Reference Architecture – Microsoft Exchange 2013 on Dell PowerEdge R730xd – 10000 Mailboxes

A Dell Reference Architecture

Dell Engineering
August 2015

Revisions

Date	Description
August 2015	Initial release

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.
Copyright ©2015 Dell Inc. All rights reserved. Dell and the Dell logo are trademarks of Dell Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

Table of contents

Executive summary

Messaging systems provide rich collaboration features, including email, calendar scheduling, voice mailbox, and collaboration from anywhere with almost any device. IT administrators face an increasingly complex task of designing, administrating and maintaining these messaging systems. With increasing storage requirements, businesses also face the challenge of implementing a suitable messaging system that meets their total cost of ownership (TCO), return on investment (ROI) and business continuity needs.

To address these challenges, the enhanced features in Microsoft Exchange Server 2013 include support for low-cost storage and a highly available disaster recovery (DR) model. This paper provides information about how these features can be used efficiently and cost-effectively with the Dell PowerEdge R730xd server. The PowerEdge R730xd is a 2U rack server that offers a balance between the dense internal storage and compute capacity to provide the performance required for enterprise applications such as Exchange. With the introduction of the internal drive tray, the PowerEdge R730xd server supports up to sixteen 3.5-inch NL-SAS drives making it a highly suitable building block for deploying Exchange 2013 for any mid-size or large businesses.

Building a scalable and highly available messaging infrastructure requires the following:

- Understanding the application architecture and its impact on the infrastructure design
- Implementing application and infrastructure design principles
- Verifying that the deployed infrastructure delivers the desired performance and capacity

This paper describes how to implement and validate a 10,000 mailbox Microsoft Exchange messaging solution in a structured way.

Section 2 provides an introduction to the Microsoft Exchange Server 2013 and the server roles.

Section 3 explains the sizing considerations and the solution components.

Section 4 explains the building block architecture for Exchange deployments and describes how the building blocks can be scaled out to support growth.

Section 5 describes the design considerations and recommendations for developing a highly available infrastructure for Exchange.

Section 6 provides a reference implementation for 10,000 mailboxes that leverages the building block architecture described in section 4 and implements the design considerations detailed in section 5.

Section 7 provides the proof points from the Exchange JetStress and Loadgen verification of the reference implementation described in section 6.

Section 8 summarizes the solution specifications for the reference implementation described in section 4.

Section 9 summarizes the benefits of building block architecture and describes how PowerEdge R730xd is the suitable building block for Exchange 2013 deployments.

1 Introduction

Exchange Server 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 provides reliable, scalable enterprise-class email with compliance and e-discovery features integrated with Microsoft SharePoint and Lync. Exchange Server supports people and organizations as their work habits evolve from a communication focus to a collaboration focus. The database design in Exchange 2013 reduces the storage I/O requirements, thereby optimizing Exchange for cost-effective, low-speed storage. The Exchange deployment must be appropriately sized not only for specific message profile requirements but also for growth and high availability requirements. The following sections provide a brief overview of Exchange 2013, describe the important considerations when sizing an Exchange deployment and provide the building block architecture that Dell developed for Exchange infrastructure.

1.1 Scope

This paper describes the design of Exchange 2013 infrastructure on PowerEdge R730xd servers. This paper proposes the PowerEdge R730xd server as a building block for Exchange and provides a sample reference implementation of the proposed block architecture for 10,000 mailboxes. The design is verified for the expected performance using both Microsoft Exchange JetStress and Loadgen tools. The paper also provides a brief overview of different user profile options.

1.2 Audience

This paper is intended for IT managers, messaging administrators and consultants interested in designing and deploying a cost-effective Exchange solution on PowerEdge R730xd servers for various user profiles. While this paper provides an overview of the PowerEdge R730xd server, the reader is expected to have sufficient understanding of Exchange 2013.

2 Microsoft Exchange 2013

In Exchange Server 2013, the five server roles from Exchange 2010 have been replaced by two main types of server roles: Mailbox server and Client Access server. The features provided by the Hub Transport server role in Exchange 2010 are handled by the Transport service—the Mailbox Transport service on Mailbox servers and the Front End Transport service on Client Access servers. The features provided by the Unified Messaging (UM) server role in Exchange 2010 are handled by the UM service. The UM service has components that run on both the Client Access server and Mailbox server. The Edge Transport server role is available with Exchange 2013 SP1. Existing Exchange 2007 or Exchange 2010 Edge can work with Exchange 2013.

2.1 Mailbox Server Role

The Mailbox server role in Exchange 2013 hosts mailboxes, mailbox databases and public folders. It also includes the Transport service and UM components. The Mailbox server handles all activity for the active mailbox databases. To achieve application high availability (HA), multiple Mailbox server roles can be clustered using the Database Availability Group (DAG) functionality.

2.2 Client Access Server Role

This server role supports messaging clients, such as Outlook, mobile cellular devices and Exchange Web Services clients. The Client Access server in Exchange Server 2013 is a thin, stateless server that functions as the front end. The Client Access server accepts all client requests, authenticates and routes/proxies the request to the Mailbox server. The mailbox   has the currently active copy of the database hosting the user's mailbox. The Client Access server provides client authentication and manages client connections through redirection and proxy functionality. In addition to all the usual client access protocols (HTTP, POP and IMAP, and SMTP), it provides network security functionalities, such as Secure Sockets Layer (SSL).

> Microsoft recommends consolidating multiple server roles as application-level entities on a single physical or virtual server.

This paper uses an Exchange multi-role server configuration in the example implementation. An Exchange multi-role server configuration consists of multiple server roles; for example, Mailbox and Client Access server roles that are collocated on a single server. This simplifies the planning and deployment and increases service availability by having a 1:1 ratio for Client Access to Mailbox. However, if Client Access and mailbox roles are collocated, a hardware load balancer must be deployed to load balance between the Client Access servers because Windows Network Load Balancing cannot be used with DAG member servers.

3 Sizing an Exchange Deployment

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 has 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 of mailboxes and size of the mailboxes and mailbox profiles, you can use the Exchange 2013 role requirements calculator. Generic sizing guidance for Exchange Server roles is outside the scope of this paper.

While the sizing activity provides the necessary capacity information for both server and storage hardware, the infrastructure and application architecture must be designed for scalability and HA. Exchange Server 2013 infrastructure can be designed in different ways based on the size of the configuration and the number of HA and DR copies required. 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 has developed a block architecture mode 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 2013.

3.1 Dell PowerEdge R730xd Server Overview

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 is a hardware building block for any mid-size or large business that provides scalability in memory density and storage capacity and IOPS performance in a dense 2U form-factor. This enables larger and more efficient databases and mail servers. The internal RAID controller provides a range of RAID levels for improved storage reliability. Major features include:

- 2 x Intel Xeon E5-2600 v3 product family
- Up to 768 GB of memory with 24 DIMMs

- Up to 96 TB maximum raw internal storage using 6 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 using PERC H330, PERC H730, PERC H730P, and External JBOD RAID support with PERC H830
- Six PCIe 3.0 expansion slots
- Dell OpenManage Essentials, Dell Management Console, Dell OpenManage Power Center, and Dell OpenManage Connections.

The PowerEdge R730xd server, when configured with sixteen 3.5-inch large farm factor 6 TB NL-SAS drives, provides up to 96 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.

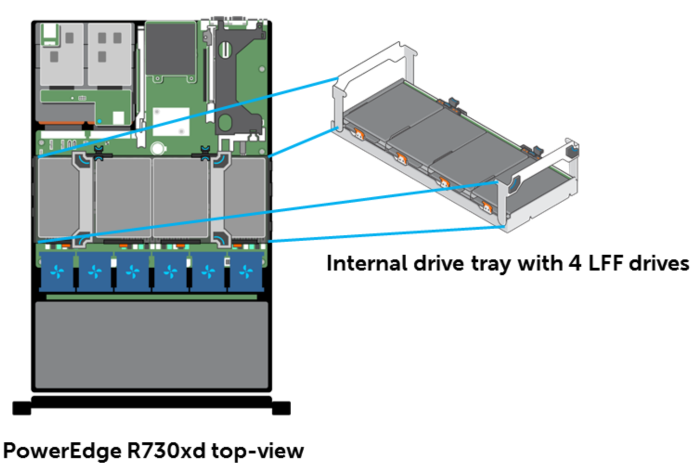

Figure 1 PowerEdge R730xd LFF chassis with an internal drive tray

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.

However, the solution presented here makes use of only twelve drives in the front bay. The configuration illustrated in Figure 3 provides the necessary storage and compute capacity for an Exchange deployment. The storage controller used in the server platform plays a critical role in delivering the performance needs of a storage-intensive application such as Exchange. The next section briefly describes Dell PERC H730P mini and its features.

3.1.1 Dell PERC H730P Mini

The PERC H730P Mini controller used in the PowerEdge R730xd server hosting the Exchange Server is an internal host-based RAID Controller card from the PERC Series 8 Family. PERC cards are built on the LSISA2208 dual-core PowerPC RAID-on-Chip (ROC) and offer unmatched I/O performance for database, applications, and streaming digital media environments.

Table 1 shows the technical specifications of the PERC H730P Mini controller. For more information, see Dell PowerEdge RAID Controller product page.

Table 1 PowerEdge RAID controller H730P Mini

Feature	**Specification**
Solution	Eight-port internal SATA+SAS solution supporting up to 32 HDDs and SSDs
Physical dimension	167.6mm (6.6 in) x 64.4mm (2.5 in) (MD2 low profile)
Connectors	Two internal mini-SAS SFF8088
Device support	Up to 32 (SAS, SATA)
Host bus type	8-lane, PCI Express 2.0 compliant
Data transfer rate	Up to 6 Gbps per port
SAS controller	LSISAS2208 dual-core PowerPC ROC
Cache size	2 GB
RAID management	Dell OpenManage Storage Services and Additional management through UEFI (HII) and CEM
Optional SSD optimization	Dell FastPath software: Delivers high IOPs performance on SSD arrays

4 Building Block Architecture

To represent the building block architecture, this paper uses the concept of a Pod, which is a standardized configuration of the minimum server and storage resources sized to meet the solution requirements, including the mailbox profile. This design requires that the configuration within each Pod in a deployed solution remains the same. To meet more mailbox requirements (of the same mailbox profile as earlier mailboxes), you need to increase the number of Pods.

Figure 2 shows the logical representation of a Pod. Each physical server has identical configurations with respect to memory, CPU, and storage components. This model can use internal disks on a dense server or an external direct-attached storage (DAS) to host mailbox databases and logs. You can increase the number of mailboxes by scaling out the building block configuration and expanding the Exchange deployment per requirements.

This paper describes a reference implementation of a building block architecture for 10,000 mailboxes. It has a single DAG spanning across two sites with the local site having two copies of Exchange mailbox databases and the remote site having a single passive copy of the Exchange databases.

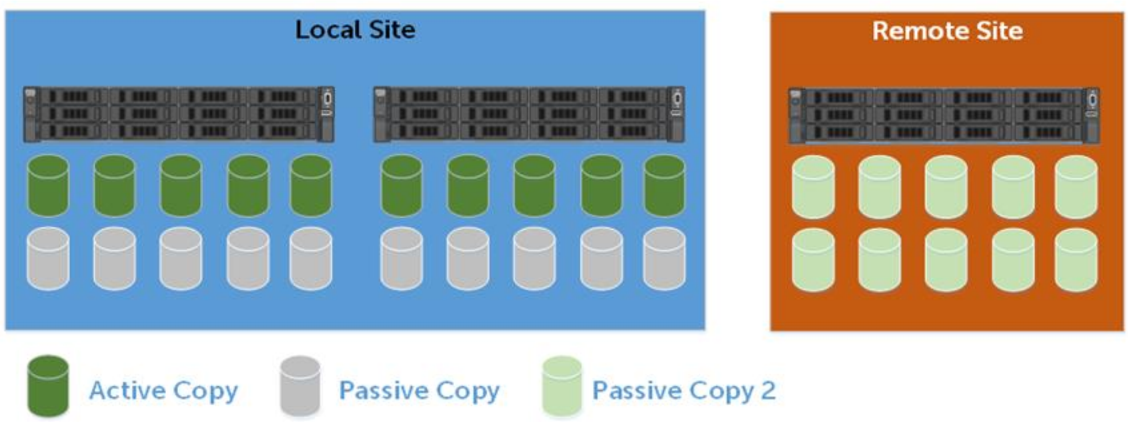

Figure 2 Pod architecture for Exchange deployment

For a site-resilient configuration, the Pod design recommends a 3-copy DAG deployment with one passive copy at the remote site. In case of a complete local site failure, the database copies on the remote site get activated and provide services to the end users.

4.1 Scaling out for growth

The Exchange Server 2013 architecture is selected depending on the user profiles, number of users and HA/DR copies required. For the three-copy DAG scenario with two active copies in the local site and one passive copy in the remote site, the building block architecture can be effectively scaled out by multiplying the number of Pods. Figure 2 shows the representation of a Pod that can be scaled out to meet the requirement.

The sizing and configuration of such a Pod is discussed in the Solution Architecture section 6 provides the specifications of the PowerEdge R730xd server used as a building block.

As shown in Figure 2, a single Pod can serve 5,000 users with 3 GB mailbox size and 150 messages per day profile. Two such configurations for 10,000 users can be stacked, provided that the message profile and mailbox size remain constant. Similarly, for 15,000 users, three such Pods can be stacked. As shown in Figure 3, 20,000 users can be accommodated by stacking four such Pods. The number of Pods that can be scaled to in a single DAG will depend on solution requirements. However, as a DAG can include a maximum of 16 servers, the number of Pods that can be scaled in a DAG is limited to five. It is recommended scaling to a maximum of four Pods (12 servers) in a three-copy DAG. This enables the fourth copy to be added in the remote site while retaining the modular design. Section 6 provides more details on this.

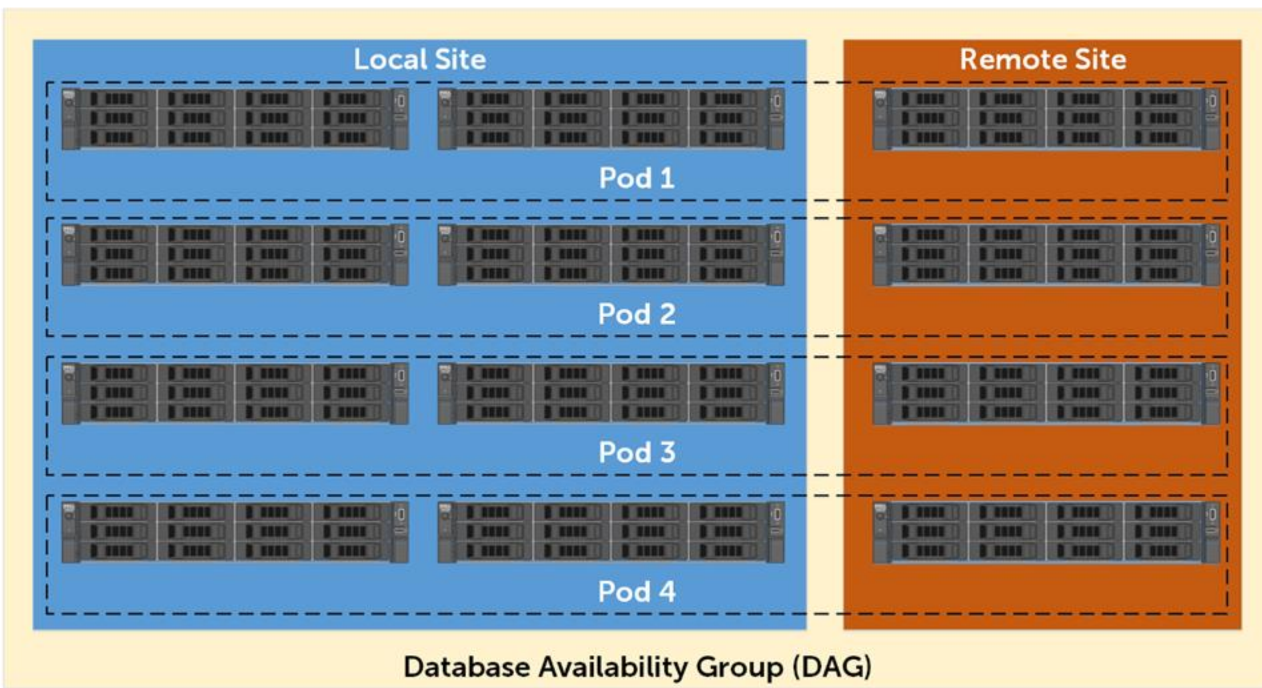

Figure 3 Scaling out the Pod architecture for growth

The preceding example assumes that servers provide adequate internal storage and processing power. As the number of servers increase, storage must be multiplied by the same number. For example, to support 5,000 users per Pod for a given mailbox profile, the server platform requires a larger amount of storage and compute capacity.

While the PowerEdge R730xd server with 12 LFF drives provides the necessary storage and compute capacity for an Exchange deployment, it is important that you design the Exchange infrastructure to support scalability, high availability and efficient management. This requires additional solution components depending on the design considerations. The following section describes these design considerations.

5 Solution Design

While designing the solution infrastructure, you must consider the best practices at each tier of the architecture. The following sections describe the design principles and how each infrastructure sub-system can be architected. These design principles, along with the component architecture, provide guidance for any Exchange deployment that uses the PowerEdge R730xd server as described in section 3.1.

5.1 Infrastructure High Availability

In a solution infrastructure, resources such as server, storage drives, network path and switches should be highly available. Infrastructure availability requires that the solution architecture builds redundancy for each component within the infrastructure. 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.

5.1.1 Storage Architecture

The PowerEdge R730xd server supports multiple internal drives. A RAID controller can be used to create RAID disks by using the internal drives. The disk mirroring (RAID 1) configuration protects the application data from disk failures. When a PowerEdge R730xd server is configured with 12 LFF drives, five RAID 1 disks can be used to store copies of Exchange databases. These RAID 1 disks can be created by using 10 disks in the front bay. The remaining two drives in the front bay can be used as Exchange restore LUN and hot spare for the RAID disks. The rear-accessible 2.5-inch drives can be configured as a RAID 1 disk and used for deploying operating system. Figure 4 illustrates this configuration.

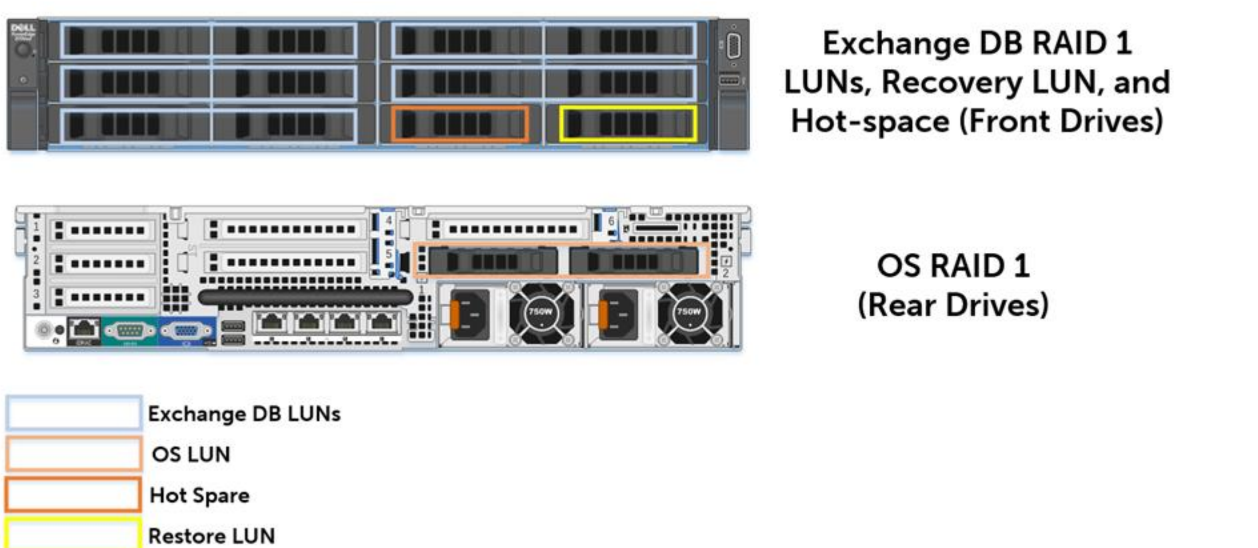

Figure 4 RAID LUN layout for the Exchange databases

With this disk layout, each RAID 1 LUN can be used to store two copies of Exchange database – one active and one passive. This is shown in Figure 5.

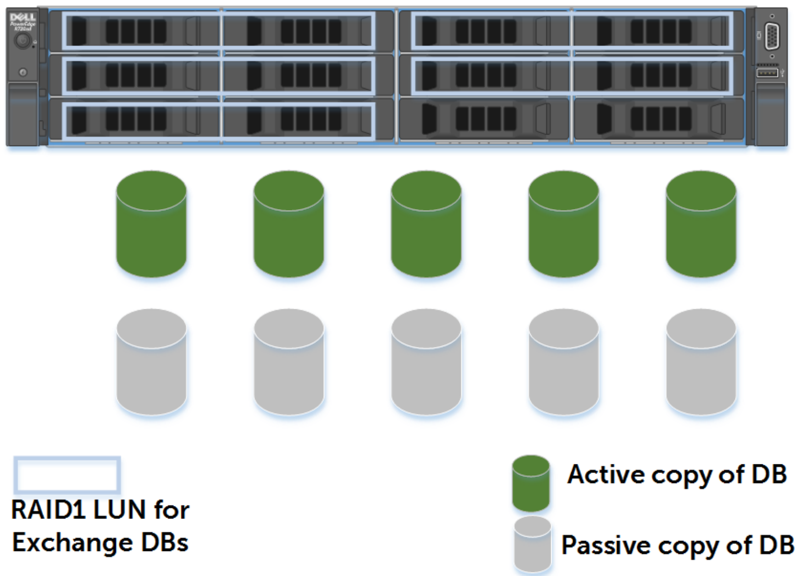

Figure 5 Exchange database layout

5.1.2 Network Architecture

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 6 illustrates this.

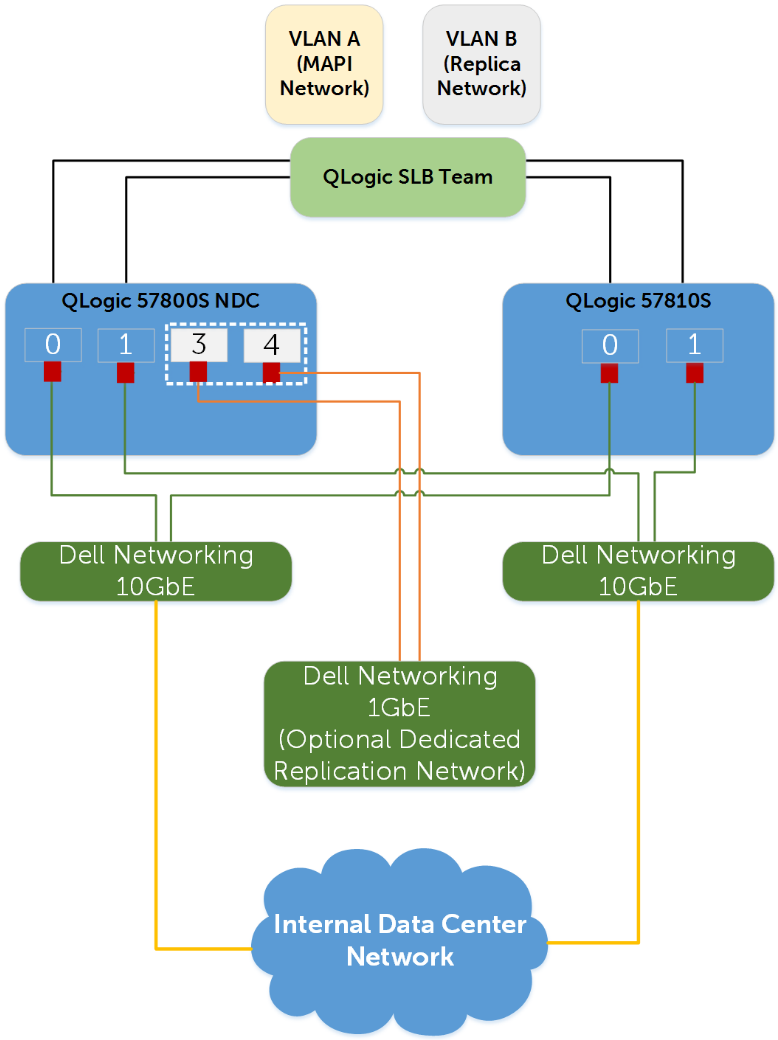

Figure 6 Network architecture for the Exchange deployment

The PowerEdge R730xd server platform supports multiple 10 GbE and/or 1G b 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 6, the QLogic 57800S 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 57800S 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.

5.2 Application-level 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 along with multiple servers hosting Exchange Mailbox and CAS server roles provides highly available messaging services to the end users. Exchange DAG provides native data protection capability. 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 CAS servers in the Exchange deployment.

6 Reference Implementation for 10,000 Mailboxes

This paper presents a reference implementation of the design principles described in section 5 while leveraging the building block architecture described in section 4. As described in section 3.1, the PowerEdge R730xd severs with 12 LFF drives is a suitable building block for Exchange 2013.

The presented solution has a 3-copy DAG layout with Exchange Servers distributed between two sites: Local and Remote. Each server node has five RAID 1 LUNs hosting one active and one passive database per LUN. Each of these databases hosts 500 users with 3 GB mailbox size each. Thus, a single server can accommodate 2,500 users during normal operating conditions. Table 2 provides the technical specifications of the PowerEdge R730xd server.

Table 2 PowerEdge R730xd specifications

Exchange Server System	**PowerEdge R730xd Server with 3.5" HDD Chassis**
CPU	2 x Intel Xeon E5-2683 v3 processor with 14-cores
Memory	Up to 192 GB DDR4
NIC	QLogic 57800S NDC QLogic 57810 PCIe Addon NIC
RAID Controller	PowerEdge RAID Controller H730P Mini
Internal Disks	2 x rear-accessible 1.2 TB SAS 2.5-inch 10K RPM disk drives (Operating System and Application) 12 x LFF 6 TB 7.2K RPM NL-SAS disk drives

Four such servers placed in the local site provide Exchange Mailbox Services for 10,000 users. The mailbox user profile that was tested had 150 messages per day or 0.121 IOPS per user, which included a 20% IO overhead.

Figure 7 represents the distribution of database copies across the DAG members. It shows a 3-copy DAG site resiliency solution with Exchange Servers hosted at both local and remote sites. The local site has one active and one passive copy of each of the databases. The remote site holds one passive copy of each of the databases. If a server fails in the local site, the databases are activated on the surviving hosts to provide mailbox service continuity. In case of a local site failure, the databases are activated in the remote site to provide the Exchange Server service.

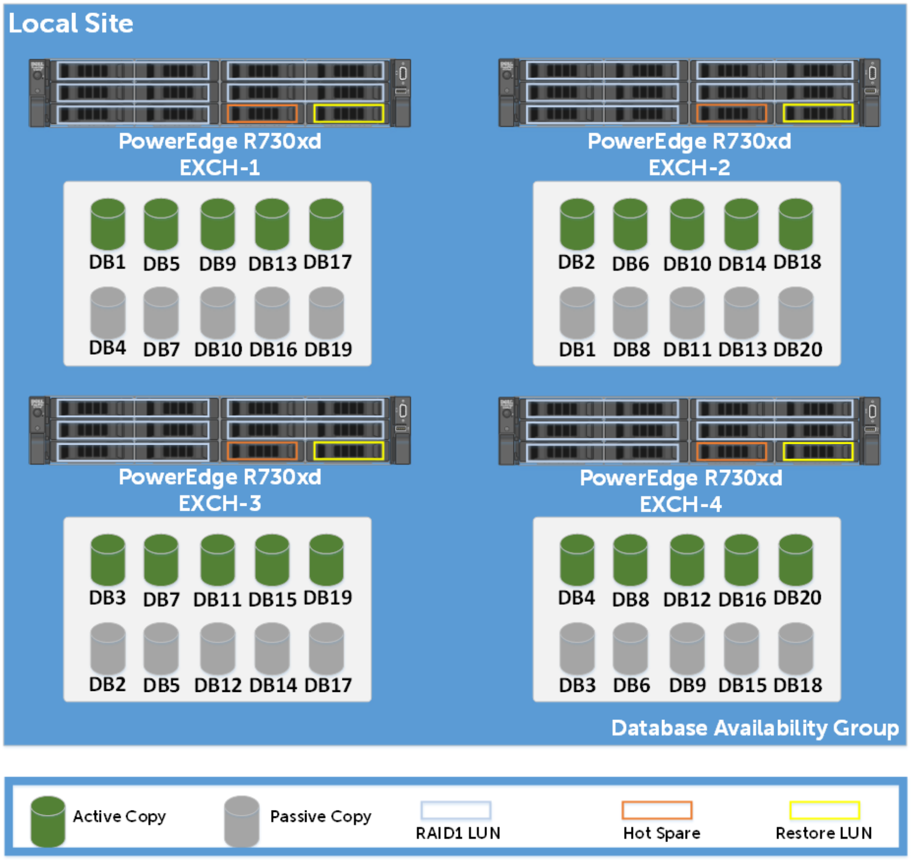

Figure 7 Exchange database layout in the reference implementation

Each Server in the local site deployment stores 10 database copies—5 active and 5 passive. The database layout is planned in a way that no single server contains both the active and passive copies of the same database instance.

6.1 Local Site Architecture

Exchange deployment depends on infrastructure services such as Active Directory (AD), Domain Name System (DNS), and load balancers and so on. The connectivity from the Exchange infrastructure to these infrastructure services and end users must be resilient and highly available. To facilitate this connectivity, this reference implementation employed Dell Networking 10 GbE Top-of-Rack (ToR) switches. This connectivity is illustrated in Figure 8.

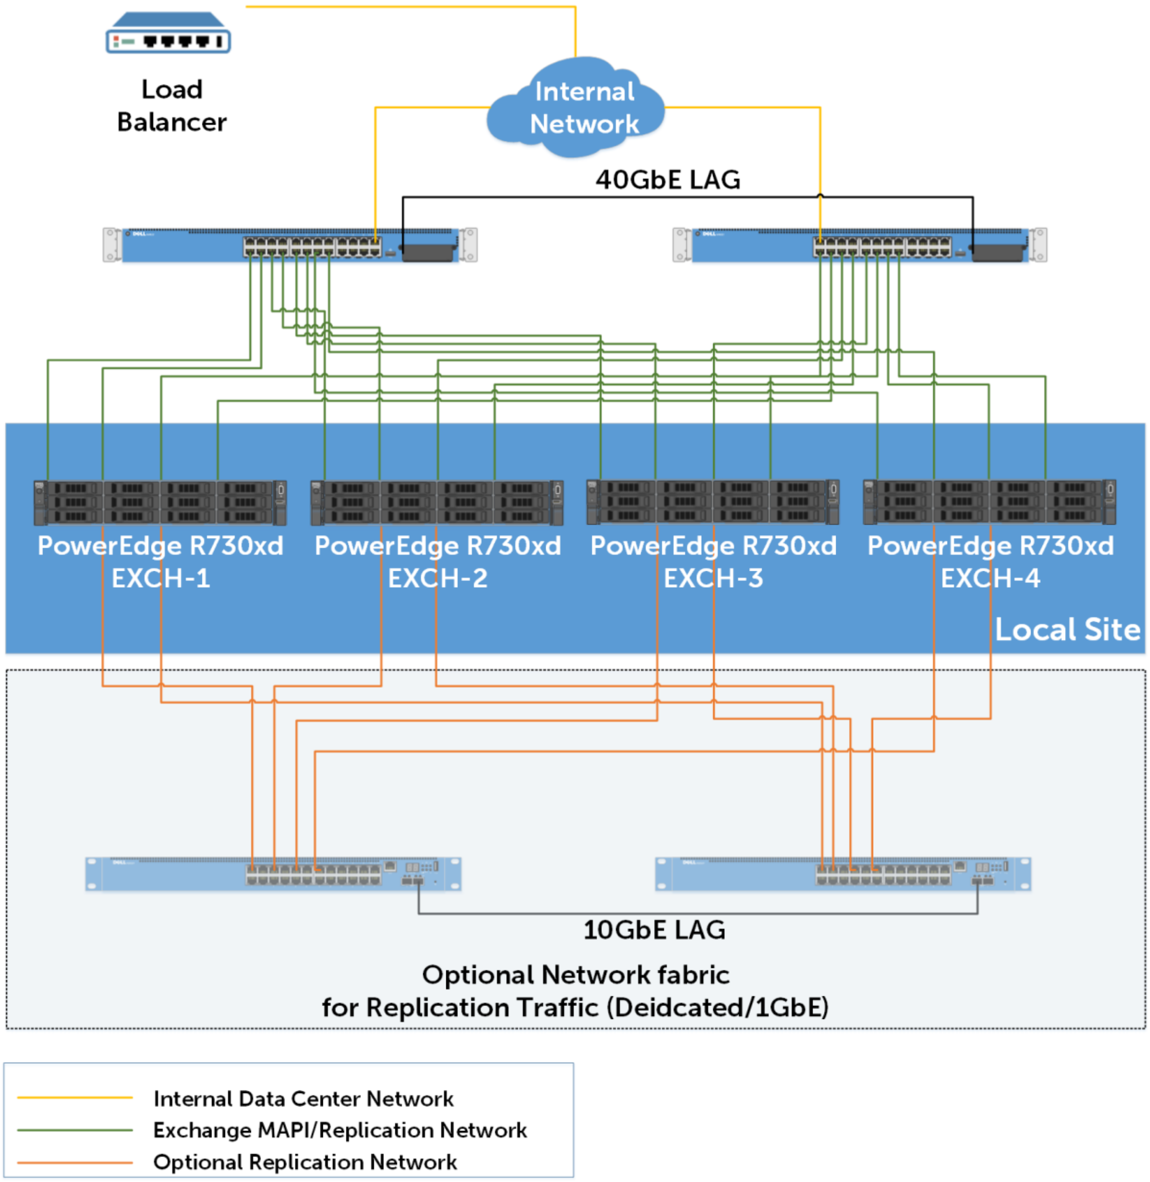

Figure 8 Complete network architecture for the Exchange reference implementation

As mentioned in Section 5.1.2, 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 8, 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 using 10 GbE ports. In this reference implementation, the MAPI and Replication traffic is separated by using VLANs. Because the QLogic 57800S 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.

Figure 9 shows the external data center architecture along with the infrastructure management services.

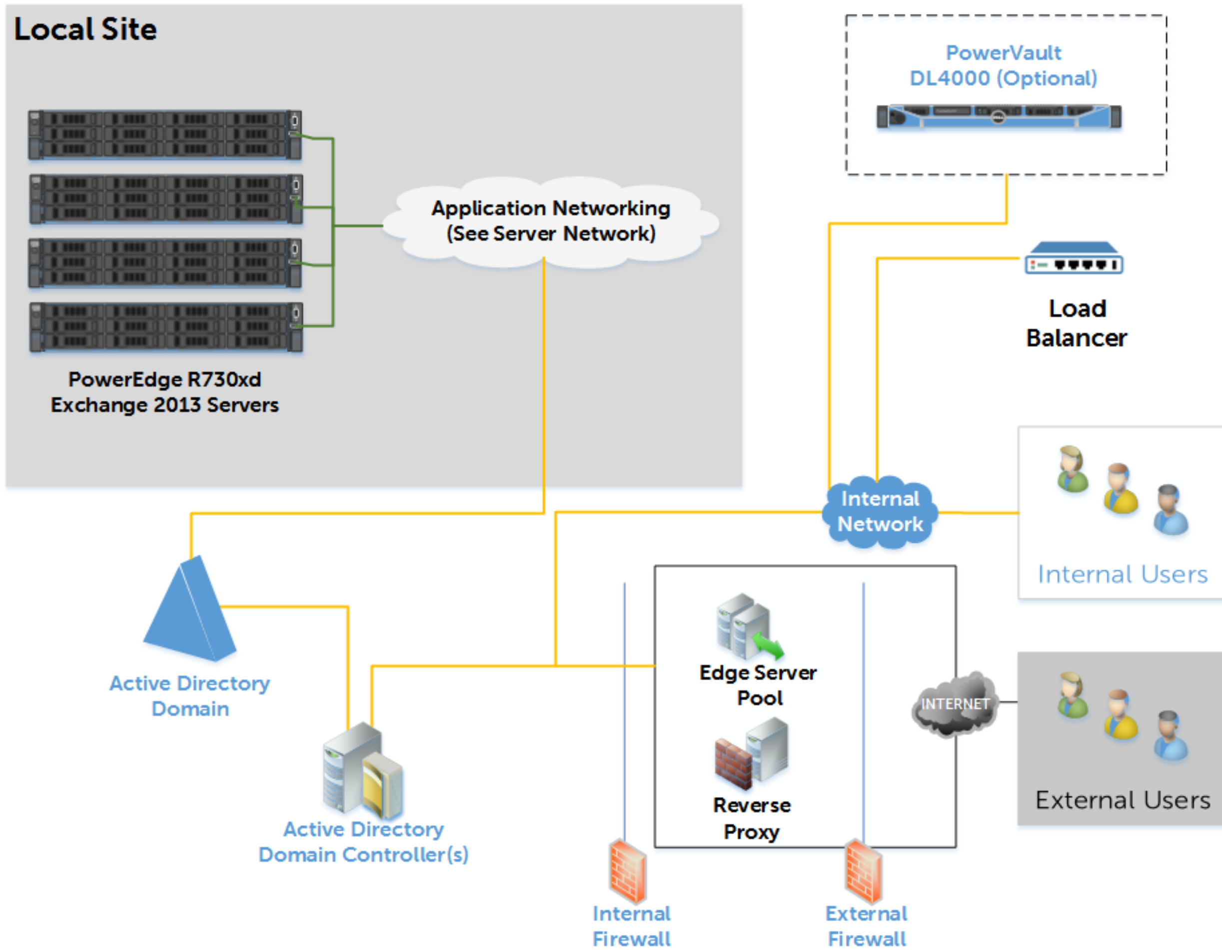

Figure 9 External data center network architecture for Exchange deployment

As shown in Figure 9, the four 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 can be configured to redirect the incoming client requests to the Exchange infrastructure. External mailbox users connect to the local site Exchange deployment over the Internet and through an Edge Server pool configured at the local site.

Also, shown in Figure 9 is an optional deployment of Dell PowerVault DL4000. PowerVault DL4000 provides an application-aware backup and recovery of Exchange data. For more information on PowerVault DL4000, see http://www.dell.com/us/business/p/dell-software-dl4000-backup-and-recovery-appliance/pd

6.2 Resilient Site Configuration

This reference implementation deploys a 3-copy DAG layout with Exchange servers distributed across two sites— two copies at the local site and one at the remote site. It is recommended that every Pod configuration has a remote server(s) hosting the passive copies of all Exchange databases at the local site. This reference implementation uses two Pods and therefore requires four servers at the local site and two at the remote site. This model provides resiliency from a complete local site failure. This is illustrated in Figure 10.

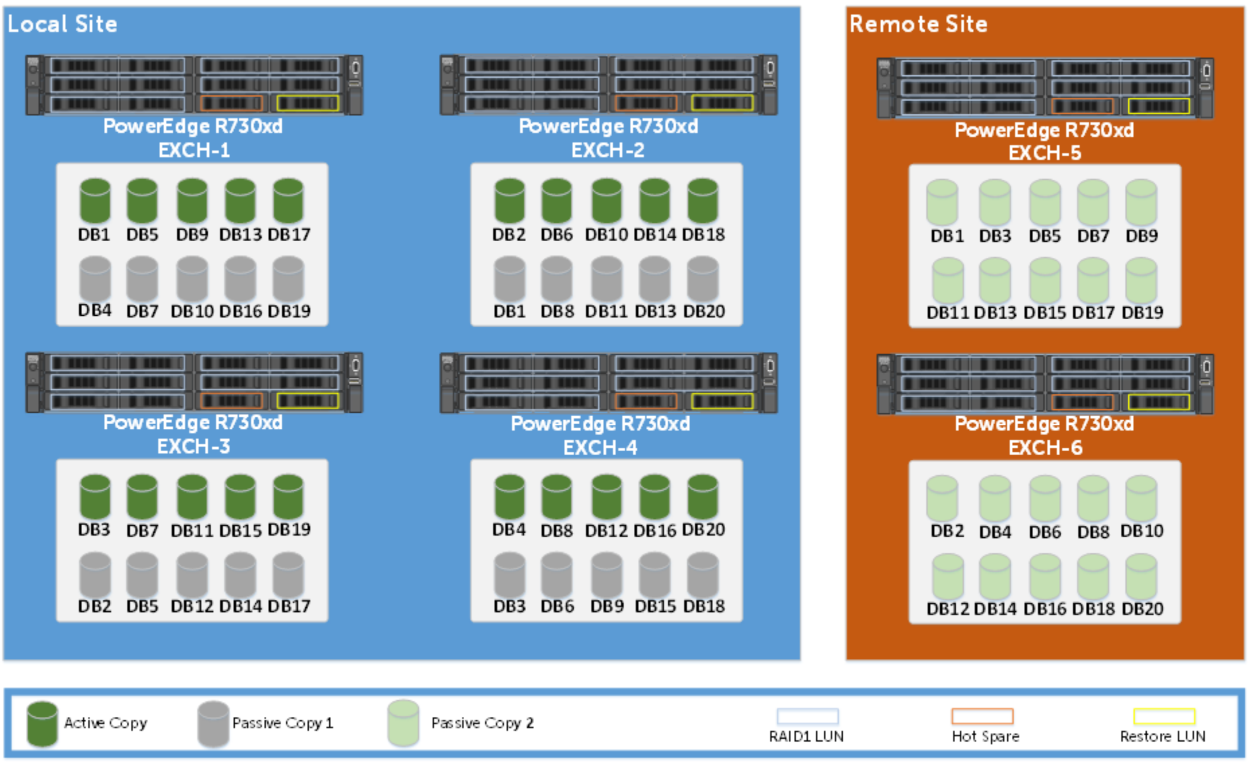

Figure 10 Exchange DAG architecture for the reference implementation

In this reference implementation for 10,000 mailboxes, each server is sized to host up to 5,000 mailboxes. Therefore, servers at the remote site can handle an 10,000 mailbox deployment in case of a local site failure.

7 Verification

This section provides proof points for the performance and resource utilization for the reference implementation for 10,000 users on PowerEdge R730xd.

7.1 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 mailbox profile. The JetStress tool measures how well a storage system can perform 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 reference implementation describes the use of internal storage offered on each PowerEdge R730xd server in the solution infrastructure. Assuming that there is a complete local site failure, the reference implementation was validated for 10,000 active users served by two PowerEdge R730xd servers at the remote site.

Each PowerEdge R730xd server at the remote site hosts 10 active databases in the local site failure scenario. Therefore, each server serves up to 5,000 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 verified 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 3.

Table 3 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
10,000	3 GB	603	923	12 x LFF 6 TB NL-SAS drives in PowerEdge R730xd	RAID 1	20 (2 databases per volume)	6 TB	150 message 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 2 was used to run the mailbox profile test with a planned target of 603 IOPS. The achieved target was 923 IOPS. The results indicated that the storage was able to exceed the target transactional IOPS, well within the latency requirements.

For more information, refer to
http://en.community.dell.com/techcenter/extras/m/white_papers/20441135

7.2 Microsoft Exchange Load Generator Verification

When verifying server and storage for Exchange Server 2013, a best practice is to simulate the worst case scenario under anticipated peak. 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. Based on the DAG configuration, one DAG node failure could be tolerated. 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.

The performance test plan intended to verify the end-to-end performance of Exchange server 2013 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. Three Load Generator instances were used to distribute the generated load across the three DAG nodes. Table 4 briefly describes the test plan for Load Generator tests.

Table 4 Exchange Loadgen Verification Tests

Tests	Test description	DAG state	Database distribution
Test 1	Peak load test with Primary site down: Simulation day: 8 hours Test length: 8 hours	Primary Site Down	Remote Site Server 1 DB1, DB3, DB5, DB7, DB9, DB11, DB13, DB15, DB17, DB19 Remote Site Server 2 DB2, DB4, DB6, DB8, DB10, DB12, DB14, DB16, DB18, DB20
Test 2	Normal load test with Primary site down: Simulation day: 4 hours Test length: 8 hours	Primary Site Down	Remote Site Server 1 DB1, DB3, DB5, DB7, DB9, DB11, DB13, DB15, DB17, DB19 Remote Site Server 2 DB2, DB4, DB6, DB8, DB10, DB12, DB14, DB16, DB18, DB20

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 using performance monitor. Table 5 provides the results captured during the load tests performed.

Table 5 Exchange Load Gen Verification Results

Performance Metrics	Thresholds	Test 1 (Peak load)	Test 2 (Normal load)
%Processor Utilization	< 70 %	23.41	12.19
I/O Database Reads (Attached) Average Latency	< 20 ms	15.44	1.47
I/O Database Writes (Attached) Average Latency	< 20 ms	0.135	0.117
I/O Log Writes Average Latency	< 20 ms	0.178	0.138
RPC Averaged Latency	< 10 ms	2.07	2.04

8 Solution Specifications

Table 6 Solution Specification for the Reference Implementation

Component	Component Details
Microsoft Exchange 2013 Multi-role Server	4 x PowerEdge R730xd 3.5-inch chassis with 12 LFF drives at the local site 2 x PowerEdge R730xd 3.5-inch chassis with 12 LFF drives at the Remote site
CPU	2 x Intel Xeon E5-2683v3 14 core processors per server
Memory	Up to 192 GB DDR4 per server
NIC	QLogic 57800S NDC QLogic 57810 PCIe Addon NIC
RAID Controller	PowerEdge RAID Controller H730P Mini
Internal Disks	2 x 1.2 TB SAS 2.5-inch 10K RPM disk drives (Operating System and Application) 12 x 6 TB Nearline SAS 3.5-inch disk drives
Network Fabric	2 x Networking 10 GbE switches for LAN connectivity

9 Solution Summary

The PowerEdge R730xd server offers an excellent balance of an ultra-dense internal storage, redundancy, and value in a compact form factor. It is a highly efficient and cost-effective hardware building block for any mid-size or large business that requires scalability in both memory density and storage capacity. It delivers enormous storage capacity & IOPS performance in a dense 2U form-factor. This enables larger and more efficient databases and mail servers.

The Exchange building block approach proposed in this paper provides a scalable architecture for Exchange deployments. The proposed Pod design includes the minimum server and storage required to meet the Exchange solution requirement, including the mailbox profile. The PowerEdge R730xd server with 12 LFF drives provides sufficient compute and storage capacity required for up to 5,000 Exchange mailboxes of 3 GB in size and 150 messages per day. The design principles described in this paper, along with the building block approach for designing Exchange infrastructure, provide a foundation for implementing scalable Exchange deployments. The reference implementation presented in this paper describes how the Pod consisting of the PowerEdge R730xd server can be scaled out to support 10,000 mailboxes in a site resilient architecture.

This paper also provides proof points for performance of Exchange server with the PowerEdge R730xd server and summarizes the benefits that Exchange can gain from the PowerEdge R730xdserver. The solution architecture for Exchange using the PowerEdge R730xd server minimizes management overhead, while delivering application performance as prescribed by Microsoft.

A Additional resources

Dell PowerEdge R730xd

Dell PowerEdge RAID Controller (PERC) H730P Mini specifications

Microsoft Exchange 2013 ESRP Program

PowerEdge R730xd 10,000 Mailbox Resiliency Exchange 2013 Storage Solution