Dell EMC Microsoft Exchange 2016 Solution

Design Guide for implementing Microsoft Exchange Server 2016 on Dell EMC R740xd servers and storage

Dell Engineering
October 2017

Design Guide

Revisions

Date	Description
October 2017	Initial release

Table of contents

Executive summary

For many organizations planning to implement Exchange Server 2016, finding cost effective solutions at the same time meeting user demand for large mailboxes is critical. Choosing the right architecture and faster time to value are also key to success. To help customers address these concerns and simplify implementation, the Dell EMC Global Solutions team has developed a tested and validated solution. The design for this solution is based on the best practices recommended by Microsoft and Dell EMC's own experience supporting hundreds of customers worldwide. The solution balances performance and cost, delivers faster time to value, and provides exceptional scalability.

The design principles and solution components for the Microsoft Exchange Solution described in this guide are suited for medium to large scale deployments. This design guide contains the following sections:

- 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 such as 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 Sample Implementations section describes sample implementations for a medium and large regional organization.
- The Verification section provides an overview of the verification process that was performed to ensure that the solution met the design requirements.

1 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 and describe the important considerations when sizing an Exchange server.

1.1 Scope

This guide describes the pre-architected solution for Microsoft Exchange on Dell EMC PowerEdge rack servers. The solution uses Dell EMC PowerEdge R740xd 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.

Microsoft Exchange Jetstress was 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.

1.2 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 R740xd servers for various user profiles. Users are expected to have sufficient understanding and knowledge of Exchange 2016.

2 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 R740xd 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 R740xd can be easily scaled up or scaled out based on the business needs.

Components for Exchange Server 2016 solution contain Dell EMC Storage products, Dell EMC networking products, Dell EMC ProDeploy services and Dell EMC ProSupport services. For network integration, Active Directory and data migration, we also have 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 Exchange Server 2016 solution.

2.1 Dell EMC PowerEdge Server

Dell EMC 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 such as 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 the 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 Dell EMC OpenManage Essentials
- **Energy efficiency** -- comprehensive optimization, 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 R740xd 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 as follows. The following simply gives the variety and range of what PowerEdge R740xd can offer. For Exchange Sever, we have selected specific configurations to

optimize performance and follow best-practices guidance given by Microsoft. Please see details of such configurations in Section 6.1 and 6.2 in this paper.

- 2 x Intel Xeon Gold 5115 CPU @ 2.40 GHz
- Up to 1536 GB of memory with 24 DIMMs
- Up to 160 TB maximum raw internal storage by using 10 TB Near-Line SAS (NL-SAS) drives
- Support for multiple internal drive configurations that include:
  - Front bays: Up to 24 x 2.5" SAS/SSD/NVMe, max 153 TB. Up to 12 x 3.5" SAS, max 120 TB
  - Mid bay: Up to 4 x 3.5" drives, max 40 TB. Up to 4 x 2.5" SAS/SSD/NVMe, max 25 TB
  - Rear bays: Up to 4 x 2.5" max 10 TB. Up to 2 x 3.5" max 20 TB
- RAID Controllers
  - Internal controllers: PERC H730p, H740p, HBA330, Software RAID (SWRAID) S140
  - Boot Optimized Storage Subsystem: HWRAID 2 x M.2 SSDs 120 GB, 240 GB
  - External PERC (RAID): H840
  - External HBAs (non-RAID): 12 Gbps SAS HBA
- 8 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 R740xd server, when configured with sixteen 3.5-inch large farm factor 10 TB NL-SAS drives, provides up to 160 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 mid bay. Figure 1 shows this configuration.

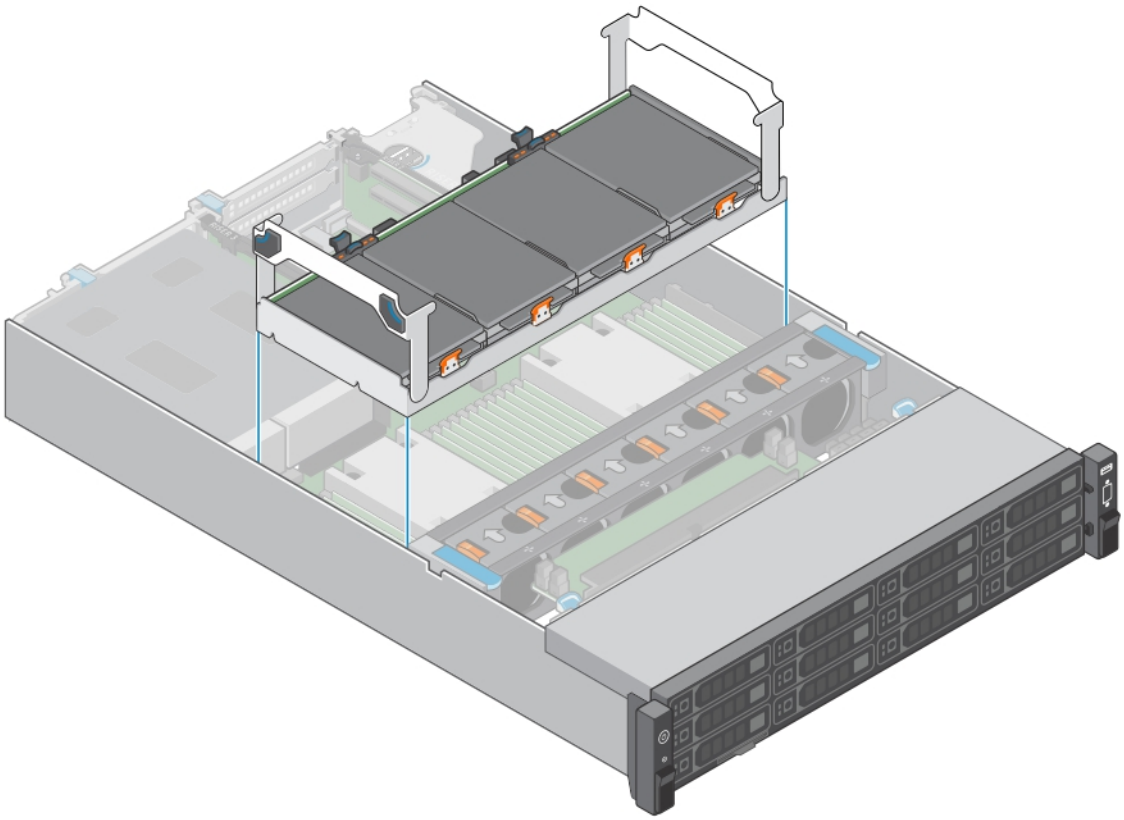

Figure 1 PowerEdge R740xd LFF chassis with mid bay

2.2 Dell EMC Storage

The Dell EMC Storage MD1400 direct-attached storage (DAS) enclosure with 12 Gb SAS throughput is specifically engineered to work with the 14th generation of PowerEdge servers that use the newest line of PowerEdge RAID Controller 10 (PERC10) 12 Gb SAS HBA cards. Scaling capacity built in the 14th generation of PowerEdge servers enables end-to-end 12 Gb solutions with exceptional storage flexibility and IO performance for applications such 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 H840 Host RAID 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 H840 Host RAID adapter. Table 1, Dell EMC Storage MD1400 features lists these features.

Table 1 Dell EMC Storage MD1400 features

Feature	Specification
Drives	Up to 12 hot-pluggable 3.5" and 2.5" drives (2.5" drives available with adapter)
Drive Performance and Capacities	3.5" NL-SAS 6 Gb HDD (7.2K): 1 TB, 2 TB, 4 TB 3.5" NL-SAS 512e 12 Gb HDD (7.2K): 6 TB, 8 TB, 10 TB 2.5” SAS 6 Gb HDD (7.2K): 500 GB 2.5” SAS 6 Gb HDD (10K): 300 GB, 600 GB, 1.2 TB, 1.8 TB 2.5” NL-SAS 12 Gb HDD (7.2K): 2 TB 2.5” SAS 12 Gb SED (15K): 600 GB 2.5” SAS 6 Gb HDD (15K): 300 GB, 600 GB 2.5” SAS 12 Gb SSD: 200 GB, 400 GB, 800 GB (WI); 200 GB, 400 GB, 800 GB, 1.6 TB (MU); 800 GB, 1.6 TB (RI)
Maximum Capacity (per enclosure)	Up to 120 TB when using 12 x 10 TB NL-SAS 3.5" HDDs
Expansion Capabilities	PERC H840 HBA enables expansion to 8 MD 1400 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 H840 (4 enclosures per port, single path), unified mode (recommended redundant path) for daisy chaining up to 4 enclosures per PERC H840 (4 enclosures connected to both ports via redundant path cabling)
Host Connectivity Split-Mode/Dual-Host Access	Split mode with dual Enclosure Management Modules (EMM) providing direct connectivity to drives 0 through 5 and a separate connectivity to drives 6 through 11
Enclosure Management Modules (EMM)	Two EMMs 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)

2.3 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 offers several 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.

2.4 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

3 Application architecture

The following sections describe the design principles of Microsoft Exchange 2016 solution and explain how the infrastructure sub-systems are architected in this solution. To demonstrate these principles, the Dell EMC Solutions Engineering Group has designed two validated email 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 20 GB. 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 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 Sample implementations in section 7.

Figure 2 shows the external data center architecture.

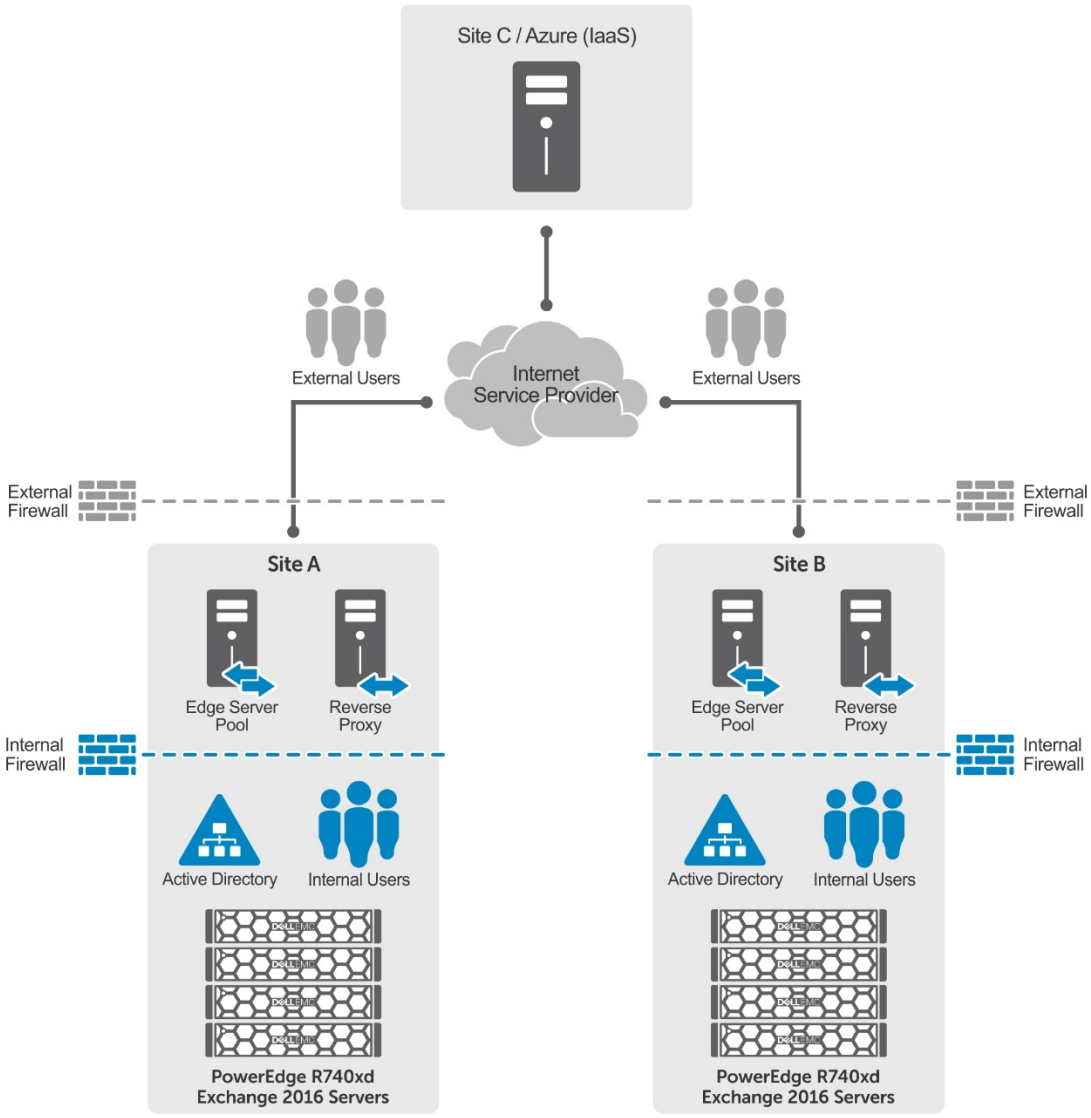

Figure 2 External data center architecture for Exchange deployment

3.1 Storage architecture

The PowerEdge R740xd server supports multiple internal drives. A RAID controller can be used to create independent RAID 0 volumes by using the internal drives. These RAID 0 volumes host Exchange databases and transaction logs.

When a PowerEdge R740xd server is configured with 16 LFF drives, 14 independent RAID0 disks can be used to store copies of Exchange databases. These RAID 0 disks can be created by using 12 disks in the front bay and 2 disks in the mid bay. 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 AutoReseed volume. Two of the rear-accessible

2.5-inch drives can be configured as a RAID 1 disk to deploy the operating system. The remaining two 2.5-inch drives can be used to store the Exchange queue database in a RAID 1 container. Figure 3 illustrates this configuration.

When configuring RAID 0 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.

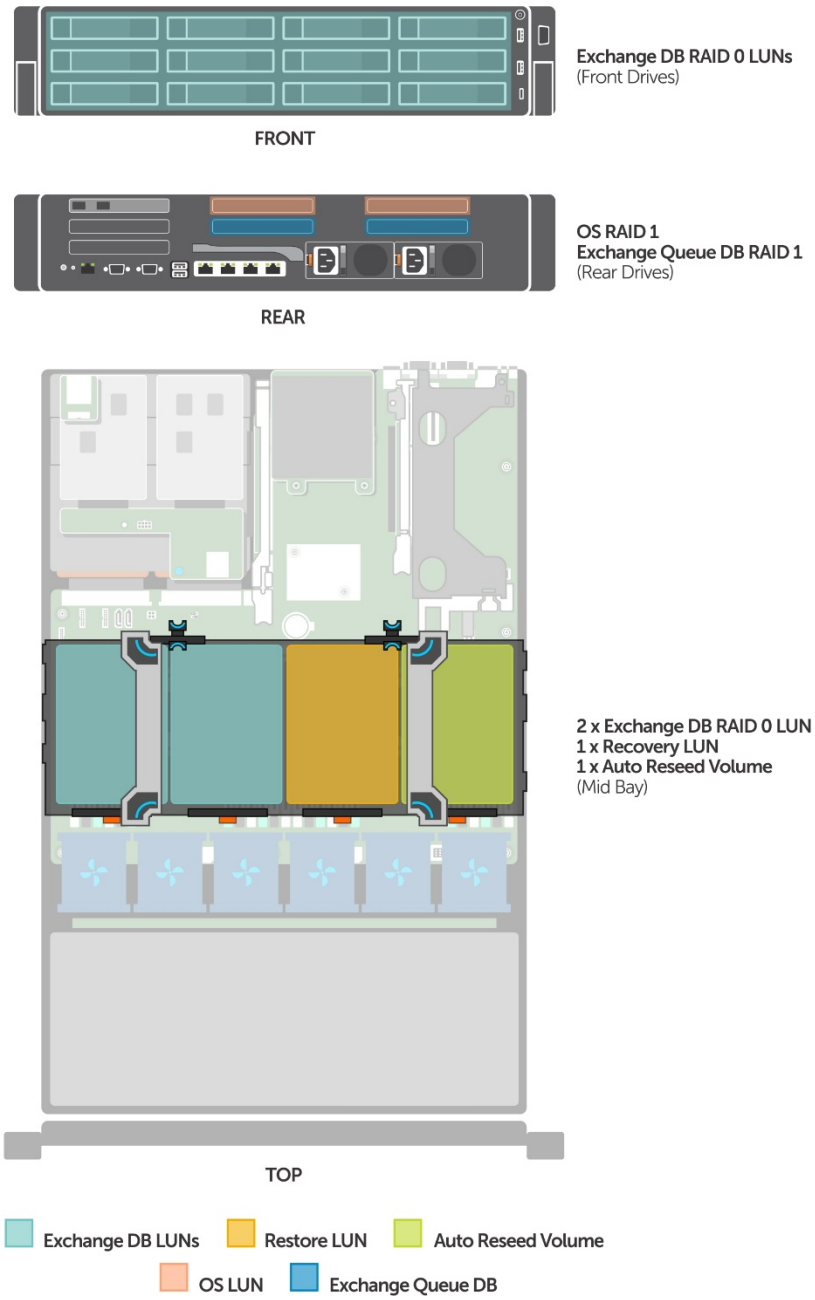

Figure 3 RAID LUN layout for the Exchange database

3.2 Network architecture

Microsoft Exchange Server 2016 network requirements are simple and easy to deploy. A single 10 GbE network can be used 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, NIC teaming is leveraged in this solution to provide multi-tier HA across the network. In addition, two 1 GbE NICs are 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.

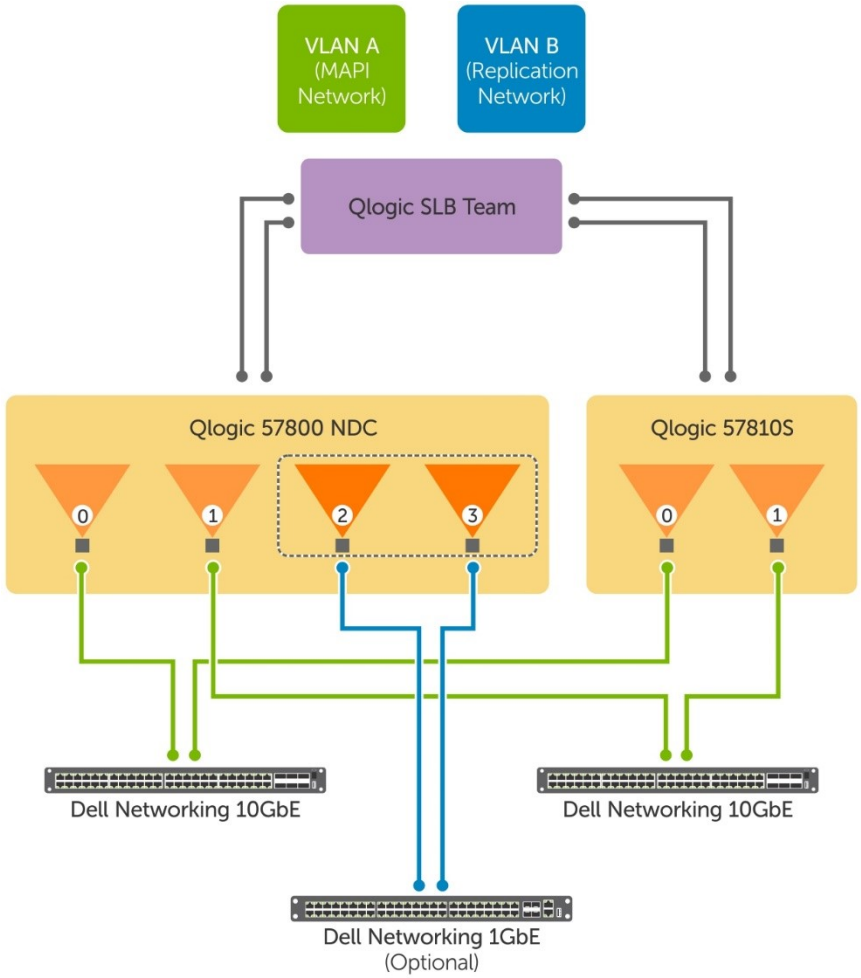

Figure 4 Network architecture for the Exchange deployment

4 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.

4.1 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 this solution, so that it can withstand server failures as well as site failures. This design is based on the building block architecture where four servers are stacked across two data centers—two in Site A and two in Site B. For more details about the design for this solution, see Comprehensive solution design.

4.1.1 Application-level High Availability

In Exchange Server 2016, DAGs are utilized to provide high availability (HA) to the mailbox databases by storing multiple passive copies of each active database. Exchange DAG also 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.

4.1.2 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. RAID1 disks provide redundancy for the operating system, the Exchange binaries, and the queue database. 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.

4.2 Application performance

In addition to HA, application performance is also critical to guarantee reliable end-user experience. Microsoft Exchange Server Jetstress 2013 Tool is used 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 is the official Microsoft tool to test both Microsoft Exchange Server 2013 and Microsoft Exchange Server 2016 storage subsystem.

4.3 Best practices

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

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

- 2U, dual socket servers (Maximum processor core count: 24)
- Maximum memory: 192 GB
- a battery-backed write cache controller
- 12 or more large form factor drive bays within the server chassis

When sizing the Exchange Storage subsystem, ensure that there are no I/O bottlenecks from an IOPS and disk latency perspective. The disk subsystem must 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, set the size of elements within a RAID stripe to 512 K for best performance.
- Ensure that each server has a single RAID1 disk pair for the Operating System, Exchange binaries, and protocol/client logs. Another RAID1 disk pair can be used to store transport database. You can configure the rest of the storage as independent RAID0 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.

5 Comprehensive solution design

A high-level diagram of the solution architecture is provided 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.

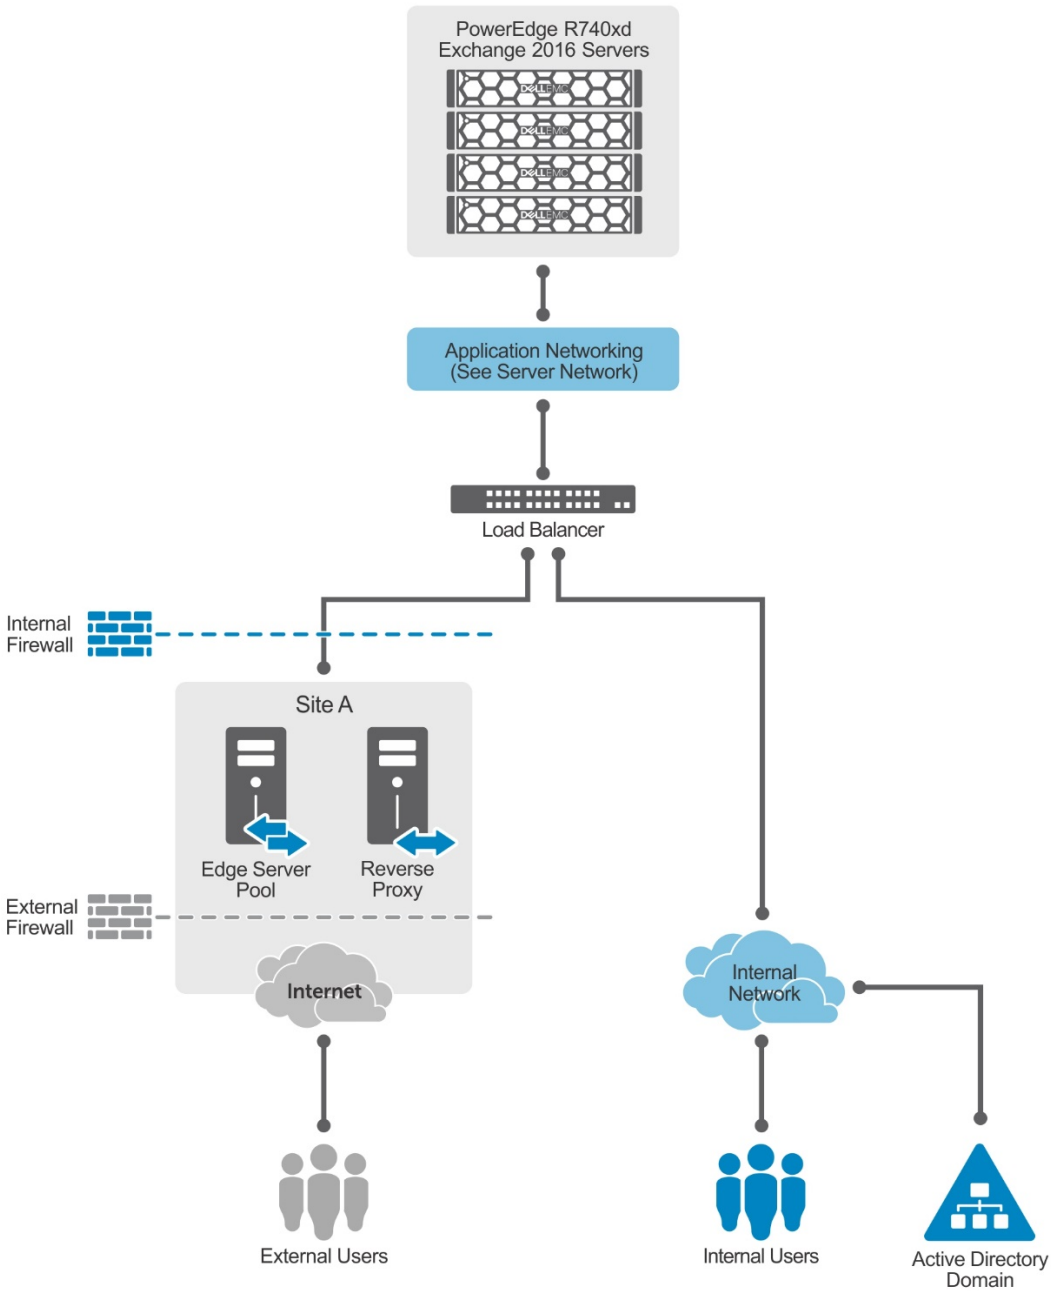

Figure 5 Site A architecture for the Exchange deployment

As shown in Figure 5, the solution design consists of multiple infrastructure components, such as server, storage, network, load balancer, and firewall reverse proxy. Every data center that has Exchange servers deployed will have the same set of infrastructure components. The PowerEdge R740xd 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.

5.1 Core components

To explain the building block architecture, this guide uses the concept of a PoD. A PoD is a standardized configuration of the minimum server and storage resources sized to meet the solution requirements of a given mailbox profile. The configuration for each member within a PoD is 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. To host mailbox databases and logs, you can use internal disks in the R740xd servers or external direct-attached storage (DAS)-- Dell EMC MD1400. You can increase the number of mailboxes by scaling out the building blocks.

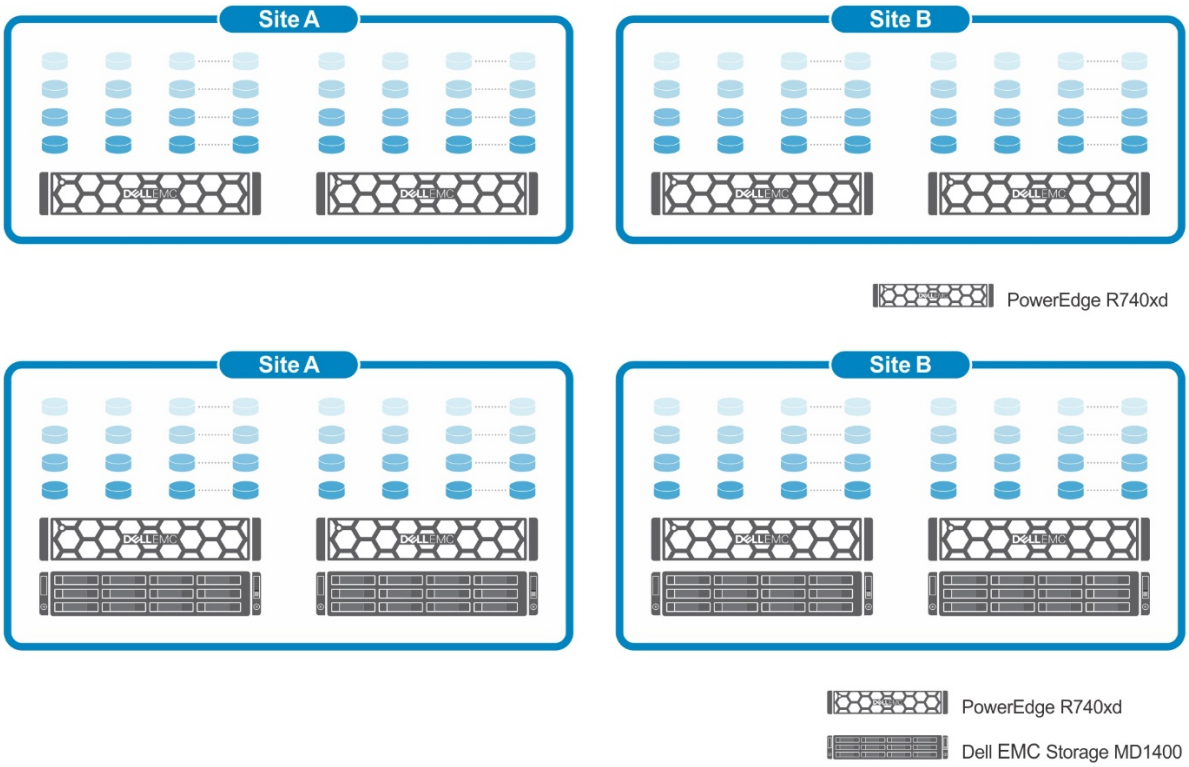

Figure 6 PoD architecture for the Exchange deployment

Figure 6 illustrates the building block architecture of Exchange Server 2016 solution. 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 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. This solution employs Dell EMC Networking 10 GbE Top-of-Rack (ToR) switches for network connectivity. The network architecture used in this solution is illustrated in Figure 7.

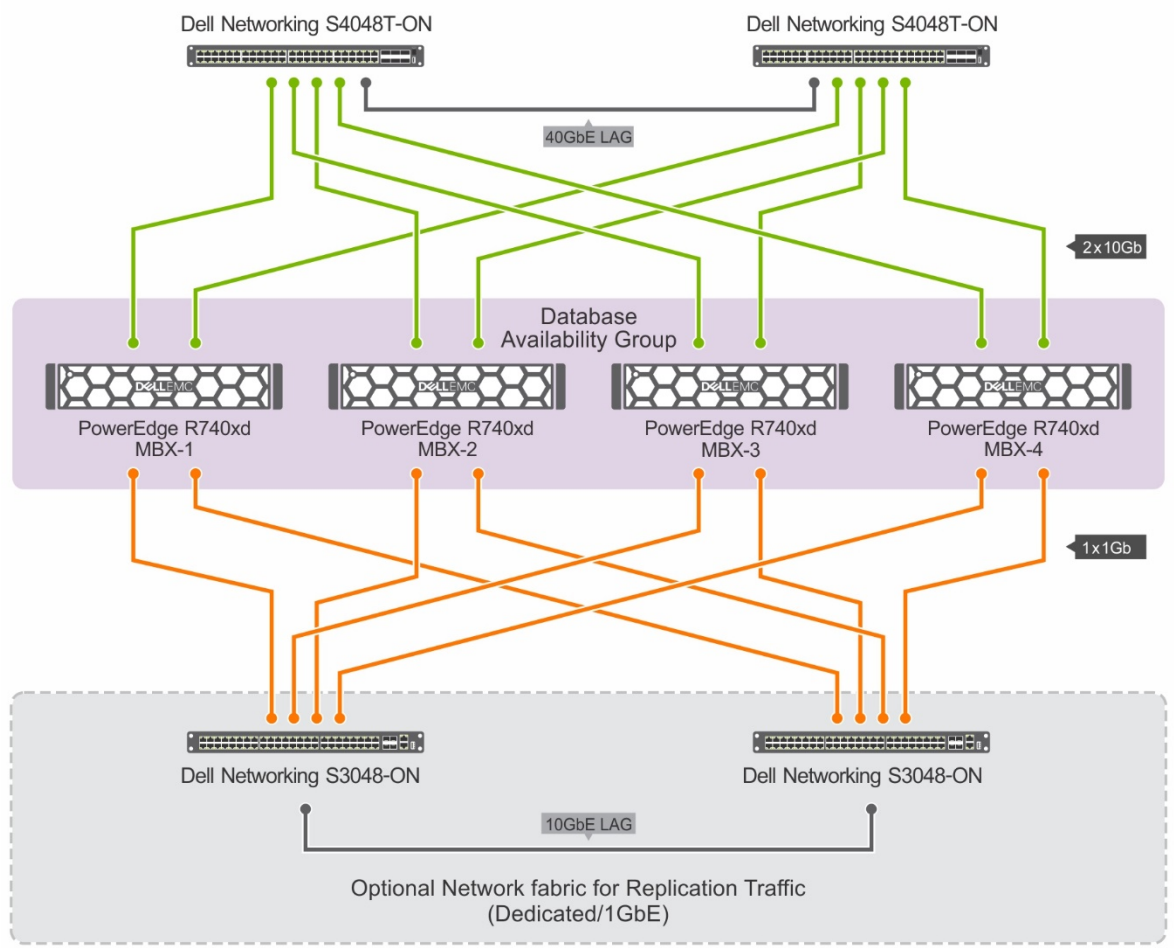

Figure 7 Detailed network architecture for the Exchange reference implementation

As mentioned in Section 2.1, a total of four 10 GbE 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 requires two additional switches to support the server connectivity.

5.2 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 in Exchange 2016. However, some organizations may require an Edge 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.

5.3 Scalability

Using PoDs as building blocks provides better scalability. As organizations grow, you can accommodate additional mailboxes by adding more PoDs. 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.

6 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 important sizing considerations are server and storage.

The following are the server sizing considerations:

- Determining the type of processor that is best suited for handling the Exchange Mailbox profile requirements
- Deciding the required size of memory 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 which can be used with both Exchange Server 2013 and Exchange Server 2016.

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 that uses the PoD concept to build and scale out an Exchange infrastructure. The subsequent sections provide an overview and benefits of using PowerEdge R740xd and the building block architecture for Exchange 2016.

6.1 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 R740xd to expand the storage capacity. To achieve 10,000 mailboxes of 35 GB each, four PoDs are used, where each PoD has four PowerEdge R740xd servers and four Dell EMC 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	**Dell EMC PowerEdge R740xd server with 3.5” HDD chassis**
CPU	2 × Intel Xeon Gold 5115 processor @ 2.40 GHz with 10-cores
Memory	Up to 96 GB DDR4
NIC	Qlogic Network adapters
RAID Controllers	1. PERC H740P Adapter Firmware version: 50.0.1-0537 Storport Driver Version 10.0.14393.351 Driver version 7.700.51.00 2. PERC H840 Adapter Firmware version: 50.0.1-0537 Storport Driver Version 10.0.14393.351 Driver version 7.700.51.00
Internal Disks	4 x 1.2 TB SAS 2.5-inch 10K RPM disk drives 2 x 1.2 TB SAS 2.5-inch 10K RPM disk drives in RAID 1 volume (Operating System and Application) 2 x 1.2 TB SAS 2.5-inch 10K RPM disk drives in RAID 1 volume (Exchange queue database)

Table 3 Storage Subsystem Configuration (PowerEdge R740xd server storage)

Storage Subsystem	**Dell EMC PowerEdge R740xd internal 3.5-inch drives**
Disks	16 x 8 TB 7.2 K RPM NL-SAS 3.5-inch disk: 14 x 8 TB 7.2 K RPM NL-SAS 3.5-inch drive in 14 x RAID 0 volumes (for DB and Log) 1 x 8 TB 7.2 K RPM NL-SAS 3.5-inch drive (for Restore LUN) 1 x 8 TB 7.2 K RPM NL-SAS 3.5-inch drive (Auto Reseed Volume)
RAID Controller	Dell EMC PowerEdge RAID Controller H740P (Firmware version: 50.0.1-0537)

Table 4 Storage Subsystem Configuration (Dell EMC Storage MD1400)

Storage Subsystem	**Dell EMC Storage MD1400 3.5-inch drive**
Disks	12 x 8 TB 7.2 K RPM NL-SAS 3.5-inch disk: 10 x 8 TB 7.2 K RPM NL-SAS 3.5-inch drive in 14 x RAID 0 volumes (for DB and Log) 1 x 8 TB 7.2 K RPM NL-SAS 3.5-inch drive (for Restore LUN) 1 x 8 TB 7.2 K RPM NL-SAS 3.5-inch drive (Auto Reseed Volume)
RAID Controller	Dell EMC PowerEdge RAID Controller H840 (Firmware version: 50.0.1-0537)

6.2 Solution for higher number of mailboxes

Organizations that need a higher number of small to medium sized 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	**Dell EMC PowerEdge R740xd server with 3.5” HDD chassis**
CPU	2 × Intel Xeon Gold 5115 processor @ 2.40 GHz with 10-cores
Memory	Up to 96 GB DDR4
NIC	Qlogic Network adapters
RAID Controllers	PERC H740P Adapter Firmware version: 50.0.1-0537 Storport Driver Version 10.0.14393.351 Driver version 7.700.51.00
Internal Disks	4 x 1.2 TB SAS 2.5-inch 10 K RPM disk drives 2 x 1.2 TB SAS 2.5-inch 10 K RPM disk drives in RAID 1 volume (Operating System and Application) 2 x 1.2 TB SAS 2.5-inch 10 K RPM disk drives in RAID 1 volume (Exchange queue database)

Table 6 Storage Subsystem Configuration (PowerEdge R740xd server storage)

Storage Subsystem	**Dell EMC PowerEdge R740xd internal 3.5-inch drives**
Disks	16 x 8 TB 7.2 K RPM NL-SAS 3.5-inch disk: 14 x 8 TB 7.2 K RPM NL-SAS 3.5-inch drive in 14 x RAID 0 volumes (for DB and Log) 1 x 8 TB 7.2 K RPM NL-SAS 3.5-inch drive (for Restore LUN) 1 x 8 TB 7.2 K RPM NL-SAS 3.5-inch drive (Auto Reseed Volume)
RAID Controller	Dell EMC PowerEdge RAID Controller H740P (Firmware version: 50.0.1-0537)

7 Sample implementations

This section outlines two sample implementations for a medium and large organization of 5,000 and 10,000 mailboxes respectively. These two customer profiles are selected for sample implementations 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.

7.1 Email solution for a medium sized organization of 5,000 employees

Leveraging the PoD architecture as explained in section 5.1, an active-active user distribution model (4 copy DAG) email solution for 5,000 users with 20 GB mailbox size each can be achieved by stacking two PoDs as shown in Figure 8.

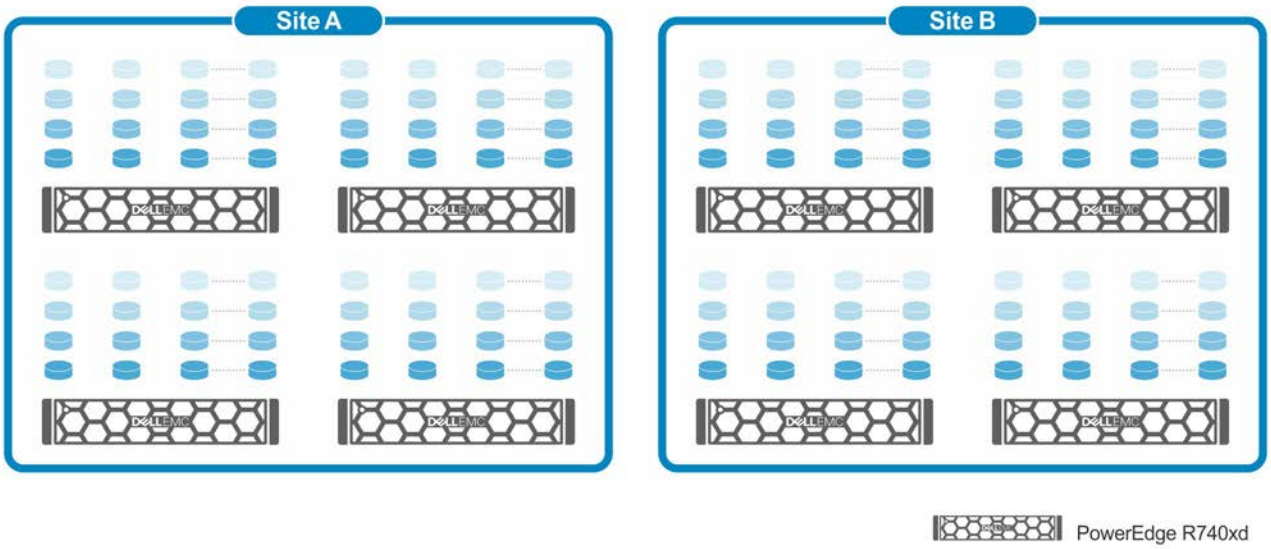

Figure 8 PoD architecture to host 5,000 mailboxes

This site resilient design that uses eight PowerEdge R740xd servers has been tested and validated in the lab by the Global Solution 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	21560 MB
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	Dell EMC PowerEdge R740xd server with 3.5” HDD chassis
CPU	2 × Intel Xeon Gold 5115 processor @ 2.40 GHz with 10-cores
Memory	Up to 96 GB DDR4
NIC	Qlogic Network adapters NetXtreme II Broadcom NetXtreme Gigabit Ethernet
RAID Controllers	PERC H740P Adapter Firmware version: 50.0.1-0537 Storport Driver Version 10.0.14393.351 Driver version 7.700.51.00
Internal Disks	4 x 1.2 TB SAS 2.5-inch 10 K RPM disk drives 2 x 1.2 TB SAS 2.5-inch 10 K RPM disk drives in RAID 1 volume (Operating System and Application) 2 x 1.2 TB SAS 2.5-inch 10 K RPM disk drives in RAID 1 volume (Exchange queue database)

Table 12 Storage subsystem configuration (PowerEdge R740xd server storage)

Storage subsystem	Dell EMC PowerEdge R740xd internal 3.5-inch drives
Disks	16 x 8 TB 7.2 K RPM NL-SAS 3.5-inch disk: 14 x 8 TB 7.2 K RPM NL-SAS 3.5-inch drive in 14 x RAID 0 volumes (for DB and Log) 1 x 8 TB 7.2 K RPM NL-SAS 3.5-inch drive (for Restore LUN) 1 x 8 TB 7.2 K RPM NL-SAS 3.5-inch drive (Auto Reseed Volume)
RAID Controller	Dell EMC PowerEdge RAID Controller H740P (Firmware version: 50.0.1-0537)

7.2 Email solution for a large regional organization of 10,000 employees

This sample implementation considers 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.

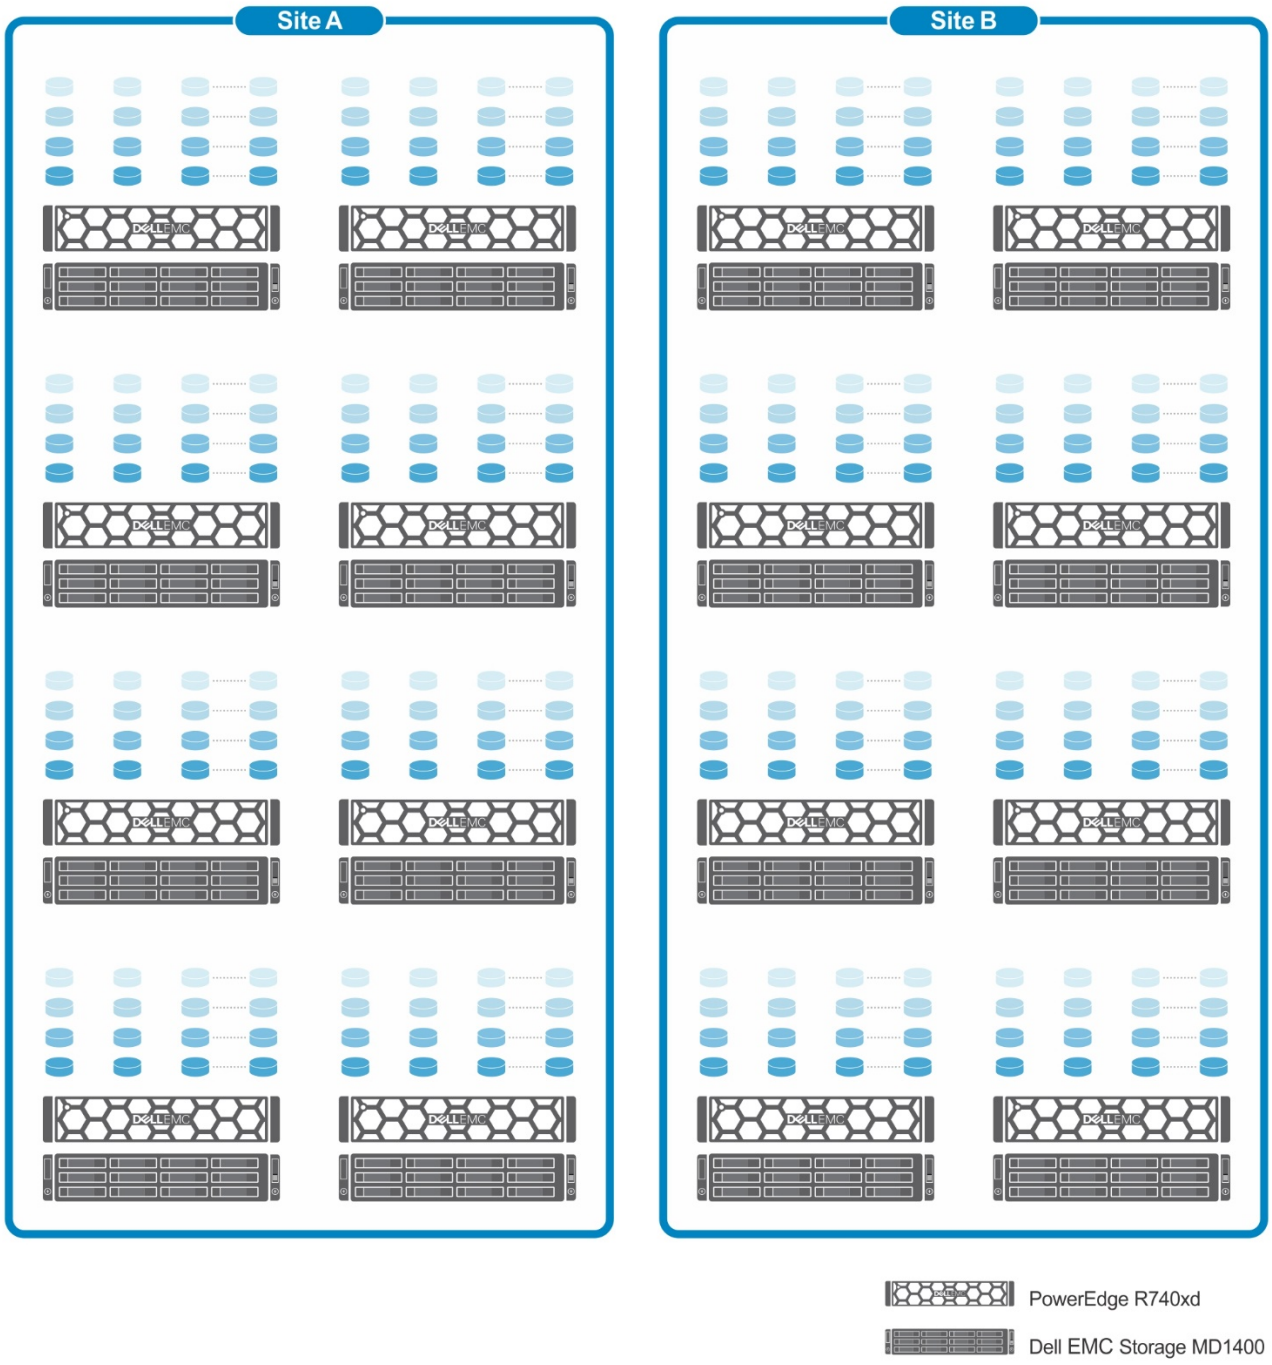

Figure 9 PoD architecture to host 10,000 mailboxes

This site resilient tested solution employs 16 PowerEdge R740xd 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)	5,000	5,000	10,000	10,000
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	37566 MB
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	1,250	384	10,000

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 (First Server Failure after SDC Activation)	56	1,458	384	10,000
Number of Active Databases (Second Server Failure after SDC Activation)	--	--	--	--

Table 17 Exchange Server configuration

Microsoft Exchange Server System	Dell EMC PowerEdge R740xd server with 3.5" HDD chassis
CPU	2 × Intel Xeon Gold 5115 processor @ 2.40 GHz with 10-cores
Memory	Up to 96 GB DDR4
NIC	Qlogic Network adapters
RAID Controllers	PERC H740P Adapter Firmware version: 50.0.1-0537 Storport Driver Version 10.0.14393.351 Driver version 7.700.51.00 PERC H840 Adapter Firmware version: 50.0.1-0537 Storport Driver Version 10.0.14393.351 Driver version 7.700.51.00
Internal Disks	4 x 1.2 TB SAS 2.5-inch 10 K RPM disk drives 2 x 1.2 TB SAS 2.5-inch 10 K RPM disk drives in RAID 1 volume (Operating System and Application) 2 x 1.2 TB SAS 2.5-inch 10 K RPM disk drives in RAID 1 volume (Exchange queue database)

Table 18 Storage Subsystem Configuration (PowerEdge R740xd server storage)

Storage subsystem	Dell EMC PowerEdge R740xd internal 3.5-inch drives
Disks	16 x 8 TB 7.2 K RPM NL-SAS 3.5-inch disk: 14 x 8 TB 7.2 K RPM NL-SAS 3.5-inch drive in 14 x RAID 0 volumes (for DB and Log) 1 x 8 TB 7.2 K RPM NL-SAS 3.5-inch drive (for Restore LUN) 1 x 8 TB 7.2 K RPM NL-SAS 3.5-inch drive (Auto Reseed Volume)

Storage subsystem	Dell EMC PowerEdge R740xd internal 3.5-inch drives
RAID Controller	Dell EMC PowerEdge RAID Controller H740P (Firmware version: 50.0.1-0537)

Table 19 Storage Subsystem Configuration (Dell EMC Storage MD1400)

Storage subsystem	Dell EMC Storage MD1400 3.5-inch drives
Disks	12 x 8 TB 7.2 K RPM NL-SAS 3.5-inch disk: 10 x 8 TB 7.2 K RPM NL-SAS 3.5-inch drive in 14 x RAID 0 volumes (for DB and Log) 1 x 8 TB 7.2 K RPM NL-SAS 3.5-inch drive (for Restore LUN) 1 x 8 TB 7.2 K RPM NL-SAS 3.5-inch drive (Auto Reseed Volume)
RAID Controller	Dell EMC PowerEdge RAID Controller H840 (Firmware version: 50.0.1-0537)

8 Verification

This section provides proof points for performance and resource utilization by using the sample implementation of 10,000 users as an example.

8.1 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 of 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 R740xd 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 R740xd server.

A PowerEdge R740xd 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 IOPS 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 Exchange 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	797	16 x LFF 8 TB NL-SAS drives in PowerEdge R740xd	RAID 0	32	8 TB	150 message per day (0.121 tested)

9 Conclusion

Microsoft Exchange solution with PowerEdge R740xd 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 R740xd server used in this solution for sample implementation offers an excellent balance between performance and management. It is a highly efficient and cost-effective hardware building block for any medium-sized 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 this solution offers a highly scalable architecture for Exchange deployments. The proposed PoD design requires minimum server and storage. Dell EMC PowerEdge R740xd server along with 16 LFF drives provides sufficient compute and storage capacity for most of the medium-sized 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 business sizes or needs.

10 Additional Resources

Dell EMC PowerEdge R740xd

PowerEdge RAID Controller H740P

PowerEdge RAID Controller H840

Dell EMC Storage MD1400 Direct Attached Storage

Exchange Solution Reviewed Program (ESRP) - Storage