# Dell™ PowerEdge™ R720xd with PERC H710P: A Balanced Configuration for Microsoft® Exchange 2010 Solutions

*A comparative analysis with PowerEdge R510 and PERC H700*

**Global Solutions Engineering**

**Dell Product Group**

March 2012| Rev 2.0

# Contents



# Tables



# Figures

# Executive summary

Microsoft® Exchange Server 2010 running on the new Dell™ PowerEdge™ R720xd rack server provides a flexible and scalable messaging platform supporting up to 130 percent more Exchange 2010 mailboxes than the comparable solution based on the previous generation PowerEdge R510. A densely integrated server and storage building block, the R720xd is ideal for building power-efficient and balanced solutions for Exchange messaging workloads.

The R720xd's internal storage options provide a level of density that previously required a standalone server and external Direct Attached Storage (DAS). The R720xd can eliminate the need for DAS, which in turn can save valuable rack space and lower power, cooling, and associated costs.

This technical white paper compares the performance of an Exchange 2010 solution on the R720xd with a Dell PowerEdge RAID Controller (PERC) H710P and the PowerEdge R510 (R510) with a PERC H700, and then presents a sample reference configuration for up to 7,500 users with a heavy I/O (~150 messages/day/mailbox) profile. Based on the results, the R720xd provides a significantly better building block for various Exchange solutions than the previous generation PowerEdge R510 with PERC H700.

Until the launch of the R720xd, the PowerEdge R510 server was the densest Dell rack-based server using internal storage. The R510 can leverage up to 12 x 3.5" internal drives, while the R720xd can scale to 24 x 2.5" internal drives. Using the Microsoft Jetstress 2010 (Jetstress) to simulate Exchange I/O workload, the R720xd showed approximately 2.5 times the number of Input / Output Operations Per Second (IOPS) achieved for RAID-1 and RAID-10 configurations and approximately twice as many IOPS achieved for RAID-0 configuration. As a result the R720xd can support up to 2.3 times the number of mailboxes per server, as compared to the R510. This performance improvement confirms the R720xd as a premiere storage-dense server building block for Exchange solutions.

# Introduction

The PowerEdge R720xd is a rack-based server with improved computing capability, more memory and twice the number of disks as the R510 in the same 2U of rack space. With multiple options for internal storage and rich configurations for processors and memory, the R720xd is a more balanced building block for Exchange solutions than the R510.

Some of the previous Dell Exchange solutions with the R510 used the PERC H700 as the storage controller. PERC H700 belongs to the PERC 7 family of RAID controllers. Along with the new generation of servers, Dell has launched a new storage controller called the PERC H710P (PERC 8 family) with improved processor and cache memory.

The following sections of this paper provide a brief introduction to the new R720xd and compare it with the previous generation R510 server. They also outline a brief comparison between the PERC H700 and the PERC H710P storage controllers. The paper then describes the test methodology used to compare the performance of Exchange with the R720xd and the R510 from the storage perspective. Jetstress tests were performed with the R510 with a H700 storage controller against the R720xd with a H710P storage controller. The experimental results suggest that the R720xd is a better building block for Exchange solutions. Based on the results, this paper proposes a sample Exchange reference configuration for up to 7,500 users spanning across three R720xd servers with 24x 2.5" internal drives in a 'Just a Bunch of Disks' (JBOD) configuration.

# Dell PowerEdge R720xd Server from an Exchange perspective

The PowerEdge R720xd is Dell's 12th generation, 2U storage dense, rack-mounted server, which can host up to 24 x 2.5" or up to 12 x 3.5" internal drives, in addition to optional 2 x 2.5" back-accessible drives. It uses Dell's new and improved PowerEdge RAID Controllers (PERC 8 family), which offer improved disk subsystem performance. The server also includes the latest generation of Intel® Xeon® Processors E5-2600 series and supports up to 768 GB of memory. Figure 1 shows a logical layout of R720xd server.

Figure 1. PowerEdge R720xd logical layout

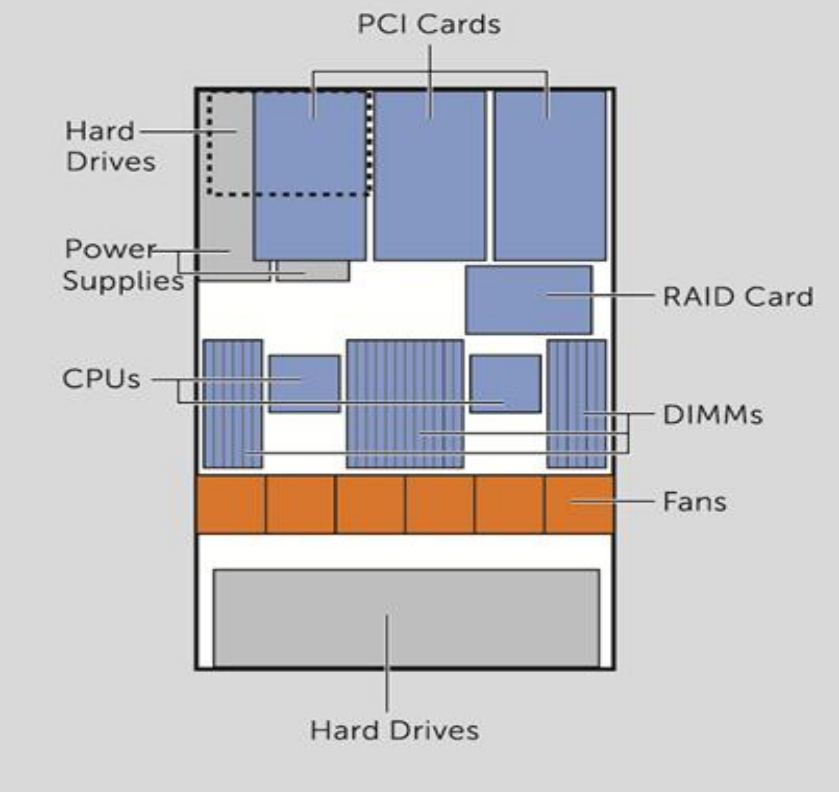


With an option of up to 24 front-loading and 2 back-accessible internal drives, the R720xd offers sufficient disk IOPS and storage capacity to host a large number of Exchange 2010 mailboxes. Improved processor performance and large server memory capacity allow Information Technology (IT) architects to design flexible Exchange solutions. For example, R720xd can be run as a multi-role server, which consolidates HUB Transport (HUB), Client Access (CAS) and Mailbox (MBX) roles on the same server. The large server memory capacity with the matching I/O performance from the storage subsystem also can accommodate more Exchange users per system.

Exchange 2010 server uses Database Availability Groups (DAG) to provide mailbox high availability. Each mailbox server within the DAG should have enough head room to absorb extra workload, should failures happen to any member server. With its improved processor, memory, RAID controller design and dense storage subsystem, the R720xd is a powerful compute and storage building block for Exchange solutions.

# Dell PowerEdge R510 and PowerEdge R720xd

This section provides a comparative overview of the R510 and the R720xd and describes feature enhancements from the PERC H710P over the PERC H700.

Table 1. Comparison of the R510 and the R720xd

| | PowerEdge R510 | PowerEdge R720xd |
|---|---|---|
| **Rack Units** | 2U | 2U |
| **CPU** | Intel Xeon processors - 5500 and 5600 Product Family | Intel Xeon E5-2600 series processors |
| **Memory slots** | Up to 8 x DDR3 DIMM slots (up to 128GB) | Up to 24 x DDR3 DIMM slots (up to 768 GB) |
| **Internal storage** | 4, 8, or 12 x 3.5" SAS, NL-SAS, SATA drives; Plus 2 x 2.5" internally cabled drive bay | 24 x 2.5" SAS, NL-SAS, SATA drives<br>12 x 3.5" SAS, NL-SAS, SATA drives<br>Plus optional 2 x 2.5" rear-accessible drive bay |
| **PCIe support** | 2.0 | 3.0 |
| **LOM** | 2 x 1GbE | Multiple NDC options:<br>minimum 4 x 1 GbE |
| **PERC support** | PERC 7 series | PERC 8 series |

The R510 is a 2-socket 2U, multi-purpose value server, offering a balance of internal storage, redundancy and value in compact 26" deep chassis. It provides a choice of chassis configurations with 4, 8 or 12 front loading drive bays, support for Six-Core and Quad-Core Intel Xeon Processors, internal capacity for up to 12 x 3.5-inch, hot-plug, 6.0-Gbps, serial-attached SCSI (SAS) or Near Line SAS (NL-SAS) hard drives, and integrated RAID support via a PERC H700 adapter.

The R720xd can support up to 33 percent more processing cores than the R510, enabling support for more Exchange users/mailboxes per server as compared to the R510. The R720xd has three times more memory slots, compared to the R510. In the case of internal storage, the R720xd supports an option for twice as many drives when compared to the R510, providing flexibility in RAID configurations. From a networking perspective, the R720xd provides more on-board network interface options than the R510; it also can include optional 2 x 1GigE and 2 x 10GigE Network Daughter Cards (NDCs). This provides flexibility in designing Exchange public, private and database networks.

For the purpose of this study, the PERC H700 was used as storage controller for the internal storage of the R510 and the PERC H710P for the internal storage of the R720xd. The following subsection compares the storage controllers.

## Dell PERC H700 and PERC H710P

The PERC H700 is the internal host-based RAID controller used to connect to the R510 backplane that supports up to 12 drive bays. The PERC H700 controller supports an 8-port LSI 2108 Chipset and 512 MB or 1GB of customized DDR2 400MHz with Error-Correcting Code (ECC) cache memory.

The PowerEdge R720xd can support the new PERC 8 series controllers. The latest generation PERC controllers use SAS RAID-on-Chip, 8-port Dell Adapter with LSI 2208 chip for even greater performance. The PERC H710P is the follow-on to the PERC H700 and was used as the internal RAID controller on the R720xd. With up to 1 GB DDR3, battery backed 1333 MHz cache memory, the PERC H710P can support RAID 0, 1, 5, 6, 10, 50, and 60 with online capacity expansion and hot swap device support. Both controllers support 6 Gbps SAS as the storage interconnects technology. Table 2 shows the PERC H710P enhancements over PERC H700.

Table 2. PERC H700 v/s PERC H710P

| Specifications | PERC H700 | PERC H710P |
|---|---|---|
| Processor | Dell adapter SAS RAID-on-Chip, 8 lanes with LSI 2108 chipset | Dell Adapter SAS RAID-on-Chip, 8-port with LSI 2208 chipset |
| Cache memory | 512MB/1GB DDR2, 800 MHz | 1GB DDR3, 1333MHz |

Due to more compute capacity and faster cache memory, PERC H710P is expected to process more I/O requests as compared to the PERC H700 and provide a better performance in terms of IOPS.

# Test methodology

Exchange 2010 Server with recent architectural enhancements supports larger mailboxes. Reduction in IOPS and sequential behavior of database access patterns with Exchange 2010, have enabled the practice of using low speed, inexpensive NL-SAS drives. Exchange 2010 deployment performs equally well with NL-SAS and SATA drives as with Fibre Channel or high RPM SAS class drives. The low speed drives are usually available in large capacities and are appropriate for larger mailboxes. Because NL-SAS is a cost effective and higher capacity drive option, the studies described in this paper utilized 2.5” NL-SAS drives.

The intent was to study and compare the performance of Exchange from the I/O perspective, with the R720xd and PERC H710P configuration against the R510 and PERC H700 configuration. To do a one-to-one comparison between the R720xd and the R510 in terms of I/O performance, we used similar configurations such that the delta would primarily be in the performance improvement of the storage subsystem. Microsoft Jetstress was used as a comparison tool to perform the studies. Both servers were prepared with hardware configurations as shown in Table 3.

Table 3. Server configurations for Jetstress tests

| | R510 | R720xd |
|---|---|---|
| **Internal Storage** | 12 x 2TB, 3.5" NL-SAS | 24 x 1TB, 2.5" NL-SAS |
| **RAID configurations** | RAID 0, RAID 1, RAID 10 | RAID 0, RAID 1, RAID 10 |
| **PERC** | H700 | H710P |

The strategy was to test whether the storage system can withstand the Exchange load for a given configuration. In order to verify this, Jetstress Disk Subsystem Throughput tests were performed. The R510 and R720xd were configured with an equal amount of memory and internal storage in order to perform a direct comparison with the throughput of the storage system. In order to observe consistency of the performance improvement from the R510 to the R720xd, the performance was measured with three RAID configurations, namely RAID-0, RAID-1 and RAID-10. For the RAID-0 test, each disk was configured as its own RAID-0 container. Thus, there were 24 x RAID-0 containers configured on the R720xd and 12 x RAID-0 containers on the R510.

**Jetstress Disk Subsystem Throughput Test** is used to determine the peak working IOPS value that the storage subsystem can handle while remaining within the disk latency targets established by Microsoft Exchange.

The objective was to observe if the R720xd disk subsystem can handle more disk IOPS than the R510. From an Exchange perspective, more IOPS per server would indicate that more mailboxes per server could be supported. The Exchange transactional I/O performance with the R720xd and the R510 were observed with the three RAID levels for the configurations listed in Table 3. The performance of the R720xd relative to the R510 was measured in terms of Average Database Transactional I/O per second, which measures the number of reads and writes performed with an Exchange database per second, averaged over the duration of test.

## Results and analysis

The R720xd achieves approximately twice as much average transactional I/O performance relative to the R510 for a RAID-0 configuration, in which each internal drive contributes as a RAID-0 volume for a single Exchange database and logs combined. The performance of the R720xd improves with RAID-1 and RAID-10 configurations, relative to the R510. Approximately 2.5 times the performance in average transactional I/O per second was observed with RAID-1 and RAID-10 configurations. Multiple disk subsystem throughput tests with the RAID configurations confirmed the performance improvements with the R720xd and provided proof points demonstrating that the R720xd with a PERC H710P can withstand higher I/O as opposed to the R510 with a PERC H700. With double the number of drives supported on the R720xd, we would expect twice the IOPS achieved: however, the tests indicated greater than 2 times improvements. This additional improvement can be attributed to the enhanced I/O subsystem on the R720xd, in particular to the new RAID controller. Figure 2 depicts the R720xd's performance improvement in terms of average transactional I/O per second relative to the R510, for all three RAID configurations.

Figure 2. Performance of the R720xd, in terms of Transactional I/O, relative to the R510

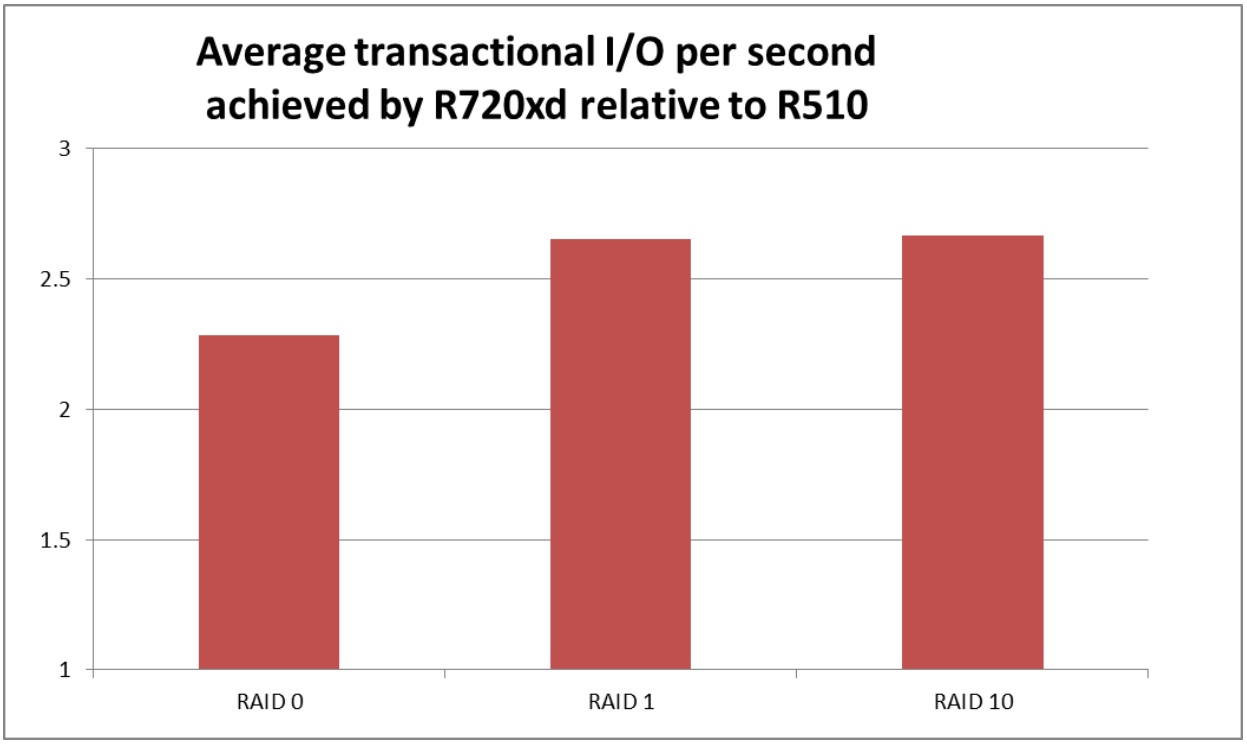


The overall results from the Jetstress disk subsystem throughput test depicted better performance with the R720xd, implying faster Exchange performance in terms of data transfers, and significant increase in IOPS per server. Based on the results of the Disk Subsystem Throughput test, mailbox profiles for the R720xd and the R510 were designed as shown in Table 4. A mailbox profile of 150 messages per day was selected which, according to Microsoft guidelines for Exchange 2010, would amount to 0.15 IOPS per mailbox. However, for our test we added an overhead of 20%, to account for a higher load during peak hours. This is also a best practice and recommendation prescribed by Microsoft.

Table 4. Exchange mailbox profiles

| Servers | Number of Mailboxes | Mailbox size (GB) | Target IOPS | Spindles | RAID configuration | Number of Exchange databases | I/O profile |
|---|---|---|---|---|---|---|---|
| R510 | 3,250 | 1.5 | 585 | 12 x 3.5” 2 TB | RAID-0 (each drive a separate volume) | 12 (1 DB per drive) | 150 messages a day (0.18 tested) |
| R720xd | 7,500 | 1.5 | 1350 | 24 x 2.5” 1TB | RAID-0 (each drive a separate volume) | 24 (1 DB per drive) | 150 messages a day (0.18 tested) |

A Jetstress mailbox profile test was executed using the parameters detailed in Table 4. The results indicated that both the R510 and the R720xd were able to meet and exceed the target average transactional IOPS. In the case of the R510, the achieved transactional IOPS were 662, which exceeded the planned target of 585 IOPS. The R720xd achieved 1436 average transactional IOPS, exceeding the planned target of 1350 IOPS. The R720xd exhibited a stellar performance in terms of average transactional IOPS, staying well within the latency requirements. The R720xd could support 2.3 times mailboxes as compared to the R510; up to 7,500 mailboxes were sustained by the R720xd, compared to 3250 on the R510.

**Jetstress Mailbox Profile Test** allows you to verify whether the storage system meets or exceeds the planned Exchange mailbox profile; meaning whether the Exchange profile is configured per the requirements of the underlying storage system.

Improvements in terms of mailbox density per server is shown in Figure 3 and performance improvement in terms of achieved average transactional IOPS for the mailbox profile shown in Table 4, is shown in Figure 4. Figure 3 gives a comparative overview of the number of mailboxes that could be supported by each server within the latency requirements. Figure 4 shows the improvement in achieved transactional IOPS for the planned mailbox profile.

Figure 3. Number of mailboxes supported by R720xd and R510

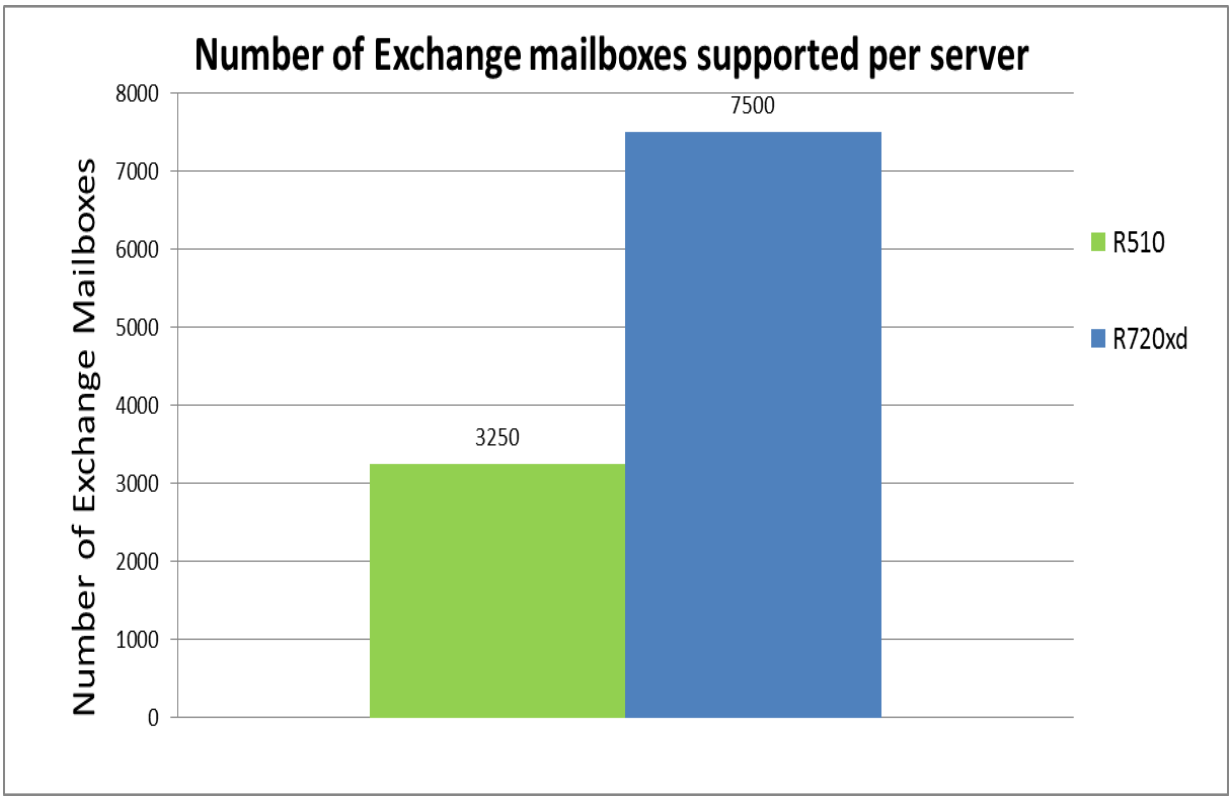

Figure 4. Achieved average transactional IOPS by R720xd and R510

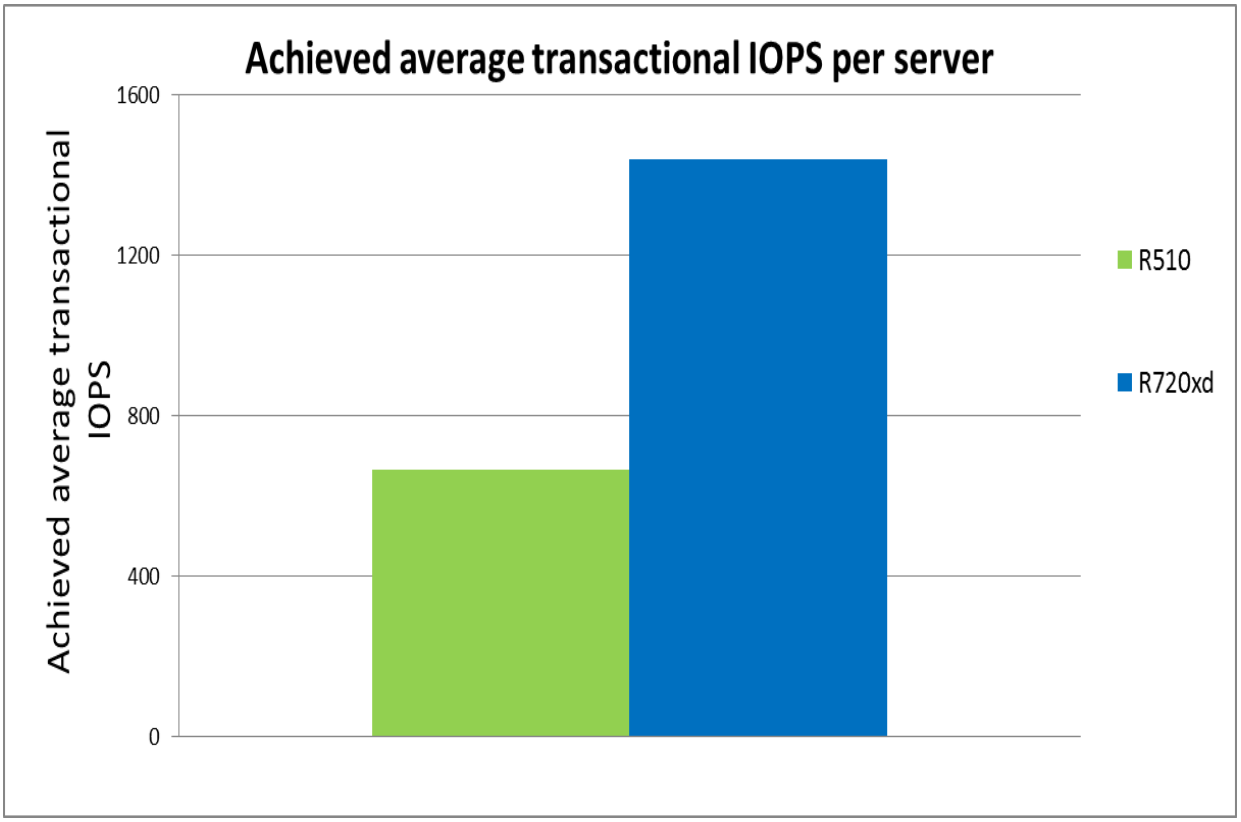


## A sample reference configuration

Microsoft Exchange 2010 Server provides application-level high availability based on the Database Availability Group (DAG). A DAG consists of multiple high-availability copies of Exchange databases across a number of servers. The number of servers determines the size of the DAG. With DAG, Exchange provides various ways to design Exchange infrastructure based on the number of high-availability database copies requirement. Reference configuration for Exchange consists of such a design with a variety of options available from server form factors, storage and disk choices.

With its new high availability feature and non-shared storage for the same database copies, Exchange 2010 enables the use of standalone servers with large internal drives or servers with an external DAS. Such a design allows an easy incremental deployment for additional mailbox growth and higher number of copies. Furthermore, multiple Exchange roles can be combined on servers with a multi-copy environment, enabling smaller and medium-sized organizations to provide redundancy of other roles along with the mailbox role in a cost-effective manner.

With its DAG architecture, Exchange 2010 now provides an option of eliminating RAID, when configured with a large number of copies of the Exchange database. This option frees up the spindles to be used for additional mailboxes or copies. Microsoft recommends using three or more Exchange database copies for non-RAID or JBOD configurations. Because the PERC controller requires a RAID configuration for the disks it manages, the same result can be achieved by using a RAID-0 configuration for each disk. Thus, on an R720xd, the result is achieved by configuring 24 RAID 0 containers.

The goal of the reference configuration is to design a server and storage-based configuration optimized and tuned to provide Exchange deployments in the simplest, most highly available and cost-effective manner. Dell recommends using at least three local copies of the mailbox database in a DAG when configuring the disks in a JBOD fashion, with each disk being its own RAID-0 volume. One such reference configuration, with three R720xd servers depicting a three-copy DAG with 7,500 mailboxes, is shown in Figure 5.

Figure 5. Sample reference configuration: 7,500 Mailboxes over three R720xd servers

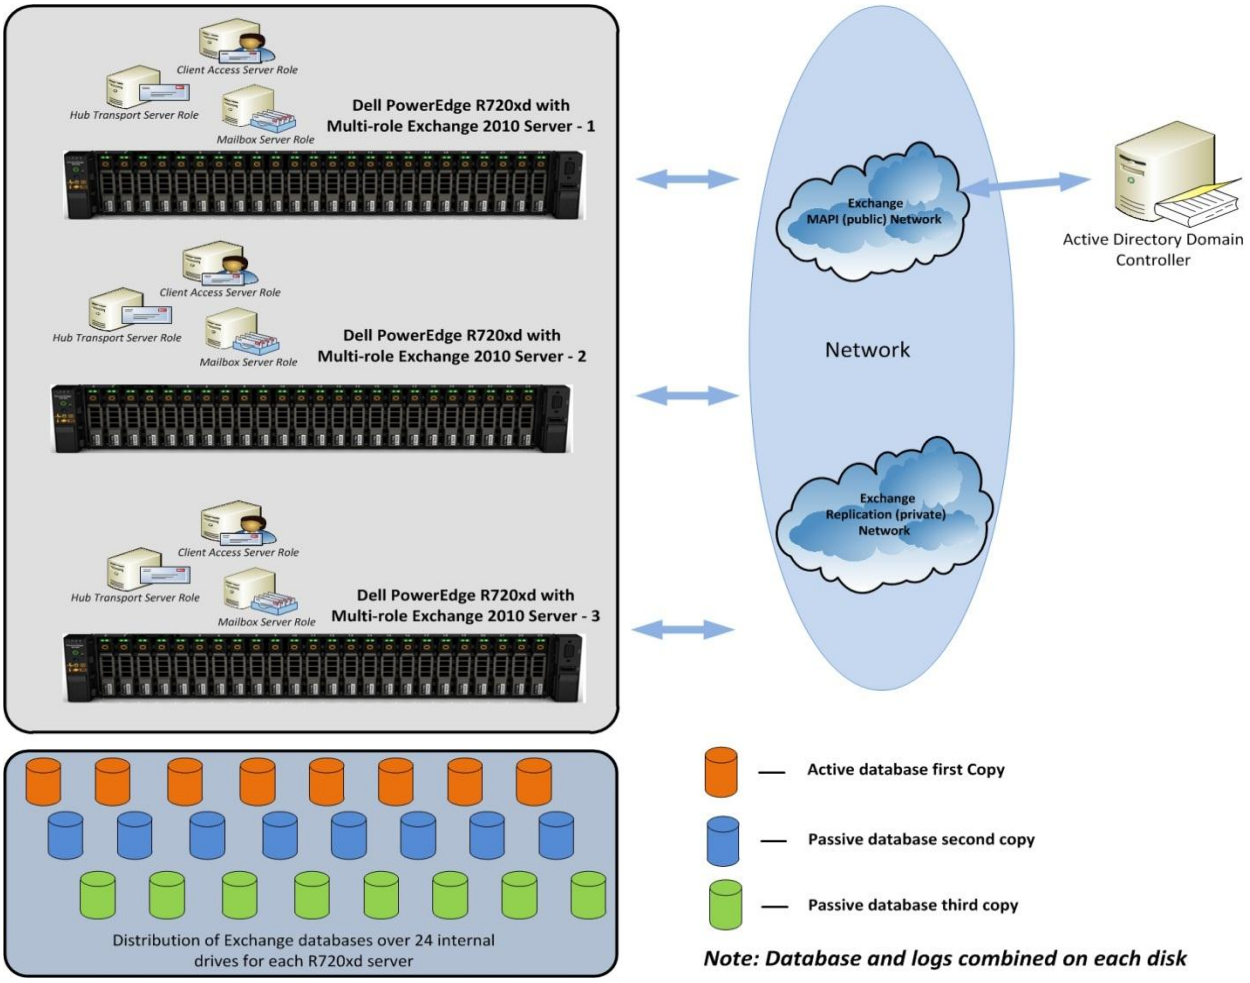


The sample reference design configuration in Figure 5 shows 7,500 heavy I/O profile users (~150 messages per day) with 1.5GB mailbox size. The rack-based 2U R720xd is used as a multi-role Exchange server, with HUB, CAS and Mailbox roles all combined on the server. Each of the 24 x 2.5" internal 1 TB NL-SAS drives are configured as their own RAID-0 container and stores an Exchange Database and the associated Log files. Additional two rear-accessible internal 2.5" 300GB SAS 10K drives are configured as RAID-1 volumes to host the operating system along with Exchange multi-role server application. For the reference configuration, three such R720xd servers are required to host three copies for each database. A hardware load balancer (not shown in Figure 5) is required to load balance the CAS role. In order to maximize the utilization of spindles and provide better failover and recovery, each database with data and log combined is hosted on a single disk. Each server consists of eight active copies and sixteen passive copies. The proposed design counts for single failure at the server level. All copies are replicated without lag to maintain full high availability.

Table 5. Sample reference configuration - solution summary

| | |
|---|---|
| **Number of mailboxes** | 7,500 |
| **Average user I/O profile (messages/day)** | 0.18 IOPS per user (~150 messages/day) |
| **Average mailbox size limit** | 1.5GB |
| **Total active/passive copies per database** | 3 copy DAG (1 active + 2 passive) |
| **Not included in this solution** | Backup and recovery infrastructure<br>UM and Edge Roles<br>Restore LUN |
| **Server Configurations** | **Detail** |
| **All-in-one (Mailbox/Hub/CAS) server (Processor/Memory)** | 3 x PowerEdge R720xd servers<br>2 x Intel Xeon E5-2600 series processors CPU, 8 cores per processor<br>64 GB of DDR3 Memory |
| **Number of DAGs** | 1 |
| **Servers per DAG** | 3 |
| **Number of Active and Passive Mailboxes per Server** | 2500 Active + 5000 Passive |
| **Hub/CAS (Client Access Server) server** | Consolidated with Mailbox Server |
| **Storage Configuration** | **Detail** |
| **Storage target** | 3 x PowerEdge R720xD servers internal drive bay with<br>24 x 2.5" NL-SAS 1 TB disks per server and<br>2 x internal drives reserved for host OS<br>Total of 72 data drives and 6 drives for OS |
| **RAID Controller** | PERC H710P,1GB DDR3, 1333 MHz |
| **Database volumes per mailbox server** | 24 Databases per server (8 Active + 16 Passive) |
| **Databases per volume** | 1 |
| **Mailboxes per database** | 313 Mailboxes |
| **Disk type** | 2.5" NL-SAS 1 TB |
| **RAID type** | RAID-0 (each drive a separate RAID-0 volume) |
| **Additional details** | Data and logs combined on each volume<br>Each HDD carries one active or passive copy in a RAID 0 with one database and log set per HDD<br>RAID stripe or element size = 256 KB or larger<br>NTFS Allocation unit size = 64KB for both DB and Log |

The objective in this configuration is to distribute smaller databases suitable for single spindles in JBOD fashion. With a higher number of database copies in the DAG, database level failover takes care of spindle failures, and no RAID configuration is required. This sample reference configuration makes use of maximum disk capacity and avoids RAID level penalties. It also considers distribution of databases, an important factor for providing faster recovery on the secondary copy.

# Conclusion

The recently launched R720xd is a well-balanced system for Exchange Solutions from a CPU, memory and internal storage perspective. The R720xd has the improved computational horsepower and an enhanced storage subsystem to support the I/O performance needed by various profiles of mailboxes. Together with the PERC H710P and 24 internal drives, the R720xd can sustain high IOPS. Jetstress test results clearly show that the R720xd with a PERC H710P significantly outperforms the previous generation storage dense standalone server, the R510 with a PERC H700, in terms of Exchange storage performance. More than twice the improvement in terms of Exchange database transactional IOPS was achieved with RAID-0, RAID-1 and RAID-10 by the R720xd.

The PowerEdge R720xd, configured with 24 x 1 TB 2.5" NL-SAS drives and a PERC H710P can support up to 130% more Microsoft Exchange 2010 mailboxes with a mailbox profile of an average 150 messages/day when compared to a PowerEdge R510 configured with 12 x 2 TB 3.5" NL-SAS drives and a PERC H700. The PowerEdge R720xd was able to support up to 7,500 mailbox users as compared to 3,250 mailbox users on the PowerEdge R510. This empirical comparison was done using the industry standard Microsoft Exchange Server Jetstress 2010 tool. This means that the PowerEdge R720xd is a storage-dense server with a balanced I/O subsystem that can support up to 2.3 times the number of mailboxes of the same I/O profile as that with PowerEdge R510.

This paper also proposes a sample reference configuration that advocates deployment of large mailboxes and inexpensive storage systems. The high availability model takes advantage of continuously replicated multiple copies and provides fast failover as well as the ability to go back to a point-in-time to handle data corruption scenarios. A number of feature enhancements make the R720xd a better and balanced building block for Exchange 2010 solutions.

# Additional resources

- Dell Exchange Solutions
- Dell PowerEdge R510
- *Dell™ PowerEdge™ R720xd: An Energy-Efficient Datacenter Building Block for Microsoft® Exchange Server 2010*
- *A Comparative Study of Microsoft® Exchange 2010 on Dell™ PowerEdge™ R720xd with Exchange 2007 on Dell™ PowerEdge™ R510*