Introduction to Microsoft® Exchange Server 2010 Sizing

Methodologies for Exchange server 2010 deployment strategies

Global Solutions Engineering

Dell

August 2012| Rev 2.7

Contents

Tables

Figures

Introduction

Over the years, messaging systems have evolved from providing basic functionality, such as email, to providing rich collaboration features, including calendar scheduling, voicemail in your inbox, instant messaging, audio-video conferencing, web conferencing and VoIP integration. The messaging requirements and needs of today's users have also changed significantly from basic desktop email access to collaboration from anywhere with any device. With the rich integration and archiving requirements, large mailboxes have become a norm in corporate and educational institutions.

In such an evolving environment, Information Technology (IT) administrators face an increasingly complex task of administering and maintaining messaging systems. Business requirements such as security, regulatory compliance, and availability further add to complexity. With the increasing storage requirements, businesses also face the enormous challenge of implementing a suitable messaging system that meets their total cost of ownership (TCO) and return on investment (ROI) targets. The performance enhancements in Microsoft® Exchange Server 2010 are designed to support low-cost storage with a new high-availability and disaster-recovery model. Exchange Server 2010 incorporates new features that enable businesses to effectively meet their challenges, without fundamentally altering the traditional methods for end users to access email. This whitepaper highlights some of the major architectural improvements and Dell hardware sizing considerations for Exchange Server roles, and shows how Dell's innovative hardware portfolio can take advantage of Exchange server 2010 enhancements through cost-effective deployment models.

Using an example solution and requirements, this paper explains a methodology for designing and sizing a solution for Exchange Server 2010. It also describes how to select and map the latest Dell servers and storage to the results obtained while applying the methodology. Later sections of the paper explain examples of reference architectures pertaining to the Dell Exchange deployment models for Exchange Sever 2010 solutions which can be derived applying the methodology. Along with reference architecture examples, the paper discusses general best practices for architecting and sizing solutions for Exchange server.

Core architectural improvements in Exchange server 2010

Exchange server 2010 offers significant improvement over Exchange 2007 and provides a cost-effective solution to overcome challenges of supporting larger mailboxes and personal archives. A number of rich features and functionalities have been added to support business needs. A few of the core architectural changes that affect the deployment and sizing strategy are described below.

Performance improvements

Exchange server 2010 has been optimized to take advantage of the low-cost "Just a Bunch of Drives" (JBOD) storage system. Using the newly enhanced Exchange Store schema, the IOPS have been significantly reduced compared to Exchange 2007 and Exchange 2003. The database pages in data structures (B-Tree) are stored in a more contiguous fashion, as opposed to random allocation as in previous versions of Exchange. This contiguity provides a sequential pattern for reads and writes and allows more efficient I/O to and from a disk subsystem. The database page size has also been increased from 8KB (in Exchange 2007) to 32KB. Multiple database page operations are coalesced to produce a single large I/O operation. All these factors, including larger I/O and page size, have contributed towards reducing the overall IOPS in Exchange server 2010. Figure-1 and Figure-2 below show

comparisons in database IOPS between Exchange 2007 and Exchange server 2010. Lower speed and larger capacity drives, such as Near-line SAS and SATA, which are optimized for sequential than random traffic, can now be deployed in an Exchange server 2010 storage subsystem without performance issues.

Figure 1. Database Read/Write I/O comparison between Exchange 2007 and Exchange server 2010 for 1,000 Heavy Online mailboxes (conducted on identical hardware configurations)

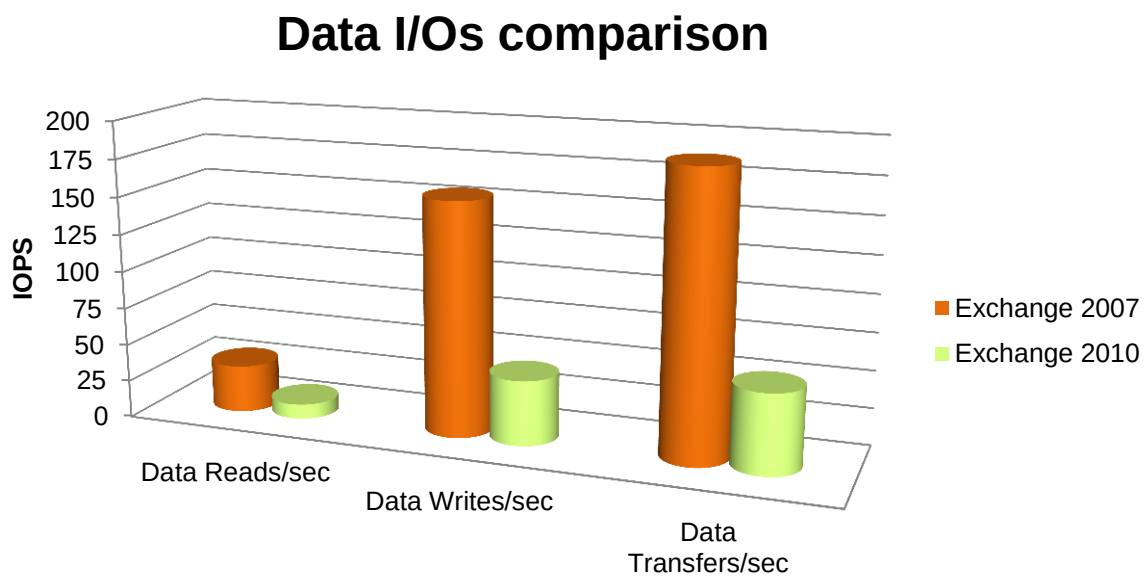

Figure 2. Database IOPS per mailbox comparison

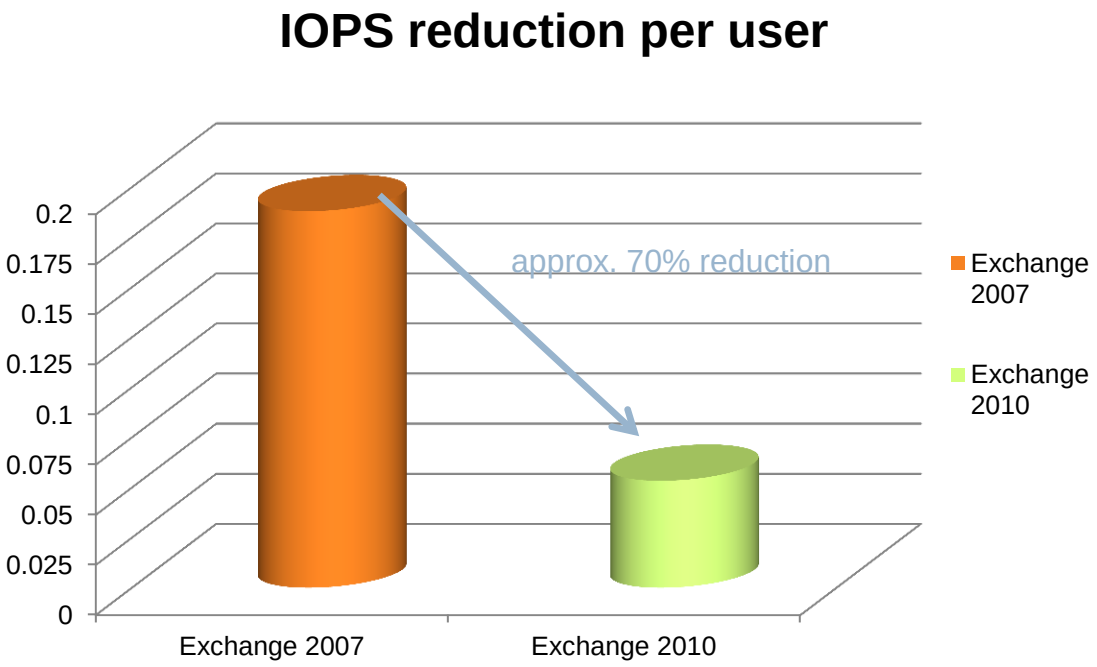

The contiguous allocation of storage data is another factor that helps maintain the reduction of I/Os. To sustain the expected level of performance, Exchange server 2010 must defragment the database as changes are made. In Exchange 2007, defragmentation is done with right merges opposite the direction of the normal database read and write I/O pattern. This methodology creates an additional performance penalty. To preserve contiguity, Exchange server 2010 has completely modified the defragmentation process to align with the left merging behavior of data reads and writes, thus optimizing performance.

Transport Resiliency

Exchange 2007 introduced *transport dumpster* for the Hub Transport Server role, which maintains a queue of messages delivered to clustered mailboxes. In the event of a database activation with lost logs, the clustered mailboxes request that all Hub servers in the Active Directory site resubmit the messages in their *transport dumpster* queues, ensuring that no in-flight messages are lost. The transport dumpster, however, only takes care of hops between Hub and Mailbox servers. Exchange server 2010 extends this high availability strategy in the transport role to include Edge Servers by introducing shadow redundancy. Shadow redundancy ensures redundancy of messages while in transit. Messages in transit remain in shadow queues until all Exchange server 2010 servers involved have acknowledged the delivery to the next hop. Shadow redundancy does not rely on the state of a transport server as long as multiple delivery paths are available, so a transport server may be brought down for upgrades because other transport servers in the topology will have a shadow copy of the messages in flight.

Continuous Availability and scalability

Microsoft Exchange server 2010 high availability leverages the cluster technologies in Continuous Cluster Replication (CCR) based on a Majority Node Set (MNS) model of Microsoft Cluster Server (MSCS). A Data Availability Group, or "DAG," as it is referred to in Exchange server 2010, is a group of 2 or more (up to 16) servers hosting mailbox databases to provide database level protection and recovery.

The DAG is a vast improvement over the Exchange 2007 two-node failover CCR model. In Exchange server 2010, a mailbox database can be hosted on any of the servers within the DAG. One or more secondary copies of the same database may be replicated to other servers (for a total of 16 copies) within the DAG. A secondary database copy gets activated in the event that the primary copy fails, allowing for more granular control over failure scenarios.

With Exchange 2007, if a database hosted on a mailbox server fails, it forces the failover of the entire mailbox server, resulting in downtime for all users on that server. In the DAG model provided by Exchange server 2010, only users attached to the database that fails are affected. Exchange server 2010 also provides incremental deployment capability, in which mailbox servers may be added to the DAG after Exchange is installed, and as the needs of the organization grow.

Extending the boundaries of high availability, a DAG can span multiple sites to provide site resiliency. So the term "Standby Cluster Replication" (SCR) in Exchange 2007 has been deprecated in Exchange server 2010 as all database copies can become active automatically within the DAG. Central to workings of an Exchange server 2010 DAG is a new component named "Active Manager." It runs on all servers within a DAG and takes either a primary or secondary role, depending on which server it is running on. The Microsoft Exchange Replication service monitoring the Extensible Storage Engine (ESE) on each server reports faults or errors to Active Manager. Active Manager initiates a switchover in the event of a database fault or failover resulting from a server component failure. The Active Manger

client on the CAS Server gets configuration changes from Active Manager so that clients are properly connected to the active copy once a switchover or a failover is complete.

Some major architectural differences between Exchange 2007 and Exchange server 2010 are summarized in Table-1.

Table 1. Exchange 2007 and Exchange server 2010 core design comparison

	Exchange 2007	Exchange server 2010
Exchange IOPS	many random IOPS	Smaller number of large sequential and coalesced IOPS
ESE page size	8KB	32KB
High availability type	LCR, CCR, SCR	DAG – enhanced availability with multiple copies extended to remote sites
Number of database copies	Max = 2	2 to 16 copies Different lag times can be set for HA or DR setup
Failover/switchover granularity	Server-level only	Server level and database level
Transport state	Transport database is stateful and causes loss of data with loss of service	Shadow redundancy – providing guaranteed delivery.
Client Access Server role (CAS)	CAS is not a middle-tier. Messaging API (MAPI) and WebDAV clients connect directly to Mailbox Server and impacted for longer time during Mailbox failure.	CAS has been modified to make a middle tier solution. All end user connects are through CAS server, providing a better user experience. CAS has also removed number of connections limitation on store.

Microsoft Exchange Server 2010 Roles

Exchange server 2010 distributed its functionality through the same five server roles used in Exchange 2007: Mailbox Server, Client Access Server, Hub Transport Server, Edge Transport Server and Unified Messaging Server. With the new architectural changes and functional design, the sizing guidelines have significantly changed for the core services. The sizing of Exchange server 2010 server roles and database storage design depend on the type of configuration and high-availability requirements.

Mailbox Server Role

The Mailbox server role provides the functionality for hosting user mailboxes, public folders, and calendar data. Exchange server 2010 introduces the concept of online archiving as a secondary mailbox to eliminate the PST allocation issues. The advanced Exchange search operation across mailbox data and the entire organization provides a better user experience for large mailboxes.

The memory and CPU sizing requirements for a Mailbox Server role depend on the number of mailboxes, type of load profile and high-availability (DAG) configuration. The Hub Transport and Client Access roles can sometimes be tied in to the Mailbox Server role in a larger distribution, for greater availability and cost benefit factors. Running multiple Exchange server roles on a single server consumes some processor and memory utilization for Hub/CAS resources in addition to the Mailbox requirements. Some sample user profiles that are also used in the Dell Exchange server 2010 Advisor tool are defined below:

- Light user: 10 send and 40 receive email operations per 8 hour work day (50 messages/day)
- Medium user: 20 send and 80 receive email operations per 8 hour work day (100 messages/day)
- Heavy user: 30 send and 120 receive email operations per 8 hour work day (150 messages/day)
- Very Heavy user: 40 send and 160 receive email operations per 8 hour work day (200 messages/day)

Client Access Server Role

The Client Access server role hosts services and functions required for supporting a variety of email clients. It hosts functionality for supporting Outlook Web App, Exchange ActiveSync client access, POP3/IMAP4 client access, Outlook Anywhere (Remote Procedure Call (RPC) over HTTP) access and new RPC service for MAPI clients. The RPC service allows internal Outlook clients to connect to a CAS server instead of directly connecting to a Mailbox Server. This results in a better client experience and negligible delay during Mailbox Server failover. Due to the connectivity through the CAS role, the number of RPC concurrent connections per Mailbox server has also increased, and can be scaled out based on the number of CAS servers in the network configuration. The CAS role also hosts Exchange Web service and auto discovery services with some added functionality enhancements in Exchange server 2010. For access to public folders, Outlook clients still connect directly to the mailbox server.

The sizing of Client Access Server in Exchange server 2010 plays an important part in hosting all external services in addition to off-loading significant processing from Mailbox server. The results have shown that the Exchange server 2010 CAS Server role has a high CPU requirement, which is directly related to the number of internal/ external connections and mail\box user load. Therefore, at least three-quarters of the Mailbox CPU resources should be allocated for the CPU resources of the CAS server. Dell PowerEdge R620 or Dell PowerEdge R720xd is a suitable server platform for hosting Exchange CAS role, and should be properly sized to handle the required Client access functions. The CAS Server requirements should be linearly scaled out with the number of Mailbox servers and cores in the configuration. A hardware load-balancer is deployed with more than one CAS server to effectively load-balance the connections and provide high availability across the configuration.

Hub Transport Server Role

The Hub Transport server role is responsible for handling all internal mail flow across Exchange roles and appropriate delivery of user messages. All incoming external SMTP email is forwarded to the Hub Transport by the Edge Transport Server role, or any messaging appliance. The Hub Transport Server role provides functionality such as journaling, server side rules, and an additional layer of antivirus/anti-spam security. Hub server must be deployed in each Active directory site containing a Mailbox server role. Because all messages in the Exchange server 2010 Organization are routed to and processed by Hub Transport, it enables enforcing messaging compliance policies at this server role.

The Hub Transport role does not have high processing and memory requirements compared to the Client Access Server role. The CPU utilization requirements tend to increase with the message size, message rate processing and antivirus configuration on the Hub role. The general recommendation is to have a Hub and Mailbox core ratio of 1:5 with antivirus configuration and 1:7 without antivirus scanning on Hub server. A minimum of two Hub Transport Servers should be deployed to satisfy high-availability requirements. When deployed independently as a separate role, the Dell PowerEdge R420 or Dell PowerEdge R320 servers are suitable server platforms for hosting the Exchange Hub Transport server role. The hardware platform should be appropriately sized for the required Hub Transport functions. Specifically, ample system resources, including processors and memory, should be allocated if Edge functions such as antivirus/anti-spam and rule based policies are configured on the Hub Server role.

It is typical to host the Hub and CAS server together on a single server. If this is done, the ratio of the combined Hub/CAS to Mailbox Server role should be 1:1. That is for every mailbox server you should have a Hub/CAS combined server.

Edge Transport Server Role

The Edge Transport server role sits in the perimeter network and provides functionality for routing external SMTP mail flow to and from the Internet. It provides antivirus/anti-spam protection, filtering capabilities, and rule based protection. The Edge Transport role is not a member of the Active Directory domain and sits outside the Active Directory framework. It uses the Active Directory Lightweight Directory Services and Edge Sync to obtain, store and synchronize the required Active Directory information for performing its functions.

The sizing of Edge Transport Server depends on the number of accepted SMTP connections and the message rate. Even though the Edge role does not have any significant CPU and memory requirements, ample resources should be available to handle the overhead associated with antivirus scanning of connections and accepted messages. Dell PowerEdge R620 or PowerEdge R720xd is a suitable platform for hosting Edge Transport role services. A minimum of two servers should be deployed to handle any unexpected failures and server down-times.

Unified Messaging Server Role

The Unified Messaging (UM) server role was introduced with Exchange 2007 to provide the functionality for integration of voice mail, along with email, into user mailboxes. Outlook Voice Access and Auto-attendant were also part of the UM functionality. Exchange server 2010 Unified Messaging Server is built on the same model with additional feature enhancements. Features like Voice Mail preview and protection provide added benefits and advantages. The voicemail preview feature enables speech-to-text format to quickly view voicemail messages like an email, without listening to all voicemails. The

voicemail protection feature uses Right Management Services to set restriction policies and ensure authorized distribution of voice mails.

The Unified Messaging Server resides between the core Exchange architecture and the Voice-over-IP (VoIP) gateway infrastructure. The VoIP gateway provides integration through IP-PBX to the Public Switching Telephony Network (PSTN). The Unified Messaging Server requires CPU processing power to handle voice connection requests for voicemail and other features, and compresses it in Windows supported format, using audio codecs. The PowerEdge R620 or the PowerEdge R720xd is a suitable rack server for hosting the UM role with a minimum of 2 to 4 Processor cores for mid-sized configuration or 8 processing cores for large enterprises.

Multi-role Exchange Server configuration

The Mailbox, Hub Transport and Client Access server roles can be collocated on a single physical or virtual server. When all these server roles are installed on a single server, the installation is referred to as multi-role Exchange Server. In a virtual environment, the Hub Transport and Client Access server roles only are typically collocated while it is recommended to have Mailbox server role installed in a virtual machine.

Exchange server 2010 sizing methodology

Each Exchange server 2010 Server role has distinct system requirements and should be sized with attention to role specific demands. Exchange server is a storage intensive workload and allows various storage options to be considered ranging from internal server storage to shared storage like Storage Area Network (SAN). The two significant sizing considerations are server and storage. A server forms a building block of the solution and needs to be sized appropriately. Server sizing considerations include deciding the type of processor that is best suited for handling the Exchange Mailbox profile requirements. It also involves deciding the size of memory required and intelligent allocation of the DIMMs to the processor memory channels so as to take advantage of full memory bandwidth. Server sizing also include selecting the right host network adapters. In case of storage sizing, selecting the right type of storage is crucial to achieve a balance between the cost and the performance of the solution. Storage sizing typically involves deciding the type of RAID, type of disks, number of disks both from an IOPS and capacity perspective and intelligently mapping Exchange databases to the storage subsystem per the solution requirements.

This section describes a step wise approach towards building Exchange Server 2010 solution ready for deploying on Dell servers and storage. The reader will go through a detailed exercise while working out the solution for an example input configuration.

Gathering requirements

Before we begin using the methodology to design the solution, some information needs to be collected. The user needs to determine mailbox profile requirements as accurately as possible because these requirements may impact all other components of the design. If Exchange is new to you, you may have to make some educated guesses. If you have an existing Exchange environment, you can use the Microsoft Exchange Server Profile Analyzer tool to assist with gathering most of this information. This information is needed and assumed available for the exercise of building the example Exchange solution, described in this section. Typically the customer requirements include the configurations related to number of mailboxes, size of the mailboxes and mailbox profile. The mailbox profile

includes the average message size and the IOPS requirements for the mailbox users. Certain mobile devices can generate additional IOPS on top of the selected user profile based on their mechanism of accessing the Exchange mailbox. Customers may want to consider these additional IOPS accounting for such mobile devices accessing the user's mailboxes. Occasionally, customer may specify the geographic location requirements. These include geographic distribution of the mailbox users and datacenters. Tables 2, 3 and 4 list the mailbox specific requirements while Table-5 lists the geographic distribution requirements.

Table 2. Mailbox count requirements

Mailbox count requirements	Value
Mailbox count (total number of mailboxes including resource mailboxes)	4000
Projected growth percent (%) in mailbox count (projected increase in mailbox count over the life of the solution)	0%
Expected mailbox concurrency % (maximum number of active mailboxes at any time)	100%

Table 3. Mailbox size requirements

Mailbox size requirements	**Value**
Average mailbox size in megabytes (MB)	1024 MB
Average mailbox archive size in MB	0 MB
Projected growth (%) in mailbox size (projected increase in mailbox size over the life of the solution)	0%
Target average mailbox size in MB	1024 MB

Table 4. Mailbox profile requirements

Mailbox profile requirements	**Value**
Target message profile (average total number of messages sent plus received per user per day)	150
Target average message size in KB	75 KB

Table 5. Geographic distribution requirements

Mailbox user site requirements	**Value**
Number of major sites containing mailbox users	1
Number of mailbox users in site 1	4000

The information in Tables 2, 3, 4 and 5 can be summarized as shown in Table-6. Henceforth Table-6 is referred to for developing the solution.

Table 6. Example solution input

Number of Mailboxes	4000
Average user I/O profile (messages/day)	0.18 IOPS (~ 150 messages per day)
Average mailbox size limit	Up to 1GB
Total active/passive copies per database	2

To summarize, a solution is required for 4000 mailbox users with heavy mailbox I/O profile and up to 1GB mailbox size. The solution demands redundancy through copies. For the scope of the exercise we do not consider backup and recovery options, disaster recovery for site resiliency and additional servers for non-mailbox server roles. However we provide guidelines to address these parameters in the solution later in the exercise.

Determining high availability strategy

When designing an Exchange server 2010 environment, many design decision points for high availability strategies impact other design components. It is recommended that you determine your high availability strategy as the first step in the design process. This section helps you to design the DAG configuration. It is highly recommended that you review the following information prior to starting this step:

- Understanding high availability factors.
- Planning for high availability and site resilience.
- Understanding backup, restore and disaster recovery.

Determine whether site resiliency is required

If you have more than one datacenter, you must decide whether to deploy Exchange infrastructure in a single datacenter or distribute it across two or more datacenters. The organization's recovery Service Level Agreements (SLAs) should define what level of service is required following a primary datacenter failure. This information should form the basis for this decision.

In this example solution, the office is located in a single geographic location, and the server infrastructure is located on the premises. There's no budget to maintain infrastructure in a second geographic location, so a site resilient deployment can't be justified. The Exchange server 2010 design will be based on a single site deployment with no site resiliency.

Determine backup and database resiliency strategy

Exchange server 2010 includes several new features and core changes that, when deployed and configured correctly, can provide native data protection that eliminates the need to make traditional data backups. Backups are traditionally used for disaster recovery, recovery of accidentally deleted items, long term data storage, and point-in-time database recovery. Exchange server 2010 can address all of these scenarios without the need for traditional backups:

- **Disaster recovery** - In the event of a hardware or software failure, multiple database copies in a DAG enable high availability with fast failover and no data loss. DAGs can be extended to multiple sites and can provide resilience against datacenter failures.
- **Recovery of accidentally deleted items** - With the new Recoverable Items folder in Exchange server 2010 and the hold policy that can be applied to it, it's possible to retain all deleted and modified data for a specified period of time, so recovery of these items is easier and faster. For more information, see Messaging Policy and Compliance, Understanding Recoverable Items, and Understanding Retention Tags and Retention Policies.
- **Long-term data storage** - Sometimes, backups also serve an archival purpose. Typically, tape is used to preserve point-in-time snapshots of data for extended periods of time as governed by compliance requirements. The new archiving, multiple mailbox search, and message retention features in Exchange server 2010 provide a mechanism to efficiently preserve data in an end-user

accessible manner for extended periods of time. For more information, see Understanding Personal Archives, Understanding Multi-Mailbox Search, and Understanding Retention Tags and Retention Policies.

- **Point-in-time database snapshot** - If a past point-in-time copy of mailbox data is a requirement for your organization, Exchange provides the ability to create a lagged copy in a DAG environment. This can be useful in the rare event that there's a logical corruption that replicates across the databases in the DAG, resulting in a need to return to a previous point in time. It may also be useful if an administrator accidentally deletes mailboxes or user data.

There are technical reasons and several issues that you should consider before using the features built into Exchange server 2010 as a replacement for traditional backups. Prior to making this decision, see Understanding Backup, Restore and Disaster Recovery.

In this example, maintaining tape backups has been difficult and testing and validating restore procedures hasn't occurred on a regular basis. Using Exchange native data protection through DAG in place of traditional backups as a database resiliency strategy would be an improvement. However, due to a limited budget, the two Mailbox servers are required to support a minimum of two database copies. The customer is strongly urged to consider implementing a backup solution.

Determine backup solution

The example solution in consideration does not prescribe any particular backup implementation. In order to design backup solution, refer to the 'Considerations for backup' section. The example solution here, however, considers space for restore LUNs while sizing the storage.

Determine number of database copies required

The next important decision when defining your database resiliency strategy is to determine the number of database copies to deploy. We strongly recommend deploying a minimum of three copies of a mailbox database on your local site before eliminating traditional forms of protection for the database, such as Redundant Array of Independent Disks (RAID) or traditional VSS-based backups.

For this example a decision is made to deploy two copies of each database. This enables high availability during a single server failure or maintenance event and provides a basic level of protection at a reasonable cost. Since you only have two copies of the data, a RAID configuration for disk drives has been selected to protect the storage from a single disk failure.

Determine number of Mailbox servers per DAG

In this step, you need to determine the minimum number of Mailbox servers required to support the DAG design. This number may be different from the number of servers required to support the workload, so the final decision on the number of servers is made in a later step.

Minimum of two mailbox servers are required to support a DAG configuration with two database copies.

Once the design decision points have been considered and finalized, we begin building the solution by calculating the detailed requirements for CPU, memory and storage.

Mailbox role megacycles calculations

For identifying the right processor, the Exchange server 2010 megacycles requirement must be calculated and compared with the megacycles of the processor offered by the server. Table-7 shows a list of megacycles required for active and passive copies based on the mailbox I/O requirement.

Table 7. Megacycle estimates[1]

Messages sent or received per mailbox per day	Megacycles per mailbox for active mailbox database	Megacycles per mailbox for remote passive mailbox database	Megacycles per mailbox for local passive mailbox
50	1	0.1	0.15
100	2	0.2	0.3
150	3	0.3	0.45
200	4	0.4	0.6

Continuing with the previous example of 4000 heavy users and assuming one active and one local passive copy, the total megacycles required as referenced from Table-7 would be calculated as follows:

Total Megacycles required = total active mailbox megacycles required + total local passive megacycles required + total remote passive megacycles required

Microsoft recommends increasing the megacycles per active mailbox by ten percent for each additional database copy after the one active copy. The example in consideration has two copy DAG; implying one additional copy. While calculating the total megacycles we consider the worst case scenario where 100% of the active mailbox copies are hosted on a single server.

Total active mailbox megacycles required = [Total number of active mailboxes x Megacycles per mailbox for active mailbox database (shown in Table-7)] x [1 + (number of additional copies x 10 / 100)]

Total megacycles required = [(4000 x 3) x 1.1 for increase of 10% for one additional copy] + 0 + 0

= 13200

Because the megacycle requirements are based on a baseline server and processor model (HP DL380 G5 x5470 3.33 GHz - 8 cores), the available megacycles for the target server need to be adjusted against the baseline. To do this, independent performance benchmarks maintained by Standard Performance Evaluation Corporation (SPEC) are used. The megacycles calculated in this section are used while mapping a Dell server from a processing capacity perspective.

[1] http://technet.microsoft.com/en-us/library/ee712771.aspx

Mailbox role memory calculations

Sizing memory correctly is an important step in designing a healthy Exchange environment. The Extensible Storage Engine (ESE) uses a database cache to reduce I/O operations. In general, the more database cache available, the less I/O generated on an Exchange server 2010 mailbox server. However, there's a point at which adding database caches no longer results in a significant reduction in IOPS. Therefore, adding large amounts of physical memory to your Exchange server without determining the optimal amount of database cache required may result in higher costs with minimal performance benefit.

Calculate required database cache

Determining Mailbox memory involves first computing the optimal amount of database cache based on the Table-8.

Table 8. Database cache per user[2]

Messages sent or received per mailbox per day (about 75 KB average message size)	Database cache per user (MB)
50	3 MB
100	6 MB
150	9 MB
200	12 MB

Determine total memory requirements per mailbox or multi-role server for the worst-case failure scenario

For this solution, the database cache required per user is 9 MB. Once the database cache required for the environment has been calculated, the next step involves calculating the minimum amount of memory required per mailbox server for the worst case failure scenario. In this example, assuming 4,000 mailboxes active on a single server, the memory requirement would be calculated as follows:

Total database cache required per server = Number of active mailboxes per server x average database cache per Mailbox

Total database cache required per server = (4000 x (9/1024))

= 36 GB

[2] http://technet.microsoft.com/en-us/library/ee832793.aspx

Once the total database cache required per server is computed, the following table is referenced to determine the recommended total physical server memory.

Table 9. Memory Requirements[3]

Server physical memory (RAM)	Database cache size (Mailbox role only)	Database cache size: Multiple-role (for example, Mailbox + Hub Transport)
2GB	512 MB	Not supported
4GB	1 GB	Not supported
8GB	3.6 GB	2 GB
16GB	10.4 GB	8 GB
24GB	17.6 GB	14 GB
32GB	24.4 GB	20 GB
48GB	39.2 GB	32 GB
64GB	53.6 GB	44 GB
96GB	82.4 GB	68 GB

Using Table-9, the recommended physical memory configuration to support 36 GB of database cache would be **48 GB** for a mailbox server role and **64 GB** for a multi-role server. Thus, in a case of a two-copy DAG configuration with two servers, each server will require 64 GB of memory if configured as a multi-role Exchange server.

Note: in the above example it is assumed that there is a single profile of 4,000 heavy mailbox users. In the event if there are multiple tiers of mailbox users with each tier having different mailbox profile, then each tier will have to have its database cache required and should be calculated separately. The average database cache per mailbox would then be calculated by adding up the database cache required across all the tiers and dividing it by the total number of mailbox users.

Mailbox role storage capacity calculations

Many factors influence the storage capacity requirements for the Mailbox server role. For additional information, we recommend that you review Understanding Mailbox Database and Log Capacity Factors. The following steps outline how to calculate mailbox capacity requirements. These requirements will then be used to make decisions about which storage solution options meet the capacity requirements.

Calculate mailbox size on disk

Before attempting to determine what your total storage requirements are, you should know what the mailbox size on disk will be. A full mailbox with a 1GB quota requires more than 1 GB of disk space because you have to account for the prohibit send/receive limit, the number of messages the user sends or receives per day, the Deleted Items folder retention window (with or without calendar version logging and single item recovery enabled.) You can use this information to do the calculations manually. The following calculations are used to determine the mailbox size on disk for single tier scenario with a total of 4000 mailboxes.

[3] http://technet.microsoft.com/en-us/library/ee832793.aspx

Mailbox profile details:
1GB mailbox quota, 150 messages per day message profile, 75 KB average message size. Refer Table 6 for more details.

Whitespace = Number of messages per day x Average message size

Whitespace = 150 messages per day × 75 ÷ 1024 MB

= 11 MB

Dumpster Size = (Daily Incoming/Outgoing email x Average Message Size x Deleted Item Retention Window) + (Mailbox quota Size x 0.012) + (Mailbox quota Size x 0.03)

When single item recovery is enabled, there is an additional 1.2 percent increase in the size of the mailbox for a 14-day deleted item retention window and for calendar version logging data, there is an additional 3 percent increase in the size of the mailbox.

Dumpster size = (150 messages per day × 75 ÷ 1024 MB × 14 days) + (1024 MB × 0.012) + (1024 MB × 0.03)

= 197 MB

Mailbox size on disk = Mailbox limit + Whitespace + Dumpster

Mailbox size on disk = 1024 MB + 11 MB + 197 MB

= 1232 MB

Average size on disk = 1232 MB.

If the customer requirements state 'N' tiers then the mailbox size on the disk should be computed for each tier and the average mailbox size on disk should then be calculated as follows:

Average mailbox size on disk =
[(Mailbox size on disk for tier 1 × number of mailbox users in tier 1) + (Mailbox size on disk for tier 2 × number of mailbox users in tier 2) +…+(Mailbox size on disk for tier N × number of mailbox users in tier N)] ÷ (Total number of mailbox users)

Calculate total database storage capacity required

In this step, the high level storage capacity required for all mailbox databases is determined. The calculated capacity includes database size, catalog index size, and 20 percent free space.
To determine the storage capacity required for all databases, use the following formulas:

Mailbox tier details: 1024 MB mailbox quota, 150 messages per day message profile, 75 KB average message size.

Total Database capacity required = (number of mailboxes × mailbox size on disk × database overhead growth factor) × (20% data overhead)

Total Database capacity required = (4000 × 1232 × 1) × 1.2

= (5913600 /1024) GB

= **5775 GB**

Database index size = 10% of total database capacity

Database index size = **578 GB**

The total capacity of the volume hosting the database is calculated as the size required for database and the size required by the database index.

Total database and index capacity required = total Database capacity + index size

= 5775 + 578

Total database and index capacity required = 6353 GB

It is recommended that the drives should not be completely filled with data. Hence as a best practice it is suggested that a 20% free space is added to the database volume capacity.

Total volume capacity required for database and index = (Total database and index capacity required ÷ 0.80) to add 20% volume free space

Total database capacity = 6353 ÷ 0.8

= **7942 GB**

Calculate transaction log storage capacity requirements

To ensure that the Mailbox server doesn't experience any outages as a result of space allocation issues, the transaction logs also need to be sized to accommodate all of the logs that will be generated during the backup set. Provided that this architecture is leveraging the mailbox resiliency and single item recovery features as the backup architecture, the log capacity should allocate for three times the daily log generation rate in the event that a failed copy isn't repaired for three days. (Any failed copy prevents log truncation from occurring.) In the event that the server isn't back online within three days, you would want to temporarily remove the copy to allow truncation to occur.
To determine the storage capacity required for all transaction logs, use the following formulas:

Mailbox tier details: 1024 MB mailbox quota, 150 messages per day message profile, 75 KB average message size.

Log files size[4] = (log file size × number of logs per mailbox per day× number of days required to replace failed infrastructure × number of mailbox users) + (1% mailbox move overhead)

The log file size is considered to be 1MB in capacity. Number of logs per mailbox per day is 30 as per the Microsoft's recommendation and the number of days to replace the failed infrastructure is considered as 3.

[4] http://technet.microsoft.com/en-us/library/ee832796.aspx

Log files size = (1 MB × 30 × 3 × 4000) + (4000 × 0.01 × 1024)

Log files size = 400960 MB = (400960 ÷ 1024) GB = **392 GB**

Similar to the case of database size, 20% free space for the volume is considered.

Total log volume capacity = log files size ÷ 0.80 to add 20% volume free space

Total log capacity = 392 ÷ 0.80 = **490 GB**

Determine total storage capacity requirements

Table-10 summarizes the high level storage capacity requirements for this solution. In a later step, you will use this information to make decisions about which storage solution to deploy.

Table 10. Mailbox Database capacity for 4000 mailboxes

Summary of storage capacity requirements Disk space requirements	Value
Average mailbox size on disk (MB)	1232
Database and index space required (GB)	7942
Log space required (GB)	490
Total space required (GB) (space for database and logs)	8432

Mailbox role IOPS calculations

Because it's one of the key transactional I/O metrics needed for adequately sizing storage, you should understand the amount of database I/O per second (IOPS) consumed by each mailbox user. Pure sequential I/O operations aren't factored in the IOPS per Mailbox server calculation because storage subsystems can handle sequential I/O much more efficiently than random I/O. These operations include background database maintenance, log transactional I/O, and log replication I/O. In this step, you calculate the total IOPS required to support all mailbox users, using the following:

To determine the IOPS profile for a different message profile, see the table "Database cache and estimated IOPS per mailbox based on message activity in Understanding Database and Log Performance Factors.

Total required IOPS = IOPS per mailbox user × number of mailboxes × I/O overhead factor

Total required IOPS = 0.15 × 4000 × 1.2

= 720

Average IOPS per mailbox = 720 ÷ 4000 = 0.18

The high level storage IOPS requirements are approximately 615. When choosing a storage solution, ensure that the solution meets this requirement.

Subsystem requirements for other Exchange server 2010 roles

Once the Mailbox role hardware requirements have been identified, we can proceed to size the hardware requirements for Hub Transport and Client Access server roles. The subsystem requirement of primary interest here are CPU and memory. Microsoft has published recommended guidance for CPU and memory sizing in Hub Transport and Client Access roles in the form of ratios in Understanding Server Role Ratios and Exchange Performance.

Table 11. Mailbox to (Hub Transport + Client Access) Processor Core Ratio

Number of Client Access and Hub Transport servers	Recommended processor core ratio
Mailbox server role: Client Access and Hub Transport combined server roles	1:1
Mailbox server role: Client Access	4:3
Mailbox server role: Client Access	7:1 (without antivirus scanning on Hub Transport) 5:1 (with antivirus scanning on Hub Transport)

Table 12. Memory configurations for Exchange servers based on installed server roles

Memory configurations for Exchange server 2010 servers based on installed server roles Exchange server 2010 server role	Minimum supported	Recommended maximum per core
Hub Transport server role	4 GB	1 GB
Client Access server role	4 GB	2 GB
Client Access and Hub Transport combined server role (Client Access and Hub Transport server roles running on the same physical server)	4 GB	2 GB

In a traditional Exchange deployment you may deploy the Client Access, Hub Transport, and Mailbox server roles on different physical servers. However, there are a number of reasons why you may want to combine the Client Access and Hub Transport roles on the same physical server or virtual machine. There are also a number of scenarios where deploying the Client Access, Hub Transport, and Mailbox server roles on the same physical server makes sense as well. For more information please see Understanding Multiple Server Role Configurations in Capacity Planning.

The impressive performance of newer processors can result in underutilized processor capacity on mailbox servers. This is especially true when selecting servers with internal disk used for Exchange databases. In most cases you will run out of available I/O before utilizing all of the processor capacity. We anticipate that we will be able to deploy the Client Access and Hub Transport roles on the same physical server (with rich processing capabilities) as the Mailbox role to utilize remaining processor capacity.

Deployment models

Exchange Server 2010 infrastructure can be designed in different ways, based on the size of the configuration and the number of high-availability/disaster recovery copies required. Customers also have a variety of server form factor, storage, and disk options. To provide the most customer value, Dell has developed three basic architecture models that balance performance, cost and power utilization factors. This section describes how to map the sizing methodology described previously to Dell hardware. For examples of specific deployment models, refer to here.

Simple Distributed Model

With its highly available functionality and non-shared storage for same database copies, Exchange Server 2010 uses standalone servers with large internal drives, or servers with direct attached storage (DAS), as viable options for mailbox servers. Such a design allows an easy incremental deployment to provide additional mailbox capacity and a higher number of copies. Additionally, multiple Exchange roles can be combined on servers in a multi-copy environment, to enable small and medium sized organizations to provide role redundancy along with providing mailbox capacity in a cost-effective manner.

The goal of the Simple Distributed model is to design a server and storage infrastructure optimized and tuned to provide Exchange server 2010 in a simple, highly available, and cost-effective manner; in this model, RAID is optional, depending on the number of database copies in a DAG. The database-level failover available in Exchange server 2010 with a higher number of copies can address spindle failure, and eliminate the need for RAID configuration. RAID 1 based configurations fit this model as well, where a single database can be mapped to its corresponding log and to a RAID 1 volume. Therefore, in the event of a RAID failure multiple databases or large number of mailbox users are not impacted. Storage configurations with RAID can become cost prohibitive when creating more than two local copy DAG solutions. Thus, depending on the number of local copies required, a customer has the option use storage solutions with or without RAID. When selecting a configuration without RAID, Dell recommends using at least three local copies of a mailbox database in DAG with either internal storage or direct-attached storage (such as PowerVault™ MD1200 or PowerVault MD1220). There are two options for building the simple distributed model: Internal Storage Distributed Model and External Storage Distributed Model.

Internal storage distributed model

This model makes use of the internal disks on a dense server as a building block to host the mailbox database and logs. Along with hosting the mailbox, this server can optionally host multiple roles, such as Hub and CAS in a multi-role Exchange server configuration. Growth is accomplished by scaling out the building block configuration, and expanding the Exchange deployment per requirements. This option typically meets the needs of small to medium organizations, or satellite offices requiring a cost-effective messaging solution with high availability, ease of management, and the ability to expand. Given that large capacity of NL-SAS/SATA drives and easy scale-out options are readily available, this model can also be a good fit for larger organizations that deploy smaller mailbox sizes in a dedicated Exchange environment, independent of their SAN.

In this section, the sample solution shown in Table-13 illustrates how the results obtained using the sizing methodology can be utilized to choose the Dell hardware that best suits the solution. The model includes 4000 heavy I/O profile (~150 messages per day) users, with up to 1GB per mailbox in a 2 copy DAG configuration.

Table 13. Solution Requirements

Number of Mailboxes	4000
Average user I/O profile (messages/day)	0.18 IOPS (~ 150 messages per day)
Average mailbox size limit	Up to 1GB
Total active/passive copies per database	2
Not included in this solution	Backup and recovery infrastructure Disaster recovery or site resiliency Additional server for Hub/CAS UM and Edge roles

We begin with designing the DAG configuration. The solution requires two copies of the each database, indicating a two copy DAG; at least two members (servers) are required to host a two copy DAG. We start sizing the solution with the minimum number of servers to satisfy the solution requirements, starting with two PowerEdge servers. In order to select specific servers, information on the suitable processor family type and memory configuration is needed.

The target server is a PowerEdge R720xd. 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" front-loading internal drives, in addition to optional 2 x 2.5" back-accessible drives. With these options, the R720xd offers sufficient disk IOPS and storage capacity to host a large number of Exchange server 2010 mailboxes. Improved processor performance and large server memory capacity allow IT architects to design flexible Exchange server 2010 solutions. For example, the R720xd can 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 considerable Exchange users per system. Table-14 shows an overview of the R720xd features.

Table 14. PowerEdge R720xd features

Parameters	PowerEdge R720xd details
Rack Units	2U
CPU	Intel® Xeon® E5-2600 series processors
Memory slots	Up to 24 x DDR3 DIMM slots (up to 768 GB)
Internal storage	24 x 2.5" SAS, NL-SAS, SATA drives 12 x 3.5" SAS, NL-SAS, SATA drives Plus optional 2 x 2.5" rear-accessible drive bay
PCIe support	3.0
LOM	Multiple NDC options: minimum 4 x 1 GbE
PERC support	PERC 8 series

Mapping processor requirements

We begin by identifying the processor requirements and assume that our target server is a Dell PowerEdge R720xd with dual E5-2630 2.3 GHz processors, for a total of 12 cores in the server. For this server configuration, we reference the SPEC website to get the SPECint_rate2006 results. The results obtained from SPEC website were as follows:

SPECint_rate2006 value = 434

SPECint_rate2006 value per processor core = 434 ÷ 12 = 36.167

The megacycles available on the target system (R720xd) can be calculated by using the formula:

Megacycles Available per core = (R720xd core value ÷ baseline core value) x Hertz per core of baseline

= (R720xd core value ÷ DL380 core value) x Hertz per core of DL380

= [(434 ÷ 12) ÷ (150 ÷ 8)] x 3330

= [36.167 ÷ 18.75] x 3330

= 6424

Megacycles per R720xd (Total) = R720xd megacycles Available per core x Number of cores per R720xd

= 6424 x 12

= 77078

As recommended by Microsoft guidelines, for solutions leveraging mailbox resiliency, the design assumptions states that, for standalone servers with only the mailbox role, the recommendation is to not exceed 70% CPU utilization during peak period. For solutions leveraging mailbox resiliency, it is recommended that the server configuration not exceed 80% CPU utilization after a single or double member server failure with the assumption that we want each server to be a multi-role server; the adjusted available megacycles per server are calculated as follows:

Adjusted available megacycles per R720xd server = 77078 x 0.80

= 61663 (>13200)

For multi-role servers, it is recommend that 50% of the adjusted available megacycles be allocated to the mailbox server role and the remaining 50% be allocated to Client Access and Hub Transport as per the processor core ratio defined for mailbox server role to the combined Hub/CAS role. Refer 'Subsystem requirements for other Exchange server 2010 roles' section for more details.

Megacycles allocated to the mailbox role in a multi-role Exchange server configuration = 50% of adjusted available megacycles

= 61663 x 0.50

= 30832 (>13200)

At 80% CPU utilization, the megacycles value computed for the mailbox role in a multi-role Exchange server configuration exceeds the overall megacycles requirement of 13200 as calculated in the 'Mailbox role megacycles calculations' section. The CPU requirements are sufficient and one PowerEdge R720xd can accommodate 4000 heavy profile users communicating 150 messages per day.

In this case, it is determined that 4000 heavy profile users can be accommodated on a single PowerEdge R720xd in a Multi-Role configuration. Thus, two PowerEdge R720xd servers with dual E5-2630 processors satisfy the CPU requirements for our 2 copy DAG example solution.

Mapping memory requirements

The memory requirement is calculated using the methodology outlined in the 'Mailbox role memory calculations' section. A total of 64 GB of memory is sufficient per server to host the multi-role Exchange configuration with all 4,000 mailboxes active. The PowerEdge R720xd is a two socket server with EP processors; it has four channels per socket, with three slots per channel. It is recommended to have at least one DIMM per channel, to maximize the memory bandwidth. Therefore, one of the following memory configurations should be sufficient for the solution requirements:

- 16 x 4 GB DIMMs: 2 DIMMs per channel.
- 8 x 8 GB DIMMs: 1 DIMM per channel.

Mapping storage requirements

PowerEdge R720xd is a storage dense server that provides multiple options for internal storage. It can host up to 24x 2.5" drives or up to 12 x 3.5" drives. We consider R720xd as a server plus storage combination option that can host the multi-role Exchange configuration; we can then evaluate its capability to host 4000 mailboxes in a two copy DAG. We validated that the selected building block is capable of meeting the IOPS requirements. This section describes how the results obtained for the storage and IOPS calculations can be mapped to the Dell PowerEdge R720xd.

Mapping storage capacity calculations

DAG provides a replication-based recovery mechanism at the database level. In order to have resiliency at the disk level, RAID-based replication is required. RAID-1 configuration is selected for each storage volume. A storage capacity requirement of about 8.5 TB is derived using the information in section 'Mailbox role storage capacity calculations'. This storage capacity requirement includes the space for the database, indexing, and log and required free space.

In order to simplify the storage mapping calculations, we use a capacity requirement of 10 TB per copy in the DAG. Since we have a two copy, two-member DAG we need additional storage; therefore the storage has to be sized with a total capacity of 20 TB. Considering that we have a RAID-1 configuration and 1 TB NL-SAS drives, at least 40x 1 TB NL-SAS drives are required for the storage capacity. We have used two R720xd servers in the two-copy DAG, and each one can host 24x 1 TB NL-SAS drives; therefore a total of forty eight drives are available and that satisfies our storage capacity requirements. The R720xd can satisfy the storage requirements, and can be used as server plus storage combination for our solution.

Mapping the IOPS capacity requirements

First, we determine the IOPS required based on the number of active mailboxes and the profiles as explained in 'Mailbox role IOPS calculations' section. The Jetstress 2010 mailbox profile test is performed to verify if the given storage size meets the mailbox profile requirements. The Jetstress

Mailbox Profile Test allows you to verify whether the storage system meets or exceeds the planned Exchange mailbox profile; in other words, it can test whether the Exchange profile is configured using the requirements of the underlying storage system. Table-15 lists the mailbox profile, and the test configuration obtained for the storage configuration and Table-16 describes the results obtained from the Jetstress mailbox profile test.

Table 15. Mailbox profile for Jetstress mailbox profile tests

Servers	Number of Mailboxes	Mailbox size (MB)	Target IOPS	Spindles	RAID configuration	Number of active Exchange databases	I/O profile
R720xd	4,000	1024	720	24 x 2.5" 1TB	RAID-1	10 (1 DB per RAID volume)	150 messages a day

Table 16. Jetstress mailbox profile test results

	Target	Achieved
Transactional IOPS	720	1094
Database Read Latency	< 20 milliseconds	~13 milliseconds
Log Write Latency	< 10 milliseconds	~2.5 milliseconds

The Jetstress test yielded **1094** achieved transactional I/O per second, satisfying the recommended latency requirements. R720xd exceeds the storage IOPS requirements, and proves that the internal storage is capable of hosting the Exchange database configuration under consideration. The additional headroom in the IOPS indicates that there is a possibility of hosting more mailboxes. However because from a capacity perspective we don't have much additional free space available, any additional mailboxes added would have been at the expense of reducing the mailbox size. Additionally, as a best practice it is also recommended to account for a RAID rebuild factor while calculating the required IOPS. From the results we seem to have an additional 50% IOPS available, which seems to provide more than sufficient headroom to account for any performance impact due to a RAID rebuild.

One logical storage volume is 1TB RAID-1 volume using two 1 TB NL-SAS physical drives; there are 20 volumes distributed over the two R720xd servers. Each volume is capable of hosting a database, as well as the logs. To simplify the Exchange database mappings, we allocated one database per volume resulting in a total of 20 databases; of these databases, 10 are active copies and 10 are passive copies. The active and passive database copies are distributed over two R720xd servers, so that each server hosts 5 active and 5 passive database copies. For storage capacity, it is recommended to include additional drives as hot spare volumes and restore LUNs. After determining the number of databases and the drives required, we have four spindles per server that can be utilized for hot spare volumes and restore LUNs. Two drives with RAID-0 configuration can be used for Restore LUNs, and two hot spare drives are allocated per server in the solution.

From a capacity perspective, we can safely allocate 400 mailboxes of 1.3 GB size each to one database. One R720xd is capable of supporting 4000 mailbox users with heavy mailbox I/O profile. As a result we distribute 2000 active and 2000 passive users per server in two-copy two-member DAG. In worst case scenario, if one server in the DAG experiences failure, the other server can host 4000 active users.

Based on the calculations described in the sizing methodology for various server roles, the solution is described in Table-17.

Table 17. Recommended Solution

Server Configurations	Detail
Multi-role (Mailbox/Hub/CAS) server	2 x PowerEdge R720xd server with 2.5" drive bays 2 x Intel Xeon processor E5-2600 product family (8 cores) 64 GB of DDR3 RAM
Number of DAGs	1
Servers per DAG	2
Number of Active and Passive Mailboxes per Server	2,000 active and 2,000 passive
Storage Configuration	**Detail**
Storage hardware	2x PowerEdge R720xd server internal storage 4 drives for Operating System (2 each server) 40 drives for data and log (20 each server) 4 drives for Hot Spares (2 per server) 4 drives for Restore LUNs (2 per server) PERC H710P—RAID controller
Data volumes per mailbox server	10 (5 active + 5 passive)
Databases per volume	1
Mailboxes per database	400
Disk type	2.5" NL-SAS 7.2K rpm 1TB for data 2.5" SAS 10K 300GB for OS
RAID type	RAID 1—Operating System volumes RAID 1—Exchange data volumes RAID 0—Restore LUN
Additional details	Databases and logs on same container Stripe size = 256KB or greater NTFS allocation unit size = 64KB Redundant switch configuration; Private and Public networks configured as VLANs. Hardware load balancer required for balancing load among CAS servers.

Note: DAG can be easily extended to a remote site (if site-resiliency is required). Dell recommends having at least two copies for local site in a Simple Distributed Configuration.

To summarize the solution: 2U PowerEdge R720xd with 24 internal 2.5" HDD used as a multi-role Exchange server. It has two 2.5" Serial Attached SCSI (SAS) internal (hidden) drives that are configured as RAID 1 to host the operating system, and the twenty four 2.5" drives used for the mailbox database and logs. For this configuration, a total of two multi-role Exchange servers are required to host two copies for each database. To make maximum use of spindles and provide better failover and recovery, each database with the data and log combined is hosted on a RAID-1 volume consisting two 2.5" 1TB drives. Each server has five active and five passive copies, for a total of 10 active and 10 passive copies in the complete configuration. Figure-3 shows the logical representation of the reference configuration. A hardware load balancer (not shown in Figure-3) is required to load balance the CAS.

The design accounts for a single failure on a server and at the database level. All copies are replicated without lag to maintain full high-availability.

Figure 3. Internal storage distributed model example

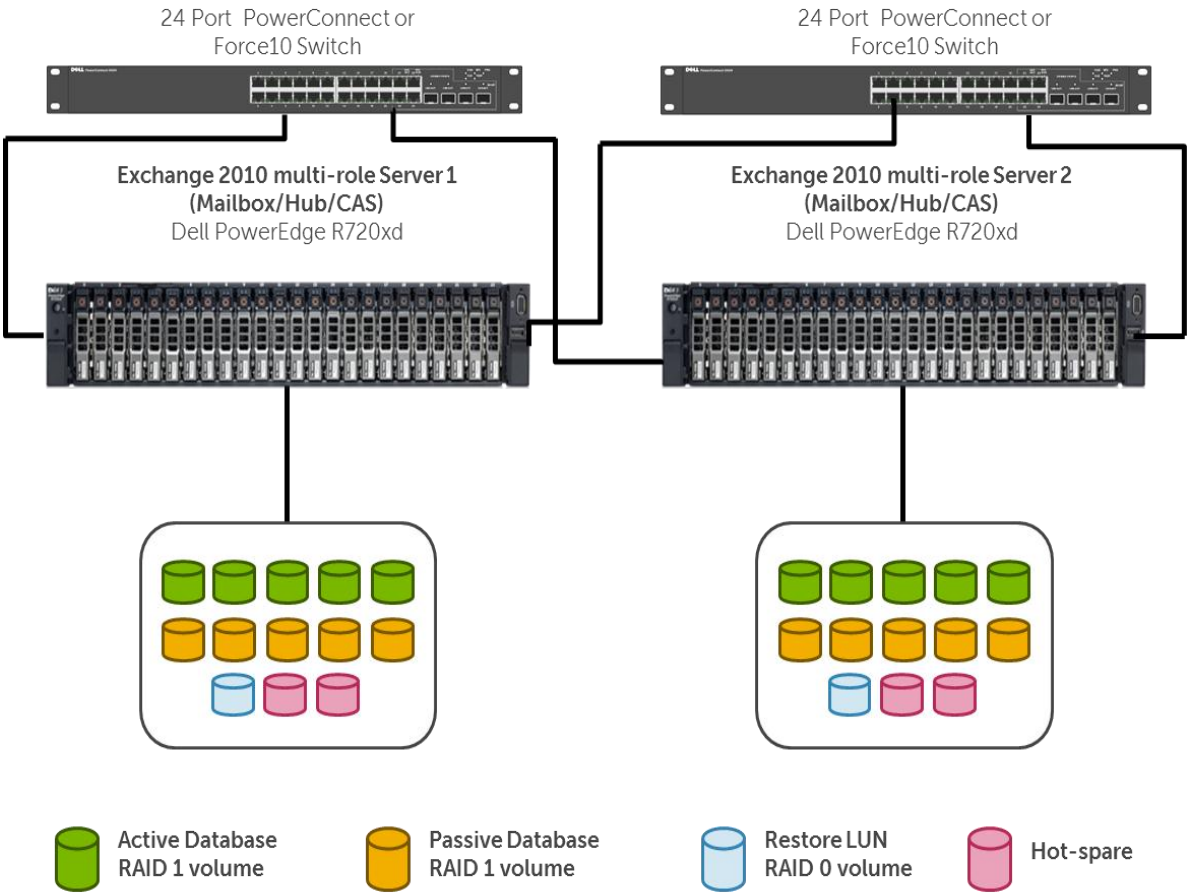

External storage distributed model

The External Storage Distributed model uses disk enclosures that are attached as Direct Attached Storage (DAS) using a RAID controller to the mailbox server, or multi-role Exchange server. Compared to the Internal Storage Distributed Model, this model allows for additional growth in terms of both capacity and I/O per server by increasing the number of arrays that are connected to each server. This model can make maximum use of the multiple storage enclosures before scaling out the configuration for future expansion needs. This option fits the needs of large or growing organizations that require larger mailboxes, and large I/O requirements.

One example solution is described here. The solution requirements are listed in Table-18. The same sizing methodology is applied to these requirements in order to map the best Dell hardware to the needs of the solution. The reference design uses 5000 users who each have a 3GB mailbox, with a medium I/O profile of 100 messages per day, configured in a three copy DAG environment.

Table 18. Solution Requirements

Number of Mailboxes	5000
Average user I/O profile (messages/day)	0.12 IOPS (~ 100 messages per day)
Average mailbox size limit	Up to 3GB
Total active/passive copies per database	3 (2 local + 1 remote)

This solution requires that the database copies are distributed between a central and a remote site. Based on the number of mailbox users and mailbox size, external storage can be selected for this solution. A PowerVault MD1200 enclosure can host up to 12x3.5" 3TB drives, and is considered a building block of this solution to host the storage. Applying the sizing methodology to the solution requirements, the three copy DAG is distributed over the central and the remote site and the PowerEdge R620 satisfies CPU and memory requirements. A solution design with these requirements is described in Table-19.

Table 19. Recommended Solution

Server Configurations	**Detail**
Multi-role (Mailbox/Hub/CAS) server	3x PowerEdge R620 Servers (2 local and 1 remote) 2 x Intel Xeon processor E5-2600 product family (8 cores) 48GB of DDR3 RAM
Number of DAGs	1
Servers per DAG	3
Number of Active and Passive Mailboxes per Server	Local site: 2500 Active and 2500 Passive Remote site: 5000 Passive
Storage Configuration	**Detail**
Storage hardware	2x PowerVault MD1200 (2 per server; 48 local + 24 remote) 60x drives for data and log (20 per server; 40 local + 20 remote) 6 x drives for OS (2 per server; 4 local + 2 remote) 3 Restore LUN (1 per Server) 3 Hot Spare (1 per server) PERC H810—RAID controller
Data volumes per mailbox server	11
Databases per volume	2
Mailboxes per database	227
Disk type	3.5" NL-SAS 7.2K rpm 3TB for data 2.5" SAS 10K rpm 300GB
RAID type	RAID 1— Exchange data and log volumes RAID 1 – Operating System RAID 0 –Restore LUN
Additional details	Databases and logs on same containers Stripe size = 256KB NTFS allocation unit size = 64KB Redundant switch configuration; Private and Public networks configured as VLANs. Hardware load balancer required for balancing load among CAS servers.

The reference design has a 3 copy DAG across two sites – central and remote. There are a total of three PowerEdge R620 1U multi-role Exchange servers in a DAG, with each of the servers connecting to two Dell PowerVault MD 1200 arrays as shown in Figure-4. The connection between the R620 and the first MD1200 in the cascade is through the latest Dell PowerEdge RAID Controller (PERC) family. Both the internal front facing drives and external drives on the MD1200 are used to store the Exchange data and logs. Each server is capable of hosting eleven RAID-1 volumes, and each volume hosts two mailbox databases; each database has 227 mailboxes. So approximately 2500 active and 2500 passive mailbox users are hosted per server at the central site; the R620 server at the remote site is capable of hosting 5000 passive mailboxes over 22 databases.

Figure 4. External Storage Distributed Model Example

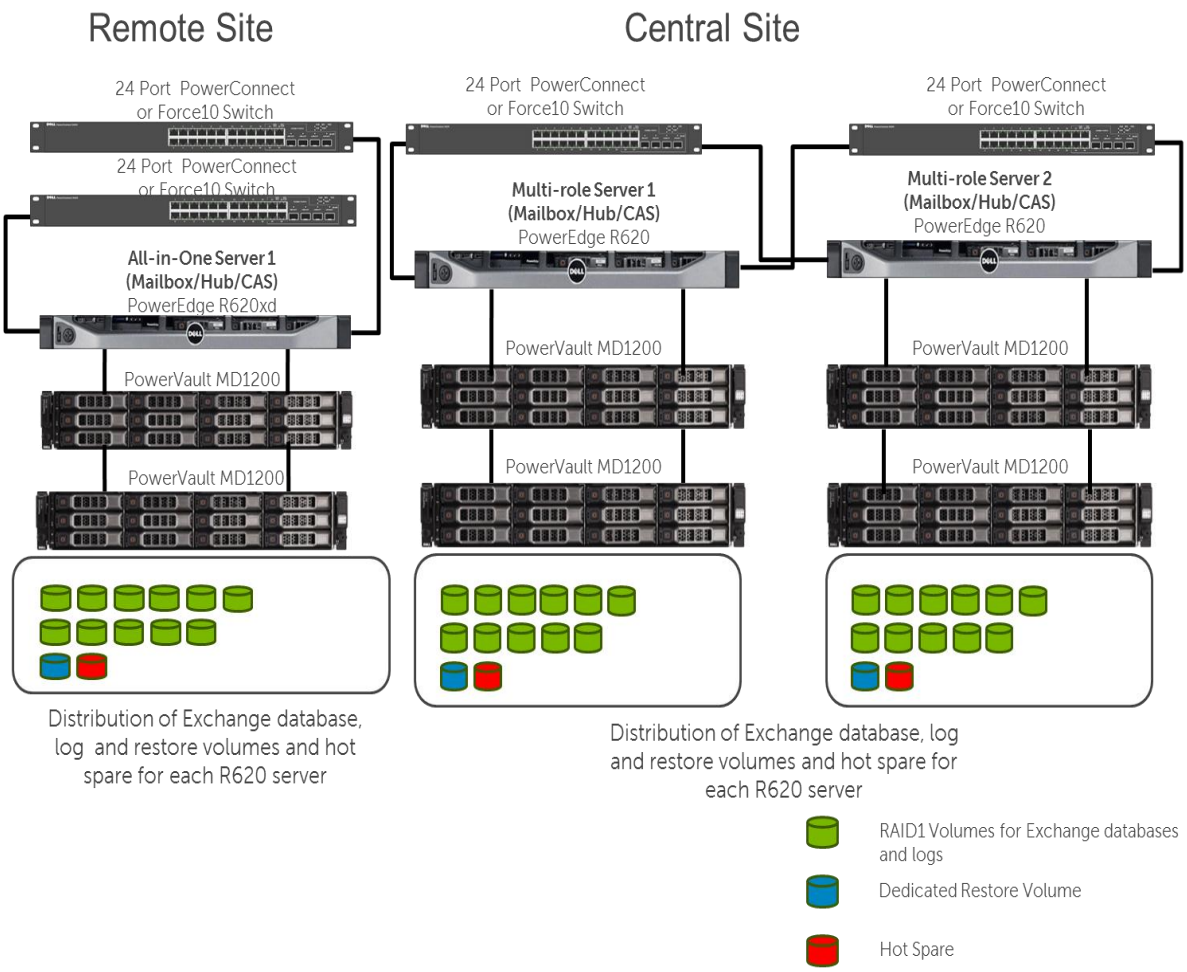

Note: The remote site hosts only passive database copies.

Each volume has two drives: 2 x 3TB 7.2K rpm NL-SAS drives in a RAID-1 configuration for the data and logs. The design allows for one failure, and is sized to have sufficient resources in terms of memory, processor, and storage I/O per server to handle a failure where 1 of the 2 servers go down and all 2600 users become active on a single server. A sample reference design model for the Simple Distributed Model using External Storage is shown in Figure-4.

A solution combining the design for internal and external storage can be achieved by leveraging PowerEdge R720xd servers and PowerVault MD series storage. Such hybrid architecture within the Simple Distributed deployment model can be used for solution requirements demanding large storage capacity and high mailbox IO. More reference configurations of the Simple distributed configuration model are available on online.

Agile Consolidated Model

This architecture is designed for the dynamic data center. It assumes that the organization values server and storage consolidation and standardization over building unique infrastructures for each application. In addition, this model can use virtualization for further consolidation, and the deployment flexibility that comes with abstracting the application and OS from the hardware infrastructure. Dell SAN products, including EqualLogic™ PS arrays and Dell Compellent™ Series 40 arrays, provide the storage foundation for this model.

An example of the Agile Consolidated Model is detailed in this section. The solution requirements are listed in Table-20. The reference design uses 5,000 users who each have a 2GB mailbox and a medium user I/O profile (~100 messages per day) in a two-copy DAG environment.

Table 20. Solution Requirements

Number of mailboxes	5,000
Average user I/O profile (messages/day)	Up to 0.12 (~ 100 messages per day)
Average mailbox size limit	Up to 2 GB
Total active/passive copies per database	2
Not included in this solution	Disaster recovery and site resiliency Backup infrastructure Additional server for Hub/CAS UM and Edge roles

With a lower number of copies and copy failover being very expensive in a SAN environment, it is necessary to deploy and use RAID as the first line of defense in case of disk failure. The Mailbox, CAS and Hub roles are separated out onto different servers or VMs in the case of virtualization, to provide role service availability.

Applying the sizing methodology to the solution requirements, the reference design uses 2 Dell PowerEdge R720xd servers running either VMware ® ESX or Microsoft Hyper-V® hypervisors. Each virtualized server hosts two virtual machines for the Mailbox role, and 2 virtual machines for the combined Hub/CAS role. Microsoft recommends that the two mailbox servers in the same DAG are not hosted on the same hypervisor host. Since there are multiple mailbox servers on a single host, the two DAGs of 2500 users are configured to each have 2 mailbox servers spread across the virtualized hosts. The storage pieces of the solution are two Compellent Series 40 storage arrays arranged in two iSCSI SANs. The rationale behind having two single SC40 in a non-cluster configuration for a high availability solution such as this is that the solution leverages Exchange server 2010 native mailbox resiliency also known as DAG to provide availability.

Each SAN has one Compellent Series 40 array configured with 1x 2U 2.5” drive enclosure cabinet and 1 x 2U 3.5” drive enclosure cabinet. In this setup, the arrays use the Compellent-recommended storage

profile for the best capacity and performance. During normal operation each Mailbox Server hosts 1250 active mailboxes, and each Compellent array hosts 2500 active and 2500 passive mailboxes. The configuration details for these servers and storage systems are provided in Table-21 below. This design is capable of sustaining one server and one storage array failure without disrupting client communication.

Table 21. Agile consolidated recommended Solution

Server Configurations	Detail
ESX/Hyper-V (Host) Server (Processor/Memory)	**Server 1 Config:** 1 x PowerEdge R720 rack server 2 x eight-core processors and 96 GB of RAM **Server 2 Config:** 1 x PowerEdge R720 rack server 2 x eight-core processors and 96 GB of RAM
Mailbox (MBX) Server Virtual Machine (Processor/Memory)	2 x VM on each Host server; 4vCPU/32 GB virtual memory
Hub/CAS server VM	2x VM on each Host server; 4vCPU/8 GB virtual memory
Number of DAGs	2
Servers VM per DAG	2
Number of mailboxes per server VM	1250 active and 1250 passive
Storage Configuration	**Detail**
Storage hardware	**Server 1 Config:** 1 x Dell Compellent Series 40 (11@2TB NL-SAS drives) + (11@900GB 10K SAS drives) **Server 2 Config:** 1 x Dell Compellent Series 40 (11@2TB NL-SAS drives) + (11@900 GB 10K SAS drives)
Data volumes per mailbox server	4
Databases per volume	1
Mailboxes per database	625
Disk type	3.5” 7.2K RPM NLSAS – 3 TB 2.5” 10K SAS – 900 GB
RAID type	Dell Compellent Recommended Storage Profile
Additional details	Databases and logs on same container NTFS allocation unit size = 64 KB 3 hot spares drives 1 hot spare drive per SAN enclosure) Redundant switch configuration; Private and Public networks configured as VLANs.

A sample reference design of an Agile Consolidated Model is shown in Figure-5. A third party load balancer (not shown in Figure-5) for a remote site can also be considered for load balancing for CAS roles.

Figure 5. Agile consolidated model example

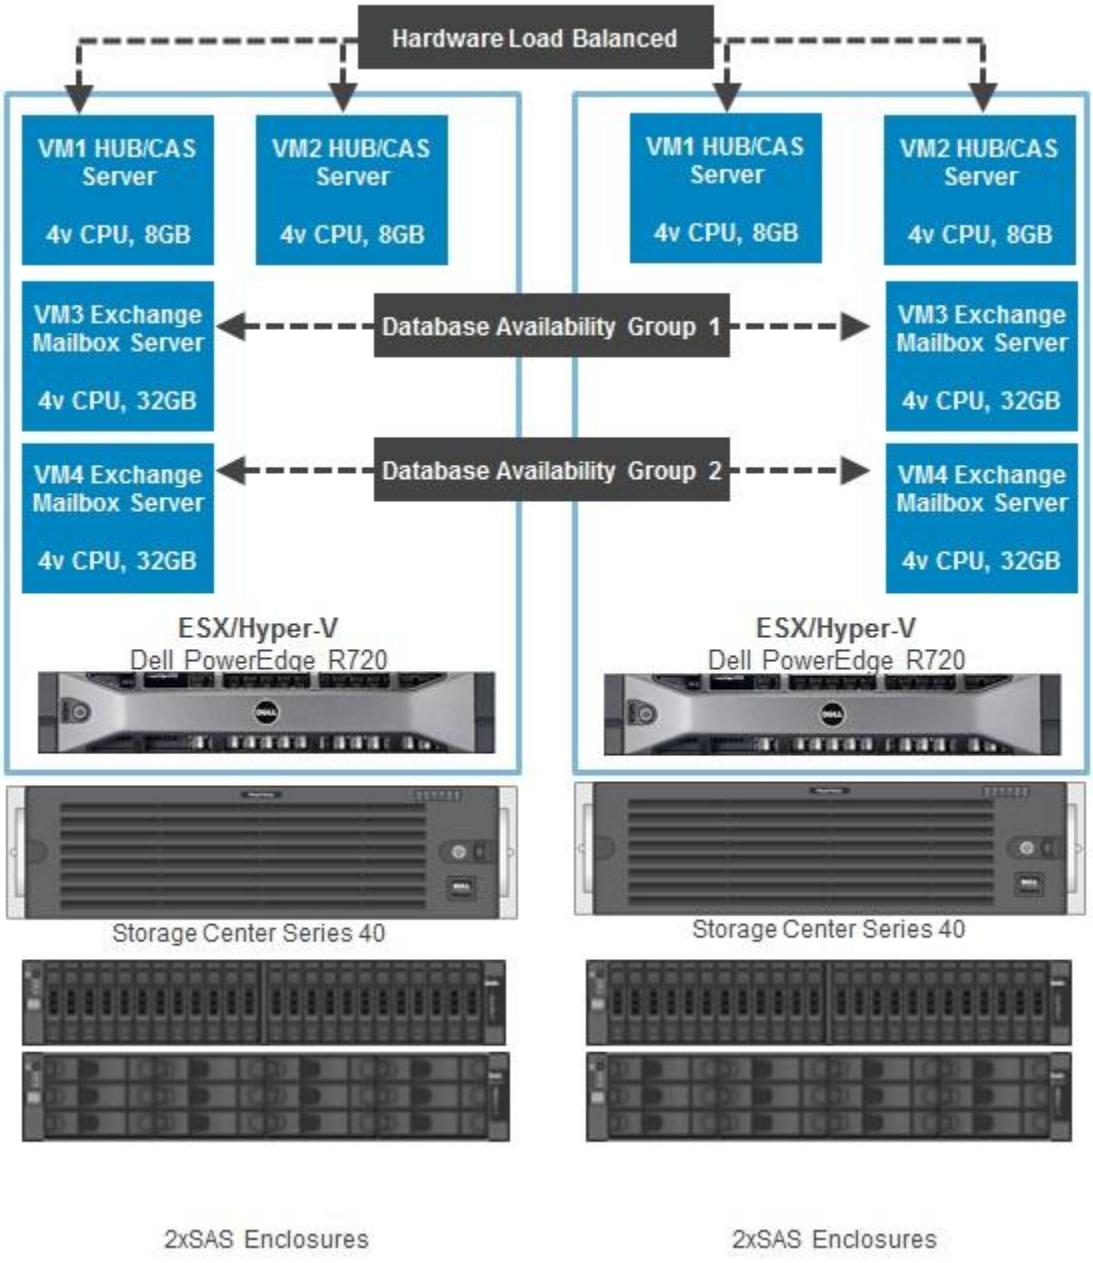

For more examples of Agile Consolidated deployment model refer to the following:

- Exchange Tested Solutions Program whitepaper for 9000 users: http://www.microsoft.com/en-us/download/details.aspx?id=8886.
- More reference configurations of the Agile Consolidated model are available on online.
- For sizing Microsoft Exchange server 2010 on Dell EqualLogic PS6100 and PS4100 Series Arrays refer this paper.

Small and branch office configurations

The following configurations are designed for organizations that have branch offices with up to 1,000 mailboxes. This model leverages tower, or rack servers, such as the Dell™ PowerEdge™ T420 to help provide reliable, Exchange Server 2010 solutions at a reduced cost. These configurations are single server solutions only, no high availability, for organizations that are cost conscious and rely on backup for data resiliency.

The Small Branch Office architecture model uses the internal disks on a dense server as a building block to host the mailbox database and logs. In addition to hosting the Mailbox server role, this server building block can optionally host multiple roles such as Hub and CAS in a multi-role server configuration. This option typically meets the needs of small organizations, or satellite offices requiring a cost-effective messaging solution with high availability, ease of management, and expansion capability. An example of this type of solution is described below. A reference design was created to meet the solution requirements listed in Table-22.

Table 22. Small and Brach office solution requirements

Number of mailboxes	500
Average user I/O profile (messages/day)	.18 IOPS (~150 messages/day)
Average mailbox size limit	4GB
Not included in this solution	High Availability Site Resiliency UM or Edge role

The solution does not need a DAG for the number of mailbox users. In order to provide high availability, the solution depends on back-up services. Applying the sizing methodology to the solution results in needing one PowerEdge™ T420 to provide a reliable, Exchange Server 2010 solution at a reduced cost that uses only internal storage and leverages virtualization for consolidation. An outline of the proposed solution is in Table-23.

Table 23. Small and branch office configuration: solution summary

Server Configurations	**Detail**
Multi-role Exchange (Mailbox/Hub/CAS) server	PowerEdge T420 with dual Intel Xeon processor E5-2400 product family 2 X Quad-core processor with 24GB RAM
Storage Configuration	**Detail**
Storage hardware	PowerEdge T420 with 2.5” internal drives 10 drives for Data 2 drives for OS 2 drives for Log 1 drive for Restore LUN 1 drive for Hotspare PERC H710P RAID Controller

Database volume	5 x RAID-1 volumes
Databases	3 per database volume for a total of 15 databases
Mailboxes per database	34
Log volume	1 x RAID-1 volume
Disk type	1 TB 2.5" 7.2K NL-SAS drive for Data, Log, Restore and Hotspare 300 GB 10K SAS drive for OS
Additional details	Data and Log on separate volumes RAID stripe or element size = 256KB or larger NTFS allocation unit size = 64KB for both database and logs

The proposed reference architecture has been sized to provide mailbox size of up to 2 GB each. The Mailbox I/O profile is sized for up to 150 messages per mailbox per day with a 20% I/O headroom to accommodate up to 0.18 IOPS per mailbox. In addition, this solution offers flexibility by having optional disk space available for backup to disk using backup software such as Dell AppAssure backup software. Figure-6 shows the storage layout within a PowerEdge T420 server.

Figure 6. Small and branch office example

For a virtualized environment scenario for small and branch office deployment solution, the readers can refer to the Exchange server 2010 Tested Solutions for 500 mailbox users in a single site running Hyper-V on Dell server solution here.

A recommended mapping of Dell systems usage to the proposed deployment models is described in Table-24.

Table 24. Recommended Dell system mappings to deployment models

Servers	Storage				Recommended Deployment Model
	Internal	External			
		DAS (PowerVault)	SAN		
			EqualLogic	Compellent	
PowerEdge R720xd	✓	✓			SDM
PowerEdge R720		✓	✓	✓	SDM / ACM
PowerEdge R620		✓	✓	✓	SDM / ACM
PowerEdge R520	✓				SBOM / SDM
PowerEdge T620	✓				SDM / SBOM
PowerEdge T420	✓				SBOM
PowerEdge M620			✓	✓	ACM
PowerEdge M520			✓	✓	ACM

SDM = Simple Distributed Model
ACM = Agile Consolidated Model
SBOM = Small and Branch Office Model

Best practices

This section includes best practice information that isn't typically connected to specific customer requirements, but is critical to both the design and validation of the solution.

Considerations for high availability

- In single-site scenarios, deploy a minimum of a two-node DAG with at least two copies of each mailbox database.
- In a multi-site scenario, deploy at least three mailbox servers as well as three copies of each mailbox database; two at the primary site and one at the secondary site.

Considerations for storage design

- Isolate the Microsoft Exchange Server database and log workload from other I/O intensive applications; assign the Exchange workload to its own set of disks.
- Deploy each DAG copy on its own set of physical disks and separate storage controller; place active and passive copies of the database into separate volumes.
- Design the storage for the active and passive copies of the mailboxes to be identical in terms of capacity and performance.
- Virtual disks that dynamically expand are not supported by Exchange.
- Just a bunch of disks (JBOD) or non-RAID configurations are only an option when you have three or more database copies for the Exchange Mailbox Database (.edb) and log volume.
- PowerEdge RAID Controller (PERC) configuration settings:
    - RAID stripe/element size should be 256 KB or greater.
    - Ensure the RAID cache is enabled and battery backup is fully charged.
    - Read policy: adaptive read ahead.
    - Write policy: write back (with battery backup).
- NTFS allocation unit size should be set to 64KB.
- Physical Disk cache should be disabled.
- For detailed guidelines on understanding the storage configuration refer here.

Considerations for Exchange database and log files

- File placement of the database to ensure log isolation
    - Standalone (Non DAG): database (.edb) file and logs should be isolated and placed on different volumes backed by different physical disks.
    - High availability (DAG): isolation of logs and databases is not required.

- File placement of the database files per volume, or log streams per volume
  - Stand alone: based on the backup methodology.
  - High availability: when using JBOD, divide a single disk into two volumes, one for the database and one for the log stream.
- Database size:
  - Stand alone: It is recommended to have a database size of 200 gigabytes (GB) or less, and to provision the database LUNs for an additional 20 percent of calculated maximum database size.
  - High availability: While a Database size of up to 16 TB is supported. It is recommended to have a database size of 2 TB or less, and to provision for the database LUNs for an additional 20 percent of calculated maximum database size.
- Log truncation:
  - Stand alone: use backups for log truncation such as when circular logging disabled. It is recommended to provision for three days of log generation capacity.
  - High availability: enable circular logging for deployments that use Exchange server 2010 data protection features.

Considerations for virtualization

It is recommended that in a virtualized environment, the ratio of vCPUs to logical cores should be 1:1 when CPU hyper-threading is enabled. When hyper-threading is disabled, the ratio of vCPUs to physical cores should also be 1:1. A 10% overhead on the CPU should be taken into account for virtualized environments. It should be noted that the enhancements provided by virtualized execution environments, like dynamic memory reservation and CPU affinity, should be disabled while running Exchange mailbox server. For solution validation, it is recommended that the validation tool, for example Jetstress, should use the guest operating system from within the virtual machine. General guidelines for running Exchange server in virtualized environment are as follows:

- The storage used by Exchange should be hosted in disk spindles that are separate from the storage that hosts the guest virtual machine's operating system.
- Plan for a dedicated network interface card (NIC) to manage the Hyper-V root server; this card should not be connected to a local Hyper-V virtual switch.
- If using live migration, plan for a dedicated NIC of 1 GB or higher due to the large amount of data that is moved across the network.
- If Internet SCSI (iSCSI) storage is being used, select dedicated, separate NICs for it.
- Memory oversubscription, or dynamic adjustment of virtual machine memory, must be disabled for production Exchange servers.
- Mailbox servers that are members of the same Database Availability Groups (DAGs) should not be deployed on the same root.
- Configure the iSCSI initiator on the host and not on the guest.
- Snapshots, differencing/delta disks is not recommended for production environments

Considerations for networking

- Separate DAG replication traffic from the client MAPI traffic on different physical adapters; it can share the same physical switch with different vLANs.
- For Exchange databases residing on SAN, it is recommended to design separate network infrastructures to isolate the LAN traffic from the SAN traffic (iSCSI).
- Implement redundant components (switches, ISLs, network adapters) to provision a resilient network infrastructure between the endpoints (stack, LAG, load balancing or network card teaming).
- Ensure that the flow control is enabled for the switch ports hosting the EqualLogic PS Series array controller connections and the network adapters dedicated to iSCSI SAN traffic.
- It is recommended to disable spanning tree from the switch ports hosting the PS Series array controller connections, and enable 'Portfast' instead.
- Ensure that jumbo frames (large MTU) are enabled for the switch ports hosting the PS Series array controller connections.

Considerations for backup

To back up and restore Exchange server 2010, you must use an Exchange-aware application that supports the VSS writer for Exchange server 2010 such as Windows Server Backup with the VSS plug-in, Microsoft System Center Data Protection Manager, or a third-party Exchange-aware VSS-based application.

For larger organizations, to back up passive mailbox database copies or to conduct incremental or differential backups, you need either the Microsoft System Center Data Protection Manager or a third-party Exchange-aware VSS-based application.

For smaller organizations on a limited budget, using the Exchange server 2010 VSS plug-in for Windows Server Backup can be a cost-effective way to back up and restore Exchange data. There are several limitations to this solution:

- You can only back up standalone (not replicated), or active database copies. The passive copy of the database cannot be backed up.
- Backups made with Windows Server Backup occur at the volume level. To back up a database and its log stream, you must back up the entire volume that contains the database and logs. You can't back up any data without backing up the entire volume that contains the data.
- The backup must run locally on the server being backed up, and you can't use the plug-in to make remote VSS backups.
- The backup can be created on a local drive or on a remote network share.
- Only full backups can be made. Log truncation occurs only after a successful completion of a full backup of a volume containing an Exchange database.

Windows Server Backup (with the VSS plug-in) is viable backup solution for smaller organizations since there isn't a requirement for point-in-time recovery with minimal data loss. Instead, any issues with multiple copies of a database won't result in total data loss. Weekly full backups of all Exchange databases can be made every Saturday evening, and data can be backed up to a remote share located

on a file server that has sufficient capacity to hold one full weekly backup. The file server should also have sufficient capacity for a recovery volume.

It is strongly recommended to deploy a minimum of three copies of a mailbox database at the local site before eliminating traditional forms of database protection, such as Redundant Array of Independent Disks (RAID) or traditional VSS-based backups.

Considerations for server configurations

The type of memory selected and how the memory is populated in a system will have a major impact on the performance of the system. To achieve maximum memory performance while running Exchange server, the memory channels on the servers should be configured with right number and size of DIMMs. To extract all available performance out of the memory system, populate all memory channels evenly with the same rank, size, and speed of DIMMs. Servers using the Intel Xeon E5-2600 processors have four memory channels as compared to the Intel Xeon E5-2500 processor based systems, which have three memory channels. It should be noted that, for most system configurations, one DIMM per channel and two DIMMs per channel memory configurations are more optimal than three DIMMs per channel, due to the fact that memory frequency will be reduced when all slots in a given memory channel are populated.

When mixing DIMM sizes, all memory should be chosen based on identical rank structure, and all memory channels should be populated with an identical mix of DIMM types. It is recommended that, Registered DIMMs (RDIMMs) should be preferred over Unregistered DIMMs (UDIMMs) in most cases, due to the greater selection, higher frequency, and greater capacity capabilities of RDIMMs. More details on memory configuration options for 12th generation Dell servers can be found here.

While running Exchange server, the administrator should keep a tap on CPU utilization targets. The following tables describe the peak CPU utilization targets for normal operations, and for a site server failure or server maintenance condition.

Table 25. Server utilization targets

Target server CPU utilization design assumption	Normal Operating Values	Operating values on node failure
Mailbox servers	<70%	<80%
Client Access servers	<70%	<80%
Transport servers	<70%	<80%
Multiple server roles (Client Access, Hub Transport, and Mailbox servers)	<70%	<80%
Multiple server roles (Client Access and Hub Transport servers)	<70%	<80%

Considerations for storage configurations

The following table summarizes the data configuration and I/O assumptions made when designing the storage configuration for the Exchange solution.

Table 26. Data configuration and I/O assumptions

Data configuration assumption	Value/description
Data overhead factor	20%
Mailbox moves per week	1%
Dedicated maintenance or restore logical unit number (LUN)	No
LUN free space	20%
Log shipping compression enabled	Yes
Log shipping encryption enabled	Yes
I/O configuration assumption	**Value/description**
I/O overhead factor	20%
Additional I/O requirements	None

Conclusion

The Exchange server 2010 architecture has been modified to support the deployment of large mailboxes, and inexpensive storage systems. The high-availability model has been redefined to take advantage of continuously replicated multiple copies, and to provide near instant switchover as well as the ability to go back to a point-in-time in response to data corruption. A number of security and functional enhancements in the core architecture provide a more reliable and better end-user experience. This paper provides details of the core architectural changes that impact an Exchange server 2010 solution design, and sever and storage sizing guidelines to consider before deploying Exchange server 2010 in your infrastructure. It also presents Dell deployment models to help customers make a balanced decision based on their environment and requirements.

Dell PowerEdge Servers along with Dell PowerVault, Dell EqualLogic and Dell Compellent Storage systems, provide the appropriate hardware platforms for deploying Exchange Server 2010 messaging solutions. For detailed sizing of Microsoft Exchange server 2010, Dell provides the Exchange server 2010 Advisor Tool (available at www.dell.com/exchange) Dell Services includes assessment, design and implementation tailored for those messaging deployments. Dell also offers end-to-end Exchange messaging solutions that include offerings for security, backup and recovery.