Design Recommendations
For a Highly Available Exchange 2007 Environment

EMC Proven™ Professional Knowledge Sharing

Douglas Collimore
EMC Corporation
Collimore_douglas@emc.com

Table of Contents

Table of Figures

Disclaimer: The views, processes or methodologies published in this article are those of the authors. They do not necessarily reflect EMC Corporation's views, processes or methodologies.

Introduction

Today, the term SOLUTION implies an application that performs a specific process. For example, Microsoft Exchange is an application that empowers collaboration among users through email, unified messaging, file sharing, etc. This is important to individuals and corporations throughout the world but is simply a point solution, a problem solved through the use of a specific software application and/or hardware.

The problem with point solutions is that they solve a specific customer requirement but the solution itself does not necessarily take into account enterprise needs and requirements. A specific example of this directly impacts Exchange 2007. A point solution might utilize Continuous Cluster Replication (CCR) for local high availability of the Exchange database. However, additional requirements might include an online, zero-impact backup of the same database(s), a remote replication technology to enable disaster recovery of the entire production Exchange environment, and/or an archive and compliance solution that must be available regardless of where the Exchange servers are operating (local or remote).

A complete enterprise solution supporting Exchange 2007 is critical to provide the customer with the ability to address the following real-world requirements:

- 5 nines local availability
- Site failover during a disaster (flood, fire, etc.) and eventual failback
- No data loss
- Quick recovery upon failure
- Archiving for extended periods (7 years or more)
- Compliance (meeting Sarbanes-Oxley, HIPAA, and other requirements)
- Integration of multiple vendor technologies

Obviously there are many more, but the idea is that there are many different requirements to meet many different business needs. No application or point solution can address each and every one. This requires the creation of an enterprise solution, one that takes into account ALL the requirements and builds a solution that will meet each one utilizing one or many hardware and software products.

This article addresses one solution in a high availability and disaster recovery Exchange 2007 infrastructure within a virtualized environment. The design criteria used to establish requirements for this solution included:

- Recovery Point Objective (RPO), both local and remote
- Recovery Time Objective (RTO), both local and remote
- Daily backup and maintenance windows
- Archiving
- Maximum utilization of server and network hardware
- Support for 10K users or more

This article describes each requirement, how it impacts the solution (performance and administrative overhead for example), and its relevance to the overall solution.

1.0 Requirements

We must collaboratively define customer requirements before we even consider an integrated solution. A design is useless without a written record of what the business customer expects for a service level. Without it, you will never know if you are meeting your customers' expectations.

Today's business clients demand more and more from their IT infrastructure. We used typical customer requirements to establish the design criteria for a complete solution. By using requirements that are asked of us every day, we can have a valuable discussion about design criteria.

We established design criteria based on the enterprise level requirements that follow:

- Support for 10K users or more
- Recovery Point Objective-Local RPO of zero implying a synchronous solution that does NOT take consider data corruption of the Exchange databases. This implies that the data on the production storage groups and logs is replicated in its current state to a set of remote disks synchronously.
- Remote RPO of less than 5 minutes for the Exchange databases implying an Asynchronous solution at a distance greater than 100 km
- Local and Remote RTO of under 4 hours
- Daily backups (RTO of less than 4 hours impact to the production Exchange server though the backups can take longer) while also completing online maintenance during off hour windows
- Archiving of all mail over 90 days old with a retention policy of 7 years or more
- Maximum utilization of server and network hardware

We assume that you have a basic understanding of collecting customer requirements, Exchange storage and performance design, and virtualization. For the sake of brevity, we will not discuss these areas in detail, but will reference articles and resources that you can use to gain additional knowledge in these areas.

2.0 Designing the Solution

A solution has many components. Business Continuance (BC) and Disaster Recovery (DR) have their own hardware and software requirements. Archiving uses an entirely different set of hardware and software applications, as do compliance and virtualization. The foundation can include the virtualization layer and the production Exchange environment.

A true solution does not necessarily focus on particular hardware or software. Collecting business requirements to meet the customers' needs is the critical step in designing a solution. Once you have requirements, you can begin adding hardware and software to meet them. Please keep in mind that the products you select to build your solution will affect performance. Not all hardware/software is created equally so you may have elongated backups, fewer supported users, or you may not meet your RPO and/or RTO objectives. Figure 1 depicts the general architecture for a design that could support the requirements previously discussed.

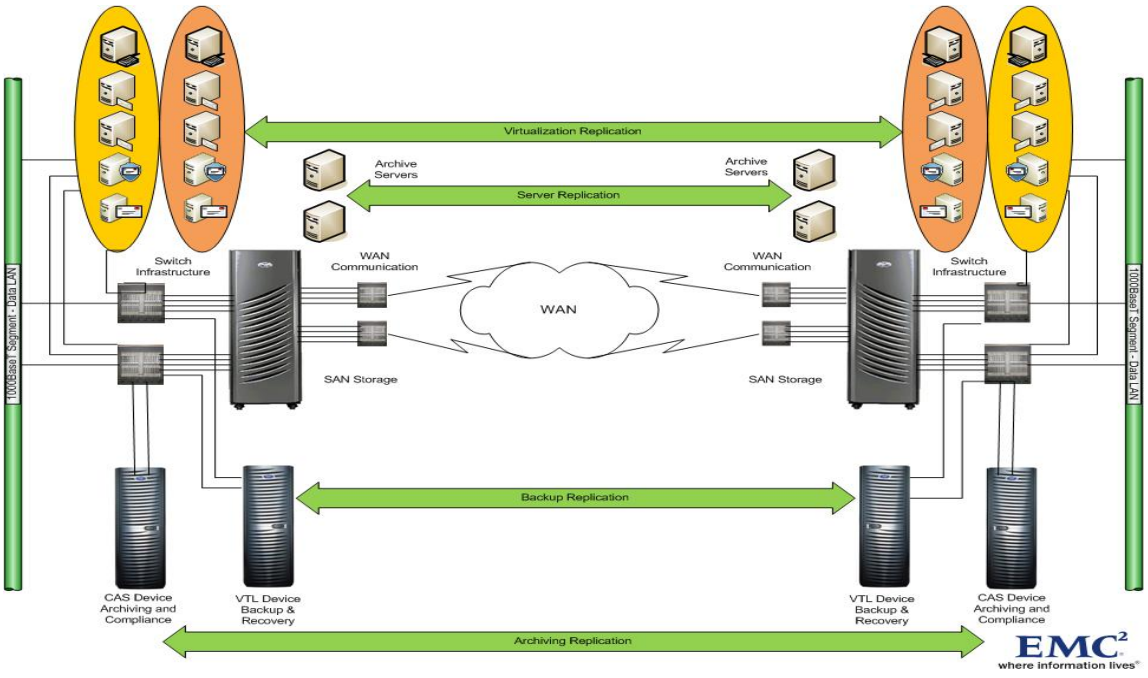

Figure 1. High Level Infrastructure Diagram

Now, let's review each requirement and discuss what it means, how it impacts the solution, and why it is important.

Support for 10K Users or More

User support size is a critical requirement when designing an enterprise solution. Performance issues in an application or infrastructure may not be recognized in a low impact environment. Flaws generally become apparent only when systems are stressed. The easiest way to stress an environment is to increase the active user count thereby increasing the IO, memory utilization, and associated network traffic.

With Exchange 2007, 10,000 users should require between 2 and 3 mail servers and associated support servers (AD, Hub Transport, etc.). This infrastructure needs servers with sufficient memory, a properly sized disk subsystem and network bandwidth to maintain low latency thus enhancing the client experience.

A limited number of tools were available to size Exchange 2003 when it was first available. As an Exchange administrator, you were often at the mercy of your server or storage vendor's opinion on exactly what was required to support your performance requirements. Sizing improved as we gained knowledge and performed testing to validate the application's scalability. Microsoft continued to build sizing tools and calculators making Exchange design more data driven.

Microsoft developed new sizing calculators for Exchange 2007. Whereas the sizing calculator for Exchange 2003 was a single Excel worksheet with limited variable input, the new calculator accepts a host of criteria ranging from user IO requirements, to message transfers and client mode type. The result of this extensive tool is more accurate information for sizing not only the storage infrastructure, but also the servers. Figure 2 depicts the sizing calculator input page.

Microsoft developed the Exchange Solution Reviewed Program (ESRP) (http://technet.microsoft.com/en-us/exchange/bb412164.aspx) to size storage subsystems. Once you have used the sizing calculator and obtained a user IO requirement, access the link to conduct research to design your Exchange 2007 environment. Every major storage vendor has tested different configurations to support a multitude of user scenarios. Your user and IO requirements may not exactly meet the listed test results from your vendor of choice. Select a configuration that has successfully passed the required JetStress testing that most closely resembles your needs. If will offer a baseline to begin sizing the remainder of your environment.

Do NOT just accept these test results. Make sure that you test your servers in your environment to obtain a true baseline measurement before releasing your mail servers to production. To ensure your servers and storage will perform as required, you should ALWAYS run a series of JetStress and LoadGen tests before deploying to production.

Jetstress is a tool created for Exchange 2003 (and has since been updated for E2K7) that tests your storage subsystem. It offers the ability to run stress and IO tests on your Exchange servers against the storage LUNs you have created to support the storage group and logs. Jetstress results give you a glimpse into how your disks will perform in production. Use the results as a baseline after implementing your servers to see if you have a performance issue.

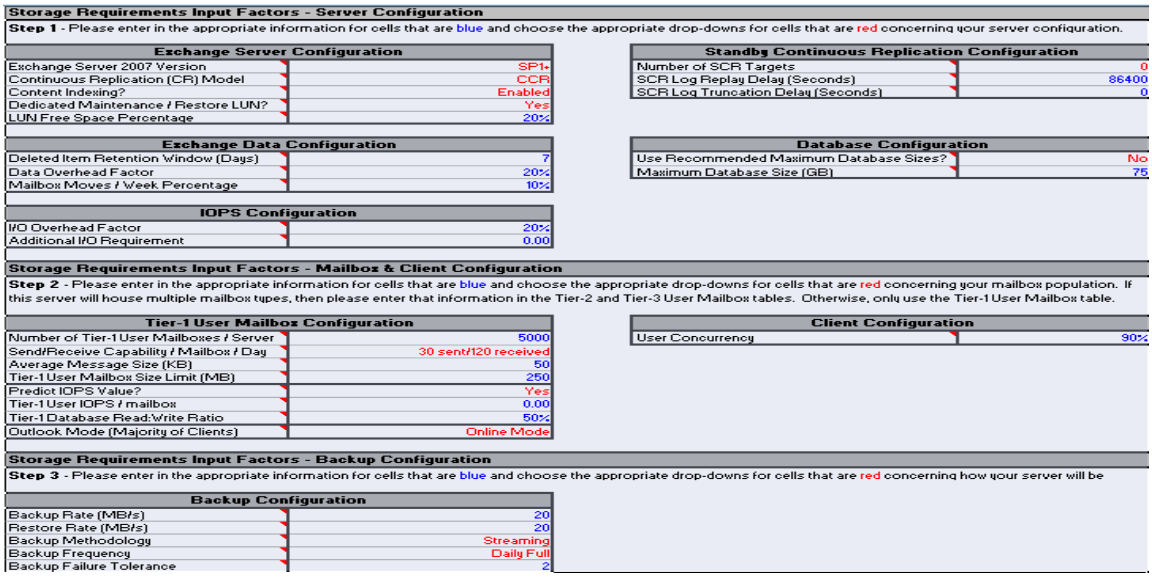

Storage Requirements Input Factors - Server Configuration

Step 1 - Please enter in the appropriate information for cells that are blue and choose the appropriate drop-downs for cells that are red concerning your server configuration.

Exchange Server Configuration	
Exchange Server 2007 Version	SP1
Continuous Replication (CR) Model	CCR
Content Indexing?	Enabled
Dedicated Maintenance / Restore LUN?	Yes
LUN Free Space Percentage	20%

Standby Continuous Replication Configuration	
Number of SCR Targets	0
SCR Log Replay Delay (Seconds)	86400
SCR Log Truncation Delay (Seconds)	0

Exchange Data Configuration	
Deleted Item Retention Window (Days)	7
Data Overhead Factor	20%
Mailbox Moves / Week Percentage	10%

Database Configuration	
Use Recommended Maximum Database Sizes?	No
Maximum Database Size (GB)	75

IOPS Configuration	
I/O Overhead Factor	20%
Additional I/O Requirement	0.00

Storage Requirements Input Factors - Mailbox & Client Configuration

Step 2 - Please enter in the appropriate information for cells that are blue and choose the appropriate drop-downs for cells that are red concerning your mailbox population. If this server will house multiple mailbox types, then please enter that information in the Tier-2 and Tier-3 User Mailbox tables. Otherwise, only use the Tier-1 User Mailbox table.

Tier-1 User Mailbox Configuration	
Number of Tier-1 User Mailboxes / Server	5000
Send/Receive Capability / Mailbox / Day	30 sent/120 received
Average Message Size (KB)	50
Tier-1 User Mailbox Size Limit (MB)	250
Predict IOPS Value?	Yes
Tier-1 User IOPS / mailbox	0.00
Tier-1 Database Read:Write Ratio	50%
Outlook Mode (Majority of Clients)	Online Mode

Client Configuration	
User Concurrency	90%

Storage Requirements Input Factors - Backup Configuration

Step 3 - Please enter in the appropriate information for cells that are blue and choose the appropriate drop-downs for cells that are red concerning how your server will be

Backup Configuration	
Backup Rate (MB/s)	20
Restore Rate (MB/s)	20
Backup Methodology	Streaming
Backup Frequency	Daily Full
Backup Failure Tolerance	2

Figure 2. E2K7 Sizing Calculator Input Screen

The LoadGen (http://www.microsoft.com/downloads/details.aspx?FamilyID=C7534CB2-4234-41F1-B0B2-ED9F376B4CB2&displaylang=en) utility is used to stress test the Exchange design. Whereas JetStress tests your storage, LoadGen tests the impact of MAPI clients against your Exchange servers. Jetstress results offer an idea as to how your storage will react with a user load, LoadGen testing shows you how your Exchange servers will react to Exchange email loads.

As you can see, both tests are required to validate your E2K7 design. Having real data that simulates your design will help you to analyze your environment once in production. Without a baseline, you have nothing more than an educated guess as to how your environment is performing and scaling.

Recovery Point Objectives

Today's enterprise solutions require a level of resiliency not seen previously in IT infrastructures. The Recovery Point Objective (RPO) is a key requirement to build the proper design. It reflects the service level agreement between the architects and the lines of business that they are supporting. If you experience an issue and need to recover, an Exchange design isn't worth much if you can't get your data back in the agreed timeframe.

The key to selecting an appropriate Recovery Point Objective is to discuss these key points with business stakeholders:

- How much data loss can the business accept?
- Must data availability remain the same both locally and remotely?
- How much are you willing to spend to build an infrastructure to meet the first two requirements?
- How difficult is it to support the solution(s)?

Today, a typical RPO for local recovery (high availability) in an Exchange enterprise is between 0 (zero data loss) and 1 minute. An agreed upon RPO for remote recovery

(disaster recovery) is commonly between 0 (zero data loss) and 15 minutes. Zero data loss means different things to different people. If you were to ask a Microsoft person what zero data loss means in an Exchange environment, they would tell you all data would be intact AND consistent on the secondary/remote storage. A storage vendor implementing a synchronous replication solution would say zero data loss would guarantee that the data on the secondary/remote storage would be identical with the data on the primary production volumes even if the Exchange data was not consistent.

These different viewpoints result from who has control of the environment and who has responsibility for each part of the solution. While a storage vendor can efficiently replicate all the data residing on disk, it really has no way of manipulating the application and verifying the actual data as it relates to Exchange. If Exchange is corrupt on the production database, any synchronous (and depending on the RPO of an asynchronous solution) solution is going to replicate that same corrupt data to the secondary storage (local or remote). The replication process could use a Microsoft tool to verify the data before it is sent across the wire, but the latency invoked on a synchronous solution by such a process would be too great to meet Microsoft's stringent requirements.

Microsoft developed clustering solutions to alleviate the chance that corrupt data is replicated. They guarantee that data is consistent by performing a consistency check on the data/logs before applying it to a backup database volume. This is NOT a synchronous solution, however. There will be some data lag using a log shipping type of methodology, even if the link between the primary and secondary storage is fast.

As an architect, you are responsible to educate your customer about the different technologies available to them, the cost to implement AND support the solution(s), and the value that each brings to the client's recovery objectives. This means that you must be able to use layman's terms to describe the replication methodology chosen to support the requirements, the effect of the replication technology (local and/or remote) on the Exchange data, the recovery point of the Exchange application, the administrative overhead to support the solution(s), how the solution meets the business objectives, and the monetary cost to test, install, and support the solution.

You have not completed your job if you are unable to reach an agreement with business stakeholders about requirements and business objectives. The first time there is an outage and data is lost, you will be held accountable because you did not meet the business' RPO. Conversely, helping the business to understand the solution and recoverability of the Exchange data ensures an agreement that the RPO to support the business will be met at a specific cost level in an agreed timeframe.

Recovery Time Objective

A discussion of the Exchange RPO is incomplete without also discussing how much time each solution will take to recover (Recovery Time Objective or RTO). Suppose you developed a solution to support near zero data loss but it takes 8 hours to actually recover and have Exchange back online. The business may not be able to tolerate the outage. The solution's RTO is the time it takes to recover both the Exchange data and the Exchange environment to an operational state. In some instances, dial-tone only would suffice to meet the RTO requirement. In other cases, full Exchange functionality might be required to meet the RTO of the business.

As with the RPO discussed in the previous section, there are a few key points that need to be discussed with the stakeholders of the business with regard to the Recovery Time Objective:

- How much of an outage can the business endure in the event of a local failure or site outage (local RTO may be different from a full site failure and associated remote RTO)?
- How much will the recovery cost both monetarily and in man-hours?
- How simple or difficult is it to implement the recovery?
- More importantly, how difficult is it to fail back to the primary server or site once the failure has been rectified?

Typically, there are two cost categories. Hard costs reflect the amount of money in hardware, man-hours, etc., that are incurred during the recovery period; soft costs reflect

the real or imagined monetary business loss incurred during the outage. Neither can be omitted from the RTO discussion.

For example, a financial business may use Blackberry devices to communicate trades from a stock floor. An Exchange outage would prevent this process from occurring and result in traders having to use another method that is probably slower. This can detrimentally impact the business.

A soft cost might be attached to this type of outage. Let's assume a cost of $1 million dollars an hour in reduced transaction completion time. In addition, the hard costs might require 8 men, 4 hours to recover the Exchange servers and a weekend outage with an additional 12 hours to failback to the primary servers/location once the issue is resolved. This single outage would have a real and imagined cost of over $16 million dollars! Suffice to say, a solution that cost $4 million to implement but saved 12 hours of work would pay for itself with just one outage.

As an architect, you must clearly understand the functionality of Exchange 2007 and its newly built-in technologies to support high availability and disaster recovery. Other vendors technologies must also be familiar as they may offer similar or better protection, but for a reduced cost or administrative overhead. The most difficult responsibility is explaining how each solution or technology impacts not only the Exchange environment but the entire enterprise infrastructure.

It would be great if we, as architects, had unlimited budgets and time to build a robust solution that could account for every type of anomaly. However, time, cost, complexity, performance impact, and most importantly, the business impact of an outage are the key variables driving us to our end state solution.

Backups and Online Maintenance

When sizing servers and storage for Exchange, Microsoft recommends collecting Perfmon statistics from your current mail system. Using your busiest 2 hour window, calculate what is required to support Exchange 2007 from a performance perspective.

This methodology usually works quite well. Most enterprise environments claim to be 24x7x365 shops but typically, the majority of email traffic can be seen during a standard business day taking into account extra time between the East and West Coast of the United States. This leaves a quiet time at night to perform backups and the required Exchange online maintenance.

Unlike IO patterns viewed during standard email operations, these two maintenance functions can bring an Exchange server and its storage to a crawl as both like to use as much memory, CPU, network bandwidth, and storage IO as they can. The problem is exacerbated when you add multiple servers performing the same maintenance functions at the same time against the same storage system and network. If you happen to be a real 24x7 environment, this can be detrimental to the user experience AND more importantly, your service level agreement to the business. The decision, in this case, is whether or not to increase server memory, storage spindle count, CPU size and quantity, or a combination to generate enough performance headroom to help eradicate the sluggishness of the environment during these maintenance periods.

Let's discuss backups first. Unlike Exchange 5.5 and 2000 days where backup to tape and sometimes expensive disk was the only supported backup solution, today's technologies offer many different solutions to the backup conundrum. Backups are a primary piece of the high availability equation. Server performance is affected during backups, so any design decision about how we will backup Exchange and, more importantly, how we will recover it needs to be made based upon the previously discussed RPO and RTO requirements.

Many people speak of backup and their backup window in the truest sense; how much data are we backing up and how long will it take. Unfortunately, recovering the backed-up data and the time it takes to complete the recovery is more important than the actual backup. When discussing backups, the following discussion points are critical to the conversation. None can be left out:

- RPO-how much data loss can we expect if we have to recover from a backup?
- RTO-how long will it take to backup our data and how long will it take to restore the data during an outage?

- How can we offload the backup and restore processes from our production servers?
- How often do I need to backup? How often CAN I backup?
- Do we have Microsoft support for the backup/restore solution?

We can derive many different solutions depending on the answers to these questions. Let's assume that an enterprise Exchange environment with minimal RPO and RTO (<30 minutes, < 4 hours for backup or recovery) are required for local recovery. This is a typical answer when having this type of discussion with business stakeholders. So, how are we able to take backups of very large Exchange databases quickly, recover these backups just as quickly, without severely impacting the Exchange server's performance?

Offloading backups to another server alleviates the performance impact of a standard streaming backup. Software VSS is one method of offloading the backup, though the software vendors are still struggling with this for Exchange 2007. Hardware VSS is another method of creating copies or clones of the Exchange databases. These backup applications create point-in-time images that are checked for consistency and then, upon successful completion of the consistency check, truncate the Exchange logs much like the normal tape backup process would.

There are two differences between a hardware VSS backup solution and the usual streaming tape solution. The hardware VSS solution only impacts the Exchange server when the clone has been completed and the replication between the production LUNs and the clone relationship needs to be halted or split. The synchronization of the production LUNs and the clones is handled by the storage system, alleviating the performance impact to the Exchange mail server. Breaking the synchronization relationship requires a VSS call to queue Exchange IO so we have a valid and consistent clone image. The maximum time frame for this process is 10 seconds though it typically takes less time, much less than the 4 to 8 hour impact of a standard tape backup. Second, the recovery from disk using current technology offers almost instantaneous restores and high speed log replays, further shrinking the recovery timeframe.

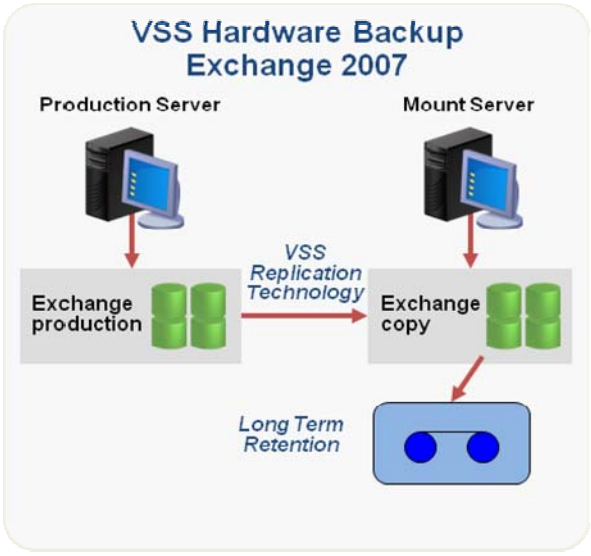

Figure 3. High Level Hardware VSS Backup Diagram

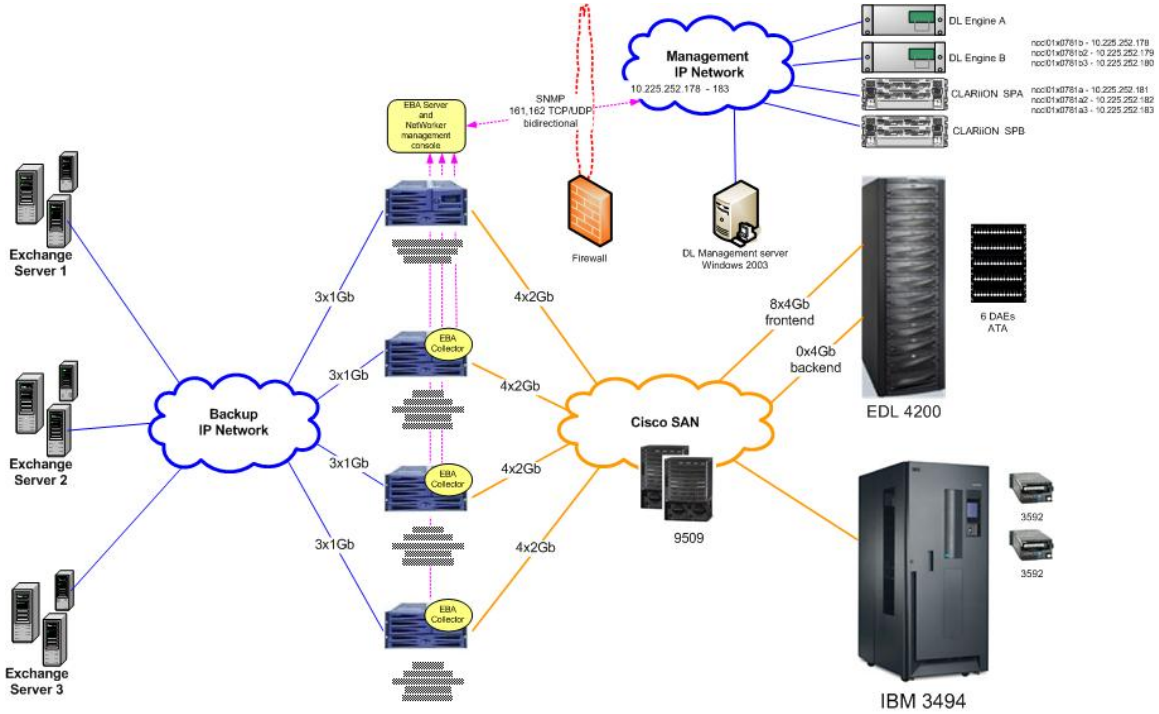

Figure 4. Typical Virtual Tape Backup Infrastructure

By completing the backup process to disk in a non-impacting method, you shorten your backup window, allowing other maintenance functions to occur. Having made a VSS clone of the production database and logs meets the requirement for a supported Microsoft backup. You can also mount the clone on a non-production server and then copy it to some form of inexpensive media designed for long term retention (tape or ATA disk). Figure 3 depicts a typical hardware VSS backup flow diagram.

The ability to complete online maintenance in the time allotted is enhanced by extending the backup window on a non-production server and removing a key performance impacting requirement from the Exchange servers. Remember, our goal is to reduce CPU and production storage utilization to enhance the user experience in a true 24 hour mail environment.

Microsoft recommends completing online maintenance (OLM) tasks, on each storage group, every two days and a minimum of at least once per week to guarantee a well-performing Exchange server. In an enterprise consisting of multiple Exchange mail servers with many storage groups (up to 50), this can be a daunting, if not impossible, task. This function can be scheduled and not all servers have to run online maintenance every day. In fact, alternating OLM on your servers so half run during one day's maintenance window while the other half runs on the following day can ensure that the process completes on a regular basis.

Storage is another avenue that can offer increased performance for online maintenance. OLM is a highly disk intensive process. While we usually size storage to meet a user IO requirement, a review of performance monitor statistics-specifically Disks Transfers/sec-will show that spindle utilization is unusually high, higher than your busiest user window, during the online maintenance and backup periods. We have already reduced the backup performance overhead by shifting the process to a mount host. Increasing hardware performance by adding memory, CPUs (both quantity and speed), and more importantly, physical disk spindles to the storage group and log disks will decrease the time needed to complete online maintenance. Your environment will vary.

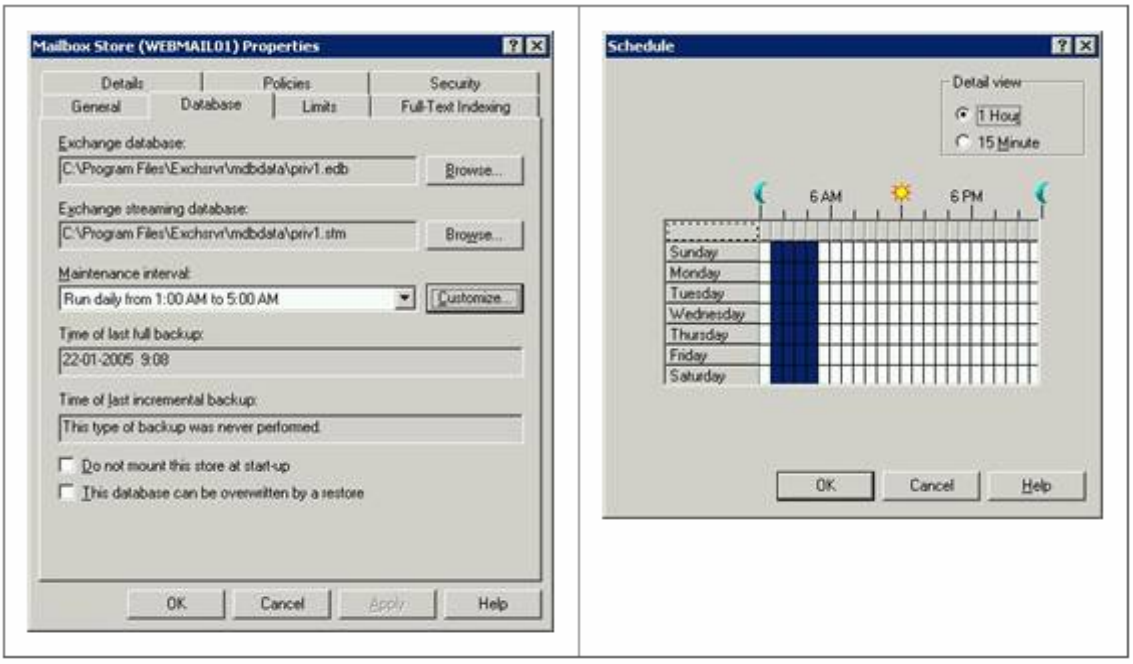

Figure 5. Online Maintenance Scheduling Window

There are also many registry settings and processes that can be tweaked to increase the backup performance and online maintenance. As always, check Microsoft best practices before manipulating any registry or application settings. Use the various analysis tools available on the Microsoft Exchange website to check your servers and make sure they have the proper settings (boot, registry, maintenance) and meet the minimum performance and sizing specifications.

Archiving

The advent of email archiving serves a multitude of purposes. It assists in meeting compliance requirements, reduces the cost and risk of legal discovery, and enables corporate governance. Records management further decreases management costs of data against the increase in mailbox quotas.

From a performance and email management perspective, archiving of email decreases the size of the Tier 1 Exchange databases as older data is shifted to less expensive storage with a short-cut pointer back to the user's mailbox. Left undeleted from the archive, this solution creates a user perception of an unlimited mailbox quota. This helps reduce the backup window of the production Exchange data as less data needs to be backed up on a daily basis off the current production volumes. We can decrease the cost of ownership by backing up archived data less frequently and on less expensive media. Finally, it improves storage management and operational efficiency, while allowing for predictable mail storage growth.

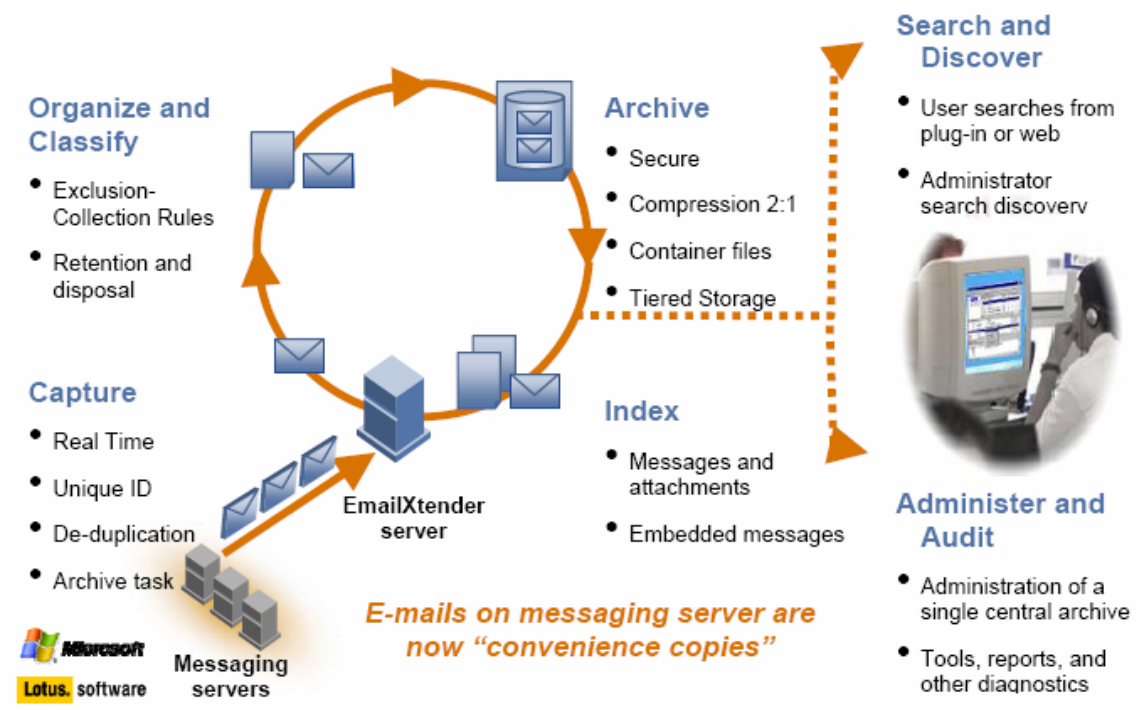

Figure 6. Archiving Process Schematic [1]

Reducing the size of the Exchange databases can significantly improve your backup timeframe allowing for more frequent backups and/or more time to perform other maintenance tasks. For example, let's say you have 10 storage groups with a single database per storage group on each of 5 Exchange Servers. Each storage group LUN and associated log represents approximately 100 GB of data for back up. That is a Terabyte of data per server times 5 servers requiring back up and storage every day. Depending on your backup methodology, streaming this to tape could take hours and impact your online maintenance window. Using VSS backup, you can reduce the time required to backup, but you are still backing up the same amount of data. It still requires vast amounts of costly Tier 1 storage even if it's not impacting the production server.

Now, let's assume that we implement an aggressive archive solution (all mail older than 90 days gets archived). We could reduce the size of the Exchange databases up to 70%! Even if we conservatively estimate the database size reduction at 30%, that is over 1.5 Terabytes that no longer has to get backed up on tier 1 storage, though models and experience show the reductions will be far greater.

However, archiving is accompanied by a performance cost. The archive engine that crawls through the databases and checks for mail that meets the archive criteria is expensive from an IO perspective. Just as Blackberry usage or online maintenance functionality adds to performance degradation of the disk subsystem, archiving your mail will take up precious IO cycles on your disk and CPUs. When designing the storage layout, expect a 30% hit on database read IO for archive functionality. You need to check with the archive manufacturer on related performance impacts to server CPU and memory.

Along with performance degradation, you also have to consider that much of your mail data is now stored on secondary storage once it is archived. While it does not need to be backed up as frequently, you do need to back it up. This amount of data will be a much larger amount of data than the original Exchange stores as retention policies can span 7 years or more. When backing the data to tape for long term retention, you do not impact performance on the production Exchange servers. This is a significant advantage. You can also isolate this Tier's backup process so the production environment is not detrimentally affected.

You must carefully consider the backup infrastructure for the archive portion of the design. It is preferable to replicate the data to an offsite location and back it up once it has been transferred (if offsite retention is a requirement). This helps to meet two standard requirements: offsite data storage and isolated backup of Tier 2 data. In the event an offsite location is not available, try to isolate the backup stream by using separate storage and backup hardware from the Exchange environment. This includes using different backup servers, isolated network bandwidth (if backup will be over an IP link) from the primary Exchange servers, and separate physical media to store the archived data. Figure 7 depicts an archive environment with a backup infrastructure.

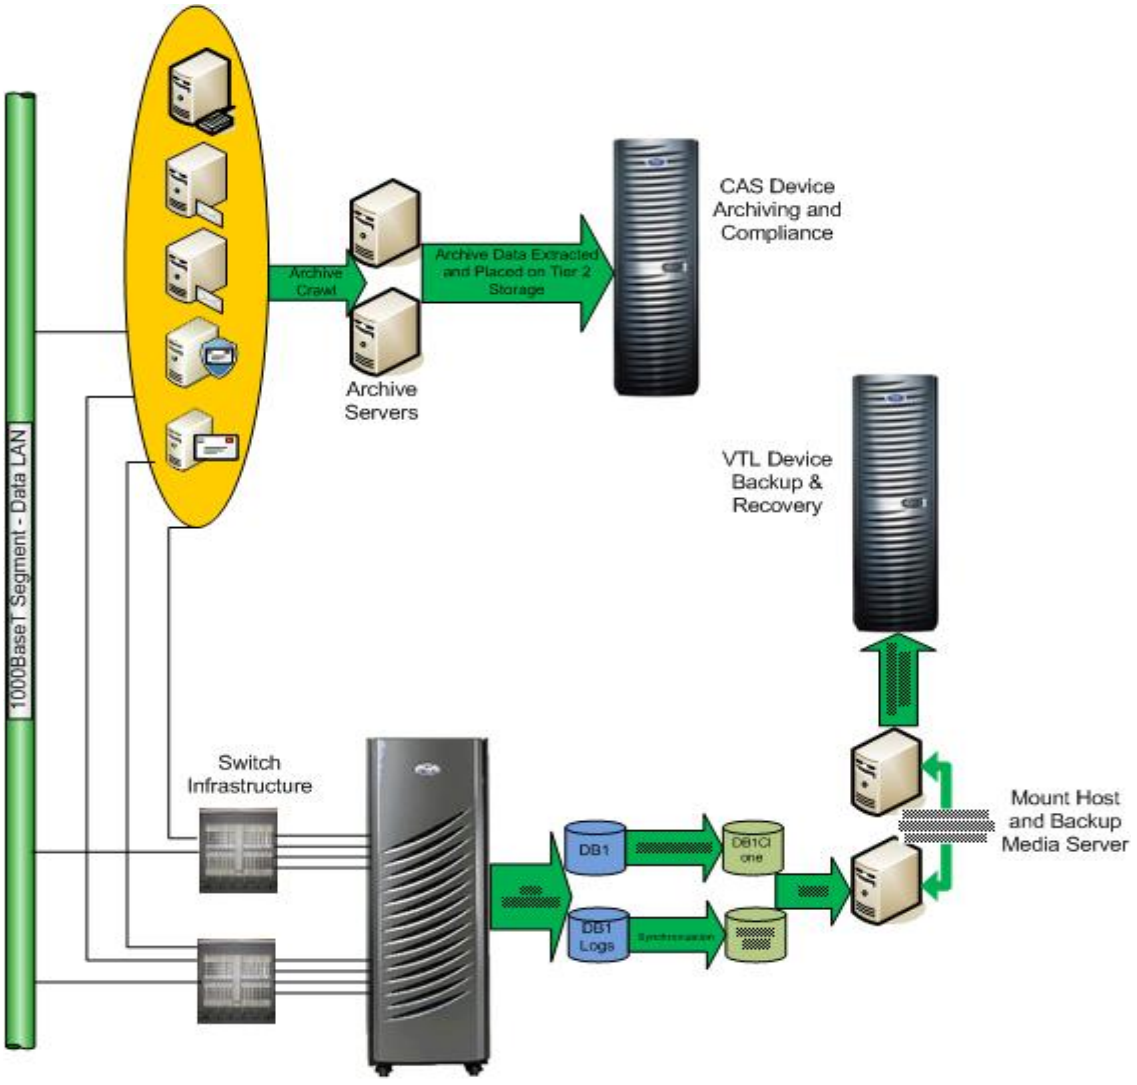

Figure 7. Archive Infrastructure

Virtualization of the Design

The foundation of the enterprise solution incorporates the physical servers, the operating system (OS), the email application, and the SAN storage. This is a stable basis upon which to build the infrastructure. Clustering the servers adds high availability but also a level of complexity. Cloning LUNs within the storage array adds business continuance and also enhances the high availability model. Archiving eases administration overhead and reduces Tier 1 storage capacity needs. Backup adds resilience and recoverability. There are certain features that help build a robust, HA/BC/DR model even within the OS and email application.

Many companies have an issue with the cost of this base foundation and its maintenance, particularly as cost assumes many different forms. There are:

- administrative costs: manpower, training, administration, software, and support;
- datacenter costs: power, cooling, real-estate; application costs-licensing fees;
- infrastructure costs: servers, network, high availability, disaster recovery, and backup.

We can overcome underutilized hardware and reduce the number of physical objects that require maintenance. Server consolidation reduces tile space, power, cooling, and network connectivity costs. The tiered storage model has already shown significant savings in the enterprise space. Storage Virtualization eases administration costs and increases storage utilization.

Some advantages of virtualizing an Exchange 2007 environment include, but are not limited to:

- Improved Administration model and tools
- Easier deployments and cost reduction associated with operational efficiencies
- Increased High-Availability through virtualized server transport
- Simplified Disaster Recovery model
- Improved Exchange performance through the use of a 64 bit platform and increase RAM availability
- Increased server utilization including maximized CUP use and memory resources. This includes leveraging hardware to incorporate server roles such as Active Directory, domain controllers, hub transport servers, DNS, DHCP, and more

Cost reduction is not the only reason to virtualize an Exchange environment. Server virtualization can offer increased high availability without having to add the complexity of clustered Windows servers. VMWare's VMotion technology allows for the uninterrupted transfer of an application from one physical server running a virtual server instance to another while maintaining a fully operational application. The client never knows the application has moved between physical servers. This allows for significant changes in server maintenance and total uptime.

You have the ability to take a virtual image and move it between physical servers, much like Microsoft Cluster Service. However, you are not limited to the same data center. For instance, build an Exchange server within a virtual server environment and make a copy of the image. Keep a local copy and replicate one to a distant datacenter. In the event the local server goes down, you only need to restart the copied image on another physical server and Exchange returns to operational status.

Lose an entire Exchange environment and you can fail over rather seamlessly to a remote site, again by simply bringing up the virtual images and pointing the application to the replicated databases and logs. DR recovery is a little more complicated than flipping a switch, but is still rather simplistic compared to some methods of DR for physical servers.

Typically, when building an Exchange server, we size the hardware to meet performance requirements. This means we purchase enough CPU, memory, and network cards to provide optimum performance while staying under the performance and latency watermarks set forth by Microsoft. We must take care to size virtual servers correctly to maintain the same performance achieved when running in a physical only environment. With the advent of Exchange 2007 and a 64 bit operating environment, duo-core CPUs, and more memory than was ever available to the Windows OS, we can now size servers to handle more than one Exchange instance.

Now, one might ask, instead of putting 2 virtual servers on one physical server and then running Exchange, why not just use 2 smaller servers? It is a frequently debated question. Smaller blade servers might offer decreased data center space, power, and cooling usage. However, this does not take advantage of other virtualization benefits.

These include moving the Exchange application between servers while maintaining uptime. In addition, just because we use a smaller server such as a blade server, there is no guarantee that we will maximize utilization of all components. Tests have shown that the HBA connecting the blade server frame can be a performance bottleneck as it is shared across many servers. CPU and memory may still be under-utilized. This negates any advantage a consolidation to blade servers may offer.

But what if a large, conventional server could operate 2 or 3 Exchange servers while maximizing utilization of memory, CPU, network, and storage? You would enjoy the cost savings associated with consolidation while still maintaining a robust Exchange environment.

There are many more reasons while you would virtualize an Exchange server and as we learn more about virtualization and consolidation, this list will continue to grow. It only takes thought, imagination, and a lot of testing to leverage a virtual infrastructure to your advantage.

Conclusion

This document has focused on the key areas an Exchange architect must focus on when initiating an Exchange 2007 infrastructure design. It is meant to be high-level as each area could have its own article. In fact, many do. However, design is easy if you ask the correct questions when working with the business unit you support. In fact, part of the fun of architecting solutions is that as long as you meet the requirements and stay within a set budget, it is up to you to decide what technologies to use to meet the criteria.

There are many technologies that can and will co-exist in any solution. The architect's responsibility is to ensure that each piece of the solution works with the other pieces. This can only be assured through experience and testing. Validation is a key component of solution design; and is done in a lab prior to deployment into production.

Remember, a simple solution is easier to install and support than a complex one. The more moving parts a solution has, the more complex it becomes. Do not get so caught up in the technology that you lose sight of the solution end state. Many great designs have failed to meet their potential because there were so many features/functions added that weren't part of the requirements. Adding unnecessary functionality might just be the proverbial straw that impacts performance negatively. Collect the requirements, design to the requirements, build the solution, test it, and deploy to production.

Appendix

The following hyperlinks will direct you to documents and websites offering knowledge and tools to assist you in developing a robust Exchange 2007 design.

Microsoft

Exchange Homepage http://www.microsoft.com/exchange/default.mspx

Exchange Team Blog http://msexchangeteam.com/

Microsoft Tools for Exchange Sizing and Analysis such as Jetstress and LoadGen http://technet.microsoft.com/en-us/exchange/bb330849.aspx?wt.svl=2007resources

EMC

Exchange Solution Homepage
https://powerlink.emc.com/nsepn/webapps/btg548664833igtcuup4826/km/appmanager/km/secureDesktop?_nfpb=true&_pageLabel=image6b&internalId=0b014066801ea5b8&_irrt=true

VMWare

Virtualizing Exchange 2007
http://www.vmware.com/landing_pages/exchange_solution.html