

Exchange 2010 Tested Solutions: 500 Mailboxes in a Single Site Running Hyper-V on Dell Servers

Rob Simpson, Program Manager, Microsoft Exchange Server; Akshai Parthasarathy, Systems Engineer, Dell; Casey Birch, Product Marketing Manager for Exchange Solutions, Dell

December 2010

In Exchange 2010 Tested Solutions, Microsoft and participating server, storage, and network partners examine common customer scenarios and key design decision points facing customers who plan to deploy Microsoft Exchange Server 2010. Through this series of white papers, we provide examples of well-designed, cost-effective Exchange 2010 solutions deployed on hardware offered by some of our server, storage, and network partners.

You can download this document from the [Microsoft Download Center](#).

Applies to:

Microsoft Exchange Server 2010 release to manufacturing (RTM)

Microsoft Exchange Server 2010 with Service Pack 1 (SP1)

Windows Server 2008 R2

Windows Server 2008 R2 Hyper-V

Table of Contents

-
- [Solution Overview of Requirements](#)
- [Customer Requirements](#)
 - [Mailbox Profile Requirements](#)
 - [Geographic Location Requirements](#)
 - [Server and Data Protection Requirements](#)
- [Design Assumptions](#)
 - [Server Configuration Assumptions](#)
 - [Storage Configuration Assumptions](#)
- [Solution Design](#)
 - [Determine High Availability Strategy](#)
 - [Estimate Mailbox Storage Capacity Requirements](#)
 - [Estimate Mailbox I/O Requirements](#)

- [Determine Storage Type](#)
- [Estimate Mailbox Memory Requirements](#)
- [Estimate Mailbox CPU Requirements](#)
- [Determine Placement of Exchange Server Roles](#)
- [Determine Load Balancing Strategy](#)
- [Determine Whether Server Virtualization Will Be Used](#)
- [Determine Whether Client Access and Hub Transport Server Roles Will Be Deployed in Separate Virtual Machines](#)
- [Determine Server Model for Hyper-V Root Server](#)
- [Calculate CPU Capacity of Root Server Model](#)
- [Determine CPU Capacity of Virtual Machines](#)
- [Determine Number of Mailbox Server Virtual Machines Required](#)
- [Determine Memory Required per Mailbox Server Virtual Machine](#)
- [Determine Number of Client Access and Hub Transport Combination Virtual Machines Required](#)
- [Determine Memory Required per Combined Client Access and Hub Transport Virtual Machines](#)
- [Determine Virtual Machine Distribution](#)
- [Determine Memory Required per Root Server](#)
- [Design Database Availability Groups](#)
- [Design Database Copy Layout](#)
- [Determine Storage Configuration](#)
- [Design Client Access Array and Load Balancing](#)
- [Determine Placement of the File Share Witness](#)
- [Solution Overview](#)
 - [Logical Layout](#)
 - [Physical Layout](#)
 - [Server Hardware Summary](#)
 - [Client Access and Hub Transport Server Configuration](#)
 - [Mailbox Server Configuration](#)
 - [Database Layout](#)
 - [Storage Configuration](#)
- [Solution Variation](#)
- [Solution Validation Methodology](#)
 - [Storage Design Validation Methodology](#)
 - [Server Design Validation](#)
 - [Functional Validation Tests](#)

- [Solution Validation Results](#)
 - [Functional Validation Results](#)
 - [Storage Design Validation Results](#)
 - [Server Design Validation Results](#)

This document provides an example of how to design, test, and validate an Exchange Server 2010 solution for environments with 500 mailboxes or less deployed on Dell server and storage solutions. One of the key challenges with designing smaller Exchange 2010 environments is being able to provide large mailboxes that are highly available at a reasonable price point. Another key challenge is examining the current server and storage options available and making the right hardware choices that provide the best value over the anticipated life of the solution. Following the step-by-step methodology in this document, we walk through the important design decision points that help address these key challenges while ensuring that the customer's core business requirements are met. After we have determined the optimal solution for this customer, the solution undergoes a standard validation process to ensure that it holds up under simulated production workloads for normal operating, maintenance, and failure scenarios.

[Return to top](#)

Solution Overview of Requirements

The following tables summarize the key Exchange and hardware components of this solution. The remainder of this section discusses in detail the requirements for:

- Mailbox profile
- Geographic location
- Server and data protection

Exchange components

Exchange component	Value or description
Target mailbox count	500
Target mailbox size	5 gigabytes (GB)
Target message profile	150 messages per day
Database copy count	2
Volume Shadow Copy Service (VSS) backup	Windows Server (VSS plug-in)
Site resiliency	No
Virtualization	Hyper-V
Exchange server count	4 virtual machines (VMs)

Exchange component	Value or description
Physical server count	2

Hardware components

Hardware component	Value or description
Server partner	Dell
Server model	PowerEdge T610
Server type	Tower
Processor	Intel Xeon 5550
Storage partner	Dell
Storage type	Local disk
Disk type	1 terabyte 7.2 kilobyte (KB) Serial ATA (SATA) 3.5"

[Return to top](#)

Customer Requirements

One of the most important first steps in Exchange solution design is to accurately summarize the business and technical requirements that are critical to making the correct design decisions. The following sections outline the customer requirements for this solution.

Mailbox Profile Requirements

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. The following tables summarize the mailbox profile requirements for this solution.

Mailbox count requirements

Mailbox count requirements	Value
Mailbox count (total number of mailboxes including resource mailboxes)	500
Projected growth percent (%) in mailbox count (projected increase in mailbox count over the life of the solution)	0%

Mailbox count requirements	Value
Expected mailbox concurrency % (maximum number of active mailboxes at any time)	100%

Mailbox size requirements

Mailbox size requirements	Value
Average mailbox size in megabytes (MB)	5000 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	5000 MB

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
% in MAPI cached mode	100%
% in MAPI online mode	0%
% in Outlook Anywhere cached mode	0%
% in Outlook Web App (Microsoft Office Outlook Web Access in Exchange 2007 and previous versions)	0%
% in Exchange Active Sync	0%

[Return to top](#)

Geographic Location Requirements

Understanding the distribution of mailbox users and datacenters is important when making design decisions about high availability and site resiliency.

The following table outlines the geographic distribution of people who will be using the Exchange system.

Geographic distribution of people

Mailbox user site requirements	Value
Number of major sites containing mailbox users	1
Number of mailbox users in site 1	500

The following table outlines the geographic distribution of datacenters that could potentially support the Exchange e-mail infrastructure.

Geographic distribution of datacenters

Datacenter site requirements	Value or description
Total number of datacenters	1
Number of active mailboxes in proximity to datacenter 1	500
Requirement for Exchange to reside in more than one datacenter	No

[Return to top](#)

Server and Data Protection Requirements

It's also important to define server and data protection requirements for the environment because these requirements will support design decisions about high availability and site resiliency.

The following table identifies server protection requirements.

Server protection requirements

Server protection requirement	Value or description
Number of simultaneous server or VM failures within site	1
Number of simultaneous server or VM failures during site failure	Not applicable

The following table identifies data protection requirements.

Data protection requirements

Data protection requirement	Value or description
Requirement to maintain a backup of the Exchange databases outside of the Exchange environment (for example, third-party backup)	Yes

Data protection requirement	Value or description
solution)	
Requirement to maintain copies of the Exchange databases within the Exchange environment (for example, Exchange native data protection)	Yes
Requirement to maintain multiple copies of mailbox data in the primary datacenter	Yes
Requirement to maintain multiple copies of mailbox data in a secondary datacenter	No
Requirement to maintain a lagged copy of any Exchange databases	No
Lagged copy period in days	Not applicable
Target number of database copies	Minimum of 2
Deleted Items folder retention window in days	14 days

[Return to top](#)

Design Assumptions

This section includes information that isn't typically collected as part of customer requirements, but is critical to both the design and the approach to validating the design.

Server Configuration Assumptions

The following table describes the peak CPU utilization targets for normal operating conditions, and for site server failure or server maintenance conditions.

Server utilization targets

Target server CPU utilization design assumption	Value
Normal operating for Mailbox servers	<70%
Normal operating for Client Access servers	<70%
Normal operating for Hub Transport servers	<70%
Normal operating for multiple server roles (Client Access, Hub Transport, and Mailbox servers)	<70%
Normal operating for multiple server roles	<70%

Target server CPU utilization design assumption	Value
(Client Access and Hub Transport servers)	
Node failure for Mailbox servers	<80%
Node failure for Client Access servers	<80%
Node failure for Hub Transport servers	<80%
Node failure for multiple server roles (Client Access, Hub Transport, and Mailbox servers)	<80%
Node failure for multiple server roles (Client Access and Hub Transport servers)	<80%

[Return to top](#)

Storage Configuration Assumptions

The following tables summarize some data configuration and input/output (I/O) assumptions made when designing the storage configuration.

Data configuration assumptions

Data configuration assumption	Value or 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 assumptions

I/O configuration assumption	Value or description
I/O overhead factor	20%
Additional I/O requirements	None

[Return to top](#)

Solution Design

The following section provides a step-by-step methodology used to design this solution. This methodology takes customer requirements and design assumptions and walks through the key design decision points that need to be made when designing an Exchange 2010 environment.

Determine High Availability Strategy

When designing an Exchange 2010 environment, many design decision points for high availability strategies impact other design components. We recommend that you determine your high availability strategy as the first step in the design process. We highly recommend 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

Step 1: 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.

Design Decision Point

In this 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 2010 design will be based on a single site deployment with no site resiliency.

Step 2: Determine backup and database resiliency strategy

Exchange 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 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 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 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 2010 as a replacement for traditional backups. Prior to making this decision, see Understanding Backup, Restore and Disaster Recovery.

Design Decision Point

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 in place of traditional backups as a database resiliency strategy would be an improvement. However, due to a limited budget, the three Mailbox servers required to support a minimum of three database copies can't be deployed. The customer is strongly urged to consider implementing a backup solution.

Step 3: Determine backup solution

To back up and restore Exchange 2010, you must use an Exchange-aware application that supports the VSS writer for Exchange 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 smaller organizations on a limited budget, using the Exchange 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 of this solution:

- You can only back up stand-alone (not replicated) or active database copies. You can't back up the passive copy of the database.
- Backups made with Windows Server Backup occur at volume level. To back up a database and its log stream, you must back up the entire volume containing the database and logs. You can't back up any data without backing up the entire volume containing the data.
- The backup must be 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.

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

For additional information, see [Using Windows Server Backup to Back Up and Restore Exchange Data](#).

Design Decision Point

The decision is made to use Windows Server Backup (with the VSS plug-in) as the backup solution. There isn't a concern with point-in-time recovery with minimal data loss. Instead, the focus is on ensuring that any issues with multiple copies of a database won't result in total data loss. Weekly full backups of all Exchange databases will be made every Saturday evening. Data will 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 also has sufficient capacity for a recovery volume. Additional information about backup design and capacity planning is provided in later steps.

Step 4: 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 before eliminating traditional forms of protection for the database, such as Redundant Array of Independent Disks (RAID) or traditional VSS-based backups.

For additional information, see [Understanding Mailbox Database Copies](#).

Design Decision Point

The 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. You could deploy three database copies, but that would require a minimum of three physical Mailbox servers and would likely exceed the current budget. To save costs, you could deploy the database copies on a RAID-0 or just a bunch of disks (JBOD) disk configuration, but because you only have two copies of the data, we recommend that RAID be used to protect your storage from a single disk failure. Because only two database copies are used, a backup solution using Windows Server Backup (with the VSS plug-in) will also be implemented.

Step 5: Determine number of Mailbox servers per DAG

In this step, you need to determine the minimum number of Mailbox servers required to support the database availability group (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.

Design Decision Point

A minimum of two Mailbox servers is required to support a DAG configuration with two database copies.

[Return to top](#)

Estimate Mailbox Storage Capacity Requirements

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. A later section covers additional calculations required to properly design the storage layout on the chosen storage platform.

Microsoft has created a Mailbox Server Role Requirements Calculator that will do most of this work for you. To download the calculator, see [E2010 Mailbox Server Role Requirements Calculator](#). For additional information about using the calculator, see [Exchange 2010 Mailbox Server Role Requirements Calculator](#).

Step 1: 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 1-GB 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), and the average database daily variations per mailbox. The Mailbox server role requirements calculator does these calculations for you. You can also use the following information to do the calculations manually.

The following calculations are used to determine the mailbox size on disk for this solution:

- Whitespace = 150 messages per day \times 75 \div 1024 MB = 11 MB
- Dumpster = (150 messages per day \times 75 \div 1024 MB \times 14 days) + (5120 MB \times 0.012) + (5120 MB \times 0.058) = 512 MB
- Mailbox size on disk = mailbox size quota + whitespace + dumpster
= 5000 + 11 + 512
= 5523 MB

Step 2: Calculate database storage capacity requirements

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:

- Database size = (number of mailboxes \times mailbox size on disk \times database growth factor) + (20% data overhead)
= (500 \times 5523 \times 1) + (552300)
= 3313800 MB
= 3236 GB
- Catalog index size = 10% of database size
= 3236 \times 0.10

=323 GB

- Total database capacity = (database size) + (index size) + 20% volume free space
= (3236 + 232) ÷ 0.8
= 4449 GB

Step 3: Calculate transaction log storage capacity requirements

To ensure that the Mailbox server doesn't sustain any outages as a result of space allocation issues, the transaction logs also need to be sized to accommodate all the logs that will be generated during the backup set. Because a weekly full backup will be made, the log capacity should allocate for seven times the daily log generation rate.

To determine the storage capacity required for all logs, use the following formula. The calculated capacity includes log files size, mailbox move overhead, and 20 percent free space:

- Log files size = (log file size × number of logs per mailbox per day × number of days between log truncations × number of mailbox users) + (1% mailbox move overhead)
= (1 MB × 30 × 7 × 500) + (500 × 0.01 × 5000 MB)
= 130000 MB = 127 GB

Step 4: Determine total storage capacity requirements

The following table 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. You will then take a closer look at specific storage requirements in later steps.

Summary of storage capacity requirements

Disk space requirements	Value
Database capacity required (GB)	4449 GB
Log capacity required (GB)	127 GB
Total capacity required (GB)	4576 GB
Total capacity required for 2 database copies (GB)	9152 GB
Total capacity required for 2 database copies (terabytes)	8.9 terabytes

The high level storage capacity requirements are approximately 9 terabytes. When choosing a storage solution, ensure that the solution meets this capacity requirement.

[Return to top](#)

Estimate Mailbox I/O Requirements

When designing an Exchange environment, you need an understanding of database and log performance factors. We recommend that you review Understanding Database and Log Performance Factors.

Calculate total mailbox I/O requirements

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:

- Total required IOPS = IOPS per mailbox user × number of mailboxes × I/O overhead factor
- Estimated IOPS per mailbox user = 0.15

 **Note:**

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.

- Number of mailbox users = 500
- I/O overhead factor = 20%
- Total required IOPS = $0.15 \times 500 \times 1.20 = 90$

The high level storage IOPS requirements are approximately 90. When choosing a storage solution, ensure that the solution meets this requirement.

[Return to top](#)

Determine Storage Type

Exchange 2010 includes improvements in performance, reliability, and high availability that enable organizations to run Exchange on a wide range of storage options.

When examining the storage options available, being able to balance the performance, capacity, manageability, and cost requirements is essential to achieving a successful storage solution for Exchange.

For more information about choosing a storage solution for Exchange 2010, see Mailbox Server Storage Design.

Step 1: Determine whether you prefer an internal or external storage solution

A number of server models on the market today support from 8 through 16 internal disks. These servers are a fit for some Exchange deployments and provide a solid solution at a low price point.

If your storage capacity and I/O requirements are met with internal storage and you don't have a specific requirement to use external storage, you should consider using server models with an internal disk for Exchange deployments. If your storage and I/O requirements are higher or your organization has an existing investment in SANs, you should examine larger external direct-attached storage (DAS) or SAN solutions.

Design Decision Point

Because only 6.6 terabytes of storage capacity and 90 IOPS is required in this example, you will explore server models with adequate internal disk space to meet these requirements.

Step 2: Determine RAID level

Because you're deploying only two copies of each database, we recommend that you protect each of the database copies with some level of RAID. In this case, there are two options:

- **RAID-10** This option provides the best performance and has a minimal rebuild performance impact, but it requires double the required capacity on disk.
- **RAID-5** This option provides reasonable performance and doesn't require a doubling of storage, but it has a strong performance impact during a rebuild following a disk failure or replacement.

Design Decision Point

Both options work. For optimal performance, you could select RAID-10, but this option likely requires two terabyte disks. Because cost is a main factor in design decisions, select RAID-5. You can deploy smaller, less expensive disks and still support a mailbox size of about 5 GB. You can mitigate the RAID-5 rebuild performance impact by moving all active databases to the alternate Mailbox server in the event of disk failure or replacement.

[Return to top](#)

Estimate Mailbox Memory Requirements

Sizing memory correctly is an important step in designing a healthy Exchange environment. We recommend that you review Understanding Memory Configurations and Exchange Performance and Understanding the Mailbox Database Cache.

Calculate required database cache

The Extensible Storage Engine (ESE) uses database cache to reduce I/O operations. In general, the more database cache available, the less I/O generated on an Exchange 2010 Mailbox server. However, there's a point where adding additional database cache 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.

The IOPS estimates that you completed in a previous step assume a minimum amount of database cache per mailbox. These minimum amounts are summarized in the table "Estimated IOPS per mailbox based on message activity and mailbox database cache" in Understanding the Mailbox Database Cache.

The following table outlines the database cache per user for various message profiles.

Database cache per user

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

In this step, you determine high level memory requirements for the entire environment. In a later step, you use this result to determine the amount of physical memory needed for each Mailbox server. Use the following information:

- Total database cache = database cache per mailbox user x total number of mailbox users
- Database cache per mailbox user = 9 MB
- Number of mailbox users = 500
- Total database cache = 9×500
= 4500 MB
= 4.4 GB

The total database cache requirements for the environment are about 5 GB.

[Return to top](#)

Estimate Mailbox CPU Requirements

Mailbox server capacity planning has changed significantly from previous versions of Exchange due to the new mailbox database resiliency model provided in Exchange 2010. For additional information, see Mailbox Server Processor Capacity Planning.

In the following steps, you calculate the high level megacycle requirements for active and passive database copies. These requirements will be used in a later step to determine the number of Mailbox servers needed to support the workload. Note that the number of Mailbox servers required also depends on the Mailbox server resiliency model and database copy layout.

Using megacycle requirements to determine the number of mailbox users that an Exchange Mailbox server can support isn't an exact science. A number of factors can result in unexpected megacycle results in test and production environments. Megacycles should only be used to approximate the number of mailbox users that an Exchange Mailbox server can support. It's always better to be conservative rather than aggressive during the capacity planning portion of the design process.

The following calculations are based on published megacycle estimates as summarized in the following table.

Megacycle estimates

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

Step 1: Calculate active mailbox CPU requirements

In this step, you calculate the megacycles required to support the active database copies, using the following:

- Active mailbox megacycles required = profile specific megacycles × number of mailbox users
= 3 × 500
= 1500

Step 2: Calculate active mailbox remote database copy CPU requirements

In a design with two copies of each database, there is processor overhead associated with shipping logs required to maintain the database copy on the remote server. This overhead is typically 10 percent of the active mailbox megacycles for each remote copy being serviced. Calculate the requirements, using the following:

- Remote copy megacycles required = profile specific megacycles × number of mailbox users × number of remote copies
= 0.3 × 500 × 1
= 150

Step 3: Calculate local passive mailbox CPU requirements

In a design with two copies of each database, there is processor overhead associated with maintaining the local passive copies of each database. In this step, you calculate the high level megacycles required to support local passive database copies. You refine these numbers in a later step so that they match the server resiliency strategy and database copy layout. Calculate the requirements, using the following:

- Passive mailbox megacycles required = profile specific megacycles × number of mailbox users × number of passive copies
= 0.45 × 500 × 1
= 225

Step 4: Calculate total CPU requirements

Calculate the total requirements, using the following:

- Total megacycles required = active mailbox + remote copies + local passive copies
= 1500 + 150 + 225
= 1875

The total megacycles required to support the environment are approximately 1,875.

[Return to top](#)

Determine Placement of Exchange Server Roles

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

Design Decision Point

Because of the small number of mailbox users, it isn't cost effective to deploy separate physical servers for Client Access and Hub Transport server roles. The decision is to deploy the Client Access, Hub Transport, and Mailbox server roles on the same physical servers.

[Return to top](#)

Determine Load Balancing Strategy

Load balancing Client Access servers in small environments can be a design challenge. Entry level hardware load balancers can be expensive and usually aren't a cost-effective solution for 500 mailbox users. You can choose to deploy a Windows Network Load Balancing software solution, which is an installable component of the Windows Server operating system. However, Windows Network Load Balancing can't be enabled on any server that has Windows failover clustering enabled. When you add a Mailbox server to a DAG, Windows failover clustering is enabled. Therefore, you can't use Windows Network Load Balancing to load balance a Client Access server role installed on the same server as a Mailbox server role that's a member of a DAG. Virtualization can be used to solve this issue. By deploying Client Access and Mailbox server roles in separate VMs, you can enable Windows Network Load Balancing on the Client Access server VM without conflicting with the Windows failover clustering components on the Mailbox server VM.

Design Decision Point

Because there isn't funding to purchase a hardware load balancer, the decision is to explore options for using Windows Network Load Balancing.

[Return to top](#)

Determine Whether Server Virtualization Will Be Used

Several factors are important when considering server virtualization for Exchange. For more information about supported configurations for virtualization, see Exchange 2010 System Requirements.

The main reasons customers use virtualization with Exchange are as follows:

- If you expect server capacity to be underutilized and anticipate better utilization, you may purchase fewer servers as a result of virtualization.
- You may want to use Windows Network Load Balancing when deploying Client Access, Hub Transport, and Mailbox server roles on the same physical server.
- If your organization is using virtualization in all server infrastructure, you may want to use virtualization with Exchange, to be in alignment with corporate standard policy.

Design Decision Point

The decision is to use virtualization for Client Access, Hub Transport, and Mailbox server roles, and Windows Network Load Balancing to balance client load across the Client Access servers. Microsoft Hyper-V is used as the virtualization platform.

[Return to top](#)

Determine Whether Client Access and Hub Transport Server Roles Will Be Deployed in Separate Virtual Machines

When using virtualization for the Client Access and Hub Transport server roles, you may consider deploying both roles on the same VM. This approach reduces the number of VMs to manage, the number of server operating systems to update, and the number of Windows and Exchange licenses you need to purchase. Another benefit to combining the Client Access and Hub Transport server roles is to simplify the design process. When deploying roles in isolation, we recommend that you deploy one Hub Transport server logical processor for every four Mailbox server logical processors, and that you deploy three Client Access server logical processors for every four Mailbox server logical processors. This can be confusing, especially when you have to provide sufficient Client Access and Hub Transport servers during multiple VM or physical server failures or maintenance scenarios. When deploying Client Access, Hub Transport, and Mailbox servers on like physical servers or like VMs, you can deploy one server with the Client Access and Hub Transport server roles for every one Mailbox server in the site.

Design Decision Point

The decision is to locate the Hub Transport and Client Access server roles together in the same VM. The Mailbox server role is deployed separately in a second VM. This reduces the number of VMs and operating systems to manage, and makes it much easier to plan for server resiliency.

[Return to top](#)

Determine Server Model for Hyper-V Root Server

You can use the following steps to determine the server model for the Hyper-V root server.

Step 1: Identify preferred server vendor

In this example, the preferred server vendor is Dell. Dell, Inc. is a leading IT infrastructure and services company with a broad portfolio of servers, storage, networking products, and comprehensive service offerings. Dell also provides testing, best practices, and architecture guidance specifically for Exchange 2010 and other Microsoft-based solutions in the unified communications and collaboration stack such as Microsoft Office SharePoint Server and Office Communications Server.

Dell PowerEdge 11G servers are designed to enhance your IT experience. Dell is changing the way you interact with technology, with easier management, more reliability, and better energy efficiency. The Lifecycle Controller for advanced manageability, interactive LCD screens for system monitoring, dual embedded hypervisors for added redundancy, and FlexMem Bridge technology for maximum memory access help you better utilize the features in Exchange.

Step 2: Review available options from Dell

The Dell server portfolio includes several models that could be considered for this implementation.

Option 1: [Dell PowerEdge T610 Tower Server](#)

In a previous step, you explored options that have internal storage to reduce cost and complexity. The Dell PowerEdge T610, which contains eight internal hard disk drives, is a good choice. The T610 is an Intel-based tower system with two sockets and up to eight internal drives. Twelve dual inline memory module (DIMM) slots provide adequate memory for small to medium sized Mailbox and multiple role servers.

Dell PowerEdge T610

Components	Description
Processors (x2)	Latest quad-core or six-core Intel Xeon 5500 and 5600 series processors
Form factor	Tower or 5U rack-mountable
Memory:	Up to 192 GB (12 DIMM slots/6 per-processor): 1 GB/2 GB/4 GB/8 GB/16 GB DDR3 800 MHz, 1066 MHz, or 1333 MHz
Drives	8 x 2.5" hard disk drive option or 8 x 3.5" hard disk drive option
I/O slots	2 PCIe x8 + 3 PCIe x4 G2

Option 2: [Dell PowerEdge T710 Tower Server](#)

For larger numbers of mailboxes in a similar form factor, the Dell PowerEdge T710 is another Intel-based, two socket tower system. It has 18 DIMM slots and the option of up to 16 internal hard disk drives. The expanded memory and number of drives of the T710 scale the capabilities of the tower platform to larger Exchange environments, or even a virtualization server.

Dell PowerEdge T710

Components	Description
Processors (x2)	Latest quad-core or six-core Intel Xeon 5500 and 5600 series processors
Form factor	Tower or 5U rack-mountable
Memory	Up to 192 GB (18 DIMM slots): 1 GB/2 GB/4 GB/8 GB/16 GB DDR3 800 MHz, 1066 MHz, or 1333 MHz
Drives	16 x 2.5" hard disk drive option or 8 x 3.5" hard disk drive option
I/O slots	1 PCIe x16 + 4 PCIe x8 + 1 PCIe x4 (all G2)

Option 3: [Dell PowerEdge R510 Rack Mounted Server](#)

The Dell PowerEdge R510 is an ideal choice for scenarios where internal storage is desired and the datacenter has standard racks for mounting servers. The R510 is a two socket Intel system with eight DIMMs and up to 14 internal drives. For this scenario, an eight drive configuration could work well.

Dell PowerEdge R510

Components	Description
Processors (x2)	Latest quad-core or six-core Intel Xeon 5500 and 5600 series processors
Form factor	2U rack
Memory	Up to 128 GB (8 DIMM slots): 1 GB/2 GB/4 GB/8 GB/16 GB DDR3 800 MHz, 1066 MHz, or 1333 MHz
Drives	Up to four cabled 3.5" SAS or SATA drives in 4 hard drive chassis Up to eight hot-swap 2.5"/3.5" SAS, SATA, or solid-state drives (SSDs) in 8 hard drive chassis Up to 12 hot-swap 2.5"/3.5" SAS, SATA, or SSD drives in 12 hard drive chassis with 2 additional 2.5" internal cabled hard drives
I/O slots	4 PCIe G2 slots: One x8 slot Two x4 slot (both with x8 connectors)

Components	Description
	One storage x4 slot (with x8 connector)

Step 3: Select a server model

For this scenario, the choice is the Dell PowerEdge T610 with Intel X5500 processors. The tower form factor is used, because there aren't standard racks in a datacenter in which to mount servers and storage units. The T710 is more expensive due to its extended capabilities in RAM and internal storage, which aren't needed because these servers will be dedicated to hosting Exchange 2010 VMs. The eight internal drives of the T610 are adequate to meet the storage capacity and performance requirements.

[Return to top](#)

Calculate CPU Capacity of Root Server Model

In previous steps, you calculated the megacycles required to support the number of active mailbox users. In the following steps, you determine how many available megacycles the server model and processor can support, to determine the number of active mailboxes each server can support.

Step 1: Determine benchmark value for server and processor

Because the megacycle requirements are based on a baseline server and processor model, you need to adjust the available megacycles for the server against the baseline. To do this, independent performance benchmarks maintained by Standard Performance Evaluation Corporation (SPEC) are used. SPEC is a non-profit corporation formed to establish, maintain, and endorse a standardized set of relevant benchmarks that can be applied to the newest generation of high-performance computers.

To obtain the benchmark value for a server and processor, see [Standard Performance Evaluation Corporation](#), search for the processor, under SPECint_rate2006, find the server model you have chosen, and record the result.

To help simplify the process of obtaining the benchmark value for your server and processor, we recommend you use the [Exchange Processor Query tool](#). This tool automates the manual steps to determine your planned processor's SPECint 2006 Rate Value. To run the tool, your computer must be connected to the Internet. The tool uses your planned processor model as input, and then runs a query against the spec.org site, returning all test result data for that specific processor model. The tool also calculates an average SPECint 2006 Rate Value based on the number of processors planned to be used in each Mailbox server. Use the following calculations:

- Processor and server platform = X5500 on Dell T610
- SPECint_rate2006 value = 241
- SPECint_rate2006 value per processor core = $241 \div 8$
= 30.12

Step 2: Calculate adjusted megacycles

In previous steps, you calculated the required megacycles for the entire environment based on megacycle per mailbox estimates. Those estimates were measured on a baseline system (HP DL380 G5 x5470 3.33 GHz, 8 cores) that has a SPECint_rate2006 value of 150 (for an 8 core server), or 18.75 per core.

In this step, you need to adjust the available megacycles for the chosen server and processor against the baseline processor so that the required megacycles can be used for capacity planning.

To determine the megacycles of the T610 platform, use the following formula:

- Adjusted megacycles per core = (new platform per core value) × (hertz per core of new platform) ÷ (baseline per core value)
= (30.12 × 3330) ÷ 18.75
= 5349
- Adjusted megacycles per server = adjusted megacycles per core × number of cores
= 5349 × 8
= 42794

Step 3: Adjust available megacycles for virtualization overhead

When deploying VMs on the root server, you need to account for megacycles required to support the hypervisor and virtualization stack. This overhead varies from server to server and under different workloads. A conservative estimate of 10 percent of available megacycles will be used. Use the following calculation:

- Adjusted available megacycles = usable megacycles × 0.90
= 42794 × 0.90
= 38515

Each server has a usable capacity for VMs of 38,515 megacycles.

The usable capacity per logical processor is 4,814 megacycles.

[Return to top](#)

Determine CPU Capacity of Virtual Machines

Now that you know the megacycles of the root server, you can calculate the megacycles of each VM. These values will be used to determine how many VMs are required and how many mailboxes will be hosted by each VM.

Step 1: Calculate available megacycles per virtual machine

In this step, you determine how many megacycles are available for each VM deployed on the root server. Because the server has eight logical processors, plan to deploy two VMs per server, each with four virtual processors. Use the following calculation:

- Available megacycles per VM = adjusted available megacycles per server ÷ number of VMs

$$\begin{aligned} &= 38515 \div 2 \\ &= 19257 \end{aligned}$$

Step 2: Determine target available megacycles per virtual machine

Because the design assumptions state not to exceed 80 percent processor utilization, in this step, you adjust the available megacycles to reflect the 80 percent target. Use the following calculation:

- Target available megacycles = available megacycles × target maximum processor utilization
= 19257×0.80
= 15406

[Return to top](#)

Determine Number of Mailbox Server Virtual Machines Required

You can use the following steps to determine the number of Mailbox server VMs required.

Step 1: Determine minimum number of servers required to support the active mailbox count

To determine the minimum number of servers required to support the active mailbox count, use the following calculation:

- Minimum number of servers = required megacycles ÷ available megacycles per Mailbox server role VM
= $1875 \div 15406$
= 0.12

Based on processor capacity, a minimum of one server is required to support the anticipated peak workload during normal operating conditions.

Step 2: Determine number of servers required to support the mailbox resiliency strategy

In a previous step, you determined that the environment needed to handle the simultaneous failure of one server. Use the following calculation:

- Required number of servers = minimum number of servers + number of failed servers allowed
= $1 + 1$
= 2

Based on processor capacity, a minimum of two servers is required to support the anticipated peak workload during a single server maintenance or failure event.

[Return to top](#)

Determine Memory Required per Mailbox Server Virtual Machine

You can use the following steps to determine the memory required for each Mailbox server VM.

Step 1: Determine database cache requirements per server for worst case failure scenario

In a previous step, you determined that the database cache requirements for all mailboxes was 4.4 GB, and the average cache required per active mailbox was 9 MB.

You need to design for the worst case failure scenario, so calculate based on active mailboxes residing on one of two Mailbox servers. Use the following calculation:

- Memory required for database cache = number of active mailboxes × average cache per mailbox
= 500 × 9 MB
= 4500 MB
= 4.4 GB

Step 2: Determine total memory requirements per server for worst case failure scenario

In this step, reference the following table to determine the recommended memory configuration.

Memory requirements

Server physical memory (RAM)	Database cache size (Mailbox server role only)
2 GB	512 MB
4 GB	1 GB
8 GB	3.6 GB
16 GB	10.4 GB
24 GB	17.6 GB

The recommended memory configuration to support 4.4 GB of database cache is 16 GB.

[Return to top](#)

Determine Number of Client Access and Hub Transport Combination Virtual Machines Required

In a previous step, you determined that two Mailbox server VMs were required. We recommend that you deploy one combination Client Access and Hub Transport server VM for every one Mailbox server VM. Therefore, the design will have two Client Access and Hub Transport server VMs.

Server role configuration ratios

Server role configuration	Recommended processor core ratio
Mailbox server role: Client Access and Hub Transport combined server role	1:1

[Return to top](#)

Determine Memory Required per Combined Client Access and Hub Transport Virtual Machines

To determine the memory configuration for the combined Client Access and Hub Transport server role VM, reference the following table.

Memory configurations for Exchange 2010 servers based on installed server roles

Exchange 2010 server role	Minimum supported	Recommended maximum
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 per core


Because the combined Client Access and Hub Transport server role VM has four virtual processors, allocate 8 GB of memory to each combined Client Access and Hub Transport server role VM.

[Return to top](#)

Determine Virtual Machine Distribution

When deciding which VMs to host on which root server, your main goal should be to eliminate single points of failure. Don't locate both Client Access and Hub Transport server role VMs on the same root server, and don't locate both Mailbox server role VMs on the same root server.

Virtual machine distribution (incorrect)



The correct distribution is one Client Access and Hub Transport server role VM on each of the physical host servers and one Mailbox server role VM on each of the physical host servers.

Virtual machine distribution (correct)



[Return to top](#)

Determine Memory Required per Root Server

To determine the memory required for each root server, use the following calculation:

- Root server memory = Client Access and Hub Transport server role VM memory + Mailbox server role VM memory
= 8 GB + 16 GB
= 24 GB

[Return to top](#)

Design Database Availability Groups

In this scenario, there are only two Mailbox servers being deployed. Therefore, only a single DAG is required.

[Return to top](#)

Design Database Copy Layout

You can use the following steps to design a database copy layout.

Step 1: Determine number of unique Exchange databases in the DAG

The easiest way to determine the optimal number of Exchange databases to deploy is to use the Exchange 2010 Mailbox Server Role Requirements Calculator. To download the calculator, see [E2010 Mailbox Server Role Requirements Calculator](#). For additional information about using the storage calculator, see [Exchange 2010 Mailbox Server Role Requirements Calculator](#). Enter the appropriate information on the input worksheet and then select **Yes** for **Automatically Calculate Number of Unique Databases / DAG**.

Database configuration

Database Configuration	Value
Maximum Database Size Configuration	Default
Maximum Database Size (GB)	0
Automatically Calculate Number of Databases / DAG	Yes
Custom Number of Databases / DAG	2

In the Database Configuration table, the recommended number of databases appears. In this scenario, the calculator recommends a minimum of two databases.

Recommended number of databases

Database Configuration	
Number of Databases / DAG	2
Recommended Number of Mailboxes / Database	250

Design Decision Point

With two databases, the size of each database will be 1,378 GB. Because there will be only two copies of each database, you want to ensure that you can seed a failed database copy as quickly as possible. Increasing the database count to four will reduce the database size to 689 GB and allow a reseed to complete in half the time. Increasing the database count from two to four doesn't significantly increase administrative effort.

Step 2: Determine number of database copies per Mailbox server

Because there are four unique databases and two copies of each database, there are a total of eight database copies in the DAG. There are two servers in the DAG, and each server will host eight database copies.

Step 3: Determine database layout during normal operating conditions

For this solution, the database layout is fairly simple. During normal operating conditions, each of the two Mailbox servers will host 50 percent of the active databases and 50 percent of the passive databases.

The database copy layout for normal operating conditions is summarized in the following table.

Database copy layout during normal operating conditions

Database	MBX1	MBX2
DB1	C1	C2
DB2	C1	C2
DB3	C2	C1

Database	MBX1	MBX2
DB4	C2	C1

In the preceding table, the following applies:

- C1 = activation preference value of 1, which is the active copy during normal operations
- C2 = activation preference value of 2, which is the passive copy during normal operations

Step 4: Determine database layout during server failure and maintenance conditions

During a server maintenance scenario or server failure event, all of the active database copies will reside on one of the two Mailbox servers. The passive copy of all databases will reside on the server that was taken down for maintenance or failed.

The database copy layout for a server maintenance or server failure event (impacting MBX2) is summarized in the following table.

Database copy layout during server maintenance or failure

Database	MBX1	MBX2
DB1	C1	C2
DB2	C1	C2
DB3	C2	C1
DB4	C2	C1

In the preceding table, the following applies:

- C1 = activation preference value of 1, which is the active copy during normal operations
- C2 = activation preference value of 2, which is the passive copy during normal operations

[Return to top](#)

Determine Storage Configuration

A well designed storage solution is a critical aspect of a successful Exchange 2010 Mailbox server role deployment. For more information, see Mailbox Server Storage Design.

For planning Exchange 2010 storage configurations, we recommend that you use the [Exchange 2010 Mailbox Server Role Requirements Calculator](#).

Determine disk configuration

This solution uses Windows Server Backup with the VSS plug-in as the backup solution. We recommend that a single LUN per database model be implemented. Because the Dell T610 only

has six disks available for Exchange databases, don't worry about isolating a set of spindles for each database.

Design Decision Point

The decision is to configure the six disks (1 terabyte NL SAS 7,200 rpm) in a single disk group. Four virtual disks of equal size are configured in the disk group. The stripe size is set to 256 KB. In a RAID-5 configuration, each virtual disk contains 1,164 MB of usable space, which equals about 4.5 terabytes of usable space on each server.

With four databases, the LUN size required to support databases and logs is 1,164 GB. The total space required for the four LUNs is 4,654 GB, so capacity requirements align with the actual capacity available.

Disk configuration

Disk space requirements	Database	Server	DAG	Environment
Database space required	827 GB	3,307 GB	6,614 GB	6,614 GB
Log space required	21 GB	86 GB	172 GB	172 GB
Database log LUN space required	1,164 GB	4,654 GB	9,309 GB	9,309 GB

[Return to top](#)

Design Client Access Array and Load Balancing

In Exchange 2010, the remote procedure call (RPC) Client Access service and the Exchange Address Book service were introduced to improve the mailbox users experience when the active mailbox database copy is moved to another Mailbox server (for example, during mailbox database failures and maintenance events). The connection endpoints for mailbox access from Microsoft Outlook and other MAPI clients have been moved from the Mailbox server role to the Client Access server role. Therefore, both internal and external Outlook connections must now be load balanced across all Client Access servers in the site to achieve fault tolerance. To associate the MAPI endpoint with a group of Client Access servers rather than a specific Client Access server, you can define a Client Access server array. You can only configure one array per site, and an array can't span more than one Active Directory site. For more information, see Understanding RPC Client Access and Understanding Load Balancing in Exchange 2010.

Design Decision Point

In a previous step, the decision was to use Windows Network Load Balancing to distribute client load across the two Client Access servers. Because the solution is deployed in a single site, the two servers will be added to a Client Access server array defined for that site. A single internal namespace will be associated with the Client Access server array and a corresponding Domain Name System (DNS) entry that points to the IP address of the Windows Network Load Balancing array. There is no reverse proxy or network address translation (NAT) deployed. Windows

Network Load Balancing will be set up to use source IP affinity as the load balancing configuration for Outlook client traffic.

[Return to top](#)

Determine Placement of the File Share Witness

In Exchange 2010, the DAG uses a minimal set of components from Windows failover clustering. One of those components is the quorum resource, which provides a means for arbitration when determining cluster state and making membership decisions. It is critical that each DAG member have a consistent view of how the DAG's underlying cluster is configured. The quorum acts as the definitive repository for all configuration information relating to the cluster. The quorum is also used as a tiebreaker to avoid *split brain syndrome*. Split brain syndrome is a condition that occurs when DAG members can't communicate with each other but are available and running. Split brain syndrome is prevented by always requiring a majority of the DAG members (and in the case of DAGs with an even number of members, the DAG witness server) to be available and interacting for the DAG to be operational.

A *witness server* is a server outside of a DAG that hosts the file share witness, which is used to achieve and maintain quorum when the DAG has an even number of members. DAGs with an odd number of members don't use a witness server. Upon creation of a DAG, the file share witness is added by default to a Hub Transport server (that doesn't have the Mailbox server role installed) in the same site as the first member of the DAG. If your Hub Transport server is running in a VM that resides on the same root server as VMs running the Mailbox server role, we recommend that you move the location of the file share witness to another highly available server. You can move the file share witness to a domain controller, but because of security implications, do this only as a last resort.

Design Decision Point

There is only one additional server in the environment: a file and print server, which is also used to host Exchange VSS-based backups. The file and print server is reasonably stable and is managed by the same administrator who supports the Exchange servers, so it's a good choice for the location of the file share witness.

[Return to top](#)

Solution Overview

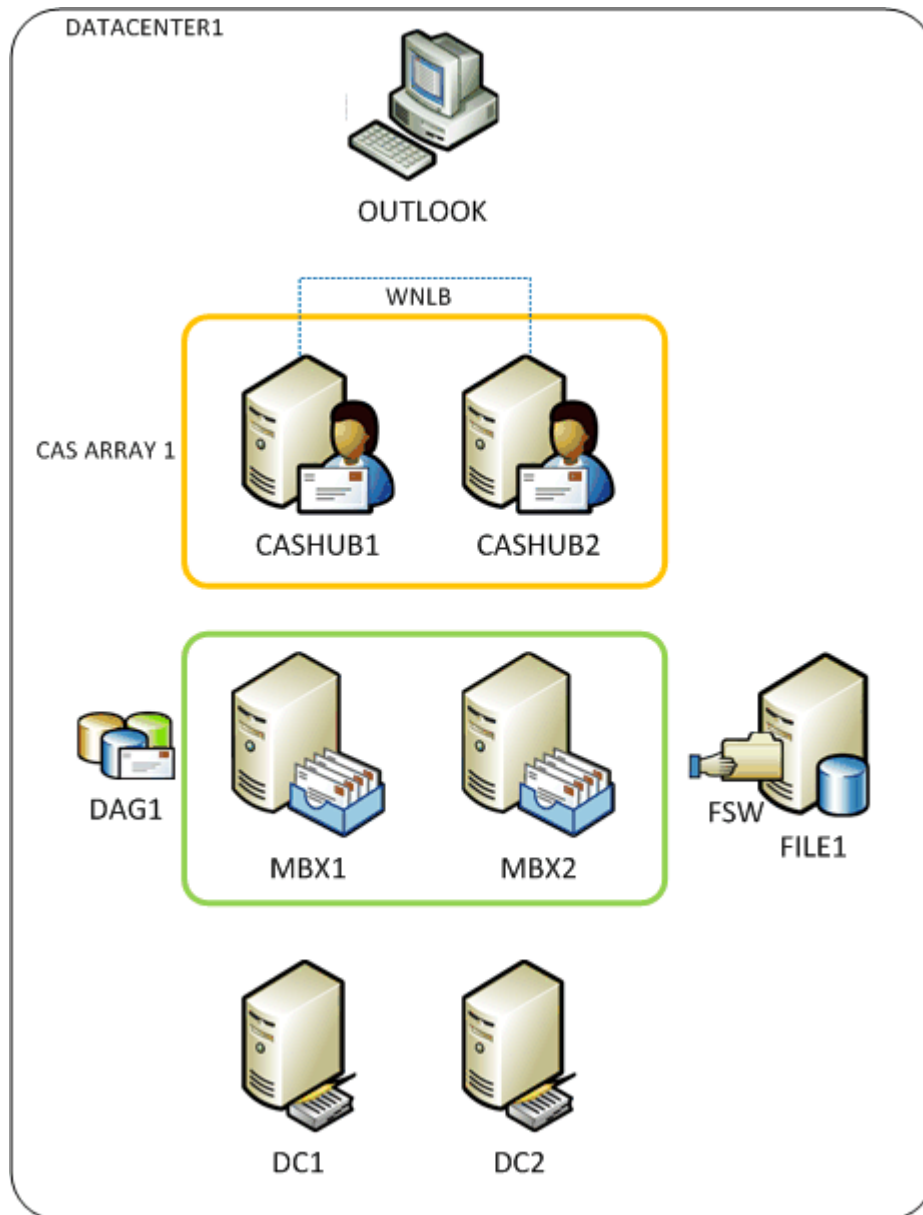
The previous section provided information about the design decisions made by the customer as they completed the Exchange 2010 solution. The following section provides an overview of the selected solution.

Logical Layout

This solution consists of four Exchange VMs located in a single site. Two VMs are running both the Client Access server role and the Hub Transport server role. The Client Access server roles are load balanced using Windows Network Load Balancing running within the VMs. The other two VMs are running the Mailbox server role. The two Mailbox server roles are members of a single

DAG. The file share witness for the DAG is located on a file server used as the storage location for VSS backups (Windows Server Backup with the Exchange 2010 VSS plug-in). The following diagram summarizes the logical layout of this solution.

Logical layout




[Return to top](#)

Physical Layout

The four VMs are deployed on two physical Dell T610 servers. To ensure that the physical server isn't a single point of failure, each server hosts one VM running the Client Access and Hub

Transport server roles and one VM running the Mailbox server role. The DAG has two networks, one public network and one network dedicated for database replication traffic. The following diagram summarizes the physical layout for this solution.

Logical layout



[Return to top](#)

Server Hardware Summary

The following table summarizes the physical server hardware used in this solution.

Server hardware

Component	Description
Server vendor	Dell, Inc
Server model	PowerEdge T610
Processor	Intel Xeon 5550
Chipset	Intel 5520 (Tylersburg)
Memory	32 GB
Operating system	Windows Server 2008 R2
Virtualization	Microsoft Hyper-V
Internal disk	8 x 3.5" 1 terabyte NL SAS
Operating system disk configuration	2 disks in single disk group with 1 virtual disk (RAID-1)
Exchange disk configuration	6 disks in single disk group with 4 virtual disks

Component	Description
	(RAID-5)
RAID controller	Perc H700 with 512 MB battery-backed cache
Network interface	Intel PRO/1000 PT dual port server adapter, gigabit, copper, PCI-E x4
Power	Two hot-plug redundant power supply units

[Return to top](#)

Client Access and Hub Transport Server Configuration

The following table summarizes the combined Client Access and Hub Transport server configuration used in this solution.

Client Access and Hub Transport server configuration

Component	Description
Physical or virtual	Hyper-V VM
Virtual processors	4
Memory	8 GB
Storage	Virtual hard drive on root server operating system volume
Operating system	Windows Server 2008 R2
Exchange version	Exchange Server 2010 Standard Edition
Exchange patch level	Exchange 2010 Update Rollup 2

[Return to top](#)

Mailbox Server Configuration

The following table summarizes the Mailbox server configuration used in this solution.

Mailbox server configuration

Component	Description
Physical or virtual	Hyper-V VM
Virtual processors	4
Memory	8 GB

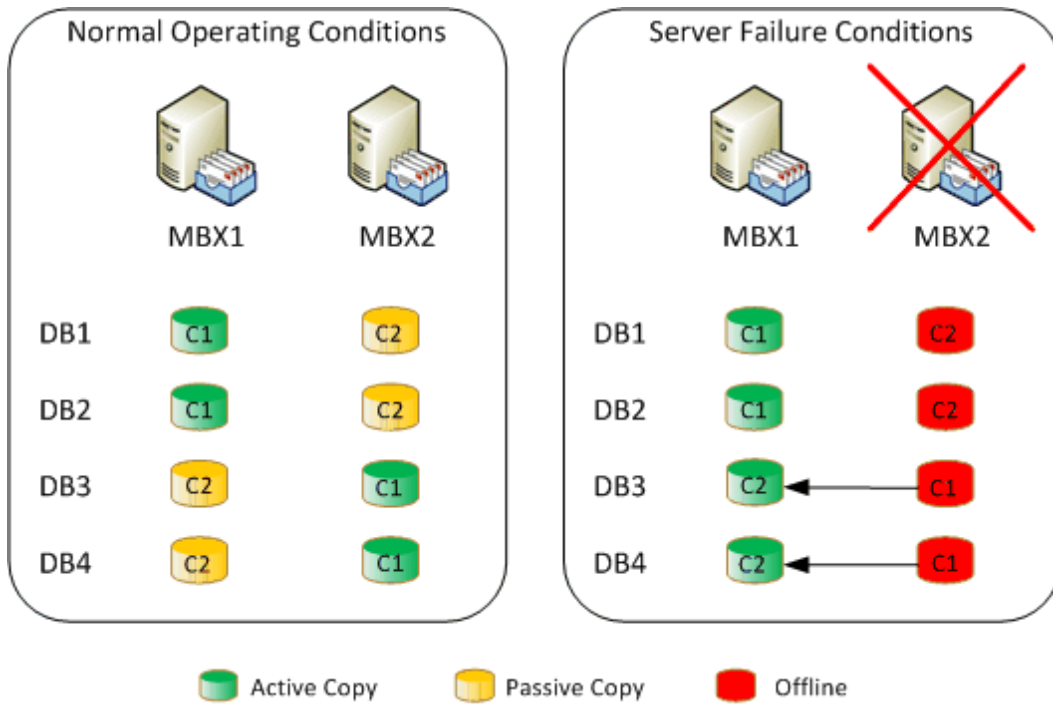
Component	Description
Storage	Virtual hard drive on root server operating system volume
Pass-through storage	4 volumes x 1164 GB
Operating system	Windows Server 2008 R2
Exchange version	Exchange Server 2010 Standard Edition
Exchange patch level	Exchange 2010 Update Rollup 2
Third-party software	None

[Return to top](#)

Database Layout

There are four unique databases and two copies of each database. Each server has two database copies with an activation preference of 1 and two database copies with an activation preference of 2. During normal operating conditions, each server has two active database copies and two passive database copies. During a server failure event or server maintenance event, all four databases are active on the surviving server. The database copy layout for normal operating conditions and single server failure conditions is summarized in the following diagram.

Database layout




[Return to top](#)

Storage Configuration

The Dell T610 has capacity for eight spindles. Disk 0 and Disk 4 are 500-GB SAS and configured in a disk group with a single virtual disk used for the operating system. The other six disks are 1-terabyte SAS and configured in a single disk group with four virtual disks of 1,164 GB each to support the four Exchange mailbox databases. The storage configuration is summarized in the following diagram.

Storage configuration




[Return to top](#)

Solution Variation

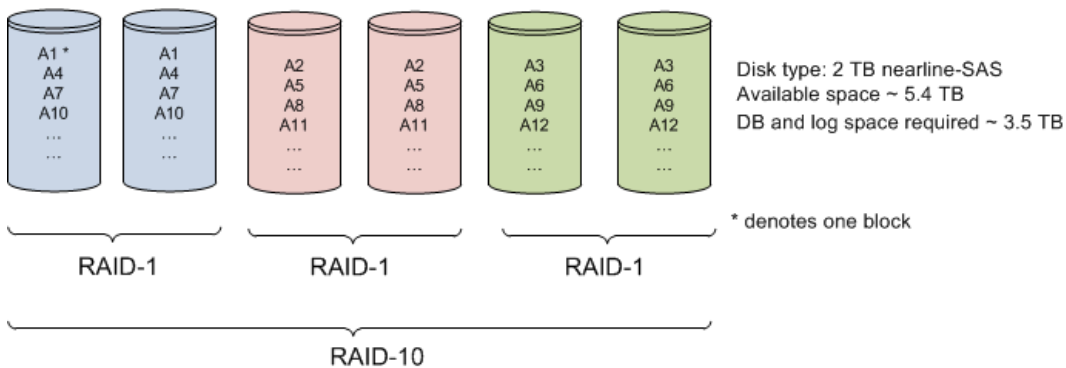
Although using a RAID-5 disk configuration makes sense from a cost perspective, the major downside to this configuration is performance during rebuilds following replacement of a failed disk. This issue can be mitigated by running all active databases on one server while the disk rebuilds on the other server. If your budget allows, we highly recommend that you consider deploying a RAID-10 disk configuration. To get the required capacity for RAID-10, use 2-terabyte NL SAS disks instead of 1-terabyte disks. The following figure shows a comparison of the two configurations.

RAID-5 vs. RAID-10 configuration

RAID-5 configuration



RAID-10 configuration (variation)



[Return to top](#)

Solution Validation Methodology

Prior to deploying an Exchange solution in a production environment, validate that the solution was designed, sized, and configured properly. This validation must include functional testing to ensure that the system is operating as desired as well as performance testing to ensure that the system can handle the desired user load. This section describes the approach and test methodology used to validate server and storage design for this solution. In particular, the following tests will be defined in detail:

- Performance tests
 - Storage performance validation (Jetstress)
 - Server performance validation (Loadgen)
- Functional tests
 - Database switchover validation
 - Server switchover validation
 - Server failover validation

[Return to top](#)

Storage Design Validation Methodology

The level of performance and reliability of the storage subsystem connected to the Exchange Mailbox server role has a significant impact on the overall health of the Exchange deployment. Additionally, poor storage performance will result in high transaction latency, primarily reflected in poor client experience when accessing the Exchange system. To ensure the best possible client experience, validate storage sizing and configuration via the method described in this section.

Tool Set

For validating Exchange storage sizing and configuration, we recommend the Microsoft Exchange Server Jetstress tool. The Jetstress tool is designed to simulate an Exchange I/O workload at the database level by interacting directly with the ESE, which is also known as Jet. The ESE is the database technology that Exchange uses to store messaging data on the Mailbox server role. Jetstress can be configured to test the maximum I/O throughput available to your storage subsystem within the required performance constraints of Exchange. Or, Jetstress can accept a target profile of user count and per-user IOPS, and validate that the storage subsystem is capable of maintaining an acceptable level of performance with the target profile. Test duration is adjustable and can be run for a minimal period of time to validate adequate performance or for an extended period of time to additionally validate storage subsystem reliability.

The Jetstress tool can be obtained from the Microsoft Download Center at the following locations:

- [Microsoft Exchange Server Jetstress 2010 \(64 bit\)](#)
- [Microsoft Exchange Server Jetstress 2010 \(32 bit\)](#)

The documentation included with the Jetstress installer describes how to configure and execute a Jetstress validation test on your server hardware.

Approach to Storage Validation

There are two main types of storage configurations:

- Direct-attached storage (DAS) or internal disk scenarios
- Storage area network (SAN) scenarios

With DAS or internal disk scenarios, there's only one server accessing the disk subsystem, so the performance capabilities of the storage subsystem can be validated in isolation.

In SAN scenarios, the storage utilized by the solution may be shared by many servers and the infrastructure that connects the servers to the storage may also be a shared dependency. This requires additional testing, as the impact of other servers on the shared infrastructure must be adequately simulated to validate performance and functionality.

Test Cases for Storage Validation

The following storage validation test cases were executed against the solution and should be considered as a starting point for storage validation. Specific deployments may have other validation requirements that can be met with additional testing, so this list isn't intended to be exhaustive.

- **Validation of worst case database switchover scenario** In this test case, the level of I/O is expected to be serviced by the storage subsystem in a worst case switchover scenario (largest possible number of active copies on fewest servers). Depending on whether the storage subsystem is DAS or SAN, this test may be required to run on multiple hosts to ensure that the end-to-end solution load on the storage subsystem can be sustained.
- **Validation of storage performance under storage failure and recovery scenario (for example, failed disk replacement and rebuild)** In this test case, the performance of the storage subsystem during a failure and rebuild scenario is evaluated to ensure that the necessary level of performance is maintained for optimal Exchange client experience. The same caveat applies for a DAS vs. SAN deployment: If multiple hosts are dependent on a shared storage subsystem, the test must include load from these hosts to simulate the entire effect of the failure and rebuild.

Analyzing the Results

The Jetstress tool produces a report file after each test is completed. To help you analyze the report, use the guidelines in Jetstress 2010 Test Summary Reports.

Specifically, you should use the guidelines in the following table when you examine data in the Test Results table of the report.

Jetstress results analysis

Performance counter instance	Guidelines for performance test
I/O Database Reads Average Latency (msec)	The average value should be less than 20 milliseconds (msec) (0.020 seconds), and the maximum values should be less than 50 msec.
I/O Log Writes Average Latency (msec)	Log disk writes are sequential, so average write latencies should be less than 10 msec, with a maximum of no more than 50 msec.
%Processor Time	Average should be less than 80%, and the maximum should be less than 90%.
Transition Pages Repurposed/sec (Windows Server 2003, Windows Server 2008, Windows Server 2008 R2)	Average should be less than 100.

The report file shows various categories of I/O performed by the Exchange system:

- **Transactional I/O Performance** This table reports I/O that represents user activity against the database (for example, Outlook generated I/O). This data is generated by subtracting background maintenance I/O and log replication I/O from the total I/O measured during the test. This data provides the actual database IOPS generated along with I/O latency measurements required to determine whether a Jetstress performance test passed or failed.
- **Background Database Maintenance I/O Performance** This table reports the I/O generated due to ongoing ESE database background maintenance.

- **Log Replication I/O Performance** This table reports the I/O generated from simulated log replication.
- **Total I/O Performance** This table reports the total I/O generated during the Jetstress test.

[Return to top](#)

Server Design Validation

After the performance and reliability of the storage subsystem is validated, ensure that all of the components in the messaging system are validated together for functionality, performance, and scalability. This means moving up in the stack to validate client software interaction with the Exchange product as well as any server-side products that interact with Exchange. To ensure that the end-to-end client experience is acceptable and that the entire solution can sustain the desired user load, the method described in this section can be applied for server design validation.

Tool Set

For validation of end-to-end solution performance and scalability, we recommend the Microsoft Exchange Server Load Generator tool (Loadgen). Loadgen is designed to produce a simulated client workload against an Exchange deployment. This workload can be used to evaluate the performance of the Exchange system, and can also be used to evaluate the effect of various configuration changes on the overall solution while the system is under load. Loadgen is capable of simulating Microsoft Office Outlook 2007 (online and cached), Office Outlook 2003 (online and cached), POP3, IMAP4, SMTP, ActiveSync, and Outlook Web App (known in Exchange 2007 and earlier versions as Outlook Web Access) client activity. It can be used to generate a single protocol workload, or these client protocols can be combined to generate a multiple protocol workload.

You can get the Loadgen tool from the Microsoft Download Center at the following locations:

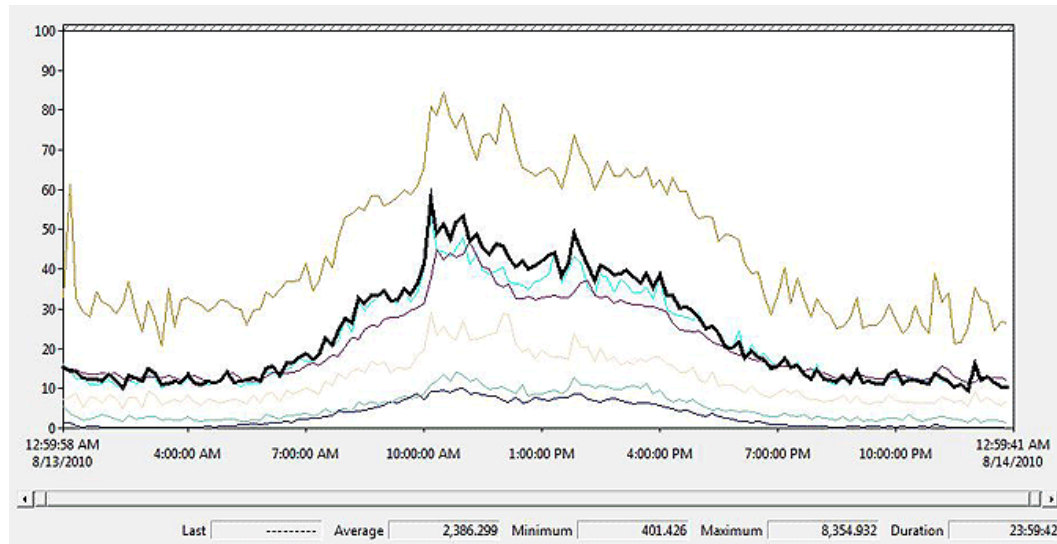
- [Exchange Load Generator 2010 \(64 bit\)](#)
- [Exchange Load Generator 2010 \(32 bit\)](#)

The documentation included with the Loadgen installer describes how to configure and execute a Loadgen test against an Exchange deployment.

Approach to Server Validation

When validating your server design, test the worst case scenario under anticipated peak workload. Based on a number of data sets from Microsoft IT and other customers, peak load is generally equal to 2x the average workload throughout the remainder of the work day. This is referred to as the peak-to-average workload ratio.

Peak load



In this Performance Monitor snapshot, which displays various counters that represent the amount of Exchange work being performed over time on a production Mailbox server, the average value for RPC operations per second (the highlighted line) is about 2,386 when averaged across the entire day. The average for this counter during the peak period from 10:00 through 11:00 is about 4,971, giving a peak-to-average ratio of 2.08.

To ensure that the Exchange solution is capable of sustaining the workload generated during the peak average, modify Loadgen settings to generate a constant amount of load at the peak average level, rather than spreading out the workload over the entire simulated work day. Loadgen task-based simulation modules (like the Outlook simulation modules) utilize a task profile that defines the number of times each task will occur for an average user within a simulated day.

The total number of tasks that need to run during a simulated day is calculated as the number of users multiplied by the sum of task counts in the configured task profile. Loadgen then determines the rate at which it should run tasks for the configured set of users by dividing the total number of tasks to run in the simulated day by the simulated day length. For example, if Loadgen needs to run 1,000,000 tasks in a simulated day, and a simulated day is equal to 8 hours (28,800 seconds), Loadgen must run $1,000,000 \div 28,800 = 34.72$ tasks per second to meet the required workload definition. To increase the amount of load to the desired peak average, divide the default simulated day length (8 hours) by the peak-to-average ratio (2) and use this as the new simulated day length.

Using the task rate example again, $1,000,000 \div 14,400 = 69.44$ tasks per second. This reduces the simulated day length by half, which results in doubling the actual workload run against the server and achieving our goal of a peak average workload. You don't adjust the run length duration of the test in the Loadgen configuration. The run length duration specifies the duration of the test and doesn't affect the rate at which tasks will be run against the Exchange server.

Test Cases for Server Design Validation

The following server design validation test cases were executed against the solution and should be considered as a starting point for server design validation. Specific deployments may have other validation requirements that can be met with additional testing, so this list isn't intended to be exhaustive:

- **Normal operating conditions** In this test case, the basic design of the solution is validated with all components in their normal operating state (no failures simulated). The desired workload is generated against the solution, and the overall performance of the solution is validated against the metrics that follow.
- **Single server failure or single server maintenance (in site)** In this test case, a single server is taken down to simulate either an unexpected failure of the server or a planned maintenance operation for the server. The workload that would normally be handled by the unavailable server is now handled by other servers in the solution topology, and the overall performance of the solution is validated.

Test Execution and Data Collection

Exchange performance data has some natural variation within test runs and among test runs. We recommend that you take the average of multiple runs to smooth out this variation. For Exchange tested solutions, a minimum of three separate test runs with durations of eight hours was completed. Performance data was collected for the full eight-hour duration of the test.

Performance summary data was taken from a three to four hour stable period (excluding the first two hours of the test and the last hour of the test). For each Exchange server role, performance summary data was averaged between servers for each test run, providing a single average value for each data point. The values for each run were then averaged, providing a single data point for all servers of a like server role across all test runs.

Validation of Expected Load

Before you look at any performance counters or start your performance validation analysis, verify that the workload you expected to run matched the workload that you actually ran. Although there are many ways to determine whether the simulated workload matched the expected workload, the easiest and most consistent way is to look at the message delivery rate.

Calculating Expected Peak Message Delivery Rate

Every message profile consists of the sum of the average number of messages sent per day and the average number of messages received per day. To calculate the message delivery rate, select the average number of messages received per day from the following table.

Peak message delivery rate

Message profile	Messages sent per day	Messages received per day
50	10	40
100	20	80

Message profile	Messages sent per day	Messages received per day
150	30	120
200	40	160

This example assumes that each Mailbox server has 5,000 active mailboxes with a 120 messages per day profile, as shown in the following table.

Peak message delivery rate for 5,000 active mailboxes

Description	Calculation	Value
Message profile	Number of messages received per day	120
Mailbox server profile	Number of active mailboxes per Mailbox server	5000
Total messages received per day per Mailbox server	5000×120	600000
Total messages received per second per Mailbox server	$600000 \div 28800$	20.83
Total messages adjusted for peak load	20.83×2	41.67

You expect 41.67 messages per second delivered on each Mailbox server running 5,000 active mailboxes with a message profile of 120 messages per day during peak load.

Measuring Actual Message Delivery Rate

The actual message delivery rate can be measured using the following counter on each Mailbox server: `MSExchangeIS Mailbox(_Total)\Messages Delivered/sec`. If the measured message delivery rate is within one or two messages per second of the target message delivery rate, you can be confident that the desired load profile was run successfully.

Server Validation: Performance and Health Criteria

This section describes the Performance Monitor counters and thresholds used to determine whether the Exchange environment was sized properly and is able to run in a healthy state during extended periods of peak workload. For more information about counters relevant to Exchange performance, see [Performance and Scalability Counters and Thresholds](#).

Hyper-V Root Servers

To validate the performance and health criteria of a Hyper-V root server and the applications running within VMs, you should have a basic understanding of the Hyper-V architecture and how that impacts performance monitoring.

Hyper-V has three main components: the virtualization stack, the hypervisor, and devices. The virtualization stack handles emulated devices, manages VMs, and services I/O. The hypervisor schedules virtual processors, manages interrupts, services timers, and controls other chip-level functions. The hypervisor doesn't handle devices or I/O (for example, there are no hypervisor drivers). The devices are part of the root server or installed in guest servers as part of integration services. Because the root server has a full view of the system and controls the VMs, it also provides monitoring information via Windows Management Instrumentation (WMI) and performance counters.

Processor

When validating physical processor utilization on the root server (or within the guest VM), the standard Processor\% Processor Time counter isn't very useful.

Instead, you can examine the Hyper-V Hypervisor Logical Processor\% Total Run Time counter. This counter shows the percentage of processor time spent in guest and hypervisor runtime and should be used to measure the total processor utilization for the hypervisor and all VMs running on the root server. This counter shouldn't exceed 80 percent or whatever the maximum utilization target you have designed for.

Counter	Target
Hyper-V Hypervisor Logical Processor\% Total Run Time	<80%

If you're interested in what percentage of processor time is spent servicing the guest VMs, you can examine the Hyper-V Hypervisor Logical Processor\% Guest Run Time counter. If you're interested in what percentage of processor time is spent in hypervisor, you can look at the Hyper-V Hypervisor Logical Processor\% Hypervisor Run Time counter. This counter should be below 5 percent. The Hyper-V Hypervisor Root Virtual Processor\% Guest Run Time counter shows the percentage of processor time spent in the virtualization stack. This counter should also be below 5 percent. These two counters can be used to determine what percentage of your available physical processor time is being used to support virtualization.

Counter	Target
Hyper-V Hypervisor Logical Processor\% Guest Run Time	<80%
Hyper-V Hypervisor Logical Processor\% Hypervisor Run Time	<5%
Hyper-V Hypervisor Root Virtual Processor\% Guest Run Time	<5%

Memory

You need to ensure that your Hyper-V root server has enough memory to support the memory allocated to VMs. Hyper-V automatically reserves 512 MB (this may vary with different Hyper-V releases) for the root operating system. If you don't have enough memory, Hyper-V will prevent the last VM from starting. In general, don't worry about validating the memory on a Hyper-V root server. Be more concerned with ensuring that sufficient memory is allocated to the VMs to support the Exchange roles.

Application Health

An easy way to determine whether all the VMs are in a healthy state is to look at the Hyper-V Virtual Machine Health Summary counters.

Counter	Target
Hyper-V Virtual Machine Health Summary\Health OK	1
Hyper-V Virtual Machine Health Summary\Health Critical	0

Mailbox Servers

When validating whether a Mailbox server was properly sized, focus on processor, memory, storage, and Exchange application health. This section describes the approach to validating each of these components.

Processor

During the design process, you calculated the adjusted megacycle capacity of the server or processor platform. You then determined the maximum number of active mailboxes that could be supported by the server without exceeding 80 percent of the available megacycle capacity. You also determined what the projected CPU utilization should be during normal operating conditions and during various server maintenance or failure scenarios.

During the validation process, verify that the worst case scenario workload doesn't exceed 80 percent of the available megacycles. Also, verify that actual CPU utilization is close to the expected CPU utilization during normal operating conditions and during various server maintenance or failure scenarios.

For physical Exchange deployments, use the Processor(_Total)\% Processor Time counter and verify that this counter is less than 80 percent on average.

Counter	Target
Processor(_Total)\% Processor Time	<80%

For virtual Exchange deployments, the Processor(_Total)\% Processor Time counter is measured within the VM. In this case, the counter isn't measuring the physical CPU utilization. It's measuring the utilization of the virtual CPU provided by the hypervisor. Therefore, it doesn't provide an accurate reading of the physical processor and shouldn't be used for design validation

purposes. For more information, see [Hyper-V: Clocks lie... which performance counters can you trust](#).

For validating Exchange deployments running on Microsoft Hyper-V, use the Hyper-V Hypervisor Virtual Processor\% Guest Run Time counter. This provides a more accurate value for the amount of physical CPU being utilized by the guest operating system. This counter should be less than 80 percent on average.

Counter	Target
Hyper-V Hypervisor Virtual Processor\% Guest Run Time	<80%

Memory

During the design process, you calculated the amount of database cache required to support the maximum number of active databases on each Mailbox server. You then determined the optimal physical memory configuration to support the database cache and system memory requirements.

Validating whether an Exchange Mailbox server has sufficient memory to support the target workload isn't a simple task. Using available memory counters to view how much physical memory is remaining isn't helpful because the memory manager in Exchange is designed to use almost all of the available physical memory. The information store (store.exe) reserves a large portion of physical memory for database cache. The database cache is used to store database pages in memory. When a page is accessed in memory, the information doesn't have to be retrieved from disk, reducing read I/O. The database cache is also used to optimize write I/O.

When a database page is modified (known as a *dirty page*), the page stays in cache for a period of time. The longer it stays in cache, the better the chance that the page will be modified multiple times before those changes are written to the disk. Keeping dirty pages in cache also causes multiple pages to be written to the disk in the same operation (known as *write coalescing*). Exchange uses as much of the available memory in the system as possible, which is why there aren't large amounts of available memory on an Exchange Mailbox server.

It may not be easy to know whether the memory configuration on your Exchange Mailbox server is undersized. For the most part, the Mailbox server will still function, but your I/O profile may be much higher than expected. Higher I/O can lead to higher disk read and write latencies, which may impact application health and client user experience. In the results section, there isn't any reference to memory counters. Potential memory issues will be identified in the storage validation and application health result sections, where memory-related issues are more easily detected.

Storage

If you have performance issues with your Exchange Mailbox server, those issues may be storage-related issues. Storage issues may be caused by having an insufficient number of disks to support the target I/O requirements, having overloaded or poorly designed storage connectivity infrastructure, or by factors that change the target I/O profile like insufficient memory, as discussed previously.

The first step in storage validation is to verify that the database latencies are below the target thresholds. In previous releases, logical disk counters determined disk read and write latency. In Exchange 2010, the Exchange Mailbox server that you are monitoring is likely to have a mix of active and passive mailbox database copies. The I/O characteristics of active and passive database copies are different. Because the size of the I/O is much larger on passive copies, there are typically much higher latencies on passive copies. Latency targets for passive databases are 200 msec, which is 10 times higher than targets on active database copies. This isn't much of a concern because high latencies on passive databases have no impact on client experience. But if you are using the traditional logical disk counters to measure latencies, you must review the individual volumes and separate volumes containing active and passive databases. Instead, we recommend that you use the new MExchange Database counters in Exchange 2010.

When validating latencies on Exchange 2010 Mailbox servers, we recommend you use the counters in the following table for active databases.

Counter	Target
MExchange Database\I/O Database Reads (Attached) Average Latency	<20 msec
MExchange Database\I/O Database Writes (Attached) Average Latency	<20 msec
MExchange Database\IO Log Writes Average Latency	<1 msec

We recommend that you use the counters in the following table for passive databases.

Counter	Target
MExchange Database\I/O Database Reads (Recovery) Average Latency	<200 msec
MExchange Database\I/O Database Writes (Recovery) Average Latency	<200 msec
MExchange Database\IO Log Read Average Latency	<200 msec

 **Note:**

To view these counters in Performance Monitor, you must enable the advanced database counters. For more information, see [How to Enable Extended ESE Performance Counters](#).

When you're validating disk latencies for Exchange deployments running on Microsoft Hyper-V, be aware that the I/O Database Average Latency counters (as with many time-based counters) may not be accurate because the concept of time within the VM is different than on the physical

server. The following example shows that the I/O Database Reads (Attached) Average Latency is 22.8 in the VM and 17.3 on a physical server for the same simulated workload. If the values of time-based counters are over the target thresholds, your server may be running correctly. Review all health criteria to make a decision regarding server health when your Mailbox server role is deployed within a VM.

Virtual Mailbox server

Counter	Target
MSExchange Database	Information Store
I/O Database Reads (Attached) / Average Latency	22.792
I/O Database Reads (Attached) / sec	17.693
I/O Database Reads (Recovery) / Average Latency	34.215
I/O Database Writes (Recovery) / sec	10.829
I/O Database Writes (Attached) / Average Latency	0.944
I/O Database Writes (Attached) / sec	10.184
MSExchangeIS	
RPC Averaged Latency	1.966
RPC Operations / sec	334.371
RPC Packets / sec	180.656
MSExchangeIS Mailbox	_Total
Messages Delivered / sec	2.062
Messages Sent / sec	0.511

Physical Mailbox server

Counter	Target
MSExchange Database	Information Store
I/O Database Reads (Attached) / Average Latency	17.250
I/O Database Reads (Attached) / sec	18.131
I/O Database Reads (Recovery) / Average Latency	27.758

Counter	Target
I/O Database Writes (Recovery) / sec	8.483
I/O Database Writes (Attached) / Average Latency	0.411
I/O Database Writes (Attached) / sec	10.963
MSExchangeIS	
RPC Averaged Latency	1.695
RPC Operations / sec	341.139
RPC Packets / sec	183.360
MSExchangeIS Mailbox	
	_Total
Messages Delivered / sec	2.065
Messages Sent / sec	0.514

Virtual Mailbox server and Physical Mailbox server

Counter	Virtual Mailbox server	Physical Mailbox server
MSExchange Database/		
I/O Database Reads (Attached) / Average Latency	22.792	17.250
I/O Database Reads (Attached) / sec	17.693	18.131
I/O Database Reads (Recovery) / Average Latency	34.215	27.758
I/O Database Writes (Recovery) / sec	10.829	8.483
I/O Database Writes (Attached) / Average Latency	0.944	0.411
I/O Database Writes (Attached) / sec	10.184	10.963
MSExchangeIS		
RPC Averaged Latency	1.966	1.695
RPC Operations / sec	334.371	341.139
RPC Packets / sec	180.656	183.360

Counter	Virtual Mailbox server	Physical Mailbox server
MSExchangeIS Mailbox		
Messages Delivered / sec	2.062	2.065
Messages Sent / sec	0.511	0.514

In addition to disk latencies, review the Database\Database Page Fault Stalls/sec counter. This counter indicates the rate of page faults that can't be serviced because there are no pages available for allocation from the database cache. This counter should be 0 on a healthy server.

Counter	Target
Database\Database Page Fault Stalls/sec	<1

Also, review the Database\Log Record Stalls/sec counter, which indicates the number of log records that can't be added to the log buffers per second because the log buffers are full. This counter should average less than 10.

Counter	Target
Database\Log Record Stalls/sec	<10

Exchange Application Health

Even if there are no obvious issues with processor, memory, and disk, we recommend that you monitor the standard application health counters to ensure that the Exchange Mailbox server is in a healthy state.

The MSExchangeIS\RPC Averaged Latency counter provides the best indication of whether other counters with high database latencies are actually impacting Exchange health and client experience. Often, high RPC averaged latencies are associated with a high number of RPC requests, which should be less than 70 at all times.

Counter	Target
MSExchangeIS\RPC Averaged Latency	<10 msec on average
MSExchangeIS\RPC Requests	<70 at all times

Next, make sure that the transport layer is healthy. Any issues in transport or issues downstream of transport affecting the transport layer can be detected with the MSExchangeIS Mailbox(_Total)\Messages Queued for Submission counter. This counter should be less than 50 at all times. There may be temporary increases in this counter, but the counter value shouldn't grow over time and shouldn't be sustained for more than 15 minutes.

Counter	Target
MSExchangeIS Mailbox(_Total)\Messages Queued for Submission	<50 at all times

Next, ensure that maintenance of the database copies is in a healthy state. Any issues with log shipping or log replay can be identified using the MSExchange Replication(*)\CopyQueueLength and MSExchange Replication(*)\ReplayQueueLength counters. The copy queue length shows the number of transaction log files waiting to be copied to the passive copy log file folder and should be less than 1 at all times. The replay queue length shows the number of transaction log files waiting to be replayed into the passive copy and should be less than 5. Higher values don't impact client experience, but result in longer store mount times when a handoff, failover, or activation is performed.

Counter	Target
MSExchange Replication(*)\CopyQueueLength	<1
MSExchange Replication(*)\ReplayQueueLength	<5

Client Access Servers

To determine whether a Client Access server is healthy, review processor, memory, and application health. For an extended list of important counters, see Client Access Server Counters.

Processor

For physical Exchange deployments, use the Processor(_Total)\% Processor Time counter. This counter should be less than 80 percent on average.

Counter	Target
Processor(_Total)\% Processor Time	<80%

For validating Exchange deployments running on Microsoft Hyper-V, use the Hyper-V Hypervisor Virtual Processor\% Guest Run Time counter. This provides an accurate value for the amount of physical CPU being utilized by the guest operating system. This counter should be less than 80 percent on average.

Counter	Target
Hyper-V Hypervisor Virtual Processor\% Guest Run Time	<80%

Application Health

To determine whether the MAPI client experience is acceptable, use the MExchange RpcClientAccess\RPC Averaged Latency counter. This counter should be below 250 msec. High latencies can be associated with a large number of RPC requests. The MExchange RpcClientAccess\RPC Requests counter should be below 40 on average.

Counter	Target
MExchange RpcClientAccess\RPC Averaged Latency	<250 msec
MExchange RpcClientAccess\RPC Requests	<40

Transport Servers

To determine whether a transport server is healthy, review processor, disk, and application health. For an extended list of important counters, see Transport Server Counters.

Processor

For physical Exchange deployments, use the Processor(_Total)\% Processor Time counter. This counter should be less than 80 percent on average.

Counter	Target
Processor(_Total)\% Processor Time	<80%

For validating Exchange deployments running on Microsoft Hyper-V, use the Hyper-V Hypervisor Virtual Processor\% Guest Run Time counter. This provides an accurate value for the amount of physical CPU being utilized by the guest operating system. This counter should be less than 80 percent on average.

Counter	Target
Hyper-V Hypervisor Virtual Processor\% Guest Run Time	<80%

Disk

To determine whether disk performance is acceptable, use the Logical Disk(*)\Avg. Disk sec/Read and Write counters for the volumes containing the transport logs and database. Both of these counters should be less than 20 msec.

Counter	Target
Logical Disk(*)\Avg. Disk sec/Read	<20 msec
Logical Disk(*)\Avg. Disk sec/Write	<20 msec

Application Health

To determine whether a Hub Transport server is sized properly and running in a healthy state, examine the MExchangeTransport Queues counters outlined in the following table. All of these queues will have messages at various times. You want to ensure that the queue length isn't sustained and growing over a period of time. If larger queue lengths occur, this could indicate an overloaded Hub Transport server. Or, there may be network issues or an overloaded Mailbox server that's unable to receive new messages. You will need to check other components of the Exchange environment to verify.

Counter	Target
MExchangeTransport Queues(_total)\Aggregate Delivery	<3000
MExchangeTransport Queues(_total)\Active Remote Delivery Queue Length	<250
MExchangeTransport Queues(_total)\Active Mailbox Delivery Queue Length	<250
MExchangeTransport Queues(_total)\Retry Mailbox Delivery Queue Length	<100
MExchangeTransport Queues(_total)\Submission Queue Length	<100

[Return to top](#)

Functional Validation Tests

You can use the information in the following sections for functional validation tests.

Database Switchover Validation

A database switchover is the process by which an individual active database is switched over to another database copy (a passive copy), and that database copy is made the new active database copy. Database switchovers can happen both within and across datacenters. A database switchover can be performed by using the Exchange Management Console (EMC) or the Exchange Management Shell.

To validate that a passive copy of a database can be successfully activated on another server, run the following command.

```
Move-ActiveMailboxDatabase <DatabaseName> -ActivateOnServer  
<TargetServer>
```

Success criteria: The active mailbox database is mounted on the specified target server. This result can be confirmed by running the following command.

```
Get-MailboxDatabaseCopyStatus <DatabaseName>
```

Server Switchover Validation

A server switchover is the process by which all active databases on a DAG member are activated on one or more other DAG members. Like database switchovers, a server switchover can occur both within a datacenter and across datacenters, and it can be initiated by using both the EMC and the Shell.

1. To validate that all passive copies of databases on a server can be successfully activated on other servers hosting a passive copy, run the following command.

```
Get-MailboxDatabase -Server <ActiveMailboxServer> | Move-ActiveMailboxDatabase -ActivateOnServer <TargetServer>
```

Success criteria: The active mailbox databases are mounted on the specified target server. This can be confirmed by running the following command.

```
Get-MailboxDatabaseCopyStatus <DatabaseName>
```

2. To validate that one copy of each of the active databases will be successfully activated on another Mailbox server hosting passive copies of the databases, shut down the server by performing the following action.

Turn off the current active server.

Success criteria: The active mailbox databases are mounted on another Mailbox server in the DAG. This can be confirmed by running the following command.

```
Get-MailboxDatabaseCopyStatus <DatabaseName>
```

Server Failover Validation

A server failover occurs when the DAG member can no longer service the MAPI network, or when the Cluster service on a DAG member can no longer contact the remaining DAG members.

To validate that one copy of each of the active databases will be successfully activated on another Mailbox server hosting passive copies of the databases, turn off the server by performing one of the following actions:

- Press and hold the power button on the server until the server turns off.
- Pull the power cables from the server, which results in the server turning off.

Success criteria: The active mailbox databases are mounted on another Mailbox server in the DAG. This can be confirmed by running the following command.

```
Get-MailboxDatabase -Server <MailboxServer> | Get-MailboxDatabaseCopyStatus
```

[Return to top](#)

Test Facility

Testing was conducted at the [Microsoft Enterprise Engineering Center](#), a state-of-the-art enterprise solutions validation laboratory on the Microsoft main campus in Redmond, Washington.

With more than 125 million dollars in hardware and with ongoing strong partnerships with the industry's leading original equipment manufacturers (OEMs), virtually any production environment can be replicated at the EEC. The EEC offers an environment that enables extensive collaboration among customers, partners, and Microsoft product engineers. This helps ensure that Microsoft end-to-end solutions will meet the high expectations of customers.

[Return to top](#)

Solution Validation Results

The following section summarizes the results of the functional and performance validation tests.

Functional Validation Results

The following table summarizes the functional validation test results.

Functional validation results

Test case	Result	Comments
Database switchover	Successful	Completed without errors
Server switchover	Successful	Completed without errors
Server failure	Successful	Completed without errors

[Return to top](#)

Storage Design Validation Results

The following tables summarize the Jetstress storage validation results. This solution achieved higher than target transactional I/O while maintaining database latencies well under the 20 msec target.

Overall test result	Pass
---------------------	------

Total transactional I/O

Total transactional I/O	Result
Target transactional I/O per second	90
Achieved transactional I/O per second	133

Transactional I/O performance: database reads

Database	I/O database reads per second	I/O database reads average latency
DB1	19.5	12.3
DB2	20.0	11.0
DB3	20.1	11.3
DB4	20.0	14.1

Transactional I/O performance: database writes

Database	I/O database writes per second	I/O database writes average latency
DB1	13.3	1.1
DB2	13.7	1.1
DB3	13.7	1.1
DB4	13.5	1.1

Transactional I/O performance: log writes

Database	I/O log writes per second	I/O database writes average latency
DB1	13.2	0.4
DB2	13.4	0.4
DB3	13.3	0.4
DB4	13.3	0.4

[Return to top](#)

Server Design Validation Results

The following sections summarize the server design validation results for the test cases.

Test Case: Normal Operating Conditions

The first test case represents peak workload during normal operating conditions. Normal operating conditions refer to a state where all of the active and passive databases reside on the servers they were planned to run on. Because this test case doesn't represent the worst case workload, it isn't the key performance validation test. It provides a good indication of how this environment should run outside of a server failure or maintenance event. In this case, each

Mailbox server is running two active and two passive databases. Test results are provided for both the planned message profile of 150 as well as a lower message profile of 100.

Validation of Expected Load

The message delivery rate verifies that tested workload matched the target workload. For both the 150 and 100 message profiles, the actual message delivery rate is on target.

Counter	Target	Tested result (150 messages per day)	Target	Tested result (100 messages per day)
Message Delivery Rate / Mailbox Server	2.08	2.07	1.39	1.38

Validation of Mailbox Servers

The following tables show the validation of Mailbox servers.

Processor

Processor utilization is low, as expected.

Counter	Target	Tested result (150 messages per day)	Tested result (100 messages per day)
Hyper-V Hypervisor Virtual Processor\% Guest Run Time	<70%	13.2	10.1

Storage

The storage results are good. The average read latency for the active databases is 20.4, which is at the top end of the target. This indicates that you wouldn't want to run more mailbox users or a higher message profile on the current storage solution.

Counter	Target	Tested result (150 messages per day)	Tested result (100 messages per day)
MSEExchange Database\I/O Database Reads (Attached) Average Latency	<20 msec	20.4	18.6
MSEExchange Database\I/O Database Writes (Attached) Average Latency	<20 msec <Reads average	0.8	0.8
Database\Database Page Fault Stalls/sec	0	0	0

Counter	Target	Tested result (150 messages per day)	Tested result (100 messages per day)
MSExchange Database\IO Log Writes Average Latency	<20 msec	0.4	<1
Database\Log Record Stalls/sec	0	0	0
MSExchange Database\IO Database Reads (Recovery) Average Latency	<200 msec	3.6	10.5
MSExchange Database\IO Database Writes (Recovery) Average Latency	<200 msec	1.2	2.6
MSExchange Database\IO Log Read Average Latency	<200 msec	0.5	<1

Application Health

Exchange is healthy, and all the counters used to determine application health are well under target values.

Counter	Target	Tested result (150 messages per day)	Tested result (100 messages per day)
MSExchange\IS\RPC Requests	<70	0.3	0.2
MSExchange\IS\RPC Averaged Latency	<10 msec	1.4	1.3
MSExchange\IS\Mailbox(_Total)\Messages Queued for Submission	0	0.4	0.3
MSExchange\Replication(*)\CopyQueueLength	<1	<1	<1
MSExchange\Replication(*)\ReplayQueueLength	<5	1.6	1.3

Validation of Client Access and Hub Transport Servers

The following tables show the validation of the Client Access and Hub Transport servers.

Processor

Processor utilization is low, as expected.

Counter	Target	Tested result (150 messages per day)	Tested result (100 messages per day)
Hyper-V Hypervisor Virtual Processor\% Guest Run Time	<70%	6.3	4.8

Storage

The storage results look good. The very low latencies should have no impact on message transport.

Counter	Target	Tested result (150 messages per day)	Tested result (100 messages per day)
Logical/Physical Disk(*)\Avg. Disk sec/Read	<20 msec	0.001	0.002
Logical/Physical Disk(*)\Avg. Disk sec/Write	<20 msec	0.000	0.001

Application Health

The low RPC Averaged Latency values confirm a healthy Client Access server with no impact on client experience.

Counter	Target	Tested result (150 messages per day)	Tested result (100 messages per day)
MSExchange RpcClientAccess\RPC Averaged Latency	<250 msec	2.75	2.5
MSExchange RpcClientAccess\RPC Requests	<40	0.3	0.2

Hub Transport Server Health

The Transport Queue counters are all well under target, confirming that the Hub Transport server is healthy, and able to process and deliver the required messages.

Counter	Target	Tested result (150 messages per day)	Tested result (100 messages per day)
MSExchangeTransport Queues(_total)\Aggregate Delivery Queue Length (All Queues)	<3000	<1	<1
MSExchangeTransport Queues(_total)\Active Remote Delivery Queue Length	<250	0	0
MSExchangeTransport Queues(_total)\Active Mailbox Delivery Queue Length	<250	<1	<1
MSExchangeTransport Queues(_total)\Submission Queue Length	<100	0	0
MSExchangeTransport Queues(_total)\Retry Mailbox Delivery Queue Length	<100	<1	<1

Validation of Root Server Health

The following tables show the validation of root server health.

Processor

As expected, the processor utilization is very low and well under target thresholds.

Counter	Target	Tested result (150 messages per day)	Tested result (100 messages per day)
Hyper-V Hypervisor Logical Processor(_total)\% Guest Run Time	<75%	10.7	8.8
Hyper-V Hypervisor Logical Processor(_total)\% Hypervisor Run Time	<5%	0.9	0.9
Hyper-V Hypervisor Logical Processor(_total)\% Total Run Time	<80%	11.6	9.7

Counter	Target	Tested result (150 messages per day)	Tested result (100 messages per day)
Hyper-V Hypervisor Root Virtual Processor(_total)\% Guest Run Time	<5%	1.7	1.7

Application Health

The Virtual Machine Health Summary counter indicates that all VMs are in a healthy state.

Counter	Target	Tested Result (150 messages per day)	Tested Result (100 messages per day)
Hyper-V Virtual Machine Health Summary\Health Critical	0	0	0

Test Case: Single Server Failure or Single Server Maintenance (In Site)

The second test case represents peak workload during a failure or maintenance event, where one of the two servers is no longer operational, and all four databases are active on the surviving server. This test case represents the worst case workload, and therefore is considered the key performance validation test. Test results for both the planned message profile of 150 as well as a lower message profile of 100 are provided.

Validation of Expected Load

Message delivery rate verifies that tested workload matched the target workload. For both the 150 and 100 message profiles, the actual message delivery rate is on target.

Counter	Target	Tested result (150 messages per day)	Target	Tested result (100 messages per day)
Message Delivery Rate / Mailbox Server	4.2	4.2	2.8	2.8

Validation of Mailbox Servers

The following tables show the validation of Mailbox servers.

Processor

Processor utilization is low, as expected.

Counter	Target	Tested result (150 messages per day)	Tested result (100 messages per day)
Hyper-V Hypervisor Virtual Processor% Guest Run Time	<70%	18.5	13.6

Storage

The storage results are good. The I/O database read and write average latencies are well within the 20 msec target, and database page fault stalls and log record stalls are both 0, as expected.

Counter	Target	Tested result (150 messages per day)	Tested result (100 messages per day)
MSExchange Database\I/O Database Reads (Attached) Average Latency	<20 msec	17	16
MSExchange Database\I/O Database Writes (Attached) Average Latency	<20 msec <Reads average	0.7	1
Database\Database Page Fault Stalls/sec	0	0	0
MSExchange Database\I/O Log Writes Average Latency	<20 msec	0.3	<1
Database\Log Record Stalls/sec	0	0	0

Application Health

Exchange is healthy, and all the counters used to determine application health are well under target values.

Counter	Target	Tested result (150 messages per day)	Tested result (100 messages per day)
MSExchangeIS\RPC Requests	<70	0.52	0.7
MSExchangeIS\RPC Averaged Latency	<10 msec	1.4	1.5
MSExchangeIS Mailbox(_Total)\Messages Queued for Submission	0	17	14.8

Counter	Target	Tested result (150 messages per day)	Tested result (100 messages per day)
MSExchange Replication(*)\CopyQueueLength	<1	<1	<1

Validation of Client Access and Hub Transport Servers

The following tables show the validation of Client Access and Hub Transport servers.

Processor

Processor utilization is low, as expected.

Counter	Target	Tested result (150 messages per day)	Tested result (100 messages per day)
Hyper-V Hypervisor Virtual Processor\ Guest Run Time	<70%	12.3	8.5

Storage

The storage results are good. The very low latencies should have no impact on message transport.

Counter	Target	Tested result (150 messages per day)	Tested result (100 messages per day)
Logical/Physical Disk(*)\Avg. Disk sec/Read	<20 msec	0.000	0.004
Logical/Physical Disk(*)\Avg. Disk sec/Write	<20 msec	0.001	0.001

Application Health

The low RPC Averaged Latency values confirm a healthy Client Access server with no impact on client experience.

Counter	Target	Tested result (150 messages per day)	Tested result (100 messages per day)
MSExchange RpcClientAccess\ RPC Averaged Latency	<250 msec	5.2	5.1

Counter	Target	Tested result (150 messages per day)	Tested result (100 messages per day)
MSExchange RpcClientAccess\RPC Requests	<40	0.6	0.4

Hub Transport Server Health

The Transport Queue counters are all well under target, confirming that the Hub Transport server is healthy, and able to process and deliver the required messages.

Counter	Target	Tested result (150 messages per day)	Tested result (100 messages per day)
\MSExchangeTransport Queues(_total)\Aggregate Delivery Queue Length (All Queues)	<3000	<1	<1
\MSExchangeTransport Queues(_total)\Active Remote Delivery Queue Length	<250	0	0
\MSExchangeTransport Queues(_total)\Active Mailbox Delivery Queue Length	<250	<1	<1
\MSExchangeTransport Queues(_total)\Submission Queue Length	<100	<1	<1
\MSExchangeTransport Queues(_total)\Retry Mailbox Delivery Queue Length	<100	<1	<1

Validation of Root Server Health

The following tables show the validation of root server health.

Processor

As expected, the processor utilization is quite low and well under target thresholds.

Counter	Target	Tested result (150 messages per day)	Tested result (100 messages per day)
Hyper-V Hypervisor Logical Processor(_total)\% Guest Run	<75%	16.6	12.1

Counter	Target	Tested result (150 messages per day)	Tested result (100 messages per day)
Time			
Hyper-V Hypervisor Logical Processor(_total)\% Hypervisor Run Time	<5%	1	0.9
Hyper-V Hypervisor LogicalProcessor(_total)\% Total Run Time	<80%	17.6	13
Hyper-V Hypervisor Root Virtual Processor(_total)\% Guest Run Time	<5%	1.8	1.3

Application Health

The Virtual Machine Health Summary counter indicates that all VMs are in a healthy state.

Counter	Target	Tested result (150 messages per day)	Tested result (100 messages per day)
Hyper-V Virtual Machine Health Summary\Health Critical	0	0	0

[Return to top](#)

#Conclusion

This white paper provides an example of how to design, test, and validate an Exchange 2010 solution for customer environments with 500 mailboxes or less deployed on Dell server and storage solutions. The step-by-step methodology in this document walks through the important design decision points that help address key challenges while ensuring that core business requirements are met.

[Return to top](#)

For the complete Exchange 2010 documentation, see [Exchange Server 2010](#). For an in-depth discussion of the key concepts used in design decision points throughout this document, see the following topics in the Exchange Library:

- Understanding High Availability and Site Resilience

#Con

- Planning for High Availability and Site Resilience
- Understanding Database Availability Groups
- Understanding Backup, Restore and Disaster Recovery

For more information, see:

- [Dell and Exchange Server 2010](#)
- [PowerEdge Servers](#)
- [Dell PowerEdge T610](#)
- [Exchange 2010 Advisor](#)
- [Optimizing Microsoft Exchange Server 2010 Deployments on Dell Servers and Storage](#)
- (Download) [Exchange 2010 migration: Legacy HP ProLiant DL385 to Dell PowerEdge R510](#)
- [Dell Enterprise Technology Center](#)

This document is provided "as-is." Information and views expressed in this document, including URL and other Internet Web site references, may change without notice. You bear the risk of using it.

This document does not provide you with any legal rights to any intellectual property in any Microsoft product. You may copy and use this document for your internal, reference purposes.

[Return to top](#)