

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 medium to large scale deployment. This solution uses PowerEdge R730xd server 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 © 2016-2017 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

EXECUTIVE SUMMARY	5
Introduction	5
Scope	5
Audience	6
SYSTEM COMPONENTS	6
Dell PowerEdge Server	6
Dell EMC Storage.....	8
Dell EMC Networking	9
Microsoft Exchange Server 2016	9
Application Architecture.....	10
Storage Architecture.....	12
Network Architecture	14
DESIGN PRINCIPLES	15
High Availability	15
Application Performance	15
Best Practices	16
COMPREHENSIVE SOLUTION DESIGN	17
Core Components	18
Auxiliary Components	19
Scalability	20
SOLUTION SIZING.....	20
Solution for Higher Capacity Mailboxes.....	20
Solution for Higher Number of Mailboxes.....	22
SAMPLE IMPLEMENTATIONS.....	23
Email Solution for a Medium Sized Organization of 5,000 Employees.....	23
Email Solution for a Large Regional Organization of 10,000 Employees	26

VERIFICATION 30

 Microsoft Jetstress Verification.....30

 Microsoft Exchange Load Generator Verification31

CONCLUSION 32

ADDITIONAL RESOURCES 32

EXECUTIVE SUMMARY

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. The Dell EMC Global Solutions Engineering team has developed a pre-architected Ready Bundle for Exchange Server to help customers simplify implementation. 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.

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

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

This technical guide describes design principles and solution components for the Ready Bundle for Exchange Server. Please contact your sales representatives about where to buy.

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 solution sizing and provides design guidance for higher capacity mailboxes and larger number of mailboxes.

The Sample Implementations section describes two sample implementations and provides proof points from the Exchange JetStress and Load Gen verification of the reference implementation.

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 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 design of Exchange 2016 solution on PowerEdge rack servers. In this guide, we have used the PowerEdge R730xd server as a building block, as it meets the requirements of Microsoft Exchange preferred architecture. Two scenarios have been considered to explain the design--5,000 and 10,000 mailbox implementations. Customers needing larger implementations can also use the building blocks to scale up and scale out as per their business needs. However, for implementations with

medium-to-high complexity, Dell EMC recommends that the customers engage with Dell EMC directly about their design needs in addition to using this guide.

We used Microsoft Exchange JetStress and Loadgen tools to verify the block design that is used as a foundation for our Exchange 2016 solution. We have also provided common user profile options.

AUDIENCE

This guide is intended for IT managers, messaging administrators and consultants interested in designing and deploying a cost-effective Exchange 2016 solution on PowerEdge R730xd servers for various user profiles. Users are expected to have sufficient understanding and 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. Dell EMC PowerEdge R730xd server is a 2U rack server that offers a balance between dense internal storage and compute capacity. It supports up to sixteen 3.5-inch NL-SAS drives, making it an ideal building block for Exchange for any medium size or large organization. Exchange solutions built on PowerEdge R730xd can be easily scale up or scale out based on the business needs.

Components for the Ready Bundle for Exchange Server 2016 also include Dell EMC Storage products, Dell networking products, Dell ProDeploy services and Dell ProSupport services. For network integration, Active Directory and data migration, we also have Dell Consulting Services for Exchange. Please contact your sales and services representatives for more details. This guide will focus on the system design.

The following subsections describe the hardware and software components of the Ready Bundle for Exchange Server 2016.

DELL POWEREDGE SERVER

Dell PowerEdge Servers are built to support the work that IT organizations do. They are engineered to handle the most demanding business applications and are designed with specific features to better run workloads like HPC, collaboration, database, ERP, business intelligence and data warehousing.

As the foundation for a complete and adaptive IT solution, Dell EMC PowerEdge servers deliver exceptional performance and management advantages that power the business applications that our customers run most.

Combined with the innovative OpenManage systems management portfolio and industry-leading workload solutions, PowerEdge servers provide technology that is intelligent, yet simple, giving you the power to do more in even the most complex environments.

The latest generation of servers responds to customer needs in the following areas:

- Memory capacity and scalability—much larger memory footprints
- Virtualization performance—more processor cores and denser memory
- Systems management—complete lifecycle management by using iDRAC with Lifecycle Controller and monitoring and updating capabilities by using OpenManage Essentials
- Energy efficiency—comprehensive optimizations, including Dell EMC OpenManage Power Center
- Infrastructure flexibility—innovations like select network adapters, offering more and better I/O options
- Reliability—even more RAS features, including a failsafe hypervisor option on most servers

PowerEdge R730xd is a 2-socket CPU, 2U, multi-purpose server, offering an excellent balance of ultra-dense internal storage, redundancy, and value in a compact form factor. It provides exceptional storage capacity, memory scalability and IOPS performance. The internal RAID controller provides a range of RAID levels for improved storage reliability. Major features are:

- 2 x Intel Xeon E5-2600 v3/v4 product family
- Up to 1536 GB of memory with 24 DIMMs
- Up to 128 TB maximum raw internal storage using 8 TB Near-Line SAS (NL-SAS) drives
- Support for multiple internal drive configurations that include:
 - Up to 16 (3.5-inch LFF disk drives) with 12 in the front bay and 4 in the internal tray plus two 2.5-inch SFF back-accessible drives
 - Up to 12 (3.5-inch LFF disk drives) in the front bay plus two 2.5-inch SFF back-accessible drives
 - 24 (2.5-inch SFF disk drives) front loading drive bays plus two 2.5-inch SFF back-accessible drives
- Integrated RAID support by using PERC H330, PERC H730, PERC H730P, and External JBOD RAID support with PERC H830
- 6 x PCIe 3.0 expansion slots
- Dell EMC OpenManage Essentials, Dell EMC Management Console, Dell EMC OpenManage Power Center, and Dell EMC OpenManage Connections

The PowerEdge R730xd server, when configured with sixteen 3.5-inch large form factor 8 TB NL-SAS drives, provides up to 128 TB raw storage capacity. In this configuration, 12 drives are installed in the front drive bay of the chassis and 4 additional drives are installed in the internal drive tray. Figure 1 shows this configuration.




Figure 1. PowerEdge R730xd LFF chassis with an internal drive tray

Note: The PowerEdge R730xd chassis with 16 LFF drives can be configured to use Intel Xeon E5-2600 family of processors with a maximum of 10 cores per processor.

DELL EMC STORAGE

The Dell EMC Storage MD1400 direct-attached storage (DAS) enclosure with 12Gb SAS throughput is specifically engineered to work with the 13th generation of PowerEdge servers that use the newest line of PowerEdge RAID Controller 9 (PERC9) 12 Gb SAS HBA cards. Scaling capacity built in the 13th generation of PowerEdge servers enables end-to-end 12Gb solutions with exceptional storage flexibility and IO performance for such applications as:

- High-performance databases
- Streaming digital media
- Storage-intensive applications

The Dell EMC Storage MD1400 direct-attached storage offers seamless expansion for PowerEdge servers with the PERC H830 HostRAID adapter. It provides customers the flexibility to expand storage as their business grows. You can easily expand your server capacity with twelve 3.5-inch SAS HDDs in a 2U array and up to 8 arrays with a single PERC H830 Host RAID adapter. Table1 lists the features of MD1400 direct-attached storage.

Table 1. Dell Storage MD1400 features

FEATURE	SPECIFICATION
Drives	Up to 12 hot-pluggable 3.5" and 2.5" drives (2.5" available with adapter)
Drive Performance and Capacities	3.5" NL-SAS 12Gb HDD (7.2K): 1TB, 2TB, 4TB, 6TB, 8TB 2.5" SAS 6Gb HDD (7.2K): 500GB 2.5" SAS 6Gb HDD (10K): 300GB, 600GB, 1.2TB, 1.8TB 2.5" NL-SAS 12Gb HDD (7.2K): 2TB 2.5" SAS 12Gb SED (15K): 600GB 2.5" SAS 6Gb HDD (15K): 300GB, 600GB 2.5" SAS 12Gb SSD: 200GB, 400GB, 800GB (WI); 200GB, 400GB, 800GB, 1.6TB (MU); 800GB, 1.6TB (RI)
Maximum Capacity (per enclosure)	Up to 96TB when using 12 x 8TB 7.2K RPM 3.5" SAS HDDs
Expansion Capabilities	PERC H830 HBA enables expansion to 8 enclosures, PCIe 3.0, 12 Gbps SAS, dual-port, 4 ports per enclosure
Host Connectivity Unified Mode	Unified mode (single path) for daisy chaining of up to 8 enclosures per PERC H830 (4 enclosures per port, single path) Unified mode (recommended redundant path) for daisy chaining up to 4 enclosures per PERC H830 (4 enclosures connected to both ports via redundant path cabling)
Host Connectivity Split-Mode/Dual-Host Access	Split mode with dual Enclosure Management Modules (EEM) providing direct connectivity to drives 0 through 5 and a separate connectivity to drives 6 through 11
Enclosure Management Modules (EMM)	2 EMM provide redundant enclosure management capability
RAID Levels	0, 1, 5, 6, 10, 50 and 60
Connectivity per EMM	4 mini-SAS HD connector for connection to the host or expansion
Service Management	USB mini-B connector (for factory use only)

DELL EMC NETWORKING

Dell EMC Networking offers S-Series and N-Series high density 100M/1G/10G/40GbE top-of-rack (ToR) switches specially built for applications in high-performance data centers and computing environments. S-Series S4048T-ON is a 10G 10BaseT 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 (RJ45) ports as well as six 40GbE QSFP+ up-links to conserve valuable rack space and simplify the migration to 40Gbps in the data center core. Each 40GbE QSFP+ up-link can also support four 10GbE (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/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 the Ready Bundle for Exchange Server 2016, and explain how the infrastructure sub-systems are architected. To demonstrate these principles, the Dell EMC Global Solutions Engineering Group has designed two Ready Bundle solutions for reference purposes.

Medium sized organization of 5,000 employees

This organization has 5,000 employees in different regions with most of the workforce based out of two main facilities—some are in branch offices and many are remote workers. The design enables each employee to have a mailbox of 20GB. This site resilient solution spread across two data centers (Site A and Site B) makes use of active-active user distribution model.

Large regional organization of 10,000 employees

This organization has 10,000 employees based in many locations across different regions, including several main facilities and a dozen satellite offices as well as many remote workers. The design enables each employee to have a higher capacity mailbox of 35 GB. This site resilient solution spread across two data centers (Site A and Site B) makes use of active-active user distribution model.

Both the above cases are detailed under the section Sample Implementations.

Figure 2 shows the external data center architecture.

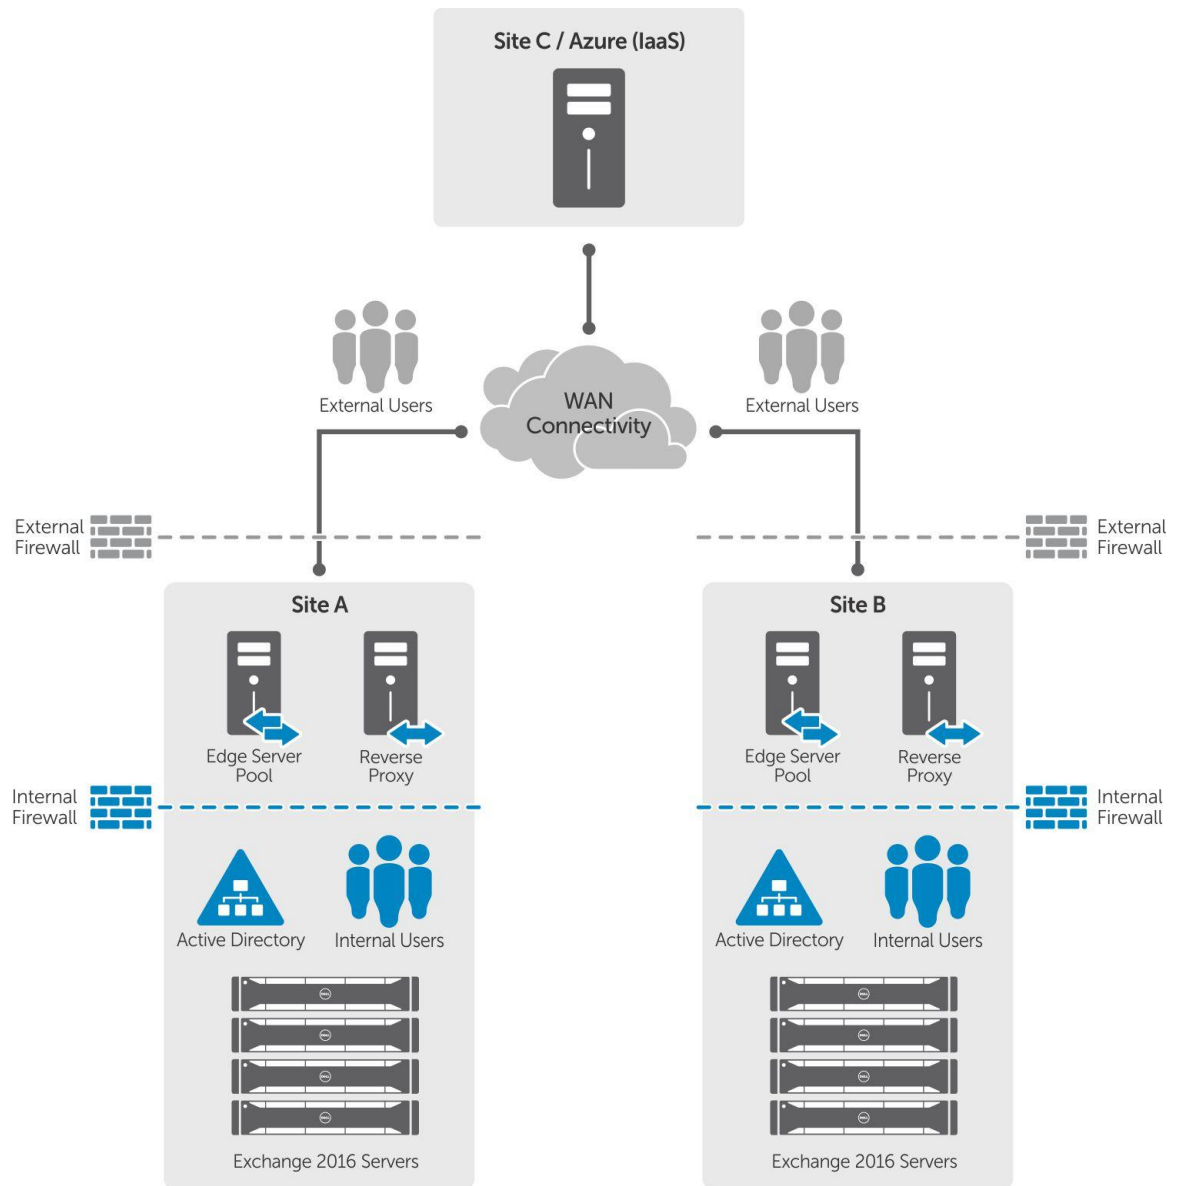


Figure 2. External data center architecture for Exchange deployment

STORAGE ARCHITECTURE

The PowerEdge R730xd server supports multiple internal drives. A RAID controller can be used to create independent RAID0 volumes by using the internal drives. These RAID0 volumes host Exchange databases and transaction logs.

When a PowerEdge R730xd server is configured with 16 LFF drives, fourteen independent RAID0 disks can be used to store copies of Exchange databases. These RAID0 disks can be created by using 12 disks in the front bay and 2 disks in the internal drive tray. One of the remaining two drives in the internal drive tray can be used for Exchange restore LUN and the other can be used for the RAID0 disks. The rear-accessible 2.5-inch drives can be configured as a RAID1 disk to deploy the operating system. Figure 3 illustrates this configuration.

While configuring RAID0 volumes, ensure that the Write Policy and Read Policy are set to Write Back and Read Ahead respectively and the Disk Cache Policy is disabled.

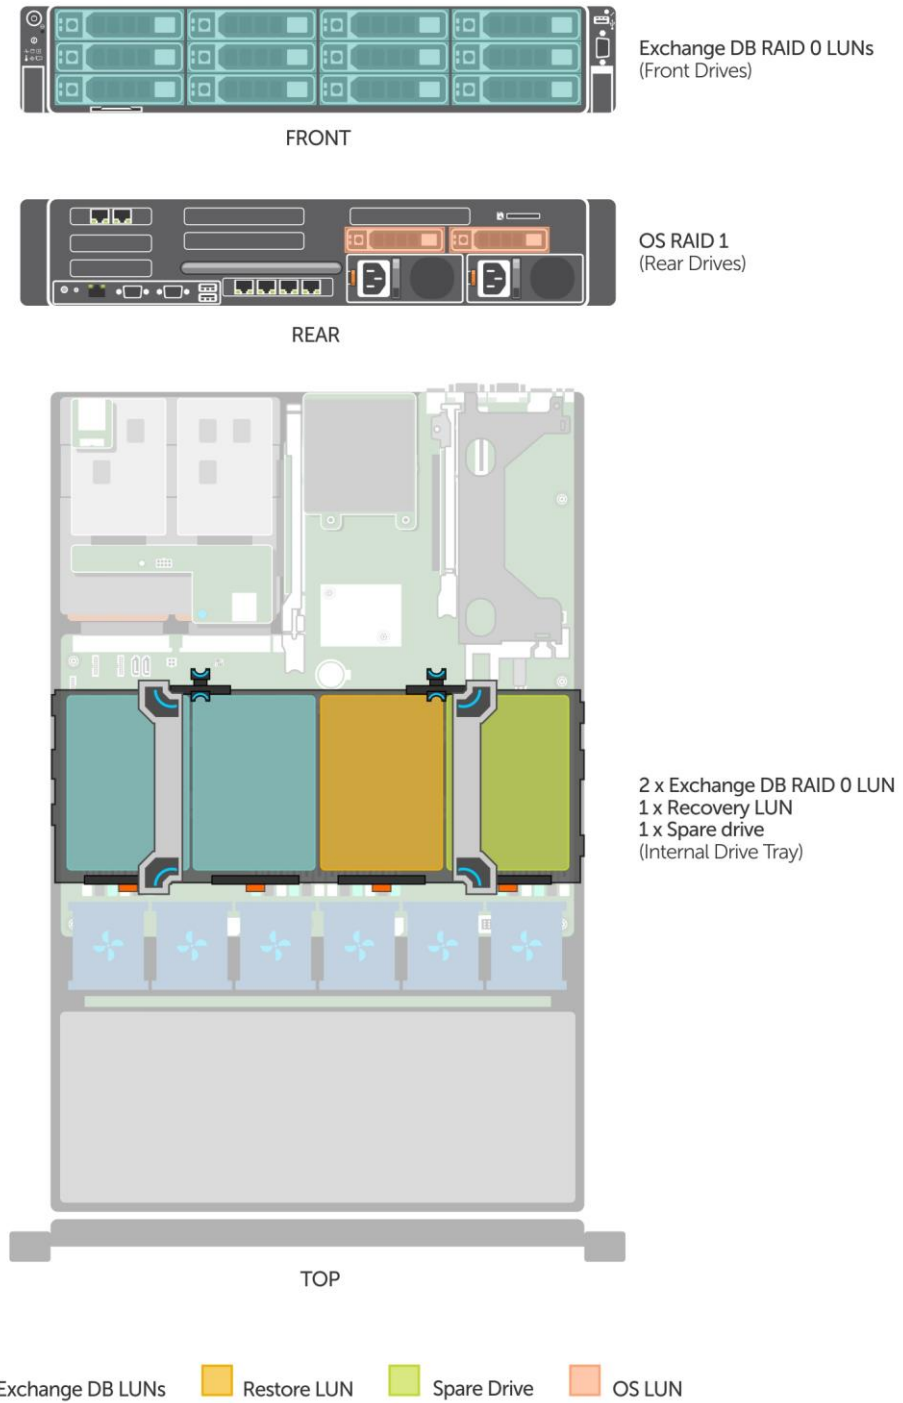


Figure 3. RAID LUN layout for the Exchange databases

NETWORK ARCHITECTURE

Microsoft Exchange Server 2016 network requirements are simple and easy to deploy. A single 10 GbE network can be leveraged for both client traffic and database replication. Even though Microsoft Exchange Server 2016 PA does not leverage NIC teaming and a dedicated network for replication, in our solution we have leveraged NIC teaming to provide multi-tier HA across the network. In addition, two 1GbE NICs are also available which can be used for database replication 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 end user and client connectivity traffic and replication traffic between DAG members. To provide redundancy and HA for network connections, multiple network adapters employed in the server can be connected to redundant Top-of-Rack (ToR) network switches. Figure 4 illustrates this.

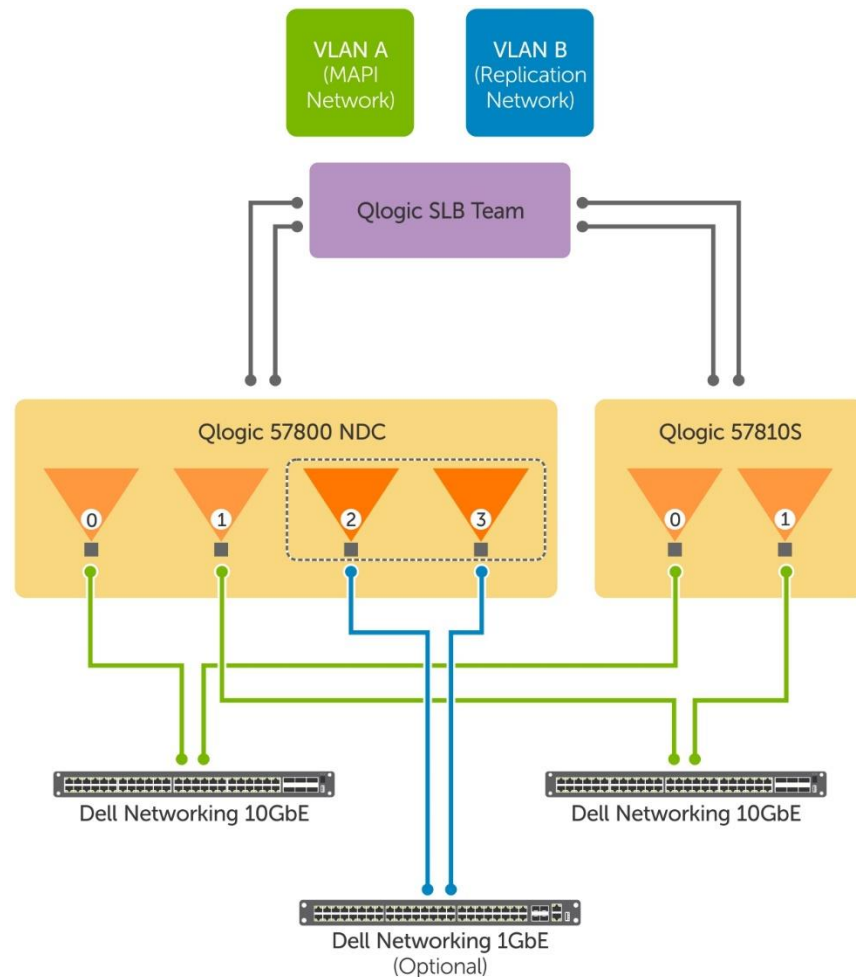


Figure 4. Network architecture for the Exchange deployment

The PowerEdge R730xd server platform supports multiple 10 GbE and/or 1 Gb network adapters. By using 10 GbE network adapters, Virtual Networks (VLANs) can be used to isolate the MAPI network from the replication traffic. As shown in Figure 4, the QLogic 57800C Network Daughter Card (NDC) provides two 10 GbE and two 1 Gb network ports. An additional QLogic 57810S can be employed to provide redundant network paths. The 10 GbE ports from the NDC and an add-on NIC can be teamed to provide an aggregate bandwidth of 40 Gbps per server. By using the Smart Load Balancing (SLB) teaming mode, all network paths remain active during normal operation. Optionally, the 1 Gb network ports in the QLogic 57800C NDC can be used to provide dedicated replication network between the Exchange mailbox servers. This eliminates the need for sharing the network bandwidth between MAPI and replication networks.

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 could provide HA and Disaster Recovery (DR) capabilities.

We build high availability into the Ready Bundle for Exchange Server 2016 so it can withstand server failures as well as site failure. Our design is based on the PoD concept where four servers are stacked across two Data Centers—two in the Primary Data Center (PDC) and two in the Secondary Data Center (SDC). This design is in line with Microsoft Preferred Architecture (PA) wherever possible.

Application-level High Availability

The pools of application servers should be utilized to provide a highly available service. The DAGs in Exchange provide HA 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 load balance the client requests to the Mailbox servers in the Exchange deployment.

Infrastructure-level High Availability

In a solution infrastructure, resources such as server, storage, network path and switches should be highly available. Application availability requires that the infrastructure builds redundancy for each component. RAID disks for storage provide data HA. Multiple network adapters and switches connecting to the Exchange infrastructure and the data center network build resiliency in network connectivity. Both Storage and Network Architecture explained earlier in this guide take into account Infrastructure-level high availability in the design.

APPLICATION PERFORMANCE

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

BEST PRACTICES

Microsoft recommends disabling Logical Processor/Hyper Threading when deploying Exchange on physical servers and configuring the System Profile Settings to Performance in the BIOS.

Microsoft PA (Preferred Architecture) recommends deploying Exchange on physical servers. The recommended server configurations are:

- 2U, dual socket servers (20-24 cores)
- up to 96 GB of memory
- a battery-backed write cache controller
- 12 or more large form factor drive bays within the server chassis

When 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 will be activated in the surviving data center to provide email services to the users.

COMPREHENSIVE SOLUTION DESIGN

A high-level diagram of the solution architecture is illustrated in Figure 2. It shows a customer scenario that consists of two data centers (Site A and Site B) 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 shows the Site A data center architecture along with the infrastructure management services.

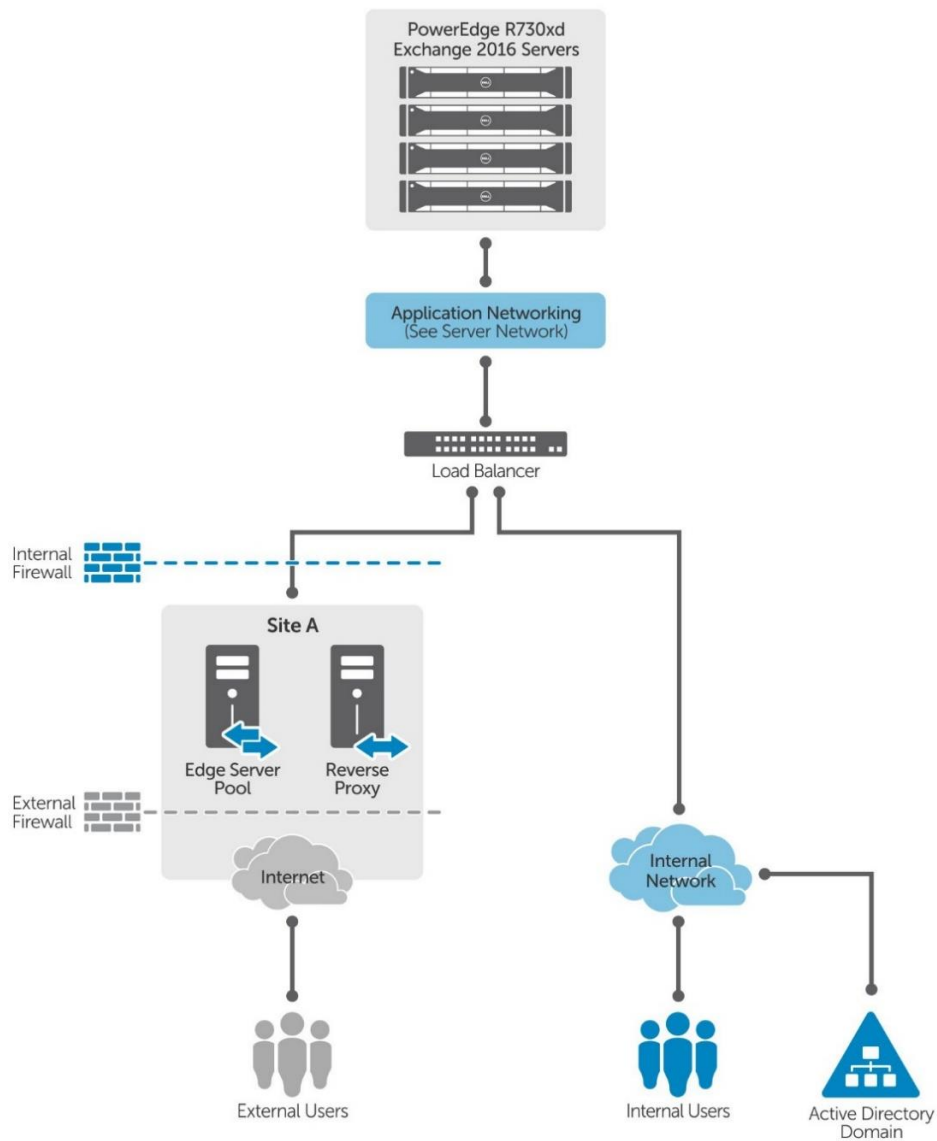


Figure 5. Site A data center architecture for Exchange deployment

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 will have the same set of infrastructure components. The PowerEdge R730xd servers 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 and 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 and through an Edge Server pool configured at the site.

CORE COMPONENTS

To explain the building block architecture, this guide uses the concept of a PoD, which is a standardized configuration of the minimum server and storage resources sized to meet the solution requirements based on a given mailbox profile. The configuration for each member within a PoD should be identical. To support larger number of mailboxes (of the same mailbox profile), just increase the number of PoDs.

Figure 6 shows the logical representation of a PoD. Each physical server has identical configurations with respect to memory, CPU and storage. In this model, you can use internal disks in the R730xd servers or external direct-attached storage (DAS) MD1400 to host mailbox databases and logs. You can increase the number of mailboxes by scaling out the building blocks.



Figure 6. PoD architecture for Exchange deployment

Figure 6 illustrates the building block architecture of our Ready Bundle for Exchange Server 2016. It has a single DAG spanning across two sites with an active-active user distribution model. Each database has four copies across two sites, and the design is site resilient.

For a site-resilient configuration, Dell MEC 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), load balancers, and so on. The connectivity from the Exchange infrastructure to these infrastructure services and end users must be resilient and highly

available. For our Ready Bundle, we employed Dell EMC Networking 10 GbE Top-of-Rack (ToR) switches. The networking connectivity is illustrated in Figure 7.

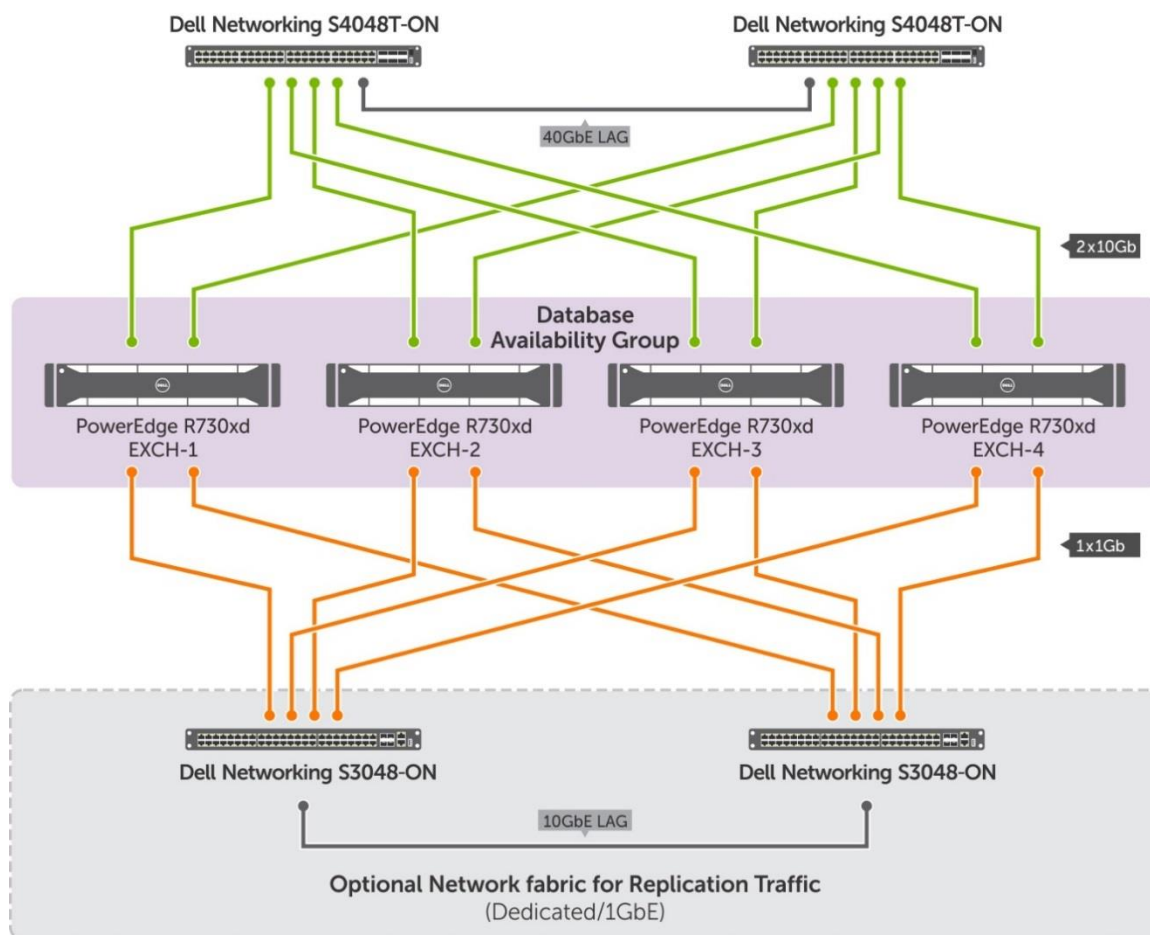


Figure 7. Detailed network architecture for the Exchange reference implementation

As mentioned in the System Components section, a total of four 10GbE network ports are available per server. These network ports are teamed to provide an aggregate network bandwidth of 40 Gbps per server. As shown in Figure 7, 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. Because the QLogic 57800C NDC provides two 1 Gb ports along with the 10 GbE ports, it is possible to create a separate network fabric for replication traffic. This, however, is not cost effective as it involves two additional switches to support the server connectivity.

AUXILIARY COMPONENTS

After the consolidation of server roles, Exchange 2016 has only two server roles – Mailbox and Edge Transport. Like in 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/services based on their needs.

SCALABILITY

Using PoDs as building blocks increases scalability. As organizations grow, you can accommodate additional mailboxes simply by adding one or more PoDs. Given that 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 first DAG to create additional DAGs. The size of the additional DAGs could be identical to the first DAG or different from the first 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.

To calculate the CPU, memory and storage sizing for a specific number and size of mailboxes and profiles, you can use the Exchange Server Role Requirements Calculator version 7.9.

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 block architecture model that uses the PoD concept to build and scale out an Exchange infrastructure. The subsequent sections provide an overview and benefits of the PowerEdge R730xd server and the building block architecture for Exchange 2016.

SOLUTION FOR HIGHER CAPACITY MAILBOXES

Over the years, storage cost has declined and Exchange supports the use of Direct Attached Storage (DAS), enabling large mailboxes to meet the ever-increasing business needs.

This section outlines the site resilient email solution for a medium size organization of 10,000 employees with 35 GB per mailbox. Dell EMC Storage MD1400 is attached directly to a PowerEdge R730xd to expand storage capacity. To achieve 10,000 mailboxes of 35 GB each, we stack four PoDs where each PoD has four PowerEdge R730xd servers and four Dell Storage MD1400 attached directly to each server.

Based on the solution requirements, Table 2, 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 2. Exchange Server configuration

MICROSOFT EXCHANGE SERVER SYSTEM	POWEREDGE R730XD SERVER WITH 3.5" HDD CHASSIS
CPU	2 × Intel Xeon E5-2660 v4 processor with 10-cores
Memory	96 GB DDR4
NIC	QLogic BCM57810
RAID Controllers	PowerEdge RAID Controller H730P Mini Firmware version: 25.3.0.0016 Storport Driver Version 6.3.9600.18229 Driver version 6.600.21.08 PowerEdge RAID Controller H830P Firmware version: 25.3.0.0016 Storport Driver Version 6.3.9600.18229 Driver version 6.600.21.08
Internal Disks	2 x 1.2 TB SAS 2.5-inch 10K RPM disk drives (Operating System and Application)

Table 3. Storage Subsystem configuration (PowerEdge R730xd server storage)

STORAGE SYSTEM	DELL POWEREDGE R730XD INTERNAL 3.5-INCH DRIVES
Disks	16 x 8 TB 7.2K RPM NL-SAS 3.5-inch disk: 14 x 8 TB 7.2K RPM NL-SAS 3.5-inch drive in 14 x RAID 0 (for DB and Log) 1 x 8 TB 7.2K RPM NL-SAS 3.5-inch drive (for Restore LUN) 1 x 8 TB 7.2K RPM NL-SAS 3.5-inch drive (spare drive)
RAID Controller	PowerEdge RAID Controller H730P Mini (Firmware version: 25.3.0.0016)

Table 4. Storage Subsystem configuration (Dell EMC Storage MD1400)

STORAGE SYSTEM	STORAGE MD1400 3.5-INCH DRIVE
Disks	12 x 8 TB 7.2K RPM NL-SAS 3.5-inch disk: 10 x 8 TB 7.2K RPM NL-SAS 3.5-inch drive in 10 x RAID 0 (for DB and Log) 1 x 8 TB 7.2K RPM NL-SAS 3.5-inch drive (for Restore LUN) 1 x 8 TB 7.2K RPM NL-SAS 3.5-inch drive (spare drive)
RAID Controller	Dell EMC PowerEdge RAID Controller H830 (Firmware version: 25.3.0.0016)

SOLUTION FOR HIGHER NUMBER OF MAILBOXES

Organizations that need a higher number of small to medium size mailboxes can leverage higher compute and memory in the server configuration with lower capacity drives.

As an example, the PoD based building block architecture can be leveraged to attain an email solution with 10 GB mailbox size, where each user can send/receive 100 messages per day. In this site resilient solution, four PoDs are stacked, so that email services can be provided to 20,000 users in an active-active distribution model. The following tables provide more information about the server and storage configuration. The firmware and driver versions are also provided for the tested solution.

Table 5. Exchange Server configuration

MICROSOFT EXCHANGE SERVER SYSTEM	POWEREDGE R730XD SERVER WITH 3.5" HDD CHASSIS
CPU	2 x Intel Xeon E5-2660 v4 processor with 10 cores
Memory	96 GB DDR4
NIC	QLogic BCM57810
RAID Controllers	PowerEdge RAID Controller H730P Mini Firmware version: 25.3.0.0016 Storport Driver Version 6.3.9600.18229 Driver version 6.600.21.08
Internal Disks	2 x 1.2 TB SAS 2.5-inch 10K RPM disk drives (Operating System and Application)

Table 6. Storage Subsystem configuration (PowerEdge R730xd server storage)

STORAGE SYSTEM	POWEREDGE R730XD INTERNAL 3.5-INCH DRIVES
Disks	16 x 8 TB 7.2K RPM NL-SAS 3.5-inch disk: 14 x 8 TB 7.2K RPM NL-SAS 3.5-inch drive in 14 x RAID 0 (for DB and Log) 1 x 8 TB 7.2K RPM NL-SAS 3.5-inch drive (for Restore LUN) 1 x 8 TB 7.2K RPM NL-SAS 3.5-inch drive (spare drive)
RAID Controller	PowerEdge RAID Controller H730P Mini (Firmware version: 25.3.0.0016)

SAMPLE IMPLEMENTATIONS

This section outlines two sample implementations for a medium and large organization of 5,000 and 10,000 mailboxes respectively. We select these two customer profiles because they represent many customers of similar needs. Larger or smaller implementations could also be achieved following the design principles illustrated in these two samples.

EMAIL SOLUTION FOR A MEDIUM SIZED ORGANIZATION OF 5,000 EMPLOYEES

Leveraging the PoD architecture as explained in the Core Components section, an active-active user distribution model (4 copy DAG) email solution for 5,000 users can be achieved by stacking two PoDs as shown in Figure 8.




Figure 8. PoD architecture to host 5,000 mailboxes

This site resilient design using eight PowerEdge R730xd servers has been tested and validated in the lab by the Global Solutions Engineering group. Table 7, 8, 9, 10, 11 and 12 list the Environment Configuration, User Mailbox Configuration, Database Configuration and Activation Scenarios.

Table 7. Environment configuration

ENVIRONMENT CONFIGURATION	PER DATA CENTER 1 (SITE A)	PER DATA CENTER 2 (SITE B)	PER DAG	PER ENVIRONMENT
Number of Active Mailboxes (Normal Run Time)	2500	2500	5000	5000
Number of Mailbox Servers per DAG	4	4	8	8
Number of Lagged Copy Servers per DAG	0	0	0	0

Table 8. User Mailbox configuration

USER MAILBOX CONFIGURATION	
Number of User Mailboxes per Environment	5000
Number of Mailboxes per Database	45
User Mailbox Size within Database	21560MB
Transaction Logs Generated per Mailbox per Day	39
IOPS Profile per Mailbox	0.10
Read:Write Ratio per Mailbox	3:2

Table 9. Database configuration

Database Configuration	
Number of Databases per DAG	112
Recommended Number of Mailboxes per Database	45
Available Database Cache / Mailbox	9.66 MB

Table 10. 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)	28	1250	112	5000
Number of Active Databases (First Server Failure after SDC Activation)	38	1696	112	5000
Number of Active Databases (Second Server Failure after SDC Activation)	--	--	--	--

Table 11. Exchange Server configuration

MICROSOFT EXCHANGE SERVER SYSTEM	POWEREDGE R730XD SERVER WITH 3.5" HDD CHASSIS
CPU	2 x Intel Xeon E5-2660 v4 processor with 10-cores
Memory	96 GB DDR4
NIC	QLogic BCM57810
RAID Controllers	PowerEdge RAID Controller H730P Mini Firmware version: 25.3.0.0016 Storport Driver Version 6.3.9600.18229 Driver version 6.600.21.08 PowerEdge RAID Controller H830P Firmware version: 25.3.0.0016 Storport Driver Version 6.3.9600.18229 Driver version 6.600.21.08
Internal Disks	2 x 1.2 TB SAS 2.5-inch 10K RPM disk drives (Operating System and Application)

Table 12. Storage Subsystem configuration (PowerEdge R730xd server storage)

STORAGE SYSTEM	POWEREDGE R730XD INTERNAL 3.5-INCH DRIVES
Disks	16 x 8 TB 7.2K RPM NL-SAS 3.5-inch disk: 14 x 8 TB 7.2K RPM NL-SAS 3.5-inch drive in 14 x RAID 0 (for DB and Log) 1 x 8 TB 7.2K RPM NL-SAS 3.5-inch drive (for Restore LUN) 1 x 8 TB 7.2K RPM NL-SAS 3.5-inch drive (spare drive)
RAID Controller	PowerEdge RAID Controller H730P Mini (Firmware version: 25.3.0.0016)

EMAIL SOLUTION FOR A LARGE REGIONAL ORGANIZATION OF 10,000 EMPLOYEES

For this sample implementation, we have considered an organization of 10,000 employees based in many locations across different regions, including several main facilities and a dozen satellite offices as well as many remote workers. In this design, each employee can have a mailbox of 35 GB. This site resilient solution spread across two data centers (Site A and Site B) makes use of active-active user distribution model in a PoD architecture. Figure 9 shows how four PoDs are stacked to host 10,000 mailboxes.



Figure 9. PoD architecture to host 10,000 mailboxes

This site resilient tested solution employs 16 PowerEdge R730xd servers and the same number of Dell EMC Storage MD1400 attached to each server across two sites. Environment Configuration, User Mailbox Configuration, Database Configuration and Activation Scenarios are listed in the following tables.

Table 13. Environment configuration

ENVIRONMENT CONFIGURATION	PER DATA CENTER 1 (SITE A)	PER DATA CENTER 2 (SITE B)	PER DAG	PER ENVIRONMENT
Number of Active Mailboxes (Normal Run Time)	5000	5000	10000	10000
Number of Mailbox Servers per DAG	8	8	16	16
Number of Lagged Copy Servers per DAG	0	0	0	0

Table 14. User mailbox configuration

USER MAILBOX CONFIGURATION	
Number of User Mailboxes per Environment	10,000
Number of Mailboxes per Database	26
User Mailbox Size within Database	37566MB
Transaction Logs Generated per Mailbox per Day	39
IOPS Profile per Mailbox	0.10
Read:Write Ratio per Mailbox	3:2

Table 15. Database configuration

DATABASE CONFIGURATION	
Number of Databases per DAG	384
Recommended Number of Mailboxes per Database	26
Available Database Cache / Mailbox	11.23 MB

Table 16. 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)	48	1250	384	10,000
Number of Active Databases (First Server Failure after SDC Activation)	56	1458	384	10,000
Number of Active Databases (Second Server Failure after SDC Activation)	--	--	--	--

Table 17. Exchange Server configuration

MICROSOFT EXCHANGE SERVER SYSTEM	POWEREDGE R730XD SERVER WITH 3.5" HDD CHASSIS
CPU	2 x Intel Xeon E5-2660 v4 processor with 10-cores
Memory	96 GB DDR4
NIC	QLogic BCM57810
RAID Controllers	PowerEdge RAID Controller H730P Mini Firmware version: 25.3.0.0016 Storport Driver Version 6.3.9600.18229 Driver version 6.600.21.08 PowerEdge RAID Controller H830P Firmware version: 25.3.0.0016 Storport Driver Version 6.3.9600.18229 Driver version 6.600.21.08
Internal Disks	2 x 1.2 TB SAS 2.5-inch 10K RPM disk drives (Operating System and Application)

Table 18. Storage Subsystem configuration (PowerEdge R730xd server storage)

STORAGE SYSTEM	POWEREDGE R730XD INTERNAL 3.5-INCH DRIVES
Disks	16 x 8 TB 7.2K RPM NL-SAS 3.5-inch disk: 14 x 8 TB 7.2K RPM NL-SAS 3.5-inch drive in 14 x RAID 0 (for DB and Log) 1 x 8 TB 7.2K RPM NL-SAS 3.5-inch drive (for Restore LUN) 1 x 8 TB 7.2K RPM NL-SAS 3.5-inch drive (spare drive)
RAID Controller	PowerEdge RAID Controller H730P Mini (Firmware version: 25.3.0.0016)

Table 19. Storage Subsystem configuration (Storage MD1400)

STORAGE SYSTEM	STORAGE MD1400 3.5-INCH DRIVE
Disks	12 x 8 TB 7.2K RPM NL-SAS 3.5-inch disk: 10 x 8 TB 7.2K RPM NL-SAS 3.5-inch drive in 10 x RAID 0 (for DB and Log) 1 x 8 TB 7.2K RPM NL-SAS 3.5-inch drive (for Restore LUN) 1 x 8 TB 7.2K RPM NL-SAS 3.5-inch drive (spare drive)
RAID Controller	PowerEdge RAID Controller H830 (Firmware version: 25.3.0.0016)

VERIFICATION

This section provides proof points for performance and resource utilization by using the sample implementation of 10,000 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. This sample implementation describes the use of internal storage offered on each PowerEdge R730xd server in the solution infrastructure. Assuming that there is a site failure along with a single server failure in the surviving site, the sample implementation was validated for 1,429 active users served by one PowerEdge R730xd server.

A PowerEdge R730xd server hosts 32 active databases in the server failure scenario, thus capable of supporting up to 1,429 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 20.

Table 20. Microsoft JetStress verification results

NUMBER OF MAILBOXES	MAILBOX SIZE (GB)	TARGET IOPS	ACHIEVED IOPS	STORAGE	RAID TYPE	NUMBER OF EXCHANGE DATABASES	VOLUME SIZE	I/O PROFILE
1,429	20 GB	172.9	911	16 x LFF 8 TB NL-SAS drives in PowerEdge R730xd	RAID 0	32	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 20 was used to run the mailbox profile test with a planned target of 172.9 IOPS. The achieved target was 911 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, refer to Microsoft Exchange 2016 Storage Reviewed Program solution for 10,000 20GB 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 converged 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, and 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 PowerEdge R730xd 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 21 briefly describes the test plan for Load Generator tests.

Table 21. Exchange Load Gen 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	32
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	32

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 22 provides the results captured during the load tests performed.

Table 22. Exchange Load Gen verification results

PERFORMANCE METRICS	THRESHOLDS	TEST 1 (PEAK LOAD)	TEST 2 (NORMAL LOAD)
%Processor Utilization	< 70 %	15%	9%
I/O Database Reads (Attached) Average Latency	< 20 ms	1.67	2.87
I/O Database Writes (Attached) Average Latency	< 20 ms	0.13	0.158
I/O Log Writes Average Latency	< 20 ms	0.116	0.102
RPC Average Latency	< 10 ms	4.18	2.41

CONCLUSION

Dell EMC Ready Bundle for Exchange Server 2016 provides reliable and cost effective email solutions. It is designed following industry best practices. It is tested and validated to reduce design cycle time and implementation risks.

The PowerEdge R730xd server used in this guide for the sample implementation offers an excellent balance between performance and management. It is a highly efficient and cost-effective hardware building block for any medium size or large organization. It delivers impressive storage capacity & IOPS performance in a dense 2U form-factor, making it ideal for Exchange deployments.

The PoD design used in Dell EMC's Ready Bundle for Exchange Server 2016 offers a highly scalable architecture for Exchange deployments. The proposed PoD design consists of minimum server and storage required. Dell EMC PowerEdge R730xd server with 16 LFF drives provides sufficient compute and storage capacity for most of the medium size and large organizations. Dell EMC Storage MD1400 direct-attached storage provides additional storage capacity for implementations with larger mailbox sizes. The PoDs can be scaled up and scaled out efficiently to meet different size or needs.

ADDITIONAL RESOURCES

Dell EMC PowerEdge R730xd

Dell EMC PowerEdge RAID Controller H730P

Dell EMC PowerEdge RAID Controller H830

Dell EMC Storage MD1400 Direct-Attached Storage

Exchange Solution Reviewed Program (ESRP) - Storage