Marvell® Platform and Methodology for Product Evaluations

Delivering Transparent and Consistent Performance and Functional Test Results



- Marvell's **QLogic FC HBAs**
- **FastLinQ Ethernet Adapters**

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 Marvell 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.

Marvell performs testing to evaluate its QLogic® FC HBAs (Host Bus Adapters) and FastLinQ Ethernet adapters from Marvell capabilities and performance in real-world enterprise application environments. Marvell strives to be impartial, transparent, and consistent when testing its adapter products and those of its competitors. This technology brief defines the Marvell common performance and functional evaluation platform and methodology.

MARVELL'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 realworld application requirements and results. This document presents benchmark best practices using commonly available platforms to drive I/O performance for often-used enterprise applications.

Marvell validates the capability of its QLogic FC HBAs and FastLinQ Ethernet 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). Marvell'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 MARVELL 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 SPx[1]
- Red Hat® Enterprise Linux® 6.0x[1]

SERVERS – MINIMUM CONFIGURATION

- 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
- HPE® 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 Marvell 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, Marvell 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 Marvell 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

The Marvell test defines two workload environments:

- 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

Each test 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:

- **Seq Read:** large-block for high-throughput applications
- **Seq Write:** medium-block for applications using file transfer or daytoday OS operations
- **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 20xx.xx.xx[1]
- Linux: Medusa, vx.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

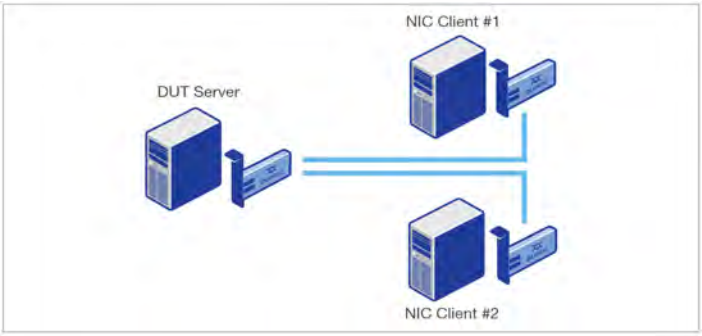


Figure 2. NIC Test Setup

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 vx.xx[1]
- Linux: IXIA Endpoint vx.xx[1]

ETHERNET TESTING (NIC)

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, the Marvell 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.

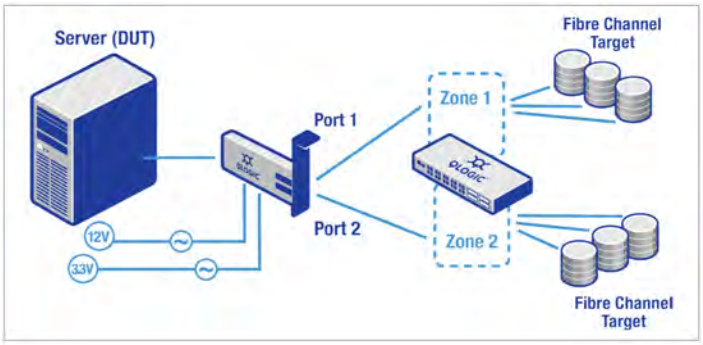


Figure 3. Setup for Power Measurement Testing

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

APPLICATION WORKLOAD TESTING

While I/O subsystem testing characterizes maximum performance for the I/O subsystem as a standalone element, this test 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 Marvell'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 "allor- 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

SUMMARY

The Marvell 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. Marvell delivers consistent, transparent test results. Marvell 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 Marvell and competitors' products and guide your purchasing choices for IT infrastructure.

TRUSTED SOLUTIONS

Marvell 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.

MARVELL

To deliver the data infrastructure technology that connects the world, we're building solutions on the most powerful foundation: our partnerships with our customers. Trusted by the world's leading technology companies for 25 years, we move, store, process and secure the world's data with semiconductor solutions designed for our customers' current needs and future ambitions. Through a process of deep collaboration and transparency, we're ultimately changing the way tomorrow's enterprise, cloud, automotive, and carrier architectures transform—for the better.

Copyright © 2020-2021 Marvell. All rights reserved. Marvell and the Marvell logo are trademarks of Marvell or its affiliates. Please visit www.marvell.com for a complete list of Marvell trademarks. Other names and brands may be claimed as the property of others.