



SnapCenter® Software 4.2

Concepts Guide

August 2019 | 215-14389_A0
doctrcomments@netapp.com

 **NetApp®**

Contents

Deciding whether to read about SnapCenter Concepts information	6
SnapCenter overview	7
SnapCenter architecture	8
SnapCenter components	9
SnapCenter Server	10
SnapCenter plug-ins	10
SnapCenter repository	12
SnapCenter basics	14
Security	15
Configuration Checker	15
Resources, resource groups, and policies	16
SnapCenter role-based access control (RBAC)	17
Types of role-based access control in SnapCenter	17
Role-based access control permissions and roles	18
Pre-defined SnapCenter roles and permissions	19
Prescripts and postscripts	22
SnapCenter REST APIs	23
SnapCenter backup workflow overview	24
SnapCenter Plug-in for Microsoft Exchange Server overview	25
What you can do with SnapCenter Plug-in for Microsoft Exchange Server	25
Defining a backup strategy for Exchange Server resources	25
Types of backups supported	26
Backup schedules for database plug-ins	26
Number of backup jobs needed for databases	27
Backup naming conventions	27
Backup retention options	27
How long to retain transaction log backups on the source storage volume	28
Defining a restore strategy for Exchange databases	28
Sources for a restore operation	28
Types of restore operations	28
SnapCenter Plug-in for Microsoft SQL Server overview	30
What you can do with the SnapCenter Plug-in for Microsoft SQL Server	30
SnapCenter Plug-in for Microsoft SQL Server features	31
Support for Asymmetric LUN Mapping in Windows clusters	32
Defining a backup strategy for SQL Server resources	33
Type of backups supported	33
Backup schedules for database plug-ins	34
Number of backup jobs needed for databases	35
Backup naming conventions	35
Backup retention options	35

How long to retain transaction log backups on the source storage system ...	36
Multiple databases on the same volume	36
Backup copy verification using the primary or secondary storage volume	36
When to schedule verification jobs	36
Defining a restoration strategy for SQL Server	36
Sources and destinations for a restore operation	37
SQL Server recovery models supported by SnapCenter	37
Types of restore operations	37
Defining a cloning strategy for SQL Server	39
Limitations of clone operations	40
Types of clone operations	40
SnapCenter Plug-in for Microsoft Windows overview	41
What you can do with the SnapCenter Plug-in for Microsoft Windows	41
SnapCenter Plug-in for Windows features	41
How SnapCenter backs up Windows file systems	42
Defining a backup strategy for Windows file systems	43
Backup schedules for Windows file systems	43
Number of backups needed for Windows file systems	44
Backup naming convention	44
Backup retention options	44
Sources and destinations of clones for Windows file systems	44
SnapCenter Plug-in for Oracle Database overview	46
What you can do with the SnapCenter Plug-in for Oracle Database	46
SnapCenter Plug-in for Oracle Database features	46
Defining a backup strategy for Oracle databases	47
Supported Oracle database configurations for backups	48
Types of backup supported	49
Understanding the discovery of Oracle databases	49
Preferred nodes in RAC setup	50
Backup cataloging with Oracle Recovery Manager	51
Backup schedules for database plug-ins	52
Backup naming conventions	53
Backup retention options	53
Backup copy verification using the primary or secondary storage volume	53
Defining a restore and recovery strategy for Oracle databases	54
Types of Oracle database backups supported for restore and recovery operations	54
Types of restore methods supported for Oracle databases	55
Types of restore operations supported for Oracle databases	57
Types of recovery operations supported for Oracle databases	57
Limitations related to restore and recovery operations	57
Sources and destinations for restore operations	58
Defining a clone strategy for Oracle databases	58

Types of Oracle database backups supported for cloning	58
Types of cloning supported	59
Clone naming conventions	59
Limitations of clone operations	59
SnapCenter Plug-in for SAP HANA Database overview	60
What you can do using the SnapCenter Plug-in for SAP HANA Database	60
SnapCenter Plug-in for SAP HANA Database features	60
Defining a backup strategy for SAP HANA databases	61
Type of backups supported	62
How SnapCenter Plug-in for SAP HANA Database uses consistency group Snapshot copies	62
How SnapCenter manages housekeeping of log and data backups	62
Considerations for determining backup schedules for SAP HANA database	62
Number of backup jobs needed for SAP HANA databases	63
Backup naming conventions	63
Types of restore strategies	63
SnapCenter Custom Plug-ins overview	64
What you can do with the SnapCenter Custom Plug-ins	64
SnapCenter Custom Plug-ins features	64
Defining a backup strategy	65
Backup schedules of custom plug-in resources	66
Number of backup jobs needed	66
Backup naming conventions	66
Types of restore strategies	67
SnapCenter Plug-in for VMware vSphere	68
What you can do with SnapCenter Plug-in for VMware vSphere	68
SnapCenter Plug-in for VMware vSphere features	69
When to use the SnapCenter GUI and the vCenter GUI	70
Restoring VMs, VMDKs, files, and folders from backups	71
VSC host migration to SnapCenter	72
Copyright	73
Trademark	74
How to send comments about documentation and receive update notifications	75
Index	76

Deciding whether to read about SnapCenter Concepts information

This information describes the concepts related to SnapCenter, including architecture, features, and concepts related to the SnapCenter plug-ins.

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

- High level overview of how to get started with SnapCenter in the Getting Started section in the SnapCenter graphical user interface (GUI).
- SnapCenter Server and plug-in installation and setup
[*Installing and setting up SnapCenter*](#)
[*Getting Started*](#)
- Other SnapCenter Data Protection Guides for specific types of resources, such as Microsoft SQL Server, Oracle, Windows file systems, and custom plug-ins
- SnapCenter PowerShell cmdlets or Linux commands
[*SnapCenter Software 4.2 Windows Cmdlet Reference Guide*](#)
[*SnapCenter Software 4.2 Linux Command Reference Guide*](#)
- SnapCenter administration, including dashboards, reporting capabilities, and REST APIs, and managing licenses, storage connections, and the SnapCenter Server repository
[*Performing administrative tasks*](#)

SnapCenter overview

SnapCenter Software is a simple, centralized, scalable platform that provides application-consistent data protection for applications, databases, host file systems, and VMs running on ONTAP systems anywhere in the Hybrid Cloud.

SnapCenter leverages NetApp Snapshot, SnapRestore, FlexClone, SnapMirror, and SnapVault technologies to provide the following:

- Fast, space-efficient, application-consistent, disk-based backups
- Rapid, granular restore and application-consistent recovery
- Quick, space-efficient cloning

SnapCenter includes both the rSnapCenter Server and individual lightweight plug-ins. You can automate deployment of plug-ins to remote application hosts, schedule backup, verification, and clone operations, and monitor all data protection operations.

SnapCenter can be deployed in the following ways:

- On premise to protect the following:
 - Data that is on ONTAP FAS or AFF primary systems and replicated to ONTAP FAS or AFF secondary systems
 - Data that is on ONTAP Select primary systems
- On premise in a Hybrid Cloud to protect the following:
 - Data that is on ONTAP FAS or AFF primary systems and replicated to Cloud Volumes ONTAP or NetApp Private Storage secondary systems
- In a public cloud to protect the following:
 - Data that is on Cloud Volumes ONTAP (formerly ONTAP Cloud) primary systems

SnapCenter includes the following key features:

- Centralized, application-consistent data protection
Data protection is supported for Microsoft Exchange Server, Microsoft SQL Server, Oracle Databases on Linux, SAP HANA database, and Windows Host Filesystems running on ONTAP systems.
Data protection is also supported for other standard or custom applications and databases by providing a framework to create user-defined SnapCenter plug-ins. This enables data protection for other applications and databases from the same single-pane-of-glass. By leveraging this framework, NetApp has released SnapCenter custom plug-ins for IBM DB2, MongoDB, MySQL etc. on the NetApp Automation Store.
- Policy-based backups
Policy-based backups leverage NetApp Snapshot copy technology to create fast, space-efficient, application-consistent, disk-based backups. Optionally, you can automate protection of these backups to secondary storage by updates to existing protection relationships.
- Back ups for multiple resources
You can back up multiple resources (applications, databases, or host file systems) of the same type, at the same time, by using SnapCenter resource groups.
- Restore and recovery

SnapCenter provides rapid, granular restores of backups and application-consistent, time-based recovery. You can restore from any destination in the Hybrid Cloud.

- **Cloning**
SnapCenter provides quick, space-efficient, application-consistent cloning, which enables accelerated software development. You can clone on any destination in the Hybrid Cloud.
- **Single user management graphical user interface (GUI)**
The SnapCenter GUI provides a single, one-stop interface for managing backups and clones of a resource in any destination in the Hybrid Cloud.
- **REST APIs, Windows cmdlets, Linux commands**
SnapCenter includes REST APIs for most functionality for integration with any orchestration software, and use of Windows PowerShell cmdlets and a Linux command-line interface.
- **Centralized data protection Dashboard and reporting**
- **Role-Based Access Control (RBAC) for security and delegation.**
- **Repository database with High Availability**
SnapCenter provides a built-in repository database with High Availability to store all backup metadata.
- **Automated push install of plug-ins**
You can automate a remote push of SnapCenter plug-ins from the SnapCenter Server host to application hosts.
- **Load balancing and High Availability**
Load balancing and High Availability for the SnapCenter Server is provided by an integration with Application Request Routing (ARR) and Microsoft Windows Network Load Balancing (NLB), with support for horizontal scaling.

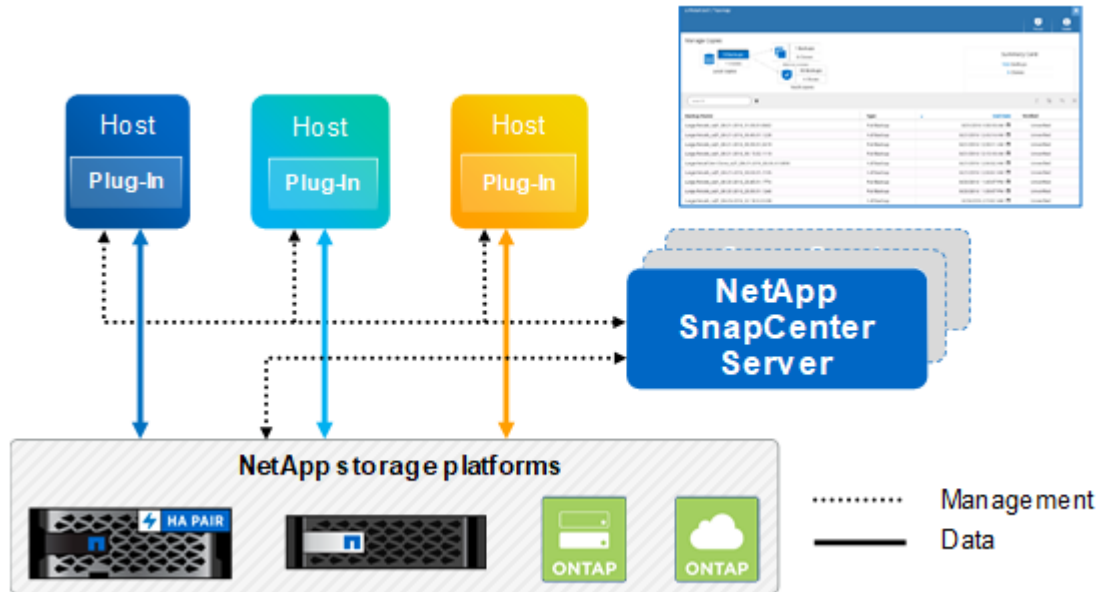
The SnapCenter 4.1.1 documentation has information on protecting virtualized databases and file systems using the SnapCenter 4.1.1 Plug-in for VMware vSphere. The NetApp Data Broker, SnapCenter Plug-in for VMware vSphere, documentation has information on protecting virtualized databases and file systems using the Linux-based NetApp Data Broker virtual appliance (Open Virtual Appliance format) for SnapCenter 4.2.

- [*Deployment Guide for SnapCenter Plug-in for VMware vSphere*](#)
- [*Data Protection Guide for VMs, Datastores, and VMDKs using the SnapCenter Plug-in for VMware vSphere*](#)
- [*NetApp Data Broker Release Notes*](#)

SnapCenter architecture

The SnapCenter platform is based on a multitiered architecture that includes a centralized management server (SnapCenter Server) and a SnapCenter plug-in host.

SnapCenter supports multisite data center. The SnapCenter Server and the plug-in host can be at different geographical locations.



SnapCenter components

SnapCenter consists of the SnapCenter Server and SnapCenter plug-ins. You should install only the plug-ins that are appropriate for the data you want to protect.

SnapCenter includes the following components:

- SnapCenter Server
- SnapCenter Plug-ins Package for Windows, which includes the following plug-ins:
 - SnapCenter Plug-in for Microsoft SQL Server
 - SnapCenter Plug-in for Microsoft Windows
 - SnapCenter Plug-in for Microsoft Exchange Server
 - SnapCenter Plug-in for SAP HANA Database
- SnapCenter Plug-ins Package for Linux, which includes the following plug-ins:
 - SnapCenter Plug-in for Oracle Database
 - SnapCenter Plug-in for UNIX
 - **Note:** SnapCenter Plug-in for UNIX is not a standalone plug-in and cannot be installed independently. This plug-in should always be installed with SnapCenter Plug-in for Oracle Database.
 - SnapCenter Plug-in for SAP HANA Database
- SnapCenter Custom Plug-ins

Custom plug-ins are community-supported and can be downloaded from the [NetApp ToolChest](#) or [NetApp Storage Automation Store](#).
- NetApp Data Broker
 - SnapCenter Plug-in for VMware vSphere

SnapCenter Server

The SnapCenter Server includes a web server, a centralized HTML5-based user interface, PowerShell cmdlets, REST APIs, and the SnapCenter repository.

SnapCenter enables load balancing, high availability, and horizontal scaling across multiple SnapCenter Servers within a single user interface. You can accomplish high availability by using Network Load Balancing (NLB) and Application Request Routing (ARR) with SnapCenter. For larger environments with thousands of hosts, adding multiple SnapCenter Servers can help balance the load.

- If you are using the SnapCenter Plug-ins Package for Windows, the host agent runs on the SnapCenter Server and Windows plug-in host. The host agent executes the schedules natively on the remote Windows host, or for Microsoft SQL Servers, the schedule is executed on the local SQL instance.

The SnapCenter Server communicates with the Windows plug-ins through the host agent.

- If you are using the SnapCenter Plug-ins Package for Linux, schedules are executed on the SnapCenter Server as Windows task schedules.
 - For SnapCenter Plug-in for Oracle Database, the host agent that runs on the SnapCenter Server host communicates with the SnapCenter Plug-in Loader (SPL) that runs on the Linux host to perform different data protection operations.
 - For SnapCenter Plug-in for SAP HANA Database and SnapCenter Custom Plug-ins, the SnapCenter Server communicates with these plug-ins through the SCCore agent that runs on the host.

The SnapCenter Server and plug-ins communicate with the host agent using HTTPS.

Information about SnapCenter operations is stored in the SnapCenter repository.

SnapCenter plug-ins

Each SnapCenter plug-in supports specific environments, databases, and applications.

This Plug-in...	Is included in this install package...	Requires these other plug-ins...	Is installed on this location...	And supports Windows or Linux...
Plug-in for SQL Server	Plug-ins Package for Windows	Plug-in for Windows	SQL Server host	Windows
Plug-in for Windows	Plug-ins Package for Windows		Windows host	Windows
Plug-in for Exchange	Plug-ins Package for Windows	Plug-in for Windows	Exchange Server host	Windows
Plug-in for Oracle Database	Plug-ins Package for Linux	Plug-in for UNIX	Oracle host	Linux
Plug-in for SAP HANA Database	Plug-ins Package for Linux and Plug-ins Package for Windows	Plug-in for UNIX or Plug-in for Windows	HDBSQL client host	Linux or Windows

This Plug-in...	Is included in this install package...	Requires these other plug-ins...	Is installed on this location...	And supports Windows or Linux...
Custom Plug-ins	<i>NetApp ToolChest</i>	For file system backups, Plug-in for Windows	Custom application host	Linux or Windows
Plug-in for VMware vSphere	NetApp Data Broker .ova file		Linux VM host	Linux or Windows

The SnapCenter 4.1.1 documentation has information on protecting virtualized databases and file systems using the SnapCenter 4.1.1 Plug-in for VMware vSphere. The NetApp Data Broker, SnapCenter Plug-in for VMware vSphere, documentation has information on protecting virtualized databases and file systems using the Linux-based NetApp Data Broker virtual appliance (Open Virtual Appliance format) for SnapCenter 4.2.

[Data Protection Guide for VMs, Datastores, and VMDKs using the SnapCenter Plug-in for VMware vSphere](#)

SnapCenter Plug-in for Microsoft SQL Server features

- Automates application-aware backup, restore, and clone operations for Microsoft SQL Server databases in your SnapCenter environment.
- Supports Microsoft SQL Server databases on VMDK and RDM LUNs when you also deploy the NetApp Data Broker virtual appliance and enable SnapCenter Plug-in for VMware vSphere.
- Supports provisioning SMB shares only. Support is not provided for backing up SQL Server databases on SMB shares.
- Supports importing backups from SnapManager for Microsoft SQL Server to SnapCenter.

SnapCenter Plug-in for Microsoft Windows features

- Enables application-aware data protection for other plug-ins that are running in Windows hosts in your SnapCenter environment.
- Automates application-aware backup, restore, and clone operations for Microsoft file systems in your SnapCenter environment.
- Supports storage provisioning, Snapshot copy consistency, and space reclamation for Windows hosts.

Note: The Plug-in for Windows provisions SMB shares and Windows file systems on physical and RDM LUNs but does not support backup operations for Windows file systems on SMB shares.

SnapCenter Plug-in for Microsoft Exchange Server features

- Automates application-aware backup and restore operations for Microsoft Exchange Server databases and Database Availability Groups (DAGs) in your SnapCenter environment.
- Supports virtualized Exchange Servers on RDM LUNs when you also deploy the NetApp Data Broker virtual appliance and enable SnapCenter Plug-in for VMware vSphere.
- Supports migrating backups from SnapManager 7.x to SnapCenter.

SnapCenter Plug-in for Oracle Database features

- Automates application-aware backup, restore, recovery, verify, mount, unmount, and clone operations for Oracle databases in your SnapCenter environment.
- Supports Oracle databases for SAP, however, SAP BR*Tools integration is not provided.
- Supports importing backups from SnapManager for Oracle and SnapManager for SAP to SnapCenter.

SnapCenter Plug-in for UNIX features

- Enables the Plug-in for Oracle Database to perform data protection operations on Oracle databases by handling the underlying host storage stack on Linux systems.
- Supports Network File System (NFS) and storage area network (SAN) protocols on a storage system that is running ONTAP.
- Supports Oracle databases on VMDK and RDM LUNs when you also deploy the NetApp Data Broker virtual appliance and enable SnapCenter Plug-in for VMware vSphere.

SnapCenter Plug-in for SAP HANA Database features

- Automates application-aware backup, restore, and cloning of SAP HANA databases in your SnapCenter environment.

SnapCenter Custom Plug-ins features

- Supports custom plug-ins to manage applications or databases that are not supported by other SnapCenter plug-ins. Custom plug-ins are not provided as part of the SnapCenter installation.
- Supports creating mirror copies of backup sets on another volume and performing disk-to-disk backup replication.
- Supports both Windows and Linux environments. In Windows environments, custom applications via custom plug-ins can optionally utilize SnapCenter Plug-in for Microsoft Windows to take file system consistent backups.

MySQL, DB2, and MongoDB custom plug-in samples for SnapCenter Software can be downloaded from the [NetApp ToolChest](#). You can create your own custom plug-ins by referring to the developer's guide for creating custom plug-ins.

Note: MySQL, DB2, and MongoDB custom plug-ins are supported via the NetApp communities only.

NetApp supports the capability to create and use custom plug-ins; however, the custom plug-ins you create are not supported by NetApp.

SnapCenter repository

The SnapCenter repository, sometimes referred to as the *NSM database*, stores information and metadata for every SnapCenter operation.

The MySQL Server repository database is installed by default when you install the SnapCenter Server. If MySQL Server is already installed and you are doing a fresh installation of SnapCenter Server, you should uninstall the MySQL Server.

SnapCenter supports MySQL Server 5.7.25 or later as the SnapCenter repository database. If you were using an earlier version of MySQL Server with an earlier version of SnapCenter, during SnapCenter upgrade, the MySQL Server is upgraded to 5.7.25.

The SnapCenter repository stores the following information and metadata:

- Backup, clone, restore, and verification metadata
- Reporting, job, and event information
- Host and plug-in information
- Role, user, and permission details
- Storage system connection information

SnapCenter can back up its own repository by using the SnapCenter repository management features. The *Administration Guide* or the *Windows Cmdlet Reference Guide* contain details.

The NetApp Interoperability Matrix Tool (IMT) contains the latest information about the MySQL supported versions.

[*NetApp Interoperability Matrix Tool*](#)

SnapCenter basics

As you start using SnapCenter, it is important that you understand some of the basic terms that you can expect to see throughout the SnapCenter user interface. This is especially useful for previous SnapManager users, because SnapCenter introduces some new concepts and terminology, and some changes to workflows.

Remote plug-in installation

Remote plug-in installation enables you to install the plug-ins on a host or cluster from the SnapCenter user interface. After initiating the installation, the remote install process happens seamlessly, without any input from you.

[Installing and setting up SnapCenter](#)

Role-based access control

Role-based access control (RBAC) enables the SnapCenter administrator to delegate access to the different SnapCenter and plug-in operations. Without RBAC setup, users will not be able to access SnapCenter or perform the necessary operations.

Setting up RBAC follows this basic workflow:

1. The SnapCenter administrator chooses an existing role or creates a new one.
2. The SnapCenter administrator adds users or groups of users to that role.
3. The SnapCenter administrator adds resources for a user to manage, or the user must create new resources.

[Installing and setting up SnapCenter](#)

SnapCenter repository and log backup folder

The SnapCenter repository contains information and metadata about every aspect of SnapCenter Software. You configure the SnapCenter repository during the SnapCenter Server installation. You can find the following types of information in the SnapCenter repository:

- Backup, restore, clone, and verification-related metadata
- Report, job, and event information
- Host and plug-in information
- Role, user, and permission details
- Storage system connection information

If you are using SnapCenter Plug-in for Microsoft SQL Server, the log backup folder also contains temporary `.trb` files created during log backup jobs. You must configure the log backup folder before performing your first backup job.

Note: In SnapManager for SQL, both backup metadata and `.trb` files are stored in the `SnapInfo` directory.

Security

SnapCenter employs strict security and authentication features to enable you to keep your data secure.

SnapCenter includes the following security features:

- All communication to SnapCenter uses HTTP over SSL (HTTPS).
- All credentials in SnapCenter are protected using Advanced Encryption Standard (AES) encryption.
- SnapCenter uses security algorithms that are compliant with the Federal Information Processing Standard (FIPS).
- SnapCenter 4.1.1 supports Transport Layer Security (TLS) 1.2 communication with ONTAP. You can also use TLS 1.2 communication between clients and servers.
- SnapCenter is installed inside your company's firewall to enable access to the SnapCenter Server and to enable communication between the SnapCenter Server and the plug-ins.
- SnapCenter API and operation access uses tokens, which expire after 24 hours. Tokens are also encrypted with AES encryption.
- SnapCenter integrates with Windows Active Directory for login and role-based access control (RBAC) that govern access permissions.
- SnapCenter PowerShell cmdlets are session secured.
- After a default period of 15 minutes of inactivity, SnapCenter warns you that you will be logged out in 5 minutes. After 20 minutes of inactivity, SnapCenter logs you out, and you must log in again. You can modify the log out period.
- Login is temporarily disabled after 5 or more incorrect login attempts.

Related information

[NetApp Knowledgebase Answer 1087701: How to configure the supported SSL Cipher Suite](#)

Configuration Checker

The Configuration Checker enables you to validate the configuration of SnapCenter Server and the plug-in hosts. The Configuration Checker identifies the issues in your environment and provides recommendations, corrective actions, and notifications to resolve the issues.

After installing SnapCenter Server when you log in for the first time, a default configuration checker schedule is created with the following characteristics:

- You cannot create additional schedules for SnapCenter Server.
- You cannot delete the default schedule.
- Only a SnapCenter administrator can modify or disable the default schedule.

After you add a plug-in host to the SnapCenter Server, Configuration Checker is triggered and alerts are generated. You can create a Configuration Checker schedule for the plug-in host, which you can modify, delete, or disable.

The *Installation and Setup Guide* contains more information.

[Installing and setting up SnapCenter](#)

Resources, resource groups, and policies

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, Windows file systems, or file shares that you back up or clone with SnapCenter.

However, depending on your environment, resources might be database instances, Microsoft SQL Server availability groups, Oracle databases, Oracle RAC databases, Windows file systems, or a group of custom applications.
- A SnapCenter *resource group*, is typically a collection of resources on a host or cluster. However, a resource group can contain a single resource.

When you perform an operation on a resource group, you perform that operation on all 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 configure scheduled backups for single resources and resource groups.

You should use a database plug-in to back up databases, a file system plug-in to back up file systems, and the Plug-in for VMware vSphere, which is part of the NetApp Data Broker virtual appliance, to backup VMs and datastores.
- 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.

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:



The following image illustrates the relationship between resources, resource groups, and policies for Windows file systems:



SnapCenter role-based access control (RBAC)

SnapCenter RBAC enables you to delegate control of SnapCenter resources to different users or groups of users. You can create and modify roles, and add resource access to users at any time, but when you are setting up SnapCenter for the first time, you should at least add Active Directory users or group to roles, and then add resource access to those users or groups.

Note that you cannot use SnapCenter to create user or group accounts. You must create those in Active Directory in the operating system or database.

Types of role-based access control in SnapCenter

SnapCenter role-based access control (RBAC) and ONTAP permissions enable SnapCenter administrators to create roles and set access permissions. This centrally managed access empowers application administrators to work securely within delegated environments.

SnapCenter uses the following types of role-based access control:

- SnapCenter RBAC
- SnapCenter plug-in RBAC (for some plug-ins)
- Application-level RBAC
- ONTAP permissions

SnapCenter RBAC

Roles and permissions

SnapCenter ships with predefined roles with permissions already assigned. You can assign users or groups of users to these roles. You can also create new roles and manage permissions and users.

- Assigning permissions to users or groups

You can assign permissions to users or groups to access SnapCenter objects such as hosts, storage connections, and resource groups. You cannot change the permissions of the SnapCenterAdmin role.

You can assign RBAC permissions to users and groups within the same forest and to users belonging to different forests. You cannot assign RBAC permissions to users belonging to nested groups across forests.

Note: If you create a custom role, it must contain all of the permissions of the SnapCenter Admin role. If you only copy some of the permissions, for example, Host add or Host remove, you cannot perform those operations.

Authentication

Users are required to provide authentication during login, through the graphical user interface (GUI) or using PowerShell cmdlets. If users are members of more than one role, after entering login credentials, they are prompted to specify the role they want to use. Users are also required to provide authentication to run the APIs.

Application-level RBAC

SnapCenter uses credentials to verify that authorized SnapCenter users also have application-level permissions.

For example, if you want to perform Snapshot copy and data protection operations in a SQL Server environment, you must set credentials with the proper Windows or SQL credentials. The SnapCenter Server authenticates the credentials set using either method. If you want to perform Snapshot copy and data protection operations in a Windows file system environment on ONTAP storage, the SnapCenter admin role must have admin privileges on the Windows host.

Similarly, if you want to perform data protection operations on an Oracle database and if the operating system (OS) authentication is disabled in the database host, you must set credentials with the Oracle database or Oracle ASM credentials. The SnapCenter Server authenticates the credentials set using one of these methods depending on the operation.

SnapCenter Plug-in for VMware vSphere RBAC

If you are using the Plug-in for VMware vSphere for VM-consistent data protection, the vCenter Server provides an additional level of RBAC. SnapCenter Plug-in for VMware vSphere supports both vCenter Server RBAC and Data ONTAP RBAC.

Note: The Plug-in for VMware vSphere is provided by the NetApp Data Broker virtual appliance. The NetApp Data Broker, SnapCenter Plug-in for VMware vSphere, documentation has more information.

[Data Protection Guide for VMs, Datastores, and VMDKs using the SnapCenter Plug-in for VMware vSphere](#)

ONTAP permissions

You must have vsadmin account permissions to access the storage system. A list of minimum required ONTAP privileges is included in the SnapCenter Server installation guide.

Role-based access control permissions and roles

SnapCenter role-based access control (RBAC) enables you to create roles and add permissions to those roles, and then assign users or groups of users to the roles. This enables SnapCenter administrators to create a centrally managed environment, while application administrators can manage data protection jobs. SnapCenter ships with some predefined roles and permissions.

SnapCenter roles

SnapCenter ships with the following predefined roles. You can either assign users and groups to these roles or create new roles.

When you assign a role to a user, only jobs that are relevant to that user are visible in the Jobs page unless you assigned the SnapCenter Admin role.

- App Backup and Clone Admin
- Backup and Clone Viewer
- Infrastructure Admin
- SnapCenterAdmin

SnapCenter Plug-in for VMware vSphere roles

For managing VM-consistent data protection of VMs, VMDKs, and datastores, the following roles are created in vCenter by the NetApp Data Broker virtual appliance when the SnapCenter Plug-in for VMware vSphere is enabled:

- SCV Administrator

- SCV View
- SCV Backup
- SCV Restore
- SCV Guest File Restore

The NetApp Data Broker, SnapCenter Plug-in for VMware vSphere, documentation has more information.

[Data Protection Guide for VMs, Datastores, and VMDKs using the SnapCenter Plug-in for VMware vSphere](#)

Best Practice: NetApp recommends that you create one ONTAP role for SnapCenter Plug-in for VMware vSphere operations and assign it all the required privileges.

SnapCenter permissions

SnapCenter provides the following permissions:

- Resource Group
- Policy
- Backup
- Host
- Storage Connection
- Clone
- Provision (only for Microsoft SQL database)
- Dashboard
- Reports
- Restore
 - Full Volume Restore (only for Custom Plug-ins)

Resource

Plug-in privileges are required from the administrator for non-administrators to perform resource discovery operation.

Plug-in Install\Uninstall

Note: When you enable Plug-in Installation permissions, you must also modify the Host permission to enable reads and updates.

- Migration
- Mount (only for Oracle database)
- Unmount (only for Oracle database)

Pre-defined SnapCenter roles and permissions

SnapCenter ships with pre-defined roles, each with a set of permissions already enabled. When setting up and administering role-based access control (RBAC), you can either use these pre-defined roles or create new ones.

SnapCenter includes the following pre-defined roles:

- SnapCenter Admin role
- Backup and Clone Viewer role
- App Backup and Clone Admin role
- Infrastructure Admin role

When you add a user to a role, you must assign either the StorageConnection permission to enable storage virtual machine (SVM) communication, or assign an SVM to the user to enable permission to use the SVM. The Storage Connection permission enables users to create SVM connections.

For example, a user with the SnapCenter Admin role can create SVM connections and assign them to a user with the App Backup and Clone Admin role, which by default does not have permission to create or edit SVM connections. Without an SVM connection, users cannot complete any backup, clone, or restore operations.

SnapCenter Admin role

The SnapCenter Admin role has all permissions enabled. You cannot modify the permissions for this role. You can add users and groups to the role or remove them.

App Backup and Clone Admin role

The App Backup and Clone Admin role has the permissions required to perform administrative actions for application backups and clone-related tasks. This role does not have permissions for host management, provisioning, storage connection management, or remote installation.

Permissions	Enabled	Create	Read	Update	Delete
Resource Group	Not applicable	Yes	Yes	Yes	Yes
Policy	Not applicable	Yes	Yes	Yes	Yes
Backup	Not applicable	Yes	Yes	Yes	Yes
Host	Not applicable	Yes	Yes	Yes	Yes
Storage Connection	Not applicable	No	Yes	No	No
Clone	Not applicable	Yes	Yes	Yes	Yes
Provision	Not applicable	No	Yes	No	No
Dashboard	Yes	Not applicable	Not applicable	Not applicable	Not applicable
Reports	Yes	Not applicable	Not applicable	Not applicable	Not applicable
Restore	Yes	Not applicable	Not applicable	Not applicable	Not applicable
Resource	Yes	Yes	Yes	Yes	Yes
Plug-in Install/Uninstall	No	Not applicable		Not applicable	Not applicable
Migration	No	Not applicable	Not applicable	Not applicable	Not applicable
Mount	Yes	Yes	Not applicable	Not applicable	Not applicable

Permissions	Enabled	Create	Read	Update	Delete
Unmount	Yes	Yes	Not applicable	Not applicable	Not applicable
Full Volume Restore	No	No	Not applicable	Not applicable	Not applicable

Backup and Clone Viewer role

The Backup and Clone Viewer role has read-only view of all permissions. This role also has permissions enabled for discovery, reporting, and access to the Dashboard.

Permissions	Enabled	Create	Read	Update	Delete
Resource Group	Not applicable	No	Yes	No	No
Policy	Not applicable	No	Yes	No	No
Backup	Not applicable	No	Yes	No	No
Host	Not applicable	No	Yes	No	No
Storage Connection	Not applicable	No	Yes	No	No
Clone	Not applicable	No	Yes	No	No
Provision	Not applicable	No	Yes	No	No
Dashboard	Yes	Not applicable	Not applicable	Not applicable	Not applicable
Reports	Yes	Not applicable	Not applicable	Not applicable	Not applicable
Restore	No	No	Not applicable	Not applicable	Not applicable
Resource	No	No	Yes	Yes	No
Plug-in Install/Uninstall	No	Not applicable	Not applicable	Not applicable	Not applicable
Migration	No	Not applicable	Not applicable	Not applicable	Not applicable
Mount	Yes	Not applicable	Not applicable	Not applicable	Not applicable
Unmount	Yes	Not applicable	Not applicable	Not applicable	Not applicable
Full Volume Restore	No	Not applicable	Not applicable	Not applicable	Not applicable

Infrastructure Admin role

The Infrastructure Admin role has permissions enabled for host management, storage management, provisioning, resource groups, remote installation reports, and access to the Dashboard.

Permissions	Enabled	Create	Read	Update	Delete
Resource Group	Not applicable	Yes	Yes	Yes	Yes
Policy	Not applicable	No	Yes	Yes	Yes
Backup	Not applicable	Yes	Yes	Yes	Yes
Host	Not applicable	Yes	Yes	Yes	Yes
Storage Connection	Not applicable	Yes	Yes	Yes	Yes
Clone	Not applicable	No	Yes	No	No
Provision	Not applicable	Yes	Yes	Yes	Yes
Dashboard	Yes	Not applicable	Not applicable	Not applicable	Not applicable
Reports	Yes	Not applicable	Not applicable	Not applicable	Not applicable
Restore	Yes	Not applicable	Not applicable	Not applicable	Not applicable
Resource	Yes	Yes	Yes	Yes	Yes
Plug-in Install/Uninstall	Yes	Not applicable	Not applicable	Not applicable	Not applicable
Migration	No	Not applicable	Not applicable	Not applicable	Not applicable
Mount	No	Not applicable	Not applicable	Not applicable	Not applicable
Unmount	No	Not applicable	Not applicable	Not applicable	Not applicable
Full Volume Restore	No	No	Not applicable	Not applicable	Not applicable

Prescripts and postscripts

You can use custom prescripts and postscripts as part of your data protection operations. These scripts enable automation either before your data protection job or after. For example, you might include a script that automatically notifies you of data protection job failures or warnings. Before you

set up your prescripts and postscripts, you should understand some of the requirements for creating these scripts.

Supported script types

The following type of scripts are supported:

- Batch files
- PowerShell scripts
- Perl scripts

Script path location

All prescripts and postscripts that are run as part of SnapCenter operations, on nonvirtualized and on virtualized storage systems, are executed on the plug-in host. Therefore, the scripts must be located on the plug-in host or on a SMB share accessible by the plug-in host.

Note: The script path is validated at the time the script is executed.

Where to specify scripts

Scripts are specified in backup policies. When a backup job is started, the policy automatically associates the script with the resources being backed up. When you create a backup policy, you can specify the prescript and pscript arguments.

To specify multiple scripts, press **Enter** after each script path to list each script on a separate line. Semicolons (;) are not allowed. You can specify multiple prescripts and multiple postscripts. A single script can be coded as both a prescript and a postscript and can call other scripts.

Script timeouts

The timeout for backup scripts is 15 minutes and cannot be modified.

Script output

The default directory for the prescripts and postscripts output files is `Windows\System32`.

SnapCenter REST APIs

You can use REST APIs to perform several SnapCenter management operations. REST APIs are exposed through the Swagger web page. You can access the Swagger web page to display the SnapCenter Server or SnapCenter for VMware REST API documentation, as well as to manually issue an API call. You can use REST APIs to help manage your SnapCenter Server or your SnapCenter vSphere host.

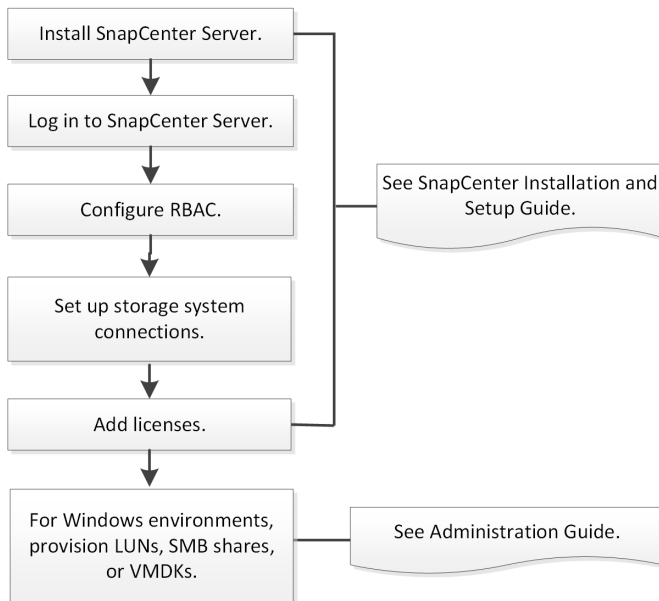
The REST APIs for...	Are located in...
SnapCenter Server	<code>https://<SnapCenter_IP_address_or_name>:8146/swagger/</code>
SnapCenter Plug-in for VMware vSphere	<code>https://<OVA_IP_address_or_host_name>:8144/api/swagger-ui.html#</code>

Performing administrative tasks

The NetApp Data Broker, SnapCenter Plug-in for VMware vSphere, documentation has more information.

SnapCenter backup workflow overview

After you initially install SnapCenter Server, you must perform several tasks in a particular sequence to use the product for the first time. After the initial setup, you can perform most of these tasks at any time, but it is important that you follow the correct sequence.



SnapCenter Plug-in for Microsoft Exchange Server overview

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

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

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

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

Related information

[Protecting Microsoft Exchange Server databases](#)

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

You can use the Plug-in for Exchange to back up and restore Exchange Server databases.

- View and manage an active inventory of Exchange Database Availability Groups (DAGs), databases, and replica sets
- Define policies that provide the protection settings for backup automation
- Assign policies to resource groups
- Protect individual DAGs and databases
- Back up primary and secondary Exchange mailbox databases
- Restore databases from primary and secondary backups

Defining a backup strategy for Exchange 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 your databases. Your Service Level Agreement (SLA), Recovery Time Objective (RTO), and Recovery Point Objective (RPO) largely determine your backup strategy.

About this task

An SLA defines the level of service 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.

Perform the following steps for each of your databases.

Steps

1. Decide the type of backup you require.
2. Determine when you should back up your databases.
3. Decide how many backup jobs you require.
4. Decide how to name your backups.
5. Determine how long you want to retain backup copies on the source storage system.
6. Determine how long you want to retain transaction log backups on the source storage system.
7. Decide whether you want to use SnapMirror disaster recovery or SnapVault archiving technology in conjunction with the SnapCenter Plug-in for Microsoft Exchange Server.

Types of backups supported

Backing up Exchange mailboxes using SnapCenter requires that you choose the resource type, such as databases and Database Availability Groups (DAG). Snapshot copy technology is leveraged to create online, read-only copies of the volumes on which the resources reside.

Backup type	Description
Full and log backup	<p>Backs up the databases and all transaction logs, including the truncated logs.</p> <p>After a full backup is complete, the Exchange Server truncates the transaction logs that are already committed to the database.</p> <p>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>
Full backup	<p>Backs up databases and transaction logs.</p> <p>The truncated transaction logs are not backed up.</p>
Log backup	<p>Backs up all the transaction logs. The truncated logs that are already committed to the database are not backed up.</p> <p>If you schedule frequent transaction log backups between full database backups, you can choose granular recovery points.</p>

Backup schedules for database plug-ins

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

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

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 the 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.

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.

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

SnapCenter Plug-in for Microsoft Exchange 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 Exchange took a full plus transaction log backup at 8:00 a.m. and another full plus transaction log 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 Exchange can perform *point-in-time* restore operations only, which restore a database to the time that Plug-in for Exchange 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.

Defining a restore strategy for Exchange databases

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

Steps

1. Determine whether you want to restore a backup of an active or passive copy.
2. Determine the requirements for restoring the database.
3. Decide the type of restore you require.

Sources for a restore operation

You can restore an Exchange Server database from a backup copy on primary storage.

Sources for a restore operation

You can restore databases from primary storage only.

Types of restore operations

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

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

Restore up to the minute

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

1. Restores the databases from the full database backup that you select.
2. Applies all the transaction logs that were backed up, as well as any new logs that were created since the most recent backup.
Transaction logs are moved ahead and applied to any selected databases.

Exchange creates a new log chain after a restore completes.

Best Practice: It is recommended that you perform a new full and log backup after a restore completes.

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

After you perform an up-to-the-minute restore, the backup you used for the restore is available only for point-in-time restore operations.

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.

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.

SnapCenter Plug-in for Microsoft SQL Server overview

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

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

SQL DBA

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

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

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

Related information

[Protecting Microsoft SQL Server databases](#)

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, including databases, instances, or Availability Groups
You can restore master system databases, msdb system databases, and model system databases. 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
- Schedule backup operations and clone operations
- Monitor backup operations, restore operations, and clone operations

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

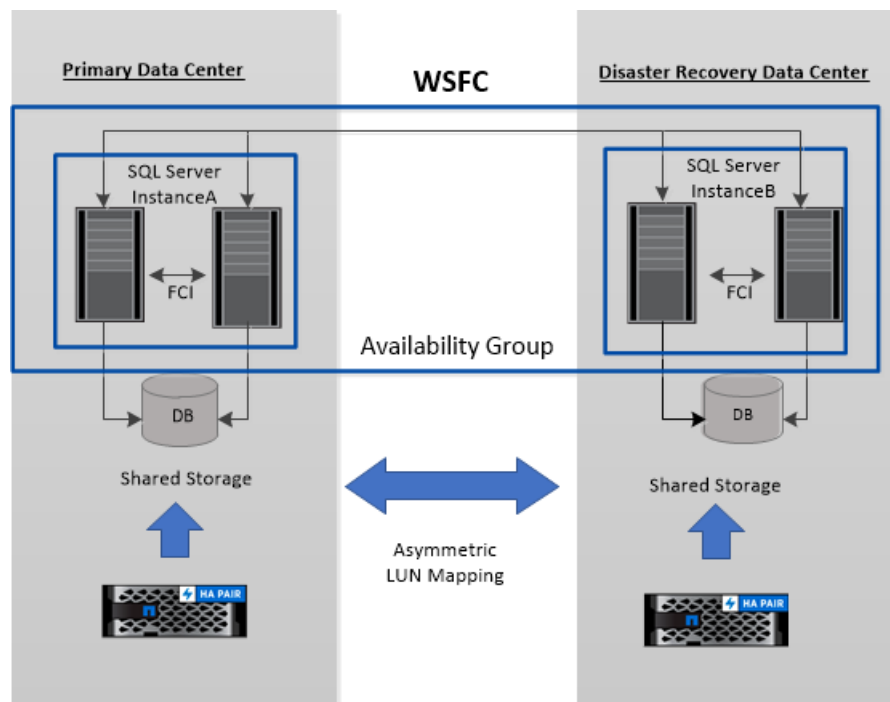
Support for Asymmetric LUN Mapping in Windows clusters

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.

Defining 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.

About this task

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.

Perform the following steps for each of your databases.

Steps

1. Decide the type of backup you require.
2. Determine when you should back up your databases.
3. Decide how many backup jobs you require.
4. Decide how to name your backups.
5. Determine when you should verify backup copies.
6. Decide whether you want to verify backup copies using the source volume or a destination volume.
7. Determine for how long you want to retain backup copies on the source storage system and the SnapMirror destination.
8. Determine for how long you want to retain transaction log backups on the source storage system.
9. Decide whether you want to use SnapMirror disaster recovery or SnapVault archiving technology in conjunction with the SnapCenter Plug-in for Microsoft SQL Server.

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>	Backs up the system database files.
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 database plug-ins

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

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

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.

Note: 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

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.

Defining a restoration strategy for SQL Server

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

Steps

1. Determine the source and destination for the restore operation.
2. Determine the requirements for restoring the database.
3. Determine the SQL Server recovery models.
4. Decide the type of restore you require.

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.
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.

Note: 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)

Note: 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).

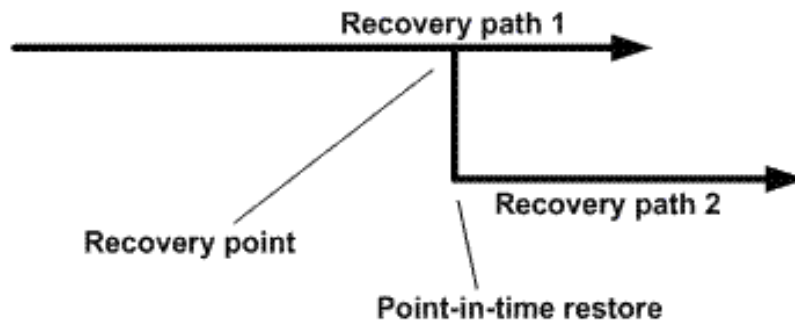
Restore to a previous point in time

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.

Note: 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.

Defining a cloning strategy for SQL Server

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

Steps

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 195447733 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.

SnapCenter Plug-in for Microsoft Windows overview

The SnapCenter Plug-in for Microsoft Windows is a host-side component of the NetApp SnapCenter Software that enables application-aware data protection management of Microsoft file system resources. In addition, it provides storage provisioning, Snapshot copy consistency, and space reclamation for Windows file systems. The Plug-in for Windows automates file system backup, restore, and cloning operations in your SnapCenter environment.

When the Plug-in for Windows 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 archival or standards compliance.

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

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

Related information

[Protecting Microsoft Windows file systems](#)

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

When the Plug-in for Windows is installed in your environment, you can use SnapCenter to back up, restore, and clone Windows file systems. You can also perform tasks supporting those operations.

- Discover resources
- Back up Windows file systems
- Schedule backup operations
- Restore file system backups
- Clone file system backups
- Monitor backup, restore, and clone operations

Note: The Plug-in for Windows does not support backup and restore of file systems on SMB shares.

SnapCenter Plug-in for Windows features

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

The Plug-in for Windows 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 file system backups, configure policy-based backup retention, and set up restore operations. You can also proactively monitor your file system environment by configuring SnapCenter to send email alerts.

Nondisruptive NetApp Snapshot copy technology

The Plug-in for Windows uses NetApp Snapshot copy technology. This enables you to back up file systems in seconds and restore them quickly without taking host offline. Snapshot copies consume minimal storage space.

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

- Backup, restore, and clone workflow support
- RBAC-supported security and centralized role delegation
- Creation of space-efficient copies of production file systems for testing or data extraction by using NetApp FlexClone technology
For FlexClone licensing information, see licensing requirements in the installation information.
[*Installing and setting up SnapCenter*](#)
- Ability to run multiple backups at the same time across multiple servers
- PowerShell cmdlets for scripting of backup, restore, and clone operations
- Support for backup of file systems and virtual machine disks (VMDKs)
- Support for physical and virtualized infrastructures
- Support for iSCSI, Fibre Channel, FCoE, raw device mapping (RDM), Asymmetric LUN Mapping (ALM), VMDK over NFS and VMFS, and virtual FC

How SnapCenter backs up Windows file systems

SnapCenter uses Snapshot copy technology to back up Windows file system resources that reside on LUNs, CSVs (cluster shared volumes), RDM (raw device mapping) volumes, ALM (asymmetric LUN mapping) in Windows clusters, and VMDKs based on VMFS/NFS (VMware Virtual Machine File System using NFS).

SnapCenter creates backups by creating Snapshot copies of the file systems. Federated backups, in which a volume contains LUNs from multiple hosts, are faster and more efficient than backups of each individual LUN because only one Snapshot copy of the volume is created compared to individual Snapshots of each file system.

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 host server for which the backup was created.

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

Note: If a Windows file system contains a database, then backing up the file system is not the same as backing up the database. To back up a database, you must use one of the database plug-ins.

Defining a backup strategy for Windows file systems

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

About this task

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

Perform the following steps for each of your Windows file systems.

Steps

1. Determine when you should back up your file systems.
2. Decide how many backups you require.
3. Decide how to tag your backups.
4. Determine for how long you want to retain backup copies on the source storage system and the SnapMirror destination.
5. Decide whether you want to use NetApp SnapMirror disaster recovery or SnapVault archiving technology in conjunction with the SnapCenter Plug-in for Microsoft Windows.
6. Determine the supported prescripts and postscripts.

Backup schedules for Windows file systems

Backup frequency 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.

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. For example, you might configure the backup frequency as hourly, daily, weekly, or monthly, or you can specify **None** which makes the policy 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 backups needed for Windows file systems

Factors that determine the number of backups that you need include the size of the Windows file system, the number of volumes used, the rate of change of the file system, and your Service Level Agreement (SLA).

Backup naming convention

Windows file system backups use the default Snapshot copy 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.

When creating a backup, you can also add a descriptive tag to help identify the backup. In contrast, if you want to use a customized backup naming convention, you need to rename the backup after the backup operation is complete.

Backup retention options

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.

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

Sources and destinations of clones for Windows file systems

You can clone a file system backup from primary storage or secondary storage. You also can choose the destination that supports your requirements; either the original backup location or a different

destination on the same host or on a different host. The destination must be on the same volume as the clone source backup.

Clone destination	Description
Original, source, location	By default, SnapCenter stores the clone on the same location and the same host as the backup being cloned.
Different location	You can store the clone on a different location on the same host or on a different host. The host must have a configured connection to the storage virtual machine (SVM).

You can rename the clone after the clone operation is complete.

SnapCenter Plug-in for Oracle Database overview

The SnapCenter Plug-in for Oracle Database is a host-side component of the NetApp SnapCenter Software that enables application-aware data protection management of Oracle databases.

The Plug-in for Oracle Database automates the backup, cataloging and uncataloging with Oracle RMAN, verification, mounting, unmounting, restore, recovery, and cloning of Oracle databases in your SnapCenter environment.

The Plug-in for Oracle Database installs SnapCenter Plug-in for UNIX to perform all the data protection operations.

You can use the Plug-in for Oracle Database to manage backups of Oracle databases running SAP applications. However, SAP BR*Tools integration is not supported.

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

Related information

[Protecting Oracle databases](#)

What you can do with the SnapCenter Plug-in for Oracle Database

You can use the Plug-in for Oracle Database to back up, verify, restore, recover, mount, unmount, and clone Oracle databases and their resources. You can also catalog or uncatalog the database backups with Oracle Recovery Manager (RMAN).

- Back up datafiles, control files, and archive log files.
Backup is supported only at container database (CDB) level.
- Restore and recovery of databases, CDBs, and pluggable databases (PDBs).
Incomplete recovery of PDBs are not supported.
- Create clones of production databases up to a point-in-time.
Cloning is supported only at CDB level.
- Verify backups immediately.
- Mount and unmount data and log backups for recovery operation.
- Schedule backup and verification operations.
- Monitor all operations.
- View reports for backup, restore, and clone operations.

SnapCenter Plug-in for Oracle Database features

The Plug-in for Oracle Database integrates with the Oracle database on the Linux host and with NetApp technologies on the storage system.

Unified graphical user interface

The SnapCenter interface provides standardization and consistency across plug-ins and environments. The SnapCenter interface enables you to complete consistent backup, restore, recovery, and clone operations across plug-ins, use centralized reporting, use at-a-

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

Automated central administration

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

Nondisruptive NetApp Snapshot copy technology

SnapCenter uses NetApp Snapshot copy technology with the Plug-in for Oracle Database and Plug-in for UNIX to back up databases. Snapshot copies consume minimal storage space.

The Plug-in for Oracle Database also offers the following benefits:

- Support for backup, restore, clone, mount, unmount, and verification workflows
- Automatically discover Oracle databases configured on the host
- Support for cataloging and uncataloging using Oracle Recovery Manager (RMAN)
- RBAC-supported security and centralized role delegation
You can also set the credentials so that the authorized SnapCenter users have application-level permissions.
- Creation of space-efficient and point-in-time copies of production databases for testing or data extraction by using NetApp FlexClone technology
A FlexClone license is required on the storage system where you want create the clone.
- Support for consistency group (CG) feature of ONTAP as part of creating backups in SAN and ASM environments
- Nondisruptive and automated backup verification
- Capability to run multiple backups simultaneously across multiple database hosts
In a single operation, Snapshot copies are consolidated when databases in a single host share the same volume.
- Support for physical and virtualized infrastructures
- Support for NFS, iSCSI, Fibre Channel (FC), RDM, VMDK over NFS and VMFS, and ASM over NFS and SAN
- Support for the Selective LUN Map (SLM) feature of ONTAP
Enabled by default, the SLM feature periodically discovers the LUNs that do not have optimized paths and fixes them. You can configure SLM by modifying the parameters in the `scu.properties` file located at `/var/opt/snapcenter/scu/etc`.
 - You can disable this by setting the value of `ENABLE_LUNPATH_MONITORING` to **false**.
 - You can specify the frequency in which the LUN paths will be fixed automatically by assigning the value (in hours) to `LUNPATH_MONITORING_INTERVAL`.

For information about SLM, see the ONTAP administration information.

[ONTAP 9 SAN Administration Guide](#)

Defining a backup strategy for Oracle databases

Defining a backup strategy before you create your backup jobs ensures 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.

About this task

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

Perform the following steps for each of your databases.

Steps

1. Determine the supported database configurations.
2. Decide the type of backup that you require.
3. Decide on which node in a RAC environment you want to create your backup.
4. Determine when you should back up your databases.
5. Decide how to name your backups.
6. Decide if you want to catalog or uncatalog your backups.
7. Determine the retention period for the Snapshot copies on the source storage system and the SnapMirror destination.

SnapMirror retention is implicitly defined based on the Snapshot copy retention of the primary storage. SnapVault retention has to be defined explicitly on the ONTAP storage system.
8. Decide whether you want to verify the backup copies using the primary or secondary storage volume.
9. Determine the supported prescripts and postscripts.
10. Decide whether you want to use NetApp SnapMirror technology for replication or NetApp SnapVault technology for long term retention.

Supported Oracle database configurations for backups

SnapCenter supports backup of different Oracle database configurations.

- Oracle Standalone
- Oracle Real Application Clusters (RAC)
- Oracle Standalone Legacy
- Oracle Standalone Container Database (CDB)
- Oracle Data Guard standby
You can create only offline (mount or shutdown) backups of Data Guard standby databases.
- Oracle Active Data Guard standby
You can create data only online backups of Active Data Guard standby databases.
- Oracle database on Automatic Storage Management (ASM)
You can perform backup, restore, and clone operations on Oracle databases on ASM, with or without ASMLib. However, ASM on Raw device mapping (RDM) and virtual machine disk (VMDK) is not supported.

Note: Oracle ASM Filter Driver (ASMFID), PDB migration, and PDB cloning are not supported.

For the latest information about supported versions, see the NetApp Interoperability Matrix Tool (IMT).

[NetApp Interoperability Matrix Tool](#)

Note: Before creating a backup of Data Guard standby or Active Data Guard standby database, the managed recovery process (MRP) is stopped and once the backup is created, MRP is started.

Types of backup supported

Backup type specifies the type of backup that you want to create. SnapCenter supports online and offline backup types for Oracle databases.

Online backup

A backup that is created when the database is in the online state is called an *online backup*. Also called a *hot backup*, an online backup enables you to create a backup of the database without shutting it down.

As part of online backup, you can create a backup of the following files:

- Datafiles and control files only
- Archive log files only (the database is not brought to backup mode in this scenario)
- Full database that includes datafiles, control files, and archive log files

Offline backup

A backup created when the database is either in a mounted or shutdown state is called an *offline backup*. An offline backup is also called a *cold backup*. You can include only datafiles and control files in offline backups. You can create either an offline mount or offline shutdown backup.

- When creating an offline mount backup, you must ensure that the database is in a mounted state. If the database is in any other state, the backup operation fails.
- When creating an offline shutdown backup, the database can be in any state. The database state is changed to the required state to create a backup. After creating the backup, the database state is reverted to the original state.

Understanding the discovery of Oracle databases

“Resources” are Oracle databases on the host that are maintained by SnapCenter. You can add these databases to resource groups to perform data protection operations after you discover the databases that are available. You should be aware of the process that SnapCenter follows to discover different types and versions of Oracle databases.

For Oracle versions 11g to 12cR1	For Oracle versions 12cR2 to 18c
<p>RAC database: The RAC databases are discovered only on the basis of <code>/etc/oratab</code> entries.</p> <p>You should have the database entries in the <code>/etc/oratab</code> file.</p>	<p>RAC database: The RAC databases are discovered using the <code>srvctl config</code> command.</p>

For Oracle versions 11g to 12cR1	For Oracle versions 12cR2 to 18c
<p>Standalone: The standalone databases are discovered only on the basis of <code>/etc/oratab</code> entries.</p> <p>You should have the database entries in the <code>/etc/oratab</code> file.</p>	<p>Standalone: The standalone databases are discovered based on the entries in the <code>/etc/oratab</code> file and the output of the <code>srvctl config</code> command.</p>
<p>ASM: The ASM instance entry should be available in the <code>/etc/oratab</code> file.</p>	<p>ASM: The ASM instance entry need not be in the <code>/etc/oratab</code> file.</p>
<p>RAC One Node: The RAC One Node databases are discovered only on the basis of <code>/etc/oratab</code> entries. The databases should be either in <code>nomount</code>, <code>mount</code>, or <code>open</code> state.</p> <p>You should have the database entries in the <code>/etc/oratab</code> file.</p> <p>The RAC One Node database status will be marked as <code>renamed</code> or <code>deleted</code> if the database is already discovered and backups are associated with the database.</p> <p>You should perform the following steps if the database is relocated:</p> <ol style="list-style-type: none"> 1. Manually add the relocated database entry in the <code>/etc/oratab</code> file on the failed-over RAC node. 2. Manually refresh the resources. 3. Select the RAC One Node database from the resource page, and then click Database Settings. 4. Configure the database to set the preferred cluster nodes to the RAC node currently hosting the database. 5. Perform the SnapCenter operations. 	<p>RAC One Node: The RAC One Node databases are discovered using the <code>srvctl config</code> command only. The databases should be either in <code>nomount</code>, <code>mount</code>, or <code>open</code> state.</p> <p>The RAC One Node database status will be marked as <code>renamed</code> or <code>deleted</code> if the database is already discovered and backups are associated with the database.</p> <p>You should perform the following steps if the database is relocated:</p> <ol style="list-style-type: none"> 1. Manually refresh the resources. 2. Select the RAC One Node database from the resource page, and then click Database Settings. 3. Configure the database to set the preferred cluster nodes to the RAC node currently hosting the database. 4. Perform the SnapCenter operations.

Note:

- If there are any Oracle 12cR2 and 18c database entries in the `/etc/oratab` file and the same database is registered with the `srvctl config` command, SnapCenter will eliminate the duplicate database entries.
- If there are stale database entries, the database will be discovered but the database will be unreachable and the status will be offline.

Preferred nodes in RAC setup

In Oracle Real Application Clusters (RAC) setup, you can specify the preferred nodes on which the backup operation will be performed. If you do not specify the preferred node, SnapCenter automatically assigns a node as the preferred node and backup is created on that node.

The preferred nodes might be one or all of the cluster nodes where the RAC database instances are present. The backup operation will be triggered only on these preferred nodes in the order of the preference.

For example, the RAC database `cdbrac` has three instances: `cdbrac1` on `node1`, `cdbrac2` on `node2`, and `cdbrac3` on `node3`. The `node1` and `node2` instances are configured to be the preferred nodes, with `node2` as the first preference and `node1` as the second preference. When you perform a backup operation, the operation is first attempted on `node2` because it is the first preferred node. If `node2` is not in the state to back up, which could be due to multiple reasons such as the plug-in agent is not running on the host, the database instance on the host is not in the required state for the specified backup type, or the database instance on `node2` in a FlexASM configuration is not being served by the local ASM instance; then the operation will be attempted on `node1`. The `node3` will not be used for backup because it is not on the list of preferred nodes.

Details about specifying preferred nodes are in the Administration Guide.

Performing administrative tasks

Required database state

The RAC database instances on the preferred nodes must be in the required state for the backup to finish successfully:

- One of the RAC database instances in the configured preferred nodes must be in the open state to create an online backup.
- One of the RAC database instances in the configured preferred nodes must be in the mount state, and all other instances, including other preferred nodes, must be in the mount state or lower to create an offline mount backup.
- RAC database instances can be in any state, but you must specify the preferred nodes to create an offline shutdown backup.

Backup cataloging with Oracle Recovery Manager

The backups of Oracle databases can be cataloged with Oracle Recovery Manager (RMAN) to store the backup information in the Oracle RMAN repository. These cataloged backups can be used later for block-level restore or tablespace point-in-time recovery operations. When you do not need these cataloged backups, you can remove the catalog information.

The database must be in mounted or higher state for cataloging. You can perform cataloging on data backups, archive log backups, and full backups. If cataloging is enabled for a backup of a resource group that has multiple databases, cataloging is performed for each database. For Oracle RAC databases, cataloging will be performed on the preferred node where the database is at least in mounted state.

Note: If you want to catalog backups of a RAC database, ensure that no other job is running for that database. If another job is running, the cataloging operation fails instead of getting queued.

By default, the target database control file is used for cataloging. If you want to add external catalog database, you can configure it by specifying the credential and Transparent Network Substrate (TNS) name of the external catalog using the Database Settings wizard from the SnapCenter GUI. You can also configure the external catalog database from the CLI by running the `Configure-SmOracleDatabase` command with the `-OracleRmanCatalogCredentialName` and `-OracleRmanCatalogTnsName` options.

If you enabled the cataloging option while creating an Oracle backup policy from the SnapCenter graphical user interface (GUI), the backups are cataloged using Oracle RMAN as a part of the backup operation. You can also perform deferred cataloging of backups by running the `Catalog-SmBackupWithOracleRMAN` command. After cataloging the backups, you can run the `Get-SmBackupDetails` command to obtain the cataloged backup information such as the tag for cataloged datafiles, the control file catalog path, and the cataloged archive log locations.

If the ASM disk group name is greater than or equal to 16 characters, from SnapCenter 3.0, the naming format used for the backup is `SC_HASHCODEoFDISKGROUP_DBSID_BACKUPID`. However, if

the disk group name is less than 16 characters, the naming format used for the backup is `DISKGROUPNAME_DBSID_BACKUPID`, which is the same format used in SnapCenter 2.0.

Note: The `HASHCODEofDISKGROUP` is an automatically generated number (2 to 10 digit) unique for each ASM disk group.

You can perform crosschecks to update outdated RMAN repository information about backups whose repository records do not match their physical status. For example, if a user removes archived logs from disk with an operating system command, the control file still indicates that the logs are on disk, when in fact they are not. The crosscheck operation enables you to update the control file with the information. You can enable crosscheck by running the `Set-SmConfigSettings` command and assigning the value `TRUE` to the `ENABLE_CROSSCHECK` parameter. The default value is set to `FALSE`.

```
sccli Set-SmConfigSettings -ConfigSettingsType Plugin -PluginCode SCO -
ConfigSettings "KEY=ENABLE_CROSSCHECK, VALUE=TRUE"
```

You can remove the catalog information by running the `Uncatalog-SmBackupWithOracleRMAN` command. You cannot remove the catalog information using the SnapCenter GUI. However, information of a cataloged backup is removed while deleting the backup or while deleting the retention and resource group associated with that cataloged backup.

Note: When you force a deletion of the SnapCenter host, the information of the cataloged backups associated with that host are not removed. You must remove information of all the cataloged backups for that host before forcing the deletion of the host.

The information regarding the parameters that can be used with the command and their descriptions can be obtained by running `help command_name`. Alternatively, you can also refer to the *Command Reference Guide*.

[SnapCenter Software 4.2 Linux Command Reference Guide](#)

Backup schedules for database plug-ins

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**.

Backup naming conventions

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

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.

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

Backup copy verification using the primary or secondary storage volume

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.

Defining a restore and recovery strategy for Oracle databases

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

Steps

1. Determine the backups that can be used for restore and recovery operations.
2. Know the limitations that are related to the restore and recovery operations.
3. Determine the restore methods that are supported.
4. Decide the type of restore and recovery operations that you want to perform.
5. Determine the source and destination for the restore operation.

Types of Oracle database backups supported for restore and recovery operations

Depending on your version of Oracle Database, SnapCenter supports restore and recovery of different types of backups.

Oracle 11g

- Standalone online data backup
- Standalone offline shutdown data backup
- Standalone offline mount data backup
- Standalone full backup
- Offline (mount or shutdown) backups of Data Guard standby databases
- Data-only online backups of Active Data Guard standby databases
 - Note:** You cannot perform recovery of Active Data Guard standby databases.
- Online data backups, online full backups, offline mount backups, and offline shutdown backups in a Real Application Clusters (RAC) configuration
- Online data backups, online full backups, offline mount backups, and offline shutdown backups in an Automatic Storage Management (ASM) configuration

Oracle 12c (legacy)

- Standalone online data backup
- Standalone offline shutdown data backup
- Standalone offline mount data backup
- Standalone full backup
- Offline (mount or shutdown) backups of Data Guard standby databases
- Data-only online backups of Active Data Guard standby databases

Note: You cannot perform recovery of Active Data Guard standby databases.

- Online data backups, online full backups, offline mount backups, and offline shutdown backups in a RAC configuration
- Online data backups, online full backups, offline mount backups, and offline shutdown backups in an ASM configuration

Oracle 12c (CDB)

- Online data backup
- Offline shutdown data backup
- Offline mount data backup
- Full backup
- Offline (mount or shutdown) backups of Data Guard standby databases
- Data-only online backups of Active Data Guard standby databases

Note: You cannot perform recovery of Active Data Guard standby databases.

- Online data backups, online full backups, offline mount backups, and offline shutdown backups in RAC configuration
- Online data backups, online full backups, offline mount backups, and offline shutdown backups in ASM configuration

Types of restore methods supported for Oracle databases

SnapCenter supports connect-and-copy or in-place restore for Oracle databases. During a restore operation, SnapCenter determines the restore method that is appropriate for the file system to be used for restore without any data loss.

Note: SnapCenter does not support volume-based SnapRestore.

Connect-and-copy restore

If the database layout differs from the backup or if there are any new files after the backup was created, connect-and-copy restore is performed. In the connect-and-copy restore method, the following tasks are performed:

1. The volume is cloned from the Snapshot copy and the file system stack is built on the host using the cloned LUNs or volumes.
2. The files are copied from the cloned file systems to the original file systems.
3. The cloned file systems are then unmounted from the host and the cloned volumes are deleted from ONTAP.

In-place restore

If the database layout is similar to the backup and has not undergone any configuration change on the storage and database stack, in-place restore is performed, wherein the restore of file or LUN is performed on ONTAP. SnapCenter supports only Single File SnapRestore (SFSR) as part of the in-place restore method.

Note: Data ONTAP 8.3 or later supports in-place restore from secondary location.

If you want to perform in-place restore on the database, ensure that you have only datafiles on the ASM disk group. You must create a backup after any changes are made to the ASM disk group or in the physical structure of the database. After performing in-place restore, the disk group will contain the same number datafiles as at the time of backup.

Performing In-place restore on ASM RAC

In SnapCenter, the node on which you perform restore is termed as primary node and all other nodes of the RAC on which ASM disk group resides are called peer nodes. SnapCenter changes the state of ASM disk group to dismount on all the nodes where the ASM disk group is in mount state before performing the storage restore operation. After the storage restore is complete, SnapCenter changes the state of ASM disk group as it was before the restore operation.

In SAN environments, SnapCenter removes devices from all the peer nodes and performs LUN unmap operation before storage restore operation. After storage restore operation, SnapCenter performs LUN map operation and constructs devices on all the peer nodes. In a SAN environment if the Oracle RAC ASM layout is residing on LUNs, then while restoring SnapCenter performs LUN unmap, LUN restore, and LUN map operations on all the nodes of the RAC cluster where the ASM disk group resides. Before restoring even if all the initiators of the RAC nodes were not used for the LUNs, after restoring SnapCenter creates a new iGroup with all the initiators of all the RAC nodes.

- If there is any failure during prerestore activity on peer nodes, SnapCenter automatically rolls back the ASM disk group state as it was before performing restore on peer nodes on which prerestore operation was successful. Rollback is not supported for the primary and the peer node on which the operation failed. Before attempting another restore you must manually fix the issue on the peer node and bring the ASM disk group on the primary node back to mount state.
- If there is any failure during restore activity, then the restore operation fails and no roll back is performed. Before attempting another restore, you must manually fix the storage restore issue and bring the ASM disk group on the primary node back to mount state.
- If there is any failure during postrestore activity on any of the peer nodes, SnapCenter continues with the restore operation on the other peer nodes. You must manually fix the post restore issue on the peer node.

The in-place restore will be applied automatically when disk group or mount point matches the following criteria:

- No new datafiles are added after backup (foreign file check)
- No addition, deletion, or recreation of ASM disk or LUN after backup (ASM disk group structural change check)
- No addition, deletion, or recreation of LUNs to LVM disk group (LVM disk group structural change check)

Note: You can also forcefully enable in-place restore either using GUI, SnapCenter CLI, or PowerShell cmdlet to override the foreign file check and LVM disk group structural change check.

Even if you have forcefully enabled in-place restore, SnapCenter performs connect-and-copy restore in the following scenarios:

- Restore from secondary storage system and if Data ONTAP is earlier than 8.3
- Restore of ASM disk groups present on nodes of an Oracle RAC setup on which database instance is not configured
- In Oracle RAC setup, on any of the peer nodes if the ASM instance or the cluster instance is not running or if the peer node is down
- Restore of control files only
- Restore a subset of tablespaces residing on a ASM disk group

- Disk group is shared between data files, sp file, and password file
- SnapCenter Plug-in Loader (SPL) service is not installed or not running on the remote node in a RAC environment
- New nodes are added to the Oracle RAC and the SnapCenter Server is not aware of the newly added nodes

Types of restore operations supported for Oracle databases

SnapCenter enables you to perform different types of restore operations for Oracle databases.

Before restoring the database, backups are validated to identify whether any files are missing when compared to the actual database files.

Full restore

- Restores only the datafiles
- Restores only the control files
- Restores the datafiles and control files
- Restores datafiles, control files, and redo log files in Data Guard standby and Active Data Guard standby databases

Partial restore

- Restores only the selected tablespaces
- Restores only the selected pluggable databases (PDBs)
- Restores only the selected tablespaces of a PDB

Types of recovery operations supported for Oracle databases

SnapCenter enables you to perform different types of recovery operations for Oracle databases.

- The database up to the last transaction (all logs)
- The database up to a specific system change number (SCN)
- The database up to a specific date and time
You must specify the date and time for recovery based on the database host's time zone.

SnapCenter also provides the No recovery option for Oracle databases.

Note: The plug-in for Oracle database does not support recovery if you have restored using a backup that was created with the database role as standby. You must always perform manual recovery for physical standby databases.

Limitations related to restore and recovery operations

Before you perform restore and recovery operations, you must be aware of the limitations.

If you are using any version of Oracle from 11.2.0.4 to 12.1.0.1, the restore operation will be in hung state when you run the `renamedg` command. You can apply the Oracle patch 195447733 to fix this issue.

The following restore and recovery operations are not supported:

- Restore and recovery of tablespaces of the root container database (CDB)

- Restore of temporary tablespaces and temporary tablespaces associated with PDBs
- Restore and recovery of tablespaces from multiple PDBs simultaneously
- Restore of log backups
- Restore of backups to a different location
- Restore with recovery of tablespace and PDBs up to a specific SCN and date
- Restore of redo log files in any configuration other than Data Guard standby or Active Data Guard standby databases
- Restore of SPFILE and Password file
- When you perform a restore operation on a database that was re-created using the preexisting database name on the same host, was managed by SnapCenter, and had valid backups, the restore operation overwrites the newly created database files even though the DBIDs are different. This can be avoided by performing either of following actions:
 - Discover the SnapCenter resources after the database is re-created
 - Create a backup of the re-created database

Sources and destinations for restore operations

You can restore an Oracle database from a backup copy on either primary storage or secondary storage. You can only restore databases to the same location on the same database instance. However, in Real Application Cluster (RAC) setup, you can restore databases to other nodes.

Sources for restore operations

You can restore databases from a backup on primary storage or secondary storage. If you want to restore from a backup on the secondary storage in a multiple mirror configuration, you can select the secondary storage mirror as the source.

Destinations for restore operations

You can only restore databases to the same location on the same database instance.

In a RAC setup, you can restore RAC databases from any nodes in the cluster.

Defining a clone strategy for Oracle databases

Defining a strategy before cloning your database ensures that the cloning operation is successful.

Steps

1. Determine the backups that can be used for cloning.
2. Decide the type of cloning that you require.

Types of Oracle database backups supported for cloning

SnapCenter supports cloning of different types of backups of Oracle databases.

Oracle 11g, Oracle 12c (legacy), and Oracle 12c (CDB)

- Online data backup
- Online full backup

- Offline mount backup
- Offline shutdown backup
- Backups of Data Guard standby databases and Active Data Guard standby databases
- Online data backups, online full backups, offline mount backups, and offline shutdown backups in a Real Application Clusters (RAC) configuration
- Online data backups, online full backups, offline mount backups, and offline shutdown backups in an Automatic Storage Management (ASM) configuration

Important: Oracle ASM configuration is not supported if `user_friendly_names` option in the multipath configuration file is set to yes and aliases or symbolic links are defined for the Oracle ASM disks using the udev rules file.

Note: Cloning of archive log backups is not supported.

Types of cloning supported

In an Oracle database environment, SnapCenter supports cloning of a database backup. You can clone the backup from primary and secondary storage systems.

The SnapCenter Server uses NetApp FlexClone technology to clone backups.

Clone naming conventions

From SnapCenter 3.0, the naming convention used for clones of file systems is different from the clones of ASM disk groups.

- The naming convention for SAN or NFS file systems is *FileSystemNameofsource`database`_CLONESID*.
- The naming convention for ASM disk groups is *SC_HASHCODEofDISKGROUP_CLONESID*. *HASHCODEofDISKGROUP* is an automatically generated number (2 to 10 digits) that is unique for each ASM disk group.

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 195447733 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.

SnapCenter Plug-in for SAP HANA Database overview

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

SnapCenter supports single container and multitenant database containers (MDC): single tenant types of SAP HANA database.

When the Plug-in for SAP HANA Database is installed, you can use SnapCenter with NetApp SnapMirror technology to create mirror copies of backup sets on another volume. You can also use the plug-in with NetApp SnapVault technology to perform disk-to-disk backup replication for standards compliance. You can use the plug-in in both Windows and Linux environments.

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

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

Related information

[Protecting SAP HANA databases](#)

What you can do using the SnapCenter Plug-in for SAP HANA Database

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

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

SnapCenter Plug-in for SAP HANA Database features

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

Unified graphical user interface

The SnapCenter interface provides standardization and consistency across plug-ins and environments. The SnapCenter interface enables you to complete consistent backup, restore, and clone operations across plug-ins, use centralized reporting, use at-a-glance

dashboard views, set up role-based access control (RBAC), and monitor jobs across all plug-ins.

Automated central administration

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

Nondisruptive NetApp Snapshot copy technology

SnapCenter uses NetApp Snapshot copy technology with the Plug-in for SAP HANA Database to back up resources.

Using the Plug-in for SAP HANA Database also offers the following benefits:

- Support for backup, restore, and clone workflows
- RBAC-supported security and centralized role delegation
You can also set the credentials so that the authorized SnapCenter users have application-level permissions.
- Creation of space-efficient and point-in-time copies of resources for testing or data extraction by using NetApp FlexClone technology
A FlexClone license is required on the storage system where you want to create the clone.
- Support for the consistency group (CG) Snapshot copy feature of ONTAP as part of creating backups.
- Capability to run multiple backups simultaneously across multiple resource hosts
In a single operation, Snapshot copies are consolidated when resources in a single host share the same volume.
- Capability to create Snapshot copies using external commands.
- Support for file-based backup

Defining a backup strategy for SAP HANA databases

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

About this task

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

Steps

1. Determine when you should back up your resources.
2. Decide how many backup jobs you require.
3. Decide how to name your backups.
4. Decide whether you want to create a Snapshot copy-based policy to back up application-consistent Snapshot copies of the database.

5. Decide whether you want to verify the integrity of the database.
6. Decide whether you want to use NetApp SnapMirror technology for replication or NetApp SnapVault technology for long-term retention.
7. Determine the retention period for the Snapshot copies on the source storage system and the SnapMirror destination.
8. Determine whether you want to run any commands before or after the backup operation and provide a prescript or postscript.

Type of backups supported

Backup type specifies the type of backup that you want to create. SnapCenter supports File-Based Backup and Snapshot copy-based backup types for SAP HANA databases.

File-Based Backup

File-Based Backups verify the integrity of the database. You can schedule the file-based backup operation to occur at specific intervals. You cannot restore and clone File-Based backups from SnapCenter.

Snapshot copy based backup

Snapshot copy-based backups leverage NetApp Snapshot copy technology to create online, read-only copies of the volumes on which the SAP HANA databases reside.

How SnapCenter Plug-in for SAP HANA Database uses consistency group Snapshot copies

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

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

How SnapCenter manages housekeeping of log and data backups

SnapCenter manages the housekeeping of log and data backups on the storage system and file system levels, and within the SAP HANA backup catalog.

The Snapshot copies on the primary or secondary storage and their corresponding entries in the SAP HANA catalog are deleted based on the retention settings. The SAP HANA catalog entries are also deleted during backup and resource group deletion.

Considerations for determining backup schedules for SAP HANA database

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

Backup schedules have two parts, as follows:

- Backup frequency (how often backups are to be performed)
Backup frequency, also called *schedule type* for some plug-ins, is part of a policy configuration. For example, you might configure the backup frequency as hourly, daily, weekly, or monthly.

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

Number of backup jobs needed for SAP HANA databases

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

Backup naming conventions

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.

Types of restore strategies

You must define a strategy before you restore your resource so that you can perform restore operations successfully. SnapCenter enables you to perform two types of restore operations.

Complete resource restore

- Restores all volumes, qtrees, and LUNs of a resource
Complete resource restore operations perform a volume-based Snapshot copy restore (VBSR). VBSR is a disruptive operation that might impact other applications with hosted data on that volume. Therefore, you must have a separate RBAC permission to perform this operation.

File level restore

- Restores files from volumes, qtrees, or directories
- Restores only the selected LUNs

SnapCenter Custom Plug-ins overview

You can develop custom plug-ins for applications that you use and then use SnapCenter to backup, restore, or clone these applications. Like other SnapCenter plug-ins, your custom plug-ins act as host-side components of the NetApp SnapCenter Software, enabling application-aware data protection and management of resources.

When Custom Plug-ins are installed, you can use SnapCenter with NetApp SnapMirror technology to create mirror copies of backup sets on another volume and use NetApp SnapVault technology to perform disk-to-disk backup replication. The Custom Plug-ins can be used in both Windows and Linux environments.

NetApp provides MySQL and DB2 custom plug-ins with SnapCenter Software 2.0 and later and MongoDB custom plug-in from 3.0 and later. These plug-ins can be downloaded from the [NetApp Storage Automation Store](#).

Note: MySQL, DB2, and MongoDB custom plug-ins are supported via the NetApp communities only.

You can create your own custom plug-ins by referring to the Developer's Guide for Creating Custom Plug-ins.

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

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

Related information

[Protecting custom applications](#)

What you can do with the SnapCenter Custom Plug-ins

You can use the SnapCenter Custom Plug-ins for data protection operations.

- Add resources such as databases, instances, documents, or tablespaces.
- Create backups.
- Restore from backups.
- Clone backups.
- Schedule backup operations.
- Monitor backup, restore, and clone operations.
- View reports for backup, restore, and clone operations.

SnapCenter Custom Plug-ins features

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

Unified graphical user interface

The SnapCenter interface provides standardization and consistency across plug-ins and environments. The SnapCenter interface enables you to complete consistent backup, restore, recovery, and clone operations across plug-ins, use centralized reporting, use at-a-

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

Automated central administration

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

Nondisruptive NetApp Snapshot copy technology

SnapCenter uses NetApp Snapshot copy technology with the SnapCenter Custom Plug-ins to back up resources. Snapshot copies consume minimal storage space.

Using the Custom Plug-ins feature also offers the following benefits:

- Support for backup, restore, and clone workflows
- RBAC-supported security and centralized role delegation
You can also set the credentials so that the authorized SnapCenter users have application-level permissions.
- Creation of space-efficient and point-in-time copies of resources for testing or data extraction by using NetApp FlexClone technology
A FlexClone license is required on the storage system where you want create the clone.
- Support for the consistency group (CG) Snapshot copy feature of ONTAP as part of creating backups.
- Capability to run multiple backups simultaneously across multiple resource hosts
In a single operation, Snapshot copies are consolidated when resources in a single host share the same volume.
- Capability to create Snapshot copy using external commands.
- Capability to create file system consistent Snapshot copies in Windows environments.

Defining a backup strategy

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

About this task

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

Steps

1. Determine when you should back up your resources.
2. Decide how many backup jobs you require.
3. Decide how to name your backups.
4. Decide if you want Consistency Group Snapshot copies and decide on appropriate options for deleting Consistency Group Snapshot copies.

5. Decide whether you want to use NetApp SnapMirror technology for replication or NetApp SnapVault technology for long term retention.
6. Determine the retention period for the Snapshot copies on the source storage system and the SnapMirror destination.
7. Determine if you want to run any commands before or after the backup operation and provide a prescript or postscript.

Backup schedules of custom plug-in resources

The most critical factor in determining a backup schedule is the rate of change for the resource. The more often you back up your resources, the fewer archive logs SnapCenter has to use for restoring, which can result in faster restore operations.

You might back up a heavily used resource every hour, while you might back up a rarely used resource once a day. Other factors include the importance of the resource to your organization, your service-level agreement (SLA) and your recovery point objective (RPO).

SLA defines the level of service expected and addresses many service-related issues, including the availability and performance of service. RPO defines the strategy for the age of the files that must be recovered from backup storage for regular operations to resume after a failure. SLA and RPO contribute to the data protection strategy.

Backup schedules have two parts, as follows:

- Backup frequency
Backup frequency (how often backups are to be performed), also called schedule type for some plug-ins, is part of a policy configuration. For example, you might configure the backup frequency as hourly, daily, weekly or monthly. You can access policies in the SnapCenter GUI by clicking **Settings > Policies**.
- Backup schedules
Backup schedules (exactly when backups are to be performed) are part of a resource or resource group configuration. For example, if you have a resource group that has a policy configured for weekly backups, you might configure the schedule to back up every Thursday at 10:00 p.m. You can access resource group schedules in the SnapCenter GUI by clicking **Resources**, then selecting the appropriate plug-in, and clicking **View > Resource Group**.

Number of backup jobs needed

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

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

Backup naming conventions

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:

```
dtst1_mach1x88_03-12-2015_23.17.26
```

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

- dtst1 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.

Types of restore strategies

You must define a strategy before you restore your resource so that you can perform restore operations successfully. SnapCenter enables you to perform two types of restore operations.

Complete resource restore

- Restores all volumes, qtrees, and LUNs of a resource
Complete resource restore operations perform a volume-based Snapshot copy restore (VBSR). VBSR is a disruptive operation that might impact other applications with hosted data on that volume. Therefore, you must have a separate RBAC permission to perform this operation.

File level restore

- Restores files from volumes, qtrees, or directories
- Restores only the selected LUNs

SnapCenter Plug-in for VMware vSphere

For SnapCenter 4.2 and later, the SnapCenter Plug-in for VMware vSphere is deployed as part of the NetApp Data Broker virtual appliance Linux-based VM. It provides a vSphere web client GUI on vCenter to protect VMware virtual machines (VMs) and datastores, and supports SnapCenter application-specific plug-ins in protecting virtualized databases on primary and secondary storage.

[Data Protection Guide for VMs, Datastores, and VMDKs using the SnapCenter Plug-in for VMware vSphere](#)

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

What you can do with SnapCenter Plug-in for VMware vSphere

You can protect virtualized databases, and protect VMs and datastores.

Using the SnapCenter GUI

- Back up virtualized databases
The Plug-in for VMware vSphere enables other SnapCenter plug-ins to create application-consistent backups of virtualized applications (virtualized MS-SQL, Oracle, and Exchange databases). To take application-consistent backups of virtualized applications, you must deploy the NetApp Data Broker virtual appliance and enable the SnapCenter Plug-in for VMware vSphere in addition to installing the appropriate application-specific SnapCenter plug-ins. You use the SnapCenter GUI to perform your data protection operations.

Note: SnapCenter does not support single Snapshot copies of databases and VMs together.

Using the vSphere web client GUI

- Back up VMs, VMDKs, and datastores
 - You can back up VMs, underlying VMDKs, and datastores using the vSphere web client for SnapCenter Plug-in for VMware vSphere. When you back up a datastore, you back up all the VMs in that datastore.
 - You can back up on demand or according to a defined protection schedule.
 - You can define custom prescripts and postscripts before or after backup jobs to enable automation or custom action depending on your data protection needs.
 - You can create mirror copies of backups on another volume that has a SnapMirror relationship to the primary backup, or perform a disk-to-disk backup replication on another volume that has a SnapVault relationship to the primary backup volume, in addition to the backup Snapshot copy.
 - You can select crash-consistent or VM-consistent backup Snapshot copies when you define the policy.

Note: The Plug-in for VMware vSphere backs up VMFS and NFS datastores; it does not back up VSAN or VVOL datastores.

- Restore VMs from backups
You can restore an entire VM to the original ESXi host, not to an alternate ESXi host. When you restore a VM, you overwrite the existing content with the backup copy that you select, and also

unregister the original VM. You can restore data from either a primary backup or from a secondary backup on a different volume.

- **Restore VMDKs from backups**
You can restore a virtual disk of a VM to the original VM or to an alternate datastore.
- **Restore files and folders from a VM guest OS**
You can restore one or more files, or one or more folders, in a virtual disk.
 - Note:** You can restore files or folders from a Windows guest OS only.
 - Note:** You cannot search for a file or folder to be restored; before you begin the guest file restore operation you must know in which backup and VMDK the file or folder resides.
- **Attach or detach a virtual disk**
You can attach one or more virtual disks from a primary or secondary backup to the parent VM (the same VM that the virtual disk was originally associated with) or an alternate VM. This makes it easier to restore one or more individual files from a drive instead of restoring the entire drive. You can detach the virtual disk after you have restored the files you need.
- **Mount or unmount a datastore**
You can mount a datastore from a backup if you want to access files in the backup. You can either mount the backup to the same ESXi host where the backup was created or to an alternate ESXi host that has the same type of VM and host configurations. You can unmount a datastore backup when you no longer need to access the files in the datastore.
- **Monitor and report data protection operations on VMs and datastores**
 - The web client interface in vCenter displays an overview of the status of completed VM and datastore backups (primary and secondary), mount, and restore jobs. The vCenter interface also displays the number of protected and unprotected VMs.
 - You can monitor a VM or datastore job and view details associated with each job. You can also download logs for individual jobs for effective trouble-shooting in case of job failure.
 - You can view and download (HTML or CSV format) backup, restore, mount and protected/unprotected VM reports, and customize the reports by applying filters like time range, job status type, resource groups, and policies.
- **Control and coordinate user access in vCenter**
The Plug-in for VMware vSphere creates plug-in specific RBAC roles (SCV Administrator, SCV View, SCV Backup, SCV Restore, and SCV Guest File Restore) to control and coordinate access of users logging in to the vCenter interface. This provides better control and prevents inadvertent manipulation or unauthorized access of resources.

SnapCenter Plug-in for VMware vSphere features

The Plug-in for VMware vSphere integrates with vCenter and with NetApp technologies on the storage system.

Unified graphical user interface in vCenter

The VMware vSphere web client interface in vCenter enables you to complete consistent backup and restore operations on VMs, VMDKs, and datastores, and use centralized reporting, use at-a-glance dashboard views, set up role-based access control (RBAC), and monitor vCenter jobs.

Automated central administration

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

Nondisruptive NetApp Snapshot copy technology

The Plug-in for VMware vSphere uses NetApp Snapshot copy technology. This enables you to back up VMs, VMDKs, and datastores in seconds and restore VMs quickly without taking a host offline. Snapshot copies consume minimal storage space.

The Plug-in for VMware vSphere also offers the following benefits:

- Support for operations that were initialized by other plug-ins for virtualized resources and support for backup, restore, attach, detach, mount, and unmount operations on VMs, VMDKs, and databases in vCenter.
- Automatically discover VMs and databases configured on vCenter
- vCenter and SnapCenter RBAC-supported security and centralized role delegation
- Support for virtualized infrastructures

When to use the SnapCenter GUI and the vCenter GUI

SnapCenter Plug-in for VMware vSphere is different from other SnapCenter plug-ins because you use the web client GUI in vCenter for all backup and restore operations for VMs, VMDKs, and datastores. For all other plug-ins, you use the SnapCenter GUI for backup and restore operations. You can also use the Dashboard in the vCenter web client GUI to monitor the list of protected and unprotected VMs.

Note: The Plug-in for VMware vSphere supports the vCenter web client. It does not support vCenter thick clients.

To work with the Plug-in for VMware vSphere to protect VMs and datastores, you use the VMware vSphere web client interface in vCenter. The web client GUI integrates with NetApp Snapshot copy technology on the storage system. This enables you to back up VMs and datastores in seconds and restore VMs without taking a host offline.

Use this GUI..	To perform these operations...	And to access these backups...
Plug-in for VMware vSphere web client GUI in vCenter	VM and datastore backup VMDK attach and detach Datastore mount and unmount VM restore VMDK restore Guest file and folder restore	Backups of VMs and datastores performed by using the Plug-in for VMware vSphere GUI in vCenter.
SnapCenter GUI	Backup and restore of virtualized databases and applications, including protecting Microsoft SQL Server databases, Microsoft Exchange databases, and Oracle databases.	Backups performed by using the SnapCenter GUI.

Note: You must use the SnapCenter web client GUI in vCenter to perform SnapCenter operations. Although it is possible to perform some operations using VMware tools, for example, mounting or renaming a datastore, those operations will not be registered in the SnapCenter repository and, therefore, will not be recognized.

Note: SnapCenter does not support single Snapshot copies of databases and VMs together. Backups for VMs and databases must be scheduled and run independently, which creates separate

Snapshot copies, even if the databases and VMs are hosted in the same volume. Application backups must be scheduled by using the SnapCenter GUI; VM and datastore backups must be scheduled by using the SnapCenter web client GUI in vCenter.

Both vCenter and SnapCenter job monitors display all SnapCenter jobs, regardless of which GUI was used to start the job.

Example for protecting a Microsoft SQL Server database that is running on VMs

To perform an application-consistent backup of a Microsoft SQL Server database that is running on VMs, you use the SnapCenter GUI to start a backup of the database.

SnapCenter uses the SnapCenter Plug-in for Microsoft SQL Server and the SnapCenter Plug-in for Microsoft Windows to perform the operation, and uses the Plug-in for VMware vSphere to communicate with vCenter.

You also use the SnapCenter GUI to access or restore the backup.

Example for protecting an Oracle database that is running on VMs

To perform an application-consistent backup of an Oracle database that is running on VMs, you use the SnapCenter GUI to start a backup of the database.

SnapCenter uses the SnapCenter Plug-in for Oracle Database and the SnapCenter Plug-in for UNIX to perform the operation, and uses the Plug-in for VMware vSphere to communicate with vCenter.

You also use the SnapCenter GUI to access or restore the backup.

Example for protecting a VM

To perform a host-level backup of a VM in which a database application is running (not an application-consistent backup), you use the SnapCenter VMware vSphere web client GUI in vCenter to start the backup.

SnapCenter uses the Plug-in for VMware vSphere to perform the operation.

You also use the VMware vSphere web client GUI in vCenter to access or restore the backup.

Restoring VMs, VMDKs, files, and folders from backups

You can restore VMs and VMDKs from primary or secondary backups. VMs are always restored to the original host and datastore; VMDKs can be restored to either the original or an alternate datastore. You cannot use SnapCenter Plug-in for VMware vSphere to restore a datastore, only the individual VMs in the datastore. You can also restore individual files and folders in a guest file restore session, which attaches a backup copy of a virtual disk and then restores the selected files or folders.

VSC host migration to SnapCenter

You must migrate Virtual Storage Console for VMware vSphere (VSC) hosts and backup jobs to the SnapCenter Plug-in for VMware vSphere that is enabled by the NetApp Data Broker virtual appliance.

In VSC 7.0 and later, backup and restore functionality of VMs, VMDKs, and datastores is provided by the SnapCenter Plug-in for VMware vSphere. You can use the VSC 7.0 Appliance for storage management operations.

There are two migration options:

- Migrating from SnapCenter to the virtual appliance for SnapCenter Plug-in for VMware vSphere
This includes migrating application-consistent backups and VM-consistent backups performed by VSC in conjunction with SnapCenter.
You use SnapCenter PowerShell cmdlets to perform this migration. The NetApp Data Broker, SnapCenter Plug-in for VMware vSphere, documentation has more information.
[Data Protection Guide for VMs, Datastores, and VMDKs using the SnapCenter Plug-in for VMware vSphere](#)
- Migrating from VSC to the virtual appliance for SnapCenter Plug-in for VMware vSphere
This includes migrating backups performed by VSC with SnapManager for Virtual Infrastructure (SMVI).
You use the *NetApp Import Utility for SnapCenter and Virtual Storage Console* in the NetApp Support ToolChest to perform this migration.
[NetApp ToolChest: NetApp Import Utility for SnapCenter and Virtual Storage Console](#)

Copyright

Copyright © 2019 NetApp, Inc. All rights reserved. Printed in the U.S.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b).

Trademark

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

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

How to send comments about documentation and receive update notifications

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email.

[*doccomments@netapp.com*](mailto:doccomments@netapp.com)

To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

You can also contact us in the following ways:

- NetApp, Inc., 1395 Crossman Ave., Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

Index

A

- architecture
 - description of SnapCenter Server [10](#)
 - list of SnapCenter components [9](#)
 - overview of SnapCenter plug-ins [10](#)
 - SnapCenter diagram [8](#)
- ASM disk group clones
 - naming conventions [59](#)
- assign database
 - recovery models [37](#)
- Asymmetric LUN Mapping
 - support for [32](#)
- authentication
 - encryption [15](#)
- automation
 - scripting for in SnapCenter [22](#)

B

- backing up resources
 - determining how often [26, 34, 43, 52](#)
 - naming convention for copies [27, 35, 44, 53, 63, 66](#)
 - number of backups needed [44](#)
 - number of jobs needed [27, 35, 66](#)
 - retaining transaction log backups [28, 36](#)
 - scheduling verification [36](#)
 - strategy [25, 33, 47](#)
- backup copies
 - verifying [36, 53](#)
- backup jobs
 - factors that determine number needed [63](#)
- backup operations
 - optimum number of [27, 35, 44, 66](#)
- backup schedules
 - considerations for determining [62](#)
- backup strategies
 - defining [61](#)
 - defining for Windows file systems [43](#)
- backup types
 - modes [49](#)
 - offline backup described [49](#)
 - online backup described [49](#)
- backups
 - defining a strategy [65](#)
 - how SnapCenter creates backups of Windows file systems [42](#)
 - Oracle database configurations supported for backup [48](#)
 - resources, resource groups, and policies overview [16](#)
 - retention options [27, 35, 44, 53](#)
 - types of [49](#)
 - when to perform for custom plug-ins [66](#)
- backups of Oracle databases
 - supported for cloning [58](#)
 - supported for recovery operations [54](#)
 - supported for restore operations [54](#)
- best practices
 - restore strategy [28](#)

C

- catalogs
 - backup cataloging with Oracle RMAN [51](#)
- clone operations
 - limitations [40, 59](#)
 - resources, resource groups, and policies overview [16](#)
- clones
 - defining a cloning strategy [39](#)
 - defining a strategy [58](#)
 - limitations while creating [40, 59](#)
 - sources and destinations for [44](#)
 - supported types for Oracle databases [59](#)
 - types of clone operations [40](#)
- comments
 - how to send feedback about documentation [75](#)
- complete resource restore operations
 - explained [63, 67](#)
- considerations
 - for determining backups schedules [62](#)
- consistency group Snapshot copies
 - SAP HANA resource groups [62](#)
- copy-only backups
 - types of [33](#)
- crosscheck
 - update RMAN [51](#)
- custom plug-ins
 - features [64](#)
 - tasks you can perform using [64](#)

D

- database backups
 - supported types [62](#)
- database state
 - required for backup in RAC setup [50](#)
- databases
 - defining a cloning strategy [39](#)
 - destinations for a restore operation [37, 58](#)
 - discovering [49](#)
 - on the same volume [36](#)
 - restoration strategy [28, 36](#)
 - sources for a restore operation [28, 37, 58](#)
 - types of clone operations [40](#)
 - types of restore operations [28, 37](#)
- databases, Oracle
 - defining a clone strategy [58](#)
- defining a backup strategy
 - for Oracle databases [47](#)
- documentation
 - how to receive automatic notification of changes to [75](#)
 - how to send feedback about [75](#)
 - use of Concepts Guide [6](#)

E

- Exchange backup types

- full backups [26](#)
 - list of [26](#)
 - log backups [26](#)
 - transaction log backups [26](#)
- Exchange Server
 - database restoration strategy [28](#)
- F**
- features
 - Custom Plug-ins [64](#)
 - explained [60](#)
 - Plug-in for Oracle Database [46](#)
- feedback
 - how to send comments about documentation [75](#)
- file level restore operations
 - explained [63, 67](#)
- file names
 - choosing for backup copies [27, 35, 53, 63, 66](#)
 - for backup copies [44](#)
- file system clones
 - naming conventions [59](#)
- file-based backups
 - support for [62](#)
- G**
- getting started
 - with SnapCenter [24](#)
- I**
- identify
 - available databases [49](#)
 - discovering
 - database type [49](#)
 - database version [49](#)
- information
 - how to send feedback about improving documentation [75](#)
- J**
- jobs
 - number of backups needed [27, 35, 44](#)
- jobs, backup
 - defining a backup strategy before creating [43](#)
- L**
- limitations
 - clone operations [40, 59](#)
- logs
 - housekeeping of log and data backups [62](#)
 - retention of [28, 36](#)
 - types of Exchange backups [26](#)
 - types of SQL Server backups [33](#)
- M**
- MySQL Server
 - identifying a repository for SnapCenter metadata [12](#)
- N**
- naming conventions
 - differences between file system clones and ASM disk group clones [59](#)
- NSM database
 - repository for SnapCenter metadata [12](#)
- O**
- offline backups
 - described [49](#)
- online backups
 - described [49](#)
- Oracle databases
 - defining a clone strategy [58](#)
 - defining a restore and recovery strategy [54](#)
 - supported configurations for backup [48](#)
- Oracle RMAN
 - backup cataloging [51](#)
- overview
 - configuration check [15](#)
- P**
- permissions
 - associated with SnapCenter pre-defined roles [19](#)
- Plug-in for Exchange
 - tasks you can perform using [25](#)
- Plug-in for Microsoft SQL Server
 - features [31](#)
 - tasks you can perform using [30](#)
- Plug-in for Oracle database
 - backups supported for cloning [58](#)
- Plug-in for Oracle Database
 - features [46](#)
 - tasks you can perform using [46](#)
- Plug-in for SAP HANA Database
 - tasks you can perform using [60](#)
- policies
 - definition of [14](#)
 - SnapCenter scripts [22](#)
 - used in SnapCenter data protection [16](#)
- postscripts
 - supported by SnapCenter [22](#)
- preferred nodes
 - RAC setup [50](#)
- prescripts
 - supported by SnapCenter [22](#)
- R**
- RAC setup
 - preferred nodes [50](#)
- recovery
 - types of [57](#)
- recovery models
 - assign database to [37](#)
 - SQL Server [37](#)
- recovery operations

- limitations [57](#)
- recovery types
 - all logs [57](#)
 - date and time [57](#)
 - SCN [57](#)
- repositories
 - SnapCenter metadata [12](#)
- resource groups
 - used in SnapCenter data protection [16](#)
- resources
 - backing up Oracle catalogs [51](#)
 - considerations for determining backup schedules for [62](#)
 - definition of [14](#)
 - strategy for backing up [65](#)
 - types of restore operations for [63](#), [67](#)
 - used in SnapCenter data protection [16](#)
 - when to back up for custom plug-ins [66](#)
- resources, Exchange backup
 - types of [26](#)
- resources, SQL Server backup
 - types of [33](#)
- restore Exchange Server database
 - defining a strategy [28](#)
 - strategy [28](#)
- restore methods
 - connect-and-copy [55](#)
 - in-place [55](#)
 - types of [55](#)
- restore operations
 - limitations [57](#)
 - possible sources and destinations [58](#)
 - resources, resource groups, and policies overview [16](#)
 - strategy for Oracle databases [54](#)
 - types of [57](#), [63](#), [67](#)
- restore SQL Server database
 - defining a strategy [36](#)
 - strategy [36](#)
- restoring databases
 - destinations for a restore operation [37](#), [58](#)
 - sources for a restore operation [28](#), [37](#), [58](#)
 - types of restore operations [28](#), [37](#)
- retention
 - options for backups [27](#), [35](#), [44](#), [53](#)
- role-based access control (RBAC)
 - basics of setting up [14](#)
 - SnapCenter pre-defined roles and permissions [19](#)
- roles
 - pre-defined for SnapCenter [19](#)
- RPO (Recover Point Objective)
 - definition of [26](#), [34](#), [43](#), [52](#)

S

- same volume
 - databases [36](#)
- schedules
 - definition of frequency versus schedule [26](#), [34](#), [43](#), [52](#)
- scheduling
 - determining when to back up resources [26](#), [34](#), [43](#), [52](#)
- security

- overview [15](#)
- SLA (Service Level Agreement)
 - definition of [26](#), [34](#), [43](#), [52](#)
- SnapCenter
 - architecture [8](#)
 - basics and terminology [14](#)
 - components [9](#)
 - differences with SnapManager [14](#)
 - overview [7](#)
 - plug-ins overview [10](#)
 - Server overview [10](#)
 - using for the first time [24](#)
- SnapCenter Custom Plug-ins
 - features [64](#)
 - overview [64](#)
 - tasks you can perform using [64](#)
- SnapCenter Plug-in for Exchange Server
 - overview [25](#)
- SnapCenter Plug-in for Microsoft SQL Server
 - features [31](#)
 - overview [30](#)
 - tasks you can perform using [30](#)
- SnapCenter Plug-in for Microsoft Windows
 - description of how back ups are created [42](#)
 - features [41](#)
 - overview [41](#)
 - tasks you can perform [41](#)
- SnapCenter Plug-in for Oracle Database
 - overview of [46](#)
 - tasks you can perform using [46](#)
- SnapCenter Plug-in for SAP HANA Database
 - consistency group Snapshot copies [62](#)
 - overview [60](#)
 - WAFI sync for resource groups [62](#)
- SnapCenter Plug-in for VMware vSphere
 - which GUI to use [70](#)
- Snapshot copy-based backups
 - support for [62](#)
- SQL policy
 - retention of [36](#)
- SQL Server
 - database restoration strategy [36](#)
 - recovery models [37](#)
- SQL Server backup types
 - copy-only backups [33](#)
 - full backups [33](#)
 - list of [33](#)
 - log backups [33](#)
 - transaction log backups [33](#)
- strategy
 - backing up resources [25](#), [33](#)
- suggestions
 - how to send feedback about documentation [75](#)

T

- transaction logs
 - retaining backups [28](#), [36](#)
- Twitter
 - how to receive automatic notification of documentation changes [75](#)

V

- verifying backups
 - on primary or secondary storage [36, 53](#)
 - on SnapMirror or SnapVault storage [36, 53](#)
 - scheduling [36](#)
- VMs
 - which GUI to use [70](#)
- VMware vCenter
 - which GUI to use [70](#)

W

- WAFS sync
 - SAP HANA resource groups [62](#)
- Windows file systems
 - defining a backup strategy for [43](#)
 - sources and destinations of clones [44](#)
- workflows
 - SnapCenter backup overview [24](#)