Optimizing Storage Performance, Provisioning and Manageability for Microsoft® SQL Server

Network Appliance, Inc.

February 2006, TR-3453

Executive Summary

Microsoft SQL Server has become very popular as a database environment driving a broad range of mid-level business applications, ranging from CRM to eCommerce and from ERP to supply chain integration. The combination of affordability, performance, ease of deployment, and sophisticated analytical features has led to a significant proliferation of SQL Server databases in many organizations. And in many cases, the applications driven by these data bases have become a critical component of business operations.

Unfortunately, in many cases these databases are stored in native internal or direct-attached storage spread over dozens of Windows servers – a storage environment which can not deliver the performance, scalability, availability and manageability required by mission critical data today. In most cases, the problems are compounded by the fact that server administrators, not database administrators (DBAs), are responsible for managing these database environments.

What is required is a consolidated networked storage environment, well-matched to the price/performance characteristics of the database servers, which can meet the demanding service level requirements of the business, and dramatically simplify the management of the distributed SQL Server database environment.

This white paper describes the profound SQL Server benefits available from storage consolidation using best-in-class IP SAN (iSCSI) storage.

Table of Contents

1. Microsoft SQL Server 2005

Nearly all of today's business applications are data-centric, requiring fast and reliable access to intelligent information architectures that can often be provided only by a high-performance relational database system. In its latest release Microsoft SQL Server 2005 offers significant architectural enhancements in performance, scalability, availability, and security. As such, SQL Server environments can be expected to continue to proliferate and become more complex as they are deployed to address increasingly complex data challenges.

Organizations face numerous data challenges today. Employees need to make faster and more data-driven decisions, developers need to be more productive and flexible, and managers need to reduce their overall information technology (IT) budgets even as they scale their data infrastructure to meet ever-increasing demands. SQL Server 2005 is designed to help enterprises address these specific challenges by.delivering increased security, scalability, and availability to enterprise data and analytical applications, while making them easier to build, deploy and manage.

Extending the strengths of SQL Server 2000, SQL Server 2005 provides an integrated data management and analysis solution that enables IT staff to:

- Build, deploy, and manage enterprise applications that are more secure, scalable, and reliable.
- Maximize IT productivity by reducing the complexity of developing and supporting database applications.
- Share data across multiple platforms, applications, and devices to make it easier to connect internal and external systems.
- Control costs without sacrificing performance, availability, scalability, or security.

2. Storage Challenges for SQL Server Environments

As the consumption of data increases, so do the demands placed on IT professionals. Database and storage administrators are confronted with numerous challenges when attempting to "right-size" their database storage for effective space and throughput utilization.

Many organizations add numerous RAID devices to their environments, creating many "islands of storage" with utilization rates estimated by industry analysts at roughly 30% to 40%. Consequently, storage is fragmented and management costs increase significantly with the addition of every new database server. Considerable space and throughput horsepower is also lost, greatly reducing the return of the overall storage investment.

Direct-attached storage in particular rapidly becomes a serious business risk as the rate of data growth and the number of servers increase – driven by lack of capacity and performance scalability, complexity of data protection and cost of operational support.

Traditional SANs have helped to consolidate RAID devices into a central location leveraged by many servers. Still, even in many SAN environments, adding more disks to boost IOPS and more effectively distribute database workload remains inseparable from adding more dedicated space to a given database environment. Sizing for growth and throughput of heavily loaded online transaction processing and decision support systems remains a real challenge for database administrators. Forecasting space requirements for enterprise database applications is often more of an art than a science.

With traditional storage environments, careful planning to correctly size throughput and space utilization is needed prior to the deployment of an application. Frequently, despite meticulous and careful planning, projected sizing proves to be either grossly over or under specified for the real world space, growth, and access patterns of the databases that backend today's critical ERP, CRM, and eCommerce applications. To

compound the problem, making subsequent changes to the database storage layout is cumbersome and comes at the price of reducing system and application availability.

DBAs also encounter weighty challenges when it comes to patches, upgrades, or modifications to their database environments. Despite careful planning and investments in testing environments, sizable risks in upgrading and modifying production environments remain. This challenge becomes greatly magnified when database transactions are tied directly to revenue generation. Downtime to make adjustments to storage layout and run unpredictable upgrades can be very disruptive to business operations.

NetApp enterprise-class storage provides the optimum storage environment for SQL Server by making storage provisioning fast, efficient and flexible; by integrating tightly with the familiar Microsoft Windows and SQL Server admin interfaces; and by simplifying and automating backup, restore and disaster recovery operations for SQL Server. In addition, the ability to deploy NetApp storage in either a Fibre Channel SAN or an affordable IP (iSCSI) SAN environment minimizes risk and provides unmatched flexibility.

3. IP SAN Topology for SQL Server Consolidation

Although SANs can provide significant advantages for SQL Server environments, most are still using direct-attached storage (SQL Server requires a block interface to storage - NAS is not a supported option). There are two primary reasons for this predominance of direct-attached storage:

1. Most SQL Server instances run on small affordable Windows server platforms, for which the cost of traditional SAN infrastructure is usually viewed as prohibitive.
2. SQL Server is most often deployed outside of the large enterprise core data center – for example in distributed environments such as regional, departmental and remote data centers. It is also very popular with smaller and medium size enterprises. In each of these environments the perceived complexity of Fibre Channel fabrics and the lack of appropriate admin skills is a major consideration.

However, with the emergence of iSCSI this situation has changed.

iSCSI is a standard SCSI block storage protocol designed to operate over TCP/IP rather than over Fibre Channel. iSCSI-based IP SANs address both of the issues listed above. They deliver lower total cost of ownership than traditional SANs through lower cost of acquisition and lower complexity; and require no additional expertise – every company has TCP/IP networking expertise and all Windows admins know how to configure Ethernet services.

IP SANs provide all the benefits of a SAN, and more than enough performance to support typical SQL Server environments (see the References section for pointers to performance comparisons).

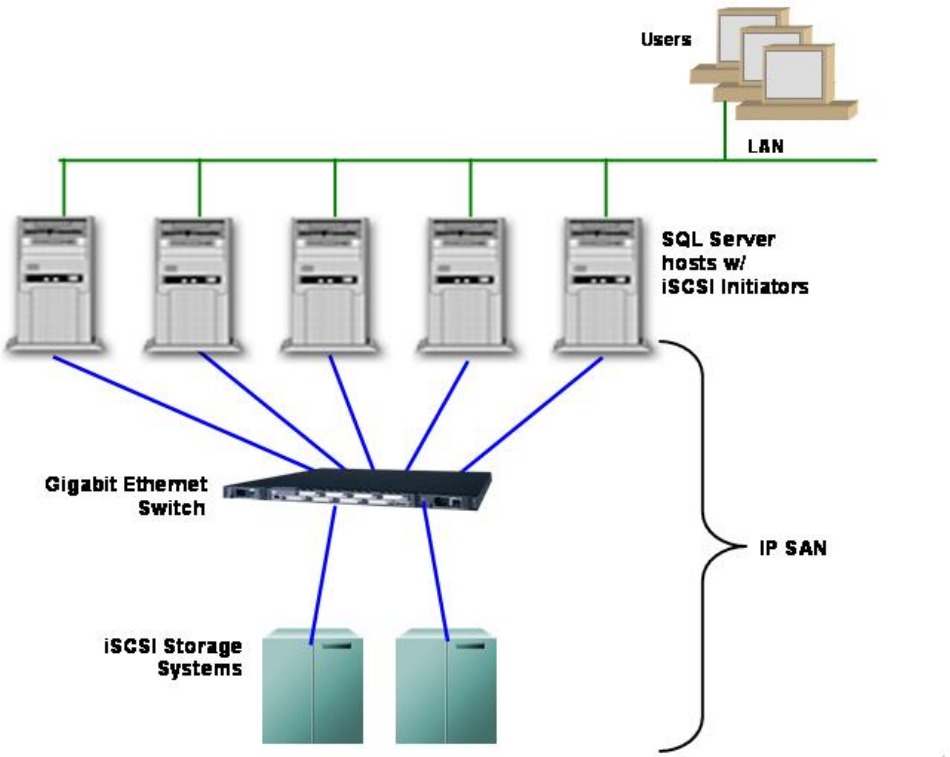

Figure 1) Typical IP SAN

A typical "IP SAN" environment is shown in Figure 1.

Each of the SQL Server host systems are provisioned with an "iSCSI initiator". This is most often simply a software driver that Microsoft supplies as part of the Windows operating which works with the server's built-in Gigabit Ethernet ports; or with an external TOE card (a network interface card with built-in TCP/IP offload chips). Alternatively you can use iSCSI Host Bus Adapters which come with their own drivers (in much the same way you would with Fibre Channel.) Most people find Microsoft's built-in driver perfectly acceptable, except when they want to do boot-from-SAN – which currently requires a TOE card or iSCSI HBA.

The Gigabit Ethernet switching environment can either be dedicated switches in the data center, completely separated from the data communications network, or it can be a subnet of the corporate Gigabit Ethernet switching environment – in which case traffic is usually isolated and protected by VLANs.

The storage systems present virtual disks (LUNs) to the host systems by implementing what is known as an iSCSI target. They are usually just referred to as iSCSI-native storage and Microsoft publishes a list of systems they have tested and qualified in the Microsoft Catalog on their website. Zoning, LUN masking and host authentication is easily accomplished using widely understood Ethernet admin tools and capabilities.

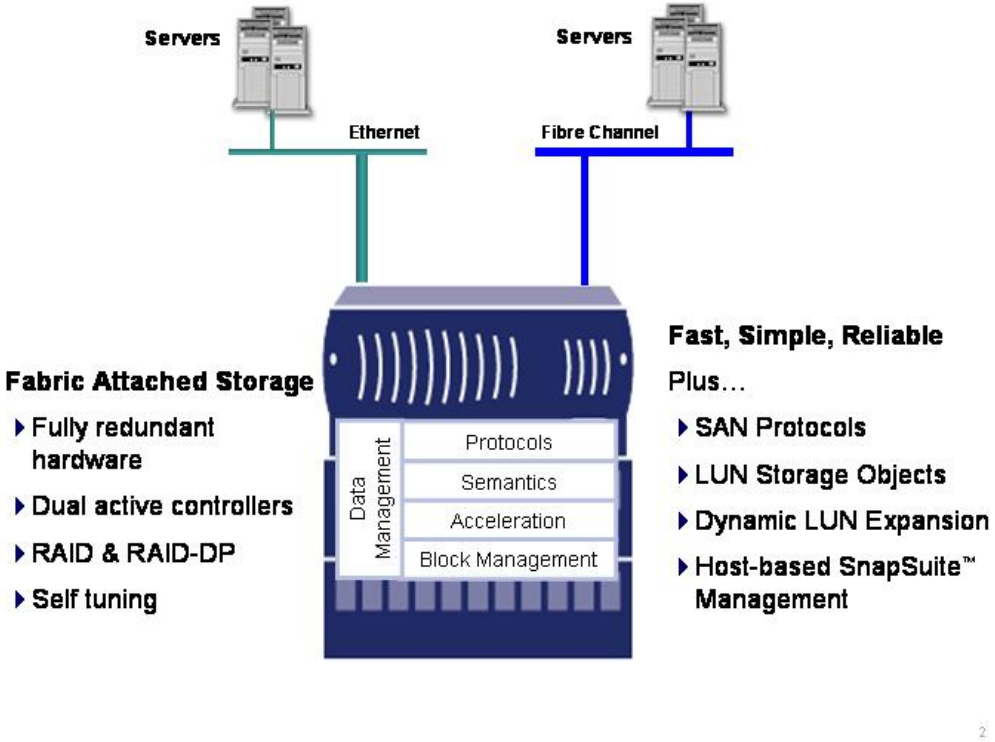

Figure 2) NetApp storage architecture

NetApp has the industry's broadest installed base of iSCSI storage systems, and the industry's broadest product portfolio. All of NetApp's primary storage (FAS Series), secondary storage (Nearstore Series) and gateway (V-Series) systems support the iSCSI protocol.

NetApp's clear differentiation comes from the storage architecture NetApp has been perfecting for 10 years. NetApp primary storage systems deliver all the features you would expect from enterprise-capable arrays: fully redundant hardware, dual active controllers with failover, Fibre Channel and/or SATA disk drives; automatic RAID storage, was well as "phone home" capabilities. However the real magic is in the built-in software which underlies each of these systems.

The heart of our systems is our storage operating system (DataONTAP) which is specifically designed to offload host-side data management complexity. This built-in software provides a sophisticated storage virtualization engine (FlexVol and FlexClone), and state-of-the-art data protection and management capabilities (SnapShot, SnapResore, SnapMirror, SnapVault, SnapLock) from a plug-and-play, self-tuning, highly reliable, modular storage platform.

In 2002, we expanded the architecture to support not only network-attached storage environments, but also SAN environments – with support for FC and IP SAN protocols and LUN storage objects. The result is a more flexible SAN storage solution, complemented by host-side plug-ins (SnapDrive and SnapManager) to take maximum advantage of the NetApp data management capabilities. Our systems can equally efficiently support Fibre Channel SAN, IP SAN, NAS, or all three simultaneously – with management through a single unified management interface.

NetApp's data management software stack enables customers to gain unbeatable automation, productivity and flexibility benefits.

- SnapShots (a patented space-efficient, high performance technology) enable you to completely eliminate backup windows
- SnapRestore significantly boosts admin and end-user productivity through almost instantaneous data restores
- SnapMirror (asynchronous or synchronous) makes Disaster Recovery feasible and affordable for your entire global IT infrastructure
- SnapDrive enables one-click storage provisioning integrated with your host OS environments
- SnapManager for SQL Server integrates SnapShot, SnapRestore and SnapMirror with SQL Server, ensuring data integrity for automated, “minimal-downtime” data protection and disaster recovery

For the remainder of this paper we will focus more specifically how some of these capabilities can optimize storage performance and provisioning, simplify and automate data protection, and significantly reduce the administrative overhead associated with your SQL Server environment.

4. Performance and Provisioning Optimization

Although SANs provide significantly better performance and provisioning flexibility than direct-attached storage, traditional disk arrays often have significant limitations:

- LUNs are typically associated with a single RAID group – therefore performance is limited by the number of disks in the RAID group.
- Admins need to guess maximum size of each LUN and allocate all that space up front. Resizing is both disruptive and complex.

Latest release of DataONTAP (7G) addresses both of these issues, revolutionizing the landscape for customers looking to leverage the maximum spindle count on their SAN environments without forgoing control of space management. With DataONTAP 7G, NetApp made significant advances in array-based virtualization. Before describing these capabilities, it is useful to define some new terms.

An ***aggregate*** is logical entity containing a pool of one or more RAID or RAID-DP groups.

A ***FlexVol*** is a logical storage entity residing within an aggregate. FlexVols can be sized appropriately anytime for the unique requirements of each SQL Server database instance. In addition, they can be grown and shrunk on-the-fly with no impact to system and database availability.

A ***FlexClone*** is generated from a Snapshot copy of a FlexVol, providing a transparent writable copy of its ancestor or parent. The underlying data of the FlexClone, less its changes, requires no immediate additional space because it physically points to the underlying blocks in the ancestor. As the data in the FlexClone begins to diverge from its ancestor, additional space will be occupied to hold the related changes.

A FlexClone can also be “split” and become an entirely new physical copy of its ancestor. One of the most powerful benefits of the FlexClone split is that it can occur while the clone is mounted and being written to by a DBMS, like SQL Server. In the case of SQL Server, an administrator may have several attached LUNs that reside in one or more FlexVols and be actively servicing thousands of transactions while a split procedure is taking place, without losing data. Unlike traditional environments in which databases have to complete a physical copy before they can be mounted again, database administrators can quickly clone a given database environment, mount it for service, and then trigger a split to reallocate the database’s underlying blocks to another volume, all while servicing requests from applications and users. Split processes can also be scheduled and executed during off hours to minimize the resource utilization impacts of copying data to an entirely new volume.

4.1 FlexVol in the SQL Server Environment

SQL Server administrators are confronted with a variety of challenges:

- Space management and predictive growth
- I/O sizing for performance
- Managing database upgrades and patches
- Managing data-centric application upgrades like SAP, Siebel, and Remedy

A FlexVol can be grown or shrunk at will, lending a greater level of granularity in managing the inevitable changes in database space requirements. If volumes are too large for one database, the unused space can quickly be reallocated to another server as needed. By contrast, particularly in direct-attached storage environments, space allocated to a SQL Server remains dedicated regardless of whether space or throughput is utilized. As more databases are added to an environment in this manner, storage asset utilization decreases.

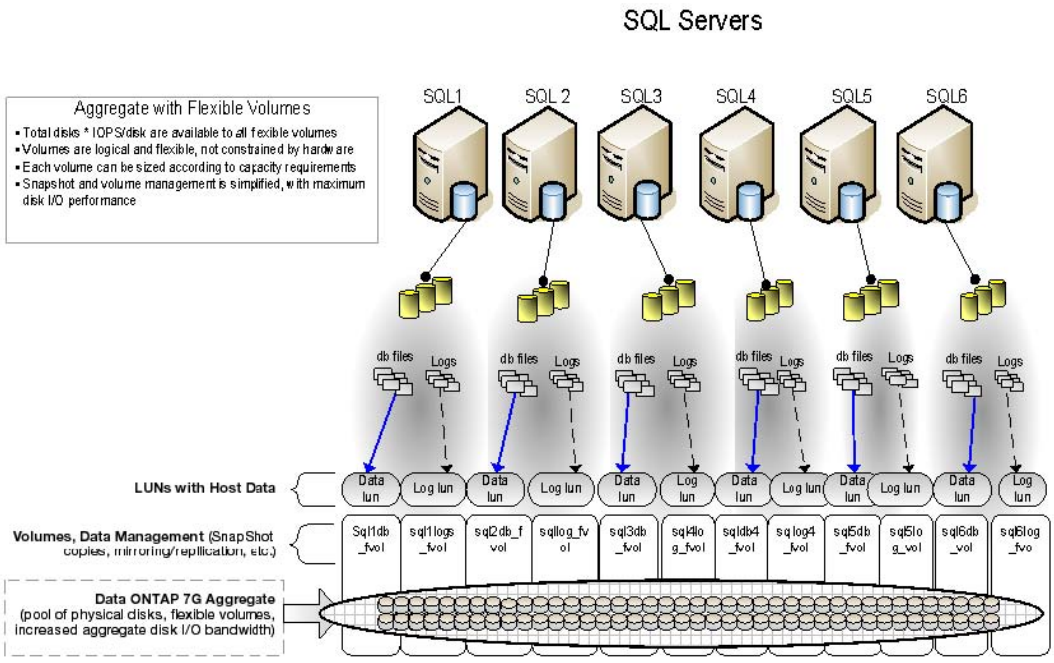

Figure 3) FlexVol for SQL Server.

FlexVols allow administrators to aggregate massive spindle counts across all their SQL Server databases, effectively boosting performance and throughput to unprecedented levels without sacrificing control over space management. Consider Figure 3, in which six SQL Server databases have consolidated all their volumes into 12 FlexVols contained within one large 100+ disk aggregate.

Each database leverages the power and reliability of over 100 disk spindles and multiple RAID groups, transcending many of the inherent performance and reliability challenges of direct-attached storage or even traditional SAN environments. If, for example, the databases on the servers SQL1 through SQL6 required additional space for any of their databases, their respective LUNs can easily be grown on-the-fly (using SnapDrive) with no impact to database availability. The previous challenges of storage inflexibility and low utilization associated with dedicated physical volumes have been virtually eliminated. Hence, FlexVols allow

administrators to quickly and effectively scale their infrastructure up and out, with virtually no impact to system availability. Because space and IOPS reside in a shared environment, storage asset utilization quickly surpasses that of traditional volume environments.

4.2 FlexClone in the SQL Server Environment

FlexClone leverages NetApp's unique Snapshot technology to enable DBAs to create instantaneous database clones, providing a replica of production data in a risk-free designated "sandbox." A FlexClone is created in seconds, whether the clone command is executed on a small 10MB volume or an enormous 16TB volume. Using conventional methods and traditional storage, a DBA would be forced to recover an entire backup from tapes or nearline devices into a test environment. When databases grow to several hundred gigabytes or more, organizations are forced to create an entirely new instance of equal or greater storage capacity in the test environment to effectively test upgrades and modifications before they go to production. The conventional method is time consuming, more costly, and less efficient than using a FlexClone. Figure 4 compares tape, mirroring, and FlexClone technologies, demonstating how FlexClone dramatically shortens test/development cycle time.

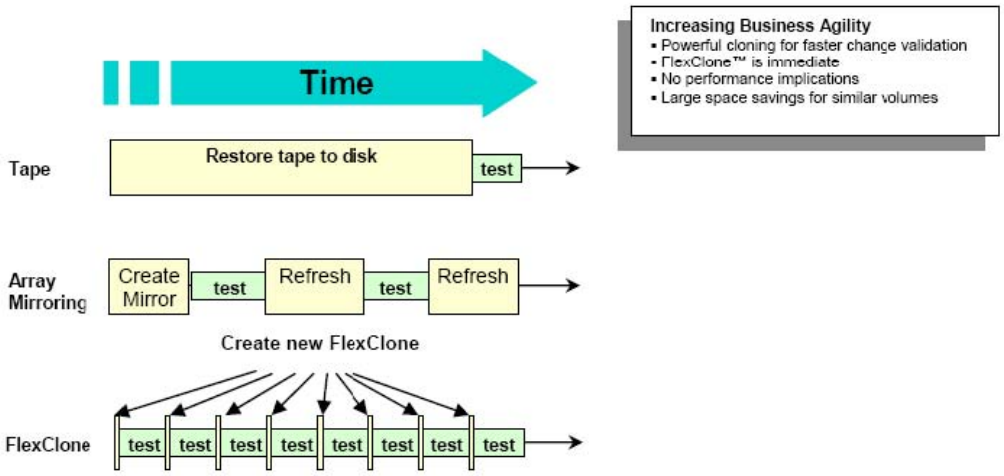

Figure 4) Tape, mirroring, and FlexClone

In the time it takes to create one tape-generated test environment, eight clone environments were generated and tested using FlexClone. While array mirroring provides a slightly better recovery rate than that of tape, it is still significantly less efficient than using FlexClone for testing purposes.

FlexClone increase ROI of the overall storage investment by providing fast and dynamic instances of production data without the awkward, expensive, and time-consuming recovery processes associated with conventional data recovery. Database upgrades, schema modifications, and patch testing can all be quickly tested against a FlexClone, enabling the IT organization to quickly investigate the repercussions of making changes, without adversely affecting the production environment. Unpredictable upgrades, patches, and modifications are quickly and economically validated, thereby minimizing the risks of such changes. FlexClones enable DBAs to play out many "what if" scenarios against production data until they produce expected or acceptable results.

5. Simplifying and Automating Data Management

Having investigated the significant benefits of NetApp's unique FlexVol and FlexClone features for SQL Server environments, we will now look at SnapDrive and SnapManager for SQL Server –host-side software which tightly integrates the Windows and SQL Server administrative interface with NetApp's built-in SnapShot, SnapRestore and SnapMirror technology.

5.1 About NetApp Snapshots

As mentioned earlier NetApp Snapshot technology is a patented space-efficient, high performance method of creating point-in-time replicas of data on disk. They provide a means of creating a backup image of data at a particular point in time. Network Appliance Snapshot backups occur in a matter of seconds and typically each copy consumes only the amount of data that has changed since the last copy was created. Thus NetApp Snapshots consume minimal disk space while providing up to 255 online point-in-time images. A related capability, SnapRestore, can recover a previous point-in-time image almost instantaneously.

These features, together with two host-side software plug-ins can dramatically simplify and automate the backup and recovery of SQL Server databases. Figure 5 illustrates a typical configuration.

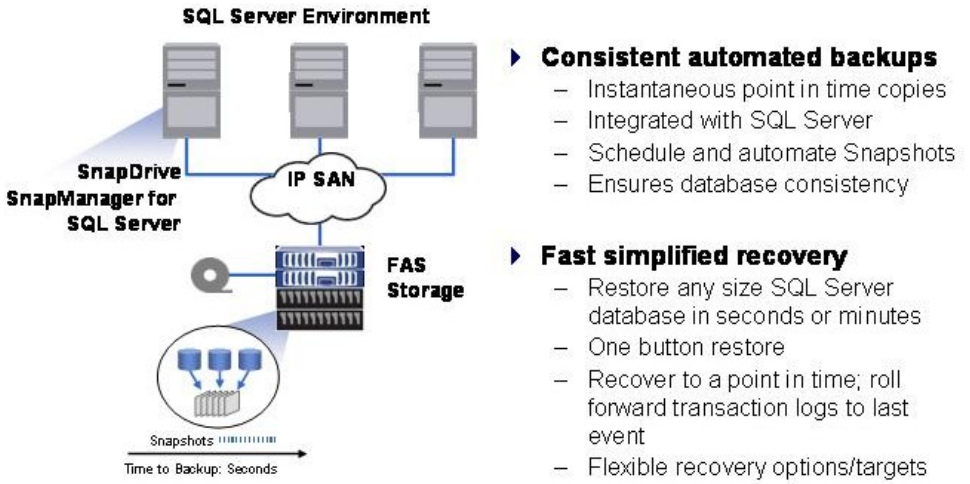

Figure 5) Simplifying and Automating SQL Server Backup and Recovery

5.2 About SnapManager for SQL Server

SnapManager for SQL Server integrates Microsoft's storage administration interface for SQL Server with NetApp's built-in data management features. This makes it possible to accomplish the instantaneous backup and very fast recovery of entire SQL Server databases and full text indexes. SnapManager for SQL Server includes "easy migration" wizards to move databases from direct-attached storage to Fibre Channel SAN and IP SAN environments. Its features include an easy-to-use, intuitive graphical user interface and rich backup scheduling and reporting capabilities.

In accomplishing a database backup, SnapManager for SQL Server ensures that:

- Snapshot copies will be created of all volumes used by the databases being backed up.
- System databases will be backed up using conventional streaming-based backup.

- User databases placed in a LUN that also contains system databases will be backed up using streaming-based backup.
- Backup of transaction logs will always use streaming-based backup to provide point-in-time restores.

In restoring a database, SnapManager for SQL Server follows the following steps:

- Trigger SnapDrive to request a Single_File_SnapRestore to restore all LUNs used by the database(s).
- Restore selected databases out of many sharing one or two LUNs. (The most efficient configuration is one user database per LUN.)
- Request SQL Server to restore the databases.
- Request SQL Server to apply selected backed-up transaction logs. This is configurable and can conduct either point-in-time or up-to-the-minute restores.

In conjunction with SnapMirror® for wide area data replication, SnapManager for SQL Server also makes it possible to create affordable automated disaster recovery for SQL Server environments.

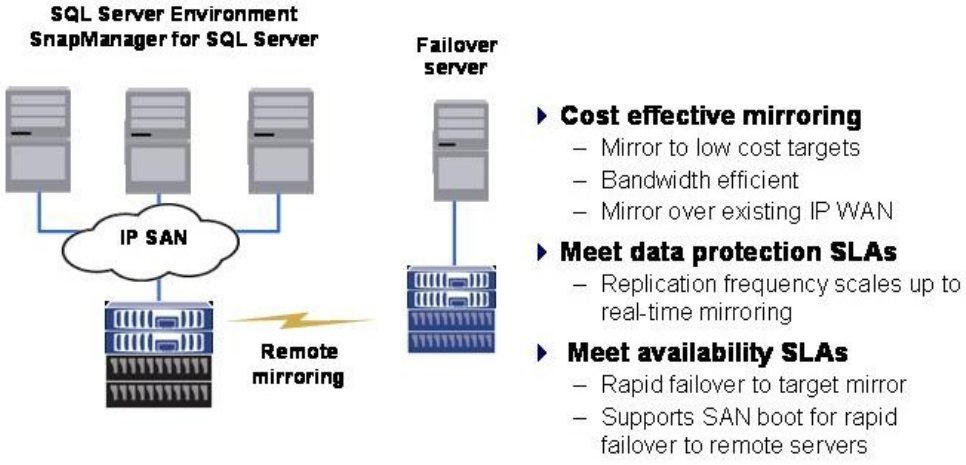

Figure 6) Typical Disaster recovery configuration

Figure 6 shows a typical scenario.

6. Summary

NetApp IP SAN storage solutions for SQL Server provide an enterprise-capable consolidated SAN environment, well-matched to the price/performance characteristics of the database servers, which can meet the demanding service level requirements of the business, and dramatically simplify the management of the distributed SQL Server database environment. With unique features like Snapshot, FlexVol, FlexClone, SnapDrive and SnapManager for SQL Server, NetApp can optimize storage performance and provisioning, simplify and automate data protection, and significantly reduce the administrative overhead associated with your SQL Server environment.

7. References and Acknowledgements

This white paper borrowed heavily from two technical reports:

TR3410: Data ONTAP™ 7G: FlexVol™ and FlexClone™ for Microsoft® SQL Server by Jonas Irwin, Network Appliance
TR 3431: Best Practices Guide: Microsoft® SQL Server 2000 / 2005 and NetApp Solution by Jonas Irwin, Network Appliance

See also:

TR3373: Data ONTAP™ 7G—The Ideal Platform for Database Applications.
ESG Validation Study – NetApp iSCSI SAN Solution

www.netapp.com