

CPM – User's Guide

V2.2.0

Contents

1	Introduction to CPM	5
1.1	Purchasing CPM on the AWS Marketplace	6
1.2	CPM Architecture	7
1.3	CPM Server Instance	8
1.4	CPM Technology.....	10
1.5	Browser Support	10
2	Configuring CPM	11
2.1	Instance ID and License Agreement.....	12
2.2	Root User.....	12
2.3	Defining a Time Zone and Data Volume	13
2.4	Complete Remaining Fields in CPM Configuration	14
2.5	Registering and Finalizing the Configuration	16
2.6	Configuration Troubleshooting.....	17
2.7	Modifying the Configuration of a CPM Server	18
2.8	Configuring CPM in Silent Mode	18
3	Start Using CPM	20
3.1	Main Screen	20
3.2	Associating an AWS Account.....	22
4	Defining Backup Policies.....	24
4.1	Schedules	24
4.2	Policies.....	26
5	Consistent Backup.....	32
5.1	Crash-Consistent Backup.....	32
5.2	Application-Consistent Backup	32
5.3	CPM and a Point in Time	32
5.4	Summary or What Type of Backup to Choose	33
6	Windows Instances Backup	34
6.1	Configuring CPM Thin Backup Agent	34
6.2	Using VSS.....	38
6.3	Using Backup Scripts on Windows	42
7	Linux/Unix Instances Backup	45
7.1	Connecting to the CPM Server	45
7.2	Backup scripts	45
8	Additional Backup Topics	50
8.1	CPM in a VPC Environment	50
8.2	Backup when an Instance is Stopped.....	50
8.3	The Freezer.....	51
8.4	Backing up Independent Volumes	51

9	Performing Recovery.....	52
9.1	Recovery AWS credentials	52
9.2	Instance Recovery	53
9.3	Volume Recovery	60
9.4	RDS Database Recovery	62
9.5	Aurora Cluster Recovery	63
9.6	Redshift Cluster Recovery	64
10	Disaster Recovery (DR).....	66
10.1	Configuring DR	66
10.2	About the DR Process.....	67
10.3	DR and mixed-region policies	67
10.4	Planning your DR Solution.....	68
10.5	DR Recovery	69
10.6	DR Monitoring and Troubleshooting	72
11	Cross-Account DR, Backup and Recovery	75
11.1	Snapshot Vaulting	75
11.2	Configuring Cross-Account Backup	75
11.3	Cross-Account DR and Clean-Up	76
11.4	Cross-Account with Cross-Region	76
11.5	Cross-Account Recovery.....	77
11.6	Cross-Account Backup and Cost.....	77
12	File-level Recovery	78
13	Tag-based Backup Management.....	80
13.1	The “cpm backup” Tag	80
13.2	Tag Scanning.....	82
13.3	Pitfalls and Troubleshooting	83
14	Security Concerns and Best Practices	85
14.1	CPM Server.....	85
14.2	Best Security Practices for CPM	86
14.3	Using IAM	87
14.4	Thin Backup Agent	92
15	Alerts, Notifications and Reporting.....	93
15.1	Alerts	93
15.2	Pull Alerts	93
15.3	Using SNS.....	95
15.4	Push Alerts	96
15.5	Daily Summary	97
15.6	Raw Reporting Data	98
15.7	Usage Reports	100
16	CPM User Management	101
16.1	Independent Users.....	101

16.2	Managed Users	101
16.3	User definitions	102
16.4	Delegates.....	103
16.5	Usage Reports	105
16.6	Audit Reports	105
17	CPM IdP Integration	106
17.1	Configuring IdPs to Work with CPM.....	106
17.2	Configuring Groups and Group Permissions on the CPM Side	109
17.3	Configuring Groups on the IdP Side	111
17.4	CPM Login Using IdP Credentials	113
17.5	Configuring CPM to Work with Active Directory / AD FS	125
17.6	Configuring an AD FS User Claim	127

1 Introduction to CPM

CPM – Cloud Protection Manager – is an enterprise-class backup, recovery and disaster recovery solution for the Amazon Web Services (AWS). Designed from the ground up to support AWS, CPM uses cloud native technologies (e.g. EBS snapshots) to provide unmatched backup and, more importantly, restore capabilities in AWS.

CPM 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 CPM. Using CPM, 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.

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
- Redshift Clusters

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

- Use a Windows to achieve application consistency using 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
- 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 CPM on the AWS Marketplace

CPM 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 CPM software is a fixed monthly price which varies between the different CPM editions.

To see the different features for each edition, along with pricing and details, go to our [pricing & purchase page](#) on the N2WS web site. Once you subscribe to one of CPM's editions, you can launch a CPM Server instance and begin protecting your AWS environment. Only one CPM Server per subscription will actually perform backup. If you run additional instances, they will only perform recovery operations (see section 1.3.3).

1.1.1 Moving between CPM Editions

If you are already subscribed and using one CPM 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 before proceeding. You can delete that snapshot once your new CPM Server is up and running. The data volume is typically named **CPM Cloud Protection Manager Data**.

1. Terminate your existing CPM instance. It is recommended to do so while no backup is running.
2. Unsubscribe from your current CPM 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 CPM 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 CPM 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 CPM edition.

1.1.2 Downgrading

If you moved to a lower CPM edition, you may find yourself in a situation where you exceed the resources your new edition allows. For example, you used CPM Advanced Edition and you moved to CPM Standard Edition, which allows fewer instances. CPM 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 CPM 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, CPM will automatically resume normal operations.

1.2 CPM Architecture

The CPM 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 and Redshift clusters. Except in cases where the user chooses to install our Thin Backup Agent for Windows Servers, CPM does not directly access your instances. Access is performed by the agent, or by a script that the user provides, which performs application quiescence.

CPM consists of three parts, all of which reside on the CPM 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 CPM server

The CPM architecture is shown in Figure 1-1. CPM Server is an EC2 instance inside the cloud, but it also connects to the AWS infrastructure to manage the backup of other instances. CPM does not need to communicate or interfere in any way with the operation of other instances. The only case where CPM 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 Volume Shadow Copy Service (VSS) or script support for application quiescence, you will need to install the CPM Thin Backup Agent. The agent will get its configuration from the CPM server, using the HTTPS protocol.

CPM Solution Architecture




Figure 1-1

1.3 CPM Server Instance

The CPM 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 CPM's perspective, the root device is dispensable. You can always terminate your CPM 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 CPM Server instance via SSH, N2WS expects the CPM Server instance will be used as a virtual appliance. N2WS 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 CPM, N2WS will not be able to provide you with support. Our first requirement will be for you to launch a clean CPM server.

Note: Remember that all your changes in the OS will be wiped out as soon as you upgrade to a new release of CPM, 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. N2WS recommends that you consult [N2WS support](#) before doing so.

1.3.2 Backing up the CPM Server

CPM 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 CPM servers. What you need to do, is to back up only the data volume, and to launch a fresh CPM server and connect it to a recovered data volume (see section 10.4.3).

1.3.3 CPM Server with HTTP Proxy

CPM needs connectivity to AWS endpoints to be able to use AWS APIs. This requires Internet connectivity. If you need CPM 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. In any event, proxy settings can be modified at any time in the general settings screen in the main CPM application.

1.3.4 Multiple CPM Servers

If you are trying to launch multiple CPM 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 CPM server should be enough to back up your entire EC2 environment. If you need more resources, you should upgrade to a higher edition of CPM. If you do need to use more than one CPM server in your account, contact [N2WS support](#).

1.3.5 Upgrading the CPM Server Instance

At certain times, you may need to terminate the current CPM Server instance and start a fresh one. The typical scenario is upgrading to a new CPM image.

To upgrade/restart the CPM Server Instance:

1. Launch a new CPM Server instance in the same region and AZ as the old one. You can launch the instance using the [Your Software](#) page on the AWS web site.

To determine the AZ of the new instance or to launch it in a Virtual Private Cloud (VPC) subnet, launch the instance using the EC2 console rather than using the 1-click option.
2. Terminate the old instance, preferably while no backup is being performed. Wait until it is in the **terminated** state.

Recommended: Go to the Volumes view in the AWS Management Console and create a snapshot of the CPM data volume. The volume is typically named **CPM Cloud Protection Manager Data**. The snapshot is only needed in the event there is a problem with the upgrade process and it can be deleted afterwards.

3. When the new instance is in the **running** state, connect to it with a browser using HTTPS.
4. Approve the exception to the SSL certificate.
5. Choose **Use Existing Data Volume** and paste in your AWS credentials.
6. Select your old data volume from the list of volumes to complete the configuration process.
Operations will resume automatically.

If you are using backup scripts that utilize SSH, you may need to login to the CPM Server once and run the scripts manually, so the use of the private key will be approved.

1.4 CPM Technology

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

1.5 Browser Support

Most interactions with the CPM server are done via a web browser. Since CPM uses modern web technologies, you will need your browser to be enabled for Java Script. CPM supports Firefox, Safari, Google Chrome, and Microsoft Internet Explorer (version 9 and newer). CPM will not work for IE versions 8 and older. Other browsers are not supported.

2 Configuring CPM

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

- When a new CPM 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 CPM Server, the same certificate will be used also for the main CPM application.
- Every CPM 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 CPM.

When configuring the CPM server, define the following settings:

- AWS Credentials for the CPM root user
- Time zone for the server
- Whether to create a new CPM data volume, or attach an existing one from a previous CPM server
- Proxy settings. Configure proxy settings in case the CPM 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.
- Whether to upload an SSL certificate and a private key for the CPM 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 N2WS to provide quicker and enhanced support. Registration information is not shared.

For the configuration process to work, as well as for normal CPM operations, CPM needs to have outbound connectivity to the Internet, for the HTTPS protocol. Assuming the CPM 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:

- Approve the end-user license agreement
- Define the root user name, email, and password
- Define the time zone of the CPM Server
- Fill in the rest of the information needed to complete the configuration process

2.1 Instance ID and License Agreement

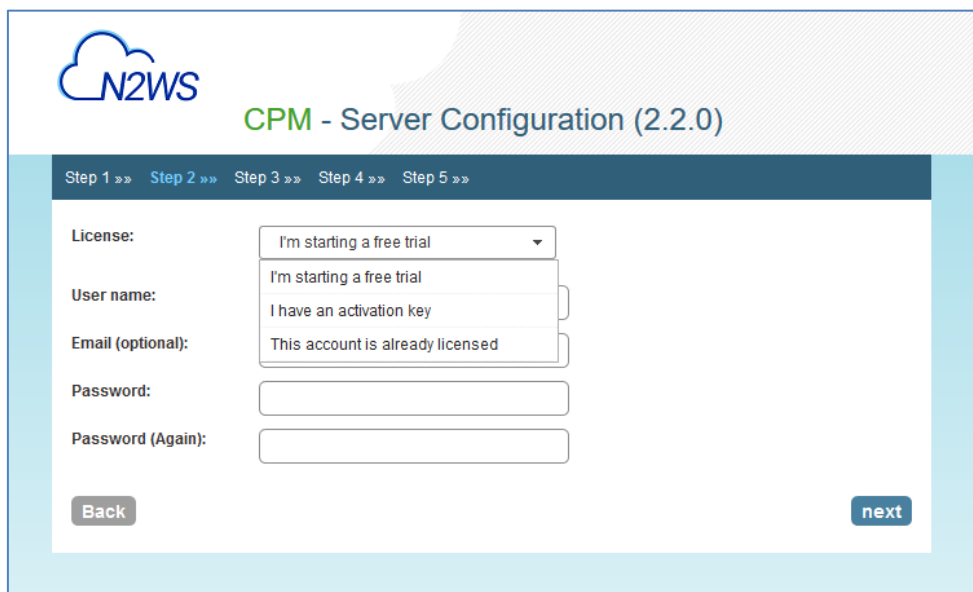
To initially be identified as the owner of this instance, you are required to type or paste the CPM server instance ID. This is just a security precaution. **In the first step of the configuration process**, you will also be required to approve the end-user license agreement.

2.2 Root User

The root user is the user who controls all the operations of the CPM server. Root user credentials are used to log on to the system and operate it. As in Figure 2-1, you need to define the root user name, email, and password. This is the second 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.

Also, if using the Free Trial or Bring Your Own License (BYOL) Edition, the **License** field is presented. Select **I'm starting a free trial** for a free trial. Alternatively, if your organization purchased a license directly from N2W Software, additional instructions are shown.

Note: Passwords: N2WS does not enforce any password policy, however, it is recommended to use passwords that are difficult to guess and that are changed from time to time.



The screenshot displays the N2WS CPM - Server Configuration (2.2.0) interface. At the top, the N2WS logo and title are visible. Below the title, a progress bar indicates the current step: Step 1 >> Step 2 >>> Step 3 >> Step 4 >> Step 5 >>. The main form area contains the following fields and options:

- License:** A dropdown menu with the following options: "I'm starting a free trial" (selected), "I'm starting a free trial", "I have an activation key", and "This account is already licensed".
- User name:** A text input field.
- Email (optional):** A text input field.
- Password:** A text input field.
- Password (Again):** A text input field.

At the bottom of the form, there are two buttons: "Back" and "next".

Figure 2-1

2.3 Defining a Time Zone and Data Volume

In the **third step of the configuration process**, define the time zone of the CPM Server. Choose whether to create a new data volume, or use an existing one, and you need to enter your AWS credentials that will be used for the data volume setup process. Additionally, you can configure proxy settings for the CPM server.

As you will see in section 4.1.2, all scheduling of backup is done according to the local time of the CPM Server. You will see all time fields displayed by local time, however, all time fields are stored in the CPM 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 in Figure 2-2, the choice of new or existing data volume is done here. Actual configuration of the volume will be done 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 CPM 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 14.3). 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 CPM 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 Proxy Settings

If the CPM server needs an HTTP proxy to connect to the Internet, in the **Connect via web proxy** drop-down list, choose **Enabled**. Define the proxy address, port, user, and password. The proxy settings will be kept as the default for the main application.

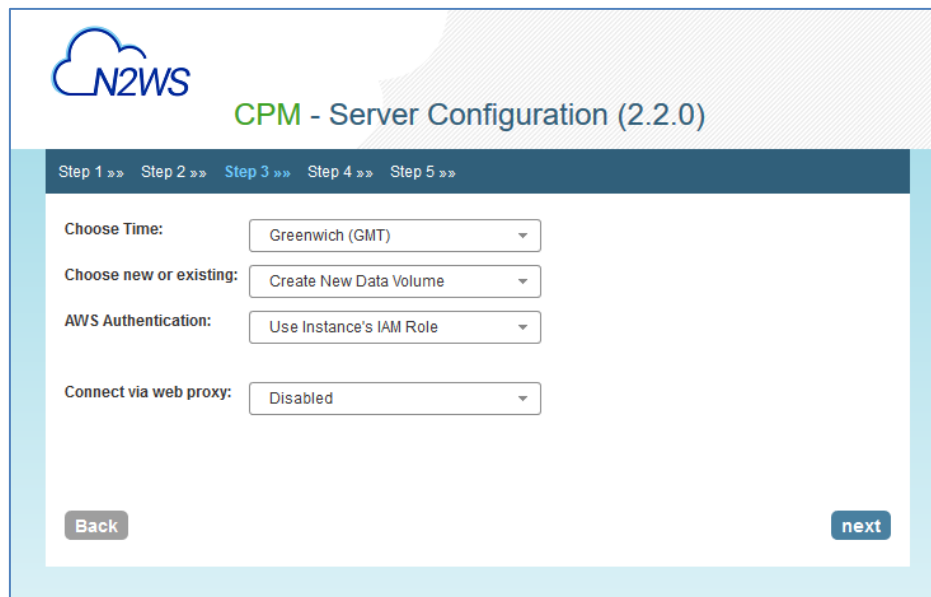


Figure 2-2

2.4 Complete Remaining Fields in CPM Configuration

In the fourth step, you will fill in the rest of the information needed for the configuration of the CPM Server.

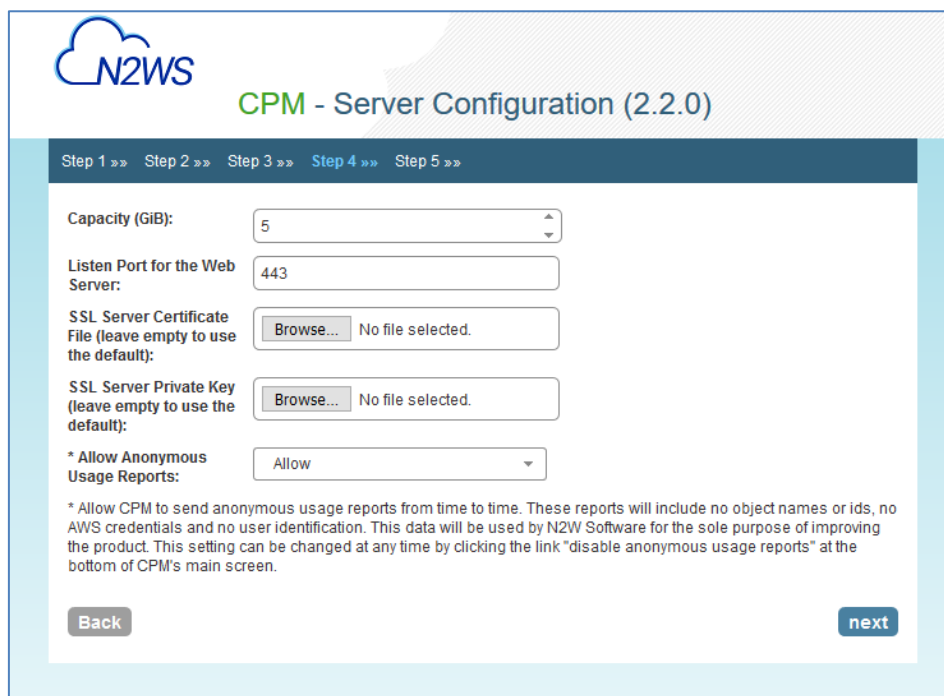


Figure 2-3

First thing you need is to finish configuring your data volume.

- If you chose to create a new volume in the previous step, you will see the screen as in Figure 2-3.

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

2.4.1 New Data Volume

When creating a new data volume, the only thing you need to define is the capacity of the created volume. The volume is going to contain the database of CPM'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 CPM instance. It will be named **CPM Cloud Protection Manager Data**. During the configuration process, the volume will be created and attached to the instance. The CPM database will be created on it.

2.4.2 Existing Data Volume

The Existing data volume option is used if:

- You have already run CPM and terminated the old CPM server, but now wish to continue where you stopped.
- You are upgrading to new CPM releases.
- You are changing some of the configuration details.

The select box for choosing the volumes will show all available EBS volumes in the same AZ as the CPM 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 name that was not created by a CPM server for an existing data volume, the application will *not* work.

2.4.3 Web Server Settings

Port 443 is the default port for the HTTPS protocol, which is used by the CPM 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.

2.4.4 Anonymous Usage Reports

Leaving the Anonymous Usage Reports value as **Allowed** enables CPM 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 CPM names of objects names, such as, policy and schedule.

It contains only details like:

- How many policies run on a CPM server
- How many instances per policy
- How many volumes
- What the scheduling is, etc....

You can change this setting at any time in the links at the bottom of CPM'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.

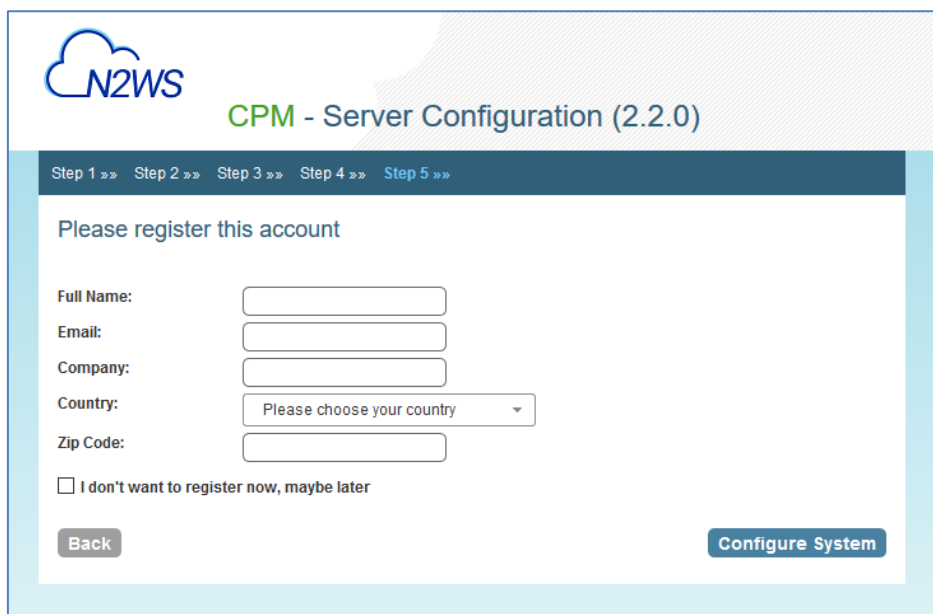


Figure 2-4

Click **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 successful configuration notification page opens.



Figure 2-5

Click the **here** link. After a few seconds, you are redirected to the login screen of the CPM application. If you are not redirected, refresh the browser manually. If you are still not redirected, reboot the CPM server via AWS Management Console or another management tool, 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 click **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 CPM 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 CPM 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 CPM 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 CPM Server

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

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

The process to make these changes is to terminate the current CPM server instance and create a new one. After you terminate the CPM 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 CPM 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, CPM 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 CPM server. It will contain the username. You can connect to CPM server using SSH to view this file (see section 7.1).

2.8 Configuring CPM in Silent Mode

From version 2.1.0, there is an option to configure CPM using a special “user data” script. The **user data** script is a configuration in `ini` file format, stating the configuration of the new CPM instance.

Create the **user data** file with `CPMCONFIG` in the first line, `[SERVER]` in the second line, followed by the configuration details.

CPM assumes that the CPM 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 cpm 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>

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>
```

Additionally, if you need the CPM 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 GUI. It can be used for automation of a Disaster Recovery (DR) server recovery. Additionally, if you state a volume ID from another AZ, CPM will attempt to create a snapshot of that volume and migrate it to the AZ of the new CPM server. This option can be used in a high availability setup.

Note: You are not required to click 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 CPM

3.1 Main Screen

As soon as you log on to CPM with the root user credentials you created during configuration, you are redirected to the main screen. CPM is a very simple application to work with. The user interface is simple, intuitive, and user-friendly. Most operations are only one mouse-click away from the main screen.




Figure 3-1

As you can see in Figure 3-1, the main screen is divided by five tabs:

- **Backup Monitor** – Here you will see all your backups. For each backup you can see the start and end times, policy, status and DR status. All operations regarding a backup are present in this tab: viewing the list of snapshots, opening the backup log, recovering from a backup, and moving it to the freezer (see section 8.3).

Sometimes you have many backups and are looking for a specific one. You can filter by policy and status, sort by all relevant columns, or browse between pages. You can also choose how many records to view in one page.

- **Policies** – Backup Policies defined in the system. From this tab you can create, modify, configure and delete backup policies.
- **Schedules** – Backup Schedules can be created, configured and deleted in this tab. You attach a schedule to a policy in the policy definition screen.
- **Agents** – Thin Backup Agents that are connected to this CPM server can be viewed here. Currently, Thin Backup Agents are needed only when application consistency is needed for Windows Servers. In any other case, the backup is done agent-less.

- **Freezer** – The freezer is a place where you can keep backups indefinitely. When you identify a backup that is worth keeping (e.g. a successful backup of a clean system right after an upgrade), you can move it to the freezer. Elements in the freezer will not be deleted by the automatic cleanup process.
- **Recovery Monitor** - This tab will contain records for all recovery operations. Each recovery record will contain a time stamp of the recovery operation, the backup is was recovered from and additional information. Recovery records are automatically deleted as the backups are.

In addition to the tabs, you have a logout link at the top right corner of the screen, and a top panel of buttons:

- **Main** – Brings you back to the main screen from wherever you are or reload the whole page.
- **Manage AWS Accounts** – Depending on the edition of CPM you subscribed to, you can define one or more AWS accounts to work with. These accounts contain the objects (instances, EBS volumes, RDS databases, Aurora clusters and Redshift clusters) you may wish to back up. Each backup policy is associated with a single AWS account.
- **Change Password** – Changes the password for the logged-in user.
- **Notifications** - Define notifications and alerts.
- **Manage Users** – Depending on the CPM edition you subscribed to, if you are the root user, click the Manage Users button to create and manage users. Managing includes the ability to:
 - Delete users.
 - Reset passwords.
 - Download usage reports.
- **General Settings** – Contains some settings you can control, including tag scan settings, when to run cleanup, and how long to save deleted records and user audit logs.

At the bottom of the screen you can find a few useful links:

- To view the license agreement.
- To download the Thin Backup Agent.
- To enable or disable sending anonymous usage reports.
- To download the CPM logs as a tarball, in case you need to send to our support team.
- To enter a new activation key. If a special permission is required in addition to the default permissions of your CPM edition, N2W Software can issue you an activation key.
- To download a backup view or snapshot view raw report in CSV format.
- To download usage reports.
- To download user audit reports.

- To register the CPM instance account with N2W Software. It is recommended that you register if you did not do so during configuration. Registering enables N2WS to provide enhanced support.
- To go to the **cpm patches** page to install patches.
- To send configurations to agents.

3.2 Associating an AWS Account

To associate an AWS account, 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 CPM server instance or cross-account roles.

To manage your users and roles and obtain AWS credentials:

1. Go to the IAM console at <https://console.aws.amazon.com/iam/home?#users>.
2. Click **Manage AWS Accounts**.
3. Click **Add New Account**.
4. Type a unique and meaningful name for the account.

3.2.1 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 is the default.
- **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 CPM to delete snapshots of this account implies that the presented IAM credentials do not have the permission to delete snapshots.

3.2.2 Authentication

CPM Supports three methods of authentication:

- **IAM User** - Authentication using IAM credentials, access and secret keys.
- **CPM Instance IAM Role** – If an IAM role was assigned to the CPM server at launch time, you can use that IAM role to manage backups in the same AWS account the CPM server is in. Only the root/admin CPM user is allowed to use the IAM role.
- **Assume Role** – This type of authentication requires another AWS account already configured in CPM. 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 CPM:

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

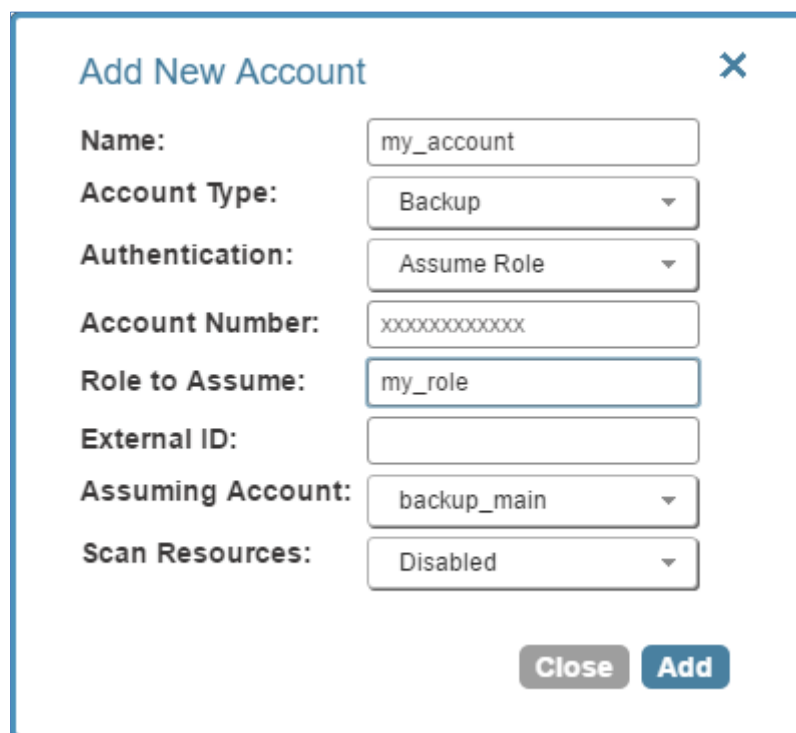


Figure 3-2

6. If you are the root user and have managed users defined, an additional selection list will appear enabling you to select the user.
7. In the **Scan Resources** list, choose whether the current account will be included in scan tags performed by the system. Once **Scan Resources** is **Enabled**, you may choose in which region to scan resources. By default, CPM will scan all regions, but you can disable any region which is not relevant to your deployment.

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

4 Defining Backup Policies

The backbone of the CPM 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 and 2012 only).
- Whether backup is performed via a backup agent (Windows only).
- The retry policy in case of failure.
- DR settings for the policy.

The following sections explain the stages for defining a policy.

4.1 Schedules

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

- Schedules are defined separately from policies.
- One schedule can be associated with several policies.
- Multiple schedules can be associated with the same policy.

4.1.1 Defining Schedules

All backup times are derived from the start time.

To define a schedule:

1. In the main screen, click the **Schedules** tab and then click **New Schedule**.
2. Type a name for the schedule and an optional description.
3. In the **Repeats Every** list, select the frequency of the backups for this schedule. The possible units are months, weeks, days, hours, and minutes.
4. In the **Start Time** list, select the schedule start time.
 - 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 the start time to XX:17.
 - The date can also be set. The default is the current day.

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

5. In the **End Time** list, select when the schedule will expire. By default, it is never. Furthermore, you can define which weekdays the schedule will be active on.

For the root/admin user, if you have created additional managed users, you will be able to select to which user the schedule belongs.

4.1.2 Scheduling and Time Zones

When you configure a CPM server, its time zone is set (see section 2.3). All time values which are shown in CPM's management application are in the time zone of the CPM server.

Important: Even when you are backing up instances that are in different time zones, the scheduled backup time is always according to the CPM server's local time.

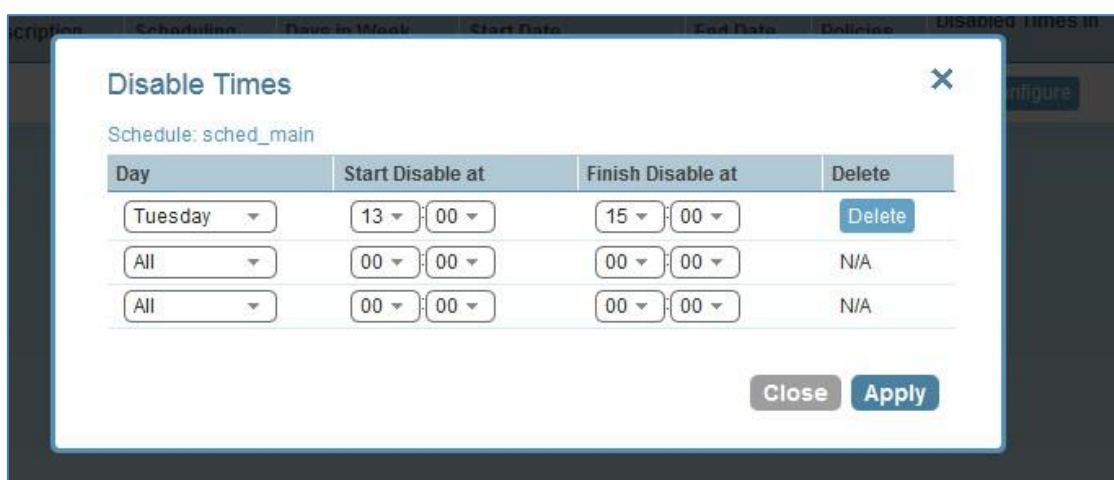
In CPM's database, times are saved in UTC time zone (Greenwich). So, if, at a later stage, you start a new CPM server instance, configure it to a different time zone, and use the same CPM data volume as before, it will still perform backup at the same times as before.

4.1.3 Disabled Times

After defining a schedule, you can set specific times when the schedule should not start a backup. For example, you want backup 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.

To define disabled times:

1. In the **Disabled Times in Day** column of the **Schedules** tab, click the **Configure** button.
2. As shown in Figure 4-1, add, edit, or remove multiple disabled times.



Day	Start Disable at	Finish Disable at	Delete
Tuesday	13:00	15:00	Delete
All	00:00	00:00	N/A
All	00:00	00:00	N/A

Close Apply

Figure 4-1

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.

Beware not to create contradictions within a schedule's definition:

- It is possible to define a schedule that will never start backups.
- 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.

4.2 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 associating schedules to the policy)

4.2.1 Creating a New Policy

To define a new policy:

1. Go to the **Policies** tab and click **New Policy**. The **Policy** window 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, can select the policy owner in the **User** box.
4. If you have more than one account, select the account the policy is associated with in the **Account** list. The account cannot be modified after the policy is already created.
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. In the **Generations to Save** list, select the number of backups to keep for this policy. Older backups will be automatically deleted by CPM.
 - If you define a daily backup and leave the value of **Generations to Save** at 30, this will give you the ability to recover from backups up to 30 days ago.
 - If you define an hourly backup, this will give you the ability to recover from backups up to 30 hours ago.
7. In the **Description** box, optionally type a description.

Note: As a user, you need to balance the amount of time you want to be able to go back and recover from Recovery Point Objective (RPO), and the cost of keeping more snapshots. Sometimes you will want to trade off the frequency of backups, and the number of generations. Consider what best suits your needs.

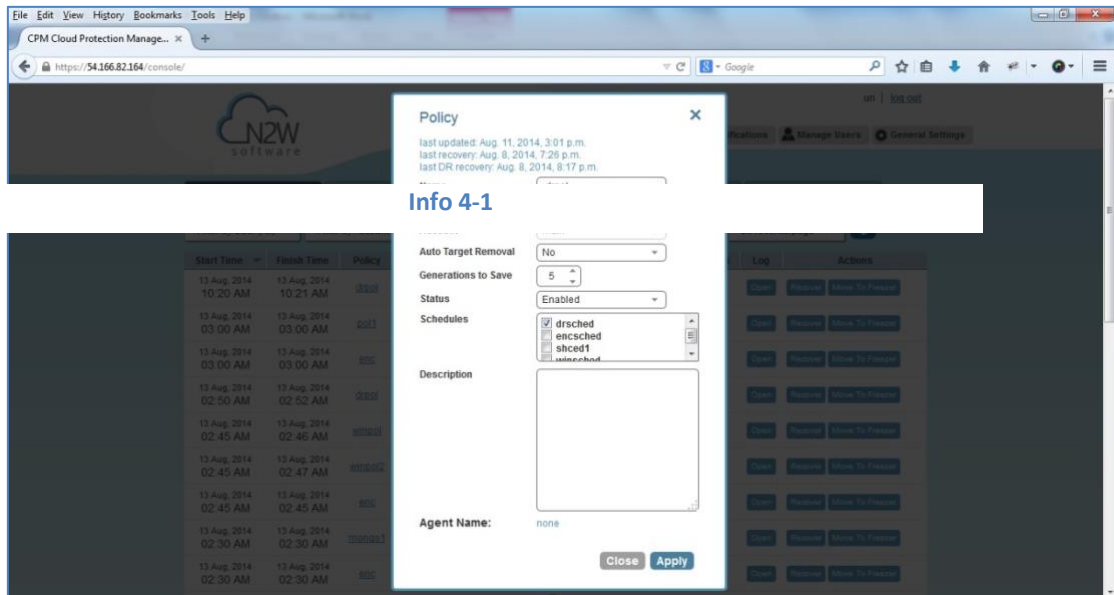


Figure 4-2

8. Click **Apply**. The new policy is included in the list of policies in the **Policies** tab.

4.2.2 Adding Backup Targets

Backup targets define what a policy is going to back up. You define backup targets by clicking the **Backup Targets** button of a policy in the **Policies** tab. You have multiple 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.
- **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 CPM 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.
- **Redshift Clusters** – You can use CPM 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.

From the Backup Targets screen, click one of the **Add** buttons to add a specific type of backup targets to the policy:

- When adding backup targets, 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 another region to see the objects in it.
- If there are many objects, you have the ability to filter, sort, or browse between pages.
- For each backup target, you can see the number of policies it is already in (**Policies** column). If the number is larger than zero, click it to see which policies it is in. See Figure 4-3.

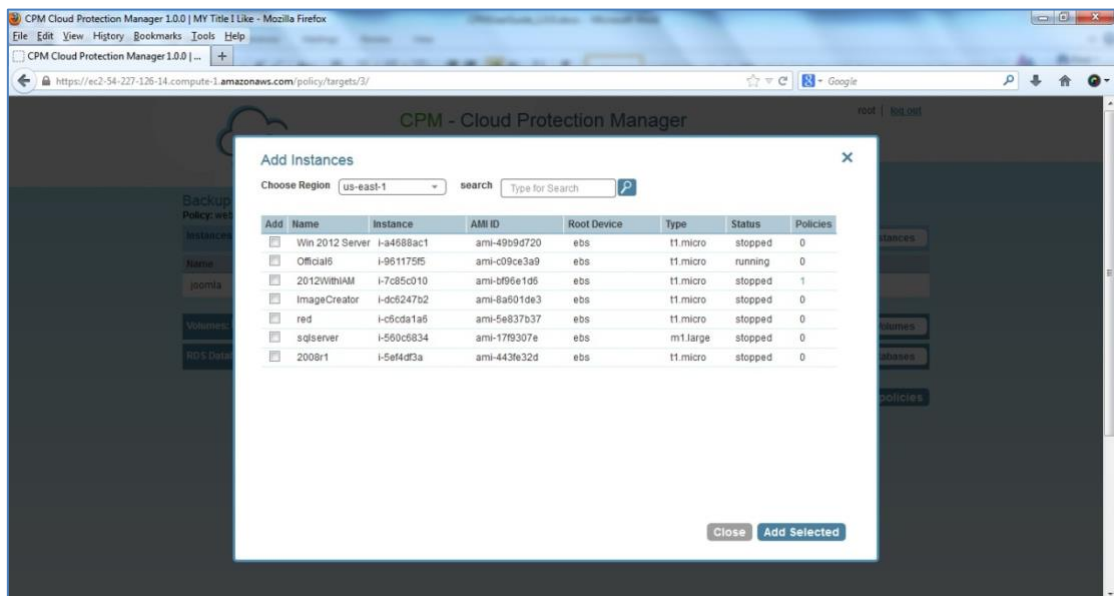


Figure 4-3

To add a backup target to the policy:

1. Select the **Add** check box of the target, or targets.
2. Click **Add Selected**. The selected objects are added to the policy's backup target list.
Repeat as many times as needed.
3. Click **Close** when finished.

4.2.3 Instance Configuration

In the case of EC2 instances, you can set options for any instance.

To configure an instance:

1. Select a policy.
2. In the **Backup Targets** screen, select an instance.
3. Click **Configure**.



Figure 4-4

4. 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, CPM will back-up all the attached storage of the instance, including volumes that are added over time.
5. In the **Backup Options** list, choose whether to:
 - Take only snapshots (the default for Linux-based instances)
 - Take an initial AMI and then snapshots (the default for Windows-based instances)
 - Just schedule AMI creation

4.2.4 AMI Creation

If you choose to just create AMIs:

- CPM 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 CPM 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. CPM will overcome this issue by scheduling a retry, which will usually succeed. However, it is recommended to avoid such scheduling conflicts.

Initial/Single AMI

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

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

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.

Note: If you move a backup into the Freezer (see section 8.3), initial AMIs will not move with it, so make sure you have an AMI ready. If the original policy exists, you can grab the AMI ID from the snapshot list of one of the backups of the policy.

4.2.5 More Options

To see more policy options, click **More Options** for a policy in the **Policies** tab.

- Backup scripts refers to those running on the CPM server (see chapter 7):
 - **Backup Scripts** – This option is turned off by default. Change to **Activate** to activate backup scripts.
 - **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.
 - **Scripts Output** – CPM 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 CPM 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 CPM. If it gets even larger, it can cause crashes in CPM processes. To avoid the risk of too much data going to `stderr`, redirect the output elsewhere.

- **Backup is 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:

- **it finishes with no 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.
- **snapshots succeed. Even if scripts or VSS fail** – 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** – Determines how much time CPM 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.
 - **Number of Failures to Trigger Alert** – The minimum number of failures to trigger an alert.

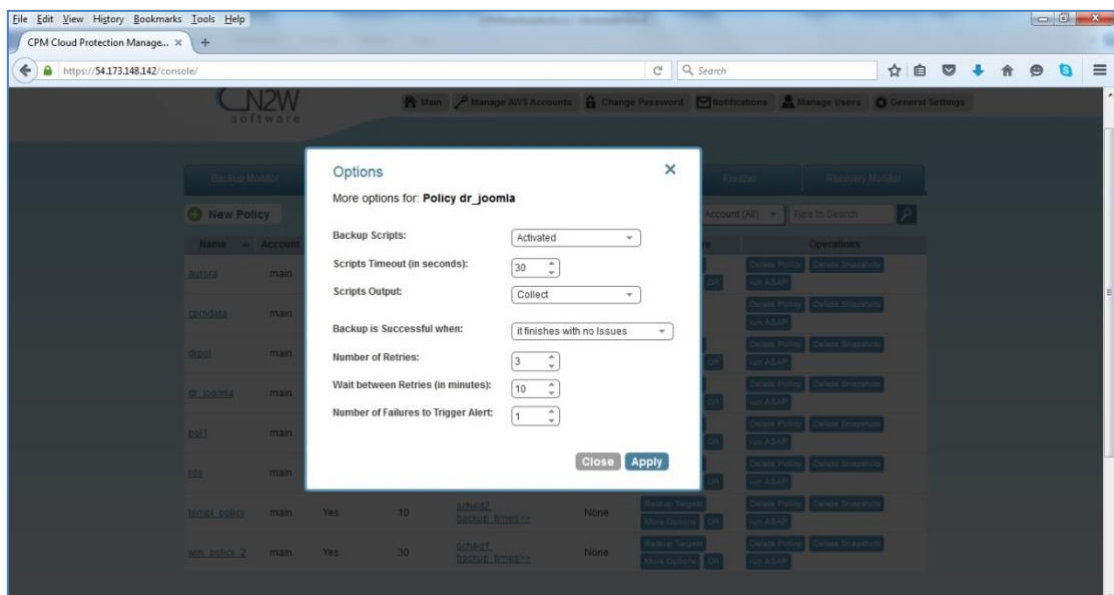


Figure 4-5

5 Consistent Backup

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

5.1 Crash-Consistent Backup

By default, snapshots taken using CPM 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 CPM 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.

In the case of taking snapshots of multiple volumes, which is probably the most common case, it would be preferable for all the volumes to be at the exact same point in time. Unfortunately, AWS does not currently support such an option. Therefore, the best CPM can offer is taking the snapshots of multiple volumes in very close succession. In most cases, it will not make a difference, but in cases where exact point in time across volumes/disks is needed, only backup scripts or VSS can achieve this goal. If the backup script of a backup policy flushes and locks all volumes in a synchronized manner, snapshots of

this policy will reflect an exact point in time. Using VSS achieves this goal, since VSS by definition performs shadow copies of multiple volumes in a synchronized manner. By freezing applications that use multiple volumes, like a database which has a volume for data and a separate volume for transaction logs, you can also achieve the goal of backing up multiple volumes at a single point in time.

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 CPM 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, CPM needs an AMI to start a recovery of a Windows instance. CPM 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, CPM 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 CPM 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 CPM Thin Backup Agent. 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 (see section 4.2.2), the next step is to configure its agent by associating it with a policy.

To associate an agent with a policy:

1. In the instance target line, select the **Configure** check box. The **Policy Instance and Volume Configuration** screen opens.

Policy Instance and Volume Configuration
Policy: win_policy_2 **Backup From:** i-400ad1fe

Which volumes:

Enabled	Name	Volume ID	Capacity	Type	IOPS	Encrypted	Zone	Status	Policies
<input type="checkbox"/>	empty	vol-f73c3a1b	20 GiB	gp2	60	no	us-east-1b	in-use	
<input type="checkbox"/>	empty	vol-f13c3a1d	30 GiB	gp2	90	no	us-east-1b	in-use	
<input type="checkbox"/>	empty	vol-c03c3a2c	20 GiB	gp2	60	no	us-east-1b	in-use	
<input type="checkbox"/>	empty	vol-be3c3a52	8 GiB	gp2	24	no	us-east-1b	in-use	

Backup Options:

Reboot:

Application-consistent backup:

Enable VSS on Agent:

Volumes for shadow copies (leave empty for all volumes):

Backup Scripts:

Scripts Timeout (in seconds):

Scripts Output:

Figure 6-1

2. In the **Application-consistent backup** list, select **Enabled**. The fields relevant for configuring an application aware backup will appear:

- **Enable VSS on Agent** – 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.

- **Volumes for shadow copies** – (This option is used only if VSS is enabled.) If you leave this field empty, VSS will create shadow copies of all of the volumes of this instance. If you want it to create shadows for only part of the volumes, you can type in drive letters with commas between them, e.g. **C;**, **D:**. For more information about VSS, see chapter 6.
- **Backup Scripts** – Whether to enable running backup scripts locally on the Windows instance.
- **Scripts Timeout** – The time given for a script to run before the CPM terminates it.
- **Script Output** – Whether to capture the output of the scripts as a log. It will capture anything the script printed to the `stderr` socket. The log will be viewable from the recovery panel screen.

6.1.2 Support for Version 1.8.0 Agents

To upgrade the agent, uninstall the previous 1.8.0 agent first.

Note: Future versions of the CPM agent will support upgrade rather than uninstall-install.

After upgrading from CPM version 1.8.0 to CPM 2.0.0:

- The old agent will continue to work as before for existing policies.

- New policies will only work with the new agents.
- New agents are on an instance level.
- Old agents used to be configured from the **More Options** screen and were on a policy level and not on the instance level.

After upgrading to CPM 2.0.0, you can open the **More Options** screen and see the agent configuration as before:

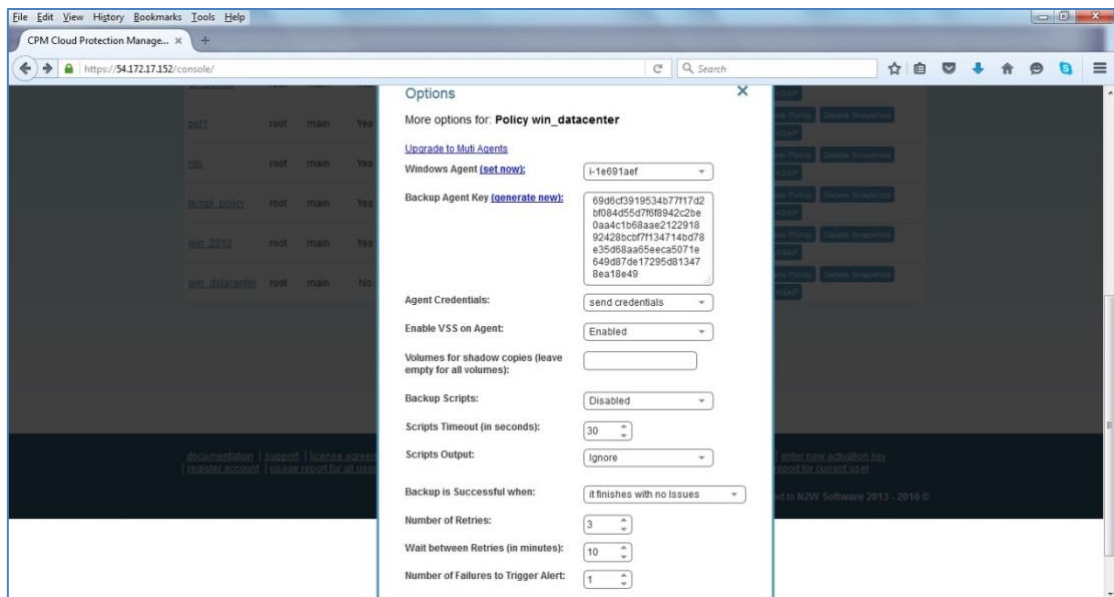


Figure 6-2

To upgrade an existing policy to the new agent, click the link **Upgrade to Multi Agents** and the policy will upgrade and will configure the agent instance.

Note: This action is irreversible. Make sure you install the 2.0.0 agent as the old agent will stop working.

6.1.3 Installing the Agent

You can download the installation package of the agent from the link **download thin backup agent** at the bottom of CPM's main screen. It will download a standard Windows `msi` package. The agent can be installed on any Windows 2003, 2008, 2012, or 2016 instance, 32 or 64-bit. After accepting the license agreement, the Setup screen opens.




Figure 6-3

The required fields are:

- The address of the CPM 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 CPM 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 agent screen in the CPM. It may take a few minutes until the CPM 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 to make a copy of the `cpmagent.cfg` file, which resides in the CPM Agent installation folder.
2. If the address or port of the CPM Server had changed, edit the agent configuration file manually. Make the change after the equation sign.
3. After making the changes, restart the **CPM 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

If the Windows instance the agent is installed on can reach the CPM server only through a proxy, CPM agent supports such a configuration.

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 CPM 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.

CPM supports VSS for backup on Windows Servers 2008 or 2012 *only*. Trying to run VSS on older 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 CPM's Use of VSS

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

- When CPM initiates a backup, it **requests** the CPM 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 CPM 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 CPM agent will request the VSS Provider to delete the shadow copies. The CPM 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 in Figure 6-4.

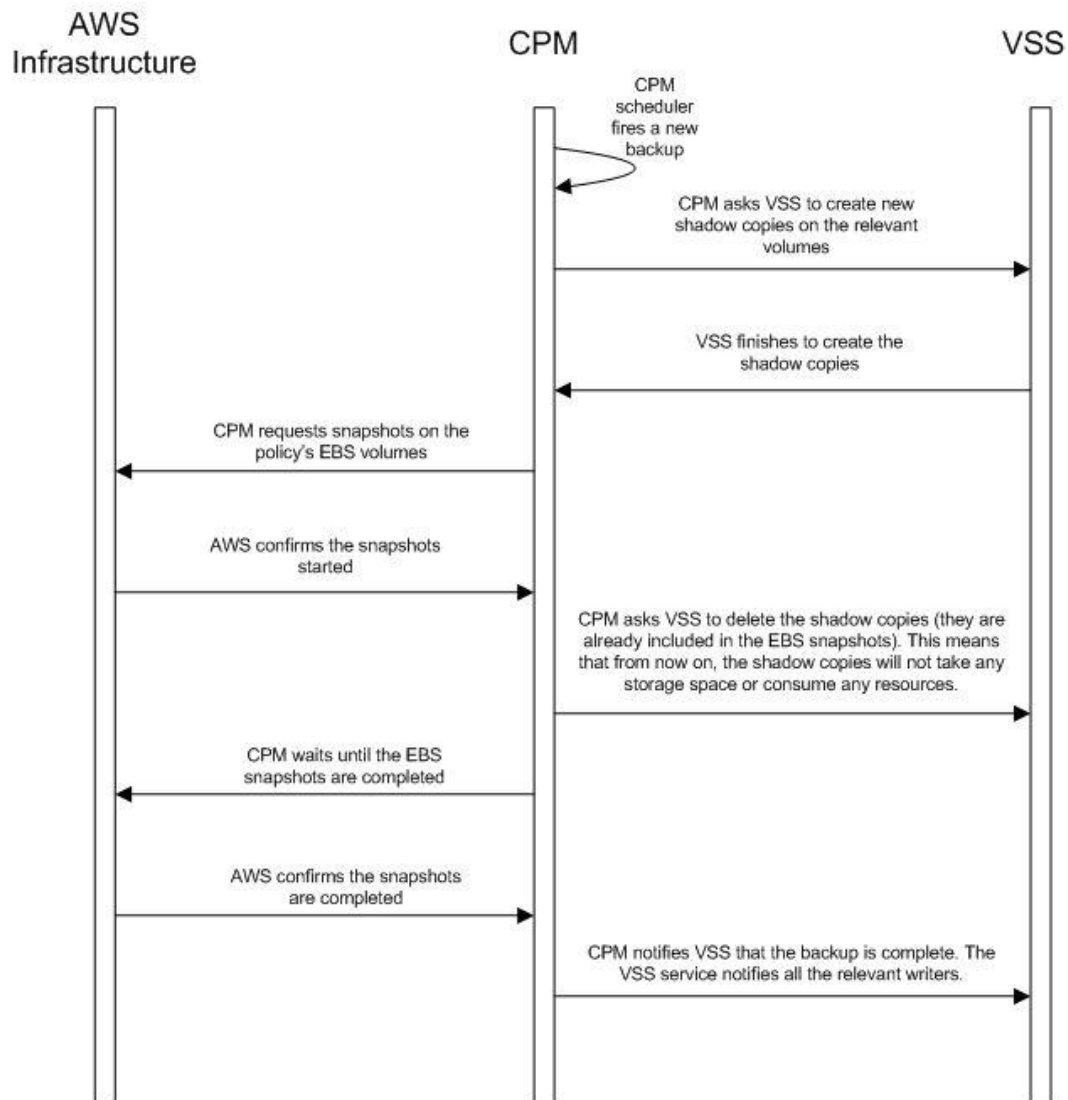


Figure 6-4

6.2.2 Configuring VSS

By default, VSS is enabled when a CPM 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:

In the Instance and Volume configuration screen, change the value of **Volumes for shadow copies**.

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 GUI.

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 CPM'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 see the errors that VSS encountered, look in the backup log.

- To see the output of the exact VSS error, click **Recover**.
- To view the VSS Disk Shadow log, click 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 CPM. 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 CPM 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 CPM 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 CPM is covered in chapter 9. 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, click 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 CPM, 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 CPM Thin Backup Agent.
- If the CPM 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, CPM 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.
- Scripts are launched by CPM 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 need to exit with the code 0 when they succeed, or 1 (or another non-zero code) when they fail.

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, CPM 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 CPM Server. See sections 7.2.4 and 4.2.5.

7 Linux/Unix Instances Backup

Making application-consistent backup of Linux instances does not require any agent installation. Since the CPM 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 done 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 CPM Server

In order to create, test, and install backup scripts, you will need to connect to the CPM 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 CPM server instance. As long as your key is not compromised, no unauthorized person will be able to connect to the CPM server.

With `cpmuser`, you will be able to copy (using secure copy), create, and test your scripts. `cpmuser` is the same user CPM will use to run the scripts. If you need to edit your scripts on the CPM 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 CPM 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 data volume of CPM, and will remain there even if you terminate the CPM 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 (e.g. `before_policy1_v11.5.bash`).
- 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. CPM 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 `xfss`.

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 CPM Server, see sections 4.2.3 and 4.2.5.

7.2.5 Troubleshooting and Debugging Backup Scripts

You can use the output collected by CPM to debug backup scripts. However, the recommended way is to run them independently of CPM, on the CPM Server machine using SSH. You can then see their outputs and fix what is needed. Once the scripts work correctly, you can start using them with CPM. 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 CPM 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 CPM Server, and then the SSH connection will not be necessary.

After that script is executed, the CPM 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

fi
```

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

- Create additional Linux users in the CPM 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 Additional Backup Topics

8.1 CPM in a VPC Environment

The CPM Server runs in a VPC, except in old environments utilizing EC2 Classic. For CPM 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 CPM 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 CPM Server.
- If a Thin Backup Agent is used in a Windows backed-up instance, the agent will need *outbound* connectivity to the CPM Server.

8.2 Backup when an Instance is Stopped

CPM 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 (see section 4.2.3), 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 clicking 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.

8.3 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.

To move a backup to the Freezer:

1. In the Backup Monitor tab of the main screen, select the backup and click **Move to Freezer**.
2. Type a unique name and an optional description. You can later search and filter frozen backups using as keywords the name or description.

After a backup is in the Freezer:

- It will only be deleted if you do so explicitly.
- 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.

8.4 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 Performing Recovery

CPM offers several options for data recovery. Since all CPM backup is based on AWS's snapshot technology, CPM can offer rapid recovery of instances, volumes, and databases. When you click **Recover** for a certain backup, you are directed to the recovery panel screen. The recovery panel screen includes:

- Links to recover the backed-up instances
- Links to recover independent volumes and databases
- 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 drop-down menu to choose in which region to perform the recovery.
- If you have cross-account functionality enabled for your CPM 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 recovery panel apply to all the recovery operations that you initiate from this screen.




Figure 9-1

Recommendation: N2WS 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.

9.1 Recovery AWS credentials

All recovery screens have a check box at the bottom labelled **Use account AWS Credentials**. By default, the recovery operation the AWS credentials that were used for backup will be used also for recovery. You can deselect it 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 14.3. When using custom credentials, CPM verifies that these credentials actually belong to the recovery account.

9.2 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 10.5.1)
- 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.

CPM 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, CPM 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 is divided to **Basic Options** and **Advanced Options**.

9.2.1 Basic Options

The basic options, shown in Figure 9-2, are:

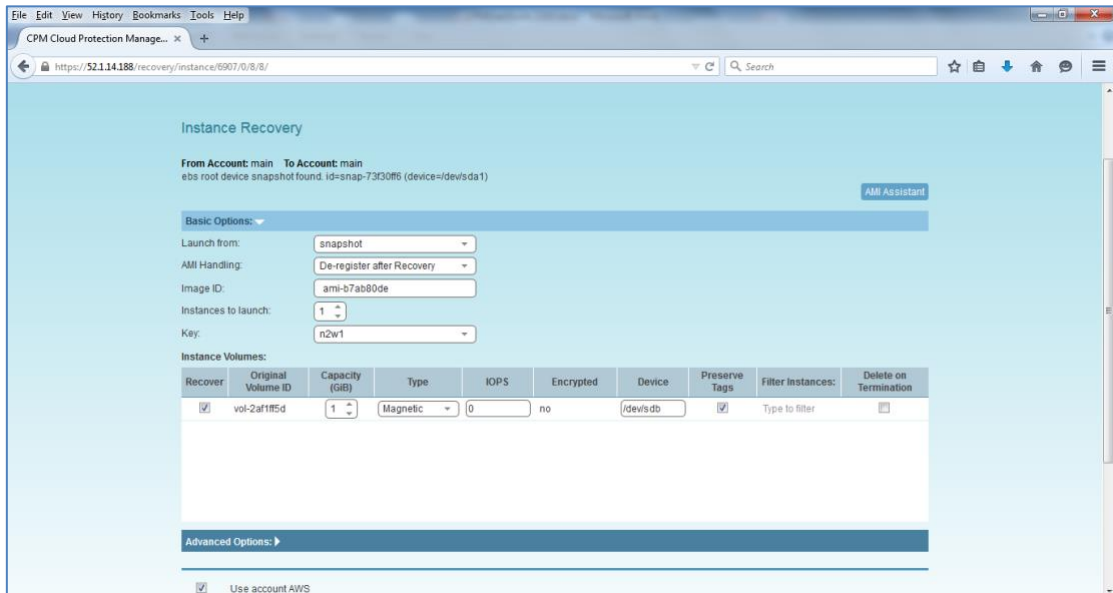
- **Launch From** – Whether to launch the boot device (image) from an existing image or a snapshot. The **snapshot** option is available only if this is an EBS-based instance, and a snapshot of the boot device is available in this backup.

- **AMI Handling** – This option is relevant only if **Launch From** is set to **snapshot**.

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 8.3).
- **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 CPM. If you wish to keep using this image, move the backup to the Freezer.
- **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 CPM 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 CPM. You can search for it with the AWS Management Console.
- **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** – The key (or 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), or to decrypt the Administrator password (Windows).
- **Instance volumes** – 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. You can make adjustments to the volumes, as follows:
 - To exclude a volume, deselect **Recover**.
 - Enlarge capacity of the volume.
 - Change the device.

- Change IOPS.
- Exclude any tags associated with the volume, such as its name
- For instances recovered from a snapshot, delete the volume on termination of the instance ().



Instance Recovery

From Account: main To Account: main
 ebs root device snapshot found. id=snap-73c086 (device=/dev/sda1)

Basic Options:

Launch from:

AMI Handling:

Image ID:

Instances to launch:

Key:

Instance Volumes:

Recover	Original Volume ID	Capacity (GB)	Type	IOPS	Encrypted	Device	Preserve Tags	Filter Instances:	Delete on Termination
<input checked="" type="checkbox"/>	vol-2af185d	<input type="text" value="1"/>	<input type="text" value="Magnetic"/>	<input type="text" value="0"/>	<input type="text" value="no"/>	<input type="text" value="/dev/sdb"/>	<input checked="" type="checkbox"/>	Type to filter	<input type="checkbox"/>

Advanced Options:

☒ Use account AWS

Figure 9-2

9.2.2 Advanced Options

Advanced options include the remaining options as shown in Figure 9-3:

- **Ephemeral Storage** – Add ephemeral drives to the new instance. The number of ephemeral storage devices you can use depends on the instance type. See <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/InstanceStorage.html>
 - Add Ephemeral storage in the format `<device name>:<virtual name>`, for example: `xvdb:ephemeral10`.
 - Add a new line for each device.
- **Architecture** – Options are:
 - **i386** – which is X86 – 32-bit
 - **x86_64** – which is X86 – 64-bit

The default will be the architecture of the backed-up instance.

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.

- **Placement** – Determines 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 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 VPC Subnet** – If you have VPC subnets defined in your account, this option is available.
 - **VPC** – You choose the VPC the instance is to be recovered to. By default, it will contain the VPC the original instance belonged to.
 - **VPC Subnet ID** – This will hold all the subnets in the currently selected VPC.
 - **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, CPM 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.
- **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.
 - **Placement Group** - Choose the placement group from the list.
- **Availability Zone** – This option is only visible if you chose **By Availability Zone** in **Placement**. 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.
- **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.
- **Security Groups** – Which security groups will 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. Every time you toggle the **Placement** option between **By Availability Zone** and **By VPC Subnet**, the list of security groups will be updated, and the previous selected items will not be saved. This field also has a filter to help you find the security group that you need.

- **Enable User Data** – Whether to use user data for this instance launch. If selected, another option appears: **User Data**.
- **User Data** – The text of the user data. Special encoding or using a file as the source is not currently supported from within CPM.

- **Preserve Tags** – By default, all the tags that were associated with the backed-up instance at the time of the backup, such as the instance's name, will also be associated with the new instance/s.
- **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.
- **Shutdown Behavior** – 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.
- **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 CPM. 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 CPM. Change this option only if you know exactly which RAM Disk you need. Choosing the wrong one will result in a failure.
- **Allow Monitoring** – Select if monitoring should be allowed for the new instance. The value in the backed-up instance is the default.
- **Instance Profile ARN** – The ARN of the instance role (IAM Role) for the instance. To find the ARN, click the Role name in IAM Management Console and click the **Summary** tab. The default will be the instance role of the backed-up instance, if it had one.
- **EBS Optimized** – Select to launch an EBS Optimized instance. The value from the backed-up instance is the default.
- **Tenancy** – Choose the tenancy option for this instance.

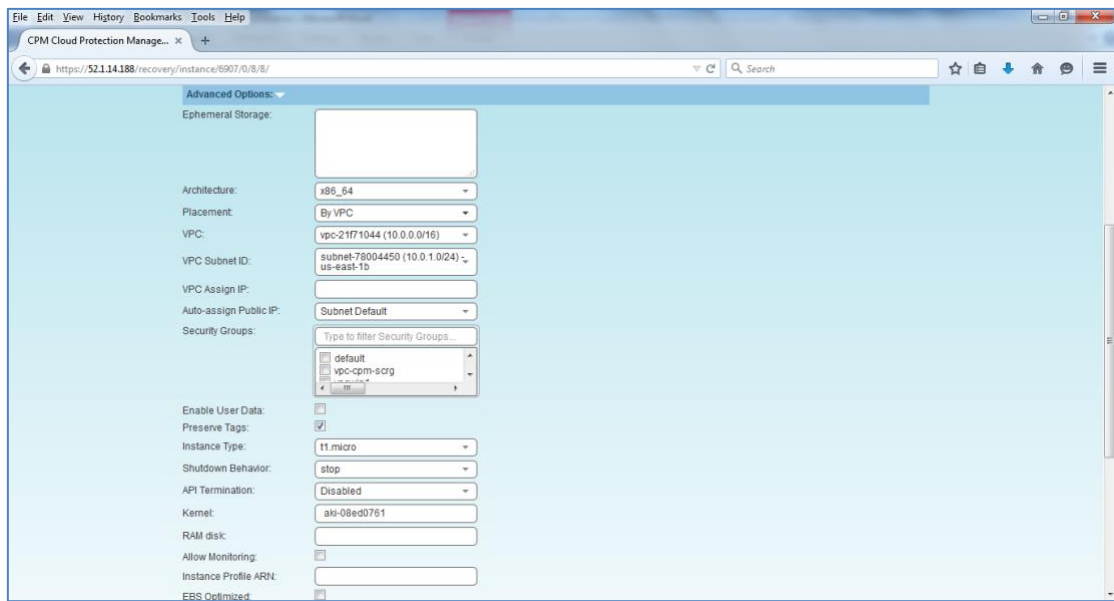


Figure 9-3

To complete the recovery operation, click **Recover Instance** and then confirm. If there are errors that CPM 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.

9.2.3 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. CPM 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, CPM will remember the details of the original AMI.

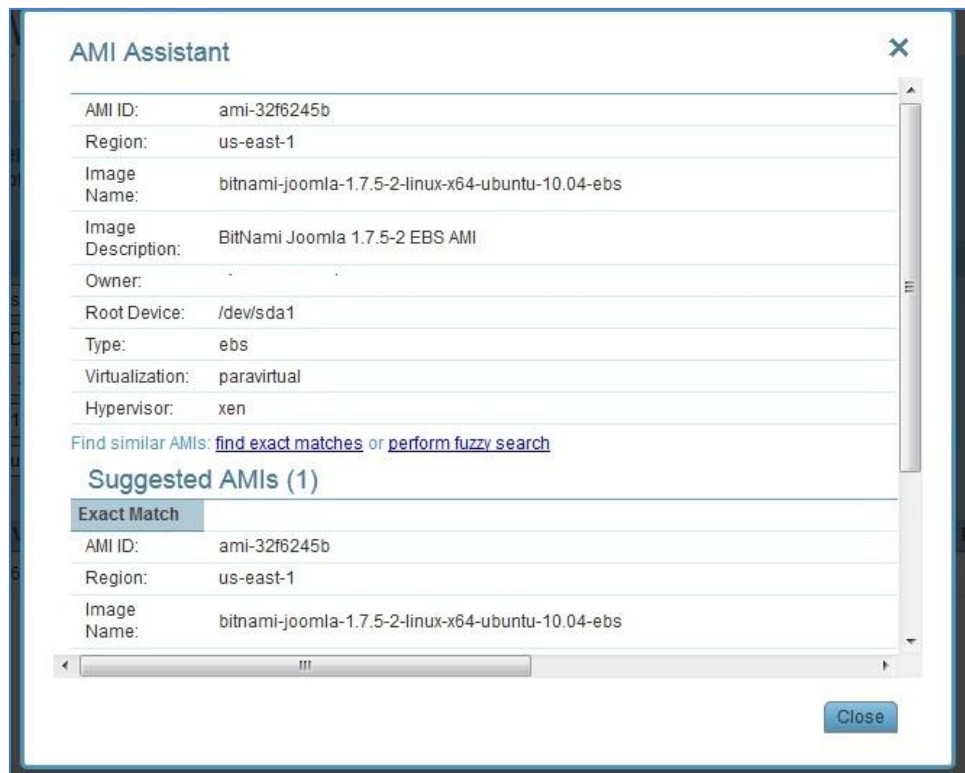


Figure 9-4

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

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

To find AMIs with properties that are exactly like the original, click **find exact matches**.

If the **find exact matches** search does not find matches, click **perform fuzzy search** 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.

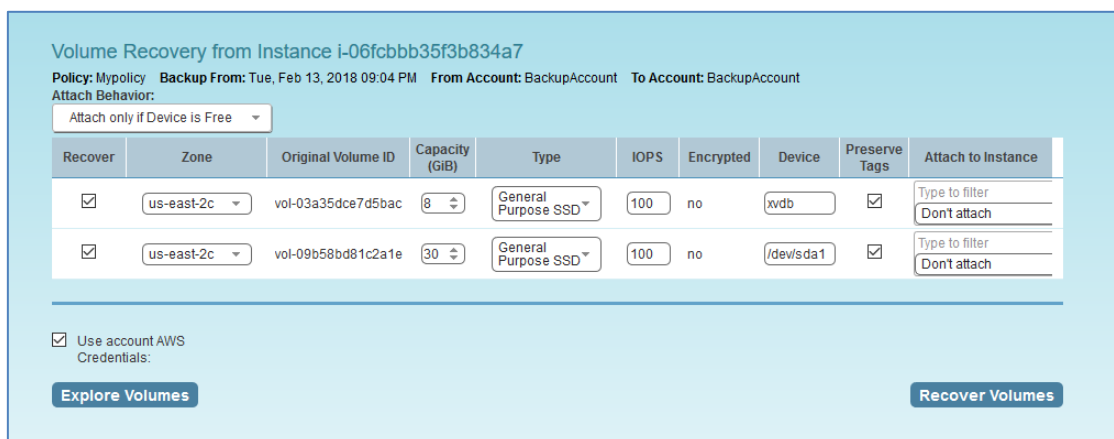
- You want to recover an instance by launching it from an image, but you want to find a newer version of the image. The fuzzy search will help you.
- You are using DR (see chapter 10) and you need to recover the instance in a different region. You may want to find the matching AMI in the target region 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.

9.3 Volume Recovery

Volume recovery means creating EBS volumes out of snapshots. In CPM, 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. Go to the Recovery Panel screen.
2. Next to an instance backup, click **Volumes Only**. The screen in Figure 9-5 opens.



Volume Recovery from Instance i-06fcbbb35f3b834a7

Policy: Mypolicy Backup From: Tue, Feb 13, 2018 09:04 PM From Account: BackupAccount To Account: BackupAccount

Attach Behavior: ☐ Attach only if Device is Free

Recover	Zone	Original Volume ID	Capacity (GiB)	Type	IOPS	Encrypted	Device	Preserve Tags	Attach to Instance
<input checked="" type="checkbox"/>	us-east-2c	vol-03a35dce7d5bac	8	General Purpose SSD	100	no	/dev/sda1	<input checked="" type="checkbox"/>	Type to filter Don't attach
<input checked="" type="checkbox"/>	us-east-2c	vol-09b58bd81c2a1e	30	General Purpose SSD	100	no	/dev/sda1	<input checked="" type="checkbox"/>	Type to filter Don't attach

☒ Use account AWS Credentials:

[Explore Volumes](#) [Recover Volumes](#)

Figure 9-5

3. Change the fields as needed:
 - **Recover** – Selected by default. Deselect if you do not want that volume recovered.
 - **Zone** – AZ. The default is the original zone of the backed-up 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.

- **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 **Zone** will refresh the content of this list.
- **Attach Behavior** – 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. CPM 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. CPM will delete them after detaching them from the target instance.

As with other recovery screens, you can choose to use different AWS credentials for the recovery operation. After clicking **Recover Volumes** and confirming, if there was a logical error in a field that CPM detected, you will be returned to the screen with an error notification. If not, you will be redirected back to the recovery panel screen with a message regarding the status of the operation.

To recover independent volumes:

Click the **Recover Independent Volumes** button at the top right of the recovery panel screen. Then a similar recover volumes screen with instance volumes opens.

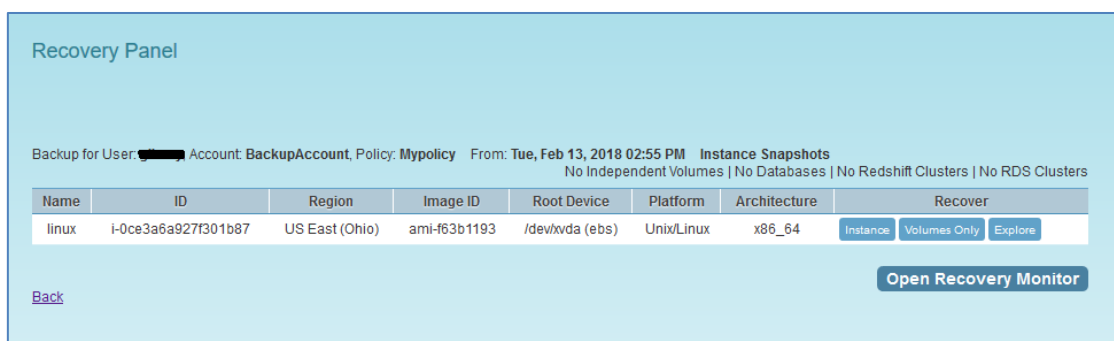


Figure 9-6

9.4 RDS Database Recovery

When a backup includes snapshots of RDS databases, the button **Recover Databases** appears on the top right corner of the recovery panel screen.



Figure 9-7

Click the **Recover Databases** button to bring you to the RDS Database Recovery screen, as shown in Figure 9-8.

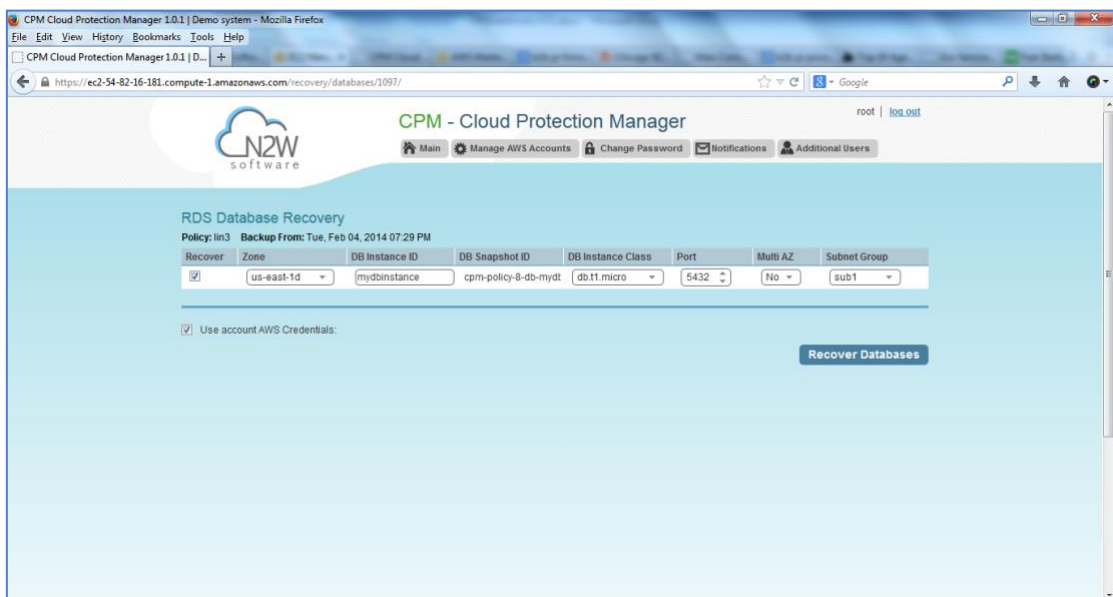


Figure 9-8

In this screen you will see a list of all RDS databases in the current backup. You can change the following options:

- **Recover** – Deselect the check box 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 CPM. 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 Snapshot ID** – Displays the snapshot ID.
- **DB Instance Class** – The default is the original class, but you can choose another.
- **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.

As in other types of recovery, you can choose to use different AWS credentials.

9.5 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, CPM will recover the DB cluster and then will create the DB instances for it.



Figure 9-9

In the Recovery Panel, click the **Recover Aurora Clusters** button to reach the **Aurora Clusters Recovery** screen:



Figure 9-10

In this screen all Aurora clusters that were backed up are listed. You can change the following options:

- **Recover** – Deselect 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.
- **DB Instance ID** – The default will be 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. CPM 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.

- **DB 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.

9.6 Redshift Cluster Recovery

When a backup includes snapshots of Redshift clusters, the button **Recover Redshift Clusters** appears on the top right corner of the recovery panel screen.



Figure 9-11

In the Recovery Panel, click the Recover Redshift Clusters button to open the Redshift Cluster Recovery screen, as shown in Figure 9-10.



Redshift Clusters Recovery
Policy: pol_test Backup From: Tue, Feb 24, 2015 01:54 PM

Recover	Zone	Cluster ID	Cluster Snapshot ID	Node Type	Nodes	Port	Subnet Group
<input checked="" type="checkbox"/>	us-east-1b	redredred	ram-policy-5-cluster-	dw2.large	1	5439	default

☒ Use account AWS Credentials

Recover Clusters

Figure 9-12

All Redshift clusters in the current backup are listed. You can change the following options:

- **Recover** – Deselect 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 CPM. You can always recover from the snapshot using AWS Management Console.

- **Cluster ID** – The default will be 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** – 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.

As in other types of recovery, you can choose to use different AWS credentials.

10 Disaster Recovery (DR)

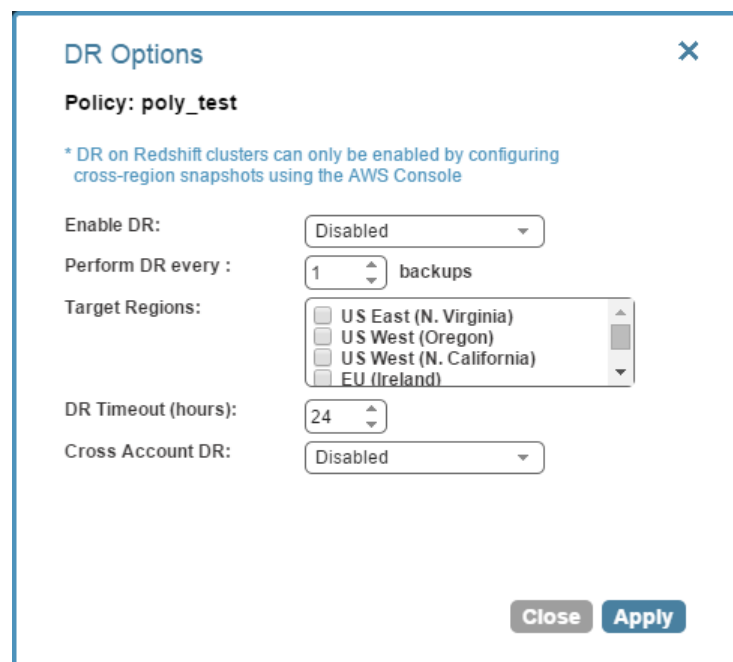
CPM's DR (Disaster Recovery) solution allows you to recover your data and servers in case of a disaster. DR will help you recover your data for whatever reason your system was taken out of service.

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 CPM uses EBS snapshots you will be able to recover your EC2 servers to other AZs. CPM's 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 CPM to take RDS snapshots, those snapshots will also be copied and will be available in other regions.

- **RDS Aurora Clusters:** DR across regions of RDS Aurora Clusters is not supported currently. N2WS plans to add DR support in future releases.
- **Redshift Clusters:** Currently CPM 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.

10.1 Configuring DR

After defining a policy, click the **DR** button under the **Configure** column in the **Policies** tab of the main screen.



DR Options

Policy: **poly_test**

* DR on Redshift clusters can only be enabled by configuring cross-region snapshots using the AWS Console

Enable DR: Disabled

Perform DR every : 1 backups

Target Regions:

- ☐ US East (N. Virginia)
- ☐ US West (Oregon)
- ☐ US West (N. California)
- ☐ EU (Ireland)

DR Timeout (hours): 24

Cross Account DR: Disabled

Close Apply

Figure 10-1

In the DR Options screen, configure the following:

- **Enable DR** – Whether DR is enabled for this policy. By default, DR is disabled.
- **Perform DR Every** – Frequency of performing DR in terms of backups. The default is to copy snapshots of all backups to other regions. To reduce costs, you may want to reduce the frequency. See section 10.4 below for considerations in planning DR.
- **Target Regions** – Which region or regions you want to copy the snapshots of the policy to.
- **DR Timeout (hours)** – How long CPM 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. CPM 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, CPM 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 volume, however, you may want to increase or decrease the time.

10.2 About the DR Process

Thing to know about the DR process:

- CPM's DR process runs in the background.
- It starts when the backup process is finished. CPM determines then if DR should run and kicks off the process.
- CPM 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 CPM 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 10.4.2.
- CPM 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.

10.3 DR and mixed-region policies

CPM 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 to that region, which are not already in this 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.

10.4 Planning your DR Solution

10.4.1 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 1 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.

10.4.2 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 CPM:

- In the definition of your policy, select the frequency in **Perform DR every....** 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.

10.4.3 Performing DR on the CPM Server (The `cpmdata` Policy)

To perform DR recovery, you will need your CPM 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 CPM server itself is down. You may want to copy your CPM database across regions as well. Generally, it is not a bad idea to place your CPM server in a different region than your other production data. CPM 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 CPM server happens to be in that zone, it will not be available.

To make it easy and safe to back up the CPM server database, there is a special policy named `cpmdata`. Although CPM supports managing multiple AWS accounts, the only account that can back up the CPM 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 in a consistent manner.

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.

10.5 DR Recovery

DR recovery is similar to regular recovery with a few differences, as shown in Figure 10-2:

- When you click the **Recover** button 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.

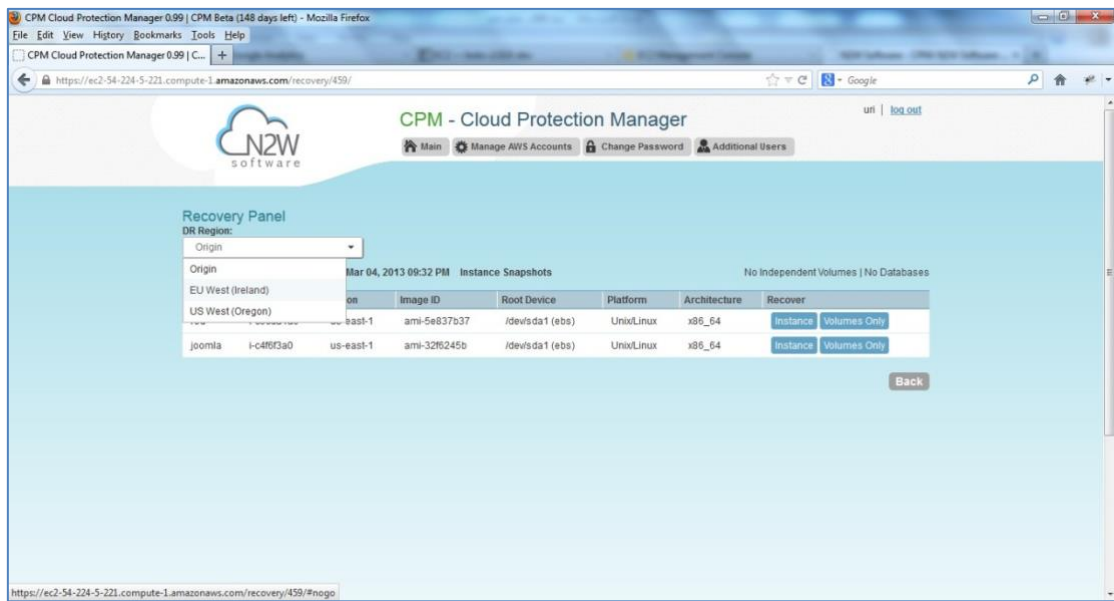


Figure 10-2

- 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: N2WS 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.

10.5.1 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 9.2.3).

- **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, CPM will remember the security groups of the original instance and use them as default. In DR recovery, CPM 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, CPM 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 9.2.3). 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.

10.5.2 DR of Encrypted Volumes, AMIs and RDS Instances

CPM supports DR of encrypted EBS volumes. If you are using KMS keys for encryption, CPM will seek a KMS key in the target region, which has the same alias.

To configure your cross-region DR:

1. Create a matching-alias key in the source and in the remote region for CPM to use automatically in the DR copy process.
2. If a matching key is not found in the target region, the DR process will fail.
3. If the key uses the default encryption, then it will be copied to the other region with the default encryption key as well.

CPM supports copy of AMIs with encrypted volumes with the same logic it uses for volumes.

CPM supports cross-region DR of encrypted RDS databases.

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

10.5.3 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 CPM Server, as it is down with all the rest of your instances. Here is Disaster Recovery step by step:

1. With AWS Management Console:
 - a. Find the latest snapshot of your `cpmdata` policy by filtering snapshots with the string `cpmdata`. CPM 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-clicking 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 CPM 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 CPM server:
 - a. Connect to the newly created instance using HTTPS.
 - b. Approve the SSL certificate exception.
Assuming the original instance still exists, CPM 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 [CPM Basic 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 CPM server, you can perform any recovery you need at the target (current) region.
 - a. Select the backup you want to recover.
 - b. Click **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.

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

10.6 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 backup log where every stage and operation during DR is recorded:

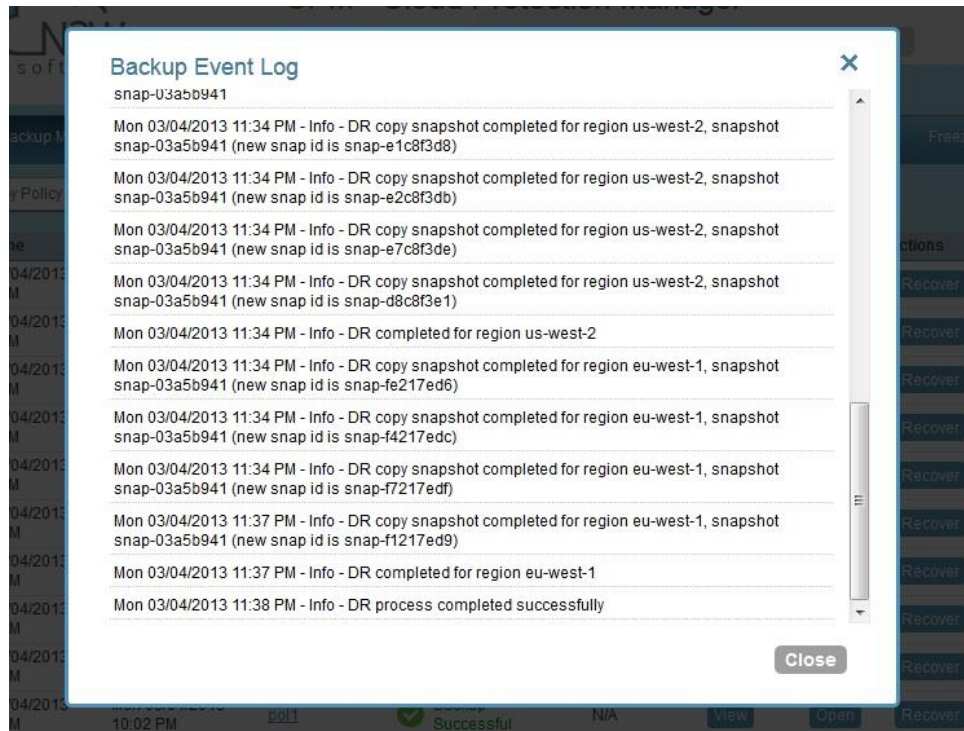


Figure 10-3

You can also view DR snapshot IDs and statuses in the snapshots screen of the backup:

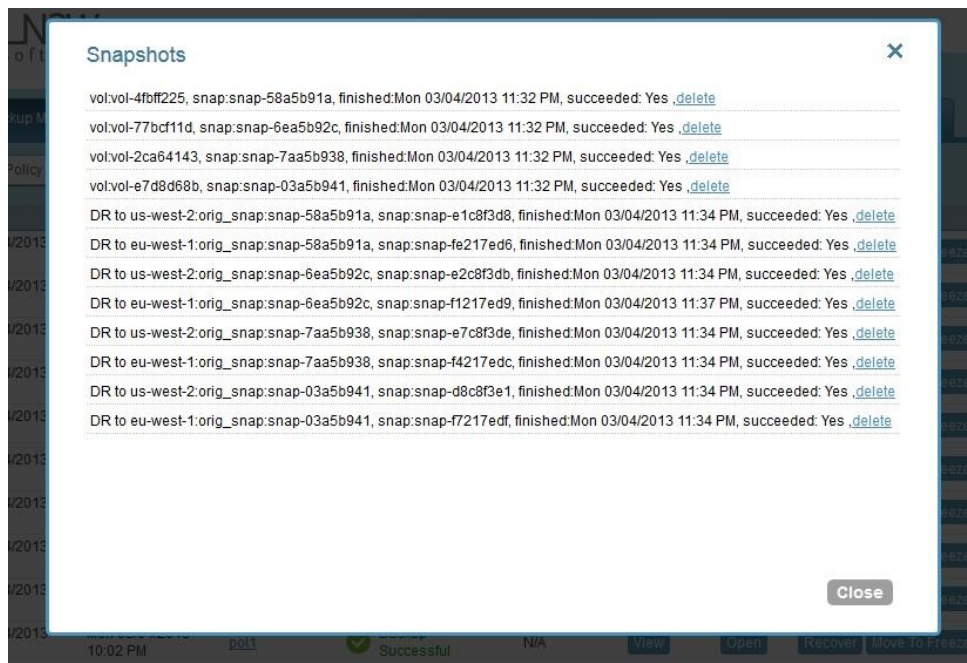


Figure 10-4

Every DR snapshot is displayed with region information and the IDs of both the original and the copied snapshots.

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, CPM 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.

11 Cross-Account DR, Backup and Recovery

Available only in Advanced and Enterprise Editions, CPM's 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. CPM 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 (excluding Aurora).

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

11.1 Snapshot Vaulting

CPM 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 CPM can have limited permissions that do not allow snapshot deletion.

CPM will tag outdated snapshots instead of actually deleting them, allowing an authorized user to delete them separately using the EC2 console or a script. 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.

11.2 Configuring Cross-Account Backup

You configure cross-account DR from the DR screen of the policy:

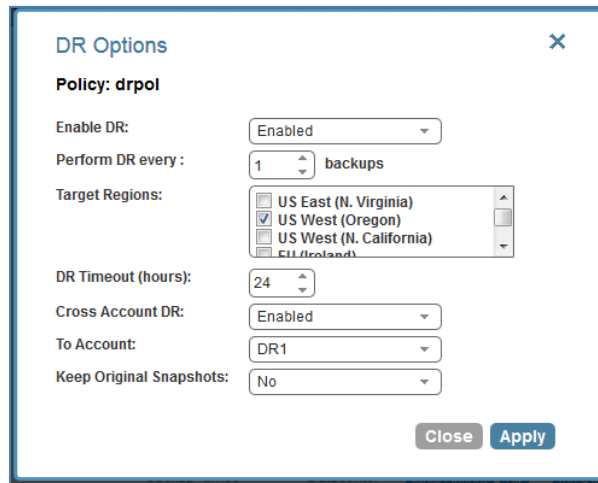


Figure 11-1

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

Once you set the **Cross Account DR** field to **Enabled**, the other fields will become visible:

- **To Account** – Which account to copy the snapshots to. This account needs to be defined as a **DR Account** in the **Manage AWS Accounts** screen.
- **Keep Original Snapshots** – Whether to keep the snapshots both in the original and the DR accounts or to delete the original snapshots once they are copied to the DR account. This can save cost by not paying to store all snapshots twice.

11.3 Cross-Account DR and Clean-Up

CPM 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, CPM will clean up snapshots copied to other accounts as well. However, if you do not wish for CPM 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, CPM 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 CPM.

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.

11.4 Cross-Account with Cross-Region

If you configure the backup policy to copy snapshots across accounts as well as across regions, CPM will combine: it will copy to the other account and to other regions. So, you can potentially copy snapshots

to regions and accounts. It is important to know exactly what you are doing and not let the cost of these actions to be too high.

11.5 Cross-Account Recovery

If you have cross-account functionality enabled in your CPM license, and even if you actually configured CPM to copy snapshots between accounts, you can recover across accounts. This is already mentioned in the recovery chapter (see chapter 9). You need to choose which account to recover the resource (EC2 instance, EBS volume or RDS database) to. 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. CPM will utilize the AWS **share snapshot** option to enable recovering resources across accounts.

11.6 Cross-Account Backup and Cost

Note: Currently, copying snapshots between accounts is not incremental by nature. Unlike creating regular snapshots or copying snapshots between regions, copying snapshots between accounts will copy the entire volume every time. This can have a considerable effect on cost. Be sure to configure the backup policies according to business needs.

12 File-level Recovery

In version 2.0.0, CPM introduced file-level recovery. CPM 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 CPM, 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.

CPM now automates this entire process. Click **Explore** (see Figure 12-1) either from the recovery panel screen for an instance, or from the volume recovery screen for a specific volume. CPM 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.

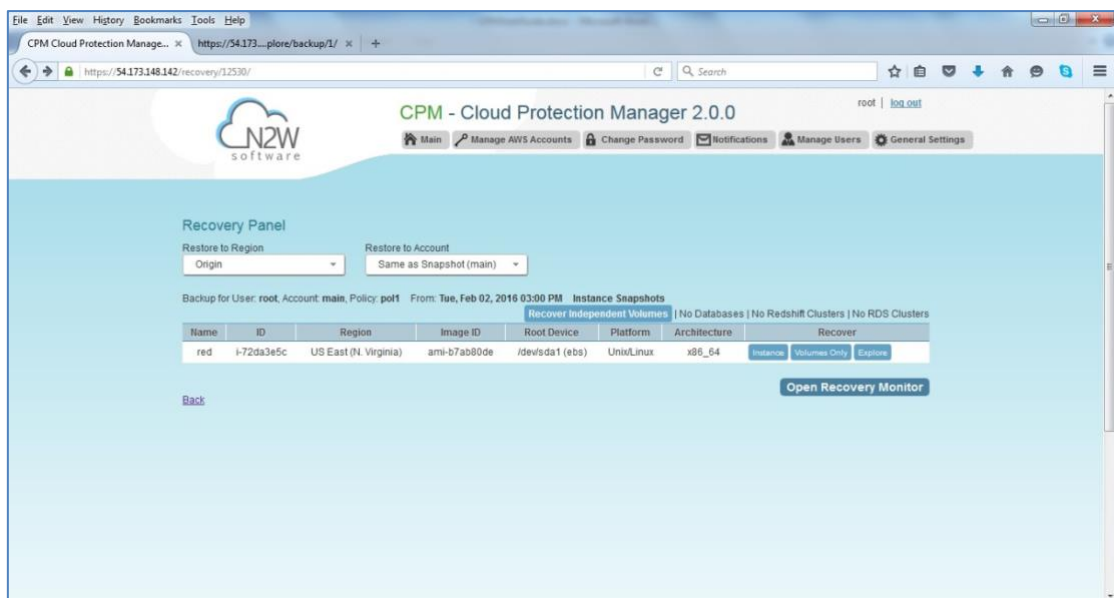


Figure 12-1

On the right side of any file or folder, there is a green download icon that will download the file or folder. Folders are downloaded as uncompressed zip files.

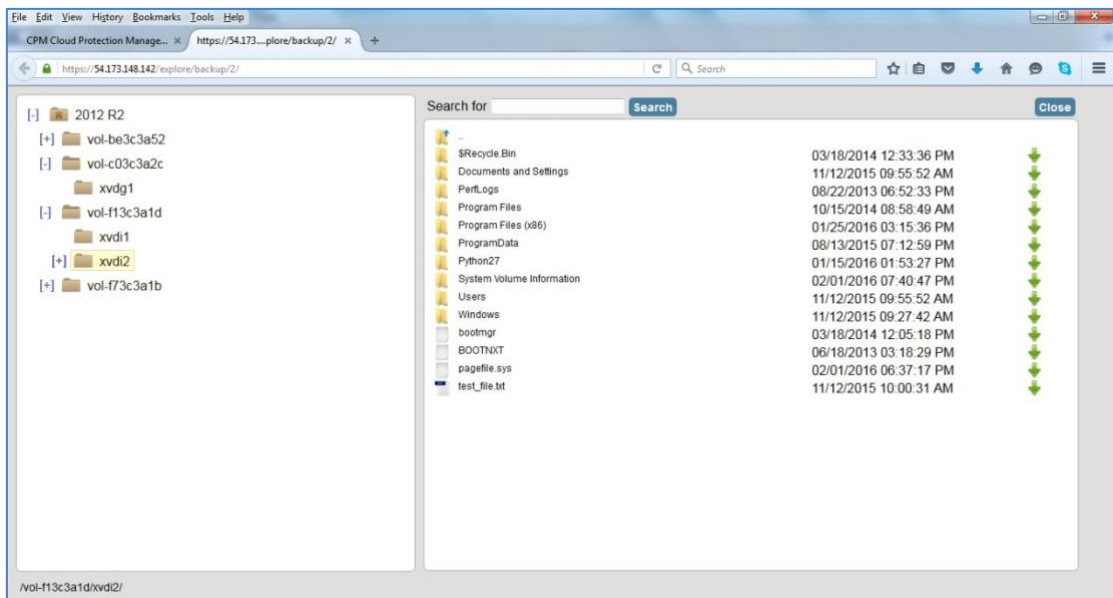


Figure 12-2

To perform these operations, CPM needs to be able to use AWS credentials belonging to the CPM server instance account, with sufficient permissions to create and attach volumes. By default, CPM 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 CPM to recover volumes in the background and attach them to the CPM server. There are a few limitations:

- Explore works only on simple volumes. LVM and Windows dynamic disks are currently not supported. Additionally, disks defined with Microsoft Storage Spaces are not supported.
- Explore works only on snapshots taken in the same region the CPM server is in.
- Explore for encrypted volumes will only work at the same account the CPM Server instance is in. Cross-account explore of encrypted volumes is not supported.

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

13 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.

CPM allows configuring backup using automation tools by utilizing AWS tags. By tagging a resource (EC2 instance, EBS volume, RDS instance, Aurora Cluster or Redshift cluster), CPM can be notified what to do with this resource, and there is no need to use the GUI. Since tagging is a basic functionality of AWS, it can be easily done via the API and scripts.

13.1 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). CPM will identify this tag and parse its content. In this tag you will be able to specify whether to:

- Remove this resource 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.

13.1.1 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

13.1.2 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 CPM 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 done 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.

13.1.3 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, CPM 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.
- Additional configurations need to be done manually, and not with the tag.

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

13.1.4 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 CPM to remove this resource from all policies.

13.2 Tag Scanning

Tag scanning can only be controlled by the admin/root user. When scanning 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. CPM will automatically scan resources in all AWS regions.

In the **General Settings** screen you can enable or disable tag scanning, and you can set the interval in hours for automatic scans. You also have the **Scan Now** button to initiate a tag scan immediately.

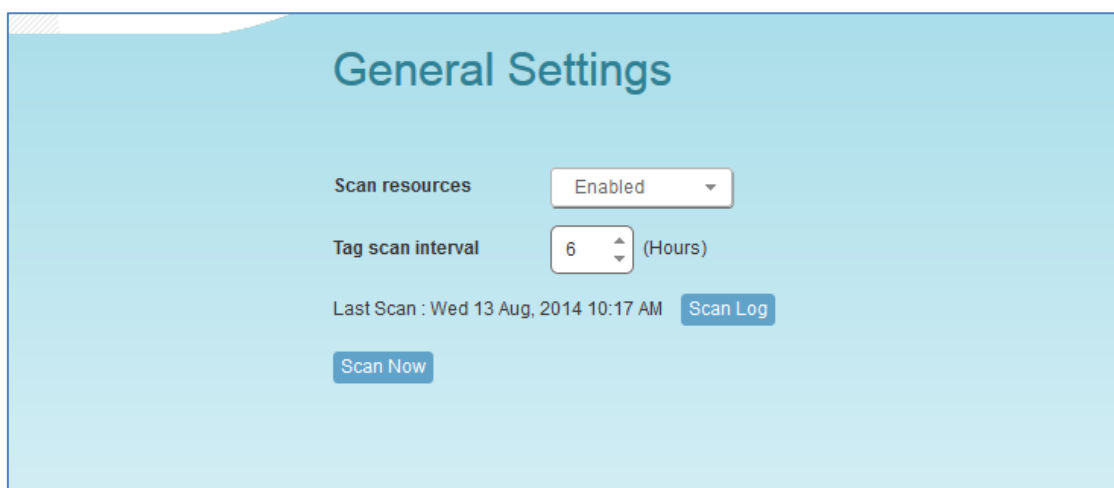


Figure 13-1

Note: Even if scanning is disabled, clicking **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 set **Scan Resources** for the relevant AWS accounts.

13.3 Pitfalls and Troubleshooting

13.3.1 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. CPM 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, CPM 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 (see section 15.5), any policies with no backup targets will be listed.

- If the same AWS account is added as multiple accounts in CPM, the same tags can be scanned multiple times, and the behavior can become unpredictable. N2WS generally discourages this practice. It is better to define an account once, and then allow delegates (see section 16.4) access to it. If you added the same AWS account multiple times (even for different users), make sure only one of the accounts in CPM has **Scan Resources** enabled.

13.3.2 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 (e.g. 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 CPM – defined manually or scanned previously.

14 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.

CPM strives to be as secure as the cloud it is in. It has many features that provide you with a secure solution.

14.1 CPM Server

CPM Server’s security features are:

- Since you are the one who launches the CPM 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 CPM 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 CPM server using HTTPS or SSH, both secure protocols, which means that all communication to and from CPM is encrypted. Also, when connecting to AWS endpoints, CPM will verify that the SSL server-side certificates are valid.
- Every CPM 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 CPM’s database.
- CPM supports using different AWS credentials for backup and recovery.
- CPM Server supports IAM Roles. If the CPM 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 CPM instance account to all of the other AWS accounts you manage and not type AWS credentials at all.
- To manage CPM, you need to authenticate using a username and password.
- CPM allows creating multiple users to separately manage the backup of different AWS accounts, except in the Basic Edition.

14.2 Best Security Practices for CPM

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 CPM a bit cumbersome, so it is your decision whether to implement them or not.

14.2.1 Avoid using AWS Credentials

By using the CPM 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.

14.2.2 Credentials Rotation

Assuming you have to use AWS credentials, you should follow AWS practices. It is recommended to 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 CPM. Click on the account name in **Manage AWS Accounts**, and modify the access and secret keys.

14.2.3 Passwords

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

14.2.4 Security Groups

Since CPM server is an instance in your account, you can define and configure its security groups. Even though CPM 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 CPM Thin Backup Agent installed on them
- You will also need to allow SSH access to create and maintain backup scripts.
- Blocking anyone else will make CPM 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.

14.3 Using IAM

CPM keeps your AWS credentials safe. However, it is preferable to use IAM roles and not use credentials at all. Additionally, CPM 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 CPM 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.

An IAM role can also be used in the CPM Server (for the account the CPM Server was launched in) and for instances running CPM 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/>

14.3.1 CPM Server Configuration Process

AWS credentials in the CPM configuration process are only used for configuring the new server. However, if you want to use IAM credentials for the CPM configuration process, or to use the IAM role associated with the CPM Server instance, its IAM policy should enable CPM 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 CPM Server instance, you will need the following policy and the policies for CPM Server's normal operations, as described in the next section.

Minimal IAM Policy for CPM 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"
}
]
}

```

14.3.2 CPM Server IAM Settings

You can use the CPM 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 CPM 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.

Minimal IAM Policy for Backup Only

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "ec2:CopySnapshot",
        "ec2:CopyImage",
        "ec2>CreateImage",
        "ec2>CreateSnapshot",
        "ec2>CreateTags",
        "ec2>DeleteSnapshot",
        "ec2:DescribeAddresses",
        "ec2:DescribeAvailabilityZones",
        "ec2:DescribeImageAttribute",
        "ec2:DescribeImages",
        "ec2:DescribeInstanceAttribute",
        "ec2:DescribeInstanceStatus",
        "ec2:DescribeInstances",
        "ec2:DescribeKeyPairs",
        "ec2:DescribePlacementGroups",

```

```

    "ec2:DescribeRegions",
    "ec2:DescribeSecurityGroups",
    "ec2:DescribeSnapshotAttribute",
    "ec2:DescribeSnapshots",
    "ec2:DescribeSubnets",
    "ec2:DescribeTags",
    "ec2:DescribeVolumeAttribute",
    "ec2:DescribeVolumeStatus",
    "ec2:DescribeVolumes",
    "ec2:DescribeVpcs",
    "ec2:ModifyImageAttribute",
    "ec2:ModifySnapshotAttribute",
    "ec2:ResetSnapshotAttribute",
    "kms:ListKeys",
    "kms:ListAliases"
  ],
  "Sid": "Stmt1374236955000",
  "Resource": [
    "*"
  ],
  "Effect": "Allow"
},
{
  "Action": [
    "rds:CreateDBSnapshot",
    "rds>DeleteDBSnapshot",
    "rds:DescribeDBInstances",
    "rds:DescribeDBClusters",
    "rds:DescribeDBParameterGroups",
    "rds:DescribeDBParameters",
    "rds:DescribeDBSnapshots",
    "rds:DescribeDBSubnetGroups",
    "rds:ListTagsForResource",
    "rds:CopyDBSnapshot",
    "redshift:DescribeClusters"
  ],
  "Sid": "Stmt1374237153000",
  "Resource": [
    "*"
  ],
  "Effect": "Allow"
}
]
}

```

Minimal IAM Policy for Recovery

This policy should be attached to the IAM user or role in addition to the one in the previous section:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "ec2:AllocateAddress",
        "ec2:AssociateAddress",
        "ec2:AttachVolume",
        "ec2:CreateImage",

```

```

    "ec2:CreateVolume",
    "ec2:DeleteVolume",
    "ec2:DeregisterImage",
    "ec2:DescribeImages",
    "ec2:DetachVolume",
    "ec2:ModifyImageAttribute",
    "ec2:ModifyInstanceAttribute",
    "ec2:ModifySnapshotAttribute",
    "ec2:ModifyVolumeAttribute",
    "ec2:RegisterImage",
    "ec2:RunInstances",
    "ec2:StartInstances",
    "ec2:StopInstances",
    "iam:PassRole"
  ],
  "Sid": "Stmt1374243096000",
  "Resource": [
    "*"
  ],
  "Effect": "Allow"
},
{
  "Action": [
    "rds:RestoreDBInstanceFromDBSnapshot",
    "rds:RestoreDBClusterFromSnapshot"
  ],
  "Sid": "Stmt1374243250000",
  "Resource": [
    "*"
  ],
  "Effect": "Allow"
}
]
}

```

IAM Policy for SNS Notifications

If you want to use CPM's alerts and notifications using an IAM user or role, you will also need an IAM policy to allow interacting with SNS:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "sns:CreateTopic",
        "sns:GetTopicAttributes",
        "sns:Publish",
        "sns:SetTopicAttributes",
        "sns:Subscribe"
      ],
      "Sid": "Stmt1374246783000",
      "Resource": [
        "*"
      ],
      "Effect": "Allow"
    }
  ]
}

```

```
}

```

If you want to be even stricter, you can define a policy that can only publish, and only to the relevant topics.

Redshift

To add the ability to manage Redshift Cluster snapshots, either create a new policy or add the following permissions to your backup policy:

```
{
  "Sid": "Stmt1425805298000",
  "Effect": "Allow",
  "Action": [
    "redshift:CopyClusterSnapshot",
    "redshift:CreateClusterSnapshot",
    "redshift:CreateTags",
    "redshift>DeleteClusterSnapshot",
    "redshift:DescribeClusterParameterGroups",
    "redshift:DescribeClusterParameters",
    "redshift:DescribeClusterSnapshots",
    "redshift:DescribeClusterSubnetGroups",
    "redshift:DescribeClusters",
    "redshift:DescribeTags",
    "redshift:RestoreFromClusterSnapshot"
  ],
  "Resource": [
    "*"
  ]
}
```

The last permission is used to recover Redshift clusters from snapshots. You can add this specific permission to your recovery IAM policy instead.

Cross-Account Backup & Recovery

To enable CPM to back-up and recover across accounts, the accounts that snapshots are copied from or recovered from must have the following IAM permissions:

```
"ec2:ModifyImageAttribute",
"ec2:ModifySnapshotAttribute",
```

14.3.3 CPM Agent IAM Role

If you are using CPM agents in your environment and do not wish the CPM Server to actually send credentials to them, you can associate the Windows instance the CPM agent is on with an IAM role at launch time. The IAM role needs less permissions than CPM Server:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "ec2:CreateSnapshot",
        "ec2:CreateTags",
        "ec2:DescribeAvailabilityZones",
        "ec2:DescribeInstanceAttribute",
        "ec2:DescribeInstanceStatus",
```

```
"ec2:DescribeInstances",
"ec2:DescribeRegions",
"ec2:DescribeSnapshotAttribute",
"ec2:DescribeSnapshots",
"ec2:DescribeVolumeAttribute",
"ec2:DescribeVolumeStatus",
"ec2:DescribeVolumes",
"ec2:ModifySnapshotAttribute"
],
"Sid": "Stmt1374250341000",
"Resource": [
  "*"
],
"Effect": "Allow"
},
{
  "Action": [
    "rds:CreateDBSnapshot",
    "rds:DescribeDBInstances",
    "rds:DescribeDBSnapshots"
  ],
  "Sid": "Stmt1374250440000",
  "Resource": [
    "*"
  ],
  "Effect": "Allow"
}
]
```

Note: If you do not use CPM with RDS at all, you can omit all RDS permissions from your IAM policies.

14.4 Thin Backup Agent

The CPM Thin Backup Agent is used for Windows instances that need to perform application quiescence using VSS or backup scripts. The agent communicates with the CPM Server using the HTTPS protocol.

No sensitive information passes between the backup agent and the CPM Server.

15 Alerts, Notifications and Reporting

CPM 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, CPM allows sending alerts, notifications and even raw reporting data. 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, CPM has an answer for that.

15.1 Alerts

Alerts are notifications about issues in your CPM backup solution. Whenever a policy fails, in backup or DR, an alert is issued so you will know this policy is not functioning properly. 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).

15.2 Pull Alerts

If you wish to integrate CPM with 3rd party monitoring solutions, CPM allows API access to pull alerts out of CPM. A monitoring solution can call this API to check if CPM 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 CPM server to use an alternate port (not 443), you will need to use that port for this API call as well. CPM requires an authentication key from the caller. Every CPM 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.

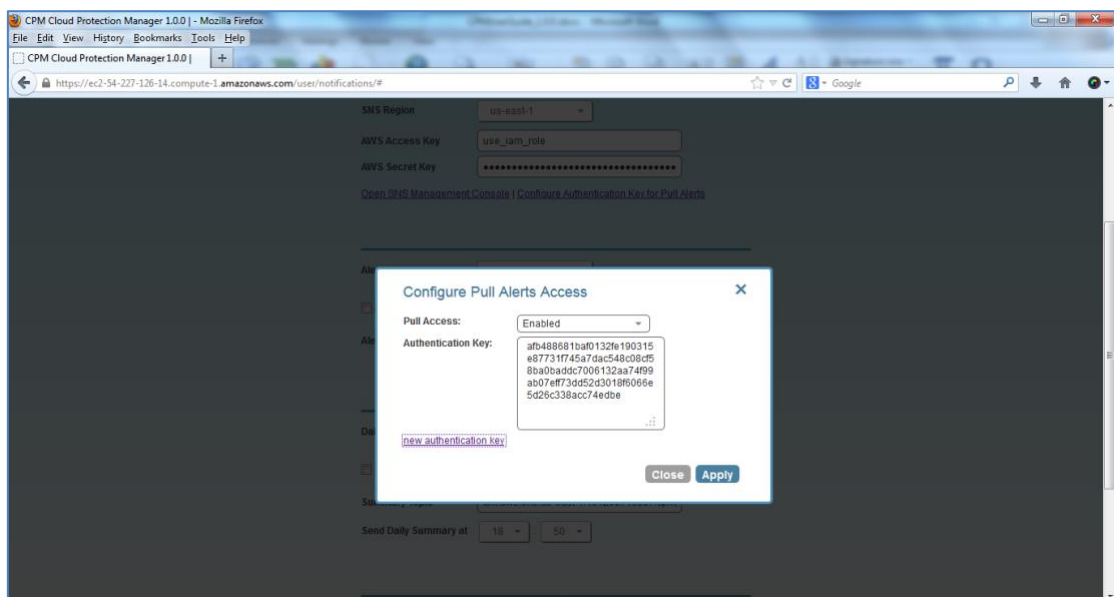


Figure 15-1

To configure an API call:

1. At the top of any screen, click the **Notifications** button.
2. In the notifications screen, click the **Configure API Authentication Key** link.
3. In the popup screen, select **Enable** or **Disable** in the **API Access** list.
4. To generate an authentication key, click **new authentication key** (see Figure 15-1).
5. To overwrite any key in the **Authentication Key** box, click **new authentication key**.
6. After enabling and setting the key, you can use the API call to get all alerts:

`https://<your CPM Server address>:<your port>/agentapi/get_cpm_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
a74f99ab07eff736477dca86b460a4b1a7bfe826e16fdbbc'
>>> 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

[{'category': u'Backup', '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',
'severity': u'E', 'title': u'Policy win_server Backup Failure',
>alert_time': u'2013-07-20 06:00:03', 'policy': {'name':
u'win_server'}}, {'category': u'Backup', '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

15.3 Using SNS

CPM 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/>.

CPM 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 CPM's notifications screen.

For the small volume of SNS messages CPM uses, there is usually no cost or it is negligible. For SNS pricing see <https://aws.amazon.com/sns/pricing/>.

15.3.1 Configuring SNS

To configure CPM for SNS, click the **Notification** button on any screen.

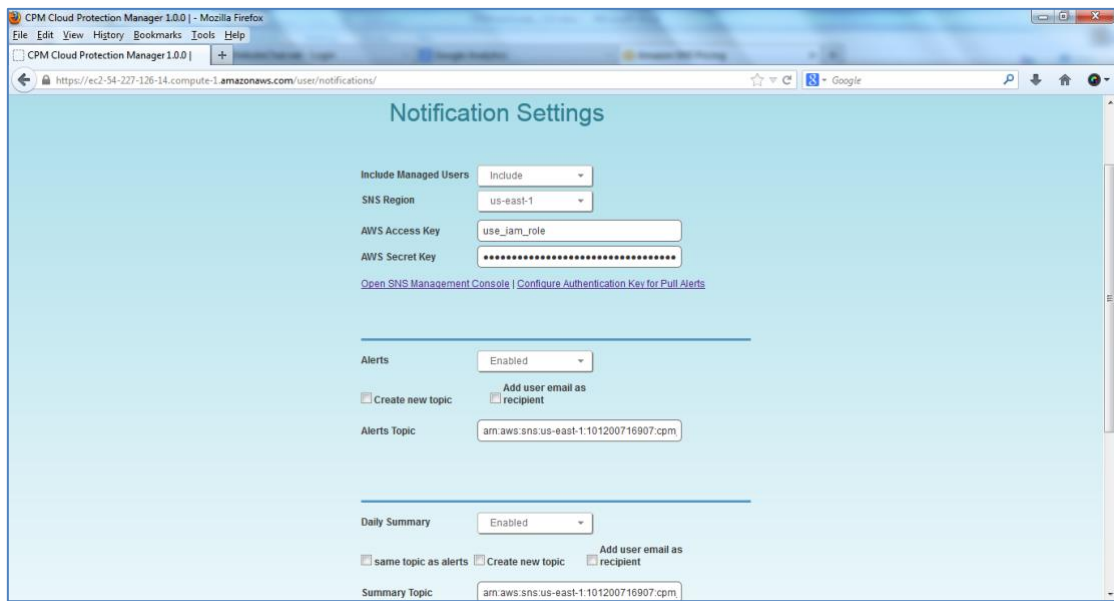


Figure 15-2

The Notifications screen appears as shown in **Error! Reference source not found..**

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 14.3). To use the CPM 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 also choose whether to include or exclude alerts about managed users (see section 16.2).
- SNS is used both for push alerts and for sending a daily summary.

15.4 Push Alerts

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

To enable push alerts:

1. Set **Alerts** to **Enabled**.
2. Either paste in the topic's ARN that you copied from the SNS tab of the AWS Management Console, or request CPM to create the topic for you and add the user's email as a recipient (optional).

Each recipient will receive a message requesting subscription confirmation before receiving alerts.

15.5 Daily Summary

Daily summary is a message that is sent once a day, summarizing all current alerts 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 CPM crashed altogether, and your monitoring solution did not report it, you will notice daily summaries 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.

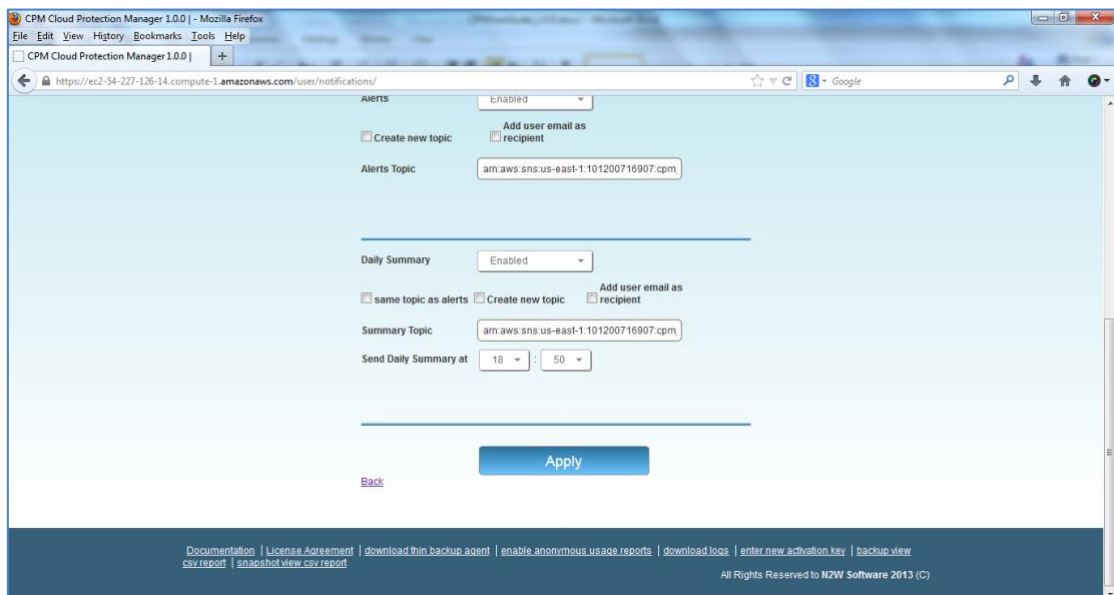


Figure 15-3

While configuring SNS, as shown in Figure 15-3, you can also configure the Daily Summary.

To configure the Daily Summary:

1. In the Notification Settings screen, select **Enabled** in the **Daily Summary** list.
2. Use one of following options for defining the Daily Summary topic:
 - If you have Alerts configured and you want to use the same SNS topic for summaries, select the **same topic as alerts** check box.
 - To create a new topic, select the **Create new topic** check box, and complete the next screen.
 - Type or paste an ARN in the **Summary Topic** box.

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 SMS, but to get the daily summary by email only. The display name of the topic also appears in the message (in emails it appears as the sender name), so with separate topics it is easier to know which is which.

3. To add a recipient, select the **Add user email as recipient** check box, and complete the next screen.
4. In the **Send Daily Summary at** lists, select the hour and minutes to send the notification.

15.6 Raw Reporting Data

In the future, N2WS plans on adding a full-scale reporting module to CPM. In the meantime, you can get two raw reports that you can download 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 CPM.

To download the reports, click the **backup view csv report** or **snapshot view csv report** link at the bottom of CPM's main screen.

15.6.1 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.
- **User** – Name of the User if the system has multiple users and the user downloading the report is root.
- **Account** – Name 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.

15.6.2 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.
- **Region** – AWS region.
- **Type** – Type of snapshot: EBS, RDS or EBS Copy, which is a DR copied snapshot.
- **Volume/DB** – AWS ID of the backed up EBS volume or RDS database.
- **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.
- **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.

15.6.3 Keeping Records after Deletion

By default, when a backup is marked for deletion, it will be deleted right away from the CPM database, and therefore not appear in the reports. There are exceptions, such as if CPM 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 save records for a period of time:

1. On the CPM server, create a file containing only the number of days to save records in
`/cpmdata/conf/num_days_to_keep_backup_records`.
2. Set the file permissions to allow all users read access.

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, you need to put 90 in the file.

3. A typical way to create it is as follows:

```
echo 90 > /cpmdata/conf/num_days_to_keep_backup_records
```

To see how to login to the CPM Server instance using SSH, see section 7.1.

Note: Keeping backups for long periods of time can cause the CPM database to grow and therefore affect the size you need to allocate for CPM's data volume. N2WS estimates that every GiB will accommodate the backup of 10 instances. N2WS 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 volume 3 GiB, for 20 make the volume 6 GiB, etc.


15.7 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 consuming. This can be helpful for inter-user accounting.

- For each user, there is a link **usage report for current user**.
- For the root user, there is also a link **usage report for all users** which will give all the breakdown of usage between all the users on the CPM server.

16 CPM User Management

CPM 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 CPM you are subscribed to). Additional users are helpful if you are a managed service provider, in need of managing multiple customers from one CPM 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 account/s.



User Name	User Type	Accounts	Policies	Num Frozen Items	Authentication	Actions
root	admin/root	2	1	0	local	usage report, audit report, delegates
okta_user_all_attr_is_set	independent	0	0	0	Identity Provider	usage report, audit report, delete, delegates
okta_user_attr_from_grp	independent	0	0	0	Identity Provider	usage report, audit report, delete, delegates
okta_user_permi_valid_1	independent	1	0	0	Identity Provider	usage report, audit report, delete, delegates
okta_usr_override_grp_attr	independent	0	0	0	Identity Provider	usage report, audit report, delete, delegates
cpm_user_1	managed	0	0	0	local	usage report, audit report, reset password, delete, delegates

Figure 16-1

There are two types of users you can define: independent users and managed users

16.1 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 does.

- Log-in to CPM
- Create their own accounts
- Manage their backup

16.2 Managed Users

Managed Users are users who can log on and manage their backup environment, or the root/admin 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 & perform recovery. Furthermore, the root user can receive alerts and notifications on behalf of managed users, although manage users can also define notifications and get them directly. To create a managed user, click the **Add New User** button in the **Manage Users** screen, and fill in the type as **Managed**. If the root user does not want managed users to login at all, they should not receive any credentials.

16.3 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 CPM via IdP integration (see chapter 17) cannot be edited, only deleted.

To define users:

1. If you are the root or admin user, at the top of any CPM screen, click the **Manage Users** button. The **Manage Users** screen opens.
2. Click the **Add New User** button.
3. In the **User name**, **Email** and **Password** boxes, type the relevant information.
4. In the **User Type** list, select the user type. For type details, see sections 16.1 and 16.2.

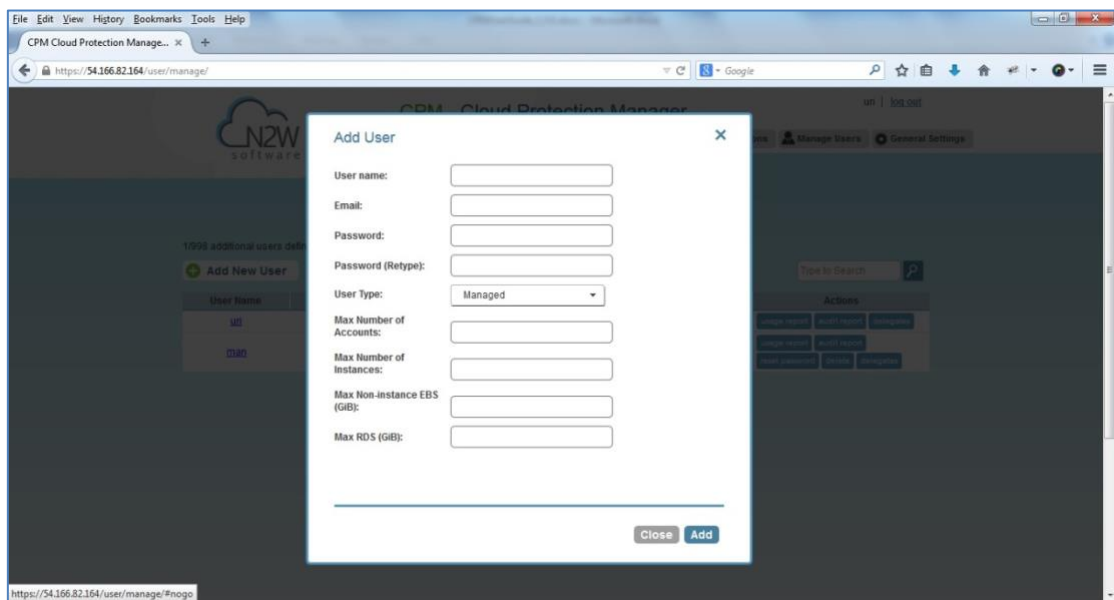


Figure 16-2

5. In the **Max Number of Accounts**, **Max Number of Instances**, **Max Non-instance EBS (GiB)**, and **Max RDS (GiB)** boxes, type the value for the respective resource limitation.

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 CPM edition used.

16.4 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.

For each user, whether it is the root user, an independent user or a managed user, there is a button **delegates** that redirects to the delegates screen for that user:

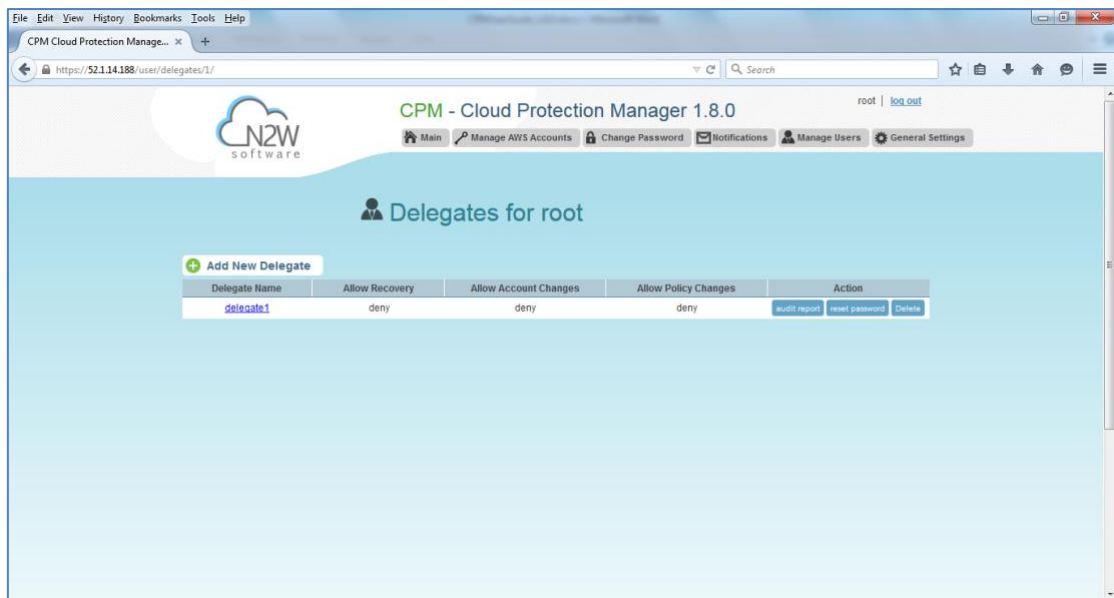


Figure 16-3

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

To add a new delegate:

1. Select a user.

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

2. Click the **Add New Delegate** button.
3. In the **User name** list, select the new delegate.

The user is added as a delegate with the following permissions set to **deny**:

- **Allow Recovery** – Perform recovery operations
 - **Allow Account Changes** – Add and remove AWS accounts, edit accounts, modify credentials
 - **Allow Backup Changes** – Change policies and their schedule and add and remove backup targets
4. Edit the delegate to set the above permissions to **allow**.

The default **allow** permissions are:

- Viewing the settings.
- Viewing the environment.
- Monitoring backups.

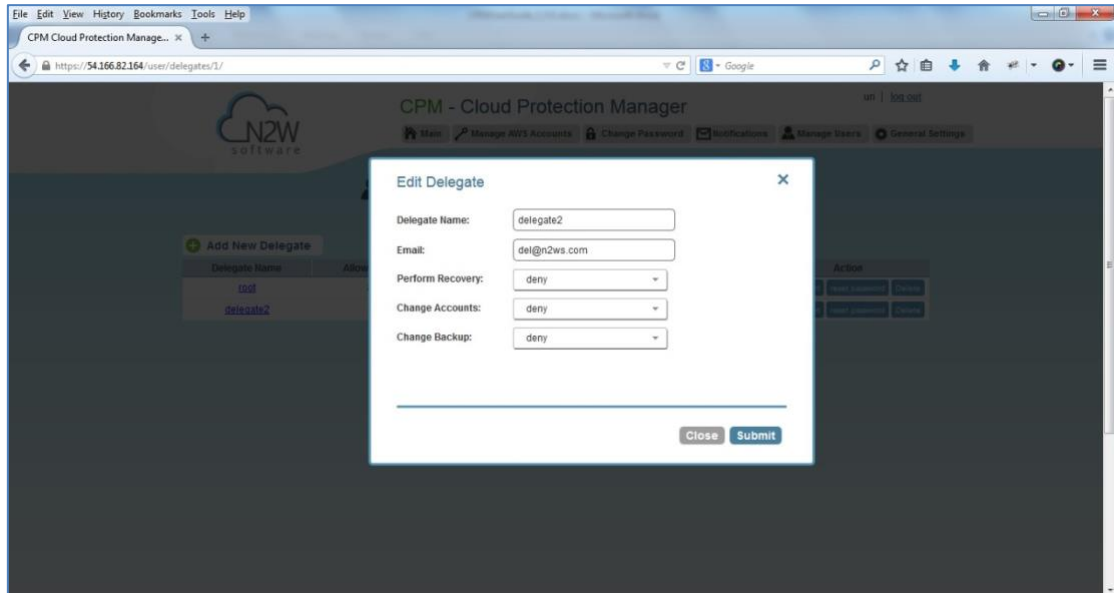


Figure 16-4

In a separate button in the delegates screen, the root user can reset passwords for delegates.

16.4.1 Delegate Permissions

There are three permissions for delegates:

- **Allow recovery** – Can perform recovery operations
- **Allow Account Changes** – Can add and remove AWS accounts as well as edit accounts and modify credentials
- **Allow Policy Changes**- 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.

16.5 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.

16.6 Audit Reports

CPM 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 **General Settings** screen.
- Download audit reports for specific users or delegates by clicking **audit report** in the users or delegates screen.
- Download the audit report for all users by clicking the link **audit report for all users** at the bottom of CPM's 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.

17 CPM IdP Integration

CPM 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 CPM User Management capabilities described above.
- IdP users are users whose credentials are received from the organization's IdP. CPM can be configured to allow users in the organization's IdP system to login to CPM using their IdP credentials. Integration with IdP systems is performed using the SAML 2.0 protocol.
- CPM supports Active Directory 2012 and 2016.

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

Configuring CPM to work with IdP consists of the following:

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

17.1 Configuring IdPs to Work with CPM

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

17.1.1 Prerequisite to IdP Integration with CPM

Prior to configuring CPM to work with an IdP system, it is required that CPM 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 CPM as a new IdP application, verify that:

- The default Name **ID** format used in SAML requests is set to **Unspecified**, or modify the default CPM configuration as per section on CPM configuration below.
- The X509 certificate Secure hash algorithm is set to SHA-256.
- The following URL values are used:
 - Note:** <CPM_ADDRESS> is either the DNS name or the IP address of the CPM Server.
 - **Entity ID** - https://<CPM_ADDRESS>/remote_auth/metadata
 - **Sign in response** - https://<CPM_ADDRESS>/remote_auth/complete_login/
 - **Sign out response** - https://<CPM_ADDRESS>/remote_auth/complete_logout/

As part of configuring CPM as a new IdP application, the IdP system will request a file containing the CPM x509 certificate. The certificate file can be obtained from the CPM **General Settings** page in the

Identify Provider Configuration section. Click the **Download CPM's certificate file** button and choose a location to save the file. See section 17.1.2.

If configuring CPM to work with Microsoft Active Directory/AD FS, refer to section 17.4.1.

17.1.2 Configuring CPM for IdP Integration

Configuring CPM to work with the organization's IdP is done by going to the CPM **General Settings** page and setting **Identity Provider** to `enabled`. Once enabled, several IdP-related parameters are presented (see Figure 17-1).

If configuring CPM for integration with Microsoft Active Directory/AD FS, refer to section 17.5.

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

CPM Server Identifier ▶

Backup Tag Scan ▶

Cleanup ▶

Proxy ▶

AWS Account for File Level Recovery ▶

Identity Provider ▼

Identity Provider:

Enabled ▼

Clear Fields

CPM IP or DNS:

172.31.43.104 ▼

Entity ID:

Identity Provider identifier (URI)

Sign in URL:

Authentication request target (URL)

Sign out URL:

Logout request target (URL)

NameID format:

Unspecified ▼

x509 cert:

Choose File

No file chosen

Uploaded file: okta.cert

Download CPM's certificate

Download CPM's metadata

Test connection...

+ Add New Group

Name	Enabled	Actions
default_independent_users	Yes	
default_managed_users	Yes	
default_root_delegates	Yes	
default_root_delegates_readonly	Yes	
disabled_grp	No	<div>Delete</div>
my_independent	Yes	<div>Delete</div>

Apply

Figure 17-1

- **Identity Provider** – Enables/disables access for IdP users.
- **CPM IP or DNS** – The IP Address or DNS name of the CPM server.
- **Entity ID** – The IdP **Identity Provider Identifier** 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 CPM 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 CPM will redirect users once they logout of CPM. Consult the IdP system's documentation.
- **NameID format** – The format of the SAML **NameID** element.

- **X509 Cert** – The X509 certificate is provided by the IdP system for uploading. Consult the IdP system's documentation about obtaining their x509 certificate.

Once all the parameters have been entered, click the **Test connection** . . . button to test the connection between CPM and the IdP.

17.2 Configuring Groups and Group Permissions on the CPM Side

Groups and the permissions assigned to groups are configured in CPM. When an IdP user logs into CPM 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 CPM group. IdP users who do not belong to a group, even if they have user-specific permissions as detailed below, cannot log on to CPM. Logon by IdP users who do not belong to a group will be failed with an appropriate error message.

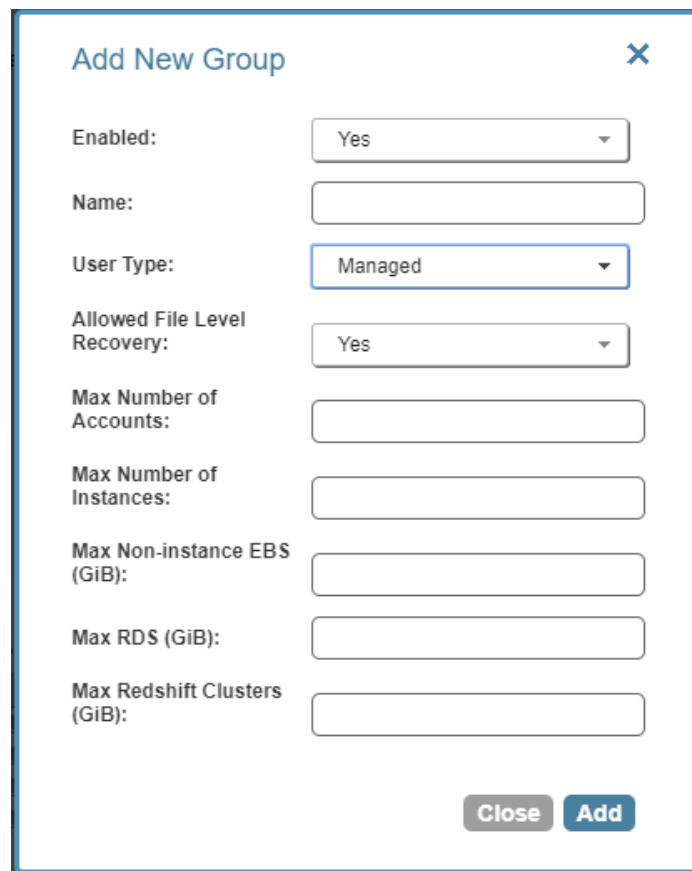
CPM comes with 4 pre-defined groups and additional groups can be created and removed easily in the **Identify Provider Configuration** section of the **CPM General Settings** page (see Figure 17-1).

Note: The default groups cannot be modified or deleted. To see the permission settings assigned to the default groups, click the group name.

To add a new group:

Click the **Add New Group** button. The add group screen will appear.

Note: The group permission settings essentially mirror the user permissions detailed in chapter 16.



Add New Group [X]

Enabled: Yes [v]

Name: []

User Type: Managed [v]

Allowed File Level Recovery: Yes [v]

Max Number of Accounts: []

Max Number of Instances: []

Max Non-instance EBS (GiB): []

Max RDS (GiB): []

Max Redshift Clusters (GiB): []

[Close] [Add]

Figure 17-2

- **Enabled** – When set to **No**, users belonging to the group will not be able to log on to CPM.
- **Name** – Name of the group.
- **User Type** – For details, see chapter 16.
 - Managed
 - Independent
 - 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 CPM. After creation, the Original Username cannot be modified.

- For User Type **Managed**:
 - **Allowed File Level Recovery** – When set to **Yes**, members of the group can use the file-level recovery feature.
 - **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** – 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** – The maximum number of Gigabytes of RDS databases that users belonging to this group can manage.
- **Max Redshift Clusters** – The maximum number of Gigabytes of Redshift clusters that users belonging to this group can manage.
- For User Type **Delegate**:
 - Original Username – User name of delegate.
 - Perform Recover – Whether the delegate can initiate a recovery.
 - Change Accounts – Whether the delegate can make changes to an account.
 - Change Backup – Whether the delegate can make changes to a backup.

17.3 Configuring Groups on the IdP Side

IdPs indicate a user's group membership to CPM 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 CPM related groups and non-CPM related groups.

Additionally, the names of the group users are assigned to in the IdP must be of the form `cpm_<GROUP_NAME_IN_CPM>` (e.g. `cpm_mygroup` where `mygroup` is the name of a group that was created in CPM). The `<GROUP_NAME_IN_CPM>` part of the name must match the name of a group in CPM (see section 17.3). For example, to give IdP users permissions of the CPM group `default_managed_users`:

1. The relevant users must be members of an IdP group called `cpm_default_managed_users`
2. The IdP must have an outgoing claimed called `cpm_user_groups` and 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`.

Note: An IdP user logging onto CPM can belong to only one CPM group, i.e. of all the groups listed in the `cpm_user_groups` claim, only one can be a CPM group, such as `cmp_mygroup`. If an IdP user is a member of more than one CPM group, the log on will fail with a message indicating the user belongs to more than one CPM group.

17.3.1 Understanding CPM User Permissions

A user logged into the CPM system can have several types of permissions. This section discusses the different types of permissions as they are applied to CPM IdP integration. For full treatment of the meanings of these permissions, see sections 16.3 and 16.4. To override CPM group permissions on a per user basis, see section 0

General User Attributes

Attribute Name	Mandatory (Y/N)	Meaning	Valid Values
user_type	N	Type of user.	<ul style="list-style-type: none"> Managed Independent Delegate
user_name	N	Username in CPM.	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 CPM file-level restore feature.	yes, no
max_accounts	N	The number of AWS accounts the user can manage in CPM. Varies by CPM license type.	Number between 1 and max licensed
max_instances	N	The number of instances the user can backup. Varies by CPM license type.	Number between 1 and max licensed
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

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 CPM restore operations.	yes, no
allow_account_changes	N	Whether the user can manage CPM 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 CPM. Additionally, it is possible to assign CPM permission at the level of individual IdP users as described in 0. 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. CPM 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_ac  
count_changes=yes;allow_backup_changes=yes;allow_file_level_restore=no;ma  
x_accounts=1;max_instances=2;max_independent_ebs_gib=3;max_rds_gib=4;max_  
redshift_gib=5;original_username=robi@stam
```

17.3.2 Overriding Group Settings at the User Level

Users get the CPM 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 CPM 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 CPM with permissions that are different from the group's permissions.

If configuring Microsoft Active Directory/AD FS, refer to section 17.6 for details.

17.4 CPM Login Using IdP Credentials

In order to use IdP credentials to log on to CPM, users need to select the **Sign in with: Identity Provider** option on the CPM Logon screen (see Figure 17-3).

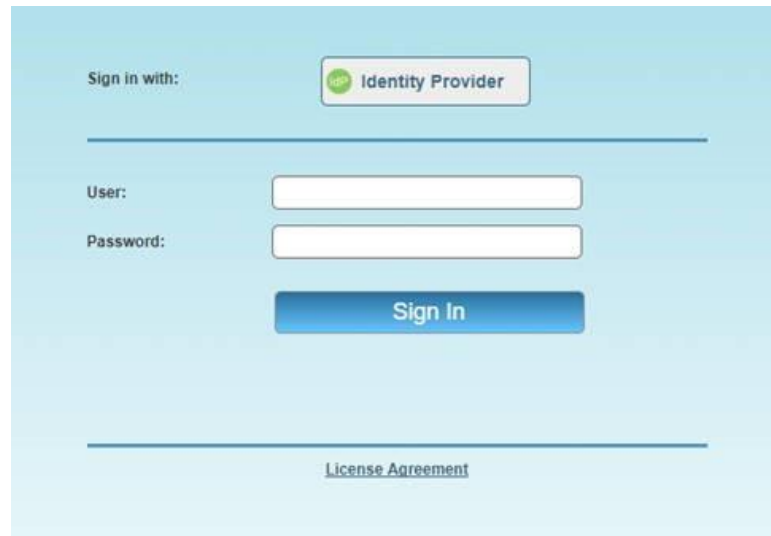


Figure 17-3

Clicking the **Identity Provider** button will redirect the user to the organization's IdP system using SAML.

Note: To log on to CPM as root, log on with the standard user and password option.

17.4.1 Configuring AD/AD FS for Integration with CPM

To enable CPM to integrate with AD/AD FS, CPM 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, click **Relying Party Trusts**.

In the right pane, click **Add Relying Party Trust**. ... The Wizard opens.

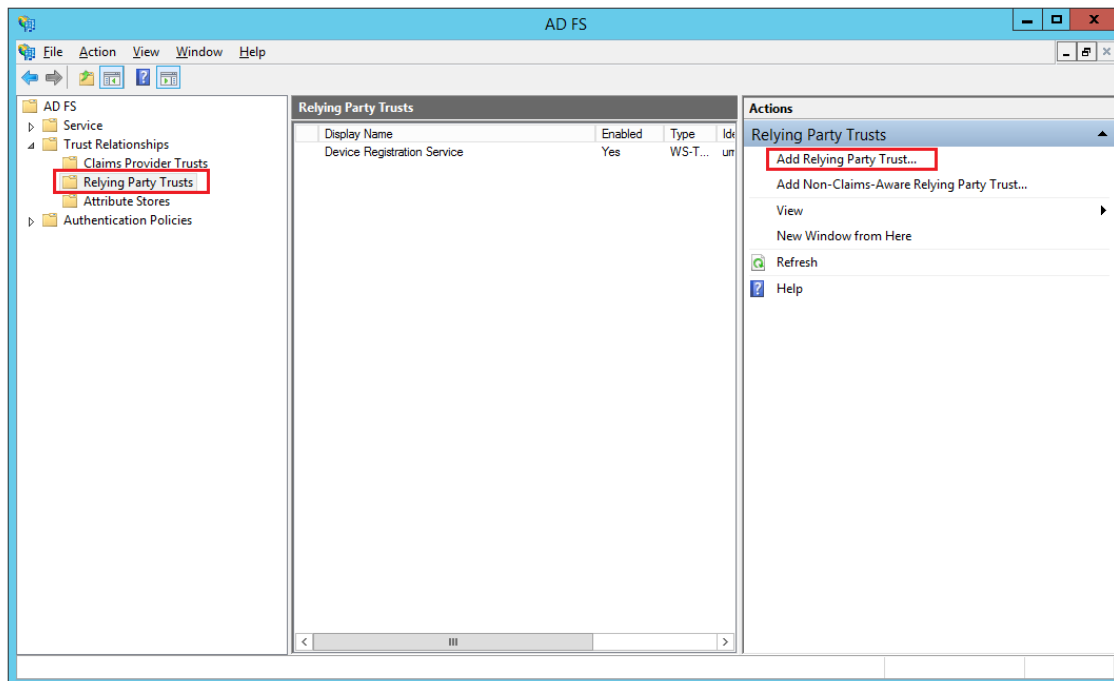
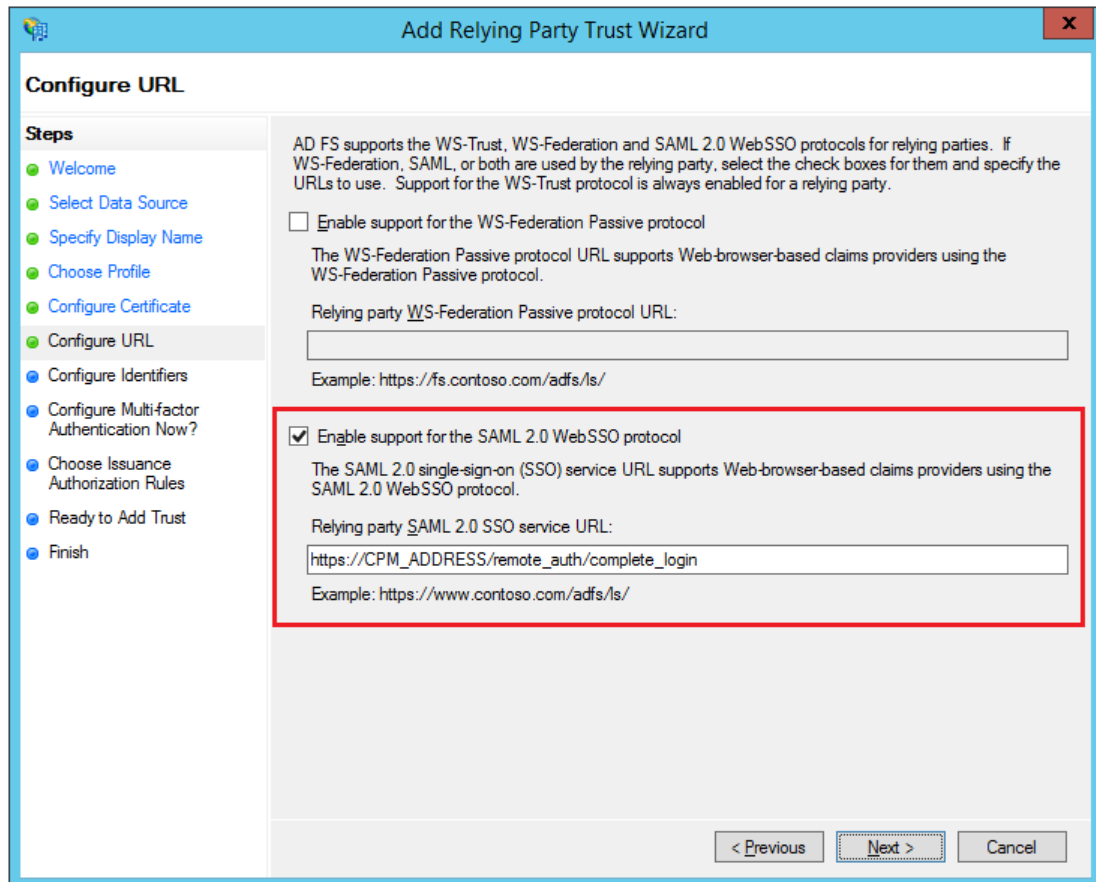


Figure 17-4

2. Click **Start**.
3. Click the **Enter data about the relying party manually** option.
4. Click **Next**.
5. On the **Welcome** screen, type the display name for CPM (e.g. CPM by N2WS), and click **Next**.
6. On the **Choose Profile** screen, click the **AD FS profile** option, and then click **Next**.
7. Skip the **Configure Certificate** screen by clicking **Next**.
8. On the **Configure URL** screen:
 - a. Select the **Enable support for SAML 2.0 WebSSO protocol** check box.
 - b. In the **Relying Party SAML 2.0 SSO Service URL** box, type `https://` followed by the CPM 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/`
9. Click **Next**.
10. In the **Configure Identifiers** screen, type `https://` followed by the CPM DNS name or IP address, and then followed by `/remote_auth/metadata` in the **Relying party trust identifier** box.



Add Relying Party Trust Wizard

Configure URL

Steps

- Welcome
- Select Data Source
- Specify Display Name
- Choose Profile
- Configure Certificate
- Configure URL**
- Configure Identifiers
- Configure Multi-factor Authentication Now?
- Choose Issuance Authorization Rules
- Ready to Add Trust
- Finish

AD FS supports the WS-Trust, WS-Federation and SAML 2.0 WebSSO protocols for relying parties. If WS-Federation, SAML, or both are used by the relying party, select the check boxes for them and specify the URLs to use. Support for the WS-Trust protocol is always enabled for a relying party.

☐ Enable support for the WS-Federation Passive protocol

The WS-Federation Passive protocol URL supports Web-browser-based claims providers using the WS-Federation Passive protocol.

Relying party WS-Federation Passive protocol URL:

Example: <https://fs.contoso.com/adfs/ls/>

☒ Enable support for the SAML 2.0 WebSSO protocol

The SAML 2.0 single-sign-on (SSO) service URL supports Web-browser-based claims providers using the SAML 2.0 WebSSO protocol.

Relying party SAML 2.0 SSO service URL:

https://CPM_ADDRESS/remote_auth/complete_login

Example: <https://www.contoso.com/adfs/ls/>

< Previous Next > Cancel

Figure 17-5

For example, the resulting string might look like:

https://ec2-123-245-789.aws.com/remote_auth/metadata

11. Click **Add** on the right.

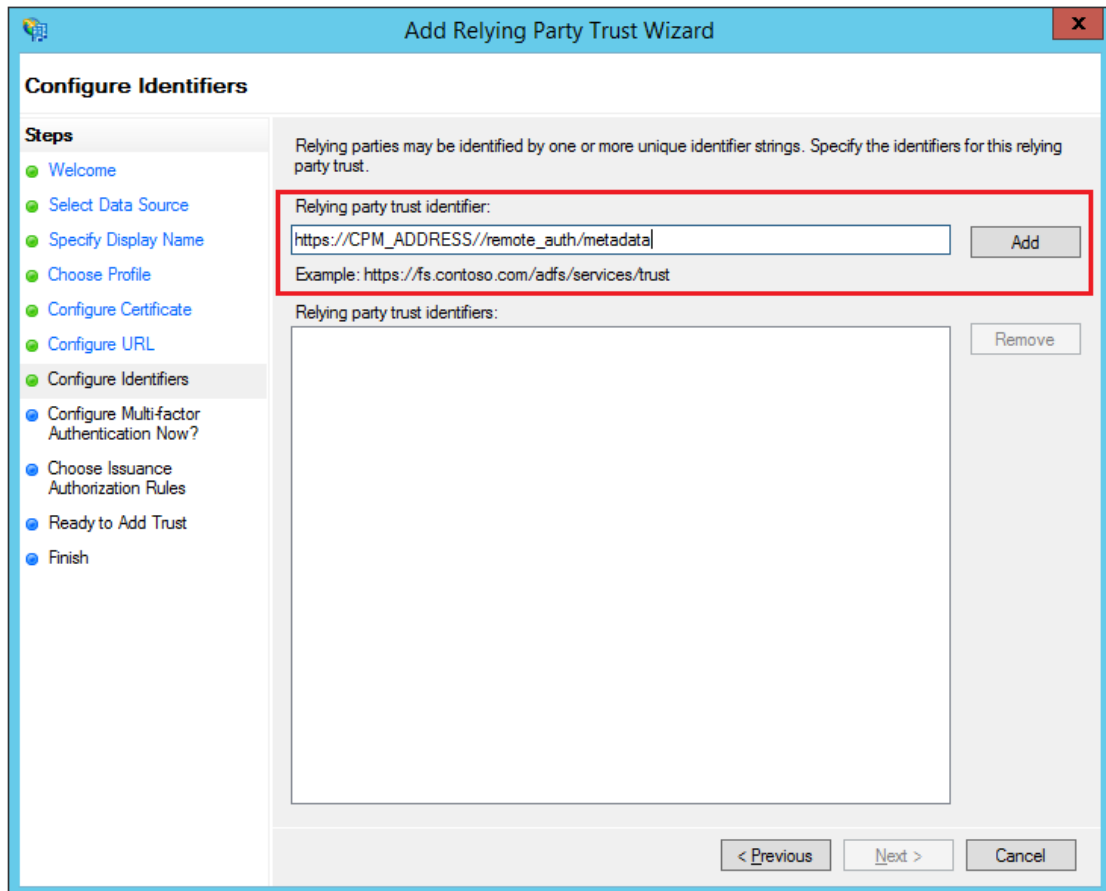


Figure 17-6

12. Click **Next**.
13. 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 click **Next**.
14. On the **Issuance Authorization Rules** screen, click the **Permit all users to access this relying party** option, and click **Next**.
15. On the **Ready to Add Trust** screen, review the setting of the **Relying party trust** configured with the Wizard. When finished, click **Next**.
16. On the **Finish** screen of the Wizard, click **Close**. There is no need to click the **Open the Edit Claim Rules** dialogue for this relying party trust when the wizard closes option.

17.4.2 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-click the CPM line under **Relying Party Trust**, and select **Properties**.

On the screen that opens, select the **Endpoints** tab, and click **Add SAML....**

- In the **Edit Endpoint** screen, select **SAML Logout** from the **Endpoint type** list.

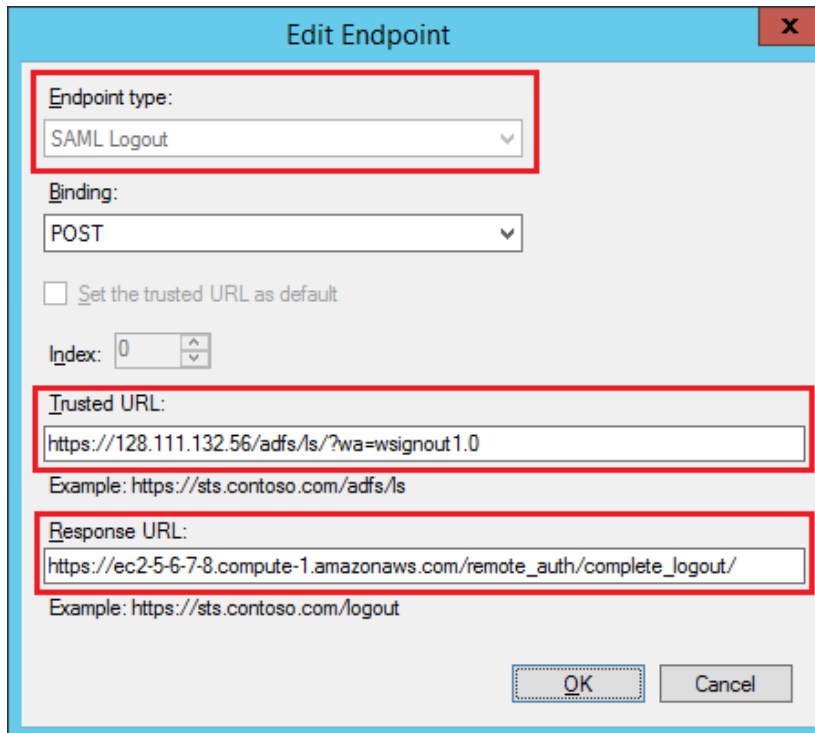


Figure 17-7

- 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`).
- In the **Response URL:** box, type DNS name or IP address of the CPM server followed by `/remote_auth/complete_logout/` (e.g. `https://ec2-123-245-789.aws.com/remote_auth/complete_logout/`).
- Click **OK**.
- Go to the **Advanced** tab, and in the **Secure hash algorithm** list, select **SHA-256**. Click **Apply**.

17.4.3 Installing the CPM Certificate

In order for CPM to work with AD FS the X.509 certificate used by CPM needs to be added to the AD FS **Trusted Root Certification Authorities** list. If you installed your own certificate in CPM when you first configured CPM (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 CPM 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:

- Go to the **Signature** tab under properties, and click **Add....**
- In the **File** box at the bottom of the screen, type the name of the file containing the CPM x.509 certificate. This will be either:

- a. The root certificate you installed in CPM 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 CPM created when it was first configured.

To obtain a copy of the certificate being used by CPM, either the one you originally installed or the one CPM created, click the **Download CPM's certificate file** button in the Active Directory Configuration section of the CPM General Settings screen (see Figure 17-13).

3. Once you have entered the name of the file, click **Open**.

The CPM certificate is now visible in the center pane in the **Signature** tab.

4. In the center pane of the **Signature** tab, double click the CPM certificate.
5. Under the **General** tab, click **Install Certificate....**
6. In the **Certificate Import Wizard** screen, click the **Local Machine** option, and click **Next**.
7. Click the **Place all certificates in the following store** option, click **Browse...**, and then select the **Trusted Root Certification Authorities** store. Click **OK**.
8. Click **Next**.
9. Click **Finish**. Then click **OK** on the pop-up screen, click **OK** on the **General** tab, and click **OK** on the **Properties** screen.

The next step is to create a Name ID claim in AD FS.

17.4.4 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 CPM's party (e.g. CPM by N2WS).
3. In the right pane, click **Edit Claim Rules...**

4. In the **Edit Claim Rules** screen, click **Add Rule**.

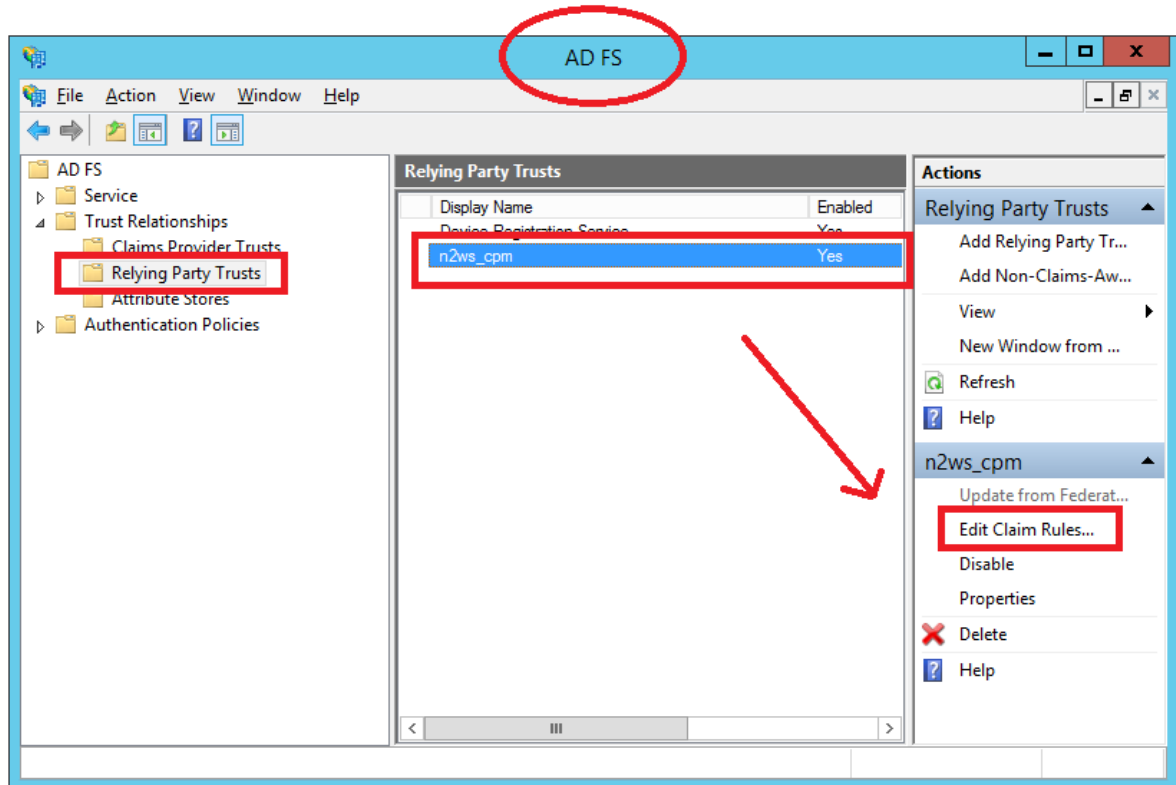


Figure 17-8

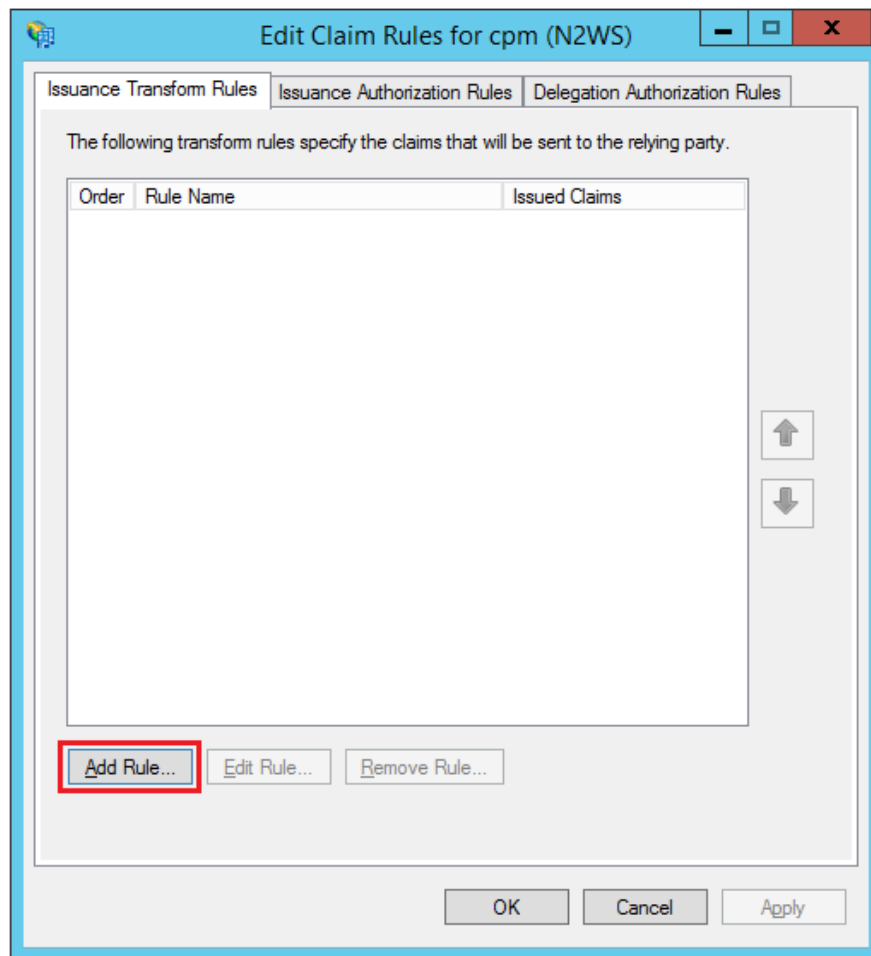
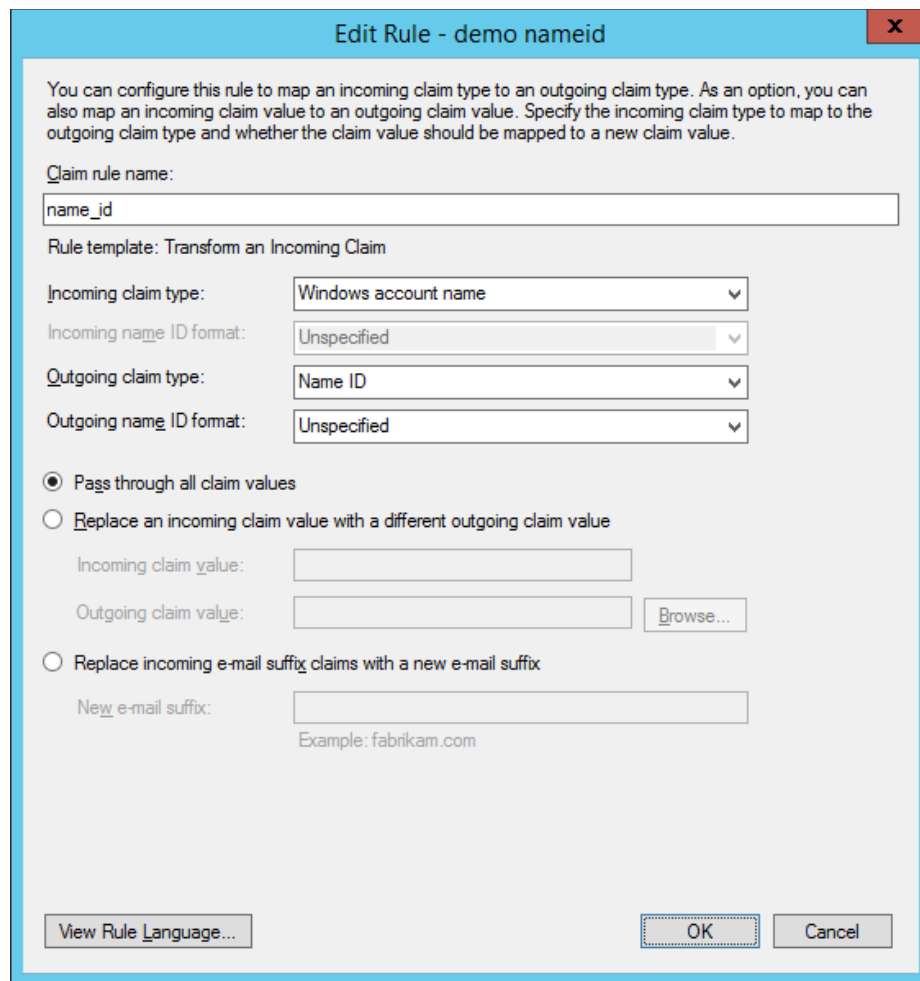


Figure 17-9

5. In the **Claim rule template** list, select **Transform an Incoming Claim** and click **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. Click the **Pass through all claim values** option.
 - f. Click **OK**.



Edit Rule - demo nameid

You can configure this rule to map an incoming claim type to an outgoing claim type. As an option, you can also map an incoming claim value to an outgoing claim value. Specify the incoming claim type to map to the outgoing claim type and whether the claim value should be mapped to a new claim value.

Claim rule name:

Rule template: Transform an Incoming Claim

Incoming claim type:

Incoming name ID format:

Outgoing claim type:

Outgoing name ID format:

☒ Pass through all claim values

☐ Replace an incoming claim value with a different outgoing claim value

Incoming claim value:

Outgoing claim value:

☐ Replace incoming e-mail suffix claims with a new e-mail suffix

New e-mail suffix:

Example: fabrikam.com

Figure 17-10

The next step is to add a Token-Groups claim.

17.4.5 Adding a Token-Groups Claim

An ADFS Token-Groups claim must be configured so that AD FS will send CPM the list of groups a user is a member of. To configure the Token Groups claim, perform steps 1 and 2 of the Configuring Name ID Claim process in section 17.4.4. Then continue as follows:

1. In the **Claim rule template** list, select **Send LDAP Attributes as Claims** and click **Next**.

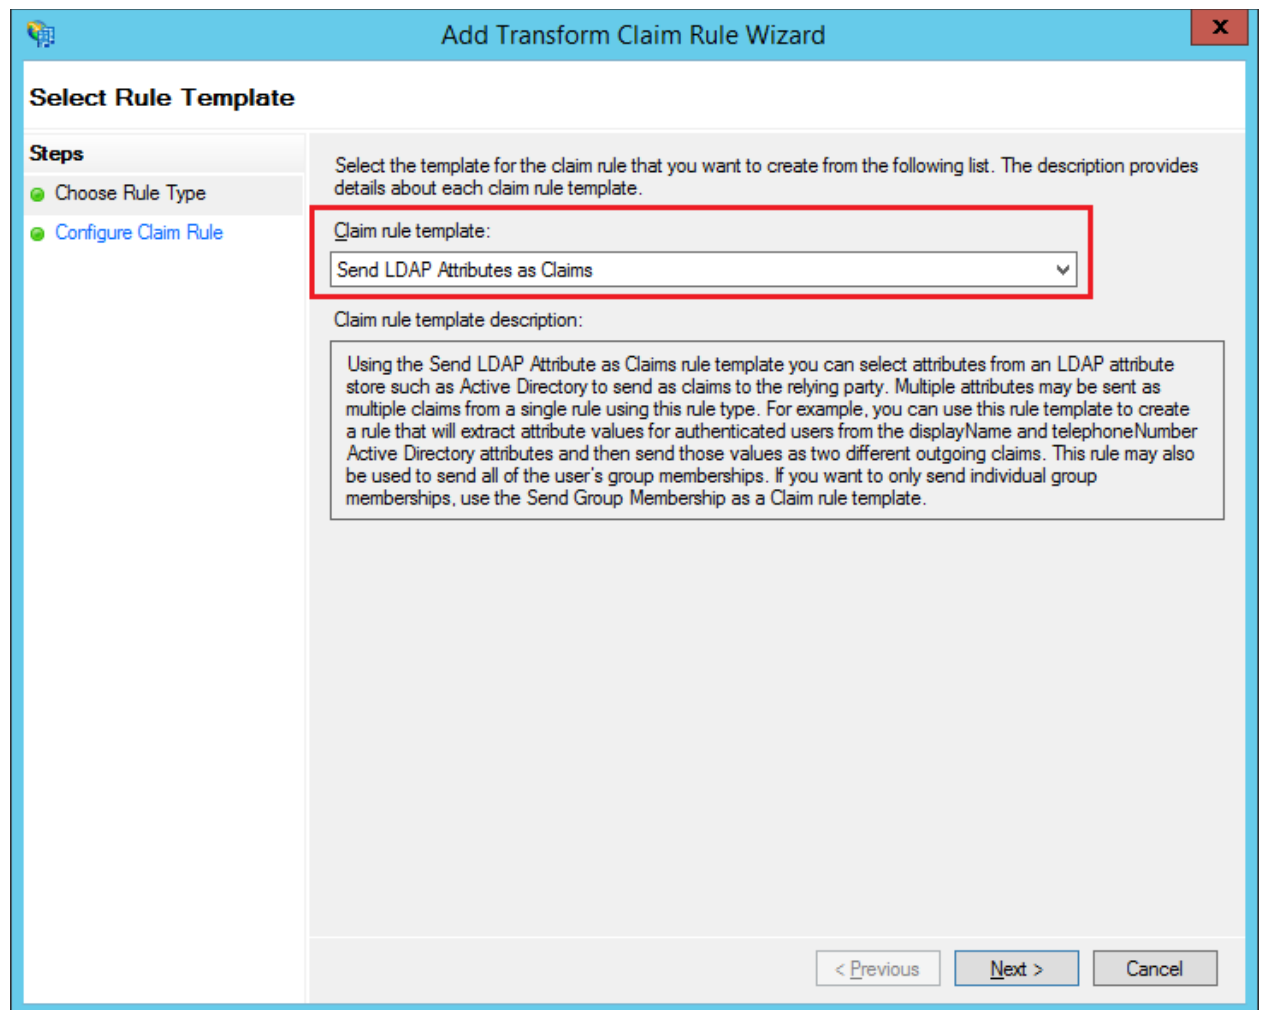


Figure 17-11

2. In the **Claim rule name** box, type a name for the rule you are creating.
 - a. In the **Attribute store** list, select **Active Directory**.
 - b. In the **Mapping of LDAP attributes to outgoing claim types** table:
 - i. In the left column (**LDAP Attribute**), select **Token-Groups - Unqualified Names**.
 - ii. In the right column (**Outgoing Claim Type**), type `cpm_user_groups`.

Edit Rule - user permissions claim
X

You can configure this rule to send the values of LDAP attributes as claims. Select an attribute store from which to extract LDAP attributes. Specify how the attributes will map to the outgoing claim types that will be issued from the rule.

Claim rule name:

Rule template: Send LDAP Attributes as Claims

Attribute store:

Mapping of LDAP attributes to outgoing claim types:

	LDAP Attribute (Select or type to add more)	Outgoing Claim Type (Select or type to add more)
	Token-Groups - Unqualified Names	cpm_user_groups
	msDS-cloudExtensionAttribute1	cpm_user_permissions
▶▶		

Figure 17-12

17.4.6 Testing the Connection

At this point AD FS has been configured to work with CPM. It is now possible to perform a connectivity test between CPM and AD FS.

To test the connection between CPM and AD FS:

1. Go to the CPM **General Settings** page.
2. Click **Test connection....**
3. Type a valid AD username and password on the logon page.
4. Click **Sign in**.

17.5 Configuring CPM to Work with Active Directory / AD FS

To configure CPM to work with the organization's AD server:

1. Go to the CPM **General Settings** page.
2. Select **Identity Provider Configuration**.
3. In the **Identity Provider** list, select **Enabled**. Several IdP related parameters are presented.



Name	Enabled	Actions
default_independent_users	Yes	
default_managed_users	Yes	
default_root_delegates	Yes	
default_root_delegates_readonly	Yes	
disabled_grp	No	<button>Delete</button>
my_independent	Yes	<button>Delete</button>

Figure 17-13

4. In the **Entity ID** box, type the AD FS **Federation Service Identifier**, as configured in AD FS. See Figure 17-14 to locate this parameter in AD FS.

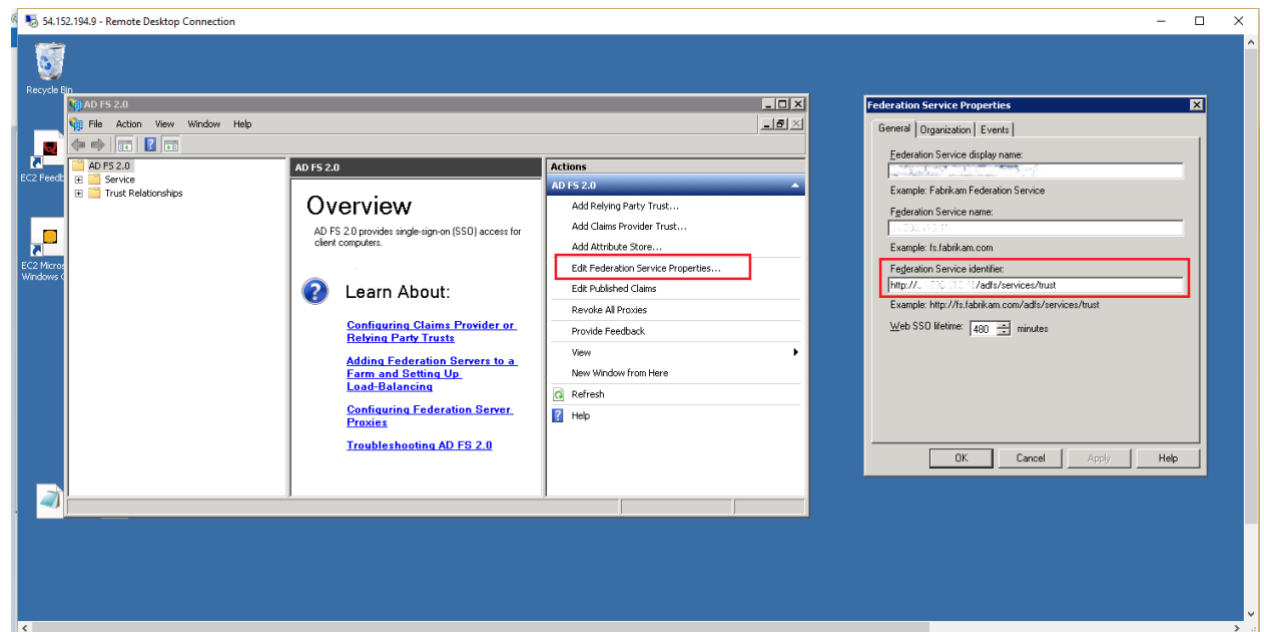


Figure 17-14

5. In the **Sign in URL** box, type the URL to which CPM 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 Figure 17-14 to locate this information in AD FS.




Figure 17-15

6. 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 Figure 17-16:

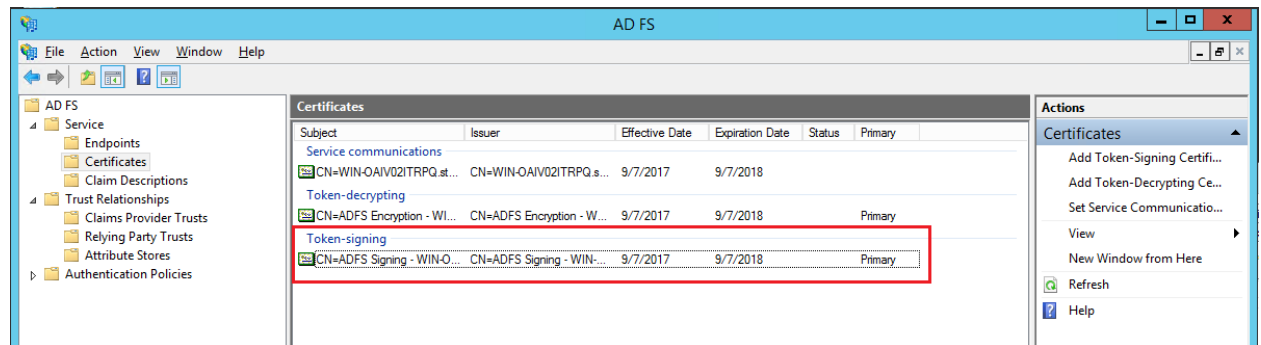


Figure 17-16

7. To export the certificate:

- Double click the **Token signing** field to open the **Certificate** screen.
- Click the **Details** tab, and click **Copy to File . . .** on the bottom right.
- Click **Next** to continue with the Certificate Export Wizard.
- Click the **Base-64 Encoded X.509 (.cer)** option, and click **Next**.
- Type a name for the exported file, and click **Next**.
- Click **Finish**.

Once all the parameters have been entered, click the **Test connection . . .** button to verify the connection between CPM and the IdP.

17.6 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:

- Open the AD FS management console.
- In the main page of the management console, in the left pane, select **Relying Party Trusts**.

3. Select CPM's party (e.g. CPM by N2WS) in the middle pane, and in the right pane, click **Edit Claim Rules**.

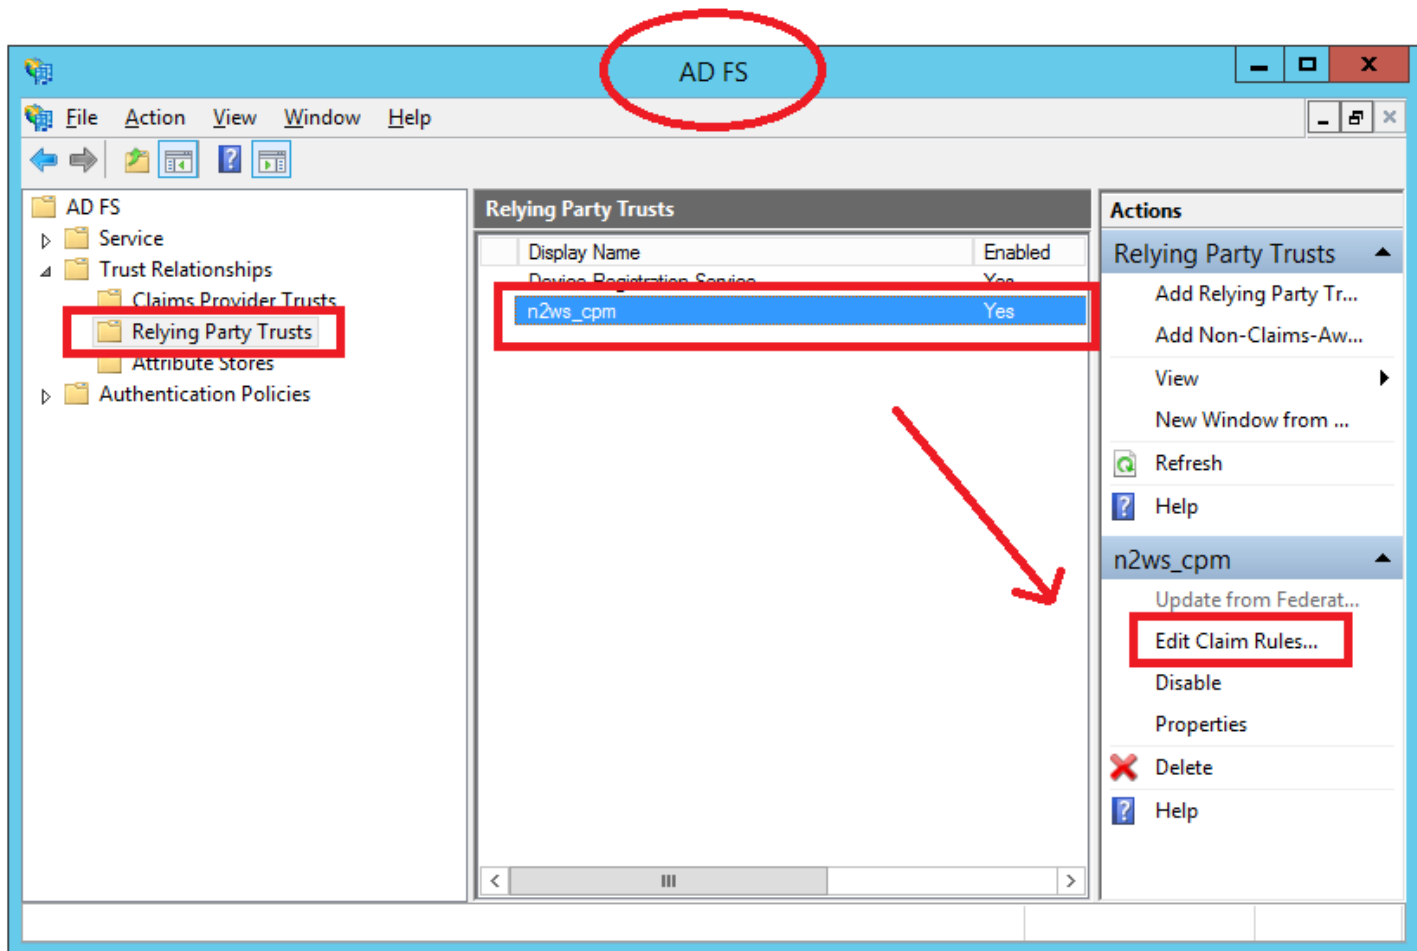


Figure 17-17

4. In the **Edit Claim Rules** screen, click **Add Rule**.

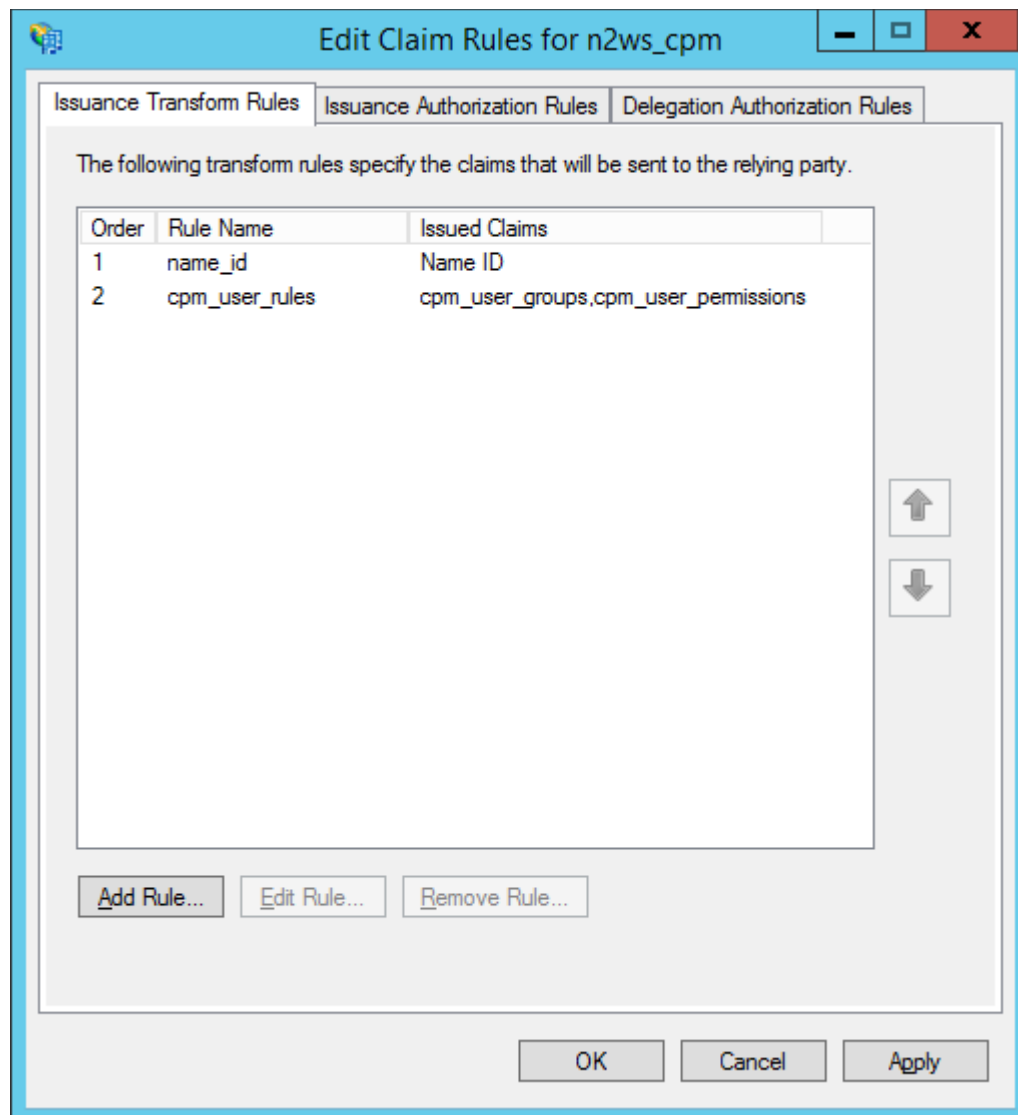


Figure 17-18




Figure 17-19

5. In the **Add Transform Claim Rule Wizard** screen, select **Send LDAP Attributes as Claims** in the **Claim rule template** list, and click **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`.

Edit Rule - user permissions claim
X

You can configure this rule to send the values of LDAP attributes as claims. Select an attribute store from which to extract LDAP attributes. Specify how the attributes will map to the outgoing claim types that will be issued from the rule.

Claim rule name:

Rule template: Send LDAP Attributes as Claims

Attribute store:

Mapping of LDAP attributes to outgoing claim types:

	LDAP Attribute (Select or type to add more)		Outgoing Claim Type (Select or type to add more)
	Token-Groups - Unqualified Names		cpm_user_groups
	msDS-cloudExtensionAttribute1		cpm_user_permissions
▶▶			

Figure 17-20

7. Click **OK** to create the rule.

Once the user-level claim is enabled, the user will be able to log on to CPM with permissions that are different from the group's permissions.