THE TOLLY GROUP

No. 206132 May 2006

BlueArc Corp.

Titan 2200 Network Storage System

Throughput Evaluation under Mixed I/O Scenarios

TOLLY 17 YEARS 1989-2006 Test Summary

Premise: *At high performance levels and big bandwidths, even well balanced general-purpose servers suffer from bus contention, resource bottlenecks, and software overhead. The Titan 2200's massively parallel FPGA-based architecture is optimized for data movement and implementing advanced features at wire speed. Potential buyers will find that the Titan 2200 delivers heightened performance, unlike general-purpose NAS products that suffer from bottlenecks.*

BlueArc Corp. commissioned The Tolly Group to examine the performance of the company's Titan 2200, a purpose-built, network-attached storage (NAS) solution that the company claims is the fastest on the market and also the most scalable, able to consolidate and manage up to 512 terabytes of data in a single storage pool and file system.

Tolly Group engineers focused tests on the average throughput of a single Titan 2200 node handling I/O requests between 48 client PCs. The back-end SAN storage was comprised of multiple Engenio Information Technologies, Inc., Engenio 2882 storage systems, consolidated into a single storage pool, name space and file system. Throughput was measured for a variety of read, write and mixed read/write operations from back-end disk, not cache. Tests were conducted in January 2006.

Tests show that a single Titan 2200 can deliver uniform performance even as NFS record sizes being accessed by clients increase from 4K bytes to 64K bytes.

Test Highlights

- Delivers average sustained throughput of 718 MB/sec from a single Titan when handling a mix of read/write I/O operations from 48 clients to disk
- Achieves sustained average read throughput of 656 MB/sec to a single Titan, when handling read operations from 48 clients through to disk, not from cache
- Sustains write throughput of 450 MB/sec on average to a single Titan, when handling writes from 48 clients to disk
- For higher throughput, clustering is supported with global name space and will be tested in a separate report

BlueArc Titan 2200 — Average I/O Throughput During Initial Read/Reread and Random Read/Write Operations
as Reported by IOzone 3.257

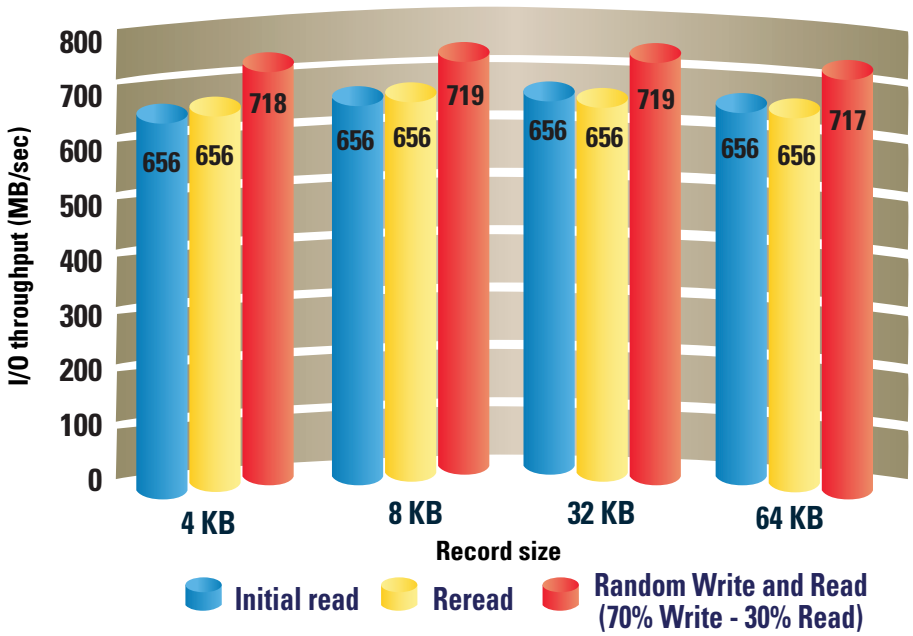

Source: The Tolly Group, January 2006 *Figure 1*

Executive Summary

In the world of NAS, solutions typically have been designed around an adaptation

of a general purpose server, operating system and file system into a file server. The BlueArc Titan 2200 represents an alternative design approach.

BlueArc's Titan 2200 architecture was developed using field programmable gate arrays (FPGAs). It employs many FPGAs with specific tasks to allow for parallelism and optimize performance. An FPGA is a specially made digital semiconductor that contains an array of logic gates that can be programmed by the design engineer (and updated later on-site with downloadable configuration files) for a specific application.

The Titan 2200 employs 12 high-density FPGAs with up to 32 gigabytes of distributed memory and multiple pipelines. This enables it to simultaneously read/write data using native CIFS and NFS file-based protocols, the Internet Small Computer System Interface (iSCSI) block-level storage networking protocol, and Network Data Management Protocol (NDMP) for standards-based backups.

For read and re-read operations, the Titan 2200 delivered 656 MB/sec of average sustained throughput to disk. Here, too, we see steady, uniform throughput as record sizes increased. Even during the reread scenario, throughput fluctuations were less than 1%. (See Figure 1.)

When handling random read/write operations, the Titan 2200 delivered an average of 718 MB/sec of sustained throughput over the six bonded 1-Gbps Ethernet connections tested. Again, throughput did not waver as record sizes of 4K, 8K, 32K and 64K bytes were tested. (See Figure 1.)

Engineers looked at initial write and re-write operations and found that the Titan 2200 sustained throughput of 450 MB/sec on average. Even as the record size increased from 4 KB, to 8KB, 32 KB and finally, 64 KB, the initial write throughput never wavered by more than 3%. On the

BlueArc Titan 2200
Average I/O Throughput During Initial Write/Rewrite Operations
as Reported by IOzone 3.257

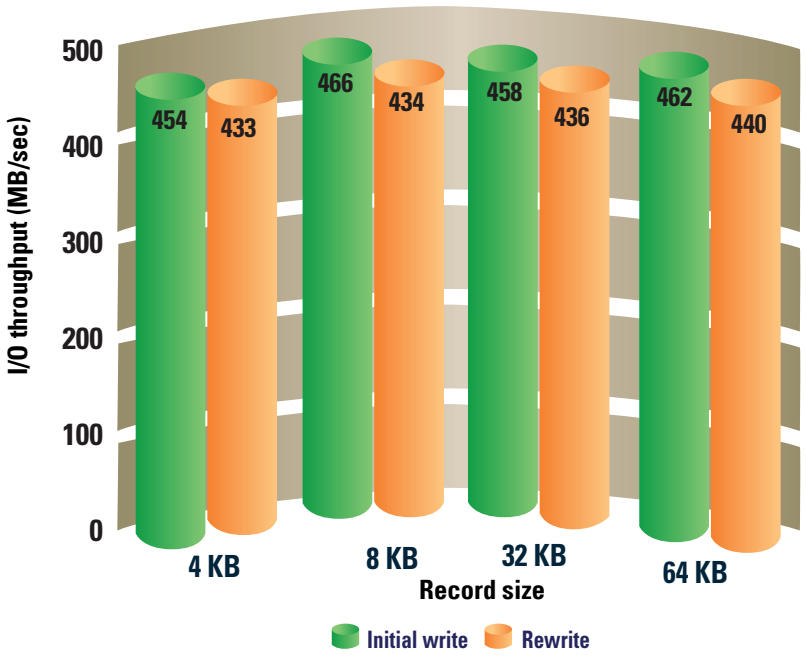

Source: The Tolly Group, January 2006 *Figure 2*

rewrite scenario, the Titan 2200 averaged I/O throughput of 436 MB/sec, never wavering more than 1% as record sizes increased. (See Figure 2.)

Tests show that the Titan 2200 can support significant numbers of connections, meaning users can have a lot of application tiers contending for attention at the same time with no degradation in throughput or I/O. This performance profile makes it ideal for high performance applications or consolidation of multiple network storage

BlueArc Titan 2200
Single-Node Throughput Performance (MB/sec)
as Reported by IOzone 3.257

I/O Operation	Record size			
	4 KB	8 KB	32 KB	64 KB
Initial write	454.1	466.4	457.7	462.4
Rewrite	433.3	434.4	435.9	440.0
Initial read	656.1	656.1	656.3	656.3
Reread	656.3	656.2	656.3	656.1
Random Write and Read (70% Write - 30% Read)	718.3	718.6	718.9	717.2

Source: The Tolly Group, January 2006 *Figure 3*

servers, into a faster centralized Titan solution. Titan can also be clustered for high-availability applications or additional throughput. Although not tested, it is also designed for high transactional loads per SPECsfs, a standard benchmark on *www.spec.org*.

Test Configuration & Methodology

The Tolly Group engineers tested BlueArc's Titan 2200 Network Storage System as a single-node connected to a pool of disk storage. Test scenarios simulated 48 client computers accessing the disk storage through the Titan 2200. As shown in Figure 4, 48 client PCs acting as load generators were connected to a Cisco Catalyst 6500 switch using Gigabit Ethernet (GbE) links. The Titan 2200 was connected to the Cisco Catalyst switch using six GbE links, aggregated using standard IEEE 802.3ad. On the back-end SAN, the Titan 2200 was in turn connected to two Brocade Silkworm 200E Fiber Channel switches using two 4 Gbps Fiber Channel links to each switch. The pool of disk storage consisted of 10 Engenio 2882 with dual RAID controller storage systems each configured with an expansion shelf to hold a total of 28 146-GB Fibre Channel hard drives (2 sets of 12+1 drives in RAID-5 with 1 drive in Hot Standby). Each Brocade Fibre Channel switch was connected to a RAID controller on each storage system through a 2 Gbps Fibre Channel link.

The client computers (acting as load generators) were generic 1U PCs (distributed by Amax) based on a Supermicro X5DPA-GG motherboard, each configured with two Intel 2.8GHz Xeon CPUs, 1GB DDR-266 RAM, and two integrated Intel Gigabit LAN 8254X controllers. On each client a minimalized version of Red Hat Fedora Core 4 (kernel, 2.6.14-1.1656_FC4smp) was installed.

Iozone version 3.257 was used as the test tool to simulate various small and large sequential and random I/O operations. Each load generator launched its own Iozone thread, generating 48 threads in parallel. Each thread reads/writes its own test file of 4 GB to or from the disk storage using various record sizes of 4/8/32/64 KB at a time. The test file size of 4 GB was chosen to be four times the memory size of each load generator to ensure that the data was not cached locally or on the file server, and to ensure that the performance measured was the worst-case performance a customer would experience when reading and writing to disk.

To start the test, the test file of 4 GB was written to the disk storage using

BlueArc Corp.

Titan 2200 Network Attached Storage (NAS) System

Throughput under Mixed I/O Scenarios

BlueArc
Titan 2200
Product Specifications*

Hardware Features

- Purpose-built hardware-based network storage server
  - Supports up to 200,000 IOPS and over 700MB/sec throughput
  - File systems up to 256TB or 512TB with global name space
- Hardware embedded file system and operating system ensures performance and scalability
- Total of 33.8GB of distributed memory
- Four-slot modular chassis with 40GB/sec passive backplane
  - Network Interface Module
    - 6 wire-speed Gigabit Ethernet ports
    - Link aggregation and Jumbo Frames support
    - Four Ethernet side-band management ports
  - Two File System Modules with 2GB of NVRAM
  - Storage Interface Module
    - Four 4GB Fibre Channel ports
    - Two 10GbE interconnect ports for active/active clustering
    - 8GB of dedicated storage sector cache
- Fibre Channel SAN backend with modular storage
  - Parallel hardware-based RAID striping for high-speed storage access
  - Supports both Fibre Channel and Serial ATA storage

File & Operating System Features

- Upgradeable hardware-based operating system & 64-bit file system
  - Supports standard NFS v2 and V3, CIFS, iSCSI, NDMP, and FTP protocols
  - Managed through CLI, GUI (HTTPS), SSL, SSH, SNMP, NIS, and WINS
  - Up to 1,024 snapshots or 1 per minute
  - Virtual volumes & quotas
  - Parallel RAID striping
  - Accelerated data copy
  - Anti-virus scanning
  - Virtualized storage pools for scalable file systems
- Advanced Features
  - Cluster name space
  - Virtual servers
  - Transparent Data Migrator
  - Incremental Block Replication
  - Synchronous Volume Mirroring
  - WORM file system

For more information contact:
BlueArc Corporation
50 Rio Robles
San Jose, CA 95134
www.bluearc.com

**Vendor-supplied information not verified by The Tolly Group*

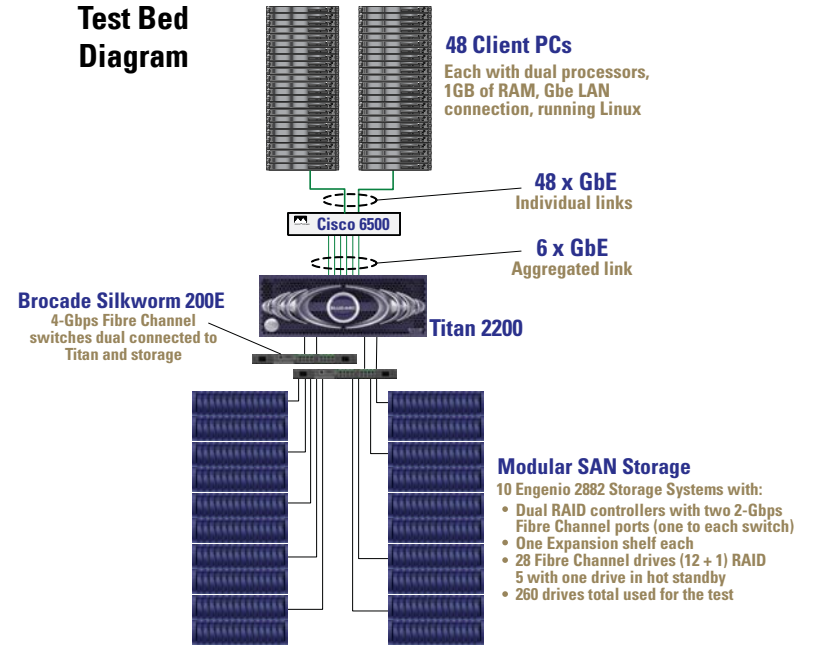

Source: The Tolly Group, January 2006 *Figure 4*

sequential write operation. Once the write operation was completed, the test file was rewritten to the disk storage using the sequential rewrite operation. After the rewrite, the test file previously written to the disk storage was read back using "sequential read" operation. Following the read operation, the test file was re-read from the disk storage using the "sequential reread" operation. Finally, each load generator simulated a random read/write with 70% write and 30% read operations. This series of tests was repeated for record sizes of 4-, 8-, 32- and 64-KBytes. Iozone measured the I/O throughput for each read/write scenario for each record size.

The Tolly Group gratefully acknowledges the providers of test equipment used in this project.

Vendor	Product	Web address
Freeware	Iozone ver. 3.257	http://www.iozone.org

Terms of Usage

USE THIS DOCUMENT ONLY IF YOU AGREE TO THE TERMS LISTED HEREIN.

This document is provided, free-of-charge, to help you understand whether a given product, technology or service merits additional investigation for your particular needs. Any decision to purchase must be based on your own assessment of suitability.

This evaluation was focused on illustrating specific features and/or performance of the product(s) and was conducted under controlled, laboratory conditions and certain tests may have been tailored to reflect performance under ideal conditions; performance may vary under real-world conditions. Users should run tests based on their own real-world scenarios to validate performance for their own networks. Commercially reasonable efforts were made to ensure the accuracy of the data contained herein but errors and/or oversights can occur.

The test/audit documented herein may also rely on various test tools the accuracy of which is beyond our control. Furthermore, the document relies on certain representations by the sponsor that are beyond our control to verify. Among these is that the software/hardware tested is production or production track and is, or will be, available in equivalent or better form to commercial customers.

The Tolly Group provides a fee-based service to assist users in understanding the applicability of a given test scenario to their specific needs. Contact us for information.

When foreign translations exist, the English document is considered authoritative. To assure accuracy, only use documents downloaded directly from The Tolly Group's Web site.

Project Profile

Sponsor: BlueArc Corp.

Document number: 206132

Product class: Network-attached storage solution

Products under test:

- BlueArc Titan 2200 Network Storage (NAS) System

Product status:

- Generally available

Testing window: January 2006

For more information on this document, or other services offered by The Tolly Group, visit our World Wide Web site at http://www.tolly.com, send E-mail to sales@tolly.com, call (561) 391-5610.

The Tolly Group doc. 206132 rev. clk 15 May 06