



Back up PostgreSQL resources

SnapCenter Software 6.0

NetApp
July 23, 2024

Table of Contents

- Back up PostgreSQL resources 1
 - Back up PostgreSQL resources 1
 - Discover the clusters automatically 1
 - Add resources manually to the plug-in host 2
 - Create backup policies for PostgreSQL 3
 - Create resource groups and attach policies 5
 - Back up PostgreSQL 9
 - Back up resource groups 12
 - Create a storage system connection and a credential using PowerShell cmdlets for PostgreSQL 13
 - Back up clusters using PowerShell cmdlets 14
 - Monitor backup operations 16
 - Cancel backup operations for PostgreSQL 17
 - View PostgreSQL backups and clones in the Topology page 18

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:

[PostgreSQL Backup workflow] | [../media/db2_backup_workflow.gif](#)

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 below, 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.



Primary SnapLock settings are managed in SnapCenter backup policy and the secondary SnapLock settings are managed by ONTAP.

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 description.
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 Snapshot page, specify the schedule type 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.

Schedule frequency

Select how often you want the schedules to occur in the policy. The specific times are set at backup job creation enabling you to stagger your start times.

On demand

Hourly

Daily

Weekly

Monthly



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

7. In the Snapshot settings section, specify the number of snapshots that you want to keep.
8. In the Retention page, 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>



For Snapshot copy-based backups, 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.

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

- In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
- In the Resources page, click **New Resource Group**.
- 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>

For this field...	Do this...
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.

Parameter	Setting	Description
ARCHIVE_LOG_EXT	file_extension	Specifies the archive log file extension length. For example, if the archive log is log_backup_0_0_0_0.1615185519429 and if the file_extension value is 5, then the extension of the log will retain 5 digits, which is 16151.
ARCHIVE_LOG_RECURSIVE_SEARCH	(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**.

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:

For Windows: `C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc\allowed_commands_list.txt`



For Linux: `/var/opt/snapcenter/scc/allowed_commands_list.txt`



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

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 .
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:
 - 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. Select  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* 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`

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**.

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 <https://docs.netapp.com/us-en/snapcenter-cmdlets/index.html#SnapCenter Software Cmdlet Reference Guide>].

Back up clusters using PowerShell cmdlets

Backing up a cluster includes establishing a connection with the SnapCenter Server, adding resources, adding a policy, creating a backup resource group, and backing up.

Before you begin

- You must have prepared the PowerShell environment to execute the PowerShell cmdlets.
- You must have added the storage system connection and created a credential.

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_data
01";"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](#).



Monitor backup operations





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.
Activity pane	<ol style="list-style-type: none"> a. After initiating the backup operation, click  on the Activity pane to view the five most recent operations. b. Select the operation. c. 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

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, either select the resource or resource group from the **View** drop-down list.
3. 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 .

Copyright information

Copyright © 2024 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.