

# Dell EMC PowerVault ME4 Series and Microsoft Exchange Server 2016

## Abstract

This document provides best practices for Microsoft® Exchange Server 2016 when using a Dell EMC™ PowerVault™ ME4 Series storage array.

September 2018

## Revisions

Date	Description
September 2018	Initial release

## Acknowledgements

Author: Damon Zaylskie

The information in this publication is provided “as is.” Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

© 2018 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.

Dell believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

# Table of contents

Revisions.....	2
Acknowledgements.....	2
Executive summary.....	4
<b>1 Predeployment.....</b>	<b>5</b>
1.1 Storage virtualization options for ME4 Series .....	5
1.2 Understanding Exchange I/O .....	5
1.3 Exchange 2016 architecture changes .....	5
<b>2 Configuration.....</b>	<b>7</b>
2.1 Disk layout .....	7
2.2 Disk groups and pools .....	8
2.3 SSD read cache.....	9
2.4 Data tiering .....	9
2.5 Resilient File System for Exchange 2016 database volumes (JBOD) .....	9
2.6 Default allocation unit size.....	9
2.7 Partition type.....	10
2.8 Mailbox server sizing considerations.....	10
2.8.1 Archive mailboxes.....	11
2.9 Online database maintenance.....	11
2.9.1 Online defragmentation .....	11
2.9.2 Online database scanning (checksumming) .....	12
2.9.3 Database fragmentation .....	12
<b>3 Exchange performance monitoring.....</b>	<b>13</b>
3.1 Monitoring the database defragmentation process .....	13
<b>4 Troubleshooting.....</b>	<b>14</b>
4.1 ME4 Series performance monitoring.....	14
4.2 Exchange Server 2013/2016 Diagnostic Service logs .....	14
4.3 Microsoft Exchange Troubleshooting Assistant (2007–2010).....	15
4.4 Microsoft Exchange Best Practices Analyzer (2007–2010) .....	15
<b>5 Recovery databases.....</b>	<b>16</b>
5.1 Single item recovery .....	16
<b>6 Disaster recovery with database copies and DAGs.....</b>	<b>17</b>
<b>A Technical support and resources.....</b>	<b>18</b>
A.1 Related documentation.....	18

## Executive summary

This document provides best practices for Microsoft® Exchange Server 2016 when using a Dell EMC™ PowerVault™ ME4 Series storage array. It does not include sizing, performance, or design guidance, but provides information about the features and benefits of using ME4 Series arrays for Exchange data storage.

# 1 Predeployment

## 1.1 Storage virtualization options for ME4 Series

Dell EMC PowerVault ME4 Series arrays support a variety of options for storage configurations and can improve the performance and availability of storage for Microsoft Exchange Server. These options are detailed in this paper to help you understand the benefits of each.

ME4 Series storage also supports several connectivity options which are discussed in this paper but not from a performance perspective. The scope of this document does not include sizing or performance data. The storage design will be driven by business requirements for performance and availability.

## 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 section 2.8, 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 section 3, Exchange performance monitoring.

## 1.3 Exchange 2016 architecture changes

With the decreasing cost of CPU hardware, the constraint of expensive server hardware has been alleviated. Exchange Server 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 server and Edge Transport server.

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

**Client Access services:** These provide authentication, limited re-direction, and proxy services offering the usual client access protocols: HTTP, POP, IMAP and SMTP.

**Mailbox services:** These include 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 introduced with Exchange Server 2016 that are notable for storage considerations include the following:

**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 5,000 mailboxes to 10,000 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 the following:

**Reduced IOPS compared to Exchange Server 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 Server 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 Server 2010), making recovery or activation much easier.

## 2 Configuration

### 2.1 Disk layout

Microsoft provides the following storage configuration best practices for Exchange 2016 in the TechNet blog, [The Exchange 2016 Preferred Architecture](#).

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.

Determining the storage layout before the installation of Microsoft Exchange Server is an important step because 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.

Table 1 shows a sample disk layout based on best practices.

Table 1 Sample disk layout

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

When using Exchange Server 2016 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 drive location 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 drives are used to provide storage for Exchange databases. Therefore, this Microsoft guidance may not apply for ME Series SAN volumes. This is due to the following reasons:

- ME4 Series arrays support a combination of HDD and SSD drives. This can improve performance by using a mix of SSD and HDD drives and tiering data.
- ME4 Series arrays support a large number of RAID configurations for flexibility to help choose a balance between redundancy and utilization efficiency. The number of drives required can be greatly reduced when using a more efficient RAID technology such as ADAPT RAID.
- Because each volume can be striped across multiple drives in a pool, a single volume can greatly exceed the performance and capacity of a single drive. Therefore, multiple databases can be created on a single volume, which can simplify management.

- Disk sparing is performed at the array level. This limits contact and access requirements for servers. All drives are hot-pluggable and will automatically spare and rebuild in case of replacement or failure. Spares are shared among disk groups to reduce the number of spares required.
- ME4 Series arrays support SAS, Fibre Channel, and iSCSI in the same platform. Based on requirements, one or more of these may be used to access the same volumes. This can provide greater accessibility and performance to any volume than by using local direct attach to a single disk, and can accelerate backups.
- ME4 Series arrays can host multiple servers on the same array. If choosing to share the drives with other applications, this can be done quickly and easily.

Exchange Server 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 to 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 for database and log volumes.

Table 2 shows a sample disk layout based on best practices using mount points.

Table 2 Sample disk layout 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 and Database 1 logs
VMP	C:\DB\Database2	Database 2 and Database 2 logs
VMP	C:\DB\Database3	Database 3 and Database 3 logs

## 2.2 Disk groups and pools

The design of the ME4 Series array is based around groups of disks. There are two types of storage available in the ME4 Series array: virtual and linear storage.

Linear storage maps directly to a single drive or multiple disks, but has less flexibility when it comes to making changes.

The newer method is virtual storage. Virtual storage starts with disk groups of the chosen RAID type. The groups of disks are added to a pool, and the volumes are then created from these pools. This allows groups of disks to change in number and scale larger, providing potential for growth with less effort. Existing groups and pools can be expanded.

By design, an ME4 Series system in virtual storage mode will have at least one storage pool. In systems with more than 12 drives, it may have two pools, one for each controller. This is based on how the ME4 Series array handles volume management. In a two-copy DAG configuration, this can provide I/O isolation if sufficient drives are available.



At a minimum, the database and log volumes should be balanced across the pools (see Table 3). Each pool should contain a relatively equal number of log and database copies. This will help balance the I/O across the storage. Using the built-in performance monitoring will provide insight into how the storage traffic is balanced.

Table 3 Balanced pool placement of database and log volumes

Controller A	Controller B
Database1&Logs1	Database2&Logs2
Database3&Logs3	Database4&Logs4

## 2.3 SSD read cache

When SSD drives are added to an HDD pool, they can be used as a read cache. When the drives are configured as a read-cache disk group, they will intelligently cache pages from the entire pool.

## 2.4 Data tiering

ME4 Series storage supports data tiering when using virtual storage.

Refer to the following information from the ME4 Series *Administrator's Guide* (available on [Dell.com/support](http://Dell.com/support)):

Automated tiered storage (ATS) is a virtual storage feature that automatically moves data residing in one class of disks to a more appropriate class of disks based on data access patterns. Frequently accessed, hot data can move to disks with higher performance, while infrequently accessed, cool data can move to disks with higher capacity, lower performance, and lower costs.

Each virtual disk group, depending on the type of disks it uses, is automatically assigned to one of the following tiers:

- **Performance:** This highest tier uses SSDs.
- **Standard:** This middle tier uses enterprise-class HDD SAS disks, which provide good performance with mid-level cost and capacity.
- **Archive:** This lowest tier uses midline HDD SAS disks, which provide the lowest performance with the lowest cost and highest capacity.

## 2.5 Resilient File System for Exchange 2016 database volumes (JBOD)

In [The Exchange 2016 Preferred Architecture](#), Microsoft recommends using Resilient File System (ReFS) for JBOD volumes and not SAN volumes. ReFS should not be used for Exchange 2016 database and log volumes on ME4 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 Auto-Reseed formats the disks with ReFS.

## 2.6 Default allocation unit size

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

## 2.7 Partition type

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

## 2.8 Mailbox server sizing considerations

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

Tools like Jetstress, Profile Analyzer, and LoadGen (described in this section) can help determine the best settings for the type of mailbox user, mailbox size, 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 have been substantially lowered since Exchange 2013. The overall I/O requirement for Exchange 2016 mailboxes has decreased by over 50 percent from Exchange 2010 (see Figure 1).

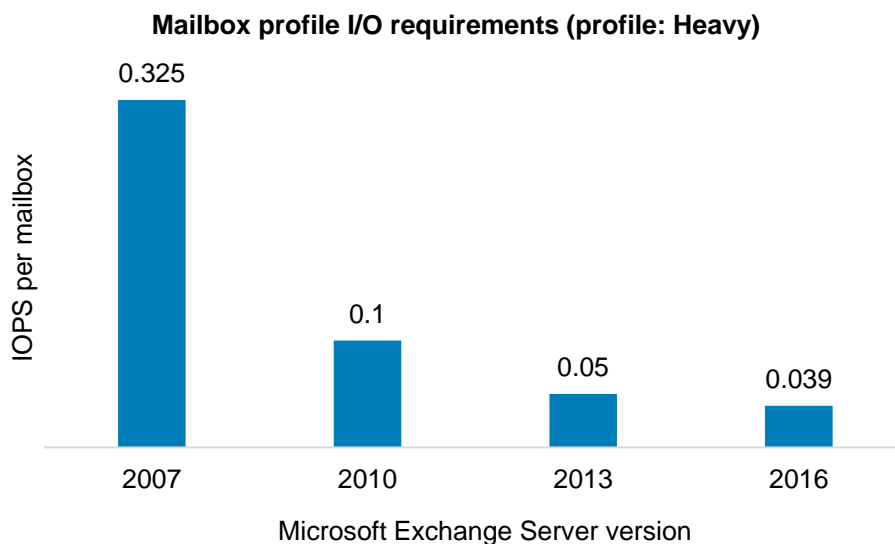


Figure 1 IOPS-per-mailbox requirement per Exchange Server version

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).

**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 \(64 bit\)](#).

The [Exchange Server Role Requirements Calculator](#), a spreadsheet-based calculator, is very helpful in designing the best environment (both storage and servers) 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: core count and clock rate
- 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

### 2.8.1 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 the ME4 Series tiered storage architecture to store archive mailboxes on a separate volume utilizing less expensive disks. Performance for archive mailboxes may have different requirements that will allow larger, cheaper drives to be used.

---

**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.

---

## 2.9 Online database maintenance

Database maintenance for Exchange Server 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 5 MB-per-sec/copy to 1 MB-per-sec/copy.

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

### 2.9.1 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.

## 2.9.2 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 time periods 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 500 GB 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).

## 2.9.3 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 Exchange performance 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 performance counters listed in Table 4 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 (1 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.

Table 4 Performance counters in Performance Monitor

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 Monitoring the database defragmentation process

In 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:** This shows the background database defragmentation tasks that are currently executing.

**Database ==> Defragmentation Tasks Pending:** This shows the background database defragmentation tasks that are currently pending

**Database ==> Instances \ Database Maintenance Duration:** This shows the number of hours that have passed since maintenance last completed for this database.

**Database ==> Instances \ Database Maintenance Pages Bad Checksums:** This shows the number of non-correctable page checksums encountered during a database maintenance pass.

## 4 Troubleshooting

### 4.1 ME4 Series performance monitoring

The ME4 Series arrays have an interactive performance monitor built-in. It can display data throughput at any level, from a single disk to a controller. Disk groups and pools can be monitored as well. This tool can be used to show volume performance and throughput for troubleshooting or monitoring purposes. Capturing a baseline can help determine changes in disk usage over time.

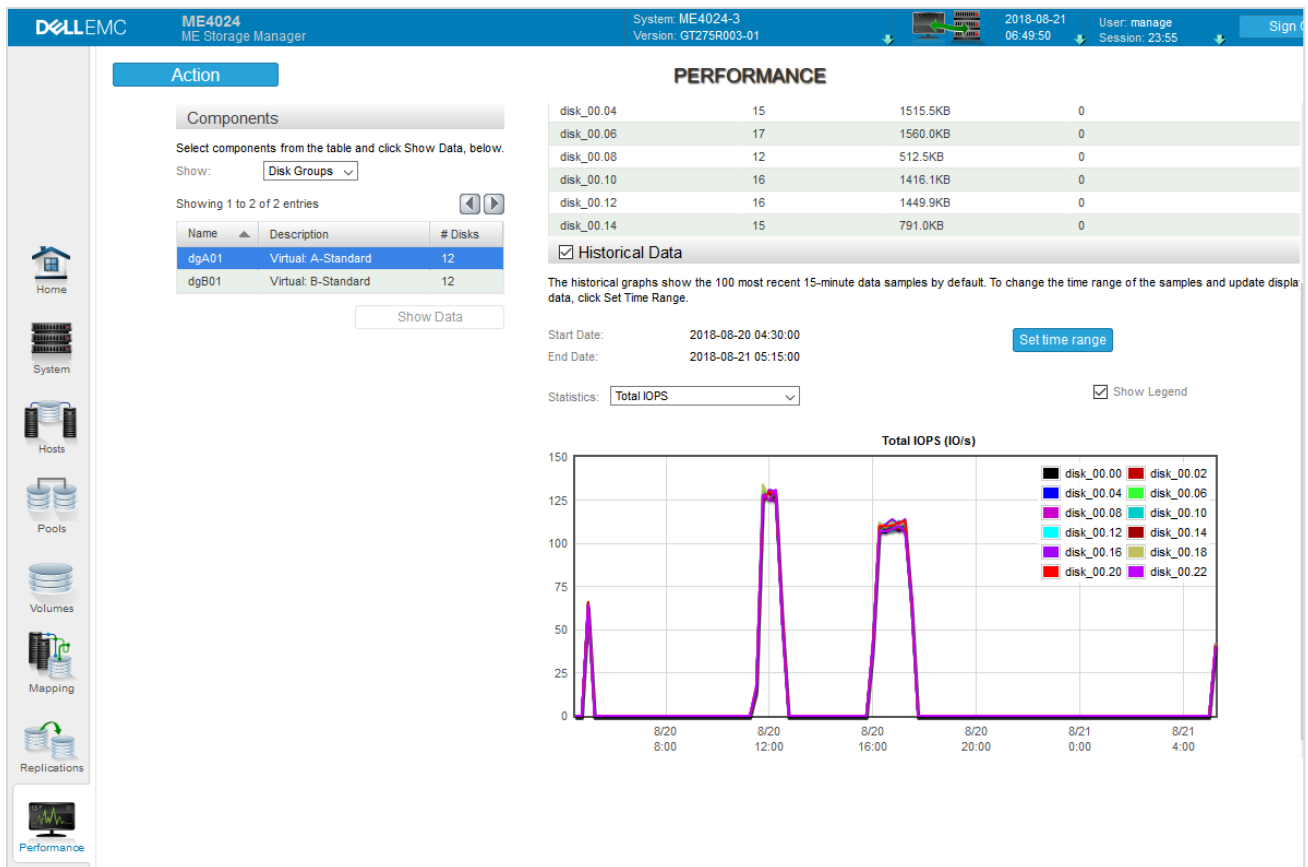


Figure 2 ME4 Series performance monitoring

### 4.2 Exchange Server 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 Exchange Managed Availability services. By default, seven days of these logs are stored, taking up 5 GB (5,120 MB) 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 appropriately 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 troubleshooting tool, see the ITPro Today article, [Quickly Get the Data You Need for Troubleshooting Exchange Server 2013 Problems](#).

### 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, the registry, the metabase, and the 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 best practice to run this analyzer in the existing Exchange 2007 or 2010 environment prior to moving to Exchange Server 2013 to ensure the best migration experience.

## 5 Recovery databases

A recovery database is an Exchange Server 2013/2016 feature that replaces the Recovery Storage Group (RSG) found in previous versions of Exchange Server. A recovery database is a special kind of mailbox database that allows mounting a restored mailbox database 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 Server 2010/2013/2016 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.

---

### 5.1 Single item recovery

Microsoft provides the following information on single item recovery for Exchange 2013/2016 in the TechNet article, [Recoverable Items folder](#).

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



## 6 Disaster recovery with database copies and DAGs

In the event of a hardware or software failure, multiple database copies in a 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 Server 2016 mailbox database can be created on multiple Mailbox servers, provided the servers are grouped into a DAG. However, the round-trip network latency must not be greater than 250 milliseconds.

ME4 Series storage can provide excellent scale and performance, which enables additional ME4 Series arrays to host DAG copies and local or remote locations to aid in data recovery.

## A Technical support and resources

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

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

### A.1 Related documentation

The following ME4 Series publications and additional resources are available at [Dell.com/support](https://dell.com/support).

- Administrator's Guide
- Deployment Guide
- CLI Guide
- Owner's Manual
- Support Matrix

Additionally, see the following Microsoft publications and articles:

- [Exchange Team Blog](#)
- [Exchange Server 2013](#) documentation
- [Exchange Server 2016](#) documentation