



Dell EMC DSS 7000/DSS 7500 100,000 Mailbox Resiliency Microsoft Exchange 2016 Storage Solution

Tested with ESRP – Storage Version 4.0
Tested Date: December 2016

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.

Copyright © 2016 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners. Published in the USA [02/17] [White Paper]

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

Table of Contents

1	Overview.....	4
1.1	Disclaimer	4
2	Features	5
3	Solution Components.....	6
3.1	DSS 7000/DSS 7500	6
3.2	MegaRAID 9361-8i RAID Controller.....	8
4	Solution Description	11
4.1	Failure and Recovery Scenarios	14
4.2	Storage Sizing	15
4.3	Recommended Hardware Configuration.....	15
5	Targeted Customer Profile	17
5.1	Tested User Profile.....	17
5.2	Tested Deployment	17
5.3	Best Practices.....	19
5.4	Backup Strategy	20
6	Test Result Summary.....	21
6.1	Reliability Tests	21
6.2	Storage Performance Test Result Report	21
6.2.1	Individual Server Metrics	21
6.2.2	Aggregate Performance Across Servers/DAGs Metrics	24
6.3	Database Backup/Recovery Performance	24
6.3.1	Database Backup Test Result Report	24
6.3.2	Soft Recovery Test Result Report	24
7	Detailed Test Results	25
7.1	How to view Jetstress Reports	25
8	Conclusion.....	26
9	Additional Information.....	27

1 Overview

This document provides information about Dell EMC's storage solution for Microsoft Exchange Server. This solution is based on the *Microsoft Exchange Solution Reviewed Program (ESRP) – Storage program v4.0*. For any questions or comments regarding the contents of this document, see Additional Information.

The ESRP – Storage program was developed by Microsoft Corporation to provide a common storage testing framework for vendors to provide storage solutions for Microsoft Exchange Server. For more information about the Microsoft ESRP — Storage program, see <http://technet.microsoft.com/en-us/exchange/ff182054.aspx>.

This technical white paper discusses Dell EMC's solution for 100,000 Exchange mailboxes with 70 GB size per mailbox, supporting up to 150 messages per day in a four-copy DAG. The solution uses the Dell EMC DSS 7000/7500 chassis/sleds for the Exchange mailbox server role and uses the internal storage of DSS7000 for storing the Exchange mailbox databases and transaction logs.

1.1 Disclaimer

This document has been produced independently of Microsoft Corporation. Microsoft Corporation expressly disclaims responsibility for, and makes no warranty, express or implied, with respect to the accuracy of the contents of this document.

The information in this document represents the current view of Dell EMC on the issues discussed as of the date of publication. Due to changing market conditions, it should not be interpreted to be a commitment on the part of Dell EMC and Dell EMC cannot guarantee the accuracy of any information presented after the date of publication.

2 Features

This technical white paper describes a tested and validated storage solution for a 100,000 mailbox Exchange 2016 site-resilient environment by using Database Availability Group (DAG). A DAG is a high availability (HA) mechanism in Microsoft Exchange 2016 that supports multiple copies (up to 16) of Exchange database. There can be only one active copy of a given Exchange 2016 database at any given time and mail clients access this copy. Any changes to the active copy are replicated to passive copies on other member servers within the same DAG. For Exchange 2016 databases and logs, all hosts within a DAG are configured to be identical in terms of storage resources. The active and passive copies do not share any storage resources, as they reside on their own dedicated storage resources.

Microsoft preferred architecture for Exchange 2016 recommends having an Active/Active multi-site deployment with four database copies equally distributed between both sites – two copies in each site. High availability is built into the tested deployment so that it can withstand server failures as well as site failure. This design is in line with Microsoft Preferred Architecture (PA) wherever possible. The building block architecture contains 16 servers stacked across two Data Centers spread across two sites—one each in Site A and Site B. The tested environment simulates up to 7,600 users with 70 GB mailbox size (20 GB Primary mailbox and 50 GB Personal archive) and 150 messages a day or 0.121 IO operations per second (IOPS) per user, including a 20% IO headroom.

In this solution, the DSS 7000/DSS 7500 sleds with 3.5-inch drives are configured for the Mailbox Server role. Built to meet the large storage requirements, the Dell EMC Datacenter Scalable Solutions DSS 7000 packs up to ninety hot-serviceable 3.5" drives and two high-performance server nodes into an ultra-dense 4U chassis. Capable of delivering up to 720 terabytes of storage per system, the DSS 7000 efficiently tackles the most demanding storage environments. The number of active databases hosted on one DSS7500 sled during normal run time is 21. In the event of a server failure in one of the sites, the surviving DSS 7500 servers in the site hosts 24 active databases.

In the event of a site failure, each DSS 7500 sled in the available site hosts 42 active databases. In case of a server failure in the only available site, each server hosts 48 active databases. Following are the major features of the server/storage system:

- DSS 7000 offers massive object or block level data storage with up to ninety hot-serviceable 3.5" drives and with two high-performance server nodes (DSS 7500), It is a highly versatile, ultra-dense storage server built to provide extremely low dollar-per-GB costs for large scale data centers.
- The DSS 7000 storage server benefits from all the features of the Intel E5-2600v4 processor family as well as an extremely dense and energy efficient shared infrastructure chassis.
- Drive Bay (Storage)
 - Up to 90x 3.5" hot-swappable SATA/NL-SAS HDD drives
- Host-based RAID options with MegaRAID 9361-8i

3 Solution Components

The solution employs a combination of building blocks containing DSS 7000 chassis and DSS 7500 sleds that are capable of meeting the high performance requirements of messaging deployments. The solution is for up to 100,000 mailboxes with 70 GB mailbox size—inclusive of 20 GB primary mailbox and 50 GB personal archive. The following subsections describe the hardware components that are part of this Exchange solution:

Figure 1 DSS 7000 Chassis



Figure 2 DSS 7000/DSS 7500 Rear View



3.1 DSS 7000/DSS 7500

DSS 7000/DSS 7500 is a 4U 90x drives ultra-dense storage server. These systems consist of DSS 7000, which is the Chassis and 2x DSS 7500 sleds with a highly expandable memory and impressive I/O capabilities. This standard width, 4U system offers the lowest cost per GB in the industry today. The DSS 7000 is now available to customers worldwide.

Note: The DSS 7000 chassis does not fit in a standard 1200 mm server rack. Please contact your ESI sales associate to discuss the supported rack options and for more details on this platform.

DSS 7000/DSS 7500 can readily handle data-intensive applications that require a large storage capacity and I/O performance, such as email. It delivers the performance and availability

required for mission-critical email and is a great hardware building block for super-large organizations with huge mailboxes.

The Dell EMC DSS 7000/DSS 7500 has the following key areas of differentiation:

- Ultra density: Up to 90x 3.5" HDDs and two independent 2S sleds in a 4U form factor
 - A selection of high performance E5-2600 v4 series processors
 - Flexible backplane enables customizable HDD combinations
- Systems management with support for IPMI 2.0, BMC, remote KVM, and LDAP
- Redundant hot-plug 1100 W AC power supply for added reliability
- Serviceable nodes: ability to service one node while the other is running. Both nodes are working independently and control a maximum of 45 HDDs each. Shutting down one node for service does not impact the other node.
- Each sled has two boot drives and 4 PCIe Gen 3 slots
- Up to 384 GB of memory with 12 DDR4 DIMM slots per server node

DSS 7000/DSS 7500 configured with the 3.5-inch large form factor drives is used as part of this solution.

Table 1 DSS 7000/7500 Technical Specifications

Feature	Specification
Form Factor	4U rack
Processors	Intel Xeon E5-2600 v4 product family processors
Processor Sockets	2 per node
Front Side Bus	2x QPI links; Up to 9.6 GT/s
Cache	2.5 MB per core Up to 14 cores
Chipset	Intel C610 series chipset
Memory	Up to 384 GB (12 DIMM Slots) DDR4 <ul style="list-style-type: none"> - 16 GB and 32 GB DIMMs - 2400 MHz - RDIMM
Min/Max RAM	64 GB/384 GB per node
RAID Controller	MegaRAID 9361-8i
I/O Slots	4 PCIe Gen 3 slots: 2x FHHL (full height) 2x HHHL (low profile)
Drive Bay (Storage)	Up to 90x 3.5" hot-swappable SATA/NL-SAS HDD drives
Maximum Internal Storage	720 TB (8TB drives) with 3.5" HDDs
HDD	SATA, NL-SAS
Embedded LOM/NIC	4x 1 GbE (2 x Dual-port BCM5720)

Power Supply	<ul style="list-style-type: none"> - 1100 W Redundant AC PSU, 80+ Platinum Efficiency - Support 100 V+
Systems Management	Systems management with support for IPMI 2.0, BMC, remote KVM, and LDAP
Rack Support	ReadyRails sliding rails with cable management arm for 4-post racks (optional adapter brackets required for threaded hole racks)
Operating Systems	<ul style="list-style-type: none"> - Microsoft® Windows® Server 2012 R2 - Red Hat® Enterprise Linux 6.5, 7.x - Ubuntu® 14.04 - CentOS 6.6, 7.0 - SLES 11 x64 Enterprise SP3

3.2 MegaRAID 9361-8i RAID Controller

MegaRAID 12 Gb/s SAS solutions are designed to deliver the performance and data protection capabilities required for the most demanding next-generation server and storage platforms. Offering up to double the data transfer rate of 6 Gb/s SAS solutions and 12 Gb/s SAS allows the SAS infrastructure to deliver bandwidth that can fully utilize that of PCI Express® 3.0 with a single controller card. MegaRAID 9361-8i, with eight internal ports, delivers two 1.2 GHz PowerPC processor cores and a 72-bit DDRIII interface that delivers 1 GB cache memory. Powered by the LSI SAS 3108 dual-core ROC, the 9361-8i controller includes the latest PCI Express 3.0 and 12 Gb/s SAS technology. It provides support for increased scaling of hard disk drives (HDDs) and maximizes the performance benefits of solid state drives (SSDs).

MegaRAID 9361-8i is used in DSS 7000/DSS 7500 that hosts the Exchange Server. It is the internal host-based RAID Controller card from Broadcom. These cards, built on the LSI SAS 3108 dual core RAID on Chip (ROC) offer unmatched I/O performance for databases, applications and streaming digital media environments.

For more information about recommended hardware specifications, see Section 4.3.

Table 2 shows the technical specifications of MegaRAID SAS 9361-8i.

Table 2 MegaRAID 9361-8i Technical Specifications

Feature	Specification
Solution provided	Eight-port internal 12 Gb dual core ROC-based SAS solution for data center, cloud, and performance-hungry applications using up to 128 SAS or SATA devices
Physical dimensions	MD2 low profile (6.6" X 2.536")
Connectors	Two mini-SAS SFF8643 internal connectors (horizontal mount)
Device support	Up to 128 SAS and/or SATA devices
Host bus type	x8 lane PCI Express 3.0 compliant
Data transfer rates	Up to 12 Gb/s per port
I/O Processor / SAS controller	LSI SAS 3108 dual core RAID on Chip (ROC)
Cache memory	2 GB 1866 MT/s DDR3 SDRAM
MegaRAID management suite	CTRL-R (BIOS configuration utility) MegaRAID storage manager StorCLI (command-line interface) HII (UEFI Human Interface Infrastructure)
Operating temperature	Maximum ambient: Controller Card: 55°C, with optional CacheVault accessory (LSICVM02): 55°C
Key RAID data protection features	<ul style="list-style-type: none"> • RAID levels 0, 1, 5, and 6 • RAID spans 10, 50, and 6 • Online Capacity Expansion (OCE) • Online RAID Level Migration (RLM) • Auto resume after loss of system power during array rebuild or reconstruction (RLM) • Single Controller Multipathing • Load Balancing • Configurable stripe size up to 1MB • Fast initialization for quick array setup • Consistency check for background data integrity • SSD Support with SSD Guard technology • Patrol read for media scanning and repairing • 64 logical drive support • DDF compliant Configuration on Disk (COD) • S.M.A.R.T support • Global and dedicated Hot Spare with Revertible Hot Spare support <ul style="list-style-type: none"> - Automatic rebuild - Enclosure affinity - Emergency SATA hot spare for SAS arrays • Enclosure management <ul style="list-style-type: none"> - SES (inband) - SGPIO (sideband) • Databolt bandwidth optimizer technology support for compatible expander-based enclosures • Shielded state drive diagnostic technology

Operating voltage	+3.3V, +12V
SSD optimization	<ul style="list-style-type: none"> • MegaRAID CacheCode Pro 2.0 Software leverages SSDs in front of HDD volumes to create high-capacity, high-performance controller cache pools (Optional Upgrade) • MegaRAID Fast Path Software provides high-performance I/O acceleration for SSD arrays connected to 6 Gb/s MegaRAID SATA+SAS controllers (Included)
OS support	<ul style="list-style-type: none"> • Extensive support includes Microsoft Windows Server 2012/8 & 7/2008/Vista/2003, Linux, Solaris (x86), FreeBSD, VMware and more.

For more information about recommended hardware specifications, see Section 4.3.

4 Solution Description

In this solution, Dell EMC DSS 7000/DSS 7500 with 3.5-inch LFF drives are used as Mailbox Servers. The DSS 7000 / DSS 7500 combination provides SAS-based internal storage with RAID. The solution uses 90x 3.5-inch LFF 7.2KRPM NL-SAS disks across one DSS 7000 chassis with two DSS 7500 sleds in the chassis. The solution also uses four back-accessible 1.8-inch SSDs—two each across DSS 7500 sleds. The disk layout is listed below:

- Two rear-accessible SSDs (in RAID 1 container) for the operating system, application files and Exchange Transport database per DSS 7500 sled
- 42 disk drives (in RAID 0 containers) for the Exchange database and its transaction logs
- One disk drive for Restore LUN
- One disk drive for Auto-reseed volume
- One disk drive as a spare drive

The solution has a 4-copy DAG Layout (Active/Active) with Exchange Servers distributed between two sites: Site A and Site B. Each server node has 42-RAID 0 LUNs hosting one active and three passive databases per LUN during normal operating conditions. Each of these databases hosts 23 users with 70 GB mailbox size. Thus, a single server can accommodate 475 mailboxes in normal runtime. Sixteen such servers in one DAG (8 each in Site A and Site B) provide Exchange Mailbox Services for 7,600 users. The solution can be scaled out for 100,000 users as outlined in this white paper by implementing additional DAGs. Thus, to host 100,000 mailboxes the solution requires 14 DAGs. Dell recommends not implementing both DSS 7500 sleds from a single DSS 7000 chassis in the same DAG to avoid single point of failure at the chassis level. In the event of a server failure in one of the sites, the corresponding passive database copies in the surviving hosts get activated. The design also considers the site failure and the first server failure in the surviving site, which is considered as the worst case scenario. In this case, seven surviving hosts in the available site can accommodate 7600 mailboxes (1,086 mailboxes per server across 48 databases), and this scenario has been tested.

Active Database Activation Configuration (DC1 Site Failure) / DAG	/ Datacenter 2 Active Server	Active Mailboxes / Server	/ Datacenter 2	Total Active Mailboxes in DC2
Number of Active Databases (Secondary Datacenter (DC2) Activation)	42	950	336	7600
Number of Active Databases (First Server Failure after SDC Activation)	48	1086	336	7600

The mailbox user profile that was tested had 150 messages per day or 0.121 IOPS per user, which included a 20% IO overhead.

Figure 3 and 4 represent the distribution of database copies across the DAG members in Site A and Site B. It shows a 4-copy Active/Active DAG site resilient solution with Exchange Servers hosted at both sites.

Figure 3 Database Availability Group architectural diagram – Site A



Figure 4 Database Availability Group architectural diagram – Site B



4.1 Failure and Recovery Scenarios

Figure 3 and 4 show a diagram of the solution at Site A and Site B respectively. There are eight servers each in Site A and Site B. A single server failure in Site A or Site B activates the passive copies of the impacted databases. If there is a complete site failure, then the passive copies in the surviving site get activated, and the users connect to their databases on the surviving site. When one site is completely unavailable, and one of the servers in the surviving site fails, the corresponding passive database copies in the surviving hosts get activated as shown in Figure 5. This scenario is simulated in the test and considered the worst-case failure. Thus, each host is designed in a way that any one server is capable of holding the additional load. Each server is capable of handling the load for 1,086 mailboxes. Therefore, with seven servers, all 7,600 mailboxes can be managed without compromising on performance. Having fourteen such DAGs with 16 servers per DAG can host 100,000 mailboxes as outlined in this white paper.

Figure 5 Worst-case failure scenario



Click the Attachments icon () in the left pane of Adobe Reader to view the database distribution across servers. Double click the item to open the image. Microsoft Exchange 2013 Server Role Requirements Calculator can be used to derive the database distribution including the active and passive copies across servers located in both Site A and Site B. The database distribution follows a particular pattern to ensure that if a server fails, the passive copies are activated on the remaining hosts and the load on each host machine is evenly distributed.

4.2 Storage Sizing

Selecting the right storage is crucial to achieve a balance between cost and performance. The storage size and design should be based on the type of RAID, type of disk drives and the number of disk drives—from both capacity and IOPS perspectives. The storage design also depends on the actual size of mailbox on the disk drive, the content indexing space, and the required log space.

Use Microsoft Exchange 2013 Server Role Requirements Calculator to derive the required IOPS for a particular user profile. Figure 6 shows the Mailbox Calculator output for 7,600 users with 150 messages per day profile. The recommended IOPS per server is 229. Microsoft Exchange Jetstress tool verifies if the storage subsystem meets the targeted IOPS requirement. For more information, see Section 5.

Figure 6 Recommended IOPS from the Microsoft Exchange 2013 Server Role Requirements Calculator

Host IO and Throughput Requirements	/ Database	/ Server	/ DAG	/ Environment
Total Database Required IOPS	3	229	3666	3666
Total Log Required IOPS	1	48	775	775
Database Read I/O Percentage	60%	—	—	—
Background Database Maintenance Throughput Requirements	1.0 MB/s	84 MB/s	1344 MB/s	1344 MB/s

4.3 Recommended Hardware Configuration

Based on the solution requirements described in the earlier sections. Table 3 and Table 4 provide more information about the server and storage configuration. The firmware and driver versions are also provided for the tested solution.

Table 3 Exchange Server Configuration

Microsoft Exchange Server System	DSS 7000 /DSS 7500 with 3.5-inch drives
CPU	2 x Intel Xeon E5-2650 v4 processor with 12-cores
Memory	Up to 96 GB DDR4
NIC	BCM5720 Gigabit Ethernet PCIe
RAID Controller	MegaRAID SAS 9361-8i C600/X79 series chipset SATA RAID Controller Firmware version: 4.650.00-6121 Storport Driver Version 6.3.9600.18514 Driver version 6.11
Internal Disks	2 x 1.2 TB SATA Solid State Drives (SSD) for Operating System and Application

Table 4 Storage Subsystem configuration

Storage Subsystem	DSS 7000 / DSS 7500 Internal 3.5-inch drives
Disks	<p>45 x 8TB 7.2K RPM NL-SAS 3.5-inch disk:</p> <ul style="list-style-type: none"> • 42 x 8TB 7.2K RPM NL-SAS 3.5-inch drive in 14 x RAID 0 volumes (for DB and Log) • 1 x 8TB 7.2K RPM NL-SAS 3.5-inch drive (for Restore LUN) • 1 x 8TB 7.2K RPM NL-SAS 3.5-inch drive (for Auto-Reseed volume) • 1 x 8TB 7.2K RPM NL-SAS 3.5-inch drive (spare drive)
RAID Controller	MegaRAID SAS 9361-8i C600/X79 series chipset SATA RAID Controller

5 Targeted Customer Profile

This solution is intended for very large organizations hosting up to 100,000 Exchange 2016 mailboxes. The configuration used for testing was as follows:

- Number of mailboxes: 7,600
- Number of sites: 2 (Site A & Site B)
- Number of servers in each site: 8 in Site A and 8 in Site B
- User IO profile: 150 messages sent and received or 0.121 IOPS per mailbox (This includes 20% IO overhead factor)
- 70 GB Mailbox quota per mailbox (20 GB Primary Mailbox, 50 GB Personal Archive)
- 24x7 Background Database Maintenance enabled
- Database Availability Group (DAG) for Mailbox Resiliency (4 copies simulated—1 Active, 3 Passive)

5.1 Tested User Profile

The tested user profile had 0.121 IOPS per user with a 70 GB mailbox size. This equates to 150 messages (sent or received) per mailbox per day and accounts for an additional 20% I/O overhead. Additional applications such as certain mobile messaging applications can increase the IOPS profile of a user by three to four times. In such cases, it is recommended to recalculate the IOPS requirements while designing the Exchange solution.

5.2 Tested Deployment

The tested deployment simulated a site level failure scenario where one member of the available site (Site A or Site B) was also completely unavailable and the passive copies on the surviving DAG members were activated to provide mailbox service continuity. Therefore, the IOPS was simulated for 1,086 mailboxes on the same Exchange 2016 Server. The target IOPS for the given profile was 229. The achieved IOPS was 2429—much higher than the target—and the solution still maintained read and write latencies well within the recommended thresholds. The following tables summarize the testing environment:

Table 5 Simulated Exchange Configuration

Feature	Specification
Number of Sites	2 (Site A and Site B)
Number of servers per DAG	16 (8 in Site A and 8 in Site B), 7 hosting active mailboxes in test
Number of active mailboxes per server	475 (during normal operations) 950 (during site failure) 1,086 (in tested failover during worst-case failure)
Number of databases per host	48
Number of copies per database	4 (2 in Site A and 2 in Site B)

Number of mailboxes per database	23
Simulated profile: IOPS per mailbox	0.121 (150 messages/day) This includes 20% IO overhead factor
Database/Log LUN size	7.27 TB
Number of LUNs per server	42 (24 tested)
Number of DBs per LUN	2
Background database maintenance (BDM)	Tested with BDM enabled
Total database size for performance testing	2.91 TB per DB 139.68 TB total
Percentage of storage capacity used by the Exchange database	139.68 TB / 174.48 TB 80.05%

Table 6 Storage and Server Hardware

Feature	Specification
Storage connectivity (Fiber Channel, SAS, SATA, iSCSI)	SAS
Storage model and OS/firmware revision	Dell EMC DSS 7000/DSS 7500 with MegaRAID SAS 9361-8i attached to internal drives in DSS 7000 chassis Firmware 4.650.00-6121
Cache Memory	1 GB- MegaRAID SAS 9361-8i
Number of storage controllers	1 attached to internal drives in Dell EMC DSS 7000 chassis 1 attached to internal drives in Dell EMC DSS 7500 sled
Connectors	2 (Two mini-SAS SFF8643 internal connectors (horizontal mount))
Data Transfer Rates	Up to 12 Gb/s per port
Switch type/model/firmware revision	NA
HBA model and firmware	MegaRAID SAS 9361-8i Firmware 4.650.00-6121
Number of HBAs per host	2
Host server type	Dell EMC DSS 7000/DSS 7500 2 x Intel Xeon processor 96 GB RAM
Total number of disks tested in the solution	168 (24 per server)

Maximum number of spindles that can be hosted in the storage	90 x 3.5" drives per DSS 7000 2 x 2.5" per DSS 7500
--	--

Table 7 Storage and Server Software

Feature	Specification
HBA driver	MegaRAID SAS 9361-8i driver version 6.11
HBA QueueTarget Setting	N/A
HBA QueueDepth Setting	N/A
Multipathing	N/A
Host OS	Windows Server 2012 R2 Data Center X64 Edition
ESE.dll file version	15.1.669.32
Replication solution name/version	N/A

Table 8 Storage Disk Configuration (Mailbox Store Disks)

Feature	Specification
Disk type, speed and firmware revision	DELL EMC 7.2K 3.5" RPM 8TB NL-SAS Model – ST8000NM0055
Raw capacity per disk (TB)	7.27 TB
Number of physical disks in the test	168 (24 per Server)
Total raw storage capacity (TB)	1221.36 TB (174.48 TB per Server)
RAID level	RAID 0
Number of disks per LUN	1
Total formatted capacity	7.27 TB per LUN 174.48 TB per server
Storage capacity utilization	$174.48/1221.36=14.293\%$ Formatted capacity/Total raw capacity
Database capacity utilization	$(2.91 \text{ TB} \times 48)/174.48 \text{ TB}=80.05\%$ Database size/Total formatted capacity

5.3 Best Practices

Exchange Server 2010, 2013 and 2016 overcome the memory limitations of earlier Exchange versions by providing support as a 64-bit application. On Windows Server 2012 R2 x64 Edition, about 4 TB of addressable memory is available for kernel mode and user mode applications. Both the application and kernel have sufficient memory for operations, allowing the Extensible Storage Engine (ESE) in Exchange Server 2016 to utilize more memory to buffer data pages. The result is a reduction in the number of I/O operations, specifically the read operations required for the disk subsystem. The total number of database disk I/O operations for a given

user load depends on the available system memory. For a given load, the total database disk IOPS decreases over a period with increase in system memory. This decrease in database IOPS is primarily caused by a decrease in database reads.

Make sure that there are no I/O bottlenecks from an IOPS and disk latency perspective while sizing the Exchange Storage subsystem. The disk subsystem should be capable of supporting both the capacity and I/O throughput demands of the application. The following best practices are recommended to improve the I/O subsystem performance:

- For the Exchange 2016 database, the size of elements within a RAID stripe be set to 512K for best performance.
- Each server have a single RAID1 disk pair for the Operating System, Exchange binaries, protocol/client logs and transport database. Rest of the storage can be configured as independent RAID0 volumes.
- Each disk that houses an Exchange database be formatted with ReFS (with integrity feature disabled) and the DAG be configured such that AutoReseed formats the disk with ReFS.
- Average database read latencies (Avg. Disk sec/Read) should not exceed 20 milliseconds. Exchange Server 2016 storage latencies are most often related to the number of disk drives available for a given workload. Windows Performance Monitor may be used to monitor Exchange Server 2016 database counters.
- Sharing the spindles that host the Exchange Database and log with any other application or operating system is not recommended. This is because Sharing Exchange 2016 storage resources with other applications may negatively affect the performance of Exchange 2016 server.

5.4 Backup Strategy

To protect email data from potential data loss, having a well designed and implemented backup solution is critical. Depending on environmental requirements, different backup strategies may be implemented, such as backup to tape or LAN/SAN-based backup. In this solution, DAG is used to maintain a passive database copy on a separate storage system. This passive copy of the database may be backed up to a tape or disk drive.

The log replay test was used to measure the maximum rate at which the log files can be replayed on the passive copies. This is used to determine the restore times and also the database write throughput that can be achieved during a log recovery.

6 Test Result Summary

This section provides a high-level summary of the test data from Microsoft Exchange Jetstress as part of the ESRP requirements and the link to the detailed HTML reports, which are generated by the ESRP testing framework.

6.1 Reliability Tests

Reliability tests are run for 24 hours and the goal is to verify if the storage can handle a high I/O load for a long period. After the stress test, both log and database files are analyzed for integrity to make sure that there is no database/log corruption.

The following list provides an overview of any errors reported during testing.

- Any errors reported in the saved event log file? No
- Any errors reported during the database and log checksum process? No

6.2 Storage Performance Test Result Report

The storage performance test is designed to exercise the storage with maximum sustainable Exchange I/O for four hours. The test shows how long it takes the storage to respond to an I/O user load. The data here is the sum of all of the logical disks I/Os and the average of all the logical disks I/O latency in the four hour test duration. The achieved IOPS was around 2429.

As part of the ESRP framework, the Stress Test was also performed. The duration of the test was 24 hours with a target IOPS of 0.121 per user or 229 IOPS per server. The achieved IOPS was around 2438 per server. This was well above the target IOPS. The Stress Test Result Report is provided for reference.

6.2.1 Individual Server Metrics

Table 9 shows the sum of I/O across Mailbox databases and the average latency across all databases on a per server basis.

Table 9 Individual Server Metrics

Server 1: DSS-1A

Database I/O	
Target Disk Transfers/sec	132
Database Disks Transfers/sec	2377
Database Disks Reads/sec	1696
Database Disks Writes/sec	681
Average Database Disk Read Latency (ms)	13.98
Average Database Disk Write Latency (ms)	0.266
Transaction Log I/O	
Log Disks Writes/sec	173.65
Average Log Disk Write Latency (ms)	0.128

Server 2: DSS-1B

Database I/O	
Target Disk Transfers/sec	132
Database Disks Transfers/sec	2361
Database Disks Reads/sec	1685
Database Disks Writes/sec	676
Average Database Disk Read Latency (ms)	14.01
Average Database Disk Write Latency (ms)	0.263
Transaction Log I/O	
Log Disks Writes/sec	171.93
Average Log Disk Write Latency (ms)	0.129

Server 3: DSS-2A

Database I/O	
Target Disk Transfers/sec	132
Database Disks Transfers/sec	2482
Database Disks Reads/sec	1770
Database Disks Writes/sec	712
Average Database Disk Read Latency (ms)	13.25
Average Database Disk Write Latency (ms)	0.261
Transaction Log I/O	
Log Disks Writes/sec	183.99
Average Log Disk Write Latency (ms)	0.119

Server 4: DSS-2B

Database I/O	
Target Disk Transfers/sec	132
Database Disks Transfers/sec	2449
Database Disks Reads/sec	1747
Database Disks Writes/sec	702
Average Database Disk Read Latency (ms)	13.47
Average Database Disk Write Latency (ms)	0.267
Transaction Log I/O	
Log Disks Writes/sec	179.61
Average Log Disk Write Latency (ms)	0.128

Server 5: DSS-3A

Database I/O	
Target Disk Transfers/sec	132
Database Disks Transfers/sec	2405
Database Disks Reads/sec	1704
Database Disks Writes/sec	701
Average Database Disk Read Latency (ms)	13.61
Average Database Disk Write Latency (ms)	0.266
Transaction Log I/O	
Log Disks Writes/sec	165.65
Average Log Disk Write Latency (ms)	0.120

Server 6: DSS-3B

Database I/O	
Target Disk Transfers/sec	132
Database Disks Transfers/sec	2406
Database Disks Reads/sec	1705
Database Disks Writes/sec	701
Average Database Disk Read Latency (ms)	13.78
Average Database Disk Write Latency (ms)	0.266
Transaction Log I/O	
Log Disks Writes/sec	165.54
Average Log Disk Write Latency (ms)	0.120

Server 7: DSS-4A

Database I/O	
Target Disk Transfers/sec	132
Database Disks Transfers/sec	2524
Database Disks Reads/sec	1790
Database Disks Writes/sec	734
Average Database Disk Read Latency (ms)	13.07
Average Database Disk Write Latency (ms)	0.258
Transaction Log I/O	
Log Disks Writes/sec	173.83
Average Log Disk Write Latency (ms)	0.11

6.2.2 Aggregate Performance Across Servers/DAGs Metrics

Table 10 shows the aggregated results of I/O and the average latency across servers in the solution.

Table 10 Aggregate Performance Metrics across all Servers

Database I/O	
Database Disks Transfers/sec	17004
Database Disks Reads/sec	12097
Database Disks Writes/sec	4907
Average Database Disk Read Latency (ms)	13.59
Average Database Disk Write Latency (ms)	0.264
Transaction Log I/O	
Log Disks Writes/sec	1214.2
Average Log Disk Write Latency (ms)	0.122

6.3 Database Backup/Recovery Performance

The database backup test measures the sequential read rate of the database files, and the soft recovery test measures the recovery/replay performance (playing transaction logs in to the database).

6.3.1 Database Backup Test Result Report

The test is to measure the maximum rate at which databases could be backed up through VSS. The following table shows the average rate for a single database file:

Table 11 Database Backup Test Metrics

MB read/sec per database	25.37
MB read/sec total per server	1217.78

6.3.2 Soft Recovery Test Result Report

The test is to measure the maximum rate at which the log files can be played on the passive copies. The following table shows the average rate for 505 log files played in a single storage group. Each log file is 1 MB in size.

Table 12 Soft Recovery Test metrics

Average number of log files played	508
Average time to play one Log file (sec)	5.67

7 Detailed Test Results

Detailed Jetstress test results for all seven mailbox servers are attached to this PDF file in the form of HTML reports. Report types are as follows:

- 4-hour performance test
- Checksum for 4-hour performance test
- 24-hour stress test
- Checksum for 24-hour stress test
- Database backup test
- Soft recovery test

7.1 How to view Jetstress Reports

Click the Attachments icon in the left pane of Adobe Reader to view the list of Jetstress reports (HTML files). Double click each item to open the corresponding report in your browser.

8 Conclusion

This ESRP document presents a tested and validated Exchange solution for 100,000 mailboxes with 70 GB mailbox size supporting up to 150 messages per day in a four-copy DAG. The solution uses the Dell EMC DSS 7000 / DSS 7500 for the Exchange mailbox server role and uses its internal storage for storing the Exchange mailbox databases and transaction logs.

Testing was carried out as part of the ESRP test framework by using Microsoft Exchange Server 2013 Jetstress. The test results showed that the proposed solution is more than capable of delivering the IOPS and meeting the capacity requirements to support 100,000 mailboxes with the described mailbox profile.

This document is developed by storage solution provider Dell EMC and reviewed by the Microsoft Exchange Product team. The test results and data presented in this document are based on the tests introduced in the ESRP test framework. Customers should not quote the data directly for pre-deployment verification. It is necessary to go through the exercises to validate the storage design for a specific customer environment.

The ESRP program is not designed to be a benchmarking program, and the tests are not designed to obtain the maximum throughput for a given solution. Rather, the tests focus on obtaining recommendations from vendors for Exchange application. The data presented in this document should not be used for direct comparisons among solutions.

9 Additional Information

1. **Dell.com/support** is focused on meeting customer requirements with proven services.
2. **DellTechCenter.com** is an IT Community where you can connect with Dell EMC customers and Dell EMC employees to share knowledge, best practices and information about Dell EMC products and installations.
3. Referenced or recommended Dell EMC publications:
 - a. Dell EMC Solutions for Unified Communications and Collaboration
 - b. DSS 7000/DSS 7500 Owner's Manual