QLogic Platform and Methodology for Product Evaluations

Delivering Transparent and Consistent Performance and Functional Test Results

Evaluate Adapter Capabilities and Performance Using Real-World Workloads

OVERVIEW

When building the infrastructure to support the enterprise data center, IT professionals conduct detailed, consistent assessments of product features and performance. The QLogic® Common Performance Test Platform can be used to validate performance benchmark results and provide IT administrators and system architects visibility into product performance in a clear and unbiased manner. QLogic performs testing to evaluate the adapter capabilities and the adapter's performance in real-world enterprise application environments. QLogic strives to be impartial, transparent, and consistent when testing the company's products and those of its competitors. This technology brief defines the QLogic common performance and functional evaluation platform and methodology.

QLOGIC'S APPROACH TO BENCHMARK TESTING

When properly structured, benchmark tests enable IT professionals to fairly compare I/O adapters under test. To be truly useful, benchmark testing should utilize the adapter under the kinds of workloads that are typically encountered in enterprise applications. Some benchmark results that appear to show only IOPS and use very small block sizes can be misleading and do not reflect a more complete set of real-world application requirements and results. This document presents benchmark best practices using commonly available platforms to drive I/O performance for often-used enterprise applications.

QLogic validates the capability of the I/O adapter performing the benchmark testing by placing a maximum load on the adapter. This testing method allows head-to-head performance evaluations and verifies the true capability of the Device Under Test (DUT). QLogic's lab environment was created to leverage a concentrated set of high-performance infrastructure. The DUT is driven to perform at its maximum capability to match or exceed the real-world stress test conditions. Thus, the results from this testing demonstrate real-world enterprise application performance of the DUT.

DEFINITION OF THE QLOGIC COMMON TEST PLATFORM

Servers – Minimum Configuration

- Intel® Xeon® E5-2600 8 core processors, 20MB cache (2.90 GHz)
- 32GB RAM (running at 1066MHz)
- Windows Server® 2008 R2 SP*x*[1]
- Red Hat® Enterprise Linux® 6.0*x*[1]

Emulated Storage Targets

- Targets are RAM based:
  - Fibre Channel and Fibre Channel over Ethernet (FCoE): RamSan® or SANBlaze®
  - iSCSI: Linux iSCSI Enterprise Target Daemon (IETD) using RAMDisk configuration
- Target throughput capability is set to exceed the initiator's maximum throughput
- Target IOPS performance is set to exceed the initiator's maximum IOPS performance

OEM-Equivalent Storage Array Platform

- Generic Fibre Channel storage arrays[1] (see actual test report for specific configuration information)
  - One target port and two target controllers
  - 10K and 15K revolutions per minute (RPM) drives

Switches

Sufficient ports to connect the initiator and targets to maximize performance

OEM-Equivalent Server Platform

- Dell® PowerEdge® R720
- HP® ProLiant® DL380p G8
- IBM® System x3650 M4
- Oracle® Sun Fire X4270 M3

1 See test report for latest OS, test platform, and firmware version numbers.

PLATFORM DEFINITION

The QLogic Common Performance Test Platform is designed to provide above-optimal capability to ensure there are no limitations encountered from the surrounding hardware test environment when testing for performance. Additionally, QLogic has created a set of detailed test plans to facilitate the validation of feature, function, and performance characteristics of the company's products and solutions. The QLogic Common Performance Test Platform uses the hardware and operating system (OS) configuration shown in the sidebar. Unless otherwise specifically mentioned, all products being tested are configured with their standard out-of-box settings, without tuning or changing the configurations for specific tests, which can potentially bias the results.

WORKLOAD DEFINITIONS

QLogic defines two workload environments for testing:

- I/O subsystem performance testing using emulated storage targets and test applications, such as IOmeter and Medusa, that are designed to create and consume storage protocol traffic
- Application workload testing using a storage array platform and test applications, such as ORION, JetStress, and transaction processing benchmarks, that mimic the operational environment of a specific class of applications

I/O SUBSYSTEM TESTING

During each test performed, QLogic gathers throughput (Mbps), IOPS, and CPU utilization data at the adapter DUT.

Three I/O workload types are used in testing: sequential read, sequential write, and random read-write mix. These workloads were selected to represent three common real-world deployments:

1. **Seq Read:** large-block for high-throughput applications
2. **Seq Write:** medium-block for applications using file transfer or day-today OS operations
3. **Sequential R/W:** small-block for IOPS-intensive applications

Details on the workload types are shown in the following table.

Typical Real-World SAN Workloads		
Workload Type	Block Size (KB)	Read:Write Ratio
Sequential Read	256 and greater	100%:0%
Sequential Write	64	0%:100%
Sequential Read/Write	4 – 32	50%:50%

Typical real-world deployments for the various block sizes are: database online transaction processing (OLTP); data warehousing, data recovery, and backups (use sequential reads); data backup, data de-duplication, file transfer or day-to-day OS operations (use sequential write workload types); and Microsoft Exchange Server® database OLTP transaction workloads (use random read-write).

Typical Real-World LAN Traffic	
MTU Size	1,500B
Direction	TX, RX, Bi-directional
Threads	Multiple

The test's objectives are selected to enable the best understanding of the adapter's performance and environmental usage with key indicators such as:

- Response time
- Throughput
- Resource utilization
- Maximum user load
- Business-related metrics

STORAGE PROTOCOL TESTING (FIBRE CHANNEL, FCOE, AND ISCSI)

Figure 1 shows the standard test configuration for storage protocol testing. The OSs used the following test tools:

- Windows®: IOmeter version 20*xx.xx.xx*[1]
- Linux: Medusa, v*x.x.xxx*[1]

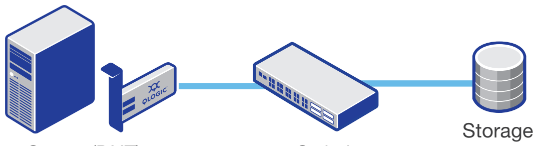

Figure 1. Standard Storage Test Configuration

ETHERNET TESTING (NIC)

Figure 2 shows the NIC testing setup, which includes the following:

- DUT server with the host NIC adapter under test
- Two client (non-DUT) servers using the adapters as NIC clients

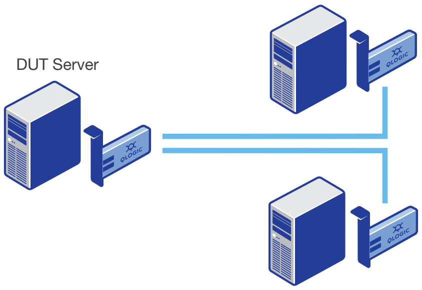

Figure 2. NIC Test Setup

The connections are made to mitigate any limitations from the NIC client, either with the server or adapter.

- A single port from each NIC client's adapter is connected to one of the DUT's NIC adapter ports
- An individual server and adapter using only a single port connected to each of the DUT adapter's ports provides sufficient traffic load to fully stress the dual-port capability of the DUT

A switch is not used in the configuration to ensure that performance of the adapter under test (DUT) is isolated. This method of evaluation testing gives the purest performance indicator of the DUT adapter's capabilities.

The test tools used by each OS are the following:

- Windows: IXIA® Endpoint v*x.xx*[1]
- Linux: IXIA Endpoint v*x.xx*[1]

1 See test report for latest OS, test platform, and firmware version numbers.

POWER MEASUREMENT TESTING

Getting an accurate power measurement of an adapter can be challenging. One approach is to measure the overall power at a system level. Another method is to measure power at a device level. The system level testing depends on the power supply of the host server, memory usage, loading of memory, the memory controller, and device driver use as well as power leakage. Accurately isolating the adapter's impact given the number of variables can be difficult. Therefore, QLogic's approach to power measurement is to create a setup that electrically isolates the adapter from the system's power, while still exercising a representative fullbandwidth load. This approach allows the adapter's power consumption to be measured in an idle and loaded state. The test setup allows power measurements to be made for the most transparent and accurate competitive comparisons between various adapters.

Figure 3 shows the test setup that uses a standup PCIe® header to isolate the adapter power from the system. The power to the adapter is supplied by separate 12V and 3.3V power supplies with a common reference to ground. The amp meter measures the current draw used to determine the power consumption of the device under test (DUT). The DUT ports are connected to a Fibre Channel Switch. The switch is zoned to separate the initiator ports and is connected to appropriately sized targets with enough LUNs to ensure the DUT is fully exercised to capacity. Test runs are performed under active and idle states to determine power consumption of the DUT.

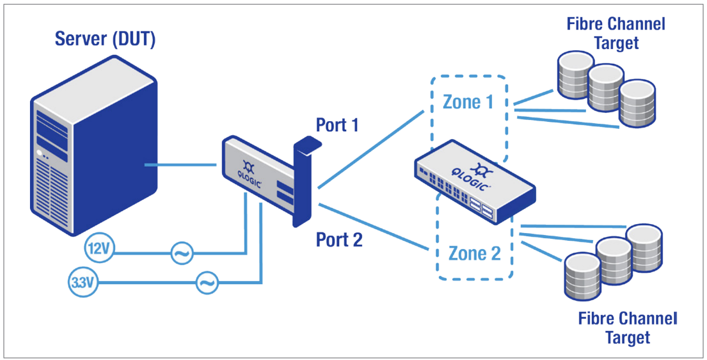

Figure 3. Setup for Power Measurement Testing

APPLICATION WORKLOAD TESTING

While I/O subsystem testing characterizes maximum performance for the I/O subsystem as a standalone element, QLogic uses specific, industry standard test applications—such as Oracle's ORION, Microsoft's JetStress, and transaction processing benchmarks—to test end-to-end server, network, and storage performance of applications as they operate in an infrastructure that includes QLogic's DUT. Tests of this type are designed around the concept of a "transaction," which may comprise multiple steps or work components, each of which could contain one or more I/Os.

For example, consider an order entry process for online shopping. On submission of the order, the database must update the following records: inventory, accounts receivable, customer profile, customer account, and other records. A "transaction" oriented system sets a transaction point before the first update and after the last one, so the transaction is an "all-or-nothing" operation update of the files. As such, the transaction may contain multiple I/O requests for the storage subsystem to process.

Oracle's ORION workload tool enables architects to effectively develop a workload that can mimic and stress a storage array in the same manner as applications designed with an Oracle backend database. Microsoft® provides the Jetstress 2010 tool for simulating Microsoft Exchange Server 2010 database I/O load and for validating hardware deployments against design targets. The Transaction Processing Performance Council defines transaction processing and database benchmarks that can be used to determine performance characteristics for various simulated OLTP environments.

Each test's objectives are selected to enable the best understanding of application performance and environmental usage under various user loads with key indicators such as:

- Transactions Per Second (TPS)
- Transaction Latency
- IOPS

DISTRIBUTED SHARED SERVER-BASED CACHE TESTING

Server-based caching is taking a leading role in accelerating application performance because it brings critical hot data inside the server and closest to the application. QLogic FabricCache™ 10000 Series 8Gb Fibre Channel Adapters are new, ground-breaking, enterprise-ready application performance acceleration solutions that bring shared, server-based caching to the SAN. QLogic's patent-pending FabricCache technology is exceptionally simple to deploy and manage, and transforms single-server, captive cache into a consolidated, shared, performance-enhancing resource across servers. The result is transparent, adapter-based caching that simplifies deployment, lowering total cost of ownership (TCO).

QLogic FabricCache 100000 Series Adapters are purposely designed to address the high I/O demands and distributed nature of clustered applications and virtualized environments. The adapters integrate a flash-based cache with a Fibre Channel Adapter that uses the existing SAN infrastructure to create a shared cache resource distributed over multiple servers.

To characterize the performance of the FabricCache solution, QLogic has performed a series of tests using both single-server and multipleserver environments. The FabricCache solution features a shared caching capability that enables each QLogic 10000 Series Adapter to share its caching resources with the other QLogic 10000 Series Adapters in the cluster. The adapters access the cache owner's resources remotely over the SAN. As such, QLogic has also provided a test that characterizes local caching versus remote caching SAN cache performance.

SINGLE-SERVER FABRICCACHE PERFORMANCE TESTING

The single-server test setup shown in Figure 4 was configured with an Intel Xeon server connected through the QLogic 10000 Series Adapter to one of the OEM-equivalent storage arrays, through a standard Fibre Channel Switch. Benchmark Factory™ for Databases TPC-E runs were executed with SSD cache levels varying in size from 0 to 100 percent of the actual SQL database size.

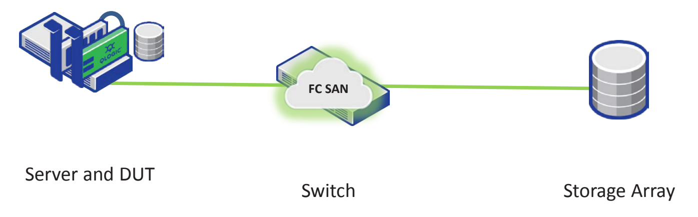

Figure 4. Single-Server FabricCache Testing Setup

MULTISERVER (4) FABRICCACHE PERFORMANCE TESTING

The four-server test setup shown in Figure 5 was configured with four Intel Xeon servers that were each connected with a QLogic 10000 Series Adapter to one of the OEM-equivalent storage arrays, through the same standard Fibre Channel Switch. Benchmark Factory for Databases TPC-E workload queries were executed in parallel with SSD cache levels varying in size from 0 to 100 percent of the SQL database size. The simultaneous runs emulate a multiple application environment, where numerous servers are running I/O traffic to a single storage array.

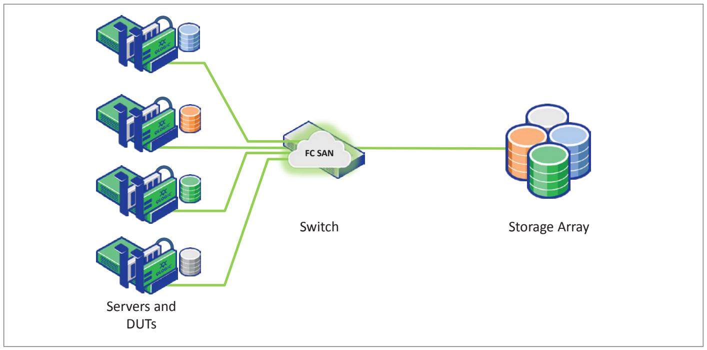

Figure 5. Multiserver (4) FabricCache Testing Setup

LOCAL CACHE VERSUS REMOTE CACHE EFFICIENCY OF FABRICCACHE REMOTE ACCESS

Figure 6 identifies the configuration for testing the transaction processing difference of local versus remote SAN cache (where a LUN cache is located in a FabricCache Adapter residing on another server). Again, Benchmark Factory for Databases TPC-E runs were executed with SSD cache levels varying in size from 0 to 100 percent of the SQL database size.

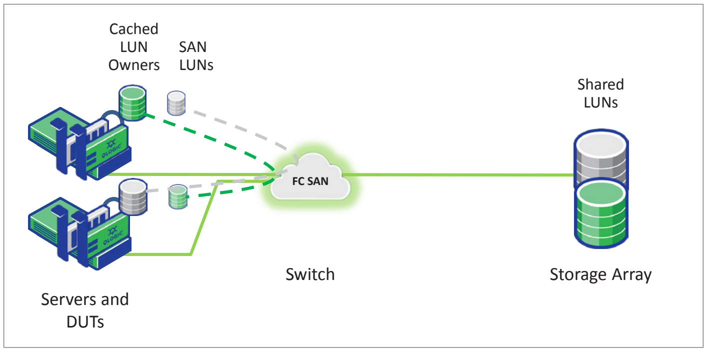

Figure 6. Local Versus Remote Caching Testing Setup

SUMMARY

The QLogic common performance and functional test platform can be used to validate benchmark results and provide IT administrators and system architects the ability to recreate and validate published results. QLogic delivers consistent, transparent test results. QLogic publishes data from its tests in several formats, including white papers, competitive briefs, deployment guides, and technology briefs. These materials can be used to evaluate QLogic and competitors' products and guide your purchasing choices for IT infrastructure.

TRUSTED SOLUTIONS

QLogic is a global leader and technology innovator in high-performance server and storage networking connectivity and application acceleration solutions. The company's leadership in product design and maturity of software stack make it the top choice of leading OEMs, including Cisco, Dell, EMC, Hitachi Data Systems, HP, IBM, Lenovo, NetApp, and Oracle, as well as channel partners worldwide for their virtualized, converged, and cloud environment solutions.

LEARN MORE

QLogic offers complete solutions to some of the most complex issues facing the data center. To read more, go to www.qlogic.com, click the **Solutions** tab, and then select **QLogic Solutions**.

QLogic has created online demonstrations that showcase QLogic technologies, such as NIC Partitioning and the application of quality of service (QoS) in a VMware® ESXi v5.x environment using a QLogic QLE8242 Converged Network Adapter. In addition, demonstrations of technology performance or proof of concepts can be configured for your specific environment. To register for a demonstration of QLogic product capabilities, visit http://www.qlogic.com/solutions/Pages/qos.aspx, and then click **Schedule** a future 'live' demo.

ABOUT CAVIUM

Cavium, Inc. (NASDAQ: CAVM), offers a broad portfolio of infrastructure solutions for compute, security, storage, switching, connectivity and baseband processing. Cavium's highly integrated multi-core SoC products deliver software compatible solutions across low to high performance points enabling secure and intelligent functionality in Enterprise, Data Center and Service Provider Equipment. Cavium processors and solutions are supported by an extensive ecosystem of operating systems, tools, application stacks, hardware reference designs and other products. Cavium is headquartered in San Jose, CA with design centers in California, Massachusetts, India, Israel, China and Taiwan.

Follow us:

Corporate Headquarters Cavium, Inc. 2315 N. First Street San Jose, CA 95131 408-943-7100

International Offices UK | Ireland | Germany | France | India | Japan | China | Hong Kong | Singapore | Taiwan | Israel