Black Book

ixia

Edition 10

Converged Data Center

http://www.ixiacom.com/blackbook

June 2014

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, please contact us at ProductMgmtBooklets@ixiacom.com.

Your feedback is greatly appreciated!

Contents

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.

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 Data Center

Test Methodologies

The tests in this booklet detail methodologies to verify the performance of converged data center networking technologies. Subjects include data center Bridging (DCB), fibre channel over Ethernet (FCoE), fibre channel, storage I/O performance, layer 2 multipath (L2MP) fabrics, and virtualization.

Introduction to Converged Data Center

The growth of cloud computing, server virtualization, applications such as enterprise resource planning (ERP), customer relationship management (CRM), and individual productivity applications are creating explosive data growth. Besides, data mining applications and social networking sites are growing and gaining popularity everywhere.. All of these emerging technologies and usage patterns are causing aunthetic and abundant data.

These trends are driving the need for highly flexible, scalable, and low latency data storage architectures. The resulting requirements for these innumerable applications and the exponential growth in data directly impact the data center. As a result, data centers are expanding intensely in size and complexity and are becoming much more costly to build and maintain. Some data centers are so large and require so much power, that many regions of the United States have moratoriums on the construction of data centers. These developments of the digital age are causing data center architects to take notice.

Over the last several years, there are discussions within the data center community to focus on lower cost, lower power, and simpler designs. With respect to data center infrastructure, better power and cooling designs are being architected. In the computing sphere, blade server design has started to take hold – accentuating their lower cost and power consumption. On the networking front, there has also been a very significant effort undertaken to optimize the infrastructure to improve raw capacity, increase redundancy, expand scalability, and reduce latency. To accomplish this expansive and ambitious goal, new technologies have been introduced across multiple domains from the server virtualization to converged storage to cloud federation.

To truly realize a converged data center to optimally support cloud-based services, each of these new technologies must be thoroughly tested in silos, then validated end-to-end as a whole system, prior to deployment. This BlackBook is structured to guide the user through a set of test methodologies in each of the areas that make up the end-to-end converged data center infrastructure. Future editions of this BlackBook will continue to add more advanced test scenarios.

Introduction to Data Center Bridging and FCoE

Converged storage technologies aim to unify various networks used in the data center to reduce both operational complexity and upfront cost. The migration of Storage Area Network (SAN) and Local Area Network (LAN) traffic onto a single transport benefits data centers in many ways. It reduces the initial capital outlays for equipment and lowers the cost of maintaining and cooling data centers. The purpose is achieved by reducing the amount of cabling, adapters, and fabric switches required to support both network traffic and storage traffic.

Fibre Channel (FC) has been the de-facto standard for SANs in the data center. FC provides high reliability, performance, and network intelligence for low latency, high bandwidth applications. At the same time, Ethernet has been the standard for network traffic in the data center just as everywhere else. Ethernet's ubiquity and massive growth has also brought about an overall lower cost for Ethernet infrastructure. In many data centers, there is also a third network for inter-processor communication that is used for high performance clustering. These three different networks require different server cards, different cables, and different expertise. The requirement for three different types of cards, which with redundancy can mean a total of six interface cards, in each associated server resulted power hungry boxes. This requirement also conflicted with the computing industry's effort to transition to lower power blade servers. The requirement for three different types of cables for each of the thousands of servers in the data centers created a massive cabling requirement. Therefore, a consolidation and convergence of fibre channel and Ethernet into a unified framework is a critical, step towards energy efficient high performing data centers.

Fibre Channel over Ethernet Technology

To date, fibre channel has dominated data centers as the technology of choice for storage area networks. However, fibre channel is built on an independent infrastructure that is separate from common Ethernet-based enterprise data communication networks. With the significant growth of server virtualization and 10 Gbps Ethernet's increasing availability and its resonating potential for CAPEX/OPEX savings, a strong push has emerged to consolidate data center technologies into a converged I/O fabric using Ethernet as the common transport media. To truly gain industry acceptance, consonant standards are necessary. To address this need, T11 has defined standards and open proposals to encapsulate fibre channel frames over Ethernet 802.3 MAC frames, thus seamlessly multiplexing storage traffic across 802.1 bridged networks. This is the concept of fibre channel over Ethernet (FCoE).

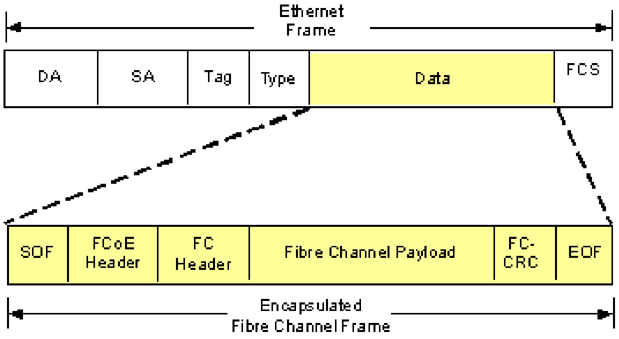

Figure 1. FCoE frame format

Simply providing a bridged network to transport FCoE frames is insufficient in data center environments. This is because fibre channel is a lossless technology, while Ethernet is a lossy transport built on the fundamental of low cost best-effort data communication. Enhancements must be incorporated into today's bridges before data centers can deploy FCoE commercially, and more importantly, reliably. To address this need, the IEEE 802.1working group initiated the Data Center Bridging (DCB) Task Group to define protocols to utilize Ethernet for lossless transmission. Congestion notification (802.1Qau), enhanced transmission selection (802.1Qaz), and priority-based flow control (802.1Qbb) are three key working documents in the IEEE 802.1 for DCB that work cohesively to achieve lossless Ethernet.

The Ethernet flow control mechanism (IEEE 802.3x) provides a means of avoiding traffic congestion at the link level regardless of the type of traffic flows or class of packets being forwarded. Ethernet's legacy pause option causes all traffic on a link to stop. Link sharing, however, is critical for I/O consolidation converged data center. Link sharing operations must manage traffic flow rates, traffic queues and latency based on traffic types.

PFC pauses traffic based on user priorities or classes of service. A physical link is divided into eight virtual links (Priority-based Flow Control) with PFC providing the capability to use Pause on a single virtual link without affecting traffic on the other virtual links. Enabling Pause on a per-user-priority basis allows administrators to create lossless links for traffic requiring no-drop service, such as fibre channel over Ethernet, while retaining packet-drop congestion management for LAN traffic.

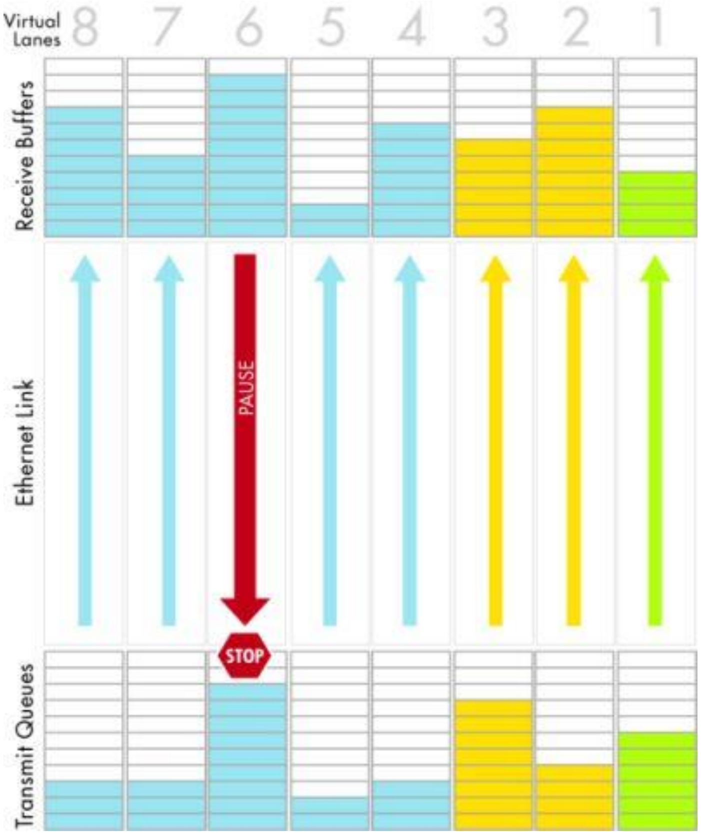

Figure 2. Priority-based Flow Control

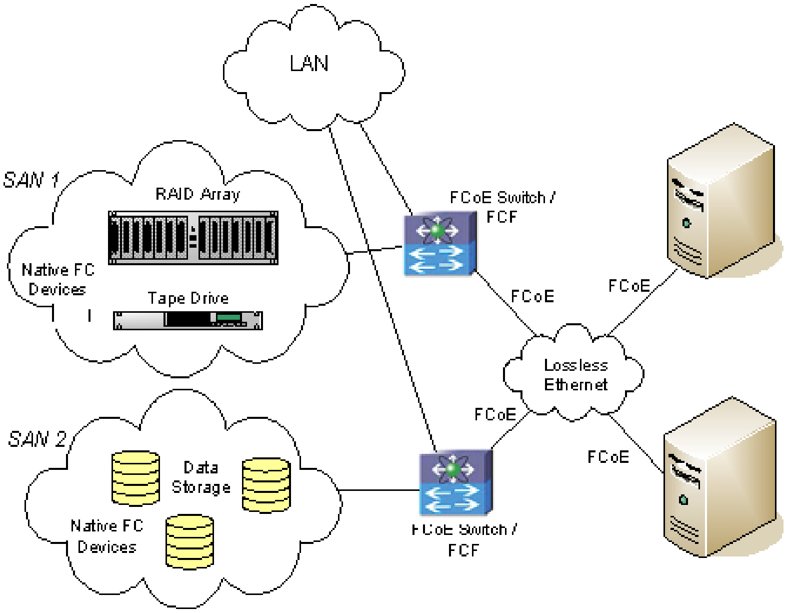

Figure 3. Sample FCoE/native FC data center

The success of FCoE is rooted in its congruous functional model architecture, in which the original fibre channel stacks from FC-2 onward remain unchanged. Simply, FC-0 and FC-1 are replaced by the 802.3 Ethernet PHY and MAC. With this concordant approach, FCoE inherits all link initialization and maintenance capabilities without modifications to the well-defined and mature fibre channel operations – a key to successful migration. When an FCoE client (ENode) connects to an FCoE Forwarder (FCF), the ENode's Virtual N_Port (VN_Port) logs in to the fabric by transmitting a fabric login (FLOGI) message, encapsulated in a standard 802.3 MAC frame, to the FCF's Virtual F_Port (VF_Port), as if it was a native fibre channel link. On receiving the FLOGI, the FCF validates its resources and, if conditions allow, returns an LS_ACC, encapsulated in an 802.3 MAC frame, back to the ENode's VN_Port.

While this FCoE functional model operates well, the addition of dynamic maintenance capabilities enhances the manageability of fibre channel, which is a connection oriented point-to-point technology over Ethernet, which is a connectionless multipoint access technology. The FCoE initialization protocol (FIP) delivers this enhancement by providing dynamic discovery, initialization, and maintenance capabilities to FCoE. With FIP, ENodes and FCFs can automatically discover each other, initialize a link, and perform maintenance functions throughout the lifetime of an FCoE connection.

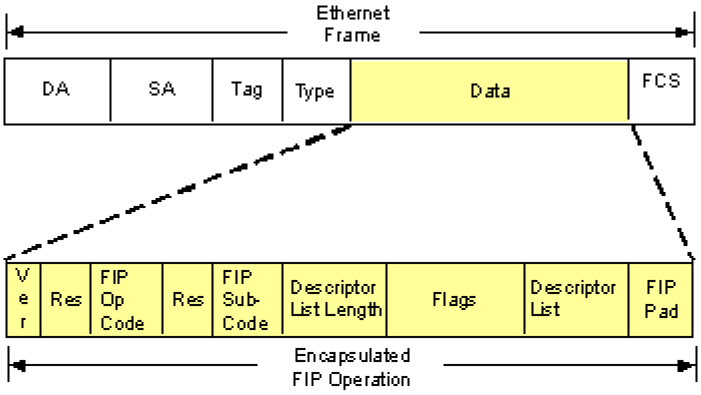

Figure 4. FIP frame format

While the FCoE and FIP protocols establish a promising foundation, it is just the foundation. Two additional areas require attention to ensure a successful migration to converged converged data center: FCoE/native FC interoperability and storage/LAN traffic multiplexing.

Fibre Channel Technology

Similar to the OSI 7 layer model, or the more practical TCP/IP layer model, fibre channel is also defined in several layers of functional blocks. Fibre channel is built on five distinct levels, from FC-0 to FC-4. Above FC-4 is where the upper layer protocols such as SCSI sit.

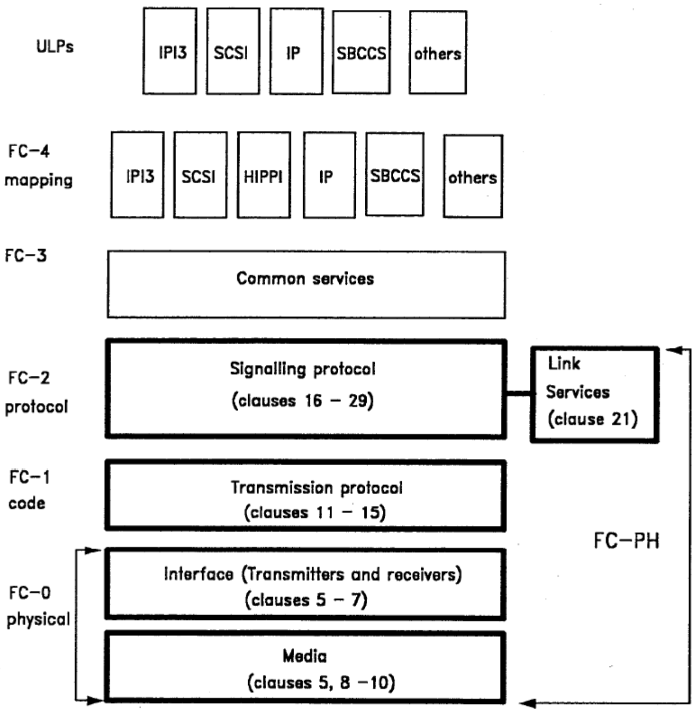

Figure 5. Fibre Channel structure

FC-0: media and interface

Several physical layer requirements are specified by FC-0, from operating speeds to optical/electrical characteristics to cable types. The most commonly used connector type is the SFP/SFP+ transceiver, which provides an interface for LC fiber cables, most commonly at 850 nm (multi mode). The serial data stream transmission speeds vary from, most commonly, 1, 2, 4, and 8 Gbps[1] to date. Recently, 16 Gbps has also been introduced mostly in the R&D phase at the time of this writing. The transmitter takes the 10-bit encoded transmission characters from FC-1, and serializes it into a 1/0 binary signal. The receiver digitizes the incoming analog binary signal into a 10-bit character, and passes it up to FC-1 for decoding. In addition, the standard imposes a bit error rate of 10^{-12} BER.

FC-1: transmission encoding/decoding

[1] The 'base' transmission rate is 1.0625 Gbps, where transmitters can transmit at the following multiples of the base rate: 1/8, 1/4, 1/2, 1, 2, 4, 8, and 16. So, the transmission rate of '4 Gbps fibre channel' is 4.25 Gbps. In addition, FC-PH specifies the transmission speed in baud rates (for example, 1.0625 Gbaud), which happens to be the same as bps (bits per second) because the modulation technique results in 1 bit per baud.

Similar to Gigabit Ethernet (except 1000BASE-T), fibre channel also uses an 8b/10b line coding scheme. While different line code schemes are used by several different speeds and flavors of Ethernet, fibre channel FC-1 uses the 8b/10b line code to map 8-bit symbols to 10-bit encoded symbols. For the same reasons as Ethernet, such encoding allows the bit stream to remain DC-balanced over time, minimizing low frequency content and thereby minimizing signal distortion across AC-coupled links. Such design also allows the clock to be locally recoverable as the line code scheme ensures sufficient state changes over a stream of bits.

In standardization, the notation convention Zxx.y is used to describe the content of the unencoded FC-1 transmission character. Z identifies whether the character carries data information, noted by D, or special control information, noted by K. xx identifies the decimal value of bit 0 (LSB) to bit 5 of the character, and y identifies bit 6 to 8 (MSB) of the character. The set of predefined transmission characters (four-character Words) that perform special functions and carry special meanings are called Ordered Sets. Ordered Sets are noted by K28.5 followed by three Dxx.y characters, which will not identify any valid data characters, and they identify frame boundaries (that is, SOF, EOF), but also signal flow control primitives as well as other link states. The information the Ordered Sets carry are very specific, in which it specifies not only the start or end of a frame, but also the class of the frame, as well as whether it is the first frame in a sequence or not. As for the Primitive Signals, there are two defined: Idle, to keep the link initialized while no data is on the wire, and R_RDY to indicate the local Port is ready to receive another frame over the link. The Primitive Sequences for the most part signals the local link status but could also engage in the Link Reset protocol, and a few other Arbitrated Loop specific operations.

FC-2: link services and protocols

Fibre channel nodes (servers, disk drives) communicate with each other through local N_Ports. As a point-to-point connection-oriented technology designed to transport storage transactions with lossless delivery, an N_Port goes through several login processes with the next-hop fabric switch as well as the far end N_Port before storage transactions (for example, SCSI transactions) can take place. Throughout the login processes (that is, FLOGI/NPIV FDISC, PLOGI, PRLI), an FCID is assigned, credits are reserved, class of service is determined, and internal images are created to allow a process on the source N_Port (for example, SCSI Initiator) to communicate with a process on the destination N_Port (for example, SCSI Target), with the assurance that network resource is available over the end-to-end path to carry the information being requested and served.

When an N_Port is first connected to a fabric switch port (F_Port), it performs a fabric login (FLOGI) by using the Extended Link Services such that a set of resources on the fabric can be set aside for the N_Port, and the operating Service Parameters such as credits and data field receive size are established. Many of the Service Parameters are critical to the operations of the link, perhaps the two most important parameters are buffer-to-buffer credits and the set of timeout values. These two parameters ultimately dictate the performance of the fabric and the error recovery abilities of the fabric.

In data communications in general, two common types of credits management exist—source based and destination based. In source based, the source requests the amount of resource it needs to transmit the data it has. If accepted, the fabric and the destination N_Port will need to ensure that it guarantees the requested resources. In destination based, the receiver indicates the amount of receiver buffers it has available to receive data, and the receiver signals its ability to receive the next frame after clearing out the most recently used receiver buffer. In Fibre Channel, Class 3[2] devices exclusively use the destination based flow control mechanism, called the buffer-to-buffer credit. In buffer-to-buffer credit, during the FLOGI request and reply process, the N_Port and the F_Port each indicate its local BB_Credit to its peer, which becomes the 'allocated BB_Credit' for the peer. Each time a local buffer is cleared after receiving a frame, the local port transmits an R_RDY primitive to inform its peer that it is ready to receive the next frame. To ensure that it is aware of the peer's ability to receive additional frames, the local port also maintains a local BB_Credit_CNT, which tracks the number of unacknowledged frames awaiting the R_RDY Primitives. As such, a local port always attempt to maintain BB_Credit_CNT between 0 and the allocated BB_Credit. BB_Credit_CNT is reset to zero at the end of a Fabric Login or relogin, or any Primitive Sequence Protocol.

The described flow control method operates well and efficiently assuming the transmitted frames are not lost, and the R_RDY primitives are not lost. In any network disruption, frame/primitive loss is possible. In such scenarios, the result is a mismatch between what the local port and the remote port in terms of number of credits available. To recover from this situation, the Fibre Channel standard also defines procedures to verify and recover lost credits after a number of frames/R_RDYs are transmitted by the local port. The BB_Credit Recovery process is achieved by two primitives, BB_SCs and BB_SCr, and three parameters, BB_SC_N, BB_RDY_N, and BB_FRM_N. From the perspective of the data frame transmitter, the BB_SCs primitive is transmitted by the local port when $2^{BB\_SC\_N}$ frames have been transmitted since the previous BB_SCs primitive, and the BB_SCr primitive is transmitted by the remote port when $2^{BB\_SC\_N}$ R_RDY primitives have been transmitted since its previous BB_SCr primitive. BB_SC_N is a configurable parameter that is exchanged during the FLOGI process, conveyed in the Common Service Parameters fields. BB_RDY_N is a local counter that increments every time an R_RDY primitive is received, and the BB_FRM_N is a local counter that increments every time a data frame is received. As such, the number of BB_Credits lost locally and remotely can be calculated as follows:

On receiving a BB_SCr: BB_Credits lost locally = $(2^{BB\_SC\_N} - BB\_RDY\_N)$ modulo $2^{BB\_SC\_N}$

On receiving a BB_SCs: BB_Credits lost remotely = $(2^{BB\_SC\_N} - BB\_FRM\_N)$ modulo $2^{BB\_SC\_N}$

If the BB_Credits lost locally is non-zero, BB_Credit_CNT is decremented by BB_Credits to correctly reflect the actual number of credits available at the remote port, and BB_RDY_N is reset to zero.

[2] In Fibre Channel, device Class distinguishes the level of data delivery integrity required for an application. Class 2 requires a device to report situations where a frame is not delivered (through F_BSY or F_RJT), and acknowledges each frame received (through ACK). Class 3 does not report or acknowledge these conditions.

On accepting the FLOGI, the N_Port is assigned an FCID. Once logged in, the N_Port registers a set of information such as WWN with a Name Server (typically on the fabric switch itself). The N_Port can later query the Name Server to retrieve information about other N_Ports on the fabric for further communication. Next, for the local N_Port, for example, a server, to communicate with a remote N_Port, for example, a disk drive, the N_Port performs a port login (PLOGI) with the remote N_Port. During this process, information such as credit and operating parameters are established, providing the foundation for end to end communication between the N_Ports. Finally, for an upper layer protocol like SCSI to read and write data from and to the disk drive, the local N_Port on the server must perform a process login (PRLI) with the disk drive N_Port, such that an image behind each N_Port is established, allowing application specific data to be multiplexed over the end-to-end N_Port-to-N_Port link.

FC-3: Common Services

FC-3 specifies a set of common services and functions to be used across multiple N_Ports on a single FC Node.

FC-4: Upper Level Protocol (ULP) Mapping

Upper level protocols such as SCSI (Small Computer Systems Interface) exchange information with remote endpoints using services provided by FC-2. FC-4 provides such functions by mapping ULP command sets (for example, SCSI READ) to fibre channel ports.

ULP: SCSI over Fibre Channel

Small Computer Systems Interface (SCSI) was drafted in the early 1980s when computers and computer peripherals were far from mainstream. During that nascent computer age, the definition fit very well. Since then, SCSI has evolved to keep pace with the ever-increasing architectures, options, performance, and stability of computing systems that are now the expected norm.

SCSI, in its barest form, is a standard mechanism for connecting peripherals (disks, tape drives, CD-ROMs) to a computer through an SCSI controller. It is now an extensive command set that covers an enormous range of topologies and interfaces from parallel SCSI, USB, and FireWire to Fibre Channel.

SCSI provides a method for allowing access to multiple devices on the same interface simultaneously. One device on a SCSI bus can communicate directly to another without going through the DMA controller, using only the SCSI controller. Thus, a SCSI hard drive can communicate directly with a SCSI CD ROM device. Part of what makes the SCSI implementations so versatile is the sprawling command set that supports it as shown in Table 1. ANSI SCSI groups these commands under the following designations:

SCSI Commands

Command	Description
SCSI Block Command (SBC)	This is designed for working with devices that access data in a random (non-sequential) format. There is also a 'reduced' subset of commands related to this (RBC).
SCSI Stream Command (SSC)	SCSI stream commands are used to access sequential data devices such as tape drives.
SCSI Controller Command (SCC)	SCSI controller commands are used by devices to communicate with SCSI RAID arrays.
Multimedia Commands (MMC)	Multimedia commands are for devices such as DVD drives.
SCSI Graphics Commands (SGC)	Graphic commands are used to communicate with printers.
Media Changer Commands (SMC)	For devices such as CD ROM jukeboxes.
Enclosure Services Commands (ESC)	SCSI commands to communicate with intelligent enclosure devices.
Object-based Storage Commands (OSD)	SCSI commands for object based devices.
Management Server Commands (MSC)	SCSI commands for management services of SCSI devices.

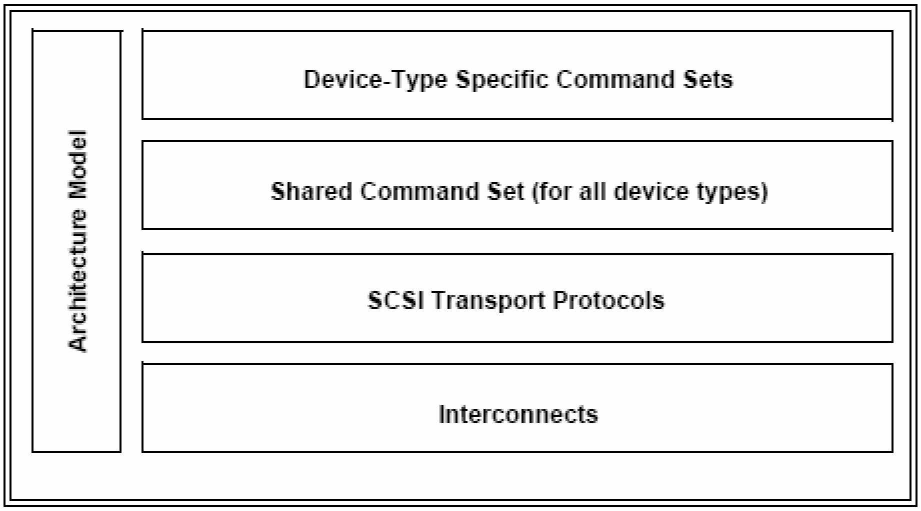
Architecture Model
Device-Type Specific Command Sets
Shared Command Set (for all device types)
SCSI Transport Protocols
Interconnects

Figure 6. Fibre Channel protocol layers

Service interfaces between SCSI entities are represented by the client-server model shown in the following figure. The dashed horizontal lines denote the request and corresponding response from the perspective of the client and server. The actual transaction path through the service delivery subsystem is shown by the solid lines.

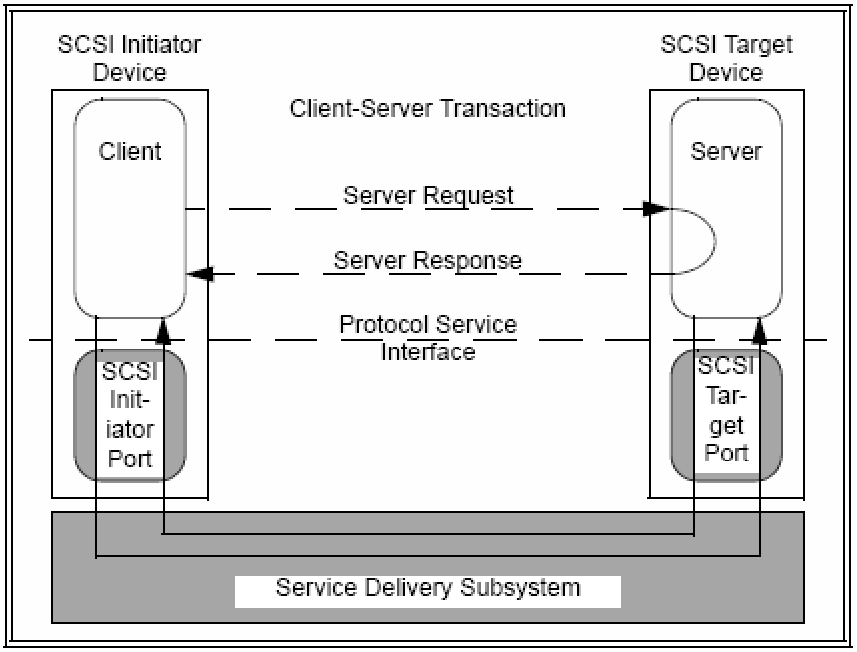

Figure 7. SCI client/server model

This architectural enhancement introduced the concept of a distributed network like fibre channel protocol to expand the service delivery subsystem into a fabric and introduced the capability to have the ports, called node port (N_Port) that connects to either another N_Port or another fabric port (F_Port).

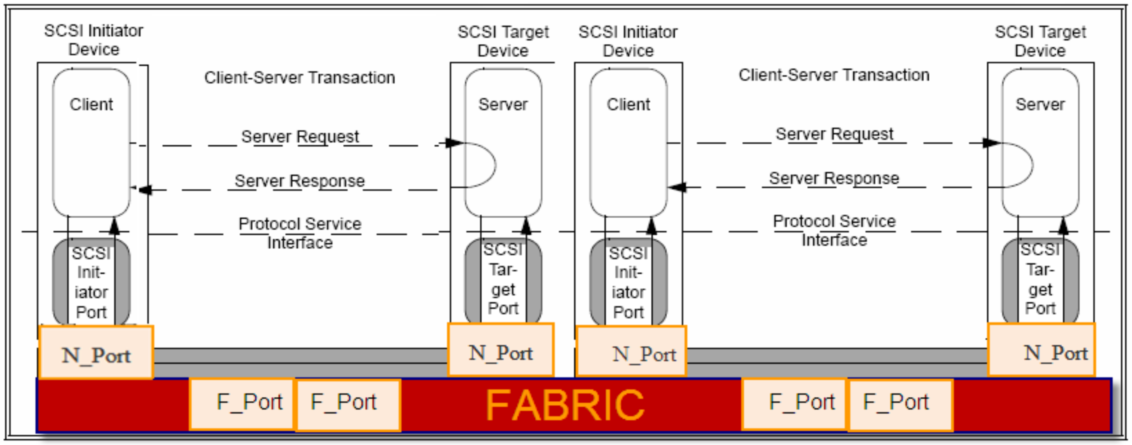

Figure 8. Node and fabric ports

In this way, protocol adoption in the SCSI protocol allowed the fibre channel protocol to adopt the SCSI protocol and create a fabric network as shown in the preceding figure with a vastly increased number of devices in the network that have the ability to communicate with each other simultaneously.

As any historic observer of telecom evolution will testify, successful adoption of a new technology requires a conscious effort to support backward compatibility to aid a smooth migration. For the converged data center, a conscious decision was made to allow FCoE and native fibre channel links to coexist. Fibre channel forwarders (FCFs) serve this function—bridging these two worlds that were traditionally disparate.

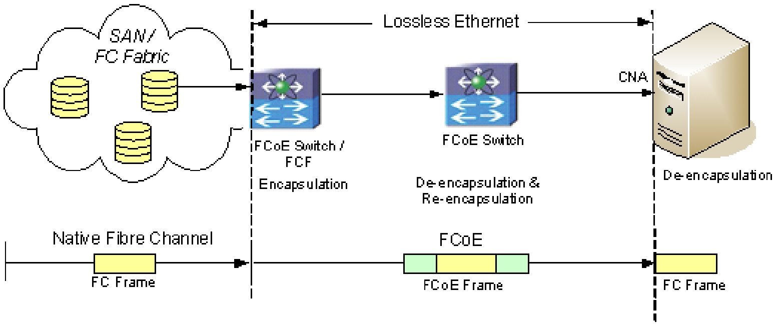

Figure 9. Bridging FCoE and native FC

Finally, to achieve an architecture capable of delivering both storage (for example, SCSI) and LAN traffic properly, the network must support each segment's traffic requirements when resources are abundant, and then instantaneously enforce strict traffic policing when resources are limited.

To summarize, several protocols and technologies must coexist harmoniously to structure a robust Ethernet based data center. Individually, these sectors fall into five categories: FCoE/FIP, Lossless Ethernet, fibre channel, end-to-end (FCoE to FC) converged traffic, and I/O performance. Collectively, the result is a converged enhanced/converged data center.

Test Case: DCBX Functionality

Overview

DCBX enables FCoE and lossless Ethernet devices to advertise, discover and configure data center bridging parameters with their peers. Such capability is beneficial but also critical in a converged data center network as it guarantees that an end-to-end network path will be set up with the necessary storage and LAN traffic resources.

The exchange of DCBX information is described in IEEE 802.1AB LLDP (Link Layer Discovery Protocol), which is a layer 2 protocol that allows bridges to advertise local parameters and information. While LLDP accommodates multiple LAN neighbors, DCBX restricts it such that only one neighbor is accepted at any given time. When a lossless Ethernet capable device such as a CNA or FCF connects to the network, it transmits LLDPDUs periodically from its LLDP-enabled ports. If DCBX is also enabled, the device's LLDPDUs include DCBX Control TLVs as well as supported feature TLVs. When the device receives an LLDPDU with DCBX TLVs, it checks the feature against its state machines, and decides to use local or peer configuration accordingly. After a feature reaches a stable state, the application that relies on that feature can begin its operations.

Objective

The objective of this test is to verify the ability of a lossless Ethernet capable device to advertise, discover, and configure the following DCBX TLVs: Priority Group TLV, PFC TLV and, if supported, the FCoE TLV. The presence of LLDP Mandatory TLVs and the DCBX Control TLV is also verified, as well as the advertisement and acknowledgment of the sequence number in the DCBX Control TLV.

Setup

One Ixia port is used in this test.

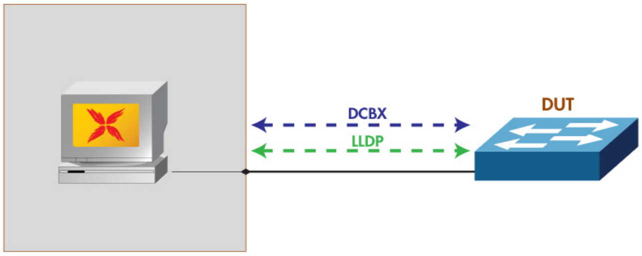

Figure 10. DCBX functionality test topology

Step-by-Step Instructions

1. Reserve one Ixia port.

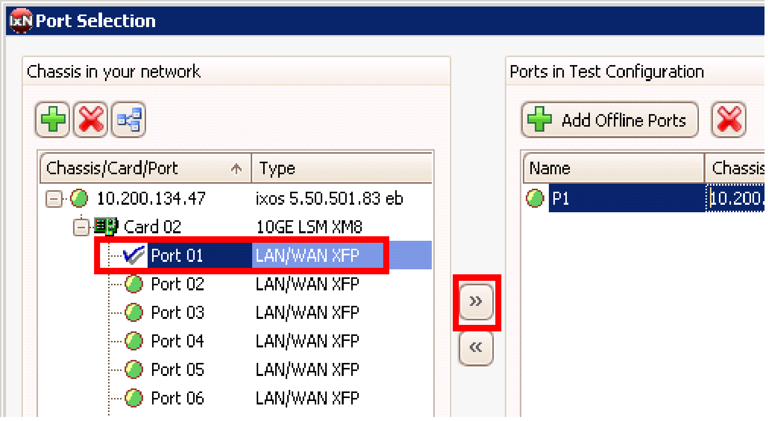

Port Selection

Chassis in your network

Chassis/Card/Port	Type
10.200.134.47	ixos 5.50.501.83 eb
Card 02	10GE LSM XM8
Port 01	LAN/WAN XFP
Port 02	LAN/WAN XFP
Port 03	LAN/WAN XFP
Port 04	LAN/WAN XFP
Port 05	LAN/WAN XFP
Port 06	LAN/WAN XFP

Ports in Test Configuration

Name	Chassis
P1	10.200.

Figure 11. DCBX functionality – port selection

2. Open the IxNetwork **Protocol Wizards** window, and invoke the **DCBX** wizard.

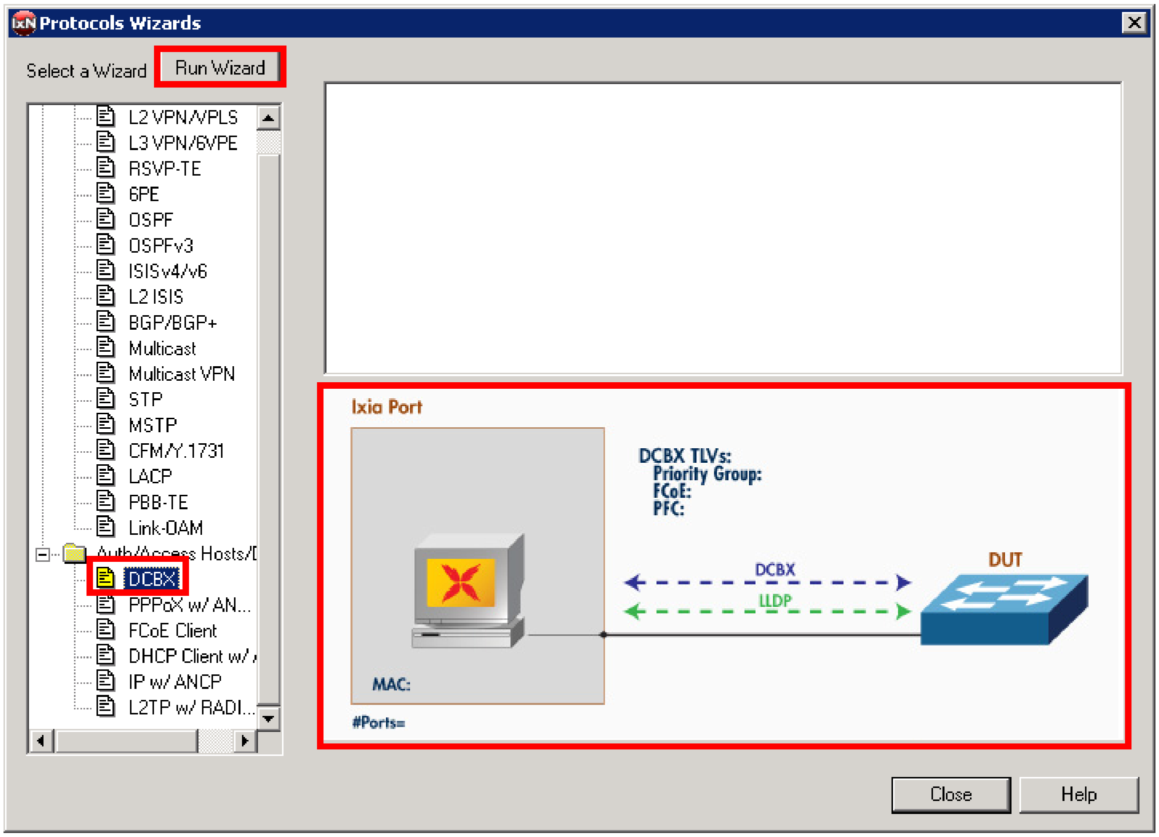

Protocols Wizards

Figure 12. DCBX functionality – DCBX protocol wizard

3. **DCBX Wizard – Port Selection**: Select Ixia port 1.

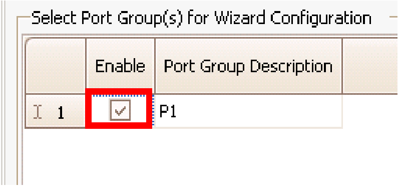

Select Port Group(s) for Wizard Configuration

	Enable	Port Group Description
1	☑	P1

Figure 13. DCBX functionality - DCBX wizard - port selection

4. **DCBX Wizard – MAC**: Set the starting MAC Address value and its increment behavior across ports.

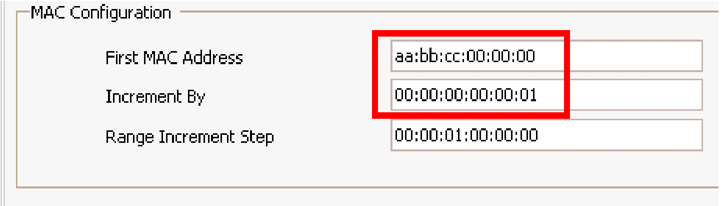

MAC Configuration

First MAC Address: aa:bb:cc:00:00:00

Increment By: 00:00:00:00:00:01

Range Increment Step: 00:00:01:00:00:00

Figure 14. DCBX functionality - DCBX wizard - MAC

5. **DCBX Wizard – LLDP & DCBX Settings**: Select the **Enable DCBX** check box and select **IEEE 1.01** as the DCBX **Subtype**. 1.01 is also known as version CEE. If the DUT only supports version pre-CEE, then set the **Subtype** to **IEEE 1.00**. IEEE 1.00 supports additional TLVs, such as the FCoE and LAN Logical Link Status TLVs. Optionally, change the default values for the LLDP parameters.

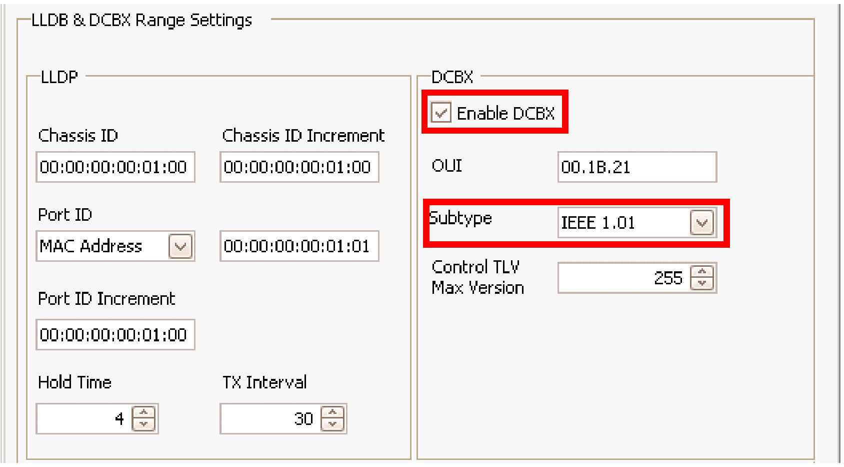

LLDB & DCBX Range Settings

LLDP

Chassis ID: 00:00:00:00:01:00

Chassis ID Increment: 00:00:00:00:01:00

Port ID: MAC Address

00:00:00:00:01:01

Port ID Increment: 00:00:00:00:01:00

Hold Time: 4

TX Interval: 30

DCBX

☑ Enable DCBX

OUI: 00.1B.21

Subtype: IEEE 1.01

Control TLV Max Version: 255

Figure 15. DCBX functionality - DCBX wizard - LLDP & DCBX settings

6. **DCBX Wizard – DCBX TLV Options**: Click **Append** to add DCBX feature TLVs and set the TLV parameters such as User Priority map for the TLVs. By default, the Willing bit is set for each DCBX feature TLV, so IxNetwork will negotiate to the DUT's DCBX settings if the Willing bit on the DUT is off, which is often the case.

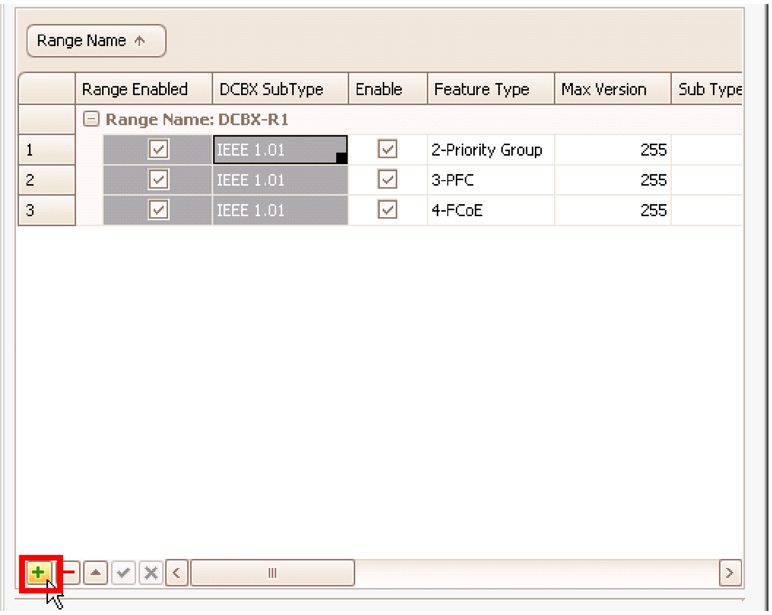

Range Name ↑

	Range Enabled	DCBX SubType	Enable	Feature Type	Max Version	Sub Type
	⊟ Range Name: DCBX-R1					
1	☑	IEEE 1.01	☑	2-Priority Group	255	
2	☑	IEEE 1.01	☑	3-PFC	255	
3	☑	IEEE 1.01	☑	4-FCoE	255	

Figure 16. DCBX functionality - DCBX wizard - DCBX TLVs

a. (Optional) **DCBX TLV Options – Priority Group TLV:** Click the **Priority Group Configuration** cell to open the **Priority Group Configuration** dialog.

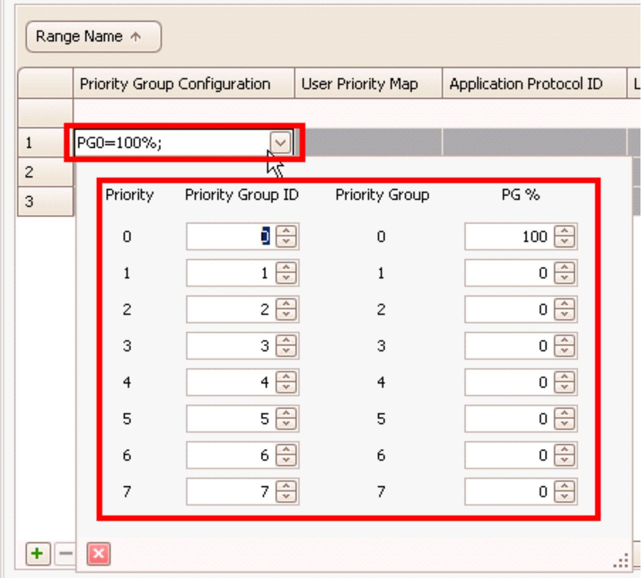

Range Name ↑

	Priority Group Configuration	User Priority Map	Application Protocol ID	L
1	PG0=100%;			
2				
3				

Priority	Priority Group ID	Priority Group	PG %
0	0	0	100
1	1	1	0
2	2	2	0
3	3	3	0
4	4	4	0
5	5	5	0
6	6	6	0
7	7	7	0

Figure 17. DCBX functionality – DCBX wizard – Priority Group TLV

b. (Optional) **DCBX TLV Options – PFC TLV:** Click the **User Priority Map** cell to open the **User Priority Map** dialog.

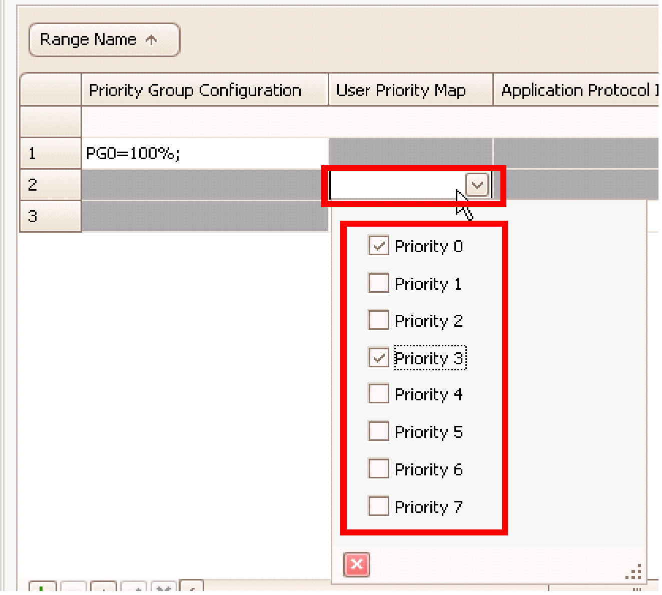

	Priority Group Configuration	User Priority Map	Application Protocol I
1	PG0=100%;		
2			
3			

☑ Priority 0
☐ Priority 1
☐ Priority 2
☑ Priority 3
☐ Priority 4
☐ Priority 5
☐ Priority 6
☐ Priority 7

Figure 18. DCBX functionality - DCBX wizard - PFC TLV

c. (Optional) **DCBX TLV Options – FCoE TLV:** Click the **Application Protocol ID** cell to open the **Application Protocol ID** dialog.

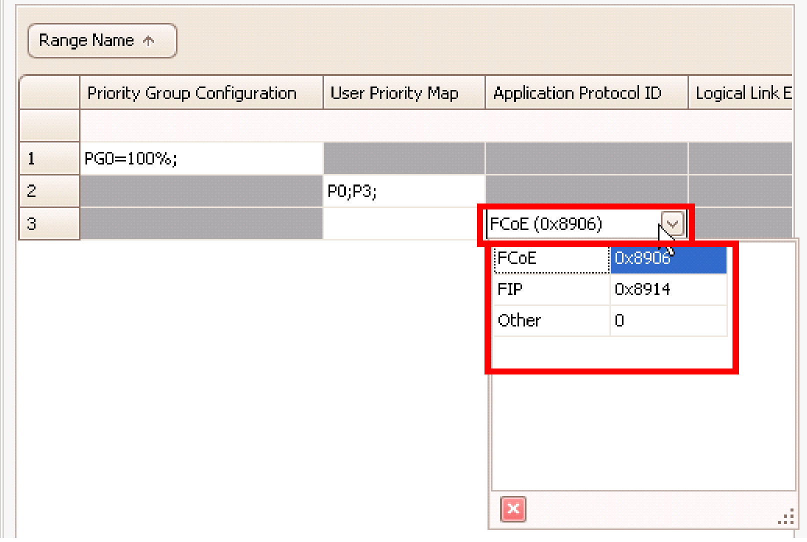

	Priority Group Configuration	User Priority Map	Application Protocol ID	Logical Link E
1	PG0=100%;			
2		P0;P3;		
3			FCoE (0x8906)	

FCoE	0x8906
FIP	0x8914
Other	0

Figure 19. DCBX functionality - DCBX wizard - Application Protocol ID TLV

7. **DCBX Wizard – Name**: Assign a name to this test configuration, and click **Generate and Overwrite Existing Configuration** to apply the configurations defined in this test.

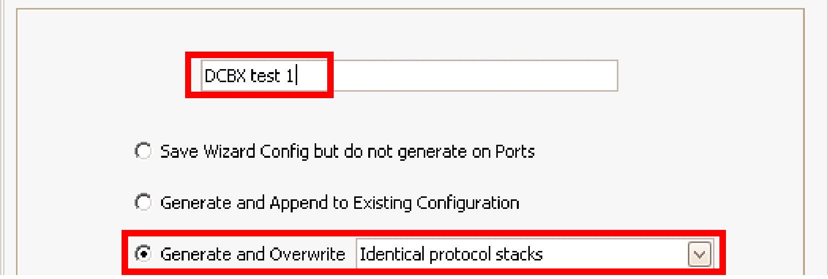

DCBX test 1

○ Save Wizard Config but do not generate on Ports

○ Generate and Append to Existing Configuration

● Generate and Overwrite Identical protocol stacks

Figure 20. DCBX functionality - DCBX wizard - name

8. Review the configuration by validating the emulated LLDP/DCBX agents created, as well as the TLVs. IxNetwork provides tabs for several categories to organize the configuration: **All**, **LLDP**, and **DCBX TLVs**.

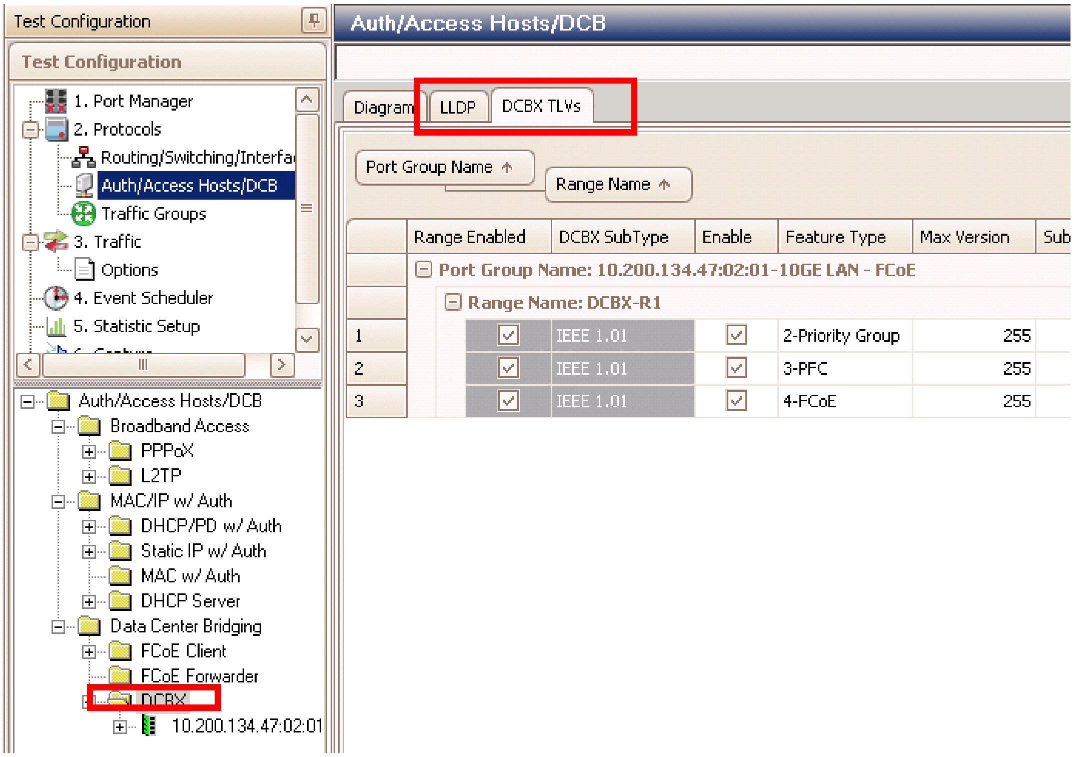

Port Group Name: 10.200.134.47:02:01-10GE LAN - FCoE

Range Name: DCBX-R1

	Range Enabled	DCBX SubType	Enable	Feature Type	Max Version	Sub
1	☑	IEEE 1.01	☑	2-Priority Group	255	
2	☑	IEEE 1.01	☑	3-PFC	255	
3	☑	IEEE 1.01	☑	4-FCoE	255	

Figure 21. DCBX functionality - DCBX configuration GUI view

9. Start the emulation and verify test results.

Test Variables

Functional Variable	Description
Number of Ports	LLDP is a single-session per-port protocol. To stress the DUT's LLDP/DCBX process, increase the **Number of Ports**. The configuration process is unchanged, because the DCBX wizard configures multiple ports simultaneously.
Fast Initialization	In the Fast Initialization phase, DCBX allows a device to transmit five LLDPDUs per second when the device is first brought up (for example, link up). Select the **Fast Init Enable** check box to enable the fast initialization function.
Enable flag	An **Enable** flag is available for each DCBX feature, indicating whether to signal an advertised DCBX feature TLV. Enable or disable the **Feature Enable** flag to verify the DUT's handling of the Enable flag. IxNetwork allows this flag to be enabled or disabled on the fly.
Willing flag	The **Willing** flag is available for each DCBX feature, enabling a device to use its peer's configuration if the peer is not willing to change its feature parameters. Select or clear the **Willing** check box for each feature TLV to verify the DUT's handling of the Willing flag. IxNetwork allows this flag to be enabled or disabled on the fly.
Error flag	The **Error** flag is available for each DCBX feature, and allows IxNetwork to forcefully signal an error for the selected DCBX feature to verify the DUT's response. Select or clear the **Error Override** check box for each TLV to verify the DUT's handling of the Error flag. IxNetwork allows this flag to be enabled or disabled on the fly.
LLDP Transmit Interval	The default LLDPDU transmit interval is 30 seconds. Change the **Tx Interval** value to use other LLDPDU transmit intervals.
Enable FCoE	DCBX advertises an Application Protocol ID TLV that is used to exchange the priority value used for FCoE operations. Set **Enable FCoE** to validate the interworking between DCBX and FCoE. (See the FCoE test cases for configuration instructions.)

Result Analysis

Use the IxNetwork **Statistics** view to verify the following LLDP and DCBX packets statistics:

- **LLDP Neighbor Count**: 1

  Only one LLDP neighbor allowed on each port.

- **LLDP Rx**: 1 every 30 seconds

  By default, LLDPDU is transmitted once every 30 seconds. If **Fast Initialization** is enabled on the DUT, 5 LLDPDUs will be transmitted per second.

- **LLDP Age Out Rx**: 0

  LLDPDUs should be transmitted periodically during the test; no age out condition should occur.

- **DCBX Rx**: 1 every 30 seconds

  By default, LLDPDU is transmitted once every 30 seconds. If **Fast Initialization** is enabled on the DUT, there will be 5 LLDPDUs transmitted per second. Each LLDPDU should include the DCBX OUI if the DUT is completely configured before starting the test.

- **DCBX Control TLV Rx**: 1 every 30 seconds

  By default, LLDPDU is transmitted once every 30 seconds. If **Fast Initialization** is enabled on the DUT, there will be 5 LLDPDUs transmitted per second. Each LLDPDU should include the DCBX Control TLV if the DUT is completely configured before starting the test.

- **DCBX Priority Groups TLV Rx**: 1 every 30 seconds

  By default, LLDPDU is transmitted once every 30 seconds. If **Fast Initialization** is enabled on the DUT, there will be 5 LLDPDUs transmitted per second. Each LLDPDU should include the DCBX Priority Groups TLV if the DUT is completely configured before starting the test.

- **DCBX PFC TLV Rx:** 1 every 30 seconds

  By default, LLDPDU is transmitted once every 30 seconds. If **Fast Initialization** is enabled on the DUT, then there will be 5 LLDPDUs transmitted per second. Each LLDPDU should include the DCBX Priority Groups TLV if the DUT is completely configured before starting the test.

- **DCBX FCoE TLV Rx**: 1 every 30 seconds

  By default, LLDPDU is transmitted once every 30 seconds. If **Fast Initialization** is enabled on the DUT, then there will be 5 LLDPDUs transmitted per second. Each

LLDPDU should include the DCBX Application Protocol ID TLV if the DUT is completely configured before starting the test.

- **DCBX Mismatches Detected**: 0

  There should be no mismatch set for any feature. If this value is non-zero, it indicates the number of configuration mismatches detected between the DUT and IxNetwork during feature negotiation.

- **DCBX Errors Detected**: 0

  There should be no error set for any feature. If this value is non-zero, it indicates the number of errors detected between the DUT and IxNetwork during feature negotiation.

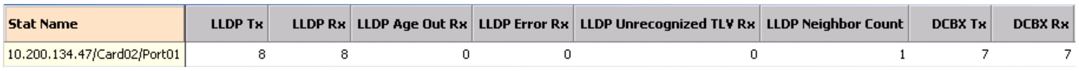

Stat Name	LLDP Tx	LLDP Rx	LLDP Age Out Rx	LLDP Error Rx	LLDP Unrecognized TLV Rx	LLDP Neighbor Count	DCBX Tx	DCBX Rx
10.200.134.47/Card02/Port01	8	8	0	0	0	1	7	7

Figure 22. DCBX functionality - DCBX aggregate statistics view

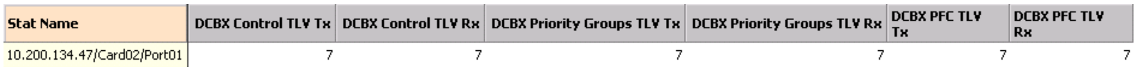

Stat Name	DCBX Control TLV Tx	DCBX Control TLV Rx	DCBX Priority Groups TLV Tx	DCBX Priority Groups TLV Rx	DCBX PFC TLV Tx	DCBX PFC TLV Rx
10.200.134.47/Card02/Port01	7	7	7	7	7	7

Figure 23. DCBX functionality - DCBX aggregate statistics view (cont'd)

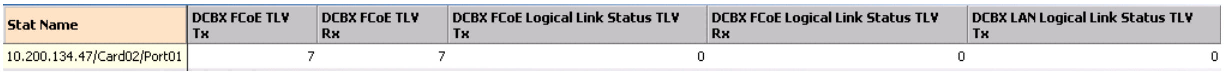

Stat Name	DCBX FCoE TLV Tx	DCBX FCoE TLV Rx	DCBX FCoE Logical Link Status TLV Tx	DCBX FCoE Logical Link Status TLV Rx	DCBX LAN Logical Link Status TLV Tx
10.200.134.47/Card02/Port01	7	7	0	0	0

Figure 24. DCBX functionality - DCBX aggregate statistics view (cont'd)

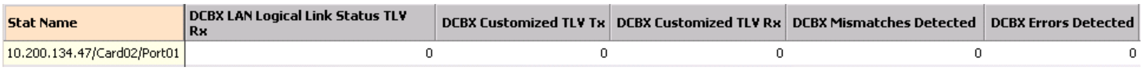

Stat Name	DCBX LAN Logical Link Status TLV Rx	DCBX Customized TLV Tx	DCBX Customized TLV Rx	DCBX Mismatches Detected	DCBX Errors Detected
10.200.134.47/Card02/Port01	0	0	0	0	0

Figure 25. DCBX functionality - DCBX aggregate statistics view (cont'd)

Use the **IxNetwork per TLV** statistics view to verify DCBX exchanges for each DCBX TLV. Right click from the aggregate statistics above and select **DrillDown per TLV**.

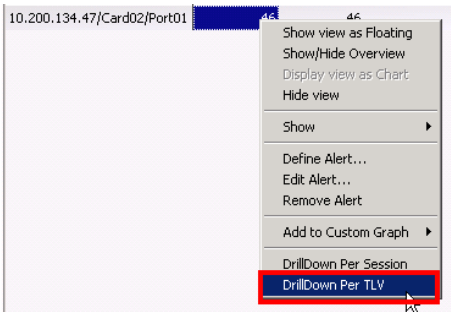

Figure 26. DCBX functionality - DrillDown per TLV

- **Remote Operating Version**: see description

  The displayed value should be equal to the value configured on the DUT for each DCBX feature TLV.

- **Remote Max Version**: see description

  The displayed value should be equal to the value configured on the DUT for each DCBX feature TLV.

- **Remote Enable**: yes

  The enable flag should be set for each DCBX feature TLV configured on the DUT.

- **Remote Willing**: see description

  The displayed value should be equal to the value configured on the DUT for each DCBX feature TLV.

- **Remote Error**: no

  There should be no errors detected for each DCBX feature TLV in this test.

- **Local State**: FEATURE_USE_PEER_CONFIG

  IxNetwork should negotiate to FEATURE_USE_PEER_CONFIG if Willing is disabled on the DUT, which is often the case.

- **Remote TLV Data**: see description

  The displayed value should be equal to the value configured on the DUT for each DCBX feature TLV.

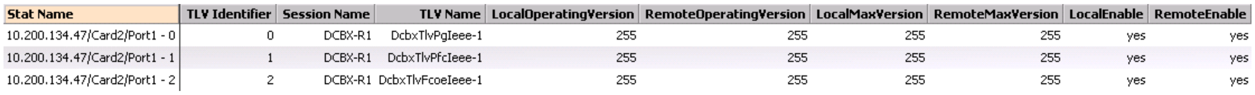

Stat Name	TLV Identifier	Session Name	TLV Name	LocalOperatingVersion	RemoteOperatingVersion	LocalMaxVersion	RemoteMaxVersion	LocalEnable	RemoteEnable
10.200.134.47/Card2/Port1 - 0	0	DCBX-R1	DcbxTlvPgIeee-1	255	255	255	255	yes	yes
10.200.134.47/Card2/Port1 - 1	1	DCBX-R1	DcbxTlvPfcIeee-1	255	255	255	255	yes	yes
10.200.134.47/Card2/Port1 - 2	2	DCBX-R1	DcbxTlvFcoeIeee-1	255	255	255	255	yes	yes

Figure 27. DCBX functionality - per-TLV statistics

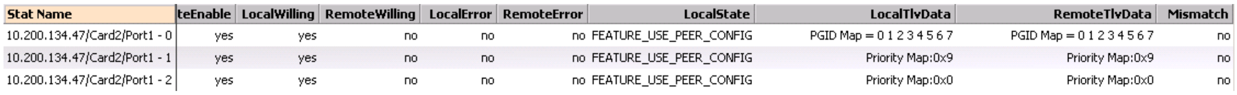

Stat Name	teEnable	LocalWilling	RemoteWilling	LocalError	RemoteError	LocalState	LocalTlvData	RemoteTlvData	Mismatch
10.200.134.47/Card2/Port1 - 0	yes	yes	no	no	no	FEATURE_USE_PEER_CONFIG	PGID Map = 0 1 2 3 4 5 6 7	PGID Map = 0 1 2 3 4 5 6 7	no
10.200.134.47/Card2/Port1 - 1	yes	yes	no	no	no	FEATURE_USE_PEER_CONFIG	Priority Map:0x9	Priority Map:0x9	no
10.200.134.47/Card2/Port1 - 2	yes	yes	no	no	no	FEATURE_USE_PEER_CONFIG	Priority Map:0x0	Priority Map:0x0	no

Figure 28. DCBX functionality - per-TLV statistics (cont'd)

Conclusions

By validating the statistics and TLV exchanges using the features above, the DUT has been proven capable of properly advertising, discovering and negotiating LLDP and DCBX feature TLVs for the specified topology and test configuration.

Test Case: FCoE Fabric Login Stress Test (with NPIV)

Overview

Fabric login is the most common protocol operation in a fibre channel environment. When connecting an FCoE ENode to an FCF, the first action performed at the FC level is a fabric login (FLOGI). Without a successful login, no real fibre channel communication can be performed. Therefore, FLOGI is the most critical functionality to verify because this is the foundation that all further FC data exchange is built on. In a large data center, the number of FLOGIs processed by the FCF can reach hundreds and even thousands per port.

For each VN_Port to instantiate, the ENode's FCoE controller on an Ethernet port initiates one FLOGI to an FCF capable of instantiating VF_Ports. Each VN_Port can be further virtualized by acquiring additional port IDs for each VN_Port. This uses a process called N_Port virtualization using the F_Port Discovery (FDISC) request/reply sequence.

Objective

The objective of this test is to verify the ability of an FCF (that is, the device under test or DUT) to accept a high number of VN_Port instantiation requests. Both FLOGI instantiated and FDISC instantiated VN_Ports will be enabled and configured to request logins. FLOGIs will first be transmitted from each emulated FLOGI VN_Port. As the FCF accepts each FLOGI VN_Port, FDISCs will subsequently be transmitted from each emulated FDISC VN_Port.

The FCF should accept all login requests using an LS_ACC, and properly assign non-overlapping FC IDs, used as S_ID (source ID) and/or D_ID (destination ID) to each VN_Port.

Setup

Two Ixia ports are used in this test, one on the initiator side, and one on the target side.

The objective of this test is to instantiate 10,000 VN_Ports.

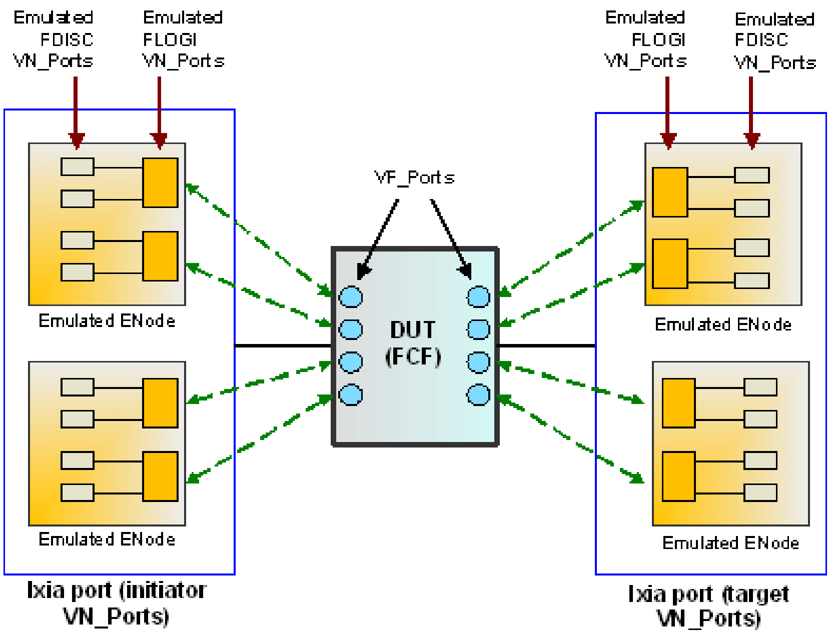

Figure 29. FCoE fabric login stress test topology

Step-by-Step Instructions

1. Reserve two ports in IxNetwork.

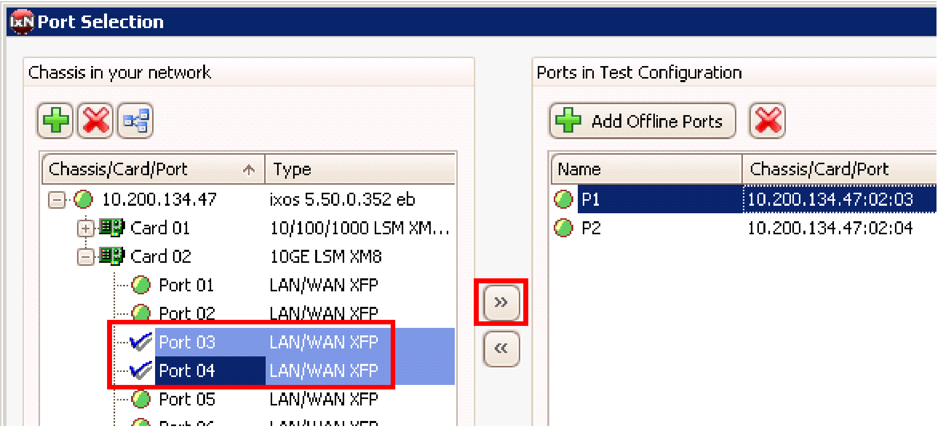

Chassis/Card/Port	Type
10.200.134.47	ixos 5.50.0.352 eb
Card 01	10/100/1000 LSM XM...
Card 02	10GE LSM XM8
Port 01	LAN/WAN XFP
Port 02	LAN/WAN XFP
✓ Port 03	LAN/WAN XFP
✓ Port 04	LAN/WAN XFP
Port 05	LAN/WAN XFP

Name	Chassis/Card/Port
P1	10.200.134.47:02:03
P2	10.200.134.47:02:04

Figure 30. FCoE fabric login stress test – reserving two ports

2. Set the port **Type** to either **10GE LAN – FCoE** or **10GE WAN – FCoE**.

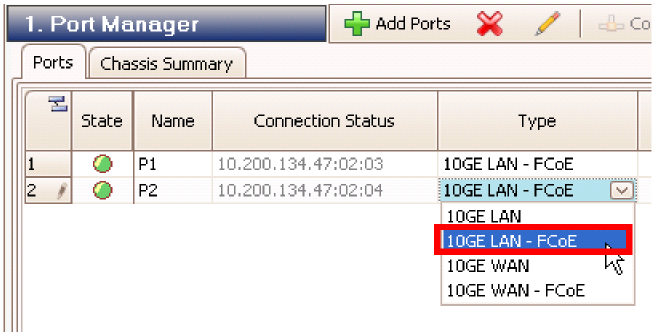

	State	Name	Connection Status	Type
1	●	P1	10.200.134.47:02:03	10GE LAN - FCoE
2	●	P2	10.200.134.47:02:04	10GE LAN - FCoE

Figure 31. FCoE fabric login stress test – setting port type to FCoE

3. (Optional) Enable DCBX. See the test case ‘DCBX Functionality’ for DCBX configuration procedures.

4. Open the IxNetwork **Protocol Wizards** window, and run the **FCoE/FC Node Wizard**.

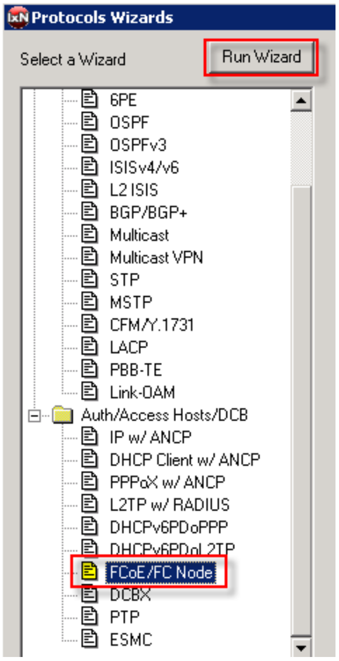

Figure 32. FCoE fabric login stress test - FCoE/FC Node protocol wizard

5. **FCoE/FC Node Wizard – Port Selection**: Configure Ixia port 1 to be the **Initiator** side port, and Ixia port 2 to be the **Target** side port.

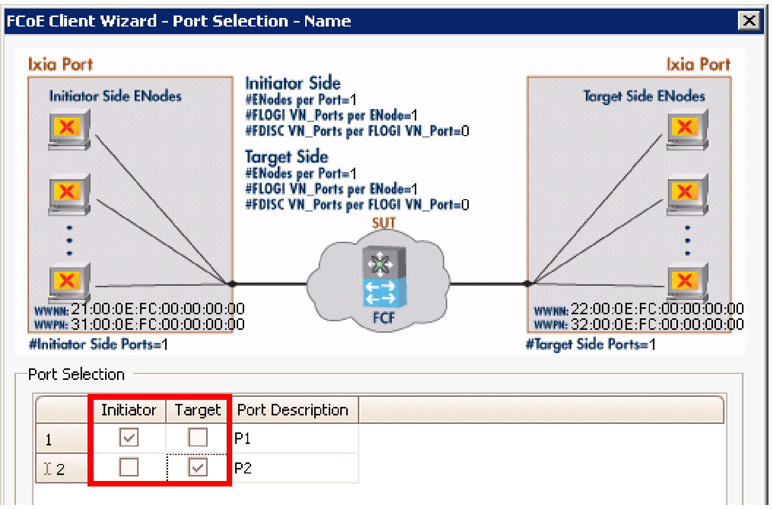

FCoE Client Wizard - Port Selection - Name

Initiator Side
#ENodes per Port=1
#FLOGI VN_Ports per ENode=1
#FDISC VN_Ports per FLOGI VN_Port=0

Target Side
#ENodes per Port=1
#FLOGI VN_Ports per ENode=1
#FDISC VN_Ports per FLOGI VN_Port=0

WWNN: 21:00:0E:FC:00:00:00:00
WWPN: 31:00:0E:FC:00:00:00:00
#Initiator Side Ports=1

WWNN: 22:00:0E:FC:00:00:00:00
WWPN: 32:00:0E:FC:00:00:00:00
#Target Side Ports=1

Port Selection

	Initiator	Target	Port Description
1	☑	□	P1
2	□	☑	P2

Figure 33. FCoE Fabric Login stress test – FCoE/FC Node Wizard – port selection

6. **FCoE/FC Node Wizard – Flow Control**: Leave the default values as this is beyond the scope of the test.

7. **FCoE/FC Node Wizard – MAC**: Set the starting MAC Address value and its increment behavior across ENodes. The **First MAC Address** value is the MAC address assigned to the first ENode. The **Increment By** value is the MAC address increment step for all subsequent ENodes on the same port. The **Range Increment Step** value is the MAC address increment step for ENodes across ports.

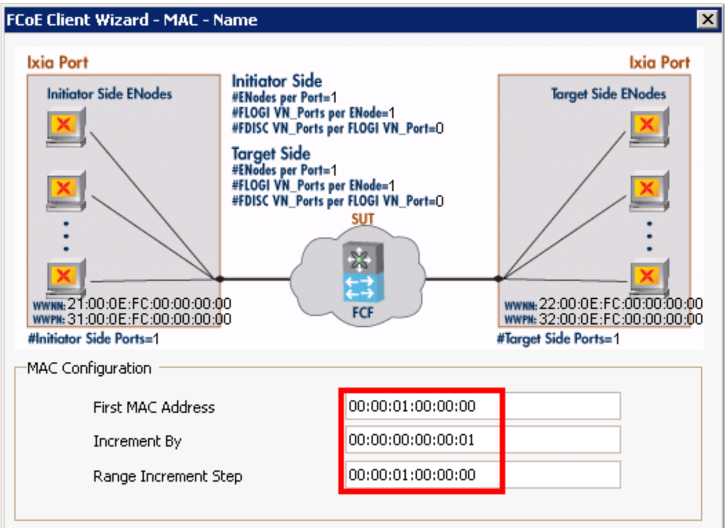

FCoE Client Wizard - MAC - Name

Initiator Side
#ENodes per Port=1
#FLOGI VN_Ports per ENode=1
#FDISC VN_Ports per FLOGI VN_Port=0

Target Side
#ENodes per Port=1
#FLOGI VN_Ports per ENode=1
#FDISC VN_Ports per FLOGI VN_Port=0

WWNN: 21:00:0E:FC:00:00:00:00
WWPN: 31:00:0E:FC:00:00:00:00
#Initiator Side Ports=1

WWNN: 22:00:0E:FC:00:00:00:00
WWPN: 32:00:0E:FC:00:00:00:00
#Target Side Ports=1

MAC Configuration

First MAC Address: 00:00:01:00:00:00
Increment By: 00:00:00:00:00:01
Range Increment Step: 00:00:01:00:00:00

Figure 34. FCoE Fabric Login stress test – FCoE/FC Node Wizard – MAC

8. **FCoE/FC Node Wizard – VLAN**: Leave the default values as this is beyond the scope of the test.

9. **FCoE/FC Node Wizard – FIP**: Configure the FCoE/FIP global protocol features. Select the **Enable Name Server Registration** and **Perform PLOGI** check boxes. Clear the **Enable FCoE Initialization Protocol (FIP)** check box.

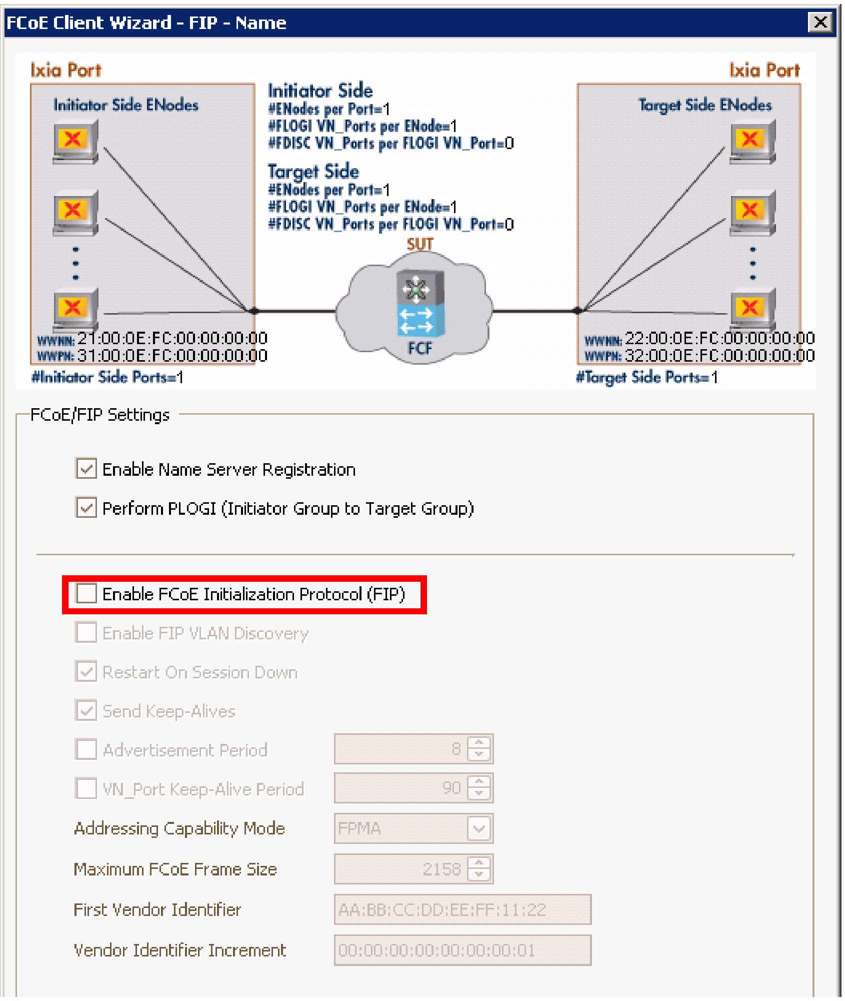
FCoE Client Wizard - FIP - Name
Ixia Port
Initiator Side ENodes
Initiator Side
#ENodes per Port=1
#FLOGI VN_Ports per ENode=1
#FDISC VN_Ports per FLOGI VN_Port=0
Target Side
#ENodes per Port=1
#FLOGI VN_Ports per ENode=1
#FDISC VN_Ports per FLOGI VN_Port=0
SUT
FCF
Ixia Port
Target Side ENodes
WWNN: 21:00:0E:FC:00:00:00:00
WWPN: 31:00:0E:FC:00:00:00:00
#Initiator Side Ports=1
WWNN: 22:00:0E:FC:00:00:00:00
WWPN: 32:00:0E:FC:00:00:00:00
#Target Side Ports=1

FCoE/FIP Settings

☑ Enable Name Server Registration

☑ Perform PLOGI (Initiator Group to Target Group)

☐ Enable FCoE Initialization Protocol (FIP)

☐ Enable FIP VLAN Discovery

☑ Restart On Session Down

☑ Send Keep-Alives

☐ Advertisement Period: 8

☐ VN_Port Keep-Alive Period: 90

Addressing Capability Mode: FPMA

Maximum FCoE Frame Size: 2158

First Vendor Identifier: AA:BB:CC:DD:EE:FF:11:22

Vendor Identifier Increment: 00:00:00:00:00:00:00:01

Figure 35. FCoE Fabric Login stress test – FCoE/FC Node Wizard – FIP

10. **FCoE/FC Node Wizard – Left** (Initiator) **Side VN_Ports**: Configure the number of ENodes, FLOGI VN_Ports and FDISC VN_Ports, as well as the OUI, Node and Port WWNs associated with each VN_Port.

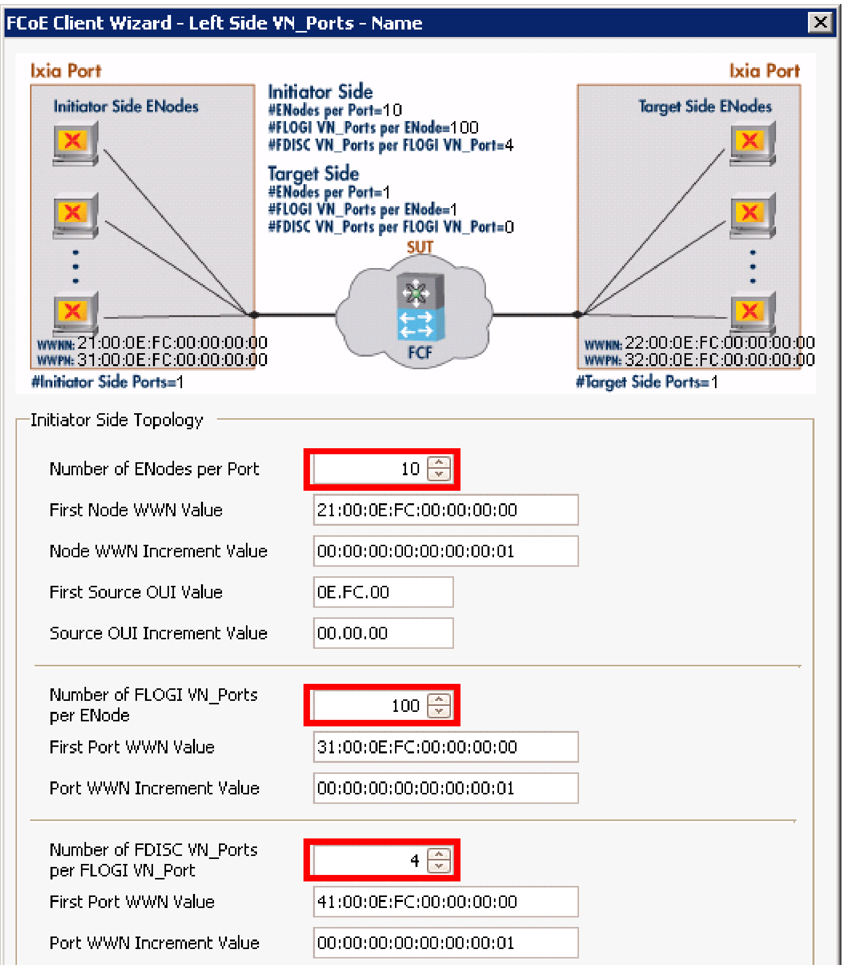
FCoE Client Wizard - Left Side VN_Ports - Name
Ixia Port
Initiator Side ENodes
Initiator Side
#ENodes per Port=10
#FLOGI VN_Ports per ENode=100
#FDISC VN_Ports per FLOGI VN_Port=4
Target Side
#ENodes per Port=1
#FLOGI VN_Ports per ENode=1
#FDISC VN_Ports per FLOGI VN_Port=0
SUT
FCF
Ixia Port
Target Side ENodes
WWNN: 21:00:0E:FC:00:00:00:00
WWPN: 31:00:0E:FC:00:00:00:00
#Initiator Side Ports=1
WWNN: 22:00:0E:FC:00:00:00:00
WWPN: 32:00:0E:FC:00:00:00:00
#Target Side Ports=1

Initiator Side Topology

Number of ENodes per Port	10
First Node WWN Value	21:00:0E:FC:00:00:00:00
Node WWN Increment Value	00:00:00:00:00:00:00:01
First Source OUI Value	0E.FC.00
Source OUI Increment Value	00.00.00
Number of FLOGI VN_Ports per ENode	100
First Port WWN Value	31:00:0E:FC:00:00:00:00
Port WWN Increment Value	00:00:00:00:00:00:00:01
Number of FDISC VN_Ports per FLOGI VN_Port	4
First Port WWN Value	41:00:0E:FC:00:00:00:00
Port WWN Increment Value	00:00:00:00:00:00:00:01

Figure 36. FCoE fabric login stress test – FCoE/FC Node Wizard – initiator side VN_Ports

To configure the target for this test, set the following values:

a. **Number of ENodes per Port**: *10*

b. **Number of FLOGI VN_Ports per ENode**: *100*

c. **Number of FDISC VN_Ports per FLOGI VN_Port**: *4*

11. **FCoE/FC Node Wizard – Right** (Target) **Side VN_Ports**: Configure a topology on the target side that is symmetric to the initiator side by simply selecting the **Same as Initiator Side Topology** check box. IxNetwork will automatically assign OUI, Node and Port WWNs that are globally unique.

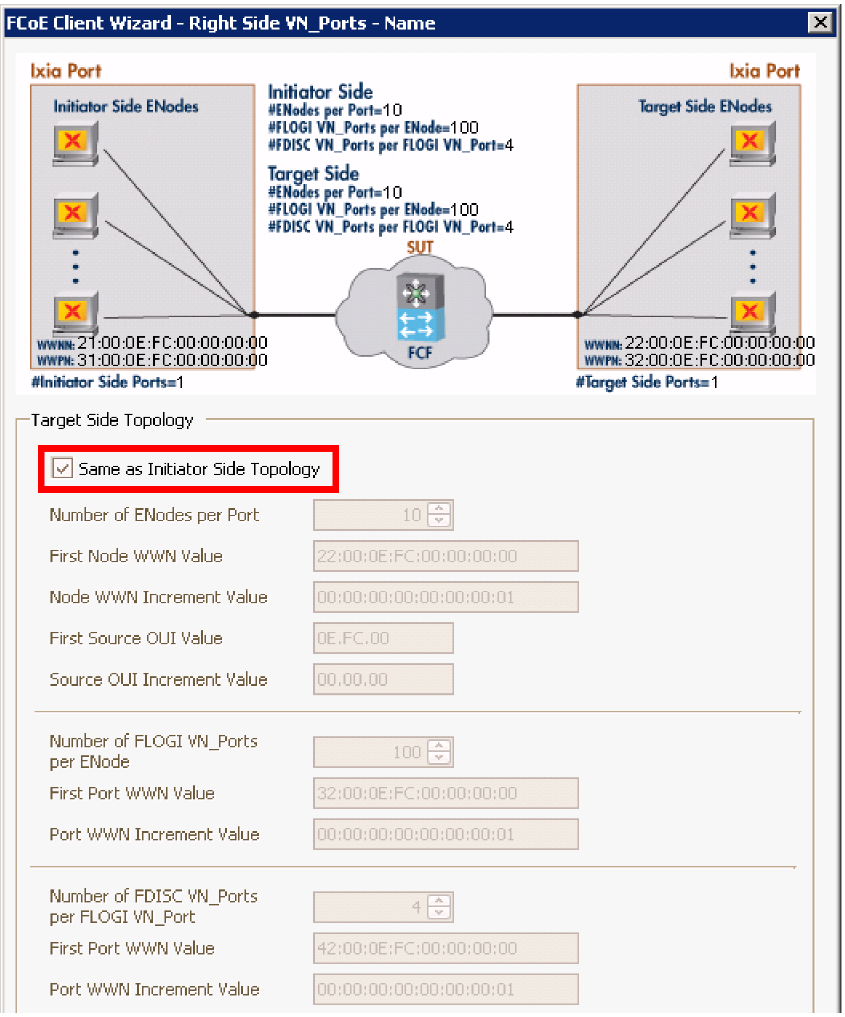
FCoE Client Wizard - Right Side VN_Ports - Name
Ixia Port
Initiator Side ENodes
Initiator Side
#ENodes per Port=10
#FLOGI VN_Ports per ENode=100
#FDISC VN_Ports per FLOGI VN_Port=4
Target Side
#ENodes per Port=10
#FLOGI VN_Ports per ENode=100
#FDISC VN_Ports per FLOGI VN_Port=4
SUT
FCF
Ixia Port
Target Side ENodes
WWNN: 21:00:0E:FC:00:00:00:00
WWPN: 31:00:0E:FC:00:00:00:00
#Initiator Side Ports=1
WWNN: 22:00:0E:FC:00:00:00:00
WWPN: 32:00:0E:FC:00:00:00:00
#Target Side Ports=1

Target Side Topology

☑ Same as Initiator Side Topology

Number of ENodes per Port	10
First Node WWN Value	22:00:0E:FC:00:00:00:00
Node WWN Increment Value	00:00:00:00:00:00:00:01
First Source OUI Value	0E.FC.00
Source OUI Increment Value	00.00.00
Number of FLOGI VN_Ports per ENode	100
First Port WWN Value	32:00:0E:FC:00:00:00:00
Port WWN Increment Value	00:00:00:00:00:00:00:01
Number of FDISC VN_Ports per FLOGI VN_Port	4
First Port WWN Value	42:00:0E:FC:00:00:00:00
Port WWN Increment Value	00:00:00:00:00:00:00:01

Figure 37. FCoE fabric login stress test – FCoE/FC Node Wizard – target side VN_Ports

12. FCoE/FC Node Wizard – Name: Assign a name to this test configuration, and click **Generate and Overwrite Existing Configuration** to apply the configurations defined in this test.

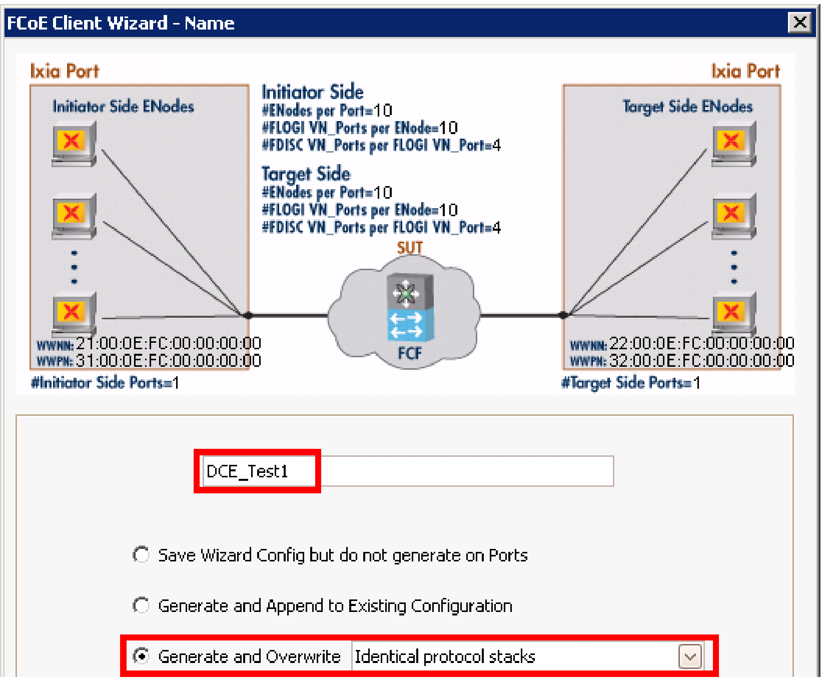

Initiator Side
#ENodes per Port=10
#FLOGI VN_Ports per ENode=10
#FDISC VN_Ports per FLOGI VN_Port=4

Target Side
#ENodes per Port=10
#FLOGI VN_Ports per ENode=10
#FDISC VN_Ports per FLOGI VN_Port=4

WWNN: 21:00:0E:FC:00:00:00:00
WWPN: 31:00:0E:FC:00:00:00:00
#Initiator Side Ports=1

WWNN: 22:00:0E:FC:00:00:00:00
WWPN: 32:00:0E:FC:00:00:00:00
#Target Side Ports=1

DCE_Test1

○ Save Wizard Config but do not generate on Ports

○ Generate and Append to Existing Configuration

● Generate and Overwrite Identical protocol stacks

Figure 38. FCoE fabric login stress test – FCoE/FC Node Wizard – name

13. Review the configuration by validating the emulated ENodes, FLOGI VN_Ports and FDISC VN_Ports created.

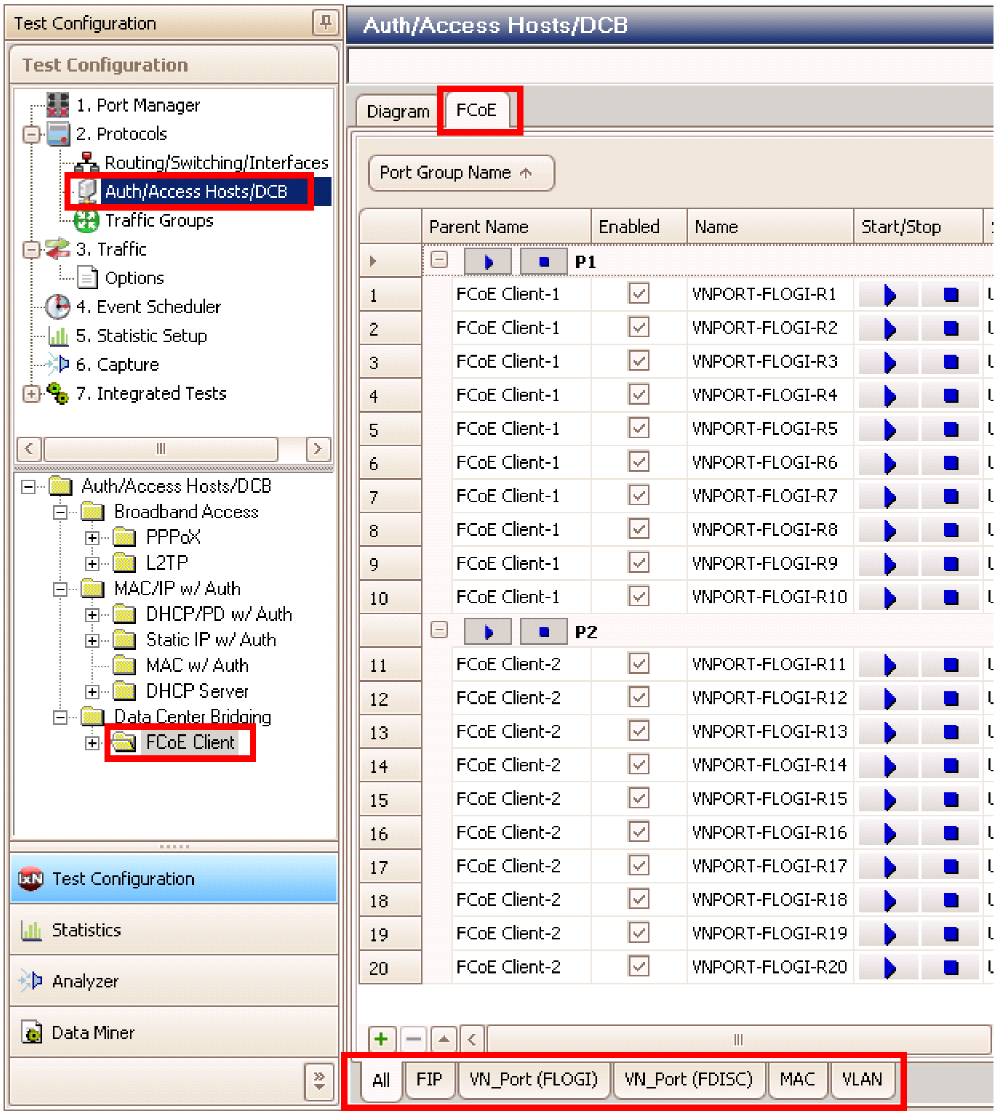

	Parent Name	Enabled	Name	Start/Stop
	⊟ ▶ ■ P1			
1	FCoE Client-1	☑	VNPORT-FLOGI-R1	▶ ■
2	FCoE Client-1	☑	VNPORT-FLOGI-R2	▶ ■
3	FCoE Client-1	☑	VNPORT-FLOGI-R3	▶ ■
4	FCoE Client-1	☑	VNPORT-FLOGI-R4	▶ ■
5	FCoE Client-1	☑	VNPORT-FLOGI-R5	▶ ■
6	FCoE Client-1	☑	VNPORT-FLOGI-R6	▶ ■
7	FCoE Client-1	☑	VNPORT-FLOGI-R7	▶ ■
8	FCoE Client-1	☑	VNPORT-FLOGI-R8	▶ ■
9	FCoE Client-1	☑	VNPORT-FLOGI-R9	▶ ■
10	FCoE Client-1	☑	VNPORT-FLOGI-R10	▶ ■
	⊟ ▶ ■ P2			
11	FCoE Client-2	☑	VNPORT-FLOGI-R11	▶ ■
12	FCoE Client-2	☑	VNPORT-FLOGI-R12	▶ ■
13	FCoE Client-2	☑	VNPORT-FLOGI-R13	▶ ■
14	FCoE Client-2	☑	VNPORT-FLOGI-R14	▶ ■
15	FCoE Client-2	☑	VNPORT-FLOGI-R15	▶ ■
16	FCoE Client-2	☑	VNPORT-FLOGI-R16	▶ ■
17	FCoE Client-2	☑	VNPORT-FLOGI-R17	▶ ■
18	FCoE Client-2	☑	VNPORT-FLOGI-R18	▶ ■
19	FCoE Client-2	☑	VNPORT-FLOGI-R19	▶ ■
20	FCoE Client-2	☑	VNPORT-FLOGI-R20	▶ ■

Figure 39. FCoE fabric login stress test – FCoE client configuration GUI view

IxNetwork provides tabs for several categories to organize the configuration: **All**, **FIP**, **VN_Port (FLOGI)**, **VN_Port (FDISC)**, **MAC**, and **VLAN**. The **All** tab should show that ten ENodes (named **VNPORT-FDISC-Rxx**) have been created on each Ixia port. The **FIP** tab should show that no ENodes have FIP enabled.

The **VN_Port (FLOGI)** tab should show each ENode has 100 VN_Ports configured. In addition, only ENodes on Port 1 have the **PLOGI Target** defined. The **VN_Port (FDISC)** tab should show each FLOGI VN_Port has four FDISC VN_Ports configured.

14. Optionally, configure the login **Setup Rate** and logout **Teardown Rate**. These values determine the rate at which FLOGIs or FDISCs are transmitted to the FCF, and the rate at which LOGOs are transmitted to the FCF.

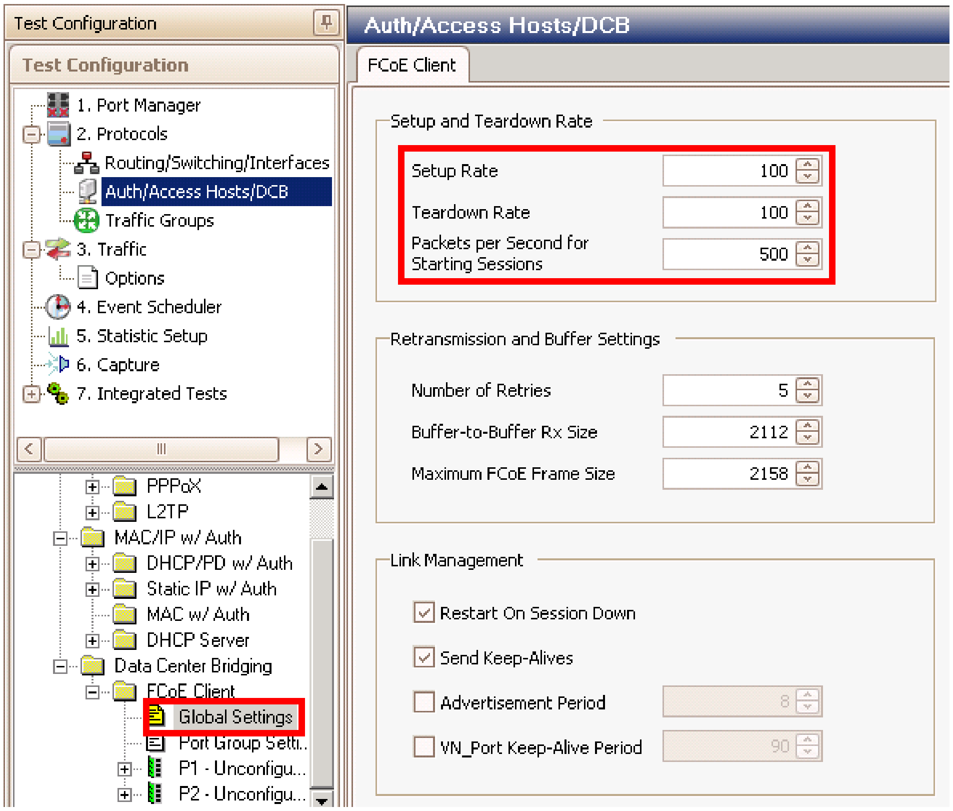

Setup and Teardown Rate

Setup Rate 100

Teardown Rate 100

Packets per Second for Starting Sessions 500

Retransmission and Buffer Settings

Number of Retries 5

Buffer-to-Buffer Rx Size 2112

Maximum FCoE Frame Size 2158

Link Management

☑ Restart On Session Down

☑ Send Keep-Alives

□ Advertisement Period 8

□ VN_Port Keep-Alive Period 90

Figure 40. FCoE fabric login stress test – login setup and logout teardown rates

15. Start the emulation and verify test results.

16. Wait until all VN_Ports are instantiated. This process can take seconds or minutes depending on the configured login setup rate.

Test Variables

Each of the following items may be used in separate test cases to test an FCF. They all use the test case developed thus far as a baseline.

Performance Variable	Description
Number of ENodes per port	Each ENode instantiates a set of FLOGI VN_Ports and FDISC VN_Ports. Increasing the **Number of ENodes per Port** will stress the DUT’s (FCF’s) ability to process FCoE login requests.
Number of FLOGI VN_Ports per ENode	A VN_Port can be instantiated by FLOGI or FDISC. Increasing the **Number of FLOGI VN_Ports per ENode** will stress the DUT’s (FCF) ability to process FLOGI requests.
Number of FDISC VN_Ports per FLOGI VN_Port	A VN_Port can be instantiated by FLOGI or FDISC. Increasing the **Number of FDISC VN_Ports per FLOGI VN_Port** will stress the DUT’s (FCF) ability to process FDISC requests.
Name Server Registration	Each VN_Port registers its attributes with the Name Server (commonly the FCF) after it logs in to the fabric. In addition, each VN_Port queries the name server for each destination VN_Port. Enabling this option puts additional stress on the DUT to process name server registrations and queries.
Perform PLOGI	PLOGI is a connection request made from a source VN_Port to a destination VN_Port. Enabling this option puts additional stress on the DUT to process PLOGI request and replies.
Setup/Teardown Rate	(Optional) The rate at which VN_Ports request to login to a fabric. Increasing this number will stress the DUT’s ability to process a large volume of Fabric Login and logout requests.

Results Analysis

Use the IxNetwork **Statistics** view to verify the following FCoE packets aggregate statistics:

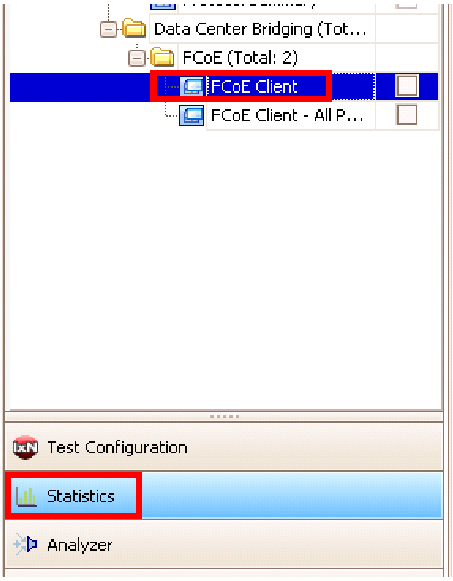

Figure 41. FCoE fabric login stress test - IxNetwork Statistics view

- **FLOGI LS_ACC Rx**: 1,000

  There are 100 FLOGI VN_Ports configured for each of the 10 ENodes. Each VN_Port is instantiated by transmitting one FLOGI. The DUT accepts each FLOGI by responding with an LS_ACC.

- **FDISC LS_ACC Rx**: 4,000

  There are 4 FDISC VN_Ports configured for each of the 1,000 FLOGI VN_Ports. Each FDISC VN_Port is instantiated by transmitting one FDISC. The DUT accepts each FDISC by responding with an LS_ACC.

- **PLOGI LS_ACC Rx**: 5,000

  There are 5,000 VN_Ports in total. Each of the 1,000 FLOGI VN_Ports transmits one PLOGI to its target VN_Port. Each of the 4,000 FDISC VN_Ports transmits one PLOGI to its target VN_Port. The target side Ixia port responds to each PLOGI with an LS_ACC. The DUT forwards each PLOGI LS_ACC to the initiator side Ixia port.

- **NS Registration Successful**: 5,000

  Each of the 5,000 VN_Ports registers with the name server (that is, the DUT in this test). The DUT acknowledges each of the 5,000 NS Registrations.

- **Interfaces Up**: 5,000

  All 5,000 VN_Ports should be instantiated.

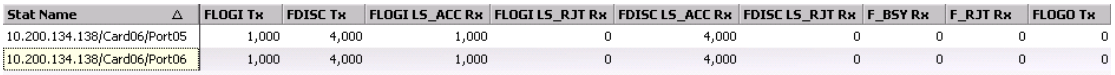

Stat Name △	FLOGI Tx	FDISC Tx	FLOGI LS_ACC Rx	FLOGI LS_RJT Rx	FDISC LS_ACC Rx	FDISC LS_RJT Rx	F_BSY Rx	F_RJT Rx	FLOGO Tx
10.200.134.138/Card06/Port05	1,000	4,000	1,000	0	4,000	0	0	0	0
10.200.134.138/Card06/Port06	1,000	4,000	1,000	0	4,000	0	0	0	0

Figure 42. FCoE fabric login stress test – FCoE client aggregate statistics view

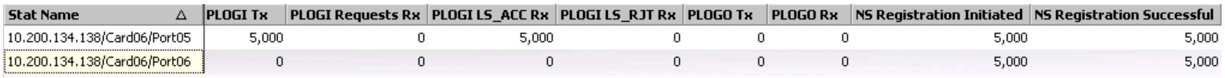

Stat Name △	PLOGI Tx	PLOGI Requests Rx	PLOGI LS_ACC Rx	PLOGI LS_RJT Rx	PLOGO Tx	PLOGO Rx	NS Registration Initiated	NS Registration Successful
10.200.134.138/Card06/Port05	5,000	0	5,000	0	0	0	5,000	5,000
10.200.134.138/Card06/Port06	0	0	0	0	0	0	5,000	5,000

Figure 43. FCoE fabric login stress test – FCoE client aggregate statistics view (cont'd)

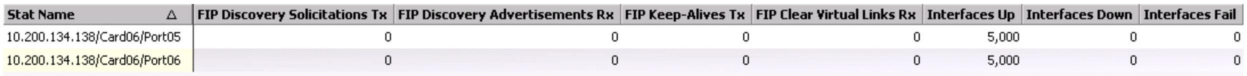

Stat Name △	FIP Discovery Solicitations Tx	FIP Discovery Advertisements Rx	FIP Keep-Alives Tx	FIP Clear Virtual Links Rx	Interfaces Up	Interfaces Down	Interfaces Fail
10.200.134.138/Card06/Port05	0	0	0	0	5,000	0	0
10.200.134.138/Card06/Port06	0	0	0	0	5,000	0	0

Figure 44. FCoE fabric login stress test – FCoE client aggregate statistics view (Cont'd)

Use the IxNetwork per session statistics view to verify the following FCoE control plane exchanges for each VN_Port. Right click from the aggregate statistics above and select **DrillDown per Session**.

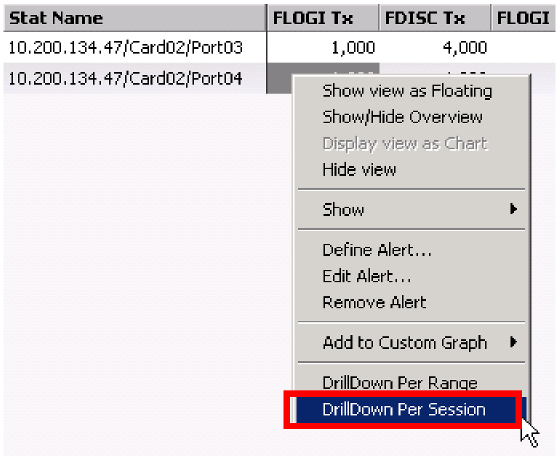

Stat Name	FLOGI Tx	FDISC Tx	FLOGI
10.200.134.47/Card02/Port03	1,000	4,000	
10.200.134.47/Card02/Port04			

Figure 45. FCoE fabric login stress test - DrillDown per Session

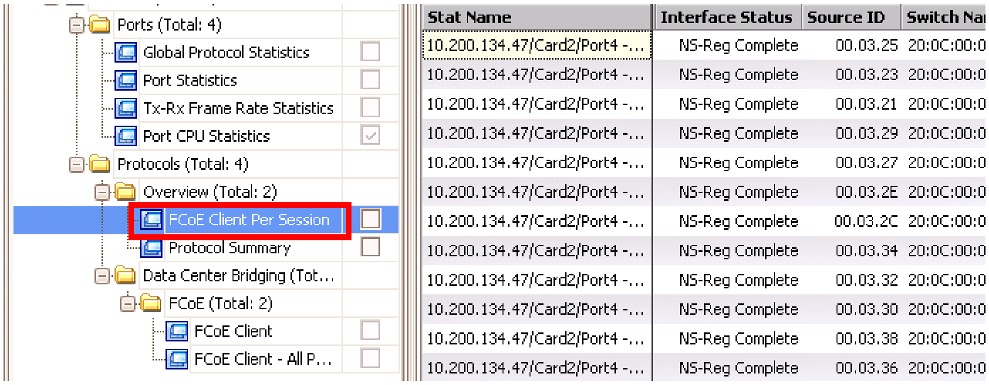

Stat Name	Interface Status	Source ID	Switch Na
10.200.134.47/Card2/Port4 -...	NS-Reg Complete	00.03.25	20:0C:00:0
10.200.134.47/Card2/Port4 -...	NS-Reg Complete	00.03.23	20:0C:00:0
10.200.134.47/Card2/Port4 -...	NS-Reg Complete	00.03.21	20:0C:00:0
10.200.134.47/Card2/Port4 -...	NS-Reg Complete	00.03.29	20:0C:00:0
10.200.134.47/Card2/Port4 -...	NS-Reg Complete	00.03.27	20:0C:00:0
10.200.134.47/Card2/Port4 -...	NS-Reg Complete	00.03.2E	20:0C:00:0
10.200.134.47/Card2/Port4 -...	NS-Reg Complete	00.03.2C	20:0C:00:0
10.200.134.47/Card2/Port4 -...	NS-Reg Complete	00.03.34	20:0C:00:0
10.200.134.47/Card2/Port4 -...	NS-Reg Complete	00.03.32	20:0C:00:0
10.200.134.47/Card2/Port4 -...	NS-Reg Complete	00.03.30	20:0C:00:0
10.200.134.47/Card2/Port4 -...	NS-Reg Complete	00.03.38	20:0C:00:0
10.200.134.47/Card2/Port4 -...	NS-Reg Complete	00.03.36	20:0C:00:0

Figure 46. FCoE fabric login stress test – FCoE client per VN_Port statistics view

- All 5,000 VN_Ports from the initiator side should have **Status** set to **PLOGI OK**. This is the last initialization operation performed by the initiator side VN_Ports.
- All 5,000 VN_Ports from the target side have **Status** set to **NS-Reg OK**. This is the last initialization operation performed by the target side VN_Ports.
- Each of the 10,000 VN_Ports is assigned a unique **Source ID**.
- Each of the 5,000 VN_Ports from the initiator side is assigned a unique **PLOGI Destination ID**. This is the ID that the initiator VN_Port uses to perform PLOGI to its target VN_Port.

*Note: use the **Filter/Sort** function to zoom in on VN_Ports that matches user defined statistics criteria*

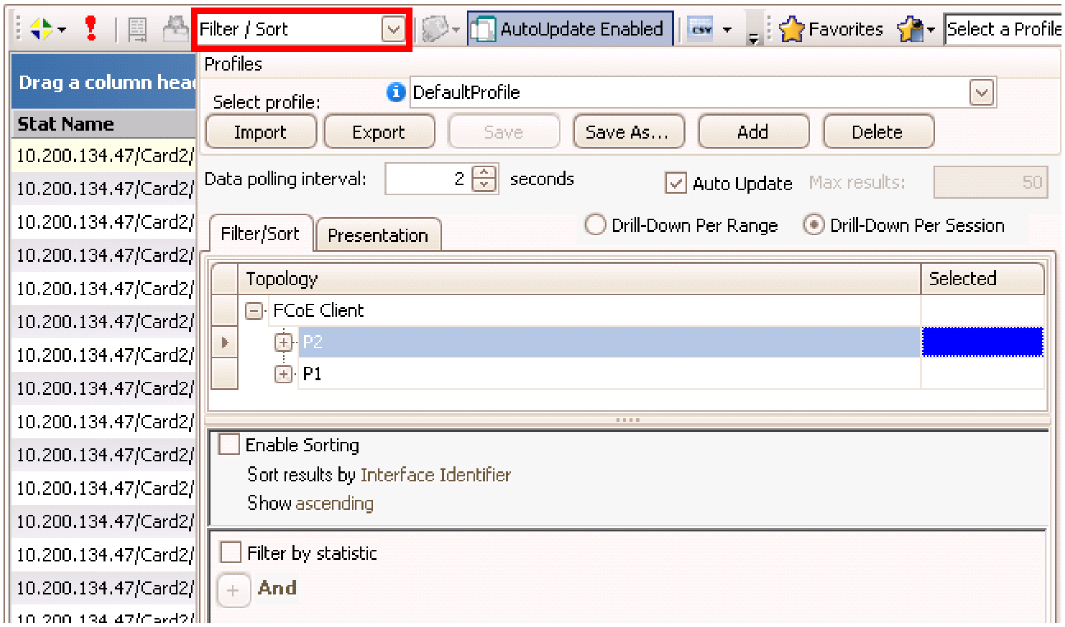

Stat Name
10.200.134.47/Card2/
10.200.134.47/Card2/
10.200.134.47/Card2/
10.200.134.47/Card2/
10.200.134.47/Card2/
10.200.134.47/Card2/
10.200.134.47/Card2/
10.200.134.47/Card2/
10.200.134.47/Card2/
10.200.134.47/Card2/
10.200.134.47/Card2/
10.200.134.47/Card2/
10.200.134.47/Card2/
10.200.134.47/Card2/
10.200.134.47/Card2/

Select profile: DefaultProfile

Data polling interval: 2 seconds ☑ Auto Update Max results: 50

○ Drill-Down Per Range ● Drill-Down Per Session

Topology	Selected
⊟ FCoE Client	
⊞ P2	■
⊞ P1	

☐ Enable Sorting

Sort results by Interface Identifier

Show ascending

☐ Filter by statistic

And

Figure 47. FCoE fabric login stress test - statistics Filter/Sort function

Conclusions

By validating the statistics and control plane exchanges using the features above, the DUT has been demonstrated to be capable of instantiating at least 5,000 VN_Ports per port, instantiated by both FLOGI and FDISC, for the specified topology.

Test Case: FIP Fabric Login Stress Test

Overview

Fabric login is the most common protocol operation in a fibre channel environment. When connecting an FCoE ENode to an FCF, the first action performed at the FC level is a fabric login (FLOGI). If the ENode and FCF have FIP enabled, then the protocol preempts fabric login by first performing the FIP discovery procedures. The FIP protocol then takes over the fabric login process by encapsulating the FLOGI request and replying in FIP frame format.

FIP offers several advantages when deployed in an FCoE network. Native FC is a point-to-point connection-oriented technology. Ethernet, on the other hand, is a multipoint connectionless technology. FIP brings manageability to fibre channel over Ethernet networks by automatically discovering neighbors' and functional roles (that is, ENode vs. FCF). In addition, it automatically logs in to an available FCF after successful discovery. In addition to other maintenance capabilities, including keepalive and clear virtual links, FIP enables FC-unaware data center bridges to snoop on control plane communications.

With or without FIP, the same number of FLOGIs is still required per VN_Port – VF_Port connection. With FIP, however, the strain imposed on the FCF is heavier. In a large data center, the number of FLOGIs to be processed by the FCF still reaches hundreds and even thousands, per port.

For each VN_Port to instantiate, the ENode's FCoE controller on an Ethernet port initiates one FLOGI to an FCF capable of instantiating VF_Ports. Each VN_Port may be further virtualized by acquiring additional port IDs for each VN_Port via a process called N_Port Virtualization, using the F_Port Discovery (FDISC) request/reply sequence.

Objective

The objective of this test is to verify the ability of an FCF (that is, the DUT) to accept a large number of VN_Port instantiation requests with FIP enabled. Both FLOGI instantiated and FDISC instantiated VN_Ports will be enabled and configured to request logins. Before logins, FIP discovery will take place in which the FCF will assign a fabric provided MAC Address (FPMA) to each VN_Port. During the FIP link instantiation phase, FLOGIs will first be transmitted from each emulated FLOGI VN_Port. As the FCF accepts each FLOGI VN_Port, FDISCs will subsequently be transmitted from each emulated FDISC VN_Port.

The FCF should assign an FPMA to all VN_Ports, and accept all login requests using an LS_ACC, and properly assign non-overlapping FC IDs (used as S_ID and/or D_ID) to each VN_Port.

Setup

Two Ixia ports are used in this test, one on the initiator side, and one on the target side.

The objective of this test is to instantiate 10,000 VN_Ports.

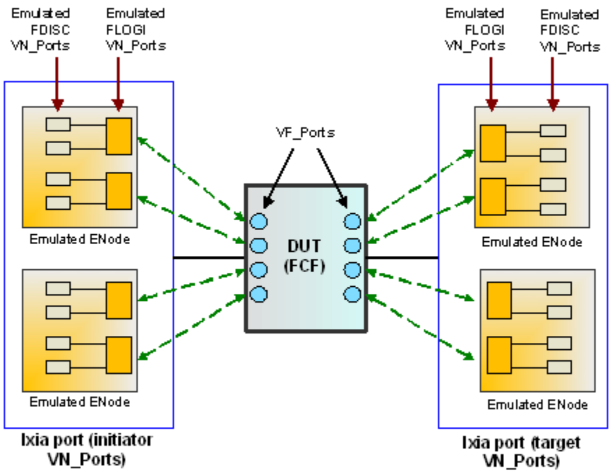

Figure 48. FIP Fabric Login stress test topology

Step-by-Step Instructions

1. Reserve two Ixia ports.

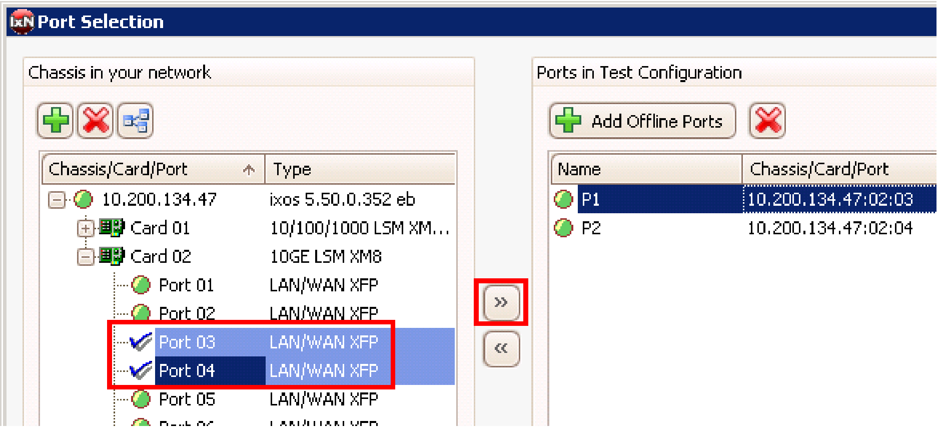

Chassis/Card/Port	Type
10.200.134.47	ixos 5.50.0.352 eb
Card 01	10/100/1000 LSM XM...
Card 02	10GE LSM XM8
Port 01	LAN/WAN XFP
Port 02	LAN/WAN XFP
Port 03	LAN/WAN XFP
Port 04	LAN/WAN XFP
Port 05	LAN/WAN XFP

Name	Chassis/Card/Port
P1	10.200.134.47:02:03
P2	10.200.134.47:02:04

Figure 49. FIP fabric login stress test – reserve two ports

2. Set the Port **Type** to either **10GE LAN – FCoE** or **10GE WAN – FCoE**.

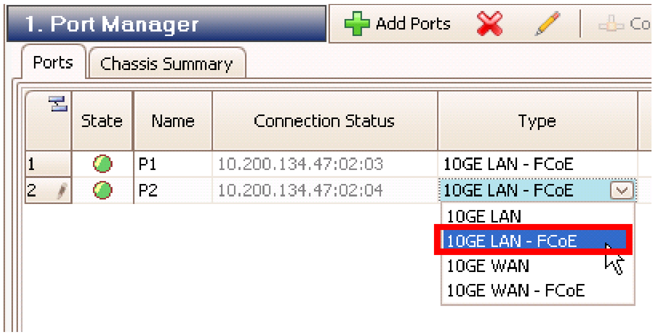

	State	Name	Connection Status	Type
1	●	P1	10.200.134.47:02:03	10GE LAN - FCoE
2	●	P2	10.200.134.47:02:04	10GE LAN - FCoE

Figure 50. FIP fabric login stress test – setting port type to FCoE

3. (Optional) Enable DCBX. See the test case ‘Test Case: DCBX Functionality’ for DCBX configuration procedures.

4. Open the IxNetwork **Protocol Wizards** window, and run the **FCoE/FC Node Wizard**.

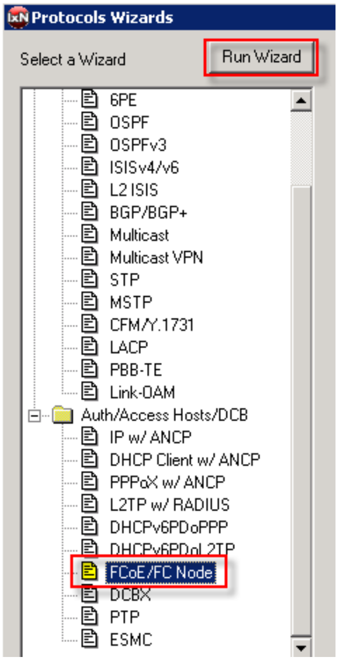

Figure 51. FIP fabric login stress test - FCoE/FC Node protocol wizard

5. **FCoE/FC Node Wizard – Port Selection**: Configure Ixia Port 1 to be the **initiator** side port, and Ixia Port 2 to be the **Target** side port.

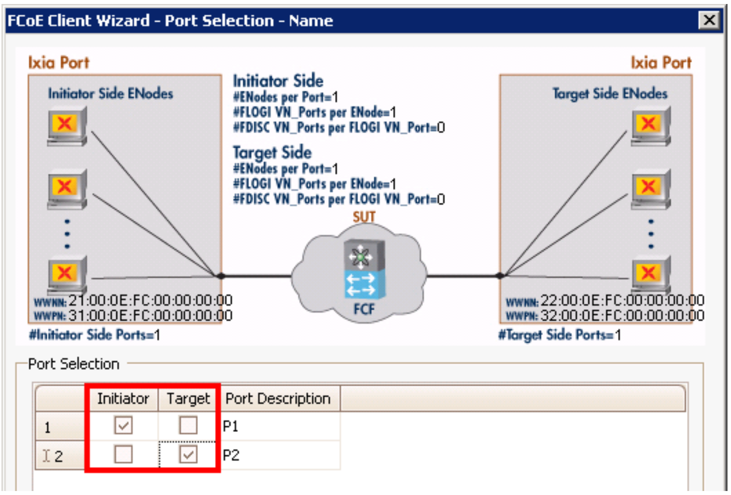
Ixia Port
Initiator Side ENodes
Initiator Side
#ENodes per Port=1
#FLOGI VN_Ports per ENode=1
#FDISC VN_Ports per FLOGI VN_Port=0
Target Side
#ENodes per Port=1
#FLOGI VN_Ports per ENode=1
#FDISC VN_Ports per FLOGI VN_Port=0
SUT
FCF
Ixia Port
Target Side ENodes
WWNN: 21:00:0E:FC:00:00:00:00
WWPN: 31:00:0E:FC:00:00:00:00
#Initiator Side Ports=1
WWNN: 22:00:0E:FC:00:00:00:00
WWPN: 32:00:0E:FC:00:00:00:00
#Target Side Ports=1

FCoE Client Wizard - Port Selection - Name

Port Selection

	Initiator	Target	Port Description
1	☑	☐	P1
2	☐	☑	P2

Figure 52. FIP Fabric Login stress test – FCoE/FC Node Wizard – port selection

6. **FCoE/FC Node Wizard – Flow Control**: Leave the default values as this is beyond the scope of the test.

7. **FCoE/FC Node Wizard – MAC**: Set the starting MAC Address value and its increment behavior across ENodes. The **First MAC Address** value is the MAC address assigned to the first ENode. The **Increment By** value is the MAC address increment step for all subsequent ENodes on the same port. The **Range Increment Step** value is the MAC address increment step for ENodes across ports.

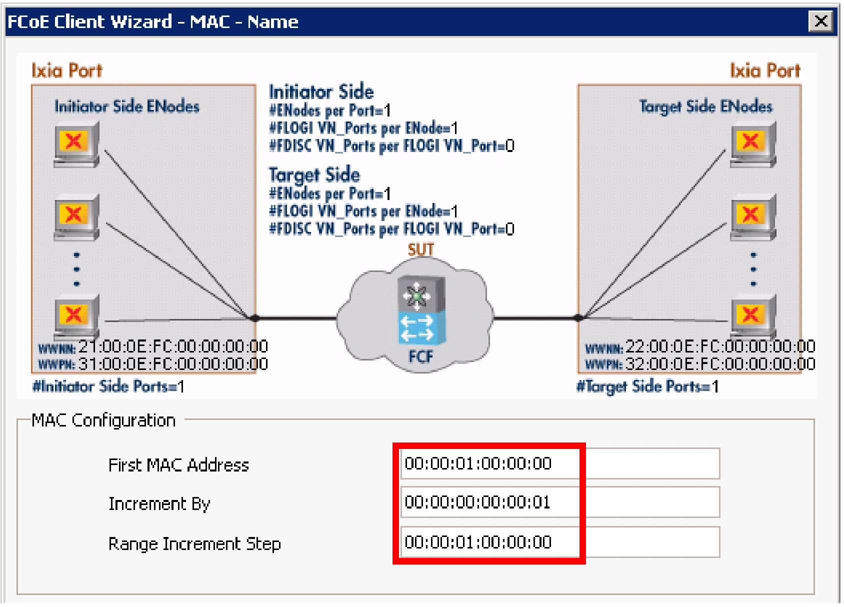
Ixia Port
Initiator Side ENodes
Initiator Side
#ENodes per Port=1
#FLOGI VN_Ports per ENode=1
#FDISC VN_Ports per FLOGI VN_Port=0
Target Side
#ENodes per Port=1
#FLOGI VN_Ports per ENode=1
#FDISC VN_Ports per FLOGI VN_Port=0
SUT
FCF
Ixia Port
Target Side ENodes
WWNN: 21:00:0E:FC:00:00:00:00
WWPN: 31:00:0E:FC:00:00:00:00
#Initiator Side Ports=1
WWNN: 22:00:0E:FC:00:00:00:00
WWPN: 32:00:0E:FC:00:00:00:00
#Target Side Ports=1

FCoE Client Wizard - MAC - Name

MAC Configuration

First MAC Address 00:00:01:00:00:00

Increment By 00:00:00:00:00:01

Range Increment Step 00:00:01:00:00:00

Figure 53. FIP fabric login stress test – FCoE/FC Node Wizard – MAC

8. **FCoE/FC Node Wizard – VLAN**: Leave the default values as this is beyond the scope of the test.

9. **FCoE/FC Node Wizard – FIP**: Configure the FCoE/FIP global protocol features.

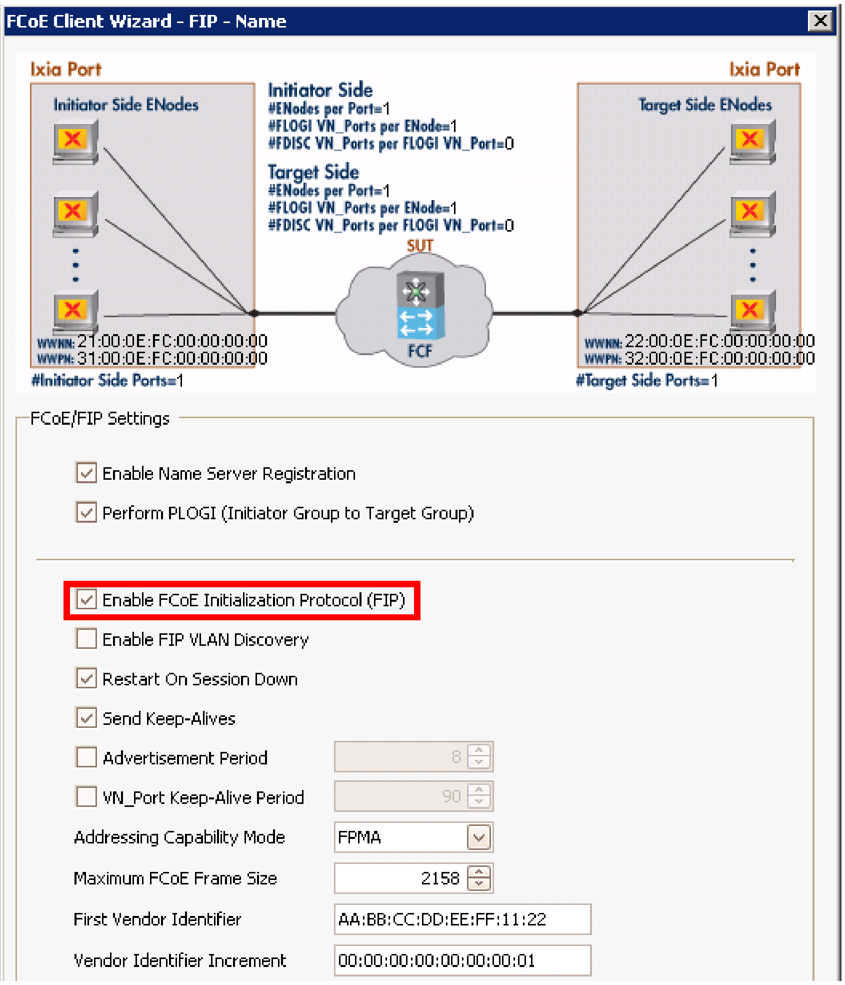

FCoE Client Wizard - FIP - Name

Ixia Port
Initiator Side ENodes
Initiator Side
#ENodes per Port=1
#FLOGI VN_Ports per ENode=1
#FDISC VN_Ports per FLOGI VN_Port=0
Target Side
#ENodes per Port=1
#FLOGI VN_Ports per ENode=1
#FDISC VN_Ports per FLOGI VN_Port=0
SUT
FCF
Ixia Port
Target Side ENodes
WWNN: 21:00:0E:FC:00:00:00:00
WWPN: 31:00:0E:FC:00:00:00:00
#Initiator Side Ports=1
WWNN: 22:00:0E:FC:00:00:00:00
WWPN: 32:00:0E:FC:00:00:00:00
#Target Side Ports=1

FCoE/FIP Settings

☑ Enable Name Server Registration

☑ Perform PLOGI (Initiator Group to Target Group)

☑ Enable FCoE Initialization Protocol (FIP)

□ Enable FIP VLAN Discovery

☑ Restart On Session Down

☑ Send Keep-Alives

□ Advertisement Period: 8

□ VN_Port Keep-Alive Period: 90

Addressing Capability Mode: FPMA

Maximum FCoE Frame Size: 2158

First Vendor Identifier: AA:BB:CC:DD:EE:FF:11:22

Vendor Identifier Increment: 00:00:00:00:00:00:00:01

Figure 54. FIP fabric login stress test – FCoE/FC Node Wizard – FIP

Select the **Enable Name Server Registration**, **Perform PLOGI**, and **Enable FCoE Initialization Protocol (FIP)** check boxes. Clear the **Enable FIP VLAN Discovery** check box. Leave the **Addressing Capability Mode** at **FPMA** (or the mode that the DUT supports), and **FIP Max FCoE Size** at *2158* (or the value that the DUT supports).

10. **FCoE/FC Node Wizard – Initiator Side VN_Ports**: Configure the number of ENodes, FLOGI VN_Ports and FDISC VN_Ports, as well as the OUI, Node and Port WWNs associated with each VN_Port.

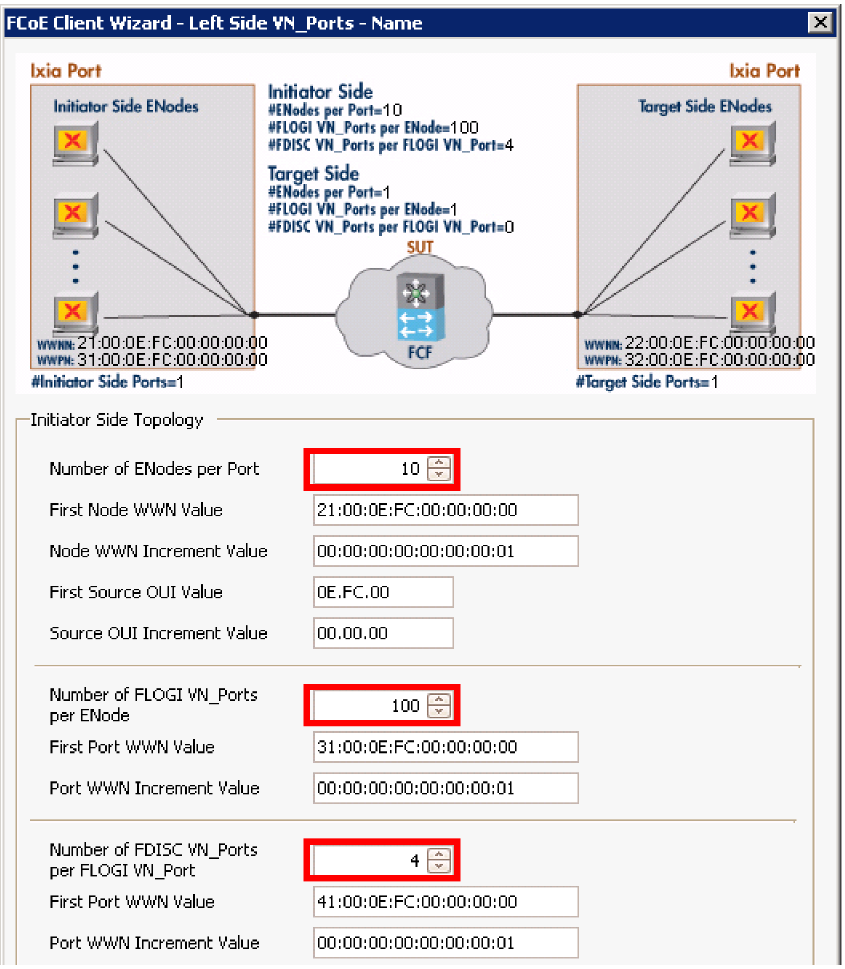

FCoE Client Wizard - Left Side VN_Ports - Name

Ixia Port
Initiator Side ENodes
WWNN: 21:00:0E:FC:00:00:00:00
WWPN: 31:00:0E:FC:00:00:00:00
#Initiator Side Ports=1
Initiator Side
#ENodes per Port=10
#FLOGI VN_Ports per ENode=100
#FDISC VN_Ports per FLOGI VN_Port=4
Target Side
#ENodes per Port=1
#FLOGI VN_Ports per ENode=1
#FDISC VN_Ports per FLOGI VN_Port=0
SUT
FCF
Ixia Port
Target Side ENodes
WWNN: 22:00:0E:FC:00:00:00:00
WWPN: 32:00:0E:FC:00:00:00:00
#Target Side Ports=1

Initiator Side Topology

Number of ENodes per Port	10
First Node WWN Value	21:00:0E:FC:00:00:00:00
Node WWN Increment Value	00:00:00:00:00:00:00:01
First Source OUI Value	0E.FC.00
Source OUI Increment Value	00.00.00
Number of FLOGI VN_Ports per ENode	100
First Port WWN Value	31:00:0E:FC:00:00:00:00
Port WWN Increment Value	00:00:00:00:00:00:00:01
Number of FDISC VN_Ports per FLOGI VN_Port	4
First Port WWN Value	41:00:0E:FC:00:00:00:00
Port WWN Increment Value	00:00:00:00:00:00:00:01

Figure 55. FIP fabric login stress test – FCoE/FC Node Wizard – initiator side VN_Ports

To configure the target for this test, set the following values:

a. **Number of ENodes per Port**: *10*

b. **Number of FLOGI VN_Ports per ENode**: *100*

c. **Number of FDISC VN_Ports per FLOGI VN_Port**: *4*

11. **FCoE/FC Node Wizard – Target Side VN_Ports**: Configure a topology on the target side that is symmetric to the initiator side by simply selecting the **Same as Initiator Side Topology** check box. IxNetwork will automatically assign OUI, Node and Port WWNs that are globally unique.

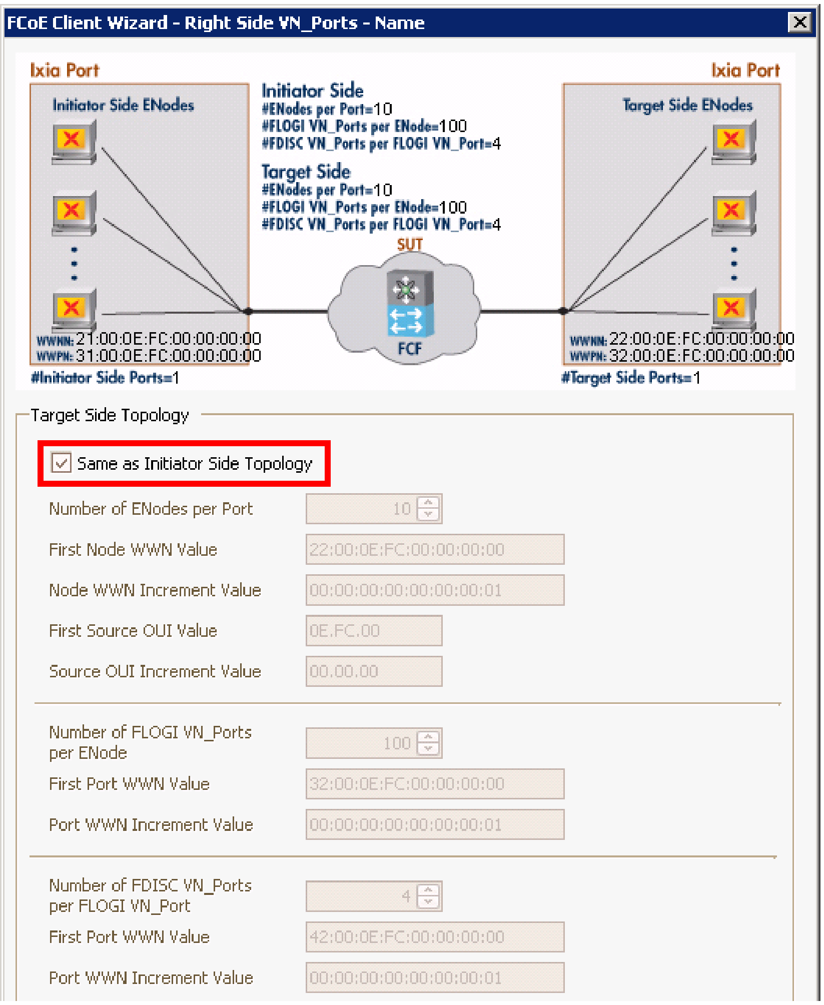

FCoE Client Wizard - Right Side VN_Ports - Name
Ixia Port
Initiator Side ENodes
Initiator Side
#ENodes per Port=10
#FLOGI VN_Ports per ENode=100
#FDISC VN_Ports per FLOGI VN_Port=4
Target Side
#ENodes per Port=10
#FLOGI VN_Ports per ENode=100
#FDISC VN_Ports per FLOGI VN_Port=4
SUT
FCF
Ixia Port
Target Side ENodes
WWNN: 21:00:0E:FC:00:00:00:00
WWPN: 31:00:0E:FC:00:00:00:00
#Initiator Side Ports=1
WWNN: 22:00:0E:FC:00:00:00:00
WWPN: 32:00:0E:FC:00:00:00:00
#Target Side Ports=1

Target Side Topology

☑ Same as Initiator Side Topology

Number of ENodes per Port: 10

First Node WWN Value: 22:00:0E:FC:00:00:00:00

Node WWN Increment Value: 00:00:00:00:00:00:00:01

First Source OUI Value: 0E.FC.00

Source OUI Increment Value: 00.00.00

Number of FLOGI VN_Ports per ENode: 100

First Port WWN Value: 32:00:0E:FC:00:00:00:00

Port WWN Increment Value: 00:00:00:00:00:00:00:01

Number of FDISC VN_Ports per FLOGI VN_Port: 4

First Port WWN Value: 42:00:0E:FC:00:00:00:00

Port WWN Increment Value: 00:00:00:00:00:00:00:01

Figure 56. FIP fabric login stress test – FCoE/FC Node Wizard – target side VN_Ports

12. FCoE/FC Node Wizard – Name: Assign a name to this test configuration, and click **Generate and Overwrite Existing Configuration** to apply the configurations defined in this test.

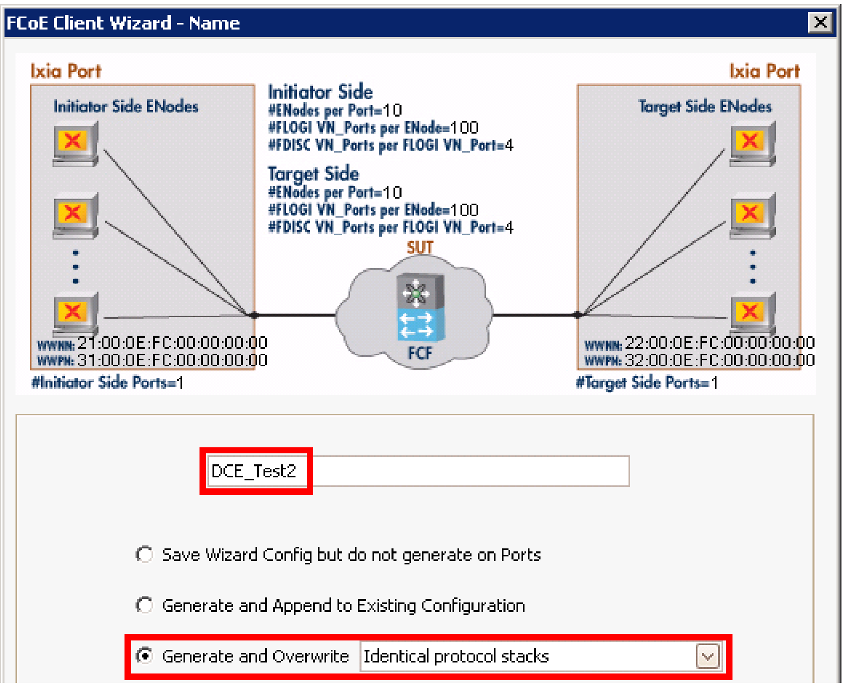

FCoE Client Wizard - Name

Ixia Port
Initiator Side ENodes
Initiator Side
#ENodes per Port=10
#FLOGI VN_Ports per ENode=100
#FDISC VN_Ports per FLOGI VN_Port=4
Target Side
#ENodes per Port=10
#FLOGI VN_Ports per ENode=100
#FDISC VN_Ports per FLOGI VN_Port=4
SUT
FCF
Ixia Port
Target Side ENodes
WWNN: 21:00:0E:FC:00:00:00:00
WWPN: 31:00:0E:FC:00:00:00:00
#Initiator Side Ports=1
WWNN: 22:00:0E:FC:00:00:00:00
WWPN: 32:00:0E:FC:00:00:00:00
#Target Side Ports=1

DCE_Test2

○ Save Wizard Config but do not generate on Ports

○ Generate and Append to Existing Configuration

● Generate and Overwrite: Identical protocol stacks

Figure 57. FIP fabric login stress test – FCoE/FC Node Wizard – name

13. Review the configuration by validating the emulated ENodes, FLOGI VN_Ports and FDISC VN_Ports created. IxNetwork provides tabs for several categories to organize the configuration: **All**, **FIP**, **VN_Port (FLOGI)**, **VN_Port (FDISC)**, **MAC**, and **VLAN**. The **All** tab should show that ten ENodes (named **VNPORT-FDISC-Rxx**) have been created on each Ixia port. The **FIP** tab should show that no ENodes have FIP enabled. The **VN_Port (FLOGI)** tab should show each ENode has 100 VN_Ports configured. In addition, only ENodes on Ixia port 1 have the **PLOGI Target** defined. The **VN_Port (FDISC)** tab should show that each FLOGI VN_Port has 4 FDISC VN_Ports configured.

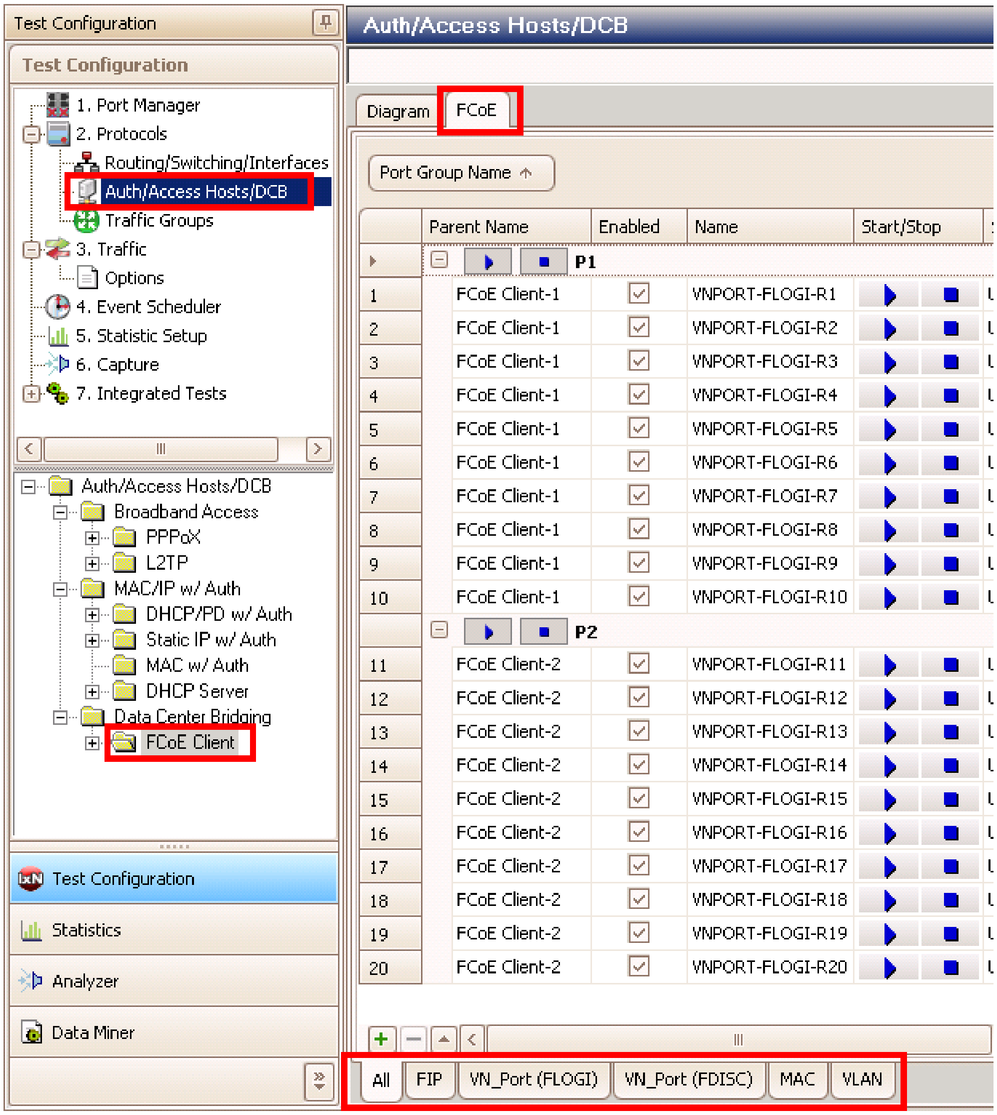

	Parent Name	Enabled	Name	Start/Stop
	P1			▶ ■
1	FCoE Client-1	☑	VNPORT-FLOGI-R1	▶ ■
2	FCoE Client-1	☑	VNPORT-FLOGI-R2	▶ ■
3	FCoE Client-1	☑	VNPORT-FLOGI-R3	▶ ■
4	FCoE Client-1	☑	VNPORT-FLOGI-R4	▶ ■
5	FCoE Client-1	☑	VNPORT-FLOGI-R5	▶ ■
6	FCoE Client-1	☑	VNPORT-FLOGI-R6	▶ ■
7	FCoE Client-1	☑	VNPORT-FLOGI-R7	▶ ■
8	FCoE Client-1	☑	VNPORT-FLOGI-R8	▶ ■
9	FCoE Client-1	☑	VNPORT-FLOGI-R9	▶ ■
10	FCoE Client-1	☑	VNPORT-FLOGI-R10	▶ ■
	P2			▶ ■
11	FCoE Client-2	☑	VNPORT-FLOGI-R11	▶ ■
12	FCoE Client-2	☑	VNPORT-FLOGI-R12	▶ ■
13	FCoE Client-2	☑	VNPORT-FLOGI-R13	▶ ■
14	FCoE Client-2	☑	VNPORT-FLOGI-R14	▶ ■
15	FCoE Client-2	☑	VNPORT-FLOGI-R15	▶ ■
16	FCoE Client-2	☑	VNPORT-FLOGI-R16	▶ ■
17	FCoE Client-2	☑	VNPORT-FLOGI-R17	▶ ■
18	FCoE Client-2	☑	VNPORT-FLOGI-R18	▶ ■
19	FCoE Client-2	☑	VNPORT-FLOGI-R19	▶ ■
20	FCoE Client-2	☑	VNPORT-FLOGI-R20	▶ ■

Figure 58. FIP Fabric Login stress test – FCoE client configuration GUI view

14. Optionally, configure the login **Setup Rate** and logout **Teardown Rate**. These values determine the rate at which FLOGIs or FDISCs are transmitted to the FCF, and the rate at which LOGOs are transmitted to the FCF.

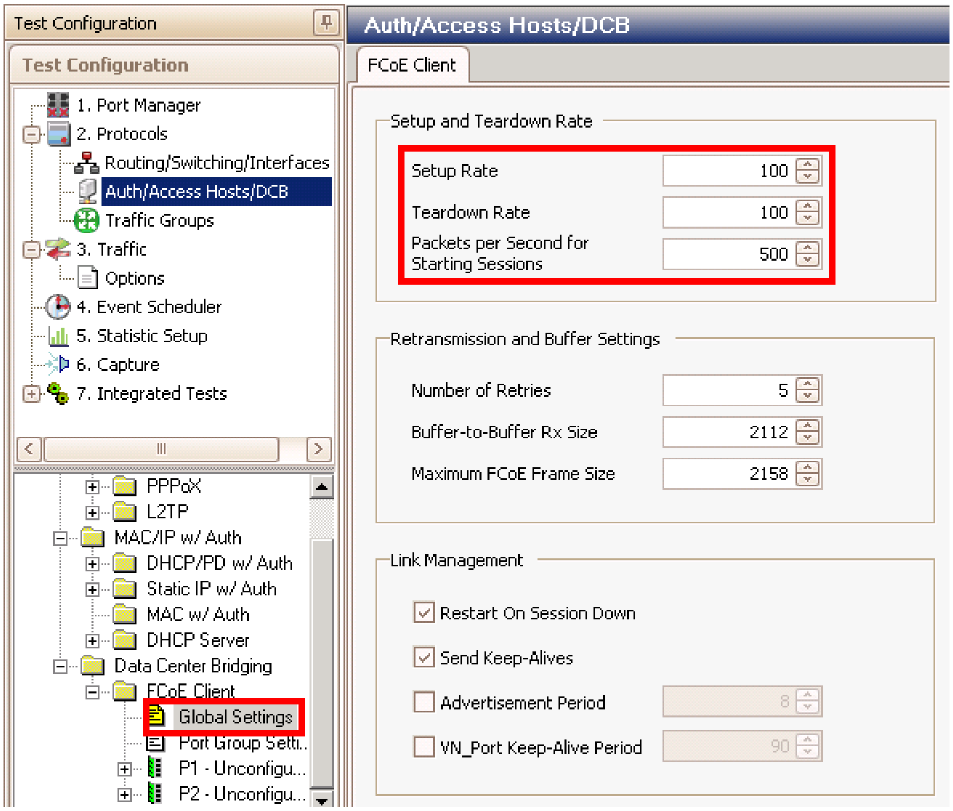

Setup and Teardown Rate

Setup Rate 100

Teardown Rate 100

Packets per Second for Starting Sessions 500

Retransmission and Buffer Settings

Number of Retries 5

Buffer-to-Buffer Rx Size 2112

Maximum FCoE Frame Size 2158

Link Management

☑ Restart On Session Down

☑ Send Keep-Alives

☐ Advertisement Period 8

☐ VN_Port Keep-Alive Period 90

Figure 59. FIP Fabric Login Stress Test – Login Setup and Logout Teardown Rates

15. Start the emulation and verify test results.

16. Wait until all VN_Ports are instantiated. This process can take seconds or minutes depending on the configured login setup rate.

Test Variables

Performance Variable	Description
Number of ENodes per Port	Each ENode instantiates a set of FLOGI VN_Ports and FDISC VN_Ports. Increasing the **Number of ENodes per Port** will stress the DUT's (FCF) ability to process FCoE login requests.
Number of FLOGI VN_Ports per ENode	A VN_Port can be instantiated by FLOGI or FDISC. Increasing the **Number of FLOGI VN_Ports per ENode** will stress the DUT's (FCF) ability to process FLOGI requests.
Number of FDISC VN_Ports per FLOGI VN_Port	A VN_Port can be instantiated by FLOGI or FDISC. Increasing the **Number of FDISC VN_Ports per FLOGI VN_Port** will stress the DUT's (FCF) ability to process FDISC requests.
Name Server Registration	Each VN_Port registers its attributes with the name server (commonly the FCF) after it logs in to the fabric. In addition, each VN_Port queries the name server for each destination VN_Port. Enabling this option puts additional stress on the DUT to process name server registrations and queries.
Perform PLOGI	PLOGI is a connection request made from a source VN_Port to a destination VN_Port. Enabling this option puts additional stress on the DUT to process PLOGI request and replies.
Enable FIP	Enabling FIP will stress the DUT's ability to process discovery solicitations and periodic keepalives from each VN_Port.
Setup/Teardown Rate	(Optional) The rate at which VN_Ports request to login to a fabric. Increasing this number will stress the DUT's ability to process a large volume of Fabric Login and logout requests.

Results Analysis

Use the IxNetwork **Statistics** view to verify the following FCoE and FIP packets statistics:

- **FIP Discovery Advertisements Rx**: *1,000*

  There are 1,000 FLOGI VN_Ports configured. Each of the 1,000 VN_Ports will transmit a **FIP Discovery Solicitation**. The DUT must respond to each solicitation with a **FIP Discovery Advertisement**.

- **FLOGI LS_ACC Rx**: *1,000*

  There are 100 FLOGI VN_Ports configured for each of the 10 ENodes. Each VN_Port is instantiated by transmitting one FLOGI. The DUT accepts each FLOGI by responding with an LS_ACC.

- **FDISC LS_ACC Rx**: *4,000*

  There are 4 FDISC VN_Ports configured for each of the 1,000 FLOGI VN_Ports. Each FDISC VN_Port is instantiated by transmitting one FDISC. The DUT accepts each FDISC by responding with an LS_ACC.

- **PLOGI LS_ACC Rx**: *5,000*

  There are 5,000 VN_Ports in total. Each of the 1,000 FLOGI VN_Ports transmits one PLOGI to its target VN_Port. Each of the 4,000 FDISC VN_Ports transmits one PLOGI to its target VN_Port. The target side Ixia port responds to each PLOGI with an LS_ACC. The DUT forwards each PLOGI LS_ACC to the initiator side Ixia port.

- **NS Registration Successful**: *5,000*

  Each of the 5,000 VN_Ports registers with the Name Server (that is, the DUT in this test). The DUT acknowledges each of the 5,000 NS Registrations.

- **Interfaces Up**: *5,000*

  All 5,000 VN_Ports should be instantiated.

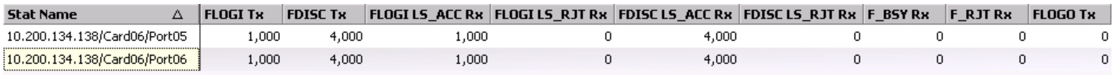

Stat Name △	FLOGI Tx	FDISC Tx	FLOGI LS_ACC Rx	FLOGI LS_RJT Rx	FDISC LS_ACC Rx	FDISC LS_RJT Rx	F_BSY Rx	F_RJT Rx	FLOGO Tx
10.200.134.138/Card06/Port05	1,000	4,000	1,000	0	4,000	0	0	0	0
10.200.134.138/Card06/Port06	1,000	4,000	1,000	0	4,000	0	0	0	0

Figure 60. FIP fabric login stress test – FCoE client aggregate statistics view

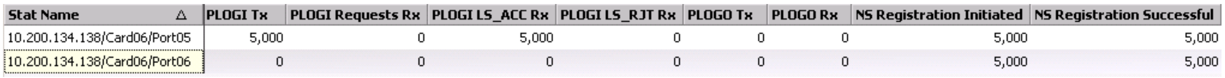

Stat Name △	PLOGI Tx	PLOGI Requests Rx	PLOGI LS_ACC Rx	PLOGI LS_RJT Rx	PLOGO Tx	PLOGO Rx	NS Registration Initiated	NS Registration Successful
10.200.134.138/Card06/Port05	5,000	0	5,000	0	0	0	5,000	5,000
10.200.134.138/Card06/Port06	0	0	0	0	0	0	5,000	5,000

Figure 61. FIP fabric login stress test – FCoE client aggregate statistics view (cont'd)

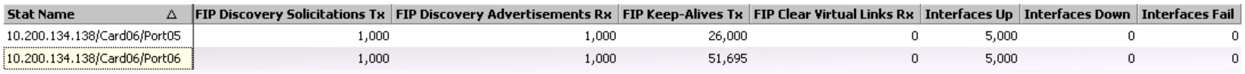

Stat Name △	FIP Discovery Solicitations Tx	FIP Discovery Advertisements Rx	FIP Keep-Alives Tx	FIP Clear Virtual Links Rx	Interfaces Up	Interfaces Down	Interfaces Fail
10.200.134.138/Card06/Port05	1,000	1,000	26,000	0	5,000	0	0
10.200.134.138/Card06/Port06	1,000	1,000	51,695	0	5,000	0	0

Figure 62. FIP fabric login stress test – FCoE client aggregate statistics view (Cont'd)

Use the IxNetwork per session statistics view to verify the following FCoE control plane exchanges for each VN_Port. Right click from the aggregate statistics above and select **DrillDown per Session**.

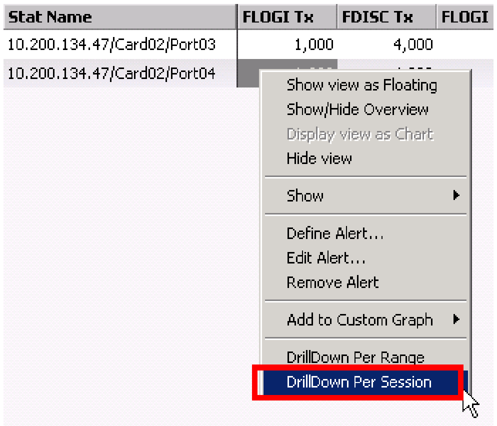

Stat Name	FLOGI Tx	FDISC Tx	FLOGI
10.200.134.47/Card02/Port03	1,000	4,000	
10.200.134.47/Card02/Port04			

Figure 63. FCoE fabric login stress test - DrillDown per Session

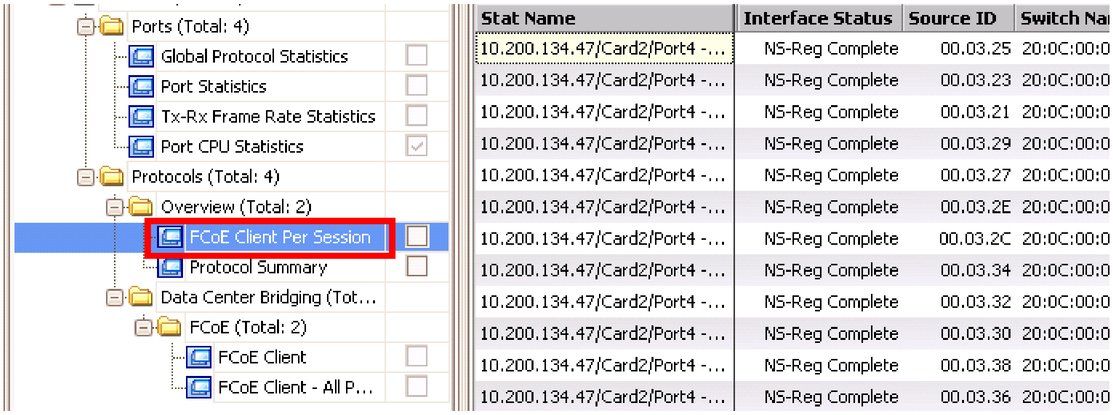

Stat Name	Interface Status	Source ID	Switch Na
10.200.134.47/Card2/Port4 -...	NS-Reg Complete	00.03.25	20:0C:00:0
10.200.134.47/Card2/Port4 -...	NS-Reg Complete	00.03.23	20:0C:00:0
10.200.134.47/Card2/Port4 -...	NS-Reg Complete	00.03.21	20:0C:00:0
10.200.134.47/Card2/Port4 -...	NS-Reg Complete	00.03.29	20:0C:00:0
10.200.134.47/Card2/Port4 -...	NS-Reg Complete	00.03.27	20:0C:00:0
10.200.134.47/Card2/Port4 -...	NS-Reg Complete	00.03.2E	20:0C:00:0
10.200.134.47/Card2/Port4 -...	NS-Reg Complete	00.03.2C	20:0C:00:0
10.200.134.47/Card2/Port4 -...	NS-Reg Complete	00.03.34	20:0C:00:0
10.200.134.47/Card2/Port4 -...	NS-Reg Complete	00.03.32	20:0C:00:0
10.200.134.47/Card2/Port4 -...	NS-Reg Complete	00.03.30	20:0C:00:0
10.200.134.47/Card2/Port4 -...	NS-Reg Complete	00.03.38	20:0C:00:0
10.200.134.47/Card2/Port4 -...	NS-Reg Complete	00.03.36	20:0C:00:0

Figure 64. FIP fabric login stress test – FCoE client per VN_Port statistics view

- All 10,000 VN_Ports have a unique **Assigned MAC** address assigned by the DUT.
- All 5,000 VN_Ports from the initiator side are associated with the same **Fabric MAC** address, but different than the second 5,000 VN_Ports.
- All 5,000 VN_Ports from the target sideare associated with the same **Fabric MAC** address, but different than the 5,000 VN_Ports from the initiator side.
- All 10,000 VN_Ports are assigned the same **FC Map**.
- All 10,000 VN_Ports are assigned the **Priority** configured on the DUT. Unless configured otherwise, the **Priority** value should be the same for all VN_Ports.
- All 10,000 VN_Ports are associated with the **Fabric Name** and **Switch Name** configured on the DUT. Unless configured otherwise, the **Fabric Name** should be the same for all VN_Ports, and the **Switch Name** should be the same for all VN_Ports.

- All 5,000 VN_Ports from the initiator side have **Status** set to **PLOGI OK**. This is the last initialization operation performed by the initiator side VN_Ports.

- All 5,000 VN_Ports from the target side have **Status** set to **NS-Reg OK**. This is the last initialization operation performed by the target side VN_Ports.

- Each 10,000 VN_Port is assigned a unique **Source ID**.

- All 5,000 VN_Ports from the initiator side are assigned a unique **PLOGI Destination ID**. This is the D_ID that the initiator VN_Port uses to perform PLOGI to its target VN_Port.

*Note: Use the **Filter/Sort** function to zoom in on VN_Ports that matches user defined statistics criteria*

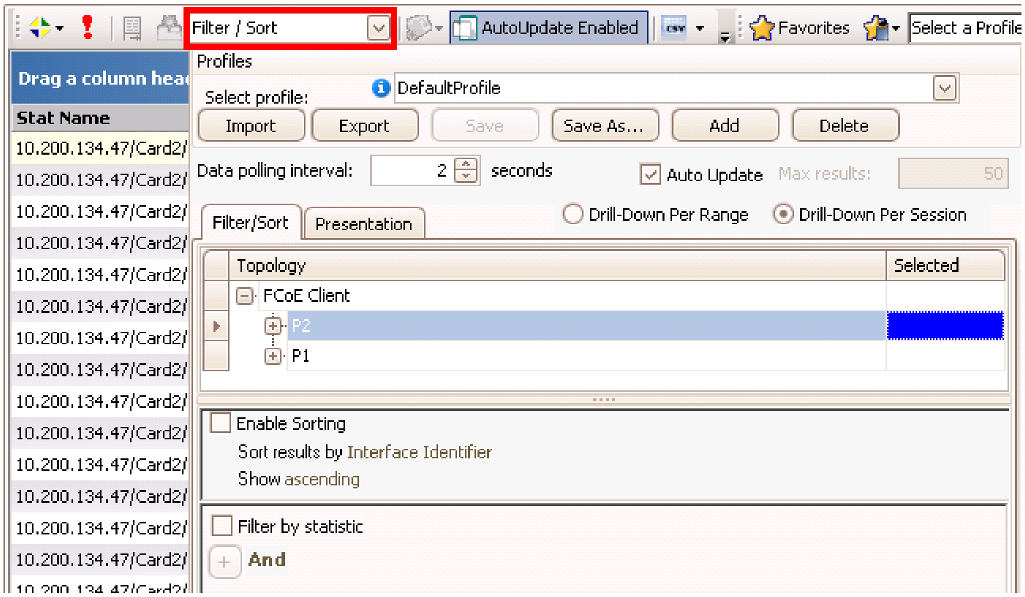

Stat Name
10.200.134.47/Card2/
10.200.134.47/Card2/
10.200.134.47/Card2/
10.200.134.47/Card2/
10.200.134.47/Card2/
10.200.134.47/Card2/
10.200.134.47/Card2/
10.200.134.47/Card2/
10.200.134.47/Card2/
10.200.134.47/Card2/
10.200.134.47/Card2/
10.200.134.47/Card2/
10.200.134.47/Card2/
10.200.134.47/Card2/

Select profile: DefaultProfile

Data polling interval: 2 seconds ☑ Auto Update Max results: 50

○ Drill-Down Per Range ● Drill-Down Per Session

□ Enable Sorting
Sort results by Interface Identifier
Show ascending

□ Filter by statistic

Figure 65. FCoE fabric login stress test - statistics Filter/Sort function

Conclusions

By validating the statistics and control plane exchanges using the features above, the DUT has been proven capable of responding to at least 1,000 FLOGI VN_Ports' Discovery Solicitations per port, and is assigning FPMA and FC MAP to at least 5,000 VN_Ports per port, for the specified topology. In addition, the DUT is capable of instantiating at least 5,000 VN_Ports per port, instantiated by both FIP FLOGI and FIP FDISC, for the specified topology. Furthermore, the DUT is capable of maintaining at least 5,000 VN_Ports per port with periodic FIP keepalives for the specified topology.

Test Case: FIP VLAN Discovery Stress Test

Overview

VLAN is the most frequently deployed network segmentation technique used in LANs today. VLANs are used on campuses, in offices, and by enterprises. As fibre channel converges with Ethernet in data centers, it will inevitably be required to integrate with VLANs as well. To make deployment manageable and maintainable, FIP provides a VLAN discovery and assignment capability that automatically configures an FCoE link on the correct VLAN so that storage traffic will be tagged according to policies.

When a capable ENode connects to the network, FIP VLAN discovery facilitates the request and retrieval of acceptable VLAN IDs from an eligible FCF. FIP carries out the remaining operations over the assigned VLANs following this crucial initial step.

When using FIP VLAN discovery, the strain imposed on the FCF will be increased during the discovery phase. Acting as a stepping stone, FIP VLAN discovery must be executed properly preceding any FIP and FCoE operations. In a large data center, the number of per port VLAN requests and assignments can reach hundreds and even thousands.

Objective

The objective of this test is to verify the ability of an FCF (that is, the DUT) to accept a large number of VN_Port instantiation requests with FIP and FIP VLAN discovery enabled. Both FLOGI instantiated and FDISC instantiated VN_Ports will be enabled and configured to request logins. Before logins, FIP discovery will take place and the FCF will assign an FPMA to each VN_Port. During the FIP Link Instantiation phase, FLOGIs will first be transmitted from each emulated FLOGI VN_Port. As the FCF accepts each FLOGI VN_Port, FDISCs will subsequently be transmitted from each emulated FDISC VN_Port.

The FCF should assign valid VLANs to all VN_Ports during the VLAN discovery phase. Once completed, the FCF should assign an FPMA to all VN_Ports, accept all login requests using an LS_ACC, and properly assign non-overlapping FC IDs (used as S_ID and/or D_ID) to each VN_Port. All operations following VLAN discovery are carried out over the assigned VLANs.

Setup

Two Ixia ports are used in this test, one on the initiator side, and one on the target side.

The objective of this test is to instantiate 10,000 VN_Ports.

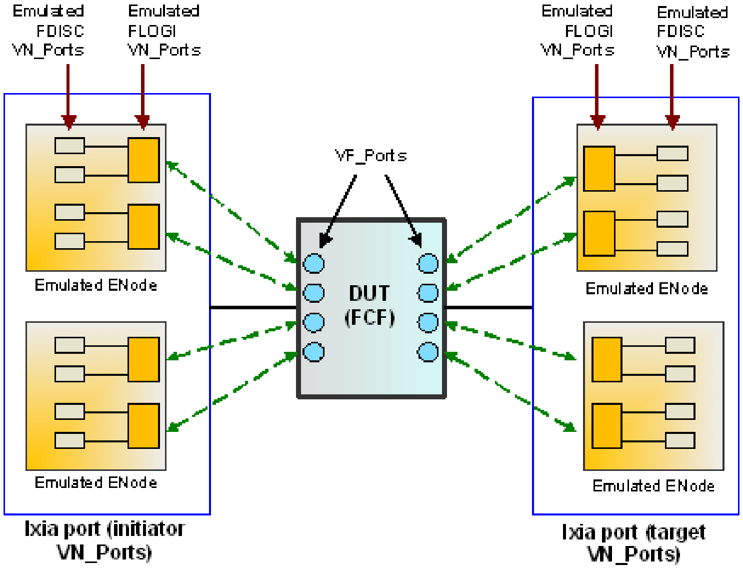

Figure 66. FIP VLAN discovery stress test topology

Step-by-Step Instructions

1. Reserve two Ixia ports.

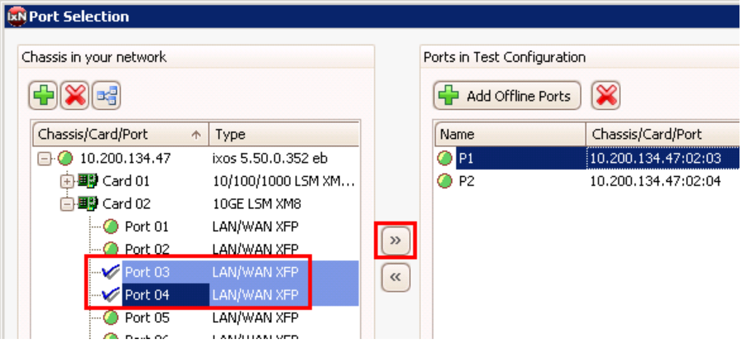

Port Selection

Chassis in your network

Chassis/Card/Port	Type
10.200.134.47	ixos 5.50.0.352 eb
Card 01	10/100/1000 LSM XM...
Card 02	10GE LSM XM8
Port 01	LAN/WAN XFP
Port 02	LAN/WAN XFP
✓ Port 03	LAN/WAN XFP
✓ Port 04	LAN/WAN XFP
Port 05	LAN/WAN XFP

Ports in Test Configuration

Name	Chassis/Card/Port
P1	10.200.134.47:02:03
P2	10.200.134.47:02:04

Figure 67. FIP VLAN discovery - port selection

2. Set the port **Type** to either **10GE LAN – FCoE** or **10GE WAN – FCoE**.

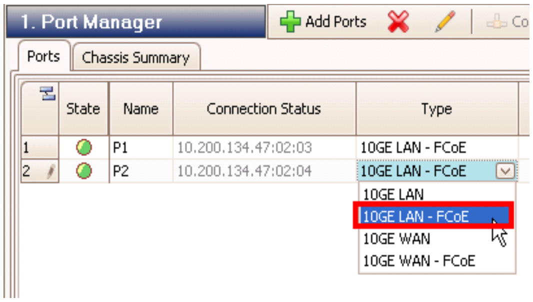

Figure 68. FIP VLAN discovery stress test – setting port type to FCoE

3. (Optional) Enable DCBX. See the test case ‘Test Case: DCBX Functionality’ for DCBX configuration procedures.

4. Open the IxNetwork **Protocol Wizards** window, and run the **FCoE/FC Node Wizard**.

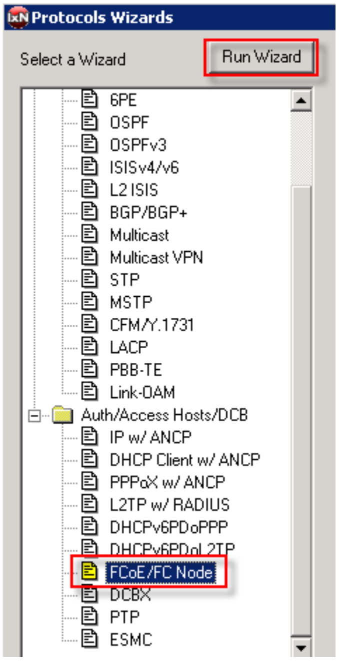

Figure 69. FIP VLAN discovery – FCoE/FC Node protocol wizard

5. **FCoE/FC Node Wizard – Port Selection**: Configure Ixia Port 1 to be the **Initiator** side port, and Ixia Port 2 to be the **Target** side port.

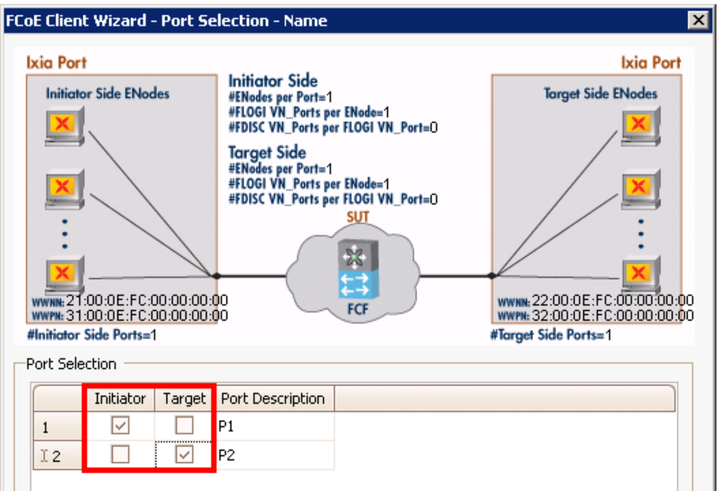
FCoE Client Wizard - Port Selection - Name
Ixia Port
Initiator Side ENodes
Initiator Side
#ENodes per Port=1
#FLOGI VN_Ports per ENode=1
#FDISC VN_Ports per FLOGI VN_Port=0
Target Side
#ENodes per Port=1
#FLOGI VN_Ports per ENode=1
#FDISC VN_Ports per FLOGI VN_Port=0
SUT
FCF
Ixia Port
Target Side ENodes
WWNN: 21:00:0E:FC:00:00:00:00
WWPN: 31:00:0E:FC:00:00:00:00
#Initiator Side Ports=1
WWNN: 22:00:0E:FC:00:00:00:00
WWPN: 32:00:0E:FC:00:00:00:00
#Target Side Ports=1

Port Selection

	Initiator	Target	Port Description
1	☑	☐	P1
2	☐	☑	P2

Figure 70. FIP VLAN discovery stress test – FCoE/FC Node Wizard – port selection

6. **FCoE/FC Node Wizard – Flow Control**: Leave the default values as this is beyond the scope of the test.

7. **FCoE/FC Node Wizard – MAC**: Set the starting MAC Address value and its increment behavior across ENodes. The **First MAC Address** value is the MAC address assigned to the first ENode. The **Increment By** value is the MAC address increment step for all subsequent ENodes on the same port. The **Range Increment Step** value is the MAC address increment step for ENodes across ports.

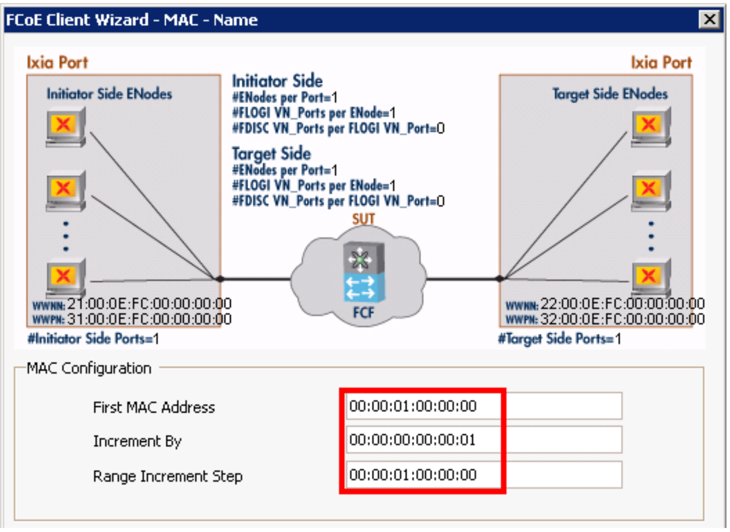
FCoE Client Wizard - MAC - Name
Ixia Port
Initiator Side ENodes
Initiator Side
#ENodes per Port=1
#FLOGI VN_Ports per ENode=1
#FDISC VN_Ports per FLOGI VN_Port=0
Target Side
#ENodes per Port=1
#FLOGI VN_Ports per ENode=1
#FDISC VN_Ports per FLOGI VN_Port=0
SUT
FCF
Ixia Port
Target Side ENodes
WWNN: 21:00:0E:FC:00:00:00:00
WWPN: 31:00:0E:FC:00:00:00:00
#Initiator Side Ports=1
WWNN: 22:00:0E:FC:00:00:00:00
WWPN: 32:00:0E:FC:00:00:00:00
#Target Side Ports=1

MAC Configuration

First MAC Address	00:00:01:00:00:00
Increment By	00:00:00:00:00:01
Range Increment Step	00:00:01:00:00:00

Figure 71. FIP VLAN discovery stress test – FCoE/FC Node Wizard – MAC

8. FCoE/FC Node Wizard – VLAN: Leave the default values as this is beyond the scope of the test.

9. **FCoE/FC Node Wizard – FIP**: Configure the FCoE/FIP global protocol features. Select the **Enable Name Server Registration**, **Perform PLOGI**, **Enable FCoE Initialization Protocol (FIP)**, and **Enable FIP VLAN Discovery** check boxes. Leave the **Addressing Capability Mode** at **FPMA** (or the mode that the DUT supports), and leave **FIP Max FCoE Size** at *2158* (or the value that the DUT supports).

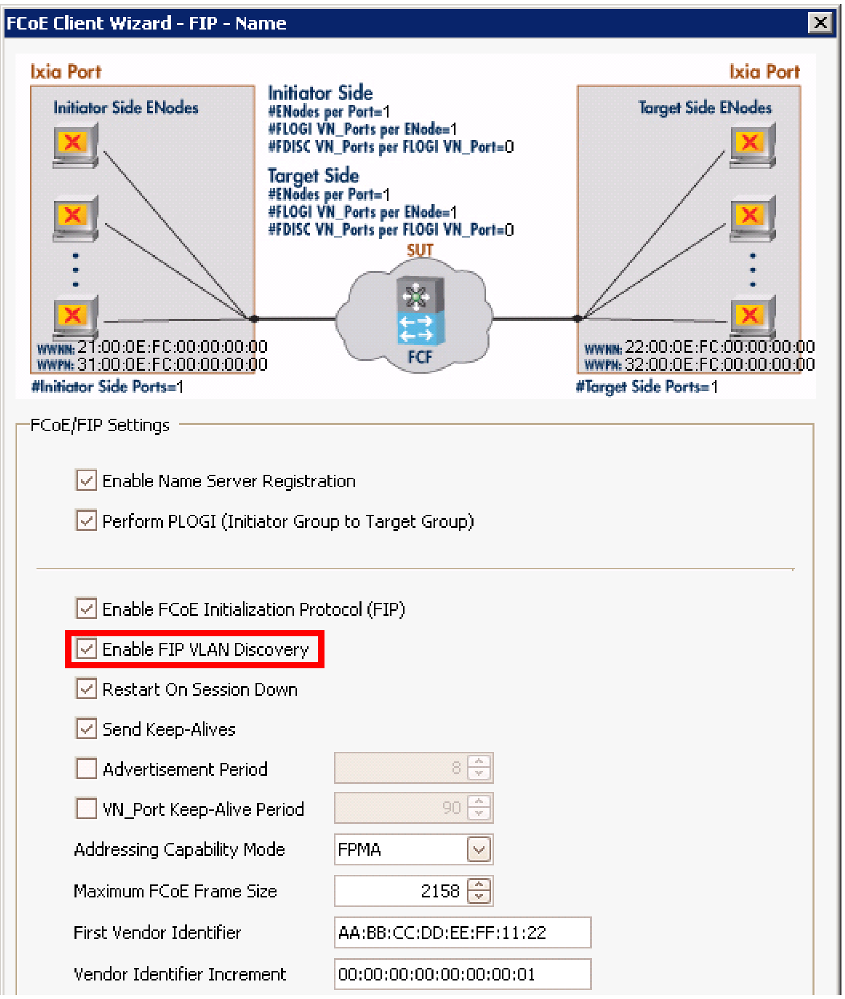
FCoE Client Wizard - FIP - Name
Ixia Port
Initiator Side ENodes
Initiator Side
#ENodes per Port=1
#FLOGI VN_Ports per ENode=1
#FDISC VN_Ports per FLOGI VN_Port=0
Target Side
#ENodes per Port=1
#FLOGI VN_Ports per ENode=1
#FDISC VN_Ports per FLOGI VN_Port=0
SUT
FCF
Ixia Port
Target Side ENodes
WWNN: 21:00:0E:FC:00:00:00:00
WWPN: 31:00:0E:FC:00:00:00:00
#Initiator Side Ports=1
WWNN: 22:00:0E:FC:00:00:00:00
WWPN: 32:00:0E:FC:00:00:00:00
#Target Side Ports=1

FCoE/FIP Settings

☑ Enable Name Server Registration

☑ Perform PLOGI (Initiator Group to Target Group)

☑ Enable FCoE Initialization Protocol (FIP)

☑ Enable FIP VLAN Discovery

☑ Restart On Session Down

☑ Send Keep-Alives

☐ Advertisement Period 8

☐ VN_Port Keep-Alive Period 90

Addressing Capability Mode FPMA

Maximum FCoE Frame Size 2158

First Vendor Identifier AA:BB:CC:DD:EE:FF:11:22

Vendor Identifier Increment 00:00:00:00:00:00:00:01

Figure 72. FIP VLAN discovery stress test – FCoE/FC Node Wizard – FIP

10. **FCoE/FC Node Wizard – Initiator Side VN_Ports**: Configure the number of ENodes, FLOGI VN_Ports and FDISC VN_Ports, as well as the OUI, Node and Port WWNs associated with each VN_Port. To configure the target for this test, set the following values:

    a. **Number of ENodes per Port**: *10***Number of FLOGI VN_Ports per ENode**: *100*

    b. **Number of FDISC VN_Ports per FLOGI VN_Port**: *4*

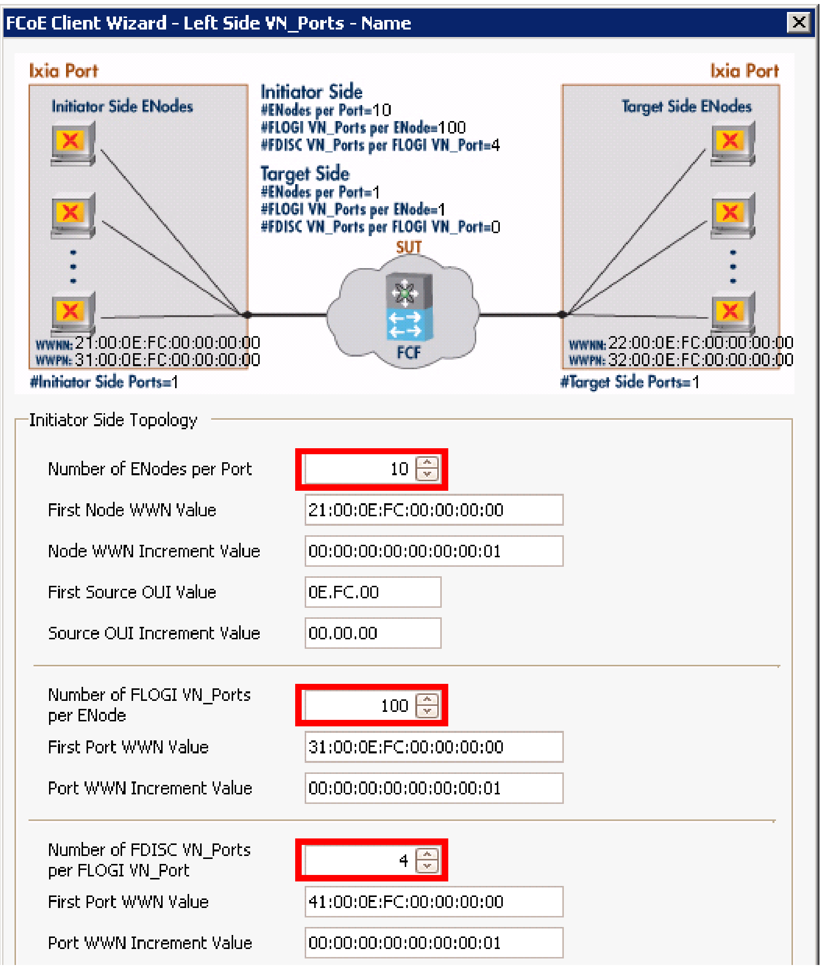

FCoE Client Wizard - Left Side VN_Ports - Name

Ixia Port
Initiator Side ENodes
Initiator Side
#ENodes per Port=10
#FLOGI VN_Ports per ENode=100
#FDISC VN_Ports per FLOGI VN_Port=4
Target Side
#ENodes per Port=1
#FLOGI VN_Ports per ENode=1
#FDISC VN_Ports per FLOGI VN_Port=0
SUT
FCF
Ixia Port
Target Side ENodes
WWNN: 21:00:0E:FC:00:00:00:00
WWPN: 31:00:0E:FC:00:00:00:00
#Initiator Side Ports=1
WWNN: 22:00:0E:FC:00:00:00:00
WWPN: 32:00:0E:FC:00:00:00:00
#Target Side Ports=1

Initiator Side Topology

Number of ENodes per Port	10
First Node WWN Value	21:00:0E:FC:00:00:00:00
Node WWN Increment Value	00:00:00:00:00:00:00:01
First Source OUI Value	0E.FC.00
Source OUI Increment Value	00.00.00
Number of FLOGI VN_Ports per ENode	100
First Port WWN Value	31:00:0E:FC:00:00:00:00
Port WWN Increment Value	00:00:00:00:00:00:00:01
Number of FDISC VN_Ports per FLOGI VN_Port	4
First Port WWN Value	41:00:0E:FC:00:00:00:00
Port WWN Increment Value	00:00:00:00:00:00:00:01

Figure 73. FIP VLAN discovery stress test – FCoE/FC Node Wizard – initiator side VN_Ports

11. **FCoE/FC Node Wizard – Target Side VN_Ports**: Configure a topology on the target side that is symmetric to the initiator side by simply selecting the **Same as Initiator Side Topology** check box. IxNetwork will automatically assign OUI, Node and Port WWNs that are globally unique.

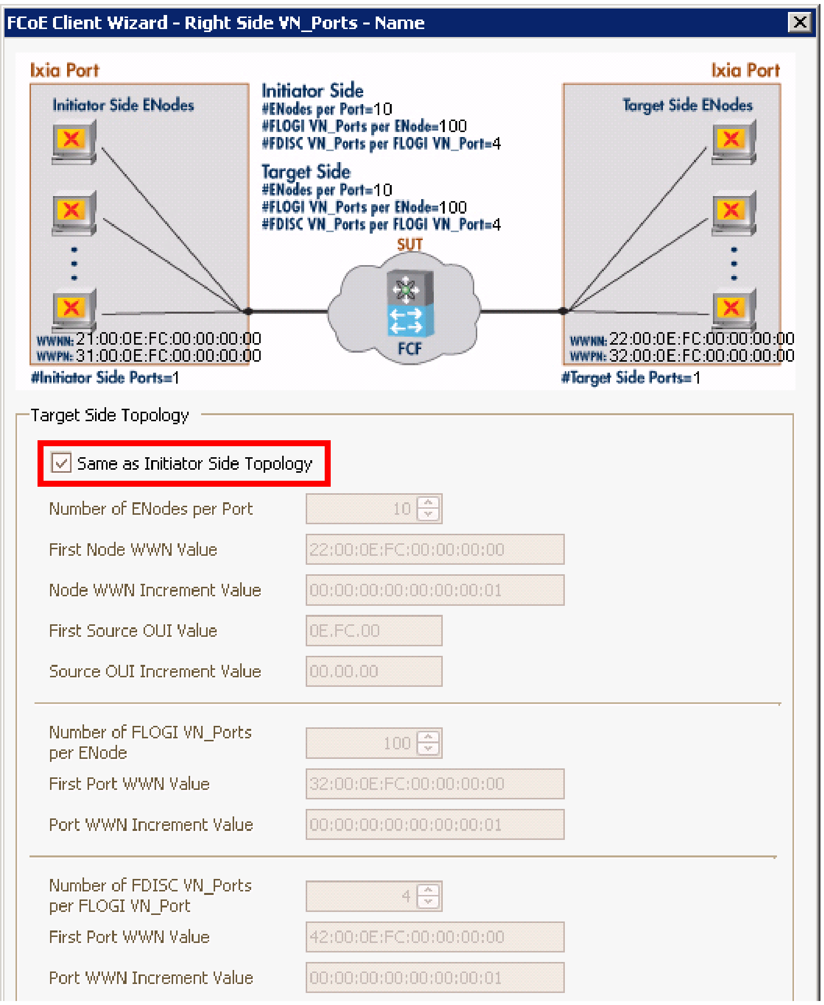
FCoE Client Wizard - Right Side VN_Ports - Name
Ixia Port
Initiator Side ENodes
Initiator Side
#ENodes per Port=10
#FLOGI VN_Ports per ENode=100
#FDISC VN_Ports per FLOGI VN_Port=4
Target Side
#ENodes per Port=10
#FLOGI VN_Ports per ENode=100
#FDISC VN_Ports per FLOGI VN_Port=4
SUT
FCF
Ixia Port
Target Side ENodes
WWNN: 21:00:0E:FC:00:00:00:00
WWPN: 31:00:0E:FC:00:00:00:00
#Initiator Side Ports=1
WWNN: 22:00:0E:FC:00:00:00:00
WWPN: 32:00:0E:FC:00:00:00:00
#Target Side Ports=1

Target Side Topology

☑ Same as Initiator Side Topology

Number of ENodes per Port	10
First Node WWN Value	22:00:0E:FC:00:00:00:00
Node WWN Increment Value	00:00:00:00:00:00:00:01
First Source OUI Value	0E.FC.00
Source OUI Increment Value	00.00.00
Number of FLOGI VN_Ports per ENode	100
First Port WWN Value	32:00:0E:FC:00:00:00:00
Port WWN Increment Value	00:00:00:00:00:00:00:01
Number of FDISC VN_Ports per FLOGI VN_Port	4
First Port WWN Value	42:00:0E:FC:00:00:00:00
Port WWN Increment Value	00:00:00:00:00:00:00:01

Figure 74. FIP VLAN discovery stress test – FCoE/FC Node Wizard – target side VN_Ports

12. **FCoE/FC Node Wizard – Name**: Assign a name to this test configuration, and click **Generate and Overwrite Existing Configuration** to apply the configurations defined in this test.

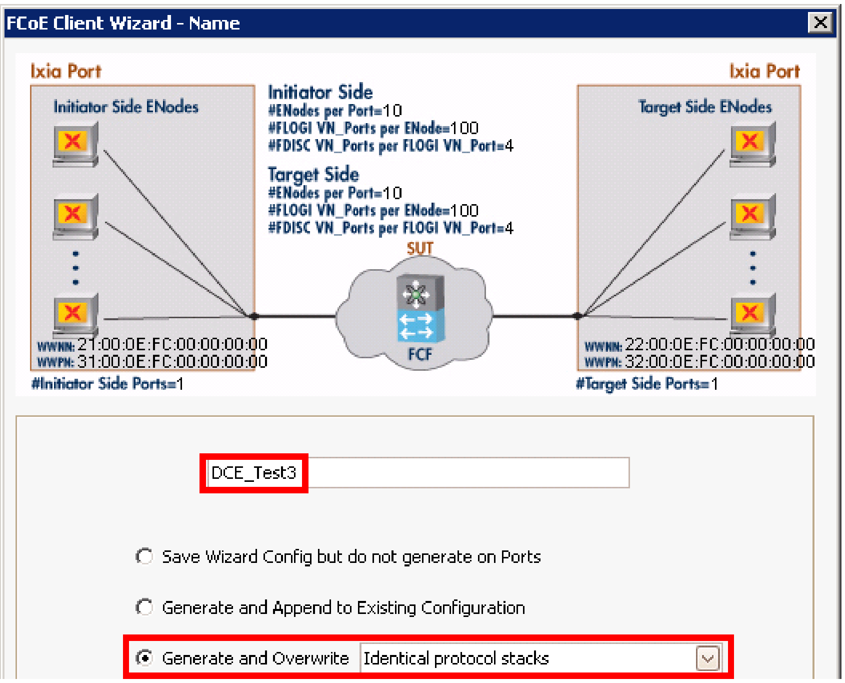
FCoE Client Wizard - Name
Ixia Port
Initiator Side ENodes
Initiator Side
#ENodes per Port=10
#FLOGI VN_Ports per ENode=100
#FDISC VN_Ports per FLOGI VN_Port=4
Target Side
#ENodes per Port=10
#FLOGI VN_Ports per ENode=100
#FDISC VN_Ports per FLOGI VN_Port=4
SUT
FCF
Ixia Port
Target Side ENodes
WWNN: 21:00:0E:FC:00:00:00:00
WWPN: 31:00:0E:FC:00:00:00:00
#Initiator Side Ports=1
WWNN: 22:00:0E:FC:00:00:00:00
WWPN: 32:00:0E:FC:00:00:00:00
#Target Side Ports=1

DCE_Test3

○ Save Wizard Config but do not generate on Ports

○ Generate and Append to Existing Configuration

● Generate and Overwrite Identical protocol stacks

Figure 75. FIP VLAN discovery stress test – FCoE/FC Node Wizard – name

13. Review the configuration by validating the emulated ENodes, FLOGI VN_Ports and FDISC VN_Ports created. IxNetwork provides tabs for several categories to organize the configuration: **All**, **FIP**, **VN_Port (FLOGI)**, **VN_Port (FDISC)**, **MAC**, and **VLAN**. The **All** tab should show that 10 ENodes (named **VNPORT-FDISC-Rxx**) have been created on each Ixia port. The **FIP** tab should show that no ENodes have FIP enabled. The **VN_Port (FLOGI)** tab should show that each ENode has 100 VN_Ports configured. In addition, only ENodes on Ixia port 1 have the **PLOGI Target** defined. The **VN_Port (FDISC)** tab should show that each FLOGI VN_Port has 4 FDISC VN_Ports configured.

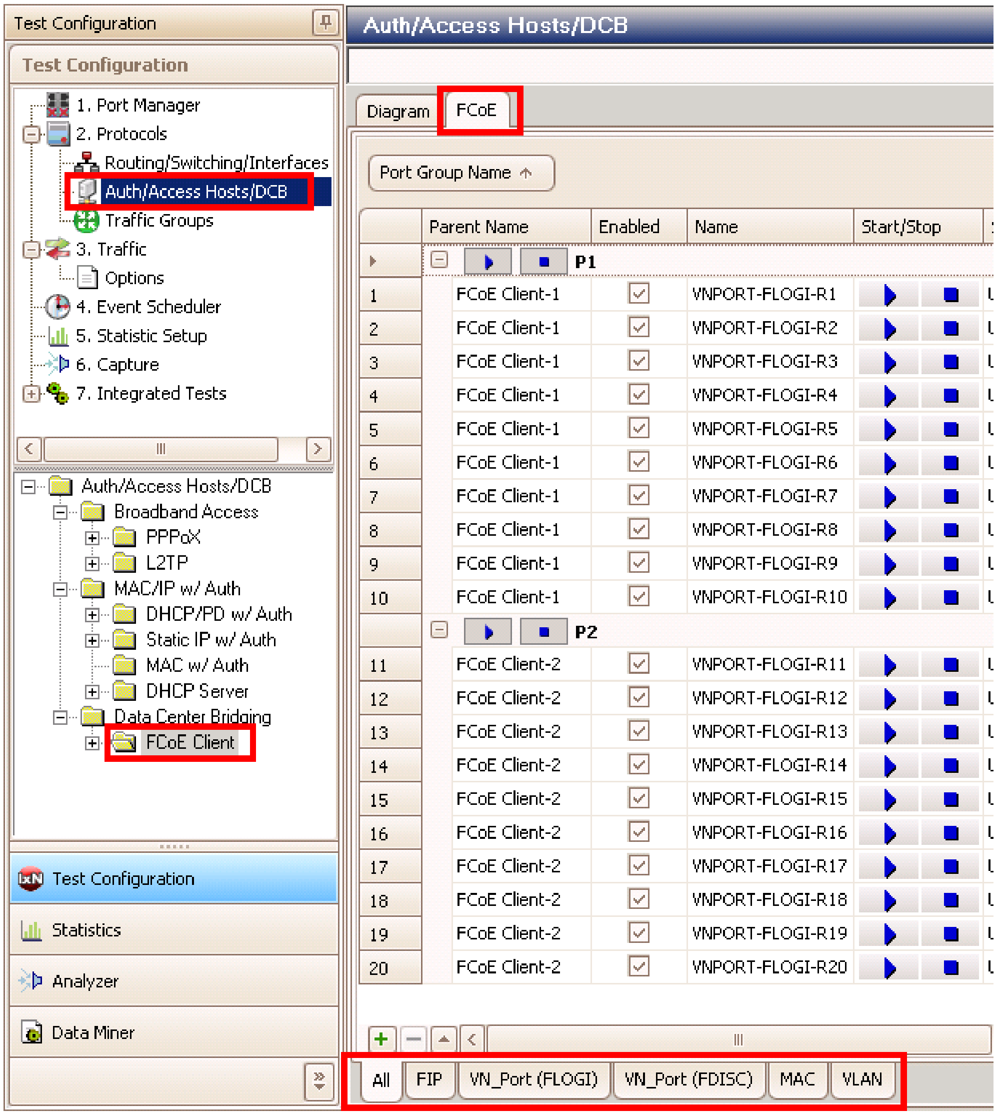

	Parent Name	Enabled	Name	Start/Stop
	P1			
1	FCoE Client-1	☑	VNPORT-FLOGI-R1	▶ ■
2	FCoE Client-1	☑	VNPORT-FLOGI-R2	▶ ■
3	FCoE Client-1	☑	VNPORT-FLOGI-R3	▶ ■
4	FCoE Client-1	☑	VNPORT-FLOGI-R4	▶ ■
5	FCoE Client-1	☑	VNPORT-FLOGI-R5	▶ ■
6	FCoE Client-1	☑	VNPORT-FLOGI-R6	▶ ■
7	FCoE Client-1	☑	VNPORT-FLOGI-R7	▶ ■
8	FCoE Client-1	☑	VNPORT-FLOGI-R8	▶ ■
9	FCoE Client-1	☑	VNPORT-FLOGI-R9	▶ ■
10	FCoE Client-1	☑	VNPORT-FLOGI-R10	▶ ■
	P2			
11	FCoE Client-2	☑	VNPORT-FLOGI-R11	▶ ■
12	FCoE Client-2	☑	VNPORT-FLOGI-R12	▶ ■
13	FCoE Client-2	☑	VNPORT-FLOGI-R13	▶ ■
14	FCoE Client-2	☑	VNPORT-FLOGI-R14	▶ ■
15	FCoE Client-2	☑	VNPORT-FLOGI-R15	▶ ■
16	FCoE Client-2	☑	VNPORT-FLOGI-R16	▶ ■
17	FCoE Client-2	☑	VNPORT-FLOGI-R17	▶ ■
18	FCoE Client-2	☑	VNPORT-FLOGI-R18	▶ ■
19	FCoE Client-2	☑	VNPORT-FLOGI-R19	▶ ■
20	FCoE Client-2	☑	VNPORT-FLOGI-R20	▶ ■

Figure 76. FIP VLAN discovery stress test – FCoE client configuration GUI view

14. Optionally, configure the login **Setup Rate** and logout **Teardown Rate**. These values determine the rate at which FLOGIs or FDISCs are transmitted to the FCF, and the rate at which LOGOs are transmitted to the FCF.

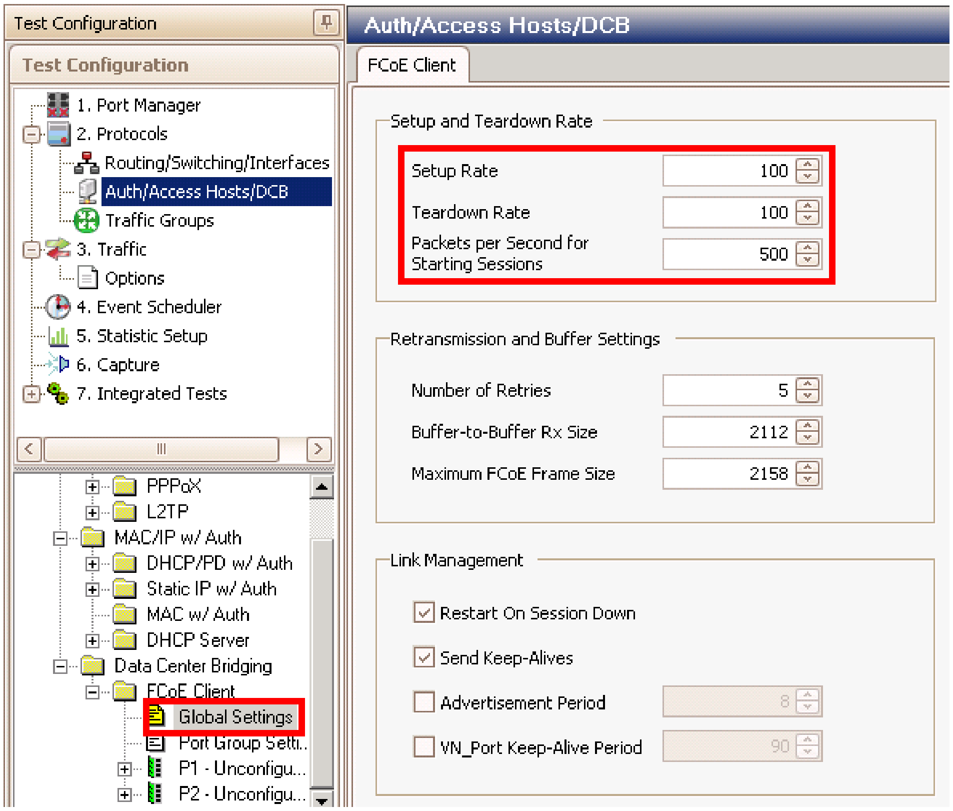

Setup and Teardown Rate

Setup Rate	100
Teardown Rate	100
Packets per Second for Starting Sessions	500

Retransmission and Buffer Settings

Number of Retries	5
Buffer-to-Buffer Rx Size	2112
Maximum FCoE Frame Size	2158

Link Management

☑ Restart On Session Down

☑ Send Keep-Alives

□ Advertisement Period 8

□ VN_Port Keep-Alive Period 90

Figure 77. FIP VLAN discovery Stress Test – Login Setup and Logout Teardown Rates

15. Start the emulation and verify test results.

16. Wait until all VN_Ports are instantiated. This process can take seconds or minutes depending on the configured login setup rate.

Test Variables

Performance Variable	Description
Number of ENodes per Port	Each ENode instantiates a set of FLOGI VN_Ports and FDISC VN_Ports. Increasing the **Number of ENodes per Port** will stress the DUT's (FCF) ability to process FCoE login requests.
Number of FLOGI VN_Ports per ENode	A VN_Port can be instantiated by FLOGI or FDISC. Increasing the **Number of FLOGI VN_Ports per ENode** will stress the DUT's (FCF) ability to process FLOGI requests.
Number of FDISC VN_Ports per FLOGI VN_Port	A VN_Port can be instantiated by FLOGI or FDISC. Increasing the **Number of FDISC VN_Ports per FLOGI VN_Port** will stress the DUT's (FCF) ability to process FDISC requests.
Name Server Registration	Each VN_Port registers its attributes with the name server (commonly the FCF) after it logs in to the fabric. In addition, each VN_Port queries the name server for each destination VN_Port. Enabling this option puts additional stress on the DUT to process name server registrations and queries.
Perform PLOGI	PLOGI is a connection request made from a source VN_Port to a destination VN_Port. Enabling this option puts additional stress on the DUT to process PLOGI request and replies.
Enable FIP	Enabling FIP will stress the DUT's ability to process discovery solicitations and periodic keepalives from each VN_Port.
Enable FIP VLAN Discovery	Enabling FIP VLAN Discovery will further stress the DUT's ability to process discovery solicitations and periodic keepalives from each VN_Port.
Setup/Teardown Rate	(Optional) The rate at which VN_Ports request to login to a fabric. Increasing this number will stress the DUT's ability to process a large volume of Fabric Login and logout requests.

Result Analysis

Use the IxNetwork **Statistics** view to verify the following FCoE and FIP packets statistics:

- **FIP Discovery Advertisements Rx**: *1,000*

  There are 1,000 FLOGI VN_Ports configured. Each of the 1,000 VN_Ports will transmit a **FIP Discovery Solicitation**. The DUT shall respond to each solicitation with a **FIP Discovery Advertisement**.

- **FLOGI LS_ACC Rx**: *1,000*

  There are 100 FLOGI VN_Ports configured for each of the 10 ENodes. Each VN_Port shall be instantiated by transmitting one FLOGI. The DUT shall accept each FLOGI by responding with an LS_ACC.

- **FDISC LS_ACC Rx**: *4,000*

  There are 4 FDISC VN_Ports configured for each of the 1,000 FLOGI VN_Ports. Each FDISC VN_Port shall be instantiated by transmitting one FDISC. The DUT shall accept each FDISC by responding with an LS_ACC.

- **PLOGI LS_ACC Rx**: *5,000*

  There are 5,000 VN_Ports in total. Each of the 1,000 FLOGI VN_Ports transmits one PLOGI to its target VN_Port. Each of the 4,000 FDISC VN_Ports transmits one PLOGI to its target VN_Port. The target side Ixia port responds to each PLOGI with an LS_ACC. The DUT shall forward each PLOGI LS_ACC to the initiator side Ixia port.

- **NS Registration Successful**: *5,000*

  Each of the 5,000 VN_Ports registers with the Name Server (that is, the DUT in this test). The DUT shall acknowledge each of the 5,000 NS Registrations.

- **Interfaces Up**: *5,000*

  All 5,000 VN_Ports shall be instantiated.

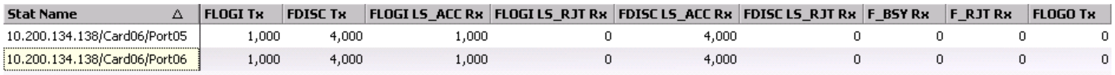

Stat Name △	FLOGI Tx	FDISC Tx	FLOGI LS_ACC Rx	FLOGI LS_RJT Rx	FDISC LS_ACC Rx	FDISC LS_RJT Rx	F_BSY Rx	F_RJT Rx	FLOGO Tx
10.200.134.138/Card06/Port05	1,000	4,000	1,000	0	4,000	0	0	0	0
10.200.134.138/Card06/Port06	1,000	4,000	1,000	0	4,000	0	0	0	0

Figure 78. FIP VLAN discovery stress test – FCoE client aggregate statistics view

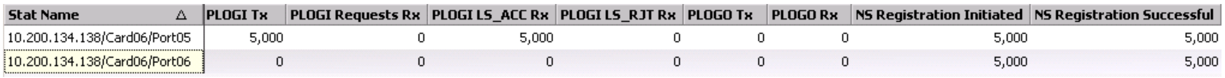

Stat Name △	PLOGI Tx	PLOGI Requests Rx	PLOGI LS_ACC Rx	PLOGI LS_RJT Rx	PLOGO Tx	PLOGO Rx	NS Registration Initiated	NS Registration Successful
10.200.134.138/Card06/Port05	5,000	0	5,000	0	0	0	5,000	5,000
10.200.134.138/Card06/Port06	0	0	0	0	0	0	5,000	5,000

Figure 79. FIP VLAN discovery stress test – FCoE client aggregate statistics view (cont'd)

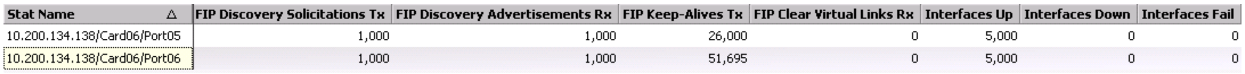

Stat Name △	FIP Discovery Solicitations Tx	FIP Discovery Advertisements Rx	FIP Keep-Alives Tx	FIP Clear Virtual Links Rx	Interfaces Up	Interfaces Down	Interfaces Fail
10.200.134.138/Card06/Port05	1,000	1,000	26,000	0	5,000	0	0
10.200.134.138/Card06/Port06	1,000	1,000	51,695	0	5,000	0	0

Figure 80. FIP VLAN discovery stress test – FCoE client aggregate statistics view (Cont'd)

Use the IxNetwork per session statistics view to verify the following FCoE control plane exchanges for each VN_Port. Right click from the aggregate statistics above and select **DrillDown per Session**.

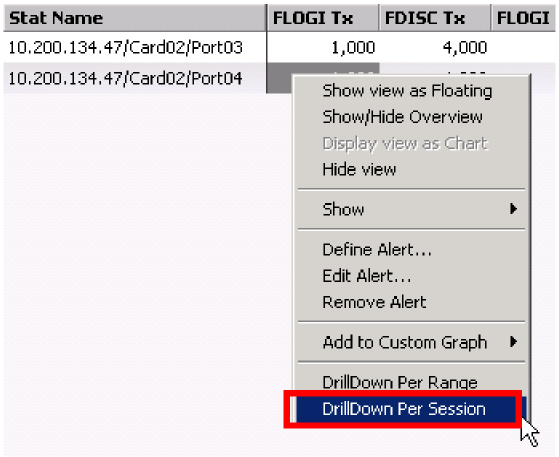

Stat Name	FLOGI Tx	FDISC Tx	FLOGI
10.200.134.47/Card02/Port03	1,000	4,000	
10.200.134.47/Card02/Port04			

Figure 81. FCoE fabric login stress test - DrillDown per Session

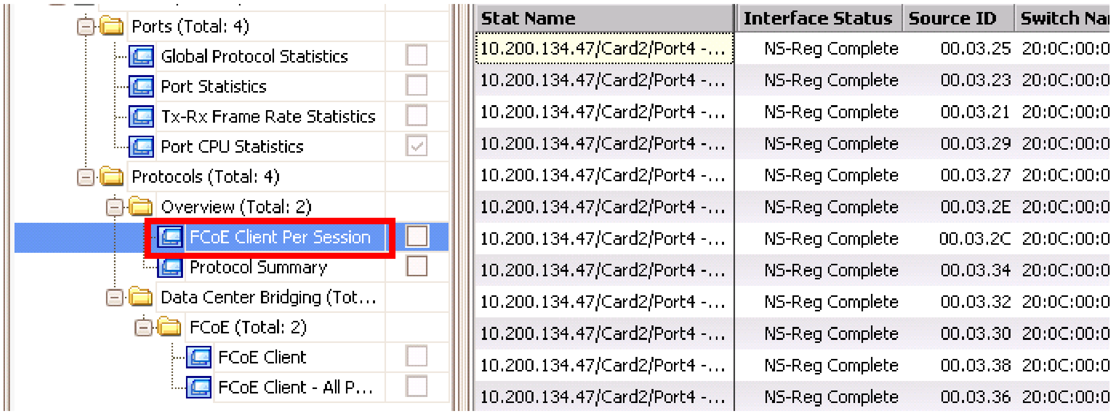

Stat Name	Interface Status	Source ID	Switch Na
10.200.134.47/Card2/Port4 -...	NS-Reg Complete	00.03.25	20:0C:00:0
10.200.134.47/Card2/Port4 -...	NS-Reg Complete	00.03.23	20:0C:00:0
10.200.134.47/Card2/Port4 -...	NS-Reg Complete	00.03.21	20:0C:00:0
10.200.134.47/Card2/Port4 -...	NS-Reg Complete	00.03.29	20:0C:00:0
10.200.134.47/Card2/Port4 -...	NS-Reg Complete	00.03.27	20:0C:00:0
10.200.134.47/Card2/Port4 -...	NS-Reg Complete	00.03.2E	20:0C:00:0
10.200.134.47/Card2/Port4 -...	NS-Reg Complete	00.03.2C	20:0C:00:0
10.200.134.47/Card2/Port4 -...	NS-Reg Complete	00.03.34	20:0C:00:0
10.200.134.47/Card2/Port4 -...	NS-Reg Complete	00.03.32	20:0C:00:0
10.200.134.47/Card2/Port4 -...	NS-Reg Complete	00.03.30	20:0C:00:0
10.200.134.47/Card2/Port4 -...	NS-Reg Complete	00.03.38	20:0C:00:0
10.200.134.47/Card2/Port4 -...	NS-Reg Complete	00.03.36	20:0C:00:0

Figure 82. FIP fabric login stress test – FCoE client per VN_Port statistics view

- All 10,000 VN_Ports have a unique **Assigned MAC** address assigned by the DUT.
- All 5,000 VN_Ports from the initiator side are associated with the same **Fabric MAC** address, but different than the second 5,000 VN_Ports.
- All 5,000 VN_Ports from the target sideare associated with the same **Fabric MAC** address, but different than the 5,000 VN_Ports from the initiator side.
- All 10,000 VN_Ports are assigned the same **FC Map**.
- All 10,000 VN_Ports are assigned the **Priority** configured on the DUT. Unless configured otherwise, the **Priority** value should be the same for all VN_Ports.

- All 10,000 VN_Ports are associated with the **Fabric Name** and **Switch Name** configured on the DUT. Unless configured otherwise, the **Fabric Name** should be the same for all VN_Ports, and the **Switch Name** should be the same for all VN_Ports.
- All 5,000 VN_Ports from the initiator side have **Status** set to **PLOGI OK**. This is the last initialization operation performed by the initiator side VN_Ports.
- All 5,000 VN_Ports from the target side have **Status** set to **NS-Reg OK**. This is the last initialization operation performed by the target side VN_Ports.
- Each 10,000 VN_Port is assigned a unique **Source ID**.
- All 5,000 VN_Ports from the initiator side are assigned a unique **PLOGI Destination ID**. This is the D_ID that the initiator VN_Port uses to perform PLOGI to its target VN_Port.
- All 2,000 FLOGI VN_Ports are assigned one or more VLAN IDs in the **Discovered VLAN IDs** column.
- All 8,000 FDISC VN_Ports are assigned the same set of VLAN IDs as their parent FLOGI VN_Ports.

*Note: Use the **Filter/Sort** function to zoom in on VN_Ports that matches user defined statistics criteria*

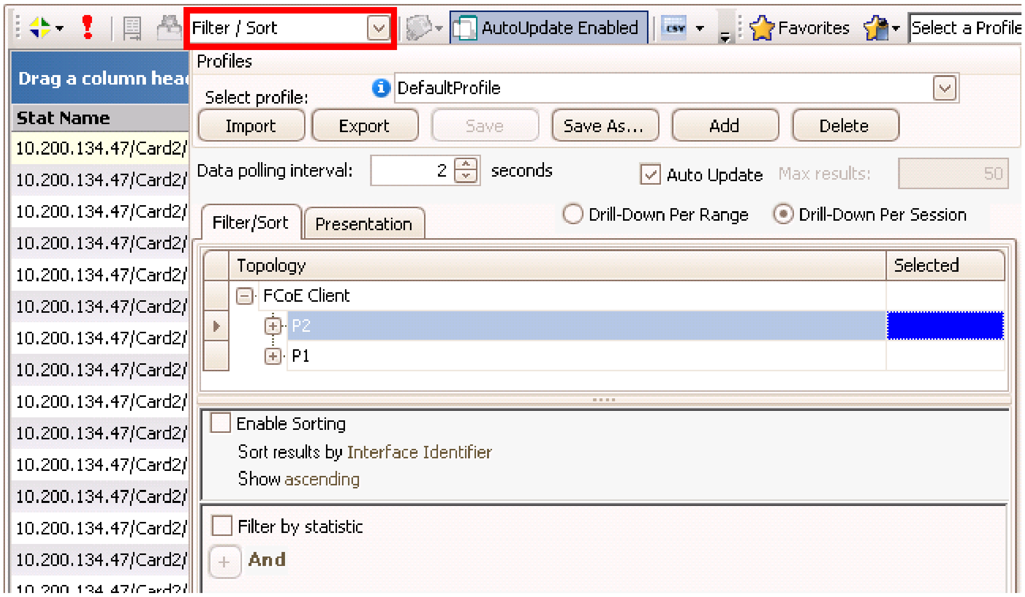

Stat Name
10.200.134.47/Card2/
10.200.134.47/Card2/
10.200.134.47/Card2/
10.200.134.47/Card2/
10.200.134.47/Card2/
10.200.134.47/Card2/
10.200.134.47/Card2/
10.200.134.47/Card2/
10.200.134.47/Card2/
10.200.134.47/Card2/
10.200.134.47/Card2/
10.200.134.47/Card2/
10.200.134.47/Card2/
10.200.134.47/Card2/

Select profile: DefaultProfile

Data polling interval: 2 seconds ☑ Auto Update Max results: 50

○ Drill-Down Per Range ● Drill-Down Per Session

□ Enable Sorting
Sort results by Interface Identifier
Show ascending

□ Filter by statistic

Figure 83. FCoE fabric login stress test - statistics Filter/Sort function

Conclusions

By validating the statistics and control plane exchanges using the features above, the DUT has been proven capable of responding to at least 1,000 FLOGI VN_Ports' VLAN Discovery requests per port, and is assigning VLAN IDs properly, in addition to FPMA and FC MAP, to at least 5,000 VN_Ports per port, for the specified topology. In addition, the DUT is capable of instantiating at least 5,000 VN_Ports per port, instantiated by both FIP FLOGI and FIP FDISC, for the specified topology. Furthermore, the DUT is capable of maintaining at least 5,000 VN_Ports per port with periodic FIP keepalives for the specified topology.

Test Case: Converged Traffic Forwarding – FCoE and LAN Traffic (with QoS)

Overview

A converged data center network will transport both FCoE and native Ethernet traffic across its diverse infrastructure. While the forwarding of native Ethernet traffic is a well established function within today's networks with the vast availability of purpose-built FPGAs, chipsets and systems, introducing FCoE traffic into the same network raises several challenges where many prototypes exist.

First, an FCoE data packet can be up to 2176 bytes without additional encapsulation, such as VLAN. Traditional Ethernet chipsets, per IEEE Std. 802.3-2005, by default do not recognize Ethernet frames with data payload greater than 1500 bytes. Therefore, an inherent risk exists that will prevent an FCoE data packet from crossing the converged data center infrastructure.

Second, each component, both internal and external, in the end-to-end data path of FCoE and/or native Ethernet traffic should be able to differentiate between these two packet types. There are many reasons why a network operator would want to distinguish FCoE packets from native Ethernet frames – for QoS policies, for example. One of the most important reasons to do so is that FCoE packets have different delivery ordering requirements than Ethernet frames. Native Ethernet forwarding requirement, with respect to mis-ordering, is first-in, first-out. The concept of packet ordering in fibre channel is based on SEQ_CNT, SEQ_ID and OX_ID, between a pair of VN_Ports identified by S_ID and D_ID.

Third, storage traffic (for example, SCSI over FCoE) is generally more critical and important than LAN Ethernet traffic. Therefore, in addition to being able to identify FCoE traffic from native Ethernet traffic, a critical capability is to be able to restrain native Ethernet traffic during congestion to ensure FCoE traffic delivery. A set of lossless Ethernet protocols are defined to achieve this requirement. One example is priority-based flow control (PFC).

Fibre channel facilitates some level of QoS differentiation when forwarding FC traffic. Preferential treatment is signified within the FC header portion of the entire packet, via the CS_CTL/Priority field. In Class 3 devices, the CS_CTL/Priority field offers two differentiation options – PREF and DSCP. PREF provides a simple preferred vs. non-preferred differentiation, while DSCP provides a multi-tier quality of service classification.

This test focuses on forwarding FCoE and native Ethernet traffic from initiators to targets, with CS_CTL/Priority enabled. Priority-based flow control is disabled in this test.

Objective

The objective of this test is to verify the ability of an FCF (that is, the DUT) to forward both FCoE and native Ethernet traffic from emulated initiators to emulated targets at line rate. Furthermore, FCoE traffic will be encoded with different QoS marking using the CS_CTL/Priority field. FCoE traffic will be transmitted from the initiator side VN_Ports to the target side VN_Ports, using the assigned fabric provided MAC addresses (FPMAs). Native Ethernet traffic will be transmitted from the initiator side ENodes to the target side ENodes, using the ENode MAC addresses.

Setup

Two Ixia ports are used in this test, one on the initiator side, and one on the target side.

The objective of this test is to instantiate 10,000 VN_Ports and forward FCoE and native Ethernet traffic from the initiator side VN_Ports and ENodes, to the target side VN_Ports and ENodes.

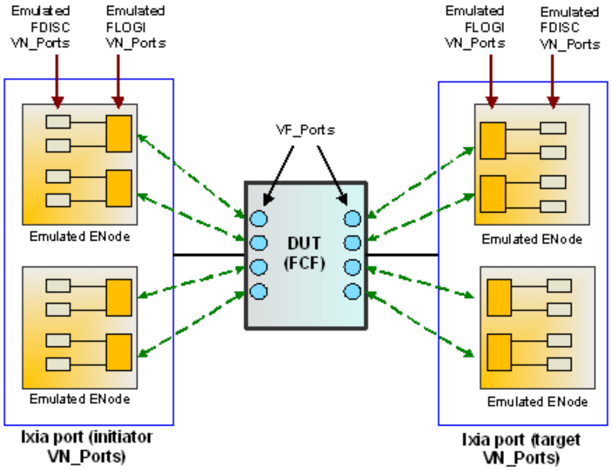

Figure 84. FCoE and native Ethernet traffic forwarding test topology

Step-by-Step Instructions

1. Complete the configurations required in the FIP Fabric Login Stress Test, described earlier in this booklet.

2. Start the **Advanced Wizard** to configure FCoE traffic.

   a. On the **Endpoints** page, set the **Type of Traffic** to **FCoE**, and select the **FCoE Clients** from **P1** on the **Source** side, and **FCoE Clients** from **P2** on the **Destination** side. These are the FCoE VN_Port endpoints. Click **Apply** when finished. Optionally, set a name in **Traffic Name**.

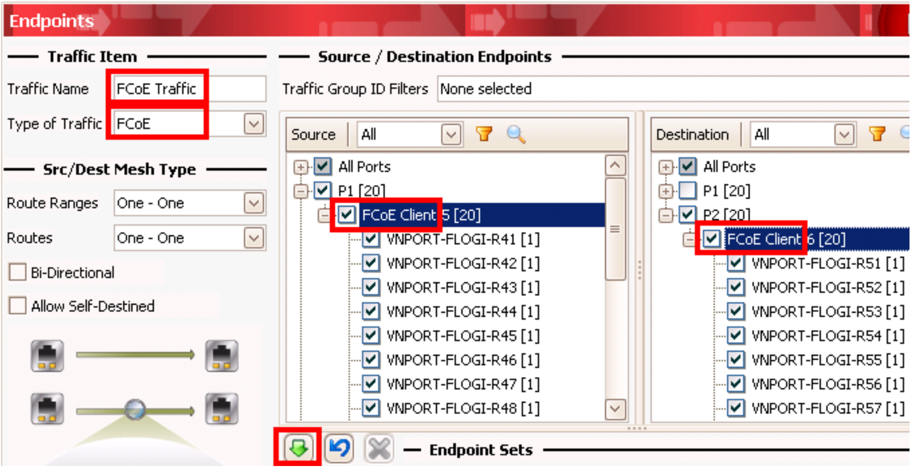

Endpoints

Traffic Item

Traffic Name: FCoE Traffic

Type of Traffic: FCoE

Src/Dest Mesh Type

Route Ranges: One - One

Routes: One - One

☐ Bi-Directional

☐ Allow Self-Destined

Source / Destination Endpoints

Traffic Group ID Filters: None selected

Source: All

- ☑ All Ports
- ☑ P1 [20]
  - ☑ FCoE Client 5 [20]
    - ☑ VNPORT-FLOGI-R41 [1]
    - ☑ VNPORT-FLOGI-R42 [1]
    - ☑ VNPORT-FLOGI-R43 [1]
    - ☑ VNPORT-FLOGI-R44 [1]
    - ☑ VNPORT-FLOGI-R45 [1]
    - ☑ VNPORT-FLOGI-R46 [1]
    - ☑ VNPORT-FLOGI-R47 [1]
    - ☑ VNPORT-FLOGI-R48 [1]

Destination: All

- ☑ All Ports
- ☐ P1 [20]
- ☑ P2 [20]
  - ☑ FCoE Client 6 [20]
    - ☑ VNPORT-FLOGI-R51 [1]
    - ☑ VNPORT-FLOGI-R52 [1]
    - ☑ VNPORT-FLOGI-R53 [1]
    - ☑ VNPORT-FLOGI-R54 [1]
    - ☑ VNPORT-FLOGI-R55 [1]
    - ☑ VNPORT-FLOGI-R56 [1]
    - ☑ VNPORT-FLOGI-R57 [1]

Endpoint Sets

Figure 85. Advanced Traffic Wizard Endpoints page – select FCoE endpoints

   b. On the **Packet/QoS** page, assign preference to the FCoE traffic using the **CS_CTL/Priority** column drop down menu, and enter the following values:

      i. CS_CTL/Priority *0x80*: PREF = *1*, DSCP = *0*

      ii. CS_CTL/Priority *0x00*: PREF = *0*, DSCP = *0*

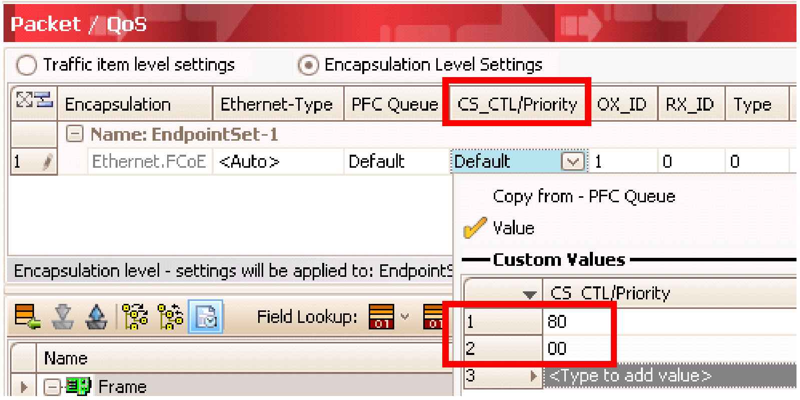

Packet / QoS

○ Traffic item level settings ● Encapsulation Level Settings

	Encapsulation	Ethernet-Type	PFC Queue	CS_CTL/Priority	OX_ID	RX_ID	Type
	Name: EndpointSet-1						
1	Ethernet.FCoE	<Auto>	Default	Default	1	0	0

Copy from - PFC Queue

✓ Value

Custom Values

	CS_CTL/Priority
1	80
2	00
3	<Type to add value>

Encapsulation level - settings will be applied to: Endpoints

Field Lookup:

Name

Frame

Figure 86. Advanced Traffic Wizard Packet/QoS page - assign PREF

By default, the CS_CTL/Priority values are cycled through the source VN_Ports. If fully mesh testing is desired, assign values so that each VN_Port transmits traffic with both CS_CTL/Priority values, select the **Fully Mesh This Field** option for the CS_CTL/Priority field in the **Packet Editor**.

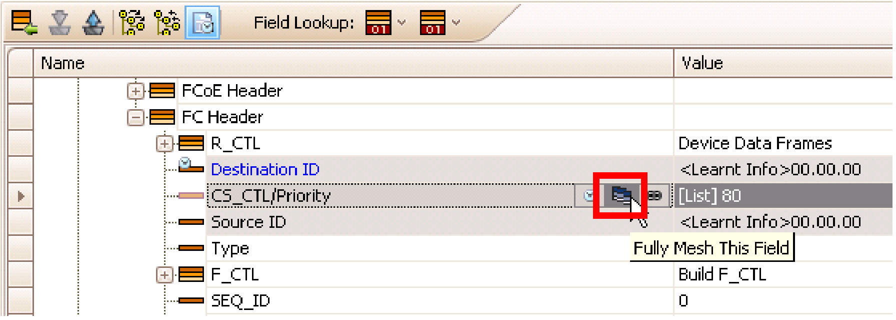

Field Lookup:

Name	Value
FCoE Header	
FC Header	
R_CTL	Device Data Frames
Destination ID	<Learnt Info>00.00.00
CS_CTL/Priority	[List] 80
Source ID	<Learnt Info>00.00.00
Type	
F_CTL	Build F_CTL
SEQ_ID	0

Figure 87. Advanced Traffic Wizard Packet/QoS page - fully mesh CS_CTL/Priority

c. On the **Flow Group Setup** page, select the check box **FCoE: CS_CTL/Priority** to create flow groups based on CS_CTL/Priority values. All packets within the same flow group can be independently rate and size varied dynamically. Using the configuration above, two flow groups will be created, one for CS_CTL/Priority 0x80 (that is, PREF = 1) and one for CS_CTL/Priority 0x00 (that is, PREF = 0).

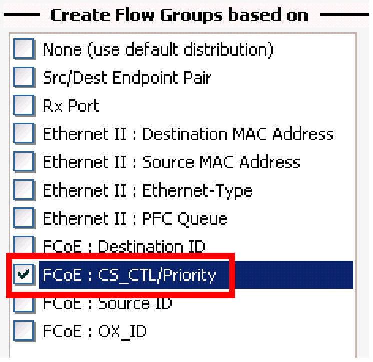

Create Flow Groups based on

- □ None (use default distribution)
- □ Src/Dest Endpoint Pair
- □ Rx Port
- □ Ethernet II : Destination MAC Address
- □ Ethernet II : Source MAC Address
- □ Ethernet II : Ethernet-Type
- □ Ethernet II : PFC Queue
- □ FCoE : Destination ID
- ☑ FCoE : CS_CTL/Priority
- □ FCoE : Source ID
- □ FCoE : OX_ID

Figure 88. Advanced Traffic Wizard Flow Group Setup page - assign flow groups

d. On the **Rate Setup** page, set the **Line Rate** to *50* %.

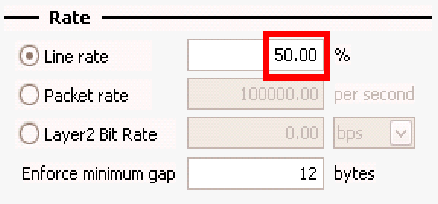

Rate

- ● Line rate 50.00 %
- ○ Packet rate 100000.00 per second
- ○ Layer2 Bit Rate 0.00 bps
- Enforce minimum gap 12 bytes

Figure 89. Advanced Traffic Wizard Rate Setup page – configure line rate

e. On the **Flow Tracking** page, select the following **Track Flows By** options: **Traffic Item**, **FCoE: Destination ID** and **FCoE: CS_CTL/Priority**. Selecting a tracking option allows traffic statistics to be drilled down to the selected option.

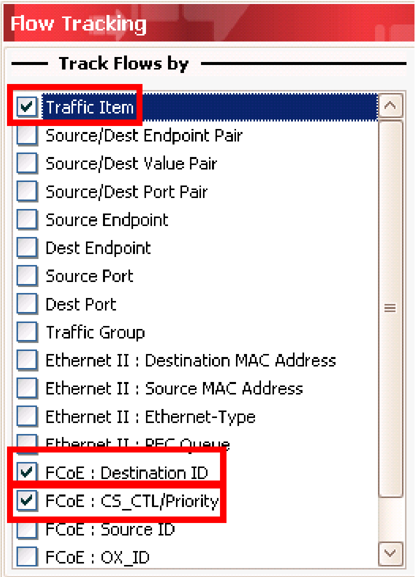

Flow Tracking

Track Flows by

☑ Traffic Item
□ Source/Dest Endpoint Pair
□ Source/Dest Value Pair
□ Source/Dest Port Pair
□ Source Endpoint
□ Dest Endpoint
□ Source Port
□ Dest Port
□ Traffic Group
□ Ethernet II : Destination MAC Address
□ Ethernet II : Source MAC Address
□ Ethernet II : Ethernet-Type
□ Ethernet II : PFC Queue
☑ FCoE : Destination ID
☑ FCoE : CS_CTL/Priority
□ FCoE : Source ID
□ FCoE : OX_ID

Figure 90. Advanced Traffic Wizard Flow Tracking page – set flow tracking options

f. (Optional) On the **Preview** page, click **View Flow Groups/Packets** to preview the content of the packets that will be transmitted to the DUT. Click **Flow Group** to view the content of each packet. There should be 5,000 packets created, with alternating CS_CTL/Priority value. If the **Fully Mesh This Field** option was selected above, then there should be 10,000 packets created, each of the 5,000 source VN_Ports should have both CS_CTL/Priority values configured.

g. (Optional) On the **Validate** page, click **Validate** to verify hardware resource is available before exiting the traffic wizard.

h. Click **Finish** to exit the traffic wizard.

3. Start the **Advanced Wizard** to configure native Ethernet traffic.

a. On the **Endpoints** page, set the **Type of Traffic** to **Ethernet/VLAN**, and select the **MAC/VLAN** from **P1** on the **Source** side, and **MAC/VLAN** from **P2** on the **Destination** side. These are the ENode Lossless Ethernet MAC endpoints (each set of FLOGI VN_Port has a unique Lossless Ethernet MAC on its parent ENode). Click **Apply** when finished. Optionally, set a name in **Traffic Name**.

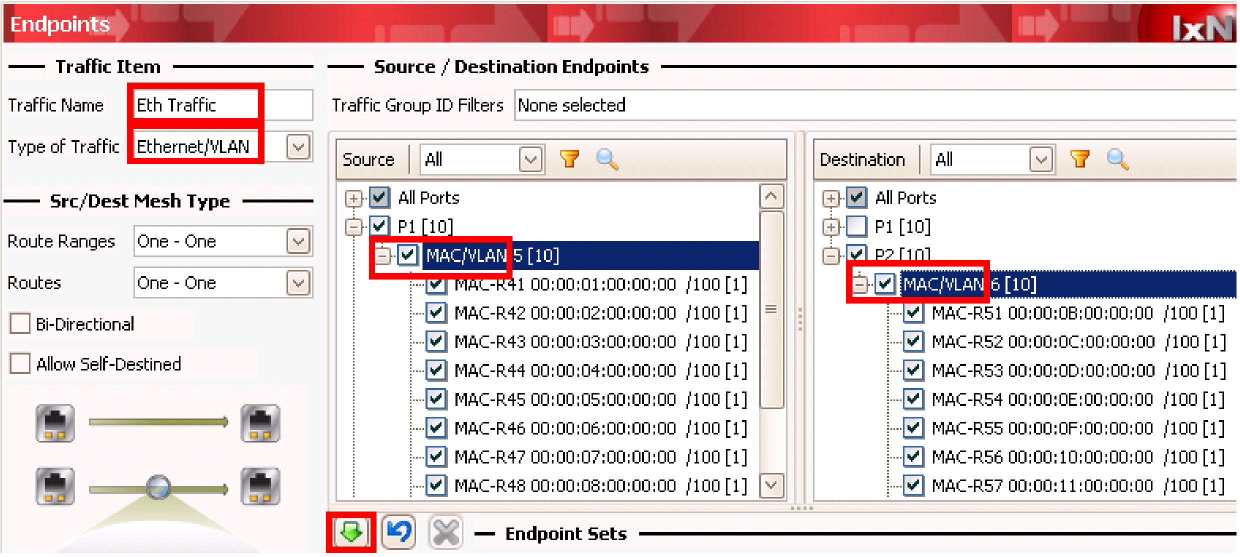

Endpoints

Traffic Item

Traffic Name: Eth Traffic

Type of Traffic: Ethernet/VLAN

Src/Dest Mesh Type

Route Ranges: One - One

Routes: One - One

□ Bi-Directional

□ Allow Self-Destined

Source / Destination Endpoints

Traffic Group ID Filters: None selected

Source: All

- ☑ All Ports
- ☑ P1 [10]
  - ☑ MAC/VLAN 5 [10]
    - ☑ MAC-R41 00:00:01:00:00:00 /100 [1]
    - ☑ MAC-R42 00:00:02:00:00:00 /100 [1]
    - ☑ MAC-R43 00:00:03:00:00:00 /100 [1]
    - ☑ MAC-R44 00:00:04:00:00:00 /100 [1]
    - ☑ MAC-R45 00:00:05:00:00:00 /100 [1]
    - ☑ MAC-R46 00:00:06:00:00:00 /100 [1]
    - ☑ MAC-R47 00:00:07:00:00:00 /100 [1]
    - ☑ MAC-R48 00:00:08:00:00:00 /100 [1]

Destination: All

- ☑ All Ports
- □ P1 [10]
- ☑ P2 [10]
  - ☑ MAC/VLAN 6 [10]
    - ☑ MAC-R51 00:00:0B:00:00:00 /100 [1]
    - ☑ MAC-R52 00:00:0C:00:00:00 /100 [1]
    - ☑ MAC-R53 00:00:0D:00:00:00 /100 [1]
    - ☑ MAC-R54 00:00:0E:00:00:00 /100 [1]
    - ☑ MAC-R55 00:00:0F:00:00:00 /100 [1]
    - ☑ MAC-R56 00:00:10:00:00:00 /100 [1]
    - ☑ MAC-R57 00:00:11:00:00:00 /100 [1]

Endpoint Sets

Figure 91. Advanced Traffic Wizard Endpoints page – select ENode endpoints

b. On the **Rate Setup** page, set the **Line Rate** to *50* %.

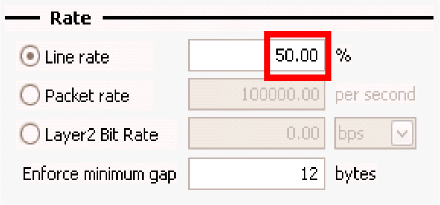

Rate

● Line rate 50.00 %

○ Packet rate 100000.00 per second

○ Layer2 Bit Rate 0.00 bps

Enforce minimum gap 12 bytes

Figure 92. Advanced Traffic Wizard Rate Setup page – configure line rate

c. On the **Flow Tracking** page, select the following **Track Flows By** options: **Traffic Item** and **Ethernet II: Destination MAC Address**. Selecting a tracking option allows traffic statistics to be drilled down to the selected option.

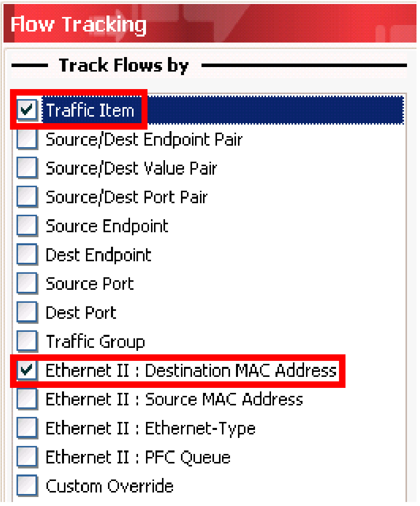

Flow Tracking

Track Flows by

- ☑ Traffic Item
- □ Source/Dest Endpoint Pair
- □ Source/Dest Value Pair
- □ Source/Dest Port Pair
- □ Source Endpoint
- □ Dest Endpoint
- □ Source Port
- □ Dest Port
- □ Traffic Group
- ☑ Ethernet II : Destination MAC Address
- □ Ethernet II : Source MAC Address
- □ Ethernet II : Ethernet-Type
- □ Ethernet II : PFC Queue
- □ Custom Override

Figure 93. Advanced Traffic Wizard Flow Tracking page – set flow tracking options

d. (Optional) On the **Preview** page, click **View Flow Groups/Packets** to preview the content of the packets that will be transmitted to the DUT. Click **Flow Group** to view the content of each packet. There should be 1,000 packets created.

e. (Optional) On the **Validate** page, click **Validate** to verify hardware resource is available before exiting the traffic wizard.

f. Click **Finish** to exit the traffic wizard.

4. Click **Apply L2-L3 Traffic** to write the traffic onto the Ixia ports. Click **Start L2-L3 Traffic** to transmit the FCoE and native Ethernet traffic.

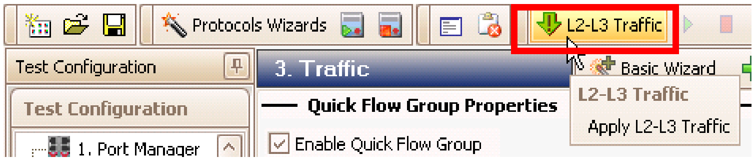

Figure 94. Traffic - apply traffic items

Test Variables

Performance Variable	Description
FIP VLAN Discovery	Enable **FIP VLAN Discovery** in the protocol wizard to exercise the DUT's traffic forwarding performance with VLANs. The IxNetwork traffic wizard will automatically insert the assigned VLAN ID (from the DUT) into packets from each lossless Ethernet MAC (that is, ENode MAC).
CS_CTL/Priority	This field conveys either CS_CTL or Priority, depending on the value of F_CTL. Class 3 devices interpret CS_CTL as PREF and DSCP settings. Assign additional DSCP values to exercise the DUT's policing algorithm.
OX_ID	Exchange identifier. Assign OX_ID values, and then track by option FCoE: OX_ID, to verify all exchanges are properly forwarded properly.
Frame size	Size of packets (in bytes) transmitted from the test port for the current traffic item or endpoint set (per user setting). The default is **fixed**, with values of 2160 bytes for FCoE traffic, and 64 bytes for native Ethernet traffic. Other options are **increment**, **random**, **IMIX** (9 defined profiles and 1 custom profile) and **quad Gaussian**.
Payload	The content (in hex) encoded in the FC data payload field of the FCoE packet. Use this to create specific payload patterns (for example, 'killer packets').
CRC Settings	The CRC setting for the Ethernet FCS field. The default is a **good** CRC value. Options are **bad** or **none**. Note: A **bad FC CRC** can be inserted using IxExplorer.
Rate	The transmission rate of the current traffic item or endpoint set. The default is **line rate**. The line rate, by default, is split evenly across all flow groups created by the traffic item or endpoint set. Options are **packet rate** and **layer2 bit rate**.
Packet transmission mode	The packet transmission pattern for the current traffic item or endpoint set. The default is **continuous**. Options are **fixed packet count**, **fixed iteration count** and **burst**.

Result Analysis

Use the IxNetwork **Statistics – Traffic Item Statistics** view to verify the following real-time FCoE traffic aggregate statistics and native Ethernet traffic aggregate statistics.

- **Rx Frames**: *equal to **Tx Frames***
- **Frames Delta**: *0*
- **Rx Frame Rate**: *equal to **Tx Frame Rate***

  Note: at very high line rates, it is normal to see a small mismatch between Tx and Rx stats, and therefore frames delta. These are not necessarily actual frame loss. Rather, it represents the packets that are still in transit across the system under test at the instance the statistics are reported.

- **Cut-Through Avg Latency**: *vendor specific*
  Verify the average latency is within the expected range. Additional latency and jitter measurement options are available under the **Traffic – Options** settings.

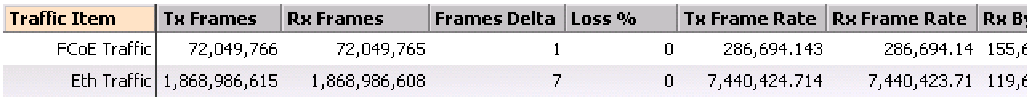

Traffic Item	Tx Frames	Rx Frames	Frames Delta	Loss %	Tx Frame Rate	Rx Frame Rate	Rx By
FCoE Traffic	72,049,766	72,049,765	1	0	286,694.143	286,694.14	155,6
Eth Traffic	1,868,986,615	1,868,986,608	7	0	7,440,424.714	7,440,423.71	119,6

Figure 95. FCoE and native Ethernet traffic forwarding – traffic item statistics

From the **Traffic Item Statistics** view, use the IxNetwork **DrillDown per FCoE: Destination ID** view to verify traffic is received properly by each target side D_ID.

- **Number of unique Destination IDs**: *5,000*
- **FCoE: Destination ID**: *equal to the target side VN_Ports' assigned Source ID identified in the FCoE/FIP per-VN_Port stats **Source ID** column*

From the **Traffic Item Statistics** view, use the IxNetwork **DrillDown per FCoE: CS_CTL/Priority** view to verify aggregate traffic statistics based on the CS_CTL/Priority values. Alternatively, from the **DrillDown per FCoE: Destination ID** view, use the IxNetwork **DrillDown per FCoE: CS_CTL/Priority** view to verify per-session traffic statistics for CS_CTL/Priority between each initiator side VN_Port and target side VN_Port pair.

- **Number of unique CS_CTL/Priority** values**:** *2*

FCoE: CS_CTL/Priority: *equal to the CS_CTL/Priority values configured in the **Advanced Wizard – Packet/QoS** page*

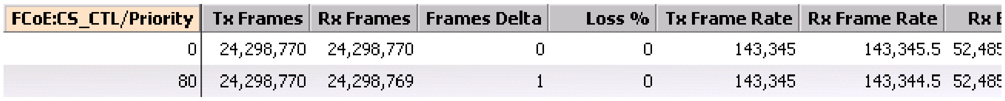

FCoE:CS_CTL/Priority	Tx Frames	Rx Frames	Frames Delta	Loss %	Tx Frame Rate	Rx Frame Rate	Rx B
0	24,298,770	24,298,770	0	0	143,345	143,345.5	52,485
80	24,298,770	24,298,769	1	0	143,345	143,344.5	52,485

Figure 96. FCoE and native Ethernet traffic forwarding – drilled down statistics per CS_CTL/Priority

Conclusions

By validating the traffic statistics using the features above, the DUT has been proven capable of forwarding FCoE traffic from 5,000 initiator side VN_Ports to 5,000 target side VN_Ports at 50% line rate, and forwarding native Ethernet traffic from 1,000 initiator side lossless Ethernet MACs (from 10 ENodes) to 1,000 target side lossless Ethernet MACs (from 10 ENodes) at 50% line rate.

Other conclusions can be drawn using different **Track Flows By** options or additional **Track Flows By** options, depending on the available tracking resources available. For example, if destination MAC address is also selected as a tracking option in the FCoE traffic item, then the mapping of target side VN_Port FPMA to each packet's Destination ID by the DUT can also be verified using the IxNetwork traffic statistics drill down feature **DrillDown per Ethernet II: Destination MAC Address**. Similarly, if the FCoE CS_CTL/Priority or Type fields are used to signify traffic with different priority values or application data type, then it is also possible to track and zoom in on these fields using the IxNetwork drill down feature.

Test Case: Converged Traffic Forwarding – FCoE and LAN Traffic (with PFC)

Overview

A converged data center network will transport both FCoE and native Ethernet traffic across its diverse infrastructure. While the forwarding of native Ethernet traffic is a well established function within today's networks with the vast availability of purpose-built FPGAs, chipsets and systems, introducing FCoE traffic into the same network raises several challenges where many prototypes exist.

First, an FCoE data packet can be up to 2176 bytes without additional encapsulation, such as VLAN. Traditional Ethernet chipsets, per IEEE Std. 802.3-2005, by default do not recognize Ethernet frames with data payload greater than 1500 bytes. Therefore, an inherent risk exists that will prevent an FCoE data packet from crossing the converged data center infrastructure.

Second, each component, both internal and external, in the end-to-end data path of FCoE and/or native Ethernet traffic should be able to differentiate between these two packet types. There are many reasons why a network operator would want to distinguish FCoE packets from native Ethernet frames – for QoS policies, for example. One of the most important reasons to do so is that FCoE packets have different delivery ordering requirements than Ethernet frames. Native Ethernet forwarding requirement, with respect to mis-ordering, is first-in, first-out. The concept of packet ordering in fibre channel is based on SEQ_CNT, SEQ_ID and OX_ID, between a pair of VN_Ports identified by S_ID and D_ID.

Third, storage traffic (for example, SCSI over FCoE) is generally more critical and important than LAN Ethernet traffic. Therefore, in addition to being able to identify FCoE traffic from native Ethernet traffic, a critical capability is to be able to pause FCoE traffic during congestion to ensure FCoE traffic is not dropped. Priority-based flow control (PFC) is one of the lossless Ethernet protocols defined to achieve this requirement.

Priority-based flow control is one of the new IEEE 802.1 protocols designed to enable lossless Ethernet. Priority-based flow control is built on the concept of the original IEEE 802.3x Flow Control, modified to operate with traffic class differentiation. When a receiving station's buffer is near exhaustion, 802.3x Flow Control allows the receiving station to request its upstream neighbor to pause transmission on the entire port, giving the receiving station an opportunity to clear its buffer. Priority-based flow control elevates this critical capability to a higher level by providing the receiving station the ability to request its upstream neighbor to pause transmission on one or more priorities (essentially virtual lanes). To achieve this, the original 802.3x flow control PAUSE frame is modified so that pause_quanta can be signaled at a per priority value, and the MAC CONTROL sublayer is enhanced with the capability to assign and throttle transmission queues on a per priority value.

This test focuses on forwarding FCoE and native Ethernet traffic from initiators to targets, with priority-based flow control enabled.

Objective

The objective of this test is to verify the ability of an FCF (that is, the DUT) to forward both FCoE and native Ethernet traffic from emulated initiators to emulated targets at line rate when its egress port is under subscribed. FCoE traffic will be assigned different PFC Queues than the native Ethernet traffic. During the test, FCoE traffic rate will be increased dynamically to oversubscribe the DUT’s egress port dynamically. The DUT is expected to issue a PFC PAUSE for the FCoE traffic’s PFC Queue. Zero packet loss is expected.

Setup

Three Ixia ports are used in this test, two on the initiator side, and one on the target side.

The objective of this test is to instantiate three VN_Ports and forward FCoE and native Ethernet traffic from the two initiator side VN_Ports and ENodes, to the target side VN_Port and ENode.

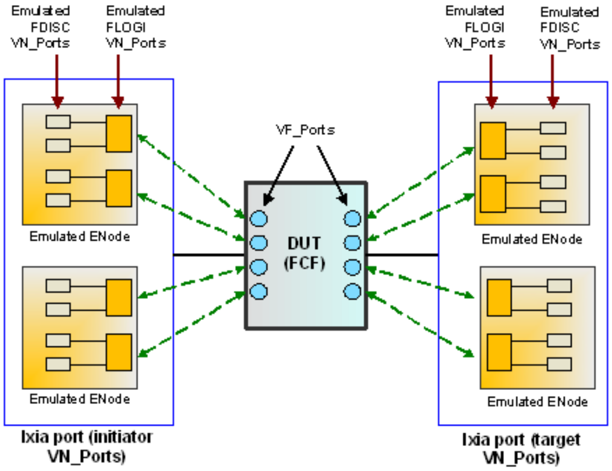

Figure 97. FCoE and native Ethernet traffic forwarding test topology

Step-by-Step Instructions

1. Complete the configurations required in the FIP Fabric Login Stress Test, described earlier in this booklet, using the following exceptions:

2. Reserve three Ixia ports

3. Assign Ixia Port 1 and Port 2 to be the **Initiator** side port, and assign Ixia Port 3 to be the **Target** side port.

   a. **FCoE/FC Node Wizard – Initiator side VN_Ports**:

      i. **Number of ENodes per Port**: *1*

      ii. **Number of FLOGI VN_Ports per ENode**: *1*

      iii. **Number of FDISC VN_Ports per FLOGI VN_Port:** *0*

   b. **FCoE/FC Node Wizard – Target side VN_Ports**:

      i. **Number of ENodes per Port**: *1*

      ii. **Number of FLOGI VN_Ports per ENode**: *1*

      iii. **Number of FDISC VN_Ports per FLOGI VN_Port:** *0*

4. By default, the PLOGI connections are mapped one-to-one from the Initiator side to the Target side VN_Ports. To map the PLOGI from Ixia Port 2's VN_Port to Ixia Port 3's VN_Port, go to the **Auth/Access Hosts/DCB – Data Center Bridging – FCoE Client** folder under **Test Configuration**, select the **VN_Port (FLOGI)** tab, and change Ixia Port 2's VN_Port's **PLOGI Target** to Ixia Port 3's VN_Port. (Such that the PLOGI Target for both Ixia Port 1 and Port 2's VN_Port is the same.)

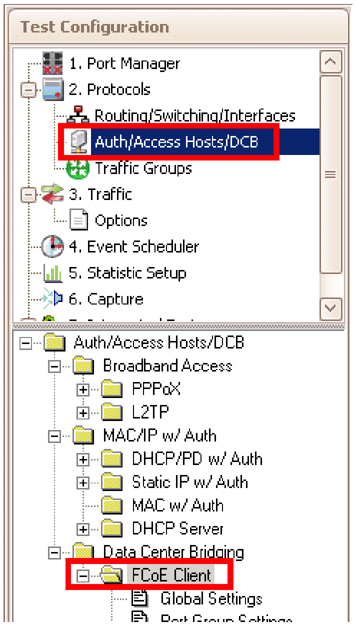

Figure 98. FCoE/FIP GUI - Test Configuration menu

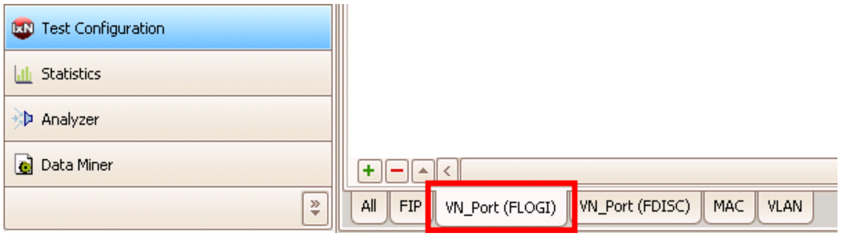

Figure 99. FCoE/FIP GUI - VN_Port (FLOGI) tab

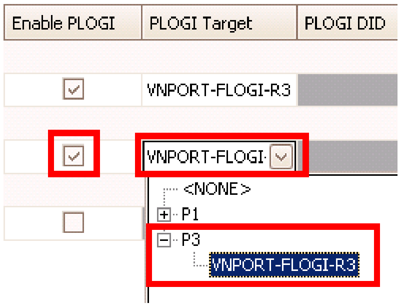

Enable PLOGI	PLOGI Target	PLOGI DID
☑	VNPORT-FLOGI-R3	
☑	VNPORT-FLOGI·	
□		

Figure 100. FCoE/FIP GUI - PLOGI Target setting

5. Start the emulation and wait until all VN_Ports are instantiated.

6. Use the IxNetwork FCoE aggregated and per-VN_Port statistics to verify all VN_Ports are properly instantiated.

7. Launch the **Advanced Wizard** to configure FCoE traffic.

   a. On the **Endpoints** page, set the **Type of Traffic** to **FCoE**, set the **Route Ranges** mesh type to **Many to Many**, and select the **FCoE Clients** from **P1** and **P2** on the **Source** side, and **FCoE Clients** from **P3** on the **Destination** side. Click **Apply** when finished. Optionally, set a name in **Traffic Name**.

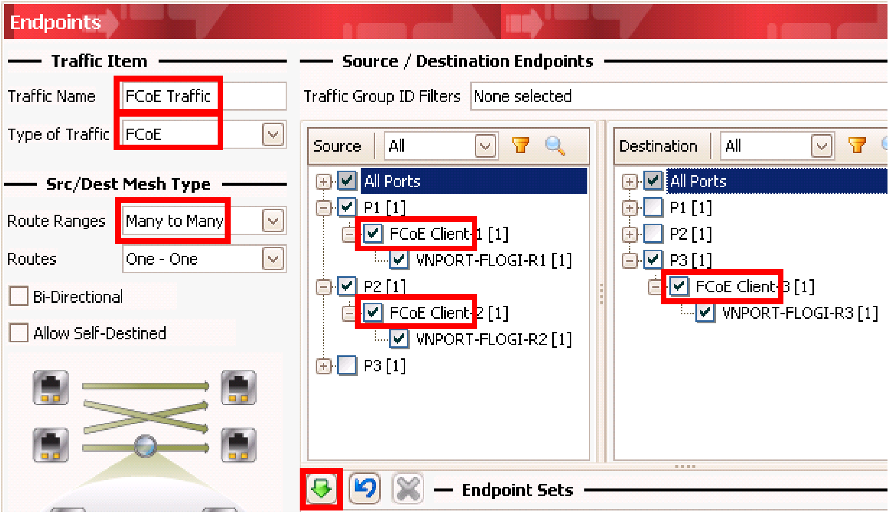

Endpoints

Traffic Item

Traffic Name: FCoE Traffic

Type of Traffic: FCoE

Src/Dest Mesh Type

Route Ranges: Many to Many

Routes: One - One

□ Bi-Directional

□ Allow Self-Destined

Source / Destination Endpoints

Traffic Group ID Filters: None selected

Source: All

```
☑ All Ports
☑ P1 [1]
   ☑ FCoE Client-1 [1]
      ☑ VNPORT-FLOGI-R1 [1]
☑ P2 [1]
   ☑ FCoE Client-2 [1]
      ☑ VNPORT-FLOGI-R2 [1]
□ P3 [1]
```

Destination: All

```
☑ All Ports
□ P1 [1]
□ P2 [1]
☑ P3 [1]
   ☑ FCoE Client-3 [1]
      ☑ VNPORT-FLOGI-R3 [1]
```

Endpoint Sets

Figure 101. Advanced Traffic Wizard Endpoints page – select FCoE endpoints

b. On the **Packet/QoS** page, assign two PFC Queues to the two initiator side VN_Ports using the **PFC Queue** column drop down menu, and enter the following values:

   i. PFC Queue 3

   ii. PFC Queue 4

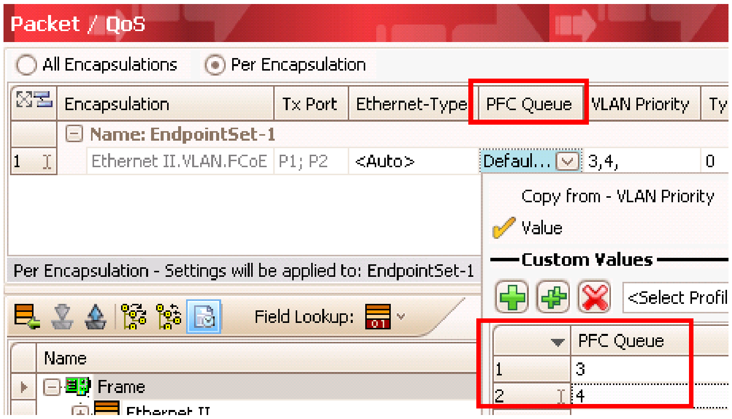

○ All Encapsulations ● Per Encapsulation

	Encapsulation	Tx Port	Ethernet-Type	PFC Queue	VLAN Priority	Ty
	Name: EndpointSet-1					
1	Ethernet II.VLAN.FCoE	P1; P2	<Auto>	Defaul...	3,4,	0

	PFC Queue
1	3
2	4

Figure 102. Advanced Traffic Wizard Packet/QoS page - assign PFC Queues

c. On the **Flow Group Setup** page, select the **Ethernet II: PFC Queue** check box to create flow groups based on PFC Queue values. All packets within the same flow group can be independently rate varied and size varied dynamically. Using the above configuration, two flow groups will be created, one for PFC Queue 3 and one for PFC Queue 4.

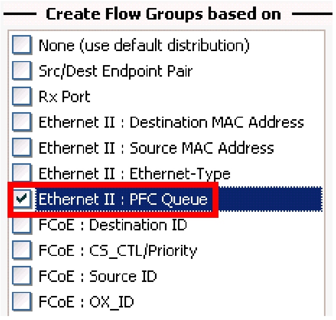

Create Flow Groups based on

- ☐ None (use default distribution)
- ☐ Src/Dest Endpoint Pair
- ☐ Rx Port
- ☐ Ethernet II : Destination MAC Address
- ☐ Ethernet II : Source MAC Address
- ☐ Ethernet II : Ethernet-Type
- ☑ Ethernet II : PFC Queue
- ☐ FCoE : Destination ID
- ☐ FCoE : CS_CTL/Priority
- ☐ FCoE : Source ID
- ☐ FCoE : OX_ID

Figure 103. Advanced Traffic Wizard Flow Group Setup page - assign flow groups

d. On the **Rate Setup** page, set the **Line Rate** to *50* %, and set the **Rate Distribution** to **Split rate evenly among ports**.

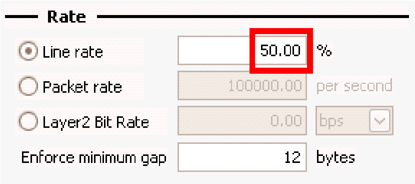

Rate

● Line rate 50.00 %

○ Packet rate 100000.00 per second

○ Layer2 Bit Rate 0.00 bps

Enforce minimum gap 12 bytes

Figure 104. Advanced Traffic Wizard Rate Setup page – configure line rate

e. On the **Flow Tracking** page, select these **Track Flows By** options: **Traffic Item**, **Ethernet II: PFC Queue**, and **FCoE: Destination ID**. Selecting a tracking option allows traffic statistics to be drilled down.

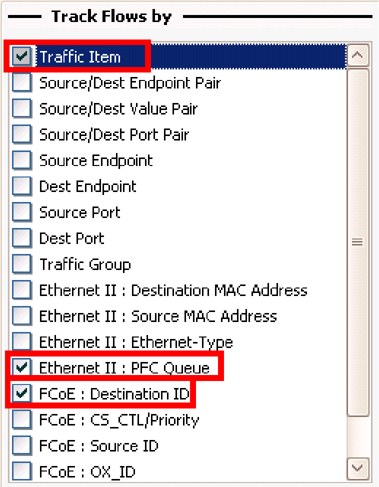

Track Flows by

☑ Traffic Item
□ Source/Dest Endpoint Pair
□ Source/Dest Value Pair
□ Source/Dest Port Pair
□ Source Endpoint
□ Dest Endpoint
□ Source Port
□ Dest Port
□ Traffic Group
□ Ethernet II : Destination MAC Address
□ Ethernet II : Source MAC Address
□ Ethernet II : Ethernet-Type
☑ Ethernet II : PFC Queue
☑ FCoE : Destination ID
□ FCoE : CS_CTL/Priority
□ FCoE : Source ID
□ FCoE : OX_ID

Figure 105. Advanced Traffic Wizard Flow Tracking page – set flow tracking options

f. (Optional) On the **Preview** page, click **View Flow Groups/Packets** to preview the content of the packets that will be transmitted to the DUT. Click **Flow Group** to view the content of each packet. There should be 2 packets created per port, with alternating PFC Queue values.

g. (Optional) On the **Validate** page, click **Validate** to verify hardware resource is available before exiting the traffic wizard.

h. Click **Finish** to exit the traffic wizard.

8. Launch the **Advanced Wizard** to configure native Ethernet traffic.

   a. On the **Endpoints** page, set the **Type of Traffic** to **Ethernet/VLAN**, and select the **MAC/VLAN** from **P1** and **P2** on the **Source** side, and **MAC/VLAN** from **P3** on the **Destination** side. These are the ENode Lossless Ethernet MAC endpoints (each set of FLOGI VN_Port has a unique Lossless Ethernet MAC on its parent ENode). Click **Apply** when finished. Optionally, set a name in **Traffic Name**.

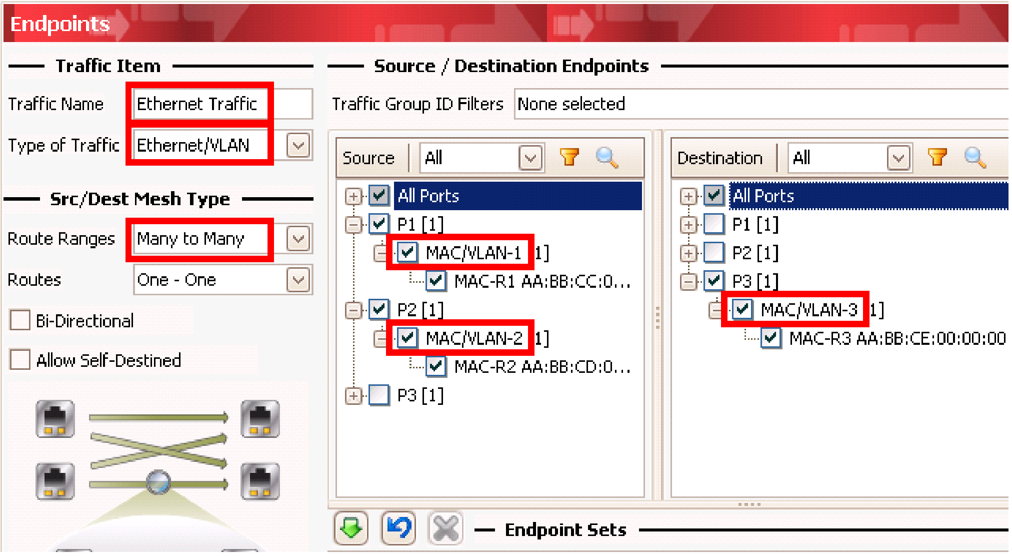

Endpoints

Traffic Item

Traffic Name: Ethernet Traffic

Type of Traffic: Ethernet/VLAN

Src/Dest Mesh Type

Route Ranges: Many to Many

Routes: One - One

□ Bi-Directional

□ Allow Self-Destined

Source / Destination Endpoints

Traffic Group ID Filters: None selected

Source: All

☑ All Ports
☑ P1 [1]
☑ MAC/VLAN-1 [1]
☑ MAC-R1 AA:BB:CC:0...
☑ P2 [1]
☑ MAC/VLAN-2 [1]
☑ MAC-R2 AA:BB:CD:0...
□ P3 [1]

Destination: All

☑ All Ports
□ P1 [1]
□ P2 [1]
☑ P3 [1]
☑ MAC/VLAN-3 [1]
☑ MAC-R3 AA:BB:CE:00:00:00

Figure 106. Advanced Traffic Wizard Endpoints page – select ENode endpoints

   b. On the **Packet/QoS** page, assign one PFC Queue to the two initiator side ENodes using the **PFC Queue** column drop down menu, and enter the following values:

      i. PFC Queue 0

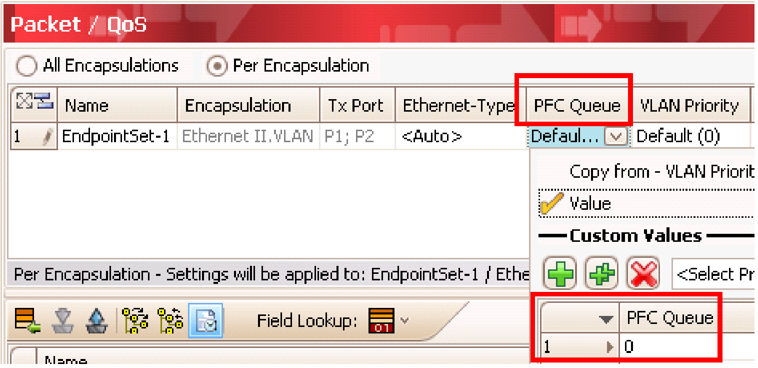

Packet / QoS

○ All Encapsulations ● Per Encapsulation

	Name	Encapsulation	Tx Port	Ethernet-Type	PFC Queue	VLAN Priority
1	EndpointSet-1	Ethernet II.VLAN	P1; P2	<Auto>	Defaul...	Default (0)

Copy from - VLAN Priorit

Value

Custom Values

<Select Pr

Per Encapsulation - Settings will be applied to: EndpointSet-1 / Ethe

	PFC Queue
1	0

Figure 107. Advancted Traffic Wizard Packet/QoS page - assign PFC Queues

c. On the **Rate Setup** page, set the **Line Rate** to *50* %, and set the **Rate Distribution** to **Split rate evenly among ports**.

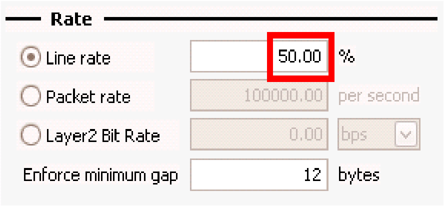

Rate

● Line rate 50.00 %

○ Packet rate 100000.00 per second

○ Layer2 Bit Rate 0.00 bps

Enforce minimum gap 12 bytes

Figure 108. Advanced Traffic Wizard Rate Setup page – configure line rate

d. On the **Flow Group Setup** page, select the **Ethernet II: PFC Queue** check box to create flow groups based on PFC Queue values. All packets within the same flow group can be independently rate varied and size varied dynamically. Using the above configuration, one flow group will be created for PFC Queue 0.

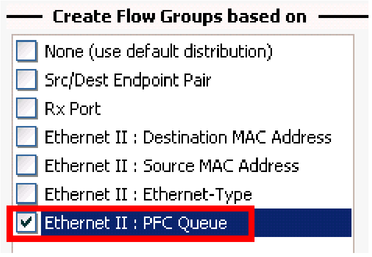

Create Flow Groups based on

☐ None (use default distribution)
☐ Src/Dest Endpoint Pair
☐ Rx Port
☐ Ethernet II : Destination MAC Address
☐ Ethernet II : Source MAC Address
☐ Ethernet II : Ethernet-Type
☑ Ethernet II : PFC Queue

Figure 109. Advanced Traffic Wizard Flow Group Setup page - assign flow groups

e. On the **Flow Tracking** page, select the following **Track Flows By** options: **Traffic Item** and **Ethernet II: Destination MAC Address**. Selecting a tracking option allows traffic statistics to be drilled down to the selected option.

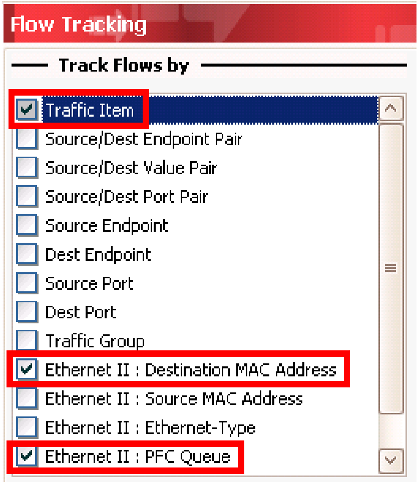

Flow Tracking

Track Flows by

☑ Traffic Item
☐ Source/Dest Endpoint Pair
☐ Source/Dest Value Pair
☐ Source/Dest Port Pair
☐ Source Endpoint
☐ Dest Endpoint
☐ Source Port
☐ Dest Port
☐ Traffic Group
☑ Ethernet II : Destination MAC Address
☐ Ethernet II : Source MAC Address
☐ Ethernet II : Ethernet-Type
☑ Ethernet II : PFC Queue

Figure 110. Advanced Traffic Wizard Flow Tracking page – set flow tracking options

f. (Optional) On the **Preview** page, click **View Flow Groups/Packets** to preview the content of the packets that will be transmitted to the DUT. Click **Flow Group** to view the content of each packet. There should be 1 packet created per port.

g. (Optional) On the **Validate** page, click **Validate** to verify hardware resource is available before exiting the traffic wizard.

h. Click **Finish** to exit the traffic wizard.

9. Click **Apply L2-L3 Traffic** to write the traffic onto the Ixia ports. Click **Start L2-L3 Traffic** to transmit the FCoE and native Ethernet traffic.

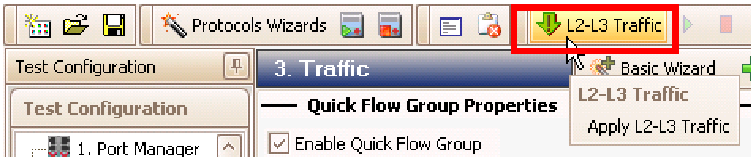

Figure 111. Traffic - apply traffic items

10. Verify part 1 of **Result Analysis** below.

11. Increase one of the FCoE PFC Queues' traffic rate to oversubscribe the DUT's output port. To do so, from the **Test Configuration** menu, select **Traffic – FCoE Traffic** (or the custom name set for the FCoE traffic item in Step 5).

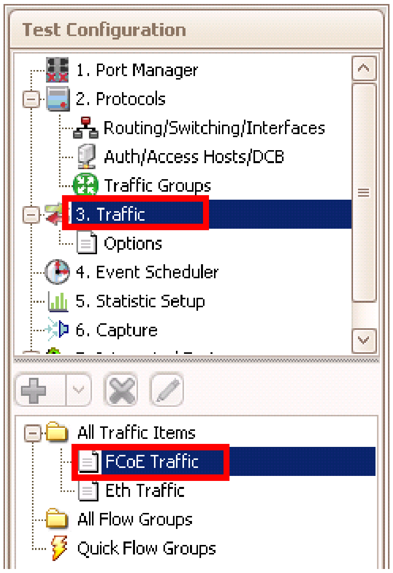

Figure 112. Traffic GUI – FCoE Traffic

From the **Flow Groups** menu, click the **Change grouping** icon and select **Group by Tx Port**.

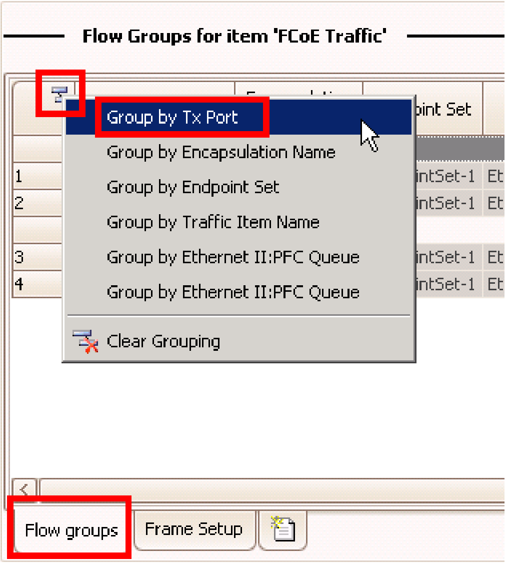
Flow Groups for item 'FCoE Traffic'
Group by Tx Port
Group by Encapsulation Name
Group by Endpoint Set
Group by Traffic Item Name
Group by Ethernet II:PFC Queue
Group by Ethernet II:PFC Queue
Clear Grouping
Flow groups
Frame Setup

Figure 113. Traffic GUI – Flow grouping

Highlight the **Ethernet II: PFC Queue 3** flow group, and increase the traffic rate to be greater than *12.50%* line rate by adjusting the traffic rate slider or manually entering in the percentage line rate value.

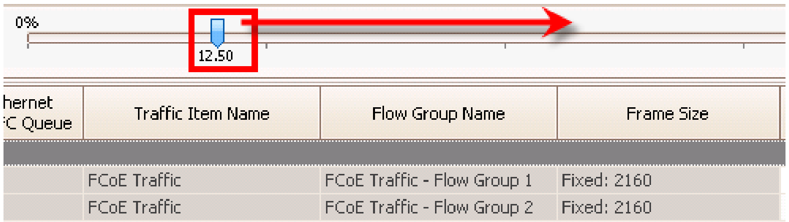

hernet C Queue	Traffic Item Name	Flow Group Name	Frame Size
	FCoE Traffic	FCoE Traffic - Flow Group 1	Fixed: 2160
	FCoE Traffic	FCoE Traffic - Flow Group 2	Fixed: 2160

Figure 114. Traffic GUI – increase FCoE flow group traffic rate

12. Verify part 2 of **Results Analysis** below.

Test Variables

Performance Variable	Description
PFC Queues per port	Up to eight PFC Queues can be selected on each Ixia port. Multiple PFC Queues per port are required to properly verify the PFC PAUSE implementation on the DUT port, as a single PFC Queue per port does not imply per-priority (or per Virtual Lane) awareness.
Frame size	Size of packets (in bytes) transmitted from the test port for the current traffic item or endpoint set (per user setting). The default is **fixed**, with values of 2160 bytes for FCoE traffic, and 64 bytes for native Ethernet traffic. Other options are **increment**, **random**, **IMIX** (9 defined profiles and 1 custom profile) and **quad Gaussian**.
Payload	The content (in hex) encoded in the FC data payload field of the FCoE packet. Use this to create specific payload patterns (for example, 'killer packets').
CRC Settings	The CRC setting for the Ethernet FCS field. The default is a **good** CRC value. Options are **bad** or **none**. Note: A **bad FC CRC** can be inserted using IxExplorer.
Rate	The transmission rate of the current traffic item or endpoint set. The default is **line rate**. The line rate, by default, is split evenly across all flow groups created by the traffic item or endpoint set. Options are **packet rate** and **layer2 bit rate**.
Packet transmission mode	The packet transmission pattern for the current traffic item or endpoint set. The default is **continuous**. Options are **fixed packet count**, **fixed iteration count** and **burst**.

Results Analysis

Part 1

Use the IxNetwork **Statistics – Traffic Item Statistics** view to verify the following real-time FCoE traffic aggregate statistics and native Ethernet traffic aggregate statistics.

- **Rx Frames**: *equal to* ***Tx Frames***
- **Frames Delta**: *0*
- **Rx Frame Rate**: *equal to* ***Tx Frame Rate***

> *Note: At very high line rates, it is normal to see a small mismatch between Tx and Rx stats, and therefore frames delta. These are not necessarily actual frame loss. Rather, it represents the packets that are still in transit across the system under test at the instance the statistics are reported.*

- **Cut-Through Avg Latency**: *vendor specific*
  Verify the average latency is within the expected range. Additional latency and jitter measurement options are available under Traffic – Options settings.

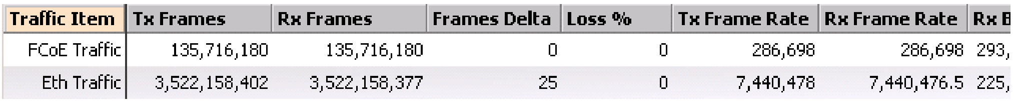

Traffic Item	Tx Frames	Rx Frames	Frames Delta	Loss %	Tx Frame Rate	Rx Frame Rate	Rx B
FCoE Traffic	135,716,180	135,716,180	0	0	286,698	286,698	293,
Eth Traffic	3,522,158,402	3,522,158,377	25	0	7,440,478	7,440,476.5	225,

Figure 115. FCoE and native Ethernet traffic forwarding – traffic item statistics

From the **Traffic Item Statistics** view, **FCoE Traffic** item, use the IxNetwork **DrillDown per Ethernet II: PFC Queue** view to verify traffic is received properly for each PFC Queue in the FCoE traffic.

- **Number of unique PFC Queues**: *2*
- **Rx Frames:** *equal to **Tx Frames** for each PFC Queue*
- **Frames Delta:** *0*
- **Rx Frame Rate:** *equal to **Tx Frame Rate***

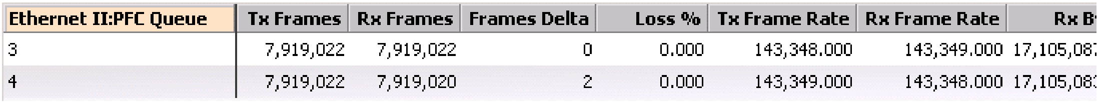

Ethernet II:PFC Queue	Tx Frames	Rx Frames	Frames Delta	Loss %	Tx Frame Rate	Rx Frame Rate	Rx B
3	7,919,022	7,919,022	0	0.000	143,348.000	143,349.000	17,105,08
4	7,919,022	7,919,020	2	0.000	143,349.000	143,348.000	17,105,08

Figure 116. FCoE and native Ethernet traffic forwarding - FCoE PFC Queues statistics

From the **Traffic Item Statistics** view, **Eth Traffic** item, use the IxNetwork **DrillDown per Ethernet II: PFC Queue** view to verify traffic is received properly for each PFC Queue in the native Ethernet traffic.

- **Number of unique PFC Queues**: *1*
- **Rx Frames:** *equal to **Tx Frames** for each PFC Queue*
- **Frames Delta:** *0*
- **Rx Frame Rate:** *equal to **Tx Frame Rate***

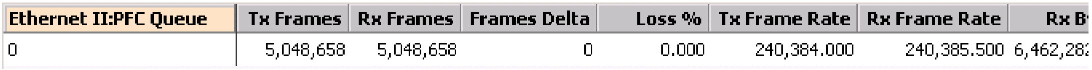

Ethernet II:PFC Queue	Tx Frames	Rx Frames	Frames Delta	Loss %	Tx Frame Rate	Rx Frame Rate	Rx B
0	5,048,658	5,048,658	0	0.000	240,384.000	240,385.500	6,462,28

Figure 117. FCoE and native Ethernet traffic forwarding - native Ethernet PFC Queues statistics

Part 2

As soon as the DUT's output port is oversubscribed, the DUT should issue a PFC PAUSE for the PFC Queues for FCoE traffic (this may require pre-configuration on the DUT). Depending on how oversubscribed the DUT's output port is, it may PAUSE only one of the two transmit ports' PFC Queues 3 and 4, or both transmit ports' PFC Queues 3 and 4. In either cases, the Tx Frame Rate and Rx Frame Rate for the FCoE traffic should be lower than those before increasing the FCoE traffic rate, or zero.

Use the IxNetwork **Statistics – Traffic Item Statistics** view to verify the following real-time FCoE traffic aggregate statistics and native Ethernet traffic aggregate statistics. (Note: it is possible to see a small number of Frames Delta while traffic is running, these are usually packets buffered by the DUT.)

- **Rx Frames**: *equal to **Tx Frames***
- **Frames Delta**: *0*
- **Rx Frame Rate**: *equal to **Tx Frame Rate***

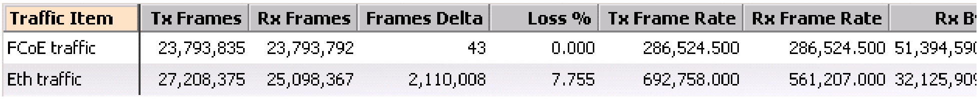

Traffic Item	Tx Frames	Rx Frames	Frames Delta	Loss %	Tx Frame Rate	Rx Frame Rate	Rx B
FCoE traffic	23,793,835	23,793,792	43	0.000	286,524.500	286,524.500	51,394,59
Eth traffic	27,208,375	25,098,367	2,110,008	7.755	692,758.000	561,207.000	32,125,90

Figure 118. FCoE and native Ethernet traffic forwarding – PFC Queues 3 and 4 are paused

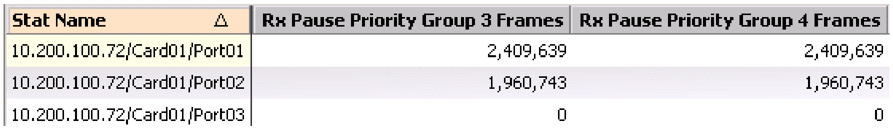

Stat Name △	Rx Pause Priority Group 3 Frames	Rx Pause Priority Group 4 Frames
10.200.100.72/Card01/Port01	2,409,639	2,409,639
10.200.100.72/Card01/Port02	1,960,743	1,960,743
10.200.100.72/Card01/Port03	0	0

Figure 119. FCoE and native Ethernet traffic forwarding – PFC Queues 3 and 4 PAUSE frames

If multiple PFC Queues are assigned to the FCoE traffic, it is possible to view which PFC Queue was paused. From the **Traffic Item Statistics** view, **Eth Traffic** item, use the IxNetwork **DrillDown per Ethernet II: PFC Queue** view to verify which PFC Queue was paused.

Conclusions

By validating the traffic statistics using the features above, the DUT has been proven capable of forwarding FCoE traffic and native Ethernet traffic on the same port at the configured line rate. In addition, the DUT is capable of detecting congestion on its output port, and issues a PFC PAUSE for one or more PFC Queues near or during congestion conditions to prevent or minimize loss of user traffic data.

Test Case: FCoE to FC Traffic Forwarding Test with End-to-End Flow Control (PFC and R_RDY)

Overview

While the converged data center network can be all Ethernet end-to-end, in most practical scenarios, the network will have both Ethernet and fibre channel components. The FCoE capable switch (that is, FCF) holds the responsibility to transport fibre channel traffic between the FCoE domain and the FC domain without loss. To accomplish the task effectively, the flow control functions in both Ethernet and fibre channel are crucial to the overall performance of the converged network and these functions must operate flawlessly at all times.

In the lossless Ethernet network, both FCoE and LAN traffic are forwarded by the FCF according to the FCoE and Ethernet forwarding rules (see **Test Case: Converged Traffic Forwarding – FCoE and LAN Traffic**). Between the lossless Ethernet network and the fibre channel network, fibre channel traffic will flow across the two networks if the traffic endpoints reside in different networks.

In an all fibre channel environment, the source endpoint transmits data packets to the destination endpoint through its directly attached switch port (that is, the F_Port). The switch then forwards the data packets to the destination port out of one of its downstream F_Ports, and returns an R_RDY to the source endpoint (that is, the N_Port) to restore the credit so that another data packet can be received. When the destination port receives the data packet, it transmits an R_RDY to its directly attached F_Port to restore the credit so another data packet can be received. In the case that the destination N_Port is not returning R_RDY fast enough (for example, a slow drain disk), the switch buffer will begin to fill up. When the buffer reaches a certain threshold as defined by BB_Credits, the switch effectively engages flow control on the source N_Port by not sending R_RDYs to the source N_Port until its buffer can accommodate another data packet. This behavior causes the source N_Port to increment BB_Credit_CNT, and when BB_Credit_CNT equals the BB_Credit, the source N_Port will stop transmitting data packets.

In an all Ethernet environment, the source endpoint transmits data packets to the destination endpoint across the network through one or more Ethernet switches (that is, IEEE 802.1Q compliant Bridges), under the forwarding rules of a Bridged network. A pure Bridged network forwards data in a multipoint connectionless environment, where incoming packets are buffered and then delivered as best-effort. Because there is no credit reservation concept in Bridged networks, it is possible for an egress port to be oversubscribed, causing traffic to be dropped. To ensure that FCoE traffic will be lossless during congestion while allowing other traffic to be delivered as best-effort, priority-based flow control (PFC) is deployed, which augments the traditional Ethernet 802.3x PAUSE flow control mechanism by enabling PAUSE on a per priority basis. In the event where the switch cannot forward data out of an egress Bridge port due to being flow controlled by the downstream port, that is, receiving PFC PAUSE frames from the downstream port, the Bridge will store additional incoming data in its egress port buffer. After the buffer reaches a certain threshold, the ingress Bridge port will start transmitting PFC PAUSE frames to its upstream port to stop additional incoming data.

In a mixed Ethernet and fibre channel environment where fibre channel traffic is transmitted from the lossless Ethernet domain to the fibre channel domain, the two flow control mechanisms described earlier operate independently within each domain, and the FCF is responsible for properly managing the congestion in both domains to ensure fibre channel traffic will remain lossless.

This test focuses on forwarding converged traffic from the lossless Ethernet cloud to FCoE switch, where the LAN traffic is forwarded to Ethernet endpoints and the FCoE traffic is forwarded to the fibre channel endpoints. R_RDY and priority-based flow control statistics are used to validate end-to-end FCoE to FC flow control. Impairment to R_RDY is then used to validate congestion management between the lossless Ethernet cloud and the fibre channel cloud.

Objective

The objective of this test is to verify the ability of an FCF (that is, the DUT) to forward converged traffic at line rate in an end-to-end environment with both lossless Ethernet and fibre channel networks. On the lossless Ethernet side, FCoE traffic will be assigned a different PFC Queue than the LAN traffic. During the test, FCoE traffic is transmitted at a higher rate than the fibre channel endpoints can drain, resulting in an oversubscription scenario. The DUT is expected to monitor the available credits on the fibre channel link, and issue PFC PAUSE frames for the FCoE traffic on the lossless Ethernet link. Zero packet loss is expected for the FCoE traffic.

Setup

Four Ixia FCoE ports and two Ixia FC ports are used in this test. The FCoE ports are on the initiator side, and the FC ports are on the target side.

The objective of this test is to instantiate both FCoE VN_Ports and FC N_Ports, and forward converged traffic from the initiator side VN_Ports and ENodes, where the LAN traffic is forwarded to other initiator side ENodes and the FCoE traffic is forwarded to the target side N_Ports.

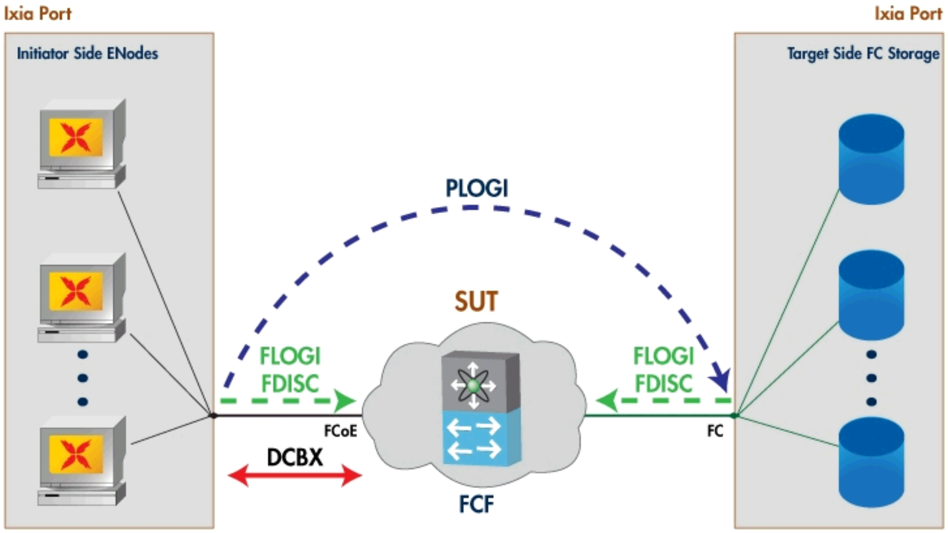

Figure 120. End-to-end FCoE to FC converged traffic forwarding test topology

Step-by-Step Instructions

1. Reserve four Ixia FCoE ports and two Ixia FC ports.

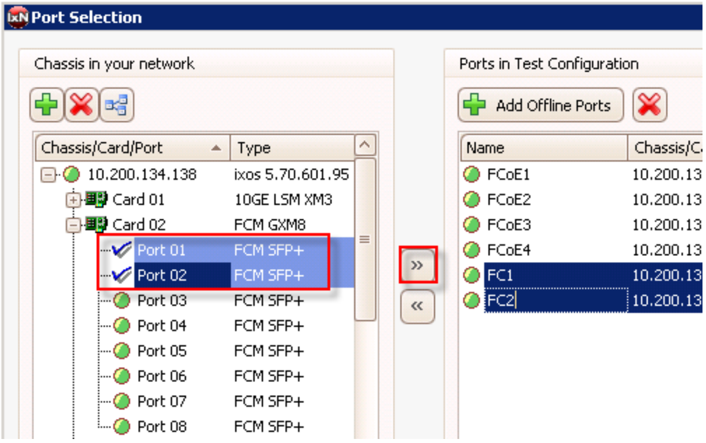

Chassis/Card/Port	Type
10.200.134.138	ixos 5.70.601.95
Card 01	10GE LSM XM3
Card 02	FCM GXM8
Port 01	FCM SFP+
Port 02	FCM SFP+
Port 03	FCM SFP+
Port 04	FCM SFP+
Port 05	FCM SFP+
Port 06	FCM SFP+
Port 07	FCM SFP+
Port 08	FCM SFP+

Name	Chassis/C
FCoE1	10.200.13
FCoE2	10.200.13
FCoE3	10.200.13
FCoE4	10.200.13
FC1	10.200.13
FC2	10.200.13

Figure 121. FCoE to FC Traffic Forwarding Test with End-to-End Flow Control (PFC and R_RDY) - port selection

2. Set the 10GE port **Type** to either **10GE LAN – FCoE** or **10GE WAN – FCoE**.

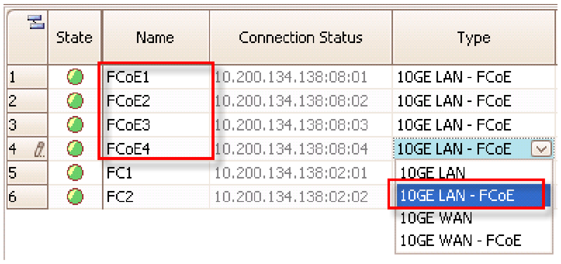

	State	Name	Connection Status	Type
1	●	FCoE1	10.200.134.138:08:01	10GE LAN - FCoE
2	●	FCoE2	10.200.134.138:08:02	10GE LAN - FCoE
3	●	FCoE3	10.200.134.138:08:03	10GE LAN - FCoE
4	●	FCoE4	10.200.134.138:08:04	10GE LAN - FCoE
5	●	FC1	10.200.134.138:02:01	
6	●	FC2	10.200.134.138:02:02	

Figure 122. FCoE to FC Traffic Forwarding Test with End-to-End Flow Control (PFC and R_RDY) – setting port type to FCoE

3. Set the FC port **Speed** to a value supported by the DUT: **2G**, **4G,** or **8G**.

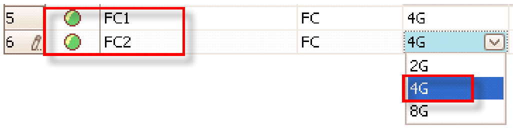

5	●	FC1	FC	4G
6	●	FC2	FC	4G

4. Open the IxNetwork **Protocol Wizards** window, and run the **FCoE/FC Node** wizard to configure the end-to-end test parameters.

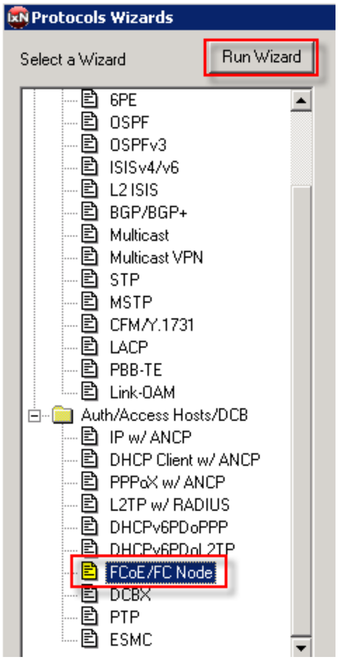

Figure 123. FCoE to FC Traffic Forwarding Test with End-to-End Flow Control (PFC and R_RDY) – FCoE/FC Node protocol wizard

5. **FCoE/FC Node Wizard – Port Selection**: Configure Ixia FCoE Ports 1 – 4 to be the **Initiator** side ports, and Ixia FC Ports 1 - 2 to be the **Target** side ports. In addition, select the **Add DCBX stack to selected Port Groups** check box to enable DCBX on the FCoE ports.

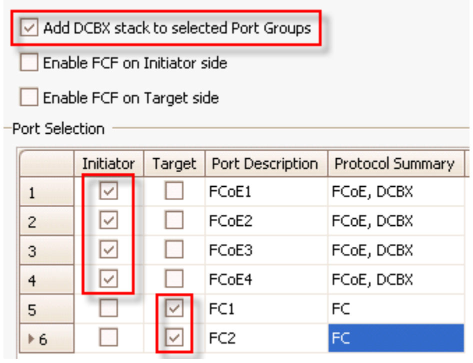

☑ Add DCBX stack to selected Port Groups

□ Enable FCF on Initiator side

□ Enable FCF on Target side

Port Selection

	Initiator	Target	Port Description	Protocol Summary
1	☑	□	FCoE1	FCoE, DCBX
2	☑	□	FCoE2	FCoE, DCBX
3	☑	□	FCoE3	FCoE, DCBX
4	☑	□	FCoE4	FCoE, DCBX
5	□	☑	FC1	FC
6	□	☑	FC2	FC

Figure 124. FCoE to FC Traffic Forwarding Test with End-to-End Flow Control (PFC and R_RDY) – FCoE/FC Node Wizard – port selection

6. **FCoE/FC Node Wizard – Flow Control**: Leave the default values.

7. **FCoE/FC Node Wizard – MAC**: Set the starting MAC Address value and its increment behavior across ENodes. The **First MAC Address** value is the MAC address assigned to the first ENode. The **Increment By** value is the MAC address increment step for all subsequent ENodes on the same port. The **Range Increment Step** value is the MAC address increment step for ENodes across ports.

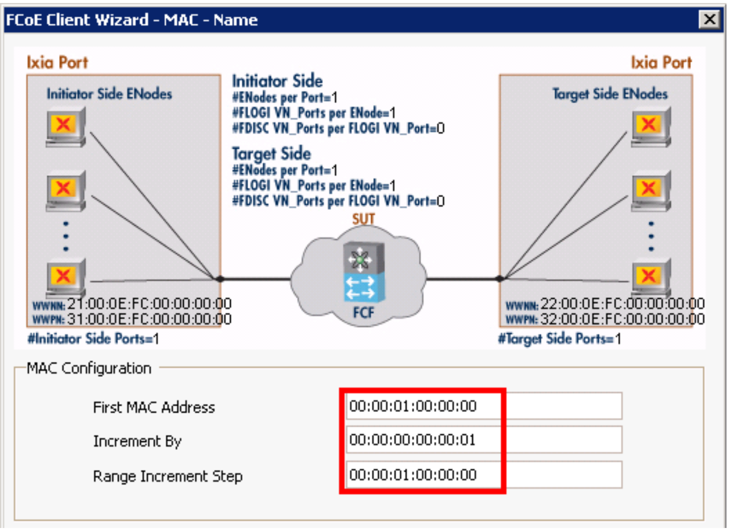

FCoE Client Wizard - MAC - Name

Initiator Side
#ENodes per Port=1
#FLOGI VN_Ports per ENode=1
#FDISC VN_Ports per FLOGI VN_Port=0

Target Side
#ENodes per Port=1
#FLOGI VN_Ports per ENode=1
#FDISC VN_Ports per FLOGI VN_Port=0

WWNN: 21:00:0E:FC:00:00:00:00
WWPN: 31:00:0E:FC:00:00:00:00
#Initiator Side Ports=1

WWNN: 22:00:0E:FC:00:00:00:00
WWPN: 32:00:0E:FC:00:00:00:00
#Target Side Ports=1

MAC Configuration

First MAC Address	00:00:01:00:00:00
Increment By	00:00:00:00:00:01
Range Increment Step	00:00:01:00:00:00

Figure 125. FCoE to FC Traffic Forwarding Test with End-to-End Flow Control (PFC and R_RDY) – MAC

8. **FCoE/FC Node Wizard – VLAN**: Leave the default values.

9. **FCoE/FC Node Wizard – LLDP and DCBX Settings**: Select the **Enable DCBX** check box and select **IEEE 1.01** as the DCBX **Subtype**. 1.01 is also known as version CEE. If the

DUT only supports version pre-CEE, set the **Subtype** to **IEEE 1.00**. IEEE 1.00 supports additional TLVs, such as the FCoE and LAN Logical Link Status TLVs. Optionally, change the default values for the LLDP parameters.

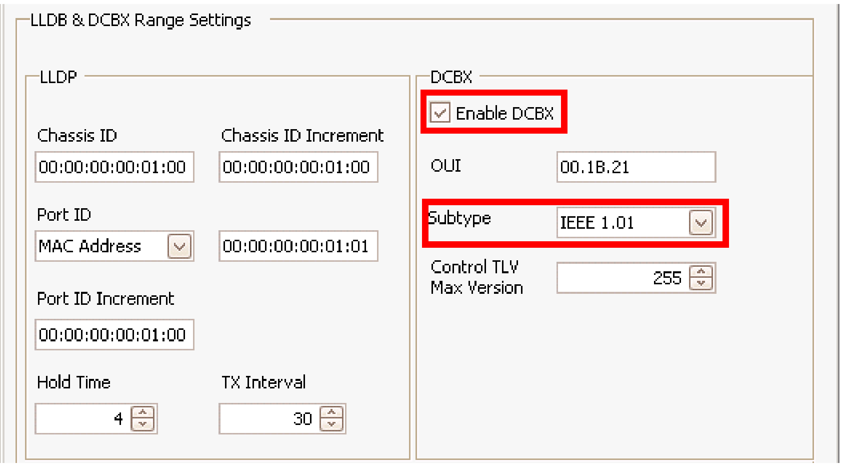

LLDB & DCBX Range Settings

LLDP

Chassis ID: 00:00:00:00:01:00

Chassis ID Increment: 00:00:00:00:01:00

Port ID: MAC Address — 00:00:00:00:01:01

Port ID Increment: 00:00:00:00:01:00

Hold Time: 4

TX Interval: 30

DCBX

☑ Enable DCBX

OUI: 00.1B.21

Subtype: IEEE 1.01

Control TLV Max Version: 255

Figure 126. FCoE to FC Traffic Forwarding Test with End-to-End Flow Control (PFC and R_RDY) - LLDP and DCBX Settings

10. **FCoE/FC Node Wizard – DCBX TLV options**: Click **Append** to add DCBX feature TLVs and set the TLV parameters such as User Priority map for the TLVs. By default, the Willing bit is set for each DCBX feature TLV, so IxNetwork will negotiate to the DUT's DCBX settings if the Willing bit on the DUT is off, which is often the case.

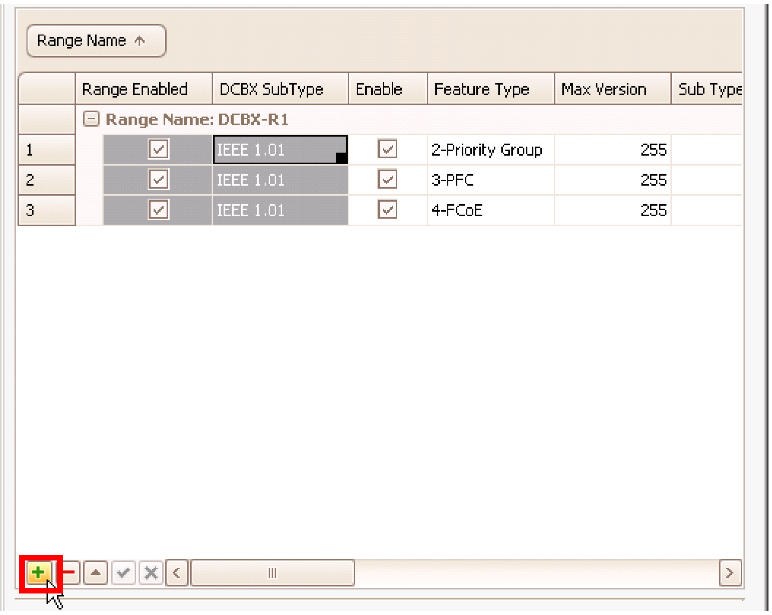

Range Name ↑

	Range Enabled	DCBX SubType	Enable	Feature Type	Max Version	Sub Type
	⊟ Range Name: DCBX-R1					
1	☑	IEEE 1.01	☑	2-Priority Group	255	
2	☑	IEEE 1.01	☑	3-PFC	255	
3	☑	IEEE 1.01	☑	4-FCoE	255	

Figure 127. FCoE to FC Traffic Forwarding Test with End-to-End Flow Control (PFC and R_RDY) – DCBX Feature TLV Settings

9. **FCoE/FC Node Wizard – FIP**: Configure the FCoE/FIP global protocol features. Select the **Enable Name Server Registration**, **Perform PLOGI**, **Enable FCoE Initialization**

Protocol (FIP), and **Enable FIP VLAN Discovery** check boxes. Leave the **Addressing Capability Mode** at **FPMA** (or the mode that the DUT supports), and leave **FIP Max FCoE Size** at *2158* (or the value that the DUT supports).
*Note: Some FCFs expect FIP VLAN Request messages to be untagged versus priority tagged, and some FCFs expect the FCoE endpoints to not propose a MAC address in FPMA mode versus proposing a non-zero MAC address in FPMA mode. Select the **Untagged VLAN Discovery** check box or **Propose MAC in FPMA** check box appropriately based on FCF implementation.*

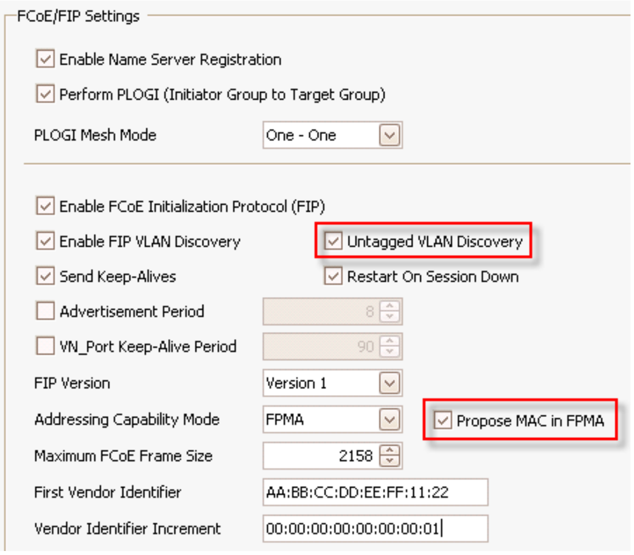

FCoE/FIP Settings

☑ Enable Name Server Registration

☑ Perform PLOGI (Initiator Group to Target Group)

PLOGI Mesh Mode: One - One

☑ Enable FCoE Initialization Protocol (FIP)

☑ Enable FIP VLAN Discovery ☑ Untagged VLAN Discovery

☑ Send Keep-Alives ☑ Restart On Session Down

□ Advertisement Period 8

□ VN_Port Keep-Alive Period 90

FIP Version: Version 1

Addressing Capability Mode: FPMA ☑ Propose MAC in FPMA

Maximum FCoE Frame Size: 2158

First Vendor Identifier: AA:BB:CC:DD:EE:FF:11:22

Vendor Identifier Increment: 00:00:00:00:00:00:00:01

Figure 128. FCoE to FC Traffic Forwarding Test with End-to-End Flow Control (PFC and R_RDY) – FIP Settings

10. **FCoE/FC Node Wizard – Initiator Side Topology**: Configure the number of ENodes, FLOGI VN_Ports, and FDISC VN_Ports, as well as the OUI, Node, and Port WWNs associated with each VN_Port. To configure the target for this test, set the following values:

    a. **Number of ENodes per Port**: *1*

    b. **Number of FLOGI VN_Ports per ENode**: *1*

    c. **Number of FDISC VN_Ports per FLOGI VN_Port**: *0*

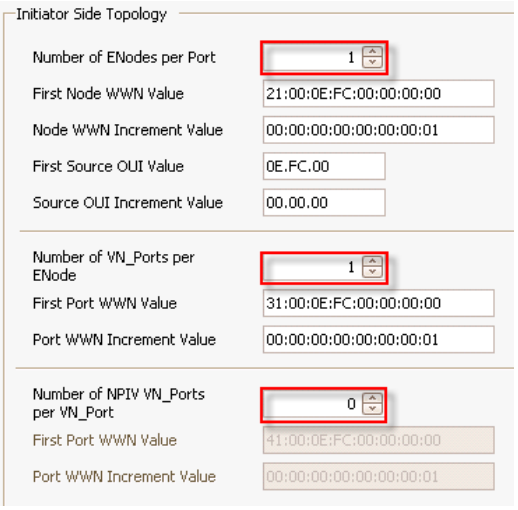

Initiator Side Topology

Number of ENodes per Port: 1

First Node WWN Value: 21:00:0E:FC:00:00:00:00

Node WWN Increment Value: 00:00:00:00:00:00:00:01

First Source OUI Value: 0E.FC.00

Source OUI Increment Value: 00.00.00

Number of VN_Ports per ENode: 1

First Port WWN Value: 31:00:0E:FC:00:00:00:00

Port WWN Increment Value: 00:00:00:00:00:00:00:01

Number of NPIV VN_Ports per VN_Port: 0

First Port WWN Value: 41:00:0E:FC:00:00:00:00

Port WWN Increment Value: 00:00:00:00:00:00:00:01

Figure 129. FCoE to FC Traffic Forwarding Test with End-to-End Flow Control (PFC and R_RDY) – initiator side topology

11. **FCoE/FC Node Wizard – Target Side Topology**: Configure the target side topology with the following values:

    a. **Number of NPIVs per N_Port**: *0*

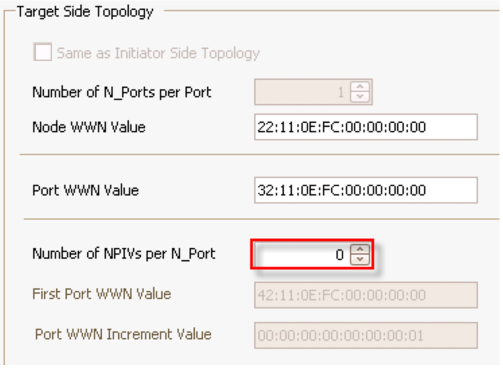

Target Side Topology

☐ Same as Initiator Side Topology

Number of N_Ports per Port: 1

Node WWN Value: 22:11:0E:FC:00:00:00:00

Port WWN Value: 32:11:0E:FC:00:00:00:00

Number of NPIVs per N_Port: 0

First Port WWN Value: 42:11:0E:FC:00:00:00:00

Port WWN Increment Value: 00:00:00:00:00:00:00:01

Figure 130. FCoE to FC Traffic Forwarding Test with End-to-End Flow Control (PFC and R_RDY) – target side topology

12. **FCoE/FC Node Wizard – Fibre Channel**: Configure mandatory fibre channel specific parameters such as credit and time out values. For this test, configure the following values:

    a. **Buffer-to-Buffer Rx Size**: 2112

    b. **Buffer-to-Buffer Credit**: 16

    c. **Error_Detect_Timeout Mode**: Obtain from Login

    d. **Receiver_Transmitter_Timeout Mode**: Obtain from Login

Note: The default *R_A_TOV value is 10 seconds.*

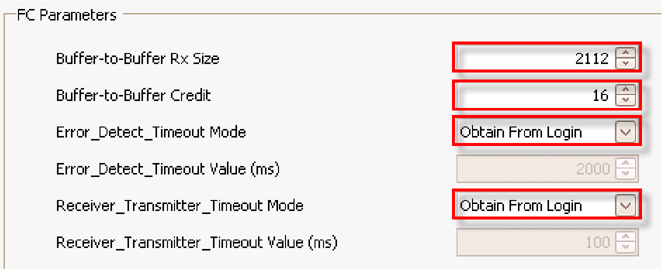

FC Parameters

Buffer-to-Buffer Rx Size: 2112

Buffer-to-Buffer Credit: 16

Error_Detect_Timeout Mode: Obtain From Login

Error_Detect_Timeout Value (ms): 2000

Receiver_Transmitter_Timeout Mode: Obtain From Login

Receiver_Transmitter_Timeout Value (ms): 100

Figure 131. FCoE to FC Traffic Forwarding Test with End-to-End Flow Control (PFC and R_RDY) – FC Parameters

13. **FCoE/FC Node Wizard – Name**: Assign a name to this test configuration, and click **Generate and Overwrite Existing Configuration** to apply the configurations defined in this test.

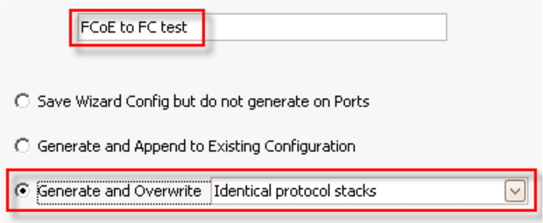

FCoE to FC test

○ Save Wizard Config but do not generate on Ports

○ Generate and Append to Existing Configuration

● Generate and Overwrite: Identical protocol stacks

Figure 132. FCoE to FC Traffic Forwarding Test with End-to-End Flow Control (PFC and R_RDY) – Name

14. From the **FCoE Client** window in the GUI, modify the PLOGI connections as follows:

    a. Port FCoE1 → Port FC1

    b. Port FCoE2 → Port FC2

    c. Port FCoE3 → Port FC1

    d. Port FCoE4 → Port FC2

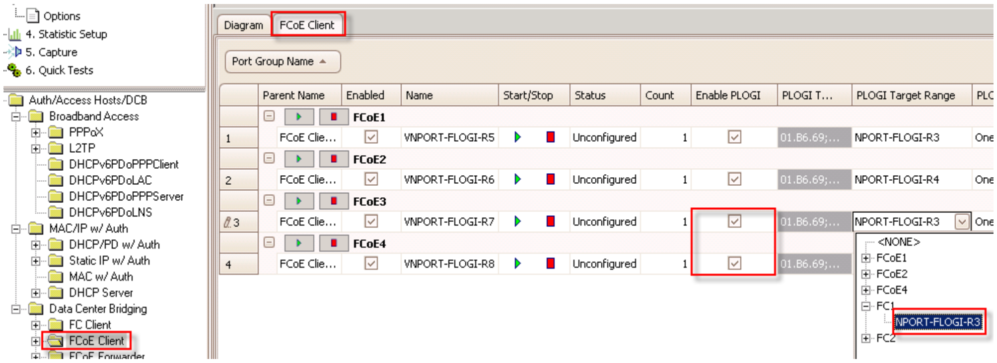

	Parent Name	Enabled	Name	Start/Stop	Status	Count	Enable PLOGI	PLOGI T…	PLOGI Target Range	PLC
	FCoE1									
1	FCoE Clie…	☑	VNPORT-FLOGI-R5	▶ ■	Unconfigured	1	☑	01.B6.69;…	NPORT-FLOGI-R3	One
	FCoE2									
2	FCoE Clie…	☑	VNPORT-FLOGI-R6	▶ ■	Unconfigured	1	☑	01.B6.69;…	NPORT-FLOGI-R4	One
	FCoE3									
3	FCoE Clie…	☑	VNPORT-FLOGI-R7	▶ ■	Unconfigured	1	☑	01.B6.69;…	NPORT-FLOGI-R3	One
	FCoE4									
4	FCoE Clie…	☑	VNPORT-FLOGI-R8	▶ ■	Unconfigured	1	☑	01.B6.69;…		

Figure 133. FCoE to FC Traffic Forwarding Test with End-to-End Flow Control (PFC and R_RDY) – FCoE Client configuration GUI view

15. Review the configuration by validating the emulated ENodes, FCoE VN_Ports, and FC N_Ports created. IxNetwork provides tabs for several categories to organize the configuration: **All**, **FIP**, **VN_Port (FLOGI)**, **VN_Port (FDISC)**, **MAC**, and **VLAN** for FCoE; **All**, **FLOGI**, **PLOGI**, **FDISC (NPIV),** and **PLOGI (NPIV)** for FC.

    a. FCoE: The **All** tab should show that 1 ENode (named **VNPORT-FLOGI-Rxx**) has been created on each Ixia FCoE port. The **FIP** tab should show that all ENodes have FIP enabled with untagged FIP VLAN Discovery. The **VN_Port (FLOGI)** tab should show that each ENode has 1 VN_Port configured. In addition, only ENodes on Ixia port 1 have the **PLOGI Target** defined. The **VN_Port (FDISC)** tab should show that each FLOGI VN_Port has no FDISC VN_Ports are configured.

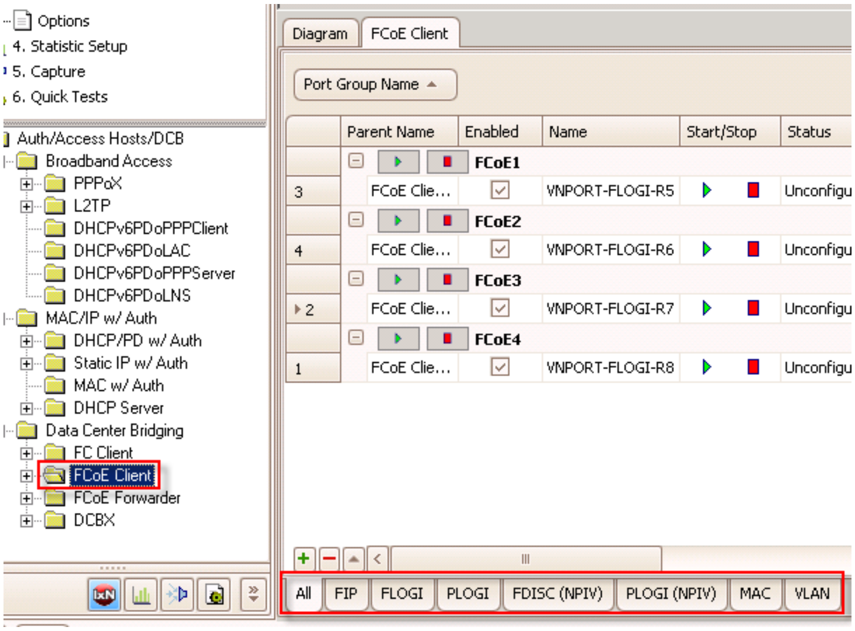

	Parent Name	Enabled	Name	Start/Stop	Status
	FCoE1				
3	FCoE Clie...	☑	VNPORT-FLOGI-R5		Unconfigu
	FCoE2				
4	FCoE Clie...	☑	VNPORT-FLOGI-R6		Unconfigu
	FCoE3				
2	FCoE Clie...	☑	VNPORT-FLOGI-R7		Unconfigu
	FCoE4				
1	FCoE Clie...	☑	VNPORT-FLOGI-R8		Unconfigu

Figure 134. FCoE to FC Traffic Forwarding Test with End-to-End Flow Control (PFC and R_RDY) – FCoE Client configuration GUI view

b. FC: The **All/FLOGI** tab should show that 1 N_Port (named **NPORT-FLOGI-Rxx**) has been created on each Ixia FC port. The **FDISC (NPIV)** tab should show that no NPIVs are configured.

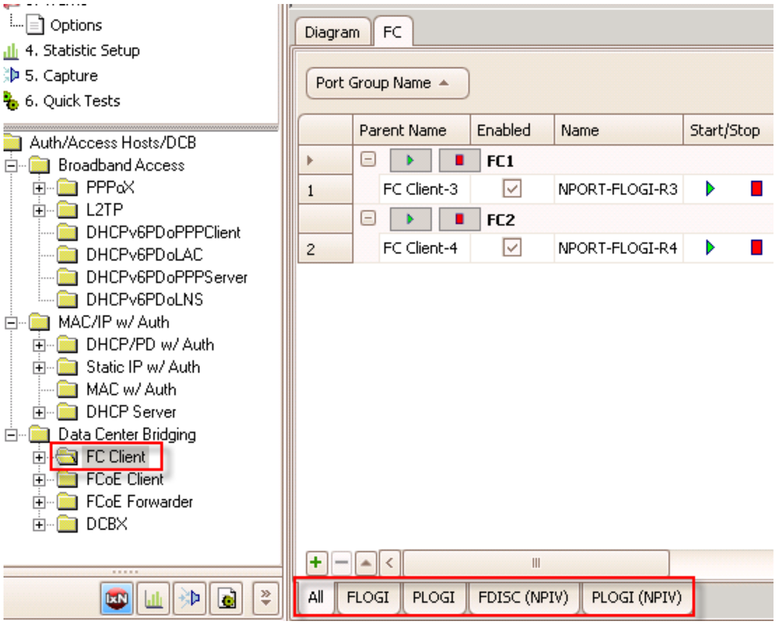

	Parent Name	Enabled	Name	Start/Stop
	FC1			
1	FC Client-3	☑	NPORT-FLOGI-R3	▶ ■
	FC2			
2	FC Client-4	☑	NPORT-FLOGI-R4	▶ ■

Figure 135. FCoE to FC Traffic Forwarding Test with End-to-End Flow Control (PFC and R_RDY) – FC client configuration GUI view

16. Optionally, configure the login **Setup Rate** and log out **Teardown Rate** for the FCoE ports and/or FC ports. These values change the rate at which FLOGIs or FDISCs are transmitted to the FCF, and the rate at which LOGOs are transmitted to the FCF. Many other optional parameters exist on this page.

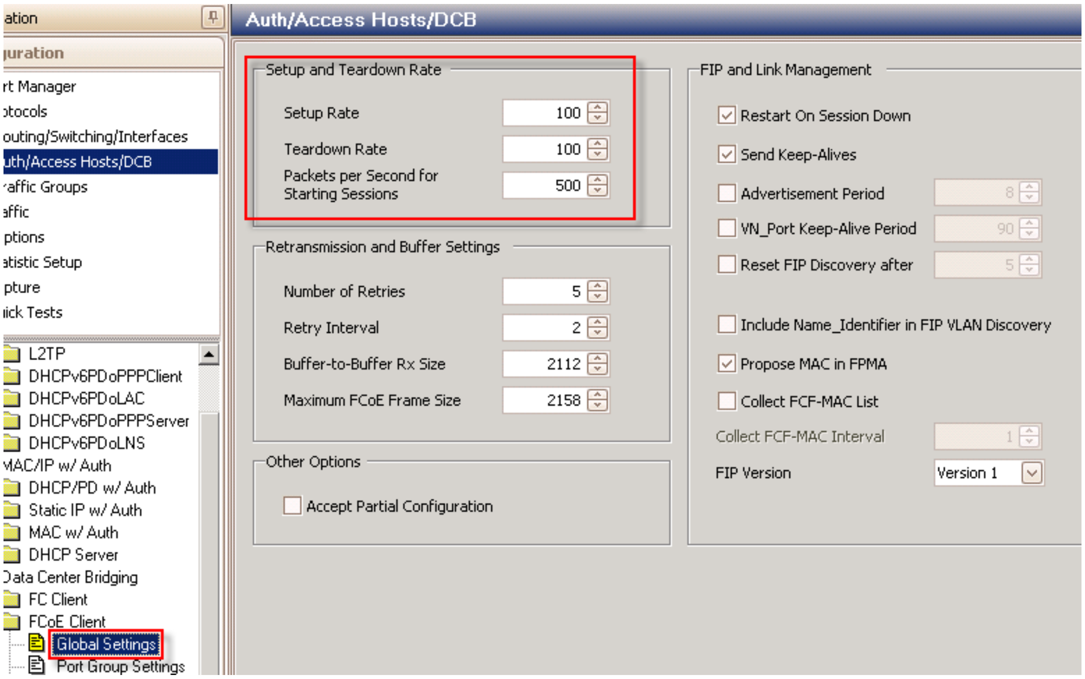

Auth/Access Hosts/DCB

Setup and Teardown Rate

Setup Rate: 100

Teardown Rate: 100

Packets per Second for Starting Sessions: 500

Retransmission and Buffer Settings

Number of Retries: 5

Retry Interval: 2

Buffer-to-Buffer Rx Size: 2112

Maximum FCoE Frame Size: 2158

Other Options

□ Accept Partial Configuration

FIP and Link Management

☑ Restart On Session Down

☑ Send Keep-Alives

□ Advertisement Period: 8

□ VN_Port Keep-Alive Period: 90

□ Reset FIP Discovery after: 5

□ Include Name_Identifier in FIP VLAN Discovery

☑ Propose MAC in FPMA

□ Collect FCF-MAC List

Collect FCF-MAC Interval: 1

FIP Version: Version 1

Figure 136. FCoE to FC Traffic Forwarding Test with End-to-End Flow Control (PFC and R_RDY) – Login Setup and Logout Teardown Rates

17. Start the emulation.

18. Wait until all FCoE VN_Ports and FC ports are instantiated. This process can take up to one minute if DCBX is enabled, as LLDPDUs are transmitted once every 30 seconds.

19. Verify Part 1 of **Results Analysis** below. Use the IxNetwork FCoE aggregated and per-VN_Port statistics to verify that all FCoE VN_Ports are properly instantiated, and use the IxNetwork FC aggregated and per-N_Port/NPIV statistics to verify that all FC N_Ports and NPIVs are properly instantiated.

20. Start the **Advanced Wizard** to configure FCoE traffic.

   a. On the **Endpoints** page, set the **Type of Traffic** to **FC**, set the **Route Ranges** mesh type to **Many to Many**, and select the **FCoE1** and **FCoE3** on the **Source** side, and **FC1** on the **Destination** side. Click **Apply** when finished. Next, add another endpoint set with **FCoE2** and **FCoE4** on the **Source** side, and **FC2** on the **Destination** side. Click **Apply** when finished. Optionally, set a name in **Traffic Name** for both endpoint sets.

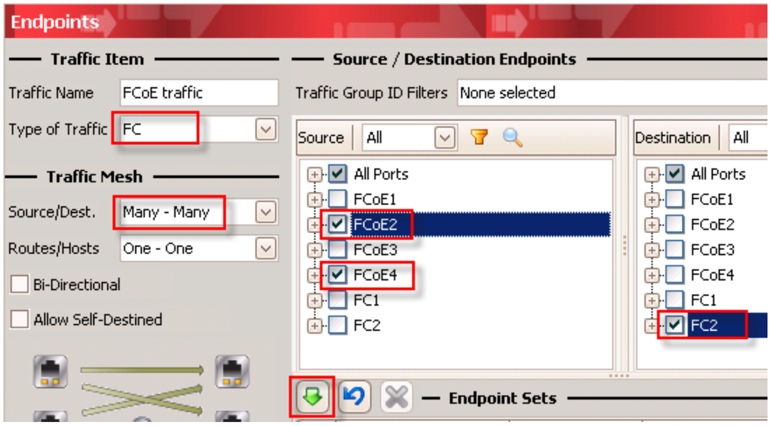

Endpoints

Traffic Item

Traffic Name: FCoE traffic

Type of Traffic: FC

Traffic Mesh

Source/Dest.: Many - Many

Routes/Hosts: One - One

□ Bi-Directional

□ Allow Self-Destined

Source / Destination Endpoints

Traffic Group ID Filters: None selected

Source: All

- ☑ All Ports
- □ FCoE1
- ☑ FCoE2
- □ FCoE3
- ☑ FCoE4
- □ FC1
- □ FC2

Destination: All

- ☑ All Ports
- □ FCoE1
- □ FCoE2
- □ FCoE3
- □ FCoE4
- □ FC1
- ☑ FC2

Endpoint Sets

Figure 137. Advanced Traffic Wizard Endpoints page – select FCoE endpoints

   b. On the **Packet/QoS** page, assign a PFC Queue to the FCoE traffic by using the **PFC Queue** column drop down menu, and select the **From – VLAN Priority** option. IxNetwork will automatically set the PFC Queue to match the VLAN Priority value assigned to FCoE by DCBX, even though this screen does not reflect the actual VLAN Priority value.

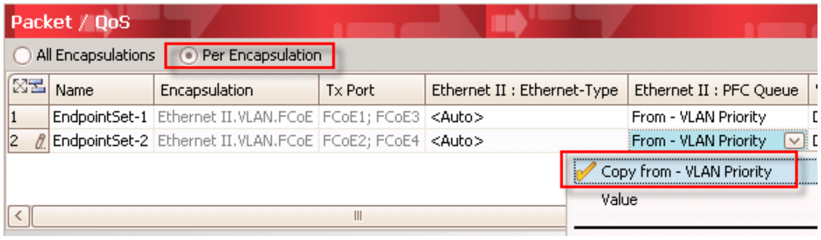

Packet / QoS

○ All Encapsulations ● Per Encapsulation

	Name	Encapsulation	Tx Port	Ethernet II : Ethernet-Type	Ethernet II : PFC Queue
1	EndpointSet-1	Ethernet II.VLAN.FCoE	FCoE1; FCoE3	<Auto>	From - VLAN Priority
2	EndpointSet-2	Ethernet II.VLAN.FCoE	FCoE2; FCoE4	<Auto>	From - VLAN Priority

Copy from - VLAN Priority

Value

Figure 138. Advanced Traffic Wizard Packet/QoS page - assign PFC Queues

c. On the **Flow Group Setup** page, click **All Encapsulations** to apply the following settings to both endpoint sets. Select the **Ethernet II: PFC Queue** check box to create flow groups based on PFC Queue values. All packets within the same flow group can be independently rate varied and size varied dynamically. Using the preceding configuration, one flow group per port will form.

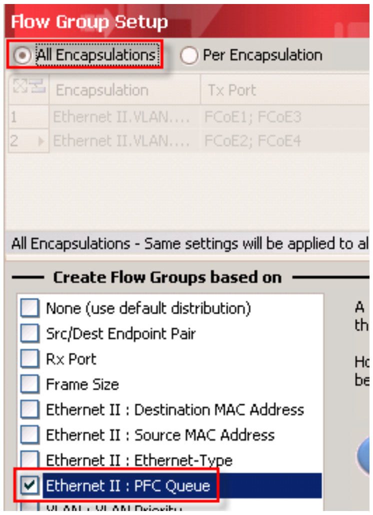
Flow Group Setup

● All Encapsulations ○ Per Encapsulation

	Encapsulation	Tx Port
1	Ethernet II.VLAN....	FCoE1; FCoE3
2	Ethernet II.VLAN....	FCoE2; FCoE4

All Encapsulations - Same settings will be applied to al

Create Flow Groups based on

- □ None (use default distribution)
- □ Src/Dest Endpoint Pair
- □ Rx Port
- □ Frame Size
- □ Ethernet II : Destination MAC Address
- □ Ethernet II : Source MAC Address
- □ Ethernet II : Ethernet-Type
- ☑ Ethernet II : PFC Queue

Figure 139. Advanced Traffic Wizard Flow Group Setup page - assign flow groups

d. On the **Frame Setup** page, set the desired **Frame Size**.

e. On the **Rate Setup** page, set the **Line Rate** to *50* %, and set the **Rate Distribution** to **Split rate evenly among ports**.

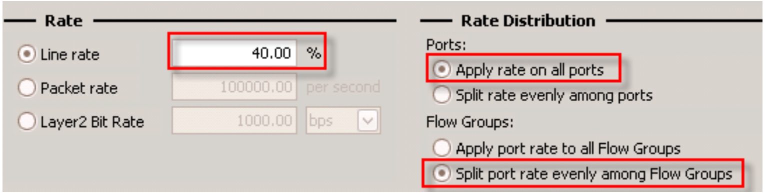
Rate

- ● Line rate 40.00 %
- ○ Packet rate 100000.00 per second
- ○ Layer2 Bit Rate 1000.00 bps

Rate Distribution

Ports:

- ● Apply rate on all ports
- ○ Split rate evenly among ports

Flow Groups:

- ○ Apply port rate to all Flow Groups
- ● Split port rate evenly among Flow Groups

Figure 140. Advanced Traffic Wizard Rate Setup page – configure line rate

f. On the **Flow Tracking** page, select these **Track Flows By** options: **Traffic Item**, **Source Port, Ethernet II: PFC Queue**, and **FCoE: Destination ID**. Selecting a tracking option allows traffic statistics to be drilled down.

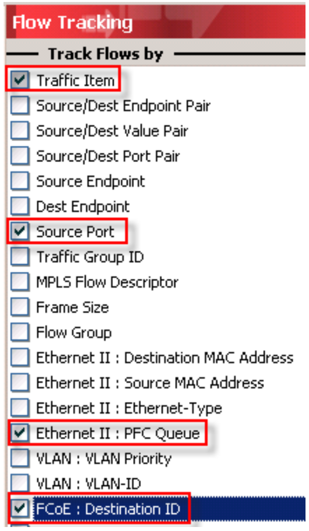

Flow Tracking

Track Flows by

☑ Traffic Item
□ Source/Dest Endpoint Pair
□ Source/Dest Value Pair
□ Source/Dest Port Pair
□ Source Endpoint
□ Dest Endpoint
☑ Source Port
□ Traffic Group ID
□ MPLS Flow Descriptor
□ Frame Size
□ Flow Group
□ Ethernet II : Destination MAC Address
□ Ethernet II : Source MAC Address
□ Ethernet II : Ethernet-Type
☑ Ethernet II : PFC Queue
□ VLAN : VLAN Priority
□ VLAN : VLAN-ID
☑ FCoE : Destination ID

Figure 141. Advanced Traffic Wizard Flow Tracking page – set flow tracking options

g. (Optional) On the **Preview** page, click **View Flow Groups/Packets** to preview the content of the packets that will be transmitted to the DUT. Click **Flow Group** to view the content of each packet. There should be one flow group per port, with **PFC Queue 3** and **VLAN Priority 3**.

h. (Optional) On the **Validate** page, click **Validate** to verify that a hardware resource is available before exiting the traffic wizard.

i. Click **Finish** to exit the traffic wizard.

21. Start the **Advanced Wizard** to configure LAN Ethernet traffic.

   j. On the **Endpoints** page, set the **Type of Traffic** to **Ethernet/VLAN**, set the **Source/Dest** field to **Fully Meshed**, and select the second **MAC/VLAN** from each of **FCoE1, FCoE2** and **FCoE3** on the **Source** side. Click **Apply** when finished. Optionally, set a name in **Traffic Name**.

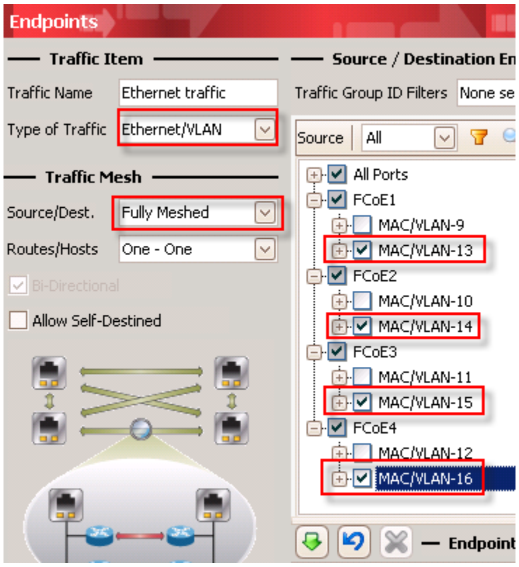

Endpoints

Traffic Item

Traffic Name: Ethernet traffic

Type of Traffic: Ethernet/VLAN

Traffic Mesh

Source/Dest.: Fully Meshed

Routes/Hosts: One - One

☑ Bi-Directional

☐ Allow Self-Destined

Source / Destination En

Traffic Group ID Filters: None se

Source: All

- ☑ All Ports
- ☑ FCoE1
  - ☐ MAC/VLAN-9
  - ☑ MAC/VLAN-13
- ☑ FCoE2
  - ☐ MAC/VLAN-10
  - ☑ MAC/VLAN-14
- ☑ FCoE3
  - ☐ MAC/VLAN-11
  - ☑ MAC/VLAN-15
- ☑ FCoE4
  - ☐ MAC/VLAN-12
  - ☑ MAC/VLAN-16

Figure 142. Advanced Traffic Wizard Endpoints page – select ENode endpoints

   k. On the **Packet/QoS** page, assign one PFC Queue to the LAN Ethernet traffic by using the **PFC Queue** column drop down menu, and enter the following values:

      i. PFC Queue 0

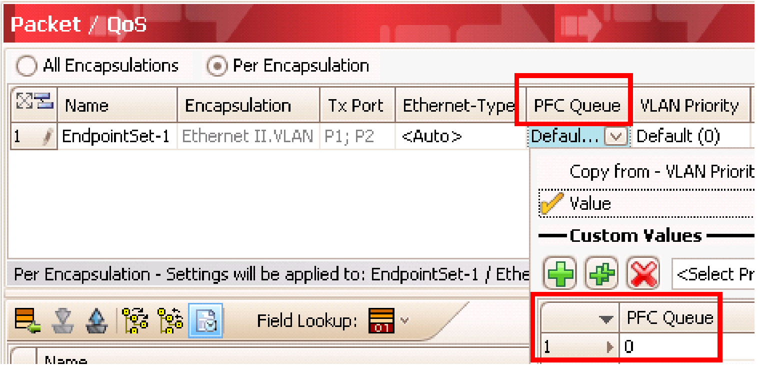

Packet / QoS

○ All Encapsulations ● Per Encapsulation

	Name	Encapsulation	Tx Port	Ethernet-Type	PFC Queue	VLAN Priority
1	EndpointSet-1	Ethernet II.VLAN	P1; P2	<Auto>	Defaul...	Default (0)

Copy from - VLAN Priorit

Value

Custom Values

<Select Pr

Per Encapsulation - Settings will be applied to: EndpointSet-1 / Ethe

	PFC Queue
1	0

Figure 143. Advancted Traffic Wizard Packet/QoS page - assign PFC Queues

   l. On the **Frame Setup** page, set the **Frame Size** to a desirable value.

m. On the **Rate Setup** page, set the **Line Rate** to *60* %, and set the **Rate Distribution** to **Split rate evenly among ports**.

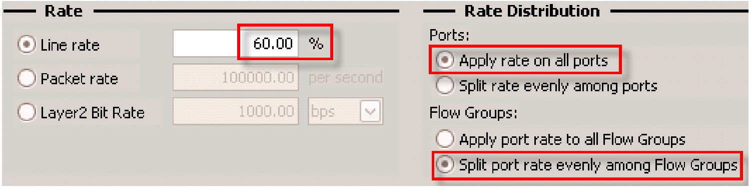

Rate

● Line rate 60.00 %

○ Packet rate 100000.00 per second

○ Layer2 Bit Rate 1000.00 bps

Rate Distribution

Ports:

● Apply rate on all ports

○ Split rate evenly among ports

Flow Groups:

○ Apply port rate to all Flow Groups

● Split port rate evenly among Flow Groups

Figure 144. Advanced Traffic Wizard Rate Setup page – configure line rate

n. On the **Flow Group Setup** page, select the **Ethernet II: PFC Queue** check box to create flow groups based on PFC Queue values. All packets within the same flow group can be independently rate varied and size varied dynamically. By using the preceding configuration, one flow group will be created for PFC Queue 0.

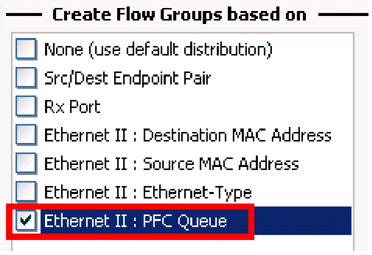

Create Flow Groups based on

☐ None (use default distribution)

☐ Src/Dest Endpoint Pair

☐ Rx Port

☐ Ethernet II : Destination MAC Address

☐ Ethernet II : Source MAC Address

☐ Ethernet II : Ethernet-Type

☑ Ethernet II : PFC Queue

Figure 145. Advanced Traffic Wizard Flow Group Setup page - assign flow groups

o. On the **Flow Tracking** page, select the following **Track Flows By** options: **Traffic Item** and **Ethernet II: Destination MAC Address**. Selecting a tracking option allows traffic statistics to be drilled down to the selected option.

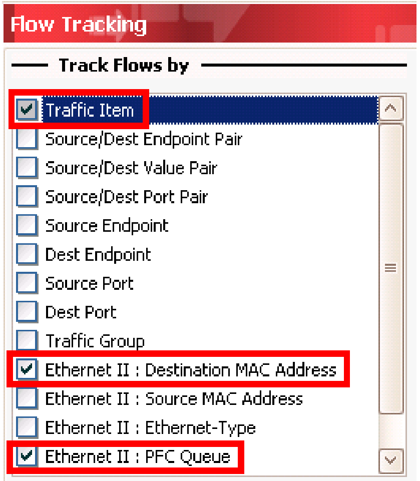

Flow Tracking

Track Flows by

☑ Traffic Item

☐ Source/Dest Endpoint Pair

☐ Source/Dest Value Pair

☐ Source/Dest Port Pair

☐ Source Endpoint

☐ Dest Endpoint

☐ Source Port

☐ Dest Port

☐ Traffic Group

☑ Ethernet II : Destination MAC Address

☐ Ethernet II : Source MAC Address

☐ Ethernet II : Ethernet-Type

☑ Ethernet II : PFC Queue

Figure 146. Advanced Traffic Wizard Flow Tracking page – set flow tracking options

p. (Optional) On the **Preview** page, click **View Flow Groups/Packets** to preview the content of the packets that will be transmitted to the DUT. Click **Flow Group** to view the content of each packet. There should be one flow group created per port.

q. (Optional) On the **Validate** page, click **Validate** to verify that a hardware resource is available before exiting the traffic wizard.

r. Click **Finish** to exit the traffic wizard.

22. Click **Apply L2-L3 Traffic** to write the traffic onto the Ixia ports. Click **Start L2-L3 Traffic** to transmit the FCoE and native Ethernet traffic.

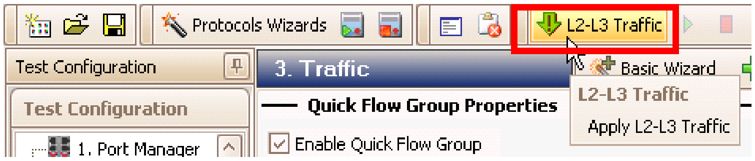

Figure 147. Traffic - apply traffic items

23. Verify part 2 of **Result Analysis** below.

Test Variables

Performance Variable	Description
R_RDY Impairment settings	The R_RDY primitive is transmitted properly by the Ixia FC ports per standard requirement. IxNetwork allows the user to intentionally mis-transmit R_RDYs and inject fixed and random delays into the transmission of R_RDYs to simulate slow-drain storage devices. These settings are in the Port Manager window.
PFC PAUSE Impairment settings	The Ixia FCoE ports reacts to the PFC PAUSE frames from the DUT per standard requirement. IxNetwork allows the user to intentionally delay the response to the incoming PFC PAUSE frame by setting the delay in integer numbers of pause_quanta units. These settings are in the Port Manager window.
Number of NPIVs per VN_Port and N_Port	Enabling NPIV to obtain additional FC_IDs per FCoE VN_Port and per FC N_Port results in additional FCoE/FC traffic endpoints, as well as sessions to maintain by the DUT.
PFC Queues per port	Up to eight PFC Queues can be selected on each Ixia port. Multiple PFC Queues per port are required to properly verify the PFC PAUSE implementation on the DUT port, as a single PFC Queue per port does not imply per-priority (or per Virtual Lane) awareness.
CRC Settings	The CRC setting for the Ethernet FCS field and FC CRC field. The default is a **good** CRC value. Options are **bad CRC** or **none**.
Rate	The transmission rate of the current traffic item or endpoint set. The default is **line rate**. The line rate, by default, is split evenly across all flow groups created by the traffic item or endpoint set. Options are **packet rate** and **layer2 bit rate**.
Packet transmission mode	The packet transmission pattern for the current traffic item or endpoint set. The default is **continuous**. Options are **fixed packet count**, **fixed iteration count**, and **burst**.

Results Analysis

Part 1

Use the IxNetwork **Statistics – Protocol Summary** view to verify the following real-time high level protocol summary statistics.

- **DCBX:** *Sessions Succeeded = 4*
- **FC Client:** *Sessions Succeeded = 2*
- **FCoE Client:** *Sessions Succeeded = 4*

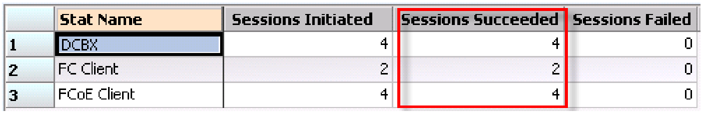

	Stat Name	Sessions Initiated	Sessions Succeeded	Sessions Failed
1	DCBX	4	4	0
2	FC Client	2	2	0
3	FCoE Client	4	4	0

Figure 148. FCoE to FC Traffic Forwarding Test with End-to-End Flow Control (PFC and R_RDY) – Protocol Summary statistics view

Right-click **DCBX** and select **DrillDown Per TLV** to retrieve detailed information about each DCBX TLV.

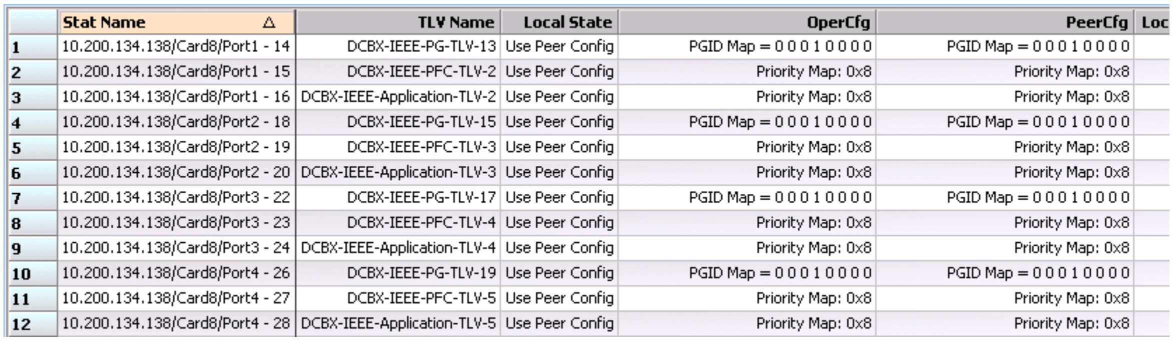

	Stat Name	TLV Name	Local State	OperCfg	PeerCfg	Loc
1	10.200.134.138/Card8/Port1 - 14	DCBX-IEEE-PG-TLV-13	Use Peer Config	PGID Map = 0 0 0 1 0 0 0 0	PGID Map = 0 0 0 1 0 0 0 0	
2	10.200.134.138/Card8/Port1 - 15	DCBX-IEEE-PFC-TLV-2	Use Peer Config	Priority Map: 0x8	Priority Map: 0x8	
3	10.200.134.138/Card8/Port1 - 16	DCBX-IEEE-Application-TLV-2	Use Peer Config	Priority Map: 0x8	Priority Map: 0x8	
4	10.200.134.138/Card8/Port2 - 18	DCBX-IEEE-PG-TLV-15	Use Peer Config	PGID Map = 0 0 0 1 0 0 0 0	PGID Map = 0 0 0 1 0 0 0 0	
5	10.200.134.138/Card8/Port2 - 19	DCBX-IEEE-PFC-TLV-3	Use Peer Config	Priority Map: 0x8	Priority Map: 0x8	
6	10.200.134.138/Card8/Port2 - 20	DCBX-IEEE-Application-TLV-3	Use Peer Config	Priority Map: 0x8	Priority Map: 0x8	
7	10.200.134.138/Card8/Port3 - 22	DCBX-IEEE-PG-TLV-17	Use Peer Config	PGID Map = 0 0 0 1 0 0 0 0	PGID Map = 0 0 0 1 0 0 0 0	
8	10.200.134.138/Card8/Port3 - 23	DCBX-IEEE-PFC-TLV-4	Use Peer Config	Priority Map: 0x8	Priority Map: 0x8	
9	10.200.134.138/Card8/Port3 - 24	DCBX-IEEE-Application-TLV-4	Use Peer Config	Priority Map: 0x8	Priority Map: 0x8	
10	10.200.134.138/Card8/Port4 - 26	DCBX-IEEE-PG-TLV-19	Use Peer Config	PGID Map = 0 0 0 1 0 0 0 0	PGID Map = 0 0 0 1 0 0 0 0	
11	10.200.134.138/Card8/Port4 - 27	DCBX-IEEE-PFC-TLV-5	Use Peer Config	Priority Map: 0x8	Priority Map: 0x8	
12	10.200.134.138/Card8/Port4 - 28	DCBX-IEEE-Application-TLV-5	Use Peer Config	Priority Map: 0x8	Priority Map: 0x8	

Figure 149. FCoE to FC Traffic Forwarding Test with End-to-End Flow Control (PFC and R_RDY) – DCBX per TLV statistics view

Right-click **FCoE Client** and select **DrillDown Per Session** to retrieve detailed information about each VN_Port. The **Interface Status** for each VN_Port should be **PLOGI Complete**, because PLOGI is the last major operation performed by the FCoE ports.

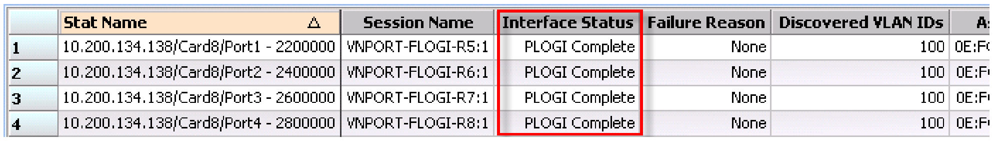

	Stat Name	Session Name	Interface Status	Failure Reason	Discovered VLAN IDs	A:
1	10.200.134.138/Card8/Port1 - 2200000	VNPORT-FLOGI-R5:1	PLOGI Complete	None	100	0E:F
2	10.200.134.138/Card8/Port2 - 2400000	VNPORT-FLOGI-R6:1	PLOGI Complete	None	100	0E:F
3	10.200.134.138/Card8/Port3 - 2600000	VNPORT-FLOGI-R7:1	PLOGI Complete	None	100	0E:F
4	10.200.134.138/Card8/Port4 - 2800000	VNPORT-FLOGI-R8:1	PLOGI Complete	None	100	0E:F

Figure 150. FCoE to FC Traffic Forwarding Test with End-to-End Flow Control (PFC and R_RDY) – FCoE per VN_Port statistics view

Right-click **FC Client** and select **DrillDown Per Session** to retrieve detailed information about each N_Port/NPIV. The **Interface Status** for each N_Port should be **NS-Reg Complete**, because Name Server Registration is the last major operation performed by the FC ports.

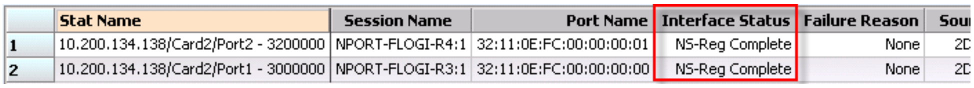

	Stat Name	Session Name	Port Name	Interface Status	Failure Reason	Sou
1	10.200.134.138/Card2/Port2 - 3200000	NPORT-FLOGI-R4:1	32:11:0E:FC:00:00:00:01	NS-Reg Complete	None	2D
2	10.200.134.138/Card2/Port1 - 3000000	NPORT-FLOGI-R3:1	32:11:0E:FC:00:00:00:00	NS-Reg Complete	None	2D

Figure 151. FCoE to FC Traffic Forwarding Test with End-to-End Flow Control (PFC and R_RDY) – FC per N_Port/NPIV statistics view

Part 2

Use the IxNetwork **Statistics – Traffic Item Statistics** view to verify the following real-time FCoE traffic aggregate statistics and native Ethernet traffic aggregate statistics.

- **Rx Frames**: *equal to **Tx Frames***
- **Frames Delta**: *0*
- **Rx Frame Rate**: *equal to **Tx Frame Rate***

  Note: At very high line rates, it is normal to see a small mismatch between Tx and Rx stats, and therefore frames delta. These are not necessarily actual frame loss. Rather, it represents the packets that are still in transit across the system under test at the instance the statistics are reported.

- **Cut-Through Avg Latency**: *vendor specific*
  Verify that the average latency is within the expected range. Additional latency and jitter measurement options are available under the Traffic – Options settings.

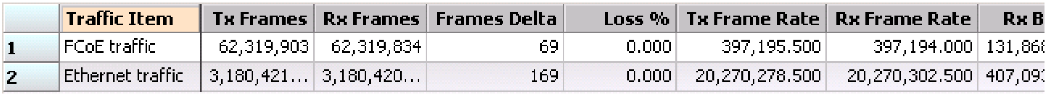

	Traffic Item	Tx Frames	Rx Frames	Frames Delta	Loss %	Tx Frame Rate	Rx Frame Rate	Rx B
1	FCoE traffic	62,319,903	62,319,834	69	0.000	397,195.500	397,194.000	131,86
2	Ethernet traffic	3,180,421...	3,180,420...	169	0.000	20,270,278.500	20,270,302.500	407,09

Figure 152. FCoE to FC Traffic Forwarding Test with End-to-End Flow Control (PFC and R_RDY) – traffic item statistics

From the **Traffic Item Statistics** view, **FCoE Traffic** item, use the IxNetwork **DrillDown per Ethernet II: PFC Queue** view to verify traffic is received properly for each PFC Queue in the FCoE traffic.

- **Number of unique PFC Queues**: 1
- **Rx Frames:** *equal to **Tx Frames** for each PFC Queue*
- **Frames Delta:** *0*
- **Rx Frame Rate:** *equal to **Tx Frame Rate***

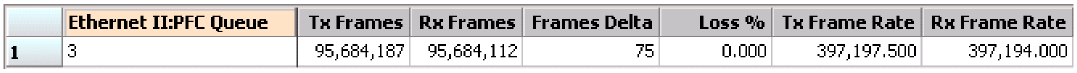

	Ethernet II:PFC Queue	Tx Frames	Rx Frames	Frames Delta	Loss %	Tx Frame Rate	Rx Frame Rate
1	3	95,684,187	95,684,112	75	0.000	397,197.500	397,194.000

Figure 153. FCoE to FC Traffic Forwarding Test with End-to-End Flow Control (PFC and R_RDY) - FCoE PFC Queues statistics

From the **FCoE Item Statistics** view, **FCoE Traffic** item, use the IxNetwork **DrillDown per Rx Port** view to verify that FCoE traffic is received properly by the fibre channel ports only.

- **Number of Rx Ports**: 2
- **Rx Frames**: equal to **Tx Frames**
- **Frames Delta:** 0
- **Rx Frame Rate:** equal to **Tx Frame Rate**

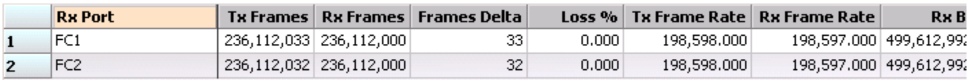

	Rx Port	Tx Frames	Rx Frames	Frames Delta	Loss %	Tx Frame Rate	Rx Frame Rate	Rx B
1	FC1	236,112,033	236,112,000	33	0.000	198,598.000	198,597.000	499,612,99
2	FC2	236,112,032	236,112,000	32	0.000	198,598.000	198,597.000	499,612,99

Figure 154. FCoE to FC Traffic Forwarding Test with End-to-End Flow Control (PFC and R_RDY) – FCoE per Rx Port statistics

From the **Traffic Item Statistics** view, **Eth Traffic** item, use the IxNetwork **DrillDown per Ethernet II: PFC Queue** view to verify that traffic is received properly for each PFC Queue in the native Ethernet traffic.

- **Number of unique PFC Queues**: *1*
- **Rx Frames:** *equal to* ***Tx Frames*** *for each PFC Queue*
- **Frames Delta:** *0*
- **Rx Frame Rate:** *equal to* ***Tx Frame Rate***

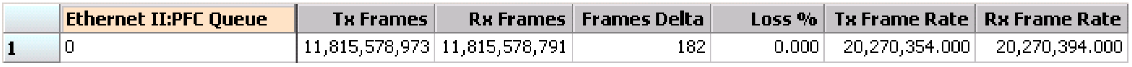

	Ethernet II:PFC Queue	Tx Frames	Rx Frames	Frames Delta	Loss %	Tx Frame Rate	Rx Frame Rate
1	0	11,815,578,973	11,815,578,791	182	0.000	20,270,354.000	20,270,394.000

Figure 155. FCoE to FC Traffic Forwarding Test with End-to-End Flow Control (PFC and R_RDY) - native Ethernet PFC Queues statistics

From the **Traffic Item Statistics** view, **Eth Traffic** item, use the IxNetwork **DrillDown per Rx Ports** view to verify that traffic is received properly by the FCoE ports only.

- **Number of unique Rx Ports**: *1*
- **Rx Frames:** *equal to* ***Tx Frames***
- **Frames Delta:** *0*
- **Rx Frame Rate:** *equal to* ***Tx Frame Rate***

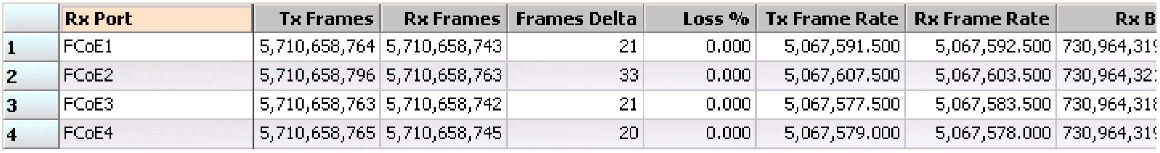

	Rx Port	Tx Frames	Rx Frames	Frames Delta	Loss %	Tx Frame Rate	Rx Frame Rate	Rx B
1	FCoE1	5,710,658,764	5,710,658,743	21	0.000	5,067,591.500	5,067,592.500	730,964,31
2	FCoE2	5,710,658,796	5,710,658,763	33	0.000	5,067,607.500	5,067,603.500	730,964,32
3	FCoE3	5,710,658,763	5,710,658,742	21	0.000	5,067,577.500	5,067,583.500	730,964,31
4	FCoE4	5,710,658,765	5,710,658,745	20	0.000	5,067,579.000	5,067,578.000	730,964,31

Figure 156. FCoE to FC Traffic Forwarding Test with End-to-End Flow Control (PFC and R_RDY) – native Ethernet per Rx Port statistics

From the **Port Statistics** view, verify that the FC ports are transmitting the same amount of R_RDYs, and that they are not receiving any R_RDYs.

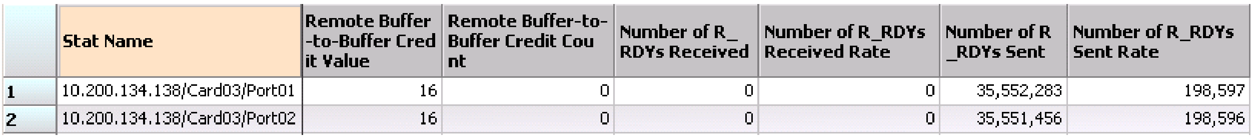

	Stat Name	Remote Buffer-to-Buffer Credit Value	Remote Buffer-to-Buffer Credit Count	Number of R_RDYs Received	Number of R_RDYs Received Rate	Number of R_RDYs Sent	Number of R_RDYs Sent Rate
1	10.200.134.138/Card03/Port01	16	0	0	0	35,552,283	198,597
2	10.200.134.138/Card03/Port02	16	0	0	0	35,551,456	198,596

Figure 157. FCoE to FC Traffic Forwarding Test with End-to-End Flow Control (PFC and R_RDY) – FC R_RDY statistics

From the **Port Statistics** view, verify that the FCoE ports are transmitting receiving PFC PAUSE frames only on Priority value 3 (that is, PFC Queue 3).

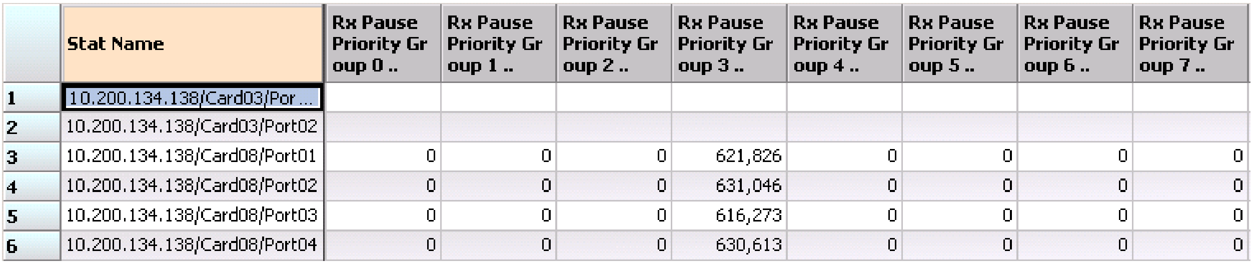

	Stat Name	Rx Pause Priority Group 0 ..	Rx Pause Priority Group 1 ..	Rx Pause Priority Group 2 ..	Rx Pause Priority Group 3 ..	Rx Pause Priority Group 4 ..	Rx Pause Priority Group 5 ..	Rx Pause Priority Group 6 ..	Rx Pause Priority Group 7 ..
1	10.200.134.138/Card03/Por ...								
2	10.200.134.138/Card03/Port02								
3	10.200.134.138/Card08/Port01	0	0	0	621,826	0	0	0	0
4	10.200.134.138/Card08/Port02	0	0	0	631,046	0	0	0	0
5	10.200.134.138/Card08/Port03	0	0	0	616,273	0	0	0	0
6	10.200.134.138/Card08/Port04	0	0	0	630,613	0	0	0	0

Figure 158. FCoE to FC Traffic Forwarding Test with End-to-End Flow Control (PFC and R_RDY) – FCoE PFC PAUSE statistics

Conclusions

By validating the traffic statistics using the preceding features, the DUT has been proven capable of forwarding converged traffic on the same port at the configured line rate, in which the fibre channel traffic are transmitted from the FCoE ports to the native fibre channel ports. In addition, the DUT is capable of properly managing the buffer-to-buffer credits on the fibre channel ports and PFC PAUSE on the FCoE ports to prevent traffic loss when the ingress rate of fibre channel traffic is higher than the egress rate.

Test Case: PFC – Testing Mixed Traffic Flow Management - IxExplorer

Overview

One of the major components of converged data center for lossless behavior is the priority-based flow control (PFC), as specified in IEEE 802.1Qbb. DCE capable switches must have the ability to forward different classes of traffic and control traffic flow based on the priority or the class of the packets forwarded. For applications such as storage, the traffic needs to be forwarded in a lossless manner, whereas most data applications can tolerate a lossy transport such as TCP or UDP where they retransmit or drop packets, respectively. To provide guaranteed bandwidth for lossless traffic in case of congestion, lossy traffic must be throttled back and managed.

This test described in this section is designed to be used with DCE capable switches and is designed to measure how PFC manages prioritized traffic. It also determines whether pause control operates on a per flow basis and does not affect the rate of the traffic that has not been paused. Priority-based flow control handles mixed traffic of different types over Ethernet while ensuring the lossless behavior for critical applications, such as data storage.

Objective

FCoE traffic should exhibit lossless behavior while lower priority IP traffic should be pause controlled by the DUT.

Setup

Create a traffic profile (as shown in the following figure) on four 10GE test ports, where the traffic destined for port 2 will create congestion (receiving 12 Gbps). Configure the DUT to guaranty lossless flow for the 8Gbps FCoE traffic set with priority of 1. Set up the rates to initially avoid any congestion and then gradually increase to the required rates as suggested for the traffic profile to observe PFC functionality.

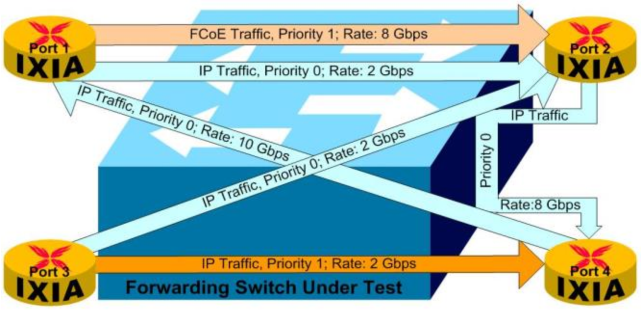
Port 1
IXIA
FCoE Traffic, Priority 1; Rate: 8 Gbps
Port 2
IXIA
IP Traffic, Priority 0; Rate: 2 Gbps
IP Traffic, Priority 0; Rate: 10 Gbps
IP Traffic
IP Traffic, Priority 0; Rate: 2 Gbps
Priority 0
Rate:8 Gbps
Port 3
IXIA
IP Traffic, Priority 1; Rate: 2 Gbps
Port 4
IXIA
Forwarding Switch Under Test

Figure 159. Test traffic configuration

Step-by-Step Instructions

Setting up the DUT

Configure the switch for priority-based flow control by giving the higher priority flows higher precedence in forwarding. In this case, priority 1 will have precedence in terms of traffic flow over the lower priority 0.

Setting up the test ports (Follow the steps in Appendix: ’PFC Getting Started Guide: Testing Mixed Traffic Flow Management’ before applying the steps listed here.)

1. Set up on all of the ports:
   a. Enable **Data Center Mode**, select **PFC (802.1Qbb)** and assign **priority group 0** to **priority 0** and **priority group 1** to **priority 1**.
   b. Set the receive mode to **packet group** with **auto-instrumentation** selected.
2. Set up streams:
   a. Port 1:
      i. Stream1:
         1. Frame Size: **1500** bytes
         2. Rate: **80%** of line rate
         3. Protocol: **FCoE**
         4. Priority Group: **1**
         5. **Dest. MAC Address**: match the **Source Address** of the target port (Test Port 2)
         6. Enable **Packet Group** and **Auto-instrumentation**
      ii. Stream2:
         1. Frame Size: **1500** bytes
         2. Rate: **20%** of line rate
         3. Protocol: **IP**
         4. Priority Group: **0**
         5. **Dest. MAC Address**: match the **Source Address** of the target port (Test Port 2)
         6. Enable **Packet Group** and **Auto-instrumentation**

b. Port 2:

  i. Stream1:

    1. Frame Size: **1500** bytes
    2. Rate: **80%** of line rate
    3. Protocol: **IP**
    4. Priority Group: **0**
    5. **Dest. MAC Address**: match the Source Address of the target port (Test Port 4)
    6. Enable **Packet Group** and **Auto-instrumentation**

c. Port 3:

  i. Stream1:

    1. Frame Size: **1500** bytes
    2. Rate: **20%** of line rate
    3. Protocol: **IP**
    4. Priority Group: **0**
    5. **Dest. MAC Address**: match the Source Address of the target port (Test Port 2)
    6. Enable **Packet Group** and **Auto-instrumentation**

  ii. Stream2:

    1. Frame Size: **1500** bytes
    2. Rate: **20%** of line rate
    3. Protocol: **IP**
    4. Priority Group: **1**
    5. **Dest. MAC Address**: match the Source Address of the target port (Test Port 4)
    6. Enable **Packet Group** and **Auto-instrumentation**

d. Port 4:

  i. Stream1:

    1. Frame Size: **1500** bytes
    2. Rate: **100%** of line rate
    3. Protocol: **IP**
    4. Priority Group: **0**
    5. **Dest. MAC Address**: match the Source Address of the target port (Test Port 1)
    6. Enable **Packet Group** and **Auto-instrumentation**

3. Create **Packet Group Statistic** views to monitor the received bit rate on each flow of each port.
4. Begin testing by **starting transmit** on all of the ports and **monitoring the bit rate** of all flows.

Test Variables

Parameter	Description
Transmit rate	• Reduce the total rate transmitted by test port 1 to 50% (40% and 10% for stream 1 and 2, respectively) of the line rate initially and observe no rate loss because there is no congestion. • Then, increase the total rate to 100% (80% and 20% for stream 1 and 2, respectively) to observe PFC functionality of the DUT.
Priorities and Priority Groups	• The number of priorities in the traffic profile and the DUT can be scaled up to test the extended capabilities of the DUT.

Test Tool Variables

Parameter	Description
Transmit rate	• Start the flow rates from ½ of the final value to avoid congestion and then increase it to observer the behavior when congestion occurs.
Priorities and Priority Groups	• The number of priorities in the traffic profile and the DUT can be scaled up to test the extended capabilities of the DUT.

DUT Test Variables

Parameter	Description
Priorities and Priority Groups	• The number of priorities in the traffic profile and the DUT can be scaled up to test the extended capabilities of the DUT.

Results Analysis

Using **Packet Group Statistic** views, the bit rate for different flows can be examined to determine that there has been no rate drop (no pause control) when there is no congestion and a rate drop (pause control) for the lower priority flows when there is congestion.

Even when pause control is in effect, port 2 should still receive about 10 Gbps. The following figure shows how the forwarding switch fabric reacts when Ixia test port 2 is congested with 12 Gbps of traffic load. This traffic is composed of 8 Gbps of FCoE and 4 Gbps of Ethernet. Only the Ethernet traffic with priority of zero is paused by the switch fabric. After Ixia port 1 and port 3 receive PFC pause frames from DUT for the lower priority (priority 0) traffic destined for the congested port, standard Ethernet traffic is throttled down from the offered 2Gbps to 1Gbps. FCoE traffic with a priority of 1 on Ixia test port 1 is allowed to flow by the switch fabric at the throughput target rate of 8 Gbps to Ixia test port 2. Because there is no other congestion created and thus no further PFC pause control needed, ports 1, 2 and 3 all end up receiving full line rate 10 Gbps of traffic.

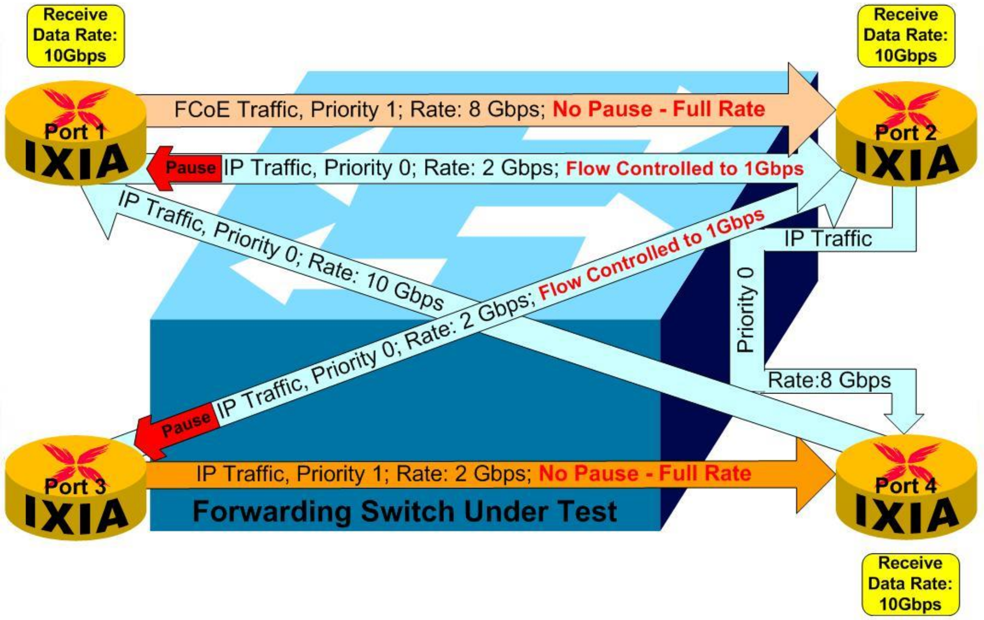
Receive Data Rate: 10Gbps
Receive Data Rate: 10Gbps
Port 1
IXIA
Port 2
IXIA
FCoE Traffic, Priority 1; Rate: 8 Gbps; No Pause - Full Rate
Pause
IP Traffic, Priority 0; Rate: 2 Gbps; Flow Controlled to 1Gbps
IP Traffic, Priority 0; Rate: 10 Gbps
IP Traffic, Priority 0; Rate: 2 Gbps; Flow Controlled to 1Gbps
Pause
IP Traffic
Priority 0
Rate:8 Gbps
Port 3
IXIA
IP Traffic, Priority 1; Rate: 2 Gbps; No Pause - Full Rate
Forwarding Switch Under Test
Port 4
IXIA
Receive Data Rate: 10Gbps

Figure 160. Priority-based Flow Control guarantees the FCoE throughput @ 8Gbps

Troubleshooting and Diagnostics

Issue	Troubleshooting Solution
Loss of throughput on FCoE traffic	Check DUT settings to ensure: • PFC mode is enabled • Correct mapping of PFC priorities Check test ports to ensure: • Correct priority mapping • Correct source and destination address configuration of streams
Loss of throughput on non-congested ports	Check DUT settings to ensure: • PFC mode is enabled instead of normal pause control Check test ports to ensure: • Correct priority mapping • Correct source and destination address configuration of streams
Dropped frames	Check DUT settings to ensure: • PFC mode is enabled Check test ports to ensure: • Correct source and destination address configuration of streams
Receive data rate below 10 Gbps	Account for packet overhead and Check DUT settings to ensure: • PFC mode is enabled Check test ports to ensure: • Correct priority mapping • Correct source and destination address configuration of streams

Conclusions

This test determined that DCE capable switches use PFC to manage prioritized traffic. It also determines whether pause control operates on a per flow basis and does not affect the rate of the traffic that has not been paused. Priority-based flow control handles mixed traffic of different types over Ethernet while ensuring the lossless behavior for critical applications, such as data storage.

PFC Getting Started Guide: Testing Mixed Traffic Flow Management - IxExplorer

1. Traffic generation and receiving test ports need to be set to **Data Center Mode** in the Port Properties dialog to take advantage of DCE features such as PFC.

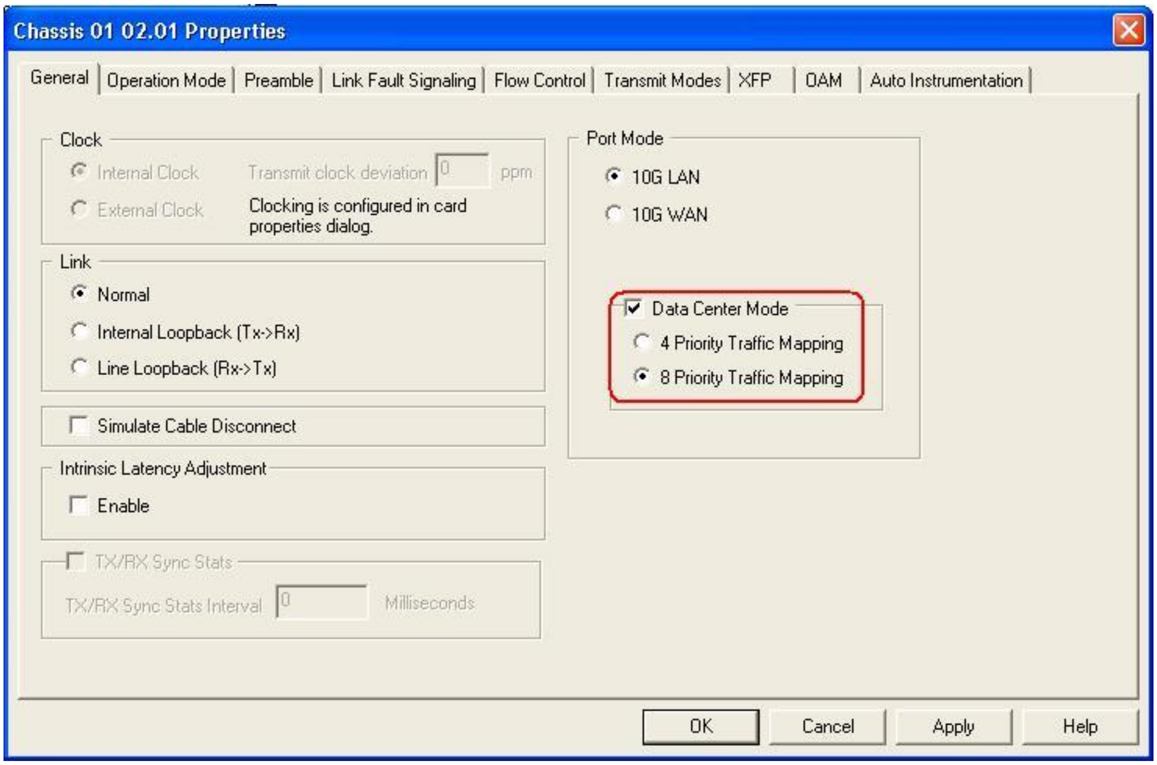

Chassis 01 02.01 Properties

Clock
● Internal Clock Transmit clock deviation 0 ppm
○ External Clock
Clocking is configured in card properties dialog.

Link
● Normal
○ Internal Loopback (Tx->Rx)
○ Line Loopback (Rx->Tx)

□ Simulate Cable Disconnect

Intrinsic Latency Adjustment
□ Enable

□ TX/RX Sync Stats
TX/RX Sync Stats Interval 0 Milliseconds

Port Mode
● 10G LAN
○ 10G WAN

☑ Data Center Mode
○ 4 Priority Traffic Mapping
● 8 Priority Traffic Mapping

Figure 161. Setting the ports to Data Center Mode

2. Priority groups need to be mapped into the priorities. Priority groups are used internally to allow assigning the same priority to several streams, flows or traffic types.

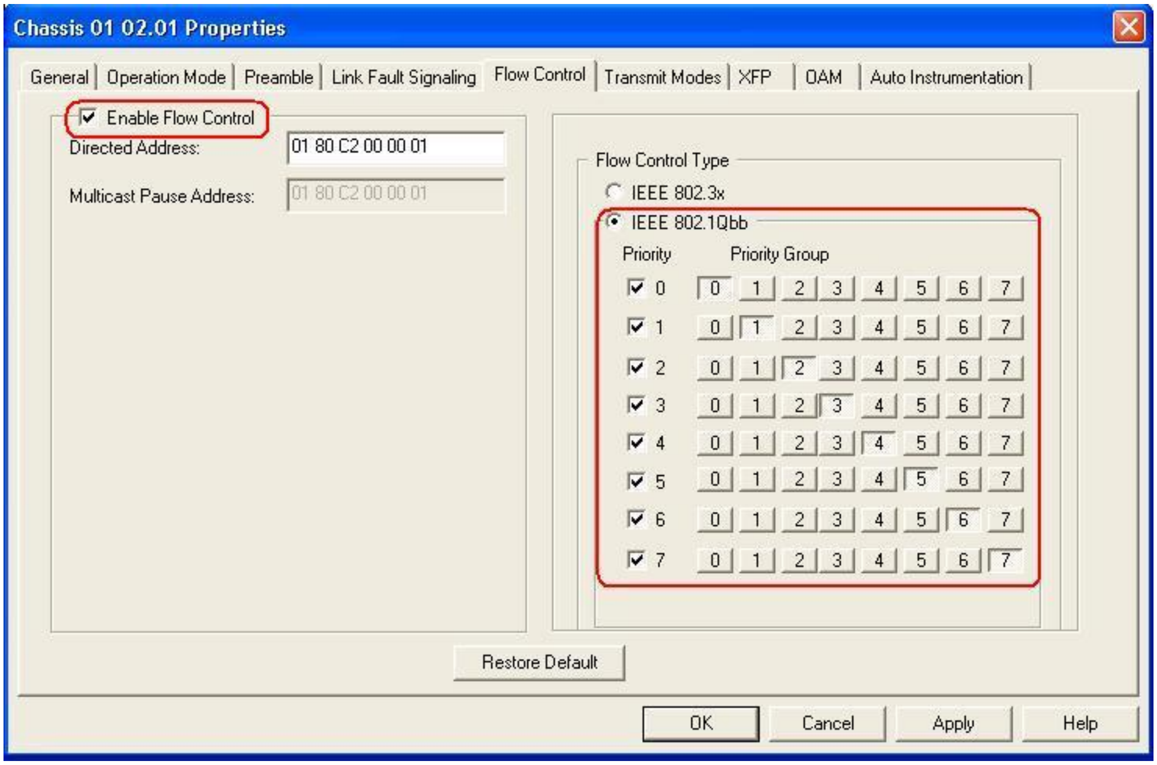

☑ Enable Flow Control

Directed Address: 01 80 C2 00 00 01

Multicast Pause Address: 01 80 C2 00 00 01

Flow Control Type

○ IEEE 802.3x

● IEEE 802.1Qbb

Priority	Priority Group
☑ 0	0
☑ 1	1
☑ 2	2
☑ 3	3
☑ 4	4
☑ 5	5
☑ 6	6
☑ 7	7

Figure 162. Mapping priorities to Priority Groups on each port

3. Streams need to be setup with the appropriate priority group for this PFC test. The transmit mode is automatically set to **Advanced Scheduler** when in DCE mode. **Error! Reference ource not found.** shows two streams that create a mixed traffic profile (Ethernet and FCoE).

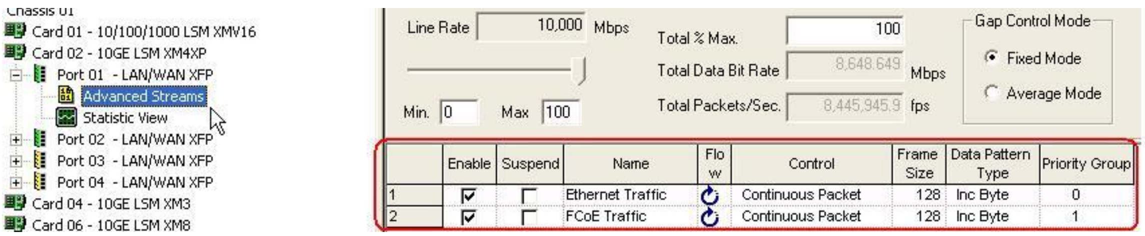

Line Rate 10,000 Mbps

Min. 0 Max 100

Total % Max. 100

Total Data Bit Rate 8,648,649 Mbps

Total Packets/Sec. 8,445,945.9 fps

Gap Control Mode

● Fixed Mode

○ Average Mode

	Enable	Suspend	Name	Flow	Control	Frame Size	Data Pattern Type	Priority Group
1	☑	□	Ethernet Traffic	↻	Continuous Packet	128	Inc Byte	0
2	☑	□	FCoE Traffic	↻	Continuous Packet	128	Inc Byte	1

Figure 163. Create two streams

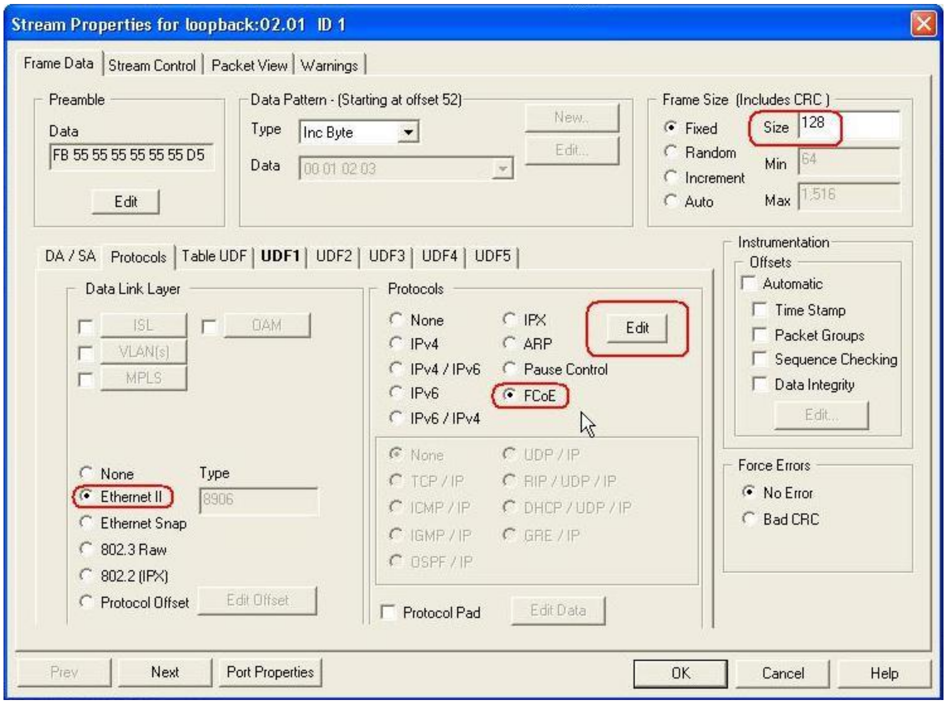

Preamble
Data: FB 55 55 55 55 55 55 D5

Data Pattern - (Starting at offset 52)
Type: Inc Byte
Data: 00 01 02 03

Frame Size (Includes CRC)
● Fixed ○ Random ○ Increment ○ Auto
Size: 128
Min: 64
Max: 1,516

Data Link Layer
○ None
● Ethernet II
○ Ethernet Snap
○ 802.3 Raw
○ 802.2 (IPX)
○ Protocol Offset
Type: 8906

Protocols
○ None ○ IPX
○ IPv4 ○ ARP
○ IPv4 / IPv6 ○ Pause Control
○ IPv6 ● FCoE
○ IPv6 / IPv4

● None ○ UDP / IP
○ TCP / IP ○ RIP / UDP / IP
○ ICMP / IP ○ DHCP / UDP / IP
○ IGMP / IP ○ GRE / IP
○ OSPF / IP

□ Protocol Pad

Instrumentation
Offsets
□ Automatic
□ Time Stamp
□ Packet Groups
□ Sequence Checking
□ Data Integrity

Force Errors
● No Error
○ Bad CRC

Figure 164. Edit streams: frame size and protocol

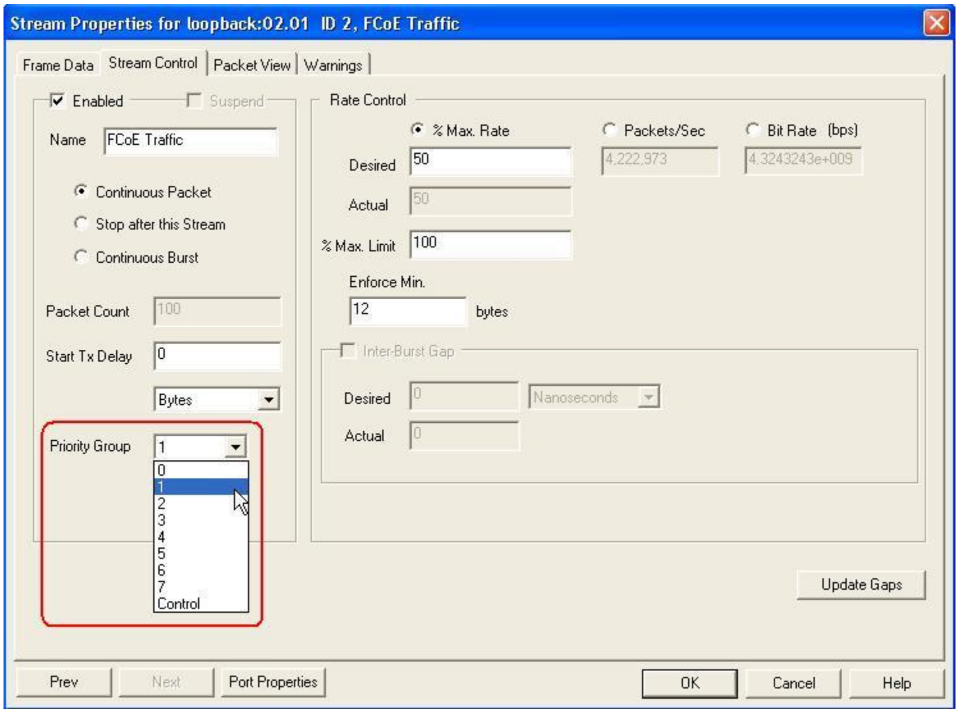

☑ Enabled □ Suspend
Name: FCoE Traffic
● Continuous Packet
○ Stop after this Stream
○ Continuous Burst
Packet Count: 100
Start Tx Delay: 0
Bytes
Priority Group: 1
0
1
2
3
4
5
6
7
Control

Rate Control
● % Max. Rate ○ Packets/Sec ○ Bit Rate (bps)
Desired: 50 4,222,973 4.3243243e+009
Actual: 50
% Max. Limit: 100
Enforce Min.: 12 bytes

□ Inter-Burst Gap
Desired: 0 Nanoseconds
Actual: 0

Figure 165. Set the stream priorities

4. Create the **Packet Group Statistics** view for monitoring all of the flows.

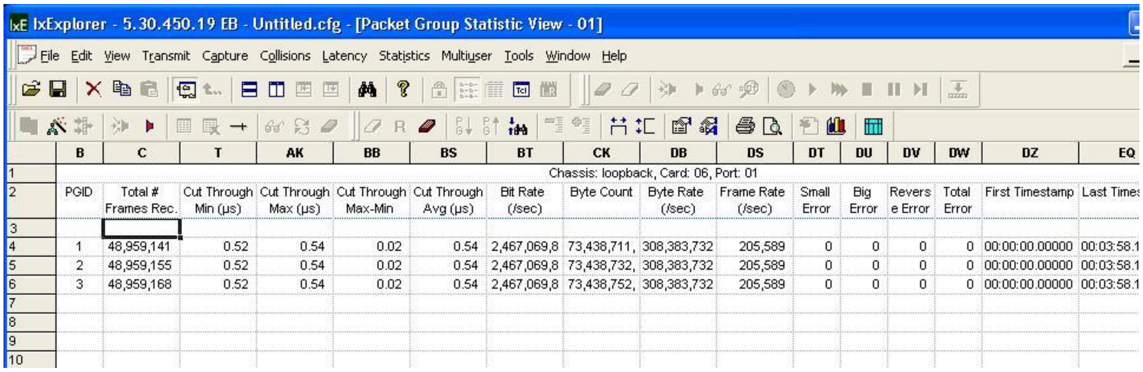

IxExplorer - 5.30.450.19 EB - Untitled.cfg - [Packet Group Statistic View - 01]

	B	C	T	AK	BB	BS	BT	CK	DB	DS	DT	DU	DV	DW	DZ	EQ
1	Chassis: loopback, Card: 06, Port: 01															
2	PGID	Total # Frames Rec.	Cut Through Min (µs)	Cut Through Max (µs)	Cut Through Max-Min	Cut Through Avg (µs)	Bit Rate (/sec)	Byte Count	Byte Rate (/sec)	Frame Rate (/sec)	Small Error	Big Error	Revers e Error	Total Error	First Timestamp	Last Times
3																
4	1	48,959,141	0.52	0.54	0.02	0.54	2,467,069,8	73,438,711,	308,383,732	205,589	0	0	0	0	00:00:00.00000	00:03:58.1
5	2	48,959,155	0.52	0.54	0.02	0.54	2,467,069,8	73,438,732,	308,383,732	205,589	0	0	0	0	00:00:00.00000	00:03:58.1
6	3	48,959,168	0.52	0.54	0.02	0.54	2,467,069,8	73,438,752,	308,383,732	205,589	0	0	0	0	00:00:00.00000	00:03:58.1
7																
8																
9																
10																

Figure 166. Create Packet Group Statistic views on all port

Test Case: PFC– Measuring DUT's PFC Response Time and Quanta Accuracy - IxExplorer

Overview

This test generates a traffic flow from a test port that is forwarded through the DUT to another test port. The receiving test port will send a PFC pause control (PC) packet to the DUT. The DUT behavior will be measured in terms of the response time (the time between receiving the PC packet and the last packet of the flow forwarded before the pause takes effect) and verification of the pause duration, as specified by the quanta in the PFC PC packet.

Objective

To characterize the behavior and accuracy of the DUT when PFC pause control occurs.

Setup

Create a traffic profile, as shown in the following figure, on two 10GE test ports where IP traffic with priority 0 will be forwarded from test port 1 to test port 2 through the DUT. In addition, IP traffic with priority 3 will be forwarded from port 2 to port 1. On port 2, create a packet stream for pause control and another packet stream to be used as a trigger/marker for the capture filter on the same port. Configure the DUT with PFC capability enabled.

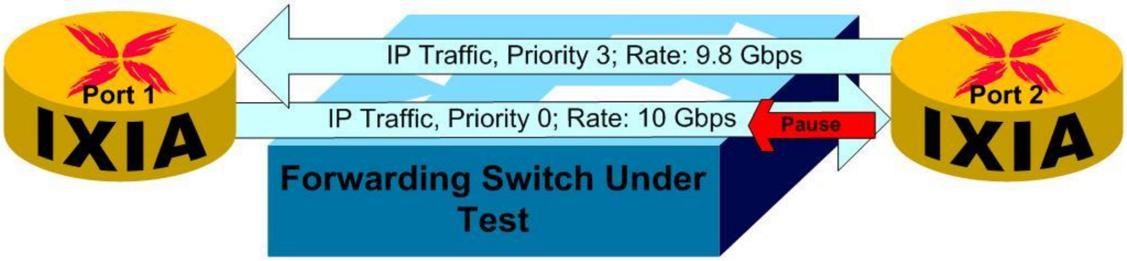
Port 1
IXIA
IP Traffic, Priority 3; Rate: 9.8 Gbps
IP Traffic, Priority 0; Rate: 10 Gbps
Pause
Port 2
IXIA
Forwarding Switch Under Test

Figure 167. Test traffic configuration

Step-by-Step Instructions

(Follow the instructions in Appendix: *PFC Getting Started Guide: Measuring DUT's PFC Response Time and Quanta Accuracy* before using the instructions in this section).

Setting up the DUT

Configure the switch for priority-based flow control.

Setting up the test ports

1. Set up on all of the ports:

    a. Enable **Data Center Mode**, select **PFC (802.1Qbb)** and assign priority **group 0** to **priority 0** and **priority group 3** to **priority 3**.

2. Set up streams:

    a. Port 1:

        i. Stream1:

            1. Frame Size: **70** bytes
            2. Rate: **100%** of line rate
            3. Protocol: **IP**
            4. Priority Group: **0**
            5. **Dest. MAC Address**: match the **Source Address** of the target port (Test Port 2)
            6. Control: **Continuous Packet**

    b. Port 2:

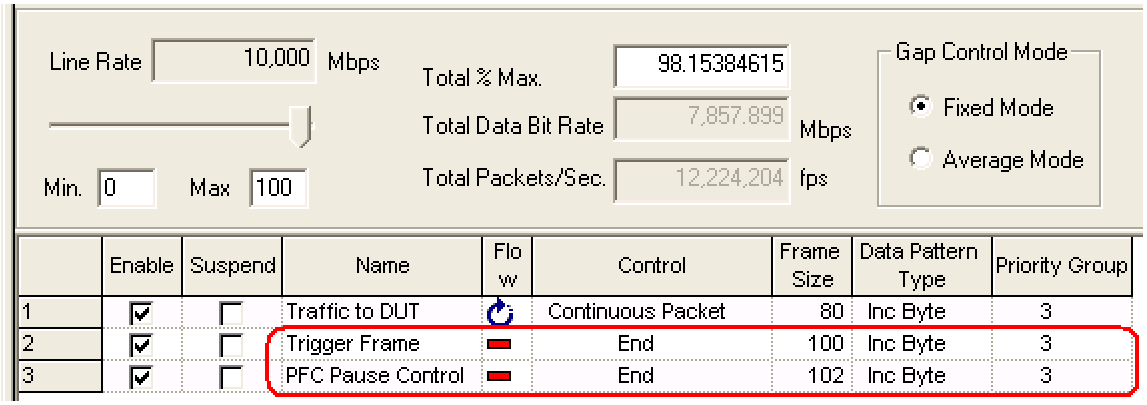

Line Rate 10,000 Mbps

Min. 0 Max 100

Total % Max. 98.15384615

Total Data Bit Rate 7,857.899 Mbps

Total Packets/Sec. 12,224,204 fps

Gap Control Mode

● Fixed Mode

○ Average Mode

	Enable	Suspend	Name	Flow	Control	Frame Size	Data Pattern Type	Priority Group
1	☑	□	Traffic to DUT	↻	Continuous Packet	80	Inc Byte	3
2	☑	□	Trigger Frame	🟥	End	100	Inc Byte	3
3	☑	□	PFC Pause Control	🟥	End	102	Inc Byte	3

Figure 168. Port 2 stream set up

i. Stream1(traffic to DUT):

   1. Frame Size: **80** bytes
   2. Rate: **97%** of line rate
   3. Protocol: **IP**
   4. Priority Group: **3**
   5. **Dest. MAC Address**: match the Source Address of the target port (Test Port 1)
   6. Control: **Continuous Packet**

ii. Stream2 (Trigger Frame):

   1. Frame Size: **100** bytes
   2. Rate: **1%** of line rate
   3. Protocol: **IP**
   4. Priority Group: **3**
   5. **Dest. MAC Address** = Source Address of the port itself. This packet will be looped back (returned) to the port by the DUT and be used as a capture trigger.
   6. Control: **END**
      - Packet count: **1**

iii. Stream3:

   1. Frame Size: **102** bytes
   2. Rate: **1%** of line rate
   3. Protocol: **IP**
      - **Pause Control**
        i. **802.1Qbb**: **Priority 0**; quanta = **40000**
   4. Priority Group: **3**
   5. **Dest. MAC Address**: Matching the Source Address of the target port (Test Port 1)

6. Control: **END**
   - Packet count: **1**

3. Set capture filter trigger to packets size of **100** bytes and the capture filter to packet size of **70** bytes.

4. Begin testing by
   a. **Starting transmit** on port1
   b. **Start capture** on port 2
   c. **Start transmit** on port 2
   d. **Open** and **view** capture for the results

Test Variables

Parameter	Description
Frame size	• Try different frame sizes for the flow from port 1 to port 2 to characterize the response time over different packet sizes.
Pause duration	• Try different pause durations.
Pause Control Vector	• Try multiple/simultaneous pause operations for different priorities.

Test Tool Variables

Parameter	Description
Transmit rate	• Start the flow rates from ½ of the final value and increase it gradually to observer the behavior of the DUT in managing the receive rate.
Pause duration	• Try different pause durations.
Priorities and Priority Groups	• The number of priorities in the traffic profile and the DUT can be scaled up to test the extended capabilities of the DUT.

DUT Test Variables

Parameter	Description
Priorities and Priority Groups	• The number of priorities in the traffic profile and the DUT can be scaled up to test the extended capabilities of the DUT. • Change priority policies (the order of priority) and observe if there are any affects on the response times.

Results Analysis

Using the capture view, the response time can be measured from the time the trigger packet was received and the time of last packet before pause period being observed by the DUT. The trigger packet is one packet ahead of the pause control packet leaving port 2. Its arrival at port 2 will be very close to the time of pause packet leaving the same port. The next packet will have the pause duration value as a relative measurement to the previous packet. All of the measurements are represented in dd:hh:mm:ss format (seconds are reported in decimal fractions).

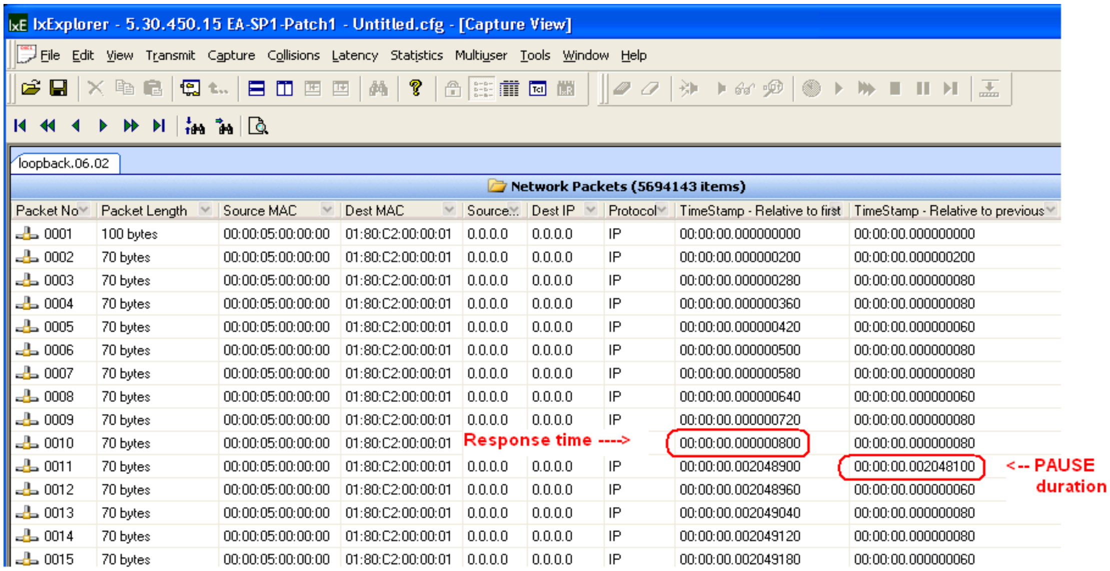

Network Packets (5694143 items)

Packet No	Packet Length	Source MAC	Dest MAC	Source...	Dest IP	Protocol	TimeStamp - Relative to first	TimeStamp - Relative to previous
0001	100 bytes	00:00:05:00:00:00	01:80:C2:00:00:01	0.0.0.0	0.0.0.0	IP	00:00:00.000000000	00:00:00.000000000
0002	70 bytes	00:00:05:00:00:00	01:80:C2:00:00:01	0.0.0.0	0.0.0.0	IP	00:00:00.000000200	00:00:00.000000200
0003	70 bytes	00:00:05:00:00:00	01:80:C2:00:00:01	0.0.0.0	0.0.0.0	IP	00:00:00.000000280	00:00:00.000000080
0004	70 bytes	00:00:05:00:00:00	01:80:C2:00:00:01	0.0.0.0	0.0.0.0	IP	00:00:00.000000360	00:00:00.000000080
0005	70 bytes	00:00:05:00:00:00	01:80:C2:00:00:01	0.0.0.0	0.0.0.0	IP	00:00:00.000000420	00:00:00.000000060
0006	70 bytes	00:00:05:00:00:00	01:80:C2:00:00:01	0.0.0.0	0.0.0.0	IP	00:00:00.000000500	00:00:00.000000080
0007	70 bytes	00:00:05:00:00:00	01:80:C2:00:00:01	0.0.0.0	0.0.0.0	IP	00:00:00.000000580	00:00:00.000000080
0008	70 bytes	00:00:05:00:00:00	01:80:C2:00:00:01	0.0.0.0	0.0.0.0	IP	00:00:00.000000640	00:00:00.000000060
0009	70 bytes	00:00:05:00:00:00	01:80:C2:00:00:01	0.0.0.0	0.0.0.0	IP	00:00:00.000000720	00:00:00.000000080
0010	70 bytes	00:00:05:00:00:00	01:80:C2:00:00:01	Response time --->			00:00:00.000000800	00:00:00.000000080
0011	70 bytes	00:00:05:00:00:00	01:80:C2:00:00:01	0.0.0.0	0.0.0.0	IP	00:00:00.002048900	00:00:00.002048100 <-- PAUSE duration
0012	70 bytes	00:00:05:00:00:00	01:80:C2:00:00:01	0.0.0.0	0.0.0.0	IP	00:00:00.002048960	00:00:00.000000060
0013	70 bytes	00:00:05:00:00:00	01:80:C2:00:00:01	0.0.0.0	0.0.0.0	IP	00:00:00.002049040	00:00:00.000000080
0014	70 bytes	00:00:05:00:00:00	01:80:C2:00:00:01	0.0.0.0	0.0.0.0	IP	00:00:00.002049120	00:00:00.000000080
0015	70 bytes	00:00:05:00:00:00	01:80:C2:00:00:01	0.0.0.0	0.0.0.0	IP	00:00:00.002049180	00:00:00.000000060

Figure 169. Capture timestamp values can be viewed as "relative to first" or as "relative to previous" packet

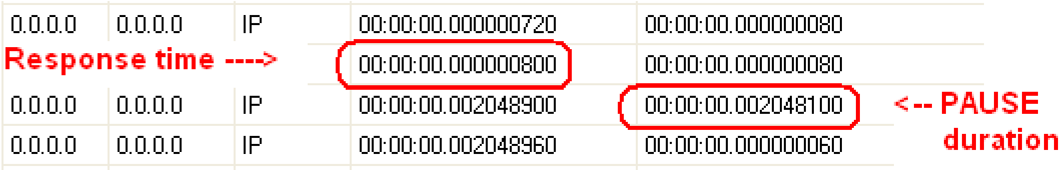

0.0.0.0	0.0.0.0	IP	00:00:00.000000720	00:00:00.000000080
Response time --->			00:00:00.000000800	00:00:00.000000080
0.0.0.0	0.0.0.0	IP	00:00:00.002048900	00:00:00.002048100 <-- PAUSE duration
0.0.0.0	0.0.0.0	IP	00:00:00.002048960	00:00:00.000000060

Figure 170. Response time and pause duration can be measured in the capture view

Troubleshooting and Diagnostics

Issue	Troubleshooting Solution
No traffic flow arriving at port 2 or port 1 (via DUT)	Check DUT settings to ensure: • Traffic ports forwarding is enabled • PFC mode is enabled • Correct mapping of PFC priorities Check test ports to ensure: • Correct source and destination address configuration of streams • Correct priority mapping
Traffic does not resume after pause quanta expires	Check DUT settings to: • Ensure PFC mode is enabled • Check stats on ingress and egress ports for packets received and forwarded Check test ports to ensure: • Check stream settings for pause control quanta value, packet stream control (END) and number of packets to be sent
Unexpected pause duration	Check DUT settings to ensure: • Check stats on ingress and egress ports for packets received and forwarded to make sure of non-congested situation Check test ports to ensure: • Check stream settings for pause control quanta value, packet stream control (END) and number of packets to be sent

Conclusions

This exercise shows how to measure the response time and pause duration time for PFC pause operation. The response time is the time it takes for the receiving port (the port that receives the PC) to stop sending packets from the flow which was paused. The duration of the pause is the time specified in the pause quanta sent as a parameter. The traffic flow is resumed for the paused flow once the pause quantum has expired.

PFC Getting Started Guide: Measuring DUT’s PFC Response Time and Quanta Accuracy - IxExplorer

1. Traffic generation and receiving test ports need to be set to **Data Center Mode** in the port properties dialog to take advantage of DCE features such as PFC.

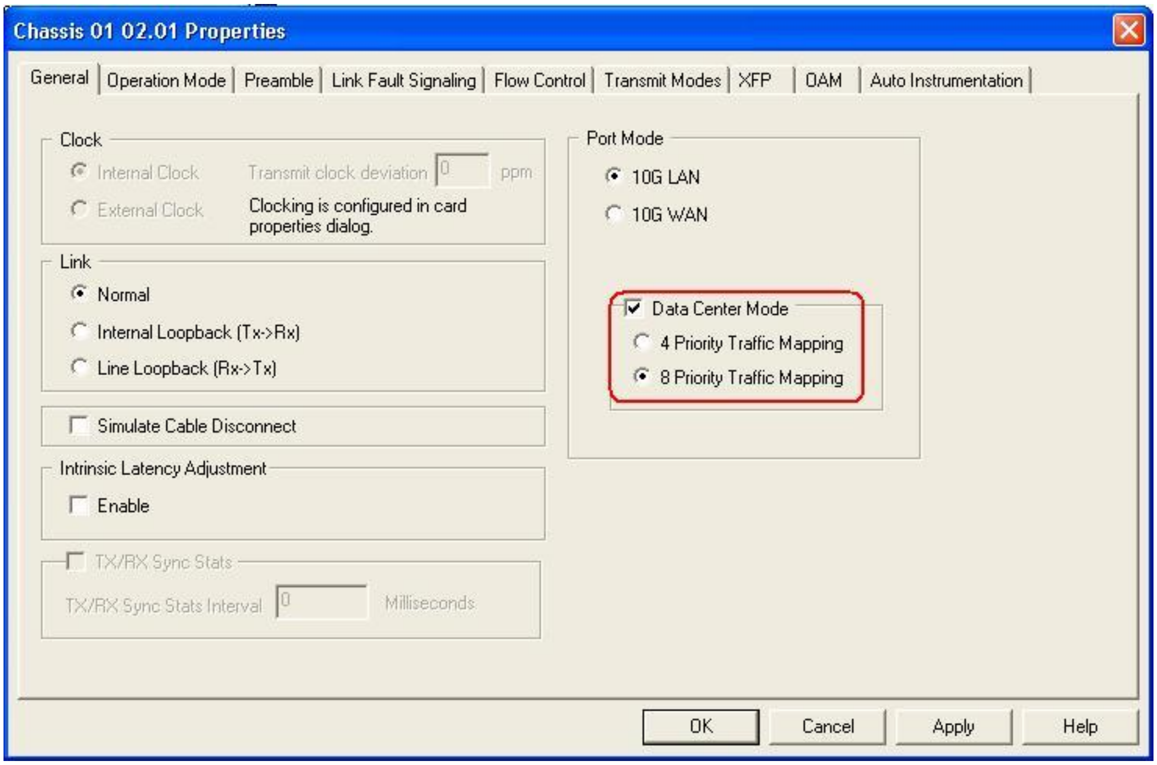

Clock: ● Internal Clock ○ External Clock; Transmit clock deviation 0 ppm; Clocking is configured in card properties dialog.

Link: ● Normal ○ Internal Loopback (Tx->Rx) ○ Line Loopback (Rx->Tx)

□ Simulate Cable Disconnect

Intrinsic Latency Adjustment: □ Enable

□ TX/RX Sync Stats; TX/RX Sync Stats Interval 0 Milliseconds

Port Mode: ● 10G LAN ○ 10G WAN

☑ Data Center Mode: ○ 4 Priority Traffic Mapping ● 8 Priority Traffic Mapping

Figure 171. Setting the ports to Data Center Mode

2. Priority groups need to be mapped into the priorities. Priority groups are used internally to allow assigning the same priority to several streams, flows or traffic types.

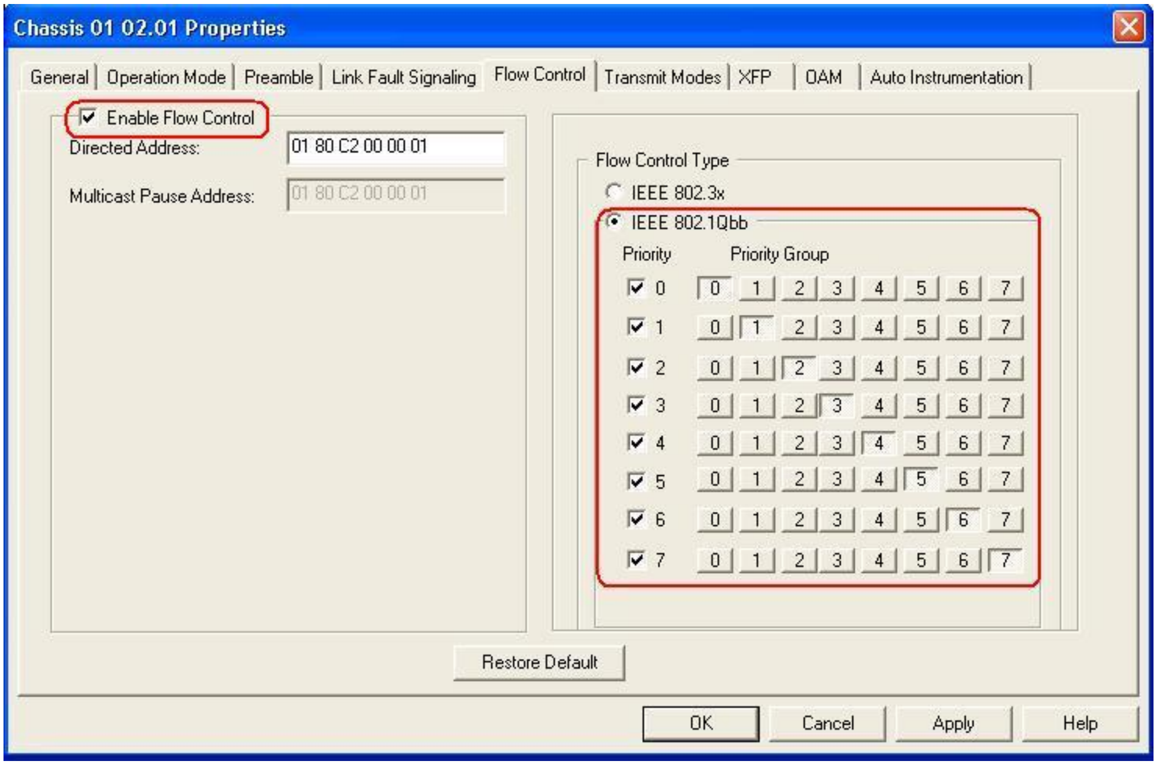

Chassis 01 02.01 Properties

☑ Enable Flow Control

Directed Address: 01 80 C2 00 00 01

Multicast Pause Address: 01 80 C2 00 00 01

Flow Control Type

○ IEEE 802.3x

● IEEE 802.1Qbb

Priority	Priority Group
☑ 0	0
☑ 1	1
☑ 2	2
☑ 3	3
☑ 4	4
☑ 5	5
☑ 6	6
☑ 7	7

Figure 172. Mapping priorities to Priority Groups on each port

3. Streams need to be setup with the appropriate priority group for this PFC test. The transmit mode is automatically set to **Advanced Scheduler** when in DCE mode. **Error! Reference ource not found.** Shows two streams that create a mixed traffic profile (Ethernet and FCoE).

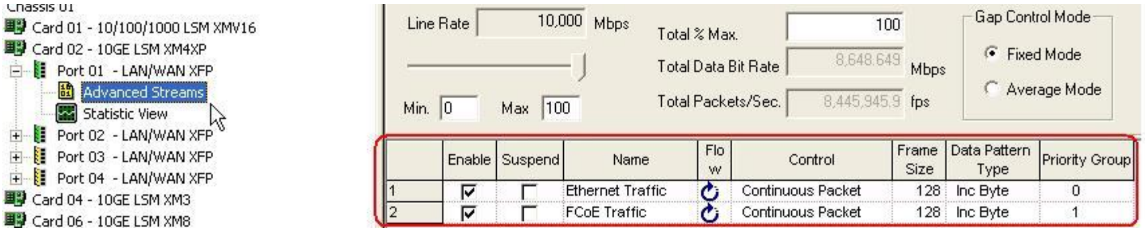

Chassis 01
- Card 01 - 10/100/1000 LSM XMV16
- Card 02 - 10GE LSM XM4XP
  - Port 01 - LAN/WAN XFP
    - Advanced Streams
    - Statistic View
  - Port 02 - LAN/WAN XFP
  - Port 03 - LAN/WAN XFP
  - Port 04 - LAN/WAN XFP
- Card 04 - 10GE LSM XM3
- Card 06 - 10GE LSM XM8

Line Rate 10,000 Mbps

Min. 0 Max 100

Total % Max. 100

Total Data Bit Rate 8,648.649 Mbps

Total Packets/Sec. 8,445,945.9 fps

Gap Control Mode

● Fixed Mode

○ Average Mode

	Enable	Suspend	Name	Flow	Control	Frame Size	Data Pattern Type	Priority Group
1	☑	□	Ethernet Traffic		Continuous Packet	128	Inc Byte	0
2	☑	□	FCoE Traffic		Continuous Packet	128	Inc Byte	1

Figure 173. Create two streams

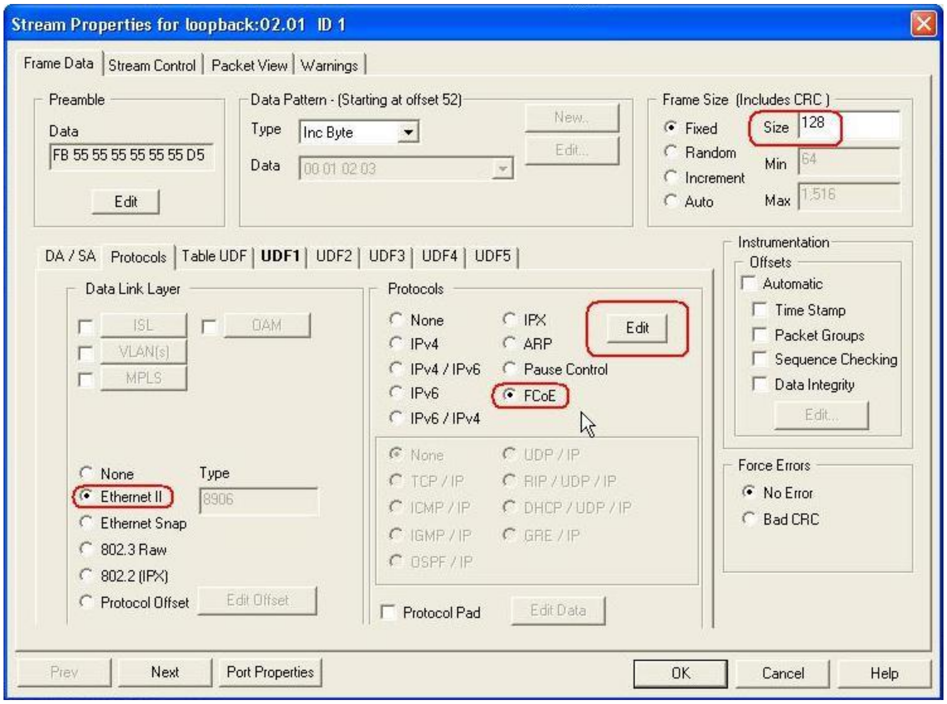

Stream Properties for loopback:02.01 ID 1

Frame Data | Stream Control | Packet View | Warnings

Preamble
Data: FB 55 55 55 55 55 55 D5

Data Pattern - (Starting at offset 52)
Type: Inc Byte
Data: 00 01 02 03

Frame Size (Includes CRC)
● Fixed ○ Random ○ Increment ○ Auto
Size: 128
Min: 64
Max: 1,516

DA / SA | Protocols | Table UDF | UDF1 | UDF2 | UDF3 | UDF4 | UDF5

Data Link Layer
□ ISL □ OAM
□ VLAN(s)
□ MPLS
○ None ● Ethernet II ○ Ethernet Snap ○ 802.3 Raw ○ 802.2 (IPX) ○ Protocol Offset
Type: 8906

Protocols
○ None ○ IPv4 ○ IPv4 / IPv6 ○ IPv6 ○ IPv6 / IPv4 ○ IPX ○ ARP ○ Pause Control ● FCoE
● None ○ TCP / IP ○ ICMP / IP ○ IGMP / IP ○ OSPF / IP ○ UDP / IP ○ RIP / UDP / IP ○ DHCP / UDP / IP ○ GRE / IP
□ Protocol Pad

Instrumentation
Offsets
□ Automatic
□ Time Stamp
□ Packet Groups
□ Sequence Checking
□ Data Integrity

Force Errors
● No Error ○ Bad CRC

Figure 174. Edit streams: frame size and protocol

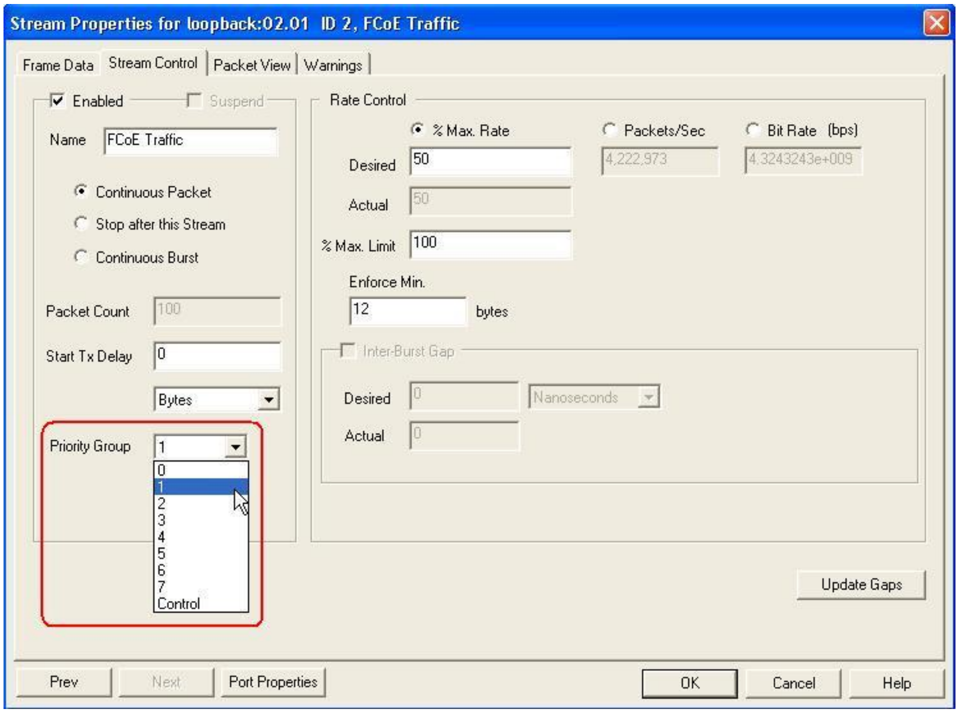

Stream Properties for loopback:02.01 ID 2, FCoE Traffic

Frame Data | Stream Control | Packet View | Warnings

☑ Enabled □ Suspend
Name: FCoE Traffic
● Continuous Packet
○ Stop after this Stream
○ Continuous Burst
Packet Count: 100
Start Tx Delay: 0
Bytes
Priority Group: 1
0
1
2
3
4
5
6
7
Control

Rate Control
● % Max. Rate ○ Packets/Sec ○ Bit Rate (bps)
Desired: 50 | 4,222,973 | 4.3243243e+009
Actual: 50
% Max. Limit: 100
Enforce Min.: 12 bytes
□ Inter-Burst Gap
Desired: 0 Nanoseconds
Actual: 0

Figure 175. Set the stream priorities

4. Create a PFC **Pause Control** packet.

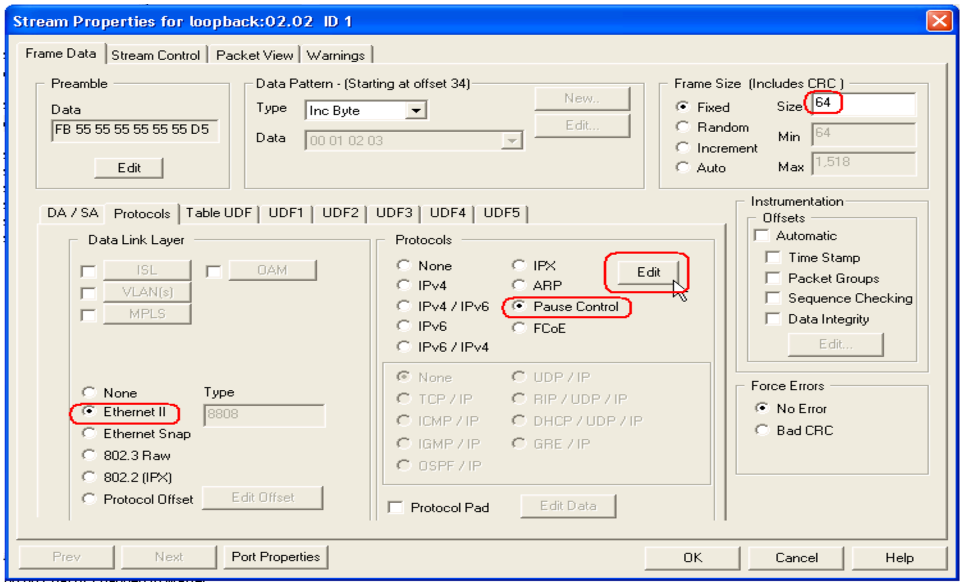

Stream Properties for loopback:02.02 ID 1

Frame Data | Stream Control | Packet View | Warnings

Preamble
Data: FB 55 55 55 55 55 55 D5

Data Pattern - (Starting at offset 34)
Type: Inc Byte
Data: 00 01 02 03

Frame Size (Includes CRC)
● Fixed ○ Random ○ Increment ○ Auto
Size: 64
Min: 64
Max: 1,518

DA / SA | Protocols | Table UDF | UDF1 | UDF2 | UDF3 | UDF4 | UDF5

Data Link Layer
□ ISL □ VLAN(s) □ MPLS □ OAM
○ None ● Ethernet II ○ Ethernet Snap ○ 802.3 Raw ○ 802.2 (IPX) ○ Protocol Offset
Type: 8808

Protocols
○ None ○ IPv4 ○ IPv4 / IPv6 ○ IPv6 ○ IPv6 / IPv4 ○ IPX ○ ARP ● Pause Control ○ FCoE
● None ○ TCP / IP ○ ICMP / IP ○ IGMP / IP ○ OSPF / IP ○ UDP / IP ○ RIP / UDP / IP ○ DHCP / UDP / IP ○ GRE / IP
□ Protocol Pad

Instrumentation
Offsets
□ Automatic
□ Time Stamp □ Packet Groups □ Sequence Checking □ Data Integrity

Force Errors
● No Error ○ Bad CRC

Figure 176. Set protocol to Pause Control, and then click Edit

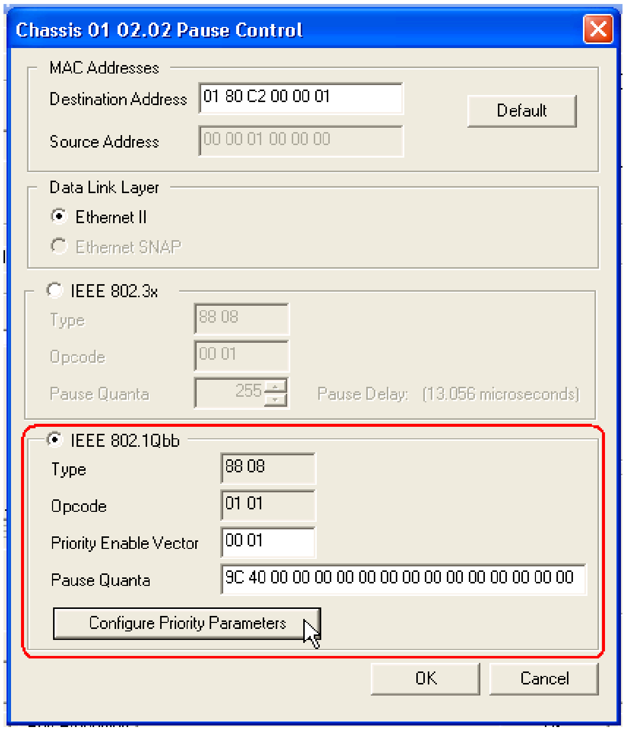

Chassis 01 02.02 Pause Control

MAC Addresses
Destination Address: 01 80 C2 00 00 01
Source Address: 00 00 01 00 00 00

Data Link Layer
● Ethernet II
○ Ethernet SNAP

○ IEEE 802.3x
Type: 88 08
Opcode: 00 01
Pause Quanta: 255
Pause Delay: (13.056 microseconds)

● IEEE 802.1Qbb
Type: 88 08
Opcode: 01 01
Priority Enable Vector: 00 01
Pause Quanta: 9C 40 00 00 00 00 00 00 00 00 00 00 00 00 00 00

Figure 177. Select IEEE 802.1Qbb and select Configure Priority Parameters

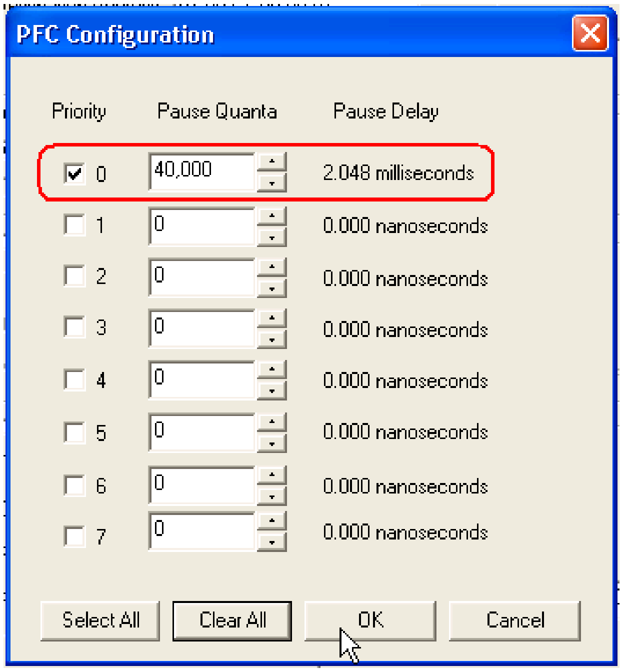

PFC Configuration

Priority	Pause Quanta	Pause Delay
☑ 0	40,000	2.048 milliseconds
□ 1	0	0.000 nanoseconds
□ 2	0	0.000 nanoseconds
□ 3	0	0.000 nanoseconds
□ 4	0	0.000 nanoseconds
□ 5	0	0.000 nanoseconds
□ 6	0	0.000 nanoseconds
□ 7	0	0.000 nanoseconds

Figure 178. Enable priority 0 with Pause Quanta of 40,000 (~ 2ms)

5. Set up the capture trigger and filter for measuring response time and pause duration.

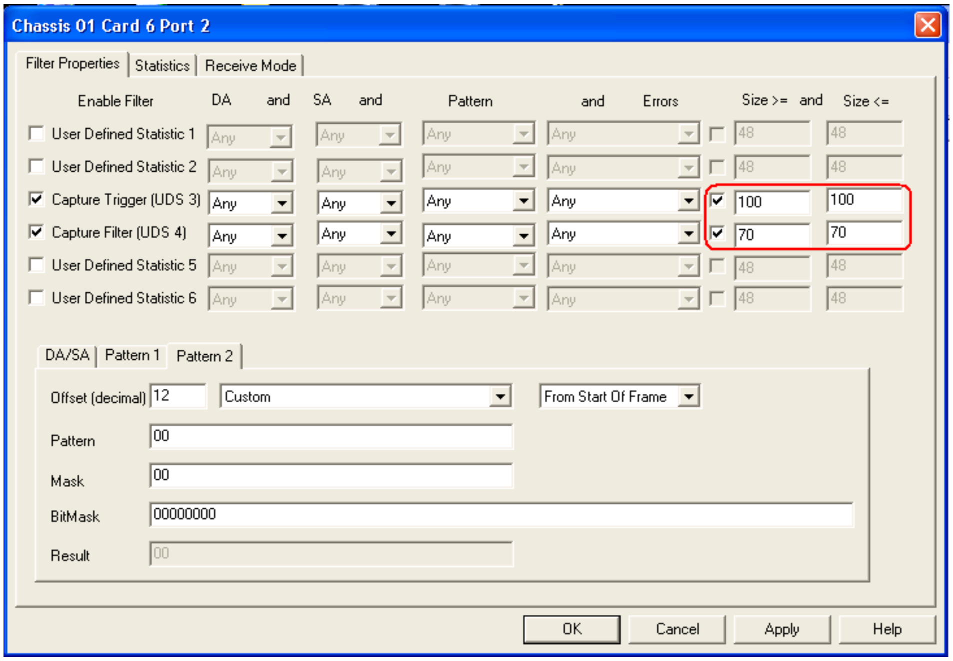

Chassis 01 Card 6 Port 2

Enable Filter	DA	and SA	and Pattern	and Errors		Size >=	and Size <=
□ User Defined Statistic 1	Any	Any	Any	Any	□	48	48
□ User Defined Statistic 2	Any	Any	Any	Any	□	48	48
☑ Capture Trigger (UDS 3)	Any	Any	Any	Any	☑	100	100
☑ Capture Filter (UDS 4)	Any	Any	Any	Any	☑	70	70
□ User Defined Statistic 5	Any	Any	Any	Any	□	48	48
□ User Defined Statistic 6	Any	Any	Any	Any	□	48	48

Offset (decimal) 12 Custom From Start Of Frame

Pattern 00

Mask 00

BitMask 00000000

Result 00

Figure 179. Set up capture filter to trigger on the 'trigger packet' and filter on the paused flow packets.

Terminology

Term	Description
FC	Fibre Channel
FCoE	Fibre Channel over Ethernet
Initiator	Device such as a HBA that acts as a SCSI initiator of commands
Target	Devices such as a disk drive that receive commands from a initiator
LUN	Logical Unit
Throughput	Amount of data that is being transferred between host and target
VN_Port	Virtual node port
VF_Port	Virtual fabric port
FCoE LEP	FCoE Link End Point
Busy	A response condition of the target wherein it is unable to execute the command requested by the initiator due to lack of resources
Dropped Frame	A condition where in a read command, the Target sends out the frame but it does not get to the Initiator. Similarly, on a write command, the Initiator sends the frame to the target but the target does not receive the frame

Introduction to Data Center Virtualization Infrastructure

The emergence of virtualization technology during the last decade has led to a paradigm shift in the way data centers are designed and managed. Server hardware has become an unseen commodity – application servers exist on a virtualized platform that can be migrated from one machine to another in the blink of an eye. Application services running in a virtualized environment instantly become fault-tolerant, scalable and easily upgradeable to any hardware platform - giving IT managers the ability to focus on delivering an optimal experience for their users.

The power of the virtualized data center brings new challenges and opportunities for technology vendors and IT managers. To enhance the performance and security of virtualized environments, technology vendors are building an entirely new class of virtualized network equipment (virtual NICs, virtual switches) that crosses traditional hardware/software boundaries. Testing these devices poses new challenges because most test tools are not capable of bridging this gap.

Because many virtual machines can co-exist on a single hardware platform, the number of IP-based network assets can increase dramatically. At the same time, these virtualized assets are sharing limited resources such as CPU, memory, and network interfaces. Without careful planning and testing, some virtual machines may end up starved of these critical resources. Effective management of the virtualized network infrastructure can only be achieved with tools and methodologies designed with virtualization in mind.

Virtualization Technology Overview

Virtualization technology allows a single computer to run multiple operating environments simulataneously. The key software component that enables virtualization is called the *hypervisor,* which separates the guest operating systems from the host's hardware. This provides a number of advantages in addition to consolidating multiple operating systems and applications onto a single system.

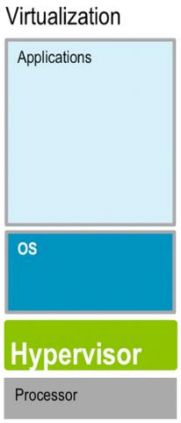

Flexibility

By abstracting the underlying hardware from guest operating systems, the hypervisor presents the same kind of hardware to each guest, regardless of the underlying devices. This means that administrators no longer have to devote hours to hunting down the right device drivers for each OS/hardware combination. It also means that a virtual machine can be moved from one kind of hardware to another without issue. IT engineers are free to install whichever kind of hardware suits their needs and continue to use their existing VMs on those new systems.

Efficiency

Because of the strict OS, software and hardware compatibility requirements of many applications, enterprise IT departments would find themselves dedicating an entire server to a single application that did not fully utilize the server's resources. By migrating little- or infrequently-used applications to virtual machines, IT administrators can dramatically reduce the number of physical servers under management. In addition, a single new multi-core, multi-processor server with the same aggregate capability as several older servers can often use less power and cooling than one older server saving on critical energy costs.

Maintainability

As a result of the efficiency offered by virtualization technology, the number of physical systems under management by an enterprise can be significantly reduced. This lowers the overall cost of managing the datacenter in both manpower and fixed costs. Additionally, the unification to a single, virtualized, hardware platform for all guest operating systems and applications makes management of datacenter resources far easier that it has been in the past.

Reliability

One of the most unique benefits offered by virtualization technology is that of *live migration.* Live migration is the ability of a virtualization system to move a running VM from one physical host to another without interruption. This enables a whole new level of fault tolerance to the datacenter that can be enabled for any application without being specifically designed for fault-tolerant or redundant behavior. Server engineers can now quickly move all running VMs from a failing blade to other blades in the system while they swap out the bad hardware. In the future, entire datacenters of VMs will be kept synchronized between multiple locations so that operation can be assumed by another location in case of an outage at the main site.

Dynamic Scalability

Live migration and the ability to quickly clone VMs also enable a new level of scalability that was previously unattainable in data centers. Data center administrators can configure virtualized environments to dynamically scale up, or down, based on usage. For example, if a datacenter is hosting a website on a VM and the website is suddenly barraged with a large number of requests, the virtualization management system can dynamically trigger replication of the HTTP server’s VM to other hardware that is either free or running lower-priority applications. This enables the datacenter to handle bursts of computing requirements on very short notice. Alternatively, when the datacenter is idling during low-demand periods, the management system can move VMs onto as few hardware platforms as possible so that unused servers can be placed in a low-power state to conserve energy.

Network Infrastructure in a Virtualized Data Center

Virtual Switches

When many virtual machines are running side-by-side on the same physical server, they are connected to each other, and the external network, via a *virtual switch*. The virtual switch acts just as a physical switch would—it monitors traffic on each of its virtual ports (one for each virtual NIC attached to a VM), keeps track of MAC addresses that are directly attached, forwards locally-bound traffic directly to the correct virtual port, and sends traffic destined for foreign MAC addresses to the upstream switch through the server's physical NIC.

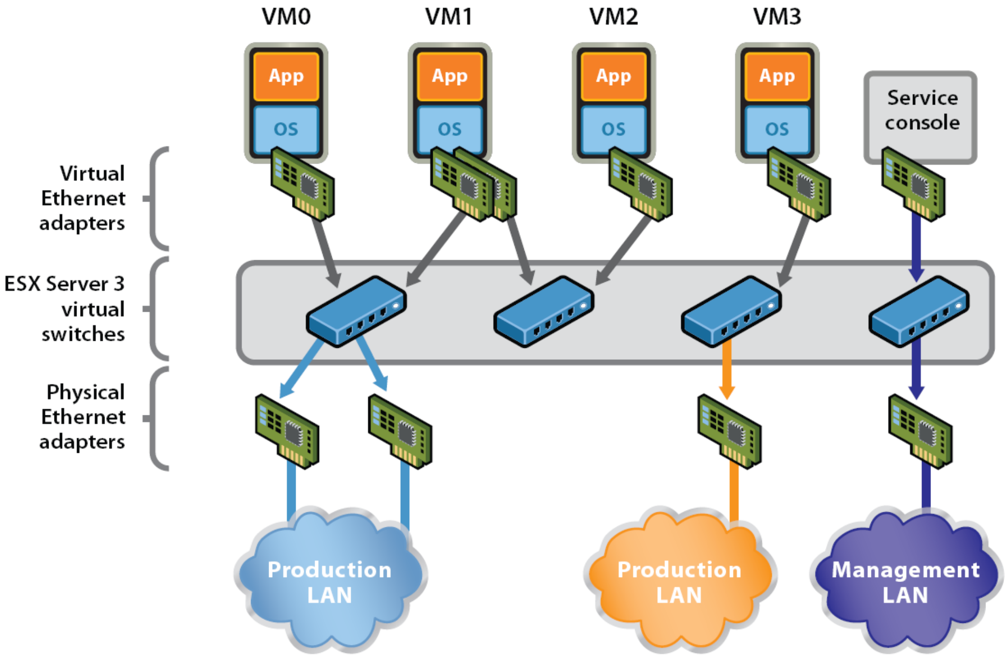

Figure 180. Diagram of a virtual switch in a typical deployment scenario.

There are a number of configurable options on the standard virtual switch included with VMware ESX server:

- Number of virtual ports
- Port groups and per-port-group VLAN tagging
- Per-port-group bandwidth limitations
- Uplink NIC teaming for redundancy & load balancing

Distributed Virtual Switches

In the most recent version of VMware, two new virtual switches are available for VM administrators—the VMware Distributed Virtual Switch and the Cisco Nexus 1000V virtual switch. These new virtual switches allow the switching infrastructure to span multiple hosts so that administrators can set global policies that affect all VMs and hosts on the distributed switching system. VLANs and port policies can be managed centrally so that when VMs move to other hosts, the configuration is kept in its original form. The following figure from VMware highlights the key differences:

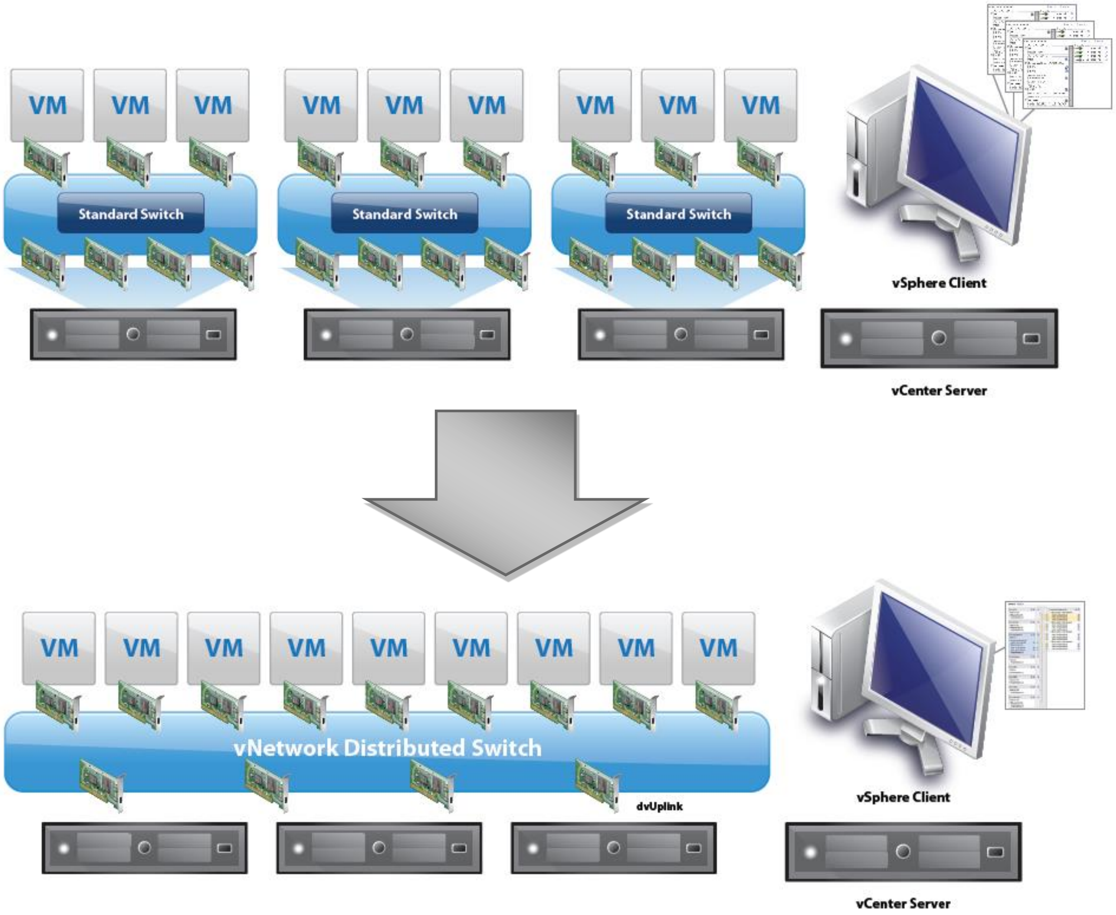

Figure 181. Traditional virtual switches vs. distributed virtual switches.

Cloud Computing

What is Cloud Computing?

Simply put, cloud computing is the use of computing resources located outside the end-user's datacenter. Many companies like Amazon, AT&T, Rackspace, Savvis, Terremark and others now offer computing platforms on a metered, utility-style pricing plan. Enterprises can access these resources to add dynamic capacity to their existing datacenters by turning up VMs on cloud datacenters to handle high-demand periods. By offering datacenter infrastructure as a service, enterprises have access to computing resources that were previously out-of-reach.

Computing-as-a-Service

There are many cloud acronyms being used today to describe to describe the kinds of cloud-based offerings that are being marketed:

SaaS – Software-as-a-Service

Companies like Salesforce.com and Google Docs offer entire application suites that are hosted on their own servers and accessed through standard web browsers.

IaaS – Infrastructure-as-a-Service

Services like Amazon's EC2 offer complete datacenter infrastructure – VMs hosted on high-end machines with full access to other VMs and the Internet. Extensions to this service offer *private cloud* access by bringing the remotely-hosted VMs onto the enterprise network via a VPN connection.

PaaS – Platform-as-a-Service

Some web services and SaaS products can be written to utilize cloud-located resources without needing a full VM and operating system. Services like Google's AppEngine and Microsoft's Azure enable Python and .NET applications to run on Google and Microsoft servers. These kinds of services are popular for web-based widgets like Facebook applications.

‘Scaling-Up’ vs. ‘Scaling-Out’

In a traditional datacenter, a user who required more computing power would ‘scale-up’ their resources by buying additional hardware to meet their needs. In the cloud-computing world, enterprises have the option of ‘scaling-out’ to utilize hardware outside their own datacenter. If the resources are needed for a long term or have specific latency or privacy requirements, users may still opt to scale ‘up.’ In other circumstances, scaling ‘out’ offers many economic advantages.

Ixia's IxVM

Ixia's IxVM products test virtualized assets with the same trusted tools that have been used for testing hardware-based networks and network equipment for over a decade. The components of the IxVM solution include:

IxVM Virtual Test Ports – The IxVM Virtual Port is a software-based version of a traditional Ixia port that provides a solution for testing virtualized environments. IxVM virtual test ports can be deployed directly on to existing servers or virtual machines running common Linux operating systems such as RedHat Enterprise Linux. After deploying the IxVM virtual test port software, these VMs gain the ability to send and receive L2-7 traffic just like a standard Ixia port. In the virtual port concept, a standard *baremetal* server or VM represents a *card* and NICs/vNICs represent *ports*. Software-based stream and protocol engines are installed on the virtual ports to provide the complete L2-7 traffic generation capabilities offered by Ixia’s hardware-based product portfolio.

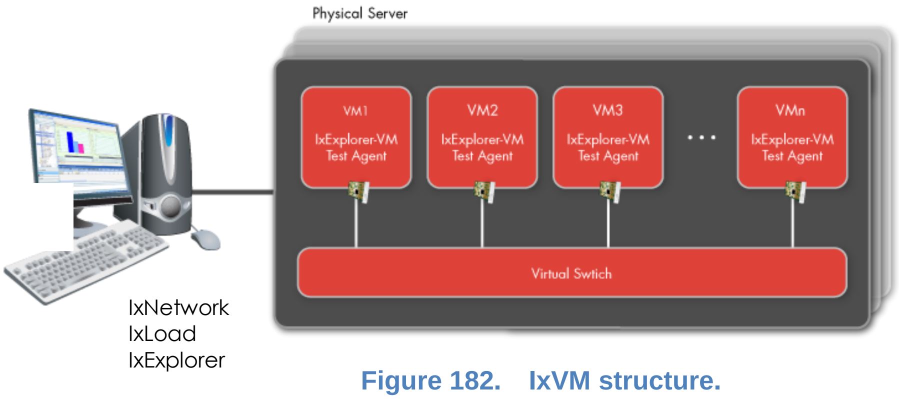

Figure 182. IxVM structure.

- **IxNetwork/VM –** IxNetwork offers the advanced capability for users to quickly build large-scale emulated network topologies using real-world routing and switching protocols and then build complex traffic topologies on those emulated topologies. IxNetwork’s large library of L2-3 protocols allows advanced testing of multicast scenarios using IGMP, convergence testing for measuring outage duration during live migration, LACP, and other kinds of failures for example.

- **IxLoad/VM** – IxLoad enables L4-7 application testing including a full range of data, file transfer, and VoIP testing allowing users to test a wide variety of virtualized application architectures. IxLoad/VM will also add new capabilities for testing client disk I/O performance which will enable complete performance testing of both the HBA and NIC aspects of 10G converged network adapters (CNAs).

Test Case: Basic Unicast Testing

Overview

This test case will describe the basic setup and steps required to use IxNetwork inside a virtualized environment. After the basic setup is complete, you can then build more sophisticated test cases in the same way as you would in a traditional hardware-based switch testing environment.

Objective

The objective of this test is to build a full-mesh of traffic between a number of virtual machines on two separate servers to exercise the vSwitch on each server, the distributed virtual switch and the physical switch connecting the two servers. After this basic test scenario is set up, it can be used to confirm full-mesh connectivity and measure the performance of virtualized infrastructure.

Setup

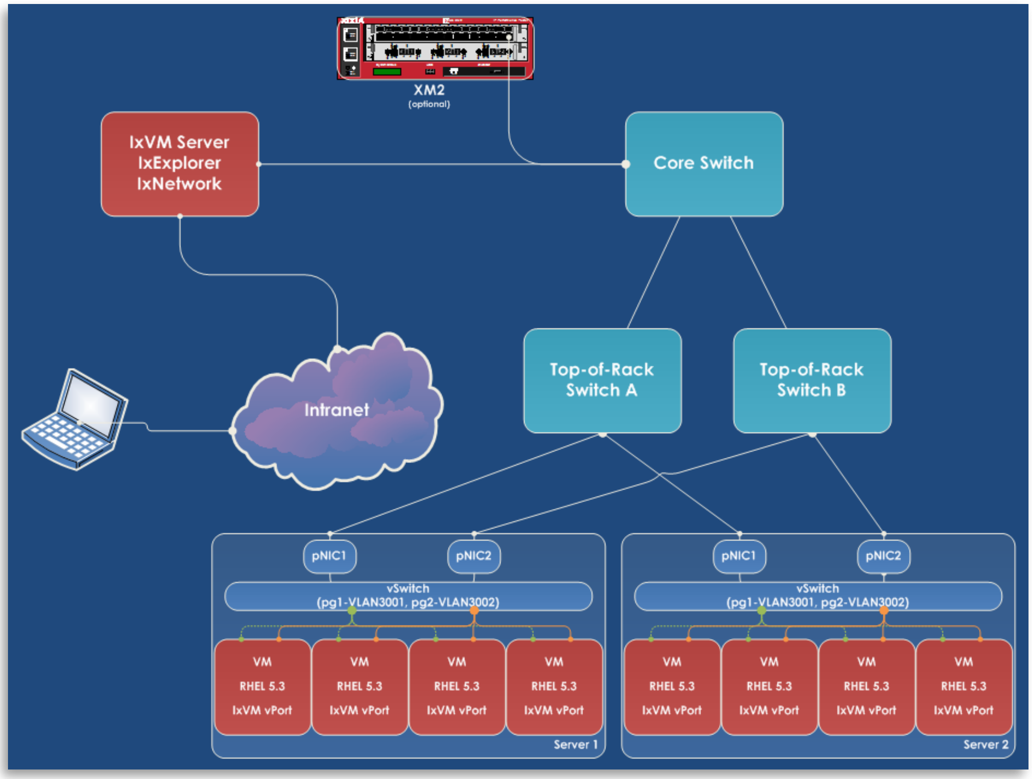

Figure 183. Basic vSwitch testing topology

In the following figure, the IxVM virtual port software has been deployed on to four virtual machines on each VMware ESX host. The virtual port software can be deployed in several different ways, depending on the requirements of the user:

IxVM virtual ports can be installed on existing virtual machines alongside other software that may be running on those VMs

A master VM image can be created with an IxVM virtual port and then cloned to other VMs and hosts

The IxVM virtual appliance can be downloaded and deployed with the IxVM virtual port software already installed.

After deploying the virtual port software, IxExplorer is used to add the virtual cards and ports to the IxVM virtual chassis so that they can be used in the examples below.

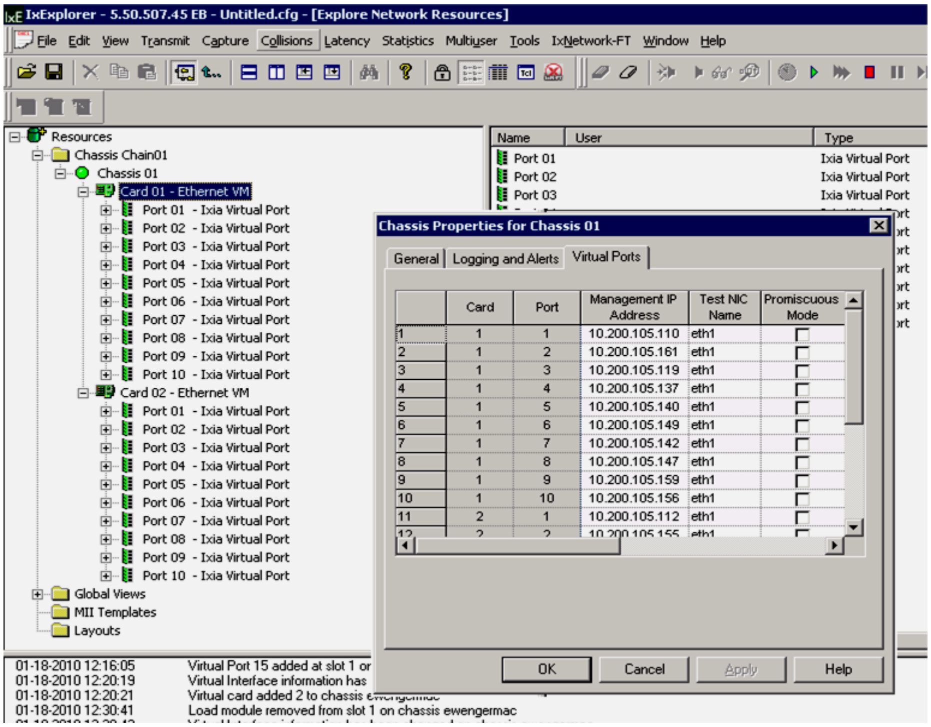
IxExplorer - 5.50.507.45 EB - Untitled.cfg - [Explore Network Resources]
File Edit View Transmit Capture Collisions Latency Statistics Multiuser Tools IxNetwork-FT Window Help
Resources
Chassis Chain01
Chassis 01
Card 01 - Ethernet VM
Port 01 - Ixia Virtual Port
Port 02 - Ixia Virtual Port
Port 03 - Ixia Virtual Port
Port 04 - Ixia Virtual Port
Port 05 - Ixia Virtual Port
Port 06 - Ixia Virtual Port
Port 07 - Ixia Virtual Port
Port 08 - Ixia Virtual Port
Port 09 - Ixia Virtual Port
Port 10 - Ixia Virtual Port
Card 02 - Ethernet VM
Port 01 - Ixia Virtual Port
Port 02 - Ixia Virtual Port
Port 03 - Ixia Virtual Port
Port 04 - Ixia Virtual Port
Port 05 - Ixia Virtual Port
Port 06 - Ixia Virtual Port
Port 07 - Ixia Virtual Port
Port 08 - Ixia Virtual Port
Port 09 - Ixia Virtual Port
Port 10 - Ixia Virtual Port
Global Views
MII Templates
Layouts
Name
User
Type
Port 01
Ixia Virtual Port
Port 02
Ixia Virtual Port
Port 03
Ixia Virtual Port
Chassis Properties for Chassis 01
General
Logging and Alerts
Virtual Ports
OK
Cancel
Apply
Help

	Card	Port	Management IP Address	Test NIC Name	Promiscuous Mode
1	1	1	10.200.105.110	eth1	□
2	1	2	10.200.105.161	eth1	□
3	1	3	10.200.105.119	eth1	□
4	1	4	10.200.105.137	eth1	□
5	1	5	10.200.105.140	eth1	□
6	1	6	10.200.105.149	eth1	□
7	1	7	10.200.105.142	eth1	□
8	1	8	10.200.105.147	eth1	□
9	1	9	10.200.105.159	eth1	□
10	1	10	10.200.105.156	eth1	□
11	2	1	10.200.105.112	eth1	□
12	2	2	10.200.105.155	eth1	□

```
01-18-2010 12:16:05     Virtual Port 15 added at slot 1 or
01-18-2010 12:20:19     Virtual Interface information has
01-18-2010 12:20:21     Virtual card added 2 to chassis
01-18-2010 12:30:41     Load module removed from slot 1 on chassis ewengermac
```

Figure 184. IxVM Virtual Ports in IxExplorer.

Although a chassis is shown in the figure, it is not used in the tests described below. IxVM ports and the Ixia hardware chassis ports are fully compatible and can be used seamlessly together in any test configuration.

Step-by-Step Instructions

1. Add the IxVM virtual ports in IxNetwork. You can add all ports in one step by selecting the chassis (*localhost* below) and then select the right-pointing arrow in the middle to add new ports and assign the selected ports.

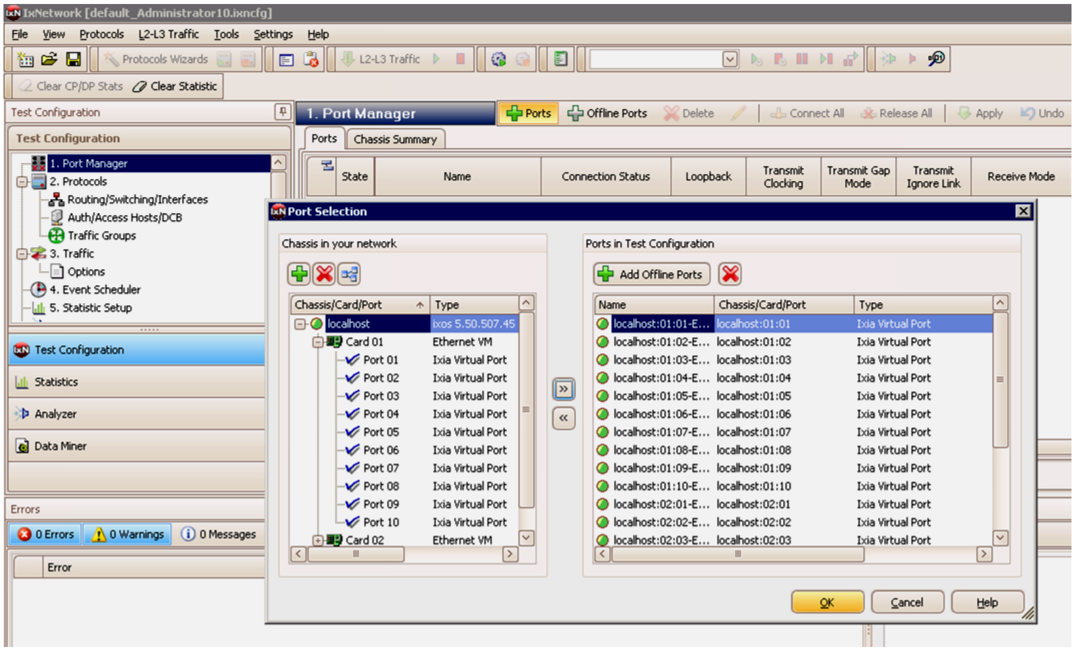

Chassis/Card/Port	Type
localhost	ixos 5.50.507.45
Card 01	Ethernet VM
Port 01	Ixia Virtual Port
Port 02	Ixia Virtual Port
Port 03	Ixia Virtual Port
Port 04	Ixia Virtual Port
Port 05	Ixia Virtual Port
Port 06	Ixia Virtual Port
Port 07	Ixia Virtual Port
Port 08	Ixia Virtual Port
Port 09	Ixia Virtual Port
Port 10	Ixia Virtual Port
Card 02	Ethernet VM

Name	Chassis/Card/Port	Type
localhost:01:01-E...	localhost:01:01	Ixia Virtual Port
localhost:01:02-E...	localhost:01:02	Ixia Virtual Port
localhost:01:03-E...	localhost:01:03	Ixia Virtual Port
localhost:01:04-E...	localhost:01:04	Ixia Virtual Port
localhost:01:05-E...	localhost:01:05	Ixia Virtual Port
localhost:01:06-E...	localhost:01:06	Ixia Virtual Port
localhost:01:07-E...	localhost:01:07	Ixia Virtual Port
localhost:01:08-E...	localhost:01:08	Ixia Virtual Port
localhost:01:09-E...	localhost:01:09	Ixia Virtual Port
localhost:01:10-E...	localhost:01:10	Ixia Virtual Port
localhost:02:01-E...	localhost:02:01	Ixia Virtual Port
localhost:02:02-E...	localhost:02:02	Ixia Virtual Port
localhost:02:03-E...	localhost:02:03	Ixia Virtual Port

Figure 185. Adding Ports

2. Enable the ARP protocol for all ports in the test. You can select the entire column by clicking on the column header and then enable all ports by right-clicking and selecting **Enable Selection**.

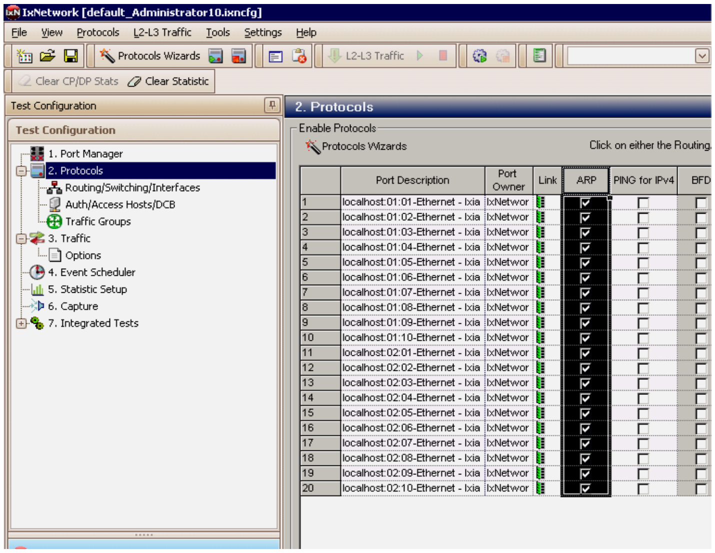

	Port Description	Port Owner	Link	ARP	PING for IPv4	BFD
1	localhost:01:01-Ethernet - Ixia	IxNetwor		☑	□	□
2	localhost:01:02-Ethernet - Ixia	IxNetwor		☑	□	□
3	localhost:01:03-Ethernet - Ixia	IxNetwor		☑	□	□
4	localhost:01:04-Ethernet - Ixia	IxNetwor		☑	□	□
5	localhost:01:05-Ethernet - Ixia	IxNetwor		☑	□	□
6	localhost:01:06-Ethernet - Ixia	IxNetwor		☑	□	□
7	localhost:01:07-Ethernet - Ixia	IxNetwor		☑	□	□
8	localhost:01:08-Ethernet - Ixia	IxNetwor		☑	□	□
9	localhost:01:09-Ethernet - Ixia	IxNetwor		☑	□	□
10	localhost:01:10-Ethernet - Ixia	IxNetwor		☑	□	□
11	localhost:02:01-Ethernet - Ixia	IxNetwor		☑	□	□
12	localhost:02:02-Ethernet - Ixia	IxNetwor		☑	□	□
13	localhost:02:03-Ethernet - Ixia	IxNetwor		☑	□	□
14	localhost:02:04-Ethernet - Ixia	IxNetwor		☑	□	□
15	localhost:02:05-Ethernet - Ixia	IxNetwor		☑	□	□
16	localhost:02:06-Ethernet - Ixia	IxNetwor		☑	□	□
17	localhost:02:07-Ethernet - Ixia	IxNetwor		☑	□	□
18	localhost:02:08-Ethernet - Ixia	IxNetwor		☑	□	□
19	localhost:02:09-Ethernet - Ixia	IxNetwor		☑	□	□
20	localhost:02:10-Ethernet - Ixia	IxNetwor		☑	□	□

Figure 186. ARP Enabled

3. Create IP protocol interfaces on each virtual port. A very simple IP scheme for a simple L2-3 network is shown in the following figure. To configure these addresses, first enter **1.1.1.1** as the IP address for the first row. Next, select the entire IP address column and right-click to choose **Incremen**t. Next, enter **1.1.1.1** in the *Gateway* column and right-click to choose **Same**. After that is complete, make sure to change the gateway in the first row to **1.1.1.2** because it cannot act as its own gateway. Finally, enable all interfaces.

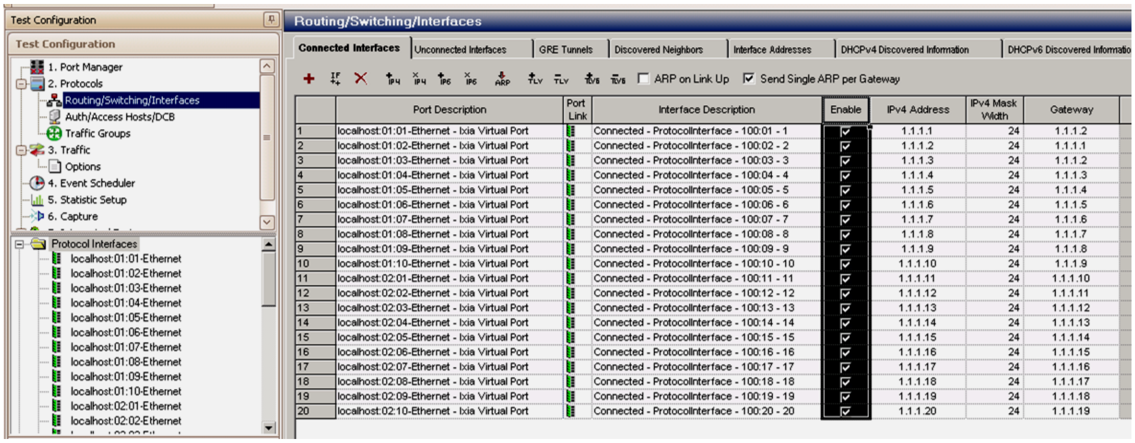

☐ ARP on Link Up ☑ Send Single ARP per Gateway

	Port Description	Port Link	Interface Description	Enable	IPv4 Address	IPv4 Mask Width	Gateway
1	localhost:01:01-Ethernet - Ixia Virtual Port		Connected - ProtocolInterface - 100:01 - 1	☑	1.1.1.1	24	1.1.1.2
2	localhost:01:02-Ethernet - Ixia Virtual Port		Connected - ProtocolInterface - 100:02 - 2	☑	1.1.1.2	24	1.1.1.1
3	localhost:01:03-Ethernet - Ixia Virtual Port		Connected - ProtocolInterface - 100:03 - 3	☑	1.1.1.3	24	1.1.1.2
4	localhost:01:04-Ethernet - Ixia Virtual Port		Connected - ProtocolInterface - 100:04 - 4	☑	1.1.1.4	24	1.1.1.3
5	localhost:01:05-Ethernet - Ixia Virtual Port		Connected - ProtocolInterface - 100:05 - 5	☑	1.1.1.5	24	1.1.1.4
6	localhost:01:06-Ethernet - Ixia Virtual Port		Connected - ProtocolInterface - 100:06 - 6	☑	1.1.1.6	24	1.1.1.5
7	localhost:01:07-Ethernet - Ixia Virtual Port		Connected - ProtocolInterface - 100:07 - 7	☑	1.1.1.7	24	1.1.1.6
8	localhost:01:08-Ethernet - Ixia Virtual Port		Connected - ProtocolInterface - 100:08 - 8	☑	1.1.1.8	24	1.1.1.7
9	localhost:01:09-Ethernet - Ixia Virtual Port		Connected - ProtocolInterface - 100:09 - 9	☑	1.1.1.9	24	1.1.1.8
10	localhost:01:10-Ethernet - Ixia Virtual Port		Connected - ProtocolInterface - 100:10 - 10	☑	1.1.1.10	24	1.1.1.9
11	localhost:02:01-Ethernet - Ixia Virtual Port		Connected - ProtocolInterface - 100:11 - 11	☑	1.1.1.11	24	1.1.1.10
12	localhost:02:02-Ethernet - Ixia Virtual Port		Connected - ProtocolInterface - 100:12 - 12	☑	1.1.1.12	24	1.1.1.11
13	localhost:02:03-Ethernet - Ixia Virtual Port		Connected - ProtocolInterface - 100:13 - 13	☑	1.1.1.13	24	1.1.1.12
14	localhost:02:04-Ethernet - Ixia Virtual Port		Connected - ProtocolInterface - 100:14 - 14	☑	1.1.1.14	24	1.1.1.13
15	localhost:02:05-Ethernet - Ixia Virtual Port		Connected - ProtocolInterface - 100:15 - 15	☑	1.1.1.15	24	1.1.1.14
16	localhost:02:06-Ethernet - Ixia Virtual Port		Connected - ProtocolInterface - 100:16 - 16	☑	1.1.1.16	24	1.1.1.15
17	localhost:02:07-Ethernet - Ixia Virtual Port		Connected - ProtocolInterface - 100:17 - 17	☑	1.1.1.17	24	1.1.1.16
18	localhost:02:08-Ethernet - Ixia Virtual Port		Connected - ProtocolInterface - 100:18 - 18	☑	1.1.1.18	24	1.1.1.17
19	localhost:02:09-Ethernet - Ixia Virtual Port		Connected - ProtocolInterface - 100:19 - 19	☑	1.1.1.19	24	1.1.1.18
20	localhost:02:10-Ethernet - Ixia Virtual Port		Connected - ProtocolInterface - 100:20 - 20	☑	1.1.1.20	24	1.1.1.19

Figure 187. Interfaces Enabled

4. Move to step 3, *Traffic*, in the Test Configuration and start the *Basic Traffic Wizard* by clicking Basic Wizard. Use the Traffic Wizard to create a fully meshed traffic topology.

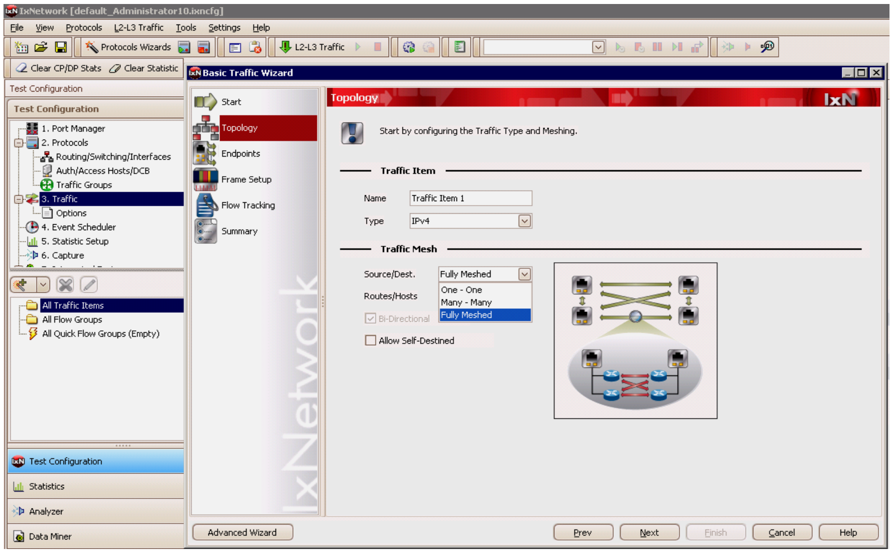

Topology

Start by configuring the Traffic Type and Meshing.

Traffic Item

Name: Traffic Item 1

Type: IPv4

Traffic Mesh

Source/Dest.: Fully Meshed (One - One, Many - Many, Fully Meshed)

Routes/Hosts

☑ Bi-Directional

☐ Allow Self-Destined

Figure 188. Selecting a Fully Meshed Topology

5. Select **All Ports** to ensure that all IxVM ports are part of the full mesh traffic item.

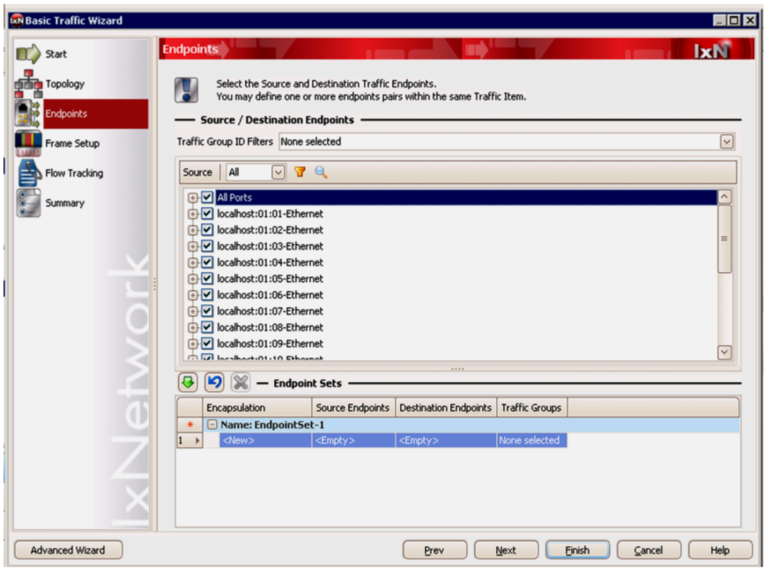

Endpoints

Select the Source and Destination Traffic Endpoints.
You may define one or more endpoints pairs within the same Traffic Item.

Source / Destination Endpoints

Traffic Group ID Filters: None selected

Source: All

- ☑ All Ports
- ☑ localhost:01:01-Ethernet
- ☑ localhost:01:02-Ethernet
- ☑ localhost:01:03-Ethernet
- ☑ localhost:01:04-Ethernet
- ☑ localhost:01:05-Ethernet
- ☑ localhost:01:06-Ethernet
- ☑ localhost:01:07-Ethernet
- ☑ localhost:01:08-Ethernet
- ☑ localhost:01:09-Ethernet
- ☑ localhost:01:10-Ethernet

Endpoint Sets

	Encapsulation	Source Endpoints	Destination Endpoints	Traffic Groups
	Name: EndpointSet-1			
1	<New>	<Empty>	<Empty>	None selected

Figure 189. Selecting Ports

6. Configure the frame size and rate. In this case, 1000-byte frames will be sent at 100 frames per second. In VM-to-VM testing, % *Line Rate* does not apply because the vNIC does not have a fixed throughput.

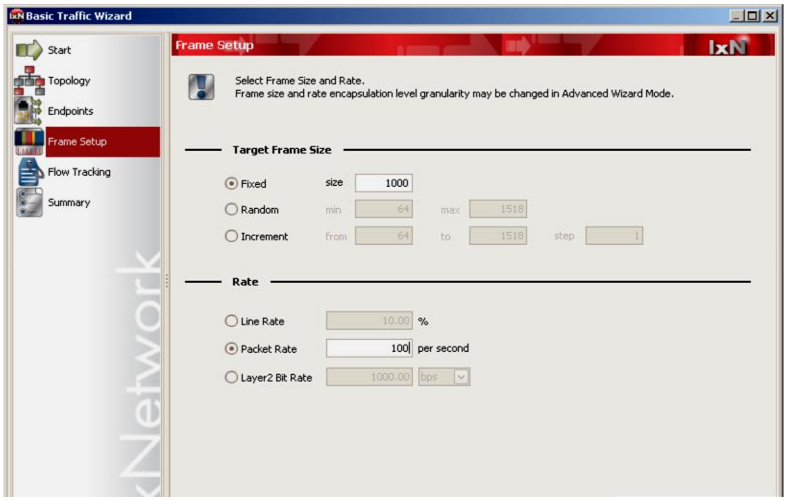

Frame Setup

Select Frame Size and Rate.
Frame size and rate encapsulation level granularity may be changed in Advanced Wizard Mode.

Target Frame Size

- ● Fixed size 1000
- ○ Random min 64 max 1518
- ○ Increment from 64 to 1518 step 1

Rate

- ○ Line Rate 10.00 %
- ● Packet Rate 100 per second
- ○ Layer2 Bit Rate 1000.00 bps

Figure 190. Frame Size and Rate

7. Select the necessary tracking fields to get the level of detail required in the statistics view. The options selected in the screenshot below will allow the tester to view traffic stats aggregated by the source/destination port pairs and source and destination IP addresses.

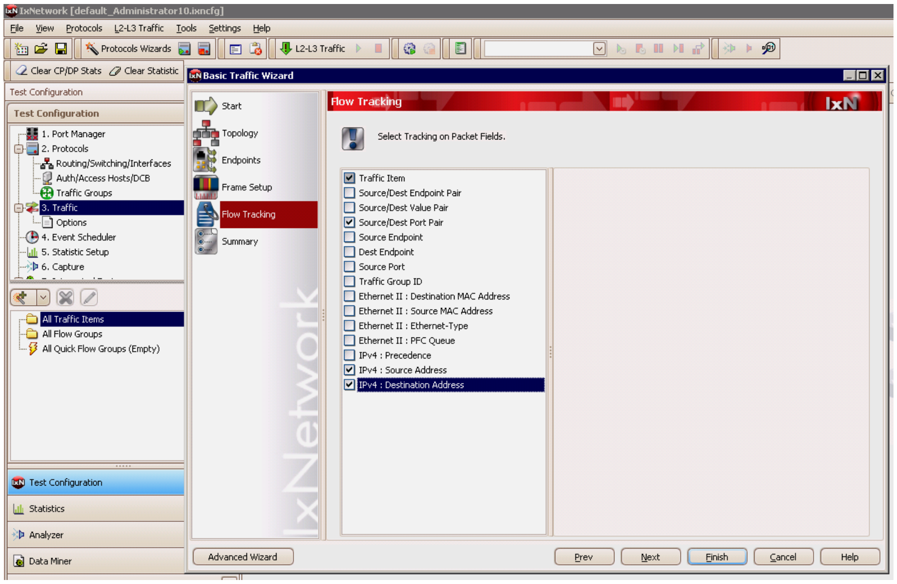

Flow Tracking

Select Tracking on Packet Fields.

- ☑ Traffic Item
- ☐ Source/Dest Endpoint Pair
- ☐ Source/Dest Value Pair
- ☑ Source/Dest Port Pair
- ☐ Source Endpoint
- ☐ Dest Endpoint
- ☐ Source Port
- ☐ Traffic Group ID
- ☐ Ethernet II : Destination MAC Address
- ☐ Ethernet II : Source MAC Address
- ☐ Ethernet II : Ethernet-Type
- ☐ Ethernet II : PFC Queue
- ☐ IPv4 : Precedence
- ☑ IPv4 : Source Address
- ☑ IPv4 : Destination Address

Figure 191. Flow Tracking

8. Finish the wizard and review the test configuration. Each rounded rectangle represents one of the ports used in the test. By selecting each port, you can see where its traffic is destined or originating from to validate the test topology.

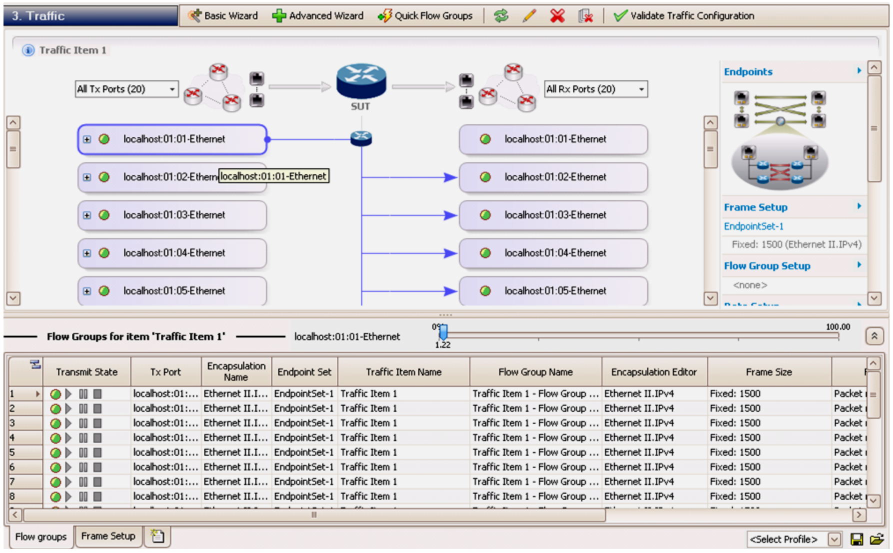

Traffic Item 1

All Tx Ports (20) → SUT → All Rx Ports (20)

Tx	Rx
localhost:01:01-Ethernet	localhost:01:01-Ethernet
localhost:01:02-Ethernet	localhost:01:02-Ethernet
localhost:01:03-Ethernet	localhost:01:03-Ethernet
localhost:01:04-Ethernet	localhost:01:04-Ethernet
localhost:01:05-Ethernet	localhost:01:05-Ethernet

Endpoints

Frame Setup

EndpointSet-1

Fixed: 1500 (Ethernet II.IPv4)

Flow Group Setup

<none>

Flow Groups for item 'Traffic Item 1' — localhost:01:01-Ethernet

	Tx Port	Encapsulation Name	Endpoint Set	Traffic Item Name	Flow Group Name	Encapsulation Editor	Frame Size
1	localhost:01:...	Ethernet II.I...	EndpointSet-1	Traffic Item 1	Traffic Item 1 - Flow Group ...	Ethernet II.IPv4	Fixed: 1500
2	localhost:01:...	Ethernet II.I...	EndpointSet-1	Traffic Item 1	Traffic Item 1 - Flow Group ...	Ethernet II.IPv4	Fixed: 1500
3	localhost:01:...	Ethernet II.I...	EndpointSet-1	Traffic Item 1	Traffic Item 1 - Flow Group ...	Ethernet II.IPv4	Fixed: 1500
4	localhost:01:...	Ethernet II.I...	EndpointSet-1	Traffic Item 1	Traffic Item 1 - Flow Group ...	Ethernet II.IPv4	Fixed: 1500
5	localhost:01:...	Ethernet II.I...	EndpointSet-1	Traffic Item 1	Traffic Item 1 - Flow Group ...	Ethernet II.IPv4	Fixed: 1500
6	localhost:01:...	Ethernet II.I...	EndpointSet-1	Traffic Item 1	Traffic Item 1 - Flow Group ...	Ethernet II.IPv4	Fixed: 1500
7	localhost:01:...	Ethernet II.I...	EndpointSet-1	Traffic Item 1	Traffic Item 1 - Flow Group ...	Ethernet II.IPv4	Fixed: 1500
8	localhost:01:...	Ethernet II.I...	EndpointSet-1	Traffic Item 1	Traffic Item 1 - Flow Group ...	Ethernet II.IPv4	Fixed: 1500

Figure 192. Test Configuration

9. Click **Apply Traffic** (L2-L3 Traffic), then click ▶ to run the test. Select **Statistics** to monitor the statistics in real-time.

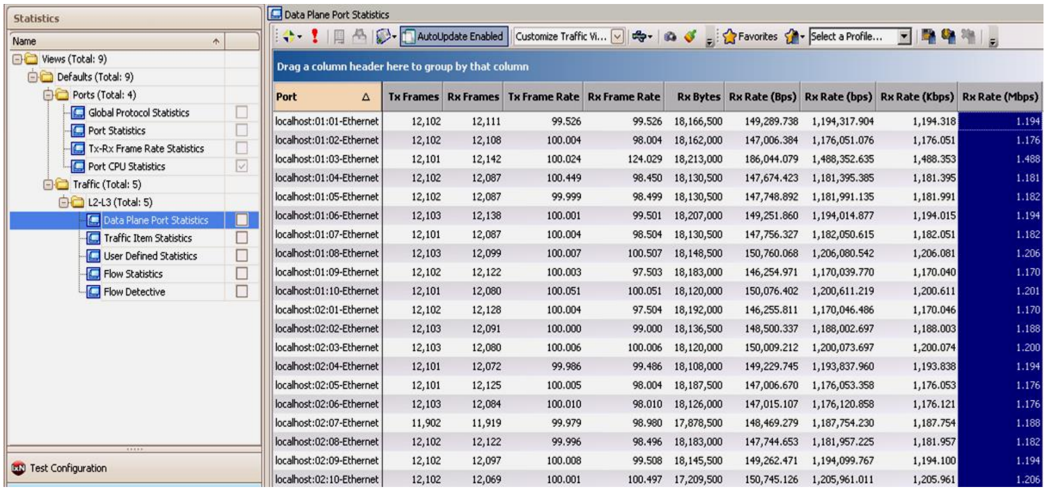

Port	Tx Frames	Rx Frames	Tx Frame Rate	Rx Frame Rate	Rx Bytes	Rx Rate (Bps)	Rx Rate (bps)	Rx Rate (Kbps)	Rx Rate (Mbps)
localhost:01:01-Ethernet	12,102	12,111	99.526	99.526	18,166,500	149,289.738	1,194,317.904	1,194.318	1.194
localhost:01:02-Ethernet	12,102	12,108	100.004	98.004	18,162,000	147,006.384	1,176,051.076	1,176.051	1.176
localhost:01:03-Ethernet	12,101	12,142	100.024	124.029	18,213,000	186,044.079	1,488,352.635	1,488.353	1.488
localhost:01:04-Ethernet	12,102	12,087	100.449	98.450	18,130,500	147,674.423	1,181,395.385	1,181.395	1.181
localhost:01:05-Ethernet	12,102	12,087	99.999	98.499	18,130,500	147,748.892	1,181,991.135	1,181.991	1.182
localhost:01:06-Ethernet	12,103	12,138	100.001	99.501	18,207,000	149,251.860	1,194,014.877	1,194.015	1.194
localhost:01:07-Ethernet	12,101	12,087	100.004	98.504	18,130,500	147,756.327	1,182,050.615	1,182.051	1.182
localhost:01:08-Ethernet	12,103	12,099	100.007	100.507	18,148,500	150,760.068	1,206,080.542	1,206.081	1.206
localhost:01:09-Ethernet	12,102	12,122	100.003	97.503	18,183,000	146,254.971	1,170,039.770	1,170.040	1.170
localhost:01:10-Ethernet	12,101	12,080	100.051	100.051	18,120,000	150,076.402	1,200,611.219	1,200.611	1.201
localhost:02:01-Ethernet	12,102	12,128	100.004	97.504	18,192,000	146,255.811	1,170,046.486	1,170.046	1.170
localhost:02:02-Ethernet	12,103	12,091	100.000	99.000	18,136,500	148,500.337	1,188,002.697	1,188.003	1.188
localhost:02:03-Ethernet	12,103	12,080	100.006	100.006	18,120,000	150,009.212	1,200,073.697	1,200.074	1.200
localhost:02:04-Ethernet	12,101	12,072	99.986	99.486	18,108,000	149,229.745	1,193,837.960	1,193.838	1.194
localhost:02:05-Ethernet	12,101	12,125	100.005	98.004	18,187,500	147,006.670	1,176,053.358	1,176.053	1.176
localhost:02:06-Ethernet	12,103	12,084	100.010	98.010	18,126,000	147,015.107	1,176,120.858	1,176.121	1.176
localhost:02:07-Ethernet	11,902	11,919	99.979	98.980	17,878,500	148,469.279	1,187,754.230	1,187.754	1.188
localhost:02:08-Ethernet	12,102	12,122	99.996	98.496	18,183,000	147,744.653	1,181,957.225	1,181.957	1.182
localhost:02:09-Ethernet	12,102	12,097	100.008	99.508	18,145,500	149,262.471	1,194,099.767	1,194.100	1.194
localhost:02:10-Ethernet	12,102	12,069	100.001	100.497	17,209,500	150,745.126	1,205,961.011	1,205.961	1.206

Figure 193. Statistics

Test Variables

Once the basic test has been set up, you can explore other, more complex, test cases. For example, you could use IxNetwork's Integrated Test Facility to run the RFC2544 test suite.

Result Analysis

You should note that latency measurements between PCs and/or VMs will be less accurate than hardware-based tests because the clocks cannot be accurately synchronized as they would be in a dedicated hardware system like Ixia's X-series chassis. In software-based environments, synchronization is limited to millisecond-level accuracy.

Test Case: Multicast Testing

Objective

In this test case, we will validate the basic multicast functionality of a virtual switch or distributed virtual switch. Like a standard L2 physical switch, the VMware virtual switch ‘snoops’ IGMP protocol queries to understand which VMs are subscribing to various multicast flows offered in the system. The vSwitch then replicates those flows only to the VMs that have requested membership in those groups. In this test, we will configure all VMs on all hosts to join a single multicast group and then send multicast traffic from a single port to ensure that it is being received by all hosts.

Setup

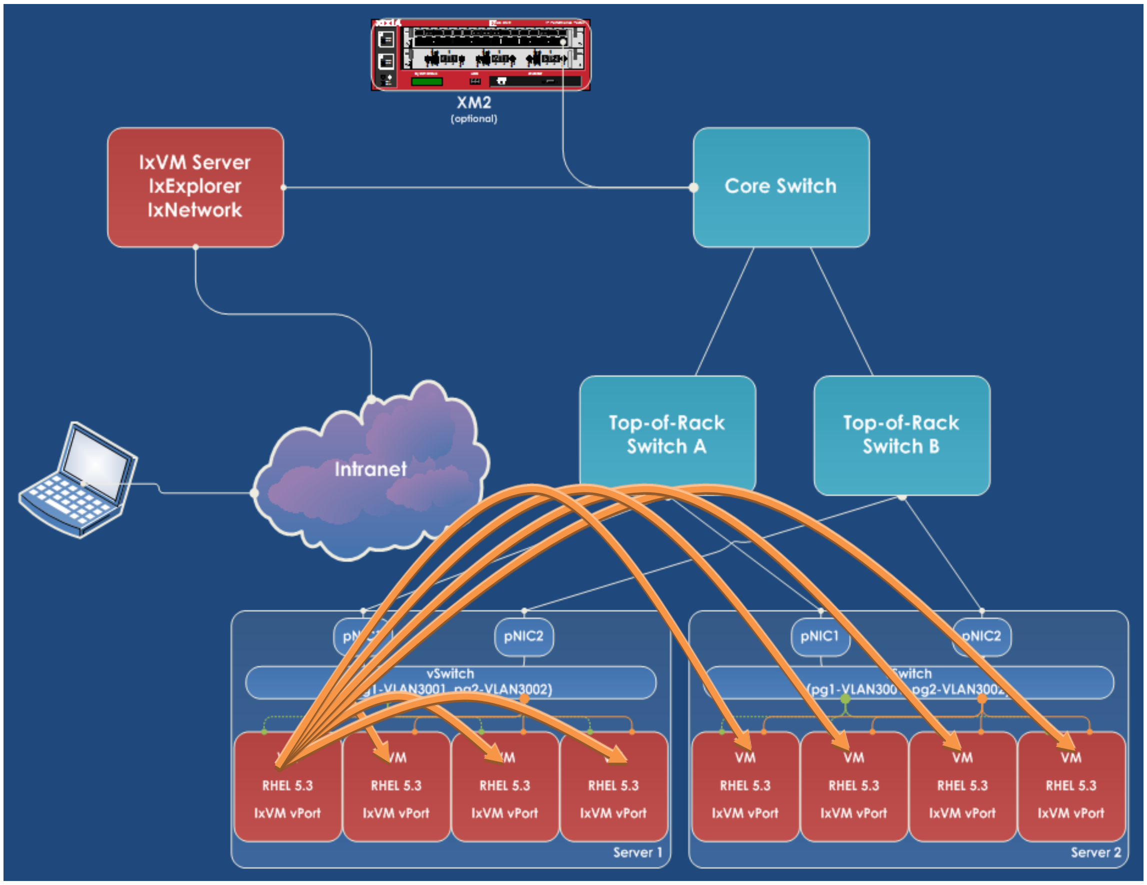

Figure 194. Multicast vSwitch testing topology.

Step-by-Step Instructions

1. Add the IxVM virtual ports in IxNetwork. You can add all ports in one step by selecting the chassis (*localhost* below) and then select the right-pointing arrow in the middle to add new ports and assign the selected ports.

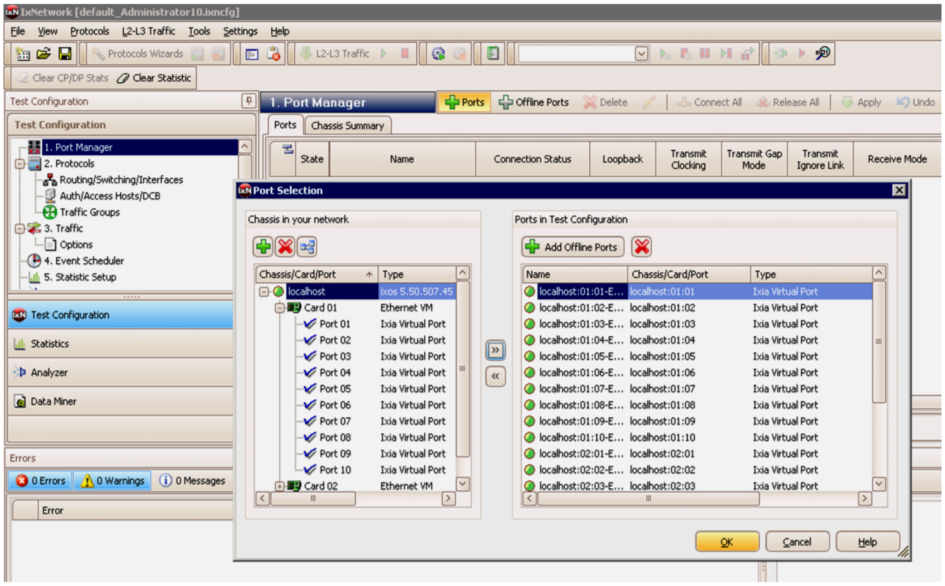

Port Selection

Chassis in your network

Chassis/Card/Port	Type
localhost	ixos 5.50.507.45
Card 01	Ethernet VM
Port 01	Ixia Virtual Port
Port 02	Ixia Virtual Port
Port 03	Ixia Virtual Port
Port 04	Ixia Virtual Port
Port 05	Ixia Virtual Port
Port 06	Ixia Virtual Port
Port 07	Ixia Virtual Port
Port 08	Ixia Virtual Port
Port 09	Ixia Virtual Port
Port 10	Ixia Virtual Port
Card 02	Ethernet VM

Ports in Test Configuration

Name	Chassis/Card/Port	Type
localhost:01:01-E...	localhost:01:01	Ixia Virtual Port
localhost:01:02-E...	localhost:01:02	Ixia Virtual Port
localhost:01:03-E...	localhost:01:03	Ixia Virtual Port
localhost:01:04-E...	localhost:01:04	Ixia Virtual Port
localhost:01:05-E...	localhost:01:05	Ixia Virtual Port
localhost:01:06-E...	localhost:01:06	Ixia Virtual Port
localhost:01:07-E...	localhost:01:07	Ixia Virtual Port
localhost:01:08-E...	localhost:01:08	Ixia Virtual Port
localhost:01:09-E...	localhost:01:09	Ixia Virtual Port
localhost:01:10-E...	localhost:01:10	Ixia Virtual Port
localhost:02:01-E...	localhost:02:01	Ixia Virtual Port
localhost:02:02-E...	localhost:02:02	Ixia Virtual Port
localhost:02:03-E...	localhost:02:03	Ixia Virtual Port

Figure 195. Adding Ports

2. Enable the **ARP** and **IGMP** protocols for all ports in the test. You can select the entire column by clicking on the column header and then enable all ports by right-clicking and selecting **Enable Selection**.

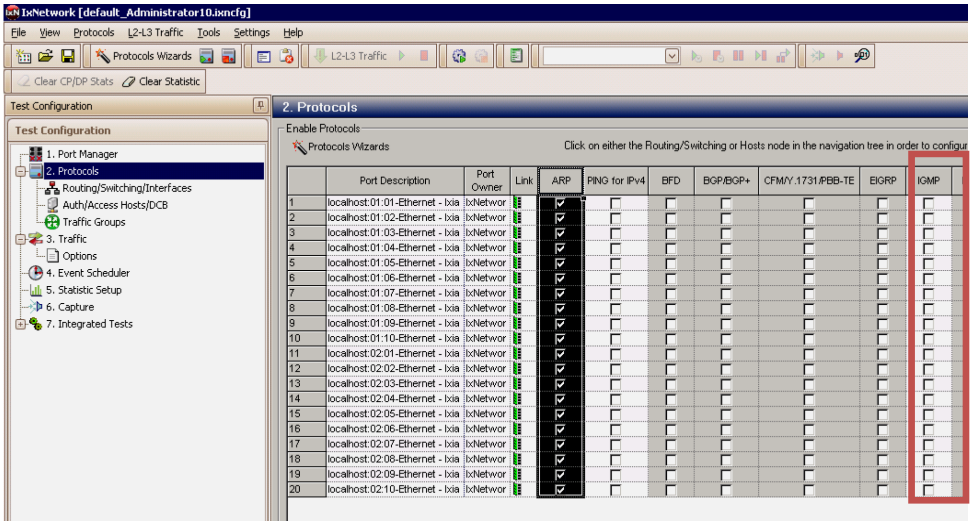

2. Protocols

Enable Protocols

Click on either the Routing/Switching or Hosts node in the navigation tree in order to configur

	Port Description	Port Owner	Link	ARP	PING for IPv4	BFD	BGP/BGP+	CFM/Y.1731/PBB-TE	EIGRP	IGMP
1	localhost:01:01-Ethernet - Ixia	IxNetwor		☑	□	□	□	□	□	□
2	localhost:01:02-Ethernet - Ixia	IxNetwor		☑	□	□	□	□	□	□
3	localhost:01:03-Ethernet - Ixia	IxNetwor		☑	□	□	□	□	□	□
4	localhost:01:04-Ethernet - Ixia	IxNetwor		☑	□	□	□	□	□	□
5	localhost:01:05-Ethernet - Ixia	IxNetwor		☑	□	□	□	□	□	□
6	localhost:01:06-Ethernet - Ixia	IxNetwor		☑	□	□	□	□	□	□
7	localhost:01:07-Ethernet - Ixia	IxNetwor		☑	□	□	□	□	□	□
8	localhost:01:08-Ethernet - Ixia	IxNetwor		☑	□	□	□	□	□	□
9	localhost:01:09-Ethernet - Ixia	IxNetwor		☑	□	□	□	□	□	□
10	localhost:01:10-Ethernet - Ixia	IxNetwor		☑	□	□	□	□	□	□
11	localhost:02:01-Ethernet - Ixia	IxNetwor		☑	□	□	□	□	□	□
12	localhost:02:02-Ethernet - Ixia	IxNetwor		☑	□	□	□	□	□	□
13	localhost:02:03-Ethernet - Ixia	IxNetwor		☑	□	□	□	□	□	□
14	localhost:02:04-Ethernet - Ixia	IxNetwor		☑	□	□	□	□	□	□
15	localhost:02:05-Ethernet - Ixia	IxNetwor		☑	□	□	□	□	□	□
16	localhost:02:06-Ethernet - Ixia	IxNetwor		☑	□	□	□	□	□	□
17	localhost:02:07-Ethernet - Ixia	IxNetwor		☑	□	□	□	□	□	□
18	localhost:02:08-Ethernet - Ixia	IxNetwor		☑	□	□	□	□	□	□
19	localhost:02:09-Ethernet - Ixia	IxNetwor		☑	□	□	□	□	□	□
20	localhost:02:10-Ethernet - Ixia	IxNetwor		☑	□	□	□	□	□	□

Figure 196. Enabling ARP and IGMP

3. Create IP protocol interfaces on each virtual port. A very simple IP scheme for a simple L2-3 network is shown in the following figure. To configure these addresses, first enter **1.1.1.1** as the IP address for the first row. Next, select the entire IP address column and right-click to choose **Increment**. Next, enter **1.1.1.1** in the *Gateway* column and right-click to choose **Same**. After that is complete, make sure to change the gateway in the first row to **1.1.1.2** because it cannot act as its own gateway. Finally, enable all the interfaces.

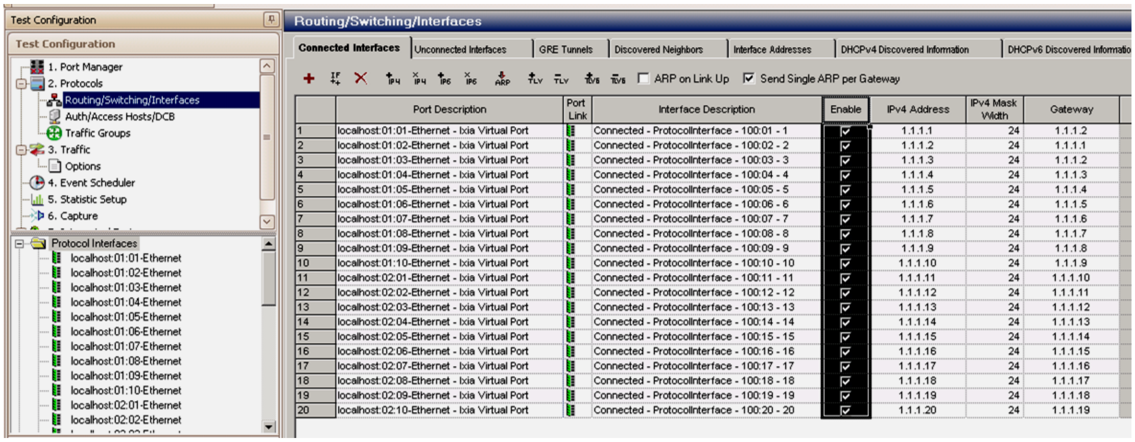

	Port Description	Port Link	Interface Description	Enable	IPv4 Address	IPv4 Mask Width	Gateway
1	localhost:01:01-Ethernet - Ixia Virtual Port		Connected - ProtocolInterface - 100:01 - 1	☑	1.1.1.1	24	1.1.1.2
2	localhost:01:02-Ethernet - Ixia Virtual Port		Connected - ProtocolInterface - 100:02 - 2	☑	1.1.1.2	24	1.1.1.1
3	localhost:01:03-Ethernet - Ixia Virtual Port		Connected - ProtocolInterface - 100:03 - 3	☑	1.1.1.3	24	1.1.1.2
4	localhost:01:04-Ethernet - Ixia Virtual Port		Connected - ProtocolInterface - 100:04 - 4	☑	1.1.1.4	24	1.1.1.3
5	localhost:01:05-Ethernet - Ixia Virtual Port		Connected - ProtocolInterface - 100:05 - 5	☑	1.1.1.5	24	1.1.1.4
6	localhost:01:06-Ethernet - Ixia Virtual Port		Connected - ProtocolInterface - 100:06 - 6	☑	1.1.1.6	24	1.1.1.5
7	localhost:01:07-Ethernet - Ixia Virtual Port		Connected - ProtocolInterface - 100:07 - 7	☑	1.1.1.7	24	1.1.1.6
8	localhost:01:08-Ethernet - Ixia Virtual Port		Connected - ProtocolInterface - 100:08 - 8	☑	1.1.1.8	24	1.1.1.7
9	localhost:01:09-Ethernet - Ixia Virtual Port		Connected - ProtocolInterface - 100:09 - 9	☑	1.1.1.9	24	1.1.1.8
10	localhost:01:10-Ethernet - Ixia Virtual Port		Connected - ProtocolInterface - 100:10 - 10	☑	1.1.1.10	24	1.1.1.9
11	localhost:02:01-Ethernet - Ixia Virtual Port		Connected - ProtocolInterface - 100:11 - 11	☑	1.1.1.11	24	1.1.1.10
12	localhost:02:02-Ethernet - Ixia Virtual Port		Connected - ProtocolInterface - 100:12 - 12	☑	1.1.1.12	24	1.1.1.11
13	localhost:02:03-Ethernet - Ixia Virtual Port		Connected - ProtocolInterface - 100:13 - 13	☑	1.1.1.13	24	1.1.1.12
14	localhost:02:04-Ethernet - Ixia Virtual Port		Connected - ProtocolInterface - 100:14 - 14	☑	1.1.1.14	24	1.1.1.13
15	localhost:02:05-Ethernet - Ixia Virtual Port		Connected - ProtocolInterface - 100:15 - 15	☑	1.1.1.15	24	1.1.1.14
16	localhost:02:06-Ethernet - Ixia Virtual Port		Connected - ProtocolInterface - 100:16 - 16	☑	1.1.1.16	24	1.1.1.15
17	localhost:02:07-Ethernet - Ixia Virtual Port		Connected - ProtocolInterface - 100:17 - 17	☑	1.1.1.17	24	1.1.1.16
18	localhost:02:08-Ethernet - Ixia Virtual Port		Connected - ProtocolInterface - 100:18 - 18	☑	1.1.1.18	24	1.1.1.17
19	localhost:02:09-Ethernet - Ixia Virtual Port		Connected - ProtocolInterface - 100:19 - 19	☑	1.1.1.19	24	1.1.1.18
20	localhost:02:10-Ethernet - Ixia Virtual Port		Connected - ProtocolInterface - 100:20 - 20	☑	1.1.1.20	24	1.1.1.19

Figure 197. Enabling Interfaces

4. Next, switch to the *Routing/Switching/Protocols* → *IGMP* folder to configure IGMP. On the first row of the table, use the dropdown in the *Protocol Interfaces* field to choose the previously-configured protocol interface. Select the entire column and use **Same** to automatically select the first protocol interface for each row. Using the same grid management techniques, select the correct IGMP *version* and one (1) *Group Range* for each port. Finally, enable **IGMP** on the interfaces created in step 3.

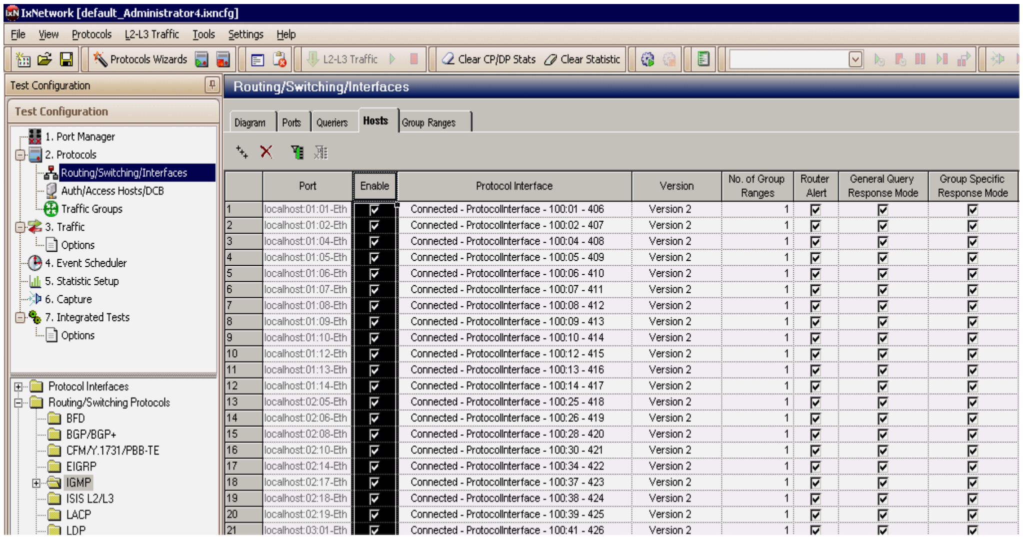

	Port	Enable	Protocol Interface	Version	No. of Group Ranges	Router Alert	General Query Response Mode	Group Specific Response Mode
1	localhost:01:01-Eth	☑	Connected - ProtocolInterface - 100:01 - 406	Version 2	1	☑	☑	☑
2	localhost:01:02-Eth	☑	Connected - ProtocolInterface - 100:02 - 407	Version 2	1	☑	☑	☑
3	localhost:01:04-Eth	☑	Connected - ProtocolInterface - 100:04 - 408	Version 2	1	☑	☑	☑
4	localhost:01:05-Eth	☑	Connected - ProtocolInterface - 100:05 - 409	Version 2	1	☑	☑	☑
5	localhost:01:06-Eth	☑	Connected - ProtocolInterface - 100:06 - 410	Version 2	1	☑	☑	☑
6	localhost:01:07-Eth	☑	Connected - ProtocolInterface - 100:07 - 411	Version 2	1	☑	☑	☑
7	localhost:01:08-Eth	☑	Connected - ProtocolInterface - 100:08 - 412	Version 2	1	☑	☑	☑
8	localhost:01:09-Eth	☑	Connected - ProtocolInterface - 100:09 - 413	Version 2	1	☑	☑	☑
9	localhost:01:10-Eth	☑	Connected - ProtocolInterface - 100:10 - 414	Version 2	1	☑	☑	☑
10	localhost:01:12-Eth	☑	Connected - ProtocolInterface - 100:12 - 415	Version 2	1	☑	☑	☑
11	localhost:01:13-Eth	☑	Connected - ProtocolInterface - 100:13 - 416	Version 2	1	☑	☑	☑
12	localhost:01:14-Eth	☑	Connected - ProtocolInterface - 100:14 - 417	Version 2	1	☑	☑	☑
13	localhost:02:05-Eth	☑	Connected - ProtocolInterface - 100:25 - 418	Version 2	1	☑	☑	☑
14	localhost:02:06-Eth	☑	Connected - ProtocolInterface - 100:26 - 419	Version 2	1	☑	☑	☑
15	localhost:02:08-Eth	☑	Connected - ProtocolInterface - 100:28 - 420	Version 2	1	☑	☑	☑
16	localhost:02:10-Eth	☑	Connected - ProtocolInterface - 100:30 - 421	Version 2	1	☑	☑	☑
17	localhost:02:14-Eth	☑	Connected - ProtocolInterface - 100:34 - 422	Version 2	1	☑	☑	☑
18	localhost:02:17-Eth	☑	Connected - ProtocolInterface - 100:37 - 423	Version 2	1	☑	☑	☑
19	localhost:02:18-Eth	☑	Connected - ProtocolInterface - 100:38 - 424	Version 2	1	☑	☑	☑
20	localhost:02:19-Eth	☑	Connected - ProtocolInterface - 100:39 - 425	Version 2	1	☑	☑	☑
21	localhost:03:01-Eth	☑	Connected - ProtocolInterface - 100:41 - 426	Version 2	1	☑	☑	☑

Figure 198. Enabling Ports

5. Under the *Group Ranges* tab, enter a single *Group ID* of **229.1.1.1** and enable it for all rows.

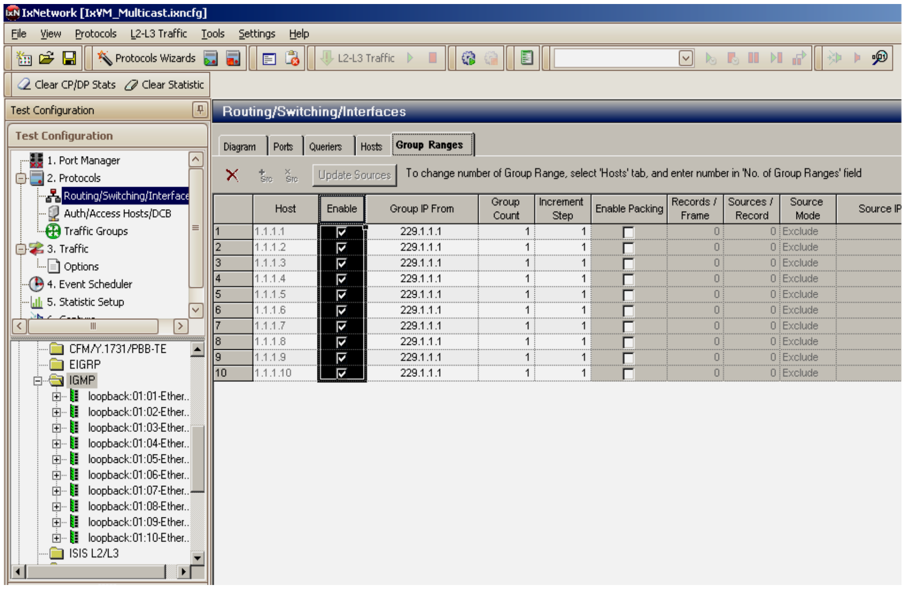

To change number of Group Range, select 'Hosts' tab, and enter number in 'No. of Group Ranges' field

	Host	Enable	Group IP From	Group Count	Increment Step	Enable Packing	Records / Frame	Sources / Record	Source Mode	Source IP
1	1.1.1.1	☑	229.1.1.1	1	1	□	0	0	Exclude	
2	1.1.1.2	☑	229.1.1.1	1	1	□	0	0	Exclude	
3	1.1.1.3	☑	229.1.1.1	1	1	□	0	0	Exclude	
4	1.1.1.4	☑	229.1.1.1	1	1	□	0	0	Exclude	
5	1.1.1.5	☑	229.1.1.1	1	1	□	0	0	Exclude	
6	1.1.1.6	☑	229.1.1.1	1	1	□	0	0	Exclude	
7	1.1.1.7	☑	229.1.1.1	1	1	□	0	0	Exclude	
8	1.1.1.8	☑	229.1.1.1	1	1	□	0	0	Exclude	
9	1.1.1.9	☑	229.1.1.1	1	1	□	0	0	Exclude	
10	1.1.1.10	☑	229.1.1.1	1	1	□	0	0	Exclude	

Figure 199. Enabling Group Ranges

6. Move to step 3, *Traffic*, in the Test Configuration and start the *Basic Traffic Wizard by clicking* Basic Wizard. Use the Traffic Wizard to create a many-to-many traffic topology.

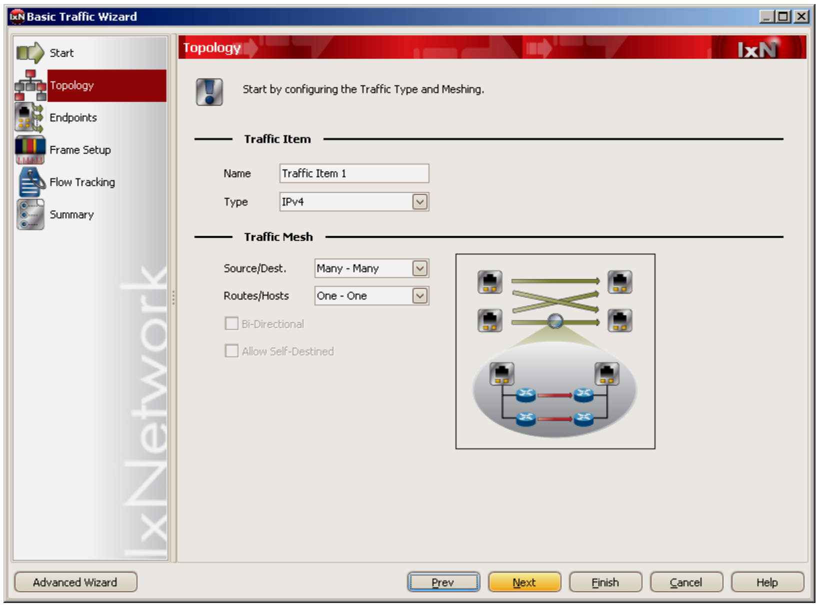

Start by configuring the Traffic Type and Meshing.

Traffic Item

Name: Traffic Item 1

Type: IPv4

Traffic Mesh

Source/Dest.: Many - Many

Routes/Hosts: One - One

□ Bi-Directional

□ Allow Self-Destined

Figure 200. Many-to-Many Topology

7. For this test, select the first port on the left hand column and then choose the multicast icon on the right side (). This will bring up the *Multicast Endpoint Selection* window where you can select the group range configured previously.

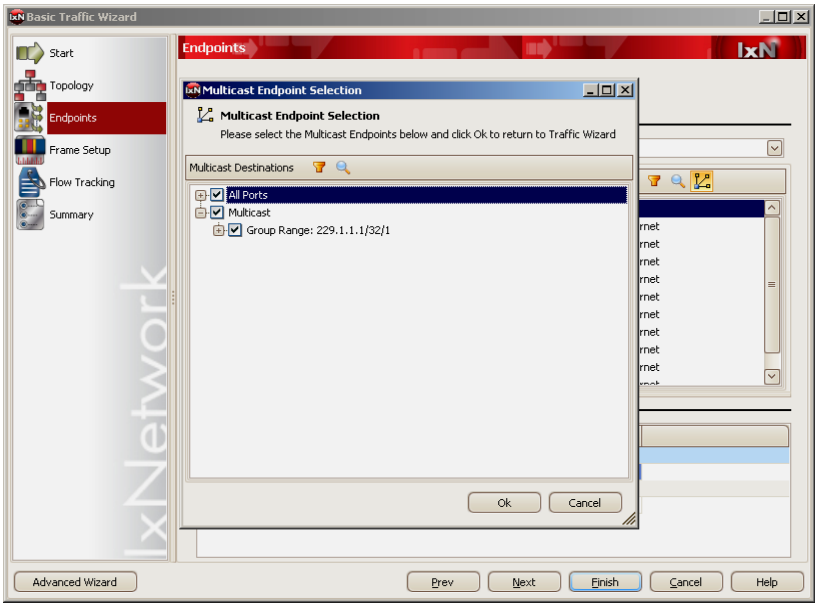

Multicast Endpoint Selection

Please select the Multicast Endpoints below and click Ok to return to Traffic Wizard

Multicast Destinations

- ☑ All Ports
- ☑ Multicast
  - ☑ Group Range: 229.1.1.1/32/1

Figure 201. Selecting Multicast Endpoints

8. Confirm that the selecting the multicast group range automatically selected all ports joined to that range.

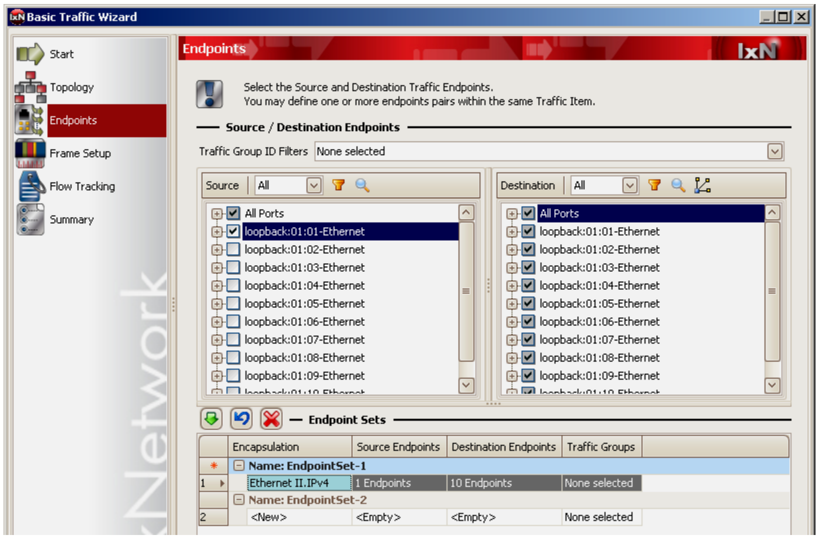

Endpoints

Select the Source and Destination Traffic Endpoints.
You may define one or more endpoints pairs within the same Traffic Item.

Source / Destination Endpoints

Traffic Group ID Filters: None selected

Source: All

- ☑ All Ports
- ☑ loopback:01:01-Ethernet
- □ loopback:01:02-Ethernet
- □ loopback:01:03-Ethernet
- □ loopback:01:04-Ethernet
- □ loopback:01:05-Ethernet
- □ loopback:01:06-Ethernet
- □ loopback:01:07-Ethernet
- □ loopback:01:08-Ethernet
- □ loopback:01:09-Ethernet
- □ loopback:01:10-Ethernet

Destination: All

- ☑ All Ports
- ☑ loopback:01:01-Ethernet
- ☑ loopback:01:02-Ethernet
- ☑ loopback:01:03-Ethernet
- ☑ loopback:01:04-Ethernet
- ☑ loopback:01:05-Ethernet
- ☑ loopback:01:06-Ethernet
- ☑ loopback:01:07-Ethernet
- ☑ loopback:01:08-Ethernet
- ☑ loopback:01:09-Ethernet
- ☑ loopback:01:10-Ethernet

Endpoint Sets

	Encapsulation	Source Endpoints	Destination Endpoints	Traffic Groups
*	Name: EndpointSet-1			
1	Ethernet II.IPv4	1 Endpoints	10 Endpoints	None selected
	Name: EndpointSet-2			
2	<New>	<Empty>	<Empty>	None selected

Figure 202. Confirming Port Selection

9. Configure the frame size and rate. In this case, 1000-byte frames will be sent at 100 frames per second. In VM-to-VM testing, % *Line Rate* does not apply because the vNIC does not have a fixed throughput.

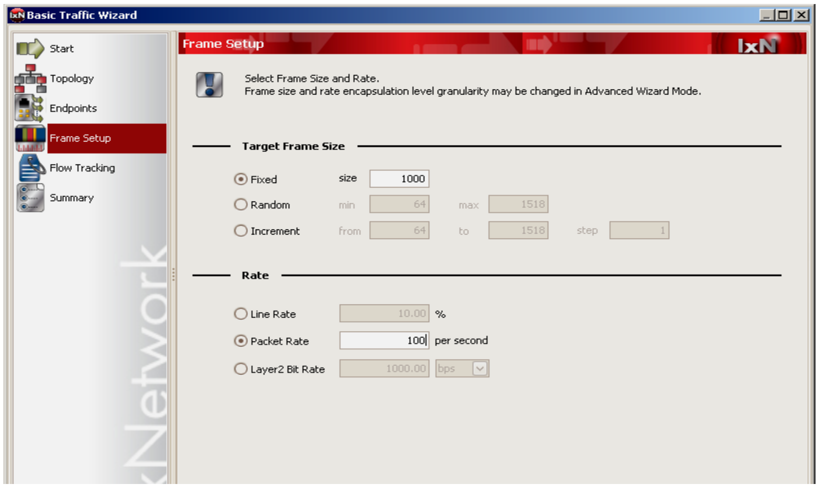

Frame Setup

Select Frame Size and Rate.
Frame size and rate encapsulation level granularity may be changed in Advanced Wizard Mode.

Target Frame Size

● Fixed size 1000
○ Random min 64 max 1518
○ Increment from 64 to 1518 step 1

Rate

○ Line Rate 10.00 %
● Packet Rate 100 per second
○ Layer2 Bit Rate 1000.00 bps

Figure 203. Frame Size and Rate

10. Select the necessary tracking fields to get the level of detail required in the statistics view. The options selected in the screenshot below will display the traffic stats aggregated by the source/destination port pairs and source and destination IP addresses.

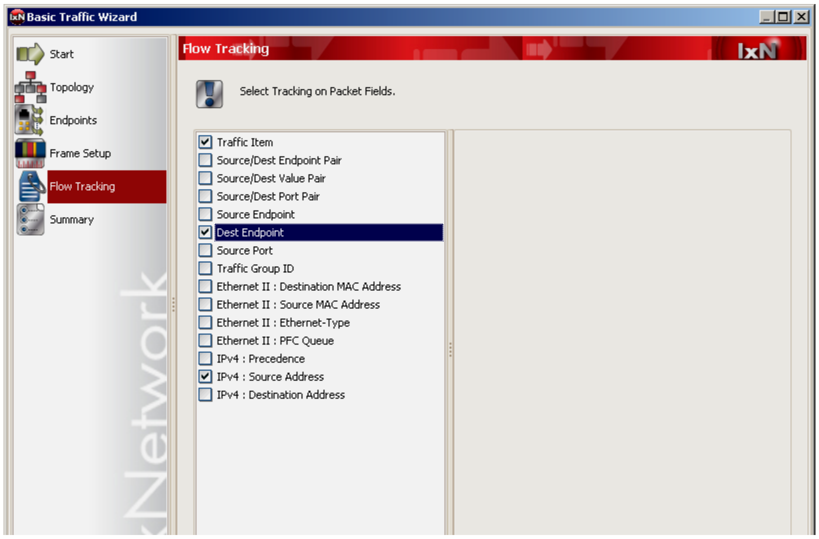

Flow Tracking

Select Tracking on Packet Fields.

☑ Traffic Item
☐ Source/Dest Endpoint Pair
☐ Source/Dest Value Pair
☐ Source/Dest Port Pair
☐ Source Endpoint
☑ Dest Endpoint
☐ Source Port
☐ Traffic Group ID
☐ Ethernet II : Destination MAC Address
☐ Ethernet II : Source MAC Address
☐ Ethernet II : Ethernet-Type
☐ Ethernet II : PFC Queue
☐ IPv4 : Precedence
☑ IPv4 : Source Address
☐ IPv4 : Destination Address

Figure 204. Flow Tracking

11. Finish the wizard and review the test configuration. Each rounded rectangle represents one of the ports used in the test. By selecting each port, you can see where its traffic is destined or originating from to validate the test topology.

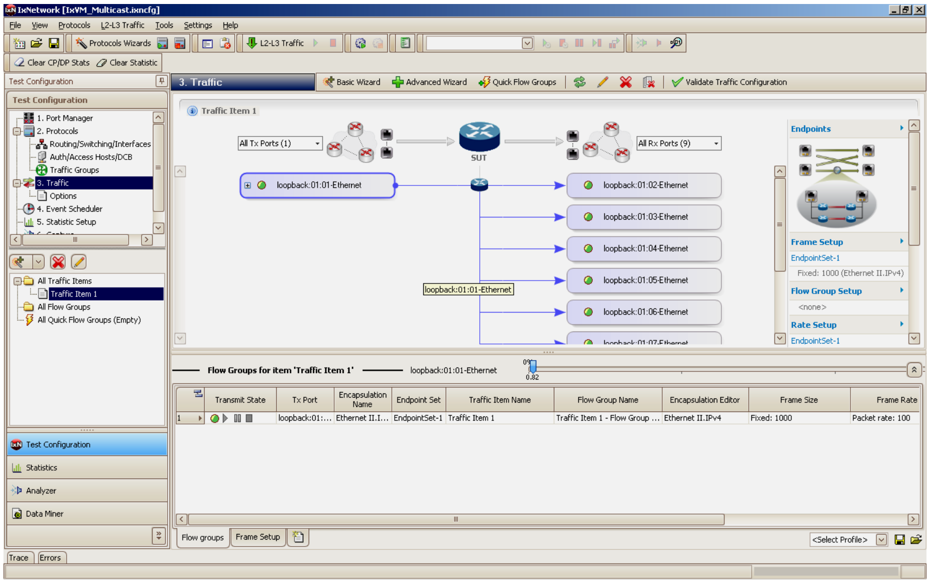

Flow Groups for item 'Traffic Item 1' loopback:01:01-Ethernet 0.82

	Transmit State	Tx Port	Encapsulation Name	Endpoint Set	Traffic Item Name	Flow Group Name	Encapsulation Editor	Frame Size	Frame Rate
1	● ▶ ⏸ ■	loopback:01:...	Ethernet II.I...	EndpointSet-1	Traffic Item 1	Traffic Item 1 - Flow Group ...	Ethernet II.IPv4	Fixed: 1000	Packet rate: 100

Figure 205. Test Topology

12. Click **Apply Traffic** (L2-L3 Traffic), and then click (▶) to run the test. Select **Statistics** to monitor the statistics in real-time. The key objective of this test is to confirm that all VM ports are receiving the multicast traffic properly.

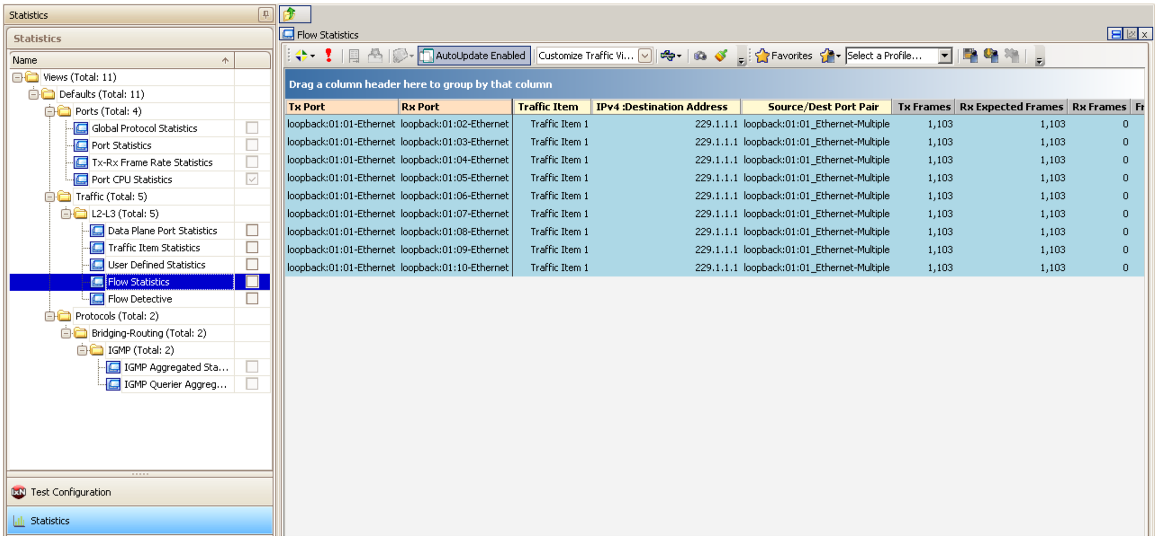

Tx Port	Rx Port	Traffic Item	IPv4 :Destination Address	Source/Dest Port Pair	Tx Frames	Rx Expected Frames	Rx Frames	Fr
loopback:01:01-Ethernet	loopback:01:02-Ethernet	Traffic Item 1	229.1.1.1	loopback:01:01_Ethernet-Multiple	1,103	1,103	0	
loopback:01:01-Ethernet	loopback:01:03-Ethernet	Traffic Item 1	229.1.1.1	loopback:01:01_Ethernet-Multiple	1,103	1,103	0	
loopback:01:01-Ethernet	loopback:01:04-Ethernet	Traffic Item 1	229.1.1.1	loopback:01:01_Ethernet-Multiple	1,103	1,103	0	
loopback:01:01-Ethernet	loopback:01:05-Ethernet	Traffic Item 1	229.1.1.1	loopback:01:01_Ethernet-Multiple	1,103	1,103	0	
loopback:01:01-Ethernet	loopback:01:06-Ethernet	Traffic Item 1	229.1.1.1	loopback:01:01_Ethernet-Multiple	1,103	1,103	0	
loopback:01:01-Ethernet	loopback:01:07-Ethernet	Traffic Item 1	229.1.1.1	loopback:01:01_Ethernet-Multiple	1,103	1,103	0	
loopback:01:01-Ethernet	loopback:01:08-Ethernet	Traffic Item 1	229.1.1.1	loopback:01:01_Ethernet-Multiple	1,103	1,103	0	
loopback:01:01-Ethernet	loopback:01:09-Ethernet	Traffic Item 1	229.1.1.1	loopback:01:01_Ethernet-Multiple	1,103	1,103	0	
loopback:01:01-Ethernet	loopback:01:10-Ethernet	Traffic Item 1	229.1.1.1	loopback:01:01_Ethernet-Multiple	1,103	1,103	0	

Figure 206. Statistics

Test Variables

Among the many possible variations of this test would be to use the chassis shown in the topology figure to generate the multicast traffic to verify that the VMs are able to join multicast groups external to the local L2 fabric.

Introduction to Network Virtulization Overlay Technology

Data center started with standalone physical servers. As demand starts to increase, additional units of servers are added to the racks. Traffic from the cluster of servers is aggregated into Top of the Rack (ToR) and/or End of Row (EoR) switches. The increasing demand and dependency on Internet resources such as web presence; social networking sites and video streaming services are driving the next evolution of the virtualized data center, thanks to the server virtualization. Server virtualization has caused a major paradigm shift, and delivered many benefits including:

- Enabling multi-tenancy, also known as cloud
- Elastic provisioning of resources, ability to quickly spin up machines
- Spanning across physical boundaries (racks, pods, and so on.)
- Isolation from other tenants

One key characteristic of server virtualization is the mobility of VMs from one server—this is known as VM migration—to another to provide:

- The ability to move VMs to address maintenance (software or hardware changes)
- The ability to address capacity, moving a VM to a server with more resources
- Disaster recovery – moving VMs away from a threat
- Cloud bursting – the ability to leverage resources in a another data center for temporary peaks

In today's data center designs, the IP addressing and VLAN are often assigned to servers based on their physical rack locations. As virtualization became prominent in data center network designs, the location based addressing limited the mobile characteristics of virtual machines. Another complication resulting from the increasing deployment of VMs is the explosive growth in the number of MAC addresses in the network as one physical server can now have 100s MAC running in a hypervisor. This growth troubled traditional layer 2 network designs, because the need to learn station MAC addresses end-to-end across the network implied an extremely large Filtering Database (that is, MAC table) across all switches. As a result, in March 2012, the IETF began work called 'Network Virtualization over Layer 3 (NVO3)'. As the IEEE NVO3 working group set out to address the issues created by server virtualization, several key problems needed to be overcome. First is the significant increase in the number of MAC addresses in the network, because each VM is addressed with its own MAC and a virtualized server can have as many as twenty VMs or possibly more. This changes a rack of twenty servers from twenty MACs to four hundred or potentially more per rack.

VLAN scale represents a second major challenge to be overcome. A data center may need thousands of VLANs to partition the traffic according to the specific group that the VM belongs to. Standard VLAN addressing is limited to 4094. Virtual machines cannot be moved across layer 3 boundaries, so if routing used to scale the datacenter can place limitations on VM mobility.

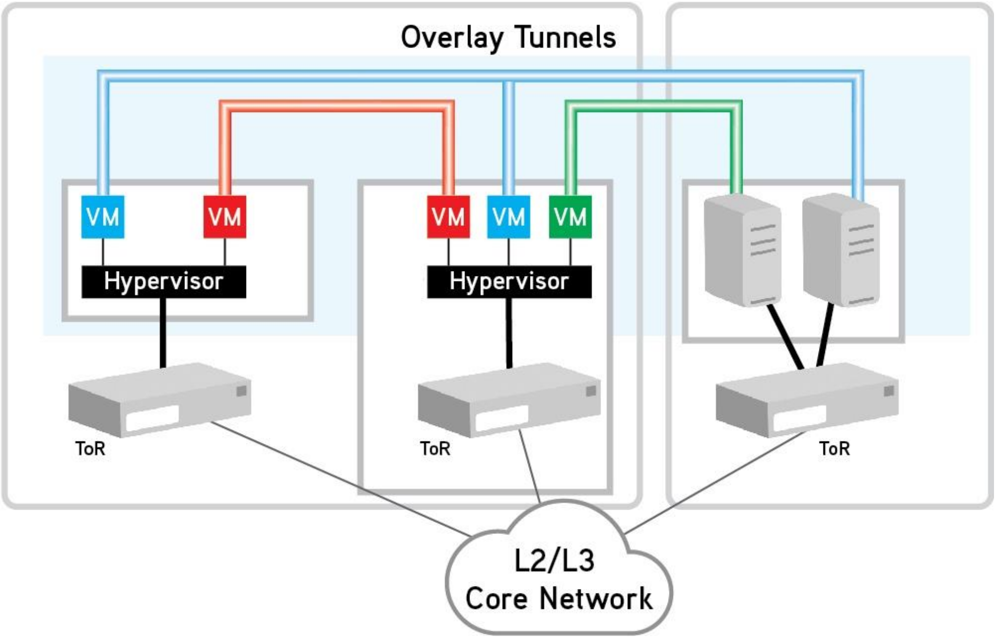

A third issue, as previously mentioned, is Spanning Tree Protocol (STP) creating a loop free topology by blocking paths and creating tree based forwarding paths, both of which are non-optimal. And finally, in a true multi-tenancy environment, tenants need their own isolated network domains including overlapping address space. There are some 802.1 (layer 2) based methods to work around some of these problems, including using Provider Back Bone Bridging (PBB). Yet many operators have already moved to layer 3 within the data center and need a solution.

This working group has addressed the problems and developed protocols among others as solutions to this problem. One such solution called **Virtual eXtensible Local Area Network (VXLAN)** provides an encapsulation scheme to address the challenges described above.

VXLAN

Virtual eXtensible Local Area Network (VXLAN) is an IP-based overlay tunneling technology. A VXLAN tunnel is formed between VXLAN Tunnel Endpoints (VTEPs), and is identified by the VXLAN Network Identifier (VNI). The VTEP is generally a virtual switch in the hypervisor, but can also be a VXLAN-capable physical switch that provides access between VXLAN domains and non-VXLAN domains. The VNI is a 24-bit field, which means up to 16 million unique VXLAN Segments can be supported. A VXLAN Segment signifies a private endpoint network, which can be a tenant, or one of several private networks of a tenant. VXLAN supports private tenant communication using the tunnels built on a per VXLAN Segment basis, and provisions are in place for unicast, multicast and broadcast traffic.

VXLAN stack and header:

Outer Ethernet Header
(optional) VLAN Header
Outer IP Header
Outer UDP Header
VXLAN header
Inner Ethernet Header
(optional) VLAN header
Payload (followed by new Frame Check Sequence (FCS)

Protocol stack with VXLAN

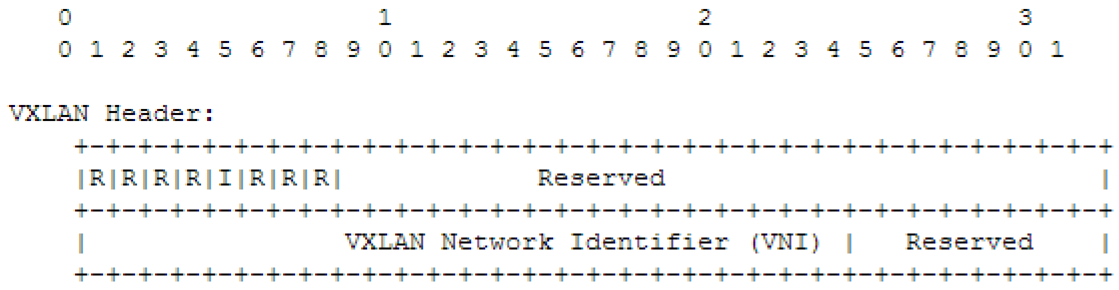

```
   0                   1                   2                   3
   0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1

VXLAN Header:
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |R|R|R|R|I|R|R|R|           Reserved                            |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                VXLAN Network Identifier (VNI) |   Reserved    |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

VXLAN Header Format

When the VM transmits data, the packet is sent to the VTEP, to which the transmitting VM is attached. The ingress VTEP identifies the destination address of the packet, and sets the proper outer destination address, which can identify a unicast endpoint or a multicast group. The receiving VTEP(s) reviews the VNI, and decapsulates the packet for forwarding, if the VNI is valid. Once the packet is decapsulated, the egress VTEP forwards the original data packet to the proper endpoint(s) attached to it.

If the VM packet is unicast and the destination is known, the ingress VTEP selects the egress VTEP. The outer IP is the unicast destination IP of the egress VTEP. If the VM packet is unicast and the destination is unknown, or if the VM packet is multicast or broadcast, the ingress VTEP forwards the packet to all remote VTEPs of the VNI. The outer MAC and IP are set to identify a multicast group as done in traditional layer 3 networks.

In summary, VXLAN offers great advantages in a data center network. VXLAN solves some of the biggest limitations data center networks face today, like multi-tenancy, addressing freedom, and mobility of VMs. It also offers a scalable tunneling technology that runs over existing layer 2 and layer 3 infrastructures. But along with these benefits, there are also limitation/challenges for VXLAN:

- VXLAN is an overlay network protocol. This means, it has no knowledge of the underlying network in terms of traffic congestion and QoS. It has to rely on the underlying routing and switch protocol to send the packets across one VTEP to other VTEP reliably.
- VXLAN uses a multicast protocol to learn unknown destination, as well as MAC learning frames. This impacts the performance of the existing physical underlay network.
- The VM end-point, once configured on a VXLAN segment, cannot communicate with any other non-VXLAN network.
- While scalability can be easily achieved in the physical VXLAN gateway, the same level of functionality and scalability must also be implemented in the vSwitch component of the Hypervisor to provide the VXLAN reachability to the VMs.

Test Case: Validate Traffic Forwarding over Virtual eXtensible Local Area Network (VXLAN) Gateway device

Objective

This test case, validates the core functionality of a VXLAN Gateway and its ability to encapsulate and decapsulate host traffic. Also, it measures, if DUT is correctly forwarding the traffic for each VXLAN segment, that is, VNI.

Setup

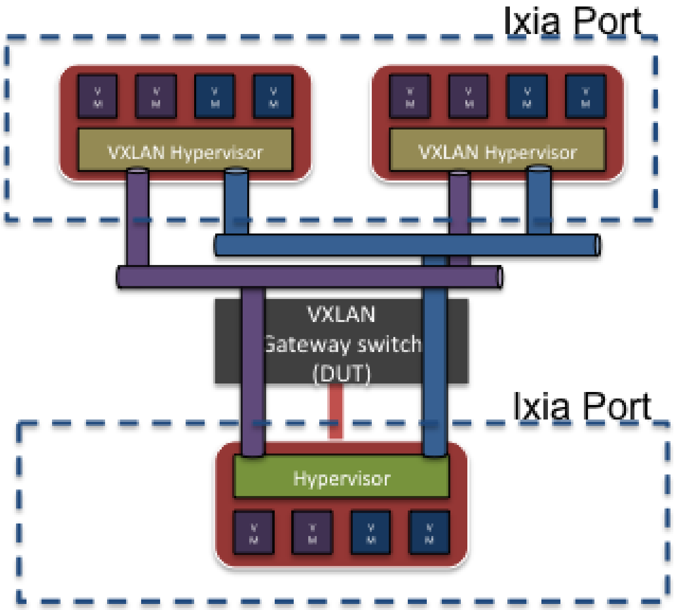

Figure 207. Data Center ToR Switch running VXLAN

Ixia Port emulates VTEP Gateway Switch and VM host. VMs are configured as IPv4 host to run traffic over the VXLAN Gateway DUT.

Step-by-Step Instructions

1. Click **Add Ports** to open the **Port Selection** Window. Click **Add Chassis** and enter the IP Address or name for your IXIA chassis. Click **OK** to accept. Expand chassis and cards and select two test ports you want to use in this test. Click **Add ports** and then **OK** to add the ports to your test configuration.

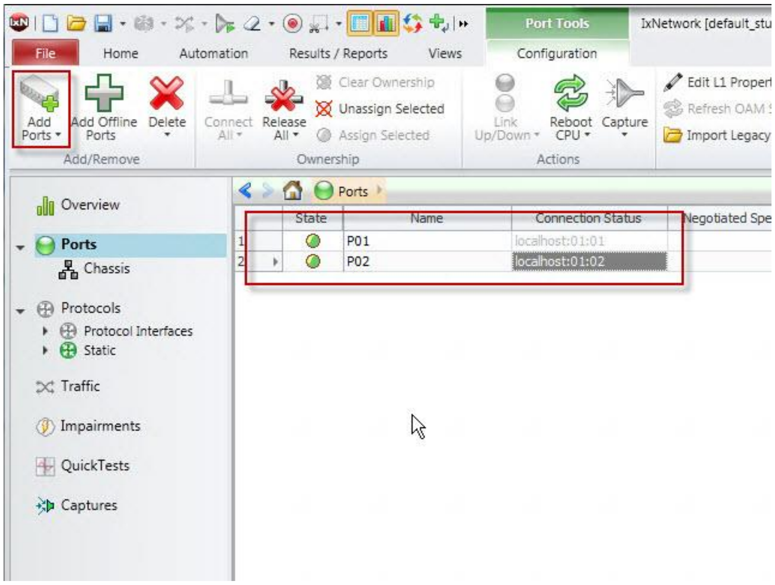

	State	Name	Connection Status	Negotiated Spe
1	●	P01	localhost:01:01	
2	●	P02	localhost:01:02	

Figure 208. Adding ports to a new configuration

2. Click the **Scenario** view, and then click **Add Topolgy** to open the **Protocol Wizard** dialog. Select **Port P01** and **Append Ports**. Click **Next**.

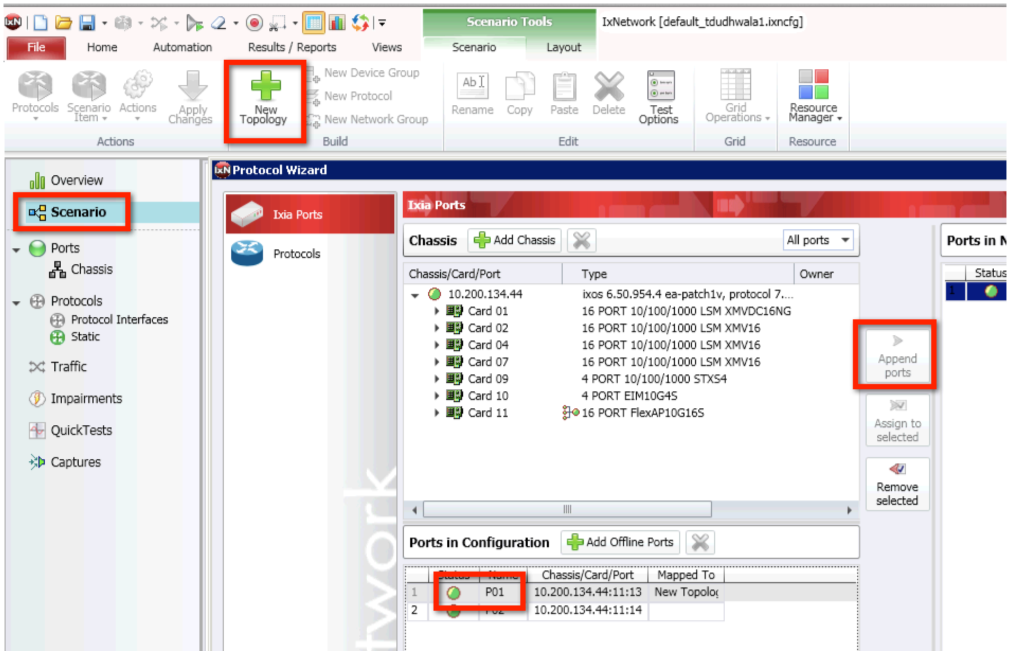

Chassis/Card/Port	Type	Owner
10.200.134.44	ixos 6.50.954.4 ea-patch1v, protocol 7....	
Card 01	16 PORT 10/100/1000 LSM XMVDC16NG	
Card 02	16 PORT 10/100/1000 LSM XMV16	
Card 04	16 PORT 10/100/1000 LSM XMV16	
Card 07	16 PORT 10/100/1000 LSM XMV16	
Card 09	4 PORT 10/100/1000 STXS4	
Card 10	4 PORT EIM10G4S	
Card 11	16 PORT FlexAP10G16S	

	Status	Name	Chassis/Card/Port	Mapped To
1	●	P01	10.200.134.44:11:13	New Topolo
2	●	P02	10.200.134.44:11:14	

Figure 209. Selecting VXLAN Port in Protocol Wizard

3. Select the **VXLAN** checkbox and click **Finish**. This action adds 10x VXLAN VTEP Gateways and 10x IPv4 Hosts behind each gateway, total 100 hosts.

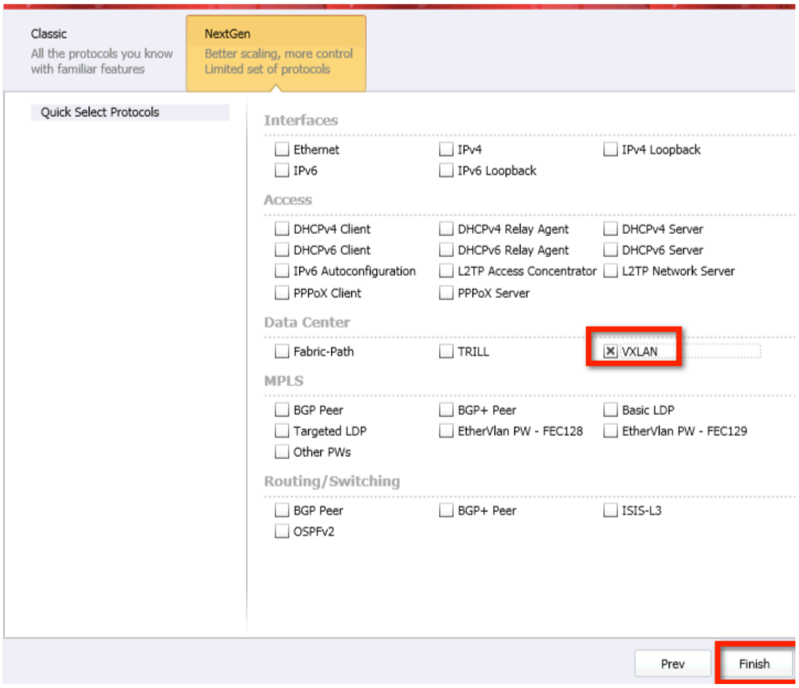

Classic
All the protocols you know with familiar features

NextGen
Better scaling, more control
Limited set of protocols

Quick Select Protocols

Interfaces
□ Ethernet □ IPv4 □ IPv4 Loopback
□ IPv6 □ IPv6 Loopback

Access
□ DHCPv4 Client □ DHCPv4 Relay Agent □ DHCPv4 Server
□ DHCPv6 Client □ DHCPv6 Relay Agent □ DHCPv6 Server
□ IPv6 Autoconfiguration □ L2TP Access Concentrator □ L2TP Network Server
□ PPPoX Client □ PPPoX Server

Data Center
□ Fabric-Path □ TRILL ☒ VXLAN

MPLS
□ BGP Peer □ BGP+ Peer □ Basic LDP
□ Targeted LDP □ EtherVlan PW - FEC128 □ EtherVlan PW - FEC129
□ Other PWs

Routing/Switching
□ BGP Peer □ BGP+ Peer □ ISIS-L3
□ OSPFv2

Prev Finish

Figure 210. Adding VXLAN on Port P01

4. Click **10x** box in the **Scenario** window and change the value to **2**. This action changes the number of **VTEP devices** from 10 to 2. Click **OK**.

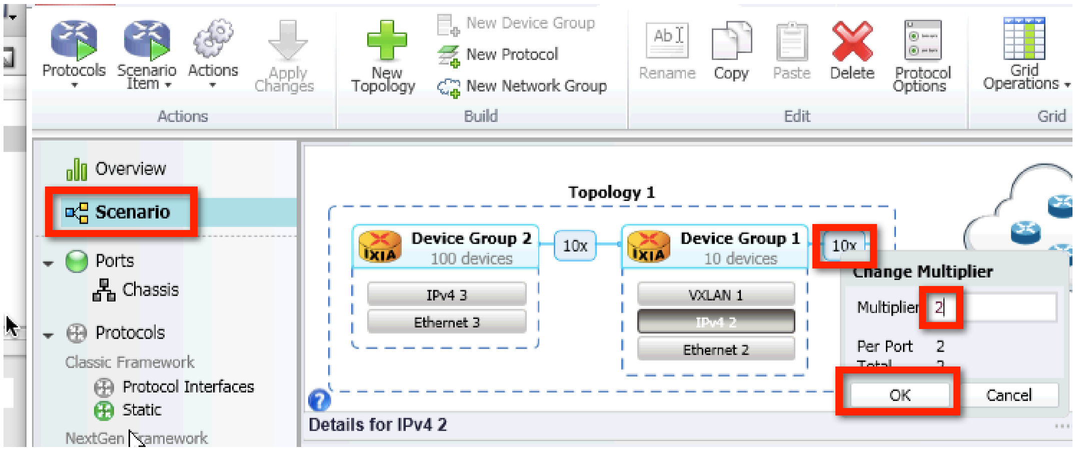

Change Multiplier
Multiplier 2
Per Port 2
Total 2
OK Cancel

Figure 211. Changin number of VTEP per port to 2

5. Click **VXLAN Device Group,** and then click **Custom Ratios** in the ribbon. Change the multiplier value between IP and VXLAN field to **2**. This action changes the number of **VNIs per VTEP** from 1 to 2, total 4 VNIs per port. Click **Close**.

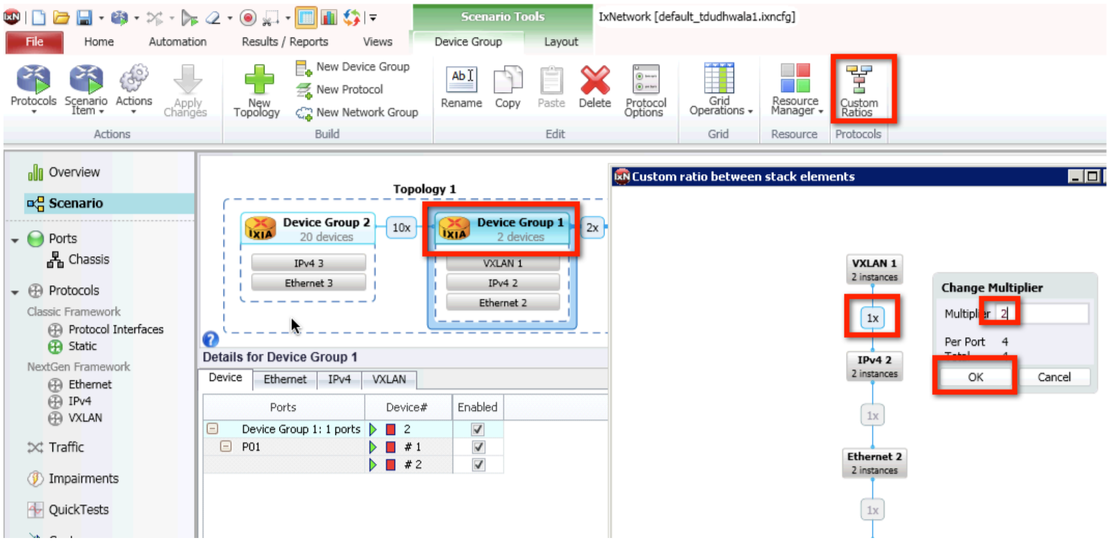

Figure 212. Chaning number of VNIs per VTEP to 2

6. Click **IPv4** in **VXLAN Device Group** to configure the **VTEP IP address**. Click the **Address** and enter **Start Value** of **101.1.0.1. Click OK.** Similarly, set the **Gateway IP** to **101.1.0.100.**

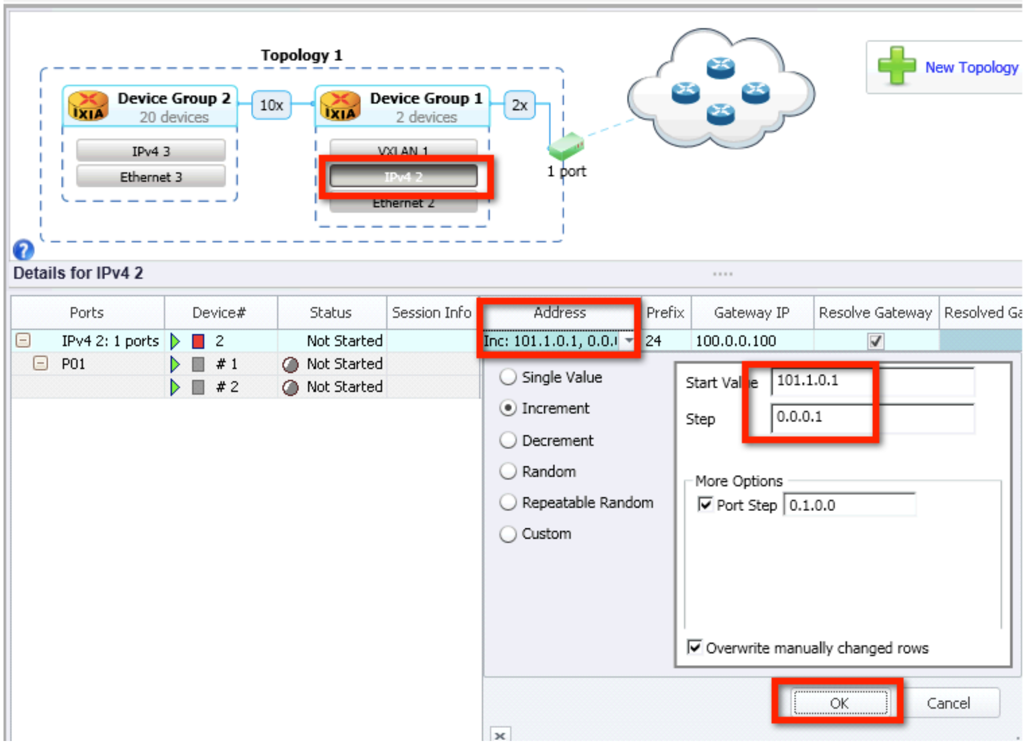

Ports	Device#	Status	Session Info	Address	Prefix	Gateway IP	Resolve Gateway	Resolved Ga
IPv4 2: 1 ports	2	Not Started		Inc: 101.1.0.1, 0.0.0	24	100.0.0.100	☑	
P01	# 1	Not Started						
	# 2	Not Started						

● Increment

Start Value 101.1.0.1

Step 0.0.0.1

☑ Port Step 0.1.0.0

☑ Overwrite manually changed rows

Figure 213. Setting VTEP IP address

7. Click **VXLAN** in **Scenario** View to change the VTEP VNI configuration. Click **VNI** and configure the **Start Value** to 1000 in Increment. Click **OK**.

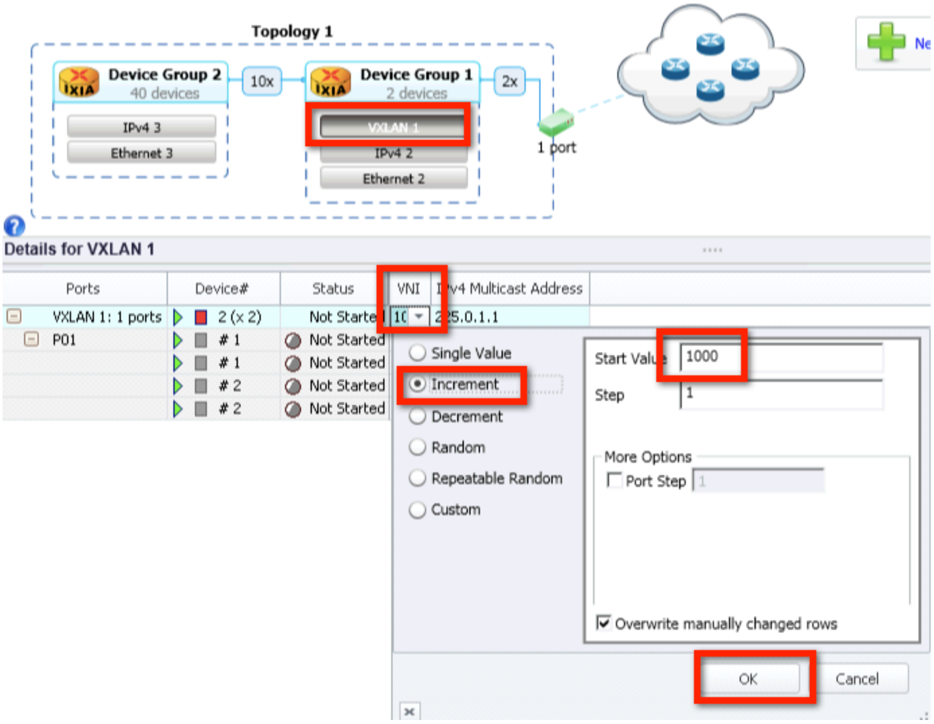

Ports	Device#	Status	VNI	IPv4 Multicast Address
⊟ VXLAN 1: 1 ports	2 (x 2)	Not Started	10	225.0.1.1
⊟ P01	# 1	Not Started		
	# 1	Not Started		
	# 2	Not Started		
	# 2	Not Started		

○ Single Value
● Increment
○ Decrement
○ Random
○ Repeatable Random
○ Custom

Start Value: 1000
Step: 1
More Options
□ Port Step: 1
☑ Overwrite manually changed rows

OK Cancel

Figure 214. Configuring VNIs for each VTEP

8. Click **IPv4 Multicast Address** to configure multicast group for each VNI. Configure Start value of **225.0.1.1** in **Increment** view and click **OK**.

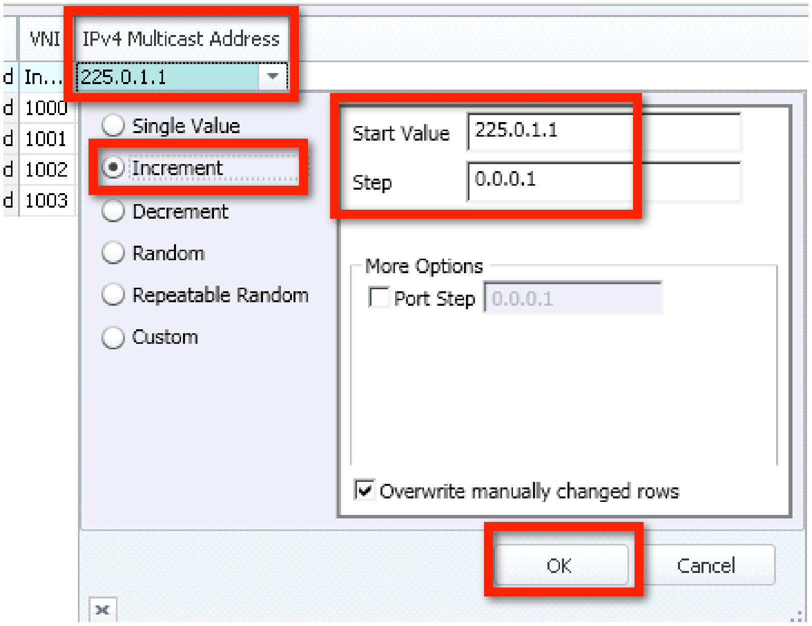

VNI	IPv4 Multicast Address
In…	225.0.1.1
1000	
1001	
1002	
1003	

○ Single Value
● Increment
○ Decrement
○ Random
○ Repeatable Random
○ Custom

Start Value: 225.0.1.1
Step: 0.0.0.1
More Options
□ Port Step: 0.0.0.1
☑ Overwrite manually changed rows

OK Cancel

Figure 215. Configuring Multicast Group for each VNI

9. Click **10x** multiplier to change the number of **Host per VNI**. Configure the Multiplier value of **5** and click **OK**.

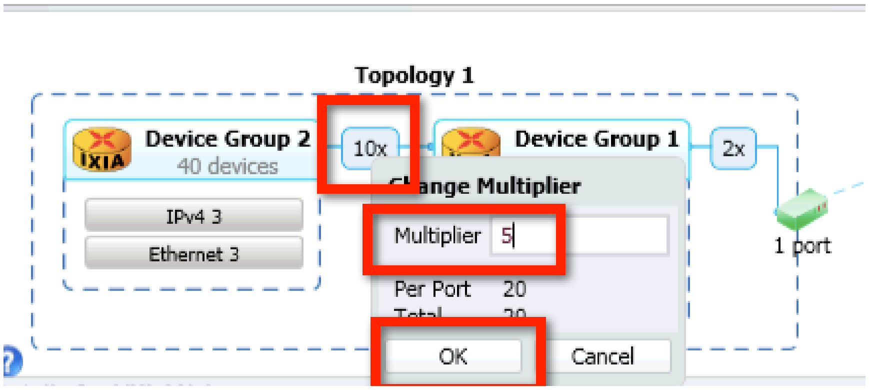

Change Multiplier

Multiplier 5

Per Port 20

Total 20

Figure 216. Configuring number of Host per VNI

10. Click **IPv4** in Device Group2**.** Select **Address** and configure **Start Value** of **200.0.0.1** in **Increment** view and click **OK**.

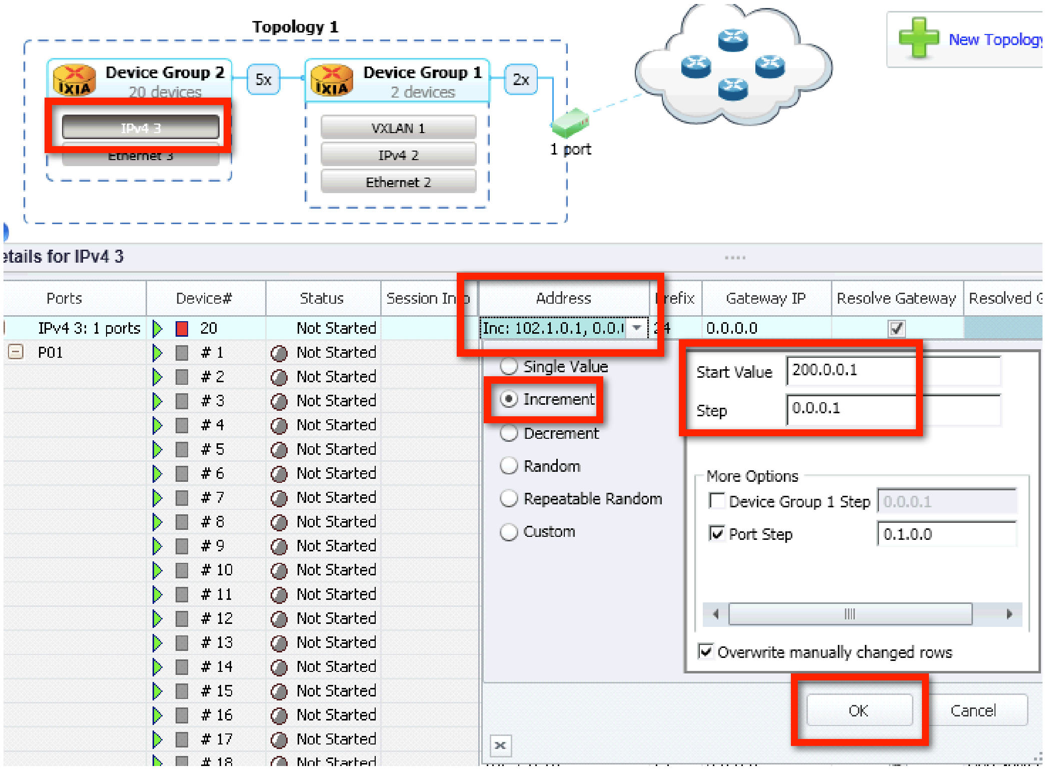

Details for IPv4 3

Ports	Device#	Status	Session Info	Address	Prefix	Gateway IP	Resolve Gateway	Resolved G
IPv4 3: 1 ports	20	Not Started		Inc: 102.1.0.1, 0.0.(	24	0.0.0.0	☑	
P01	# 1	Not Started						
	# 2	Not Started						
	# 3	Not Started						
	# 4	Not Started						
	# 5	Not Started						
	# 6	Not Started						
	# 7	Not Started						
	# 8	Not Started						
	# 9	Not Started						
	# 10	Not Started						
	# 11	Not Started						
	# 12	Not Started						
	# 13	Not Started						
	# 14	Not Started						
	# 15	Not Started						
	# 16	Not Started						
	# 17	Not Started						
	# 18	Not Started						

○ Single Value
● Increment
○ Decrement
○ Random
○ Repeatable Random
○ Custom

Start Value 200.0.0.1
Step 0.0.0.1

More Options
□ Device Group 1 Step 0.0.0.1
☑ Port Step 0.1.0.0

☑ Overwrite manually changed rows

Figure 217. Configuring Host IP address

11. Select the **Gateway IP** and configure the **Start Value of 200.0.0.100** in **Increment** view and click **OK**.

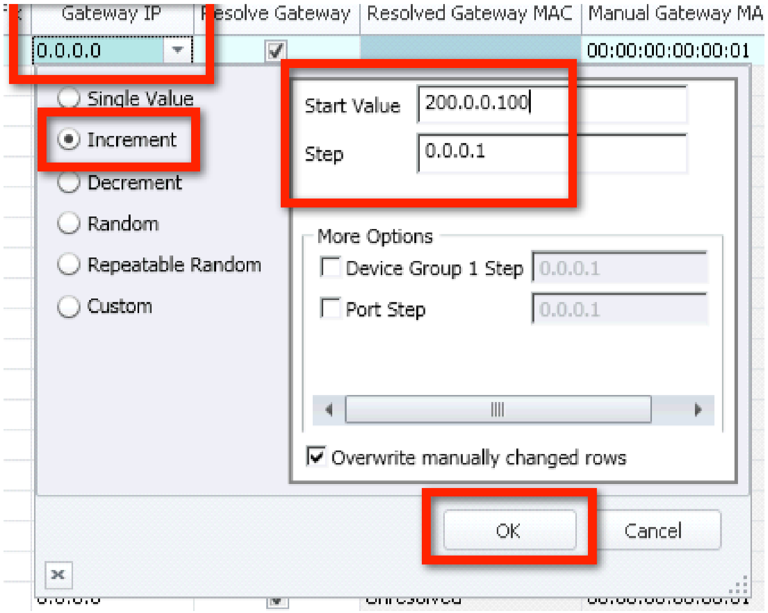

Gateway IP: 0.0.0.0

○ Single Value
● Increment
○ Decrement
○ Random
○ Repeatable Random
○ Custom

Start Value: 200.0.0.100
Step: 0.0.0.1

More Options
□ Device Group 1 Step: 0.0.0.1
□ Port Step: 0.0.0.1

☑ Overwrite manually changed rows

Figure 218. Configuring Host IP Gateway address

12. Click the **Scenario** view, and then click **Add Topolgy** to open the **Protocol Wizard** dialog. Click **Port P02,** and then click **Append Ports**. Click **Next**.

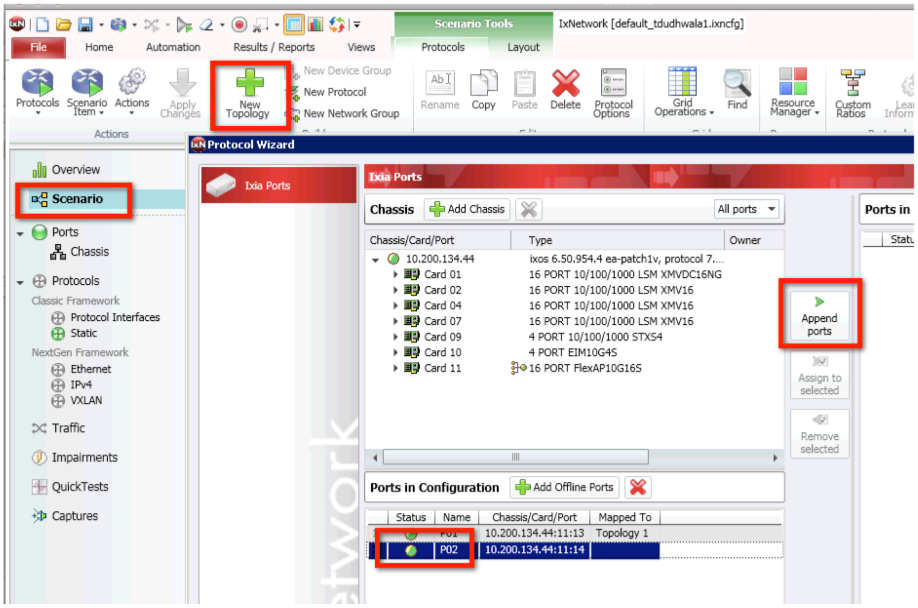

Chassis/Card/Port	Type	Owner
10.200.134.44	ixos 6.50.954.4 ea-patch1v, protocol 7....	
Card 01	16 PORT 10/100/1000 LSM XMVDC16NG	
Card 02	16 PORT 10/100/1000 LSM XMV16	
Card 04	16 PORT 10/100/1000 LSM XMV16	
Card 07	16 PORT 10/100/1000 LSM XMV16	
Card 09	4 PORT 10/100/1000 STXS4	
Card 10	4 PORT EIM10G4S	
Card 11	16 PORT FlexAP10G16S	

Ports in Configuration

Status	Name	Chassis/Card/Port	Mapped To
	P01	10.200.134.44:11:13	Topology 1
	P02	10.200.134.44:11:14	

Figure 219. Selecting Host port in Protocol Wizard

13. Select the **IPv4** checkbox and click **Finish**.

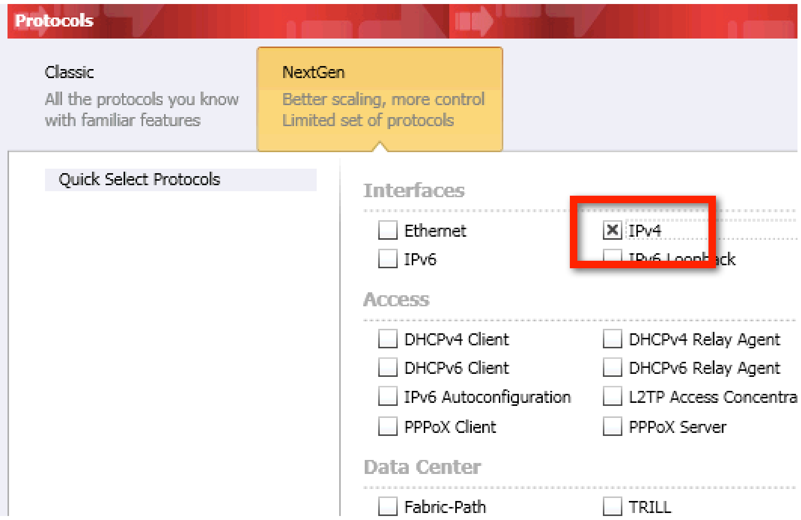

Protocols
Classic
All the protocols you know with familiar features
NextGen
Better scaling, more control Limited set of protocols
Quick Select Protocols

Interfaces

☐ Ethernet ☒ IPv4
☐ IPv6 ☐ IPv6 Loopback

Access

☐ DHCPv4 Client ☐ DHCPv4 Relay Agent
☐ DHCPv6 Client ☐ DHCPv6 Relay Agent
☐ IPv6 Autoconfiguration ☐ L2TP Access Concentra
☐ PPPoX Client ☐ PPPoX Server

Data Center

☐ Fabric-Path ☐ TRILL

Figure 220. Selecting IPv4 host

14. Select **10x** multiplier in **Topology2** and change the value to **20**. This action sets the number of host to 20 to match with host on the VXLAN port. Click **OK.**

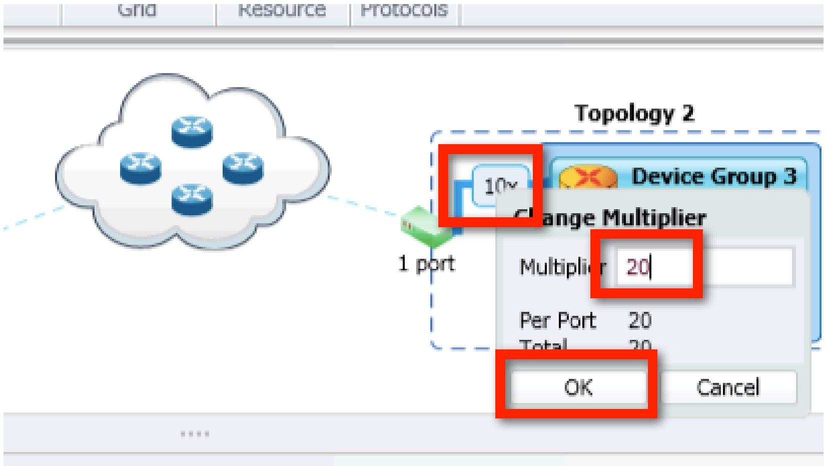

Change Multiplier

Multiplier 20

Per Port 20

Total 20

Figure 221. Configuring number of Host

15. Click **IPv4** in Device Group3**.** Click the **Address** and configure **Start Value** of **200.0.0.100** in **Increment** view. Click **OK**.

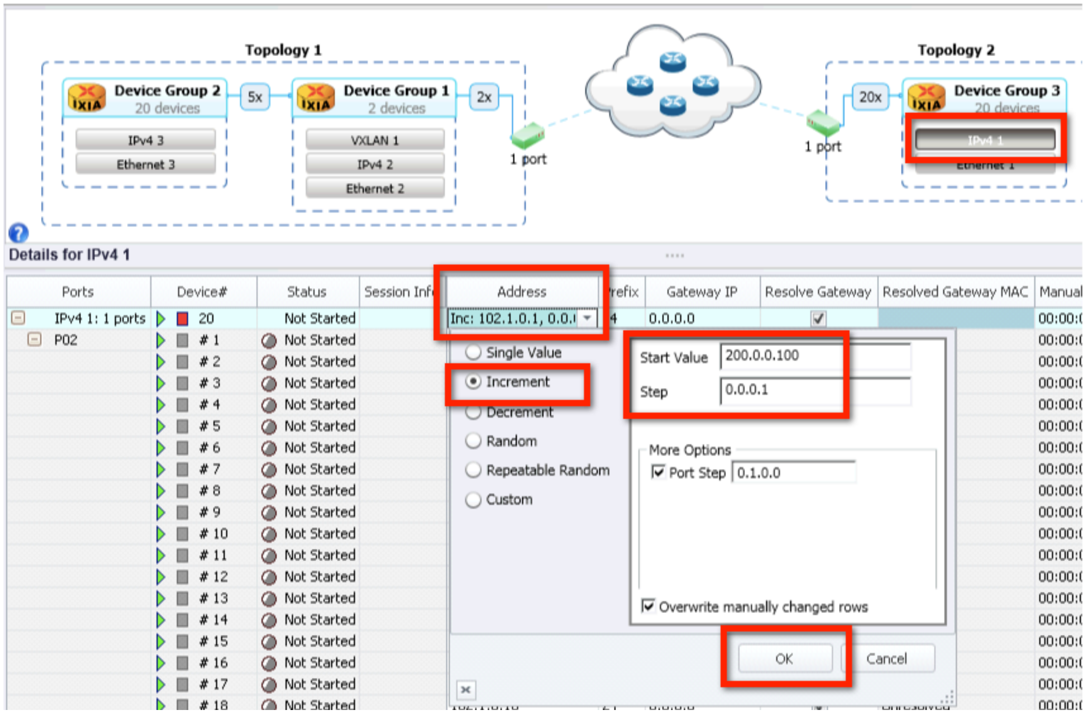

Details for IPv4 1

Ports	Device#	Status	Session Info	Address	Prefix	Gateway IP	Resolve Gateway	Resolved Gateway MAC	Manual
IPv4 1: 1 ports	20	Not Started		Inc: 102.1.0.1, 0.0.(	4	0.0.0.0	☑		00:00:(
P02	# 1	Not Started							00:00:(
	# 2	Not Started							00:00:(
	# 3	Not Started							00:00:(
	# 4	Not Started							00:00:(
	# 5	Not Started							00:00:(
	# 6	Not Started							00:00:(
	# 7	Not Started							00:00:(
	# 8	Not Started							00:00:(
	# 9	Not Started							00:00:(
	# 10	Not Started							00:00:(
	# 11	Not Started							00:00:(
	# 12	Not Started							00:00:(
	# 13	Not Started							00:00:(
	# 14	Not Started							00:00:(
	# 15	Not Started							00:00:(
	# 16	Not Started							00:00:(
	# 17	Not Started							00:00:(
	# 18	Not Started							00:00:(

○ Single Value
● Increment
○ Decrement
○ Random
○ Repeatable Random
○ Custom

Start Value: 200.0.0.100
Step: 0.0.0.1

More Options
☑ Port Step: 0.1.0.0

☑ Overwrite manually changed rows

OK | Cancel

Figure 222. Configuring Host IP address

16. Click the **Gateway IP** and configure the **Start Value of 200.0.0.1** in **Increment** view. Click **OK**.

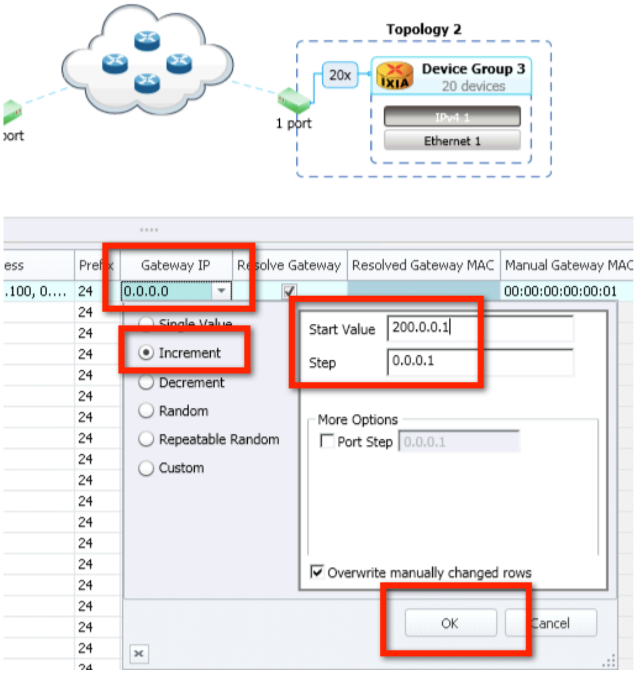

ess	Prefix	Gateway IP	Resolve Gateway	Resolved Gateway MAC	Manual Gateway MAC
.100, 0....	24	0.0.0.0	☑		00:00:00:00:00:01
	24				
	24				
	24				
	24				
	24				
	24				
	24				
	24				
	24				
	24				
	24				
	24				
	24				
	24				
	24				
	24				
	24				
	24				

○ Single Value
● Increment
○ Decrement
○ Random
○ Repeatable Random
○ Custom

Start Value: 200.0.0.1
Step: 0.0.0.1

More Options
☐ Port Step: 0.0.0.1

☑ Overwrite manually changed rows

OK | Cancel

Figure 223. Configuring Host IP Gateway address

17. Click **Scenario** in the left pane to view complete topology.

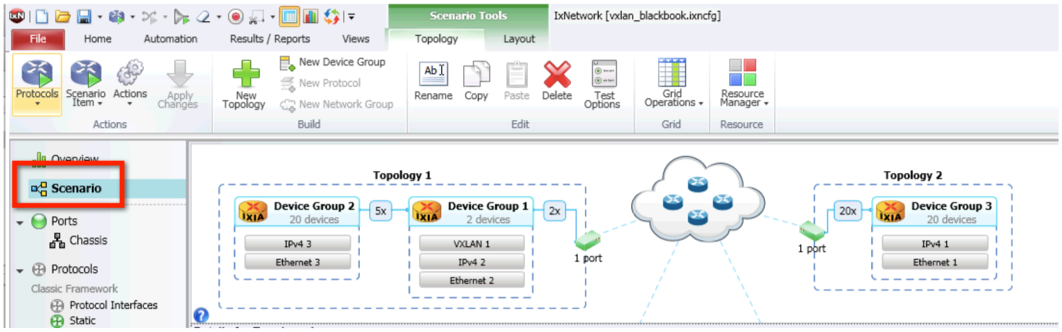

Figure 224. Expanding the ISIS L2/L3 protocol tree

18. Click **Start All Protocols** and wait for the VXLAN and IPv4 Host protocol to come up green.

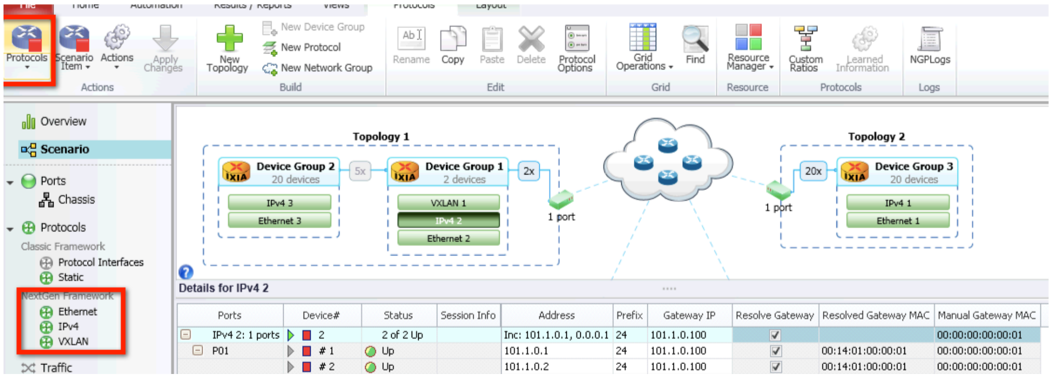

Details for IPv4 2

Ports	Device#	Status	Session Info	Address	Prefix	Gateway IP	Resolve Gateway	Resolved Gateway MAC	Manual Gateway MAC
⊟ IPv4 2: 1 ports	▷ 🟥 2	2 of 2 Up		Inc: 101.1.0.1, 0.0.0.1	24	101.1.0.100	☑		00:00:00:00:00:01
⊟ P01	▷ 🟥 # 1	● Up		101.1.0.1	24	101.1.0.100	☑	00:14:01:00:00:01	00:00:00:00:00:01
	▷ 🟥 # 2	● Up		101.1.0.2	24	101.1.0.100	☑	00:14:01:00:00:01	00:00:00:00:00:01

Figure 225. Starting All Protocols

19. Click the **VXLAN** stack and verify that the **Resolved Gateway MAC** value matches the **DUT** MAC address.

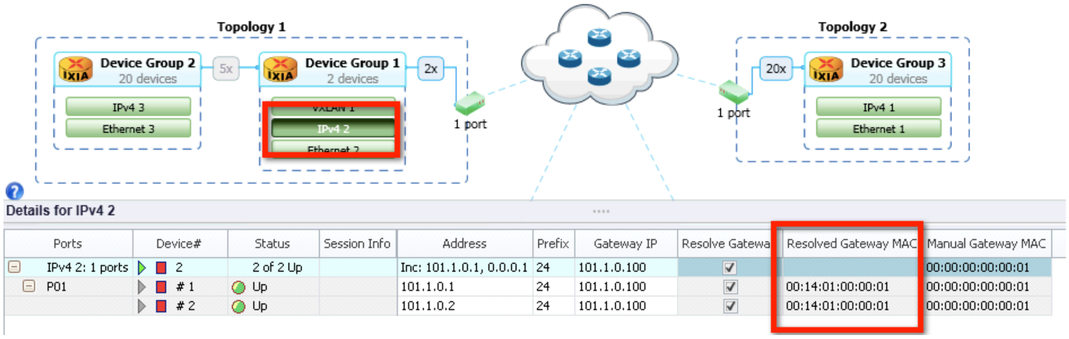

Details for IPv4 2

Ports	Device#	Status	Session Info	Address	Prefix	Gateway IP	Resolve Gatewa	Resolved Gateway MAC	Manual Gateway MAC
⊟ IPv4 2: 1 ports	▷ 🟥 2	2 of 2 Up		Inc: 101.1.0.1, 0.0.0.1	24	101.1.0.100	☑		00:00:00:00:00:01
⊟ P01	▷ 🟥 # 1	● Up		101.1.0.1	24	101.1.0.100	☑	00:14:01:00:00:01	00:00:00:00:00:01
	▷ 🟥 # 2	● Up		101.1.0.2	24	101.1.0.100	☑	00:14:01:00:00:01	00:00:00:00:00:01

Figure 226. Verifying VTEP IP information

20. Click **IPv4** stack and verify that the **Resolved Gateway MAC** value matches the MAC address of Ixia port P02.

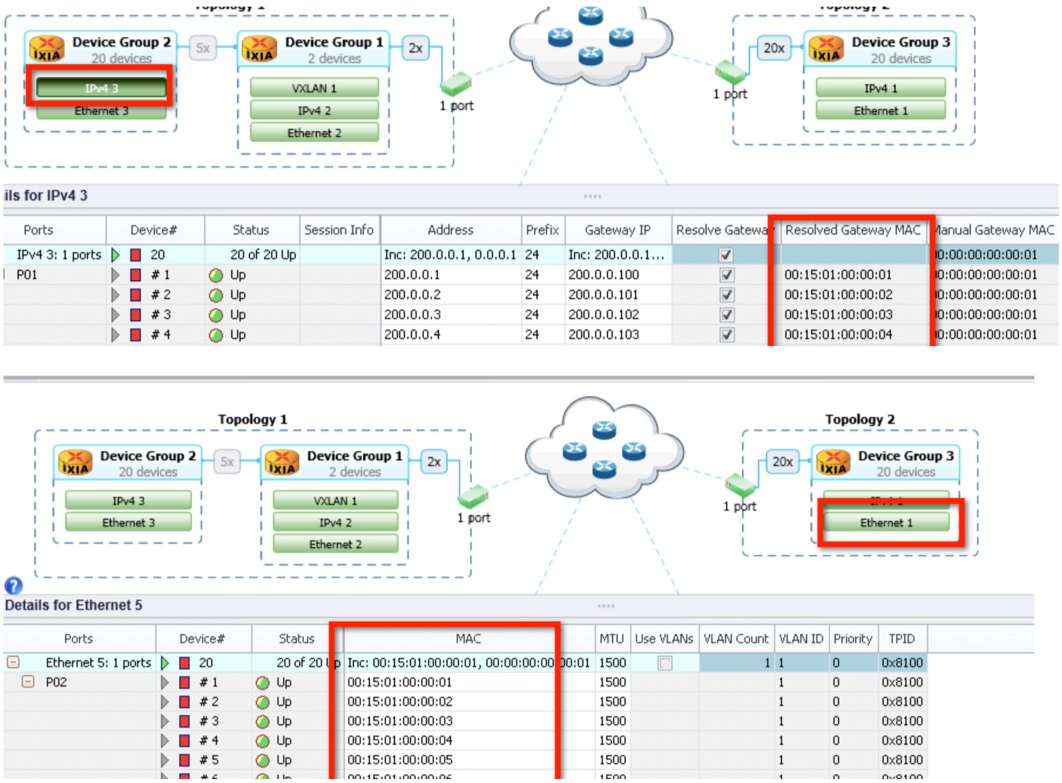

ils for IPv4 3

Ports	Device#	Status	Session Info	Address	Prefix	Gateway IP	Resolve Gatewa	Resolved Gateway MAC	lanual Gateway MAC
IPv4 3: 1 ports	20	20 of 20 Up		Inc: 200.0.0.1, 0.0.0.1	24	Inc: 200.0.0.1...	☑		0:00:00:00:00:01
P01	# 1	Up		200.0.0.1	24	200.0.0.100	☑	00:15:01:00:00:01	0:00:00:00:00:01
	# 2	Up		200.0.0.2	24	200.0.0.101	☑	00:15:01:00:00:02	0:00:00:00:00:01
	# 3	Up		200.0.0.3	24	200.0.0.102	☑	00:15:01:00:00:03	0:00:00:00:00:01
	# 4	Up		200.0.0.4	24	200.0.0.103	☑	00:15:01:00:00:04	0:00:00:00:00:01

Details for Ethernet 5

Ports	Device#	Status	MAC	MTU	Use VLANs	VLAN Count	VLAN ID	Priority	TPID
Ethernet 5: 1 ports	20	20 of 20 Up	Inc: 00:15:01:00:00:01, 00:00:00:00:00:01	1500	☐	1	1	0	0x8100
P02	# 1	Up	00:15:01:00:00:01	1500			1	0	0x8100
	# 2	Up	00:15:01:00:00:02	1500			1	0	0x8100
	# 3	Up	00:15:01:00:00:03	1500			1	0	0x8100
	# 4	Up	00:15:01:00:00:04	1500			1	0	0x8100
	# 5	Up	00:15:01:00:00:05	1500			1	0	0x8100

Figure 227. Verifying Host IP information

21. Click **Traffic** in the left pane and click **Add L2-3 Traffic Items**.

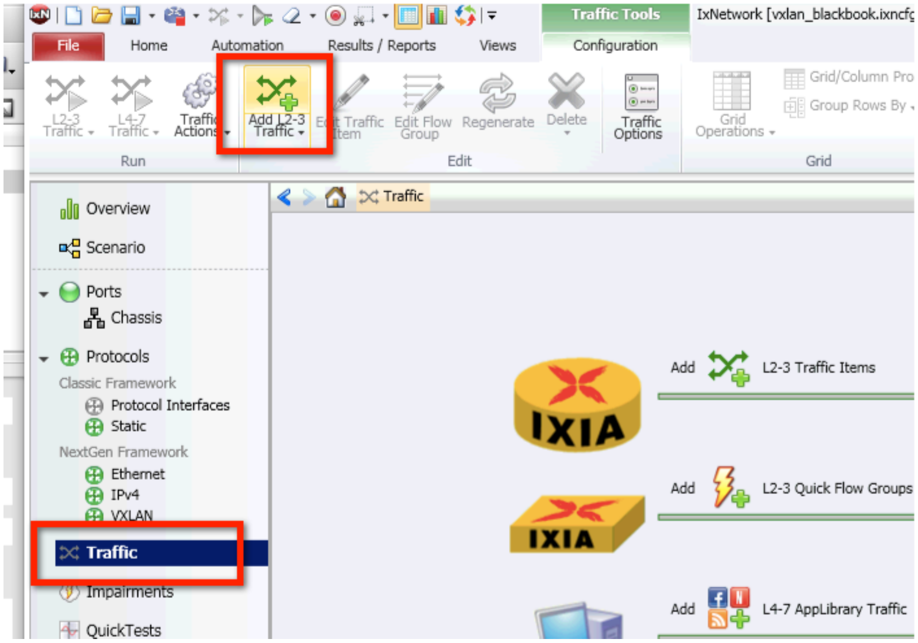

Figure 228. Configuring Traffic

22. In the **Type of Traffic** dropdown list, select **IPv4.** Select **Bi-Directional** checkbox. In the **Source** dropdown list, select **IPv4** inside the **Device Group2** of **Topology1**. In the **Destination** dropdown list, select **IPv4** in **Device Group3** of **Topology2**.

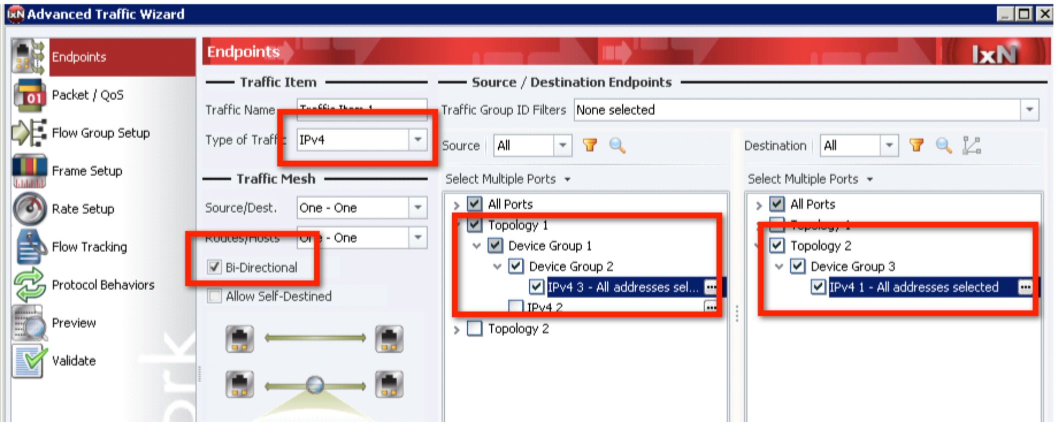

Endpoints

Traffic Item

Traffic Name: Traffic Item 1

Type of Traffic: IPv4

Traffic Mesh

Source/Dest.: One - One

Routes/Hosts: One - One

☑ Bi-Directional

☐ Allow Self-Destined

Source / Destination Endpoints

Traffic Group ID Filters: None selected

Source: All

Select Multiple Ports

- ☑ All Ports
- ☑ Topology 1
  - ☑ Device Group 1
    - ☑ Device Group 2
      - ☑ IPv4 3 - All addresses sel...
      - ☐ IPv4 2
- ☐ Topology 2

Destination: All

Select Multiple Ports

- ☑ All Ports
- ☐ Topology 1
- ☑ Topology 2
  - ☑ Device Group 3
    - ☑ IPv4 1 - All addresses selected

Figure 229. Selecting IPv4 from VTEP and remote Host port

23. In the left pane, click **Frame Setup** and change the **Fixed** frame size value to **1500**.

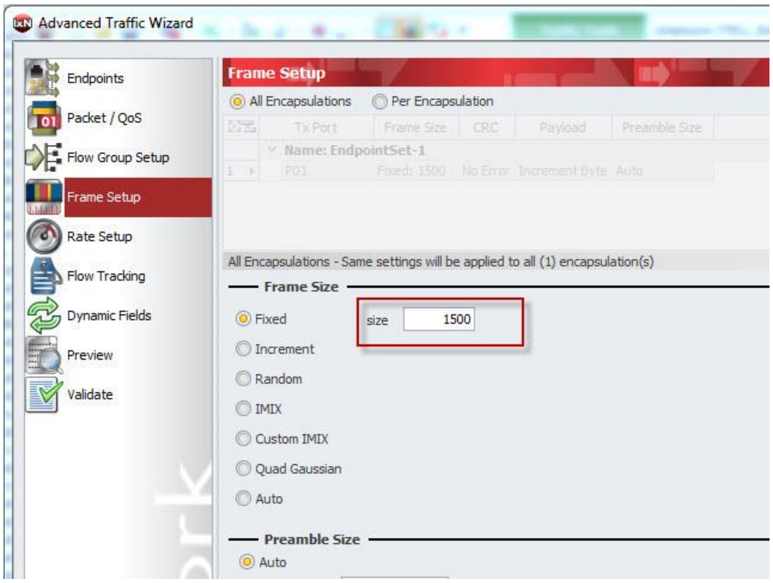

Frame Setup

● All Encapsulations ○ Per Encapsulation

Name: EndpointSet-1

1 P01 Fixed: 1500 No Error Increment Byte Auto

All Encapsulations - Same settings will be applied to all (1) encapsulation(s)

Frame Size

● Fixed size 1500

○ Increment

○ Random

○ IMIX

○ Custom IMIX

○ Quad Gaussian

○ Auto

Preamble Size

● Auto

Figure 230. Changing the fixed frame size value

24. In the left pane, click **Rate Setup** and change the **Line rate** percentage value to 100.

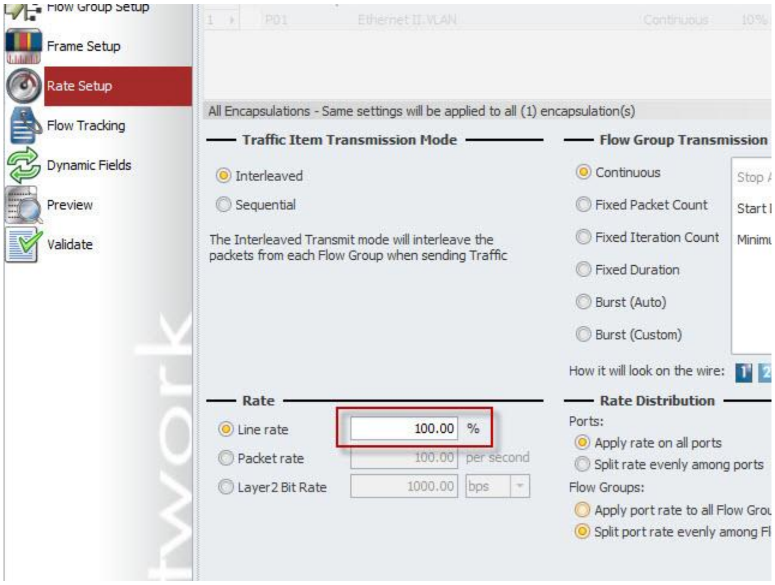

All Encapsulations - Same settings will be applied to all (1) encapsulation(s)

Traffic Item Transmission Mode

● Interleaved
○ Sequential

The Interleaved Transmit mode will interleave the packets from each Flow Group when sending Traffic

Flow Group Transmission

● Continuous
○ Fixed Packet Count
○ Fixed Iteration Count
○ Fixed Duration
○ Burst (Auto)
○ Burst (Custom)

Rate

● Line rate 100.00 %
○ Packet rate 100.00 per second
○ Layer2 Bit Rate 1000.00 bps

Rate Distribution

Ports:
● Apply rate on all ports
○ Split rate evenly among ports

Flow Groups:
● Apply port rate to all Flow Grou
● Split port rate evenly among Fl

Figure 231. Changing the line rate throughput value

25. In the left pane, click **Flow Tracking** and select **Traffic Item** tracking options. Select **IPv4: Destination Address(1) and VXLAN: VNI** tracking options. These options help you to drill down for more details dynamically on VXLAN and IPv4 Host endpoint. Click **Finish** to close the **Advanced Traffic Wizard**.

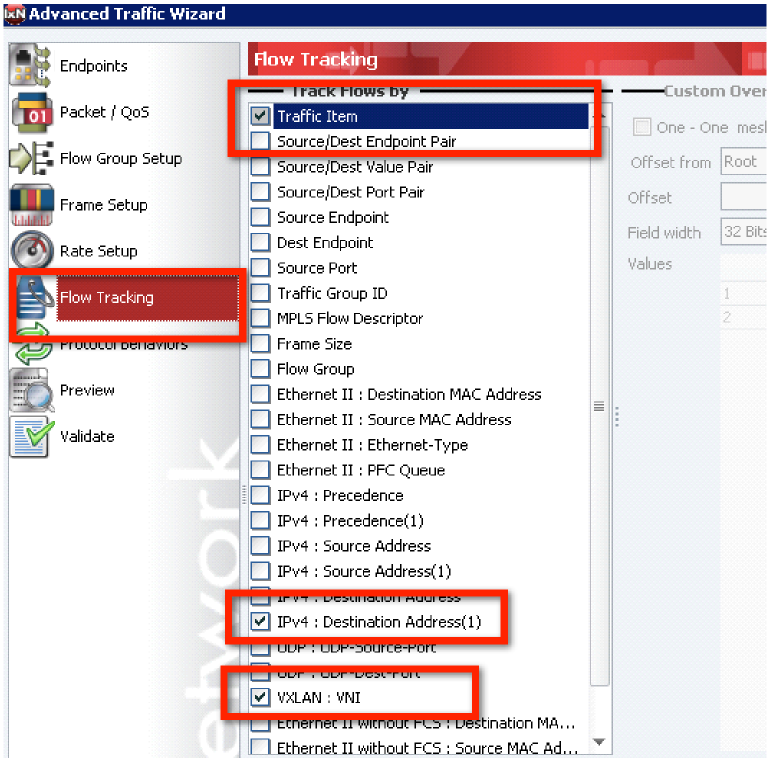

Track Flows by

☑ Traffic Item
□ Source/Dest Endpoint Pair
□ Source/Dest Value Pair
□ Source/Dest Port Pair
□ Source Endpoint
□ Dest Endpoint
□ Source Port
□ Traffic Group ID
□ MPLS Flow Descriptor
□ Frame Size
□ Flow Group
□ Ethernet II : Destination MAC Address
□ Ethernet II : Source MAC Address
□ Ethernet II : Ethernet-Type
□ Ethernet II : PFC Queue
□ IPv4 : Precedence
□ IPv4 : Precedence(1)
□ IPv4 : Source Address
□ IPv4 : Source Address(1)
□ IPv4 : Destination Address
☑ IPv4 : Destination Address(1)
□ UDP : UDP-Source-Port
□ UDP : UDP-Dest-Port
☑ VXLAN : VNI
□ Ethernet II without FCS : Destination MA...
□ Ethernet II without FCS : Source MAC Ad...

Figure 232. Selecting flow tracking options

Test Variables

Performance Variable	Description
No of IPs in VXLAN stack	Use this test variable to increase the count of VTEPs
Multiplier between VXLAN and IP stack	Use this test variable to increase the count of VNIs per VTEPs to more then 1
Multiplier between VXLAN Device Group and IP DeviceGroup	Use this test variable to increase the count of IP Host per VNIs
DHCP Client	Add DHCP Client device instead of IPv4 behind VTEP Gateway to emulate DHCP client Host.
IPv6 Client	Add IPv6 device instead of IPv4 behind VTEP Gateway to emulate IPv6 Host.

Troubleshooting and Diagnostics

Issue	Troubleshooting Solution
IPv4 Host ARP does not resolve	Check DUT settings to ensure: • VNIs configuration matches with IxNetwork • Multicast groups configuration matches with IxNetwork • UDP port configuration matches with IxNetwork, look for global protocol settings • Host IP address doesn’t conflict with VTEP IP address

Results Analysis

1. Click **Start All Traffic** and, then click **L2-L3 Test Summary Statistics**.

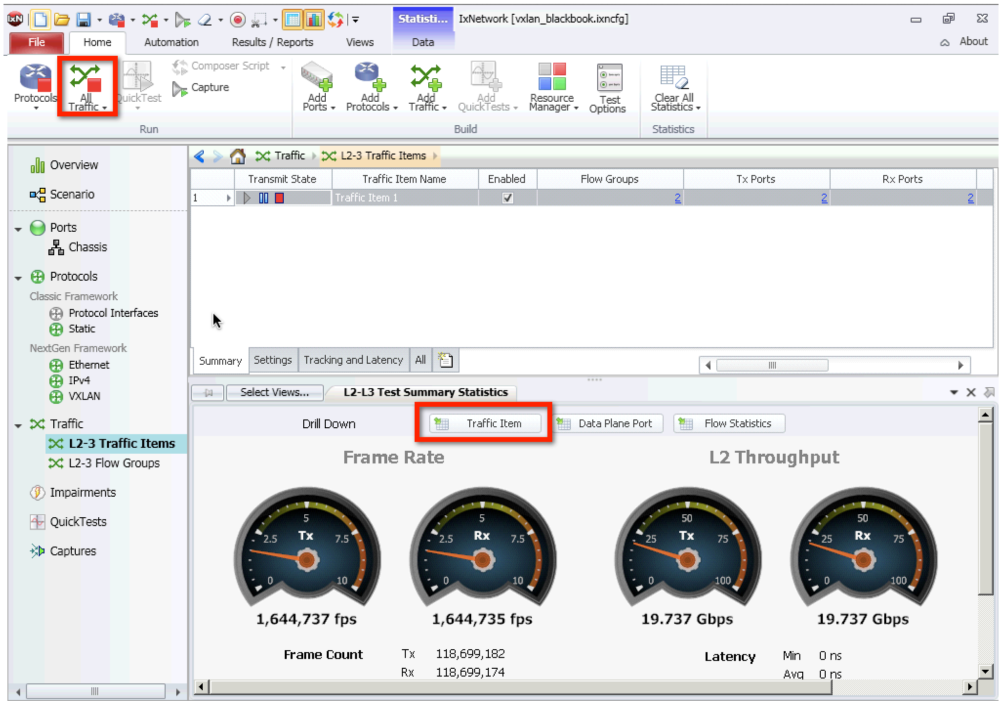

	Transmit State	Traffic Item Name	Enabled	Flow Groups	Tx Ports	Rx Ports
1		Traffic Item 1	☑	2	2	2

1,644,737 fps
1,644,735 fps
19.737 Gbps
19.737 Gbps

Frame Count Tx 118,699,182
Rx 118,699,174

Latency Min 0 ns
Avg 0 ns

Figure 233. Viewing L2-L3 Test Summary Statistics

2. Select **Traffic Item 1** from the **Traffic Item Statistics** view. Right-click to open the context menu, and then click **Drill down per VXLAN: VNI** to see throughput and packet loss statistics on a per VNI basis in the **User Defined Statistics** view.

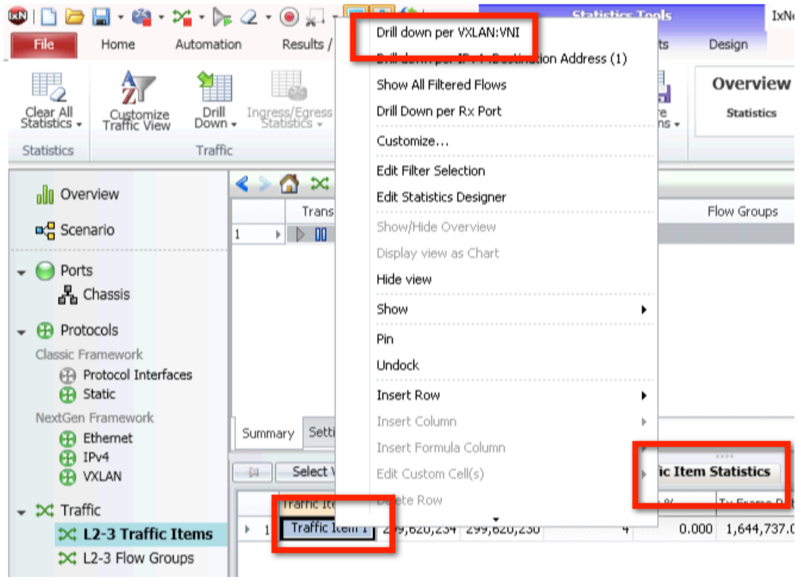

Figure 234. Reviewing Traffic Items Statistics

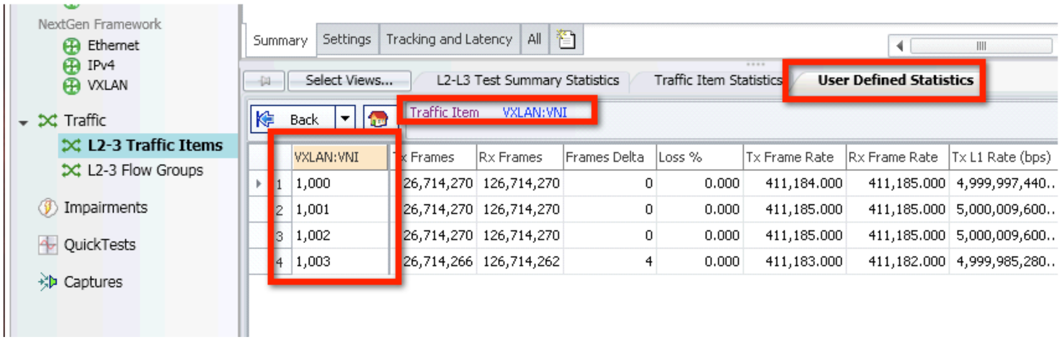

	VXLAN:VNI	x Frames	Rx Frames	Frames Delta	Loss %	Tx Frame Rate	Rx Frame Rate	Tx L1 Rate (bps)
1	1,000	26,714,270	126,714,270	0	0.000	411,184.000	411,185.000	4,999,997,440..
2	1,001	26,714,270	126,714,270	0	0.000	411,185.000	411,185.000	5,000,009,600..
3	1,002	26,714,270	126,714,270	0	0.000	411,185.000	411,185.000	5,000,009,600..
4	1,003	26,714,266	126,714,262	4	0.000	411,183.000	411,182.000	4,999,985,280..

Figure 235. Reviewing per-VNI statistics

3. Select **VNI 1000** from the **User Defined Statistics** view. Right-click to open the context menu, and the click **Drill down per IPv4: Destination Address (1)** to see throughput and packet loss statistics on a per IPv4 Host per VNI bases in the **User Defined Statistics** view.

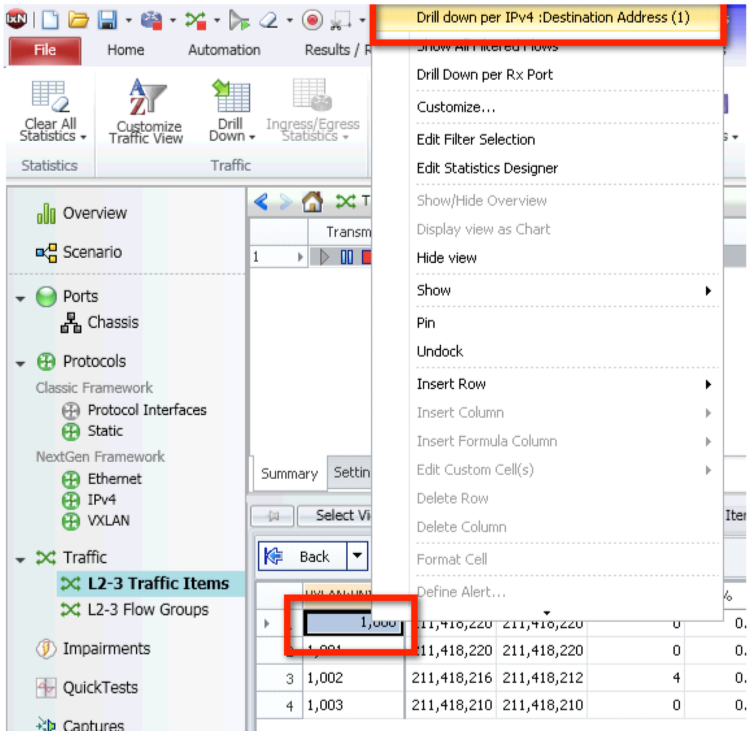

	1,000	211,418,220	211,418,220	0	0.
	1,001	211,418,220	211,418,220	0	0.
3	1,002	211,418,216	211,418,212	4	0.
4	1,003	211,418,210	211,418,210	0	0.

Figure 236. Selecting per-IPv4 Destination statistics

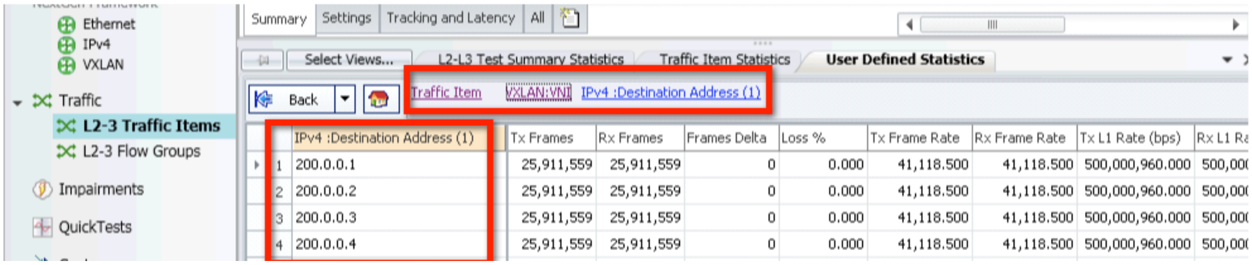

	IPv4 :Destination Address (1)	Tx Frames	Rx Frames	Frames Delta	Loss %	Tx Frame Rate	Rx Frame Rate	Tx L1 Rate (bps)	Rx L1 Ra
1	200.0.0.1	25,911,559	25,911,559	0	0.000	41,118.500	41,118.500	500,000,960.000	500,00
2	200.0.0.2	25,911,559	25,911,559	0	0.000	41,118.500	41,118.500	500,000,960.000	500,00
3	200.0.0.3	25,911,559	25,911,559	0	0.000	41,118.500	41,118.500	500,000,960.000	500,00
4	200.0.0.4	25,911,559	25,911,559	0	0.000	41,118.500	41,118.500	500,000,960.000	500,00

Figure 237. Reviewing per-IPv4 Destination statistics

Conclusions

The IxNetwork VXLAN Emulation enables the configuration of complex VTEP topologies along with all of the VNI and various Host protocols per VNI, such as IPv4, IPv6, and DHCP. The flexible flow tracking in the **Advanced Traffic Wizard** along with the real time drill down capability in statistics views allows the user to quickly isolate flows that are not being forwarded correctly by the DUT and detect the root cause by associating them with individual VNI within.

Test Case: VXLAN Unicast VTEP Reachability and Forwarding

Objective

This test case validates the core functionality of a VXLAN gateway and its ability to encapsulate and decapsulate host traffic. Also, it measures, if the DUT is correctly forwarding the traffic for each VXLAN segment, that is, VNI. This test case differs from the previous one in that instead of joining multicast groups to determine reachability of remote VTEPs, the test tool is programmed with the static IP address of the peer to which it wants to encapsulate traffic.

Setup

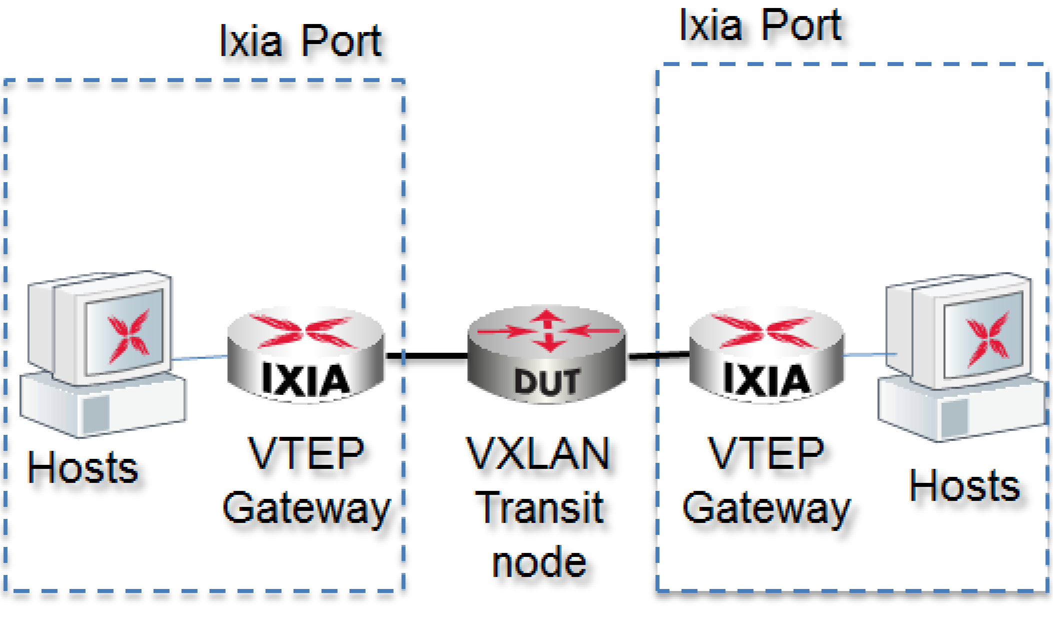

Figure 238. Data Center Transit Switch running VXLAN

Ixia ports emulate VTEP Gateways and simulated hosts behind the gateway. The DUT forms VXLAN Segments over the VXLAN domain to securely bridge the services running between VXLAN-attached hosts.

Step-by-Step Instructions

1. Click **Add Ports** to open the **Port Selection** Window. Click **Add Chassis** and enter the IP Address or name for your IXIA chassis. Click **OK** to accept. Expand chassis and cards and select two test ports you want to use in this test. Click **Add ports** and then **OK** to add the ports to your test configuration.

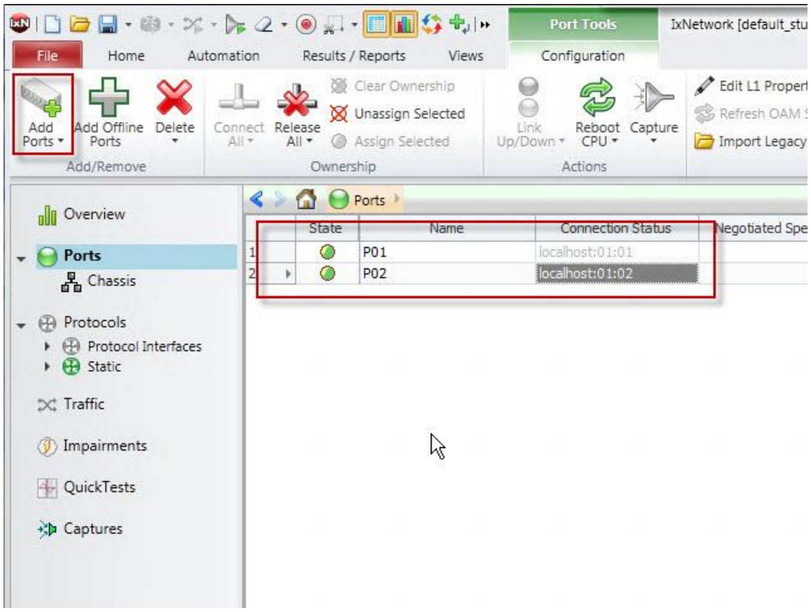

	State	Name	Connection Status	Negotiated Spe
1	●	P01	localhost:01:01	
2	●	P02	localhost:01:02	

Figure 239. Adding ports to a new configuration

2. Click the **Scenario** view, and then click **Add Topolgy** to open the **Protocol Wizard** dialog. Select **Port P01** and **Append Ports**. Click **Next**.

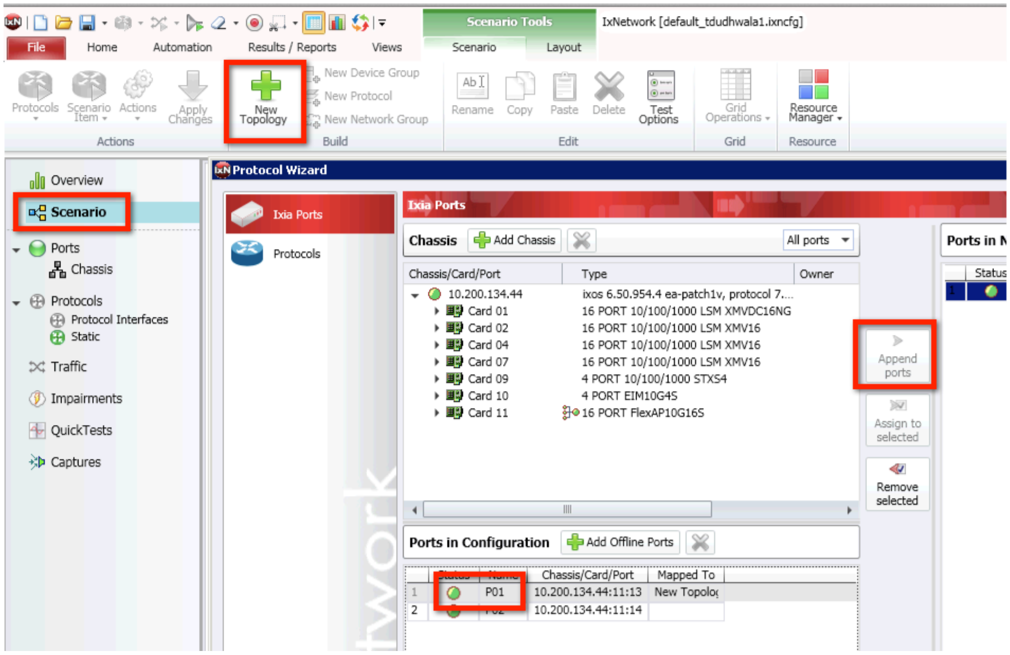

Ixia Ports

Chassis/Card/Port	Type	Owner
10.200.134.44	ixos 6.50.954.4 ea-patch1v, protocol 7....	
Card 01	16 PORT 10/100/1000 LSM XMVDC16NG	
Card 02	16 PORT 10/100/1000 LSM XMV16	
Card 04	16 PORT 10/100/1000 LSM XMV16	
Card 07	16 PORT 10/100/1000 LSM XMV16	
Card 09	4 PORT 10/100/1000 STXS4	
Card 10	4 PORT EIM10G4S	
Card 11	16 PORT FlexAP10G16S	

Ports in Configuration

	Status	Name	Chassis/Card/Port	Mapped To
1	●	P01	10.200.134.44:11:13	New Topolo
2	●	P02	10.200.134.44:11:14	

Figure 240. Selecting VXLAN Port in Protocol Wizard

3. Select the **VXLAN** checkbox and click **Finish**. This action adds 10x VXLAN VTEP Gateways and 10x IPv4 Hosts behind each gateway, total 100 hosts.

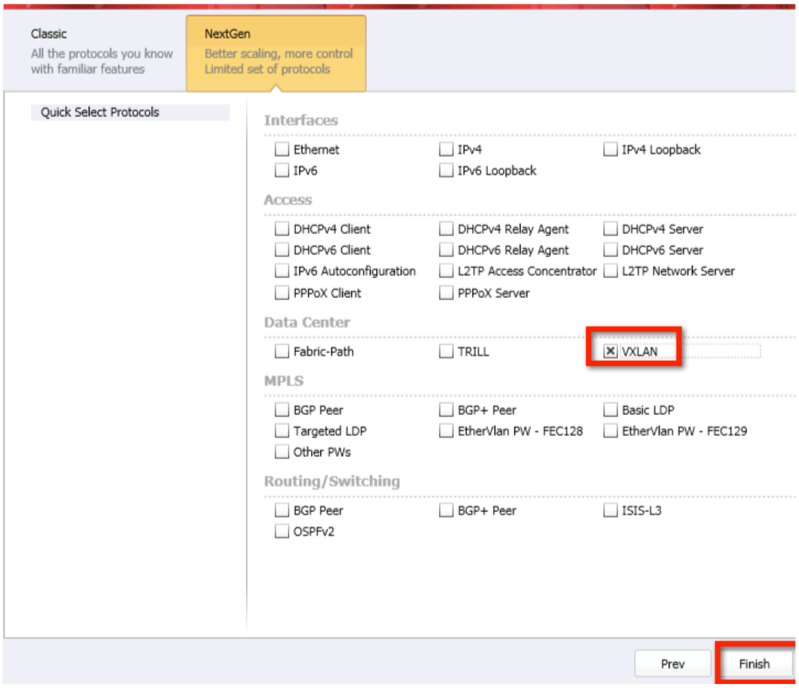
Classic
All the protocols you know with familiar features
NextGen
Better scaling, more control
Limited set of protocols
Quick Select Protocols
Prev
Finish

Interfaces

☐ Ethernet ☐ IPv4 ☐ IPv4 Loopback
☐ IPv6 ☐ IPv6 Loopback

Access

☐ DHCPv4 Client ☐ DHCPv4 Relay Agent ☐ DHCPv4 Server
☐ DHCPv6 Client ☐ DHCPv6 Relay Agent ☐ DHCPv6 Server
☐ IPv6 Autoconfiguration ☐ L2TP Access Concentrator ☐ L2TP Network Server
☐ PPPoX Client ☐ PPPoX Server

Data Center

☐ Fabric-Path ☐ TRILL ☒ VXLAN

MPLS

☐ BGP Peer ☐ BGP+ Peer ☐ Basic LDP
☐ Targeted LDP ☐ EtherVlan PW - FEC128 ☐ EtherVlan PW - FEC129
☐ Other PWs

Routing/Switching

☐ BGP Peer ☐ BGP+ Peer ☐ ISIS-L3
☐ OSPFv2

Figure 241. Adding VXLAN on Port P01

4. Click **10x** box in the **Scenario** window and change the value to **1**. This action changes the number of **VTEP devices** from 10 to 1. Click **OK**.

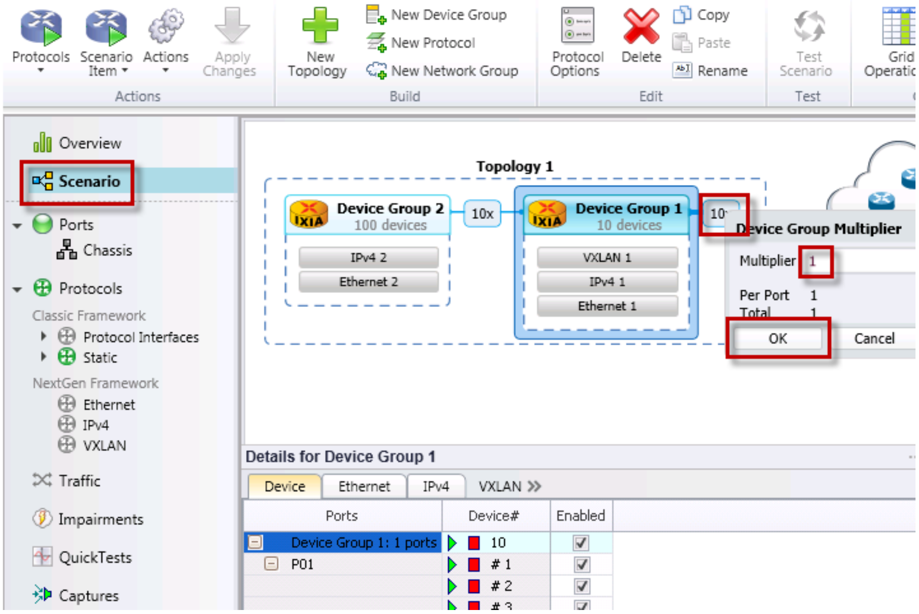

Device Group Multiplier

Multiplier 1

Per Port 1

Total 1

Details for Device Group 1

Device | Ethernet | IPv4 | VXLAN

Ports	Device#	Enabled
⊟ Device Group 1: 1 ports	▶ 🟥 10	☑
⊟ P01	▶ 🟥 # 1	☑
	▶ 🟥 # 2	☑
	▶ 🟥 # 3	☑

Figure 242. Changing number of VTEP per port to 1

5. Click **10x** multiplier to change the number of **Host per VNI**. Configure the Multiplier value of **2** and click **OK**.

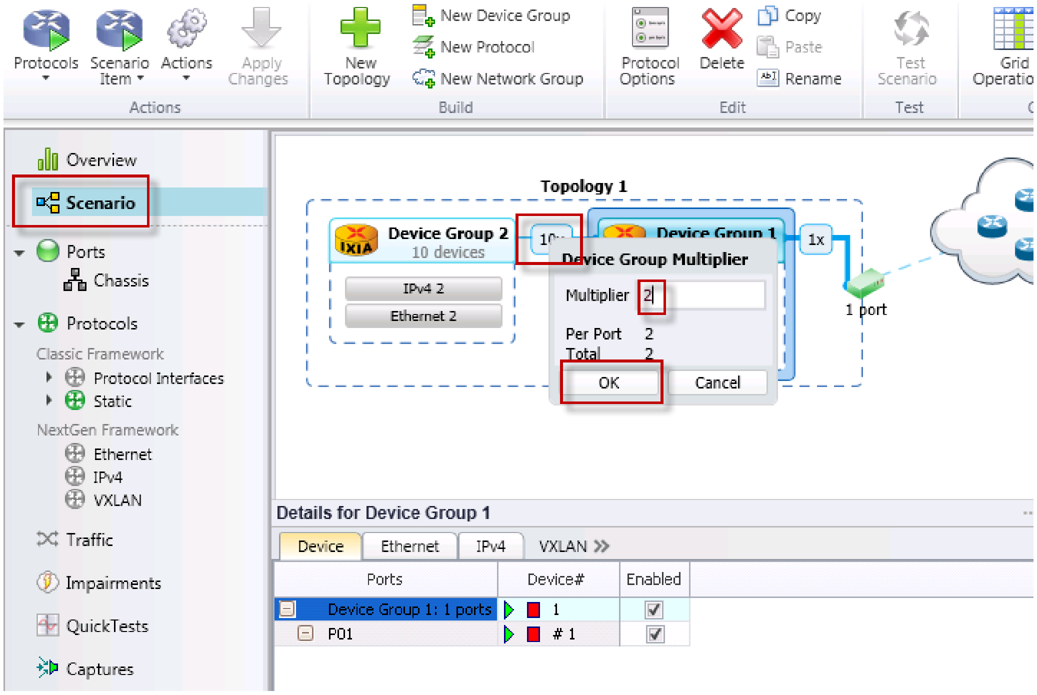

Device Group Multiplier

Multiplier: 2
Per Port: 2
Total: 2

OK Cancel

Ports	Device#	Enabled
Device Group 1: 1 ports	1	☑
P01	# 1	☑

Figure 243. Configuring number of Host per VNI

6. Click **IPv4** in **VXLAN Device Group** to configure the **VTEP IP address**. Click the **Address** and enter **Start Value** of **20.0.0.1.** Click **OK.** Similarly, set the **Gateway IP** to **20.0.0.101.**

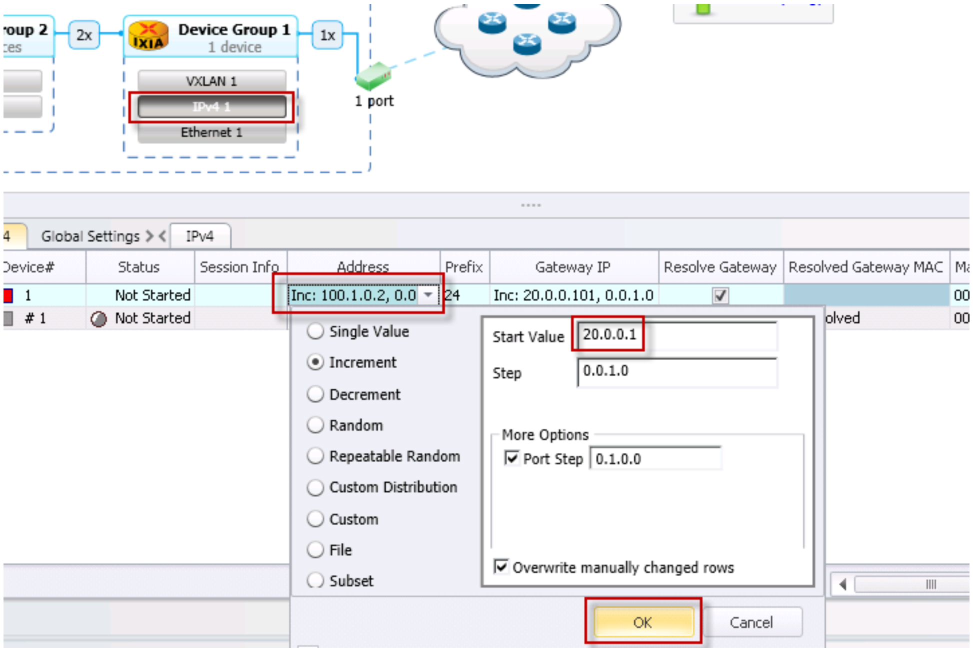

Device#	Status	Session Info	Address	Prefix	Gateway IP	Resolve Gateway	Resolved Gateway MAC	Ma
1	Not Started		Inc: 100.1.0.2, 0.0	24	Inc: 20.0.0.101, 0.0.1.0	☑		00
# 1	Not Started						olved	00

○ Single Value
● Increment
○ Decrement
○ Random
○ Repeatable Random
○ Custom Distribution
○ Custom
○ File
○ Subset

Start Value: 20.0.0.1
Step: 0.0.1.0

More Options
☑ Port Step 0.1.0.0

☑ Overwrite manually changed rows

OK Cancel

Figure 244. Setting VTEP IP address

7. Click **VXLAN** in **Scenario** View to change the VTEP VNI configuration. Click **VNI** and configure the **Start Value** to 1001 in Increment. Click **OK**.

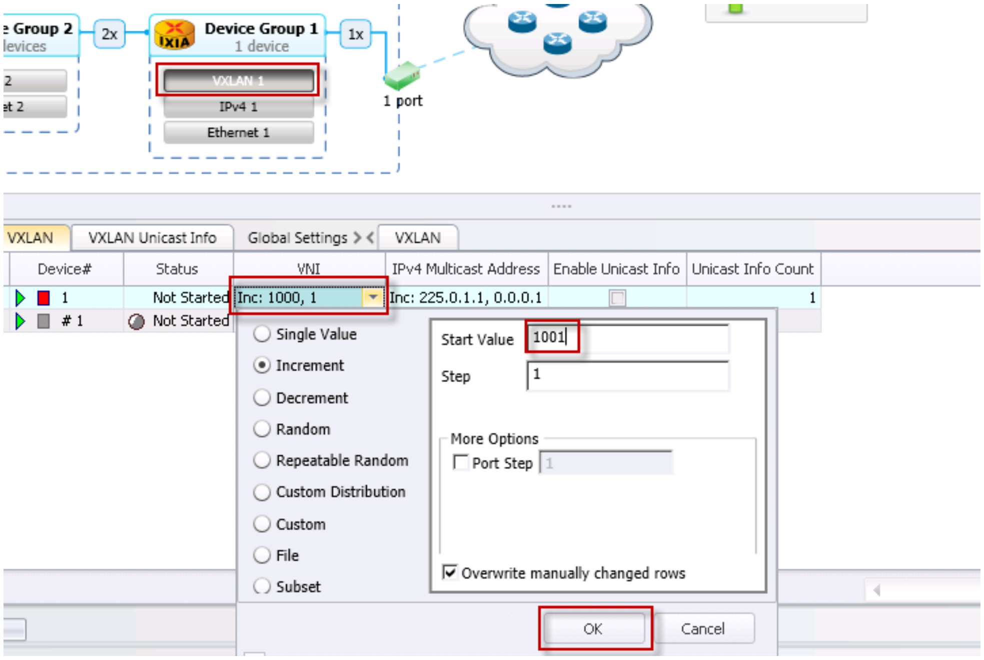

Device#	Status	VNI	IPv4 Multicast Address	Enable Unicast Info	Unicast Info Count
1	Not Started	Inc: 1000, 1	Inc: 225.0.1.1, 0.0.0.1	□	1
# 1	Not Started				

○ Single Value
● Increment
○ Decrement
○ Random
○ Repeatable Random
○ Custom Distribution
○ Custom
○ File
○ Subset

Start Value 1001
Step 1
More Options
□ Port Step 1
☑ Overwrite manually changed rows
OK Cancel

Figure 245. Configuring VNIs for each VTEP

8. Click **Enable Unicast Info** checkbox and change **Unicast Info Count** value to **2**.

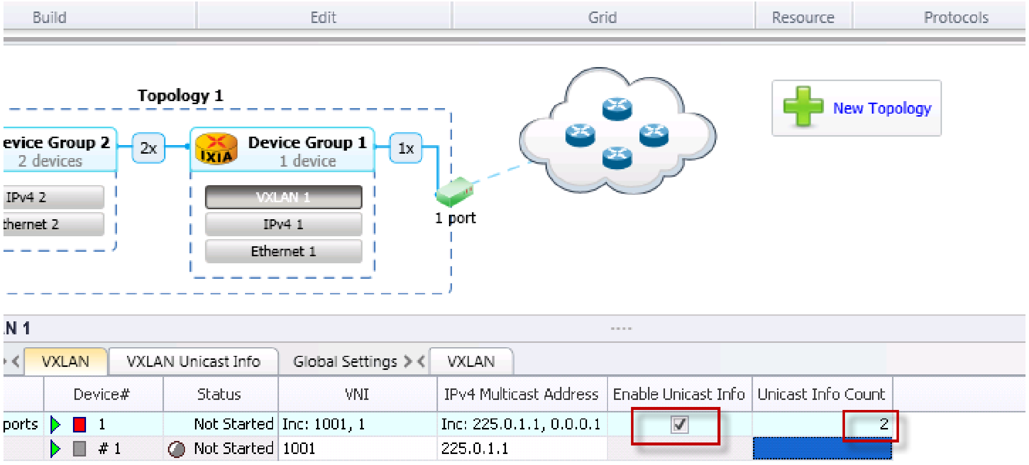

Device#	Status	VNI	IPv4 Multicast Address	Enable Unicast Info	Unicast Info Count
1	Not Started	Inc: 1001, 1	Inc: 225.0.1.1, 0.0.0.1	☑	2
# 1	Not Started	1001	225.0.1.1		

Figure 246. Enabling VXLAN Unicast Info instead of multicast

9. Select the **VXLAN Unicast Info** tab and select the pattern dropdown for **Remote VTEP Unicast IPv4**. Specify a **Single Value** radio button option **Start Value** of **20.0.0.101** and click **OK** to accept.

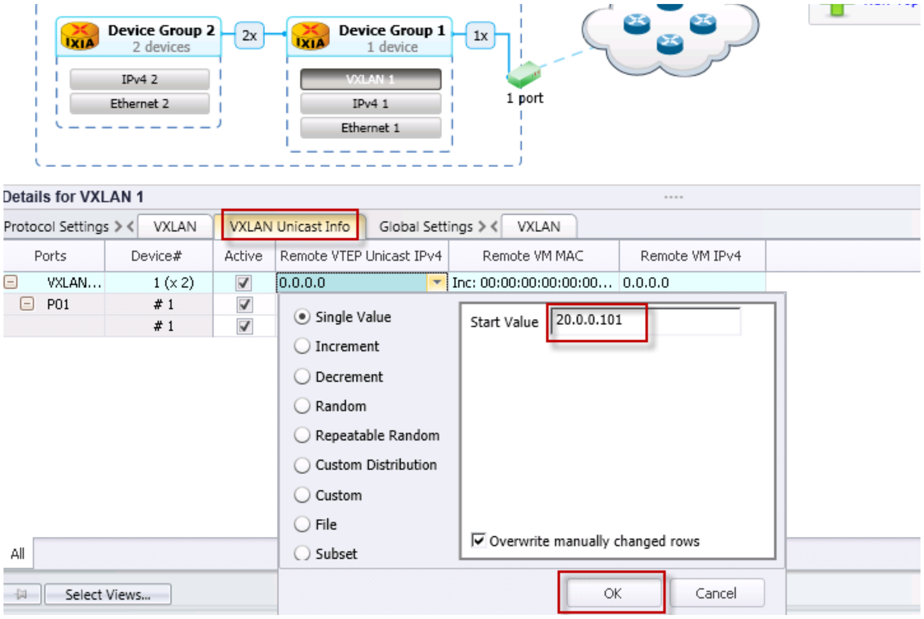

Ports	Device#	Active	Remote VTEP Unicast IPv4	Remote VM MAC	Remote VM IPv4
VXLAN...	1 (x 2)	☑	0.0.0.0	Inc: 00:00:00:00:00:00...	0.0.0.0
P01	# 1	☑			
	# 1	☑			

● Single Value
○ Increment
○ Decrement
○ Random
○ Repeatable Random
○ Custom Distribution
○ Custom
○ File
○ Subset

Start Value 20.0.0.101

☑ Overwrite manually changed rows

Figure 247. Configuring Remote VTEP Unicast IPv4 address

10. Select the **VXLAN Unicast Info** tab and select the pattern dropdown for **Remote VM MAC**. Specify an Increment radio button option **Start Value** of **00:14:01:00:00:01** and a **Step** value of **00:00:00:00:00:01**. Click **OK** to accept.

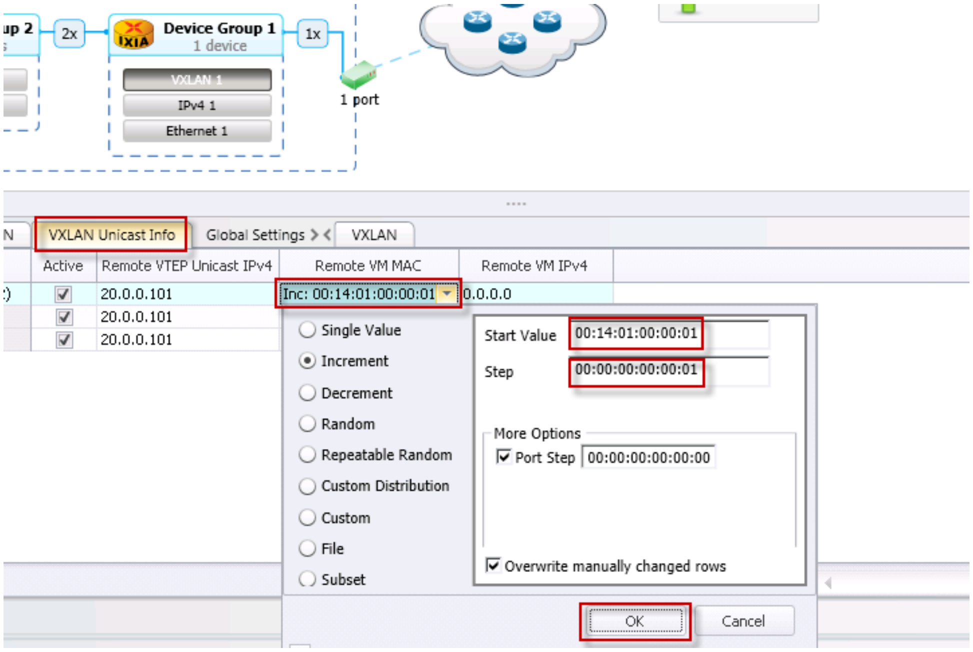

Active	Remote VTEP Unicast IPv4	Remote VM MAC	Remote VM IPv4
☑	20.0.0.101	Inc: 00:14:01:00:00:01	0.0.0.0
☑	20.0.0.101		
☑	20.0.0.101		

○ Single Value
● Increment
○ Decrement
○ Random
○ Repeatable Random
○ Custom Distribution
○ Custom
○ File
○ Subset

Start Value 00:14:01:00:00:01

Step 00:00:00:00:00:01

More Options

☑ Port Step 00:00:00:00:00:00

☑ Overwrite manually changed rows

Figure 248. Configuring Remote VM MAC address

11. Select the VXLAN Unicast Info tab and select the pattern dropdown for **Remote VM IPv4**. Specify an increment radio button option **Start Value** of **30.0.0.101** and a **Step** value of **0.0.0.1**. Click **OK** to accept.

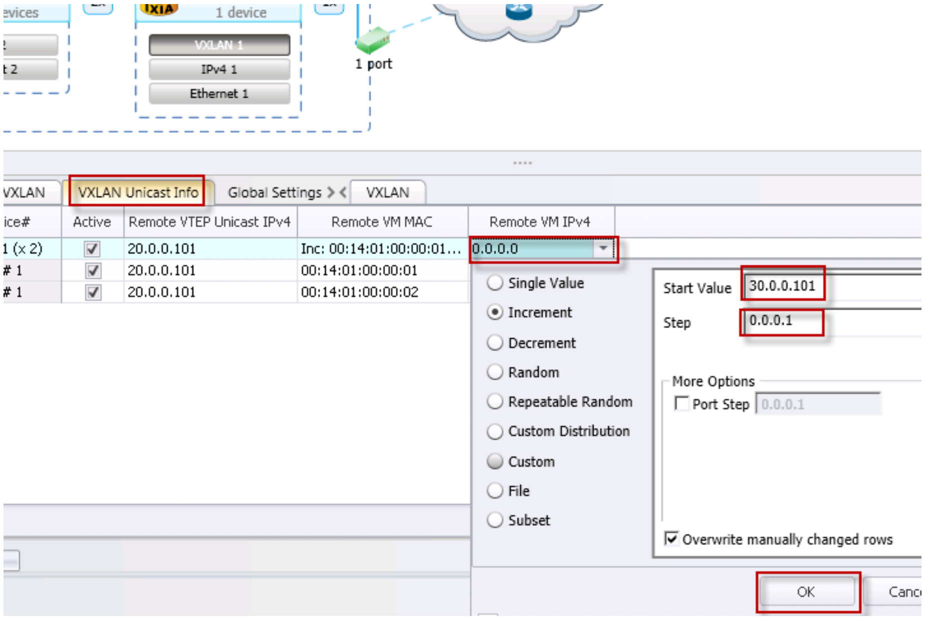

ice#	Active	Remote VTEP Unicast IPv4	Remote VM MAC	Remote VM IPv4
1 (x 2)	☑	20.0.0.101	Inc: 00:14:01:00:00:01...	0.0.0.0
# 1	☑	20.0.0.101	00:14:01:00:00:01	
# 1	☑	20.0.0.101	00:14:01:00:00:02	

○ Single Value
● Increment
○ Decrement
○ Random
○ Repeatable Random
○ Custom Distribution
○ Custom
○ File
○ Subset

Start Value 30.0.0.101
Step 0.0.0.1

More Options
☐ Port Step 0.0.0.1

☑ Overwrite manually changed rows

OK Canc

Figure 249. Configuring Remote VM IPv4 address

12. Click **IPv4** in Device Group2**.** Select **Address** and configure **Start Value** of **30.0.0.1** and a **Step** value in **Increment** view and click **OK**.

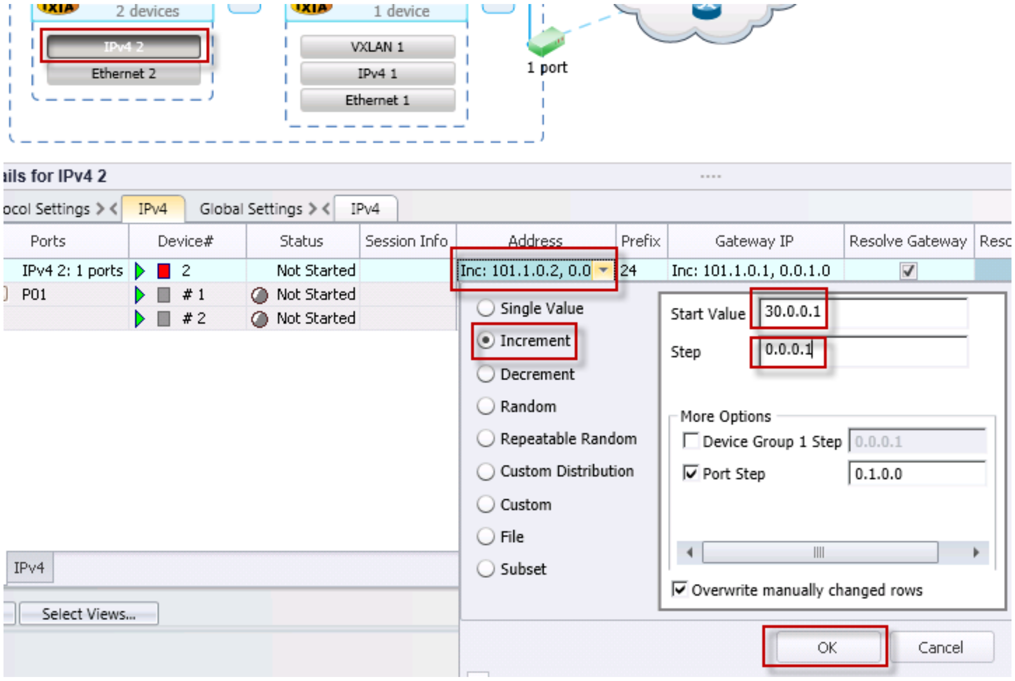

ils for IPv4 2

ocol Settings | IPv4 | Global Settings | IPv4

Ports	Device#	Status	Session Info	Address	Prefix	Gateway IP	Resolve Gateway	Reso
IPv4 2: 1 ports	▶ 🟥 2	Not Started		Inc: 101.1.0.2, 0.0	24	Inc: 101.1.0.1, 0.0.1.0	☑	
P01	▶ # 1	Not Started						
	▶ # 2	Not Started						

○ Single Value
● Increment
○ Decrement
○ Random
○ Repeatable Random
○ Custom Distribution
○ Custom
○ File
○ Subset

Start Value 30.0.0.1
Step 0.0.0.1

More Options
☐ Device Group 1 Step 0.0.0.1
☑ Port Step 0.1.0.0

☑ Overwrite manually changed rows

IPv4

Select Views...

OK Cancel

Figure 250. Configuring Host IP address

13. Select the **Gateway IP** and configure the **Start Value of 30.0.0.101** in **Increment** view and click **OK**.

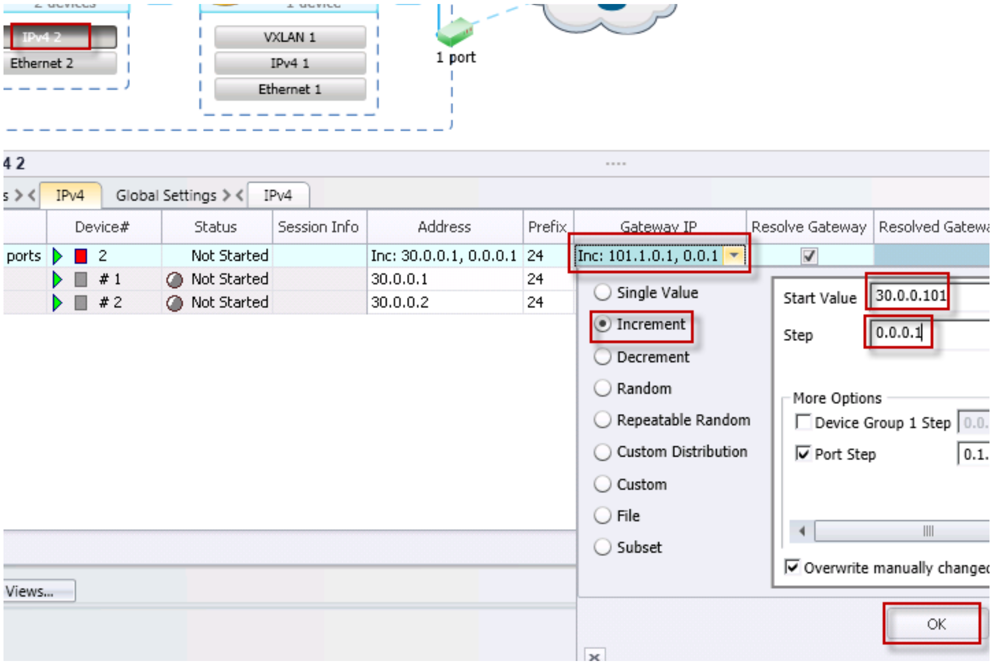

		Device#	Status	Session Info	Address	Prefix	Gateway IP	Resolve Gateway	Resolved Gatewa
ports	▶ 🟥	2	Not Started		Inc: 30.0.0.1, 0.0.0.1	24	Inc: 101.1.0.1, 0.0.1	☑	
	▶	# 1	Not Started		30.0.0.1	24			
	▶	# 2	Not Started		30.0.0.2	24			

○ Single Value
● Increment
○ Decrement
○ Random
○ Repeatable Random
○ Custom Distribution
○ Custom
○ File
○ Subset

Start Value: 30.0.0.101
Step: 0.0.0.1

More Options
□ Device Group 1 Step 0.0.
☑ Port Step 0.1.

☑ Overwrite manually change

Figure 251. Configuring Host IP Gateway address

14. Click the **Scenario** view, and then click **Add Topolgy** to open the **Protocol Wizard** dialog. Click **Port P02,** and then click **Append Ports**. Click **Next**.

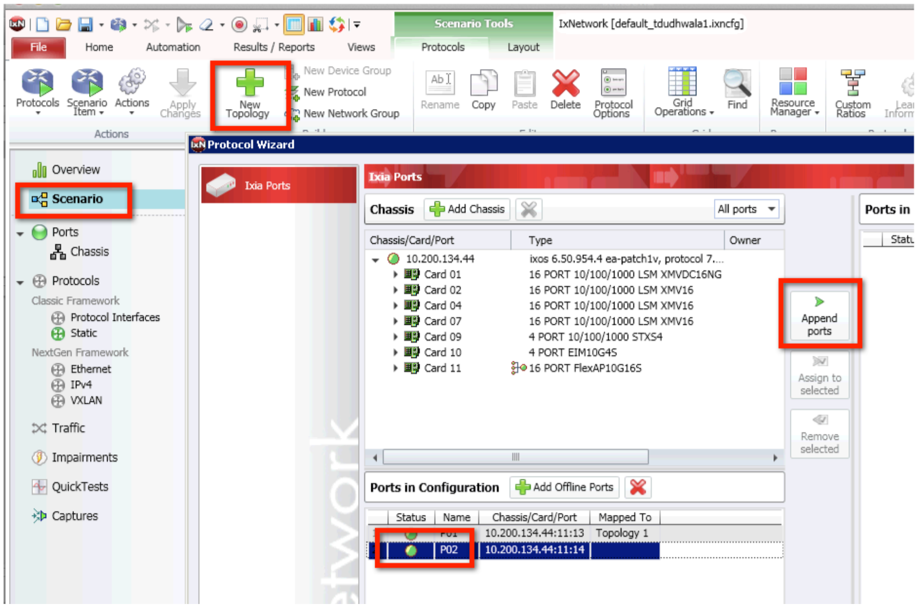

Chassis/Card/Port	Type	Owner
10.200.134.44	ixos 6.50.954.4 ea-patch1v, protocol 7....	
Card 01	16 PORT 10/100/1000 LSM XMVDC16NG	
Card 02	16 PORT 10/100/1000 LSM XMV16	
Card 04	16 PORT 10/100/1000 LSM XMV16	
Card 07	16 PORT 10/100/1000 LSM XMV16	
Card 09	4 PORT 10/100/1000 STXS4	
Card 10	4 PORT EIM10G4S	
Card 11	16 PORT FlexAP10G16S	

Status	Name	Chassis/Card/Port	Mapped To
●	P01	10.200.134.44:11:13	Topology 1
●	P02	10.200.134.44:11:14	

Figure 252. Selecting VXLAN port in Protocol Wizard

15. Select the **VXLAN** checkbox and click **Finish**. This action adds 10x VXLAN VTEP Gateways and 10x IPv4 Hosts behind each gateway, total 100 hosts.

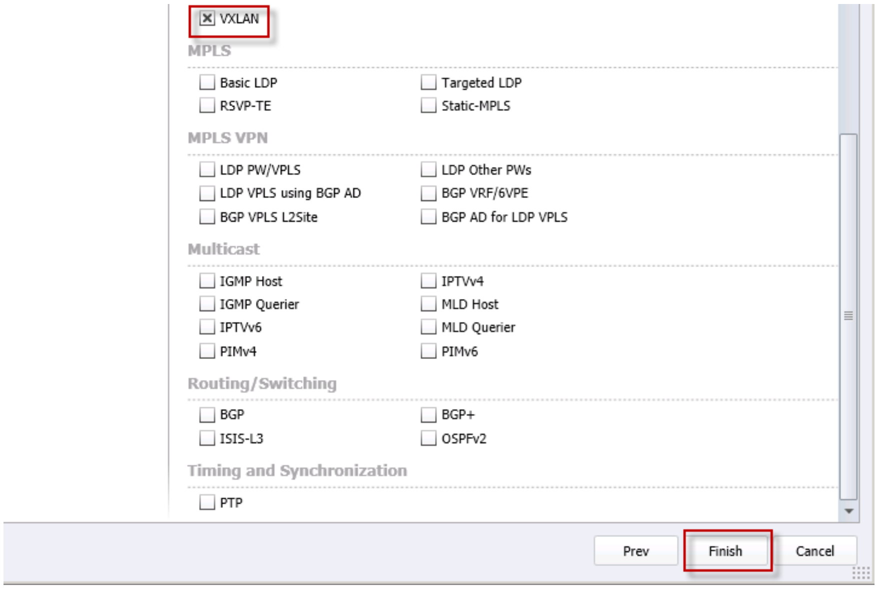

☒ VXLAN

MPLS

☐ Basic LDP ☐ Targeted LDP
☐ RSVP-TE ☐ Static-MPLS

MPLS VPN

☐ LDP PW/VPLS ☐ LDP Other PWs
☐ LDP VPLS using BGP AD ☐ BGP VRF/6VPE
☐ BGP VPLS L2Site ☐ BGP AD for LDP VPLS

Multicast

☐ IGMP Host ☐ IPTVv4
☐ IGMP Querier ☐ MLD Host
☐ IPTVv6 ☐ MLD Querier
☐ PIMv4 ☐ PIMv6

Routing/Switching

☐ BGP ☐ BGP+
☐ ISIS-L3 ☐ OSPFv2

Timing and Synchronization

☐ PTP

Figure 253. Adding VXLAN on Port P02

16. Click **10x** box in the **Scenario** window and change the value to **1**. This action changes the number of **VTEP devices** from 10 to 1. Click **OK**.

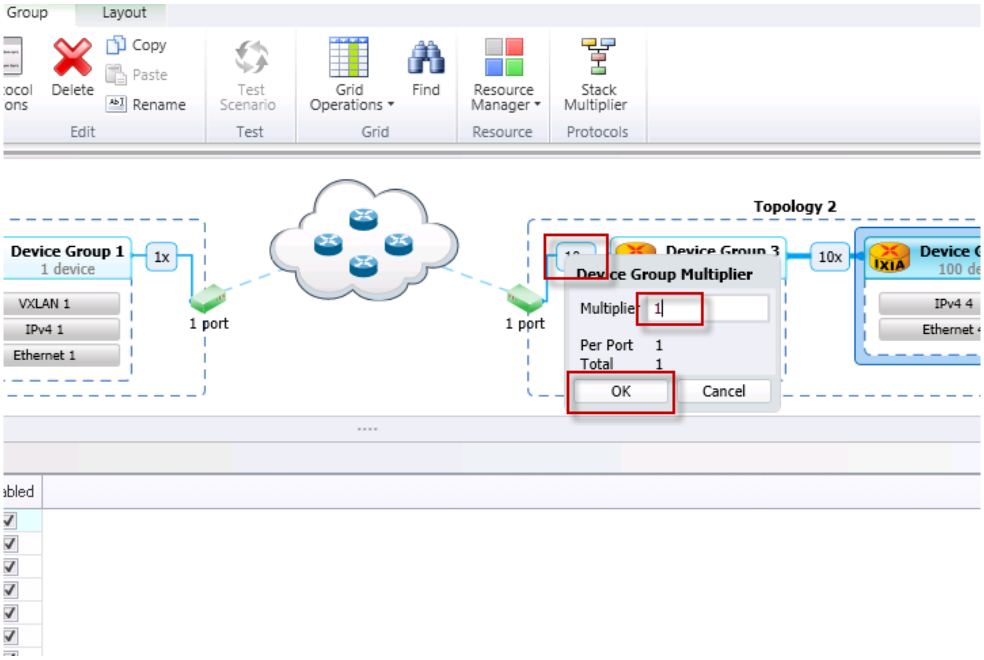

Device Group Multiplier

Multiplier 1

Per Port 1
Total 1

Figure 254. Changing number of VTEP per port to 1

17. Click **10x** multiplier to change the number of **Host per VNI**. Configure the Multiplier value of **2** and click **OK**.

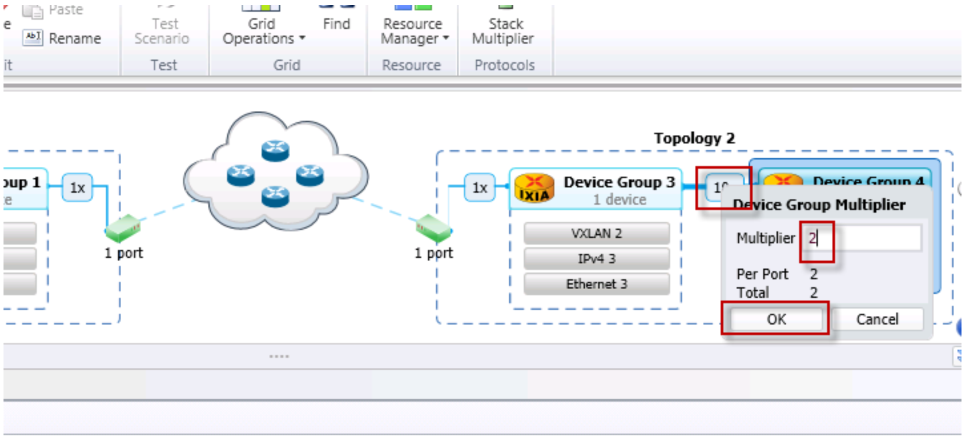

Device Group Multiplier

Multiplier 2

Per Port 2

Total 2

OK

Figure 255. Configuring number of Host per VNI

18. Click **IPv4** in **VXLAN Device Group** to configure the **VTEP IP address**. Click the **Address** and enter **Start Value** of **20.0.0.101.** Click **OK.**

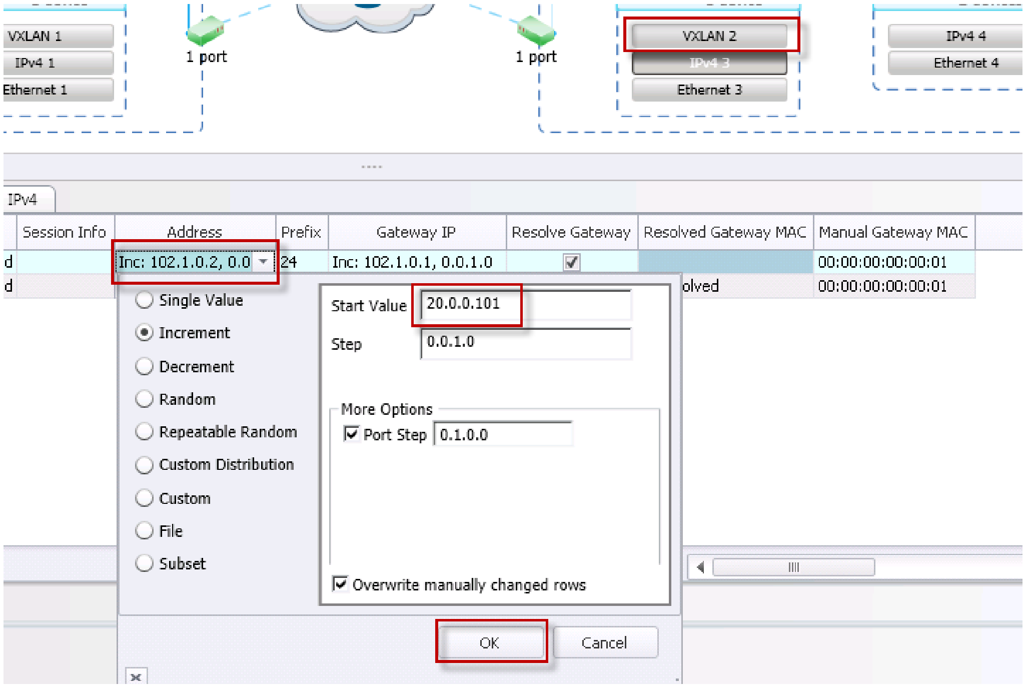

	Session Info	Address	Prefix	Gateway IP	Resolve Gateway	Resolved Gateway MAC	Manual Gateway MAC
d		Inc: 102.1.0.2, 0.0	24	Inc: 102.1.0.1, 0.0.1.0	☑		00:00:00:00:00:01
d						olved	00:00:00:00:00:01

○ Single Value
● Increment
○ Decrement
○ Random
○ Repeatable Random
○ Custom Distribution
○ Custom
○ File
○ Subset

Start Value 20.0.0.101

Step 0.0.1.0

More Options

☑ Port Step 0.1.0.0

☑ Overwrite manually changed rows

OK Cancel

Figure 256. Configuring VTEP IPv4 address

19. Similarly, set the **Gateway IP** to **20.0.0.1.** Click **OK.**

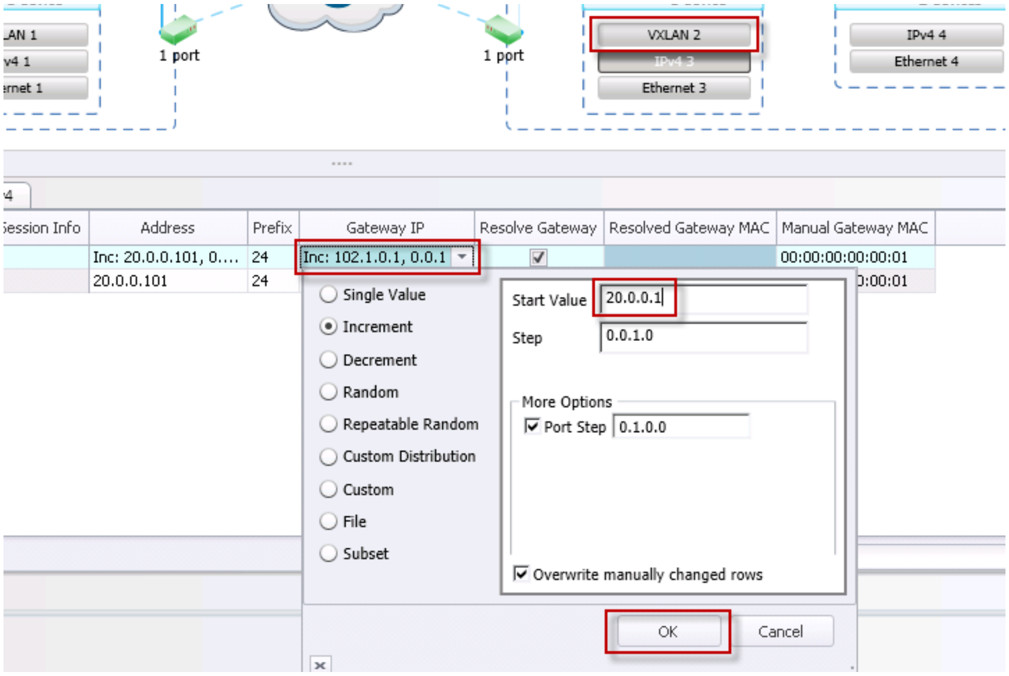

Session Info	Address	Prefix	Gateway IP	Resolve Gateway	Resolved Gateway MAC	Manual Gateway MAC
	Inc: 20.0.0.101, 0....	24	Inc: 102.1.0.1, 0.0.1	☑		00:00:00:00:00:01
	20.0.0.101	24				0:00:01

● Increment

Start Value 20.0.0.1

Step 0.0.1.0

More Options

☑ Port Step 0.1.0.0

☑ Overwrite manually changed rows

Figure 257. Configuring VTEP Gateway IP

20. Click **VXLAN** in **Scenario** View to change the VTEP VNI configuration. Click **VNI** and configure the **Start Value** to 1001 in Increment. Click **OK**.

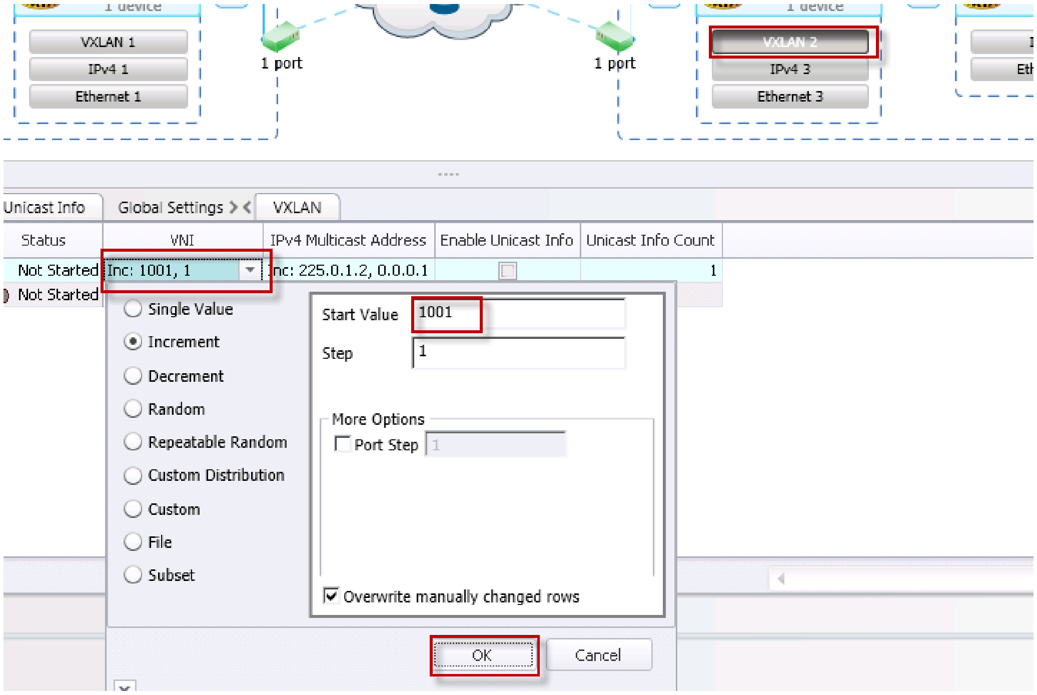

Status	VNI	IPv4 Multicast Address	Enable Unicast Info	Unicast Info Count
Not Started	Inc: 1001, 1	Inc: 225.0.1.2, 0.0.0.1	□	1
Not Started				

● Increment

Start Value 1001

Step 1

More Options

□ Port Step 1

☑ Overwrite manually changed rows

Figure 258. Configuring VTEP VNI

21. Click **Enable Unicast Info** checkbox and change **Unicast Info Count** value to **2**.

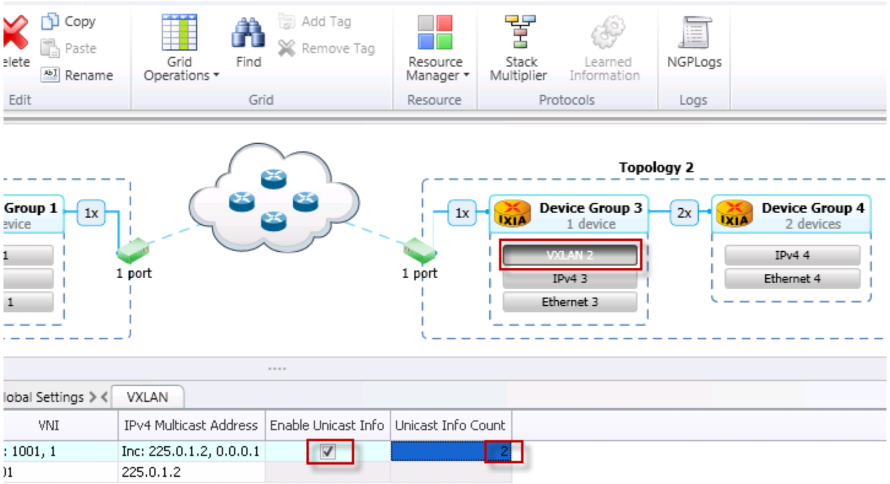

VNI	IPv4 Multicast Address	Enable Unicast Info	Unicast Info Count
: 1001, 1	Inc: 225.0.1.2, 0.0.0.1	☑	2
)1	225.0.1.2		

Figure 259. Enabling VXLAN Unicast Info instead of multicast

22. Select the **VXLAN Unicast Info** tab and select the pattern dropdown for **Remote VTEP Unicast IPv4**. Specify a **Single Value** radio button option **Start Value** of **20.0.0.1** and click **OK** to accept.

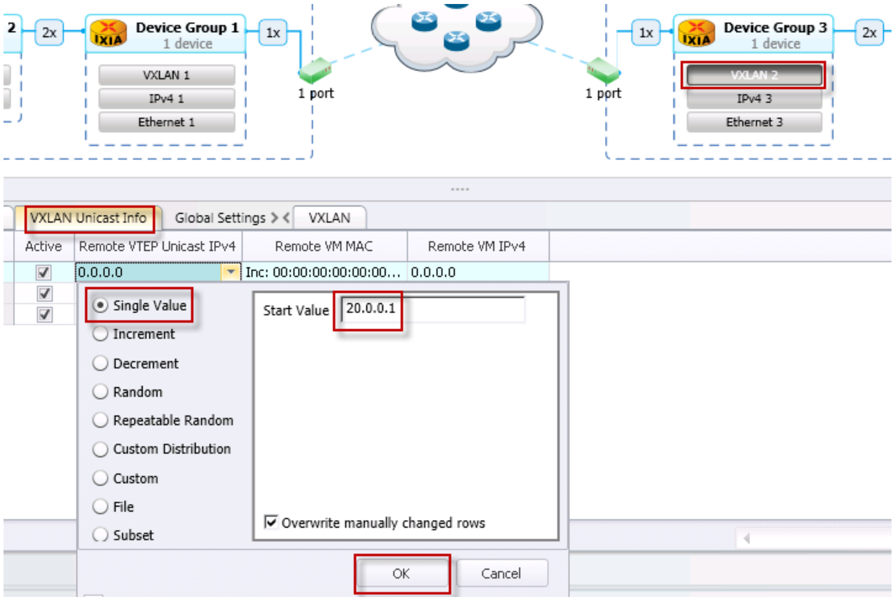

Active	Remote VTEP Unicast IPv4	Remote VM MAC	Remote VM IPv4
☑	0.0.0.0	Inc: 00:00:00:00:00:00...	0.0.0.0
☑			
☑			

● Single Value
○ Increment
○ Decrement
○ Random
○ Repeatable Random
○ Custom Distribution
○ Custom
○ File
○ Subset

Start Value 20.0.0.1

☑ Overwrite manually changed rows

OK Cancel

Figure 260. Configuring Remote VTEP Unicast IPv4 address

23. Select the **VXLAN Unicast Info** tab and select the pattern dropdown for **Remote VM MAC**. Specify an Increment radio button option **Start Value** of **00:12:01:00:00:01** and a **Step** value of **00:00:00:00:00:01**. Click **OK** to accept.

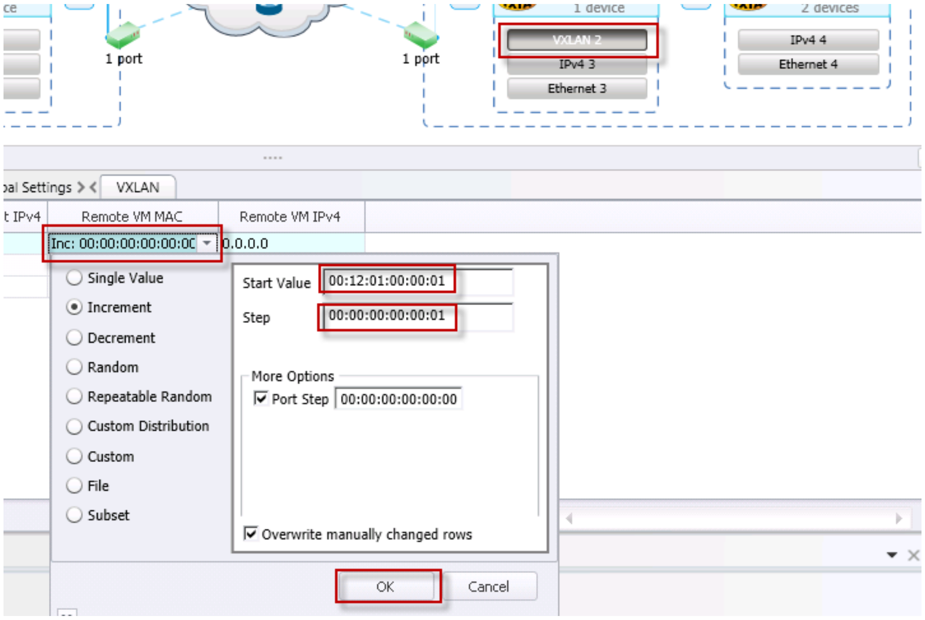

Inc: 00:00:00:00:00:0C | 0.0.0.0

○ Single Value
● Increment
○ Decrement
○ Random
○ Repeatable Random
○ Custom Distribution
○ Custom
○ File
○ Subset

Start Value 00:12:01:00:00:01
Step 00:00:00:00:00:01
More Options
☑ Port Step 00:00:00:00:00:00
☑ Overwrite manually changed rows

Figure 261. Configuring Remote VM MAC address

24. Select the VXLAN Unicast Info tab and select the pattern dropdown for **Remote VM IPv4**. Specify an increment radio button option **Start Value** of **30.0.0.1** and a **Step** value of **0.0.0.1**. Click **OK** to accept.

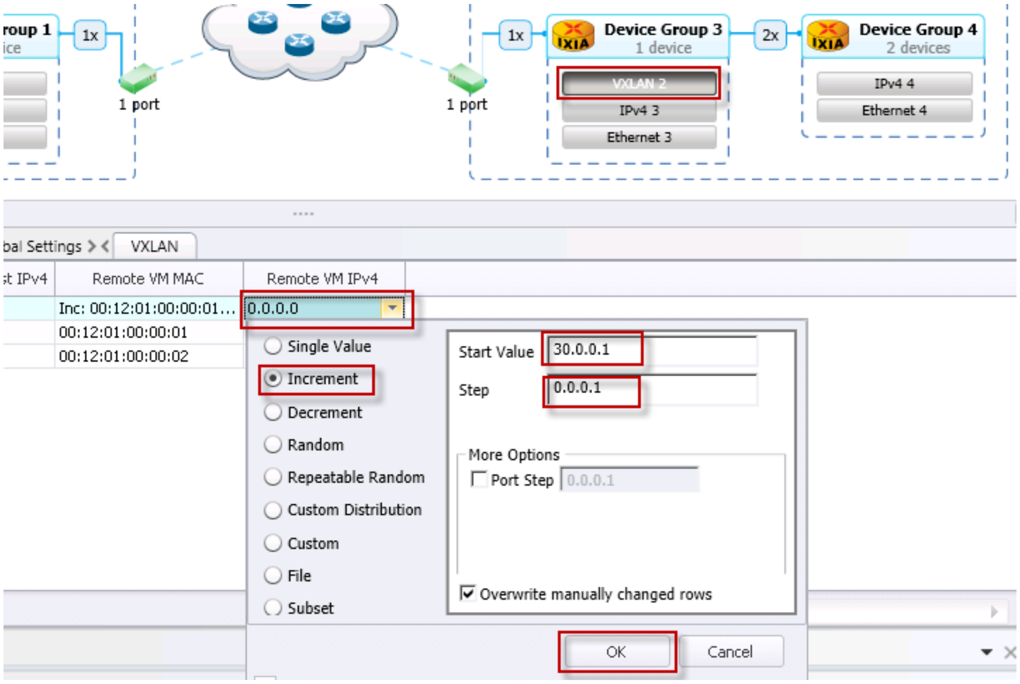

Inc: 00:12:01:00:00:01... | 0.0.0.0
00:12:01:00:00:01
00:12:01:00:00:02

○ Single Value
● Increment
○ Decrement
○ Random
○ Repeatable Random
○ Custom Distribution
○ Custom
○ File
○ Subset

Start Value 30.0.0.1
Step 0.0.0.1
More Options
☐ Port Step 0.0.0.1
☑ Overwrite manually changed rows

Figure 262. Configuring Remote VM IPv4 address

25. Click **IPv4** in Device Group3**.** Select **Address** and configure **Start Value** of **30.0.0.101** and a **Step** value in **Increment** view and click

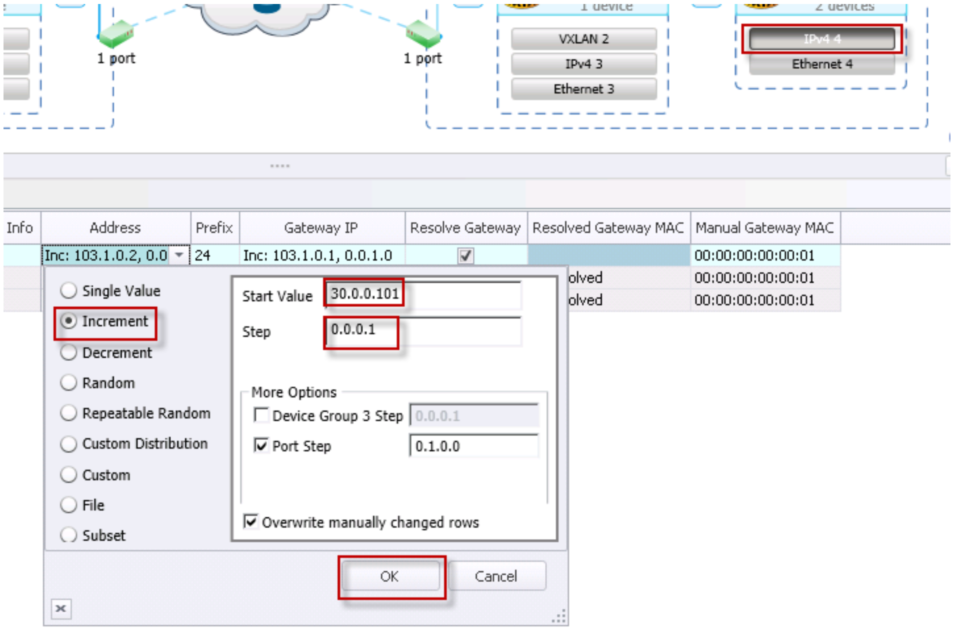

Info	Address	Prefix	Gateway IP	Resolve Gateway	Resolved Gateway MAC	Manual Gateway MAC
	Inc: 103.1.0.2, 0.0	24	Inc: 103.1.0.1, 0.0.1.0	☑		00:00:00:00:00:01
					olved	00:00:00:00:00:01
					olved	00:00:00:00:00:01

○ Single Value
● Increment
○ Decrement
○ Random
○ Repeatable Random
○ Custom Distribution
○ Custom
○ File
○ Subset

Start Value 30.0.0.101
Step 0.0.0.1

More Options
☐ Device Group 3 Step 0.0.0.1
☑ Port Step 0.1.0.0

☑ Overwrite manually changed rows

OK Cancel

Figure 263. Configuring Host IP Address

26. Select the **Gateway IP** and configure the **Start Value of 30.0.0.1** in **Increment** view and click **OK**.

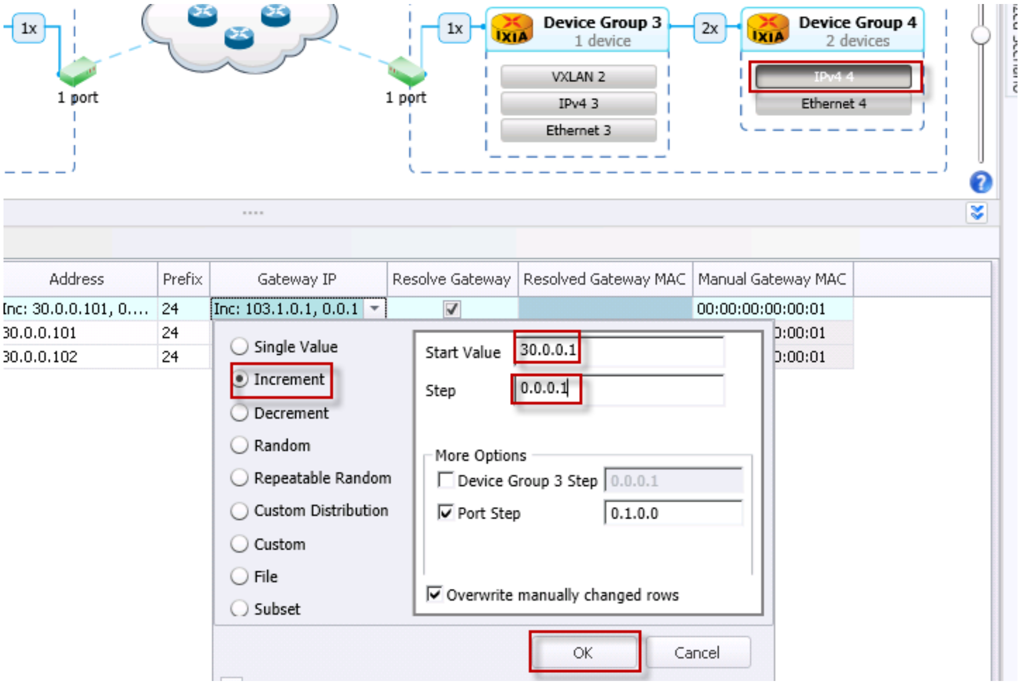

Address	Prefix	Gateway IP	Resolve Gateway	Resolved Gateway MAC	Manual Gateway MAC
Inc: 30.0.0.101, 0.…	24	Inc: 103.1.0.1, 0.0.1	☑		00:00:00:00:00:01
30.0.0.101	24				):00:01
30.0.0.102	24				):00:01

○ Single Value
● Increment
○ Decrement
○ Random
○ Repeatable Random
○ Custom Distribution
○ Custom
○ File
○ Subset

Start Value 30.0.0.1
Step 0.0.0.1

More Options
☐ Device Group 3 Step 0.0.0.1
☑ Port Step 0.1.0.0

☑ Overwrite manually changed rows

OK Cancel

Figure 264. Configuring Host Gateway IP

27. Click the **Protocol Options** button and select **VXLAN** under **Protocol Options**. Change the **Outer IGMP Mode** pattern dropdown value to **DoNotSend** to disable the sending of IGMP Joins and then click **Close** to finish.

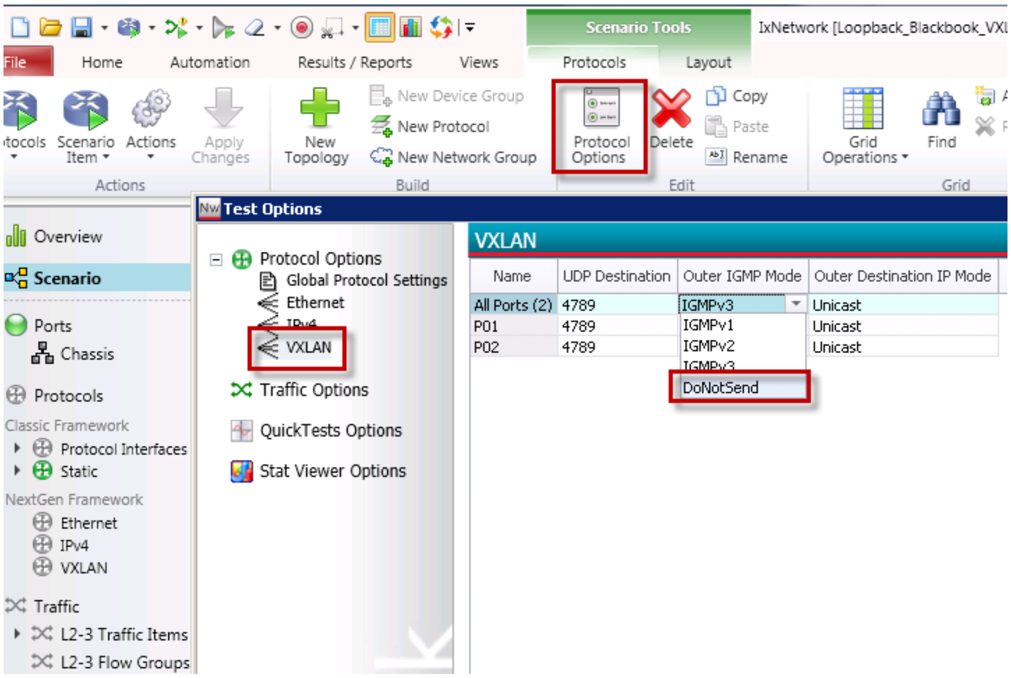

Name	UDP Destination	Outer IGMP Mode	Outer Destination IP Mode
All Ports (2)	4789	IGMPv3	Unicast
P01	4789		Unicast
P02	4789		Unicast

Figure 265. Disable the sending of IGMP Joins

28. Click **Scenario** in the left pane to view complete topology.

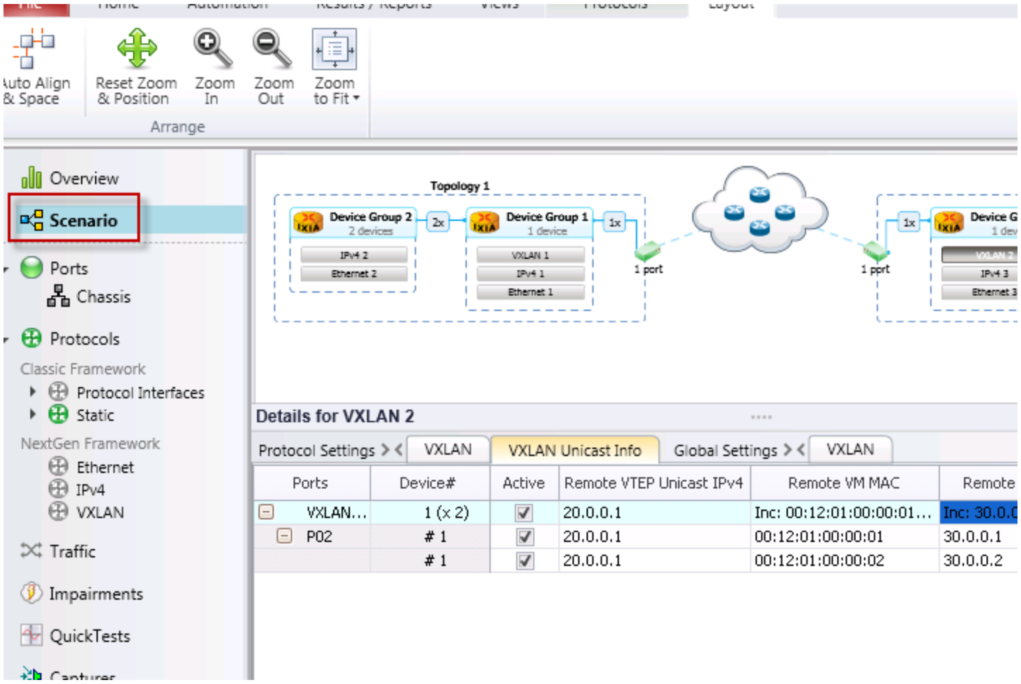

Ports	Device#	Active	Remote VTEP Unicast IPv4	Remote VM MAC	Remote
⊟ VXLAN...	1 (x 2)	☑	20.0.0.1	Inc: 00:12:01:00:00:01...	Inc: 30.0.0
⊟ P02	# 1	☑	20.0.0.1	00:12:01:00:00:01	30.0.0.1
	# 1	☑	20.0.0.1	00:12:01:00:00:02	30.0.0.2

Figure 266. Reviewing Scenario Editor Topology

29. Click **Start All Protocols** and wait for the VXLAN and IPv4 Host protocol to come up green.

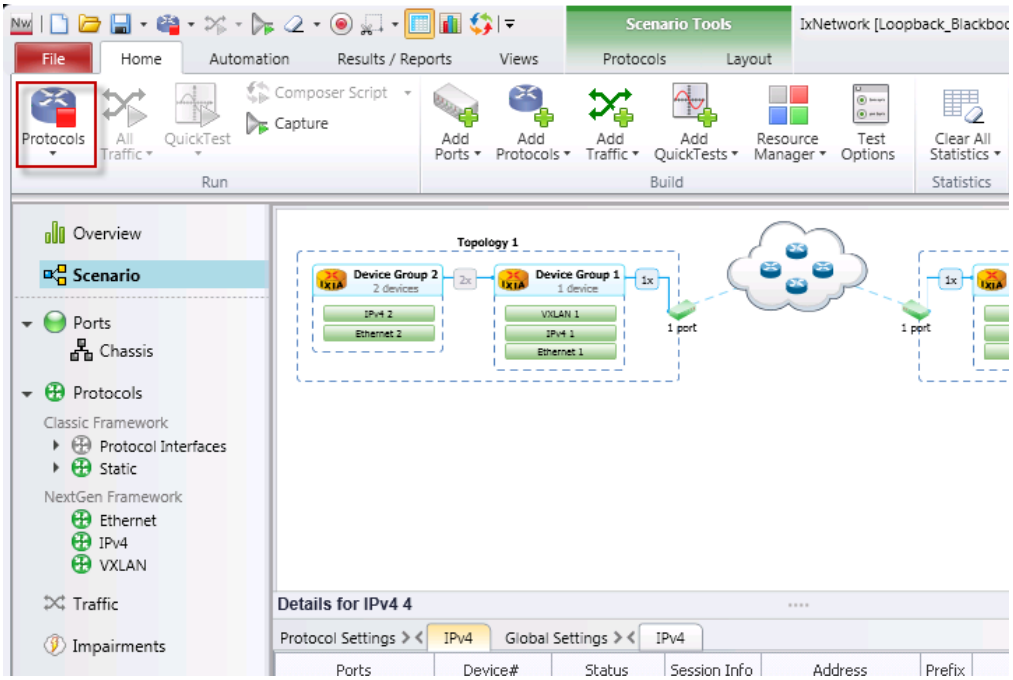

Figure 267. Starting All Protocols

30. Click **Traffic** in the left pane and click **Add L2-3 Traffic Items**.

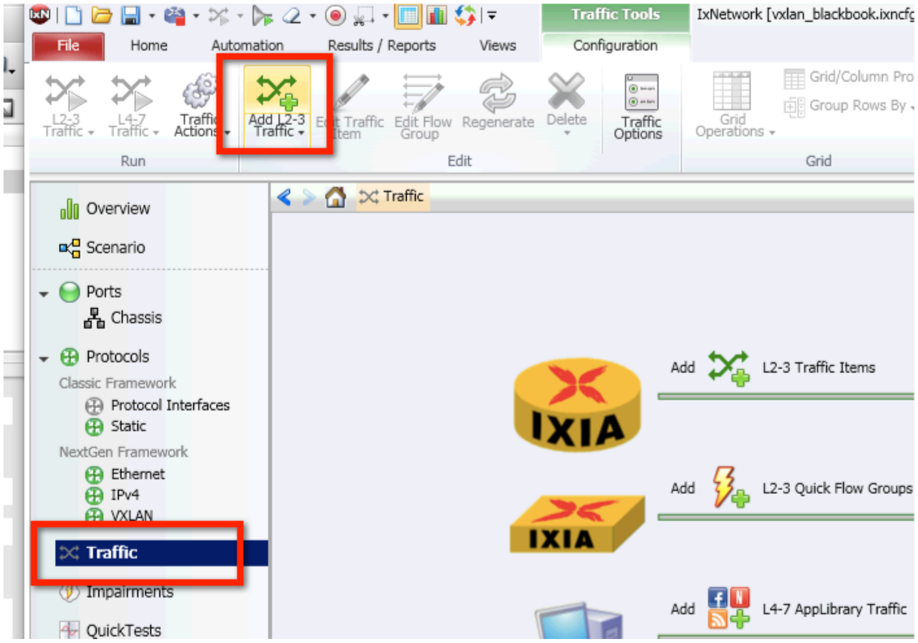

Figure 268. Configuring Traffic

31. In the **Type of Traffic** dropdown list, select **IPv4.** Select **Bi-Directional** checkbox. In the **Source** dropdown list, select **IPv4** inside the **Device Group2** of **Topology1**. In the **Destination** dropdown list, select **IPv4** in **Device Group4** of **Topology2**.

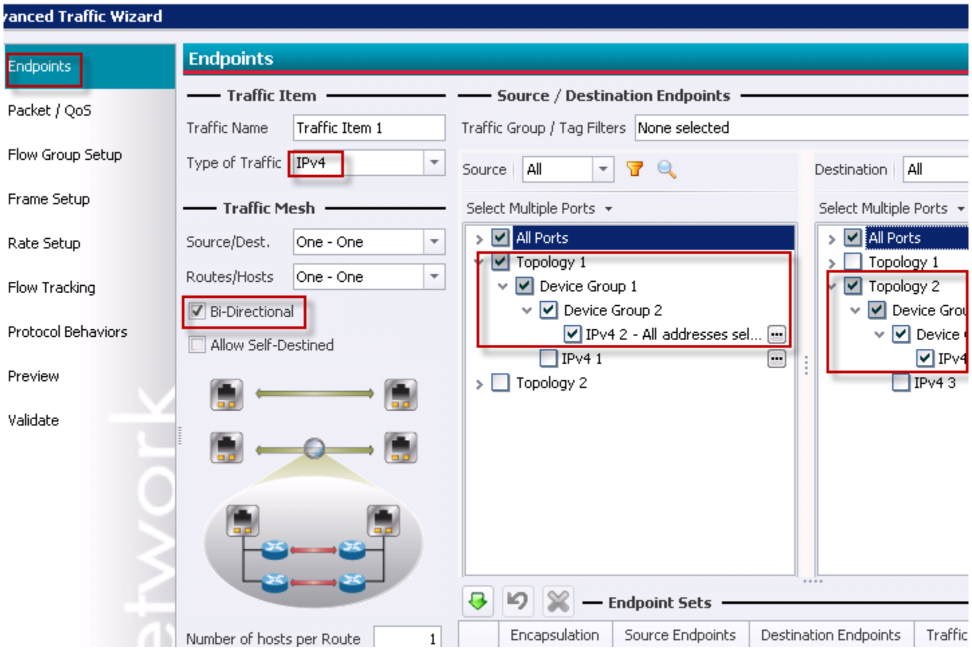

Traffic Name: Traffic Item 1

Type of Traffic: IPv4

Source/Dest.: One - One

Routes/Hosts: One - One

☑ Bi-Directional

□ Allow Self-Destined

Number of hosts per Route: 1

Traffic Group / Tag Filters: None selected

Source: All

- ☑ All Ports
- ☑ Topology 1
  - ☑ Device Group 1
  - ☑ Device Group 2
    - ☑ IPv4 2 - All addresses sel...
    - □ IPv4 1
- □ Topology 2

Destination: All

- ☑ All Ports
- □ Topology 1
- ☑ Topology 2
  - ☑ Device Grou
    - ☑ Device
      - ☑ IPv4
      - □ IPv4 3

Figure 269. Selecting IPv4 from VTEP and remote Host port

32. In the left pane, click **Frame Setup** and change the **Fixed** frame size value to **1500**.

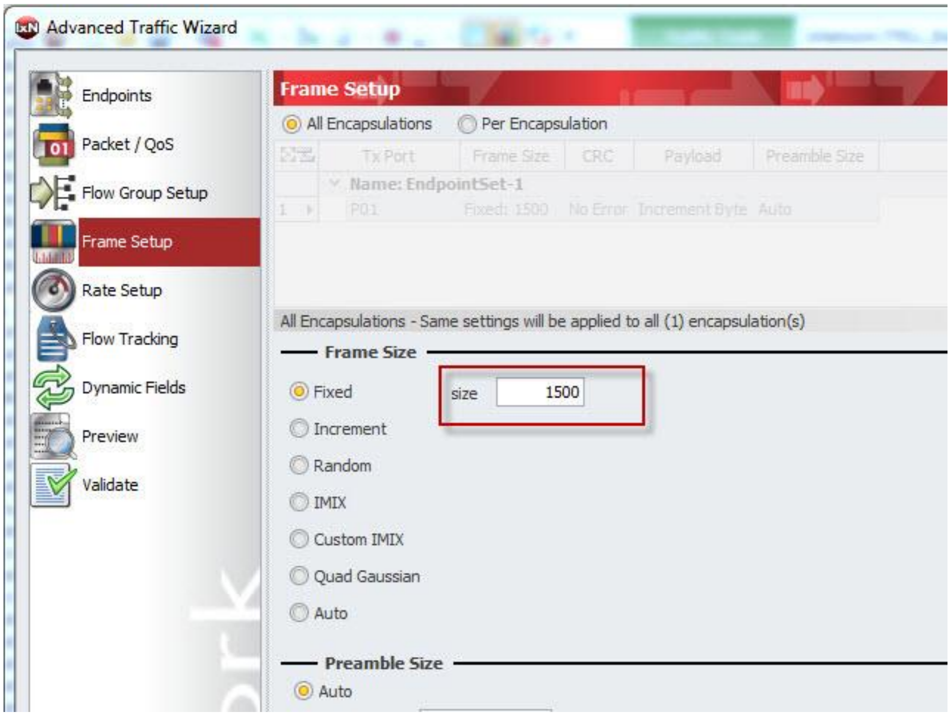

● All Encapsulations ○ Per Encapsulation

Tx Port	Frame Size	CRC	Payload	Preamble Size
Name: EndpointSet-1				
1 P01	Fixed: 1500	No Error	Increment Byte	Auto

All Encapsulations - Same settings will be applied to all (1) encapsulation(s)

Frame Size

● Fixed size 1500

○ Increment

○ Random

○ IMIX

○ Custom IMIX

○ Quad Gaussian

○ Auto

Preamble Size

● Auto

Figure 270. Changing the fixed frame size value

33. In the left pane, click **Rate Setup** and change the **Line rate** percentage value to 100.

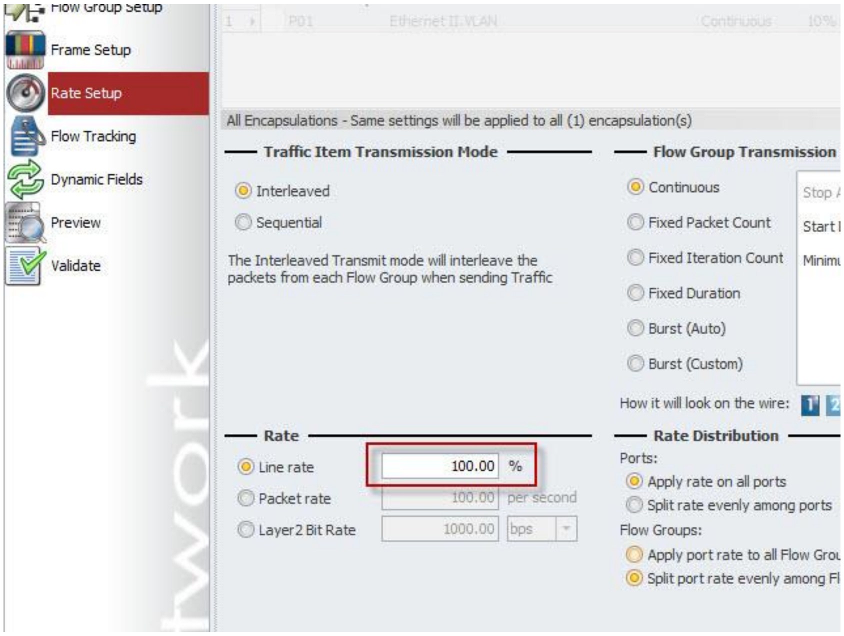
Flow Group Setup
Frame Setup
Rate Setup
Flow Tracking
Dynamic Fields
Preview
Validate
All Encapsulations - Same settings will be applied to all (1) encapsulation(s)
Traffic Item Transmission Mode
● Interleaved
○ Sequential
The Interleaved Transmit mode will interleave the packets from each Flow Group when sending Traffic
Flow Group Transmission
● Continuous
○ Fixed Packet Count
○ Fixed Iteration Count
○ Fixed Duration
○ Burst (Auto)
○ Burst (Custom)
How it will look on the wire:
Rate
● Line rate 100.00 %
○ Packet rate 100.00 per second
○ Layer2 Bit Rate 1000.00 bps
Rate Distribution
Ports:
● Apply rate on all ports
○ Split rate evenly among ports
Flow Groups:
Apply port rate to all Flow Grou
Split port rate evenly among Fl

Figure 271. Changing the line rate throughput value

34. In the left pane, click **Flow Tracking** and select **Traffic Item** tracking options. Select **IPv4: Destination Address(1) and VXLAN: VNI** tracking options. These options help you to drill down for more details dynamically on VXLAN and IPv4 Host endpoint. Click **Finish** to close the **Advanced Traffic Wizard**.

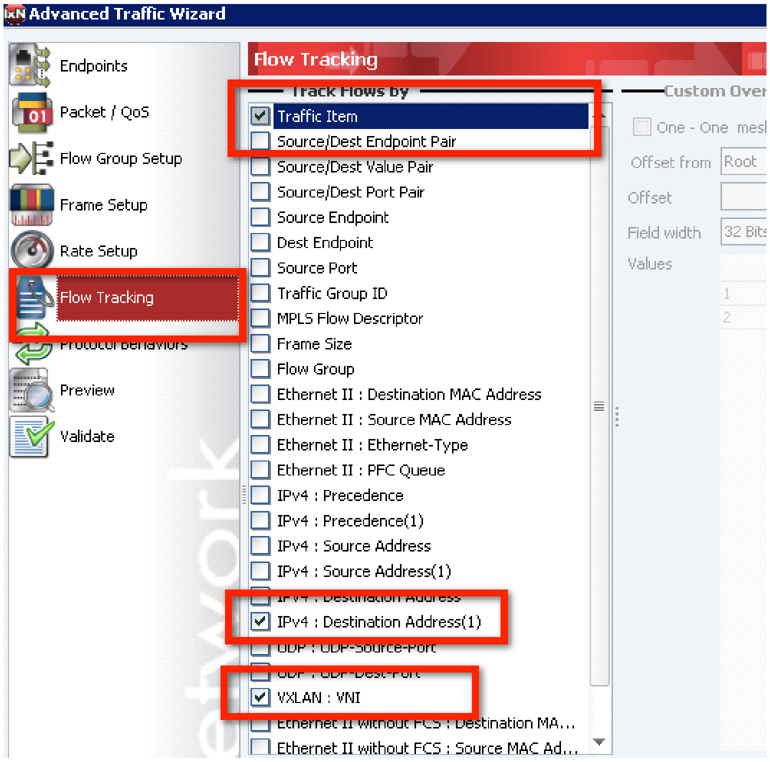

Figure 272. Selecting flow tracking options

Test Variables

Performance Variable	Description
No of IPs in VXLAN stack	Use this test variable to increase the count of VTEPs
Multiplier between VXLAN and IP stack	Use this test variable to increase the count of VNIs per VTEPs to more then 1
Multiplier between VXLAN Device Group and IP DeviceGroup	Use this test variable to increase the count of IP Host per VNIs
DHCP Client	Add DHCP Client device instead of IPv4 behind VTEP Gateway to emulate DHCP client Host.
IPv6 Client	Add IPv6 device instead of IPv4 behind VTEP Gateway to emulate IPv6 Host.

Troubleshooting and Diagnostics

Issue	Troubleshooting Solution
IPv4 Host ARP does not resolve	Check DUT settings to ensure: • VNIs configuration matches with IxNetwork • Unicast IP and MAC configuration matches with IxNetwork • UDP port configuration matches with IxNetwork, look for global protocol settings • Host IP address doesn’t conflict with VTEP IP address

Results Analysis

1. Click **Start All Traffic** and, then click **L2-L3 Test Summary Statistics**.

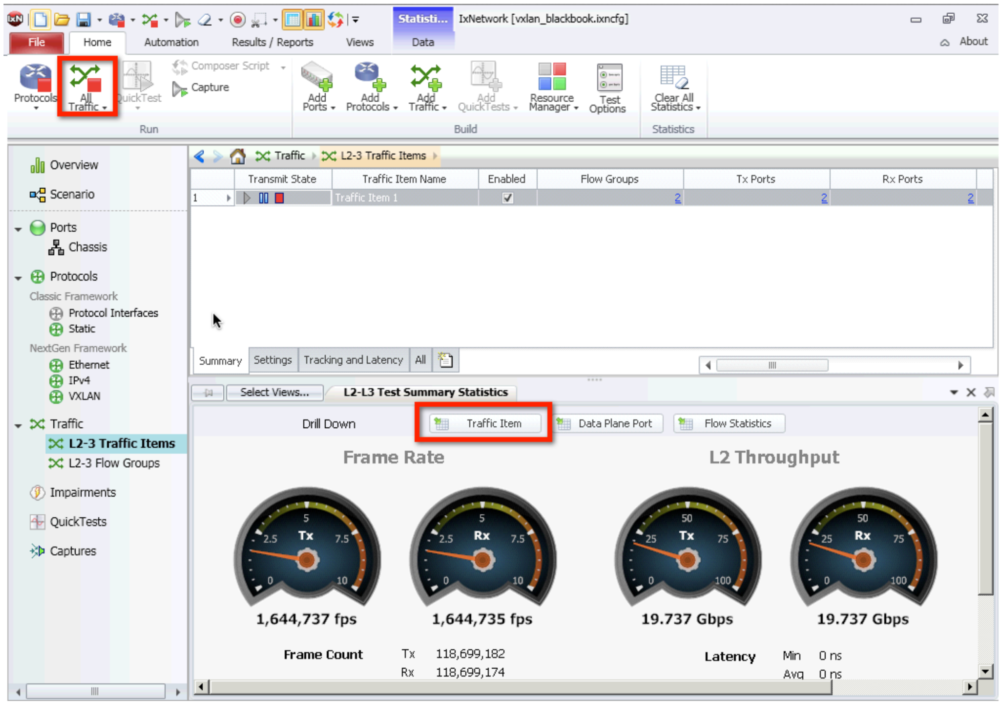

	Transmit State	Traffic Item Name	Enabled	Flow Groups	Tx Ports	Rx Ports
1		Traffic Item 1	✓	2	2	2

```
Frame Rate                          L2 Throughput
Tx 1,644,737 fps   Rx 1,644,735 fps    Tx 19.737 Gbps   Rx 19.737 Gbps

Frame Count   Tx   118,699,182          Latency   Min   0 ns
              Rx   118,699,174                    Avg   0 ns
```

Figure 273. Viewing L2-L3 Test Summary Statistics

2. Select **Traffic Item 1** from the **Traffic Item Statistics** view. Right-click to open the context menu, and then click **Drill down per VXLAN: VNI** to see throughput and packet loss statistics on a per VNI basis in the **User Defined Statistics** view.

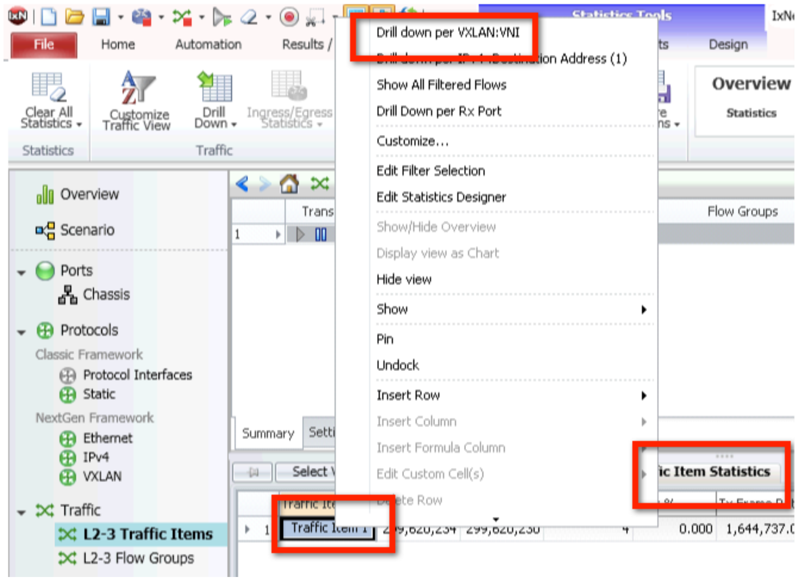

Figure 274. Reviewing Traffic Items Statistics

3. Select **VNI 1001** from the **User Defined Statistics** view. Right-click to open the context menu, and the click **Drill down per IPv4: Destination Address (1)** to see throughput and packet loss statistics on a per IPv4 Host per VNI bases in the **User Defined Statistics** view.

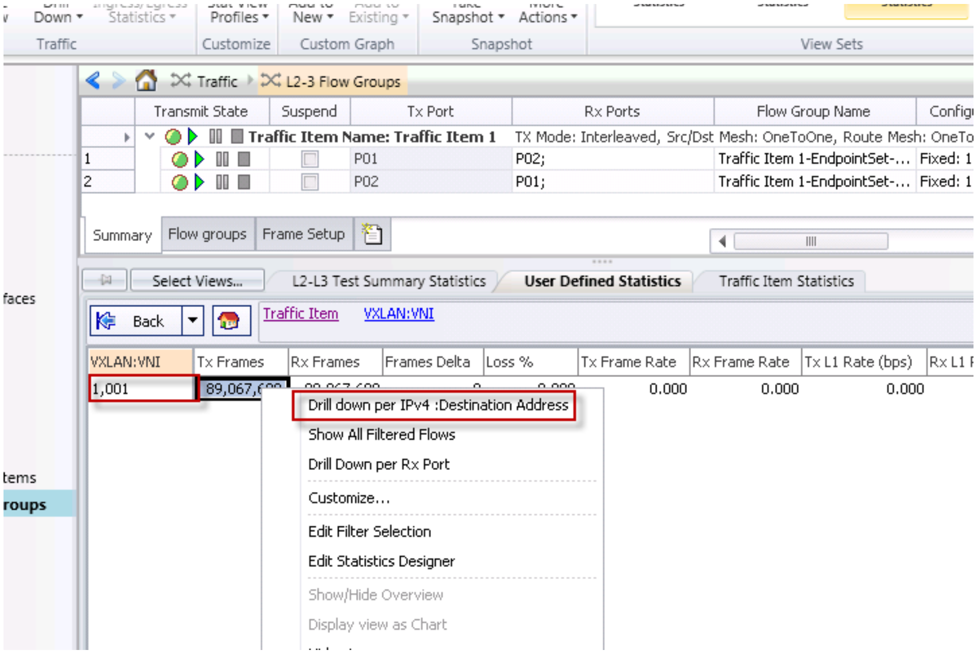

Figure 275. Selecting per-IPv4 Destination statistics

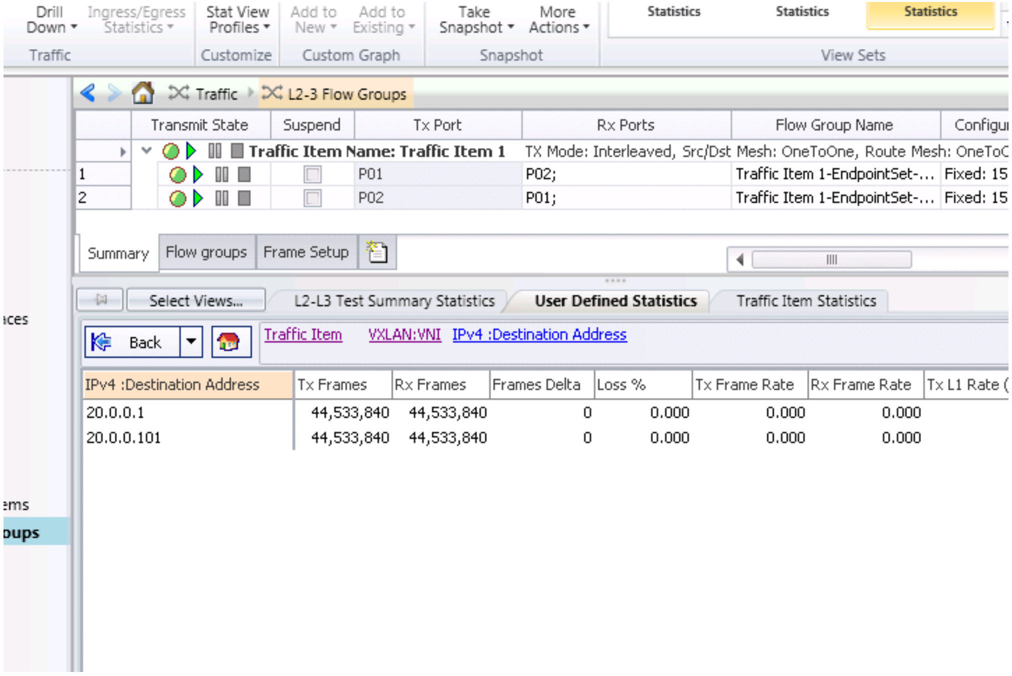

IPv4 :Destination Address	Tx Frames	Rx Frames	Frames Delta	Loss %	Tx Frame Rate	Rx Frame Rate	Tx L1 Rate (
20.0.0.1	44,533,840	44,533,840	0	0.000	0.000	0.000	
20.0.0.101	44,533,840	44,533,840	0	0.000	0.000	0.000	

Figure 276. Reviewing per-IPv4 Destination statistics

Conclusions

The IxNetwork VXLAN Emulation enables the configuration of complex VTEP topologies along with all of the VNI and various Host protocols per VNI, such as IPv4, IPv6, and DHCP. VTEP Discovery can now be accomplished by **Unicast packets** in between VTEP gateways in addition to Multicast packets from the original specification. Using Unicast discovery allows VXLAN to be supported on top of underlays that do not have multicast capability. The flexible flow tracking in the **Advanced Traffic Wizard**, along with the real time drill down capability in statistics views, allows the user to quickly isolate flows that are not being forwarded correctly by the DUT and detect the root cause by associating them with individual VNI within.

Introduction to Layer 2 Multi-Path (L2MP)

Data center operators are faced with the challenge of expanding infrastructure capacity while minimizing data center footprint and cost. Both government departments and commercial organizations have initiated aggressive plans to consolidate data center resources and slash the number of sites required. Layer 2 Multipath (L2MP) / Equal Cost Multi Path (ECMP) technologies such as IETF TRILL, IEEE SPBM, and Cisco FabricPath have shown a lot of interest, because they pave the way for a two-tier flat Ethernet infrastructure that is layer 2 based, highly scalable, highly redundant with active-active multipath support.

L2MP protocols provide loop-free connectivity like the well-known Spanning Tree Protocols (STP), but address many limitations of STP when supporting a large scale next generation data center fabric. L2MP protocols alleviate core switches from having to learn the MAC address of every endpoint machine, both physical and virtual, at every hop of the network. Also, with L2MP, every server in the data center can receive non-blocking access to any other server in the fabric in the shortest path possible. In addition, every link in the network can be in active state carrying live traffic, greatly optimizing the resources available while preventing congestion and loops.

Test Case: Layer 2 Traffic Forwarding over Transparent Interconnection of Lots of Links (TRILL)

Objective

In this test case, we will validate the core functionality of a TRILL RBridge to negotiate TRILL protocol and compute reachability to Unicast MAC destinations, Multicast MAC destinations, IPv4 and IPv6 Multicast destinations. Traffic is then sourced from a set of emulated customer sites, on the first port connected to the DUT to a Campus network consisting of RBridges and end stations, emulated on the second port connected to the DUT.

Setup

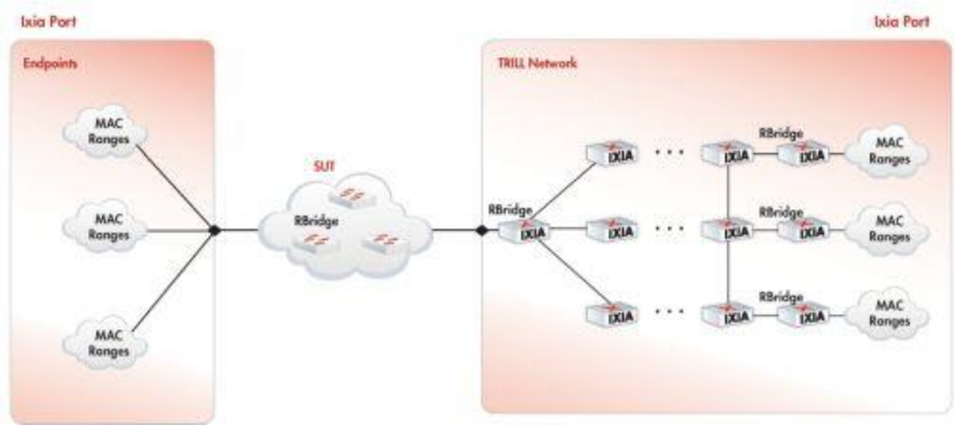

Figure 277. Campus Network running TRILL Layer 2 fabric

Step-by-Step Instructions

1. Click **Add Ports** to open the **Port Selection** Window. Click **Add Chassis** and enter the IP Address or name for your IXIA chassis. Click **OK** to accept. Expand chassis and cards and select two test ports you want to use in this test. Click **Add ports** and then **OK** to add the ports to your test configuration.

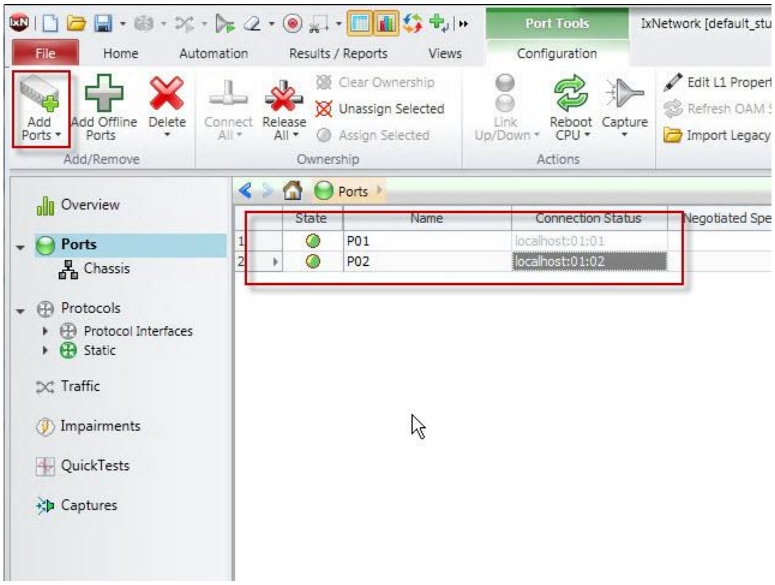

	State	Name	Connection Status	Negotiated Spe
1	●	P01	localhost:01:01	
2	●	P02	localhost:01:02	

Figure 278. Adding ports to a new configuration

2. Click the **Protocols** view, and then click **Add Protocols** to open the **Protocol Wizards** dialog. Select **TRILL ISIS** in Routing/Switching and click **Run Wizard**.

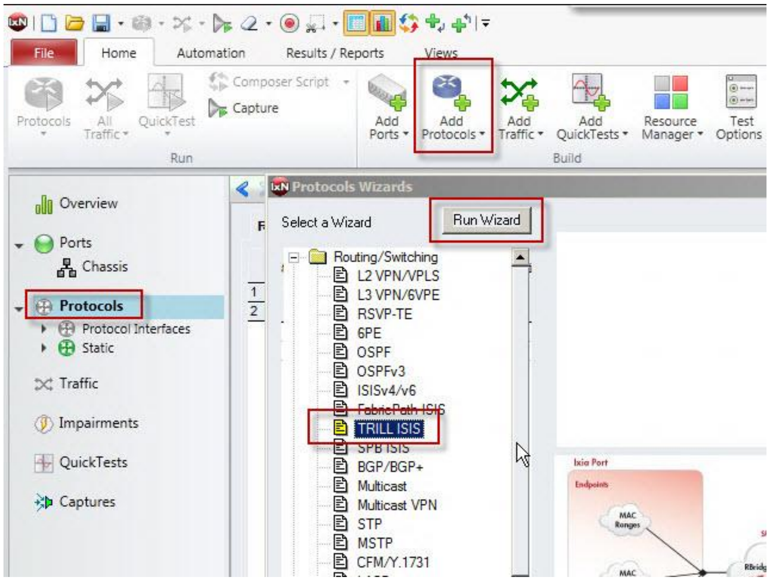

Figure 279. Selecting TRILL ISIS Protocol Wizard

3. sSelect the **Customer Side** checkbox for port 1 to emulate MAC stations outside of the TRILL campus. Select the **Core Side** checkbox for port 2 to emulate the TRILL campus behind the ingress DUT.

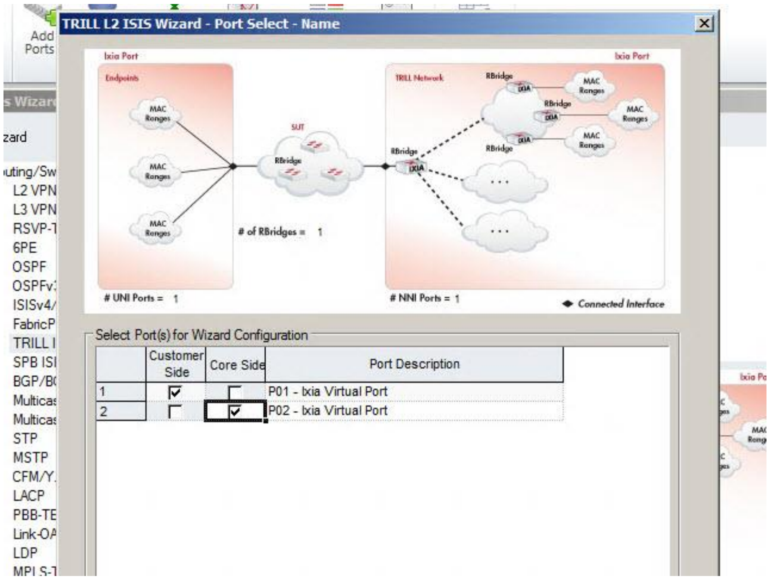

Select Port(s) for Wizard Configuration

	Customer Side	Core Side	Port Description
1	☑	□	P01 - Ixia Virtual Port
2	□	☑	P02 - Ixia Virtual Port

Figure 280. Selecting TRILL port roles for Customer and Core sides

4. Click **Next.**

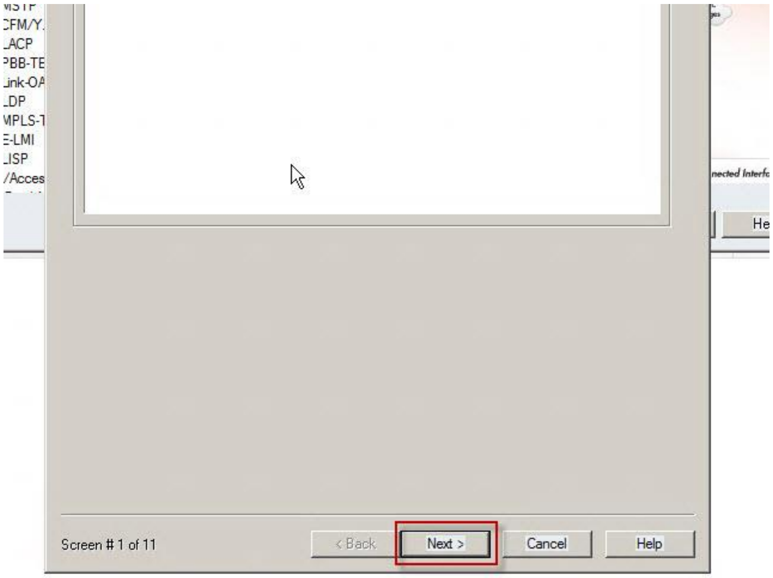

Screen # 1 of 11

Figure 281. Clicking Next to advance to next wizard screen

5. Increase the **No of RBridges** per port value to **2**. The default behavior is to **Auto Generate System Id and Nickname** for each Rbridge so leave this as selected.

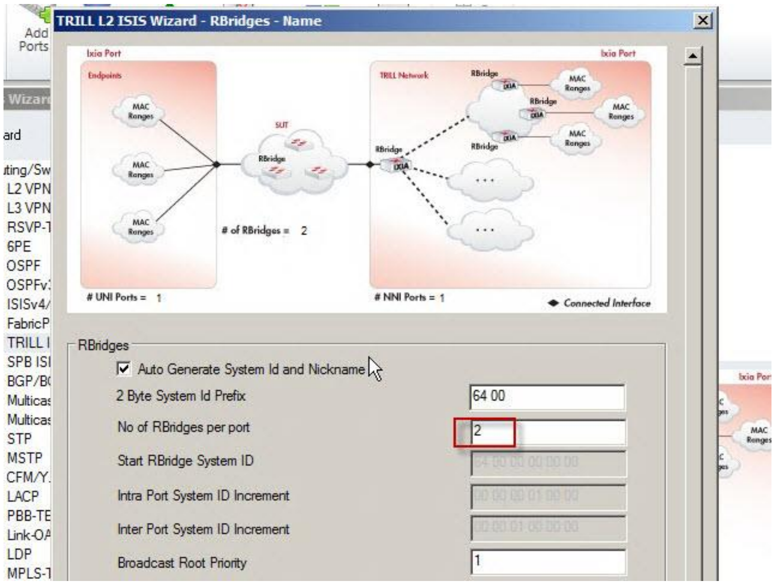

TRILL L2 ISIS Wizard - RBridges - Name

RBridges

☑ Auto Generate System Id and Nickname

2 Byte System Id Prefix: 64 00

No of RBridges per port: 2

Start RBridge System ID: 64 00 00 00 00 00

Intra Port System ID Increment: 00 00 00 01 00 00

Inter Port System ID Increment: 00 00 01 00 00 00

Broadcast Root Priority: 1

Figure 282. Quickly increasing the number of emulated RBridges on a port

6. Select the **Enable 3-Way Handshake** checkbox for point-to-point connections to the DUT. **Enable VLAN** checkbox is automatically selected. Enter a **VLAN ID** value of **101**. Click **Next.**

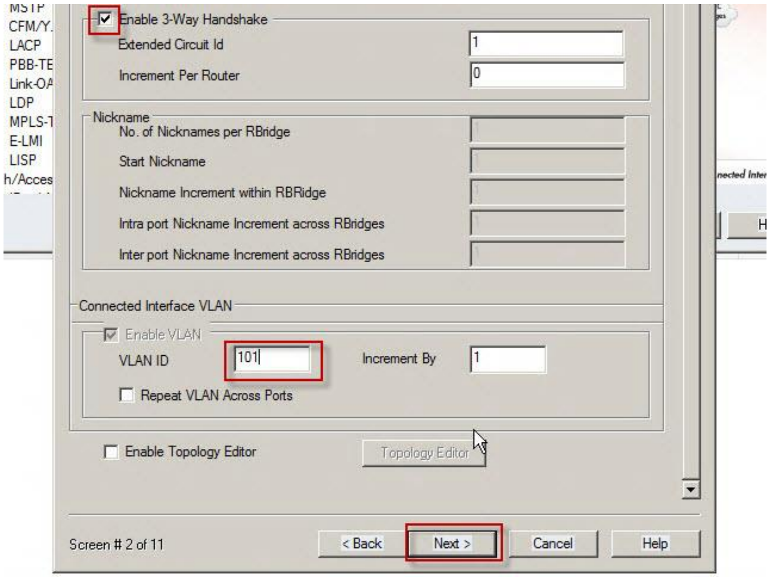

☑ Enable 3-Way Handshake

Extended Circuit Id: 1

Increment Per Router: 0

Nickname

No. of Nicknames per RBridge: 1

Start Nickname: 1

Nickname Increment within RBRidge: 1

Intra port Nickname Increment across RBridges: 1

Inter port Nickname Increment across RBridges: 1

Connected Interface VLAN

☑ Enable VLAN

VLAN ID: 101

Increment By: 1

☐ Repeat VLAN Across Ports

☐ Enable Topology Editor

Screen # 2 of 11

Figure 283. Enabling a point-to-point connection with DUT over VLAN

7. Enter 201for **Start VLAN ID**. Increase the **No of Campus VLANs per RBridge** to **2** to indicate the range of Interested VLANs.

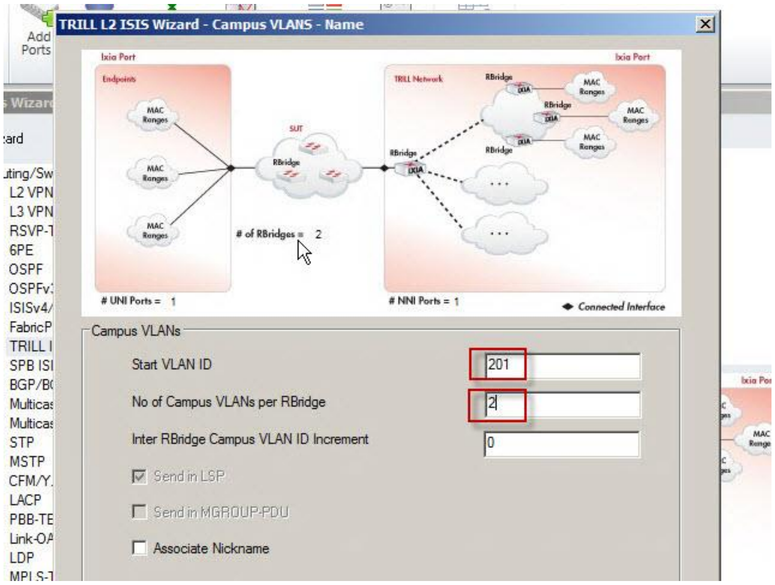

Campus VLANs

Start VLAN ID: 201

No of Campus VLANs per RBridge: 2

Inter RBridge Campus VLAN ID Increment: 0

☑ Send in LSP

□ Send in MGROUP-PDU

□ Associate Nickname

Figure 284. Configuring Interested VLANs within emulated RBridge Campus

8. Click **Next.**

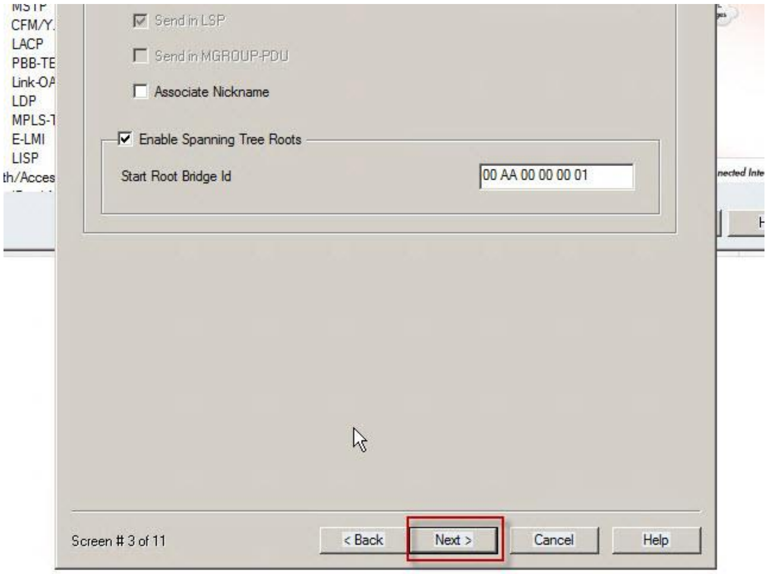

☑ Send in LSP

□ Send in MGROUP-PDU

□ Associate Nickname

☑ Enable Spanning Tree Roots

Start Root Bridge Id: 00 AA 00 00 00 01

Screen # 3 of 11

Figure 285. Accepting default Start Root Bridge Id

9. Select the **Advertise Unicast MACs for Campus** checkbox. Increase the **No of MACs per VLAN** value to 2. Change an octect in the **Start MAC Address** to make it a unique start of range value.

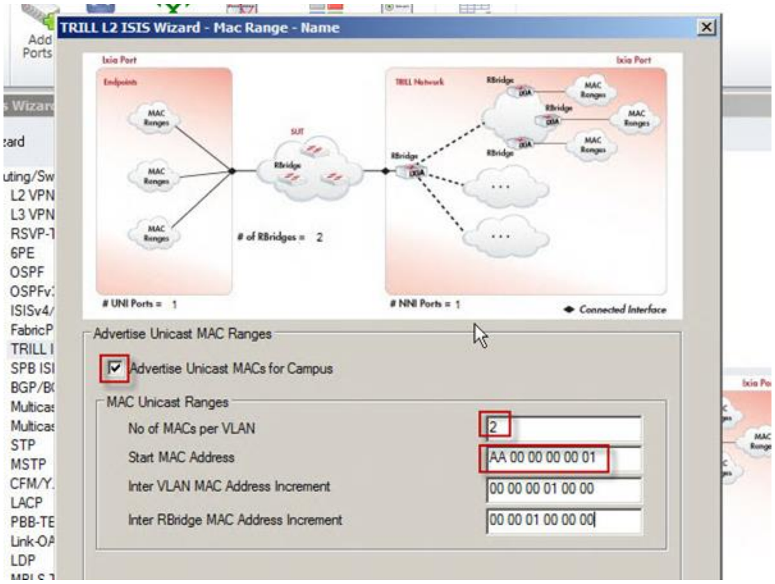

TRILL L2 ISIS Wizard - Mac Range - Name

Advertise Unicast MAC Ranges

☑ Advertise Unicast MACs for Campus

MAC Unicast Ranges

No of MACs per VLAN: 2

Start MAC Address: AA 00 00 00 00 01

Inter VLAN MAC Address Increment: 00 00 00 01 00 00

Inter RBridge MAC Address Increment: 00 00 01 00 00 00

Figure 286. Advertising unique Unicast MAC range values

10. Click **Next.**

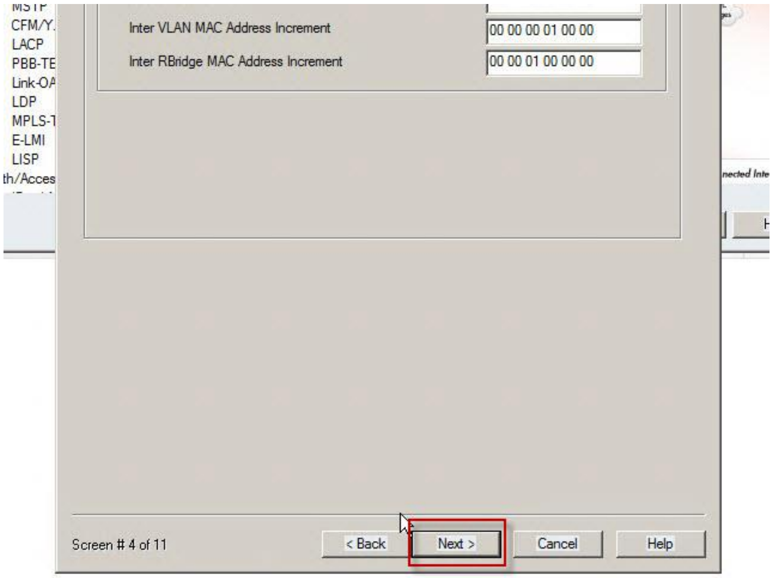

Inter VLAN MAC Address Increment: 00 00 00 01 00 00

Inter RBridge MAC Address Increment: 00 00 01 00 00 00

Screen # 4 of 11

Figure 287. Accepting the Unicast MAC range values

11. Select the **Advertise Multicast Group MACs for Campus VLANs** checkbox. Increase the **No of Multicast Group Receivers per VLAN** value to **2**.

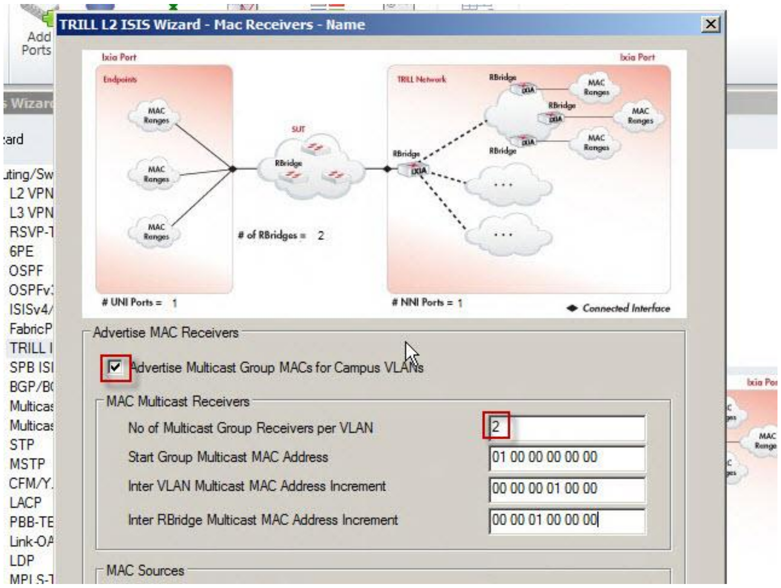

TRILL L2 ISIS Wizard - Mac Receivers - Name

Advertise MAC Receivers

☑ Advertise Multicast Group MACs for Campus VLANs

MAC Multicast Receivers

No of Multicast Group Receivers per VLAN: 2

Start Group Multicast MAC Address: 01 00 00 00 00 00

Inter VLAN Multicast MAC Address Increment: 00 00 00 01 00 00

Inter RBridge Multicast MAC Address Increment: 00 00 01 00 00 00

MAC Sources

Figure 288. Advertising MAC Multicast Receivers

12. Increase the **Number of Sources per Group** value to **2.** Click **Next**.

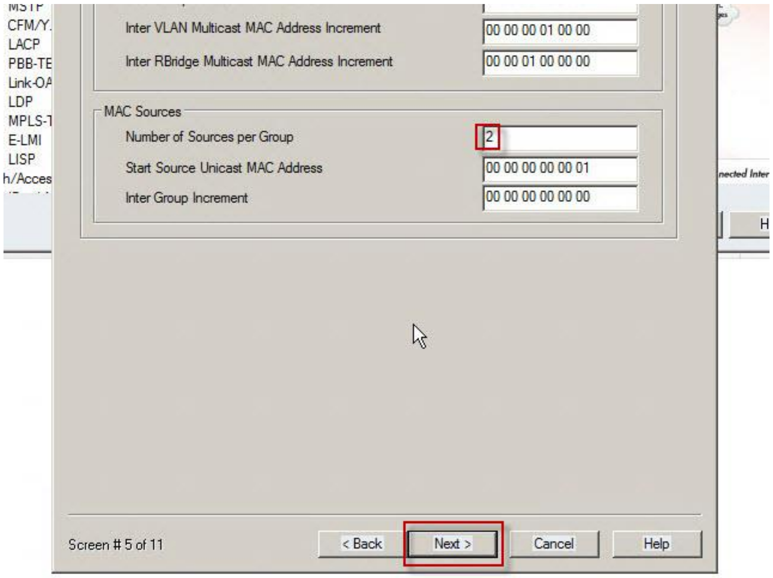

Inter VLAN Multicast MAC Address Increment: 00 00 00 01 00 00

Inter RBridge Multicast MAC Address Increment: 00 00 01 00 00 00

MAC Sources

Number of Sources per Group: 2

Start Source Unicast MAC Address: 00 00 00 00 00 01

Inter Group Increment: 00 00 00 00 00 00

Screen # 5 of 11

Figure 289. Accepting multiple MAC Sources per Group

13. Select the **Advertise IPv4 Multicast Groups for Campus VLANs** checkbox. Increase the **No of Multicast IPv4 Group Receivers per VLAN** value to **2**.

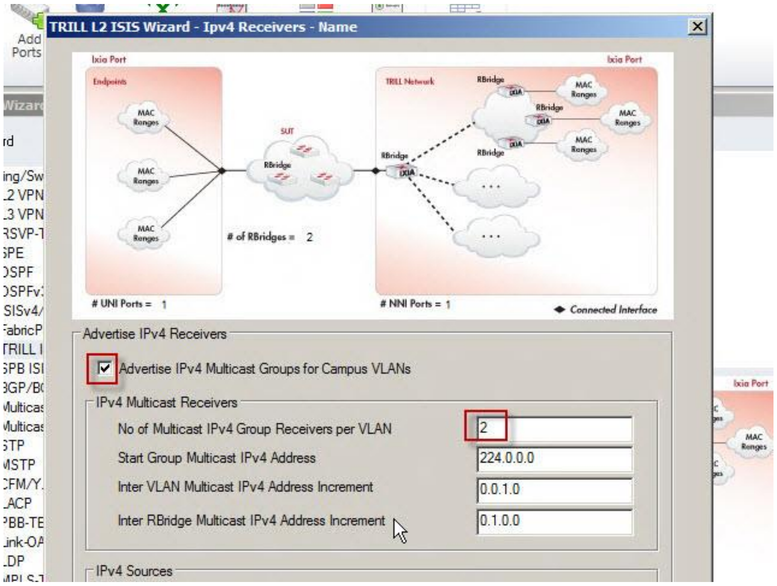

TRILL L2 ISIS Wizard - Ipv4 Receivers - Name

Advertise IPv4 Receivers

☑ Advertise IPv4 Multicast Groups for Campus VLANs

IPv4 Multicast Receivers

No of Multicast IPv4 Group Receivers per VLAN: 2

Start Group Multicast IPv4 Address: 224.0.0.0

Inter VLAN Multicast IPv4 Address Increment: 0.0.1.0

Inter RBridge Multicast IPv4 Address Increment: 0.1.0.0

IPv4 Sources

Figure 290. Advertising IPv4 Multicast Receivers

14. Increase the **Number of Sources per Group** value to **2.** Click **Next.**

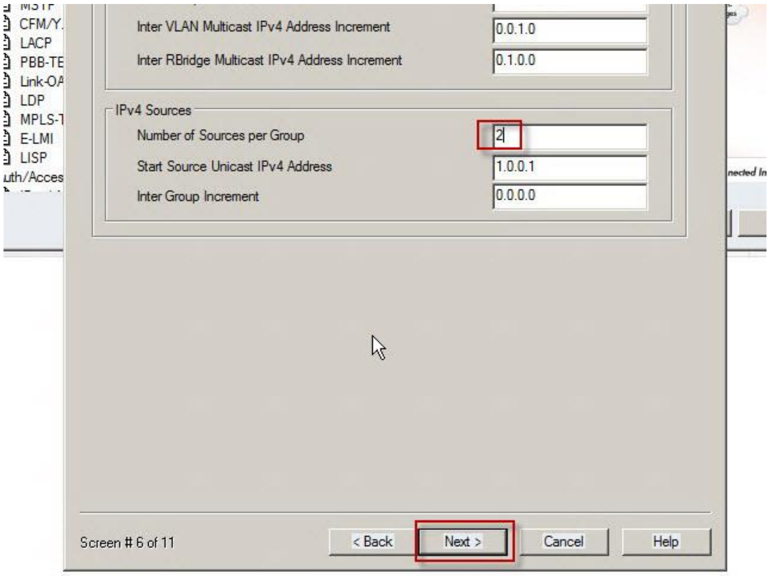

Inter VLAN Multicast IPv4 Address Increment: 0.0.1.0

Inter RBridge Multicast IPv4 Address Increment: 0.1.0.0

IPv4 Sources

Number of Sources per Group: 2

Start Source Unicast IPv4 Address: 1.0.0.1

Inter Group Increment: 0.0.0.0

Screen # 6 of 11

Figure 291. Accepting multiple IPv4 Sources per Group

15. Select the **Advertise IPv6 Multicast Groups for Campus VLANs** checkbox. Increase the **No of Multicast IPv6 Group Receivers per VLAN** to **2**.

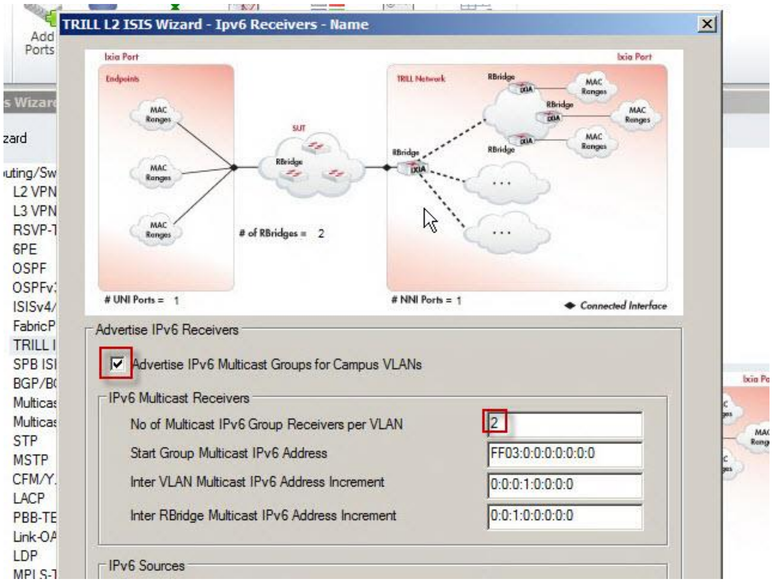

Advertise IPv6 Receivers

☑ Advertise IPv6 Multicast Groups for Campus VLANs

IPv6 Multicast Receivers

No of Multicast IPv6 Group Receivers per VLAN: 2
Start Group Multicast IPv6 Address: FF03:0:0:0:0:0:0:0
Inter VLAN Multicast IPv6 Address Increment: 0:0:0:1:0:0:0:0
Inter RBridge Multicast IPv6 Address Increment: 0:0:1:0:0:0:0:0

IPv6 Sources

Figure 292. Advertising IPv6 Multicast Receivers

16. Increase the **Number of Sources per Group** value to **2.** Click **Next.**

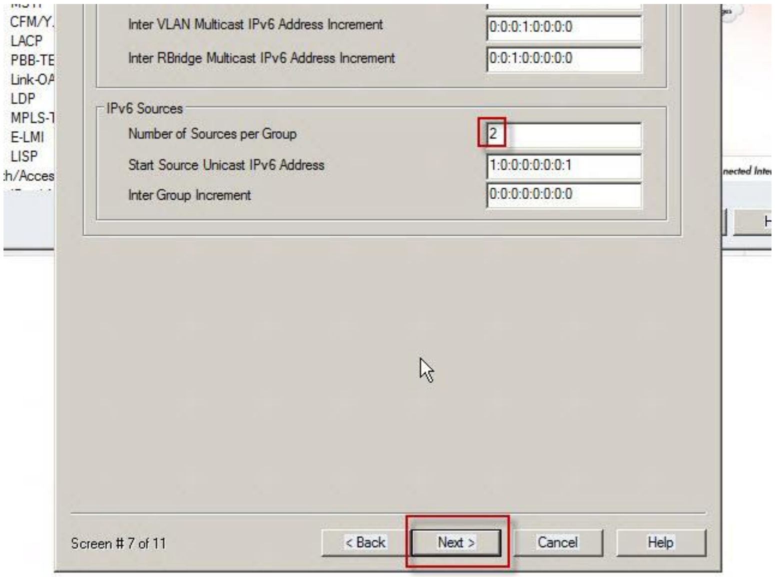
Inter VLAN Multicast IPv6 Address Increment: 0:0:0:1:0:0:0:0
Inter RBridge Multicast IPv6 Address Increment: 0:0:1:0:0:0:0:0

IPv6 Sources

Number of Sources per Group: 2
Start Source Unicast IPv6 Address: 1:0:0:0:0:0:0:1
Inter Group Increment: 0:0:0:0:0:0:0:0

Screen # 7 of 11

Figure 293. Accepting multiple IPv6 Sources per Group

17. Select the **Advertise L2 Network Ranges** checkbox. The default behavior is to advertise a 3 rows x 3 columns grid of RBridges within the Campus network. Leave the default behavior to **Auto Generate Host Name** as selected.

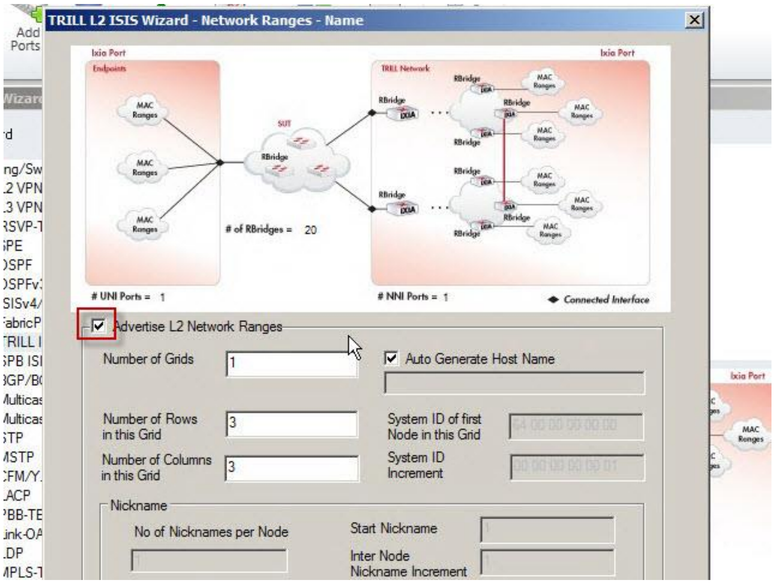

☑ Advertise L2 Network Ranges

Number of Grids: 1

☑ Auto Generate Host Name

Number of Rows in this Grid: 3

System ID of first Node in this Grid: 64 00 00 00 00 00

Number of Columns in this Grid: 3

System ID Increment: 00 00 00 00 00 01

Nickname

No of Nicknames per Node: 1

Start Nickname: 1

Inter Node Nickname Increment: 1

Figure 294. Advertising a 3 rows x 3 columns grid of RBridges

18. Select the **Shared** radio button to indicate that the **Topology Behind all other Ports** are reachable from each Emulated RBridge. Each RBridge simulates the links connecting to every RBridge on every core port. This feature is very useful for testing Equal Cost Multi-Path (ECMP) when multiple test ports are used to emulate core connections to the same grid. Click **Next.**

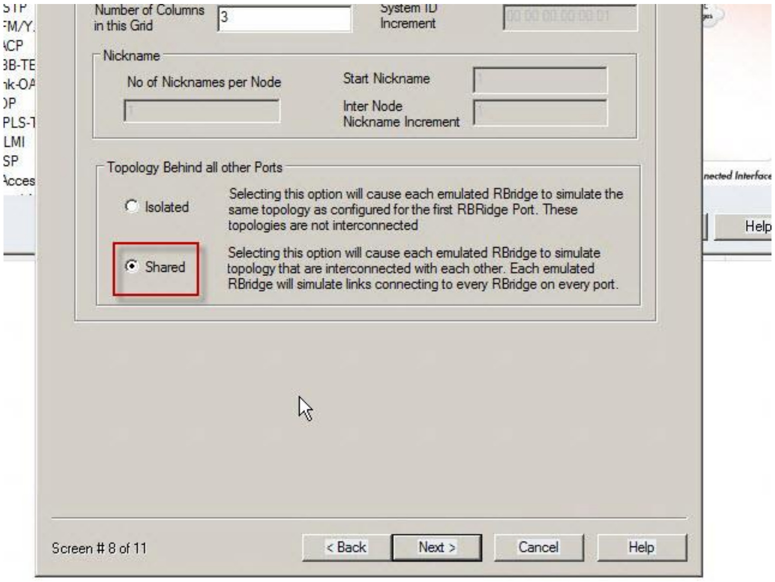

Number of Columns in this Grid: 3

System ID Increment: 00 00 00 00 00 01

Nickname

No of Nicknames per Node: 1

Start Nickname: 1

Inter Node Nickname Increment

Topology Behind all other Ports

○ Isolated — Selecting this option will cause each emulated RBridge to simulate the same topology as configured for the first RBridge Port. These topologies are not interconnected

● Shared — Selecting this option will cause each emulated RBridge to simulate topology that are interconnected with each other. Each emulated RBridge will simulate links connecting to every RBridge on every port.

Figure 295. Sharing Topology between each RBridge on every core port

19. Select the **Advertise Campus VLANs** checkbox. Increase the **No of Campus VLANs per Node** value to **2**. Enter 301 for **Start Vlan Id.** Select the **Advertise Unicast MACs** checkbox. Increase the **No of Unicast MACs per VLAN** value to **2**. Change an octet in the **Start Unicast MAC Address** to make it a unique start of range value.

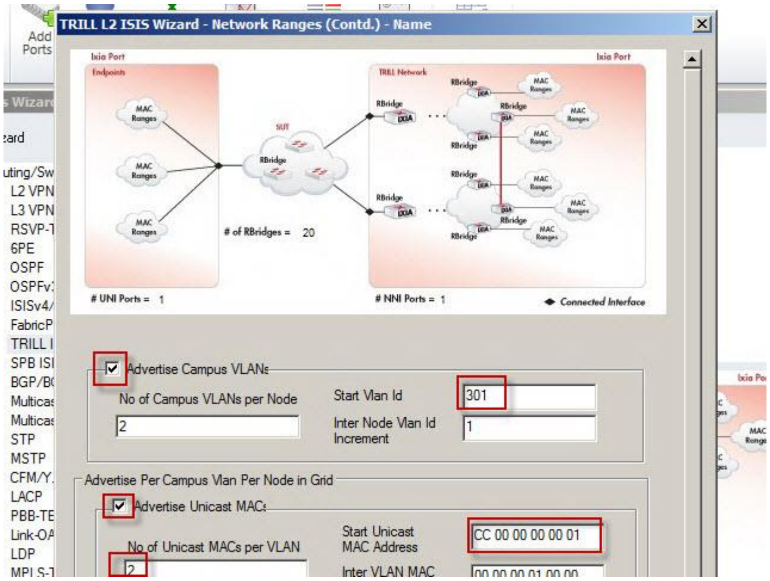

☑ Advertise Campus VLANs

No of Campus VLANs per Node: 2

Start Vlan Id: 301

Inter Node Vlan Id Increment: 1

Advertise Per Campus Vlan Per Node in Grid

☑ Advertise Unicast MACs

No of Unicast MACs per VLAN: 2

Start Unicast MAC Address: CC 00 00 00 00 01

Inter VLAN MAC: 00 00 00 01 00 00

Figure 296. Advertising Campus VLAN Unicast MAC ranges

20. Select the **Advertise Multicast MACs** checkbox. Increase the **No of Multicast MACs per VLAN** value to **2**. Increase the **No of Unicast Sources per Multicast MAC** value to **2**. Select the **Advertise IPv4 Groups** checkbox. Increase the **No of IPv4 Groups per VLAN** value to **2**. Increase the **No of Sources per IPv4 Group Address** value to 2. Select the **Advertise IPv6 Groups** checkbox. Increase the **No of IPv6 Groups per VLAN** value to **2**. Increase the **No of Sources per IPv6 Group Address** value to **2**. Click Next.

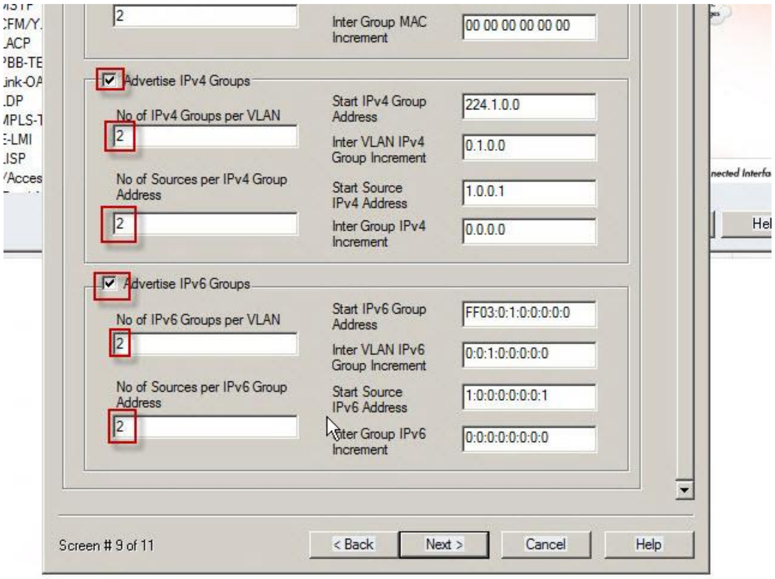

2

Inter Group MAC Increment: 00 00 00 00 00 00

☑ Advertise IPv4 Groups

No of IPv4 Groups per VLAN: 2

No of Sources per IPv4 Group Address: 2

Start IPv4 Group Address: 224.1.0.0

Inter VLAN IPv4 Group Increment: 0.1.0.0

Start Source IPv4 Address: 1.0.0.1

Inter Group IPv4 Increment: 0.0.0.0

☑ Advertise IPv6 Groups

No of IPv6 Groups per VLAN: 2

No of Sources per IPv6 Group Address: 2

Start IPv6 Group Address: FF03:0:1:0:0:0:0:0

Inter VLAN IPv6 Group Increment: 0:0:1:0:0:0:0:0

Start Source IPv6 Address: 1:0:0:0:0:0:0:1

Inter Group IPv6 Increment: 0:0:0:0:0:0:0:0

Screen # 9 of 11

Figure 297. Advertising Campus VLAN Multicast MAC, IPv4, IPv6 ranges

21. Select the **Use Vlans configured on PE Side** checkbox to mirror the VLANs emulated on the Core Side automatically over to the Customer Side.

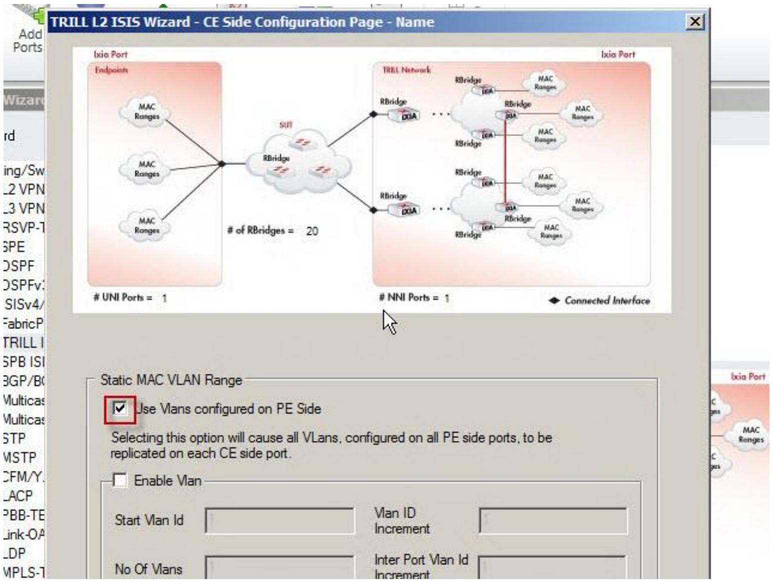

Static MAC VLAN Range

☑ Use Vlans configured on PE Side

Selecting this option will cause all VLans, configured on all PE side ports, to be replicated on each CE side port.

□ Enable Vlan

Start Vlan Id | Vlan ID Increment

No Of Vlans | Inter Port Vlan Id Increment

Figure 298. Mirroring Core Side VLANs on the Customer Side

22. Change an Octet in the **Start MAC Address** to make it a unique start of range value. Increase the **No of MACs Per Vlan** value to **2**. Click Next.

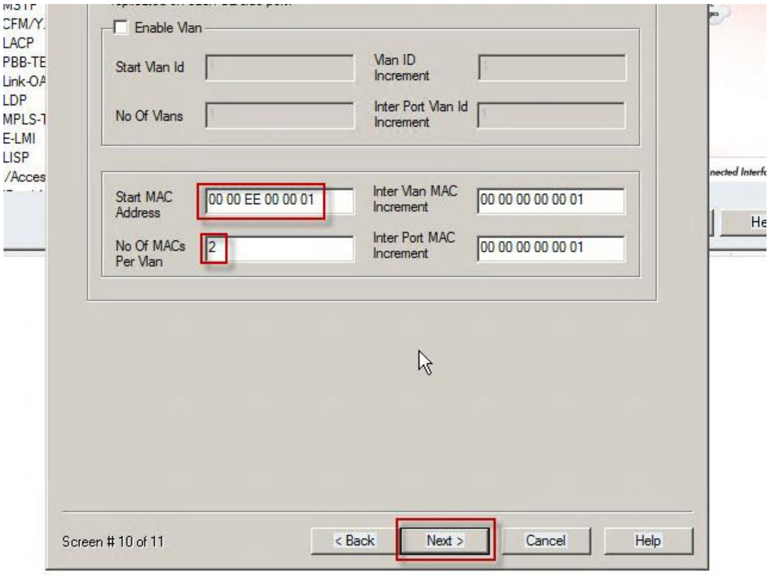

□ Enable Vlan

Start Vlan Id | Vlan ID Increment

No Of Vlans | Inter Port Vlan Id Increment

Start MAC Address: 00 00 EE 00 00 01

Inter Vlan MAC Increment: 00 00 00 00 00 01

No Of MACs Per Vlan: 2

Inter Port MAC Increment: 00 00 00 00 00 01

Screen # 10 of 11

Figure 299. Configuring static MAC ranges on the Customer Side

23. Enter a unique **Name** value for the configuration wizard run to save it for future use. Click **Generate and Overwrite Existing Configuration** to ensure new information propogates to the port.

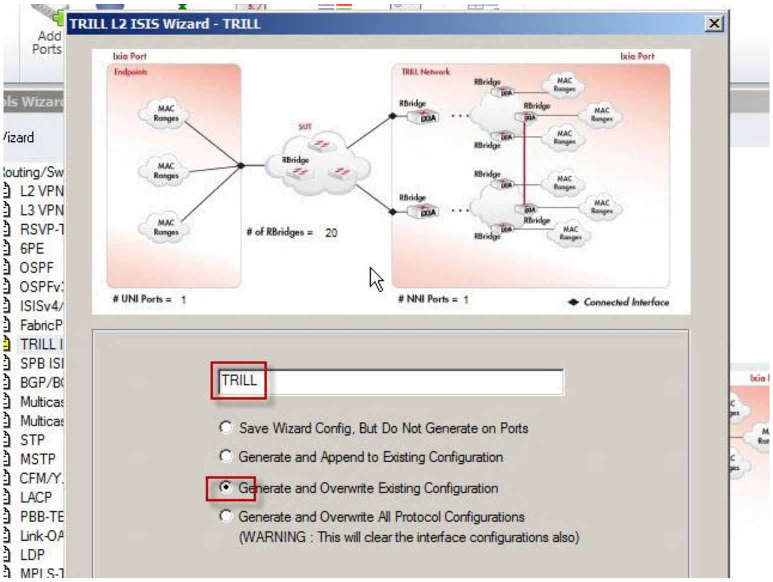

TRILL

○ Save Wizard Config, But Do Not Generate on Ports
○ Generate and Append to Existing Configuration
● Generate and Overwrite Existing Configuration
○ Generate and Overwrite All Protocol Configurations (WARNING : This will clear the interface configurations also)

Figure 300. Saving Current, Generating and Overwriting Existing Configuration

24. Click **Finish** to complete the wizard configuration. Click **Close** to leave the **Protocol Wizards** selector and return to the main window.

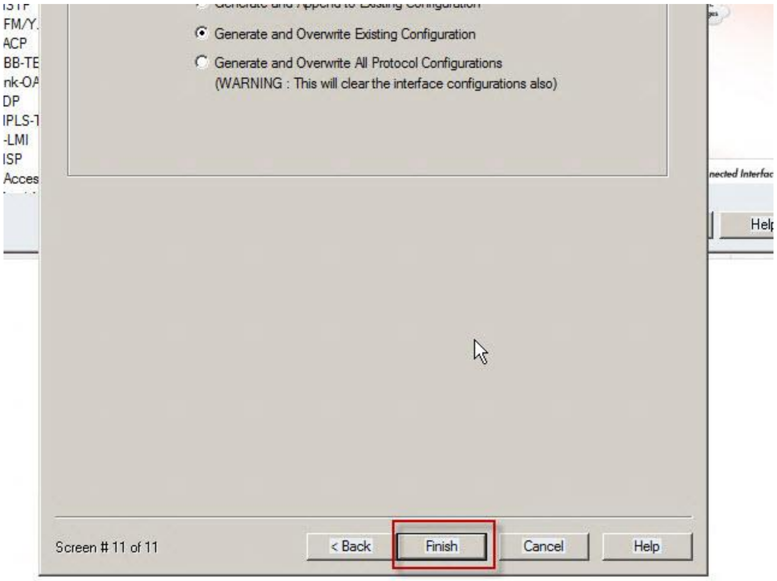

● Generate and Overwrite Existing Configuration
○ Generate and Overwrite All Protocol Configurations (WARNING : This will clear the interface configurations also)

Screen # 11 of 11

Figure 301. Finishing the Protocol Wizard configuration

25. Click Protocols > ISIS L2/L3 in the left pane to view details about the protocol interfaces and emulation ranges associated with TRILL on the core port. Select **TRILL** in the **Show tabs relevant to** dropdown box. This action reduces the set of tabs displayed for the ISIS protocol to those related to TRILL as opposed to other L2MP technologies.

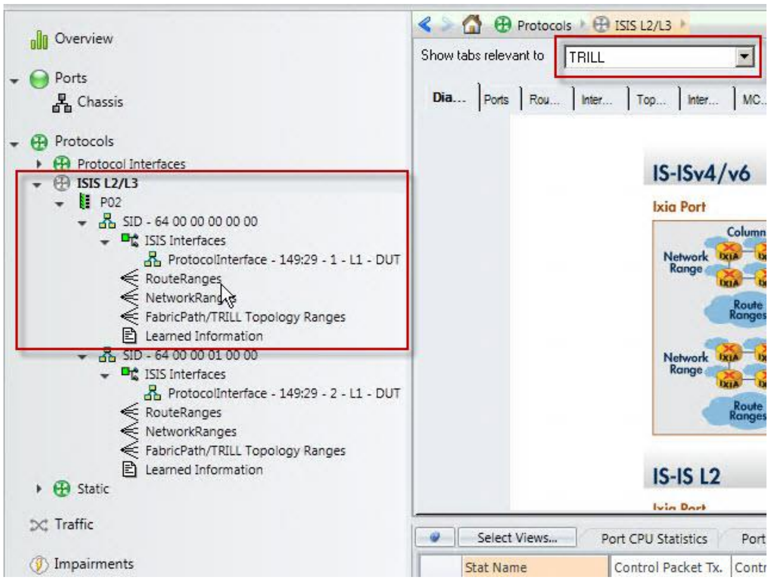

Show tabs relevant to: TRILL

Figure 302. Expanding the ISIS L2/L3 protocol tree

26. Click the **Ports** tab to review the number of RBridges and Emulation type details generated by the protocol wizard.

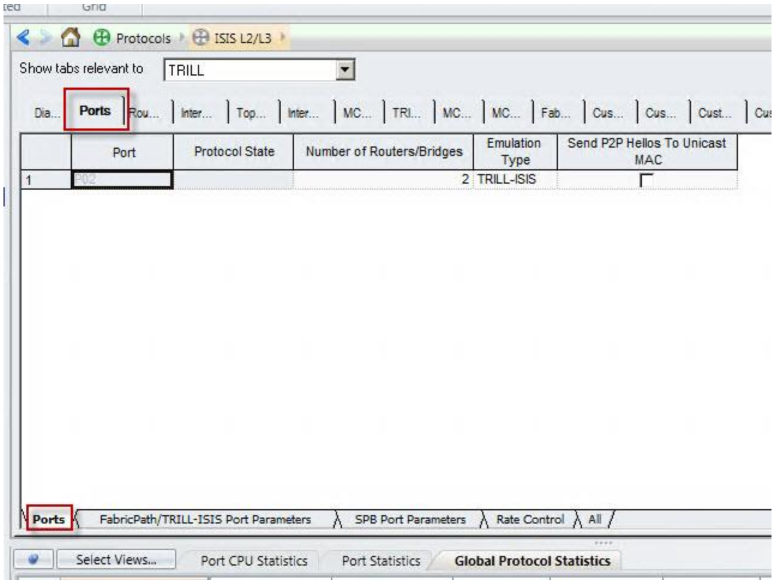

Show tabs relevant to: TRILL

	Port	Protocol State	Number of Routers/Bridges	Emulation Type	Send P2P Hellos To Unicast MAC
1	P02		2	TRILL-ISIS	☐

Figure 303. Exploring the ISIS L2/3 ports tab details

27. Click the **Router/Bridge** tab to review the system ID and hostname details generated by the protocol wizard.

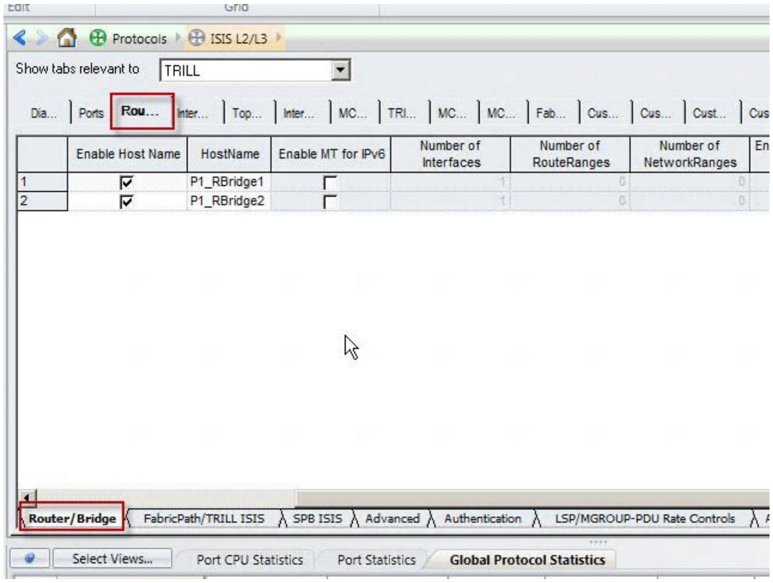

Show tabs relevant to TRILL

	Enable Host Name	HostName	Enable MT for IPv6	Number of Interfaces	Number of RouteRanges	Number of NetworkRanges	En
1	☑	P1_RBridge1	□	1	0	0	
2	☑	P1_RBridge2	□	1	0	0	

Figure 304. Exploring the ISIS L2/3 Router/Bridge tab details

28. Scroll to the far right of the **Router/Bridge** tab and clear the **Discard LSPs** checkboxes for your emulated RBridges. Clear this option allows you to use the Learned Information feature to review the topology information shared by the DUT.

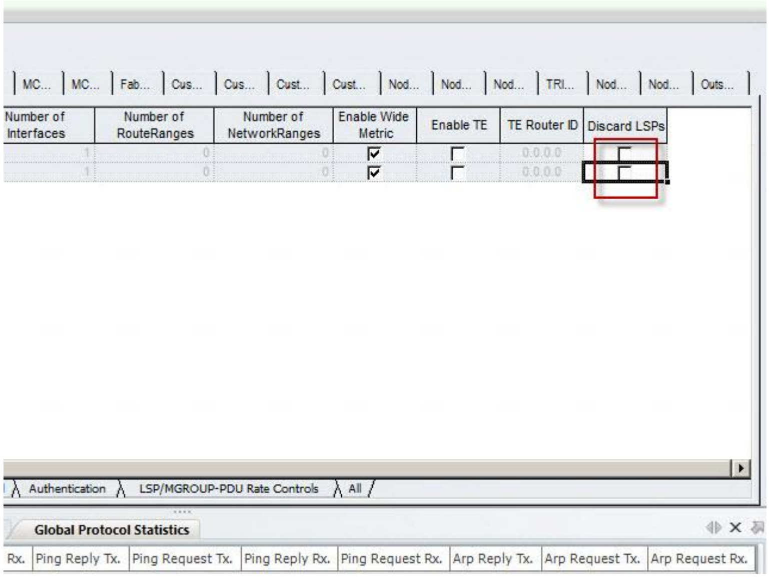

Number of Interfaces	Number of RouteRanges	Number of NetworkRanges	Enable Wide Metric	Enable TE	TE Router ID	Discard LSPs
1	0	0	☑	□	0.0.0.0	□
1	0	0	☑	□	0.0.0.0	□

Figure 305. Unchecking the Discard LSPs option in the Router/Bridge tab

29. Click the **Interfaces** tab to review the metrics, timers, and link type details generated by the protocol wizard.

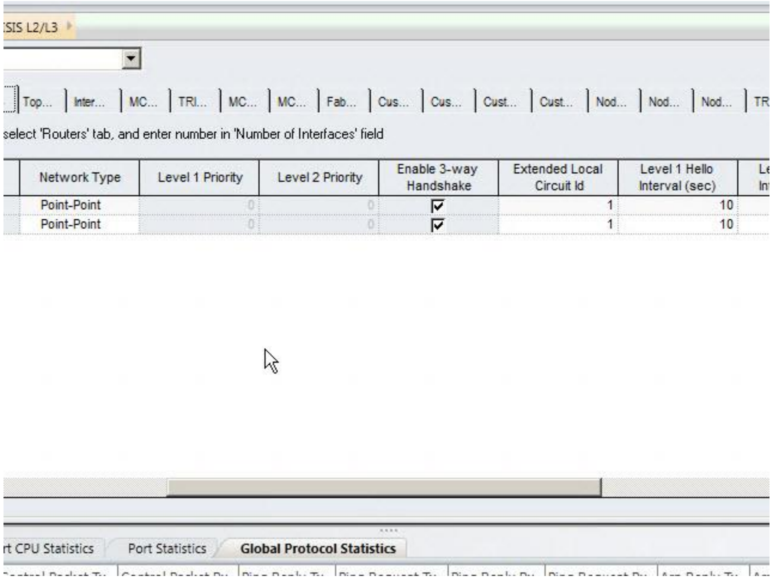

select 'Routers' tab, and enter number in 'Number of Interfaces' field

Network Type	Level 1 Priority	Level 2 Priority	Enable 3-way Handshake	Extended Local Circuit Id	Level 1 Hello Interval (sec)	Le In
Point-Point	0	0	☑	1	10	
Point-Point	0	0	☑	1	10	

Figure 306. Exploring the ISIS L2/L3 Interfaces tab details

30. Click the **FabricPath/TRILL Topology Ranges** tab to review the Interested VLAN range counts and tree calculation details generated by the protocol wizard.

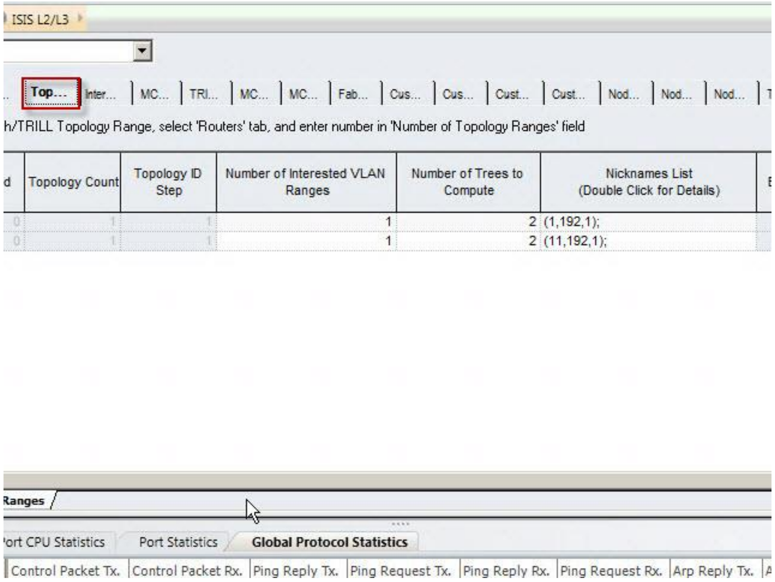

h/TRILL Topology Range, select 'Routers' tab, and enter number in 'Number of Topology Ranges' field

d	Topology Count	Topology ID Step	Number of Interested VLAN Ranges	Number of Trees to Compute	Nicknames List (Double Click for Details)	E
0	1	1	1	2	(1,192,1);	
0	1	1	1	2	(11,192,1);	

Figure 307. Exploring the ISIS L2/3 FabricPath/TRILL Topology Ranges tab details

31. Click the **FabricPath/TRILL Interested VLAN Ranges** tab to review campus VLAN details generated by the Protocol Wizard.

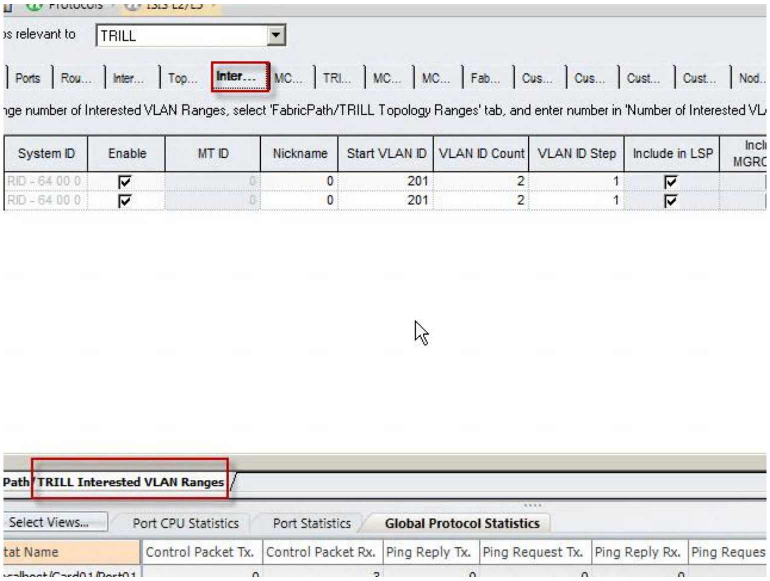

ıs relevant to: TRILL

ıge number of Interested VLAN Ranges, select 'FabricPath/TRILL Topology Ranges' tab, and enter number in 'Number of Interested VL

System ID	Enable	MT ID	Nickname	Start VLAN ID	VLAN ID Count	VLAN ID Step	Include in LSP	Incl MGRO
RID - 64 00 0	☑	0	0	201	2	1	☑	
RID - 64 00 0	☑	0	0	201	2	1	☑	

tat Name	Control Packet Tx.	Control Packet Rx.	Ping Reply Tx.	Ping Request Tx.	Ping Reply Rx.	Ping Reques

Figure 308. Exploring the ISIS L2/3 FabricPath/TRILL Interested VLAN Ranges tab details

32. Click the **FabricPath/TRILL Multicast MAC Ranges** tab to review the multicast destinations and unicast sources details generated by the Protocol Wizard.

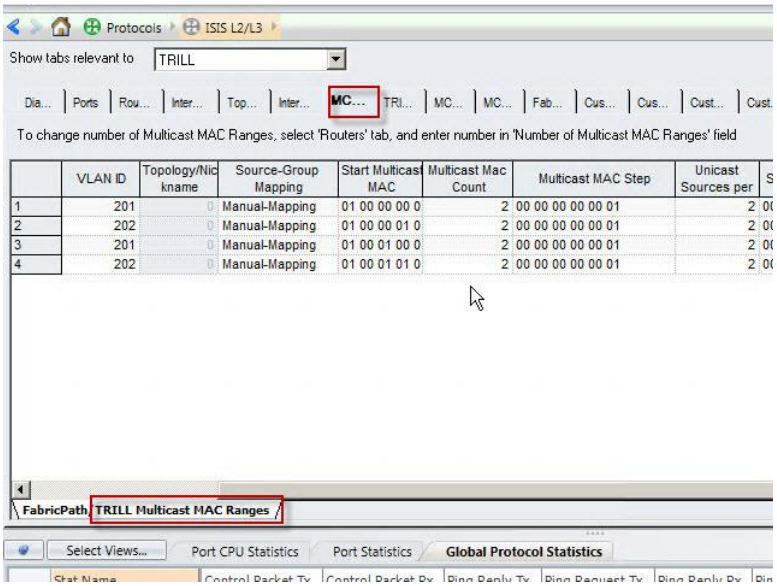

Show tabs relevant to: TRILL

To change number of Multicast MAC Ranges, select 'Routers' tab, and enter number in 'Number of Multicast MAC Ranges' field

	VLAN ID	Topology/Nickname	Source-Group Mapping	Start Multicast MAC	Multicast Mac Count	Multicast MAC Step	Unicast Sources per	S
1	201	0	Manual-Mapping	01 00 00 00 0	2	00 00 00 00 00 01	2	00
2	202	0	Manual-Mapping	01 00 00 01 0	2	00 00 00 00 00 01	2	00
3	201	0	Manual-Mapping	01 00 01 00 0	2	00 00 00 00 00 01	2	00
4	202	0	Manual-Mapping	01 00 01 01 0	2	00 00 00 00 00 01	2	00

Stat Name	Control Packet Tx	Control Packet Rx	Ping Reply Tx	Ping Request Tx	Ping Reply Rx	Pin

Figure 309. Exploring the ISIS L2/3 FabricPath/TRILL Multicast Ranges tab details

33. Click the **TRILL Unicast MAC Ranges** tab to review the unicast destination details generated by the Protocol Wizard.

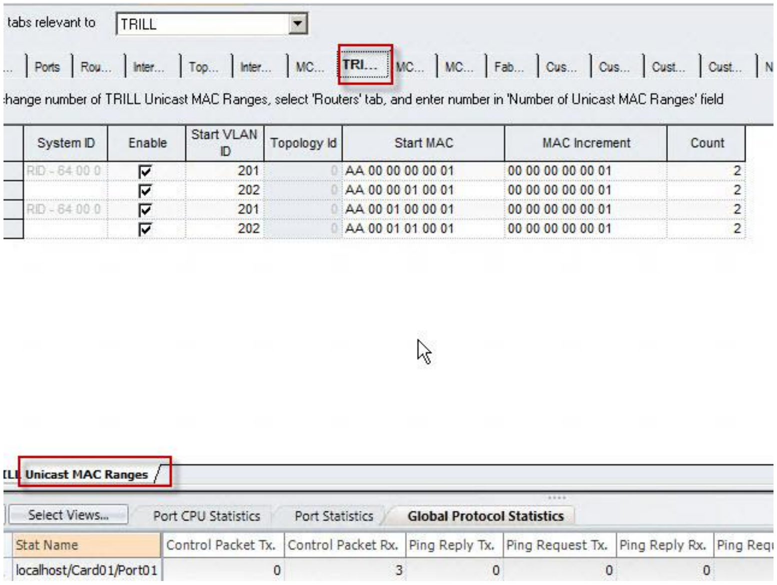

hange number of TRILL Unicast MAC Ranges, select 'Routers' tab, and enter number in 'Number of Unicast MAC Ranges' field

System ID	Enable	Start VLAN ID	Topology Id	Start MAC	MAC Increment	Count
RID - 64 00 0	☑	201	0	AA 00 00 00 00 01	00 00 00 00 00 01	2
	☑	202	0	AA 00 00 01 00 01	00 00 00 00 00 01	2
RID - 64 00 0	☑	201	0	AA 00 01 00 00 01	00 00 00 00 00 01	2
	☑	202	0	AA 00 01 01 00 01	00 00 00 00 00 01	2

Stat Name	Control Packet Tx.	Control Packet Rx.	Ping Reply Tx.	Ping Request Tx.	Ping Reply Rx.	Ping Requ
localhost/Card01/Port01	0	3	0	0	0	

Figure 310. Exploring the TRILL Unicast MAC Ranges tab details

34. Click the **FabricPath/TRILL Multicast IPv4 Group Ranges** tab to review multicast destinations and unicast sources details generated by the Protocol Wizard.

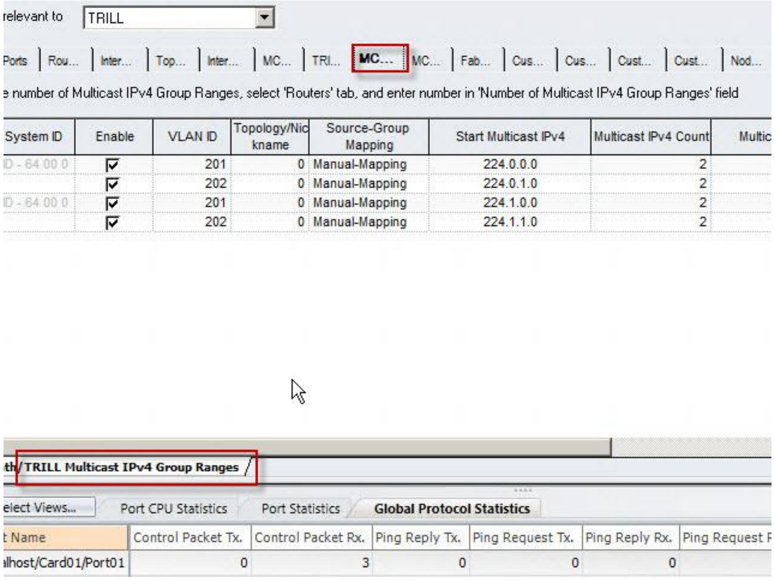

e number of Multicast IPv4 Group Ranges, select 'Routers' tab, and enter number in 'Number of Multicast IPv4 Group Ranges' field

System ID	Enable	VLAN ID	Topology/Nickname	Source-Group Mapping	Start Multicast IPv4	Multicast IPv4 Count	Multic
D - 64 00 0	☑	201	0	Manual-Mapping	224.0.0.0	2	
	☑	202	0	Manual-Mapping	224.0.1.0	2	
D - 64 00 0	☑	201	0	Manual-Mapping	224.1.0.0	2	
	☑	202	0	Manual-Mapping	224.1.1.0	2	

t Name	Control Packet Tx.	Control Packet Rx.	Ping Reply Tx.	Ping Request Tx.	Ping Reply Rx.	Ping Request F
alhost/Card01/Port01	0	3	0	0	0	

Figure 311. Exploring the FabricPath/TRILL Multicast IPv4 Group Ranges tab details

35. Click the **FabricPath/TRILL Multicast IPv6 Group Ranges** tab to review multicast destinations and unicast sources details generate by the Protocol Wizard.

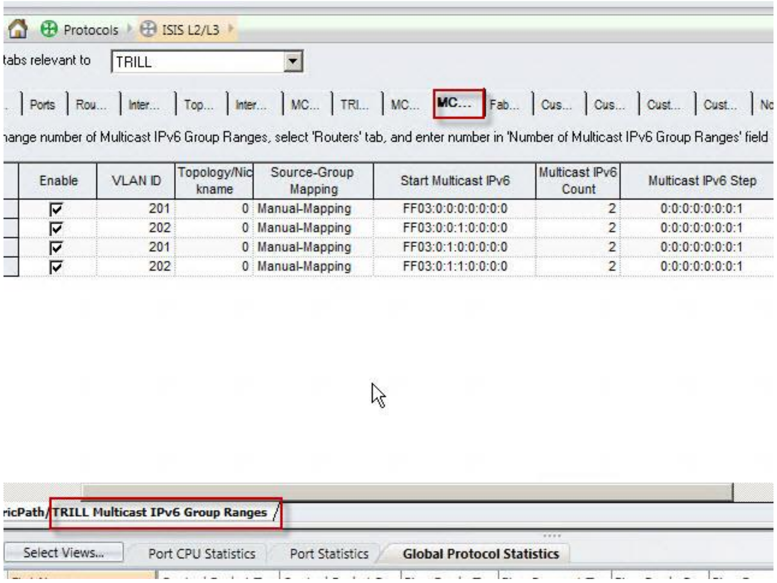

tabs relevant to TRILL

hange number of Multicast IPv6 Group Ranges, select 'Routers' tab, and enter number in 'Number of Multicast IPv6 Group Ranges' field

	Enable	VLAN ID	Topology/Nickname	Source-Group Mapping	Start Multicast IPv6	Multicast IPv6 Count	Multicast IPv6 Step
	☑	201	0	Manual-Mapping	FF03:0:0:0:0:0:0:0	2	0:0:0:0:0:0:0:1
	☑	202	0	Manual-Mapping	FF03:0:0:1:0:0:0:0	2	0:0:0:0:0:0:0:1
	☑	201	0	Manual-Mapping	FF03:0:1:0:0:0:0:0	2	0:0:0:0:0:0:0:1
	☑	202	0	Manual-Mapping	FF03:0:1:1:0:0:0:0	2	0:0:0:0:0:0:0:1

Figure 312. Exploring the FabricPath/TRILL Multicast IPv6 Group Ranges tab details

36. Click the **FabricPath/TRILL L2 Network Ranges** tab to review the Rows x Columns Grid of RBridges topology generated by the Protocol Wizard.

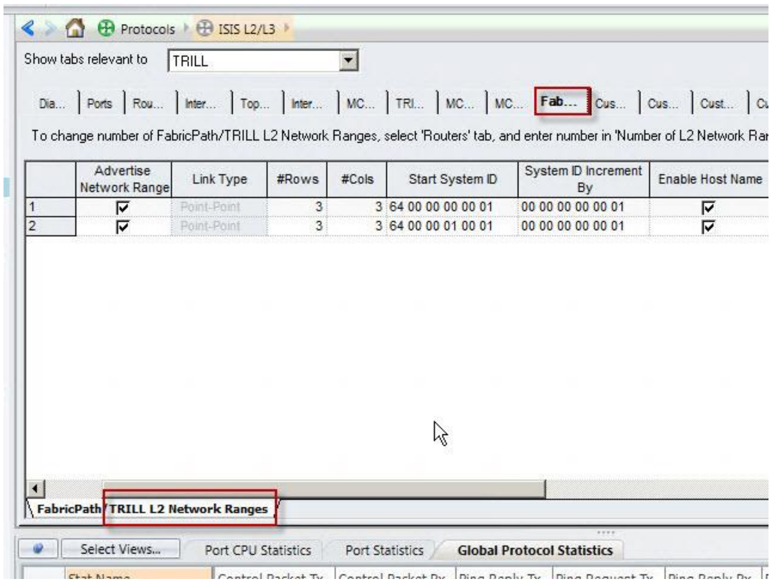

Show tabs relevant to TRILL

To change number of FabricPath/TRILL L2 Network Ranges, select 'Routers' tab, and enter number in 'Number of L2 Network Rar

	Advertise Network Range	Link Type	#Rows	#Cols	Start System ID	System ID Increment By	Enable Host Name
1	☑	Point-Point	3	3	64 00 00 00 00 01	00 00 00 00 00 01	☑
2	☑	Point-Point	3	3	64 00 00 01 00 01	00 00 00 00 00 01	☑

Figure 313. Exploring the FabricPath/TRILL L2 Network Ranges tab details

37. Click the **FabricPath/TRILL Node Topology Ranges** tab to review the number of Interested VLAN ranges and trees to compute details generated by the Protocol Wizard for use by RBridges in the rows x columns grid topology.

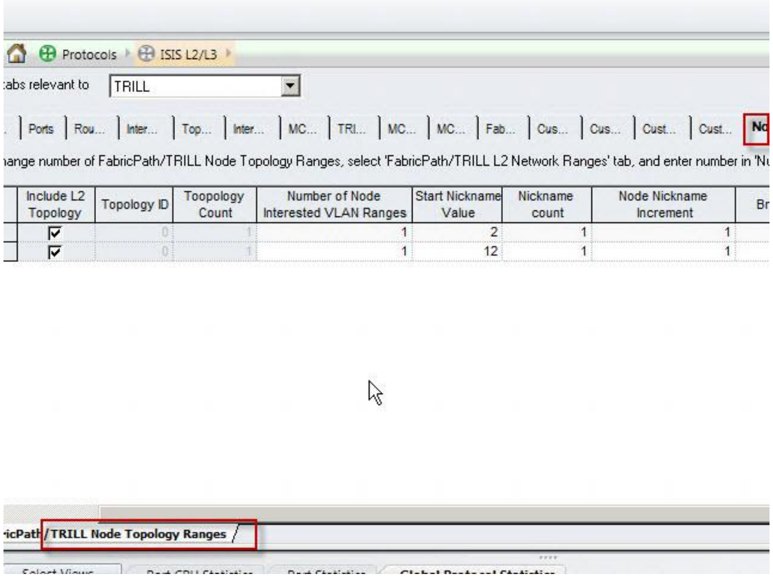

abs relevant to TRILL

ange number of FabricPath/TRILL Node Topology Ranges, select 'FabricPath/TRILL L2 Network Ranges' tab, and enter number in 'Nu

Include L2 Topology	Topology ID	Toopology Count	Number of Node Interested VLAN Ranges	Start Nickname Value	Nickname count	Node Nickname Increment	Br
☑	0	1	1	2	1	1	
☑	0	1	1	12	1	1	

Figure 314. Exploring the FabricPath/TRILL Node Topology Ranges tab details

38. Click the **FabricPath/TRILL Node Interested VLAN Ranges** tab to review the campus VLANs details generated by the Protocol Wizard for use by RBridges in the rows x columns grid topology.

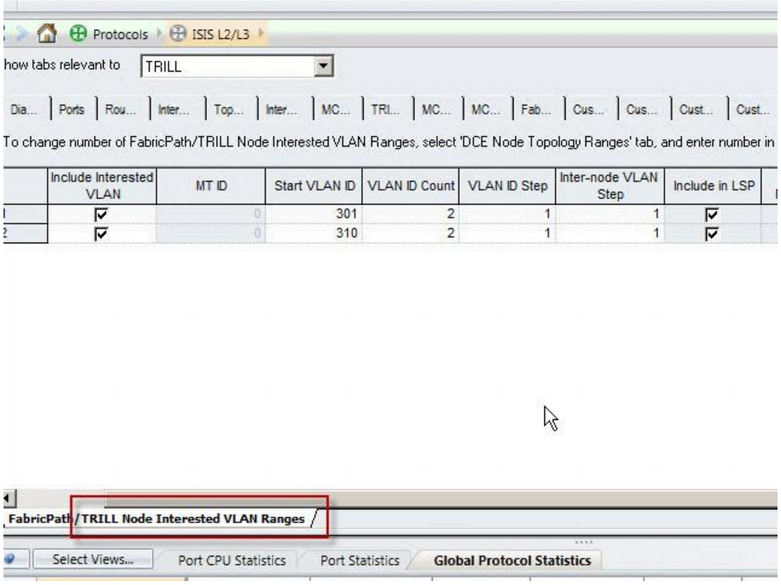

how tabs relevant to TRILL

To change number of FabricPath/TRILL Node Interested VLAN Ranges, select 'DCE Node Topology Ranges' tab, and enter number in

	Include Interested VLAN	MT ID	Start VLAN ID	VLAN ID Count	VLAN ID Step	Inter-node VLAN Step	Include in LSP
1	☑	0	301	2	1	1	☑
2	☑	0	310	2	1	1	☑

Figure 315. Exploring the FabricPath/TRILL Node Interested VLAN Ranges tab details

39. Click the **FabricPath/TRILL Node MAC Groups** tab to review the multicast destinations and unicast sources details generated by the Protocol Wizard for use by RBridges in the rows x columns grid topology.

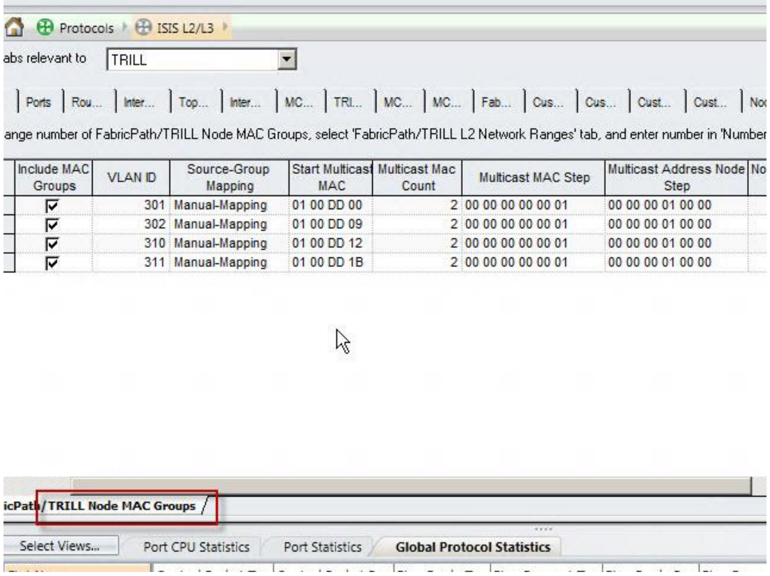

ange number of FabricPath/TRILL Node MAC Groups, select 'FabricPath/TRILL L2 Network Ranges' tab, and enter number in 'Number

Include MAC Groups	VLAN ID	Source-Group Mapping	Start Multicast MAC	Multicast Mac Count	Multicast MAC Step	Multicast Address Node Step	No
☑	301	Manual-Mapping	01 00 DD 00	2	00 00 00 00 00 01	00 00 00 01 00 00	
☑	302	Manual-Mapping	01 00 DD 09	2	00 00 00 00 00 01	00 00 00 01 00 00	
☑	310	Manual-Mapping	01 00 DD 12	2	00 00 00 00 00 01	00 00 00 01 00 00	
☑	311	Manual-Mapping	01 00 DD 1B	2	00 00 00 00 00 01	00 00 00 01 00 00	

Figure 316. Exploring the FabricPath/TRILL Node MAC Groups tab details

40. Click the **TRILL Node MAC Ranges** tab to review the unicast destinations details generated by the Protocol Wizard for use by RBridges in the rows x columns grid topology.

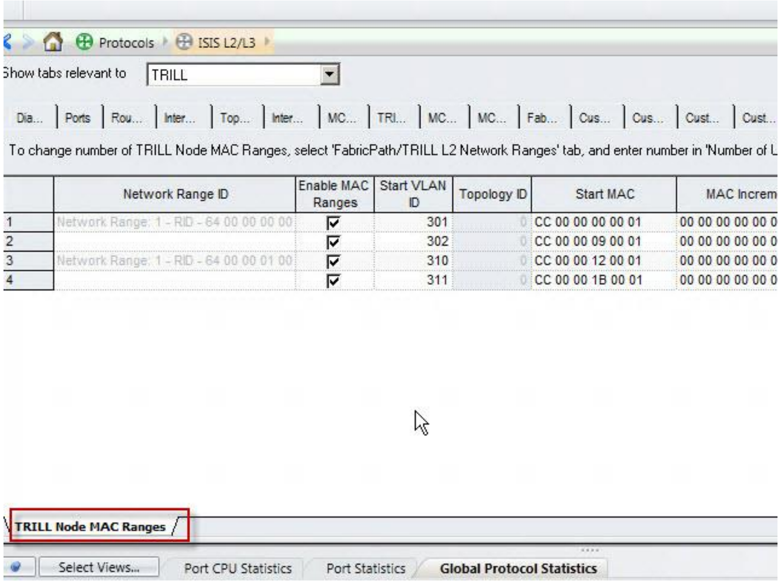

To change number of TRILL Node MAC Ranges, select 'FabricPath/TRILL L2 Network Ranges' tab, and enter number in 'Number of L

	Network Range ID	Enable MAC Ranges	Start VLAN ID	Topology ID	Start MAC	MAC Increm
1	Network Range: 1 - RID - 64 00 00 00 00	☑	301	0	CC 00 00 00 00 01	00 00 00 00 00 0
2		☑	302	0	CC 00 00 09 00 01	00 00 00 00 00 0
3	Network Range: 1 - RID - 64 00 00 01 00	☑	310	0	CC 00 00 12 00 01	00 00 00 00 00 0
4		☑	311	0	CC 00 00 1B 00 01	00 00 00 00 00 0

Figure 317. Exploring the TRILL Node MAC Ranges tab details

41. Click the **FabricPath/TRILL Node IPv4 Groups** tab to review the multicast destinations and unicast sources details generated by the Protocol Wizard for use by RBridges in the rows x columns grid topology.

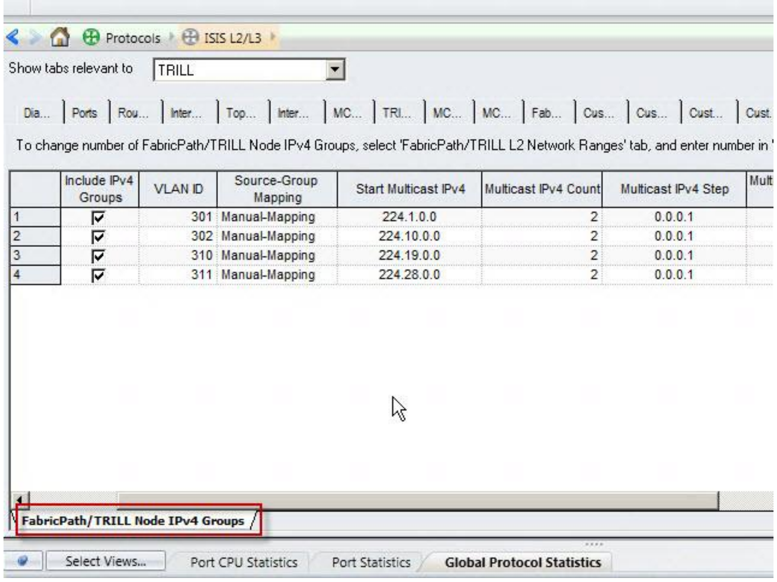

Show tabs relevant to TRILL

To change number of FabricPath/TRILL Node IPv4 Groups, select 'FabricPath/TRILL L2 Network Ranges' tab, and enter number in '

	Include IPv4 Groups	VLAN ID	Source-Group Mapping	Start Multicast IPv4	Multicast IPv4 Count	Multicast IPv4 Step	Mult
1	☑	301	Manual-Mapping	224.1.0.0	2	0.0.0.1	
2	☑	302	Manual-Mapping	224.10.0.0	2	0.0.0.1	
3	☑	310	Manual-Mapping	224.19.0.0	2	0.0.0.1	
4	☑	311	Manual-Mapping	224.28.0.0	2	0.0.0.1	

Figure 318. Exploring the FabricPath/TRILL Node IPv4 Groups tab details

42. Click the **FabricPath/TRILL Node IPv6 Groups** tab to review the multicast destinations and unicast sources details generated by the Protocol Wizard for use by Rbridges in the rows x columns grid topology.

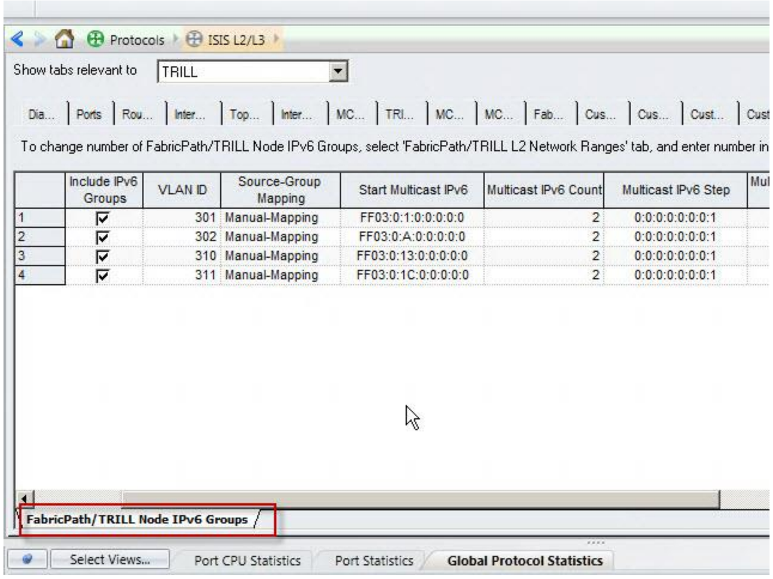

Show tabs relevant to TRILL

To change number of FabricPath/TRILL Node IPv6 Groups, select 'FabricPath/TRILL L2 Network Ranges' tab, and enter number in

	Include IPv6 Groups	VLAN ID	Source-Group Mapping	Start Multicast IPv6	Multicast IPv6 Count	Multicast IPv6 Step	Mul
1	☑	301	Manual-Mapping	FF03:0:1:0:0:0:0:0	2	0:0:0:0:0:0:0:1	
2	☑	302	Manual-Mapping	FF03:0:A:0:0:0:0:0	2	0:0:0:0:0:0:0:1	
3	☑	310	Manual-Mapping	FF03:0:13:0:0:0:0:0	2	0:0:0:0:0:0:0:1	
4	☑	311	Manual-Mapping	FF03:0:1C:0:0:0:0:0	2	0:0:0:0:0:0:0:1	

Figure 319. Exploring the FabricPath/TRILL Node IPv6 Groups tab details

43. Click **Static** > **LANs** in the left pane.

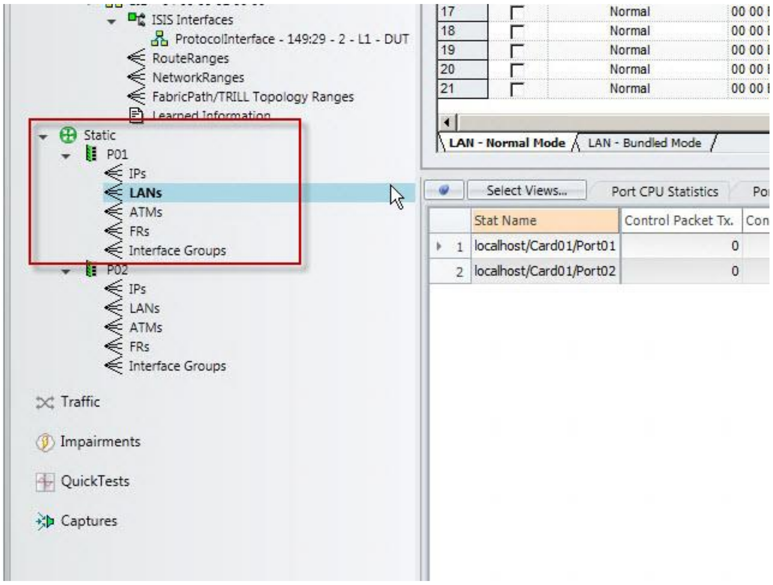

17	☐	Normal	00 00 I
18	☐	Normal	00 00 I
19	☐	Normal	00 00 I
20	☐	Normal	00 00 I
21	☐	Normal	00 00 I

	Stat Name	Control Packet Tx.	Con
1	localhost/Card01/Port01	0	
2	localhost/Card01/Port02	0	

Figure 320. Expanding the Static Protocols tree

44. Click the **LAN – Normal Mode** tab to review the VLAN configuration mirrored to the Customer Side port and the associated MAC sources generated by the Protocol Wizard.

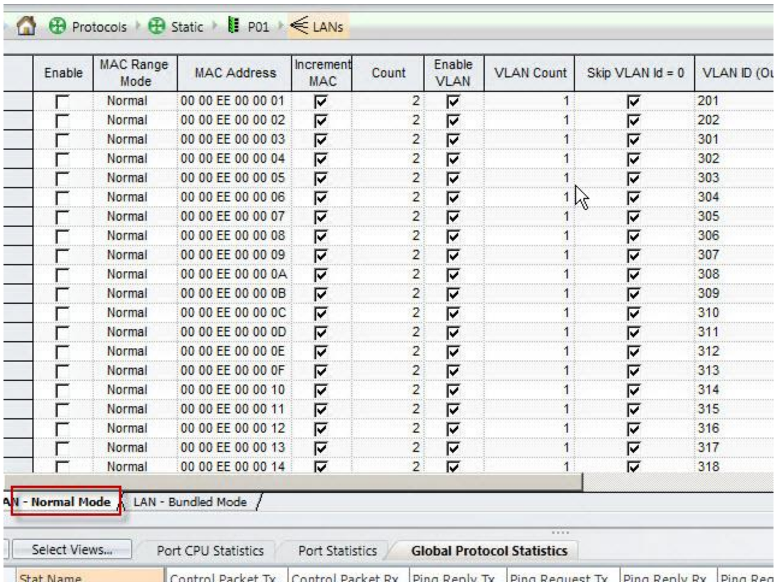

Enable	MAC Range Mode	MAC Address	Increment MAC	Count	Enable VLAN	VLAN Count	Skip VLAN Id = 0	VLAN ID (Ou
☐	Normal	00 00 EE 00 00 01	☑	2	☑	1	☑	201
☐	Normal	00 00 EE 00 00 02	☑	2	☑	1	☑	202
☐	Normal	00 00 EE 00 00 03	☑	2	☑	1	☑	301
☐	Normal	00 00 EE 00 00 04	☑	2	☑	1	☑	302
☐	Normal	00 00 EE 00 00 05	☑	2	☑	1	☑	303
☐	Normal	00 00 EE 00 00 06	☑	2	☑	1	☑	304
☐	Normal	00 00 EE 00 00 07	☑	2	☑	1	☑	305
☐	Normal	00 00 EE 00 00 08	☑	2	☑	1	☑	306
☐	Normal	00 00 EE 00 00 09	☑	2	☑	1	☑	307
☐	Normal	00 00 EE 00 00 0A	☑	2	☑	1	☑	308
☐	Normal	00 00 EE 00 00 0B	☑	2	☑	1	☑	309
☐	Normal	00 00 EE 00 00 0C	☑	2	☑	1	☑	310
☐	Normal	00 00 EE 00 00 0D	☑	2	☑	1	☑	311
☐	Normal	00 00 EE 00 00 0E	☑	2	☑	1	☑	312
☐	Normal	00 00 EE 00 00 0F	☑	2	☑	1	☑	313
☐	Normal	00 00 EE 00 00 10	☑	2	☑	1	☑	314
☐	Normal	00 00 EE 00 00 11	☑	2	☑	1	☑	315
☐	Normal	00 00 EE 00 00 12	☑	2	☑	1	☑	316
☐	Normal	00 00 EE 00 00 13	☑	2	☑	1	☑	317
☐	Normal	00 00 EE 00 00 14	☑	2	☑	1	☑	318

Stat Name	Control Packet Tx	Control Packet Rx	Ping Reply Tx	Ping Request Tx	Ping Reply Rx	Ping Req

Figure 321. Exploring the LAN–Normal Mode tab details

45. Click **Start All Protocols** and wait for the ISIS L2/L3 protocol to come up green.

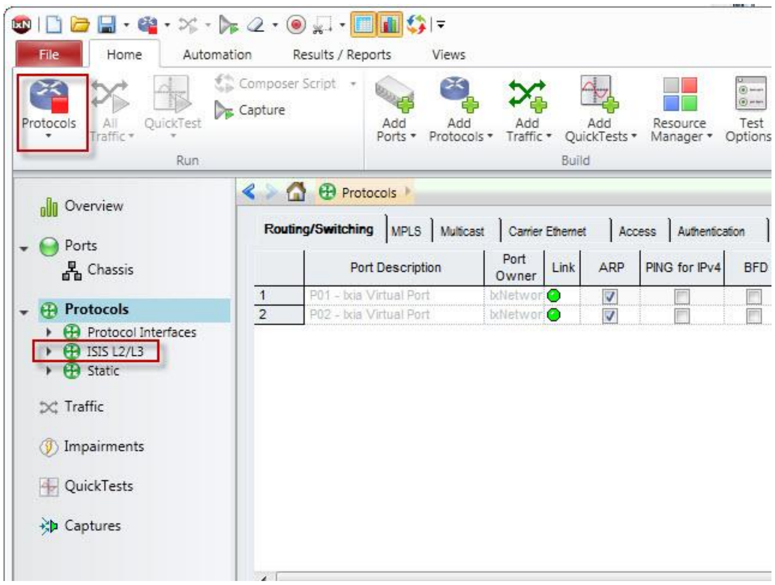

	Port Description	Port Owner	Link	ARP	PING for IPv4	BFD
1	P01 - Ixia Virtual Port	IxNetwor	●	☑	□	□
2	P02 - Ixia Virtual Port	IxNetwor	●	☑	□	□

Figure 322. Starting All Protocols

46. Click the **ISIS Aggregated Statistics** tab and scroll to the right to verify the **L1 Full State Count** value matches the **L1 Neighbors** value and both are non-zero.

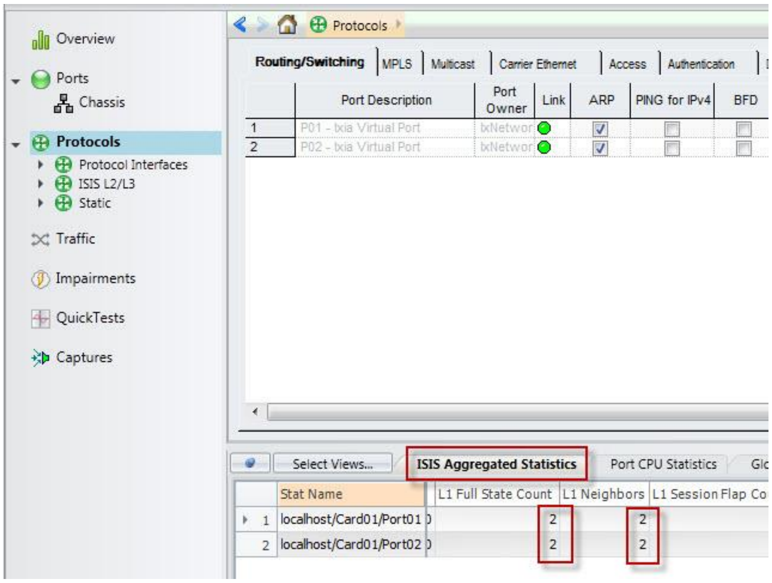

	Port Description	Port Owner	Link	ARP	PING for IPv4	BFD
1	P01 - Ixia Virtual Port	IxNetwor	●	☑	□	□
2	P02 - Ixia Virtual Port	IxNetwor	●	☑	□	□

	Stat Name	L1 Full State Count	L1 Neighbors	L1 Session Flap Co
1	localhost/Card01/Port01	2	2	
2	localhost/Card01/Port02	2	2	

Figure 323. Verifying ISIS Level 1 statistics

47. Click **Traffic** in the left plane and then click **Add L2-3 Traffic Items** to open the **Advance d Traffic Wizard**.

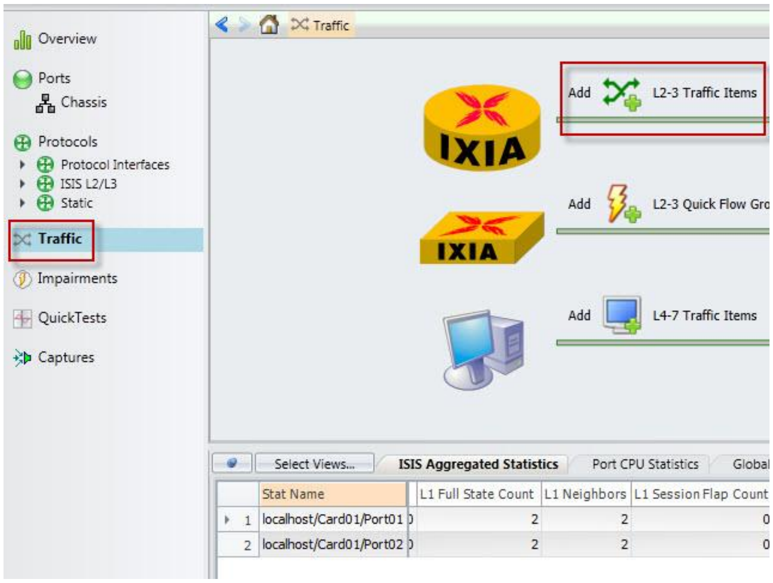

	Stat Name	L1 Full State Count	L1 Neighbors	L1 Session Flap Count
1	localhost/Card01/Port01	2	2	0
2	localhost/Card01/Port02	2	2	0

Figure 324. Selecting Add L2-3 Traffic Items from the Traffic view

48. In the **Type of Traffic** dropdown list, select **Ethernet/VLAN.** In the **Source** dropdown list, select **All**. Expand the first port and select the **Static > Ethernet** MAC sources corresponding to the Customer Side.

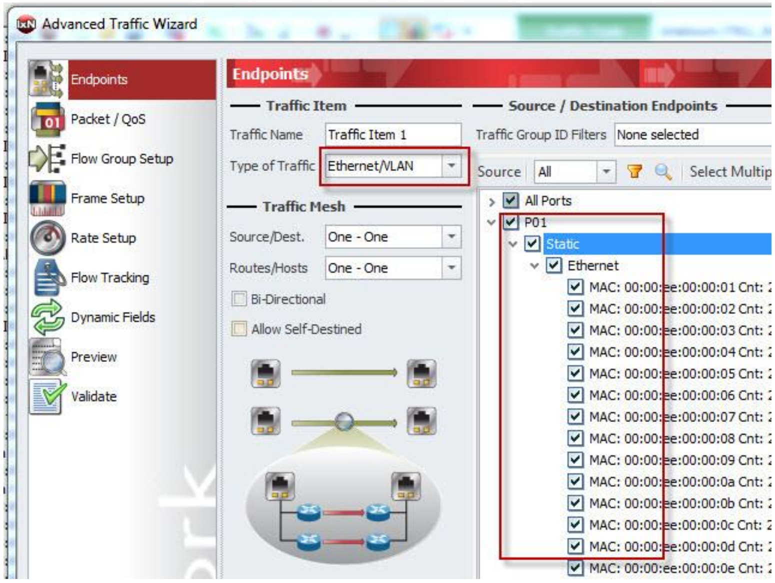

Traffic Name: Traffic Item 1

Type of Traffic: Ethernet/VLAN

Source/Dest.: One - One

Routes/Hosts: One - One

□ Bi-Directional

□ Allow Self-Destined

Traffic Group ID Filters: None selected

Source: All

- ☑ All Ports
  - ☑ P01
    - ☑ Static
      - ☑ Ethernet
        - ☑ MAC: 00:00:ee:00:00:01 Cnt:
        - ☑ MAC: 00:00:ee:00:00:02 Cnt:
        - ☑ MAC: 00:00:ee:00:00:03 Cnt:
        - ☑ MAC: 00:00:ee:00:00:04 Cnt:
        - ☑ MAC: 00:00:ee:00:00:05 Cnt:
        - ☑ MAC: 00:00:ee:00:00:06 Cnt:
        - ☑ MAC: 00:00:ee:00:00:07 Cnt:
        - ☑ MAC: 00:00:ee:00:00:08 Cnt:
        - ☑ MAC: 00:00:ee:00:00:09 Cnt:
        - ☑ MAC: 00:00:ee:00:00:0a Cnt:
        - ☑ MAC: 00:00:ee:00:00:0b Cnt:
        - ☑ MAC: 00:00:ee:00:00:0c Cnt:
        - ☑ MAC: 00:00:ee:00:00:0d Cnt:
        - ☑ MAC: 00:00:ee:00:00:0e Cnt:

Figure 325. Selecting Static Ethernet sources on Customer Side port

49. Select **TRILL** from the **Destination** dropdown box. Expand the second port and select **ISIS L2/3 > ISIS Unicast MAC Ranges** destinations corresponding to the Core Side.

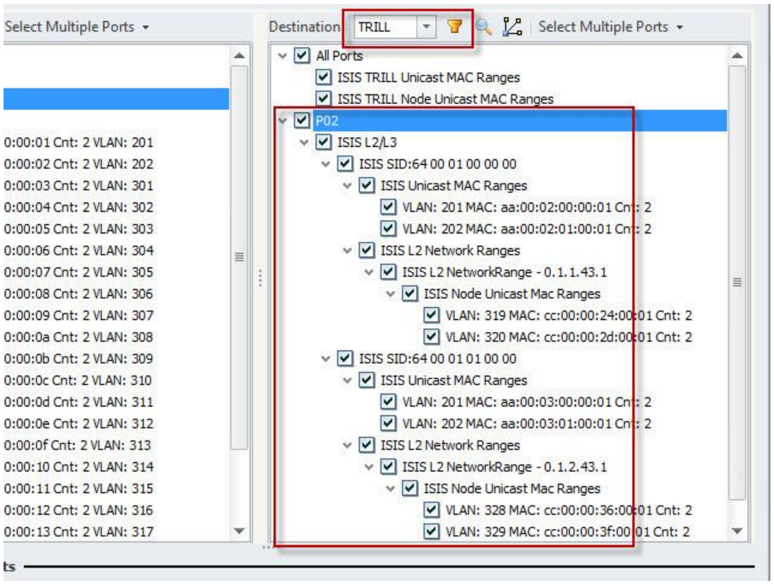

0:00:01 Cnt: 2 VLAN: 201
0:00:02 Cnt: 2 VLAN: 202
0:00:03 Cnt: 2 VLAN: 301
0:00:04 Cnt: 2 VLAN: 302
0:00:05 Cnt: 2 VLAN: 303
0:00:06 Cnt: 2 VLAN: 304
0:00:07 Cnt: 2 VLAN: 305
0:00:08 Cnt: 2 VLAN: 306
0:00:09 Cnt: 2 VLAN: 307
0:00:0a Cnt: 2 VLAN: 308
0:00:0b Cnt: 2 VLAN: 309
0:00:0c Cnt: 2 VLAN: 310
0:00:0d Cnt: 2 VLAN: 311
0:00:0e Cnt: 2 VLAN: 312
0:00:0f Cnt: 2 VLAN: 313
0:00:10 Cnt: 2 VLAN: 314
0:00:11 Cnt: 2 VLAN: 315
0:00:12 Cnt: 2 VLAN: 316
0:00:13 Cnt: 2 VLAN: 317

- ☑ All Ports
  - ☑ ISIS TRILL Unicast MAC Ranges
  - ☑ ISIS TRILL Node Unicast MAC Ranges
- ☑ P02
  - ☑ ISIS L2/L3
    - ☑ ISIS SID:64 00 01 00 00 00
      - ☑ ISIS Unicast MAC Ranges
        - ☑ VLAN: 201 MAC: aa:00:02:00:00:01 Cnt: 2
        - ☑ VLAN: 202 MAC: aa:00:02:01:00:01 Cnt: 2
      - ☑ ISIS L2 Network Ranges
        - ☑ ISIS L2 NetworkRange - 0.1.1.43.1
          - ☑ ISIS Node Unicast Mac Ranges
            - ☑ VLAN: 319 MAC: cc:00:00:24:00:01 Cnt: 2
            - ☑ VLAN: 320 MAC: cc:00:00:2d:00:01 Cnt: 2
    - ☑ ISIS SID:64 00 01 01 00 00
      - ☑ ISIS Unicast MAC Ranges
        - ☑ VLAN: 201 MAC: aa:00:03:00:00:01 Cnt: 2
        - ☑ VLAN: 202 MAC: aa:00:03:01:00:01 Cnt: 2
      - ☑ ISIS L2 Network Ranges
        - ☑ ISIS L2 NetworkRange - 0.1.2.43.1
          - ☑ ISIS Node Unicast Mac Ranges
            - ☑ VLAN: 328 MAC: cc:00:00:36:00:01 Cnt: 2
            - ☑ VLAN: 329 MAC: cc:00:00:3f:00:01 Cnt: 2

ts

Figure 326. Selecting ISIS Unicast MAC Ranges on Core Side Port

50. In the left pane, click **Frame Setup** and change the **Fixed** frame size value to **1500**.

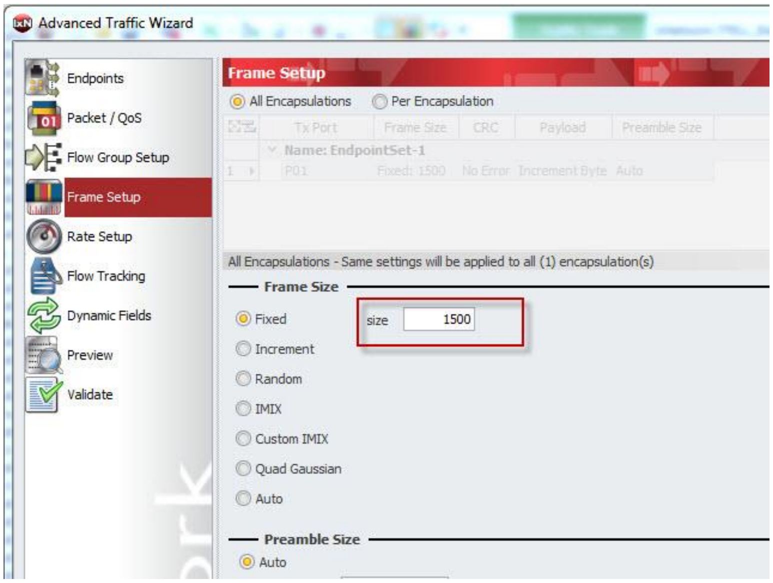

Frame Setup

● All Encapsulations ○ Per Encapsulation

	Tx Port	Frame Size	CRC	Payload	Preamble Size
Name: EndpointSet-1					
1	P01	Fixed: 1500	No Error	Increment Byte	Auto

All Encapsulations - Same settings will be applied to all (1) encapsulation(s)

Frame Size

● Fixed size 1500
○ Increment
○ Random
○ IMIX
○ Custom IMIX
○ Quad Gaussian
○ Auto

Preamble Size

● Auto

Figure 327. Changing the fixed frame size value

51. In the left pane, click **Rate Setup** and change the **Line rate** percentage value to 100.

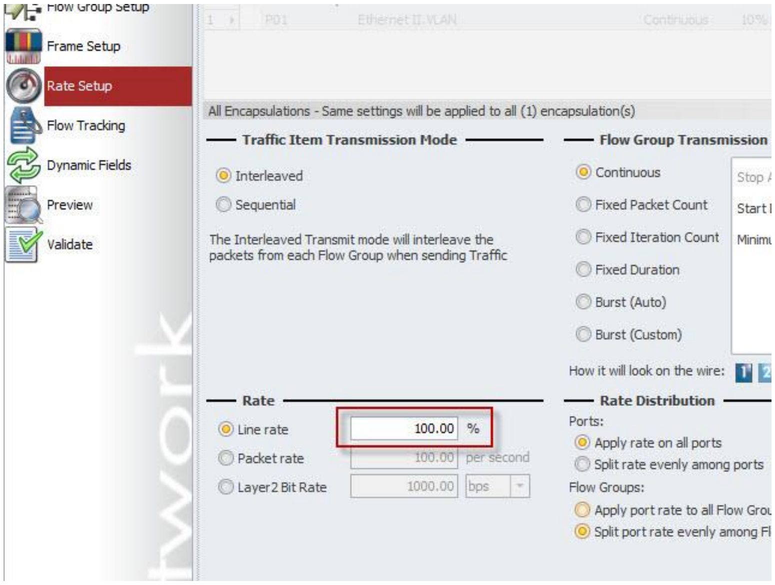

All Encapsulations - Same settings will be applied to all (1) encapsulation(s)

Traffic Item Transmission Mode

● Interleaved
○ Sequential

The Interleaved Transmit mode will interleave the packets from each Flow Group when sending Traffic

Flow Group Transmission

● Continuous
○ Fixed Packet Count
○ Fixed Iteration Count
○ Fixed Duration
○ Burst (Auto)
○ Burst (Custom)

How it will look on the wire:

Rate

● Line rate 100.00 %
○ Packet rate 100.00 per second
○ Layer2 Bit Rate 1000.00 bps

Rate Distribution

Ports:
● Apply rate on all ports
○ Split rate evenly among ports

Flow Groups:
○ Apply port rate to all Flow Grou
● Split port rate evenly among Fl

Figure 328. Changing the line rate throughput value

52. In the left pane, click **Flow Tracking** and select **Traffic Item** and **Source/Dest Endpoint Pair** tracking options. Select **VLAN: VLAN ID**, **TRILL: Egress RBridge Nickname, and TRILL: Ingress RBridge Nickname** tracking options. These options help you to drill down for more details dynamically on TRILL endpoint types. Click **Finish** to close the **Advanced Traffic Wizard**.

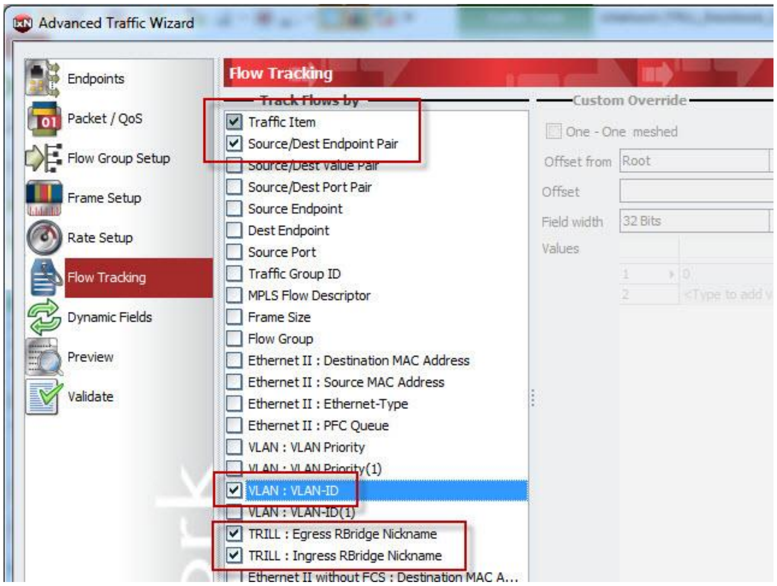

Track Flows by

☑ Traffic Item
☑ Source/Dest Endpoint Pair
☐ Source/Dest Value Pair
☐ Source/Dest Port Pair
☐ Source Endpoint
☐ Dest Endpoint
☐ Source Port
☐ Traffic Group ID
☐ MPLS Flow Descriptor
☐ Frame Size
☐ Flow Group
☐ Ethernet II : Destination MAC Address
☐ Ethernet II : Source MAC Address
☐ Ethernet II : Ethernet-Type
☐ Ethernet II : PFC Queue
☐ VLAN : VLAN Priority
☐ VLAN : VLAN Priority(1)
☑ VLAN : VLAN-ID
☐ VLAN : VLAN-ID(1)
☑ TRILL : Egress RBridge Nickname
☑ TRILL : Ingress RBridge Nickname
☐ Ethernet II without FCS : Destination MAC A...

Custom Override

☐ One - One meshed

Offset from Root
Offset
Field width 32 Bits
Values

Figure 329. Selecting flow tracking options

Test Variables

Performance Variable	Description
No of Campus VLANS per RBridge	Use this test variable to increase the count of Interested VLANs emulated for each directly connected RBridge
No of MACs per VLAN	Use this test variable to increase the count of end stations emulated for each campus VLAN
No of Multicast Group Receivers per VLAN	Use this test variable to increase the count of MAC muticast receivers emulated for each campus VLAN
No of Multicast IPv4 Group Receivers per VLAN	Use this test variable to increase the count of IPv4 multicast receivers emulated for each campus VLAN
No of Mulicast IPv6 Group Receivers per VLAN	Use this test variable to increase the count of IPv6 multicast receivers emulated for each campus VLAN.
Number of Grids	Use this test variable to increase the count of RBridge matrix topologies emulated within the Campus.
Number of Rows in this Grid	Use this test variable to increase the row count within an emulated matrix topology of RBridges.
Number of Columns in this Grid	Use this test variable to increase the column count within an emulated matrix topology of RBridges.

Troubleshooting and Diagnostics

Issue	Troubleshooting Solution
Loss of throughput on VLAN traffic	Check DUT settings to ensure: • Campus VLAN ID is included Interested VLANs • Multicast Destination Trees have been correctly computed Check test ports to ensure: • Correct priority mapping • Correct source unicast address configuration • Correct mulitiicast destination address configuration

Results Analysis

1. Click **Start All Traffic** and select **Traffic Item 1** from the **Traffic Item Statistics** view. Right-click to open the context menu and select **Drill down per VLAN: VLAN-ID** to see throughput and packet loss statistics on a per vlan basis in the **User Defined Statistics** view.

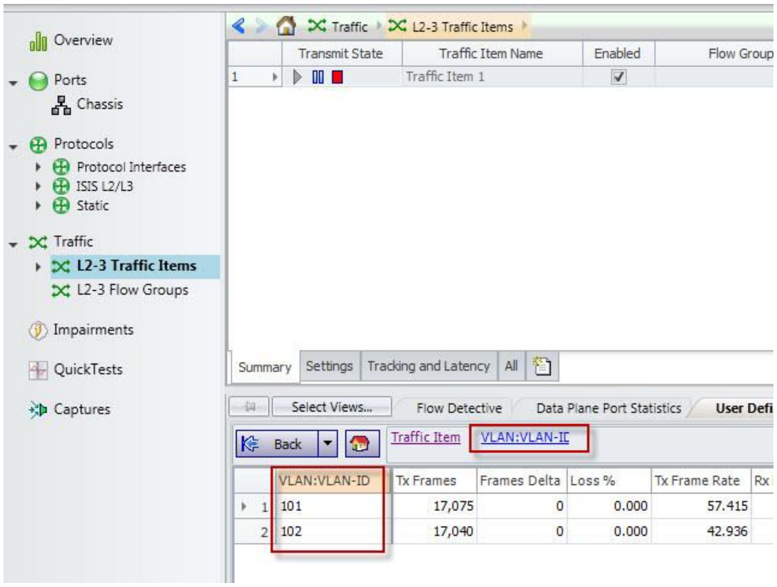

	Transmit State	Traffic Item Name	Enabled	Flow Group
1		Traffic Item 1	☑	

	VLAN:VLAN-ID	Tx Frames	Frames Delta	Loss %	Tx Frame Rate	Rx
1	101	17,075	0	0.000	57.415	
2	102	17,040	0	0.000	42.936	

Figure 330. Reviewing per vlan statistics in User Defined Statistics view

2. Select **Traffic Item 1** from the **Traffic Item Statistics** view. Right-click to open context menu and select **Drill down per TRILL: Ingress RBridge Nickname** to see throughput and packet loss statistics on a per ingress RBridge basis in the **User Defined Statistics** view.

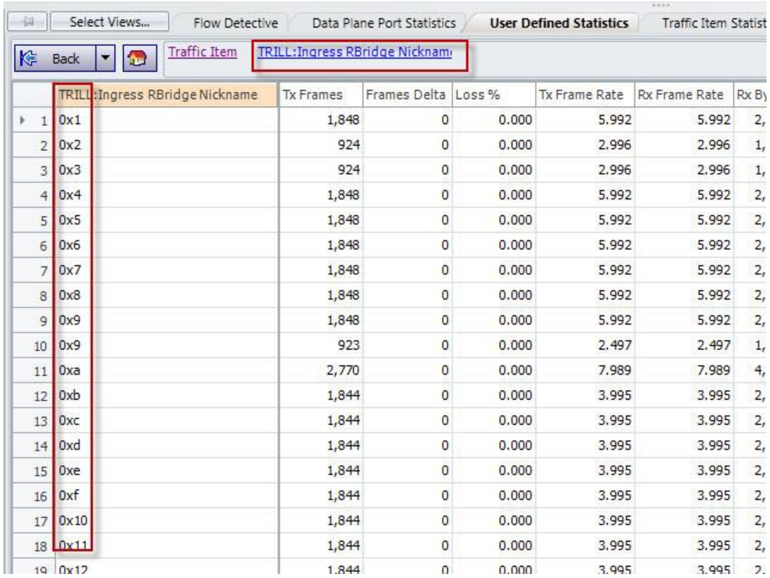

	TRILL:Ingress RBridge Nickname	Tx Frames	Frames Delta	Loss %	Tx Frame Rate	Rx Frame Rate	Rx By
1	0x1	1,848	0	0.000	5.992	5.992	2,
2	0x2	924	0	0.000	2.996	2.996	1,
3	0x3	924	0	0.000	2.996	2.996	1,
4	0x4	1,848	0	0.000	5.992	5.992	2,
5	0x5	1,848	0	0.000	5.992	5.992	2,
6	0x6	1,848	0	0.000	5.992	5.992	2,
7	0x7	1,848	0	0.000	5.992	5.992	2,
8	0x8	1,848	0	0.000	5.992	5.992	2,
9	0x9	1,848	0	0.000	5.992	5.992	2,
10	0x9	923	0	0.000	2.497	2.497	1,
11	0xa	2,770	0	0.000	7.989	7.989	4,
12	0xb	1,844	0	0.000	3.995	3.995	2,
13	0xc	1,844	0	0.000	3.995	3.995	2,
14	0xd	1,844	0	0.000	3.995	3.995	2,
15	0xe	1,844	0	0.000	3.995	3.995	2,
16	0xf	1,844	0	0.000	3.995	3.995	2,
17	0x10	1,844	0	0.000	3.995	3.995	2,
18	0x11	1,844	0	0.000	3.995	3.995	2,
19	0x12	1,844	0	0.000	3.995	3.995	2,

Figure 331. Reviewing per ingress RBridge statistics in User Defined Statistics view

3. Expand the **ISIS L2/3** protocol tree for the Core Side port. Select **Learned Information** and then click the **Refresh** button to review all of the reachability information shared by the DUT.

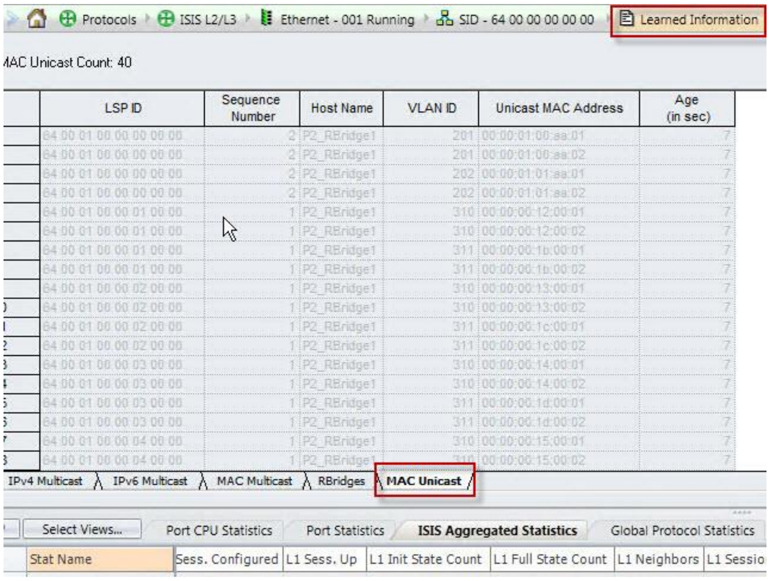

MAC Unicast Count: 40

LSP ID	Sequence Number	Host Name	VLAN ID	Unicast MAC Address	Age (in sec)
64 00 01 00 00 00 00 00	2	P2_RBridge1	201	00:00:01:00:aa:01	7
64 00 01 00 00 00 00 00	2	P2_RBridge1	201	00:00:01:00:aa:02	7
64 00 01 00 00 00 00 00	2	P2_RBridge1	202	00:00:01:01:aa:01	7
64 00 01 00 00 00 00 00	2	P2_RBridge1	202	00:00:01:01:aa:02	7
64 00 01 00 00 01 00 00	1	P2_RBridge1	310	00:00:00:12:00:01	7
64 00 01 00 00 01 00 00	1	P2_RBridge1	310	00:00:00:12:00:02	7
64 00 01 00 00 01 00 00	1	P2_RBridge1	311	00:00:00:1b:00:01	7
64 00 01 00 00 01 00 00	1	P2_RBridge1	311	00:00:00:1b:00:02	7
64 00 01 00 00 02 00 00	1	P2_RBridge1	310	00:00:00:13:00:01	7
64 00 01 00 00 02 00 00	1	P2_RBridge1	310	00:00:00:13:00:02	7
64 00 01 00 00 02 00 00	1	P2_RBridge1	311	00:00:00:1c:00:01	7
64 00 01 00 00 02 00 00	1	P2_RBridge1	311	00:00:00:1c:00:02	7
64 00 01 00 00 03 00 00	1	P2_RBridge1	310	00:00:00:14:00:01	7
64 00 01 00 00 03 00 00	1	P2_RBridge1	310	00:00:00:14:00:02	7
64 00 01 00 00 03 00 00	1	P2_RBridge1	311	00:00:00:1d:00:01	7
64 00 01 00 00 03 00 00	1	P2_RBridge1	311	00:00:00:1d:00:02	7
64 00 01 00 00 04 00 00	1	P2_RBridge1	310	00:00:00:15:00:01	7
64 00 01 00 00 04 00 00	1	P2_RBridge1	310	00:00:00:15:00:02	7

Figure 332. Reviewing learned information advertised by DUT

Conclusions

The IxNetwork TRILL Protocol Wizard enables the configuration of complex RBridge topologies along with all of the campus vlans, unicast, and multicast stations associated with large scale TRILL deployment. The flexible flow tracking in the **Advanced Traffic Wizard** along with the real time drill down capability in statistics views allows the user to quickly isolate flows that are not being forwarded correctly by the DUT and detect the root cause by associating them with individual RBridges within the emulated topology.

Introduction to Storage Technologies

Controlling the ever-growing data is still the greatest challenge as far as storage is concerned. Data is a critical component of any enterprise or organization and it increases every day. In many businesses, the volume of data is amplifying by twice every year. Another trend being observed is the growth of all non-structured data (file-based) that exceeds structured data (block-based). Balance between structured data entered in databases or ERP applications and non-structured data resulting from e-mails, Office files, presentations, videos and so on, has shifted.

Handling the massive increase in data is the prime task. The end users must possess quick access to data in addition to ensuring safe backup as a contingency plan. There are many options for data storage, but no storage system providing single solution for all use cases. This section focusses on the technologies, benefits, and methodologies to validate different storage options matching various enterprises and converged data center requirements.

Types of Storage

There are three types of storage: block, file, and object. Each type offers their own advantages and has their own use cases.

Types of Storage	
Block Storage	Block storage gives access to the 'bare metal'. There is no concept of 'files' at this level. There are just evenly sized blocks of data. Generally, using block storage offers the best performance, but it is quite simple. Database, Exchange, some VMware, and Server Boot/VDI often take advantage of block storage systems.
File Storage	File storage provides simple access to a file system. This is the most familiar kind of storage–it is what we interact with most on a daily basis. Users of file storage have access to files and can read and write to either the whole file or a part of it. File systems are what operating systems provide on all of our personal computers. In a shared environment, file storage is often seen as a network drive.
Object Storage	Object storage is probably the least familiar type of storage to most people. Object storage does not provide access to raw blocks of data. It doesn't offer file-based access. Object storage provides access to whole objects, or blobs of data and generally does so with an API specific to that system. Unlike file storage, object storage generally does not allow the ability to write to one part of a file. Objects must be updated as a whole unit. Example of object orinteated storage is Amazon S3.

Types of Storage Architecures

There are five key storage architectures in play in today's networks. Each option is suitable for different application requirements and offers varying degree of scalability, reliability, performance, availability, affordability and manageability.

- Direct attached storage (DAS) – the simplest type of data storage, located in or attached directly to a server
- Network attached storage (NAS) – a dedicated file server attached to a local area network, running an operating system that is dedicated specifically to file serving
- Storage area network (SAN) – a dedicated network for storage of traffic between servers and a disk storage array or tape device
- Hybrid or Unified storage – a combination of a NAS gateway with a SAN and a hybrid NAS/SAN, with a simple management interface to hide the complexity
- Cloud storage - is a model of networked online storage where data is stored in virtualized pools of storage which are generally hosted by third party converged data centers

Architecure	Application	Advantages	I/O Protocols/ Transport
Direct Attached Storage (DAS)	Local operating system	Useful for small, predictable storage No network or sharing is required	I/O – SCSI Transport media could be any (that is, Fibre Channel, SCSI, SSA, and Ethernet).
Network Attach Stroage (NAS)	Primarily targeted at storing and sharing files Useful for data backup or simple storage	Simple to configure and maintain Best for low-volume file sharing between multiple peer clients which are less sensitive to response time	I/O – SMB (CIFS), NFS Transport Media – TCP/IP Ethernet
Storage Area Network (SAN)	Mission critical and IO intensive applications Backup and restore Business Continuance	High performance, scalability and availability	I/O – SCSI Transport Media – Fibre Channel and Ethernet (iSCSI and FCoE)
Hybrid or Unfied Storage	When an organization has existing investment in SAN and wanting to expand	Convergence and reduced operation cost Ease of maintance	I/O – SCSI, SMB (CIFS) and NFS Transport – Both Ethernet and FC
Cloud Storage	Highly fault tolerant through redundancy and	Pay only for storage actually used	I/O – HTTP based

Architecure	Application	Advantages	I/O Protocols/ Transport
	distribution of data	Storage maintenance tasks, such as backup, data replication, and purchasing additional storage devices are offloaded to the responsibility of the cloud provider	RESTful AP Transport - Ethernet

Key Storage Performance Indicators

There are three independent metrics of storage performance:

- IOPS (I/O transactions per second)
- Bandwidth and
- IO response time.

Understanding the relationships between these metrics is the key to understanding and validating storage performance.

Bandwidth is really just a limitation of the design or networking technology that are used to connect storage, 10G Ethernet vs 4G Fibre Channel. It is the maximum number of bytes that can be moved in a specific time period. IOPS are nothing more than the number of I/O transactions that can be performed in a single second. Determining the maximum theoretical IOPS for a given transfer size is as simple as dividing the maximum bandwidth by the transfer size.

There is tight coupling between response time and concurrency. Ultimately, the limitation of IOPS performance is the ability of a system to handle outstanding I/Os concurrently. Once that limit is reached, the I/Os get clogged up and the response time increases rapidly. This is the reason a common tactic to increase storage performance has been to simply add disks – each additional disk increases the concurrent I/O capabilities.

Interestingly, the IOPS performance isn’t limited by the response time. Lower response times merely allow a given level of IOPS to be achieved at lower levels of concurrency. There are practical limits on the level of concurrency that can be achieved by the interfaces to the storage (for example, the execution throttle setting in an HBA), and many applications have fairly low levels of concurrent I/O, but the response time by itself does not limit the IOPS. This is why even though Flash media has a higher response time than RAM, Flash systems that handle a high level of concurrent I/O can achieve as good as or better IOPS performance.

Test Case: Validate the iSCSI I/O Performance of SAN Array

Overview

iSCSI (Internet Small Computer System Interface), an Internet Protocol (IP)-based storage networking standard for linking data storage facilities. Unlike Fibre channel which requires a dedicated adapter (that is, FC HBA) and a dedicated network infrastructure, iSCSI runs over TCP/IP allowing it to easily operate over WANs and the IP networks that already exist in enterprise networks. iSCSI maps SCSI commands and data over the TCP protocol, where traditionally the same SCSI commands were sent over FC.

The 10 GigE and higher speed enhancements to Ethernet and emergence of the Data Center Bridging (DCB) standards makes iSCSI a stronger storage player in the data centers because it mitigates the performance and reliability concerns of Ethernet.

iSCSI has expanded from $18 million in revenue in 2003 to $3 billion in 2011, according to market research firm IDC. As iSCSI SAN arrays grow in popularity and start to adopt 10G support, it is crical to prove that these devices are capable of handling the various storage demands of enterprise applications. This is achieved by benchmarking the performance and scalabiltiy limits of as these SAN arrays with varying IO parmeters – Read/Write IO propositions, block size of individual iSCSI commands, sequential or random nature of the IO seeks are key parameters that can have direct influence on storage performance.

Objective

The objective of this test is to validate the IO performance of a SAN array that is acting as the iSCSI target. Ixia IxLoad emulates the iSCSI initiators that will simulate the IO workloads. The Ixia initiators first discovers the different Target groups and the Logical Units present under them. Once they are discovered, the IXIA initiators run IO test involving SCSI read and write commands with varying command block sizes, Random and Sequential ness of Reads and Writes and percentage of Reads and writes. The IO performances are monitored for each of these configuration changes.

The test results measure the peak IO performance, the drill down IO performance of Read and Write commands. It will also monitor the average latencies, SCSI protocol status replies including errors and the pattern validations to verify data integrity.

Setup

One Ixia port connected against an External SAN Array Target.

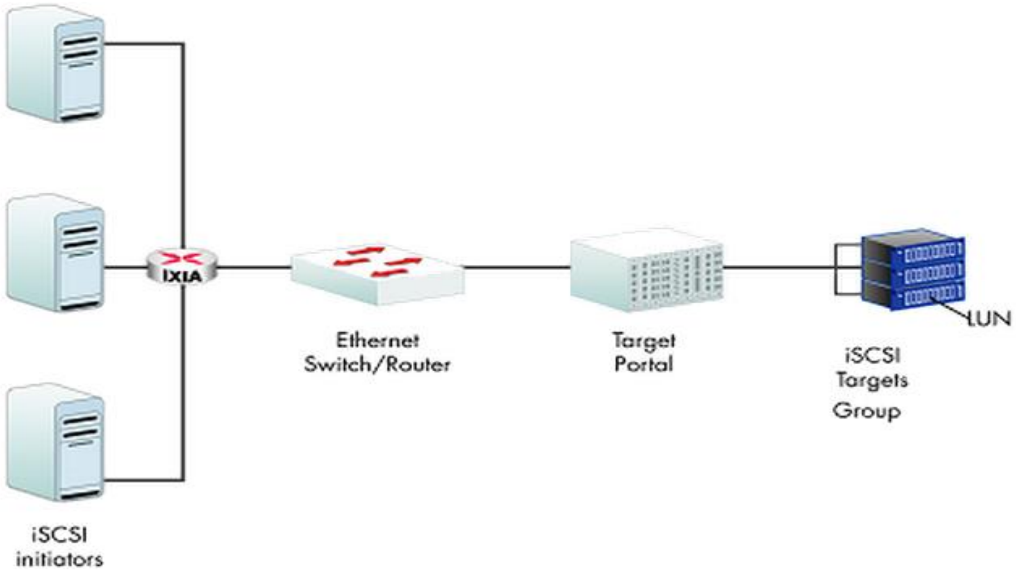

Figure 333. iSCSI test setup with IxLoad iSCSI initiators and external iSCSI target.

Target DUT Configurations (Optional):

This optional section details the target configurations that are assumed in the test case description. The tester must ensure that their DUT is configured correctly for the testing.

The **target Portal** of the device in this test case is **1.1.1.10**

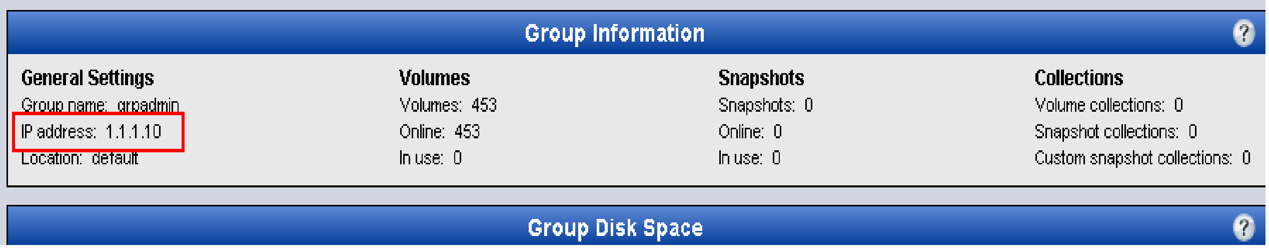

General Settings	Volumes	Snapshots	Collections
Group name: grpadmin	Volumes: 453	Snapshots: 0	Volume collections: 0
IP address: 1.1.1.10	Online: 453	Online: 0	Snapshot collections: 0
Location: default	In use: 0	In use: 0	Custom snapshot collections: 0

Figure 334. Target Portal of the Device Under Test(DUT)

The device used has **453 online volumes** where each volume actually relates to a Target Group and under it one single Logical Unit.

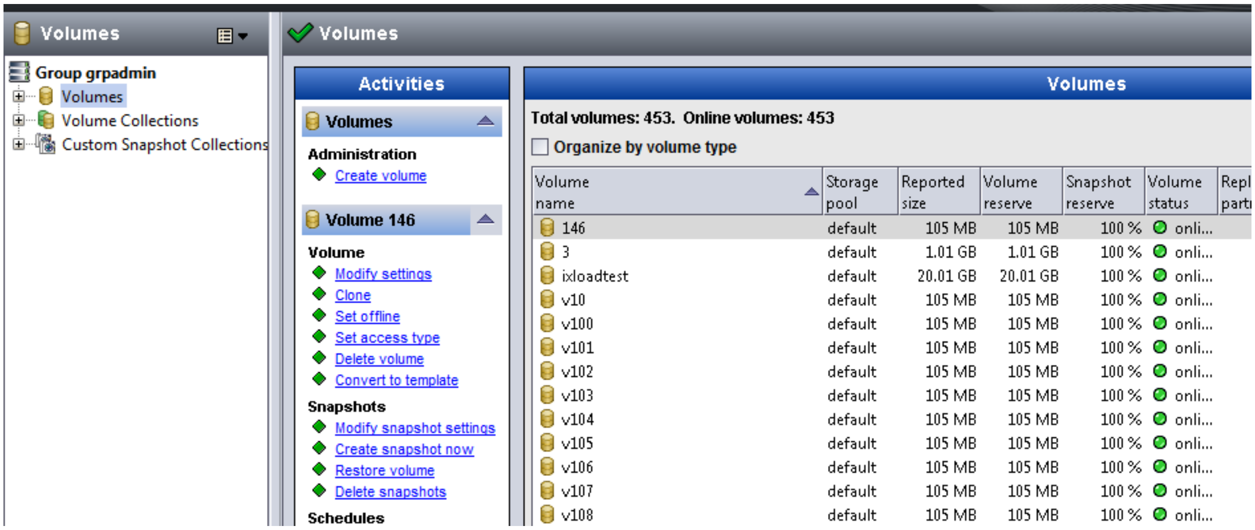

Total volumes: 453. Online volumes: 453

☐ Organize by volume type

Volume name	Storage pool	Reported size	Volume reserve	Snapshot reserve	Volume status	Repl partn
146	default	105 MB	105 MB	100 %	onli...	
3	default	1.01 GB	1.01 GB	100 %	onli...	
ixloadtest	default	20.01 GB	20.01 GB	100 %	onli...	
v10	default	105 MB	105 MB	100 %	onli...	
v100	default	105 MB	105 MB	100 %	onli...	
v101	default	105 MB	105 MB	100 %	onli...	
v102	default	105 MB	105 MB	100 %	onli...	
v103	default	105 MB	105 MB	100 %	onli...	
v104	default	105 MB	105 MB	100 %	onli...	
v105	default	105 MB	105 MB	100 %	onli...	
v106	default	105 MB	105 MB	100 %	onli...	
v107	default	105 MB	105 MB	100 %	onli...	
v108	default	105 MB	105 MB	100 %	onli...	

Figure 335. The diferent Target Groups/LUN's configured in the device

Some of the target groups have been assigned **CHAP** Authentication while others have no authentiactions. All have a common CHAP username of '**ixia**' and password of '**ixia**'. All the volumes allow access to any IPs that begin with **1.1.*.**

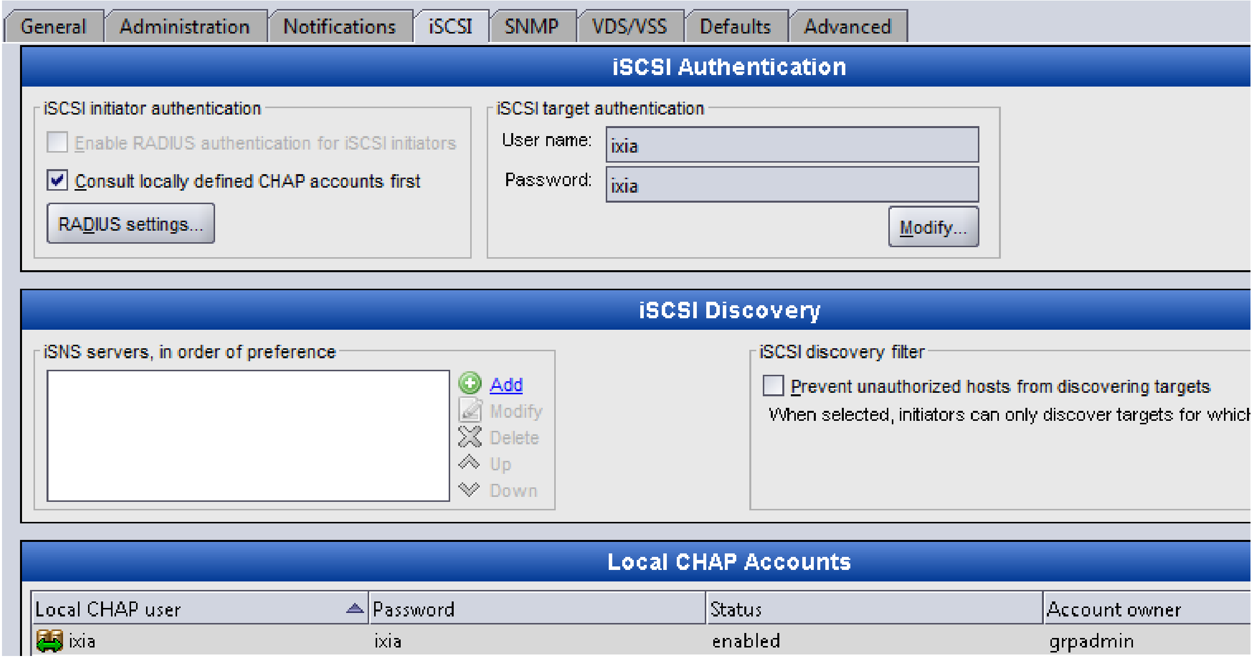

iSCSI initiator authentication

☐ Enable RADIUS authentication for iSCSI initiators

☑ Consult locally defined CHAP accounts first

iSCSI target authentication

User name: ixia

Password: ixia

iSNS servers, in order of preference

iSCSI discovery filter

☐ Prevent unauthorized hosts from discovering targets

When selected, initiators can only discover targets for which

Local CHAP user	Password	Status	Account owner
ixia	ixia	enabled	grpadmin

Figure 336. The Authentication configured at the Target Portal for Discovery

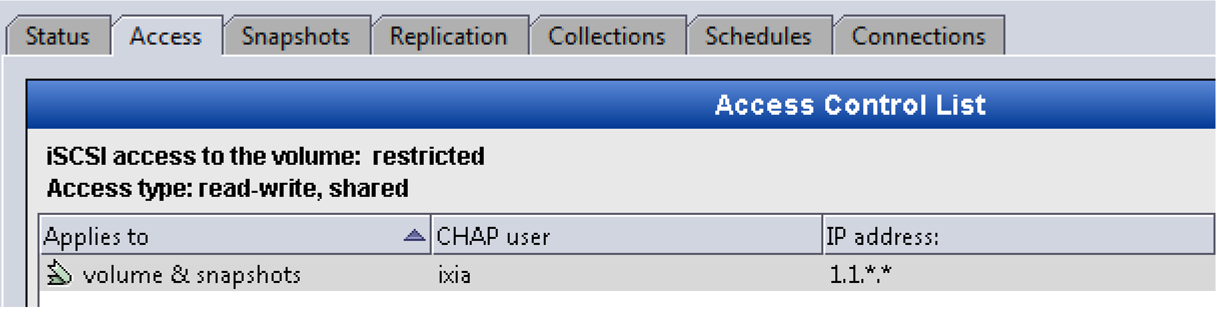

iSCSI access to the volume: restricted
Access type: read-write, shared

Applies to	CHAP user	IP address:
volume & snapshots	ixia	1.1.*.*

Figure 337. The IP access configured at the Target Group.

Step-by-Step Instructions

1. Start IxLoad. In the main window, the **Scenario Editor** window appears. All tests are configured here.

2. Add a NetTraffic at the Originate side.

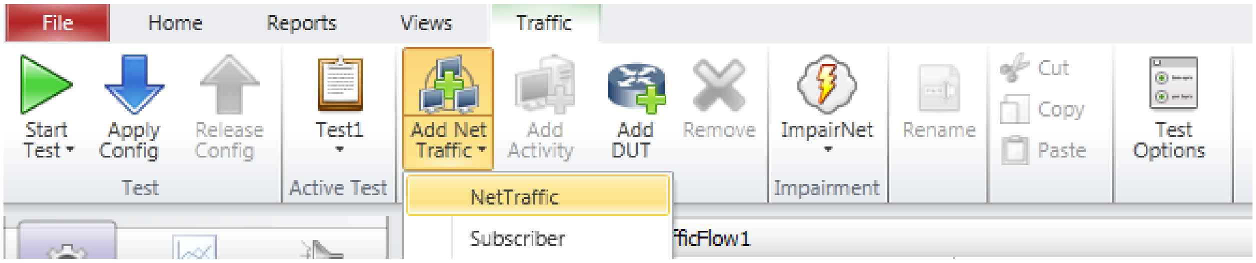

Figure 338. IxLoad NetTraffic addition

3. Configure the Initiator network with a total of 100 IP addresses and gateway so that the Initiators can reach the target. The IPs must belong to **1.1.*** network because this IP has access to the volumes (see Figure 264). In this testcase, the Ixia port is connected to the DUT through an L2 switch. Hence we do not require toprovide any gateway IP.

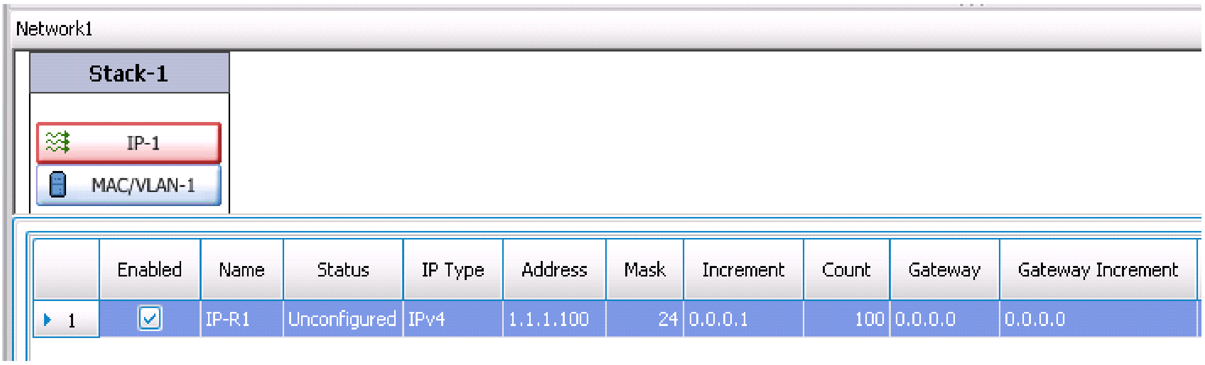

	Enabled	Name	Status	IP Type	Address	Mask	Increment	Count	Gateway	Gateway Increment
1	☑	IP-R1	Unconfigured	IPv4	1.1.1.100	24	0.0.0.1	100	0.0.0.0	0.0.0.0

Figure 339. Setting the Correct IP address at the client side

4. Verify that for the first iteration, the MSS is set to 1460 and MTU to 1500.

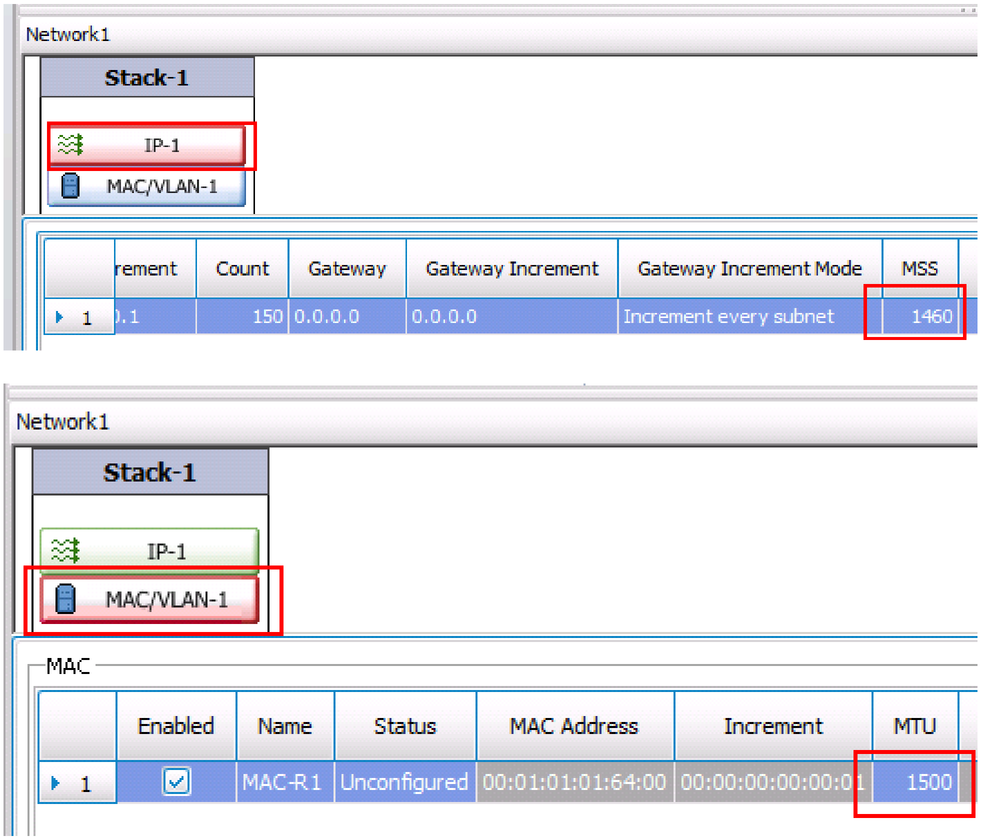

	rement	Count	Gateway	Gateway Increment	Gateway Increment Mode	MSS
1	).1	150	0.0.0.0	0.0.0.0	Increment every subnet	1460

	Enabled	Name	Status	MAC Address	Increment	MTU
1	☑	MAC-R1	Unconfigured	00:01:01:01:64:00	00:00:00:00:00:01	1500

Figure 340. Configuring the MSS and MTU

5. Add an iSCSI Client (Initiator) activity under the NetTraffic.

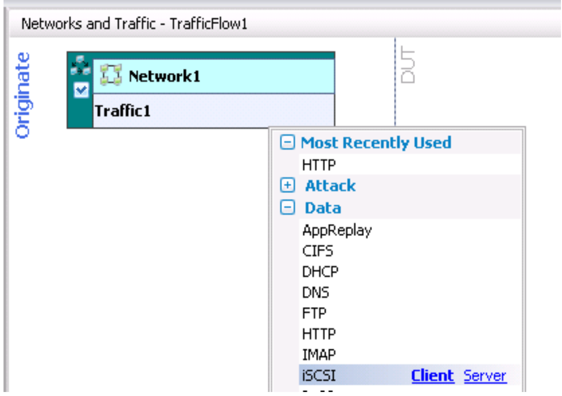

Figure 341. Add the iSCSI Client plugin under the Network1

6. Configure the iSCSI Initiator activity to discover the different Target groups and the Logical Units present on the external target (DUT) by selecting the **Discovery Options** tab in Settings. Select the **Perform Discovery** checkbox.

    a. In the **Portal Address** field enter the portal IP of the Target Portal of the DUT. As shown in DUT configuration, the Portal IP is **1.1.1.10**

    b. In the **Authentication Settings,** set the correct **Authentication Type**, **User Name** and **Password**. Some of the device's Target Groups do not have Authentiaction enabled, hence set as **CHAP,None.** With this setting, discovery can discover all the target groups under the target portal irrespective of it having authentication enabled or not.

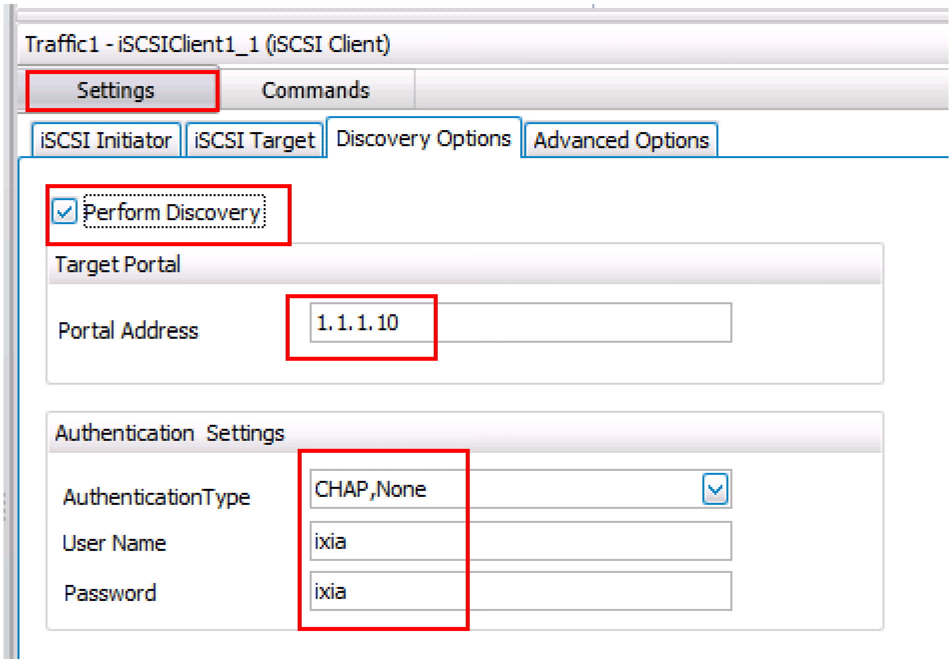

☑ Perform Discovery

Target Portal

Portal Address: 1.1.1.10

Authentication Settings

AuthenticationType: CHAP,None

User Name: ixia

Password: ixia

Figure 342. Setting the correct Target Portal IP and the Target Authentications for discovering the target groups.

7. Add Chassis, using the **Add Chassis** function and provide the **chassis IP address** that is connected to the DUT.

8. Assign the correct **card** and **port** for the chassis to the NetTraffic.

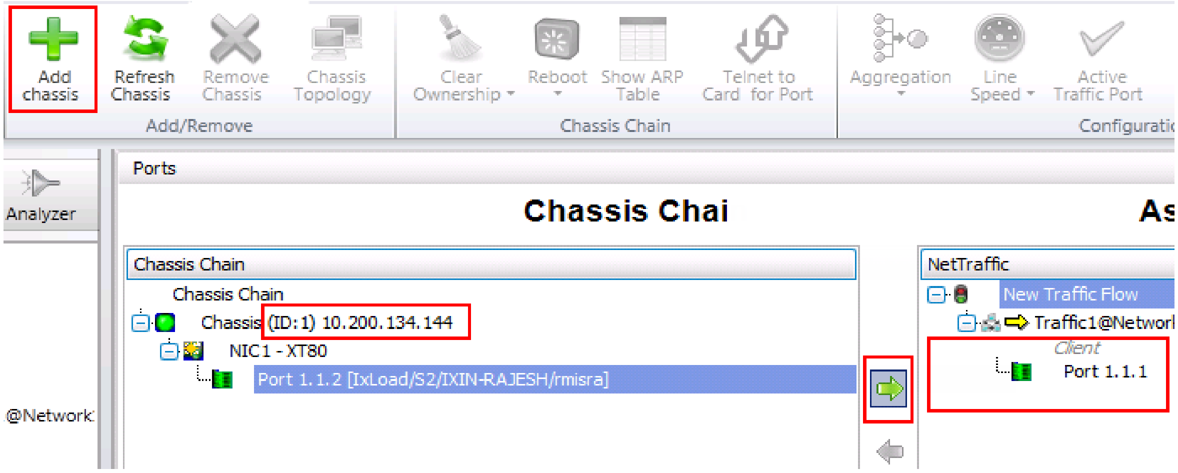

Figure 343. Assigning the Chassis, Card and port to the Network..

9. Start the Discovery by selecting the **Apply Config** button.

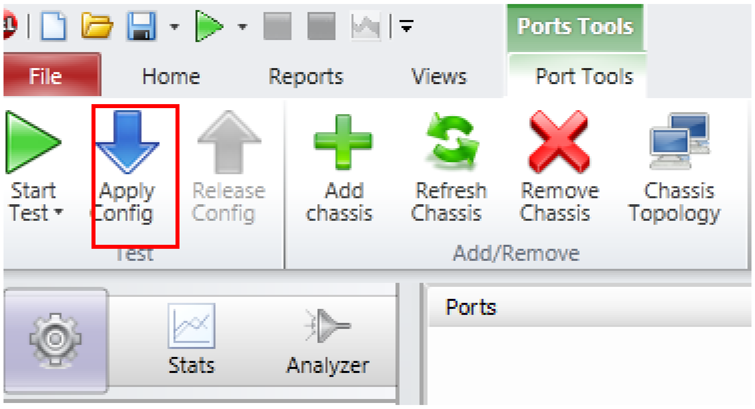

Figure 344. Apply configurations for the discovery.

10. The completion of Discovery is denoted by a successful operation of the Apply Config action as seen in the logs. After completion, click **Settings** > **iSCSI Target** tab to ensure that all the Discovered Targets and the underlying LUNs are visible.

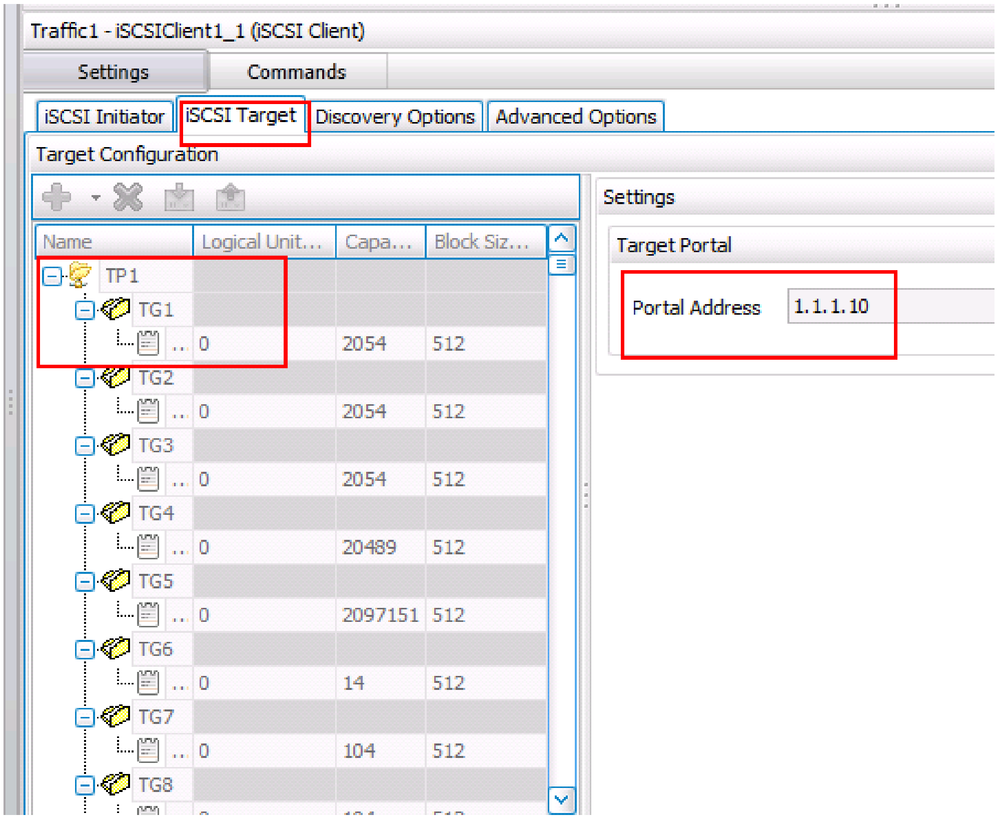

Name	Logical Unit...	Capa...	Block Siz...
TP1			
TG1			
..	0	2054	512
TG2			
..	0	2054	512
TG3			
..	0	2054	512
TG4			
..	0	20489	512
TG5			
..	0	2097151	512
TG6			
..	0	14	512
TG7			
..	0	104	512
TG8			

Portal Address 1.1.1.10

Figure 345. Verify that the Discovery has been successful and the Targets and LU's are being displayed.

11. Once the Discovery is complete, move to the next phase, which is to stress all the discover targets with simultaneous IO operations from the emulated iSCSI initiators. To configure the initiators click **Settings** > **iSCSI Initiator** tab.

    a. Set the pipeline to **2048**. It is safe to set it to a high value, because the pipeline or queue depth is controlled by the iSCSI target.

    b. It is advisable to set Initial R2T to NO, becase we perform both read and write operations. Note: Setting Immediate Data to 'Yes' can further increase IO performance. However, this setting can sometimes be counter-productive resulting in decrease in the performance over time due to heavy stress on the device.

    c. Verify the **Initiator Configuration** values. These do not need to be changed for an IO performance test. For throughput test, increasing the **MaxRcvDataSegmentLength** or **FIrstBirstLength** can be helpful.

    d. Enable Pattern Matching check box must be cleared, because the test consists of a random percentage of Reads and Writes. Select Pattern Matching if it is assured that the particular blocks are already pre-written with a known pattern for data integrity validation.

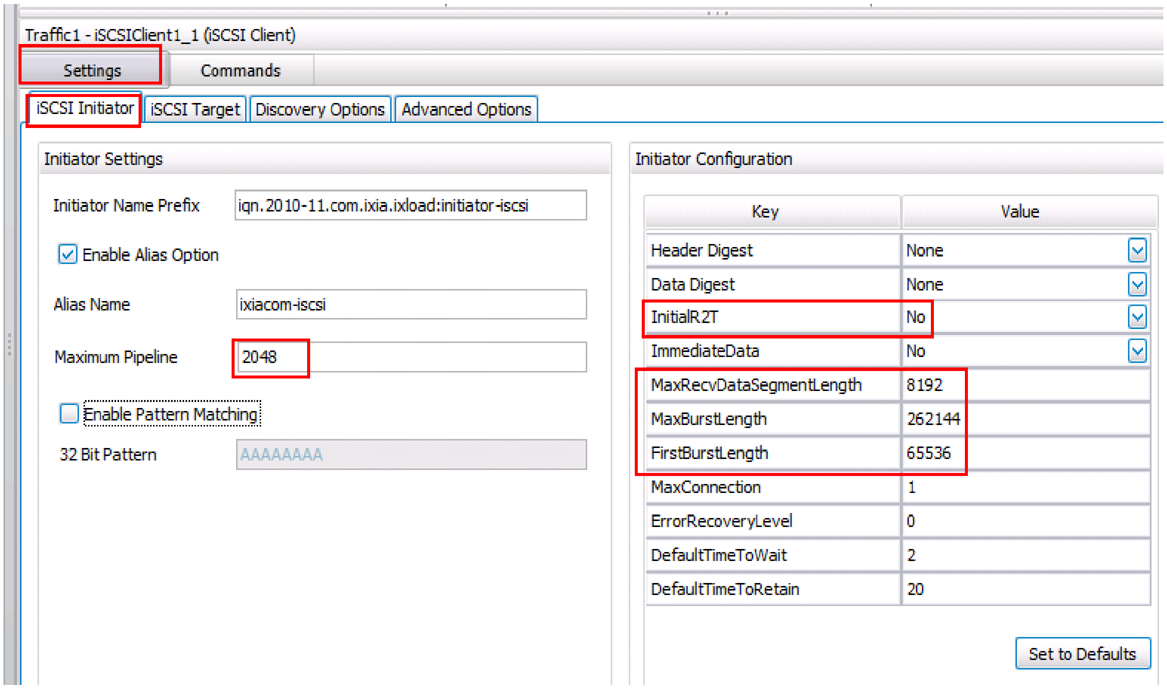

Initiator Name Prefix: iqn.2010-11.com.ixia.ixload:initiator-iscsi

☑ Enable Alias Option

Alias Name: ixiacom-iscsi

Maximum Pipeline: 2048

☐ Enable Pattern Matching

32 Bit Pattern: AAAAAAAA

Key	Value
Header Digest	None
Data Digest	None
InitialR2T	No
ImmediateData	No
MaxRecvDataSegmentLength	8192
MaxBurstLength	262144
FirstBurstLength	65536
MaxConnection	1
ErrorRecoveryLevel	0
DefaultTimeToWait	2
DefaultTimeToRetain	20

Figure 346. Configure Queue Depth and other initiator side configurations that will be exchanged with the Target.

12. Add the **Login** command in the **Command** dialoge. At the portal name field, add the portal name as displayed in the **Settings** > **iSCSI Target** tab. For this case, it is being showed as **TP1.**

13. This test is designed to login to 300 of the total targets presents in the device and simultaneously do IO to the first LUN present under each target. This will be achieved by simulating as many numbers of users as the number of targets. By setting **($user-id)** as the Target Name we ensure that each simulated user actually logs in to a different target. [Note: In IxLoad, each user is identified by a number starting from 0. So each Simulated User (Initiator) actually tries to log into a different Target identified by the index of the Simulated User]

14. Set the **Authentication** to match what is configured on the external target (DUT) for all the target groups.

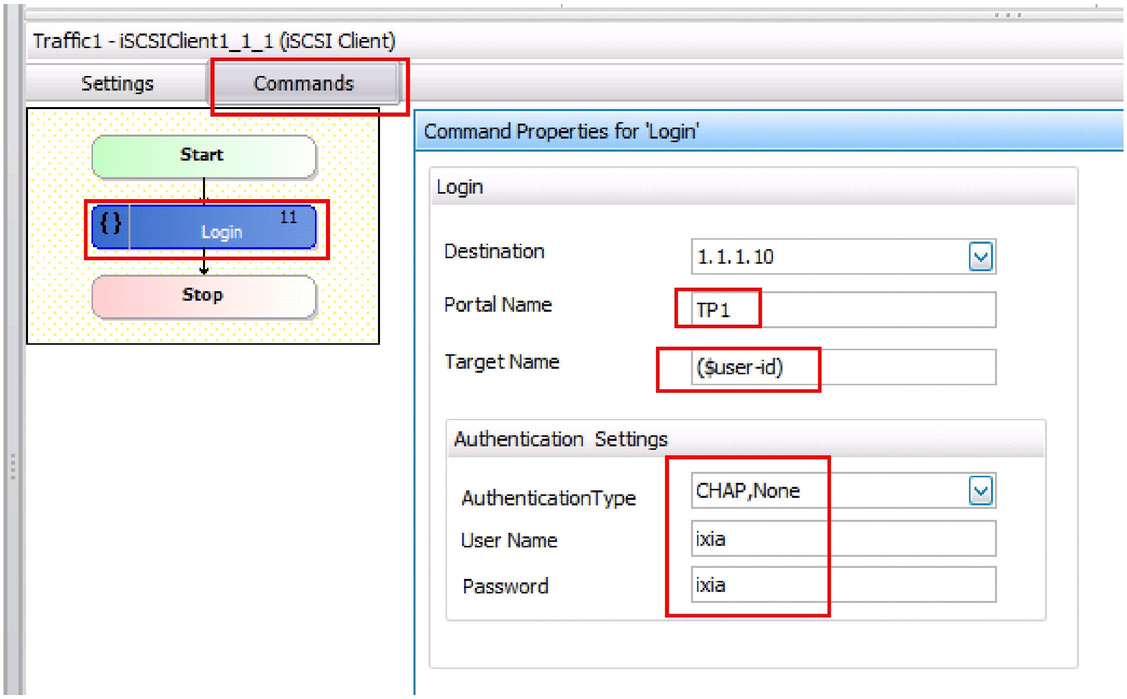
Traffic1 - iSCSIClient1_1_1 (iSCSI Client)
Settings
Commands
Start
Login
11
Stop
Command Properties for 'Login'
Login

Destination 1.1.1.10
Portal Name TP1
Target Name ($user-id)

Authentication Settings
AuthenticationType CHAP,None
User Name ixia
Password ixia

Figure 347. Set the Target Portal Name, Target Group and Authentication

15. Add the **IO** command. This action performs read and write operations. This particular device has only one logical Unit under each target and hence need to access the first LUN only. Set the LUN index as **0,** here ‘0’ indicates the first LUN. If in case there were multiple LUNs under each target, use ($All) token to access all the LUNs or ($start-end) token to access a specific range of LUNs.

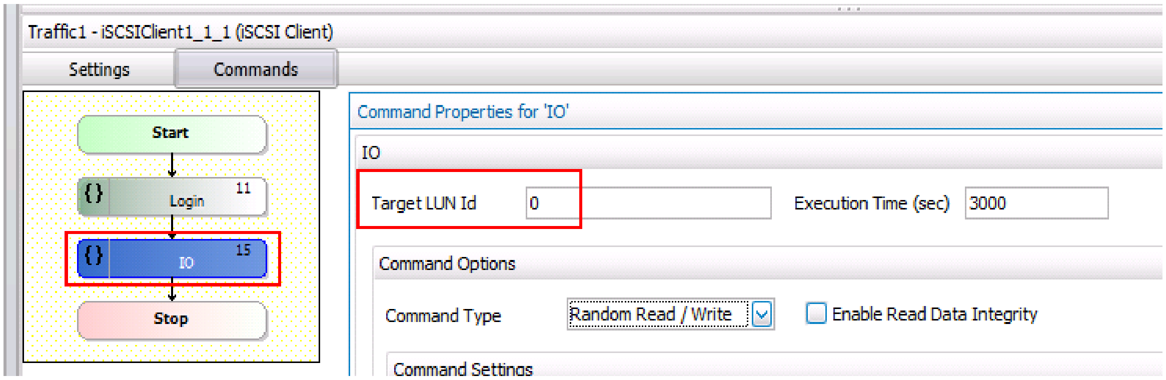
Traffic1 - iSCSIClient1_1_1 (iSCSI Client)
Settings
Commands
Start
Login
11
IO
15
Stop
Command Properties for 'IO'
IO
Command Options
Command Settings

Target LUN Id 0
Execution Time (sec) 3000
Command Type Random Read / Write
☐ Enable Read Data Integrity

Figure 348. Set the LUN within the target

16. To achive continuous IO, set the test **Execution Time** as **3000** seconds. This action ensures that post log on, the IO is repeated for 3000 seconds before the Log out happens. The test is directed to have high **IO** performace; however the configurations still must be a mix of reads and writes to emulate the most realistic load conditions. In this case, we are configuring 70% reads and 30% writes. Set the command as **Random Read / Write** with the **Read Percentage** marked as **70%.**
    [Note: IO performance is highest with Reads as it is a less costly operation for the target device in comparison to wirte]

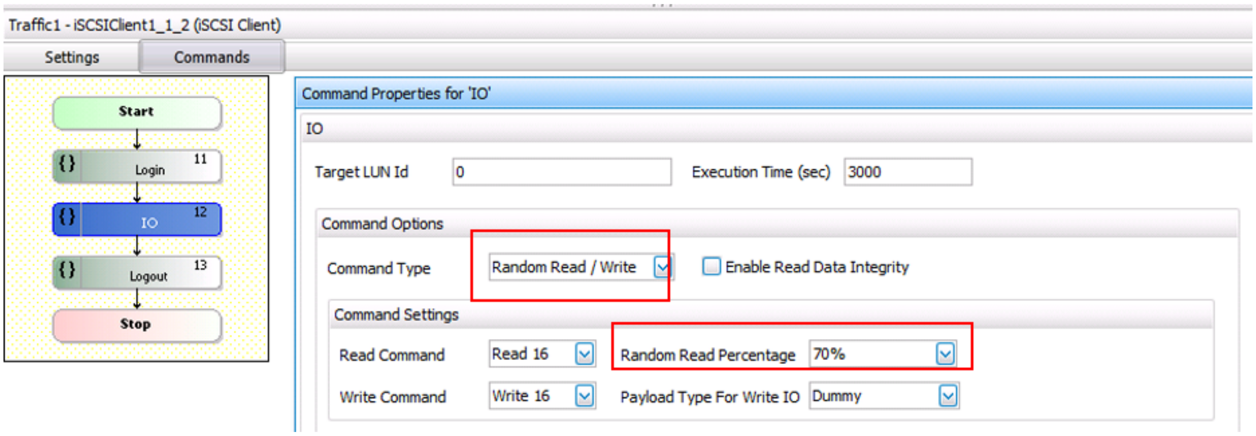

Traffic1 - iSCSIClient1_1_2 (iSCSI Client)

Command Properties for 'IO'

IO

Target LUN Id: 0
Execution Time (sec): 3000

Command Options

Command Type: Random Read / Write
□ Enable Read Data Integrity

Command Settings

Read Command: Read 16
Random Read Percentage: 70%
Write Command: Write 16
Payload Type For Write IO: Dummy

Figure 349. Set the Command Type and the percentage

17. The Payload type must be set to **Dummy,** because integrity check is not performed while measuring the IO performance. Each IO should start at offset **0** and continue till **4 MB**. This is done to ensure that we remain in the cache of the device to optimize the max achievable IO. (Note: Device cache varies for different devices.) The individual block length is **512 Bytes** that is the minimum block size that a single iSCSI read/write command can carry. To truely characterize the IO performance of the device, repeat the test with varying block sizes.

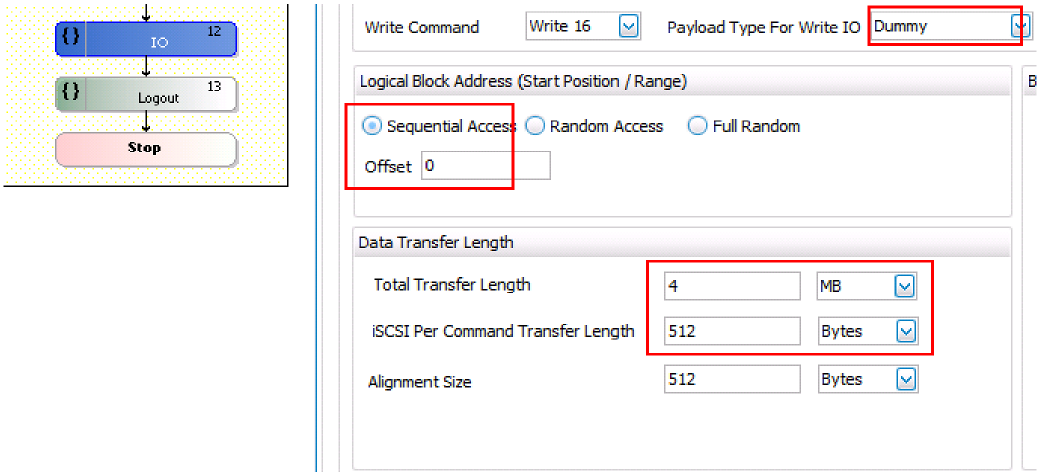

Write Command: Write 16
Payload Type For Write IO: Dummy

Logical Block Address (Start Position / Range)

● Sequential Access ○ Random Access ○ Full Random

Offset: 0

Data Transfer Length

Total Transfer Length: 4 MB
iSCSI Per Command Transfer Length: 512 Bytes
Alignment Size: 512 Bytes

Figure 350. Set the Payload type, IO length and total Bytes per iSCSI read or write

18. IxLoad emulated Initiators are quick in logging in and generating traffic. Simultaneous login to several Target Groups followed by a very high burst of IO can sometimes lead to login failures because the device becomes busy in servicing the IOs. The solution is to introduce a **Think** command (with static or random duration) before and after Login. This action enables enough sleep time for all the logins to be completed before IO burst starts.

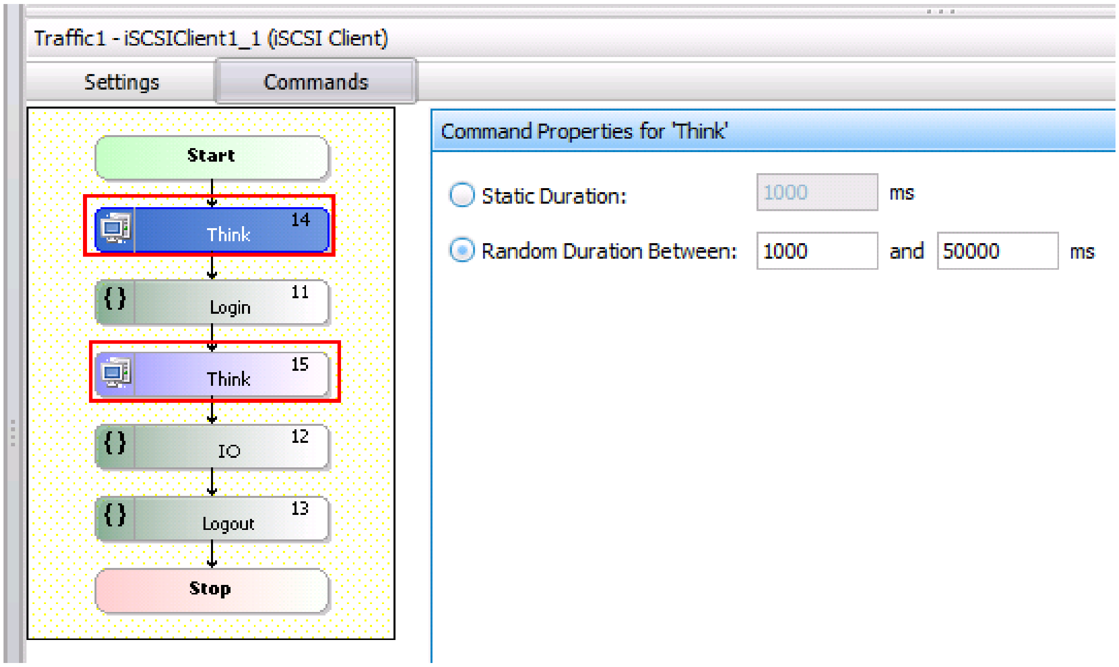

Command Properties for 'Think'

○ Static Duration: 1000 ms

● Random Duration Between: 1000 and 50000 ms

Figure 351. Set the think duration to random to emulate real life scenario

19. Add a **Logout** command at the end to complete the command list.

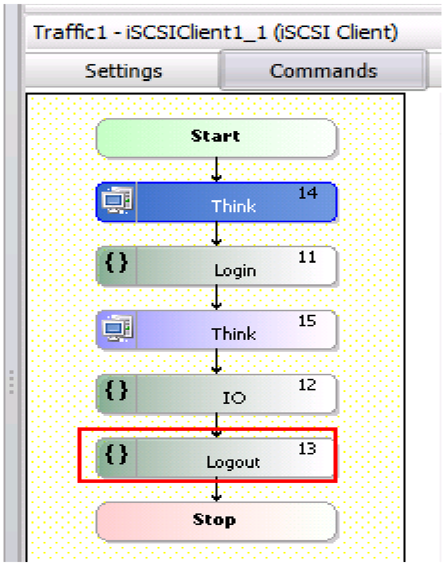

Figure 352. Add the logout command

20. Having setup the iSCSI initiators and the IO workload profile, you can now configure the test objective. In the **Timeline and Objective,** set the objective type as **Simulated User** and Objective Value as **300**. This action ensures that only 300 users or initiators are simulated as there are 300 targets present in the device and each target has a unique user logged in. Refer Step :14 where we have added the **Target Name** as **($user-id)**

21. Set the timeline to **10 Minutes** or more so that you can collect the data for a sufficiently longer duration.

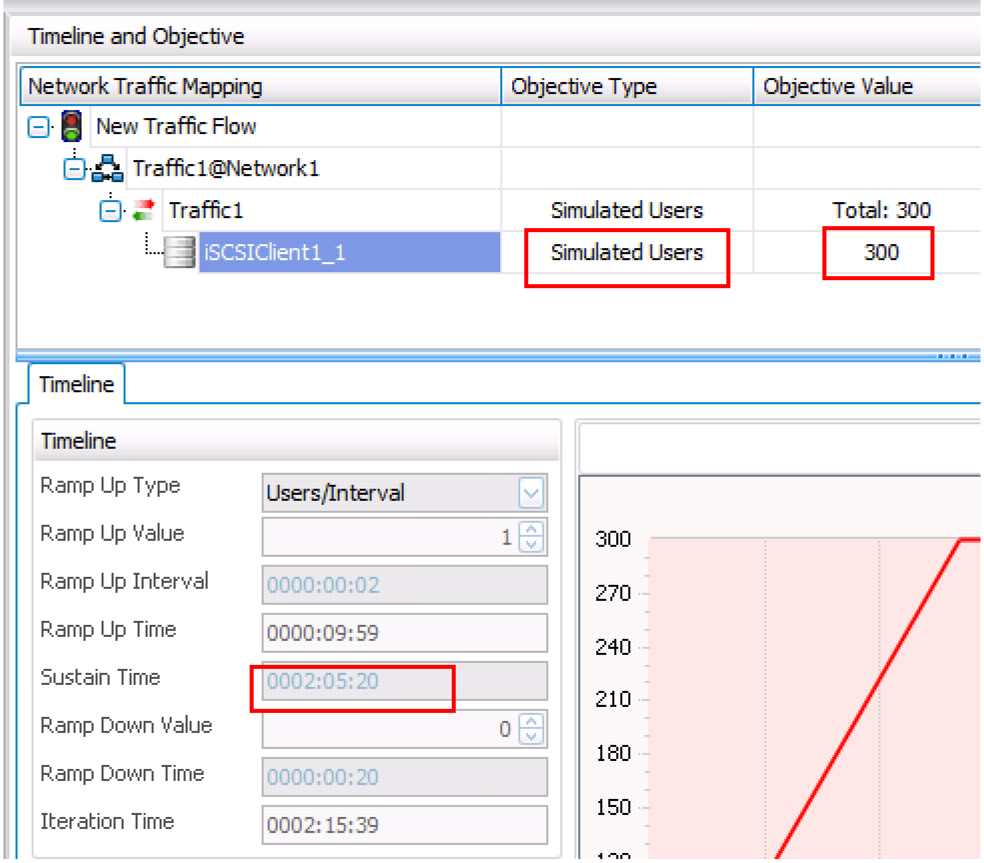

Timeline and Objective

Network Traffic Mapping	Objective Type	Objective Value
New Traffic Flow		
Traffic1@Network1		
Traffic1	Simulated Users	Total: 300
iSCSIClient1_1	Simulated Users	300

Timeline

Timeline

Ramp Up Type	Users/Interval
Ramp Up Value	1
Ramp Up Interval	0000:00:02
Ramp Up Time	0000:09:59
Sustain Time	0002:05:20
Ramp Down Value	0
Ramp Down Time	0000:00:20
Iteration Time	0002:15:39

Figure 353. Set the Objective type and the sustain time

22. Run the test for few minutes to allow the performance to reach a steady-state. Steady state is referred as **Sustain** duration in the test. Continue to monitor the DUT for the IO rate and any failure/error counters. See the **Results Analysis** section below for important statistics and diagnostics information.

    Interpretation of result statistics can sometimes be difficult, deducing what they mean under different circumstances. The **Results Analysis** section below provides a diagnostics-based approach to highlight some common scenarios, the statistics being reported, and how to interpret them.

23. Iterate through the test varying the test variable described in the table below to determine the the IO performance characterization of the SAN Array.

a. You can start and stop the test tool intermediate to a test cycle, or wait for it to be gracefully stopped using the test controls shown here.

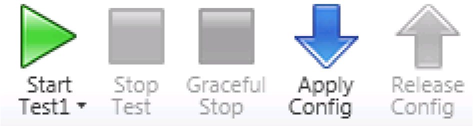

Test Variables

Functional Variable	Description
Block Length per iSCSI command	The iSCSI block length determines the IO rate. A block length of 1MB causes a much lower iSCSI IO rate than a block length of 512 Bytes.
Read/Write Percentage	For a target, write is a costlier operation than read. Hence a change in read/write percentage impacts the performance.
Sequential and Random Read and Write.	This test is configured to have **Reads** and **Writes** be chosen **randomly** similar to the real network. However doing Reads after Writes or vice versa can increase the IO performance, because there is no extra task on IxLoad for randomizations. IO command settings like 'Read then Write', 'Write then Read' and so on are available in the **Command Type** in the IO command.
Logical Block Address(Start position/ Range)	The **LBA start position** is present in the IO command. You can adjust the IO start position to be random per IO start offset (Random Access) or to be random per SCSI command start offset (Full Random).
MTU/MSS	The **MSS** is present in the Network config in the **IP** tab. The **MTU** setting is present in the **MAC** tab of the Network configuration. Increasing these values enables the iSCSI initiators to work with Jumbo frames, which can create throughput performance.
Allignement Size	The Alllignment Size setting is present in the **IO** command. You can use this setting to create Holes between 'Reads' or 'Writes'. Example 'Write' IO command with Block size of 512Bytes and Allignment of 1024Bytes creates write to the alternate blocks and create a hole of (1024 - 512) = 512Bytes for each Writes.
Data integrity validation or Enable Pattern Matching	The filed is present in the **Settings** dialogue in the **iSCSI initiator** tab**.** Select this field to verify that the data written is correct; however when selected, it impacts the tester performance. This flag is most effective if the LUNs are already written with a known pattern.

Functional Variable	Description
Synthetic Pattern	The Synthetic Pattern is present in IO command under **Payload Type for Write IO**. The Synthetic Pattern writes the blocks with random Data patterns and has the provision to later verify the data patterns.
Block Address Overflow handling	The **Block address Overflow handling** is present in the IO command. This option helps in handling the 'read' or 'write' if incase the offset has crossed the total length of the LUN. Once this is selected, the offset is re-adjusted to the start or the length of transfer is truncated to ensure that the IO is still executed.

Following varialble chart details common block lengths and read/write percentages commonly used to verify storage IO performances.

Read Size	Write Size	Read %age	Sequential %age
512Bytes	N/A	100	100
N/A	512B	0	100
1KB	1KB	60	40
8KB	8KB	70	60
64KB	64KB	65	60
256KB	256KB	65	70

Result Analysis

IxLoad Statviewer publishes several stats for analysis and debugs. For an IO performance test, to get the overview, see the total transactions, total throughputs, total number of sessions successful ,the number of reads and writes sent/succeded and the command latency stats.

For debugging purposes see the Async Logout, SCSI Protocol Counters, Login redirections, Command failures, TCP failures and the latency statisitics.

Use the IxLoad statistic view to verify the following results:

- **Total Transactions**: Observe the **iSCSI Initiator – Objectives** view. This observation shows the max IOPS that the device can achieve and the IOPS that it maintains over a period of time.

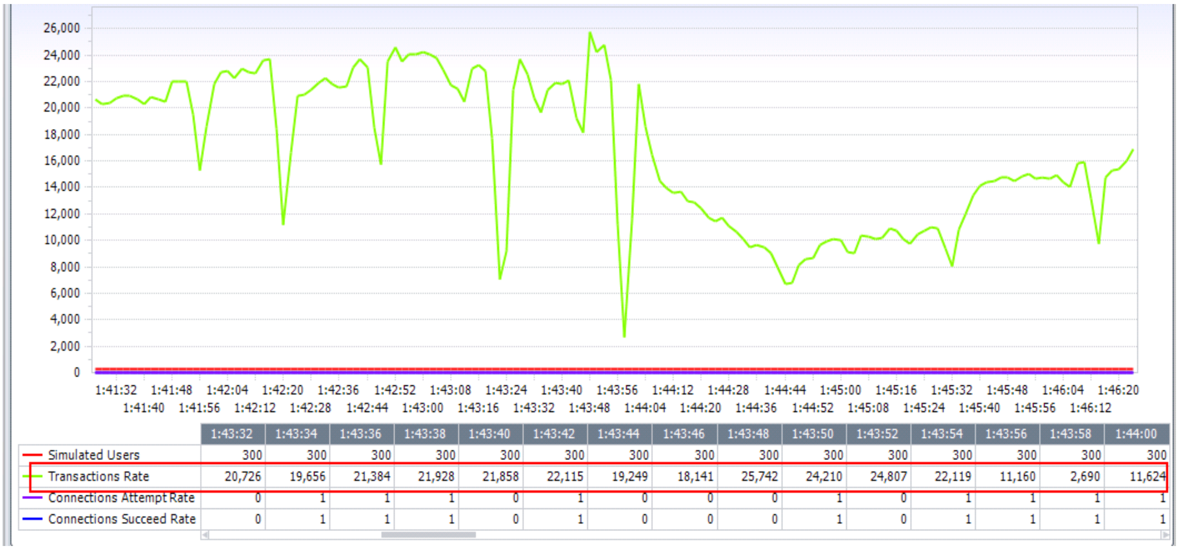

	1:43:32	1:43:34	1:43:36	1:43:38	1:43:40	1:43:42	1:43:44	1:43:46	1:43:48	1:43:50	1:43:52	1:43:54	1:43:56	1:43:58	1:44:00
Simulated Users	300	300	300	300	300	300	300	300	300	300	300	300	300	300	300
Transactions Rate	20,726	19,656	21,384	21,928	21,858	22,115	19,249	18,141	25,742	24,210	24,807	22,119	11,160	2,690	11,624
Connections Attempt Rate	0	1	1	1	0	1	0	0	0	1	0	1	1	1	1
Connections Succeed Rate	0	1	1	1	0	1	0	0	0	1	0	1	1	1	1

Figure 354. The Transaction rate objective showing the total IOps with respect to time.

- **iSCSI Initiator-Login Statisitcs:** This view must show that most of the Logins are succesfull. Some logins or re-logins might fail over a period of time as the Device becomes busy in servicing the IO.

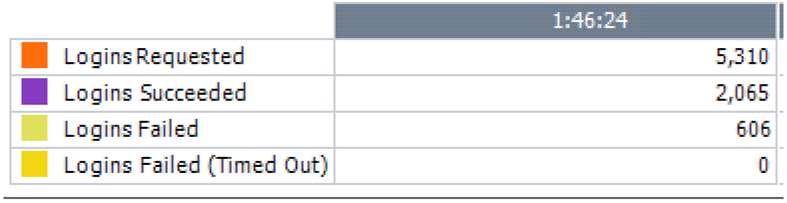

	1:46:24
🟧 Logins Requested	5,310
🟪 Logins Succeeded	2,065
🟨 Logins Failed	606
🟨 Logins Failed (Timed Out)	0

Figure 355. Logins succeded and failed.

- **iSCSI Initiator-Total Throughput:** This displays the total Tx and Rx throughput

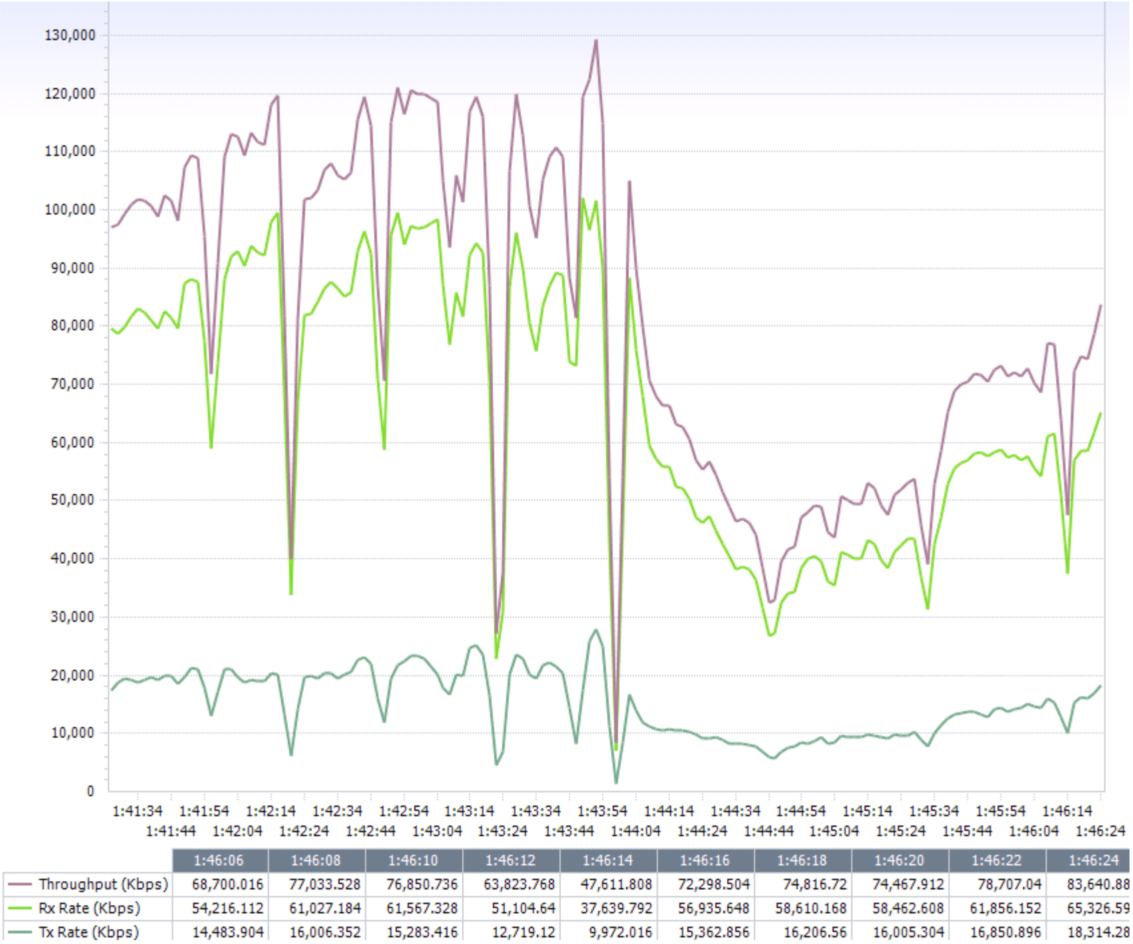

	1:46:06	1:46:08	1:46:10	1:46:12	1:46:14	1:46:16	1:46:18	1:46:20	1:46:22	1:46:24
— Throughput (Kbps)	68,700.016	77,033.528	76,850.736	63,823.768	47,611.808	72,298.504	74,816.72	74,467.912	78,707.04	83,640.88
— Rx Rate (Kbps)	54,216.112	61,027.184	61,567.328	51,104.64	37,639.792	56,935.648	58,610.168	58,462.608	61,856.152	65,326.59
— Tx Rate (Kbps)	14,483.904	16,006.352	15,283.416	12,719.12	9,972.016	15,362.856	16,206.56	16,005.304	16,850.896	18,314.28

Figure 356. Tx and Rx throughput .

- **iSCSI Initiator-Read IO and Write IO:** Check these view to verify the total reads sent, succeded, and failed.

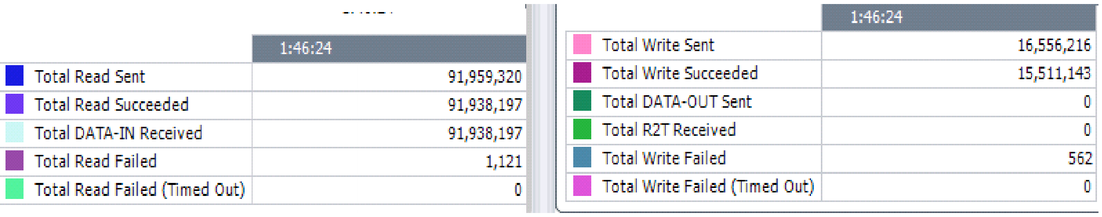

	1:46:24
■ Total Read Sent	91,959,320
■ Total Read Succeeded	91,938,197
■ Total DATA-IN Received	91,938,197
■ Total Read Failed	1,121
■ Total Read Failed (Timed Out)	0

	1:46:24
■ Total Write Sent	16,556,216
■ Total Write Succeeded	15,511,143
■ Total DATA-OUT Sent	0
■ Total R2T Received	0
■ Total Write Failed	562
■ Total Write Failed (Timed Out)	0

Figure 357. Reads and Write ssent/succeeded/failed

- **iSCSI Initiator Command Latency:** This view showcases the average Laetncies of the reads and writes. The Average Latency can actually pin point the element that is slowing the IOPS down. For example, if the write latency is much higher than the read latency (As in this case hightlighed), then this might be the cause of low IO performance.

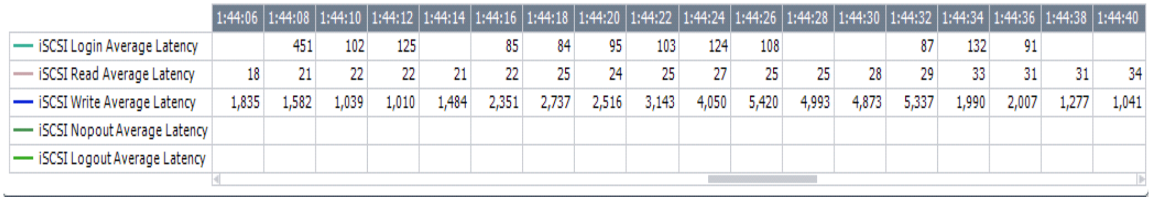

	1:44:06	1:44:08	1:44:10	1:44:12	1:44:14	1:44:16	1:44:18	1:44:20	1:44:22	1:44:24	1:44:26	1:44:28	1:44:30	1:44:32	1:44:34	1:44:36	1:44:38	1:44:40
— iSCSI Login Average Latency		451	102	125		85	84	95	103	124	108			87	132	91		
— iSCSI Read Average Latency	18	21	22	22	21	22	25	24	25	27	25	25	28	29	33	31	31	34
— iSCSI Write Average Latency	1,835	1,582	1,039	1,010	1,484	2,351	2,737	2,516	3,143	4,050	5,420	4,993	4,873	5,337	1,990	2,007	1,277	1,041
— iSCSI Nopout Average Latency																		
— iSCSI Logout Average Latency																		

Figure 358. Average Latencies of Individual commands

- **iSCSI Initiator-Login Redirections:** This shows how many logins were actually successfully redirected to the Target Group. An initiator generally tries to login to the Target Portal and the target portal then redirects it to the right target group.

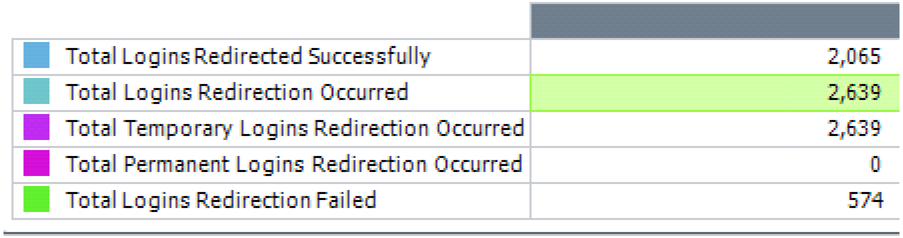

🟦 Total Logins Redirected Successfully	2,065
🟦 Total Logins Redirection Occurred	2,639
🟪 Total Temporary Logins Redirection Occurred	2,639
🟪 Total Permanent Logins Redirection Occurred	0
🟩 Total Logins Redirection Failed	574

Figure 359. Login Redirections.

- **iSCSI Initiator Async events:** This shows the number of Async events like 1,2, or 3 that has occurred at the device. We might see some ASYNC logouts at high loads. **iSCSI Initiator-SNACKS:** This shows how many times the Device has sent a SNACK back to the initiator
- **iSCSI Initiator –Protocol counter:** This shows all the different SCSI protocol errors that could have occurred.
- **iSCSI Initiator-Status Counter:** This shows all the SCSI status that has been recived for the SCSI requests sent from the initiator.

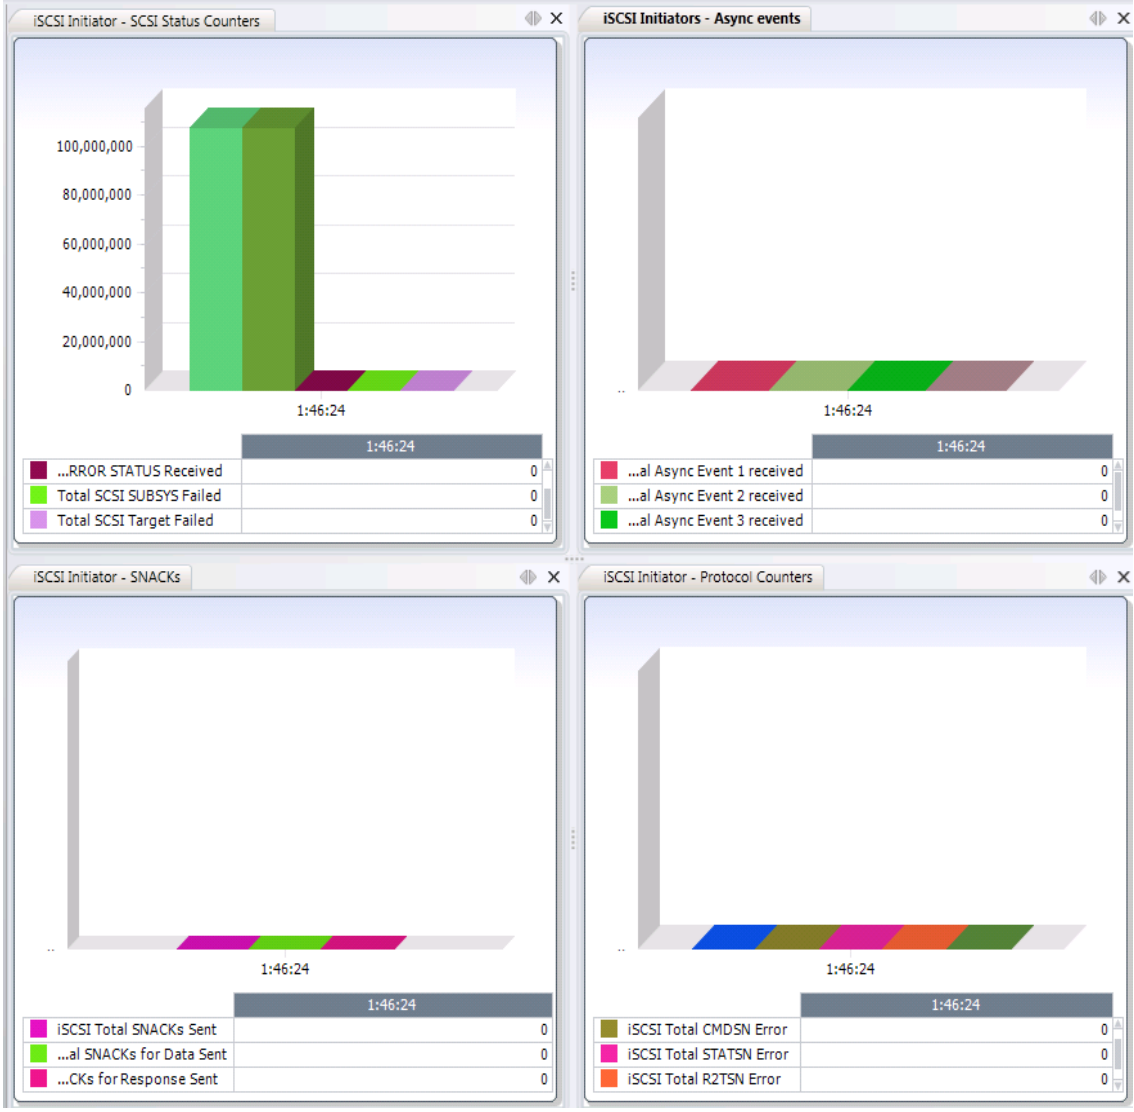
iSCSI Initiator - SCSI Status Counters
100,000,000
80,000,000
60,000,000
40,000,000
20,000,000
0
1:46:24
iSCSI Initiators - Async events
1:46:24
iSCSI Initiator - SNACKs
1:46:24
iSCSI Initiator - Protocol Counters
1:46:24

	1:46:24
...RROR STATUS Received	0
Total SCSI SUBSYS Failed	0
Total SCSI Target Failed	0

	1:46:24
...al Async Event 1 received	0
...al Async Event 2 received	0
...al Async Event 3 received	0

	1:46:24
iSCSI Total SNACKs Sent	0
...al SNACKs for Data Sent	0
...CKs for Response Sent	0

	1:46:24
iSCSI Total CMDSN Error	0
iSCSI Total STATSN Error	0
iSCSI Total R2TSN Error	0

Figure 360. iSCSI protocol related and debug stats

- **iSCSI initiator – Pattern mismatch count:** This actually shows the number of times the SCSI Read command did not find the pattern it was expecting. This stat is only valid if the test was run with **Enable Pattern Matching** checkbox selected.
- **iSCSI Initiator – TCP Connections:** This gives a cumulative count of the total exchanges involved in the TCP hand shakes and the connection closures.

Use the IxLoad Results CSV to verify the following results:

IxLoad results are present in a version specific result folder and the same can be found at “File -> Preferences”. Access the last run test’s result folder and open the xls file iscsi_Client_-_Default_CSV_Logs_iSCSIClient1_1_Traffic1@Network1.csv in the stored location.

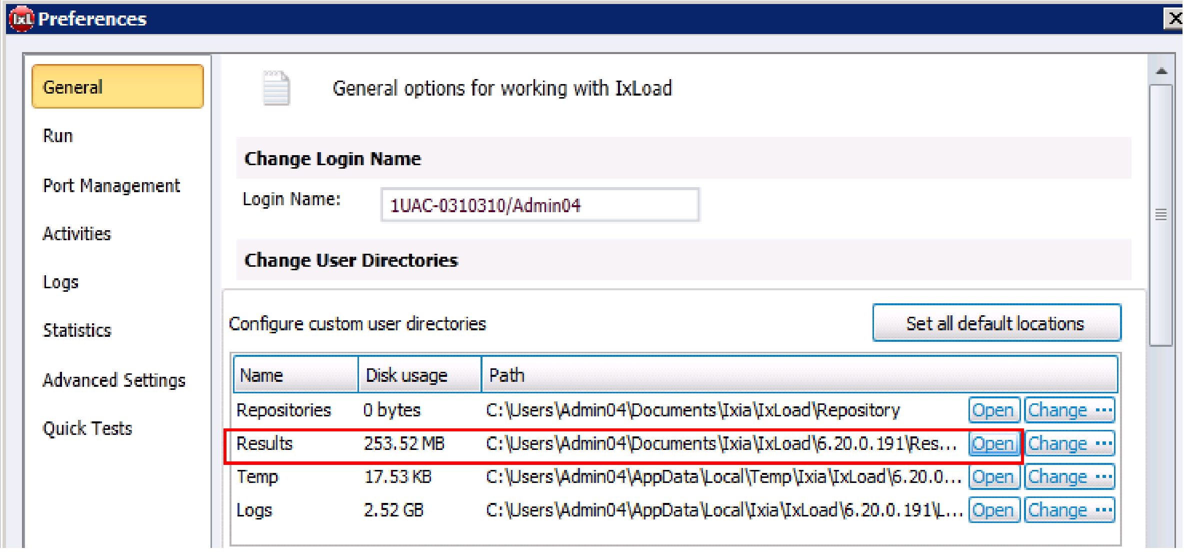

General options for working with IxLoad

Change Login Name

Login Name: 1UAC-0310310/Admin04

Change User Directories

Configure custom user directories

Name	Disk usage	Path		
Repositories	0 bytes	C:\Users\Admin04\Documents\Ixia\IxLoad\Repository	Open	Change ···
Results	253.52 MB	C:\Users\Admin04\Documents\Ixia\IxLoad\6.20.0.191\Res...	Open	Change ···
Temp	17.53 KB	C:\Users\Admin04\AppData\Local\Temp\Ixia\IxLoad\6.20.0...	Open	Change ···
Logs	2.52 GB	C:\Users\Admin04\AppData\Local\Ixia\IxLoad\6.20.0.191\L...	Open	Change ···

Figure 361. IxLoad results folder

- Select the Column **iSCSI Total Read Sent/s** and **iSCSI Total Write Sent/s** and plot a graph of the two column values with respect to time.

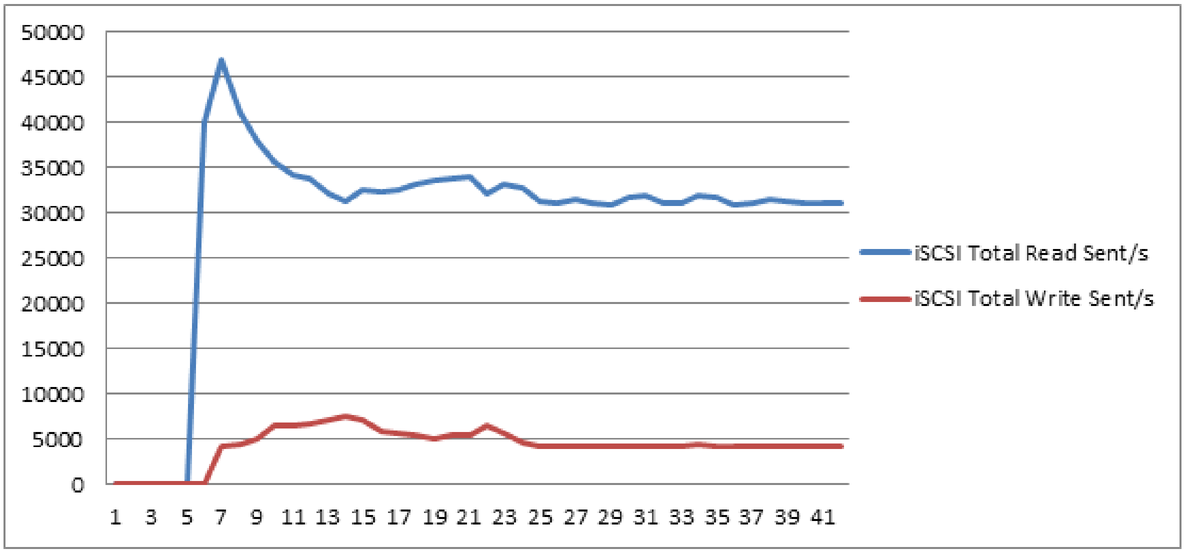

Figure 362. Read and write Sent rate

- Select the Column **iSCSI Total Read Succeded/s** and **iSCSI Total Write Succeded/s** and plot a graph of the two column values with respect to time.

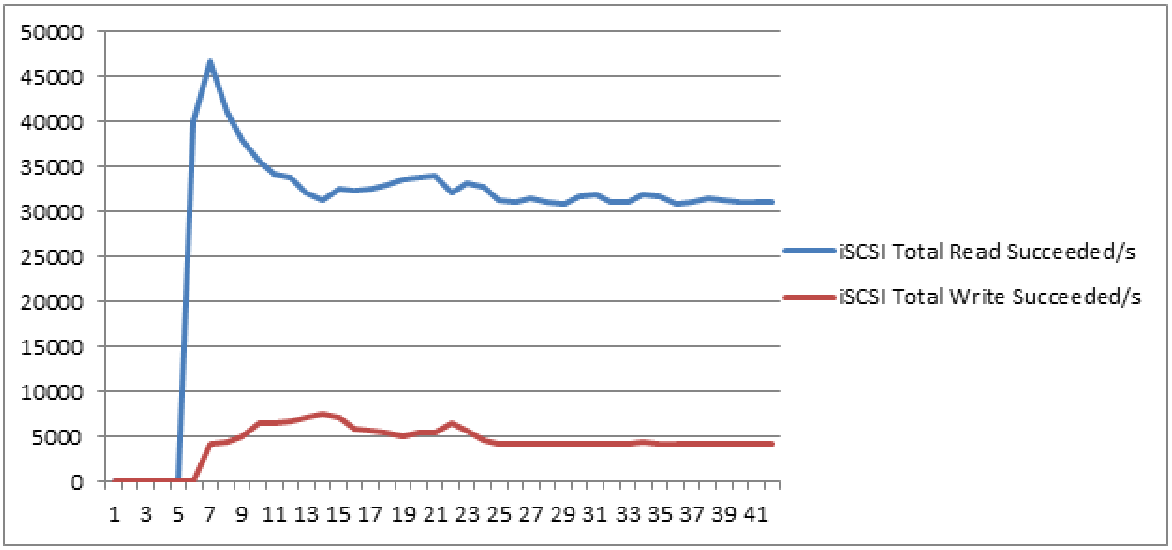

Figure 363. Read and write Succeded rate

- Select the Column **iSCSI Total Write Sent/s** to **Write Succeded/s** and **iSCSI Total Read sent/s** and **Read Succeded/s** and plot a graph of the two column values with respect to time. This displays if some of reads/ writes are getting delayed.

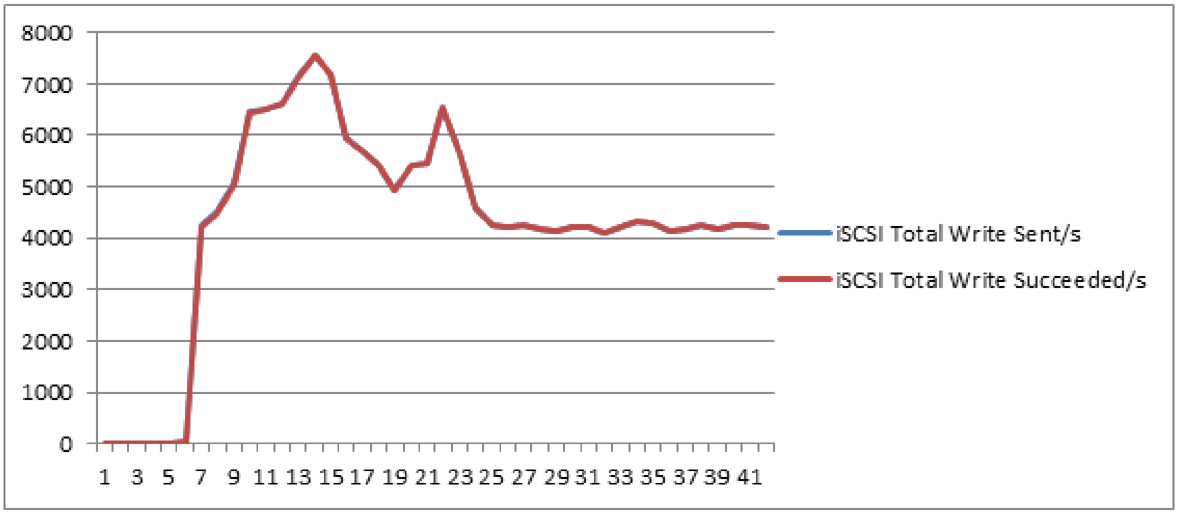

Figure 364. Total Write Sent and succeded per second rate.

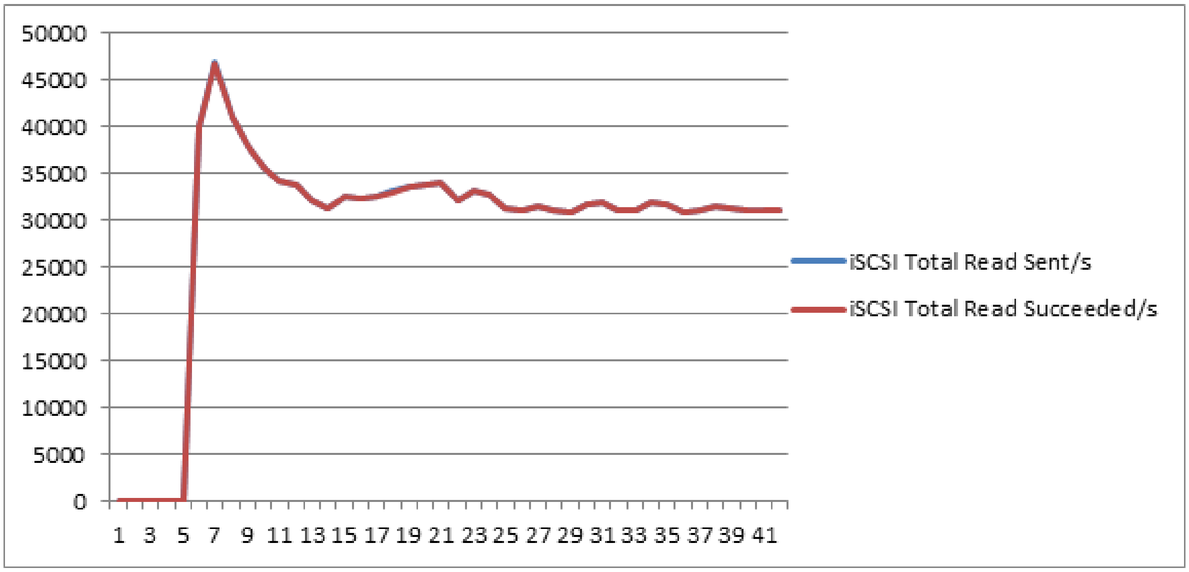

Figure 365. Total Read Sent and succeded per second rate.

- Similar graphs can be obtained in CSV based on the requirements. The other time graphs that can be plot is "**Total Data-IN bytes received/s**" and "**Total Data-Out Bytes Sent/s**"
- If incase there are failures then total reads failed/s or total writes failed/s can also be plotted.
- Total Read Sent to Write Sent or total Read Succeded to Writes Succeded can show as the %age of reads and writes present in the IO.

Conclusions

The iSCSI IO performance test shows the max IO that the device can reach when stressed with realistic IO loads. The emulated scenario is such that multiple users simultaneously try to login to the same device as happens in real life. Varying the different test variables like the block sizes, the sequential or randomness, the read percentages or the MTU/MSS can characterize the IO performances at different conditions or application behaviour and also assess the impact of a particular parameter on the IO performance.

Test Case: Validate Priority Flow Control Performance with Converged Application and Storage Traffic

Overview

Converged data center networks will transport both application and storage traffic over Ethernet infrastructure. However traditional Ethernet is designed to be a best-effort network that may drop packets when the network or devices become busy to handle oversubcribed traffic forwarding requirements. In IP networks, transport reliability has traditionally been the responsibility of the transport protocols, such as the Transmission Control Protocol (TCP), with the trade-off being higher complexity, greater processing overhead and the resulting impact on performance and throughput.

Storage traffic performance in a data center network is generally more critical and important than application LAN Ethernet traffic. One area of evolution for Ethernet is to add extensions to the existing protocol suite to provide reliability without incurring the penalties of TCP. With the move to 10 Gbit/s and faster transmission rates, there is also a desire for higher granularity in control of bandwidth allocation and to ensure it is used more effectively. Beyond the benefits to traditional application traffic, these enhancements would make Ethernet a more viable transport for storage and server cluster traffic.

To meet these goals, standards are being developed to make Ethernet a lossless envrionement. To achieve this environment, data center devices must identify storage traffic (iSCSI or FCoE) from native Ethernet application traffic and be able to pause storage traffic during congestion to ensure storage traffic is not dropped.

Priority-based flow control is one of the IEEE 802.1 protocols designed to enable lossless Ethernet. Priority-based flow control is built on the concept of the original IEEE 802.3x Flow Control, modified to operate with traffic class differentiation. When a receiving station's buffer is near exhaustion, 802.3x Flow Control allows the receiving station to request its upstream neighbor to pause transmission on the entire port, giving the receiving station an opportunity to clear its buffer. Priority-based flow control elevates this critical capability to a higher level by providing the receiving station the ability to request its upstream neighbor to pause transmission on one or more priorities (essentially virtual lanes). To achieve this, the original 802.3x flow control PAUSE frame is modified so that pause_quanta can be signaled at a per priority value, and the MAC CONTROL sublayer is enhanced with the capability to assign and throttle transmission queues on a per priority value.

This test focuses on forwarding stateful storage and native application Ethernet traffic from initiators and clients to targets and servers, with priority-based flow control enabled.

Objective

The objective for this test is to ensure that Storage traffic is prioritized over the LAN application traffic. IxLoad can simulate stateful Storage and LAN traffic through a device or terminating a strorage device. In this test case, we are validating a storage switch (DUT) and the switch interfaces DCB values are set such that they enable priority flow control and ETS on the Storage traffic, thus making it lossless infastrucutre like the FC. In addition, the following points should also be observed and assessed:

1. During congestion, pause frames are sent only on those traffic classes on which Priority Flow control is enabled.
2. The priority percentage set using the Priority Groups is respected.
3. The LAN traffic is able to take the additional bandwidth, if the storage traffic has freed up some bandwidth.
4. The TCP failures and retries should be considerably less on the storage traffic.

Setup

The setup requires three Ixia ports (two acting as client and one acting as server) and three DCB ports of the switch.

Ixia port 1(**Client/Initiator**) generating the priority 6 storage and untagged LAN traffic connected to the 10G interface of the DCB Switch

Ixia port 2 (**Client/Initiator**) generating the priority 7 storage and untagged LAN traffic connected to a 10G interface of the DCB switch

Ixia port 2(**Server/Target**) terminating both the storage and the untagged application traffic on the LAN

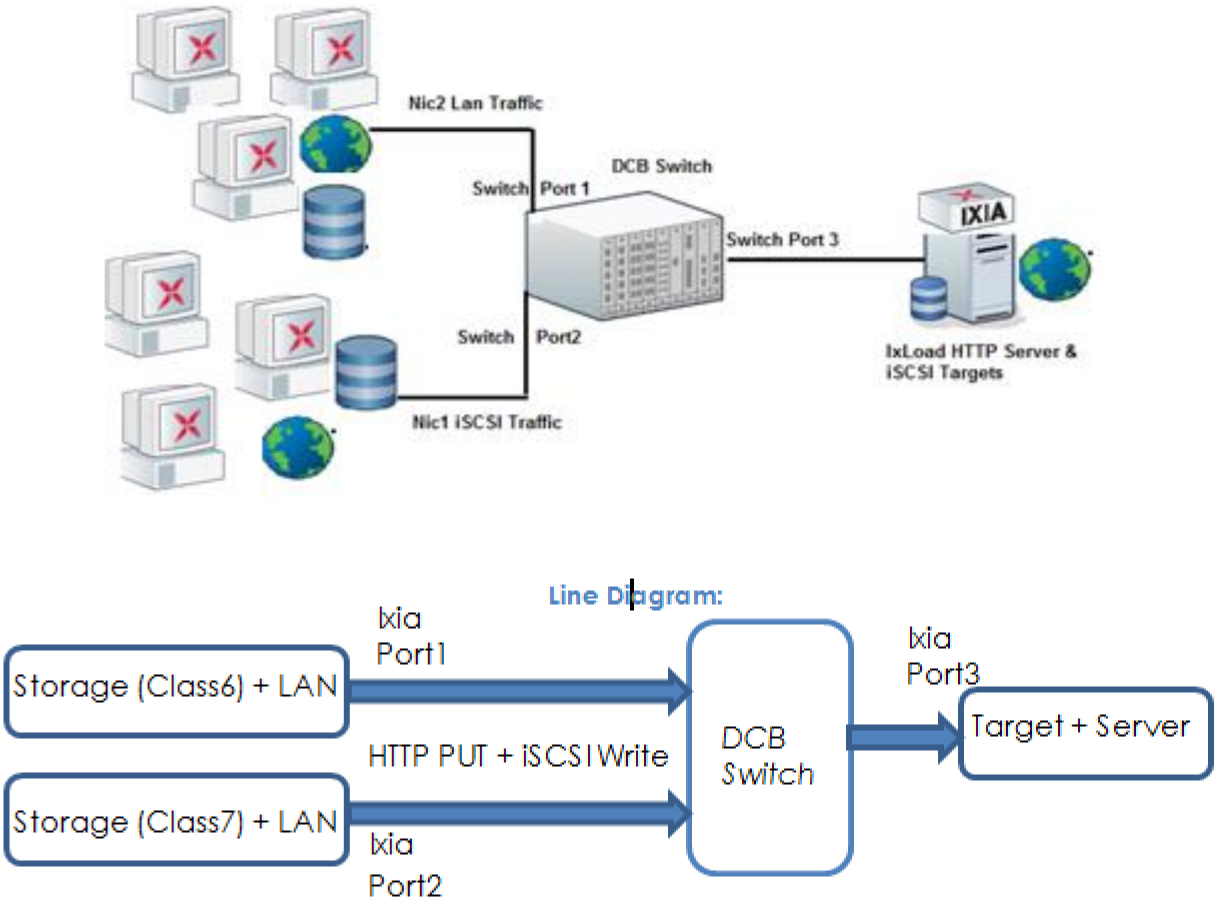

Figure 366. Test setup with Ixia iSCSI initiators, LAN Clients and Ixia iSCSI targets and LAN servers with the DCB switch in the middle.

Two 10Gig Ixia Links are pushing Tx Traffic (HTTP Put and iSCSI Write) to the receiver that has only one 10Gig link to receive.

Step-by-Step Instructions

1. Start IxLoad. In the main window, the **Scenario Editor** window appears. All test configurations are performed here.

2. Add a NetTraffic at the **Originate** and the **Terminate** side. The **Originate** Side serves as the traffic generator for the iSCSI Initiators and HTTP Clients and the **Terminate** Side acts as the HTTP Server and iSCSI Target.

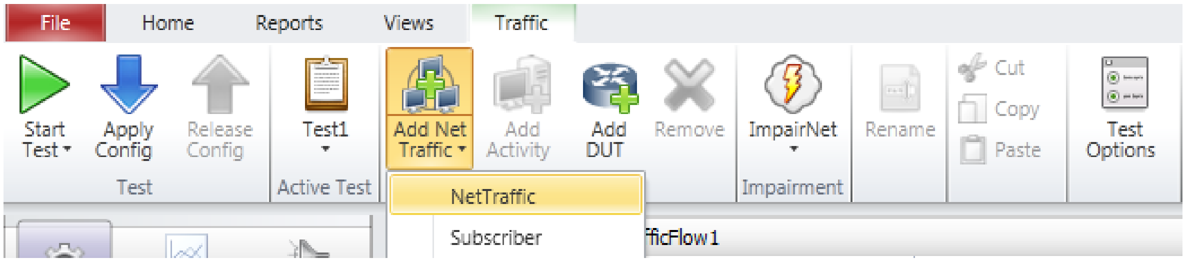

Figure 367. IxLoad NetTraffic addition

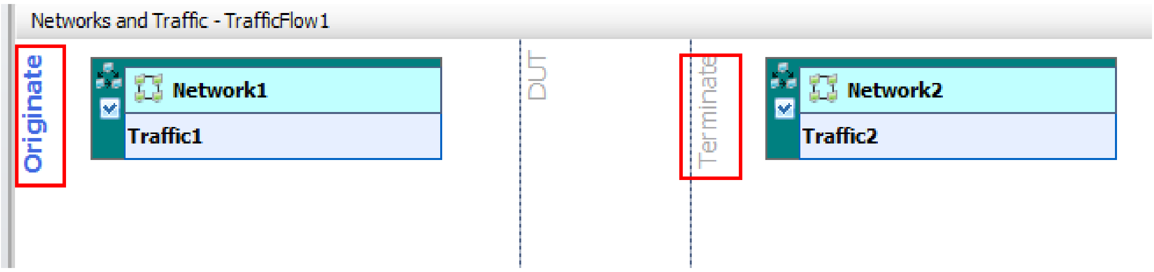

Figure 368. Adding NetTraffic at both Originate and Terminate side.

3. Select the Terminate side Network **Network2.** It displays the default IP and the MAC/VLAN stack that is already added. Change the IP count from 100 to 1. Ixia emulated iSCSI target and HTTP server islistening on this IP.

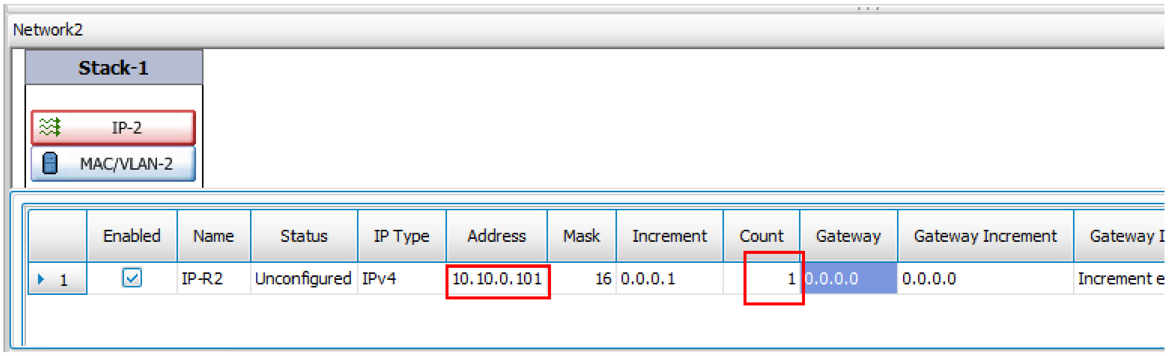

	Enabled	Name	Status	IP Type	Address	Mask	Increment	Count	Gateway	Gateway Increment	Gateway I
1	☑	IP-R2	Unconfigured	IPv4	10.10.0.101	16	0.0.0.1	1	0.0.0.0	0.0.0.0	Increment e

Figure 369. IP configurations at Network2

4. The Strorage switch (DUT) is set as trunk with VLAN ID 2 added to it. Hence configure the clients and server within the same VLAN. Select the **MAC/VLAN** stack and enable **VLAN.** Set the First ID as 2 and the **Increment By** to 0.

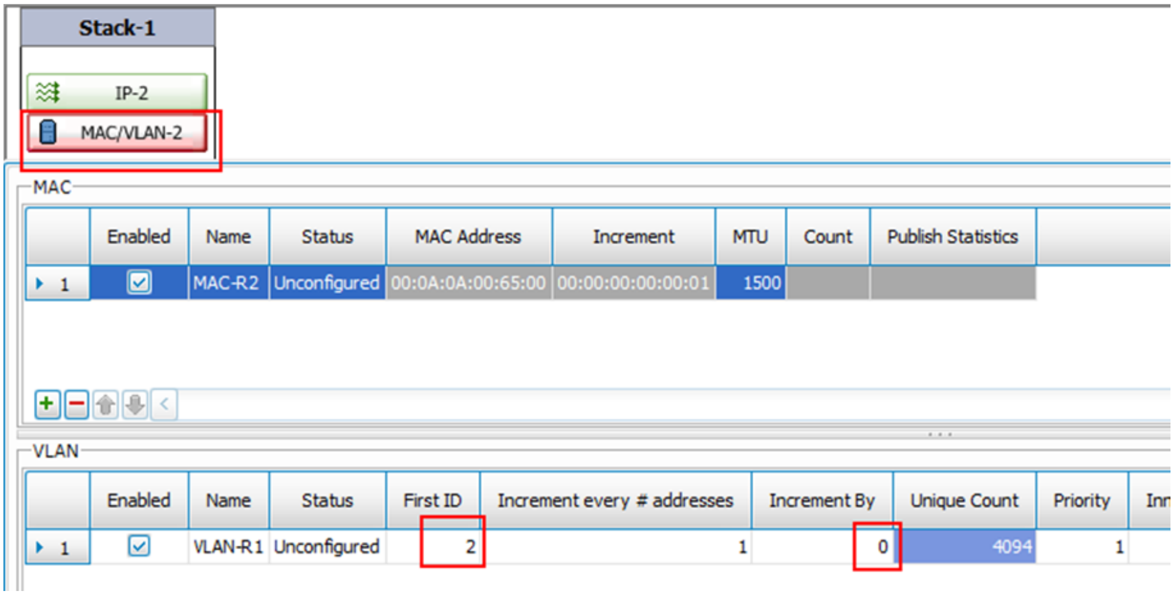

MAC

	Enabled	Name	Status	MAC Address	Increment	MTU	Count	Publish Statistics
1	☑	MAC-R2	Unconfigured	00:0A:0A:00:65:00	00:00:00:00:00:01	1500		

VLAN

	Enabled	Name	Status	First ID	Increment every # addresses	Increment By	Unique Count	Priority	Inr
1	☑	VLAN-R1	Unconfigured	2	1	0	4094	1	

Figure 370. MAC/VLAN configurations at the Server side

5. Add a DCBx stack to the interface emulating the iSCSI target and HTTP server, so that the correct TLVs are exchanged to negotiate the priority. To add the DCBx stack, select the **MAC/VLAN** stack, right-click on it and select **Add above -> DCBX** to add a new stack.

6. After adding, configure the stacks properly to have the right values under each TLV. While on DCBx stack setting, select the **DCBx TLVs** tab, and then click the **IEEE1.01 TLVs** tab at the bottom, because the DUT supports the IEEE version of the data center bridging speficication. By default, the Priority Group TLV is added for each IP range. [Note: For this scenario, we are only sending majorityof the traffic (HTTP PUT and iSCSI Writes) towards the target. Hence the Terminate side DCBx session is not mandatory. We can even have plain IP at the Terminate side. We are adding DCB at both client and target side for consistency].

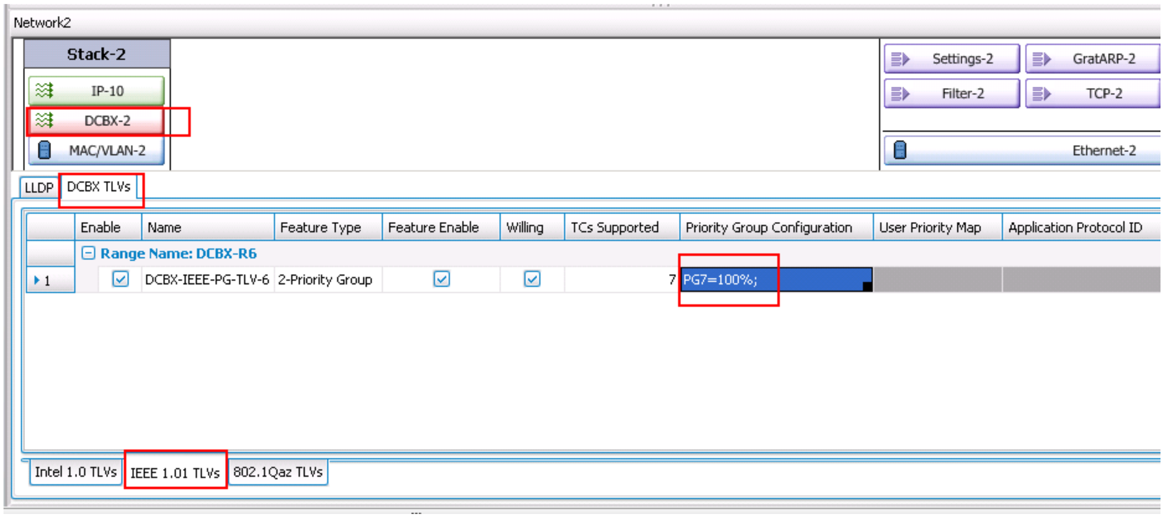

	Enable	Name	Feature Type	Feature Enable	Willing	TCs Supported	Priority Group Configuration	User Priority Map	Application Protocol ID
	Range Name: DCBX-R6								
▸ 1	☑	DCBX-IEEE-PG-TLV-6	2-Priority Group	☑	☑	7	PG7=100%;		

Figure 371. Selecting the IEEE 1.01 TLV

7. Add one more TLV using the **Add Range** button on the menu bar**.** The new TLVs automatically possess feature type as **PFC**.[Note: IxLoad also supports Application TLV, but this particular switch’s firmware version does not recongnise iSCSI TLV].

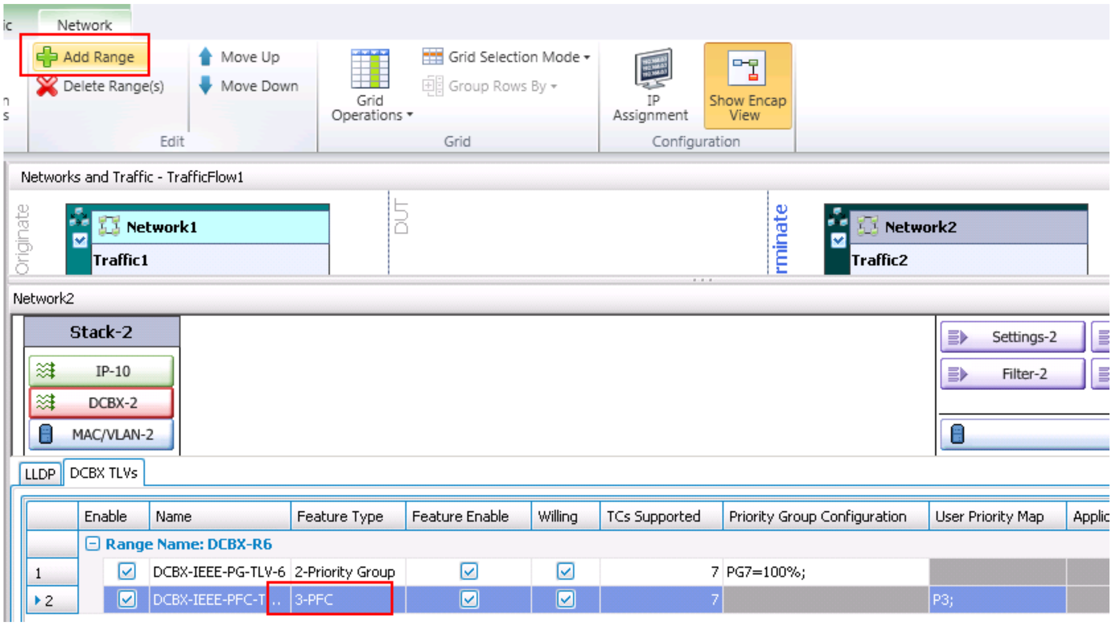

	Enable	Name	Feature Type	Feature Enable	Willing	TCs Supported	Priority Group Configuration	User Priority Map	Applic
	Range Name: DCBX-R6								
1	☑	DCBX-IEEE-PG-TLV-6	2-Priority Group	☑	☑	7	PG7=100%;		
▸ 2	☑	DCBX-IEEE-PFC-T..	3-PFC	☑	☑	7		P3;	

Figure 372. Adding the PFC TLV

8. This DCBx session is for the Rx traffic (From Server to Client) from the Targets back to the client.
   a. In all TLVs ensure the **Willing Flag** checkbox is selected.. This selection ensures that the Ixia ports take the switche’s configurations. [Note: Ixia ports are willing the setting Priority Group percentage is not mandatory.]
   b. Edit the **PFC TLV** to turn on PFC on both Priority **6** and **7** in the **User Priority Map** column.

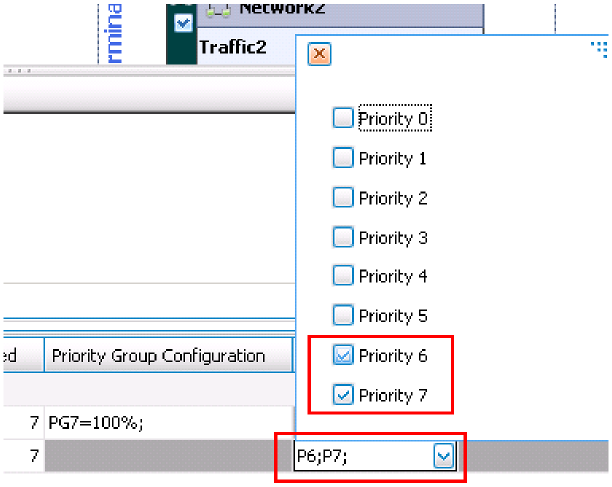

Network2

Traffic2

☐ Priority 0
☐ Priority 1
☐ Priority 2
☐ Priority 3
☐ Priority 4
☐ Priority 5
☑ Priority 6
☑ Priority 7

ed	Priority Group Configuration
7	PG7=100%;
7	P6;P7;

Figure 373. Editing the PFC TLV

c. Change the Priority Group percentages to give 40% to class 6 and 40% to Class 7 and the remaing portion to be untagged.

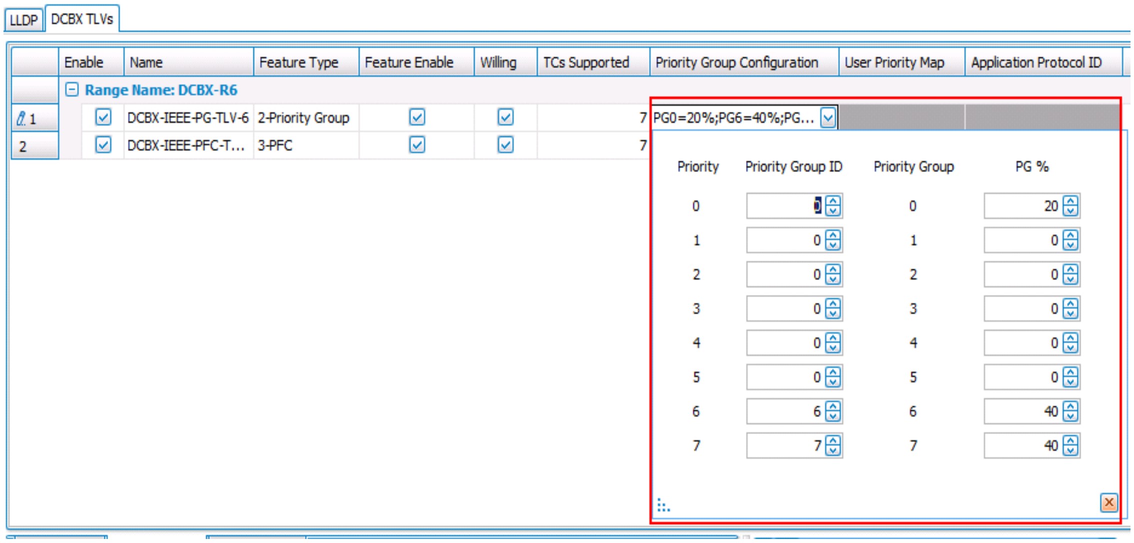

LLDP | DCBX TLVs

	Enable	Name	Feature Type	Feature Enable	Willing	TCs Supported	Priority Group Configuration	User Priority Map	Application Protocol ID
	Range Name: DCBX-R6								
1	☑	DCBX-IEEE-PG-TLV-6	2-Priority Group	☑	☑	7	PG0=20%;PG6=40%;PG...		
2	☑	DCBX-IEEE-PFC-T...	3-PFC	☑	☑	7			

Priority	Priority Group ID	Priority Group	PG %
0	0	0	20
1	0	1	0
2	0	2	0
3	0	3	0
4	0	4	0
5	0	5	0
6	6	6	40
7	7	7	40

Figure 374. Editing the Priority Group TLV

9. Click the '+' button available in **Traffic2** to add the HTTP Server and iSCSI Server (Target) at the Terminate side NetTraffic.

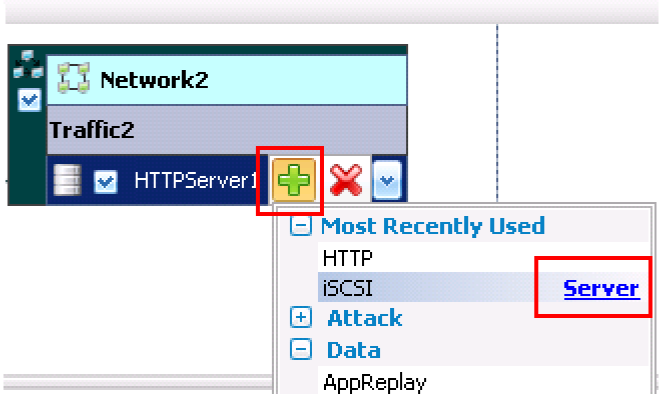

Figure 375. Adding the HTTP and iSCSI servers at the "Terminate" Side.

10. Click the **iSCSIServer1**. Click the **Shared Target Pool** tab and set the LUN capacity to **100000 MB**. This makes the LUN big enough to handle large reads.

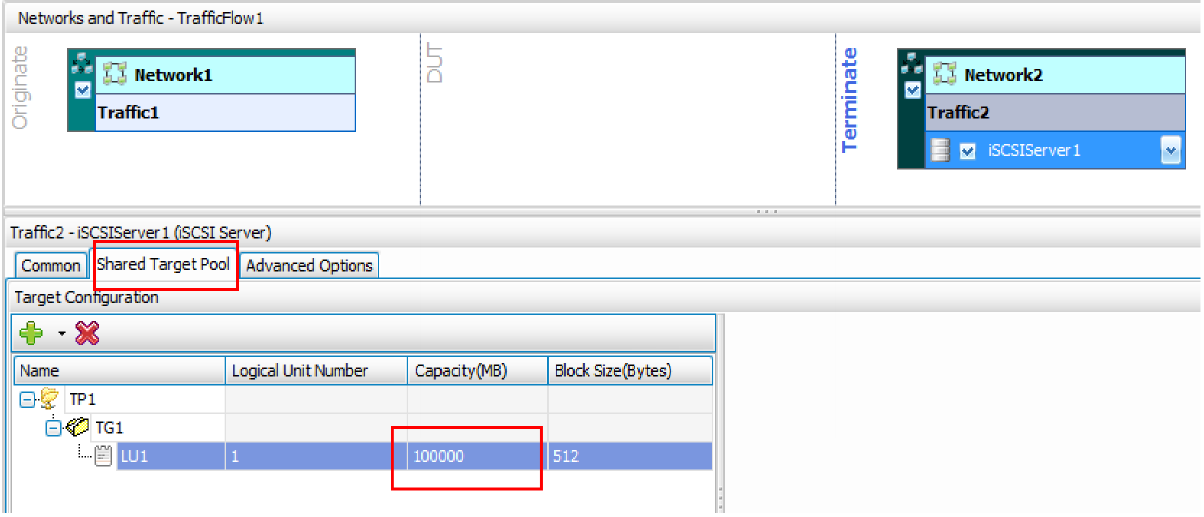

Name	Logical Unit Number	Capacity(MB)	Block Size(Bytes)
TP1			
TG1			
LU1	1	100000	512

Figure 376. Setting the Target's LU length to 100000 MB.

11. Still in the iSCSI Server configuration, click the **Advanced Options** tab and set the priority as 7. This setting ensures that all the Tx traffic from the Ixia server is marked with the COS bit of **7**. [Note: This test does not deal with a lot of Tx traffic from the server, hence it is expected that the Switch port does not send any pause frames to the Ixia server].

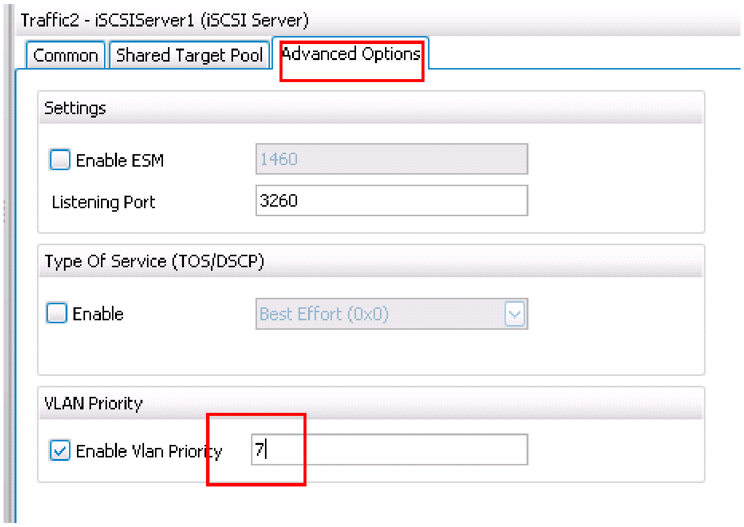

Traffic2 - iSCSIServer1 (iSCSI Server)

Settings

☐ Enable ESM 1460

Listening Port 3260

Type Of Service (TOS/DSCP)

☐ Enable Best Effort (0x0)

VLAN Priority

☑ Enable Vlan Priority 7

Figure 377. Setting the Vlan priority of traffic from the iSCSI target to 7

12. This concludes the Server/Target side configurations. Let us shift to the Originate side to complete the HTTP Client and iSCSI Initator configurations.
 Click the Originate **Network1** to display the default IP and VLAN stack. Add another IP range to this. The second IP range carries traffic directed towards the switch port 2, whereas the first range one carries traffic for the switch port 1.

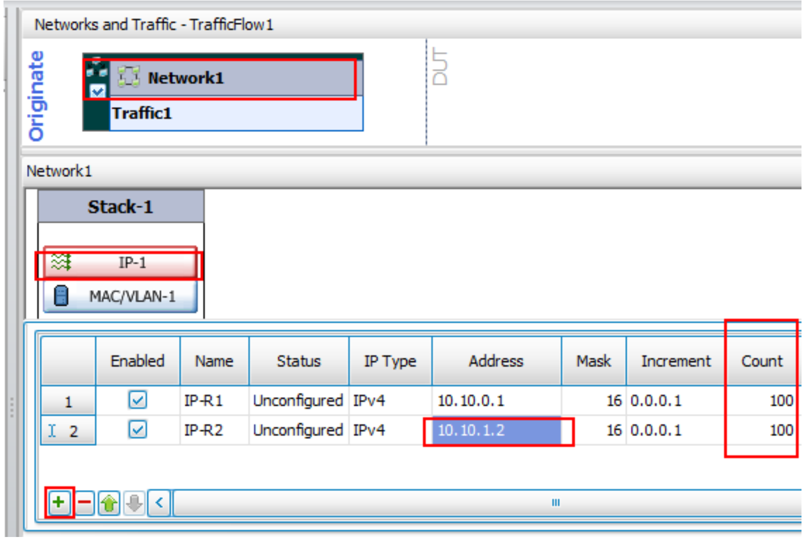

Networks and Traffic - TrafficFlow1

	Enabled	Name	Status	IP Type	Address	Mask	Increment	Count
1	☑	IP-R1	Unconfigured	IPv4	10.10.0.1	16	0.0.0.1	100
2	☑	IP-R2	Unconfigured	IPv4	10.10.1.2	16	0.0.0.1	100

Figure 378. Adding the second IP range at the Originate side

13. Click the **MAC/VLAN** stack and enable **VLAN**. Set the VLAN as **2**. The Switch interface is configured as trunk and VLAN ID 2 is allowed on this trunk.

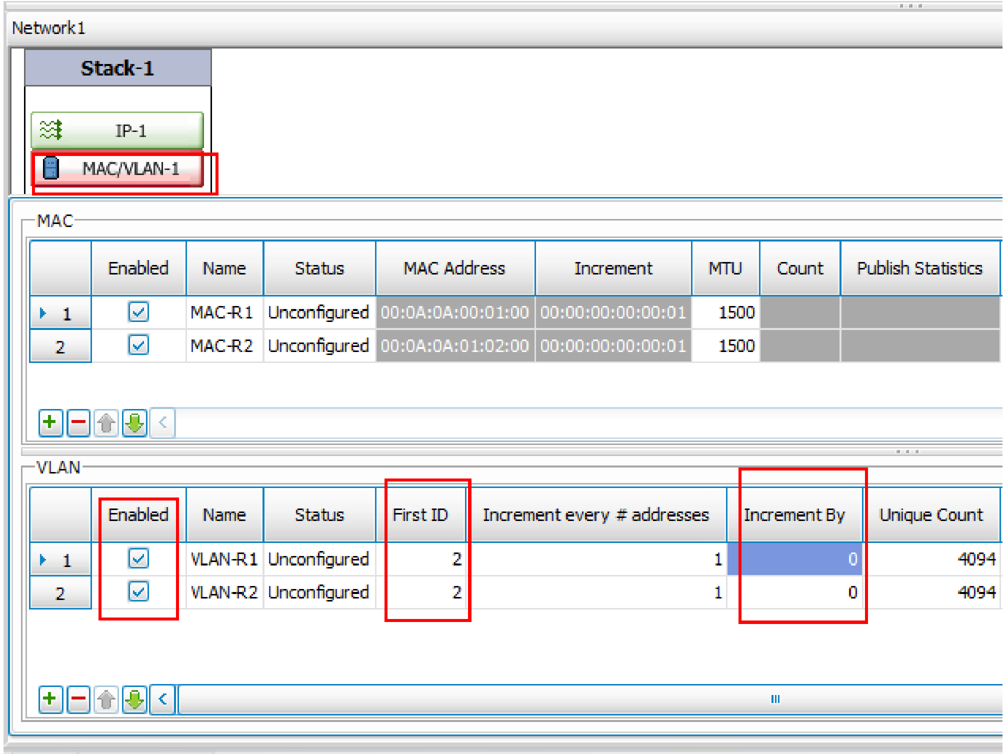

MAC

	Enabled	Name	Status	MAC Address	Increment	MTU	Count	Publish Statistics
1	☑	MAC-R1	Unconfigured	00:0A:0A:00:01:00	00:00:00:00:00:01	1500		
2	☑	MAC-R2	Unconfigured	00:0A:0A:01:02:00	00:00:00:00:00:01	1500		

VLAN

	Enabled	Name	Status	First ID	Increment every # addresses	Increment By	Unique Count
1	☑	VLAN-R1	Unconfigured	2	1	0	4094
2	☑	VLAN-R2	Unconfigured	2	1	0	4094

Figure 379. Set the vlan id of both the ranges to 2.

14. Add a DCBx stack on the **Originate** side emulating the HTTP Clients and iSCSI initiators so that the correct TLVs are exchanged to negotiate the priority. To add the DCBx stack, click the **MAC/VLAN** stack, right click on it and select **Add above -> DCBX** to add a new stack.

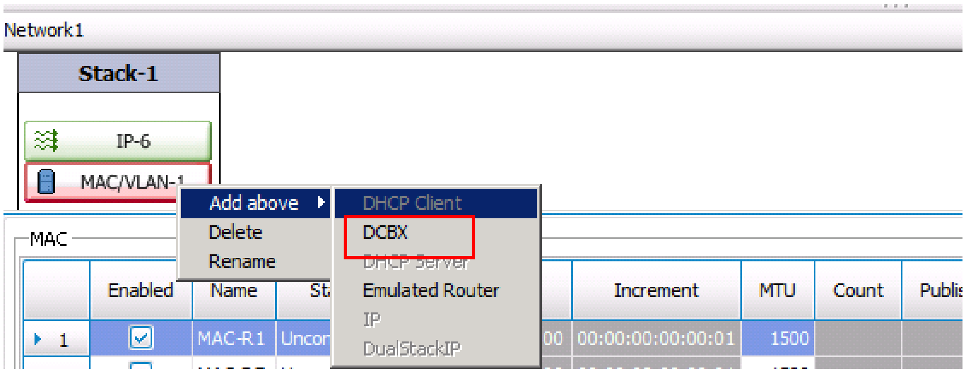

	Enabled	Name	St	Increment	MTU	Count	Publis
1	☑	MAC-R1	Uncon	00:00:00:00:00:01	1500		

Figure 380. Adding the DCBx stack on top of the MAC/VLAN

15. After adding, configure the stack properly to have the right values under each TLV. While on DCBx stack settings click the **DCBx TLVs** tab, and then click the **IEEE1.01 TLVs** tab on the bottom. By default the Priority Group TLV will be already added for each range.

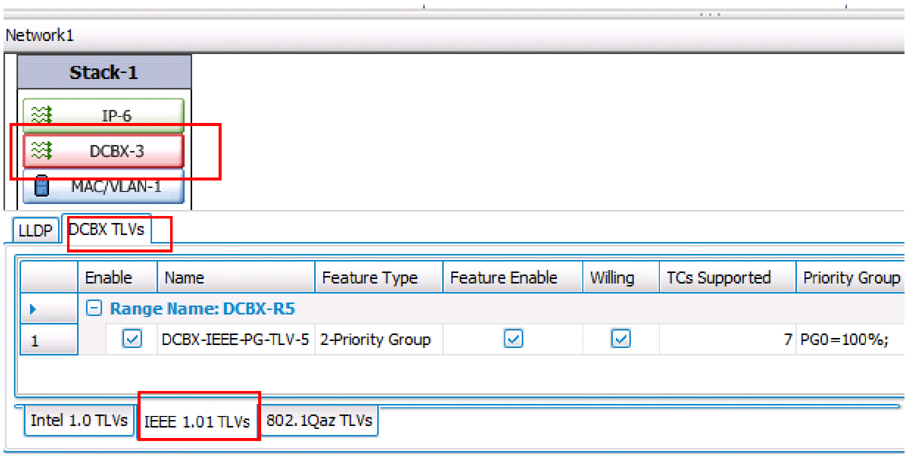

	Enable	Name	Feature Type	Feature Enable	Willing	TCs Supported	Priority Group
	Range Name: DCBX-R5						
1	☑	DCBX-IEEE-PG-TLV-5	2-Priority Group	☑	☑	7	PG0=100%;

Figure 381. Select the 1EEE 1.01 TLV under the DCBx TLVs

16. Add one more TLV using the **Add Range** button on the menu bar. The new TLVs will be automatically be of type **PFC**. If not change the feature type to PFC.
[Note: IxLoad also supports Application TLV but this particular switch’s firmware version does not recongnise iSCSI TLV].

17. Add the **PFC TLV** for both the ranges. All the TLV’s will have the default values.

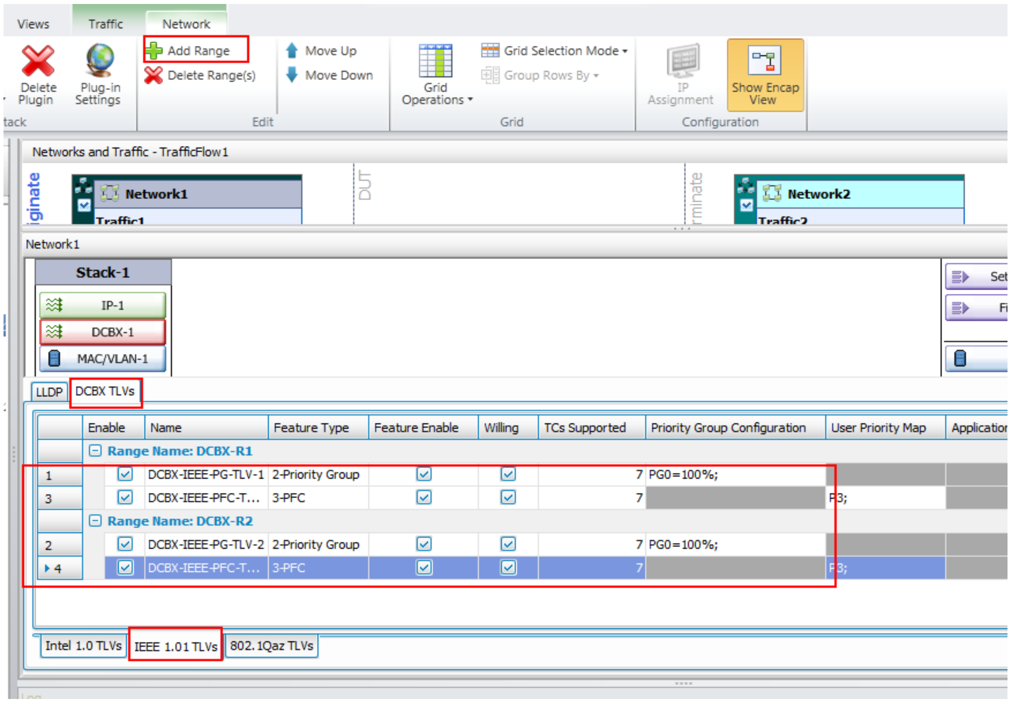

	Enable	Name	Feature Type	Feature Enable	Willing	TCs Supported	Priority Group Configuration	User Priority Map	Applicatio
	Range Name: DCBX-R1								
1	☑	DCBX-IEEE-PG-TLV-1	2-Priority Group	☑	☑	7	PG0=100%;		
3	☑	DCBX-IEEE-PFC-T...	3-PFC	☑	☑	7		P3;	
	Range Name: DCBX-R2								
2	☑	DCBX-IEEE-PG-TLV-2	2-Priority Group	☑	☑	7	PG0=100%;		
4	☑	DCBX-IEEE-PFC-T...	3-PFC	☑	☑	7		P3;	

Figure 382. PFC and Priority Group TLV’s already added with default priority values.

18. Now both the DCBx ranges will take with them a combination of iSCSI and HTTP (HTTP is emulating application traffic on the LAN) traffic types. Since the switch does not support iSCSI TLV we will add VLAN priorities to the storage traffics to have them in different traffic classes. The traffic through Ixia port 1 will carry storage traffic with VLAN priority 6 and traffic through Ixia port 2 will carry storage traffic with VLAN priority 7. Both port 1 and port 2 will also carry some LAN traffic.
    a. In all the TLVs ensure the **Willing Flag** is. This selection ensures that the Ixia ports use the switche’s configuration. [Note: Setting Priority Group percentage is not absolutely mandatory, because Ixia ports are willing.
    b. Edit the **PFC TLV** of first range to set it to **Priority 6** in the **User Priority Map** column.

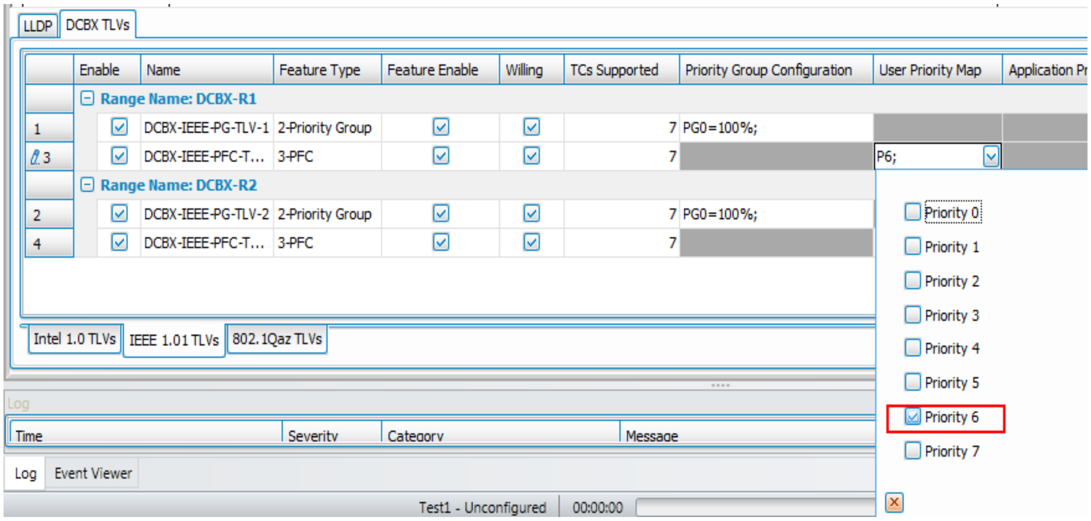

	Enable	Name	Feature Type	Feature Enable	Willing	TCs Supported	Priority Group Configuration	User Priority Map	Application Pr
	⊟ Range Name: DCBX-R1								
1	☑	DCBX-IEEE-PG-TLV-1	2-Priority Group	☑	☑	7	PG0=100%;		
3	☑	DCBX-IEEE-PFC-T...	3-PFC	☑	☑	7		P6;	
	⊟ Range Name: DCBX-R2								
2	☑	DCBX-IEEE-PG-TLV-2	2-Priority Group	☑	☑	7	PG0=100%;		
4	☑	DCBX-IEEE-PFC-T...	3-PFC	☑	☑	7			

□ Priority 0
□ Priority 1
□ Priority 2
□ Priority 3
□ Priority 4
□ Priority 5
☑ Priority 6
□ Priority 7

Test1 - Unconfigured 00:00:00

Figure 383. Setting priority of the First range

c. Edit the **PFC TLV** of the second range and set it to **Priority** 7 in the **User Priority Map** column.
d. Change the Priority Group (PG) percentages to give 40% to class 6 and 40% to Class 7 and the remaining 20% as untagged.

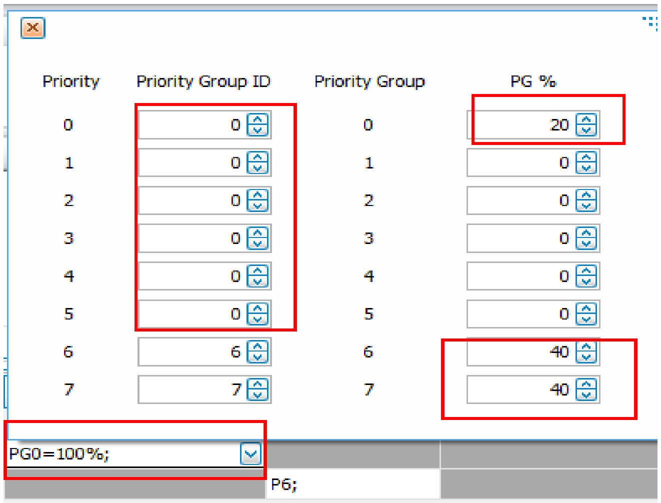

Priority	Priority Group ID	Priority Group	PG %
0	0	0	20
1	0	1	0
2	0	2	0
3	0	3	0
4	0	4	0
5	0	5	0
6	6	6	40
7	7	7	40

PG0=100%;

P6;

Figure 384. Setting the priority group of the first range.

e. Perform the same actions for the second range. After completion, the DCBX range appears as follows:

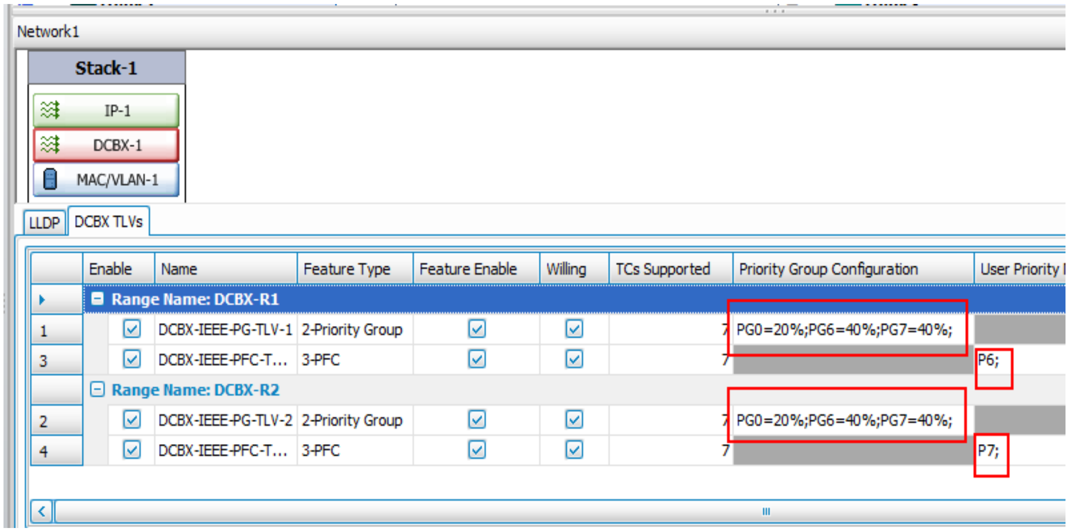

	Enable	Name	Feature Type	Feature Enable	Willing	TCs Supported	Priority Group Configuration	User Priority I
	Range Name: DCBX-R1							
1	☑	DCBX-IEEE-PG-TLV-1	2-Priority Group	☑	☑	7	PG0=20%;PG6=40%;PG7=40%;	
3	☑	DCBX-IEEE-PFC-T...	3-PFC	☑	☑	7		P6;
	Range Name: DCBX-R2							
2	☑	DCBX-IEEE-PG-TLV-2	2-Priority Group	☑	☑	7	PG0=20%;PG6=40%;PG7=40%;	
4	☑	DCBX-IEEE-PFC-T...	3-PFC	☑	☑	7		P7;

Figure 385. The finalized configurations on both the ranges.

This configuration is similar to the configuration on Switches PFC map that are applied on these interfaces:

The switch PFC map on all the interfaces does the following functions:

- Enabled PFC on class 6 and allocate 40% of the total bandwidth
- Enabled PFC on class 7 and allocate 40% of the total bandwidth
- Do not enable PFC on class 0(untagged) and allocate 20% of the total bandwidth.

19. Configuring the **HTTP Client** traffic.

a) Add an HTTP Client at the Originate Side.

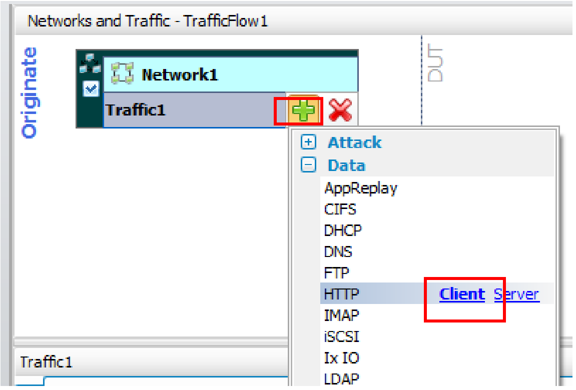

Figure 386. Adding the HTTP client.

b) Click the newly added HTTP client. Add a **PUT** command in the **HTTP** client. This PUT helps to generate a high volume of Tx traffic on the Switch interface. This action is needed to simulate congestion behaviour.

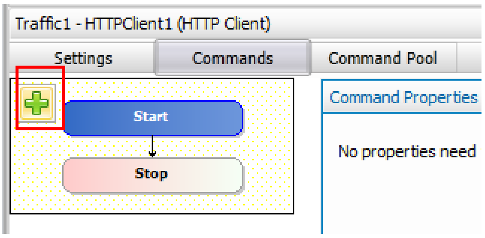

No properties need

Figure 387. Adding the put activity.

c) In the destination drop down, select the HTTP server that is already added at the Terminate side.

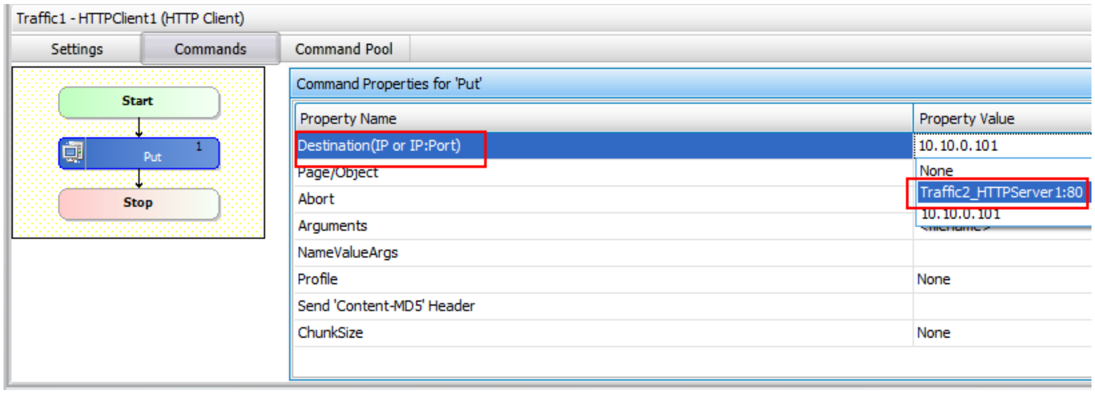

Command Properties for 'Put'

Property Name	Property Value
Destination(IP or IP:Port)	10.10.0.101 None Traffic2_HTTPServer1:80 10.10.0.101 <filename>
Page/Object	
Abort	
Arguments	
NameValueArgs	
Profile	None
Send 'Content-MD5' Header	
ChunkSize	None

Figure 388. Set the HTTP server as the Destination of the HTTP client.

d) In the **Arguments** of the **PUT** command, add any file that is of size 512KB or more. This file is sent as a PUT argument and generates high volume of Tx traffic.

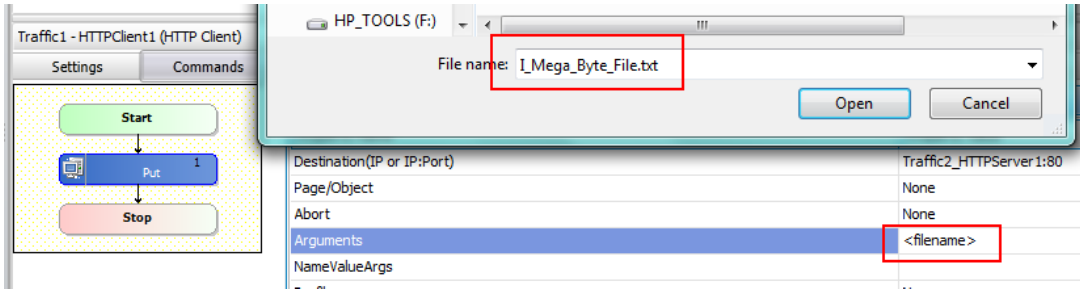

File name: I_Mega_Byte_File.txt

Destination(IP or IP:Port)	Traffic2_HTTPServer1:80
Page/Object	None
Abort	None
Arguments	<filename>
NameValueArgs	

Figure 389. Add a file as argument to the PUT command.

20. Configuring the iSCSI Initiator.

a. Add the **iSCSI Client** to the same **Originate** site NetTraffic.

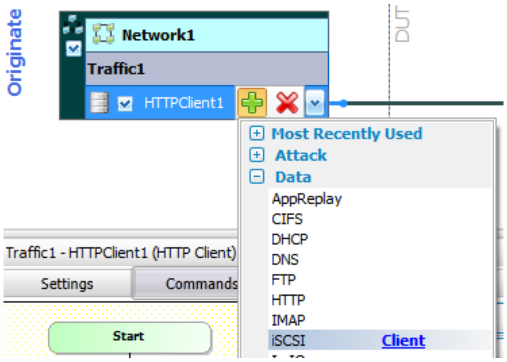

Figure 390. Adding the iSCSI client.

b. Select the iSCSIClient activity and click the ➕ button to add the **Login, IO, and Logout** commands in the iSCSI client activity.

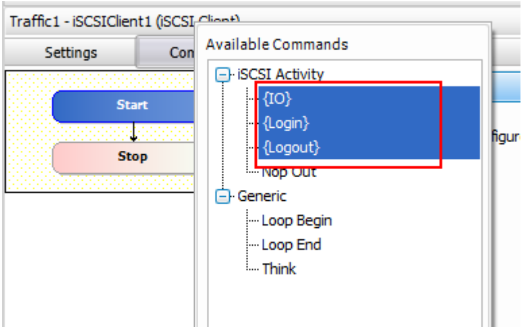

Figure 391. Adding the Login, IO and Logout command.

c. In the **Login** command select the destination as the already added iSCSI server and set the **Target Name** as **TG1**. This will select the default target that has already been added in the Ixia emulate iSCSI target.

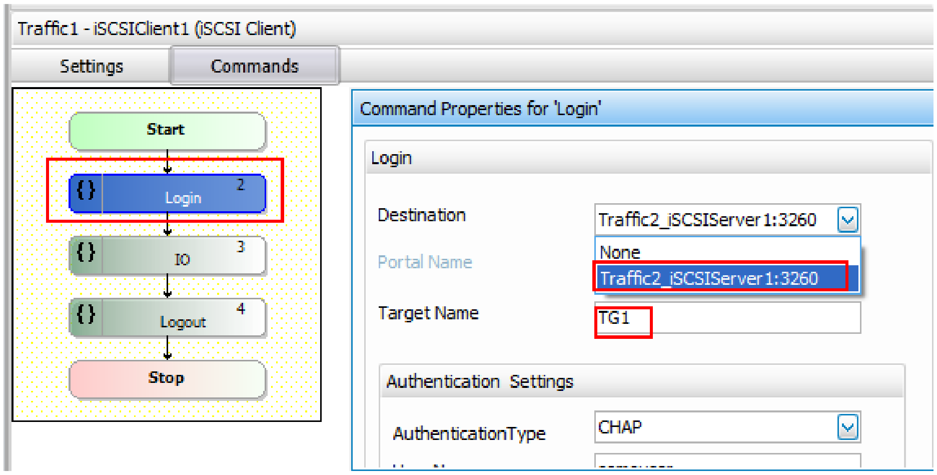

Destination: Traffic2_iSCSIServer1:3260

Target Name: TG1

AuthenticationType: CHAP

Figure 392. Selecting the Ixia iSCSI server as the target and setting the target name as TG1.

d. To configure the **{IO}** command, set the **Target LUN ID** as **1** and the **Command Type** as **Write**. The LUN ID of 1 is the default LUN that is added in the Ixia emulated iSCSI target in TG1. The write command generates the storage Tx traffic.

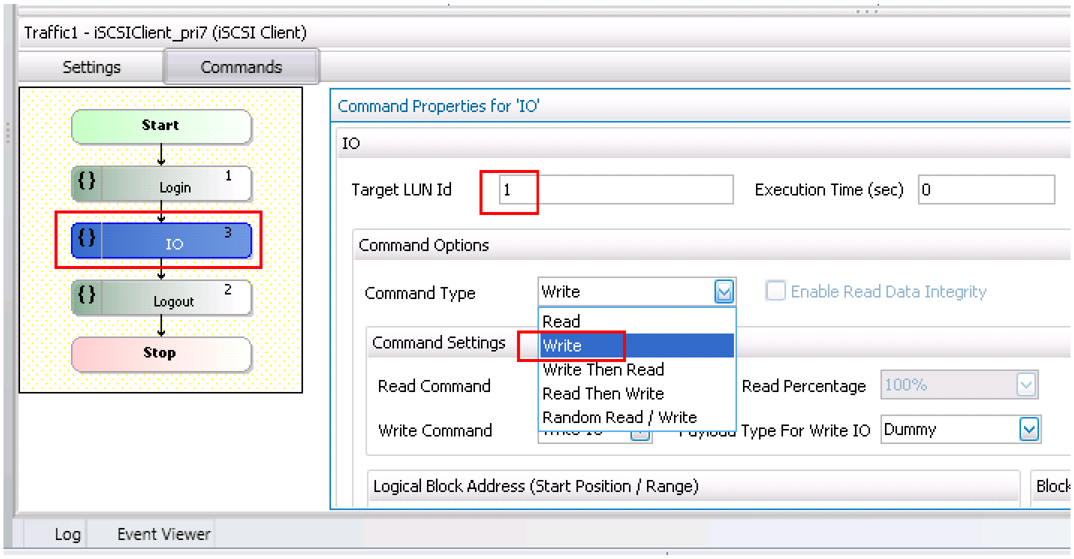

Traffic1 - iSCSIClient_pri7 (iSCSI Client)

Settings | Commands

Command Properties for 'IO'

IO

Target LUN Id: 1

Execution Time (sec): 0

Command Options

Command Type: Write

□ Enable Read Data Integrity

Command Settings

Read Command

Read Percentage: 100%

Write Command

Payload Type For Write IO: Dummy

Logical Block Address (Start Position / Range)

Figure 393. Setting the IO type as Write and the LUN id as 1.

e. In the **Data Transfer Length** section set the total transfer as **1 GB** and the **iSCSI Per Command Transfer Length** to **1MB**. This setting ensures that the iSCSI clients write large amount of storage traffic.

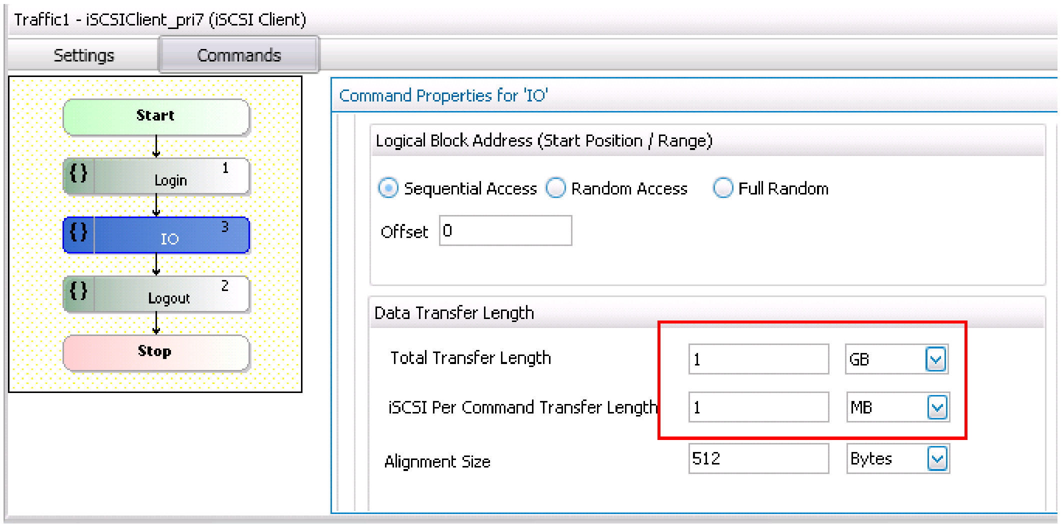

Traffic1 - iSCSIClient_pri7 (iSCSI Client)

Settings | Commands

Command Properties for 'IO'

Logical Block Address (Start Position / Range)

● Sequential Access ○ Random Access ○ Full Random

Offset: 0

Data Transfer Length

Total Transfer Length: 1 GB

iSCSI Per Command Transfer Length: 1 MB

Alignment Size: 512 Bytes

Figure 394. Setting the total data transfer length per IO and the iSCSI per command transfer length.

f. Click Settings > Advanced Options. Set the **VLAN priority** of the iSCSI traffic to **6**. This sets the COS bit in the VLAN header of the iSCSI packets. The PFC on Traffic Class 6 is selected at the switch ports; hence this traffic will have PFC enabled.

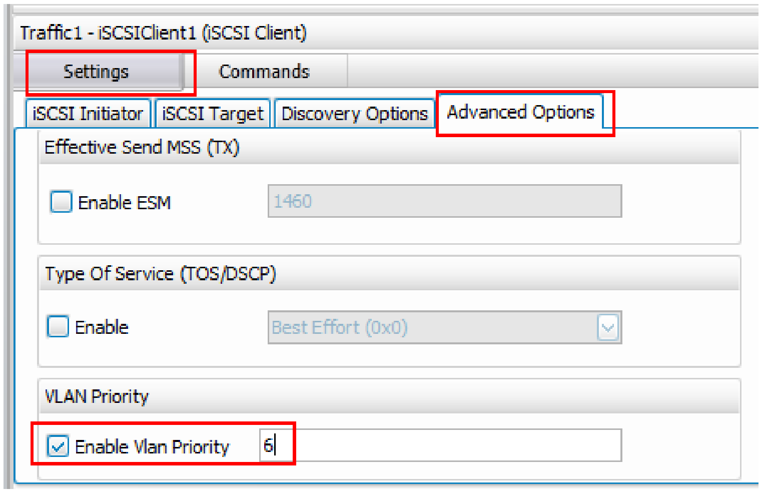

Traffic1 - iSCSIClient1 (iSCSI Client)

Settings | Commands

iSCSI Initiator | iSCSI Target | Discovery Options | Advanced Options

Effective Send MSS (TX)

☐ Enable ESM 1460

Type Of Service (TOS/DSCP)

☐ Enable Best Effort (0x0)

VLAN Priority

☑ Enable Vlan Priority 6

Figure 395. Setting the IO type as Write and the LUN id as 1.

21. Create a second storage activity emulating Traffic Class of 7. Select the iSCSI Client activity previously created and right click to **Copy** the iSCSI Client activity.
    a. Now click the **Traffic1** area of NetTraffic, right click, and then click **Paste**.

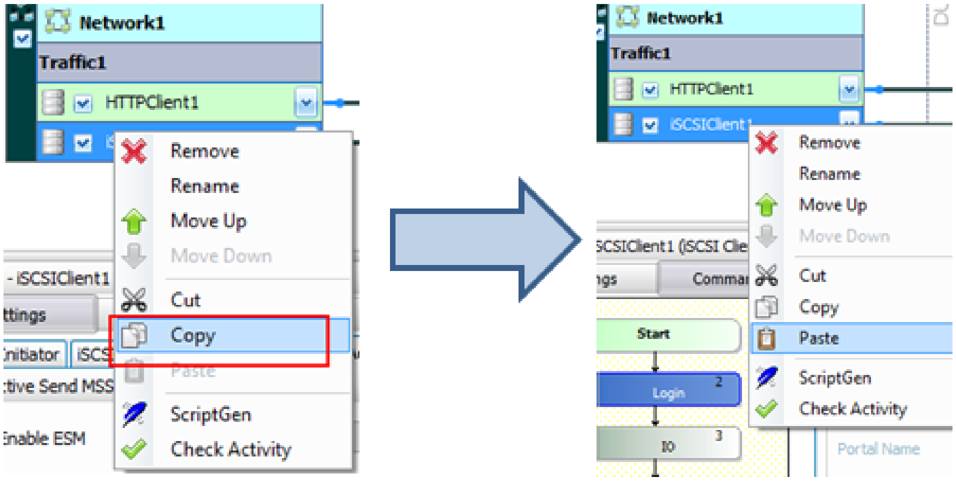

Figure 396. Copy and paste of the iSCSI client activities

b. This action creates two similar iSCSI activities. Rename each of them with the first one named as ‘**iSCSI Client_pri6**’ and the second as ‘**iSCSIClient_pri7**’

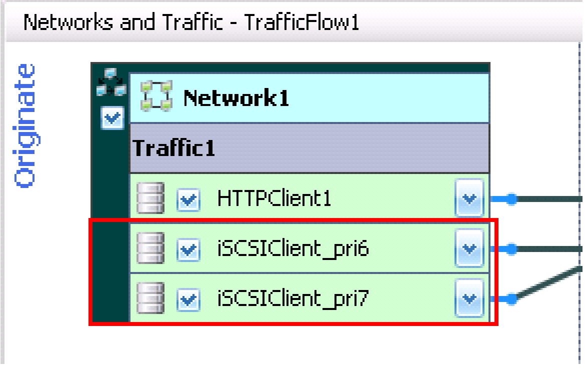

Figure 397. Renaming the activities to signify the traffic class that they are carrying.

c. For the second activity, set the **VLAN priority** to **7** in the **Advanced Options** tab. The two VLAN priorities basically create two different traffic classes on which PFC runs.

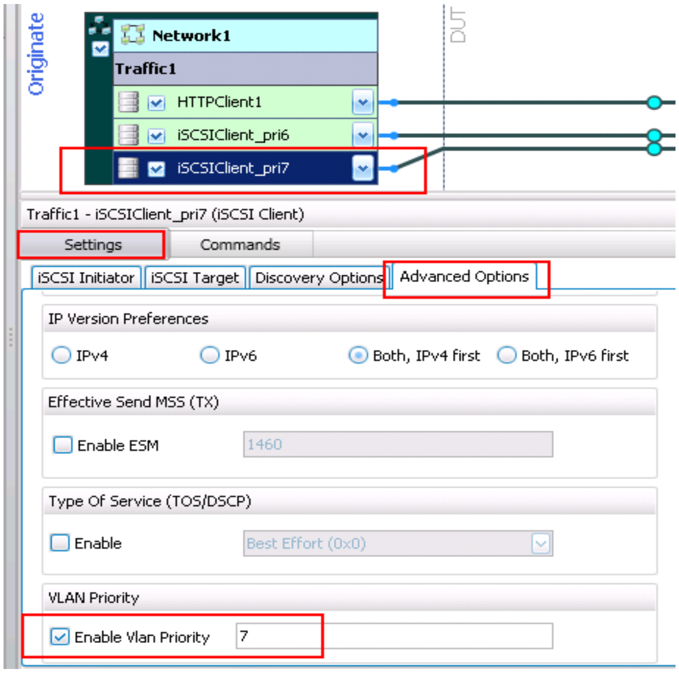

Traffic1 - iSCSIClient_pri7 (iSCSI Client)

IP Version Preferences

○ IPv4 ○ IPv6 ● Both, IPv4 first ○ Both, IPv6 first

Effective Send MSS (TX)

□ Enable ESM 1460

Type Of Service (TOS/DSCP)

□ Enable Best Effort (0x0)

VLAN Priority

☑ Enable Vlan Priority 7

Figure 398. Setting the priority 7 storage traffic.

22. Divide the two IP ranges so added at the client side such that each range carries only one priority calss of traffic. Map each iSCSCI activity to each of the IP ranges and later map range to a different Ixia port.
Activiy to IP range mapping is done by clicking the **IP Mappings** tab when Traffic1 is highlighted.Then using the checkboxes select the first range exclusively for iSCSI priority 6 traffic and second range for iSCSI priority 7 traffic. [Note: Both the ports carry some HTTP traffic]

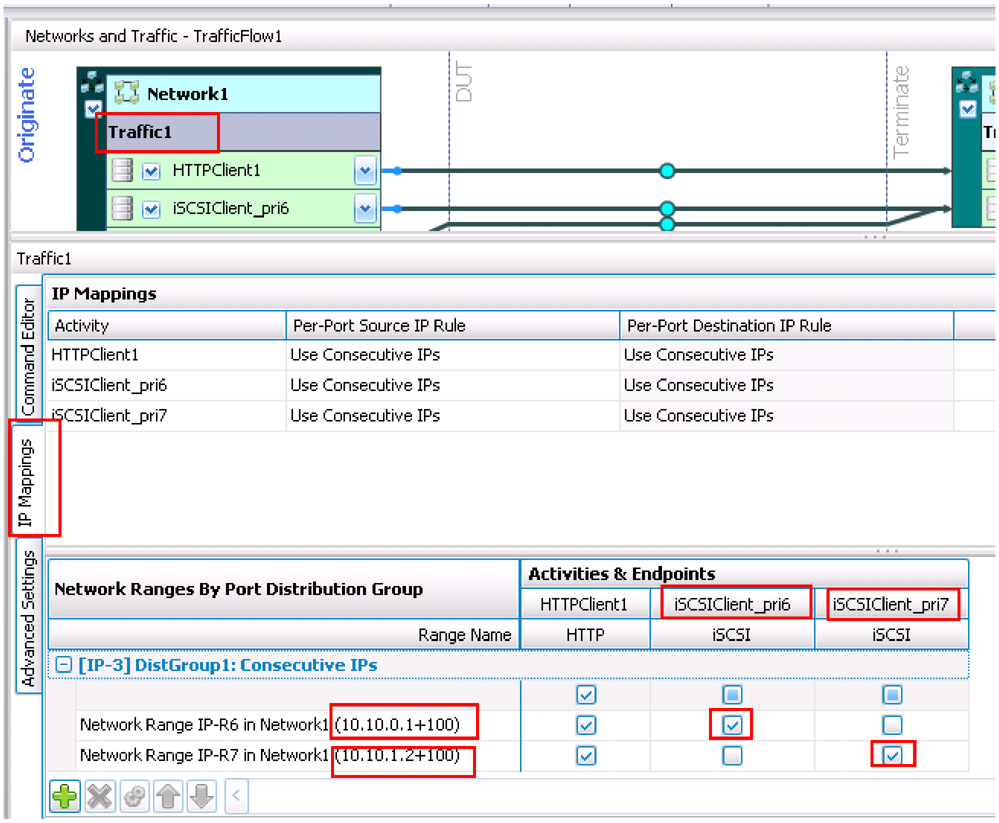

IP Mappings

Activity	Per-Port Source IP Rule	Per-Port Destination IP Rule
HTTPClient1	Use Consecutive IPs	Use Consecutive IPs
iSCSIClient_pri6	Use Consecutive IPs	Use Consecutive IPs
iSCSIClient_pri7	Use Consecutive IPs	Use Consecutive IPs

Network Ranges By Port Distribution Group	Activities & Endpoints		
	HTTPClient1	iSCSIClient_pri6	iSCSIClient_pri7
Range Name	HTTP	iSCSI	iSCSI
⊟ [IP-3] DistGroup1: Consecutive IPs			
	☑	▣	▣
Network Range IP-R6 in Network1 (10.10.0.1+100)	☑	☑	☐
Network Range IP-R7 in Network1 (10.10.1.2+100)	☑	☐	☑

Figure 399. Mapping the ranges to correct class of storage traffic.

23. Now to add the chassis and test ports, click **Ports** on the left navigation pane. On the context sentivie toolbar, click **Add Chassis** button and provide the **Chassis IP address** that is connected to the Switch (DUT).

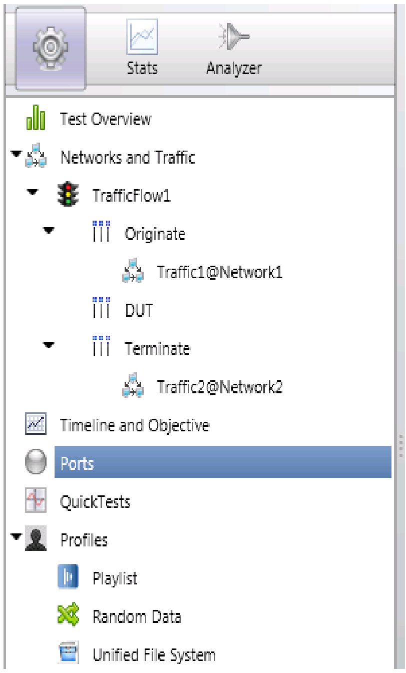

Figure 400. Adding Ixia chassis ports to the test.

24. Assign the correct **card** and **port** from the chassis to the NetTraffic. Here is the configuration and the interfaces that you must add to each NetTraffic.

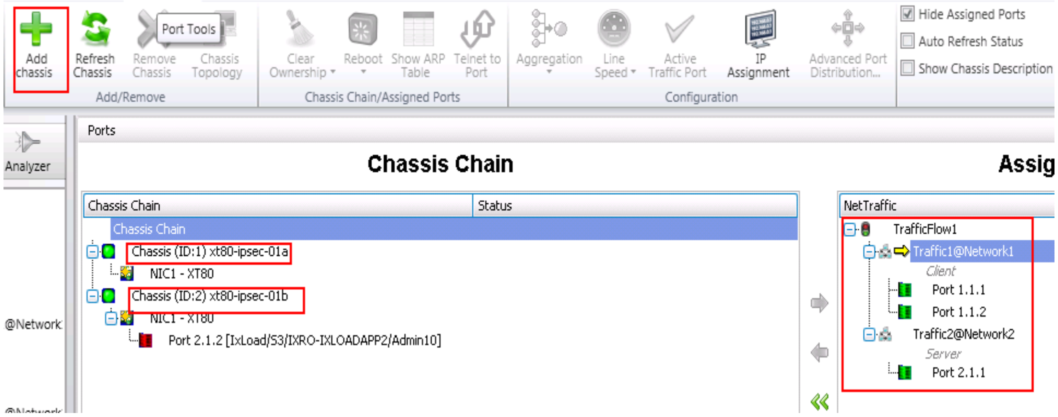

Figure 401. Adding Ixia chassis to the test

a. Let us explore the port assignments to understand their connectivity to the Switch.

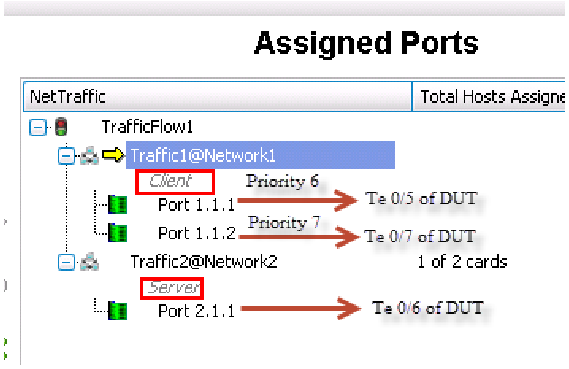

Assigned Ports

NetTraffic	Total Hosts Assigne
TrafficFlow1	
Traffic1@Network1	
Port 1.1.1	
Port 1.1.2	
Traffic2@Network2	1 of 2 cards
Port 2.1.1	

Figure 402. Representation of Ixia port to Switch interface connection.

The PFC Settings on each of the interfaces on the switch are as follows:

- Enabled PFC on class 6 and allocate 40% of the total bandwidth
- Enabled PFC on class 7 and allocate 40% of the total bandwidth
- Do not enable PFC on class 0(untagged) and allocate 20% of the total bandwidth.

25. Now that the iSCSI Initiators, HTTP Clients and the corresponding Servers are configured, it is time to set the Objective of each of the traffic in **Timeline and Objective** dialog slected from the left navigation pane.

    a. Configure **iSCSIClient_pri6** to generate **7Gbps** of storage Tx traffic. Similarly, configure **iSCSIClient_pri7** to generate **7Gbps** of Tx traffic and set **HTTPClient** (untagged) to generate **6Gbps** of Tx traffic. This action ensures a congestion is simulated as the total Tx traffic(HTTP Puts and iSCSI Writes) of 20Gbps is being generated by the two 10Gig ports, whereas the server has only one 10Gig port to service all of them.
    b. Set the **Sustain Time** to **60** minutes or more. This setting ensures the test runs for a sufficient amount of time for reliable data collection.

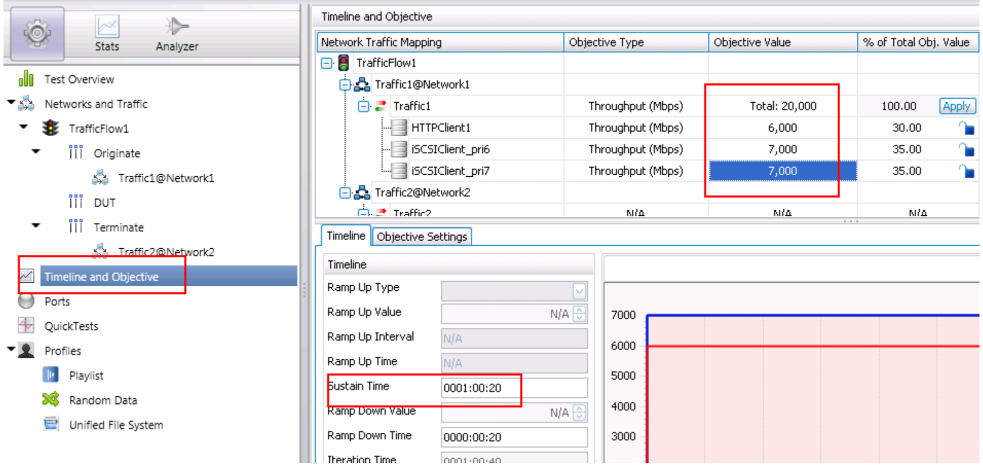

Timeline and Objective

Network Traffic Mapping	Objective Type	Objective Value	% of Total Obj. Value
TrafficFlow1			
Traffic1@Network1			
Traffic1	Throughput (Mbps)	Total: 20,000	100.00 Apply
HTTPClient1	Throughput (Mbps)	6,000	30.00
iSCSIClient_pri6	Throughput (Mbps)	7,000	35.00
iSCSIClient_pri7	Throughput (Mbps)	7,000	35.00
Traffic2@Network2			
Traffic2	N/A	N/A	N/A

Timeline | Objective Settings

Timeline

Ramp Up Type	
Ramp Up Value	N/A
Ramp Up Interval	N/A
Ramp Up Time	N/A
Sustain Time	0001:00:20
Ramp Down Value	N/A
Ramp Down Time	0000:00:20
Iteration Time	0001:00:40

Figure 403. Setting test timeline and objective and the sustain time.

26. Click **Test Options.** In the pop-up dialog select the **Enable Network Diagnostics** checkbox. This selection allows to publish the DCBx related stats.

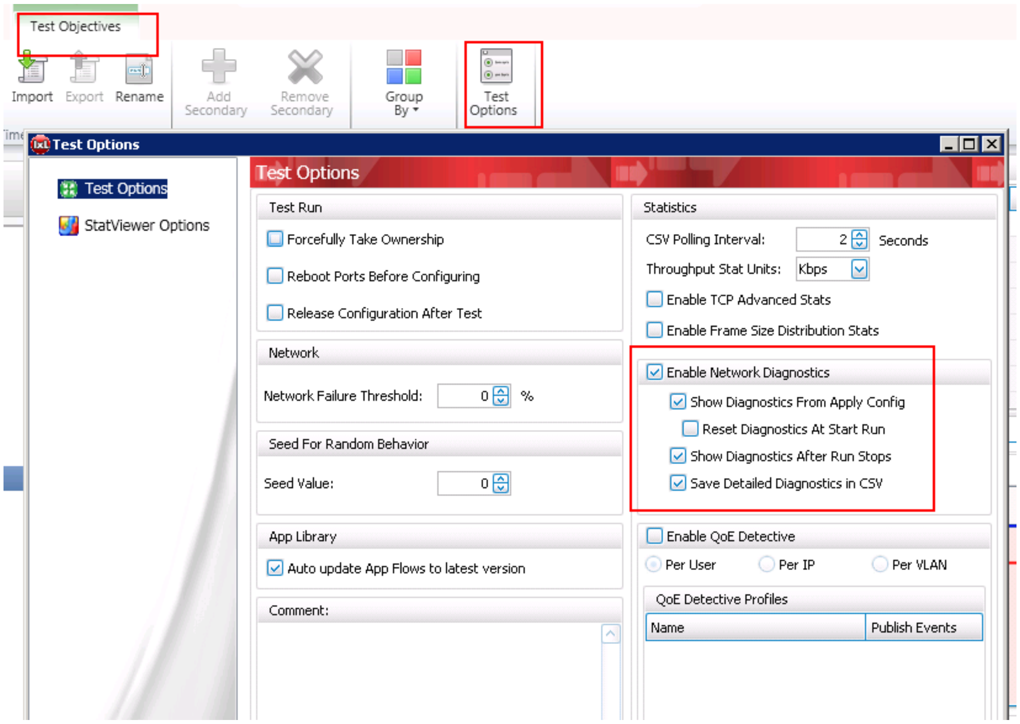

Test Options

Test Options
StatViewer Options

Test Options

Test Run
☐ Forcefully Take Ownership
☐ Reboot Ports Before Configuring
☐ Release Configuration After Test

Network
Network Failure Threshold: 0 %

Seed For Random Behavior
Seed Value: 0

App Library
☑ Auto update App Flows to latest version

Comment:

Statistics
CSV Polling Interval: 2 Seconds
Throughput Stat Units: Kbps
☐ Enable TCP Advanced Stats
☐ Enable Frame Size Distribution Stats

☑ Enable Network Diagnostics
☑ Show Diagnostics From Apply Config
☐ Reset Diagnostics At Start Run
☑ Show Diagnostics After Run Stops
☑ Save Detailed Diagnostics in CSV

☐ Enable QoE Detective
○ Per User ○ Per IP ○ Per VLAN

QoE Detective Profiles

Name	Publish Events

Figure 404. Enabling network diagnostics to enable DCB and PFC stats in stat viewer

27. Click the Start Test1 button to start the test.

28. Run the test for a few minutes to allow the performance to reach a steady-state. Steady state is referred as **Sustain** duration in the test. Continue to monitor the DUT for any failure/error counters. See the **Results Analysis** section below for important statistics and diagnostics information.

    [Note: Interpretation of result statistics can sometimes be difficult, deducing what they mean under different circumstances. The **Results Analysis** section below provides a diagnostics-based approach to highlight some common scenarios, the statistics being reported, and how to interpret them.]

29. Iterate through the test varying the test variable described in the table below to determine the the DCB characterization of the switch.

    a. The test tool can be started and stopped in the middle of a test cycle, or wait for it to be gracefully stopped using the test controls shown here.

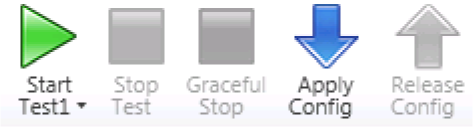

Test Variables

Functional Variable	Description
CEE Maps in switch interface	Change the CEE Map of the switch interfaces to different priority group percentages during the test run time and check how it affects the traffic. Refer to Result section **Change the CEE Map on the switch interface during test run** for more details.
Change objective on the fly	IxLoad objectives can be changed on the fly. Reduce the Storage traffic during test run and check if the LAN application traffic is able to take up the additional bandwidth. Later increase the storage traffic to the previous value and recheck if the LAN traffic is reduced again. Refer to result section **Change Objective on the Fly** for more details.
Remove DCBx	Check the effect of running the same test without DCBx by reverting back to plain IP. Refer to result section **Remove DCBx from the Client side** for more details.
Introduce Rx Traffic	Check the effect of Rx traffic. Reconfigure the ISCSI clients to have ‘Reads’ and the HTTP clients to have ‘GET’ and observe the difference of introducing Rx traffic on the switch.

Result Analysis

IxLoad Statviewer publishes several stats for analysis and debugs. For a DCBx test, to get the overview see the DCBx stats, PFC stats total throughputs, and total TCP failures.

Use the IxLoad statistic view to verify the following results:

- Check the **DCBx statistics** to verify that all Ixia ports can successfully complete the DCBx session.

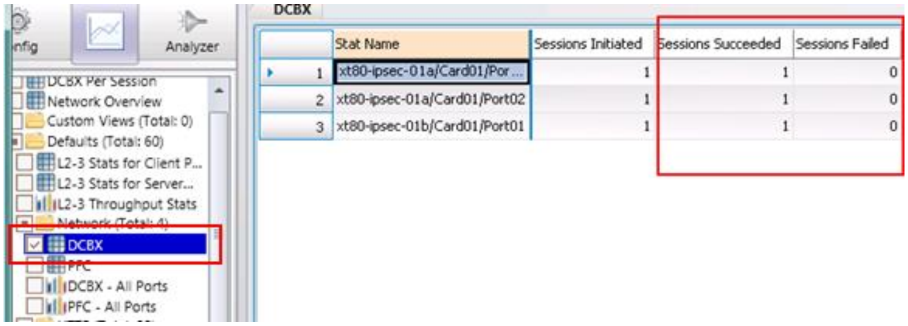

	Stat Name	Sessions Initiated	Sessions Succeeded	Sessions Failed
1	xt80-ipsec-01a/Card01/Por...	1	1	0
2	xt80-ipsec-01a/Card01/Port02	1	1	0
3	xt80-ipsec-01b/Card01/Port01	1	1	0

Figure 405. DCBx stat showing all the DCB sessions successful.

- The PFC statistics view must show pause frames are received only at TC6 on the port running class 6 storage traffic and TC7 on the port running class 7 storage trffic as shown below. The pause should not be received on any other classes.

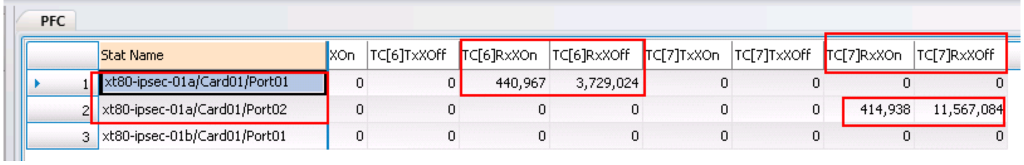

	Stat Name	XOn	TC[6]TxXOff	TC[6]RxXOn	TC[6]RxXOff	TC[7]TxXOn	TC[7]TxXOff	TC[7]RxXOn	TC[7]RxXOff
1	xt80-ipsec-01a/Card01/Port01	0	0	440,967	3,729,024	0	0	0	0
2	xt80-ipsec-01a/Card01/Port02	0	0	0	0	0	0	414,938	11,567,084
3	xt80-ipsec-01b/Card01/Port01	0	0	0	0	0	0	0	0

Figure 406. PFC stat showing pause On and Off frames received from the switch on the Traffic Class 6 on port 01 and traffic class 07 on port 02.

- **PFC – All Ports** statistics view. Right click the stats and click **Drill Down Per session**.

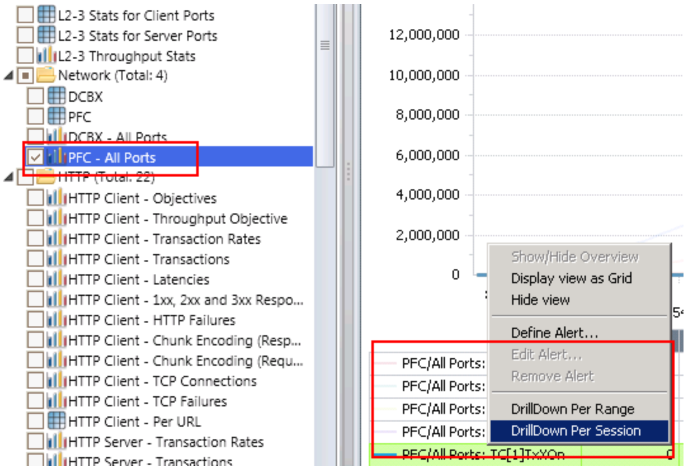
L2-3 Stats for Client Ports
L2-3 Stats for Server Ports
L2-3 Throughput Stats
Network (Total: 4)
DCBX
PFC
DCBX - All Ports
☑ PFC - All Ports
HTTP (Total: 22)
HTTP Client - Objectives
HTTP Client - Throughput Objective
HTTP Client - Transaction Rates
HTTP Client - Transactions
HTTP Client - Latencies
HTTP Client - 1xx, 2xx and 3xx Respo...
HTTP Client - HTTP Failures
HTTP Client - Chunk Encoding (Resp...
HTTP Client - Chunk Encoding (Requ...
HTTP Client - TCP Connections
HTTP Client - TCP Failures
HTTP Client - Per URL
HTTP Server - Transaction Rates
HTTP Server - Transactions
12,000,000
10,000,000
8,000,000
6,000,000
4,000,000
2,000,000
0
Show/Hide Overview
Display view as Grid
Hide view
Define Alert...
Edit Alert...
Remove Alert
DrillDown Per Range
DrillDown Per Session
PFC/All Ports:
PFC/All Ports:
PFC/All Ports:
PFC/All Ports:
PFC/All Ports: TC[1]TxXOn

Figure 407. Drilling down on “PFC-All ports”

- The Drill down per Session must show each of the DCBx TLV negotiatied. Ensure that the Peer Config is similar to the Local Config (Since theDCBx ranges have been set to “Willing” at Ixia side). Also the PGID map must show correct percentages for the traffic maps and the PFC must reflect the correct traffic classes on which it is enabled.

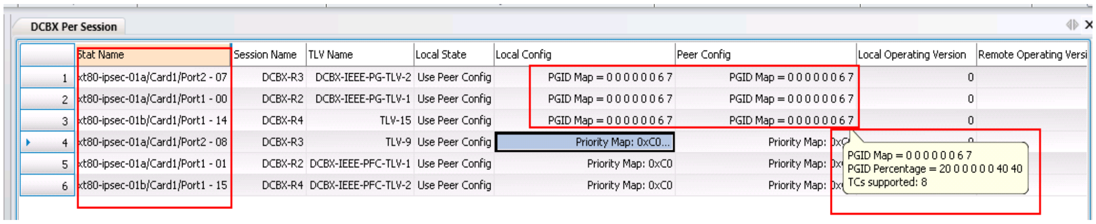

DCBX Per Session

	Stat Name	Session Name	TLV Name	Local State	Local Config	Peer Config	Local Operating Version	Remote Operating Versi
1	xt80-ipsec-01a/Card1/Port2 - 07	DCBX-R3	DCBX-IEEE-PG-TLV-2	Use Peer Config	PGID Map = 0 0 0 0 0 0 6 7	PGID Map = 0 0 0 0 0 0 6 7	0	
2	xt80-ipsec-01a/Card1/Port1 - 00	DCBX-R2	DCBX-IEEE-PG-TLV-1	Use Peer Config	PGID Map = 0 0 0 0 0 0 6 7	PGID Map = 0 0 0 0 0 0 6 7	0	
3	xt80-ipsec-01b/Card1/Port1 - 14	DCBX-R4	TLV-15	Use Peer Config	PGID Map = 0 0 0 0 0 0 6 7	PGID Map = 0 0 0 0 0 0 6 7	0	
4	xt80-ipsec-01a/Card1/Port2 - 08	DCBX-R3	TLV-9	Use Peer Config	Priority Map: 0xC0...	Priority Map: 0xC	0	
5	xt80-ipsec-01a/Card1/Port1 - 01	DCBX-R2	DCBX-IEEE-PFC-TLV-1	Use Peer Config	Priority Map: 0xC0	Priority Map: 0x		
6	xt80-ipsec-01b/Card1/Port1 - 15	DCBX-R4	DCBX-IEEE-PFC-TLV-2	Use Peer Config	Priority Map: 0xC0	Priority Map: 0x		

PGID Map = 0 0 0 0 0 0 6 7
PGID Percentage = 20 0 0 0 0 0 40 40
TCs supported: 8

Figure 408. Drill down per session gives extensive details about each DCB session, per TLV per Port.

- Further LLDP/DCB level debuggins can be done by checking the **DCBx All Ports** statistics view.

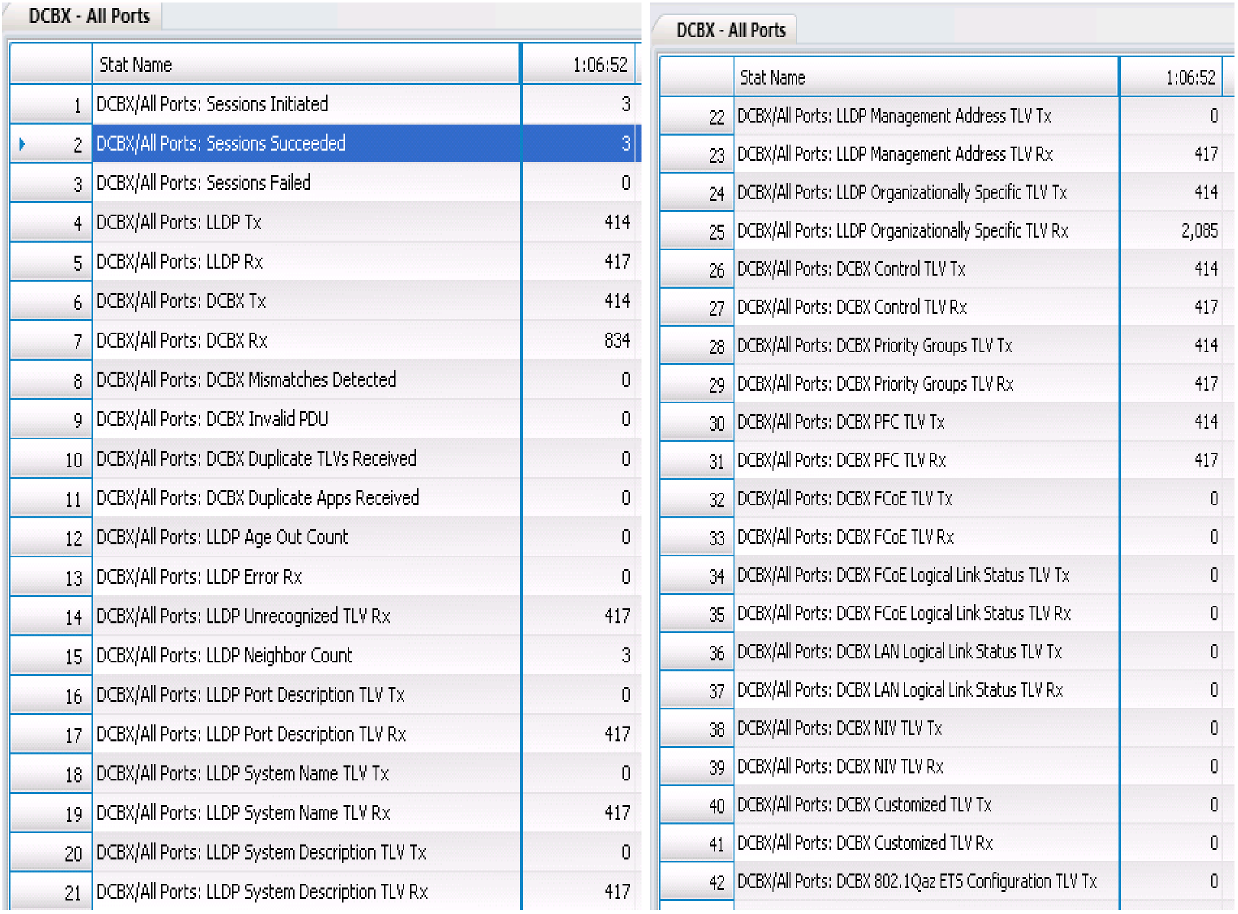

DCBX - All Ports

	Stat Name	1:06:52
1	DCBX/All Ports: Sessions Initiated	3
2	DCBX/All Ports: Sessions Succeeded	3
3	DCBX/All Ports: Sessions Failed	0
4	DCBX/All Ports: LLDP Tx	414
5	DCBX/All Ports: LLDP Rx	417
6	DCBX/All Ports: DCBX Tx	414
7	DCBX/All Ports: DCBX Rx	834
8	DCBX/All Ports: DCBX Mismatches Detected	0
9	DCBX/All Ports: DCBX Invalid PDU	0
10	DCBX/All Ports: DCBX Duplicate TLVs Received	0
11	DCBX/All Ports: DCBX Duplicate Apps Received	0
12	DCBX/All Ports: LLDP Age Out Count	0
13	DCBX/All Ports: LLDP Error Rx	0
14	DCBX/All Ports: LLDP Unrecognized TLV Rx	417
15	DCBX/All Ports: LLDP Neighbor Count	3
16	DCBX/All Ports: LLDP Port Description TLV Tx	0
17	DCBX/All Ports: LLDP Port Description TLV Rx	417
18	DCBX/All Ports: LLDP System Name TLV Tx	0
19	DCBX/All Ports: LLDP System Name TLV Rx	417
20	DCBX/All Ports: LLDP System Description TLV Tx	0
21	DCBX/All Ports: LLDP System Description TLV Rx	417

DCBX - All Ports

	Stat Name	1:06:52
22	DCBX/All Ports: LLDP Management Address TLV Tx	0
23	DCBX/All Ports: LLDP Management Address TLV Rx	417
24	DCBX/All Ports: LLDP Organizationally Specific TLV Tx	414
25	DCBX/All Ports: LLDP Organizationally Specific TLV Rx	2,085
26	DCBX/All Ports: DCBX Control TLV Tx	414
27	DCBX/All Ports: DCBX Control TLV Rx	417
28	DCBX/All Ports: DCBX Priority Groups TLV Tx	414
29	DCBX/All Ports: DCBX Priority Groups TLV Rx	417
30	DCBX/All Ports: DCBX PFC TLV Tx	414
31	DCBX/All Ports: DCBX PFC TLV Rx	417
32	DCBX/All Ports: DCBX FCoE TLV Tx	0
33	DCBX/All Ports: DCBX FCoE TLV Rx	0
34	DCBX/All Ports: DCBX FCoE Logical Link Status TLV Tx	0
35	DCBX/All Ports: DCBX FCoE Logical Link Status TLV Rx	0
36	DCBX/All Ports: DCBX LAN Logical Link Status TLV Tx	0
37	DCBX/All Ports: DCBX LAN Logical Link Status TLV Rx	0
38	DCBX/All Ports: DCBX NIV TLV Tx	0
39	DCBX/All Ports: DCBX NIV TLV Rx	0
40	DCBX/All Ports: DCBX Customized TLV Tx	0
41	DCBX/All Ports: DCBX Customized TLV Rx	0
42	DCBX/All Ports: DCBX 802.1Qaz ETS Configuration TLV Tx	0

Figure 409. Several DCBx statistics available for debug

- Select the **HTTP Client – Throughput Objective** and **iSCSI Initiator – Throughput Objective** statistics view. The view must show very high storage traffic whereas the LAN traffic is limited to around 1 Gig. (Because the present map gurantees 80% to class 6 and class 7 traffic).

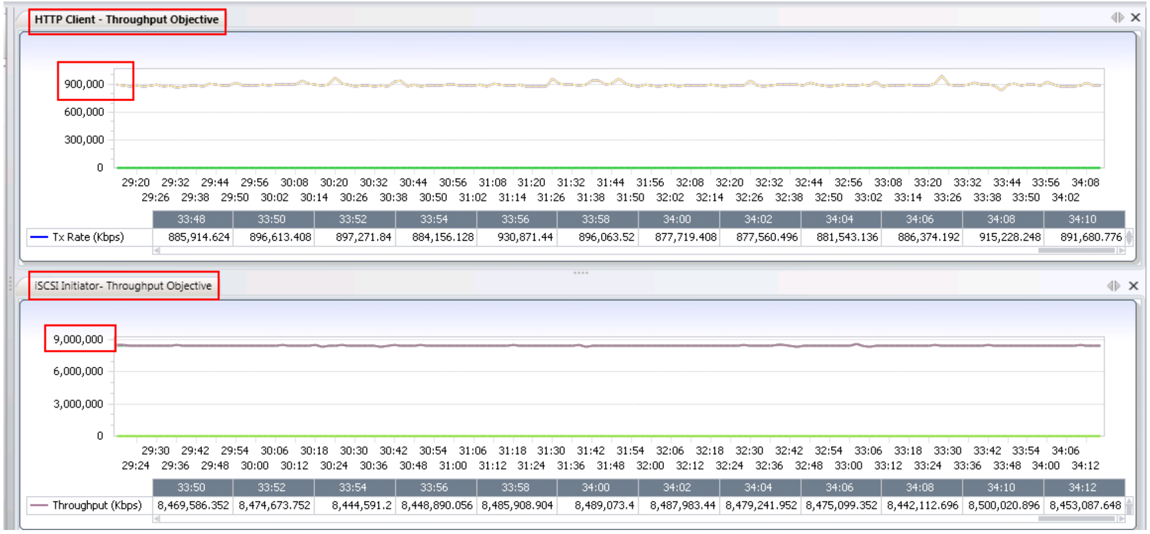

	33:48	33:50	33:52	33:54	33:56	33:58	34:00	34:02	34:04	34:06	34:08	34:10
Tx Rate (Kbps)	885,914.624	896,613.408	897,271.84	884,156.128	930,871.44	896,063.52	877,719.408	877,560.496	881,543.136	886,374.192	915,228.248	891,680.776

	33:50	33:52	33:54	33:56	33:58	34:00	34:02	34:04	34:06	34:08	34:10	34:12
Throughput (Kbps)	8,469,586.352	8,474,673.752	8,444,591.2	8,448,890.056	8,485,908.904	8,489,073.4	8,487,983.44	8,479,241.952	8,475,099.352	8,442,112.696	8,500,020.896	8,453,087.648

Figure 410. Throughput comparision showing very high and stable Storage thtoughput and low HTTP throughput

- Select the **HTTP Client – TCP Failures** and the **iSCSI Initiators – TCP failures** statistics view. It shows very high around 8Million retries of HTTP traffic wheeas the iSCSI retries is limited to 300.

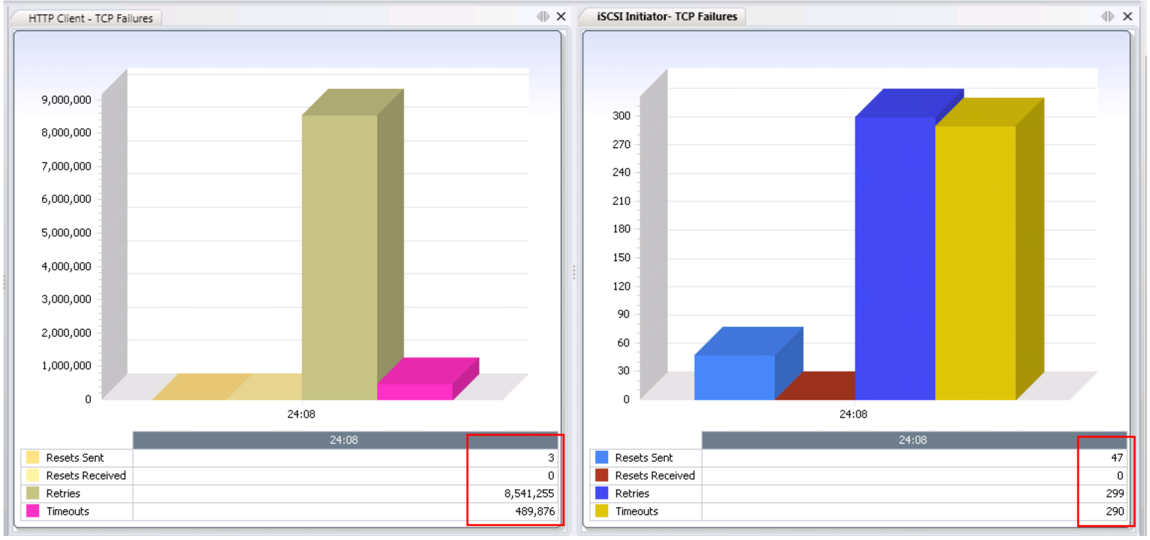

	24:08
Resets Sent	3
Resets Received	0
Retries	8,541,255
Timeouts	489,876

	24:08
Resets Sent	47
Resets Received	0
Retries	299
Timeouts	290

Figure 411. HTTP and iSCSI TCP Failures stat comparision shows that while HTTP had over 8Million retries the iSCSI has only 300 retries.

Change the CEE Map on the switch interface during test run.

Remove the earlier map pfc4 and add a new map PFC_no_class_6

Let the new CEE Map **PFC_no_class_6** disable PFC from class 6. Set the PFC only on Class 7 with 20% bandwidth allocated to the class 7 traffic in the ETS and the rest bandwidth shared between class 6 and the LAN traffic.

The storage traffic running on class 6 and 7 immediately goes down and the HTTP traffic jumps to higher throughput.

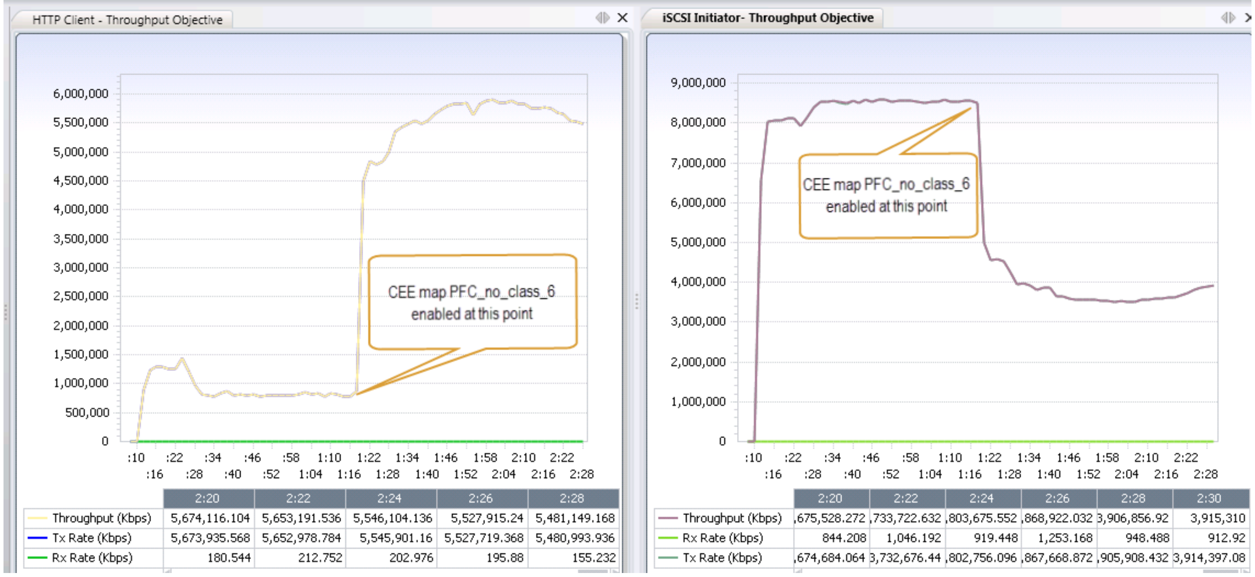

CEE map PFC_no_class_6 enabled at this point

CEE map PFC_no_class_6 enabled at this point

	2:20	2:22	2:24	2:26	2:28
Throughput (Kbps)	5,674,116.104	5,653,191.536	5,546,104.136	5,527,915.24	5,481,149.168
Tx Rate (Kbps)	5,673,935.568	5,652,978.784	5,545,901.16	5,527,719.368	5,480,993.936
Rx Rate (Kbps)	180.544	212.752	202.976	195.88	155.232

	2:20	2:22	2:24	2:26	2:28	2:30
Throughput (Kbps)	,675,528.272	,733,722.632	,803,675.552	,868,922.032	3,906,856.92	3,915,310
Rx Rate (Kbps)	844.208	1,046.192	919.448	1,253.168	948.488	912.92
Tx Rate (Kbps)	,674,684.064	3,732,676.44	,802,756.096	,867,668.872	,905,908.432	3,914,397.08

Figure 412. HTTP throughput immediately picks up as PFC is disabled in traffic class 6. The storage traffic goes down.

Re-enable the original PFC that used to:

- Enabled PFC on class 6 and allocate 40% of the total bandwidth
- Enabled PFC on class 7 and allocate 40% of the total bandwidth
- Do not enable PFC on class 0(untagged) and allocate 20% of the total bandwidth.

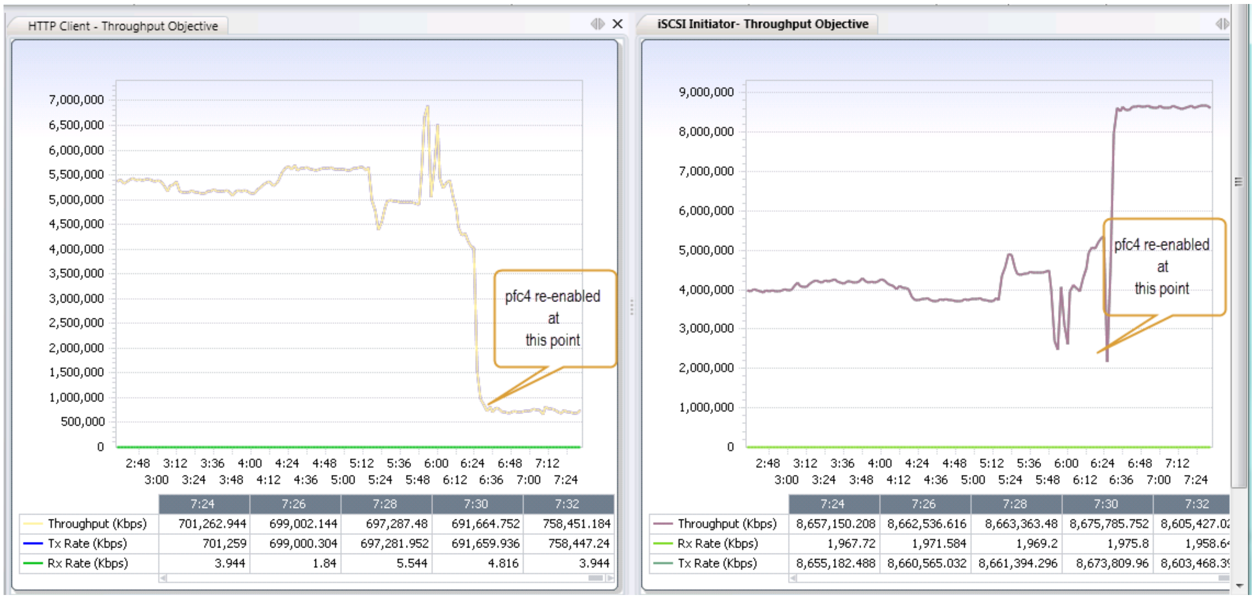

	7:24	7:26	7:28	7:30	7:32
Throughput (Kbps)	701,262.944	699,002.144	697,287.48	691,664.752	758,451.184
Tx Rate (Kbps)	701,259	699,000.304	697,281.952	691,659.936	758,447.24
Rx Rate (Kbps)	3.944	1.84	5.544	4.816	3.944

	7:24	7:26	7:28	7:30	7:32
Throughput (Kbps)	8,657,150.208	8,662,536.616	8,663,363.48	8,675,785.752	8,605,427.0
Rx Rate (Kbps)	1,967.72	1,971.584	1,969.2	1,975.8	1,958.6
Tx Rate (Kbps)	8,655,182.488	8,660,565.032	8,661,394.296	8,673,809.96	8,603,468.3

Figure 413. Storage traffic moves back to its original values the momet traffic class 6 has pfc re-enabled.

Change Objective on the Fly.

Reduce the iSCSI objective on-the-fly (while the test is running) by reducing the storage Objective to 100 Mbps for both priorities.

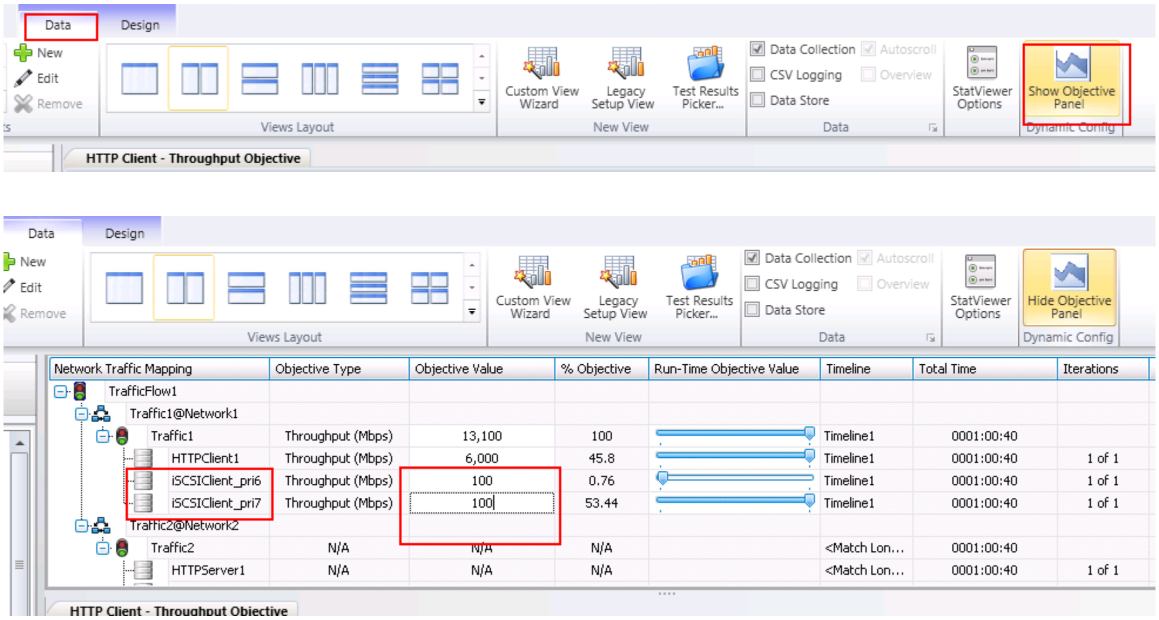

Network Traffic Mapping	Objective Type	Objective Value	% Objective	Run-Time Objective Value	Timeline	Total Time	Iterations
TrafficFlow1							
Traffic1@Network1							
Traffic1	Throughput (Mbps)	13,100	100		Timeline1	0001:00:40	
HTTPClient1	Throughput (Mbps)	6,000	45.8		Timeline1	0001:00:40	1 of 1
iSCSIClient_pri6	Throughput (Mbps)	100	0.76		Timeline1	0001:00:40	1 of 1
iSCSIClient_pri7	Throughput (Mbps)	100	53.44		Timeline1	0001:00:40	1 of 1
Traffic2@Network2							
Traffic2	N/A	N/A	N/A		<Match Lon…	0001:00:40	
HTTPServer1	N/A	N/A	N/A		<Match Lon…	0001:00:40	1 of 1

Figure 414. Changing the objective of the test on-the-fly during test run.

The HTTP objective immediately should reach its peak of 6000 Mbps.

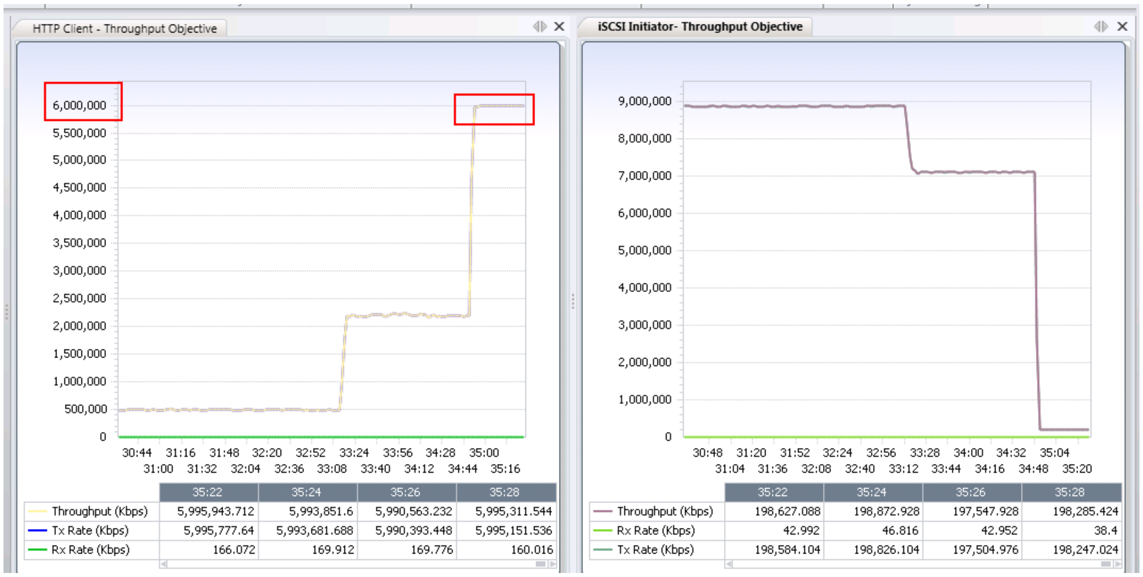

	35:22	35:24	35:26	35:28
Throughput (Kbps)	5,995,943.712	5,993,851.6	5,990,563.232	5,995,311.544
Tx Rate (Kbps)	5,995,777.64	5,993,681.688	5,990,393.448	5,995,151.536
Rx Rate (Kbps)	166.072	169.912	169.776	160.016

	35:22	35:24	35:26	35:28
Throughput (Kbps)	198,627.088	198,872.928	197,547.928	198,285.424
Rx Rate (Kbps)	42.992	46.816	42.952	38.4
Tx Rate (Kbps)	198,584.104	198,826.104	197,504.976	198,247.024

Figure 415. HTTP traffic reaches its peak of 6 Gbps.

Revert back to the previous value using the same “on-the-fly” and the Storage must again go back to the 8400 Mbps that it was achieving earlier.

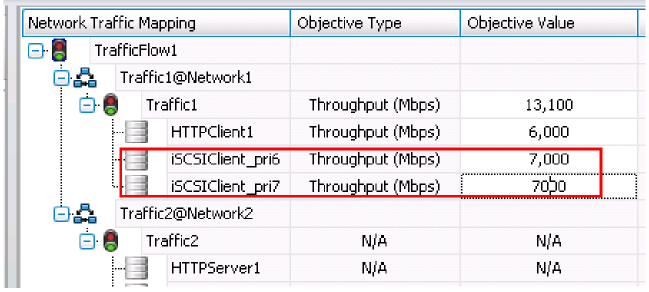

Network Traffic Mapping	Objective Type	Objective Value
TrafficFlow1		
Traffic1@Network1		
Traffic1	Throughput (Mbps)	13,100
HTTPClient1	Throughput (Mbps)	6,000
iSCSIClient_pri6	Throughput (Mbps)	7,000
iSCSIClient_pri7	Throughput (Mbps)	7000
Traffic2@Network2		
Traffic2	N/A	N/A
HTTPServer1	N/A	N/A

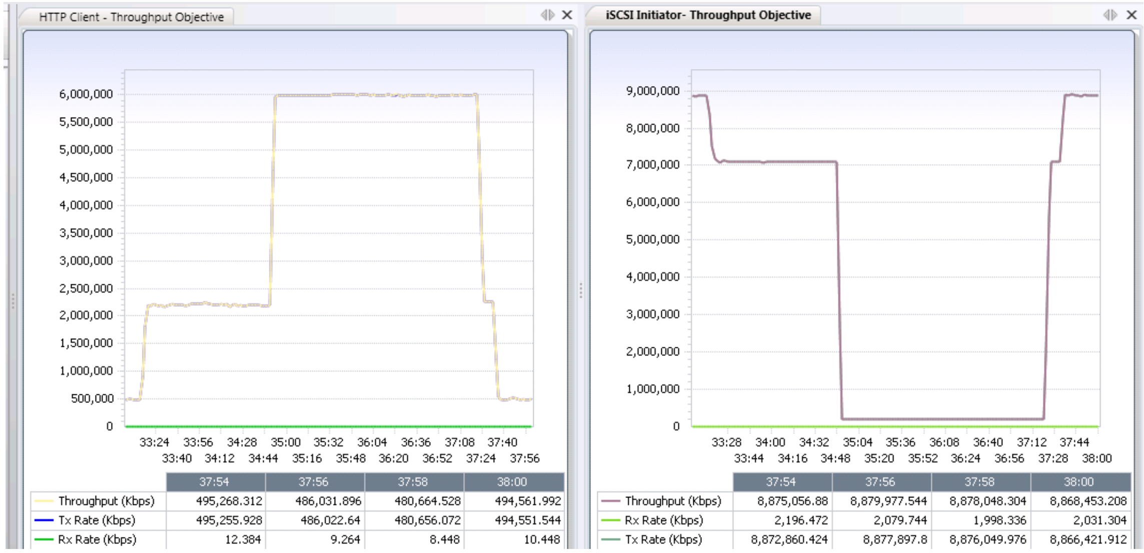

	37:54	37:56	37:58	38:00
Throughput (Kbps)	495,268.312	486,031.896	480,664.528	494,561.992
Tx Rate (Kbps)	495,255.928	486,022.64	480,656.072	494,551.544
Rx Rate (Kbps)	12.384	9.264	8.448	10.448

	37:54	37:56	37:58	38:00
Throughput (Kbps)	8,875,056.88	8,879,977.544	8,878,048.304	8,868,453.208
Rx Rate (Kbps)	2,196.472	2,079.744	1,998.336	2,031.304
Tx Rate (Kbps)	8,872,860.424	8,877,897.8	8,876,049.976	8,866,421.912

Figure 416. The moment storage traffic is re-increased , it achieves its previous high value. Indicating that PFC isworking all the time.

Remove DCBx from the Client side:

To show the effectiveness of DCBx, remove the DCBx completely from the storage and HTTP traffic. This is achieved by first deleting the DCBx stack at the client side and later adding two IP ranges on the IP stack.

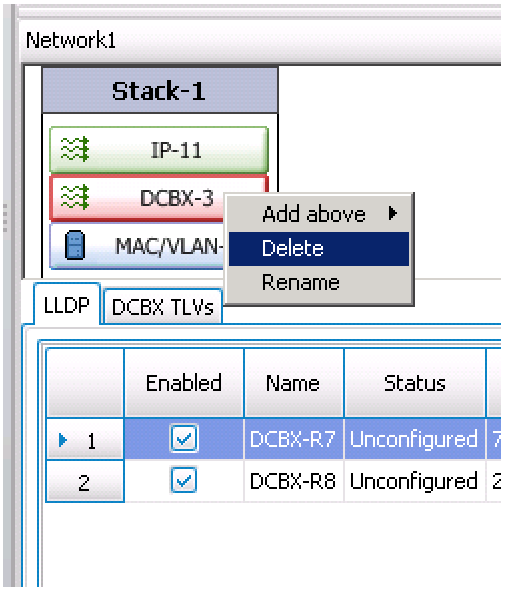

	Enabled	Name	Status	
1	☑	DCBX-R7	Unconfigured	7
2	☑	DCBX-R8	Unconfigured	2

Figure 417. Removing DCBx at the client side.

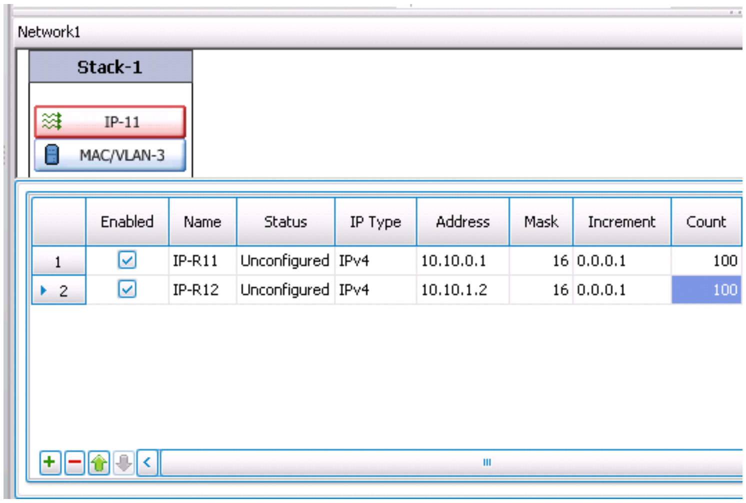

	Enabled	Name	Status	IP Type	Address	Mask	Increment	Count
1	☑	IP-R11	Unconfigured	IPv4	10.10.0.1	16	0.0.0.1	100
2	☑	IP-R12	Unconfigured	IPv4	10.10.1.2	16	0.0.0.1	100

Figure 418. Reverting back to plain IP.

- Storage Traffic shows huge amount of retries 55Million within 18 minutes.

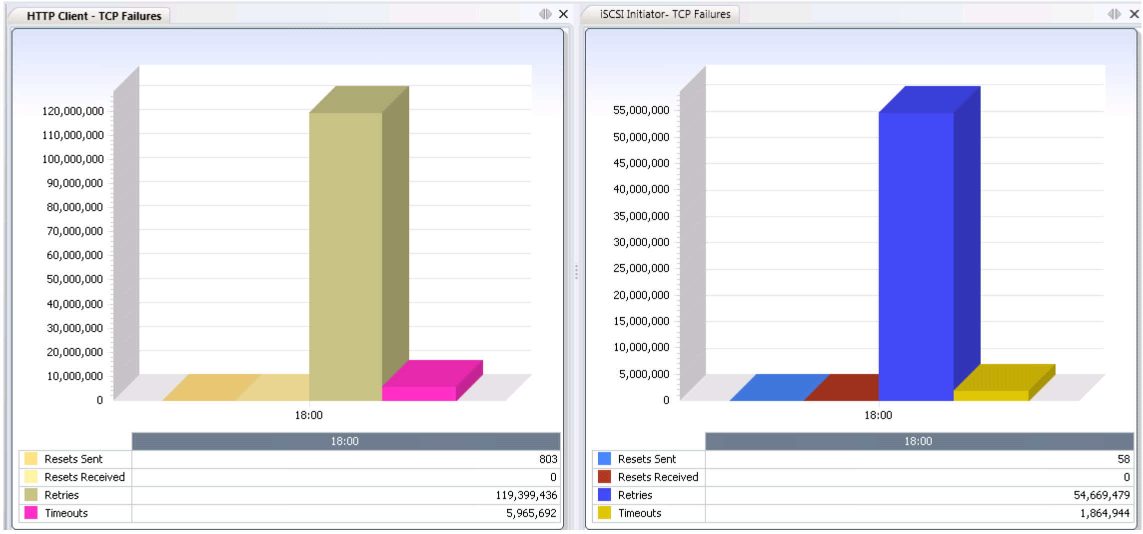

HTTP Client - TCP Failures

	18:00
Resets Sent	803
Resets Received	0
Retries	119,399,436
Timeouts	5,965,692

iSCSI Initiator- TCP Failures

	18:00
Resets Sent	58
Resets Received	0
Retries	54,669,479
Timeouts	1,864,944

Figure 419. Both HTTP and iSCSI now shows huge amount of retries of the order of billions within a short span of 18 minutes.

- The Throughput distribution between Storage and HTTP is unpredictable and jittery.

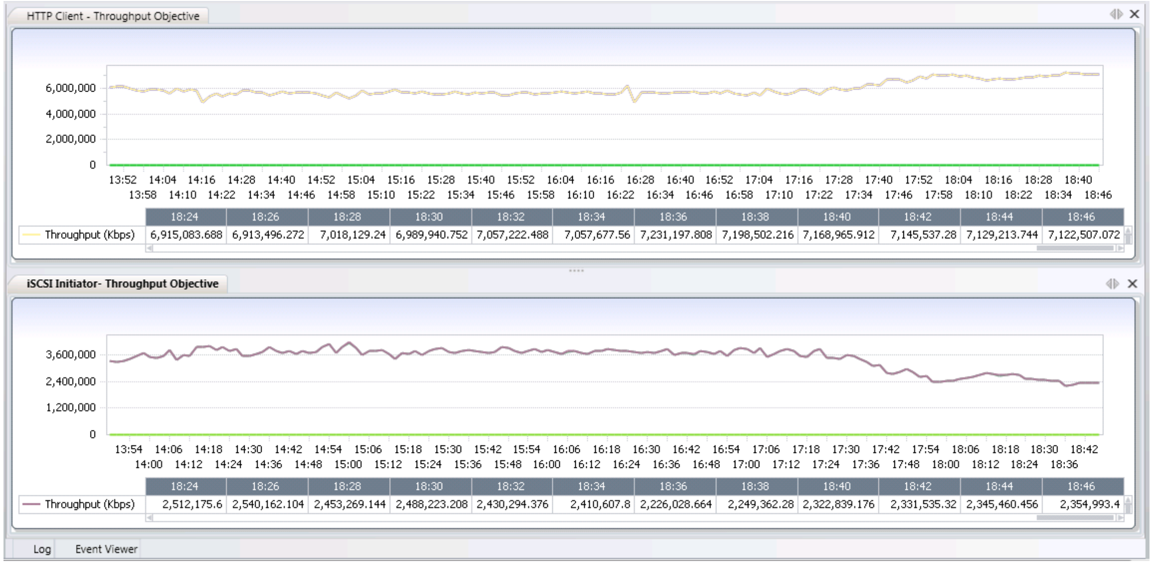

HTTP Client - Throughput Objective

	18:24	18:26	18:28	18:30	18:32	18:34	18:36	18:38	18:40	18:42	18:44	18:46
Throughput (Kbps)	6,915,083.688	6,913,496.272	7,018,129.24	6,989,940.752	7,057,222.488	7,057,677.56	7,231,197.808	7,198,502.216	7,168,965.912	7,145,537.28	7,129,213.744	7,122,507.072

iSCSI Initiator- Throughput Objective

	18:24	18:26	18:28	18:30	18:32	18:34	18:36	18:38	18:40	18:42	18:44	18:46
Throughput (Kbps)	2,512,175.6	2,540,162.104	2,453,269.144	2,488,223.208	2,430,294.376	2,410,607.8	2,226,028.664	2,249,362.28	2,322,839.176	2,331,535.32	2,345,460.456	2,354,993.4

Figure 420. The throughput distribution with both HTTP and Storage traffic is not unpredictable and jittery.

Conclusions

The test showcased the efficiency of the DCB switch in handling the storage traffic in an Ethernet environment.

At very high congestion, the storage traffic always got the required priority and remained lossless. The priority group settings were respected and the traffic class on which PFC was enabnled generated the pause frames and was also allocated the sufficient bandwidth as configured in the priority group settings. Apart from that when the storage traffic rate was reduced, the DCB switch allowed the LAN traffic to occupy the rest of the bandwidth. The bandwidth allocations when changed on-the-fly the same was reflected immediately.

Overall this test, if successful, can provide sufficient confidence to the IT personnel for moving the storage server networks from FC to FCoE or Ethernet without worrying about the effects on end users and enterprise applications.

Test Case: Cloud Performance Testing

Objective

This test case, measures the forwarding performance of a data center top-of-rack (ToR) switch or network using simulated data center North-South and East-West traffic profiles.

Setup

North-South HTTP traffic is set up on one north and one south port. The request is 83 bytes and the response is 305 bytes.

North-South YouTube traffic uses the same north and south ports as the HTTP traffic. The request is 500. The response traffic is further broken down into a 5/2/1 percentage breakdown of 1518, 512 and 64 bytes.

East-West database traffic is set up as a request/response. 64-byte requests are sent out and three different sized responses are returned (64, 1518 and 9216 bytes). A total of 2 ports are used for east-west traffic. One port is set as east and one port is set as west. The response traffic is further broken down with weights of 1/2/1 for 64/1518/9216 byte frames for the three response sizes.

East-West iSCSI traffic is set up as a request/response with the same east and west port used in each direction. The request is 64 bytes and the response is 9216 bytes.

East-West Microsoft Exchange traffic is set up with the same east and west port. The request and response are both 1518 bytes and set at 70% of line rate

Each direction sends at 70% of line rate.

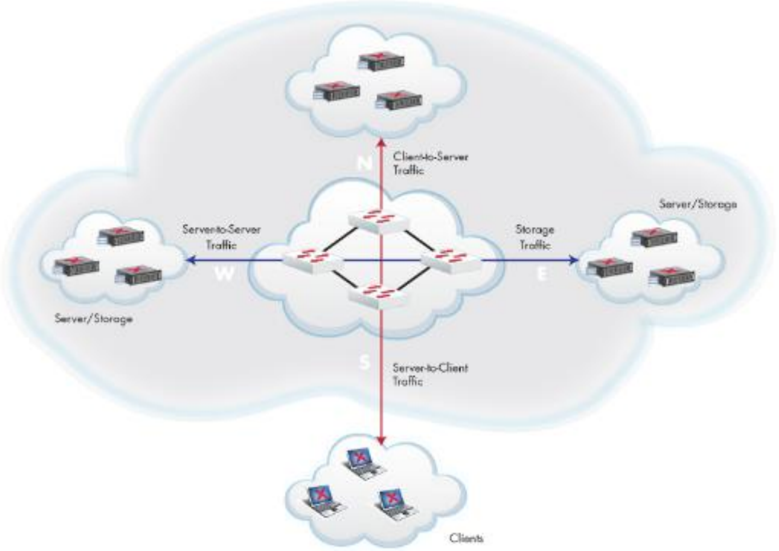

Figure 421. North-South client-server traffic and East-West server/storage traffic across a Data Center Fabric

Step-by-Step Instructions

1. Click **Add Ports** to open the **Port Selection** Window.
2. Click **Add Chassis** and enter the IP Address or name for your IXIA chassis.
3. Click **OK** to accept. Expand chassis and cards and select **four test ports** you want to use in this test.
4. Click **Add ports** and then **OK** to add the ports to your test configuration.

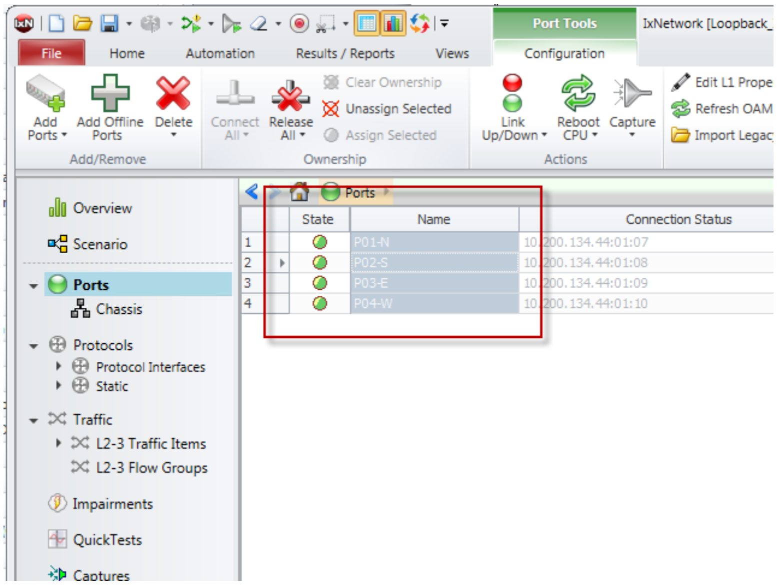

	State	Name	Connection Status
1	●	P01-N	10.200.134.44:01:07
2	●	P02-S	10.200.134.44:01:08
3	●	P03-E	10.200.134.44:01:09
4	●	P04-W	10.200.134.44:01:10

Figure 422. Adding four ports for N, S, E, W traffic

5. Click **QuickTests** in the left pane.
6. Click **Add QuickTests** in the ribbon to open the **Test Selection** wizard.

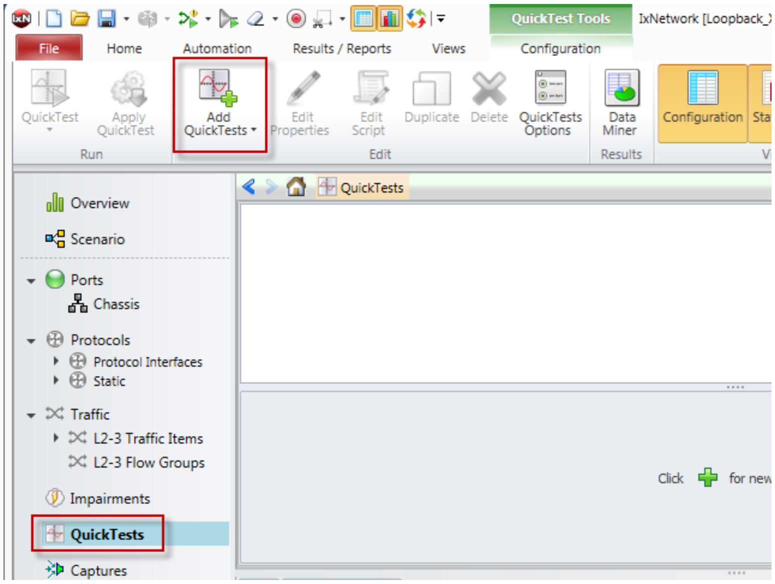

Figure 423. Adding a new QuickTest

7. Expand the **Converged Data Center** folder and select **CloudPerf** test.
8. Click the **Next** button to continue.

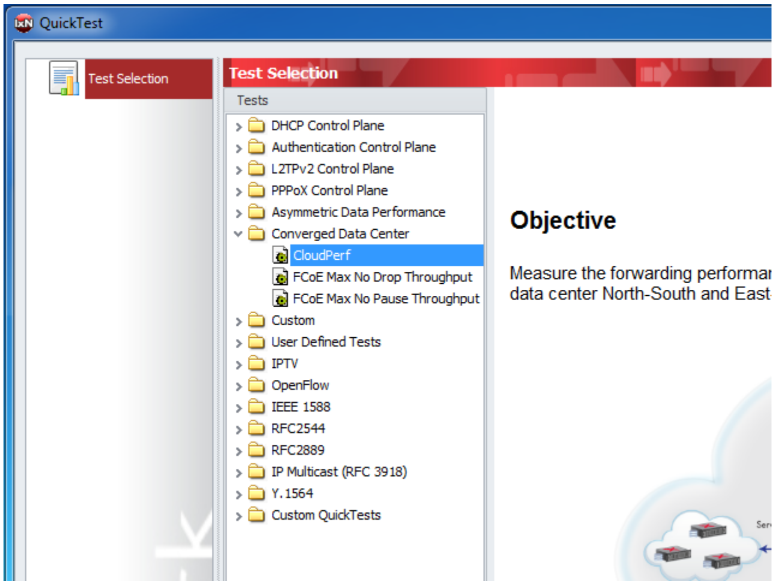

- DHCP Control Plane
- Authentication Control Plane
- L2TPv2 Control Plane
- PPPoX Control Plane
- Asymmetric Data Performance
- Converged Data Center
  - CloudPerf
  - FCoE Max No Drop Throughput
  - FCoE Max No Pause Throughput
- Custom
- User Defined Tests
- IPTV
- OpenFlow
- IEEE 1588
- RFC2544
- RFC2889
- IP Multicast (RFC 3918)
- Y.1564
- Custom QuickTests

Objective

Measure the forwarding performar data center North-South and East

Figure 424. Selecting Converged Data Center CloudPerf QuickTest

9. In the **Ports** view, make sure the **Include in Test** checkbox is selected for each port.
10. Change the **Port Role** for the **four test ports** so that there is one each of **North**, **South**, **East**, and **West** roles.
11. Click the **Next** button to continue.

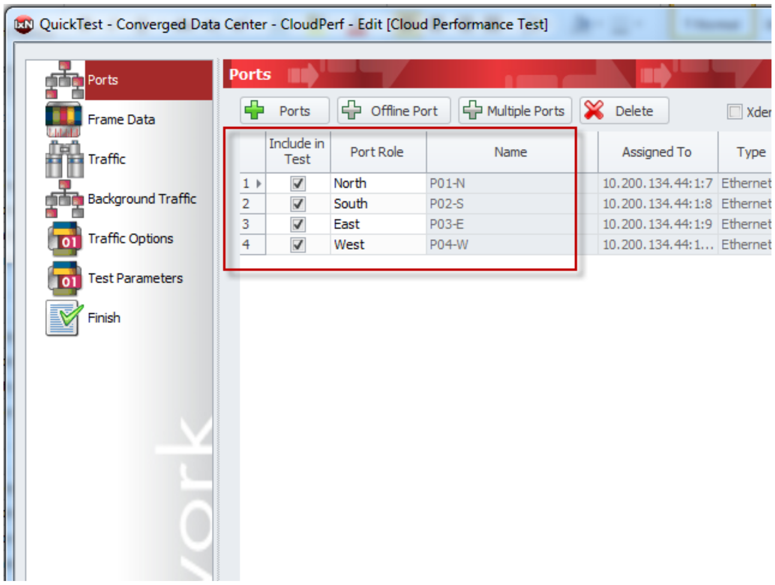

	Include in Test	Port Role	Name	Assigned To	Type
1	☑	North	P01-N	10.200.134.44:1:7	Ethernet
2	☑	South	P02-S	10.200.134.44:1:8	Ethernet
3	☑	East	P03-E	10.200.134.44:1:9	Ethernet
4	☑	West	P04-W	10.200.134.44:1...	Ethernet

Figure 425. Assigning N, S, E, W port roles

12. In the **Frame Data** view, change the **Select type of traffic:** dropdown box value to **MAC**.
13. Click **Next** button to continue.

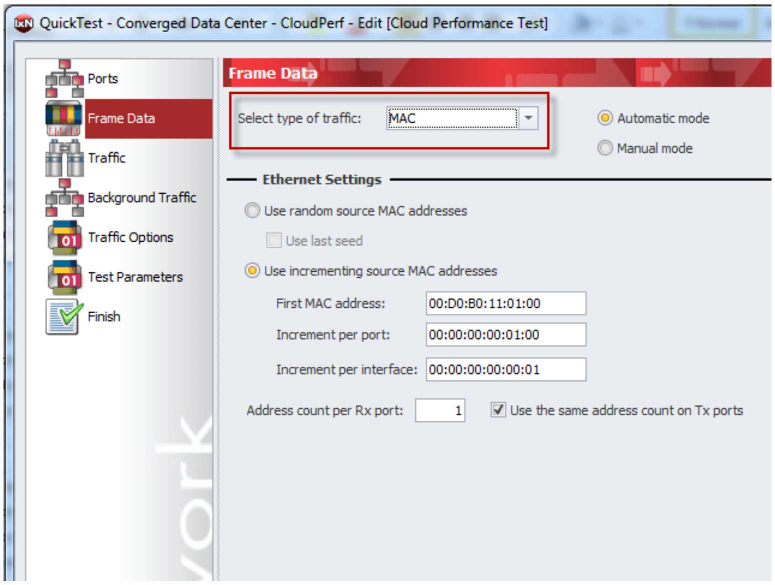

QuickTest - Converged Data Center - CloudPerf - Edit [Cloud Performance Test]

Frame Data

Select type of traffic: MAC

● Automatic mode
○ Manual mode

Ethernet Settings

○ Use random source MAC addresses
□ Use last seed
● Use incrementing source MAC addresses

First MAC address: 00:D0:B0:11:01:00
Increment per port: 00:00:00:00:01:00
Increment per interface: 00:00:00:00:00:01

Address count per Rx port: 1 ☑ Use the same address count on Tx ports

Figure 426. Selecting Frame Data traffic type

14. Click the **Add new** button to open the **Add new traffic profile** panel.
15. Select **Profile direction** value **North-South**.
16. Select **Packet type** value **HTTP**.
17. Select **North role port** checkbox in **Response ports** list.
18. Select **South role port** checkbox in **Request ports** list.

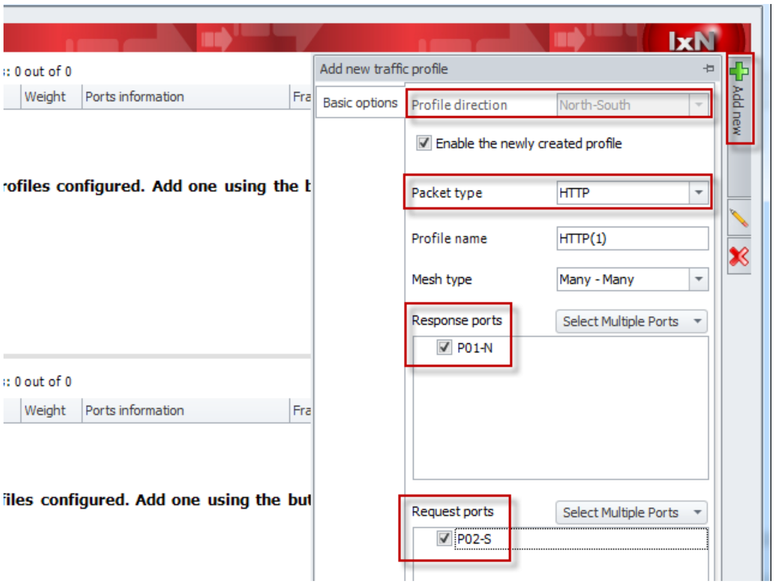

Add new traffic profile

Basic options

Profile direction: North-South
☑ Enable the newly created profile
Packet type: HTTP
Profile name: HTTP(1)
Mesh type: Many - Many
Response ports: Select Multiple Ports
☑ P01-N
Request ports: Select Multiple Ports
☑ P02-S

Figure 427. Adding a new traffic profile

19. Select **Configure advanced options** checkbox to see more configuration tabs.

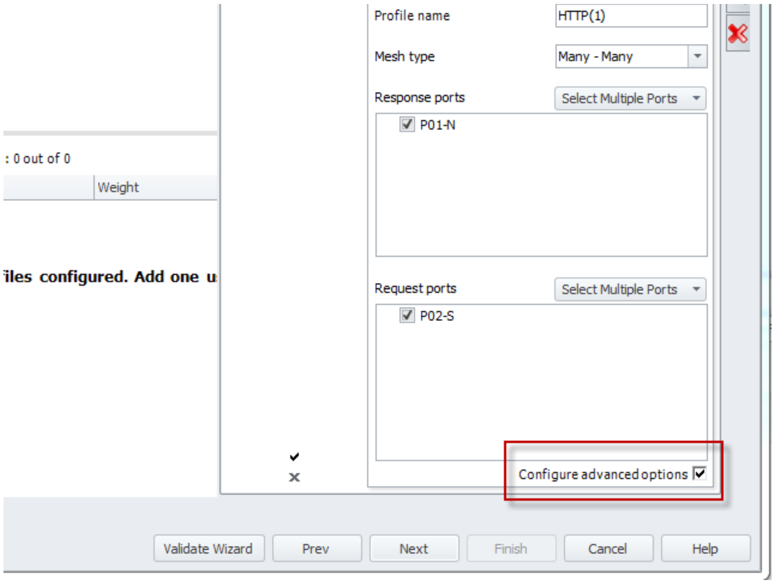

Profile name: HTTP(1)

Mesh type: Many - Many

Response ports: Select Multiple Ports
☑ P01-N

Request ports: Select Multiple Ports
☑ P02-S

Configure advanced options ☑

Figure 428. Configuring advanced options

20. Select **Frame size settings** tab.
21. Select **Request size** value **83**.
22. Select **Response size** value **305**.

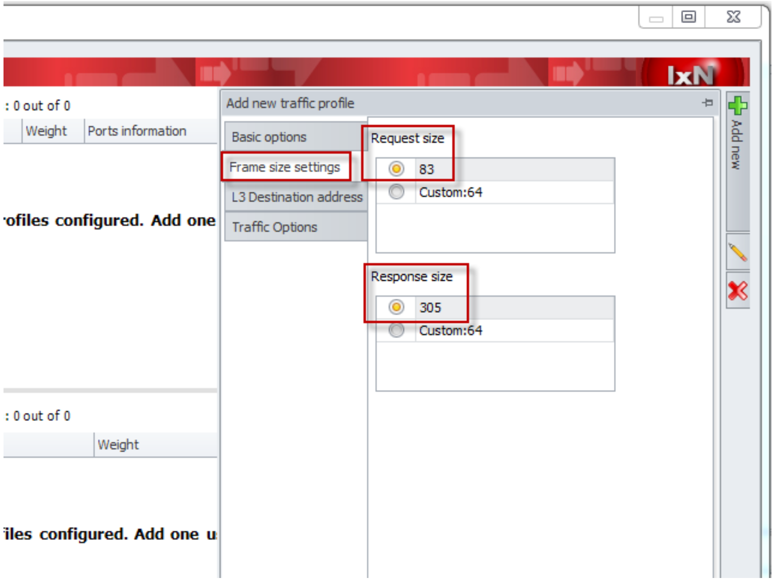

Add new traffic profile

Request size
● 83
○ Custom:64

Response size
● 305
○ Custom:64

Figure 429. Setting request and response frame sizes

23. Click the **Traffic options** tab.
24. Set the **Weight:** value to **1**.

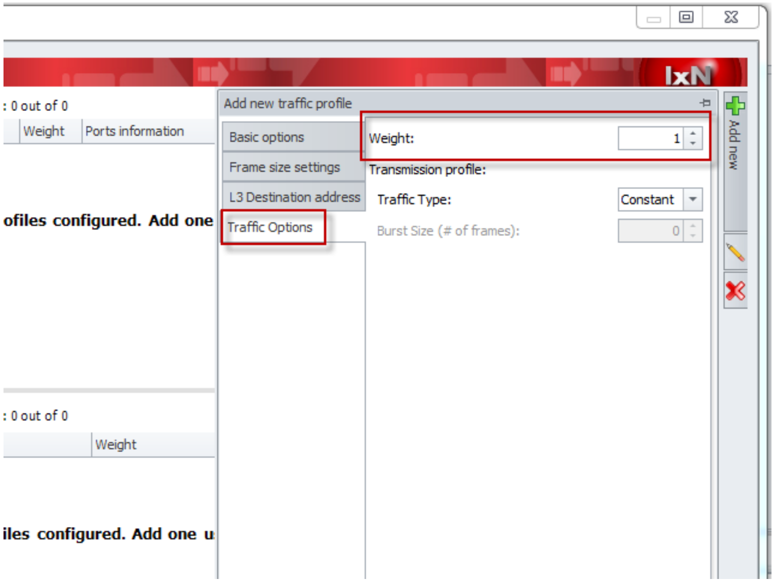

Add new traffic profile

Weight: 1

Transmission profile:

Traffic Type: Constant

Burst Size (# of frames): 0

Figure 430. Setting traffic profile weighting

25. Select **checkmark** icon at bottom of **Add new traffic profile** panel to accept the configuration and add the profile to the list.

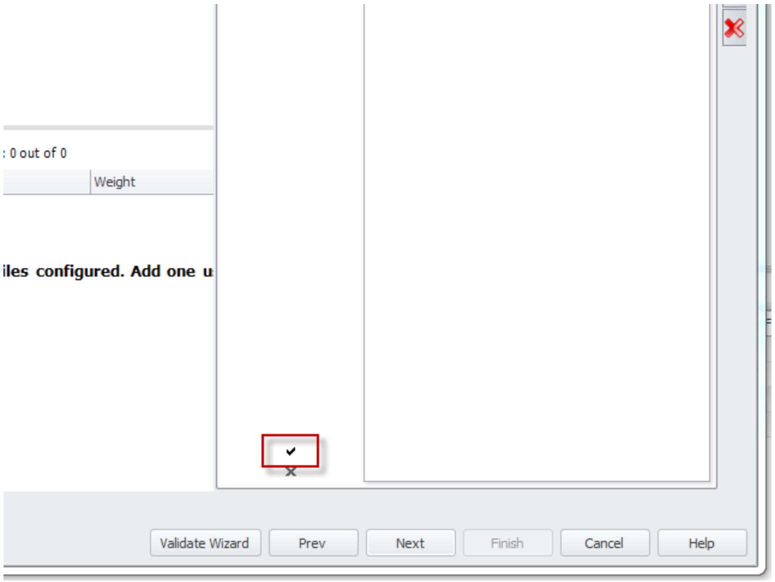

Figure 431. Accepting new traffic profile

26. Click the **Add new** button to open the **Add new traffic profile** panel.
27. Select **Profile direction** value **North-South**.
28. Select **Packet type** value **Youtube**.
29. Select **North role port** checkbox in **Response ports** list.
30. Select **South role port** checkbox in **Request ports** list.

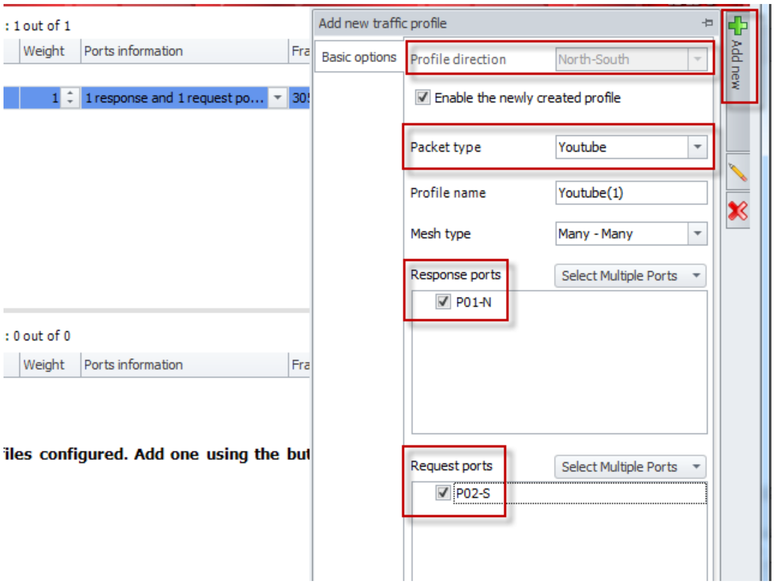
Add new traffic profile

Basic options

Profile direction: North-South

☑ Enable the newly created profile

Packet type: Youtube

Profile name: Youtube(1)

Mesh type: Many - Many

Response ports: Select Multiple Ports
☑ P01-N

Request ports: Select Multiple Ports
☑ P02-S

Figure 432. Adding new traffic profile

31. Select **Configure advanced options** checkbox to see more configuration tabs.
32. Click **Frame size settings** tab.
33. Click **Request size** value as **500**.
34. Click **Response size** value as **IMIX**.
35. Select **Size/Weight** pairs of **1518/5**, **512/2**, and **64/1**.

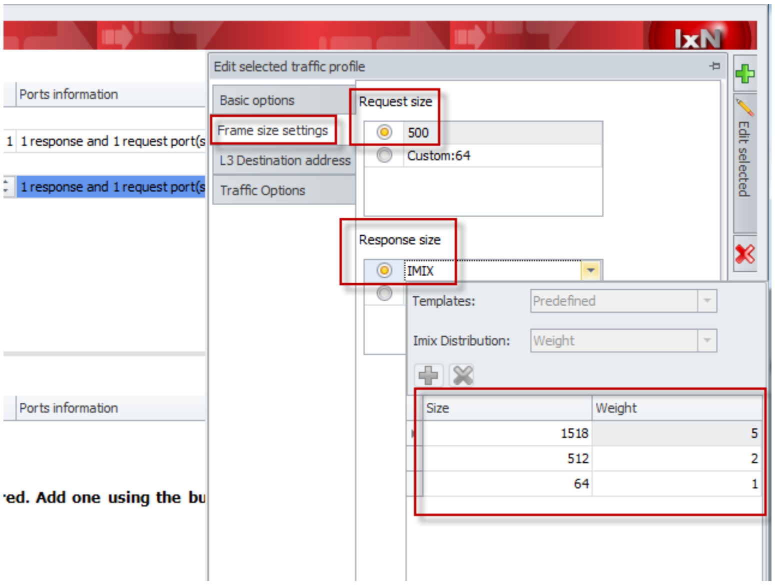
Edit selected traffic profile

Basic options | Frame size settings | L3 Destination address | Traffic Options

Request size
● 500
○ Custom:64

Response size
● IMIX
○

Templates: Predefined

Imix Distribution: Weight

Size	Weight
1518	5
512	2
64	1

Figure 433. Setting request and response frame sizes

36. Click the **Traffic options** tab.

37. Set the **Weight:** value to **1**.
38. Select **checkmark** icon at bottom of **Add new traffic profile** panel to accept the configuration and add the profile to the list.

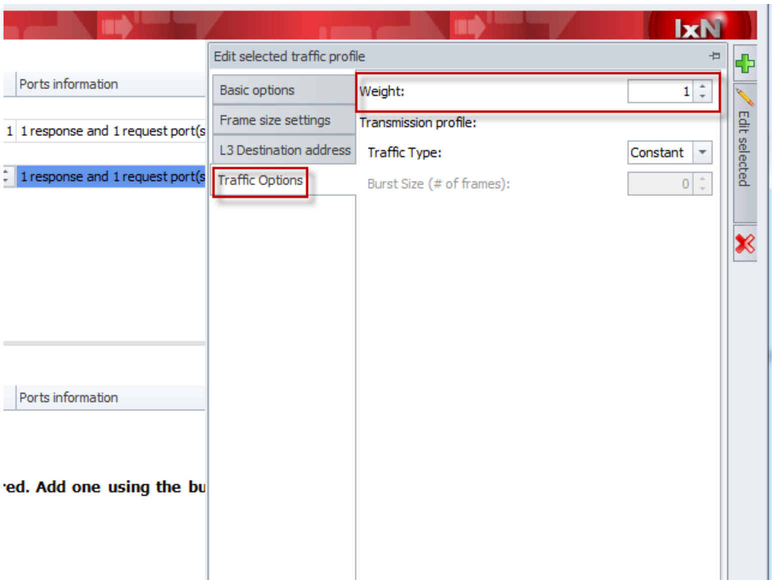

Weight: 1

Transmission profile:

Traffic Type: Constant

Burst Size (# of frames): 0

Figure 434. Setting traffic profile weighting

39. Click the **Add new** button to open the **Add new traffic profile** panel.
40. Select the **Profile direction** value as **East-West**.
41. Select the **Packet type** value as **Database**.
42. Select **West role port** checkbox in **Response ports** list.
43. Select **East role port** checkbox in **Request ports** list.

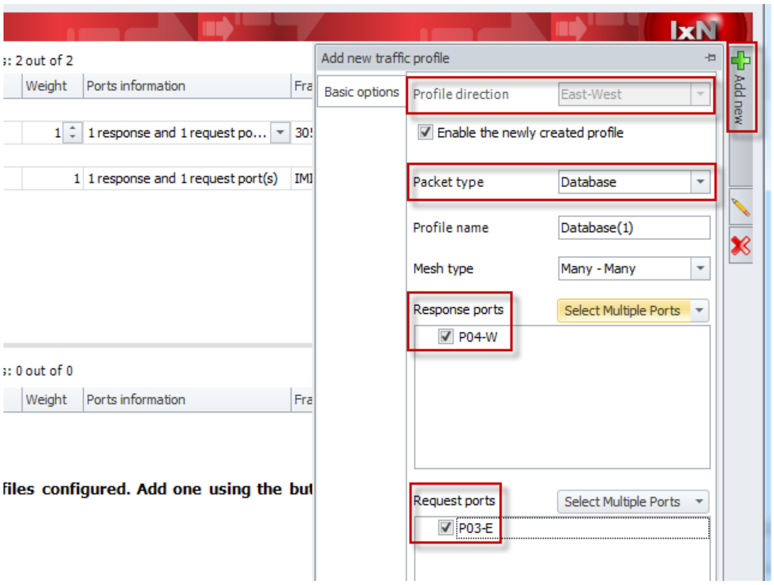

Profile direction: East-West

☑ Enable the newly created profile

Packet type: Database

Profile name: Database(1)

Mesh type: Many - Many

Response ports: Select Multiple Ports
☑ P04-W

Request ports: Select Multiple Ports
☑ P03-E

Figure 435. Adding new traffic profile

44. Select **Configure advanced options** checkbox to see more configuration tabs.
45. Click the **Frame size settings** tab.
46. Click the **Request size** value as **64**.
47. Click the **Response size** value as **Custom: Imix**.
48. Select **Size/Weight** pairs of **64/1**, **1518/2**, and **9216/1**.

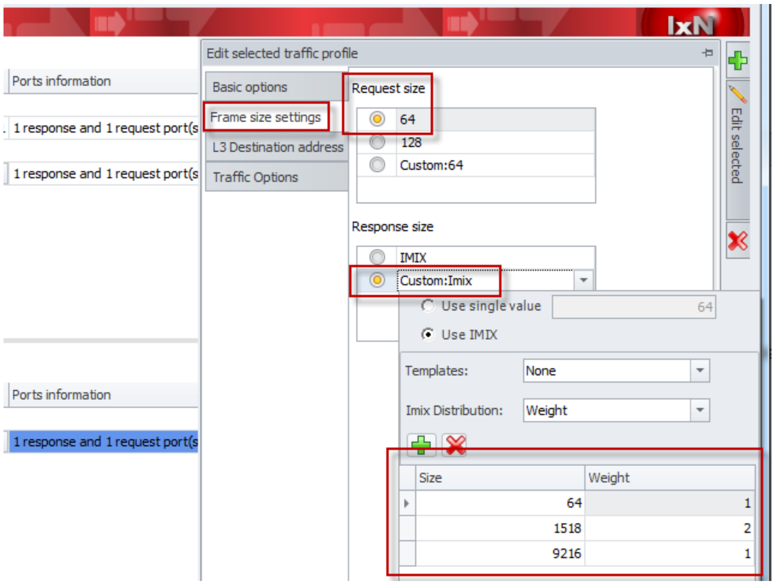
IxN
Edit selected traffic profile
Ports information
1 response and 1 request port(s
1 response and 1 request port(s
Ports information
1 response and 1 request port(s
Basic options
Frame size settings
L3 Destination address
Traffic Options
Edit selected

Request size

- ● 64
- ○ 128
- ○ Custom:64

Response size

- ○ IMIX
- ● Custom:Imix
  - ○ Use single value 64
  - ● Use IMIX

Templates: None

Imix Distribution: Weight

Size	Weight
64	1
1518	2
9216	1

Figure 436. Setting request and response frame sizes

49. Click the **Traffic options** tab.
50. Set the **Weight:** value to **1**.
51. Select **checkmark** icon at bottom of **Add new traffic profile** panel to accept the configuration and add the profile to the list.

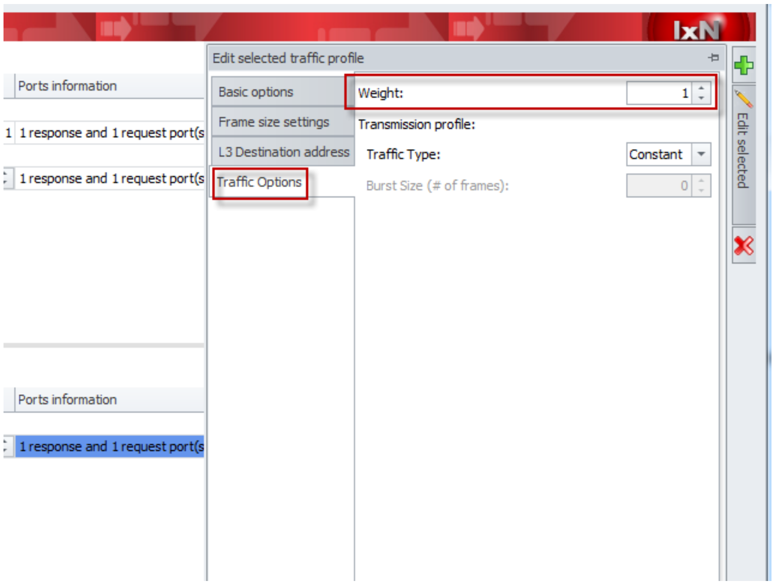
IxN
Edit selected traffic profile
Ports information
1 response and 1 request port(s
1 response and 1 request port(s
Ports information
1 response and 1 request port(s
Basic options
Frame size settings
L3 Destination address
Traffic Options
Edit selected

Weight: 1

Transmission profile:

Traffic Type: Constant

Burst Size (# of frames): 0

Figure 437. Setting traffic profile weighting

52. Click the **Add new** button to open the **Add new traffic profile** panel.
53. Select the **Profile direction** value as **East-West**.
54. Select the **Packet type** value as **iSCSI**.
55. Select **West role port** checkbox in the **Response ports** list.
56. Select **East role port** checkbox in the **Request ports** list.

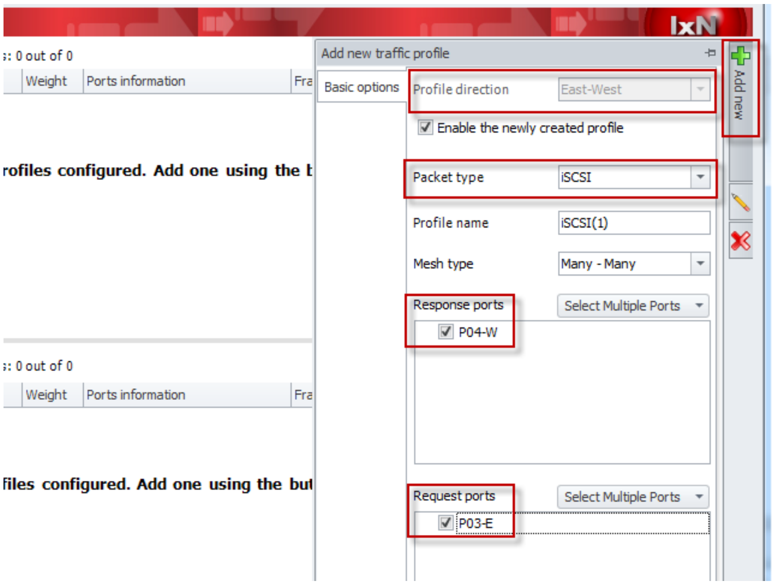

Profile direction: East-West

☑ Enable the newly created profile

Packet type: iSCSI

Profile name: iSCSI(1)

Mesh type: Many - Many

Response ports: Select Multiple Ports
☑ P04-W

Request ports: Select Multiple Ports
☑ P03-E

Figure 438. Adding a new traffic profile

57. Select **Configure advanced options** checkbox to see more configuration tabs.
58. Click the **Frame size settings** tab.
59. Click the **Request size** value as **64**.
60. Click the **Response size** value as **9216**.

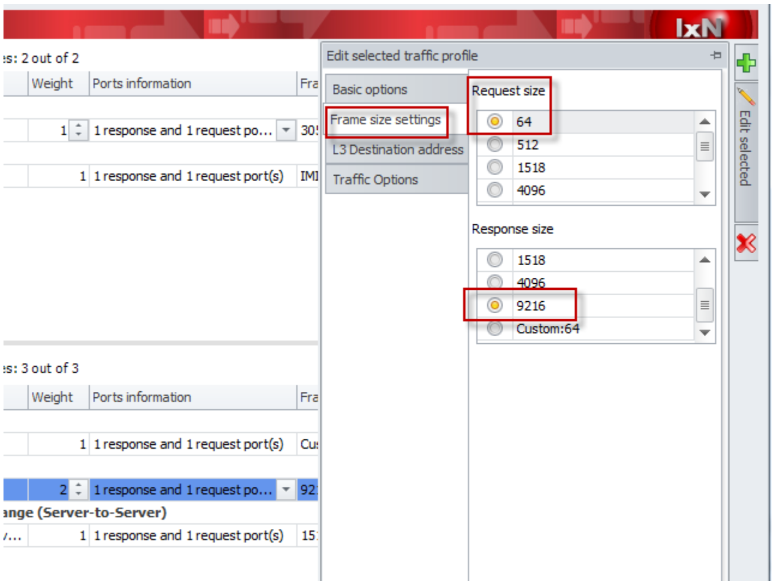

Request size
● 64
○ 512
○ 1518
○ 4096

Response size
○ 1518
○ 4096
● 9216
○ Custom:64

Figure 439. Setting request and response frame sizes

61. Click the **Traffic options** tab.
62. Set the **Weight:** value to **2**.
63. Select the **checkmark** icon at bottom of **Add new traffic profile** panel to accept the configuration and add the profile to the list.

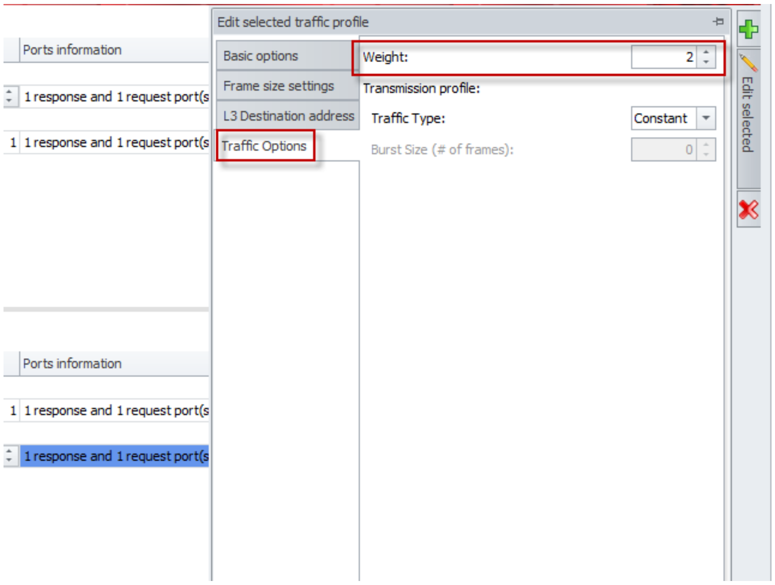

Edit selected traffic profile

Weight: 2

Transmission profile:

Traffic Type: Constant

Burst Size (# of frames): 0

Ports information
1 response and 1 request port(s
1 1 response and 1 request port(s
Ports information
1 1 response and 1 request port(s
1 response and 1 request port(s
Basic options
Frame size settings
L3 Destination address
Traffic Options
Edit selected

Figure 440. Setting traffic profile weighting

64. Click the **Add new** button to open the **Add new traffic profile** panel.
65. Select the **Profile direction** value as **East-West**.
66. Select the **Packet type** value as **Microsoft Exchange (Server-to-Server)**.
67. Select the **West role port** checkbox in the **Response ports** list.
68. Select the **East role port** checkbox in the **Request ports** list.

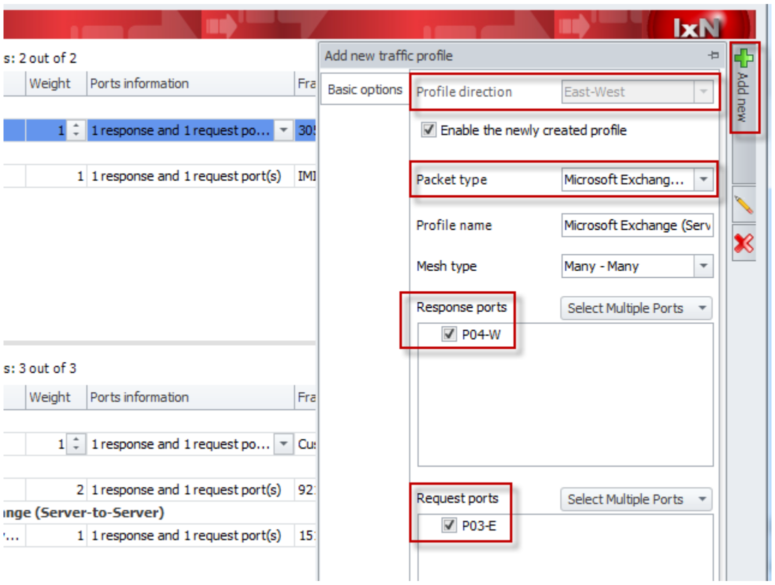

Add new traffic profile

Profile direction East-West

☑ Enable the newly created profile

Packet type Microsoft Exchang…

Profile name Microsoft Exchange (Serv

Mesh type Many - Many

Response ports Select Multiple Ports

☑ P04-W

Request ports Select Multiple Ports

☑ P03-E

IxN
s: 2 out of 2
Weight
Ports information
Fra
1 1 response and 1 request po… 30!
1 1 response and 1 request port(s) IMI
s: 3 out of 3
Weight
Ports information
Fra
1 1 response and 1 request po… Cu:
2 1 response and 1 request port(s) 92
ınge (Server-to-Server)
'… 1 1 response and 1 request port(s) 15
Basic options
Add new

Figure 441. Adding a new traffic profile

69. Select the **Configure advanced options** checkbox to see more configuration tabs.
70. Click the **Frame size settings** tab.
71. Click the **Request size** value as **1518**.
72. Click the **Response size** value as **1518**.

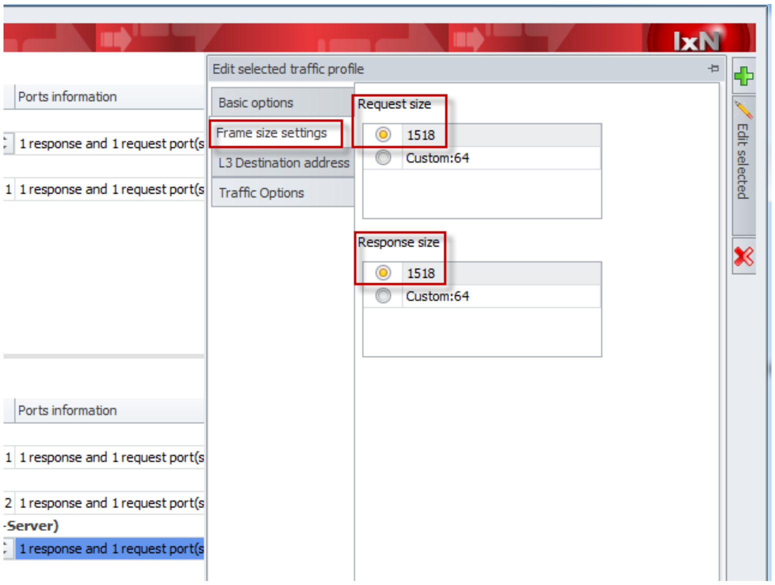
IxN
Edit selected traffic profile
Ports information
1 response and 1 request port(s
1 1 response and 1 request port(s
Basic options
Frame size settings
L3 Destination address
Traffic Options
Ports information
1 1 response and 1 request port(s
2 1 response and 1 request port(s
-Server)
1 response and 1 request port(s
Edit selected

Request size
● 1518
○ Custom:64

Response size
● 1518
○ Custom:64

Figure 442. Setting request and response frame sizes

73. Click the **Traffic options** tab.
74. Set the **Weight:** value as **1**.
75. Select the **checkmark** icon at the bottom of **Add new traffic profile** panel to accept the configuration and add the profile to the list.
76. Click the **Next** button to continue.

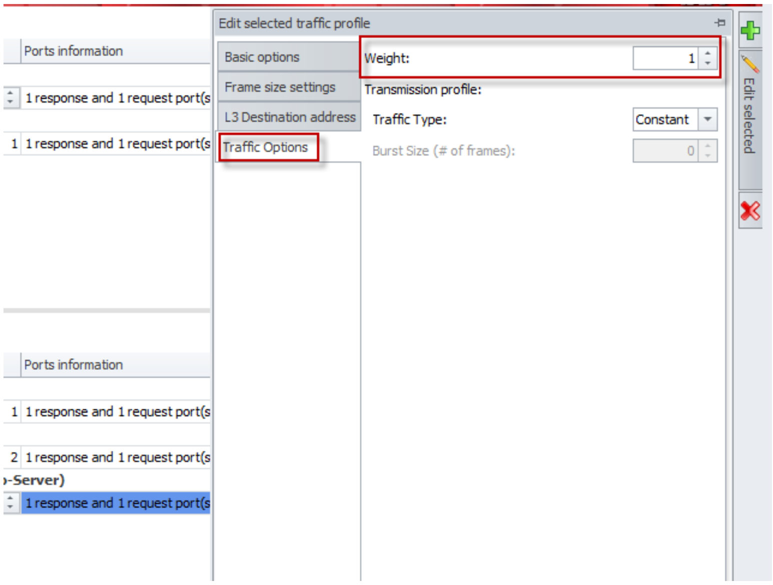
Edit selected traffic profile
Ports information
1 response and 1 request port(s
1 1 response and 1 request port(s
Basic options
Frame size settings
L3 Destination address
Traffic Options
Ports information
1 1 response and 1 request port(s
2 1 response and 1 request port(s
-Server)
1 response and 1 request port(s
Edit selected

Weight: 1
Transmission profile:
Traffic Type: Constant
Burst Size (# of frames): 0

Figure 443. Setting traffic profile weighting

77. (**Optional**) Traffic Items pre-existing on ports not used for N, S, E, W roles may be included in the configuration as background traffic. For this example, background traffic is not specifically needed. Background traffic can be used to place additional stress on the DUT on other ports, while measuring performance of N-S and E-W profiles on the primary test ports.
78. Click the **Next** button to continue.

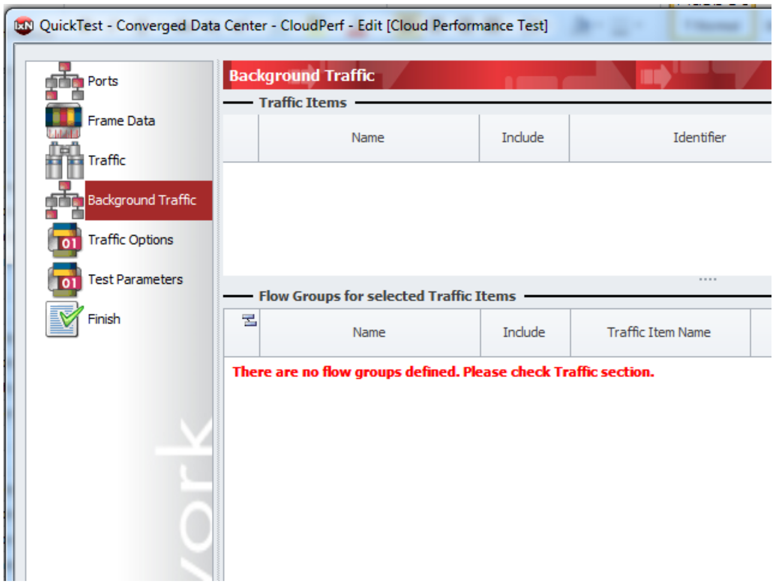
QuickTest - Converged Data Center - CloudPerf - Edit [Cloud Performance Test]
Ports
Frame Data
Traffic
Background Traffic
Traffic Options
Test Parameters
Finish

Background Traffic

Traffic Items

	Name	Include	Identifier

Flow Groups for selected Traffic Items

	Name	Include	Traffic Item Name

There are no flow groups defined. Please check Traffic section.

Figure 444. Background traffic must be run on additional separate ports

79. In the **Learning Frames** section, select the **Frequency** as **Once Per Test**.

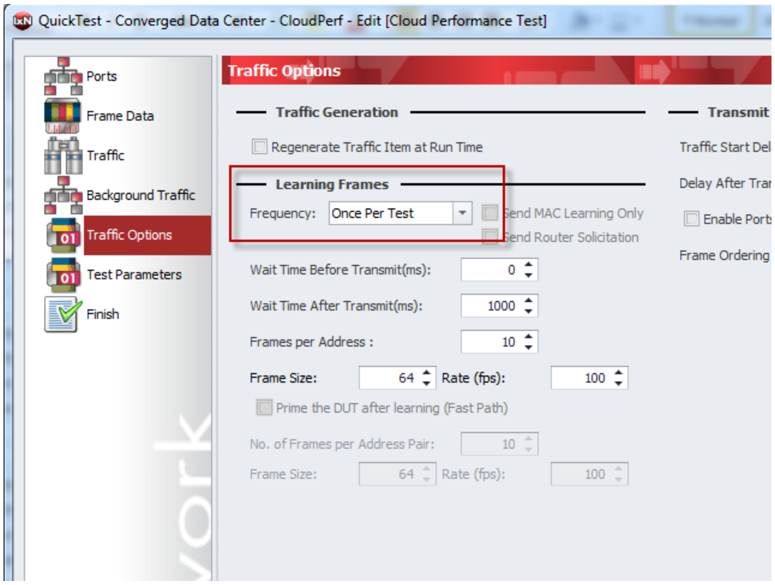
QuickTest - Converged Data Center - CloudPerf - Edit [Cloud Performance Test]
Ports
Frame Data
Traffic
Background Traffic
Traffic Options
Test Parameters
Finish

Traffic Options

Traffic Generation

□ Regenerate Traffic Item at Run Time

Learning Frames

Frequency: Once Per Test

□ Send MAC Learning Only

□ Send Router Solicitation

Wait Time Before Transmit(ms): 0

Wait Time After Transmit(ms): 1000

Frames per Address : 10

Frame Size: 64 Rate (fps): 100

□ Prime the DUT after learning (Fast Path)

No. of Frames per Address Pair: 10

Frame Size: 64 Rate (fps): 100

Transmit

Traffic Start Del

Delay After Trar

□ Enable Ports

Frame Ordering

Figure 445. Setting learning frames frequency

80. Set the **Transmit Traffic Start Delay (s)** value to **2**.
81. Set the **Transmit Delay After Transmit (s)** value to **2**.
82. Click the **Next** button to continue.

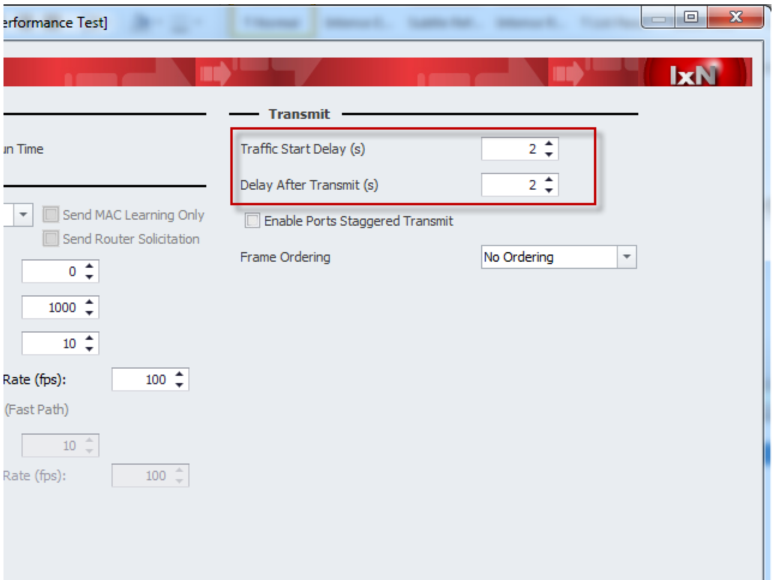
erformance Test]
IxN
Transmit
Traffic Start Delay (s) 2
Delay After Transmit (s) 2
□ Enable Ports Staggered Transmit
Frame Ordering No Ordering
in Time
□ Send MAC Learning Only
□ Send Router Solicitation
0
1000
10
Rate (fps): 100
(Fast Path)
10
Rate (fps): 100

Figure 446. Setting transmit delays

83. Select the **Calculate Latency** checkbox and set value to **Cut Through**.

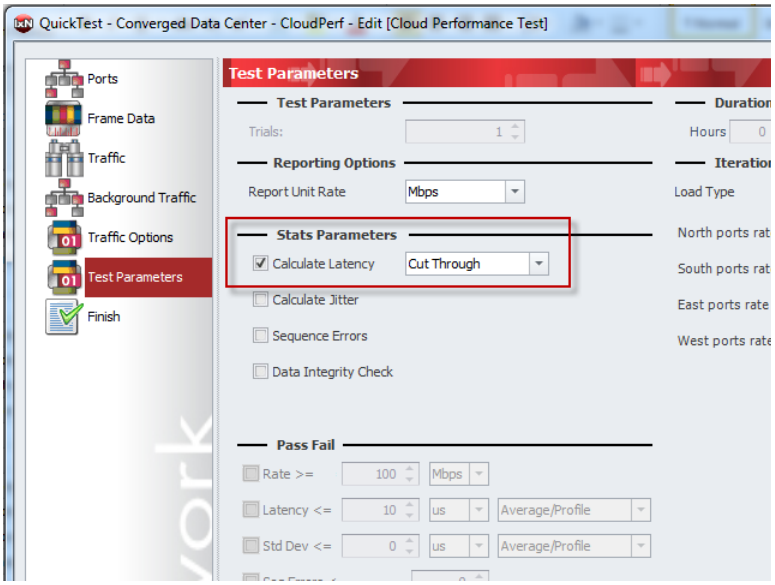

Figure 447. Enabling latency calculation in statistics

84. Set **North ports rate (%Line rate)** value to **70**.
85. Set **South ports rate (%Line rate)** value to **70**.
86. Set **East ports rate (%Line rate)** value to **70**.
87. Set **West ports rate (%Line rate)** value to **70**.
88. Click **Next** button to continue.

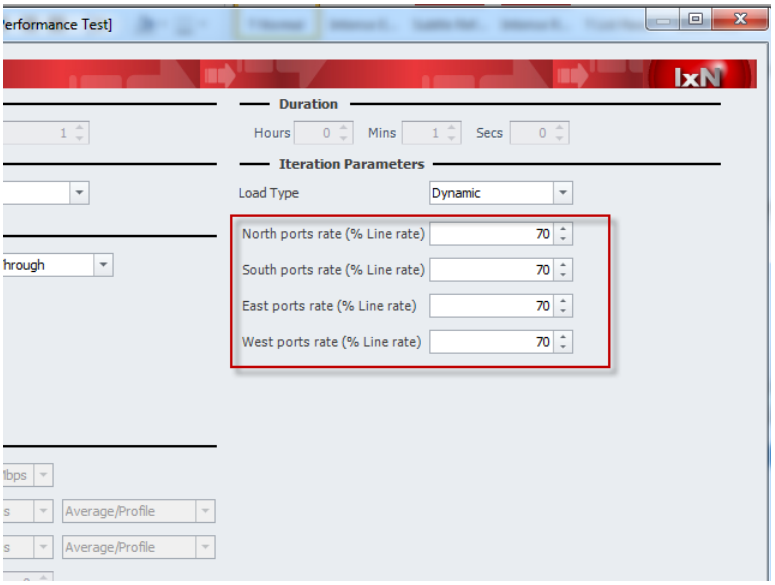

Load Type Dynamic

North ports rate (% Line rate) 70

South ports rate (% Line rate) 70

East ports rate (% Line rate) 70

West ports rate (% Line rate) 70

Figure 448. Setting N, S, E, W traffic rates

89. In the Configuration section, enter the **Name** as **Cloud Performance Test**.
90. Click the **Finish** button to exit the wizard.

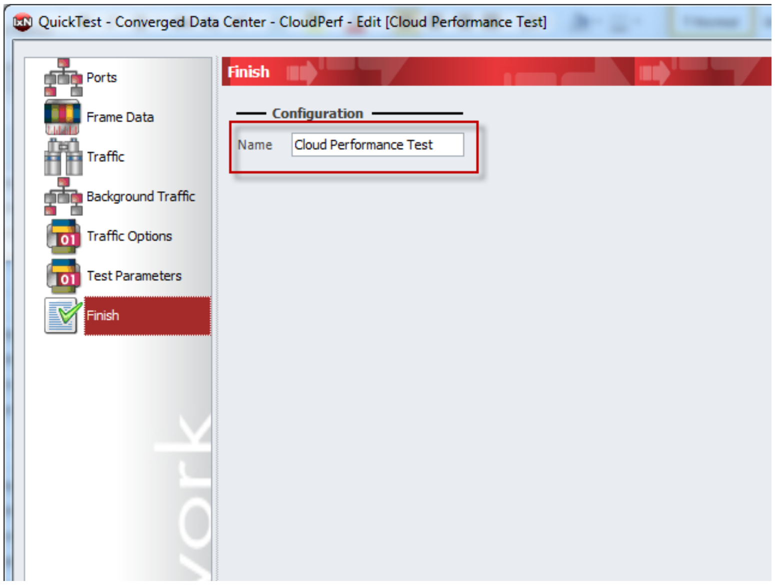

Name Cloud Performance Test

Figure 449. Naming a QuickTest

91. Click the **Cloud Performance** play button to begin executing the QuickTest case.

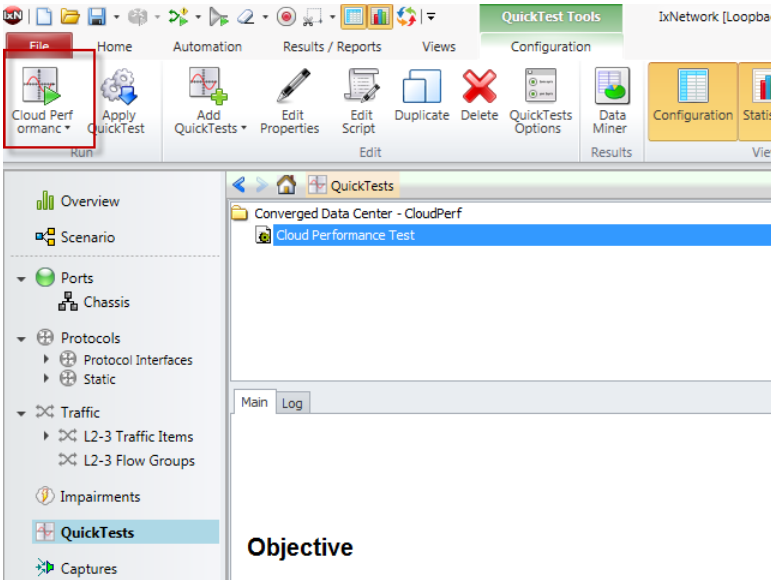

Figure 450. Starting a QuickTest run

92. (**Optional**) Click the **Dynamic Rate Control** tab in order to change values for **North %**, **South %**, **East %**, and **West %** port rates on the fly.

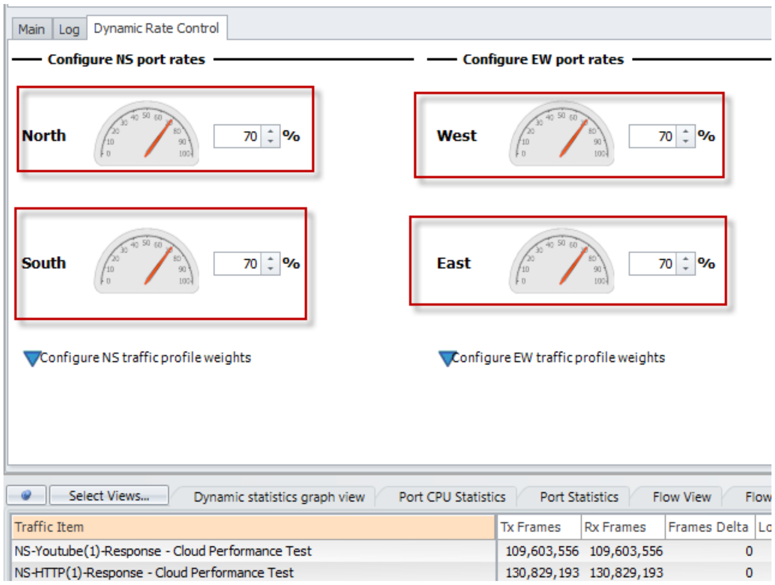

North 70 %

West 70 %

South 70 %

East 70 %

Traffic Item	Tx Frames	Rx Frames	Frames Delta	Lo
NS-Youtube(1)-Response - Cloud Performance Test	109,603,556	109,603,556	0	
NS-HTTP(1)-Response - Cloud Performance Test	130,829,193	130,829,193	0	

Figure 451. Dyanically changing N, S, E, W traffic rates

93. (**Optional**) Expand the **Configure NS traffic profile weights** to view and change the North-South traffic type weightings on the fly.
94. (**Optional**) Expand the **Configure EW traffic profile weights** to view and change the East-West traffic type weightings on the fly.

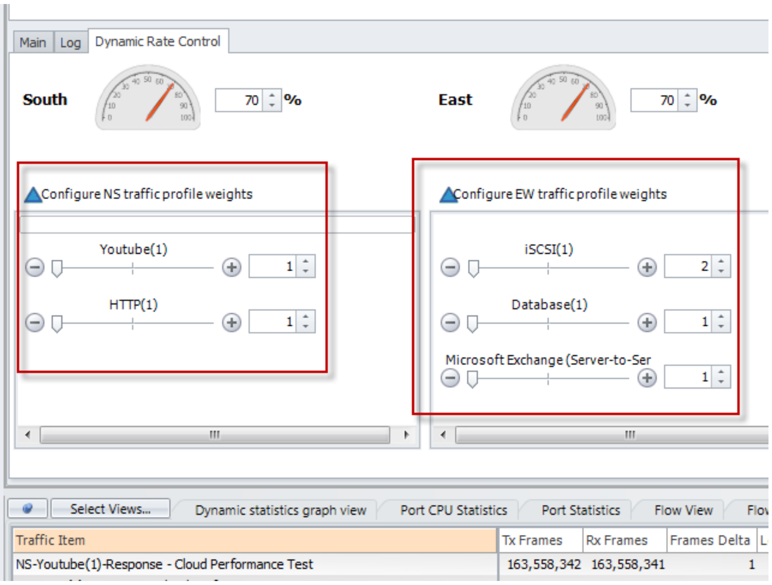

South 70 %

East 70 %

Configure NS traffic profile weights

Youtube(1) 1

HTTP(1) 1

Configure EW traffic profile weights

iSCSI(1) 2

Database(1) 1

Microsoft Exchange (Server-to-Ser 1

Traffic Item	Tx Frames	Rx Frames	Frames Delta	L
NS-Youtube(1)-Response - Cloud Performance Test	163,558,342	163,558,341	1	

Figure 452. Dynamically changing N-S, E-W traffic profile weightings

Test Variables

Performance Variable	Description
North ports rate (%Line rate)	Use this test variable to dynamically vary the total TX throughput generated by North role port(s)
South ports rate (%Line rate)	Use this test variable to dynamically vary the total TX throughput generated by South role port(s)
East ports rate (%Line rate)	Use this test variable to dynamically vary the total TX throughput generated by East role port(s)
West ports rate (%Line rate)	Use this test variable to dynamically vary the total TX throughput generated by West role port(s)
NS traffic profile weights	Use this test variable to dynamically vary the relative weighting of various applicaton traffic flows that make up the North-South TX throughput generated.
EW traffic profile weights	Use this test variable to dynamically vary the relative weighting of various applicaton traffic flows that make up the East-West TX throughput generated.

Troubleshooting and Diagnostics

Issue	Troubleshooting Solution
Loss of throughput on profile traffic	Check DUT settings to ensure: • Correct bandwidth profiles/traffic shaping are enabled and match test port profiles Check test ports to ensure: • Correct application profile weightings • Correct application request and response sizes

Results Analysis

1. Select **Dyamic statistcs graph view** to display a real time graph of throughput and latency values for each traffic profile.

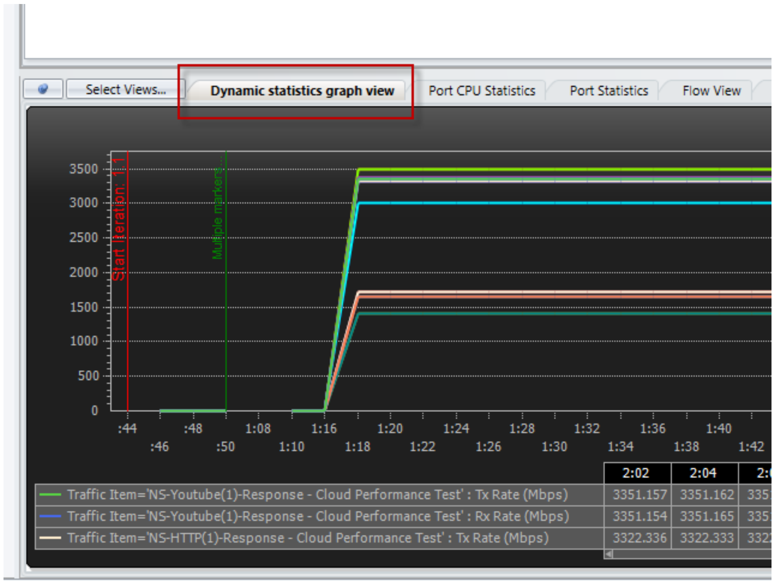

	2:02	2:04	2:
🟩 Traffic Item='NS-Youtube(1)-Response - Cloud Performance Test' : Tx Rate (Mbps)	3351.157	3351.162	335
🟦 Traffic Item='NS-Youtube(1)-Response - Cloud Performance Test' : Rx Rate (Mbps)	3351.154	3351.165	335
⬜ Traffic Item='NS-HTTP(1)-Response - Cloud Performance Test' : Tx Rate (Mbps)	3322.336	3322.333	332

Figure 453. Reviewing real time statistics in the Dynamic statistics graph view

2. Click the **Traffic Item Statistics** view to display detailed packet counts, thoughput and latency values fore each traffic profile.

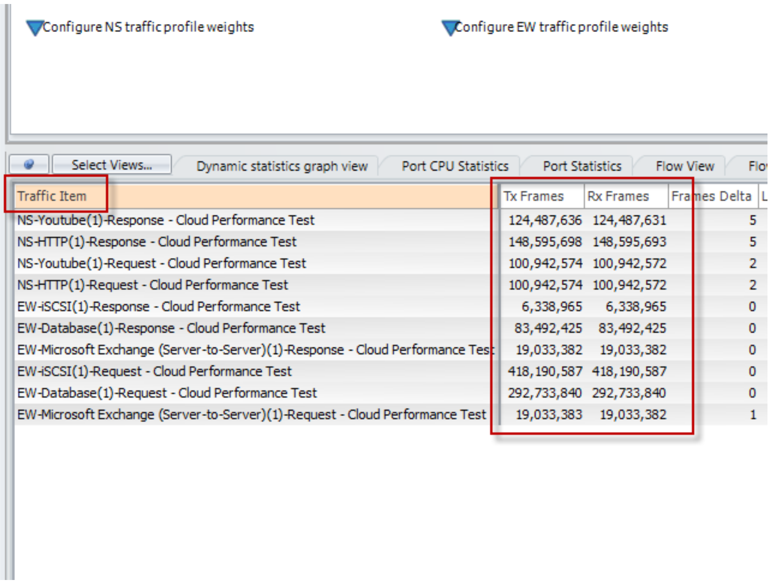
Configure NS traffic profile weights
Configure EW traffic profile weights
Select Views...
Dynamic statistics graph view
Port CPU Statistics
Port Statistics
Flow View
Flo

Traffic Item	Tx Frames	Rx Frames	Frames Delta	L
NS-Youtube(1)-Response - Cloud Performance Test	124,487,636	124,487,631	5	
NS-HTTP(1)-Response - Cloud Performance Test	148,595,698	148,595,693	5	
NS-Youtube(1)-Request - Cloud Performance Test	100,942,574	100,942,572	2	
NS-HTTP(1)-Request - Cloud Performance Test	100,942,574	100,942,572	2	
EW-iSCSI(1)-Response - Cloud Performance Test	6,338,965	6,338,965	0	
EW-Database(1)-Response - Cloud Performance Test	83,492,425	83,492,425	0	
EW-Microsoft Exchange (Server-to-Server)(1)-Response - Cloud Performance Test	19,033,382	19,033,382	0	
EW-iSCSI(1)-Request - Cloud Performance Test	418,190,587	418,190,587	0	
EW-Database(1)-Request - Cloud Performance Test	292,733,840	292,733,840	0	
EW-Microsoft Exchange (Server-to-Server)(1)-Request - Cloud Performance Test	19,033,383	19,033,382	1	

Figure 454. Reviewing real time statistics in the Traffic Item Statistics view

3. Click the **Cloud Performance Test** stop button to terminate the test.

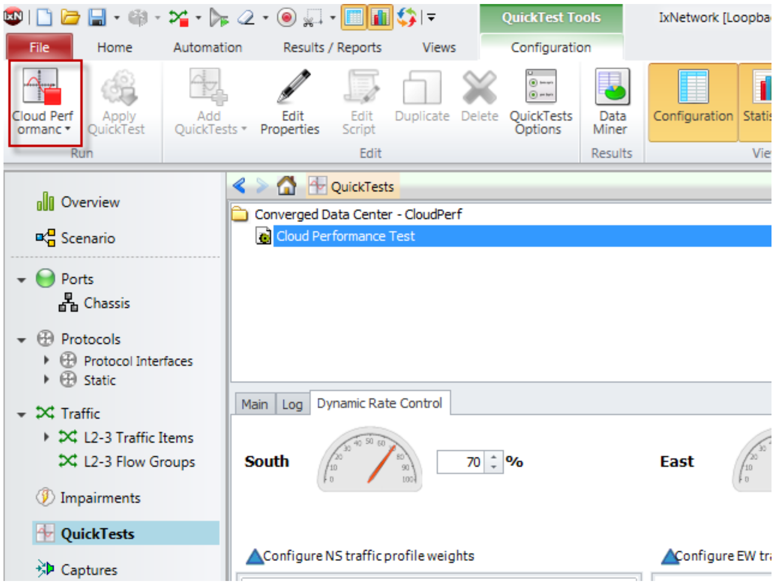

Figure 455. Stopping a QuickTest

4. The **Data Miner** results window opens with a final snapshot of the traffic profile statistics that can be used for post processing or comparison against previous run results.

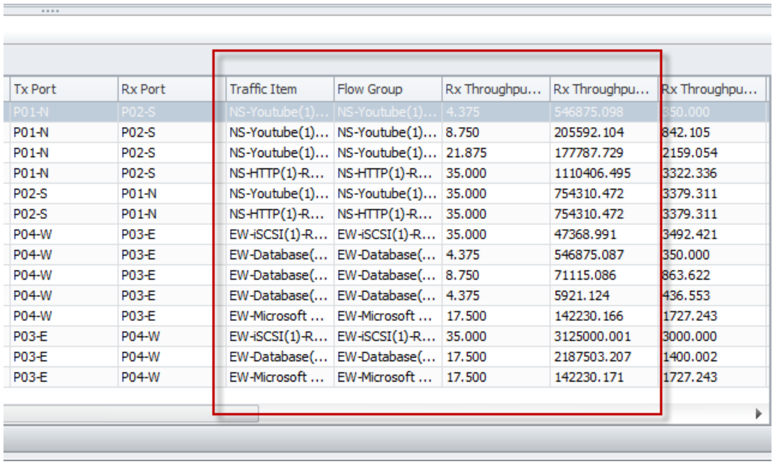

Tx Port	Rx Port	Traffic Item	Flow Group	Rx Throughpu...	Rx Throughpu...	Rx Throughpu...
P01-N	P02-S	NS-Youtube(1)...	NS-Youtube(1)...	4.375	546875.098	350.000
P01-N	P02-S	NS-Youtube(1)...	NS-Youtube(1)...	8.750	205592.104	842.105
P01-N	P02-S	NS-Youtube(1)...	NS-Youtube(1)...	21.875	177787.729	2159.054
P01-N	P02-S	NS-HTTP(1)-R...	NS-HTTP(1)-R...	35.000	1110406.495	3322.336
P02-S	P01-N	NS-Youtube(1)...	NS-Youtube(1)...	35.000	754310.472	3379.311
P02-S	P01-N	NS-HTTP(1)-R...	NS-HTTP(1)-R...	35.000	754310.472	3379.311
P04-W	P03-E	EW-iSCSI(1)-R...	EW-iSCSI(1)-R...	35.000	47368.991	3492.421
P04-W	P03-E	EW-Database(...	EW-Database(...	4.375	546875.087	350.000
P04-W	P03-E	EW-Database(...	EW-Database(...	8.750	71115.086	863.622
P04-W	P03-E	EW-Database(...	EW-Database(...	4.375	5921.124	436.553
P04-W	P03-E	EW-Microsoft ...	EW-Microsoft ...	17.500	142230.166	1727.243
P03-E	P04-W	EW-iSCSI(1)-R...	EW-iSCSI(1)-R...	35.000	3125000.001	3000.000
P03-E	P04-W	EW-Database(...	EW-Database(...	17.500	2187503.207	1400.002
P03-E	P04-W	EW-Microsoft ...	EW-Microsoft ...	17.500	142230.171	1727.243

Figure 456. Reviewing end of test run results in Data Miner

Conclusions

The Cloud Perf QuickTest determines the traffic delivery performance of a data center fabric in forwarding a variety of north/south and east-west traffic in cloud computing applications.

Test Case: Adding Impairment Delay to Data Center Testing

Overview

When executing lab network testing for the datacenter, the tester is striving to achieve a realistic reproduction of live networks within the lab. Many times this test consists of a good mix of background traffic, protocol test traffic and fail-over testing, but the underlying network is pristine and contains no impairments. All production networks contain impairments such as delay and lost packets. The addition of this realism is often overlooked and is the missing link in the creation of a realistic test environment. A key impairment that operators of datacenters are concerned about is delay due to the distance between the primary and backup datacenters.

Data centers have focused on fibre channel as the technology of choice for storage area networks. Data centers also have the need for fail over to remote back up centers. This requirement is many times mandated by regulation or law. As the backup center has to be remote, it introduces distance delay that must be accounted for in the overall test plan. The addition of Ixia’s Network Emulators brings this realism to the lab, enabling more realistic testing. The result is a better understanding and characterization of system and application performance in the event of data center migration or fail over scenarios.

Objective

Demonstrate how delay impacts the system and application performance in the event of a data center fail-over, where the distance between the primary and remote datacenter is 200 KM. The following diagram illustrates the datacenter architecture.

Fibre Channel Backup Architecture

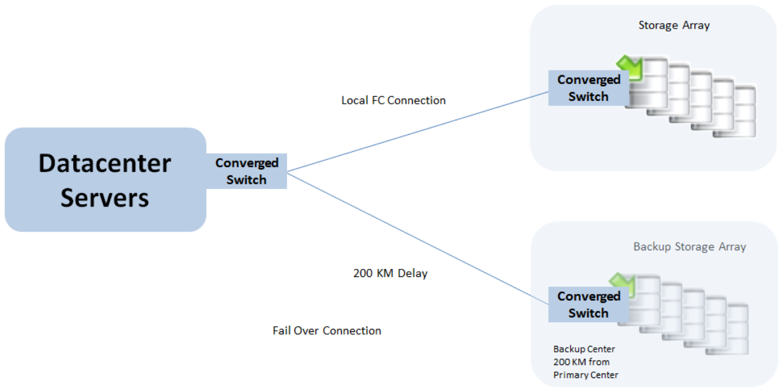

Setup

This test is very simple to set up and consists of only a few steps.

1. Determine performance level objectives.
    - Define the minimum acceptable throughput (bits/sec) expectations for the remote datacenter.
    - Define minimum acceptable application transaction response times (ms or sec).
2. Run the benchmark test to determine performance level without impairment.
3. Position the Network Emulator in-line between the two systems, where delay testing is desired.
4. Configure the Network Emulator to emulate delay. The diagram below illustrates this network setup.

Fibre Channel Backup Test Setup

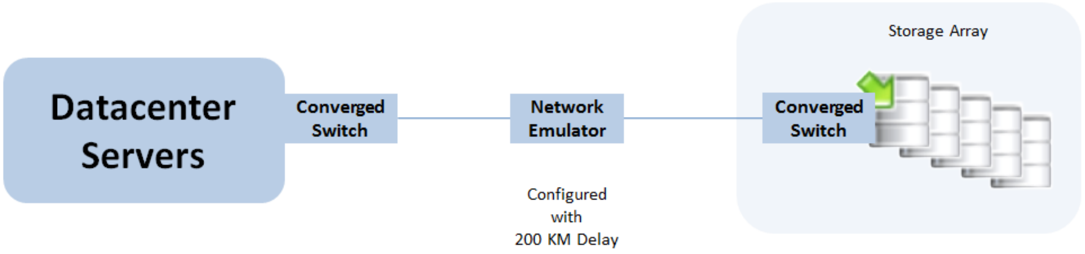

Step-by-Step Instructions

1. Position the Network Emulator in-line between the two systems, where delay testing is desired.

2. Log on to the Network Emulator. The welcome window appears.

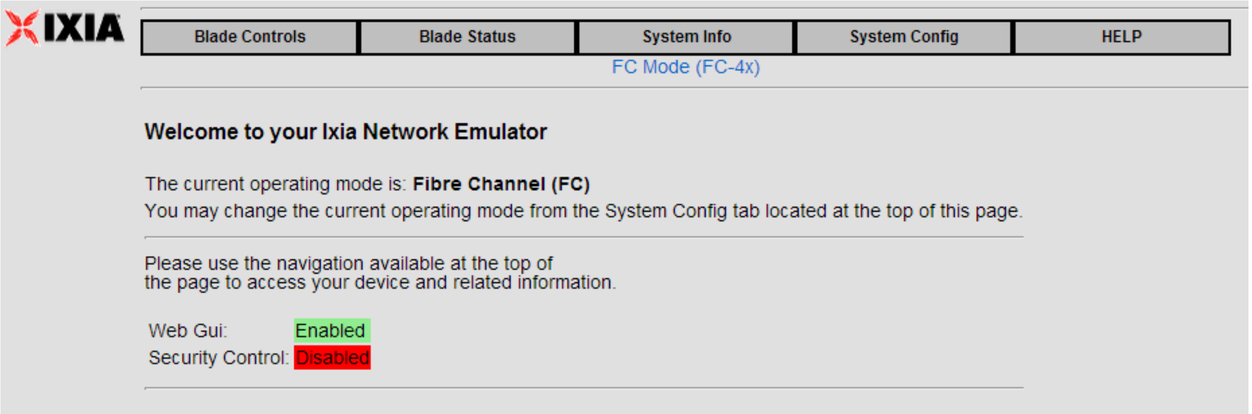

FC Mode (FC-4x)

Welcome to your Ixia Network Emulator

The current operating mode is: **Fibre Channel (FC)**
You may change the current operating mode from the System Config tab located at the top of this page.

Please use the navigation available at the top of
the page to access your device and related information.

Web Gui: Enabled
Security Control: Disabled

3. Click the **Blade Controls** tab. Click **Blade 1,** if not selected and set the **Delay** to 200 km. The delay amount is automatically calculated. The unit of measurement can be in km, ms, or ns. Click the **Set Delay** button to activate the delay. Traffic now has the correct amount of delay to emulate a datacenter that is at a 200 km distance.

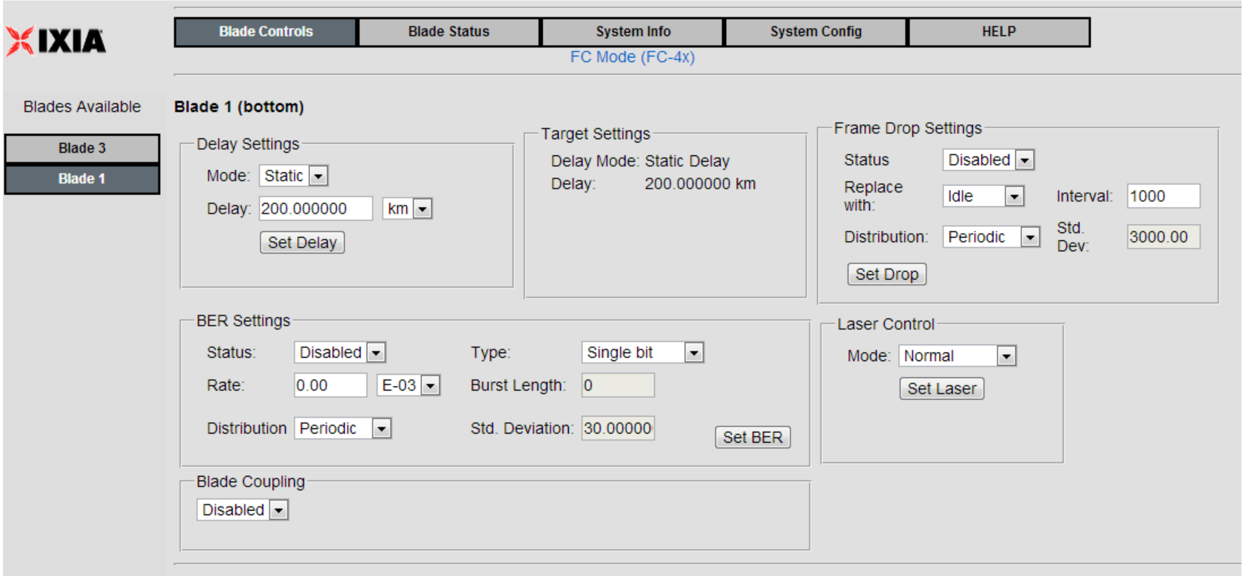

FC Mode (FC-4x)

Blade 1 (bottom)

Delay Settings
Mode: Static
Delay: 200.000000 km

Target Settings
Delay Mode: Static Delay
Delay: 200.000000 km

Frame Drop Settings
Status: Disabled
Replace with: Idle
Interval: 1000
Distribution: Periodic
Std. Dev: 3000.00

BER Settings
Status: Disabled
Type: Single bit
Rate: 0.00 E-03
Burst Length: 0
Distribution Periodic
Std. Deviation: 30.00000

Laser Control
Mode: Normal

Blade Coupling
Disabled

4. Click Blade 3 and repeat the DELAY configuration. BLADE 3 configuration is required, so delay from remote is calculated.

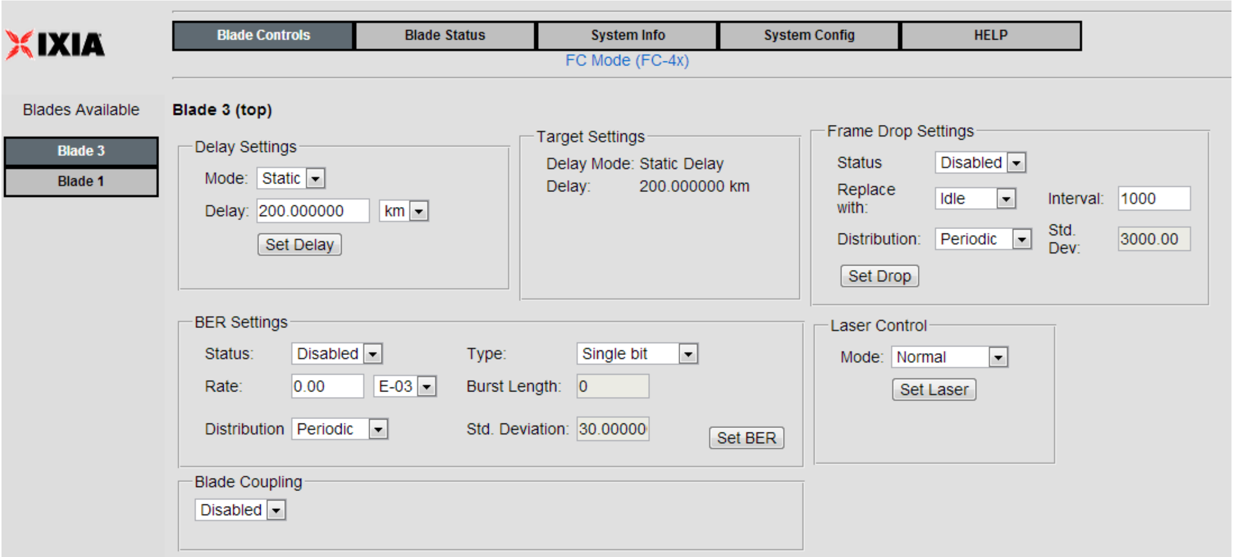

FC Mode (FC-4x)

Blade 3 (top)

Delay Settings
Mode: Static
Delay: 200.000000 km

Target Settings
Delay Mode: Static Delay
Delay: 200.000000 km

Frame Drop Settings
Status: Disabled
Replace with: Idle
Interval: 1000
Distribution: Periodic
Std. Dev: 3000.00

BER Settings
Status: Disabled
Type: Single bit
Rate: 0.00 E-03
Burst Length: 0
Distribution: Periodic
Std. Deviation: 30.00000

Laser Control
Mode: Normal

Blade Coupling
Disabled

Test Variables

Test tool variables

Parameter	Description
Delay	Set appropriate amount of delay
Mode	Static

Results

When delay is added to the benchmark test, verify that overall throughput goals are still achievable when the network accurately reflects the distance delay. Additionally, characterize the impact on any specific application and verify the application's transaction response time goals are still met.

Conclusions

Network Emulation is normally the missing link to realistic network testing and is often ignored in the overall test planning. The procedures described above show a simple method of adding distance delay to the datacenter failover test. If the network is 10G Ethernet, then this test can also be performed with Ixia's 10G impairment solutions. After adding Network Emulation to the datacenter test system, the network reflects a real world environment and reflects better conditions that are found when the product is deployed. The addition of Ixia's ImpairNet or Network Emulator impairment devices can bring this realism to the lab.

Test Case: Performance Measurement of a NAS Target

Overview:

The storage world is changing. Treating Storage as the cold realm of large data sets or corporate backup plans does not work anymore. Today, storage means our photos, movies, emails, presentations, novels, music, and business-critical data. Validated storage saves a lifetime of memories. The following testcase demonstrates different methods in which a large NAS storage can be validated. The testcase covers end to end testing, where we first setup the NAS filer by creating the needed files and folders in it, and later accessing them with different types of workload patterns.

Objective

The test measures the max transactions per second and max throughput per second that can be achieved against a NAS filer.

Setup

The setup requires at least one client port. The CIFS client traffic reaches the DUT either directly or through a switch.

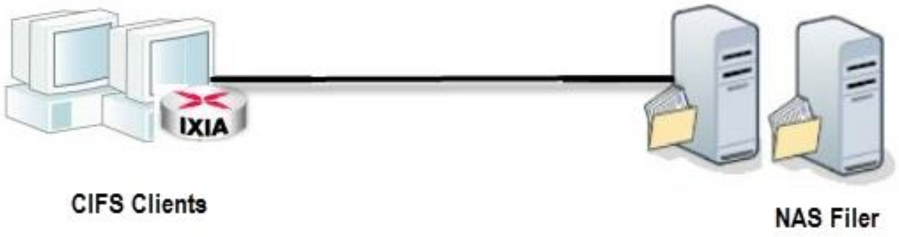

Figure 457. NAS Test Setup

Test Variables

Test Tool Variables

The following test configuration parameters provide the flexibility to create the traffic profile that a device experiences in a production network.

Table - CIFS configuration parameters

Parameters	Description
CIFS clients	1 IP address for setup phase and many more IP address in the Run phase
CIFS client parameters	NTLMv2 Kerberos Authentication Command Chunk Size Lock On or Off Types of Locks Domain-Names
TCP parameters	TCP RX and TX buffer at 4096 bytes
CIFS client command list	• SessionSetup, WriteToFile, ReadFromFile commands • SessionSetup command with several username and passwords entered through playlist or using sequence generators • WriteToFile with different file sizes and types. • Varying Fileoffses in WriteToFile • ReadFromFile with different read offsets • ReadFromFile with different File sizes • Playlist to create different file names, file length, file offsets.

DUT Setup

1. Before running an NAS test, make sure that the target folder in the NAS server has sufficient access privileges. The below screenshots guide you through the process of folder sharing in a windows computer running CIFSv2. Similar processes can be involved to do sharing in linux based systems.

2. **Right-Click** the folder that you want to share and click **Properties**. The **Properties** dialog opens. Click the **Sharing** tab. In the **Network File and Folder Sharing** click the **Share** button(marked in red). Clicking the **Share** button opens a new window that shows the user names having access to the folders. You can either allow everyone or selected users. Also, ensure that the allowed users have both read and write permissions.

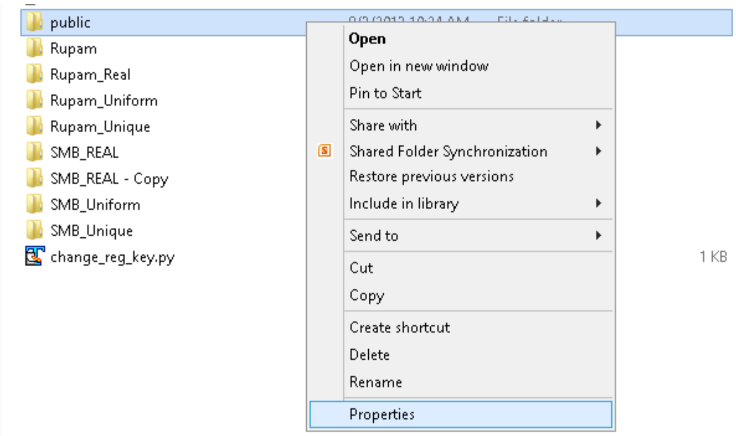

3. To verify that the folder is properly shared, type the ip of the external server in the **Run** command of windows. Once authenticated, it must ideally display the folder shared recently at this target.

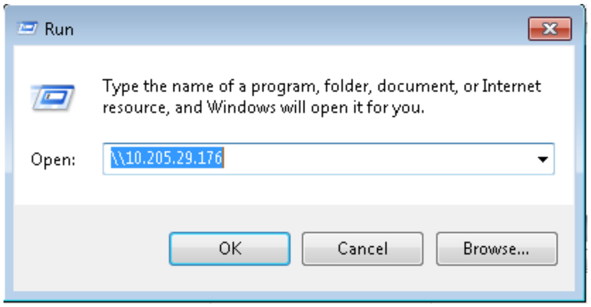

Type the name of a program, folder, document, or Internet resource, and Windows will open it for you.

Open: \\10.205.29.176

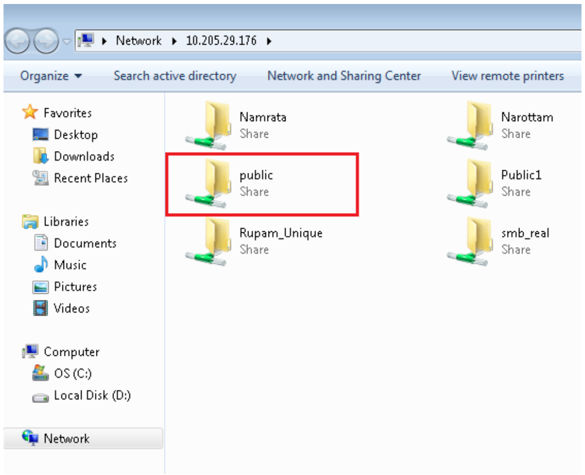

Step-by-Step Instructions

1. Add the client **NetTraffic** object. Configure the **Client** network with total IP count, gateway, and VLAN, if used.

   For a step-by-step workflow, see Appendix A.

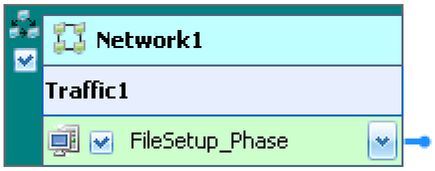

2. The TCP parameters that are used for a specific test type are important when optimizing the test tool. Refer to the **Test Variables** section to set the correct TCP parameters.

There are several other parameters that can be changed. Leave them at their defaults values unless you need to change them for testing requirements.

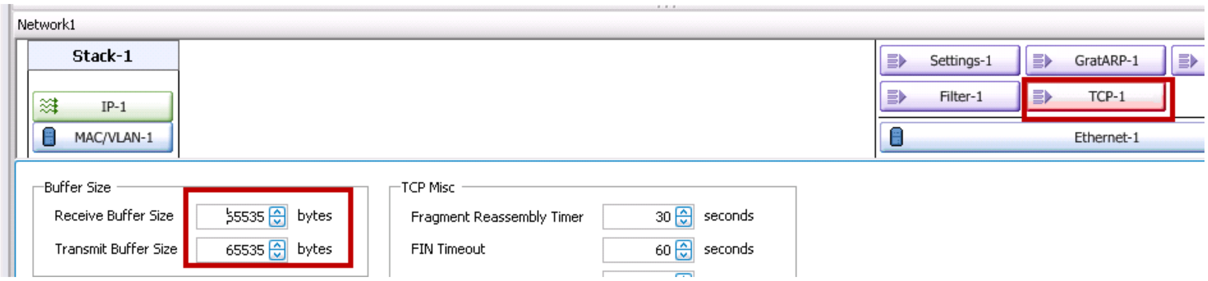

Buffer Size

Receive Buffer Size 65535 bytes

Transmit Buffer Size 65535 bytes

TCP Misc

Fragment Reassembly Timer 30 seconds

FIN Timeout 60 seconds

Figure 458. TCP Buffer Settings Dialogue

3. Configure the CIFS **client**. Add the CIFS **Client Activity** to the client **NetTraffic**.

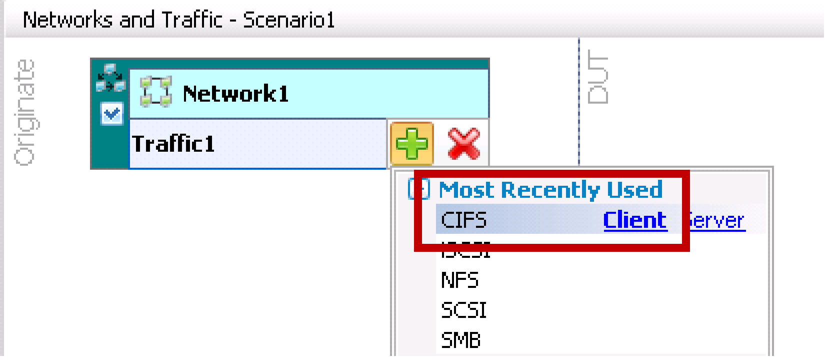

Figure 459. Adding the CIFS plugin

4. Click the **Settings** tab. Change the **CIFS version** to ***CIFSv2*** *and* ***Authentication Mechanism*** *to* ***NTLM****.*

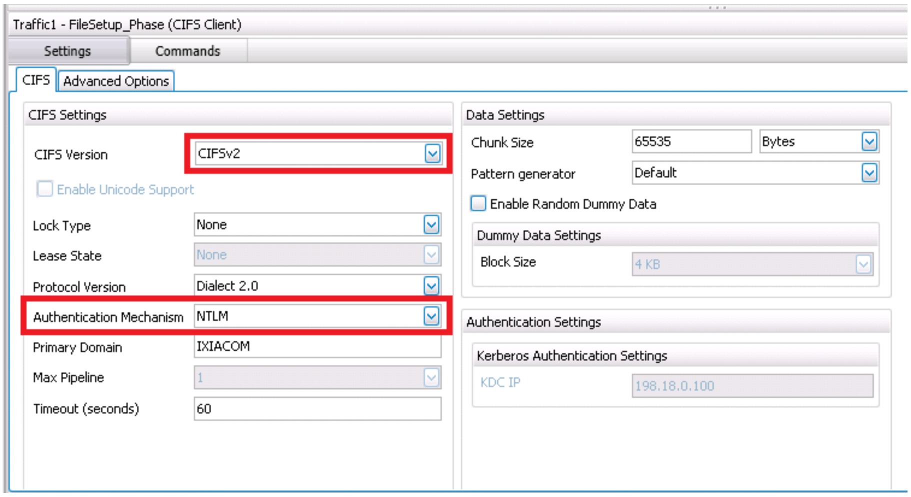

Traffic1 - FileSetup_Phase (CIFS Client)

CIFS Settings

CIFS Version: CIFSv2

□ Enable Unicode Support

Lock Type: None

Lease State: None

Protocol Version: Dialect 2.0

Authentication Mechanism: NTLM

Primary Domain: IXIACOM

Max Pipeline: 1

Timeout (seconds): 60

Data Settings

Chunk Size: 65535 Bytes

Pattern generator: Default

□ Enable Random Dummy Data

Dummy Data Settings

Block Size: 4 KB

Authentication Settings

Kerberos Authentication Settings

KDC IP: 198.18.0.100

Figure 460. Setting Authentication to NTLM

5. Having setup the client networks and the traffic profile, now configure the filesystem at the server. For this particular test, create files in the shared folder at the server. Create the following file distribution in the server that generally resembles a common distribution of small enterprise filer.

File Count	**10**	**90**	**899**	**3000**	**100**	**100**	**20**
File Size	1KB	10KB	50KB	100KB	1MB	100MB	1GB

The file size distribution created at client

There are a total of 4220 files created on the external server, which will be accessed later through other cifs activity.

6. Use the command SessionSetup to create a session with the external server. For the **Server IP** field, set the ip of the external server that is running CIFS**.** Set the **Username** and **Password** correctly to enable the user to gain access to the external server.

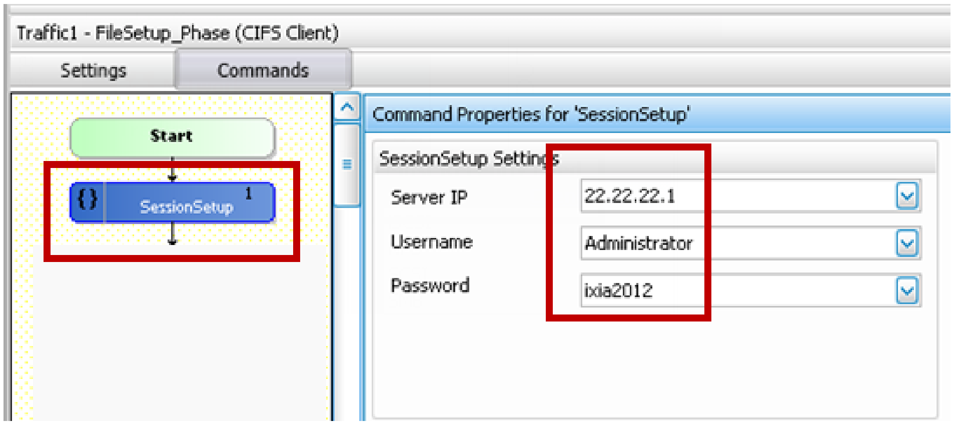

Traffic1 - FileSetup_Phase (CIFS Client)

Command Properties for 'SessionSetup'

SessionSetup Settings

Server IP: 22.22.22.1

Username: Administrator

Password: ixia2012

Figure 461. Adding “Session Setup” command that will log into the NAS

7. As discussed in the setup phase, create all the files that will be later accessed. The first loop creates the first 10 files of size 1KB each. Set the loop count as **10**.

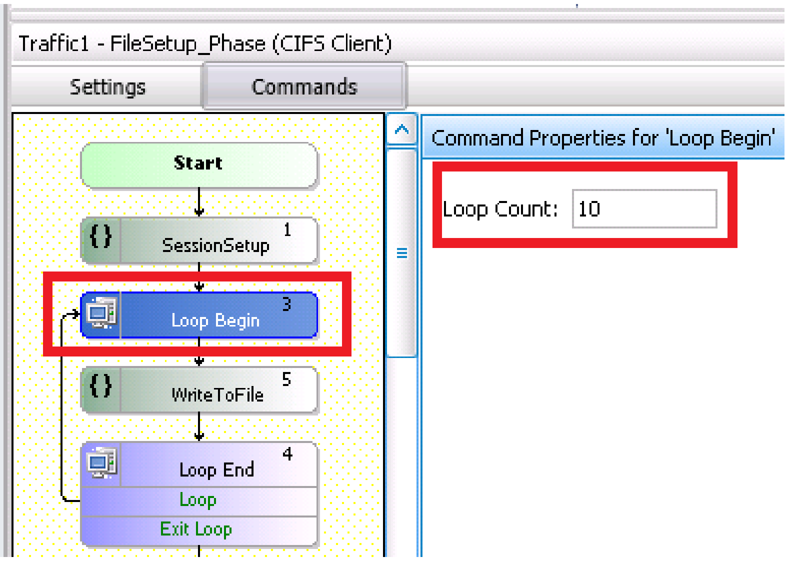

Traffic1 - FileSetup_Phase (CIFS Client)

Command Properties for 'Loop Begin'

Loop Count: 10

Figure 462. Adding loop command . This loops in the next command for exact number of time.

8. Add a **WriteToFile** command from the command dropdown. In the **WriteToFile** Cmd provide the target folder and the name of the files to be created. Use the sequence generator to create ten different files.

   In this example, **public** is the name of the shared folder. 10 loops are added. The **file[0-9].txt** creates 10 files at the server. Loop 1 → file0.txt, loop2 -> file1.txt, loop 2 -> file2.txt……loop 10 -> file10.txt.
9. Set the filesize as 1KB.

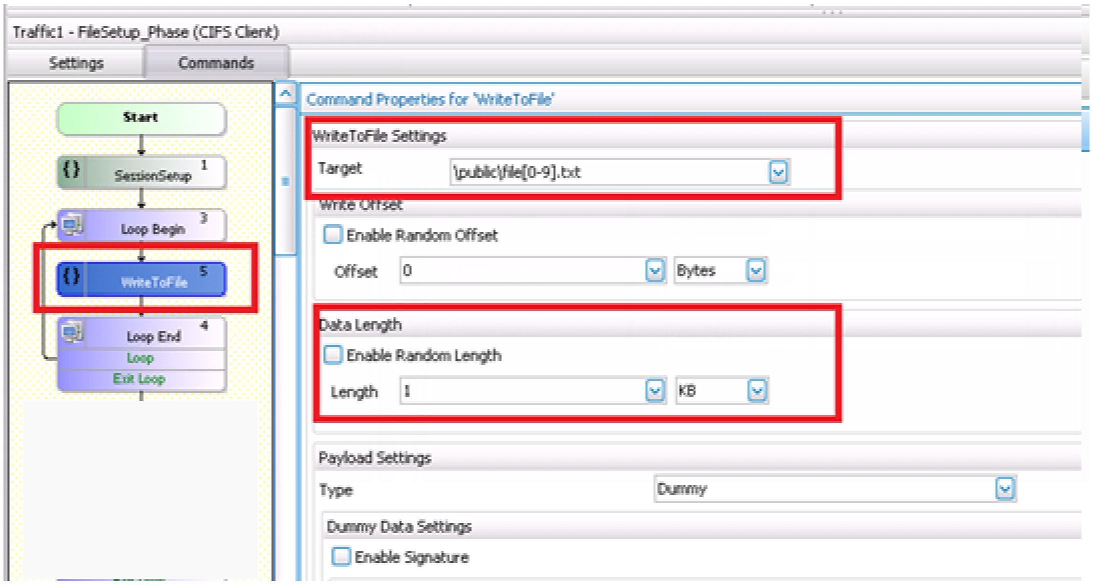
Traffic1 - FileSetup_Phase (CIFS Client)

Command Properties for 'WriteToFile'

WriteToFile Settings

Target \public\file[0-9].txt

Write Offset

□ Enable Random Offset

Offset 0 Bytes

Data Length

□ Enable Random Length

Length 1 KB

Payload Settings

Type Dummy

Dummy Data Settings

□ Enable Signature

Figure 463. The WriteToFile command creates and writes a pre-determined length of data to the file.

10. Similarly add another Loop. Begin with a count of 90 after the end of the first loop.

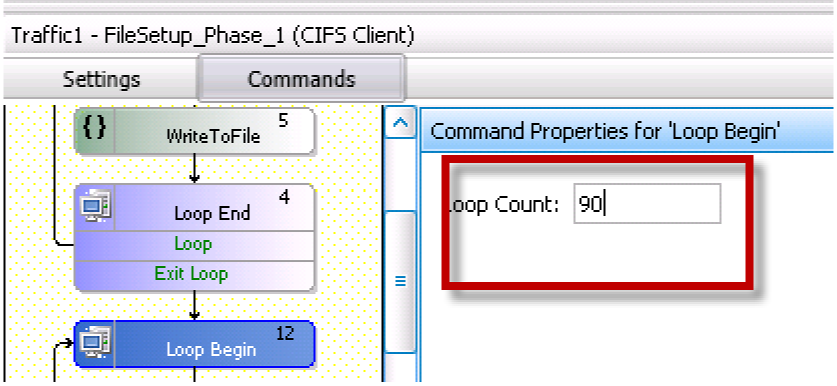
Traffic1 - FileSetup_Phase_1 (CIFS Client)

Command Properties for 'Loop Begin'

oop Count: 90

The next set of files will be if count 90

11. Add another **writetofile** command.

12. See the target folder and file as \public\file[10-99].txt . Use a different sequence generator range to enable to create file10.txt, file11.txt …..file99.txt.

13. Set the data length as 10 Kb.

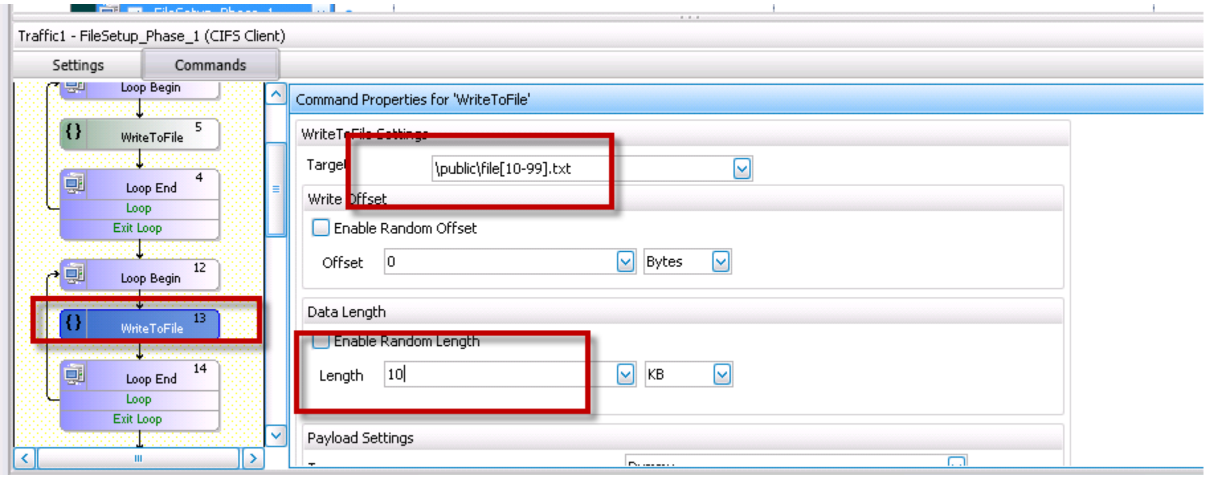

Command Properties for 'WriteToFile'

WriteToFile Settings

Target: \public\file[10-99].txt

Write Offset

□ Enable Random Offset

Offset: 0 Bytes

Data Length

□ Enable Random Length

Length: 10 KB

Payload Settings

Figure 464. Configuring the WriteToFile for the files of length 10KB

14. Repeat similar process for the next loop. This time the loop is of 900 and the command file size has changed to 50KB.

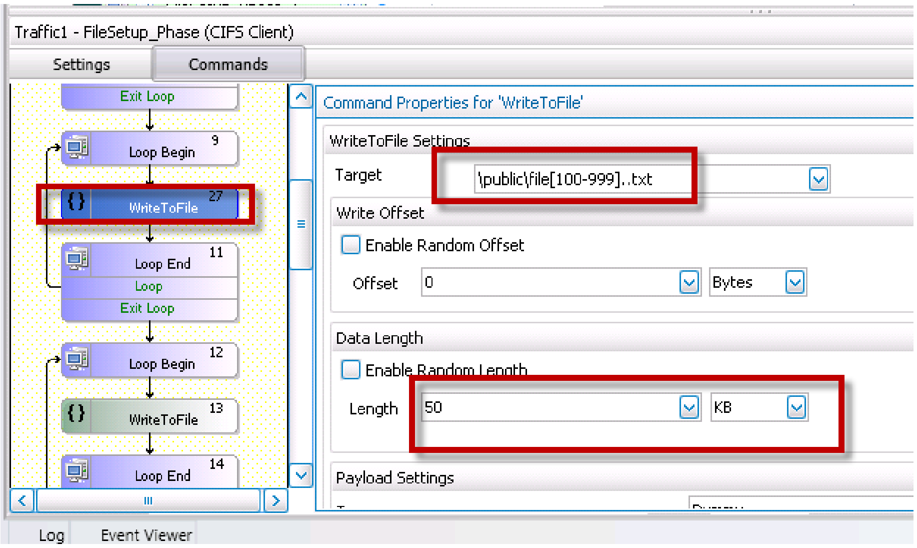

Command Properties for 'WriteToFile'

WriteToFile Settings

Target: \public\file[100-999]..txt

Write Offset

□ Enable Random Offset

Offset: 0 Bytes

Data Length

□ Enable Random Length

Length: 50 KB

Payload Settings

Figure 465. Configuring the WriteToFile for the files of length 50KB

15. Repeat similar process for the next loop to create 3000 files of size 100 KB.

   a. Creating **3000** files of **100KB**

      i. Add loop begin and set loop count as 3000

ii. In the **WriteToFile** command, use the sequence generator as file[1000-4000] and set Daat Length as 100 KB.

b. Creating **100** files of **1MB**

i. Add loop begin and set loop count as 100

ii. In the **WriteToFile** command, use the sequence generator as file[4001-4100].txt and set Data Length as 1MB.

c. Creating **100** files of **10 MB**

i. Add loop begin and set loop count as 100

ii. In the **WriteToFile** command, use the sequence generator as file[4101-4200] txt and set Data Length as 10MB.

d. Creating **20 file** of **1GB** each

i. Add loop begin and set loop count as 20

ii. In the **WriteToFile** command, use the sequence generator as file[4201-4220] .txt and set Data Length as 1GB.

16. Set the **Objective** as **Simulated User** with value as 1. Only one iteration of the command list creates 4220 files at the server, because loops are used for each write to file command.

File Count	**10**	**90**	**899**	**3000**	**100**	**100**	**20**
File Size	1KB	10KB	50KB	100KB	1MB	100MB	1GB

The file count and file lenth created at the target

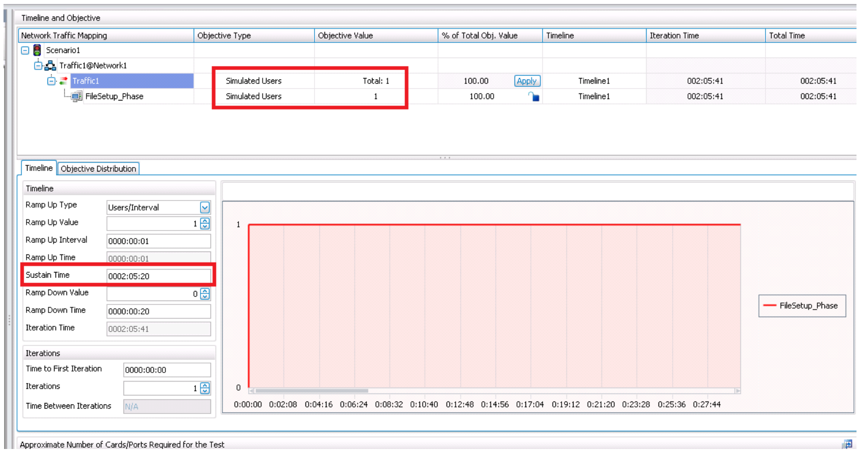

Timeline and Objective

Network Traffic Mapping	Objective Type	Objective Value	% of Total Obj. Value		Timeline	Iteration Time	Total Time
Scenario1							
Traffic1@Network1							
Traffic1	Simulated Users	Total: 1	100.00	Apply	Timeline1	002:05:41	002:05:41
FileSetup_Phase	Simulated Users	1	100.00		Timeline1	002:05:41	002:05:41

Timeline | Objective Distribution

Timeline

Ramp Up Type: Users/Interval

Ramp Up Value: 1

Ramp Up Interval: 0000:00:01

Ramp Up Time: 0000:00:01

Sustain Time: 0002:05:20

Ramp Down Value: 0

Ramp Down Time: 0000:00:20

Iteration Time: 0002:05:41

Iterations

Time to First Iteration: 0000:00:00

Iterations: 1

Time Between Iterations: N/A

Approximate Number of Cards/Ports Required for the Test

Figure 466. Only one user will create all the files at the NAS

17. Set the global loop as 1. Ceating the filesystem at the server side is a one time activity, so run the iteration only once.

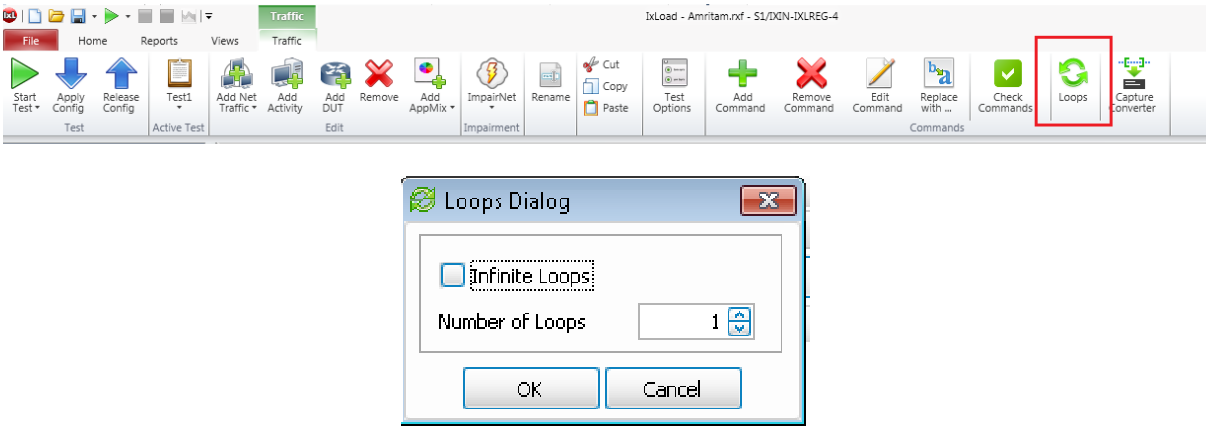

☐ Infinite Loops

Number of Loops 1

Figure 467. Disabling loop in the activity level.

18. Add the port connected to the NAS target. You can add the relevant ixia ports using the **Add Chassis**.

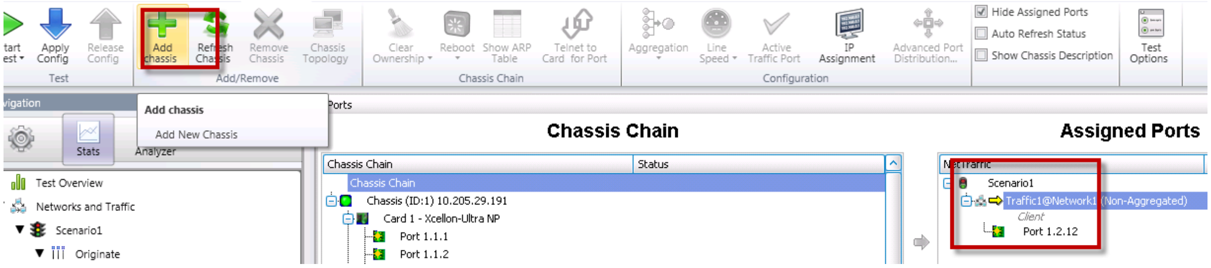

Figure 468. Add ports to the test.

19. Click the **Start Test** button to run the test.

    After the end of first run, ensure that all the files and folders are created successfully on the server side. Also ensure that there is no error in the stats.
    As the NAS server side, check and ensure that the files are created with the naming convention and data length.

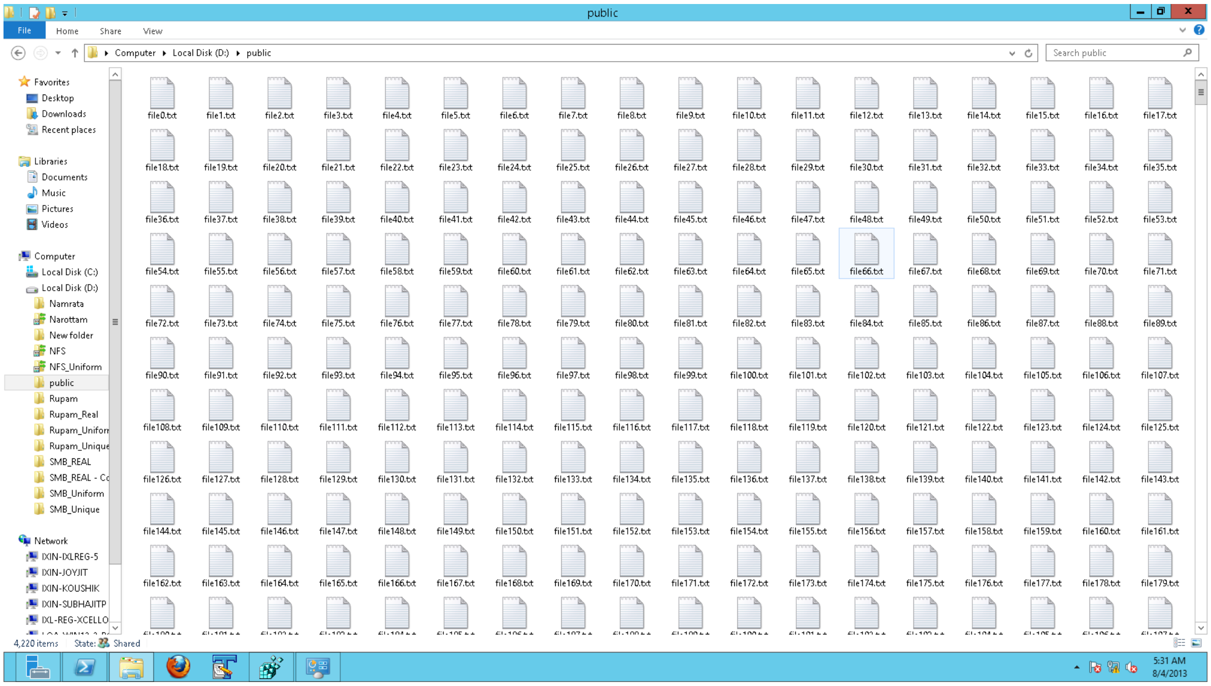

Figure 469. Several files getting created at the NAS

20. Once the setup phase is complete, start creating the testcases against the server. The first test involves accessing only the smaller sized files with sizes lesser than 100 KB. Disable the original activity that was used for setup and add a new CIFS client activity.

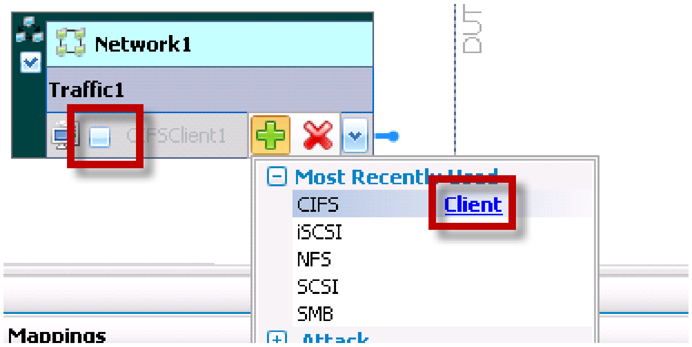

Figure 470. The setup phase is complete and the filer is ready. Now the real test validates the IOps and Tput performance of DUT.

21. For the New activity, add a **SessionSetup** command that has the information to log on to the server.

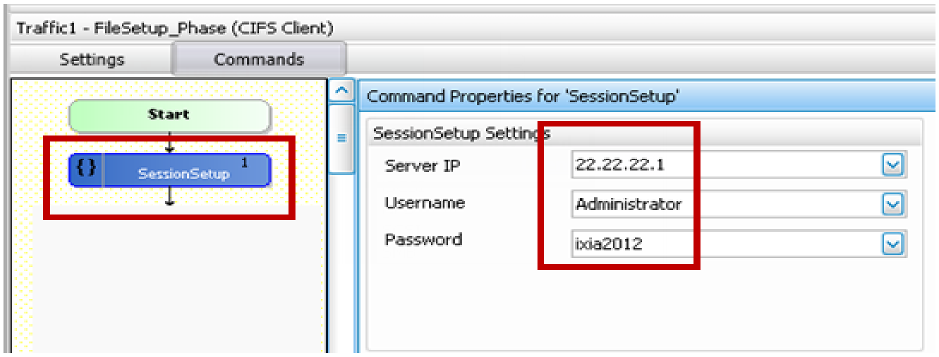
Traffic1 - FileSetup_Phase (CIFS Client)
Settings
Commands
Start
SessionSetup
1
Command Properties for 'SessionSetup'
SessionSetup Settings

Server IP 22.22.22.1

Username Administrator

Password ixia2012

22. The session setup command is followed by a **ReadFromFile** command. This command reads the files from the server.

23. Set the file name as \public\file($user-id).txt. In IxLoad, the token $user-id is replaced by a unique number. So each user has a unique number starting from 0. It means, if the test had 10 users user1 will acces file0.txt, user2 will access file1.txt user10 will acces file10.txt . There by, all the files are accessed and read simultaneously.

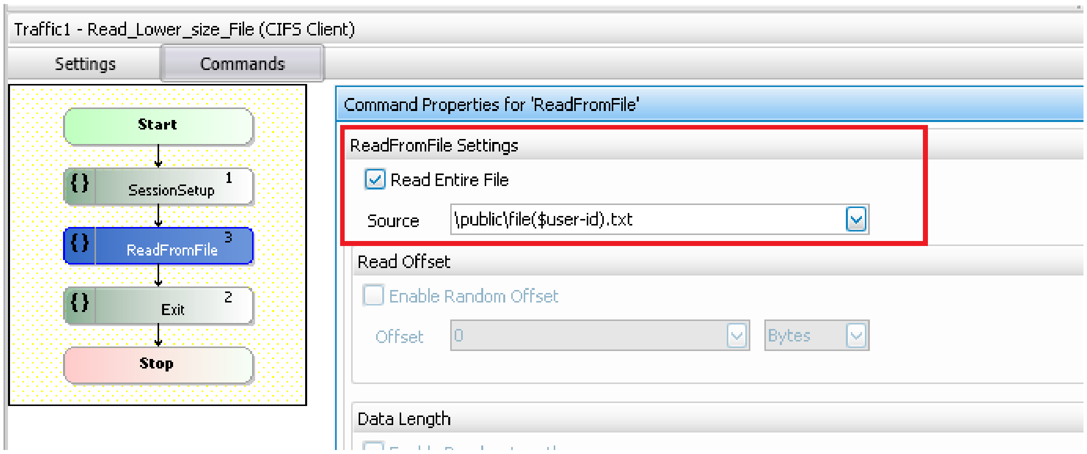

☑ Read Entire File

Source \public\file($user-id).txt

☐ Enable Random Offset

Offset 0 Bytes

Figure 471. Configuring the ReadFromFile command that reads back the files from target

There are 4000 files created in the server whose size are less than 100 KB. So set the objective as Simulated User = 4000. Use the timeline as desired.

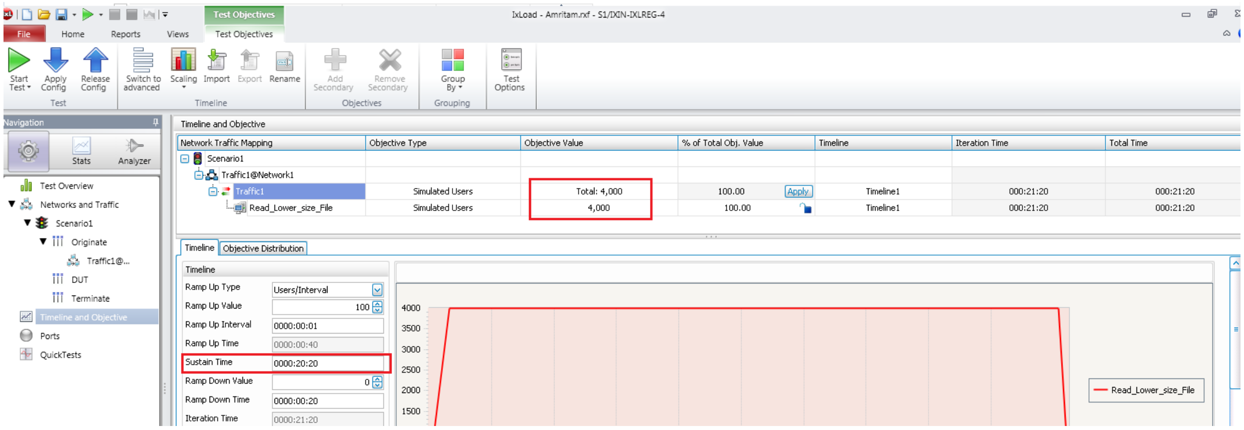

Network Traffic Mapping	Objective Type	Objective Value	% of Total Obj. Value		Timeline	Iteration Time	Total Time
Scenario1							
Traffic1@Network1							
Traffic1	Simulated Users	Total: 4,000	100.00	Apply	Timeline1	000:21:20	000:21:20
Read_Lower_size_File	Simulated Users	4,000	100.00		Timeline1	000:21:20	000:21:20

Timeline	
Ramp Up Type	Users/Interval
Ramp Up Value	100
Ramp Up Interval	0000:00:01
Ramp Up Time	0000:00:40
Sustain Time	0000:20:20
Ramp Down Value	0
Ramp Down Time	0000:00:20
Iteration Time	0000:21:20

Figure 472. Configuring the Timeline for test with 4000 users

24. Run the test with the above configurations against the same NAS.

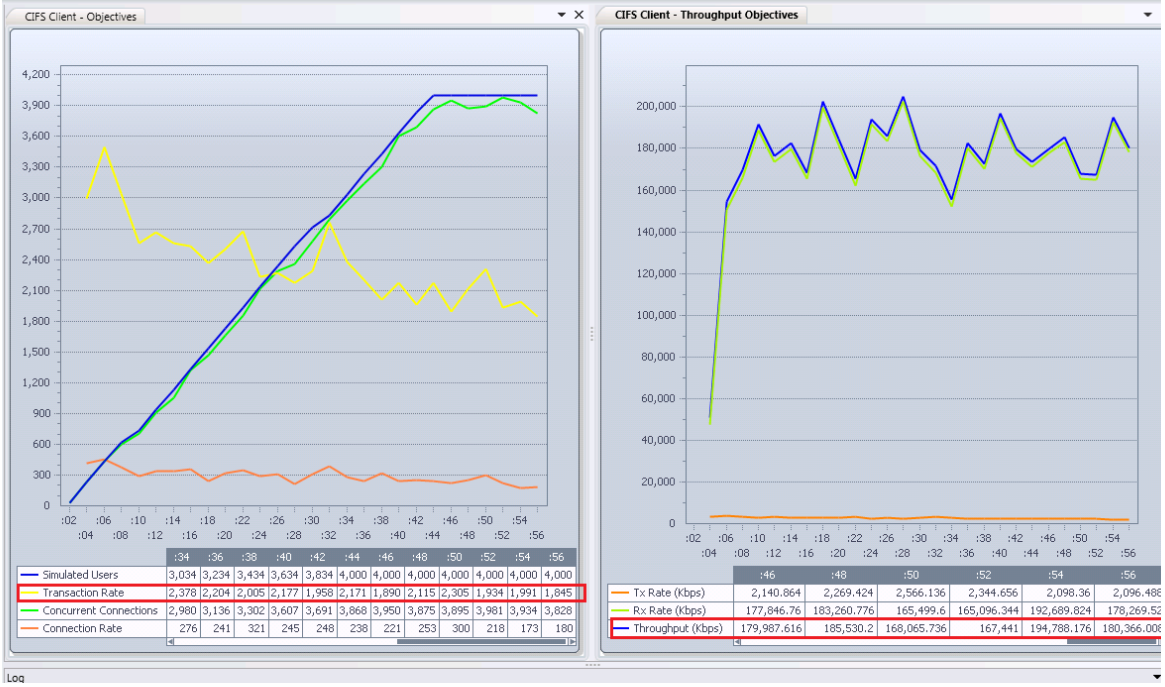

	:34	:36	:38	:40	:42	:44	:46	:48	:50	:52	:54	:56
Simulated Users	3,034	3,234	3,434	3,634	3,834	4,000	4,000	4,000	4,000	4,000	4,000	4,000
Transaction Rate	2,378	2,204	2,005	2,177	1,958	2,171	1,890	2,115	2,305	1,934	1,991	1,845
Concurrent Connections	2,980	3,136	3,302	3,607	3,691	3,868	3,950	3,875	3,895	3,981	3,934	3,828
Connection Rate	276	241	321	245	248	238	221	253	300	218	173	180

	:46	:48	:50	:52	:54	:56
Tx Rate (Kbps)	2,140.864	2,269.424	2,566.136	2,344.656	2,098.36	2,096.488
Rx Rate (Kbps)	177,846.76	183,260.776	165,499.6	165,096.344	192,689.824	178,269.52
Throughput (Kbps)	179,987.616	185,530.2	168,065.736	167,441	194,788.176	180,366.008

Figure 473. The max transaction and throughput achieved at

25. Rerun the test now with a **WriteToFIle** command so that it can also edit the files.

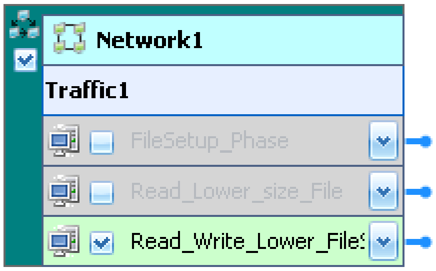

- Network1
- Traffic1
- □ FileSetup_Phase
- □ Read_Lower_size_File
- ☑ Read_Write_Lower_File$

Figure 474. Add another activity disabling the previous

26. The command sequence must be similar to previous one except that a WriteToFile command after ReadFromFile command is added. And use similar kind of $user-id token. The intention is to write 10KB to each file.

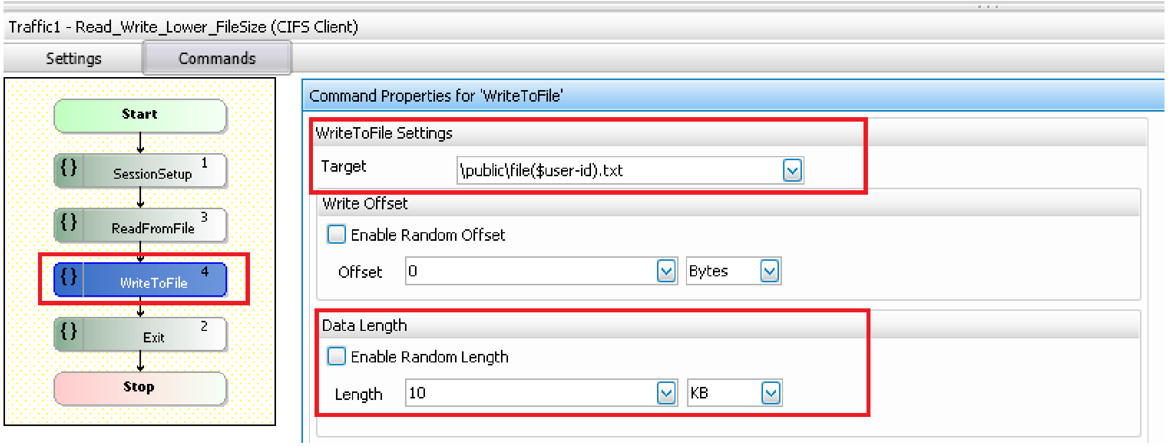

Command Properties for 'WriteToFile'

WriteToFile Settings

Target \public\file($user-id).txt

Write Offset

□ Enable Random Offset

Offset 0 Bytes

Data Length

□ Enable Random Length

Length 10 KB

Figure 475. The WriteToFile will edit the files.

27. Check the commands that are successful and failed in the **CIFSv2 Client Commands** stats.

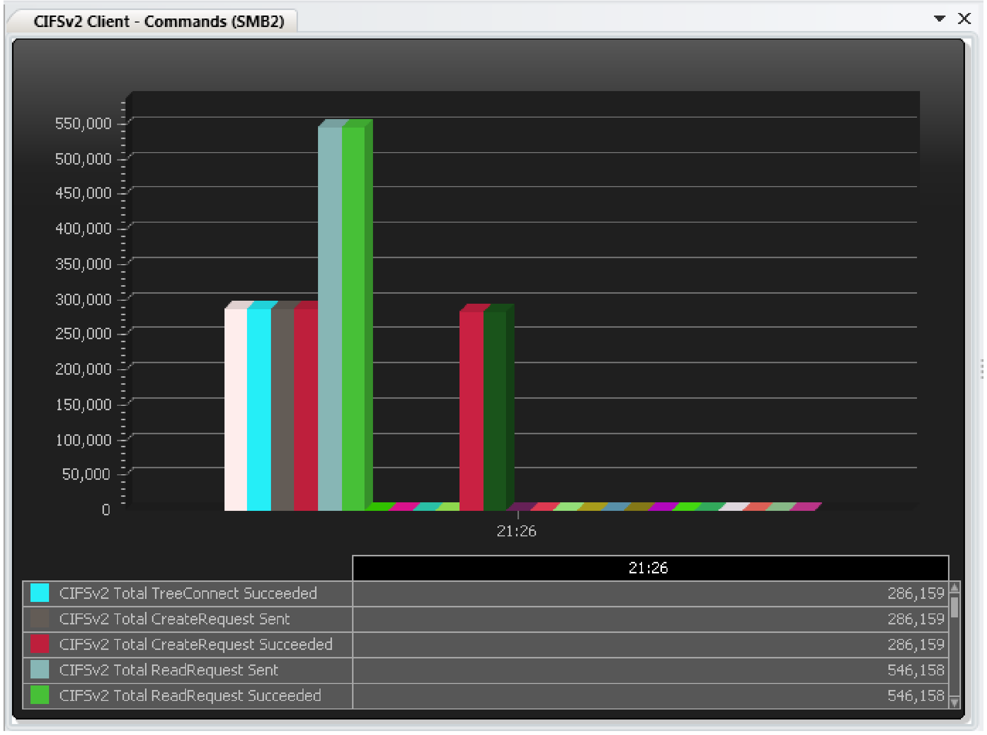

	21:26
🟦 CIFSv2 Total TreeConnect Succeeded	286,159
🟫 CIFSv2 Total CreateRequest Sent	286,159
🟥 CIFSv2 Total CreateRequest Succeeded	286,159
⬜ CIFSv2 Total ReadRequest Sent	546,158
🟩 CIFSv2 Total ReadRequest Succeeded	546,158

Figure 476. CIFSv2 Client command Stats

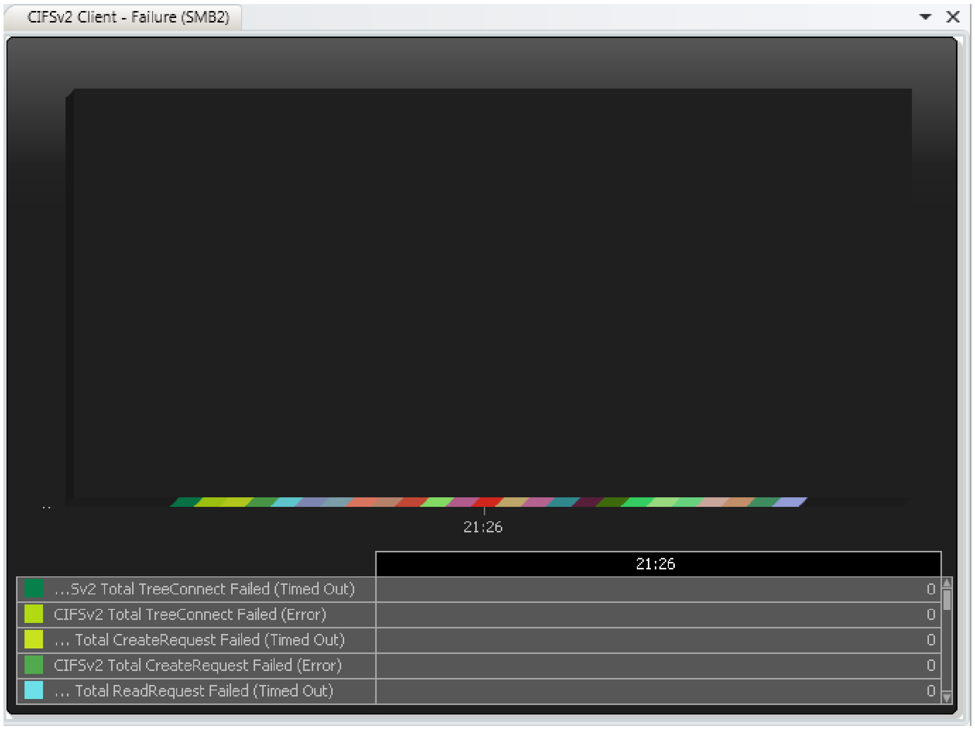

	21:26
🟩 ...Sv2 Total TreeConnect Failed (Timed Out)	0
🟨 CIFSv2 Total TreeConnect Failed (Error)	0
🟨 ... Total CreateRequest Failed (Timed Out)	0
🟩 CIFSv2 Total CreateRequest Failed (Error)	0
🟦 ... Total ReadRequest Failed (Timed Out)	0

Figure 477. CIFSv2 Failure stats

28. Rerun the test, but now the objective is to access all the files.

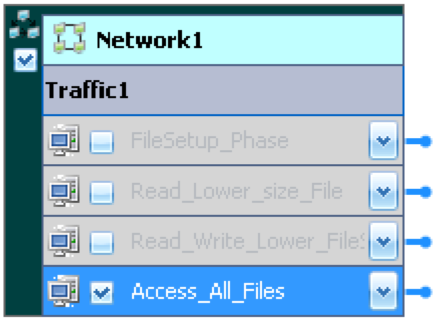

29. The command sequence for this activity must be same as the previous one. Change the objective value to the number of files present in the shared folder and that is 4220 . Using this, all the users can access their respective files.

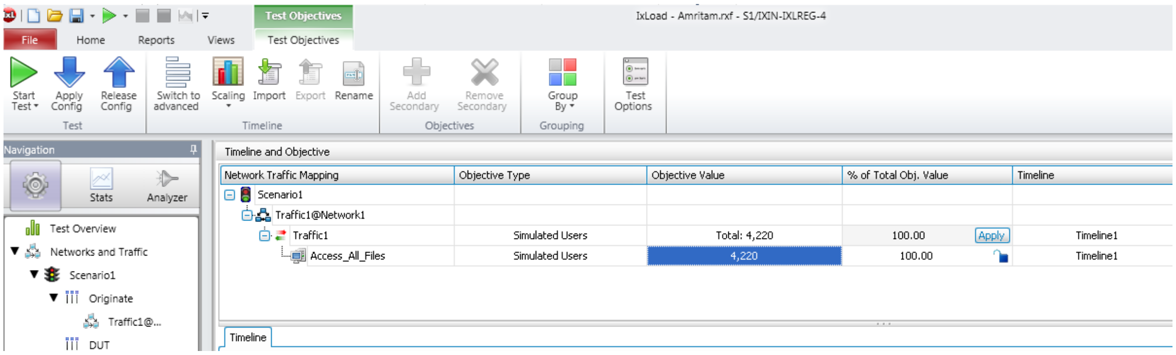

Network Traffic Mapping	Objective Type	Objective Value	% of Total Obj. Value	Timeline
Scenario1				
Traffic1@Network1				
Traffic1	Simulated Users	Total: 4,220	100.00 Apply	Timeline1
Access_All_Files	Simulated Users	4,220	100.00	Timeline1

Figure 478. Changing Simulated User to 4220

Results

The steady-state throughput performance, the max connections per second and the latency can be obtained in the respective views in the **Statistics** viewer.

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/