



Protect PostgreSQL

SnapCenter software

NetApp

February 20, 2026

Table of Contents

- Protect PostgreSQL 1
 - SnapCenter Plug-in for PostgreSQL 1
 - SnapCenter Plug-in for PostgreSQL overview 1
 - What you can do using the SnapCenter Plug-in for PostgreSQL 1
 - SnapCenter Plug-in for PostgreSQL features 1
 - Storage types supported by SnapCenter Plug-in for PostgreSQL 2
 - Minimum ONTAP privileges required for PostgreSQL plug-in 3
 - Prepare storage systems for SnapMirror and SnapVault replication for PostgreSQL 5
 - Backup strategy for PostgreSQL 6
 - Restore and recovery strategy for PostgreSQL 8
 - Prepare to install the SnapCenter Plug-in for PostgreSQL 10
 - Installation workflow of SnapCenter Plug-in for PostgreSQL 10
 - Prerequisites to add hosts and install SnapCenter Plug-in for PostgreSQL 10
 - Host requirements to install SnapCenter Plug-ins Package for Windows 13
 - Host requirements for installing the SnapCenter Plug-ins Package for Linux 14
 - Set up credentials for the SnapCenter Plug-in for PostgreSQL 15
 - Configure gMSA on Windows Server 2016 or later 17
 - Install the SnapCenter Plug-in for PostgreSQL 18
 - Configure CA Certificate 24
 - Prepare for data protection 31
 - Prerequisites for using the SnapCenter Plug-in for PostgreSQL 31
 - How resources, resource groups, and policies are used for protecting PostgreSQL 31
 - Back up PostgreSQL resources 32
 - Back up PostgreSQL resources 32
 - Discover the clusters automatically 34
 - Add resources manually to the plug-in host 34
 - Create backup policies for PostgreSQL 36
 - Create resource groups and attach policies 39
 - Create resource groups and enable secondary protection for PostgreSQL resources on ASA r2 systems 43
 - Create a storage system connection and a credential using PowerShell cmdlets for PostgreSQL 45
 - Back up PostgreSQL 47
 - Back up resource groups 52
 - Monitor PostgreSQL backup operations 53
 - Cancel backup operations for PostgreSQL 54
 - View PostgreSQL backups and clones in the Topology page 55
 - Restore PostgreSQL 56
 - Restore workflow 56
 - Restore and recover a manually added resource backup 57
 - Restore and recover an auto discovered cluster backup 61
 - Restore resources using PowerShell cmdlets 63
 - Monitor PostgreSQL restore operations 66
 - Clone PostgreSQL resource backups 67

Clone workflow	67
Clone a PostgreSQL backup	68
Monitor PostgreSQL clone operations	71
Split a clone	72
Delete or split PostgreSQL cluster clones after upgrading SnapCenter	73

Protect PostgreSQL

SnapCenter Plug-in for PostgreSQL

SnapCenter Plug-in for PostgreSQL overview

The SnapCenter Plug-in for PostgreSQL cluster is a host-side component of the NetApp SnapCenter software that enables application-aware data protection management of PostgreSQL clusters. The Plug-in for PostgreSQL cluster automates the backup, restore, and cloning of PostgreSQL clusters in your SnapCenter environment.

SnapCenter supports single cluster and multi cluster PostgreSQL setups. You can use the Plug-in for PostgreSQL Clusters in both Linux and Windows environments. In Windows environments, PostgreSQL will be supported as manual resource.

When the Plug-in for PostgreSQL cluster is installed, you can use SnapCenter with NetApp SnapMirror technology to create mirror copies of backup sets on another volume. You can also use the plug-in with NetApp SnapVault technology to perform disk-to-disk backup replication for standards compliance.

SnapCenter Plug-in for PostgreSQL supports NFS and SAN on ONTAP and Azure NetApp File storage layouts.

VMDK, vVol, and RDM virtual storage layout is supported.

What you can do using the SnapCenter Plug-in for PostgreSQL

When you install the Plug-in for PostgreSQL cluster in your environment, you can use SnapCenter to back up, restore, and clone PostgreSQL clusters and their resources. You can also perform tasks supporting those operations.

- Add clusters.
- Create backups.
- Restore from backups.
- Clone backups.
- Schedule backup operations.
- Monitor backup, restore, and clone operations.
- View reports for backup, restore, and clone operations.

SnapCenter Plug-in for PostgreSQL features

SnapCenter integrates with the plug-in application and with NetApp technologies on the storage system. To work with the Plug-in for PostgreSQL Cluster, you use the SnapCenter graphical user interface.

- **Unified graphical user interface**

The SnapCenter interface provides standardization and consistency across plug-ins and environments. The SnapCenter interface enables you to complete consistent backup, restore, and clone operations

across plug-ins, use centralized reporting, use at-a-glance dashboard views, set up role-based access control (RBAC), and monitor jobs across all plug-ins.

- **Automated central administration**

You can schedule backup operations, configure policy-based backup retention, and perform restore operations. You can also proactively monitor your environment by configuring SnapCenter to send email alerts.

- **Nondisruptive NetApp snapshot copy technology**

SnapCenter uses NetApp snapshot technology with the Plug-in for PostgreSQL cluster to back up resources.

Using the Plug-in for PostgreSQL also offers the following benefits:

- Support for backup, restore, and clone workflows
- RBAC-supported security and centralized role delegation

You can also set the credentials so that the authorized SnapCenter users have application-level permissions.

- Creation of space-efficient and point-in-time copies of resources for testing or data extraction by using NetApp FlexClone technology

A FlexClone license is required on the storage system where you want to create the clone.

- Support for the consistency group (CG) snapshot feature of ONTAP as part of creating backups.
- Capability to run multiple backups simultaneously across multiple resource hosts

In a single operation, snapshots are consolidated when resources in a single host share the same volume.

- Capability to create snapshots using external commands.
- Support for Linux LVM on XFS file system.

Storage types supported by SnapCenter Plug-in for PostgreSQL

SnapCenter supports a wide range of storage types on both physical machines and virtual machines (VMs). You must verify the support for your storage type before installing SnapCenter Plug-in for PostgreSQL.

Machine	Storage type
Physical server	<ul style="list-style-type: none">• FC-connected LUNs• iSCSI-connected LUNs• NFS-connected volumes

Machine	Storage type
VMware ESXi	<ul style="list-style-type: none"> • RDM LUNs connected by an FC or iSCSI ESXi HBAScanning of host bus adapters (HBAs) might take long time to complete because SnapCenter scans all the host bus adaptors present in the host. <p>You can edit the LinuxConfig.pm file located at <i>/opt/NetApp/snapcenter/spl/plugins/scu/scu/core/modules/SCU/Config</i> to set the value of the SCSI_HOSTS_OPTIMIZED_RESCAN parameter to 1 to rescan only those HBAs that are listed in HBA_DRIVER_NAMES.</p> <ul style="list-style-type: none"> • iSCSI LUNs connected directly to the guest system by the iSCSI initiator • VMDKs on NFS datastores • VMDKs on VMFS • NFS volumes connected directly to the guest system • vVol datastores on both NFS and SAN <p>vVol datastore can only be provisioned with ONTAP Tools for VMware vSphere.</p>

Minimum ONTAP privileges required for PostgreSQL plug-in

The minimum ONTAP privileges that are required vary according to the SnapCenter plug-ins you are using for data protection.

- All-access commands: Minimum privileges required for ONTAP 9.12.1 and later
 - event generate-autosupport-log
 - job history show
 - job stop
 - lun
 - lun create
 - lun create
 - lun create
 - lun delete
 - lun igroup add
 - lun igroup create
 - lun igroup delete
 - lun igroup rename

- lun igroup rename
- lun igroup show
- lun mapping add-reporting-nodes
- lun mapping create
- lun mapping delete
- lun mapping remove-reporting-nodes
- lun mapping show
- lun modify
- lun move-in-volume
- lun offline
- lun online
- lun persistent-reservation clear
- lun resize
- lun serial
- lun show
- snapmirror policy add-rule
- snapmirror policy modify-rule
- snapmirror policy remove-rule
- snapmirror policy show
- snapmirror restore
- snapmirror show
- snapmirror show-history
- snapmirror update
- snapmirror update-ls-set
- snapmirror list-destinations
- version
- volume clone create
- volume clone show
- volume clone split start
- volume clone split stop
- volume create
- volume destroy
- volume file clone create
- volume file show-disk-usage
- volume offline
- volume online
- volume modify

- volume qtree create
- volume qtree delete
- volume qtree modify
- volume qtree show
- volume restrict
- volume show
- volume snapshot create
- volume snapshot delete
- volume snapshot modify
- volume snapshot modify-snaplock-expiry-time
- volume snapshot rename
- volume snapshot restore
- volume snapshot restore-file
- volume snapshot show
- volume unmount
- vservers cifs
- vservers cifs share create
- vservers cifs share delete
- vservers cifs shadowcopy show
- vservers cifs share show
- vservers cifs show
- vservers export-policy
- vservers export-policy create
- vservers export-policy delete
- vservers export-policy rule create
- vservers export-policy rule show
- vservers export-policy show
- vservers iscsi
- vservers iscsi connection show
- vservers show
- Read-only commands: Minimum privileges required for ONTAP 8.3.0 and later
 - network interface
 - network interface show
 - vservers

Prepare storage systems for SnapMirror and SnapVault replication for PostgreSQL

You can use a SnapCenter plug-in with ONTAP SnapMirror technology to create mirror

copies of backup sets on another volume, and with ONTAP SnapVault technology to perform disk-to-disk backup replication for standards compliance and other governance-related purposes. Before you perform these tasks, you must configure a data-protection relationship between the source and destination volumes and initialize the relationship.

SnapCenter performs the updates to SnapMirror and SnapVault after it completes the Snapshot operation. SnapMirror and SnapVault updates are performed as part of the SnapCenter job; do not create a separate ONTAP schedule.



If you are coming to SnapCenter from a NetApp SnapManager product and are satisfied with the data protection relationships you have configured, you can skip this section.

A data protection relationship replicates data on primary storage (the source volume) to secondary storage (the destination volume). When you initialize the relationship, ONTAP transfers the data blocks referenced on the source volume to the destination volume.



SnapCenter does not support cascade relationships between SnapMirror and SnapVault volumes (**Primary > Mirror > Vault**). You should use fanout relationships.

SnapCenter supports the management of version-flexible SnapMirror relationships. For details about version-flexible SnapMirror relationships and how to set them up, see the [ONTAP documentation](#).

Backup strategy for PostgreSQL

Define a backup strategy for PostgreSQL

Defining a backup strategy before you create your backup jobs helps you to have the backups that you require to successfully restore or clone your resources. Your service-level agreement (SLA), recovery time objective (RTO), and recovery point objective (RPO) largely determine your backup strategy.

About this task

An SLA defines the level of service that is expected and addresses many service-related issues, including the availability and performance of the service. RTO is the time by which a business process must be restored after a disruption in service. RPO defines the strategy for the age of the files that must be recovered from backup storage for regular operations to resume after a failure. SLA, RTO, and RPO contribute to the data protection strategy.

Steps

1. Determine when you should back up your resources.
2. Decide how many backup jobs you require.
3. Decide how to name your backups.
4. Decide whether you want to create a Snapshot copy-based policy to back up application-consistent snapshots of the cluster.
5. Decide whether you want to use NetApp SnapMirror technology for replication or NetApp SnapVault technology for long-term retention.
6. Determine the retention period for the snapshots on the source storage system and the SnapMirror destination.

7. Determine whether you want to run any commands before or after the backup operation and provide a prescript or postscript.

Automatic discovery of resources on Linux host

Resources are PostgreSQL clusters and instances on the Linux host that are managed by SnapCenter. After installing the SnapCenter Plug-in for PostgreSQL plug-in, the PostgreSQL clusters from all the instances on that Linux host are automatically discovered and displayed in the Resources page.

Type of backups supported

Backup type specifies the type of backup that you want to create. SnapCenter supports snapshot copy-based backup type for PostgreSQL clusters.

Snapshot copy based backup

Snapshot copy-based backups leverage NetApp snapshot technology to create online, read-only copies of the volumes on which the PostgreSQL clusters reside.

How SnapCenter Plug-in for PostgreSQL uses consistency group snapshots

You can use the plug-in to create consistency group snapshots for resource groups. A consistency group is a container that can house multiple volumes so that you can manage them as one entity. A consistency group is simultaneous snapshots of multiple volumes, providing consistent copies of a group of volumes.

You can also specify the wait time for the storage controller to consistently group snapshots. The available wait time options are **Urgent**, **Medium**, and **Relaxed**. You can also enable or disable Write Anywhere File Layout (WAFL) sync during consistent group snapshot operation. WAFL sync improves the performance of a consistency group snapshot.

How SnapCenter manages housekeeping of data backups

SnapCenter manages the housekeeping of data backups on the storage system and file system levels.

The snapshots on the primary or secondary storage and their corresponding entries in the PostgreSQL catalog are deleted based on the retention settings.

Considerations for determining backup schedules for PostgreSQL

The most critical factor in determining a backup schedule is the rate of change for the resource. You might back up a heavily used resource every hour, while you might back up a rarely used resource once a day. Other factors include the importance of the resource to your organization, your service-level agreement (SLA) and your recovery point objective (RPO).

Backup schedules have two parts, as follows:

- Backup frequency (how often backups are to be performed)

Backup frequency, also called schedule type for some plug-ins, is part of a policy configuration. For example, you might configure the backup frequency as hourly, daily, weekly, or monthly.

- Backup schedules (exactly when backups are to be performed)

Backup schedules are part of a resource or resource group configuration. For example, if you have a resource group that has a policy configured for weekly backups, you might configure the schedule to back up every Thursday at 10:00 p.m.

Number of backup jobs needed for PostgreSQL

Factors that determine the number of backup jobs that you need include the size of the resource, the number of volumes used, the rate of change of the resource, and your Service Level Agreement (SLA).

Backup naming conventions for Plug-in for PostgreSQL clusters

You can either use the default Snapshot naming convention or use a customized naming convention. The default backup naming convention adds a timestamp to Snapshot names that helps you identify when the copies were created.

The Snapshot uses the following default naming convention:

```
resourcegroupname_hostname_timestamp
```

You should name your backup resource groups logically, as in the following example:

```
dts1_mach1x88_03-12-2015_23.17.26
```

In this example, the syntax elements have the following meanings:

- *dts1* is the resource group name.
- *mach1x88* is the host name.
- *03-12-2015_23.17.26* is the date and timestamp.

Alternatively, you can specify the Snapshot name format while protecting resources or resource groups by selecting **Use custom name format for Snapshot copy**. For example, `customtext_resourcegroup_policy_hostname` or `resourcegroup_hostname`. By default, the time stamp suffix is added to the Snapshot name.

Restore and recovery strategy for PostgreSQL

Define a restore and recovery strategy for PostgreSQL resources

You must define a strategy before you restore and recover your cluster so that you can perform restore and recovery operations successfully.



Only manual recovery of cluster is supported.

Steps

1. Determine the restore strategies supported for manually added PostgreSQL resources
2. Determine the restore strategies supported for auto discovered PostgreSQL clusters
3. Decide the type of recovery operations that you want to perform.

Types of restore strategies supported for manually added PostgreSQL resources

You must define a strategy before you can successfully perform restore operations using SnapCenter.



You cannot recover manually added PostgreSQL resources.

Complete resource restore

- Restores all volumes, qtrees, and LUNs of a resource



If the resource contains volumes or qtrees, the snapshots taken after the snapshot selected for restore on such volumes or qtrees are deleted and cannot be recovered. Also, if any other resource is hosted on the same volumes or qtrees, then that resource is also deleted.

NOTE: Plug-in for PostgreSQL creates a backup_label and tablespace_map in `/<OS_temp_folder>/postgresql_sc_recovery<Restore_JobId>/_` folder to help recover manually .

Type of restore strategy supported for automatically discovered PostgreSQL

You must define a strategy before you can successfully perform restore operations using SnapCenter.

Complete resource restore is the restore strategy supported for automatically discovered PostgreSQL clusters. This restores all the volumes, qtrees, and LUNs of a resource.

Types of restore operations for auto discovered PostgreSQL

SnapCenter Plug-in for PostgreSQL supports Single File SnapRestore, and connect-and-copy restore types for automatically discovered PostgreSQL clusters.

Single File SnapRestore is performed in NFS environments for the following scenarios:

- If only the **Complete Resource** option is selected
- When the backup selected is from a SnapMirror or SnapVault secondary location, and the **Complete Resource** option is selected

Single File SnapRestore is performed in SAN environments for the following scenarios:

- If only the **Complete Resource** option is selected
- When the backup is selected from a SnapMirror or SnapVault secondary location, and the **Complete Resource** option is selected

Types of recovery operations supported for PostgreSQL clusters

SnapCenter enables you to perform different types of recovery operations for PostgreSQL clusters.

- Recover the cluster up to the most recent state
- Recover the cluster up to a specific point in time


You must specify the date and time for recovery.

SnapCenter also provides the No recovery option for PostgreSQL clusters.

Prepare to install the SnapCenter Plug-in for PostgreSQL

Installation workflow of SnapCenter Plug-in for PostgreSQL

You should install and set up the SnapCenter Plug-in for PostgreSQL if you want to protect PostgreSQL clusters.



Prerequisites to add hosts and install SnapCenter Plug-in for PostgreSQL

Before you add a host and install the plug-in packages, you must complete all the requirements. SnapCenter Plug-in for PostgreSQL is available in both Windows and Linux environments.

- You must have installed Java 11 on your host.



IBM Java is not supported on Windows and Linux host.

- For Windows, plug-in Creator Service should be running using the “LocalSystem” windows user, which is the default behavior when Plug-in for PostgreSQL is installed as domain administrator.
- When installing a plug-in on a Windows host, if you specify a credential that is not built-in or if the user belongs to a local workgroup user, you must disable UAC on the host. SnapCenter Plug-in for Microsoft

Windows will be deployed by default with the PostgreSQL plug-in on Windows hosts.

- SnapCenter Server should have access to the 8145 or custom port of Plug-in for PostgreSQL host.

Windows hosts

- You must have a domain user with local administrator privileges with local login permissions on the remote host.
- While installing Plug-in for PostgreSQL on a Windows host, SnapCenter Plug-in for Microsoft Windows is installed automatically.
- You must have enabled the password-based SSH connection for the root or non-root user.
- You must have installed Java 11 on your Windows host.

[Download JAVA for All Operating Systems](#)

Linux hosts

- You must have enabled the password-based SSH connection for the root or non-root user.
- You must have installed Java 11 on your Linux host.

[Download JAVA for All Operating Systems](#)

- For PostgreSQL clusters that are running on a Linux host, while installing Plug-in for PostgreSQL, SnapCenter Plug-in for UNIX is installed automatically.
- You should have **bash** as the default shell for plug-in installation.

Supplemental commands

To run a supplemental command on the SnapCenter Plug-in for PostgreSQL, you must include it in the *allowed_commands.config* file.

- Default location on the Windows host: *C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc\allowed_commands.config*
- Default location on the Linux host: */opt/NetApp/snapcenter/scc/etc/allowed_commands.config*

To allow supplemental commands on the plug-in host, open *allowed_commands.config* file in an editor. Enter each command on a separate line and the commands are not case sensitive. Ensure that you specify the fully qualified pathname and enclose the pathname in quotation marks (") if it contains spaces.

For example:

```
command: mount command: umount command: "C:\Program Files\NetApp\SnapCreator commands\sdcli.exe"  
command: myscript.bat
```

If the *allowed_commands.config* file is not present, the commands or script execution will be blocked and the workflow will fail with the following error:

```
"[/mnt/mount -a] execution not allowed. Authorize by adding the command in the file %s on the plugin host."
```

If the command or script is not present in the *allowed_commands.config*, the command or script execution will be blocked and the workflow will fail with the following error:

```
"[/mnt/mount -a] execution not allowed. Authorize by adding the command in the file %s on the plugin host."
```



You should not use a wildcard entry (*) to allow all commands.

Configure sudo privileges for non-root users for Linux host

SnapCenter allows a non-root user to install the SnapCenter Plug-ins Package for Linux and to start the plug-in process. The plug-in processes will be running as an effective non-root user. You should configure sudo privileges for the non-root user to provide access to several paths.

What you will need

- Sudo version 1.8.7 or later.
- If the umask is 0027, ensure that the java folder and all the files inside should have permission of 555. Otherwise plug-in installation might fail.
- For the non-root user, ensure that the name of the non-root user and the user's group should be the same.
- Edit the `/etc/ssh/sshd_config` file to configure the message authentication code algorithms: MACs hmac-sha2-256 and MACs hmac-sha2-512.

Restart the sshd service after updating the configuration file.

Example:

```
#Port 22
#AddressFamily any
#ListenAddress 0.0.0.0
#ListenAddress ::
#Legacy changes
#KexAlgorithms diffie-hellman-group1-sha1
#Ciphers aes128-cbc
#The default requires explicit activation of protocol
Protocol 2
HostKey/etc/ssh/ssh_host_rsa_key
MACs hmac-sha2-256
```

About this task

You should configure sudo privileges for the non-root user to provide access to the following paths:

- `/home/LINUX_USER/.sc_netapp/snapcenter_linux_host_plugin.bin`
- `/custom_location/NetApp/snapcenter/spl/installation/plugins/uninstall`
- `/custom_location/NetApp/snapcenter/spl/bin/spl`

Steps

1. Log in to the Linux host on which you want to install the SnapCenter Plug-ins Package for Linux.
2. Add the following lines to the `/etc/sudoers` file by using the visudo Linux utility.

```

Cmnd_Alias HPPLCMD = sha224:checksum_value== /home/
LINUX_USER/.sc_netapp/snapcenter_linux_host_plugin.bin,
/opt/NetApp/snapcenter/spl/installation/plugins/uninstall,
/opt/NetApp/snapcenter/spl/bin/spl, /opt/NetApp/snapcenter/scc/bin/scc
Cmnd_Alias PRECHECKCMD = sha224:checksum_value== /home/
LINUX_USER/.sc_netapp/Linux_Prechecks.sh
Cmnd_Alias CONFIGCHECKCMD = sha224:checksum_value==
/opt/NetApp/snapcenter/spl/plugins/scu/scucore/configurationcheck/Config
_Check.sh
Cmnd_Alias SCCMD = sha224:checksum_value==
/opt/NetApp/snapcenter/spl/bin/sc_command_executor
Cmnd_Alias SCCMDEXECUTOR =checksum_value==
/opt/NetApp/snapcenter/scc/bin/sccCommandExecutor
LINUX_USER ALL=(ALL) NOPASSWD:SETENV: HPPLCMD, PRECHECKCMD,
CONFIGCHECKCMD, SCCMDEXECUTOR, SCCMD
Defaults: LINUX_USER env_keep += "IATEMPDIR"
Defaults: LINUX_USER env_keep += "JAVA_HOME"
Defaults: LINUX_USER !visiblepw
Defaults: LINUX_USER !requiretty

```

LINUX_USER is the name of the non-root user that you created.

You can obtain the *checksum_value* from the **sc_unix_plugins_checksum.txt** file, which is located at:

- `C:\ProgramData\NetApp\SnapCenter\Package Repository\sc_unix_plugins_checksum.txt` _if SnapCenter Server is installed on Windows host.
- `/opt/NetApp/snapcenter/SnapManagerWeb/Repository/sc_unix_plugins_checksum.txt` _if SnapCenter Server in installed on Linux host.




The example should be used only as a reference for creating your own data.

Host requirements to install SnapCenter Plug-ins Package for Windows

Before you install the SnapCenter Plug-ins Package for Windows, you should be familiar with some basic host system space requirements and sizing requirements.


Item	Requirements
Operating Systems	Microsoft Windows For the latest information about supported versions, see the NetApp Interoperability Matrix Tool .
Minimum RAM for the SnapCenter plug-in on host	1 GB

Item	Requirements
Minimum install and log space for the SnapCenter plug-in on host	5 GB  You should allocate sufficient disk space and monitor the storage consumption by the logs folder. The log space required varies depending on the number of the entities to be protected and the frequency of data protection operations. If there is no sufficient disk space, the logs will not be created for the recently run operations.
Required software packages	<ul style="list-style-type: none"> • ASP.NET Core Runtime 8.0.12 (and all subsequent 8.0.x patches) Hosting Bundle • PowerShell Core 7.4.2 <p>For .NET specific troubleshooting information, see SnapCenter upgrade or install fails for legacy systems that do not have internet connectivity.</p>

Host requirements for installing the SnapCenter Plug-ins Package for Linux

Before you install the SnapCenter Plug-ins Package for Linux, you should be familiar with some basic host system space and sizing requirements.

Item	Requirements
Operating systems	<ul style="list-style-type: none"> • Red Hat Enterprise Linux • SUSE Linux Enterprise Server (SLES) <p>For the latest information about supported versions, see the NetApp Interoperability Matrix Tool.</p>
Minimum RAM for the SnapCenter plug-in on host	1 GB

Item	Requirements
Minimum install and log space for the SnapCenter plug-in on host	2 GB <div style="border-left: 1px solid #ccc; padding-left: 10px; margin-left: 20px;">  <p>You should allocate sufficient disk space and monitor the storage consumption by the logs folder. The log space required varies, depending on the number of the entities to be protected and the frequency of data protection operations. If there is no sufficient disk space, the logs will not be created for the recently run operations.</p> </div>
Required software packages	Java 11 Oracle Java and OpenJDK <p>If you have upgraded JAVA to the latest version, you must ensure that the JAVA_HOME option located at <code>/var/opt/snapcenter/spl/etc/spl.properties</code> is set to the correct JAVA version and the correct path.</p>

Set up credentials for the SnapCenter Plug-in for PostgreSQL

SnapCenter uses credentials to authenticate users for SnapCenter operations. You should create credentials for installing SnapCenter plug-ins and additional credentials for performing data protection operations on clusters or Windows file systems.

About this task

- Linux hosts

You must set up credentials for installing plug-ins on Linux hosts.

You must set up the credentials for the root user or for a non-root user who has sudo privileges to install and start the plug-in process.

Best Practice: Although you are allowed to create credentials for Linux after deploying hosts and installing plug-ins, the best practice is to create credentials after you add SVMs, before you deploy hosts and install plug-ins.

- Windows hosts

You must set up Windows credentials before installing plug-ins.


You must set up the credentials with administrator privileges, including administrator rights on the remote host.

If you set up credentials for individual resource groups and the username does not have full admin privileges, you must assign at least the resource group and backup privileges to the username.

Steps

1. In the left navigation pane, click **Settings**.
2. In the Settings page, click **Credential**.
3. Click **New**.
4. In the Credential page, specify the information required for configuring credentials:

For this field...	Do this...
Credential name	Enter a name for the credentials.
User name	<p>Enter the user name and password that are to be used for authentication.</p> <ul style="list-style-type: none"> • Domain administrator or any member of the administrator group <p>Specify the domain administrator or any member of the administrator group on the system on which you are installing the SnapCenter plug-in. Valid formats for the Username field are:</p> <ul style="list-style-type: none"> ◦ <i>NetBIOS\UserName</i> ◦ <i>Domain FQDN\UserName</i> <ul style="list-style-type: none"> • Local administrator (for workgroups only) <p>For systems that belong to a workgroup, specify the built-in local administrator on the system on which you are installing the SnapCenter plug-in. You can specify a local user account that belongs to the local administrators group if the user account has elevated privileges or the User Access control feature is disabled on the host system. The valid format for the Username field is: <i>UserName</i></p> <p>Do not use double quotes (") or backtick (`) in the passwords. You should not use the less than (<) and exclamation (!) symbols together in passwords. For example, <i>lessthan<!10</i>, <i>lessthan10<!</i>, <i>backtick`12</i>.</p>
Password	Enter the password used for authentication.
Authentication Mode	Select the authentication mode that you want to use.

For this field...	Do this...
Use sudo privileges	<p>Select the Use sudo privileges check box if you are creating credentials for a non-root user.</p> <p> Applicable to Linux users only.</p>

5. Click **OK**.

After you finish setting up the credentials, you might want to assign credential maintenance to a user or group of users in the User and Access page.

Configure gMSA on Windows Server 2016 or later

Windows Server 2016 or later enables you to create a group Managed Service Account (gMSA) that provides automated service account password management from a managed domain account.

Before you begin

- You should have a Windows Server 2016 or later domain controller.
- You should have a Windows Server 2016 or later host, which is a member of the domain.

Steps

1. Create a KDS root key to generate unique passwords for each object in your gMSA.
2. For each domain, run the following command from the Windows domain controller: `Add-KDSRootKey -EffectiveImmediately`
3. Create and configure your gMSA:
 - a. Create a user group account in the following format:

```
domainName\accountName$
```

- b. Add computer objects to the group.
- c. Use the user group you just created to create the gMSA.

For example,

```
New-ADServiceAccount -name <ServiceAccountName> -DNSHostName <fqdn>
-PrincipalsAllowedToRetrieveManagedPassword <group>
-ServicePrincipalNames <SPN1,SPN2,...>
```

- d. Run `Get-ADServiceAccount` command to verify the service account.
4. Configure the gMSA on your hosts:
 - a. Enable the Active Directory module for Windows PowerShell on the host where you want to use the gMSA account.

To do this, run the following command from PowerShell:

```
PS C:\> Get-WindowsFeature AD-Domain-Services

Display Name                               Name                               Install
State
-----
-----

[ ] Active Directory Domain Services      AD-Domain-Services      Available

PS C:\> Install-WindowsFeature AD-DOMAIN-SERVICES

Success Restart Needed Exit Code      Feature Result
-----
True      No                Success      {Active Directory Domain
Services, Active ...
WARNING: Windows automatic updating is not enabled. To ensure that
your newly-installed role or feature is
automatically updated, turn on Windows Update.
```

- b. Restart your host.
 - c. Install the gMSA on your host by running the following command from the PowerShell command prompt: `Install-AdServiceAccount <gMSA>`
 - d. Verify your gMSA account by running the following command: `Test-AdServiceAccount <gMSA>`
5. Assign the administrative privileges to the configured gMSA on the host.
 6. Add the Windows host by specifying the configured gMSA account in the SnapCenter Server.

SnapCenter Server will install the selected plug-ins on the host and the specified gMSA will be used as the service log on account during the plug-in installation.

Install the SnapCenter Plug-in for PostgreSQL

Add hosts and install plug-in packages on remote hosts

You must use the SnapCenter Add Host page to add hosts, and then install the plug-ins packages. The plug-ins are automatically installed on the remote hosts. You can add the host and install plug-in packages for an individual host.

Before you begin

- If the operating system of the SnapCenter Server host is Windows 2019 and the operating system of the plug-in host is Windows 2022, you should perform the following:
 - Upgrade to Windows Server 2019 (OS Build 17763.5936) or later
 - Upgrade to Windows Server 2022 (OS Build 20348.2402) or later

- You must be a user that is assigned to a role that has the plug-in install and uninstall permissions, such as the SnapCenter Admin role.
- When installing a plug-in on a Windows host, if you specify a credential that is not built-in, or if the user belongs to a local workgroup user, you must disable UAC on the host.
- You should ensure that the message queueing service is running.
- The administration documentation contains information about managing hosts.
- If you are using group Managed Service Account (gMSA), you should configure gMSA with administrative privileges.


[Configure group Managed Service Account on Windows Server 2016 or later for PostgreSQL](#)


About this task

- You cannot add a SnapCenter Server as a plug-in host to another SnapCenter Server.

Steps

1. In the left navigation pane, click **Hosts**.
2. Verify that the **Managed Hosts** tab is selected at the top.
3. Click **Add**.
4. In the Hosts page, perform the following actions:


For this field...	Do this...
Host Type	Select the type of host: <ul style="list-style-type: none"> • Windows • Linux <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;">  <p>The Plug-in for PostgreSQL is installed on the PostgreSQL client host, and this host can be on either a Windows system or a Linux system.</p> </div>
Host name	Enter the communication host name. Enter the fully qualified domain name (FQDN) or the IP address of the host. SnapCenter depends on the proper configuration of the DNS. Therefore, the best practice is to enter the FQDN.



For this field...	Do this...
Credentials	<p>Either select the credential name that you created or create new credentials. The credential must have administrative rights on the remote host. For details, see the information about creating credentials.</p> <p>You can view details about the credentials by positioning your cursor over the credential name that you provided.</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;">  The credentials authentication mode is determined by the host type that you specify in the Add Host wizard. </div>

5. In the Select Plug-ins to Install section, select the plug-ins to install.

While using the REST API to install Plug-in for PostgreSQL, you must pass the version as 3.0. For example, PostgreSQL:3.0

6. (Optional) Click **More Options**.

For this field...	Do this...
Port	<p>Either retain the default port number or specify the port number. The default port number is 8145. If the SnapCenter Server was installed on a custom port, that port number will be displayed as the default port.</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;">  If you manually installed the plug-ins and specified a custom port, you must specify the same port. Otherwise, the operation fails. </div>
Installation Path	<p>The Plug-in for PostgreSQL is installed on the PostgreSQL client host, and this host can be on either a Windows system or a Linux system.</p> <ul style="list-style-type: none"> • For the SnapCenter Plug-ins Package for Windows, the default path is C:\Program Files\NetApp\SnapCenter. Optionally, you can customize the path. • For the SnapCenter Plug-ins Package for Linux, the default path is /opt/NetApp/snapcenter. Optionally, you can customize the path.

For this field...	Do this...
Skip preinstall checks	Select this check box if you already installed the plug-ins manually and you do not want to validate whether the host meets the requirements for installing the plug-in.
Add all hosts in the cluster	Select this check box to add all the cluster nodes.
Use group Managed Service Account (gMSA) to run the plug-in services	<p>For Windows host, select this check box if you want to use group Managed Service Account (gMSA) to run the plug-in services.</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> <p> Provide the gMSA name in the following format: domainName\accountName\$.</p> <p> gMSA will be used as a log on service account only for SnapCenter Plug-in for Windows service.</p> </div>

7. Click **Submit**.

If you have not selected the Skip prechecks checkbox, the host is validated to verify whether the host meets the requirements for installing the plug-in. The disk space, RAM, PowerShell version, .NET version, location (for Windows plug-ins), and Java version (for Linux plug-ins) are validated against the minimum requirements. If the minimum requirements are not met, appropriate error or warning messages are displayed.

If the error is related to disk space or RAM, you can update the web.config file located at C:\Program Files\NetApp\SnapCenter WebApp to modify the default values. If the error is related to other parameters, you must fix the issue.



In an HA setup, if you are updating web.config file, you must update the file on both nodes.

8. If host type is Linux, verify the fingerprint, and then click **Confirm and Submit**.

In a cluster setup, you should verify the fingerprint of each of the nodes in the cluster.



Fingerprint verification is mandatory even if the same host was added earlier to SnapCenter and the fingerprint was confirmed.

9. Monitor the installation progress.

- For Windows plug-in, the install and upgrade logs are located at: *C:\Windows\SnapCenter plugin\Install<JOBID>_*
- For Linux plug-in, the install logs are located at: */var/opt/snapcenter/logs/SnapCenter_Linux_Host_Plug-in_Install<JOBID>.log_* and the upgrade logs are located at: */var/opt/snapcenter/logs/SnapCenter_Linux_Host_Plug-in_Upgrade<JOBID>.log_*

Install SnapCenter Plug-in Packages for Linux or Windows on multiple remote hosts by using cmdlets

You can install the SnapCenter Plug-in Packages for Linux or Windows on multiple hosts simultaneously by using the `Install-SmHostPackage` PowerShell cmdlet.

Before you begin

You must have logged in to SnapCenter as a domain user with local administrator rights on each host on which you want to install the plug-in package.

Steps

1. Launch PowerShell.
2. On the SnapCenter Server host, establish a session using the `Open-SmConnection` cmdlet, and then enter your credentials.
3. Install the plug-in on multiple hosts using the `Install-SmHostPackage` cmdlet and the required parameters.

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running `Get-Help command_name`. Alternatively, you can also refer to the [SnapCenter Software Cmdlet Reference Guide](#).

You can use the `-skipprecheck` option when you have installed the plug-ins manually and do not want to validate whether the host meets the requirements to install the plug-in.

4. Enter your credentials for remote installation.

Install the SnapCenter Plug-in for PostgreSQL on Linux hosts by using the command-line interface

You should install the SnapCenter Plug-in for PostgreSQL cluster by using the SnapCenter user interface (UI). If your environment does not allow remote installation of the plug-in from the SnapCenter UI, you can install the Plug-in for PostgreSQL cluster either in console mode or in silent mode by using the command-line interface (CLI).

Before you begin

- You should install the Plug-in for PostgreSQL cluster on each of the Linux host where the PostgreSQL client resides.
- The Linux host on which you are installing the SnapCenter Plug-in for PostgreSQL cluster must meet the dependent software, cluster, and operating system requirements.

The [Interoperability Matrix Tool \(IMT\)](#) contains the latest information about the supported configurations.

- The SnapCenter Plug-in for PostgreSQL cluster is part of SnapCenter Plug-ins Package for Linux. Before you install SnapCenter Plug-ins Package for Linux, you should have already installed SnapCenter on a Windows host.

Steps

1. Copy the SnapCenter Plug-ins Package for Linux installation file (`snapcenter_linux_host_plugin.bin`) from `C:\ProgramData\NetApp\SnapCenter\Package Repository` to the host where you want to install the Plug-in for PostgreSQL.

You can access this path from the host where the SnapCenter Server is installed.

2. From the command prompt, navigate to the directory where you copied the installation file.

3. Install the plug-in: `path_to_installation_bin_file/snapcenter_linux_host_plugin.bin -i silent -DPORT=port_number_for_host -DSERVER_IP=server_name_or_ip_address -DSERVER_HTTPS_PORT=port_number_for_server`

- -DPORT specifies the SMCORE HTTPS communication port.
- -DSERVER_IP specifies the SnapCenter Server IP address.
- -DSERVER_HTTPS_PORT specifies the SnapCenter Server HTTPS port.
- -DUSER_INSTALL_DIR specifies the directory where you want to install the SnapCenter Plug-ins Package for Linux.
- DINSTALL_LOG_NAME specifies the name of the log file.

```
/tmp/sc-plugin-installer/snapcenter_linux_host_plugin.bin -i silent
-DPORT=8145 -DSERVER_IP=scserver.domain.com -DSERVER_HTTPS_PORT=8146
-DUSER_INSTALL_DIR=/opt
-DINSTALL_LOG_NAME=SnapCenter_Linux_Host_Plugin_Install_2.log
-DCHOSEN_FEATURE_LIST=CUSTOM
```

4. Edit the `/<installation directory>/NetApp/snapcenter/scc/etc/SC_SMS_Services.properties` file, and then add the `PLUGINS_ENABLED = PostgreSQL:3.0` parameter.
5. Add the host to the SnapCenter Server using the `Add-Smhost` cmdlet and the required parameters.






The information regarding the parameters that can be used with the command and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the [SnapCenter Software Cmdlet Reference Guide](#).

Monitor the status of installing Plug-in for PostgreSQL

You can monitor the progress of SnapCenter plug-in package installation by using the Jobs page. You might want to check the progress of installation to determine when it is complete or if there is an issue.

About this task

The following icons appear on the Jobs page and indicate the state of the operation:

-  In progress
-  Completed successfully
-  Failed
-  Completed with warnings or could not start due to warnings
-  Queued

Steps

1. In the left navigation pane, click **Monitor**.
2. In the **Monitor** page, click **Jobs**.
3. In the **Jobs** page, to filter the list so that only plug-in installation operations are listed, do the following:

- a. Click **Filter**.
 - b. Optional: Specify the start and end date.
 - c. From the Type drop-down menu, select **Plug-in installation**.
 - d. From the Status drop-down menu, select the installation status.
 - e. Click **Apply**.
4. Select the installation job and click **Details** to view the job details.
 5. In the **Job Details** page, click **View logs**.

Configure CA Certificate

Generate CA Certificate CSR file

You can generate a Certificate Signing Request (CSR) and import the certificate that can be obtained from a Certificate Authority (CA) using the generated CSR. The certificate will have a private key associated with it.

CSR is a block of encoded text that is given to an authorized certificate vendor to procure the signed CA certificate.



CA Certificate RSA key length must be minimum 3072 bits.

For information to generate a CSR, see [How to generate CA Certificate CSR file](#).



If you own the CA certificate for your domain (*.domain.company.com) or your system (machine1.domain.company.com), you can skip generating the CA Certificate CSR file. You can deploy the existing CA certificate with SnapCenter.

For cluster configurations, the cluster name (virtual cluster FQDN), and the respective host names should be mentioned in the CA certificate. The certificate can be updated by filling the Subject Alternative Name (SAN) field before procuring the certificate. For a wild card certificate (*.domain.company.com), the certificate will contain all the hostnames of the domain implicitly.

Import CA certificates

You must import the CA certificates to the SnapCenter Server and the Windows host plug-ins using the Microsoft management console (MMC).

Steps

1. Go to the Microsoft management console (MMC), and then click **File > Add/Remove Snapin**.
2. In the Add or Remove Snap-ins window, select **Certificates** and then click **Add**.
3. In the Certificates snap-in window, select the **Computer account** option, and then click **Finish**.
4. Click **Console Root > Certificates – Local Computer > Trusted Root Certification Authorities > Certificates**.
5. Right-click on the folder “Trusted Root Certification Authorities”, and then select **All Tasks > Import** to start the import wizard.
6. Complete the wizard, as follows:

In this wizard window...	Do the following...
Import Private Key	Select the option Yes , import the private key, and then click Next .
Import File Format	Make no changes; click Next .
Security	Specify the new password to be used for the exported certificate, and then click Next .
Completing the Certificate Import Wizard	Review the summary, and then click Finish to start the import.



Importing certificate should be bundled with the private key (supported formats are: *.pfx, *.p12, and *.p7b).

7. Repeat Step 5 for the “Personal” folder.

Get the CA certificate thumbprint

A certificate thumbprint is a hexadecimal string that identifies a certificate. A thumbprint is calculated from the content of the certificate using a thumbprint algorithm.

Steps

1. Perform the following on the GUI:
 - a. Double-click the certificate.
 - b. In the Certificate dialog box, click the **Details** tab.
 - c. Scroll through the list of fields and click **Thumbprint**.
 - d. Copy the hexadecimal characters from the box.
 - e. Remove the spaces between the hexadecimal numbers.

For example, if the thumbprint is: "a9 09 50 2d d8 2a e4 14 33 e6 f8 38 86 b0 0d 42 77 a3 2a 7b", after removing the spaces, it will be: "a909502dd82ae41433e6f83886b00d4277a32a7b".

2. Perform the following from PowerShell:
 - a. Run the following command to list the thumbprint of the installed certificate and identify the recently installed certificate by the subject name.

```
Get-ChildItem -Path Cert:\LocalMachine\My
```

- b. Copy the thumbprint.

Configure CA certificate with Windows host plug-in services

You should configure the CA certificate with Windows host plug-in services to activate the installed digital certificate.

Perform the following steps on the SnapCenter Server and all the plug-in hosts where CA certificates are

already deployed.

Steps

1. Remove the existing certificate binding with SMCore default port 8145, by running the following command:

```
> netsh http delete sslcert ipport=0.0.0.0:_{SMCore Port}
```

For example:

```
> netsh http delete sslcert ipport=0.0.0.0:8145
```

2. Bind the newly installed certificate with the Windows host plug-in services, by running the following commands:

```
> $cert = "_{certificate thumbprint}_"  
> $guid = [guid]::NewGuid().ToString("B")  
> netsh http add sslcert ipport=0.0.0.0:_{SMCore Port}_ certhash=$cert  
appid="$guid"
```

For example:

```
> $cert = "a909502dd82ae41433e6f83886b00d4277a32a7b"  
> $guid = [guid]::NewGuid().ToString("B")  
> netsh http add sslcert ipport=0.0.0.0:_{SMCore Port}_ certhash=$cert  
appid="$guid"
```

Configure the CA Certificate for the SnapCenter PostgreSQL Plug-ins service on Linux host

You should manage the password of the plug-ins keystore and its certificate, configure the CA certificate, configure root or intermediate certificates to the plug-ins trust-store, and configure CA signed key pair to plug-ins trust-store with SnapCenter plug-ins service to activate the installed digital certificate.

The plug-ins uses the file 'keystore.jks', which is located at `/opt/NetApp/snapcenter/scc/etc` both as its trust-store and key-store.

Manage password for plug-in keystore and alias of the CA signed key pair in use

Steps

1. You can retrieve plug-in keystore default password from plug-in agent property file.

It is the value corresponding to the key 'KEYSTORE_PASS'.

2. Change the keystore password:

```
keytool -storepasswd -keystore keystore.jks
```

3. Change the password for all aliases of private key entries in the keystore to the same password used for the keystore:

```
keytool -keypasswd -alias "alias_name_in_cert" -keystore keystore.jks
```

Update the same for the key `KEYSTORE_PASS` in `agent.properties` file.

4. Restart the service after changing the password.



Password for plug-in keystore and for all the associated alias password of the private key should be same.

Configure root or intermediate certificates to plug-in trust-store

You should configure the root or intermediate certificates without the private key to plug-in trust-store.

Steps

1. Navigate to the folder containing the plug-in keystore: `/opt/NetApp/snapcenter/scc/etc`.
2. Locate the file 'keystore.jks'.
3. List the added certificates in the keystore:

```
keytool -list -v -keystore keystore.jks
```

4. Add a root or intermediate certificate:

```
keytool -import -trustcacerts -alias myRootCA -file  
/root/USERTrustRSA_Root.cer -keystore keystore.jks
```

5. Restart the service after configuring the root or intermediate certificates to plug-in trust-store.



You should add the root CA certificate and then the intermediate CA certificates.

Configure CA signed key pair to plug-in trust-store

You should configure the CA signed key pair to the plug-in trust-store.

Steps

1. Navigate to the folder containing the plug-in keystore `/opt/NetApp/snapcenter/scc/etc`.
2. Locate the file 'keystore.jks'.
3. List the added certificates in the keystore:

```
keytool -list -v -keystore keystore.jks
```

4. Add the CA certificate having both private and public key.

```
keytool -importkeystore -srckeystore /root/snapcenter.ssl.test.netapp.com.pfx
-srcstoretype pkcs12 -destkeystore keystore.jks -deststoretype JKS
```

5. List the added certificates in the keystore.

```
keytool -list -v -keystore keystore.jks
```

6. Verify that the keystore contains the alias corresponding to the new CA certificate, which was added to the keystore.
7. Change the added private key password for CA certificate to the keystore password.

Default plug-in keystore password is the value of the key `KEYSTORE_PASS` in `agent.properties` file.

```
keytool -keypasswd -alias "alias_name_in_CA_cert" -keystore keystore.jks
```

8. If the alias name in the CA certificate is long and contains space or special characters ("*", ",",), change the alias name to a simple name:

```
keytool -changealias -alias "long_alias_name" -destalias "simple_alias"
-keystore keystore.jks
```

9. Configure the alias name from CA certificate in `agent.properties` file.

Update this value against the key `SCC_CERTIFICATE_ALIAS`.

10. Restart the service after configuring the CA signed key pair to plug-in trust-store.

Configure certificate revocation list (CRL) for plug-ins

About this task

- SnapCenter Plug-ins will search for the CRL files in a pre-configured directory.
- Default directory for the CRL files for SnapCenter Plug-ins is 'opt/NetApp/snapcenter/scc/etc/crl'.

Steps

1. You can modify and update the default directory in `agent.properties` file against the key `CRL_PATH`.

You can place more than one CRL file in this directory. The incoming certificates will be verified against each CRL.

Configure the CA Certificate for the SnapCenter PostgreSQL Plug-ins service on Windows host

You should manage the password of the plug-ins keystore and its certificate, configure the CA certificate, configure root or intermediate certificates to the plug-ins trust-store, and configure CA signed key pair to plug-ins trust-store with SnapCenter plug-ins service to activate the installed digital certificate.

The plug-ins uses the file `keystore.jks`, which is located at `C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc` both as its trust-store and key-store.

Manage password for plug-in keystore and alias of the CA signed key pair in use

Steps

1. You can retrieve plug-in keystore default password from plug-in agent property file.

It is the value corresponding to the key `KEYSTORE_PASS`.

2. Change the keystore password:

```
keytool -storepasswd -keystore keystore.jks
```



If the "keytool" command is not recognized on the Windows command prompt, replace the keytool command with its complete path.

```
C:\Program Files\Java\<jdk_version>\bin\keytool.exe" -storepasswd -keystore keystore.jks
```

3. Change the password for all aliases of private key entries in the keystore to the same password used for the keystore:

```
keytool -keypasswd -alias "alias_name_in_cert" -keystore keystore.jks
```

Update the same for the key `KEYSTORE_PASS` in `agent.properties` file.

4. Restart the service after changing the password.



Password for plug-in keystore and for all the associated alias password of the private key should be same.

Configure root or intermediate certificates to plug-in trust-store

You should configure the root or intermediate certificates without the private key to plug-in trust-store.

Steps

1. Navigate to the folder containing the plug-in keystore `C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc`
2. Locate the file 'keystore.jks'.
3. List the added certificates in the keystore:

```
keytool -list -v -keystore keystore.jks
```

4. Add a root or intermediate certificate:

```
keytool -import -trustcacerts -alias myRootCA -file /root/USERTrustRSA_Root.cer -keystore keystore.jks
```

5. Restart the service after configuring the root or intermediate certificates to plug-in trust-store.



You should add the root CA certificate and then the intermediate CA certificates.

Configure CA signed key pair to plug-in trust-store

You should configure the CA signed key pair to the plug-in trust-store.

Steps

1. Navigate to the folder containing the plug-in keystore *C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc*
2. Locate the file *keystore.jks*.
3. List the added certificates in the keystore:

```
keytool -list -v -keystore keystore.jks
```

4. Add the CA certificate having both private and public key.

```
keytool -importkeystore -srckeystore /root/snapcenter.ssl.test.netapp.com.pfx -srcstoretype pkcs12 -destkeystore keystore.jks -deststoretype JKS
```

5. List the added certificates in the keystore.

```
keytool -list -v -keystore keystore.jks
```

6. Verify that the keystore contains the alias corresponding to the new CA certificate, which was added to the keystore.

7. Change the added private key password for CA certificate to the keystore password.

Default plug-in keystore password is the value of the key `KEYSTORE_PASS` in `agent.properties` file.

```
keytool -keypasswd -alias "alias_name_in_CA_cert" -keystore keystore.jks
```

8. Configure the alias name from CA certificate in `agent.properties` file.

Update this value against the key `SCC_CERTIFICATE_ALIAS`.

9. Restart the service after configuring the CA signed key pair to plug-in trust-store.

Configure certificate revocation list (CRL) for SnapCenter plug-ins

About this task

- To download the latest CRL file for the related CA certificate see [How to update certificate revocation list file in SnapCenter CA Certificate](#).
- SnapCenter plug-ins will search for the CRL files in a pre-configured directory.
- Default directory for the CRL files for SnapCenter plug-ins is *'C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc\crl'*.

Steps

1. You can modify and update the default directory in `agent.properties` file against the key `CRL_PATH`.
2. You can place more than one CRL file in this directory.

The incoming certificates will be verified against each CRL.

Enable CA Certificates for plug-ins

You should configure the CA certificates and deploy the CA certificates in the SnapCenter Server and the corresponding plug-in hosts. You should enable the CA certificate validation for the plug-ins.

Before you begin

- You can enable or disable the CA certificates using the run *Set-SmCertificateSettings* cmdlet.
- You can display the certificate status for the plug-ins using the *Get-SmCertificateSettings*.





The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the [SnapCenter Software Cmdlet Reference Guide](#).

Steps

1. In the left navigation pane, click **Hosts**.
2. In the Hosts page, click **Managed Hosts**.
3. Select single or multiple plug-in hosts.
4. Click **More options**.
5. Select **Enable Certificate Validation**.

After you finish

The Managed Hosts tab host displays a padlock and the color of the padlock indicates the status of the connection between SnapCenter Server and the plug-in host.

-  indicates that the CA certificate is neither enabled nor assigned to the plug-in host.
-  indicates that the CA certificate is successfully validated.
-  indicates that the CA certificate could not be validated.
-  indicates that the connection information could not be retrieved.



When the status is yellow or green, the data protection operations completes successfully.

Prepare for data protection

Prerequisites for using the SnapCenter Plug-in for PostgreSQL

Before you use SnapCenter Plug-in for PostgreSQL, the SnapCenter administrator must install and configure the SnapCenter Server and perform the prerequisite tasks.

- Install and configure SnapCenter Server.
- Log in to SnapCenter Server.
- Configure the SnapCenter environment by adding storage system connections and creating credentials, if applicable.
- Install Java 11 on your Linux or Windows host.

You must set the Java path in the environmental path variable of the host machine.

- Set up SnapMirror and SnapVault, if you want backup replication.

How resources, resource groups, and policies are used for protecting PostgreSQL

Before you use SnapCenter, it is helpful to understand basic concepts related to the

backup, clone, and restore operations you want to perform. You interact with resources, resource groups, and policies for different operations.

- Resources are typically PostgreSQL clusters that you back up or clone with SnapCenter.
- A SnapCenter resource group, is a collection of resources on a host.

When you perform an operation on a resource group, you perform that operation on the resources defined in the resource group according to the schedule you specify for the resource group.

You can back up on demand a single resource or a resource group. You also can perform scheduled backups for single resources and resource groups.

- The policies specify the backup frequency, replication, scripts, and other characteristics of data protection operations.

When you create a resource group, you select one or more policies for that group. You can also select a policy when you perform a backup on demand for a single resource.

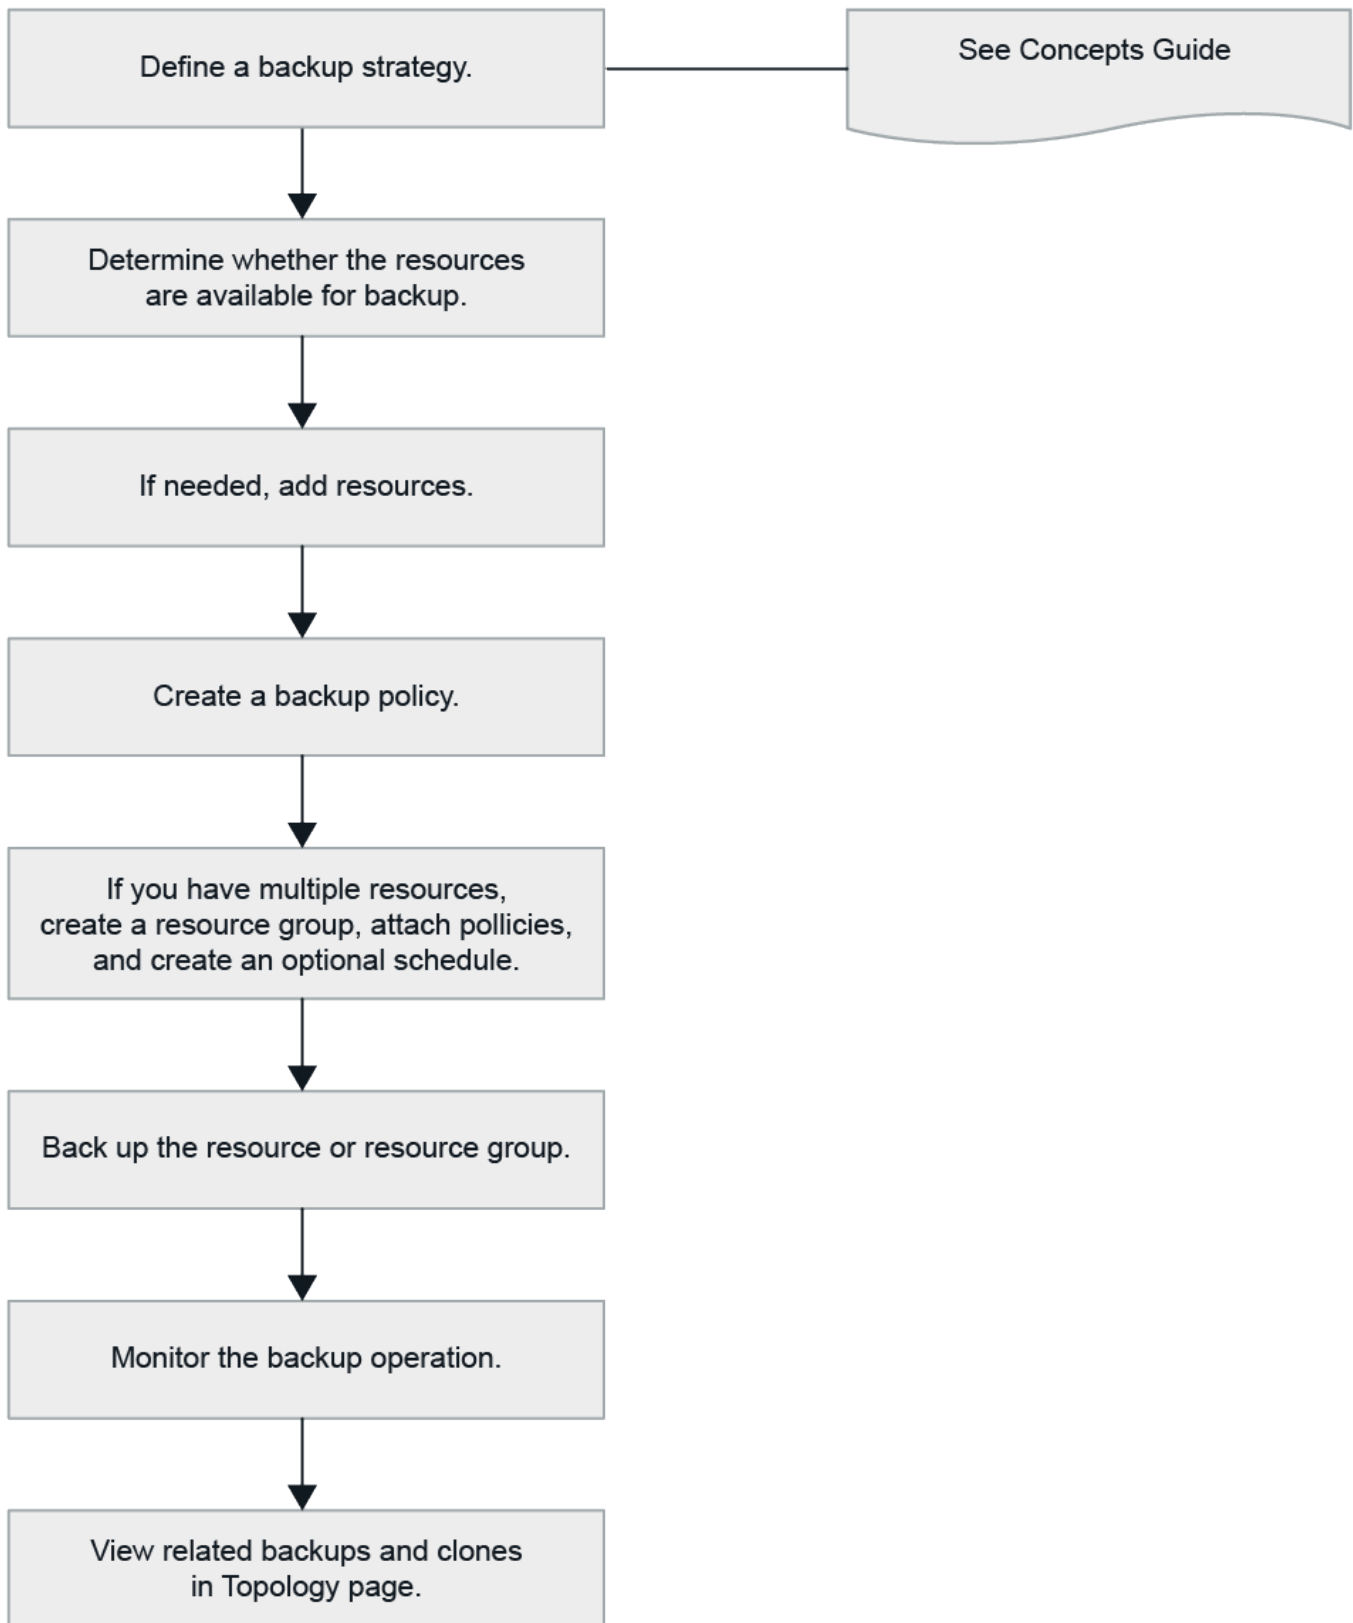
Think of a resource group as defining what you want to protect and when you want to protect it in terms of day and time. Think of a policy as defining how you want to protect it. If you are backing up all clusters, for example, you might create a resource group that includes all of the clusters in the host. You could then attach two policies to the resource group: a daily policy and an hourly policy. When you create the resource group and attach the policies, you might configure the resource group to perform a full backup daily.

Back up PostgreSQL resources

Back up PostgreSQL resources

You can either create a backup of a resource (cluster) or resource group. The backup workflow includes planning, identifying the clusters for backup, managing backup policies, creating resource groups and attaching policies, creating backups, and monitoring the operations.

The following workflow shows the sequence in which you must perform the backup operation:



You can also use PowerShell cmdlets manually or in scripts to perform backup, restore, and clone operations. The SnapCenter cmdlet help and the cmdlet reference information contain more information about PowerShell cmdlets. [SnapCenter Software Cmdlet Reference Guide](#).

Discover the clusters automatically

Resources are PostgreSQL clusters on the Linux host that are managed by SnapCenter. You can add the resources to resource groups to perform data protection operations after you discover the PostgreSQL clusters that are available.

Before you begin


- You must have already completed tasks such as installing the SnapCenter Server, adding hosts, and setting up the storage system connections.
- SnapCenter Plug-in for PostgreSQL does not support automatic discovery of the resources residing on RDM/VMDK virtual environments.

About this task

- After installing the plug-in, all the clusters on that Linux host are automatically discovered and displayed on the Resources page.
- Only clusters are auto-discovered.

The automatically discovered resources cannot be modified or deleted.

Steps

1. In the left navigation pane, click **Resources**, and then select the Plug-in for PostgreSQL from the list.
2. In the Resources page select the resource type from the View list.
3. (Optional) Click , and then select the host name.

You can then click  to close the filter pane.

4. Click **Refresh Resources** to discover the resources available on the host.

The resources are displayed along with information such as resource type, host name, associated resource groups, backup type, policies and overall status.

- If the cluster is on a NetApp storage and not protected, then Not protected is displayed in the Overall Status column.
- If the cluster is on a NetApp storage system and protected, and if there is no backup operation performed, then Backup not run is displayed in the Overall Status column. The status will otherwise change to Backup failed or Backup succeeded based on the last backup status.



You must refresh the resources if the clusters are renamed outside of SnapCenter.

Add resources manually to the plug-in host

Automatic discovery is not supported on Windows host. You must add Postgresql cluster resources manually.

Before you begin

- You must have completed tasks such as installing the SnapCenter Server, adding hosts, and setting up storage system connections.

About this task

Automatic discovery is not supported for the following configurations:

- RDM and VMDK layouts

Steps

1. In the left navigation pane, select the SnapCenter Plug-in for Postgresql from the drop-down list, and then click **Resources**.
2. In the Resources page, click **Add Postgresql resources**.
3. In the Provide Resource Details page, perform the following actions:

For this field...	Do this...
Name	Specify the cluster name.
Host Name	Enter the host name.
Type	Select cluster.
Instance	Specify the name of the instance, which is the parent of the cluster.
Credentials	Select the credentials or add information for the credential. This is optional.

4. In the Provide Storage Footprint page, select a storage type and choose one or more volumes, LUNs, and qtrees, and then click **Save**.

Optional: You can click the  icon to add more volumes, LUNs, and qtrees from other storage systems.

5. Optional: In the Resource Settings page, for resources on the Windows host, enter custom key-value pairs for PostgreSQL plug-in
6. Review the summary, and then click **Finish**.

The clusters are displayed along with information such as the host name, associated resource groups and policies, and overall status

If you want to provide users access to resources, you must assign the resources to the users. This enables users to perform the actions for which they have permissions on the assets that are assigned to them.

[Add a user or group and assign role and assets](#)

After you finish

- After adding the clusters, you can modify the PostgreSQL cluster details.
- The migrated resources (tablespace and clusters) from SnapCenter 5.0 will be tagged as PostgreSQL cluster type in SnapCenter 6.0.
- When you modify the manually added resources that are migrated from SnapCenter 5.0 or earlier, do the following in the **Resource Settings** page for custom key value pairs:
 - Specify the term "PORT" in the **Name** field.

- Specify the port number in the **Value** field.

Create backup policies for PostgreSQL

Before you use SnapCenter to back up PostgreSQL resources, you must create a backup policy for the resource or resource group that you want to back up. A backup policy is a set of rules that governs how you manage, schedule, and retain backups.

Before you begin

- You must have defined your backup strategy.

For details, see the information about defining a data protection strategy for PostgreSQL clusters.

- You must have prepared for data protection by completing tasks such as installing SnapCenter, adding hosts, setting up storage system connections, and adding resources.
- The SnapCenter administrator must have assigned the SVMs for both the source and destination volumes to you if you are replicating snapshots to a mirror or vault.

Additionally, you can specify replication, script, and application settings in the policy. These options saves time when you want to reuse the policy for another resource group.

About this task

- SnapLock
 - If 'Retain the backup copies for a specific number of days' option is selected, then the SnapLock retention period must be lesser than or equal to the mentioned retention days.
 - Specifying a snapshot locking period prevents deletion of the snapshots until the retention period expires. This could lead to retaining a larger number of snapshots than the count specified in the policy.
 - For ONTAP 9.12.1 and below version, the clones created from the SnapLock Vault snapshots as part of restore will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.

Steps

1. In the left navigation pane, click **Settings**.
2. In the Settings page, click **Policies**.
3. Click **New**.
4. In the Name page, enter the policy name and details.
5. In the Policy type page, perform the following:
 - a. Select storage type.
 - b. In the **Custom backup settings** section, provide any specific backup settings that have to be passed to the plug-in in key-value format.

You can provide multiple key-values to be passed to the plug-in.

6. In the Backup and Replication page, perform the following actions:
 - a. Specify the schedule frequency by selecting **On demand**, **Hourly**, **Daily**, **Weekly**, or **Monthly**.



You can specify the schedule (start date, end date, and frequency) for the backup operation while creating a resource group. This enables you to create resource groups that share the same policy and backup frequency, but also enables you to assign different backup schedules to each policy.



If you have scheduled for 2:00 a.m., the schedule will not be triggered during daylight saving time (DST).

- b. In the Snapshot settings section, specify the retention settings for the backup type and the schedule type selected in the **Backup Type** page:

If you want to...	Then...
Keep a certain number of Snapshots	<p>Select Copies to keep, and then specify the number of Snapshots that you want to keep.</p> <p>If the number of Snapshots exceeds the specified number, the Snapshots are deleted with the oldest copies deleted first.</p> <div data-bbox="898 940 954 999" data-label="Image"></div> <p>You must set the retention count to 2 or higher if you plan to enable SnapVault replication. If you set the retention count to 1, the retention operation might fail because the first Snapshot is the reference Snapshot for the SnapVault relationship until a newer Snapshot is replicated to the target.</p> <div data-bbox="898 1213 954 1272" data-label="Image"></div> <p>The maximum retention value is 1018. Backups will fail if retention is set to a value higher than what ONTAP version supports.</p>
Keep the Snapshots for a certain number of days	<p>Select Retain copies for, and then specify the number of days for which you want to keep the Snapshots before deleting them.</p>
Snapshot copy locking period	<p>Select Snapshot copy locking period, and specify days, months, or years.</p> <p>SnapLock retention period should be less than 100 years.</p>

7. Select a Policy label.



You can assign SnapMirror labels to primary snapshots for remote replication, allowing the primary snapshots to offload the snapshot replication operation from SnapCenter to ONTAP secondary systems. This can be done without enabling SnapMirror or SnapVault option in the policy page.

8. In the Select secondary replication options section, select one or both of the following secondary replication options:

For this field...	Do this...
Update SnapMirror after creating a local Snapshot copy	<p>Select this field to create mirror copies of the backup sets on another volume (SnapMirror replication).</p> <p>If the protection relationship in ONTAP is of type Mirror and Vault and if you select only this option, Snapshot created on the primary will not be transferred to the destination, but will be listed in the destination. If this Snapshot is selected from the destination to perform a restore operation, then the following error message is displayed: Secondary Location is not available for the selected vaulted/mirrored backup.</p> <p>During secondary replication, the SnapLock expiry time loads the primary SnapLock expiry time.</p> <p>Clicking the Refresh button in the Topology page refreshes the secondary and primary SnapLock expiry time that are retrieved from ONTAP.</p> <p>See View PostgreSQL resource related backups and clones in the Topology page.</p>

For this field...	Do this...
Update SnapVault after creating a local Snapshot copy	<p>Select this option to perform disk-to-disk backup replication (SnapVault backups).</p> <p>During secondary replication, the SnapLock expiry time loads the primary SnapLock expiry time. Clicking the Refresh button in the Topology page refreshes the secondary and primary SnapLock expiry time that are retrieved from ONTAP.</p> <p>When SnapLock is configured only on the secondary from ONTAP known as SnapLock Vault, clicking the Refresh button in the Topology page refreshes the locking period on the secondary that is retrieved from ONTAP.</p> <p>For more information on SnapLock Vault see Commit Snapshots to WORM on a vault destination</p> <p>See View PostgreSQL resource related backups and clones in the Topology page.</p>
Error retry count	Enter the maximum number of replication attempts that can be allowed before the operation stops.



You should configure SnapMirror retention policy in ONTAP for the secondary storage to avoid reaching the maximum limit of Snapshots on the secondary storage.

9. Review the summary, and then click **Finish**.

Create resource groups and attach policies


A resource group is the container to which you must add resources that you want to back up and protect. A resource group enables you to back up all the data that is associated with a given application simultaneously. A resource group is required for any data protection job. You must also attach one or more policies to the resource group to define the type of data protection job that you want to perform.

About this task

- For ONTAP 9.12.1 and below version, the clones created from the SnapLock Vault snapshots as part of restore will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, click **New Resource Group**.
3. In the Name page, perform the following actions:

For this field...	Do this...
Name	<p>Enter a name for the resource group.</p> <p> The resource group name should not exceed 250 characters.</p>
Tags	<p>Enter one or more labels that will help you later search for the resource group.</p> <p>For example, if you add HR as a tag to multiple resource groups, you can later find all resource groups associated with the HR tag.</p>
Use custom name format for snapshot copy	<p>Select this check box, and enter a custom name format that you want to use for the snapshot name.</p> <p>For example, customtext_resource_group_policy_hostname or resource_group_hostname. By default, a timestamp is appended to the snapshot name.</p>


- In the Resources page, select a host name from the **Host** drop-down list and resource type from the **Resource Type** drop-down list.

This helps to filter information on the screen.

- Select the resources from the **Available Resources** section, and then click the right arrow to move them to the **Selected Resources** section.
- In the Application Settings page, do the following:
 - Click the **Backups** arrow to set additional backup options:

Enable consistency group backup and perform the following tasks:

For this field...	Do this...
Afford time to wait for Consistency Group snapshot operation to complete	<p>Select Urgent, Medium, or Relaxed to specify the wait time for snapshot operation to complete.</p> <p>Urgent = 5 seconds, Medium = 7 seconds, and Relaxed = 20 seconds.</p>
Disable WAFL Sync	Select this to avoid forcing a WAFL consistency point.



- b. Click the **Scripts** arrow and enter the pre and post commands for quiesce, snapshot, and unquiesce operations. You can also enter the pre commands to be executed before exiting in the event of a failure.
- c. Click the **Custom Configurations** arrow and enter the custom key-value pairs required for all data protection operations using this resource.

Parameter	Setting	Description
ARCHIVE_LOG_ENABLE	(Y/N)	Enables the archive log management to delete the archive logs.
ARCHIVE_LOG_RETENTION	number_of_days	Specifies the number of days the archive logs are retained. This setting must be equal to or greater than NTAP_SNAPSHOT_RETENTIONS.
ARCHIVE_LOG_DIR	change_info_directory/logs	Specifies the path to the directory that contains the archive logs.
ARCHIVE_LOG_EXT	file_extension	Specifies the archive log file extension. For example, if the archive log is log_backup_0_1.log, specify .log as the file extension.

Parameter	Setting	Description
ARCHIVE_LOG_RECURSIVE_SE ARCH	(Y/N)	Enables the management of archive logs within subdirectories. You should use this parameter if the archive logs are located under subdirectories.



The custom key-value pairs are supported for PostgreSQL Linux plug-in systems and not supported for PostgreSQL cluster registered as a centralized windows plug-in.

d. Click the **Snapshot Copy Tool** arrow to select the tool to create snapshots:

If you want...	Then...
SnapCenter to use the plug-in for Windows and put the file system into a consistent state before creating a snapshot. For Linux resources, this option is not applicable.	Select SnapCenter with File System Consistency .
SnapCenter to create a storage level snapshot	Select SnapCenter without File System Consistency .
To enter the command to be executed on the host to create snapshot copies.	Select Other , and then enter the command to be executed on the host to create a snapshot.

7. In the Policies page, perform the following steps:

a. Select one or more policies from the drop-down list.



You can also create a policy by clicking  .

The policies are listed in the Configure schedules for selected policies section.

b.

In the Configure Schedules column, click  for the policy you want to configure.

c. In the Add schedules for policy *policy_name* dialog box, configure the schedule, and then click **OK**.

Where, *policy_name* is the name of the policy that you have selected.

The configured schedules are listed in the **Applied Schedules** column.

Third party backup schedules are not supported when they overlap with SnapCenter backup schedules.

8. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

You must also specify the sender and receiver email addresses, and the subject of the email. The SMTP

server must be configured in **Settings > Global Settings**.

9. Review the summary, and then click **Finish**.

Create resource groups and enable secondary protection for PostgreSQL resources on ASA r2 systems

You should create the resource group to add the resources that are on ASA r2 systems. You can also provision the secondary protection while creating the resource group.

Before you begin

- You should ensure that you are not adding both ONTAP 9.x resources and ASA r2 resources to the same resource group.
- You should ensure that you do not have a database with both ONTAP 9.x resources and ASA r2 resources.

About this task

- The secondary protection is available only if the logged-in user is assigned to the role that has the **SecondaryProtection** capability enabled.
- If you enabled secondary protection, the resource group is put into maintenance mode while creating the primary and secondary consistency groups. After the primary and secondary consistency groups are created, the resource group is put out of maintenance mode.
- SnapCenter does not support secondary protection for a clone resource.

Steps

1. In the left navigation pane, select **Resources**, and the appropriate plug-in from the list.
2. In the Resources page, click **New Resource Group**.
3. In the Name page, perform the following actions:
 - a. Enter a name for the resource group in the Name field.



The resource group name should not exceed 250 characters.

- b. Enter one or more labels in the Tag field to help you search for the resource group later.

For example, if you add HR as a tag to multiple resource groups, you can later find all resource groups associated with the HR tag.

- c. Select this check box, and enter a custom name format that you want to use for the Snapshot name.

For example, customtext_resource group_policy_hostname or resource group_hostname. By default, a timestamp is appended to the Snapshot name.

- d. Specify the destinations of the archive log files that you do not want to back up.



You should use the exact same destination as it was set in the application including prefix, if needed.

4. In the Resources page, select the database host name from the **Host** drop-down list.




The resources are listed in the Available Resources section only if the resource is discovered successfully. If you have recently added resources, they will appear on the list of available resources only after you refresh your resource list.

5. Select the ASA r2 resources from the Available Resources section and move them to the Selected Resources section.
6. In the Application Settings page, select the backup option.
7. In the Policies page, perform the following steps:
 - a. Select one or more policies from the drop-down list.



You can also create a policy by clicking .

In the Configure schedules for selected policies section, the selected policies are listed.

- b. Click  in the Configure Schedules column for the policy for which you want to configure a schedule.
 - c. In the Add schedules for policy *policy_name* window, configure the schedule, and then click **OK**.

Where, *policy_name* is the name of the policy that you have selected.

The configured schedules are listed in the Applied Schedules column.

Third party backup schedules are not supported when they overlap with SnapCenter backup schedules.

8. If the secondary protection is enabled for the policy that you have selected, then Secondary Protection page is displayed and you need to perform the following steps:
 - a. Select the type of the replication policy.



Synchronous replication policy is not supported.


- b. Specify the consistency group suffix that you want to use.
 - c. From the Destination Cluster and Destination SVM drop-downs select the peered cluster and SVM that you want to use.



The cluster and SVM peering is not supported by SnapCenter. You should use System Manager or ONTAP CLIs to perform cluster and SVM peering.



If the resources are already protected outside of SnapCenter, those resources will be displayed in the Secondary Protected Resources section.

1. On the Verification page, perform the following steps:
 - a. Click **Load locators** to load the SnapMirror or SnapVault volumes to perform verification on secondary storage.
 - b. Click  in the Configure Schedules column to configure the verification schedule for all the schedule types of the policy.

c. In the Add Verification Schedules policy_name dialog box, perform the following actions:

If you want to...	Do this...
Run verification after backup	Select Run verification after backup .
Schedule a verification	Select Run scheduled verification and then select the schedule type from the drop-down list.

d. Select **Verify on secondary location** to verify your backups on secondary storage system.

e. Click **OK**.

The configured verification schedules are listed in the Applied Schedules column.

2. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

You must also specify the sender and receiver email addresses, and the subject of the email. If you want to attach the report of the operation performed on the resource group, select **Attach Job Report**.



For email notification, you must have specified the SMTP server details using either the GUI or the PowerShell command Set-SmSmtServer.

3. Review the summary, and then click **Finish**.

Create a storage system connection and a credential using PowerShell cmdlets for PostgreSQL

You must create a storage virtual machine (SVM) connection and a credential before using PowerShell cmdlets to back up, restore, or clone PostgreSQL clusters.

Before you begin

- You should have prepared the PowerShell environment to execute the PowerShell cmdlets.
- You should have the required permissions in the Infrastructure Admin role to create storage connections.
- You should ensure that the plug-in installations are not in progress.

Host plug-in installations must not be in progress while adding a storage system connection because the host cache might not be updated and clusters status might be displayed in the SnapCenter GUI as “Not available for backup” or “Not on NetApp storage”.

- Storage system names should be unique.

SnapCenter does not support multiple storage systems with the same name on different clusters. Each storage system that is supported by SnapCenter should have a unique name and a unique data LIF IP address.

Steps

1. Initiate a PowerShell Core connection session by using the Open-SmConnection cmdlet.

```
PS C:\> Open-SmConnection
```

2. Create a new connection to the storage system by using the Add-SmStorageConnection cmdlet.

```
PS C:\> Add-SmStorageConnection -Storage test_vs1 -Protocol https  
-Timeout 60
```

3. Create a new credential by using the Add-SmCredential cmdlet.

This example shows how to create a new credential named FinanceAdmin with Windows credentials:

```
PS C:> Add-SmCredential -Name FinanceAdmin -AuthMode Windows  
-Credential sddev\administrator
```

4. Add the PostgreSQL communication host to SnapCenter Server.

```
PS C:> Add-SmHost -HostName 10.232.204.61 -OSType Windows -RunAsName  
FinanceAdmin -PluginCode PostgreSQL
```

5. Install the package and the SnapCenter Plug-in for PostgreSQL on the host.

For Linux:

```
PS C:> Install-SmHostPackage -HostNames 10.232.204.61 -ApplicationCode  
PostgreSQL
```

For Windows:

```
Install-SmHostPackage -HostNames 10.232.204.61 -ApplicationCode  
PostgreSQL -FileSystemCode scw -RunAsName FinanceAdmin
```

6. Set the path to the SQLLIB.

For Windows, PostgreSQL plug-in will use the default path for SQLLIB folder: "C:\Program Files\IBM\SQLLIB\BIN"

If you want to override the default path, use the following command.

```
PS C:> Set-SmConfigSettings -Plugin -HostName 10.232.204.61 -PluginCode  
PostgreSQL -configSettings @{"PostgreSQL_SQLLIB_CMD" =  
"<custom_path>\IBM\SQLLIB\BIN"}
```

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the [SnapCenter Software Cmdlet Reference Guide](#).

Back up PostgreSQL

If a resource is not yet part of any resource group, you can back up the resource from the Resources page.

Before you begin

- You must have created a backup policy.
- If you want to back up a resource that has a SnapMirror relationship with a secondary storage, the ONTAP role assigned to the storage user should include the “snapmirror all” privilege. However, if you are using the “vsadmin” role, then the “snapmirror all” privilege is not required.
- For Snapshot copy-based backup operation, ensure that all the tenant clusters are valid and active.
- For pre and post commands for quiesce, Snapshot, and unquiesce operations, you should check if the commands exist in the command list available on the plug-in host from the following paths:
 - Default location on the Windows host: *C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc\allowed_commands.config*
 - Default location on the Linux host: */opt/NetApp/snapcenter/scc/etc/allowed_commands.config*





If the commands do not exist in the command list, then the operation will fail.

SnapCenter UI

Steps

1. In the left navigation pane, select **Resources**, and then select the appropriate plug-in from the list.
2. In the Resource page, filter resources from the **View** drop-down list based on resource type.

Select , and then select the host name and the resource type to filter the resources. You can then select  to close the filter pane.

3. Select the resource that you want to back up.
4. In the Resource page, select **Use custom name format for Snapshot copy**, and then enter a custom name format that you want to use for the Snapshot name.

For example, *customtext_policy_hostname* or *resource_hostname*. By default, a timestamp is appended to the Snapshot name.

5. In the Application Settings page, do the following:
 - Select the **Backups** arrow to set additional backup options:

Enable consistency group backup, if needed, and perform the following tasks:

For this field...	Do this...
Afford time to wait for "Consistency Group Snapshot" operation to complete	Select Urgent , or Medium , or Relaxed to specify the wait time for Snapshot operation to finish. Urgent = 5 seconds, Medium = 7 seconds, and Relaxed = 20 seconds.
Disable WAFL Sync	Select this to avoid forcing a WAFL consistency point.


- Select the **Scripts** arrow to run pre and post commands for quiesce, Snapshot, and unquiesce operations.

You can also run pre commands before exiting the backup operation. Prescripts and postscripts are run in the SnapCenter Server.

- Select the **Custom Configurations** arrow, and then enter the custom value pairs required for all jobs using this resource.
- Select the **Snapshot Copy Tool** arrow to select the tool to create Snapshots:

If you want...	Then...
SnapCenter to create a storage-level Snapshot	Select SnapCenter without File System Consistency .
SnapCenter to use the plug-in for Windows to put the file system into a consistent state and then create a Snapshot	Select SnapCenter with File System Consistency .


If you want...	Then...
To enter the command to create a Snapshot	Select Other , and then enter the command to create a Snapshot.



6. In the Policies page, perform the following steps:
- Select one or more policies from the drop-down list.

 You can also create a policy by clicking .

In the Configure schedules for selected policies section, the selected policies are listed.

- Select  in the Configure Schedules column for the policy for which you want to configure a schedule.
- In the Add schedules for policy *policy_name* dialog box, configure the schedule, and then select **OK**.

policy_name is the name of the policy that you selected.

The configured schedules are listed in the Applied Schedules column.

7. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

You must also specify the sender and receiver email addresses, and the subject of the email. SMTP must also be configured in **Settings > Global Settings**.

8. Review the summary, and then select **Finish**.

The resources topology page is displayed.

9. Select **Back up Now**.

10. In the Backup page, perform the following steps:

- a. If you applied multiple policies to the resource, from the **Policy** drop-down list, select the policy that you want to use for backup.

If the policy selected for the on-demand backup is associated with a backup schedule, the on-demand backups will be retained based on the retention settings specified for the schedule type.

- b. Select **Backup**.

11. Monitor the operation progress by clicking **Monitor > Jobs**.

- In MetroCluster configurations, SnapCenter might not be able to detect a protection relationship after a failover.

For information, see: [Unable to detect SnapMirror or SnapVault relationship after MetroCluster failover](#)

- If you are backing up application data on VMDKs and the Java heap size for the SnapCenter Plug-in for VMware vSphere is not large enough, the backup might fail.

To increase the Java heap size, locate the script file `/opt/netapp/init_scripts/scvservice`. In that script, the `do_start method` command starts the SnapCenter VMware plug-in service. Update that command to the following: `Java -jar -Xmx8192M -Xms4096M`

PowerShell cmdlets

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user by using the `Open-SmConnection` cmdlet.

```
PS C:\> Open-SmConnection
```

The username and password prompt is displayed.

2. Add manual resources by using the `Add-SmResources` cmdlet.

This example shows how to add a PostgreSQL instance:

```
PS C:\> Add-SmResource -HostName 10.32.212.13 -PluginCode PostgreSQL
-ResourceType Instance -ResourceName postgresqlinst1
-StorageFootPrint
(@{"VolumeName"="winpostgresql01_data01";"LUNName"="winpostgresql01_
data01";"StorageSystem"="scsnfssvm"}) -MountPoints "D:\"
```

3. Create a backup policy by using the `Add-SmPolicy` cmdlet.
4. Protect the resource or add a new resource group to SnapCenter by using the `Add-SmResourceGroup` cmdlet.
5. Initiate a new backup job by using the `New-SmBackup` cmdlet.

This example shows how to backup a resource group:

```
C:\PS> New-SMBackup -ResourceGroupName 'ResourceGroup_wback-up-  
clusters-using-powershell-cmdlets-postgresql.adocith_Resources'  
-Policy postgresql_policy1
```

This example backs up a protected resource:

```
C:\PS> New-SMBackup -Resources  
@{"Host"="10.232.204.42";"Uid"="MDC\SID";"PluginName"="postgresql"}  
-Policy postgresql_policy2
```

6. Monitor the job status (running, completed, or failed) by using the `Get-smJobSummaryReport` cmdlet.

```
PS C:\> Get-smJobSummaryReport -JobID 123
```

7. Monitor the backup job details like backup ID, backup name to perform restore or clone operation by using the `Get-SmBackupReport` cmdlet.

```

PS C:\> Get-SmBackupReport -JobId 351
Output:
BackedUpObjects           : {DB1}
FailedObjects             : {}
IsScheduled               : False
HasMetadata               : False
SmBackupId                : 269
SmJobId                   : 2361
StartDateTime             : 10/4/2016 11:20:45 PM
EndDateTime               : 10/4/2016 11:21:32 PM
Duration                  : 00:00:46.2536470
CreatedDateTime           : 10/4/2016 11:21:09 PM
Status                    : Completed
ProtectionGroupName       : Verify_ASUP_Message_windows
SmProtectionGroupId       : 211
PolicyName                : test2
SmPolicyId                : 20
BackupName                : Verify_ASUP_Message_windows_scc54_10-04-
2016_23.20.46.2758
VerificationStatus        : NotVerified
VerificationStatuses      :
SmJobError                :
BackupType                 : SCC_BACKUP
CatalogingStatus          : NotApplicable
CatalogingStatuses        :
ReportDataCreatedDateTime :

```

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the [SnapCenter Software Cmdlet Reference Guide](#).

Back up resource groups

A resource group is a collection of resources on a host. A backup operation on the resource group is performed on all resources defined in the resource group.

Before you begin



- You must have created a resource group with a policy attached.
- If you want to back up a resource that has a SnapMirror relationship with a secondary storage, the ONTAP role assigned to the storage user should include the “snapmirror all” privilege. However, if you are using the “vsadmin” role, then the “snapmirror all” privilege is not required.

About this task

You can back up a resource group on demand from the Resources page. If a resource group has a policy attached and a schedule configured, then backups occur automatically according to the schedule.

Steps

1. In the left navigation pane, select **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select **Resource Group** from the **View** list.

You can search the resource group either by entering the resource group name in the search box or by selecting , and then selecting the tag. You can then select  to close the filter pane.

3. In the Resource Groups page, select the resource group that you want to back up, and then select **Back up Now**.
4. In the Backup page, perform the following steps:
 - a. If you associated multiple policies with the resource group, from the **Policy** drop-down list, select the policy that you want to use for backup.

If the policy selected for the on-demand backup is associated with a backup schedule, the on-demand backups will be retained based on the retention settings specified for the schedule type.







- b. Select **Backup**.
5. Monitor the operation progress by selecting **Monitor > Jobs**.

Monitor PostgreSQL backup operations


You can monitor the progress of different backup operations by using the SnapCenterJobs page. You might want to check the progress to determine when it is complete or if there is an issue.

About this task

The following icons appear on the Jobs page and indicate the corresponding state of the operations:


-  In progress
-  Completed successfully
-  Failed
-  Completed with warnings or could not start due to warnings
-  Queued
-  Canceled

Steps

1. In the left navigation pane, click **Monitor**.
2. In the Monitor page, click **Jobs**.
3. In the Jobs page, perform the following steps:
 - a. Click  to filter the list so that only backup operations are listed.
 - b. Specify the start and end dates.
 - c. From the **Type** drop-down list, select **Backup**.
 - d. From the **Status** drop-down, select the backup status.
 - e. Click **Apply** to view the operations completed successfully.

4. Select a backup job, and then click **Details** to view the job details.



Though the backup job status displays , when you click on job details you might see that some of the child tasks of the backup operation are still in progress or marked with warning signs.

5. In the Job Details page, click **View logs**.


The **View logs** button displays the detailed logs for the selected operation.

Monitor data protection operations on PostgreSQL clusters in the Activity pane

The Activity pane displays the five most recent operations performed. The Activity pane also displays when the operation was initiated and the status of the operation.

The Activity pane displays information regarding backup, restore, clone, and scheduled backup operations.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. Click  on the Activity pane to view the five most recent operations.

When you click one of the operations, the operation details are listed in the **Job Details** page.

Cancel backup operations for PostgreSQL

You can cancel backup operations that are queued.


What you will need

- You must be logged in as the SnapCenter Admin or job owner to cancel operations.
- You can cancel a backup operation from either the **Monitor** page or the **Activity** pane.
- You cannot cancel a running backup operation.
- You can use the SnapCenter GUI, PowerShell cmdlets, or CLI commands to cancel the backup operations.
- The **Cancel Job** button is disabled for operations that cannot be canceled.
- If you selected **All members of this role can see and operate on other members objects** in Users\Groups page while creating a role, you can cancel the queued backup operations of other members while using that role.

Steps

1. Perform one of the following actions:

From the...	Action
Monitor page	<ol style="list-style-type: none">a. In the left navigation pane, click Monitor > Jobs.b. Select the operation, and then click Cancel Job.

From the...	Action
Activity pane	<ol style="list-style-type: none"> After initiating the backup operation, click  on the Activity pane to view the five most recent operations. Select the operation. In the Job Details page, click Cancel Job.




The operation is canceled, and the resource is reverted to the previous state.

View PostgreSQL backups and clones in the Topology page

When you are preparing to back up or clone a resource, you might find it helpful to view a graphical representation of all backups and clones on the primary and secondary storage.

About this task

You can review the following icons in the Manage Copies view to determine whether the backups and clones are available on the primary or secondary storage (Mirror copies or Vault copies).

-  displays the number of backups and clones that are available on the primary storage.
-  displays the number of backups and clones that are mirrored on the secondary storage using SnapMirror technology.
-  displays the number of backups and clones that are replicated on the secondary storage using SnapVault technology.



The number of backups displayed includes the backups deleted from the secondary storage. For example, if you created 6 backups using a policy to retain only 4 backups, the number of backups displayed is 6.



Clones of a backup of a version-flexible mirror on a mirror-vault type volume are displayed in the topology view, but the mirror backup count in the topology view does not include the version-flexible backup.

In the Topology page, you can see all of the backups and clones that are available for the selected resource or resource group. You can view the details of those backups and clones, and then select them to perform data protection operations.

Steps

- In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
- In the Resources page, either select the resource or resource group from the **View** drop-down list.
- Select the resource either from the resource details view or from the resource group details view.

If the resource is protected, the topology page of the selected resource is displayed.

4. Review the **Summary card** to see a summary of the number of backups and clones available on the primary and secondary storage.

The **Summary Card** section displays the total number of Snapshot copy-based backups, and clones.

Clicking the **Refresh** button starts a query of the storage to display an accurate count.

If SnapLock enabled backup is taken, then clicking the **Refresh** button refreshes the primary and secondary SnapLock expiry time retrieved from ONTAP. A weekly schedule also refreshes the primary and secondary SnapLock expiry time retrieved from ONTAP.

When the application resource is spread across multiple volumes, the SnapLock expiry time for the backup will be the longest SnapLock expiry time that is set for a Snapshot in a volume. The longest SnapLock expiry time is retrieved from ONTAP.

After on demand backup, by clicking the **Refresh** button refreshes the details of backup or clone.



5. In the Manage Copies view, click **Backups** or **Clones** from the primary or secondary storage to see details of a backup or clone.

The details of the backups and clones are displayed in a table format.

6. Select the backup from the table, and then click the data protection icons to perform restore, clone, and delete operations.



You cannot rename or delete backups that are on the secondary storage.


7. If you want to delete a clone, select the clone from the table, and then click .
8. If you want to split a clone, select the clone from the table, and then click .

Restore PostgreSQL

Restore workflow

The restore and recovery workflow includes planning, performing the restore operations, and monitoring the operations.

The following workflow shows the sequence in which you must perform the restore operation:



You can also use PowerShell cmdlets manually or in scripts to perform backup, restore, and clone operations. The SnapCenter cmdlet help and the cmdlet reference information contain detailed information about PowerShell cmdlets.

[SnapCenter Software Cmdlet Reference Guide](#).

Restore and recover a manually added resource backup

You can use SnapCenter to restore and recover data from one or more backups.

Before you begin

- You must have backed up the resource or resource groups.
- You must have canceled any backup operation that is currently in progress for the resource or resource group that you want to restore.
- For pre restore, post restore, mount, and unmount commands, you should check if the commands exist in the command list available on the plug-in host from the following paths:
 - Default location on the Windows host: *C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc\allowed_commands.config*
 - Default location on the Linux host: */opt/NetApp/snapcenter/scc/etc/allowed_commands.config*



If the commands do not exist in the command list, then the operation will fail.

About this task

- For ONTAP 9.12.1 and below version, the clones created from the SnapLock Vault Snapshots as part of restore will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.

SnapCenter UI

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, filter resources from the **View** drop-down list based on resource type.

The resources are displayed along with type, host, associated resource groups and policies, and status.




Although a backup might be for a resource group, when you restore, you must select the individual resources you want to restore.

If the resource is not protected, “Not protected” is displayed in the Overall Status column. This can mean either that the resource is not protected, or that the resource was backed up by a different user.

3. Select the resource, or select a resource group and then select a resource in that group.

The resource topology page is displayed.

4. From the Manage Copies view, select **Backups** either from the primary or secondary (mirrored or vaulted) storage systems.

5. In the Primary backup(s) table, select the backup that you want to restore from, and then click .



Backup Name	End Date
rg1_scipr0191685001_01-05-2017_01.35.06.6463	1/5/2017 1:35:27 AM

6. In the Restore Scope page, select **Complete Resource**.

- a. If you select **Complete Resource**, all of the configured data volumes of the PostgreSQL cluster are restored.

If the resource contains volumes or qtrees, the Snapshots taken after the Snapshot selected for restore on such volumes or qtrees are deleted and cannot be recovered. Also, if any other resource is hosted on same volumes or qtrees, then that resource is also deleted.

You can select multiple LUNs.



If you select **All**, all the files on the volumes, qtrees, or LUNs are restored.

7. In the Pre ops page, enter pre restore and unmount commands to run before performing a restore job.

Unmount commands are not available for auto discovered resources.

8. In the Post ops page, enter mount and post restore commands to run after performing a restore job.

Mount commands are not available for auto discovered resources.

9. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

You must also specify the sender and receiver email addresses and the subject of the email. SMTP must also be configured on the **Settings > Global Settings** page.

10. Review the summary, and then click **Finish**.
11. Monitor the operation progress by clicking **Monitor > Jobs**.

PowerShell cmdlets

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user by using the `Open-SmConnection` cmdlet.

```
PS C:\> Open-Smconnection
```

2. Retrieve the information about the one or more backups that you want to restore by using the `Get-SmBackup` and `Get-SmBackupReport` cmdlets.

This example displays information about all available backups:

```
PS C:\> Get-SmBackup
```

BackupId	BackupName	BackupTime	BackupType
-----	-----	-----	-----
1	Payroll Dataset_vise-f6_08...	8/4/2015	11:02:32
AM	Full Backup		
2	Payroll Dataset_vise-f6_08...	8/4/2015	11:23:17
AM			

This example displays detailed information about the backup from January 29th 2015 to February 3rd, 2015:

```
PS C:\> Get-SmBackupReport -FromDate "1/29/2015" -ToDate "2/3/2015"
```

```
SmBackupId      : 113
SmJobId         : 2032
StartDateTime   : 2/2/2015 6:57:03 AM
EndDateTime     : 2/2/2015 6:57:11 AM
Duration        : 00:00:07.3060000
CreatedDateTime : 2/2/2015 6:57:23 AM
Status          : Completed
ProtectionGroupName : Clone
SmProtectionGroupId : 34
PolicyName      : Vault
SmPolicyId      : 18
BackupName      : Clone_SCSPR0019366001_02-02-2015_06.57.08
VerificationStatus : NotVerified

SmBackupId      : 114
SmJobId         : 2183
StartDateTime   : 2/2/2015 1:02:41 PM
EndDateTime     : 2/2/2015 1:02:38 PM
Duration        : -00:00:03.2300000
CreatedDateTime : 2/2/2015 1:02:53 PM
Status          : Completed
ProtectionGroupName : Clone
SmProtectionGroupId : 34
PolicyName      : Vault
SmPolicyId      : 18
BackupName      : Clone_SCSPR0019366001_02-02-2015_13.02.45
VerificationStatus : NotVerified
```

3. Restore data from the backup by using the Restore-SmBackup cmdlet.

```

Restore-SmBackup -PluginCode 'DummyPlugin' -AppObjectId
'scc54.sscore.test.com\DummyPlugin\NTP\DB1' -BackupId 269
-Confirm:$false
output:
Name                : Restore
'scc54.sscore.test.com\DummyPlugin\NTP\DB1'
Id                  : 2368
StartTime           : 10/4/2016 11:22:02 PM
EndTime             :
IsCancellable       : False
IsRestartable       : False
IsCompleted         : False
IsVisible           : True
IsScheduled         : False
PercentageCompleted : 0
Description         :
Status              : Queued
Owner               :
Error               :
Priority            : None
Tasks               : {}
ParentJobID        : 0
EventId            : 0
JobTypeId           :
ApisJobKey         :
ObjectId           : 0
PluginCode         : NONE
PluginName         :

```

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the [SnapCenter Software Cmdlet Reference Guide](#).

Restore and recover an auto discovered cluster backup

You can use SnapCenter to restore and recover data from one or more backups.

Before you begin

- You must have backed up the resource or resource groups.
- You must have canceled any backup operation that is currently in progress for the resource or resource group that you want to restore.
- For pre restore, post restore, mount, and unmount commands, you should check if the commands exist in the command list available on the plug-in host from the following paths:
 - Default location on the Windows host: *C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in*

Creator\etc\allowed_commands.config

- Default location on the Linux host: /opt/NetApp/snapcenter/scc/etc/allowed_commands.config



If the commands do not exist in the command list, then the operation will fail.

About this task

- File-based backup copies cannot be restored from SnapCenter.
- For Auto-discovered resources, restore is supported with SFSR.
- Auto-recovery is not supported.
- For ONTAP 9.12.1 and below version, the clones created from the SnapLock Vault snapshots as part of restore will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, filter resources from the **View** drop-down list based on resource type.

The resources are displayed along with type, host, associated resource groups and policies, and status.




Although a backup might be for a resource group, when you restore, you must select the individual resources you want to restore.

If the resource is not protected, “Not protected” is displayed in the Overall Status column. This can mean either that the resource is not protected, or that the resource was backed up by a different user.

3. Select the resource, or select a resource group and then select a resource in that group.

The resource topology page is displayed.

4. From the Manage Copies view, select **Backups** either from the primary or secondary (mirrored or vaulted) storage systems.

5. In the Primary backup(s) table, select the backup that you want to restore from, and then click .



Backup Name	End Date
rg1_scpr0191685001_01-05-2017_01.35.06.6463	1/5/2017 1:35:27 AM

6. In the Restore Scope page, select **Complete Resource** to restore the configured data volumes of the PostgreSQL cluster.
7. In the Recovery scope page, select one of the following options:

If you...	Do this...
-----------	------------

Want to recover as close as possible to the current time	Select Recover to most recent state . For single container resources specify one or more log and catalog backup locations.
Want to recover to the specified point in time	Select Recover to point in time . a. Enter date and time. Enter date and time. For example, the PostgreSQL Linux host is located in Sunnyvale, CA and the user in Raleigh, NC is recovering the logs in to SnapCenter. If the user wants to perform a recovery to 5 a.m. Sunnyvale, CA, then the user has to set the browser time zone to the PostgreSQL Linux host time zone, which is GMT-07:00 and specify the date and time as 5:00 a.m.
Do not want to recover	Select No recovery .



You cannot recover manually added PostgreSQL resources.



SnapCenter Plug-in for PostgreSQL creates a `backup_label` and `tablespace_map` in `/<OS_temp_folder>/postgresql_sc_recovery<Restore_JobId>/_` folder to help recover manually.

- In the Pre ops page, enter pre restore and unmount commands to run before performing a restore job.
Unmount commands are not available for auto discovered resources.
- In the Post ops page, enter mount and post restore commands to run after performing a restore job.
Mount commands are not available for auto discovered resources.
- In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

You must also specify the sender and receiver email addresses and the subject of the email. SMTP must also be configured on the **Settings > Global Settings** page.
- Review the summary, and then click **Finish**.
- Monitor the operation progress by clicking **Monitor > Jobs**.

Restore resources using PowerShell cmdlets

Restoring a resource backup includes initiating a connection session with the SnapCenter Server, listing the backups and retrieving backup information, and restoring a backup.

You must have prepared the PowerShell environment to execute the PowerShell cmdlets.

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user by using the `Open-SmConnection` cmdlet.

```
PS C:\> Open-Smconnection
```

2. Retrieve the information about the one or more backups that you want to restore by using the `Get-SmBackup` and `Get-SmBackupReport` cmdlets.

This example displays information about all available backups:

```
PS C:\> Get-SmBackup
```

BackupId	BackupName	BackupTime
1	Payroll Dataset_vise-f6_08...	8/4/2015 11:02:32 AM
2	Payroll Dataset_vise-f6_08...	8/4/2015 11:23:17 AM

This example displays detailed information about the backup from January 29th 2015 to February 3rd, 2015:

```

PS C:\> Get-SmBackupReport -FromDate "1/29/2015" -ToDate "2/3/2015"

SmBackupId      : 113
SmJobId         : 2032
StartDateTime   : 2/2/2015 6:57:03 AM
EndDateTime     : 2/2/2015 6:57:11 AM
Duration        : 00:00:07.3060000
CreatedDateTime : 2/2/2015 6:57:23 AM
Status          : Completed
ProtectionGroupName : Clone
SmProtectionGroupId : 34
PolicyName      : Vault
SmPolicyId      : 18
BackupName      : Clone_SCSPR0019366001_02-02-2015_06.57.08
VerificationStatus : NotVerified

SmBackupId      : 114
SmJobId         : 2183
StartDateTime   : 2/2/2015 1:02:41 PM
EndDateTime     : 2/2/2015 1:02:38 PM
Duration        : -00:00:03.2300000
CreatedDateTime : 2/2/2015 1:02:53 PM
Status          : Completed
ProtectionGroupName : Clone
SmProtectionGroupId : 34
PolicyName      : Vault
SmPolicyId      : 18
BackupName      : Clone_SCSPR0019366001_02-02-2015_13.02.45
VerificationStatus : NotVerified

```

3. Restore data from the backup by using the Restore-SmBackup cmdlet.

```

Restore-SmBackup -PluginCode 'DummyPlugin' -AppObjectId
'scc54.sscore.test.com\DummyPlugin\NTP\DB1' -BackupId 269
-Confirm:$false
output:
Name                : Restore
'scc54.sscore.test.com\DummyPlugin\NTP\DB1'
Id                  : 2368
StartTime           : 10/4/2016 11:22:02 PM
EndTime             :
IsCancellable       : False
IsRestartable      : False
IsCompleted         : False
IsVisible           : True
IsScheduled         : False
PercentageCompleted : 0
Description         :
Status              : Queued
Owner               :
Error               :
Priority             : None
Tasks               : {}
ParentJobID        : 0
EventId            : 0
JobTypeId           :
ApisJobKey         :
ObjectId           : 0
PluginCode         : NONE
PluginName         :

```

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the [SnapCenter Software Cmdlet Reference Guide](#).

Monitor PostgreSQL restore operations






You can monitor the progress of different SnapCenter restore operations by using the Jobs page. You might want to check the progress of an operation to determine when it is complete or if there is an issue.

About this task


Post-restore states describe the conditions of the resource after a restore operation and any further restore actions that you can take.

The following icons appear on the Jobs page, and indicate the state of the operation:

-  In progress

-  Completed successfully
-  Failed
-  Completed with warnings or could not start due to warnings
-  Queued
-  Canceled

Steps

1. In the left navigation pane, click **Monitor**.
2. In the **Monitor** page, click **Jobs**.
3. In the **Jobs** page, perform the following steps:
 - a. Click  to filter the list so that only restore operations are listed.
 - b. Specify the start and end dates.
 - c. From the **Type** drop-down list, select **Restore**.
 - d. From the **Status** drop-down list, select the restore status.
 - e. Click **Apply** to view the operations that have been completed successfully.
4. Select the restore job, and then click **Details** to view the job details.
5. In the **Job Details** page, click **View logs**.

The **View logs** button displays the detailed logs for the selected operation.

Clone PostgreSQL resource backups


Clone workflow

The clone workflow includes performing the clone operation and monitoring the operation.

About this task

- You can clone on the source PostgreSQL server.
- You might clone resource backups for the following reasons:
 - To test functionality that has to be implemented using the current resource structure and content during application development cycles
 - For data extraction and manipulation tools when populating data warehouses
 - To recover data that was mistakenly deleted or changed

The following workflow shows the sequence in which you must perform the clone operation:



You can also use PowerShell cmdlets manually or in scripts to perform backup, restore, and clone operations. The SnapCenter cmdlet help and the cmdlet reference information contain detailed information about PowerShell cmdlets.

Clone a PostgreSQL backup

You can use SnapCenter to clone a backup. You can clone from primary or secondary backup.

Before you begin

- You should have backed up the resources or resource group.
- You should ensure that the aggregates hosting the volumes should be in the assigned aggregates list of the storage virtual machine (SVM).
- For pre clone or post clone commands, you should check if the commands exist in the command list available on the plug-in host from the following paths:
 - Default location on the Windows host: *C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc\allowed_commands.config*
 - Default location on the Linux host: */opt/NetApp/snapcenter/scc/etc/allowed_commands.config*



If the commands do not exist in the command list, then the operation will fail.

About this task

- For information about FlexClone volume split operations, see, [Split a FlexClone volume from its parent volume.](#)
- For ONTAP 9.12.1 and below version, the clones created from the SnapLock Vault Snapshots as part of restore will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.

SnapCenter UI

Steps


1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, filter resources from the **View** drop-down list based on resource type.

The resources are displayed along with information such as type, host, associated resource groups and policies, and status.

3. Select the resource or resource group.

You must select a resource if you select a resource group.

The resource or resource group topology page is displayed.

4. From the Manage Copies view, select **Backups** either from the primary or secondary (mirrored or vaulted) storage systems.
5. Select the data backup from the table, and then click .
6. In the Location page, perform the following actions:

For this field...	Do this...
Clone server	Choose a host on which the clone should be created.
Target Port	Enter the target PostgreSQL target port to clone from the existing backups.
NFS Export IP Address	Enter IP addresses or the host names on which the cloned volumes will be exported. This is applicable only to NFS storage type resource.
Capacity Pool Max. Throughput (MiB/s)	Enter the maximum throughput of a capacity pool. This is applicable only for ANF storage type resource.

7. In the Scripts page, perform the following steps:



The scripts are run on the plug-in host.

- a. Enter the commands for pre clone or post clone that should be run before or after the clone operation, respectively.
 - Pre clone command: delete existing clusters with the same name
 - Post clone command: verify a cluster or start a cluster.

b. Enter the mount command to mount a file system to a host.

Mount command for a volume or qtree on a Linux machine:

Example for NFS: `mount VSERVER_DATA_IP:%VOLUME_NAME_Clone /mnt`

8. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

You must also specify the sender and receiver email addresses, and the subject of the email.

9. Review the summary, and then click **Finish**.

10. Monitor the operation progress by clicking **Monitor > Jobs**.

PowerShell cmdlets

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user by using the `Open-SmConnection` cmdlet.

```
PS C:\> Open-SmConnection
```

2. Retrieve the backups to perform the clone operation by using the `Get-SmBackup` cmdlet.

This example shows that two backups are available for cloning:

```
C:\PS> Get-SmBackup

      BackupId                BackupName
-----
BackupTime                    BackupType
-----
1                               Payroll Dataset_vise-f6_08...
8/4/2015 11:02:32 AM          Full Backup
2                               Payroll Dataset_vise-f6_08...
8/4/2015 11:23:17 AM
```

3. Initiate a clone operation from an existing backup and specify the NFS export IP addresses on which the cloned volumes are exported.

This example shows that the backup to be cloned has an `NFSEXPORIPs` address of 10.32.212.14:

For PostgreSQL cluster:

```
PS C:\> New-SmClone -AppPluginCode PostgreSQL -BackupName "
scpostgresql01_ openenglab_netapp_com_PostgreSQL_postgres_5432_06-
26-2024_00_33_41_1570" -Resources @{"Host"="
10.32.212.13";"Uid"="postgres_5432"} -port 2345 -CloneToHost
10.32.212.14
```



If NFSExportIPs is not specified, the default is exported to the clone target host.

4. Verify that the backups were cloned successfully by using the Get-SmCloneReport cmdlet to view the clone job details.

You can view details such as clone ID, start date and time, end date and time.

```
PS C:\> Get-SmCloneReport -JobId 186







SmCloneId           : 1
SmJobId              : 186
StartDateTime       : 8/3/2015 2:43:02 PM
EndDateTime         : 8/3/2015 2:44:08 PM
Duration             : 00:01:06.6760000
Status               : Completed
ProtectionGroupName : Draper
SmProtectionGroupId : 4
PolicyName           : OnDemand_Clone
SmPolicyId           : 4
BackupPolicyName     : OnDemand_Full_Log
SmBackupPolicyId     : 1
CloneHostName       : SCSPR0054212005.mycompany.com
CloneHostId         : 4
CloneName            : Draper__clone__08-03-2015_14.43.53
SourceResources      : {Don, Betty, Bobby, Sally}
ClonedResources      : {Don_DRAPER, Betty_DRAPER, Bobby_DRAPER,
Sally_DRAPER}
SmJobError           :
```

Monitor PostgreSQL clone operations


You can monitor the progress of SnapCenter clone operations by using the Jobs page. You might want to check the progress of an operation to determine when it is complete or if there is an issue.

About this task

The following icons appear on the Jobs page, and indicate the state of the operation:

-  In progress
-  Completed successfully
-  Failed
-  Completed with warnings or could not start due to warnings
-  Queued
-  Canceled

Steps

1. In the left navigation pane, click **Monitor**.
2. In the **Monitor** page, click **Jobs**.
3. In the **Jobs** page, perform the following steps:
 - a. Click  to filter the list so that only clone operations are listed.
 - b. Specify the start and end dates.
 - c. From the **Type** drop-down list, select **Clone**.
 - d. From the **Status** drop-down list, select the clone status.
 - e. Click **Apply** to view the operations that are completed successfully.
4. Select the clone job, and then click **Details** to view the job details.
5. In the Job Details page, click **View logs**.

Split a clone

You can use SnapCenter to split a cloned resource from the parent resource. The clone that is split becomes independent of the parent resource.

About this task

- You cannot perform the clone split operation on an intermediate clone.

For example, after you create clone1 from a database backup, you can create a backup of clone1, and then clone this backup (clone2). After you create clone2, clone1 is an intermediate clone, and you cannot perform the clone split operation on clone1. However, you can perform the clone split operation on clone2.

After splitting clone2, you can perform the clone split operation on clone1 because clone1 is no longer the intermediate clone.

- When you split a clone, the backup copies and clone jobs of the clone are deleted.
- For information about FlexClone volume split operations, see, [Split a FlexClone volume from its parent volume](#).
- Ensure that the volume or aggregate on the storage system is online.


Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the **Resources** page, select the appropriate option from the View list:

Option	Description
For database applications	Select Database from the View list.
For file systems	Select Path from the View list.

3. Select the appropriate resource from the list.

The resource topology page is displayed.

4. From the **Manage Copies** view, select the cloned resource (for example, the database or LUN), and then click .
5. Review the estimated size of the clone that is to be split and the required space available on the aggregate, and then click **Start**.
6. Monitor the operation progress by clicking **Monitor > Jobs**.

The clone split operation stops responding if the SMCORE service restarts. You should run the Stop-SmJob cmdlet to stop the clone split operation, and then retry the clone split operation.

If you want a longer poll time or shorter poll time to check whether the clone is split or not, you can change the value of *CloneSplitStatusCheckPollTime* parameter in *SMCoreServiceHost.exe.config* file to set the time interval for SMCORE to poll for the status of the clone split operation. The value is in milliseconds and the default value is 5 minutes.

For example:

```
<add key="CloneSplitStatusCheckPollTime" value="300000" />
```

The clone split start operation fails if backup, restore, or another clone split is in progress. You should restart the clone split operation only after the running operations are complete.

Related information

[SnapCenter clone or verification fails with aggregate does not exist](#)

Delete or split PostgreSQL cluster clones after upgrading SnapCenter

After upgrading to SnapCenter 4.3, you will no longer see the clones. You can delete the clone or split the clones from the Topology page of the resource from which the clones were created.

About this task



If you want to locate the storage footprint of the hidden clones, run the following command: `Get-SmClone -ListStorageFootprint`

Steps

1. Delete the backups of the cloned resources by using the `remove-smbbackup` cmdlet.
2. Delete the resource group of the cloned resources by using the `remove-smresourcegroup` cmdlet.

3. Remove the protection of the cloned resource by using the `remove-smprotectresource` cmdlet.
4. Select the parent resource from the Resources page.

The resource topology page is displayed.

5. From the Manage Copies view, select the clones either from the primary or secondary (mirrored or replicated) storage systems.
6. Select the clones, and then click  to delete clones or click  to split the clones.
7. Click **OK**.

Copyright information

Copyright © 2026 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.