THE WINDOWS MANAGER'S GUIDE TO **iSCSI SANs**

INSIDE:

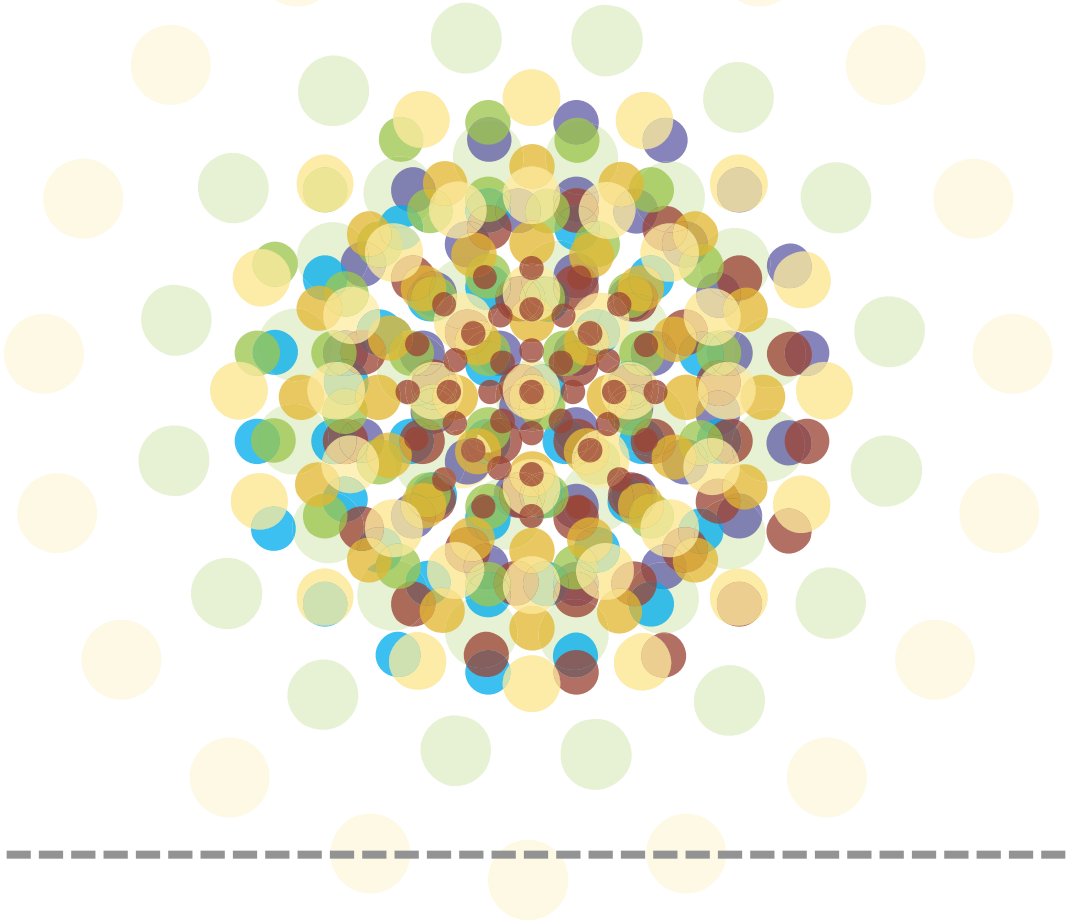

CHAPTER 3

Reducing Exchange Server complexity with SANs

Reducing Exchange Server complexity with SANs

The goal of reducing Exchange Server complexity by leveraging storage platform capabilities makes a lot of sense to many organizations. But the choice between DAS or SAN with Exchange Server 2007 is not merely a question of where to place information stores.

BY MARK ARNOLD

ALTHOUGH MICROSOFT deploys Exchange Server 2007 on direct-attached disks, many prefer iSCSI or Fibre Channel (FC) storage area networks (SANs) as a storage architecture. The main reason is that they offer superior backup and recovery abilities and disk-performance improvements. Of these two methods, iSCSI SANs are preferred over FC SANs because they are more cost efficient and scalable. Because of this, iSCSI SANs are gaining popularity in Exchange Server 2007 infrastructures.

A single Exchange 2007 server can host nearly 4,000 mailboxes and allows each user to retain a gigabyte of email within their mailboxes, making SANs on Exchange 2007 more suitable than direct-attached storage (DAS). Four terabytes of email might be attached to that particular server.

The recommended size of an information store is currently 200 GB. Therefore, five stores give you a full terabyte of capability, even if you don't initially use that much. Provisioning

enough disk capacity from the start requires an enormous expense that can put many organizations off and lead others to place short-sighted size limitations on the user base.

Direct-attached disks were previously considered a viable solution for Exchange Server 2007. However, a large server requires a disk array that has a framework very similar to a SAN.

In a 2007 survey of more than 700 Exchange administrators, SANs were ranked as the main storage solution—nabbing 58%. iSCSI SANs are gaining popularity within the storage market with 11% use, while 9% of respondents chose DAS (Figure 1).

DECISION FACTORS: DISKS AND RAID

When a large organization decides to migrate to Exchange 2007 from Lotus Notes or other legacy platforms, they must choose whether to procure a new SAN, leverage an existing SAN or use DAS. Perhaps the most important factor in deciding whether to choose a DAS versus SAN is the number of input/output operations per second (IOPS) that messaging architects have calculated.

The only way to deliver high message throughput—little or no queued messages—and high resiliency is to use a sufficient number of high-capacity and high-speed disks. At best, a server with direct-attached disks will only offer a couple dozen disks for storage. It won't be long before someone sees the unused space and asks what your data recovery plan is if a server should fail. It's unlikely that plugging a new server into the existing disk array will result in instantaneous recovery.

There also is the issue of disk resiliency. Simply put, disks fail. Disks designed for servers are as likely to fail as those on the SAN, but server disks will likely do more damage.

A DAS may have, for example, 20 disks. About 10 of those disks will be invisible—wrapped up in a RAID 1 or RAID 10 configuration.

For SAN, the story is radically different. RAID 6 is a highly efficient solution for an organization looking

FIGURE 1: *Exchange administrators rank SANs as the main email storage management solution.*

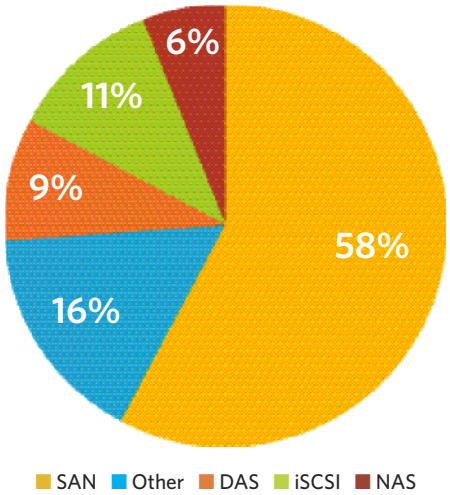

SOURCE: TECHTARGET

to gain the most from an expensive storage system. The useable number of disks from those initial 20 will reduce by three or four. Two disks will be used for RAID resiliency and one or two, possibly more, will be used for hot spares. Instead of 10 useable disks (1.4 TB with 144 GB disks), there is 2.3 TB available space. This means that there is no loss in disk resiliency or performance.

DISK CAPACITY

Although there is a practical limitation to the number of disks you can put in both a SAN and a server, the number of disks a SAN can support far exceeds what a server supports. Procuring larger disks can be tempting. On a SAN, the number of disks is also restricted, meaning that more small disks can be used.

Large disks cause disk latency as the heads travel over the network, whereas smaller disks drop smaller packets of information and move on to the next job.

A poorly designed SAN can suffer just as much from a small disk/large disk problem; however, you're more likely to see a poorly designed DAS solution than a poorly designed SAN. A cleverly designed iSCSI SAN might also implement slower, higher-capacity SATA disks. An accelerator card can increase the available memory working set and reduce the number of disk reads.

EXCHANGE BACKUPS AND HIGH AVAILABILITY

After deciding that you need a SAN for your Exchange Server 2007 deployment, you must decide how to provide users with server resiliency. Previously, clustering would be used. In such a case, two or more nodes are linked into a cluster and share access to a common data set. Only one node possesses a given piece of data at any

VSS backups and Exchange 2007

THE VOLUME SHADOW COPY SERVICE (VSS) must validate the integrity of a database backup after it has been taken. This enables you to rapidly restore a database, as needed. If the verification hadn't been run, the restore wouldn't take more time. You wouldn't know, however, if the VSS backup was consistent. Database mounting would also take longer because the database must be checked before it can mount. For more information on this, click here.

one time; when one node fails, another takes over.

This idea currently exists in the form of local continuous replication (LCR), cluster continuous replication (CCR) and standby continuous replication (SCR), each of which has benefits and drawbacks. Among the drawbacks is the complete duplication of required storage. For each gigabyte of mail in a live store, there is another gigabyte in a passive store waiting to come online if the primary fails. How does a SAN address this challenge?

SAN offers a mature solution—snapshots and rollback points. If you want to restore an Exchange database using DAS, you need to either have a backup file or tapes. Conventional streaming backups take about an hour per store to secure the data, and about three hours to recover data.

This is different on a SAN. Every 30 minutes to an hour, the SAN can process a snapshot of the data on disk. A "photograph" of the disk map is taken, and future writes to the information store go to a different location on the disk.

If a store rollback is required, an Exchange administrator can do this successfully since the data exists on disk. Only a handful of log files must be processed against the store. An administrator knows that he can get the database back online and process email within minutes of initiating a restore.

EXCHANGE 2007 HIGH AVAILABILITY

Exchange Server 2007 has four built-in high availability or storage replication features: SCC, SCR, LCR and CCR. Many organizations, however, have SANs across multiple locations and replicate data between them.

Single copy cluster (SCC)

- Organizations run more servers and are consolidating and virtualizing infrastructures to run on a common platform. Therefore, they're more likely to have several identical hardware platforms in their data centers.
- For email, maintaining an additional server that is completely dedicated to a specific Exchange server doesn't make the most of expensive IT assets. It is more sensible to nominate a server that can assume the role of six or more different servers if one fails.
- A single copy of an Exchange database doesn't protect that database; it only protects from a failed server node. This is hardly an all-encompassing resiliency solution.

Local continuous replication (LCR)

- LCR protects against the physical corruption of an Exchange database by hosting another copy of the store on another set of disks. It doubles

storage requirements, but doesn't enhance server resiliency (Figure 2).

Standby continuous replication (SCR)

■ SCR is the same as LCR in terms of storage growth; however, storage resides on a different server. Failovers are handled manually, but the target server can receive database replication from any number of sources—up to the licensed capacity of the target.

Cluster continuous replication (CCR)

■ CCR presents an interesting hybrid solution—automatic failover as well as cross-site and dual-server resiliency. The remaining issue is storage duplication.

■ CCR benefits from the ability to back up the passive node, rather than affecting the active node. This enables CCR to run a wider availability window for users working extended hours.

■ Microsoft recommends deploying CCR in the same data center so that it can be regarded as a high-availability solution, not a disaster recovery solution. Microsoft recommends SCR for disaster recovery.

Each Exchange resiliency solution has advantages and shortcomings. When using an iSCSI SAN, LCR and CCR might not be considered as high-availability or disaster recovery solutions. This is because a SAN's snap-

FIGURE 2: *Local continuous replication can protect an Exchange database.*

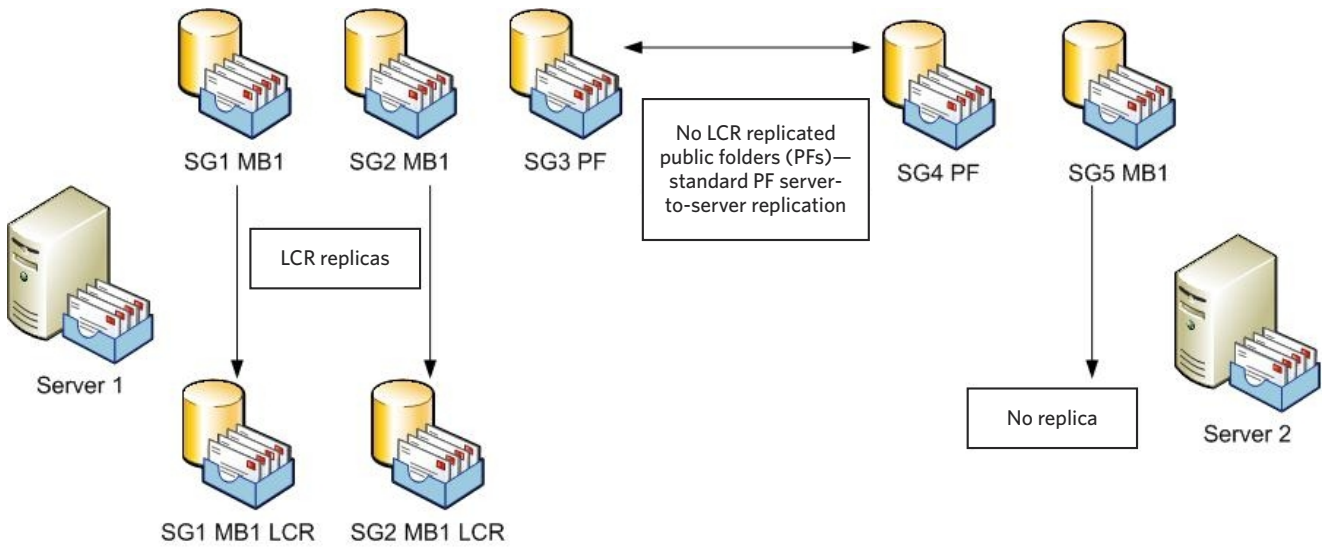

shot capabilities can bring the information store back online in the same or less time than Exchange can when using these replication technologies.

CCR is an excellent solution to use on the SAN because, although the SAN can replicate LUNs that host information stores, it can't assist an Exchange server in achieving high availability. Letting the Exchange server handle data replication is the best solution, but this changes in disaster recovery situations. In the event of a failure, and remembering that SCR activation is a manual process, users can be moved to the new server via database portability.

PHYSICAL DEPLOYMENT

Going into Disk Manager will present several disks representing LUNs that a storage administrator has configured. While there are too many iSCSI SAN architectures to cover individually, some vendors might configure a pair of disks into a RAID 1 set. From there, they create several LUNs on that disk pair to obtain best use of the available physical space.

Other vendors take a large number of disks and make a single data set available. Then, they create LUNs that are presented to the server. How each vendor makes disks resilient varies. The storage administrators in your company will give you the necessary guidance on which disks to use for stores and logs. For large deployments, you will probably see a dedicated LUN per store. Additionally, several transaction logs might be hosted on a single LUN. In all cases, storage administrators should have ensured that databases and logs are on different spindles, SAN "heads" or controllers to provide optimum performance and resiliency.

The connection from the server to the SAN is a critical element in

Database portability

DATABASE PORTABILITY is the ability to mount an Exchange information store database on a completely different server. However, in some cases—if the primary server has failed, for example—it may be necessary to mount the store on a spare server and give users access to their email. For more information on database portability, **click here**. For specific instructions on how to point all users in one store or server to a new server using the MoveMailbox command with the ConfigurationOnly parameter **click here**.

Exchange deployment. An iSCSI SAN can and will have multiple paths between the server and SAN. Microsoft iSCSI Initiator software supports Microsoft Multipath I/O (MPIO), which establishes two connections from the server to the storage (Figure 3).

An iSCSI SAN accesses storage in two ways. iSCSI may use either a standard network interface card (NIC) or a card that works in the same way as a Host Bus Adapter (HBA) card. A standard NIC is significantly less expensive than an HBA, but there are times when this card is beneficial.

DISASTER RECOVERY (DR)

Hosting the server's operating system on the SAN allows that system to be replicated to another storage platform. It also enables a block-level backup of the Exchange server and replicates that data to a disaster recovery site. If a disaster occurs and the data center is inaccessible, the spare servers that are maintained in the secondary data center are brought into service by connecting to the LUNs on the recovery SAN. iSCSI SAN management software can be pre-configured to boot a specified server from a replicated LUN. ■

FIGURE 3: *An iSCSI SAN has multiple paths between the server and the SAN.*

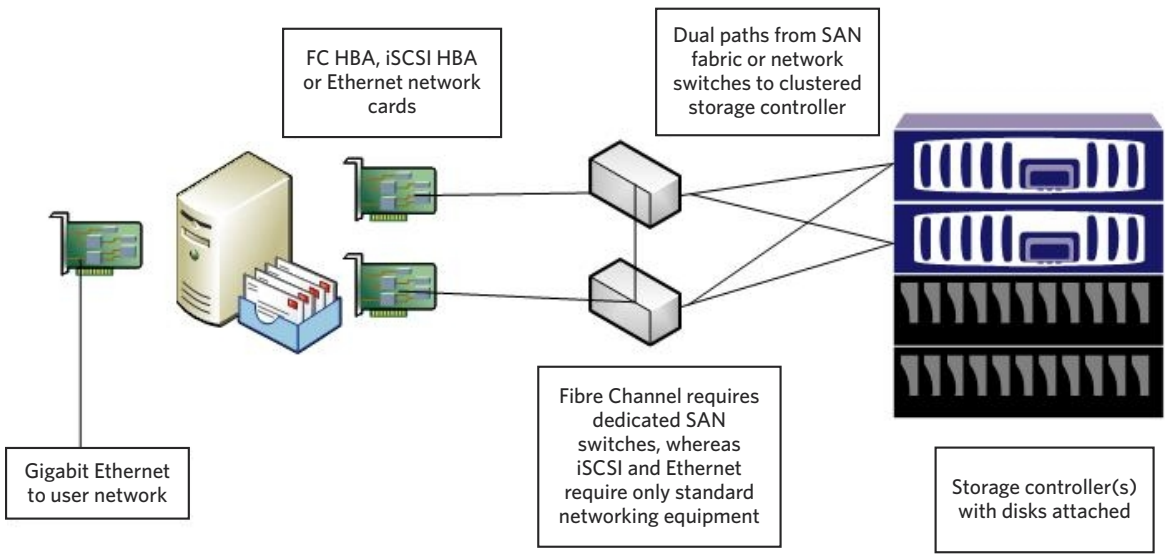

ABOUT THE AUTHOR:

Mark Arnold, *CMCSE+M, Microsoft MVP, has been a consultant in the storage industry for many years. Mark has assisted in designs of SAN-based Exchange servers and other Microsoft infrastructure implementations. He has been a Microsoft MVP in the Exchange discipline since 2001, contributes to various Microsoft-focused technology websites and can be found in the Exchange newsgroups and other Exchange forums.*

SearchExchange.com

The Windows Manager's Guide to iSCSI SANs is part of the Windows Manager's Guide series.

Marilyn Cohodas
Editorial Director
mcohodas@techtarget.com

Michelle Boisvert
Editor
mboisvert@techtarget.com

Martha Moore
Copy Editor
mmoore@techtarget.com

Linda Koury
Art Director of Digital Content
lkoury@techtarget.com

Jonathan Brown
Publisher
jebrown@techtarget.com

Matt Salamone
Senior Director of Sales
msalamone@techtarget.com