



# Dell SC Series Storage and Microsoft Exchange Server 2016 Best Practices

Dell Engineering  
August 2016

## Revisions

Date	Description
October 2012	Initial release
April 2013	Update
September 2015	Updated with new SC Series products and best practices updates
August 2016	Updated for Exchange Server 2016

## Acknowledgements

Engineering: Mark Boeser

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

© 2012 - 2016 Dell Inc. All rights reserved. Dell, the DELL logo, and the DELL badge are trademarks of Dell Inc. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell disclaims any proprietary interest in the marks and names of others.



# Table of contents

Revisions.....	2
Acknowledgements.....	2
Executive summary .....	5
1 Predeployment.....	6
1.1 Understanding storage virtualization.....	6
1.2 Understanding Exchange I/O.....	6
1.3 Exchange 2016 architecture changes .....	7
2 Configuration .....	9
2.1 Disk layout .....	9
2.2 Storage Center page size considerations .....	10
2.3 Storage Profiles and Data Progression.....	11
2.3.1 Microsoft recommends ReFS for Exchange 2016 database volumes (JBOD).....	11
2.3.2 Default allocation unit size .....	11
2.3.3 Partition type.....	11
2.4 Mailbox server sizing considerations.....	11
2.4.1 Archive mailboxes .....	13
2.5 Online database maintenance.....	13
2.5.1 Online defragmentation .....	13
2.5.2 Online database scanning (checksumming) .....	14
2.5.3 Database fragmentation .....	14
3 Performance and monitoring.....	15
3.1 Monitoring.....	15
3.1.1 Monitoring the database defragmentation process .....	15
3.1.2 Monitoring with System Center Operations Manager .....	16
4 Troubleshooting.....	17
4.1 Enterprise Manager/Dell Storage Manager .....	17
4.2 Microsoft Exchange 2013/2016 Diagnostic Service logs.....	17
4.3 Microsoft Exchange Troubleshooting Assistant (2007–2010).....	18
4.4 Microsoft Exchange Best Practices Analyzer (2007–2010).....	18
5 Backup and recovery .....	19
5.1 Backups and Instant Replays.....	19



5.1.1	Backup schedules .....	19
5.1.2	Data Instant Replays (snapshots).....	19
5.1.3	Database verification .....	19
6	Recovery.....	20
6.1	Recovery databases .....	20
6.1.1	Single item recovery .....	20
7	Disaster recovery .....	21
7.1	Database copies and database availability groups .....	21
A	Technical support and resources .....	22
A.1	Related documentation .....	22



## Executive summary

The purpose of this document is to provide additional guidelines when using Microsoft® Exchange Server 2016 with Dell™ SC Series storage.

These guidelines should be evaluated thoroughly as every environment configuration is different. They should not be construed as final recommendations in the configuration of Dell SC Series or Microsoft Exchange environments.



# 1 Predeployment

## 1.1 Understanding storage virtualization

Traditional SAN solutions require a LUN to be created by carving out a specific amount of disk space from a specified set of spindles. This process is time-consuming and hard to manage because it involves keeping track of which servers are mapped to which storage by which worldwide name (WWN). This often requires administrators to manage the storage configuration using a spreadsheet or other means outside of the storage management interface.

Dell SC Series storage virtualizes storage at the disk level and uses all available spindles as a single virtual disk pool. It can create volumes in which all of the drives are used together, causing performance to improve because the blocks of data are written in parallel to all managed drives at once.




Figure 1 Traditional disk mapping compared to the Dell SC Series approach

## 1.2 Understanding Exchange I/O

The SAN configuration is an important part of any application configuration, and this is especially true with Exchange Server. Understanding how Exchange Server works with storage helps administrators make sure that systems run in their most capable state. To ensure that Exchange Server will run in its optimal environment, performing some simple tests can determine whether a server and disk subsystem can provide the necessary performance.

Several tools exist to put a load against and test the performance of Exchange Server and disk storage, including Exchange Load Generator (LoadGen) and Jetstress. Each of these has the capability to simulate Exchange I/O patterns as well as the client experience, which can provide the estimated performance numbers to expect from the disk subsystem. LoadGen and Jetstress are available from Microsoft as free downloads and are discussed further in the section: Mailbox server sizing considerations.



Another useful tool is the Windows Performance Monitor, which can help define a baseline and show how the application may perform in the current environment. This tool is discussed further in the section, Performance and monitoring.

## 1.3 Exchange 2016 architecture changes

With CPU hardware generally less expensive than in the past, the constraint of expensive server hardware has been alleviated. Exchange 2016 takes advantage of this with a primary design goal of simplicity in scale and hardware utilization. The number of server roles has been reduced to two: Mailbox and Edge Transport server roles.

The new 2016 **Mailbox server role** includes all server components from Exchange 2013 Mailbox and Client Access roles:

- **Client Access services:** provide authentication, limited re-direction, and proxy services offering the usual client access protocols: HTTP, POP, IMAP and SMTP.
- **Mailbox services:** including the back-end client access protocols, Transport service, Mailbox databases, as well as Unified Messaging. The Mailbox server manages all active mailboxes on that server.

Other new or enhanced features that are notable for storage considerations are:

### **In-place Archiving, retention, and eDiscovery:**

- Public folder support for In-Place eDiscovery and In-Place Hold
- Compliance Search – available only in Exchange Management Shell (EMS)

### **Improved performance and scalability:**

- Search architecture redesigned as asynchronous
- Improved search scalability from 5K mailboxes to 10K mailboxes, or unlimited in EMS

To provide Exchange Native Data Protection, Exchange 2016 continues to use database availability groups (DAGs) and mailbox database copies, along with features such as single item recovery, retention policies, lagged database copies, and others. The high availability platform, the Exchange Information Store, and the Extensible Storage Engine (ESE) have all been enhanced to provide greater availability, easier management, and reduced costs.

With respect to storage, these enhancements include:

- **Reduced IOPS compared to Exchange 2013:** A reduction in IOPS/mailbox size enables larger disks to be better utilized, providing capacity and IOPS as efficiently as possible.
- **Multiple databases per volume:** This enables multiple databases (mixtures of active and passive copies) to be hosted on the same volume and is another enhancement that allows larger disks to be used.



- **Automatic Reseed for DAS disk failures:** This provides a quick restore to database redundancy after a DAS disk failure. If a physical disk fails, the database copy stored on that disk is copied from the active database copy to a spare physical DAS disk on the same server. If multiple database copies were stored on the failed disk, they can all be automatically reseeded on a spare disk. This enables faster reseeds because the active databases are likely to be on multiple servers and the data is copied in parallel.
- **Automatic recovery from storage failures:** This allows the system to recover from failures that affect resiliency or redundancy. Exchange 2013 includes recovery behaviors for long I/O times, excessive memory consumption by the Microsoft Exchange Replication service (MSEExchangeRepl.exe), and also for severe cases in which the system is in such a bad state that threads cannot be scheduled.
- **DAG lagged copy enhancements:** Lagged copies can now care for themselves to a certain extent using automatic log play down. In addition, lagged copies can leverage Safety.Net (previously Transport Dumpster in Exchange 2010), making recovery or activation much easier.





## 2 Configuration

### 2.1 Disk layout

Microsoft provides the following storage configuration best practices for Exchange 2016:

Redundant Array of Independent Disks (RAID) is often used to both improve the performance characteristics of individual disks (by striping data across several disks) as well as to provide protection from individual disk failures. With the advancements in Exchange 2016 high availability, RAID is not a required component for Exchange 2016 storage design. However, RAID is still an essential component of Exchange 2016 storage design for standalone servers as well as solutions that require storage fault tolerance.<sup>1</sup>

Determining the storage layout before the installation of Microsoft Exchange Server is an important step as it can have direct impact on performance when using other disk solutions.

With Exchange Server 2016, due to the reduced IOPS required, Microsoft allows placement of logs and databases on the same volume for DAG-protected databases. The Jet database (EDB) activity resembles sequential reading and writing to 32 KB blocks. The transaction logs see 100 percent sequential reads and writes.

The following table shows a sample disk layout based on best practices.

Drive	Recommended configuration	Contents
C:	DAS/SAN	Windows, Exchange binaries
D:	DAS/SAN	Pagefile
E:	SAN	Database 1 / Logs 1
F:	SAN	Database 2 / Logs 2

When using Exchange Server 2016 Database Availability Groups (DAGs) to create a highly resilient database infrastructure, Microsoft "Preferred Architecture" guidance discusses distributing three copies plus a lagged copy across DAG members, and utilizing the same spindle on each of four servers to host an active copy, as well as two passive copies and a lagged copy of each of the four server's databases.<sup>1</sup> Microsoft design guidance is specifically for JBOD (non-SAN) environments where larger, slower single spindles are

---

<sup>1</sup> ["The Exchange 2016 Preferred Architecture"](#) by Ross Smith IV, dated October 12, 2015



used to provide storage for Exchange databases. Therefore, this Microsoft guidance does not apply for Dell SC Series SAN volumes. This is due to the following reasons:

1. Because RAID is inherent in SC Series arrays, 3-copy DAGs are not an absolute requirement for resilience. 2-copy DAGs are more common.
2. When taking volume Replays with Replay Manager, an additional lagged database copy is not needed.
3. The risk of database reseeding is very low when compared to a single disk JBOD. If a disk drive in an SC Series array fails, a SAN spare is automatically used and the server and application are unaware of any hardware change.
4. The main reason Microsoft best practices puts these multiple databases on a disk is to allow for faster reseeding of the storage space from multiple sources. Since this is not a factor for SC Series volumes, Dell does not recommend this as a best practice.

Exchange 2016 supports 100 databases per server with up to 16 copies of each to other mailbox servers in a DAG organization. A best practice is to minimize the number of databases, as few as recovery objectives will allow. As the number of databases increases, so does the I/O required supporting the additional data streams. This can have a negative impact on the I/O load of a system.

As environments grow larger and larger, it becomes common to run out of drive letters for the volumes. For the purpose of scalability, it may be suitable to use Windows mount points (VMP) for database and log volumes.

The following table shows a sample disk layout based on best practices using mount points.

Drive	Recommended configuration	Contents
C:	DAS/SAN	Windows, Exchange binaries
D:	DAS (if available)	Pagefile
C:\DB	SAN	Mount points
VMP	C:\DB\Database1	Database 1 / Database 1 logs
VMP	C:\DB\Database2	Database 2 / Database 2 logs
VMP	C:\DB\Database3	Database 3 / Database 3 logs

## 2.2 Storage Center page size considerations

By default, the Dell Storage Center writes data to individual 2MB pages. Exchange 2016 writes database pages of 32KB each. Write concatenation will attempt to put as many I/Os together to fill a page where possible.

It is very likely that during the course of a day, many pieces of the database will be accessed and changed. That being considered, it is also important to understand that a single block being altered constitutes a



change of that 2MB page where it is located on the Storage Center. Potentially, that could be the only block on that 2MB page that was changed. This can result in larger-than-expected Replay sizes.

Dell SC Series customers who experience large Replays (75 percent or greater of database size) can consider the 512K page option (not available on the SCv2xxx Series array) for their Exchange volumes on the Storage Center. This will reduce the amount of blocks that can write to a page, thus reducing the Replay sizes. The tradeoff with this smaller page size is reduced storage capacity.

## 2.3 Storage Profiles and Data Progression

Storage Profiles were introduced in Storage Center 4.2. For SC Series models, the **Recommended** storage profile will provide the best Data Progression configuration for most implementations. Essentially, this storage profile will write new data to RAID10 with Replay data being placed on RAID5–9.

For SCv2xxx models, either the **Balanced** or **Maximize Efficiency** Storage Profiles are recommended. Maximize Efficiency provides the most available storage capacity using RAID5/6, while Balanced provides RAID10 performance on all writes and RAID5/6 for Replays. SCv2xxx models only provide Data Progression between RAID levels when using Balanced.

### 2.3.1 Microsoft recommends ReFS for Exchange 2016 database volumes (JBOD)

This recommendation is for JBOD and not SAN volumes, and should not be used for Exchange 2016 database and log volumes on SC Series arrays: "Each disk that houses an Exchange database is formatted with ReFS (with the integrity feature disabled) and the DAG is configured such that AutoReseed formats the disks with ReFS. " – From published Microsoft guidance: [The Exchange 2016 Preferred Architecture](#)

### 2.3.2 Default allocation unit size

Exchange reads and writes to the database in 32KB chunks. Formatting volumes with the correct default allocation unit size is also an important best practice. The default allocation unit size of 64KB is recommended.

### 2.3.3 Partition type

GPT (GUID Partition Table) is recommended for Exchange 2016 volumes. Volumes larger than 2TB will require that GPT is used to access the entire size of the volume. MBR (Master Boot Record), the default partition type, is also supported but the partition size is limited to less than 2TB.

## 2.4 Mailbox server sizing considerations

Understanding user patterns is an important part of determining sizing parameters for any Exchange Server.

Specifically, tools like Jetstress, Profile Analyzer, and LoadGen (described below) can help determine the best settings for the type of mailbox user, mailbox sizes, and server configuration. By having information such as mailbox count, required mailbox size, send/receive statistics, and average message size, Jetstress



can use these parameters to run a performance test to simulate how the specific configuration would perform.

Servers should be sized before they are moved into production. This process should include a Jetstress test to determine if the hardware and disk configurations are suitable for the type and amount of traffic to be expected. Once a successful Jetstress test is completed, Exchange Server can be installed and client testing can be performed using LoadGen. LoadGen simulates actual client traffic operations such as sending and receiving mail, creating calendar appointments, and accessing public folder data. A number of other tests are available as part of the profile configuration and can be customized as necessary.

Mailbox I/O requirements substantially lowered since Exchange 2013. The overall I/O requirement for Exchange 2016 mailboxes has decreased by over 50 percent from Exchange 2010.




Figure 2 Change in IOPS per mailbox requirements by version of Exchange Server

Microsoft provides the following information on simulation and analysis tools for Exchange (links provided):

**Jetstress 2013** should be used to simulate disk I/O load on a test server running Exchange 2013 or 2016 to verify the performance and stability of the disk subsystem before putting the server into a production environment. Find more information on [Microsoft.com](https://www.microsoft.com/jetstress).

**Exchange Load Generator** is a simulation tool used to measure the impact of MAPI, OWA, ActiveSync, IMAP, POP, and SMTP clients on Exchange 2013–2016 servers. Find more information on the page, [Exchange Load Generator 2013 \(64 bit\)](https://www.microsoft.com/exchangeloadgenerator).

**Exchange Profile Analyzer** collects estimated statistical information from a single mailbox store or across an Exchange Server 2007 or 2010 organization. The collected data can be used to analyze the performance and sizing of a server with mailboxes that will be migrated to Exchange 2016. [Microsoft Exchange Server Profile Analyzer \(64 bit\)](https://www.microsoft.com/exchange-profile-analyzer).

**Exchange Server Role Requirements Calculator**, a spreadsheet-based calculator, is very helpful in designing the best environment (both storage and hardware) based on criteria including:

- User profile: message profile, mailbox size, and number of users
- High availability architecture: number of database copies planned to deploy, whether the solution will be site resilient or not, and desired number of mailbox servers
- Server CPU platform
- Storage architecture: disk capacity/type and storage solution
- Backup architecture: whether to use hardware/software VSS and specify the frequency of the backups, or leverage the Exchange native data protection features
- Network architecture: utilization, throughput, and latency aspects

Download the [Exchange Server Role Requirements Calculator](#).

### 2.4.1 Archive mailboxes

Each mailbox in the database can have an archive mailbox. This allows users to have an archive repository that does not impact existing mailbox quotas and allows them to retain data for longer periods of time. Archive mailboxes can help eliminate PST files and keep information all contained within Exchange Server. It is important to note that an archive mailbox is a separate object from the user mailbox.

Beginning with Exchange 2010 Service Pack 1 (SP1), it is possible to provision a user's personal archive on the same mailbox database as the user's primary mailbox, on another mailbox database on the same mailbox server, or on a mailbox database on another mailbox server in the same Active Directory site. This feature offers the flexibility to utilize SC Series tiered storage architecture to store archive mailboxes on a separate volume utilizing a tier-3 RAID 5 storage profile, putting this archive data on less expensive disks. Instant Replays can be done much less frequently as this data changes less frequently.

**Note:** Archive mailboxes are only available while a user is online. They are not available if a user is running in cached mode and is not connected to Exchange.

## 2.5 Online database maintenance

Database maintenance for Exchange 2013/2016 has also been modified to work more efficiently. One of these changes regards background database maintenance (BDM), which is now throttled back from 5MB-per-sec/copy to 1MB-per-sec/copy.

In addition, the updates to ESE can reduce the cost of maintaining and managing a database. Database maintenance is comprised of storage mailbox maintenance and ESE database maintenance.

### 2.5.1 Online defragmentation

Online defragmentation now runs 24x7 in the background by default. There are no configurable settings with the default feature. Exchange monitors the database as it is being used and small changes are made over time to keep it defragmented for space and contiguity. Online defragmentation is also throttled so that it does not negatively impact client performance.



## 2.5.2 Online database scanning (checksumming)

The checksumming process can run in two different modes on the active database copies:

**Default option:** This has checksumming run in the background 24x7. This option is intended for most databases which can require longer periods of time to complete a checksum. Exchange scans the full database no more than once a day and will generate an alert if it does not finish scanning within seven days.

**Second option:** This runs as the last task in the custom-scheduled mailbox database maintenance process. The amount of time it runs can be configured by altering the schedule calendar, which is recommended for databases that are less than 500GB in size. This allows these databases to finish scanning quickly, having less impact to the continuous BDM process. The schedule can be modified in the database properties under the maintenance section.

For more information, see [technet.microsoft.com](http://technet.microsoft.com).

**Note:** Online database maintenance can have a major impact on replay sizes as the online defragmentation process moves data around and pages are cleared. Although disabling online database maintenance completely is not recommended, some customers have experienced relief in specific extreme situations by throttling the online database maintenance process back to one day per week.

## 2.5.3 Database fragmentation

Defragmentation of the information storage database occurs after information within the database is deleted, leaving a data-free page still within the file. Defragmentation has performance implications if not addressed properly, so it is advantageous to use the built-in information store maintenance.

Fragmentation is limited and contiguity is maintained through the background processes that run to manage the databases.



## 3 Performance and monitoring

### 3.1 Monitoring

Performance Monitor is a helpful tool for monitoring the overall operations of Exchange server. A number of specific counters are good indicators of performance bottlenecks. The following performance counters are helpful in monitoring and troubleshooting. The thresholds provided are based on standard industry guidance.

By default, the real-time view of Performance Monitor shows the last 1:40 (one minute, 40 seconds), measuring every two seconds and taking the average of the current measurement and last measurement.

If a performance issue is suspected, it is a good idea to set up Performance Monitor to record key counters during the period of peak usage when the problem occurs.

Physical disk counters	Memory counters	Exchange counters
Avg. Queue Length	Available MBytes	MSExchangeIS Store – RPC Average Latency
Avg. Disk Sec/Transfer	Pages/sec	Database - Database Cache % Hits
Disk Transfers/sec	Processor counters	
% Disk Idle	% Processor Time	
Avg. Disk Bytes/Transfer	Network counters	
Avg. Disk Sec/Read	Bytes Total/sec	
Avg. Disk Sec/Write		

#### 3.1.1 Monitoring the database defragmentation process

In Microsoft Exchange Server 2016, the following performance counters for monitoring the behavior of database defragmentation have been added for use with Performance Monitor:

- **Database ==> Instances \ Defragmentation tasks:** Shows the background database defragmentation tasks that are currently executing
- **Database ==> Defragmentation Tasks Pending:** Shows the background database defragmentation tasks that are currently pending
- **Database ==> Instances \ Database Maintenance Duration:** the number of hours that have passed since maintenance last completed for this database
- **Database ==> Instances \ Database Maintenance Pages Bad Checksums:** the number of non-correctable page checksums encountered during a database maintenance pass



### 3.1.2 Monitoring with System Center Operations Manager

Dell has developed a System Center Operations Manager (SCOM) Management Pack that provides the capabilities to manage the Storage Center within a Windows environment. There is also a custom Management Pack available for Exchange 2013, which uses Exchange “Managed Availability” to also run with Exchange 2016.

For more information on the SCOM Management Pack for Storage Center, see the [SC Series support site](#) (authentication required).





## 4 Troubleshooting

### 4.1 Enterprise Manager/Dell Storage Manager

Use Dell Enterprise Manager/Dell Storage Manager to monitor, troubleshoot and manage Exchange workloads along with other application workloads on Dell Storage systems in your organization. Dell SC Series has had the Enterprise Manager services for many years, and recently has been re-launched as Dell Storage Manager to manage all Dell Storage offerings.




Figure 3 Dell Enterprise Manager charting of running application workloads

### 4.2 Microsoft Exchange 2013/2016 Diagnostic Service logs

Since the release of Exchange Server 2013 version CU6, the Exchange Diagnostic Service constantly records Exchange performance counter information and stores them in local log files. These log files are part of the new Exchange Managed Availability services. By default, seven days of these logs are stored, taking up 5GB (5120MB) of disk space per day on the location where Exchange binaries are installed.

To store these log files, either ensure that there is enough capacity or edit the **Microsoft.Exchange.Diagnostics.Service.exe.config** file as follows to conform to the capacity available:

```
<add Name="DailyPerformanceLogs" LogDataLoss="True"
      MaxSize="5120" MaxSizeDatacenter="2048" MaxAge="7.00:00:00"
      CheckInterval="08:00:00" />
```

Set the MaxSize and MaxAge parameters appropriate to the capacity available on the disk. Exchange will delete the data as soon as it reaches the first of those two values. Microsoft Support recommends leaving these settings at the default to collect the appropriate logging required for detailed troubleshooting operations.

For reference information on setting up and working with this new troubleshooting tool, see the [WindowsITPro article](#).

## 4.3 Microsoft Exchange Troubleshooting Assistant (2007–2010)

The Exchange Troubleshooting Assistant is an integrated component of the Exchange Management Console Toolbox in Exchange Server 2007–2010. It programmatically executes a set of troubleshooting steps to identify the root cause of performance, mail flow, and database mounting issues.

This tool automatically determines what set of data is required to troubleshoot the identified symptoms and collects configuration data, performance counters, event logs, and live tracing information from an Exchange server and other appropriate sources. It also analyzes each subsystem to determine individual bottlenecks and component failures, and then aggregates the information to provide root cause analysis. The output of this report can assist customers and their hardware vendors in troubleshooting Exchange-related problems prior to migration to Exchange Server 2013 or greater.

## 4.4 Microsoft Exchange Best Practices Analyzer (2007–2010)

The Exchange Best Practices Analyzer is available with Exchange Server 2007–2010. It is installed during Exchange Setup and can be run from the Exchange Management Console Toolbox.

This tool programmatically collects settings and values from data repositories such as Active Directory, registry, metabase, and performance monitor. Once collected, a set of comprehensive, best-practice rules are applied to the topology. Administrators running this tool will get a detailed report listing the recommendations that could be made to the environment to achieve greater performance, scalability and uptime.

It is a recommended best practice to run this in the existing Exchange 2007 or 2010 environment prior to moving to Exchange Server 2013 to have the best migration experience.

## 5 Backup and recovery

### 5.1 Backups and Instant Replays

#### 5.1.1 Backup schedules

A standard backup schedule should be established for every Exchange environment. This backup schedule should include at least one full backup per week. If an environment is small, or if time permits, a daily full backup would be optimal.

With Exchange 2013/2016, Microsoft recommends using database copies as a means to eliminate regular backups. By having at least three copies of each mailbox database, Microsoft considers these to be highly available. In addition, Microsoft recommends having several lag copies of each database available in case there is a need for a point-in-time copy (logical corruption).

#### 5.1.2 Data Instant Replays (snapshots)

SC Series Data Instant Replays with Replay Manager are recommended for use with Microsoft Servers. Replays make an easy means to recover individual mailboxes or entire databases very quickly. However, they are not intended to replace regular backup processes because some verticals have specific regulatory requirements relating to data retention.

Business requirements will dictate the backup schedules to be utilized. As a rule of thumb, Exchange VSS snapshots using Replay Manager Services for Microsoft Servers can be taken in 15-minute increments. Beyond that, the options are limitless concerning the flexibility of snapshot schedules. It is important to note that if the database consistency check will run with the backup set, it may not be possible to complete the snapshot and consistency check before the next scheduled job runs when scheduling around the 15-minute window.

By design, VSS data has to be captured within 10 seconds or the snapshot will fail. Since Exchange 2013/2016 allows databases and logs to reside on the same volume, there are fewer volumes that need to be part of a backup set. It is recommended to include no more than 12 mailbox databases per VSS backup set. This will allow the snapshot to complete within the 10-second window.

#### 5.1.3 Database verification

To run the database verification process as required by Microsoft, Replay Manager for Microsoft Exchange Server provides a verification service that can be used to mount the database Replay after a successful snapshot to a utility server. The database verification process is a time-, processor-, and memory-intensive operation. It is a critical part of the backup process, but due to the amount of time that it can take, customers who do frequent database snapshots will typically only run verification on the database(s) once a day to check for consistency.

## 6 Recovery

### 6.1 Recovery databases

A recovery database is an Exchange 2013/2016 feature that replaces the Recovery Storage Group (RSG) found in previous versions of Exchange. A recovery database is a special kind of mailbox database that allows mounting a restored mailbox database (using a Replay) and extracting data from the restored database as part of a recovery operation.

The Restore-Mailbox cmdlet can be used to extract data from a recovery database. After extraction, the data can be exported to a folder or merged into an existing mailbox. Recovery databases enable recovery of data from a backup or copy of a database without disturbing user access to current data.

Before using a recovery database, there are certain requirements that must be met. For instance, a recovery database can be used for Exchange 2016/2013/2010 mailbox databases only.

**Note:** The target mailbox used for data merges and extraction must be in the same Active Directory forest as the database mounted in the recovery database.

#### 6.1.1 Single item recovery

Microsoft provides the following information on single item recovery for Exchange 2013/2016.<sup>2</sup>

To protect from accidental or malicious deletion and to facilitate discovery efforts commonly undertaken before or during litigation or investigations, Microsoft Exchange Server 2013/2016 uses the Recoverable Items folder. The Recoverable Items folder replaces the feature known as the dumpster in Exchange Server 2007. The Recoverable Items folder is used by the following Exchange features:

- Deleted item retention
- Single item recovery
- In-Place Hold
- Litigation hold
- Mailbox audit logging
- Calendar logging

For more information on single item recovery, see the [Exchange help article](#).

---

<sup>2</sup> "Recoverable Items folder." Exchange 2013 Help. February 27, 2015. [http://technet.microsoft.com/en-us/library/ee364755\(v=exchg.150\).aspx](http://technet.microsoft.com/en-us/library/ee364755(v=exchg.150).aspx).



## 7 Disaster recovery

VSS replays can be replicated using the Storage Center Remote Data Instant Replay. This would allow data (as current as the last Replay available) to be brought online at the main site or disaster recovery site.

If Exchange Server boots from SAN, Replays from the boot volume and data volumes can be easily mapped at the disaster recovery site, provided that duplicate hardware is used at the site. In order for the operating system to boot properly, the volumes must be mapped to hardware that matches the original server hardware at the main site. Additional steps for recovery beyond the scope of this document would have to be taken for systems with dissimilar hardware at the main site and disaster recovery site.

It is important to understand the overall disaster recovery strategy around core services such as directory services and DNS. Exchange relies heavily on Active Directory being available as well as DNS services. Preparations must be taken to ensure that these dependencies are online and operational at the disaster recovery site before the Exchange Server and its services can come online. Although Exchange data will be available, it is equally important that the infrastructure is in place to allow for a smooth cutover, allowing users to access Exchange soon after a disaster occurs.

### 7.1 Database copies and database availability groups

In the event of a hardware or software failure, multiple database copies in a database availability group (DAG) enable high availability with fast failover and no data loss. This eliminates end-user downtime and lost productivity that make up a significant cost of recovering from a past point-in-time backup to disk or tape. DAGs can be extended to multiple sites and can provide resilience against data center failures as well.

While all Mailbox servers in a DAG must be in the same Active Directory domain, up to 16 copies of an Exchange 2016 mailbox database can be created on multiple Mailbox servers, provided the servers are grouped into a database availability group (DAG). However, the round-trip network latency must not be greater than 250 milliseconds (ms).

## A Technical support and resources

[Dell.com/support](http://Dell.com/support) is focused on meeting customer needs with proven services and support.

[Dell TechCenter](#) is an online technical community where IT professionals have access to numerous resources for Dell software, hardware and services.

[Storage Solutions Technical Documents](#) on Dell TechCenter provide expertise that helps to ensure customer success on Dell Storage platforms.

### A.1 Related documentation

Table 1 lists the referenced or recommended resources related to this document.

Table 1 Referenced or recommended resources

Vendor	Resource
Dell	<a href="#">Sizing and Best Practices for Deploying Microsoft Exchange Server 2013 with Dell Compellent Storage Arrays</a> - (applies to Exchange Server 2016 as well)
Dell	<a href="#">Dell Compellent Replay Manager 7.0 and Exchange 2013 Demo Video</a>
Dell	<i>Compellent Replay Manager Best Practices Guide</i> available in the <i>Replay Manager Admin Guide</i> on the <a href="#">Knowledge Center</a>
Microsoft	<a href="#">Exchange Team Blog</a>
Microsoft	<a href="#">Exchange Server 2013</a> documentation
Microsoft	<a href="#">Exchange Server 2016</a> documentation

