



SnapCenter Plug-in for Microsoft Exchange Server concepts

SnapCenter Software 4.5

NetApp
February 12, 2024

This PDF was generated from https://docs.netapp.com/us-en/snapcenter-45/protect-sce/concept_snapcenter_plug_in_for_exchange_server_overview.html on February 12, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- SnapCenter Plug-in for Microsoft Exchange Server concepts 1
 - SnapCenter Plug-in for Microsoft Exchange Server overview 1
 - What you can do with SnapCenter Plug-in for Microsoft Exchange Server 1
 - Storage types supported by SnapCenter Plug-in for Microsoft Windows and for Microsoft Exchange Server 1
 - Minimum ONTAP privileges required 3
 - Prepare storage systems for SnapMirror and SnapVault replication 7
 - Define a backup strategy for Exchange Server resources 8
 - Define a restore strategy for Exchange databases 11

SnapCenter Plug-in for Microsoft Exchange Server concepts

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.

If you want to restore and recover mails or mailbox instead of the complete Exchange Database, you can use the Single Mailbox Recovery (SMBR) software. NetApp® Single Mailbox Recovery has come to the end of availability (EOA) on May 12, 2023. NetApp will continue to support customers that have purchased mailbox capacity, maintenance, and support through marketing part numbers introduced on June 24, 2020, for the duration of the support entitlement.

NetApp Single Mailbox Recovery is a partner product provided by Ontrack. Ontrack PowerControls offers capabilities that are similar to those of NetApp Single Mailbox Recovery. Customers can procure new Ontrack PowerControls software licenses and Ontrack PowerControls maintenance and support renewals from Ontrack (through licensingteam@ontrack.com) for granular mailbox recovery.

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

Storage types supported by SnapCenter Plug-in for Microsoft Windows and for Microsoft Exchange Server

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

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

Machine	Storage type	Provision using	Support notes
Physical server	FC-connected LUNs	SnapCenter graphical user interface (GUI) or PowerShell cmdlets	
Physical server	iSCSI-connected LUNs	SnapCenter GUI or PowerShell cmdlets	
VMware VM	RDM LUNs connected by an FC or iSCSI HBA	PowerShell cmdlets	Physical compatibility only  VMDKs are not supported.
VMware VM	iSCSI LUNs connected directly to the guest system by the iSCSI initiator	SnapCenter GUI or PowerShell cmdlets	 VMDKs are not supported.
Hyper-V VM	Virtual FC (vFC) LUNs connected by a virtual Fibre Channel Switch	SnapCenter GUI or PowerShell cmdlets	You must use Hyper-V Manager to provision Virtual FC (vFC) LUNs connected by a virtual Fibre Channel Switch.  Hyper-V pass through disks and backing up databases on VHD(x) that are provisioned on NetApp storage are not supported.

Machine	Storage type	Provision using	Support notes
Hyper-V VM	iSCSI LUNs connected directly to the guest system by the iSCSI initiator	SnapCenter GUI or PowerShell cmdlets	 <p>Hyper-V pass through disks and backing up databases on VHD(x) that are provisioned on NetApp storage are not supported.</p>

Minimum ONTAP privileges required

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

All-access commands: Minimum privileges required for ONTAP 8.2.x and later
event generate-autosupport-log
job history show
job stop

All-access commands: Minimum privileges required for ONTAP 8.2.x and later

lun

lun create

lun delete

lun igroup add

lun igroup create

lun igroup delete

lun igroup rename

lun igroup show

lun mapping add-reporting-nodes

lun mapping create

lun mapping delete

lun mapping remove-reporting-nodes

lun mapping show

lun modify

lun move-in-volume

lun offline

lun online

lun persistent-reservation clear

lun resize

lun serial

lun show

All-access commands: Minimum privileges required for ONTAP 8.2.x and later

snapmirror policy add-rule

snapmirror policy modify-rule

snapmirror policy remove-rule

snapmirror policy show

snapmirror restore

snapmirror show

snapmirror show-history

snapmirror update

snapmirror update-ls-set

snapmirror list-destinations

version

All-access commands: Minimum privileges required for ONTAP 8.2.x and later

volume clone create

volume clone show

volume clone split start

volume clone split stop

volume create

volume destroy

volume file clone create

volume file show-disk-usage

volume offline

volume online

volume modify

volume qtree create

volume qtree delete

volume qtree modify

volume qtree show

volume restrict

volume show

volume snapshot create

volume snapshot delete

volume snapshot modify

volume snapshot rename

volume snapshot restore

volume snapshot restore-file

volume snapshot show

volume unmount

All-access commands: Minimum privileges required for ONTAP 8.2.x and later

vserver cifs

vserver cifs share create

vserver cifs share delete

vserver cifs shadowcopy show

vserver cifs share show

vserver cifs show

vserver export-policy

vserver export-policy create

vserver export-policy delete

vserver export-policy rule create

vserver export-policy rule show

vserver export-policy show

vserver iscsi

vserver iscsi connection show

vserver show

Read-only commands: Minimum privileges required for ONTAP 8.2.x and later

network interface

network interface show

vserver

Prepare storage systems for SnapMirror and SnapVault replication

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



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

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



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

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

Define 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.

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.

Types of backups supported for Exchange database

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>

Backup type	Description
Log backup	<p>Backs up all the transaction logs.</p> <p>The truncated logs that are already committed to the database are not backed up. 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 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.

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.



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

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

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.

Define a restore strategy for Exchange databases

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

Sources for a restore operation in Exchange Server

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

You can restore databases from primary storage only.

Types of restore operations supported for Exchange Server

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.

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

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

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

Trademark information

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