Fibre Channel and iSCSI Performance Comparison for DSS Workloads Using SQL Server 2005

Yogesh Manocha, and Ian Breitner

Network Appliance, Inc., May 2006, TR-3476

Abstract

This report details a series of tests performed by Microsoft and Network Appliance designed to compare the performance of a large-scale, Microsoft® SQL Server 2005 Decision Support System (DSS) workload over Fibre Channel and iSCSI protocols.

Table of Contents

Acknowledgements

Network Appliance would like to extend grateful thanks to Microsoft for their participation in the tests conducted for this technical report. Particular thanks to the SQL Server Performance Team in Redmond, WA, for their expertise and intensive support during the testing.

1. Introduction

For many years, large enterprises have relied on Fibre Channel (FCP) as the default storage area network (SAN) protocol for high-performance application and database workloads. As enterprises look to drive down costs and reduce the total cost of ownership (TCO) of IT environments, many have sought alternative, lower-cost mechanisms to provide similar scalability to that of conventional FCP deployments. iSCSI leverages the ubiquitous features of Ethernet topologies and TCP/IP by encapsulating the SCSI command set over IP. While iSCSI technology may seem like a new concept, it really represents the marriage of two rock-solid, proven communication protocols: IP and SCSI. As network speeds increase and the industry moves toward wide-scale adoption of 10 GbE networks, combining these two mature technologies represents a natural progression in the evolution of storage networks.

Microsoft and other major software vendors have declared support for iSCSI and even provide free-of-charge iSCSI software initiators that leverage commodity network interface cards and existing Ethernet infrastructure to provide blocks-based storage connectivity. The cost of commodity Ethernet hardware is significantly less than that of FCP components and switches. Manageability is also greatly simplified, because IP is easily one of the most well-understood protocols in the IT industry. However, questions remain about the scalability, performance, and reliability of iSCSI with real-world applications and database workloads.

1.1 Purpose and Scope

This paper provides customers with quantifiable metrics of iSCSI performance relative to FCP in a large-scale Microsoft SQL Server 2005 Decision Support System (DSS) and demonstrates how iSCSI is appropriate for high-performance applications and database workloads. The systems tested in this paper ran identical SQL Server DSS workloads, first using FCP and then using iSCSI, after which the results were compared. At the conclusion of this report, readers will have a good sense of the performance and scalability of iSCSI in a full-scale database deployment. The scalability, features, and robust characteristics of Microsoft SQL Server 2005 and NetApp storage systems are also covered in this report.

2. NetApp Simplifies SQL Server Data

Highly scalable Network Appliance™ storage subsystems provide simple-to-manage, cost-effective, and reliable storage solutions for mission-critical high-availability database infrastructures. In addition, NetApp storage solutions improve the usage of storage resources and offer excellent data protection and rapid application testing.

NetApp storage solutions provide advanced backup and recovery capabilities critical for keeping SQL highly available. SnapManager® for SQL Server dramatically reduces SQL Server backup times—from hours to seconds—and makes every backup a full backup. Backups are based on NetApp Snapshot™ copies, enabling simultaneous backups of multiple databases of any size while requiring minimal disk space for each additional full backup. SnapManager for SQL Server enables organizations to recover from SQL Server outages in minutes, not hours or days, making it one of the fastest backup and recovery solutions available and delivering an industry-leading combination of availability, scalability, and reliability for SQL Server environments.

2.1 High Availability

NetApp storage systems are uniquely scalable and highly available, with the ability to scale from 4TB up to 504TB. NetApp fabric-attached storage (FAS) systems offer a rich set of features with inherent redundancy of robust hardware components. Clustered FAS systems are available for even higher levels of availability. Innovations like RAID-DP™ (Double Parity RAID) enhance availability by protecting against dual-disk drive failures that may occur during RAID reconstructs. The SnapDrive® family of products offers Multipath Input/Output (MPIO) solutions for both iSCSI and FCP protocols, providing redundant paths from hosts to FAS systems.

2.2 Improved Performance and Asset Utilization

NetApp FlexVol™ technology, introduced in Data ONTAP® 7G, makes it possible to optimize storage utilization by logically partitioning NetApp storage arrays into smaller virtualized volumes (FlexVol volumes). Many SQL Server environments support small databases with very high I/O requirements, often requiring many more disk drives for throughput than are required for space. FlexVol technology enables the creation of a base aggregate container with many disk drives and logically partitions the aggregate into smaller virtualized volume structures. By combining the same number of disks that were previously sliced up across many conventional volumes, much higher throughput and storage asset utilization can be achieved while still maintaining the benefits of discrete dedicated volumes, because all of the volumes share the benefit of having access to all of the drives for improved performance. FlexVol volumes are quickly created, expanded, or reduced in mere seconds, regardless of size.

3. Microsoft SQL Server 2005

Microsoft SQL Server is used to support businesses in a number of mission-critical database processing functions, including online transaction processing (OLTP) and Decision Support System workloads. This paper focuses on comparing the performance of iSCSI relative to FCP in DSS environments. DSS workloads are highly resource intensive, and can potentially saturate both system CPUs and disk bandwidth.

The database engine in Microsoft SQL Server 2005 has been extensively reworked to improve both absolute performance and multiprocessor scaling in DSS environments, without the need for an administrator to manually tune the system. The following sections summarize these improvements.

3.1 Memory Management

Memory management is a critical part of the server operation. Enhancements in this area improve overall server performance, particularly in DSS environments. Several major improvements in SQL Server memory management contribute to improved memory support for complex, long-running, resource-intensive queries critical in DSS workloads.

- Uniform Memory Management. The uniform memory management framework in the SQL Server OS layer provides common memory brokerage between different components of the SQL Server, improving performance and providing flexible operations under a variety of memory pressures. One particular aspect of this framework, uniform dynamic caching, improves the behavior and locality of the internal caches, resulting in performance improvements in a wide range of operating conditions.
- NUMA Aware. SQL Server 2005 is non-uniform memory access (NUMA) architecture aware, which enables it to take advantage of memory locality, thereby supporting greater scalability.
- Dynamic Memory Management. SQL Server 2005 supports dynamic management of conventional, locked, and large-page memory. The new version of SQL Server also supports Hot Add Memory, reducing the need for reboots.
- Memory tracking. Major enhancements in memory tracking, including tracking of internal memory allocations between components, as well as external operating system-wide memory events, provide superior supportability features in memory management for SQL Server 2005.

3.2 Query Processing

SQL Server 2005 provides major advances in query processing capability for decision support applications. The query processor of SQL Server encompasses the execution environment, query optimizer, and query executor.

- Execution Environment. Controlling the query compilation and plan caching (one of the major enhancements of the execution environment) is a new query plan stability feature. This feature allows query hints to be attached to queries at runtime, even when the application cannot be modified. This is useful for tuning decision support applications purchased from Independent Software Vendors (ISVs).
- Query Optimizer. Responsible for finding the best plan for executing a query, the optimizer has been significantly improved for SQL Server 2005 for DSS environments. Improvements include:
  - Full optimization capability for partitioned tables, including partition elimination capability that can avoid reading partitions that can be shown not to contain rows that contribute to the query result.
  - A new date correlation optimization feature that allows many-fold speedups for join queries across tables with correlated date-time columns when there is a range filter on one of the date-time columns.
  - The statistics subsystem of the optimizer provides improved statistics loading capability, loading the right statistics at the right time, every time, for improved cost estimation and better query plans. A new string summary statistics feature allows precise estimation of the selectivity of LIKE conditions that involve patterns with wildcards in any location, such as the string '%john%smith%'. Traditional histogram statistics also have improved accuracy.
  - The optimizer provides another plan stability feature called plan forcing that enables the user to force the optimizer to choose a particular plan for a query. This can help solve the problem where a handful of queries in a large workload slow down after upgrade to a new service pack or release.
- Query Executor. Responsible for running the plan generated by the optimizer to solve a query, the SQL Server 2005 executor has been enhanced in a number of ways, including:
  - Full support for partitioned tables
  - New query execution operators to support T-SQL language extensions
  - Improved hash join performance
  - Improved nested loop performance
  - Improved bitmap join optimizations

4. Test Environment

The following sections detail the server and storage configurations used in the FCP and iSCSI performance testing comparisons.

4.1 Servers

The 4-way server used for testing is configured with Microsoft Windows® Server 2003 Enterprise Edition for 64-bit Itanium-based systems. For full details, see Table 1.

COMPONENT	DETAILS	
Operating System	Microsoft Windows Server 2003 Enterprise Edition for 64-bit Itanium-based Systems	
System Type	Itanium-based	
Database Servers	Microsoft SQL Server 2005 Enterprise Edition	
Page File Space	2.00 GB	
Processor	4 * 1500 MHz	
Total Physical Memory	16 GB	
Total Virtual Memory	17.57 GB	
Protocol	iSCSI (MPIO)	FCP (MPIO)
Storage Interconnect	Microsoft iSCSI Software Initiator 2.01 and Intel MT/1000 NIC Cards	Emulex Lightpulse 9002 HBAs (Queue depth set to the maximum of 254)
Storage Network	1Gb Ethernet	2Gb FC-AL

Table 1) Server Configuration

4.2 Storage

All potential bottlenecks outside the server space were identified and eliminated prior to testing. To ensure the validity of the test results, the storage architecture is designed in such a way as to not introduce any potential bottlenecks with respect to I/O demands and scaling capabilities.

COMPONENT	DETAILS
Operating System	Data ONTAP 7.1
Processor	Two Network Appliance FAS3050 storage systems
Disks	Eighty-four 72 GB 15K RPM Drives

LUNs	Eight 150 GB	
Files	8	
Protocol	iSCSI	FCP
Storage Interconnect	Emulex LP100i-D1 iSCSI target cards	Onboard QLogic 2362 FCP target cards
Storage Network	1 Gb Ethernet	2 Gb FC-AL

Table 2) FAS3050 Storage Systems Configuration

Test Setup

Two NetApp FAS3050 systems are connected directly to the server, eliminating the need for FC or Ethernet switching during the tests (to simplify the test environment, switches are not used in this deployment). Both storage systems are running Data ONTAP 7.1 and configured with forty-two 72GB 15K RPM disks. Of these, 36 disks are configured to host four LUNs for data, with the remaining 6 disks configured as dual-parity disks for RAID protection (see Figure 3 and the appendix for details). Using FlexVol technology, data is then automatically striped across all of the disks without the need for manual operator intervention.

NetApp Data ONTAP LUNs are designed to be protocol independent and quickly interchangeable between FCP and iSCSI. This flexibility greatly helped accelerate the protocol performance comparison testing. For a schematic of the test deployments, see Figures 1 and 2.

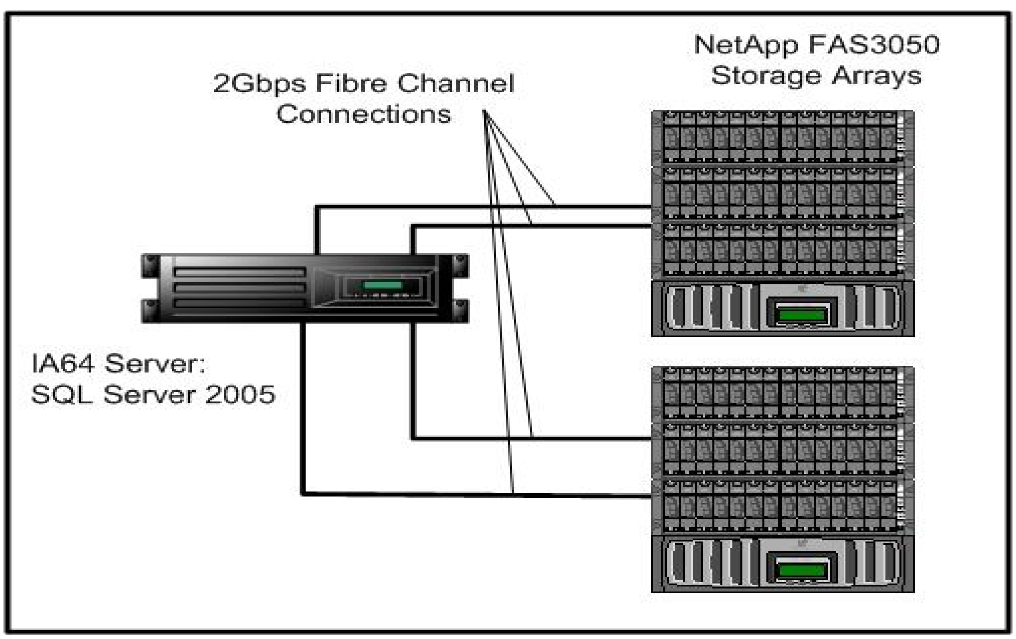

Figure 1) Schematic of a 4-way server and NetApp direct-attached storage for the FCP test environment

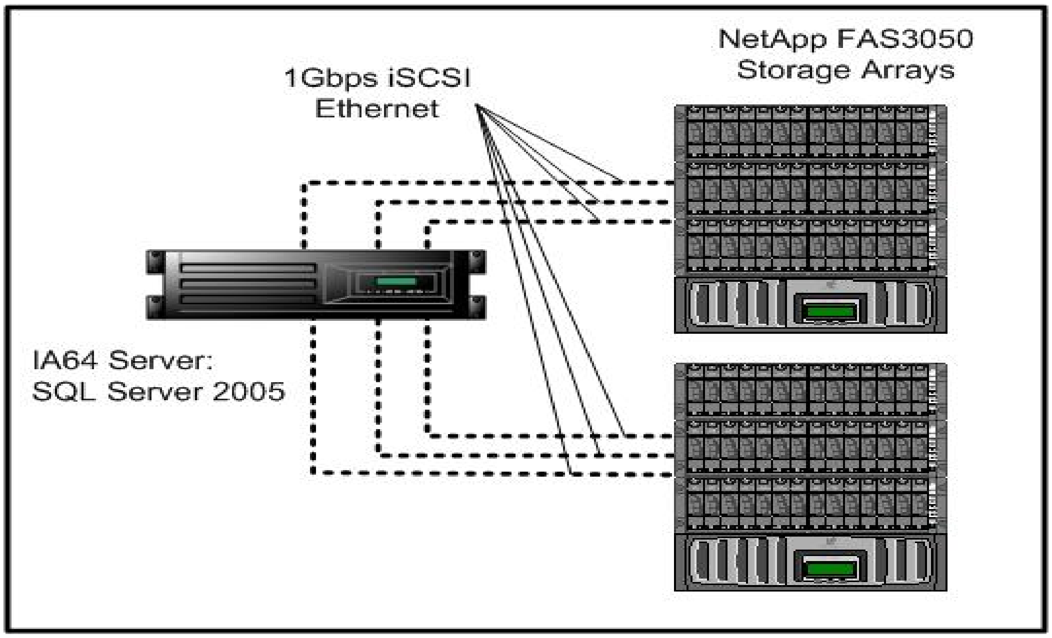

Figure 2) Schematic of a 4-way server and NetApp direct-attached storage for the iSCSI test environment

Four 2-Gbps FCP and six 1-Gbps iSCSI connections are used for FCP and iSCSI configurations respectively. Preliminary tests quickly determine that the available bandwidth in each configuration will not introduce a bottleneck in any of the tests.

NetApp FAS3050 Storage System

The FAS3050 storage appliance has 4 GB of RAM and 512 MB of NVRAM. The new Direct Memory Access (DMA) capability of the NVRAM-5 card enables it to read directly from main memory without using CPU cycles. Also, the substantial capacity and battery-backed resiliency of the NVRAM 5 card means that large amounts of write data can be immediately and safely acknowledged by the storage subsystem without engaging physical disks in the process. Writes are then flushed to disk in optimized batches as needed, thus improving efficiency and performance. Table 3 shows FAS3050 system specifications.

COMONENT	**DETAILS**
Processor	2 X 2.8GHz Intel Xeon
Memory	4GB
NVRAM	512MB NVRAM - 5
PCI-X 100 Slots	4

Max Disks	336

Table 3) NetApp FAS3050 System Specifications

4.3 Database Layout

SQL Server 2005 Test Runs on the 4-Way System

Each of the two FAS3050 systems is configured with four LUNs of 150 GB each, which are located in a single FlexVol volume on each storage system. Using multiple data files, the SQL Server DSS database is distributed across all LUNs to obtain optimal performance for the 4-way server system. Because the tempdb database provides a “scratch-pad” area for queries and sorting activity, it is critical that tempdb reads and writes are not inhibited by I/O limitations. As such, the tempdb database is also split into 16 data files distributed uniformly over each of the 8 data LUNs. The size of each tempdb data file is also expanded prior to the test to prevent costly file auto-growth activities from occurring during the DSS workload. For more information about optimizing tempdb, please refer to the appendix.

Figure 3 shows the physical and logical storage layout used on one (netapp1) of the two FAS3050 storage systems. The same configuration is repeated on the second storage system (netapp2).

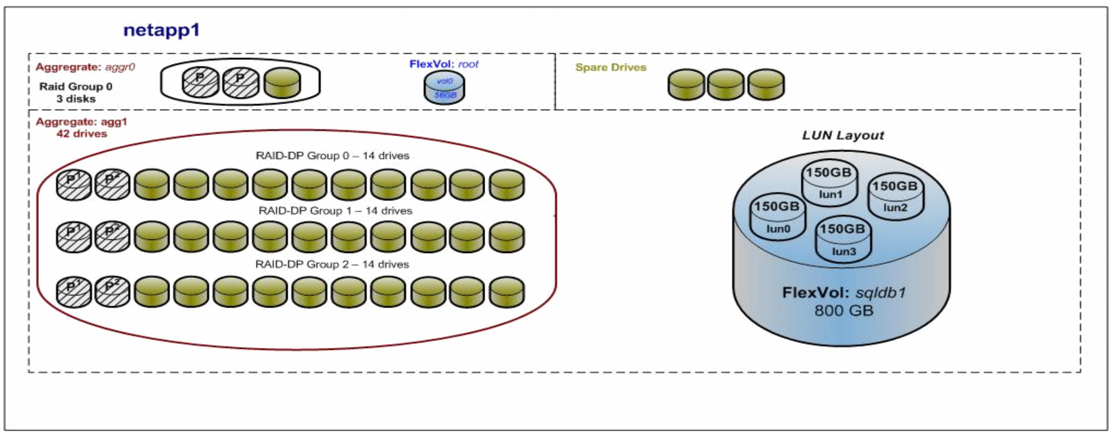

Figure 3) Disk layout for both FAS3050 storage systems

For the FCP and iSCSI performance test runs on the 4-way system, a total of eight 150-GB LUNs are provisioned, each containing one DSS SQL Server database file, as shown in Table 4.

STORAGE APPLIANCE	LUN LOCATION	LUN SIZE	DATABASE FILE
Netapp1	/vol/sqldb1/lun0	150 GB	(Database file 1)
	/vol/sqldb1/lun1	150 GB	(Database file 3)

	/vol/sqldb1/lun2	150 GB	(Database file 5)
	/vol/sqldb1/lun3	150 GB	(Database file 7)
Netapp2	/vol/sqldb2/lun0	150 GB	(Database file 2)
	/vol/sqldb2/lun1	150 GB	(Database file 4)
	/vol/sqldb2/lun2	150 GB	(Database file 6)
	/vol/sqldb2/lun3	150 GB	(Database file 8)

Table 4) Database Layout

4.4 Performance

The NetApp FAS3050 storage systems provided the high-performance and low-latency access to storage required for this series of tests. Using the single 4-way Windows 2003 Server, this configuration ultimately achieves throughput totaling 620 MB per second on FCP and 464 MB per second on iSCSI during large table scans. To ensure the performance differences measured are a result of moving from FCP to iSCSI connections, default storage array settings are used for all tests (with a few minor protocol-specific exceptions discussed later). Because SQL Server 2005 already generates a sufficiently deep pre-fetch, the NetApp storage system pre-fetch is disabled during sequential read by setting the volume option minra to on, as shown below.

The volume option no_atime_update is also turned on to prevent the update of the access time on an inode during file reads.

The following Data ONTAP commands are used to implement these changes to the volume options:

```
Netapp1> vol options sqldb1 minra on
Netapp1> vol options sqldb1 no_atime_update on
Netapp2> vol options sqldb2 minra on
Netapp2> vol options sqldb2 no_atime_update on
```

4.5 SQL Server 2005 Tuning Options

SQL Server 2005 performs most of the necessary tuning automatically and dynamically configures its parameters based on usage of system resources. There are, however, options that can be configured to maximize performance in specific environments. For both the iSCSI and FCP test runs, the configurations shown in Table 5 are used.

PARAMETER NAME	VALUE	DESCRIPTION OF PARAMETER

Affinity Mask	15	Affinity Mask specifies a mask that SQL Server uses to associate its threads to the processors on the system.
Light Weight Pooling	1	Lightweight pooling reduces excessive context switching and reduces processor overhead.
Max Worker Threads	1024	The Max Worker Threads value specifies the number of worker threads servicing requests from clients.
Max Server Memory	15000 MB	Max Server Memory is the maximum amount of memory assigned to SQL Server. The value was set to approximately 95% of Total RAM (16 GB).
Min Server Memory	14000 MB	Min Server Memory was set to approximately 90% of Total RAM (16 GB) to improve performance of queries that use hashing or sorting operations.
Recovery Interval	32767 minutes	Recovery Interval was set to a high value to avoid excessive checkpoints.

Table 5) SQL Server Parameter Tuning

4.6 FCP Tuning

The Emulex Fibre Channel HBA default value of 32 for queue depth is too low for heavy database workloads. The QueueDepth value is increased to 256 for the Emulex HBA driver. The Emulex parameter QueueDepth value is set by using the LputilNT.exe program supplied with the driver.

The following Windows registry value is also increased appropriately:

HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\lpxnds\Parameters\

Device\NumberOfRequests

NumberOfRequests is a generic Windows parameter that specifies the maximum number of outstanding requests allowed from any initiator. To maximize performance, the NumberOfRequests registry parameter is increased to 256.

4.7 iSCSI Tuning

Server

The Microsoft iSCSI Software Initiator version 2.01 is used during the iSCSI test runs. MPIO support included in the Microsoft iSCSI Software Initiator enables aggregation of multiple iSCSI sessions from the server to both storage systems. The round-robin policy is selected to make all paths active and send I/O in a round-robin fashion for load balancing. In the Fibre Channel test configuration, multiple paths are used by statically assigning each LUN to a path.

The Microsoft iSCSI Software Initiator registry entry MaxPendingRequest, which controls the maximum number of outstanding requests, is increased to 256. The registry path for MaxPendingRequest may be different from server to server. As an example, here is the registry key path for the MaxPendingRequest parameter on the server used for these tests:

\MyComputer\HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\class\

{4D36E97B-E325-11CE-BFC1-08002BE10318}\0008\Parameters

The best way to find this key is to execute a find operation in regedit from the root (My Computer) of the tree.

4.8 Storage Systems

The iSCSI option max_ios_per_session, which specifies the maximum outstanding commands per session, is changed from the default value of 128 to 256 for database workloads.

The following Data ONTAP command is used to change the default value:

```
options iscsi.max_ios_per_session 256
```

Emulex LP100i-D1 iSCSI target cards are used in the FAS3050 storage systems for iSCSI connections to offload some of the iSCSI protocol processing, which helped to improve storage systems CPU resource utilization.

5. Database Design and Purpose

The type of DSS database tested is representative of databases found in many customer environments and is designed for tracking sales, customer, supply-chain, and product lifecycle trends. The central charter of a DSS database is to help organizations increase profitability by analyzing trends and correlations over long periods of time. The test database was fully normalized and fully indexed on primary and foreign keys. The size of the database, including tables and indexes, is 560 GB on disk.

6. DSS Testing Queries

A number of typical decision support queries were run to stress the system and to evaluate the performance of FCP and iSCSI protocols. These queries, all of which differ from one another, are chosen for their complexity and diversity in terms of data access patterns and query parameters, and because they access a large proportion of the available data. All of these queries benefit from effective memory management. All queries were run as separate job requests. The set of queries included the following:

Profit Enhancement (PE). This query provides the total increase in profits, if certain discounts had not been offered on products sold during a specified time period. The PE query assists in determining future product discounts.

Stress characteristics:

- I/O intensive
- Intensive scan indexing required

Cost-Efficient Suppliers (CE). This query lists the supplier's details for purchasing a specified part, from lowest to highest cost. The CE query identifies the supplier that offers a specific part at the lowest cost, and can help a business identify the appropriate supplier in a given geographic region.

Stress characteristics:

- Complex query
- Query plan and optimization critical

Specified Product Profit (SP). This query provides annual profit summaries for a specified product, itemized by geographic region. The SP query helps a company to determine its most profitable product by geographic region.

Stress characteristics:

- Complex query

- Query plan and optimization critical

Pricing Summary (PS). This query provides a total count and total price of all products sold during a specified time period. The time period was selected such that approximately 95% of the table was scanned. The PS query measures the total amount of business sales during the specified time frame.

Stress characteristics:

- Floating point calculations
- Expression evaluation performance critical

Unshipped Order Summary (UO). This query was used to extract a summary of a specified number of highest revenue generating orders not shipped by specified date. For example, if N=20, the query must return the highest 20 returns (unless fewer than 20 rows qualify for the SQL statement, in which case all rows must be returned). This query assists in identifying the shipment priority to meet company revenue goals.

Stress characteristics:

- Random reads
- Nested loop and hash joins critical

Late Received Orders (LR). This query provides a summary breakdown of products that were shipped to the customer prior to the promised delivery date, but that were not received on time. The LR query, which summarizes late orders by shipping method, helps in selecting the best shipping method to ensure future on-time deliveries.

Stress characteristics:

- Random reads
- Nested loop and hash joins critical

7. Test Results: SQL Server 2005 iSCSI/FCP Performance Comparisons

For each DSS query type, tests are first run using FCP and then iSCSI with the configurations described in the preceding sections. The following sections show a three-part analysis consisting of query response times, average read MB/sec, and server CPU categorized by the DSS query types shown in Table 6.

QUERY NAME	**WORKLOAD CHARACTERISTIC**
Pricing Summary (PS)	Intensive floating point calculations
Profit Enhancements (PE) and Unshipped Orders (UO)	Intensive joins
Cost-Efficient Suppliers(CE) and Specified Product Profit (SP)	Complex queries
Late Received Orders (LR)	High random reads

Table 6) Recap of Query Types

7.1 Comparative Execution Times

Figure 4 shows a comparative measure of execution times aggregated together and puts iSCSI within 12-13% of the execution times of FCP. The total execution time for all DSS queries using FCP is 2131.1 seconds; iSCSI performs at 2389.9 seconds.

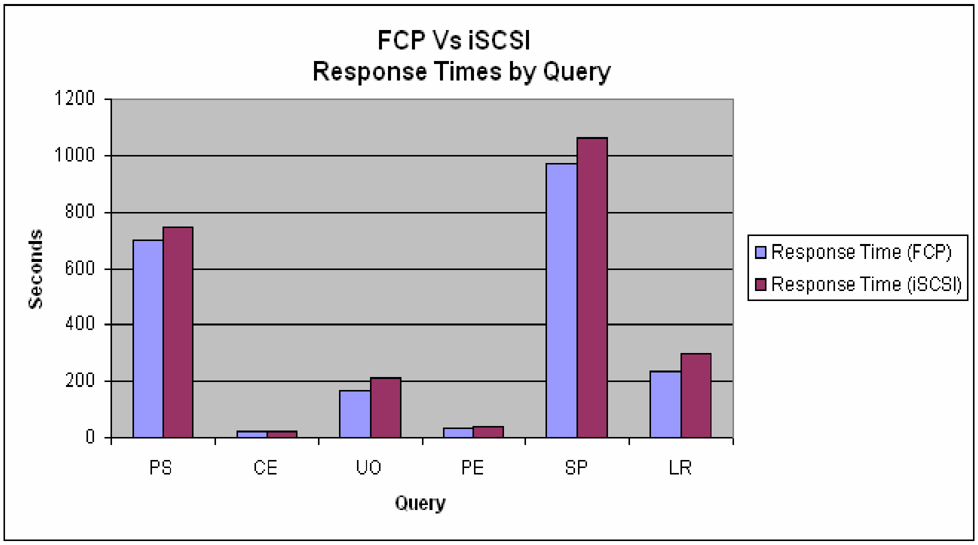

Figure 4) Comparison of execution times for queries

7.2 Disk Read Performance

Figure 5 shows a comparison of total disk read throughput (MB/second) between iSCSI and FCP connected storage. For most query types, FCP tests show a slightly higher disk read throughput. SQL Server 2005 generates, on average, 23% higher disk read throughput (measured in MB/second) using FCP compared to the same work load on iSCSI. The disk read performance difference between FCP and iSCSI is due to several factors, including the offload of FCP-specific processing to HBAs.

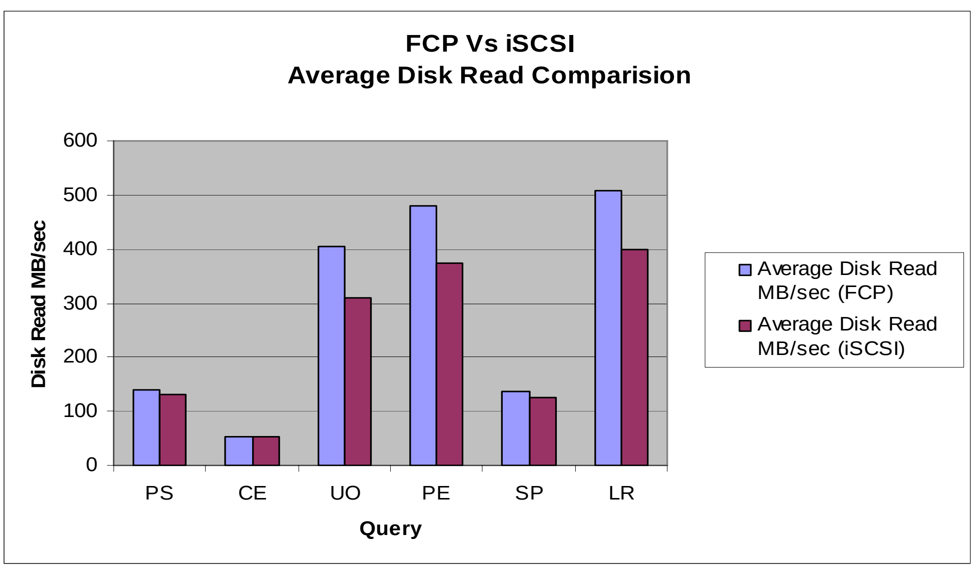

Figure 5) Comparison of disk read times

7.3 Processor Utilization

Figure 6 shows the processor utilization on the 4-way server with FCP and iSCSI protocols. Overall, system resource utilization for iSCSI using the software initiator adds a small overhead of approximately 11%, when compared to the same workload on FCP. Tests show the average host processor utilization for iSCSI workloads using a software initiator is either similar to or slightly higher than for FCP workloads, as shown in Figure 6. It should be noted that most organizations do not run systems at full CPU saturation.

Typically, for commercial database deployments, servers are configured such that only 60-70% of system resources are used. Therefore an increase in CPU usage does not necessarily result in proportional server performance degradation.

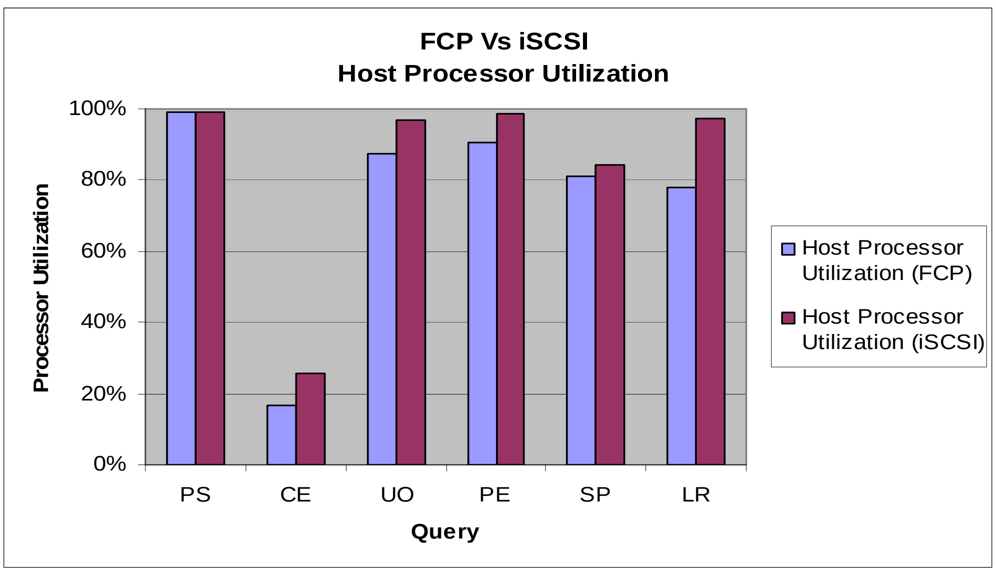

Figure 6) Comparison of processor utilization

8. Summary

Organizations looking to reduce the costs and complexity associated with conventional SAN storage should use the data in this document to gauge how iSCSI compares to Fibre Channel when running large-scale DSS workloads. This paper highlights several DSS workload types (sequential and random) and demonstrates that iSCSI is a highly scalable and appropriate storage protocol for DSS workloads, delivering high throughput and fast response times similar to FCP.

The key take-away point that emerges from the data presented in this report is that iSCSI is just as viable of a storage interconnect as FCP for large-scale DSS workloads. In both the iSCSI and FCP test runs, performance was inhibited only by available server CPU resources while sufficient unused bandwidth remained in both the storage interconnects and the NetApp storage systems. Perhaps most notable is that on average, both iSCSI and FCP performed at approximately 78% of their total available bandwidth as configured in the test scenarios. Also of note is the 11% increase in iSCSI CPU utilization and 12% latency disparity between iSCSI and FCP, which can be attributed to the the use of an iSCSI software initiator. Running the same performance tests using an offload engine such as an iSCSI HBA would likely close this gap. In most environments, however, the Microsoft iSCSI Software Initiator is more than adequate to drive high-performance DSS workloads at a fraction of the cost of FCP.

Regardless of the storage interconnect an organization chooses to deploy, NetApp storage systems are uniquely positioned to bring unparalleled flexibility and scalability to enterprise-class SQL Server environments. NetApp storage systems provide a single integrated data management and protection environment, supporting customers' high-performance storage needs across a wide array of platforms, applications, and geographies.

9. Appendices

9.1 Storage System Layout

Figures 7 and 8 show the physical and logical layouts of the storage systems for hosting the SQL Server database files. All storage resources were provided by NetApp, as detailed in Table 3.

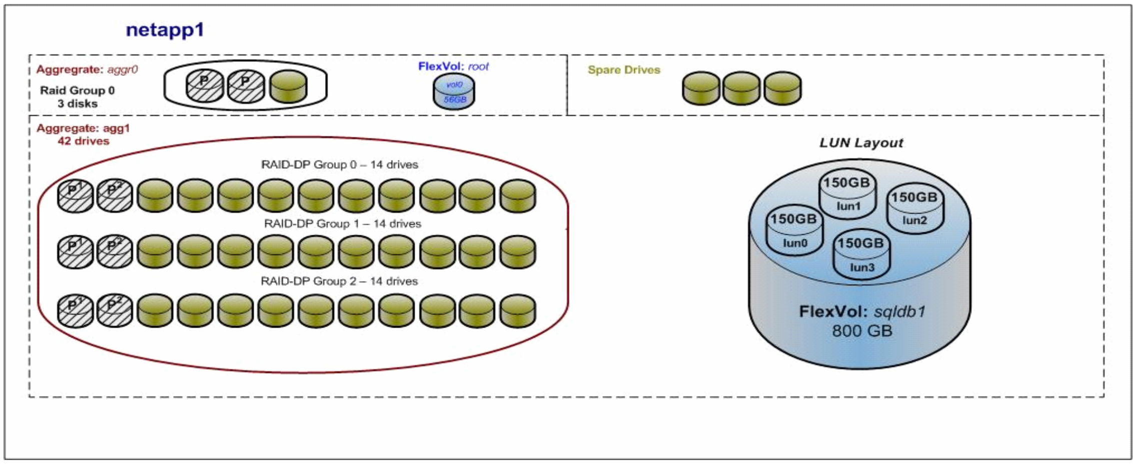

Figure 7) FAS3050-1 disk layout

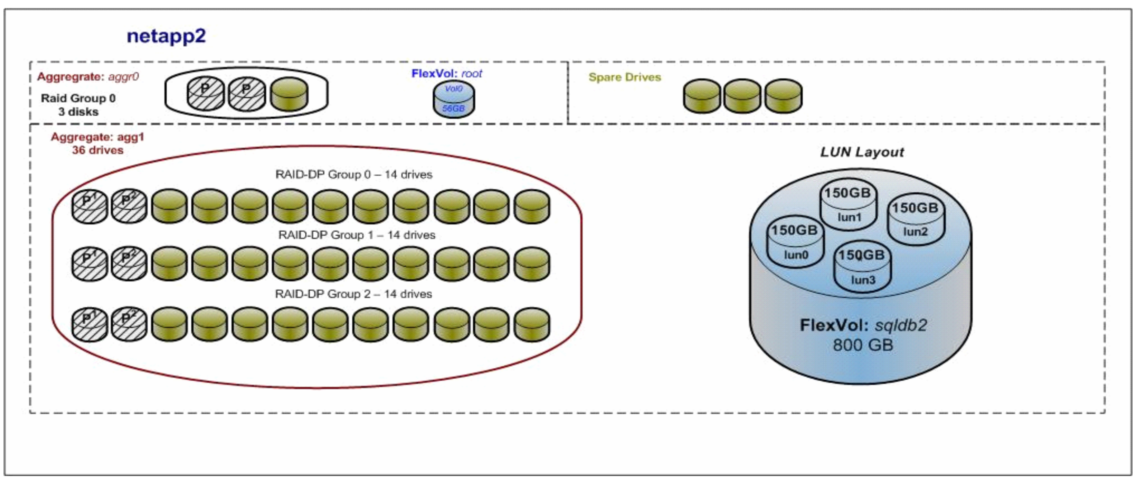

Figure 8) FAS3050-2 disk layout

10. Related Links

See the following resources for further information:

- 10 Reasons to Choose NetApp Storage Consolidation for SQL Server: http://www.netapp.com/ftp/sc-sql-10reasons.pdf
- NetApp SnapManager for SQL Server: http://www.netapp.com/products/software/snapmanager-sql.html
- Microsoft SQL Server 2005 Partner Solution Case Study: http://www.netapp.com/ftp/netapp-msft-tpcc-study.pdf
- Microsoft SQL Server 2005 Decision Support Scalability Improvements: http://www.netapp.com/partners/docs/msft-sqlserver2005.pdf
- TPC-C Benchmark Using IBM xSeries servers, Microsoft SQL Server 2005 and NetApp FAS3000 Series Storage: http://www.netapp.com/partners/microsoft/tpc-c-results.html
- For the latest information about Microsoft SQL Server 2005, see the SQL Server site: http://www.microsoft.com/sql/default.mspx
- Optimizing tempdb Performance: http://msdn2.microsoft.com/en-us/library/ms175527(SQL.90).aspx
- NetApp Technical Report about Best Practices for Microsoft SQL Server: http://www.netapp.com/library/tr/3431.pdf
- NetApp Technical Report about FlexVol Volumes and FlexClone™ Volumes: http://www.netapp.com/library/tr/3410.pdf
- NetApp Technical Report about Data ONTAP 7G: http://www.netapp.com/library/tr/3373.pdf

[1] A LUN (logical unit) is a portion of available storage that is configured and presented to the Windows operating system as a physical disk drive.