

Cisco SN 5420 Storage Router



Executive Summary

One of the hottest topics in the Storage Industry is the suitability of iSCSI for traditional storage applications. iSCSI is the encapsulation of SCSI commands inside TCP/IP. Some storage experts question whether iSCSI adds significant latency to storage access because of TCP/IP processing. Testing of the Cisco SN 5420, the industry's first iSCSI networking product, indicates that TCP/IP processing adds minimal additional latency to overall storage access. To measure the latency, Cisco developed a test environment that measured the number of read I/O operations per second (IOPS) using various block sizes that are widely used by database applications. Comparison tests were run with iSCSI attached and Fibre Channel attached storage. The test results show the SN 5420, which implements the iSCSI protocol, is suitable for traditional storage applications.

Test Methodology

The term *latency* is used in numerous ways when describing computer system delays. In this document, latency refers to the time required to request a block of data from the disk drive and transfer it back to the application buffer in the host or server.

Latency in read operations to disk can be hidden by read-ahead, operating system (OS) buffering and the read cache RAM buffers located on the disk drives. The test code was specifically written to diminish these effects. To minimize the effect of read-ahead, both at an OS and disk drive level, the test program used a pseudo-random number generator to select a new seek position before each read operation. To remove any OS buffering, all tests were conducted on a Linux "raw" device, which bypasses the OS buffers. To help ensure that the influence of the cache on the disk drive was kept to a minimum, a 2 GB section of the drive was used that diminished the effect of any read caching done by the 4 MB of RAM cache on the drive.

Figure 1 iSCSI Test Setup





Figure 2 Fibre Channel Test Setup



The tests were run on a 933 MHz single-processor Intel Pentium III-based system. The PC motherboard was an Intel STL2 server board, which has two 64 bit/66 MHz PCI card slots and four 32 bit/33 MHz slots. The Fibre Channel tests used a Qlogic 2100 HBA and the iSCSI tests used a 3COM 3C985B-SX Gigabit Ethernet Network Interface Card (NIC). Both the Qlogic HBA and the 3COM Gigabit Ethernet NIC are based on similar serial gigabit technology with comparable raw throughput rates. Only one card at a time was plugged into the motherboard and the card was plugged into a 64 bit/66 MHz PCI slot to ensure maximum performance.

Tests were run on three block sizes: 1 KB, 8 KB, and 64 KB. These sizes were selected because they are the most common block sizes used by database applications. On each test, 100,000 seek and read operations were performed. The results have the average of 10 runs at each block size.

The operating system used for the tests was RedHat Linux version 7.1 using the unmodified RedHat supplied 2.4.2 Linux kernel. The 2.4.2 kernel supported all devices and protocols except the iSCSI protocol. The Cisco iSCSI driver, version 1.8.7, was used for the iSCSI tests. This driver was loaded as a module, so the Linux kernel did not have to be modified. The driver and the test program were the only software loaded on the system that was not part of the RedHat 7.1 distribution. No RedHat patches were applied to the system.

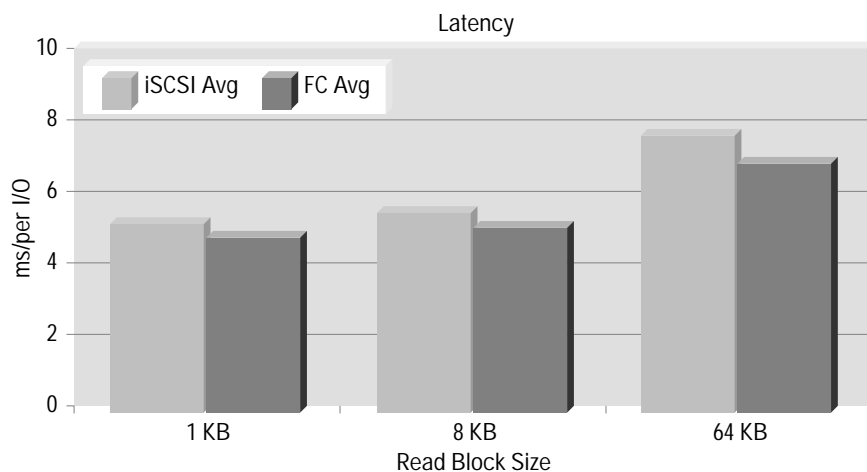
A single Seagate 15,000 RPM Fibre Channel disk drive was used for all the tests. A Vixel 2100 Fibre Channel hub was used between the Cisco SN 5420 and the Fibre Channel enclosure to convert the multimode fiber interface on the Cisco SN 5420 to the copper interface on the fibre-channel disk enclosure. Figure 1 shows the iSCSI setup with the Cisco SN 5420, and Figure 2 shows the Fibre Channel test setup.

Test Results

Chart 1 shows the average latency for each block size tested. The latency differences between the direct Fibre Channel-attached disk tests and the iSCSI tests using the Cisco SN 5420 are minimal. In all test cases, either technology could be used for most applications. In future iSCSI initiator and target products, Cisco expects that the latency will be reduced even further and any latency differences between Fibre Channel and iSCSI-attached storage will be insignificant.




Chart 1 shows iSCSI latency is a small portion of the overall latency.



Databases are one of the most latency-sensitive applications, and the primary measurement for disk performance of a database is the number of I/O per second that the disk system can sustain. Chart 2 shows the average number of I/O per second for both the Fibre Channel tests and the Cisco SN 5420 or iSCSI tests. Like the latency measurements, this chart shows that there is only a minimal reduction in I/O per second for each of the block sizes between the Cisco SN 5420 and the Fibre Channel-connected storage. The chart shows that the Cisco SN 5420 is well suited for database operations.

Chart 2 shows iSCSI is valid for consolidating storage and business continuance.



Summary

Cisco measured the latency of both iSCSI attached and Fibre Channel attached storage and found that the SN 5420 is suitable for traditional storage applications. The Cisco SN 5420 provides an extension of the existing Fibre Channel SAN using the iSCSI protocol. This extension allows for storage consolidation, remote data access, and the ability to do remote backup. For storage consolidation, the SN 5420 cost-effectively extends the benefits of a Fibre Channel SAN to servers throughout your enterprise using TCP/IP. All of these features of the SN 5420 allow for a lower cost of storage management without a significant impact to system performance or latency.



Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems Europe
11, Rue Camille Desmoulins
92782 Issy-les-Moulineaux
Cedex 9
France
www-europe.cisco.com
Tel: 33 1 58 04 60 00
Fax: 33 1 58 04 61 00

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems Australia, Pty., Ltd
Level 9, 80 Pacific Highway
P.O. Box 469
North Sydney
NSW 2060 Australia
www.cisco.com
Tel: +61 2 8448 7100
Fax: +61 2 9957 4350

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the

Cisco Web site at www.cisco.com/go/offices

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia
Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia
Ireland • Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru
Philippines • Poland • Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa
Spain • Sweden • Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

Copyright © 2001, Cisco Systems, Inc. All rights reserved. Printed in the USA. AccessPath, AtmDirector, Browse with Me, CCIP, CCSI, CD-PAC, *CiscoLink*, the *Cisco Powered* Network logo, Cisco Systems Networking Academy, the Cisco Systems Networking Academy logo, Fast Step, Follow Me Browsing, FormShare, FrameShare, GigaStack, IGX, Internet Quotient, IP/VC, iQ Breakthrough, iQ Expertise, iQ FastTrack, the iQ logo, iQ Net Readiness Scorecard, MGX, the Networkers logo, *Packet*, RateMUX, ScriptBuilder, ScriptShare, SlideCast, SMARTnet, TransPath, Unity, Voice LAN, Wavelength Router, and WebViewer are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, Discover All That's Possible, and Empowering the Internet Generation are service marks of Cisco Systems, Inc.; and Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, the Cisco IOS logo, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Enterprise/Solver, EtherChannel, EtherSwitch, FastHub, FastSwitch, IOS, IP/TV, LightStream, MICA, Network Registrar, PIX, Post-Routing, Pre-Routing, Registrar, StrataView Plus, Stratm, SwitchProbe, TeleRouter, and VCO are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

All other trademarks mentioned in this document or Web site are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0108R) 09/01 LW2678