



Technical Report

# Microsoft Exchange Server 2016/2013 and SnapManager for Exchange

## Best Practices Guide for Clustered Data ONTAP

Niyaz Mohamed, Krishna Vasudevan, NetApp  
March 2016 | TR-4221

## TABLE OF CONTENTS

<b>1</b>	<b>Executive Summary</b> .....	<b>5</b>
1.1	Purpose and Scope .....	5
1.2	Intended Audience .....	5
<b>2</b>	<b>Introduction to Clustered Data ONTAP</b> .....	<b>5</b>
<b>3</b>	<b>Microsoft Exchange Server Architecture</b> .....	<b>6</b>
3.1	Database Availability Groups .....	6
3.2	In-Place Archiving .....	7
<b>4</b>	<b>Microsoft Exchange Server—Planning Considerations</b> .....	<b>7</b>
4.1	System Requirements for Exchange 2016.....	7
4.2	System Requirements for Exchange 2013.....	7
<b>5</b>	<b>NetApp Storage Options for Microsoft Exchange Server</b> .....	<b>7</b>
<b>6</b>	<b>NetApp Storage Efficiency Technologies</b> .....	<b>8</b>
6.1	NetApp RAID DP Data Protection Technology .....	8
6.2	Snapshot.....	9
6.3	Thin Provisioning .....	9
6.4	Space Guarantee.....	9
6.5	Space Reclamation.....	10
6.6	Fractional Reserve.....	10
6.7	Autodelete and Autosize .....	10
6.8	Best Practice Configurations When Using Thin Provisioning for Microsoft Exchange Server 2016/2013 Environments.....	12
6.9	Monitoring .....	13
6.10	NetApp FlexClone.....	14
6.11	NetApp Deduplication .....	14
6.12	NetApp Compression.....	15
<b>7</b>	<b>Designing Storage Efficiency for Microsoft Exchange Server 2016/2013</b> .....	<b>16</b>
<b>8</b>	<b>NetApp Solution for Microsoft Exchange Server 2016/2013</b> .....	<b>17</b>
8.1	NetApp Storage Software and Tools.....	17
<b>9</b>	<b>Backup and Recovery</b> .....	<b>18</b>
9.1	Overview of SnapManager for Exchange Server .....	18
9.2	SnapManager for Exchange Server Architecture.....	19
9.3	SnapManager for Exchange Server Installation Considerations .....	19

9.4	Migrating Microsoft Exchange Data to NetApp Storage .....	20
9.5	Layout Recommendation .....	20
9.6	SnapManager Service Account.....	21
9.7	Prerequisites for Migrating Microsoft Exchange Server Mailbox Databases .....	22
9.8	Backup Best Practices .....	23
9.9	Backup .....	23
9.10	Gapless Backup .....	24
9.11	Server Backup and Frequent Recovery Point .....	26
9.12	Snapshot Retention Guidelines.....	29
9.13	Single Mailbox and Item-Level Recovery .....	30
9.14	Single Mailbox Recovery 7.1 and 7.2 Administrative Server.....	31
9.15	Troubleshooting .....	31
<b>10</b>	<b>High Availability .....</b>	<b>31</b>
10.1	Single-Site Scenario .....	31
10.2	Multisite Scenario.....	32
<b>11</b>	<b>Sizing and Storage Layout for Microsoft Exchange Server 2016/2013 .....</b>	<b>33</b>
11.1	Aggregate Recommendations.....	33
11.2	Storage Virtual Machine Recommendations .....	33
11.3	SAN LIF Types and Design.....	34
11.4	Volume Planning and Layout .....	34
11.5	Capacity Planning .....	37
<b>12</b>	<b>Performance.....</b>	<b>37</b>
12.1	NetApp Flash Cache Intelligent Data Caching.....	38
12.2	NetApp Flash Pool Intelligent Data Caching .....	38
12.3	SATA Performance Considerations .....	39
12.4	Database Sizing Considerations .....	39
12.5	Aggregate Sizing and Configuration Considerations.....	40
12.6	Volume Configuration Considerations.....	40
<b>13</b>	<b>Storage QoS and Nondisruptive Operations .....</b>	<b>40</b>
13.1	Storage Quality of Service .....	41
13.2	Nondisruptive Operations.....	41
<b>14</b>	<b>NetApp MetroCluster Software .....</b>	<b>41</b>
<b>15</b>	<b>Virtualization .....</b>	<b>42</b>
15.1	Microsoft Support for Exchange 2013 in Virtualized Environments.....	42

<b>16 Disaster Recovery .....</b>	<b>42</b>
<b>17 Summary .....</b>	<b>42</b>
<b>References.....</b>	<b>43</b>
<b>Version History .....</b>	<b>43</b>

**LIST OF TABLES**

Table 1) Option 1: volume guarantee set to none.....	12
Table 2) Option 2: using autogrow or autodelete.....	12
Table 3) Storage efficiency principles.....	16
Table 4) Nine gapless backup groups.....	26
Table 5) Three gapless backup groups.....	26
Table 6) Flash Cache options.....	38

**LIST OF FIGURES**

Figure 1) SnapManager for Exchange Server architecture.....	19
Figure 2) Gapless backup example.....	24
Figure 3) Sequence of operation.....	25
Figure 4) Backup schedule.....	26
Figure 5) Integration of SMBR with NetApp storage solution.....	30

# 1 Executive Summary

Many organizations have come to rely on Microsoft Exchange Server to facilitate critical business e-mail communication, group scheduling, and calendaring on a 24/7 basis. System failures could result in unacceptable operational and financial losses. Because of the increasing importance of Microsoft Exchange Server, data protection, disaster recovery, and high availability are of increasing concern. Companies require quick recovery with little or no data loss.

With the rapid growth of Microsoft Exchange Server databases, it is increasingly difficult to complete time-consuming backup operations quickly. When an outage occurs, it can take days to restore service from slower media such as tape, even if all of the backup tapes are available and error free. NetApp offers a comprehensive suite of hardware and software solutions that enable an organization to keep pace with the increasing data availability demands of an ever-expanding Microsoft Exchange Server environment. The solutions also scale to accommodate future needs while reducing cost and complexity.

NetApp® SnapManager® for Microsoft Exchange (SME) versions 7.0 and 7.1 are available for Microsoft Exchange Server 2013, and SnapManager for Microsoft Exchange version 7.2 is available for Microsoft Exchange Server 2016. SME is tightly integrated with Microsoft Exchange Server, which enables consistent online backups of your Microsoft Exchange environment while leveraging NetApp Snapshot® technology.

SME is a Volume Shadow Copy Service (VSS) requestor, which means that it uses the Microsoft VSS framework to initiate backups. SME provides a complementary feature set for Microsoft Exchange Server 2013 and the Microsoft Exchange Server 2016 data replication features. SME works with database availability group (DAG) databases on servers participating in a DAG and also on standalone servers, providing a rich feature set to leverage these new technologies.

## 1.1 Purpose and Scope

The success or failure of any software or infrastructure deployment hinges on making the proper design and architecture decisions in the planning phase. This guide provides recommended best practices for deploying Microsoft Exchange Server 2016/2013 and using SnapManager versions 7.x for Microsoft Exchange with a NetApp storage system running the NetApp clustered Data ONTAP® storage operating system and any supporting software. Organizations that want to get the most out of their NetApp storage investment for Microsoft Exchange Server will benefit from the recommendations in this report.

## 1.2 Intended Audience

This best practice guide is intended for experienced Microsoft Exchange Server administrators who are familiar with the installation and administration of NetApp SnapDrive® for Windows data management software, SME, and Data ONTAP.

In addition, readers should be well versed in the concepts surrounding Microsoft Exchange Server storage architecture, administration, backup, and restore. The recommendations in this document are best practices to assist with the design, implementation, and configuration of SME in Windows Server 2008 R2, Windows Server 2012, and Windows Server 2012 R2 environments with Microsoft Exchange Server 2016/2013.

# 2 Introduction to Clustered Data ONTAP

Storage controllers running clustered Data ONTAP are referred to as nodes. These nodes are joined into a clustered system. The nodes in the cluster continuously communicate with each other, coordinate cluster activities, and transparently move data between nodes.

Although the basic unit of a cluster is the node, nodes are added to the cluster as part of a high-availability (HA) pair. As with Data ONTAP operating in 7-Mode, HA pairs enable high availability by communicating with each other over an HA interconnect (separate from the dedicated cluster network) and by maintaining redundant connections to the HA pair's disks. Also, as is the case for Data ONTAP operating in 7-Mode, disks are not shared between HA pairs, although shelves might contain disks that belong to either member of an HA pair.

Clusters are administered on a whole-cluster rather than a per-node basis, and data is served from one or more storage virtual machines (SVMs, formerly called Vservers). Each SVM is configured to own storage, in the form of volumes provisioned from a physical aggregate, and logical interfaces (LIFs) assigned either to a physical Ethernet network or to Fibre Channel (FC) target ports. Logical disks (LUNs) are created inside an SVM's volumes and mapped to hosts to provide them with storage space. SVMs are node independent and cluster based; they can make use of physical resources such as volumes or network ports anywhere in the cluster.

For more details, see [TR-4080: Best Practices for Scalable SAN in Clustered Data ONTAP 8.3, which provides detailed best practices](#) for leveraging the HA and data mobility features of clustered Data ONTAP.

### 3 Microsoft Exchange Server Architecture

In Exchange Server 2016, a single building block provides the client access services and the high availability architecture necessary for the enterprise messaging environment. The mailbox server role performs the following functions:

- It handles the logic to route protocol requests to the correct destination endpoint.
- It hosts all of the components and/or protocols that process, render, and store data.
- It contains the client-access services that accept client connections. These services are routed through a local or remote proxy to the back-end services on the mailbox server that hosts the active database containing the user's mailbox.

As is the case for Exchange Server 2013, mailbox servers can be added to a DAG, thereby forming an HA unit that can be deployed in one or more datacenters. DAGs in Exchange Server 2016 have several significant enhancements:

- The parameter `DatabaseAvailabilityGroupIpAddresses` is no longer required when creating a DAG. By default, the failover cluster is created without an administrative access point.
- The parameter `Replay Lag Manager` is enabled by default.
- Database failover times are reduced by 33% when compared with Exchange Server 2013.

For more information about the architecture, refer to the [Microsoft Exchange Architecture site](#).

Microsoft Exchange Server 2013 includes the following server roles:

- [Client access server \(CAS\)](#). All client traffic connects to the now stateless CAS server.
- **Mailbox servers.** Maintain mailbox store databases, client access protocols, transport service, and unified messaging components.

In Exchange Server 2013, client access servers make up the CAS array, while mailbox servers comprise the DAG.

#### 3.1 Database Availability Groups

Microsoft Exchange Server 2016/2013 uses a [DAG](#) with a built-in log shipping feature called continuous replication with Microsoft clustering of nonshared storage. A DAG is a group of up to 16 nodes that provide automatic database-level recovery from failures that affect individual servers or databases.

## 3.2 In-Place Archiving

An archive mailbox is an additional mailbox associated with a user's primary mailbox. This new mailbox is provisioned automatically for the user when the administrator enables the personal archive feature.

After the archive mailbox has been associated with the user account, mail can be moved by the user into the personal archive by dragging and dropping mail items or automatically through retention policies.

# 4 Microsoft Exchange Server—Planning Considerations

## 4.1 System Requirements for Exchange 2016

This section describes the [system requirements](#) and [prerequisites](#) for Microsoft Exchange Server 2016 on NetApp storage systems:

- Windows Server 2012 or Windows Server 2012 R2 is required.
- The minimum and maximum page-file size must be set to the physical RAM size plus 10MB.
- Memory requirements vary depending on which Exchange roles are installed. For example, 8GB of memory is the minimum size for the mailbox role.
- [Disk space](#) depends on the roles installed. At least 30GB is required on the drive where Exchange is installed, and 500MB of free space is required on the drive that stores the message queue database.
- Disk partitions must be formatted as NTFS file systems.

## 4.2 System Requirements for Exchange 2013

This section describes the [system requirements](#) for Microsoft Exchange Server 2013 on NetApp storage systems:

- Windows Server 2008 R2 or later, Windows Server 2012, or Windows Server 2012 R2 is required.
- The minimum and maximum page-file size must be set to the physical RAM size plus 10MB.
- Memory requirements vary depending on which Exchange roles are installed.
- [Disk space](#) depends on the roles installed. At least 500MB of free space is required on the drive that stores the message queue database.
- Disk partitions must be formatted as NTFS file systems.

**Note:** Windows Server 2012 R2 is only supported with Microsoft Exchange 2013 SP1 or later.

For more detailed requirements, refer to the Microsoft TechNet article [Exchange 2013 Storage Configuration Options](#).

# 5 NetApp Storage Options for Microsoft Exchange Server

Many factors can affect your decision when selecting storage systems for Microsoft Exchange Server. To control costs, Microsoft recommends using SATA disks in simple JBOD (just a bunch of disks) directly attached storage enclosures. However, this recommendation might not be appropriate for all environments. Organizations that implement Exchange Server must understand the pros and cons of the different storage options to make a decision that is best for their environment.

One of the first questions that a company might ask when planning a storage platform for Exchange Server is "How much is the storage solution going to cost?" Although cost is a valid concern, the better question might be "Are we making the best decision for the money for a solution that meets our performance needs?" Organizations should consider the following criteria when selecting a storage platform for Exchange Server:

- The total size of the installation in terms of users and capacity

- The degree to which infrastructure can be shared across multiple applications
- Expected growth rates
- The fundamental need for server virtualization
- Backup and recovery requirements
- HA requirements
- Disaster recovery (DR) requirements

NetApp recommends the following storage options to address these criteria:

- Option 1: A fabric-attached storage (FAS) solution
- Option 2: An All Flash FAS solution
- Option 3: A hybrid solution that combines FAS systems and NetApp E-Series storage systems
- Option 4: An E-Series storage solution

This technical report describes best practices when deploying Microsoft Exchange Server 2016/2013 with Option 1 and Option 2.

## 6 NetApp Storage Efficiency Technologies

The NetApp strategy for storage efficiency is based on the built-in foundation of storage virtualization and unified storage provided by its core Data ONTAP operating system and the NetApp WAFL<sup>®</sup> (Write Anywhere File Layout) file system. Unlike other options, the NetApp technologies in the FAS and the NetApp FlexArray<sup>®</sup> storage virtualization feature have storage efficiency built into their core.

Customers who already have other vendors' storage systems and disk shelves can still leverage all of the storage-saving features that come with the NetApp FAS system simply by using FlexArray. This benefit again aligns with the NetApp philosophy of storage efficiency (helping increase storage use and decrease storage cost). It does so because customers can continue to use their existing third-party storage infrastructure and disk shelves and yet save more by leveraging the storage-efficiency technologies inherent in FlexArray.

NetApp offers the following technologies that increase storage efficiency.

### 6.1 NetApp RAID DP Data Protection Technology

NetApp RAID DP<sup>®</sup> technology safeguards against double-disk failure and delivers high performance. RAID DP is integrated with the WAFL file system to prevent the dedicated parity drives from becoming a performance bottleneck. RAID DP makes serial ATA (SATA) disks an option for your enterprise storage. Microsoft Exchange administrators can use less-expensive SATA without worrying about data loss, and they can also lower their storage acquisition costs.

RAID 4, RAID 5, RAID 6, and RAID DP all leverage parity so that data is not lost because of drive failures. These RAID options offer much better storage efficiency when compared with mirroring, but most RAID implementations have a drawback that affects write operations. This drawback is that performing a write operation requires multiple disk reads to regenerate the parity data. This requirement is commonly called the RAID penalty.

However, leveraging Data ONTAP does not incur a RAID penalty because WAFL is integrated with the RAID layer. Write operations, including parity generation, are coalesced in RAM and prepared as a complete RAID stripe. There is no need to perform a read to service a write, which means that Data ONTAP and WAFL avoid the RAID penalty.

Concerning reliability, statistically, RAID DP offers better protection than RAID mirroring because of the demands on disks during a RAID rebuild. With a mirrored RAID set, the risk of data loss from a disk failing

while rebuilding to its partner in the RAID set is much greater than the risk of a triple-disk failure in a RAID DP set.

## 6.2 Snapshot

NetApp Snapshot technology provides low-cost, fast-backup, point-in-time copies of the file system (volume) or LUN by preserving Data ONTAP architecture WAFL consistency points.

There is no performance penalty for creating Snapshot copies because data is never moved, as is the case with other copy-out technologies. The cost for Snapshot copies is limited to block-level changes rather than the whole backup, as is the case for mirror copies. Snapshot copies can reduce storage costs for backup and restore purposes and enable efficient data management possibilities.

## 6.3 Thin Provisioning

Thin provisioning in a shared storage environment optimizes the use of available storage. This method relies on the on-demand allocation of data blocks rather than the traditional method of allocating all of the blocks up front. This method eliminates almost all white space and helps avoid poor utilization rates. Flexible volumes (NetApp FlexVol<sup>®</sup> volumes) are the enabling technology behind NetApp thin provisioning, which can be thought of as the virtualization layer of Data ONTAP.

When a LUN is created, specific blocks are not dedicated out of the NetApp volume for the LUN or for Snapshot copies of the LUN. Instead, the blocks are allocated from the NetApp aggregate when the data is actually written. This approach allows the administrator to provision more storage space, as seen from the connected servers, than is actually physically present in the storage system.

When storage consumption is unpredictable or highly volatile, you should reduce the level of storage overcommitment so that storage is available for any growth spikes. Consider limiting storage commitment to 100%—no overcommitment—and using the trending functionality to determine how much, if any, overcommitment is acceptable.

Overcommitment of storage must be carefully considered and managed for a mission-critical application, such as Microsoft Exchange, for which even a minimal outage cannot be tolerated. In such a case, it is best to monitor storage consumption trends to determine how much, if any, overcommitment is acceptable.

If it takes a very long time to procure new storage, storage overcommitment thresholds should be adjusted accordingly. The overcommitment threshold should alert administrators early enough to allow them to procure and install new storage.

When configuring a Microsoft Exchange environment for thin provisioning, there is a chance that a LUN might go offline when there is not enough space to write further data. Use volume autogrow or volume move as mitigation mechanisms to safely enable thin provisioning and higher storage utilization.

## 6.4 Space Guarantee

The space guarantee function enables thin provisioning. Space guarantees can be set at the volume or the LUN level, depending on the space guarantee requirements of the application. Typically, if the space guarantee at the volume level is set to `volume`, the amount of space required by the flexible volume or the FlexVol volume is always available from its aggregate. This is the default setting for FlexVol volumes. When the space guarantee is set to `volume`, space is reserved from the aggregate's available space during volume creation.

When space guarantee is set to `none`, the volume reserves no space from the aggregate during volume creation. Space is first taken from the aggregate when data is actually written to the volume. Write operations to space-reserved LUNs in a volume with `guarantee=none` fail if the containing aggregate does not have enough available space.

LUN reservation enables space in the volume for the LUN, but `guarantee=none` does not enable space in the aggregate for the volume. When the space guarantee for the volume is set to `file`, the aggregate enables space for completely rewriting LUNs that have space reservation enabled.

#### Best Practice

NetApp recommends using thin provisioning in Microsoft Exchange environments to increase space utilization and to reduce the overall storage requirement when using the space guarantee function.

## 6.5 Space Reclamation

Space reclamation can be initiated from time to time to recover the unused space in a LUN. From the host's perspective, space reclamation to improve space utilization is coordinated with space consumption on the NetApp storage controllers. The SnapDrive Space Reclaimer can be used between the host and the controller to free up these deleted blocks, thus reducing storage utilization in the LUN and in Snapshot copies. Storage space can be reclaimed at the storage level by using the SnapDrive > Start Space Reclaimer option. The Data ONTAP PowerShell toolkit commandlet `invoke-nchostvolumespacereclaim` can also be used to perform space reclamation.

## 6.6 Fractional Reserve

Fractional reserve is a volume option that determines how much space Data ONTAP reserves for Snapshot overwrite data for LUNs to be used after all of the other space in the volume is used. Using fractional reserve is generally extremely wasteful because it is very unlikely that every byte in the database volume would be overwritten. There is no reason to reserve space for an event that rarely happens.

#### Best Practice

NetApp recommends setting the file space reserve (FSR) to 0 in Microsoft Exchange environments.

## 6.7 Autodelete and Autosize

The autosize volume setting (available in Data ONTAP 7.1 and later) defines whether a volume automatically grows to avoid filling up to capacity. You can define how quickly a volume grows with the `-i` option. The default growth increment is 5% of volume size at creation. It is also possible to define how large the volume can grow with the `-m` option. If volume autosize is enabled, the default maximum size is 120% of the original volume size. The following text provides an example:

```
vol autosize vol0 -m 1500g -i 1g on
vol status -v vol0
Volume autosize settings:
    state=on
    maximum-size=1500GB
    increment-size=1GB
```

#### Best Practices

- NetApp recommends planning for additional buffer space when using thin provisioning for Microsoft Exchange Server 2016/2013 environments.
- NetApp recommends prioritizing autosize over autodelete because deletions occur at the Data ONTAP level.
- NetApp recommends using the volume move functionality in clustered Data ONTAP to move volumes nondisruptively from an aggregate running out of space or requiring more IOPS than can be satisfied within the required response time.

Use of autodelete should be the lowest priority. When autodelete functionality is desired, NetApp recommends that the Snapshot backup retention functionality in SnapManager for Exchange be used instead of the Microsoft Exchange Server backup set—unaware autodelete feature in Data ONTAP. When properly configured, SnapManager deletes Snapshot copies as determined by the retention policy. With SME version 7.x, if Data ONTAP autodelete is used, SnapManager detects the orphaned Snapshot copies on the next run and honors the retention logic.

Snapshot copies should be deleted only by using SnapManager, with either the retention wizard or the delete backup wizard. In such scenarios, NetApp recommends using autosize. However, this function might fail because of space constraints in the aggregate and must be properly monitored by using Operations Manager. For volume autosize to work, the containing aggregate must have enough space (at least 1.2 times of the volume size).

Both autodelete and autosize work at the volume level, not on individual LUNs. This means that LUNs do not automatically grow and must be handled separately with different commands. NetApp SnapDrive for Windows (SDW) can be used to make more space available for the LUN.

The autodelete volume setting (available in Data ONTAP 7.1 and later) allows Data ONTAP to delete Snapshot copies if a threshold, called a trigger, is met. This setting can be configured so that Snapshot copies are automatically deleted when one of the following conditions is met:

- **Volume.** The volume is nearly full. This is reported in the first line report for each volume in the `df` command.  
**Note:** The volume can be full even though there might still be space in the `snap_reserve` areas.
- **snap\_reserve.** The snap reserve space is nearly full.  
**Note:** Snap reserve is automatically disabled by SnapDrive, so NetApp recommends not using this trigger type.
- **space\_reserve.** The overwrite-reserved space is full. This is the space determined by the LUNs with space reservations enabled and the `fractional_reserve` option. The reserve space is never filled until both the volume and the `snap_reserve` areas are full.

**Note:** The `df` command is available when NetApp storage is accessed through the CLI.

```
6240b> df
File system      Kbytes      used  avail      capacity  Mounted on
/vol/vol10/     1407415772  12094808  73618924  1%        /vol/vol10/
/vol/vol10/.snapshot  74074512   455588   73618924  1%        /vol/vol10/.snapshot
```

## Best Practices

- NetApp recommends using autogrow instead of autodelete. When using autodelete, set the autodelete trigger to `volume`.
- Always use NetApp SnapManager retention capabilities. NetApp recommends not using scripts to remove the log snapshots (especially `eloginfo`) because doing so can trigger an orphaned snapshot situation.

The order in which Snapshot copies are deleted is determined by the following three options:

- **delete\_order.** This option determines whether the oldest or the newest Snapshot copies are deleted first.
- **defer\_deleted.** This option allows the user to define a group of Snapshot copies that are deleted first when no other Snapshot copies are available. You can defer the deletion of user-created Snapshot copies, scheduled Snapshot copies, or Snapshot copies beginning with a configurable prefix.
- **Commitment.** This option determines how Snapshot copies used for NetApp SnapMirror® data replication software and dump operations are handled. If set to `try`, this option deletes these

Snapshot copies only if they are not locked. If this option is set to `disrupt`, these Snapshot copies are deleted even if they are locked.

### Best Practice

When using SnapMirror or NetApp SnapVault® products for replicating Microsoft Exchange Server 2016/2013 databases, NetApp recommends not using the `disrupt` option for commitment. This is because SnapMirror baseline Snapshot copies can be destroyed by autodelete even though they will always be the last Snapshot copies deleted. In many configurations, deleting the last SnapMirror Snapshot copy is not desired because a new full baseline copy is required to resume mirroring operations. If, for example, the source and destination are at different sites, recreating this baseline can be a time-consuming and costly process.

## 6.8 Best Practice Configurations When Using Thin Provisioning for Microsoft Exchange Server 2016/2013 Environments

There are many ways to configure a NetApp storage appliance for LUN thin provisioning; each has advantages and disadvantages. It is possible to have thinly provisioned volumes and volumes that are not thinly provisioned on the same storage system or even on the same aggregate. The settings listed in Table 1 are best practice configurations when using thin provisioning for Microsoft Exchange Server 2016/2013.

Table 1) Option 1: volume guarantee set to none.

Setting	Configuration
Volume guarantee	None
LUN reservation	Enabled
<code>fractional_reserve</code>	0%
<code>snap_reserve</code>	0%
Autodelete	<code>volume / oldest_first</code>
Autosize	Off
<code>try_first</code>	<code>snap_delete</code>

The advantage of this configuration is that free space in the aggregate is used as a shared pool of free space. The disadvantages of this configuration are the high level of dependency between volumes and that the level of thin provisioning cannot easily be tuned on an individual volume basis.

When this configuration is used, the total size of the volumes is greater than the actual storage available in the host aggregate. With this configuration, storage administrators can generally size the volume so that they must manage and monitor only the used space in the aggregate. This option does not affect the space for hosting the live data, but rather allows the backup space to dynamically change.

Table 2 lists the settings for using the `autogrow` or `autodelete` function.

Table 2) Option 2: using `autogrow` or `autodelete`.

Setting	Configuration
Volume guarantee	Volume
LUN reservation	Disabled

Setting	Configuration
fractional_reserve	0%
snap_reserve	0%
Autodelete	volume / oldest_first
Autosize	On
try_first	autogrow

This configuration allows the administrator to finely tune the level of thin provisioning for Microsoft Exchange Server 2016/2013 environments. With this configuration, the volume size defines or guarantees space that is only available to LUNs within that volume. The aggregate provides a shared storage pool of available space for all of the volumes contained in it.

If the LUNs or Snapshot copies require more space than is available in the volume, the volumes automatically grow, taking more space from the containing aggregate. In addition, the advantage of having the LUN space reservation disabled is that Snapshot copies can then use the space that is not needed by the LUNs. The LUNs themselves are also not in danger of running out of space because the autodelete feature removes the Snapshot copies that are consuming space.

This is an ideal setting for many migrations in which Snapshot copy space is high during the initial mailbox moves, but tapers off in the following months and years when more space is required in the database to store mail.

**Note:** Snapshot copies used to create NetApp FlexClone® volumes are not deleted by the autodelete option.

#### Best Practice

NetApp recommends using autogrow for most common deployment configurations.

## 6.9 Monitoring

When NetApp efficiency features are used, the volumes should be appropriately sized so that autosize and/or autodelete policies are not triggered unless there is an abnormal rate of change or a problem with Snapshot copy retention. NetApp recommends NetApp OnCommand® Unified Manager Core Package management software, which includes [Operations Manager](#), to monitor Microsoft Exchange volumes for these events and to send notifications to the storage administration team to follow up with the Microsoft Exchange administration team. Simple Network Management Protocol (SNMP) can also be used to monitor these events.

After the storage administration team receives notification of a volume autogrow or Snapshot autodelete event, NetApp recommends that the team examine the affected storage controllers. The team should then follow up with the Microsoft Exchange administration team for further administrative actions.

A typical cause of volume autosize events is that the rate of change greatly surpassed the rate of change assumption used in sizing the volume. Adding more Microsoft Exchange mailboxes beyond the original database design parameters or e-mail storms can cause increased data change rates. Another cause for volume autosize events is that older Snapshot copies created by SME are not deleted. As Snapshot copies age, they can grow in size and consume more capacity than originally allocated in the volume.

A typical cause of SME not deleting backups is that SME backups are failing. By default, SME does not delete Snapshot copies of older SME backup sets if the backup fails. Another cause for SME not deleting backups is that the SME backup retention policies are not being enforced correctly because Snapshot copies were manually removed outside of SME on the controller.

The health of SME can be monitored by monitoring for SME event IDs and by the enhanced enterprise monitoring functionality in SME versions 7.x and later. To monitor the health of SME retention external to SME, the number of Snapshot copies and SnapInfo directories should be calculated for a specific Microsoft Exchange Server. For example, if the SME retention policy for a particular server is 10 backups online (`-RetainBackups 10` parameter in the `new-backup` command), there are then 10 SME Snapshot copies in each Microsoft Exchange database volume (with a prefix of `exchsnap`). There are 20 SME Snapshot copies in each Microsoft Exchange transaction log volume (10 with a prefix of `exchsnap` and 10 with a prefix of `eloginfo`).

If the SME retention policy for a particular server is 10 days of backups online (`-RetainDays 10` parameter in the `new-backup` command), SME Snapshot copies will be available in the Microsoft Exchange database and transaction log volumes no older than 10 days.

An alternate way of calculating SME retention when using the `-RetainDays backup` parameter is to multiply the number of days that you keep backups online by the number of backups taken each day. If one backup per day is taken, then there would be 10 SnapManager for Exchange Snapshot copies in each Microsoft Exchange database volume. There would be 20 SME Snapshot copies in each Microsoft Exchange transaction log volume.

To monitor the health of SnapManager retention, use Windows PowerShell commands from the [Data ONTAP PowerShell Toolkit](#) and native Windows PowerShell commands.

## 6.10 NetApp FlexClone

A FlexClone volume is a writable point-in-time Snapshot copy of a FlexVol volume or another FlexClone volume. FlexClone uses space very efficiently, leveraging the Data ONTAP architecture to store only data that changes between the parent and the clone. FlexClone volumes are appropriate for any situation in which testing or development occurs or any situation in which progress is made by locking in incremental improvements. These volumes are also appropriate for any situation in which there is a desire to distribute data in a changeable form without endangering the integrity of the original. A common scenario is to use FlexClone in an environment before committing a Microsoft Exchange Server 2016/2013 cumulative update or hotfix into production.

FlexClone technology can be leveraged at both the primary storage system and the SnapMirror destinations for effective use of resources. FlexClone can also be used for disaster-recovery testing without affecting the operational continuity of the Microsoft Exchange Server 2016/2013 environment.

### Best Practice

Use SnapDrive to connect to the required Snapshot copy and automatically execute the FlexClone operation.

## 6.11 NetApp Deduplication

The deduplication process stores only unique blocks of data in the volume and creates additional metadata in this process.

Each 4KB block in the storage system has a digital fingerprint that is compared to other fingerprints on the volume. If two fingerprints are found to be the same, a byte-for-byte comparison is made of all bytes in the block. If they are an exact match, the duplicate block is discarded and the space is reclaimed.

The core enabling technology of deduplication is fingerprints. When deduplication runs for the first time on a FlexVol volume, it scans the blocks and creates a fingerprint database that contains a sorted list of all fingerprints for used blocks in the flexible volume.

Deduplication consumes system resources and can alter the data layout on disk. Because of the application I/O pattern and the effect of deduplication on the data layout, the read and write I/O

performance can vary. Deduplication might also have a positive effect on performance because the system memory and NetApp Flash Cache™ modules are both aware of the deduplicated blocks. As a block is read, it is inserted into memory or cache. If this deduplicated block is accessed by another operation, it will be accessed from physical memory as opposed to the spinning disk, resulting in much improved access time.

**Note:** Setting `read_realloc` to `on` for a volume that has enabled deduplication neither affects performance nor reduces storage efficiency.

**Note:** Deduplication is transparent to Microsoft Exchange, and the block changes are not recognized by Microsoft Exchange. Therefore, the Microsoft Exchange database remains unchanged in size from the host's perspective, even though there are capacity savings at the volume level.

**Note:** Tests have shown space savings on Microsoft Exchange Server 2016/2013 databases in the 15 to 35% range.

#### Best Practices

- Deduplication rates cannot be predicted and should not be used when sizing capacity.
- Deduplication can provide additional capacity for user growth and/or increased Snapshot retention. NetApp recommends deduplication for database volumes, but not for transaction log volumes.
- Turn scheduled deduplication on and schedule it for off-peak hours (late at night).

To configure deduplication, see the [Clustered Data ONTAP 8.3 Logical Storage Management Guide](#).

## 6.12 NetApp Compression

Every e-mail that is retained in the mailbox database is also stored in the database copies. NetApp inline compression and Snapshot capabilities help reduce the overall storage footprint without compromising I/O performance. The value of compression is not limited to an All Flash FAS environment. Hybrid options such as NetApp Flash Pool™ aggregates also benefit from compression. Data ONTAP 8.3.1 and later introduced adaptive compression, an inline compression method that works with blocks of varying sizes. The performance effect is minimal, and enabling compression can improve overall performance in some cases.

The compression method available in Data ONTAP 8.3 and earlier is still available and is now called secondary compression. There is no single best practice for the use of compression; the best option depends on business practices.

Adaptive compression has been thoroughly tested with Exchange workloads, and the performance effect has been found to be negligible, providing the sizing is accurate. This is true even in an all-flash environment in which latency is measured in microseconds, showcasing the substantial space savings available in the era of multiple DAG copies.

In initial testing, some customers have reported a performance increase with the use of compression. This is the result of compression effectively increasing the amount of Flash Pool caching available to the database. Secondary compression can be enabled on volumes hosting archive mailboxes, which hold archived mail items for users, to achieve maximum space efficiency.

Data ONTAP manages physical blocks in 4KB units. Therefore, the maximum possible compression ratio would be 2:1 with a typical Exchange database. Early testing with real customer data has shown compression ratios approaching this level and, in certain cases, far beyond.

## Version-Specific Differences in Data ONTAP Data Compression

Different versions of Data ONTAP employ data compression in different ways:

- In Data ONTAP 8.2, the use of compression on a volume prevents data from being cached by Flash Pool.
- In Data ONTAP 8.3, compressed blocks are eligible for read but not write Flash Pool data caching.

- In Data ONTAP 8.3.1, compressed blocks are eligible for both read and write Flash Pool data caching.
- Flash Cache can be used with compressed volumes, but the data is stored in the flash layer in an uncompressed format and does not see any performance gains.

#### Best Practices

- With Data ONTAP 8.3.1 or later, enable adaptive compression for all database volumes for the simplest data compression method on All Flash FAS (enabled by default) and for hybrid aggregates.
- In an All Flash FAS environment, enable inline compression with inline zeroing and deduplication for greater space savings.

## 7 Designing Storage Efficiency for Microsoft Exchange Server 2016/2013

Table 3 lists the principles for designing storage efficiency for Microsoft Exchange Server 2016/2013.

Table 3) Storage efficiency principles.

Principle 1	
Statement	Employ RAID DP.
Rationale	RAID DP provides double-disk failure protection in a RAID group with minimal storage needs.
Implications	RAID DP provides a high level of disk failure fault tolerance without sacrificing performance and storage efficiency.
Principle 2	
Statement	Use SATA disks wherever appropriate.
Rationale	SATA disks can be a good fit in environments with high-capacity requirements. The higher read latency associated with SATA drives can be reduced by using the NetApp Flash Cache card.
Implications	Economical, high-capacity disks have better mapping to Microsoft Exchange environments with larger mailbox sizes and high-capacity demands.
Principle 3	
Statement	Thin provision the storage with NetApp FlexVol volumes.
Rationale	Thin provisioning in a Microsoft Exchange environment increases storage utilization, reduces complexity by provisioning the sized LUNs to the host, and reduces overall storage requirements. As utilization grows, storage can be added without affecting Microsoft Exchange.
Implications	Storage silos can be removed.
Principle 4	
Statement	Whenever appropriate, use server virtualization technologies with NetApp storage to gain additional savings.
Rationale	NetApp deduplication and cloning technology is extremely effective in virtual infrastructures for the guest machine's operating system footprint, which provides additional storage savings.

Implications	NetApp deduplication and cloning technology enables a more dynamic and storage-efficient infrastructure.
<b>Principle 5</b>	
Statement	Using deduplication and inline compression in Microsoft Exchange Server 2016/2013 environments can reduce the amount of storage used.
Rationale	NetApp provides an in-place deduplication strategy, applicable for both primary and secondary storage.
Implications	Depending on the message profile and attachments, using deduplication and inline compression can save anywhere from 15% to 40% of storage space.

## 8 NetApp Solution for Microsoft Exchange Server 2016/2013

### 8.1 NetApp Storage Software and Tools

This section describes the NetApp storage software and tools used in this solution:

- **NetApp Windows Host Utilities Kit.** Installation of the Host Utilities Kit sets timeout and other operating system–specific values to their recommended defaults. This kit includes utilities for examining the LUNs provided by NetApp storage, whether clustered or operating in 7-Mode. This kit should be used in both physical and virtual environments because it helps to align the master boot record for the Microsoft virtual hard disk (VHD) file layout, preventing it from getting out of alignment with the underlying NetApp LUN. This is very important for optimal I/O performance. This kit is not necessary on Windows Server 2012 or Windows Server 2012 R2 when the Data ONTAP Device Specific Module (DSM) is installed.
- **Microsoft Windows and native multipath input/output (MPIO).** To operate as intended, clustered Data ONTAP requires MPIO and Asymmetric Logical Unit Access (ALUA). For Microsoft Windows 2008, Windows Server 2012, and Windows 2012 R2, these features are natively supported whenever the MPIO feature is installed.

When using the iSCSI protocol, you must configure multipath support on iSCSI devices in the MPIO Properties administrative application. Navigate to the Discover Multi-Paths pane, select Add Support for iSCSI Devices, and click Add. It is also necessary to create multiple sessions from the host initiators to the target iSCSI LIFs on the clustered Data ONTAP system. You can do so by using the native iSCSI initiator. Select Enable Multi-Path when logging on to a target. Set the default load balance policy of all the devices to the least queue depth or round-robin.

Sessions can also be managed by using the NetApp SnapDrive iSCSI management pane. This is the preferred method, because SnapDrive remembers which target logical interfaces already have an established session and preselects an unused target portal.

- **Data ONTAP DSM.** The Data ONTAP DSM supports attachment to a Data ONTAP cluster beginning with version 3.5. Consult the [NetApp Interoperability Matrix Tool](#) for current information on supported configurations.

Windows 2008 and Windows 2012 support MPIO and ALUA natively. However, the Data ONTAP DSM should be used over the native MPIO implementation when accessing NetApp LUNs to discover which paths are direct and indirect so that traffic is routed appropriately. DSM also has the advantage of a GUI in the Microsoft Management Console that correctly displays LIF and SVM names, making management simpler and more intuitive.

During installation, the Data ONTAP DSM sets a number of Windows registry values to optimize performance and provide correct behavior during failover.

- **NetApp SnapDrive for Windows.** SDW performs the same essential functions for clustered Data ONTAP as it does for Data ONTAP operating in 7-Mode:

- SDW includes a VSS hardware provider to verify that Snapshot copies of an NTFS LUN are consistent. SME provides consistency at the application level.
- SDW enables the management of iSCSI sessions by using the iSCSI management pane. This functionality is very useful now that multiple sessions are required for selection of direct paths during failover and volume move scenarios.
- SDW enables on-the-fly LUN resizing and cloning.

**Note:** NetApp SnapDrive can be used in conjunction with clustered Data ONTAP beginning with version 6.4.

The unit of management when using NetApp SnapDrive with clustered Data ONTAP is at the level of individual SVMs, not at the node or cluster level. As virtual storage servers, SVMs are meant to provide a secure multitenancy environment in which specific volumes, network interfaces, target ports, and management users are logically grouped. Therefore, a host connected to multiple SVMs that are part of the same cluster and access the same physical disks and network interfaces would not have visibility into the fact that they are logical entities distributed across physically separate but interconnected nodes.

When first accessing an SVM with NetApp SDW, you must access the Transport Protocol settings under the server list in the SnapDrive management pane. HTTP and HTTPS are the supported protocols with SnapDrive in clustered Data ONTAP.

## 9 Backup and Recovery

### 9.1 Overview of SnapManager for Exchange Server

SnapManager for Exchange provides an integrated data management solution for Microsoft Exchange Server 2016/2013 that enhances the availability, scalability, and reliability of Microsoft Exchange databases. SME provides rapid online backup and restoration of databases, along with local or remote backup set mirroring for disaster recovery.

SME uses online Snapshot technologies that are part of Data ONTAP. It integrates with Microsoft Exchange backup and restores APIs and the VSS. SnapManager for Exchange can use SnapMirror to support disaster recovery even if native Microsoft Exchange DAG replication is used.

SME provides the following data-management capabilities:

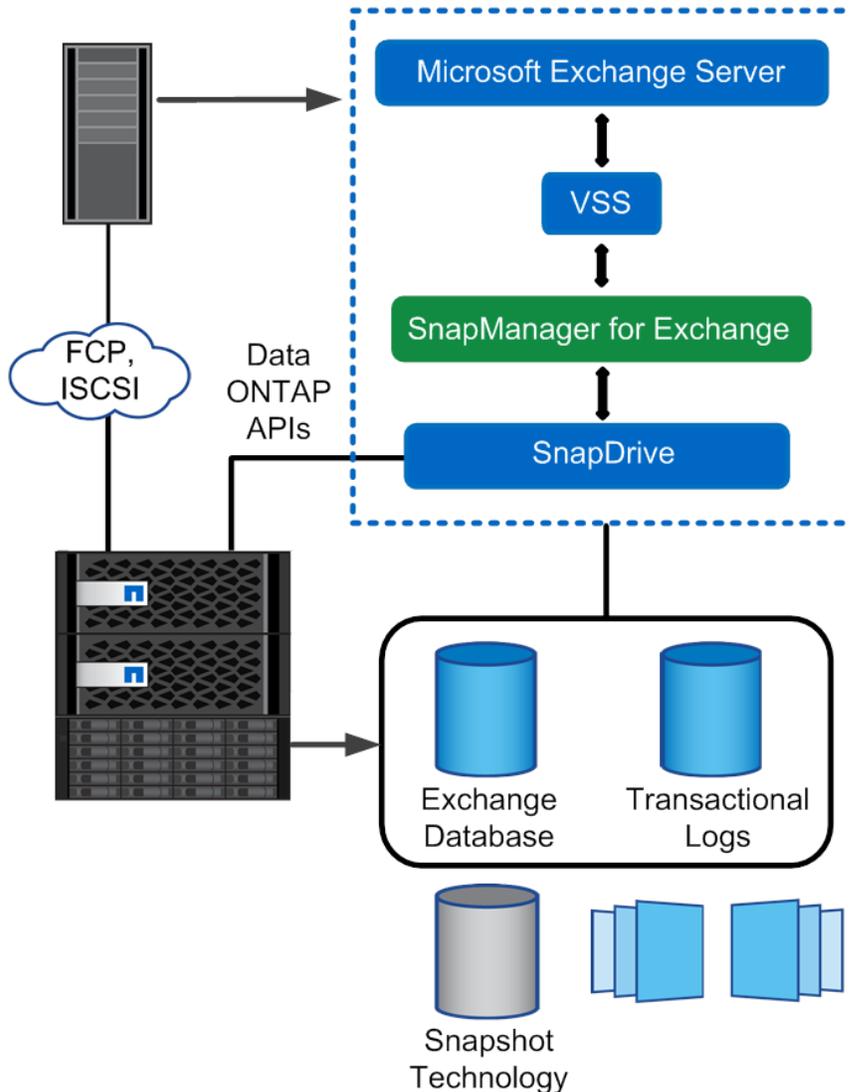
- Migration of Microsoft Exchange data from local or remote storage onto NetApp LUNs
- Application-consistent backups of Microsoft Exchange databases and transaction logs from NetApp LUNs
- Verification of Microsoft Exchange databases and transaction logs in backup sets
- Management of backup sets
- Archiving of backup sets
- Restoration of Microsoft Exchange databases and transaction logs from previously created backup sets, providing a lower recovery time objective (RTO) and more frequent recovery point objectives (RPOs)
- Capability to reseed database copies in DAG environments and prevent full reseed of a replica database across the network
- Support for IP-less DAG in Exchange 2013 and 2016 environments
- Capability to restore a passive database copy without having to reseed it across the network by using the database Reseed wizard
- Native SnapVault integration without using Protection Manager
- RBAC support for service account
- Retention enhancements

## 9.2 SnapManager for Exchange Server Architecture

SnapManager for Microsoft Exchange versions 7.0 and 7.1 support Microsoft Exchange Server 2013, and SnapManager for Microsoft Exchange version 7.2 supports Microsoft Exchange Server 2016. SME is tightly integrated with Microsoft Exchange, which allows consistent online backups of Microsoft Exchange environments while leveraging NetApp Snapshot technology. SME is a VSS requestor, meaning that it uses the VSS framework supported by Microsoft to initiate backups. SME works with the DAG, providing the ability to back up and restore data from both active database copies and passive database copies.

Figure 1 shows the SnapManager for Exchange Server architecture. For more information about VSS, refer to the [Volume Shadow Copy Service Overview](#) on the Microsoft Developer Network.

Figure 1) SnapManager for Exchange Server architecture.



## 9.3 SnapManager for Exchange Server Installation Considerations

For information about compatible versions of SnapManager for Exchange, SDW, and Data ONTAP, refer to the [NetApp Interoperability Matrix Tool](#).

Before upgrading SnapManager for Exchange, perform the following tasks:

1. Back up the operating system installation on the Microsoft Exchange Server. This process includes backing up all of the server system state, which consists of the registry, the boot files, and the COM+ class registry.
2. Back up the data on the local drives on the Microsoft Exchange Server.
3. Back up the boot and system drives.

**Note:** NetApp provides backup tools for physical (Syncsort) and virtual servers (VSC/SMHV).

## 9.4 Migrating Microsoft Exchange Data to NetApp Storage

The process of migrating Microsoft Exchange databases and transaction log files from one location to another can be a time-consuming and lengthy process. You must take many manual steps so that the Microsoft Exchange database files are in the proper state to be moved. In addition, you must perform more manual steps to bring those files back online for handling Microsoft Exchange traffic. SME automates the entire migration process, eliminating any potential user errors. After the data is migrated, SME automatically mounts the Microsoft Exchange data files and allows Microsoft Exchange to continue serving e-mail.

## 9.5 Layout Recommendation

Storage layout for Microsoft Exchange data should be designed for optimal performance, high availability, and data protection before the actual migration of the Microsoft Exchange databases is carried out. The best practices for designing the storage layout for Microsoft Exchange environments are discussed in section 11, “Sizing and Storage Layout for Microsoft Exchange Server 2016/2013.” For additional design information, see the [Installation and Setup Guide for SnapManager for Exchange](#).

The following paragraphs present a comprehensive list of guidelines that must be followed when designing the Microsoft Exchange data layout on NetApp storage so that SME functions correctly.

### Best Practice

Keep Microsoft Exchange Server 2016/2013 databases on individual LUNs on separate volumes.

Placing databases on individual LUNs on separate volumes enables fast and granular recovery. However, the databases can also be placed on individual LUNs on the same volume. This arrangement allows you to restore the databases individually because SnapManager for Exchange uses LUN restore for restore operations. Keeping multiple databases on the respective individual LUNs on a single volume also provides slightly higher deduplication rates because multiple databases are hosted on a single volume. Preferably, the first of these two options should be used. The second option should be considered for cases in which the volume count might cause problems.

**Note:** For Microsoft Exchange Server 2016/2013, multiple mailbox databases cannot be placed on the same LUN for SnapManager for Exchange to function.

Typically, backups are run on one of the passive database copies in DAG environments. Make sure that the transaction log LUN on the database copy where backups are run with longer retention is correctly sized by taking the backup retention into account. Not doing so can lead to situations in which the backup retention is different for other copies of the database, and the administrator might change the Snapshot retention without considering the initial volume sizing. However, if the gapless backup feature is used, be sure to correctly size the transaction log LUN on all of the database copies in which backup will be run with longer retention.

For additional information on Exchange data layout and sizing, refer to section 11, “Sizing and Storage Layout for Microsoft Exchange Server 2016/2013.”

The SnapInfo directory in SME is a central repository containing two types of data:

- Backup metadata
- Transaction log backups

In SME, if the SnapInfo directory is placed in the same path as the log folder for a database, then SME creates an NTFS hard link for each log file in the SnapInfo directory while backing up that particular database. This action not only saves space but also makes the log backups quicker.

#### Best Practices

- Place the Microsoft Exchange transaction log files and the SnapInfo directory in the same LUN
- Place Microsoft Exchange transaction log files into as few flexible volumes as possible, as appropriate for the business use and the RPO and RTO.
- NetApp recommends using NTFS hard links by placing the SnapInfo directory on the same LUN as the transaction log directory whenever possible. This arrangement increases storage utilization, eliminates the physical copy overhead incurred on the Microsoft Exchange Server, and increases backup performance.
- When separate LUNs are used for the Microsoft Exchange transaction log files and the SnapInfo directory, place those LUNs in the same volume. Both of these LUNs have a similar I/O profile, allowing them to share the same volume. In disaster-recovery scenarios, having the entire log set for Microsoft Exchange on the same volume helps to achieve SLAs. Placing the SnapInfo directory in a LUN different from the transaction logs requires additional I/O during backup to physically move the log files that are scheduled to be truncated.

**Note:** In environments with high LUN counts, transaction logs for multiple mailbox databases can be placed on a single LUN. When deploying LUNs in this manner, NetApp recommends limiting the number of log streams per LUN to 5 to 10.

## 9.6 SnapManager Service Account

Microsoft Exchange Server uses role-based access control (RBAC) permissions. For more information, see [Understanding Role-Based Access Control](#).

SME 7.1 and later understand the Microsoft Exchange RBAC model and impersonate the service account permissions. This feature provides administrators with more flexibility, allowing them to define broad categories of permissions or to narrow the scope as much as they want.

To use SME, you must have a service logon account with appropriate permissions. At a minimum, the following permissions are required:

- Unlike earlier versions of SME, an SME 7.1 or later service account must have a minimum set of role entries assigned to perform backup and restore operations.
- The SnapManager for Exchange Service must be a member of the Microsoft Exchange Server's local administrators group,
- The SME service account used for backup and restore of Microsoft Exchange 2010, 2013, or 2016 must have the following role entries:
  - `get-exchangeserver`
  - `Get-Mailboxserver`
  - `get-mailboxdatabase`
  - `get-publicfolderdatabase`
  - `Get-DatabaseAvailabilityGroup`
  - `Get-mailboxstatistics`
  - `Get-MailboxDatabaseCopyStatus`
  - `Set-MailboxDatabase`

- Set-PublicFolderDatabase
- New-MailboxDatabase
- New-PublicFolderDatabase
- Set-mailboxdatabase -allowfilerestore:\$true
- Set-ADServerSettings
- Mount-Database
- Dismount-Database
- Move-DatabasePath -ConfigurationOnly:\$true
- Remove-MailboxDatabaseCopy
- Suspend-MailboxDatabaseCopy
- Update-MailboxDatabaseCopy
- Resume-MailboxDatabaseCopy
- Remove-MailboxDatabaseCopy
- Remove-PublicFolderDatabase
- Add-MailboxDatabaseCopy
- Set-MailboxDatabaseCopy
- Move-ActiveMailboxDatabase

Based on this requirement, different role groups can be created with a different set of permissions in order to implement RBAC in SME.

For more information about creating role groups and role scopes, see the following articles:

- [Create a Role Group](#)
- [Understanding Management Role Scopes](#)

## 9.7 Prerequisites for Migrating Microsoft Exchange Server Mailbox Databases

In SME versions 7.x, the DAG is a unit of management: All nodes of a DAG can be managed by registering the DAG with SME versions 7.x. Optionally, each DAG node can also be managed at the server level.

Before migrating the Microsoft Exchange Server mailbox databases on a DAG, make sure that NetApp storage is provisioned on each server that will have a database copy that uses the same path specified during migration. Install SnapManager for Exchange on all of the member servers of the DAG before migrating the mailbox databases.

If SnapManager for Exchange is not installed on a member server of the DAG, all databases hosted by that member server are not migrated by SnapManager, including both the active and the passive databases. Also, SnapManager is not able to back up databases on that server.

Make sure that the database replication status is healthy before migration.

In situations in which not all member servers of a DAG use NetApp storage to store Microsoft Exchange data, two key questions arise:

- What is the deployment strategy?
- What is the licensing policy?

SnapManager for Exchange can be installed in a mixed-vendor storage environment that contains both NetApp and third-party storage. Within this environment, the following requirements must be met:

- For versions earlier than SDW 7.1, SnapDrive must be installed on all nodes in a Microsoft failover cluster, even if the nodes are connected to third-party storage. With SDW 7.1, however, SnapDrive should be installed only on DAG nodes that are connected to NetApp storage.
- It is not necessary to have SnapManager installed on every node. In this configuration, the backup must be made at the server level.

In this scenario, SME cannot be used to connect to the DAG. The user must connect to individual mailbox servers on which SME is installed and perform tasks related to SME, including the migration of mailbox databases. This operation must be performed individually on every member node that uses NetApp storage and has SME installed on it.

Furthermore, in this scenario, SME can be licensed on only those nodes on which it is installed. However, there is a risk involved for situations in which multiple mailbox databases reside on the DAG and the backup policy dictates that backups should be created on the active server. In the event of database failover for one of the many databases that reside in the DAG, a DAG member server that does not have SME might become the active server. In that case, even if the member server is connected to NetApp storage, SME cannot be used on it.

**Note:** When a Microsoft Exchange 2016/2013 DAG is deployed, the sector sizes of the volumes hosting the databases and log files must be the same across all of the nodes in the DAG. This requirement is outlined in [Exchange 2013 Storage Configuration Options](#).

#### Best Practices

- On Microsoft Exchange Server 2016/2013, use the same drive letters or mount points for the Microsoft Exchange data LUNs on all nodes of a DAG.
- Install SnapManager for Exchange and SDW on all member servers of the DAG. If SnapManager for Exchange is not installed on all nodes, SME cannot migrate databases by using the configuration wizard.

## 9.8 Backup Best Practices

The storage design requires careful planning in order to meet the customer's backup frequency. This section discusses the concepts related to backing up Microsoft Exchange data by using SnapManager for Exchange.

## 9.9 Backup

It is important to consider the following factors for planning a backup strategy for Microsoft Exchange data in the organization:

- **Organization SLA.** This parameter determines the frequency and the type of backups.
- **Backup retention planning.** This parameter determines whether backup sets must be retained on the primary or the secondary site.
- **Backup verification policy.** This parameter determines when backup verification is engaged.

The time required to restore Microsoft Exchange data during an outage depends on the number of transaction logs that must be replayed. Therefore, reducing the number of transaction logs that must be replayed when restoring from a full backup is important. The only way to reduce the number is to create more frequent backups.

An important consideration for continued data protection in a Microsoft Exchange environment is to verify that backups are completed successfully as planned. Doing so calls for active monitoring of the backup jobs being run by SME. There are two ways to monitor the health of backup jobs:

- **The SME notification system.** This feature in SME allows administrators to receive detailed e-mails for failures in executing operations in SME. E-mail notification can be configured by using the SME configuration wizard. See the [SnapManager for Exchange Administration Guide](#) for more details.
- **Monitoring host-side events.** SME logs many events, each with its own event ID. These events are logged in the Windows application event log on the Microsoft Exchange Server. Monitoring some of the critical events by using scripts, Microsoft System Center Operations Manager, Microsoft Operations Manager, or other options helps alert administrators to failures encountered by SME.

RPOs have become a defining part of any data protection plan for Microsoft Exchange. The ability to have a near-zero RPO is highly desirable for Microsoft Exchange administrators because it minimizes the amount of data that is lost between the last fully verified backup set and the point of failure.

To help achieve the desired SLA and RPO times, SME has frequent recovery points (FRPs). FRPs are optimized backup sets that are created through SME. The backup sets contain only the transaction log files that have been created since the last full backup or FRP backup that was created. The transaction log files are hard-linked or copied into the SnapInfo directory, and then a Snapshot copy is created. An FRP backup set contains a small amount of information. Backups can be created as often as every 10 minutes. Having a higher frequency of FRP backups reduces RPO times.

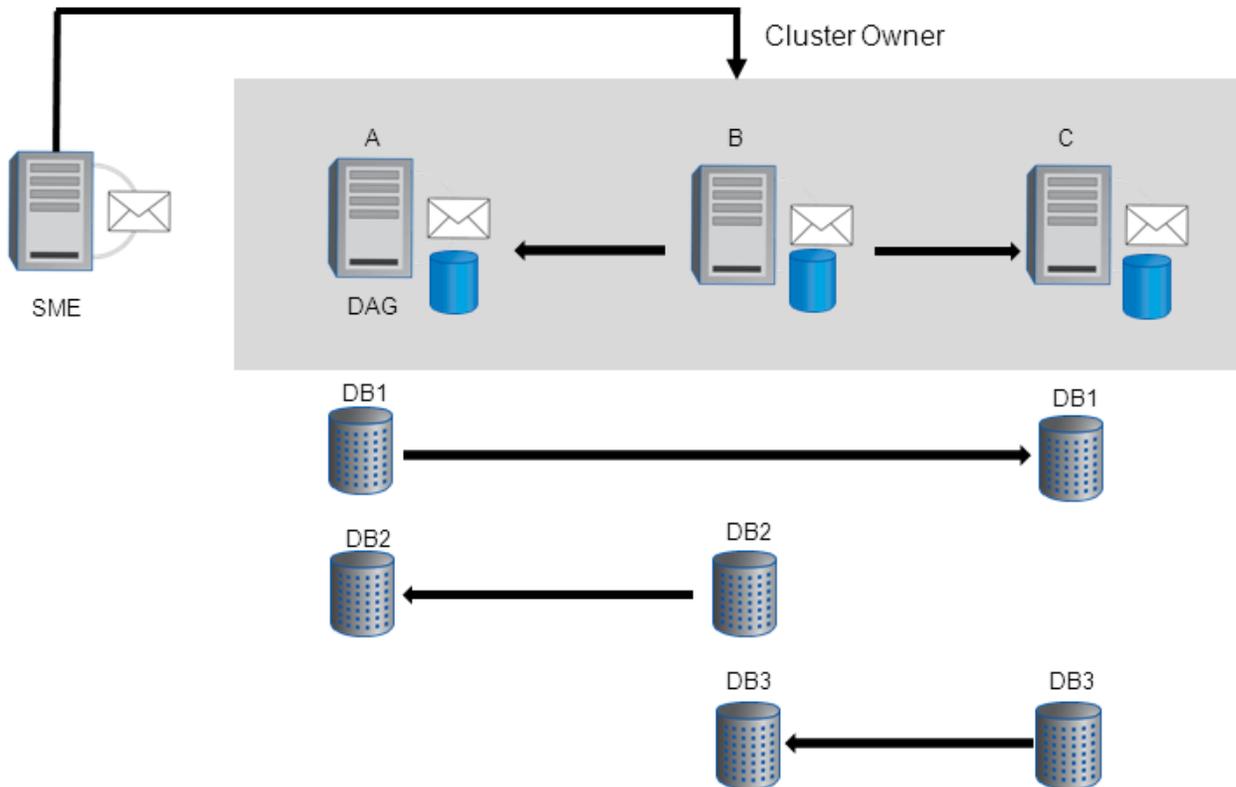
## 9.10 Gapless Backup

The gapless backup feature is designed so that a Snapshot copy that is older than the most recent full backup, which truncates the transaction logs, can use up-to-the-minute restore (roll-forward recovery).

### SnapManager for Exchange Gapless Backup Example

Figure 2 shows a typical workflow for gapless backup.

Figure 2) Gapless backup example.



The following list explains key points in Figure 2:

- This is a three-node DAG, and node B is the owner node of the DAG
- Each database has two copies. The arrows indicate the log-shipping direction and point to the passive database copy.
- SME is shown running on a client machine, although it can be run from any of the DAG nodes

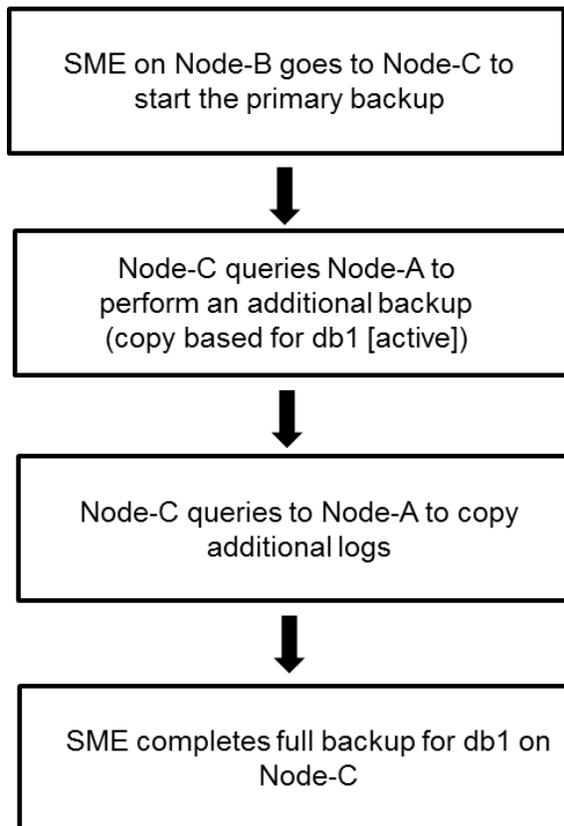
The administrator selects a full backup on all the databases. Additional remote copy backup must be selected to perform a copy backup on all passive database copies.

**Note:** If the DAG consists of up to nine nodes, choose the gapless backup strategy based on the SME remote additional copy backup feature. Remote additional copy backup greatly simplifies the administration of backups.

## Sequence of Operation

When SME connects to Node B, which is the DAG owner, it receives the list of databases to back up and then starts the backup operation, as shown in Figure 3.

Figure 3) Sequence of operation.



In the end, there is a backup for all six database copies, and all backups can perform an up-to-the minute restore. Because the job was initiated on Node B, the overall operation resides on Node B. The database layout plays a major role in the backup completion time for gapless backup. Make sure that the databases are placed in an optimal manner to achieve the best results while using gapless backup.

For example, the database layout listed in Table 4 shows nine backup groups created for gapless backup subjobs.

Table 4) Nine gapless backup groups.

DAG Node 1	DAG Node 2	DAG Node 3
DB1 (Active)	DB1	DB1
DB2	DB2 (Active)	DB2
DB3	DB3	DB3 (Active)

The database layout listed in Table 5 shows three backup groups created for gapless backup sub-jobs.

Table 5) Three gapless backup groups.

DAG Node 1	DAG Node 2	DAG Node 3
DB1 (Active)	DB1	DB1
DB2 (Active)	DB2	DB2
DB3 (Active)	DB3	DB3

For more information about the optimal layout, refer to [SnapManager 7.1 Microsoft Exchange Server documentation](#) and [SnapManager 7.2 Microsoft Exchange Server documentation](#)

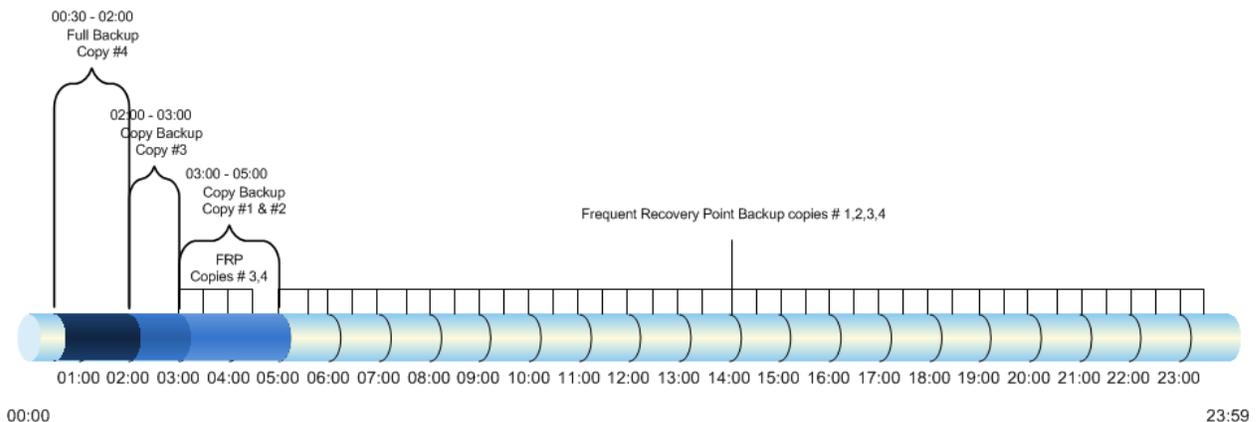
## 9.11 Server Backup and Frequent Recovery Point

A smaller RPO can be achieved in larger environments with many databases and server nodes in a DAG. The sample configuration shown in Figure 4 has been successfully implemented in large customer environments. The strategy uses a full backup on one copy of the database, which does not have to be the active database copy. The remaining database copies run a copy backup, which does not truncate logs. All databases participate in a frequent recovery point backup every 30 minutes. If one copy of the database fails, the rapid-reseed process can be used to perform a nondisruptive restore.

The backups happen on the following schedule:

- 00:30–02:00 database copy #4 full backup
- 02:00–03:00 database copy #3 copy backup
- 03:00–04:00 database copy #1 and database copy #2 copy backup
- 03:00–05:00 database copy #3 and database copy #4 FRP backup
- 05:00–00:00 all database copies FRP backup every 30 minutes

Figure 4) Backup schedule.



It is important that a copy backup is scheduled at least 10 minutes after the full backup so that the log truncation activity from the full backup successfully replicates to each database copy.

**Note:** If the DAG is larger than nine nodes, choose the Server Backup and Frequent Recovery Point Backup strategy.

**Note:** A database can also be restored from a full backup with FRP or a copy backup with FRP using SME.

#### Best Practice

NetApp recommends using server-based full and copy backups with frequent recovery points to achieve the smallest RPO.

## IP-less DAG

Microsoft Exchange Server 2013 SP1 or later supports IP-less DAG. A DAG without an Administrative Access Point reduces the complexity and simplifies DAG management.

DAGs without cluster administrative access points have the following characteristics:

- There is no IP address assigned to the cluster or DAG and therefore no IP address resource in the cluster core resource group.
- There is no network name assigned to the cluster and therefore no network name resource in the cluster core resource group.
- The name of the cluster or DAG is not registered in the Domain Name System, and it is not resolvable on the network.
- A cluster name object is not created in Active Directory.
- The cluster cannot be managed by using the Failover Cluster Management tool. It must be managed using Windows PowerShell, and PowerShell cmdlets must be run against individual cluster members.

SME 7.2 supports IP-less DAGs and enables you to connect to DAGs without an admin access point running on Exchange Server 2016/2013 environments.

## Database Verification

Starting with Microsoft Exchange 2010, database verification is not required for databases with [two or more database copies in a DAG](#). However, Microsoft recommends verifying the transaction logs for each of the database copies. By default, SME verifies logs during restore operations. When a DAG backup is performed with SME, verification is off by default. You can monitor verification jobs running both locally and remotely through the main SME management console.

A single Microsoft Exchange mailbox server can run only one verification process at a time on a particular verification server. A verification server can simultaneously run one verification job from each Microsoft Exchange mailbox server. More than one verification server can be used to simultaneously verify more than one backup job on a single Microsoft Exchange mailbox server. Many customers use virtual machines to offload verification.

SnapManager for Exchange also supports the verification of backups on SnapMirror destinations and SnapVault secondary locations. This arrangement offloads the read I/O from the production database servicing users.

## Recovery

The ability to recover Microsoft Exchange databases is a critical operation for a Microsoft Exchange administrator. SME restore functionality allows you to recover Exchange databases and transaction logs from backups that it created or from the SnapVault archive. SME has two types of restore operations:

- **Up-to-the-minute.** Selected by default, an up-to-the-minute restore replays any necessary and available transaction logs from the backup set and from the transaction log directory and applies them to the database. A contiguous set of transaction logs is required for an up-to-the-minute restore to succeed.
- **Point-in-time.** This option allows you to restore your Microsoft Exchange data to a chosen point in time. Any Microsoft Exchange data past that point is not restored. This option is particularly useful when restoring to a point before an event such as data corruption has occurred. A point-in-time restore replays and applies to the database only those transaction logs that existed in the active file system when the backup was created up to the specified point in time. All transaction logs beyond that time are discarded.

**Note:** If the most recent backup is unverified, the verification can be done prior to the recovery, or the recovery can continue without verification (for a quicker recovery).

#### Best Practice

When performing an up-to-the-minute restore, restore from your most recent backup to minimize the number of transaction logs that must be replayed.

## Reseeding Passive Copies in a DAG Setup

This section discusses best practices and solutions when using SME to restore the passive copies of Microsoft Exchange databases.

One of the key challenges in a DAG environment is to minimize reseeding of database copies in the event of a database failure. The rapid reseed functionality using Snapshot copies is much faster than the out-of-box reseeding solution from Microsoft. When using the reseed feature from SME 7.1 and later, network resources are not consumed when a Snapshot copy of the failed database is restored.

Consider a scenario in which a passive copy of the database has been in a failed state for a while with no replication enabled between it and the active database copy. Using SME 7.1 and later, you can reseed a passive database copy that is in a nonhealthy state (`Failed`, `FailedandSuspended`, or `Suspended`) to restore it to the latest Snapshot copy and place the passive database copy in a healthy state.

Reseeding a database copy covers two different scenarios:

- For databases residing on individual LUNs on separate volumes, the reseed workflow performs a volume-based Snapshot restore.
- For multiple databases residing on their respective individual LUNs on a single volume, the reseed workflow performs a LUN restore.

The active copy of the database is not dismounted during the reseed of the replica database copy. The purpose of the replica database restore without dismounting the database is to reseed the database very quickly by using a Snapshot copy rather than manually seeding through the network.

The purpose of a passive database restore operation is to reseed the database more quickly using a Snapshot copy, based on the following reasons:

- The passive copy is corrupt, the exchange resynchronization fails, and the passive copy is unrecoverable without a reseed.
- Some environments with a slow WAN link across DAG nodes might require a long time to reseed a database copy.
- Restoring from a backup of a currently passive copy will recover the passive copy without changing the (newer) content of the active copy.

Before performing the reseed operation from SME 7.1 and later, one must confirm that all of the services associated with Microsoft Exchange are running.

## Best Practice

NetApp recommends configuring databases on individual LUNs on separate volumes to leverage fast and granular reseed functionality.

## 9.12 Snapshot Retention Guidelines

### Primary Storage

The RPO determines how frequently a backup is taken. NetApp flexible volumes running in clustered Data ONTAP can store a maximum of 255 Snapshot copies per flexible volume. The amount of storage needed for Snapshot copies depends on the rate of change.

Consult a local NetApp Exchange expert or your NetApp partner to provide accurate volume sizing and layout for Microsoft Exchange environments.

### Secondary Storage

SME backups can be archived with SnapVault. Volumes can be copied for space-efficient, read-only, disk-to-disk backup. SME 7.1 and later have native SnapVault integration.

**Note:** SnapVault updating is an asynchronous operation. An update is triggered after an SME backup is completed. A schedule can be created to fit the user's requirements.

## Best Practices

- The IP address and the credentials of the secondary SVM must be added in the transport protocol settings in SDW.
- Verify that all of the volumes that are part of a dataset have SnapVault enabled. These volumes include those containing the database file, the log files, and the snapinfo files.
- Do not add the cluster management IP address to the transport protocol settings in SDW because doing so might cause failures.

### Long-Term Archiving to Tape

There are two ways to archive SME backups to tape for long-term storage. You can use an NDMP-based backup to copy the LUN in the Snapshot copy that was created by SME to tape. Or you can mount the LUNs in the Snapshot copy created by the SME backup and then stream the LUNs to tape.

## Best Practices

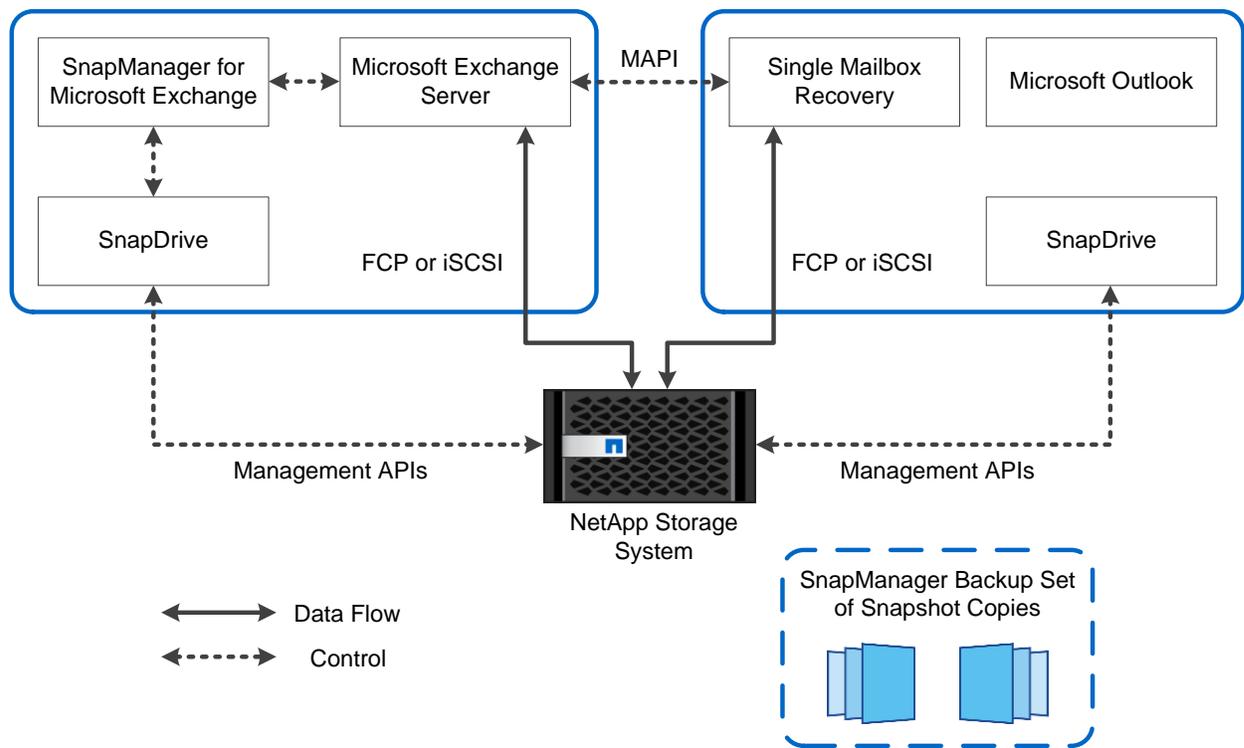
- Use the business requirements established by the Microsoft Exchange stakeholders to help determine the number of Snapshot copies to retain online.
- Use SME-native integration to archive SME backup sets from primary to secondary storage.
- License the controller with FlexClone so that there are no busy Snapshot copies if you mount LUNs in the Snapshot copy created by the SME backup to archive the SME backup to tape.
- When performing an up-to-the-minute restore, restore from the most recent backup because fewer logs must be applied to the recovered database. If the most recent backup has not been verified, the administrator can either verify prior to recovery (longer recovery) or optionally override SME verification and recover from an unverified backup (quickest recovery).

### 9.13 Single Mailbox and Item-Level Recovery

The Single Mailbox Recovery (SMBR) application (versions 7.1 and 7.2) integrates with the NetApp storage solution for Microsoft Exchange Server 2016/2013 data to retrieve individual messages, mailboxes, and attachments without disrupting server availability. SMBR can be used to rapidly locate and restore items, either directly to an existing mailbox on the primary Microsoft Exchange Server or to an offline Microsoft Outlook Personal Storage Table file. SMBR locates and restores Exchange mail items without the assistance of the Microsoft Exchange Server, thereby eliminating the need for a separate recovery server. SMBR versions 7.1 and 7.2 can also be used to recover public folder mailboxes, including organizational forms. Figure 5 shows the integration of SMBR with NetApp storage.

**Note:** SMBR can be launched through SnapManager for Exchange by using the Run SMBR option in the Action pane. SMBR can then locate and then restore items at any level of granularity directly to an existing mailbox on the Microsoft Exchange Server.

Figure 5) Integration of SMBR with NetApp storage solution.



#### Best Practices

- Install and configure Microsoft Outlook 32-bit version.
- Dedicate SMBR on a server with the following configuration:
  - A dedicated computer with at least a 2GB processor and 4GB of RAM running Windows Server 2008 or Windows Server 2012 with the latest service packs installed
  - A high-speed drive to access the archived data
- Operate the SMBR application directly or through RDP (remote desktop) instead of using terminal services or other remote access applications.
- Verify that the log files for all archived information stores (EDB and log files) are available to SMBR during operation.

## 9.14 Single Mailbox Recovery 7.1 and 7.2 Administrative Server

The NetApp Single Mailbox Recovery 7.1 and 7.2 Administrative Server (SMAS) application is a framework that can centralize services for multiple clients. The application provides both client and server support for NetApp SMBR 7.1 and 7.2 and NetApp SMBR extract wizard users. SMAS assigns recovery access permissions for mailboxes, centralizes administration of certain application settings, and provides log-file auditing services for SMBR and SMBR extract wizard clients.

When SMBR is launched, it attempts to connect to SMAS if SMAS is activated. SMAS is located automatically by using either a service connection point in Microsoft Active Directory or a server whose details have been manually provided.

SMBR is designed to run from Windows Server and uses native Microsoft Messaging Application Program Interface (MAPI) protocols to communicate with Microsoft Exchange Server. For MAPI to initialize properly, Microsoft Outlook must be installed and configured on the SMBR host server before the host server is connected to the separate Microsoft Exchange Server.

SMBR and SMAS are not directly integrated with SnapDrive, which connects and disconnects to and from SnapManager to implement Snapshot copies of the Microsoft Exchange LUN. Therefore, connecting or disconnecting LUN Snapshot copies is a manual process that must be carried out by an administrator on either the SMAS server or another server with SnapDrive installed. However, since both SnapDrive and SMBR have command-line interfaces, it is possible for a customer to script actions by using both products.

**Note:** In SMBR, access control is applied to mailboxes through the SMBR Administrative Server (SMBR-AS). NetApp strongly recommends that SMBR-AS be used along with SMBR.

SMBR-AS is an add-on option and must be activated separately by updating the `license.ini` file.

## 9.15 Troubleshooting

SnapManager for Exchange reports list the step-by-step details of every SME operation that is performed, its final status, and any error messages that are encountered during the operation. The SME Report Directory provides subfolders that group the reports for each operation type.

The following troubleshooting steps can be followed to gather additional information:

1. Enable debug logging on all nodes.
2. Identify which operation on which node failed, based on the SME operation sequence.
3. Go to the node with the failure and find the backup report and debug log under `\Backup\Server_name\`, and `\Debug\Server_name\`.
4. Use the server-level backup report and debug log to find the root cause of the problem.

## 10 High Availability

This section describes Microsoft Exchange 2016/2013 DAG deployment scenarios.

### 10.1 Single-Site Scenario

Deploying a two-node DAG with a minimum of two copies of each mailbox database in a single site is best suited for companies that want to achieve server-level and application-level redundancy. In this situation, deploying a two-node DAG using RAID DP provides not only server-level and application-level redundancy, but also protection against double disk failure.

Adding SnapManager for Exchange in a single-site scenario enables point-in-time restores without the added capacity requirements and complexity of a DAG copy. The reseed functionality in SME 7.1 and

later allows the database copies to be in a healthy state and reduces the RTO for failed databases, enabling resiliency at all times.

## 10.2 Multisite Scenario

Extending a DAG across multiple data centers provides high availability for servers and storage components and adds site resiliency. When planning a multisite scenario, NetApp recommends having at least three mailbox servers as well as three copies of each mailbox database: two in the primary site and one in the secondary site. Adding at least two copies in both primary and secondary sites provides not only site resiliency but also high availability in each site. Using the reseed functionality in SME 7.1 and later allows the database copies to be in a healthy state and reduces the RTO for failed databases, enabling resiliency at all times.

For additional information on DAG layout planning, refer to the Microsoft TechNet article [Database Availability Groups](#).

When designing the storage layout and data protection for a DAG scenario, use the following design considerations and best practices.

### Deployment Best Practice

In a multisite scenario, it is a best practice to deploy at least three mailbox servers and three copies of each mailbox database, two in the primary site and one in the secondary site. Adding at least two copies in both the primary and secondary sites provides not only site resiliency but also high availability in each site.

### Storage Design Best Practices

- Design identical storage for active and passive copies of the mailboxes in terms of capacity and performance
- Provision the active and passive LUNs identically with regard to path, capacity, and performance
- Place flexible volumes for active and passive databases onto separate aggregates that are connected to separate SVMs. If a single aggregate is lost, only the database copies on that aggregate are affected.

### Volume Separation Best Practice

Place active and passive copies of the database into separate volumes.

### Backup Best Practices

- Perform a SnapManager for Exchange full backup on one copy of the database and a copy-only backup on the rest of the database copies.
- Verification of database backups is not required if Microsoft Exchange 2016/2013 is in a DAG configuration with at least two copies of the databases, with Microsoft Exchange background database maintenance enabled.
- Verification of database backups and transaction log backups is required if Microsoft Exchange 2016/2013 is in a standalone (non-DAG) configuration.
- In Microsoft Exchange 2016/2013 standalone environments that use SnapMirror, configure database backup and transaction log backup verification to occur on the SnapMirror destination storage.

# 11 Sizing and Storage Layout for Microsoft Exchange Server 2016/2013

This section describes the sizing and storage layout recommendations for Microsoft Exchange Server 2016/2013.

## 11.1 Aggregate Recommendations

Fewer, larger aggregates maximize performance. However, such a configuration might not meet the data availability requirements set forth in the SLA agreement. In Microsoft Exchange Server 2016/2013 environments with multiple database copies, Microsoft no longer requires separating database and transaction log files to separate sets of disks. This means that database and transaction log volumes can be placed in the same aggregate. Each database copy of the same database must be placed in a separate aggregate that is built using separate disk spindles.

### Best Practices

- For optimal storage performance, NetApp recommends having at least 10% free space available in an aggregate hosting Microsoft Exchange data.
- On controllers that are dedicated to Microsoft Exchange deployment, set the global `waf1.optimize_write_once` flag to off to optimize random workloads. The flag must be set before Microsoft Exchange aggregates are created. If Microsoft Exchange aggregates are already present, the `reallocate -A` command must be run on each Microsoft Exchange aggregate. This is a time-consuming process that can affect performance during the reallocation scan.
- Set `waf1_max_write_alloc_blocks` to 256 on the aggregates to increase the sequential read performance chain length.
- Configure two or more RAID groups per aggregate. All RAID groups in an aggregate should have a similar number of disks. The recommended range for RAID group size is between 12 and 20.

## 11.2 Storage Virtual Machine Recommendations

The SVM is the secure logical storage partition through which data is accessed in clustered Data ONTAP. A cluster serves data through at least one and possibly multiple SVMs. An SVM is a logical abstraction that represents a set of physical resources of the cluster. Data volumes and LIFs are created and assigned to an SVM and can reside on any node in the cluster to which the SVM has been given access. Each SVM operates as a separate and distinct entity with its own security domain.

The benefit of creating and dedicating an SVM for Microsoft Exchange workloads on the cluster is that it enables the delegation of data management to the Microsoft Exchange team that is directly responsible for the data being stored. Microsoft Exchange administrators can be given the autonomy to control the datasets belonging to Microsoft Exchange Server, while not having administrative rights to SVMs hosting other application workloads or responsibility for overall cluster administration. When a Microsoft Exchange workload is split into isolated SVMs, QoS policies can be put into place to provide performance isolation at the SVM or workload level.

### Best Practices

- Place Microsoft Exchange data on dedicated SVMs.
- Make sure that each SVM used for Microsoft Exchange workloads has a management interface in addition to interfaces that serve data by using block protocols.
- Distribute the database copies onto separate SVMs to provide higher availability, performance, and resiliency.

## 11.3 SAN LIF Types and Design

There are multiple LIF types, and, from a functional perspective, LIFs can be organized in the following way:

- **Cluster and node management LIFs.** Assist in storage cluster management.
- **SVM management LIFs.** Allow storage virtual machines to perform functions such as snapshot creation or volume resizing.
- **Data LIFs.** Carry FC, iSCSI or CIFS data.

The [official Data ONTAP documentation](#) has more complete information on this topic.

LIF design in a SAN environment for Microsoft Exchange is relatively simple for one primary reason: multipathing. All modern SAN implementations allow a client to access data over multiple network paths and select the best path or paths for access. As a result, with respect to LIF design, performance is simpler to address because SAN clients will automatically load balance I/O across the best available paths.

If a path becomes unavailable, the client will automatically select a different path, and the resulting simplicity of design makes SAN LIFs generally more manageable. This does not mean that a SAN environment is generally more easily managed, because there are many other aspects of SAN storage that are much more complicated than NFS. It simply means that SAN LIF design is easier.

Bandwidth is the most important consideration when evaluating LIF performance in a SAN environment. For example, a four-node Data ONTAP cluster with two 16GB FC ports per node allows up to 32GB of bandwidth from each node. I/O is automatically balanced between ports, and all I/O is directed to the most optimal path. SAN LIFs do not fail over. If a SAN LIF fails, the client's multipathing capability will detect the loss of a path and select a different LIF. SAN LIFs must be created on the LUN-owning node and the owning node's HA partner.

### Best Practices

- Make sure that NetApp storage controllers do not receive Ethernet flow control packets. You can generally do so by setting the switch ports into which the controller is attached, but some switch hardware has limitations that might require client-side changes instead.
- A NIC on a NetApp system should not receive flow-control requests. The method to achieve this situation varies based on the network switch manufacturer. In most cases, flow control on an Ethernet switch can be set to Receive Desired or Receive On, which means that a flow control request is not forwarded to the storage controller.
- Jumbo frames are desirable but not required with 1GB Ethernet.
- Jumbo frames are required for maximum performance with 10GB Ethernet.
- NetApp recommends implementing jumbo frames when possible, both to realize any potential performance benefits and to future-proof the solution.

## 11.4 Volume Planning and Layout

Data ONTAP enables the creation of flexible volumes for managing data without the need to assign physical disks to the volumes. Instead, the flexible volumes enjoy performance benefits from a larger pool of physical disks called an aggregate. This process results in the following additional benefits for Microsoft Exchange Server 2016/2013 environments:

- A large number of volumes can be created, all with independent Snapshot copy schedules and SnapMirror policies.
- All volumes can be managed independently while receiving the maximum I/O benefit of a much larger pool of disks.

## Best Practices

- NetApp recommends separating database and transaction logs from different servers into separate volumes to prevent a potential busy Snapshot copy problem. Utilizing separate volumes for each server reduces complexity because there is no concern for Snapshot copy schedules overlapping different servers.
- NetApp recommends having at least 10% free space available in a volume hosting Microsoft Exchange data.
- NetApp recommends placing each database in a separate volume with copies of the same database isolated in separate aggregates.

Volume sizing is different for transaction logs and database volumes. Transaction log sizing involves calculating the size of the transaction log LUN(s) in the volume, adding space for the Snapshot retention length, and including 10% free space in the volume.

### Transaction Log Volume

- The transaction log LUN size is 71GB
- 833 users send 100 messages per day.
- 83,300 messages \* 75KB = 5.95GB per day
- A snapshot retention time of 7 days + 3 days of fault tolerance = 10 days
- 10 days \* 5.95GB = 59.5GB
- The log volume = (59.5GB + 71GB / (1-.1)) or 145GB
- 10 days of Snapshot copies + transaction log LUN plus 10% free space in the volume

### Database Volume

- The database LUN size = 2691GB (1826GB database)
- Sizing the database LUN for quota includes the maximum mailbox size, deleted items in the dumpster, calendar, three days of incoming mail, and whitespace in the database.
- A Snapshot retention time of 7 days + 3 days of fault tolerance = 10 days
- The daily change rate is 5% \* 10 days (common rates are in the 2% to 8% range).
- 2691 + ( 1826 \* 50%) = 3604GB

### LUN Planning and Layout

A database and its corresponding transaction log must be placed on separate LUNs for SnapManager for Exchange. In environments with high LUN counts, transaction logs for multiple mailbox databases can be consolidated on a single LUN. NetApp recommends limiting the number of transaction log streams per LUN from 5 to 10.

## Best Practices

- When creating LUNs, use volume mount points. There are a finite number of drive letters, and each database path in a DAG must be the same on every server that has a copy of that database.
- Place each database on a separate LUN in a separate volume.
- Use larger databases. Microsoft supports up to 16TB databases. NetApp recommends using fewer databases (2TB at a minimum).

Do not create mount points for additional LUNs on another LUN that holds a Microsoft Exchange Server 2016/2013 database or create any files or folders in the root folder where the mount points are created. If you have to complete a restore of a database residing on a LUN with volume mount points, the restore operation removes any mount points that were created after the backup, disrupting access to the data on

the mounted volumes referenced by these volume mount points. SnapManager for Exchange does not allow users to store files or to back up databases on an NTFS volume that has mount points.

**Note:** Do not place databases or transaction logs on a mount point root volume.

It is a NetApp best practice to place the transaction logs and database files on separate LUNs. These calculations are for the primary active database and its corresponding transaction log files. Each additional copy of the database requires a multiple of the sizing. The same calculations can be used to estimate the size of the archive database and its corresponding transaction log files.

## Database

The database LUN houses 5% free disk space, the database itself, and the content index files.

- The MBXSize is the MBXLimit plus the dumpster.
- The MBXLimit is the stated maximum mailbox size; in our example case, this is 2GB.
- Calculate the space consumed in the dumpster, which also includes space consumed by enabling both single-item recovery and calendar version storage.
  - Single item recovery =  $\text{MBXLimit} * 0.012$  (1.2%)
  - Calendar version storage =  $\text{MBXLimit} * 0.03$  (3%)

**Example:** Assuming a 2GB mailbox and a default 14-day retention:

- Dumpster =  $\text{SingleItemRecovery} + \text{CalendarVersionStore} + (\#\text{Messages} * \text{MessageSize} * (\text{DeletedItemRetention} + 1 [\text{today}]))$
- Dumpster =  $24.6 + 61.4 + (((100 * 75\text{KB} * (14 + 1)) / 1024)$   
 $= 24.6 + 61.4 + 109.8\text{MB} = 195.8\text{MB}$
- The database size is the MBXSize multiplied by the number of users.
- The DB size + overhead adds 0% to the database size (compared to 20% in Microsoft Exchange 2010).
- The DB LUN is calculated by adding the DB size + overhead to the content index, while padding in 5% free disk space.

**Example:**

- DB size of 1826GB
- DB size + overhead = 1826GB
- ContentIndex =  $1826 * 40\% = 365.2\text{GB}$  (versus 10% in Microsoft Exchange 2010)
- $(1826 + 730) / (1 - .05) = 2691\text{GB}$

## Transaction Logs

The transaction logs are 1MB in size. They must include transaction logs generated by the users from move mailbox requests and the backup fault tolerance window. By default, Microsoft assumes a three-day backup fault tolerance window.

- User logs (calnumusertlogs):  $\text{LogGen} * \text{users} * \text{Datagrowth}$   
 $\text{LogGen} = 10$  for each 50 messages per day in a user profile  
Users: An example 100 messages per day 20,000-users configuration would be:  
 $20 * 20,000 = 400,000$  logs
- Move mailbox (logdiskspacereqmove) =  $(\text{users} * 1\%) * (\text{MBXSize} / 1024)$   
 $(20,000 * 1\%) * (2244 / 1024)$   
 $200 * 2.191$   
438GB

- Log backup (logdiskspacereqbackup)  
UserLogs \* BackupFailTol (default 3)  
 $400,000 * 3 = 1,200,000$  logs
- Total log disk space (totlogdiskspace)  
Log backup + move mailbox  
 $(1,200,000 / 1024) + 438\text{GB}$   
1609GB
- Add 5% free disk space  
 $1609\text{GB} / (1-.05) = 1694\text{GB}$
- Divide by the number of databases, in this case, 24  
 $1694\text{GB} / 24 = 71\text{GB}$  per transaction log LUN

**Note:** Most of this space is because of move mailbox and large mailbox sizes.

## 11.5 Capacity Planning

A properly sized Microsoft Exchange environment meets or exceeds the customer SLA. To properly size an environment, information from the customer environment is collected, and tools are used to convert that information into a physical storage recommendation.

Two primary tools should be used when planning a Microsoft Exchange environment for a customer:

- **The Microsoft Exchange 2013 Server Role Requirements Calculator.** The latest version, V7.8, was introduced to support Exchange Server 2016. You can select the appropriate version using the dropdown option in the Input tab.
- **The NetApp Exchange Sizing and Performance Modeling (SPM) tool.** Work with your local NetApp partner or your NetApp representative for proper sizing.

The sizing information provided by these tools is an important component for planning a Microsoft Exchange environment, and the information provides a framework for database layout and LUN requirements. The Microsoft storage calculator cannot accurately make recommendations on proprietary storage technology because the storage design largely depends on the type of storage array being utilized. When sizing Microsoft Exchange Server deployments using NetApp storage, it is important to use the NetApp SPM tool with the data from the Microsoft Exchange 2016/2013 Mailbox Server Role Requirements Calculator.

### Best Practice

Consult a local NetApp Exchange expert or your NetApp partner to assist in accurately sizing Microsoft Exchange Server 2016/2013. Use the NetApp SPM tool to size all Microsoft Exchange server deployments using NetApp storage.

## 12 Performance

Accurately sizing NetApp storage controllers for Microsoft Exchange workloads is essential for good Microsoft Exchange performance and so that Microsoft Exchange service levels are met. Consult a local NetApp Exchange expert to provide accurate performance sizing. These experts can also provide capacity requirements in the previous section and layout for Microsoft Exchange environments using the Microsoft Exchange 2016/2013 storage calculator as the input.

## 12.1 NetApp Flash Cache Intelligent Data Caching

Flash Cache has been the leading flash-based technology used in Microsoft Exchange deployments for one simple reason. Most databases are limited by random read latency. Using Flash Cache is the simplest method to accelerate random read performance.

Flash Cache is a read cache that can be installed on certain models of NetApp storage controllers. Flash Cache does, however, have the limitation that it is tied to the particular node that hosts the PCIe card containing the flash memory. As spindle sizes increase, customers deploy fewer spindles. This situation increases the risk that controller failure will result in a period of performance degradation while the Flash Cache card on the takeover node warms up. Until warmup is complete, there can be significant extra I/O on the spinning media. For this reason, NetApp Flash Pool intelligent data caching is the preferred flash technology because the flash layer follows the spinning media during takeover. No warm up time is required, and the cache does not go cold.

Table 6 lists the setting options when using Flash Cache.

Table 6) Flash Cache options.

Options	Setting
<code>flexscale.enable</code>	On
<code>flexscale.lopri_blocks</code>	Off
<code>flexscale.normal_data_blocks</code>	On
<code>flexscale.pcs_high_res</code>	Off
<code>flexscale.pcs_size</code>	0GB
<code>flexscale.readahead_blocks</code>	Off
<code>flexscale.rewarm</code>	On

**Note:** NetApp recommends using the same value for local and partner for all options.

## 12.2 NetApp Flash Pool Intelligent Data Caching

Flash Pool combines solid-state disks and traditional hard disk drives into a single aggregate. Flash Pool improves the latency of random reads. Flash Pool is also a cost-saving technology because a small Flash Pool allocation can replace a large number of spinning drives.

Write caching, or more specifically overwrite caching, is a further benefit of Flash Pool that is not realized with Flash Cache. Using Flash Pool for write caching does not directly affect write performance because writes commit to NVRAM/NVMEM. From a latency perspective, I/O is complete when data is journaled into NVRAM/NVMEM.

The type of media on which an inbound write is subsequently stored does not affect performance by itself. There can, however, be an indirect benefit to write performance if Flash Pool write caching reduces pressure on spinning media, leading to a general improvement in I/O performance for the entire array. For Flash Pool, however, persistent caching occurs at the aggregate level. Therefore, hot blocks are written to the SSDs and are not affected by a storage failover event or a giveback event.

NetApp makes the following recommendations when using Flash Pool:

- Use the default Flash Pool policy, which includes both random read and random write caching
- With Data ONTAP 8.3 and later releases, Flash Pool is also an option with Microsoft Exchange. This option adds the benefit of inline compression because Flash Pool is optimized for SSDs.

For more information, see [TR-4070: Flash Pool Design and Implementation Guide](#).

#### Best Practice

It is a NetApp best practice to use one of the flash options for Microsoft Exchange workloads, particularly if SATA disks are used, with appropriate sizing.

### 12.3 SATA Performance Considerations

SATA-based deployments of Microsoft Exchange must take into account that SATA drives have a lower I/O profile than SAS and FC disks. The I/O profile of a 7,200-RPM SATA drive is approximately 45 to 55 IOPS at a 20ms response time.

Microsoft Exchange 2016/2013 utilizes background database maintenance (BDM) to maintain database consistency. BDM applies a per-database performance tax on the storage system that must be taken into account when sizing the storage for Microsoft Exchange. Having fewer, larger databases in the Microsoft Exchange database design helps reduce the amount of background database maintenance I/O, which in some cases can exceed the transactional I/O generated by users. This is typically seen in designs with a large number of small databases. Such configurations invoke more 256KB sequential-read I/O that is then added to the 32KB random I/O. This addition produces a large load as an effect of I/O coalescing on the host side.

In Microsoft Exchange 2016/2013, the BDM process consumes considerably less bandwidth and produces fewer IOPS compared to that of Microsoft Exchange 2010. BDM is now throttled back from 5MB per sec per copy to ~1MB per sec per copy in production and 7.5MB per sec per copy to ~2.25MB per sec per copy in Jetstress.

To help improve the storage efficiency and read I/O performance and latency of SATA-based deployments, either Flash Cache or Flash Pool should be used. Both options require fewer SATA disks in SATA-based deployments because a portion of the Microsoft Exchange database working set is cached, which greatly reduces the amount of read I/O on the SATA disk. NetApp recommends using Flash Cache or Flash Pool along with SATA for deployments that exceed 1,000 mailboxes or when SATA-based designs are bound by performance instead of capacity.

#### Best Practice

NetApp recommends using flash options when placing Microsoft Exchange Server 2016/2013 database files on SATA physical disk drives.

### 12.4 Database Sizing Considerations

Using a smaller number of larger databases can help reduce the amount of background database maintenance I/O as well as the complexity of the storage design. NetApp recommends using a database size of at least 2TB with at least two copies in a DAG, or 200GB for non-DAG databases. Two terabytes is a practical database size that can be restored in minutes with SnapManager for Exchange. Microsoft recommends 2TB as a maximum size, but Microsoft supports up to 16TB for databases on both the Standard and Enterprise Server editions.

## Best Practice

NetApp recommends using a database size of at least 2TB. With the protection of clustered storage controllers, RAID DP, and Snapshot backups, many customers have deployed databases larger than 2TB, which significantly reduces the database maintenance I/O.

## 12.5 Aggregate Sizing and Configuration Considerations

Aggregates are sized for performance and storage capacity in storage designs that support Microsoft Exchange workloads as well as maintain data protection for Microsoft Exchange data.

As mentioned in the Best Practices table in section 11.4, “Volume Planning and Layout,” Microsoft Exchange databases and transaction logs can be placed on the same aggregate. There is a marginal benefit to locating transaction logs and databases on separate aggregates. However, putting DAG database copies on separate aggregates and/or controllers enables at least one copy of the Microsoft Exchange data to exist if an aggregate is lost because of a catastrophic failure. Placing database copies on separate aggregates also helps isolate background database maintenance I/O to the aggregates in which the database copies are located.

If Microsoft Exchange is virtualized, place the Microsoft Exchange virtual machines on a separate aggregate from the Microsoft Exchange data. Determine that there are at least two spare disks per controller and that the Data ONTAP option `disk.maint_center.enable` is enabled. It is on by default but requires two hot spares. Maintenance Center is a feature in Data ONTAP that can prefail a disk if the disk does not pass a certain number of diagnostic tests.

The aggregates for Microsoft Exchange should be configured for RAID DP. This arrangement enables maximum data protection for the Microsoft Exchange data so that an aggregate can survive a double disk failure in any RAID group of that aggregate.

The RAID group size of the aggregate affects the level of data protection, speed of recovery, and available data storage space. Configuring an optimum RAID group size for an aggregate requires a trade-off of factors. Adding more data disks to a RAID group increases the striping of data across those disks, which typically improves I/O performance. In addition, a smaller percentage of disks is used for parity rather than data. However, with more disks in a RAID group, there is a greater risk that one of the disks might fail. For that reason, NetApp recommends using the default RAID group size when the Microsoft Exchange aggregate is created, because doing so balances storage efficiency and performance.

Microsoft Exchange workloads can run effectively on both 32-bit and 64-bit aggregates if the aggregates and controller heads are properly sized. NetApp recommends using 64-bit aggregates to support a Microsoft Exchange workload only in configurations that are supported in the NetApp Exchange sizing tool. Doing so enables the storage to be properly sized for the anticipated Microsoft Exchange workload. NetApp recommends consulting a local NetApp Exchange expert for accurate performance sizing for Microsoft Exchange environments.

## 12.6 Volume Configuration Considerations

NetApp recommends setting the volume option `read_realloc` on each database volume. This setting is particularly helpful in environments with many databases and the corresponding sequential read because of the background database maintenance.

## 13 Storage QoS and Nondisruptive Operations

This section describes the issues surrounding storage quality of service (QoS) and nondisruptive operations.

## 13.1 Storage Quality of Service

Data ONTAP 8.2 introduces QoS, which can help you manage the risks that accompany meeting performance objectives for expensive Microsoft Exchange workloads. QoS allows you to limit the amount of I/O sent to an SVM, a flexible volume, a LUN, or a file. An entire SVM, or a group of volumes or LUNs within an SVM that holds Microsoft Exchange data, can be dynamically assigned to a policy group, which specifies a throughput limit (defined in terms of IOPS or MB/sec). This limit can be used to reactively or proactively throttle non-Microsoft Exchange workloads and prevent them from affecting Microsoft Exchange workloads. In addition, rogue applications that consume too many system resources can also be rapidly identified and corralled by setting IOPS or MB/sec limits, enabling continuous functioning of Microsoft Exchange Server data.

From the performance perspective of a Microsoft Exchange workload, maximum IOPS and throughput levels can be set per SVM by using QoS policy groups. QoS can be used to set service-level objectives defining “not to exceed” performance at the file, LUN, volume, or SVM level. Doing so is particularly useful for passive database copies in a remote data center. The QoS policies can be removed during a data center failover. When the remote data center becomes active, latencies must be low.

Note the following requirements when assigning a LUN to a policy group:

- The LUN must be contained by the SVM to which the policy group belongs.
- If you assign a LUN to a policy group, then you cannot assign the LUNs containing a volume or SVM to a policy group.

### Best Practice

Apply QoS on SVMs or volumes that do not host critical Microsoft Exchange databases and transaction logs to define isolation boundaries between those workloads. Doing so also prevents the boundaries from interfering with the higher-priority Microsoft Exchange traffic.

For more details, see to the white paper [Clustered Data ONTAP Quality of Service](#).

## 13.2 Nondisruptive Operations

Clustered Data ONTAP is designed to enable users to manage, upgrade, and service their storage infrastructure without disruption. A NetApp clustered Data ONTAP system can nondisruptively fail over network connections while a server is executing live read and write I/O to the volume. Moving the volume to a different aggregate and failing back on demand can be done without interrupting data access and without affecting performance for Microsoft Exchange workloads. Make sure that the SME backups are not triggered when nondisruptive operations are already in progress.

The main consideration when moving a Microsoft Exchange data volume is the host-side timeout value. The cutover window defined for a volume move with a Microsoft Exchange data volume should not exceed the expected timeout value on the host side. NetApp recommends that the cutover window not exceed 120 seconds for volumes in Microsoft Exchange workload solutions. When a Microsoft Exchange data volume is moved, ALUA is used for optimized access to the volume. For a more comprehensive overview and for best practices, see [TR-4080: Best Practices for Scalable SAN in Clustered Data ONTAP 8.3.1](#).

## 14 NetApp MetroCluster Software

NetApp MetroCluster™ high-availability (HA) and disaster recovery storage software provides both local-failover and site-failover capabilities with clustered Data ONTAP by providing RAID-level replication of data between sites. For Microsoft Exchange environments in which Microsoft Exchange DAG replication is leveraged, MetroCluster provides four copies of the Microsoft Exchange data. NetApp recommends

considering the multiplication or duplication of copies of the Microsoft Exchange databases (both DAG copies and storage-level copies). MetroCluster is perfectly suited for standalone Microsoft Exchange Deployments over DAG.

For more information, see [TR-3548: MetroCluster Version 8.2.1 Best Practices for Implementation](#).

## 15 Virtualization

### 15.1 Microsoft Support for Exchange 2013 in Virtualized Environments

The documentation for supporting Exchange 2016/2013 in virtualized environments can be found in the Microsoft TechNet article [Exchange 2013 Virtualization and Exchange 2016 Virtualization](#).

Here is a high-level list of some important considerations:

- All Microsoft Exchange 2016/2013 server roles are supported in a virtual machine
- Microsoft Exchange Server 2016/2013 virtual machines (including Microsoft Exchange Mailbox virtual machines that are part of a DAG) can be combined with host-based failover clustering and migration technology. This benefit is true as long as the virtual machines are configured so that they do not save and restore state on disk when moved or taken offline
- All storage used by a Microsoft Exchange guest machine for storing Microsoft Exchange data must be block-level storage because Microsoft Exchange 2016/2013 does not support using network-attached storage (NAS) volumes. NAS storage that is presented to the guest as block-level storage through the hypervisor is also not supported.
- Microsoft does not support the use of virtual machine Snapshot copies of Microsoft Exchange virtual machines.

**Note:** Currently, SME supports raw device mapping only through FC/iSCSI initiators or through LUNs attached directly to the guest with iSCSI initiators.

## 16 Disaster Recovery

E-mails are an integral part of modern businesses. Achieving the required level of availability and recovery of Microsoft Exchange from failure is a big challenge that is critical to organizations.

For information about how to design a disaster-recovery model for Microsoft Exchange Server by using the comprehensive NetApp suite of hardware and software solutions, refer to [TR-4332: Disaster Recovery Solution for Microsoft Exchange on Clustered Data ONTAP](#).

**Note:** The same approach can be used for Microsoft Exchange Server 2016.

The key topics highlighted in the solution are the local and secondary site-availability options, encompassing the disaster-recovery options available with Microsoft Exchange and NetApp storage.

## 17 Summary

Microsoft Exchange Server 2016/2013 is not a one-size-fits-all application. Multiple configuration options are available to suit the needs of almost any customer. NetApp storage appliances and data management software are built in a similar fashion, providing users with the flexibility to manage Microsoft Exchange data in a manner that best meets their business requirements. With high-performance, easy-to-manage storage appliances and robust software offerings, NetApp offers the flexible storage and data management solutions to support Microsoft Exchange Server 2016/2013 enterprise messaging systems.

The best practices and recommendations set forth in this guide are also not a one-size-fits-all solution. This document contains a collection of best practices and recommendations that provide a guideline to

plan, deploy, and manage Microsoft Exchange data. This guideline enables a highly available, easy-to-manage Microsoft Exchange environment that meets SLAs. Consult with a local NetApp Exchange expert when planning and deploying Microsoft Exchange environments onto NetApp storage. NetApp Exchange experts can quickly identify the needs and demands of any Microsoft Exchange environment and adjust the storage solution accordingly.

## References

- Exchange 2016 System Requirements  
[https://technet.microsoft.com/en-us/library/aa996719\(v=exchg.160\).aspx](https://technet.microsoft.com/en-us/library/aa996719(v=exchg.160).aspx)
- Exchange 2013 System Requirements  
[http://technet.microsoft.com/en-us/library/aa996719\(v=exchg.150\).aspx](http://technet.microsoft.com/en-us/library/aa996719(v=exchg.150).aspx)
- Logical Storage Management Guide  
<http://support.netapp.com/documentation/docweb/index.html?productID=61651>
- Volume Shadow Copy Service Overview  
[http://msdn.microsoft.com/en-us/library/aa384649\(v=VS.85\).aspx](http://msdn.microsoft.com/en-us/library/aa384649(v=VS.85).aspx)
- Exchange Storage Configurations Options  
[http://technet.microsoft.com/en-us/library/ee832792\(v=exchg.150\).aspx](http://technet.microsoft.com/en-us/library/ee832792(v=exchg.150).aspx)
- Exchange 2016/2013 Mailbox Server Role Requirements Calculator  
<http://gallery.technet.microsoft.com/office/Exchange-2013-Server-Role-f8a61780>
- Data ONTAP Documentation  
<http://now.netapp.com/NOW/knowledge/docs/docs.cgi>
- Exchange 2013 High Availability and Site Resiliency  
[http://technet.microsoft.com/en-us/library/dd638137\(v=exchg.150\).aspx](http://technet.microsoft.com/en-us/library/dd638137(v=exchg.150).aspx)
- Backup, Restore, and Disaster Recovery  
[http://technet.microsoft.com/en-us/library/dd876874\(v=exchg.150\).aspx#SerRec](http://technet.microsoft.com/en-us/library/dd876874(v=exchg.150).aspx#SerRec)
- Exchange 2016 Architecture  
[https://technet.microsoft.com/en-us/library/jj150491\(v=exchg.160\).aspx](https://technet.microsoft.com/en-us/library/jj150491(v=exchg.160).aspx)
- TR-4332: Disaster Recovery Solution for Microsoft Exchange on Clustered Data ONTAP  
<http://www.netapp.com/us/system/pdf-reader.aspx?pdfuri=tcm:10-129197-16&m=TR.4332.pdf>

## Version History

Version	Date	Document Version History
Version 2.0	March 2016	Updated with Exchange 2016, Flash Cache, Flash Pool, and AFF details
Version 1.1	November 2014	Updated with MetroCluster and SME 7.1 features
Version 1.0	September 2013	Initial release

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

## Copyright Information

Copyright © 1994–2016 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.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

## Trademark Information

NetApp, the NetApp logo, Go Further, Faster, AltaVault, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Fitness, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANshare, SANtricity, SecureShare, Simplicity, Simulate ONTAP, SnapCenter, SnapCopy, Snap Creator, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, SolidFire, StorageGRID, Tech OnTap, Unbound Cloud, WAFL, and other names are trademarks or registered trademarks of NetApp Inc., in the United States and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the web at <http://www.netapp.com/us/legal/netapptmlist.aspx>.