

Microsoft® SQL Server™ 2005

SQL Server 2005: Decision Support Scalability Improvements

Microsoft Corporation

Published: May 2005

Abstract

The new release of the Microsoft® SQL Server™ contains numerous enhancements designed to improve both overall performance and the ability to scale effectively as more storage systems are added to the configuration. This paper details two series of performance tests, run jointly by Microsoft and Network Appliance™, which were designed to determine the effectiveness of the decision support system (DSS) enhancements in SQL Server 2005. Testing reveals significant gains in performance of SQL Server 2005 when compared with SQL Server 2000, as well as dramatic increases in scalability when more CPU resources are added to the configuration.

This is a preliminary document and may be changed substantially prior to final commercial release of the software described herein.

The information contained in this document represents the current view of Microsoft Corporation on the issues discussed as of the date of publication. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information presented after the date of publication.

This White Paper is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS, IMPLIED OR STATUTORY, AS TO THE INFORMATION IN THIS DOCUMENT.

Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Microsoft Corporation.

Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

© 2005 Microsoft Corporation. All rights reserved.

Microsoft, SQL Server and Windows are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

The names of actual companies and products mentioned herein may be the trademarks of their respective owners.

Contents

Acknowledgements	1
Introduction	2
SQL Server 2005 Database Engine Enhancements	2
Memory Management	2
Query Processing	3
Test Environment	5
Servers	5
Storage	6
Test Setup	6
Performance	7
Network Appliance Value Add	8
Database Design	8
DSS Testing Queries	9
Test Results	11
SQL Server 2000 and SQL Server 2005 Performance Comparisons	11
Comparative Execution Times	11
Disk Read Performance	12
Processor Utilization	12
SQL Server 2005 Scaling Testing	13
Comparative Execution Times	13
Disk Read Performance	14
Processor Utilization	15
Summary	16
Appendix	17
Storage System Layout	17
Related Links	18

Acknowledgements

The SQL Server team would like to extend grateful thanks to Network Appliance™ for their loan of storage hardware for the tests conducted in this white paper. Particular thanks to the Database Performance Engineering Team in Raleigh, NC, for their expertise and intensive support during testing, as well as for their input in the write-up of this paper.



Introduction

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 (DSS) workloads. The speed with which query results are returned is directly impacted by several key factors related to the:

- **Database**—including the nature of the queries (DSS or OLTP, for example) run against the database, the size and complexity of the database, and the type of processing being performed (such as joins or aggregations);
- **Hardware**—commonly the speed and number of CPUs as well as the amount of memory; more rarely, the number of disks;
- **Software**—the efficiency of the application algorithms used to execute the queries.

In this paper, we focus on performance testing in DSS environments. DSS workloads, unlike OLTP, are highly resource intensive (see Table 1). Characterized by long transactions with complex queries (either ad hoc or programmed, as with online analytical processing), DSS queries touch large amounts of data (often terabytes in size). As a consequence, DSS queries can potentially saturate both system CPUs and disk bandwidth.

Table 1. Characteristics of DSS (in comparison with OLTP)

Decision Support Systems (including OLAP)	OLTP (Online transaction processing)
<ul style="list-style-type: none">• Long transactions• Complex queries• Large amounts of data accessed• Combines data from different sources• Resource intensive• Indexing enables higher performance	<ul style="list-style-type: none">• Short transactions• Limited number of standardized queries• Small amounts of data accessed• Uses data from only one source• Not resource intensive• Maintenance costs incurred with indexing

SQL Server 2005 Database Engine Enhancements

The database engine in the new Microsoft SQL Server has been extensively re-worked to improve both absolute performance and multiprocessor scaling in DSS environments, without the need for the administrator to manually tune the system. The details of these improvements are summarized below.

Memory Management

Memory management is a critical part of server operations. Improvements in this area enhance overall server performance, particularly in DSS environments. There are several major improvements in SQL Server memory management that contribute to improved memory support for the complex, long-running, and resource-intensive queries critical to DSS workloads.

- **Uniform Memory Management.** The uniform memory management framework in the SQL Server operating system layer provides common memory brokerage between different components of SQL Server, improving performance and providing flexible operations under a variety of memory pressures. One particular aspect of this framework, uniform dynamic caching, improves behavior/locality of the internal caches, thereby enhancing performance in a wide range of operating conditions.
- **NUMA Aware.** SQL Server 2005 is non-uniform memory access (NUMA) architecture aware. This improvement enables the new SQL Server to take advantage of memory locality, thereby conferring 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, such as external operating system-wide memory events and tracking of internal memory allocations between components, provide SQL Server 2005 with superior supportability features in memory management.

Query Processing

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

- **Execution Environment.** The execution environment controls query compilation and plan caching. One of the major enhancements to the execution environment is a new 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 SQL Server 2005 query optimizer has been significantly improved for DSS environments. Specific improvements include:
 - Full optimization capability for partitioned tables, including partition elimination capabilities. This capability precludes the need to read partitions which contain rows unrelated to the query results.
 - A new date-time correlation optimization feature. When there is a range filter on one of the data-time columns, this feature facilitates accelerated join queries across tables with correlated date-time columns.
 - Enhancements to the statistics subsystem. More effective statistic loading capabilities now load the appropriate statistic at the right time, every time, thereby improving cost estimation and query plans. In addition, a new string summary statistics feature allows precise estimation of the selectivity of LIKE conditions involving patterns with wildcards in any location (as with, for example, the string '%john%smith%'). Traditional histogram statistics have also improved accuracy.

- Plan forcing. Another plan stability feature, plan forcing enables the user to direct the optimizer to choose a specific query plan. This is particularly useful in the event of a new release or service pack upgrade, after which a few queries in a large workload may slow performance.
- **Query Executor.** Responsible for running the query plan generated by the optimizer, the SQL Server 2005 executor has been enhanced in a number of ways. These include:
 - 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.

The SQL Server 2005 query processing capabilities for decision support applications represent a tremendous advance over the already strong decision support features of SQL Server 2000.

Test Environment

The following sections detail the server and storage configurations used in the SQL Server performance testing comparisons. For a schematic of the test deployments, see Figure 1.

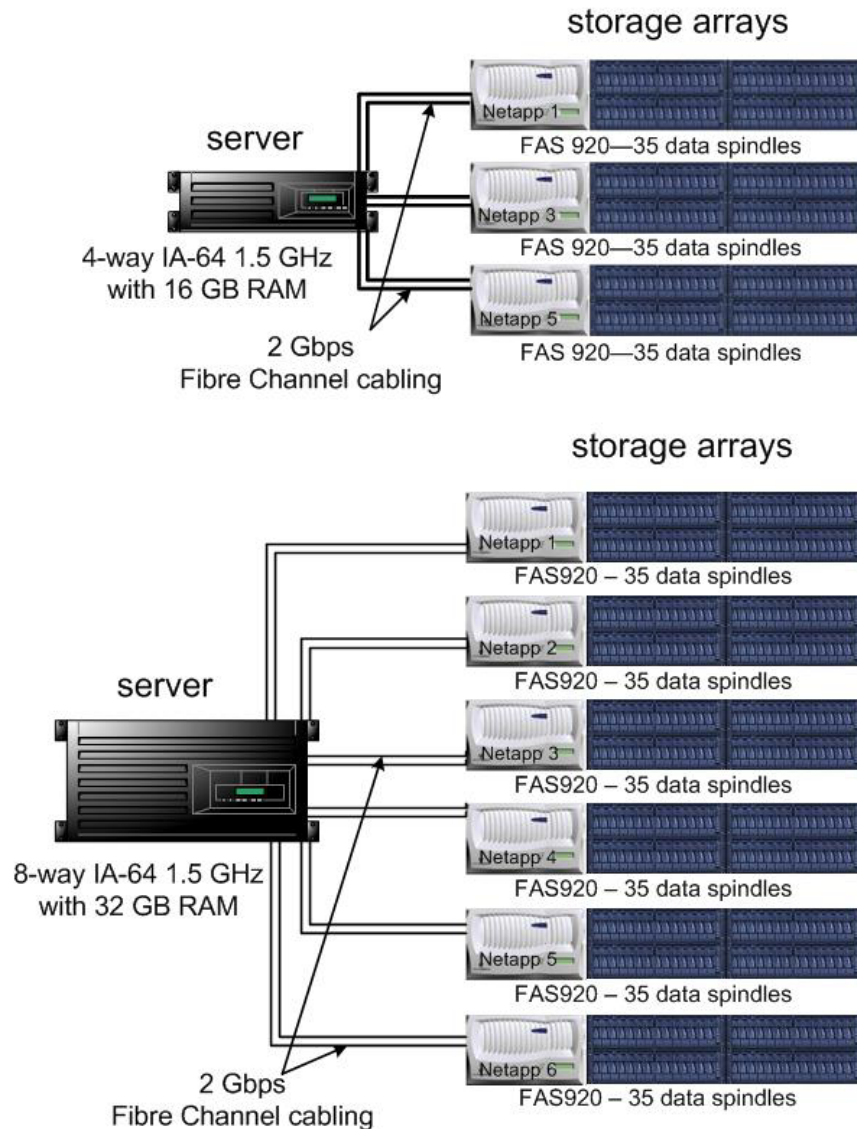


Figure 1. Schematic of 4- and 8-way servers and externally-attached Network Appliance storage.

Servers

All test systems ran Microsoft® Windows Server™ 2003 Datacenter Edition, which provides 64-bit capabilities. Two sets of database performance comparisons were made: First, between SQL Server 2000 and SQL Server 2005 Beta 2, and second, between four-processor and eight-processor server configurations both running SQL Server 2005. For full details, see Table 2.

Table 2. Server Configuration

Component	Details	
Operating System	Microsoft Windows Server 2003, Datacenter Edition for 64-bit Itanium-based Systems	
Version	5.2.3790 Service Pack 1, v.1433, build 3790	
System Type	Itanium™ based	
Database Servers	<ul style="list-style-type: none"> • Microsoft SQL Server 2000 Enterprise Edition, SP4 build • Microsoft SQL Server 2005 post Beta 2, build 1099 	
Page File Space	2.00 GB	
	4-way	8-way
Processor	4 * 1500 MHz	8 * 1500 MHz
Total Physical Memory	16 GB	32 GB
Total Virtual Memory	17.57 GB	33.57 GB

Storage

It was critical both that SQL Server performance was truly stressed, and that potential bottlenecks outside the server space were identified and eliminated prior to testing. While the servers had finite CPU and memory resources (essentially, they were fixed variables), it was important to use a storage subsystem architecture that would not introduce any I/O bottlenecks or scaling constraints.

Test Setup

Network Appliance FAS920 Fabric-Attached Storage systems provided externally-attached storage for the test setup. (To simplify the test environment, switches were not used in this deployment.) The NetApp storage solution is very scalable, both at the four and eight-processor server configurations. As storage demands increased, the modular design and integrated management GUI made adding storage on the fly fast and simple.

All storage systems ran Data ONTAP™ version 6.5.2. Each storage system was configured with 56 x 72GB 15K RPM disks; of these, 35 spindles were configured to host four LUNs¹ for data, and three spindles were configured as parity disks for RAID protection (see Figure 2 for an example of the configuration). More spindles were configured on the storage subsystems than would normally be required to support the workloads generated. Note that the storage system automatically distributes the data across all of the disks in the volume without the need for operator intervention.

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

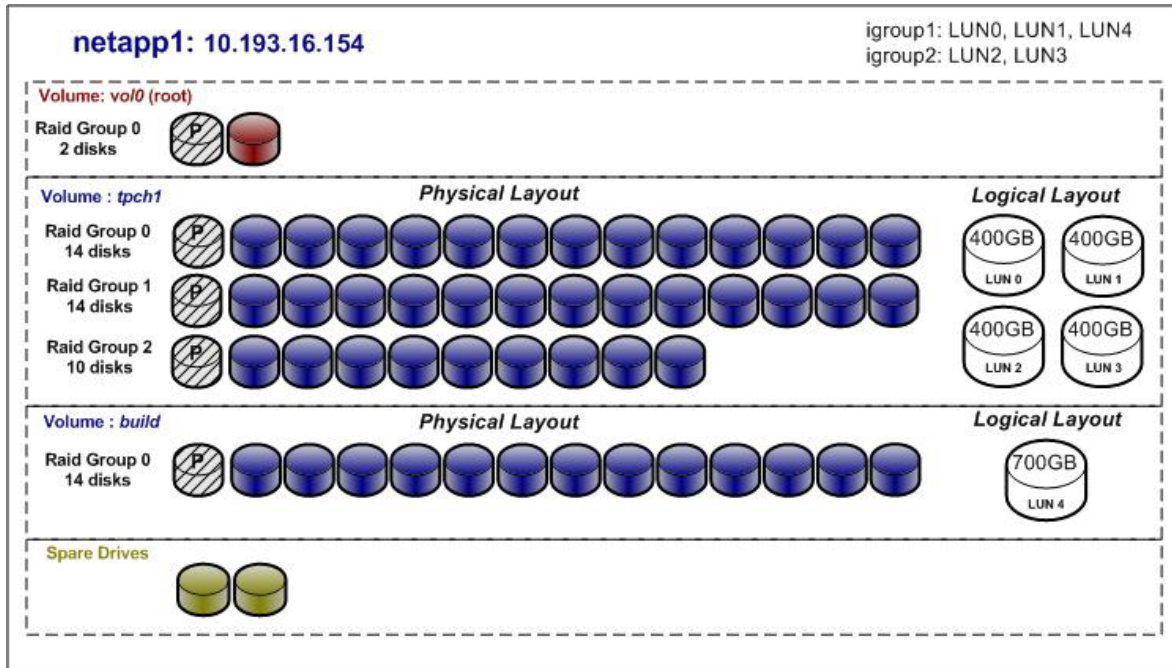


Figure 2: Disk layout for storage array FAS920-1

For the SQL Server 2000 and SQL Server 2005 test runs on the four-processor system, a total of 12*400GB LUNs were created, each containing two SQL Server database files. For the eight-processor system, a total of 24*400GB LUNs were created, each also containing two SQL Server database files. See Table 3 for full details.

Table 3. Storage Configuration

Component	Details	
Operating System	Data ONTAP 6.5.2	
Storage Interconnect	Emulex LightPulse 9002 and 9802 HBAs (Queue depth set to the maximum of 254.)	
Storage Data Link Rate	2 Gbps	
	4-way	8-way
Processor	3 * Network Appliance FAS920	6 * Network Appliance FAS920
Disks	114 * 72GB 15K RPM Drives	228 * 72GB 15K RPM Drives
LUNs	12 * 400GB	24 * 400GB
Database Files	24	48

Performance

The NetApp storage systems provided the high-performance and low-latency access to storage needed to demonstrate SQL Server 2005's abilities. These systems each delivered up to 250 megabytes per second I/O each during large table scans.

In order to ensure that any observed performance gains were the result of changes to SQL Server only, default storage array settings were used for all tests. The only exception to this was to *disable* pre-fetch during sequential read, since both SQL Server 2000 and SQL Server 2005 already generate sufficiently deep pre-fetch, enabling full system performance.

Network Appliance Value Add

Network Appliance storage solutions provide advanced backup and recovery capabilities critical for keeping SQL Server databases highly available. SnapManager® for SQL Server reduces backup times using backups based on NetApp® Snapshot™ copies. This technology allows simultaneous backups of multiple databases (of any size), and requires minimal disk space for each additional full backup. NetApp's scalable SAN storage subsystems provide simple-to-manage, cost-effective, and reliable solutions for database infrastructures requiring mission-critical high availability solutions. NetApp claims its storage solutions deliver availability in excess of 99.997%, as measured across its entire installed base. In addition, NetApp SAN solutions improve storage resource utilization, offer excellent data protection, and facilitate application testing.

Database Design

The test database was fully normalized and fully indexed on primary and foreign keys. The size of the database, including tables and indexes, was 560 Gigabytes on disk.

DSS Testing Queries

A number of typical decision support queries were run to stress the system and to evaluate changes as the system scaled. These queries, all of which differ from one another, were chosen for their complexity, their diversity in terms of data access patterns and query parameters, and because they access a large proportion of the available data. All queries were run as separate job requests. Note that all of these queries benefit from effective memory management.

Profit Enhancement (PE): This query provides the total increase in profits, had certain discounts 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 offering 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

High-Value Customers (HV): This query lists the top 'N' number of customers based on their order history. The HV query ranks customers based on the quantity of their orders, thereby helping to identify the most valuable customers for the company.

Stress characteristics:

- High number of random reads
- Requires intensive nested loop joins

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:

- Numerous floating point calculations

- Expression evaluation performance critical

Unshipped Order Summary (UO): This query is 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
- Requires intensive nested loop and hash joins

Late Received Orders (LR): This query provides a summary breakdown of products shipped to the customer prior to promised delivery date, but which 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
- Requires intensive nested loop and hash joins

Test Results

The following graphs summarize the results of the decision support system testing performed on SQL Server 2000 and SQL Server 2005.

SQL Server 2000 and SQL Server 2005 Performance Comparisons

For each query, tests were first run on SQL Server 2000 using a four-processor IA-64 platform, and then on SQL Server 2005 using the same physical configuration.

Comparative Execution Times

Figure 3 shows the query response times for High Value Customers (intensive joins), Specified Product Profit (complex query), Pricing Summary (intensive floating point calculations), and Unshipped Orders (intensive joins), and Late Received Orders (high random reads). In all cases, response times for queries run on SQL Server 2005 were faster than those run on SQL Server 2000, indicating significant performance improvements. The response times for the remaining two variables, Cost-Efficient Suppliers and Profit Enhancement (not shown) are not significantly different, although in both cases, the SQL Server 2005 numbers are incrementally better.

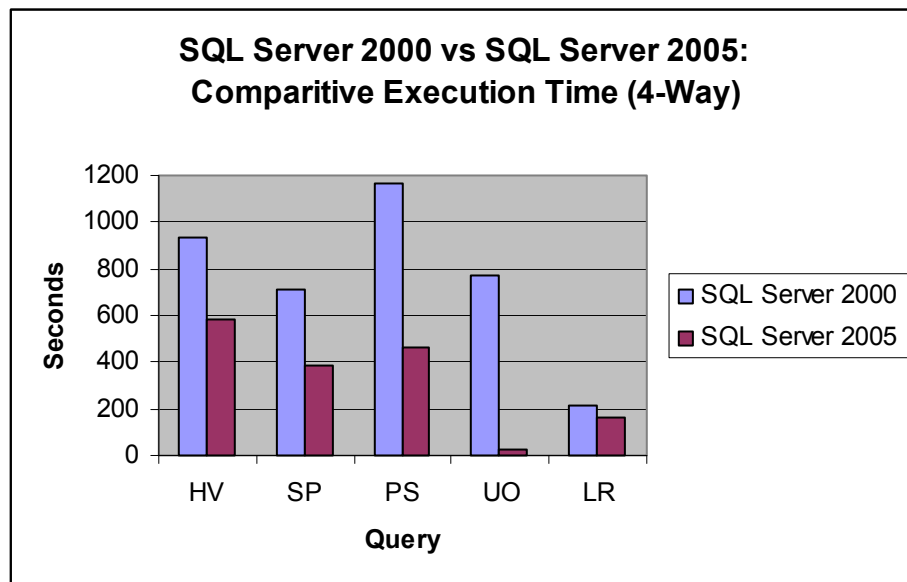


Figure 3: Comparative Response Times for HV, SP, PS, UO and LR Queries

Note that the combination of SQL Server 2005 and NetApp storage is particularly powerful. While processing full table scans, both SQL Server 2000 and SQL Server 2005 were essentially I/O constrained. However, SQL Server 2005 was able to structure the data and the I/O requests more efficiently, delivering a throughput boost on disk-limited queries between 4% and 20%.

Disk Read Performance

Figure 4 compares total disk reads (megabytes/second) for SQL Server 2000 and SQL Server 2005. For all query types, SQL Server 2005 disk reads are either equivalent or dramatically faster than the equivalent test in SQL Server 2000.

The increase in disk reads was especially high for the Unshipped Order query, demonstrating very effective utilization of the disk subsystems. In addition to query execution improvements, the query optimizer generated a very effective plan utilizing data correlation optimizations.

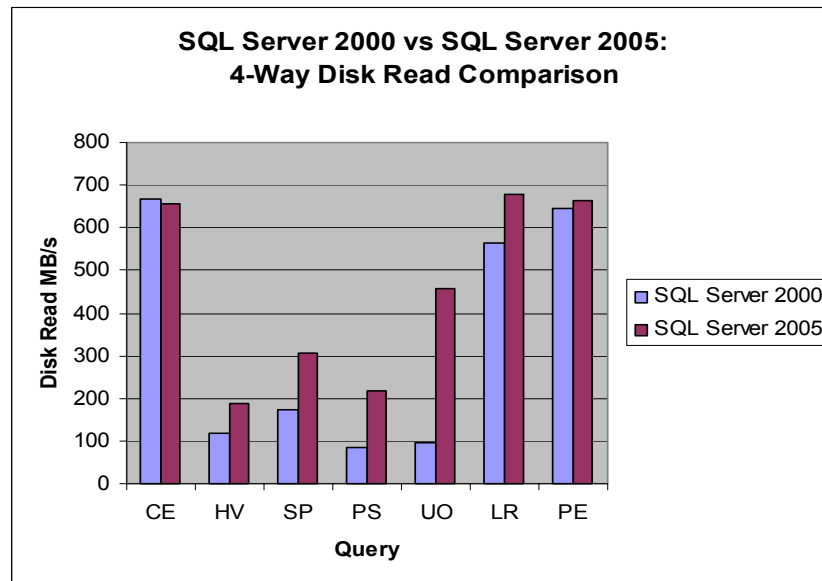


Figure 4: Comparison of Total Disk I/O

Processor Utilization

SQL Server 2005 delivers a faster query response than SQL Server 2000, but at what cost? The next round of DSS testing looks at system processor consumption—as measured by percent processor utilization—to determine whether SQL Server 2005 uses more system resources than its predecessor.

Figure 5 shows the results of four-way processor utilization. Overall, there is no significant change in CPU consumption between SQL Server 2005 and SQL Server 2000, indicating that the performance gains in SQL Server 2005 do not increase system resource costs.

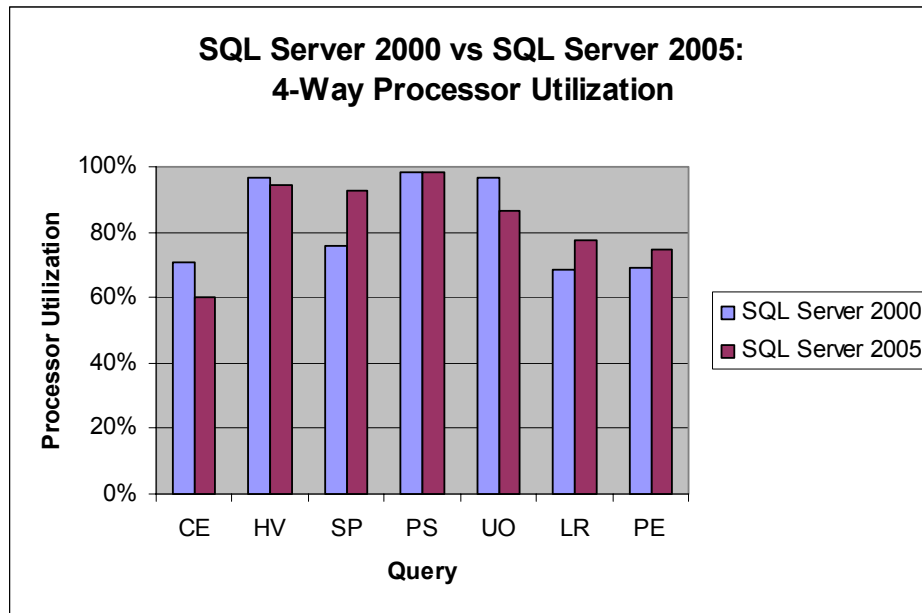


Figure 5: Comparative CPU Consumption

SQL Server 2005 Scaling Testing

In order to determine the performance impact of scaling from four processors to eight, additional SQL Server 2005 tests were run. Note that the eight-processor configuration was performed with storage arrays scaled up from three to six FAS920 storage systems.

Comparative Execution Times

The decrease in response time in Figures 6a and 6b demonstrates that as the resources were doubled, the overall system responded with an almost linear improvement (1.9 times) in system scalability.

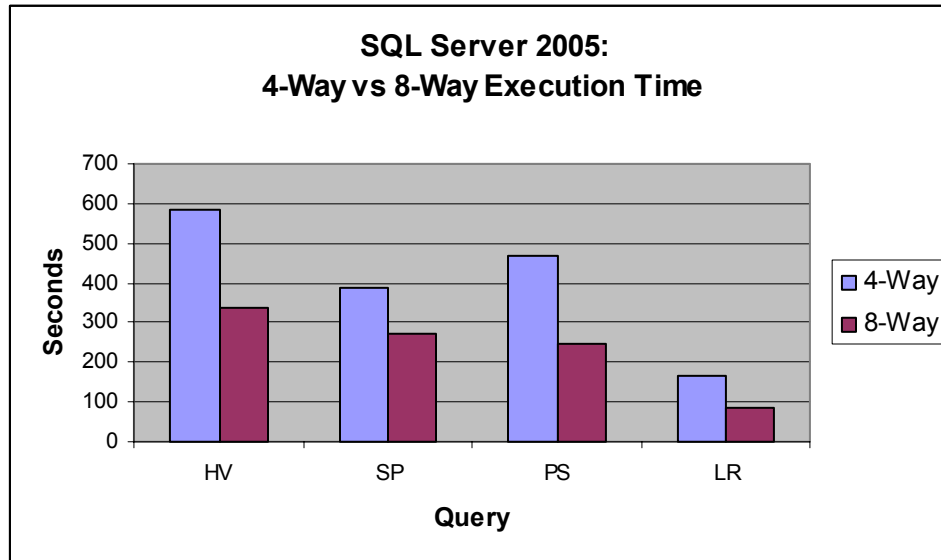


Figure 6a: 4-Processor / 8-Processor Execution Times for HV, SP, PS and LR Queries

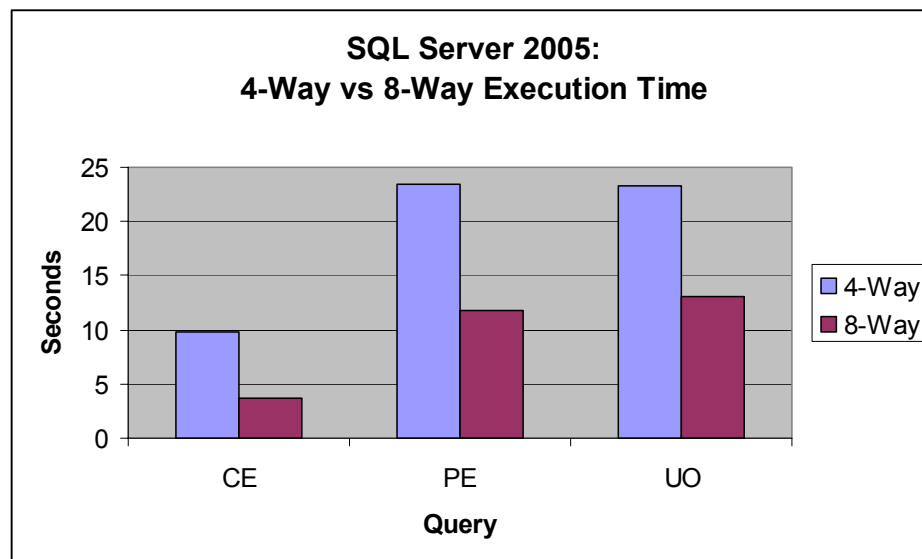


Figure 6b: 4-Processor / 8-Processor Execution Times for CE, PE and UO Queries

Disk Read Performance

How is system throughput impacted by scaling from four to eight processors? As seen in Figure 7, storage system throughput, measured as disk reads (megabytes/second), increases as the number of processors increases.

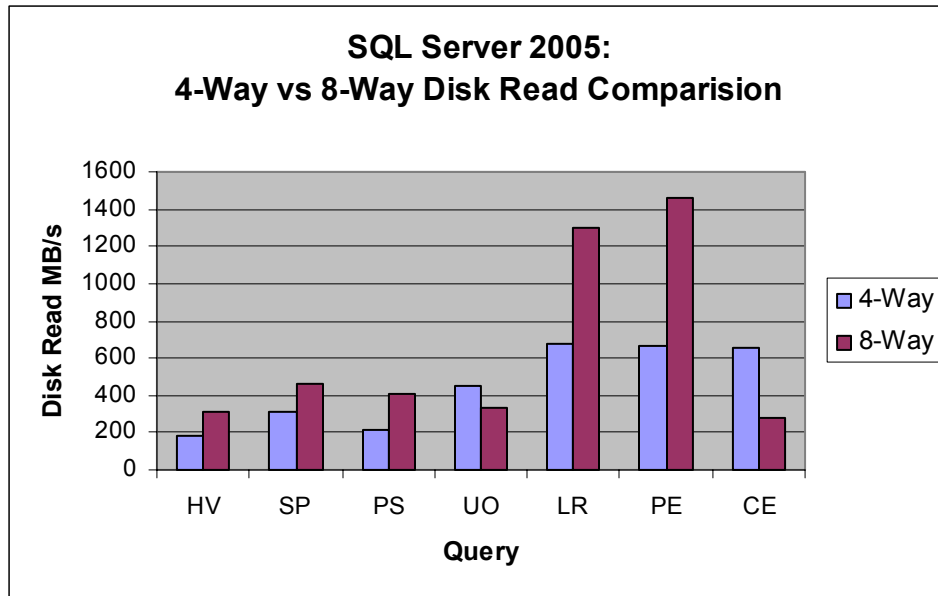


Figure 7: 4-Processor / 8-Processor Disk Read Scaling

Processor Utilization

Although query response time decreases as more system resources are added, how effective is SQL Server 2005 in making use of the overall system configuration? Figure 8 shows that CPU utilization of the four-processor and eight-processor systems are very similar, indicating that SQL Server 2005 uses system resources effectively.

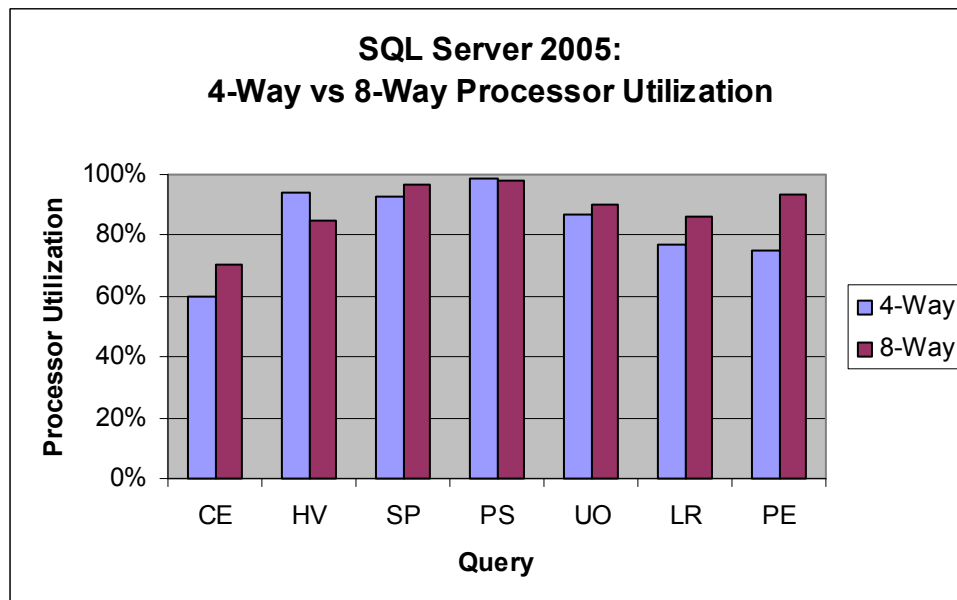


Figure 8: 4-Processor / 8-Processor Comparative CPU Usage

Summary

This paper details a number of improvements in SQL Server 2005 that enhance the performance and scalability of decision support system queries.

A series of DSS queries were run against SQL Server 2000 and SQL Server 2005 to determine the effects of these changes on performance and resource utilization. Among the key results are the following:

- SQL Server 2005 has significantly improved the response time for certain DSS-type queries compared with SQL Server 2000.
- SQL Server 2005 and SQL Server 2000 consume equivalent amounts of CPU for the same types of queries.
- SQL Server 2005 generates equivalent or faster disk reads for the given queries than does SQL Server 2000.

A second series of tests compared the scalability of SQL Server 2005 with four or eight processors. Among the key results were the following:

- The speed with which query results are returned increases linearly with additional CPU.
- Overall, disk reads are faster when scaling from four to eight processors.
- System resource usage is very similar for both four and eight processor configurations.

Based on these performance tests, the results indicate that for complex DSS queries, SQL Server 2005 performance and scalability exceeds that of SQL Server 2000. Additionally, these results demonstrate that SQL Server 2005 can make effective use of industry-leading block-based storage technologies like the Network Appliance FAS920 Storage System.

Appendix

Storage System Layout

The following Figures show the physical and logical layouts of the storage systems for hosting the SQL Server database files. All storage was provided by Network Appliance, as detailed in Table 3. Similar configuration to Figure 11 was used on other three FAS920 storage systems to perform 8-processor server configuration test runs.

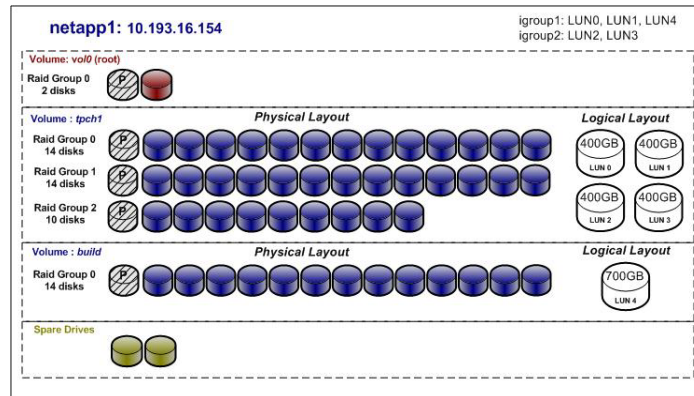


Figure 9: FAS920-1 Disk Layout

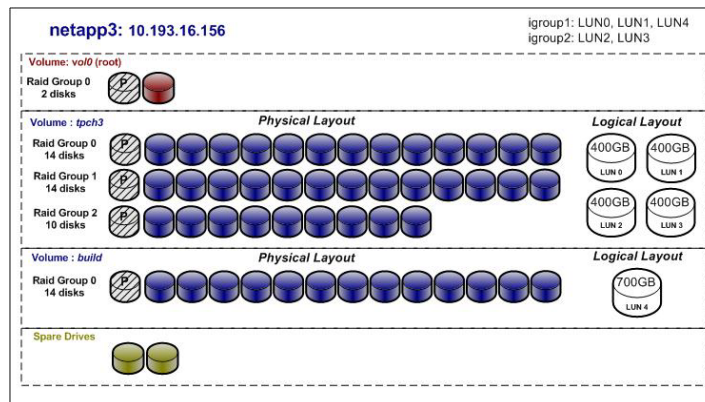


Figure 10: FAS920-3 Disk Layout

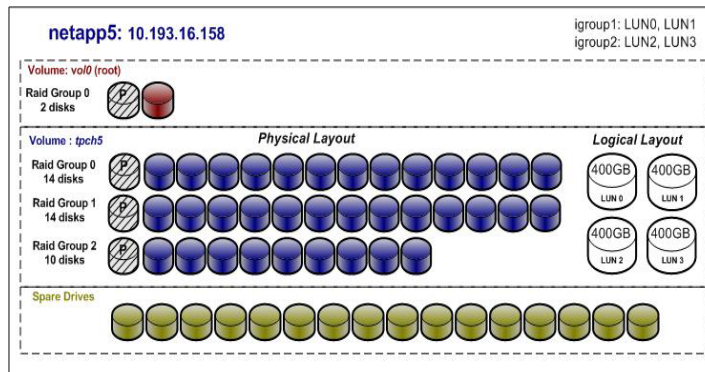


Figure 11: FAS920-5 Disk Layout

Related Links

See the following resources for further information:

- For the latest information about SQL Server 2005, see the Microsoft SQL Server site at www.microsoft.com/sql/2005/.
- NetApp storage systems information at <http://www.netapp.com/products/filer/index.html>
- NetApp SnapManager for SQL Server at <http://www.netapp.com/products/software/snapmanager-sql.html>