



MICROSOFT EXCHANGE 2007

CHANGING THE STORAGE LANDSCAPE FOR MESSAGING

Matt Baker
Storage Technology Strategist
Dell, Inc.

Abstract

Microsoft® Exchange 2007 includes new features and capabilities that are beginning to help reduce the costs associated with e-mail storage. At the same time, Serial Attached SCSI (SAS)-based storage systems are delivering high performance and high capacity at a low cost. These innovations combine to enable an affordable increase in mailbox sizes, and new productivity for e-mail users. Together, they're changing the economics of messaging storage.

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

Dell, PowerVault and PowerEdge are trademarks of Dell Inc. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell disclaims proprietary interest in the marks and names of others.

©Copyright 2007 Dell Inc. All rights reserved. Reproduction in any manner whatsoever without the express written permission of Dell Inc. is strictly forbidden. For more information, contact Dell.

Information in this document is subject to change without notice.

Storage Spending: Still Growing

The overall cost of storage, measured in dollars per GB, has been dropping, and continues to drop steadily. But storage spending continues to accelerate, driven by compliance regulations, security issues, and a new consciousness of the need for disaster recovery. These three factors are spurring data growth, requiring customers to increase their storage spend. Some estimates place this growth at 15 to 20 percent annually, just to meet the swelling wave of data.

For many organizations, storage capacity requirements are estimated to be increasing between fifty and one hundred percent annually, creating a huge gap between their required storage spend and their IT budget, and leaving them struggling to find a way to achieve balance.

Dell believes it's time to change the economics of storage, and in conjunction with partners such as Microsoft, EMC Corporation, Symantec, and CommVault, we've set a goal to do just that. One specific area of storage spending that's ripe for change is the storage associated with messaging—Microsoft Exchange storage.

E-mail's Growing Importance

Rooted in academia, e-mail started life in the business world as a tool to facilitate simple asynchronous communications and collaboration. But the convenience and power of e-mail communications caused its use to proliferate, and capability of attaching documents for review caused e-mail to evolve from a mere tool for communications to an application for collaboration. At the same time, purpose-made collaboration tools were proving complex and cumbersome, encouraging teams to make e-mail the default collaboration tool, and ultimately entrenching e-mail as the single most important tool in today's information economy.

Simply put, e-mail is now mission critical. It's the lifeblood of a global organization; for the corporate "road warrior" and for every information worker anywhere. It's hard to imagine a project or program where team members don't have e-mail. Today e-mail is far more than a communications or even a collaboration tool. E-mail has become the user's de facto personal storage system, and most users feel content that the content of their mailbox is protected and available at all times.

Taking It To The Limit

Unfortunately, the operation of large individual-user mailboxes has been both complex and expensive, and has led many a worker to dread the "size limit exceeded" warnings that have become a business reality.

The expense associated with large individual mailboxes has been largely driven by the requirement to maintain a robust shared-storage backend to support data protection and clustered operations. The pervasive use of e-mail, particularly that provided by large Web players who support free and massive e-mail quotas, puts pressure on IT departments to provide the required storage. Moreover, these large mailbox sizes are not a luxury, but actually justifiable as a valid enhancement to workers' productivity—so the pressure to grow storage cannot be ignored.

Exchange 2007 and Storage

Microsoft Exchange 2007 has the potential to change the e-mail storage equation. Developed with the goal of driving cost and complexity out of the system, it was designed to enable large, low-cost mailboxes, offer more storage options, and increase reliability of the messaging system.

In targeting cost and complexity, Microsoft noted several trends to be dealt with. First, users are sharing richer media via e-mail, consuming more capacity. The average message size is increasing, and there is demand to access e-mail from any type of device (for instance, cell phones and PDAs).

In addition to these end-user trends, Microsoft noted that the use of shared storage for e-mail stores results in high complexity and costs.

Given its goals, Microsoft sought to change a few critical architectural areas in Exchange 2007, with the focus on reducing I/O load and building in enhanced storage flexibility.

64-bit Architecture

Microsoft has launched Exchange 2007 as an x64 system only. A 64-bit architecture enables significantly large system-memory sizes. Moreover, it increases the amount of data moved with each memory copy operation, making it more efficient. Additionally, Microsoft tuned the I/O path to make the most of every operation. They increased queue depths and page sizes, and devised enhanced techniques for I/O coalescing (the grouping of I/Os before processing), with the result of reducing the I/O load associated with Exchange operations.

Microsoft states that these changes drastically reduced the number of IOPS, or I/Os per second, required per user. IOPS are the transactional operations associated with mail storage and retrieval, and reducing the IOPS required per user reduces the number of drives required to support a given set of users; which directly reduces solution costs. Combined with the availability of new storage systems, these changes can have a significant impact on the cost of an Exchange 2007-based system.

Storage Flexibility

The single biggest change in this area is the addition of new clustering options. In prior versions of Microsoft Exchange, the only option for building a highly available mail system was the utilization of a shared storage cluster. Available solutions, such as cluster-enabled server-internal RAID systems with JBODs, presented customers with significant complexity and management challenges. Fibre Channel SANs, another option, were far more reliable but generally expensive to operate and daunting to some IT shops.

Exchange 2007 offers up a number of new options that reduce the barriers to adopting high-availability Exchange systems by reducing the cost and complexity associated with the underlying storage.

Exchange 2007 Data Protection Solutions

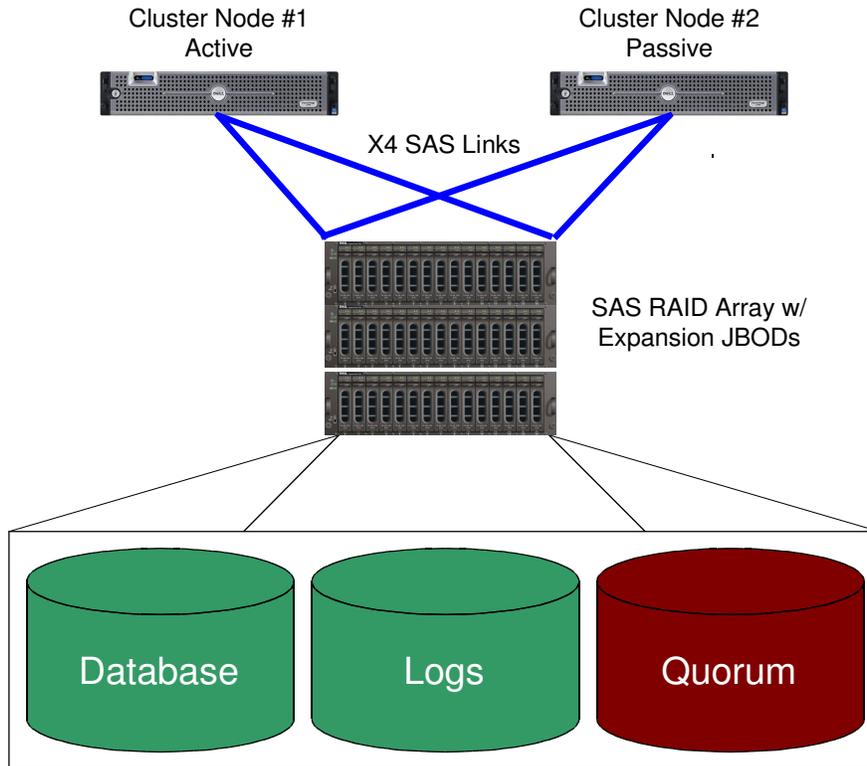
Single Copy Clusters

Single Copy Clusters, or SCC, are the traditional type of cluster system available on previous versions of exchange. SCC requires a storage system that facilitates simultaneous arbitrated access by two or more servers.

With the introduction of Service Pack 2003 for Windows Server, Microsoft eliminated support for clustered storage systems built with in-server RAID controllers. SCC requires an external storage system with RAID.

With Exchange 2007, SCC supports only Active/Passive configurations.

Single Copy Cluster

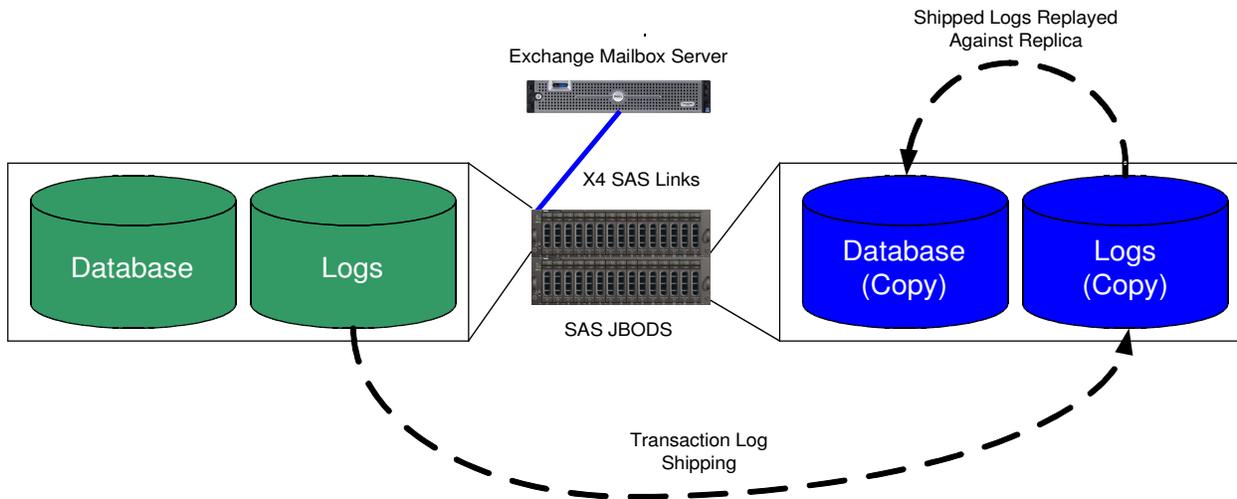


Three data-protection solutions are enabled as new features of Exchange 2007: Local Continuous Replication (LCR), and Cluster Continuous Replication (CCR), both of which are currently available; and Standby Continuous Replication (SCR), a WAN-capable solution designed to complement SCC or CCR, planned for Exchange 2007 Service Pack 1.

Local Continuous Replication

LCR, while not a cluster, facilitates local replication of mailbox data. This solution is designed to protect against data loss associated with storage-level corruptions and disk failure. With replication taking place higher in the stack, a clean copy of the data is maintained on a separate storage partition. The process is facilitated by replay of transaction logs against another shadow mailbox. In this solution, Microsoft reduced the log file size to accelerate the log shipping process and lessen the risk of data loss associated with failures taking place between log shipment operations.

Local Continuous Replication (LCR)



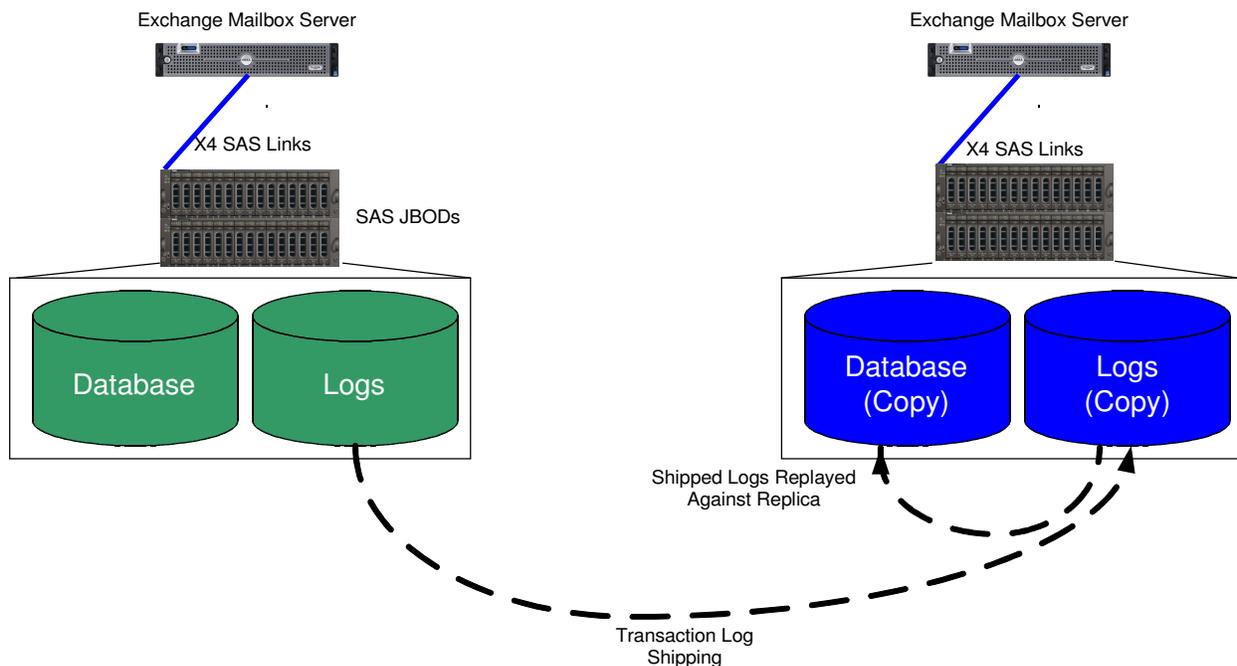
Cluster Continuous Replication

CCR is perhaps the most talked about and significant change to the clustering capabilities with Exchange 2007. Utilizing the same replication technique as LCR, CCR maintains a full copy of the active mailbox on a secondary server and storage system.

CCR offers the best of both SCC and LCR; a true high-availability system, plus protection from storage sourced errors and faults. Failover in the CCR model is automatic, similar to that offered by SCC. CCR, however, requires an external system, called a witness, to serve as the tiebreaker during failover operations. The witness server can service multiple CCR pairs and can be entirely secondary to the system on which it runs. In other words, the witness is a small process that can run on virtually any other system in the network.

Today, a CCR cluster can be stretched between two datacenters only by stretching the subnet. SCR, or Standby Continuous Replication, a feature planned for Exchange 2007 Service Pack 1, extends the CCR concept over a WAN to provide a standby copy of the data at a remote site. SCR will eliminate the stretching of the subnet and provide a third copy of the data, but will not provide automatic failover to that remote copy.

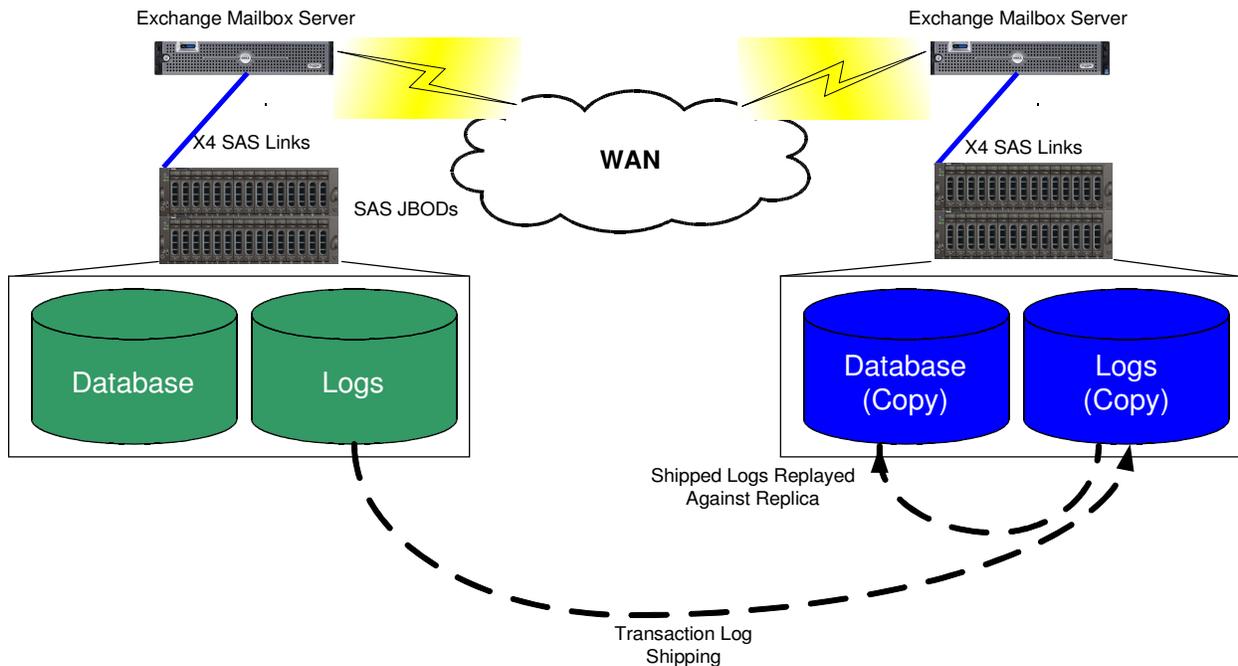
Cluster Continuous Replication (CCR)



Standby Continuous Replication

SCR Standby Continuous Replication is a feature planned for Exchange 2007 Service Pack 1. In essence, it extends the CCR concept over a WAN to provide a standby copy of the data at a remote site, to facilitate disaster recovery solutions. The SCR can protect either SCC or CCR primary solutions at the active site or even a non-clustered system at the primary active site. SCR engages only in the event of a complete failure of both clustered nodes at the primary site. Like LCR, failover to the standby SCR replica is not automatic.

Standby Continuous Replication (SCR)



Serial Attached SCSI and the New Economics of E-mail Storage

Another recent development in the storage equipment space is the introduction of Serial Attached SCSI, or SAS. SAS-based solutions provide a number of benefits over parallel SCSI, including high performance and low latency at an affordable price, improved capacity and expandability, and flexible support for heterogeneous storage opens.

SAS-based JBODs can deliver up to four times the throughput performance of previous SCSI-based JBODs, a bandwidth improvement capable of satisfying demanding e-mail applications running on a single server.

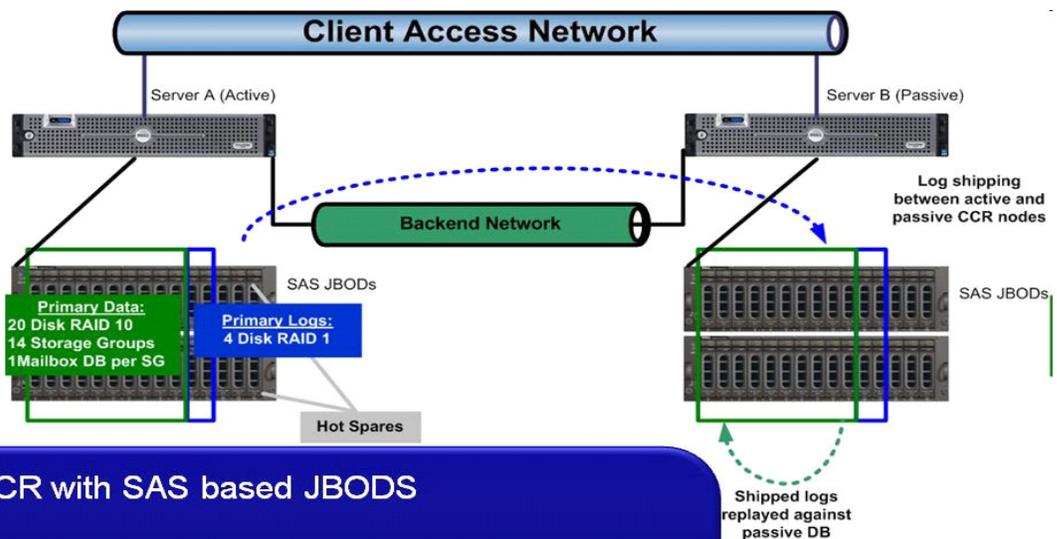
SAS includes a number of architectural enhancements that make expandability simpler and support intermixing SAS and SATA drives. Daisy-chaining drive enclosures enables deep, high-performance storage systems that are highly reliable and easy to use. SAS/SATA drive intermixing opens new opportunities for in-box storage tiering and disk-based backup.

SAS offers new levels of performance and availability at an entry price point. For example, highly available SAS RAID arrays, capable of supporting SCC clusters, for 8th and 9th generation Dell™ PowerVault™ MD3000 arrays can offer performance up to 92K IOPS, with up to 1,400 MB/s throughput; expansion up to 45 drives/13.5TB of raw storage, and optional integrates snapshots and virtual disk copy software.

Combining the Power of SAS and Exchange 2007

When SAS technology works in tandem with the advanced features of Exchange 2007, the benefits are significant. New clustering features like CCR enable highly available Exchange deployments without requiring complex and expensive shared storage systems. New SAS-based RAID arrays are changing the economics of deploying traditional SCC. Affordable, high-capacity SAS drives in expandable enclosures, coupled with the Exchange 2007's improvements to the I/O profile, are now helping to make low-cost, highly reliable, and very large mailboxes a reality.

Customers can use SAS-based JBODs and host RAID controllers to build a robust CCR solution. Figure n shows an example of how such a solution might be deployed to support 2,000 users with a 1GB quota per mailbox.



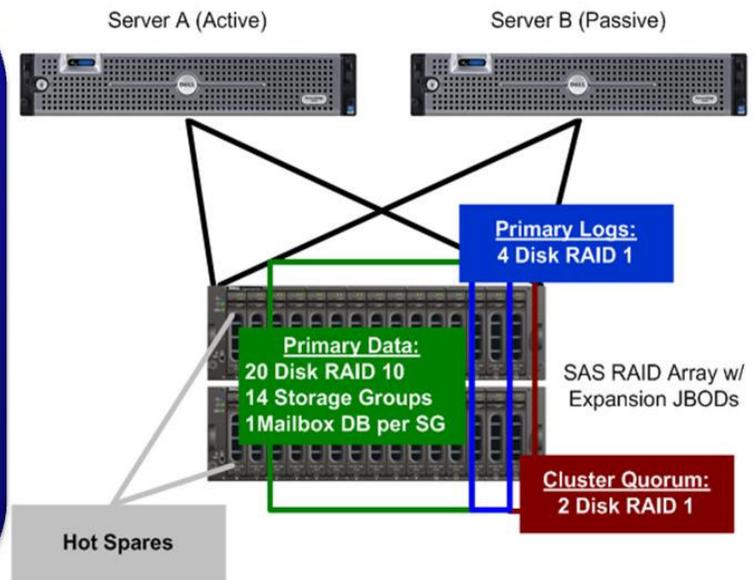
Exchange 2007 CCR with SAS based JBODs

- Cluster set supports 2000+ users with 1GB mailbox quotas
- Two Dual Socket Servers w/ Two SAS JBODs each
- 30 300GB SAS 10K HDDs per Node
 - 20 Drive RAID 10 (DB)
 - 4 Drive RAID 1 (Logs)

The Dell PowerVault™ MD3000 represents a new class of SAS-based RAID arrays designed with two-node application clustering in mind. This new class of SAS-based RAID array can significantly increase performance and reliability while helping reduce the overall cost of the solution.

Exchange 2007 SCC with SAS Based RAID Arrays

- Cluster set supports 2000+ users with 1GB mailbox quotas
- Two Dual Socket Servers w/ SAS based RAID array w/ Expansion Shelf
- 30 300GB SAS 10K HDDs
 - 20 Drive RAID 10 (DB)
 - 4 Drive RAID 1 (Logs)
 - 2 Drive RAID 1 (Quorum)



Changing the Storage Landscape

The data explosion is not likely to abate anytime soon, but significant enhancements and new features of Exchange 2007, combined with the new level of attractively priced performance and functionality of SAS-based storage systems, are helping IT shops to balance their storage needs with their budgets. The increased productivity afforded by larger mailboxes is no longer in conflict with the cost of the solution, because new Dell SAS-based solutions paired with Exchange 2007 are changing the economics of storage for messaging.