



SnapCenter® Software 4.4

Data Protection Guide

For Microsoft® SQL Server®

November 2020 | 215-14667_2020-11_en-us
doccomments@netapp.com

Contents

Deciding on whether to read the SnapCenter Data Protection Guide for Microsoft SQL..... 4

SnapCenter Plug-in for Microsoft SQL Server overview.....5

SnapCenter Plug-in for Microsoft SQL Server data protection workflow.....6

Preparing for data protection..... 8

 Prerequisites for using SnapCenter Plug-in for Microsoft SQL Server..... 8

 Storage types supported by SnapCenter Plug-ins for Microsoft Windows and for Microsoft SQL Server..... 9

 Storage layout recommendations for SnapCenter Plug-in for Microsoft SQL Server..... 10

 How resources, resource groups, and policies are used for protecting SQL Server..... 11

Logging in to SnapCenter..... 13

Backing up SQL Server resources.....15

 Configuring credentials for an individual SQL Server resource..... 16

 Determining whether resources are available for backup..... 17

 Migrating resources to NetApp storage system..... 19

 Creating backup policies for SQL Server databases.....20

 Creating resource groups and attaching policies for SQL Server.....25

 Requirements for backing up SQL resources..... 27

 Backing up SQL resources..... 28

 Backing up SQL Server resource groups.....30

 Monitoring backup operations.....31

 Monitoring operations in the Activity pane.....32

 Canceling the SnapCenter Plug-in for Microsoft SQL Server backup operations..... 32

 Viewing SQL Server backups and clones in the Topology page..... 33

Restoring SQL Server resources.....35

 Requirements for restoring a database.....35

 Restoring SQL Server databases..... 36

 Restoring an SQL Server database from secondary storage.....39

 Reseeding Availability Group databases..... 40

 Monitoring restore operations.....40

 Canceling restore operations.....41

Cloning SQL Server database resources.....43

 Cloning from a SQL Server database backup.....44

 Performing Clone Lifecycle 46

Monitoring clone operations in SnapCenter..... 47

Canceling clone operations..... 48

Splitting a clone..... 49

Backing up, restoring, cloning, and removing backups using PowerShell cmdlets..... 51

Creating a storage system connection and a credential using PowerShell cmdlets..... 51

Backing up resources using PowerShell cmdlets..... 52

Restoring and recovering resources using PowerShell cmdlets..... 53

Cloning backups using PowerShell cmdlets..... 55

Removing backups using PowerShell cmdlets..... 57

Cleaning up the secondary backup count using PowerShell cmdlets..... 58

Managing policies..... 59

Detaching policies..... 59

Modifying policies..... 60

Deleting policies..... 60

Managing resource groups..... 62

Stopping and resuming operations on resource groups..... 62

Deleting resource groups..... 63

Managing backups..... 64

Renaming backups..... 64

Deleting backups..... 64

Managing clones..... 65

Deleting clones..... 65

Copyright and trademark..... 66

Copyright..... 66

Trademark..... 66

Deciding on whether to read the SnapCenter Data Protection Guide for Microsoft SQL

This information describes how to use SnapCenter to perform backup, restore, and clone operations on Microsoft SQL Server resources.

You should read this information if you want to use SnapCenter in the following ways:

- You want to create data protection policies and resource groups for Microsoft SQL Server resources
- You want to perform backup, restore, or clone operations on Microsoft SQL Server resources using the graphical user interface (GUI)
- You want to perform backup, restore, or clone operations on Microsoft SQL Server resources using Windows PowerShell cmdlets

You should have already performed the following tasks:

- Installed SnapCenter Server and the SnapCenter Plug-ins Package for Windows
- Configured role-based access control (RBAC), storage system connections, and credentials
- Deployed the SnapCenter Plug-in for VMware vSphere and registered the plug-in with SnapCenter The SnapCenter Plug-in for VMware vSphere documentation has more information.

[*SnapCenter Plug-in for VMware vSphere Deployment Guide*](#)

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

You can also use the following information to help accomplish your data protection goals:

- SnapCenter Server and plug-in installation and setup
[*Installing and setting up SnapCenter*](#)
[*Getting Started*](#)
- SnapCenter concepts, including architecture, features, and benefits
[*Concepts*](#)
- Other SnapCenter Data Protection Guides for specific types of resources, such as Microsoft SQL Server, Oracle, Windows file systems, and custom plug-ins
- SnapCenter PowerShell cmdlets
[*SnapCenter Software 4.4 Cmdlet Reference Guide*](#)
- SnapCenter administration, including dashboards, reporting capabilities, and REST APIs, and managing licenses, storage connections, and the SnapCenter Server repository
[*Performing administrative tasks*](#)
- For SnapCenter 4.1.1 users, the SnapCenter Plug-in for VMware vSphere 4.1.1 documentation has information on protecting virtualized databases and file systems. For SnapCenter 4.2.x users, the NetApp Data Broker 1.0 and 1.0.1, documentation has information on protecting virtualized databases and file systems using the SnapCenter Plug-in for VMware vSphere that is provided by the Linux-based NetApp Data Broker virtual appliance (Open Virtual Appliance format). For SnapCenter 4.3.x users, the SnapCenter Plug-in for VMware vSphere 4.3 documentation has information on protecting virtualized databases and file systems using the Linux-based SnapCenter Plug-in for VMware vSphere virtual appliance (Open Virtual Appliance format).

[*SnapCenter Plug-in for VMware vSphere Data Protection Guide*](#)

SnapCenter Plug-in for Microsoft SQL Server overview

The SnapCenter Plug-in for Microsoft SQL Server is a host-side component of the NetApp SnapCenter Software that enables application-aware data protection management of Microsoft SQL Server databases. The Plug-in for SQL Server automates SQL Server database backup, verify, restore, and cloning operations in your SnapCenter environment.

Success story: "We have many Microsoft SQL databases in production with Agile development currently using DB copies that are over a week old and taking 5 to 6 hours to provision. With SnapCenter, cloning takes minutes to provision a Dev environment that is typically one day old."

SQL DBA

When the Plug-in for SQL Server is installed, you can use SnapCenter with NetApp SnapMirror technology to create mirror copies of backup sets on another volume and with NetApp SnapVault technology to perform disk-to-disk backup replication for standards compliance or archival purposes.

The Data Protection Guide for your plug-in has information about data protection operations.

The SnapCenter concepts documentation has information about the SnapCenter architecture, features, and benefits.

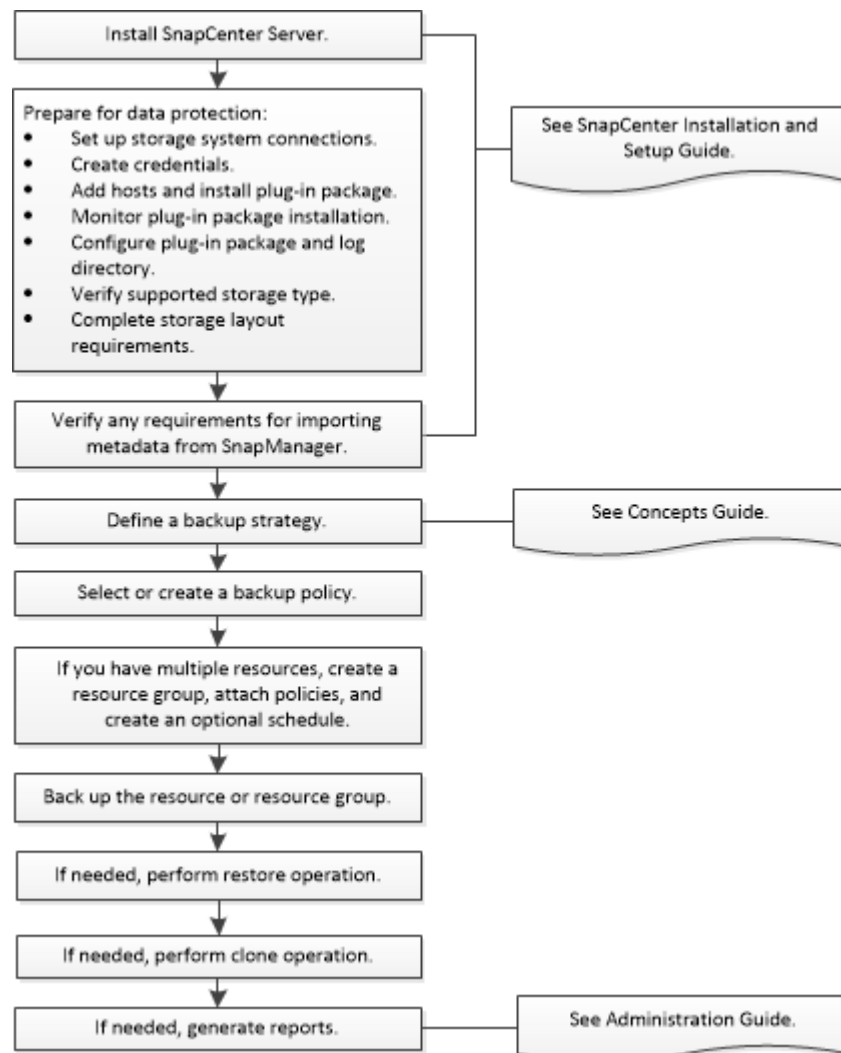
Related information

[*Concepts*](#)

[*Installing and setting up SnapCenter*](#)

SnapCenter Plug-in for Microsoft SQL Server data protection workflow

Before you can use SnapCenter to protect Microsoft SQL Server databases, your SnapCenter Server administrator must have installed the SnapCenter Plug-in Package for Windows.



Related reference

[Backing up SQL Server resources](#) on page 15

When you install the SnapCenter Plug-in for Microsoft SQL Server in your environment, you can use SnapCenter to back up the SQL Server resources.

[Restoring SQL Server resources](#) on page 35

You can use SnapCenter to restore SQL Server databases by restoring the data from one or more backups to your active file system and then recovering the database. You can also restore databases that are in Availability Groups and then add the restored databases to the Availability Group.

Before restoring an SQL Server database, you must perform several preparatory tasks.

[Cloning SQL Server database resources](#) on page 43

You must perform several tasks using SnapCenter Server before cloning database resources from a backup. Database cloning is the process of creating a point-in-time copy of a production database

or its backup set. You can clone databases to test functionality that has to be implemented using the current database structure and content during application development cycles, to use the data extraction and manipulation tools when populating data warehouses, or to recover data that was mistakenly deleted or changed.

Related information

[Installing and setting up SnapCenter](#)

[Concepts](#)

[Performing administrative tasks](#)

Preparing for data protection

Before performing any data protection operation such as backup, clone, or restore operations, you must define your strategy and set up the environment. You can also set up the SnapCenter Server to use SnapMirror and SnapVault technology.

To take advantage of SnapVault and SnapMirror technology, you must configure and initialize a data protection relationship between the source and destination volumes on the storage device. You can use NetApp System Manager or you can use the storage console command line to perform these tasks.

Prerequisites for using SnapCenter Plug-in for Microsoft SQL Server

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

- Install and configure SnapCenter Server.
- Log in to SnapCenter.
- Configure the SnapCenter environment by adding or assigning storage system connections and creating credentials.

Note: SnapCenter does not support multiple SVMs with the same name on different clusters. Each SVM supported by SnapCenter must have a unique name.

- Add hosts, install the plug-ins, discover (refresh) the resources, and configure the plug-ins.
- Move an existing Microsoft SQL Server database from a local disk to a NetApp LUN or vice versa by running `Invoke-SmConfigureResources`.

For information to run the cmdlet, see the information on Windows cmdlets.

[SnapCenter Software 4.4 Cmdlet Reference Guide](#)

- If you are using SnapCenter Server to protect SQL databases that reside on VMware RDM LUNs or VMDKs, you must deploy the SnapCenter Plug-in for VMware vSphere and register the plug-in with SnapCenter. The SnapCenter Plug-in for VMware vSphere documentation has more information.

[SnapCenter Plug-in for VMware vSphere Deployment Guide](#)

- Perform host-side storage provisioning using the SnapCenter Plug-in for Microsoft Windows.
- Move existing databases onto NetApp storage.
For details, see SnapCenter importing information.
- Set up SnapMirror and SnapVault relationships, if you want backup replication.

For details, see SnapCenter installation information.

For SnapCenter 4.1.1 users, the SnapCenter Plug-in for VMware vSphere 4.1.1 documentation has information on protecting virtualized databases and file systems. For SnapCenter 4.2.x users, the NetApp Data Broker 1.0 and 1.0.1, documentation has information on protecting virtualized databases and file systems using the SnapCenter Plug-in for VMware vSphere that is provided by the Linux-based NetApp Data Broker virtual appliance (Open Virtual Appliance format). For SnapCenter 4.3.x users, the SnapCenter Plug-in for VMware vSphere 4.3 documentation has information on protecting virtualized databases and file systems using the Linux-based SnapCenter Plug-in for VMware vSphere virtual appliance (Open Virtual Appliance format).

[SnapCenter Plug-in for VMware vSphere Data Protection Guide](#)

Related information

[Installing and setting up SnapCenter](#)

Storage types supported by SnapCenter Plug-ins for Microsoft Windows and for Microsoft SQL Server

SnapCenter supports a wide range of storage types on both physical machines and virtual machines. You must verify whether support is available for your storage type before installing the package for your host.

SnapCenter provisioning and data protection support is available on Windows Server. For the latest information about supported versions, see the NetApp Interoperability Matrix Tool (IMT).

[NetApp Interoperability Matrix Tool](#)

Machine	Storage type	Provision using	Support notes
Physical server	FC-connected LUNs	SnapCenter graphical user interface (GUI) or PowerShell cmdlets	
	iSCSI-connected LUNs	SnapCenter GUI or PowerShell cmdlets	
	SMB3 (CIFS) shares residing on a storage virtual machine (SVM)	SnapCenter GUI or PowerShell cmdlets	Support for provisioning only. You cannot use SnapCenter to back up any data or shares using the SMB protocol.
VMware VM	RDM LUNs connected by an FC or iSCSI HBA	PowerShell cmdlets	
	iSCSI LUNs connected directly to the guest system by the iSCSI initiator	SnapCenter GUI or PowerShell cmdlets	
	Virtual Machine File Systems (VMFS) on VMFS or NFS datastores	VMware vSphere or VSC cloning utility	
	A guest system connected to SMB3 shares residing on an SVM	SnapCenter GUI or PowerShell cmdlets	Support for provisioning only. You cannot use SnapCenter to back up any data or shares using the SMB protocol.
Hyper-V VM	Virtual FC (vFC) LUNs connected by a virtual Fibre Channel Switch	SnapCenter GUI or PowerShell cmdlets	You must use Hyper-V Manager to provision Virtual FC (vFC) LUNs connected by a virtual Fibre Channel Switch.
	iSCSI LUNs connected directly to the guest system by the iSCSI initiator	SnapCenter GUI or PowerShell cmdlets	
	A guest system connected to SMB3 shares residing on an SVM	SnapCenter GUI or PowerShell cmdlets	Support for provisioning only. You cannot use SnapCenter to back up any data or shares using the SMB protocol.
	Note: Hyper-V pass through disks and backing up databases on VHD(x) that are provisioned on NetApp storage are not supported.		

Related information

[Installing and setting up SnapCenter](#)

Storage layout recommendations for SnapCenter Plug-in for Microsoft SQL Server

A well-designed storage layout allows SnapCenter Server to back up your databases to meet your recovery objectives. You should consider several factors while defining your storage layout, including the size of the database, the rate of change of the database, and the frequency with which you perform backups.

The following sections define the storage layout recommendations and restrictions for LUNs and virtual machine disks (VMDKs) with SnapCenter Plug-in for Microsoft SQL Server installed in your environment.

In this case, LUNs can include VMware RDM disks and iSCSI direct-attached LUNs that are mapped to the guest.

Note: VMDK and physical raw device mapping (RDM) disk mapped to a single VM guest is not supported.

LUN and VMDK requirements

You can optionally use dedicated LUNs or VMDKs for optimum performance and management for the following databases:


- Master and model system databases
- Tempdb
- User database files (.mdf and .ndf)
- User database transaction log files (.ldf)
- Log directory

To restore large databases, the best practice is to use dedicated LUNs or VMDKs. The time taken to restore a complete LUN or VMDK is less than the time taken to restore the individual files that are stored in the LUN or VMDK.


For the log directory, you should create a separate LUN or VMDK so that there is sufficient free space in the data or log file disks.

LUN and VMDK sample layouts


The following graphic shows how you can configure the storage layout for large databases on LUNs:




The following graphic shows how you can configure the storage layout for medium or small databases on LUNs:



The following graphic shows how you can configure the storage layout for large databases on VMDKs:



The following graphic shows how you can configure the storage layout for medium or small databases on VMDKs:



How resources, resource groups, and policies are used for protecting SQL Server

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 databases, database instances, or Microsoft SQL Server availability groups that you back up or clone with SnapCenter.
- A SnapCenter *resource group*, is a collection of resources on a host or cluster.

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 resource groups were formerly known as *datasets*.

- The *policies* specify the backup frequency, copy retention, 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 databases or backing up all file systems of a host, for example, you might create a resource group that includes all the databases or all the file systems 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 and another schedule that performs log backups hourly.

The following image illustrates the relationship between resources, resource groups, and policies for databases:



Logging in to SnapCenter

SnapCenter supports role-based access control (RBAC). SnapCenter admin assigns roles and resources through SnapCenter RBAC to either a user in workgroup or active directory, or to groups in active directory. The RBAC user can now log in to SnapCenter with the assigned roles.

Before you begin

- You should enable Windows Process Activation Service (WAS) in Windows Server Manager.
- If you want to use Internet Explorer as the browser to log in to the SnapCenter Server, you should ensure that the Protected Mode in Internet Explorer is disabled.

About this task

During installation, the SnapCenter Server Install wizard creates a shortcut and places it on the desktop and in the Start menu of the host where SnapCenter is installed. Additionally, at the end of the installation, the Install wizard displays the SnapCenter URL based on the information that you provided during installation, which you can copy if you want to log in from a remote system.



Attention: If you have multiple tabs open in your web browser, closing just the SnapCenter browser tab does not log you out of SnapCenter. To end your connection with SnapCenter, you must log out of SnapCenter either by clicking the **Sign out** button, or by closing the entire web browser.

Best Practice: For security reasons, it is recommended that you do not enable your browser to save your SnapCenter password.

The default GUI URL is a secure connection to the default port 8146 on the server where the SnapCenter Server is installed (`https://server:8146`). If you provided a different server port during the SnapCenter installation, that port is used instead.

For Network Load Balance (NLB) deployment, you must access SnapCenter using the NLB cluster IP (`https://NLB_Cluster_IP:8146`). If you do not see the SnapCenter UI when you navigate to `https://NLB_Cluster_IP:8146` in Internet Explorer (IE), you must add the NLB IP address as a trusted site in IE on each plug-in host, or you must disable IE Enhanced Security on each plug-in host.

- *SnapCenter in an HA configuration with Application Request Routing (ARR) enabled exhibits backup jobs in a perpetually 'running' state.*
- *Unable to access cluster IP address from outside network*

In addition to using the SnapCenter GUI, you can use PowerShell cmdlets to create scripts to perform configuration, backup, and restore operations. Some cmdlets might have changed with each SnapCenter release. The SnapCenter cmdlet or SnapCenter CLI documentation has the details.

Note: If you are logging in to SnapCenter for the first time, you must log in using the credentials that you provided during the install process.

Steps

1. Launch SnapCenter from the shortcut located on your local host desktop, or from the URL provided at the end of the installation, or from the URL provided by your SnapCenter administrator.
2. Enter user credentials.

To specify the following...	Use one of these formats...
Domain administrator	<i>NetBIOS\UserName</i> <i>UserName@UPN suffix</i> For example, username@netapp.com <i>Domain FQDN\UserName</i>
Local administrator	<i>UserName</i>

3. If you are assigned more than one role, from the **Role** box, select the role that you want to use for this login session.
Your current user and associated role are shown in the upper right of SnapCenter after you are logged in.

Result

If you are using SnapCenter for the first time, the Storage Systems page is displayed, and the Get Started pane is expanded.

If the logging fails with the error that site cannot be reached, you should map the SSL certificate to SnapCenter.

[Site cannot be reached](#)

After logging to SnapCenter Server for the first time, refresh the resources list.

After logging to SnapCenter Server for the first time, the SnapCenter Server Configuration Checker schedule is created. The default values are Weekly and Every Sunday at 11:59 pm. To modify the schedule or run the SnapCenter Server schedule, click **Settings > Scheduled Configuration Checker**.

After you finish

If you have untrusted Active Directory domains that you want SnapCenter to support, you must register those domains with SnapCenter before configuring the roles for the users on untrusted domains. The administration documentation has more details.

[Performing administrative tasks](#)

Related information

[SnapCenter Software 4.4 Cmdlet Reference Guide](#)

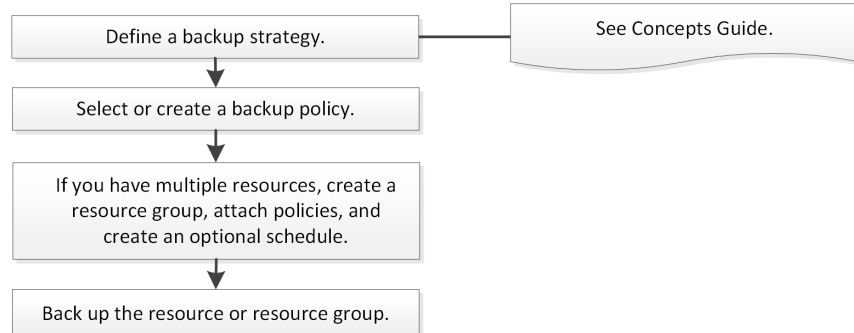
Backing up SQL Server resources

When you install the SnapCenter Plug-in for Microsoft SQL Server in your environment, you can use SnapCenter to back up the SQL Server resources.

You can schedule multiple backups to run across servers simultaneously.

Backup and restore operations cannot be performed simultaneously on the same resource.

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



Note: The Backup Now, Restore, Manage Backups, and Clone options on the Resources page are disabled if you select a non-NetApp LUN, a database that is corrupted, or a database that is being restored.

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

[SnapCenter Software 4.4 Cmdlet Reference Guide](#)

How SnapCenter backs up databases

SnapCenter uses Snapshot copy technology to back up the SQL Server databases that reside on LUNs or VMDKs. SnapCenter creates the backup by creating Snapshot copies of the databases.

When you select a database for a full database backup from the Resources page, SnapCenter automatically selects all the other databases that reside on the same storage volume. If the LUN or VMDK stores only a single database, you can clear or reselect the database individually. If the LUN or VMDK houses multiple databases, you must clear or reselect the databases as a group.

All the databases that reside on a single volume are backed up concurrently using Snapshot copies. If the maximum number of concurrent backup databases is 35, and if more than 35 databases reside in a storage volume, then the total number of Snapshot copies that are created equals the number of databases divided by 35.

Note: You can configure the maximum number of databases for each Snapshot copy in the backup policy.

When SnapCenter creates a Snapshot copy, the entire storage system volume is captured in the Snapshot copy. However, the backup is valid only for the SQL host server for which the backup was created.

If data from other SQL host servers resides on the same volume, this data cannot be restored from the Snapshot copy.

Related tasks

[Logging in to SnapCenter](#) on page 13

SnapCenter supports role-based access control (RBAC). SnapCenter admin assigns roles and resources through SnapCenter RBAC to either a user in workgroup or active directory, or to groups in active directory. The RBAC user can now log in to SnapCenter with the assigned roles.

Backing up resources using PowerShell cmdlets on page 52

You can use the PowerShell cmdlets to backup SQL Server databases or Windows file systems. This would include backing up a SQL Server database or Windows file system includes establishing a connection with the SnapCenter Server, discovering the SQL Server database instances or Windows file systems, adding a policy, creating a backup resource group, backing up, and verifying the backup.

Related information

Quiesce or grouping resources operations fail

Configuring credentials for an individual SQL Server resource

You can configure credentials to perform data protection jobs on individual SQL Server resources for each user. While you can configure the credentials globally, you might want to do this only for a particular resource.

About this task

- If you are using Windows credentials for authentication, you must set up your credential before installing plug-ins.
However, if you are using an SQL Server instance for authentication, you must add the credential after installing plug-ins.
- If you have enabled SQL authentication while setting up the credentials, the discovered instance or database is shown with a red color lock icon.
If the lock icon appears, you must specify the instance or database credentials to successfully add the instance or database to a resource group.
- You must assign the credential to a role-based access control (RBAC) user without sysadmin access when the following conditions are met:
 - The credential is assigned to an SQL instance.
 - The SQL instance or host is assigned to an RBAC user.

The user must have both the resource group and backup privileges.

Steps

1. In the left navigation pane, click **Settings**.
2. In the **Settings** page, click **Credential**.
3. To add a new credential, click **New**.
4. In the **Credential** page, configure the credentials:



For this field...	Do this...
Credential name	Enter a name for the credentials.

For this field...	Do this...
Username	<p>Enter the user name used for SQL Server authentication.</p> <ul style="list-style-type: none"> Domain administrator or any member of the administrator group 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 <i>Username</i> field are: <code>NetBIOS\UserName</code> <code>Domain FQDN\UserName</code> Local administrator (for workgroups only) 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 <i>Username</i> field is: <code>UserName</code>
Password	Enter the password used for authentication.
Authentication mode	<p>Select the SQL Server authentication mode.</p> <p>You can also choose Windows authentication if the Windows user has sysadmin privileges on the SQL server.</p>
Host	Select the host.
SQL Server Instance	Select the SQL Server Instance.

5. Click **OK** to add the credential.

6. In the left navigation pane, click **Resources**.

7. In the **Resources** page, select **Instance** from the **View** list.

- Click , and then select the host name to filter the instances.
- Click  to close the filter pane.

Note: The credential option does not apply to databases and availability groups.

8. In the **Instance Protect** page, protect the instance, and if required, click **Configure Credentials**.

If the user who is logged in to the SnapCenter Server does not have access to SnapCenter Plug-in for Microsoft SQL Server, then the user has to configure the credentials.

9. Click **Refresh Resources**.

Determining whether resources are available for backup

Resources are the databases, application instances, Availability Groups, and similar components that are maintained by the plug-ins you have installed. You can add those resources to resource groups so that you can perform data protection jobs, but first you must identify which resources you have available. Determining available resources also verifies that the plug-in installation has completed successfully.

Before you begin

- You must have already completed tasks such as installing SnapCenter Server, adding hosts, creating storage system connections, and adding credentials.

- To discover the Microsoft SQL databases, one of the following conditions should be met.
 - The user who has logged into SnapCenter should have the required permissions (sysadmin) on the Microsoft SQL Server.
 - The user that was used to add the plug-in host to SnapCenter Server should have the required permissions (sysadmin) on the Microsoft SQL Server.
 - If the above two conditions are not met, in the SnapCenter Server you should configure the user that has the required permissions (sysadmin) on the Microsoft SQL Server. The user should be configured at the Microsoft SQL Server instance level and the user can be a SQL or Windows user.
- To discover the Microsoft SQL databases in a Windows cluster, you must unblock the Failover Cluster Instance (FCI) TCP/IP port.
- If databases reside on VMware RDM LUNs or VMDKs, you must deploy the SnapCenter Plug-in for VMware vSphere and register the plug-in with SnapCenter. The SnapCenter Plug-in for VMware vSphere documentation has more information.

[SnapCenter Plug-in for VMware vSphere Data Protection Guide](#)



About this task

You cannot back up databases when the **Overall Status** option in the Details page is set to `Not available for backup`. The **Overall Status** option is set to `Not available for backup` when any of the following is true:

- Databases are not on a NetApp LUN.
- Databases are not in normal state.
Databases are not in normal state when they are offline, restoring, recovery pending, suspect, and so on.
- Databases have insufficient privileges.
For example, if a user has only view access to the database, files and properties of the database cannot be identified and hence cannot be backed up.

Note: SnapCenter can backup only the primary database if you have a availability group configuration on SQL Server Standard Edition.

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 **Database**, or **Instance**, or **Availability Group**, from the **View** drop-down list.
Click  and select the host name and the SQL Server Instance to filter the resources. You can then click  to close the filter pane.
3. Click **Refresh Resources**.
The newly added, renamed, or deleted resources are updated to the SnapCenter Server inventory.

Result

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

- If the database is on a non NetApp storage, `Not available for backup` is displayed in the Overall Status column.
You cannot perform data protection operations on a database that is on a non NetApp storage.
- If the database is on a NetApp storage and not protected, `Not protected` is displayed in the Overall Status column.

- If the database is on a NetApp storage system and protected, the user interface displays `Backup not run` message in the Overall Status column.
- If the database is on a NetApp storage system and protected and if the backup is triggered for the database, the user interface displays `Backup succeeded` message in the Overall Status column.

Note: If you have enabled an SQL authentication while setting up the credentials, the discovered instance or database is shown with a red color lock icon. If the lock icon appears, you must specify the instance or database credentials for successfully adding the instance or database to a resource group.

After you finish

After the SnapCenter administrator assigns the resources to a RBAC user, the RBAC user must log in and click **Refresh Resources** to see the latest Overall Status of the resources.

Migrating resources to NetApp storage system

After you have provisioned your NetApp storage system using SnapCenter Plug-in for Microsoft Windows, you can migrate your resources to the NetApp storage system or from one NetApp LUN to another NetApp LUN using either the SnapCenter graphical user interface (GUI) or using the PowerShell cmdlets.

Before you begin

- You must have added storage systems to SnapCenter Server.
- You must have refreshed (discovered) the SQL Server resources.

About this task

Most of the fields on these wizard pages are self-explanatory. The following information describes some of the fields for which you might require guidance.

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 **Database** or **Instance** from the **View** drop-down list.
3. Select either the database or the instance from the list and click **Migrate**.
4. In the **Resources** page, perform the following actions:

For this field...	Do this...
Database Name (optional)	If you have selected an instance for migration, you must select the databases of that instance from the Databases drop-down list.
Choose Destinations	Select the target location for data and log files. The data and log files are moved to Data and Log folder respectively under the selected NetApp drive. If any folder in the folder structure is not present, then a folder is created, and the resource is migrated.
Show database file details (optional)	Select this option when you want to migrate multiple files of a single database. Note: This option is not displayed when you select the Instance resource.
Options	Select Delete copy of Migrated Database at Original Location to delete copy of database from the source. Optional: RUN UPDATE STATISTICS on tables before detaching the database.

5. In the **Verify** page, perform the following actions:

For this field...	Do this...
Database Consistency Check Options	Select Run before to check the integrity of the database before migration. Select Run after to check the integrity of the database after migration.
DBCC CHECKDB options	Select PHYSICAL_ONLY option to limit the integrity check to the physical structure of the database and to detect torn pages, checksum failures, and common hardware failures that impact the database. Select NO_INFOMSGS option to suppress all of the informational messages. Select ALL_ERRORMSGs option to display all of the reported errors per object. Select NOINDEX option if you do not want to check nonclustered indexes. The SQL Server database uses Microsoft SQL Server Database Consistency Checker (DBCC) to check the logical and physical integrity of the objects in the database. Note: You might want to select this option to decrease the execution time. Select TABLOCK option to limit the checks and obtain locks instead of using an internal database Snapshot copy.

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

Related information

[SnapCenter Software 4.4 Cmdlet Reference Guide](#)

Creating backup policies for SQL Server databases

You can create a backup policy for the resource or the resource group before you use SnapCenter to back up SQL Server resources, or you can create a backup policy at the time you create a resource group or backup a single resource.

Before you begin

- You must have defined your data protection strategy.
For details, see the information about defining a data protection strategy for SQL databases.
[Concepts](#)
- You must have prepared for data protection by completing tasks such as installing SnapCenter, adding hosts, identifying resources, and creating storage system connections.
- You must have configured the host log directory for log backup.
- You must have refreshed (discovered) the SQL Server resources.
- If you are replicating Snapshot copies to a mirror or vault, the SnapCenter administrator must have assigned the storage virtual machines (SVMs) for both the source volumes and destination volumes to you.
For information about how administrators assign resources to users, see the SnapCenter installation information.
- If you want to run the PowerShell scripts in prescripts and postscripts, you should set the value of the `usePowershellProcessforScripts` parameter to **true** in the `web.config` file.
The default value is **false**.

About this task

A backup policy is a set of rules that governs how you manage and retain backups, and how frequently the resource or resource group is backed up. Additionally, you can specify replication and script settings. Specifying options in a policy saves time when you want to reuse the policy for another resource group.

Most of the fields on these wizard pages are self-explanatory. The following information describes some of the fields for which you might require guidance.

Steps

1. In the left navigation pane, click **Settings**.
2. In the **Settings** page, click **Policies**.
3. Click **New**.
4. On the **Name** page, enter the policy name and description.
5. On the **Backup Type** page, perform the following steps:
 - a. Choose the backup type:

If you want to...	Do this...
Back up the database files and truncate the transaction logs	<ol style="list-style-type: none"> i. Select Full backup and Log backup. ii. Enter the maximum number of databases that should be backed up for each Snapshot copy. Note: You must increase this value if you want to run multiple backup operations concurrently.
Back up the database files	<ol style="list-style-type: none"> i. Select Full backup. ii. Enter the maximum number of databases that should be backed up for each Snapshot copy. Note: You must increase this value if you want to run multiple backup operations concurrently.
Back up the transaction logs	Select Log backup .

- b. If you are backing up your resources by using another backup application, select **Copy only backup**.

Keeping the transaction logs intact allows any backup application to restore the databases. You typically should not use the copy only option in any other circumstance.

Note: Microsoft SQL does not support the **Copy only backup** option together with the **Full backup and Log backup** option for secondary storage.

- c. In the **Availability Group Settings** section, perform the following actions:

For this field...	Do this...
Backup on preferred backup replica only	Select this option to backup only on preferred backup replica. The preferred backup replica is decided by the backup preferences configured for the AG in the SQL Server.
Select replicas for backup	Choose the primary AG replica or the secondary AG replica for the backup.
Backup priority (Minimum and Maximum backup priority)	Specify a minimum backup priority number, and a maximum backup priority number that decide the AG replica for backup. For example, you can have a minimum priority of 10 and a maximum priority of 50. In this case, all the AG replicas with a priority more than 10 and less than 50 are considered for backup.

Note: In cluster configurations, the backups are retained at each node of the cluster according to the retention settings set in the policy. If the owner node of the AG

changes, the backups are taken according to the retention settings and the backups of the previous owner node will be retained. The retention for AG is applicable only at the node level.

- d. If you want to schedule the backup that you want to create with this policy, specify the schedule type by selecting either **On demand**, **Hourly**, **Daily**, **Weekly**, or **Monthly**.

You can select one schedule type for a 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

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

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

- 6. On the **Retention** page, depending on the backup type selected in the backup type page, perform one or more of the following actions:
 - a. In the Retention settings for the up-to-the-minute restore operation section, perform one of the following actions:

If you want to...	Do this...
Retain only a specific number of Snapshot copies	Select the Keep log backups applicable to last <number> full backups option, and specify the number of Snapshot copies to be retained. If you near this limit, you might want to delete older copies.
Retain the backup copies for a specific number of days	Select the Keep log backups applicable to last <number> days of full backups option, and specify the number of days to keep the log backup copies.

- b. In the Full backup retentions settings section for the On Demand retention settings, perform the following actions:

For this field...	Do this...
Total Snapshot copies to keep	<p>If you want to specify the number of Snapshot copies to keep, select Total Snapshot copies to keep.</p> <p>If the number of Snapshot copies exceeds the specified number, the Snapshot copies are deleted with the oldest copies deleted first.</p> <p>Note: The maximum retention value is 1018 for resources on ONTAP 9.4 or later, and 254 for resources on ONTAP 9.3 or earlier. Backups will fail if retention is set to a value higher than what the underlying ONTAP version supports.</p> <p>Important: 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 copy is the reference Snapshot copy for the SnapVault relationship until a newer Snapshot copy is replicated to the target.</p>
Keep Snapshot copies for	If you want specify the number of days for which you want to keep the Snapshot copies before deleting them, select Keep Snapshot copies for .

- c. In the Full backup retentions settings section for the Hourly, Daily, Weekly and Monthly retention settings, specify the retention settings for the schedule type selected on Backup Type page.

For this field...	Do this...
Total Snapshot copies to keep	<p>If you want to specify the number of Snapshot copies to keep, select Total Snapshot copies to keep.</p> <p>If the number of Snapshot copies exceeds the specified number, the Snapshot copies are deleted with the oldest copies deleted first.</p> <p>Important: 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 copy is the reference Snapshot copy for the SnapVault relationship until a newer Snapshot copy is replicated to the target.</p>
Keep Snapshot copies for	If you want specify the number of days for which you want to keep the Snapshot copies before deleting them, select Keep Snapshot copies for .

The log Snapshot copy retention is set to 7 days by default. Use `Set-SmPolicy` cmdlet to change the log Snapshot copy retention.

This example sets the log Snapshot copy retention to 2:

```
Set-SmPolicy -PolicyName 'newpol' -PolicyType 'Backup' -PluginPolicyType 'SCSQL' -
sqlbackuptype 'FullBackupAndLogBackup' -RetentionSettings
@{BackupType='DATA';ScheduleType='Hourly';RetentionCount=2},@{BackupType='LOG_SNAPSHOT
';ScheduleType='None';RetentionCount=2},@{BackupType='LOG';ScheduleType='Hourly';Reten
tionCount=2} -scheduletype 'Hourly'
```

7. On the **Replication** page, specify replication to the secondary storage system:

For this field...	Do this...
Update SnapMirror after creating a local Snapshot copy	Select this option to create mirror copies of backup sets on another volume (SnapMirror).

For this field...	Do this...
Update SnapVault after creating a Snapshot copy	Select this option to perform disk-to-disk backup replication.
Secondary policy label	<p>Select a Snapshot label.</p> <p>Depending on the Snapshot copy label that you select, ONTAP applies the secondary Snapshot copy retention policy that matches the label.</p> <p>Note: If you have selected Update SnapMirror after creating a local Snapshot copy, you can optionally specify the secondary policy label. However, if you have selected Update SnapVault after creating a local Snapshot copy, you should specify the secondary policy label.</p>
Error retry count	Enter the number of replication attempts that should occur before the process halts.

8. Optional: On the **Script** page, enter the path and the arguments of the prescript or postscript that should be run before or after the backup operation, respectively.

For example, you can run a script to update SNMP traps, automate alerts, and send logs.

Note: You must configure the SnapMirror retention policy in ONTAP so that the secondary storage does not reach the maximum limit of Snapshot copies.

9. On the **Verification** page, perform the following steps:
- In the **Run verification for following backup schedules** section, select the schedule frequency.
 - In the **Database consistency check options** section, perform the following actions:

For this field...	Do this...
Limit the integrity structure to physical structure of the database (PHYSICAL_ONLY)	Select Limit the integrity structure to physical structure of the database (PHYSICAL_ONLY) to limit the integrity check to the physical structure of the database and to detect torn pages, checksum failures, and common hardware failures that impact the database.
Suppress all information messages (NO_INFOMSGS)	Select Suppress all information messages (NO_INFOMSGS) to suppress all informational messages. Selected by default.
Display all reported error messages per object (ALL_ERRORMSGs)	Select Display all reported error messages per object (ALL_ERRORMSGs) to display all the reported errors per object.
Do not check nonclustered indexes (NOINDEX)	<p>Select Do not check nonclustered indexes (NOINDEX) if you do not want to check nonclustered indexes.</p> <p>The SQL Server database uses Microsoft SQL Server Database Consistency Checker (DBCC) to check the logical and physical integrity of the objects in the database.</p>

For this field...	Do this...
Limit the checks and obtain the looks instead of using an internal database Snapshot copy (TABLOCK)	Select Limit the checks and obtain the looks instead of using an internal database Snapshot copy (TABLOCK) to limit the checks and obtain locks instead of using an internal database Snapshot copy.

- c. In the **Log Backup** section, select **Verify log backup upon completion** to verify the log backup upon completion.
- d. In the **Verification script settings** section, enter the path and the arguments of the prescript or postscript that should be run before or after the verification operation, respectively.

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

Related information

[Performing administrative tasks](#)

[Installing and setting up SnapCenter](#)

[SnapCenter retains Snapshot copies of the database](#)

Creating resource groups and attaching policies for SQL Server

A resource group is a container to which you add resources that you want to back up and protect together. A resource group enables you to back up all of 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

You can protect resources individually without creating a new resource group. You can take backups on the protected resource.

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 **Database** from the **View** list.

Note: If you have recently added a resource to SnapCenter, click **Refresh Resources** to view the newly added resource.

3. Click **New Resource Group**.
4. On the **Name** page, perform the following actions:

For this field...	Do this...
Name	Enter the resource group name. Note: The resource group name should not exceed 250 characters.
Tags	Enter one or more labels that will help you later search for the resource group. 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.

For this field...	Do this...
Use custom name format for Snapshot copy	Optional: Enter a custom Snapshot copy name and format. For example, <i>customtext_resourcegroup_policy_hostname</i> or <i>resourcegroup_hostname</i> . By default, a timestamp is appended to the Snapshot copy name.

5. On the **Resources** page, perform the following steps:


- a. Select the host name, resource type, and the SQL Server instance from drop-down lists to filter the list of resources.

Note: If you have recently added resources, they will appear on the list of Available Resources only after you refresh your resource list.


- b. To move resources from the **Available Resources** section to the **Selected Resources** section, perform one of the following steps:
 - Select **Autoselect all resources on same storage volume** to move all of the resources on the same volume to the Selected Resources section.
 - Select the resources from the Available Resources section and then click the right arrow to move them to the Selected Resources section.

6. On the **Policies** page, perform the following steps:

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

Note: You can also create a policy by clicking .

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

- b. In the **Configure schedules for selected policies** section, click  in the **Configure Schedules** column for the policy for which you want to configure the schedule.
- c. In the **Add schedules for policy policyname** *policy_name* dialog box, configure the schedule by specifying the start date, expiration date, and frequency, and then click **OK**. You must do this for each frequency listed in the policy. The configured schedules are listed in the Applied Schedules column in the Configure schedules for selected policies section.
- d. Select the **Microsoft SQL Server scheduler**.

You must also select a scheduler instance to associate with the scheduling policy.

If you do not select Microsoft SQL Server scheduler, the default is Microsoft Windows scheduler.

Third party backup schedules are not supported when they overlap with SnapCenter backup schedules. You should not modify the schedules from the Windows task scheduler and SQL Server Agent.

7. On the **Verification** page, perform the following steps:

- a. Select the verification server from the **Verification server** drop-down list.

The list includes all the SQL Servers added in SnapCenter. You can select multiple verification servers (local host or remote host).

Note: The verification server version should match the version and edition of the SQL server that is hosting the primary database.

- b. Click **Load locators** to load the SnapMirror and SnapVault volumes to perform verification on secondary storage.

- c. Select the policy for which you want to configure your verification schedule, and then click



- d. 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 .

- e. Click **OK**.

The configured schedules are listed in the Applied Schedules column. You can review and



then edit by clicking or delete by clicking .

8. On 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**.

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

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

Related tasks

[Creating backup policies for SQL Server databases](#) on page 20

You can create a backup policy for the resource or the resource group before you use SnapCenter to back up SQL Server resources, or you can create a backup policy at the time you create a resource group or backup a single resource.

Related information

[Performing administrative tasks](#)

Requirements for backing up SQL resources

Before you backup a SQL resource, you must ensure that several requirements are met.

- You must have migrated a resource from a non-NetApp storage system to a NetApp storage system.
- You must have created a backup policy.
- If you want to back up a resource that has a SnapMirror relationship to 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.
- The backup operation initiated by an active directory (AD) user fails if the SQL instance credential is not assigned to the AD user or group. You must assign the SQL instance credential to AD user or group from the **Settings > User Access** page.
- You must have created a resource group with a policy attached.
- If a resource group has multiple databases from different hosts, the backup operation on some hosts might be triggered late because of network issues. You should configure the value of `ForUninitializedHosts` in `web.config` by using the `Set SmConfigSettings PS cmdlet`.

Backing up SQL resources

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

Steps


1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. On the **Resources** page, select **Database**, or **Instance**, or **Availability Group** from the **View** drop-down list.
 - a. Select the database, or instance, or availability group that you want to back up.

When you take a backup of an instance, the information about the last backup status or the timestamp of that instance will not be available in the resources page.


In the topology view, you cannot differentiate whether the backup status, timestamp, or backup is for an instance or a database.
3. On the **Resources** page, select the **custom name format for Snapshot copy** check box, and then enter a custom name format that you want to use for the Snapshot copy name.

For example, *customtext__policy_hostname* or *resource_hostname*. By default, a timestamp is appended to the Snapshot copy name.
4. On the **Policies** page, perform the following tasks:

- a. In the **Policies** section, select one or more policies from the drop-down list.

You can create a policy by clicking  to start the policy wizard.


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* dialog box, configure the schedule, and then click **OK**.

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

The configured schedules are listed in the **Applied Schedules** column.
 - d. Select the **Use Microsoft SQL Server scheduler**, and then select the scheduler instance from the **Scheduler Instance** drop-down list that is associated with the scheduling policy.
5. On the **Verification** page, perform the following steps:
- a. Select the verification server from the **Verification server** drop-down list.

You can select multiple verification servers (local host or remote host).

Note: The verification server version should match the version and edition of the SQL server that is hosting the primary database.
 - b. Select **Load secondary locators to verify backups on secondary** to verify your backups on secondary storage system.
 - c. Select the policy for which you want to configure your verification schedule, and then click .
 - d. 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 .

If you want to...	Do this...
Schedule a verification	Select Run scheduled verification .

Note: If the verification server does not have a storage connection, the verification operation fails with error: Failed to mount disk.

- e. Click **OK**.

The configured schedules are listed in the Applied Schedules column.

6. On 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**.

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

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

The database topology page is displayed.

8. Click **Back up Now**.

9. On the **Backup** page, perform the following steps:

- a. If you have 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 **Verify after backup** to verify your backup.

- c. Click **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.

An implicit resource group is created. You can view this by selecting the **Resource Group** from the Asset drop-down list on the User Access page. The implicit resource group type is "Resource".

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

After you finish

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

Related tasks

[Creating backup policies for SQL Server databases](#) on page 20

You can create a backup policy for the resource or the resource group before you use SnapCenter to back up SQL Server resources, or you can create a backup policy at the time you create a resource group or backup a single resource.

[Monitoring backup operations](#) on page 31

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

[Backing up resources using PowerShell cmdlets](#) on page 52

You can use the PowerShell cmdlets to backup SQL Server databases or Windows file systems. This would include backing up a SQL Server database or Windows file system includes establishing a connection with the SnapCenter Server, discovering the SQL Server database instances or Windows file systems, adding a policy, creating a backup resource group, backing up, and verifying the backup.

Related information

[Backup operations fails with MySQL connection error because of the delay in the TCP_TIMEOUT](#)



[Backup fails with Windows scheduler error](#)

[Quiesce or grouping resources operations fail](#)

Backing up SQL Server resource groups

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, click **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 clicking , and then selecting the tag. You can then click  to close the filter pane.
3. On the **Resource Groups** page, select the resource group that you want to back up, and then click **Back up Now**.
4. On the **Backup** page, perform the following steps:
 - a. If you have 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. After backup, select **Verify** to verify the on-demand backup.
The **Verify** option in the policy applies only to scheduled jobs.
 - c. Click **Backup**.
5. Monitor the operation progress by clicking **Monitor > Jobs**.

Related tasks

[Creating backup policies for SQL Server databases](#) on page 20

You can create a backup policy for the resource or the resource group before you use SnapCenter to back up SQL Server resources, or you can create a backup policy at the time you create a resource group or backup a single resource.

[Creating resource groups and attaching policies for SQL Server](#) on page 25

A resource group is a container to which you add resources that you want to back up and protect together. A resource group enables you to back up all of 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.

[Monitoring backup operations](#) on page 31

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

[Backing up resources using PowerShell cmdlets](#) on page 52

You can use the PowerShell cmdlets to backup SQL Server databases or Windows file systems. This would include backing up a SQL Server database or Windows file system includes establishing a connection with the SnapCenter Server, discovering the SQL Server database instances or Windows file systems, adding a policy, creating a backup resource group, backing up, and verifying the backup.

Related information

[Backup operations fails with MySQL connection error because of the delay in the TCP_TIMEOUT](#)







[Backup fails with Windows scheduler error](#)

Monitoring backup operations


You can monitor the progress of different backup operations by using the SnapCenter Jobs 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. Optional: On 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.

Note: 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. Optional: On the **Job Details** page, click **View logs**.

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


Monitoring operations 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.

About this task

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.

Canceling the SnapCenter Plug-in for Microsoft SQL Server backup operations

You can cancel backup operations that are running, queued, or non-responsive. When you cancel a backup operation, the SnapCenter Server stops the operation and removes all the Snapshot copies from the storage if the backup created is not registered with SnapCenter Server. If the backup is already registered with SnapCenter Server, it will not roll back the already created Snapshot copy even after the cancellation is triggered.

Before you begin


- You must be logged in as the SnapCenter Admin or job owner to cancel restore operations.

About this task

- You can cancel only the log or full backup operations that are queued or running.
- You cannot cancel the operation after the verification has started.
If you cancel the operation before verification, the operation is canceled, and the verification operation will not be performed.
- You can cancel a backup operation from either the Monitor page or the Activity pane.
- In addition to using the SnapCenter GUI, you can use PowerShell cmdlets to cancel 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.

Step

Perform one of the following actions:

From the...	Action
Monitor page	a. In the left navigation pane, click Monitor > Jobs . b. Select the job and click Cancel Job .
Activity pane	a. After initiating the backup job, 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 .

Result

The operation is canceled, and the resource is reverted to the previous state. If the operation you canceled is non-responsive in the canceling or running state, you should run the `Cancel-SmJob -JobID <int> -Force` cmdlet to forcefully stop the backup operation.

Related information

[SnapCenter Software 4.4 Cmdlet Reference Guide](#)




Viewing SQL Server 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

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.

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.
 - If you have upgraded from SnapCenter 1.1, the clones on the secondary (mirror or vault) are not displayed under Mirror copies or Vault copies in the Topology page. All of the clones created using SnapCenter 1.1 are displayed under the Local copies in SnapCenter 3.0.

Note: 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.

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 selected is a cloned database, protect the cloned database, source of the clone is displayed in the Topology page. Click **Details** to view the backup used to clone.

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 backups and clones.

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


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, rename, and delete operations.

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

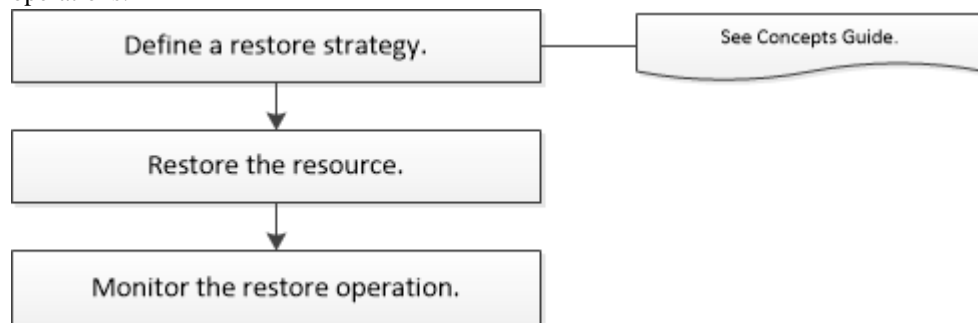
7. Select a clone from the table and click **Clone Split**.

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

Restoring SQL Server resources

You can use SnapCenter to restore SQL Server databases by restoring the data from one or more backups to your active file system and then recovering the database. You can also restore databases that are in Availability Groups and then add the restored databases to the Availability Group. Before restoring an SQL Server database, you must perform several preparatory tasks.

The following workflow shows the sequence in which you must perform the database restoration operations:



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

[SnapCenter Software 4.4 Cmdlet Reference Guide](#)

Related tasks

[Restoring an SQL Server database from secondary storage](#) on page 39

You can restore the backed-up SQL Server databases from the physical LUNs (RDM, iSCSI, or FCP) on a secondary storage system. The Restore feature is a multiphase process that copies all of the data and the log pages from a specified SQL Server backup residing on the secondary storage system to a specified database.

[Restoring and recovering resources using PowerShell cmdlets](#) on page 53

Restoring and recovering a SQL Server database or Windows file system includes initiating a connection session with the SnapCenter Server, listing the backups and retrieving backup information, and restoring a backup.

Related information

[Restore operation might fail on Windows 2008 R2](#)

Requirements for restoring a database

Before you restore a SQL Server database from a SnapCenter Plug-in for Microsoft SQL Server backup, you must ensure that several requirements are met.

- The SQL Server instance must be online and running before you can restore a database. This applies to both user database restore operations and system database restore operations.
- SnapCenter operations that are scheduled to run against the SQL Server data you are restoring must be disabled, including any jobs scheduled on remote management or remote verification servers.
- If system databases are not functional, you must first rebuild the system databases using a SQL Server utility.
- If you are installing the plug-in, ensure that you grant permissions for other roles to restore the Availability Group (AG) backups.

Restoring AG fails when one of the following conditions are met:

- If the plug-in is installed by RBAC user and an admin tries to restore an AG backup
- If the plug-in is installed by an admin and a RBAC user tries to restore an AG backup
- If you are restoring custom log directory backups to an alternate host, the SnapCenter Server and the plug-in host must have the same SnapCenter version installed.
- You must have installed Microsoft hotfix, KB2887595. The Microsoft Support Site contains more information about KB2887595.

[Microsoft Support Article 2887595: Windows RT 8.1, Windows 8.1, and Windows Server 2012 R2 update rollup: November 2013](#)

- You must have backed up the resource groups or database.
- If you are replicating Snapshot copies to a mirror or vault, the SnapCenter administrator must have assigned you the storage virtual machines (SVMs) for both the source volumes and destination volumes.

For information about how administrators assign resources to users, see the SnapCenter installation information.

- All backup and clone jobs must be stopped before restoring the database.
- The restore operation might timeout if the database size is in terabytes (TB).

You must increase the value of the RESTTimeout parameter of SnapCenter Server to 20000000 ms by running the following command: `Set-SmConfigSettings -Agent -configSettings @{ "RESTTimeout" = "2000000" }`. According to the size of the database, the timeout value can be changed and the maximum value that you can set is 2147483648.

If you want to restore while the databases are online, the online restore option should be enabled in the Restore page.

Related tasks

[Stopping and resuming operations on resource groups](#) on page 62

You can temporarily disable scheduled operations from starting on a resource group. Later when you want, you can enable those operations.


Restoring SQL Server databases


You can use SnapCenter to restore backed-up SQL Server databases. Database restoration is a multiphase process that copies all of the data and log pages from a specified SQL Server backup to a specified database.

About this task

- You can restore the backed-up SQL Server databases to a different SQL Server instance on the same host where the backup was created.
You can use SnapCenter to restore the backed-up SQL Server databases to an alternate path so that you do not replace a production version.
- SnapCenter can restore databases in a Windows cluster without taking the SQL Server cluster group offline.
- If a cluster failure (a cluster group move operation) occurs during a restore operation (for example, if the node that owns the resources goes down), you must reconnect to the SQL Server instance, and then restart the restore operation.
- You cannot restore the database when the users or the SQL Server Agent jobs are accessing the database.
- You cannot restore system databases to an alternate path.
- Most of the fields on the Restore wizard pages are self-explanatory. The following information describes fields for which you might need guidance.

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 either **Database** or **Resource Group** from the **View** list.
- 3. Select the database or the resource group from the list.
The topology page is displayed.
- 4. From the **Manage Copies** view, select **Backups** from the storage system.
- 5. Select the backup from the table, and then click the  icon.



- 6. On the **Restore Scope** page, select one of the following options:

Option	Description
Restore the database to the same host where the backup was created	Select this option if you want to restore the database to the same SQL server where the backups are taken.
Restore the database to an alternate host	<p>Select this option if you want the database to be restored to a different SQL server in the same or different host where backups are taken.</p> <p>Select a host name, provide a database name (optional), select an instance, and specify the restore paths.</p> <p>Note: The file extension provided in the alternate path must be same as the file extension of the original database file.</p> <p>If the Restore the database to an alternate host option is not displayed in the Restore Scope page, clear the browser cache.</p>
Restore the database using existing database files	<p>Select this option if you want the database to be restored to an alternate SQL Server in the same or different host where backups are taken. Database files should be already present on the given existing file paths.</p> <p>Select a host name, provide a database name (optional), select an instance, and specify the restore paths.</p>

- 7. On the **Recovery Scope** page, select one of the following options:

Option	Description
None	Select None when you need to restore only the full backup without any logs.
All log backups	Select All log backups up-to-the-minute backup restore operation to restore all of the available log backups after the full backup.

Option	Description
By log backups until	Select By log backups to perform a point-in-time restore operation, which restores the database based on backup logs until the backup log with the selected date.
By specific date until	Select By specific date until to specify the date and time after which transaction logs are not applied to the restored database. This point-in-time restore operation halts the restoration of transaction log entries that were recorded after the specified date and time.
Use custom log directory	If you have selected All log backups , By log backups , or By specific date until and the logs are located at a custom location, select Use custom log directory , and then specify the log location. Note: The custom log directory is not supported for availability group database.

8. On the **Pre Ops** page, perform the following steps:
 - a. On the **Pre Restore Options** page, select one of the following options:
 - Select **Overwrite the database with same name during restore** to restore the database with the same name.
 - Select **Retain SQL database replication settings** to restore the database and retain the existing replication settings.
 - Select **Create transaction log backup before restore** to create a transaction log before the restore operation begins.
 - Select **Quit restore if transaction log backup before restore fails** to abort the restore operation if the transaction log backup fails.
 - b. Specify optional scripts to run before performing a restore job.
For example, you can run a script to update SNMP traps, automate alerts, send logs, and so on.
9. On the **Post Ops** page, perform the following steps:
 - a. In the **Choose database state after restore completes** section, select one of the following options:
 - Select **Operational, but unavailable for restoring additional transaction logs** if you are restoring all of the necessary backups now.
This is the default behavior, which leaves the database ready for use by rolling back the uncommitted transactions. You cannot restore additional transaction logs until you create a backup.
 - Select **Non-operational, but available for restoring additional transactional logs** to leave the database non-operational without rolling back the uncommitted transactions. Additional transaction logs can be restored. You cannot use the database until it is recovered.
 - Select **Read-only mode, available for restoring additional transactional logs** to leave the database in read-only mode.
This option undoes uncommitted transactions, but saves the undone actions in a standby file so that recovery effects can be reverted.
If the `Undo directory` option is enabled, more transaction logs are restored. If the restore operation for the transaction log is unsuccessful, the changes can be rolled back. The SQL Server documentation contains more information.
 - b. Specify optional scripts to run after performing a restore job.

For example, you can run a script to update SNMP traps, automate alerts, send logs, and so on.

10. On 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.

11. Review the summary, and then click **Finish**.
12. Monitor the restore process by using the **Monitor > Jobs** page.

Related tasks

[Restoring and recovering resources using PowerShell cmdlets](#) on page 53

Restoring and recovering a SQL Server database or Windows file system includes initiating a connection session with the SnapCenter Server, listing the backups and retrieving backup information, and restoring a backup.

[Restoring an SQL Server database from secondary storage](#) on page 39

You can restore the backed-up SQL Server databases from the physical LUNs (RDM, iSCSI, or FCP) on a secondary storage system. The Restore feature is a multiphase process that copies all of the data and the log pages from a specified SQL Server backup residing on the secondary storage system to a specified database.

Related information

[Installing and setting up SnapCenter](#)

Restoring an SQL Server database from secondary storage

You can restore the backed-up SQL Server databases from the physical LUNs (RDM, iSCSI, or FCP) on a secondary storage system. The Restore feature is a multiphase process that copies all of the data and the log pages from a specified SQL Server backup residing on the secondary storage system to a specified database.

Before you begin


You must have replicated the Snapshot copies from primary to secondary storage system.

You must ensure that the SnapCenter Server and the plug-in host are able to connect to the secondary storage system.

About this task

Most of the fields on the Restore wizard pages are explained in the basic restore process. The following information describes some of the fields for which you might need guidance.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. On the **Resources** page, select **Database** or **Resource Group** from the **View** drop-down list.
3. Select the database or resource group.
The database or resource group topology page is displayed.
4. In the **Manage Copies** section, select **Backups** from the secondary storage system (mirrored or vault).
5. Select the backup from the list, and then click .
6. On the **Location** page, choose the destination volume for restoring selected resource.
7. Complete the **Restore** wizard, review the summary, and then click **Finish**.

After you finish

If you restored a database to a different path that is shared by other databases, you should perform a full backup and backup verification to confirm that your restored database is free of physical-level corruption.

Reseeding Availability Group databases

Reseed is an option to restore Availability Group (AG) databases. If a secondary database gets out of synchronization with the primary database in an AG, you can reseed the secondary database.

Before you begin

- You must have created backup of secondary AG database that you want to restore.
- The SnapCenter Server and the plug-in host must have the same SnapCenter version installed.

About this task

You cannot perform reseed operation on primary databases.

You cannot perform a reseed operation if the replica database is removed from the availability group. When the replica is removed, the reseed operation fails.

While running the reseed operation on SQL Availability Group database, you should not trigger log backups on the replica databases of that availability group database. If you trigger log backups during reseed operation, the reseed operation fails with The mirror database, "database_name" has insufficient transaction log data to preserve the log backup chain of the principal database error message.

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 **Database** from the **View** list.
3. Select secondary AG database from the list.
4. Click **Reseed**.
5. Monitor the operation progress by clicking **Monitor > Jobs**.






Monitoring 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. Optional: On the **Jobs** page, perform the following steps:
 - a. Click  to filter the list so that only restore operations are listed.
 - b. Optional: 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. Optional: On the **Job Details** page, click **View logs**.

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

After the volume based restore operation, the backup metadata is deleted from the SnapCenter repository but the backup catalog entries remain in SAP HANA catalog. Though the backup job status displays  , you should click on job details to see the warning sign of some of the child tasks. Click on the warning sign and delete the indicated backup catalog entries.

Canceling restore operations

You can cancel restore jobs that are queued.

Before you begin


- You must be logged in as the SnapCenter Admin or job owner to cancel restore operations.

About this task

- You can cancel a restore operation from either the Monitor page or the Activity pane.
- You cannot cancel a running restore operation.
- You can use the SnapCenter GUI, PowerShell cmdlets, or CLI commands to cancel the restore operations.
- The **Cancel Job** button is disabled for restore 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.

Step

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 job and click Cancel Job.
Activity pane	<ol style="list-style-type: none">a. After initiating the restore 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.

Related information

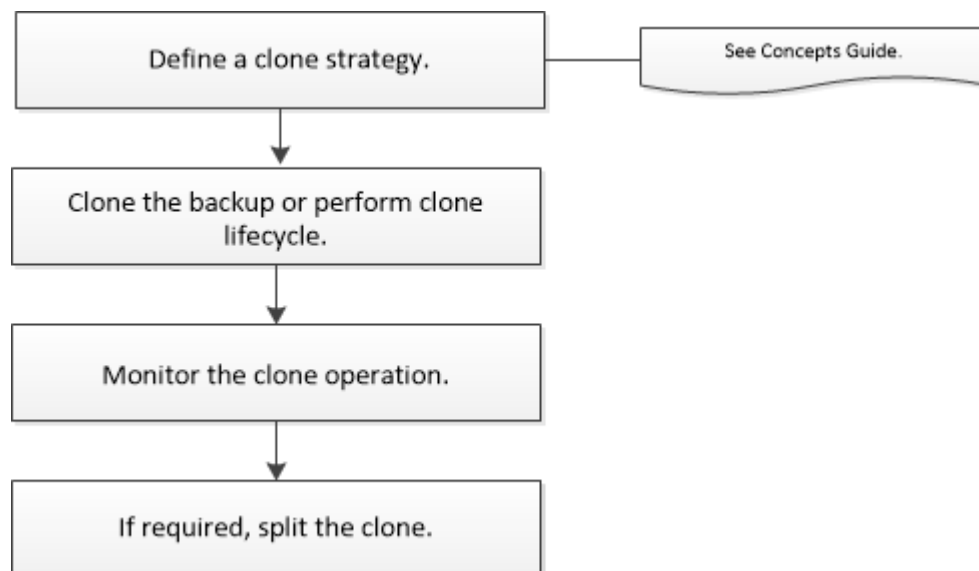
[*SnapCenter Software 4.4 Cmdlet Reference Guide*](#)

Cloning SQL Server database resources

You must perform several tasks using SnapCenter Server before cloning database resources from a backup. Database cloning is the process of creating a point-in-time copy of a production database or its backup set. You can clone databases to test functionality that has to be implemented using the current database structure and content during application development cycles, to use the data extraction and manipulation tools when populating data warehouses, or to recover data that was mistakenly deleted or changed.

A database cloning operation generates reports based on the job IDs.

The following workflow shows the sequence in which you must perform the cloning operations:



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

[SnapCenter Software 4.4 Cmdlet Reference Guide](#)

Related concepts

[Backing up, restoring, cloning, and removing backups using PowerShell cmdlets](#) on page 51
The SnapCenter Plug-in for Microsoft SQL Server and the SnapCenter Plug-in for Microsoft Windows include PowerShell cmdlets for scripting of backup, restore, and clone operations.

Related tasks

[Cloning from a SQL Server database backup](#) on page 44
You can use SnapCenter to clone a SQL Server database backup. If you want to access or restore an older version of the data, you can clone database backups on demand.

[Performing Clone Lifecycle](#) on page 46
Using SnapCenter, you can create clones from a resource group or database. You can either perform on-demand clone or you can schedule recurring clone operations of a resource group or database. If you clone a backup periodically, you can use the clone to develop applications, populate data, or recover data.

[Monitoring clone operations in SnapCenter](#) on page 47

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.

Related information

Clone operation might fail or take longer time to complete with default TCP_TIMEOUT value

Cloning from a SQL Server database backup

You can use SnapCenter to clone a SQL Server database backup. If you want to access or restore an older version of the data, you can clone database backups on demand.

Before you begin

- You should have prepared for data protection by completing tasks such as adding hosts, identifying resources, and creating storage system connections.
- You should have backed up databases or resource groups.
- The protection type such as mirror, vault, or mirror-vault for data LUN and log LUN should be same to discover secondary locators during cloning to an alternate host using log backups.
- If the mounted clone drive cannot be found during a SnapCenter clone operation, you should change the CloneRetryTimeout parameter of SnapCenter Server to **300**.
- You should ensure that the aggregates hosting the volumes should be in the assigned aggregates list of the storage virtual machine (SVM).

About this task


- While cloning to a standalone database instance, ensure that the mount point path exists and it is a dedicated disk.
- While cloning to a Failover Cluster Instance (FCI), ensure that the mount points exists, it is a shared disk, and the path and the FCI should belong to the same SQL resource group.
- Ensure that there is only one vFC or FC initiator attached to each host. This is because, SnapCenter supports only one initiator per host.
- If the source database or the target instance is on a cluster shared volume (csv), then the cloned database will be on the csv.

Note: For virtual environments (VMDK/RDM), ensure that the mount point is a dedicated disk.

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 either **Database** or **Resource Group** from the **View** list.

Note: Cloning of a backup of an instance is not supported.

3. Select the database or resource group.
4. From the **Manage Copies view** page, select the backup either from primary or secondary (mirrored or vaulted) storage system.
5. Select the backup, and then click .
6. On the **Clone Options** page, perform the following actions:

For this field...	Do this...
Clone server	Choose a host on which the clone should be created.
Clone instance	Choose a clone instance to which you want to clone the database backup. This SQL instance must be located in the specified clone server.

For this field...	Do this...
Clone suffix	Enter a suffix that will be appended to the clone file name to identify that the database is a clone. For example, "db1_clone". If you are cloning to the same location as the original database, you must provide a suffix to differentiate the cloned database from the original database. Otherwise, the operation fails.
"Auto assign mount point" or "Auto assign volume mount point under path"	Choose whether to automatically assign a mount point or a volume mount point under a path. Auto assign volume mount point under path: The mount point under a path allows you to provide a specific directory. The mount points will be created within that directory. Before you choose this option, you must ensure that the directory is empty. If there is a database in the directory, the database will be in an invalid state after the mount operation.

7. On the **Logs** page, select one of the following options:

For this field...	Do this...
None	Choose this option when you want to clone only the full backup without any logs.
All log backups	Choose this option to clone all the available log backups dated after the full backup.
By log backups until	Choose this option to clone the database based on the backup logs that were created up to the backup log with the selected date.
By specific date until	Specify the date and time after which the transaction logs are not applied to the cloned database. This point-in-time clone halts the clone of the transaction log entries that were recorded after the specified date and time.

8. On the **Script** page, enter the script timeout, path, and the arguments of the prescript or postscript that should be run before or after the clone operation, respectively.
For example, you can run a script to update SNMP traps, automate alerts, send logs, and so on.
The default script timeout is 60 seconds.
9. On 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 restore operation performed, select **Attach Job Report**.

Note: For email notification, you must have specified the SMTP server details using the either the GUI or the PowerShell command `Set-SmSmtServer`.
10. Review the summary, and then click **Finish**.
11. Monitor the operation progress by clicking **Monitor > Jobs**.

After you finish

After the clone is created, you should never rename it.

Related tasks

[Cloning backups using PowerShell cmdlets](#) on page 55

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

Related reference

[Backing up SQL Server resources](#) on page 15

When you install the SnapCenter Plug-in for Microsoft SQL Server in your environment, you can use SnapCenter to back up the SQL Server resources.

Related information

[Clone operation might fail or take longer time to complete with default TCP_TIMEOUT value](#)

[The failover cluster instance database clone fails](#)

Performing Clone Lifecycle

Using SnapCenter, you can create clones from a resource group or database. You can either perform on-demand clone or you can schedule recurring clone operations of a resource group or database . If you clone a backup periodically, you can use the clone to develop applications, populate data, or recover data.

About this task

SnapCenter enables you to schedule multiple clone operations to run simultaneously across multiple servers.

- While cloning to a standalone database instance, ensure that the mount point path exists and it is a dedicated disk.
- While cloning to a Failover Cluster Instance (FCI), ensure that the mount points exists, it is a shared disk, and the path and the FCI should belong to the same SQL resource group.
- If the source database or the target instance is on a cluster shared volume (csv), then the cloned database will be on the csv.

Note: For virtual environments (VMDK/RDM), ensure that the mount point is a dedicated disk.

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 either **Database** or **Resource Group** from the **View** list.
3. Select the resource group or database, and then click **Clone Lifecycle**.
4. In the **Options** page, perform the following actions:

For this field...	Do this...
Clone job name	Specify the clone life cycle job name that helps in monitoring and modifying the clone life cycle job.
Clone server	Choose the host on which the clone should be placed.
Clone instance	Choose the clone instance to which you want to clone the database. This SQL instance must be located in the specified clone server.

For this field...	Do this...
Clone suffix	<p>Enter a suffix that will be appended to the clone database to identify that it is a clone.</p> <p>Each SQL instance that is used to create a clone resource group must have a unique database name. For example, if the clone resource group contains a source database "db1" from an SQL instance "inst1", and if "db1" is cloned to "inst1", then the clone database name should be "db1__clone". "__clone" is a mandatory user-defined suffix because the database is cloned to the same instance. If "db1" is cloned to the SQL instance "inst2", then the clone database name can remain "db1" (the suffix is optional) because the database is cloned to a different instance.</p>
Auto assign mount point or Auto assign volume mount point under path	<p>Choose whether to automatically assign a mount point or volume mount point under a path.</p> <p>Choosing to auto assign a volume mount point under a path enables you to provide a specific directory. The mount points will be created within that directory. Before you choose this option, you must ensure that the directory is empty. If there is a database in the directory, the database will be in an invalid state after the mount operation.</p>

- In the **Location** page, select a storage location to create a clone.
- In the **Script** page, enter the path and the arguments of the prescript or postscript that should be run before or after the clone operation, respectively.
For example, you can run a script to update SNMP traps, automate alerts, send logs, and so on.
The default script timeout is 60 seconds.
- In the **Schedule** page, perform one of the following actions:
 - Select **Run now** if you want to execute the clone job immediately.
 - Select **Configure schedule** when you want to determine how frequently the clone operation should occur, when the clone schedule should start, on which day the clone operation should occur, when the schedule should expire, and whether the clones have to be deleted after the schedule expires.
- 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 restore operation performed, select **Attach Job Report**.
Note: For email notification, you must have specified the SMTP server details using the either the GUI or the PowerShell command `Set-SmSmtServer`.
- Review the summary, and then click **Finish**.

After you finish


You should monitor the cloning process using the **Monitor > Jobs** page.






Monitoring clone operations in SnapCenter

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**.

Canceling clone operations

You can cancel clone operations that are queued.

Before you begin


- You must be logged in as the SnapCenter Admin or job owner to cancel operations.

About this task

- You can cancel a clone operation from either the Monitor page or the Activity pane.
- You cannot cancel a running clone operation.
- You can use the SnapCenter GUI, PowerShell cmdlets, or CLI commands to cancel the clone operations.
- 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.

Step

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 click Cancel Job.
Activity pane	<ol style="list-style-type: none">a. After initiating the clone 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.

Related information

[SnapCenter Software 4.4 Cmdlet Reference Guide](#)

Splitting 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.
- Logical Storage Management Guide* has more information about clone split operation limitations.
[ONTAP 9 Logical Storage Management Guide](#)
- Ensure that the volume or aggregate on the storage system is online.

Steps

- In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
- On 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.

- Select the appropriate resource from the list.
The resource topology page is displayed.
- From the **Manage Copies** view, select the cloned resource (for example, the database or LUN), and then click .
- Review the estimated size of the clone that is to be split and the required space available on the aggregate, and then click **Start**.
- Monitor the operation progress by clicking **Monitor** > **Jobs**.
The clone split operation stops responding if the SMCore service restarts and the databases on which the clone split operation was performed are listed as clones in the Resources page. You should run the `stop-smJob` cmdlet to stop the clone split operation, and then retry the clone split operation.

After you finish

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](#)

Backing up, restoring, cloning, and removing backups using PowerShell cmdlets

The SnapCenter Plug-in for Microsoft SQL Server and the SnapCenter Plug-in for Microsoft Windows include PowerShell cmdlets for scripting of backup, restore, and clone operations.

The following are common tasks you might perform using PowerShell cmdlets:

- Preparing the PowerShell environment
- Creating a storage system connection and a credential
- Backing up SQL Server databases or Windows file systems
- Restoring and recovering SQL Server databases
- Restoring Windows file systems
- Cloning backups for SQL Server databases or Windows file systems
- Removing backups
- Cleaning up secondary backup counts

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 *Cmdlet Reference Guide*.

[SnapCenter Software 4.4 Cmdlet Reference Guide](#)

Creating a storage system connection and a credential using PowerShell cmdlets

You must create a storage virtual machine (SVM) connection and a credential before using PowerShell cmdlets to perform data protection operations.

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 databases 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.

About this task

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 *Cmdlet Reference Guide*.

[SnapCenter Software 4.4 Cmdlet Reference Guide](#)

Steps

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

This example opens a PowerShell session:

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

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

This example creates a new storage system connection:

```
PS C:\> Add-SmStorageConnection -Storage test_vsl -Protocol Https  
-Timeout 60
```

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

This example creates a new credential named `FinanceAdmin` with Windows credentials:

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

Backing up resources using PowerShell cmdlets

You can use the PowerShell cmdlets to backup SQL Server databases or Windows file systems. This would include backing up a SQL Server database or Windows file system includes establishing a connection with the SnapCenter Server, discovering the SQL Server database instances or Windows file systems, adding a policy, creating a backup resource group, backing up, and verifying the backup.

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.
- You must have added hosts and discovered resources.

About this task

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 *Cmdlet Reference Guide*.

[SnapCenter Software 4.4 Cmdlet Reference Guide](#)

Steps

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

```
Open-smconnection -SMSbaseurl https://snapctr.demo.netapp.com:8146
```

The username and password prompt is displayed.

2. Create a backup policy by using the `Add-SmPolicy` cmdlet.

This example creates a new backup policy with a SQL backup type of `FullBackup`:

```
PS C:\> Add-SmPolicy -PolicyName TESTPolicy  
-PluginPolicyType SCSQL -PolicyType Backup  
-SqlBackupType FullBackup -Verbose
```

This example creates a new backup policy with a Windows file system backup type of `CrashConsistent`:

```
PS C:\> Add-SmPolicy -PolicyName FileSystemBackupPolicy  
-PluginPolicyType SCW -PolicyType Backup  
-ScwBackupType CrashConsistent -Verbose
```

3. Discover host resources by using the `Get-SmResources` cmdlet.

This example discovers the resources for the Microsoft SQL plug-in on the specified host:

Backing up, restoring, cloning, and removing backups using PowerShell cmdlets

```
C:\PS>PS C:\> Get-SmResources -HostName vise-f6.sddev.mycompany.com
-PluginCode SCSQL
```

This example discovers the resources for Windows file systems on the specified host:

```
C:\PS>PS C:\> Get-SmResources -HostName vise2-f6.sddev.mycompany.com
-PluginCode SCW
```

4. Add a new resource group to SnapCenter by using the Add-SmResourceGroup cmdlet.

This example creates a new SQL database backup resource group with the specified policy and resources:

```
PS C:\> Add-SmResourceGroup -ResourceGroupName AccountingResource
-Resources @{"Host"="visef6.org.com";
"Type"="SQL Database";"Names"="vise-f6\PayrollDatabase"}
-Policies "BackupPolicy"
```

This example creates a new Windows file system backup resource group with the specified policy and resources:

```
PS C:\> Add-SmResourceGroup -ResourceGroupName EngineeringResource
-PluginCode SCW -Resources @{"Host"="WIN-VOK20IKID5I";
"Type"="Windows Filesystem";"Names"="E:\"}
-Policies "EngineeringBackupPolicy"
```

5. Initiate a new Snapshot copy job by using the New-SmBackup cmdlet.

```
PS C:\> New-SmBackup -ResourceGroupName PayrollDataset -Policy FinancePolicy
```

6. View the status of the backup job by using the Get-SmBackupReport cmdlet.

This example displays a job summary report of all jobs that were run on the specified date:

```
PS C:\> Get-SmJobSummaryReport -Date ?1/27/2016?
```

Restoring and recovering resources using PowerShell cmdlets

Restoring and recovering a SQL Server database or Windows file system includes initiating a connection session with the SnapCenter Server, listing the backups and retrieving backup information, and restoring a backup.

Before you begin

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

About this task

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 *Cmdlet Reference Guide*.

[SnapCenter Software 4.4 Cmdlet Reference Guide](#)

Steps

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

```
Open-smconnection -SMSbaseurl https://snapctr.demo.netapp.com:8146
```

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:

Backing up, restoring, cloning, and removing backups using PowerShell cmdlets

```
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

Secondary_SCSPR0019366001_01-15-2015_06.49.08:

```
PS C:\> Get-SmBackupReport
-BackupName Secondary_SCSPR0019366001_01-15-2015_06.49.08

BackedUpObjects : {TestDB1, TestDB2, TestDB3, TestDB4...}
FailedObjects : {}
BackupType : Full Backup
IsScheduled : False
SmBackupId : 52
SmJobId : 585
StartDateTime : 1/15/2015 6:49:07 AM
EndDateTime : 1/15/2015 6:49:21 AM
Duration : 00:00:13.8370000
CreatedDateTime : 1/15/2015 6:49:18 AM
Status : Completed
ProtectionGroupName : Secondary
SmProtectionGroupId : 5
PolicyName : Vault
SmPolicyId : 18
BackupName : Secondary_SCSPR0019366001_01-15-2015_06.49.08
VerificationStatus : NotVerified
```

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

```
C:\PS>PS C:\> Restore-SmBackup -PluginCode SCSQL
-AppObjectId 'vise-f6\PayrollDatabase'
-BackupName 'NetApp_PayrollDataset_Backup Policy_
vise-f6_NetApp_08-07-2015_08.48.59.6962'
-RestoreWhenOnline
```

Name	: Restore 'vise-f6\PayrollDatabase'
Id	: 199
StartTime	: 8/7/2015 9:21:36 AM
EndTime	:
IsCancellable	: False
IsRestartable	: False
IsCompleted	: False
IsVisible	: False
IsScheduled	: False
PercentageCompleted	: 0
Description	:
Status	: Queued
Owner	:
Error	:
Priority	: None
Tasks	: {}
ParentJobID	: 0
EventId	: 0

```
Restore-SmBackup -PluginCode SCSQL -AppObjectId 'scspr0270378001\abc' -BackupName 'sc
spr0270378001_abc_scspr0270378001_07-25-2017_04.51.10.5795' -AlternatePath @{Source='D:
\data\abc.mdf';Destinatio
n='D:\data\bharathaewf123.mdf'},@{Source='D:\log\bharath_log.ldf';Destination='D:\log
\bharathaef_log123.ldf'} -SQLInstan
ceName 'scspr0273089004' -DatabaseName 'abc123adwqa1231' -ExistingFiles
```

Cloning backups using PowerShell cmdlets

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

Before you begin

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

About this task

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 *Cmdlet Reference Guide*.

[SnapCenter Software 4.4 Cmdlet Reference Guide](#)

Steps

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

```
Open-SmConnection -SMSbaseurl https://snapctr.demo.netapp.com:8146
```

2. List the backups that can be cloned by using the `Get-SmBackup` or `Get-SmResourceGroup` cmdlet.

This example displays information about all available backups:

```
C:\PS>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 information about a specified resource group, its resources, and associated policies:

```
PS C:\> Get-SmResourceGroup -ListResources -ListPolicies
```

```
Description :
CreationTime : 8/4/2015 3:44:05 PM
ModificationTime : 8/4/2015 3:44:05 PM
EnableEmail : False
EmailSMTPServer :
EmailFrom :
EmailTo :
EmailSubject :
EnableSysLog : False
ProtectionGroupType : Backup
EnableAsupOnFailure : False
Policies : {FinancePolicy}
HostResourceMapping : {}
Configuration : SMCoreContracts.SmCloneConfiguration
LastBackupStatus :
VerificationServer :
EmailBody :
EmailNotificationPreference : Never
VerificationServerInfo : SMCoreContracts.SmVerificationServerInfo
SchedulerSQLInstance :
CustomText :
CustomSnapshotFormat :
SearchResources : False
```

```
ByPassCredential : False
IsCustomSnapshot :
MaintenanceStatus : Production
PluginProtectionGroupTypes : {SMSQL}
Name : Payrolldataset
Type : Group
Id : 1
Host :
UserName :
Passphrase :
Deleted : False
Auth : SMCoreContracts.SmAuth
IsClone : False
CloneLevel : 0
ApplySnapvaultUpdate : False
ApplyRetention : False
RetentionCount : 0
RetentionDays : 0
ApplySnapMirrorUpdate : False
SnapVaultLabel :
MirrorVaultUpdateRetryCount : 7
AppPolicies : {}
Description : FinancePolicy
PreScriptPath :
PreScriptArguments :
PostScriptPath :
PostScriptArguments :
ScriptTimeout : 60000
DateModified : 8/4/2015 3:43:30 PM
DateCreated : 8/4/2015 3:43:30 PM
Schedule : SMCoreContracts.SmSchedule
PolicyType : Backup
PluginPolicyType : SMSQL
Name : FinancePolicy
Type :
Id : 1
Host :
UserName :
Passphrase :
Deleted : False
Auth : SMCoreContracts.SmAuth
IsClone : False
CloneLevel : 0
clab-a13-13.sddev.lab.netapp.com
DatabaseGUID :
SQLInstance : clab-a13-13
DbStatus : AutoClosed
DbAccess : eUndefined
IsSystemDb : False
IsSimpleRecoveryMode : False
IsSelectable : True
SqlDbFileGroups : {}
SqlDbLogFiles : {}
AppFileStorageGroups : {}
LogDirectory :
AgName :
Version :
VolumeGroupIndex : -1
IsSecondary : False
Name : TEST
Type : SQL Database
Id : clab-a13-13\TEST
Host : clab-a13-13.sddev.mycompany.com
UserName :
Passphrase :
Deleted : False
Auth : SMCoreContracts.SmAuth
IsClone : False
```


3. Initiate a clone operation from an existing backup by using the `New-SmClone` cmdlet.

This example creates a clone from a specified backup with all logs:

```
PS C:\> New-SmClone
-BackupName payroll_dataset_vise-f3_08-05-2015_15.28.28.9774
-Resources @{"Host"="vise-f3.sddev.mycompany.com";
"Type"="SQL Database";"Names"="vise-f3\SQLExpress\payroll"}
-CloneToInstance vise-f3\sqlexpress -AutoAssignMountPoint
-Suffix _clonefrombackup
-LogRestoreType All -Policy clonefromprimary_ondemand

PS C:\> New-SmBackup -ResourceGroupName PayrollDataset -Policy FinancePolicy
```

This example creates a clone to a specified Microsoft SQL Server instance:

```
PS C:\> New-SmClone
-BackupName "BackupDS1_NY-VM-SC-SQL_12-08-2015_09.00.24.8367"
-Resources @{"host"="ny-vm-sc-sql";"Type"="SQL Database";
"Names"="ny-vm-sc-sql\AdventureWorks2012_data"}
-AppPluginCode SMSQL -CloneToInstance "ny-vm-sc-sql"
-Suffix _CLPOSH -AutoAssignMountPoint
-AssignMountPointUnderPath "C:\SCMounts"
```

4. View the status of the clone job by using the `Get-SmCloneReport` cmdlet.

This example displays a clone report for the specified job ID:

```
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}
```

Removing backups using PowerShell cmdlets

You can use the `Remove-SmBackup` cmdlet to delete backups if you no longer require them for other data protection operations.

Before you begin

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

About this task

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 *Cmdlet Reference Guide*.

[SnapCenter Software 4.4 Cmdlet Reference Guide](#)

Steps

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

```
Open-SmConnection -SMSbaseurl https:\\snapctr.demo.netapp.com:8146/
```

2. Delete one or more backup using the `Remove-SmBackup` cmdlet.
This example deletes two backups using their backup IDs:

```
Remove-SmBackup -BackupIds 3,4  
Remove-SmBackup  
Are you sure want to remove the backup(s).  
[Y] Yes [A] Yes to All [N] No [L] No to All [S] Suspend [?] Help (default is "Y"):
```

Related information

[SnapCenter Software 4.4 Cmdlet Reference Guide](#)

Cleaning up the secondary backup count using PowerShell cmdlets

You can use the `Remove-SmBackup` cmdlet to clean up the backup count for secondary backups that have no Snapshot copies. You might want use this cmdlet when the total Snapshot copies displayed in the Manage Copies topology do not match the secondary storage Snapshot copy retention setting.

Before you begin

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

About this task

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 *Cmdlet Reference Guide*.

[SnapCenter Software 4.4 Cmdlet Reference Guide](#)

Steps

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

```
Open-SmConnection -SMSbaseurl https:\\snapctr.demo.netapp.com:8146/
```

2. Clean up secondary backups count using the `-CleanupSecondaryBackups` parameter.
This example cleans up the backup count for secondary backups with no Snapshot copies:

```
Remove-SmBackup -CleanupSecondaryBackups  
Remove-SmBackup  
Are you sure want to remove the backup(s).  
[Y] Yes [A] Yes to All [N] No [L] No to All [S] Suspend [?] Help (default is "Y"):
```

Related information

[SnapCenter Software 4.4 Cmdlet Reference Guide](#)

Managing policies

You can create, copy, modify, view, and delete backup policies.

About this task

You can perform the following tasks related to policies:

- Create a policy.
- Modify a policy.

Note: You can change the schedule type for a policy only after you detach the policy. To change the schedule you must modify the resource group.

- Copy a policy by accepting the default name or typing a new name.
- Detach a policy from a resource group.
- Delete a policy.

Related tasks

[Creating backup policies for SQL Server databases](#) on page 20

You can create a backup policy for the resource or the resource group before you use SnapCenter to back up SQL Server resources, or you can create a backup policy at the time you create a resource group or backup a single resource.

Detaching policies

You can detach policies from a resource or resource group any time that you no longer want those policies to govern data protection for the resources. You must detach a policy before you can delete it or before you modify the schedule type.

About this task



Attention: You cannot detach a policy from a resource or resource group if it is the only policy attached.

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 **Resource Group** from the **View** list.
3. Select the resource group, and then click **Modify Resource Group**.
4. In the **Policies** page of the **Modify Resource Group** wizard, from the drop-down list, clear the check mark next to the policies you want to detach.

Note: You cannot detach the last remaining policy from an individual resource because every resource must have at least one policy attached. Therefore, if you want to delete or modify the only policy attached to a resource, you must perform the following:

- a. Attach a second placeholder policy.
 - b. Detach the original policy from the resource.
5. Make any additional modifications to the resource group in the rest of the wizard, and then click **Finish**.

Modifying policies

You can modify the replication options, Snapshot copy retention settings, error retry count, or scripts information while a policy is attached to a resource or resource group. You can modify the schedule type (frequency) only after you detach a policy.

About this task

Modifying the schedule type in a policy requires additional steps because the SnapCenter Server registers the schedule type only at the time the policy is attached to a resource or resource group.

If you want to...	Then...
Add an additional schedule type	<p>Create a new policy and attach it to the necessary resources or resource groups.</p> <p>For example, if a resource group policy specifies only hourly backups and you want to add daily backups also, you can create a policy with a daily schedule type and add it to the resource group. The resource group would then have two policies: hourly and daily.</p>
Remove or change a schedule type	<ol style="list-style-type: none">1. Detach the policy from every resource and resource group that uses that policy.2. Modify the schedule type.3. Attach the policy again to all the resources and resource groups. <p>For example, if a policy specifies hourly backups and you want to change that to daily backups, you must detach the policy first.</p> <p>Note: You cannot detach the last remaining policy from a individual resource because every resource must have at least one policy attached. Therefore, if you want to modify the schedule type of the only policy attached to a resource, you must perform the following:</p> <ol style="list-style-type: none">1. Attach a second placeholder policy.2. Detach the original policy from every resource and resource group that uses that policy.3. Modify the schedule type.4. Attach the modified policy again to all the resources and resource groups.5. Detach the placeholder policy.

Steps

1. In the left navigation pane, click **Settings**.
2. In the **Settings** page, click **Policies**.
3. Select the policy, and then click **Modify**.
4. Modify the information, and then click **Finish**.

Deleting policies

If you no longer require policies, you might want to delete them.

Before you begin

You must have detached the policy from resource groups if the policy is associated with any resource group.

Steps

1. In the left navigation pane, click **Settings**.
2. In the **Settings** page, click **Policies**.
3. Select the policy, and then click **Delete**.
4. Click **Yes**.

Managing resource groups

You can create, modify, and delete resource groups. You can also perform backup and verification operations on resource groups.

About this task

You can perform the following tasks related to resource groups:

- Create a resource group.
- Modify a resource group by selecting the resource group and clicking **Modify Resource Group** to edit the information you provided while creating the resource group.

Note: You can change the schedule while modifying the resource group. However, to change the schedule type you must modify the policy.

Note: If you remove resources from a resource group, the backup retention settings defined in the policies currently attached to the resource group will continue to be applied to the removed resources.

- Create a backup of a resource group.
- Create a clone of a backup.
You can clone from the existing backups of SQL, Oracle, Windows file systems, custom applications, and SAP HANA database resources or resource groups.
- Create a clone of a resource group.
This operation is supported only for SQL resource groups (which contains only databases). You can configure a schedule for cloning a resource group (clone lifecycle).
- Prevent scheduled operations on resource groups from starting.
- Delete a resource group.

Related tasks

[Creating resource groups and attaching policies for SQL Server](#) on page 25

A resource group is a container to which you add resources that you want to back up and protect together. A resource group enables you to back up all of 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.

[Backing up SQL Server resource groups](#) on page 30

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.

Stopping and resuming operations on resource groups

You can temporarily disable scheduled operations from starting on a resource group. Later when you want, you can enable those 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, select **Resource Group** from the **View** list.
3. Select the resource group and click **Maintenance**.
4. Click **OK**.

After you finish

If you want to resume operations on the resource group that you had put on maintenance mode, select the resource group and click **Production**.

Deleting resource groups

You can delete a resource group if you no longer need to protect the resources in the resource group. You must ensure that resource groups are deleted before you remove plug-ins from SnapCenter.

Before you begin

If you are managing SQL or Windows file system resources, you must have manually deleted all clones created for any of the resources in the resource group.

About this task

You can optionally force the deletion of all backups, metadata, policies, and Snapshot copies associated with the resource group.

Note: In a SnapVault relationship, the last Snapshot copy cannot be deleted; therefore, the resource group cannot be deleted. Before deleting a resource group that is part of a SnapVault relationship, you must remove the SnapVault relationship and then delete the last Snapshot copy.

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 **Resource Group** from the **View** list.
3. Select the resource group, and then click **Delete**.
4. Optional: Select the **Delete backups and detach policies associated with this Resource Group** check box to remove all backups, metadata, policies, and Snapshot copies associated with the resource group.
5. Click **OK**.

Managing backups


You can rename and delete backups. You can also delete multiple backups simultaneously.

Renaming backups

You can rename backups if you want to provide a better name to improve searchability.

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 or resource group from the list.
The resource or resource group topology page is displayed. If the resource or resource group is not configured for data protection, the Protect wizard is displayed instead of the topology page.
4. From the **Manage Copies** view, select **Backups** from the primary storage systems.
You cannot rename the backups that are on the secondary storage system.
If you have cataloged the backups of Oracle databases using Oracle Recovery Manager (RMAN), you cannot rename those cataloged backups.

5. Select the backup, and then click .
6. In the **Rename backup as** field, enter a new name and click **OK**.

Deleting backups


You can delete backups if you no longer require the backup for other data protection operations.

Before you begin

You must have deleted the associated clones before deleting a backup.

If a backup is associated with a cloned resource, you cannot delete the backup.

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 or resource group from the list.
The resource or resource group topology page is displayed.
4. From the **Manage Copies** view, select **Backups** from the primary storage systems.
You cannot delete the backups that are on the secondary storage system.
5. Select the backup, and then click .
6. Click **OK**.

Note: If you have some stale database backups in SnapCenter which do not have corresponding backups on the storage system, you must use `remove-smbbackup` command to clean up these stale backup entries. If the stale backups were cataloged, they will be uncataloged from the recovery catalog database.

Managing clones

You can view and delete clones.


Deleting clones

You can delete clones if you find them no longer necessary.

About this task

A clone that has been cloned again cannot be deleted. For example, if the production database *db1* is cloned to *db1_clone1* and subsequently cloned to *db1_clone2*, and you decide that you want to delete *db1_clone1*, you must first delete *db1_clone2*, and then delete *db1_clone1*.

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 or resource group from the list.
The resource or the resource group topology page is displayed.
4. From the **Manage Copies** view, select **Clones** either from the primary or secondary (mirrored or replicated) storage systems.
5. Select the clone, and then click .
6. Click **OK**.

After you finish

After deleting the clone, sometimes the file systems are not deleted. You must increase the value of the `CLONE_DELETE_DELAY` parameter by running the following command:

```
./sccli Set-SmConfigSettings
```

Note: The `CLONE_DELETE_DELAY` parameter specifies the number of seconds to wait after completing the deletion of application clone and before starting the deletion of file system.

After modifying the value of the parameter, restart the SnapCenter Plug-in Loader (SPL) service.

Copyright and trademark

Copyright

Copyright © 2020 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.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, 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).

Trademark

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

<http://www.netapp.com/us/legal/netapptmlist.aspx>