



PS Series Storage Arrays Delivering High Performance iSCSI SANs

Abstract

This Technical Report analyzes the performance and scalability of SAN storage devices and compares EqualLogic iSCSI products with Fibre Channel products.

Copyright © 2005 EqualLogic, Inc.

July 2005

EqualLogic is a registered trademark of EqualLogic, Inc.

All trademarks and registered trademarks mentioned herein are the property of their respective owners.

Possession, use, or copying of the documentation or the software described in this publication is authorized only under the license agreement specified herein.

EqualLogic, Inc. will not be held liable for technical or editorial errors or omissions contained herein. The information in this document is subject to change. Your performance can vary.

PS Series Firmware Version 2.1 or later.

Table of Contents

Introduction	1
Tools for Testing	1
Cached Reads Per Second	2
Cached Reads Throughput	3
Summary	4
References	4
More Information and Customer Support	5

Introduction

A Storage Area Network (SAN) encompasses the network infrastructure and storage devices that provide storage services to host systems. This Technical Report analyzes the performance and scalability of SAN storage devices and compares EqualLogic iSCSI products with Fibre Channel products, demonstrating that iSCSI performs as well as—or better than—Fibre Channel on comparable platforms. In addition, the low cost and the ease with which an iSCSI SAN can be managed makes iSCSI superior to Fibre Channel when considering the total cost of ownership.

Tools for Testing

Scaling network infrastructure is well understood in the industry. An infrastructure that includes a switched network and multiple links on host systems and storage devices allows you to scale network performance efficiently and predictably—while also providing redundancy.

In contrast, performance analysis of SAN storage devices is a complex subject. These devices contain storage network ports and disk drives, in addition to protocol processing, caching, and RAID capabilities. Depending on the workload, any of these components can represent a bottleneck at a given time.

Simple load test tools such as ioMeter are popular for analyzing storage device performance. When tests are carefully performed using these tools, the results can predict the performance of complex applications. In addition, these tools can be used to perform stress testing on selected storage device subsystems.

The following load test parameters determine which storage device subsystem to test:

- Block number range. Is the range of block numbers accessed in the test so small that all the data fits in the storage device cache, or is the range much larger than the cache?
- I/O size and pattern. Does the test do large I/O requests (for example, 128 KB) to sequential blocks or Logical Block Addresses (LBAs), or does the test do small I/O requests (for example, 512 bytes) to random LBAs?

As shown in the following table, the two load test parameters produce four distinctive tests, each of which tests a particular subsystem or set of components within the storage device.

Load Test	Component Tested
Long sequential I/O, large LBA range	Disk I/O channels and network ports.
Short random I/O, large LBA range	Disk head actuators, disk rotation speed, and I/O scheduling.
Long I/O, small LBA range	Network ports.
Short I/O, small LBA range	Protocol processing, controller CPU performance, and cache efficiency.

For most common applications, a combination of tests provides more useful results than a single test. In particular, small range (cached I/O) performance may not be a good predictor of overall performance because applications usually process more data than can be held in the cache.

However, load tests are able to analyze the efficiency of individual SAN storage device components. This Technical Report will use the results of load tests to demonstrate that devices

based on the iSCSI protocol can perform as well as or better than Fibre Channel devices. Because the issue involves the storage network protocol and not the backend disk subsystem, the small range (cached I/O) tests were deemed appropriate to use for the analysis.

Network Appliance™ published a report in which cached read tests were used to compare the performance of the Network Appliance FAS product line with the EMC® CLARiiON® product line—all Fibre Channel based storage devices. Two following cases were considered:


- Small reads (to test cached I/Os per second)
- Large reads (to test storage network throughput)

This Technical Report will use the published Network Appliance test results and compare them to the test results for EqualLogic PS Series storage arrays to demonstrate that well-designed iSCSI storage devices can deliver excellent performance that is comparable to Fibre Channel.

Cached Reads Per Second

Figure 1 compares EMC CLARiiON storage systems and Network Appliance FAS storage systems (both of which use Fibre Channel as the storage interconnect) to EqualLogic PS Series storage arrays (using iSCSI as the storage interconnect). The reads are cached 512-byte reads per second.

Figure 1: Storage System Comparison with Cached 512-byte Reads




*Note that the value shown for the FAS270 is an approximation based on the FAS 270c (dual FAS270 cluster) results.

As shown in Figure 1, the PS Series arrays (red bars), using iSCSI as the storage interconnect, outperform comparable Network Appliance storage systems (blue bars) and EMC storage systems (yellow bars), both of which use Fibre Channel as the storage interconnect.

Cached Reads Throughput

Figure 2 compares EMC CLARiiON storage systems and Network Appliance FAS storage systems (both of which use Fibre Channel as the storage interconnect) to EqualLogic PS Series storage arrays (using iSCSI as the storage interconnect). The reads are cached 128-KB reads.

Figure 2: Storage System Comparison of Throughput with Cached Reads



*Note that the value shown for the FAS270 is an approximation based on the FAS 270c (dual FAS270 cluster) results.


For large cached-read tests, the expected outcome for an efficient storage network implementation that fully utilizes the internal system bandwidth is that the observed throughput roughly equals the bandwidth of the storage network ports. As shown in Figure 2, this is indeed the measured performance in many but not all cases.

You can also judge the effectiveness of the storage network implementation in a SAN storage array by comparing the measured throughput with the theoretical wire speed of the storage network ports. Figure 3 shows this analysis.

Note that “2 Gigabit Fibre Channel” capability actually provides a user data capacity of 1700 Mb/s (212.5 MB/s). Gigabit Ethernet has a user data capacity of 1000 Mb/s (125 MB/s). Theoretical wire speed for a SAN storage array is calculated by multiplying the port speed (212.5 or 125 MB/s) by the number of SAN network ports in the array. The number of ports on each product varies, depending on the design.

Network efficiency is calculated by dividing measured throughput (from Figure 2) by theoretical wire speed and expressing the results as a percentage. Note that the iSCSI-based EqualLogic PS Series storage arrays (red bars) deliver 94% network efficiency—the best results for any of the tested storage systems.

Figure 3: Throughput as a Fraction of Wire Speed (Network Efficiency)



*Note that the value shown for the FAS270 is an approximation based on the FAS 270c (dual FAS270 cluster) results.

Summary

Cached I/O tests comparing EMC CLARiiON and Network Appliance FAS storage systems to EqualLogic PS Series storage arrays demonstrate that iSCSI performs as well as—or better than—Fibre Channel on comparable platforms. The claim that iSCSI is inherently a low performance protocol is clearly a myth. Performance is influenced more by the storage system design than by a particular choice of network technology.

When the differences in cost and ease-of-use are taken into consideration, the iSCSI protocol is clearly superior to Fibre Channel.

References

The EqualLogic performance data was obtained from measurements taken in the EqualLogic development labs and from the PS Series Data Sheets available on the EqualLogic website:

<http://www.equallogic.com/pages/products.htm>

The Network Appliance performance data was obtained from the following location:

http://www.netapp.com/tech_library/3321.html

All CLARiiON performance data was obtained from the following location:

www.emc.com/products/systems/cx_compare_357.jsp

More Information and Customer Support

More information on performance, multipathing and other topics is available at:

http://www.equallogic.com/pages/Resource_Center.htm

Visit the EqualLogic Customer Support website, where you can download the latest documentation and firmware. You can also view FAQs, the Knowledge Base, and Technical Reports and submit a service request.

Visit the EqualLogic Customer Support website, where you can download the latest documentation and firmware. You can also view FAQs, the Knowledge Base

EqualLogic PS Series storage array documentation includes the following:

- *Release Notes*. Provides the latest information about PS Series storage arrays and groups.
- *QuickStart*. Describes how to set up the hardware and start using a PS Series storage array.
- *Group Administration*. Describes how to use the Group Manager GUI to manage a PS Series group. This manual provides comprehensive information about product concepts and procedures.
- *CLI Reference*. Describes how to use the Group Manager command line interface to manage a group and individual arrays.
- *Hardware Maintenance*. Provides information on maintaining the PS Series storage array hardware.

To access the Customer Support website, from the EqualLogic website (www.equallogic.com), click **Support** and log in to a support account. If you do not have an account, create one by clicking the link under the login prompt.

To contact customer support, send e-mail to supportnp@equallogic.com. If the issue is urgent, call 1-877-887-7337 to speak with a member of the customer support team.