

# DELL EMC READY BUNDLE FOR MICROSOFT EXCHANGE

## EXCHANGE SERVER 2016

### *Design Guide*

#### **ABSTRACT**

This Design Guide describes the design principles and solution components for Dell EMC Ready Bundle for Microsoft Exchange suited for very large scale deployment. This solution uses DSS 7500 sleds on a DSS 7000 chassis as a building block.

March, 2017

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 © 2017 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. 03/17. Design Guide.

## TABLE OF CONTENTS

<b>EXECUTIVE SUMMARY .....</b>	<b>5</b>
Introduction .....	5
Scope .....	5
Audience .....	6
<b>SYSTEM COMPONENTS .....</b>	<b>6</b>
Dell EMC DSS servers .....	6
Dell EMC networking .....	8
Microsoft Exchange Server 2016 .....	8
<b>APPLICATION ARCHITECTURE .....</b>	<b>8</b>
Storage architecture .....	10
Network architecture .....	12
<b>DESIGN PRINCIPLES .....</b>	<b>13</b>
High availability .....	13
Application-level high availability .....	13
Infrastructure-level high availability .....	13
Application performance .....	13
Best practices .....	14
<b>COMPREHENSIVE SOLUTION DESIGN .....</b>	<b>15</b>
Core components .....	16
Configuring even number of DAGs .....	17
Configuring DAGs with partial number of sleds .....	18
Configuring odd number of DAGs .....	19
Optional Components .....	20
Scalability .....	21
<b>SOLUTION SIZING .....</b>	<b>21</b>
Email Solution for a large organization of 100,000 mailboxes .....	22

**VERIFICATION ..... 25**

    Microsoft Jetstress verification ..... 25

    Microsoft Exchange Load Generator verification..... 26

**CONCLUSION ..... 27**

**ADDITIONAL RESOURCES ..... 27**

## EXECUTIVE SUMMARY

Large organizations face challenges in identifying the right set of hardware to design and size large scale email system. The Dell EMC Global Solutions Engineering team has developed a pre-architected Dell EMC Ready Bundle for Microsoft Exchange. The solution uses Dell EMC DSS servers to help customers simplify implementation and choose the appropriate hardware. We base our design on best practices recommended by Microsoft and Dell EMC's own experience supporting hundreds of customers worldwide. Our solution balances performance and cost, provides faster time to value, and increases return on investment.

High availability and site resilience are essential to ensure business continuity. Business growth also requires an easily scalable solution. We take these requirements into account in our design for the Ready Bundle for Microsoft Exchange.

Microsoft Exchange Server 2016 brings a new set of technologies, features, and services to Exchange Server—the messaging platform that provides email, scheduling, and tools for custom collaboration and messaging service applications.

Sizing the Exchange Server and selecting the right hardware platform is the first step of the implementation cycle.

The design principles and solution components for the Ready Bundle for Microsoft Exchange described in this guide are suited for large scale deployments. Please contact your sales representatives for buying the solution.

- The System Components section provides an introduction to the Microsoft Exchange Server 2016 system components and the solution architecture.
- The Application Architecture section explains the application architecture and describes the storage and network architecture in detail.
- The Design Principles section explains the design principles including high availability, application performance and best practices.
- The Comprehensive Solution Design section explains the comprehensive solution design and describes how the building blocks can be scaled out to support growth.
- The Solution Sizing section describes the sizing considerations for server and storage.
- The Verification section provides an overview of the verification process that was performed to ensure that the solution met the design requirements.

## INTRODUCTION

Microsoft Exchange Server 2016 is a leading enterprise messaging system that delivers email, calendar, voice mail and contacts to users on a variety of devices through the Outlook client. Exchange Server 2016 provides reliable, scalable, enterprise-class email with compliance and e-discovery features integrated with Microsoft SharePoint and Skype for Business. Exchange Server 2016 supports people and organizations as their work habits evolve from a communication focus to a collaboration focus. The database design in Exchange Server 2016 reduces the storage I/O requirements, thereby optimizing Exchange for cost-effective, low-speed storage. Exchange deployment must be appropriately sized to meet not only specific message profile requirements, but also growth and high availability requirements. The following sections provide a brief overview of Exchange 2016, describe the important considerations when sizing an Exchange deployment and provide the building block architecture that Dell EMC has developed for Exchange infrastructure.

## SCOPE

This guide describes the pre-architected Ready Bundle for Microsoft Exchange on Dell EMC Datacenter Scalable Solutions (DSS) servers. The solution uses DSS 7500 sleds on a DSS 7000 chassis as a building block, as it can support large volumes of data. The design is explained with a sample implementation of 100,000 mailboxes. Customers needing larger implementations can also use the

building blocks to scale out as per their business needs. However, for complex implementations, Dell EMC recommends that the customers engage with Dell EMC Consulting Services directly about their design needs in addition to using this guide.

Microsoft Exchange Jetstress and LoadGen tools were used to verify the building block design that is used as a foundation for our Exchange 2016 solution. Common user profile options are also provided.

## AUDIENCE

This guide is intended for IT managers, messaging administrators, and consultants interested in designing and deploying a large-scale Exchange 2016 solution on DSS 7000/DSS 7500 for various user profiles. Users are expected to have sufficient knowledge of Exchange 2016.

## SYSTEM COMPONENTS

Users today demand larger mailboxes and faster email. Microsoft Exchange supports low-cost and large-capacity storage and has high-availability built in through the database availability group. The DSS 7000/DSS 7500 system is designed to address the needs of hyperscale customers for ultra-dense, yet relatively inexpensive object or block data storage (per TB). Capable of delivering up to 720 TB of storage per system, DSS 7000/DSS 7500 efficiently tackles the most demanding storage environments. Exchange solutions built on DSS 7000/DSS 7500 can be scaled out based on the business needs.

Components for Ready Bundle for Microsoft Exchange also include Dell EMC Networking products, Dell EMC ProDeploy services and Dell EMC ProSupport services. For network integration, Active Directory and data migration, we also have the Dell EMC Consulting Services for Exchange. Please contact your sales and services representatives for more details. This guide focuses on the system design.

The following subsections describe the hardware and software components of Ready Bundle for Microsoft Exchange.

## DELL EMC DSS SERVERS


The DSS 7000/DSS 7500 systems benefit from all the features of the Intel E5-2600 v4 processor family as well as an extremely dense and energy-efficient, shared infrastructure chassis. The combination of Intel's enhancements with the E5-2600v4 processors and the DSS 7000 ultra-dense chassis provide performance, efficiency, and flexibility in one dense package.

The DSS 7000/DSS 7500 system has the following specifications:

- Dual node 2-socket server sleds
- Intel Xeon E5-2600 v4 product family processors
- Intel C610 series chipset
- Up to 384 GB of memory with 12 DDR4 DIMM slots per server node
- Systems management with support for IPMI 2.0, BMC, remote KVM, and LDAP
- Redundant hot-plug 1100 W AC power supply for added reliability
- Two boot drives in each sled and four PCIe Gen 3 slots
- Up to 90 x 3.5" hot-swappable SATA/NL-SAS hard drives

- MegaRAID 9361-8i and the Microsemi 8805 RAID controllers
- 4 x PCIe 3.0 expansion slots

The Dell EMC DSS 7000 chassis, when configured with two DSS 7500 sleds of 45 drives each, provides up to 720 TB raw storage capacity.



**Figure 1. Rear view and front view of Dell EMC DSS 7000 chassis with DSS 7500 sleds**

The DSS 7000 is a unique platform in the Dell portfolio targeted at a specific set of use cases. Special attention must be given to the following items:

- Overall datacenter planning around rack infrastructure. The DSS 7000 chassis does not fit in a standard 1200 mm server rack. Please contact your Extreme Scale Infrastructure (ESI) sales representatives to discuss the supported rack options.
- Add on and upgrade PCIe cards
- Management
- Firmware updates
- Installation

Please contact your ESI sales representatives to discuss this platform in more depth.

## DELL EMC NETWORKING

Dell EMC Networking offers S-Series and N-Series high density 100M/1G/10G/40 GbE top-of-rack (ToR) switches specially built for applications in high-performance data centers and computing environments. S-Series S4048T-ON is a 10G BASE-T (RJ45) switch that leverages a non-blocking switching architecture; S4048T-ON delivers line-rate L2 and L3 forwarding capacity within a conservative power budget. The compact S4048T-ON design provides industry leading density of 48 dual-speed 1/10G BASE-T ports as well as six 40 GbE QSFP+ up-links to conserve valuable rack space and simplify the migration to 40 Gbps in the data center core. Each 40 GbE QSFP+ up-link can also support four 10 GbE (SFP+) ports with a breakout cable. In addition, S4048T-ON incorporates multiple architectural features that optimize data center network flexibility, efficiency and availability, including I/O panel to PSU airflow or PSU to I/O panel airflow for hot or cold aisle environments and redundant, hot-swappable power supplies and fans.

## MICROSOFT EXCHANGE SERVER 2016

Exchange Server 2016 is the latest release from Microsoft, where the number of server roles is reduced to two: Mailbox and Edge Transport. The server components from Exchange 2013 Mailbox and Client Access server roles such as Client Access protocols, Transport service, Mailbox role and Unified Messaging are combined in the Exchange 2016 mailbox server.

There are two editions of Exchange Server 2016—Exchange Server 2016 Standard Edition and Exchange Server 2016 Enterprise Edition. The number of mounted databases on each edition is as given below:

- Exchange Server 2016 Standard Edition – Up to 5 mounted databases per server
- Exchange Server 2016 Enterprise Edition – Up to 100 mounted databases per server

Key changes made in Exchange 2016 are:

- Improved performance and reliability
- Faster and more intuitive search
- 33% faster failovers than Exchange 2013
- Simplified document sharing with support for OneDrive and SharePoint 2016
- New Outlook web experience – Outlook On The Web

## APPLICATION ARCHITECTURE

The following sections describe the design principles of Ready Bundle for Microsoft Exchange, and explain how the infrastructure sub-systems are architected in this solution.

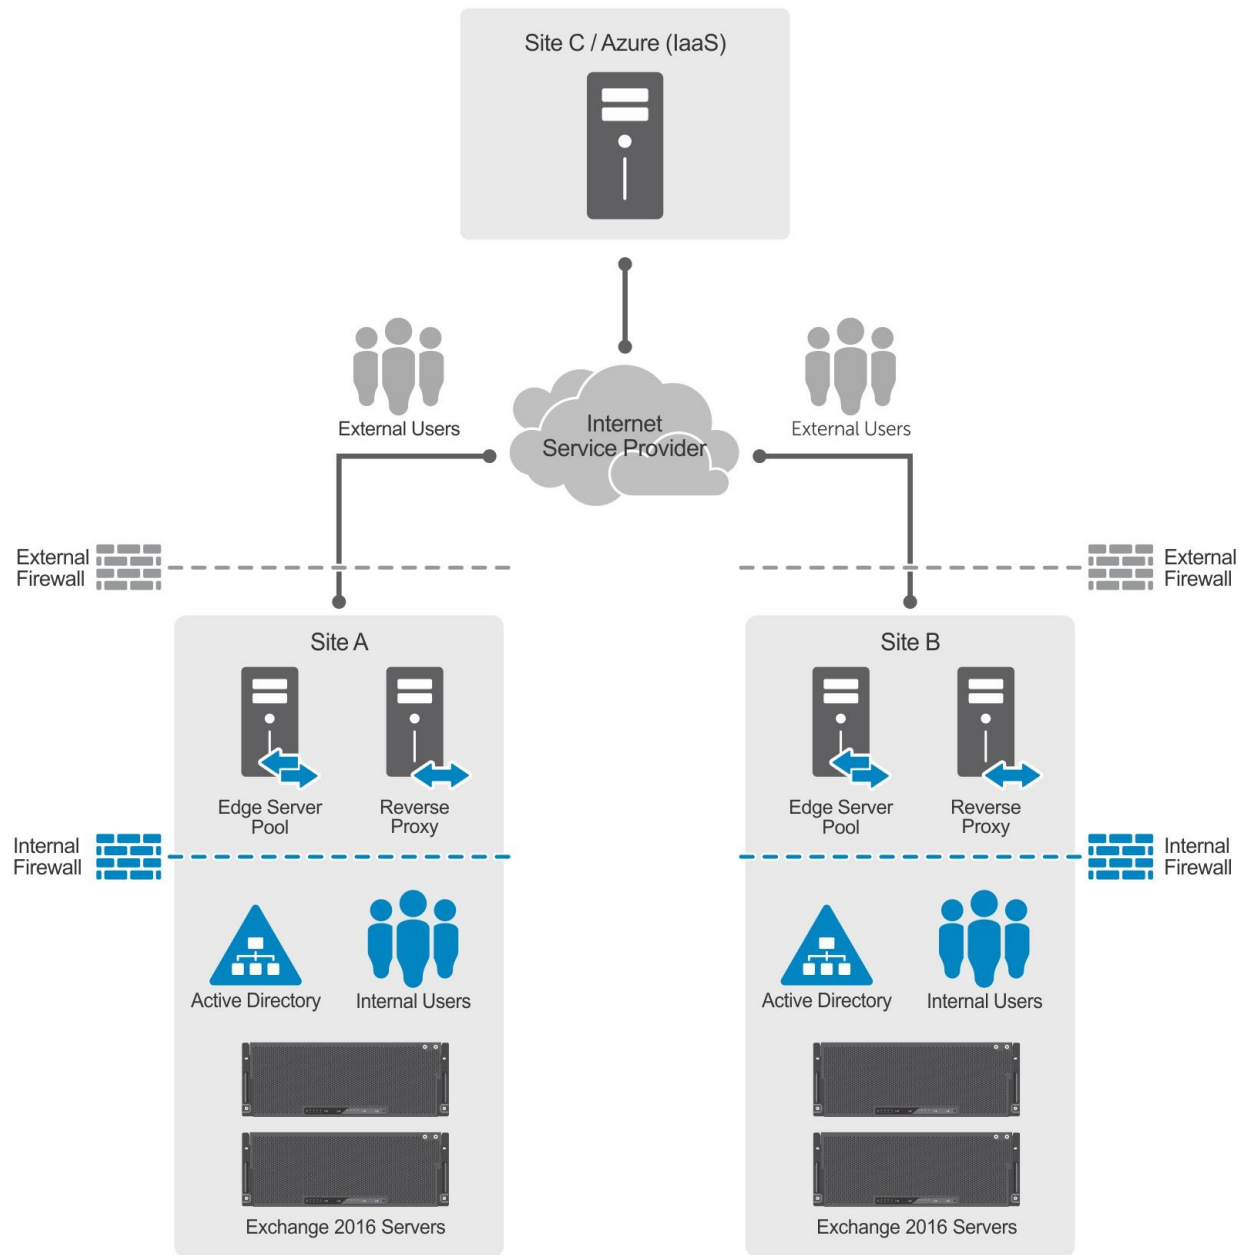
### Architecture for a large global organization of 100,000 employees

This organization has 100,000 employees based in many locations across different countries, including several main facilities and branch offices as well as many remote workers. The design enables each employee to have a higher capacity mailbox of 70 GB (20 GB primary mailbox and 50 GB personal archives) and 150 messages sent or received per day. This site-resilient solution spread across three regions with two data centers in each site (Site A and Site B) makes use of active-active user distribution model.

This case is detailed under the section Sample Implementation.



Figure 2 shows the external data center architecture.



**Figure 2. External data center architecture for Exchange deployment**

## STORAGE ARCHITECTURE

This solution employs 3.5-inch large form factor (LFF) drives as part of DSS 7000/DSS 7500 for storing the Exchange mailbox databases and transaction logs.

DSS 7000/DSS 7500 provides SAS-based internal storage with RAID. This solution uses 90 x 3.5-inch LFF 7.2 K RPM NL-SAS disks (90 drives distributed between the two sleds). The solution also uses four rear-accessible 1.8-inch SSDs—two per sled. The disk drive layout is explained 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

Figure 3 illustrates the disk drive layout.

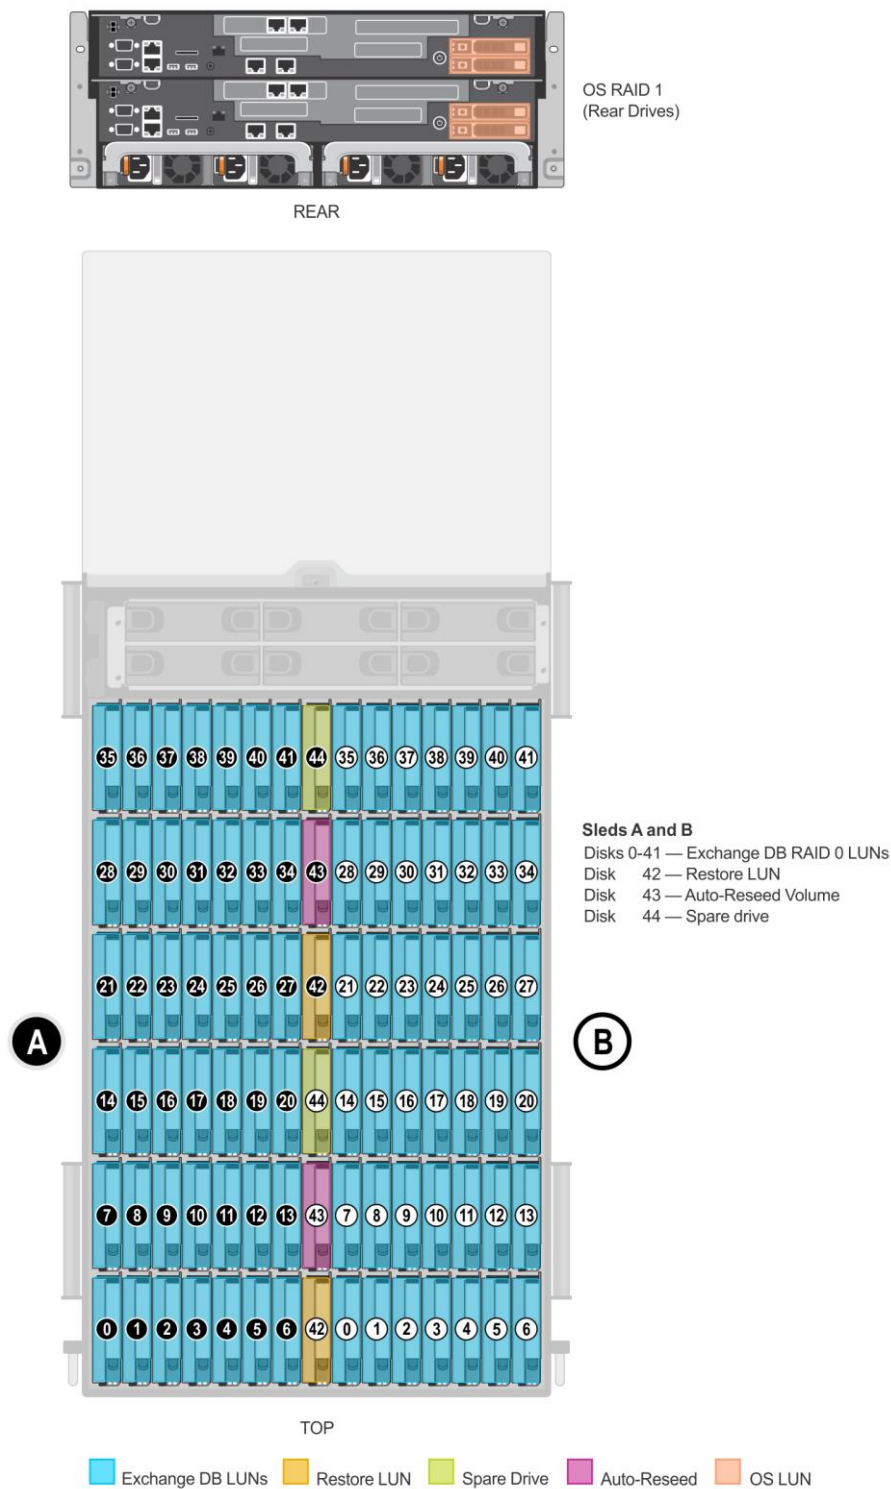


Figure 3. RAID LUN layout for the Exchange databases

## NETWORK ARCHITECTURE

Microsoft Exchange Server 2016 network requirements are simple and easy to deploy. In our solution we have leveraged NIC teaming to provide multi-tier HA across the network.

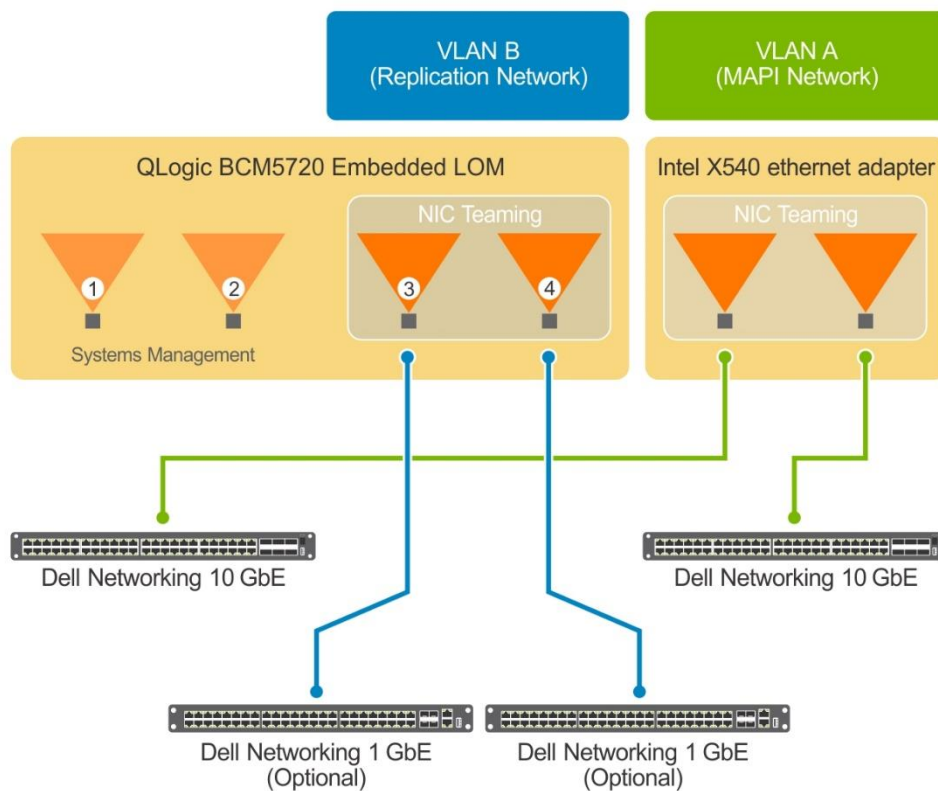
In an Exchange deployment, the servers configured for Exchange Server roles have two types of network traffic:

- Messaging Application Programming Interface (MAPI) traffic that includes client connectivity traffic
- Replication traffic between DAG members (optional)

The DSS 7500 sled supports QLogic BCM5720 embedded LOM with four 1 GbE network ports.

Two 1 GbE embedded LOM ports are used for database replication network and the other two 1 GbE embedded LOM ports can be used for system or chassis management by using BMC. The two 1 GbE ports (numbered 3 and 4) in the embedded LOM are teamed to provide redundancy.

The 10 GbE dual-port Ethernet adapter is used for MAPI traffic. The two ports are teamed to provide redundancy and HA for network connections. Figure 4 illustrates this architecture.



**Figure 4.** Network architecture for the Exchange deployment

The two 10 GbE ports can be teamed to provide an aggregate bandwidth of 20 Gbps per sled. By using the Smart Load Balancing (SLB) teaming mode, all network paths remain active during normal operation. The 1 GbE network ports in the embedded LOM can be used to provide dedicated replication network between the Exchange mailbox servers.

## DESIGN PRINCIPLES

Emails are mission critical for organizations and users. When designing an Exchange solution, we should take into account high availability (HA) and Service Level Agreement (SLA). Today's users also demand large mailboxes, making it necessary to build a solution that provides both large storage capacity and cost effectiveness.

### HIGH AVAILABILITY

Exchange 2016 has native high availability through Database Availability Groups (DAG). This provides HA and Disaster Recovery (DR) capabilities.

High availability is built into Ready Bundle for Microsoft Exchange, so that it can withstand server failures as well as site failures. This design is based on the building block architecture where 16 servers/sleds per DAG are stacked across two data centers spread across two sites—eight in Site A and eight in Site B. See the section Comprehensive Solution Design for details on the design for this solution.

#### Application-level high availability

The DAGs in Exchange Server 2016 are utilized to provide a highly available service for the Exchange mailbox databases by storing multiple copies of the active database. Deploying a DAG with multiple servers hosting the Exchange Mailbox server role provides highly available messaging services to the end users. Exchange DAG provides native data protection. This capability can be complemented by deploying a backup and recovery solution that is application-aware and can help in performing item-level recovery of the application data.

A hardware or software load balancer should be deployed to balance the load of the client requests to the Mailbox servers.

#### Infrastructure-level high availability

In a solution infrastructure, resources such as server, storage, network path and switches should be highly available. Application high availability requires that the infrastructure builds redundancy for each component. RAID 1 disk provides redundancy for the operating system and the Exchange binaries. Multiple network adapters and switches connecting to the Exchange infrastructure and the data center network build resiliency in network connectivity. The Storage and Network architectures explained in this guide take infrastructure-level high availability into account.

### APPLICATION PERFORMANCE

In addition to HA, application performance is also critical to guarantee reliable end-user experience. We have used Microsoft Exchange Server Jetstress 2013 Tool and Exchange Load Generator 2013 to ensure that the solution is appropriately sized to meet the performance needs as explained in the Verification section.

**Note:** Microsoft Exchange Server Jetstress 2013 Tool and Exchange Load Generator 2013 are the latest versions of the tools from Microsoft and they are used for testing with Microsoft Exchange Server 2013 as well as Microsoft Exchange Server 2016.

## BEST PRACTICES

Microsoft recommends disabling Logical Processor/hyper-threading when deploying Exchange Server on physical servers and configuring the System Profile Settings to Performance in BIOS.

The recommended server configurations are:

- Dual socket servers (20-24 cores)
- Up to 96 GB of memory
- A battery-backed write cache controller
- Large form factor drive bays (preferably internal storage)

When a DAG is deployed across two sites for site resilience with both sites running active copies, Microsoft recommends deploying a witness server in a third site for automatic failover. In the event of a data center failure, mailbox databases are activated in the surviving data center to provide email services to the users.

Ensure that there are no I/O bottlenecks from an IOPS and disk latency perspective while sizing the Exchange Storage subsystem. The disk subsystem must be capable of supporting both the capacity and I/O throughput demands of the application.

**Note:** Microsoft Preferred Architecture recommendations have been followed wherever applicable in this solution. As this solution does not use conventional rack servers, there are some deviations.

The following best practices are recommended to improve the I/O subsystem performance:

- For the Exchange 2016 database, set the size of elements within a RAID stripe to 512 K for best performance.
- Ensure that each server has a single RAID 1 disk pair for the Operating System, Exchange binaries, protocol/client logs and transport database. You can configure the rest of the storage as independent RAID 0 volumes.
- Format each disk that houses an Exchange database with Resilient File System (ReFS) with the integrity feature disabled and configure the DAG such that AutoReseed formats the disk with ReFS.
- Ensure that the average database read latencies (Avg. Disk sec/Read) do 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.

## COMPREHENSIVE SOLUTION DESIGN

Figure 5 shows the high-level architecture across different regions. It shows a customer scenario that consists of two data centers (Site A and Site B) in each region for Exchange Server placement and a third data center (Site C) or cloud providers such as Microsoft Azure IaaS for witness server placement with WAN links between them.

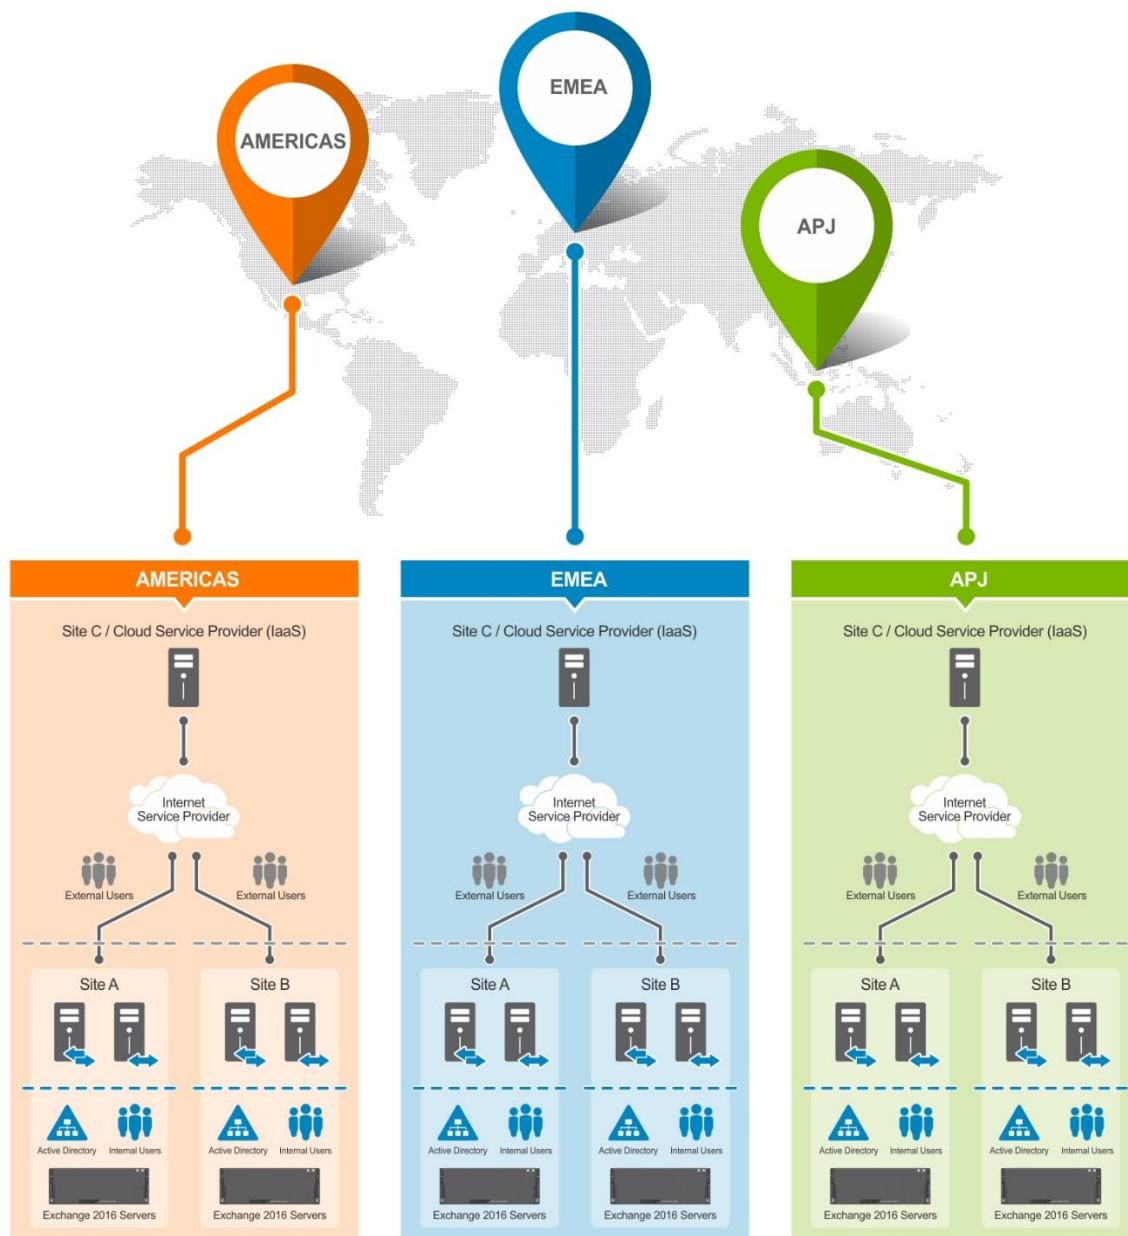


Figure 5. Data center architecture for different regions

As shown in Figure 5, the solution design consists of multiple infrastructure components such as server, storage, network, load balancer, firewall, reverse proxy, and so on. Every data center that has Exchange servers deployed has the same set of infrastructure components. The DSS 7000/DSS 7500 systems configured for Exchange deployment are connected to the internal data center network where all the other dependent infrastructure services are available. The internal data center network is a routable network and connects end users to the server infrastructure. A hardware or software load balancer is configured to redirect the incoming client requests to the Exchange infrastructure. External mailbox users connect to Exchange over the internet through an Edge Server pool configured at the site.

## CORE COMPONENTS

This solution ensures high availability and fault-tolerance by providing server-level, chassis-level, and site-level redundancies. Each physical server has identical configurations with respect to memory, processor and storage.

To ensure site resiliency and high-availability across hardware (server and chassis) and application, Dell EMC recommends that when you configure a DAG, you include servers from across chassis and sites. Members of a DAG must not be in the same physical chassis.

**Note:** Do not include both the sleds/servers of any chassis in a single DAG.



The following sub-sections explain how DAGs are configured for different number of sleds.



### Configuring DAGs with partial number of sleds

Consider a scenario where your setup requires 2.5 DAGs with a total of 40 sleds. You must ensure that each DAG has sleds included from different chassis and sites as shown in Figure 7.

In this example, DAG 1 and DAG 2 have 16 servers each configured across different chassis and sites. DAG 3 includes only eight servers and this DAG is configured to include four sleds from each site from different chassis.



Figure 7. Configuring DAGs with partial number of sleds

### Configuring odd number of DAGs

Consider a scenario where your setup requires three DAGs with a total of 48 sleds. You must ensure that all three DAGs have servers included from across chassis and sites as shown in Figure 8.

In this example, DAG 1, DAG 2, and DAG 3 have 16 servers each and these are configured from different chassis and sites.


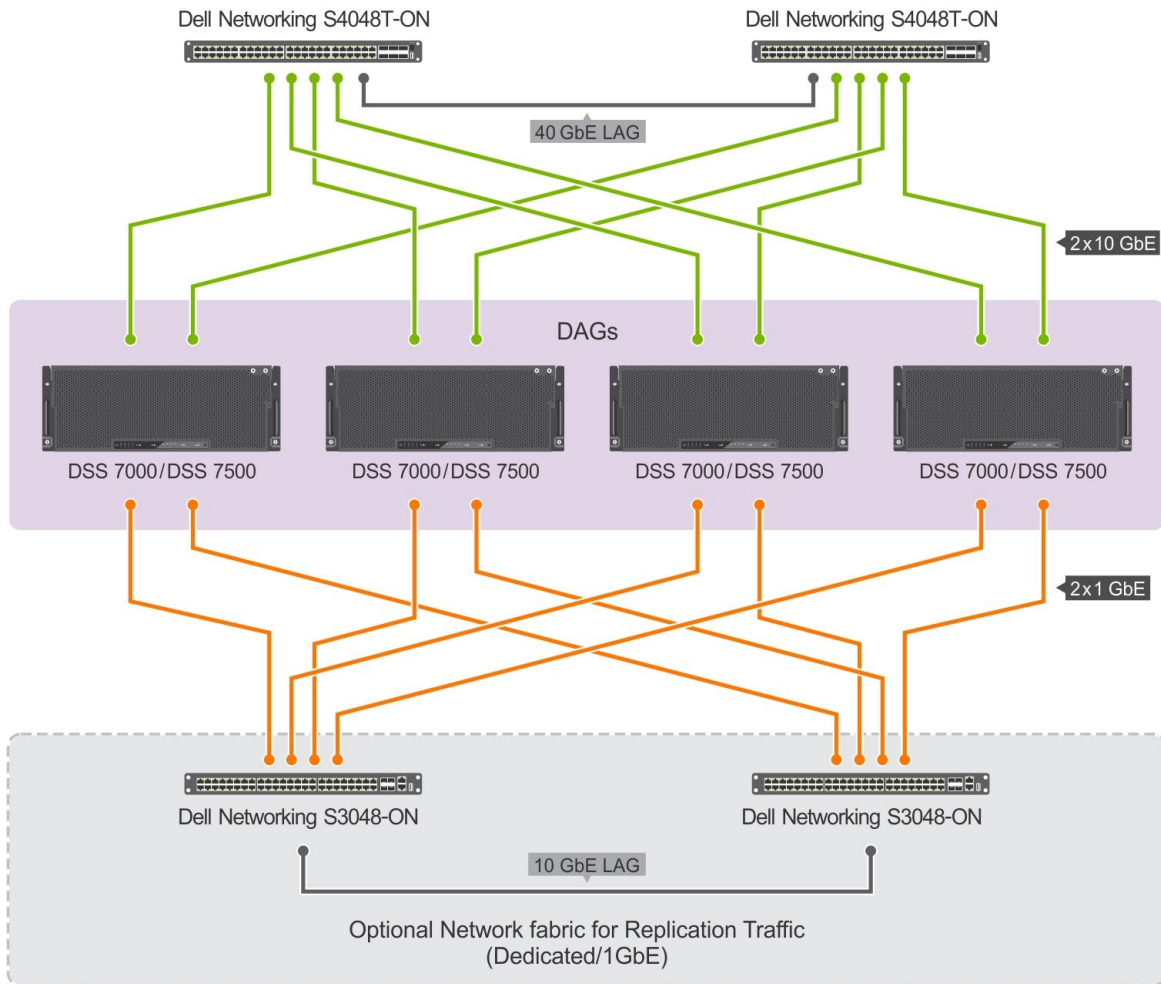


Figure 8. Configuring odd number of DAGs

For a site-resilient configuration, Dell EMC recommends a 4-copy DAG deployment with active-active distribution model. In case of a site failure, the database copies on the surviving site provide email services to the end users.

Exchange deployment depends on infrastructure services such as Active Directory (AD), Domain Name System (DNS) and load balancers. The connectivity from the Exchange infrastructure to these infrastructure services and end users must be resilient and highly available. For the Ready Bundle, Dell EMC Networking 10 GbE Top-of-Rack (ToR) switches were employed. The network architecture is illustrated in Figure 9.



**Figure 9. Detailed network architecture for the Exchange reference implementation**

As mentioned, a total of two 10 GbE network ports are available per sled. These network ports are teamed to provide an aggregate network bandwidth of 20 Gbps per server. As shown in Figure 9, these network ports are connected to two separate ToR switches. These switches are connected together with a 40 GbE Link Aggregation (LAG) and connected to the internal data center network by using 10 GbE ports. In the above example, the MAPI and replication traffic are separated by using VLANs. Using the two 10 GbE ports of the Intel X520 Ethernet adapter, it is possible to create a separate network fabric for replication traffic. A dedicated network for replication provides more efficiency as the replication network is not impacted by any issues in the MAPI network.

## OPTIONAL COMPONENTS

After the consolidation of server roles, Exchange 2016 has only two server roles – Mailbox and Edge Transport. Similar to Exchange 2013, it is not mandatory to deploy Edge Transport server. However, some organizations may require an Edge Transport server based on business needs. In these cases, the Edge server should be placed at the perimeter network (DMZ). The purpose of deploying an Edge server is to provide an additional layer of security for inbound and outbound emails and enable features such as address re-writing. Customers can choose the Edge server or third-party security appliances and services based on their needs.

## SCALABILITY

As organizations grow, you can accommodate additional mailboxes by adding more building blocks. As the maximum number of servers in a DAG is limited to 16, scaling out a solution beyond a certain scale requires additional DAGs to be deployed. You can replicate the process of building the initial DAG setup to create additional DAGs. The size of the additional DAGs could be identical to the initial DAG or different from the initial DAG.

## SOLUTION SIZING

Each Exchange Server role has distinct system requirements and must be sized according to the role-specific demands and the mailbox profile. The mailbox profile in an Exchange deployment describes the mailbox characteristics for a given mailbox size, such as number of messages per user per day and the average size of a message. Exchange Server is a storage-intensive workload and allows various storage options to be considered, ranging from internal server storage to shared storage such as Storage Area Network (SAN). The two significant sizing considerations are server and storage.

Server sizing considerations include:

- Determining the type of processor that is best suited for handling the Exchange Mailbox profile requirements
- Deciding the size of memory required and allocating the DIMMs to the processor memory channels to take advantage of full memory bandwidth
- Selecting the right host network adapters
- Selecting the right type of storage to achieve a balance between solution cost and performance. Storage sizing involves deciding the type of RAID, type of disks and number of disks—both from IOPS and capacity perspective—and intelligently mapping Exchange databases to the Storage Subsystem per the solution requirements.

**Note:** To calculate the processor, memory and storage sizing for a specific number and size of mailboxes and profiles, you can use the latest version of Exchange Server Role Requirements Calculator published by Microsoft.

Sizing provides the necessary capacity information for both server and storage hardware. Exchange Server 2016 infrastructure can be designed in different ways based on the size of the configuration and the number of HA and DR copies required. The infrastructure and application architecture must be designed for scalability and HA. Customers also have a variety of server form factor, storage, and disk options. To provide an easy-to-deploy, scalable, cost-effective and flexible solution, Dell EMC has developed a building block architecture model to build and scale out an Exchange infrastructure. The subsequent sections provide an overview and benefits of using DSS 7000/DSS 7500 and the building block architecture for Exchange 2016.

Over the years, storage cost has declined enabling large mailboxes to meet the ever-increasing business needs.

Based on the solution requirements, Table 1 and Table 2 provide more information about the Exchange server and Storage subsystem configuration. The firmware and driver versions are also provided for the tested solution.

**Table 1. Exchange Server configuration**

<b>MICROSOFT EXCHANGE SERVER</b>	<b>DSS 7500</b>
CPU	2 x Intel Xeon E5-2650 v4 processors
Memory	8 x 16 GB DDR4 at 2400 MHz
NIC	Dual-port QLogic BCM5720 embedded LOM Intel X540-T2 Ethernet Converged Network Adapter
RAID Controllers	MegaRAID SAS 9361-8i C600/X79 series chipset SATA RAID Controller
Internal Disks	2 x 1.2 TB SATA Solid State Drives (SSD)

**Table 2. Storage Subsystem configuration**

<b>STORAGE SYSTEM</b>	<b>DSS 7000 CHASSIS WITH DSS 7500 SLEDS</b>
Disks	45 x 8 TB 7.2 K RPM NL-SAS 3.5-inch disk: <ul style="list-style-type: none"> <li>• 42 x 8 TB 7.2 K RPM NL-SAS 3.5-inch drive in 14 x RAID 0 volumes (for database and logs)</li> <li>• 1 x 8 TB 7.2 K RPM NL-SAS 3.5-inch drive (for Restore LUN)</li> <li>• 1 x 8 TB 7.2 K RPM NL-SAS 3.5-inch drive (for Auto-Reseed volume)</li> <li>• 1 x 8 TB 7.2 K RPM NL-SAS 3.5-inch drive (spare drive)</li> </ul>
RAID Controller	MegaRAID SAS 9361-8i

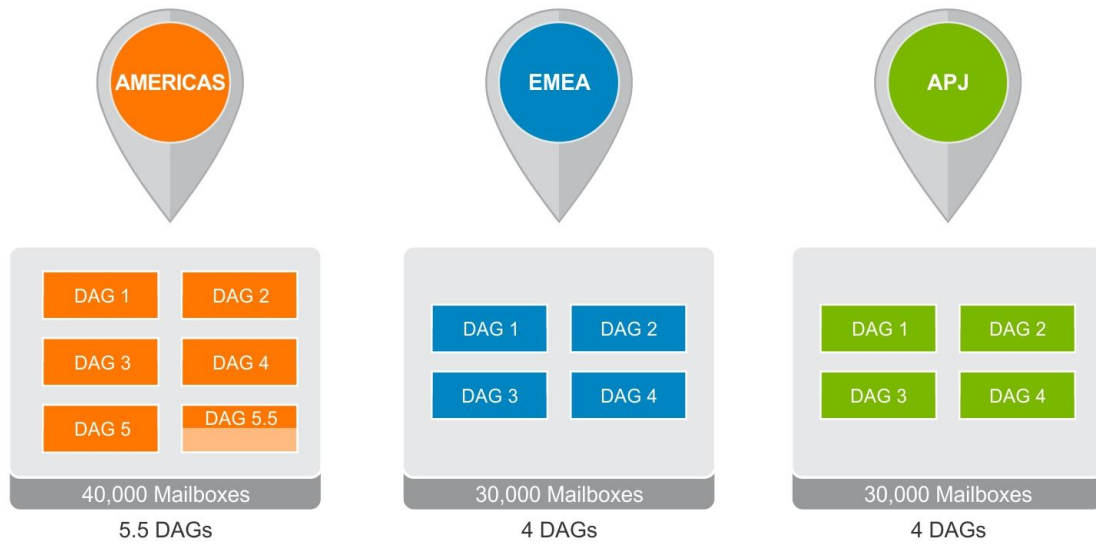
## EMAIL SOLUTION FOR A LARGE ORGANIZATION OF 100,000 MAILBOXES

Implementing a solution for 100,000 mailboxes using monolithic servers would increase the number of servers required and this in turn would increase the data center footprint. To counter this, the solution described in this guide uses the modular DSS 7000 chassis with DSS 7500 sleds where each chassis can accommodate two sleds. This model reduces the data center footprint by using lesser number of servers without compromising on the storage that is needed to support huge mailbox sizes.

Leveraging the architecture as explained in the Core Components section, an active-active user distribution model (4-copy DAG) email solution for 100,000 users is achieved by using:

- The internal drives in the DSS 7500 sleds for the Operating System and the Exchange binaries
- The internal storage of DSS 7000 chassis for storing the Exchange mailbox databases and transaction logs.

In this solution, each DAG can support up to 7600 mailboxes. If the 100,000 users are spread across three regions—North America, EMEA, and APJ, and each region requires 40,000, 30,000 and 30,000 mailboxes respectively, the number of DAGs needed to support this solution is as shown in Figure 10.



**Figure 10. Architecture to host 100,000 mailboxes**

This site-resilient design using 216 DSS 7500 sleds in 108 chassis (dual-node 44, 32, and 32 chassis in AMERICAS, EMEA, and APJ respectively) has been tested and validated in the lab by the Global Solutions Engineering group. Table 3, 4, 5 and 6 list the Environment Configuration, User Mailbox Configuration, Database Configuration and Activation Scenarios.

**Table 3. Environment configuration**

ENVIRONMENT CONFIGURATION	PER DATA CENTER 1 (SITE A) PER DAG	PER DATA CENTER 2 (SITE B) PER DAG	PER DAG	PER ENVIRONMENT
Number of Active Mailboxes (Normal Run Time)	3800	3800	7600	100,000 mailboxes using multiple DAGs
Number of Mailbox Servers per DAG	8	8	16	216

**Table 4. User mailbox configuration**

USER MAILBOX CONFIGURATION	
Number of User Mailboxes per Environment	100,000
Number of Mailboxes per Database	23
User Mailbox Size within Database	75350 MB
Transaction Logs Generated per Mailbox per Day	57
IOPS Profile per Mailbox	0.121
Read:Write Ratio per Mailbox	3:2

**Table 5. Database configuration**

DATABASE CONFIGURATION	
Number of Databases per DAG	336
Recommended Number of Mailboxes per Database	23
Available Database Cache/Mailbox	11.32 MB



**Table 6.     Activation scenarios**

ACTIVE DATABASE ACTIVATION CONFIGURATION (SITE A FAILURE) PER DAG	PER DATA CENTER 2 ACTIVE SERVER	ACTIVE MAILBOXES PER SERVER	PER DATA CENTER 2 (SITE B)	TOTAL ACTIVE MAILBOXES IN DC2 (SITE B)
Number of Active Databases (Secondary data center (DC2) Activation)	42	950	336	7600
Number of Active Databases (First Server Failure after SDC Activation)	48	1086	336	7600

## VERIFICATION

This section provides proof points for performance and resource utilization by using the sample implementation of 100,000 users as an example.

### MICROSOFT JETSTRESS VERIFICATION

The Storage Subsystem performance was verified to ensure that the storage meets the performance expectations for the given number of mailbox users and the mailbox profile. The Jetstress tool measures how well the storage system performs, and whether the Storage Subsystem meets the sizing requirements for a given Exchange mailbox profile. The Jetstress Disk Subsystem Throughput test was performed to measure how well the storage performs at peak load, while staying in the latency threshold established by Microsoft Exchange. In this sample implementation, the internal drives in the DSS 7500 sleds are used for the Operating System and the Exchange binaries and the internal storage of DSS 7000 chassis is used for storing the Exchange mailbox databases and transaction logs. Assuming that there is a site failure along with a single sled failure in the surviving site, the sample implementation was validated for 1086 active users served by one DSS 7500 sled.

A DSS 7500 sled hosts 48 active databases in the server failure scenario, thus capable of supporting up to 1086 active Exchange mailboxes. The Jetstress test was run on each of these servers and the transactional I/O performance in terms of achieved average transactional I/O per second was observed. The overall results from the Jetstress Disk Subsystem Throughput test depicted impressive performance, implying faster Exchange performance in terms of data transfers and significant increase in IOPS per server. The results of the Disk Subsystem Throughput test indicated that the storage can handle the peak load for the mailbox profile in consideration. The mailbox profile for the reference implementation used for running Jetstress Mailbox Profile test is shown in Table 7.

**Table 7.     Microsoft Jetstress verification results**

NUMBER OF MAILBOXES	MAILBOX SIZE	TARGET IOPS	ACHIEVED IOPS	STORAGE	RAID TYPE	NUMBER OF EXCHANGE DATABASES	VOLUME SIZE	I/O PROFILE
1086	70 GB	132	2429	90 x LFF 8 TB NL-SAS drives in DSS 7000 chassis	RAID 0	48	8 TB	150 messages per day (0.121 tested)

The Jetstress Mailbox Profile test verifies whether the storage system meets or exceeds the planned Exchange mailbox profile. The configuration shown in Table 7 was used to run the mailbox profile test with a planned target of 132 IOPS. The achieved target was 2429 IOPS. The results indicated that the storage was able to exceed the target transactional IOPS, which was well within the latency requirements.

For more information, see Microsoft Exchange 2016 Storage Reviewed Program solution for 100,000 70 GB mailboxes.

## MICROSOFT EXCHANGE LOAD GENERATOR VERIFICATION

When verifying server and storage for Exchange Server 2016, a best practice is to simulate the worst case scenario during peak time. Based on a number of data sets, peak load is generally twice the average workload throughout the remainder of the work day. In order to determine whether the infrastructure can handle both normal as well as peak loads, the test plan consisted of running Load Generator tests with peak and normal loads. To simulate peak load, the simulation day was set to four hours and test duration was set to eight hours. To simulate normal load, the simulation day and test duration were both set to eight hours. The LoadGen workload was configured to simulate Outlook 2007 Online with Outlook-150 profile. This profile simulated sending and receiving 150 messages per mailbox per day.

A performance test plan was designed to verify the end-to-end performance of Exchange Server 2016 on DSS 7500 with a steady state configuration (DAG with all nodes functional) as well as the failed host configuration (DAG with one failed node). Both the configurations were subjected to run under peak and normal loads. Two Load Generator instances were used to distribute the generated load across the two DAG nodes. Table 8 briefly describes the test plan for Load Generator tests.

**Table 8. Exchange LoadGen verification tests**

TESTS	TEST DESCRIPTION	DAG STATE	NUMBER OF ACTIVE DATABASES
Test 1	Peak load test with one site unavailable while one server in the surviving site is also unavailable: Simulation day: 4 hours Test length: 8 hours	One node down in the surviving site	48
Test 2	Normal load test with one site unavailable while one server in the surviving site is also unavailable: Simulation day: 8 hours Test length: 8 hours	One node down in the surviving site	48

The performance results pertaining to the middle four hours of an eight-hour test were considered in the analysis. All performance counters including processor, memory utilization and all the application-specific counters were measured by using a performance monitor.

Table 9 provides the results captured during the load tests performed.

**Table 9. Exchange Load Gen verification results**

PERFORMANCE METRICS	THRESHOLDS	TEST 1 (PEAK LOAD)	TEST 2 (NORMAL LOAD)
% Processor Utilization	< 70%	15%	8.3%
I/O Database Reads (Attached) Average Latency	< 20 ms	4.95	3.64
I/O Database Writes (Attached) Average Latency	< 20 ms	0.12	0.12
I/O Log Writes Average Latency	< 20 ms	0.122	0.114
RPC Average Latency	< 10 ms	1.91	1.88

## CONCLUSION

Ready Bundle for Microsoft Exchange with DSS 7000/DSS 7500 provides reliable and cost effective email solutions for very large enterprise customers. It is designed by following industry best practices. It is tested and validated to reduce design cycle time and implementation risks.

The Dell EMC DSS 7000/DSS 7500 system offers an excellent balance between performance and density. It is a highly efficient building block for any super-large organization. It delivers impressive storage capacity and IOPS performance in a dense 4U form factor, making it ideal for Exchange deployments.

The building block design used in Ready Bundle for Microsoft Exchange offers a highly scalable architecture for Exchange deployments. The proposed design consists of minimum server and storage required. DSS 7000/DSS 7500 can readily handle data intensive applications that require large storage capacity and I/O performance. It delivers the performance and availability required for mission-critical application such as email, and is a great hardware building block for large organizations with huge mailbox sizes.

## ADDITIONAL RESOURCES

DSS 7000/7500 Owner's Manual

MegaRAID SAS 9361-8i

Exchange Solution Reviewed Program (ESRP) - Storage

Exchange Server Role Requirements Calculator