



Protect Microsoft SQL Server databases

SnapCenter Software 4.6

NetApp
November 24, 2022

This PDF was generated from https://docs.netapp.com/us-en/snapcenter-46/protect-scsql/concept_snapcenter_plug_in_for_microsoft_sql_server_overview.html on November 24, 2022. Always check docs.netapp.com for the latest.

Table of Contents

- Protect Microsoft SQL Server databases 1
 - SnapCenter Plug-in for Microsoft SQL Server 1
 - Quick start to install SnapCenter Plug-in for SQL Server 21
 - Prepare to install the SnapCenter Plug-in for Microsoft SQL Server 26
 - Install SnapCenter Plug-in for VMware vSphere 47
 - Prepare for data protection 47
 - Back up SQL Server database, or instance, or availability group 49
 - Restore SQL Server resources 74
 - Clone SQL Server database resources 84

Protect Microsoft SQL Server databases

SnapCenter Plug-in for Microsoft SQL Server

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, verification, restore, and clone operations in your SnapCenter environment.

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.

What you can do with the SnapCenter Plug-in for Microsoft SQL Server

When the SnapCenter Plug-in for Microsoft SQL Server is installed in your environment, you can use SnapCenter to back up, restore, and clone SQL Server databases.

You can perform the following tasks that support backup operations, restore operations, and clone operations of SQL Server databases and database resources:

- Back up SQL Server databases and associated transaction logs

You cannot create a log backup for master and msdb system databases. However, you can create log backups for model system database.

- Restore database resources
 - You can restore master system databases, msdb system databases, and model system databases.
 - You cannot restore multiple databases, instances, and availability groups.
 - You cannot restore the system database to an alternate path.
- Create point-in-time clones of production databases

You cannot perform backup, restore, clone, and clone lifecycle operations on tempdb system databases.

- Verify backup operations immediately or defer verification until later

Verification of SQL Server system database is not supported. SnapCenter clones the databases to perform verification operation. SnapCenter cannot clone SQL Server system databases, and therefore verification of these databases is not supported.

- Schedule backup operations and clone operations
- Monitor backup operations, restore operations, and clone operations



The Plug-in for SQL Server does not support backup and recovery of SQL Server databases on SMB shares.

SnapCenter Plug-in for Microsoft SQL Server features

The Plug-in for SQL Server integrates with Microsoft SQL Server on the Windows host and with NetApp Snapshot copy technology on the storage system. To work with the Plug-in for SQL Server, you use the SnapCenter interface.

The Plug-in for SQL Server includes these major features:

- **Unified graphical user interface powered by SnapCenter**

The SnapCenter interface provides you with standardization and consistency across plug-ins and environments. The SnapCenter interface enables you to complete consistent backup and restore processes across plug-ins, use centralized reporting, use at-a-glance dashboard views, set up role-based access control (RBAC), and monitor jobs across all plug-ins. SnapCenter also offers centralized scheduling and policy management to support backup and clone operations.

- **Automated central administration**

You can schedule routine SQL Server backups, configure policy-based backup retention, and set up point-in-time and up-to-the-minute restore operations. You can also proactively monitor your SQL Server environment by configuring SnapCenter to send email alerts.

- **Nondisruptive NetApp Snapshot copy technology**

The Plug-in for SQL Server uses NetApp Snapshot copy technology with the NetApp SnapCenter Plug-in for Microsoft Windows. This enables you to back up databases in seconds and restore them quickly without taking SQL Server offline. Snapshot copies consume minimal storage space.

In addition to these major features, the Plug-in for SQL Server offers the following benefits:

- Backup, restore, clone, and verification workflow support
- RBAC-supported security and centralized role delegation
- Creation of space-efficient, point-in-time copies of production databases for testing or data extraction by using NetApp FlexClone technology

A FlexClone license is required on the storage system holding the clone.

- Nondisruptive and automated backup verification
- Ability to run multiple backups at the same time across multiple servers
- PowerShell cmdlets for scripting of backup, verify, restore, and clone operations
- Support for AlwaysOn Availability Groups (AGs) in SQL Server to accelerate AG setup, backup, and restore operations
- In-memory database and Buffer Pool Extension (BPE) as part of SQL Server 2014
- Support for backup of LUNs and virtual machine disks (VMDKs)
- Support for physical and virtualized infrastructures
- Support for iSCSI, Fibre Channel, FCoE, raw device mapping (RDM), and VMDK over NFS and VMFS



NAS volumes should have a default export policy in storage virtual machine (SVM).

- Support for FileStream and file group in SQL Server standalone databases.


Support for Asymmetric LUN Mapping in Windows clusters

SnapCenter Plug-in for Microsoft SQL Server supports discovery in SQL Server 2012 and later, Asymmetric LUN Mapping (ALM) configurations for high availability, and availability groups for disaster recovery. When discovering resources, SnapCenter discovers databases on local hosts and on remote hosts in ALM configurations.

An ALM configuration is a single Windows server failover cluster that contains one or more nodes in a primary data center and one or more nodes in a disaster recovery center.

Following is an example of an ALM configuration:

- Two failover cluster instances (FCI) in a multi-site datacenter
- FCI for local high availability (HA) and Availability Group (AG) for disaster recovery with a stand-alone instance at the disaster recovery site



WSFC—Windows Server Failover Cluster

The storage in the primary datacenter is shared between the FCI nodes present in the primary datacenter. The storage in the disaster recovery datacenter is shared between the FCI nodes present in the disaster recovery datacenter.

The storage on the primary datacenter is not visible to the nodes on the disaster recovery datacenter, and vice versa.



ALM architecture combines two shared storage solution used by FCI, with non-shared or dedicated storage solution used by SQL AG. The AG solution uses identical drive letters for shared disk resources across data centers. This arrangement of storage, where a cluster disk is shared between a subset of nodes within a WSFC, is referred to as ALM.


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](#).

Machine	Storage type	Provision using	Support notes
Physical server	FC-connected LUNs	SnapCenter graphical user interface (GUI) or PowerShell cmdlets	
Physical server	iSCSI-connected LUNs	SnapCenter GUI or PowerShell cmdlets	
Physical server	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	
VMware VM	iSCSI LUNs connected directly to the guest system by the iSCSI initiator	SnapCenter GUI or PowerShell cmdlets	
VMware VM	Virtual Machine File Systems (VMFS) or NFS datastores	VMware vSphere	
VMware VM	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.

Machine	Storage type	Provision using	Support notes
Hyper-V VM	Virtual FC (vFC) LUNs connected by a virtual Fibre Channel Switch	SnapCenter GUI or PowerShell cmdlets	<p>You must use Hyper-V Manager to provision Virtual FC (vFC) LUNs connected by a virtual Fibre Channel Switch.</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;">  <p>Hyper-V pass through disks and backing up databases on VHD(x) that are provisioned on NetApp storage are not supported.</p> </div>
Hyper-V VM	iSCSI LUNs connected directly to the guest system by the iSCSI initiator	SnapCenter GUI or PowerShell cmdlets	<div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;">  <p>Hyper-V pass through disks and backing up databases on VHD(x) that are provisioned on NetApp storage are not supported.</p> </div>

Machine	Storage type	Provision using	Support notes
Hyper-V VM	A guest system connected to SMB3 shares residing on an SVM	SnapCenter GUI or PowerShell cmdlets	<p>Support for provisioning only.</p> <p>You cannot use SnapCenter to back up any data or shares using the SMB protocol.</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;">  <p>Hyper-V pass through disks and backing up databases on VHD(x) that are provisioned on NetApp storage are not supported.</p> </div>

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.

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:



Minimum ONTAP privileges required for SQL plug-in

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

All-access commands: Minimum privileges required for ONTAP 8.3.0 and later
event generate-autosupport-log
job history show
job stop

All-access commands: Minimum privileges required for ONTAP 8.3.0 and later

lun

lun create

lun delete

lun igroup add

lun igroup create

lun igroup delete

lun igroup rename

lun igroup show

lun mapping add-reporting-nodes

lun mapping create

lun mapping delete

lun mapping remove-reporting-nodes

lun mapping show

lun modify

lun move-in-volume

lun offline

lun online

lun resize

lun serial

lun show

All-access commands: Minimum privileges required for ONTAP 8.3.0 and later

snapmirror policy add-rule

snapmirror policy modify-rule

snapmirror policy remove-rule

snapmirror policy show

snapmirror restore

snapmirror show

snapmirror show-history

snapmirror update

snapmirror update-ls-set

snapmirror list-destinations

version

All-access commands: Minimum privileges required for ONTAP 8.3.0 and later

volume clone create

volume clone show

volume clone split start

volume clone split stop

volume create

volume destroy

volume file clone create

volume file show-disk-usage

volume offline

volume online

volume modify

volume qtree create

volume qtree delete

volume qtree modify

volume qtree show

volume restrict

volume show

volume snapshot create

volume snapshot delete

volume snapshot modify

volume snapshot rename

volume snapshot restore

volume snapshot restore-file

volume snapshot show

volume unmount

All-access commands: Minimum privileges required for ONTAP 8.3.0 and later

vserver cifs

vserver cifs share create

vserver cifs share delete

vserver cifs shadowcopy show

vserver cifs share show

vserver cifs show

vserver export-policy

vserver export-policy create

vserver export-policy delete

vserver export-policy rule create

vserver export-policy rule show

vserver export-policy show

vserver iscsi

vserver iscsi connection show

vserver show

network interface

network interface show

vserver

metrocluster show

Prepare storage systems for SnapMirror and SnapVault replication for Plug-in for SQL server

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



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

A data protection relationship replicates data on primary storage (the source volume) to secondary storage (the

destination volume). When you initialize the relationship, ONTAP transfers the data blocks referenced on the source volume to the destination volume.



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

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



SnapCenter does not support **sync_mirror** replication.

Backup strategy for SQL Server resources

Define a backup strategy for SQL Server resources

Defining a backup strategy before you create your backup jobs helps ensure that you have the backups that you require to successfully restore or clone your databases. Your Service Level Agreement (SLA), Recovery Time Objective (RTO), and Recovery Point Objective (RPO) largely determine your backup strategy.

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

Type of backups supported

Backing up SQL Server system and user databases using SnapCenter requires that you choose the resource type, such as databases, SQL server instances, and Availability Groups (AG). Snapshot copy technology is leveraged to create online, read-only copies of the volumes on which the resources reside.

You can select the copy-only option to specify that the SQL Server does not truncate transaction logs. You should use this option when you are also managing the SQL Server with other backup applications. Keeping the transaction logs intact enables any backup application to restore the system databases. Copy-only backups are independent of the sequence of scheduled backups, and they do not affect the backup and restore procedures of the database.

Backup type	Description	Copy-only option with backup type
Full backup and log backup	<p>Backs up the system database and truncates the transaction logs.</p> <p>The SQL Server truncates the transaction logs by removing the entries that are already committed to the database.</p> <p>After the full backup is complete, this option creates a transaction log that captures transaction information. Typically, you should choose this option. However, if your backup time is short, you can choose not to run a transaction log backup with full backup.</p> <p>You cannot create a log backup for master and msdb system databases. However, you can create log backups for model system database.</p>	<p>Backs up the system database files and the transaction logs without truncating the logs.</p> <p>A copy-only backup cannot serve as a differential base or differential backup, and does not affect the differential base. Restoring a copy-only full backup is the same as restoring any other full backup.</p>
Full database backup	<p>Backs up the system database files.</p> <p>You can create full database backup for master, model, and msdb system databases.</p>	<p>Backs up the system database files.</p>
Transaction log backup	<p>Backs up the truncated transaction logs, copying only the transactions that were committed since the most recent transaction log was backed up.</p> <p>If you schedule frequent transaction log backups alongside full database backups, you can choose granular recovery points.</p>	<p>Backs up the transaction logs without truncating them.</p> <p>This backup type does not affect the sequencing of regular log backups. Copy-only log backups are useful for performing online restore operations.</p>

Backup schedules for Plug-in for SQL server

Backup frequency (schedule type) is specified in policies; a backup schedule is specified in the resource group configuration. The most critical factor in determining a backup frequency or schedule is the rate of change for the resource and the importance of the data. You might back up a heavily used resource every hour, while you might back up a rarely used resource once a day. Other factors include the importance of the resource to your organization, your Service Level Agreement (SLA), and your Recover Point Objective (RPO).

An SLA defines the level of service expected and addresses many service-related issues, including the

availability and performance of service. An RPO defines the strategy for the age of the files that must be recovered from backup storage for regular operations to resume after a failure. The SLA and RPO contribute to the data protection strategy.

Even for a heavily used resource, there is no requirement to run a full backup more than once or twice a day. For example, regular transaction log backups might be sufficient to ensure that you have the backups you need. The more often you back up your databases, the fewer transaction logs SnapCenter has to use at restore time, which can result in faster restore operations.

Backup schedules have two parts, as follows:

- Backup frequency

Backup frequency (how often backups are to be performed), called *schedule type* for some plug-ins, is part of a policy configuration. You can select hourly, daily, weekly, or monthly as the backup frequency for the policy. If you do not select any of these frequencies, then the policy created is an on-demand-only policy. You can access policies by clicking **Settings > Policies**.

- Backup schedules

Backup schedules (exactly when backups are to be performed) are part of a resource group configuration. For example, if you have a resource group that has a policy configured for weekly backups, you might configure the schedule to back up every Thursday at 10:00 PM. You can access resource group schedules by clicking **Resources > Resource Groups**.

Number of backup jobs needed for databases

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

For database backups, the number of backup jobs that you choose typically depends on the number of volumes on which you placed your databases. For example, if you placed a group of small databases on one volume and a large database on another volume, you might create one backup job for the small databases and one backup job for the large database.

Backup naming conventions for Plug-in for SQL server

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

The Snapshot copy uses the following default naming convention:

```
resourcegroupname_hostname_timestamp
```

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

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

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

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

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

Backup retention options for Plug-in for SQL Server

You can choose either the number of days for which to retain backup copies or specify the number of backup copies you want to retain, up to a ONTAP maximum of 255 copies. For example, your organization might require that you retain 10 days of backup copies or 130 backup copies.

While creating a policy, you can specify the retention options for the backup type and the schedule type.

If you set up SnapMirror replication, the retention policy is mirrored on the destination volume.

SnapCenter deletes the retained backups that have retention labels that match the schedule type. If the schedule type was changed for the resource or resource group, backups with the old schedule type label might still remain on the system.



For long-term retention of backup copies, you should use SnapVault backup.

How long to retain transaction log backups on the source storage system

SnapCenter Plug-in for Microsoft SQL Server needs transaction log backups to perform up-to-the-minute restore operations, which restore your database to a time between two full backups.

For example, if Plug-in for SQL Server took a full backup at 8:00 a.m. and another full backup at 5:00 p.m., it could use the latest transaction log backup to restore the database to any time between 8:00 a.m. and 5:00 p.m. If transaction logs are not available, Plug-in for SQL Server can perform point-in-time restore operations only, which restore a database to the time that Plug-in for SQL Server completed a full backup.

Typically, you require up-to-the-minute restore operations for only a day or two. By default, SnapCenter retains a minimum of two days.

Multiple databases on the same volume

You can put all databases on the same volume, because the backup policy has an option to set the maximum databases per backup (default value is 100).

For example, if you have 200 databases in the same volume, two Snapshot copies are created with 100 databases in each of the two Snapshot copies.

Backup copy verification using the primary or secondary storage volume for Plug-in for SQL Server

You can verify backup copies on the primary storage volume or on either the SnapMirror or SnapVault secondary storage volume. Verification using a secondary storage volume reduces load on the primary storage volume.

When you verify a backup that is either on the primary or secondary storage volume, all the primary and the secondary Snapshot copies are marked as verified.

SnapRestore license is required to verify backup copies on SnapMirror and SnapVault secondary storage volume.

When to schedule verification jobs

Although SnapCenter can verify backups immediately after it creates them, doing so can significantly increase the time required to complete the backup job and is resource intensive. Hence, it is almost always best to schedule verification in a separate job for a later time. For example, if you back up a database at 5:00 p.m. every day, you might schedule verification to occur an hour later at 6:00 p.m.

For the same reason, it is usually not necessary to run backup verification every time you perform a backup. Performing verification at regular but less frequent intervals is usually sufficient to ensure the integrity of the backup. A single verification job can verify multiple backups at the same time.

Restoration strategy for SQL Server

Define a restoration strategy for SQL Server

Defining a restoration strategy for SQL Server enables you to restore your database successfully.

Sources and destinations for a restore operation

You can restore a SQL Server database from a backup copy on either primary or secondary storage. You also can restore the database to different destinations in addition to its original location, enabling you to choose the destination that supports your requirements.

Sources for a restore operation

You can restore databases from primary or secondary storage.

Destinations for a restore operation

You can restore databases to various destinations:

Destination	Description
The original location	By default, SnapCenter restores the database to the same location on the same SQL Server instance.
A different location	You can restore the database to a different location on any SQL Server instance within the same host.

Destination	Description
Original or different location using different database names	You can restore the database with a different name to any SQL Server instance on the same host where the backup was created.



Restore to alternate host across ESX servers for SQL databases on VMDKs (NFS and VMFS datastores) is not supported.

SQL Server recovery models supported by SnapCenter

Specific recovery models are assigned to each database type by default. The SQL Server database administrator can reassign each database to a different recovery model.

SnapCenter supports three types of SQL Server recovery models:

- Simple recovery model

When you use the simple recovery model, you cannot back up the transaction logs.

- Full recovery model

When you use the full recovery model, you can restore a database to its previous state from the point of failure.

- Bulk logged recovery model

When you use the bulk logged recovery model, you must manually re-execute the bulk logged operation. You must perform the bulk logged operation if the transaction log that contains the operation's commit record has not been backed up before restore. If the bulk logged operation inserts 10 million rows in a database, and the database fails before the transaction log is backed up, then the restored database will not contain the rows that were inserted by the bulk logged operation.

Types of restore operations

You can use SnapCenter to perform different types of restore operations on SQL Server resources.

- Restore up-to-the-minute
- Restore to a previous point in time

You can restore up to the minute or restore to a previous point in time in the following situations:

- Restore from SnapMirror or SnapVault secondary storage
- Restore to alternate path (location)



SnapCenter does not support volume-based SnapRestore.

Restore up to the minute

In an up-to-the-minute restore operation (selected by default), databases are recovered up to the point of failure. SnapCenter accomplishes this by performing the following sequence:

1. Backs up the last active transaction log before restoring the database.
2. Restores the databases from the full database backup that you select.
3. Applies all the transaction logs that were not committed to the databases (including transaction logs from the backups from the time the backup was created up to the most current time).

Transaction logs are moved ahead and applied to any selected databases.

An up-to-the-minute restore operation requires a contiguous set of transaction logs.

Because the SnapCenter cannot restore SQL Server database transaction logs from log-shipping backup files (log-shipping enables you to automatically send transaction log backups from a primary database on a primary server instance to one or more secondary databases on separate secondary server instances), you are not able to perform an up-to-the-minute restore operation from the transaction log backups. For this reason, you should use the SnapCenter to back up your SQL Server database transaction log files.

If you do not need to retain up-to-the-minute restore capability for all backups, you can configure your system's transaction log backup retention through the backup policies.

Example of an up-to-the-minute restore operation

Assume that you run the SQL Server backup every day at noon, and on Wednesday at 4:00 p.m. you need to restore from a backup. For some reason, the backup from Wednesday noon failed verification, so you decide to restore from the Tuesday noon backup. After that, if the backup is restored, all the transaction logs are moved forward and applied to the restored databases, starting with those that were not committed when you created Tuesday's backup and continuing through the latest transaction log written on Wednesday at 4:00 p.m. (if the transaction logs were backed up).

Restore to a previous point in time


In a point-in-time restore operation, databases are restored only to a specific time from the past. A point-in-time restore operation occurs in the following restore situations:

- The database is restored to a given time in a backed-up transaction log.
- The database is restored, and only a subset of backed-up transaction logs are applied to it.



Restoring a database to a point in time results in a new recovery path.

The following image illustrates the issues when a point-in-time restore operation is performed:



In the image, recovery path 1 consists of a full backup followed by several transaction log backups. You restore the database to a point in time. New transaction log backups are created after the point-in-time restore operation, which results in recovery path 2. The new transaction log backups are created without creating a new full backup. Due to data corruption or other problems, you cannot restore the current database until a new full backup is created. Also, it is not possible to apply the transaction logs created in recovery path 2 to the full backup belonging to recovery path 1.

If you apply transaction log backups, you can also specify a particular date and time at which you want to stop the application of backed up transactions. To do this, you specify a date and time within the available range and the SnapCenter removes any transactions that were not committed prior to that point in time. You can use this method to restore databases to a point in time before a corruption occurred, or to recover from an accidental database or table deletion.

Example of a point-in-time restore operation

Suppose you make full database backups once at midnight and a transaction log backup every hour. The database crashes at 9:45 a.m., but you still back up the transaction logs of the failed database. You can choose from among these point-in-time restore scenarios:

- Restore the full database backup made at midnight and accept the loss of the database changes made afterward. (Option: None)
- Restore the full database backup and apply all the transaction log backups until 9:45 a.m. (Option: Log until)
- Restore the full database backup and apply transaction log backups, specifying the time you want the transactions to restore from the last set of transaction log backups. (Option: By specific time)

In this case, you would calculate the date and time at which a certain error was reported. Any transactions that were not committed prior to the date and time specified are removed.

Define a cloning strategy for SQL Server

Defining a cloning strategy enables you to clone your database successfully.

1. Review the limitations related to clone operations.
2. Decide the type of clone you require.

Limitations of clone operations

You should be aware of the limitations of clone operations before you clone the databases.

- If you are using any version of Oracle from 11.2.0.4 to 12.1.0.1, the clone operation will be in hung state when you run the *renamedg* command. You can apply the Oracle patch 19544733 to fix this issue.
- Cloning of databases from a LUN that is directly attached to a host (for instance, by using Microsoft iSCSI Initiator on a Windows host) to a VMDK or an RDM LUN on the same Windows host, or another Windows host, or vice versa, is not supported.
- The root directory of the volume mount point cannot be a shared directory.
- If you move a LUN that contains a clone to a new volume, the clone cannot be deleted.

Types of clone operations

You can use SnapCenter to clone either a SQL Server database backup or a production database.

- Clone from a database backup

The cloned database can serve as a baseline for developing new applications and help isolate application errors that occur in the production environment. The cloned database can also be used for recovery from soft database errors.

- Clone lifecycle

You can use SnapCenter to schedule recurring clone jobs that will occur when the production database is not busy.

Quick start to install SnapCenter Plug-in for SQL Server

Provides a condensed set of installation instructions for installing the SnapCenter Server and the SnapCenter Plug-in for Microsoft SQL Server.

Prepare for installation

Domain and workgroup requirements


SnapCenter Server can be installed on systems that are either in a domain or in a workgroup.

If you are using an Active Directory domain, you should use a Domain user with local administrator rights. The Domain user should be a member of the local Administrator group on the Windows host.

If you are using workgroups, you should use a local account that has local administrator rights.

License requirements

The type of licenses you install depends on your environment.

License	Where required
SnapCenter Standard controller-based	<p>Required for FAS or AFF storage controllers</p> <p>SnapCenter Standard license is a controller-based license and is included as part of the premium bundle. If you have the SnapManager Suite license, you also get the SnapCenter Standard license entitlement. If you want to install SnapCenter on a trial basis with FAS or AFF storage, you can obtain a Premium Bundle evaluation license by contacting the sales representative.</p>
SnapCenter Standard capacity-based	<p>Required with ONTAP Select and Cloud Volumes ONTAP</p> <p>If you are a Cloud Volumes ONTAP or ONTAP Select customer, you need to procure a per TB capacity-based license based on the data managed by SnapCenter. By default, SnapCenter ships a built-in 90-day 100 TB SnapCenter Standard capacity-based trial license. For other details, contact the sales representative.</p>
SnapMirror or SnapVault	<p>ONTAP</p> <p>Either SnapMirror or SnapVault license is required if replication is enabled in SnapCenter.</p>
Additional licenses (optional)	See SnapCenter licenses .
SnapCenter Standard licenses (optional)	<p>Secondary destinations</p> <div style="border-left: 1px solid #ccc; padding-left: 10px; margin-left: 20px;"> <p> It is recommended, but not required, that you add SnapCenter Standard licenses to secondary destinations. If SnapCenter Standard licenses are not enabled on secondary destinations, you cannot use SnapCenter to backup resources on the secondary destination after performing a failover operation. However, a FlexClone license is required on secondary destinations to perform clone and verification operations.</p> </div>

Additional requirements

For ONTAP and application plug-in minimum requirements see [Interoperability Matrix Tool](#)

Hosts	Minimum requirements
Operating System (64-bit)	See Interoperability Matrix Tool
CPU	<ul style="list-style-type: none"> • Server host: 4 cores • Plug-in host: 1 core
RAM	<ul style="list-style-type: none"> • Server host: 8 GB • Plug-in host: 1 GB
Hard drive space	<p>Server host:</p> <ul style="list-style-type: none"> • 4 GB for SnapCenter Server software and logs • 6 GB for SnapCenter repository • Each plug-in host: 2 GB for plug-in installation and logs, this is required only if plug-in is installed on a dedicated host.
Third-party libraries	<p>Required on SnapCenter Server host and plug-in host:</p> <ul style="list-style-type: none"> • Microsoft .NET Framework 4.5.2 or later • Windows Management Framework (WMF) 4.0 or later • PowerShell 4.0 or later
Browsers	Chrome, Internet Explorer, and Microsoft Edge

Port type	Default port
SnapCenter port	8146 (HTTPS), bidirectional, customizable, as in the URL <i>https://server:8146</i>
SnapCenter SMCORE communication port	8145 (HTTPS), bidirectional, customizable
Repository database	3306 (HTTPS), bidirectional

Port type	Default port
Windows plug-in hosts	135, 445 (TCP) In addition to ports 135 and 445, the dynamic port range specified by Microsoft should also be open. Remote install operations use the Windows Management Instrumentation (WMI) service, which dynamically searches this port range. For information on the dynamic port range supported, see Service overview and network port requirements for Windows .
SnapCenter Plug-in for Windows	8145 (HTTPS), bidirectional, customizable
ONTAP cluster or SVM communication port	443 (HTTPS), bidirectional; 80 (HTTP), bidirectional The port is used for communication between the SnapCenter Server host, plug-in host, and SVM or ONTAP Cluster.

SnapCenter Plug-in for Microsoft SQL Server requirements

You should have a user with local administrator privileges with local login permissions on the remote host. If you manage cluster nodes, you need a user with administrative privileges to all the nodes in the cluster.

You should have a user with sysadmin permissions on the SQL Server. The plug-in uses Microsoft VDI Framework, which requires sysadmin access.

If you were using SnapManager for Microsoft SQL Server and want to import data from SnapManager for Microsoft SQL Server to SnapCenter, see [Import archived backups](#)

Install SnapCenter Server

Download and install SnapCenter Server

Steps

1. Download the SnapCenter Server installation package from the [NetApp Support Site](#) and then double-click the exe.

After you initiate the installation, all the prechecks are performed and if the minimum requirements are not met appropriate error or warning messages are displayed. You can ignore the warning messages and proceed with installation; however, errors should be fixed.

2. Review the pre-populated values required for the SnapCenter Server installation and modify if required.

You do not have to specify the password for MySQL Server repository database. During SnapCenter Server installation the password is auto generated.



The special character “%” is not supported in the custom path for installation. If you include “%” in the path, installation fails.

3. Click **Install Now**.

Log in to SnapCenter

Steps

1. Launch SnapCenter from a shortcut on the host desktop or from the URL provided by the installation (*https://server:8146* for default port 8146 where SnapCenter Server is installed).
2. Enter the credentials.

For a built-in domain admin username format, use: *NetBIOS\<username>* or *<username>@<domain>* or *<DomainFQDN>\<username>*.

For a built-in local admin username format, use *<username>*.

3. Click **Sign In**.

Add a SnapCenter Standard controller-based license

Steps

1. Log in to the controller using the ONTAP command line and enter:

```
system license add -license-code <license_key>
```

2. Verify the license:

```
license show
```

Add a SnapCenter capacity-based license

Steps

1. In the SnapCenter GUI left pane, click **Settings > Software**, and then in the License section, click **+**.
2. Select one of two methods for obtaining the license:
 - Enter your NetApp Support Site login credentials to import licenses.
 - Browse to the location of the NetApp License File and click **Open**.
3. In the Notifications page of the wizard, use the default capacity threshold of 90 percent.
4. Click **Finish**.

Set up storage system connections

Steps

1. In the left pane, click **Storage Systems > New**.
2. In the Add Storage System page, perform the following:
 - a. Enter the name or IP address of the storage system.
 - b. Enter the credentials that are used to access the storage system.
 - c. Select the check boxes to enable Event Management System (EMS) and AutoSupport.

3. Click **More Options** if you want to modify the default values assigned to platform, protocol, port, and timeout.
4. Click **Submit**.

Install the Plug-in for Microsoft SQL Server

Set up Run As Credentials to install the Plug-in for Microsoft SQL Server

Steps

1. In the left pane, click **Settings > Credentials > New**.
2. Enter the credentials.

For a built-in domain admin username format, use: *NetBIOS\<username>* or *<username>@<domain>* or *<DomainFQDN>\<username>*.

For a built-in local admin username format, use *<username>*.

Add a host and install the Plug-in for Microsoft SQL Server


Steps

1. In the SnapCenter GUI left pane, click **Hosts > Managed Hosts > Add**.
2. In the Hosts page of the wizard, perform the following:
 - a. Host Type: Select Windows host type.
 - b. Host name: Use the SQL host or specify the FQDN of a dedicated Windows host.
 - c. Credentials: Select the valid credential name of the host that you created or create new credentials.
3. In the Select Plug-ins to Install section, select **Microsoft SQL Server**.
4. Click **More Options** to specify the following details:
 - a. Port: Either retain the default port number or specify the port number.
 - b. Installation Path: The default path is *C:\Program Files\NetApp\SnapCenter*. You can optionally customize the path.
 - c. Add all hosts in the cluster: Select this check box if you are using SQL in WSFC.
 - d. Skip preinstall checks: Select this check box if you already installed the plug-ins manually or you do not want to validate whether the host meets the requirements for installing the plug-in.
5. Click **Submit**.

Prepare to install the SnapCenter Plug-in for Microsoft SQL Server

Installation workflow for SnapCenter Plug-in for Microsoft SQL Server

You should install and set up the SnapCenter Plug-in for Microsoft SQL Server if you want to protect SQL Server databases.



Prerequisites to add hosts and install SnapCenter Plug-in for Microsoft SQL Server

Before you add a host and install the plug-ins packages, you must complete all the requirements.

- If you are using iSCSI, the iSCSI service must be running.
- You must have a user with local administrator privileges with local login permissions on the remote host.
- If you manage cluster nodes in SnapCenter, you must have a user with administrative privileges to all the nodes in the cluster.
- You must have a user with sysadmin permissions on the SQL Server.

SnapCenter Plug-in for Microsoft SQL Server uses Microsoft VDI Framework, which requires sysadmin access.

[Microsoft Support Article 2926557: SQL Server VDI backup and restore operations require Sysadmin privileges](#)

- When installing a plug-in on a Windows host, if you specify a credential that is not built-in or if the user belongs to a local workgroup user, you must disable UAC on the host.
- If SnapManager for Microsoft SQL Server is installed, you must have stopped or disabled the service and schedules.

If you plan to import backup or clone jobs into SnapCenter, do not uninstall SnapManager for Microsoft SQL Server.


- The host must be resolvable to the fully qualified domain name (FQDN) from the server.

If the hosts file is modified to make it resolvable and if both the short name and the FQDN are specified in the hosts file, create an entry in the SnapCenter hosts file in the following format: <ip_address>

<host_fqdn> <host_name>

Host requirements to install SnapCenter Plug-ins Package for Windows

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

Item	Requirements
Operating Systems	Microsoft Windows For the latest information about supported versions, see the NetApp Interoperability Matrix Tool .
Minimum RAM for the SnapCenter plug-in on host	1 GB
Minimum install and log space for the SnapCenter plug-in on host	5 GB  You should allocate sufficient disk space and monitor the storage consumption by the logs folder. The log space required varies depending on the number of the entities to be protected and the frequency of data protection operations. If there is no sufficient disk space, the logs will not be created for the recently run operations.
Required software packages	<ul style="list-style-type: none">• Microsoft .NET Framework 4.5.2 or later• Windows Management Framework (WMF) 4.0 or later• PowerShell 4.0 or later For the latest information about supported versions, see the NetApp Interoperability Matrix Tool .

Set up credentials for the SnapCenter Plug-ins Package for Windows

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

What you will need

- You must set up Windows credentials before installing plug-ins.
- You must set up the credentials with administrator privileges, including administrator rights on the remote host.

- SQL authentication on Windows hosts

You must set up SQL credentials after installing plug-ins.

If you are deploying SnapCenter Plug-in for Microsoft SQL Server, you must set up SQL credentials after installing plug-ins. Set up a credential for a user with SQL Server sysadmin permissions.

The SQL authentication method authenticates against a SQL Server instance. This means that a SQL Server instance must be discovered in SnapCenter. Therefore, before adding a SQL credential, you must add a host, install plug-in packages, and refresh resources. You need SQL Server authentication for performing operations such as scheduling or discovering resources.

Steps

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

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

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

5. Click **OK**.

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

Configure credentials for an individual SQL Server resource

You can configure credentials to perform data protection jobs on individual SQL Server resource 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 padlock icon.

If the padlock 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.
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 Username field are:<ul style="list-style-type: none">◦ <i>NetBIOS\UserName</i>◦ <i>Domain FQDN\UserName</i>• 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 Username field is: <i>UserName</i>

For this field...	Do this...
Password	Enter the password used for authentication.
Authentication mode	Select the SQL Server authentication mode. You can also choose Windows authentication if the Windows user has sysadmin privileges on the SQL server.
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.
 - a. Click , and then select the host name to filter the instances.
 - b. Click  to close the filter pane.
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 Plugin for Microsoft SQL Server, then the user has to configure the credentials.



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

9. Click **Refresh Resources**.

Configure gMSA on Windows Server 2012 or later

Windows Server 2012 or later enables you to create a group Managed Service Account (gMSA) that provides automated service account password management from a managed domain account.

What you will need

- You should have a Windows Server 2012 or later domain controller.
- You should have a Windows Server 2012 or later host, which is a member of the domain.

Steps

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

```
domainName\accountName$
```

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

For example,

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

- d. Run `Get-ADServiceAccount` command to verify the service account.

4. Configure the gMSA on your hosts:

- a. Enable the Active Directory module for Windows PowerShell on the host where you want to use the gMSA account.

To do this, run the following command from PowerShell:

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

Display Name                               Name                               Install
State
-----
[ ] Active Directory Domain Services      AD-Domain-Services      Available

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

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

- b. Restart your host.
- c. Install the gMSA on your host by running the following command from the PowerShell command prompt: `Install-AdServiceAccount <gMSA>`
- d. Verify your gMSA account by running the following command: `Test-AdServiceAccount <gMSA>`

5. Assign the administrative privileges to the configured gMSA on the host.

6. Add the Windows host by specifying the configured gMSA account in the SnapCenter Server.

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

Install SnapCenter Plug-in for Microsoft SQL Server

Add hosts and install the SnapCenter Plug-ins Package for Windows

You must use the SnapCenter **Add Host** page to add hosts and install the plug-ins package. The plug-ins are automatically installed on the remote hosts.

What you will need

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

[Configure group Managed Service Account on Windows Server 2012 or later for SQL](#)

About this task

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


You can add a host and install the plug-in packages either for an individual host or for a cluster. If you are installing the plug-ins on a cluster or Windows Server Failover Clustering (WSFC), the plug-ins are installed on all of the nodes of the cluster.

For information on managing hosts, see [Manage hosts](#).

Steps


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


For this field...	Do this...
Host Type	Select Windows as the host type. The SnapCenter Server adds the host, and then installs the Plug-in for Windows if the plug-in is not already installed on the host. If you select the Microsoft SQL Server option on the Plug-ins page, the SnapCenter Server installs the Plug-in for SQL Server.

For this field...	Do this...
Host name	<p>Enter the fully qualified domain name (FQDN) or the IP address of the host. IP address is supported for untrusted domain hosts only if it resolves to the FQDN.</p> <p>SnapCenter depends on the proper configuration of the DNS. Therefore, the best practice is to enter the FQDN.</p> <p>You can enter the IP addresses or FQDN of one of the following:</p> <ul style="list-style-type: none"> • Stand-alone host • WSFC If you are adding a host by using SnapCenter and the host is part of a subdomain, you must provide the FQDN.
Credentials	<p>Select the credential name that you created or create new credentials. The credential must have administrative rights on the remote host. For details, see the information about creating credentials.</p> <p>You can view details about the credentials by positioning your cursor over the credential name that you specified.</p> <div data-bbox="873 1098 928 1150" style="border: 1px solid #ccc; padding: 5px; display: inline-block;">  </div> <p>The credentials authentication mode is determined by the host type that you specify in the Add Host wizard.</p>

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

6. Click **More Options**.

For this field...	Do this...
Port	<p>Either retain the default port number or specify the port number. The default port number is 8145. If the SnapCenter Server was installed on a custom port, that port number will be displayed as the default port.</p> <div data-bbox="873 1692 928 1745" style="border: 1px solid #ccc; padding: 5px; display: inline-block;">  </div> <p>If you manually installed the plug-ins and specified a custom port, you must specify the same port. Otherwise, the operation fails.</p>

For this field...	Do this...
Installation Path	The default path is C:\Program Files\NetApp\SnapCenter. You can optionally customize the path.
Add all hosts in the cluster	Select this check box to add all of the cluster nodes in a WSFC or a SQL Availability Group. You should add all the cluster nodes by selecting the appropriate cluster check box in the GUI if you want to manage and identify multiple available SQL Availability Groups within a cluster.
Skip preinstall checks	Select this check box if you already installed the plug-ins manually and you do not want to validate whether the host meets the requirements for installing the plug-in.
Use group Managed Service Account (gMSA) to run the plug-in services	<p>Select this check box if you want to use group Managed Service Account (gMSA) to run the plug-in services.</p> <p>Provide the gMSA name in the following format: domainName\accountName\$.</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;">  <p>If the host is added with gMSA and if the gMSA has login and sys admin privileges, the gMSA will be used to connect to the SQL instance.</p> </div>

7. Click **Submit**.

8. For SQL Plug-in, select the host to configure the log directory.

- a. Click **Configure log directory** and in the Configure host log directory page, click **Browse** and complete the following steps:

Only NetApp LUNs (drives) are listed for selection. SnapCenter backs up and replicates the host log directory as part of the backup operation.

- i. Select the drive letter or mount point on the host where the host log will be stored.
- ii. Choose a subdirectory, if required.
- iii. Click **Save**.

9. Click **Submit**.

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

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



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

10. Monitor the installation progress.

Install SnapCenter Plug-in for Microsoft SQL Server on multiple remote hosts by using cmdlets

You can install the SnapCenter Plug-in for Microsoft SQL Server on multiple hosts simultaneously by using the Install-SmHostPackage PowerShell cmdlet.

What you will need

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

Steps

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

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

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

4. Enter your credentials for remote installation.

Install the SnapCenter Plug-in for Microsoft SQL Server silently from the command line

You should install SnapCenter Plug-in for Microsoft SQL Server from within the SnapCenter user interface. However, if you cannot for some reason, you can run the Plug-in for SQL Server installation program unattended in silent mode from the Windows command line.

What you will need

- You must delete the earlier version of SnapCenter Plug-in for Microsoft SQL Server before installing.

For more information, see [How to Install a SnapCenter Plug-In manually and directly from the Plug-In Host](#).

Steps

1. Validate whether `C:\temp` folder exists on the plug-in host and the logged in user has full access to it.
2. Download the Plug-in for SQL Server software from `C:\ProgramData\NetApp\SnapCenter\Package Repository`.

This path is accessible from the host where the SnapCenter Server is installed.

3. Copy the installation file to the host on which you want to install the plug-in.
4. From a Windows command prompt on the local host, navigate to the directory to which you saved the plug-in installation files.
5. Install the Plug-in for SQL Server software:

```
"snapcenter_windows_host_plugin.exe"/silent /debuglog"Debug_Log_Path"  
/log"Log_Path" BI_SNAPCENTER_PORT=Num  
SUITE_INSTALLDIR="Install_Directory_Path"  
BI_SERVICEACCOUNT=domain\administrator BI_SERVICEPWD=password  
ISFeatureInstall=SCW, SCSQL
```

Replace the placeholder values with your data

- `Debug_Log_Path` is the name and location of the suite installer log file.
- `Log_Path` is the location of the installation logs of the plug-in components (SCW, SCSQL, and SMCore).
- `Num` is the port on which SnapCenter communicates with SMCore
- `Install_Directory_Path` is the host plug-in package installation directory.
- `domain\administrator` is the SnapCenter Plug-in for Microsoft Windows web service account.

- password is the password for the SnapCenter Plug-in for Microsoft Windows web service account.
"snapcenter_windows_host_plugin.exe"/silent
/debuglog"C:\HPPW_SCSQL_Install.log" /log"C:\\" BI_SNAPCENTER_PORT=8145
SUITE_INSTALLDIR="C:\Program Files\NetApp\SnapCenter"
BI_SERVICEACCOUNT=domain\administrator BI_SERVICEPWD=password
ISFeatureInstall=SCW,SCSQL



All the parameters passed during the installation of Plug-in for SQL Server are case sensitive.

6. Monitor the Windows task scheduler, the main installation log file C:\Installdebug.log, and the additional installation files in C:\Temp.
7. Monitor the %temp% directory to verify that the msix.exe installers are installing the software without errors.



The installation of Plug-in for SQL Server registers the plug-in on the host and not on the SnapCenter Server. You can register the plug-in on the SnapCenter Server by adding the host using the SnapCenter GUI or PowerShell cmdlet. After the host is added, the plug-in is automatically discovered.

Monitor the status of installing Plug-in for SQL Server

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

About this task

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

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

Steps

1. In the left navigation pane, click **Monitor**.
2. In the Monitor page, click **Jobs**.
3. In the Jobs page, to filter the list so that only plug-in installation operations are listed, do the following:
 - a. Click **Filter**.
 - b. Optional: Specify the start and end date.
 - c. From the Type drop-down menu, select **Plug-in installation**.
 - d. From the Status drop-down menu, select the installation status.
 - e. Click **Apply**.

4. Select the installation job and click **Details** to view the job details.
5. In the Job Details page, click **View logs**.

Configure CA Certificate

Generate CA Certificate CSR file

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

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

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



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

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

Import CA certificates

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

Steps

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

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

In this wizard window...	Do the following...
Completing the Certificate Import Wizard	Review the summary, and then click Finish to start the import.



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

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

Get the CA certificate thumbprint

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

Steps

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

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

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

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

- b. Copy the thumbprint.

Configure CA certificate with Windows host plug-in services

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

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

Steps

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

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

For example:

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

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

```
> $cert = "<certificate thumbprint>"
```

```
> $guid = [guid]::NewGuid().ToString("B")
```

```
> netsh http add sslcert ipport=0.0.0.0: <SMCore Port> certhash=$cert  
appid="$guid"
```

For example:

```
> $cert = "a909502dd82ae41433e6f83886b00d4277a32a7b"  
> $guid = [guid]::NewGuid().ToString("B")  
> netsh http add sslcert ipport=0.0.0.0:8145 certhash=$cert  
appid="$guid"
```

Enable CA Certificates for plug-ins

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

What you will need

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

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




Steps

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

After you finish

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

-  indicates that the CA certificate is neither enabled nor assigned to the plug-in host.

-  indicates that the CA certificate is successfully validated.
-  indicates that the CA certificate could not be validated.
-  indicates that the connection information could not be retrieved.



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

Configure Disaster recovery

Disaster recovery of SnapCenter Plug-in for SQL Server

When the SnapCenter Plug-in for SQL Server is down, switch to a different SQL host and recover the data by performing few steps.

What you will need

- The secondary host should have the same operating system, application, and host name as the primary host.
- Push the SnapCenter Plug-in for SQL Server to an alternative host using **Add host** or **Modify host** page.

[Manage hosts](#)

Steps

1. Select the host from the **Hosts** page to modify and install the SnapCenter Plug-in for SQL Server.
2. (Optional) Replace the SnapCenter Plug-in for SQL Server configuration files from disaster recovery (DR) backup to the new machine.
3. Import Windows and SQL schedules from the SnapCenter Plug-in for SQL Server folder from the DR backup.

For more information, see the [Disaster Recovery APIs](#) video.

Storage disaster recovery (DR) for SnapCenter Plug-in for SQL Server

You can recover the SnapCenter Plug-in for SQL Server storage by enabling the DR Mode for Storage in the Global Settings page.

What you will need

- Ensure that the plug-ins are in maintenance mode.
- Break the SnapMirror/SnapVault relationship. [Breaking SnapMirror relationships](#)
- Attach the LUN from secondary to the host machine with same drive letter.
- Ensure that all the disks are connected using the same drive letters that were used prior to DR.
- Restart MSSQL server service.
- Ensure that the SQL resources are back online.

About this task

Disaster recovery (DR) is not supported on VMDK and RDM configurations.

Steps

1. In the Settings page, navigate to **Settings > Global Settings > Disaster Recovery**.
2. Select **Enable Disaster Recovery**.
3. Click **Apply**.
4. Verify whether the DR job is enabled or not by clicking **Monitor > Jobs**.

After you finish

- If new databases are created after the failover, the databases will be in non-DR mode.

The new databases will continue to operate like they did before the failover.

- The new backups that were created in DR mode will be listed under SnapMirror or SnapVault (secondary) in the Topology page.

An "i" icon is displayed next to the new backups to indicate that these backups were created during DR mode.

- You can delete the SnapCenter Plug-in for SQL Server backups that were created during failover either by using the UI or the following cmdlet: `Remove-SmBackup`
- After failover, if you want some of the resources to be in non-DR mode, use the following cmdlet: `Remove-SmResourceDRMode`

For more information refer to the [SnapCenter Software Cmdlet Reference Guide](#).

- SnapCenter Server will manage the individual storage resources (SQL databases) that are in DR or non-DR mode but not the resource group with storage resources that are in DR mode or non-DR mode.

Failback from SnapCenter Plug-in for SQL Server secondary storage to primary storage

After the SnapCenter Plug-in for SQL Server primary storage is back online, you should failback to the primary storage.

What you will need

- Place the SnapCenter Plug-in for SQL Server in **Maintenance** mode from the Managed Hosts page.
- Disconnect the secondary storage from the host and connect from the primary storage.
- To failback to the primary storage, ensure that the relationship direction remains the same as it was before the failover by performing the reverse resync operation.

To retain the roles of primary and secondary storage after the reverse resync operation, perform the reverse resync operation once again.

For more information see [Reverse resynchronizing mirror relationships](#)

- Restart MSSQL server service.
- Ensure that the SQL resources are back online.

Steps

1. In the Settings page, navigate to **Settings > Global Settings > Disaster Recovery**.
2. Unselect **Enable Disaster Recovery**.
3. Click **Apply**.
4. Verify whether the DR job is enabled or not by clicking **Monitor > Jobs**.

After you finish

- You can delete the SnapCenter Plug-in for SQL Server backups that were created during failover either by using the UI or the following cmdlet: `Remove-SmDRFailoverBackups`

Import archived backups from SnapManager for Microsoft SQL Server to SnapCenter

Import archived backups overview

Importing data from SnapManager for Microsoft SQL Server to SnapCenter enables you to continue to use your data from previous versions. You can import only those backups that were archived using SnapVault technology from SnapManager for Microsoft SQL Server to SnapCenter.

SnapCenter does not support Data ONTAP operating in 7-Mode. You can use the 7-Mode Transition Tool to migrate data and configurations that are stored on a system running Data ONTAP operating in 7-Mode to an ONTAP system.

Limitations related to the import feature

Understanding limitations before you import archived backups from SnapManager for Microsoft SQL Server to SnapCenter will help you complete the import successfully. These backups must have been archived with SnapVault technology.

- You cannot manage clones created using SnapManager for Microsoft SQL Server in SnapCenter. You must manage the clones in SnapManager for Microsoft SQL Server.
- You cannot import backups from the SnapCenter graphical user interface (GUI).
- You cannot import databases from a virtual machine with a VMDK disk.

Import archived backups

You can import only those backups that were archived using SnapVault technology from SnapManager for Microsoft SQL Server to SnapCenter from the command-line interface.

What you will need

- Both the SnapCenter Plug-in for Microsoft SQL Server and SnapManager for Microsoft SQL Server 7.x must be installed on the same host.
- For cluster support, both the plug-in and 7.x product must have been installed on each node in the cluster.
- You must have added the host or cluster to SnapCenter and discovered the resources available for backup on the host or cluster.
- You must have backed up the SnapManager for Microsoft SQL Server SnapInfo folder.

If the import operation fails, you can retrieve the backup metadata from the SnapInfo directory.

Best Practice: It is a best practice to configure the log backup folder for the SnapCenter Plug-in for Microsoft SQL Server.

For information on SnapCenter installation and configuration tasks, see the SnapCenter installation information. For information on SnapManager for Microsoft SQL Server installation and configuration tasks, see the SnapManager installation information.

After the archived backups have been imported to SnapCenter from SnapManager for Microsoft SQL Server, you can perform restore and clone operations on these backups from SnapCenter. Simultaneously, you can also continue to perform backup, restore, and clone operations using SnapManager for Microsoft SQL Server.

Steps

1. Perform the following steps to create schedules similar to SnapManager for Microsoft SQL Server schedules manually in SnapCenter:
 - a. Retrieve the schedule information from SnapManager for Microsoft SQL Server by using the `Get-SmSchedule` cmdlet.
 - b. Save the schedule information.

You can use this schedule information while creating policies in SnapCenter.

- c. Disable SnapManager for Microsoft SQL Server schedules by using the `Get-SmSchedule -Hostname <hostname> -PluginCode SMSQL -DisableCurrentGenSchedule` cmdlet.
2. Import archived backups by using the `Invoke-SmBackupMigration` cmdlet.

For example, to migrate standalone SQL host: `Invoke-SmBackupMigration -Hostname <hostname> -PluginCode SMSQL`

For example, to migrate cluster host: `Invoke-SmBackupMigration -Hostname <clustername> -PluginCode SMSQL`

View the imported backups in SnapCenter Server

After you have imported your SnapManager for Microsoft SQL Server archived backups, you can view these backups in SnapCenter from the Topology page.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. Create a resource group with all the databases.
3. Select the resource group.
4. Protect the resource group.
5. In the Topology page, from Manage Copies view, click backups from Vault copies.

The details of the backups are displayed.



When the archived Snapshot copies are migrated and they get cleaned up from the storage as part of the secondary retention, the respective primary Snapshot copies that are created in SnapManager must be cleaned up manually.

Install SnapCenter Plug-in for VMware vSphere

If your database is stored on virtual machines (VMs), or if you want to protect VMs and datastores, you must deploy the SnapCenter Plug-in for VMware vSphere virtual appliance.

For information to deploy, see [Deployment Overview](#).

Deploy CA certificate

To configure the CA Certificate with SnapCenter Plug-in for VMware vSphere, see [Create or import SSL certificate](#).

Configure the CRL file

SnapCenter Plug-in for VMware vSphere looks for the CRL files in a pre-configured directory. Default directory of the CRL files for SnapCenter Plug-in for VMware vSphere is `/opt/netapp/config/crl`.

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

Prepare for data protection

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.



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 [SnapCenter Software 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 documentation](#)

- 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.

[Import archived backups from SnapManager for Microsoft SQL Server to SnapCenter](#)

- 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 documentation](#)

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 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:



Back up SQL Server database, or instance, or availability group


Backup workflow

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:



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 [SnapCenter Software 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.



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.

Find more information

[Back up resources using PowerShell cmdlets](#)

[Quiesce or grouping resources operations fail](#)

Determine 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.

What you will need

- 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 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 condition is 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.

For more information, see [Deploy SnapCenter Plug-in for VMware vSphere](#)

- If the host is added with gMSA and if the gMSA has login and system admin privileges, the gMSA will be used to connect to the SQL instance.

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.



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.



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

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.



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

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.

Migrate 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.


What you will need

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

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.  This option is not displayed when you select the Instance resource.

For this field...	Do this...
Options	<p>Select Delete copy of Migrated Database at Original Location to delete copy of database from the source.</p> <p>Optional: RUN UPDATE STATISTICS on tables before detaching the database.</p>

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	<ul style="list-style-type: none"> • 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. <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> <div style="border: 1px solid #ccc; padding: 5px; margin: 10px 0;">  You might want to select this option to decrease the execution time. </div> <ul style="list-style-type: none"> • 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**.

Create 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.

What you will need

- You must have defined your data protection strategy.
- 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. In the Name page, enter the policy name and description.
5. In 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 transaction logs and to 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. <div data-bbox="945 1696 1003 1759"></div> <div data-bbox="1058 1684 1458 1780">You must increase this value if you want to run multiple backup operations concurrently.</div>

If you want to...	Do this...
Back up the database files	i. Select Full backup . ii. Enter the maximum number of databases that should be backed up for each Snapshot copy. Default value is 100  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.



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.  By default, the minimum priority is 1 and maximum priority is 100.



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



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.





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


6. In 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> days option, and specify the number of days 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> <div style="border: 1px solid #ccc; padding: 5px; margin: 10px 0;"> <p> 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> </div> <div style="border: 1px solid #ccc; padding: 5px; margin: 10px 0;"> <p> By default, the value of retention count is set to 2. 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> </div>
Keep Snapshot copies for	<p>If you want to specify the number of days for which you want to keep the Snapshot copies before deleting them, select Keep Snapshot copies for.</p>

- 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 in 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. If the number of Snapshot copies exceeds the specified number, the Snapshot copies are deleted with the oldest copies deleted first.</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p> You must set the retention count to 2 or higher if you plan to enable SnapVault replication. If you set the retention count to 1, the retention operation might fail because the first Snapshot copy is the reference Snapshot copy for the SnapVault relationship until a newer Snapshot copy is replicated to the target.</p> </div>
Keep Snapshot copies for	<p>If you want to specify the number of days for which you want to keep the Snapshot copies before deleting them, select Keep Snapshot copies for.</p>

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},@{BackupT
ype='LOG_SNAPSHOT';ScheduleType='None';RetentionCount=2},@{BackupType
='LOG';ScheduleType='Hourly';RetentionCount=2} -schedulescheduletype 'Hourly'
```

SnapCenter retains Snapshot copies of the database

7. In 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).
Update SnapVault after creating a Snapshot copy	Select this option to perform disk-to-disk backup replication.

For this field...	Do this...
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> <div style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;"> <p> 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> </div>
Error retry count	Enter the number of replication attempts that should occur before the process halts.

8. In 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.



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

9. In the Verification page, perform the following steps:
 - a. In the Run verification for following backup schedules section, select the schedule frequency.
 - b. 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.

For this field...	Do this...
Do not check nonclustered indexes (NOINDEX)	Select Do not check nonclustered indexes (NOINDEX) 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.
Limit the checks and obtain the locks instead of using an internal database Snapshot copy (TABLOCK)	Select Limit the checks and obtain the locks 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**.

Create resource groups and attach 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.

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.



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

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

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

5. In 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.



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. In the Policies page, perform the following steps:

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



You can also create a policy by clicking  .

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

- b. 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 *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 and rename the backup job created in Windows scheduler or SQL Server agent.

7. In 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).



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

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



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

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

Find more information

[Create backup policies for SQL Server databases](#)

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 `FMaxRetryForUninitializedHosts` in `web.config` by using the `Set-SmConfigSettings` PS cmdlet.

Back up SQL resources

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

About this task

- For Windows credentials authentication, you must set up your credential before installing the plug-ins.
- For SQL Server instance authentication, you must add the credential after installing the plug-ins.
- For gMSA authentication, you must setup gMSA while registering the host with SnapCenter in the **Add Host** or **Modify Host** page to enable and use the gMSA.
- If the host is added with gMSA and if the gMSA has login and system admin privileges, the gMSA will be used to connect to the SQL instance.

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.
 - 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. In 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. In 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. In 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).



The verification server version should be equal or above the version of the 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 .
Schedule a verification	Select Run scheduled verification .



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

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



For email notification, you must have specified the SMTP server details 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. In 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**.



You should not rename the backup job created in Windows scheduler or SQL Server agent.

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 respective user or group from 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`.

Find more information

[Create backup policies for SQL Server databases](#)

[Back up resources using PowerShell cmdlets](#)

[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](#)

Back 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. In the Resource Groups page, select the resource group that you want to back up, and then click **Back up Now**.
4. In 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**.

Find more information

[Create backup policies for SQL Server databases](#)

[Create resource groups and attach policies for SQL Server](#)

[Back up resources using PowerShell cmdlets](#)

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

[Backup fails with Windows scheduler error](#)







Monitor backup operations

Monitor SQL resources backup operations in the SnapCenter Jobs page


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

About this task


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

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

Steps

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



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

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


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

Monitor data protection operations on SQL resources in the Activity pane

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

The Activity pane displays information regarding backup, restore, clone, and scheduled backup operations. If you are using Plug-in for SQL Server or Plug-in for Exchange Server, the Activity pane also displays information about the reseed operation.

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.

Create 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.

What you will need

- 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 management LIF IP address.

Steps

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

This example opens a PowerShell session:

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

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_vs1 -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
```

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

Back up 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.

What you will need

- 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.

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:

```
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 backup 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'
```

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

Cancel 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.

What you will need


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

View 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.



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.



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  .

Remove 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.

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

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

Steps

1. 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"):
```

Clean 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 to use this cmdlet when the total Snapshot copies displayed in the Manage Copies topology do not match the secondary storage Snapshot copy retention setting.

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

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

Steps

1. 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:

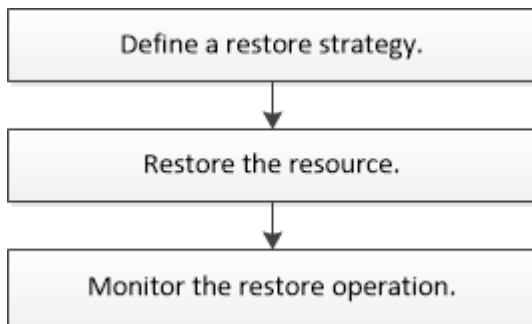
```
Remove-SmBackup -CleanupSecondaryBackups
Remove-SmBackup
Are you sure want to remove the backup(s) .
[Y] Yes [A] Yes to All [N] No [L] No to All [S] Suspend [?] Help
(default is "Y"):
```

Restore SQL Server resources

Restore workflow

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 [SnapCenter Software 4.4 Cmdlet Reference Guide](#)

Find more information

[Restore an SQL Server database from secondary storage](#)

[Restore and recover resources using PowerShell cmdlets](#)

[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 target 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" = "20000000"}`. 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.

Restore SQL Server database backups

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. In 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.

Option	Description
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> <div style="border: 1px solid gray; padding: 5px; margin: 10px 0;">  The file extension provided in the alternate path must be same as the file extension of the original database file. </div> <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.</p> <p>Database files should be already present on the given existing file paths. Select a host name, provide a database name (optional), select an instance, and specify the restore paths.</p>

7. In 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.
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	<p>Select By specific date until to specify the date and time after which transaction logs are not applied to the restored database.</p> <p>This point-in-time restore operation halts the restoration of transaction log entries that were recorded after the specified date and time.</p>

Option	Description
Use custom log directory	<p>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.</p> <div style="display: flex; align-items: center; margin-top: 10px;">  <p>The custom log directory is not supported for availability group database.</p> </div>

8. In the Pre Ops page, perform the following steps:

a. In 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. In 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. 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.

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

Find more information

[Restore and recover resources using PowerShell cmdlets](#)

[Restore an SQL Server database from secondary storage](#)

Restore 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.


What you will need

- 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.
- 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 **SnapCenter Plug-in for SQL Server** from the list.
2. In 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. In the Location page, choose the destination volume for restoring selected resource.
7. Complete the Restore wizard, review the summary, and then click **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.

Reseed 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.

What you will need

- 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 **SnapCenter Plug-in for SQL Server** 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**.

Restore resources using PowerShell cmdlets

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

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

Steps

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

```
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:

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

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

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

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

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

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

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

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

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







Monitor SQL resources restore operations

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


About this task

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

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


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

Steps

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

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



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 restore 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.

Cancel SQL resources restore operations

You can cancel restore jobs that are queued.


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

About this task

- You can cancel a queued 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 queued 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 restore 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">In the left navigation pane, click Monitor > Jobs.Select the job and click Cancel Job.
Activity pane	<ol style="list-style-type: none">After initiating the restore operation, click  on the Activity pane to view the five most recent operations.Select the operation.In the Job Details page, click Cancel Job.


Clone SQL Server database resources

Clone workflow

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 [SnapCenter Software Cmdlet Reference Guide](#)

Find more information

[Clone from a SQL Server database backup](#)

[Perform Clone Lifecycle](#)

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

Clone 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.

What you will need

- 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.



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 **SnapCenter Plug-in for SQL Server** from the list.
2. In the Resources page, select either **Database** or **Resource Group** from the **View** list.



Cloning of a backup of an instance is not supported.

Steps

1. Select the database or resource group.
2. From the Manage Copies view page, select the backup either from primary or secondary (mirrored or vaulted) storage system.
3. Select the backup, and then click .
4. In 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.
Clone suffix	Enter a suffix that will be appended to the clone file name to identify that the database is a clone. For example, <i>db1_clone</i> . 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.

5. In 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.

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

- In 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.

- 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 clone operation performed, select **Attach Job Report**.



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

For EMS, you can refer to [Manage EMS data collection](#)

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

After you finish

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

Find more information

[Back up SQL Server database, or instance, or availability group](#)

[Clone backups using PowerShell cmdlets](#)

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

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

Clone backups using PowerShell cmdlets

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

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

Steps

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

```
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       Full Backup
              11:02:32 AM
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 -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}

```

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

Perform 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.

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

What you will need

- 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.



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.
Clone suffix	Enter a suffix that will be appended to the clone database to identify that it is a clone. 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 “db1clone”. “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.

For this field...	Do this...
Auto assign mount point or Auto assign volume mount point under path	Choose whether to automatically assign a mount point or volume mount point under a path. 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.

- 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 clone operation performed, select **Attach Job Report**.



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

For EMS, you can refer to [Manage EMS data collection](#)

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

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






Monitor SQL database clone operations

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


About this task

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

- In progress

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

Steps

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

Cancel SQL resource clone operations

You can cancel clone operations that are queued.

You should be logged in as the SnapCenter Admin or job owner to cancel clone operations.


About this task

- You can cancel a queued 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 queued 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 clone 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.

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

Split a clone

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

About this task

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

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

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

- When you split a clone, the backup copies and clone jobs of the clone are deleted.
- For information about clone split operation limitations, see [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.
- In the Resources page, select the appropriate option from the View list:

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

- 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. You should run the Stop-SmJob cmdlet to stop the clone split operation, and then retry the clone split operation.

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

For example:

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

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

Find more information

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

Copyright information

Copyright © 2022 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

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

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

Trademark information

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