

A Comparative Study of Microsoft® Exchange 2010 on Dell™ PowerEdge™ R720xd with Exchange 2007 on Dell™ PowerEdge™ R510

Incentives for migrating to Exchange 2010 on Dell PowerEdge R720xd

Global Solutions Engineering

Dell Product Group



This document 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.

© 2012 Dell Inc. All rights reserved. Dell and its affiliates cannot be responsible for errors or omissions in typography or photography. Dell, the Dell logo, PowerVault, and PowerEdge are trademarks of Dell Inc. Intel and Xeon are registered trademarks of Intel Corporation in the U.S. and other countries. Microsoft, Windows, and Windows Server are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell disclaims proprietary interest in the marks and names of others.

March 2012 | Rev 1.0

Contents

- Executive summary..... 4
- Introduction 4
- Microsoft Exchange overview 5
 - Performance improvement factors 5
 - Mailbox sizing improvements 5
- PowerEdge R720xd overview 6
- Test components..... 7
 - Hardware and software settings 7
 - Microsoft Jetstress testing tool..... 8
- Test methodology and results 9
- Summary 9
- References 10

Tables

- Table 1. Microsoft Exchange 2007 message profiles 6
- Table 2. Microsoft Exchange 2010 Mailbox message profiles 6
- Table 3. Storage configuration for R510 and R720xd 8
- Table 4. Jetstress Test Results 9

Figures

- Figure 1. Dell PowerEdge R720xd 7

Executive summary

IT organizations of many corporations follow refresh cycles to keep up-to-date on the latest hardware and software technology. Implementing these refresh cycles requires knowledge of new features in the latest generation of software and how they will work with improved servers, storage and associated component hardware. Incorporating these improvements in the solution design during the refresh cycle can lead to significant consolidation benefits. This white paper focuses on performance improvements in Microsoft Exchange Server from the 2007 version to the 2010 version and the benefits of using the latest generation Dell PowerEdge servers to provide significant performance improvement per server.

Introduction

There are two distinct aspects in choosing the correct building blocks for a messaging solution. The first task is to evaluate the appropriate number of mailboxes, maximum mailbox size and mailbox profile that the solution can support given the constraints of hardware. Based on these calculations, the correct server and storage combination to support a single instance of the solution can be calculated. This calculation depends on the projected storage I/O throughput and on the CPU and memory capacities of the server to handle the workload.

Due to improvements in mailbox I/O performance in Microsoft® Exchange Server 2010, the IOPS (I/O operations per second) generated per mailbox have been reduced as compared to Exchange Server 2007. These performance improvements allow for a larger number of statistically heavier users (users with a heavier mailbox profile) to be supported on a single hardware building block.

The first task in this study is to determine the appropriate mailbox profile for a typical Exchange 2010 user. Based on the choice of the mailbox profile, the hardware building block is then chosen to determine how many such users can be supported. The choice is driven by the throughput and capacity of the storage and the performance of the server hardware to handle these users.

The latest generation Dell™ PowerEdge™ R720xd is a high storage capacity server with up to 24 front-facing, hot swappable drives and 2 optional back-accessible drives in a 2U enclosure. With the latest Dell PERC H710P storage controller, the potential throughput delivered by these drives has also improved vs. earlier generation servers. Finally, due to increased memory capacity and up to 33 percent increase in CPU cores available on the Intel® Xeon® E5-2600 series processors, the processing capability of the server has increased to handle the additional Exchange user workload.

The study shows a greater mailbox consolidation factor per server unit because of an increased number of available spindles, memory and CPU in the R720xd. Due to these improvements and the changes in the mailbox I/O profile with Exchange Server 2010, a better building block can be used to construct a complete Exchange solution with hardware consolidation and increased performance. The R720xd provides the required resources to better support a large number of mailboxes per server, without compromising on the mailbox size.

Microsoft Exchange overview

Microsoft Exchange Server is the leading enterprise level messaging software. Over the years, several new features have been incorporated to enhance the product. The core component for the mailbox server, the database engine, also has gone through several updates in order to increase efficiency, consolidation and performance.

Performance improvement factors

One of the primary enhancements to Exchange Server 2010 is an increase in the overall efficiency and performance of the database engine. The improvements to individual components of the database engine enable a significant reduction in mailbox IOPS when compared to Exchange Server 2007 and are detailed below:

- **Increased Database page size:** The amount of data stored in a single database page indicates the amount of disk operations required to retrieve information from the database file. The database page size was increased from 8 KB (Exchange Server 2007) to 32 KB (Exchange Server 2010). Increasing the amount of data stored in a single database page, means that fewer I/O operations are needed to complete the same task. There is a four-fold improvement in target IOPS and, consequently, similar improvement in the number of users supported.
- **Schema Changes:** The logical organization for the database structure has also changed to reduce the number of IOPS required for views and search. All previous versions of Exchange had a single message table per database, as well as folders table and attachments table. This database centric table structure was created with Single Instance Storage (SIS) optimization in mind. But in order to achieve the IOPS goal of Exchange Server 2010, the tables were shifted from the database level to the mailbox level. This change led to an increase in sequential I/O operations and a decrease in the total number of I/O operations.
- **Database Cache Enhancements:** Another change in the Exchange store engine behavior that leads to I/O reduction is the transaction checkpoint depth. The checkpoint depth is the maximum amount of data that can be written to transaction logs before the data must be saved to the database file. In Exchange Server 2010, a database that is part of a high available solution has a 100 MB checkpoint depth (active database copy) and 5 MB for passive database copies. This creates opportunities for larger and more sequential I/O's, while decreasing the number of times needed to synchronize pages from store cache to the database file.

Mailbox sizing improvements

One of the main components for correctly sizing a mailbox server for an Exchange Server solution is to predict the amount of disk read and write operations per second, or IOPS, required to sustain an optimal disk response time, or latency, in order to assure a reliable customer experience.

Several factors come in to play in order to accurately predict IOPS, and they can fluctuate with changes in usage and clients. The main factors that come into play are:

- **Database cache size per mailbox user:** The size of the database cache determines the amount of information (database pages) that can be stored in RAM per mailbox user. Every operation in a mailbox happens within the cache stored in RAM. If we need to read or write any item in the mailbox that is not in the cache, an I/O operation is generated to retrieve the necessary

information and to flush database pages that were previously stored in RAM. Increasing the database cache per mailbox user decreases the number of IOPS per mailbox.

- **User profile:** The user profile indicates the actual transactions being processed by the database engine. The primary information used to calculate the profile is the total number of messages sent and received per day and the average message size. Each version of Exchange Server will consider a certain amount of messages and their size as a light, medium or heavy profile and associate an IOPS number to it. The IOPS number has been calculated throughout the years by Microsoft and updated along with every version of Exchange Server.

Based on the above factors, the following user profile tables have been provided for Microsoft Exchange Server versions. As can be seen from the tables, there is not an exact correlation between the two based on the number of messages per user that the user profile can handle. Consequently, the comparison in this paper will be between the Heavy profile of 100 messages on Exchange Server 2007 and 200 messages per user on Exchange Server 2010. Due to the decreased IOPS per user, the same target IOPS per mailbox will support more users. However, due to the increase in Database Cache per user, the memory requirements for a given mailbox profile also increase and should be taken into account while configuring the physical server.

Table 1. Microsoft Exchange 2007 message profiles

User type (usage profile)	Messages Sent/received per day	Database cache per user	Estimated IOPS per user
Light	5 sent/20 received	2 MB	0.11
Heavy	20 sent/80 received	5 MB	0.32
Extra heavy	40 sent/160 received	5 MB	0.64

Table 2. Microsoft Exchange 2010 Mailbox message profiles

Messages sent/received per mailbox per day	Database cache per user (MB)	Estimated IOPS per user
100	6	0.1
200	12	0.2
300	18	0.3
400	24	0.4
500	30	0.5

In addition to the improvements in hardware as shown in the next section, the IOPS capacity per server also has greatly increased. The combination of these two improvements leads to dramatic improvements in the performance of a server running an Exchange Server 2010 as compared to Exchange Server 2007 on older generation hardware.

PowerEdge R720xd overview

The Dell PowerEdge R720xd is a 2U rack-mount, high storage-capacity server featuring the latest generation Intel processor technology and updated storage subsystem consisting of a PERC H710P storage controller and options of SATA, SAS, Nearline SAS and SSD drives.

The server can support up to two Intel Xeon E5-2600 series processors with up to 8 cores per processor. Each server has 24 DIMM slots and can support a variety of DIMM capacities ranging from 2 GB to 32 GB. The maximum RAM potential for each server is 768 GB.

The server can support both 1 GbE and 10 GbE Network Daughter Cards (NDCs) with up to 4-port 1 GbE NDCs; or 2-port 10 GbE and 2-port 1 GbE NDCs. It can support up to 12 3.5" or 24 2.5" (as shown in Figure 1) SAS, SATA, Nearline SAS or SSD drives. On the back, it has slots for 2 additional 2.5" SAS, Nearline SAS or SSD drives. The internal RAID controllers supported are PERC H310, H710 and H710P; the external RAID controller supported is PERC H810. Additional external Disk Attached Storage (DAS) is available using the Dell PowerVault™ MD1220 modules.

Figure 1. Dell PowerEdge R720xd



R720xd is an excellent choice to host Microsoft Exchange Server 2010 workloads. The increased capacities in the platform support the Exchange Server 2010 workload profile. The R720xd provides a balance between performance and capacity of both server and storage, allowing greater flexibility for designing solutions and avoiding trade-offs in number of mailboxes or mailbox size.

Test components

Hardware and software settings

In order to compare across generations of software, it is important to consider the corresponding hardware generation that would be running the older software. For the purposes of this paper, a PowerEdge R510 server was used as a test bed for Exchange Server 2007. The R510 provides a better storage capacity for a single 2U Dell server configuration in its generation of servers and provides a good basis of comparison for the R720xd running Exchange Server 2010.

Disk subsystem configuration

Table 3 lists the disk array and partition configuration for both R510 and R720xd servers. The R510 has up to 12 internal hot-swappable disks, and thus 6 RAID 1 volumes were created. The operating system and software were installed on a separate pair of internal hard drives in a RAID 1 configuration. The R720xd has up to 24 internal hot-swappable disks, and thus 12 RAID 1 volumes were created. The OS and software were installed on a separate 2 X 2.5" hard drives in a RAID 1 configuration on the back of the system.

Table 3. Storage configuration for R510 and R720xd

	PowerEdge R510	PowerEdge R720xd
Storage controller	PERC H700	PERC H710P
HDD Disks	12 x 2 TB NL-SAS 7.2K 3.5”	24 x 1 TB NL-SAS 7.2K 2.5”
RAID configuration	6 x RAID 1	12 x RAID 1

Exchange Server 2007 database setup

Exchange Server 2007 uses the concept of Storage Groups. Storage Groups are logical containers for databases and associated transaction log files. Microsoft recommends that each new database be placed in its own Storage Group. The physical placement of database files and transaction log files is a key factor in designing the disk subsystem for an Exchange Server 2007. Microsoft also recommends that database files and transaction log files be placed in different volumes. The latest recommendation for maximum database file is around 100 GB per file.

In our study, the number of users was divided into 15 databases. Each database was placed in its own Storage Group. Each disk volume hosted 3 databases. All transaction logs were placed in a single volume. Therefore, out of the 6 RAID 1 disk volumes, 5 of them hosted the 15 databases, and the remaining 1 hosted the transactional logs for each of those databases.

Exchange Server 2010 database setup

Exchange 2010 allows the combining of database and logs on the same volume, backed by the same physical disk. This recommendation is specifically for a high available configuration using a Database Availability Group (DAG). If configuring a non-highly available standalone server configuration, the best practice from recoverability perspective is to move the database and the log file from the same database to separate physical disk volumes. When configured in a highly available configuration, the best practice is for the database size to be a maximum of 2 TB.

For the purpose of this comparison a highly available configuration with a 2-copy DAG was the implemented Exchange configuration. The configuration for a single server in this environment was 12 x RAID 1 volumes across the 24 x 2.5” 1 TB NL-SAS drives. Each volume hosts a single database and its associated log volume, for a total of 12 databases.

Microsoft Jetstress testing tool

The Microsoft Jetstress tool is a tool used to simulate an Exchange Server workload. It is recommended by Microsoft in order to verify the performance and stability of a disk subsystem. It simulates the load produced by database and transaction log files against a specific number of mailbox users and profile. Two types of tests can be run using Jetstress - a **disk subsystem** test that can simulate the peak IOPS load from an Exchange server and a **mailbox profile** test that is used to simulate the load for a given number of users with a particular user profile.

Because Jetstress is a disk subsystem tool that simulates an Exchange Server workload, the amount of memory in each system is not a key factor in the testing methodology. However, the server should have sufficient memory to take full advantage of the database engine cache enhancements. In a production environment, the number of users, roles and third party software installed in the server will affect the amount of memory and performance of an Exchange Server.

Test methodology and results

As detailed above, the Jetstress Testing tool can be run in two modes: disk subsystem test and mailbox profile test. Microsoft provides recommended values for latency for disk operations, specifically the disk read latency and the log write latency. The ideal configuration of Exchange should not exceed these recommended values under the maximum load. In order to determine the IOPS generated by an Exchange server running within acceptable latency values and under maximum load, the disk subsystem test was run on both the R510 server running Exchange Server 2007 and the R720xd server running Exchange Server 2010.

The IOPS results from the disk subsystem test were then used to determine the maximum number of users that could be supported for the respective user profiles in the two Exchange setups. This calculation depends on the I/O profiles of the particular user profiles as detailed by Microsoft. However, an additional overhead of 20 percent over the I/O profile was used in the calculations. This was done based on Microsoft recommendations to consider peak load conditions when many users are simultaneously accessing their mailboxes. The derived values for maximum number of supported mailbox users are shown in the table below, along with a comparison of the user profiles for the two setups.

Table 4. Jetstress Test Results

		Exchange 2007 on R510	Exchange 2010 on R720xd
Derived values from Disk Subsystem Test	Message Profile (message/day)	100 messages / day	200 messages/day
	I/O profile (IOPs/user)	0.32	0.2
	Number of Mailbox users	1500	5000
Mailbox Profile Test Performance Results	Target Transactional IOPs	480	1200
	Achieved Transactional IOPs	503	1345
	DB read latency (ms)	17.4	17.06

To confirm the derived results for maximum number of users above, the second Jetstress test, mailbox profile test, was run. The test results, as tabulated above, show that the achieved transactional IOPS exceeded the target transactional IOPS, and the database read latencies were within the 20 ms threshold as defined by Microsoft. As can be ascertained from the table, running the latest generation of Microsoft Exchange on the latest generation Dell hardware supports up to three times the number of users; each of those mailboxes is capable of handling up to twice the number of messages as before.

Summary

The decision to upgrade the infrastructure for an IT organization is a critical one and it is important to understand the changes in both software and hardware technology. As was shown in this paper, upgrading to Exchange Server 2010 running on the latest generation Dell PowerEdge R720xd servers provides significant consolidation benefits. Each server can support up to three times as many users as with Exchange Server 2007 running on PowerEdge R510; and each user can send and receive twice as many messages as before. Using these metrics can help an IT organization achieve consolidation of server resources, while helping alleviate end user mailbox constraints.

A Comparative Study of Microsoft® Exchange 2010 on Dell™ PowerEdge™ R720xd with Exchange 2007 on Dell™ PowerEdge™ R510

References

<http://www.dell.com/exchange>

Exchange Server 2007 Mailbox Server Storage Design - [http://technet.microsoft.com/en-us/library/bb738147\(EXCHG.80\).aspx](http://technet.microsoft.com/en-us/library/bb738147(EXCHG.80).aspx)

Determining the Exchange 2010 Mailbox IOPS Profile - <http://technet.microsoft.com/en-us/library/ee832791.aspx>