

CPM – User's Guide

V1.8.0

Contents

1	Introduction to CPM	7
1.1	What is CPM?	7
1.2	What you can do with CPM	7
1.3	Purchasing CPM on the AWS Marketplace	7
1.3.1	Purchasing	7
1.3.2	Moving between CPM Editions	8
1.3.3	Downgrading	8
1.4	CPM architecture.....	8
1.5	The CPM Server Instance.....	9
1.5.1	Root Volume	10
1.5.2	Backing up the CPM Server	10
1.5.3	Multiple CPM Servers	10
1.5.4	Upgrading the CPM Server Instance.....	10
1.6	CPM Technology.....	11
1.7	Browser support	11
2	Configuring CPM.....	12
2.1	General	12
2.2	Root user	12
2.3	Defining a time zone and data volume type	13
2.4	Fourth stage of configuration.....	14
2.4.1	New data volume	15
2.4.2	Existing data volume.....	15
2.4.3	Web server settings.....	15
2.4.4	Anonymous Usage Reports	16
2.5	Registering and finalizing the configuration.....	16
2.6	Configuration troubleshooting.....	17
2.7	Modifying the configuration of a CPM Server.....	18
3	Start Using CPM.....	19
3.1	Main screen	19
3.2	Associating an AWS account	20

4	Defining Backup Policies.....	22
4.1	Schedules.....	22
4.1.1	Defining	22
4.1.2	Scheduling and Time Zones	22
4.1.3	Disabled times	23
	24
4.2	Policies.....	24
4.2.1	Creating a new policy	24
4.2.2	Adding backup targets.....	25
4.2.3	More Options	26
5	Introduction to Consistent Backup.....	30
5.1	Crash-consistent backup	30
5.2	Application-consistent backup	30
5.3	CPM and a “Point in Time”	30
5.4	Summary or “What Type of Backup to Choose”	31
5.4.1	Crash-consistent	31
5.4.2	Application-consistent.....	31
6	Windows Instances Backup	32
6.1	Introduction.....	32
6.2	Configuring CPM Thin Backup Agent.....	32
6.2.1	Associating an agent with a policy	32
6.2.2	Installing the agent	33
6.2.3	Changing Agent Configuration	33
6.2.4	Using the agent with an http proxy.....	34
6.3	Using VSS	34
6.3.1	Introduction.....	34
6.3.2	CPM’s use of VSS	35
6.3.3	Configuring VSS	36
6.3.4	Excluding and verifying VSS writers.....	36
6.3.5	Troubleshooting VSS issues	37
6.3.6	VSS Recovery	37
	39
6.4	Using backup scripts on Windows.....	39

6.4.1	“before” script	39
6.4.2	“after” script	39
6.4.3	“complete” script.....	40
6.4.4	Capturing the output of backup scripts.....	40
7	Linux/Unix Instances Backup	41
7.1	Connecting to the CPM Server	41
7.2	Backup scripts.....	41
7.2.1	General	41
7.2.2	“before” script	42
7.2.3	“after” script	42
7.2.4	“complete” script.....	42
7.2.5	Capturing the output of backup scripts.....	42
7.2.6	Troubleshooting and debugging backup scripts.....	42
7.2.7	Example backup scripts	43
7.2.8	Scripts and SSH access in a multi-user environment.....	45
8	Additional Backup Topics	46
8.1	CPM in a VPC Environment	46
8.2	Backup when an Instance is stopped	46
8.3	Backing up independent volumes	47
8.4	The Freezer	47
9	Performing Recovery	48
9.1	Recovery AWS credentials.....	49
9.2	Instance recovery	49
9.2.1	Basic options.....	49
9.2.2	Advanced options.....	51
9.2.3	AMI Assistant.....	54
9.2.4	“Workaround” recovery of a Windows instance.....	55
9.3	Volume recovery	55
9.4	RDS Database Recovery.....	58
9.5	Redshift Cluster Recovery	59
10	Disaster Recovery (DR)	62
10.1	Introduction.....	62
10.2	Configuring DR.....	62

10.3	How it actually Works?	64
10.4	DR and mixed-region policies	64
10.5	Planning your DR Solution	64
10.5.1	Considerations	64
10.5.2	Timing your DR processes	65
10.5.3	Performing DR on the CPM Server (The cpmdata Policy)	65
10.6	DR Recovery	66
10.6.1	DR Instance Recovery	66
10.6.2	A Complete Disaster Recovery Scenario	67
10.7	DR Monitoring and Troubleshooting	68
11	Cross-account DR, Backup and Recovery	70
11.1	Introduction	70
11.2	Snapshot Vaulting	70
11.3	Configuring cross-account backup	70
11.4	Cross-account DR and clean-up	71
11.5	Cross-account with cross-region	72
11.6	Cross-account recovery	72
12	Tag-based Backup Management	73
12.1	Introduction	73
12.2	The “cpm backup” tag	73
12.2.1	Adding to a policy or policies	73
12.2.2	Creating a policy from a template	73
12.2.3	Tagging a resource to be removed from all policies	74
12.3	Tag scanning	74
12.4	Pitfalls and troubleshooting	75
12.4.1	Pitfalls	75
12.4.2	Troubleshooting	75
13	Security Concerns and Best Practices	77
13.1	Introduction	77
13.2	CPM Server	77
13.3	Best security practices for CPM	78
13.3.1	Credentials rotation	78
13.3.2	Passwords	78

13.3.3	Security Groups	78
13.4	Using IAM	78
13.4.1	CPM Server Configuration Process.....	79
13.4.2	CPM Server IAM Settings.....	79
13.4.3	CPM Agent IAM Role	83
13.5	Thin Backup Agent.....	84
14	Alerts, Notifications and Reporting	85
14.1	Introduction.....	85
14.2	Alerts	85
14.3	“Pull” Alerts	85
14.4	Using SNS.....	87
14.4.1	Introduction.....	87
14.4.2	Configuring SNS	87
14.5	“Push” Alerts	88
14.6	Daily Summary.....	88
14.7	Raw Reporting Data.....	89
14.7.1	Backup view csv report.....	90
14.7.2	Snapshot view csv report	90
14.7.3	Keeping Records after Deletion.....	90
14.8	Usage Reports.....	91
15	CPM User Management	92
15.1	Independent Users	92
15.2	Managed Users.....	92
15.3	User definitions	93
15.4	Delegates	93
15.4.1	Delegate permissions	94
15.5	Usage Reports.....	95
15.6	Audit Reports.....	95

1 Introduction to CPM

1.1 What is CPM?

CPM – Cloud Protection Manager – is an enterprise-class backup, recovery & disaster recovery solution for the EC2 compute cloud. It is a software product that uses existing snapshot abilities of EBS volumes and RDS databases.

CPM is marketed as a service. When you register to use the service, you get permission to launch a virtual machine image (AMI) of an EC2 instance. After you launch the instance, and after a short configuration process, you can start backing up your data using CPM.

1.2 What you can do with CPM

Using CPM you can create backup policies and schedules. Backup policies define what you want to backup (what we call “Backup Targets”) and how. Backup targets can be of three types: EC2 instances (including attached EBS volumes or not), independent EBS volumes (regardless if they are attached, and to which instance), and RDS databases. In addition to backup targets, you also define other backup parameters: will it use an agent (Windows); will it run backup scripts; what will be the retry policy in case of a failure; how many generations of data you want to keep (older backups will be automatically deleted).

Schedules define the way you want to schedule 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 the schedule will be active on; and special times to disable it.

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.

1.3 Purchasing CPM on the AWS Marketplace

1.3.1 Purchasing

CPM comes in several different editions which represent different usage tiers of the solution. The price for using the software (CPM) is a fixed monthly price which varies between the different CPM editions.

To see the different editions with pricing and details, please go to our [pricing & purchase page](#) on N2WS’s web site. Once you subscribe to one of CPM’s editions, you can launch a CPM Server instance and start working. Only one CPM Server per subscription will actually perform backup. If you run additional instances, they will only perform recovery operations (see 1.5.3).

1.3.2 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:

- Terminate your existing CPM instance. It is recommended to do so while no backup is running.
- 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.
- It is recommended to create a snapshot of your CPM Data Volume before proceeding, just to be on the safe side. 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," so it's easy to find.
- Subscribe to the new CPM Edition and launch an instance. You need to launch the instance in the same availability zone the old one was. 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.
- During configuration choose "Use Existing Data Volume" and select the existing data volume.
- Once configuration completes, you'll continue to work with your existing configuration with the new CPM edition.

1.3.3 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 to manage the backup of 30 EC2 instances, and you moved to CPM Standard Edition, which allows only 25 instances. CPM will detect such a situation as a "compliance issue," will cease to perform backup, display a message and issue an alert specifying the problem.

To fix the problem, you can 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.

1.4 CPM architecture

CPM Server is a Linux based software appliance. It uses AWS APIs to access your AWS account. It allows managing snapshots of EBS volumes and RDS databases. Except in case of installing a Thin Backup Agent for Windows Servers, CPM does not directly access your instances. Access is done by the agent, or by a script that the user provides, which performs application quiescence.

CPM consists of three parts: a database that holds your backup related metadata; a Web/Management server that manages metadata; and a backup server that actually performs the backup operations. These components reside in the CPM server and should not concern you as a user.

The architecture of the CPM solution can be seen 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 doesn't 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. If you wish to have VSS or scripts support for application quiescence, you will need to install CPM Thin Backup Agent. The agent will get its configuration from the CPM server, using the HTTPS protocol.

CPM Solution Architecture




Figure 1-1

1.5 The 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. By definition, all persistent data and configuration reside on the data volume. From CPM's point of view, the root device is dispensable. You

can always terminate your CPM instance and launch a new one, then in a short configuration process continue working with your existing data volume.

1.5.1 Root Volume

Although you have access to the CPM Server instance by SSH, we expect you to consider the CPM Server to be a virtual appliance. We expect you not to change the OS and not to start running additional products or services on it. If you do so and it affects CPM and causes it to malfunction, we will not be able to provide you with support. Our first requirement will be for you to launch a clean CPM server. Please 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). That said, 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. We recommend you consult [N2W Software support](#) before doing so.

1.5.2 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, and that's fine. But if you create an image (AMI) of it and launch a new one later on, with the system and data volume, you will find that the new server will be dysfunctional. It will load and will allow you to perform recovery, but it will not continue performing backup. This is simply 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 10.5.3).

1.5.3 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's 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, and if you need more resources, you should just upgrade to a higher edition of CPM. If yours is a special case, and you do need to use more than one CPM server in your account, please contact [N2W Software support](#).

1.5.4 Upgrading the CPM Server Instance

In 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. Here are instructions for performing this upgrade/restart:

- Launch a new CPM Server instance in the same region and availability zone as the old one. You can launch the instance using the "[Your Software](#)" page in the AWS web site. To determine the availability zone of the new instance or to launch it in a VPC subnet, you'll need to launch the instance using the EC2 console rather than using the 1-click option.
- Terminate the old instance, preferably while no backup is being performed. Please wait until it is in "terminated" state.
- Recommended: go to the volumes view in AWS Management Console and create a snapshot of the CPM data volume. The volume is easy to find as it's typically named

"CPM Cloud Protection Manager Data." The snapshot is just for the case that there is a problem with the upgrade process and can be deleted afterwards.

- When the new instance is in "running" state, connect to it with a browser using https.
- Approve exception to the SSL certificate
- Step 3: Choose "Use Existing Data Volume," and paste in your AWS credentials.
- 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.

1.6 CPM Technology

As part of the "cloud" ecosystem, CPM relies on web technology. The management interface through which you manage your 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 with the HTTPS protocol, which means it is all encrypted. This is important, since sensitive data will be communicated to/from the CPM server: AWS credentials, CPM credentials, object ids of your AWS objects (instances, volumes, databases, images, snapshot ids etc.).

1.7 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. Most of CPM's testing has been done using Firefox, as this is our favorite browser. CPM has also been tested (and is supported) for 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

2.1 General

As with most other operations, you use a web interface to configure a new CPM Server. When launching a new CPM Server, the server will automatically create a new self-signed SSL certificate. This certificate will be used for the web application at the configuration stage. 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. This means that no two CPM servers will ever have the same certificate, and therefore it is perfectly safe to use. Since it is not signed by an external authority, you will need to approve an exception for your browser to start using CPM.

When you configure a CPM server you enter the following settings:

- Credentials for the CPM root user
- The time zone for the server
- Whether to create a new CPM data volume, or attach an existing one (from a previous CPM server)
- The port the web server will listen on. The default is 443.
- Whether to upload an SSL certificate and private key for the CPM server to use. If you provide a certificate you need to provide a key as well. The private key must not be protected by a passphrase, or the application will not work.
- Register the AWS account with N2W Software – this is mandatory only for free trials but recommended to all users. I will allow us to provide quicker and better support. This information will not be shared with anyone.

Furthermore, 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 either needs to have a public IP, an Elastic IP attached to it, or it needs connectivity via a NAT setup or Internet Gateway. If you get an unexplained error during configuration, it may be because of a connectivity issue. If such an issue happens, please check that the instance has Internet connectivity, that the DNS is configured properly and that the security group allow outbound connections for port 443 (HTTPS).

We will now go through the stages one by one.

2.2 Root user

The root user is the user that controls all the operations of the CPM server. Root user credentials are used to log in the system and to use it. As you can see in Figure 2-1, you need to define the user name, email, and password. The email may be used when defining SNS-based alerts and notification. You can then choose to automatically add this email to the SNS topic recipients. Also, If you are using the Free Trial & BYOL Edition, you will have the

“license” field. Choose “Start free trial” for a free trial, and if you purchased a license directory from N2W Software, you’ll get instructions.

Passwords: We do not enforce any password rules. However, it is recommended to use

Info 2-1

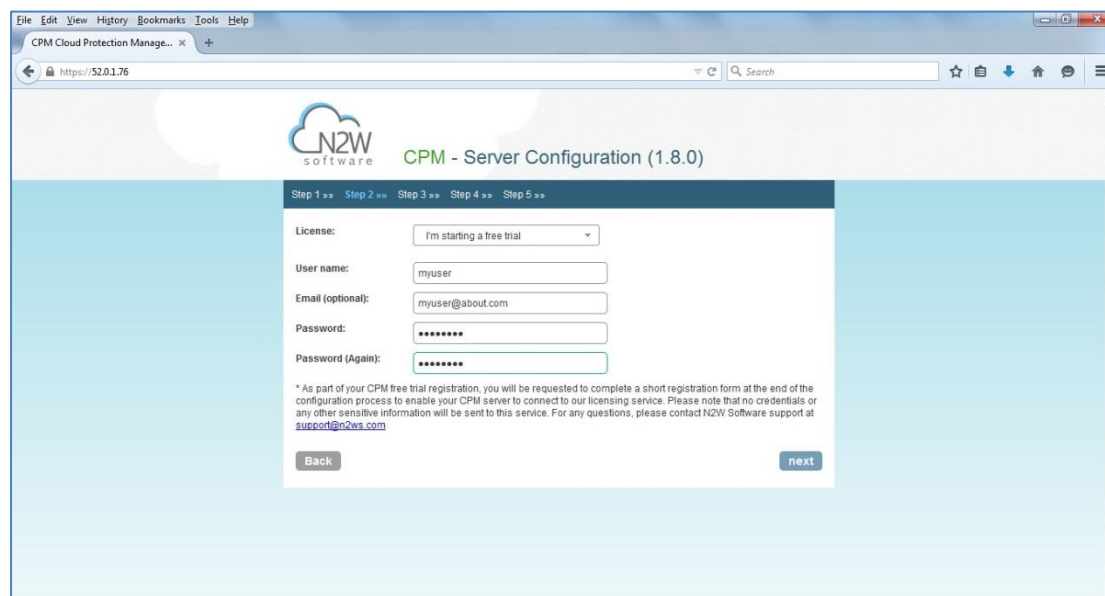


Figure 2-1

2.3 Defining a time zone and data volume type

In the second step of the configuration process, you define the time zone of the CPM Server. You 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.

As you will see later (see 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 needed to create a new EBS data volume if needed, and to attach the volume to the CPM Server instance. If you are using IAM to 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 13.4). These

credentials are not kept or logged anywhere after using them, and they are used only for the above-mentioned purpose. If you assigned an IAM Role to the CPM Server instance, and this role includes the needed permissions, you can just state “USE_IAM_ROLE” in the access key and secret key fields instead of entering actual credentials.

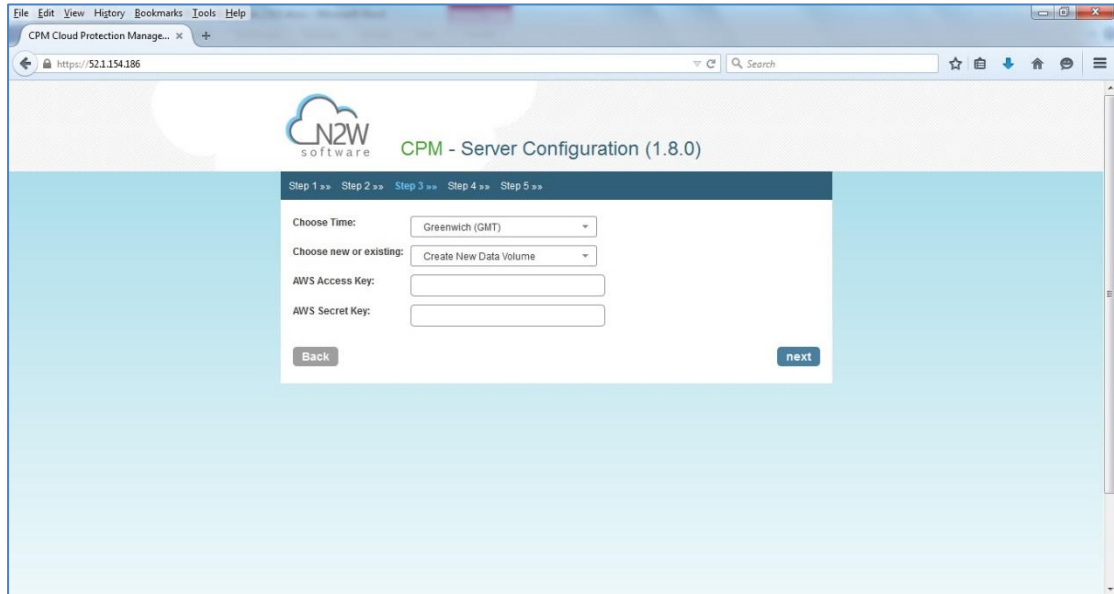


Figure 2-2

2.4 Fourth stage of configuration

At this stage (stage 4), you will fill in the rest of the information needed to configure 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, instead of the capacity field, you will see a drop-down select box, from which to choose the volume.

2.4.1 New data volume

When you choose to create 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 5Gib. This will be large enough to manage roughly 50 instances (and about 3 times that much EBS volumes). If your environment is bigger than that, make the volume bigger at about the rate of 1Gib per 10 backed-up instances.

The new volume will be automatically created in the availability zone the CPM Server instance is in. 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

You will use this option, in case you already ran CPM before, and for some reason you terminated the old CPM server, and now wish to continue where you stopped. Using an existing volume is also needed for upgrading to new CPM releases (when they will be available), and to change some configuration details.

The select box for choosing the volumes will show all available EBS volumes in the same availability zone as the CPM Server instance. It is important to create the instance in the availability zone 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 at any availability zone you require.

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

2.4.3 Web server settings

You can configure a different port for the web server. Port 443 is the standard port for the HTTPS protocol, which is used by CPM, so it is the default. If you wish you can change it to a different port. Please bear in mind that the specified port will need to be open in the instance security group/s 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 this value as “Allowed,” enables CPM to send anonymous usage data to N2W Software. This data does not contain anything identifying who the user is. No AWS account numbers or credentials. No AWS objects or ids like instances or volumes. No CPM objects names, like policy and schedule names. It contains only details like how many policies run on a CPM server, how many instances per policy, and 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 (step 4), you will the last screen that asks you to register. This is mandatory for free trials and optional for paid products.

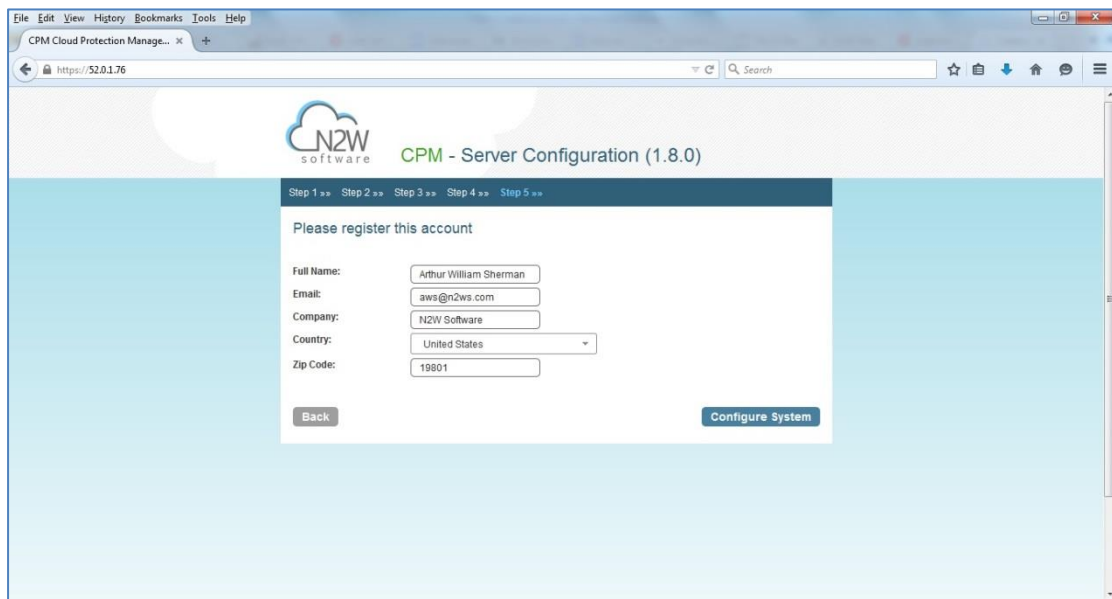


Figure 2-4

Click on “Configure System” to finalize the configuration. The configuration will take somewhere between 30 seconds and 2-3 minutes for new volumes, and usually less for attaching existing volumes. After the configuration is complete, you will be redirected to a screen that will indicate it:



Figure 2-5

After you see the success page, you know that CPM was configured correctly. Click on the “here” link and wait a few seconds. You should be redirected to the login screen of the CPM application. If, for some reason, you are not redirected, try to refresh the browser manually. If that doesn’t work, just 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 different configuration steps are checked when you click “next,” and you will usually get a clear and straightforward message indicating what went wrong. These errors are easy to correct.

A less obvious problem you may encounter is if you reach the third step and get the existing instance select box with only one value in it: “No Instances found.” This can arise from two reasons. The first is obvious - if you chose to use an existing volume, and in the CPM Server’s availability zone there are no available EBS volumes, you will get this response. In this case, you probably did not have your existing data volume in the same availability zone. To correct this, you can either terminate and relaunch the CPM server instance in the correct zone and start over the configuration process, or you can take a snapshot of the data volume, and create a volume from it in the zone the server is in. The other reason that may cause this issue is a problem with the credentials you typed in. In this case 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 this, please go back and enter the credentials correctly.

In very rare cases, you may encounter a more difficult error. It is discovered after you already approved configuring the server. In this case you will usually get a clear message regarding the nature of the problem. This type of problems 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 they need (in case they were created using IAM); and also, if you chose a bad port (e.g. the SSH port which is already in use), or if you specified an invalid SSL certificate and/or private key files.

In case you can't figure the problem, you can try again. If it persists, please contact N2W Software support (support@n2ws.com).

In any case, if the error occurred after approving the last configuration stage, it is recommended to terminate the CPM server instance, delete the new data volume (if one was already created), and 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 had already been create, 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 by terminating the current CPM server instance and creating a new one. This should not take more than a few minutes. After you terminate the CPM server, the data volume becomes available. You need to remember to launch the new server in the same availability zone.

All you need to do is to configure the server as you wish, and connect to the old data volume.

As for the CPM root user, you may change the email or the password. The username of the root user can't 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, `"/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 7.1).

3 Start Using CPM

3.1 Main screen

As soon as you log in 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.

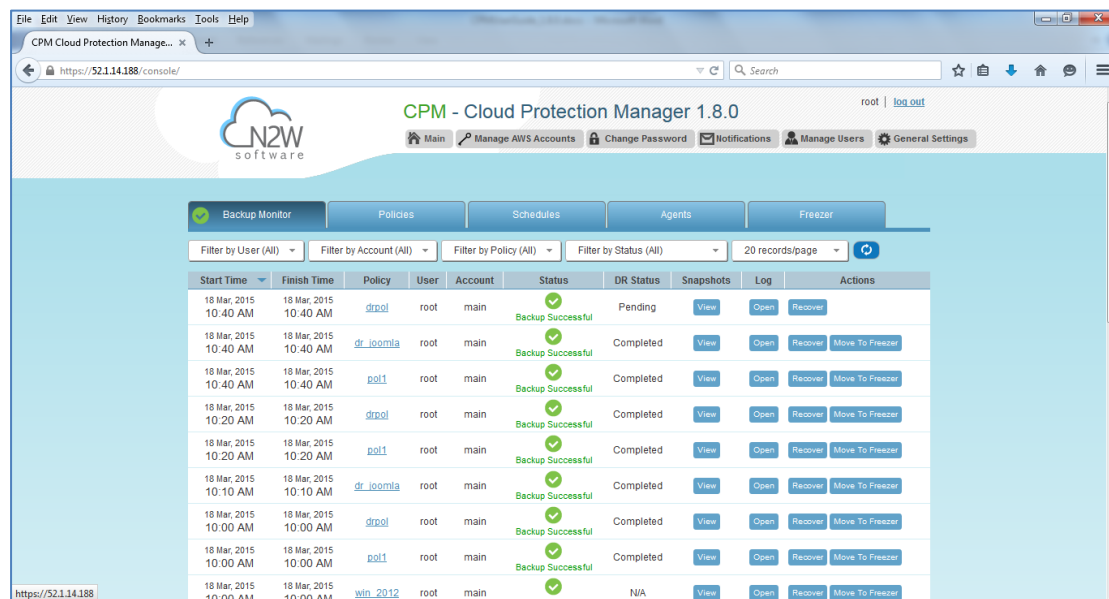


Figure 3-1

As you can see in **Error! Reference source not found.**, 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 8.4). 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.

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. It can also be used to 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 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, whether it's the root user or a different one.
- Notifications – Define notifications and alerts.
- Manage Users – Depending on the CPM edition you are registered to, you may have the ability to create additional users. By clicking this button you may create new users, or manage existing ones, i.e. delete them, reset their passwords or download a usage report. Only the root user may create and manage other users.
- General Setting – 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 tar ball (in case you need to send to our support team).
- To enter a new activation key. If any 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: if you haven't done this already during configuration, it is recommended to do so, as it will allow us to provide faster and better support.

3.2 Associating an AWS account

In order to associate an AWS account, you will need to enter AWS credentials (an access key and secret key). You can obtain your credentials in the IAM console:

<https://console.aws.amazon.com/iam/home?#users>

Click “Manage AWS Accounts” and then click “Add New Account.” Enter the credentials and a unique name for the account. The name should be something meaningful like “main,” “accounting,” “production,” etc...

If you are using Advanced or Enterprise Edition as well as a free trial, you will need to choose an account type. Backup account is an account used to perform backup and is the default. “DR Account” is an account used to copy snapshot to as part of cross-account functionality.

If you are logged in as the CPM root user (the first user defined in the configuration stage), and the CPM server instance is associated (at launch time) with an IAM role (see 13.4), you can define an account without credentials and the IAM role will be used. To do that, simply enter the string “USE_IAM_ROLE” (case doesn’t matter) as your access key as well as your secret key.

As the root user you are also able to add accounts for other managed users. If you are the root user and have managed users defined, an additional select box will be added, allowing you to select the user.

“Scan Resources” allows you to determine whether the current account will be included in scan tags performed by the system.

If this is a DR account, you choose whether this account is allowed to delete snapshot. If it is not, CPM will not delete snapshots of this account when performing cleanup of outdated backups. It will tag them instead. Not allowing CPM to delete snapshots of this account implies that the IAM credentials given do not have the permission to delete snapshots.

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 backup → schedules
- Whether to use backup scripts
- Whether VSS is activated (Windows 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

We will go through the stages of defining a policy.

4.1 Schedules

Schedules are the objects defining **when** to perform backup. Schedules are defined separately from policies. A schedule can be associated with several policies. Multiple schedules can be associated with the same policy.

4.1.1 Defining

To define a schedule click on the “Schedules” tab in the main screen, then click on “New Schedule.” You need to enter a name and an optional description. The main field, defining the behavior of the schedule, is “Repeats Every.” It defines the frequency of the backups this schedule will launch. The possible units are months, weeks, days, hours, and minutes.

The other important field is “Start Time.” Start time determines when the schedule will start. If you want a daily backup to run at 10:00 AM, you set “Repeats Every” to one day, and the start time to 10:00 AM (the date can also be set, the default is the current day). If you want an hourly backup to run at 17 minutes after the hour, you set “Repeats Every” to one hour, and the start time to XX:17. All backup times are derived from the start time. Please note: for weekly or monthly backups, the start time will also determine the day of week of the backup schedule and not the days of week check-boxes.

“End Time” is when the schedule will expire. By default it’s never, because typically schedules are not temporary. Furthermore, you can define which week days 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.

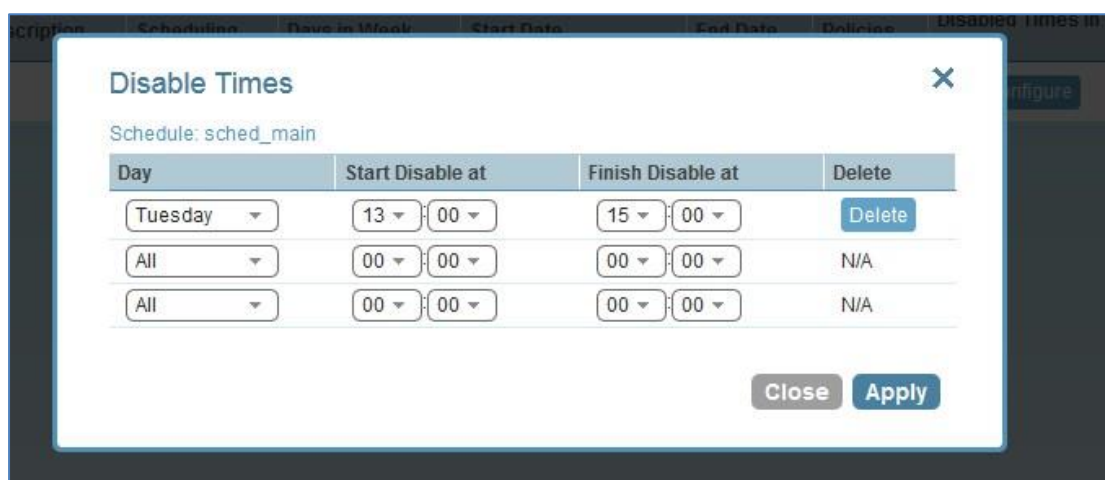
4.1.2 Scheduling and Time Zones

When you configure a CPM server, you set, among other things, its time zone (see 2.3). All time values which are displayed in CPM’s management application are in that same local time. Bear in mind, that even when you are backing up instances that are in different time

zones, the scheduled backup time is always according to 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. An example for a case when this feature is needed is if you want backup to run every hour. However, you do not want it to run Tuesdays between 01:00 PM and 3:00PM. Maybe these are critical business hours, and you do not want the application quiescence process to slow the system during that time. You can simply define that Tuesdays, between these hours, the schedule is disabled. You define disabled times by clicking the "Configure" button in the "Disabled Times in Day" column of the "Schedules" Tab. You can add as many as you want, edit, or remove them, as seen in Figure 4-1.



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

Figure 4-1



You can define a disabled time when the finish time is earlier than the start time. The meaning of disabling the schedule Mondays 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 every day between 18:00 and 6:00 will be that every day the schedule will be disabled until 6:00 and after 18:00.

Info 4-1



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 uncheck Monday from the week days. It is also possible to define "disabled times," which make sure the schedule is always disabled when it is originally supposed to run.

Warning 4-1

4.2 Policies

Policies are the main objects defining backups. With a policy you define what to backup, how to back it up, and by associating schedules, when to perform backup.

4.2.1 Creating a new policy

To define a new policy, simply go to the "Policies" tab and click "New Policy." In the policy popup window, you define a name for the policy, and an optional description. You can also set the account the policy is associated with. This is only meaningful if you have more than one account. The account can't be modified after the policy is already created. Furthermore, you decide how many generations of backup data to keep for this policy, which means that older backups will be automatically deleted by CPM. If you create one 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.



As a user, you need to balance the amount of time you want to be able to go back and recover from (RPO – Recovery Point Objective), 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.

Info 4-2

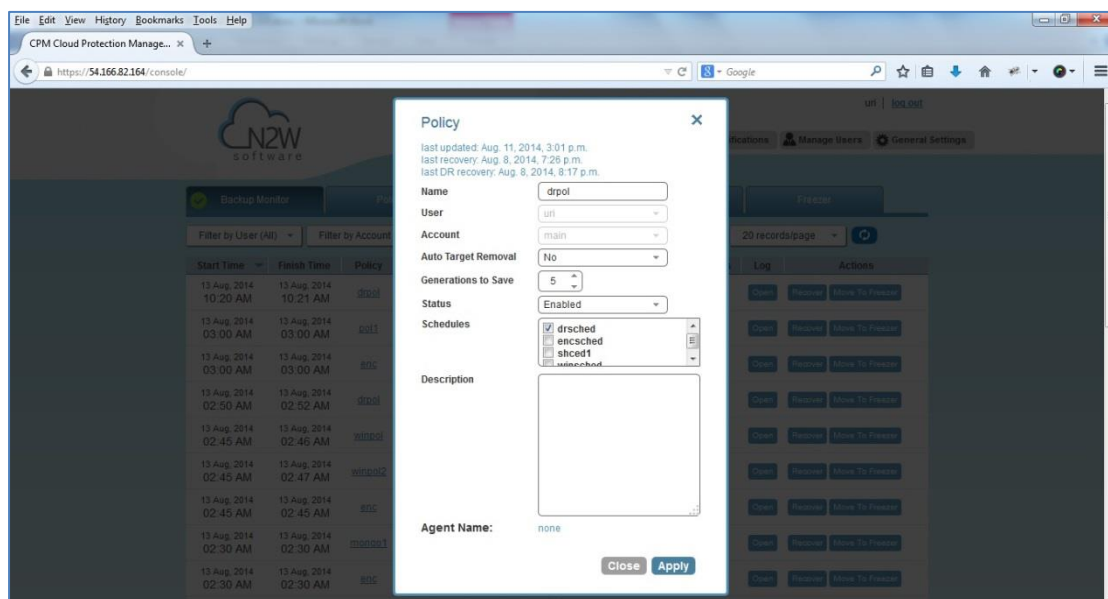


Figure 4-2

The field “Auto Target Removal” will specify whether to automatically remove resources that no longer exist. So, if you enable this removal, if an instance terminates, 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.

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

After clicking “Apply,” you can see the new policy on 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 three types of backup targets:

- **Instances** – This is the most common type. You can choose as many instances as you wish for a policy (limited by the number of instances you’re licensed to use). For each instance, you can define whether you back up all its attached volumes, some, or none.
- **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 move around between instances (e.g. cluster volumes).
- **RDS Databases** – You can use CPM to backup 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.

- **Redshift Clusters** – You can use CPM to backup 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 you can click on “Add Instances,” “Add Volumes,” “Add RDS Databases,” or “Add Redshift Clusters” to add backup targets to the policy. When adding backup targets, you 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. In case you have many objects, you have the ability to filter, sort, or browse between pages. Furthermore, for each backup target, you can see the number of policies it’s already in (“Policies” column). If the number is larger than zero, you can click on it to see which policies it’s in. You can see the selection screen for instances in Figure 4-3

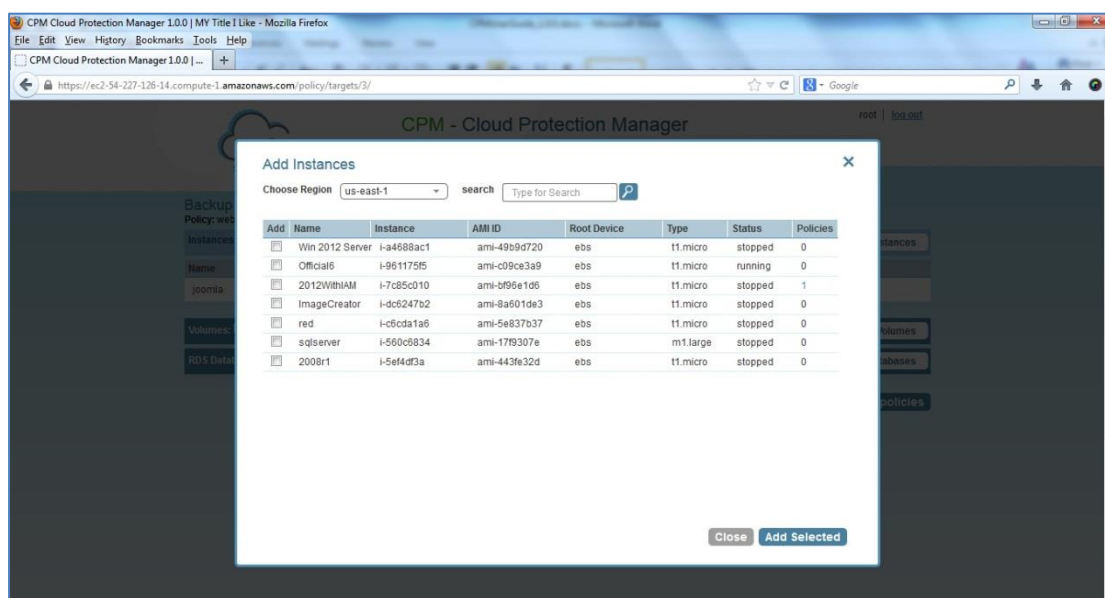


Figure 4-3

When you want to add an instance (or another type of backup target) to the policy, you simply check the “Add” checkbox of that instance (you can check more than one), and then click “Add Selected.” This operation will not close the popup window. It will remove the selected objects from the list and add them to the policy’s backup targets. You can repeat this operation as many times as you like, and click “Close” when you are done.

4.2.3 More Options

There are more options you can set for a policy. You can see them by clicking on “More Options” for a policy in the “Policies” tab. It is divided into three parts:

- The following options are available only for policies that have Windows instances in their backup targets:

- ‘Windows Agent’ – You can choose a Windows instance from the policy instances, which will serve as the policy’s backup agent. You can assign only one agent per policy, so if you want to back up a few Windows instances at the same time, you should define a separate policy for each of them. You can assign these policies the same schedule. If you want to change this option without closing the popup window, click on “set now.”
- “Backup Agent Key” – This is a generated authentication key for the agent. You will need to copy this key and paste it to the agent’s installation or configuration file. If the agent has more than one policy associated with it (you may want to back up the same instance in two different policies), the key is shared between them. If you want to generate a new backup agent key, click on “generate new.”
- “Agent Credentials” – In order to perform the backup, CPM sends the relevant AWS credentials to the agent. It sends them in an encrypted and secure way. However, if the instance the agent is on has an IAM role and that role is sufficient for performing backup (see 13.4.3), you can avoid sending credentials by choosing “Use IAM Role.”
- “Enable VSS on Agent” – By default this option is turned on. This means that VSS quiescence will be activated for this policy. 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 don’t need scripts, there is no use installing an agent, so just perform backup without one.
- “Volumes for shadow copies” – This option is used only if VSS is turned on. 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.
- The following options are related to backup scripts. They are valid for all types of instances. In Unix/Linux-based instances, the scripts will run 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 when scripts fail. 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 you need to choose “Collect.” The default option is “Ignore,” which means the output is not collected.



Note that the output of a script should typically be a few lines of output. It's okay if it is larger than that. However, if it gets really big (MBs) it can affect the performance of CPM. If it gets even larger, it can even cause crashes in CPM processes. Please make sure the scripts don't output large amounts of data to stderr. If there is any risk of this, make sure its output is redirected elsewhere.

Warning 4-2

- “Backup is Successful when...” - This indicates whether a backup needs its scripts/VSS completed, 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 for this option are:
 - “...it finishes with no Issues” – This means that 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 more strict 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 (this can happen in a stressful environment). Choosing this option accepts the assumption that most applications will recover correctly even from a crash-consistent backup.
- Retry information - the last three fields deal with what to do when a backup doesn't succeed:
 - “Number of Retries” - the maximal number of retries that can be run for each failed backup (the default is three). After this number of retries, the backup will only 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 very busy (during peak hours). In this case, it makes sense to substantially extend the time until the next retry, so there is a better chance the system will be more responsive.

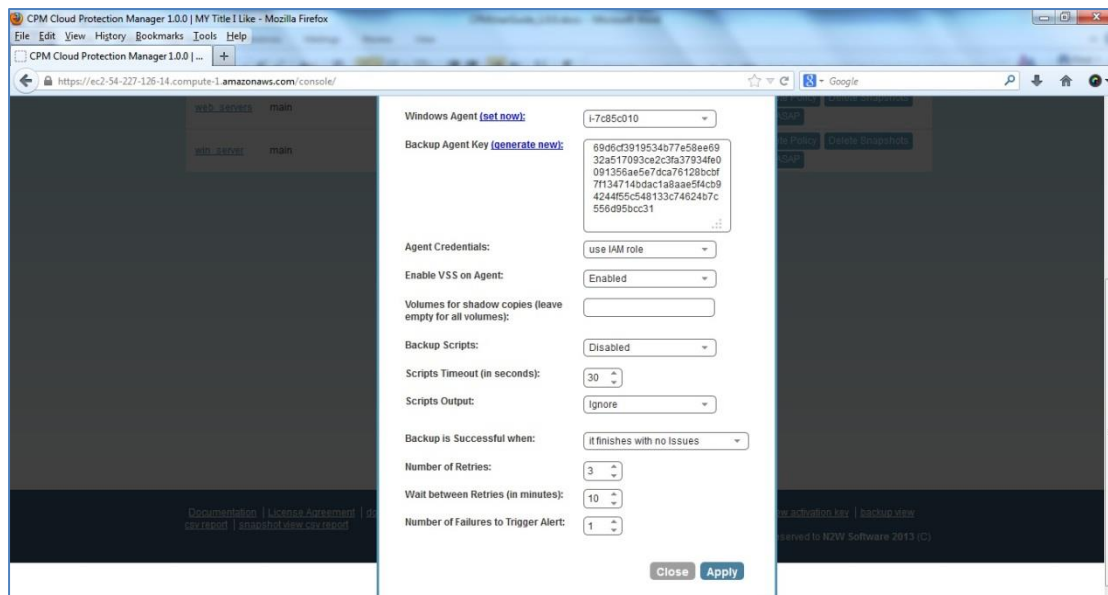


Figure 4-4

5 Introduction to Consistent Backup

This guide by no means claims to be a comprehensive guide on how to create consistent backups. It simply 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.” This means that 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 machine that had its power cord pulled out. It is 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. They have mechanisms that are intended to give answers to situations exactly like a system crash or a power outage. 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. Modern file systems are also built to deal with crashes. NTFS in Windows and EXT3 in Linux had implemented journaling, which is not unlike transaction logs in databases.

5.2 Application-consistent backup

In application-consistent backups, the application is informed it is about to be backed up. It then has time to perform operations to make sure the actual data on disk is consistent. It can ensure it’s not in the middle of a transaction. It can flush certain data that is cached in memory to the disk. It can also freeze any further activity, until notified by the backup system that it is okay to do so. Such a state is also referred to as “backup mode.” Despite what is written in the previous section, backing up an application in such a way is probably safer than depending on the robustness of the backed-up application.

There is also one more function that application-consistent backups perform especially for databases. As we said before, databases keep transaction logs. These transaction logs take storage space, so occasionally they need to be deleted. Such an operation is called “log truncation.” When can transaction logs be deleted without impairing the robustness of the database? The answer is - probably after you make sure you have a successful backup of the database. So 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 (probably the most common case), we would like all the volumes to have 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 close succession (typically only split seconds between them). 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”

There is actually no one correct answer here. It 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
- May ensure one exact point in time across multiple volumes/disks
- May automatically truncate database transaction logs

Cons:

- Requires 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
- Backup takes more resources and time

6 Windows Instances Backup

6.1 Introduction

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 later recovering instances.

In Unix/Linux instances we can back up system volumes (root devices), and later launch instances based on the snapshot of the system volume. We can simply 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, there is no way to create a launchable image (AMI) from that snapshot. The only way to create such an image for Windows Servers is to create an EBS-based image directly (in AWS Management Console you do this by clicking “Create Image (EBS AMI)”). Creating an image this way, requires rebooting, or its point of time is not guaranteed. You should always keep a recent image of your Windows servers, which can be used to restore the system. Data volumes will be recovered from the snapshots. Later we see there is a workaround for this problem (see 9.2.4).

- A good practice is creating a new image for a Windows instance whenever maintenance is being done on it, be it upgrades, installing service packs, etc. Usually these maintenance operations entail reboots anyway, so it is a good time to create an image. It's a good practice to keep the id of that image (AMI ID) close at hand, in case it is needed for recovery.

Info 6-1

6.2 Configuring CPM Thin Backup Agent

If you decide crash-consistent backup is good enough for your needs, you do not need to install any agent. However, if you wish to use VSS or run backup scripts, you will need to install CPM Thin Backup Agent. A policy can have only one agent associated with it, so if you wish to backup multiple Windows instances and use backup scripts or VSS on them, you will need to create separate policies.

6.2.1 Associating an agent with a policy

After adding your windows instance in the backup targets page (see 4.2.2), click on “More Options” from the “policies” tab (see 4.2.3). You will have a drop-down list labelled “Windows Agent.” Click on the instance id (assuming this is the only Windows instance in the policy, it will be the only one on the list) and then click “set now,” approve, and you will have a few fields added to this popup screen. Assuming this instance is not already backed up by another policy that uses the agent, the “Backup Agent Key” field should be empty. Click on “generate new,” approve, and then the field will contain the authorization code for the agent. That code will be used in the agent’s configuration. Furthermore, if the instance

the agent is on has an IAM role and you want CPM server not to Send Credentials to the agent, you can state that. See 13.4.3 for configuring IAM.

6.2.2 Installing the agent

You can download the installation package of the agent from the link “download thin backup agent” t the bottom of CPM’s main screen. It will download a standard Windows “msi” package. The agent can be installed on any “Windows 2003,” “Windows 2008” or “Windows 2012” instance, 32 or 64 bit. The installation process is very simple and takes only a few seconds. After approving the license agreement you will reach the configuration screen.

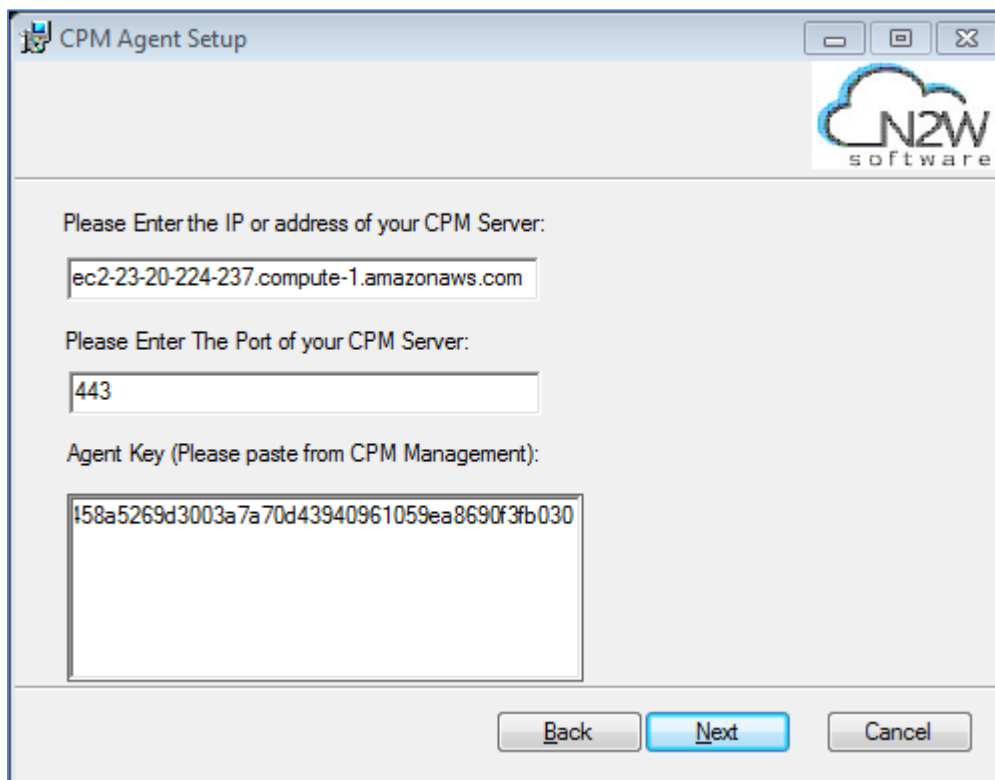


Figure 6-1

The fields are straightforward. You need the address of the CPM server. It is required that the CPM Server will be reachable from this instance. Also the agent will need to reach AWS endpoints to interact with the AWS API. The port is 443 by default (agents communicate with CPM server using the HTTPS protocol). However, if you are using a custom port for your CPM server, you will need to change the port here to the correct value. The agent key is the same key you generated in the previous section, please copy and paste it here.

After finishing the installation, the CPM agent will be a service in your Windows system. It will run automatically and you will not need to deal with it, unless you need to change its configuration.

6.2.3 Changing Agent Configuration

If you need to change the configuration of the backup agent after it has already been installed, you can do so by editing the backup agent configuration file. If the address or port

of the CPM Server had changed, or if you are changing the backup agent key, you need to edit the agent configuration file manually. The file is called “cpmagent.cfg,” (located in the CPM Agent installation folder) and it is in a simple Windows “ini” format. When you switch a value, be careful not to add any unneeded characters or new lines. Just change whatever comes after the equation sign. It is recommended to save the old file separately, in case something goes wrong. After making the required changes, you will need to restart the CPM Backup Agent service for the changes to take effect. Just restart the “CPM Agent Service” in the services management console. Another alternative would be to uninstall and reinstall the agent. This seems drastic, but actually take only a few seconds.

6.2.4 Using the agent with an http proxy

If the Windows instance the agent is installed on connects to the outside (meaning the CPM Server and the AWS endpoints) only through a proxy, CPM agent supports such a configuration. To do so, you will need to edit “cpmagent.cfg” (see previous section) and 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)>
```

If your proxy server requires authentication you add the following two lines as well:

```
proxy_user=<proxy user name>
```

```
proxy_password=<proxy password>
```

For these changes to take effect, you will need to restart the CPM Agent service from the service manager.

6.3 Using VSS

6.3.1 Introduction

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 2008 or 2012 Servers 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 “More Options” screen (see 4.2.3).

In a nutshell, any application that wishes to be “backup aware” has a component called “VSS Writer,” e.g. SQL Server has its VSS writer, the Windows Registry has its VSS Writer, NTFS (the file system) has its own writer, etc... Every vendor who would like to support copying

the actual backup data (or, in other words, 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,” which 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 to be done, 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, not unlike EBS snapshots). When the backup is complete, the writers get notified. They can then do “stuff” knowing that they have a consistent backup for the point in time of the shadow copy. As an example, Microsoft Exchange automatically truncates its transaction logs when it gets notified that a backup is complete.

6.3.2 CPM’s use of VSS

CPM uses the “System Provider” to perform shadow copies. The process is very simple. When CPM fires a backup, the CPM Thin Backup Agent asks the system provider to create shadow copies of all relevant volumes. In this case, differential copies are created, so the complete data of the volume is not copied, but only changes made since the beginning of the backup. You need to have enough space on your volumes to store the shadow copy data, though not a lot (a few hundred MBs should be enough), since CPM will keep the shadow copies for a short while before deleting them. CPM 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 6-2.

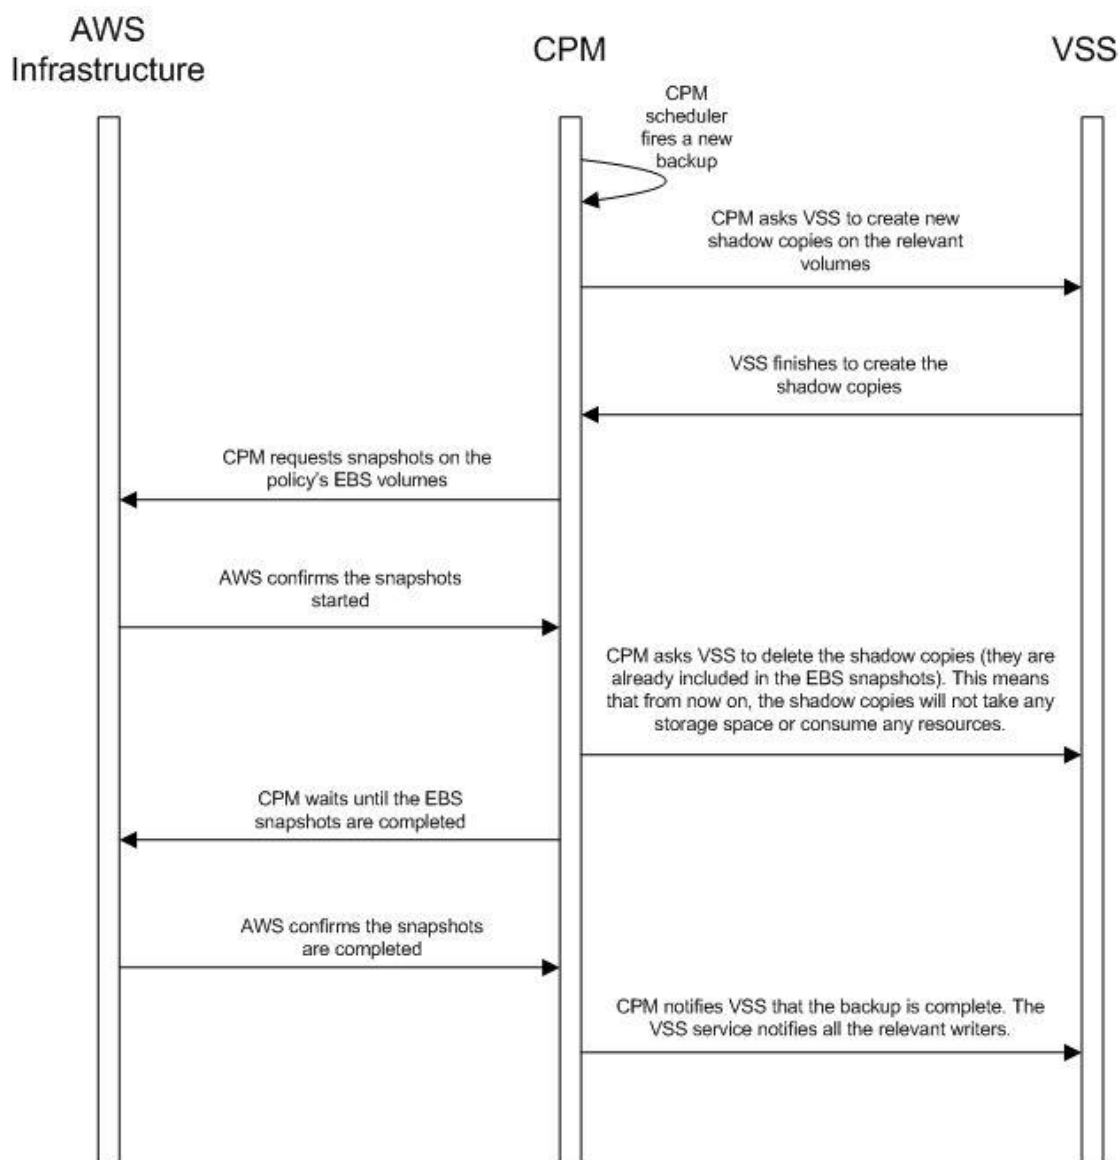


Figure 6-2

6.3.3 Configuring VSS

By default, VSS is enabled when a CPM Thin Backup Agent is associated with a policy. So 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 narrow it down. For example, since the system volume (typically C:\) can't be used to recover the instance in a regular scenario, you may want to exclude it from the backup. In this case there is no use taking a shadow copy of it; this will unnecessarily take up additional resources. To make shadow copies of only some of the volume you use, change the value of "Volumes for shadow copies" in the "More Options" screen. You need to type drive letters followed by a colon, and separate volumes with a comma, e.g. "d:,e:,f:".

6.3.4 Excluding and verifying VSS writers

In some cases you may wish to exclude writers from the backup process. These cases may include a writer that is failing the backup, or consuming too many resources, and is not

essential for the backup's consistency. There is no way in the GUI to configure this. You will need to create a text file and place it in the subfolder "scripts" under the installation folder of the Thin Backup Agent (on the backed-up instance). The file should be named "vss_exclude_writers_< policy name>.txt." The structure of the file is very simple - 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 is probably not what you want, but it can be handy sometimes for testing purposes.

You can also state a list of writers to verify. This means that the writers that you state will have to be included in the shadow copy. If not, the operation will fail. In that case, you will have to create another file at the exact same location, with the exact same format (except that "all" is not viable here). The file should be named "vss_verify_writers_< policy name>.txt."

An example for a line in any of the files:

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

6.3.5 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 state will be "Backup Partially Successful." To see the errors that VSS encountered, look in the backup log. To see the output of the exact VSS error, click on "Recover," and view the VSS Disk Shadow log by clicking its link in the recovery panel.

In most cases, VSS will work out of the box with no issues. There can be a failure from time to time due to stress on the system. If the writers don't answer to the "freeze" request fast enough, the process times out and fail. 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 its documentation: <http://technet.microsoft.com/en-us/library/cc772172%28v=ws.10%29.aspx>. You may see failures in backup because VSS times out or is having issues. You will see that the backup is in 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 itself too often, it may be worth checking that everything is working properly with your Windows server. You can check the application log in Window's Event Log once in a while. If you see VSS errors reported frequently, you should look into it. Contact N2W Software support for any questions.

6.3.6 VSS Recovery

Recovering instances using CPM is covered in chapter 9. When recovering a Windows Server that was backed up with VSS, you have an additional step to perform after recovering the

instance. The volumes from the recovery contain the shadow copies, and you will need to revert back to those shadows to get the consistent state of the data.

After you connect to the newly recovered instance, it's best to stop the services of your application, e.g. SQL Server, Exchange, SharePoint, etc. Next, open an administrator command line console and type "diskshadow." The IDs of the shadow copies made for the required backup can be viewed by clicking the "VSS DiskShadow Data" link in the recovery panel screen. 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, and you can turn the services on and resume work. If you are using the workaround recovery (see 9.2.4), and wish to recover a system disk, that disk can't be reverted to the shadow copy using this method. Generally speaking the system volume should not contain actual application data (it's not a recommended configuration). So you can skip this revert operation. If there is a problem, you can expose the system disk from the shadow and inspect its contents. You do so in the diskshadow utility by typing: expose {shadow id} volletter: You need to remember to unexpose after finishing, and to delete the shadow to avoid unnecessary resources consumption (delete shadow {shadow id}).

6.3.6.1 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, it is possible to do so. Please follow these instructions:

- Before reverting for other volumes, stop the instance; wait until it is in "stopped" state.
- Using the AWS Console, detach the EBS volume of the C: drive from the instance and attach it to another Windows instance, but as an additional disk
- Using the Windows "disk management" utility, make sure the disk is online and exposed with a drive letter.
- Go back to the process in the previous section, and revert to the snapshot of drive C (it will now have a different drive letter). Since it's now not a system volume, it is possible to do so.
- Detach the volume from the second Windows instance, re-attach to the original instance using the original device (typically /dev/sda1), and turn the recovered instance back on.



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. Furthermore, when you revert to shadows you need to do it in the right order. If you revert a volume that contains another volumes shadow data, it will delete it, and you will not be able to revert that other volume. Since this is a recovery operation, you can always start over if you encounter this issue.

Warning 6-1

6.4 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 “More Options” screen of 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. ...\\script\\cpm_user1). All scripts are named with a prefix plus the name of the policy. 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 (for example, if you want to run it with a user who has administrative rights in a domain), you can use the service manager to change the logged-on user to a different one.

All scripts need to exit with the code “0” when they succeed, or “1” (or another non-zero code) when they fail.

There are three scripts for each policy:

6.4.1 “before” script

This script is run 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. This state will stay so for a very short while, until the snapshots of the policy start.

The name of the “before” script is “before_<policy name>.<ext>”

6.4.2 “after” script

This script runs after all the snapshots of the policy start. It runs a very short time after the “before” script, typically between split-seconds and 2-3 seconds. This script should release anything that may have been frozen or locked by the “before” script. This script accepts one

argument when it runs: 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 that this is the opposite of the exit status. Think of it as an argument that is true when the “before” script succeeded.

The name of the “after” script is “after_<policy name>.<ext>”

6.4.3 “complete” script

The “complete” script runs after all snapshots are completed. It is undetermined how long it takes for snapshots to complete. However, usually it’s pretty fast, 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 along the way it will be “0.” The straightforward thing to do will be to truncate logs only if this argument is “1.”

The name of the “complete” script is “complete_<policy name>.<ext>”



When you enable backup scripts, CPM assumes you implemented all three scripts. Any missing script will be interpreted as an error and will reflect in the backup status. Sometimes you do not need all three (the “complete” script is often not needed). In cases like this, you should still write a script that does nothing but exit with the code “0,” and the policy will experience no errors.

Warning 6-2

6.4.4 Capturing the output of backup scripts

You can have the output of backup scripts collected and saved in the CPM Server. Please see 4.2.3.

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 differently (e.g. 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. The only user you can connect with to the CPM server is “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. This gives you a high level of security. 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 editors: vim or nano (simpler to use).

7.2 Backup scripts

7.2.1 General

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 need to 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 they can be executed (and have the correct permissions).



Note that having more than one script with the same base name can result in unexpected behaviour. CPM does not guarantee which script it will run, and even to run the same script every backup. Please avoid such a situation.

Warning 7-1

There are three scripts for each policy:

7.2.2 “before” script

This 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. This state will stay so for a very short while, until the snapshots of the policy are fired. One option is to issue a freeze command to a file system like xfs.

The base name of the “before” script is “before_<policy name>”

7.2.3 “after” script

This script runs after all the snapshots of the policy fire. It runs a very short time after the “before” script, typically between split-seconds and 2-3 seconds. This script should release anything that may have been frozen or locked by the “before” script. This script accepts one argument when it runs: 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 that this is the opposite of the exit status. Think of this as an argument that is true when the “before” script succeeds.

The base name of the “after” script is “after_<policy name>”

7.2.4 “complete” script

The “complete” script runs after all snapshots are completed. It is undetermined how long it will take for snapshots to complete. However, usually it’s pretty fast, 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 along the way it will be “0.” The straightforward thing to do will be to truncate logs only if this argument is “1.”

The base name of the “complete” script is “complete_<policy name>”

7.2.5 Capturing the output of backup scripts

You can have the output of backup scripts collected and saved in the CPM Server, please see 4.2.3.

7.2.6 Troubleshooting and debugging backup scripts

You can use the output collected by CPM to debug backup scripts. However, the recommended way would be to run them independently of CPM, on the CPM Server machine using SSH. You can then easily see their outputs and fix whatever is needed. Once the scripts work correctly, you can start using them with CPM. Assuming these scripts are using ssh, the first time run, you will need to approve the ssh key (by answering “yes”). You do this one time from the command line, and the scripts will run seamlessly from that point

on. Also, if you terminate your CPM Server and start a new one, you will need to run the scripts again from command line and approve the ssh key again.

7.2.7 Example backup scripts

As an example, we can look at 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, and the “complete” script will truncate binary logs older than the backup. Please note, these scripts are given as an example without warranties. Please test and make sure scripts work in your environment and do what you expect them to before actually 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 won't be necessary.

After that script is executed CPM server will start the snapshots, and after that call the next script:

after MyPolicy.bash

```
#!/bin/bash
```

```
if [ $1 -eq 0 ]; then
```

```
    echo "There was an issue running first script" 1>&2
```

```
fi
```

```
ssh -i /cpmdata/scripts/mysshkey.pem sshuser@ec2-host_name.compute-1.amazonaws.com "date +%F %H:%M:%S" > sql_backup_time; mysql -u root -p<MySQL root password> -e 'unlock tables;'"
```

```
if [ $? -gt 0 ]; then
```

```
    echo "Failed running mysql unfreeze" 1>&2
```

```
    exit 1
```

```
else
```

```
    echo "mysql unfreeze succeeded" 1>&2
```

```
fi
```

This script checks the status in the first argument and then does two things: First it saves a timestamp of the current time into a file, then it releases the lock on the MySQL table. This time stamp is at the exact point-in-time of the current backup, since it is taken when the database is frozen. After that, CPM waits for all the snapshots to succeed, and when they do, it run the last 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""'
```

What these two scripts do, is essentially to purge the binary logs, and only if the "complete" script gets "1" as the argument, indicating success. They read the time from the timestamp file and execute the purge command to purge logs earlier than the timestamps.

7.2.8 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. You need to allow users only to access their own scripts folder and not to folders of other users. This can be done by creating 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 scripts of all users. This can typically 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

CPM supports working in a VPC environment. Let's look at a few caveats:

- If the CPM Server is in a VPC, it will need outward access to the Internet (AWS endpoints), for that you will need to either attach an elastic IP to it or enable a NAT configuration. Furthermore 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 a Linux backed up instance is in a VPC and backup scripts are enabled, it will need to be able to get inward connection from the CPM Server.
- If a Windows backed up instance is in a VPC and you need to install a Thin Backup Agent, the agent will need outward connections to both the Internet (AWS endpoints) and 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 a few implications:

- If the policy has backup scripts and they try to connect to the instance, they will fail, and the backup will be in a "Backup Partially Successful" state.
- 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 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 didn't crash, backups will be consistent by definition.

It is recommended that if you are aware of an instance that will be stopped for a while, you disable the policy by clicking on 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 keep the default stricter option that treats script failure as a policy failure. This will make sure older generations of the policy, before it was stopped, will not be deleted.



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 backup to resume.

8.3 Backing up independent volumes

Backing up independent volumes in a policy is done regardless of volumes attachment state. A volume can be attached to any instance or not attached at all, and the policy will still back it up. If this policy is using backup scripts, these can be aware of the volume's state. They can, for instance, check which instance is the active node of a cluster and perform application quiescence through it.

8.4 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.

If you wish 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, including:

- An important backup of an instance we already recovered from. We want to keep the snapshots, to be able to recover the same instance again if needed.
- A backup of interest. We may want to keep the first backup after a major change in the system or after an important update.
- You want to delete a policy and only keep a backup (or a few backups) for future needs.

You move a backup to the freezer by clicking on "Move to Freezer" in the backup monitor tab of the main screen. When you move a backup to the freezer, you type a unique name and an optional description. You can later search and filter frozen backups using keywords from the name or description.

After a backup is in the freezer it will only be deleted if you do so explicitly. The automatic retention management process of CPM will not touch it. Even if you delete the whole policy, frozen backups from the policy will still remain. You can recover from a frozen backup the same way you do from a regular backup.

9 Performing Recovery

CPM offers several options for data recovery. Since all backup is based on AWS's snapshot technology, CPM can offer rapid recovery of instances, volumes, and databases. When you click on "Recover" for a backup at a certain hour, you are directed to the recovery panel screen. This screen will include the instances that were backed up with links to recover them, and links to recover independent volumes and databases. It will also include the outputs of backup scripts and VSS, if they exist. These outputs may be important as reference during a recovery operation.

Also, in the recovery panel screen you may see a drop-down menu to choose whether to perform the recovery in the original AWS region or to another region. This choice will be available if this backup includes DR to another region.

If you have cross-account functionality enabled for your CPM license, you may see two other drop-down menus. You will see "Restore to Account" field 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" to choose from which account to perform recovery.

All the choices about regions and accounts you make in the recovery panel apply to all recovery operation you initiate from this screen.

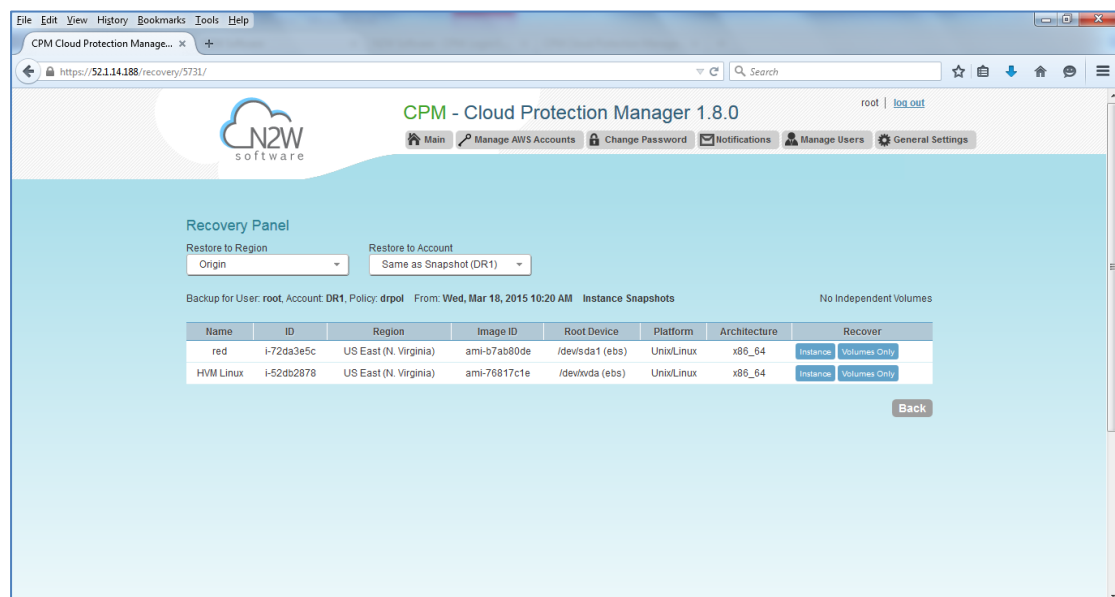


Figure 9-1

We strongly recommend you perform recovery drills from time to time to make sure your recovery scenarios work. It's not recommended to try it for the first time when your servers are down. For any policy you can see on the policy screen, when the last time recovery was performed on it. It can help you track the last time you performed a recovery drill.

9.1 Recovery AWS credentials

All recovery screens will have a checkbox at the bottom labelled “Use account AWS Credentials.” By default it is checked, and it means that for the recovery operation the AWS credentials that were used for backup will be used also for recovery. You can choose to uncheck it, and fill in different credentials for recovery. This can be useful if you choose to use IAM-created credentials for backups that do not have permissions for recovery. Please see 13.4. When using custom credentials CPM verifies these credentials actually belong to the recovery account. If they are not, the recovery operation will fail.

9.2 Instance recovery

This is the most popular option. You recover a complete instance with its data. This type of recovery can be used for many purposes, like:

- An instance crashed or got corrupted – you need to create a new one
- Creating an instance in a different availability zone
- Create an instance in a different region (see 10.6.1)
- Create an instance from a frozen image

When you recover an instance, by default you recover it with all its configuration, tags, and data, as they were at the time of the backup. However, you can change any of these elements if you wish. You can change 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 an image: its original image, or a different one. For instance-store-based or Windows instances, you will only have the image option. This means you can’t use the snapshot of the instance’s root device to launch a new instance. For EBS-based Windows Servers, there is a workaround that allows you to recover a Windows instance’s root device from snapshot (see 9.2.4).

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 them. You can also choose to enlarge their capacity, change their device name or iops value.

You can choose to preserve tags related to the instance and/or data volumes, or you can choose not to.

The instance recovery screen is divided to “Basic Options” and “Advanced Options.” This helps making the recovery process simpler.

9.2.1 Basic options

You can see basic options in **Error! Reference source not found..** The options are:

- “Launch From” – Whether to launch the boot device (image) from a snapshot or use an existing image. The “snapshot” option is available only if this is an EBS-based Linux instance, and a snapshot of the boot device is available in this backup.
- “AMI Handling” – this option is irrelevant (and therefore greyed) unless “Launch From” is set to “snapshot.” If this instance is launched from a snapshot, a new AMI image will be registered. This field handles what to do with this new image:

- “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” – In this case 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, you should move the whole backup to the freezer (see 8.4).
- “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 (for some reason) to launch the instance/s from outside CPM. Again, if you wish to keep using this image, you should to move the backup to the freezer.
- “Image ID” – This is only relevant if “Launch From” is set to “image.” By default, this will contain the original AMI ID from which the backed-up instance was launched. You can type or paste a different AMI ID here, but you can’t search AMIs from within CPM. You can search for it with a different tool (like AWS Management Console).
- “Instances to Launch” – Specifies how many instances to launch from the image. The default is one, and it’s also 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. In the message at the end of the recovery operation, you will see exactly 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 (in Linux instances), or to decrypt the Administrator password (in Windows instances).
- “Instance volumes” – All data volumes (those included in the policy excluding 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 uncheck “recover” to exclude a volume, or change capacity (only to enlarge it), device and iops. You can also decide to exclude any tags associated with the volume (like its name), or whether the volume will be deleted on termination of the instance (for instances recovered from a snapshot).

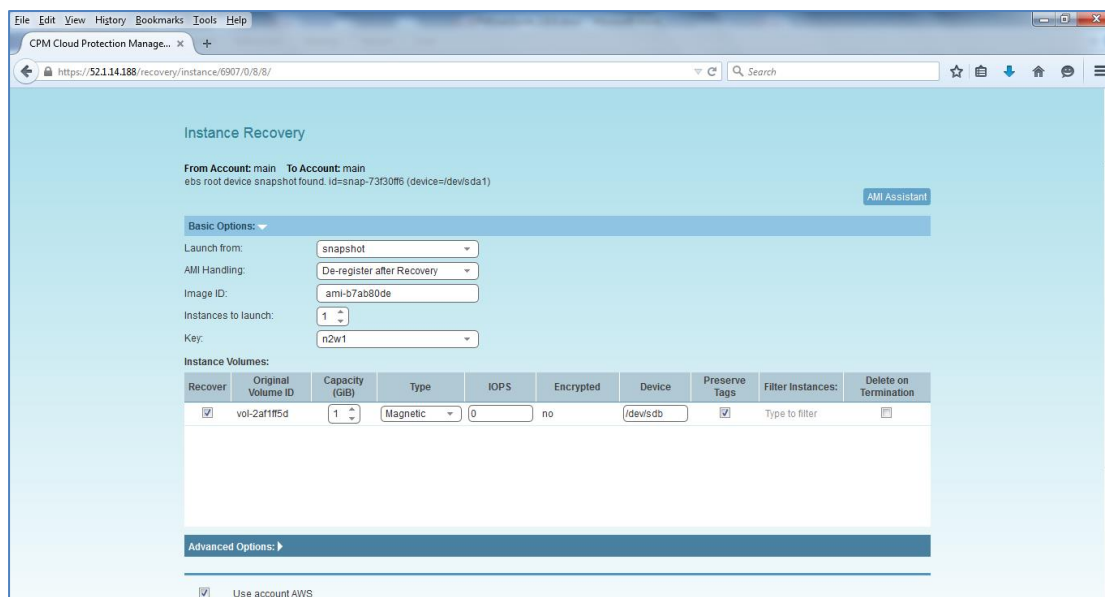


Figure 9-2

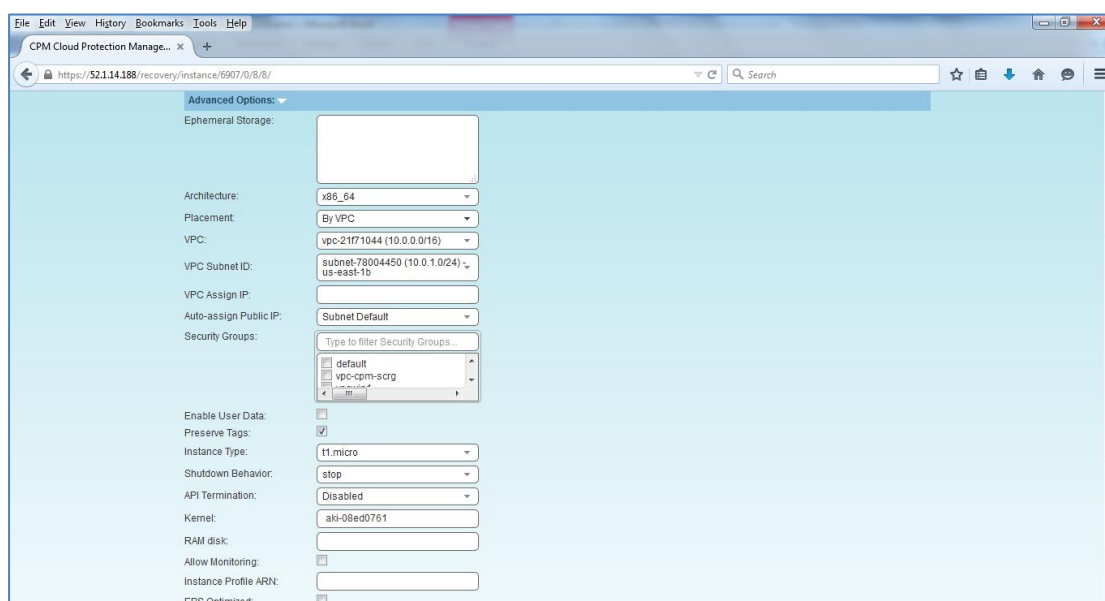
9.2.2 Advanced options

Advanced options include all the rest of the options. You can see them in **Error! Reference source not found..** The options are:

- “Architecture” – You can choose between “i386” – which is X86 - 32 bits, and “x86_64” which is X86 - 64 bits. The default will be the architecture of the backed-up instance. Note that changing this 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” – will determine 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, not all necessarily available.
 - “By Availability Zone” – This is the most basic type. It is the only one which is always available. You can choose in which availability zone to launch the instance.
 - “By VPC Subnet” – This option is only available if you have VPC subnets defined in your account.
 - “By Placement Group” – This option is only available if you have placement groups defined in your account, and this is an instance type that can be placed in a placement group (see AWS documentation for more details).
- “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.
- VPC – This option is only visible if you chose “By VPC Subnet” in “Placement.” 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 option is only visible if you chose “By VPC Subnet” in “Placement.” This will hold all the subnets in the currently selected VPC.
- “VPC Assign IP” – This option is only visible if you chose “By VPC Subnet” in “Placement.” If the backed-up instance was in a VPC subnet, the default value will be the IP assigned to the original instance. Note that if that 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 you leave this field empty, an IP address from the subnet will be automatically allocated for the new instance.
- “Auto-assign Public IP”: Will let you choose whether to assign a public IP to the new instance. This is for public subnets. By default it will behave as the subnet defines.
- “Placement Group” - This option is only visible if you chose “By Placement Group” in “Placement.” You can choose the placement group from the list.
- “Security Groups” – You can choose which security groups will be applied with the new instance. This is a multiple-choice field, which means you can choose more than one. By default, the security groups of the backed-up instance will be chosen. Please note that 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 checked items will not be saved. This field also has a filter to help you find the security group that you need, in case you have many security groups defined.
- “Enable User Data” – States whether to use user data for this instance launch. If checked, 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 this option is checked. If checked, all the tags that were associated with the backed-up instance at the time of the backup (like the instance’s name) will also be associated with the new instance/s.
- “Instance Type” – Choose the instance type of the new instance/s. By default the instance type of the backed-up instance will be chosen. If you choose an instance type that is incompatible with the image or placement method, the recovery operation will fail.
- “Shutdown Behavior” – By default it will have the value of the original instance. 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” – States whether terminating the new instance by API is enabled or not. The default value will be as the backed-up instance.
- “Kernel” – Will hold the kernel id of the backed-up instance. You can type or paste a different one. However, you can’t 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, if it had one. You can type or paste a different one. However, you can’t 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” – Is checked if monitoring should be allowed for the new instance. The default will be the value in the backed-up instance.
- “Instance Profile ARN” – The ARN of the instance role (IAM Role) for the instance. You can find the ARN by clicking on the Role name in IAM Management Console and clicking on the “Summary” tab. The default will be the instance role of the backed-up instance, if it had one.
- “EBS Optimized” – Is checked to launch an EBS Optimized instance. The default will be the value from the backed-up instance.
- “Tenancy” – Lets you choose the tenancy option for this instance.



The screenshot shows the 'Advanced Options' window in the CPM Cloud Protection Manager. The window is titled 'CPM Cloud Protection Manager' and has a search bar. The 'Advanced Options' section is expanded, showing various configuration options for instance recovery. The options are as follows:

Option	Value
Ephemeral Storage	[Empty text box]
Architecture	x86_64
Placement	By VPC
VPC	vpc-21f71044 (10.0.0.0/16)
VPC Subnet ID	subnet-78004450 (10.0.1.0/24) us-east-1b
VPC Assign IP	[Empty text box]
Auto-assign Public IP	Subnet Default
Security Groups	default, vpc-cpm-scrsg
Enable User Data	<input type="checkbox"/>
Preserve Tags	<input checked="" type="checkbox"/>
Instance Type	t1.micro
Shutdown Behavior	stop
API Termination	Disabled
Kernel	aki-08ed0761
RAM disk	[Empty text box]
Allow Monitoring	<input type="checkbox"/>
Instance Profile ARN	[Empty text box]
EBS Optimized	<input type="checkbox"/>

Figure 9-3

To complete the recovery operation, click on “Recover Instance” and then approve. If there are errors 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, and possibly find similar AMIs. CPM will record the details of the AMI when you start backing up the instance. If at that time, the AMI no longer existed, then CPM can't do anything about it. However, if the AMI gets deleted sometime after the instance started backing up, CPM will remember the details of the original AMI.

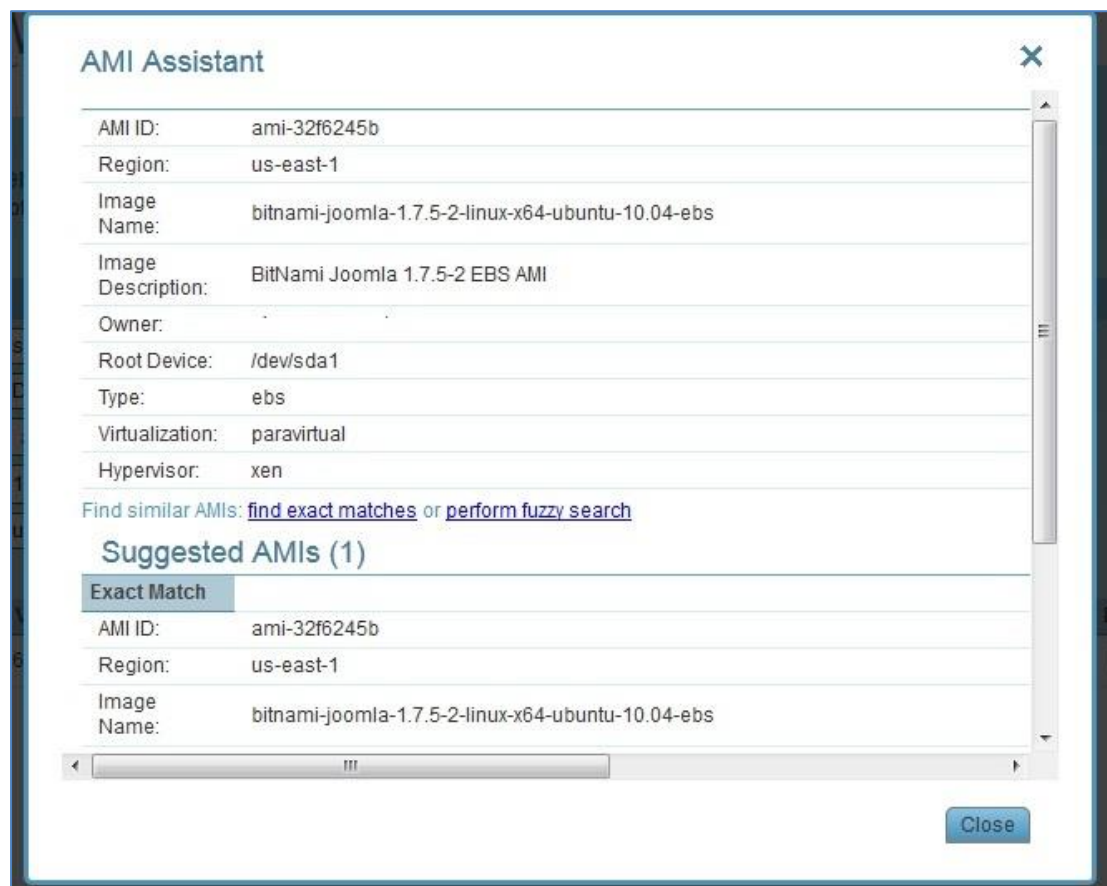


Figure 9-4

When clicking on the “AMI Assistant” button in the instance recovery screen, you will see these details. You will then be able to try and find similar AMIs. Clicking on “find exact matches” will try and find AMIs that according to their properties are exactly like the original. If that doesn't turn up anything, you can click on “perform fuzzy search” which will take a bit longer and will try and find AMIs similar to the original. That will typically turn up several AMI's, usually different versions/flavours of the same offering.

AMI Assistant can be useful for 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 (fuzzy search will help you there).

- You are using DR (see 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.

9.2.4 “Workaround” recovery of a Windows instance

If you wish to recover your Windows instance and use the snapshot of the root device to get the most recent image of it, there is a way to do so. You can recover the instance from image, then stop it and switch volumes. It requires a few simple manual steps:

- Recover the instance using “Instance Recovery” and the “launch from image” option. If recovery fails because the original image is not available, you can choose a pre-created AMI of the original instance, or find a similar one using the AWS Console: simply take a plain vanilla AMI of your Windows OS (e.g. Windows 2008 R2 Base) and paste it’s AMI ID instead of the original one (please see warning below). You can check off all the data in the volumes section, since they will be recovered later anyway.
- After the instance is in “running” state, stop it (in AWS Management Console, just right click on the instance and click “stop”). Wait until instance is in “stopped” state.
- Now, in CPM’s recovery panel,” click on “Volumes Only” for the relevant instance.
- In the volume recovery screen simply choose the correct availability zone of the new instance, and then choose the new instance in the “Attach to Instance” column. In the “Attach Behavior” drop-down list choose “Switch Attached Volumes and Delete Old Ones.” Please be very careful with this option, as it will delete the old instance’s volumes. Make sure you don’t choose the wrong instance by mistake. You can also choose “Switch Attached Volumes” and then delete the old ones manually later on.
- Click on “Recover Volumes” and it will recover the relevant volumes and attach them to the new instance.
- Start the new instance again, the root volume will be the one created from the snapshot (and all other volumes as well).



In some rare cases, instances launched from new AWS AMIs may not boot after switching the volumes. It is recommended to create an image of your original instance in advance, and keep it as a basis for recovery. In any way, do not trust any recovery process until you tested it and made sure it works!

Warning 9-1

9.3 Volume recovery

Volume recovery basically 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 independent volumes. The recovery process is basically the same.

To recover volumes belonging to an instance, simply click on “Volumes Only” next to an instance backup in the recovery panel screen (see **Error! Reference source not found.**).

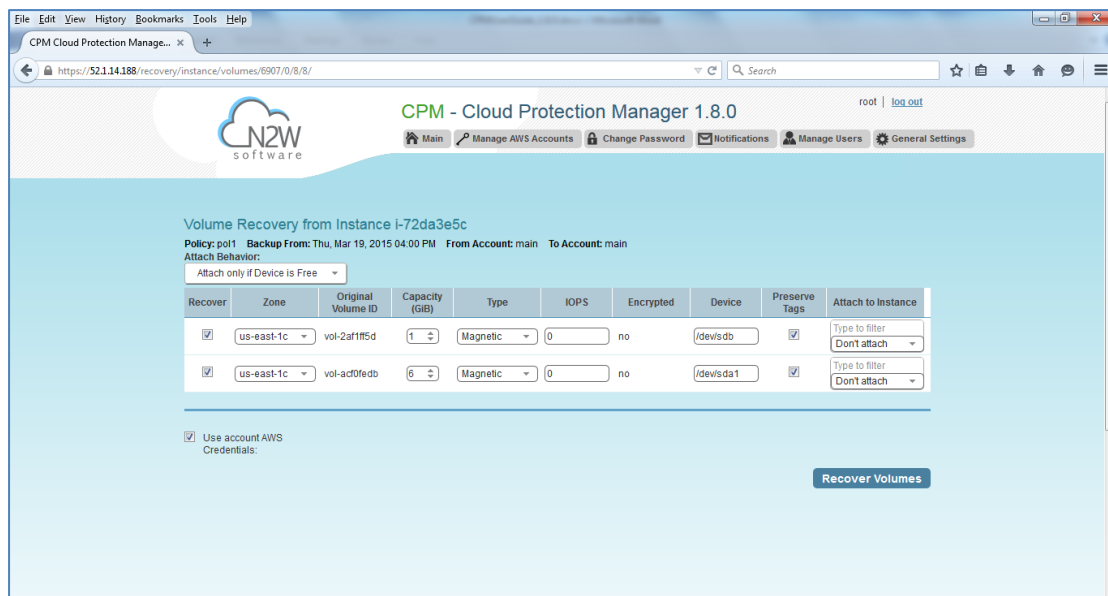


Figure 9-5

As you can see in **Error! Reference source not found.**, the volume recovery screen is straightforward. The following are the fields you can change:

- “Recover” – Checked by default. Uncheck if you don’t want that volume recovered.
- “Zone” – Availability zone. The default is the original zone of the backed-up volume.
- “Capacity” – You can choose to enlarge the capacity of a volume. You can’t make it smaller than the size of the original volume, which is also the default.
- “Type” - Lets you choose the 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. E.g. if you want to create a 100 IOPS volume, its size needs to be at least 10Gib. If you will not abide to this rule, the recovery operation will fail, and you will receive an error message.
- “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 (like the volume’s name) to the recovered volume. Default is yes.
- “Attach to Instance” – Choose whether to attach the newly recovered volume to an instance. The list holds instances that are in the same availability zone as the volume. Changing “Zone” will refresh the content of this list. This field also has a filter, to allow finding the instance easily.
- “Attach Behavior” – This applies to all the volumes you are recovering, if you choose to attach them to an instance. You can choose from these three options:

- “Attach only if Device is Free” – This means that 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” – As the previous option, this one will work only on stopped instances. This option will also delete the old volumes that are detached from the instance.



If you choose “Switch Attached Volumes and Delete Old Ones,” please make sure you don’t need the old volumes. CPM will delete them after detaching them from the target instance.

Warning 9-1

As with other recovery screens, you can choose to use different AWS credentials for the recovery operation. After clicking “Recover Volumes” and approving, 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, you simply click on the “Recover Independent Volumes” button at the top right of the recovery panel screen. This button will only be available if there are independent volumes in the current backup. After clicking, you will reach a similar recover volumes screen as with instance volumes.

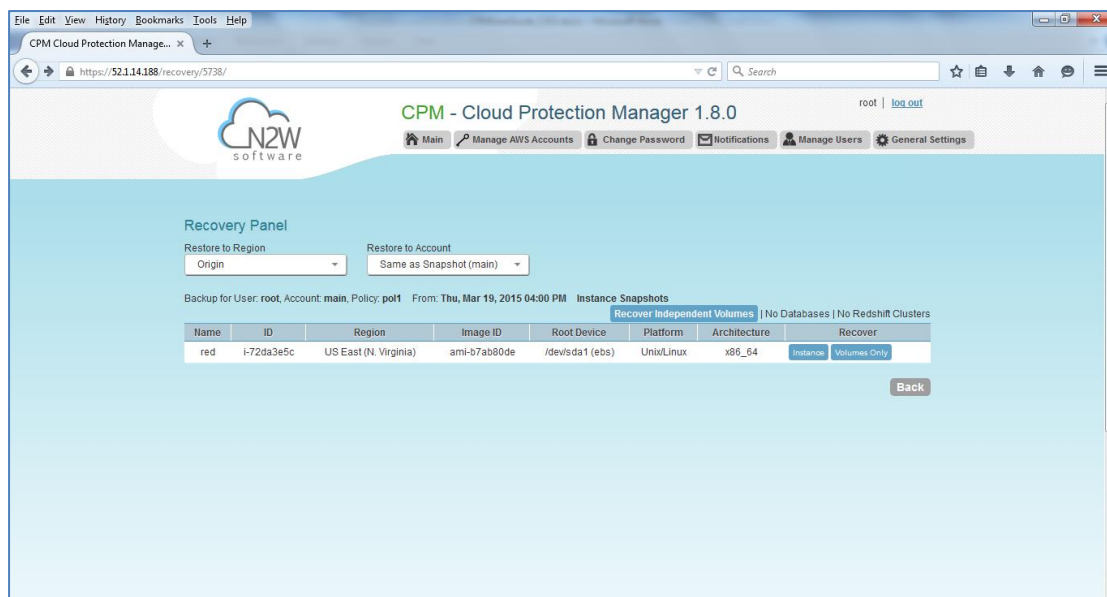


Figure 9-6

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.

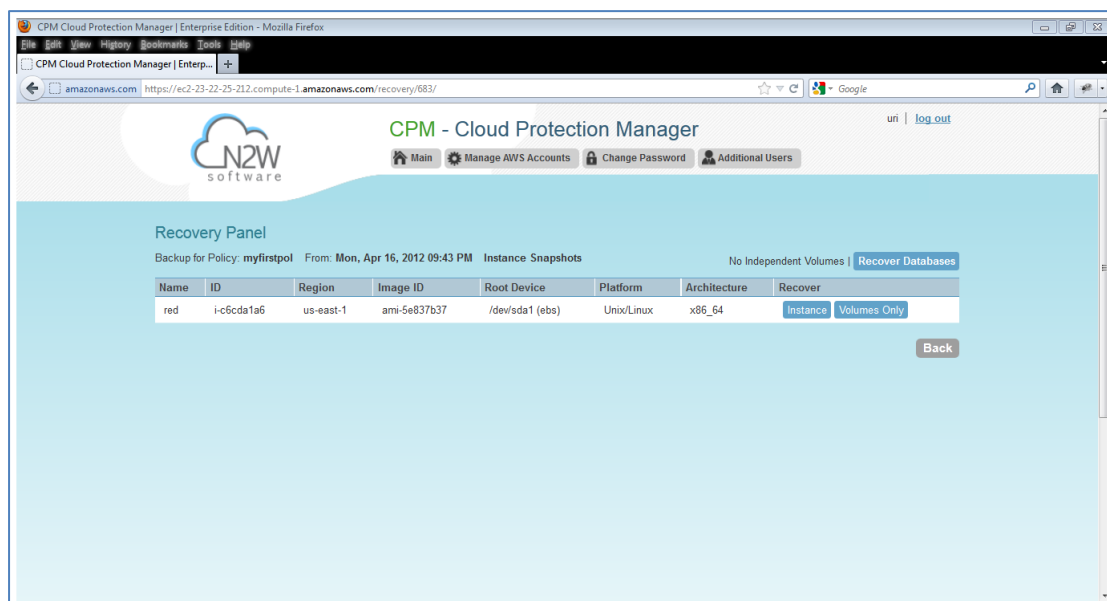


Figure 9-7

Clicking on it will bring you to the RDS Database Recovery screen, as seen in Figure 9-8.

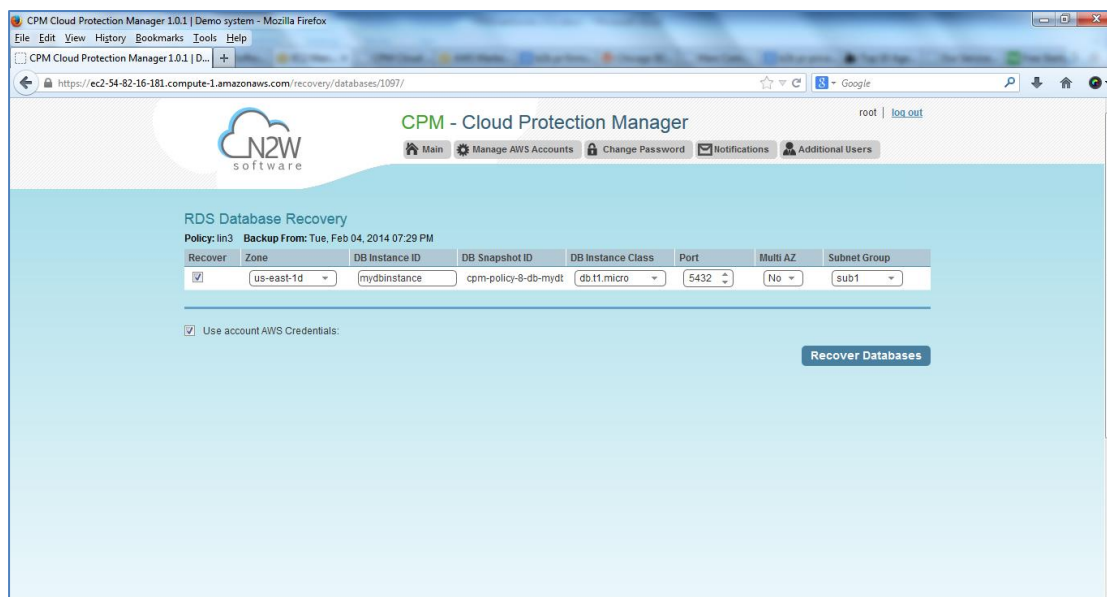


Figure 9-8

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

- “Recover” – Uncheck to not recover the current database.
- “Zone” – The availability zone 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 always recover from the snapshot using AWS Management Console.
- “DB Instance ID” – The default will be the ID of the original database. If the original database still exists, the recovery operation will fail. You can type in a new ID to recover a new database.
- “DB Snapshot ID” – This is just a display field of the snapshot ID.
- “DB Instance Class” – The default is the original class, but you can choose another.
- “Port” – You can choose the port of the database. The default is the port of the original backed-up database.
- “Multi AZ” – Determines whether to launch the database in a multi AZ configuration or not. The default will be the value from the original backed-up database.
- “Subnet Group” – Determines 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 Redshift Cluster Recovery

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

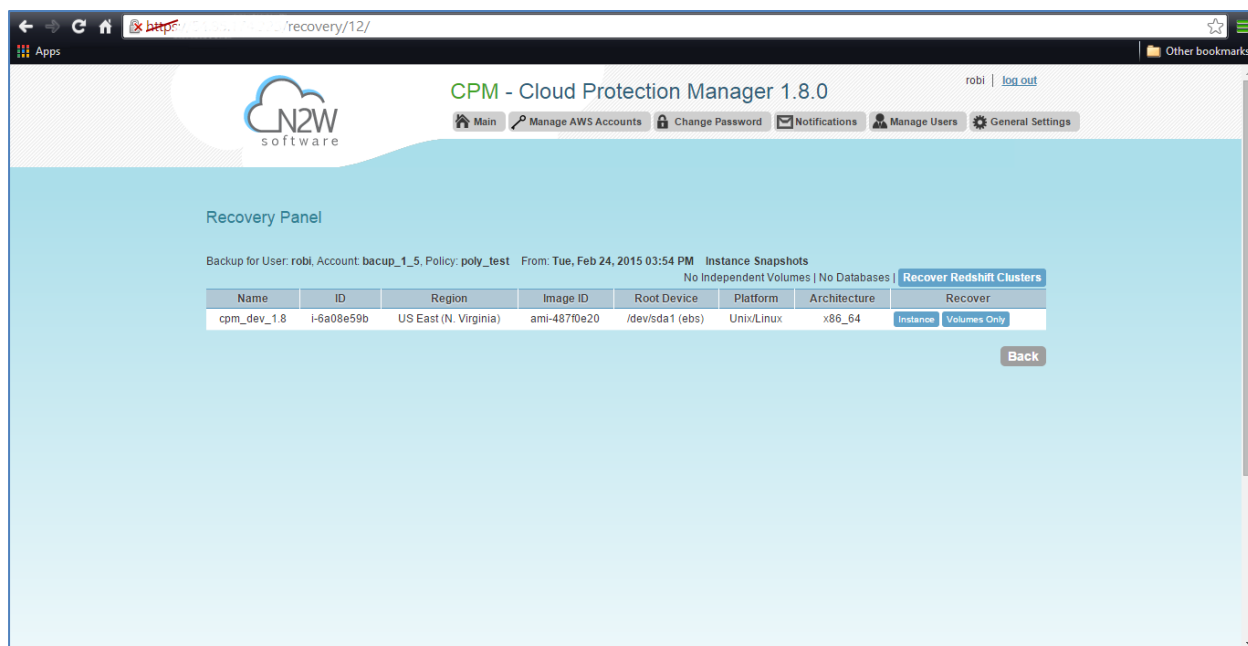


Figure 9-9

Clicking on it will bring you to the Redshift Cluster Recovery screen, as seen in Figure 9-10.

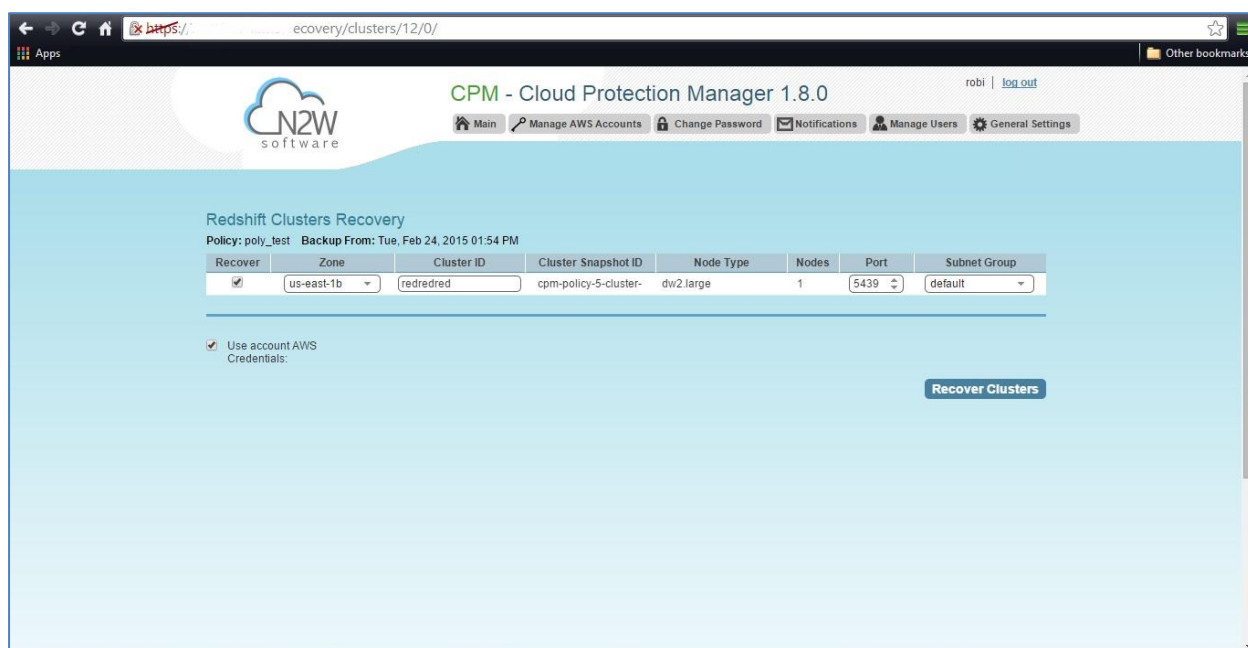


Figure 9-10

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

- “Recover” – Uncheck to not recover the current cluster.
- “Zone” – The availability zone 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. You can type in a new ID to recover a new cluster.
- “Cluster Snapshot ID” – This is just a display field of the snapshot ID.
- “Node Type” and “Nodes” – only for display. Changing these fields are not supported by AWS.
- “Port” – You can choose the port of the cluster. The default is the port of the original backed-up cluster.
- “Subnet Group” – Determines 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 and will return an error.

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

10 Disaster Recovery (DR)

10.1 Introduction

CPM's DR (disaster recovery) solution allows you to recover your data and servers in case of a disaster. A "disaster" doesn't necessarily mean a horrible man-made or natural disaster, although you'll want to be prepared for that as well. DR will also help you recover your data in case of an outage or malfunction, or for any other reason.

What does that mean in a cloud environment like EC2? Every EC2 region is divided into availability zones which use separate infrastructure (power, networking etc...). So, when you use EBS snapshots as CPM does, then by definition you will be able to recover your EC2 servers to other availability zones in case of an outage in one of the zones. 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 for the ability to migrate instances and data between regions and is not limited to the case of an outage or disaster. If you use CPM to take RDS snapshots, those snapshots will also be copied and will be available in other regions.

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 EC2 console.

10.2 Configuring DR

It is very easy setting up DR using CPM. After defining a policy (or any time after a policy started to perform backup), you can click on the "DR" button under the "Configure" column in the "Policies" tab of the main screen. It will then open a very simple popup 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

As you can see, there are only a few options to be configured:

- “Enable DR” – States if DR is enabled for this policy. By default DR is disabled.
- “Perform DR Every...” – States the frequency of performing DR in terms of backups. You may want to copy snapshots of all backups to other regions, which is the default. However, you may want to reduce the frequency for the purpose of reducing costs. Please see 10.5 below for considerations in planning DR.
- “Target Regions” – Here you choose which regions you want to copy the snapshots of the policy to. You can choose one or more.
- “DR Timeout (hours)” – DR copies data between regions over WAN (Wide Area Network), which means it can take a long time. CPM will wait on the copy processes to make sure they are completed successfully. If the DR process is not completed in a certain timeframe, CPM assumes the process is hanging, and will declare it as failed. This value can determine how long CPM will wait for the DR process on the policy. 24 hours is the default and should be enough even for a few 1TiB EBS volumes to copy, however in certain cases you may want to increase it. Also for small volumes, you may want to make it shorter. In most cases you should stay with the default value.

10.3 How it actually Works?

CPM's DR process runs in the background. It starts when the backup is finished. CPM determines then if DR should run and kicks off the process. The actual copying of snapshots can take a long time. CPM will wait until all copy operations are completed successfully before declaring the DR status as "Completed." 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, DR of the next backup had already begun before the first one completed. Although CPM can handle many DR processes in parallel, it is not recommended to take it too far. AWS limits the number of copy operations that can run in parallel to any given region, and too many processes can cause congestion and may never catch up. See 10.5.2 later on this chapter.

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 will the DR snapshots that were copied to other regions.

10.4 DR and mixed-region policies

CPM supports backup up objects from multiple regions in one policy. For most cases it would probably not be the best practice, but sometimes it's useful. When you choose a target region for DR, DR will copy all the objects from the policy to that region, which are not already in this region. For example, if you backup 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.

10.5 Planning your DR Solution

10.5.1 Considerations

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

- Copying between regions is transferring data over a WAN (Wide Area Network). It means that it will be much slower than moving data locally. As you'd expect, 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.

Let's take an extreme example: You have an instance with 4 1TiB 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 100GiB 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.

10.5.2 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's ok to say: For local recovery I want frequent backup, say four times a day, but for DR recovery it's enough for me to have a daily copy of my data. Or maybe it's enough to have it every two days. It's easy to define such a policy using CPM. There are two ways to do it:

- You can use the "Perform DR every..." definition in your policy. If the policy runs four times a day you can ask DR to run once every four backups. The DR status of all the rest will be "Skipped."
- Another method that can give you better control is to define a special policy for the DR process. If you have a "sqlserver1" policy you can define another one and name it something like "sqlserver1_dr." Define all targets and options same as the first policy, but choose a different schedule for it that runs once a day (or whatever you decide) at the time of your choice. Then define DR for the second policy. Locally it will not add any significant cost since it's all incremental, but you will get DR only once a day.

10.5.3 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. It can happen as a result of an outage (or something else). You may want to copy your CPM database across regions as well. As a general note, it's not a bad idea to place your CPM server in a different region than your other production data. It has no problem working across regions and even if you want to perform recovery because of a malfunction in only one of the availability zones 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." Bear in mind that 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 one used to create it in the first place. So you define a new policy and name it "cpmdata" (case doesn't matter) and it will automatically create a policy that backs up the CPM data volume and will do so 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 those remember that in 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. So, even if you want to recover an instance from a week ago, you should always use the latest backup of the “cpmdata” policy.

10.6 DR Recovery

DR recovery is similar to regular recovery with a few differences. First of all, when you click on 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.

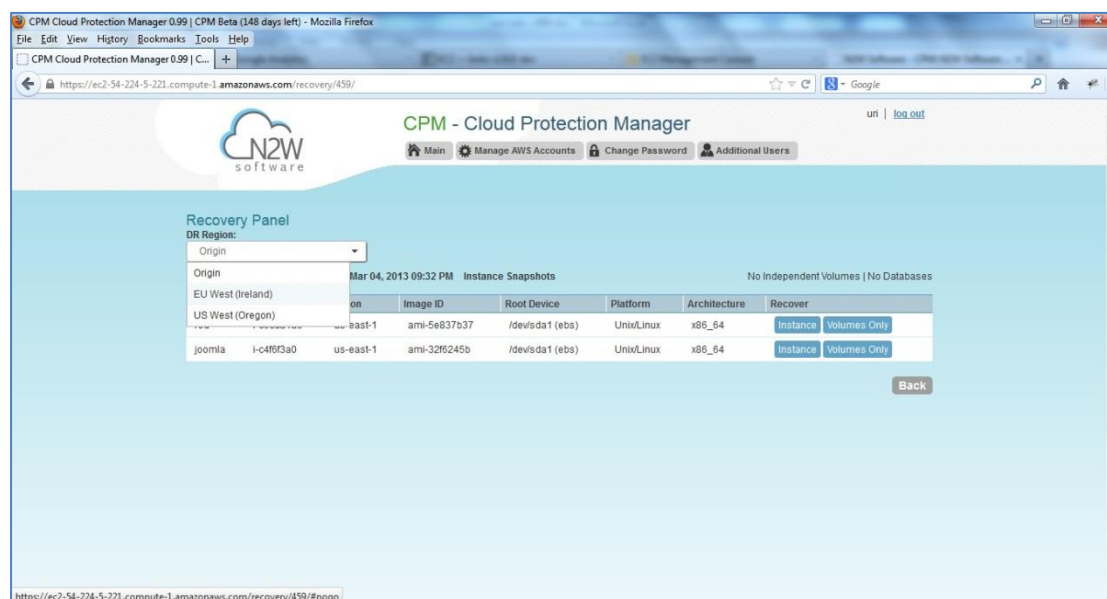


Figure 10-2

As you can see in figure 10-2, the 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. We strongly recommend you perform recovery drills from time to time to be sure your recovery scenario works. It’s not recommended to try it for the first time when your servers are down. For any policy you can see on the policy screen, when the last time DR recovery was performed on it. It can help you track when you performed a recovery drill last.

10.6.1 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: an Image ID (AMI), Key Pair, Security Groups, Kernel ID and Ram disk ID. These objects exist in the region of the original instance; they mean nothing 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 prepare yourself, you should also prepare them at your target regions and make their ids handy when you need to recover. If not, AMI Assistant can help you find a matching image (see 9.2.3).
- **Key Pair:** It's easy to create key pairs with AWS Management Console. You should have one ready so you won't need to create it when you perform recovery.
- **Security Groups:** In regular recovery CPM will remember the security groups of the original instance and use them as default. In DR recovery CPM can't 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 won't 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 9.2.3). When you find such an AMI, you can simply copy and paste its kernel ID from the AMI Assistant window.
- **RAMDisk ID:** Many instances don't 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's a value in the RAMDisk ID field.

10.6.2 A Complete Disaster Recovery Scenario

Let's assume a real disaster recovery scenario (of course with wishes that it never occurs): The region where your operation is in 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's down with all the rest of your instances. Here is what you need to do step by step:

- With AWS Management Console, find the latest snapshot of your "cpmdata" policy. You do that by filtering snapshots with the string "cpmdata" (CPM always adds the policy name to any snapshot's description). Then sort by "Started" in descending order and it's the first one on the list. You create a volume from this snapshot by right-clicking on it and choosing "Create Volume from Snapshot." You can give the new volume a name so it will be easy to find later.
- Launch a new CPM Server at the target region. You can use the "[Your Software](#)" page to launch the AWS Marketplace AMI. Please wait until you see the instance in "running" state.

- As in with regular configuration of a CPM server, you connect to the newly created instance using https. You approve the SSL certificate exception. Assuming the original instance still exists, CPM will come up in “recovery” mode, which means that the new sever will perform recovery and not backup. If you are running the BYOL edition and need an activation key, most likely you do not have a valid key and that time, and you don’t want to wait until you can acquire one from N2W Software. So please register quickly to [CPM Basic Edition](#) (it is the cheapest). In step 2. Use your own username and you can type any password. In step 3, choose the volume you just created for the CPM data volume. Afterwards, complete the configuration.
- Now with a working CPM server you can perform any recovery you need at the target (current) region. You find the backup you want to recover, click on “Recover,” then choose the target region from the dropdown list. If your new server allows backup (it can happen if you registered to a different edition or if the original one is not accessible), please bear in mind that it can start to perform backups. If that’s not what you want, it’s best to disable all policies before you start the recovery process.
- You can recover all the backed up objects that are available in the region.

10.7 DR Monitoring and Troubleshooting

DR is a pretty simple and 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. Every stage and operation during DR is recorded there:

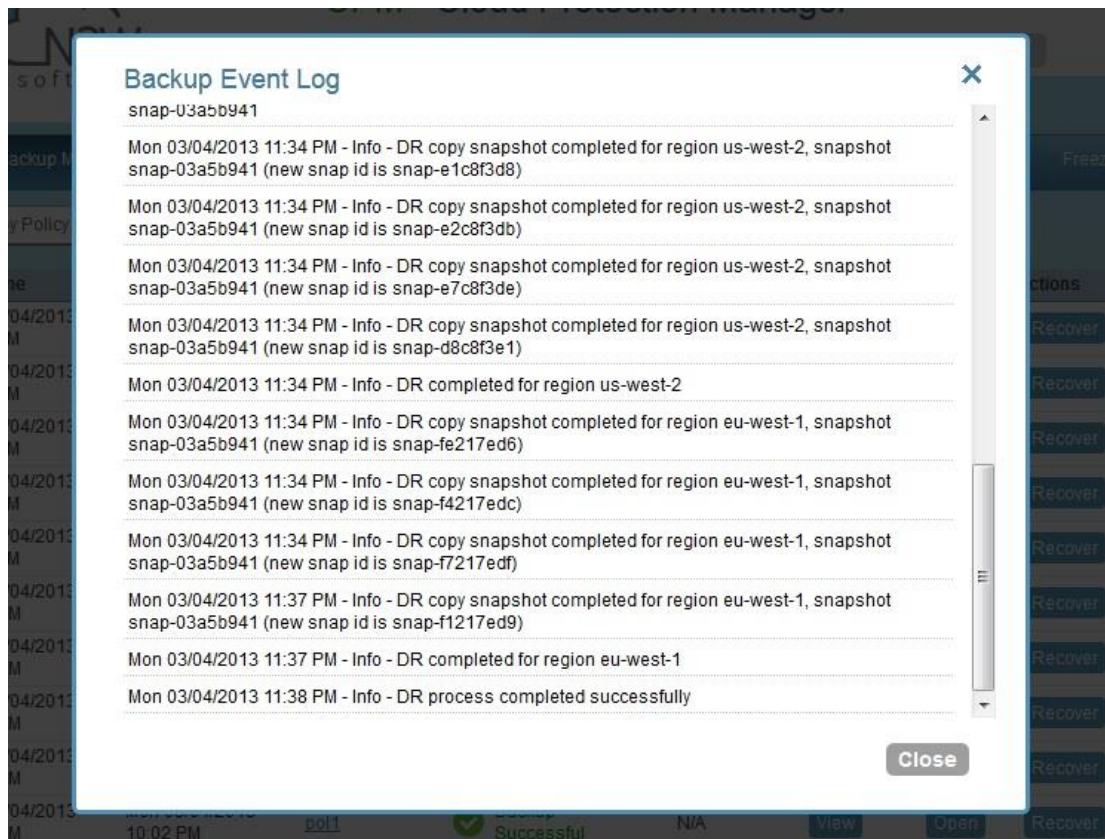


Figure 10-3

Furthermore, you can view DR snapshot ids and statuses in the snapshots screen of the backup:

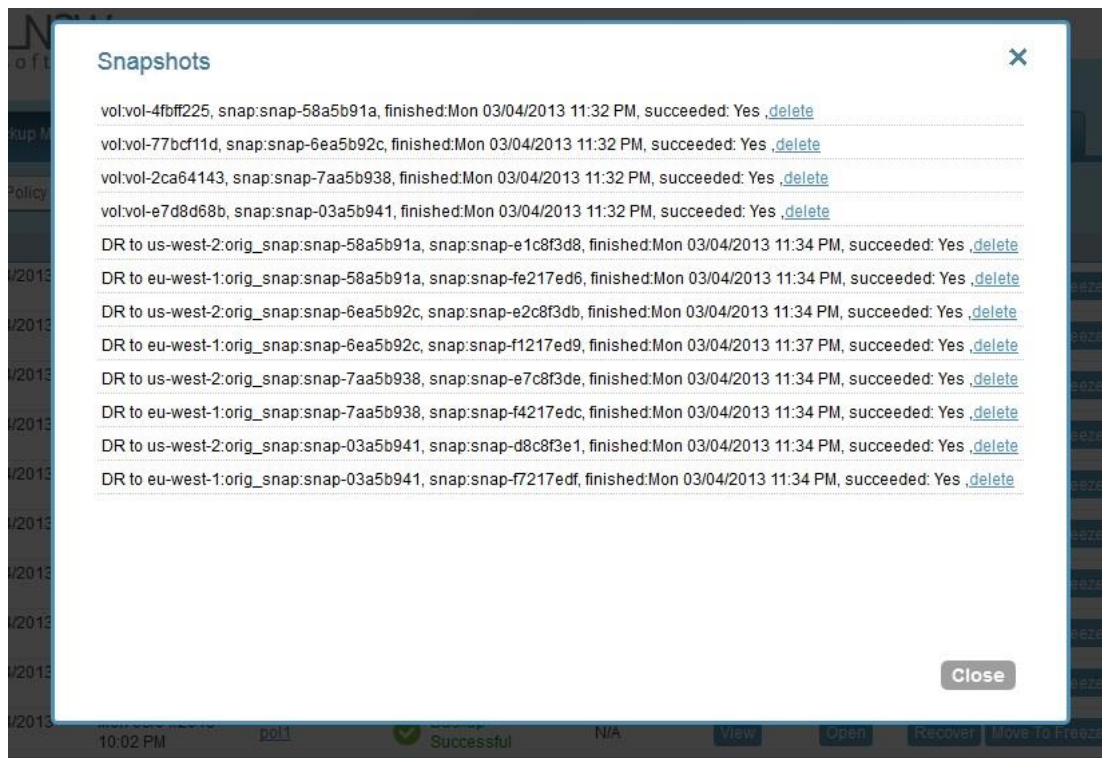


Figure 10-4

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 in case you need them, 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. Please bear in mind, that you can only copy a limited number of snapshots to a given region at one time (currently the number is 5). If it reaches the limit, CPM will wait for copy operations to finish before it continues with more of them. That is no problem, but that can affect the time it takes to complete the DR process.

11 Cross-account DR, Backup and Recovery

11.1 Introduction

Enabled only for Advanced and Enterprise Editions, CPM's cross-account functionality allows you to automatically copy snapshot between AWS accounts. This works as part of the DR module, and in concert with cross-region DR: you can copy snapshots between regions as well as between accounts and any combination of both.

In addition CPM offers cross-account recovery: you can recover resources (e.g. EC2 instances) to a different AWS account whether you copied the snapshots to that account or not.

This cross account functionality is important for security reasons. The ability to copy snapshots between regions can prove crucial if your AWS credentials had been compromised and there's 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 and EBS volumes. It is not supported for RDS databases and Redshift clusters.

Also, cross-account functionality is not enabled for encrypted EBS volumes and instances with encrypted EBS volumes.

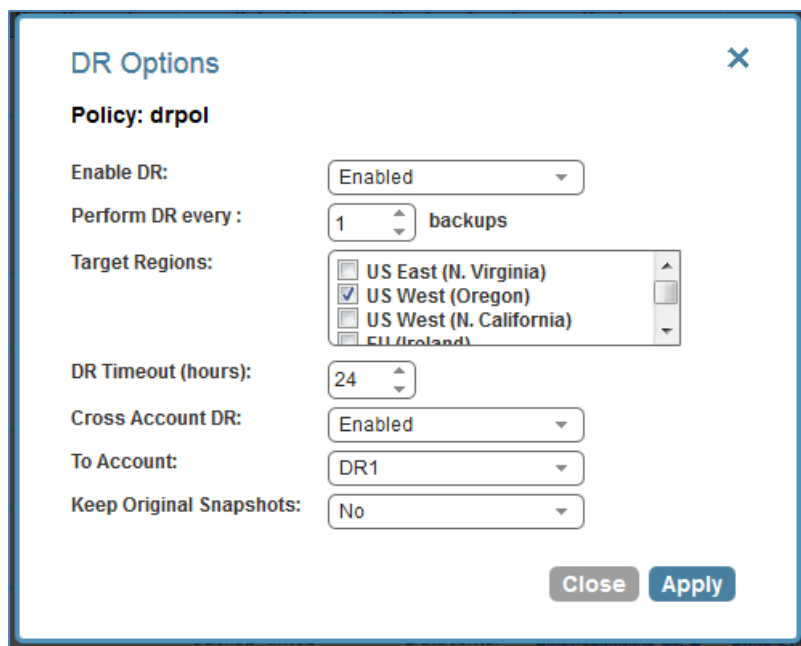
11.2 Snapshot Vaulting

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 keep them in their original account. This will allow to save storage costs (no need to store the snapshot data twice), and utilize the cross-recovery capability to recover resources from the vault account back to the original one.

11.3 Configuring cross-account backup

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



The screenshot shows a 'DR Options' dialog box with a close button (X) in the top right corner. The 'Policy' is set to 'drpol'. The 'Enable DR' dropdown is set to 'Enabled'. The 'Perform DR every' spinner is set to '1' with the unit 'backups'. The 'Target Regions' list includes 'US East (N. Virginia)', 'US West (Oregon)' (which is selected with a checkmark), 'US West (N. California)', and 'EU (Ireland)'. The 'DR Timeout (hours)' spinner is set to '24'. The 'Cross Account DR' dropdown is set to 'Enabled'. The 'To Account' dropdown is set to 'DR1'. The 'Keep Original Snapshots' dropdown is set to 'No'. At the bottom right, there are 'Close' and 'Apply' buttons.

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” – You need to choose to which account to copy the snapshots. 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.

11.4 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 do that for snapshots copied to other accounts as well. However, if you do not wish for CPM to do that, because you want to provide IAM credentials that are limited and can’t 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 add 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 you should make sure these snapshots are actually deleted from time to time. You can either run a script on a schedule, with proper permissions, or make it delete all snapshots with the tag “cpm_deleted”, or do it from time to time in EC2 console, simply filter snapshots by the tag name and delete them.

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

11.6 Cross-account recovery

If you have cross-account functionality enabled in your CPM license, and regardless 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 simply need to choose which account to recover the resource (EC2 instance or EBS volume) 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.

12 Tag-based Backup Management

12.1 Introduction

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 or RDS instance), 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.

12.2 The “cpm backup” tag

To automate backup management for a resource, you can add a tag to that resource named “cpm backup” (lower case). 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, whether to add it to a policy or list of policies, and whether to create a new policy, based on an existing one (template), and then add the resource to it.

12.2.1 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”

12.2.2 Creating a policy from a template

To create a new policy and to add the resource to it, all you need to do is to 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 and 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's 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.

12.2.3 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.

12.3 Tag scanning

Tags 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.

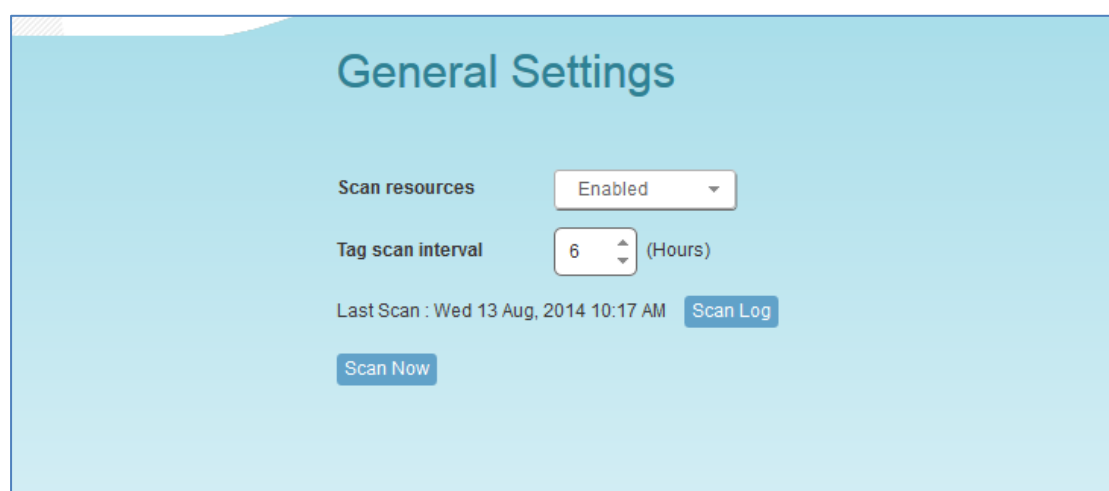


Figure 12-1

Even if scanning is disabled, clicking on "scan now" will initiate a scan.

If you do want automated scans to run, simply 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).

12.4 Pitfalls and troubleshooting

12.4.1 Pitfalls

There are potential issues you should try to avoid when managing your backup via tags, the first being not to create contradictions between the tags content and manual configuration. If you tag a resource and it's 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. Please make sure that if you rename a policy, you will also correct any relevant tag values.

When you open a policy to edit it, and this policy was created by a tag scan, you will be able to see a message at the top of the dialog window: “* This policy was automatically added by tag scan”. Please beware that 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 14.6), any policies with no backup targets will be listed.

Another issue you should try to avoid can happen if the same AWS account is added as multiple accounts in CPM. In that case, the same tags can be scanned multiple times, and the behaviour can become unpredictable. We generally discourage this practice. It is better to define an account once, and then allow other delegates (see 15.4) access to it. In any case, if you added the same AWS account multiple times (even for different users), please make sure only one of the accounts in CPM has “Scan Resources” enabled.

12.4.2 Troubleshooting

Sometimes you need to understand what happened during a tag scan, especially if the tag scan did not behave as expected (e.g. I expected a policy to be created and it didn't). In the “General Settings” screen you can view the log of the last tag scan, and you'll be able to see exactly what happened during this scan, and any problems (e.g. problem parsing the tag value) that were encountered. Furthermore, if the daily summary is enabled, any new scan results from the last day will be listed in the summary.

12.4.2.1 Ensure tag format is correct

- When listing multiple policy names, make sure they are separated by spaces.
- When creating new policy, verify using ‘:’ and not ‘;’. The syntax is “new_policy1:existing_policy1”.
- Use valid name for new policy or it won't be created (error message will be added to scan log).
- Make sure using 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.

13 Security Concerns and Best Practices

13.1 Introduction

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.

13.2 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 correct.
- Every CPM has a unique self-signed SSL certificate. Furthermore, it is 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 support IAM Roles. If the CPM Server instance is assigned an adequate IAM role at launch time, you can save yourself from typing in credentials for that account. In a single account environment you can work with CPM without typing in AWS credentials at all.
- To manage CPM you need to authenticate using a username and password.
- CPM (except in Basic Edition) allows creating multiple users to separately manage the backup of different AWS accounts.

13.3 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.

13.3.1 Credentials rotation

This is actually an AWS practice and not only related to CPM. It is recommended to rotate account credentials from time to time. See

<http://docs.amazonaws.com/AWSSecurityCredentials/1.0/AboutAWSCredentials.html#CredentialRotation>

After changing credentials in AWS you will also need to update them in CPM. You can click on the account name in "Manage AWS Accounts," and modify the access and secret keys.

13.3.2 Passwords

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

13.3.3 Security Groups

Since CPM server is an instance in your account, you can define and configure its security groups. CPM is a secure product, but you can choose to block access from unauthorized addresses. Basically you need HTTPS access (original 443 port or a custom port you decided at configuration time) from any machine which will need to open the management application, as well as from 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.

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 (like from a laptop), you will need to change the settings of the security groups. This can get a bit tedious.

13.4 Using IAM

CPM keeps your AWS credentials safe. However, if you want to minimize risk, you can use "AWS Identity and Access Management" (IAM) 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 and use them as your AWS account credentials in CPM. You will need to create a user using IAM, then attach a user policy to it, and use the policy generator to give the user custom permissions. You can also create IAM roles to be used in the CPM Server (for the account the CPM Server was launched in) and for instances running CPM Agent. For an IAM role that can perform the configuration stage as well as normal operations, you will need to combine 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 do with them. For more information about IAM, see IAM documentation:

<http://aws.amazon.com/documentation/iam/>

13.4.1 CPM Server Configuration Process

AWS credentials in the CPM configuration process are only used for configuring the new server. They are never displayed or saved anywhere. 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 see volumes instances, tags and security groups, to modify security groups, to create EBS volumes, to attach EBS volumes to instances and to create tags. Typically, 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 (in the next section).

13.4.1.1 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"
    }
  ]
}
```

13.4.2 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 need to type in 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 be able to manipulate

snapshots, but for recovery you will need to be able to create volumes, create instances and create RDS databases. Plus, you will need to be able 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 type in ad-hoc credentials when you perform a recovery operation, or else that operation will fail.

13.4.2.1 Minimal IAM policy for backup only

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "ec2:CopySnapshot",
        "ec2:CreateSnapshot",
        "ec2:CreateTags",
        "ec2>DeleteSnapshot",
        "ec2:DescribeAddresses",
        "ec2:DescribeAvailabilityZones",
        "ec2:DescribeImageAttribute",
        "ec2:DescribeImages",
        "ec2:DescribeInstanceAttribute",
        "ec2:DescribeInstanceState",
        "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:ModifySnapshotAttribute",
        "ec2:ResetSnapshotAttribute"
      ],
      "Sid": "Stmt1374236955000",
      "Resource": [
        "*"
      ],
      "Effect": "Allow"
    },
    {
      "Action": [
        "rds:CreateDBSnapshot",
        "rds:CopyDBSnapshot",
        "rds>DeleteDBSnapshot",
        "rds:DescribeDBInstances",
        "rds:DescribeDBParameterGroups",
        "rds:DescribeDBParameters",
        "rds:DescribeDBSnapshots",
        "rds:DescribeDBSubnetGroups",
        "rds:DescribeDBSecurityGroups",
        "rds:ListTagsForResource"
      ],
      "Sid": "Stmt1374237153000",
```



```

    "Resource": [
        "*"
    ],
    "Effect": "Allow"
}
]
}

```

13.4.2.2 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:DetachVolume",
        "ec2:ModifyImageAttribute",
        "ec2:ModifyVolumeAttribute",
        "ec2:RegisterImage",
        "ec2:RunInstances",
        "iam:PassRole"
      ],
      "Sid": "Stmt1374243096000",
      "Resource": [
        "*"
      ],
      "Effect": "Allow"
    },
    {
      "Action": [
        "rds:RestoreDBInstanceFromDBSnapshot"
      ],
      "Sid": "Stmt1374243250000",
      "Resource": [
        "*"
      ],
      "Effect": "Allow"
    }
  ]
}

```

13.4.2.3 IAM Policy for SNS notifications

If you want to you 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 than this, you can define a policy that can only publish, and only to the relevant topics, that is assuming the topics already exist, and subscribed to.

13.4.2.4 Redshift

To add the ability to manage Redshift Cluster snapshots you need to 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.

13.4.2.5 Cross-Account backup & recovery

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

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

13.4.3 CPM Agent IAM Role

If you are using CPM agents in your environment and don't wish CPM Server to actually send credentials to them, you can do so by associating the Windows instance the CPM agent is on with an IAM role (it has to be done at launch time). This IAM role needs even less permissions than CPM Server:

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



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

Info 13-1

13.5 Thin Backup Agent

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. It uses a RESTful API and a special authorization scheme. All communication between the backup agent and CPM server is encrypted. Authorization is done with a backup agent key. The backup agent key is unique and can be changed at any time.

Sensitive information passes between the backup agent and the CPM Server. Therefore, if you have a reason to think someone may have taken hold of the key, please change it right away. To change the key you need to go to the “More Options” screen of the policy (see 4.2.3) and click on “generate new.” Then change the key in the agent’s configuration file (see 6.2.3).

14 Alerts, Notifications and Reporting

14.1 Introduction

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 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.

14.2 Alerts

What are 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'll know this policy is not functioning properly. Later, when the policy succeeds, the alert is turned off or deleted, so you'll 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).

14.3 "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 every once and a while to check if CPM has alerts. When calling this API, the caller gets all 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. Even though this API call is read-only and can't actually change anything in your configuration, 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 user (not from independent users).

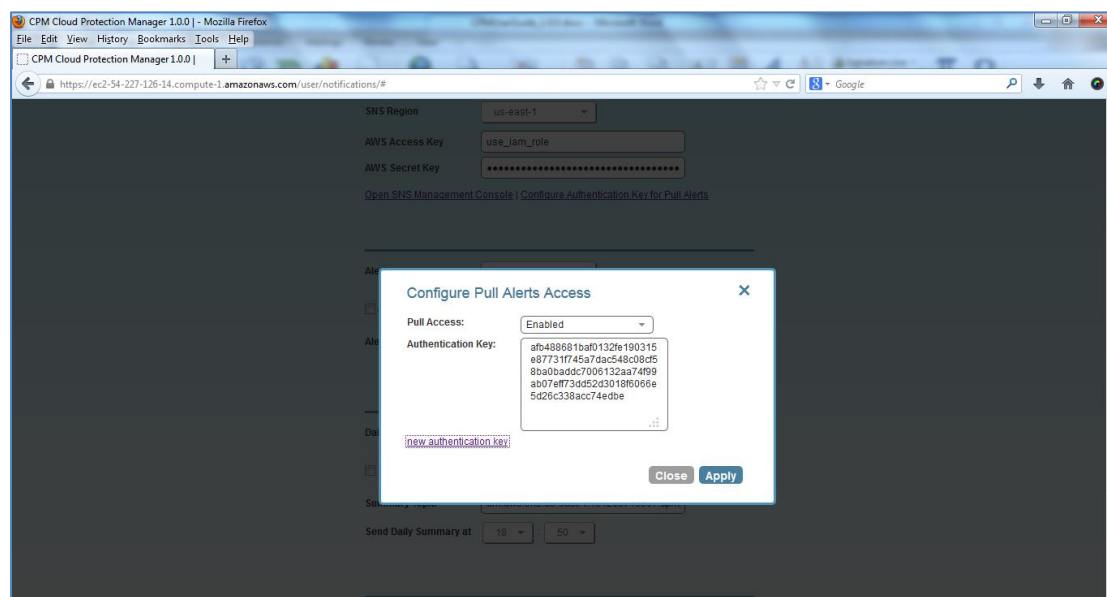


Figure 14-1

In order to configure this API call, you will need to go to the Notifications screen by clicking the “Notifications” button at the top of any screen. In the notifications screen, click on the “Configure API Authentication Key” link. In the popup screen you can enable or disable pull alerts access and you can generate an authentication key by clicking “new authentication key” (see Figure 14-1). Clicking that link while there is already an authentication key assigned, will switch it to a new one. 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/`

Let’s look at a simple python example:

```
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', u'message_body': u'Policy win_server (user:
root, account: main) - backup that started at 07/20/2013 09:00:00 AM
failed. Last successful backup was at 07/20/2013 08:00:00 AM',
u'severity': u'E', u'title': u'Policy win_server Backup Failure',
```

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

The json response is a list of alert objects, each containing the following fields: category, title, message_body, alert_time (time of the last failure) and policy.

14.4 Using SNS

14.4.1 Introduction

CPM can also push alerts to notify you of any malfunction or issue. CPM uses SNS (Simple Notification Service) for this purpose. To use it, your account needs to have SNS enabled. SNS can send push requests to email, http/s, SQS and even SMS (in some locations).

Basically, with SNS you create a topic, and for each topic there can be multiple subscribers, of multiple types (email, email-json, http/s, SQS, SMS). 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 uses SNS in a simple way. It can create the SNS topic for you and subscribe the official user email (the one entered in the configuration process). If you want to add recipients, use SMS, http or other, you can do that using the SNS Management console (part of the AWS Management console). You have a link to this console in CPM's notifications screen.

SNS can incur costs. For the small volume of messages CPM uses, it is usually free or the cost is negligible. For SNS pricing see <https://aws.amazon.com/sns/pricing/>.

14.4.2 Configuring SNS

To Configure CPM for SNS, please click on the "Notification" button on any screen.

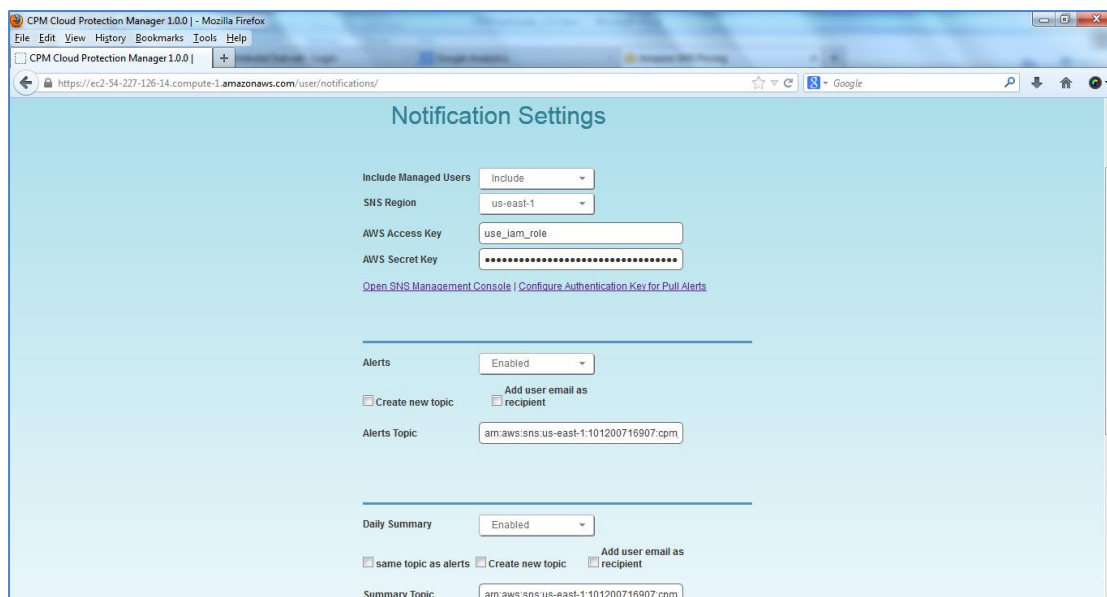


Figure 14-2

You can see the Notifications screen in Figure 14-2. In order 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 type in credentials, which can be regular credentials, IAM user (see 13.4.2.3). To use the CPM Server instance's IAM role (only for the root user), simply type in "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 15.2).

SNS is used both for "push" alerts and for sending a daily summary.

14.5 "Push" Alerts

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

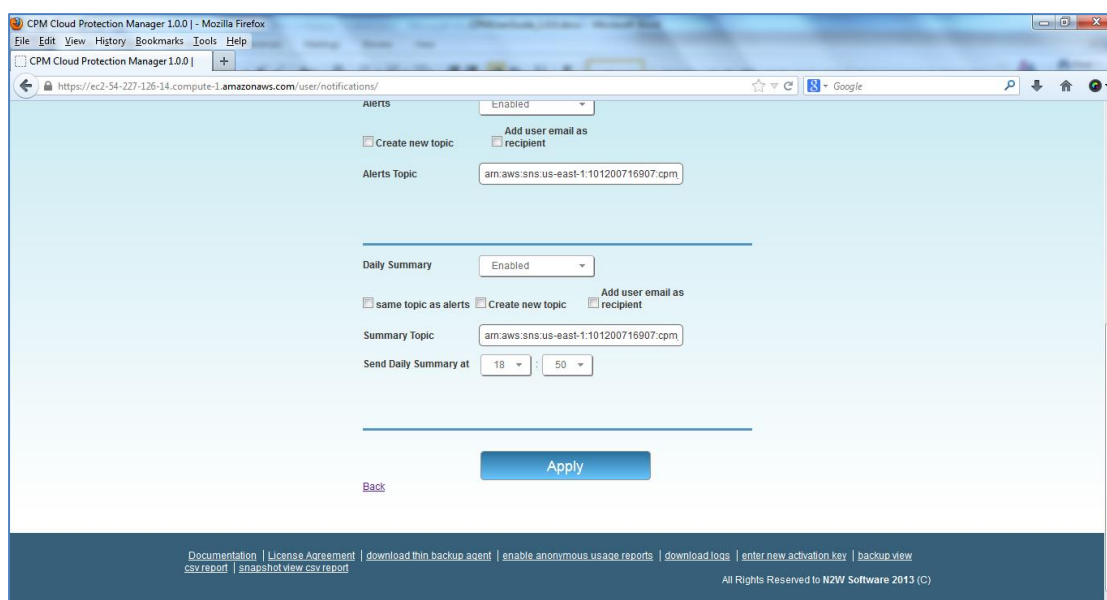
To enable alerts you need to set "Alerts" to "Enabled," then you can either paste in a topic's ARN (Amazon Resource Name) you copied from the SNS Management Console, or request CPM to create the topic for you and add the user's email as a recipient (optional).

After completing this simple step (assuming SNS was configured correctly), Alerts will be sent automatically. Make sure you add all the recipients you need. Each recipient needs to approve the subscription (you get a message for that) before starting to receive alerts.

14.6 Daily Summary

Daily summary is a message (typically an email) that is sent once a day, at an hour of your choosing, summarizing all current alerts in the system. It can be configured instead or in addition to regular alerts. It can be useful for several reasons:

- If you are experiencing issues often (hopefully this won't happen), it sometimes reduces noise to get one summary a day. Furthermore, since backup is the second line of defense, i.e. the production environment does not depend on it, some people feel they don't need to get an instant message on every backup issue that occurs.
- Even if there are no issues (hopefully this is the case most of the time), a daily summary is sent, saying all is ok. This is a positive sign you get from the system once a day. If something happens and CPM crashed altogether (and your monitoring solution did not pick up on that) you will notice daily summaries will stop.
- Daily summary contains a bit more than just alerts; they also contain a list of policies which are disabled and policies that don't have schedules assigned to them. Although neither of these cases is an error, sometimes someone can leave a policy disabled or without a schedule and forget about it, thinking that it continues to perform backup, when actually it does nothing.



The screenshot shows the 'user/notifications/' page of the CPM Cloud Protection Manager 1.0.0. It features two main sections: 'Alerts' and 'Daily Summary'. Both sections have a dropdown menu set to 'Enabled'. The 'Alerts' section includes checkboxes for 'Create new topic' and 'Add user email as recipient', and a text field for 'Alerts Topic' containing 'arn:aws:sns:us-east-1:101200716907:cpm'. The 'Daily Summary' section has similar checkboxes and a 'Summary Topic' field with the same ARN. Below these is a 'Send Daily Summary at' section with dropdowns for '18' and '50'. An 'Apply' button is at the bottom, and a 'Back' link is on the left. The footer contains links for documentation, license agreement, and other resources, along with a copyright notice for N2W Software 2013 (C).

Figure 14-3

As seen in Figure 14-3, you set an SNS topic for the daily summary as well. If you have alerts configured, you can choose to use the same SNS topic for summaries as well, or you can choose to create a new one, or you can paste an ARN of your choosing. 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's easier to know which is which.

Besides that you can choose the hour in which the daily summary will be sent.

14.7 Raw Reporting Data

In the future we plan on adding a full-scale reporting module to CPM. In the meanwhile you can get two raw reports, and you can download them 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 spread sheet or reporting tool. The two reports combined give a complete picture of backups and snapshots taken by CPM. You can download the reports by clicking one of the links: “backup view csv report” or “snapshot view csv report” at the bottom of CPM’s main screen.

14.7.1 Backup view csv report

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

- Backup ID – A unique numerical ID representing the backup
- Account – Name of the AWS account
- 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 – Says yes if this backup was a retry after failure, otherwise says no
- Marked for Deletion – Says yes if this backup was marked for deletion. If it is “yes,” the backup no longer appears in the backup monitor and is not recoverable.

14.7.2 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.
- 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
- End Time – Time the snapshot ended (start time is the start time of the backup)
- Deleted At – Time of deletion, or N/A, if the snapshot was not deleted yet

14.7.3 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; If CPM could not delete all the snapshots in a backup (e.g. a snapshot is included in an AMI and can’t be deleted). If you wish to save records a period time after they were marked for deletion it can be done. Sometimes this is needed for compliance (e.g. GCC).

This can be achieved by creating a file in a specific path on the CPM server. That file must contain only a number. That number is the number of days to save records. Please note that

the number of days is counted since the backup was created and not deleted. So, 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 that file. The path for this file is `/cpmdata/conf/num_days_to_keep_backup_records`. You need to create it and make sure all users have read permissions for it. A typical way to create it is as following:

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

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

Please bear in mind that keeping backups for longer can cause the CPM database to grow and therefore affect the size you need to allocate for CPM's data volume. We estimated that every GiB will accommodate managing the backup of 10 instances. We would say that this estimation is correct when every record is kept for around 30 days. If you want to keep records for 90 days, we would need to triple it, i.e. for 10 instances make the volume 3 GiB , for 20 6 GiB etc...

14.8 Usage Reports

In addition to 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's a link "usage report for current user." For the root user, there's also a link "usage report for all users" which will give all the breakdown of usage between all the users on the CPM server.

15 CPM User Management

CPM is built for a multi-user environment. At the configuration stage you define a user which 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.

To define additional users you need to click on the “Manage Users” button at the top of any CPM screen (available only for the root/admin user).

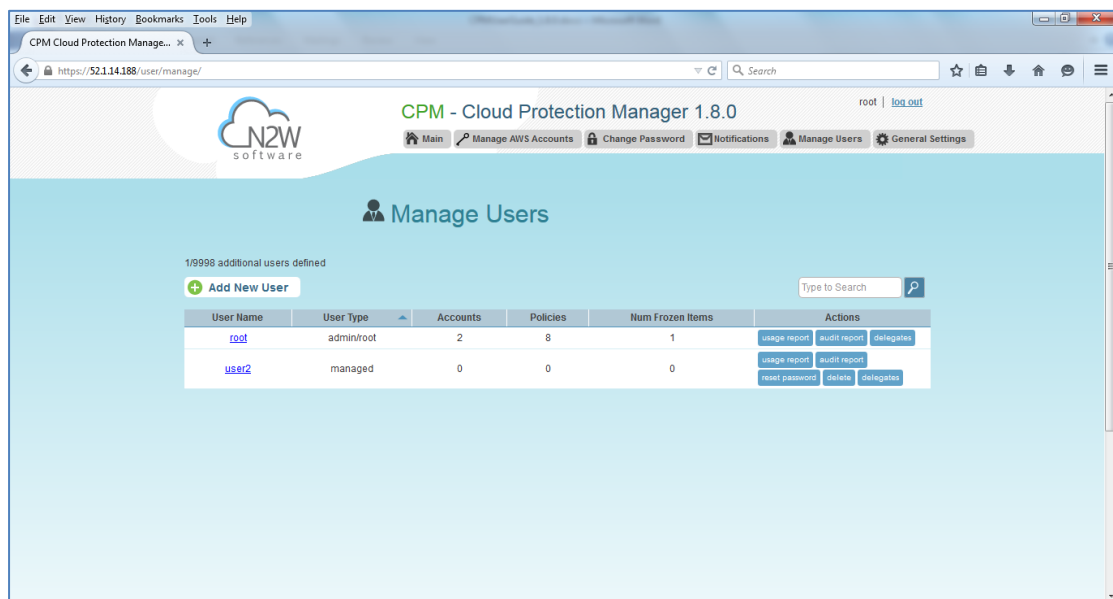


Figure 15-1

There are two types of users you can define (you can also switch types after a user is already created):

15.1 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. Independent users log-in to CPM, create their own accounts and manage their backup. To create an independent user, click on the “Add New User” button in the “Manage Users” screen, and choose the type “Independent.”

15.2 Managed Users

Managed Users are users who can log in 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 on 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, he/she should not supply any credentials to them.

15.3 User definitions

For any type of user, the root user can control several settings:

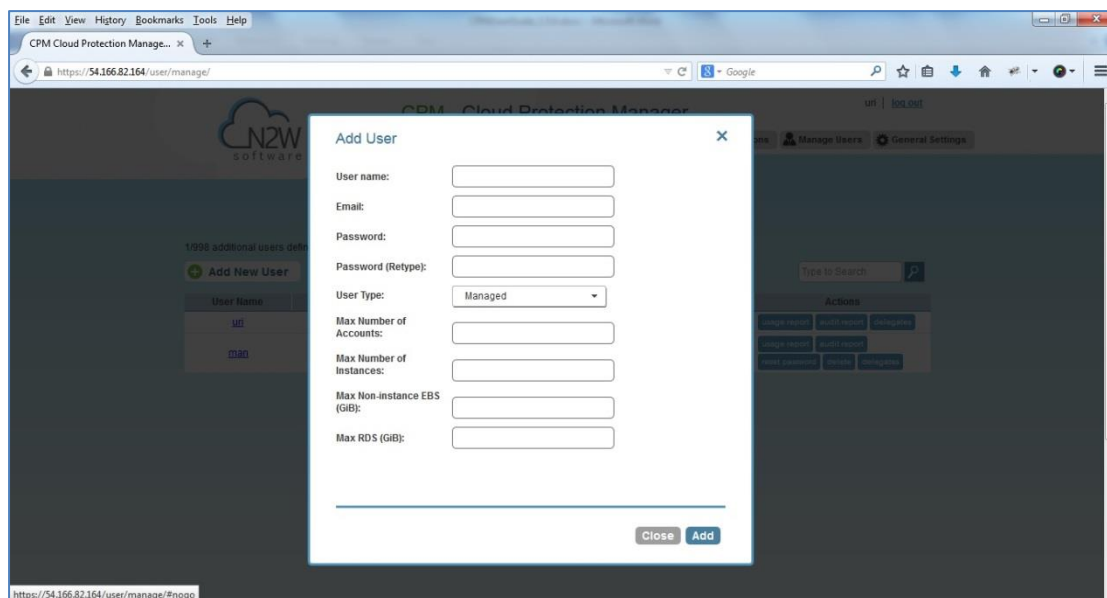


Figure 15-2

Besides the user name, email, password and user type, the root user can set limitations on the amount of resources this user is allowed to have. The limitations include number of accounts, instances and volume of non-instance storage: independent EBS volumes and RDS databases. If you leave these fields empty, there is no limitation on resources, except the system level limitations that are derived from the CPM edition used.

When editing a user, the root user can modify email, type of user and limitations. The user name cannot be modified once a user is created.

15.4 Delegates

Delegates are a special kind of user, which is managed via a separate screen. Delegates are similar to IAM users in AWS, they are credentials used to log in and access another user’s environment. That access is given with specific permissions. For each user, whether it’s the root user, an independent user or a managed user, there is a button “delegates” that redirects to the delegates screen for that user:

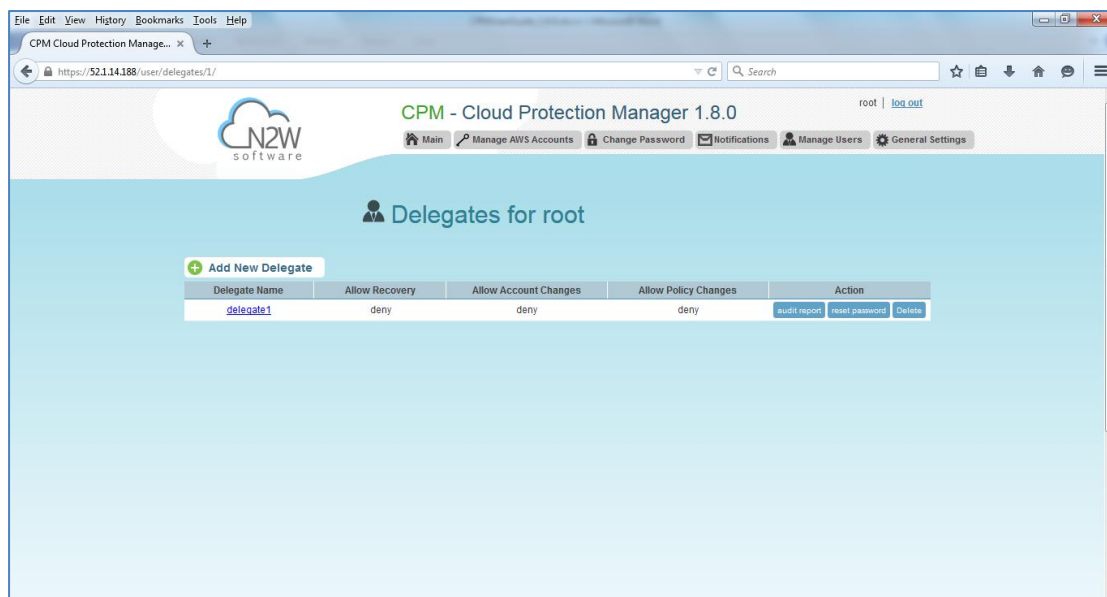


Figure 15-3

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

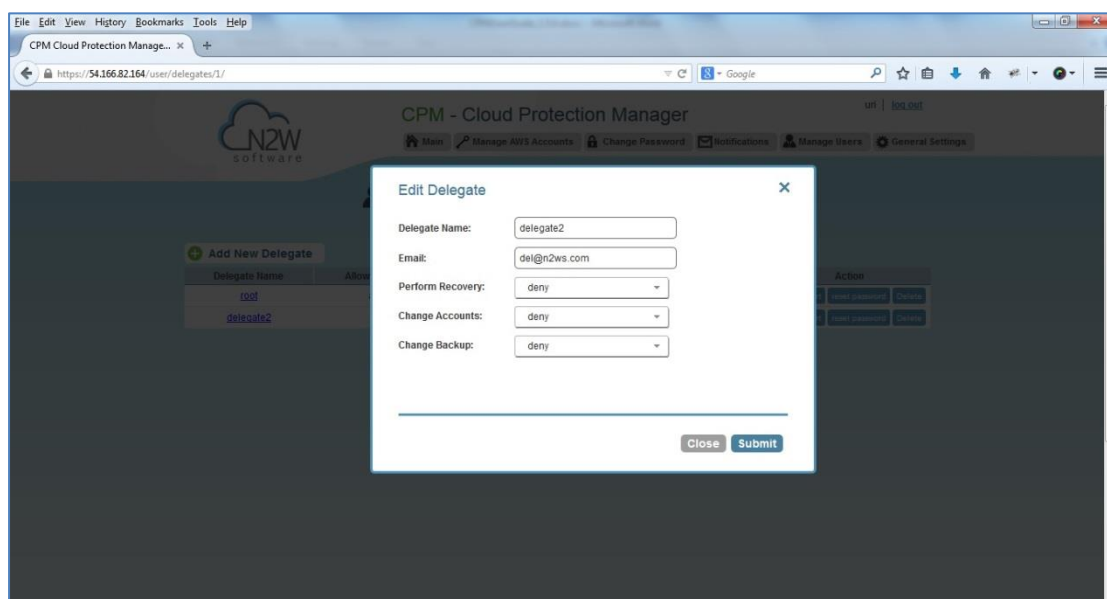


Figure 15-4

The settings include the delegate name (cannot be modified once a delegate is created), email, and permissions. In a separate button in the delegates screen, the root user can reset passwords for delegates.

15.4.1 Delegate permissions

There are three permissions for delegates: "Allow recovery," "Allow Account Changes" and "Allow Policy Changes". As a default they are all denied, which means that the delegate will only have permissions to view the settings and environment and monitor backups. Adding

“Allow Recovery” will enable the delegate to perform recovery operations. Adding “Allow Account Changes” will enable the delegate to add and remove AWS accounts as well as to edit accounts and modify credentials. “Allow Backup Changes” will allow changing policies: adding, removing and editing policies and schedule, as well as adding and removing backup targets.

Allowing all permissions will allow the delegate everything the original user does except notification settings. For delegates of the root/admin user, they won’t be able to change notification settings, general settings or manage users.

15.5 Usage Reports

The root user can also use the user management screen to download CSV usage reports for each user. This can allow accounting and billing per-user. 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.

15.6 Audit Reports

CPM will record every operation done by any user or delegate. 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 this value in the “General Settings” screen. The root user can download audit reports for specific users or delegates (by clicking on “audit report” in the users or delegates screen), or download the audit report for all users (by clicking on the link “audit report for all users” at the bottom of CPM’s main screen).

Each line in the audit report will include a timestamp, event type and a description of the exact operation. For the audit report of all users, there will also be the user who performed the operation, whether he/she is a delegate, and if so, a delegate to which user.