Technical Report

Microsoft Exchange Server 2013 and SnapManager for Exchange

Best Practices Guide for Data ONTAP Operating in 7-Mode

Niyaz Mohamed, Robert Quimbey, NetApp

September 2013 | TR-4224

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1 Executive Summary

Many organizations have come to rely on Microsoft® Exchange Server to facilitate critical business e-mail communication processes, group scheduling, and calendaring on a 24/7 basis. System failures might result in unacceptable operational and financial losses. Due to 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 Microsoft Exchange Server databases growing rapidly, 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 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, as well as scale to accommodate future needs while reducing cost and complexity.

NetApp® SnapManager® 7.0 for Microsoft Exchange (SME) is available for Microsoft Exchange Server 2013. SME is tightly integrated with Microsoft Exchange Server, which allows for consistent online backups of your 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 the new Microsoft Exchange Server 2013 data replication features. SME works with database availability group (DAG) databases on both standalone servers and servers participating in a DAG and provides 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 2013 and using SnapManager 7.0 for Microsoft Exchange with a NetApp storage system running Data ONTAP® 7-Mode and any supporting software. Organizations that want to get the most out of their NetApp storage investment for Microsoft Exchange Server will benefit from putting into practice the recommendations in this report.

1.2 Intended Audience

This paper is a best practice guide for experienced Microsoft Exchange Server administrators who are familiar with the installation and administration of SnapDrive® for Windows®, SnapManager for Exchange, and Data ONTAP.

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

2 Exchange Server 2013 Architecture

Exchange Server 2013 includes the following server roles:

- **Client access servers (CASs).** 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.

2.1 Database Availability Groups

Exchange Server 2013 uses DAG, which utilizes 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.

2.2 In-Place Archiving

An archive mailbox is an additional mailbox associated with a user's primary mailbox. This new mailbox is known as an archive mailbox and 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.

3 Exchange Server 2013 Planning Considerations

3.1 System Requirements

In this section we discuss the system requirements for Microsoft Exchange Server 2013 on NetApp storage systems:

- Windows Server 2008 R2 or Windows Server 2012.
- Minimum and maximum page file size set to physical RAM plus 10MB.
- Memory requirements vary depending on Exchange roles that are installed.
- Disk space depends upon the roles installed; at least 500MB of free space is required on the drive that stores the message queue database.
- Disk partitions formatted as NTFS file systems.

For more detailed requirements refer to the Microsoft TechNet article: Exchange 2013 Storage Configuration Options.

4 NetApp Storage Efficiency Technologies

NetApp's 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 WAFL® file system. Unlike other technologies, NetApp's technologies surrounding its FAS and V-Series product line have storage efficiency built into their core.

Customers who already have other vendors' storage systems and disk shelves can still leverage all the storage-saving features that come with the NetApp FAS system simply by using the NetApp V-Series product line. This is again in alignment with NetApp's philosophy of storage efficiency (helping increase storage use and decrease storage cost) because customers can continue to use their existing third-party storage infrastructure and disk shelves, yet save more by leveraging the storage-efficient technologies inherent to NetApp v-Series controllers.

NetApp offers the following technologies that increase storage efficiency.

4.1 RAID-DP

RAID-DP® technology safeguards against double-disk failure and delivers high performance. RAID-DP is integrated with the WAFL file system so that the dedicated parity drives don't become a performance bottleneck. RAID-DP makes SATA disks an option for your enterprise storage. Exchange administrators can use less-expensive SATA without worrying about data loss and also lower their storage acquisition costs.

Note: SyncMirror® can be used along with RAID-DP to provide a second layer of mirrored protection for a more robust disk protection strategy.

4.2 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 it is with other copy-out technologies. The cost for Snapshot copies is only at the rate of block-level changes, not 100% for each backup as with mirror copies. It can result in savings in storage costs for backup and restore purposes and opens up a number of efficient data management possibilities.

4.3 Thin Provisioning

Thin provisioning, in a shared storage environment, is a method for optimizing utilization of available storage. It relies on on-demand allocation of blocks of data versus the traditional method of allocating all of the blocks up front. This methodology eliminates almost all white space, which helps avoid poor utilization rates. Flexible volumes (FlexVol® 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, it does not dedicate specific blocks out of the NetApp volume for the LUN or for Snapshot copies of the LUN. Instead, it allocates the blocks from the NetApp aggregate when the data is actually written. This 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, it is best to 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 overcommitment is acceptable, if any.

Overcommitment of storage must be carefully considered and managed for mission-critical applications in which even a minimal outage is not tolerable. In such a case, it is best to monitor storage consumption trends to determine how much overcommitment is acceptable, if any.

If the time required to procure new storage is very long, storage overcommitment thresholds should be adjusted accordingly. The overcommitment threshold should alert administrators early enough to allow new storage to be procured and installed.

The potential risk when configuring the Exchange environment for thin provisioning is a LUN going offline when there is not enough space to write further data. Use volume autogrow as mitigation mechanism to safely allow thin provisioning and higher storage utilization.

4.4 Space Guarantee

The space guarantee is the enabler of 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 FlexVol volume is always available from its aggregate. This is the default setting for FlexVol volumes. When the space guarantee is set to "volume," the space is reserved from the aggregate's available space at volume creation time.

When the 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` will fail if the containing aggregate does not have enough available space.

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

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

4.5 Space Reclamation

Space reclamation can be initiated from time to time to recover the unused space in a LUN. Space reclamation for space utilization from a host perspective is coordinated with space consumed on the NetApp storage controllers. SnapDrive Space Reclaimer can be used to remediate between the host and the controller to free up these deleted blocks, thus reducing utilization in the LUN and in Snapshot copies. Storage space can be reclaimed at the storage level using the SnapDrive > Start Space Reclaimer option.

4.6 Fractional Reserve

Fractional reserve is a volume option that determines how much space Data ONTAP will reserve for Snapshot overwrite data for LUNs to be used after all other space in the volume is used. NetApp recommends setting FSR to 0 in Exchange environments.

4.7 Autodelete and Autosize

The autosize volume setting (available in Data ONTAP 7.1 and later) defines whether a volume should automatically grow to avoid filling up to capacity. It is possible to define how quickly the volume should grow with the “-i” option. The default growth increment is 5% of the volume size at creation. It is also possible to define how large the volume is allowed to grow with the “-m” option. If volume autosize is enabled, the default maximum size to grow to is 120% of the original volume size.

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 2013 environments.
- NetApp recommends prioritizing autosize over autodelete because deletions occur at the Data ONTAP level, and it is possible to have a backup set of a transaction log and database where one Snapshot copy has been automatically deleted or orphaned.

Autodelete should be the lowest priority, and 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 built into Data ONTAP. When properly configured, SnapManager will delete Snapshot copies as per the retention. If Data ONTAP autodelete is used, SnapManager will be unable to perform its retention duties, and backup sets could become broken. Snapshot copies should only be deleted using SnapManager, with either the retention or with the delete backup wizard. Snap autodelete needs to be off on ALL volumes managed by any SnapManager product.

In such scenarios, the recommendation is to use autosize; however, it might fail due to space constraints in the aggregate and must be properly monitored using Operations Manager. For volume autosize to work, it is mandatory that the containing aggregate has enough space (at least 1.2 times the volume size).

Both autodelete and autosize work at the volume level, and not on individual LUNs. This means that LUNs will 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 is met. This threshold is called a “trigger” and can be set so that Snapshot copies will be automatically deleted when one of the following conditions is met:

- **Volume.** The volume is nearly full. This is reported in the first line reported for each volume in the `df` command. It should be noted that 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 we recommend 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 will never be filled until both the volume and the snap_reserve areas are full.

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

```
6240b> df
File system              Kbytes          used            avail           capacity      Mounted on
/vol/vol0/               1407415772      12094808        1395320964      1%             /vol/vol0/
/vol/vol0/.snapshot      74074512        455588          73618924        1%       /vol/vol0/.snapshot
```

Best Practice

NetApp recommends using autogrow instead of autodelete. When using autodelete, set the autodelete trigger to volume.

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

- **delete_order.** This option determines whether the oldest or newest Snapshot copies should be deleted first.
- **defer_deleted.** This option allows the user to define a group of Snapshot copies that should first be deleted when no other Snapshot copies are available. It is possible to 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 SnapMirror® and dump operations should be handled. If set to “try,” it will only delete these Snapshot copies if they are not locked. If set to “`disrupt`,” these Snapshot copies will be deleted even if they are locked.

Best Practice

When using SnapMirror products or SnapVault® hardware for replicating Microsoft Exchange Server 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.

4.8 Best Practice Configurations When Using Thin Provisioning for Exchange Server 2013 Environments

There are many ways to configure the NetApp storage appliance for LUN thin provisioning; each has advantages and disadvantages. It should be noted that it is possible to have thinly provisioned volumes and non–thinly provisioned volumes on the same storage system or even the same aggregate. The following are considered to be best practice configurations when using thin provisioning for Microsoft Exchange Server 2013.

Option 1: Volume Guarantee Set to “None”

Volume guarantee	= none
LUN reservation	= enabled
fractional_reserve	= 0%
snap_reserve	= 0%
autodelete	= volume / oldest_first
autosize	= off
try_first	= snap_delete

This configuration has the advantage of the free space in the aggregate being 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 using this configuration 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 only need to manage and monitor 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.

Option 2: Using Autogrow/Autodelete

Volume guarantee	= volume
LUN reservation	= disabled
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 2013 environments. With this configuration the volume size defines or guarantees an amount of space that is only available to LUNs within that volume. The aggregate provides a shared storage pool of available space for all the volumes contained within it. If the LUNs or Snapshot copies require more space than available in the volume, the volumes will automatically grow, taking more space from the containing aggregate. Additionally, the advantage of having the LUN space reservation disabled in that case 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 will remove the Snapshot copies consuming space. This is an ideal setting for many migrations in which Snapshot copy space will be high during the initial mailbox moves, but will taper off in the months and years to come when more space is required within the database to store mail.

Note: Snapshot copies used to create FlexClone® volumes will not be deleted by the `autodelete` option.

Best Practice
NetApp recommends using autogrow for most common deployment configurations.

4.9 Monitoring

When using NetApp efficiency features, 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 OnCommand® Unified Manager Core Package management software that includes Operations Manager is the recommended tool to monitor Exchange volumes for these events and to send notifications to the storage administration team to follow up further with the Exchange administration team. SNMP can also be used to monitor these events.

After a notification for a volume autogrow or Snapshot autodelete event has been received by the storage administration team, the recommended action is for the storage administration team to examine the affected storage controllers and then follow up with the 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 additional 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 SnapManager for Exchange are not being deleted. As Snapshot copies age, they can grow in size and consume more capacity than originally allocated in the volume. A typical cause of SnapManager for Exchange not deleting backups is that SnapManager for Exchange backups are failing. By default, SnapManager for Exchange does not delete Snapshot copies of older SnapManager for Exchange backup sets if the backup fails. Another cause for SnapManager for Exchange not deleting backups is that the SnapManager for Exchange backup retention policies are not being enforced correctly because Snapshot copies were manually removed outside of SnapManager for Exchange on the controller itself.

Monitoring the health of SnapManager for Exchange can be done by monitoring for SnapManager for Exchange event IDs and the enhanced enterprise monitoring functionality in SME 7.0. To monitor the health of SnapManager for Exchange retention external to SnapManager for Exchange, the number of Snapshot copies and SnapInfo directories should be calculated for a specific Microsoft Exchange Server. For example, if the SnapManager for Exchange retention policy for a particular server is 10 backups online (-RetainBackups 10 parameter in the `new-backup` command), there will be 10 SnapManager for Exchange Snapshot copies in each Exchange database volume (with a prefix of exchsnap__) and 20 SnapManager for Exchange Snapshot copies in each Exchange transaction log volume (10 with a prefix of exchsnap__ and 10 with a prefix of eloginfo__). If the SnapManager for Exchange retention policy for a particular server is 10 days of backups online (-RetainDays 10 parameter in the `new-backup` command), there will be SnapManager for Exchange Snapshot copies in the Exchange database and transaction log volumes no older than 10 days. An alternate way of calculating SnapManager for Exchange retention when using the `–RetainDays backup` parameter is to multiply the number of days you keep backups online by the number of backups taken each day. If 1 backup per day is taken, then there would be 10 SnapManager for Exchange Snapshot copies in each Exchange database volume and 20 SnapManager for Exchange Snapshot copies in each Exchange transaction log volume.

To monitor the health of SnapManager retention, use Windows PowerShell® commands from the Data ONTAP PowerShell Toolkit and Windows native PowerShell commands.

4.10 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 great for any situation in which testing or development occurs, any situation in which progress is made by locking in incremental improvements, and any situation in which there is a desire to distribute data in 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 2013 cumulative update or hotfix into production.

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

Refer to the FlexClone section in the “Data ONTAP Administration Guide” for detailed information on how FlexClone works and on command-line references.

Best Practice
Use SnapDrive to connect to the required Snapshot copy. This will automatically execute the FlexClone operation.

4.11 NetApp Deduplication

The deduplication process only stores 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, which will be compared to other fingerprints on the volume. If two fingerprints are found to be the same, a byte-for-byte comparison is done 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. Due to 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 increase on performance due to the fact that the system memory and 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 a much improved access time.

Note: Setting `read_realloc` to “`on`” for a volume that has enabled deduplication will not affect performance, nor will it reduce storage efficiency.

Note: Deduplication is transparent to Exchange, and the block changes are not recognized by Exchange. Therefore the 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 Exchange Server 2013 databases in the 15–35% range.

Best Practices

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

5 Designing Storage Efficiency for Exchange Server 2013

Table 1 presents the principles for designing storage efficiency for Exchange Server 2013.

Table 1) Storage efficiency principles.

Principle 1	
Statement	Standardize around RAID-DP as the RAID technology.
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. Higher read latency associated with SATA drives can be reduced by using the NetApp Flash Cache card.
Implications	High-capacity, economical disks have a better mapping to Exchange environments with larger mailbox sizes and high-capacity demands.
Principle 3	
Statement	Thin provision the storage using NetApp FlexVol.
Rationale	Thin provisioning in an Exchange environment increases the storage utilization, reduces complexity by provisioning the sized LUNs to the host, and reduces the overall storage requirement. As utilization grows, storage can be added, with no impact to Exchange.
Implications	Storage silos can be removed.
Principle 4	
Statement	Wherever appropriate, use server virtualization technologies with NetApp storage to gain additional savings.
Rationale	NetApp's 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's deduplication and cloning technology enables a more dynamic and storage-efficient infrastructure.

Principle 5	
Statement	Using deduplication in Exchange Server 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, using deduplication can save anywhere from 15% to 35% of storage.

6 NetApp Solution for Microsoft Exchange Server 2013

6.1 NetApp Storage Software and Tools

NetApp Windows Host Utilities Kit. This kit should be used in both physical and virtual environments; it configures Windows Server to access virtual disks on a NetApp storage system using the Fibre Channel, iSCSI, or FCoE protocol. It also helps to align the master boot record for the Microsoft 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 when Data ONTAP DSM is installed.

MPIO. The NetApp Windows Host Utilities Kit uses the Microsoft framework for MPIO, and it helps storage providers develop multiple paths to optimize connectivity with the storage arrays.

MPIO load balancing. This type of load balancing, supported by MPIO, uses multiple data paths between server and storage to provide greater throughput of data than could be achieved with only one connection.

MPIO-based fault-tolerant failover. In this scenario, multiple data paths to the storage are configured. If one path fails, the HBA or NIC fails over to the other path and resends any outstanding I/O.

- For a server that has one or more HBAs or NICs, MPIO offers support for redundant switch fabrics or connections from the switch to the storage array.
- For a server that has more than one HBA or NIC, MPIO also offers protection against the failure of one of those adapters directly within the server.

NetApp OnCommand Host Agent. The OnCommand Host Agent gathers host-specific data and sends it to NetApp OnCommand. It reports on files, folders, drive size, HBA information, and average CPU and memory usage. It helps applications such as Exchange by showing how the server is running and tracks the hourly average of many data points.

SnapDrive for Windows. This application helps with storage provisioning and managing disks in both physical and virtual environments. SnapDrive for Windows manages the LUNs on the storage system, making them available as local disks on Windows hosts.

The following are the key features of SnapDrive for Windows:

- Enhances online storage configuration, LUN expansion, and shrinking; provides streamlined management
- Works in conjunction with NetApp SnapMirror software to facilitate disaster recovery from either asynchronously or synchronously mirrored destination volumes
- Enables management of SnapDrive for Windows on multiple hosts
- Enhances support on Microsoft failover cluster configurations
- Simplifies iSCSI session management
- Enables technology for SnapManager for Exchange

- SnapDrive in Data ONTAP 7-Mode supports RPC, HTTP, and HTTPS protocols

7 Backup and Recovery

7.1 SnapManager for Exchange Server Overview

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

SnapManager for Exchange uses online Snapshot technologies that are part of Data ONTAP. It integrates with Exchange backup and restores APIs and the Volume Shadow Copy Service (VSS). SnapManager for Exchange can use SnapMirror to support disaster recovery even if native Exchange DAG replication is leveraged.

SnapManager for Exchange provides the following data management capabilities:

- Migration of Exchange data from local or remote storage onto NetApp LUNs. Application-consistent backups of Exchange databases and transaction logs from NetApp LUNs
- Verification of Exchange databases and transaction logs in backup sets
- Management of backup sets
- Archiving of backup sets
- Restoration of Exchange databases and transaction logs from previously created backup sets providing lower RTO and more frequent recovery point objectives (RPOs)
- Capability to restore active or replica databases using rapid reseed and prevent full reseed of a replica database across network.

Some of the new features released in SnapManager for Exchange 7.0 include:

- Gapless DAG backup enhancements
- Backup retention management enhancements
- Backup management improvements for remote additional backups

7.2 SnapManager for Exchange Server Architecture

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

For more information about VSS, refer to Microsoft's Volume Shadow Copy Service Overview.

Figure 1) SnapManager for Exchange Server architecture.

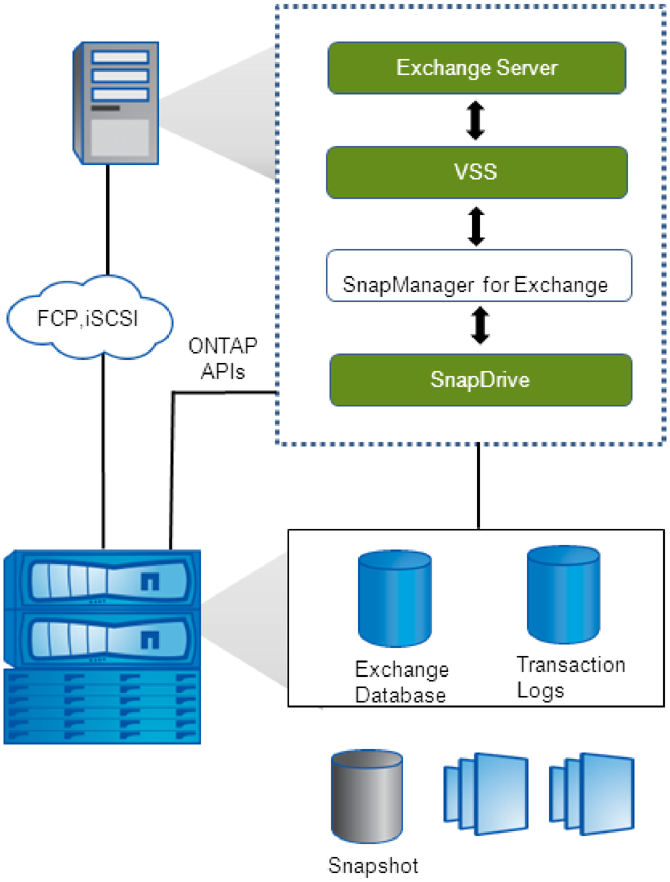

7.3 SnapManager for Exchange Server Installation and Upgrade Considerations

For information about compatible versions of SnapManager for Exchange, SnapDrive for Windows, and Data ONTAP, see the NetApp Interoperability Matrix Tool.

Before upgrading SnapManager for Exchange, consider the following steps.

1. Back up the operating system installation on the Exchange server. This 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 Exchange server.
3. Back up the boot and system drives.

7.4 Migrating Exchange Data to NetApp Storage

The process of migrating Exchange databases and transaction log files from one location to another can be a time-consuming and lengthy process. There are many manual steps that need to be taken to make sure the Exchange database files are in the proper state to be moved. Additionally, more manual steps need to be performed to bring those files back online for handling Exchange traffic. SME automates the entire migration process, eliminating any potential user errors. After the data is migrated, SME automatically mounts the Exchange data files and allows Exchange to continue to serve e-mail.

7.5 Layout Recommendation

Storage layout for Exchange data should be designed for optimal performance, high availability, and data protection before the actual migration of the Exchange databases can be carried out. The best practices for designing the storage layout for Exchange environments are discussed in Sizing and Storage Layout for Exchange Server 2013. The following sections of SnapManager 7.0 Microsoft Exchange documentation contain more details that can affect the design.

- Rules for Exchange Server databases enforced by the Configuration Wizard
- SnapManager support for NTFS volume mount points

These sections present a comprehensive list of guidelines that must be followed while designing the Exchange data layout on NetApp storage in order for SnapManager for Exchange (SME) to work successfully.

Best Practices
Exchange Server 2013 databases should be kept on individual LUNs on separate volumes.

Placing databases on individual LUNs on separate volumes enables fast and granular recovery. The rapid reseed functionality using Snapshot copies would be much faster than the out-of-box reseeding solution from Microsoft. However the databases can also be placed on individual LUNs on the same volume, which also provides the capability to restore the databases individually as SnapManager for Exchange utilizes the LUN Clone Split restore for restore operations. Keeping multiple databases on the respective individual LUNs on a single volume also provides slightly higher deduplication rates since multiple databases are hosted on a single volume. Preferably, the first of the two options should be used. The second option needs to be considered in those cases where volume count would be a challenge.

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

Typically in DAG environments, backups are run on one of the passive database copies. Make sure that the transaction log LUN on the database copy where backup is run with longer retention is correctly sized by taking the backup retention into account. If not done, it can lead to situations when the backup retention is different for other copies of the database and the administrator may change the Snapshot retention without considering the initial volume sizing.

For additional information on Exchange data layout and sizing, refer to the section 8, “Sizing and Storage Layout for Exchange Server 2013.”

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

- Backup metadata
- Transaction log backups

In SnapManager for Exchange, 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 not only saves space but also makes the log backups quicker.

Best Practices

- Place the Exchange transaction log files and the SnapInfo directory in the same LUN.
- Place Exchange transaction log files into as few flexible volumes as possible, as appropriate for the business use and RPO/RTO.
- NetApp recommends using NTFS hard links by placing the SnapInfo directory on the same LUN as the transaction log directory whenever possible. This will increase storage utilization, eliminate the physical copy overhead incurred on the Exchange Server, and increase backup performance.
- When separate LUNs are used for the Exchange transaction log files and the SnapInfo directory, place those LUNs in the same volume. Both of these LUNs will have a similar I/O profile, allowing them to share the same volume. For disaster recovery scenarios, having the entire log set for Exchange on the same volume will help achieve SLAs. Placing the SnapInfo directory in a LUN different from the transaction logs will require 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 fashion, NetApp recommends limiting the number of log streams per LUN to from 5 through 10.

7.6 SME Service Account

For SME to work, it needs to have a service logon account with appropriate permissions. This requires at least the following permissions:

- The SME service logon account must be a member of the “Organization Management” role in Exchange Server 2013.
- The SnapManager for Exchange Service must be a member of the Exchange server’s local administrators group.

7.7 Prerequisites for Migrating Exchange Server Mailbox Databases

In SME 7.0, the DAG is itself a unit of management: all the nodes of a DAG can be managed by registering the DAG with SME 7.0. Optionally, each DAG node can also be managed at the server level.

Before you migrate the Exchange Server mailbox databases on a DAG, make sure that NetApp storage is provisioned on each server that will have a database copy utilizing the same path to be specified during the 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, including both the active and passive databases, will not be migrated by SnapManager, and SnapManager will not be able to back up databases on that server.

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

There can be situations where all member servers of a DAG might not use NetApp storage to store Exchange data. In this case 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:

- All nodes in a Microsoft failover cluster must have SnapDrive installed, even if connected to third-party storage. Make sure there is proper network connectivity to the storage system from all the nodes.

- However, SnapManager does not need to be installed on every node. In this configuration, the backup must be made at the server level.

Note that in the preceding 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. This includes migration of mailbox databases as well. This operation must be performed individually on each and every member node that uses NetApp storage and has SME installed on it.

Furthermore, note that in the preceding scenario, SME can be licensed on only those nodes on which it is installed. However, there is a risk involved in those situations where 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, it is possible that a DAG member server that does not have SME can become the active server. In this case, even if the member server is connected to NetApp storage, SME cannot be used on it.

Best Practices

- On Exchange Server 2013, use the same drive letters or mount points for the Exchange data LUNs on all the nodes of a DAG.
- Install SnapManager for Exchange and SnapDrive for Windows on all member servers of the Database Availability Group (DAG). If SnapManager for Exchange is not installed on all nodes, SME will not be able to migrate databases using the configuration wizard.

7.8 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 Exchange data using SnapManager for Exchange.

7.9 Backup

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

- **Organization SLA**. This parameter will determine the frequency and the type of backups.
- **Backup retention planning**. This parameter will determine whether backup sets need to be retained on the primary or secondary site.
- **Backup verification policy**. This parameter will determine when to engage backup verification and when not to.

The time taken to restore Exchange data in the event of an outage is dependent on the number of transaction logs that need to be replayed. Hence, reducing the number of transaction logs that need to be replayed when restoring from a full backup is important. The only way to do this is to create more frequent backups.

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

- **SME notification system.** This feature in SME allows the administrators to receive detailed e-mails for failures in executing operations in SME. The e-mail notification can be configured using the SME configuration wizard. Refer to SnapManager for Exchange Installation and Administration Guide for more details.
- **Monitoring host-side events**. SME logs many events, each with its own event ID. These events get logged in the Windows application event log on the Exchange Server. Monitoring some of the critical events using scripts, SCOM, MOM, and so on will help alert administrators of failures encountered by SME.

Recovery point objectives (RPOs) have become a defining part of a data protection plan for Exchange. The ability to have a near-zero RPO is highly desirable by Exchange administrators, as it minimizes the amount of data that is lost between the last full verified backup set and the point of failure. To help achieve the desired service-level agreement (SLA) and RPO times, SME has frequent recovery points (FRPs). FRPs are optimized backup sets that are created through SME. The backup sets only contain 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. Higher frequency of FRP backups reduces RPO times.

7.10 Gapless Backup

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

Figure 2) Gapless backup example.

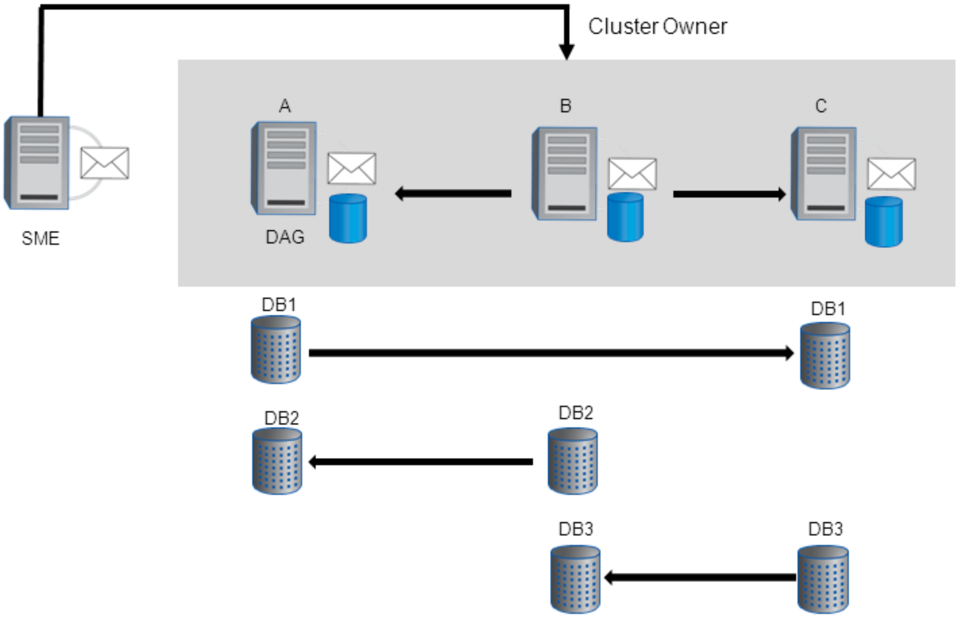

SnapManager for Exchange (SME) Gapless Backup Example

- Three-node DAG; Node B is the owner node of the DAG.
- Each database has two copies. The arrows indicate the log shipping direction pointing to the passive database copy.
- SME is shown running on a client machine though it can be run from any of the DAG nodes.

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

Sequence of Operation

When SME connects to Node B, which is the DAG owner, it gets the list of databases to back up and then starts the backup operation.

Figure 3) Sequence of operation.

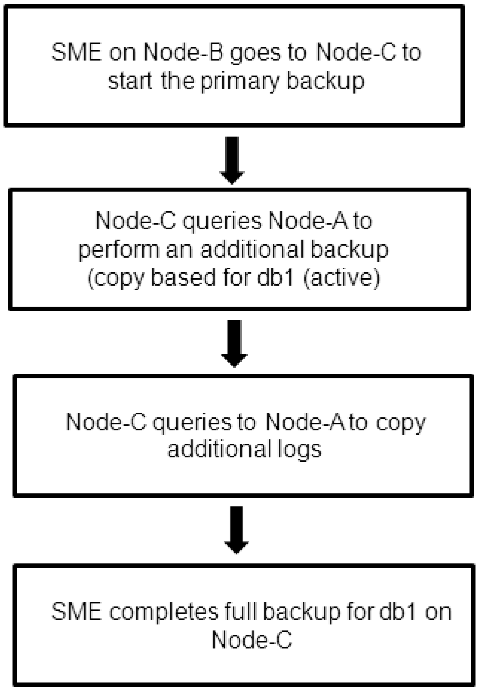
SME on Node-B goes to Node-C to start the primary backup
Node-C queries Node-A to perform an additional backup (copy based for db1 (active)
Node-C queries to Node-A to copy additional logs
SME completes full backup for db1 on Node-C

In the end you will have a backup for all six database copies, and all backups can perform an up-to-the minute restore. Since 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 the databases are placed in an optimal manner to achieve the best results while using gapless backup.

For more information about the optimal layout, refer to SnapManager 7.0 Microsoft Exchange documentation.

7.11 Server Backup and Frequent Recovery Point

A smaller recovery point objective can be achieved in larger environments with many databases and server nodes in a DAG. This is a sample configuration that was successfully implemented at a couple of large customer environments. The strategy utilizes a “full backup” on one copy of the database that does not need 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. In the event that one copy of the database fails, the rapid reseed process can be used to perform a nondisruptive restore.

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.

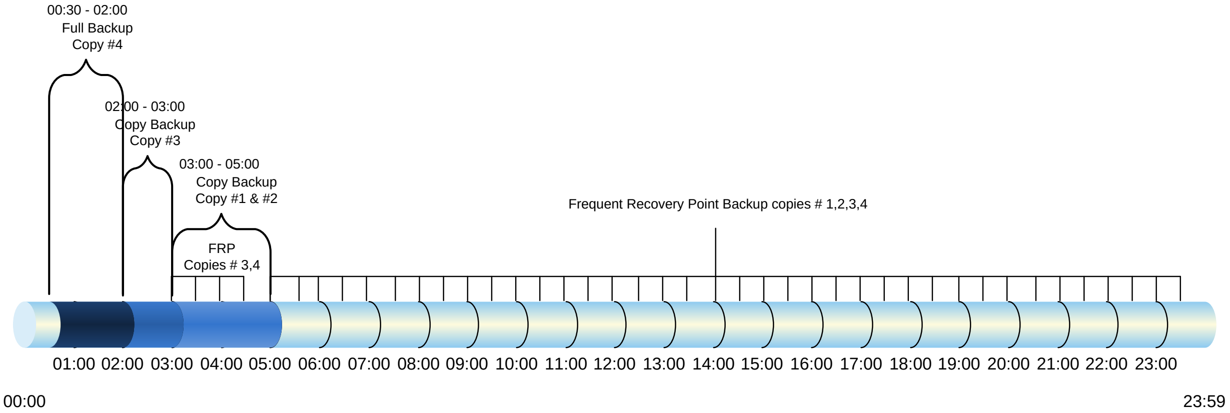

It is important that a copy backup is scheduled at least 10 minutes after the full backup. This is to make sure that the log truncation activity from the full backup has successfully replicated to each database copy.

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.

Database Verification

Database verification is not a support requirement for databases with at least two copies in a DAG. Microsoft recommends verification of the transaction logs and SnapManager for Exchange performs log verification during the restore operations. By default, when performing a DAG backup with SnapManager for Exchange, verification will be off. You can monitor both locally running and remote verification jobs through the main SnapManager for Exchange management console. A single Exchange mailbox server can only run one verification process at a time on a particular verification server. A verification server can simultaneously run one verification job from each Exchange mailbox server. More than one verification server can be used in order to simultaneously verify more than one backup job on a single Exchange mailbox server. Many customers utilize virtual machines to off-load the verification.

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

Recovery

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

- **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 Exchange data to a chosen point in time. Any Exchange data past that point is not restored. This option is particularly useful when trying to restore to a point before something such as data corruption occurred. A point-in-time restore only replays and applies to the database of 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 point in 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 (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.

Restoring Passive Copies in a DAG Setup

In this section discusses the best practices and solutions with respect to restoring Exchange data using SnapManager for Exchange.

One of the key challenges in a DAG environment is to minimize reseeding of database copies in the event of a database failure. When a Snapshot copy of the failed database is restored, the network resources are not consumed by the reseeding operation. 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.

For the replica database copy, the following steps are performed:

1. Select a Snapshot copy (latest) of the replica database and do not select the Mount after restore option. (Make sure “Mount after restore” is unchecked to restore a passive database copy. If it is checked, the passive database copy being restored will become active, potentially losing data created after the backup was created.)
2. SnapManager for Exchange suspends the replication on the replica database copy before starting the restore using the suspend-database Windows PowerShell command.
3. SME restores the selected Snapshot copy.
4. Replication is resumed on the replica copy after restore is completed using the restore-database Windows PowerShell command.
5. Restore completes with the replica database copy in healthy state.

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

- The replica database copy is corrupt, exchange resynchronization fails, and it 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.

7.12 Snapshot Retention Guidelines

Primary Storage

The recovery point objective (RPO) will guide how frequently a backup is created. NetApp flexible volumes running Data ONTAP 7-Mode can store a maximum of 255 Snapshot copies per flexible volume. The amount of storage needed for Snapshot copies will depend on the rate of change.

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

Secondary Storage

SnapManager for Exchange backups can be archived with SnapVault (volumes can be copied for space-efficient, read-only, disk-to-disk backup). SnapManager coordinates with OnCommand Unified Manager Core Package 5.2 or later, which includes the NetApp Management Console for the integration.

Long-Term Archiving to Tape

Long-term archiving of SnapManager for Exchange backups to tape can be done by using an NDMP-based backup to copy the LUN in the Snapshot copy created by SnapManager for Exchange to tape or by mounting the LUNs in the Snapshot copy created by the SnapManager for Exchange backup and then streaming to tape.

Best Practices
• Use the business requirements established by the Exchange stakeholders to help determine the number of Snapshot copies to retain online. • Use NetApp OnCommand Unified Manager Core Package 5.2 or later, which includes the NetApp Management Console to archive SnapManager for Exchange backup sets from primary to secondary storage. • If mounting LUNs in the Snapshot copy created by the SnapManager for Exchange backup to archive the SnapManager for Exchange backup to tape, license the controller with FlexClone so that there are no busy Snapshot copies. • When performing an up-to-the-minute restore, restore from the most recent backup, as less logs need to be applied to the recovered database. If the most recent backup hasn't 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).

7.13 Single Mailbox and Item-Level Recovery

The Single Mailbox Recovery (SMBR) 7.0 application integrates with the NetApp storage solution for Microsoft Exchange Server 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 Exchange server or to an offline Outlook Personal Storage Table (PST) file. SMBR locates and restores Exchange mail items without the assistance of the Exchange Server, thereby eliminating the need for a separate recovery server. SMBR 7.0 can also be used to recover public folder mailboxes, including organizational forms.

Note: Single Mailbox Recovery can be launched through SnapManager for Exchange using the Run SMBR option in the Action pane, to locate and then restore items at any level of granularity, directly to an existing mailbox on the Exchange server.

Figure 5) Integration of SMBR with NetApp storage solution.

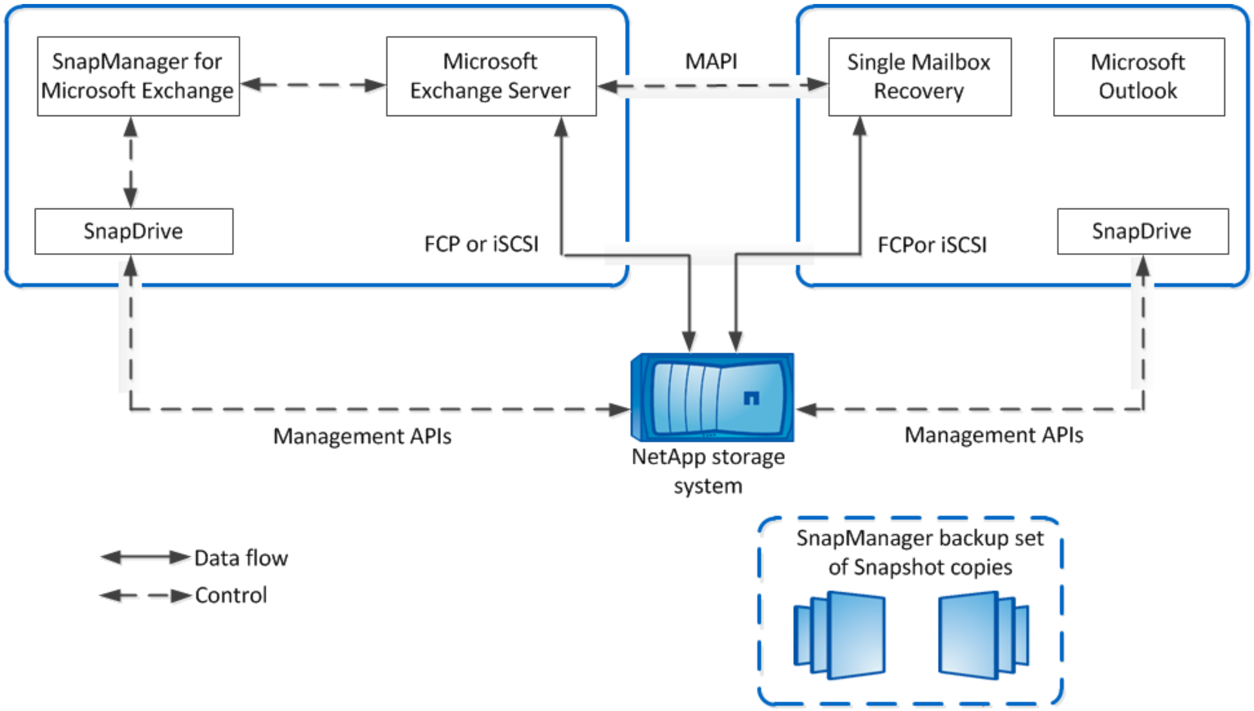

7.14 Single Mailbox Recovery 7.0 Administrative Server

The NetApp Single Mailbox Recovery 7.0 Administrative Server (SMAS) application is a framework that can centralize services for multiple clients. It provides both client and server support for NetApp Single Mailbox Recovery (SMBR) 7.0 and NetApp Single Mailbox Recovery 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 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 Exchange Server. In order 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 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 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 to mailboxes is done utilizing the SMBR-Administrative Server (SMBR-AS). It is strongly recommended that SMBR-AS be used along with SMBR.

7.15 Troubleshooting

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

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

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

8 Sizing and Storage Layout for Exchange Server 2013

8.1 Aggregate Recommendations

Fewer, larger aggregates will maximize performance; however, they might not meet the data availability requirements set forth in the SLA agreement. In Exchange Server 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.

Best Practices

- NetApp recommends having at least 10% free space available in an aggregate hosting Exchange data for optimal storage performance.
- On controllers that are dedicated for Exchange deployment, set the global wafl.optimize_write_once flag to off to optimize random workloads. The flag must be set before Exchange aggregates are created. If Exchange aggregates are already present, reallocate –A must be run on each Exchange aggregate. This is a time-consuming process that can affect performance during the reallocation scan. .

8.2 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 results in the following additional benefits for Microsoft Exchange Server 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, since there is no concern about Snapshot copy schedules overlapping different servers.
- NetApp recommends having at least 10% free space available in a volume hosting 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, plus 10% free space in the volume.

Transaction Log Volume

1. Transaction log LUN size: 71GB

   Holds 833 users sending 100 messages a day
2. Messages per day: 83,300 * 75KB = 5.95GB per day
3. Snapshot retention of 7 days with 3 days of fault tolerance is 10 days.
4. 5.95GB * 10days = 59.5GB
5. Log volume (59.5GB + 71GB / (1-.1)) or 145GB

   10 days of Snapshot copies+ transaction log LUN + 10% free space in the volume.

Database Volume

1. Database LUN size: 2691GB (1826GB database)
2. If you size the database LUN for quota, this includes; maximum mailbox size, deleted items in the dumpster, calendar, 3 days of incoming mail, and whitespace in the database.
3. Snapshot retention of 7 days with 3 days of fault tolerance is 10 days.
4. Daily change rate (5%) * 10 days (common rates are in the 2% to 8% range)
5. 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 to fewer than 10.

Best Practices
• When creating LUNs, use volume mount points. There are a finite number of drive letters, and in a DAG each database path 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 databases of at least 2TB.

Do not create mount points for additional LUNs on another LUN that holds an Exchange Server 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 would require 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 the 5% free disk space, the database itself, and the content index files.

1. The MBXSize is the MBXLimit plus the dumpster.
2. The MBXLimit is the stated maximum mailbox size, in our example case, that is 2GB.
3. 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 = MBXLimit * 0.012 (1.2%)

   Calendar version storage = MBXLimit * 0.03 (3%)

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

   Dumpster = SingleItemRecovery + CalendarVersionStore + (#Messages * MessageSize * (DeletedItemRention+1 [today]))

   Dumpster = 24.6 + 61.4 + (((100 * 75KB * (14+1))/1024)

   = 24.6 + 61.4 + 109.8MB = 195.8MB
4. The database size is the MBXSize multiplied by the number of users.
5. The 'DB size + overhead' adds 0% to the database size (vs. 20% in Exchange 2010).
6. 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% = 365.2GB (vs. 10% in Exchange 2010)

   (1826 + 730) / (1-.05) = 2691GB

Transaction Logs

The transaction logs are 1MB in size, must include transaction logs generated by the users from move mailbox requests, and the backup fault tolerance window. Microsoft by default assumes a 3-day backup fault tolerance window.

1. User logs (calcnumusertlogs): LogGen * users * Datagrowth

   LogGen = 10 for each 50 messages per day in a user profile

   Users: An example 100 message per day 20,000-user configuration would be:

   20 * 20000 = 400,000 logs
2. Move mailbox (logdiskspacereqmove)= (users * 1%) * (MBXSize /1024)

   (20000 * 1%) * (2244 /1024)

   200 * 2.191

   438GB
3. Log backup (logdiskspacereqbackup)

   UserLogs * BackupFailTol (default 3)

   400,000 * 3 = 1,200,000 logs
4. Total log disk space (totlogdiskspace)

   Log backup + move mailbox

   (1,200,000 /1024) + 438GB

   1609GB
5. Add 5% free disk space.

   1609GB / (1-.05) = 1694GB
6. Divide by the number of databases, in this case, 24.

1694GB / 24 = 71GB per transaction log LUN

Notice how the majority of this space is because of move mailbox and large mailbox sizes.

8.3 Capacity Planning

A properly sized Exchange environment will meet or exceed the customer service-level agreement (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 utilized when planning an Exchange environment for a customer:

- The Microsoft Exchange 2013 Server Role Requirements Calculator
- The NetApp Exchange Sizing Tool (SPM); work with your local NetApp partner or your NetApp representative to enable proper sizing

The sizing information provided by these tools is an important component for planning an Exchange environment and provides a framework for database layout and LUN requirements. It is important to realize that 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 Exchange server deployments using NetApp storage, it is important to use the NetApp Exchange Sizing Tool with the data from the Microsoft Exchange 2013 Mailbox Server Role Requirements Calculator.

Best Practice
Consult a local NetApp Exchange expert or your NetApp partner to assist in accurately sizing Exchange Server 2013. Use the NetApp Sizing Tool for Exchange (SPM) to size all Exchange server deployments utilizing NetApp storage.

9 Performance

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

9.1 Flash Pool

Flash Pool™ is the combination of solid-state disks and traditional hard disk drives into a single aggregate. Aggregates configured in a Flash Pool aggregate do not participate with Flash Cache, and performance is less than using Flash Cache instead. Utilizing Flash Cache is the NetApp best practice for Exchange workloads, particularly if SATA is used. Flash Cache is not mirrored to the partner controller and will require time after a controller failover to warm, slowing performance during that time. With Flash Pool the caching (to SSD) occurs at the aggregate level, and any controller failover does not require time for the cache to warm, since the cache is part of the aggregate.

9.2 Flash Accel

Flash Accel™ is a host-based cache that utilizes PCI-E or SSD devices to enhance the performance of read requests for data that is cached locally on the server. This cache is additive to the caching built into the storage, such as Flash Cache. Typical results in a Jetstress environment show 15% to 20% increase in the total achievable user transaction IOPS.

9.3 SATA Performance Considerations

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

Exchange 2013 utilizes background database maintenance (BDM) to maintain the consistency of the databases. BDM applies a per-database performance tax on the storage system that must be taken into account when sizing the storage for Exchange. Having fewer, larger databases in the 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 in which there are a large number of small databases.

To help improve the storage efficiency and read I/O performance and latency of SATA-based deployments, Flash Cache should be used. Flash Cache is a read cache that can be installed on certain models of NetApp storage controllers. Flash Cache enables fewer SATA disks to be used in SATA-based deployments, because a percentage of the Exchange database working set is cached in the Flash Cache aggregate, thus greatly reducing the amount of read I/O on the SATA disk. NetApp recommends Flash Cache and SATA for deployments exceeding 1,000 mailboxes or when SATA-based designs are bounded by performance instead of capacity.

Best Practice
NetApp recommends using Flash Cache when placing Exchange Server 2013 database files on SATA physical disk drives.

9.4 Database Sizing Considerations

Using a smaller number of larger databases can help reduce the amount of background database maintenance I/O as well as reduce 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. A database size of 2TB is a practical size that can be restored in minutes with SnapManager for Exchange. Microsoft recommends 2TB as a maximum size, but supports up to 16TB for databases on both 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.

9.5 Aggregate Sizing and Configuration Considerations

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

As mentioned earlier in this document, Exchange databases and transaction logs can be placed on the same aggregate. There is marginal benefit in 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 Exchange data to exist if an aggregate is lost due to a catastrophic failure. Placing database copies on separate aggregates also helps isolate background database maintenance I/O to the aggregates where the database copies are located.

If Exchange is virtualized, place the Exchange VMs on a separate aggregate from the 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 Exchange should be configured for RAID-DP. This enables maximum data protection for the 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. Additionally, 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. The recommendation is to use the default RAID group size when the Exchange aggregate is created, because this balances storage efficiency and performance.

Exchange workloads can run effectively on both 32-bit and 64-bit aggregates if the aggregates and controller heads are properly sized. NetApp recommends that 64-bit aggregates supporting an Exchange workload be used only in configurations that are supported in the NetApp Exchange sizing tool. This is so that the storage is properly sized for the anticipated Exchange workload. NetApp recommends consulting a local NetApp Exchange expert to provide accurate performance sizing when considering the use of 64-bit aggregates for Exchange environments.

9.6 Volume Configuration Considerations

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

10 Virtualization

10.1 Microsoft Support for Exchange 2013 in Virtualized Environments

The documentation for support for Exchange 2013 in virtualized environments can be found in the Microsoft TechNet article Exchange 2013 Virtualization.

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

- All Exchange 2013 server roles are supported in a virtual machine.
- Exchange Server 2013 virtual machines (including Exchange Mailbox virtual machines that are part of a DAG) may be combined with host-based failover clustering and migration technology as long as the virtual machines are configured such that they will not save and restore state on disk when moved or taken offline.
- All storage used by an Exchange guest machine for storage of Exchange data must be block-level storage because Exchange 2013 doesn't support the use of network-attached storage (NAS) volumes. Also, NAS storage that's presented to the guest as block-level storage through the hypervisor isn't supported.
- Microsoft does not support the use of virtual machine Snapshot copies of Exchange virtual machines.

Note: Currently, SME only supports RDMs through the ESX® iSCSI initiator, and LUNs attached directly to the guest with an iSCSI initiator.

11 High Availability

In Exchange 2013, the Database Availability Group (DAG) feature was implemented to support mailbox database resiliency, mailbox server resiliency, and site resiliency. The DAG consists of two or more servers and each server can store up to one copy of each mailbox database.

The DAG Activation Manager manages the database and mailbox failover and switchover processes. A failover is an unplanned failure and a switchover is a planned administrative activity to support maintenance activities.

The database and server failover process is an automatic process when a database or mailbox server incurs a failure. The order in which a database copy is activated is set by the administrator.

For more information on Exchange 2013 Database Availability Groups, refer to the Microsoft TechNet article Database Availability Groups.

11.1 Exchange 2013 Database Availability Group Deployment Scenarios

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- and application-level redundancy. In this situation, deploying a two-node DAG utilizing RAID-DP provides not only server- and application-level redundancy but double disk failure as well. Adding SnapManager for Exchange in a single-site scenario enables point-in-time restores without the added capacity requirements and complexity of a LAG copy.

Multisite Scenario

Extending a DAG across multiple data centers provides high availability of servers and storage components and adds site resiliency. When planning a multisite scenario, NetApp recommends 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 will provide site resiliency but also provide high availability in each site.

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 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 will provide site resiliency but also provide 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 regarding path, capacity, and performance.
- Place flexible volumes for active and passive databases onto separate aggregates. 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 Exchange 2013 is in a DAG configuration with at least two copies of the databases, with Exchange background database maintenance enabled.
- Verification of database backups and transaction log backups is required if Exchange 2013 is in a standalone (non-DAG) configuration.
- In Exchange 2013 standalone environments using SnapMirror, configure database backup and transaction log backup verification to occur on the SnapMirror destination storage.

12 Summary

Microsoft Exchange Server 2013 is not a one-size-fits-all application. Multiple configuration options are available to suit most of the needs of any customer. NetApp storage appliances and data management software are built in a similar fashion, providing users with the flexibility to manage 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 Exchange Server 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 Exchange data. This guideline enables a highly available, easy-to-manage Exchange environment that meets SLAs. Consult with a local NetApp Exchange expert when planning and deploying Exchange environments onto NetApp storage. NetApp Exchange experts can quickly identify the needs and demands of any Exchange environment and adjust the storage solution accordingly.

References

- Exchange 2013 system requirements
  http://technet.microsoft.com/en-us/library/aa996719(v=exchg.150).aspx
- Storage management guide:
  http://support.netapp.com/documentation/docweb/index.html?productID=61652&language=en-US
- Volume shadow copy service overview:
  http://msdn.microsoft.com/en-us/library/aa384649(v=VS.85).aspx
- Exchange 2013 storage configurations options:
  http://technet.microsoft.com/en-us/library/ee832792(v=exchg.150).aspx
- Exchange 2013 mailbox server role requirements calculator:
  http://gallery.technet.microsoft.com/office/Exchange-2013-Server-Role-f8a61780
- Exchange 2013 high availability and site resiliency:
  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

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.

NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.

Go further, faster®

NetApp®

www.netapp.com