



Black Book

Edition 10

Converged Network Adapters

CONVERGED NETWORK ADAPTERS (CNA)

Your feedback is welcome

Our goal in the preparation of this Black Book was to create high-value, high-quality content. Your feedback is an important ingredient that will help guide our future books.

If you have any comments regarding how we could improve the quality of this book, or suggestions for topics to be included in future Black Books, contact us at ProductMgmtBooklets@ixiacom.com.

Your feedback is greatly appreciated!

Copyright © 2014 Ixia. All rights reserved.

This publication may not be copied, in whole or in part, without Ixia's consent.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the U.S. Government is subject to the restrictions set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013 and FAR 52.227-19.

Ixia, the Ixia logo, and all Ixia brand names and product names in this document are either trademarks or registered trademarks of Ixia in the United States and/or other countries. All other trademarks belong to their respective owners. The information herein is furnished for informational use only, is subject to change by Ixia without notice, and should not be construed as a commitment by Ixia. Ixia assumes no responsibility or liability for any errors or inaccuracies contained in this publication.

Contents

How to Read this Book.....	vii
Dear Reader	viii
Introduction	1
Guide to Converged Network Adapter Technology	1
10 Gigabit Ethernet: The Foundation for Network Convergence	5
Ixia's IxChariot – Testing Performance from the User's Perspective	13
Test Case: TCP Throughput	19
Test Case: UDP Throughput.....	29
Test Case: Latency	33
Test Case: Storage I/O Performance	39
Test Case: Virtualized Environments	45
Appendix A: Emulex OneConnect UCNA Platform.....	49
Contact Ixia.....	57

How to Read this Book

The book is structured as several standalone sections that discuss test methodologies by type. Every section starts by introducing the reader to relevant information from a technology and testing perspective.

Each test case has the following organization structure:

Overview	Provides background information specific to the test case.
Objective	Describes the goal of the test.
Setup	An illustration of the test configuration highlighting the test ports, simulated elements and other details.
Step-by-Step Instructions	Detailed configuration procedures using Ixia test equipment and applications.
Test Variables	A summary of the key test parameters that affect the test's performance and scale. These can be modified to construct other tests.
Results Analysis	Provides the background useful for test result analysis, explaining the metrics and providing examples of expected results.
Troubleshooting and Diagnostics	Provides guidance on how to troubleshoot common issues.
Conclusions	Summarizes the result of the test.

Typographic Conventions

In this document, the following conventions are used to indicate items that are selected or typed by you:

- **Bold** items are those that you select or click on. It is also used to indicate text found on the current GUI screen.
- *Italicized* items are those that you type.

CONVERGED NETWORK ADAPTERS (CNA)

Dear Reader

Ixia's Black Books include a number of IP and wireless test methodologies that will help you become familiar with new technologies and the key testing issues associated with them.

The Black Books can be considered primers on technology and testing. They include test methodologies that can be used to verify device and system functionality and performance. The methodologies are universally applicable to any test equipment. Step by step instructions using Ixia's test platform and applications are used to demonstrate the test methodology.

This tenth edition of the black books includes twenty two volumes covering some key technologies and test methodologies:

Volume 1 – Higher Speed Ethernet

Volume 2 – QoS Validation

Volume 3 – Advanced MPLS

Volume 4 – LTE Evolved Packet Core

Volume 5 – Application Delivery

Volume 6 – Voice over IP

Volume 7 – Converged Data Center

Volume 8 – Test Automation

Volume 9 – Converged Network Adapters

Volume 10 – Carrier Ethernet

Volume 11 – Ethernet Synchronization

Volume 12 – IPv6 Transition Technologies

Volume 13 – Video over IP

Volume 14 – Network Security

Volume 15 – MPLS-TP

Volume 16 – Ultra Low Latency (ULL) Testing

Volume 17 – Impairments

Volume 18 – LTE Access

Volume 19 – 802.11ac Wi-Fi Benchmarking

Volume 20 – SDN/OpenFlow

Volume 21 – Network Convergence Testing

Volume 22 – Testing Contact Centers

A soft copy of each of the chapters of the books and the associated test configurations are available on Ixia's Black Book website at <http://www.ixiacom.com/blackbook>. Registration is required to access this section of the Web site.

At Ixia, we know that the networking industry is constantly moving; we aim to be your technology partner through these ebbs and flows. We hope this Black Book series provides valuable insight into the evolution of our industry as it applies to test and measurement. Keep testing hard.



Errol Ginsberg, Acting CEO

Converged Network Adapters

Test Methodologies

This document presents a thorough methodology for testing the performance characteristics of a 10 Gbps converged network adapter (CNA) in a variety of operating environments.

Introduction

In this Black Book on testing converged network adapters (CNAs), Ixia and Emulex have partnered to produce a definitive guide to the technological characteristics of CNAs, converged Ethernet technology, and the methodology for testing these new devices. The introduction provides a detailed background on FCoE and datacenter network convergence as well as a look at the IxChariot test application. The latter half of the book is devoted to a detailed test plan that will guide the reader through the key tests required to thoroughly measure the performance of a modern CNA.

Guide to Converged Network Adapter Technology

Data center objectives have always worked to improve their operational efficiency. Whether they're concerned with servers, storage or the network, their goals are clearly focused on improving the performance and manageability of operations in a highly cost-efficient manner. Data center managers are often drowning in a sprawl of network infrastructure, looking for solutions that will allow them to transition to a more dynamically provisioned network that addresses the quality and service level requirements of business applications.

The evolution of the converged network has empowered today's data center to deploy an adaptive, next-generation network. Network convergence allows companies to:

1. **Consolidate** multiple low-bandwidth links into a faster, high-bandwidth infrastructure that significantly reduces the number of switch and adapter ports and cables.
2. **Converge** networking, storage and clustering traffic into a single network, eliminating the need for dedicated networks for each traffic type, reducing the number of network ports and cables, thereby simplifying deployment and management.
3. **Virtualize** the underlying physical network infrastructure, providing service level guarantees for each type of traffic. In addition, the solution must be capable of responding to dynamic changes in network services based on the business demands of data center applications.

CNAs were introduced in 2008 as the critical enabling technology for a converged network. First-generation CNAs leverage Ethernet enhancements and fibre channel over Ethernet (FCoE) technology to provide convergence over a 10GbE infrastructure. CNAs combine the network interface card (NIC) functionality and the fibre channel HBA functionality into a single adapter.

The universal converged network adapter (UCNA) was introduced by Emulex in October, 2009. Unlike first-generation CNAs that only provide FCoE convergence, UCNAs offer high-performance offload for TCP/IP, iSCSI and FCoE protocols. They operate at 10 Gbps on a single adapter using a single network infrastructure. The Emulex OneConnect UCNA was tested with the IxChariot test application, measuring CNA performance. The detailed test plans beginning on page 13. Data Center Challenges that Drive Converged Networking

CONVERGED NETWORK ADAPTERS (CNA)

Data centers have traditionally used specialized networks to meet individual I/O connectivity requirements for networking and storage. With increasing deployments of blades and virtualized servers, data centers are facing network sprawl, primarily driven by deployments of multiple 1Gb Ethernet links. This sprawl has resulted in:

- Increased capital costs for adapters, switch ports and cables
- Increased operational costs for power, cooling and IT management

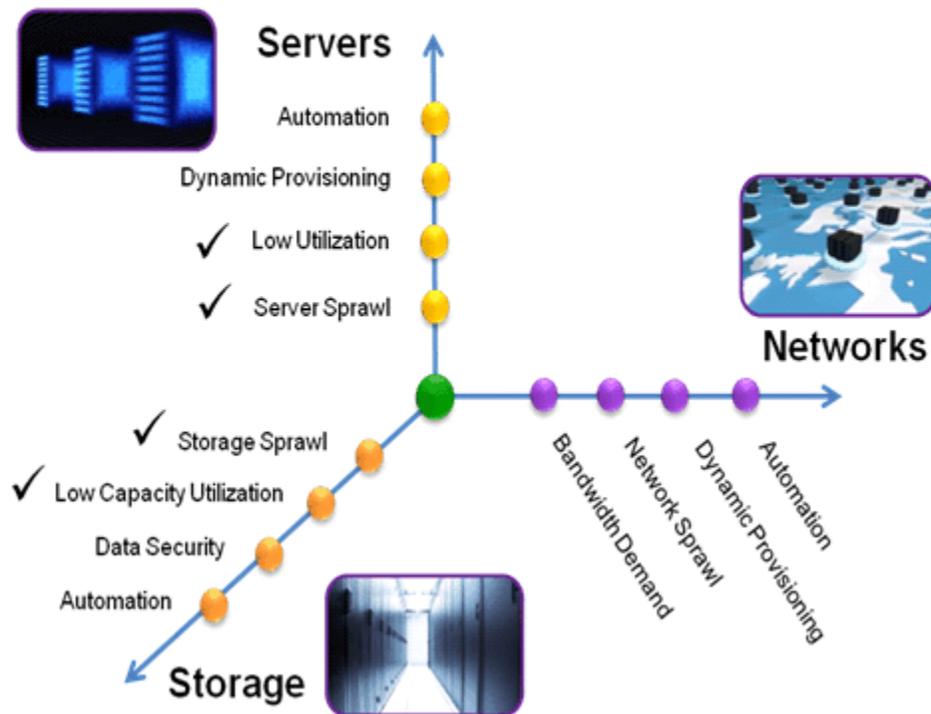


Figure 1. Data center challenges across IT segments.

While server consolidation initiatives have enabled higher efficiencies and agility in the computing infrastructure, the overall networking infrastructure has not kept pace with the changing dynamics. With virtualization, data centers are able to achieve higher levels of utilization, but this is also creating the need for higher bandwidth. Networks are still managed as individual silos, where storage and networking traffic are each carried over a dedicated infrastructure, as illustrated in Figure 2.

CONVERGED NETWORK ADAPTERS (CNA)

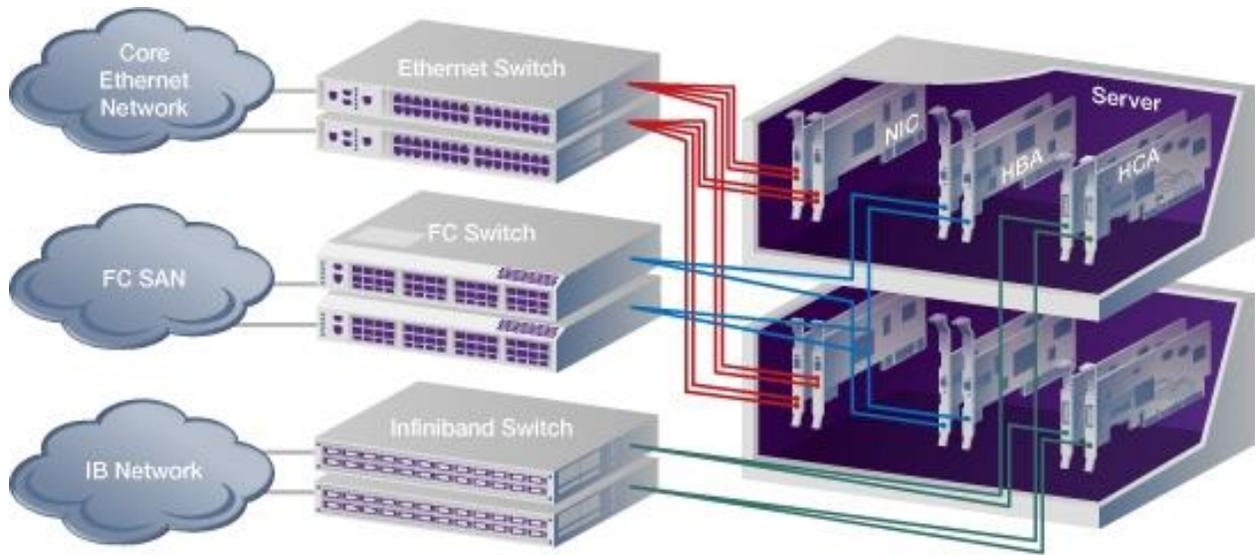


Figure 2. Traditional dedicated networks for SAN and LAN

Consolidation of these multiple networks into a common infrastructure that can be shared by multiple traffic types—the converged network—helps to overcome the challenges facing emerging networks. Once traffic is converged on a link, network links can be dynamically provisioned based on the application or business requirements, facilitating a highly responsive IT infrastructure.

10 Gigabit Ethernet: The Foundation for Network Convergence

Ethernet innovations such as lossless transmissions, efficient physical media and transceiver optimizations make 10 Gigabit Ethernet the technology of choice for converging data center networks.

The 10 Gigabit Ethernet networking standard (10GbE), ratified in 2002, initially found traction in specific applications such as inter-switch connectivity and high performance applications. The Ethernet networking technology has since witnessed several key innovations:

- Enhanced physical media
 - 10Gb/s connectivity over UTP cabling
 - 10Gb/s connectivity over direct attach twin-ax copper cabling
- Optimizations in 10Gb/s transceiver technology with the SFP+ form factor
- Support for lossless Ethernet
- New physical network designs such as top-of-rack switch architectures

A converged network based on 10GbE fully complements the data center consolidation efforts and improves the efficiency of overall operations. Leveraging 10GbE to carry data networking and storage traffic simplifies network infrastructure and reduces the number of cables, switch ports and adapters while lowering overall power, cooling and space requirements.

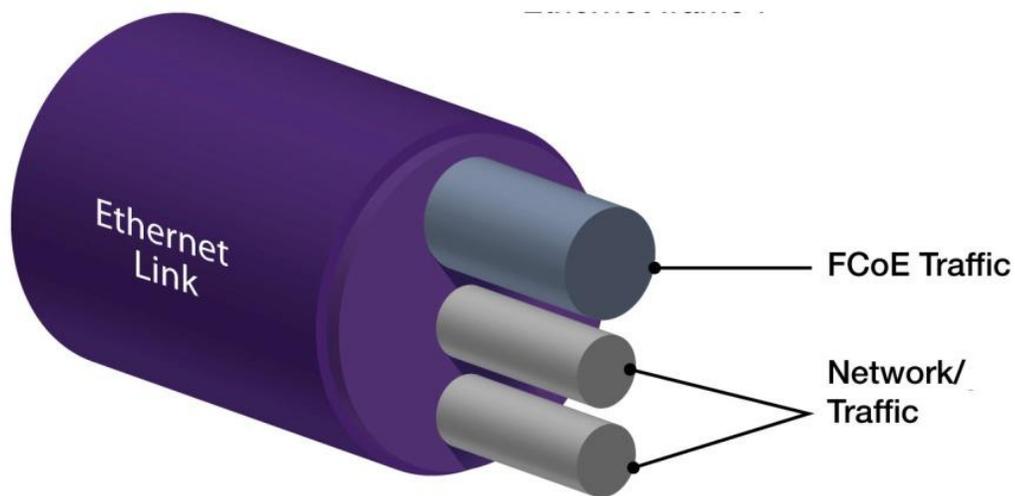


Figure 3. 10GbE enables multiple traffic type over single link

In addition to lowering costs, 10GbE enables much-needed scalability by providing additional network bandwidth. 10GbE also simplifies management by reducing the number of ports and facilitating flexible bandwidth assignments for individual traffic types.

Network Convergence with Lossless 10 Gigabit Ethernet

In order to facilitate network convergence and carry fibre channel traffic over 10GbE, Ethernet technology had to support “no-drop” or “lossless” behavior. Since its introduction, Ethernet has steadily evolved to keep pace with these changing requirements. The recent innovations that support lossless characteristic in 10GbE are the ability to:

- Isolate and prioritize different traffic types using priority flow control (PFC)
- Maintain bandwidth guarantees for multiple traffic types
- Assure that end-points and switches know about each other’s capabilities through an enhanced management protocol

In parallel with the emergence of lossless 10GbE, the emergence of newer standards such as FCoE are accelerating the adoption of Ethernet as the medium of network convergence.

Ethernet Enhancements Overview

Efforts are well underway with the Data Center Bridging Task Group of the IEEE 802.1 Working Group (LANs) to provide the necessary framework to enable 10GbE converged networking within a data center. Four key protocols are being defined or enhanced as shown in Table 1 .

Table 1. Protocol standards are enabling convergence

Protocols	Key Functionality	Business Value
Priority Flow Control (PFC) P802.1Qbb	Management of bursty, single traffic sources on a multi-protocol link.	Enables storage traffic over 10GbE link with “no-drop” in the network.
Enhanced Transmission Selection (ETS) P802.1Qaz	Bandwidth management between traffic types for multi-protocol links.	Enables bandwidth assignments per traffic type. Bandwidth is configurable on-demand.
Data Center Bridging Capabilities Exchange Protocol (DCBCXP) 802.1Qaz	Auto exchange of Ethernet parameters between peers (switch to NIC, switch to switch).	Facilitates interoperability by exchanging capabilities supported across the nodes.
Congestion Management (CM) P802.1Qau	Addresses problem of sustained congestion, driving corrective action to the edge.	Facilitates larger end-to- end deployment of network convergence.

Network Convergence with iSCSI and 10 Gigabit Ethernet

The iSCSI protocol, ratified by the Internet Engineering Task Force (IETF) in 2003, brought storage area networks within the reach of small and mid-sized businesses. The protocol encapsulates native SCSI commands over TCP/IP and transmits the packets over an Ethernet network infrastructure. The emergence of 10GbE addressed IT managers’ concerns regarding bandwidth and latency issues of 1 GbE and laid the foundation for more widespread adoption of iSCSI-based network convergence.

First-generation CNAs: FCoE and NIC

First generation CNAs leverage Ethernet enhancements and FCoE technology to provide convergence over a 10GbE infrastructure. CNAs combine NIC functionality and the fibre channel HBA functionality into a single adapter.

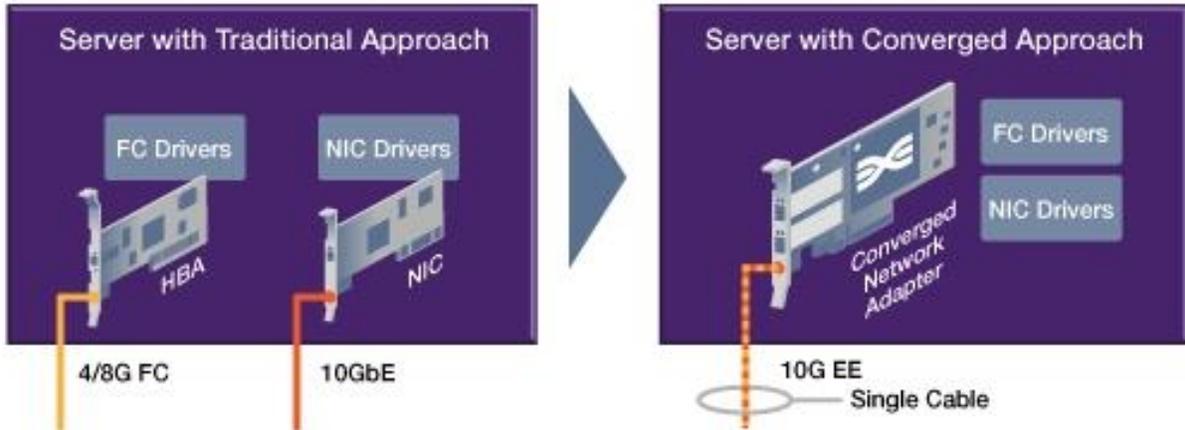


Figure 4. FCoE CNAs are provisioned with both fibre channel and NIC drivers

FCoE-enabled Enterprise Network Convergence

FCoE is a standard under development by INCITS T11 that will leverage 10GbE enhanced features for I/O consolidation in the data center. FCoE expands fibre channel into the Ethernet environment, combining two leading technologies. It provides options for SAN connectivity and networking. An FCoE-enabled converged network represents the first deployment of network convergence, as illustrated in Figure 5.

10 GIGABIT ETHERNET: THE FOUNDATION FOR NETWORK CONVERGENCE

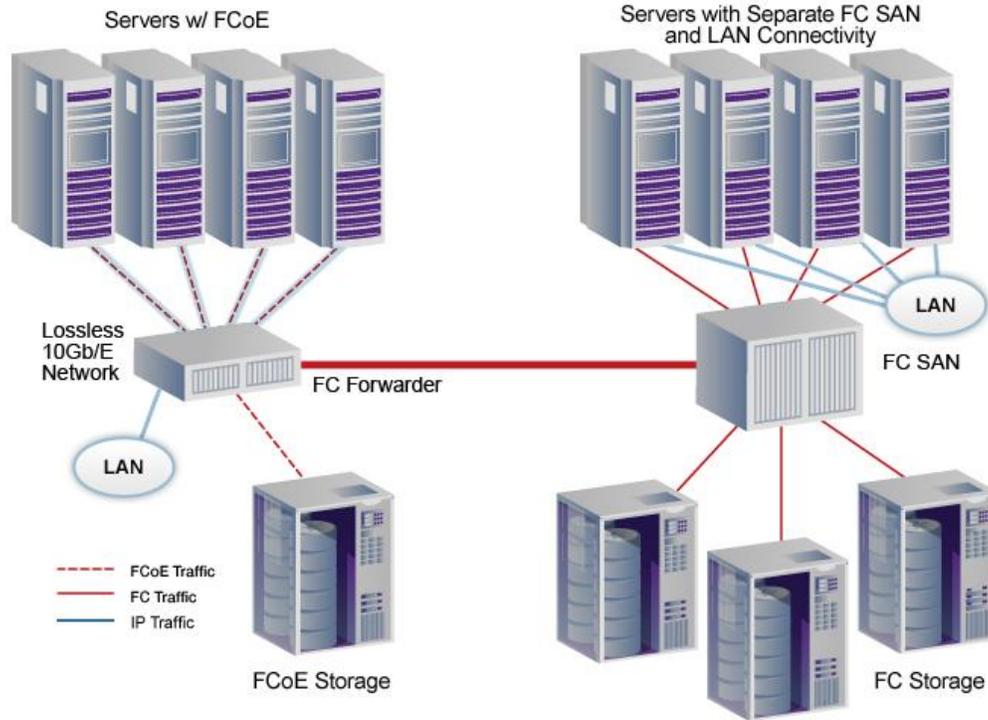


Figure 5. FCoE-enabled network convergence

FCoE is not a replacement for conventional fibre channel, but is an extension of fibre channel over a different link layer. A lossless Ethernet infrastructure that carries both fibre channel storage data as along with other data types allows simplified server connectivity, while retaining the performance and reliability required for storage transactions. Instead of provisioning a server with dual-redundant Ethernet and fibre channel ports—a total of 4 ports—blade servers can be configured with two lossless FCoE 10GbE ports. This reduction of interfaces greatly simplifies blade server deployment and ongoing cable management. The key virtue of FCoE is the streamlining of server connectivity using lossless Ethernet while retaining the channel characteristics of conventional fibre channel SANs.

Preserved Fibre Channel Characteristics

The FCoE protocol specification maps a complete fibre channel frame, including checksum and framing bits, directly into an Ethernet payload, avoiding the overhead of intermediate protocols.

This light weight encapsulation ensures that FCoE-capable Ethernet switches are not computationally intensive, allowing them to provide the high performance and low latencies of fibre channel networks. The technology fully leverages existing fibre channel constructs such as fabric login, zoning and LUN masking, while ensuring secure access to the networked storage.

N_Port ID Virtualization

N_Port ID virtualization (NPIV) is a fibre channel innovation that enables multiple addresses, referred to as virtual ports, to share a single fibre channel port when registering with the SAN

Fabric. This allows virtual machines to have individual dedicated virtual ports, limiting necessary storage access to only the required resources. Further, the ability to reinitiate a virtual port on a different server greatly enhances virtual machine mobility for load balancing, portability and disaster recovery. FCoE retains the use of NPIV to improve the flexibility and security of virtual server deployments.

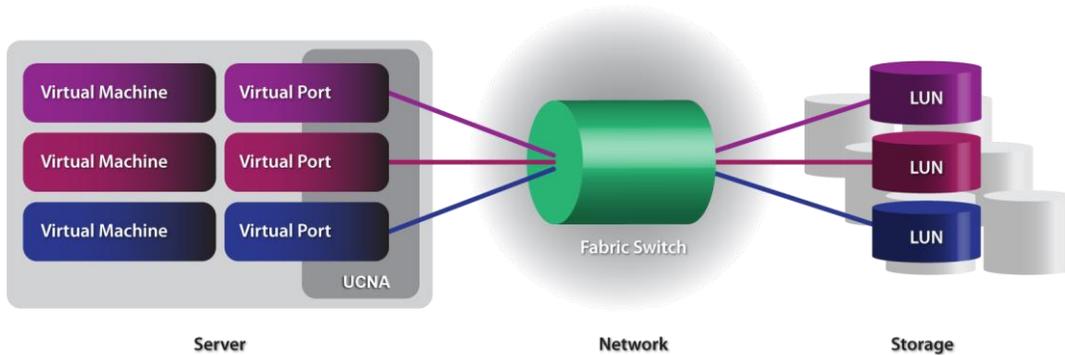


Figure 6. NPIV enables multiple addresses, referred to as virtual ports, to share a single fibre channel port when registering with the SAN Fabric.

Fibre Channel SAN Investments Protected

Isolation of storage devices in fibre channel SAN environments is provided through zoning and LUN masking techniques based on the address, known as the world wide port name, of the fibre channel port. In a virtualization environment, when multiple virtual machines share a single fibre channel port, the virtual machines are indistinguishable to the SAN fabric with respect to data isolation and quality of service (QoS) management.

One of the goals of FCoE, therefore, is to ensure protection of fibre channel investments through interoperability with existing infrastructure. The use of fibre channel as the upper layer protocol in FCoE ensures that the fibre channel SAN infrastructure can be seamlessly accessed from FCoE-enabled servers. This ensures that FCoE-enabled servers can use established fibre channel software and can transparently interoperate with existing fibre channel SAN infrastructure that includes switches, storage arrays, tape drives and data backup tools. This approach also acts as an efficient mechanism for extending fibre channel SAN infrastructure to Ethernet networks, thus expanding the accessibility and usage of existing storage investments.

Second-Generation Universal CNAs: FCoE, iSCSI, and NIC

Second generation Universal CNAs are emerging in the marketplace, providing multiple options for combining network and storage traffic over a 10GbE infrastructure. Emulex announced the first UCNA, the OneConnect UCNA platform, as a single chip design that transports all major protocols over Ethernet, including FCoE and iSCSI for block storage, NFS and CIFS for file storage, TCP/IP Local Area Network (LAN) and RDMA for clustering. For specific OneConnect

UCNA product information, refer to Appendix A: Emulex OneConnect UCNA Platform on page 49.

CNA adapter personalities allow high-performance server offloads to be used for FCoE or iSCSI connectivity. UCNA functionality is transparent to the operating system, since it presents both the NIC and iSCSI or fibre channel functionality to the OS. UCNAs offload fibre channel protocol processing from the CPU, thus providing high performance storage connectivity while improving overall CPU utilization.

Servers supporting PCI Express (x8 or higher) are capable of utilizing the full potential of dual-port 10Gb/s interfaces of the UCNA. UCNAs enable efficient consolidation by isolating and prioritizing different traffic types and enforcing traffic type specific bandwidth guarantees. UCNAs also ensure that storage traffic is delivered with low latency, ensuring no degradation in performance.

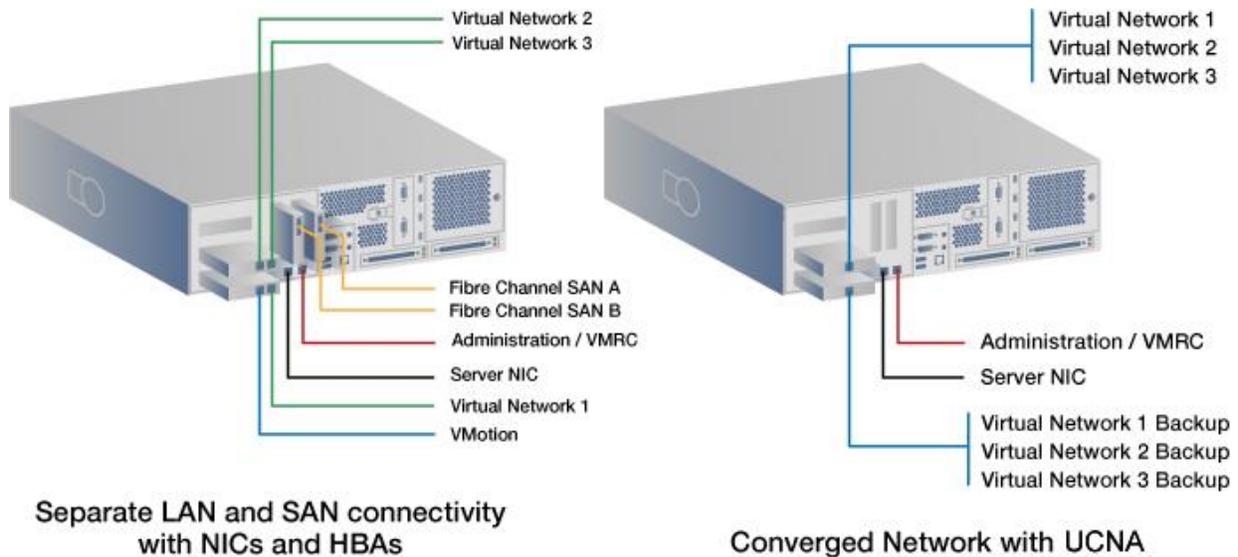


Figure 7. Converged networking with UCNA in a blade environment.

UCNA Functional Architecture

UCNAs appears as two PCI devices to the host — a network adapter and fibre channel or iSCSI adapter. Host networking and storage drivers communicate with the appropriate PCI function in the UCNA. The I/O operations are routed to the appropriate controller for processing. Networking transactions are delivered to the lossless NIC and forwarded on to the unified fabric.

In the FCoE portion of the UCNA, fibre channel frames are sent to the FCoE encapsulation engine, and then transferred to the lossless. Received traffic is handled in much the same way, with fibre channel blocks relayed to the appropriate host device driver.

10 GIGABIT ETHERNET: THE FOUNDATION FOR NETWORK CONVERGENCE

In the iSCSI portion of the UCNA, SCSI commands are exchanged between initiator and target devices over the IP network.

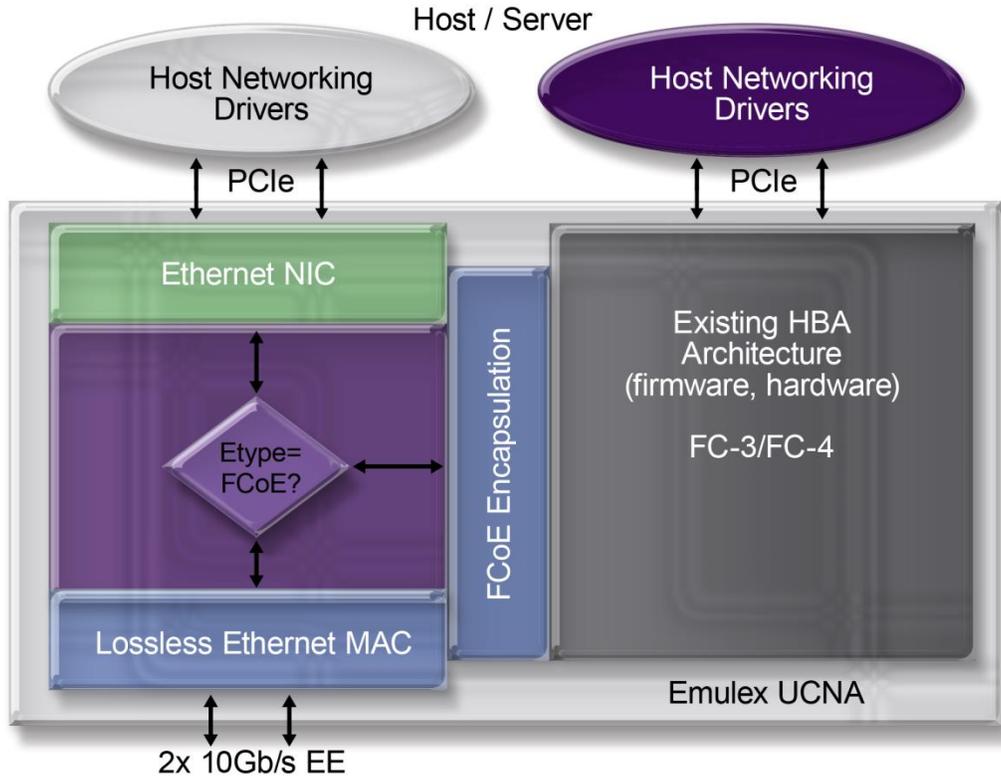


Figure 8. UCNA functional architecture

Device Manager View in the Operating System

Since the UCNA provides NIC connectivity with both fibre channel and iSCSI, the Windows device manager will show a 10Gbps NIC and fibre channel or iSCSI devices in the device tree. This level of transparency ensures minimum impact on management processes.

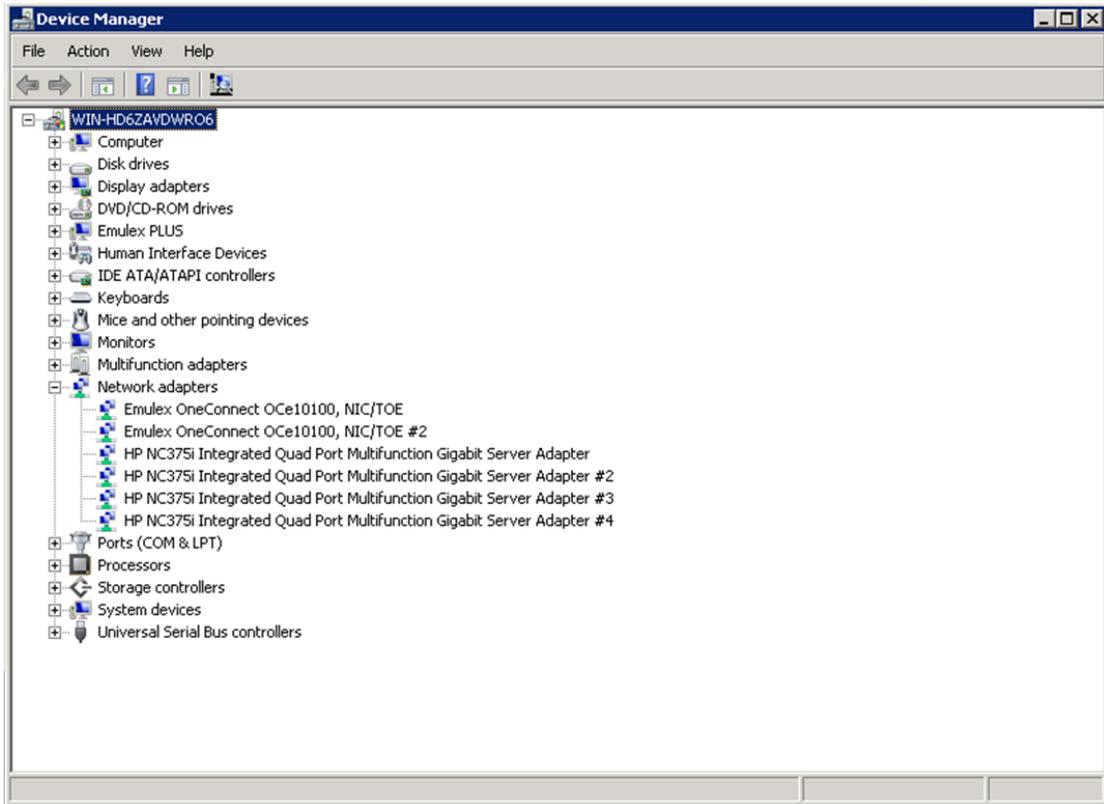


Figure 9. View of a UCNA in Windows Device Manager

Ixia’s IxChariot – Testing Performance from the User’s Perspective

IxChariot has been the industry’s leading benchmark tool for measuring network interface card performance for over a decade – since 10Mbps Ethernet adapters were first introduced on PC platforms. The key reason for the longevity of the IxChariot platform lies in the flexibility of its design and testing methodology.

IxChariot Overview

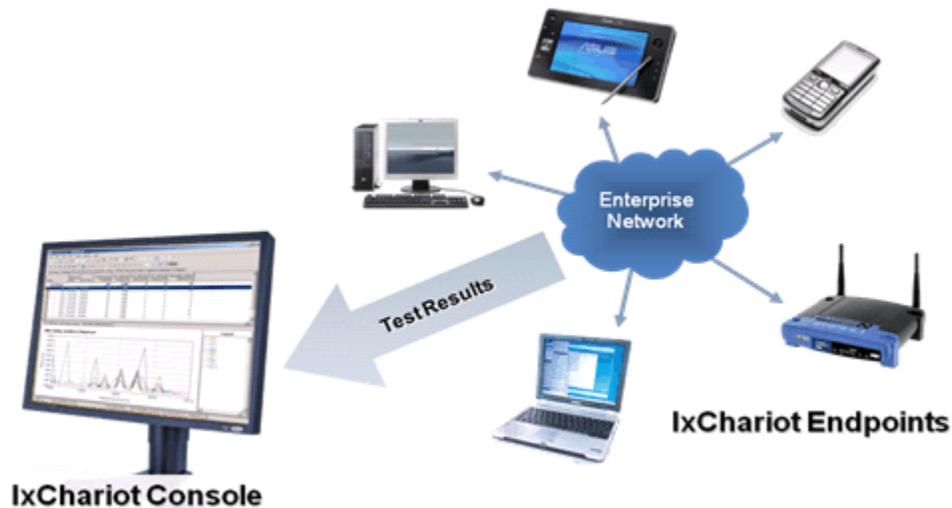


Figure 10. IxChariot Console and Endpoints

Endpoints

IxChariot conducts tests between software agents, called Ixia performance endpoints, deployed on devices connected to the network under test. The IxChariot endpoint software was designed from the ground up to be as lightweight and portable as possible, which is reflected by the fact that the endpoint software has been ported to more than 50 different platforms. These platforms include PCs, workstations, servers, virtual machines, cell phones, DSL modems, and wireless routers to name a few. The endpoint software interfaces with the same publicly available sockets APIs that are used by end-user applications.



Console

The IxChariot Console is the command and control center where tests are defined, executed and monitored by the operator. IxChariot tests can utilize any number of endpoint devices; test traffic is defined flow-by-flow. IxChariot’s flow-level granularity provides a great deal of flexibility in both test definition and result reporting. This allows tests to very closely simulate and assess performance in real-world environments.

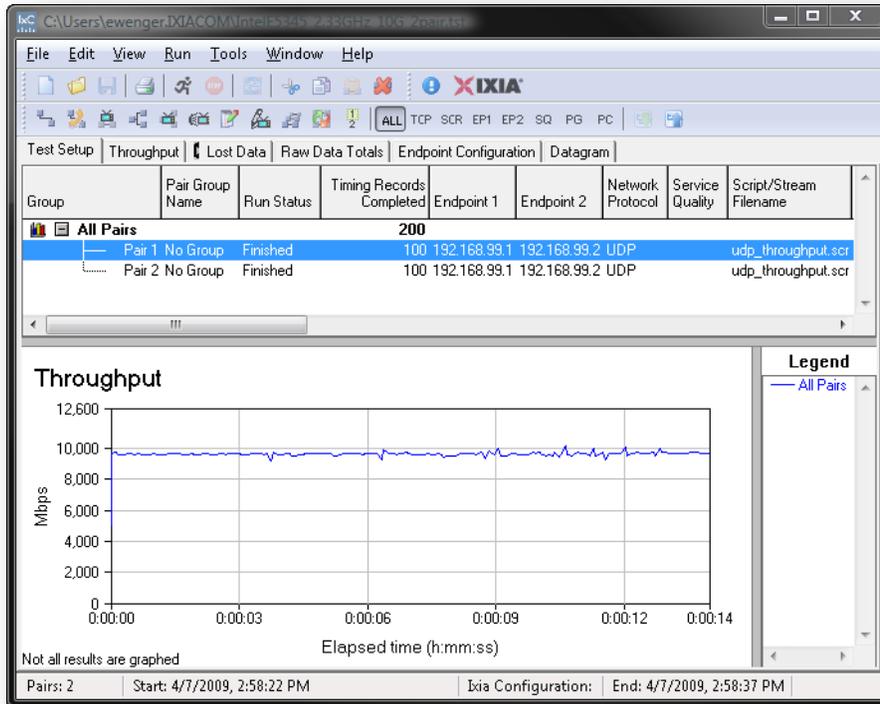


Figure 11. IxChariot Console

Pairs

IxChariot flow definitions are called *pairs*. A pair defines the following attributes of a flow:

- Endpoint IP addresses
- Simulated application profile, referred to as a script
- QoS profile
- Payload definition
- Timing information

During an IxChariot test, the console communicates with each of the endpoints in the test and downloads the configured pair definition to be run on that endpoint. The endpoint is then responsible for executing the test and returning results back to the console.

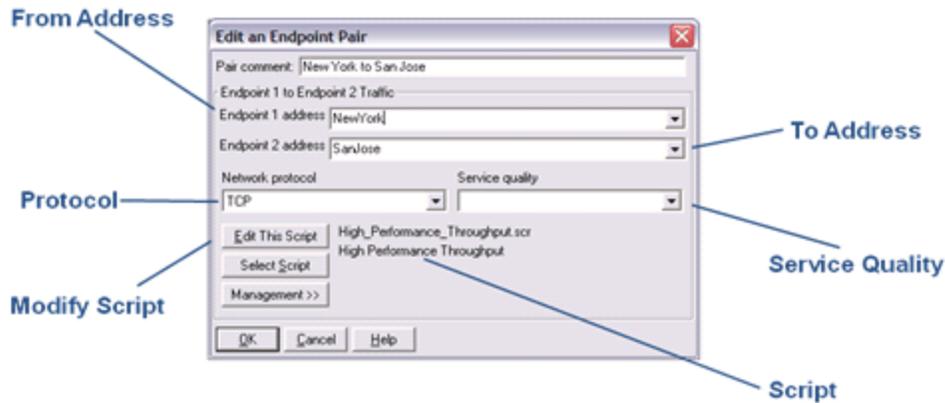


Figure 12. Pair Definition Dialog.

Scripts

An IxChariot script contains the instructions for a dialog between two endpoints. Each endpoint node in an IxChariot pair reads its half of the script to understand how to interact with its partner endpoint. As a simple example, an HTTP script appears as shown in Figure 13.

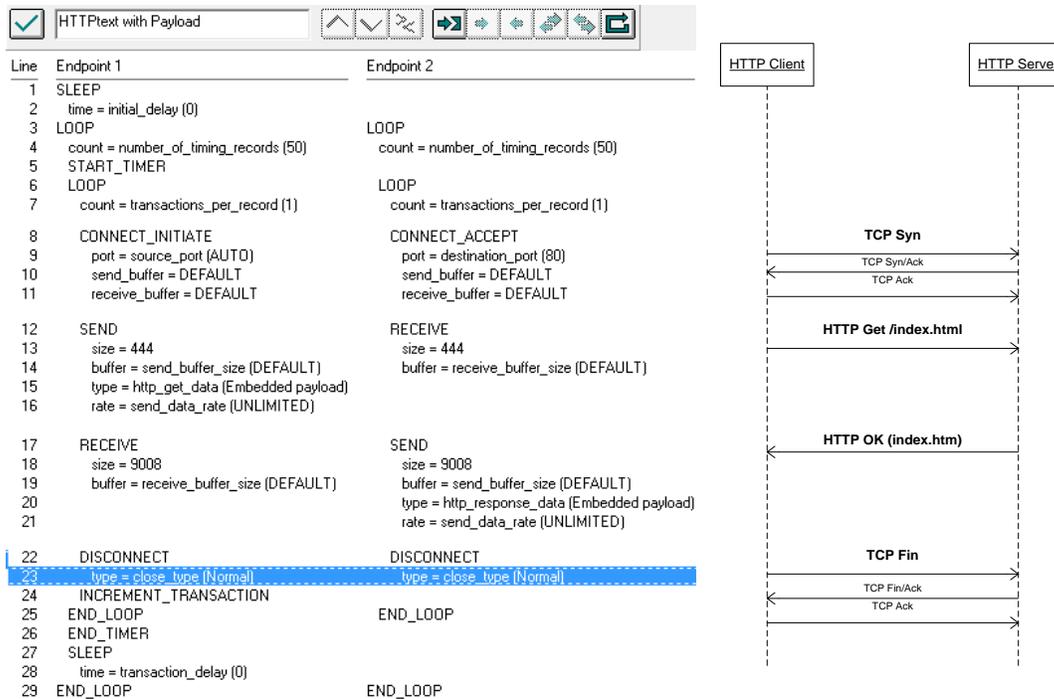


Figure 13. HTTP Script and Sequence Diagram.

This scripting mechanism makes it very easy to create multiple scenarios that allow users to understand how a network will perform with different kinds of user applications. IxChariot provides an easy to use tool that creates accurate application simulations that can be combined in a regression test bed. This is far easier than developing performance instrumentation for each

one of dozens of applications, Ixia’s IxProfile application can be used to model applications closely. IxProfile operates by monitor Windows sockets API call sequences that are used during typical application transactions, creating IxChariot scripts that model those specific transactions in great detail.



Figure 14. A few of the scripts included with IxChariot

IxChariot’s Approach to Testing Network Performance

IxChariot is used to measure a variety of different performance indicators using standard TCP/IP stacks on endpoint devices.

The optimum configuration for sending and receiving that data is dependent on test goals. For example, throughput measurements are best achieved by sending large data packets, while latency measurements are most accurate when measured with very small amounts of data.

As discussed previously, a pair executes instructions from its associated script via the endpoint operating system’s socket API. Each IxChariot pair defined on a particular endpoint will run in a separate OS thread on that node. If a test has four pairs between two endpoints, then the run-time process table on each endpoint will show that the endpoint has spawned four additional threads and each thread executes its part of the script defined for a pair. In the world of multi-core, multi-socket server platforms, it is important to note this since a single pair between two machines will only utilize a single core on each machine. To utilize multiple cores, you must define multiple pairs on each endpoint.

Throughput

IxChariot includes a number of scripts for measuring the throughput of devices or networks. These are explained in the following sections.

Basic Throughput

The basic **Throughput** script is the simplest methodology for measuring TCP throughput on networks and devices. The throughput script establishes a TCP connection from Endpoint 1 (E1) to Endpoint 2 (E2); this connection is maintained for the lifetime of the test. After establishing the connection, E1 sends large files of data (the value of the **file_size** variable, 100KB, by default) to E2 in blocks defined by the **send_buffer_size** variable – typically using the system default of 8KB – 32KB per block. The time required to send each file is recorded and converted to a throughput measurement that is graphed by the IxChariot Console. This test will continue to run until 100 files have been sent or for the defined test duration.

High-Performance Throughput

On Windows-based operating systems, the **High-Performance Throughput** script may be used to enable Winsock’s overlapped I/O mechanism, which can increase the efficiency of network transactions by allowing an application to queue multiple requests. Each request can then be handled independently by the Windows kernel. In addition to utilizing the overlapped I/O mechanism, the high-performance throughput script has a default **file_size** of 10 MB and a default **send_buffer_size** of 64KB. This script is ideal for testing TCP throughput on 100Mbps, 802.11n, and 1 Gbps networks.

Ultra-High Performance Throughput

The **Ultra-High Performance Throughput** script is designed for testing in the world of 10Gbps Ethernet and beyond. This script uses overlapped/asynchronous I/O as in the **High-Performance Throughput** script but increases the default settings for socket buffers, file sizes and send buffer sizes in order to maximize the benefit of TCP offloading technologies.

UDP Throughput

The **UDP Throughput** script was added in IxChariot 7.0 in order to provide an optimized method for measuring peak UDP throughput on the network. This script will stream UDP datagrams from E1 to E2 as quickly as the sender can put them on the network. There are a few key items that should be noted for UDP throughput testing.

- UDP traffic is more CPU intensive than TCP traffic for two reasons. First, the application writes each datagram directly to the OS; each datagram is then sent individually by the device driver. In TCP transactions, the application may send as much as 1MB of data to the OS in a single chunk and then the OS, or the NIC hardware itself, will packetize this data – a much more efficient process. The possibility of hardware offloaded packetization is the second reason.
- Because of the inefficient nature of UDP traffic generation on traditional operating systems, it is likely that many systems will not be able to generate more than 1.5 – 7 Gbps of UDP traffic using jumbo frames in a single pair on a modern server. To achieve line-rate throughput, multiple pairs must be used so that each thread can utilize an independent CPU core.
- The **send_buffer_size** variable in this script defines the UDP datagram size. As long as the entire datagram can fit in the MTU of the network packet, approximately 1460 bytes for standard framing and 9000 bytes for jumbo framing, then it will be sent as a

single packet. If the datagram size exceeds the network MTU by even a single byte, it will result in IP fragmentation that requires multiple packets per datagram, which can be a *very* inefficient process.

- It is likely that packet loss will be observed on a high-bandwidth UDP test even if the endpoints are directly connected to each other. This is typically due to the small default socket buffer size on most endpoint systems. The UDP throughput script has a 512KB socket buffer defined for receiving packets on E2, but this may need to be increased under extreme conditions.

Latency

The IxChariot **Response Time** script implements a classic ping-pong measurement of machine-to-machine round-trip time. E1 sends 100 bytes of data in a single small packet to E2, which immediately responds with another 100 byte packet back to E1. By measuring the time it takes to send a few thousand ping-pong transactions, the average round-trip latency between two endpoints can be estimated.

Application Performance

IxChariot includes more than 100 scripts based on modern Internet and enterprise network applications. These scripts can be used to measure the typical performance of these applications over a network between any E1 and E2. For example, customers who are building application servers will use application scripts matching their expected profile to understand how a specific server and network adapter combination will perform on their network. Database applications are a prime example in this category, since they are so heavily transactional in nature; a minor difference in overall system latency can translate into large amounts of idle time observed by end users.

VoIP and Video

Finally, IxChariot has unique simulation capabilities for VoIP and video traffic. IxChariot was the first tool on the market to implement accurate measurements of call quality based on the ITU G.107 E-Model specification. To measure the impact of network quality on call scores, IxChariot uses its UDP/RTP streaming capabilities. It sends simulated media frames in order to measure the jitter, latency and loss, which play a significant role in users’ experience. Some of the key factors that have led to IxChariot’s widespread use as a standard in measuring network voice quality are:

- Its ability to generate media streams and measure quality on-demand
- Quick and accurate one-way delay measurements
- Per-flow QoS settings and statistics
- Realistic background traffic generation using Ixia’s library of application scripts

Test Case: TCP Throughput

Overview

The **TCP Throughput** test is designed to measure the maximum TCP throughput and CPU effectiveness of a CNA across a wide range of TCP send buffer sizes. By testing with a range of TCP buffer sizes the performance of the CNA can be assessed with specific applications. For example, databases and application servers will utilize small buffer sizes (1K-4K), whereas file servers will more frequently utilize very large buffers (64K+).

Objective

Assess the maximum throughput and CPU utilization across the following range of buffer sizes, expressed in bytes: 64, 128, 256, 512, 1K, 2K, 4K, 8K, 32K, 64K, 128K, 1M. The end result should resemble that shown below:

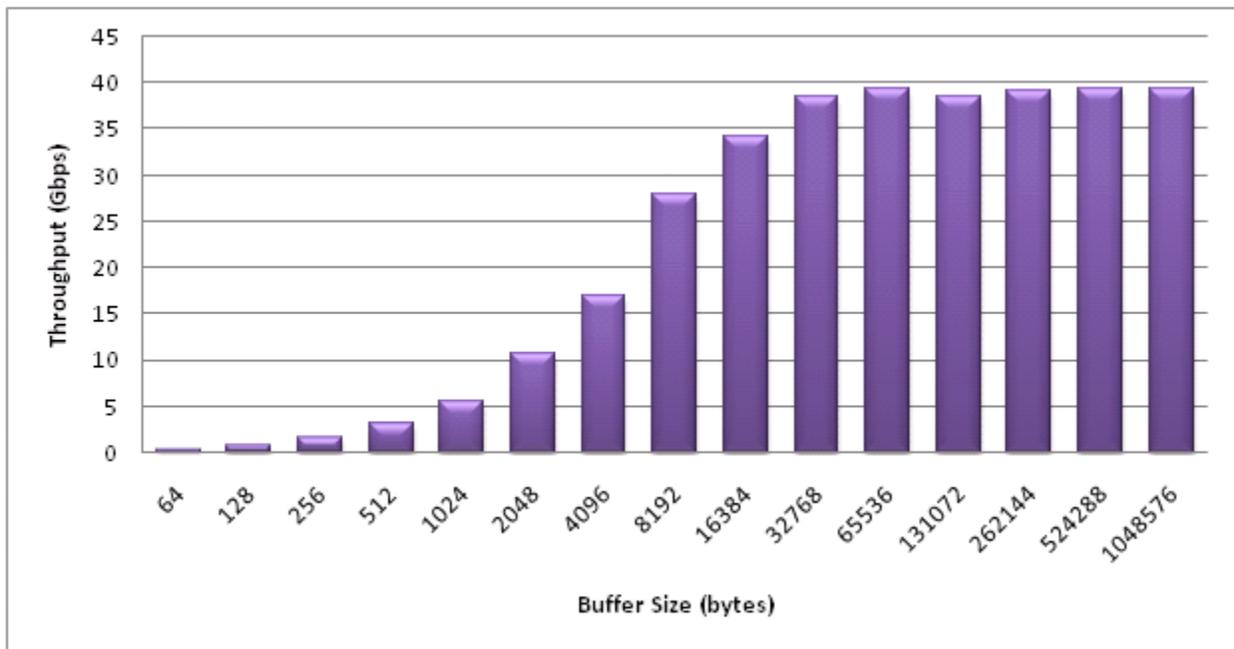


Figure 15. Emulex OneConnect 10GbE UCNA dual-port bidirectional throughput performance with jumbo frames running on Windows Server 2008 using two Nehalem 2.67 Ghz processors.

TEST CASE: TCP THROUGHPUT

Topology

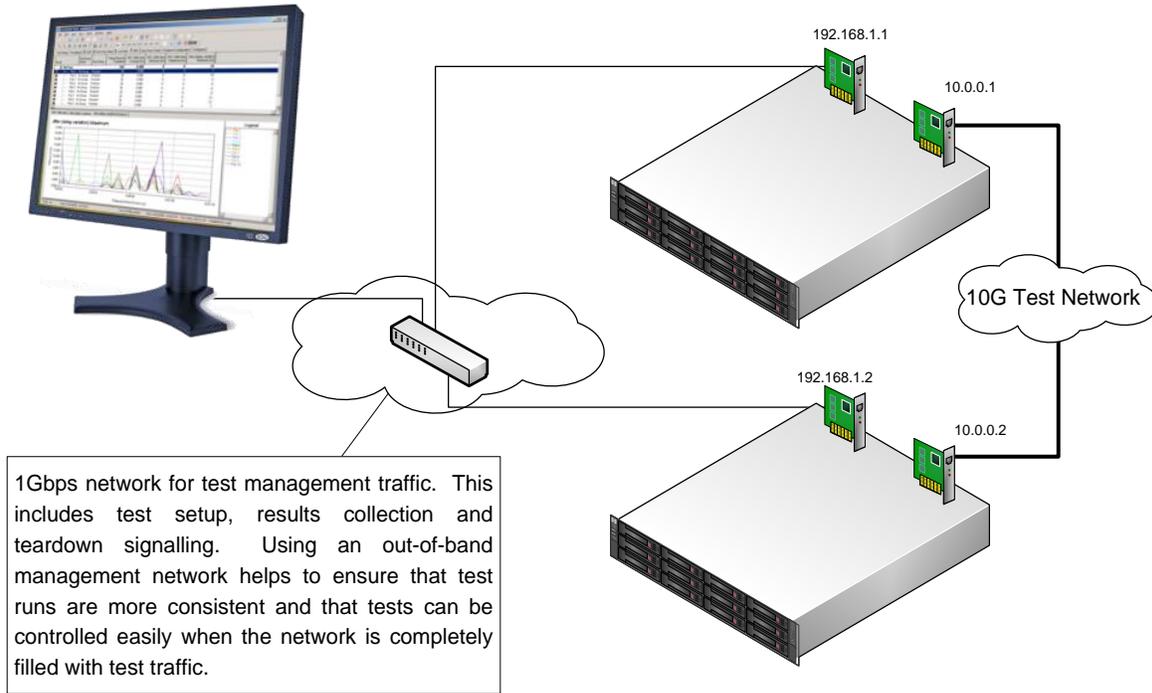


Figure 16. Basic Topology for CNA Testing.

Server Configuration

In this test, the servers should be configured for optimal TCP throughput. The general recommendations are:

- The operating system hosting the endpoint should be freshly installed and should have no unnecessary processes running, such as virus scanners or indexing utilities.
- Device drivers should be configured to utilize TCP and checksum offloading, and interrupt moderation should be configured to maximize throughput.
- When using Windows Server 2008, TCP chimney offloading and receive-side scaling should be enabled and receive window auto-tuning level should be set to normal or experimental.
- Jumbo frames should be enabled on the CNAs and on any switches included in the setup. This is not recommended as switches can introduce significant latency.

TEST CASE: TCP THROUGHPUT

Step-by-Step Instructions

1. Start the IxChariot Console...

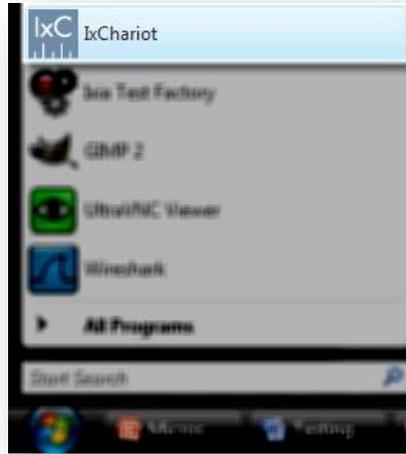


Figure 17. Start the IxChariot Console

2. Create a new pair in the blank test.

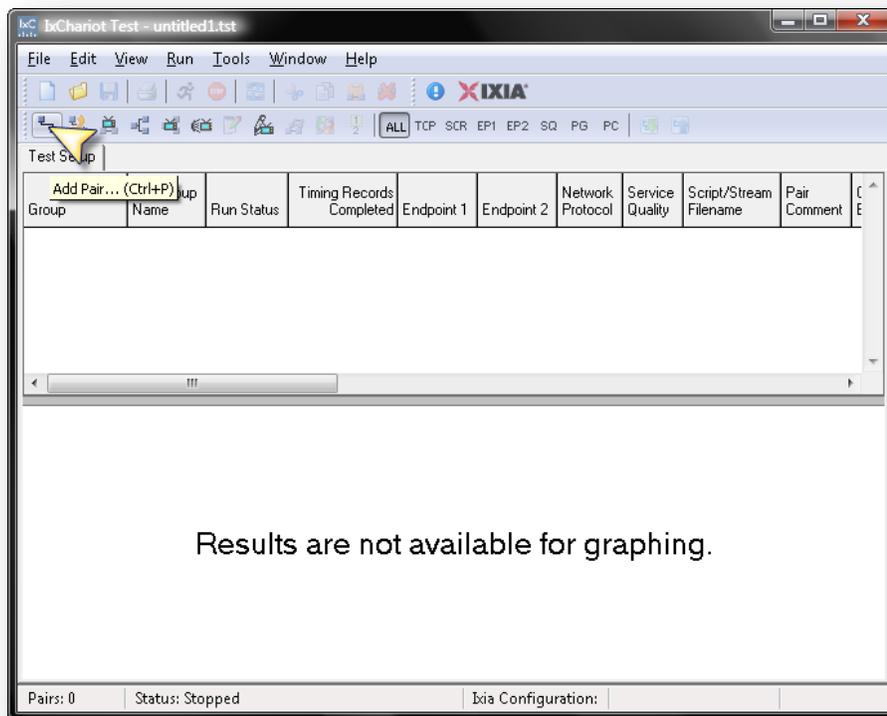


Figure 18. Create a new pair

TEST CASE: TCP THROUGHPUT

3. Configure the pair with the necessary parameters from test topology.

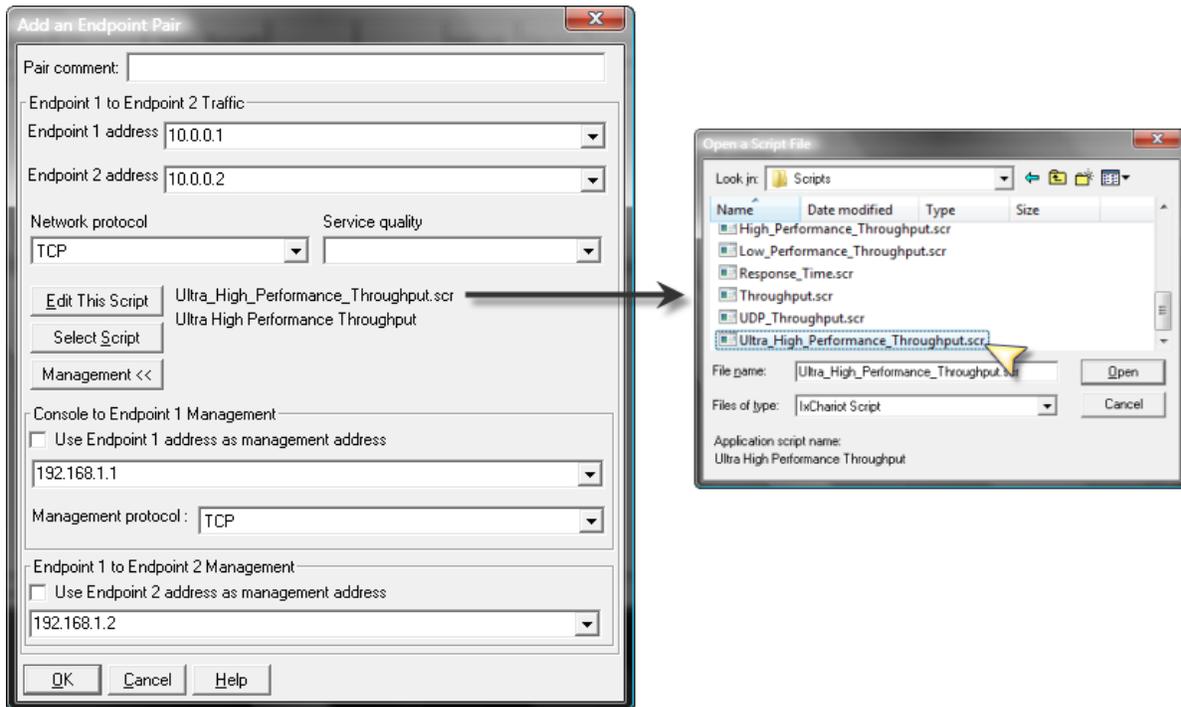


Figure 19. Add pair dialog showing script and endpoint IP addresses

TEST CASE: TCP THROUGHPUT

4. Modify the buffer size for this test run. To modify the buffer size, double-click to edit the pair and then click the **Edit script** button. This will display the dialog shown in Figure 20 below. The goal of this test is to measure the throughput and CPU utilization across a range of buffer sizes – from 64 bytes through 1MB. For smaller buffer sizes, we recommend that the record size on line 13 should be changed to a value that is 1000-10000 times the send/receive buffer size on line 14. One final recommendation for this test is to delete line 17, as this will remove an unnecessary confirmation exchange between the endpoints.

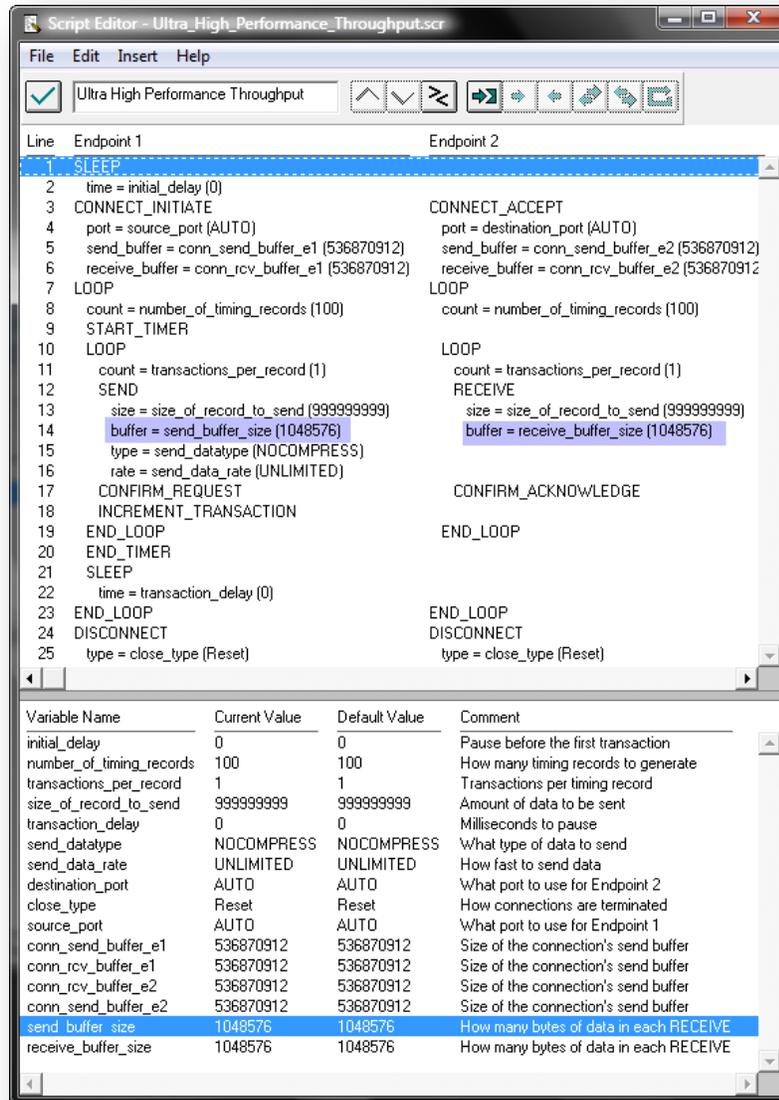


Figure 20. Editing the send buffer size.

TEST CASE: TCP THROUGHPUT

5. Set the **Run Options** for the test.
 - a. **Run for a fixed duration** – this option tells IxChariot to run this test for 1 minute.
 - b. **Batch Mode** – in this mode the endpoints will send their results back in batches to make more efficient use of CPU resources.
 - c. **Collect endpoint CPU utilization** – this option tells the endpoints to collect CPU utilization data and report it in the test results. This data is helpful for calculating the efficiency of the CNA.

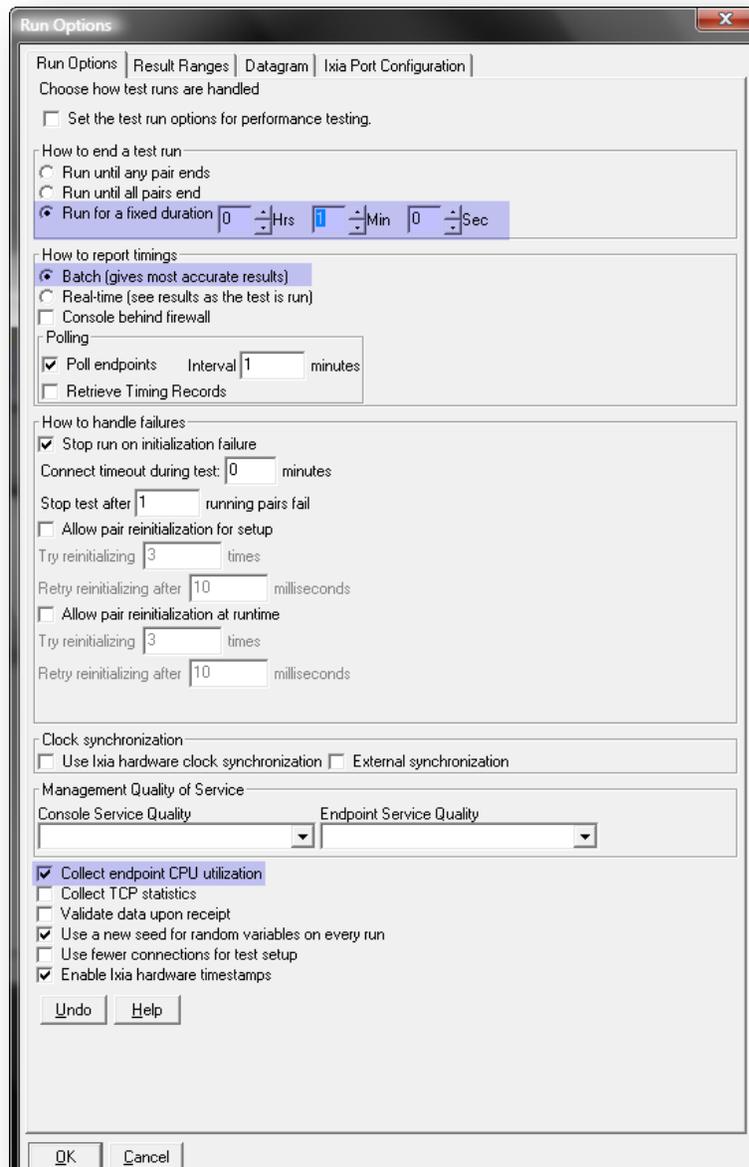


Figure 21. IxChariot Run Options dialog

TEST CASE: TCP THROUGHPUT

6. Run the test. 

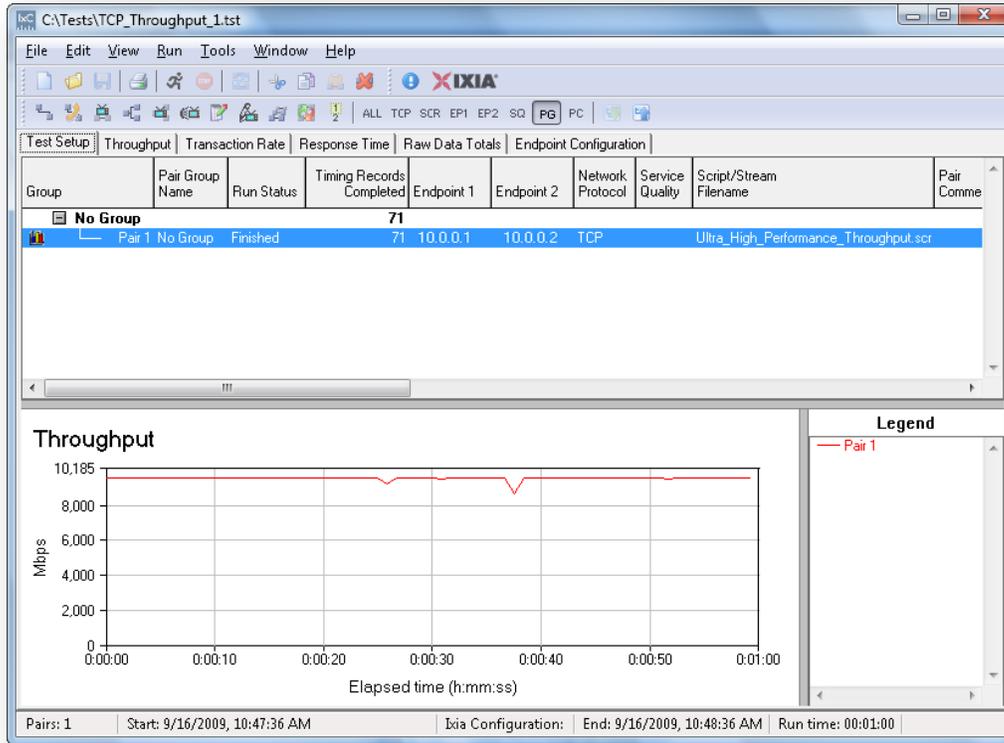


Figure 22. Test running at 10Gbps¹

Results Analysis

After running this test we now have one data point needed for building the graph shown in Figure 22. The key data points in this test are the Maximum Throughput and the CPU Utilization. Figure 23 and Figure 24 below highlight these data points as they appeared on the IxChariot Console. The results of this test indicate that the maximum throughput was 9.6 Gbps and the CPU utilization was 9% on the sender and 40% on the receiver side.

¹ Note – the results in this document are not representative of the performance of Emulex's OneConnect Universal CNA. The data contained herein were collected from a variety of different tests using equipment from multiple vendors.

TEST CASE: TCP THROUGHPUT

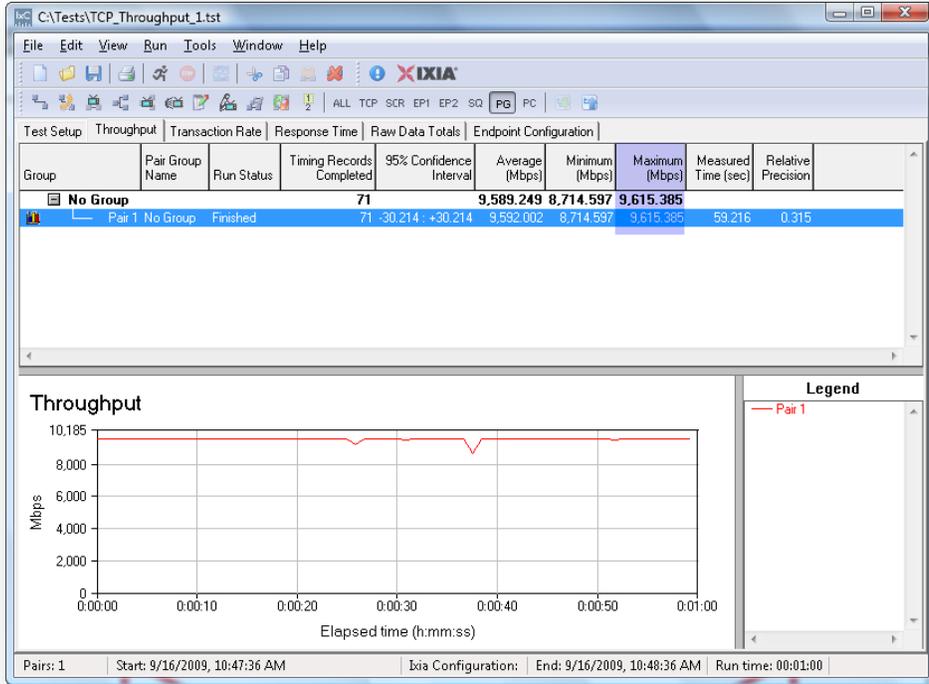


Figure 23. Maximum throughput

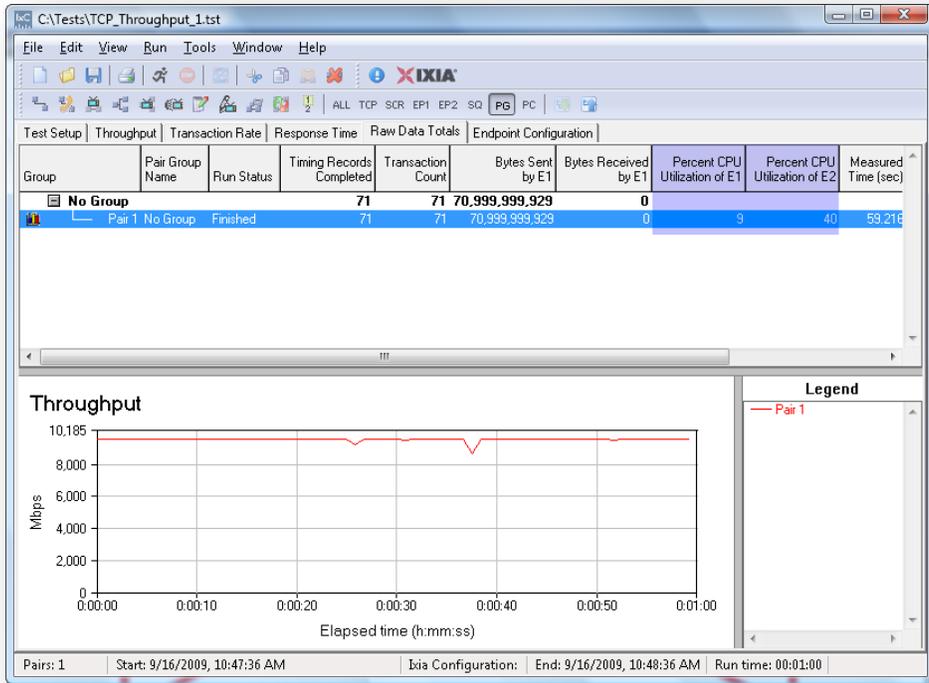


Figure 24. CPU utilization

Test Variables

Test Tool Variables

- This test should be run in both unidirectional and bidirectional mode to understand how the CNA handles full-duplex communications. In bidirectional mode, simply copy and paste the original pair (CTRL-C, CTRL-V) and then use the **Swap Endpoints** button () to reverse the direction of one pair. The user can also use the Shift+Click operation to select multiple pairs and copy or swap the entire grouping.
- In test scenarios where less than 100% of line rate is utilized, it may be interesting to create multiple pairs that allows utilization of multiple CPU cores for generating traffic. In a unidirectional test at small buffer sizes, 16 unidirectional pairs could be created using copy and paste, in order to utilize all 16 threads on dual-socket Nehalem-based servers. Likewise, in the bidirectional testing, 8 pairs could be created in each direction.
- Vary the socket buffer size using the default, 64K, 128K, and 1M. The socket buffer sizes can be adjusted on lines 5 and 6 of the script as shown in Figure 20. This is more of a functional test than a performance evaluation; it can sometimes highlight anomalies with certain value combinations.

DUT Test Variables

- Modify the device driver and operating system settings for offloading, TCP congestion algorithm, jumbo frames, VLAN tags, RSS and interrupt moderation to gain an understanding of how those settings impact performance.
- Incorporate a switched environment to simulate a typical data center network. The additional latency introduced by the switch will delay TCP acknowledgements slightly, resulting in reduced performance. This kind of evaluation can be helpful in diagnosing issues in actual deployments.
- Run this test using other operating systems. IxChariot endpoints support many server operating systems such as Windows Server 2008, RedHat and SUSE Linux, Solaris and Mac OS X. Each operating system will have a unique set of performance characteristics.

TEST CASE: TCP THROUGHPUT

Conclusions

After running the TCP throughput test sequence for each of the buffer sizes in the test plan, the following graphs can be reported to help understand how devices perform under varying conditions.

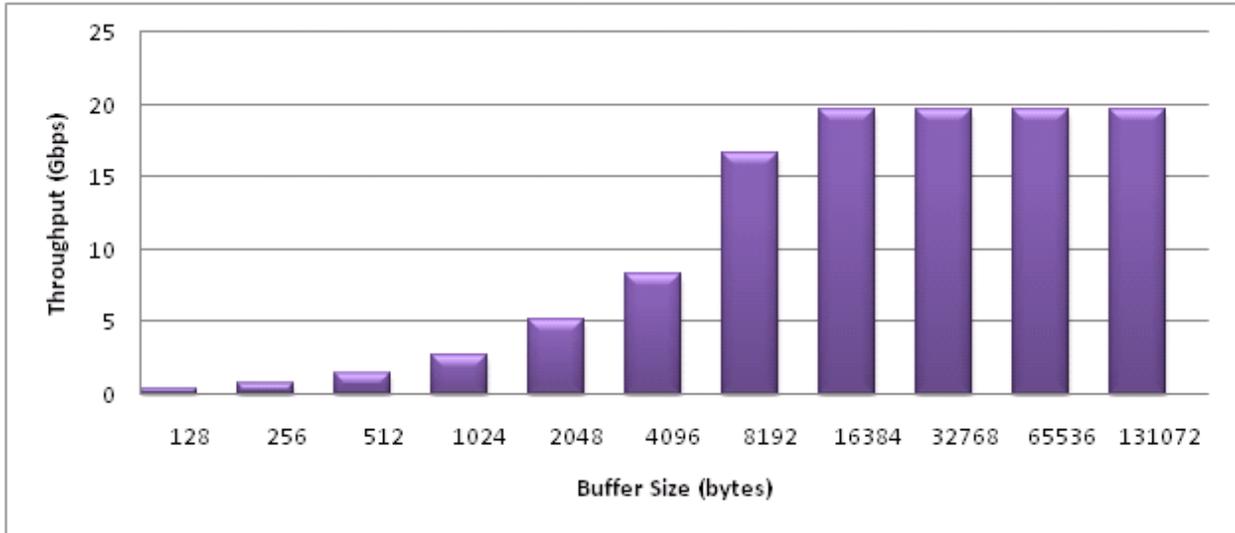


Figure 25. Emulex OneConnect 10GbE UCNA single-port bidirectional throughput performance with jumbo frames running on Windows Server 2008 using two Nehalem 2.67 Ghz processors.

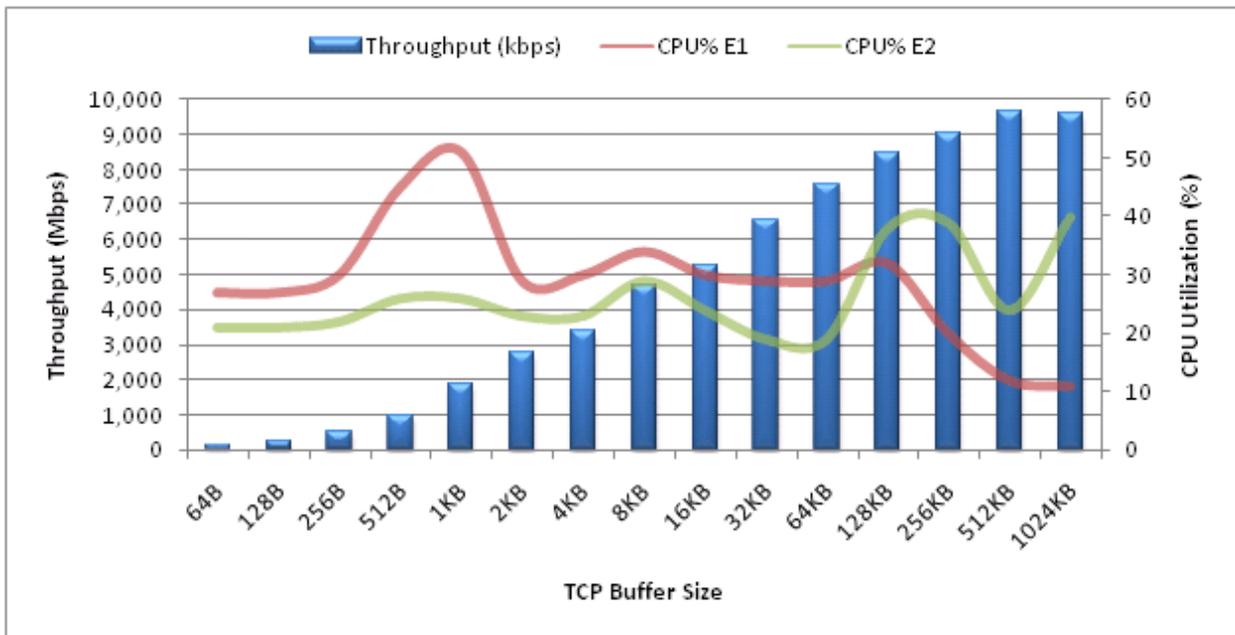


Figure 26. Throughput and CPU utilization (non-Emulex)

Test Case: UDP Throughput

Overview

Another important metric is UDP CNA throughput. Although UDP is used less frequently in enterprise environments, the use of UDP in IP-based video technologies provides compelling motivations.

Objective

The objective of this test is to measure the CNA's performance while streaming UDP traffic using a range of datagram sizes from 64 bytes through 1500 bytes.

Setup

This test case utilizes the same topology as the test. The primary difference of this test case versus that one is that typically no hardware or operating system optimizations, except UDP checksum offloading, will impact performance.

Step-by-Step Instructions

Test Setup

1. Start by building a new test and adding a pair to the test by following Steps 1-3 of the test. In Step 3, choose the **udp_throughput.scr** script and select UDP as the **Network Protocol** for the test in the **Edit Pair** dialog. The run options from Step 5 of the previous test remain the same.
2. Adjust the datagram size for packets generated by this script. The datagram size is controlled by the **send_buffer_size** on line 16 in Figure 27 below. As in the previous example, the **file_size** on line 15 should be 1,000 to 10,000 times the size of the **send_buffer_size**. The **send_buffer_size** for E1 must match the **receive_buffer_size** for E2 unless mismatched combinations are specifically being tested. In addition, make sure that the **send_buffer_size** is small enough to fit inside a single packet that includes Ethernet, IP and UDP headers, otherwise it will be fragmented by the IP stack.

TEST CASE: UDP THROUGHPUT

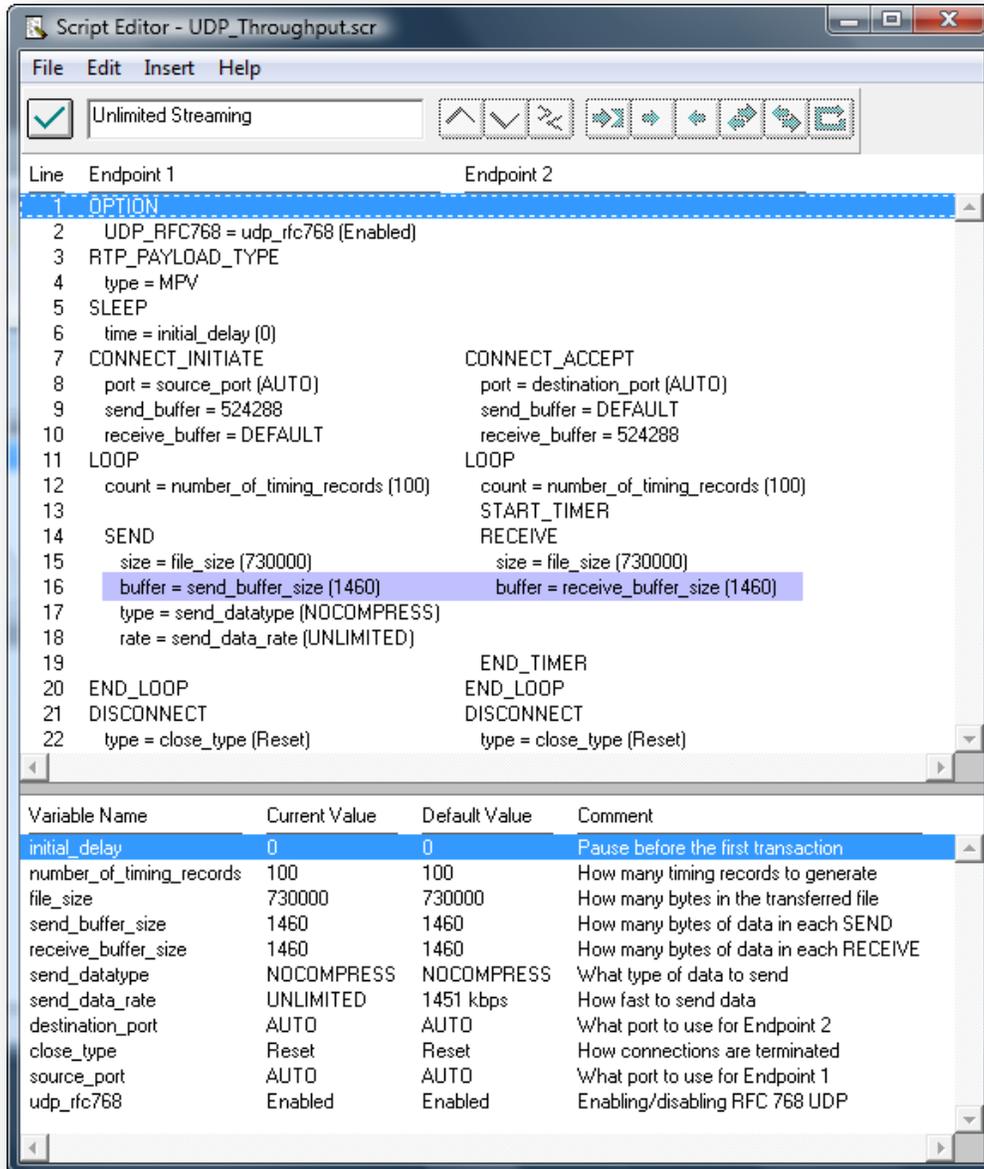


Figure 27. Setting the UDP datagram size

3. Run the test 

TEST CASE: UDP THROUGHPUT

Results Analysis

The analysis of UDP throughput tests results is quite similar to the analysis for TCP – the primary performance indicators are throughput and CPU utilization. In UDP testing, packet loss at specific bit rates may also be of interest. Using the **send_data_rate** variable on line18 of Figure 27 above, the UDP bit rate can be set to a specific value. Figure 28 below shows the packet loss statistics in the console.

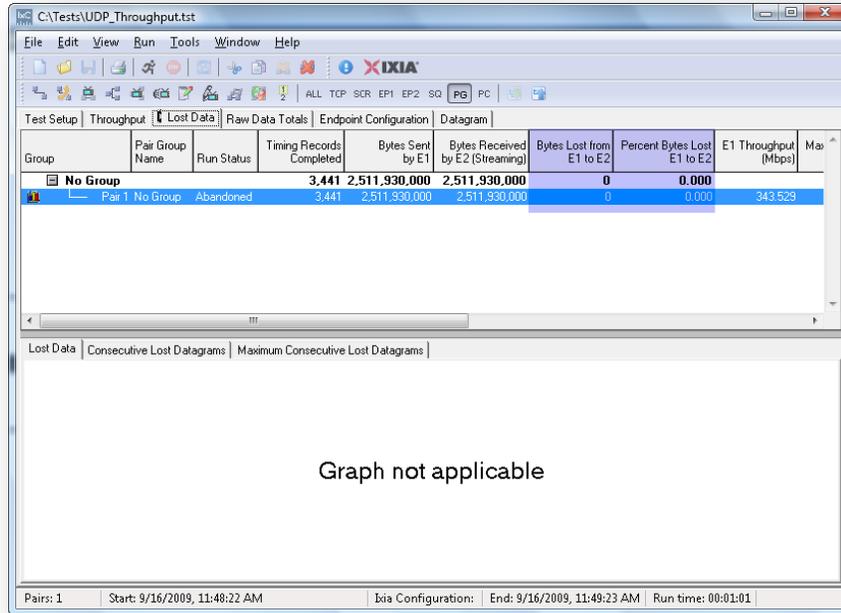


Figure 28. UDP packet loss statistics

One additional step may be used if packet loss is observed. At very high data rates, the default UDP receive buffer may be too small to hold packets during the test. Try increasing the **send_buffer** and **receive_buffer** values on lines 9 and 10 from 512KB to 1MB or higher to see if this helps.

Test Variables

Test Tool Variables

The same variations apply here as in the first test, with the additional point that increasing the number of flows to utilize multiple cores will have a significant impact on overall UDP test performance since there is no hardware-based offloading in the CNA.

Test Case: Latency

Overview

For many applications, such as databases or algorithmic trading software, latency is much more important than total throughput. This test case discovers how the CNA can impact the total response time. To do this, IxChariot sends out a series of requests and measure the time it takes to get responses for each transaction. This measures the total response time from the application layer on one endpoint to the application layer of the other endpoint, including any delays in the operating system TCP/IP stack and the I/O bus of the system. As a result, measurements of this type should primarily be used to compare multiple devices, not as a direct measurement of the device's latency.

Objective

The objective of this test is to measure the end-to-end latency or response time between two systems, providing an indirect indicator of CNA introduced latency.

Setup

This test case utilizes the same topology as the test. The primary difference between that case and this one is that typically there are no hardware or operating system optimizations, except UDP checksum offloading, that will impact the performance. The best results from this test are generally observed with UDP traffic, since offloading and interrupt moderation tend to induce greater latency in TCP transactions.

Step-by-Step Instructions

Test Setup

1. Start by building a new test and adding a pair to the test by following Steps 1-3 of the test. In Step 3, choose the **response_time.scr** script and select UDP as the **Network Protocol** for the test in the **Edit Pair** dialog. The run options from Step 5 of the previous test remain the same.

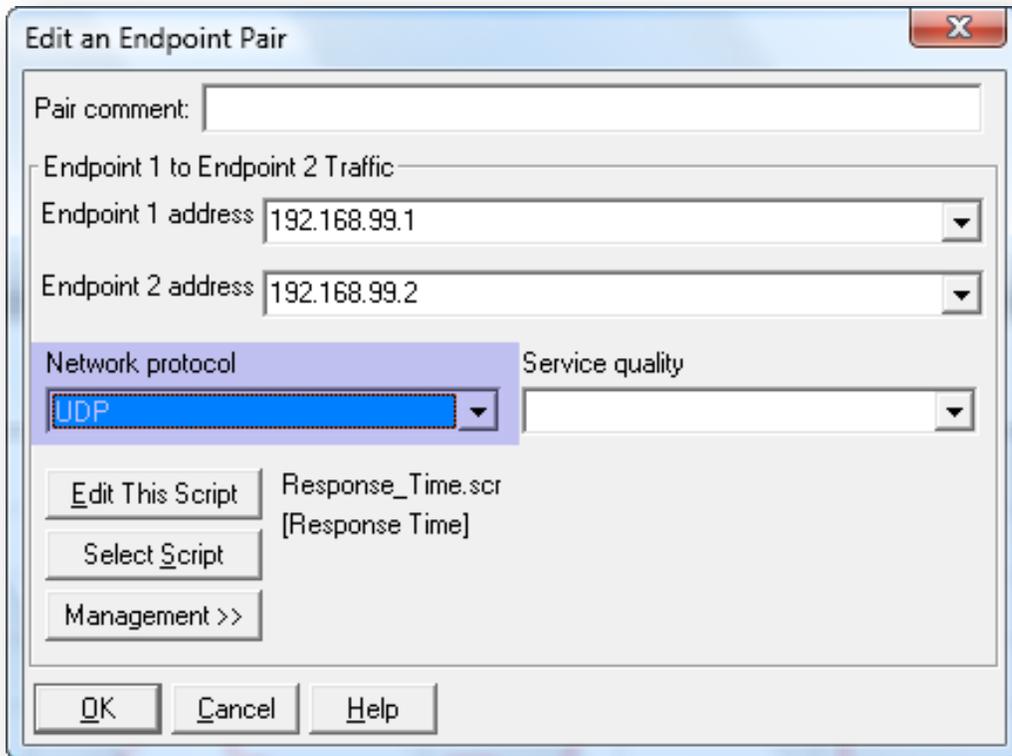


Figure 29. Select Response_Time.scr and UDP Network Protocol

TEST CASE: LATENCY

2. Modify the script (**Edit This Script**) to increase the **transactions_per_record** to 5000, as shown in Figure 30 below. This will cause a larger amount of data to be collected in each timing record, providing more accurate results.

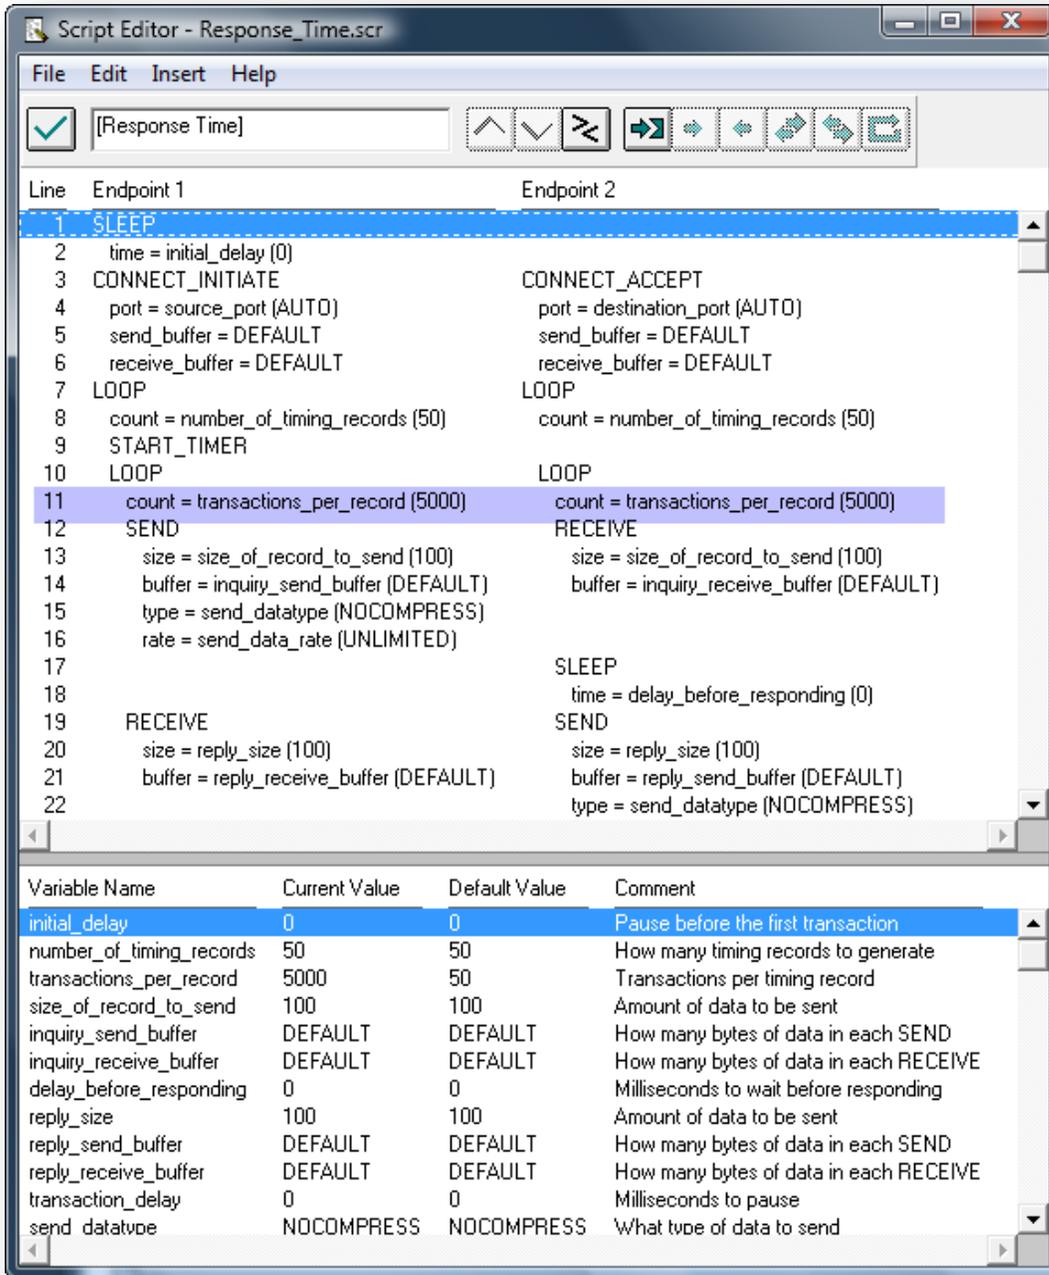


Figure 30. Increasing the `transactions_per_record` variable.

TEST CASE: LATENCY

3. One further optimization is required in the **Run Options** dialog to increase the efficiency of the test. The **UDP Window Size** value in the **Run Options Datagram** tab must be increased to 100KB as shown below.

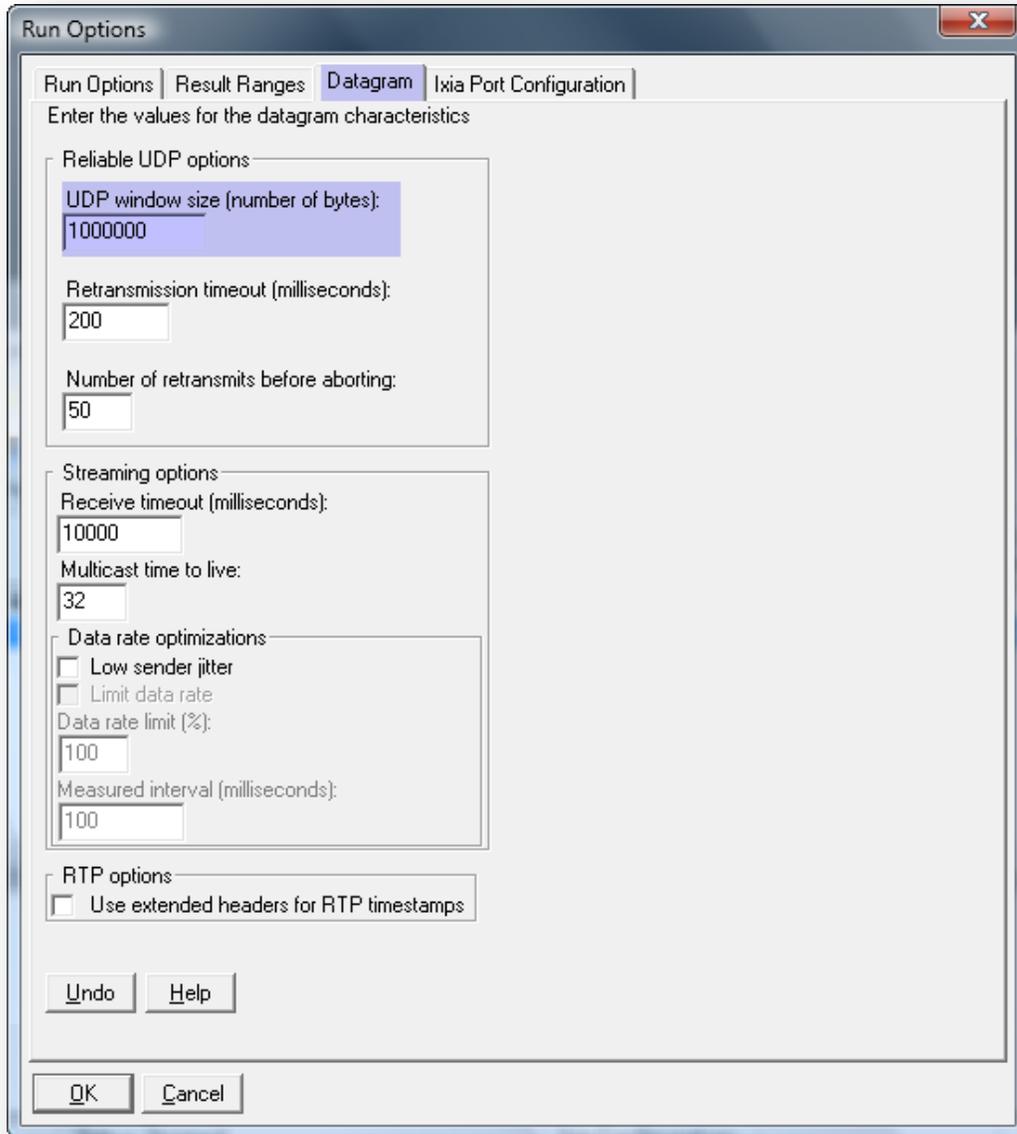


Figure 31. Changing the UDP window size

4. Run the test 

TEST CASE: LATENCY

Results Analysis

The key results from the latency test are shown in Figure 32 below. The key metric in this case is the number of transactions completed per second. In this case the test was able to complete a maximum of **14,164** transactions per second. By inverting this metric you can see that the latency, or seconds per transaction, was 1/14,164 or 70.6 microseconds. Since the test data traversed two CNAs (one on each server), twice each (transmit and receive), the latency can be divided by four, yielding 17.5 microseconds. This measures the time it takes for an application to pass a packet into or out of the CNA. Since this metric is captured at the application layer, this also includes the measurement of the operating system stack and device driver.

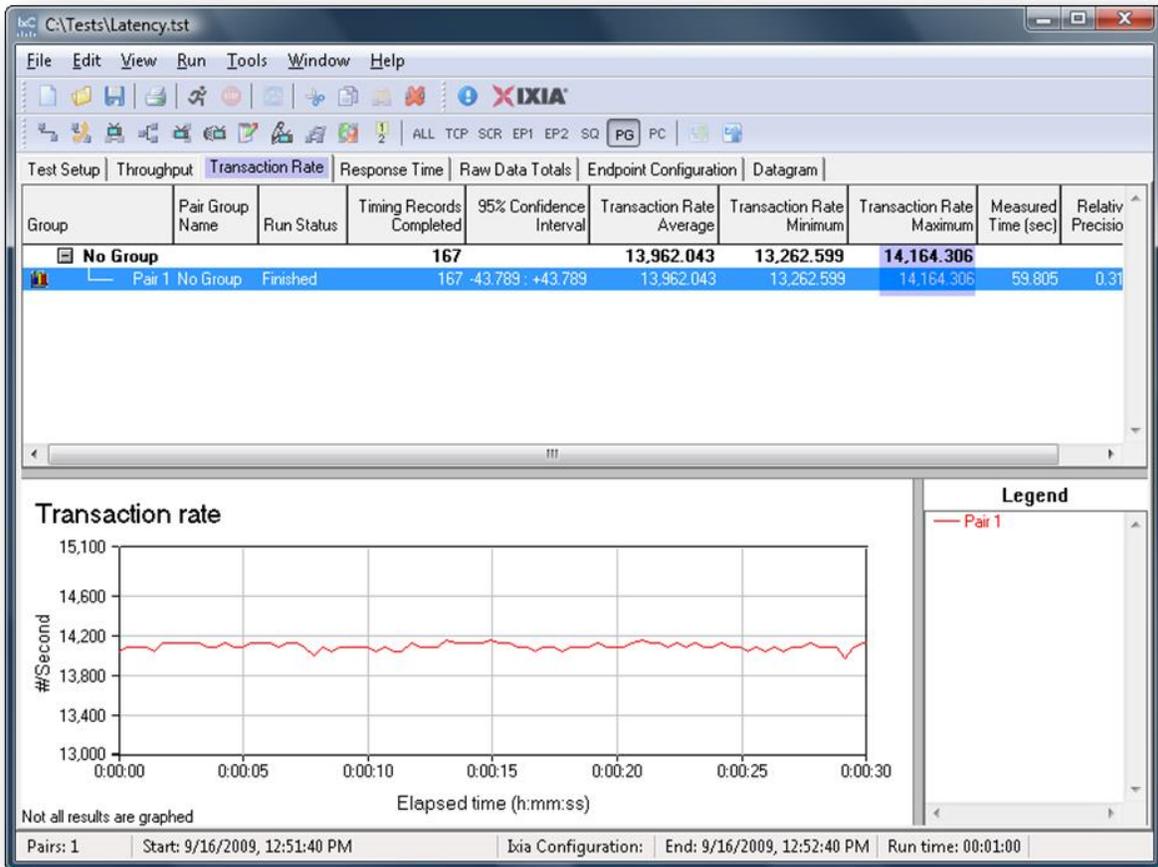


Figure 32. Results from the Latency test.

Test Variables

Test Tool Variables

A measure of response time using TCP or one of the enterprise application scripts may be desired in order to ascertain the expected performance of a CNA for certain applications.

TEST CASE: LATENCY

DUT Test Variables

Many CNAs now include settings to tune interrupt moderation for higher throughput or lower latency. This test is ideal for measuring the impact of these settings on application performance.

Conclusions

Since this metric is captured at the application layer, this also includes the measurement of the operating system stack and device driver. The quality of the device driver can impact the latency as much as the hardware; therefore, latency measured in this manner can be very important. For this reason, the latency measurement is best used as a comparison between two CNAs to gauge their relative performance.

Test Case: Storage I/O Performance

Overview

Testing the performance of the CNA storage component could very well be the subject of its own Black Book. This section of this test methodology is an overview of the test plan and methodology for storage I/O performance testing.

A flexible storage I/O test application implementation is shown in the figure below. The basic technique runs a user-space application that can open, close, read and write files on any block device or file system. This is the same kind of activity that is observed in any user application and provides the best ways of measuring device performance.

Ixia is developing a groundbreaking new test application for I/O server performance testing. As a part of its IxVM family of products, Ixia has developed a software-based implementation of its hardware test port technology that can support a full range of layer 2-7 network protocol testing originating from an x86-based server. To support I/O testing, Ixia has added a new “protocol” to its IxLoad product that can be used to emulate user applications running on a host server, or virtual machine, accessing both local and remote disks and file systems mounted in the host’s operating system.

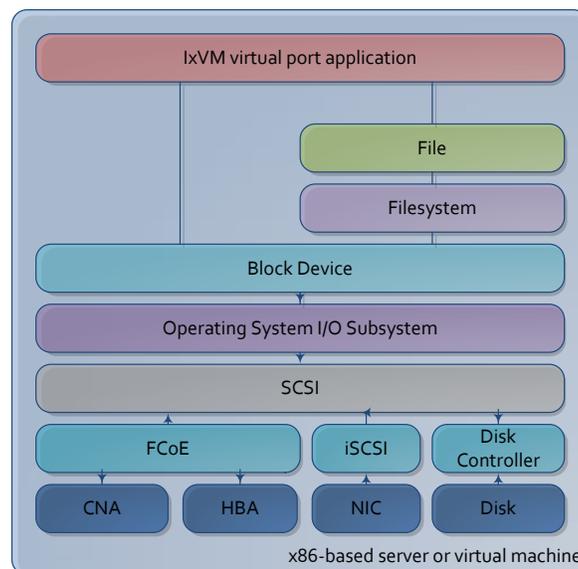


Figure 33. Emulated virtual machine

By installing the IxVM virtual port agent on a standard server or virtual machine, the IxLoad I/O plug-in can build highly sophisticated tests of the storage subsystem and SAN. Combined with IxLoad’s ability to generate high volumes of TCP traffic, this I/O plug-in can be used to create synchronized, mixed scenarios where IP and FCoE traffic are measured in coexistence with each other.

TEST CASE: STORAGE I/O PERFORMANCE

Objective

The objectives of a storage I/O test are very similar to the objectives of the IP data testing outlined in the first three test plans. The key performance indicators are:

- Throughput while reading and/or writing data to the storage device in block sizes ranging from 512B to 1MB.
- I/O operations per second for reading/writing data in the same set of block sizes.
- Latency of the system while reading/writing data at each block size
- Throughput and I/O operations per second for several application-specific mixtures of traffic with predefined read/write operations ratios over a statistical distribution of block sizes.

Setup

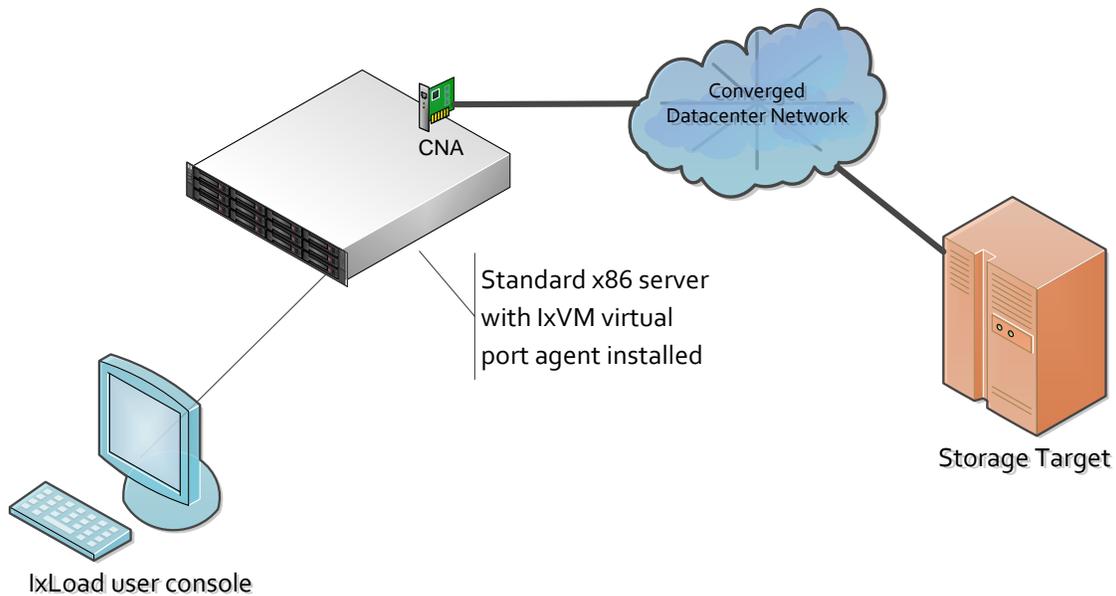


Figure 34. Typical storage I/O testing topology

Step-by-Step Instructions

1. Define target disks to be used in the defining I/O transactions.

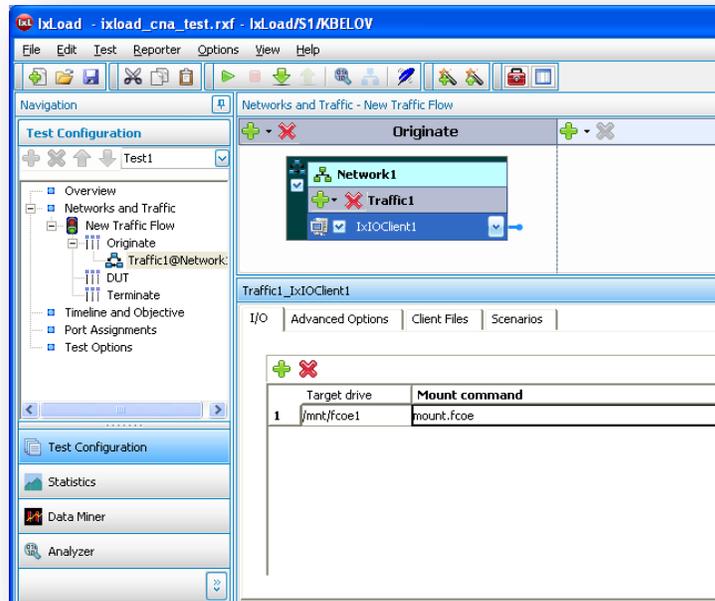


Figure 35. Specify disk targets and scripts for mounting those targets into the file system

2. Define target files or block ranges for conducting raw disk transactions.

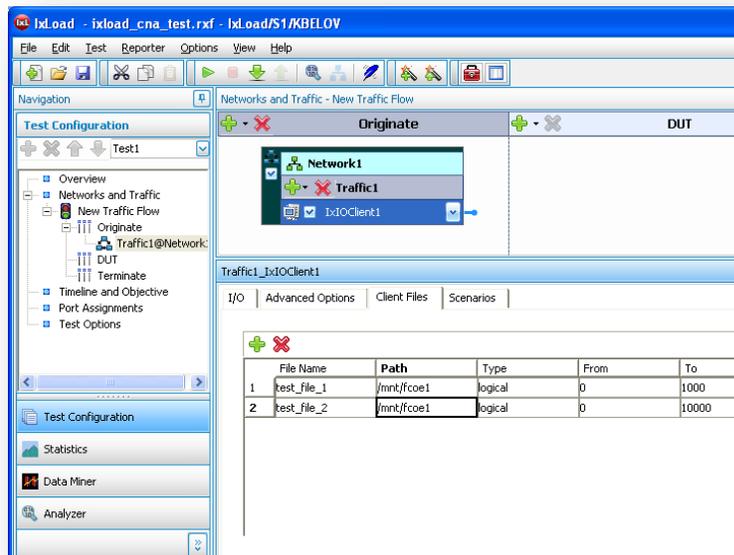


Figure 36. Defining files to be used in I/O transactions

TEST CASE: STORAGE I/O PERFORMANCE

3. Define fine-grained transactions simulating real-world I/O transactions. These transactions can be scripted using IxLoad's typical sequences and control constructs in the same manner as a real world application conducts I/O operations. In addition, configuration files from popular freeware I/O testing tools can be imported and converted into IxLoad's format.

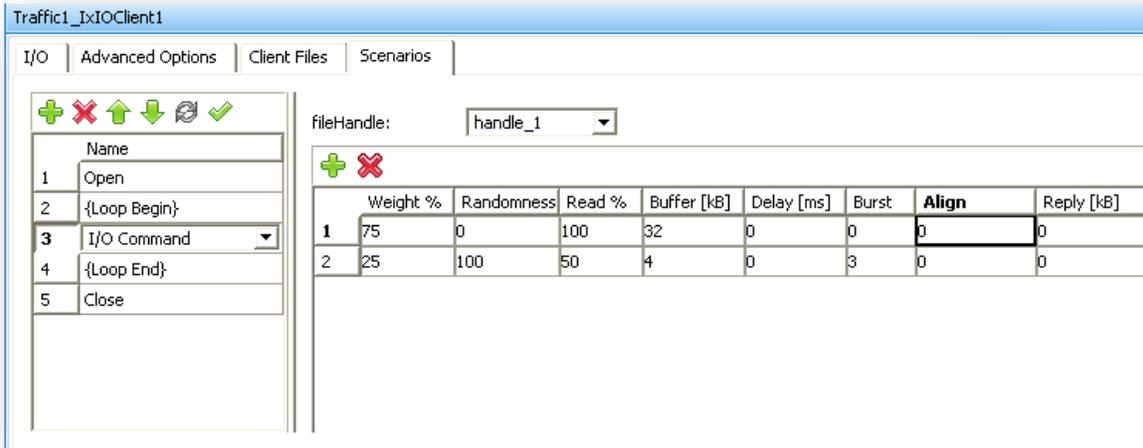


Figure 37. Defining an I/O transaction

4. Add additional application-layer protocols to the I/O traffic. IxLoad is able to statefully emulate dozens of real-world network applications.

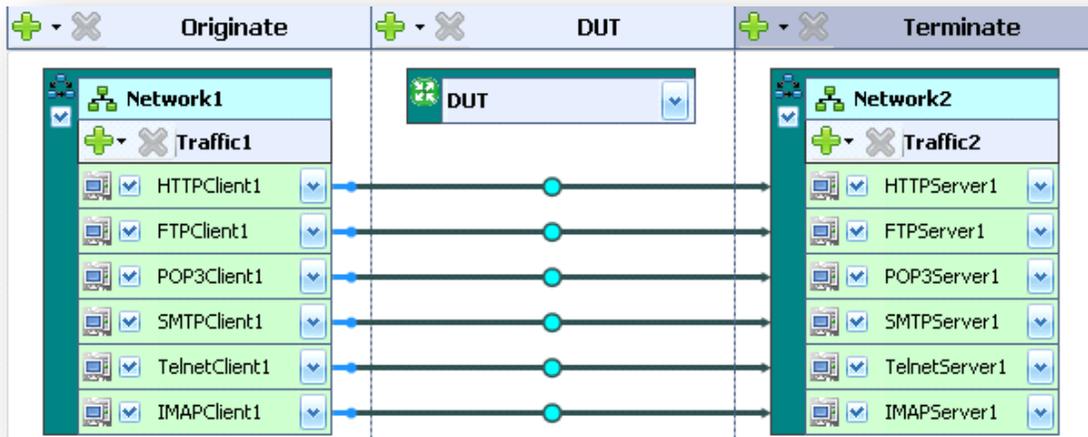


Figure 38. A selection of protocols supported by IxLoad/VM

5. Define per-activity timelines to create custom mixtures of application and I/O traffic varying over time. Mixes varying levels of HTTP and I/O traffic to understand how each type of traffic impacts the other on a PFC-enabled CEE network.

TEST CASE: STORAGE I/O PERFORMANCE



Figure 39. Example timelines available in IxLoad/VM

Results Analysis

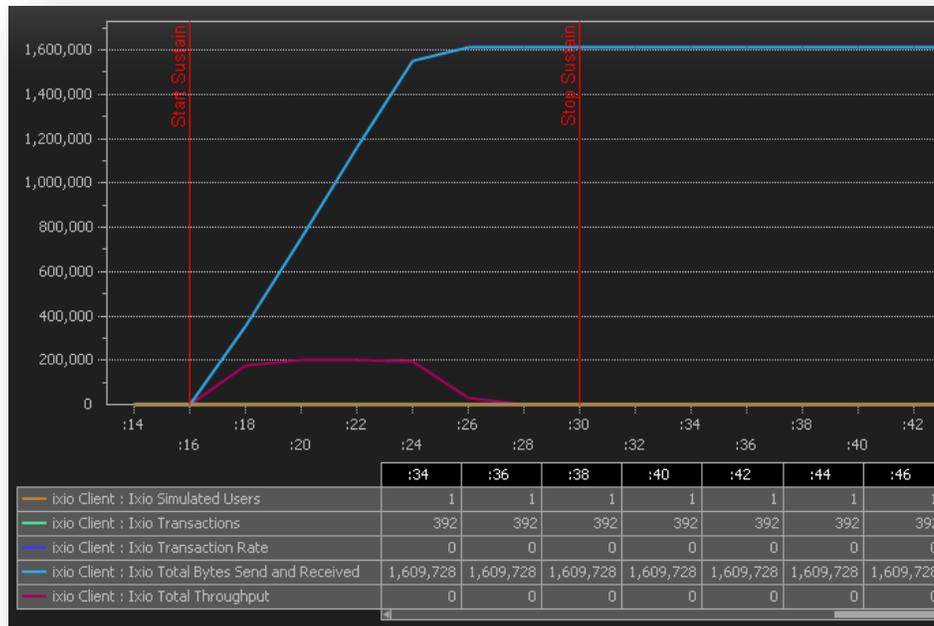


Figure 40. IxLoad I/O test results

The IxLoad I/O plug-in reports a full range of statistics covering throughput, operations per second, CPU utilization, network bytes sent and received and more. These data can be exported as graphs or in complete CSV files for user reporting.

Test Case: Virtualized Environments

Overview

Most 10Gbps CNAs found in today's datacenter are utilized in virtualized environments. There are two primary cases used in virtualized environments that are important to assess. In the first case, VMs utilize virtual NIC adapters that are connected to a virtual switch. The virtual switch then passes traffic to the CNA to be sent to other physical hosts. In the second case, the CNA is directly mapped to a specific VM so that it avoids the many layers of software in the virtual switch and operates more like a physical interface.

Objective

This test case describes the topologies needed to run the four primary benchmarks in virtualized environments.

Setup

There are two primary ways in which a 10G CNA can be used in a virtualized environment.

Shared among VMs using a vSwitch

In this model, which supports a wealth of existing hardware, a software-emulated virtual switch is used to direct packets from virtual NICs to physical NICs. Each virtual machine in this model uses a software-emulated virtual NIC that is connected to a port on the virtual switch. The virtual switch uses the physical CNA as an uplink to the enterprise network. This allows many virtual machines to easily share a limited number of CNAs. The primary disadvantage to this model is that the virtual switching process can consume a large amount of CPU time, reducing the overall efficiency of the system.

TEST CASE: VIRTUALIZED ENVIRONMENTS

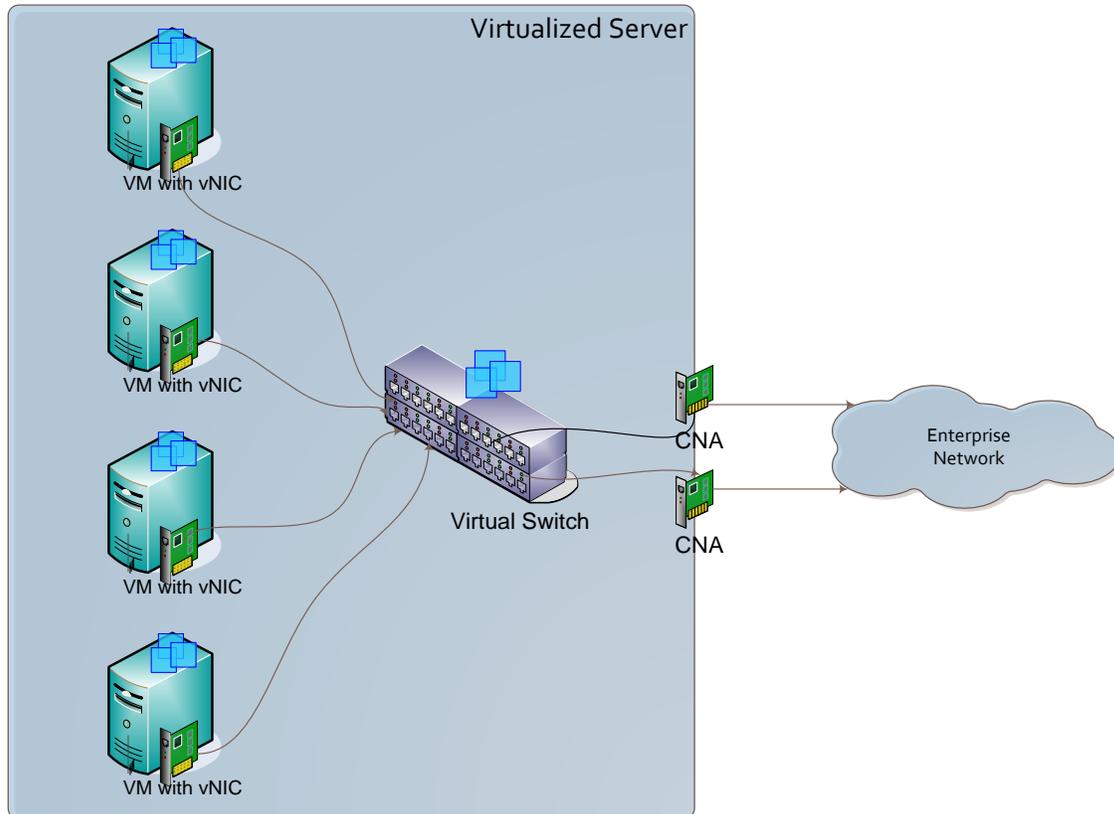


Figure 41. Virtualized topology showing the use of a virtual switch.

VMs directly connected to CNA using PCI-IOV

Next-generation CNAs use a technology called I/O virtualization (IOV) to enable VMs to directly access physical hardware. In this model, the VM has an actual device driver for the physical card; it is assigned a virtualized I/O address and IRQ for interacting with that CNA. The CNA management software can be used to assign priorities and bandwidth allowances for each virtualized slice of the CNA. This removes the processing overhead of the vSwitch and frees up significant CPU time.

TEST CASE: VIRTUALIZED ENVIRONMENTS

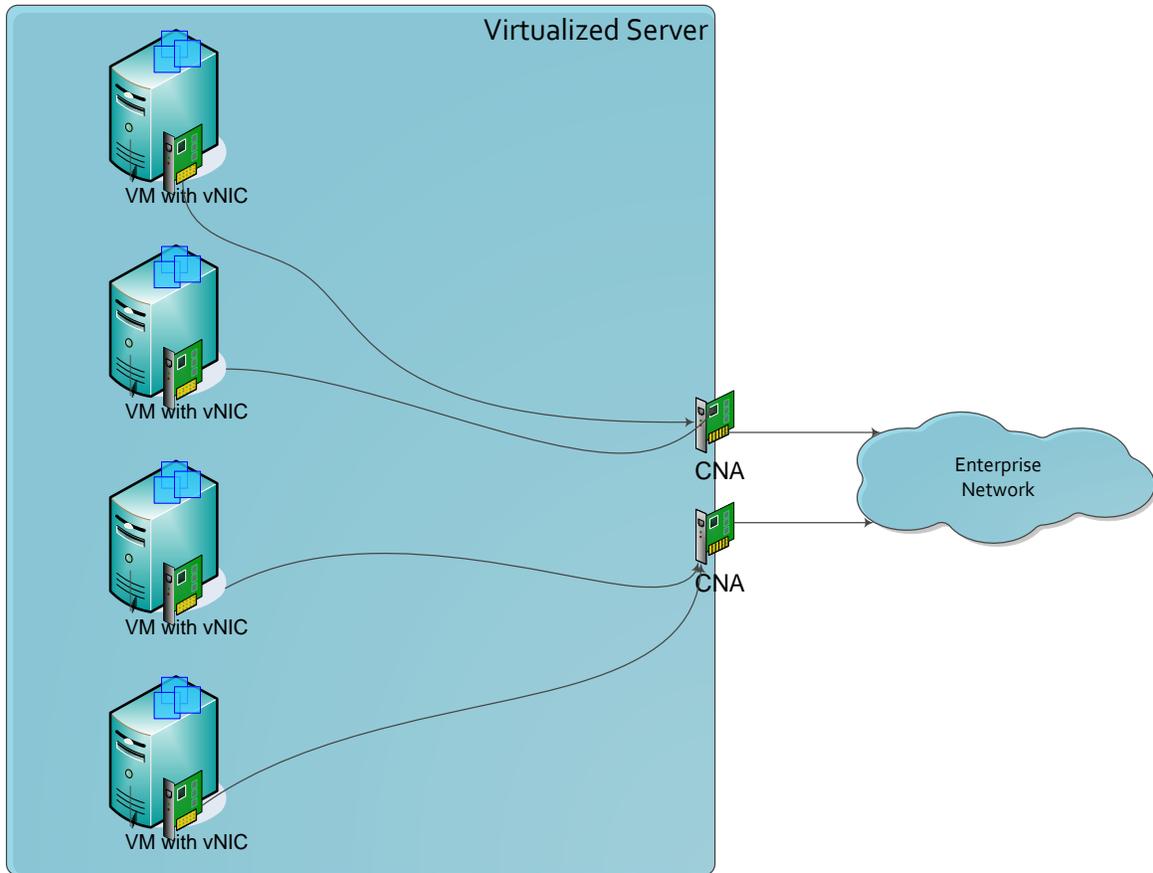


Figure 42. Virtualized topology showing the use of CNA's with SR-IOV technology

Conclusions

Since the mechanisms for these two virtualization models vary dramatically from the typical bare metal configuration, it is important to measure the performance of each.

Appendix A: Emulex OneConnect UCNA Platform

The Emulex OneConnect™ UCNA is a single-chip, high-performance 10Gbps Ethernet product that provides server connectivity for network and storage traffic with one multi-function server adapter. The payoff is a rapid return on investment (ROI) based on:

- Lower capital expense (CapEx)—single 10GbE adapter replaces storage adapter and multiple 1GbE NICs
- Lower operational expense (OpEx)—one management console for network and storage
- Higher virtualization ratios—protocol offloads save CPU cycles and optimize memory usage

Unlike first generation CNAs that only provide FCoE convergence, the Emulex OneConnect UCNA technology provides optimized performance for TCP/IP, FCoE and iSCSI protocols. Accelerators/offload engines for all supported protocols allow Emulex OneConnect UCNAs to deliver maximum performance, regardless of the mix of network traffic.

The Emulex OneConnect UCNA provides a unique advantage that enables a drastic reduction in the number of server hardware configurations and provides the commonality across several I/O connectivity architectures, thus directly contributing to the efficiency of the data center procurement and operations.

As data centers move toward a virtualized pool of resources, this commonality also enables virtualized servers to be easily migrated across hardware platforms (e.g., from rack servers to blades).



Figure 43. Emulex OneConnect OCe10000 UCNA

The OneConnect product family includes:

- OCe10102-N: 10Gb Ethernet Adapter

- OCe10102-I: 10Gb iSCSI Converged Network Adapter
- OCe10102-F: 10Gb FCoE Converged Network Adapter

OneConnection for Network and Storage

10 Gigabit Ethernet

10 Gigabit Ethernet (10GbE) has become the standard backbone for enterprise class networks. New multi-core server platforms, server virtualization and high-demand applications such as database clusters, video on demand and VOIP are all driving the need for high bandwidth. Emulex OneConnect UCNAs deliver optimized 10GbE performance that meets these requirements with support for stateless TCP offload and TCP chimney offload.

Enhanced Ethernet (EE)

Ethernet Enhancements—also known as CEE (by IBM), DCE (by Cisco) or DCB (IEEE)—enable lossless Ethernet behavior. Emulex OneConnect UCNAs allow fibre channel and Ethernet traffic to be carried over a converged lossless enhanced Ethernet network, eliminating separate infrastructures for storage and networking, and delivering significant cost savings. Emulex UCNAs also support TCP/IP offloads and bandwidth allocation to maximize performance.

Fibre Channel over Ethernet

Emulex OneConnect UCNAs support FCoE, seamlessly converging fibre channel and Ethernet while leveraging mature fibre channel management tools and processes. By consolidating traffic, FCoE simplifies network management and significantly reduces total cost of ownership (TCO).

iSCSI

Emulex OneConnect is an optimum solution for iSCSI storage deployments. Support for iSCSI offload provides performance that is superior to iSCSI solutions based on software initiators and standard NICs.

Boot from SAN

OneConnect UCNAs are supported with fibre channel and iSCSI boot from SAN. Universal boot capability allows the appropriate boot environment to be automatically selected for any OS.

Simplified Management

OneCommand Manager Application

The OneCommand™ Manager application enables management of Emulex OneConnect UCNAs and LightPulse® host bus adapters (HBAs) and CNAs throughout the data center from a centralized management console. The OneCommand manager application provides a

TEST CASE: BGP RIB-IN CONVERGENCE TEST

graphical user interface (GUI) and a scriptable command line user interface (CLI) that drives administration efficiency and business agility. Centralized discovery, monitoring, reporting and management of both local and remote adapters can be done from a secure remote client. In-depth management capabilities include remote firmware and boot code upgrades, beaconing, statistics and advanced diagnostics.

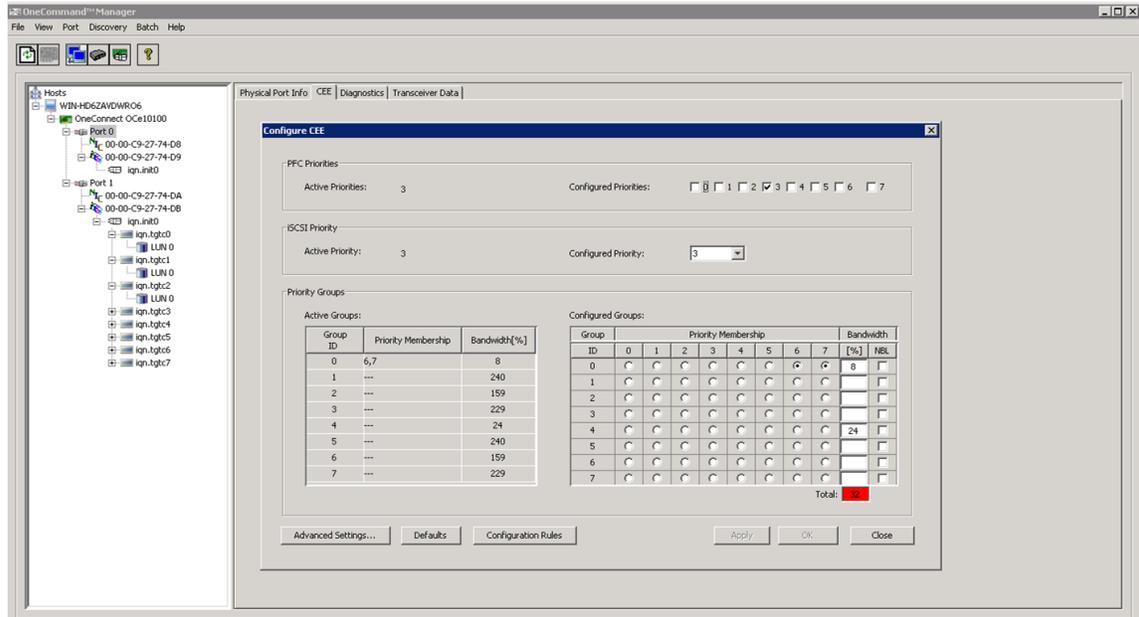


Figure 44. OneCommand Manager configures the traffic priorities and per Priority flow control

Quality of service (QoS)

Using the OneCommand Manager application, administrators can allocate portions of the 10Gb Ethernet bandwidth to network or storage traffic. Using OneConnect support for virtual ports, storage-based QoS can be individually managed for virtual machines.

Streamlined installation

A single installation of drivers and applications for Windows servers eliminates multiple reboots and ensures that each component is installed correctly and OneConnect UCNAs are ready to use.

Emulex OneConnect UCNAs: Highest Performance and Reliability

Enterprise-ready

Leveraging nine generations of advanced, field-proven fibre channel HBA technology, the Emulex OneConnect family meets the robust interoperability and reliability requirements of corporate data centers. The design leverages proven Emulex enterprise class drivers, firmware and hardware architectures, while delivering sophisticated capabilities required to manage multiple types of data flow concurrently with optimized performance.

More virtual machines per server with VEngine™ technology

Higher CPU efficiency and the highest virtual port count per VM enables Emulex OneConnect UCNAs to offload processing from the hypervisor and support more VMs per server. Dynamic NPV-based QoS bandwidth allocation based on individual VMs gives data centers the flexibility and performance they need to maximize server consolidation.

Greener data centers

The Emulex OneConnect UCNA platform delivers industry-leading performance and scalability per watt, reducing requirements for power and cooling. Protocol offload enables efficient use of computing resources, supports more VMs per CPU, and reduces the number of servers required to support data center demands.

Advanced error checking

End-to-end data protection with hardware parity, CRC, ECC and other advanced error checking and correcting ensure that data is safe from corruption.

Key Features: OneConnect OCe10000

One platform for every network connection

- 10Gbps Ethernet, iSCSI and FCoE offloads
- Simplifies I/O hardware choice for IT managers

Superior performance

- Hardware accelerators for all protocols
- High performance with high CPU efficiency

Energy efficient design

- Industry-leading performance per watt
- Complements data center “green” initiatives

TEST CASE: BGP RIB-IN CONVERGENCE TEST

Easy to deploy and manage with OneCommand™ Manager application

- One management console, many protocol services
- Integrated management of UCNAs and HBAs
- Over 7 million ports administered with Emulex management software

Key Benefits: OneConnect OCe10000

Optimized for convergence

- Maximizes I/O consolidation with high-performance 10GbE ports
- One network infrastructure reduces capital expenses (CapEx)
- One management console reduces operational expenses (OpEx)
- Leverages existing IT investments

Optimized for server virtualization with VEngine™ technology

- More virtual machines (VMs) per server with TCP/IP, iSCSI and FCoE offload
- VM-specific performance, quality of service (QoS) and management

Enterprise-Ready

- Hardware parity, CRC, ECC and other advanced error checking
- Backed by field proven Emulex reliability and support

Standards

- ANSI INCITS T11 FC-BB-5 2.0, FC-PI-2, FC-GS-4, FC-TAPE, and FCP-3
- PCI Express base spec 2.0, PCI Bus Power Management Interface, rev. 1.2, Advanced Error Reporting (AER)
- IEEE 802.3ae (10Gb/s Ethernet), 802.1q(VLAN), 802.1p(QoS/CoS), 802.3ad (Link Aggregation), 802.3x (Flow Control), 802.3ap (Backplane Ethernet for mezzanine models)
- Enhanced Ethernet (draft): Enhanced Transmission Selection, ETS(P802.1Qaz); Priority-based Flow Control, PFC (P802.1Qbb); Data Center Bridging Capabilities eXchange Protocol, CIN-DCBX and CEE-DCBX (P802.1Qaz)
- PHP hot plug-hot swap Architecture
- Dual-channel, 10Gb/s Ethernet Link speed
- PCIe Express 2.0 (x8, 5GT/s), MSI-X support
- Integrated data buffer and code space memory

TEST CASE: BGP RIB-IN CONVERGENCE TEST

Ethernet Features

- IPv4/IPv6 TCP, UDP checksum offload; Large Send Offload(LSO); Large Receive Offload; Receive Side Scaling (RSS); IPV4 TCP Chimney Offload
- VLAN insertion and extraction
- Jumbo frames up to 9000 Bytes
- Preboot eXecutive Environment (PXE) 2.0 network boot and installation support
- Interrupt coalescing
- Load balancing and failover support including adapter fault tolerance (AFT), switch fault tolerance (SFT), adaptive load balancing (ALB), teaming support and IEEE 802.3ad

fibre channel over Ethernet (FCoE) Features

- Common driver for UCNAs and HBAs
- N_Port ID Virtualization (NPIV)
- Support for FIP and FCoE Ether Types
- Fabric Provided MAC Addressing (FPMA) support
- 1024 concurrent port logins (RPIs) per adapter
- active exchanges (XRIs) per port

iSCSI Features

- Target discovery methods
- Authentication modes
- INT 13 Boot

Comprehensive OS Support

- Windows Server 2008, Windows Server 2003
- Red Hat Enterprise Linux Server
- Novell SUSE® Linux Enterprise Server
- VMware ESX

Hardware Environments

- x86, x64 processor family

Interconnect

Copper

- SFP+ Direct Attached Twin-Ax Copper interface
- Standards compliant passive and active copper cables supported up to 5m

Optical

- Optics: 10GBASE-SR short wave lasers with LC type connector

Physical Dimensions

- Low profile with standard bracket (low profile bracket available)

TEST CASE: BGP RIB-IN CONVERGENCE TEST

Power and Environmental Requirements

- Volts: +3.3, +12
- Operating temperature: 0° to 55° C (32° to 131° F)
- Storage temperature: -40° to 70° C (-40° to 158° F)
- humidity: 5% to 95% non-condensing
- UL recognized to UL 60950-1
- CUR recognized to CSA22.2, No. 60950-1-07
- Bauart-certified to EN60950-1
- FCC Rules, Part 15, Class A
- ICES-003, Class A
- EMC Directive 2004/108/EEC (CE Mark)
- EN55022, Class A
- EN55024
- Australian EMC Framework (
- VCCI (Japan), Class A
- KCC (Korea), Class A
- BSMI (Taiwan), Class A
- EU RoHS Compliant (Directive 2002/95/EC)
- China RoHS Compliant

Contact Ixia

Corporate Headquarters
Ixia Worldwide Headquarters
26601 W. Agoura Rd.
Calabasas, CA 91302
USA
+1 877 FOR IXIA (877 367 4942)
+1 818 871 1800 (International)
(FAX) +1 818 871 1805
sales@ixiacom.com

Web site: www.ixiacom.com
General: info@ixiacom.com
Investor Relations: ir@ixiacom.com
Training: training@ixiacom.com
Support: support@ixiacom.com
+1 877 367 4942
+1 818 871 1800 Option 1 (outside USA)
online support form:
<http://www.ixiacom.com/support/inquiry/>

EMEA
Ixia Technologies Europe Limited
Clarion House, Norreys Drive
Maiden Head SL6 4FL
United Kingdom
+44 1628 408750
FAX +44 1628 639916
VAT No. GB502006125
salesemea@ixiacom.com

Renewals: renewals-emea@ixiacom.com
Support: support-emea@ixiacom.com
+44 1628 408750
online support form:
<http://www.ixiacom.com/support/inquiry/?location=emea>

Ixia Asia Pacific Headquarters
21 Serangoon North Avenue 5
#04-01
Singapore 5584864
+65.6332.0125
FAX +65.6332.0127
Support-Field-Asia-Pacific@ixiacom.com

Support: Support-Field-Asia-Pacific@ixiacom.com
+1 818 871 1800 (Option 1)
online support form:
<http://www.ixiacom.com/support/inquiry/>