



Enabling 8Gbps Fibre Channel End-to-End Performance

*By Darryl Gardner and Shaluka Perera
IBM Systems and Technology Group*

Executive Overview

In these economic times companies are leveraging hardware and software for all they're worth; IT spending is being scrutinized closely and executives are demanding ROI within the first 12 months of a project. With datacenters running out of physical space, power, and cooling capacity, consolidation is the not-so-new buzzword running throughout the IT organization. Management wants to see more servers in less rack space, using less power and able to do two to three times the work of their predecessors. From consolidation comes the latest advancement: *virtualization*. To take full advantage of consolidation you must virtualize. That is the best way to consolidate all those servers into less rack space while using less energy for power and cooling.

To make this new paradigm a reality in today's data centers, users are looking closely at server performance. In years past, most applications were given their own server, but in today's virtual environments, with many applications—and users—sharing one server, performance is now more critical than ever. Just as important as processor and memory performance is storage and I/O performance. After all, even the fastest processor in the world has to wait for data from I/O devices. The faster the I/O channels, the less waiting the processor has to do. This white paper discusses performance metrics as they pertain to host and storage I/O and IBM's new 8Gbps Fibre Channel end-to-end solution and how this new solution will address the needs of the data center today.

I/O Performance Metrics

It is important to understand what we mean by performance and how it is measured. In this case, we are discussing a complete solution: from a server with an 8Gbps Fibre Channel interface to an 8Gbps Fibre Channel storage area network (SAN) to 8Gbps Fibre Channel storage.

Host and storage I/O performance is traditionally measured in one of two ways: 1) bandwidth, or the size of the "pipe" used to transport data (defined as the capacity for data transfer (data rate or throughput) measured in bits per second); and 2) IOPS or Input/Output Operations per second. Bandwidth is the most commonly used definition for describing a host bus adapter's (HBA) capability, such as 2Gbps or 4Gbps (gigabits per second).

The theory behind this is that the bigger the pipe the faster you can transfer the data. However, this is not always the case. As a matter of fact, most non-virtualized users today run applications that barely require 2Gbps rates, meaning there is no need for 4Gbps HBAs—never mind 8Gbps. What we often forget to consider is the transactions driven by an application. If you have multiple transactions at various block sizes you are mostly concerned about latency and not bandwidth. If you are running in a virtualized server environment with multiple guests and multiple applications this should be your main concern. How many I/O operations per second can I push up and down this pipe to and from my storage? In a virtualized environment, with many users and applications, it becomes easy to overload even a 4Gbps solution.

IBM 8Gbps End-to-End Solution

IBM, together with partners, has designed an 8Gbps end-to-end solution geared toward the virtual operating environment, which is becoming prevalent in today's data centers. By increasing the number of IOPS from the server through the SAN to the storage, this configuration addresses the issue of multiple applications with high transaction operational needs by relying on server virtualization. Using an IBM System x3650 M2 x-series server with a Brocade 8Gb FC Single-port HBA for IBM System x connected to an IBM System Storage® SAN24B-4 Express fabric switch and two IBM DS5300 storage arrays, we were able to achieve more than 400,000 IOPS!

Configuration

- IBM x3650 M2, 1 Quad core 2.97GHz Intel® Xeon® 5500 Series (Nehalem) processor
 - 6GB system memory
 - Microsoft® Windows® Server 2008 Enterprise Server x64
- Brocade 8Gb FC Single-port HBA for IBM System x (46M6049)
- IBM System Storage SAN24B-4 Express fabric switch
- Two IBM System Storage DS5300 storage arrays
 - Four 8Gbps FC ports (1 per controller x 4 controllers)
 - Four to eight logical drives
 - All I/O was read from cache
- IOmeter version 2006.07.27 testing platform

Figure 1 illustrates the configuration.

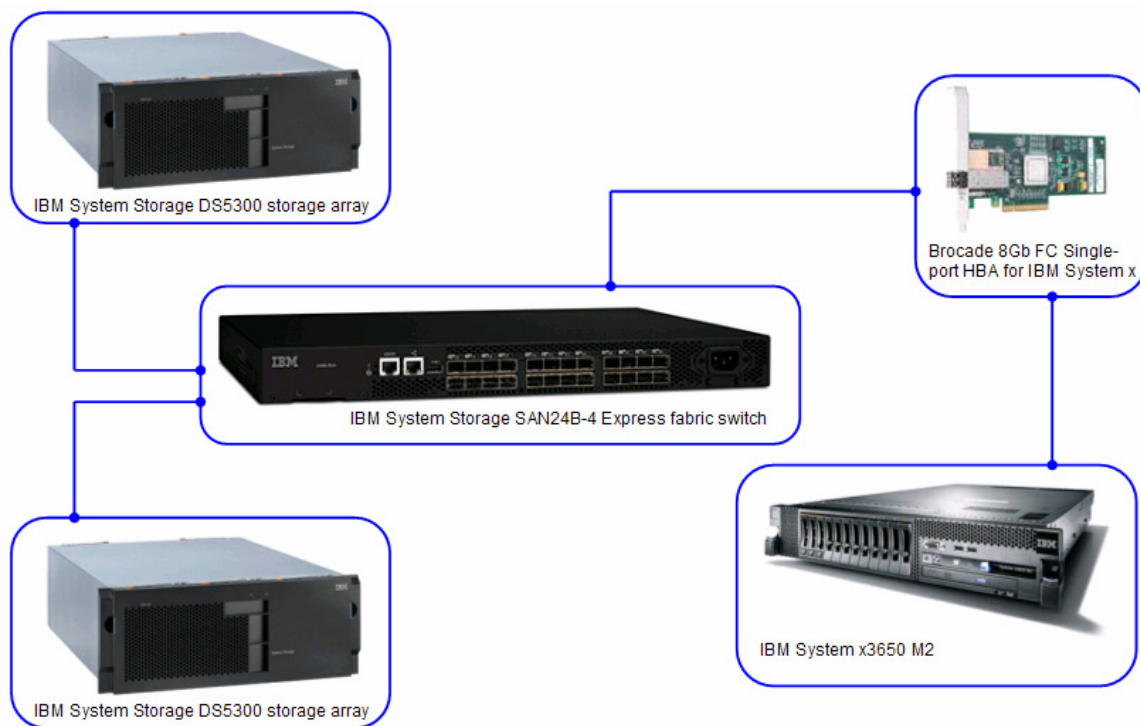


Figure 1. IBM 8GBps end-to-end solution

The Data — IOPS

The following raw numbers in *Table 1* were recorded using the configuration above.

Maximum IOPS			
Workload	Queue Depth	Transfer size	IOPS
50/50 R/W	32	512 bytes	401605
50/50 R/W	32	1KB	396112
50/50 R/W	32	2KB	412050
50/50 R/W	32	4KB	315386
50/50 R/W	32	8KB	176179
100% Write	32	512 bytes	276448
100% Read	16	512 bytes	474858

Table 1. IOPS performance

Figure 2 shows these results in graphic form.

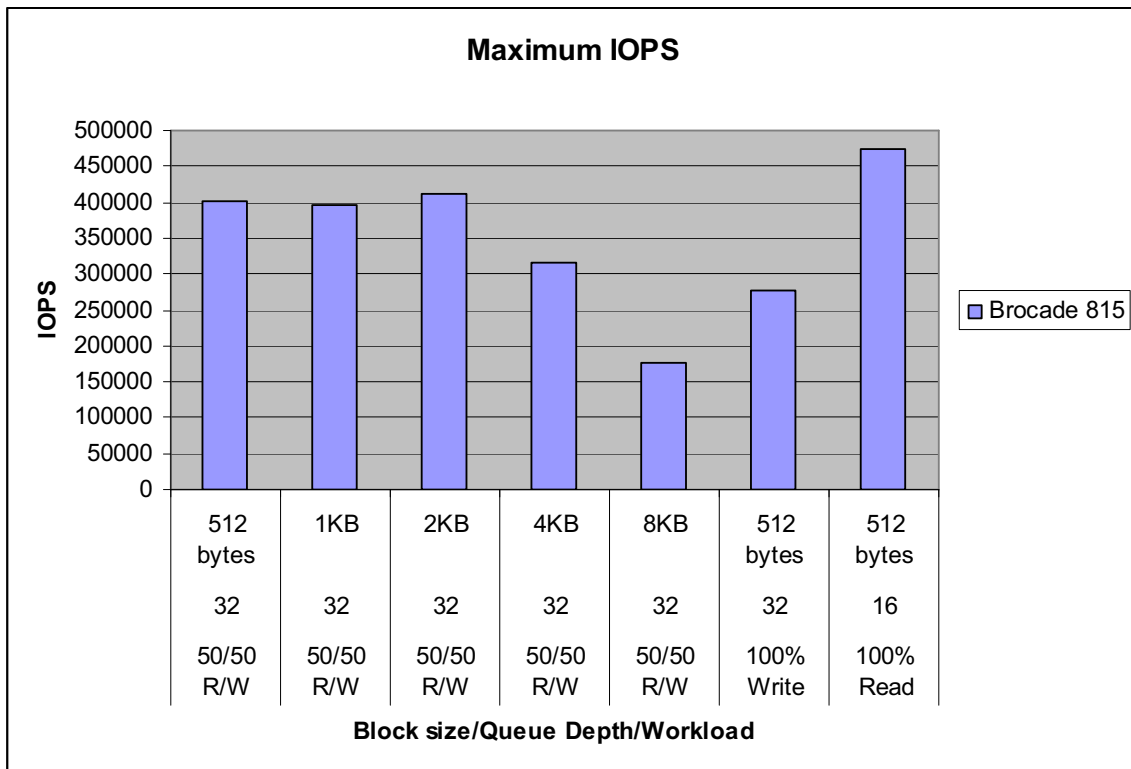


Figure 2. IOPS performance

The Data Throughput

The following raw numbers in *Table 2* were recorded using the same configuration.

Maximum Throughput					
Workload	Queue Depth	Transfer size	Read MBps	Write MBps	MBps
R/W Full duplex	160	256KB	770	764	1534
R/W Full duplex	32	256KB	786	751	1537

Table 2. Throughput performance

Figure 3 shows these results in graphic form.

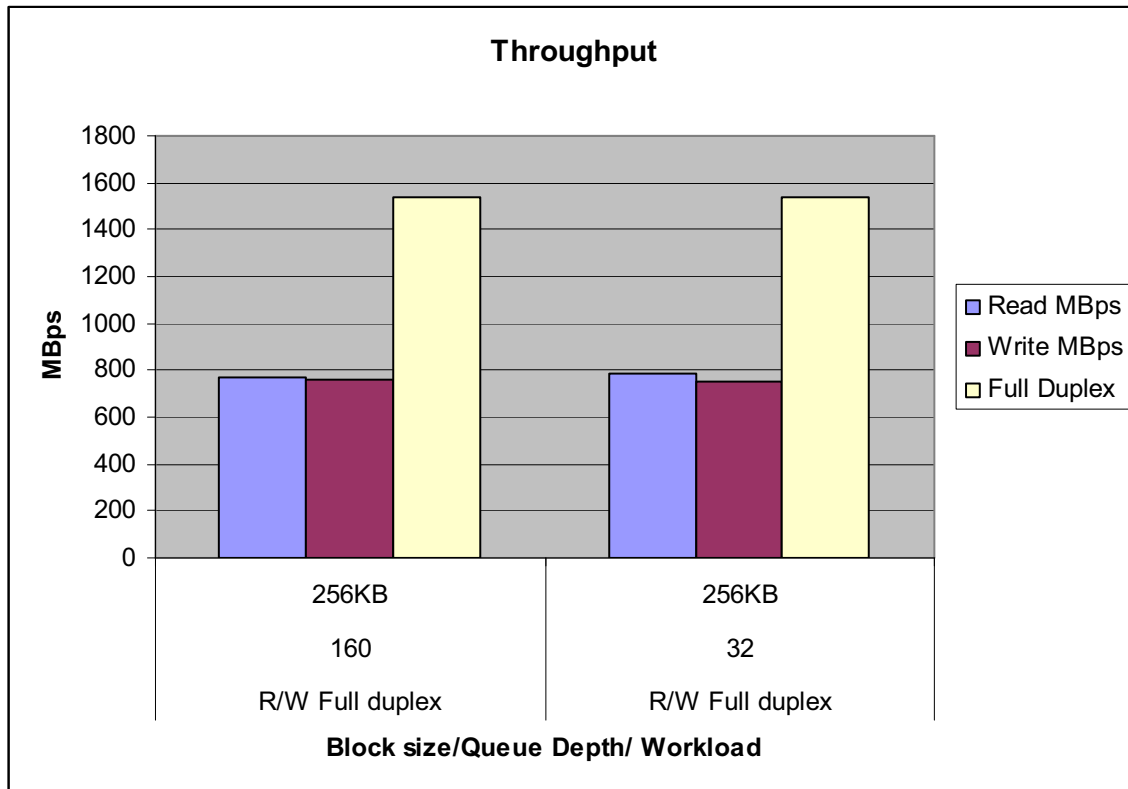


Figure 3. Throughput performance

Summary

From the first measurement of raw data, we have learned that at certain block sizes this 8Gbps Fibre Channel end to end solution can generate close to 500,000 IOPS. However, we know that in a real world environment, not many applications run with block sizes of 512 bytes. The point is that we are demonstrating the ability to address the high demands of transactional I/Os in any server environment running any OLTP applications. If you relate these results to real world applications such as Microsoft Exchange at 4KB block sizes (315,386 IOPS) or DB2 at 8KB block sizes (176,179 IOPS), these are impressive results. When configured in a virtual server environment, these results show great value in the ability to support a higher rate of transactional I/Os per second to/from the storage for any application.

The second measurement we performed was for the overall throughput. Again, if you are looking for a solution that runs at near wire speed, this is it. Where would this make sense? How about

Enabling 8Gbps Fibre Channel End-to-End Performance

Page 6

backups? There are no dedicated backup windows in today's enterprise data centers, therefore the more data you can stream down a pipe the better. Consider again a virtual server environment: 8 to 20 guest operating systems or virtual machines (VMs) all backing up through the same pipe at the same time. Using a solution that can support the required bandwidth to handle the backups simultaneously only makes sense.

Combine the power of the IBM System x3550 M2 and x3650 M2 servers with the performance of the Brocade 8Gb FC Single-port HBA, the low latency of the System Storage SAN24B-4 Express fabric switch and the IBM DS5300 storage array and you have a robust solution for any environment. Add to this solution the ability to isolate and prioritize the data flow from a virtual machine (VM) or any WWPN, through the SAN and to the storage without competing for available bandwidth and you have an *intelligent* robust solution. This type of solution, capable of accommodating wire speed data transfers to meet the backup needs of enterprise data centers and capable of handling a high rate of transactional I/O operations and isolating and prioritizing your data flow from the host to the storage, would be an attractive solution for any data center.



For More Information

IBM System x and xSeries Servers	ibm.com/systems/x
IBM System Storage	ibm.com/systems/storage
IBM BladeCenter Server and options	ibm.com/systems/bladecenter
IBM Standalone Solutions Configuration Tool (SSCT)	ibm.com/servers/eserver/xseries/library/configtools.html
IBM Electronic Service Agent	ibm.com/support/electronic
IBM ServerProven Program	ibm.com/servers/eserver/serverproven/compat/us
IBM Technical Support	ibm.com/server/support
IBM Configuration and Options Guide	ibm.com/servers/eserver/xseries/cog

Legal Information

© IBM Corporation 2009
IBM Systems and Technology Group
Dept. U2SA
3039 Cornwallis Road
Research Triangle Park, NC 27709

Produced in the USA
July 2009
All rights reserved.

For a copy of applicable product warranties, write to: Warranty Information, P.O. Box 12195, RTP, NC 27709, Attn: Dept. JDJA/B203. IBM makes no representation or warranty regarding third-party products or services including those designated as ServerProven or ClusterProven. Telephone support may be subject to additional charges. For onsite labor, IBM will attempt to diagnose and resolve the problem remotely before sending a technician.

IBM, the IBM logo, ibm.com, System Storage and System x are trademarks of IBM Corporation in the United States and/or other countries. For a list of additional IBM trademarks, please see <http://ibm.com/legal/copytrade.shtml>.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds.

Microsoft, Windows, and SQL Server are trademarks or registered trademarks of Microsoft Corporation.

Other company, product and service names may be trademarks or service marks of others.

IBM reserves the right to change specifications or other product information without notice. References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates. IBM PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions; therefore, this statement may not apply to you.

This publication may contain links to third party sites that are not under the control of or maintained by IBM. Access to any such third party site is at the user's own risk and IBM is not responsible for the accuracy or reliability of any information, data, opinions, advice or statements made on these sites. IBM provides these links merely as a convenience and the inclusion of such links does not imply an endorsement.

Information in this presentation concerning non-IBM products was obtained from the suppliers of these products, published announcement material or other publicly available sources. IBM has not tested these products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

MB, GB and TB = 1,000,000, 1,000,000,000 and 1,000,000,000,000 bytes, respectively, when referring to storage capacity. Accessible capacity is less; up to 3GB is used in service partition. Actual storage capacity will vary based upon many factors and may be less than stated.

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will depend on considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

Maximum internal hard disk and memory capacities may require the replacement of any standard hard drives and/or memory and the population of all hard disk bays and memory slots with the largest currently supported drives available. When referring to variable speed CD-ROMs, CD-Rs, CD-RWs and DVDs, actual playback speed will vary and is often less than the maximum possible.