

Continuous Cluster Replication and Direct Attached Storage: High Availability without Breaking the Bank

Missy Koslosky

Sr. Exchange Architect

3Sharp

March 2, 2009



About the Author

Missy Koslosky is a Senior Exchange Architect for 3Sharp, LLC and was previously awarded status as an Exchange MVP nine times. Missy has managed Exchange Server organizations for the federal government, designed Exchange architectures, and worked as a consultant, helping organizations architect and deploy both Exchange and Active Directory. She is a coauthor of the [Exchange Cookbook](#) (O'Reilly, 2005); she has edited various books on Exchange for Microsoft Press, O'Reilly, Digital Press, and Sybex. She has written numerous Microsoft white papers including *Supporting Regulatory Compliance with Microsoft Exchange Server 2003*, *Integrated Solutions for Regulatory Compliance with Windows Server Technologies*, and the *Solution Accelerator for Exchange Migration and Consolidation*.

About 3Sharp

3Sharp is a leading technology service provider focused on enabling its customers to get more value out of the tools they already know and use. 3Sharp's solutions begin by leveraging the extensive out-of-the-box capabilities of the Microsoft Office system, including SharePoint, the Office client applications, Exchange Server, and Office Communications Server. 3Sharp incorporates custom code, templates, workflows, and other components that integrate with its customers' existing line of business systems so that information workers can do everything they need to do without leaving familiar and productive Office environments.

3Sharp is located at 14700 NE 95th St, in Redmond, Washington and was founded in 2002 by Peter Kelly, Paul Robichaux and John Peltonen. Visit their website at www.3Sharp.com

3Sharp was contracted by Microsoft to write this paper.

All product and company names herein may be trademarks of their registered owners.

Table of Contents

- Introduction 5
 - Current Business Needs 6
 - Storage Options 6
 - Large Mailboxes 7
- Solution Simplicity..... 9
- Solution Effectiveness 10
- Capital and Maintenance Costs 12
 - Direct Attached Storage Costs 12
 - Storage Area Network Costs 13
- Power and Cooling 14
- Manageability 14
- Performance and Scalability 15
- Conclusion..... 16

Introduction

When Exchange Server 4.0™ was introduced, internal drives (or, for the well-funded, direct attached storage) were widely used. When I first started working with Exchange in the mid-1990s, many of our organization's Exchange Servers used 2GB internal drives; these servers provided file and print services, and also ran SMS. For a few of our more robust servers, we had two 2GB drives, and the database resided on one drive, the log files on the other. None of our servers used RAID drive sets, and clustering was a much-discussed dream.

When clustering was eventually introduced for Exchange, it included shared storage for the mailbox databases and only protected Exchange Server from a small subset of issues, such as blue screens. Mailbox size limits were kept small, as both backup and recovery processes needed to take place within a short time window, and Exchange Servers were limited to a single database, so there was no option to split an organization's mailboxes across multiple mailbox databases without deploying multiple physical servers.

Those days have long since passed; storage options now abound, and disk space is cheap. Over the past ten years, storage area networks (SANs) have become ubiquitous, and IT management has often been persuaded that all data should reside on the SAN, as the SAN was seen to provide an additional level of reliability. Many IT staff are now focused specifically on SAN-based storage, and these Storage Administrators believe that disk is disk, and that the SAN can meet the needs of any application. However, the idea that the SAN is the best option for Exchange Server storage needs to be revisited; there are other options available today that provide the necessary reliability at greatly reduced cost.

Microsoft has made other improvements that impact what type of storage is appropriate for use with Exchange 2007. While there was much debate over their decision to develop Exchange Server 2007 for 64-bit systems and not for 32-bit systems, the new architecture provides expanded application access to RAM, which had been a bottleneck in previous versions. By taking advantage of the additional RAM availability, input/output operations per second (IOPS) were greatly reduced, reducing the load on physical disk up to 70%. The synchronicity between reduced IOPS and cheap disk space allows organizations to rethink their approach to storage.

There are other changes that can cause organizations to rethink the way they design and deploy Exchange Server 2007. After years of offering only one native method for clustering, Microsoft decided to introduce numerous high availability options with Exchange 2007. Instead just the simple failover clustering previously available, several of these options utilize additional mailbox database instances, which help to ensure that data will be available and protected:

- **Cluster Continuous Replication** uses asynchronous log file shipping to create an additional copy of the Exchange mailbox databases on a separate physical cluster node; the passive node automatically begins to serve users if the active node fails.
- **Standby Continuous Replication (SCR)** also uses asynchronous log file shipping to create an additional copy of the Exchange mailbox databases on a separate physical server; this option is

often used when organizations wish to have a server available in an alternate physical site. Failover to an SCR node must be performed manually.

- **Local Continuous Replication (LCR)** also uses asynchronous log file shipping to create an additional copy of the Exchange mailbox databases. LCR differs from CCR and SCR in that a single physical server is used to host both database copies.
- **Single Copy Clustering (SCC)** uses multiple server nodes to connect to a single copy of the Exchange mailbox databases to provide failover clustering. With this simple failover clustering, the clustered nodes provide protection from bluescreens and machine failure. The single database copy is a single point of failure and does not protect organizations from storage failure or database corruption.

The new continuous clustering options outlined above are a radical change from previous iterations of native clustering for Exchange Server. With continuous replication, multiple instances of the mailbox databases are created and the single point of failure in the shared storage infrastructure seen in SCC is gone. There is no longer a need to invest heavily in a storage infrastructure, because the data itself is now redundant. The white paper [High Availability Choices for Exchange Server 2007: Continuous Cluster Replication or Single Copy Clustering](#) provides a more in-depth discussion of continuous clustering and how it differs from SCC.

When considering which type of clustering is appropriate for your organization, you should examine the recovery time objectives (RTO), recovery point objectives (RPO), and cost for each solution in order to determine their return on investment (ROI). For most organizations, high availability requirements for Exchange 2007 are most cost-effectively fulfilled by deploying CCR with DAS. Organizations that require high availability with an additional safety net can deploy SCR in conjunction with CCR to provide continuity of service in the face of site failures. The [Exchange 2007 Mailbox Server Role Storage Requirements Calculator Spreadsheet](#) made available by the Exchange Product Group can help you determine the specific storage requirements for your organization.

Current Business Needs

Organizations rely on e-mail services to be up and available at all times, just like their phone lines. It used to be that using SCC with a SAN for storage was the best way to approach high availability for Exchange Server, as SCC protected e-mail services and the SAN protected the data. Everything has now changed, though, and it is time to really rethink the ways an organization can cost-effectively deploy Exchange in a highly reliable fashion.

Storage Options

With the release of Exchange Server 2007, Microsoft introduced CCR in order to help organizations cost-effectively meet their data protection and storage needs. CCR was designed to provide an integrated solution that addresses failures throughout the entire stack, from storage to the application, without the need for an expensive storage infrastructure, while eliminating the need to keep mailboxes at artificially small sizes. The [Continuous Replication Deep Dive](#) white paper from Microsoft provides an in-depth examination of the inner workings of CCR.

Many Exchange Architects and Administrators are focused on ensuring that messaging services are available and fully operational and match the needs of their organization, as e-mail is accessed by their users all day long and is an integral part of their business. With previous versions of Exchange Server, the options for keeping messaging services up and running were severely limited. Previously, the best solution to provide high uptime was considered to be SCC (simply known as “clustering” in previous versions of Exchange), particularly in conjunction with a SAN, which was used to provide reliable storage. However, and SCC and SAN-based solutions only protect data from a small range of problems, such as a single disk failure, and rely upon a shared storage infrastructure to host the mailbox databases. To achieve better protection, organizations using SCC felt the need to deploy additional services, such as synchronous SAN replication. In reality, SAN replication with Exchange is expensive and unnecessary for most organizations, and highly complex to implement and manage; many organizations could not afford to invest in a second SAN to host a remote copy of their data, or could only manage do so by implementing small mailbox quotas for their users.

Because SANs are so expensive to purchase and require specific expertise to implement and maintain, organizations needed to carefully balance between the need to provide data availability protection and the need to provide adequate amounts of storage for the increasingly important business data within mailboxes. While a solution based upon SCC and SAN replication was the best that could be created in the past, there are different options available today that protect critical data for a much wider range of scenarios.

The single, critical, point of failure in SCC is the mailbox databases. While two (or more) physical servers may be clustered, there is no second copy of the mailbox databases, which is why many organizations felt the need to add SAN replication. Even with SAN replication, there are multiple problems that can occur, including logical or physical database corruption, failures during firmware upgrades, HBA failures, storage controller failures, human error, full SAN outages, and power outages. While there are mechanisms to address many of these failures, such the [Exchange Server Database Utilities \(ESEUTIL\)](#), standard backup and restore procedures, and Volume Shadow Copy Service (VSS) clones, they provide a complex patchwork of solutions instead of a single integrated and simple failover mechanism that administrators can rely upon, no matter the cause of the failure.

The best option for keeping services available to users and avoiding downtime caused by database corruption is to have a second copy of the mailbox database itself, an option that was unavailable natively before the release of Exchange Server 2007.

Large Mailboxes

Today, organizations often keep mailboxes artificially small because of the high storage costs associated with using a SAN, and the time and effort necessary to backup and restore large volumes of data. This practice is in conflict with the needs of information workers, who are no longer content to work within the confines of small mailbox sizes. Because of the smaller investment required with DAS, organizations can implement large mailboxes with more than adequate space for long-term storage. An important benefit of large mailboxes is a reduction in the use of PST files by users; other benefits are described in the white paper [Planning for Large Mailboxes with Exchange Server 2007](#). While PST files may seem

innocent enough, there are many hidden costs associated with PST files that are not always tracked, including:

- PST files are stored on local machines, few of which are regularly backed up. This leaves data at risk for loss, without recoverability.
- PST files are not supported on network storage, as described in [MS KB 297019](#).
- PST files cannot be searched from a central location, making discovery time and effort intensive
- PST files are widely dispersed and their use can be hard to control and track
- Storing data in PST files takes approximately twice as much disk space as storing data on the Exchange Server, because messages are stored in both RTF and ASCII format in PSTs
- PSTs cannot be accessed by more than one device at a time; data on the Exchange Server is searchable not only through Outlook 2007, but also by using Outlook Web Access or Outlook Mobile
- PST files are insecure; the risk of having data in a PST accessed by a malicious party is very real, and sensitive data is not encoded or safeguarded against such an intrusion

There is no need to restrict mailboxes to small sizes when using CCR and DAS; storage is inexpensive, and quick VSS backups can run against the passive mailbox databases. Users no longer need to dedicate valuable time to weeding through their mailboxes to reduce their mailbox size when CCR and DAS are used, freeing them to concentrate on their actual roles within the organization more effectively.

The need to stay connected has grown with the popularity of mobile devices and telecommuting options. Instead of simply using Outlook from the desktop to access messaging data, users now use Outlook Web Access (OWA) or Mobile Outlook on Windows Mobile in order to keep in touch and access vital information. The mobile nature of the workforce has led to the need to store more data on the server where it may be searched for and found. If data is stored in PST files instead of on the Exchange Server, it is not accessible from any source other than the computer on which it is stored.

By using CCR and DAS, organizations can cost-effectively increase mailbox size limits to ensure that messaging data is always available and accessible. Organizations should examine the mailbox statistics related to their users, including messages sent and received per day, mailbox size, and PST files on desktops associated with users to get a better understanding of the needs of their users – and then size mailboxes so that users no longer need to resort to using offline storage to hold important data. Disk space is getting less expensive each day, and with a gigabyte of storage costing around 25 cents (US) today, the cost of allocating large mailboxes to users is no longer prohibitive.

Solution Simplicity

The design and setup of Exchange Server 2007 CCR on DAS is an easy process of putting together building blocks using the same process – get a server, add an array controller and cabinet, add disks, install the first node, then repeat for the second node to achieve full redundancy. This simple solution protects organizations from a wide range of problems that can occur across the deployment stack, including physical corruption of data, hardware failures (storage or server-based), network failure, and operating system or server software failures.

When using a SAN, there are numerous operational tasks that need to be performed, including:

- Allocating and de-allocating storage pools
- Fibre Channel zoning
- Hardware VSS backup integration
- Data replication management (in cases where SAN-based data replication tools are used)
- Balancing the requirements of Exchange with other applications to achieve peak performance
- Monitoring SAN array controller performance
- Monitoring storage cache utilization
- Changing HBA queue depth configuration
- Load-balancing multi-path storage configuration

Implementing CCR on DAS requires no specialized expertise beyond that of the average Exchange Administrator, and allows organizations to unify all Exchange-related management under one team. With CCR on DAS, Exchange Administrators simply design the system, deploy it, and let it run. In contrast, when using SCC and a SAN, multiple teams must be involved, and each possible point of failure must be individually identified and analyzed, and third-party products that fill the gaps must be evaluated and implemented.

Getting the system up and running with properly allocated storage and working replication is neither quick nor simple when using SCC, SAN, and SAN replication. In order to get SCC working with a SAN, the storage environment must first be properly configured to support the cluster; this involves ensuring the components within the complex storage stack are all working correctly together. Once the storage stack is properly prepared and the cluster is formed, the storage configuration needs to be rechecked and often modified. Setting up SCC is easier than with earlier releases of Exchange clustering, but additional

“The deployment (of CCR and DAS) has been smooth. Our DAS deployment is atypical since the passive nodes are running SATA drives for the database array while the active nodes use the faster SAS drives. We were able to save money using the cheaper SATA drives for the passive node while allowing us to validate this drive technology for Exchange databases.”

-University of Miami

tools beyond those native to Exchange are still required, and once the cluster itself is up and running, the configuration process necessary for real protection is not yet complete.

Because there are so many moving parts with CCR and SAN replication, there need to be more people involved in its management – not only are Exchange Administrators required, but Storage Administrators must also be involved, and consultants from third-party vendors are likely used for support and maintenance. There have been numerous cases reported to Microsoft support where human error, errant firmware upgrades, or physical corruption has caused data loss on a SAN – and these problems were dutifully replicated to the other copy of the data.

CCR, especially with DAS, is simpler to design and implement than SCC and a SAN, with or without additional third-party tools. SCC and SAN replication can be manipulated in an attempt to simulate the availability and reliability of CCR and DAS, but achieving the same RPO and RTO as with CCR requires implementing a complex assortment of backup, replication, and storage technologies and products. Getting a meshed SCC solution working correctly is complicated in comparison to the simple implementation of the native continuous replication used with CCR.

Solution Effectiveness

CCR eliminates the largest failing of SCC and previous Exchange Server clustering options – the single point of failure for the private mailbox databases. Instead of simply providing failover clustering, CCR gives organizations real data resiliency by enabling them to have more than a single copy of an mailbox database. By building Exchange Server 2007 with the type of true clustering long available in systems like OpenVMS from HP, Microsoft has found a solution that satisfies the needs of the mission-critical application that Exchange has become.

In order to achieve the same level of protection as natively available with CCR, SCC must have additional services configured and regularly managed. VSS technology is necessary so snapshots of the environment can be taken at regular intervals to mirror the data on disk for restore (and requires RAID-10 disk configuration to achieve similar RTO as with CCR; it cannot match the RPO point of CCR– with CCR, the RPO tends to be a lost contact, task, or flag – all things that are easy to reproduce).

“Since we live in a hurricane prone area, we were looking for a high availability solution for our growing user base, but did not want to incur the expense and complexity of a SAN. We evaluated other solutions for high availability, but once CCR became available on Exchange 2007, we concluded that high availability using commodity hardware in conjunction with an Exchange-based solution was the right choice.”

- University of Miami

When using SAN replication technology, the system must be configured to maintain a second copy of the mailbox database(s) on separate, isolated storage. Some SAN replication solutions make a block-level copy of the data on disk; errors on the source disk can thus replicated to the target, leaving organizations with two problematic data sets. In contrast, the continuous replication technology used with CCR (also with LCR and SCR) uses asynchronous replication, where each transaction log is inspected for corruption multiple times (approximately 15 checks are run) and any issues can be detected prior to the transaction being applied to the passive copy. Because of the different mechanisms used to make changes to the active and passive nodes (on the active node, the transaction is written from the cache to the database, and on the passive node, changes are made by replaying the transaction logs), disk errors on one cluster node are not recreated on the other node, so each server is isolated from one another in ways not seen with SAN replication solutions.

The advantages provided by the simple building block architecture of CCR on DAS cannot be discounted; whereas a SAN is a complex network of components, DAS provides a simple, intuitive storage platform that does not require constant changes and system updates. Instead of competing for storage resources with other applications, as happens so often with a SAN, systems using DAS perform optimally due to the dedicated nature of the storage. Table 1 displays many of the differences between various SCC and CCR configurations.

		SCC + 2 VSS Clones	SCC + 2 VSS Clones + Synchronous Replication	CCR on SAN	CCR on DAS	CCR on DAS + SCR
Complexity	Hardware	FC, HBA, VSS, Switches, Management Equipment, Cabling	FC, HBA, VSS, Switches, Management Equipment, Cabling, Dark Fibre, 2x SAN, Replication	FC, HBA, Cabling	Controller	Controller, IP Network
	Software	Cluster, LUN Resources, Multipath	Cluster, LUN Resources, Multipath, IP Subnet Span (Windows 2003)	Cluster, LUN Resources, Multipath	Cluster	Cluster
	Operations	VSS/Clone Scheduling, Restore	VSS/Clone Scheduling, Restore, Replication, Dark Fibre	Same as scheduled move	Same as scheduled move	Database Portability, Server Recovery
HW cost (storage)	Per User	\$\$\$	\$\$\$\$\$\$	\$\$\$	\$	\$
	Operations	\$\$\$	\$\$\$\$	\$\$	\$	\$
Total Physical Copies		7 = 2 + 4VSS (RAID10) + 1 off array ¹	9 = 4 + 4VSS (RAID10) + 1 off array	4	4	4 + 2N

Table 1: Clustering and Storage Options Comparison

¹ SCC includes RAID10 for the VSS clones in order to match the fast recovery provided by CCR. RAID5 could also be used, but may cause reduced return to operations (RTO) times.

Capital and Maintenance Costs

DAS has substantially lower implementation and maintenance costs than that of a SAN; capital expenditures for DAS can be as low as 10 percent of the cost of other solutions. The savings are found not only in the purchase price of storage, but also in the soft costs associated with ongoing support, management, and maintenance of the systems. Instead of needing services from highly specialized Storage Consultants to implement, configure, and optimize storage, organizations are able to use their internal Exchange Administrators to deploy the systems – no specialized storage knowledge is necessary. Microsoft IT was able to reduce costs by 74% by implementing CCR clusters using DAS for storage.

Microsoft did an analysis of expected capital expenditures over three years for different clustering and disk solutions and summarized the results in a presentation they called, “Going Big! Deploying Large Mailboxes with Microsoft Exchange Server without Breaking the Bank,” which was delivered at TechEd, Exchange Connections, and other conferences in 2008. They found that, for 1GB mailboxes², a solution based upon SCC and SAN cost an average of 37USD per user, while a CCR and DAS-based solution cost an average of 23USD per user. When mailbox sizes were increased to 2GB per user, costs for the SAN-based solution were raised to 44USD per user, but the cost for the DAS-based solution shrunk even further, to 17USD per user. In order to get an idea of the cost savings that could be achieved by your organization, the [Exchange Server 2007 Storage Cost Calculator](#) and enter information specific to your environment. This tool helps you compare differing storage configurations and their respective costs.

Direct Attached Storage Costs

When looking at the total cost of ownership of storage for Exchange Server 2007, DAS can actually cost as much as quarter of what a SAN would cost. The simplicity of design, implementation, management, and maintenance associated with DAS allows the internal Exchange Administrators to provide all necessary services, instead of using expensive, high-level Storage consultants to provide implementation and configuration services and relying upon internal Storage Administrators to spend time on management for the servers. Because of the nature of Exchange Server, it has always been important for Exchange Administrators to explain the specific needs of Exchange to Storage Administrators – and miscommunications or misunderstandings between these groups can lead to mistakes that have an adverse impact upon Exchange performance and operability.

The application experts – the Exchange Administrators – understand how to best benefit their Exchange infrastructure and can make the decisions that benefit the application itself. Storage Administrators do not understand the requirements inherent in Exchange Server 2007, nor do they tend to differentiate

“Our Exchange 2003 implementation was on a shared SAN infrastructure. Avanade has seen a positive ROI when taking “meshed” SAN infrastructure costs into consideration (after moving to DAS), maximizing uptime and lowering TCO.”

-Avanade

² Based upon 4000 mailboxes per server. The CCR/SAN solution was not geographically dispersed.

storage needs between applications. For Exchange Administrators, DAS is easy to provision and manage, and does not require complex configuration or maintenance – reducing costs not only initially, but also over time.

Implementing and maintaining CCR on DAS requires no special expertise and can be set up and managed by in-house Exchange Administrators. Achieving the same level of service, recovery time, and redundancy with a SAN or NAS solution is more challenging; a complex assortment of backup, replication, storage technologies, and products are required to be interwoven. With CCR on DAS, you design, deploy, and forget – the long term ramifications of using a SAN (especially if using a SAN with SCC and data replication) are much different.

Storage Area Network Costs

There are multiple options available to consider when making design decisions about what type of storage to use with Exchange Server 2007. When factoring the return on investment (ROI) of a SAN to support CCR (or SCC), some organizations feel that they have already made significant investments in SAN technology and want to only use the SAN, because they already have it, and there is space available that might as well get used. Consider whether you can achieve a much better recovery point objective (RPO) and return to operations (RTO) with the SAN than you could with a less expensive, less complex solution.

The actual cost of the SAN for storage is not only a matter of the initial investment in the storage technology – there are also inherent costs to the setup, maintenance, and monitoring of the storage. SAN implementation generally requires professional services from the SAN vendor, and ongoing costs for maintenance and management greatly exceed those of DAS. Many operational costs are actually reduced or removed when using CCR on DAS instead of using a SAN:

When using a solution consisting of SCC and Exchange Server 2007 coupled with SAN replication technology, the hard and soft costs can be 10 to 20 times that of CCR on DAS. The resultant solution should therefore achieve RPO/RTO numbers with similar returns – 10 to 20 times better than those provided with CCR and DAS. That type of ROI isn't possible with an SCC and SAN solution, as one physical database corruption or problematic firmware upgrade can result in a blown SLA for the year. With SCC and SAN replication, you simply can't achieve 10 to 20 times the RPO/RTO of CCR and DAS, and there is a real risk of not even being able to achieve equal numbers.

Often, management believes that storage is simply storage, and that applications should not be examined individually to find the solution that best meets their needs – the decision to use the SAN for all storage has been made and there are no exceptions to that rule. This is an erroneous viewpoint, as Exchange Server 2007 has very specialized storage needs, and was optimized to take advantage of the less expensive, more flexible storage offered by DAS.

Power and Cooling

The SAN storage infrastructure, including HBAs, fibre channel switches, SAN controllers, and other components, can consume significantly more power than a simpler, DAS-based solution. While it has been argued that power and cooling costs are higher with DAS than with a SAN, an architecturally correct DAS-based storage design can be more power efficient than a SAN-based storage approach (given the same levels of data protection), due to the following factors:

- 2.5 inch 15K SAS disks (DAS) are twice as energy efficient as 3.5 inch 15K FC disks (SAN). 2.5 inch SAS disks in a RAID-5 configuration can often meet or exceed the capacity and performance requirements of the Exchange while providing a lower power footprint than 3.5 inch FC disks in the RAID-10 configuration necessary on the SAN.
- 3.5 inch 7.2K SATA disks (DAS) are 1.8 times more energy efficient than 3.5 inch 15K FC disks (SAN). For large mailboxes (over one GB in size), 3.5 inch 7.2K SATA disks in a RAID-10 configuration can meet or exceed the capacity performance requirements of Exchange while providing a lower power footprint than 3.5 inch FC disks.

Manageability

CCR and DAS are less complex and easier to manage than SCC using a SAN, especially compared to a solution that includes SAN replication or other third-party tools. With CCR and DAS, there is no special expertise required, aside from that possessed by Exchange Administrators. DAS technology has changed, and there are new storage solutions available that make managing hundreds of DAS disks simple and intuitive. With DAS, there is no need to constantly revisit storage allocation amongst applications; DAS requires minimal maintenance in order to keep its components updated – generally, only the drivers and firmware for the array controllers and the firmware for the disks need be updated. This makes for minimal time spent on system maintenance. Server and storage operations can be managed by a single team, with no need to coordinate between groups with conflicting goals.

In the past, the slow performance of parallel SCSI solutions made DAS an unattractive solution – the limited number of disks that could be utilized limited performance. Today's faster and smarter SAS disks and arrays remove those old barriers to DAS deployment, and new management tools from vendors like HP and Dell make deploying and managing DAS disks simple.

Once expensive storage is deployed, organizations tend to over manage it in order to try to get the most ROI from the significant investment they feel they have made in their infrastructure. When you have paid 2,000 USD for a 300GB FC disk, then the amount of time in man hours that will be regularly invested to manage the system effectively seem pretty reasonable. But when you're paying 200 USD for a terabyte disk that can be used in a DAS array, then investing man hours in managing that storage seems wasteful. A SAN requires more management in order to achieve a reasonable ROI, while DAS requires less management, because ROI is so easily achieved

Microsoft IT has deployed Exchange Server 2007 using DAS for storage instead of the SAN that was previously used. They are more efficient now, as they perform routine maintenance tasks at the same

time as system patching. Microsoft has found very low failure rates for the SAS disks they deployed, and learned that operations run much more smoothly when the Exchange Administrators are able to control all of the components used by Exchange Server. For more information on the deployment of Exchange Server 2007 by Microsoft IT, see [Storage Design for Exchange Server 2007: How Microsoft IT Exceeds High Availability Targets at Low Costs Based on New Storage Designs](#).

With a SAN, there are multiple components that require maintenance; HBAs have drivers and firmware that must be updated, fabric switches need firmware upgrades, storage controllers may need upgrading, and firmware for the disks may need updating – and both Storage and Exchange Administrators still need to touch every server during this process. With so many moving parts, an outage is much more likely than that of a system with minimal components. When going through newsgroup and forum posts, it is not uncommon to read about problems caused by many of these components –these issues are often caused by human error.

One of the best ways to avoid human error is to put systems in place that do not need a lot of tending on a regular basis. Storage Administrators tend to regularly change settings, often to re-allocate storage or enhance system performance. Every time a system is touched, there is a chance a mistake can happen. Many of the most successful Exchange Administrators follow a hands-off policy that often isn't possible when using a SAN for storage – they configure the system, and then let it do its job. CCR on DAS allows organizations to design the system, deploy it, and let it run.

Performance and Scalability

CCR and DAS provide enterprise-level performance and scalability. DAS provides guaranteed stable performance, as you simply first design, then deploy, and finally leave the infrastructure alone. With a SAN, there is a good chance that the storage will be changed over time, which can cause a negative impact on performance and availability.

A DAS array using 200 SAS disks per server can provide over 40,000 IOPS with over 1GB of sequential IO throughput for each Exchange Server 2007 mailbox server. DAS can use Small Form Factor (2.5") SAS disks, with a proven 0.8% Annual Failure Rate (AFR), while SAN technology today only supports Large Form Factor SAS/Fibre Channel disks, with a much higher 5% AFR.

“Our service performance is reasonable, and our ability to fail-over (especially during patch maintenance) is a fairly simple process. In almost 2 years of having CCR and DAS in Production, we have not had a service impacting outage because of the CCR and DAS implementation.”

-Avanade

Ten years ago, SAN technology was not widely deployed or understood. Today, it seems that most organizations have some type of SAN deployed. *Storage Administrator* is now a job title, instead of a side role for an IT generalist. While SAN use is widespread, it should be noted that the SAN is not always

the best solution when it comes to the needs of Exchange Server 2007, which has much different storage needs than a file system or most other applications.

The specific needs of Exchange Server do not fit into the general mindset of most Storage Administrators, who view disk as disk and do not care about the intricacies of Exchange Server storage; this can lead to under-provisioned, overburdened systems that do not adequately protect important data. Exchange Server 2007 functions best when the disk volumes are optimized to meet its needs – the disk volume hosting the transaction logs needs to support ongoing sequential writes (when using continuous replication solutions, the requirements are a little different, as each transaction is read after it is written to disk so it can be written to the passive database copy) , while the disk volume used for the mailbox databases should be configured to support the random rapid reads and writes necessary for optimal performance. While the 64-bit architecture of Exchange Server 2007 has greatly reduced I/O needs, these specific disk configurations help Exchange Server perform optimally. In many cases, organizations that choose to use a SAN for storage deploy SCC instead of CCR; however, SCC is outdated, inherently flawed technology. With SCC, each mailbox database is held on shared storage, and a problem with this single point of failure can cause a long-term outage. To alleviate this problem, many organizations choose to use SAN replication to copy the mailbox databases to a second SAN; however this type of solution has many flaws and, in many cases, is not supported by Microsoft and only by the third-party vendors involved.

Conclusion

Every minute an outage lasts decreases the productivity of workers and loses money for an organization. Eliminating single points of failure and ensuring mission critical communications services are available under any circumstances helps keep organizations running smoothly. It is understandable that some may be skeptical of the advantages of CCR and DAS over those of SCC and SAN storage; Microsoft IT, when first introduced to the concept of CCR, was not convinced that the new technology provided the best solution. They then experienced a three-day outage that was caused by the failure of a SAN storage array, and were unable to recover the mail from backup tapes; this issue was a factor in persuading them of the need for dual copies of the mailbox databases and the viability of a CCR and DAS-based solution. The outage showed where IT was not meeting the business needs of their users, and opened their minds to the new opportunity presented with CCR. Two copies of data on less expensive storage can be a more reliable solution than one copy of data on very expensive storage.

By deploying CCR and DAS, operational costs are reduced, single points of failure are eliminated, and Exchange Administrators are able to design the system that best meets the needs of the organization, deploy it, and let it run. The more hands that touch a server or system, the higher the likelihood is that human error will occur. Instead of having to update numerous system components as with a SAN, those who choose to use CCR with DAS as the storage target have fewer components in the path, and thus reduced vulnerability to outages and higher availability at a lower cost.