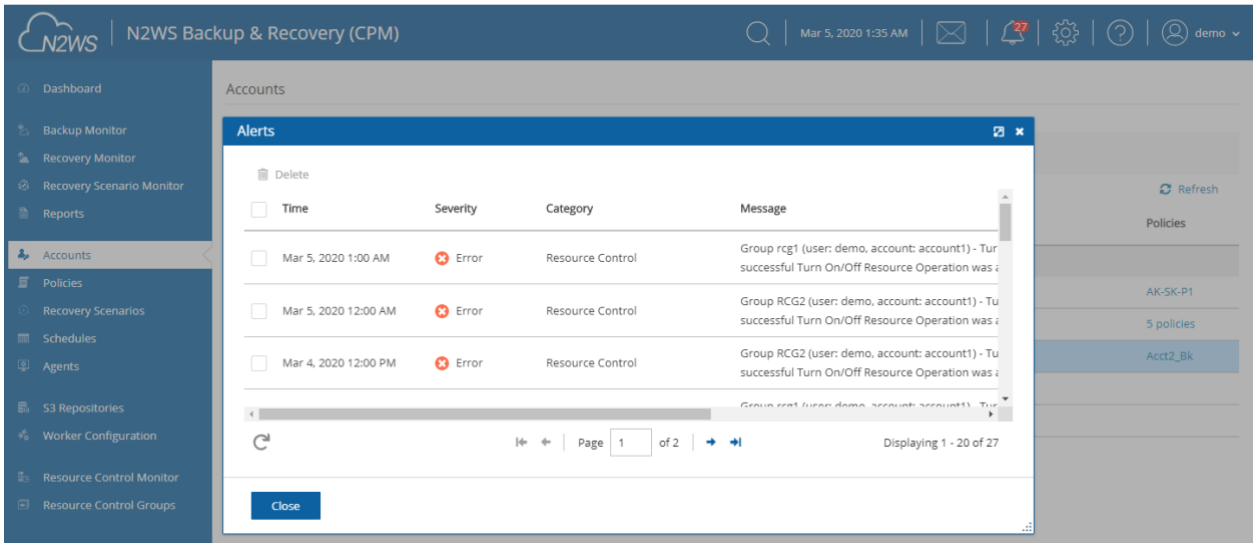

Later, when the policy succeeds, the alert is turned off or deleted, so you will know that the issue is resolved. Alerts can be issued for failures in backup and DR, as well as general system issues like license expiration, for relevant installations.
Depending on the resolution of the output device, a list of Alerts is automatically shown under the Dashboard. The Dashboard list shows the same information except for an abbreviated message and is grouped by functional category, such as Backup and Resource Control.

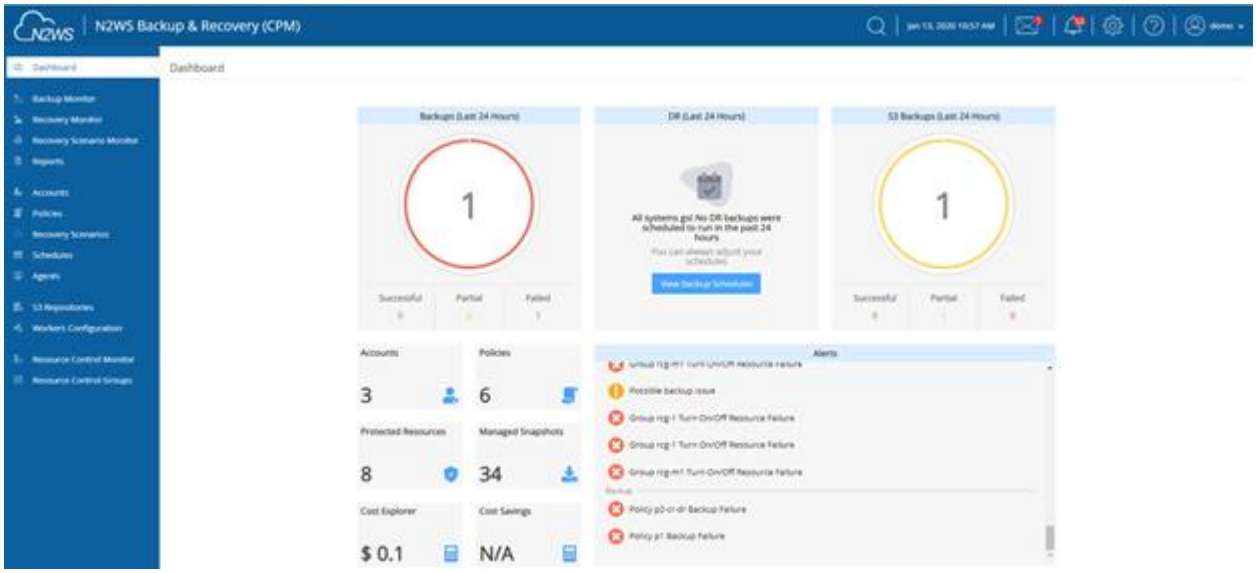

You can manage the number of Alerts shown by selecting alerts to remove in the toolbar Alerts list and then selecting **Delete**.

17.2 Pull Alerts

If you wish to integrate N2WS with 3rd party monitoring solutions, N2WS allows API access to pull alerts out of N2WS. A monitoring solution can call this API to check if N2WS has alerts. When calling this API, the caller receives the current alerts in JSON format. The call is an HTTPS call, and if you configured N2WS server to use an alternate port (not 443), you will need to use that port for this API call as well. N2WS requires an authentication key from the caller. Every N2WS user can define such a key to get the relevant alerts. The root user can also get relevant alerts from other managed users, but not from independent users.

To configure an API call:

1. In the toolbar, select **Settings** in the **User** menu.
2. In the **User Settings** panel, select the **API Access** tab.

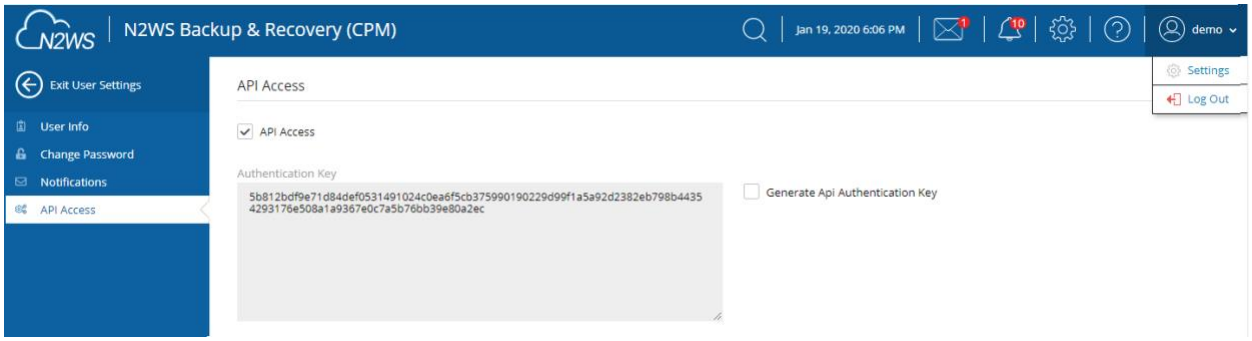

3. To enable access and generate an Authentication Key, select **API Access**.
4. To generate a new **Authentication Key** and invalidate the current, select **Generate API Authentication Key.**
5. After enabling and setting the key, you can use the API call to get all alerts: `https://{{host}}/api/alerts`

A simple example of Python is:

```
d:\tmp>python

Python 2.7.2 (default, Jun 12 2011, 15:08:59) [MSC v.1500 32 bit (Intel)]
on win32

Type "help", "copyright", "credits" or "license" for more information.

>>> import urllib2, json

>>> server_address = 'ec2-54-228-126-14.compute-1.amazonaws.com'

>>> server_port = 443

>>> authkey =
'afb488681baf0132fe190315e87731f883a7dac548c08cf58ba0baddc7006132a

a74f99ab07eff736477dca86b460a4b1a7bfe826e16fdbc'

>>> url = 'https://%s:%d/agentapi/get_cpm_alerts/' % (server_address,
server_port)

>>> url

'https://ec2-54-228-126-14.compute-
1.amazonaws.com:443/agentapi/get_cpm_alerts/'

>>> request = urllib2.Request (url)

>>> request.add_header("Authorization", authkey)

>>> handle = urllib2.urlopen (request)

>>> answer = json.load (handle)

>>> handle.close ()

>>> answer

[{u'category': u'Backup', u'message_body': u'Policy win_server (user:
root, account: main) - backup that started at 07/20/2013 09:00:00 AM
failed. Last successful backup was at 07/20/2013 08:00:00 AM',
u'severity': u'E', u'title': u'Policy win_server Backup Failure',
u'alert_time': u'2013-07-20 06:00:03', u'policy': {u'name':
u'win_server'}}, {u'category': u'Backup', u'message_body': u'Policy
web_servers (user: root, account: main) - backup that started at
07/20/2013 09:20:03 AM failed. Last successful backup was at 07/20/2013
08:30:00 AM', u'severity':u'E', u'title': u'Policy web_servers Backup
Failure', u'alert_time': u'2013-07-20 06:22:12', u'policy': {u'name':
u'web_servers'}}]

>>>
```

The JSON response is a list of alert objects, each containing the following fields:

- category
- title
- message_body
- alert_time (time of the last failure)
- policy
- severity

17.3 Using SNS

N2WS can also push alerts to notify you of any malfunction or issue via SNS. To use it, your account needs to have SNS enabled. SNS can send push requests via email, HTTP/S, SQS, and depending on location, SMS.
With SNS you create a topic, and for each topic there can be multiple subscribers and multiple protocols. Every time a notification is published to a topic, all subscribers get notified. For more information about SNS, see https://aws.amazon.com/sns/.
N2WS can create the SNS topic for you and subscribe the user email defined in the configuration phase. To add subscribers, go to the SNS Dashboard in the AWS Management console), add a recipient, and choose a protocol (SMS, HTTP, etc.), A link to this console is in the N2WS' notifications screen.
For the small volume of SNS messages N2WS uses, there is usually no cost or it is negligible. For SNS pricing see https://aws.amazon.com/sns/pricing/.

17.3.1 Configuring SNS

To configure users for SNS:

1. In the toolbar, select **Settings** in the **User** menu
2. Select the **Notifications** tab in the **User Settings** panel.
3. For each of the boxes, select a value from its list.
4. Depending on the type of credentials selected in the **Authenticate using** box, you may be prompted with additional boxes:
   - **CPM Instance IAM Role** – Requires no additional selections.
   - **Account** – Select the Account name or add a new Account by selecting + **New**.
   - **IAM User Credentials** – Enter the AWS Access and Secret keys.

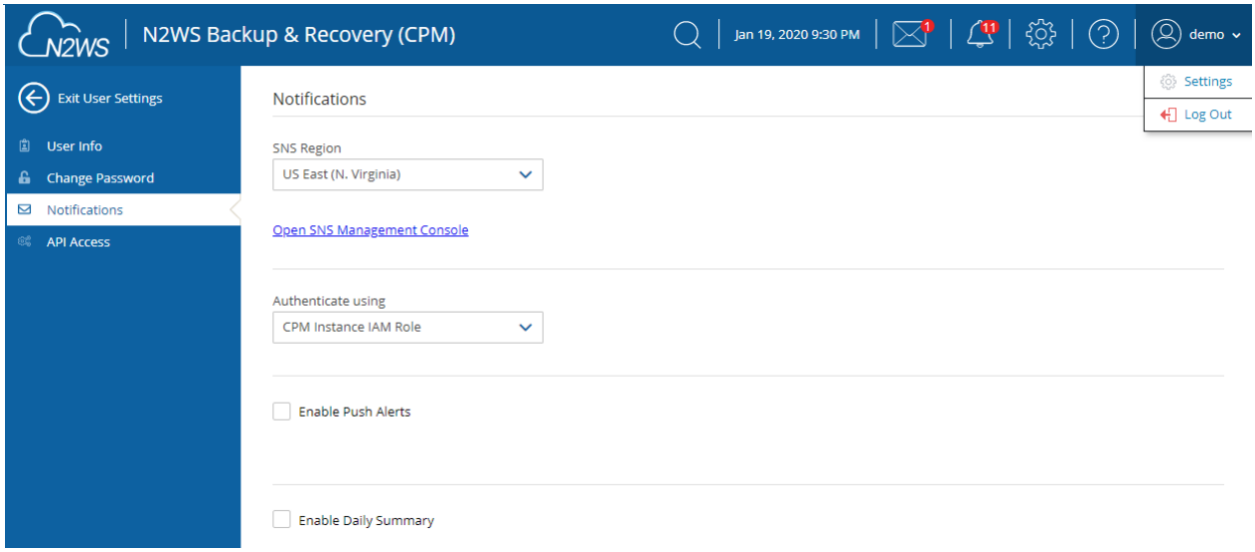

To use SNS:

- You will need to enter AWS account credentials for the SNS service.

- There is one notifications configuration per user, but there can be multiple AWS accounts (where applicable).
- SNS credentials are not tied to any of the backed-up AWS accounts. You can choose a region, and enter credentials, which can be regular credentials, IAM user. See section 16.3. To use the N2WS Server instance's IAM role (only for the root user), type `use_iam_role` for both access and secret keys.
- If you are the root (main) user, you can choose whether to include or exclude alerts about managed users. See section 18.2.
- Root/admin users, and independent users who oversee managed users, can also configure a managed user to receive alerts directly by selecting the user in the **User** list and setting the notification properties described in sections 17.4 and 17.5.
- SNS is used both for push alerts and for sending a daily summary.

17.4 Push Alerts

Push alerts use SNS to send notifications about malfunctions and issues in N2WS' operation.

To enable push alerts:

1. In the **Notifications** tab of the **User Settings** panel, select **Enable Push Alerts**.
2. Define the **Alerts Topic** by selecting one of following options in the list:
   - To create a new topic, select **Auto Generate New Topic**. Optionally, you send the topic to the user by selecting **Add User Email as Recipient**.
   - To use a current topic, select **Use Existing Topic**. Enter the name in the **Alerts Topic Name** box. Or, you can copy the topic's ARN from the **SNS** tab of the AWS Management Console and paste it in the box.
3. To have the user also receive the alert as email, select **Add User Email as Recipient**. The recipient will receive a message requesting subscription confirmation before receiving alerts.

17.5 Daily Summary

The Daily Alert Summary is a message that is sent once a day, summarizing all current alerts, and some policy warnings, in the system. It can be configured instead of, or in addition to, regular alerts. It can be useful for several reasons:

- If you are experiencing issues frequently it sometimes reduces noise to get a daily summary. Furthermore, since backup is the second line of defense, some people feel they do not need to get an instant message on every backup issue that occurs.
- Even if there are no issues, a daily summary is a reminder that all is ok. If something happens and N2WS crashed altogether, and your monitoring solution did not report it, you will notice the Daily Summary will stop.
- The Daily Summary contains a list of policies which are disabled and policies that do not have schedules assigned to them. Although neither is an error, sometimes someone can accidentally leave a policy disabled or without a schedule and not realize that it is not working.

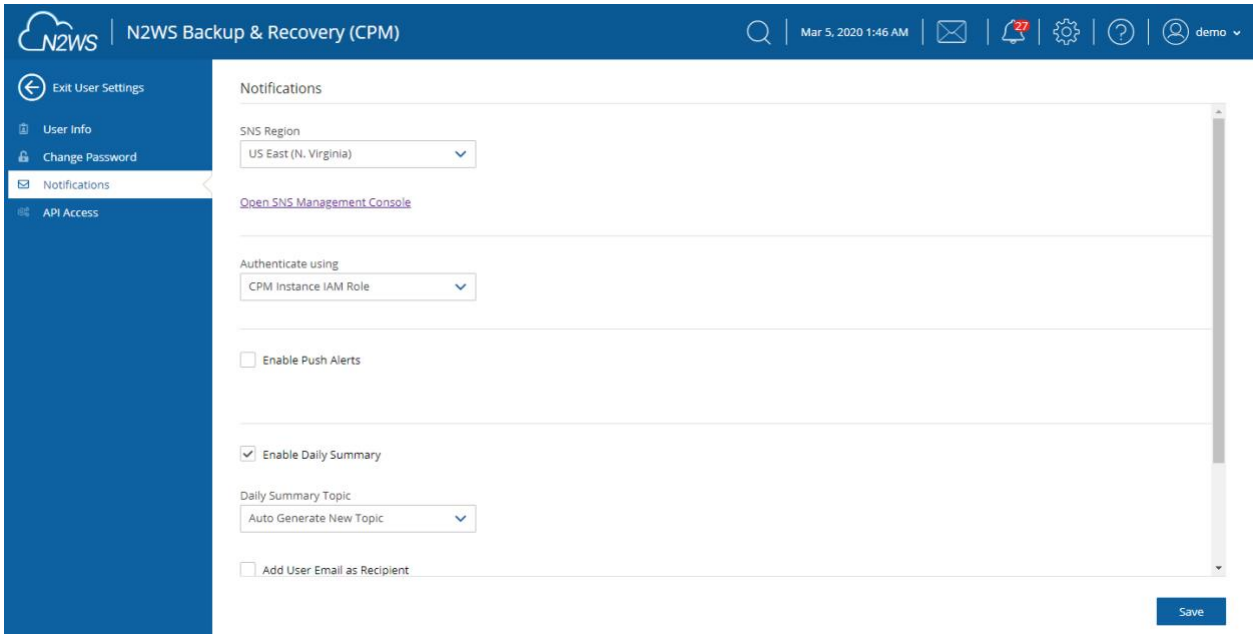

To configure the Daily Summary:

1. In the **Notifications** tab of the **User Settings**, select **Enable Daily Summary**.
2. Define the **Daily Summary Topic** by selecting one of following options in the list:
   - If you want to use the Alert topic for summaries, select **Use Same Topic as Alert**.
   - To create a new topic, select **Auto Generate New Topic**.

   Note: There is an advantage of using a separate topic since sometimes you want different recipients: It makes sense for a system admin to get alerts by SNS and to get the daily summary by email only. The display name of the topic appears in the message and in emails it appears as the sender name. With separate topics, it is easier to distinguish alerts.
3. If you selected **Auto Generate New Topic**, select **Add User Email as Recipient** to have the user also receive the summary as email.
4. In the **Send Daily Summary At** list, select the hour and minutes to send the notification.

17.6 Raw Reporting Data

You can download two raw data reports in CSV format (Comma Separated Values). These reports are for the logged-in user. For the root user, they will include also data of other managed users. These reports include all the records in the database; you can filter or create graphic reports from them by loading them to a spreadsheet or reporting tool. The two reports combined give a complete picture of backups and snapshots taken by N2WS.

To download the CSV reports:

1. In the left panel, select **Reports**.
2. For the **backup view report**, in the **Report Types** list, select **Backup**.

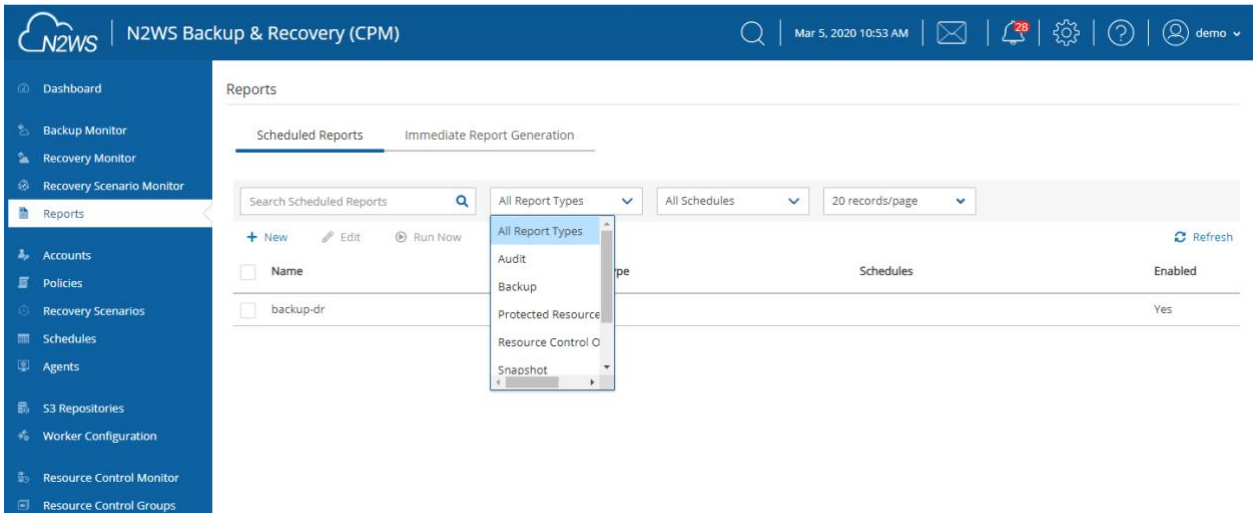

3. For the **snapshot view report**, in the **Report Types** list, select **Snapshot**.
4. For further details on Scheduled Reports, see section 17.9.1.
5. Select **Run Now**.

17.6.1 Backup View CSV Report

This report will have a record for each backup (similar to the Backup Monitor) with details for each of the backups:

- **Backup ID** – A unique numerical ID representing the backup.
- **Account** – Name of the AWS account if the system has multiple users and the user downloading the report is root.
- **AWS Account** Number –ID of the AWS account.
- **Policy** – Name of the policy.
- **Status** – Status of the backup, same is in the Backup Monitor.
- **DR Status** – Status of DR, same as in the Backup Monitor.
- **Start Time** – Time the backup started.
- **End Time** – Time the backup ended.
- **Is Retry** – **Yes** if this backup was a retry after failure, otherwise **no**.
- **Marked for Deletion** – **Yes** if this backup was marked for deletion. If **yes**, the backup no longer appears in the Backup Monitor and is not recoverable.

17.6.2 Snapshot View CSV Report

This report will have a record for each EBS or RDS snapshot in the database:

- **Backup ID** – ID of the backup the snapshot belongs to. Matches the same snapshots in the previous report.
- **Account** – Name of the AWS account.
- **AWS Account Number** – Number of the AWS account
- **Policy** – Name of the policy.
- **Status** – Backup status of success or failure.
- **Region** – AWS region.
- **Type** – Type of snapshot: EBS, RDS or EBS Copy, which is a DR copied snapshot.
- **Volume/DB/Cluster** – AWS ID of the backed up EBS volume, RDS database, or cluster.

- **Volume/DB/Cluster Name** – Name of backed up volume, database, or cluster.
- **Instance** – If this snapshot belongs to a backed up EC2 instance, the value will be the AWS ID of that instance, otherwise it will contain the string: None.
- **Instance Name** – Name of instance.
- **Snapshot ID** – AWS ID of the snapshot.
- **Succeeded** – **Yes** or **No.**
- **Start Time** – Time the snapshot started.
- **End Time** – Time the snapshot ended.
- **Deleted At** – Time of deletion, or N/A, if the snapshot was not deleted yet.

17.6.3 Keeping Records After Deletion

By default, when a backup is marked for deletion, it will be deleted right away from the N2WS database, and therefore not appear in the reports. There are exceptions, such as if N2WS could not delete all the snapshots in a backup (e.g. a snapshot is included in an AMI and cannot be deleted). Sometimes you need to save records for a period of time after they were marked for deletion for compliance, such as General Certificate of Conformity (GCC). To keep records after deletion, see section 9.4.

17.7 Usage Reports

In addition to the raw reports, you can also download CSV usage reports. A usage report for a user will give the number of AWS accounts, instance and non-instance storage this user is using. This can be helpful for inter-user accounting.

1. In the left panel, select the **Reports** tab.
   - For the **usage report (current user),** select **Usage** in the **Report Types** list and your user name in the **Users** list.
   - To get the **usage report (all users)** for the root user, select **Usage** in the **Report Types** list and **All** in the **Users** list.
2. Select **Run Now**.

17.8 Protected and Unprotected Resources Reports

The protected and unprotected resources reports provide information about the AWS resources with and without backup protection.

1. In the left panel, select the **Reports** tab.
   - For the **unprotected/protected resources report (current user),** select **Unprotected or Protected Resources** in the **Report Types** list and your user name in the **Users** list.
   - To get the **unprotected/protected resources report (all users)** for the root user, select **Unprotected or Protected Resources** in the **Report Types** list and **All** in the **Users** list.
2. Select **Run Now**.
3. When you are notified that the report has completed, check your Downloads folder.

AWS resources that are tagged with key:'**cpm backup'**, value:'**no-backup**' will be ignored. Also, see section 14.1.5.

17.8.1 Protected Resources

The protected resources report contains information about the AWS resources with backup policies.

- User Name (on all users reports)
- ID for the resource
- AWS resource name
- Region
- Polices
- Schedules

The protected resources report is available immediately for the current user or all users depending on the account type.
The protected resources report is also available as a Scheduled Report. See section 17.9.

17.8.2 Unprotected Resources

The unprotected resources report contains information about the AWS resources that do not have backup policies.

- Resource Type
- Name of resource
- Resource ID
- Region
- Partial
- Account
- User
- Count of number of unprotected resources per resource type.

17.9 Reports Page

As part of the N2W Software plan of moving toward a robust reporting module, version 3.0.0 has all Reports accessible from the **Reports** tab in the left panel.

The reports will be available in your Downloads folder. Reports are for the logged-in user. For the root user, the reports will also include the data of other managed users.

17.9.1 Scheduled Reports

Scheduled Reports allow you to create a schedule for each report. In order to receive a Scheduled Report, configure at least one recipient email address and the SES service for that email. See section 18.7.

You can run reports outside of a schedule and create ad hoc reports for download:

- In the **Scheduled Reports** tab, **Run Now** generates a defined Scheduled Report and sends emails to its recipients.

- In the **Immediate Report Generation** tab, you can define a new report for immediate download.

See section 17.9.3.

By default, the **Reports** page opens with a list for all reports which have been scheduled. To narrow the list, use the search box, or the filters for report type, user, and schedule.

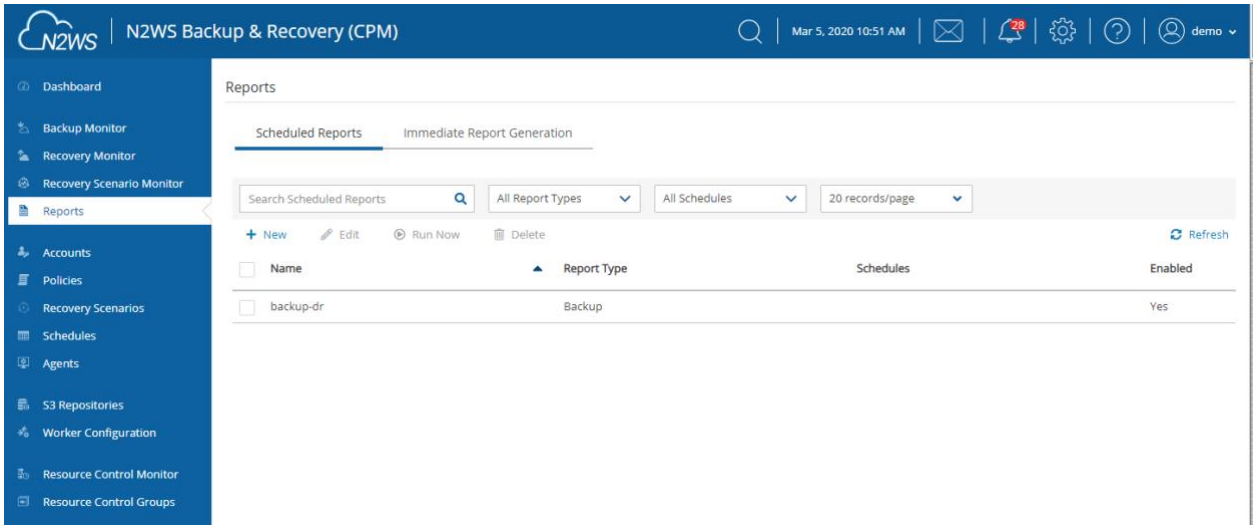

Filters are available based on the chosen Report Type. Depending on the report, you can filter the results as follows:

- **Audit** – Filter for User and From/To Date/Time
- **Backups** – Filter for Account and From/To Date/Time
- **Protected Resources Report** – Filter for User and Account
- **Resource Control Operation Report** – Filter for Account and From/To Date/Time
- **Snapshots** - Filter for Account and From/To Date/Time
- **Usage** – Filter for User and From/To Date/Time – The default is a summary report; select **Detailed** for a detailed report.
- **Unprotected Resources** - AWS resources that do not have backup policies.

17.9.2 Defining a Scheduled Report

Reports are run according to their defined schedule and immediately using **Run Now**. Schedules reports must include at least one email recipient.

To create a scheduled report:

1. Select the **Scheduled Reports** tab and then **New**.

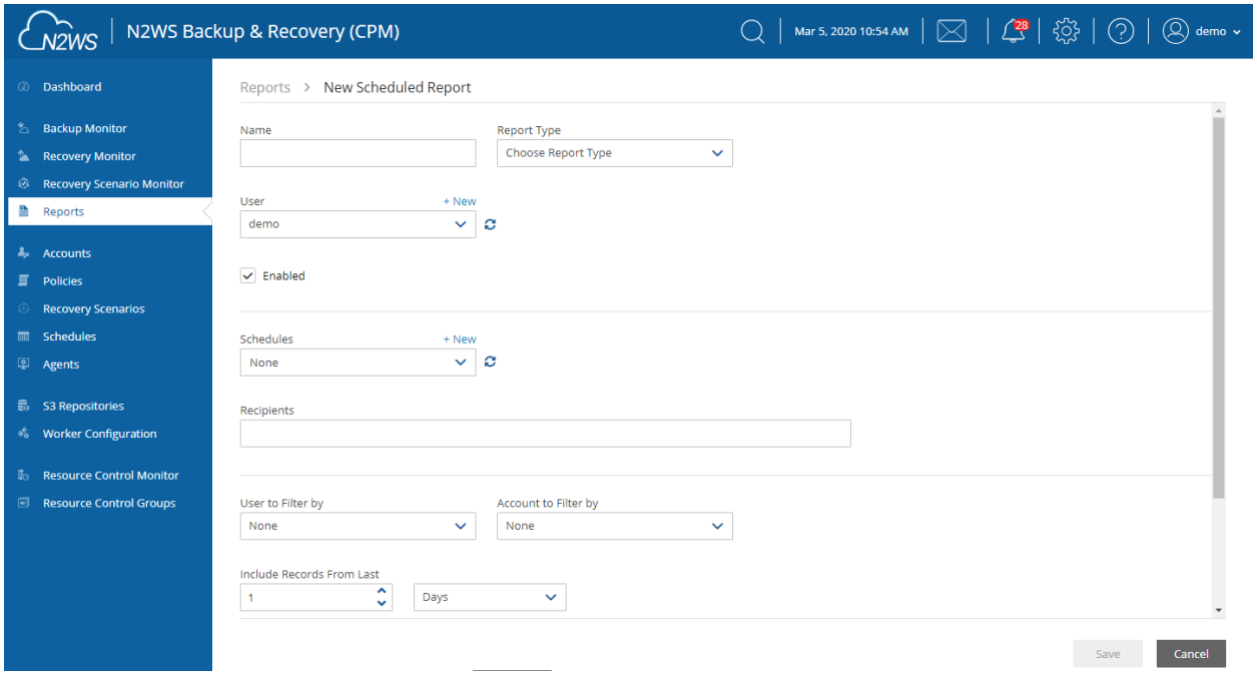

4. Enter a name for the new report and choose the **Report Type**.
5. By default, the report is enabled. To disable the Schedule Report, clear **Enabled**.
6. In the **Schedules** list, select one or more schedules. To create or edit a schedule, see section 4.1.1.

Note:	You can create a Scheduled Report without a schedule and edit the report later after creating the schedule.

7. In the **Recipients** box, enter the email address of recipients, separated by a semi-colon (';').
8. Select from the filters presented for the **Report Type**.
9. In the **Include Records From Last** boxes, you can select the number (first list) of **Days**, **Weeks**, or **Months** (last list) to include in the report. The default is all available records.
10. In the **Description** box, enter an optional informative description.
11. Select **Save**.

17.9.3 Running Reports Outside Their Schedule

To run a Scheduled Report and send emails to its recipients immediately:

In the **Scheduled Reports** tab, select the report in the list and then select **Run Now**.

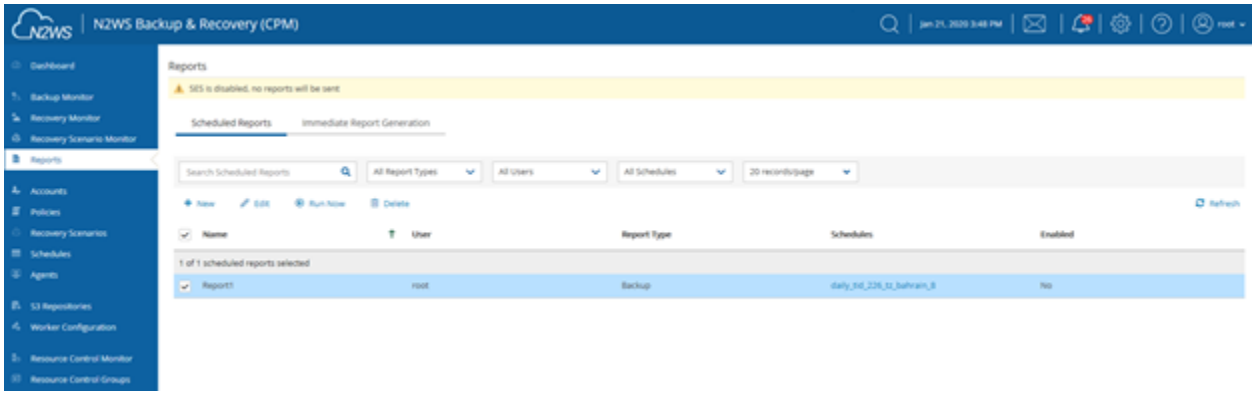

To define a new report and download it immediately:

1. Select the **Immediate Report Generation** tab.
2. Select a Report **Type** and one or more filters depending on the **Type** selected, as listed above in section 17.9.1.

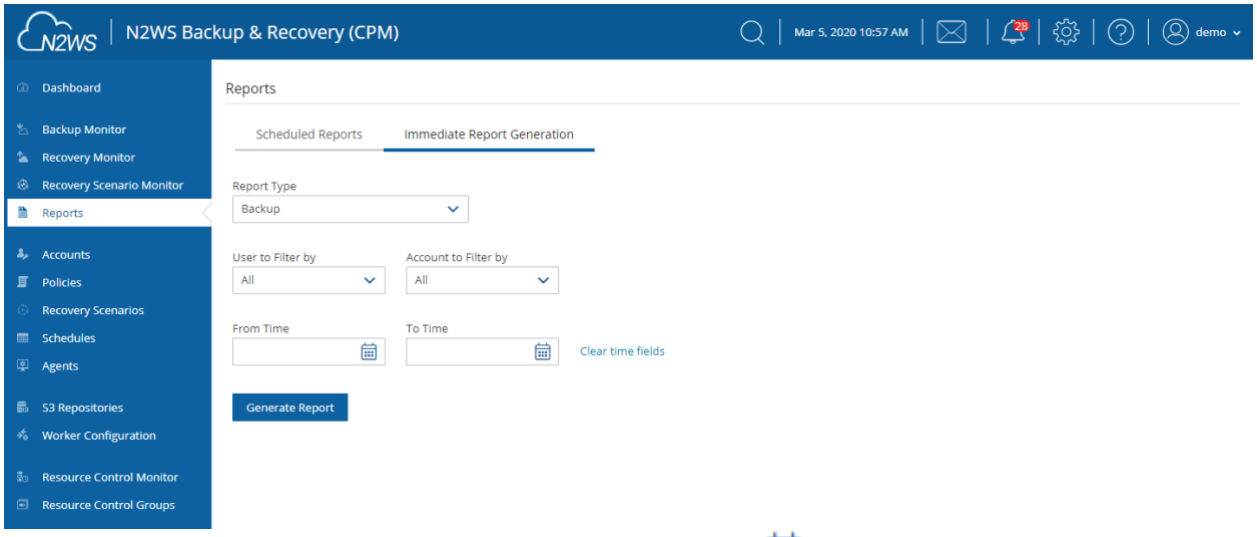

3. To filter the report data by date and time, select Calendar and choose the From and To date and time values. Select **Apply** after each definition.

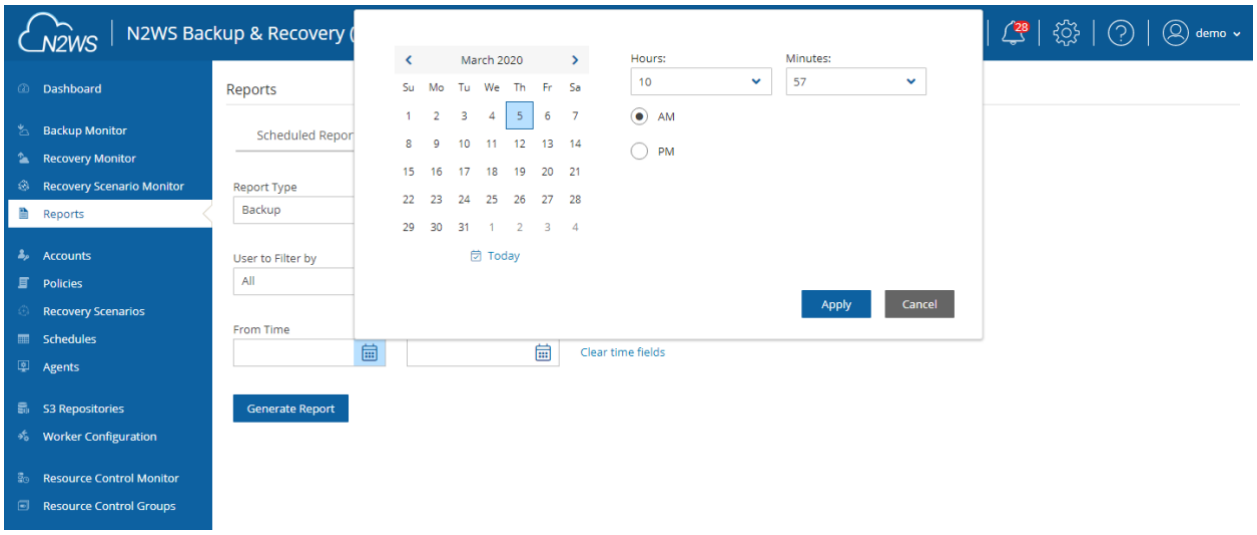

4. Select **Generate Report**.

   The output will be downloaded by your browser.

17.10 Examples of AWS Alerts

AWS uses SNS to provide a number of N2WS alert services by subscription.

17.10.1 Subscription Confirmation Alert

After subscribing to CPM Alerts in AWS, you will receive an email with a confirmation link:

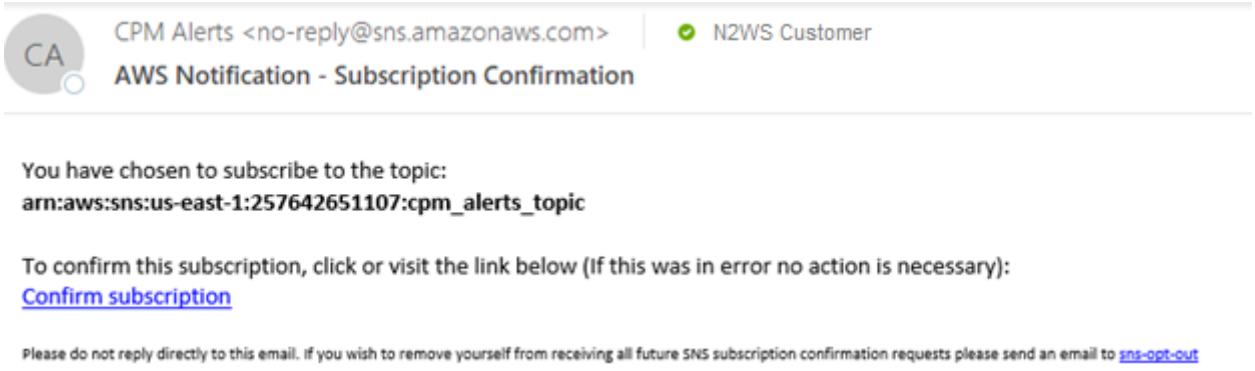

Select the **Confirm subscription** link. You will receive a subscription confirmation email:

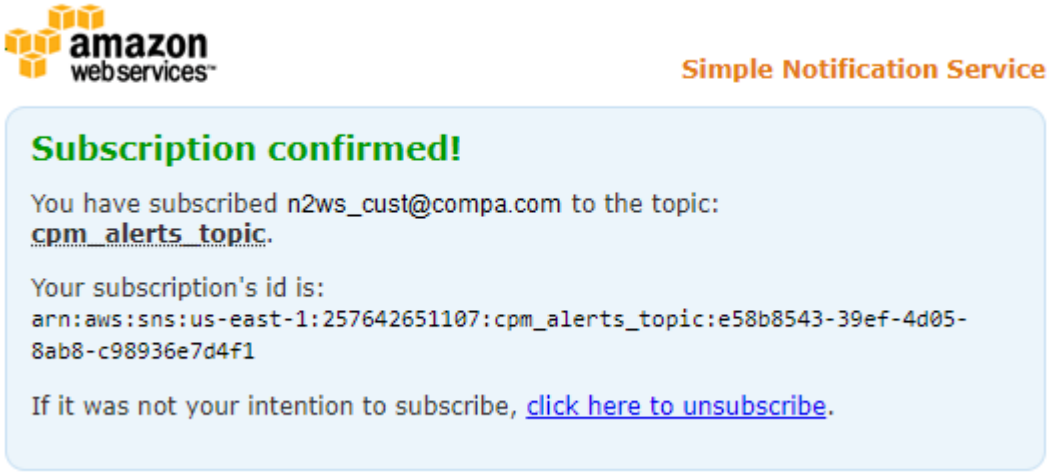

17.10.2 Daily Summary Alert

Following is an example of a CPM Daily Summary where all AWS functions were OK:

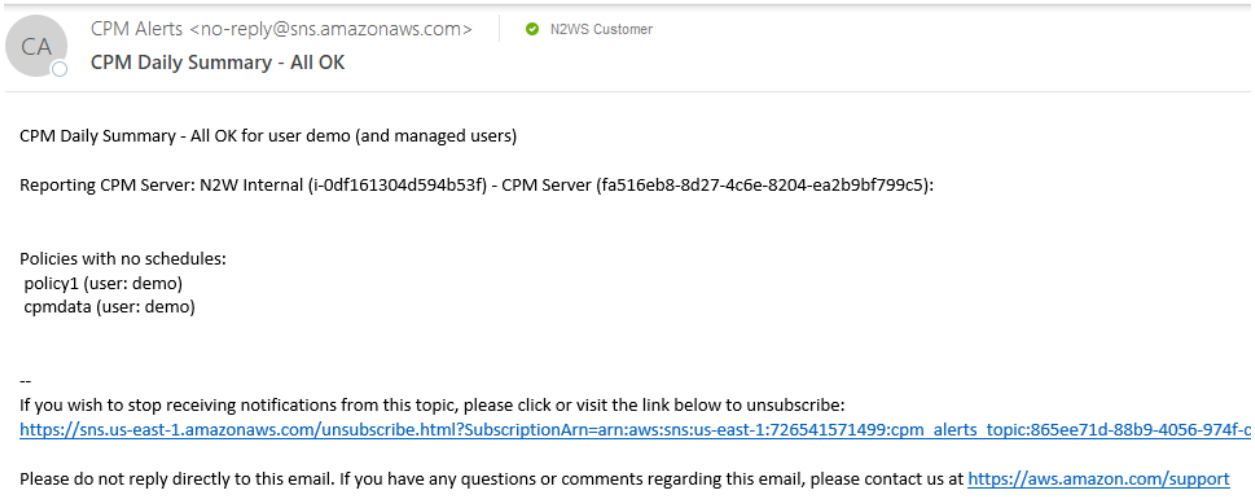

17.10.3 Unprotected Resources Alert

Following is an example of an alert that the unprotected resources report is available:

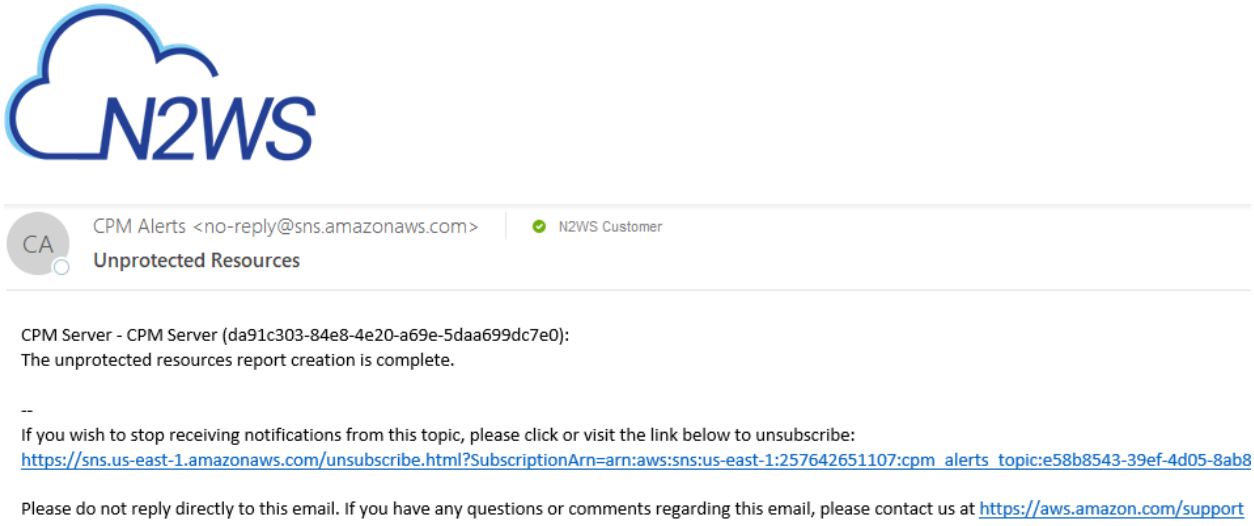

17.11 Announcements

In version 3.0.0, Announcements were introduced as a method for N2WS to communicate directly to users about non-operational topics, such as promotions and other sales-related information.

In the Toolbar, **Announcements** shows the number of unread announcements waiting in the user's mail box. In the **Announcements** inbox panel, select an announcement to open.

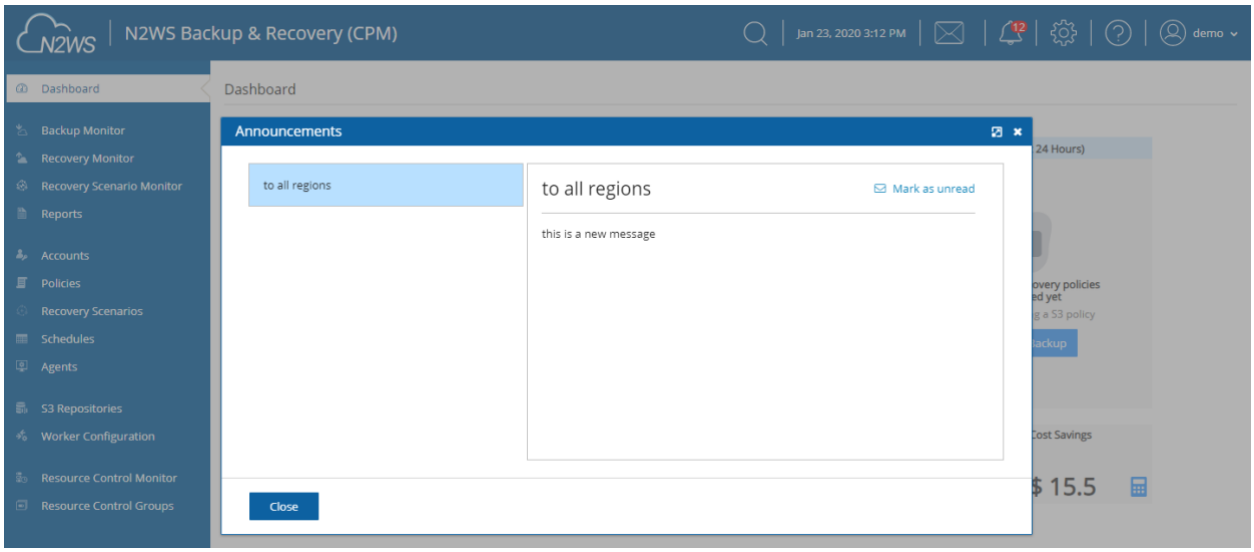

After selecting an announcement, you can reset the message status by selecting **Mark as Unread** in the upper right corner of the message.

18 N2WS User Management

N2WS is built for a multi-user environment. At the configuration stage, you define a user that is the root user. The root user can create additional users (depending on the edition of N2WS you are subscribed to). Additional users are helpful if you are a managed service provider, in need of managing multiple customers from one N2WS server or if you have different users or departments in your organization, each managing their own AWS resources. For instance, you may have a QA department, a Development Department and IT department, each with their own AWS accounts. Select **Server Settings** > **Users**.

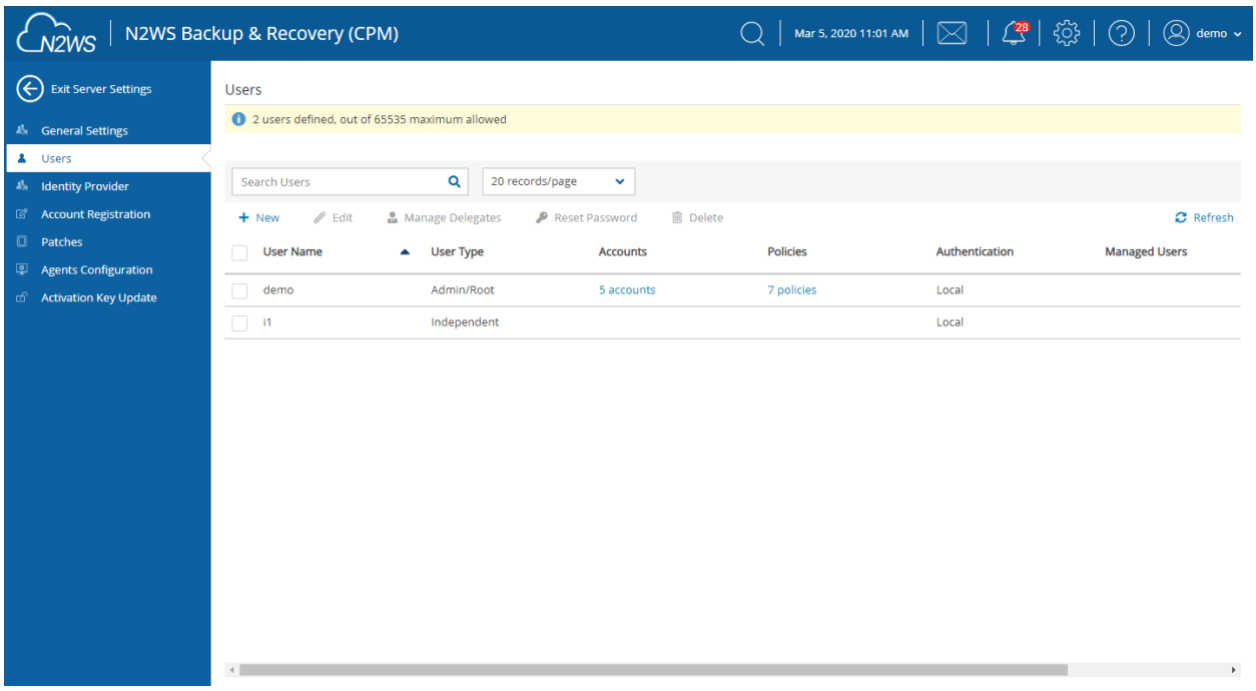

Following are the types of users you can define. Delegate users are defined after users are created.

- Independent
- Managed

18.1 Independent Users

Independent users are completely separate users. The root user can create such a user, reset its password, and delete it with all its data, but it does not manage what this user's policies and resources. Independent users can:

- Log-in to N2WS
- Create their own accounts
- Manage their backup
- Mange policies and resources of managed users that were assigned to them

Independent users can have Managed users assigned to them by the root/admin in the **Users** management screen. An Independent user can log on, manage the backup environment of their assigned Managed users, and receive alerts and notifications on their behalf.

18.2 Managed Users

Managed Users are users who can log on and manage their backup environment, or the root/admin user or independent user, can do it for them. The root user can perform all operations for managed users: add, remove and edit accounts, manage backup policies, view backups and perform recovery. Furthermore, the root user, or independent user, can receive alerts and notifications on behalf of managed users. The root/admin user can also configure notifications for any managed user and independent users can configure notifications for their managed users (section 17.3.1.) To create a managed user, select + **New** and choose **Managed** as the **User Type**. If the root user does not want managed users to login at all, they should not receive any credentials.
Managed users may be managed by Independent users. See section 18.1.

18.3 User definitions

When editing a user, the root user can modify email, password, type of user, and resource limitations.

Note:	The user name cannot be modified once a user is created.
Note:	Users who are created in N2WS via IdP integration (section 19) cannot be edited, only deleted.

To define users:

1. If you are the root or admin user, in the toolbar, select Server Settings.
2. In the left panel, select the **Users** tab. The **Users** screen opens.
3. Select + **New**.

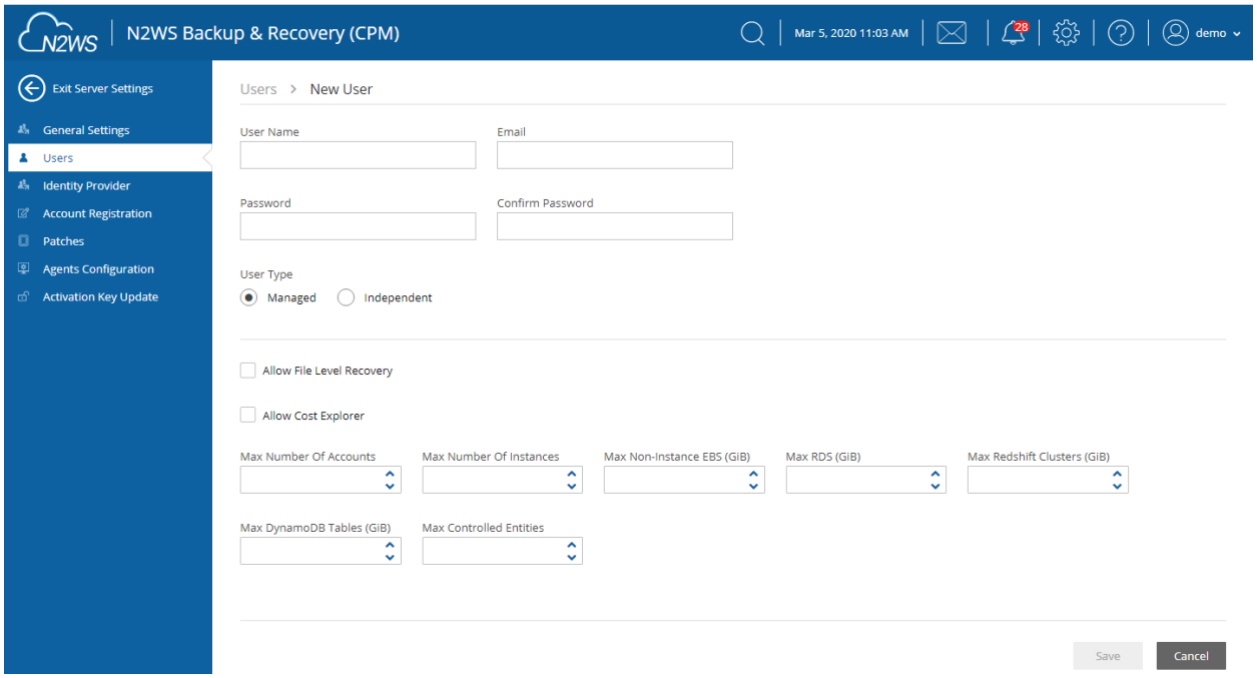

4. In the **User name**, **Email** and **Password** boxes, type the relevant information.
5. Select the **User Type** option. For type details, see sections 18.1 and 18.2.
6. If the user can recover at file level, select **Allow File Level Recovery**.
7. To enable Cost Explorer calculations:
   - Select **Allow Cost Explorer**. The default is to deny the calculations.
   - In AWS, allow the CPM Cost Explorer feature. See section 18.3.1.
   - For information about Cost Explorer, see section 25.
8. In the **Max Number of Accounts**, **Max Number of Instances, Max Non-instance EBS (GiB)**, **Max RDS (GiB), Max Redshift Clusters, Max DynamoDB Tables (GiB**), and **Max Controlled Entities** boxes, select the value for the respective resource limitation from its list.

   The value for **Max Controlled Entities** is the maximum number of allowed instances and RDS database resources.

Note:	If you leave these resource limitation fields empty, there is no limitation on resources, except the system level limitations that are derived from the licensed N2WS edition used.

18.3.1 Configuring AWS to Allow CPM Cost Explorer Calculations

In order to allow CPM Cost Explorer calculations in AWS, users must add cost allocation tags ***once***.

To activate user cost allocation tags:

1. Sign in to the AWS Management Console and open the Billing and Cost Management console at https://console.aws.amazon.com/billing/home#/.
2. In the navigation pane, choose Cost Allocation Tags.
3. Select the tags to activate:
   - **cpm_server_id**
   - **cpm_policy_name**

4. Choose **Activate.**

See http://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/activating-tags.html

18.4 Delegates

Delegates are a special kind of user, which is managed via a separate screen. Delegates are similar to IAM users in AWS:

- They have credentials used to log on and access another user's environment.
- The access is given with specific permissions.

Warning: Using IAM User credentials is not recommended as they are less secure than using IAM roles.

For each user, whether it is the root user, an independent user or a managed user, the **Manage Delegates** command in the **Users** list screen that opens the Delegates screen for that user. Selecting an existing entry in the Delegates column also opens the Delegates screen for that user

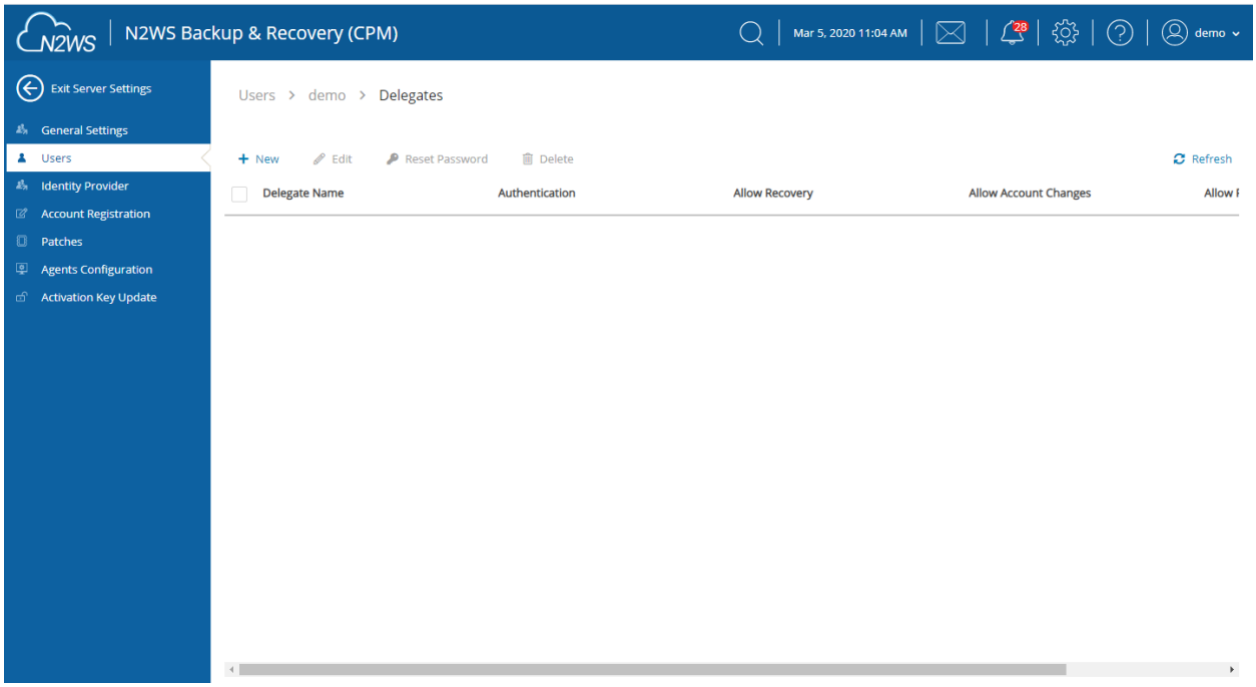

You can add as many delegates as needed for each user and also edit any delegate's settings:

To add a new delegate:

Note: Once a user is defined as a delegate, the name cannot be changed.

1. Select a user.
2. Select **Manage Delegates** and then select **New**.

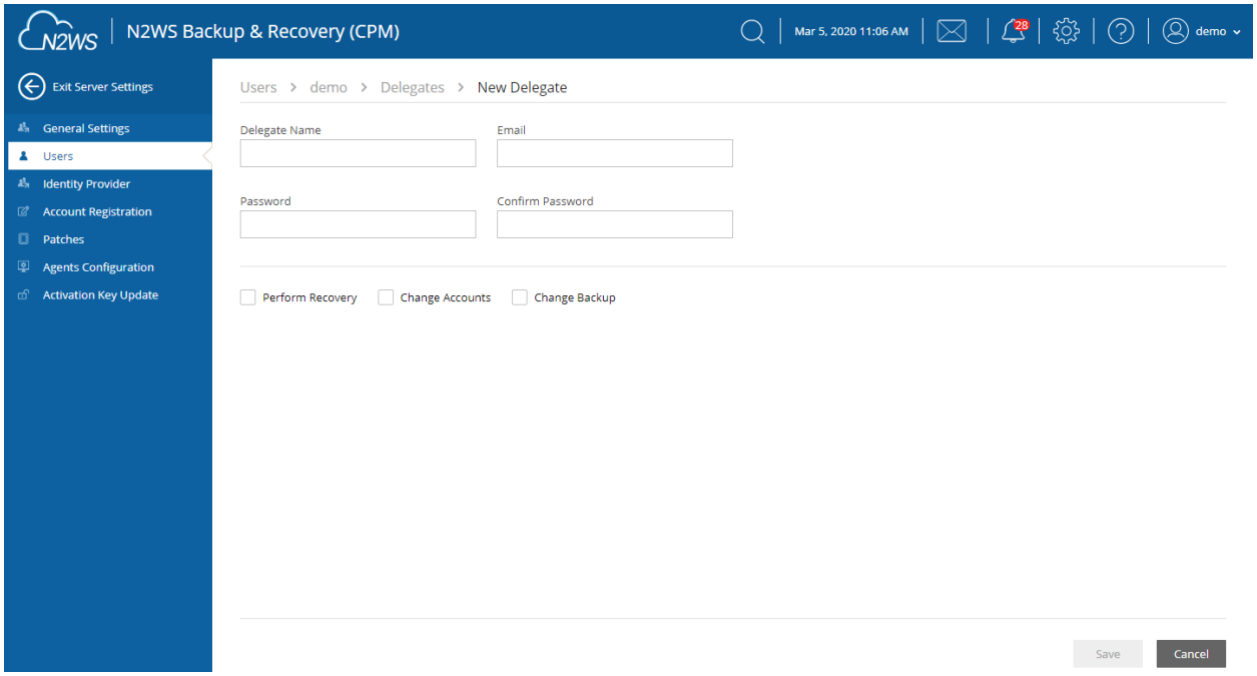

3. In the **Delegate Name** box, type the name of the new delegate.
4. Enter a valid **Email** and set the **Password**.
5. Permissions are denied by default. To allow permissions, select the relevant ones for this delegate:
   - **Perform Recovery** – Can perform recovery operations
   - **Change Accounts** – Can add and remove AWS accounts as well as edit accounts and modify credentials
   - **Change Backup** - Can change policies: adding, removing and editing policies and schedules, as well as adding and removing backup targets

   By default, all are denied, which means that the delegate will only have permissions to view the settings and environment and to monitor backups.

- Allowing all permissions will allow the delegate the permissions of the original user except for notification settings.
- For delegates of the root/admin user, they will not be able to change notification settings, General Settings, or manage users.

When in Edit mode, the root user can reset passwords for delegates.

18.5 Usage Reports

The root user can also use the user management screen to download CSV usage reports for each user, which can be used for accounting and billing. The usage report will state how many accounts this user is managing, and for each account, how many instances and non-instance storage is backed up.
Reporting is now available for daily tracking of resources that were configured as a backup target on each policy. The **Reports** tab contains two levels of detail for Usage Reports. Users can download the following Usage Reports, both of which are filterable by user and timeframe. The report can be created as a **Scheduled Report** or for **Immediate Report Generation**. In each

case, select **Detailed** for usage per account or **Anonymized** for aggregated account usage per user.

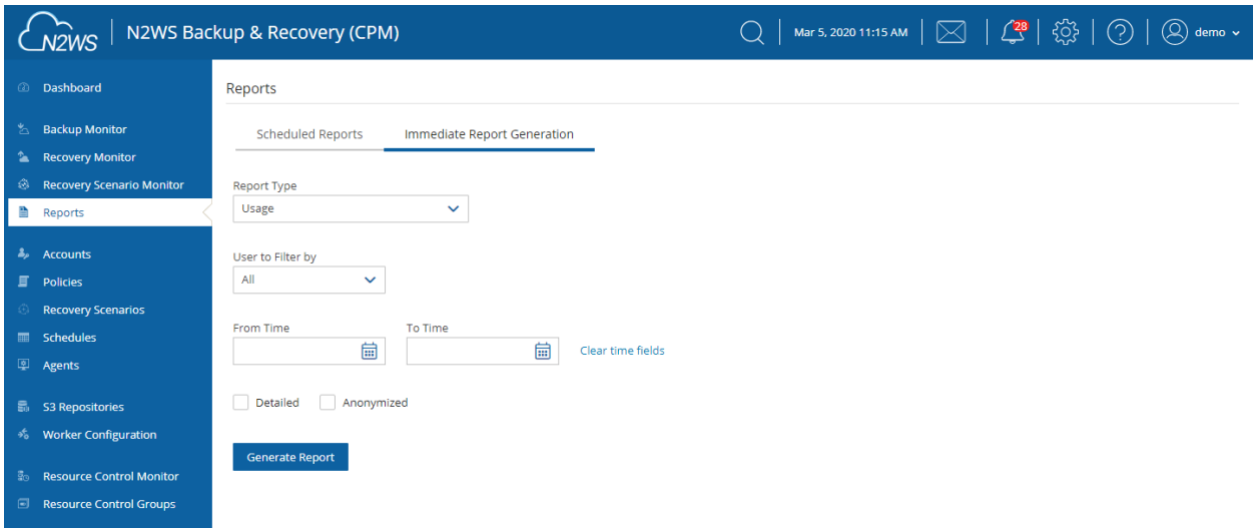

Note: Data saved to the reports is compliant with the EU's General Data Protection Regulation (GDPR).

18.6 Audit Reports

N2WS will record every operation initiated by users and delegates. This is important when the admin needs to track who performed an operation and when. By default, audit logs are kept for 30 days. The root user can:

- Modify the audit log retention value in the **Cleanup** tab of the **General Settings** screen. See section 9.4.
- Download audit reports for specific users or delegates by selecting **audit report** in the users or delegates screen.
- Download the audit report for all users by selecting the link **audit report (all users)** at the bottom of N2WS' main screen.

Included in the audit reports are:

- A timestamp.
- The event type.
- A description of the exact operation.
- In the report of all users, the user with delegate information, if any.

18.7 Configuring for SES

Amazon Simple Email Service (SES) is a cloud-based email sending service that N2WS uses to effortlessly distribute reports. The AWS SES parameters are configured in **Server Settings** > **General Settings**.

Note:	Currently, the only regions that are available for the SES service are Asia Pacific (Mumbai), Asia Pacific (Sydney), EU (Frankfurt), EU (Ireland), US East (N. Virginia), US West (Oregon).

To allow N2WS to configure the AWS SES parameters:

1. In the toolbar, select **Server Settings** > **General Settings**.
2. Select the **Simple Email Service** tab.
3. Select **Enable SES Configuration**.
4. Complete the parameters:
   - **Sender Email Address** – The 'From' e-mail address.
   - **Verify Email Address** – Select to verify address.
   - **SES Region** – Select the region for the SES service.
   - **Authentication Method** – Select a method and supply additional information if prompted:
     - **IAM User Credentials** – Enter AWS Access and Secret keys.
     - **CPM Instance IAM Role** – Additional information is not needed.
     - **Account** – In the **Account** list, select one of the CPM accounts defined in the **Accounts** tab.
5. When finished, select **Save** to confirm the parameters.

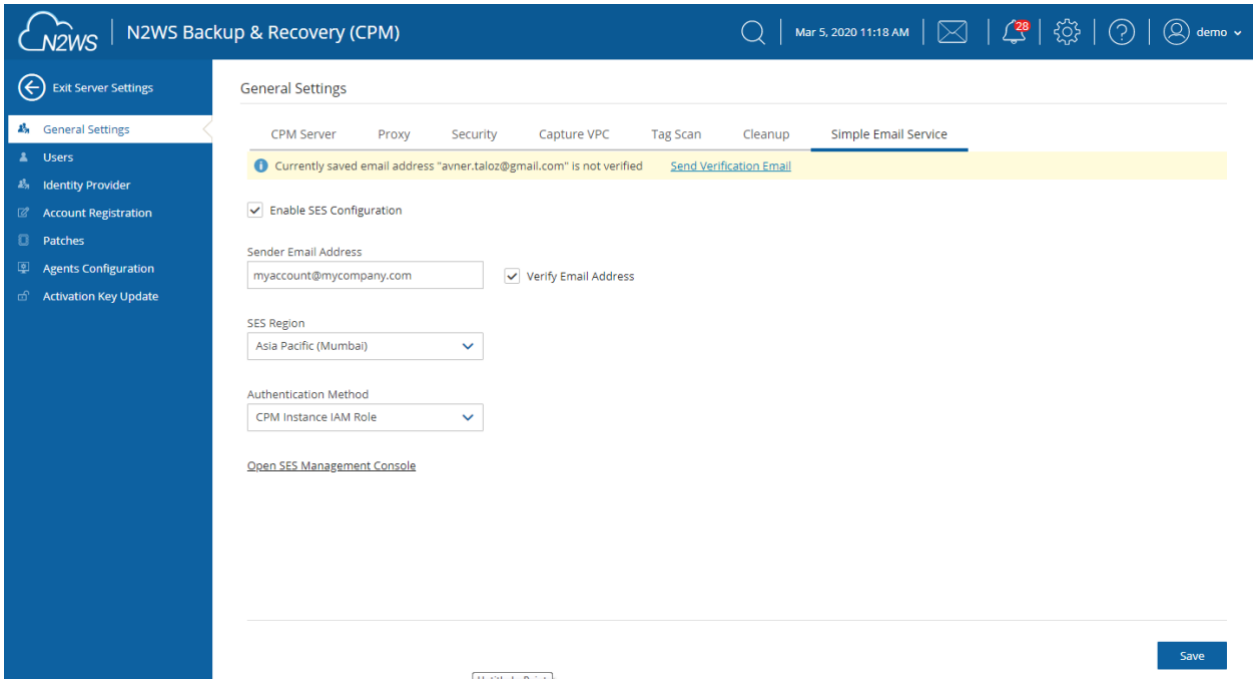

Amazon will respond with an Email Address Verification Request for the region to the defined address. The Amazon verification e-mail contains directions for completing the verification process, including the amount of time the confirmation link is valid.
Currently, the Scheduled Reports are sent using the defined SES email identity if the reports are run with **Schedules** or the **Run Now** option.

19 N2WS IdP Integration

N2WS supports users configured locally (local users) and users configured using the organization's federated identity provider (IdP).

- Local users are created and managed using the N2WS User Management capabilities described above.
- IdP users are users whose credentials are received from the organization's IdP. N2WS can be configured to allow users in the organization's IdP system to login to N2WS using their IdP credentials. Integration with IdP systems is performed using the SAML 2.0 protocol.
- N2WS supports:
  - Active Directory (AD) 2012 and 2016. If using SAML 2.0, AD 2019 also supported.
  - Azure Active Directory-based Single Sign-On (SSO)
  - IDP vendors who support SAML 2.0

Note: The N2WS root user can only login through the local user account even when N2WS is configured to work with IdP.

Configuring N2WS to work with IdP consists of the following:

- Configuring the IdP to work with N2WS
- Configuring N2WS to work with the IdP
- Configuring N2WS Groups in N2WS
- Configuring N2WS Groups and Users in IdP

19.1 Configuring IdPs to Work with N2WS

N2WS supports the SAML 2.0 protocol for integration with IdP systems. N2W Software qualifies only certain IdP systems internally, but any SAML 2.0 compliant IdP system should be able to work smoothly with N2WS.

19.1.1 Prerequisites to IdP Integration with N2WS

Prior to configuring N2WS to work with an IdP system, it is required that N2WS be configured in the IdP system as a new application. Consult the IdP system's documentation on how to configure a new application.

Note: When configuring N2WS as a new IdP application, verify that:

- The default Name **ID** format used in SAML requests is set to **Unspecified**, or modify the default N2WS configuration as per section on N2WS configuration below.
- The X509 certificate Secure hash algorithm is set to SHA-256.
- The following URL values are used:

  Note: `<N2WS_ADDRESS>` is either the DNS name or the IP address of the N2WS Server.

  - **Entity ID** - `https://<N2WS_ADDRESS>/remote_auth/metadata`

- **Sign in response** - `https://<N2WS_ADDRESS>/remote_auth/complete_login/`
- **Sign out response** - https://<N2WS_ADDRESS>/remote_auth/complete_logout/

As part of configuring N2WS as a new IdP application, the IdP system will request a file containing the N2WS x509 certificate. The certificate file can be obtained from the N2WS **Settings** screen in the **Identity Provider** tab. In the **Settings** tab, select **Download CPM Certificate** and choose a location to save the file. See section 19.1.2.
If configuring N2WS to work with Microsoft Active Directory/AD FS, refer to section 19.4.1.

19.1.2 Configuring N2WS for IdP Integration

If configuring N2WS for integration with Microsoft Active Directory/AD FS, refer to section 19.5.

To configure N2WS to work with the organization's IdP:

1. In the N2WS toolbar, select **Server Settings**.
2. In the left panel, select the **Identity Provider** tab and then select the **Settings** tab.
3. Select **Identity Provider**. The configuration parameters appear.

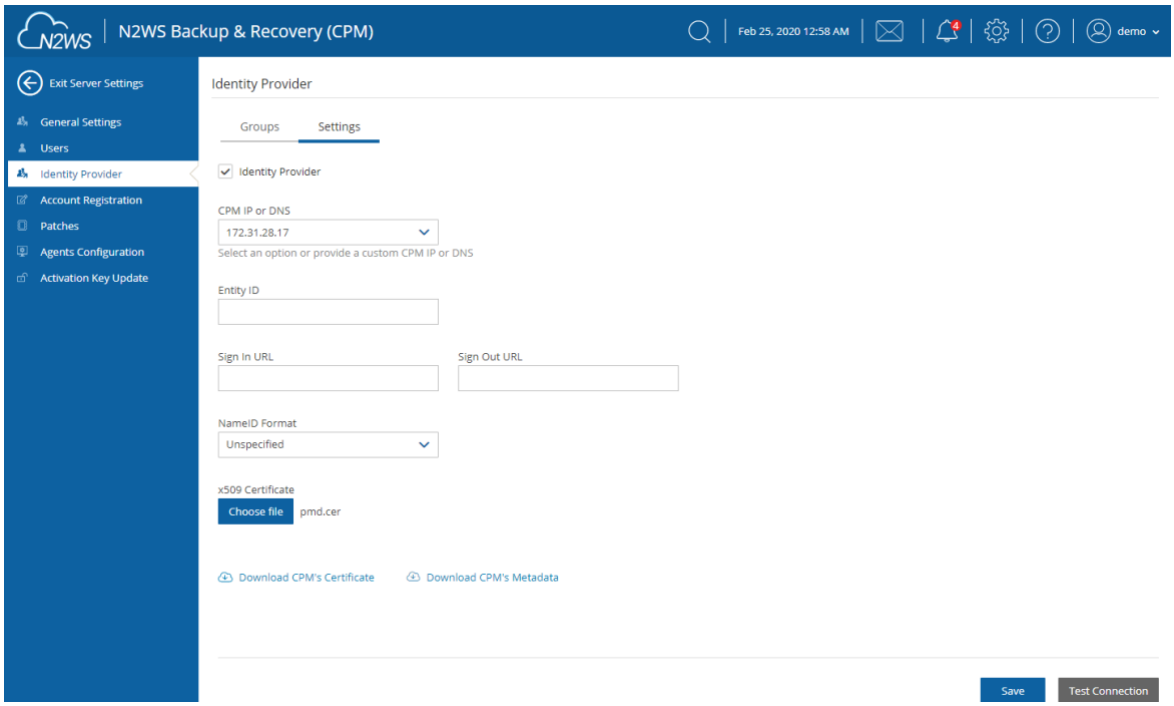

4. Complete the following:
   - **CPM IP or DNS** – The IP Address or DNS name of the N2WS server.

Note:	N2WS accepts either the IP address or DNS name in many fields. However, some IdPs require that N2WS be configured using the format used when configuring N2WS as an application in the IdP system. If the IdP uses DNS names, use DNS names in N2WS, and if the IdP uses IP address, use IP addresses in N2WS

- **Entity ID** – The Identity Provider Identifier's URI provided by the IdP system. Consult the IdP system's documentation.
- **Sign In URL** – The authentication request target is the URL, provided by the IdP system, to which N2WS will redirect users after entering their IdP credentials. Consult the IdP system's documentation.
- **Sign Out URL** – The logout request target is the URL, provided by the IdP system, to which N2WS will redirect users once they logout of N2WS. Consult the IdP system's documentation.
- **NameID format** – The format of the SAML **NameID** element.
- **X509 Certificate** – Select **Choose file** to upload the IdP's X509 certificate. Consult the IdP system's documentation about obtaining their x509 certificate.

5. Optionally, you can **Download CPM's Certificate** and **Metadata.**
6. Once all the parameters have been entered, select **Save** and then select **Test Connection** to test the connection between N2WS and the IdP.

19.2 Configuring Groups and Group Permissions on the N2WS Side

Groups and the permissions assigned to groups are configured in N2WS. When an IdP user logs into N2WS, the information about the user's group membership is received from the IdP and that group's permissions are assigned to the user.

Note:	Every IdP user must belong to a N2WS group. IdP users who do not belong to a group, even if they have user-specific permissions as detailed below, cannot log on to N2WS. Logon by IdP users who do not belong to a group will be failed with an appropriate error message.
Note:	Default groups do not appear until **Identity Provider** is enabled in the **Settings** tab.

N2WS comes with pre-defined groups named **default***:

- default_managed_users
- default_independent_users
- default_root_delegates
- default_root_delegates_readonly

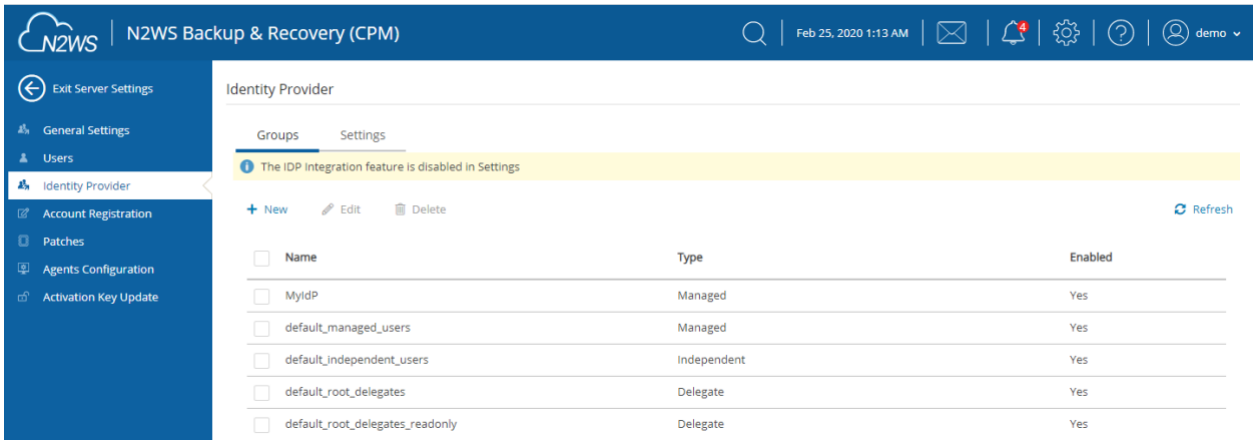

Note:	The default groups cannot be modified or deleted. To see the permission settings assigned to the default groups, select the group name.

Additional groups can be created and removed easily in the **Identity Provider** tab of the N2WS **Server Settings** screen.

To add a new group:

Note:	The group permission settings essentially mirror the user permissions detailed in section 18.

1. In the **Identity Provider** tab, select the **Groups** tab and then select **+ New**. The New IDP Group screen will appear.

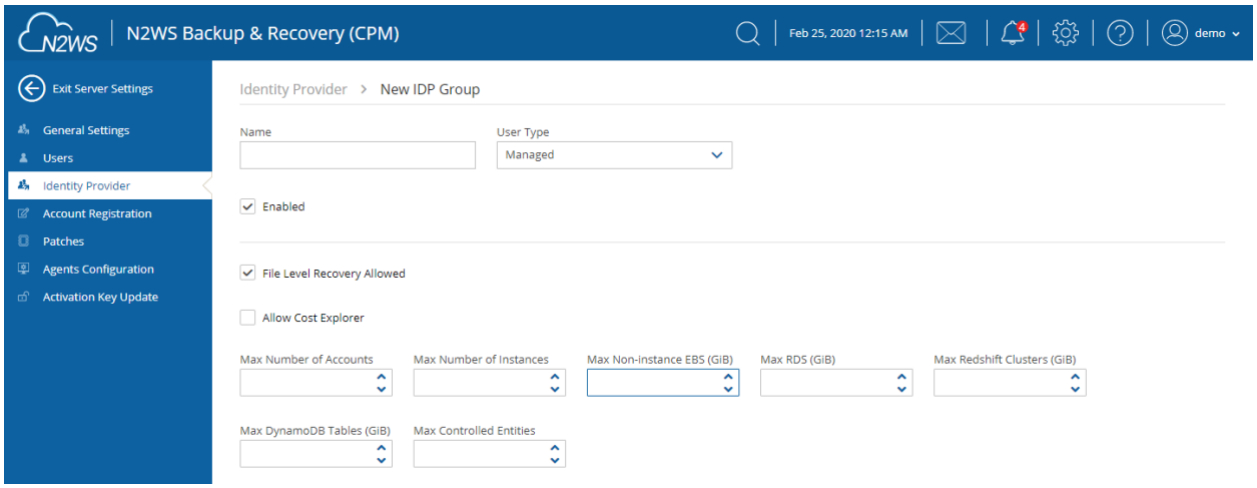

   - **Name** – Name of the group.
   - **User Type** – For details, see section 18. Parameters depend on the **User Type** selected.
     - Managed
     - Independent
     - Delegate
   - **Enabled** – When disabled, group users will not be able to log on to N2WS.
2. For User Type **Managed**:

- **File Level Recovery Allowed**– When selected**,** members of the group can use the file-level recovery feature.
- **Allow Cost Explorer** – When selected, members of group can see cost data. For Cost Explorer information, see section 25.
- **Max Number of Accounts** – The maximum number of AWS accounts users belonging to this group can manage.
- **Max Number of Instances** – The maximum number of instances users belonging to this group can manage.
- **Max Non-Instance EBS (GiB)** – The maximum number of Gigabytes of EBS storage that is not attached to EC2 instances that users belonging to this group can manage.
- **Max RDS (GiB**) – The maximum number of Gigabytes of RDS databases that users belonging to this group can manage.
- **Max Redshift Clusters (GiB)** – The maximum number of Gigabytes of Redshift clusters that users belonging to this group can manage.
- **Max DynamoDB Tables (GiB)** – The maximum number of Gigabytes of DynamoDB tables that users belonging to this group can manage.
- **Max Controlled Entities** – The maximum number of allowed entities for Resource Control.

3. For User Type **Delegate**:

Note:	When Delegate is selected, the **Original Username** to which this group is a delegate is required although the Original Username does not yet need to exist in N2WS. After creation, the Original Username cannot be modified.

- **Original Username** – User name of delegate.
- **Allow to Perform Recovery** – Whether the delegate can initiate a recovery.
- **Allow to Change Accounts** – Whether the delegate can make changes to an account.
- **Allow to Change Backup** – Whether the delegate can make changes to a backup.

19.3 Configuring Groups on the IdP Side

IdPs indicate a user's group membership to N2WS using IdP claims. Specifically, the IdP must configure an **Outgoing Claim Type** of `cpm_user_groups` whose value is set to all the groups the user is a member of, both N2WS related groups and non-N2WS related groups.

Note:	Group names on the IdP side no longer need the '`cpm`' prefix. In cases where the names of the group users are assigned to in the IdP is of the form `cpm_<GROUP_NAME_IN_N2WS>`, for example `cpm_mygroup` where `mygroup` is the name of a group that was created in N2WS, the `<GROUP_NAME_IN_N2WS>` part of the name must match the name of a group in N2WS. See section 19.2. For example, to give IdP users permissions of the N2WS group `default_managed_users`: 1. The relevant users can be members of an IdP group called `cpm_default_managed_users`. 2. The IdP must have an outgoing claimed called `cpm_user_groups`.

3. The value of the claim must include the names of all the user's groups in the IdP, which presumably includes `cpm_default_managed_users`.

Or

1. The relevant users can be members of an IdP group called `default_managed_users`.
2. The IdP must have an outgoing claimed called `cpm_user_groups`.
3. The value of the claim should not include the names of all the user's groups in the IdP, which presumably is `default_managed_users`.

Note: An IdP user logging onto N2WS can belong to only one N2WS group, i.e. of all the groups listed in the `cpm_user_groups` claim, only one can be a N2WS group, such as `cmp_mygroup`. If an IdP user is a member of more than one N2WS group, the log on will fail with a message indicating the user belongs to more than one N2WS group.

19.3.1 Understanding N2WS User Permissions

A user logged into the N2WS system can have several types of permissions. This section discusses the different types of permissions as they are applied to N2WS IdP integration. For full treatment of the meanings of these permissions, see section 16.3. To override N2WS group permissions on a per user basis, see section 19.3.2.

General User Attributes

Attribute Name	Mandatory (Y/N)	Meaning	Valid Values
user_type	N	Type of user.	• Managed • Independent • Delegate
user_name	N	Username in N2WS.	Alphanumeric string
user_email	N	User's email address.	Valid email address

Attributes for Independent and Managed Users

Attribute Name	Mandatory (Y/N)	Meaning	Valid Values
allow_file_level_recovery	N	Whether the user is allowed to use the N2WS file-level restore feature.	yes, no
max_accounts	N	The number of AWS accounts the user can manage in N2WS. Varies by N2WS license type.	Number between 1 and max licensed
max_instances	N	The number of instances the user can backup. Varies by N2WS license type.	Number between 1 and max licensed

Attribute Name	Mandatory (Y/N)	Meaning	Valid Values
max_independent_ebs_gib	N	Total size of EBS independent volumes being backed up in GiB (i.e. volumes not attached to a backed-up instance).	Number between 1 and max licensed
max_rds_gib	N	Total size of AWS RDS data being backed up in GiB	Number between 1 and max licensed
max_redshift_gib	N	Total size of AWS Redshift data being backed up in GiB	Number between 1 and max licensed
max_dynamodb_gib	N	Total size of AWS DynamoDB data being backed up in GiB.	Number between 1 and max licensed
max_controlled_entities	N	Total number of AWS resources under N2WS Resource Control.	Number between 1 and max licensed.

Attributes for Delegate Users

Attribute Name	Mandatory (Y/N)	Meaning	Valid Values
original_username	Y	The name of the user for whom **user_name** is a delegate.	Alphanumeric string
allow_recovery	N	Whether the user can perform N2WS restore operations.	yes, no
allow_account_changes	N	Whether the user can manage N2WS user accounts.	yes, no
allow_backup_changes	N	Whether the user can modify backup policies.	yes, no

All the permissions detailed above are set for a group when the group is created in N2WS. Additionally, it is possible to assign N2WS permission at the level of individual IdP users as described in 19.3.2. When there is a conflict between a user's group permissions and a user's individual permissions, the individual permissions take precedence.
A permission string consists of **key=value** pairs, with pairs separated by a semicolon.
For convenience, below is a string of all the possible security parameters. N2WS will accept a partial list consisting of any number of these parameters in any order:

```
user_type=independent;email=yeepee@redpil.com;allow_recovery=yes;allow
_account_changes=yes;allow_backup_changes=yes;allow_file_level_restore
=no;max_accounts=1;max_instances=2;max_independent_ebs_gib=3;max_rds_g
ib=4;max_redshift_gib=5;max_dynamodb_gib=5;original_username=robi@stam
```

19.3.2 Overriding Group Settings at the User Level

Users get the N2WS permissions assigned to their group. However, it is possible to give specific IdP group members permissions different from their group permissions.

To override the group permission for a specific user:

1. The IdP administrator must first enter the new permissions in an IdP user attribute associated with the user. The attribute can be an existing attribute that will now serve this role (e.g. msDS-cloudExtensionAttribute1) or a custom attribute added to the IdP user schema specifically for this purpose.

   The content of the attribute specifies the user's N2WS permissions in the `key=values` format detailed in the section above.
   - Permissions specified in the user attribute will override permissions inherited from the group.
   - Permission types not specified in the user attribute will be inherited from the group's permissions. For example, if the attribute contains only the value `max_accounts=1`, all other permissions will be inherited from the user's group permissions.
2. Once a user attribute has been configured with the correct permissions, an IdP claim rule with Outgoing Claim Type `cpm_user_permissions` must be created. The value of the claim must be mapped to the value of the attribute chosen above.
3. When the user-level claim is enabled, the user will be able to log on to N2WS with permissions that are different from the group's permissions.

If configuring Microsoft Active Directory/AD FS, refer to section 19.6 for details.

19.4 N2WS Login Using IdP Credentials

In order to use IdP credentials to log on to N2WS, users need to select the **Sign in with: Identity Provider** option on the N2WS Logon screen.

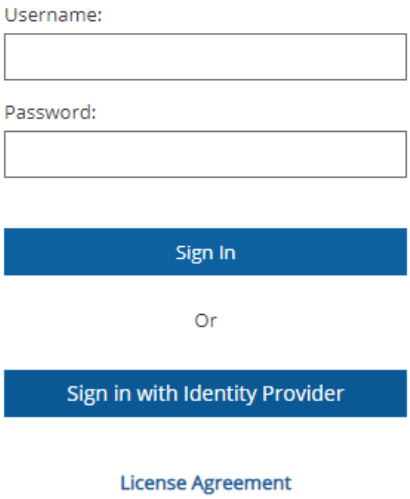

Selecting **Sign in with Identity Provider** will redirect the user to the organization's IdP system using SAML.

Note: To log on to N2WS as root, log on with the standard user and password option.

19.4.1 Configuring AD/AD FS for Integration with N2WS

To enable N2WS to integrate with AD/AD FS, N2WS must be added to AD FS as a **Relying Party Trust**.

Note:	The following AD FS screenshots are from AD 2012. The AD 2016 screens are very similar.

To run the Add Relying Party Trust Wizard:

1. In the left pane of the AD FS console, select **Relying Party Trusts**.
2. In the right pane, select **Add Relying Party Trust. . .**.. The Wizard opens.

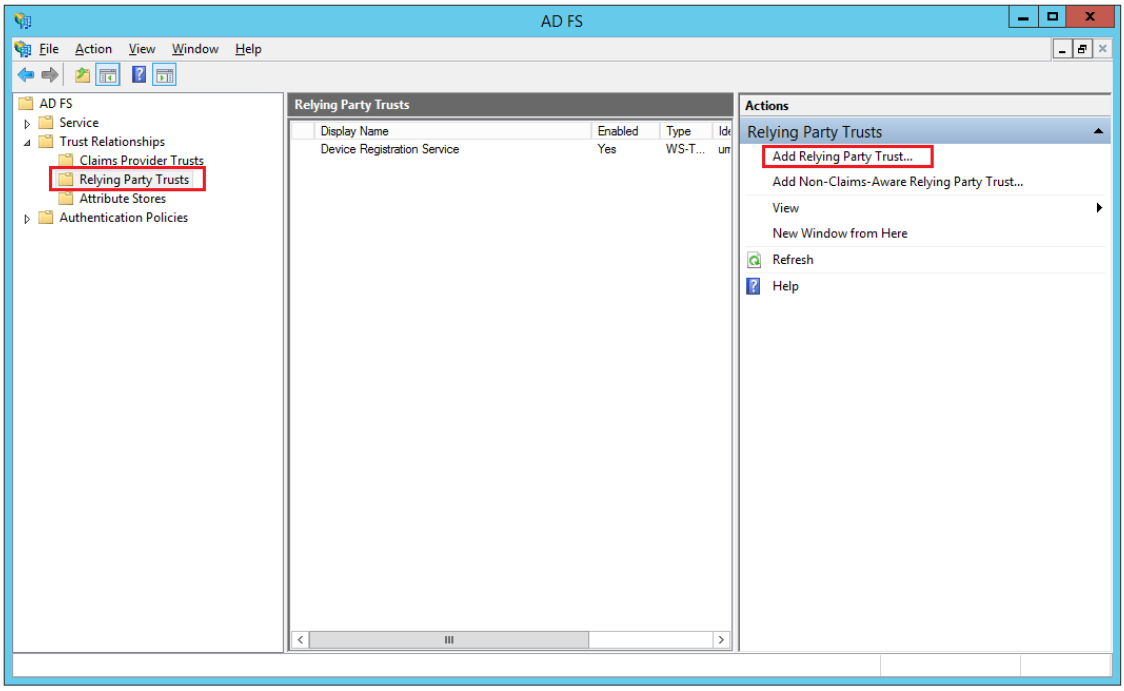

3. Select **Start**.
4. Select the **Enter data about the relying party manually** option.
5. Select **Next**.
6. On the **Welcome** screen, type the display name for N2WS (e.g. `N2WS`), and then select **Next**.
7. On the **Choose Profile** screen, select the **AD FS profile** option, and then select **Next**.
8. Skip the **Configure Certificate** screen by selecting **Next**.
9. On the **Configure URL** screen:
   a. Select **Enable support for SAML 2.0 WebSSO protocol**.
   b. In the **Relying Party SAML 2.0 SSO Service URL** box, type `https://` followed by the N2WS DNS name or IP address, and then followed by `/remote_auth/complete_login/`.
   For example, the resulting string might look like:
   `https://ec2-123-245-789.aws.com/remote_auth/complete_login/`
10. Select **Next**.

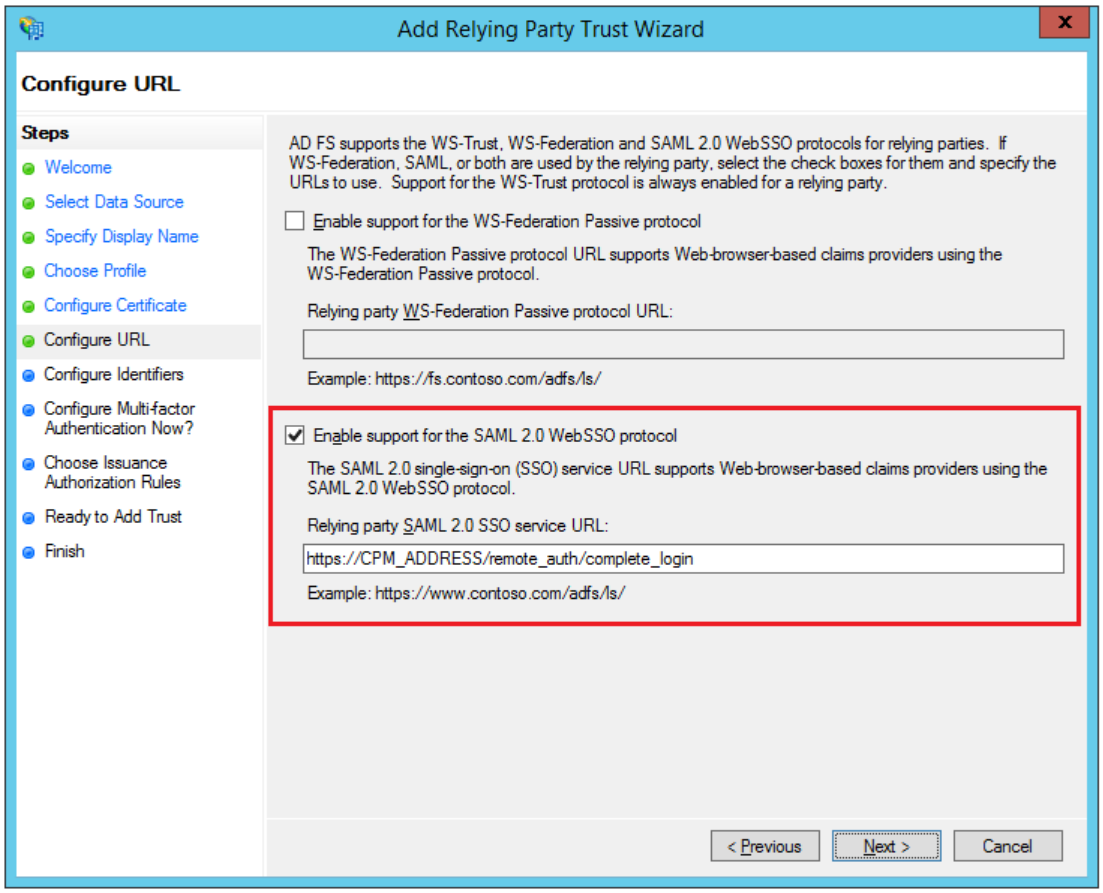

11. In the **Configure Identifiers** screen, type `https://` followed by the N2WS DNS name or IP address, and then followed by `/remote_auth/metadata` in the **Relying party trust identifier** box.

    For example, the resulting string might look like:
    `https://ec2-123-245-789.aws.com/remote_auth/metadata`
12. Select **Add** on the right.

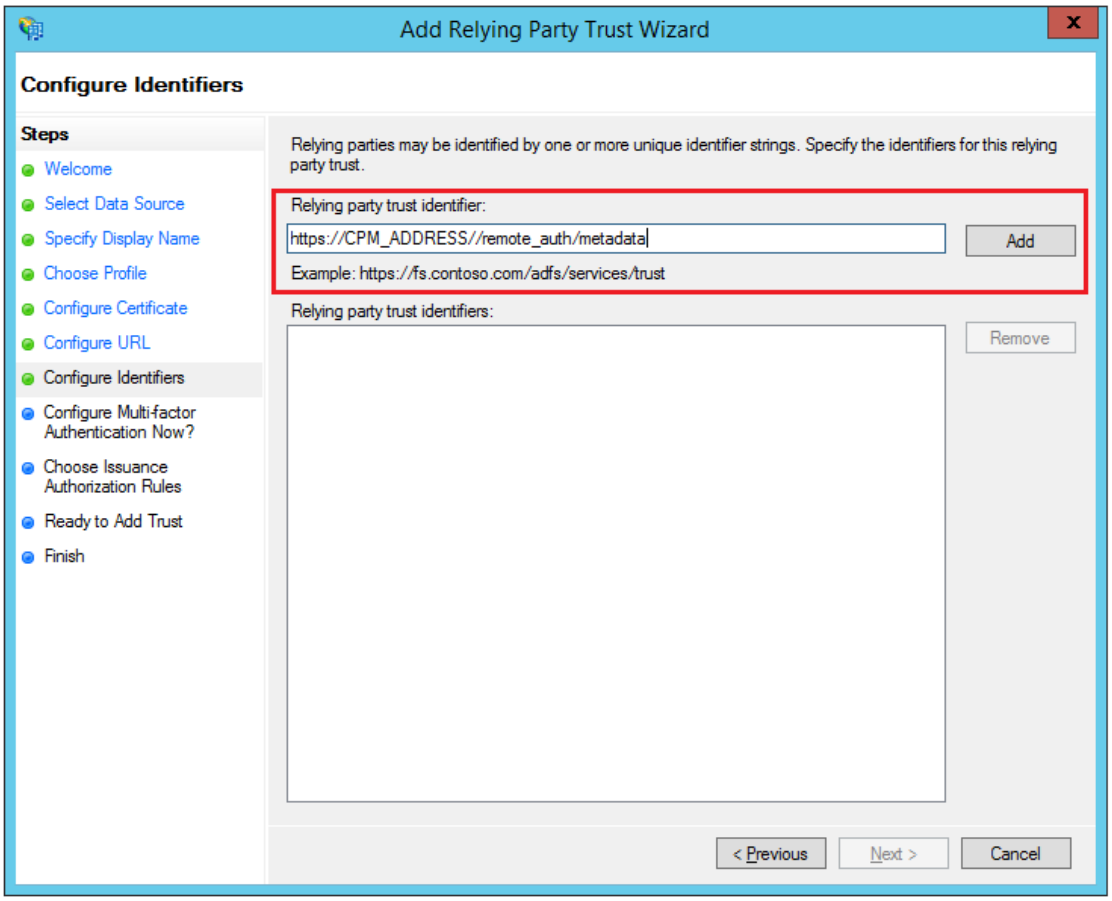

13. Select **Next**.
14. On the **Configure Multi-factor Authentication Now?** screen, select the **I do not want to configure multi-factor authentication settings for this relying party trust at this time** option, and then select **Next**.
15. On the **Issuance Authorization Rules** screen, select the **Permit all users to access this relying party** option, and then select **Next**.
16. On the **Ready to Add Trust** screen, review the setting of the **Relying party trust** configured with the Wizard. When finished, select **Next**.
17. On the **Finish** screen of the Wizard, select **Close**. There is no need to select the **Open the Edit Claim Rules dialogue for this relying party trust when the wizard closes** option.

19.4.2 Setting AD FS Properties

Once the Relying Party Trust has been configured, set the AD FS properties.

To set the AD FS properties:

1. Go back to the AD FS management console, and in the middle pane, right-select the N2WS line under **Relying Party Trust**, and then select **Properties**.
2. On the screen that opens, select the **Endpoints** tab, and then select **Add SAML…**.

In the **Edit Endpoint** screen, select **SAML Logout** from the **Endpoint type** list.

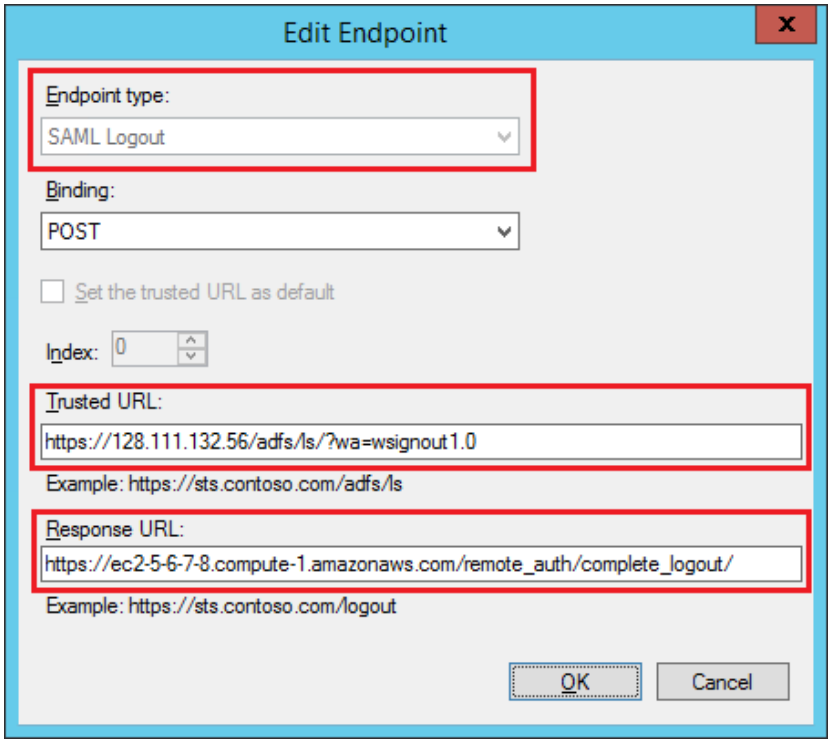

3. In the **Trusted URL:** box, type the DNS name or IP address of the AD FS server followed by `/adfs/ls/?wa=wsignout1.0` (e.g. `https://adserver.mycompany.com/adfs/ls/?wa=wsignout1.0`)
4. In the **Response URL:** box, type DNS name or IP address of the N2WS server followed by `/remote_auth/complete_logout/` (e.g. `https://ec2-123-245-789.aws.com/remote_auth/complete_logout/`).
5. Select **OK**.
6. Go to the **Advanced** tab, and in the **Secure hash algorithm** list, select **SHA-256**. Select **Apply**.

19.4.3 Installing the N2WS Certificate

In order for N2WS to work with AD FS the X.509 certificate used by N2WS needs to be added to the AD FS **Trusted Root Certification Authorities** list. If you installed your own certificate in N2WS when you first configured N2WS (as per section 2.5.3) then your certificate may already be in your AD FS root trust. Otherwise you will need to add it. If you used the certificate N2WS creates during installation, you will need to add that certificate into the AD FS **Trusted Root Certification Authorities**.

To add a root certificate to the AD FS Trusted Root Certification Authorities:

1. Go to the **Signature** tab under properties and select **Add….**
2. In the **File** box at the bottom of the screen, type the name of the file containing the N2WS x.509 certificate. This will be either:
   a. The root certificate you installed in N2WS when it was first configured as per section 2.5.3 of the User Guide, if not already in the AD FS Trusted Root Certification Authorities, or

b. The certificate N2WS created when it was first configured.

3. To obtain a copy of the certificate being used by N2WS, either the one you originally installed or the one N2WS created, select **Download N2WS's certificate file** in the Active Directory Configuration section of the N2WS **General Settings** screen.
4. Once you have entered the name of the file, select **Open**.

   The N2WS certificate is now visible in the center pane in the **Signature** tab.
5. In the center pane of the **Signature** tab, double select the N2WS certificate.
6. Under the **General** tab, select **Install Certificate….**
7. In the **Certificate Import Wizard** screen, select the **Local Machine** option, and then select **Next**.
8. Select the **Place all certificates in the following store** option, select **Browse…,** and then select the **Trusted Root Certification Authorities** store. Select **OK**.
9. Select **Next**.
10. Select **Finish**. Then select **OK** on the pop-up screen, select **OK** on the **General** tab, and then select **OK** on the **Properties** screen.

The next step is to create a Name ID claim in AD FS.

19.4.4 Creating an AD FS Name ID Claim

To create an AD FS claim:

1. Open the ADFS management console. In the main page of the management console, select **Relying Party Trusts** in the left pane.
2. In the middle **Relying Party Trust** pane, select N2WS' party (e.g. `N2WS`).
3. In the right pane, select **Edit Claim Rules…**
4. In the **Edit Claim Rules** screen, select **Add Rule**.

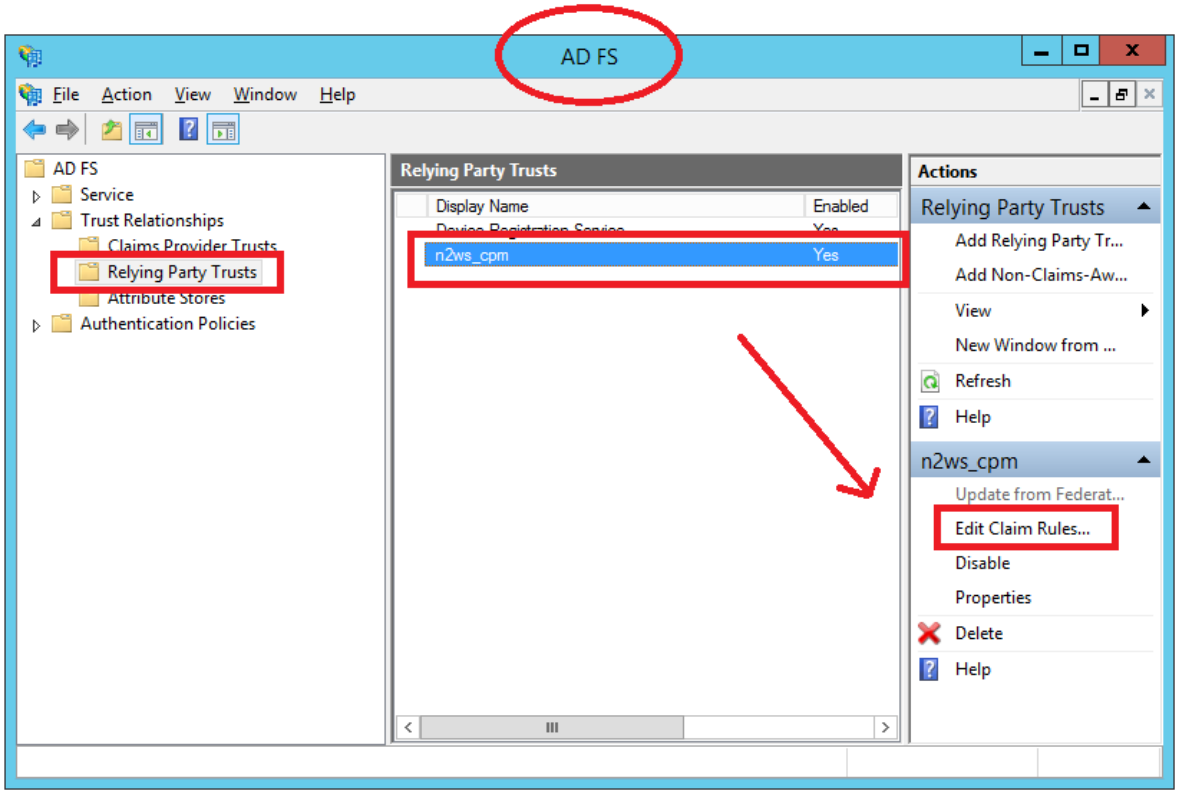

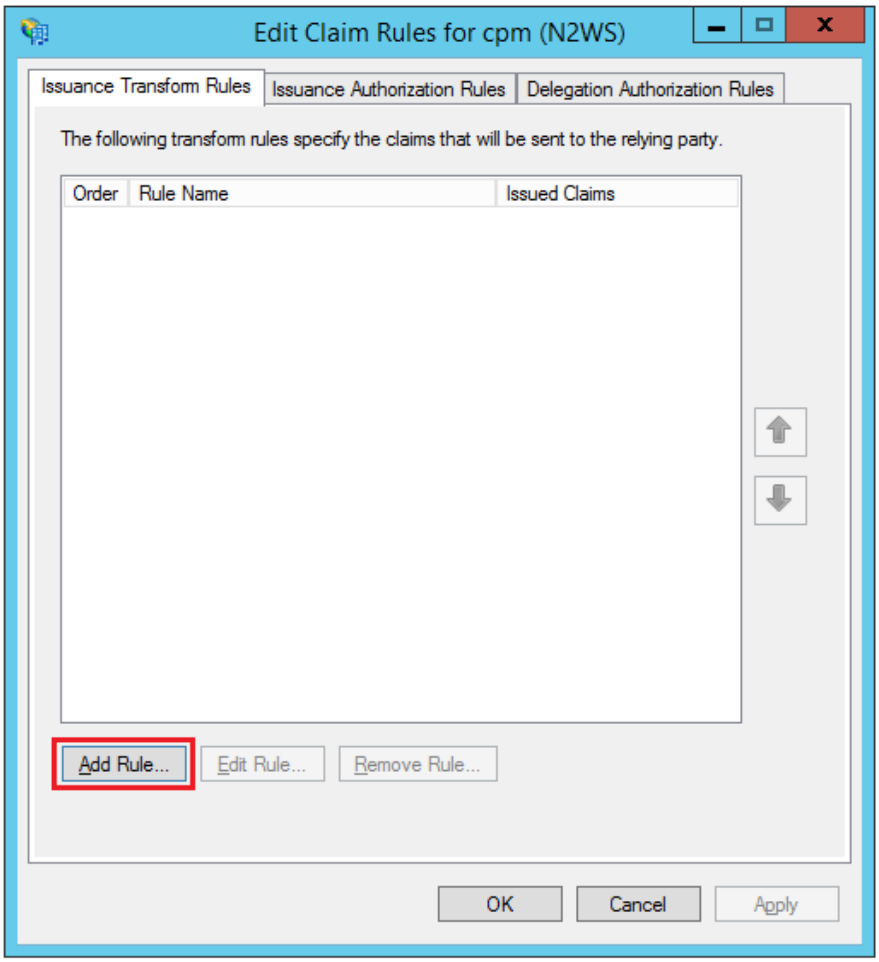

5. In the **Claim rule template** list, select **Transform an Incoming Claim** and then select **Next**.
6. Complete the **Add Transform Claim Rule Wizard** screen:
   a. In the **Claim rule name** box, type a name for the claim.
   b. In the **Incoming claim type** list, select **Windows account name**.
   c. In the **Outgoing claim type** list, select **Name ID**.
   d. In the **Outgoing name ID format** list, select **Unspecified**.
   e. Select the **Pass through all claim values** option.
   f. Select **OK**.

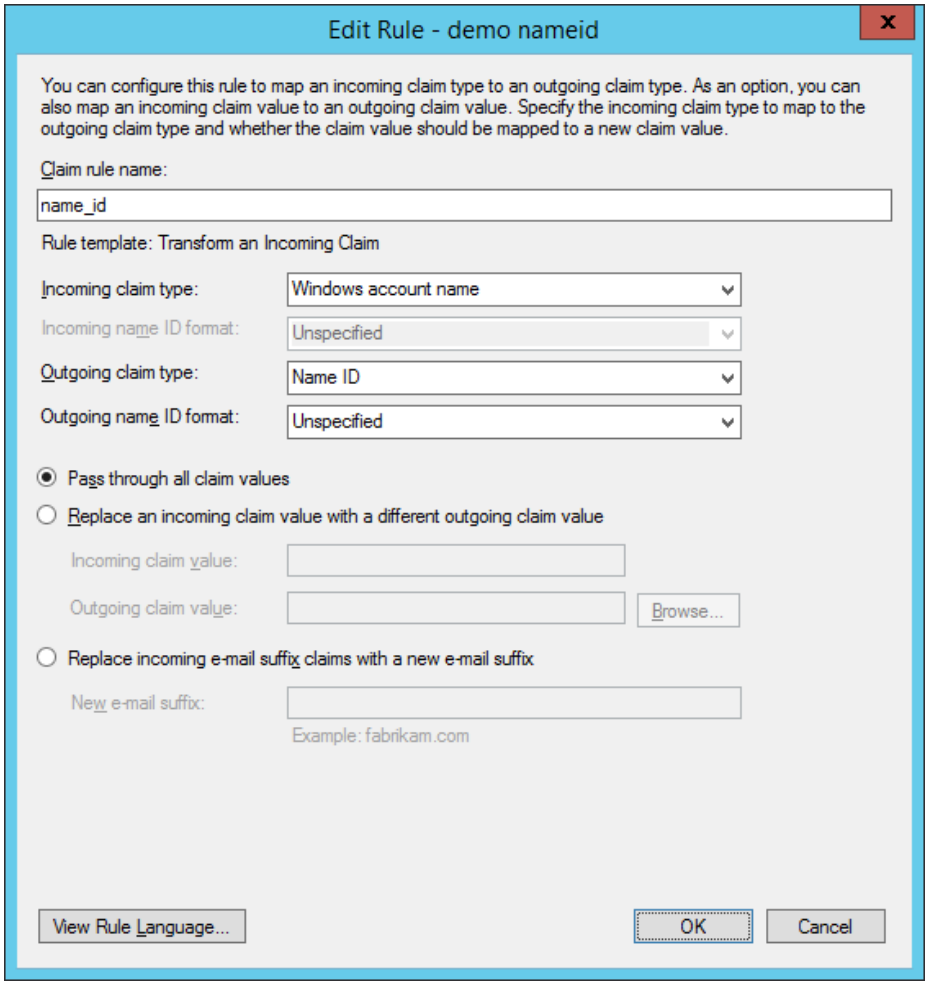

The next step is to add a Token-Groups claim.

19.4.5 Adding a Token-Group's Claim

An ADFS Token-Groups claim must be configured so that AD FS will send N2WS the list of groups a user is a member of. To configure the Token Group's claim, perform steps 1 and 2 of the Configuring Name ID Claim process in section 19.4.4. Then continue as follows:

1. In the **Claim rule template** list, select **Send LDAP Attributes as Claims** and then select **Next.**

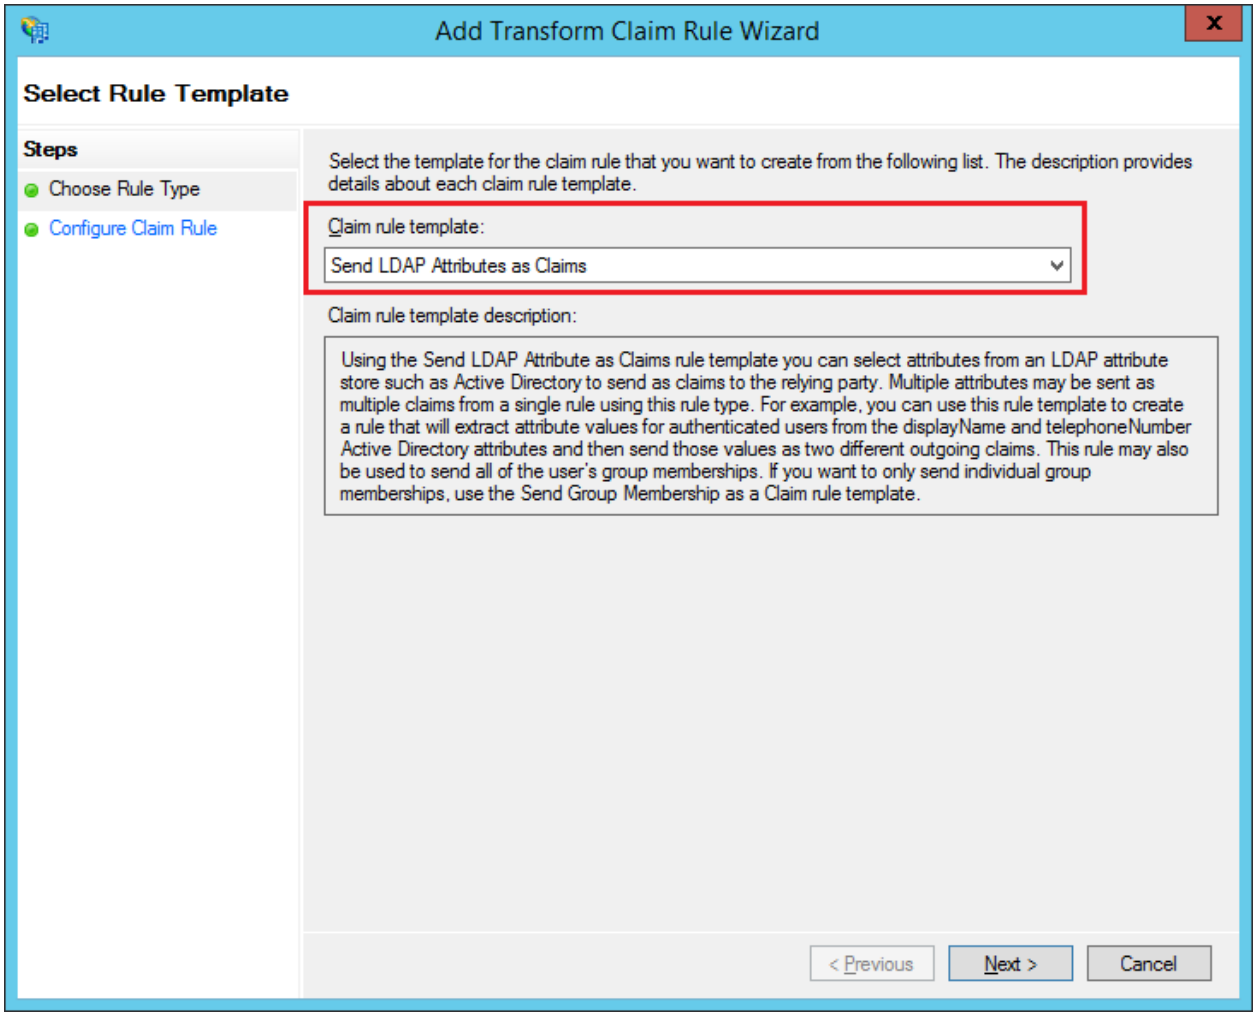

2. In the **Claim rule name** box, type a name for the rule you are creating.
3. In the **Attribute store** list, select **Active Directory**. In the **Mapping of LDAP attributes to outgoing claim types** table:
   a. In the left column (**LDAP Attribute**), select **Token-Groups - Unqualified Names**.
   b. In the right column (**Outgoing Claim Type**), type `cpm_user_groups`.

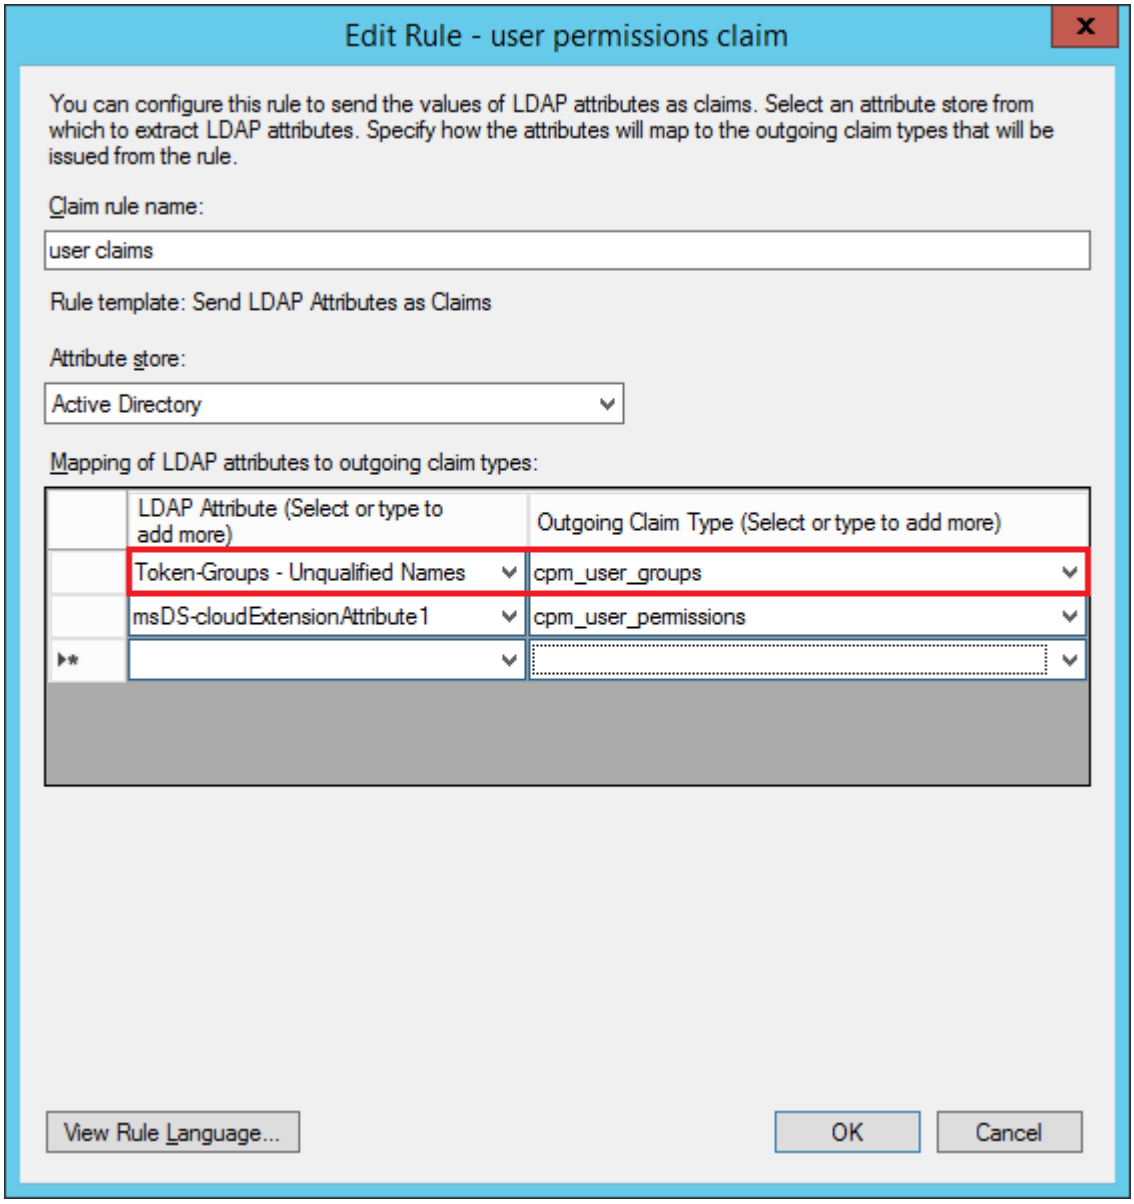

19.4.6 Testing the Connection

At this point AD FS has been configured to work with N2WS. It is now possible to perform a connectivity test between N2WS and AD FS.

To test the connection between N2WS and AD FS:

1. Go to the N2WS **General Settings** screen.
2. Select **Identity Provider**.
3. In the **Groups** tab, select an Identity Provider.
4. In the **Settings** tab, select **Test Connection.**
5. Type a valid AD username and password on the logon page.
6. Select **Sign in**.

19.5 Configuring N2WS to Work with Active Directory / AD FS

To configure N2WS to work with the organization's AD server:

1. Go to the N2WS **Server Settings** > **General Settings**.
2. Select the **Identity Provider** tab.
3. In the Identity Provider list, select a **Group.**
4. To enable the group, select the **Settings** tab and then select **Identity Provider.** Several IdP related parameters are presented.

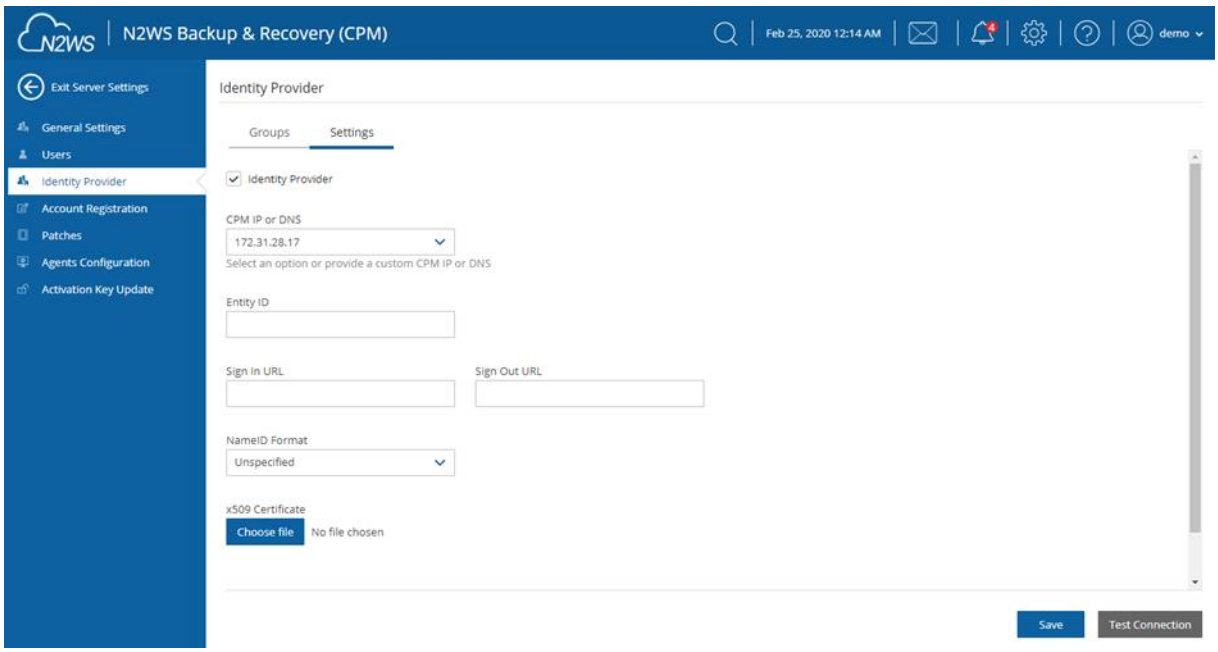

5. In the **Entity ID** box, type the AD FS **Federation Service Identifier**, as configured in AD FS. See below to locate this parameter in AD FS.

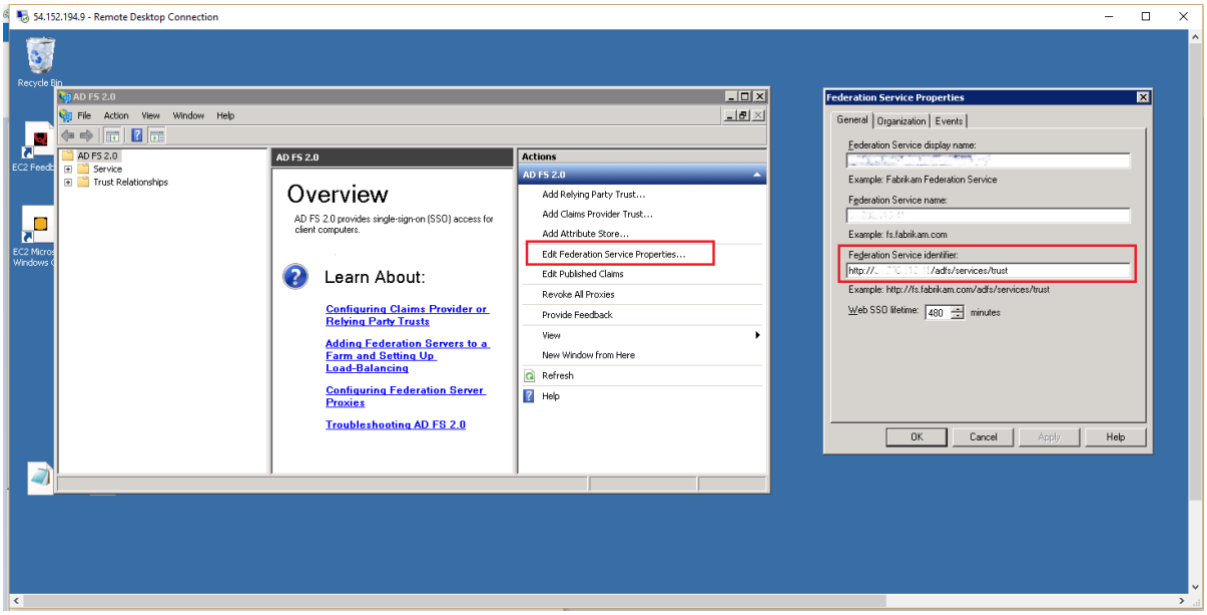

6. In the **Sign in URL** box, type the URL to which N2WS will redirect users for entering their AD credentials.

This parameter is configured as part of AD FS. The AD FS server's DNS name, or IP address, must be prepended to the URL Path listed in AD FS. See below to locate this information in AD FS.

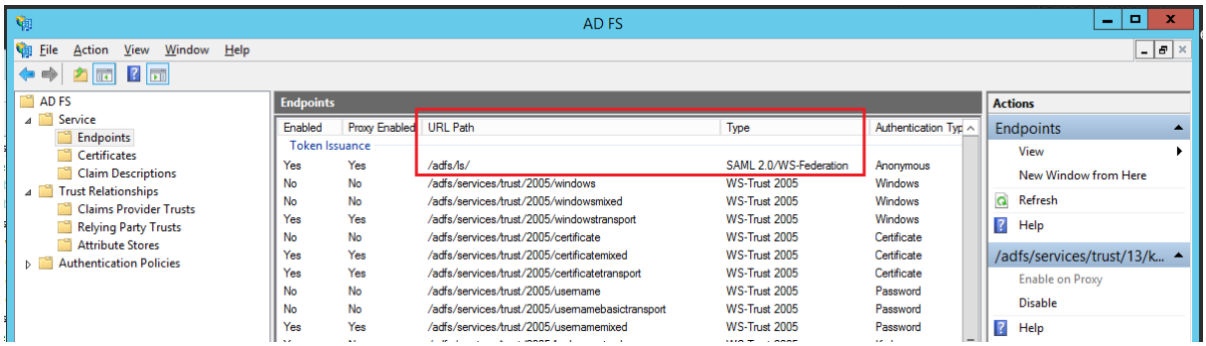

7. In the **NameID Format** list, select the format of the SAML **NameID** element.
8. In the **x509 cert** box, upload the X509 certificate of the AD FS server. The certificate file can be retrieved from the AD FS management console under **Service -> Certificates**, as shown below:

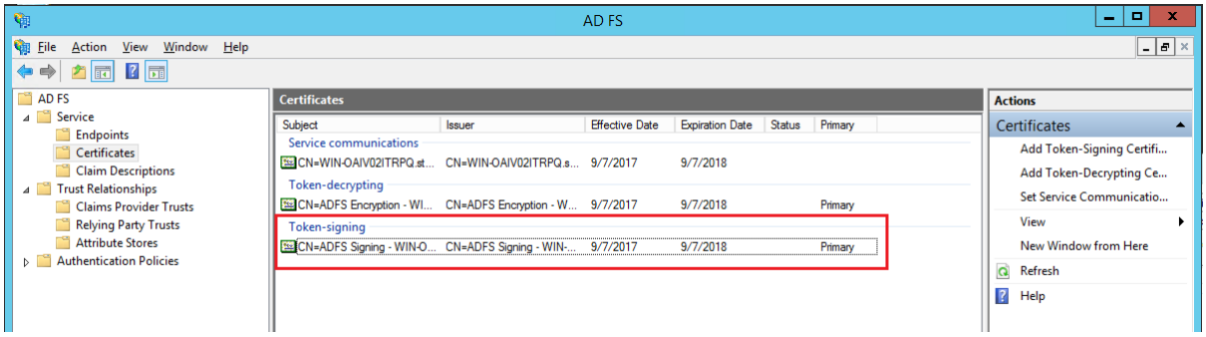

9. To export the IdP's certificate:
   a. Double select the **Token signing** field to open the **Certificate** screen.
   b. Select the **Details** tab and then select **Copy to File . . .** on the bottom right.
   c. Select **Next** to continue with the Certificate Export Wizard.
   d. Select the Base-64 Encoded X.509 (.cer) option and then select Next.
   e. Type a name for the exported file and select **Next**.
   f. Select **Finish**.
10. Once all the parameters for N2WS have been entered, select **Save** and then select **Test Connection** to verify the connection between N2WS and the IdP.

19.6 Configuring an AD FS User Claim

Once a user attribute has been configured with the correct permissions, an ADFS claim rule with **Outgoing Claim Type** `cpm_user_permissions` must be created before the user-level permissions can take effect.

To create the claim rule:

1. Open the AD FS management console.
2. In the main page of the management console, in the left pane, select **Relying Party Trusts**.

3. Select N2WS' party (e.g. `N2WS`) in the middle pane, and in the right pane, select **Edit Claim Rules**.

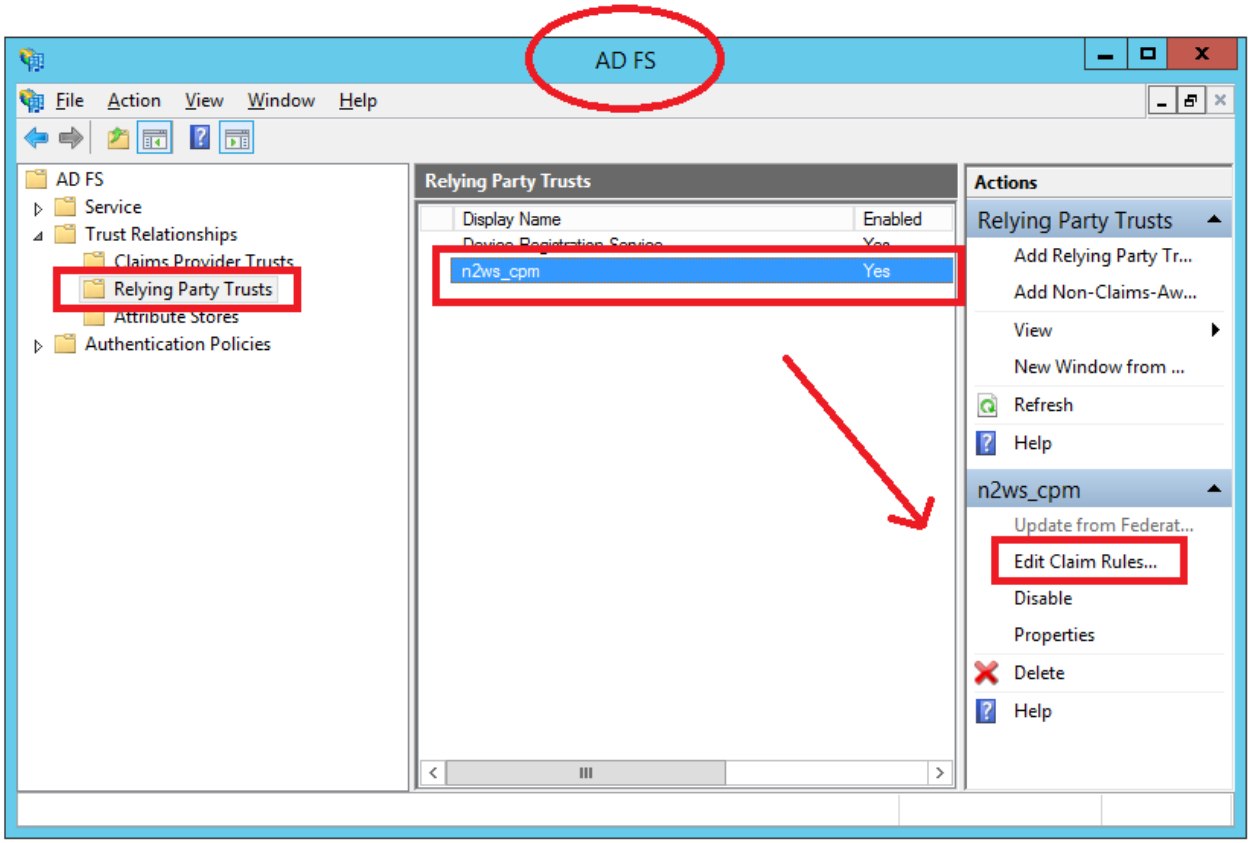

4. In the **Edit Claim Rules** screen, select **Add Rule.**

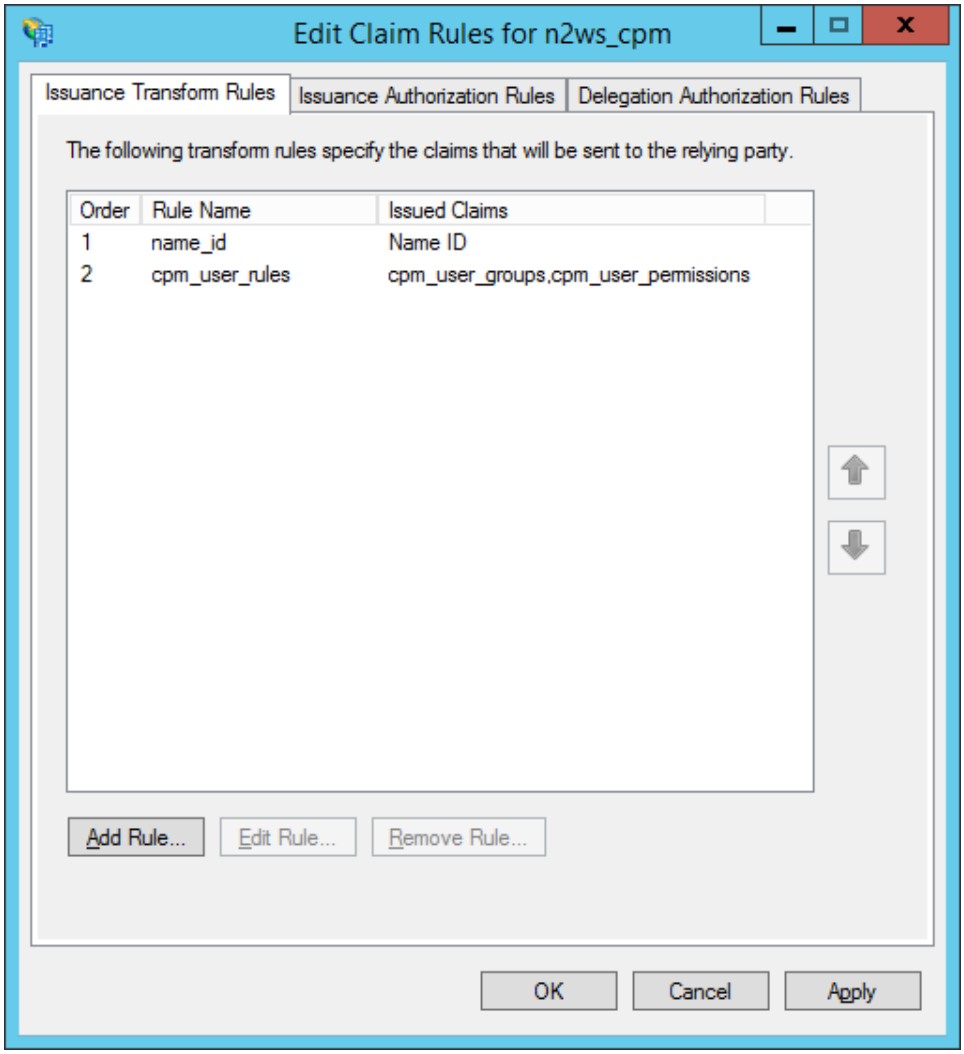

5. In the **Add Transform Claim Rule Wizard** screen, select **Send LDAP Attributes as Claims** in the **Claim rule template** list, and then select **Next**.

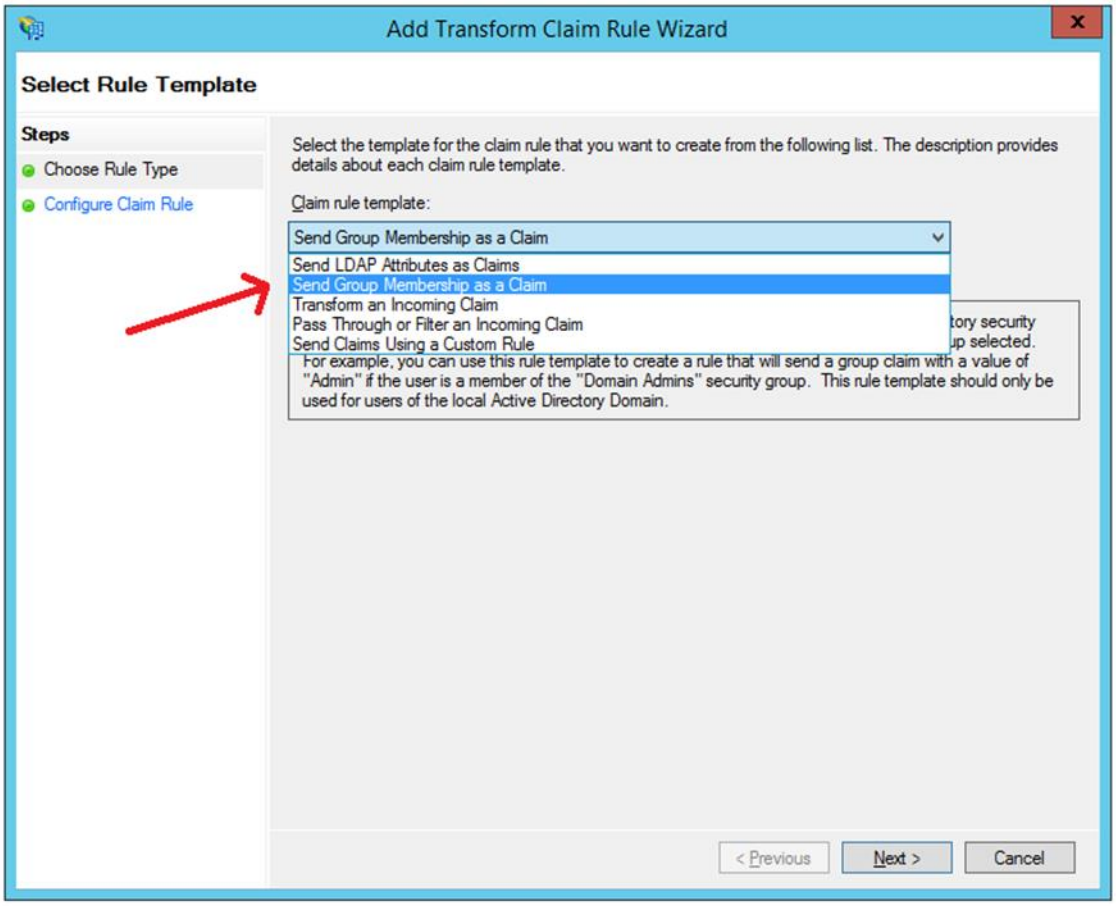

6. The **Claim Rule Wizard** opens the **Edit Rule** screen. Complete as follows:
   a. In the **Claim rule name** box, type a name for the rule you are creating.
   b. In the **Attribute store** list, select **Active Directory**.
   c. In the **Mapping of LDAP attributes to outgoing claim types** table:
      i. In the left column (**LDAP Attribute**), type the name of the user attribute containing the user permissions (e.g. `msDS-cloudExtensionAttribute1`).
      ii. In the right column (**Outgoing Claim Type**), type `cpm_user_permissions`.

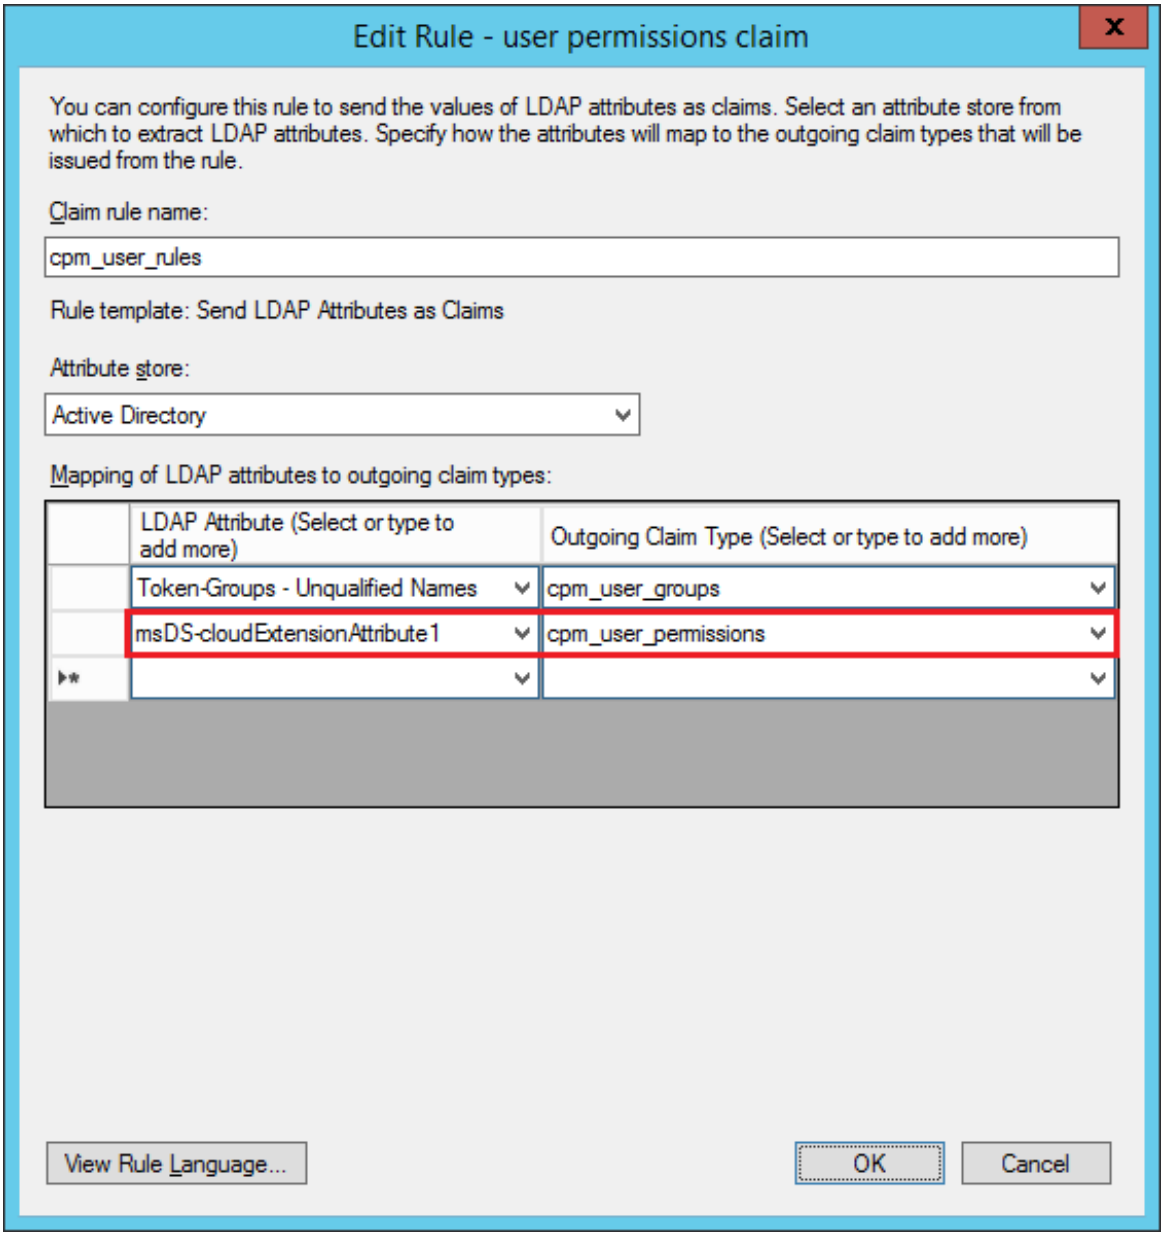

7. Select **OK** to create the rule.

Once the user-level claim is enabled, the user will be able to log on to N2WS with permissions that are different from the group's permissions.

19.7 Configuring Azure AD and N2WS IdP Settings

This section shows how to configure Microsoft Azure Active Directory and N2WS IdP settings to communicate and enable logging.

19.7.1 Azure AD Configuration

After logging in to Azure, go to **Azure Active Directory** in the left menu.

1. Start creating a new user (or use the existing user), group and application in the 'Create' menu on the right.

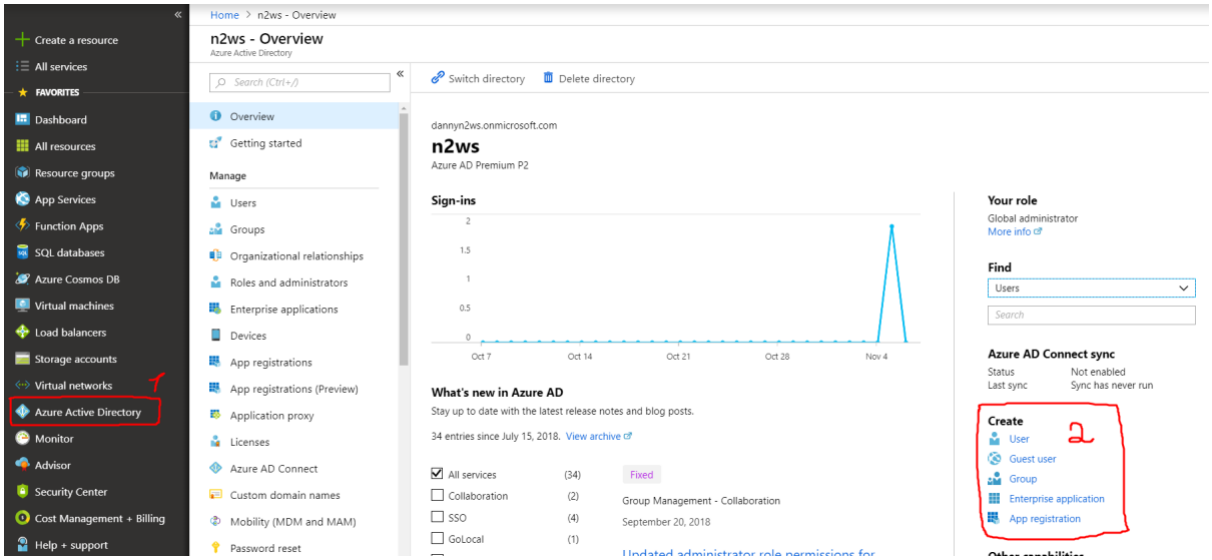

2. Create the group and assign a user:

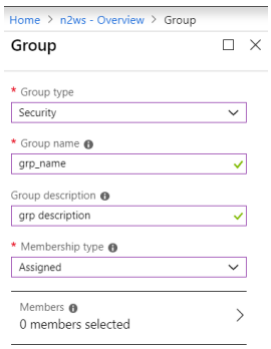

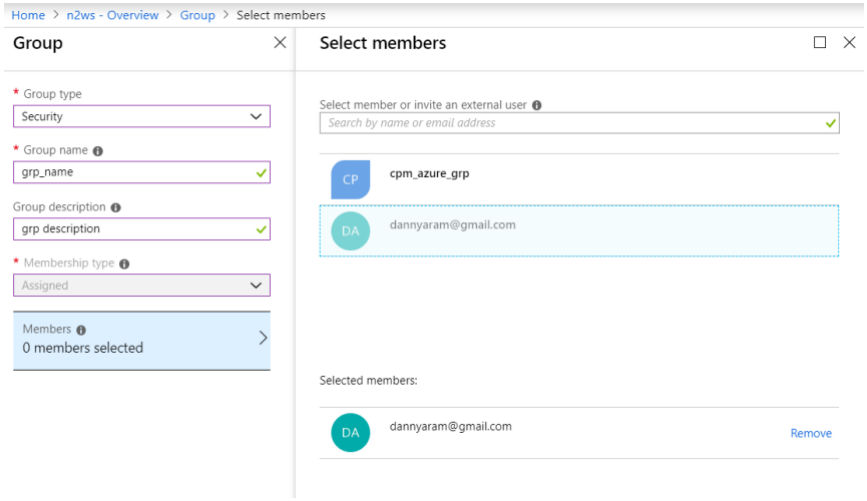

3. Create a new application and choose a Non-Gallery application. Name the application.

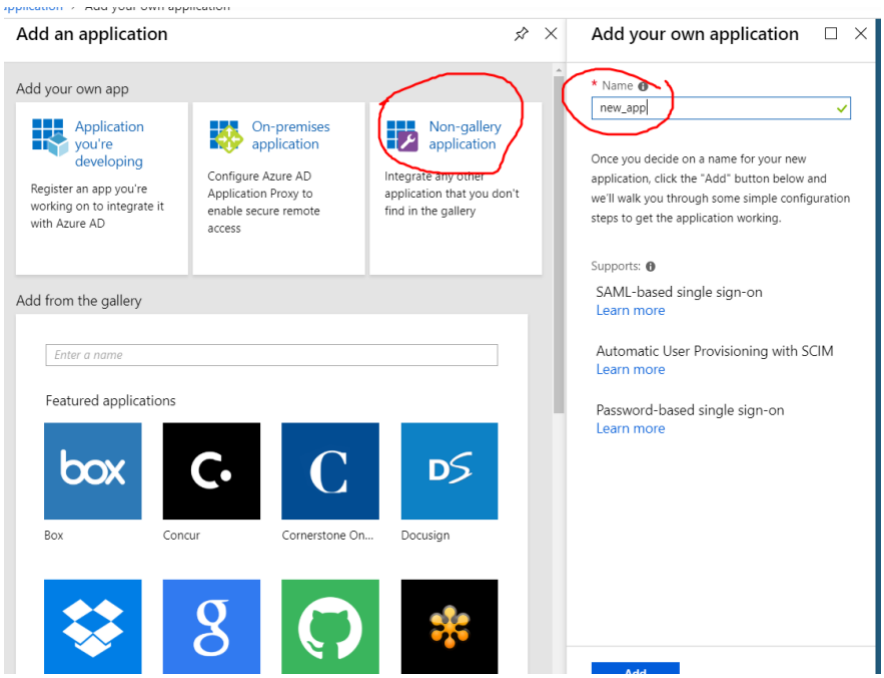

4. After naming the application, choose single sign-on and SAML.

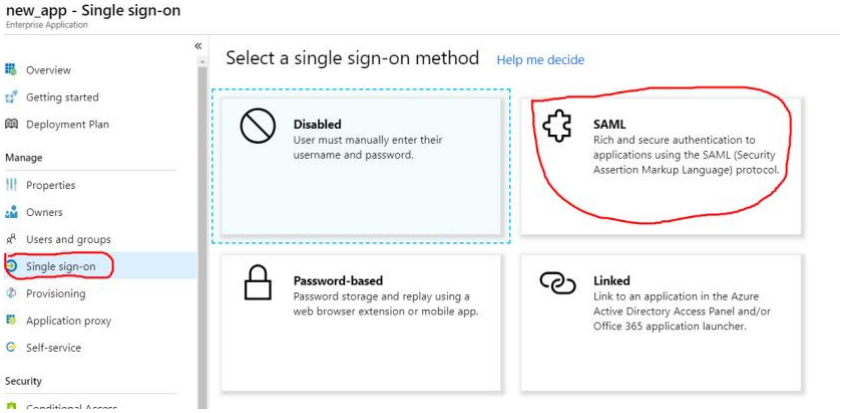

5. In the single sign-on setting, enter Identifier and reply URL using your own N2WS IP or URL.

Note:	There is a link at the top to change the appearance of this settings page.

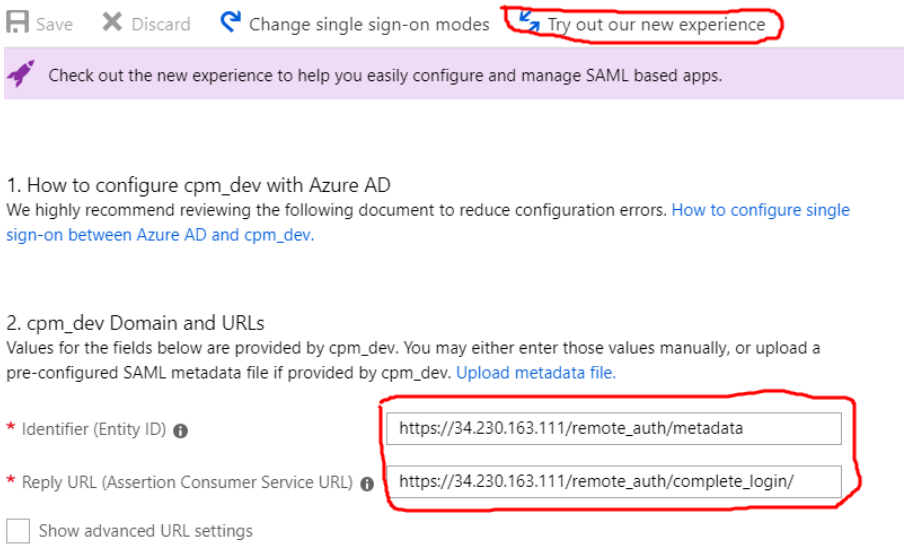

6. Ensure that the other attributes match. Download the certificate.

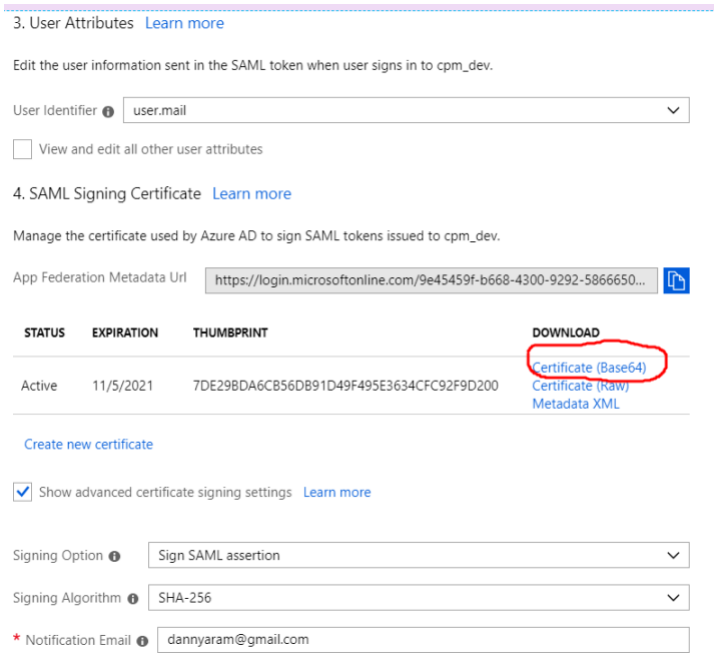

7. Save the new application.
8. In the main menu, go to **Users and groups** and select **+ Add user**.

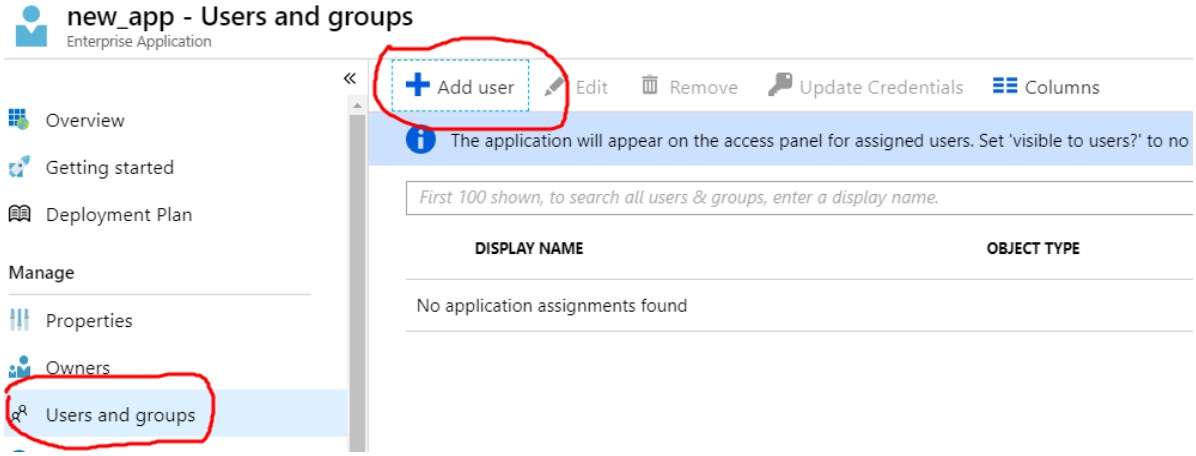

9. Select **Users and groups** and then select the group you created.

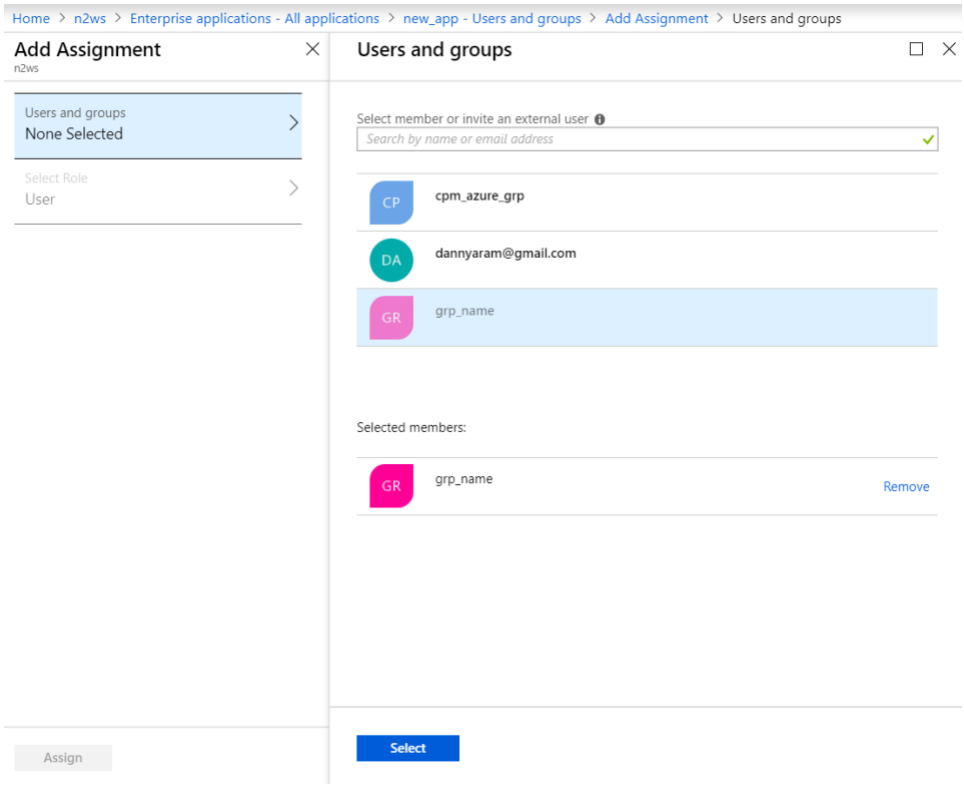

10. In your new application, choose **Single sign-on** and edit the attributes.

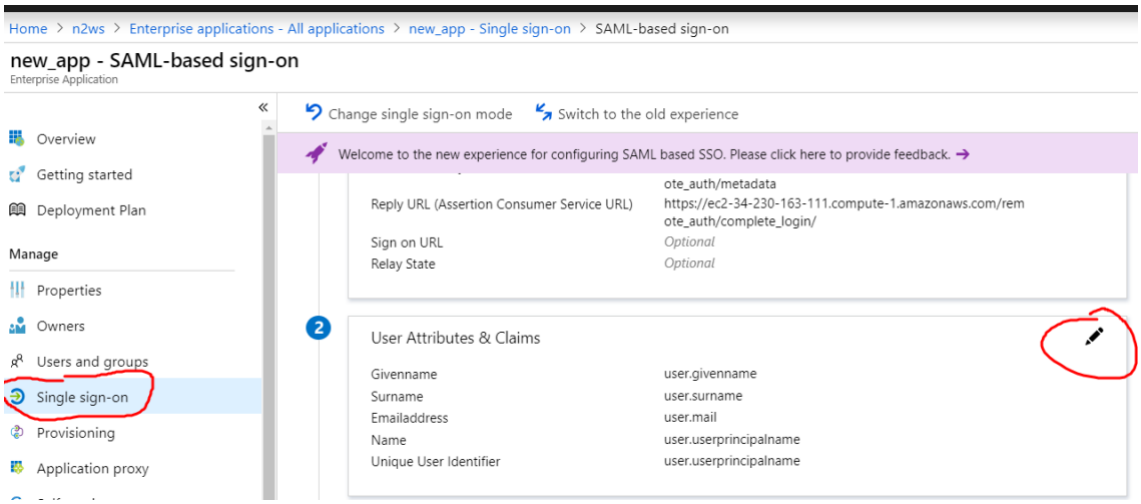

11. Choose to edit the **Name identifier** attribute and change the value to user.mail.

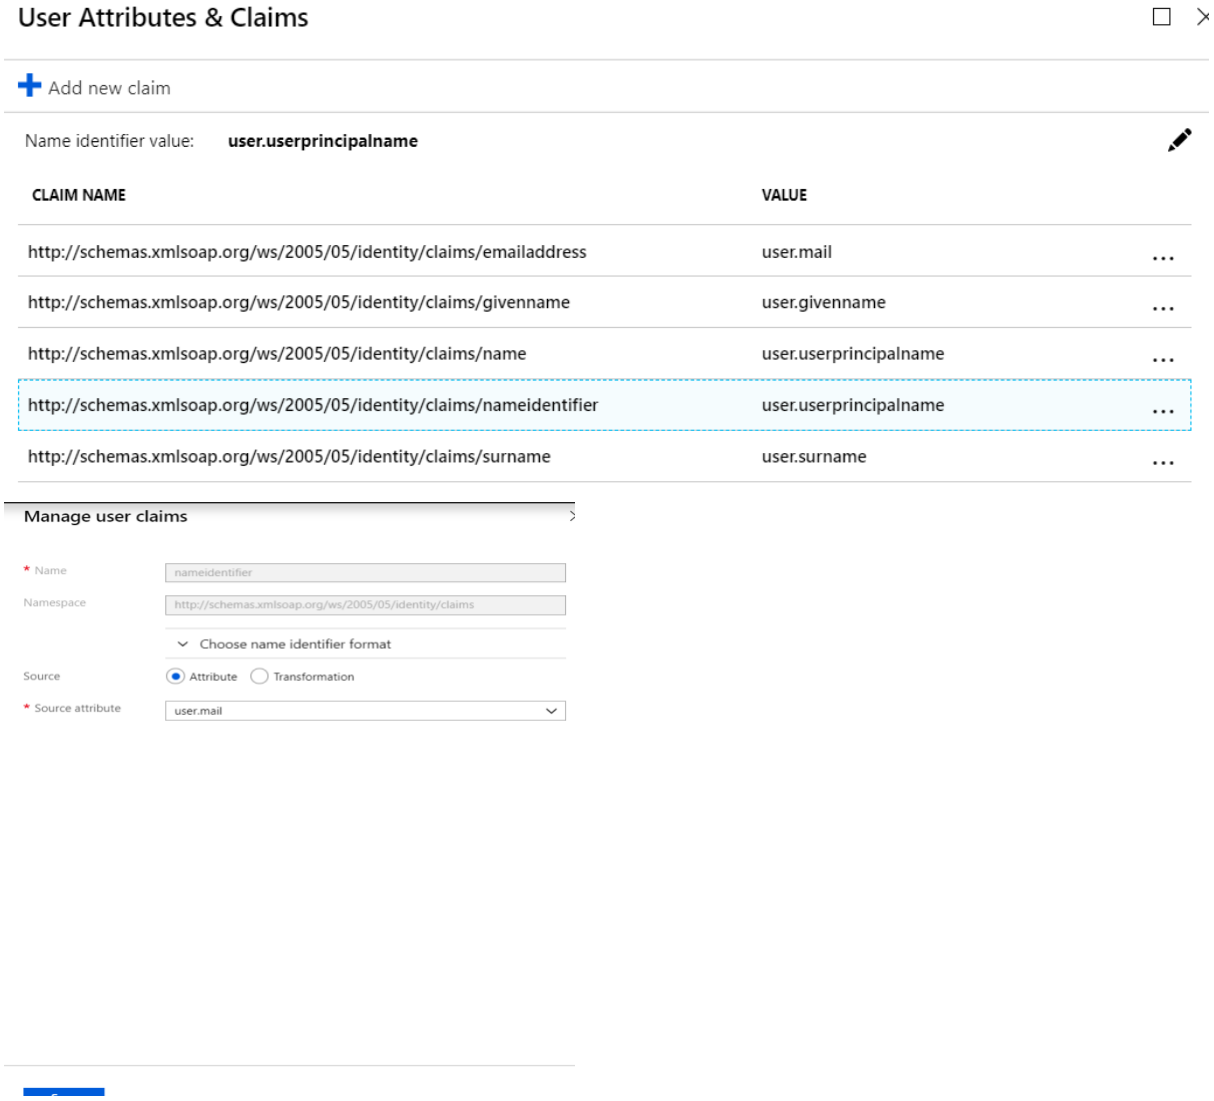

12. Add 2 new attributes.

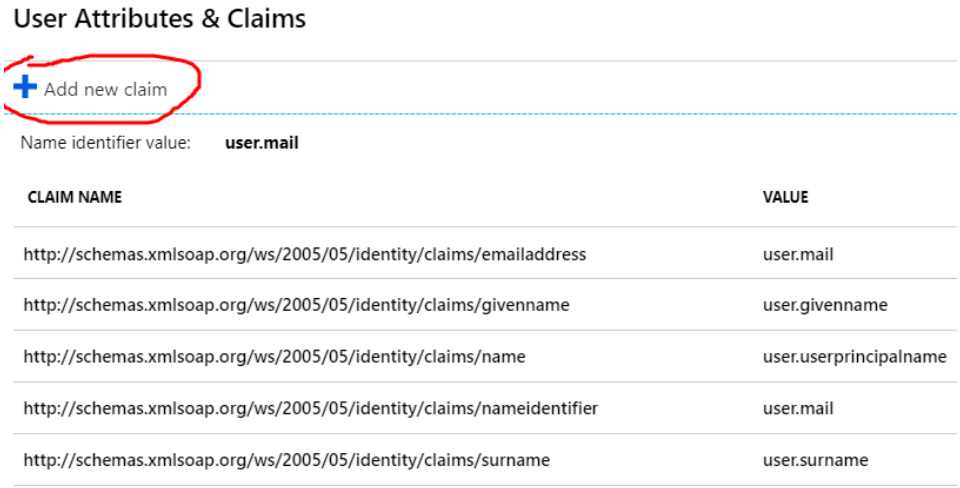

name: **cpm_user_groups**
value: **user_type=default;**
user_email=cpm@gmail.com;allow_file_level_recovery=default;
max_accounts=default;max_instances=default;max_independent_ebs_gib=default;ma

x_rds_gib=default;max_redshift_gib=default;
Change the parameters to meet N2WS needs:
name: **cpm_user_groups**
value: **cpm_<group name>**

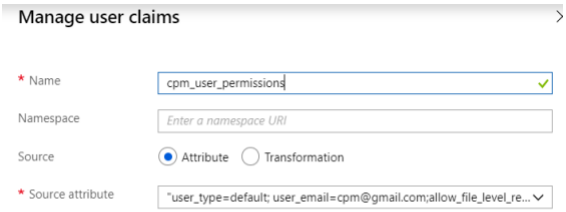

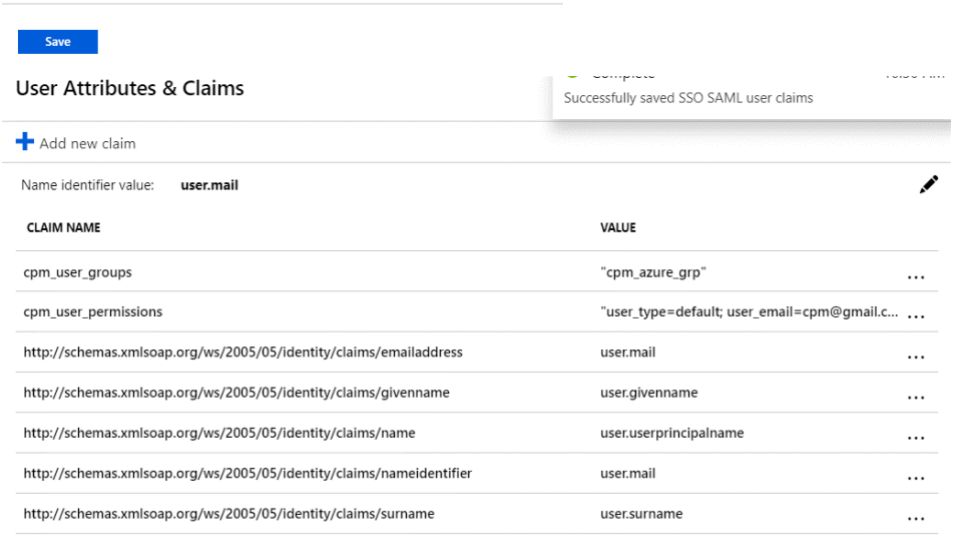

19.7.2 N2WS IdP Configuration

1. While still in the Azure AD settings, go to single sign-on and switch to new view:

2. Scroll down to section 4. These parameters will be used to configure the N2WS IdP settings.

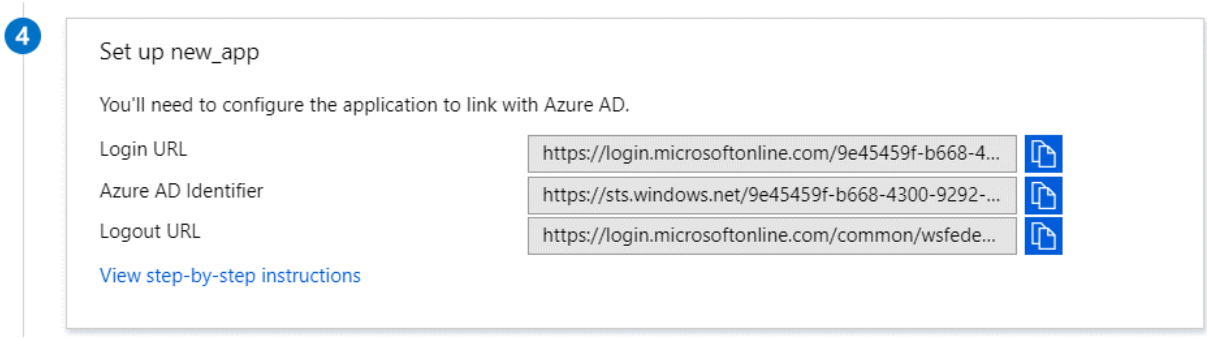

3. Switch to the N2WS **Identity Provider** page, select the provider, and select **Edit**.
4. In the **Settings** tab, complete the following:

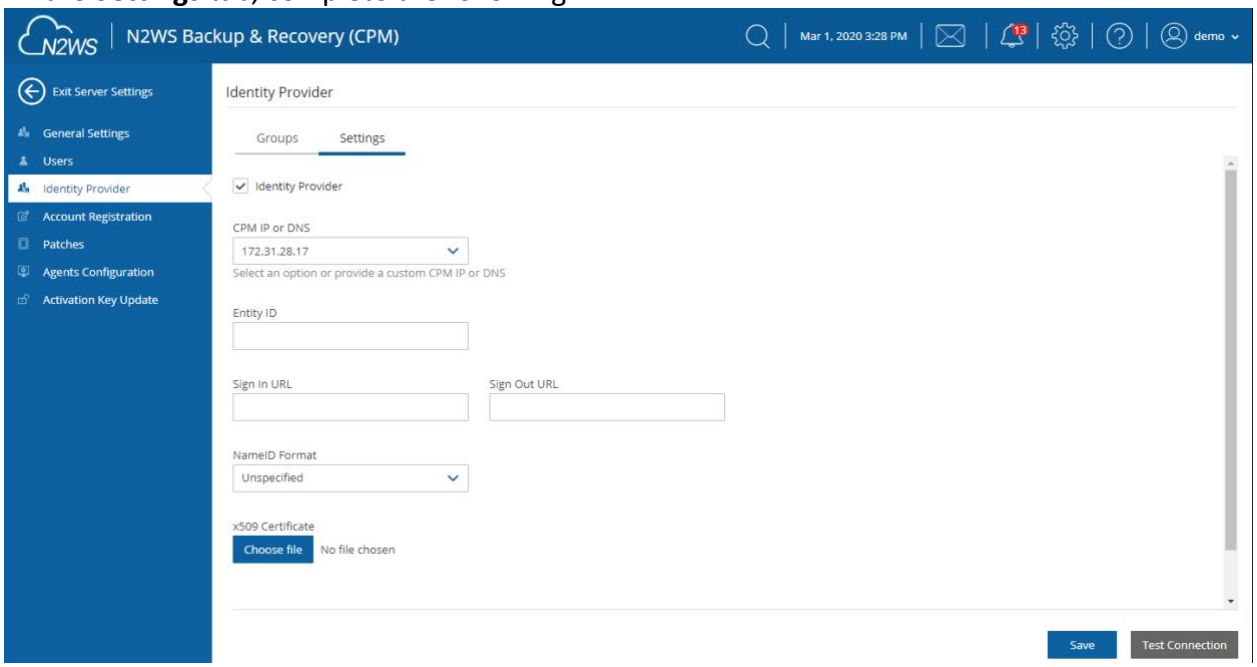

- **Entity ID -** Copy Azure AD Identifier.
- **Sign In URL** - Copy Login URL.
- **NameID format** - Select **Unspecified.**
- **x509 cert** - Upload the certificate downloaded in section 2.

5. Add a new group with the name of the group you added in the Azure Active Directory, without the **cpm** prefix. Select the **Groups** tab and then select **+ New** and add a name for the group.

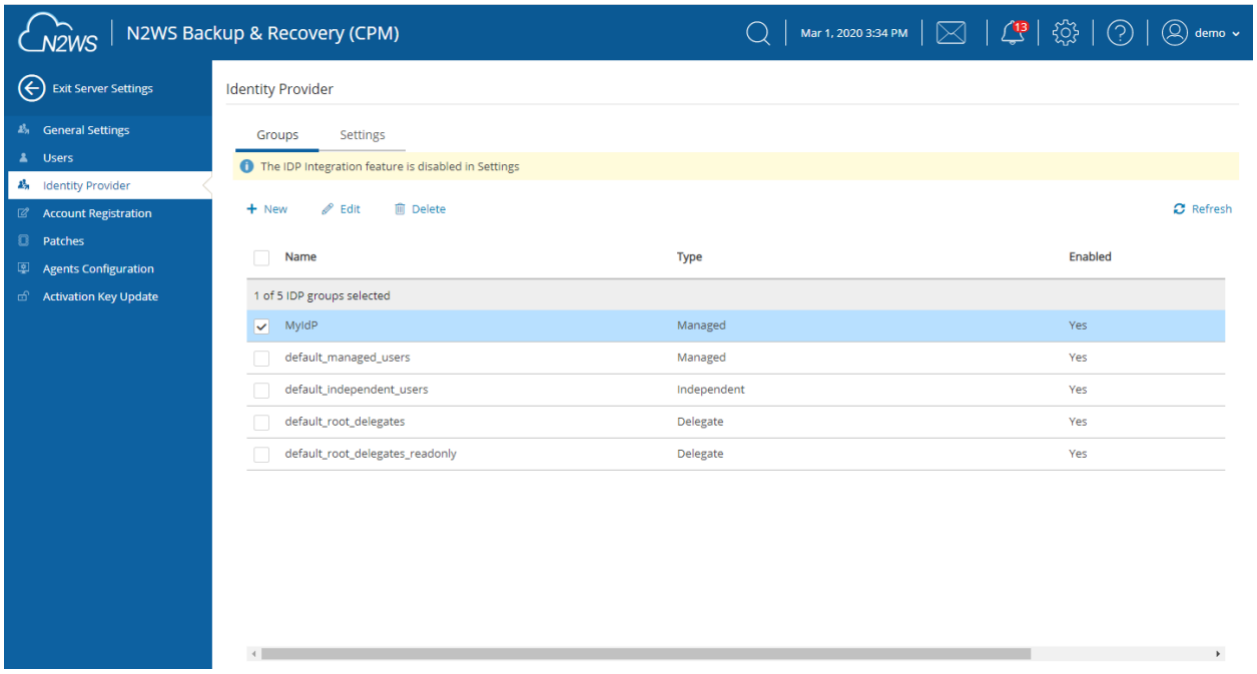

6. Select **Save**.
7. Return to the **Settings** tab and select **Test Connection**.

20 Configuring N2WS with CloudFormation

The process to configure N2WS to work with CloudFormation is a single stream that starts with subscribing to N2WS on the Amazon Marketplace and ends with configuring the N2WS server.

- N2WS provides a number of editions all of which support CloudFormation.
- An IAM role will automatically be created with minimal permissions and assigned to the N2WS instance.

1. Go to https://aws.amazon.com/marketplace
2. Search for N2WS.
3. Select CPM Edition to install:
   - Free Trial & BYOL
   - Advanced
   - Free
   - Standard
   - Enterprise

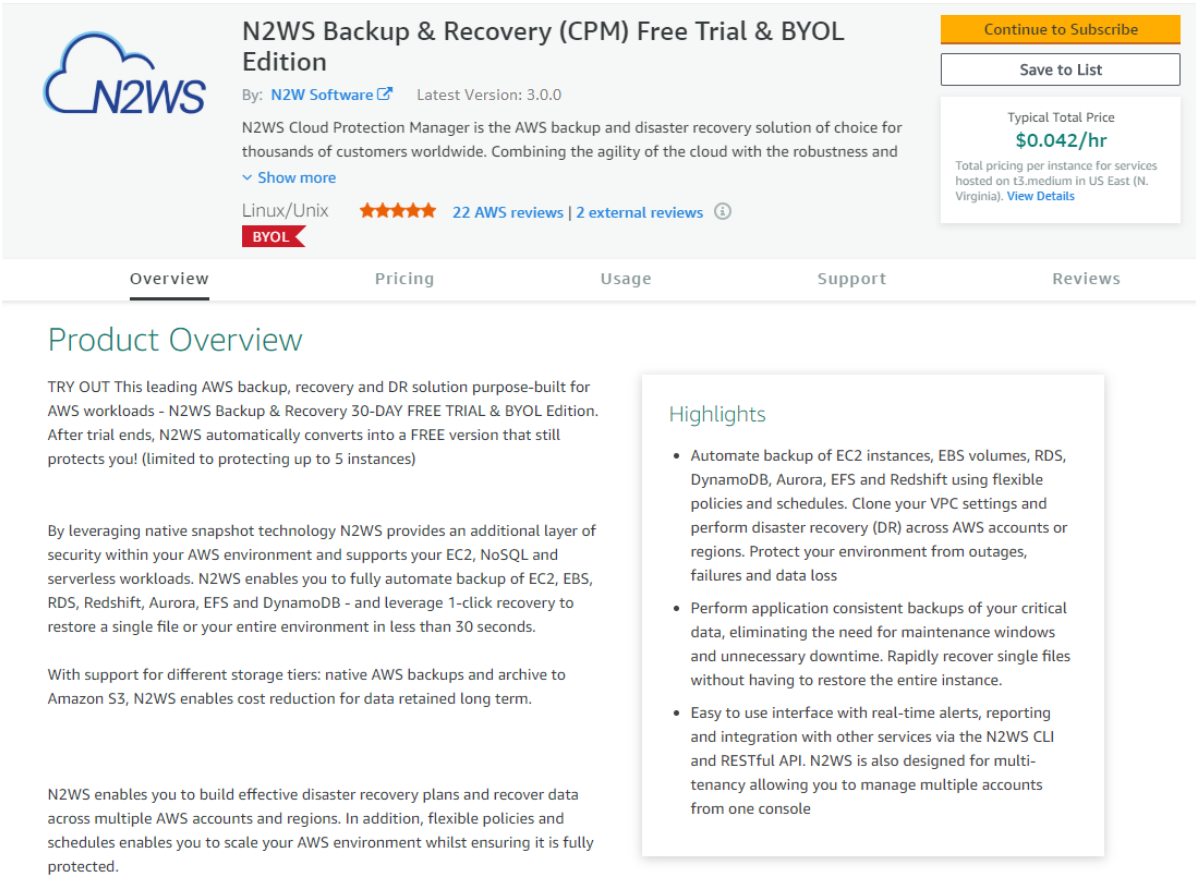

4. Select **Continue to Subscribe**. Log in and select **Accept Terms**.

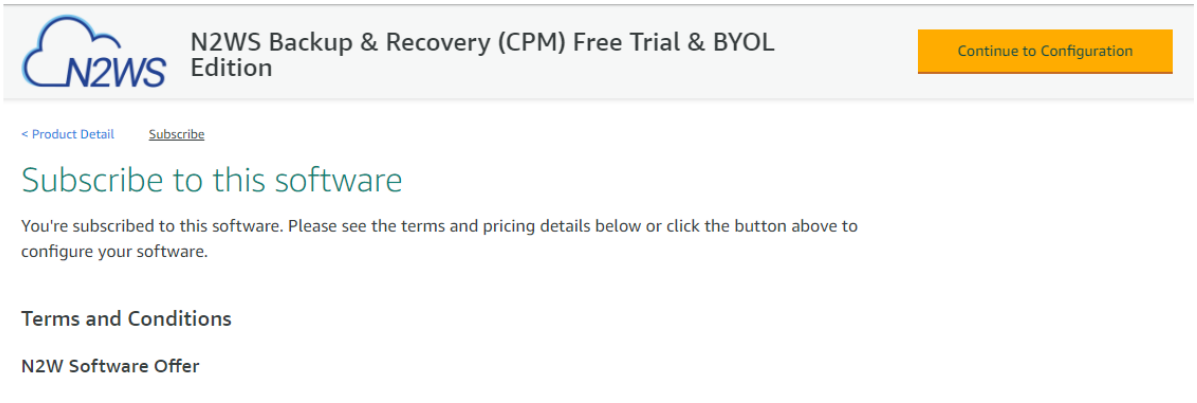

5. Select **Continue to Configuration**.

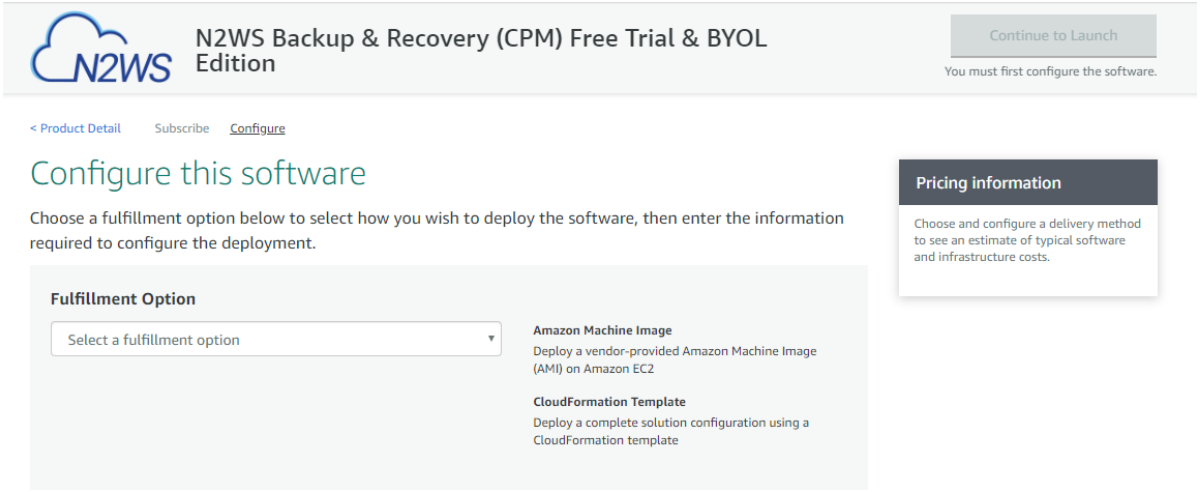

6. In the **Fulfillment Option** drop-down list, select **CloudFormation Template**. Select the relevant **Software Version** and **Region** and then select **Continue to Launch**.

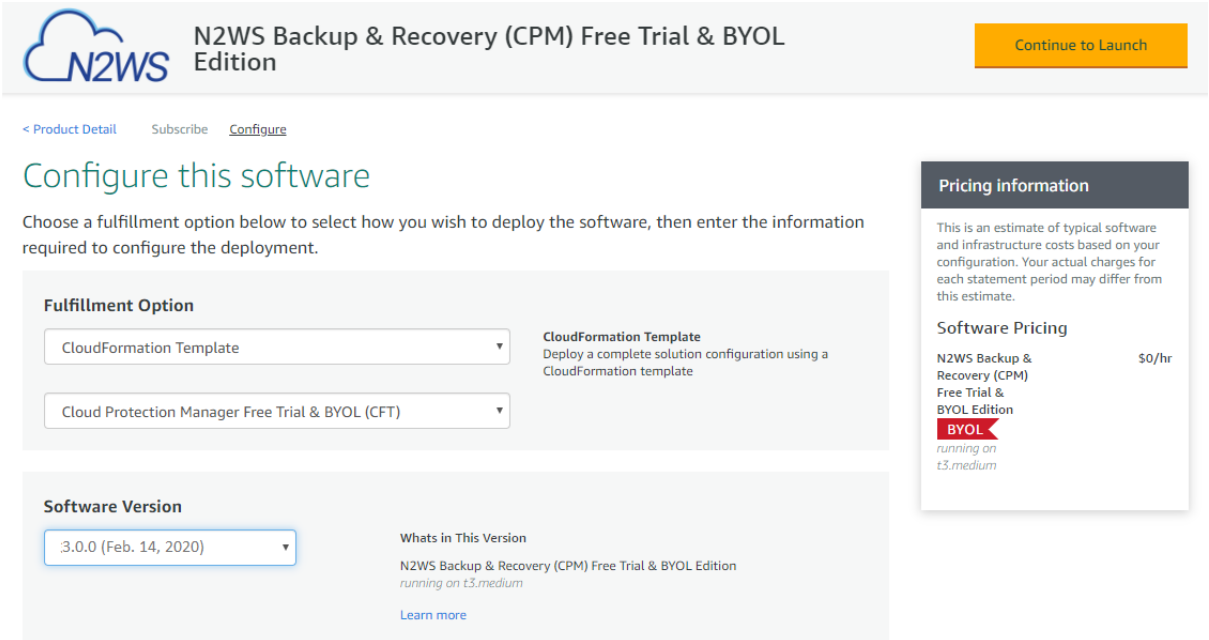

7. In the **Launch this software** page, select **Launch CloudFormation** in the **Choose Action** list and then select **Launch**.

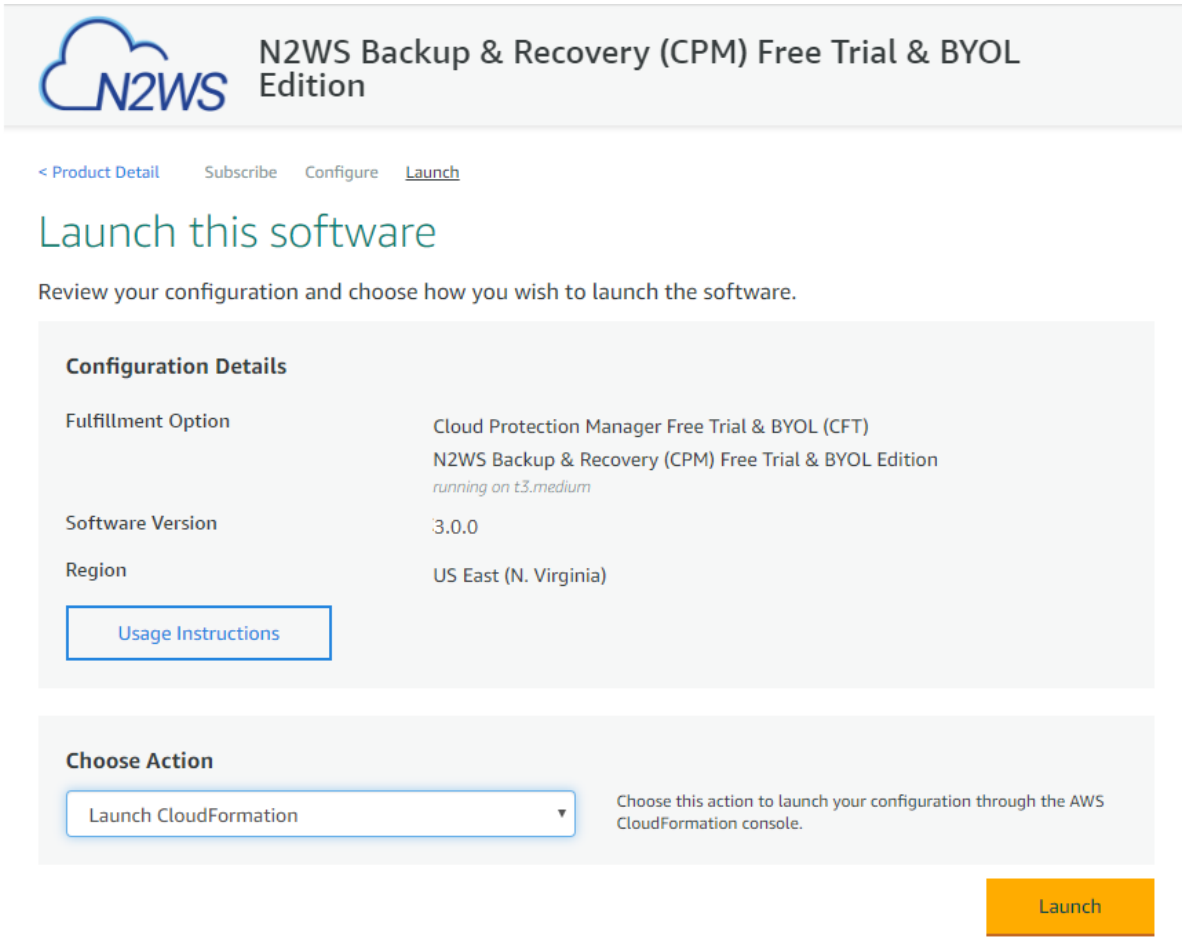

The **Create stack/Specify template** page opens.

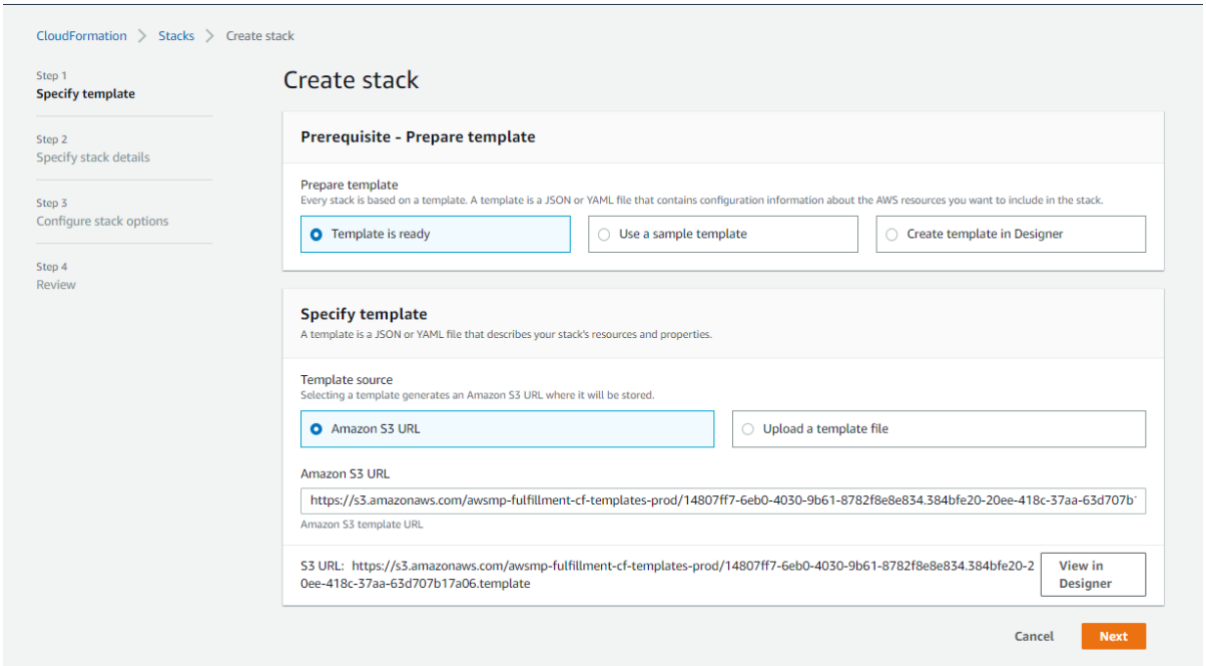

8. Under **Prepare template**, select **Template is ready**.
9. Under **Template source**, choose **Amazon S3 URL**. Select the default Amazon S3 URL and then select **Next**. The **Specify stack details** page opens.

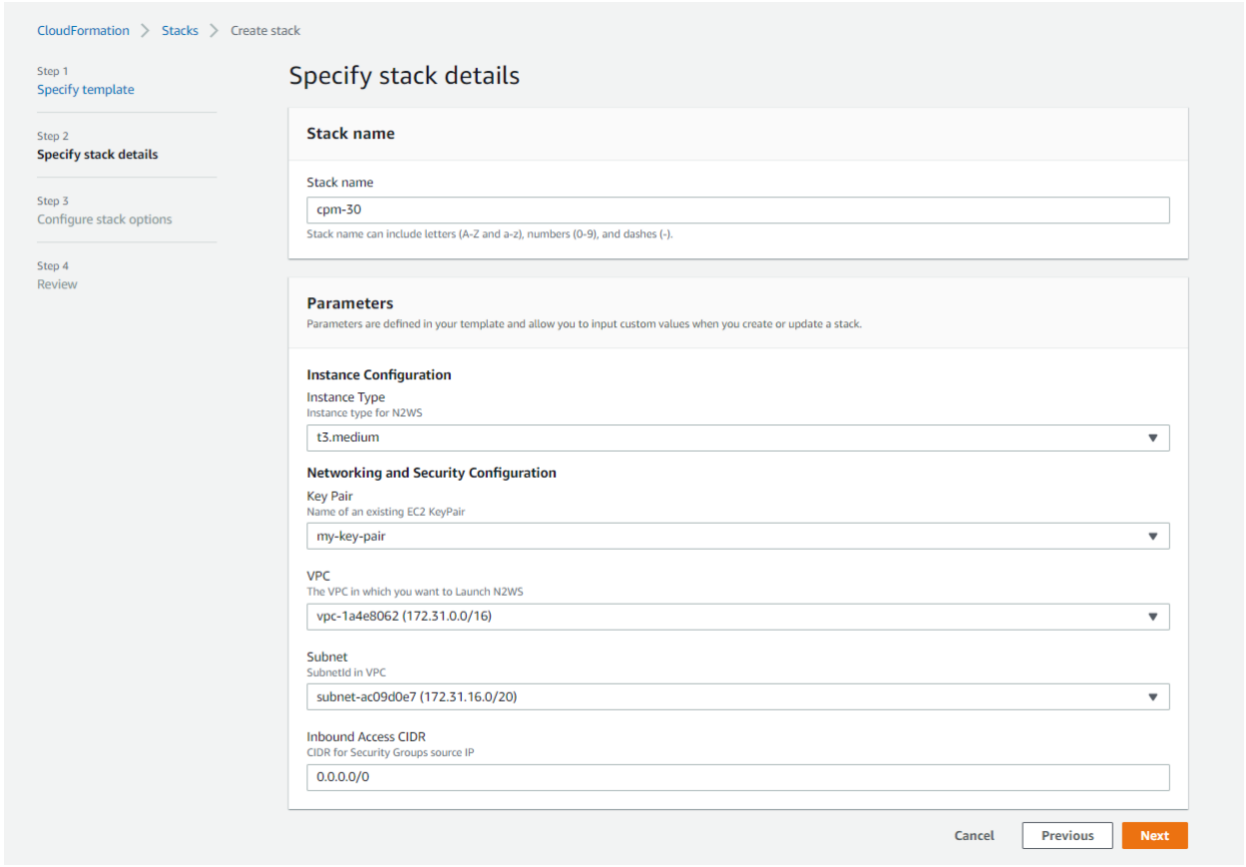

10. Complete the **Stack name** and **Parameters** sections.

    For **Inbound Access CIDR**, security groups act as a firewall for associated instances, controlling both inbound and outbound traffic at the instance level. Configuring **Inbound Access CIDR** allows you to add rules to a security group that enable you to connect to your Linux instance from your IP address using SSH:

    - If your IPv4 address is `203.0.113.25`, specify `203.0.113.25/32` to list this single IPv4 address in CIDR notation. If your company allocates addresses within a range, specify the entire range, such as `203.0.113.0/24`.
    - If you specify `0.0.0.0/0`, it will enable all IPv4 addresses to access your instance using SSH.
    - For further details, refer to “Adding a Rule for Inbound SSH Traffic to a Linux Instance” at https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/authorizing-access-to-an-instance.html

11. Select **Next**. The **Configure stack options** page opens.

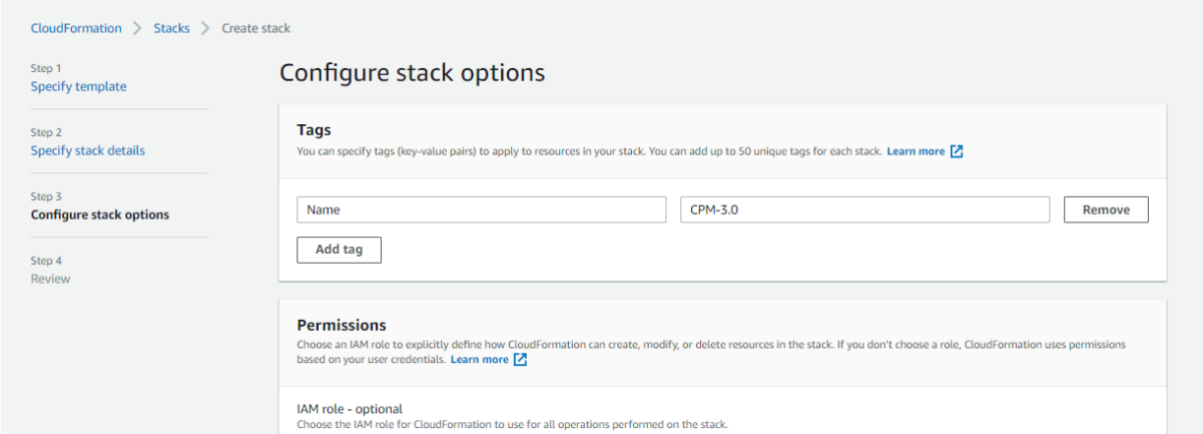

12. Complete the **stack options** and select **Next**. The **Review** page opens.

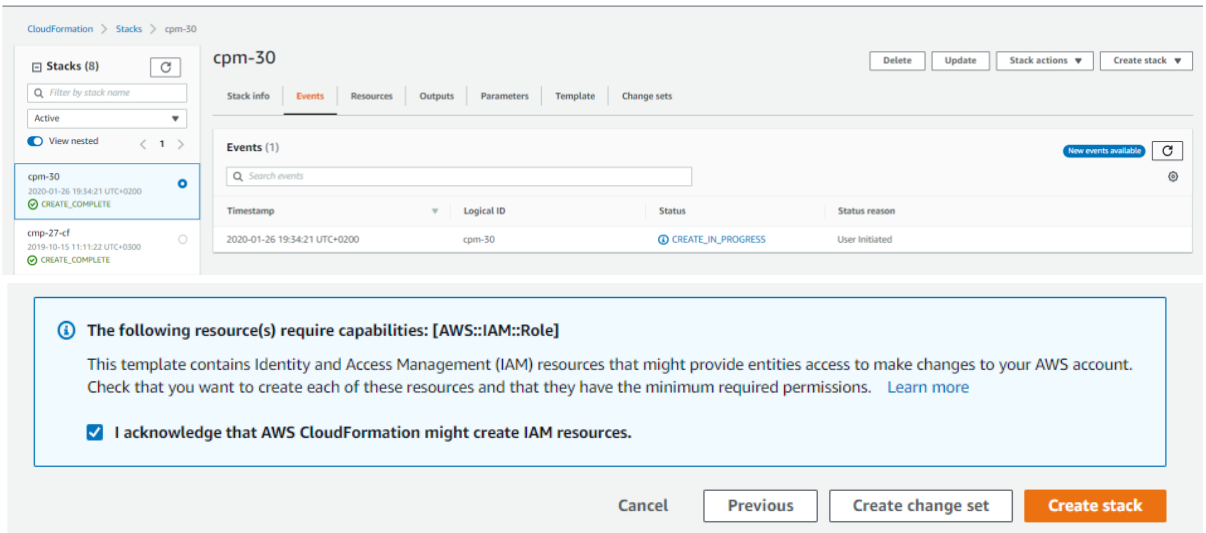

13. Select **I acknowledge that AWS CloudFormation might create IAM resources**, and then select **Create stack**. The **CloudFormation Stacks** page opens.
14. Select the new stack. The **Instances** page opens.
15. Select the instance. Copy the **Instance ID** value shown in the **Description** tab and select **Launch Instance**. The **N2WS Server Configuration** page opens.

Note:	Configure CPM with CloudFormation will fail where the requested Instance type is not supported in the requested Availability Zone. Retry your request, but do not specify an Availability Zone or choose us-east-1a, us-east-1b, us-east-1c, us-east-1d, or us-east-1f.

16. Continue configuring N2WS as in section 2.

21 Managing Snapshots with Lifecycle Policies

In addition to creating and managing EBS snapshots, N2WS can store backups in Simple Storage Service (S3) and S3 Glacier, allowing you to lower backup costs when storing backups for a prolonged amount of time. N2WS allows you to create a lifecycle policy, where older snapshots are automatically moved from high-cost to low-cost storage tiers. A typical lifecycle policy would consist of the following sequence:

1. Store daily EBS snapshots for 30 days.
2. Store one out of seven (weekly) snapshots in S3 for 3 months.
3. Finally, store a monthly snapshot in S3 Glacier for 7 years, as required by regulations.

Note: Storing snapshots in S3 in not supported for periods of less than 1 week.

Configuring a lifecycle management policy in N2WS consists the following sequence:

1. Defining how many EBS snapshots to keep.
2. Enabling and configuring Backup to S3.
3. Enabling and configuring Archive to S3 Glacier.

Refer to https://aws.amazon.com/s3/storage-classes for detailed S3 storage class information.

21.1 Using S3 with N2WS

Using the N2WS **Copy to S3** feature, you can:

- Define multiple folders, known as repositories, within a single S3 bucket
- Define the frequency with which N2WS backups are moved to a Repository in S3, similar to DR backup. For example, copy every third generation of a N2WS backup to S3.
- Define backup retention based on time and/or number of generations per Policy.
- N2WS stores backups in S3 as block-level incremental backups.

Note: Only ***one*** S3 operation is allowed for a policy at a time – Copy, Recovery, Archive, or retention Cleanup. For instance, an S3 Copy or S3 Recovery is not allowed when the S3 backup retention Cleanup is executing. If the S3 Cleanup process is running at the time of an S3 Copy or Recovery, you can abort the Cleanup process in order to allow the Copy or Recovery process to continue. See section 21.5.3.

Important: AWS Encryption at the bucket-level *must* be *enabled*.

Strongly Recommended:

- S3 buckets used by **Copy to S3** should *no*t be used by other applications**.**
- Versioning at the bucket level should be *disabled*.

Notes: Before continuing, consider the following:

- Copy to S3 currently supports only backups of Windows and Linux instances. RDS, DynamoDB, etc. are not supported.

- Independent volumes will be supported in a future release.

Note: Most N2WS operations related to the S3 repository (e.g. writing objects to S3, clean up, restoring, etc.) are performed by launching N2WS worker instances in AWS. The worker instances are terminated when their tasks are completed.

21.1.1 Limitations

Only copy of instance backups is supported.

- Copy to S3 is supported for weekly and monthly backup frequencies *only*. Daily backup copies to S3 are *not* supported.
- Copy of standalone volumes is not supported.
- Copy is not supported for other AWS resources that N2WS supports, such as RDS and Aurora.
- Snapshots consisting of 'AMI-only' cannot be copied to a S3 repository.
- The root volume of instances purchased from Amazon Marketplace, such as instances with product code, cannot be copied to S3. The data volumes of such instances, if they exist, will be copied.
- Backup records that were copied to S3 cannot be moved to Freezer.
- User cannot delete specific snapshots from S3 repository. S3 snapshots are deleted according to retention policy. In addition, users can delete all S3 snapshots of a specific policy, account or an entire repository. See below.
- A separate N2WS server, for example, one with a different "CPM Cloud Protection Manager Data" volume, cannot reconnect to an existing S3 repository.
- In order to use the Copy to S3 functionality the "`cpmdata`" policy must be enabled. See *N2WS User Guide* for details on enabling the "`cpmdata`" policy.
- Only a single S3 operation is possible on a policy at any given time. Additional executions of Copy to S3 backups will not occur if the previous execution is still running. Restore from S3 is always possible, except when Cleanup is running
- AWS accounts have a default limit to the number of instances that can be launched. Copy to S3 launches extra instances as part of its operation and may fail is the AWS quota is reached. See *N2WS User Guide* for details.
- Copy and Restore of volumes to/from regions different from where the S3 bucket resides may incur long delays and additional bandwidth charges.
- Instance names may not contain slashes (/) or backslashes (\) or the copy will fail.

21.1.2 Cost Considerations

N2W Software has the following recommendations to N2WS customers for help lowering transfer and storage costs:

- Lowering transfer fees:
  - When an 'N2WSWorker' instance is using a public IP (or NAT/IGW within a VPC) to access an S3 bucket within the same region/account, it results in network transfer fees.

- Using a VPC endpoint instead will enable instances to use their private IP to communicate with resources of other services within the AWS network, such as S3, without the cost of network transfer fees.
- For further information on how to configure N2WS with a VPC endpoint, see section Appendix A.

21.1.3 Overview of S3 and N2WS

The Copy to S3 feature is similar in many ways to the N2WS Disaster Recovery (DR) feature. When Copy to S3 is enabled for a policy, copying EBS snapshot data to S3 begins at the completion of the EBS backup, similar to the way DR works. Copy to S3 can be used simultaneously with DR feature.

21.1.4 Workflow for Using S3 with N2WS

1. Define an S3 Repository**.**
2. Define a Policy with a Schedule, as usual.
3. Configure the policy to include **Copy to S3** by selecting the **Lifecycle Management (Snapshot/S3/Glacier** tab. Turn on the **Backup to S3** toggle and complete the parameters.
4. If you are going to back up and restore S3 instances and volumes across accounts and regions, prepare a Worker Configuration using the **Worker Configuration** tab. See section 22.
5. Use the **Backup Monitor** and **Recovery Monitor,** with some additional controls, to manage S3 snapshots as usual.

21.2 The S3 Repository

21.2.1 Configuring an S3 Repository

There can be multiple repositories in a single AWS S3 bucket.

1. In N2WS, select the **S3 Repositories** tab.

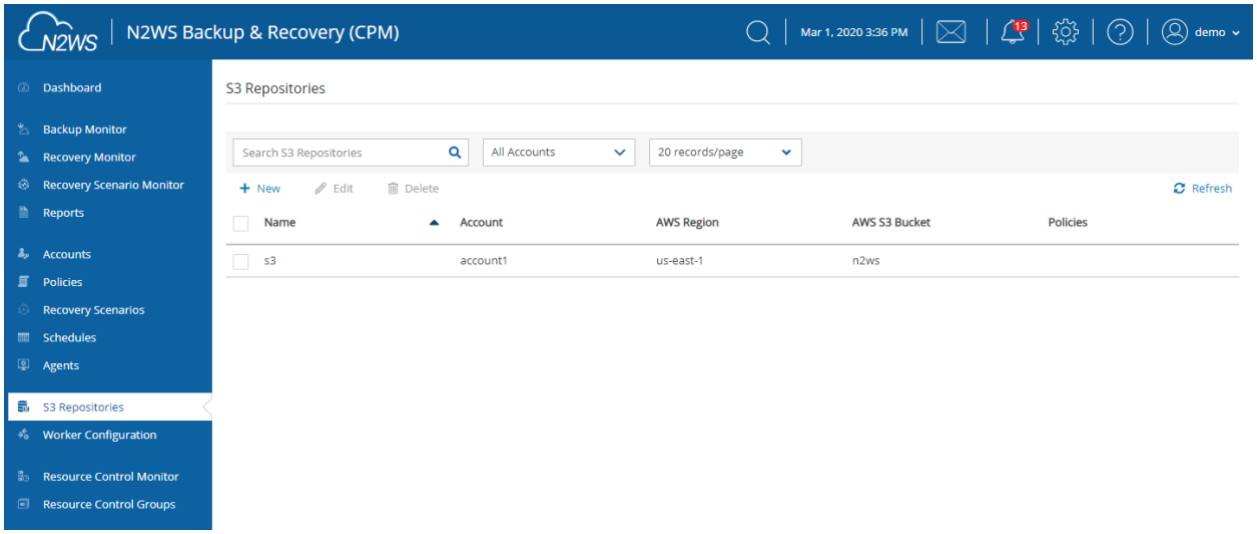

2. Select **+ New**.

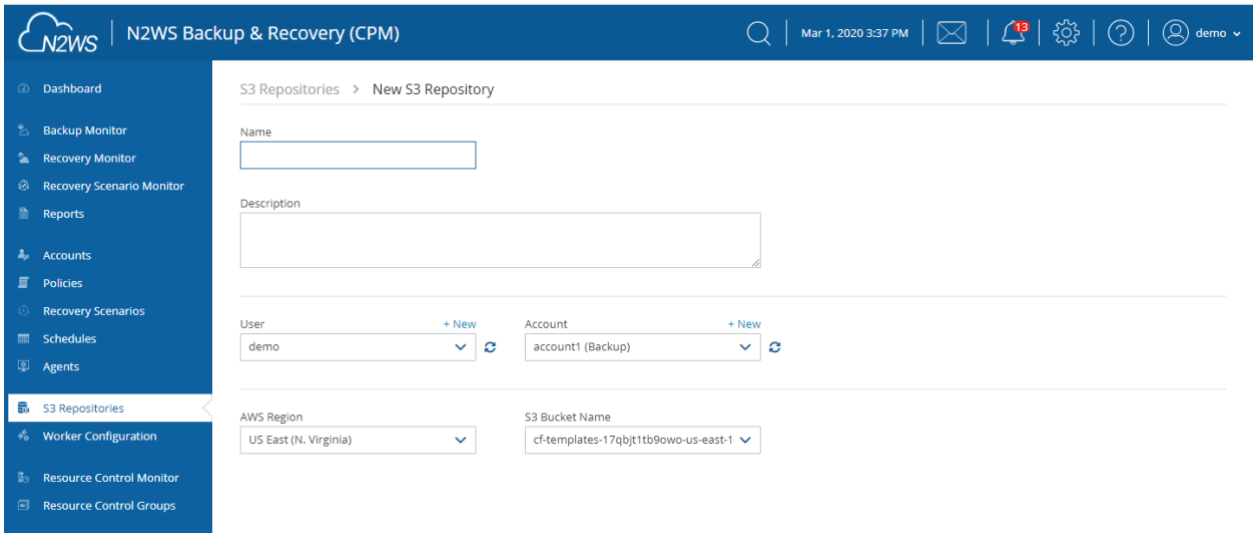

3. In the **New S3 Repository** screen, complete the following information:
   - **Name** - Type a unique name for the new repository, which will also be used as a folder name in the AWS bucket. Only alphanumeric characters and the underscore are allowed.
   - **Description** - Optional brief description of contents of repository.
   - **User** – Select the user in the list.
   - **Account** - Select the account that has access to the S3 bucket.
   - **AWS Region** - Select the region in which the S3 bucket is located.
   - **S3 Bucket Name** - Type the name of the S3 bucket that exists in this region.

   Note: AWS encryption must have been enabled for the bucket.

4. When complete, select **Save**.

21.2.2 Deleting an S3 Repository

You can delete all snapshots copied to a specific S3 repository.

Note: Deleting a repository is not possible when the repository is used by a policy. You must change any policy using the repository to a different repository before the repository can be deleted.

1. Select the **Repositories** tab and then select a repository.

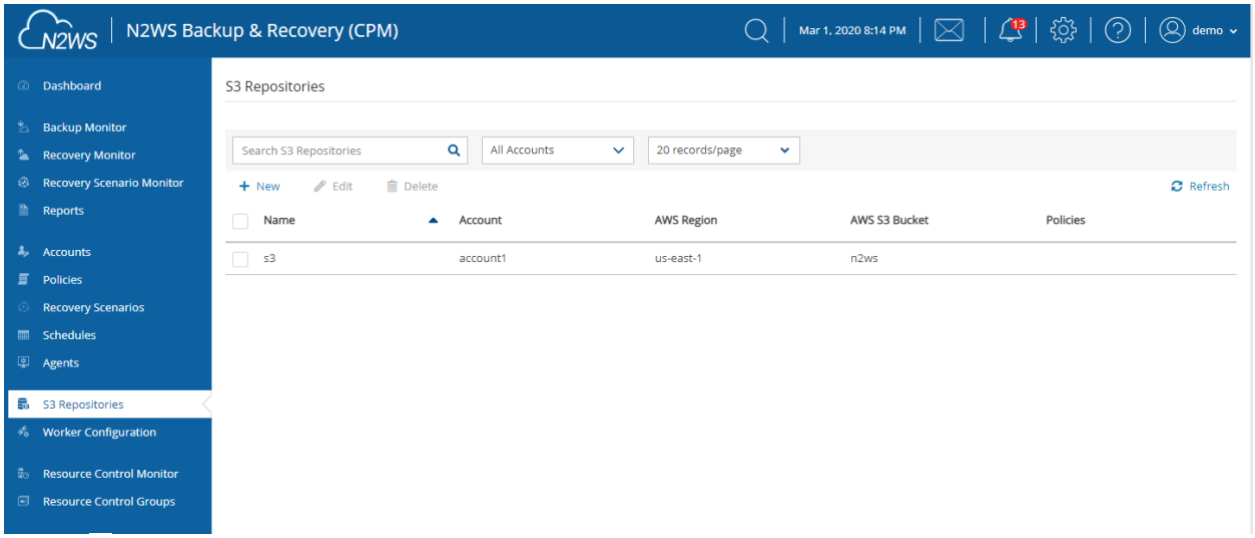

2. Select **Delete**.

21.3 The S3 Policy

21.3.1 Configuring a Policy to Backup to S3

Configuring a Policy for Copy to S3 backups includes definitions for the following:

- Name of the S3 Repository defined in N2WS.
- Interval of AWS snapshots to copy.
- Snapshot retention policy.

It is possible to retain a backup based on both time and number of generations copied. If both Time Retention (**Keep backups in S3 for at least** x time) and Generation Retention (**Keep backups in S3 for at least** x **generations**) are enabled, both constraints must be met before old snapshots are deleted or moved to Glacier, if enabled.

For example, when the automatic cleanup runs:

- If Time Retention is enabled for 7 days and Generation Retention is disabled, S3 snapshots older than 7 days are deleted or archived.
  If **Run ASAP** is executed 10 times in one day, none of the snapshots would be deleted until they are more than 7 days old.
- If Generation Retention is enabled for 4 and Time Retention is disabled, the 4 most recent S3 snapshots are saved.
- If Time Retention is enabled for 7 days and Generation Retention is enabled for 4 generations, a single S3 snapshot would be deleted, or archived, after 7 days if the number of generations had reached 5.

1. In the left panel, select the **Policies** tab.
2. Select a Policy and then select **Edit**.

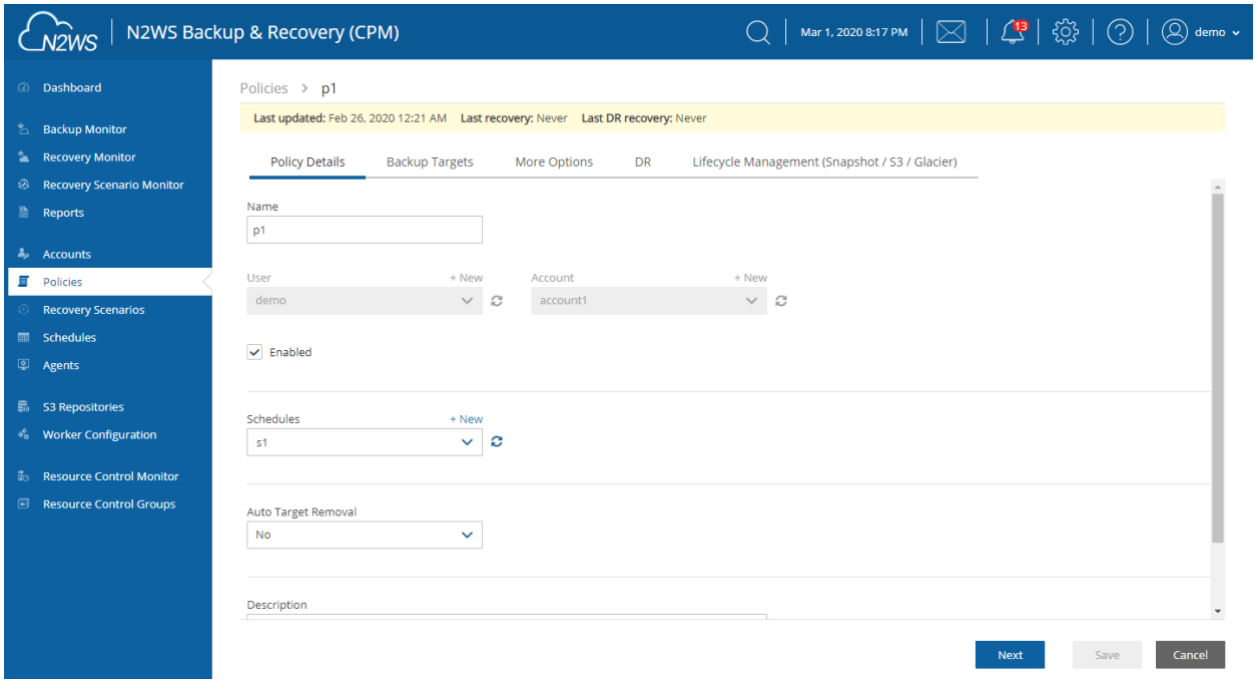

3. Select the **Lifecycle Management** tab.

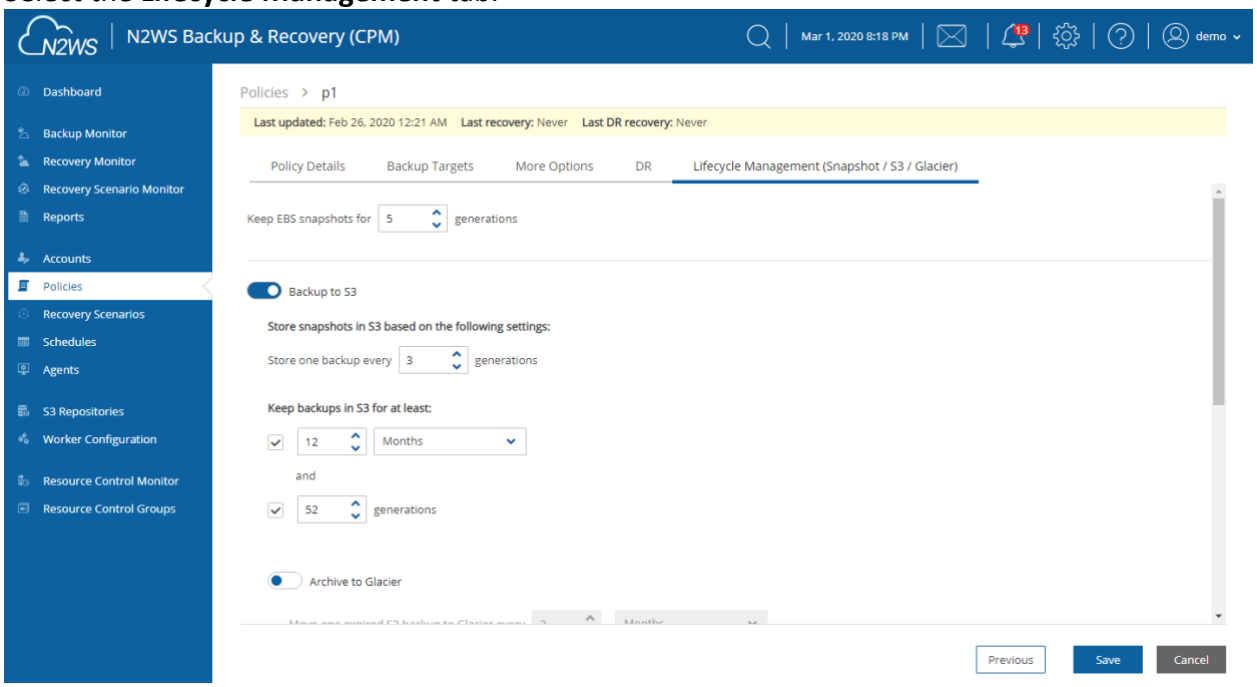

4. Complete the following fields:
   - **Backup to S3** – By default, Backup to S3 is disabled. Turn the toggle on to enable.
   - **S3 Repository** - Select the Repository in the S3 bucket to copy your backup to.
   - In the **Store one backup every** X **generations** list, select the frequency for storing backups in S3.
   - Specify retention: To retain backups based on either the length of time or the number of generations alone, clear the check box for the other parameter. By default, backups are retained in S3 based on both the length of time and number of generations.

- **Store one backup every** x **generations**: Select the number of generations between backups to move. For example, if **Store one backup every** is 3, move every 3rd N2WS backup to S3.
- **Keep backups in S3 for at least:** - Select the length of time and/or the number of generations check boxes to enable the method(s) for determining how long the backups should be kept in S3. Select the relevant parameters in the respective lists.

5. In the **Storage settings** section, choose the following parameters:
   - Select the **Target repository** in the S3 bucket to move the backup to, or select **+ New** to define a new repository. If you define a new repository, select Refresh before selecting.
   - Choose a **S3 Storage Class** that meets your needs:
     - **Standard** – (Frequent Access) - For frequent access and backups.
     - **Infrequent Access** – For data that is accessed less frequently.
     - **Intelligent Tiering** – Automatic cost optimization for S3 copy.
       Intelligent Tiering incorporates the Standard (Frequent Access) and Infrequent Access tiers. It monitors access patterns and moves objects that have not been accessed for 30 consecutive days to the Infrequent Access tier. If the data is subsequently accessed, it is automatically moved back to the Frequent Access tier. For more information about Amazon S3 Storage Classes, see https://aws.amazon.com/s3/storage-classes/.

   Notes: Storage charges:
   - S3 Infrequent Access and Intelligent Tiering have minimum storage duration charge.
   - S3 Infrequent Access has per GB retrieval fee.

   For additional information, refer to AWS S3 documentation.

6. Select **Save.**

21.3.2 Recovering an S3 Backup

You can recover an S3 backup to the same or different regions and accounts.

Note: 'Marked for deletion' snapshots can no longer be recovered.

1. Select the **Backup Monitor** tab.
2. Select a backup and then select **Recover**.
3. In the **Restore from** drop-down list of the **Recover** screen, select the name of the **S3 Repository** to recover from. If you have multiple N2WS accounts defined, you can choose a different target account to recover to.

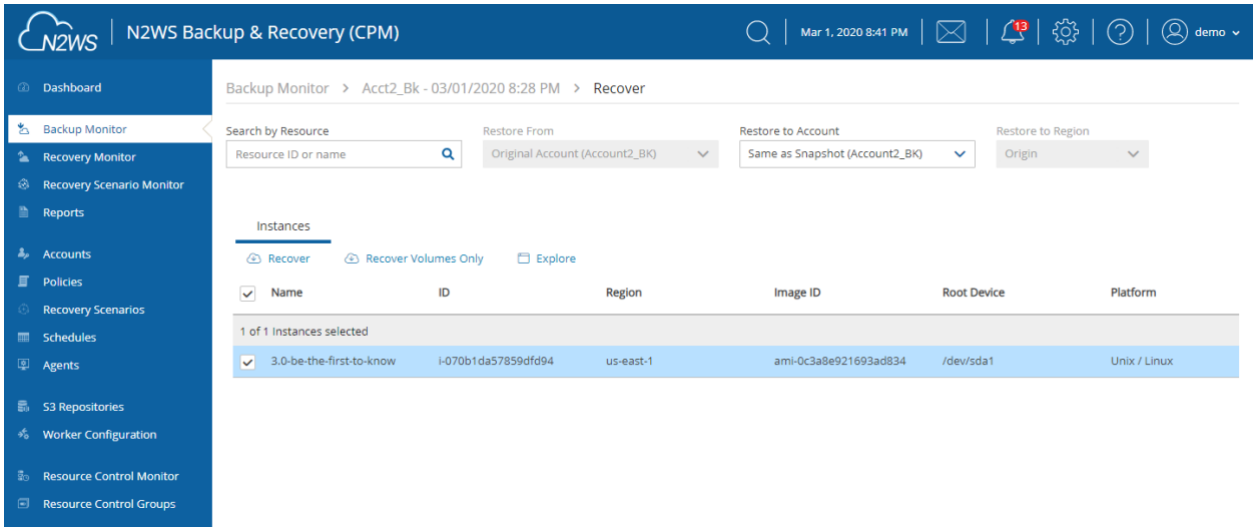

4. In the **Restore to Region** drop-down list, select the Region to restore the S3 copy to. The source Region of the S3 copy is displayed in the **Region** column.
5. If you select **Recover Volumes Only**, the **Volume Recovery** screen opens and you can:
   - **Explore** and select volumes for recovery
   - Define **Attach Behaviour**
   - Define the **AWS Credentials** for access
   - Configure a **Worker** in the **Worker Configuration** tab
   - **Clone** a **VPC**

Note: If recovering an S3 instance, you can specify the recovery encryption key:

- If **Use Default Volume Encryption Keys** is enabled, the recovered volumes will have the default key of each encrypted volume.
- If **Use Default Volume Encryption Keys** is disabled, all encrypted volumes will be recovered with the same key that was selected in the **Encryption Key** list.

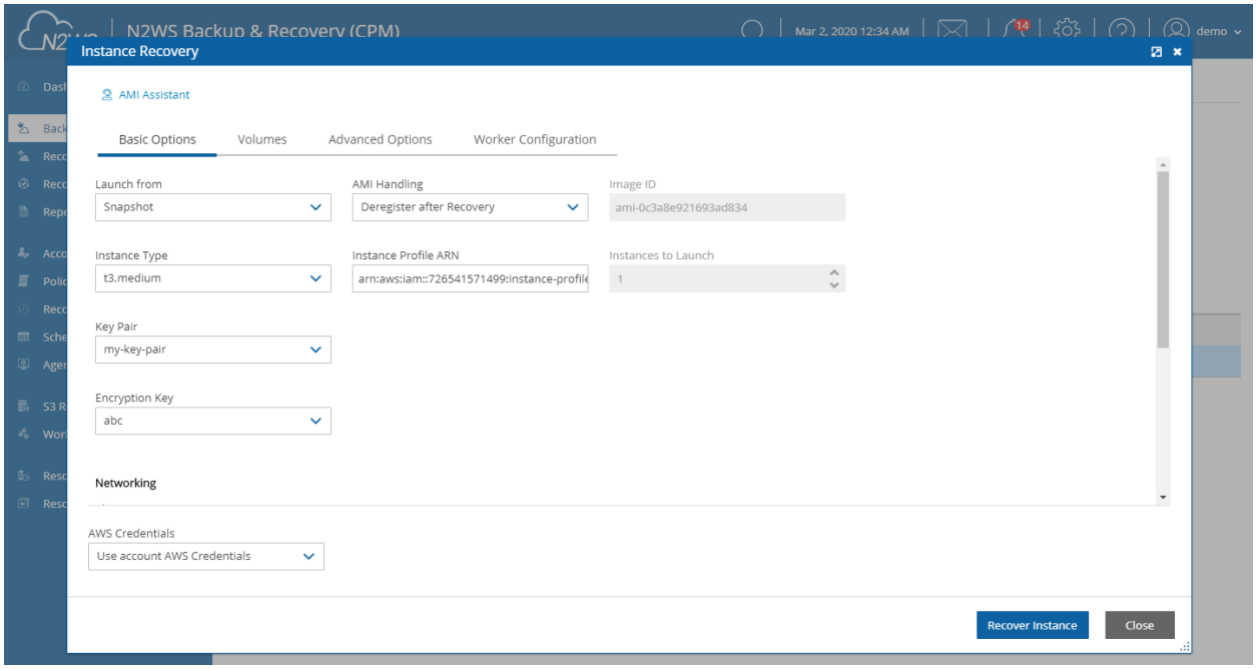

6. If you selected **Recover**, the **Instance Recovery** screen open to the **Basic Options** tab
   a. Change the **Basic Options** default values as necessary.
   b. Continue defining the Recovery options in the **Volumes** and **Advanced Options** tabs. See sections 10.2.2 and 10.2.3 respectively.
   c. If a worker has not been configured or assembled by N2WS, the **Worker Configuration** tab will appear next to the **Advanced Options**. Complete the form as necessary for the current recovery.

   Note: If you choose '**Any**' in the **Subnet** drop-down list, N2WS will automatically choose a subnet that is in the same Availability Zone as the one you are restoring to. If you choose a specific subnet that is not in the same Availability Zone as the one you are restoring to, you will have to choose a different subnet from the **Subnet** drop-down list.

   d. Select **Recover Instance**.
7. If you selected **Recover Volumes Only**:

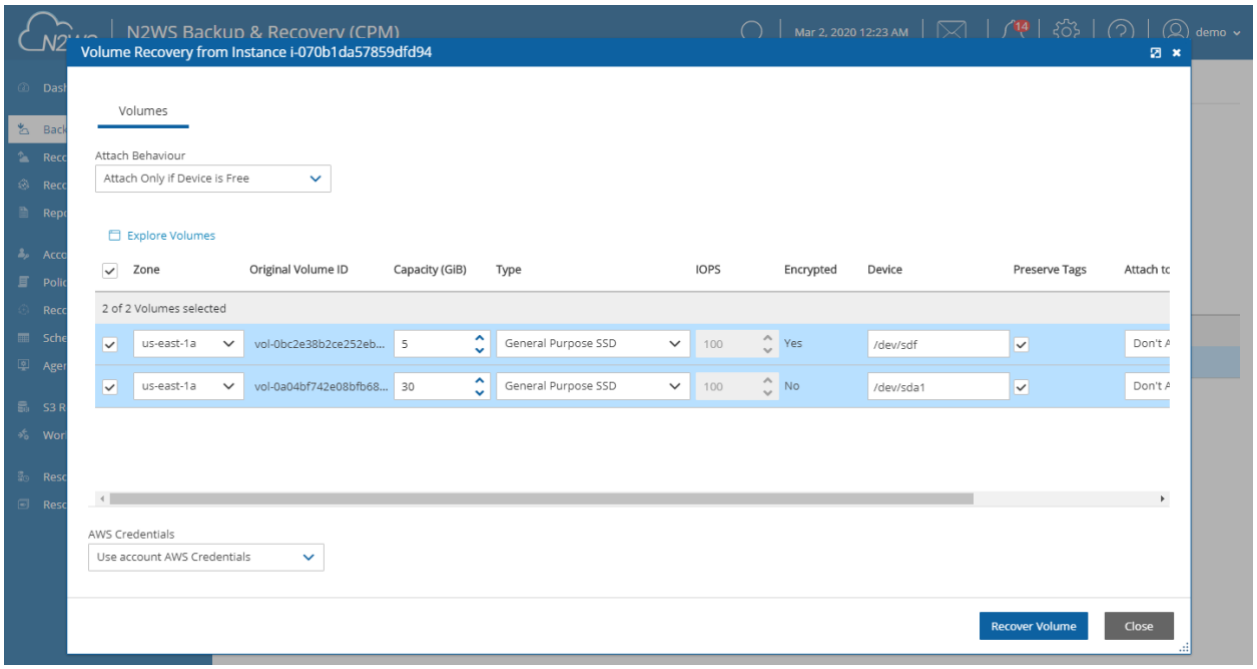

   a. Change the default values as necessary. In the **Attach Behaviour** drop-down list, select the appropriate behavior for the recovery.
      - Attach only if Device is Free
      - Switch Attached Volumes
      - Switch Attached Volumes and Delete Old Ones
   b. If a worker has not been configured or assembled by N2WS, select the **Worker Configuration** tab. Complete the form as necessary for the current recovery.

Note: If you choose '**Any**' in the **Subnet** drop-down list, N2WS will automatically choose a subnet in the same Availability Zone as the one you are restoring to. If you choose a specific subnet that is not in the same Availability Zone as the one you are restoring to, you will have to choose a different subnet from the **Subnet** drop-down list.

8. To follow the progress of the recovery, select the **Open Recovery Monitor** link in the 'Recovery started' message Recovery Started (Open Recovery Monitor) at the top right corner, or select the **Recovery Monitor** tab.

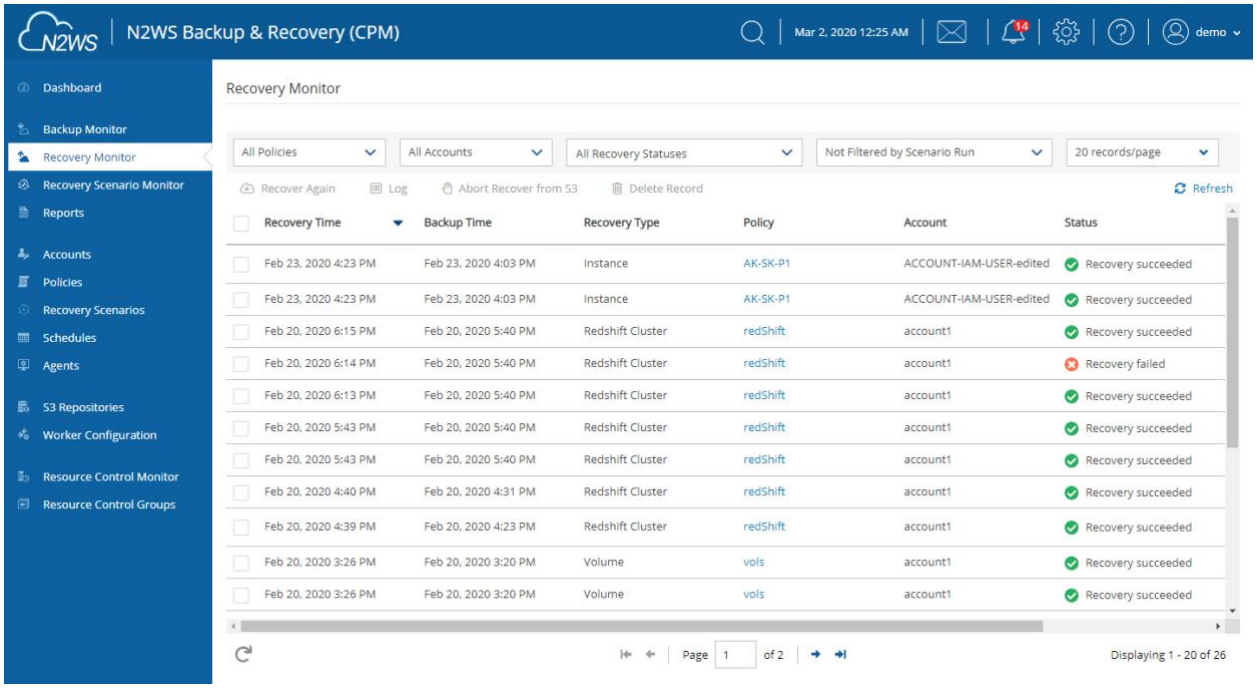

9. To abort a recovery in progress, in the **Recovery Monitor**, select the recovery item and then select **Abort Recovery from S3**.

21.3.3 Forcing a Single Full Copy

By default, Copy to S3 is performed incrementally for data modified since the previous snapshot was stored. However, you can force the copy of the full data for a single iteration to your S3 Repository. While configuring the **Backup Targets** for a policy with Copy to S3, select **Force a single full Copy**. See section 4.2.3.

Note: This option is only available for Copy to S3.

21.3.4 Changing the S3 Retention Rules for a N2WS Policy

You can set a different retention rules in each Policy.

To update the S3 retention rules for a policy:

1. In the **Policies** column, select the target policy.
2. Select the **Lifecycle Management** tab.
3. Update the **Keep backups in S3 for at least** lists for time and generations, as described in section 21.3, and select **Save**.

21.4 The Glacier Archive

21.4.1 Archiving Snapshots to S3 Glacier

Amazon S3 Glacier and S3 Glacier Deep Archive provide comprehensive security and compliance capabilities that can help meet regulatory requirements, as well as durable and extremely low-cost data archiving and long-term backup.
CPM allows customers to use the Amazon Glacier low-cost cloud storage service for data with longer retrieval times.
The CPM can now backup your data to a cold data cloud service on Amazon Glacier by moving infrequently accessed data to archival storage to save money on storage costs.

Notes: S3 is a better fit than AWS' Glacier storage where the customer requires regular or immediate access to data.

Recommendations:

- Use Amazon S3 if you need low latency or frequent access to your data.
- Use Amazon S3 Glacier if low storage cost is paramount, and you do not require millisecond access to your data.

21.4.2 Pricing

Following are some of the highlights of the Amazon pricing for Glacier:

- Amazon charges per gigabyte (GB) of data stored per month on Glacier.
- Objects that are archived to S3 Glacier and S3 Glacier Deep Archive have a minimum 90 days and 180 days of storage, respectively.
- Objects deleted before 90 days and 180 days incur a pro-rated charge equal to the storage charge for the remaining days.

For more information about S3 Glacier pricing, refer to sections 'S3 Intelligent – Tiering' / 'S3 Standard-Infrequent Access' / 'S3 One Zone - Infrequent Access' / 'S3 Glacier' / 'S3 Glacier Deep Archive' at https://aws.amazon.com/s3/pricing/

21.4.3 Configuring a Policy to Archive to S3 Glacier

To configure archiving S3 backups to Glacier:

1. From the left panel, in the **Policies** tab, select a **Policy** and then select **Edit.**
2. Select the **Lifecycle Management (Snapshot / S3 / Glacier)** tab**.** See section 21.3.
3. Turn on the **Archive to Glacier** toggle.

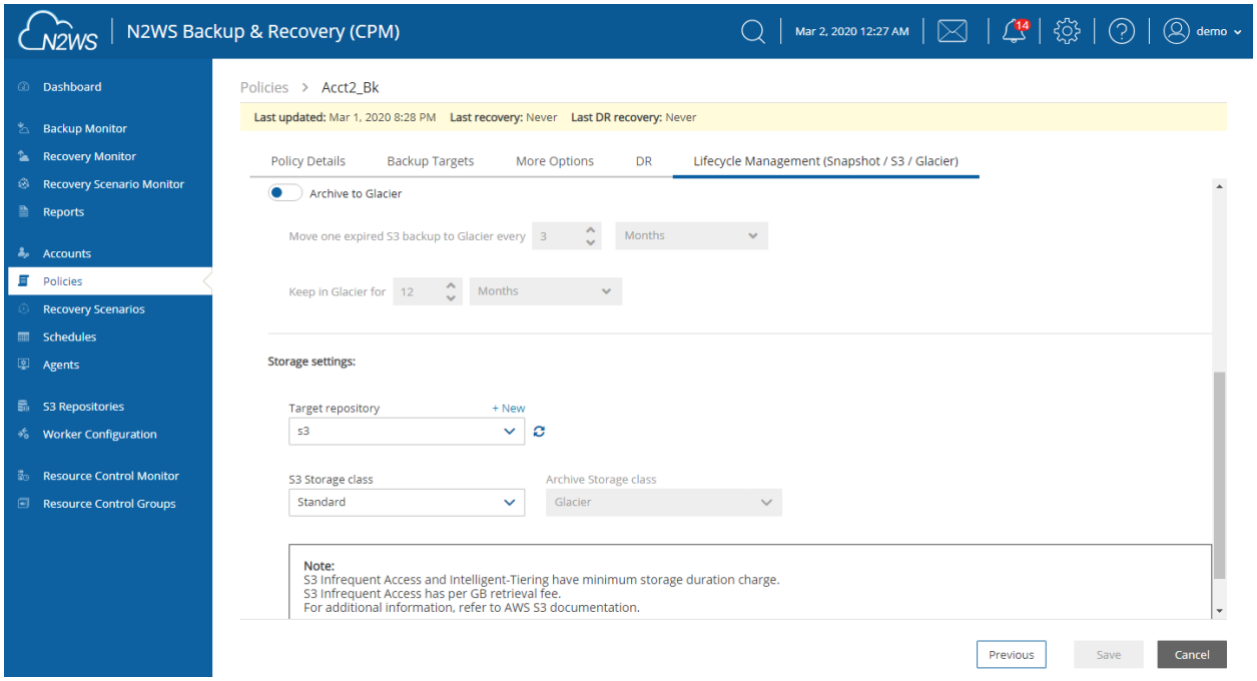

4. Complete the following parameters:
   - **Move one expired S3 backup to Glacier every** X period – Select the frequency of archiving.
   - **Keep in Glacier for** X period– Select the duration of archive in Glacier.

   Note: The duration is measured from the creation of the original EBS snapshot, not the time of archiving.

5. Select the **Archive Storage class**:
   - **Glacier** - Designed for archival data that will be rarely, if ever, accessed.
   - **Deep Archive** - Solution for storing archive data that only will be accessed in rare circumstances.

21.4.4 Recovering Snapshots from Archive

Archived snapshots cannot be recovered directly from Glacier. The data must first be copied to S3 ('retrieved') before it can be accessed.

Note: The retrieving process only runs on objects that have never been retrieved. In other words, an archived snapshot can only be retrieved once.

The process of retrieving data from archive to S3 is automatically and seamlessly managed by N2WS. However, in order to recover an archived snapshot, user should specify the following parameters:

- Retrieval tier
- Days to keep

Duration and cost of Instance recovery is determined by the retrieval tier selected. Depending on the **Retrieval option** selected, the restore completes in:

- Expedited - 1-5 minutes
- Standard - 3-5 hours
- Bulk - 5-12 hours

Note:	A typical instance backup that N2WS stores in Glacier is composed of many data objects and will probably take much longer than a few minutes.

To restore data from S3 Glacier:

1. Follow the steps for Recovering an S3 Backup. See section 21.3.2.
2. In the **Backup Monitor**, select a successful Glacier copy, and then select **Recover**.
3. In the **Restore from** drop-down list, select 'Glacier'. N2WS will copy the data from Glacier to S3 and keep it for the expiration set in **Keep data in S3 for x days**.
4. In the **Restore to Region** list, select the target region.
5. Select a **Retrieval option** (Bulk, Standard, or Expedited) and then select **Retrieve**.

21.5 Monitoring Lifecycle Activities

After a policy with Backup to S3 starts, you can:

- Follow its progress in the **Status** column the **Backup Monitor**.
- Abort the copy of snapshots to S3.
- Stop S3 and Archive operations.
- Delete S3 snapshots.

21.5.1 Viewing Status of Backups in S3 or Glacier

You can view the progress and status of S3 and archived backups in the **Backup Monitor**.

1. Select the **Backup Monitor** tab.

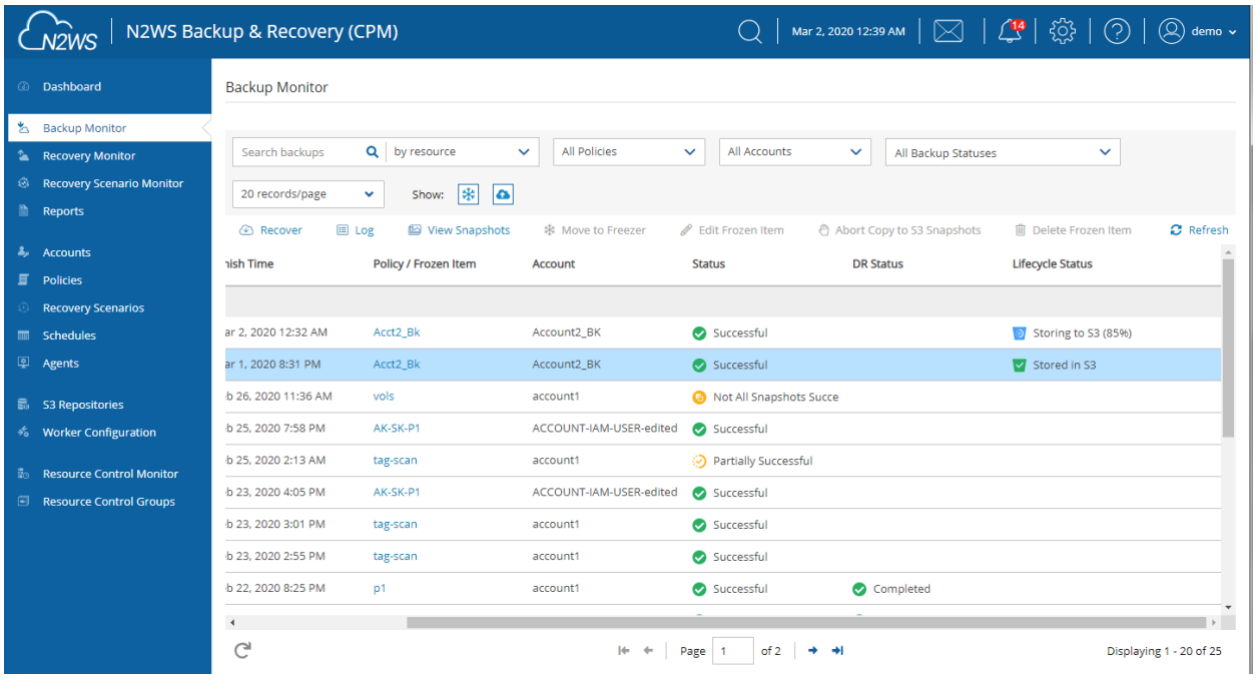

2. In the **Lifecycle Status** column, the real-time status of an S3 Copy is shown. Possible lifecycle statuses include:
   - Storing to S3 (n%)
   - Stored in S3

- Not stored in S3 – Operation failed or was aborted by user.
- Archiving
- Archived
- Marked as archived – Some or all the snapshots of the backup were not successfully moved to Archive storage, either due to the user aborting the operation or an internal failure. However, the all snapshots in the backup will be retained according to Archive retention policy, regardless of their actual storage.
- Marked for deletion – The backup was scheduled for deletion according to the retention policy and will be deleted shortly.

  Note: 'Marked for deletion' snapshots can no longer be recovered.
- Deleted from S3/Archive – Snapshots were successfully deleted from either S3 or Archive.

21.5.2 Aborting a Copy to S3 'In Progress'

The Copy to S3 portion of a Policy backup occurs after the non-S3 backups have completed.

Note: Aborting an S3 Copy does not stop the non-S3 backup portion of the policy from completing. Only the Copy to S3 portion is stopped.

To stop an S3 Copy in progress:

1. In the **Backup Monitor**, select the policy.
2. When the **Lifecycle Status** is 'Storing to S3', select Abort Copy to S3 Snapshots.

21.5.3 Stopping an S3 Cleanup in Progress

If an S3 retention Cleanup is 'In progress', in the **Policies** tab, select an S3 policy and then select **Stop S3 / Archive Operations** to stop the Cleanup**.** See the Note in section 21 for the reasons you might want to stop the S3 Cleanup.

- Stopping S3 Cleanup does *not* stop the non-S3 cleanup portion of the policy from completing. Only the S3 cleanup portion is stopped.
- Stopping S3 Cleanup of a policy containing several instances will stop the cleanup process for policy as follows:
  - CPM will perform the cleanup of the current instance according to its retention policy.
  - CPM will terminate all S3 Cleanups for the remainder of the instances in the policy.
  - CPM will set the session status to **Aborted**.
  - CPM user will get a 'S3 Cleanup of your policy aborted by user' notification by email.

To stop an S3 Cleanup in progress:

1. Determine when the S3/Archiving is taking place by going to the **Backup Monitor**
2. Select the policy and then select **Log**.
3. When the log indicates the start of the Cleanup, select **Stop S3 /Archive Operations**.

21.5.4 Deleting Copy to S3 Snapshots in a Repository

To delete only the snapshots copied to a specific S3 repository:

1. Select the **S3 Repositories** tab.

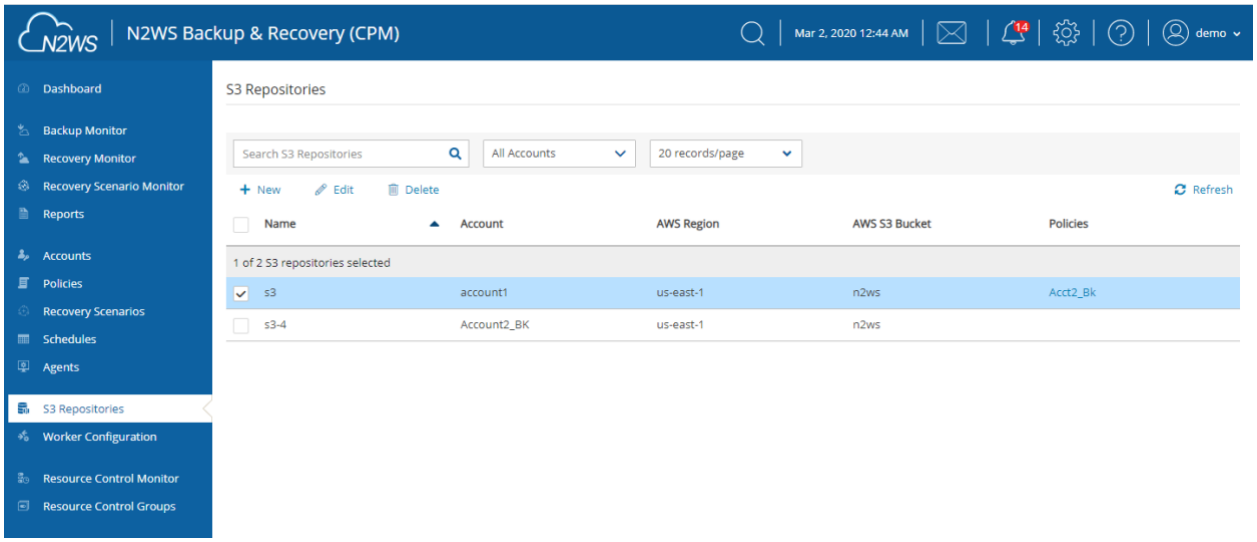

2. Select a repository, and then select **Delete**.

Note: When deleting Policies and **Snapshots** in the **Policies** tab or **Account and Data** in the **Accounts** tab, S3 copies are also deleted.

22 Configuring Workers

When N2WS copies data to or restores data from an S3 repository, or **Explores** snapshots in a region other than that of the N2WS server, it launches a temporary 'worker' instance to perform the actual work, such as writing objects into S3 or exploring snapshots.

- When performing backup operations, or **Exploring** in a non-N2WS server region, the 'worker' instance is launched in the region and account of the target instance. The backup or **Explore** 'worker' instance is configured in the **Worker Configuration** tab.
- When performing restore operations, the 'worker' instance is launched in the region and account that the backed-up instances are to be restored to. The restore 'worker' instance is selected or configured according to the following criteria:

- If a 'worker' for the target account/region combination was configured in the **Worker Configuration** screen, that 'worker' instance will be used during the restore, or during the **Explore**.
- If such a 'worker' does not exist for the target account/region combination, N2WS will attempt to launch one based on N2WS's own configuration.
- If N2WS' configuration cannot be used because the restore, or **Explore**, will be to a different account or region than N2WS', the user will be prompted during the restore to configure the 'worker'.

Note: If you plan to Copy to S3 only instances belonging to the same account and residing in the same region as that of the N2WS server, worker configuration Is not required since the worker will derive its configuration from the N2WS server instance.

Attempts to perform S3/Glacier backup and restore operations from an account/region, or to **Explore** out of the N2WS server account/region, without a valid worker configuration will fail.

You can manage workers and their configurations as well as test their communication with the CPM in the **Worker Configuration** tab:

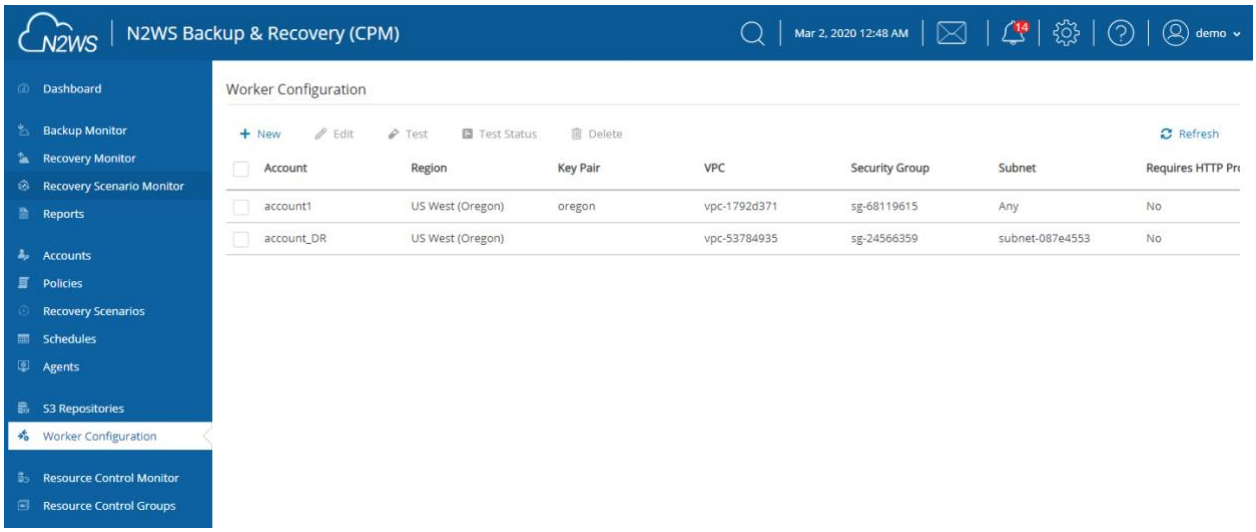

22.1 Worker Parameters

It is necessary to define a *separate* worker configuration for *each* planned account/region combination of Copy to S3 instance snapshots, or each **Explore** region that is different from the N2WS server region:

To configure S3 worker parameters:

1. Select the **Worker Configuration** tab.
2. Select **+ New**.

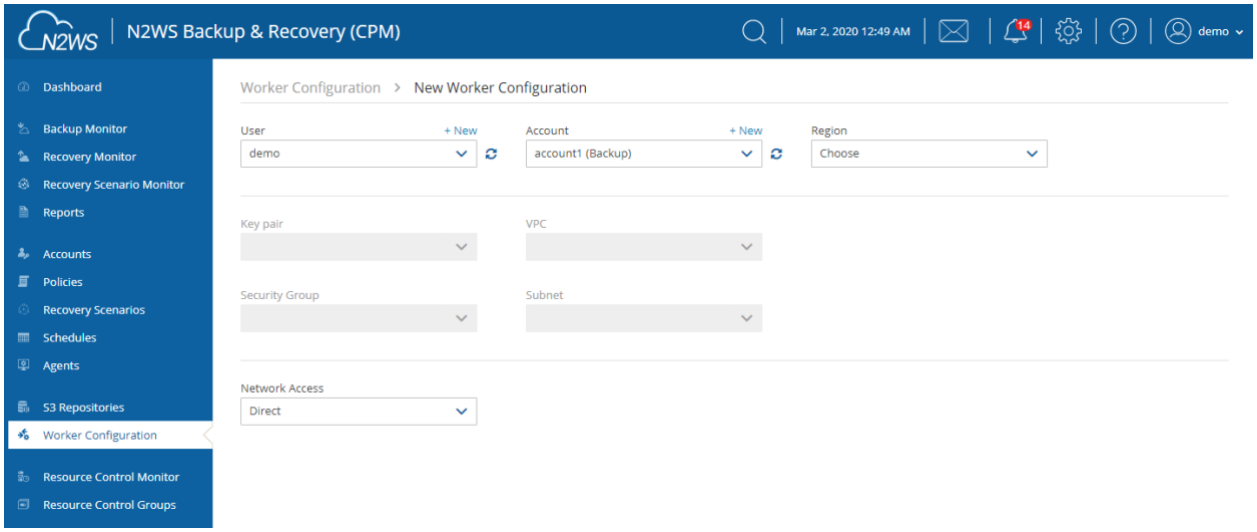

3. In the **Account** list, select the Account that the new worker is associated with.
4. In the **Region** list, select a Region. This configuration will be applied to all workers launched in this region for this account.
5. In the **Key pair** list, select a key pair. Using the default, **Don't use key pair**, disables SSH connections to this worker.
6. In the **VPC** list, select a VPC. The selected VPC must be able to access the subnet where N2WS is running as well as the S3 endpoint.
7. In the **Security Group** list, select a security group. The selected security group must allow outgoing connections to the N2WS server and to the S3 endpoint.
8. In the **Subnet** list, select a subnet, or choose **Any** to have N2WS choose a random subnet from the selected VPC.

Note:	If you choose '**Any**' in the **Subnet** drop-down list, N2WS will automatically choose a subnet that is in the same Availability Zone as the one you are restoring to. If you choose a specific subnet that is not in the same Availability Zone as the one you are restoring to, you will have to choose a different subnet from the **Subnet** drop-down list.

9. In the **Network Access** list, select a network access method.

Note:	Direct network access or indirect access via an HTTP proxy is required: • **Direct** - Select a Direct connection if no HTTP proxy is required. • **via HTTP proxy** – If an HTTP proxy is required, select and fill in the proxy values.

10. Select **Save**.
11. In the **Worker Configuration** list, test the new worker by selecting **Test**. In the **Test Status** column, you will see 'In Progress'. Select Test Status. If not 'Successful', select **Edit** and check settings.

To edit or delete a worker configuration:

1. In the **Worker Configuration** tab, select a worker.
2. Select **Delete** or **Edit**.

22.2 Testing the Configuration for a Worker

You can simulate communication between a worker and CPM before it is needed:

To test a worker configuration:

1. Select the worker in the list and then select **Test**. The 'Worker Configuration test is underway' message briefly appears at the top right and the 'In Progress' message appears in the **Test Status** column.
2. Check the results in the **Test Status** column: Successful or Failed.
3. In the case of failure, select **Test Status** to display information about the root cause.

   Information includes: Account, Region, Zone, Key Pair, VPC, Security Group, and whether a HTTP Proxy was required.

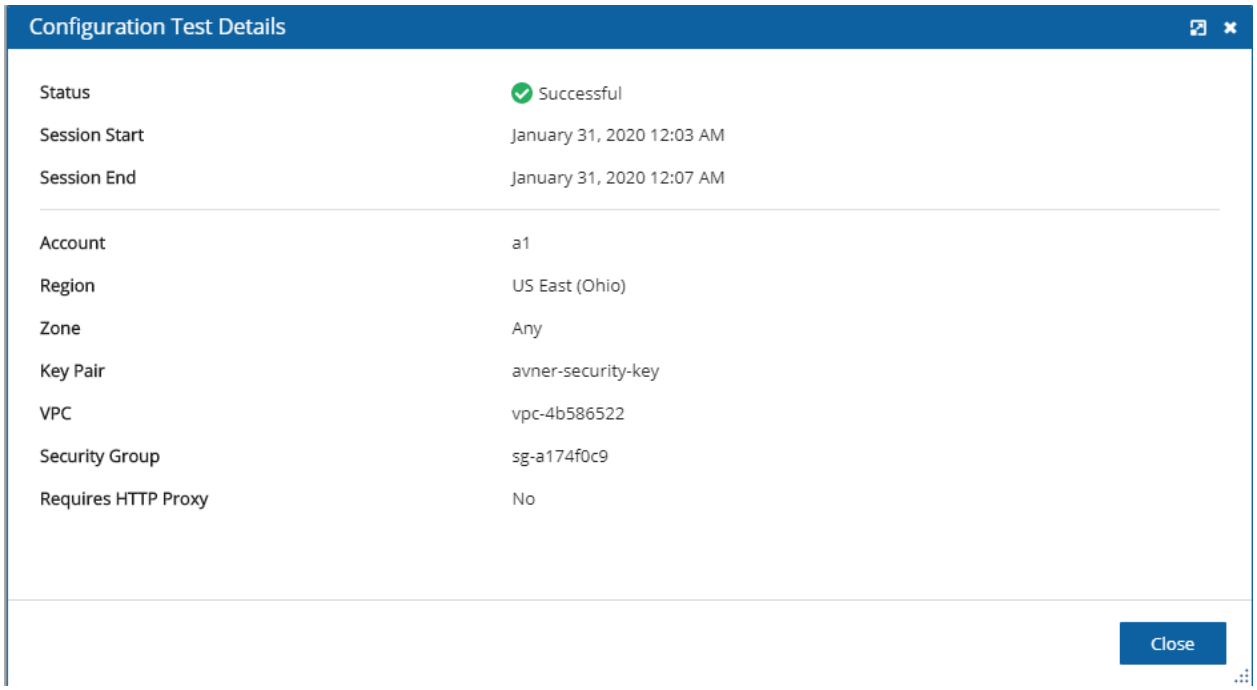

23 Capturing and Cloning in VPC Environments

Note: VPC support is not available with the Free edition of N2WS.

23.1 Overview of VPC and N2WS

VPC is an AWS service which allows the definition of virtual networks in the AWS cloud. Users can define VPCs with a network range, define subnets under them, security groups, Internet Getaways, VPN connections, and more. One of the resources of the VPC service is also called 'VPC', which is the actual virtual, isolated network.
N2WS can capture the VPC settings of user environments and clone those settings back to AWS:

- In the same region and account, for example, if the original settings were lost.
- To another region and/or account, such as in DR scenarios.
- With VPC resource properties modified in template uploaded with CloudFormation, if required.

Once enabled from **General Settings**, N2WS will automatically capture VPC settings at pre-defined intervals, such as for cleanup and tag scanning. The root/admin user can enable the feature in the **Capture VPC** tab of the **General Settings** screen and set the interval of VPC captures. VPC settings are enabled at the account level, by default, same as tag scanning.
Because VPC configuration metadata is small, VPC does not consume a lot of resources during storage of the capture. Metadata is captured incrementally. If nothing changed since the last capture, the metadata will not be captured again. This is the most common case in an ongoing system, where defined networks do not change frequently.

- Regions - N2WS will only capture VPC settings in regions that include backed-up resources. If the customer is not backing up anything in a specific region, N2WS will not try to capture the VPC settings there.
- Retention - N2WS will retain the VPC data as long as there are backups requiring it. If N2WS still holds backups from a year ago, the VPC version relevant for that time is still retained. Once there are no relevant backups, N2WS will delete the old VPC captured data.
- CloudFormation - N2WS will use the AWS CloudFormation service to clone VPCs to an AWS account. N2WS will create a CloudFormation template with the definitions for the VPC and use the template to launch a new stack and create all the VPC settings in one operation.

23.2 Features of Capturing and Cloning VPCs

The objective of Capture and Clone is to provide the ability to protect VPCs from disaster, by saving VPC configurations and allowing for recovery in any region.

- Backed up VPC entities include:
  - VPC resource configuration
  - Subnets - N2WS tries to match AZs with similar names and spread subnets in destinations in the same way as in source regions.
  - Security groups
  - DHCP Options Sets - Not supporting multi-name in domain server name.

- Route tables - Not supporting rules with entities that are specific to the source region.
- Network ACLs
- Internet Gateways, Egress Internet Gateways
- VPN Gateways

Note: The **Capture Log** in the **Capture VPC** tab of **General Settings** reports entities not captured or only partially captured.

- VPC capturing:
  - Accounts are enabled for VPC capturing by default, but this setting can be disabled as needed.
  - Captures in all regions of interest.
  - N2WS will capture and save all changes made on AWS for a user's VPCs.
  - Not supported: NAT gateways, VPC peering connections, customer gateways, VPN connections, Network interfaces, Elastic IP addresses, VPC Endpoints, VPC Endpoints services, Transit Gateways
- VPC cloning:
  - Every Account that has a VPC captured in a region can clone a version of the VPC to any destination, region, and account.
  - The subnets of the cloned VPC will be located in the destination's Availability Zone with respect to their availability in the region.
  - Users can download a template of VPC resources to manually configure and load it with AWS CloudFormation.

23.3 Configuring VPC Capturing

The root user can:

- Enable or disable automatic VPC captures for Accounts that are VPC-enabled.
- Schedule automatic capture interval.
- Initiate an ad hoc capture by selecting **Capture Now** for all VPC-enabled Accounts, even if VPC is disabled in **General Settings**.
- View the last captured VPCs in the different regions and accounts in **Capture Log.**

1. Select **Server Settings** > **General Settings**.
2. In the **Capture VPC** tab, select **Capture VPC Environments** to enable the feature.

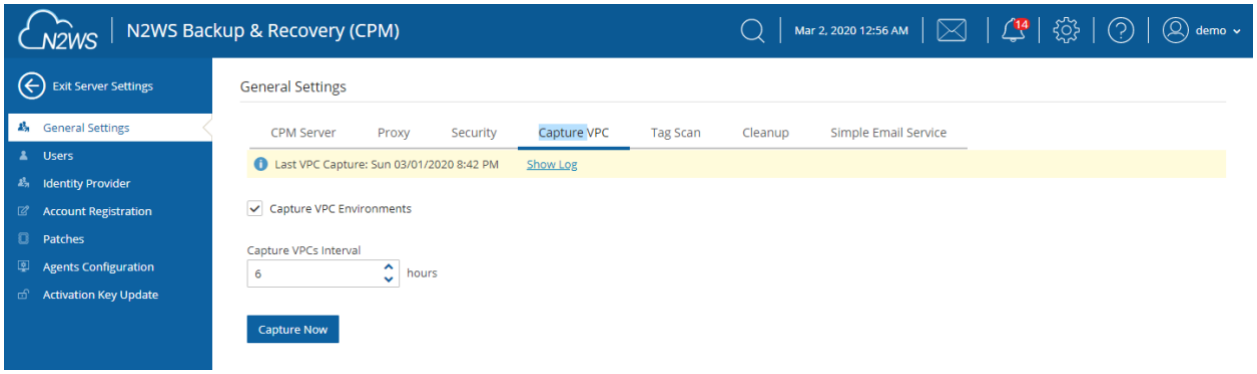

3. To change the capture frequency from the default, select a new interval from the **Capture VPCs Interval** list. Valid choices are from every hour to every 24 hours.

4. Select **Save** to update N2WS.
5. To initiate an immediate capture for all VPC-enabled Accounts regardless of server setting, select **Capture Now.**

23.4 Updating Accounts for VPC

By default, Accounts are enabled to Capture VPCs. VPCs are automatically captured for all enabled Accounts according to the interval configured in the **General Settings**. To not capture VPCs for an Account, disable the feature in the Account.

To disable, or enable, an individual account for capturing VPCs:

1. Select the **Accounts** tab and then select an Account.
2. Select **Capture VPCs** to enable:

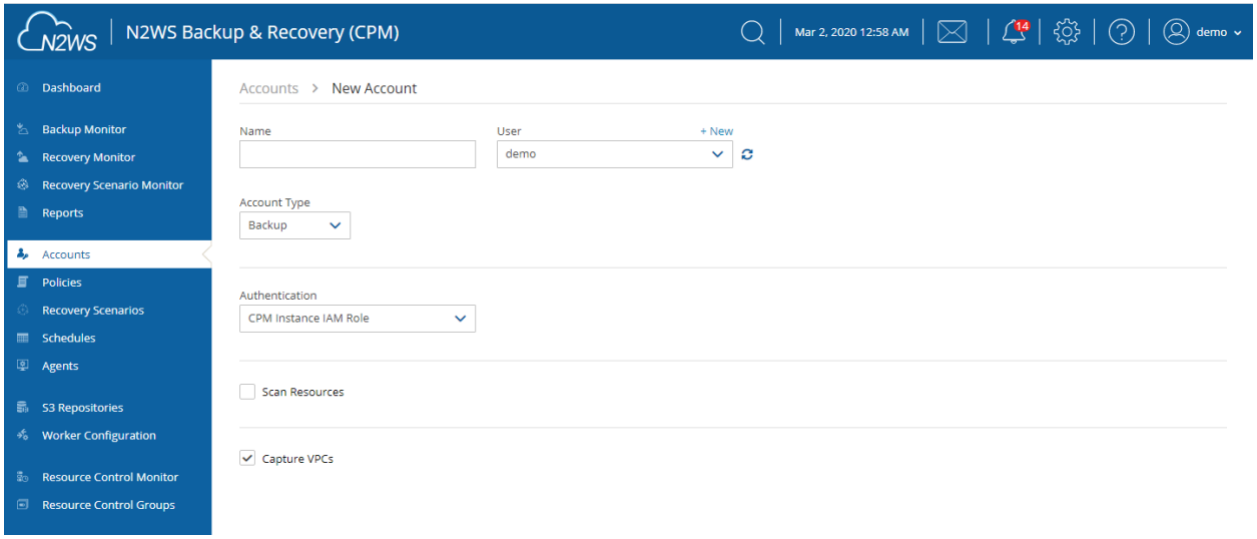

3. Select **Save**.

23.5 Cloning VPCs

The following entities are not supported:

- Cloning CIDR block IPV6 on a subnet.
- Inbound and Outbound Endpoint rules of Security Groups.
- Inbound and Outbound rules of Security Groups that refer to a security group on a different VPC.
- Route Table rules with NAT Instance as target.
- Route Table rules with NAT Gateway as target.
- Route Table rules with Network Interface as target.
- Route Table rules with VPC peering connection as target.
- Route Table rules with status 'Black Hole'.

A VPC-enabled account must have at least one policy with a backup target in order to clone VPCs.

Cloning VPCs includes the following features:

- Both cross-region and cross-account cloning are supported.

- The target clone can have a new name. The name will automatically include '(cloned)' at the end.
- During instance recovery and DR, clones may be optionally created in order to replicate a particular VPC environment before the actual instance recovery proceeds. The new instance will have the environment of the cloned VPC and will subsequently appear at the top of the target region and account list. A typical scenario might be to capture the VPC, clone the VPC for the first instance, and then apply the cloned VPC to additional instances in the region/account.
- Instances recovered into a cloned VPC destination environment will also have new default entities, such as the VPC's subnet definition and 1 or more security groups attached to the instance, regardless of the original default entities. Security groups can be changed during recovery.

When cloning VPCs to an AWS account, N2WS generates a JSON template for use with CloudFormation.

- If the size of the CloudFormation template generated will be over 50 kB, N2WS requires the use of an existing S3 Bucket in the target destination for storing the template. **There should be an S3 bucket for each combination of accounts and regions in the destination clone**. The template file in a S3 bucket will not be removed after cloning.
- In addition to having a bucket in the target region in the presented settings, you must choose that bucket when defining where to **Upload the CF template to S3**.

To clone captured VPCs:

1. Select the **Accounts** tab and then select an account.
2. Select **Clone VPC**.

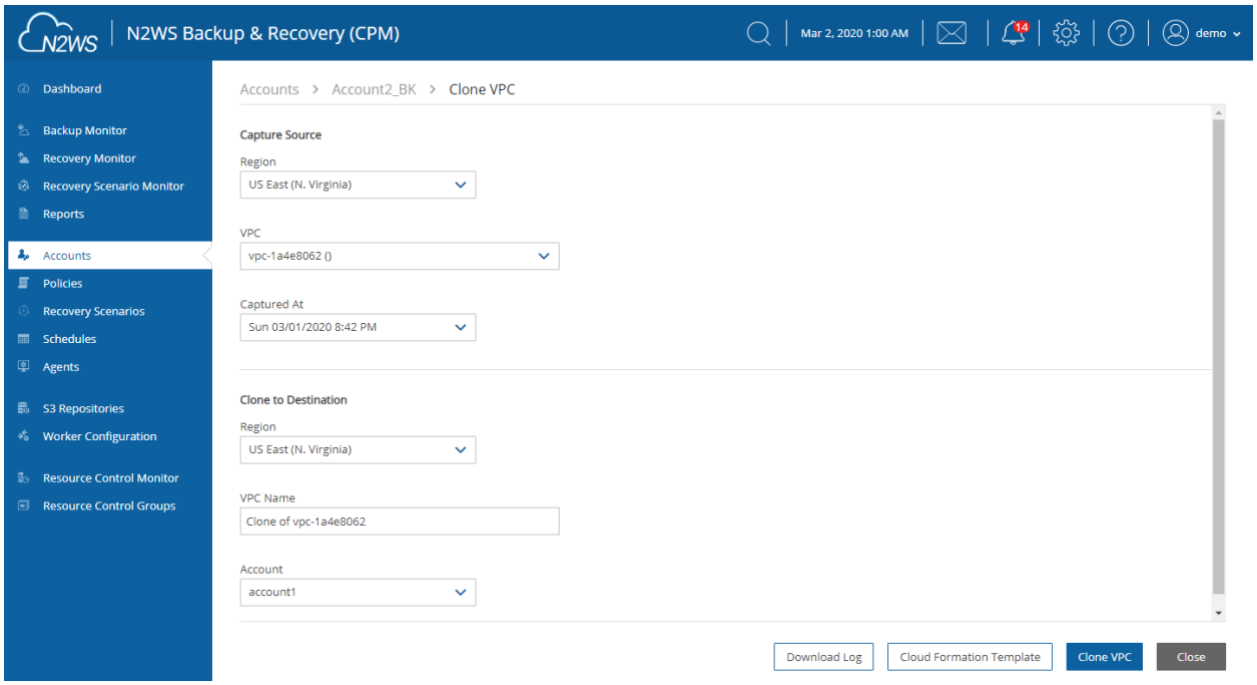

3. In the **Capture Source Region** drop-down list, select the source region of the capture to clone.
4. In the **VPC** drop-down list, select the VPC to clone.
5. In the **Captured at** drop-down list, select the date and time of the capture to clone.

6. In the **Clone to Destination Region** drop-down list, select the region to create the clone.
7. In the **VPC Name** box, a suggested name for the VPC is shown. Enter a new VPC name, if needed.
8. In the **Account** drop-down list, select the account in which to create the clone.
9. If the CF template is over 50 kB, the **CloudFormation Template** button appears. Click to download a json file with cloning information.

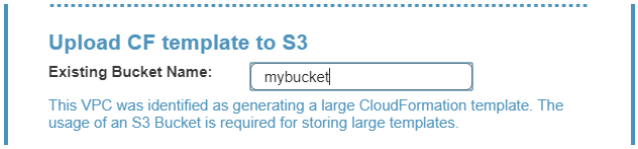

Enter an **Existing Bucket Name**.

Note: The existing bucket must be located in the selected target region.

10. Select **Clone VPC**. At the end of the cloning, a status message will appear in a box:
    - VPC was Cloned. There may be an informational message that you may need to make manual changes. Check the log for further information.
11. To view the results of the clone VPC action, select **Download Log**.

When cloning VPCs with resources not supported by N2WS, you can download the CloudFormation template for the VPC, add or modify resource information, and upload the modified template to CloudFormation manually.

To create a clone manually with CloudFormation:

1. In the **Account Clone VPC** screen, complete the fields as described above.
2. Select **CloudFormation Template** to download the CloudFormation JSON template.
3. Modify the template, as required. See section 23.5.1.
4. Manually upload the modified template with CloudFormation.

23.5.1 Example of CloudFormation Template

```
{'AWSTemplateFormatVersion': '2010-09-09',
 'Description': 'Template created by N2WS',
 'Resources': {'dopt4a7bcf33': {'DeletionPolicy': 'Retain',
                                'Properties': {'DomainName': 'ec2.internal',
                                               'DomainNameServers': ['AmazonProvidedDNS']},
                                'Type': 'AWS::EC2::DHCPOptions'},
               'dopt4a7bcf33vpc9d4bcbe6': {'DeletionPolicy': 'Retain',
                                           'Properties': {'DhcpOptionsId': {'Ref':
'dopt4a7bcf33'},
                                                          'VpcId': {'Ref': 'vpc9d4bcbe6'}},
                                           'Type': 'AWS::EC2::VPCDHCPOptionsAssociation'},
               'sgcd8af6bb': {'DeletionPolicy': 'Retain',
                              'Properties': {'GroupDescription': 'default VPC security
group',
                                             'GroupName': 'default-0',
                                             'SecurityGroupEgress': [{'CidrIp': '0.0.0.0/0',
                                                                      'IpProtocol': '-1'}],
                                             'SecurityGroupIngress': [],
                                             'Tags': [{'Key': 'cpm:original:GroupId',
                                                       'Value': 'sg-cd8af6bb'}],
                                             'VpcId': {'Ref': 'vpc9d4bcbe6'}},
                              'Type': 'AWS::EC2::SecurityGroup'},
               'vpc9d4bcbe6': {'DeletionPolicy': 'Retain',
                               'Properties': {'CidrBlock': '10.0.0.0/24',
                                              'EnableDnsHostnames': false,
                                              'EnableDnsSupport': true,
                                              'InstanceTenancy': 'default',
                                              'Tags': [{'Key': 'Name',
```

```
                                                 'Value': 'Public-VPC-for-CF'},
                                                {'Key': 'cpm:capturetime',
                                                 'Value': 'Aug 22, 2018 16:15'},
                                                {'Key': 'cpm:clonetime',
                                                 'Value': 'Aug 25, 2018 21:20'},
                                                {'Key': 'cpm:original:VpcId',
                                                 'Value': 'vpc-9d4bcbe6'},
                                                {'Key': 'cpm:original:region',
                                                 'Value': 'us-east-1'}]},
                                  'Type': 'AWS::EC2::VPC'}}}
```

24 Orchestrating Recovery Scenarios

24.1 Overview

The Recovery Scenarios feature allows N2W Software users to design an object that will automatically coordinate a sequence of recoveries for several or all backup targets of a single policy during one recovery session.

- A Recovery Scenario object is created with the saved configurations of successful backups for the particular policy.
- The user will save the recovery configuration for each selected backup target and add it to the Recovery Scenario object.
- At runtime, the user selects a successful backup record to use in the recovery.

Note: Backups in the Freezer are not recoverable as part of a Recovery Scenario.

24.2 Conditions

- During the Recovery Process:
  - All Recovery Scenario targets share the same destination account and destination region, which are set as part of the Recovery Scenario parameters.
  - Recovery Scenarios can have pre- and post- scripts which will run, respectively, prior to recovery execution and subsequent to recovery completion.
    - In case of a pre-script failure, the Recovery Scenario will not execute.
    - In case of a Recovery Scenario failure or pre-script failure, the post-script will not run.
- Every Recovery Scenario target has a sequential **Recovery Order** value within the Recovery Scenario which determines the order in which each target is recovered.
  - Execution of a target recovery within the recovery scenario is sequenced using the targets **Recovery Order** value. The target with the lowest **Recovery Order** value runs first.
  - All recovery targets sharing same **Recovery Order** value will run in an arbitrary sequence.
  - If the recovery of a target fails, the targets next in sequential order will not run, unless Recovery Scenario's **Continue recovering ignoring failures** parameter is enabled.
  - Testing: You can verify the Recovery Scenario input parameters, such as key pair, security groups, and VPC, by selecting the **Dry Run** link. You will be prompted to select a successful backup for the **Dry Run** just as with an actual **Run Scenario**.

24.3 Creating a Recovery Scenario

1. Select the **Recovery Scenarios** tab.

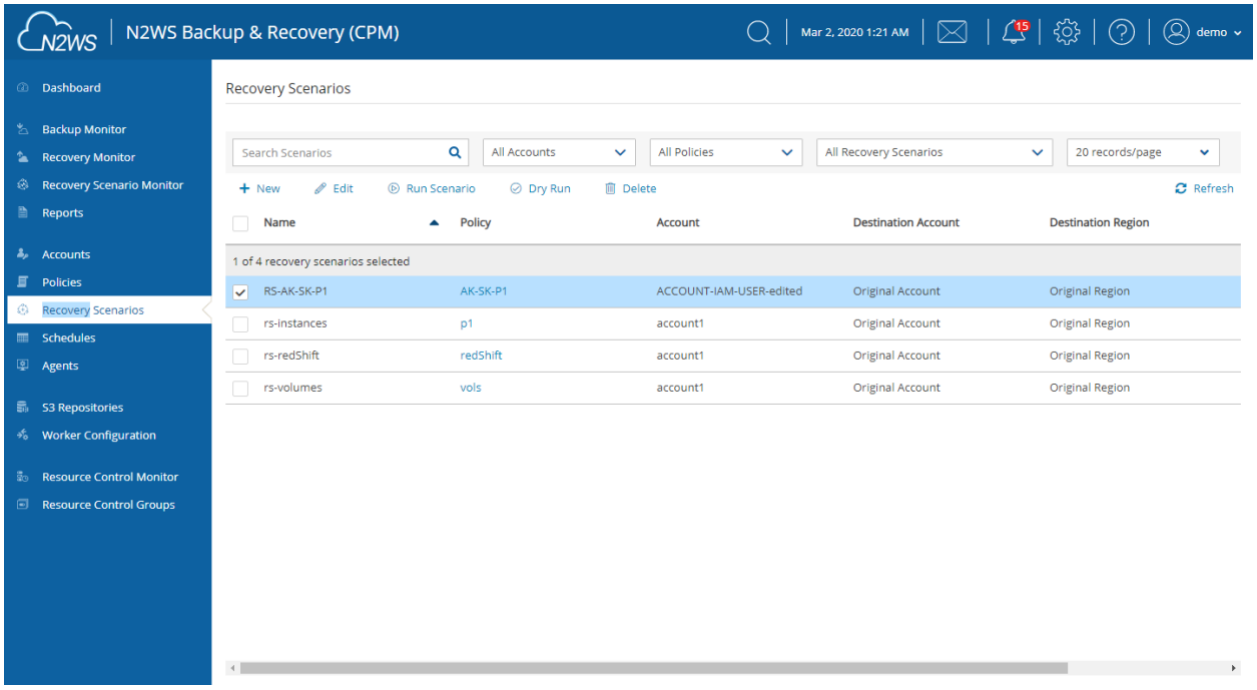

2. Select **+ New**.

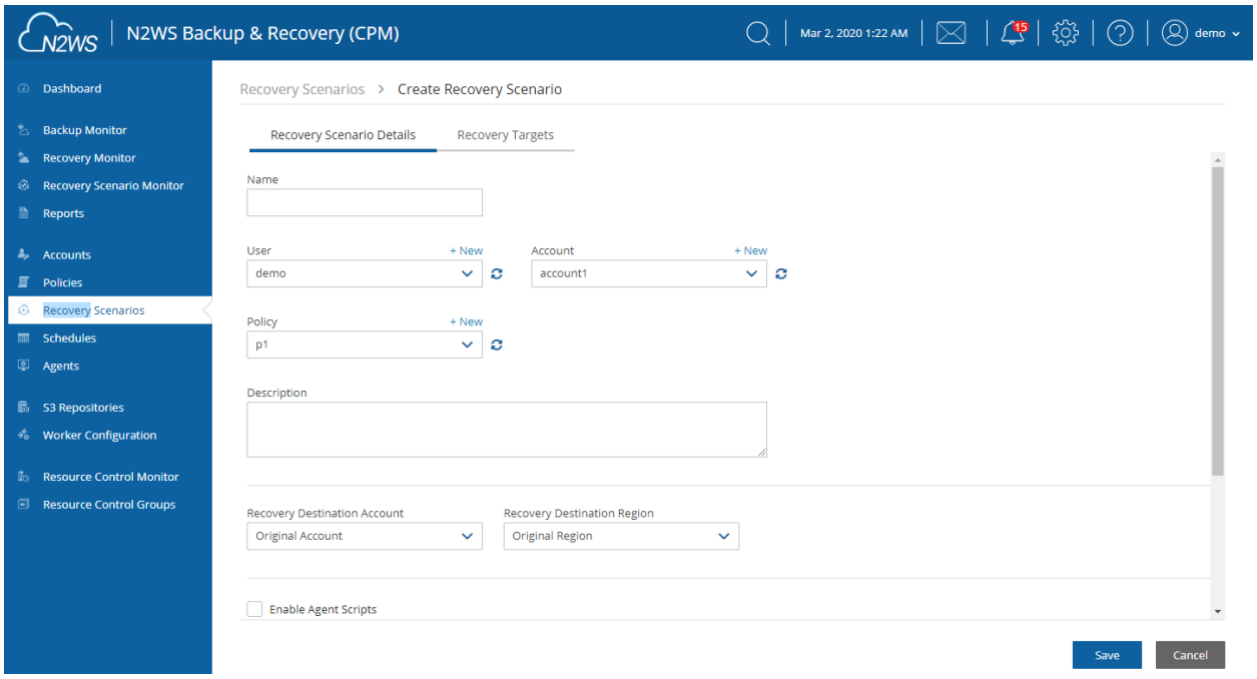

3. In the **Recovery Scenario Details** tab, complete the fields as follows:
   - **Name** – Enter a unique name.
   - **User, Account, Policy** - Select from the lists or select **+ New**. After an addition, select **Refresh**. Select the policy for which the Recovery Scenario is defined.
   - **Recovery Destination Account** and **Recovery Destination Region** – Select from the lists.
   - **Enable Agent Scripts** – Select if the Recovery Scenario will be run by a custom script. The default is *not* to run user scripts. See section 24.7.
     - Select **Agent Script Timeout** in seconds from the list. When the timeout is reached, CPM will skip the script and continue with the recovery scenario.

- **Collect Script Output** – Whether to collect script output in a log. Default is to collect.
- **Continue recovering ignoring failures** – Whether to continue the sequence of recoveries in the scenario if there is a failure. The default is to not continue the script on the failure of a recovery.

4. Select **Save**.
5. Select the **Recovery Targets** tab.

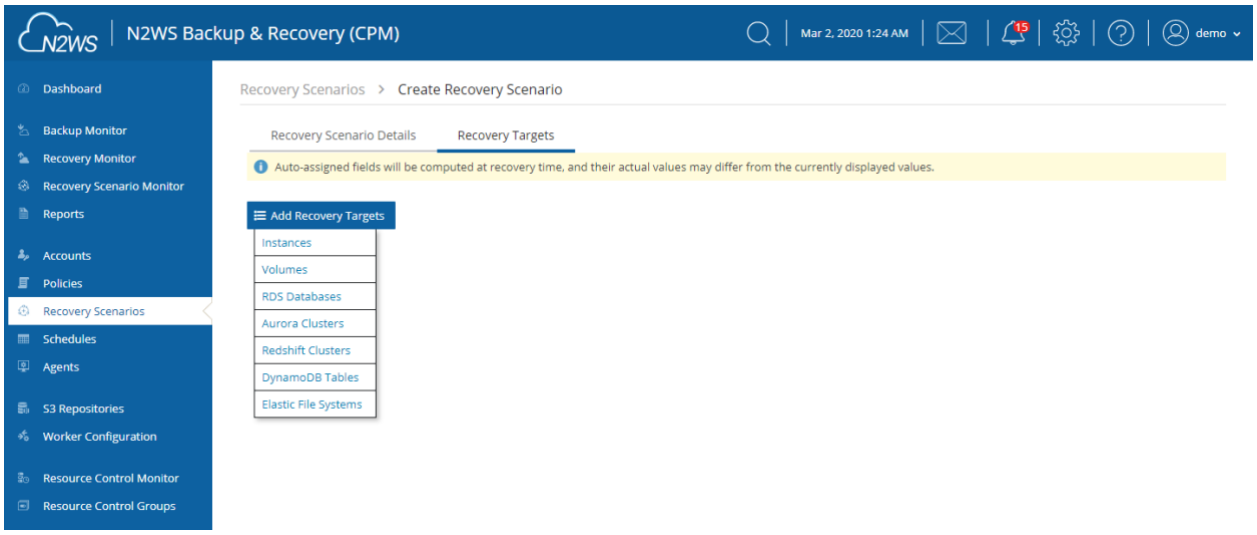

6. In the **Add Recovery Targets** menu, select a resource type from the target policy to add to the scenario.

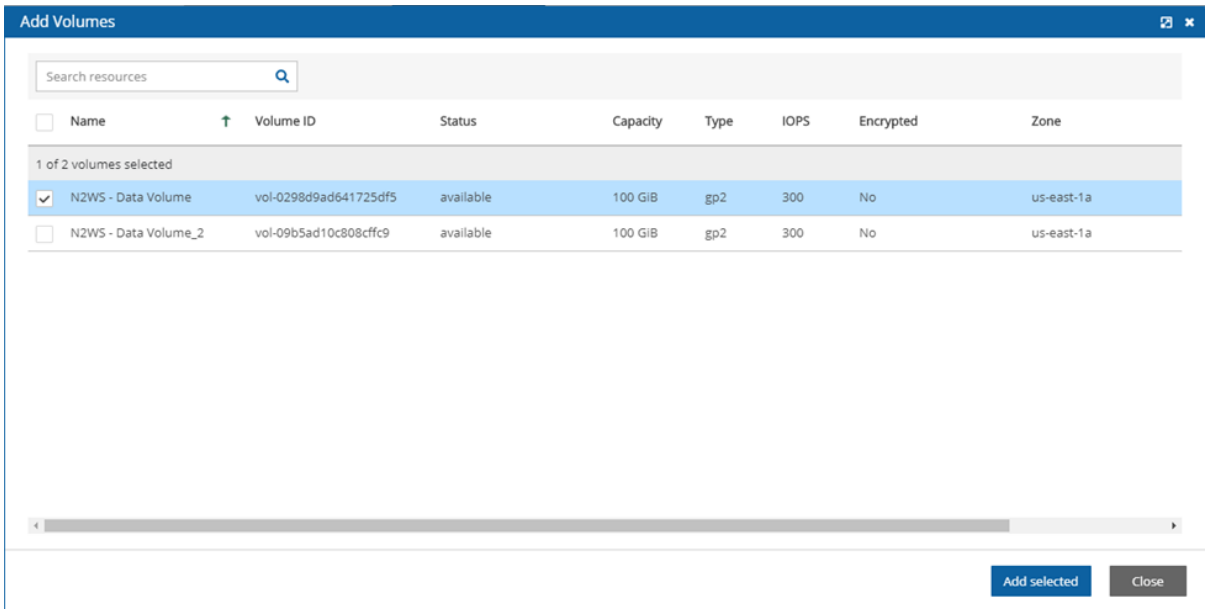

7. Select one or more **Recovery Targets** for the resource type, and then select **Add Selected**.
8. To change the **Recovery Source** for a target, select an option from the list.

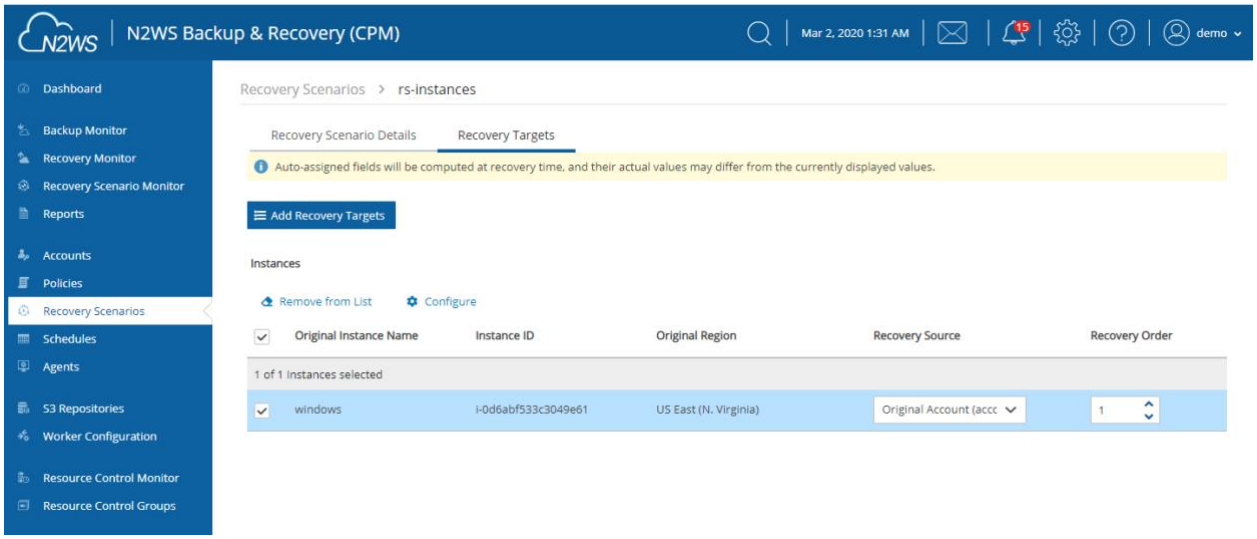

9. To change the **Recovery Order** for a target, select a value from the list.
   - Every Recovery Scenario target has a number identifying the sequential order of execution within the Recovery Scenario.
   - The execution of the recovery source within the Recovery Scenario is sequenced using the target's **Recovery Order** value. The recovery of the target with the lowest value runs first.
10. For Instance and EFS Recovery Targets, it is important to configure the recovery details for each target. Select a recovery target and then select **Configure**. See section 24.3.1.
11. When all details are complete, select **Save** in the Create Recovery Scenario screen.

24.3.1 Configuring an Instance Recovery Target

The Configuration screen opens with additional tabs:

- **Basic Options**
- A tab for the resource type, such as **Volumes**
- **Advanced Options**

Note: The configuration **Auto assigned** values may be different than the values that are shown as grayed-out. In order to be sure about a value, you need to assign it.

For each data item in the configuration tabs, assign the appropriate value. In each tab, you can customize a setting by turning off its **Auto assigned** toggle. Depending on the data item, you can:

- Select a different value from the **Custom** drop-down list.
- Enable or disable a feature.
- Enter a new value.

When finished with each tab, select **Close**.

In the **Basic Options** tab, you can configure basic recovery actions, such as whether to launch from a snapshot or image, which key pair to use, and network placement.

Note: Since not all instance types are available in all AWS regions, recovery of an instance type to a region where the type is unsupported may fail. Where the instance type is

not supported yet in an AWS region, we recommend configuring a supported **Instance Type** in the **Basic Options** parameters. See section 10.2.1.

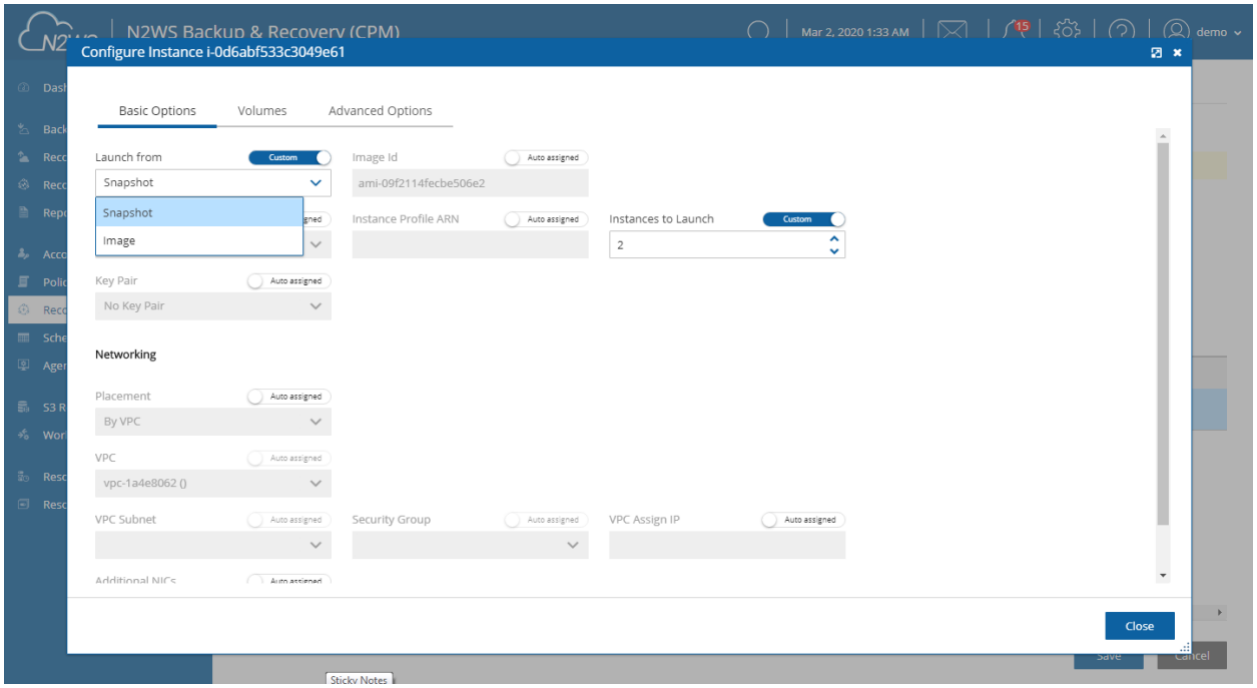

In the **Volumes** tab, you can configure device information, such as capacity and type and whether to preserve tags and delete on termination. To expand the configuration section for a volume, select the right arrow .

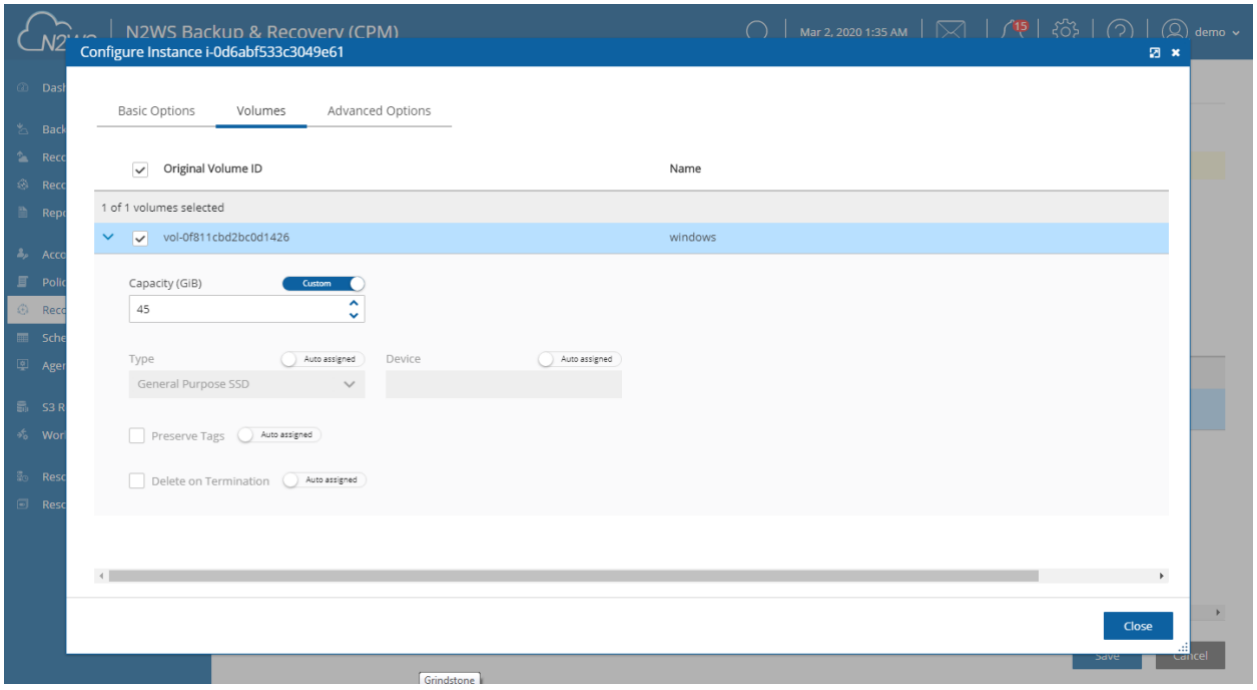

In the **Advanced Options** tab for an instance, you can customize recovery target features, such as architecture, shutdown behavior, whether to enable ENA and user data.

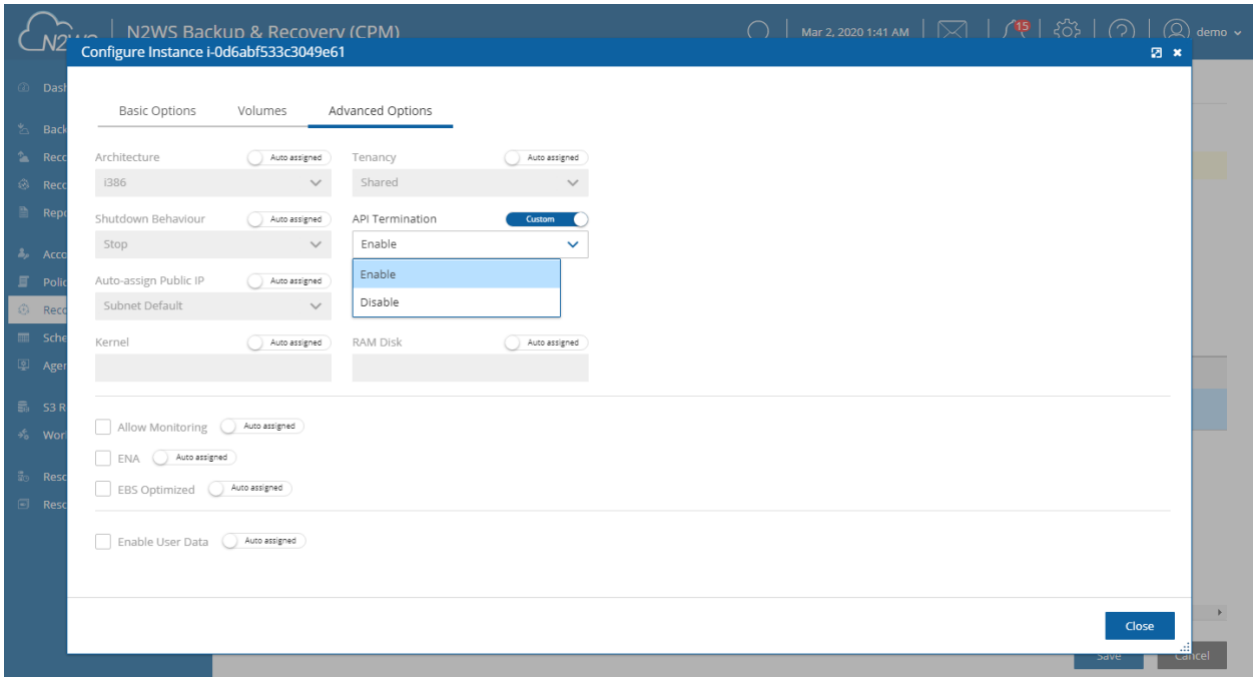

For complete details about performing an instance recovery, see section 10.2.

24.4 Testing a Recovery Scenario

The **Dry Run** option allows you to determine whether the input parameters, such as key pair, security groups and VPC, are correct for the recovery.

To test a Recovery Scenario:

1. In the **Recovery Scenarios** tab, select a Recovery Scenario and then select **Dry Run.**
2. In the list of successful backups, select one backup to recover from, and then select **Dry Run**.
3. Open the **Recovery Scenario Monitor**.
4. In the **Status** column for the Recovery Scenario, you will see a success message for the test: Recovery Successful
5. Selecting **Recoveries** brings you to the regular **Recovery Monitor**.

24.5 Managing Recovery Scenarios and Targets

To manage a Recovery Scenario object:

1. In the **Recovery Scenarios** tab, select a scenario.
2. Select Edit, **Delete**, **Run Scenario**, or **Dry Run,** as needed.

To manage targets in the scenario:

1. In the Recovery **Scenarios** tab, select a scenario, and then select **Edit**.
2. To delete a target, select the **Recovery Targets** tab, select a target, and then select **Remove from List**.
3. Depending on the resource type, the action **Configure** is available. Configure opens tabs for **Basic Options**, resource type details, and **Advanced Options**.

24.6 Running and Monitoring a Recovery Scenario

1. In the **Recovery Scenarios** tab, select a Recovery Scenario and then select **Run Scenario**.
   A list of backups, successful and unsuccessful, opens.

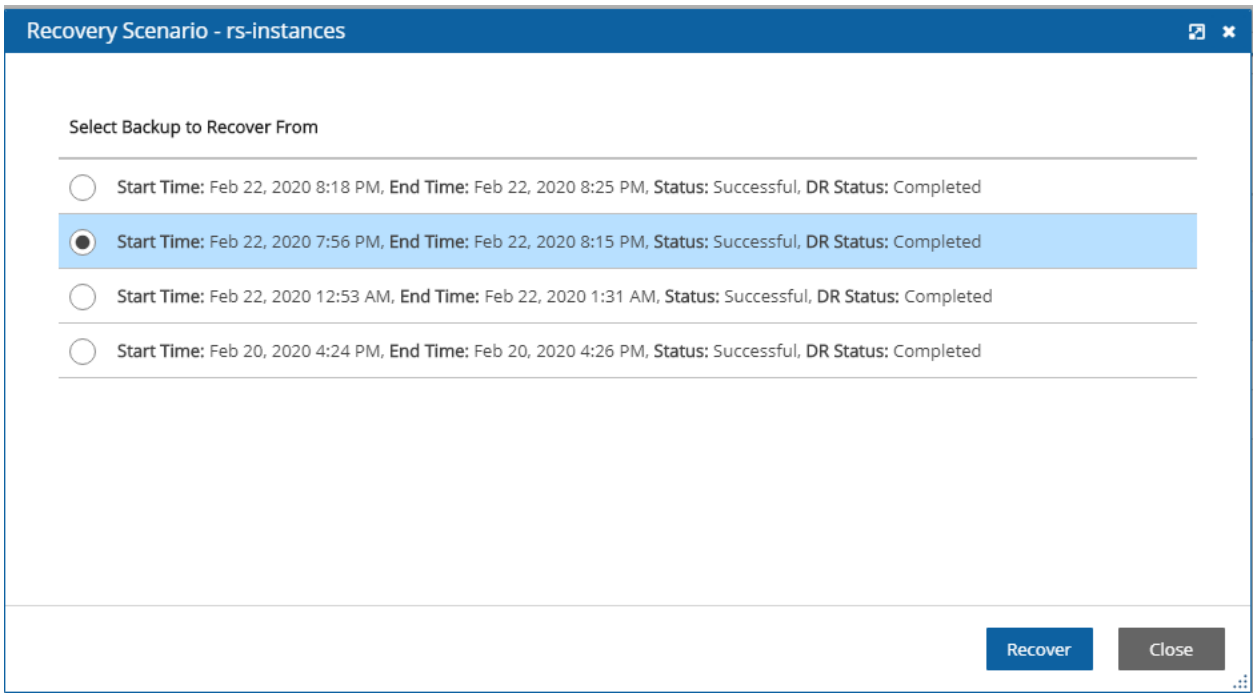

2. Select one successful backup to recover from and then select **Recover**.
   The started message opens in the top right corner:

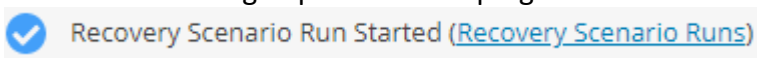

3. To open the **Recovery Scenario Monitor**, select the link, or select the **Recovery Scenario Monitor** tab.

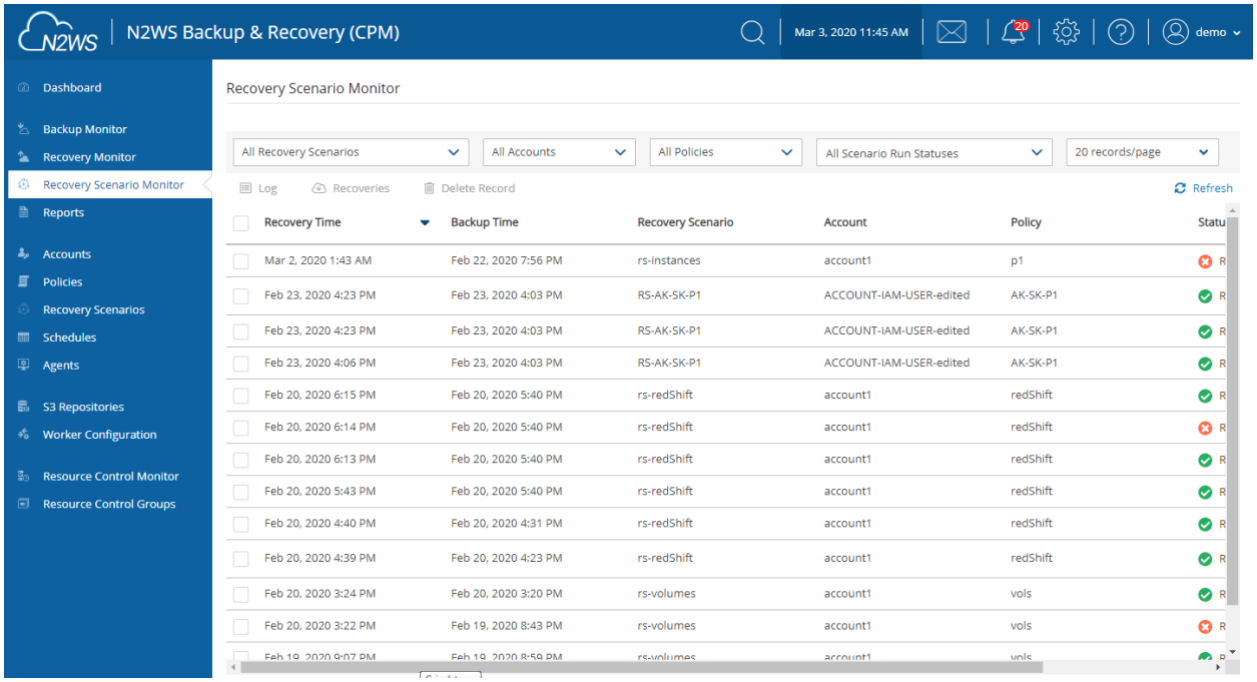

A **Status** of 'Recovery succeeded' with a test tube icon next to it indicates that the recovery was a Dry Run.

4. To view details of the recovery in the Run Log, select a **Recovery Scenario** and then select **Log**.

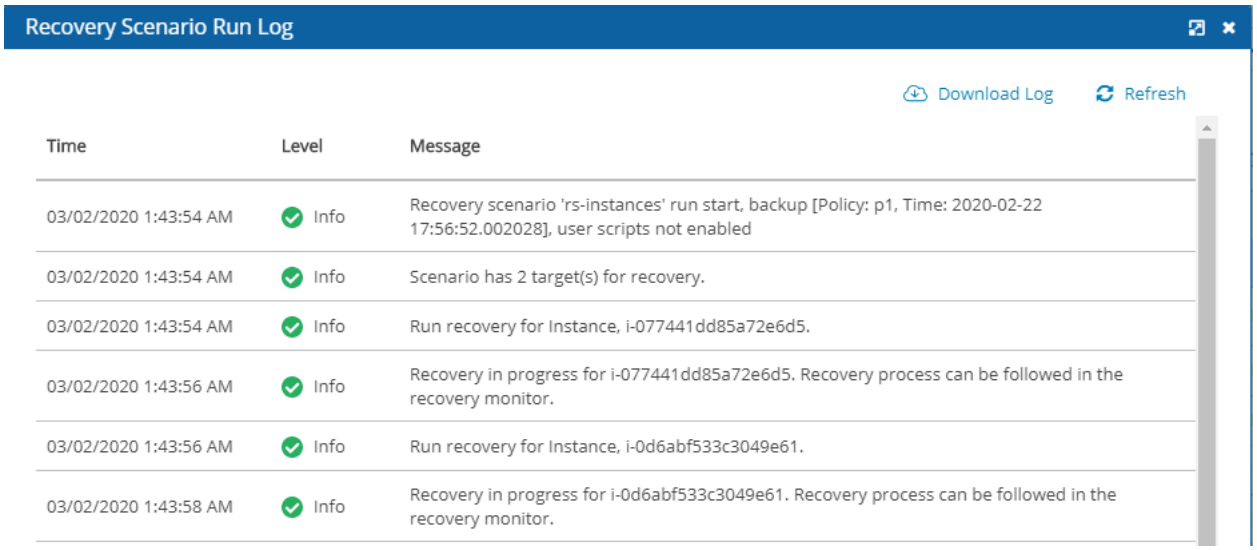

5. To delete a run record, select a scenario, and then select **Delete Record**.

Note: Deleting a run record will trigger deletion of all its target recovery records.

6. To view a live recovery, select a scenario, and then select **Recoveries**. The **Recovery Monitor** opens.

24.7 Recovery Scenario User Scripts

When **Enable Agent Scripts** is set in the **Recovery Scenario Details** tab, N2WS will run two scripts, one before and one after the recovery run:

- `before_<recovery_scenario_name>`
- `after_<recovery_scenario_name>`

A file extension is optional and, if added, may be for any interpreter.

Note: This is somewhat similar to the Linux Backup Scripts feature described in the Before Script and After Script topics, sections 6.3.1 and 6.3.2.

These scripts must be located on the N2WS server in the following folder:

- For root user: `/cpmdata/scripts/scenario`
- For non-root user: `/cpmdata/scripts/scenario/user name`

24.7.1 Before Script

The **before** script passes the following parameters, in the following order:

#	Parameter	Notes
1	Scenario Id	
2	Account Id	May be `null`, if the value is NULL.
3	Policy account Id	

#	Parameter	Notes
4	Destination region	May be `null`, if the value is NULL.

24.7.2 After Script

The **after** script passes the same parameters as the **before** with the addition of parameters for the scenario's recovery targets:

#	Param.	Notes		
1 – 4	…	Same as **before_** parameters.		
5	Target lists	Each target is represented by the following colon separated format: **RecoveryType:OriginalAwsResourceId:OriginalRegion:RecoveredAwsResourceId**		
		RecoveryType	A single character identifying resource type:	
			I	Instance
			V	Volume
			R	RDS Database
			A	RDS (Aurora) Cluster
			C	Redshift Cluster
			D	DynamoDB Table
			E	EFS
		OriginalAwsResourceId	The AWS ID of the original resource.	
		OriginalRegion	The AWS region of the original resource.	
		RecoveredAwsResourceId	The AWS ID of the recovered resource. If not recovered, then '`null`'.	

Following is an example of an **after_** script for a Recovery Scenario that was defined with 2 targets: an EC2-instance and an EC2-volume. The **after_** script passes 6 parameters, 2 of which are for the targets. The instance recovery target was *not* recovered:

```
1
null
1
null
I:i-0a87ab83ca3fa62c2:us-east-1:null
V:vol-0197aba1f7090c513:us-east-1:vol-03336f4ed151b5d29
```

25 Monitoring Costs and Savings

N2W Software customers have a single point of control and management over the procedure of backing up their cloud-based services and data stores. Monitoring the costs will help customers define backup plans that fit their budget and thereby avoid unexpected costs. N2WS provides the following services for monitoring costs:

- **Cost Explorer** – Cost of storage that was created by N2WS to hold short-term backup snapshots of the customer's cloud-based assets.
  - Allows customers to monitor the costs by the backup processes generated by N2WS.
  - Allows customers to issue monthly bills per policy backups.
  - Calculations are made for the last month by default and can be set to prior periods.
  - Breakdown of costs is found in the **Costs ($)** column of the **Policies** tab.
- **Cost Savings** – Amount of money that users can save by enabling Resource Control management.
  - Calculations are made for the next month.
  - Breakdown of savings is found in the **Cost Savings** column of the **Resource Control Groups** tab.

Notes:

- Cost Explorer support is currently limited to the AWS resource EBS.
- N2WS uses the AWS REST API for retrieving costs for specified policy. The Cost Explorer API allows us to programmatically query your data usage and compute the cost and usage data. It can take up to 48 hours for the cost increase to take effect.
- The costs will include both short-term and long-term backups (cross-region DR), but not snapshots which were copied onto cheaper media such as S3 and Glacier.

In the Dashboard screen, you can find both Cost Explorer and Cost Savings information in their respective tiles:

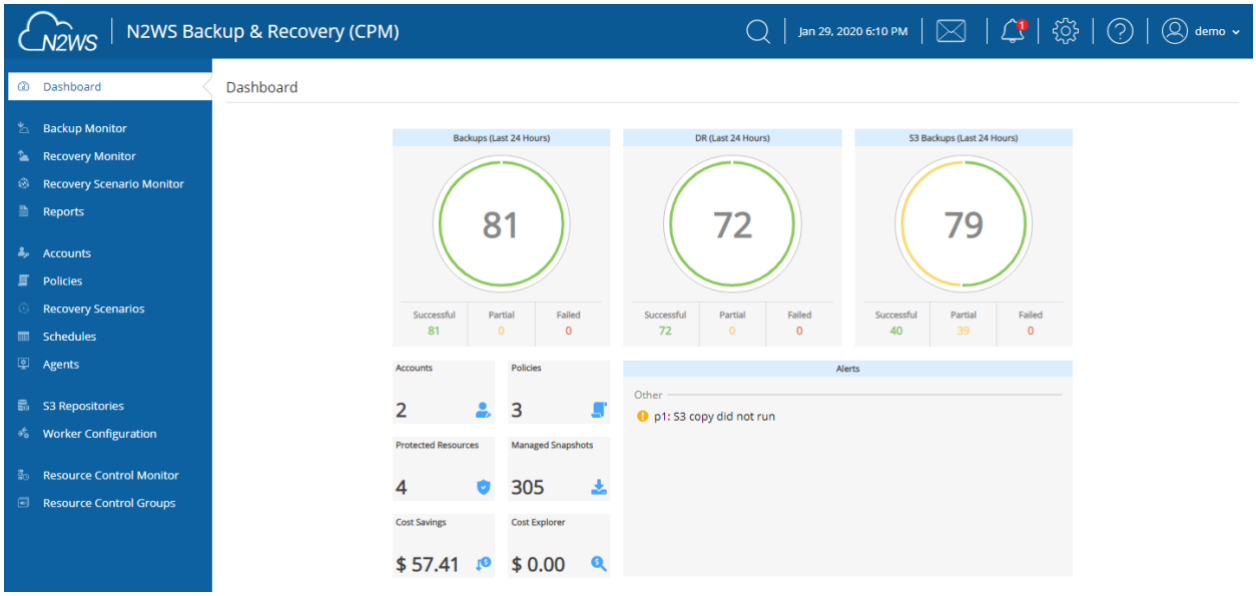

25.1 Requirements

Note:	Cost Explorer is *not* available in AWS GovCloud (US).

Following are the steps necessary for using Cost Explorer:

- In AWS, activate cost allocation tags. See section 18.3.1.
- In N2WS, enable **Allow Cost Explorer** for user. See section 18.3.

25.2 Monitoring Costs

In the **Policies** screen, you can monitor the backup costs of each policy for the last month or a different time period. The breakdown of costs per policy in dollars for the desired period appears in the **Cost ($)** column.

Notes:	Cost Explorer inclusions and exclusions: • The cost figure includes all policies that were ever backed up by N2WS and does not filter out deleted policies. The costs for policies deleted during a cost period are still included in the cost figure for that period. • N2WS Cost Explorer does *not* include the cost of cross-account DR snapshots.

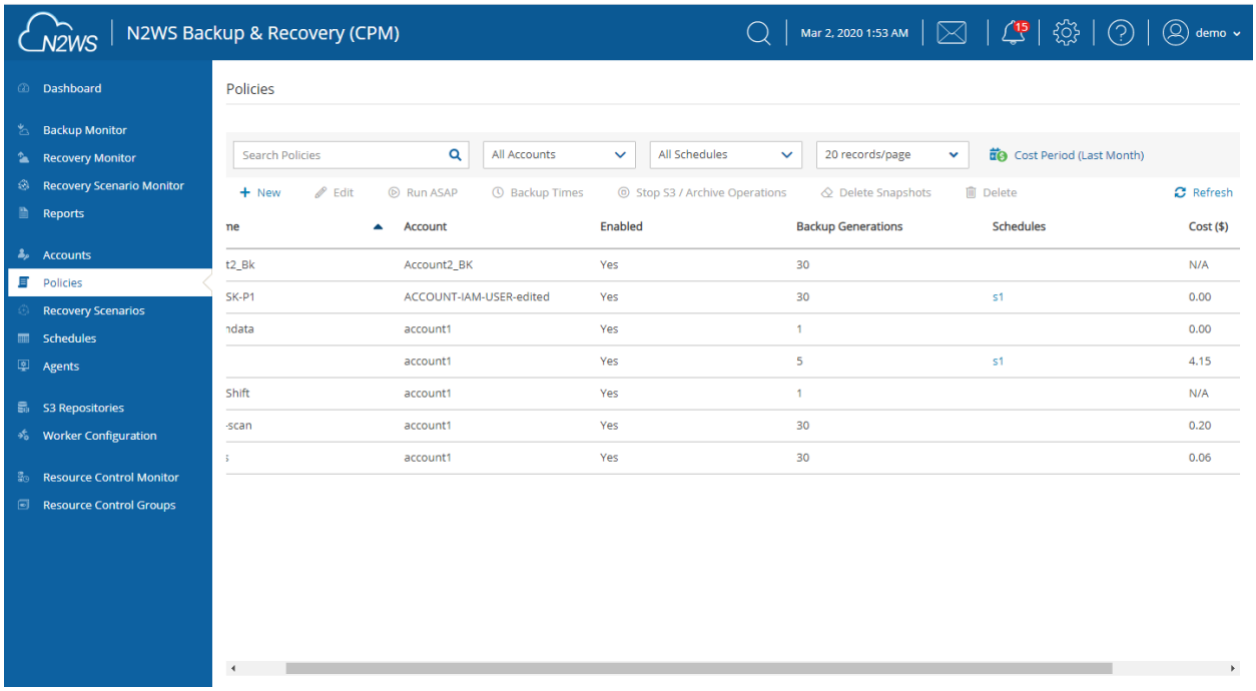

If the **Allow Cost Explorer** option is not enabled for the logged in user, or if the backup was generated within the last 24 hours, the **Cost ($)** column will show 'N/A'. To enable Cost Explorer for a user, see section 18.3.

25.2.1 Specifying a Different Time Period for Cost Calculations

You can monitor costs for a different time period by setting the **Cost Period** for all policies. The maximum period is one year. The current period is shown next to **Cost Period** in the **Policies** tab below the filters.

To specify the period for cost calculations:

1. Select the **Policies** tab and then select **Cost Period**.

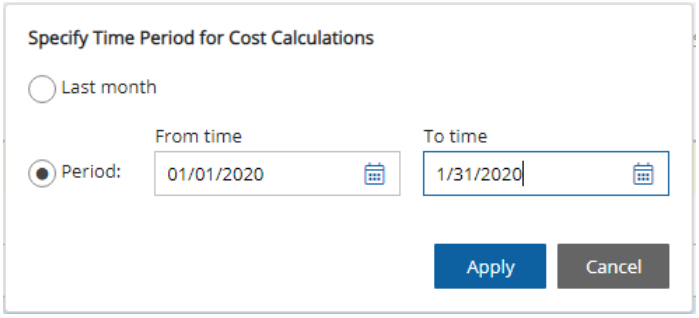

2. Select **Period**.
3. Choose the **From** and **To** dates from the calendars, selecting **Apply** after each date.

25.3 Monitoring Expected Cost Savings

In the **Resource Control Groups** tab, you can monitor the expected Cost Savings for each group, based on the schedules you have set to Turn Off an instance or an RDS database.

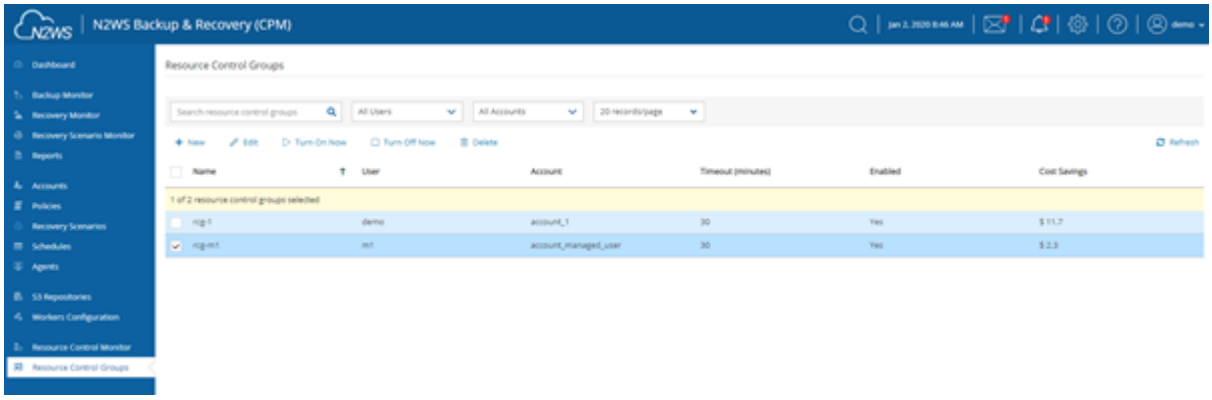

Note: When the **Operation Mode** of a Resource Control Group to **Turn Off Only**, CPM will show 'No-Data' in the **Cost Savings** column.

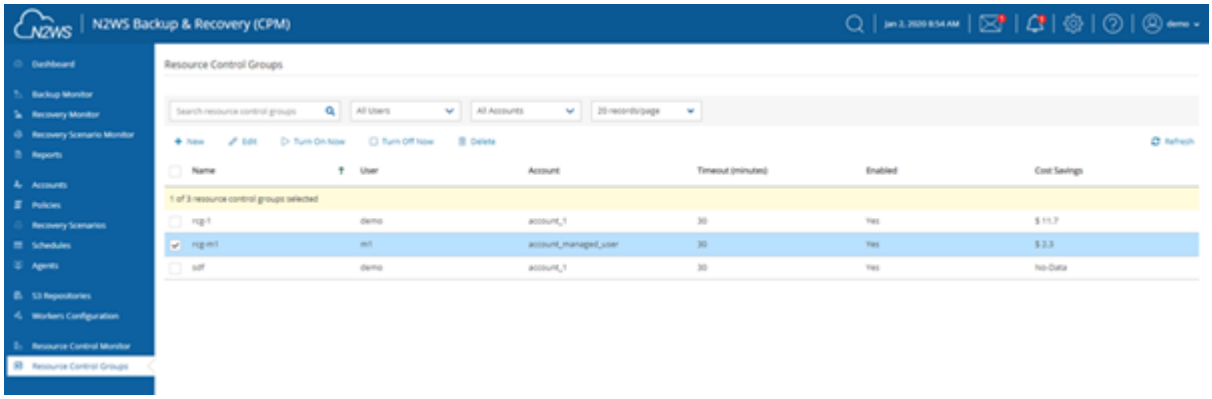

Appendix A – Recommended Configuration for Copy to S3

When 'worker' instances are using public IP, NAT, or IGW within a VPC to access S3 buckets within the same region/account, it results in network transfer fees:

https://www.linkedin.com/pulse/keep-s3-traffic-private-your-vpc-aws-travis-haag/

https://medium.com/nubego/how-to-save-money-with-aws-vpc-endpoints-9bac8ae1319c

If the bucket is in another region or in another account, the transport charges will be incurred anyway.
Using VPC endpoint enables instances to use their private IP to communicate with resources in other services, such as S3, **within the AWS network without incurring network transfer fees**.

1. To create a subnet associated with a route table that will direct connections to S3 in the same region as the VPC endpoint:
2. In AWS, create a subnet within VPC of the region.

After successful creation, the successful creation message appears.

The subnet is automatically associated with the default route table.

3. Create a new route table.

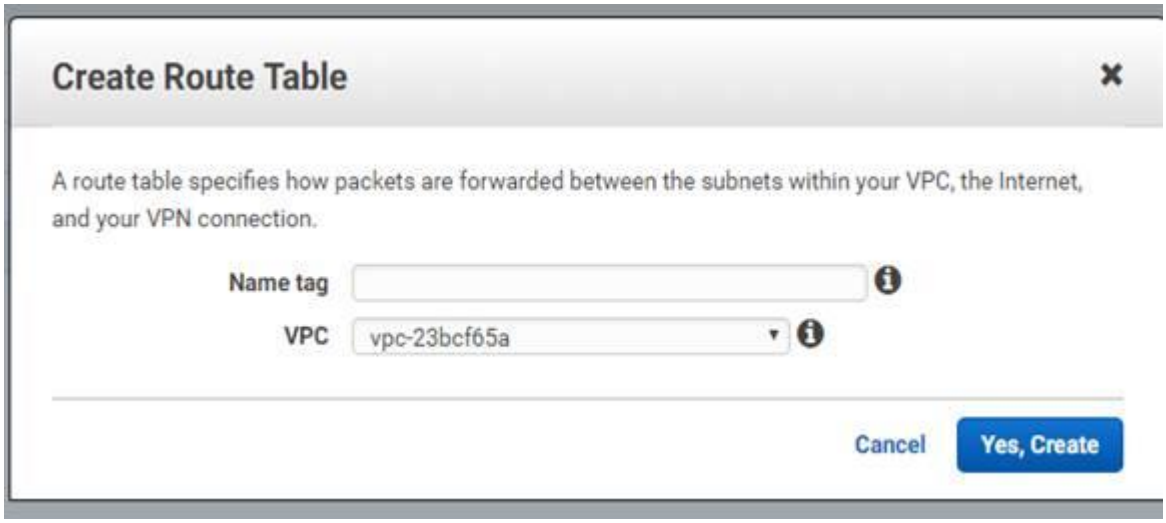

4. Change the subnet association by associating the previously created subnet with this route table.
5. Create a VPC endpoint for S3 in the region and associate it with the previously created route table.

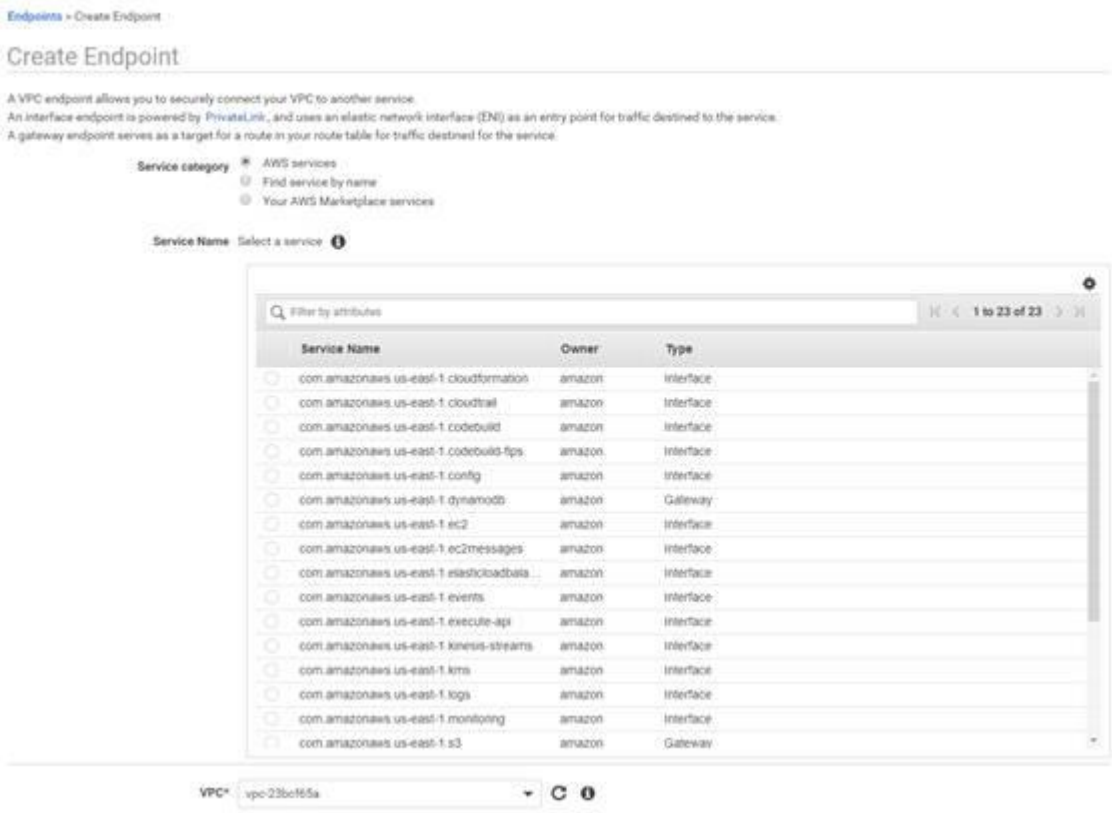

6. Choose a region.

7. Then choose the previously defined route table.

The permissions to access the bucket will be defined by the IAM policies attached to the roles of the N2WS.

8. Grant Full Access.

The route table of the subnet now looks like the following:

9. If N2WS is in a different account/region/VPC, add to the route table an Internet Gateway so the 'worker' can communicate with N2WS. Add the following rule:

The route table will look like:

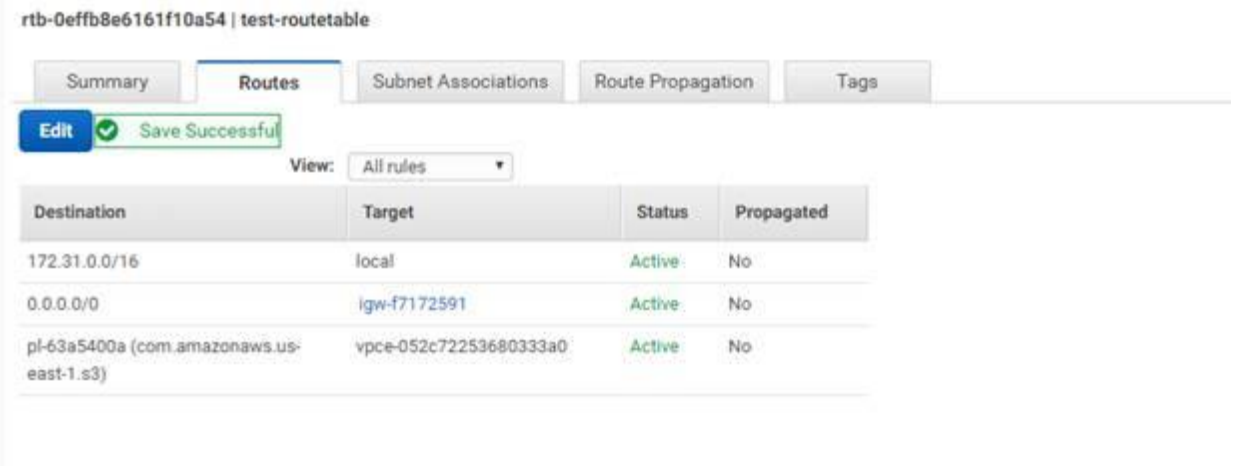

In this configuration, the connection to S3 will be routed to the VPC endpoint. See Note at the end of this section.

10. In N2WS, select the **Worker Configuration** tab, and then select **+ New**. Configure the worker to use this subnet in the specific region and the VPC where it is defined.

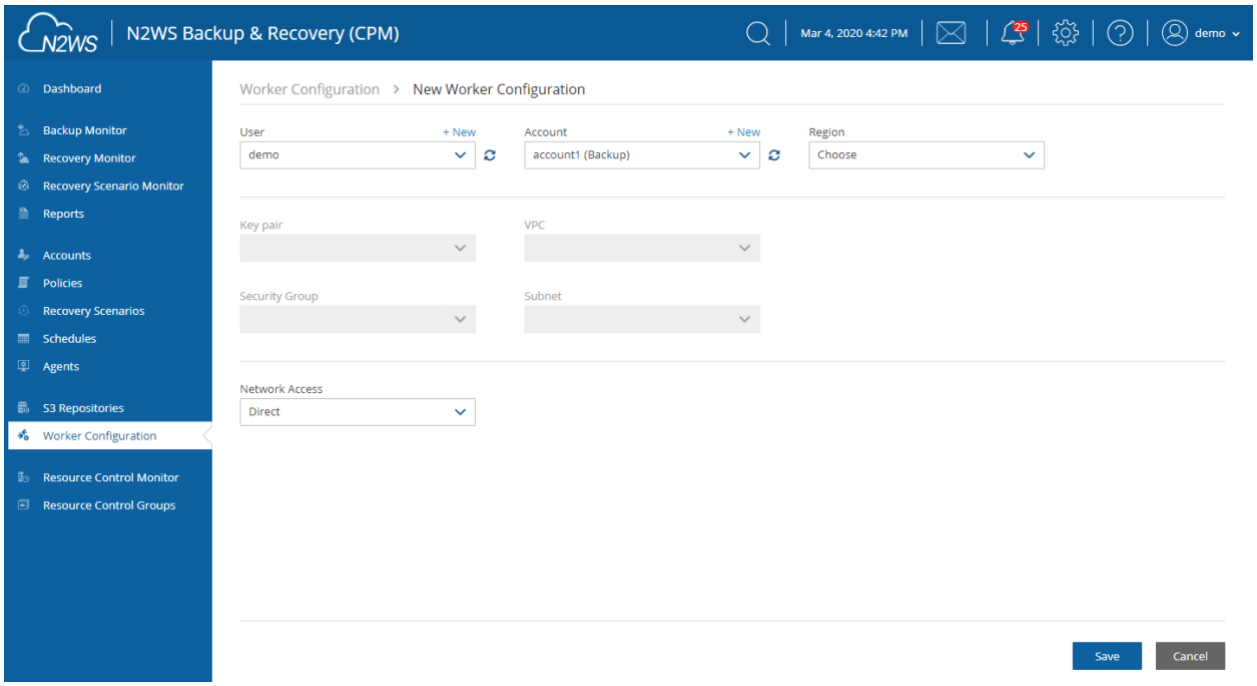

Note:

Example: An Endpoint Route in a Route Table

In this scenario, you have an existing route in your route table for all internet traffic (`0.0.0.0/0`) that points to an internet gateway. Any traffic from the subnet that's destined for another AWS service uses the internet gateway.

Destination	Target
10.0.0.0/16	Local
0.0.0.0/0	igw-1a2b3c4d

You create an endpoint to a supported AWS service, and associate your route table with the endpoint. An endpoint route is automatically added to the route table, with a destination of `pl-1a2b3c4d` (assume this represents the service to which you've created the endpoint). Now, any traffic from the subnet that's destined for that AWS service in the same region goes to the endpoint, and does not go to the internet gateway. All other internet traffic goes to your internet gateway, including traffic that's destined for other services, and destined for the AWS service in other regions.

Destination	Target
10.0.0.0/16	Local
0.0.0.0/0	igw-1a2b3c4d
pl-1a2b3c4d	vpce-11bb22cc

Example: Adjusting Your Route Tables for Endpoints

In this scenario, you have configured your route table to enable instances in your subnet to communicate with Amazon S3 buckets through an internet gateway. You've added a route with `54.123.165.0/24` as a destination (assume this is an IP address range currently within Amazon S3), and the internet gateway as the target. You then create an endpoint, and associate this route table with the endpoint. An endpoint route is automatically added to the route table. You then use the describe-prefix-lists command to view the IP address range for Amazon S3. The range is `54.123.160.0/19`, which is less specific than the range that's pointing to your internet gateway. This means that any traffic destined for the `54.123.165.0/24` IP address range continues to use the internet gateway, and does not use the endpoint (for as long as this remains the public IP address range for Amazon S3).

Destination	Target
10.0.0.0/16	Local
54.123.165.0/24	igw-1a2b3c4d
pl-1a2b3c4d	vpce-11bb22cc

To ensure that all traffic destined for Amazon S3 in the same region is routed via the endpoint, you must adjust the routes in your route table. To do this, you can delete the route to the internet gateway. Now, all traffic to Amazon S3 in the same region uses the endpoint, and the subnet that's associated with your route table is a private subnet.

For additional information about setting up VPC Gateway Endpoints, see https://docs.aws.amazon.com/vpc/latest/userguide/vpce-gateway.html

Appendix B – Agents Configuration Format

N2WS allows configuring remote and local agents from the UI. See section 6.1.2.

- The configuration in the text box needs to be in '`INI`' format.
- According to the section header, N2WS will pass the key-value pairs to the appropriate agents.
- Each agent writes the set of key-value pairs it receives for a section to its configuration file and restarts to reload the configuration.

To configure agents:

1. Select  **Server Settings** > **Agents Configuration**.
2. Write the configuration in the text box with the section header followed by its key-pair, as shown below.
3. Select **Publish**.

The following sample rules show how to configure relevant agents:

- Pass configuration to **all remote agents of** a given **policy**.
  The following will pass the key-value '`max_seconds_to_wait_for_vss=100`' to all remote agents that belong to policy by the name '`p1`':
  ```
  [policy__p1]
  max_seconds_to_wait_for_vss=100
  ```
- Pass configuration to **specific remote** agent.
  The following will pass the key-value '`max_seconds_to_wait_for_vss=100`' to the remote agent whose AWS instance ID is '`agent_id` ':
  ```
  [agent__agent_id]
  max_seconds_to_wait_for_vss=100
  ```
- Pass configuration to **all remote** agents.
  The following will pass the key-value '`max_seconds_to_wait_for_policy=600`' to all remote agents:
  ```
  [all_remote_agents]
  max_seconds_to_wait_for_policy=600
  ```
- Pass configuration to **local** agent.
  The following will pass the key-value '`max_seconds_to_wait_for_policy=600`' to the local agent:
  ```
  [local_agent]
  max_seconds_to_wait_for_policy=600
  ```

One or more instances of all of the above can be pasted together to the text box in the **Agent Configuration** screen. On **Publish**, N2WS iterates over all sections and passes the relevant configuration to each agent.