

Dell EMC Unity Storage with Microsoft Exchange Server

All-flash arrays

[Abstract](#)

This white paper provides best practices guidance for configuring Microsoft Exchange Server to perform optimally with Dell EMC Unity All-Flash arrays.

October 2020

Revisions

Date	Description
July 2017	Initial release for Dell EMC Unity OE Version 4.2
October 2020	Remove reference to Dell EMC Storage Integrator (ESI) as it is end-of-life

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.

© 2020 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
Table of contents	3
Executive summary.....	4
Audience	4
1 Storage configuration.....	5
1.1 Storage pools.....	5
2 Dell EMC Unity features.....	7
2.1 FAST VP.....	7
2.2 FAST Cache	7
2.3 Compression.....	7
2.4 Snapshot best practices	8
2.5 Data at Rest Encryption (D@RE)	8
2.6 Host I/O limits.....	8
3 Microsoft Exchange considerations.....	9
3.1 Understanding Exchange I/O.....	9
3.2 Exchange 2016 architecture changes	9
3.3 Background database maintenance (BDM).....	10
3.4 Exchange storage layout	10
4 Performance monitoring and troubleshooting	14
4.1 Monitoring	14
4.2 Troubleshooting	15
4.3 Microsoft Exchange 2013/2016 Diagnostic Service logs	16
5 Data protection - AppSync	17
5.1 Snapshots using AppSync for Exchange	17
6 Recovery	19
6.1 Using AppSync snapshots.....	19
6.2 Recovery databases.....	20
7 Disaster recovery.....	21
7.1 Database copies and database availability groups	21
A Technical support and resources	22
A.1 Related resources	22

Executive summary

This solution guide delivers straightforward guidance to customers using Dell EMC™ Unity storage systems in a Microsoft® Exchange Server® 2016 environment. The focus is on system performance and reliability while maximizing the ease of use of the automated storage features when used with Exchange Server. For general best practices on using Dell EMC Unity, please see the *Dell EMC Unity Best Practices Guide*.

These guidelines are intended to cover specific Exchange Server 2016 use cases. They are strongly recommended by Dell EMC; however, they are not strictly required.

This paper was developed using the Dell EMC Unity 450F All-Flash array, but is also applicable when using the 350F, 550F, or 650F Dell EMC Unity All-Flash arrays.

If you have questions about the applicability of these guidelines in your environment, contact your Dell EMC representative to discuss the appropriateness of the recommendations.

Audience

This document is intended for Dell EMC customers, partners, and employees who desire to learn more about best practices when configuring Microsoft Exchange Server with Dell EMC Unity All-Flash storage systems. It assumes the reader has working knowledge of Dell EMC Unity and Microsoft Exchange.

We welcome your feedback along with any recommendations for improving this document. Send comments to StorageSolutionsFeedback@dell.com.

1 Storage configuration

Dell EMC Unity is a virtually provisioned, flash optimized storage system designed for ease of use. This paper covers the all-flash array models. This section provides foundational array technologies that support the application-specific sections that follow. Additional information for this section can be found in the *Dell EMC Unity Best Practices Guide*.

Exchange Server, like all other mission critical systems requires storage to be configured properly for the specific workload being deployed. Here are some overall best practices for storage configuration on Dell EMC Unity arrays.

1.1 Storage pools

As of Dell EMC Unity OE version 4.2, Dell EMC Unity supports two different types of storage pools on all-flash storage systems: Traditional pools and dynamic pools. Traditional pools apply RAID protection to discrete groups of drives within the storage pool. Dynamic pools apply RAID to groups of drive extents from drives within the pool and allow for greater flexibility in managing and expanding the pool. Dynamic pools must be configured from all-flash drives; dynamic pools cannot be built with HDDs.

In general, it is recommended to use a small number of storage pools within Dell EMC Unity to reduce complexity and increase flexibility. However, it may be appropriate to configure additional storage pools in order to:

- Separate workloads with different I/O profiles
- Dedicate resources to meet specific performance goals
- Separate resources for multi-tenancy
- Create smaller failure domains

Additional information can be found in the [Dell EMC Unity: Dynamic Pools](#) white paper.

1.1.1 Storage pool capacity

Storage pool capacity is used for multiple purposes:

- To store all data written into storage objects (LUNs, file systems, and datastores, and VVols) in that pool
- To store data that is needed for snapshots of storage objects in the pool
- To track changes to replicated storage objects in that pool

Storage pools must maintain free capacity in order to operate properly. By default, Dell EMC Unity will raise an alert if a storage pool has less than 30% free capacity, and will begin to automatically invalidate snapshots and replication sessions if the storage pool has less than 5% free capacity. Dell EMC recommends that a storage pool always have at least 10% free capacity.

1.1.2 All-Flash pool

All-flash pools provide the highest level of performance in Dell EMC Unity. Use an all-flash pool when the application requires the highest storage performance at the lowest response time.

- Dell EMC FAST™ Cache and FAST VP are not applicable to all-flash pools.
- Compression is only supported on an all-flash pool.
- Snapshots and replication operate most efficiently in all-flash pools.
- Dell EMC recommends using only a single drive size and a single RAID width within an all-flash pool.

Example: For an all-flash pool, use only 1.6TB SAS flash 3 drives and configure them all with RAID 5 8+1.

1.1.3 Hybrid pool

Hybrid pools are not applicable when using Dell EMC Unity All-Flash arrays.

2 Dell EMC Unity features

This section describes some of the native features available on the Dell EMC Unity platform. Additional information on each of these features can be found in the *Dell EMC Unity Best Practices Guide*. Features not applicable to the all-flash array models covered by this paper are noted.

2.1 FAST VP

FAST VP accelerates performance of a specific storage pool by automatically moving data within that pool to the appropriate drive technology based on data access patterns. FAST VP is only applicable to hybrid pools within a Dell EMC Unity Hybrid flash system.

2.2 FAST Cache

FAST Cache is a single global resource that can improve performance of one or more hybrid pools within a Dell EMC Unity Hybrid flash system. FAST Cache can only be created with SAS flash 2 drives, and is only applicable to hybrid pools.

2.3 Compression

Dell EMC Unity compression is available for block LUNs and VMFS datastores in an all-flash pool starting with Dell EMC Unity OE version 4.1. Compression is available for file systems and NFS datastores in an all-flash pool starting with Dell EMC Unity OE version 4.2.

Be aware that compression increases the overall CPU load on the system when storage objects service reads or writes of compressible data, and may increase latency when accessing the data. Before enabling compression on a storage object, it is recommended to monitor the system and ensure the system has available resources to support compression. (Refer to the “Hardware Capability Guidelines” section and Table 2 in the *Dell EMC Unity Best Practices Guide*.) Enable Compression on a few storage objects at a time and then monitor the system to be sure it is still within recommended operating ranges, before enabling compression on more storage objects. Additional information regarding compression can be found in the *Dell EMC Unity: Compression* white paper.

Compression will only provide space savings if the data on the storage object is at least 25% compressible. Before enabling compression on a storage object, determine if it contains data that will compress; do not enable compression on a storage object if there will be no space savings. Contact your Dell EMC representative for tools that can analyze the data compressibility. Additional information regarding compression can be found in the *Dell EMC Unity: Compression* white paper.

2.4 Snapshot best practices

Be aware that snapshots increase the overall CPU load on the system, and increase the overall drive IOPS in the storage pool. Snapshots also use pool capacity to store the older data being tracked by the snapshot, which increases the amount of capacity used in the pool until the snapshot is deleted. Consider the overhead of snapshots when planning both performance and capacity requirements for the storage pool.

Before enabling snapshots on a storage object, it is recommended to monitor the system and ensure that existing resources can meet the additional workload requirements. (Refer to the “Hardware Capability Guidelines” section and Table 2 in the *Dell EMC Unity Best Practices Guide*.) Enable snapshots on a few storage objects at a time, and then monitor the system to be sure it is still within the recommended operating ranges before enabling more snapshots. Additional information can be found in the *Dell EMC Unity: Snapshots and Thin Clones* white paper.

2.5 Data at Rest Encryption (D@RE)

D@RE is controller based encryption that does not impact system performance; therefore Dell EMC recommends ordering Dell EMC Unity systems as encryption-enabled, if appropriate for your environment.

Note: Encryption can only be enabled at the time of system installation with the appropriate license.

If encryption is enabled, Dell EMC recommends making external backups of the encryption keys after system installation as well as immediately following any change in the system’s drive configuration (such as creating or expanding a storage pool, adding new drives or replacing a faulted drive). Additional information regarding Data at Rest Encryption can be found in the *Dell EMC Unity: Data at Rest Encryption* white paper.

2.6 Host I/O limits

Dell EMC recommends setting host I/O limits on workloads that might monopolize pool resources and starve other applications of their required performance.

Example: Limit the bandwidth available to large-block applications that may be increasing the latency on other small-block workloads.

Additional information can be found in the *Dell EMC Unity: Unisphere Overview* white paper.

3 Microsoft Exchange considerations

While Exchange Server is highly dependent on consistent I/O throughput from the back-end storage, it no longer requires as much I/O chatter as it did in past versions. Larger I/O sizes and intelligent cache and memory management on the Exchange mailbox server have altered some of the previous ideas and practices about storage requirements for this mission critical application.

3.1 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. Run performance tests to determine whether a server and disk subsystem will be efficient and ensure that Exchange Server will run in its optimal environment.

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 estimated performance numbers for the disk subsystem. LoadGen and Jetstress are available from Microsoft as free downloads and are discussed later in this document.

Another useful tool is the Windows® Performance Monitor that helps to define a baseline and show how the application may perform in the current environment. This tool is discussed further in the “Performance monitoring and troubleshooting” section.

3.2 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 2016 **Mailbox server role** includes all server components from Exchange 2013 Mailbox and Client Access roles:

- **Client Access services:** provide authentication, limited redirection, 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:** 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:** 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:** 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.

3.3 Background database maintenance (BDM)

The background database maintenance (BDM) process includes online defragmentation and online database scanning and produces a large 256 KB sequential read I/O.

The BDM process is enabled by default and runs continuously on both active and passive database copies. Scheduling the BDM process for active copies is possible but scheduling or disabling it on passive database copies is not available.

In Exchange 2016, the BDM process consumes considerably less bandwidth and produces fewer I/Os than the Exchange 2010 version. In the Exchange 2010 database availability group (DAG) configuration, the BDM process ran once a week and produced about 30 IOPS per database with a throughput of 7.5 MB/s. In Exchange 2016, this process needs to run only once in four weeks and produces about nine IOPS per database with a throughput of 2.25 MB/s per database copy.

3.4 Exchange storage 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.¹

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.

Table 1 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 is to distribute three copies plus a lagged copy across DAG members and utilize the same spindle on each of the four servers to host an active copy, two passive copies and a lagged copy of each server's databases.¹

Microsoft design guidance is specifically for JBOD (non-SAN) environments where larger, slower single spindles are used to provide storage for Exchange databases. Therefore, this Microsoft guidance does not apply for Dell EMC Unity SAN LUNs due to the following reasons:

1. Because RAID is inherent in Dell EMC Unity arrays, 3-copy DAGs are not an absolute requirement for resilience. 2-copy DAGs are more common.
2. When taking volume snapshots with AppSync™, an additional lagged database copy may not be needed.
3. The risk of database reseeding is very low when compared to a single disk JBOD. If a disk drive in a Dell EMC Unity array fails, a SAN spare is automatically used and the server and applications are unaware of any hardware change.
4. Microsoft's rationale for putting multiple databases on a disk is to allow for faster reseeding of the database from multiple sources. Since this is not a factor for Dell EMC Unity LUNs, Dell EMC 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.

¹ [“The Exchange 2016 Preferred Architecture”](#) by Ross Smith IV, dated October 12, 2015

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 volume mount points (VMP) for database and log volumes.

Table 2 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

3.4.1 Microsoft recommends ReFS for Exchange 2016 database volumes (JBOD)

Use NTFS on SAN volumes. This ReFS recommendation is for JBOD and not SAN volumes, and should not be used for Exchange 2016 database and log volumes on Dell EMC Unity 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](#).

3.4.2 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.

3.4.3 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.

3.4.4 Faster volume formatting

Dell EMC Unity volumes will quick format most efficiently in Windows Server environments if prior to formatting the following command is run at a Windows command prompt or PowerShell® command prompt:

```
C:\> fsutil behavior set DisableDeleteNotify 1
```

After formatting the volumes of the LUNs presented to the server, set the behavior back to default:

```
C:\> fsutil behavior set DisableDeleteNotify 0
```

3.4.5 Multi-path I/O (MPIO) recommended

Install Dell EMC PowerPath™ for optimal path management and maximum I/O throughput. In virtual deployments, install PowerPath on the hypervisor host. Use a minimum of two host bus adapters (HBAs) per server with each HBA connected to at least two array front-end ports to avoid a single point of failure.

3.4.6 Virtual machines (VMs) running Exchange

Block storage is recommended, and in most cases the only supported storage access method for Exchange databases, transaction logs and transport queues. Storage used by an Exchange server guest machine can use virtual storage that is a fixed size (VHD or VHDX in a Hyper-V environment), or direct access to block storage using either the pass through/RDM hypervisor option or in-guest iSCSI. On the file side, there is only one supported use case/configuration for SMB access to storage:

“Fixed VHDs may be stored on SMB 3.0 files that are backed by block-level storage if the guest machine is running on Windows Server 2012 Hyper-V (or a later version of Hyper-V). The only supported usage of SMB 3.0 file shares is for storage of fixed VHDs. Such file shares can't be used for direct storage of Exchange data. When using SMB 3.0 file shares to store fixed VHDs, the storage backing the file share should be configured for high availability to ensure the best possible availability of the Exchange service.” – From [Exchange 2016 virtualization](#) TechNet page.

Note: NFS shares or SMB implementations other than SMB 3.0 are not supported.

For VMware vSphere virtual machines, use virtual LSI Logic SAS (vSCSI) or Paravirtual SCSI (PVSCSI) adapters. Use multiple vSCSI adapters. This helps to avoid inducing queue depth saturation within the guest OS.

Add all four SCSI controllers to the virtual machine and equally distribute disks across the four available vSCSI adapters.

Place the Exchange storage on disks that are separate from the guest OS physical storage.

3.4.7 Storage design validation

3.4.7.1 Jetstress

The final step in the Exchange pre-production deployment phase is to validate the storage to ensure that it is configured properly and can sustain the load it is designed to support.

The Microsoft Jetstress tool is used to validate the Microsoft Exchange storage design. The tool simulates Microsoft Exchange I/O at the database level by interacting directly with the Extensible Storage Engine (ESE) database technology (also known as Jet) on which Microsoft Exchange is built. Jetstress can be configured to test the maximum I/O throughput available to the disk subsystem within the required performance constraints of Microsoft Exchange. Jetstress can accept a simulated profile of specific user count and IOPS per user to validate that the disk subsystem is capable of maintaining an acceptable performance level by generating the metrics defined in that profile. Dell EMC recommends Jetstress testing to validate storage reliability and performance prior to the deployment of the Microsoft Exchange production environment.

Download Jetstress 2013 (for Exchange 2013/2016 testing):

<https://www.microsoft.com/en-us/download/details.aspx?id=36849>

3.4.7.2 Exchange Solution Reviewed Program

The Exchange Solution Review Program (ESRP) is a Microsoft program for validating storage vendors' Exchange Server designs. Vendors run multiple Jetstress tests based on various performance, stress, backup-to-disk, and log file replication requirements. Microsoft reviews the results and posts approved solutions at <https://technet.microsoft.com/en-us/office/dn756396.aspx>.

4 Performance monitoring and troubleshooting

4.1 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.

Table 3 Exchange performance monitoring counters

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		

4.1.1 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: The background database defragmentation tasks that are currently executing

Database ==> Defragmentation Tasks Pending: 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

4.2 Troubleshooting

4.2.1 Unisphere Performance Dashboard

Use Dell EMC Unisphere™ Performance Dashboard to monitor and troubleshoot Exchange workloads along with other application workloads on Dell EMC Unity systems in your organization.

Historical Charts review performance statistics from prior hours/days/weeks to determine trends.

Real-time Charts allow adding specific charts to show current activity on specified storage objects; it is fully configurable.

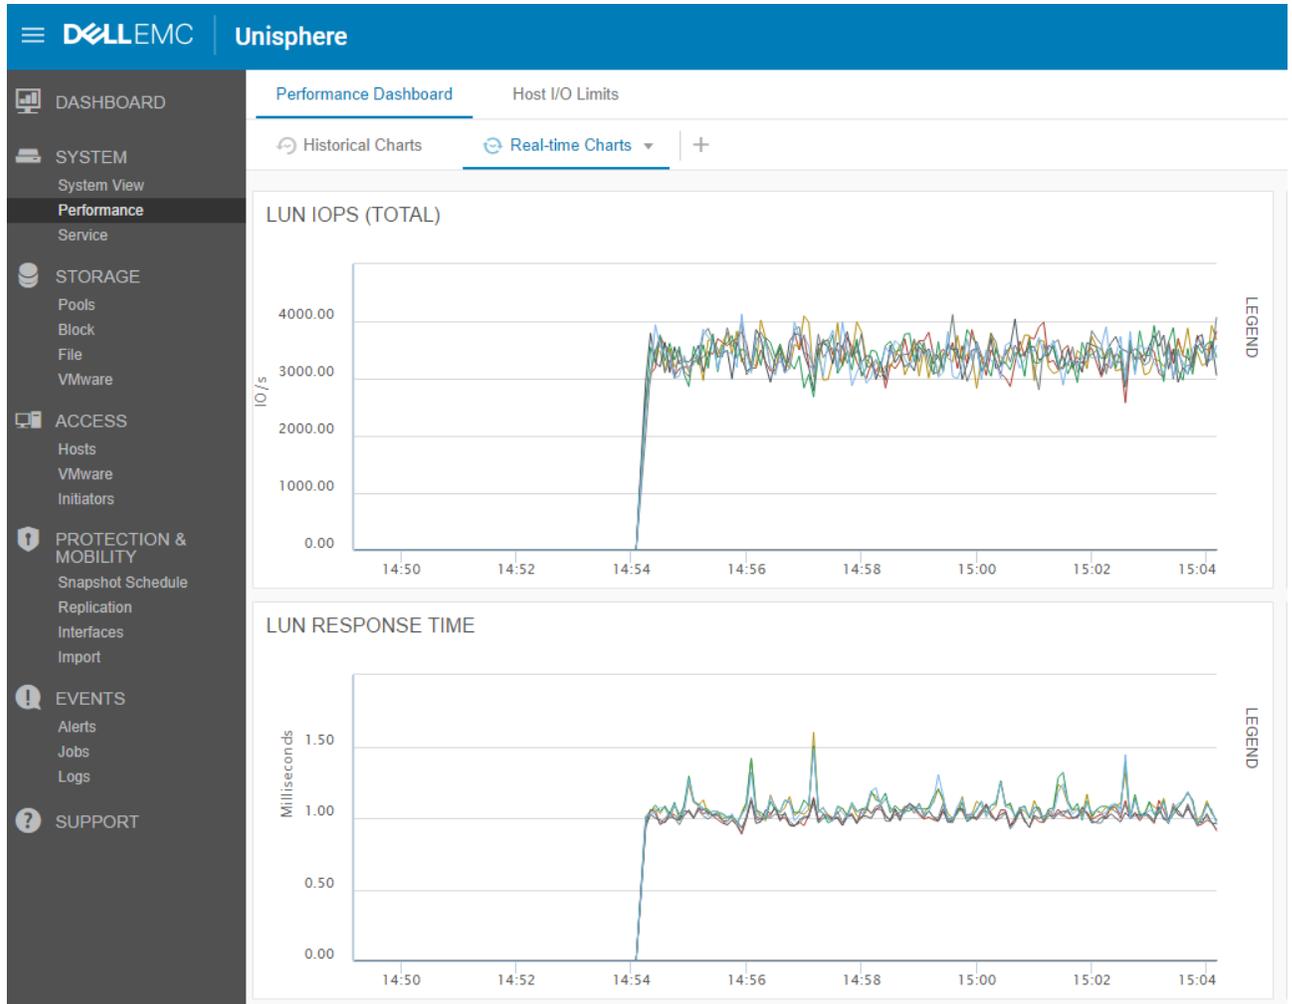


Figure 1 Dell EMC Unisphere chart while running application workloads

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

5 Data protection - AppSync

AppSync is software that enables Integrated Copy Data Management (iCDM) with Dell EMC's primary storage systems, including Dell EMC Unity.

AppSync simplifies and automates the process of creating and using snapshots of production data. By abstracting the underlying storage and replication technologies, and through application integration, AppSync empowers application owners to manage data copy needs themselves. Therefore, the storage administrator only needs to be concerned with initial setup and policy management, resulting in a more agile environment. Additional information on AppSync can be found in the [AppSync User and Administrator Guide](#) and the [AppSync Performance and Scalability Guidelines](#).

5.1 Snapshots using AppSync for Exchange

Dell EMC Unity snapshots initiated through AppSync are recommended for use with Microsoft Servers. Snapshots 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 AppSync Servers can be taken in 30-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 30-minute window.

By design, VSS data has to be captured within 10 seconds or the snapshot will fail. Since Exchange 2013, databases and logs can reside on the same volume, meaning fewer volumes 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.

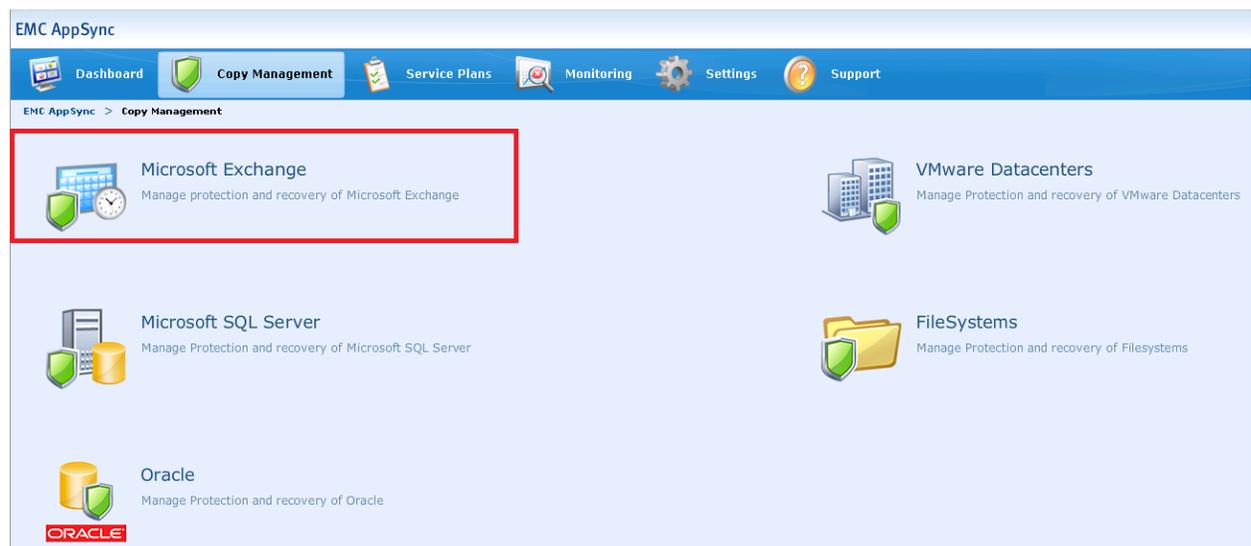


Figure 2 AppSync Microsoft Exchange Copy Management

5.1.1 Exchange Database verification

To run the database verification process as required by Microsoft, AppSync for Microsoft Exchange Server provides a verification option that can be used to mount the database volume snapshot 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.

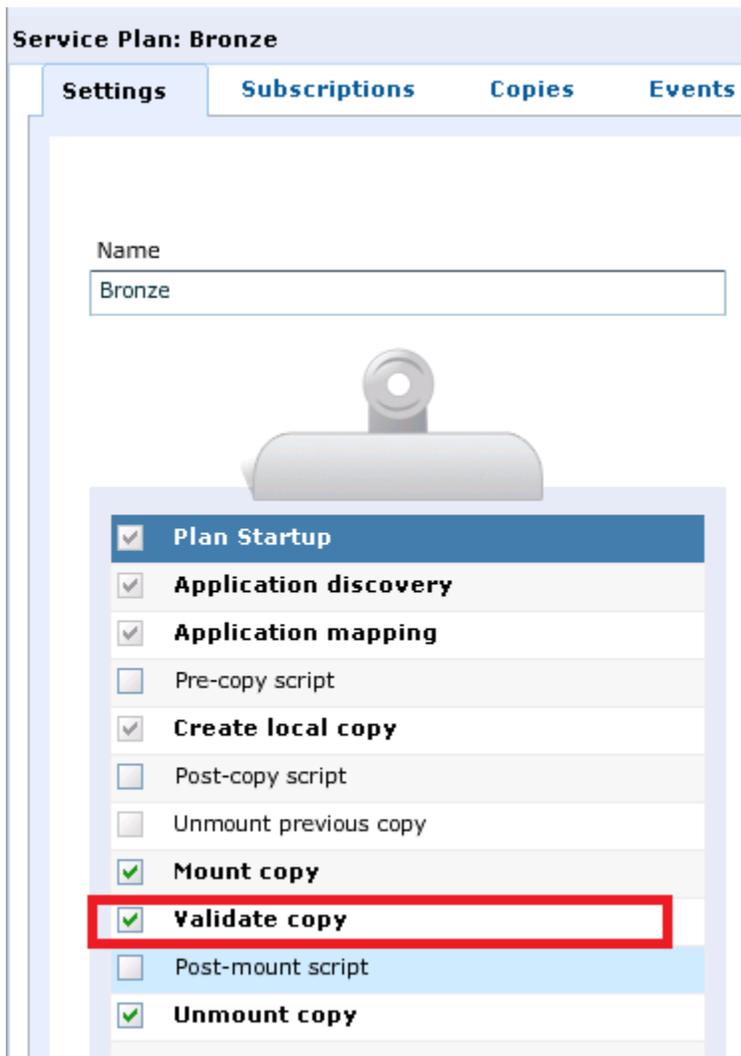


Figure 3 Selecting option for validation in AppSync Service Plan

For more information about AppSync for Microsoft Exchange, see the [AppSync support site](#).

6 Recovery

Recovery of Exchange data is typically a multi-level process to restore lost data at the database level, mailbox or even item level. The processes are well documented in the specific recovery application being used.

6.1 Using AppSync snapshots

A database level snapshot can be recovered by using the **Restore** option to either recover the database “in place” or re-direct it to an Exchange Recovery Database. Using the **Mount** option will allow the mounting of the specific snapshot to an alternate location. This can be useful to recover more granular data using 3rd-party tools such as Quest Recovery Manager for Exchange or KrollOntrack™ PowerControls™.

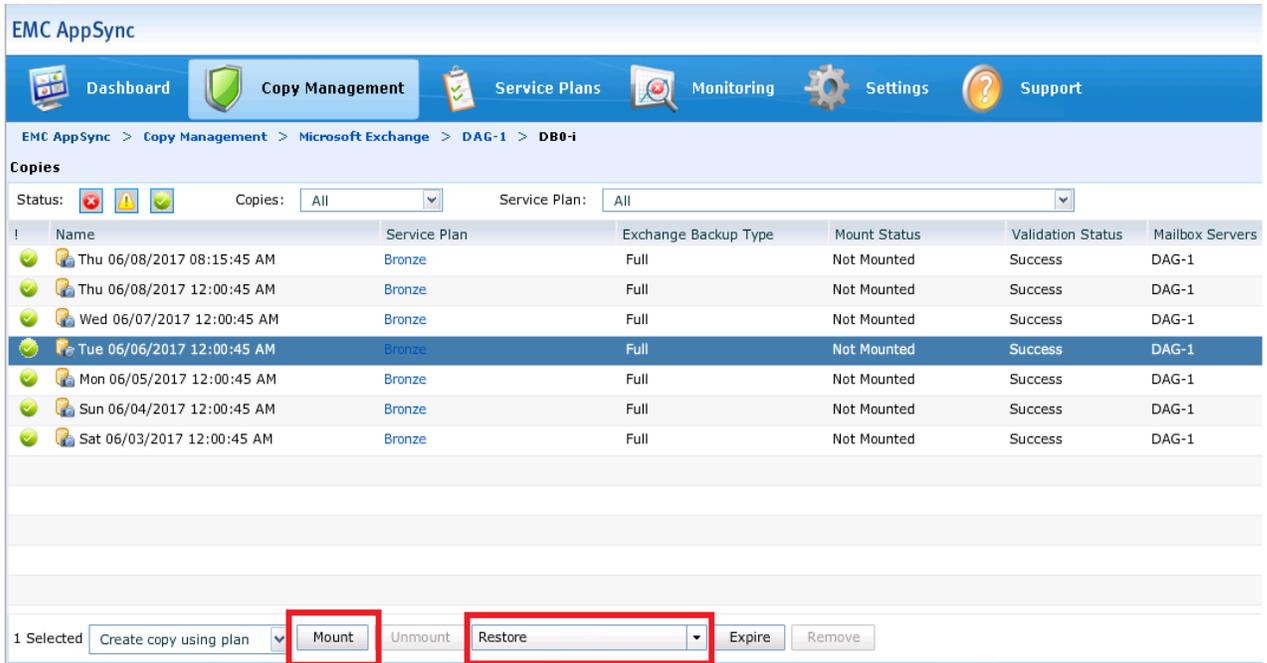


Figure 4 AppSync Copy Management screen

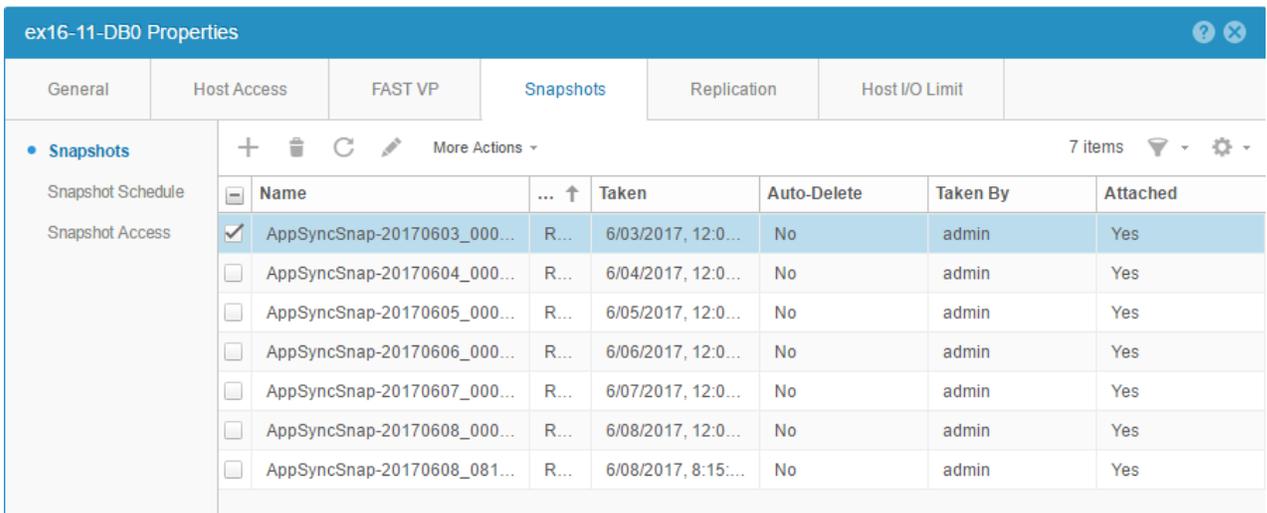


Figure 5 AppSync snapshot appearance in Unisphere

6.2 Recovery databases

A recovery database is a feature introduced with Exchange 2013 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, and 2010 mailboxes 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.

6.2.1 Single item recovery

Microsoft provides the following information on single item recovery for Exchange 2013 and 2016.²

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](#).

² "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 consistent snapshots taken with AppSync can be replicated to a DR array using Dell EMC Unity replication. This would allow data (as current as the last snapshot available) to be brought online at the main site or disaster recovery site.

If Exchange Server boots from SAN, snapshots 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 and DNS services being available. 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 DAG. However, the round-trip network latency must not be greater than 250 milliseconds (ms).

A Technical support and resources

The [Dell EMC Data Center Transformation](#) site is focused on enabling customers to gain understanding regarding the transformation of their IT infrastructure utilizing new technologies.

[Dell EMC Unity Info Hub](#) is a source providing helpful links to document and tools to customers.

Hardware and software support for the Dell EMC Unity platform is found at: [Dell EMC Support](#).

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

A.1 Related resources

Dell EMC Publications:

- Dell EMC Unity Best Practices Guide

<https://www.emc.com/collateral/white-papers/h15093-dell-emc-unity-best-practices-guide.pdf>

- Dell EMC titles on [Dell EMC Online Support](#):
 - Dell EMC Unity: Compression – Overview
 - Dell EMC Unity: Data at Rest Encryption
 - Dell EMC Unity: Dynamic Pools - Overview
 - Dell EMC Unity: Performance Metrics – A Detailed Review
 - Dell EMC Unity: Replication Technologies – A Detailed Review
 - Dell EMC Unity: Snapshots and Thin Clones – A Detailed Review
 - Dell EMC Unity: Unisphere Overview – Simplified Storage Management
 - Dell EMC Unity: Virtualization Integration - Whitepaper

[AppSync support site](#).