

Commentary

July 16, 2007

Enterprise Exchange Server 2007 Infrastructure: What Changes, What Doesn't

Microsoft Exchange Server 2007 is a powerful, functionality-rich option for today's enterprise messaging requirements. But since IT organizations have to move to a new 64-bit server environment anyway they are taking advantage of looking at their overall messaging architecture.

What changes and what doesn't change? The answer is the more things change, the more things stay the same. Enterprises will take advantage of the new functionality in Exchange 2007, but at the same time preserve the flexibility and power of their existing IT infrastructure, such as networked storage.

Start Your Planning Engines for Microsoft Exchange Server 2007

Microsoft Exchange is an essential business application. Period. Everything about Exchange is deeply interwoven into the business process fabric of any organization that uses it. No wonder then that enterprises will give the new version of Exchange — Microsoft Exchange Server 2007 — all due consideration in their planning process.

And enterprises will find a lot to like about Microsoft Exchange Server 2007 from throwing off the scalability shackles to greater protection to more user-friendly capabilities to greater deployment options. But enterprises are wrestling with what are the appropriate IT infrastructure changes aligned with Exchange 2007's native

64-bit application and new functionality. Enterprises should examine very carefully their choices as to what will remain the same in their IT infrastructure and what will change.

Please note that Microsoft's mission is to optimize the user experience with Exchange 2007. Although Microsoft provides a basic infrastructure foundation, Microsoft relies upon others to provide the deep capabilities necessary for a complete enterprise solution. Enterprises will have to plan their migration to Exchange Server 2007 carefully.

Exchange 2007's Role in Data Protection

A key focus of Exchange Server 2007 is business continuity through improved data protection for both operational recovery and disaster recovery issues. Exchange Server 2007 features two new

capabilities — Cluster Continuous Replication and Local Continuous Replication — to provide enterprises a new choice in achieving that goal of improved data protection. These new options will likely be used with or compared to other proven third party solutions to extend data protection to the required levels.

Cluster Continuous Replication (CCR)

CCR replicates both Exchange servers and Exchange databases. It delivers both service redundancy (a characteristic of traditional clusters) as well as database redundancy for high availability.

CCR is a two node system. An active node has its own storage as does the replicated passive node. (Multi-node systems will need multiple CCR instances – a key consideration in large environments.) The passive node updates by pulling full transaction logs from the active node asynchronously, so-called log shipping.

The key benefit of CCR is a low recovery time objective (RTO), i.e. high availability in terms of a downtime measured in minutes through failover to the passive node in case of a physical failure on the active side. That is really a restart rather than a restore that tape would provide that would take much longer.

Note that CCR only restarts from physical failure; it does not deal with logical failures, such as corruption of the database for any reason. Thus, even in a CCR environment, point-in-time backups (traditional or snapshot) from disk or tape provide an extra layer of both physical and logical data

protection and remain a necessity in an enterprise environment.

CCR supports different storage options including direct-attached storage (DAS) and storage area networks (SAN). DAS represents non-shared storage whereas SAN represents shared external storage —normally via Fibre Channel (FC) or IP connectivity. DAS storage was not viable for highly available deployments prior to CCR, as Exchange clusters used a shared storage model. General storage considerations — DAS vs. SAN — will be discussed later; however replication oriented features of advanced SAN storage still affect CCR planning.

CCR is an asynchronous technology, meaning a small data loss is likely during a disaster recovery scenario. Many enterprises who depend on Microsoft Exchange for critical business processes may not find that acceptable, say for revenue-producing or legally-monitored transactions (such as broker orders).

Third party synchronous technologies are available — often as a feature of a shared storage array — as an alternative and can provide a no data loss scenario. This is how many mission critical Exchange deployments are architected already, and would be just as viable in Exchange 2007.

Even if enterprises do not require synchronous protection today, the potential for a difficult move from CCR/DAS to a consolidated synchronous protection methodology is daunting. Customers need to consider flexible protection scenarios that match potential business requirement changes early in their planning cycle.

Local Continuous Replication (LCR)

LCR is similar to CCR from a database duplication perspective, but servers are not replicated so LCR is only a local solution.

LCR uses the same log shipping technique that CCR uses. If there is a catastrophic failure of the active copy of the database or the logs, an administrator could quickly activate the passive copy of the data manually. Enterprises that find CCR too expensive and do not need server or site redundancy could look at LCR as an alternative to potentially lessen traditional backup requirements.

Standby Continuous Replication (SCR)

With Microsoft already talking about Exchange 2007 SP1, a new replication option is on the horizon. While documentation on this still-not-released enhancement isn't robust, it introduces more flexibility in replicating data within Exchange by including a stand-by server (or cluster) option to replicate to. However, the fundamental choices – synchronous protection, Exchange-only or enterprise-wide replication, etc. remain. The new features in this release will improve your capabilities if CCR is for you, but probably wouldn't change your decision on whether to deploy it in the first place.

Storage Architecture Considerations: SAN vs. DAS

Technically, Exchange Servers can use direct-attached storage (DAS) in a back-to-the-future move from the generally accepted shared storage in the form of a storage area network (SAN).

As mentioned earlier, the potential of using CCR to address local failures without a shared storage cluster makes this feasible in a more highly available environment.

The move to 64-bit servers dramatically breaks the scalability and performance constraints of prior versions of Exchange. Notably large server caches (when combined with internal changes to Exchange) reduce the I/O load on storage since many more requests are now handled from server memory. (Note the impact of the I/O reduction depends heavily on the write/read ratio as read requests are what is served by large server memory.)

Customers are discussing DAS deployments because of the implied direct purchase costs savings of non-shared storage. But the same reasons that propelled the use of external shared storage in the first place still remain solidly in place.

Managing storage as a consolidated shared resource has a number of benefits (in addition to TCO) over managing a series of different application-silos as storage islands:

Improved configuration management

Economies of scale apply to the management of centralized storage over that of decentralized storage for tasks such as provisioning, tiering and planning growth. No longer must administrators "chase" problems from server to server.

Pooling of storage through an FC or iSCSI network makes growth planning easier since trying to estimate storage requirements for DAS is inexact at best. With network storage over and under estimates for storage provisioning are more easily balanced out.

Improved data protection:

Applying global practices consistently across all servers and environments for operational recovery and disaster recovery purposes ensures better business continuity.

Exchange is an essential application for most companies, but IT departments need to manage multiple key applications at the same time.

Managing pooled network storage across applications rather than managing individual application silos provides consistency that leads to efficiencies in provisioning and consolidation, business continuity and disaster recovery planning, and overall management (both proactive maintenance activities and problem resolution.)

This structured infrastructure approach is probably the largest detractor to application-specific infrastructure decision making in a large enterprise.

Enhanced scalability and performance:

Shared storage has advanced architectures that maintain high performance as environments scale. Pooled storage is better able to handle performance fluctuations more easily as Exchange storage groups are spread out over more storage system resources like disk spindles, cache etc. These means performance "hot spots" can be better balanced.

Greater flexibility:

SAN storage offers increased deployment options with powerful software for replication (e.g. synchronous mirroring as discussed prior), tiered storage and workload balancing are available. Tiered storage also be-

comes a much more viable proposition within a SAN environment.

Fundamentally, if an IT organization already deploys network storage for Exchange, protection of existing investments as well as avoidance of switching costs to a new architecture are likely to weigh heavily in favor of continuing this internal best practice.

Exchange 2007's Role in Compliance

Microsoft has introduced an increased focus on compliance with Exchange Server 2007. The three main capabilities that Exchange 2007 focuses on are:

- *Message retention* — the ability to retain e-mails for a set period of time
- *Controlled access* — prevent unauthorized disclosure or certain types of data, such as social security numbers, among other capabilities
- *Information and process integrity* — classify e-mails to prevent inappropriate distribution and to ensure the proper copying of compliance-related e-mails

These Exchange Server 2007 built-in capabilities are welcomed as *part* of a comprehensive compliance process. Depending on an individual company's compliance requirements, solutions that extend Exchange functionality will likely be deployed as well.

At the heart of its new capabilities are transport rules. These are mail flow controls that can be implemented globally across an enterprise by administrators and compliance officers using a new transport rules policy engine. Corporate policies on internal or outbound e-mail can be enforced for compliance purposes.

One example is that certain groups of employees may be prevented from exchanging e-mails with one another, say brokers and research analysts (a so-called ethical wall requirement). Another example is that e-mail that is classified based on certain topics may be automatically archived or routed to a compliance officer for review. Enterprises should find the power and flexibility in this functionality very attractive in meeting the key compliance requirements of controlled access and information and process integrity. So Exchange Server 2007 is especially strong on helping organizations ensure that they do compliance right in the first place. However, Exchange Server 2007 is not the proper venue for actually storing compliance data for e-Discovery and auditing purposes. For that an archive is necessary.

Active Archiving: A Necessity for Compliance

An active archive is a separate distinct copy of all relevant Exchange data. Two key characteristics of an archive are that it has completeness of data and that it is data-retention managed.

“Completeness” means that all the data — except that data legally deleted through policy — from all Exchange servers is available all the time. That is necessary because legally an enterprise has to be able to definitely state to a court of law that all data has been made available.

“Data retention managed” means that the retention of data is managed by tamperproof policy. For example, selected data may be put on litigation hold either at the time of ingestion into the archive or at a time designated by legal counsel. The data has then been

put under a chain of custody and cannot be altered, mutilated, or destroyed — and thus can be authenticated. Only authenticated data is admissible as evidence in a court of law.

Since full archives offer completeness and are data-retention-managed they are therefore mandatory for both compliance and governance, such as litigation holds, requirements.

Key Benefits from a Full Archive

Records management. From a records management perspective, Exchange can move compliant messages to a managed folder on an Exchange server. Exchange does this through an automated process that can scan the inbox and designated folders to retain, expire, or journal messages based upon an organization's compliance requirements. However, the best practice for managing chain of custody is to give external control of functionality to independent software rather than under the control of software where information is created and destroyed.

Journaling. Exchange Server 2007 has new and flexible journaling capabilities where journaling can be triggered per database, user, or by transport rules. Use of this helpful new capability will need to be weighed against using a full archive that can typically capture data without the performance impact of Exchange journaling. Moreover, archiving data (especially e-mail attachments) greatly reduces the data on an Exchange Server — through short-cutting/stubbing and having single instance storage (i.e., saving only one copy of an attachment) capabilities across multiple Exchange servers. This drastically reduces backup-window concerns by managing Exchange Server capacity.

Journaling to SharePoint. With Exchange 2007, the Exchange journal can be archived to a Microsoft SharePoint services site. SharePoint is a well liked collaboration tool, but it was not designed to be an active archive. Compared to purpose-built archival tools, it lacks some important features:

No capability for tiered storage. SharePoint is created for actively changing data. Archives however, can utilize less expensive tiered storage for a greatly improved TCO.

Potential object limitation problems. While not typically an issue in a collaboration environment, these limits could easily be reached in an e-mail environment with very many small objects. Purpose-built archiving products typically do not face an object limitation constraint.

Lack of short-cutting. Neither SharePoint nor Exchange offer built-in shortcutting (a.k.a. stubbing). A shortcut is a link to an object that a user on an Exchange Server can use to retrieve an object transparently that is really in an active archive. That means that the object does not have to be duplicated at both the Exchange server and archive locations — eliminating a good deal of storage overhead, especially with large attachments. This also means that within the single Outlook user interface, users can access vast amounts of content — even if some of that content physically sits in the archive — an important productivity consideration.

Mailbox Searches. Exchange Server 2007 has a strong fully indexed content search capability across all the

mailboxes on a single Exchange server. But users may turn to external software to add search across multiple Exchange servers and the ability to demonstrate completeness during a search.

Exchange Hosted Archive. Microsoft has attempted to address many of these challenges, by offering a subscriber service called Exchange Hosted Archive. If hosted solutions are a preference for your organization, this is something you can compare to third party archiving products to determine which best meets your needs.

Conclusion

Microsoft Exchange Server 2007 is very appealing in its advancements around security, user productivity (unified messaging) and flexibility. At the same time, it represents a challenge to IT organizations as they would need to fully plan both software upgrades and new features, but also a hardware upgrade to 64-bit servers.

Since they are doing a technology makeover anyway, many IT organizations are taking a close look at their overall messaging infrastructure to optimize them for ever-changing business requirements for high availability and compliance. IT wants to cut the infrastructure diamond carefully — measure many times, but cutting only once. So how do they cut the diamond? What changes and what remains the same?

Some enterprises may consider deploying Exchange's built-in replication features with cheaper, decentralized storage for potential upfront cost advantages. However, many others will see that managing storage and high availability as a global resource across multiple applications as part

of an enterprise infrastructure has benefits over trying to manage storage and data protection on an application-silo by application-silo basis.

From a compliance perspective, Exchange Server brings to the table strong capabilities (especially through the use of transport rules) to help enterprises with front-end compliance. However, an infrastructure approach to active archiving at the back-end of Exchange for compliance delivers, among other things, the necessary authentication of data that enterprises require. Active archiving is a necessary complement to Exchange 2007 for compliance purposes and also offers significant potential TCO reductions.

So the more things change, the more they stay the same. By all means, take advantage of the benefits of Exchange 2007 as appropriate, but complement Exchange 2007 with an infrastructure that ensures you can meet ever expanding business requirements. For many enterprises, that will likely include networked FC and iSCSI SANs, proven disaster recovery and replication tools and a unified, externally controlled archive.

David Hill

Analyst Name: David Hill
Topic Area: Infrastructure

Mesabi Group LLC
26 Country Lane
Westwood, MA 02090
www.mesabigroup.com

Mesabi Group LLC is an affiliate of Valley View Ventures that aims to provide thought leadership and sound advice to both vendors and users of information technology. .

Phone: (781) 326-0038
email the author: davidhill@mesabigroup.com

The information contained in this publication has been obtained from sources Mesabi Group LLC believes to be reliable, but is not warranted by Mesabi Group LLC. Commentary opinions reflect the analyst's judgment at the time and are subject to change without notice. Unless otherwise noted, the entire contents of this publication are copyrighted by Mesabi Group LLC, and may not be reproduced, stored in a retrieval system, or transmitted in any form or by any means without prior written consent by Mesabi Group LLC