



TECHNICAL REPORT

**MICROSOFT SQL SERVER 2008:
DECISION SUPPORT WORKLOADS -
1TB DATA WAREHOUSE ON IA64**

Wei Liu, Ian Breitner, NetApp

Mostafa Mokhtar, Microsoft

Yves Vinter, BULL

February 2008 | TR-3650

SUMMARY

This report details a series of tests performed by NetApp, Microsoft and BULL designed to demonstrate the high performance of a large-scale, Microsoft® SQL Server 2008 Decision Support System (DSS) workload on BULL NovaScale 5325 and NetApp FAS3070 storage systems.

TABLE OF CONTENTS

- 1. INTRODUCTION 3
 - NETAPP FAS3070 3
 - BULL NOVASCALE SERVER 4
 - MICROSOFT SQL SERVER 2008..... 5
- 2. TEST ENVIRONMENT 6
 - TEST TOPOLOGY 6
 - SERVER CONFIGURATION 7
 - STORAGE CONFIGURATION 8
 - SQL SERVER 2008 DATABASE..... 10
- 3. DSS TESTING QUERIES 11
 - COUNTRY TRADE SUMMARY (CT)..... 11
 - LARGEST BUYERS (LB) 11
 - PROFIT ENHANCEMENT (PE)..... 11
 - ADVERTISEMENT PROFIT (AP) 11
 - PRICING SUMMARY (PS)..... 12
 - UNSHIPPED ORDER SUMMARY (UO)..... 12
- 4. RESULTS AND ANALYSIS 13
 - QUERY EXECUTION TIME 13
 - DISK READ PERFORMANCE..... 14
 - QUERY CPU TIME 14
- 5. CONCLUSION 15
- ACKNOWLEDGEMENT..... 15
- APPENDIX – HARDWARE SPECIFICATION 16
 - BULL NOVASCALE 5325..... 16

1. INTRODUCTION

Microsoft® SQL Server® 2008 makes a major advance in data warehouse scalability. The database engine contains numerous enhancements designed to improve both absolute performance and multi-CPU scaling on decision support workloads. This technical report documents a series of tests, run jointly by NetApp, Microsoft and BULL, which demonstrates the high performance of SQL Server 2008's decision support workloads on a BULL IA-64 server and NetApp FAS3070 storage systems.

This is the first in a series of papers that will cover various aspects of the performance and deployment possibilities of SQL Server 2008. Several preliminary tests were performed by NetApp, Microsoft and BULL on then available CTP builds of SQL Server 2008. We intend to highlight some of the advances made in SQL Server 2008 and the impact that SQL Server 2008 has on its attached storage subsystem.

It should be noted that the testing carried out for this paper was on the CTP6 (Community Technology Preview) build of SQL Server 2008 and that the measurements and results shown in this paper rely on features in CTP6. SQL Server 2008 RTM will deliver performance which is at par or better than the results in this paper.

For this series of tests a SQL Server 2008 1TB Data Warehouse was used running on a BULL 5325 NovaScale IA64 32 CPU (64-core) Server connected via 4Gb Fibre Channel (FC) to NetApp FAS3070 Storage Controllers.

The selection of the FAS3070 storage platform and the configurations used for the SQL Server 2008 tests was determined in conjunction with the Microsoft SQL Server Performance team. The requirements were to provide a flexible yet scalable and performant storage subsystem. It was important to avoid bottlenecks in the storage system or the storage interconnects. Not only was it important to understand the storage loads generated by SQL Server 2008 under various application requirements, but more importantly, it provided an opportunity for the SQL Server Performance team to examine and address any performance bottlenecks in the SQL Server 2008 engine. Both FC and iSCSI storage interconnects were also a requirement – 4Gb for the FC interconnect and both 1Gb and 10Gb Ethernet for the iSCSI interconnect. These requirements for flexible configuration, large capacity, high throughput and capability to sustain a high IOPS rate resulted in the selection of the NetApp FAS3070 as the appropriate storage platform. These requirements are met through the FAS3070's ability to provide for the I/O and bandwidth requirements as well as support the different storage interconnects.

NETAPP FAS3070

NetApp fabric-attached storage (FAS) systems simplify data management, enabling enterprise customers to reduce costs and complexities, minimize risks, and control change. NetApp FAS systems are the most versatile storage systems in the industry for storage consolidation.

The FAS3070 addresses the core requirements of the midrange enterprise storage market, delivering a superb blend of price, performance, and scalability for SQL Server databases and business applications. The compact, modular design provides integrated FC SAN and IP SAN (iSCSI) storage with scalability to over 500 disk drives. The FAS3070 storage controller supports both FC and SATA disk drives for tiered storage. FAS3070 systems support as many as 32 FC ports or 32 Ethernet ports, including support for both 2Gb, 4Gb FC and 10 Gigabit Ethernet.

The FAS3070 runs the NetApp Data ONTAP® operating system, which is optimized for fast, efficient, and reliable data access and retention. Data ONTAP 7G dramatically simplifies common storage provisioning and management operations. LUNs and volumes created and configured using FlexVol® technology can be dynamically expanded or contracted with a single command. Host-based NetApp SnapDrive® extends this flexible storage provisioning capability to databases and applications. FlexVol also enables thin provisioning, which avoids the cost of over-provisioning and the time-consuming reconfiguration typical with other storage solutions. Another Data ONTAP 7G feature, FlexClone®, instantaneously creates cloned LUNs or volumes

without requiring additional storage. FlexClone can dramatically improve the effectiveness and productivity of application and database development and pre-deployment testing.

FAS hardware designs and the Data ONTAP operating system are tightly integrated to provide resilient system operation and high data availability. FAS systems incorporate redundant and hot-swappable components, and patented double-parity RAID-DP™ (the NetApp high-performance implementation of RAID 6) provides superior data protection with negligible impact on performance. NetApp Snapshot™ technology provides up to 255 data-in-place, point-in-time images per LUN or file system, available for near-instantaneous file-level or full data set recovery, and the minimal performance overhead makes it uniquely suited for protecting production data. Host-based SnapManager® software integrates Snapshot management with applications, ensuring consistent backup images and application-level recovery in minutes. SnapMirror® utilizes Snapshot copies to provide incremental block-level synchronous and asynchronous replication; SnapVault® uses it for block-level incremental backups to another system. Together, these SnapSuite™ products help deliver the high application-level availability that enterprises require for 24×7 operation.

BULL NOVASCALE SERVER

The NovaScale® 5325 server belongs to the NovaScale 5005 Series of high-end servers from BULL, ideally suited for Data Warehousing/BI applications, support of large data bases and HPC systems and clusters. By supporting future generations of dual-core Intel® Itanium® 2 processors, the NovaScale® 5325 protects customers' investments.

- Up to 32 dual-core Intel® Itanium® 2 processors (codenamed Montecito)
- Up to 512GB DDR200/266 SDRAM
- Up to 16x internal hot-swap disks (max storage 4.8TB)
- 40U rack

BULL NovaScale servers are a real alternative to large traditional servers in the data center. These modular, mainframe-class servers benefit from BULL's expertise in large enterprise systems. They combine computing power with cost-effective standards and support Windows®.

COMPACT DESIGN AND MODULARITY TO REDUCE TCO

The NovaScale 5325 is based on second generation modules twice as compact as the former ones. These modules can be easily combined to deliver 8- to 32-socket SMP (Symmetrical Multi-Processing) servers in a 40U rack. The packaging enables easy building: modules can be added as the enterprise grows.

ENHANCED PHYSICAL PARTITIONING FOR HIGHER FLEXIBILITY OF USE

The NovaScale 5325 provides a minimum of 4 sockets per partition, thereby supplying MISs with an unprecedented flexibility of use for large SMP servers supporting Windows environments.

OUTSTANDING PRICE/PERFORMANCE RATIO

The outstanding price/performance ratio of the NovaScale 5005 Series is achieved thanks to the low cost of ownership of standard technologies combined with their breakthrough performance and scalability.

The Intel® Itanium® architecture of BULL NovaScale servers is built on the EPIC technology (Explicitly Parallel Instruction Computing) providing the scalability and reliability required by complex business-critical and technical applications.

HIGH SCALABILITY TO ADDRESS THE ENTERPRISE PROGRESSIVE GROWTH

The high scalability of NovaScale 5325 is provided by the BULL FSS (FAME Scalability Switch) and the Intel® E8870 chipset interconnecting dual-core Intel® Itanium® 2 processors in order to deliver high performance 32-socket servers.

INVESTMENT PROTECTION

The NovaScale 5005 Series based on dual core Intel® Itanium® 2 processors offers extended investment protection by supporting future generations of the Intel® Itanium® 2 Processor Family (IPF).

MICROSOFT SQL SERVER 2008

Microsoft SQL Server 2008 includes enhancements which improve the performance of DSS workloads to meet the mission-critical needs of large enterprise customers, running in terabyte data warehouses. Here are a few highlights of new performance enhancements for DSS workloads:

- **Star Join.** With dimensionally modeled data warehouses, a big part of the workload typically consists of what are known as star join queries. These queries follow a common pattern that joins the fact table with one or several dimension tables. In addition, star join queries usually express filter conditions against the non-key columns of the dimension tables and perform an aggregation on a column of the fact table. With SQL Server 2008, customer will experience significant performance improvements for many star join queries that process a significant fraction of fact table rows. The new star join optimization uses a series of hash joins, building a hash table for each dimension table that participates. As a byproduct of building a hash table, an additional structure, a *bitmap filter*, is created. The bitmap filters are applied while scanning the fact table to eliminate almost all of the rows that would otherwise be eliminated in later joins. Eliminating rows early in the query plan with the bitmap filter is more efficient than using only the hash table.
- **Partition Table Parallelism.** Data warehouse applications typically collect large amounts of historical data in fact tables, which are often partitioned by date. In SQL Server 2005, queries that touch more than one partition use one thread and thus one processor core per partition. This sometimes limits the performance of queries that involve partitioned tables, especially when running on symmetric multiprocessor (SMP) computers with many processor cores. Partitioned table parallelism improves the performance of parallel query plans against partitioned tables by better utilizing the processing power of the existing hardware, regardless of how many partitions a query touches.
- **Partition-Aligned Indexed Views.** In a typical data warehouse scenario where a fact table is partitioned by date, indexed views (summary aggregates) are defined on this table to help speed up queries. When you switch in a new table partition, the matching partitions of the partition-aligned indexed views defined on the partitioned table switch too, and do so automatically. In addition, indexed view matching is also automatic, so no query rewrites are required to use the aggregates.
- **Few Outer Row.** It is common for DSS query plans to include nested loop joins with a very selectively filtered parallel scan on the outer side. Depending of the clustering of these few rows, this quickly leads to an unfavorable distribution of data among available threads. Additional exchange operators redistribute rows evenly among threads and significantly improve parallelism and performance for such scenarios.

For a complete list and more in-depth description of data warehousing improvements in SQL Server's relational database management system, as well as in Integration Services, Analysis Services, and Reporting Services, please refer to this online technical white paper by Microsoft:

<http://www.microsoft.com/sql/techinfo/whitepapers/SQL2008IntroDW.msp>

2. TEST ENVIRONMENT

The NetApp equipment onsite in the Microsoft SQL lab serves as a high performance, highly flexible, scalable, and reliable storage platform for SQL Server 2008 development and testing. For instance, the NetApp storage platform is being used for SQL Server 2008 development tracking, validation, performance testing and regression testing. The equipment is also intended for a variety of different SQL Server workload performance baselines and comparisons. As such the storage configurations used in this testing is over-configured ensuring that the storage or the storage inter-connect will not be a bottleneck during the test runs.

The following sections detail the test topology, server and storage configurations used, and database layout in this series of DSS workload tests.

TEST TOPOLOGY

Decision support workloads are frequently very I/O intensive. The high I/O requirements of these tests were met using NetApp FAS3070 storage systems connected to a BULL server using multiple 4 Gbps Fibre Channel interconnects.

Figure 1 below illustrates the topology of the SQL Server 2008 1TB data warehouse workload testing configuration. The BULL server was connected to 12*FAS3070 controllers using 24*4Gbps FC interconnects, 2*FC links per controller. Note: due to the choice of workload size, this topology only uses a subset of both the BULL server capability as well as the NetApp storage capability deployed in the Microsoft SQL Server lab.

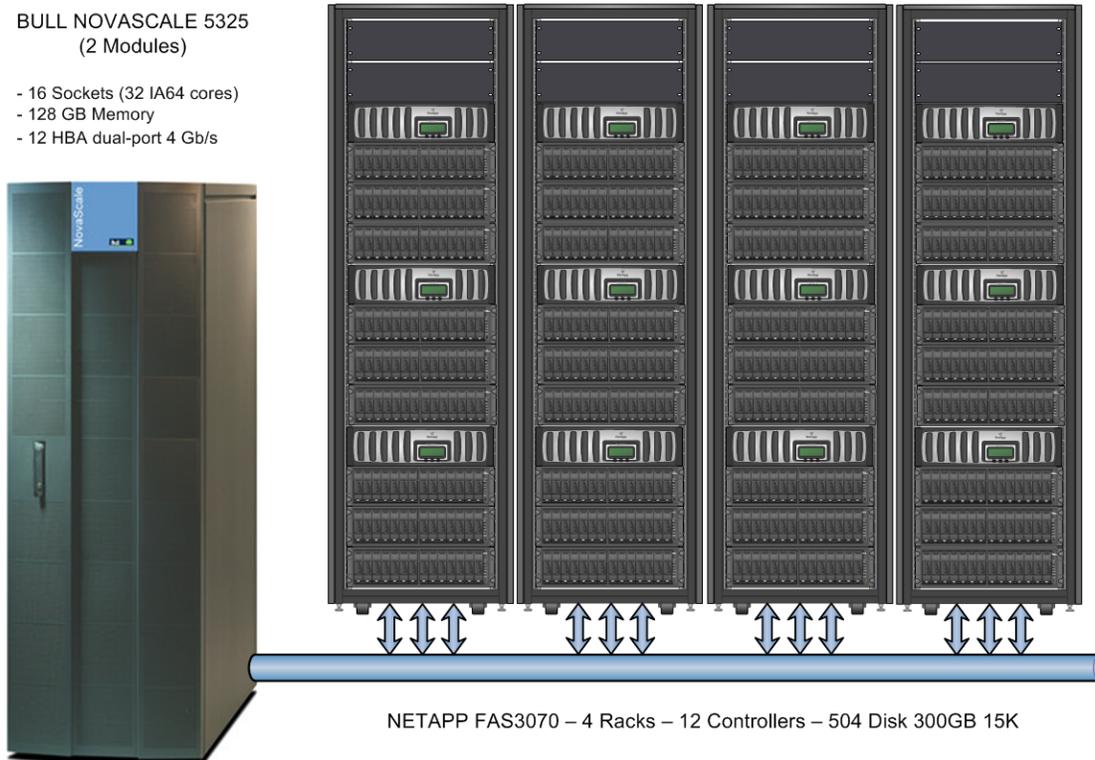


Figure 1 - Server, Storage and Fibre Channel interconnects for the 1TB Data Warehouse

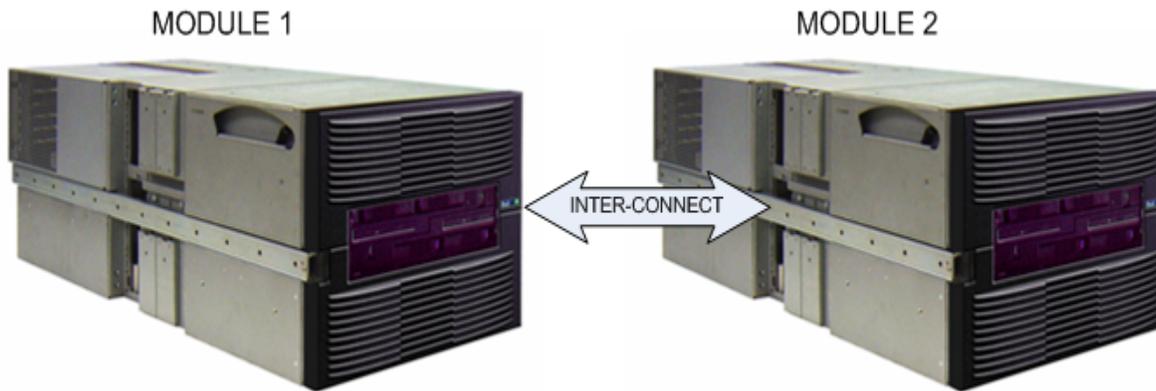


Figure 2 - Server Configuration

SERVER CONFIGURATION

The tests were conducted using 32-core BULL NovaScale Server (IA64) with 128GB RAM, running Windows 2003 Data Center Edition SP2 (build 3790) on a bi-module partition.

Figure 2 shows the BULL server configuration. The two modules shown in Figure 2 were linked through a high speed interconnect controlled by BULL Fame Scalability Switch (FSS, 25GB/s).

Each module had 8 Itanium® 2 Montecito processors (1.6 Ghz, 24MB L3 DC, Double Core), 64GB DDR 266 SDRAM, and 6 Emulex dual-port 4Gbps HBAs.

Figure 3 shows the detailed server configuration used in the 1TB data warehouse workload testing.

STORAGE FOR OPERATING SYSTEM

The Windows 2003 Server Data Center Service Pack 2 Operating System was installed on an internal RAID 1 set (two 73GB 15K Ultra 320 SCSI drives), which was managed by an LSI 22320-R RAID controller.

STORAGE FOR TRANSACTION LOGS

An Emulex LP1055DC FC HBA (Dual-port 2Gb/s) was installed and connected to 16 external FC drives (73GB 15K) to provide the storage for the SQL Server transaction logs. Windows software RAID was used to stripe the 16 drives into a RAID 0 device. Note that this configuration was specifically required by Microsoft for this test in order to deliver database only load to the storage systems. The NetApp FAS 3070 controllers have more than enough performance and capacity to handle both the database and log loads for this and other higher I/O load testing. In future tests the logs will also be housed on the storage controllers.

NETAPP STORAGE FOR DSS DATABASES

NetApp storage controllers formed the bulk of the storage for SQL Server 2008 DSS databases. Figure 3 shows 12 NetApp storage controllers (FAS3070) were directly connected to the BULL server via 12 Emulex LP111002-M4 Dual-port 4Gb/s HBAs. Thus, two 4Gbps FC links were connecting each FAS3070 controller to the BULL server.

Also shown in Figure 3 is that 8 LUNs (500GB) per controller, a total of 96 LUNs, were mounted to the BULL server for DSS databases.

The detail information about LUN and storage configuration is discussed in the next section. The database layout information is provided in the subsequent section.

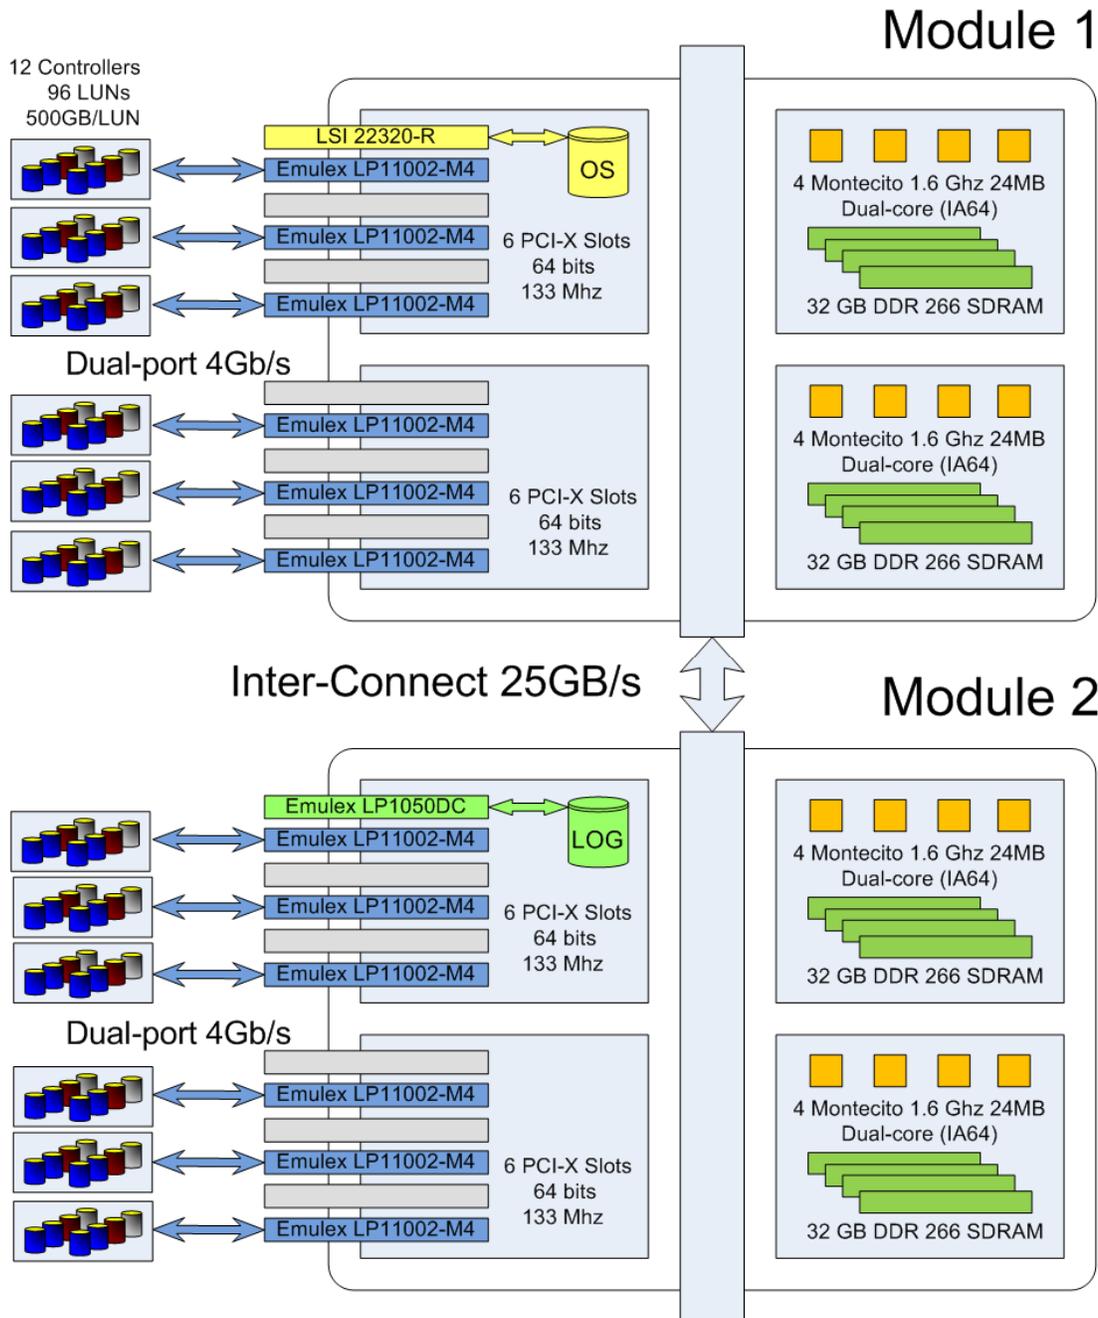


Figure 3 - Detailed server configuration. Also shown is the FC Interconnect to the storage systems.

STORAGE CONFIGURATION

The storage for these testing was provided by NetApp FAS3070 FC Storage systems (or controllers), configured as direct-attach Fibre-Channel storage systems. 12 FAS3070 controllers were used for the 32 CPU tests. The FAS3070 controllers ran version 7.2.2P3 of Data ONTAP. Each FAS3070 was configured with 42 300GB 15K RPM FC disks.

It is worth noting that the FAS3070 storage system described in this paper is used by the Microsoft SQL Server performance team for multiple parallel projects and was configured for maximum testing flexibility rather than

sized for this particular test. As a result, the tested configuration could provide more I/O throughput and storage capacity than was utilized during this test. This design allowed the storage to be provided such that there would be no bottlenecks associated with the storage system or the interconnects to that storage. This means that more controllers and more spindles were provided than would otherwise have been necessary to support the load during the test. There will be additional technical reports that will focus on deployment configurations and best practices for SQL Server 2008 with NetApp storage – these will be published closer to the RTM release of SQL Server 2008.

STORAGE VIRTUALIZATION

Data ONTAP 7G introduced powerful storage virtualization features that dramatically simplifies storage provisioning, greatly improves storage utilization, flexibility and manageability. Two of the key concepts are aggregate and FlexVol[®] volumes. An aggregate is a logical container for pools of physical disks that are organized into one or more RAID-DP groups. An aggregate is the logical layer that decouples volumes from the underlying physical storage. A FlexVol volume is a logical entity that is separated from the physical storage that contains the associated data. One or more FlexVol volumes can reside within an aggregate. A FlexVol volume can grow or shrink in size, constrained only by the hard limits of the aggregate size or the soft limits set when the volume was first created.

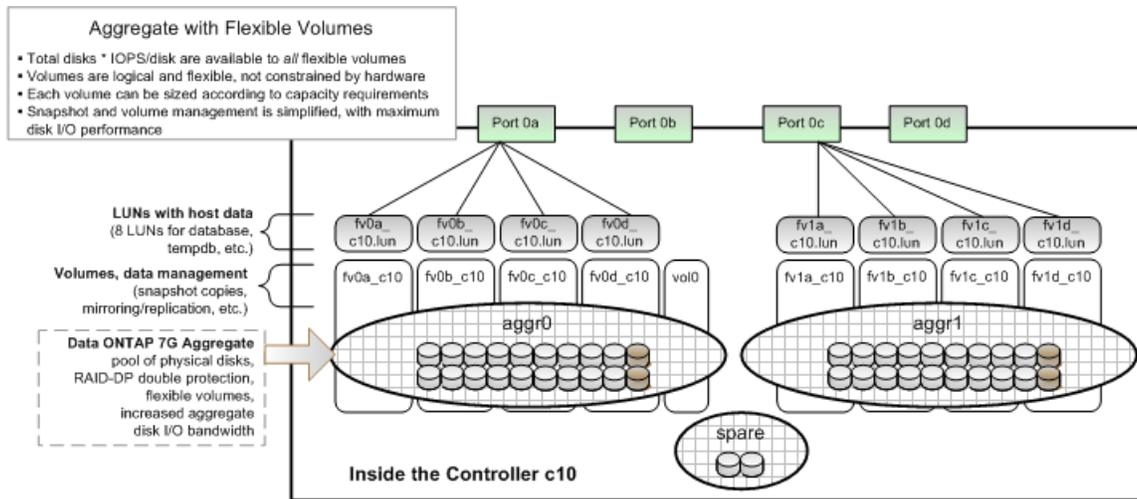


Figure 4 - FAS3070 Configuration

Figure 4 shows detailed configuration of one of the FAS3070 controllers. Note that all 12 controllers were configured in identical fashion.

A total of 40 of the 42 drives within each controller were used to form 2 aggregates, each using 20 drives. The remaining 2 drives were left as hot swappable spares. As seen in Figure 4, 2 drives within each aggregate are in dark gray color. These 2 drives were the 2 parity drives within the RAID-DP group, which protect against double disk failure.

Four FlexVol volumes were then created within each aggregate resulting in 8 volumes per controller. Note also that residing in aggr0 is a vol0, which is needed by Data ONTAP for controller management.

Within each FlexVol volume, a LUN (500GB) was created. Thus, 4 LUNs per aggregate, 8 LUNs per controller (96 LUNs total on all 12 controllers) were created and then mapped to the BULL server for storing DSS databases.

Figure 4 also illustrates that 4 LUNs residing within an aggregate were physically connected to the BULL server using one 4Gbps FC port. Therefore, 2 ports per controller were used in this test environment. Table 1 below sums up the storage configuration information.

Storage Entity	Quantity
NetApp FAS3070 controllers	12
300GB 15K 4Gbps FC disks	504 (42 per controller)
Aggregates	24 (2 per controller)
FlexVol volumes	96 (8 per controller, 4 per aggregate)
LUNs	96 (8 per controller, 4 per aggregate, 1 per volume)
4Gbps FC links	24 (2 per controller, 1 per aggregate)

Table 1 - Summary of NetApp storage used in the 1TB Data Warehouse

NETAPP SNAPSHOT

It is worth mentioning that each aggregate was configured with adequate space for NetApp Snapshots. In the testing process, NetApp Snapshots were taken after initial LUN creation and again after DSS database load completion. NetApp Snapshots enabled fast and efficient backup and restore.

Also worth mentioning is the fact that all NetApp storage configuration was carried out via easy to use scripts. The configuration of aggregates, volumes, and LUNs was done in minutes, while NetApp Snapshots were created in seconds.

SQL SERVER 2008 DATABASE

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, indexes, and backup, is 3.7 TB on disk.

DATABASE LAYOUT

Among the 96 LUNs mounted on the BULL server, 48 of them contained databases, shown in Figure 3 in blue. The other 48 LUNs contained tempdb files (in brown and gray in Figure 3), of which half (in gray) also contain a database backup.

From Figure 3 it is apparent that the large-scale DSS database was evenly distributed across all 12 controllers, available spindles and FC interconnects.

SQL SERVER 2008 TUNING OPTIONS

SQL Server 2008 performs most of the necessary tuning automatically and dynamically configures its parameters based on usage and availability of system resources. In addition, the affinity mask option (sp_configure 'affinity mask') was used to associate SQL Server threads to 32 cores on the BULL server.

3. DSS TESTING QUERIES

A number of typical decision support queries were run to stress the system and to evaluate the performance of the system. These queries, all of which differ from one another, were 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 queries were run as separate job requests. The set of queries included the following:

COUNTRY TRADE SUMMARY (CT)

This query calculates the total revenue of products sold between certain countries during a specified time period. The CT query measures the trade volume and assists in re-negotiation of shipping contracts.

Stress characteristics:

- Complex query
- Query plan and optimization critical

LARGEST BUYERS (LB)

This query is used to extract a specified number of large buyers whose purchases have exceeded certain transaction threshold. For example, if N=50, the 50 largest buyer would be returned, along with information on the customers and associated transactions. The EB query helps companies identify the customers they should value the most.

Stress characteristics:

- Intensive index scans and lookups required
- Hash Joins and Nested Loop Joins critical
- Intensive sorting

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 indexing scan required

ADVERTISEMENT PROFIT (AP)

This query measures the percent of profit that was made as result of advertisement for a given time period. The AP query helps companies gauge the return on advertising.

Stress characteristics:

- Hash Joins and Nested Loop Joins 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 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:

- Complex query
- Query plan and optimization critical
- Intensive sorting

The queries are summarized in the table below.

Stress characteristic	CT	LB	PE	AP	PS	UO
Complex query	X					X
Query plan and optimization critical	X					X
Intensive index scans and lookups required		X				
Hash Joins and Nested Loop Joins critical		X		X		
Intensive sorting		X				X
I/O intensive			X			
Intensive indexing scan required			X			
Floating point calculations					X	
Expression evaluation performance critical					X	

Table 2 - Summary of query stress characteristics

RESULTS AND ANALYSIS

For each DSS query, tests were run using the configurations described in the preceding sections. Baseline measurements were first taken on SQL Server 2005 Enterprise Edition SP2; then comparison measurements were taken on SQL Server 2008, CTP6 build.

This section shows the test results in terms of execution times, average read MB/sec, and server CPU time.

QUERY EXECUTION TIME

Figure 5 shows the relative query response times for Profit Enhancements, Advertisement Profit, Unshipped Orders, Country Trade Summary, Pricing Summary, and Largest Buyers. Each query's execution time was normalized using the baseline execution time in SQL Server 2005 SP2. In all cases, response times for queries run on SQL Server 2008 were faster than those run on SQL Server 2005, indicating significant performance improvements. Table 2 summarizes the relative performance gain.

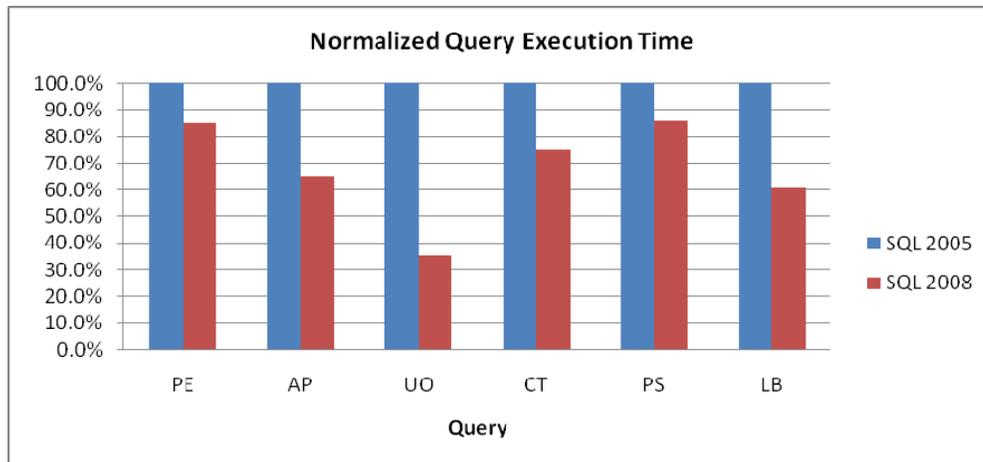


Figure 5 - Normalized Query Execution Time (lower is better)

Query Name	Performance Gain (%)
Profit Enhancements (PE)	15.8%
Advertisement Profit (AP)	36.4%
Unshipped Orders (UO)	64.3%
Country Trade Summary (CT)	25.2%
Pricing Summary (PS)	14.4%
Largest Buyers (LB)	39.1%

Table 2 - SQL Server 2008 query performance gain over SQL Server 2005

DISK READ PERFORMANCE

Figure 6 compares total disk read throughput (normalized with the values for SQL Server 2005 representing the baseline of 100%) for SQL Server 2005 and SQL Server 2008. For all queries, SQL Server 2008 disk read throughput is either equivalent to or dramatically better than that of SQL Server 2005.

The increase in disk read throughput was especially large for the Unshipped Order query, demonstrating very effective utilization of the storage subsystems. In addition to query execution improvements, the query optimizer generated a very effective plan.

The storage controllers were running under 50% utilization. It should be noted that the storage systems were over-configured for this workload – in part because the storage systems are intended to handle the larger and more complex workloads of other test scenarios and partly to ensure that storage bottlenecks would not occur during the testing. Later technical papers will provide guidance on configuring NetApp storage controllers for various workload types.

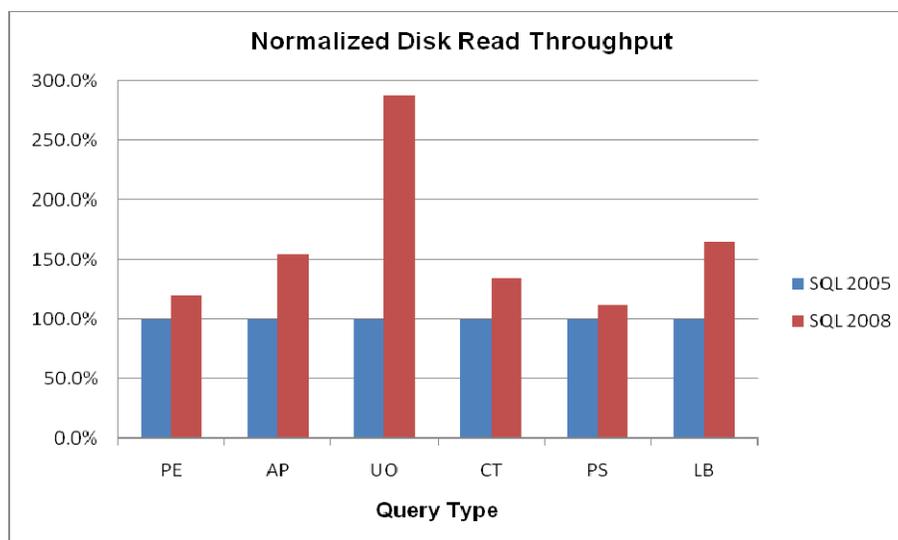


Figure 6- Query: Disk Read Throughput (higher is better)

QUERY CPU TIME

Figure 7 shows the relative query CPU time, computed by multiplying query execution time and average processor utilization for each query, and then normalized using the query CPU time in SQL Server 2005 SP2. The reduction in processor resource spent in query execution corresponds nicely with query performance gain, once again demonstrating the impressive improvement SQL Server 2008 made in DSS workload performance.

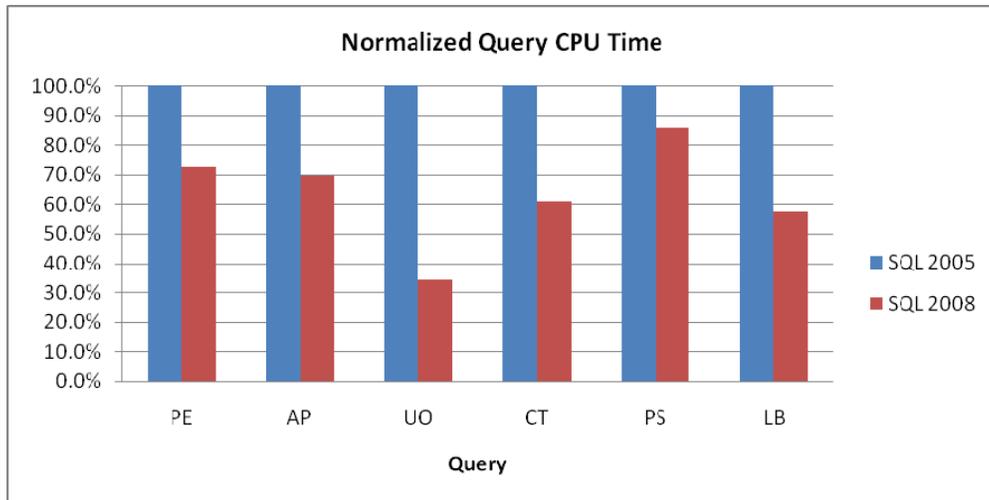


Figure 7 - Normalized query CPU time (lower is better)

4. CONCLUSION

This is the first in a series of reports that examine the high performance of SQL Server 2008 on decision support workloads using a configuration on a BULL IA-64 server and NetApp FAS3070 storage systems. From these preliminary test results, it can be seen that there have been significant performance improvements in the SQL Server 2008 engine.

The selection of the NetApp FAS3070 storage systems for this SQL Server 2008 Data Warehouse workload provided more than the required performance for this workload type with the storage utilization being under 50% of available throughput and capacity. The ability to utilize the snapshotting capability of the storage controllers to rapidly re-provision for each test run provided great value.

ACKNOWLEDGEMENT

The authors wish to thank BULL and NetApp for providing the hardware equipment made this work possible. Thanks to Microsoft SQL performance team for the opportunity to work jointly on SQL Server 2008 prior to its shipping. Special thanks to Shirley Wang and Ashit Gosalia of Microsoft SQL performance team for their tireless work and instrumental contributions.

APPENDIX – HARDWARE SPECIFICATION

BULL NOVASCALE 5325

NovaScale Design	
Platform	40U (195.5 cm H) housing: <ul style="list-style-type: none"> • NovaScale 5325 server • Platform Administration Processor (PAP) with one external USB modem • 17" SVGA LCD console with keyboard and touchpad • 2x 8-port KVM switch
Form Factor	4 interconnected rack mount modules (32U drawer)

Processors	
Number of processors	x8 (standard), up to 32 dual-core Intel® Itanium® 2
Type	<ul style="list-style-type: none"> • 1.4GHz /12 MB L3 DC • 1.6GHz /8 MB L3 DC • 1.6GHz /18 MB L3 DC • 1.6GHz /24 MB L3 DC

Architecture & Board	
Processing Architecture	SMP (Symmetrical Multi-Processing), cc Numa architecture with a 2- to 4-socket board
Chipset	Intel® E8870 chipset with FSS (BULL Fame Scalability Switch)

Physical Partitioning	
Number of Partitions	From 1 to 8 partitions

Memory	
Memory Min/Max	16GB to 512GB (128 DIMM slots) 4GB to 64GB memory per processor board (4 to 8 processor boards)
Type	DDR200/266 SDRAM – 4-way interleaved

SCSI Disk Subsystems	
Internal SCSI HDD	Up to 16x internal hot swap disks (max storage: 4.8 TB) PCI SCSI RAID Ultra 320 Dual port adapter (optional)

PCI Adapters (optional)	
Ethernet (Copper)	PCI-X (dual/quad port) / PCI Express (dual port) 10/100/1000 (RJ45)
Ethernet (Fibre)	PCI-X 1000
SCSI	PCI-X SCSI Ultra 320 (dual port) / PCI-X SCSI RAID Ultra 320 (Dual Port)
Fibre Channel	PCI-X and PCI Express (mono or dual 4Gb/s)

I/O Box (*)	
Number of I/O box	x1 I/O box for the NovaScale 5325 (plus x7 optional)
USB ports	x2
Serial port	x2
SVGA video port	1x (connected to LCD console)
Ethernet port	1x Ethernet 10/100/1000 (RJ45)
Bus slots	4x hot-plug PCI-X slots [4x 64-bit/133MHz (3 free)] 2x hot-plug PCI-Express x8 slots

DVD/CD-ROM drive	1x DVD/CD-ROM writer
------------------	----------------------

Power Supply	
Hot Swap PSU	x2 + x2 1200W redundant (per module)
Voltage	200-240 VAC
Frequency	50 to 60Hz

Ventilation	
Fan specifications	Hot swap redundant fans

Physical Specifications	
Size (HxWxD)	195 x 60 x 129.5 cm (40U rack)
Weight	650 kg (40U rack)
Operating constraints	15°C to 30°C 5°C to 50°C (non-operating mode)

Platform and System Management	
Administration and Maintenance	PAM (Platform Administration and Maintenance)
System Management	BULL NovaScale Master (BULL System Management)

Regulatory & Safety	
Regulatory compliance	UL, FCC and CE compliance, RoHS, DEEE

* When partitioned, the number of usable legacy I/Os (USB, serial, SVGA, Ethernet, DVD/CD-ROM ports) depends on the NovaScale 5325 configuration.

FAS3000 Series Technical Specifications

Technical Highlights	
Maximum RAID Group Sizes	<p>RAID-DP™ (high performance RAID 6) FC – 28 (26 data disks plus 2 parity disks) SATA – 16 (14 data disks plus 2 parity disks)</p> <p>RAID 4¹ FC – 14 (13 data disks plus 1 parity disk) SATA – 7 (6 data disks plus 1 parity disk)</p> <p>SyncMirror® (high resiliency RAID 1) Options:</p> <ol style="list-style-type: none"> 1. RAID-DP + RAID 1 2. RAID 4 + RAID 1
SAN Protocol Support	Fibre Channel Protocol (FCP) for SCSI (2Gb and 4Gb); fabric-attached and direct-attached; iSCSI
File Service Protocol Support	NFS V2/V3/V4 over UDP or TCP, PCNFSD V1/V2 for (PC) NFS client authentication, Microsoft® CIFS
Other Protocol Support	HTTP 1.0, HTTP 1.1 virtual hosts
LUNs	Up to 2,048 front-end LUNs on FAS3040 and FAS3070; up to 1,024 front-end LUNs on FAS3020
FlexVol® Volumes	Up to 500 per controller
Volume/Aggregate Size²	16TB
Snapshot™ Copies	Up to 127,000
Number of Supported SAN Hosts	Up to 4 ports per controller configured as front-end (host) connectivity ports Up to 4 directly connected servers (per controller and per active/active configuration) Up to 256 SAN connected servers (per controller and per active/active configuration)
Supported Configurations	<p><u>Highly Available Controller Configurations</u> Active/active controller with Controller Failover (CFO)³, active/active controller with stretch (non-switch) MetroCluster, active/active controller with Fabric-Attached MetroCluster⁴</p> <p><u>Highly Available Back-End (Disk) Configurations</u> Dual-path, Multipath HA Storage</p>
Reliability	Redundant hot-swappable controllers, cooling fans, power supplies, optics, and RJ-45 ports
Management	Full-duplex 10/100/1000Base-T Ethernet onboard console, diagnostic LED/LCD, Remote LAN Management (RLM) Card ⁵ , SNMP, telnet, SSH, HTTP, Web (SSL), host scripting, e-mail alerts
Security	<p>Virus Protection,</p> <p>Authentication: Kerberos (MIT/Active Directory) Authorization: LDAP, NIS Auditing: File access logging Encryption: IPSec, SSL, SSH, SMB signing Role Based Access Control (RBAC)</p>

Scalability	FAS3070		FAS3040		FAS3020	
	Active/	Single	Active/	Single	Active/	Single

	Active Configuration	Controller	Active Configuration	Controller	Active Configuration	Controller
Maximum Number of Disk Drives	504 ⁶	504	336 ⁷	336	168 ⁸	168
Maximum Drives per Back-end (Disk) FC Loop	84	84	84	84	84	84
Memory (ECC)	16GB	8GB	8GB	4GB	4GB	2GB
Nonvolatile Memory	1GB	512MB	1GB	512MB	1GB	512MB

¹ RAID-DP is the recommended configuration for drives greater than 144GB.

² Maximum volume/aggregate size is calculated using base 2 arithmetic (1TB = 2⁴⁰ bytes).

³ For more information about cluster configurations, refer to <http://now.netapp.com/NOW/knowledge/docs/ontap/rel704/html/filer/cluster/index.htm>.

⁴ Fabric-Attached MetroCluster is not currently supported on FAS3070 or FAS3040. Please contact your NetApp sales representative for details regarding future support.

⁵ RLM module is standard on FAS3070 and FAS3040 and is an optional capability on FAS3050 and FAS3020 that is included at no charge with SupportEdge Premium contract.

⁶ Available with all drive types = 1TB starting in Data ONTAP 7.2.4.

⁷ Available with all drive types = 1TB starting in Data ONTAP 7.2.4.

⁸ Available with all drive types = 500GB. The maximum number of 750GB disk drives is 112. The maximum number of 1TB disk drives is 84.

On-board I/O Ports/Slots	FAS3070		FAS3040		FAS3020	
	Active/Active Configuration	Single Controller	Active/Active Configuration	Single Controller	Active/Active Configuration	Single Controller
Fibre Channel Ports (target or initiator)	8 ports 1, 2, or 4Gb (auto-sensing)	4 ports 1, 2, or 4Gb (auto-sensing)	8 ports 1, 2, or 4Gb (auto-sensing)	4 ports 1, 2, or 4Gb (auto-sensing)	8 ports of 2Gb	4 ports of 2Gb
GbE Ports	8	4	8	4	8	4
Ultra160 SCSI Ports	N/A	N/A	N/A	N/A	2	1
I/O Expansion Slots⁹	6	3	6	3	6	3

I/O Card Support	FAS3070		FAS3040		FAS3020	
	Active/Active Configuration	Single Controller	Active/Active Configuration	Single Controller	Active/Active Configuration	Single Controller
2-port FC Disk Adapters (max.)	6	3	6	3	6	3
2-port 4Gb FC Target Adapters (max.)	4	2	4	2	4	2
4-port 4Gb FC Disk/Tape HBAs (max.)	6	3	6	3	6	3
2-port Ultra320 SCSI Tape HBA	6	3	6	3	6	3
Optional Network/SAN Cards (max.)						
2-port FC Tape Adapters	6	3	6	3	6	3
4-port 1GbE TOE (copper) cards	6	3	6	3	4	2
1-port 10GbE TOE (fiber) cards	NA	NA	NA	NA	4	2
2-port 10GbE TOE (fiber) cards¹⁰	6	3	6	3	NA	NA
2-port 1GbE (copper or fiber) cards	6	3	6	3	6	3
2-port 1Gb iSCSI (copper or fiber) HBA	6	3	6	3	6	3

Software

Operating System	Network Appliance™ Data ONTAP®
-------------------------	--------------------------------

Operating Systems Supported	Windows® 2000, Windows Server 2003, Windows XP, Linux®, Sun™ Solaris™, AIX, HP-UX, Mac OS, VMware, ESX	
Software Features	<p style="text-align: center;"><u>Standard</u></p> RAID manager, including RAID-DP Snapshot Fast Boot NIS DNS FilerView® , FlexVol FlexCache FlexShare™ Network Data Management Protocol (NDMP) Maintenance tCenter	<p style="text-align: center;"><u>Licensed</u>¹¹</p> FlexClone® FlexCache® MultiStore® SnapMirror® SnapRestore® Single Mailbox Recovery SnapVault® SnapMover® LockVault™ SnapLock® Compliance SnapLock Enterprise SyncMirror MetroCluster NearStore® Personality Option Manageability Software Application Suite: SnapManager® for Microsoft Exchange SnapManager for Microsoft SQL Server SnapManager for Oracle Server Suite: ApplianceWatch™ for HP Openview ApplianceWatch for Tivoli® SnapDrive® for UNIX® SnapDrive for Windows Data Suite: Protection Manager™¹² Virtual File Manager™ - EE Virtual File Manager - ME Storage Suite: CommandCentral Storage File Storage Resource Manager Operations Manager (formerly DFM) SAN Manager

⁹ FAS3070 and FAS3040 have PCI-e I/O slots, while FAS3020 have PCI-x I/O slots.

¹⁰ 10GbE TOE (fiber) card will be available after the FAS3070 and FAS3040 are available. Please contact Network Appliance regarding availability.

¹¹ For more information about available software, refer to www.netapp.com/products/software.html.

¹² Protection Manager was previously known as Business Continuity Option.

Disk Drive Storage Expansion Shelf Specifications	
Storage Expansion Shelves Supported (Fibre Channel, SATA)	FC: DS14mk4 FC DS14mk2 FC, DS14 (for upgrade migrations only)

	SATA: DS14mk2 AT All shelves have 14 bays for 3.5" low-profile FC or SATA disk drives
Storage Expansion Shelf to Controller Interface	Fibre Channel-Arbitrated Loop (FC-AL)
Storage Expansion Shelf Configurations	<u>New System Configurations</u> FC: DS14mk4 FC with ESH4 (Embedded Switched Hub 4) Electronics SATA: DS14mk2 AT with AT-FCX Electronics AT-FCX <u>Supported via Upgrade Migrations from Installed Systems</u> FC: DS14mk4 FC with ESH4 Electronics DS14mk2 FC with ESH4 or ESH2 Electronics DS14 with ESH2 or ESH Electronics DS14 with LRC (Loop Resiliency Circuit) Electronics (FAS3020 only) SATA: DS14mk2 AT with AT-FCX Electronics
Power Supply/Cooling Fans	Dual, redundant, hot-pluggable, integrated power supply/fan assemblies ¹³
AC Power/Max. Current¹⁴	SATA: 7200 RPM drives: 3.01 / 1.51 / 7.4A @ 100VAC / 200VAC / -40VDC FC: 10,000 RPM drives: 3.09 / 1.56 / 7.61A @ 100VAC / 200VAC / -40VDC FC: 15,000 RPM drives: 3.96 / 1.94 / 10.22? @ 100VAC / 200VAC / -40VDC
Thermal Rating¹⁵	SATA 7200 RPM drives: 1017 BTU/hr FC: 10,000 RPM drives: 1043 BTU/hr FC: 15,000 RPM drives: 1394 BTU/hr
Operating Acoustic Noise	49 dBA sound pressure (LpA) @ normal operating conditions (at 22°C and at sea level)
Dimensions (height/width/depth)	3 EIA U (5.25", 13.3 cm)/19" IEC rack-compliant (17.6", 44.7 cm)/20" (50.85 cm)
Weight	77 lb (35 kg) fully loaded

FC Switch Support—Tape/SAN	
Vendor	Model
Brocade	Silkworm 2400/2800 (TapeSAN only)
Brocade	Silkworm 3200/3250/3800/3850/3900/4100/12000/24000/200E/48000
McData	ES 3016/3032, Intrepid 6064/6140/ED-5000, Sphereon 3216/3232/4300/4400/4500/4700
Qlogic	5600/5602
Cisco	MDS 9120/9140/9216/9506/9509/9513/9124

FC = Fibre Channel; LVD = Low Voltage Differential. Visit www.netapp.com/solutions/data_protection_devices.html for current supported tape devices. Visit www.netapp.com/osn/info/config.html for current FC switch support.

NetApp Systems Cabinet	
20A and 30A single phase cabinets	20A Cabinet ? 1 controller & 13 DS14-type shelves; 2 controllers & 12 DS14-type shelves
20A Cabinet - Includes 4 power cords with NEMA L6-20 or IEC 309 or Australia AS/NZ53123-20 connectors	30A Cabinet ? 1 controller & 12 DS14-type shelves; 2

30A Cabinet – Includes 2 power cords with NEMA L6-30 or IEC 309 connectors	controllers & 11 DS-14 type shelves
AC Power	200 to 240VAC
Current	20A Cabinet - Maximum 32A total 30A Cabinet - Maximum 24A total
Dimensions (height/width/depth)	78.75" (200 cm)/23.6" (60 cm)/37.4" (95 cm)
Min. Clearances	48" (122 cm) front and rear
Shipping Height	90.6" (230 cm)
Weight (empty)	275 lb (125 kg)

¹³ For detailed disk shelf specifications, refer to <http://now.netapp.com/NOW/knowledge/docs/hardware/NetApp/site/guide11.htm#1792411>.

¹⁴ These numbers represent a single supply on, with fans on high setting at 100V, 200V, -40V DC.

¹⁵ Thermal rating, like current, is based on drive type/speed. This value is relatively constant regardless of input voltage, so only one value is required for each.

FAS System Specifications	Active/Active Configuration	Single Controller Configuration
Controller Weight	150 lb (68 kg)	75 lb (34 kg)
Controller Height	10.24" (26 cm), fits into 6U space	5.12" (13 cm), fits into 3U space
Controller Width	19" IEC rack-compliant (17.7", 44.9 cm)	
Controller Depth	24" (61 cm), 30" (76.2 cm) with cable management bracket	
Operating Temperature, Altitude, and Relative Humidity	10° C to 40° C (50° F to 104° F); at <=/ 3,000 m (at <=/ 10,000') elevation; 20% to 80% relative humidity, non-condensing (28° C wet bulb temperature)	
Non-operating Temperature and Relative Humidity	-40° C to 65° C (-40° F to 149° F); 10% to 95% relative humidity, non-condensing, in original container	
Operating Acoustic Noise	≤ 54 dBA sound pressure (LpA) @ normal operating conditions (at 22°C and at sea level)	
Min. Cabinet Clearances for Airflow	10" (25.4 cm) in front, 12" (30.5 cm) in rear	
Min. Cabinet Clearances for Service	30" (76.2 cm) in front, 30" (76.2 cm) in rear	
Compliance	RoHS-compliant	
Safety/Emissions/Immunity	<p>Safety: EN 60950, CE, CSA 60950, UL 60950, CB IEC60950-1 (all national deviations), EN60825-1, IRAM, GOST-R, BSMI</p> <p>Emissions/Immunity: FCC Part 15 Class A, ICES-03, CE, MIC, VCCI, AS/NZS CISPR 22, EN55022, EN55024, EN61000-3-2, EN61000-3-3, CoC (South Africa), BSMI</p>	

Model	Input Voltage		Active/Active			Single Controller		
			Typical		Maximum	Typical		Maximum
			One side	Both sides		One side	Both sides	
FAS3020	100-120VAC	Current, A	.2.4	4.8	6.8	1.2	2.4	3.4

		Power, W	236	472	672	118	236	336
		Thermal, BTU/hr	805	1610	2288	402.5	805	1144
	200-240VAC	Current, A	1.4	2.8	3.6	0.7	1.4	1.8
		Power, W	229	458	658	114.5	229	329
		Thermal, BTU/hr	783	1566	2244	391.5	783	1122
	-40 to -60VDC	Current, A	5.7	11.4	16.4	2.85	5.7	8.2
		Power, W	226	452	656	113	226	328
		Thermal, BTU/hr	771	1542	2236	385.5	771	1118
FAS3040	100-120VAC	Current, A	3.4	6.8	7.4	1.7	3.4	3.7
		Power, W	338	676	726	169	338	363
		Thermal, BTU/hr	1152	2304	2476	576	1152	1238
	200-240VAC	Current, A	1.9	3.8	3.8	0.95	1.9	1.9
		Power, W	330	660	716	165	330	358
		Thermal, BTU/hr	1127	2254	2442	563.5	1127	1221
	-40 to -60VDC	Current, A	7.4	14.8	15.88	3.7	7.4	7.94
		Power, W	296	592	636	148	296	318
		Thermal, BTU/hr	1011	2022	2168	505.5	1011	1084
FAS3070	100-120VAC	Current, A	3.7	7.4	8	1.85	3.7	4
		Power, W	361	722	800	180.5	361	400
		Thermal, BTU/hr	1233	2466	2730	616.5	1233	1365
	200-240VAC	Current, A	2.1	4.2	4.2	1.05	2.1	2.1
		Power, W	355	710	774	177.5	355	387
		Thermal, BTU/hr	1212	2424	2640	606	1212	1320
	-40 to -60VDC	Current, A	9.4	18.8	21.2	4.7	9.4	10.6
		Power, W	376	752	846	188	376	423
		Thermal, BTU/hr	1283	2566	2884	641.5	1283	1442

© 2008 Network Appliance, Inc. | All rights reserved. | Specifications subject to change without notice.



© 2008 NetApp, Inc. All rights reserved. Specifications subject to change without notice. NetApp, the NetApp logo, Decru, NearStore, SnapLock, and SnapVault are registered trademarks and NetApp, DataFort, FlexClone, and FlexVol are trademarks of NetApp, Inc. in the U.S. and other countries. Windows is a registered trademark of Microsoft Corporation.. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.
TR-3650