Analysis of Voice Performance between IPv4 and IPv6 Networks on IPv6 over MPLS (Cisco 6PE) Environment

Created by:

Azmat Ali

A project report part of the requirements for the degree of

Master of Science in Internetworking
(Department of Computer Science)

University of Alberta
April 2011

Abstract

2011 is turning into a crunch year for adopting the new standard for the Internet addressing, called IPv6, because IPv4 addresses have almost run out, and more businesses and institutions see the necessity of migrating to an IPv6 addressing system. An IPv6 address is 128 bits, which means that, in theory, there are 2^{128} addresses available, which represents approximately the number of grains of sand on our planet.

The Internet is a constantly evolving environment, which puts pressures on existing and evolving systems. Any protocol changes should be carefully designed and even more carefully deployed to avoid any disruption to the running system. It is no longer possible to deploy IPv6 infrastructure overnight, so engineers are considering various transition and evolution strategies.

Deploying IPv6 over Multiprotocol Label Switching (MPLS) backbones can be used for a transition from IPv4 to IPv6. Cisco 6PE (IPv6 over MPLS) enabled backbones allow IPv6 domains to communicate with each other over an MPLS IPv4 core network. This implementation requires no backbone infrastructure upgrades and no reconfiguration of core routers, because forwarding is based on labels rather than on the IP header itself.

We did a comparison of VoIP performance on IPv4 and IPv6 networks in the presence of varying levels of background UDP traffic, by using a powerful IXIA 400T traffic generator device to test performance and analyzed network traffic flows with the use of IxExplorer, which provides a high level of flexibility and functionality in protocol emulation, traffic generation, and analysis. IxExplorer is the primary controlling application for Ixia's purpose-built hardware test platform, allowing detailed configuration of protocols and analysis of test results.

Acknowledgments

This MINT Master's Capstone Project was developed in the State-Of-The-Art lab at University of Alberta to gain hands on experience and confidence. I would like to thank my supervisor Dr. Mike McGregor for his support especially when I was stuck to implement MPLS on native IPv6 environment. However, some leading vendors still say "We have no LDP over IPv6 implementation. This is under development" because LDPv6 specification still has "draft" status, and also we discussed with Cisco Support Team regarding LDPv6 verification (SR: 617148427. dated: March 16, 2011).

Therefore, I forwarded my project task to IPv6 over MPLS (Cisco 6PE) to analyze the performance differences between two primary IPv4 and IPv6 networks, although the results are not 100% perfect as we want to see with pure IPv6 network.

I needed to write to express my thanks for the MINT program as I have learnt various technologies and specially " **I**XIA **4**00T **T**raffic **G**enerator/**P**erformance **A**nalyzer " it is the very powerful device to analyze network performance and measure interface capabilities. Furthermore, in the MINT program I have learnt from highly experienced professors.

Finally, I would like to express my honest appreciation to my family, for their patience and moral support during my M.Sc. studies.

Contents

Chapter 1 (Introduction)

Internet Protocol Version 6 (IPv6) 5

Multiprotocol Label Switching (MPLS) 10

Multiprotocol Border Gateway Protocol (MP-BGP) 14

IXIA 400T IxExplorer 15

Chapter 2 (Methodology)

Equipments Details 16

Topology 17

Lab Experiment Statistics 18

Calculating Voice, Video, and FTP packet size 27

IxExplorer Operations 28

Conclusion 44

Appendix - A (Bibliography) 45

Appendix - B (Routers Configuration) 46

Introduction to IPv6:

IPv6 was proposed when it became clear that the 32-bit addressing scheme of IP version 4 (IPv4) was inadequate to meet the demands of Internet growth. IPv6 has a larger address space. The architecture of IPv6 was designed to allow existing IPv4 users to transition easily to IPv6, while providing services such as end-to-end security, Quality of Service, and globally unique addresses.

IPv6 enables the use of a global and reachable address for almost every kind of device: Computers, IP Phones, IP Faxes, TV setup boxes, Cameras, Pagers, Wireless PDAs, 802.11b devices, Cell phones, 3G IPADs, Home Networking, and Vehicles. The larger IPv6 address space allows networks to scale and provide global reachability.

The following list summarizes the characteristics of IPv6 and the improvements it can deliver:

- **Large address space:** Increased address size from 32 bits to 128
- **lined protocol header:** Improves packet-forwarding efficiency
- **Stateless autoconfiguration:** The ability for nodes to determine their own address
- **Multicast:** Increased use of efficient one-to-many communications
- **Jumbograms:** The ability to have very large packet payloads for greater efficiency
- **Network layer security:** Encryption and authentication of communications
- **Quality of service (QoS) capabilities:** QoS markings of packets and flow labels that help identify priority traffic
- **Anycast:** Redundant services using non-unique addresses
- **Mobility:** Simpler handling of mobile or roaming nodes

IPv6 Header Format

```
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|Version| Traffic Class |           Flow Label                  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         Payload Length        |  Next Header  |   Hop Limit   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
+                                                               +
|                                                               |
+                         Source Address                        +
|                                                               |
+                                                               +
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
+                                                               +
|                                                               |
+                      Destination Address                      +
|                                                               |
+                                                               +
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

[RFC - 2460]

- **Version:** ***4 bits:*** Identifies the version of the protocol. For IPv6, the version is 6.
- **Traffic Class:** ***8 bits:*** Intended for originating nodes and forwarding routers to identify and distinguish between different classes or priorities of IPv6 packets.
- **Flow Label:** ***20 bits:*** Defines how traffic is handled and identified. A flow is a sequence of packets sent either to a unicast or a multicast destination. This field identifies packets that require special handling by the IPv6 node. The following list shows the ways the field is handled if a host or router does not support flow label field functions:
  - If the packet is being sent, the field is set to zero.
  - If the packet is being received, the field is ignored.
- **Payload Length:** ***16 bits:*** Identifies the length, in octets, of the payload. This field is a 16-bit unsigned integer. The payload includes the optional extension headers as well as the upper-layer protocols, for example, TCP.
- **Next Header:** ***8 bits:*** Identifies the header immediately following the IPv6 header. The following are examples of the next header:
  - 00 = Hop-by-Hop options
  - 01 = ICMPv4
  - 04 = IP in IP (encapsulation)
  - 06 = TCP
  - 17 = UDP
  - 43 = Routing
  - 44 = Fragment
  - 50 = Encapsulating security payload
  - 51 = Authentication
  - 58 = ICMPv6 [1]

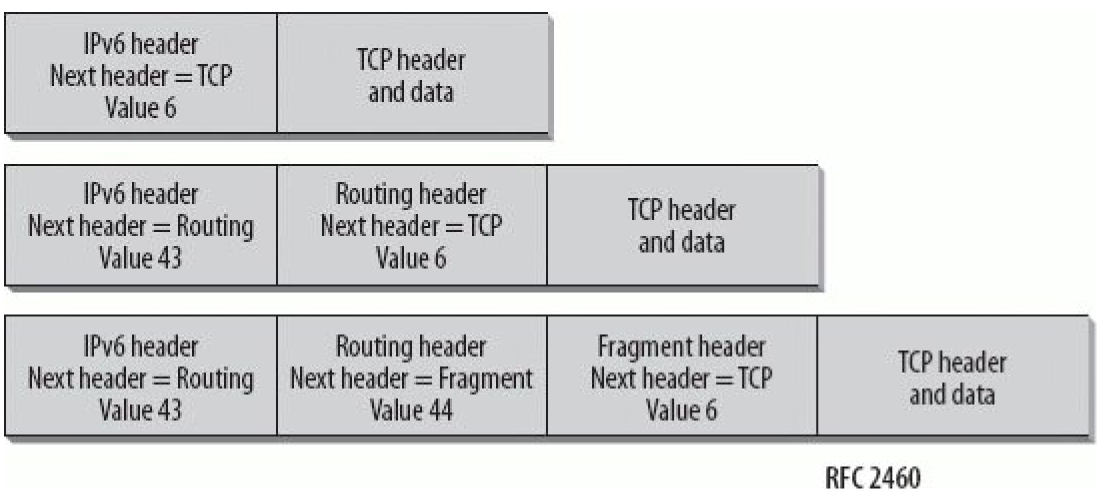

Figure: The use of Extension headers. [3]

- **Hop Limit:** ***8 bits:*** Identifies the number of network segments, also known as links or subnets, on which the packet is allowed to travel before being discarded by a router. The hop limit is set by the sending host and is used to prevent packets from endlessly circulating on an IPv6 internetwork.

  When forwarding an IPv6 packet, IPv6 routers must decrease the hop limit by 1 and must discard the IPv6 packet when the hop limit is 0.

- **Source Address:** ***128 bits:*** Identifies the IPv6 address of the original source of the IPv6 packet.
- **Destination Address:** ***128 bits:*** Identifies the IPv6 address of the intermediate or final destination of the IPv6 packet. [1]

IPv6 Addressing

IPv6 addresses are 128-bit identifiers for interfaces and sets of interfaces. These are represented as a series of 16-bit hexadecimal fields and each 16-bit block is converted to a 4-digit hexadecimal number separated by colons (**:**) in the format: x**:**x**:**x**:**x**:**x**:**x**:**x**:**x. The resulting representation is called colon-hexadecimal.
There are three types of addresses:

- **Unicast:** A unicast address uniquely identifies an interface of an IPv6 node. A packet sent to a unicast address is delivered to the interface identified by that address.
- **Anycast:** An anycast address is assigned to multiple interfaces (usually on multiple nodes). A packet sent to an anycast address is delivered to only one of these interfaces, usually the nearest one.
- **Multicast:** A multicast address identifies a group of IPv6 interfaces. A packet sent to a multicast address is processed by all members of the multicast group. [1]

IPv6 Address Notation

There are some conventions for representing IPv6 address's text strings:

- It is common for IPv6 addresses to contain successive hexadecimal fields of zeros. To make IPv6 addresses less cumbersome, two colons **(::)** may be used to compress successive hexadecimal fields of zeros at the beginning, middle, or end of an IPv6 address.
- Note that the double colon can appear only once in an address. The reason for this rule is that the computer always uses a full 128-bit binary representation of the address, even if the displayed address is simplified. When the computer finds a double colon, it expands it with as many zeros as are needed to get 128 bits. If an address had two double colons, the computer would not know how many zeros to add for each colon. So the IPv6 address ***2001:DB8:0000:0056:0000:ABCD:EF12:1234*** can be represented in the following ways (note the two possible positions for the double colon): [3]

  2001:DB8:0000:0056:0000:ABCD:EF12:1234
  2001:DB8:0:0056:0:ABCD:EF12:1234
  2001:DB8::0056:0:ABCD:EF12:1234
  2001:DB8:0:0056::ABCD:EF12:1234

IPv6 Address Types

The type of an IPv6 address is identified by the high-order bits of the address, as follows:

Address Type	Binary Prefix	Ipv6 notation
Unspecified	000…000 (128 bits)	::/128
Loopback	000…001 (128 bits)	::1/128
Multicast	11111111	FF00::/8
Link-local unicast	1111111010	FE80::/10
Site-local unicast	1111111011	FEC0::/10
Global unicast	Everything else	

Interface identifiers in IPv6 unicast addresses are used to identify interfaces on a link. They are required to be unique within a subnet prefix. It is recommended that the same interface identifier not be assigned to different nodes on a link. They also must be unique over a broader scope.

- The address 0:0:0:0:0:0:0:0 is called the unspecified address. It must never be assigned to any node. It indicates the absence of an address.
- The unicast address 0:0:0:0:0:0:0:1 is called the loopback address. It may be used by a node to send an IPv6 packet to itself. [4]

IPv6 Global Address

Aggregate-able global addresses are used on links that are aggregated upward through organizations,

3 bits	*45 bits*	*16 bits*	*64 bits*
001	**Routing Prefix**	**SLA**	**Interface ID**

- **001:** Identifies the address as being an aggregate-able global address.
- **Routing Prefix:** Included two other hierarchically structured fields named Top-Level Aggregator (TLA) and Next-Level Aggregator (NLA).
- **SLA:** Subnet ID, used by individual organizations to create their own local addressing hierarchy and to identify subnets.
- **Interface ID:** Must be unique to the link. [4]

IPv6 Link Local Address

This is use on a single link. Routers must not forward any packets with link-local source or destination addresses to other links. Link-Local addresses are designed to be used for addressing on a single link for purposes such as automatic address configuration, neighbor discovery, or when no routers are present.

10 bits	*54 bits*	*64 bits*
1111111010	000…000	**Interface ID**

IPv6 Site Local Address

This is addressing inside a site without the need for a global prefix. Routers must not forward any packets with site-local source of destination addresses outside of the site.

10 bits	*54 bits*	*64 bits*
1111111011	**Subnet ID**	**Interface ID**

The replacement for Site-Local addresses is called unique local IPv6 unicast address, or local IPv6 address for short. It is specified in RFC - 4193. These addresses are globally unique but should not be routed to the global Internet. They are designed to be used within corporate sites or confined sets of networks. [4]

Introduction to MPLS:

MPLS is a packet-forwarding technology that uses appended labels to make forwarding decisions for packets. [6]

- Within the MPLS network, the Layer 3 header analysis is done just once (when the packet enters the MPLS domain). Labels are appended to the packet, and then the packet is forwarded into the MPLS domain. [6]
- Simple label inspection integrated with CEF switching drives subsequent packet forwarding.[6]

MPLS was designed to support efficiently forwarding packets across the network core based on a simplified header. Label switching is performed regardless of the Layer 3 routing protocol.

MPLS labels typically correspond to Layer 3 destination addresses (equal to destination-based routing). Labels can also correspond to other parameters, such as quality of service (QoS), source address, or a Layer 2 circuit. An MPLS label is a short, 4-byte, fixed-length, locally significant identifier. [6]

MPLS Label Header

MPLS operates by defining a label inside MPLS "Shim header" that is placed on the packet between layer-2 and layer-3 headers. The 32-bit MPLS header is organized as in Figure.

Figure - MPLS Header [6]

Label: The header label consists of 20 bits Label, which is used to identify the Label Switched Path (LSP) to which the packet belongs in the MPLS domain. The labels on the packets are established by using Forward Equivalence Class (FEC).

EXP: The experimental 3 bits field used to handle traffic with Quality of Service (QoS).

S: 1 bit called stack field this is used to indicate the bottom of label stack.

TTL: The tail consists of 8 bits TTL (Time to Live) field that has the similar function as in the IP header.

MPLS Architecture

MPLS is a form of Internet specifications that describes interoperable mechanisms for implementing IP switching. MPLS Label Switching Routers (LSRs) switch IP datagrams over a Label Switched Path (LSP) that drives an MPLS domain. The LSP is established manually via management access or via signaling.

MPLS Architecture Control Plane:

The control plan takes care of the routing information exchange and the label exchange between adjacent devices.

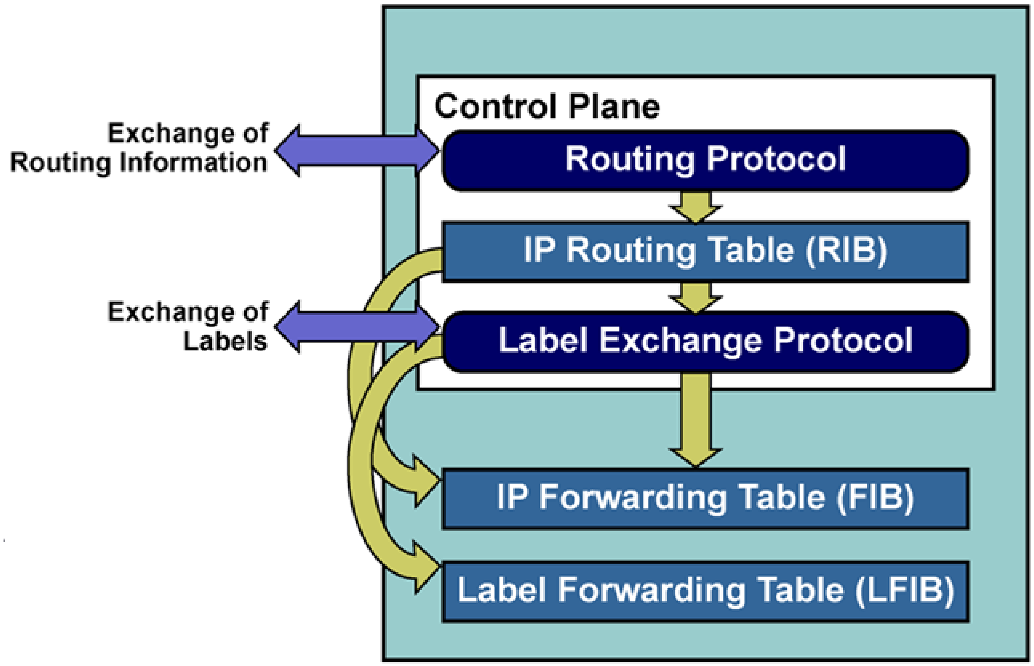

The control plan builds a routing table (Routing Information Base [RIB]) based on the routing protocol. [7] This maintains and controls the forwarding table by learning the network topology from the routing protocols such as OSPF, IS-IS, and BGP.

The control plan also builds two forwarding tables, a FIB from the information in the RIB, and a label forwarding information base (LFIB) table based on the label exchange protocol and the RIB. [6]

MPLS Architecture Data Plane:

The data plane is a simple forwarding engine that is independent of the type of routing protocol or label exchange protocol being used. The data plane forwards packets to the appropriate interface based on the information in the LFIB or the FIB tables. [6]

So, when a packet arrives at the router, the forwarding decision is taken by the data plane through consulting the forwarding table, which is maintained by control plane. The packets are then forwarded towards the appropriate node based on the forwarding decision.

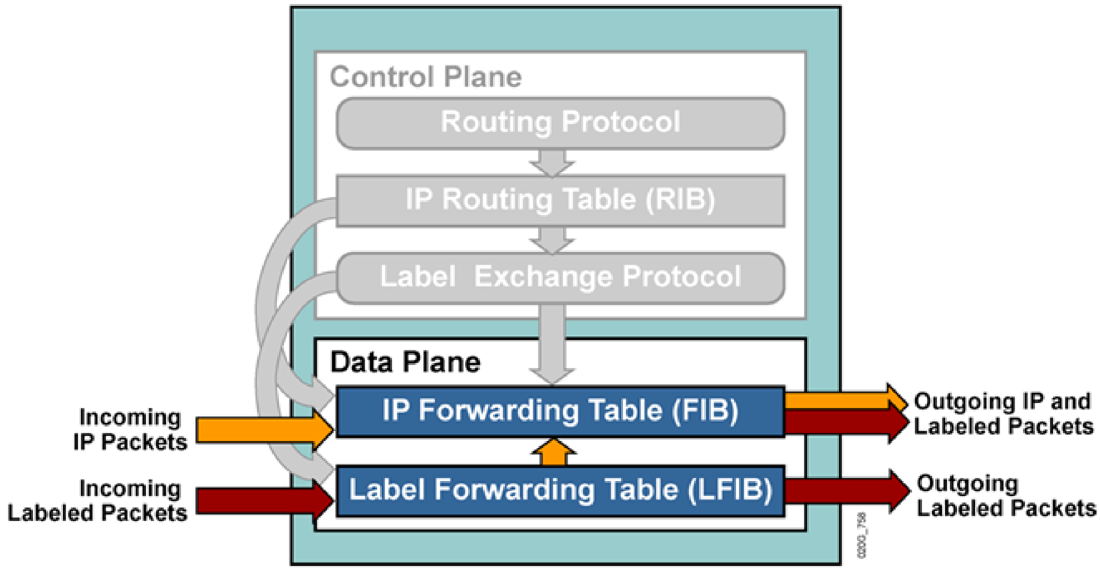

Label Switch Routers: Architecture of LSRs:

Routers that support MPLS are Label Switched Routers (LSRs) or a device that creates label distribution procedures and primarily forwards packets based on label.

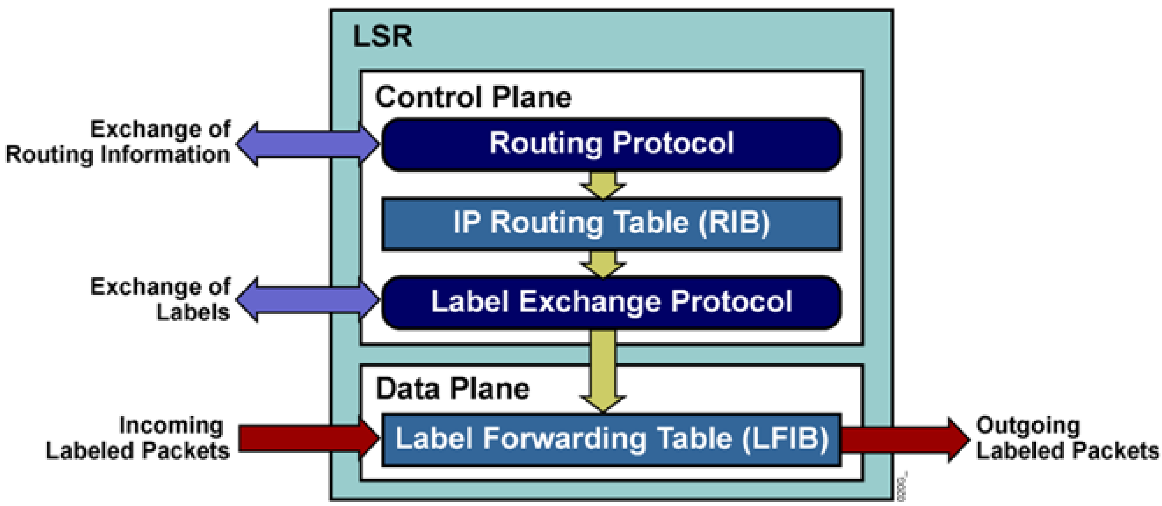

The primary LSR functions are to exchange labels with other LSRs and to forward labeled packets. Therefore, every LSR needs a Layer 3 routing protocol (for example, OSPF, EIGRP, or IS-IS), and a label exchange protocol (for example, LDP or TDP). [6]

LSRs: Architecture of Edge LSRs:

An LSR on the edge of an MPLS domain that implements label distribution procedures, forwards packets based on labels, and primarily inserts labels on packets or removes labels for non-MPLS devices. [6]

[To the service provider MPLS environment, an edge LSR is typically known as a provider edge (PE) router, and an LSR is known as a provider (P) router.]

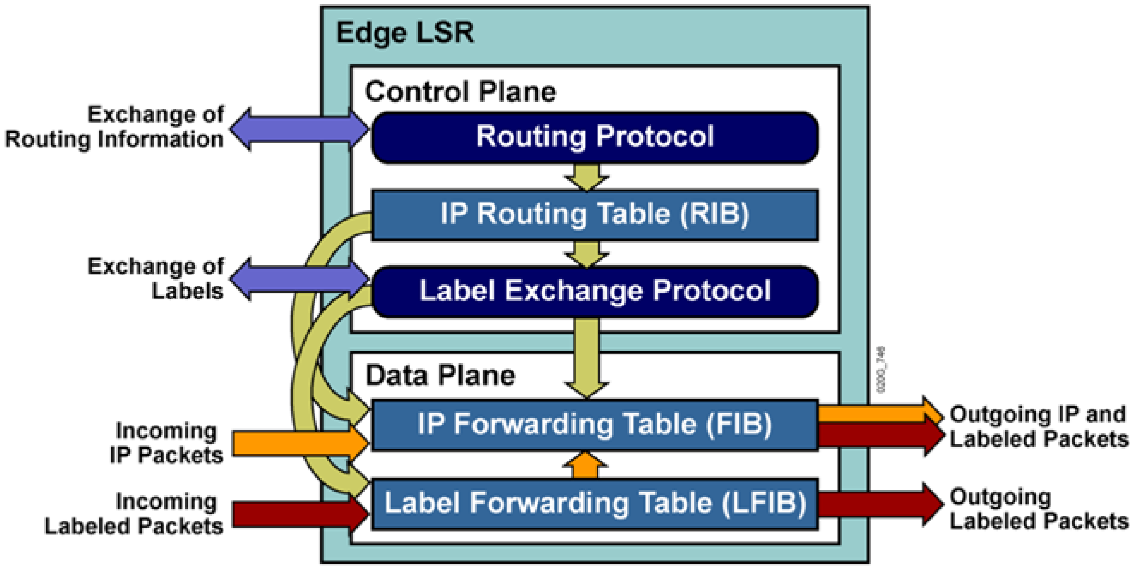

All the packets enter or exit to the MPLS domain through Edge LSR. ELSR has two roles, Ingress router and Egress router. The ingress router maps the incoming traffic into MPLS domain, while the egress router forward traffic from MPLS domain to the IP networks.

Label Distribution Protocol (LDP):

LDP is a protocol in which the label mapping information is exchanged between LSRs. It establishes and maintains labels. In addition, LDP is a new protocol that defines a set of procedures and messages by which one LSR informs another of the label bindings it has made.

Forward Equivalence Class (FEC):

The FEC represents the path that packets are going to take through the MPLS cloud based on criteria like the iBGP next-hop address (in the common case of Layer 3 MPLS VPN).
Throughout the MPLS network, all packets in an FEC are forwarded using the next-hop address for the FEC. The label value changes as the IP packet traverses the network. When a labeled packet is sent from one LSR to the next-hop LSR, the label value carried by the packet is the label value that the next-hop LSR assigned to represent the FEC of the packet. [6]

Label Switched Path (LSP):

LSP is the path groups by the signaling protocols in MPLS domain. In MPLS domain there exists the number of LSPs that originate at Ingress router and traverse one or more core LSRs and terminate at the Egress router.

Introduction to MP-BGP:

The Internet is a network of interconnected networks. A network or a set of networks under a single authority is called an Autonomous System (AS). Border Gateway Protocol (BGP) version 4 (RFC1771) is the routing protocol that is used currently to maintain connectivity between AS's. BGP is a path vector protocol. BGP relies on Transmission Control Protocol (TCP) for reliable transport. BGP multiprotocol extensions support carrying IPv6 prefixes, MPLS labels and IP multicast in addition to IPv4 prefixes.

IPv6 consists of three unicast address scopes, Link-Local, Site-Local, and Global. Link local and Site local addresses are unique but should not be routed to the internet while the link local address can be used as next hop addresses in some routing protocols.

Link local addresses are unsuitable to be used as next hop attributes in BGP - 4, due to that reason, when BGP - 4 is used to convey IPv6 reachability information it is sometimes necessary to announce a next hop attribute that consists of global address and a link-local address. [7]

MP-BGP uses two new attributes, 1) Multiprotocol - Network Layer Reachability Information (MP_REACH_NLRI) attribute, and 2) Multiprotocol - Unreachable Network Layer Reachability Information (MP_UNREACH_NLRI) attribute.

Multiprotocol Reachable NLRI - MP_REACH_NLRI: This is an optional non-transitive attribute that can be used for the following purposes:
(a) To advertise a feasible route to a peer.
(b) To permit a router to advertise the Network Layer address of the router that should be used as the next hop to the destinations listed in the Network Layer Reachability Information field of the MP_NLRI attribute.
(c) To allow a given router to report some or all of the Subnetwork Points of Attachment (SNPAs) that exist within the local system. [8]

Multiprotocol Unreachable NLRI - MP_UNREACH_NLRI: This is an optional non-transitive attribute that can be used for the purpose of withdrawing multiple unfeasible routes from service. [8]

MP-BGP already supports the IPv6 address family through the MP_REACH_NLRI and MP_UNREACH_NLRI attributes. Cisco 6PE requires that MP-BGP be further extended to be able to bind and MPLS label to the IPv6 route as per [MPLS_BGP]. [9]

The label binding information is piggybacked along the prefix advertisement in the same MP_REACH_NLRI attribute. The fact that the MP_REACH_NLRI contains a label is indicated by SAFI (Subsequent Address Family Identifier) value of 4. [9]

The Network Layer Reachability information is encoded as one or more triples of the form <length, label, prefix> whose fields are described below: [10]

```
+---------------------------+
|   Length (1 octet)        |
+---------------------------+
|   Label (3 octets)        |
+---------------------------+
.............................
+---------------------------+
|   Prefix (variable)       |
+---------------------------+
```

- **Length:** The Length field indicates the length in bits of the address prefix plus the label(s).
- **Label:** The Label field carries one or more labels (that corresponds to the stack of labels [MPLS-ENCAPS]). Each label is encoded as 3 octets, where the high-order 20 bits contain the label value, and the low order bit contains "Bottom of Stack" (as defined in MPLS-ENCAPS]).
- **Prfix:** The Prefix field contains address prefixes followed by enough trailing bits to make the end of the field fall on an octet boundary. Note that the value of trailing bits is irrelevant. [10]

Introduction to IXIA 400T IxExplorer:

As networks converge into a single network carrying voice, video, data, and wireless traffic, it is critical that leading-edge test systems be used to assess functionality, scale, and performance. The increase in complexity necessitates additional capabilities, flexibility, and power in the test equipment designed to assess network and network device performance. Such sophisticated test systems must be flexible, highly scalable, easy to use, and fully support multiple routing protocol emulations. In addition, test systems must generate line-rate traffic and analyze millions of traffic flows with comprehensive QoS analysis. [11]

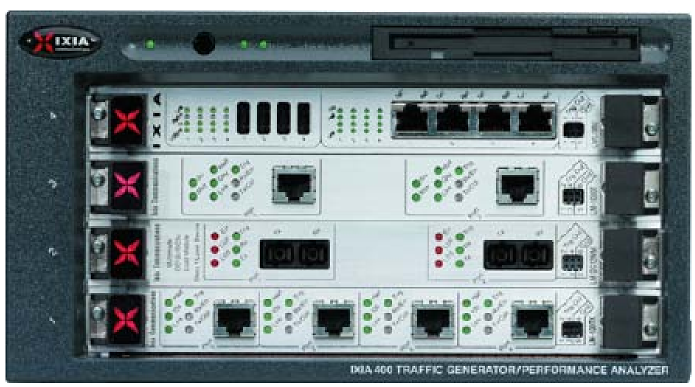

The Ixia system is the most comprehensive tool available in our MINT lab for testing multilayer 10/100 Mbps Ethernet, Ethernet Gigabit, 10 Gigabit Ethernet.

The Ixia ExExplorer software provides complete configurations, control, and monitoring of all Ixia resources in the test network, and the Tcl scripts allow one to rapidly conduct the most popular industry benchmark test.

The unit can be connected to an Ethernet network, and an administrator can remotely monitor and control it using the IxExplorer software program. Multiple users can access the unit simultaneously, splitting the ports within a chassis and controlling the activity and configuration of all ports and functions.

Equipment Details:

The generic devices are used in the project include the following:

Devices	Software Version	Device's Roles	Interfaces
2821	12.4(11)XJ4	RCE1	Gigabit Ethernet 0/0 Gigabit Ethernet 0/1
2821	12.4(11)XJ4	RCE2	Gigabit Ethernet 0/0 Gigabit Ethernet 0/1
2921	15.0(1)M4	Rack-1-RPE1	Gigabit Ethernet 0/0 Gigabit Ethernet 0/1 Gigabit Ethernet 0/2
2921	15.0(1)M4	RP1	Gigabit Ethernet 0/0 Gigabit Ethernet 0/1 Gigabit Ethernet 0/2
2921	15.0(1)M4	RP2	Gigabit Ethernet 0/0 Gigabit Ethernet 0/1 Gigabit Ethernet 0/2
2921	15.0(1)M4	Rack-3-RPE2	Gigabit Ethernet 0/0 Gigabit Ethernet 0/1 Gigabit Ethernet 0/2
2821	12.4(11)XJ4	RCE3	Gigabit Ethernet 0/0 Gigabit Ethernet 0/1
2821	12.4(11)XJ4	RCE4	Gigabit Ethernet 0/0 Gigabit Ethernet 0/1

Note: The experimental lab base on above hardware and IOS versions.

Topology:

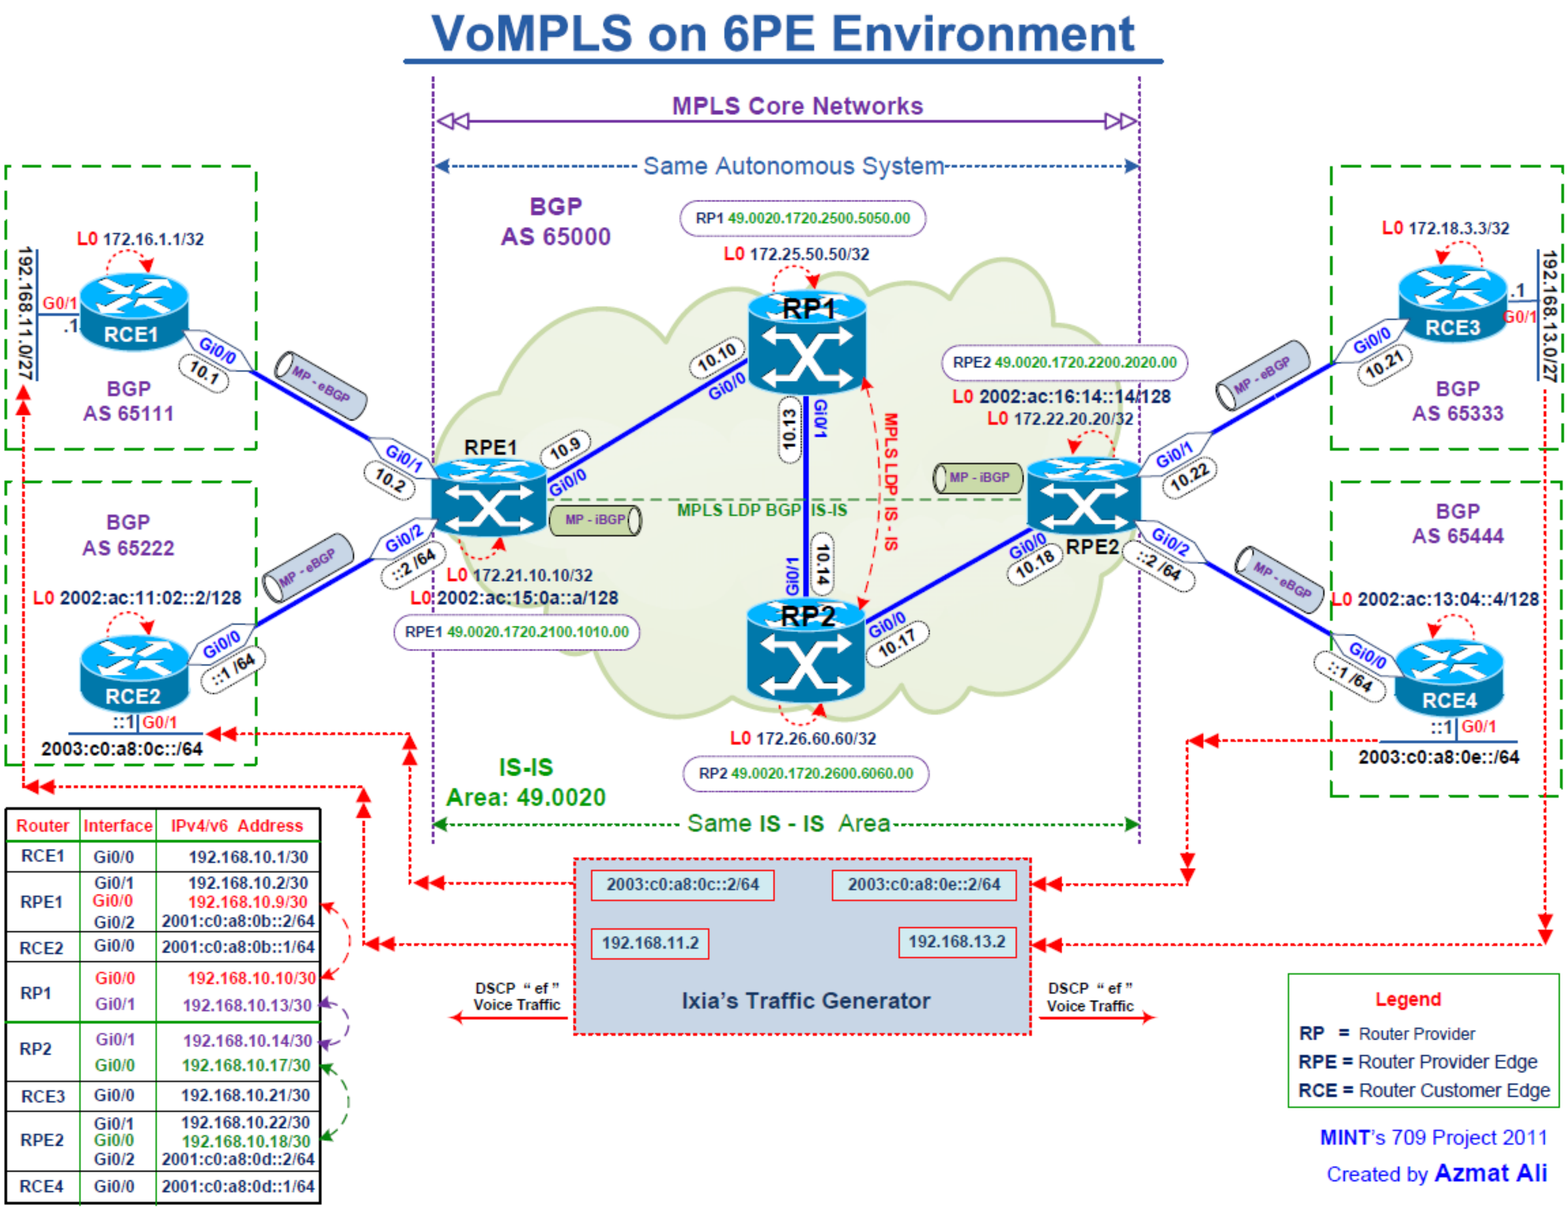

Router	Interface	IPv4/v6 Address
RCE1	Gi0/0	192.168.10.1/30
RPE1	Gi0/1 Gi0/0 Gi0/2	192.168.10.2/30 192.168.10.9/30 2001:c0:a8:0b::2/64
RCE2	Gi0/0	2001:c0:a8:0b::1/64
RP1	Gi0/0 Gi0/1	192.168.10.10/30 192.168.10.13/30
RP2	Gi0/1 Gi0/0	192.168.10.14/30 192.168.10.17/30
RCE3	Gi0/0	192.168.10.21/30
RPE2	Gi0/1 Gi0/0 Gi0/2	192.168.10.22/30 192.168.10.18/30 2001:c0:a8:0d::2/64
RCE4	Gi0/0	2001:c0:a8:0d::1/64

Lab Experiment Statistics:

The Cisco 6PE feature is particularly applicable to Service Providers who already run an MPLS network or plan to do it. One of the Cisco 6PE advantages is that there is no need to upgrade the hardware, software or configuration of the core network. MPLS has been chosen by many Service Providers as a transport mechanism to deliver services to customers. MPLS as a multi-service infrastructure technology that is able to provide layer 3 VPN, QoS, traffic engineering, fast re-routing and integration of ATM and IP switching. It is in a very natural manner that MPLS is put to contribution to ease IPv6 introduction in existing production networks. [9]

Four types of routing interactions are involved on the path to or from the IPv6 stations:

- RCE2 router running an IPv6 Multi-Protocol External Border Gateway Protocol (MP-eBGP).
- RCE2 and RPE1 router exchanging IPv6 routing information through an IPv6 External Gateway Protocol using MP- eBGP connection.
- RPE1 and RPE2 router peering together through MP - iBGP for exchanging IPv6 reachability with the MPLS cloud and performing IPv6 forwarding.
- Core MPLS/IPv4 cloud, running an IPv4 Internal Gateway Protocol Intermediate System to Intermediate System (IS - IS) to establish reachability between RPEs, an IPv4 label distribution protocol (LDP) to distribute IPv4 labels.

  RCE2 is not aware that forwarding occurs over MPLS clouds. They operate in their regular IPv6 way.

  RPs routers inside the MPLS clouds are not aware that they are switching IPv6 packets, as they only use MPLS forwarding, LDP for binding IPv4 labels and an IPv4 IGP (IS-IS) to establish internal reachability inside the MPLS cloud. Therefore, the Cisco 6PE feature does not impact these MPLS core devices.

Control Plane Function at Egress RPE2:

In terms of routing function, the egress RPE2 is responsible to learn by usual means and to advertise through MP-BGP the IPv6 prefixes in its reach: [9]

- RPE2 includes, in its IPv6 routing table, prefix RP2 reachable through connected interfaces or connected RCEs. IPv6 routing protocol running between RCEs and RPEs (MP - eBGP) to reduce configuration complexity.
- RPE2 redistribute RP2 prefixes in MP - BGP under the "IPv6+label" address family.
- MP-BGP assigns a local label to prefix RP2. This label is an aggregate label. It is not specific to a given prefix, it is chosen out of a sixteen-label pool in a round robin fashion. Consequently, egress RPE is not able to take a forwarding decision solely based on this label. The advantage of using more than one label value is that it results in better load sharing in the core of the network. [9]

- MP-BGP create a BGP update including a MP_REACH_NLRI attribute with ***AFI = 2*** (IPv6), ***SAFI = 4*** (label), ***Prefix*** = RP2, ***Label*** = L2, ***next-hop*** = **::**FFFF**:**172.21.10.10. The IPv6 next-hop address is expressed as an IPv4-mapped IPv6 address.

Where: ***AFI*** = 2 (*Address Family Identifier*)**:** IPv6 address of RPE2's loopback address: 2002:AC:16:14::14/128. ***SAFI*** = 4 (*Subsequent Address Family Identifier*)**:** Label 24. ***Label*** = L2 to RP2 .

Prefix = RP2

```
Rack-3-RPE2#sh mpls forwarding-table
Local      Outgoing    Prefix           Bytes Label   Outgoing    Next Hop
Label      Label       or Tunnel Id     Switched      interface
16         Pop Label   172.26.60.60/32  0             Gi0/0       192.168.10.17
```

next-hop = **::**FFFF**:**172.21.10.10 (MP - iBGP peering with RPE1).

The Following Statistics:

The Egress RPE2 router, prefix *2002:AC:16:14::14/128* is locally connected via interface loopack0:

```
Rack-3-RPE2#sh ipv6 route
IPv6 Routing Table - default - 9 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       D - EIGRP, EX - EIGRP external, ND - Neighbor Discovery
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
B   2001:C0:A8:B::/64 [200/0]
     via 172.21.10.10%default, indirectly connected
C   2001:C0:A8:D::/64 [0/0]
     via GigabitEthernet0/2, directly connected
L   2001:C0:A8:D::2/128 [0/0]
     via GigabitEthernet0/2, receive
B   2002:AC:11:2::2/128 [200/0]
     via 172.21.10.10%default, indirectly connected
B   2002:AC:13:4::4/128 [20/0]
     via FE80::217:E0FF:FEC0:17C0, GigabitEthernet0/2
B   2002:AC:15:A::A/128 [200/0]
     via 172.21.10.10%default, indirectly connected
LC  2002:AC:16:14::14/128 [0/0]
     via Loopback0, receive
B   2003:C0:A8:E::/64 [20/0]
     via FE80::217:E0FF:FEC0:17C0, GigabitEthernet0/2
L   FF00::/8 [0/0]
     via Null0, receive
```

And then injected in MP - BGP:

```
Rack-3-RPE2#sh bgp ipv6 unicast
BGP table version is 96, local router ID is 6.6.6.6
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*>i2001:C0:A8:B::/64
                    ::FFFF:172.21.10.10
                                             0    100      0 ?
*  2001:C0:A8:D::/64
                    2001:C0:A8:D::1          0             0 65444 ?
*>                  ::                       0         32768 ?
*>i2002:AC:11:2::2/128
                    ::FFFF:172.21.10.10
                                             0    100      0 65222 i
*> 2002:AC:13:4::4/128
                    2001:C0:A8:D::1          0             0 65444 i
*>i2002:AC:15:A::A/128
                    ::FFFF:172.21.10.10
                                             0    100      0 i
*> 2002:AC:16:14::14/128
                    ::                       0         32768 i
*> 2003:C0:A8:E::/64
   Network          Next Hop            Metric LocPrf Weight Path
                    2001:C0:A8:D::1          0             0 65444 ?
```

Aggregate label 24 is associated to this route:

```
Rack-3-RPE2#sh bgp ipv6 unicast tags
   Network          Next Hop      In tag/Out tag
   2001:C0:A8:B::/64
                    ::FFFF:172.21.10.10
                                notag/22
   2001:C0:A8:D::/64
                    2001:C0:A8:D::1 23/notag
                    ::              23/notag
   2002:AC:11:2::2/128
                    ::FFFF:172.21.10.10
                                notag/23
   2002:AC:13:4::4/128
                    2001:C0:A8:D::1 21/notag
   2002:AC:15:A::A/128
                    ::FFFF:172.21.10.10
                                notag/21
   2002:AC:16:14::14/128
                    ::              24/notag
   2003:C0:A8:E::/64
                    2001:C0:A8:D::1 19/notag
```

```
Rack-3-RPE2#sh mpls forwarding-table
Local      Outgoing    Prefix            Bytes Label   Outgoing   Next Hop
Label      Label       or Tunnel Id      Switched      interface
16         Pop Label   172.26.60.60/32   0             Gi0/0      192.168.10.17
17         17          172.25.50.50/32   0             Gi0/0      192.168.10.17
18         Pop Label   192.168.10.12/30  0             Gi0/0      192.168.10.17
19         No Label    2003:C0:A8:E::/64   \
                                         10723140      Gi0/2      FE80::217:E0FF:FEC0:17C0
20         18          192.168.10.8/30   0             Gi0/0      192.168.10.17
21         No Label    2002:AC:13:4::4/128   \
                                         24852         Gi0/2      FE80::217:E0FF:FEC0:17C0
23         No Label    2001:C0:A8:D::/64   \
                                         7068          aggregate
24         Pop Label   2002:AC:16:14::14/128    \
                                         1500          aggregate
26         19          172.21.10.10/32   0             Gi0/0      192.168.10.17
```

Control Plane Function at Ingress RPE1:

- RPE1 extract the (prefix, label, next-hop = RPE2) triple from MP_REACH_NLRI attributes included in MP - BGP updates coming from RPE2. In turn the (prefix, label) couple is installed in the MP - BGP table for AFI = 2 (IPv6) / SAFI = 4 (Label) and the IPv4 address is extracted from the "IPv4 mapped" IPv6 address which was advertised as next-hop. [9]

The following Statistics:

Prefix *2002:AC:16:14::14/128* is learnt by RPE1 router (ingress RPE) from RPE2 router:

```
Rack-1-RPE1#sh bgp ipv6 unicast
BGP table version is 14, local router ID is 5.5.5.5
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*  2001:C0:A8:B::/64
                    2001:C0:A8:B::1          0             0 65222 ?
*>                  ::                       0         32768 ?
*>i2001:C0:A8:D::/64
                    ::FFFF:172.22.20.20
                                             0    100      0 ?
*> 2002:AC:11:2::2/128
                    2001:C0:A8:B::1          0             0 65222 i
*>i2002:AC:13:4::4/128
                    ::FFFF:172.22.20.20
                                             0    100      0 65444 i
*> 2002:AC:15:A::A/128
                    ::                       0         32768 i
*>i2002:AC:16:14::14/128
                    ::FFFF:172.22.20.20
                                             0    100      0 i
*>i2003:C0:A8:E::/64
   Network          Next Hop            Metric LocPrf Weight Path
                    ::FFFF:172.22.20.20
                                             0    100      0 65444 ?
```

```
Rack-1-RPE1#sh ipv6 route
IPv6 Routing Table - default - 9 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       D - EIGRP, EX - EIGRP external, ND - Neighbor Discovery
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
C   2001:C0:A8:B::/64 [0/0]
     via GigabitEthernet0/2, directly connected
L   2001:C0:A8:B::2/128 [0/0]
     via GigabitEthernet0/2, receive
B   2001:C0:A8:D::/64 [200/0]
     via 172.22.20.20%default, indirectly connected
B   2002:AC:11:2::2/128 [20/0]
     via FE80::217:E0FF:FE69:24C0, GigabitEthernet0/2
B   2002:AC:13:4::4/128 [200/0]
     via 172.22.20.20%default, indirectly connected
LC  2002:AC:15:A::A/128 [0/0]
     via Loopback0, receive
B   2002:AC:16:14::14/128 [200/0]
     via 172.22.20.20%default, indirectly connected
B   2003:C0:A8:E::/64 [200/0]
     via 172.22.20.20%default, indirectly connected
L   FF00::/8 [0/0]
     via Null0, receive
```

```
Rack-1-RPE1#sh ipv6 cef
::/0
  no route
::/127
  discard
2001:C0:A8:B::/64
  attached to GigabitEthernet0/2
2001:C0:A8:B::1/128
  attached to GigabitEthernet0/2
2001:C0:A8:B::2/128
  receive for GigabitEthernet0/2
2001:C0:A8:D::/64
  nexthop 192.168.10.10 GigabitEthernet0/0 label 19 23
2002:AC:11:2::2/128
  nexthop FE80::217:E0FF:FE69:24C0 GigabitEthernet0/2
2002:AC:13:4::4/128
  nexthop 192.168.10.10 GigabitEthernet0/0 label 19 21
2002:AC:15:A::A/128
  receive for Loopback0
2002:AC:16:14::14/128
  nexthop 192.168.10.10 GigabitEthernet0/0 label 19 24
FE80::/10
  receive for Null0
FF00::/8
  multicast
```

Cisco 6PE - The Magic Command:

Essentially, one new command is added to advertise the "IPv6 + label" capability and to start 6PE operations.

In address-family IPv6 configuration mode, specify that "IPv6 + label" capability is advertised to a neighbor:

router(config-rotuer-af)#neighbor <ip-address> send-label [9]

A brief configuration details:

Rack-3-RPE2#sh running config
!
ipv6 cef
ip cef
!
mpls label protocol ldp
!
router bgp 65000
neighbor 172.21.10.10 remote-as 65000
 neighbor 172.21.10.10 update-source Loopback0
!
address-family ipv6
  neighbor 172.21.10.10 activate
  neighbor 172.21.10.10 send-label
 exit-address-family

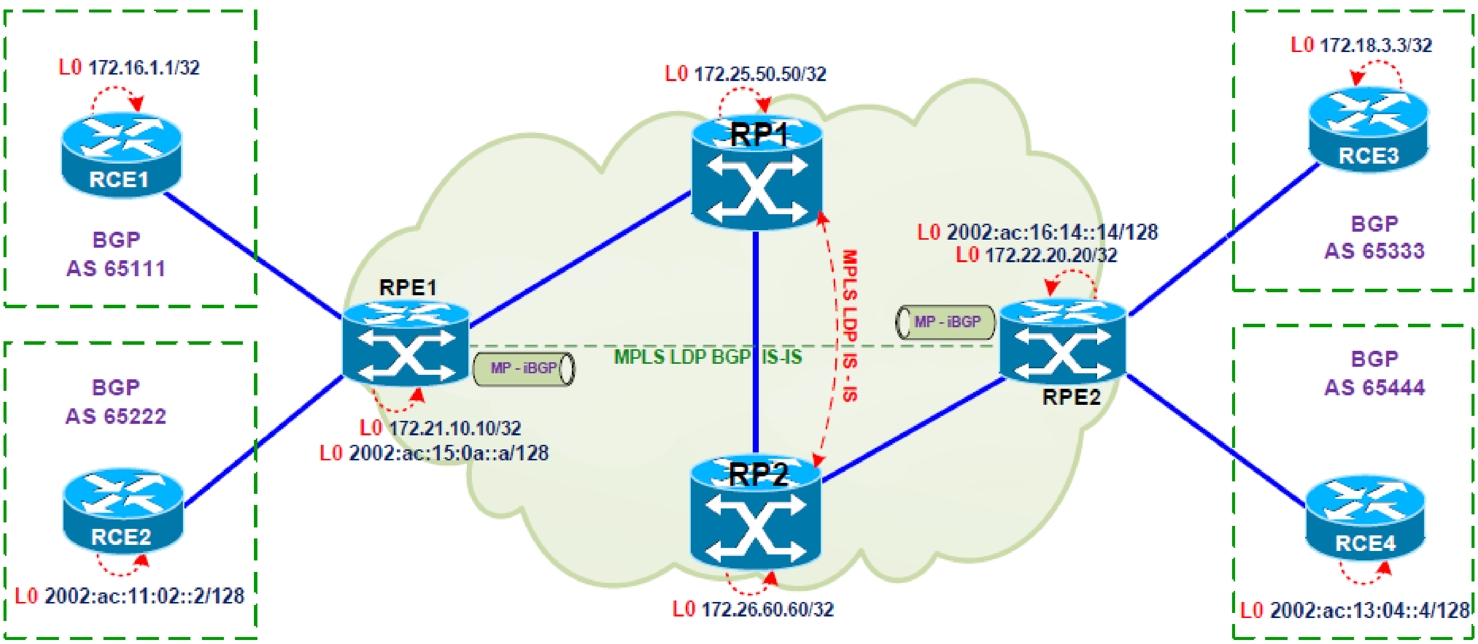

The following displays perform the details of Cisco 6PE BGP routes, which the ingress RPE1 learns from RPE2:

Show bgp ipv6 unicast 2002:AC:16:14::14/128

```
Rack-1-RPE1#sh bgp ipv6 unicast 2002:AC:16:14::14/128
BGP routing table entry for 2002:AC:16:14::14/128, version 5
Paths: (1 available, best #1, table default)
  Advertised to update-groups:
     2
  Local
    ::FFFF:172.22.20.20 (metric 24) from 172.22.20.20 (6.6.6.6)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      mpls labels in/out nolabel/24
```

On The ingress RPE1, the top of the stack label is shown by the following command, which describes the label switching operations for prefix 2002:AC:16:14::14/128:

Show bgp ipv6 unicast tags

```
Rack-1-RPE1#sh bgp ipv6 unicast tags
   Network          Next Hop      In tag/Out tag
   2001:C0:A8:B::/64
                    2001:C0:A8:B::1 22/notag
                    ::              22/notag
   2001:C0:A8:D::/64
                    ::FFFF:172.22.20.20
                                    notag/23
   2002:AC:11:2::2/128
                    2001:C0:A8:B::1 23/notag
   2002:AC:13:4::4/128
                    ::FFFF:172.22.20.20
                                    notag/21
   2002:AC:15:A::A/128
                    ::              21/notag
   2002:AC:16:14::14/128
                    ::FFFF:172.22.20.20
                                    notag/24
```

The CEF tables contain a two-label stack entry for this prefix:

Show ipv6 cef 2002:AC:16:14::14/128

```
Rack-1-RPE1#sh ipv6 cef 2002:AC:16:14::14/128
2002:AC:16:14::14/128
  nexthop 192.168.10.10 GigabitEthernet0/0 label 19 24
```

Finally, we can verify configurations from RCE1 for RCE3 through MPLS core. RCE1 and RCE3 are running IPv4 traffic only: ***RCE3's loopback address: 1.72.18.3.3/32***

```
RCE1#sh ip bgp
BGP table version is 12, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 172.16.1.1/32    0.0.0.0                  0         32768 i
*> 172.18.3.3/32    192.168.10.2                           0 65000 65333 i
*> 172.21.10.10/32  192.168.10.2             0             0 65000 i
*> 172.22.20.20/32  192.168.10.2                           0 65000 i
*  192.168.10.0/30  192.168.10.2             0             0 65000 ?
*>                  0.0.0.0                  0         32768 ?
*> 192.168.10.8/30  192.168.10.2             0             0 65000 ?
*> 192.168.10.16/30 192.168.10.2                           0 65000 ?
*> 192.168.10.20/30 192.168.10.2                           0 65000 ?
*> 192.168.13.0/27  192.168.10.2                           0 65000 65333 i
```

```
RCE1#sh ip route bgp
     192.168.13.0/27 is subnetted, 1 subnets
B       192.168.13.0 [20/0] via 192.168.10.2, 1d06h
     192.168.10.0/30 is subnetted, 4 subnets
B       192.168.10.8 [20/0] via 192.168.10.2, 1d06h
B       192.168.10.16 [20/0] via 192.168.10.2, 1d06h
B       192.168.10.20 [20/0] via 192.168.10.2, 1d06h
     172.18.0.0/32 is subnetted, 1 subnets
B       172.18.3.3 [20/0] via 192.168.10.2, 1d06h
     172.21.0.0/32 is subnetted, 1 subnets
B       172.21.10.10 [20/0] via 192.168.10.2, 1d06h
     172.22.0.0/32 is subnetted, 1 subnets
B       172.22.20.20 [20/0] via 192.168.10.2, 1d06h
```

```
RCE1#ping 172.18.3.3 rep 100

Type escape sequence to abort.
Sending 100, 100-byte ICMP Echos to 172.18.3.3, timeout is 2 seconds:
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
Success rate is 100 percent (100/100), round-trip min/avg/max = 1/1/4 ms
```

Verify configurations from RCE2 for RCE4 through MPLS core. RCE2 and RCE4 are running IPv6 traffic only: ***RCE4's loopback address: 2002:AC:13:4::4/128***

```
RCE2#sh bgp ipv6 unicast
BGP table version is 114, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*  2001:C0:A8:B::/64
                    2001:C0:A8:B::2          0             0 65000 ?
*>                  ::                       0         32768 ?
*> 2001:C0:A8:D::/64
                    2001:C0:A8:B::2                        0 65000 ?
*> 2002:AC:11:2::2/128
                    ::                       0         32768 i
*> 2002:AC:13:4::4/128
                    2001:C0:A8:B::2                        0 65000 65444 i
*> 2002:AC:15:A::A/128
                    2001:C0:A8:B::2          0             0 65000 i
*> 2002:AC:16:14::14/128
                    2001:C0:A8:B::2                        0 65000 i
*> 2003:C0:A8:E::/64
                    2001:C0:A8:B::2                        0 65000 65444 ?
```

```
RCE2#sh ipv6 route bgp
IPv6 Routing Table - 9 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       U - Per-user Static route, M - MIPv6
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
       D - EIGRP, EX - EIGRP external
B   2001:C0:A8:D::/64 [20/0]
     via FE80::523D:E5FF:FE7C:7482, GigabitEthernet0/0
B   2002:AC:13:4::4/128 [20/0]
     via FE80::523D:E5FF:FE7C:7482, GigabitEthernet0/0
B   2002:AC:15:A::A/128 [20/0]
     via FE80::523D:E5FF:FE7C:7482, GigabitEthernet0/0
B   2002:AC:16:14::14/128 [20/0]
     via FE80::523D:E5FF:FE7C:7482, GigabitEthernet0/0
B   2003:C0:A8:E::/64 [20/0]
     via FE80::523D:E5FF:FE7C:7482, GigabitEthernet0/0
```

```
RCE2#ping ipv6 2002:AC:13:4::4 rep 100

Type escape sequence to abort.
Sending 100, 100-byte ICMP Echos to 2002:AC:13:4::4, timeout is 2 seconds:
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
Success rate is 100 percent (100/100), round-trip min/avg/max = 0/1/4 ms
```

Calculating Voice, Video, and FTP packet:

Calculating Voice Packet Size for IPv4:

- If sample size is 20 ms (.02 seconds) and we're using the G.711 codec, then basic voice information requires .02 * 64000 / 8 = 160 bytes per sample.
- Ethernet II header = 18 bytes (included 4 bytes CRC)
- IPv4 header = 20 bytes.
- UDP header = 8 bytes.
- RTP header = 12 bytes.

So, calculated *VOICE* packet size = 218 bytes.

Calculating Video Packet Size for IPv4:

- MPEG-2 Transports are composed of 188 byte TS Packets, and each with a 4 byte header. TS packet payloads may contain program information as well as Packetized Elementary s (PES), typically video and audio streams. PES packets are broken into 184 byte chunks to fit into the TS packet payload. [12]
- If MPEG-2 TS packet payload 184 bytes.
- Ethernet II header = 18 bytes (included 4 bytes CRC).
- IPv4 header = 20 bytes.
- UDP header = 8 bytes.

So, calculated *VIDEO* packet size = 230

Note: RTP header didn't include on VIDEO packet, because this is an assumption for video through-put.

Calculating FTP Packet Size for IPv4:

- FTP payload size 412 bytes (this is assumption for data through-put)
- Ethernet II header = 18 bytes (included 4 bytes CRC).
- IPv4 header = 20 bytes.
- TCP header = 20 bytes.

So, calculated *FTP* packet size = 452

Total Voice Packet Size for IPv6: = 238 Bytes (IPv6 header = 40 bytes)
(Payload + Ethernet II header (include CRC) + IPv6 header + UDP header + RTP header)

Total Video Packet Size for IPv6: = 250 Bytes (IPv6 header = 40 bytes)
(Payload + Ethernet II header (include CRC) + IPv6 header + UDP header)

Total FTP Packet Size for IPv6: = 472 Bytes (IPv6 header = 40 bytes)
(Payload + Ethernet II header (include CRC) + IPv6 header + TCP header)

IxExplorer Operation:

IxEplorer provides a powerful and interactive Graphical User Interface (GUI) for managing Ixia test hardware resources.

- Complete control is provided for generation and analysis of Layer 1 - 4 traffic streams.
- True interface emulation with ARP, Ping (IPv4 and IPv6).
- Fine tune control of stream parameters such as frame size, data patterns, protocols, and rates.
- Capabilities to protocols and encapsulations, such as IPv4, IPv6, TCP, UDP, IPX, VLAN, Cisco ISL, ARP, ICMP, IGMPv1/v2/v3, MPLS, GRE, OSPF, RIP, DHCP.

Operating Mode:

- IxExplorer can be operated locally (on chassis desktop) or remotely via client software.
- On opening IxExplorer you will be prompted to enter the IP address of the chassis. For local operation acceptable addresses are:

  - 127.0.01
  - The Chassis IP address or the word: "loopback"

- For remotely operation you must enter the Chassis IP address only.

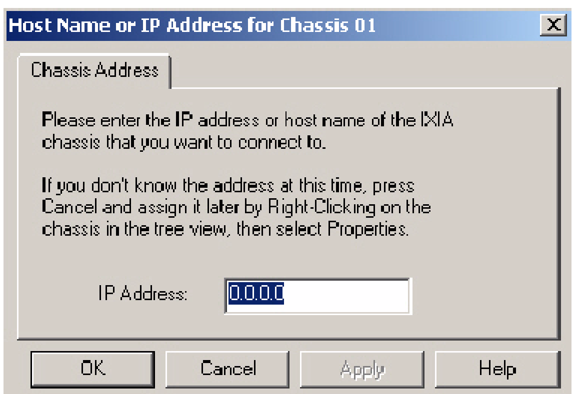

The Interface

- The interface consists of a Title Bar, Menus, Tool Bar, Transmit Bar, Windows Area, Status Window and Status Bar.

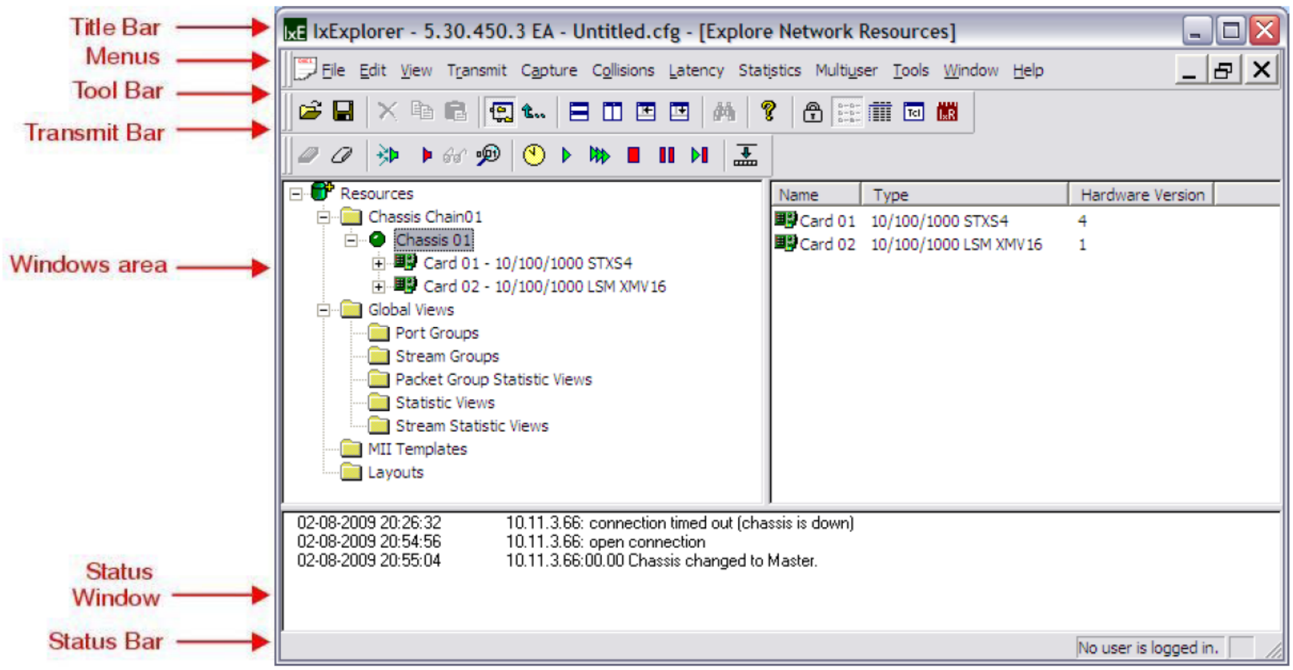

Window Area: Resource Tree

- The resource tree is a tree view of available Chassis, cards and ports available.
- The tree view displays ownership of ports in brackets.
- If port is already owned by another user the settings to adjust the packet streams will be grayed out.

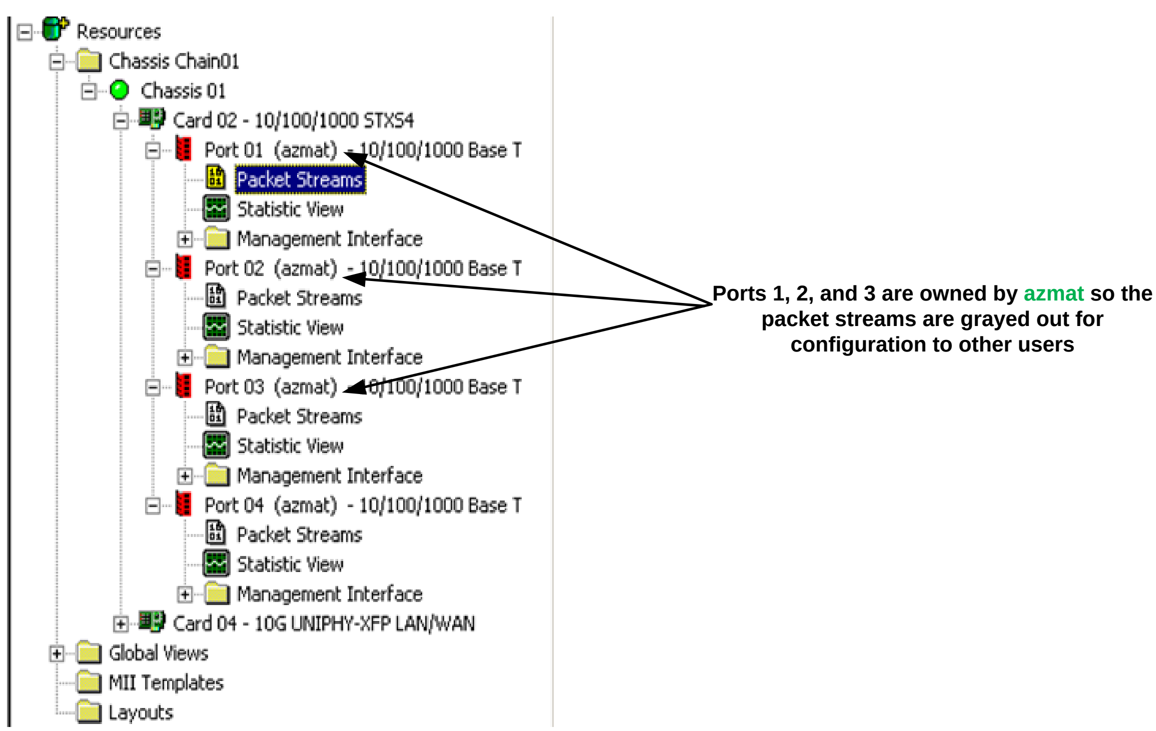

- The property window has four main tabs (Frame Data, Control, Packet View, and Warnings).

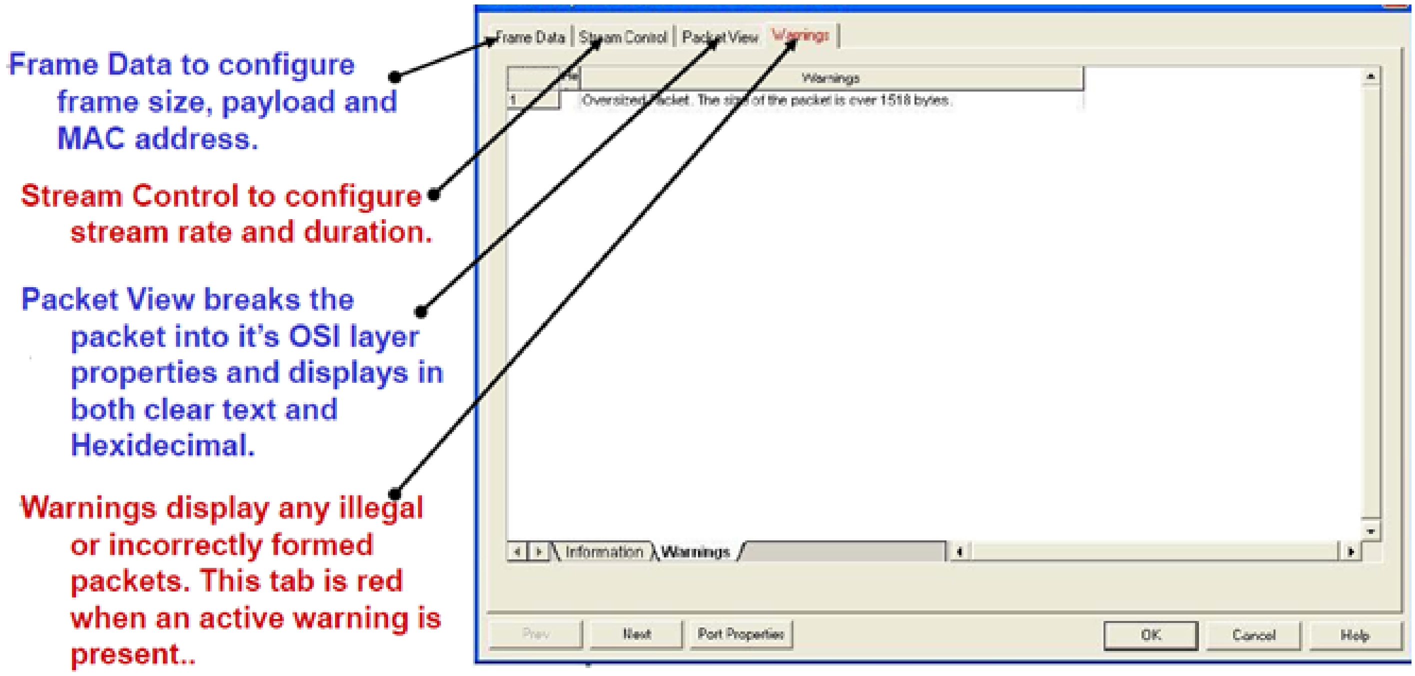

- The Frame Data tab allows to set the preamble size, MAC address, and to insert Time Stamp, Packet Group, Sequence Checking etc.

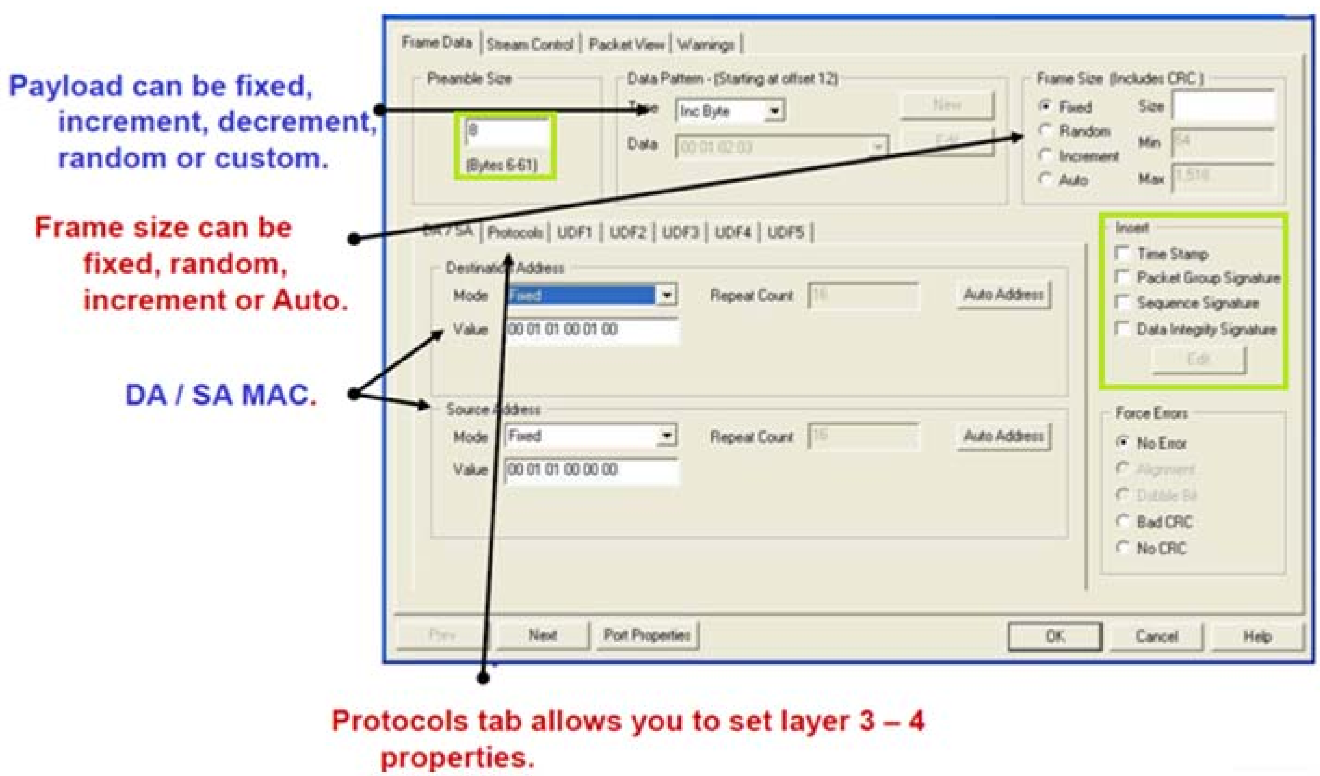

Packet Stream

- Streams are added as line items much like in an excel spread sheet. Right clicking a stream gives you the ability to copy, paste, duplicate or insert a new stream with default values.

	Name	Enable	Control	Flow	Loop Count	Suspend	Frame Size	Data Pattern Type	L3 Protocol	L4 Protocol	IP Source Address	IP Dest Address
1	VOICE IPv4	☑	End		1	☐	218	Inc Byte	IPv4	UDP	192.168.11.2	192.168.13.2
2	VOICE IPv6	☐	Disabled		1	☐	238	Inc Byte	IPv6	UDP	2003:C0:A8:C:0:0:0:2	2003:C0:A8:E:0:0:0:2

- Double clicking an individual packet brings up the stream properties window, allowing you to configure the packet properties. *(VOICE-IPv4 Frame Size)*
- The Protocols tab: Can set layer 3 to 4 parameters such as IPv4/v6, UDP etc., and Edit option allow to configure Source and Destination IP addresses.

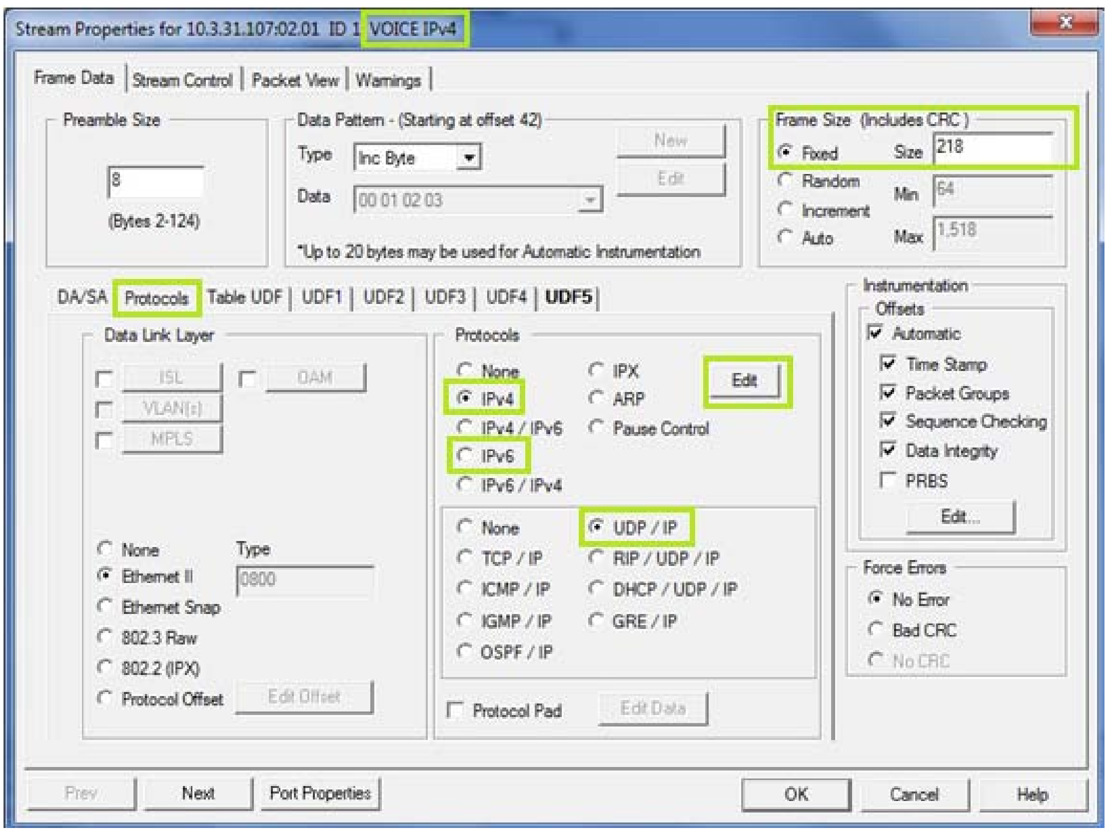

- It can see: Name of traffic, Frame Size, and Source / Destination IP and MAC Addresses are configured through Properties.

Voice Frame Size:

	Name	Enable	Control	Flow	Loop Count	Suspend	Frame Size	Data Pattern Type	L3 Protocol	L4 Protocol	IP Source Address	IP Dest Address
1	VOICE IPv4	☑	End		1	☐	218	Inc Byte	IPv4	UDP	192.168.11.2	192.168.13.2
2	VOICE IPv6	☐	Disabled		1	☐	238	Inc Byte	IPv6	UDP	2003:C0:A8:C:0:0:0:2	2003:C0:A8:E:0:0:0:2

Video Frame Size:

	Name	Enable	Control	Flow	Loop Count	Suspend	Frame Size	Data Pattern Type	L3 Protocol	L4 Protocol	IP Source Address	IP Dest Address
1	VIDEO IPv4	☑	End		1	☐	230	Inc Byte	IPv4	UDP	192.168.11.2	192.168.13.2
2	VIDEO IPv6	☐	Disabled		1	☐	250	Inc Byte	IPv6	UDP	2003:C0:A8:C:0:0:0:2	2003:C0:A8:E:0:0:0:2

FTP Frame Size:

	Name	Enable	Control	Flow	Loop Count	Suspend	Frame Size	Data Pattern Type	L3 Protocol	L4 Protocol	IP Source Address	IP Dest Address
1	FTP IPv4	☑	End		1	☐	452	Inc Byte	IPv4	TCP	192.168.11.2	192.168.13.2
2	FTP IPv6	☐	Disabled		1	☐	472	Inc Byte	IPv6	TCP	2003:C0:A8:C:0:0:0:2	2003:C0:A8:E:0:0:0:2

- The DA/SA tab: We can select ARP / Discovery mode to get destination and source MAC addresses automatically, and From Protocol Interfaces check mark can also get source IP address. The following windows are an example for both IPv4 and IPv6.

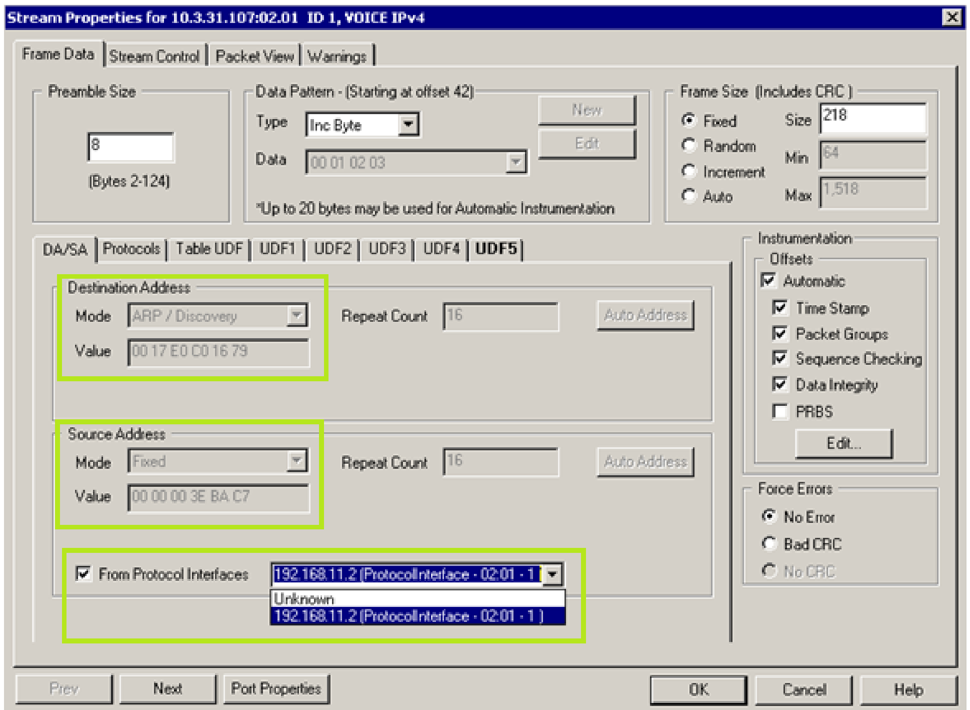

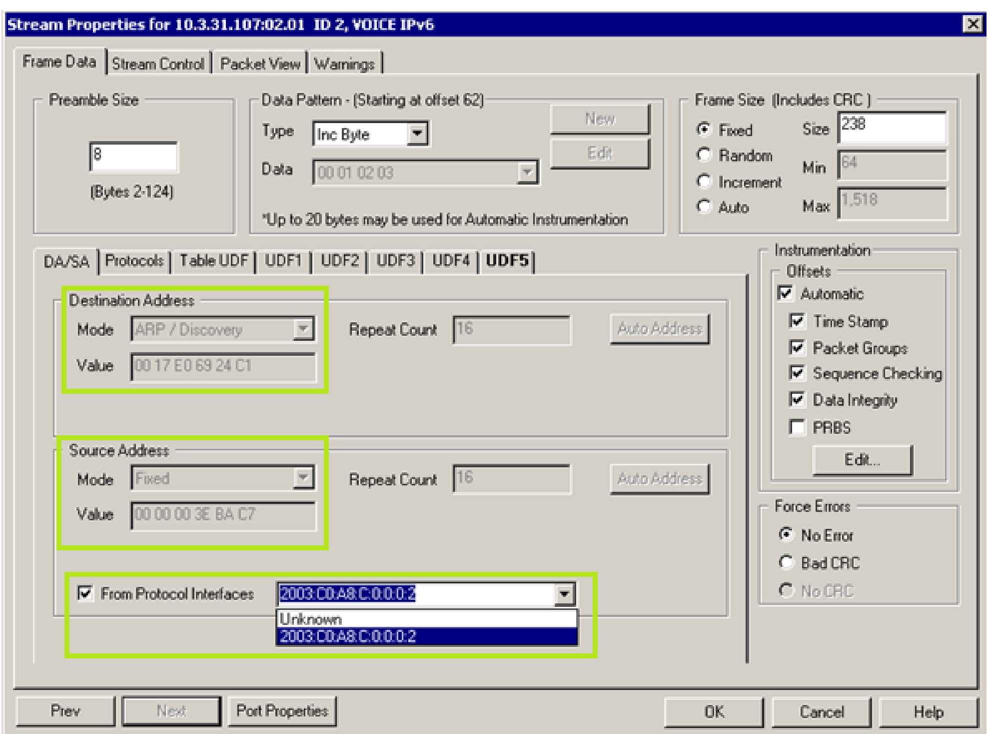

Packet Streams: Frame Data Tab / Protocols / Edit IP / Edit UDP and TCP

- The IP Header sets priority (ToS or DSCP QoS), Destination and Source IP addresses, UDP/TCP ports number, Sequence number, Acknowledgement number, Header Length, Window size, and all relevant Flags.

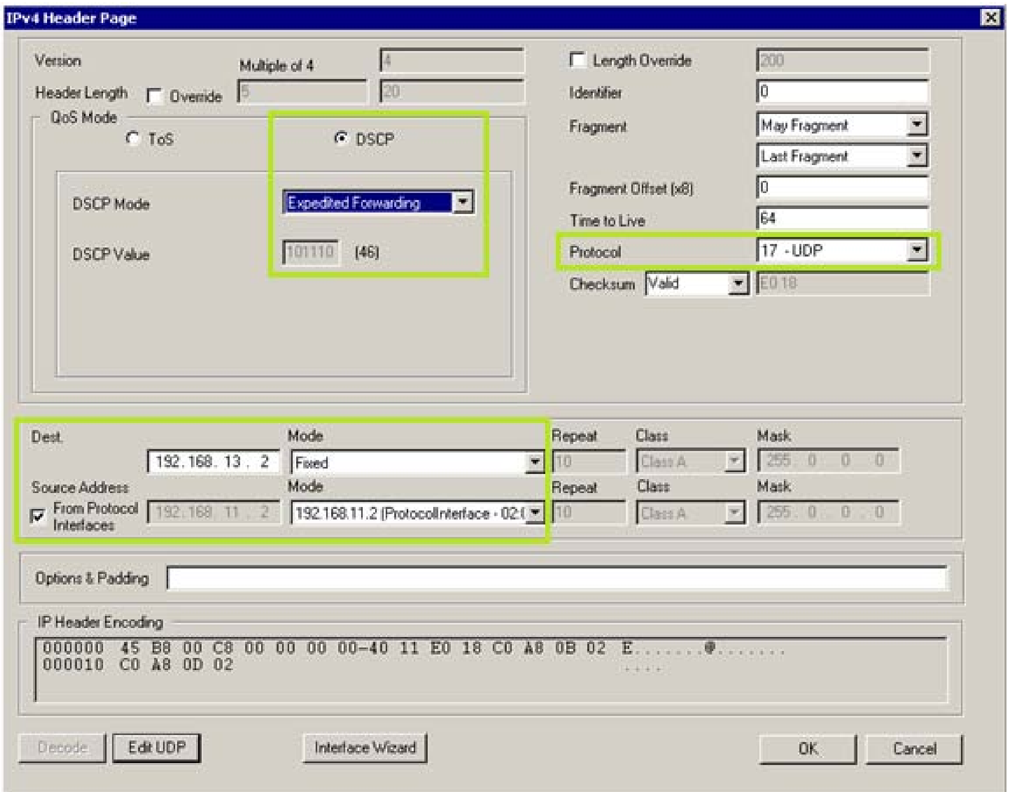

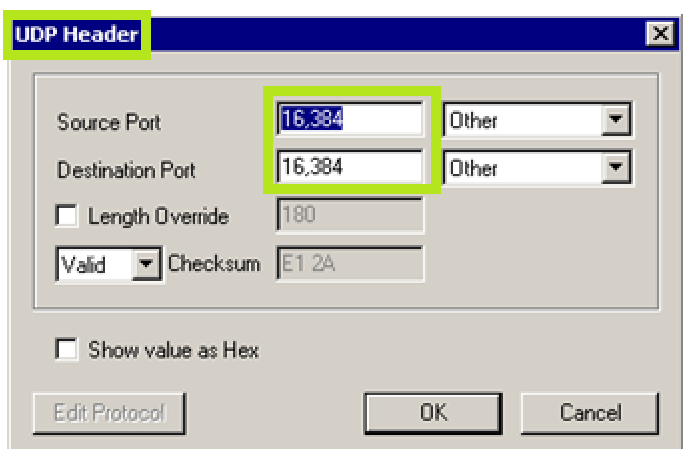

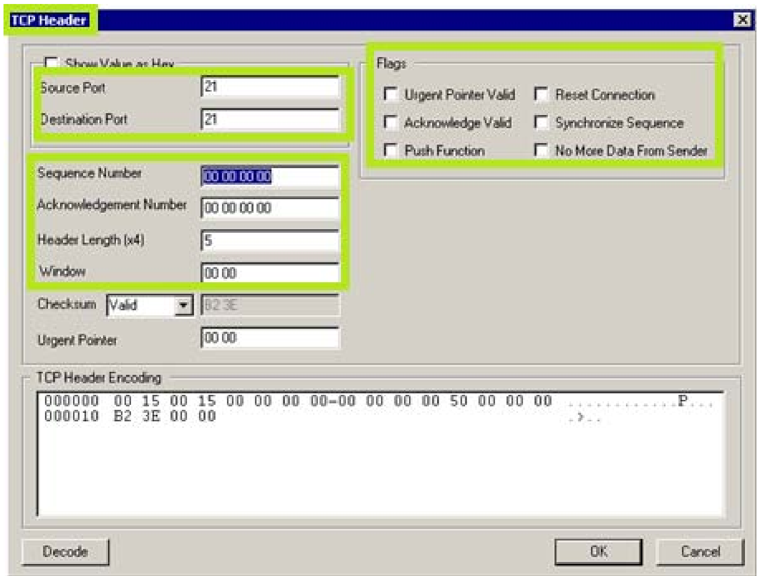

- The IPv6 Header allows one to set Traffic Class, Flow Label, Hop Limit, Destination and Source IP addresses with Prefix such as Global Unicast, Link Local etc., and all other Available Extension Headers.

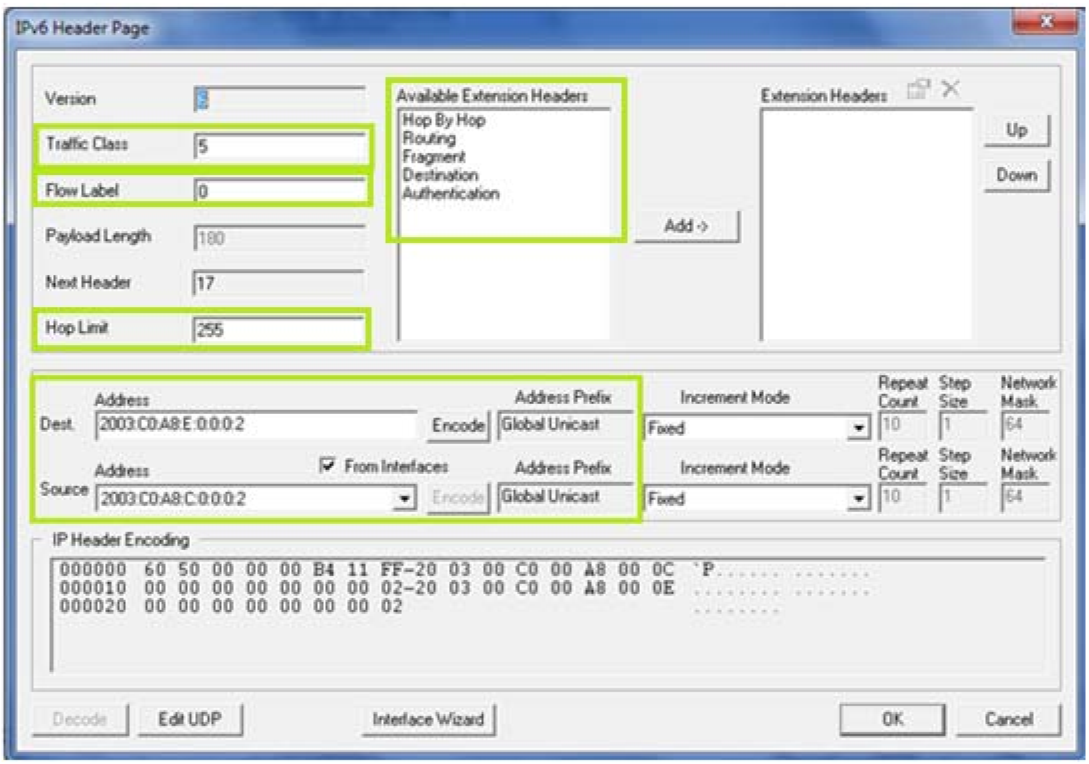

Packet Stream: Control Tab

- Control tab allows one to set parameters, such as Name of stream, Stop after this Stream (transmit one time then stop), Packet per Burst (how many packets will be transmitting at a time), and Rate Control (percentage of bandwidth for this stream).
- The following same parameters are also set for VOICE-IPv6.

- VIDEO's parameters are set ***Packets per Burst = 4000*** and ***Rate Control = 20 %*** for both IPv4 and IPv6.
- FTP's parameters are set ***Packets per Burst = 3000*** and ***Rate Control = 10 %*** for both IPv4 and IPv6.

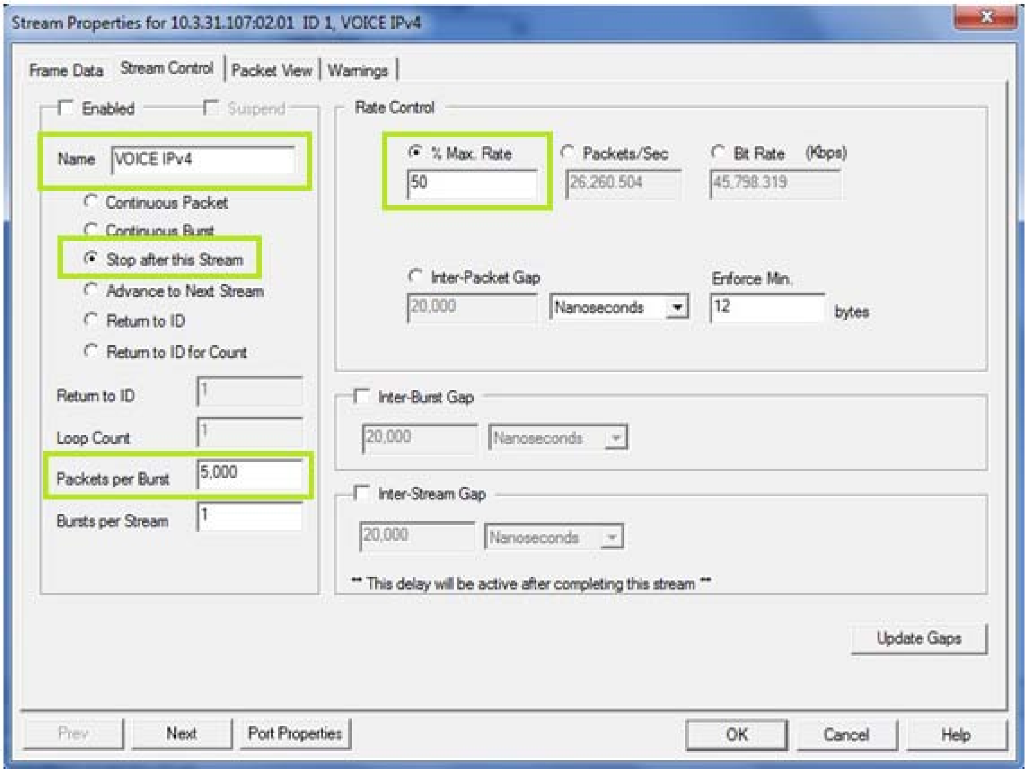

- The available stream termination controls are continuous packet, continuous burst, stop after this stream, advance to next stream return to ID and return to ID for count to set options for generating traffic.

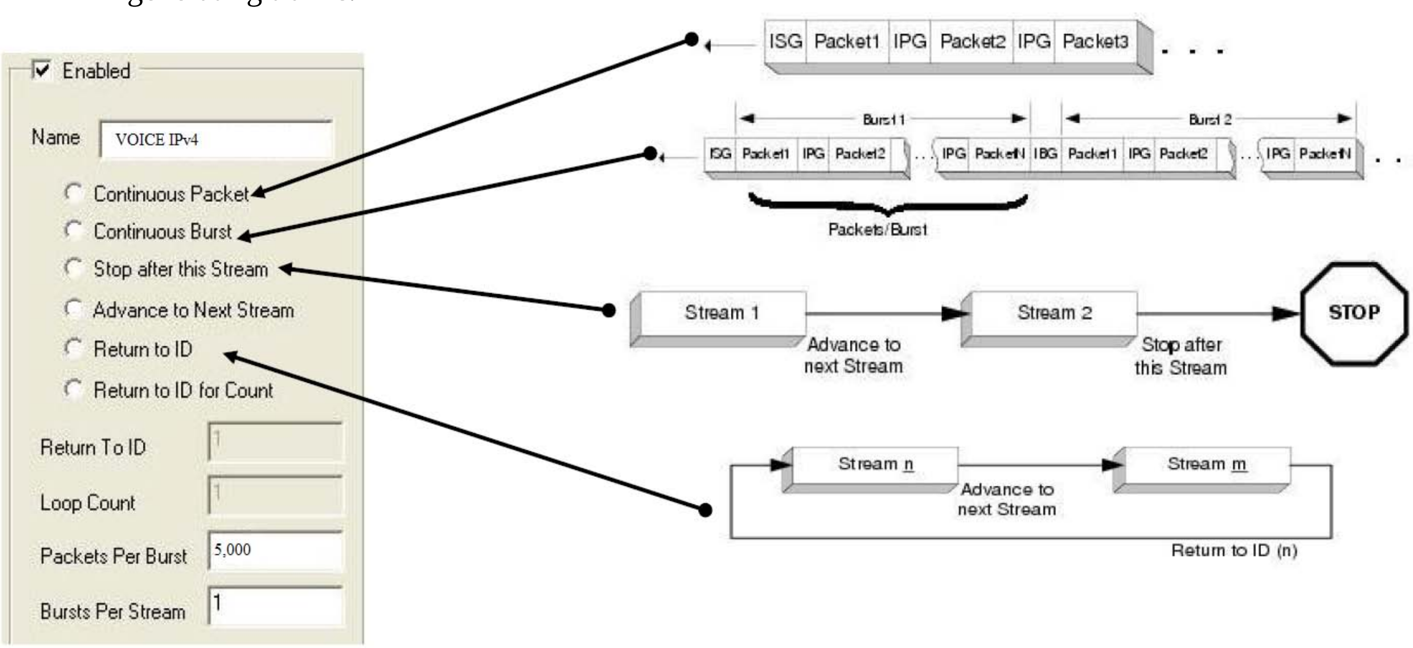

Practical Session:

- Created a single packet stream on each individual port (ports 1, 2, & 3) with the following parameters:
  - ➢ MAC Destination Address: 00 17 E0 C0 16 79 *(Router RCE1 IPv4)*
  - ➢ MAC Destination Address: 00 17 E0 69 24 C1 *(Router RCE2 IPv6)*
  - ➢ MAC Source Address: 00 00 00 3E BA C7 *(IXIA Port-**1 VOICE** IPv4 - IPv6)*
  - ➢ MAC Source Address: 00 00 00 38 70 30 *(IXIA Port-**2 VIDEO** IPv4 - IPv6)*
  - ➢ MAC Source Address: 00 00 00 3A 41 16 *(IXIA Port-**3 FTP** IPv4 - IPv6)*
  - ➢ Frames Size (Fixed): 218 Bytes *(VOICE IPv4)*
  - ➢ Frames Size (Fixed): 230 Bytes *(VIDEO IPv4)*
  - ➢ Frames Size (Fixed): 452 Bytes *(FTP IPv4)*
  - ➢ Frames Size (Fixed): 238 Bytes *(VOICE IPv6)*
  - ➢ Frames Size (Fixed): 250 Bytes *(VIDOE IPv6)*
  - ➢ Frames Size (Fixed): 472 Bytes *(FTP IPv6)*
  - ➢ Destination IP Address: 192.168.13.2 *(IXIA Port-**4** Receive Mode for all three traffic sending (ports 1, 2, & 3) of IPv4).*
  - ➢ Destination IP Address: 2003:C0:A8:0E::2 *(IXIA Port-**4** Receive Mode for all three traffic sending (ports 1, 2, & 3) of IPv6).*
  - ➢ Source IP Address: 192.168.11.2 *(Router RCE1 Interface)*
  - ➢ Source IP Address: 2003:C0:A8:0C::2 *(Router RCE2 Interface)*
  - ➢ UDP SA/DA Ports: 16384 *(VOICE for both IPv4/v6)*
  - ➢ UDP SA/DA Ports: 970 *(VIDEO for both IPv4/v6)*
  - ➢ TCP SA/DA Ports: 21 *(FTP for both IPv4/v6)*
- The following above mentioned parameters we can verify on the Stream Properties / Packet View Tab.

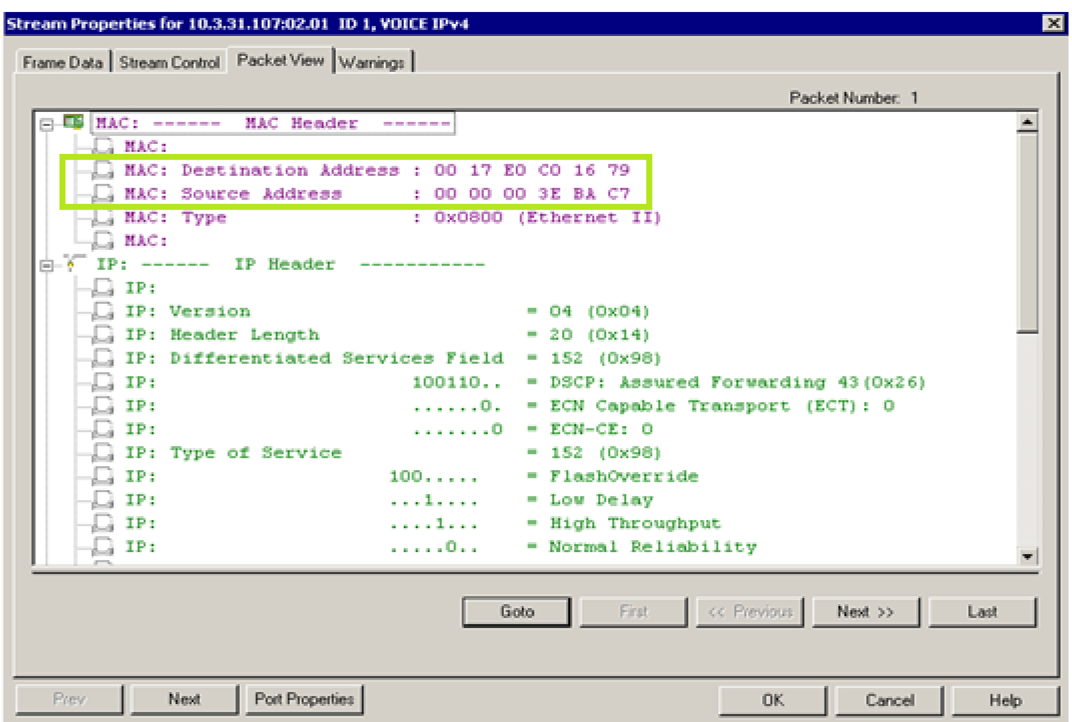

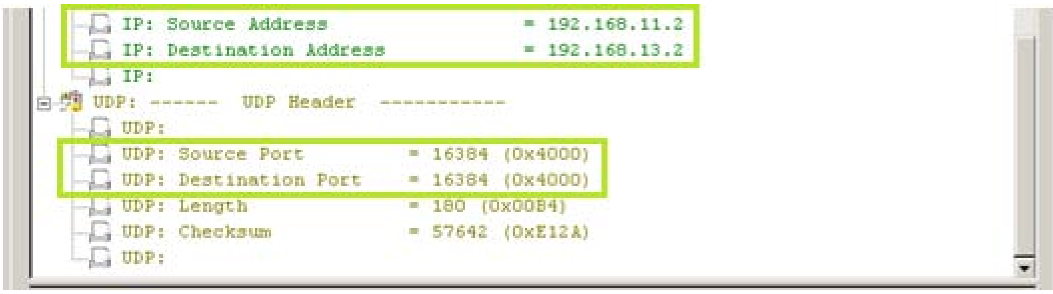

This is for IPv6 Packet View:

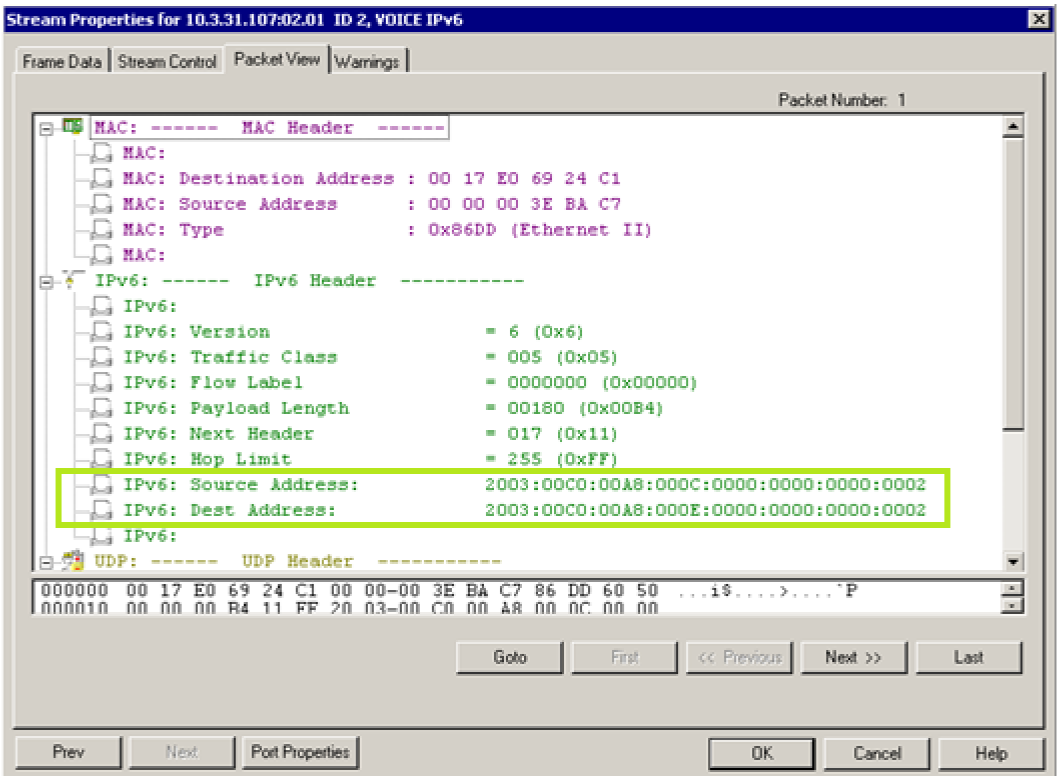

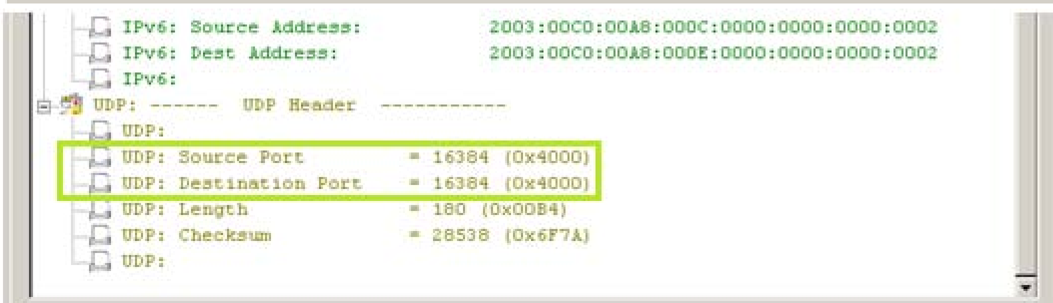

Statistic Overview

- Statistic views allows one to see the received data on a port.
- Statistics are closely tied to capture filters and received modes of the port.
- Statistics are available as total count and rate per second.
- Statistics can be used to generate logging and alerts.
- Statistics can be viewed as Graph and used to determine latency and packet sequece.
- Statistics can be viewed with group ports together at a time.

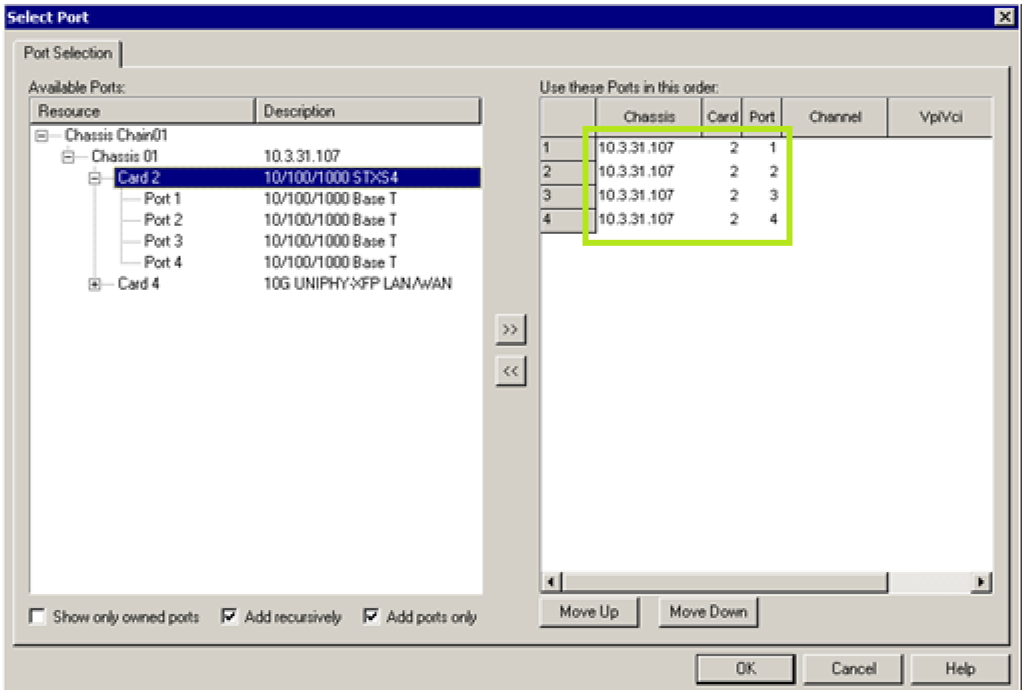

After one has added all ports on the Select Port property as above mentioned, they can be viewed on the following StatView's options.

- Now we can see multiple ports statistics simultaneously.

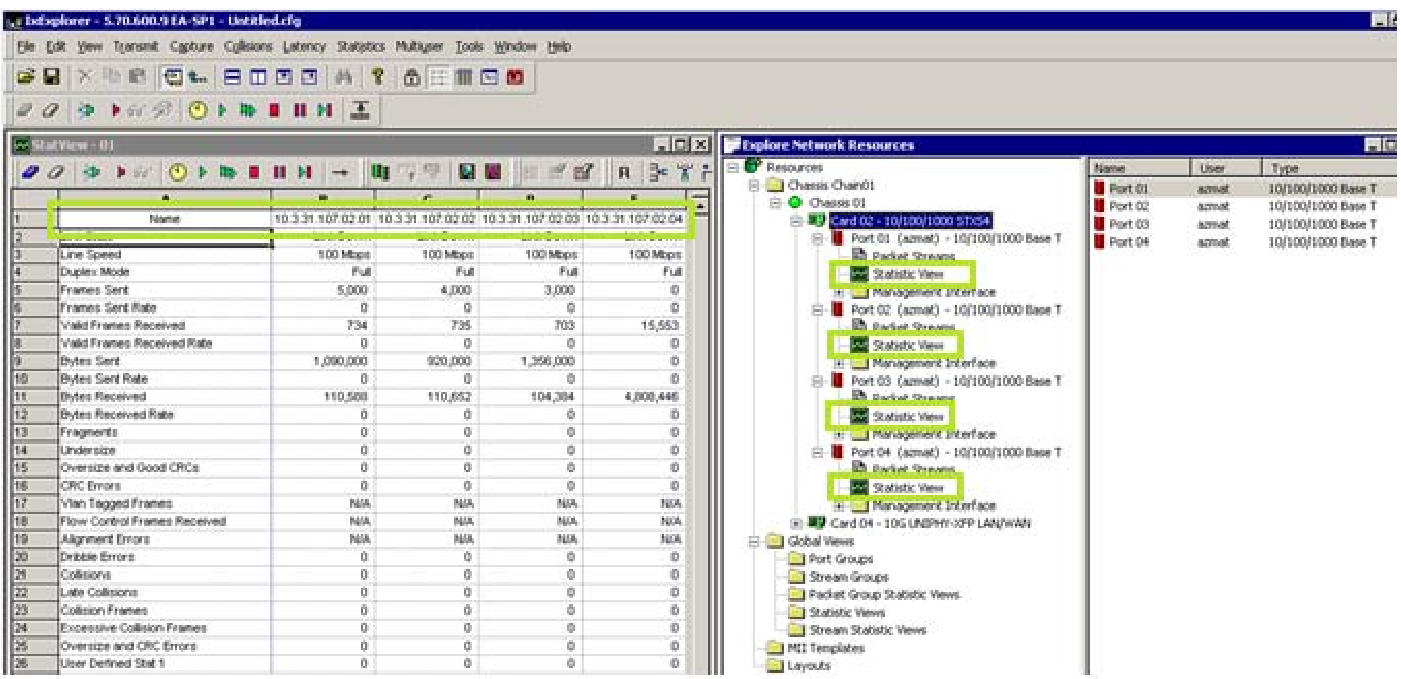

- Receive Mode on port 4 has been selected to capture all traffic from ports 1, 2, & 3 simultaneously, and saved to the External Protocol Analyzer for analysis of generated traffics.

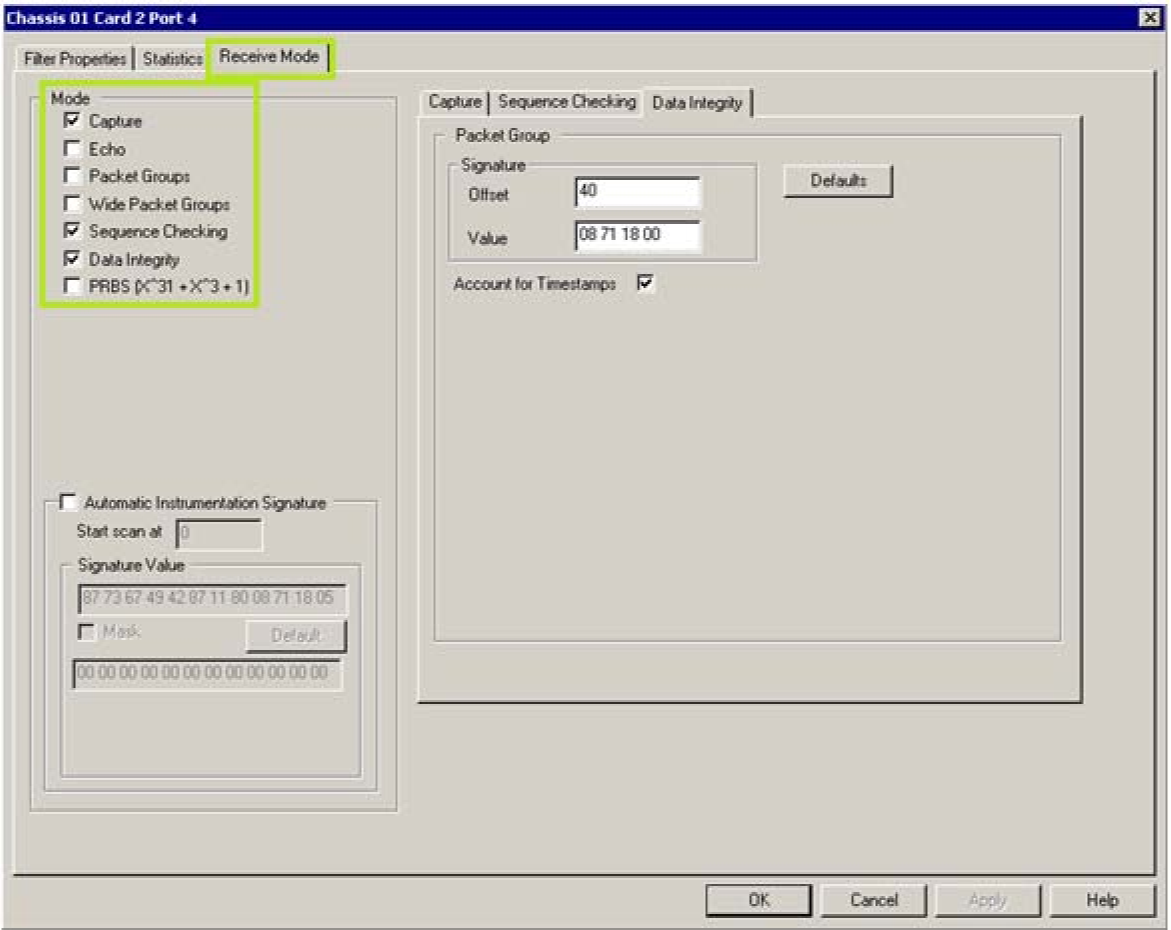

Experimental Scenario for capturing traffics:

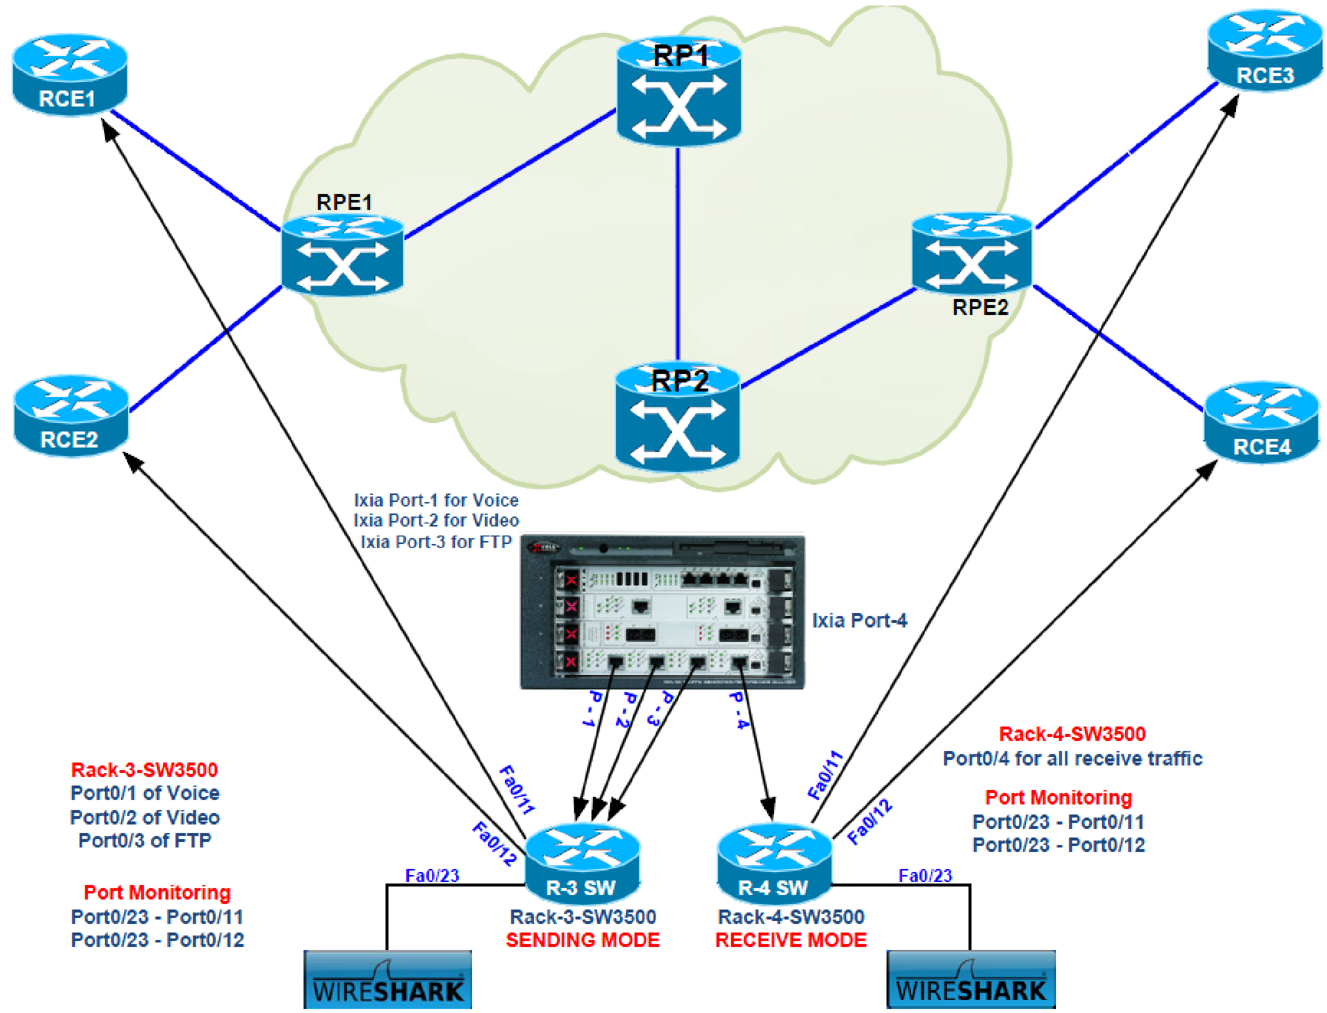

Over all packets results of IPv4 from IXIA's StatView option:

	A	B	C	D	E
1	Name	10.3.31.107:02.01	10.3.31.107:02.02	10.3.31.107:02.03	10.3.31.107:02.04
2	Link State	Link Up	Link Up	Link Up	Link Up
3	Line Speed	100 Mbps	100 Mbps	100 Mbps	100 Mbps
4	Duplex Mode	Full	Full	Full	Full
5	Frames Sent	5,000	4,000	3,000	0
6	Frames Sent Rate	0	0	0	0
7	Valid Frames Received	18	18	18	11,832
8	Valid Frames Received Rate	0	0	0	0
9	Bytes Sent	1,090,000	920,000	1,356,000	0
10	Bytes Sent Rate	0	0	0	0
11	Bytes Received	3,172	3,172	3,172	3,322,707
12	Bytes Received Rate	0	0	0	0
13	Fragments	0	0	0	0
14	Undersize	0	0	0	0
15	Oversize	0	0	0	0
16	CRC Errors	0	0	0	0

Over all packets results of IPv6 from IXIA's StatView option:

	A	B	C	D	E
1	Name	10.3.31.107:02.01	10.3.31.107:02.02	10.3.31.107:02.03	10.3.31.107:02.04
2	Link State	Link Up	Link Up	Link Up	Link Up
3	Line Speed	100 Mbps	100 Mbps	100 Mbps	100 Mbps
4	Duplex Mode	Full	Full	Full	Full
5	Frames Sent	5,000	4,000	3,000	3
6	Frames Sent Rate	0	0	0	0
7	Valid Frames Received	21	21	21	12,024
8	Valid Frames Received Rate	0	0	0	0
9	Bytes Sent	1,190,000	1,000,000	1,416,000	254
10	Bytes Sent Rate	0	0	0	0
11	Bytes Received	3,306	3,306	3,306	3,609,652
12	Bytes Received Rate	0	0	0	0
13	Fragments	0	0	0	0
14	Undersize	0	0	0	0
15	Oversize	0	0	0	0
16	CRC Errors	0	0	0	0

- The following statistics are captuered from open source the Wirehark Network Analyzer version 1.4.4.

Voice Packets Sent

Wireshark: Protocol Hierarchy Statistics

Display filter: none

Protocol	% Packets	Packets	% Bytes	Bytes	Mbit/s	End Packets	End Bytes	End Mbit/s
⊞ Frame	100.00 %	5000	100.00 %	1070000	44.971	0	0	0.000

Voice Packets Received

Wireshark: Protocol Hierarchy Statistics

Display filter: none

Protocol	% Packets	Packets	% Bytes	Bytes	Mbit/s	End Packets	End Bytes	End Mbit/s
⊞ Frame	100.00 %	4875	100.00 %	1043250	43.112	0	0	0.000

Video Packets Sent

Wireshark: Protocol Hierarchy Statistics

Display filter: none

Protocol	% Packets	Packets	% Bytes	Bytes	Mbit/s	End Packets	End Bytes	End Mbit/s
⊞ Frame	100.00 %	4000	100.00 %	904000	18.087	0	0	0.000

Video Packets Received

Wireshark: Protocol Hierarchy Statistics

Display filter: none

Protocol	% Packets	Packets	% Bytes	Bytes	Mbit/s	End Packets	End Bytes	End Mbit/s
⊞ Frame	100.00 %	3964	100.00 %	895864	17.934	0	0	0.000

FTP Packets Sent

Wireshark: Protocol Hierarchy Statistics

Display filter: none

Protocol	% Packets	Packets	% Bytes	Bytes	Mbit/s	End Packets	End Bytes	End Mbit/s
⊞ Frame	100.00 %	3000	100.00 %	1344000	9.495	0	0	0.000

FTP Packets Received

Wireshark: Protocol Hierarchy Statistics

Display filter: none

Protocol	% Packets	Packets	% Bytes	Bytes	Mbit/s	End Packets	End Bytes	End Mbit/s
⊞ Frame	100.00 %	2975	100.00 %	1332800	9.418	0	0	0.000

IPv4 Frames Lost Statistics

Name	Port - 1	Port - 2	Port - 3	Port - 4	Total
Link State	Link Up	Link Up	Link Up	Link Up	
Line Speed	100 Mbps	100 Mbps	100 Mbps	100 Mbps	
Frames Sent	5,000	4,000	3,000	0	**12,000**
Frames Received	4,875	3,964	2,975	11,814	**11,814**
Frames Lost	**125**	**36**	**25**		**186**
Bytes Sent	1,070,000	904,000	1,344,000	0	**3,318,000**
Bytes Received	1,043,250	895,864	1,332,800	3,271,914	**3,271,914**
Bytes Lost					**46,086**

IPv6 Frames Lost Statistics

Name	Port - 1	Port - 2	Port - 3	Port - 4	Total
Link State	Link Up	Link Up	Link Up	Link Up	
Line Speed	100 Mbps	100 Mbps	100 Mbps	100 Mbps	
Frames Sent	5,000	4,000	3,000	0	**12,000**
Frames Received	5,000	4,000	3,000	12,000	**12,000**
Frames Lost	**0**	**0**	**0**		**0**
Bytes Sent	1,190,000	1,000,000	1,416,000	0	**3,606,000**
Bytes Received				3,606,000	**3,606,000**
Bytes Lost	**0**	**0**	**0**		**0**

- We can see captured voice traffic on Wireshark to verify voice payload.
- I made a voice packet size 218 bytes (including Ethernet II + IP + UDP + RTP + Payload) headers with payload of 160 bytes of voice's sample.
- 18 bytes + 20 bytes + 8 bytes + 12 bytes + 160 bytes = 218 Bytes.
- Now, we can see payload 172 bytes (included Data 160 bytes + RTP 12 bytes)
- Rest of other headers used via IXIA's IxEplorer, because this is support Layer 1 - 4:

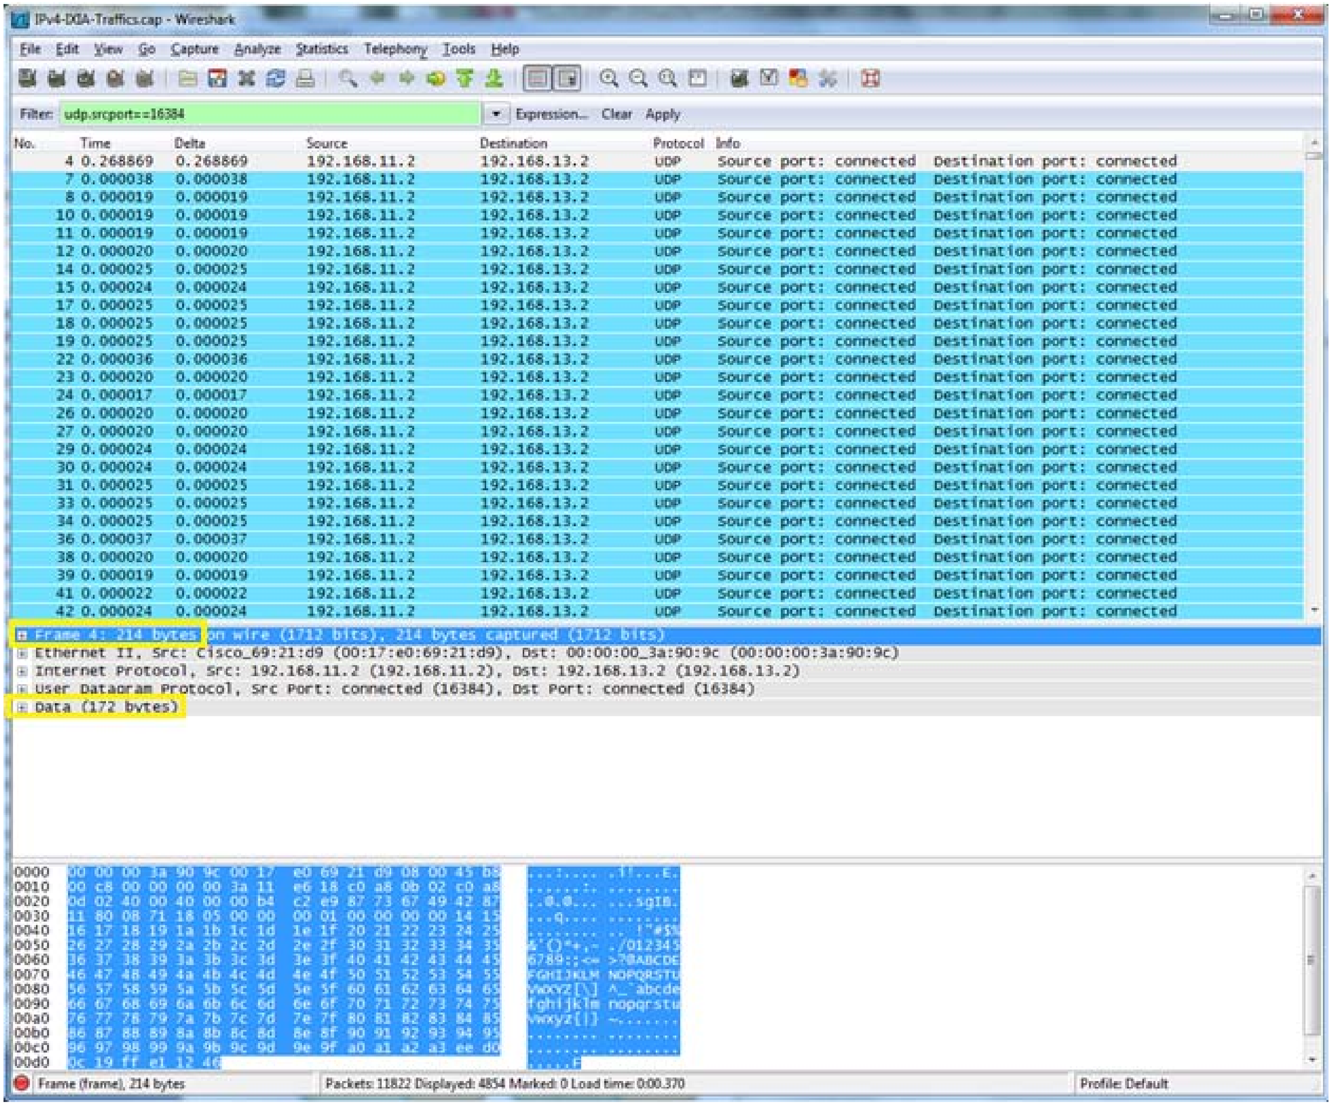

Note: Total Frames per packet shows 214 bytes rather than 218 bytes because this captured traffic exports from Ixia device to External Protocol Analyzer (Wireshark), therefore, there is no 4 bytes sequence number included.

Conclusion:

Based on given Cisco devices, IOSs version, IXIA's platform 400T with IxExplorer ver.5.70, and Configurations:

The following statistics mentioned above shows packets lost on an IPv4 network environment, while on an IPv6 network environment there is no lost packets found.

IxExplorer is capable to support up-to Layer 4, while RTP is a Session Layer (OSI Layer 5) protocol that lies on top of UDP (an OSI Layer 4 protocol). RTP is specifically designed to deliver streaming audio and video content on time and in order. Utilizing UDP for its unreliable time-conscious transmission methods, RTP ensures that packets reach the end node's application both in a timely manner and in the originally intended order.

- The following diagram shows limitation of layer protocols.

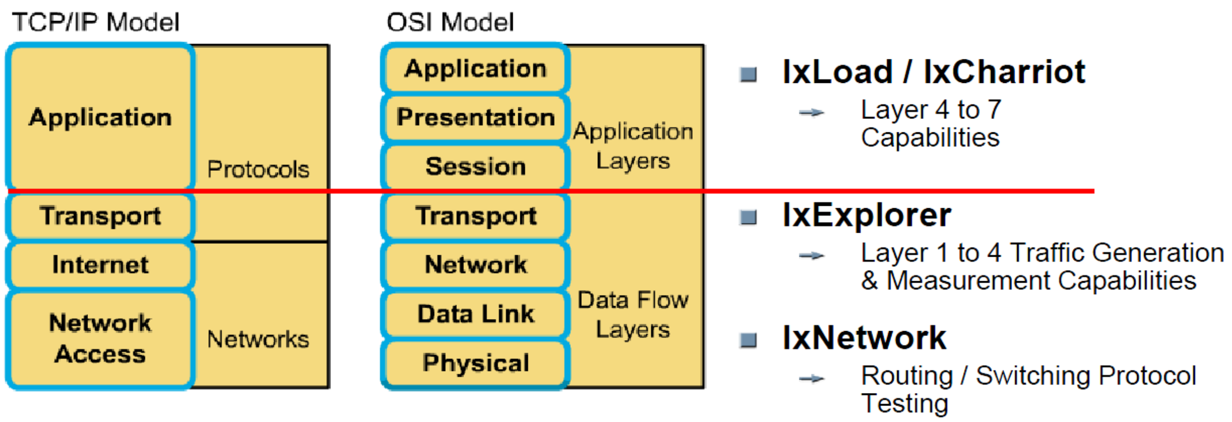

Note: I'm unable to take delay and jitter from the voice traffic, because I made voice stream, including RTP header's size on the payload to Ixia device. In addition, this is not actually voicing packet with RTP header.

Appendix - A

Bibliography:

[1] RFC 2460: Internet Protocol, Version 6 (IPv6) Specification. S. Deering , R. Hinden, December 1998.

[2] Handbook of IPv4 to IPv6 Transition: Methodologies for Institutional and Corporate Networks. John J. Amoss and Daniel Minoli.

[3] IPv6 Essentials, 2nd Edition, May 2006: Published by O'Reilly by Silvia Hagen, Figure from Chapter 2: The Structure of the IPv6 Protocol, Section 2.3 Extension Headers.

[4] IPexpert's IPv6 eBook: powered by Proctor Labs, 2005.

[5] IPv6 Tutorial Presentation 2001: RIPE 40 Meeting Prague, Czech Republic.

[6] Cisco Press: Implementing Cisco MPLS 2006, Volume 1, Version 2.2, Student Guide.

[7] RFC - 2545: Use of BGP - 4 Multiprotocol Extensions for IPv6 Inter-Domain Routing, P. Margques, F. Dupont, March 1999.

[8] RFC - 2283: Multiprotocol Extensions for BGP - 4, T. Bates, R. Chandra, D. Katz, Y. Rekhter, 1998.

[9] IPv6 over MPLS (6PE): Application Note, by Cisco System, Inc. 2002.

[10] RFC - 3107: Carrying Label Information in BGP - 4, Y. Rekhter, E. Rosen, May 2001.

[11] http://www.ixiacom.com/solutions/testing_routing_and_switching/index.php.

[12] Whitepaper - IP ing of MPEG-4: Native RTP vs MPEG-2 Transport , by Alex MacAulay, Boris Felts, Yuval Fisher, October 2005.

Appendix - B

Routers Configuration:

```
RCE1#sh run
Building configuration...

Current configuration : 2445 bytes
!
version 12.4
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
!
hostname RCE1
!
boot-start-marker
boot-end-marker
!
memory-size iomem 10
!
ip cef
!
no ip domain lookup
!
class-map match-all MARK-VIDEO
 match access-group 102
class-map match-all MARK-VOICE
 match access-group 101
class-map match-all MARK-FTP
 match protocol ftp
!
policy-map MARKING-TRAFFIC
 class MARK-VOICE
  set ip dscp ef
 class MARK-VIDEO
  set ip dscp af13
 class MARK-FTP
  set ip dscp af11
policy-map LLQ
 class MARK-VOICE
  priority percent 50
 class MARK-VIDEO
```

```
  bandwidth percent 20
 class MARK-FTP
  bandwidth percent 10
 class class-default
  fair-queue
!
!
interface Loopback0
 ip address 172.16.1.1 255.255.255.255
!
interface GigabitEthernet0/0
 description Connected to RPE1
 ip address 192.168.10.1 255.255.255.252
 duplex auto
 speed 100
!
interface GigabitEthernet0/1
 description Connected to IXIA
 ip address 192.168.11.1 255.255.255.224
 duplex auto
 speed 100
 max-reserved-bandwidth 90
 service-policy input MARKING-TRAFFIC
 service-policy output LLQ
!
router bgp 65111
 bgp router-id 1.1.1.1
 bgp log-neighbor-changes
 neighbor 192.168.10.2 remote-as 65000
 !
 address-family ipv4
 redistribute connected
 neighbor 192.168.10.2 activate
 no auto-summary
 no synchronization
 network 172.16.1.1 mask 255.255.255.255
 network 192.168.11.0 mask 255.255.255.224
 exit-address-family
!
!
access-list 101 permit udp any any range 16384 32767
access-list 101 permit udp any any precedence critical
access-list 101 permit udp any any dscp ef
```

```
access-list 102 permit ip any any dscp cs1
access-list 102 permit ip any any dscp af13
access-list 102 permit udp any eq 970 any
access-list 102 permit udp any any eq 970
!
control-plane
!
!
alias exec i show ip int bri
!
end
############################################################
RCE2#sh run
Building configuration...

Current configuration : 2721 bytes
!
version 12.4
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
!
hostname RCE2
!
boot-start-marker
boot-end-marker
!
!
memory-size iomem 10
!
ip cef
!
no ip domain lookup
ipv6 unicast-routing
ipv6 cef
!
class-map match-all MATCH-VOICE
 match access-group name VOICE-TRAFFIC
class-map match-all MATCH-VIDEO
 match access-group name VIDEO-TRAFFIC
class-map match-all MATCH-FTP
 match access-group name FTP-TRAFFIC
!
```

```
policy-map MARKING-TRAFFIC
 class MATCH-VOICE
  set dscp ef
 class MATCH-VIDEO
  set dscp af13
 class MATCH-FTP
  set dscp af11
policy-map LLQ
 class MATCH-VOICE
  priority percent 50
 class MATCH-VIDEO
  bandwidth percent 20
 class MATCH-FTP
  bandwidth percent 10
 class class-default
  fair-queue
!
interface Loopback0
 no ip address
 ipv6 address 2002:AC:11:2::2/128
 ipv6 enable
!
interface GigabitEthernet0/0
 description Connected to RPE1
duplex auto
 speed 100
 ipv6 address 2001:C0:A8:B::1/64
 ipv6 enable
!
interface GigabitEthernet0/1
 description Connected to IXIA
 no ip address
 duplex auto
 speed 100
 ipv6 address 2003:C0:A8:C::1/64
 ipv6 enable
 max-reserved-bandwidth 90
 service-policy input MARKING-TRAFFIC
 service-policy output LLQ
!
!
router bgp 65222
 bgp router-id 2.2.2.2
```

```
 no bgp default ipv4-unicast
 bgp log-neighbor-changes
 neighbor 2001:C0:A8:B::2 remote-as 65000
 !
 address-family ipv6
 neighbor 2001:C0:A8:B::2 activate
 network 2002:AC:11:2::2/128
 network 2003:C0:A8:C::1/64
 redistribute connected
 no synchronization
 exit-address-family
!
!
ipv6 access-list VIDEO-TRAFFIC
 permit ipv6 any any dscp cs1
 permit ipv6 any any dscp af13
 permit udp any any eq 970
 permit udp any eq 970 any
!
ipv6 access-list FTP-TRAFFIC
 permit tcp any eq ftp any
 permit tcp any any eq ftp
 permit ipv6 any any dscp af11
!
ipv6 access-list VOICE-TRAFFIC
 permit udp any any range 16384 32767
 permit ipv6 any any dscp ef
!
control-plane
!
!
alias exec i show ipv6 int bri
!
end
##############################################################
```

Note: Quality of Service (QoS) implemented on RCE1 and RCE2 routers only, because IXIA generated traffic from Port - 1, 2, & 3 (one way traffic), and IXIA's Port - 4 is receiving mode that is attached RCE3 and RCE4 routers.

```
Rack-1-RPE1#sh run
Building configuration...

Current configuration : 2199 bytes
!
! Last configuration change at 00:27:30 UTC Fri Apr 15 2011
!
version 15.0
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
!
hostname Rack-1-RPE1
!
boot-start-marker
boot-end-marker
!
!
memory-size iomem 10
!
ipv6 unicast-routing
ipv6 cef
ip source-route
ip cef
!
no ip domain lookup
!
mpls label protocol ldp
!
!
license udi pid CISCO2921/K9 sn FGL150811PW
license boot module c2900 technology-package datak9
!
!
interface Loopback0
 ip address 172.21.10.10 255.255.255.255
 ipv6 address 2002:AC:15:A::A/128
 ipv6 enable
!
interface GigabitEthernet0/0
 description Connected to RP1
 ip address 192.168.10.9 255.255.255.252
 ip router isis
```

```
 duplex auto
 speed auto
 mpls ip
 isis metric 20 level-2
!
interface GigabitEthernet0/1
 description Connected to RCE1
 ip address 192.168.10.2 255.255.255.252
 duplex auto
 speed auto
!
interface GigabitEthernet0/2
 description Connected to RCE2
duplex auto
 speed auto
 ipv6 address 2001:C0:A8:B::2/64
 ipv6 enable
!
router isis
 net 49.0020.1720.2100.1010.00
 is-type level-2-only
 passive-interface Loopback0
!
router bgp 65000
 bgp router-id 5.5.5.5
 bgp log-neighbor-changes
 neighbor 2001:C0:A8:B::1 remote-as 65222
 neighbor 172.22.20.20 remote-as 65000
 neighbor 172.22.20.20 update-source Loopback0
 neighbor 192.168.10.1 remote-as 65111
 !
 address-family ipv4
  no synchronization
  network 172.21.10.10 mask 255.255.255.255
  redistribute connected
  no neighbor 2001:C0:A8:B::1 activate
  neighbor 172.22.20.20 activate
  neighbor 192.168.10.1 activate
  no auto-summary
 exit-address-family
 !
 address-family ipv6
  redistribute connected
```

```
  no synchronization
  network 2002:AC:15:A::A/128
  neighbor 2001:C0:A8:B::1 activate
  neighbor 172.22.20.20 activate
  neighbor 172.22.20.20 send-label
 exit-address-family
!
ip forward-protocol nd
!
control-plane
!
end
############################################################
Rack-2-RP1#sh run
Building configuration...

Current configuration : 1430 bytes
!
! Last configuration change at 20:09:20 UTC Mon Apr 4 2011
!
version 15.0
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
!
hostname Rack-2-RP1
!
boot-start-marker
boot-end-marker
!
ip cef
!
no ip domain lookup
!
mpls label protocol ldp
!
!
license udi pid CISCO2921/K9 sn FGL150811PS
license boot module c2900 technology-package datak9
!
!
interface Loopback0
 ip address 172.25.50.50 255.255.255.255
```

```
!
interface GigabitEthernet0/0
 description Connected to RPE1
 ip address 192.168.10.10 255.255.255.252
 ip router isis
 duplex auto
 speed auto
 mpls label protocol ldp
 mpls ip
 isis metric 20 level-2
!
interface GigabitEthernet0/1
 description Connected to RP2
 ip address 192.168.10.13 255.255.255.252
 ip router isis
 duplex auto
 speed auto
 mpls label protocol ldp
 mpls ip
 isis metric 2 level-2
!
router isis
 net 49.0020.1720.2500.5050.00
 is-type level-2-only
 passive-interface Loopback0
!
ip forward-protocol nd
!
control-plane
!
end
############################################################
Rack-4-RP2#sh run
Building configuration...

Current configuration : 1427 bytes
!
! Last configuration change at 20:15:21 UTC Mon Apr 4 2011
!
version 15.0
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
```

```
!
hostname Rack-4-RP2
!
boot-start-marker
boot-end-marker
!
no aaa new-model
!
ip cef
!
no ip domain lookup
!
mpls label protocol ldp
!
!
license udi pid CISCO2921/K9 sn FGL150811PR
license boot module c2900 technology-package datak9
!
!
interface Loopback0
 ip address 172.26.60.60 255.255.255.255
!
interface GigabitEthernet0/0
 description Connected to RP1
 ip address 192.168.10.17 255.255.255.252
 ip router isis
 duplex auto
 speed auto
 mpls label protocol ldp
 mpls ip
 isis metric 2 level-2
!
interface GigabitEthernet0/1
 description Connected to RPE2
 ip address 192.168.10.14 255.255.255.252
 ip router isis
 duplex auto
 speed auto
 mpls label protocol ldp
 mpls ip
 isis metric 20 level-2
!
interface GigabitEthernet0/2
```

```
 no ip address
 shutdown
 duplex auto
 speed auto
!
router isis
 net 49.0020.1720.2600.6060.00
 is-type level-2-only
 passive-interface Loopback0
!
ip forward-protocol nd
!
control-plane
!
end
################################################################
Rack-3-RPE2#sh run
Building configuration...

Current configuration : 2185 bytes
!
! Last configuration change at 20:16:59 UTC Mon Apr 4 2011
!
version 15.0
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
!
hostname Rack-3-RPE2
!
boot-start-marker
boot-end-marker
!
!
ipv6 unicast-routing
ipv6 cef
ip cef
!
no ip domain lookup
!
mpls label protocol ldp
!
license udi pid CISCO2921/K9 sn FGL150811PU
```

```
license boot module c2900 technology-package datak9
!
interface Loopback0
 ip address 172.22.20.20 255.255.255.255
 ipv6 address 2002:AC:16:14::14/128
 ipv6 enable
!
interface GigabitEthernet0/0
 description Connected to RP2
 ip address 192.168.10.18 255.255.255.252
 ip router isis
 duplex auto
 speed auto
 mpls ip
 isis metric 20 level-2
!
interface GigabitEthernet0/1
 description Connected to RCE3
 ip address 192.168.10.22 255.255.255.252
 duplex auto
 speed auto
!
interface GigabitEthernet0/2
 description Connected to RCE4
duplex auto
 speed auto
 ipv6 address 2001:C0:A8:D::2/64
 ipv6 enable
!
router isis
 net 49.0020.1720.2200.2020.00
 is-type level-2-only
 passive-interface Loopback0
!
router bgp 65000
 bgp router-id 6.6.6.6
 bgp log-neighbor-changes
 neighbor 2001:C0:A8:D::1 remote-as 65444
 neighbor 172.21.10.10 remote-as 65000
 neighbor 172.21.10.10 update-source Loopback0
 neighbor 192.168.10.21 remote-as 65333
 !
 address-family ipv4
```

```
  no synchronization
  network 172.22.20.20 mask 255.255.255.255
  redistribute connected
  no neighbor 2001:C0:A8:D::1 activate
  neighbor 172.21.10.10 activate
  neighbor 192.168.10.21 activate
  no auto-summary
 exit-address-family
 !
 address-family ipv6
  redistribute connected
  no synchronization
  network 2002:AC:16:14::14/128
  neighbor 2001:C0:A8:D::1 activate
  neighbor 172.21.10.10 activate
  neighbor 172.21.10.10 send-label
 exit-address-family
!
ip forward-protocol nd
!
control-plane
!
end
##############################################################
RCE3#sh run
Building configuration...

Current configuration : 2314 bytes
!
version 12.4
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
!
hostname RCE3
!
boot-start-marker
boot-end-marker
!
memory-size iomem 10
!
ip cef
!
```

```
no ip domain lookup
!
interface Loopback0
 ip address 172.18.3.3 255.255.255.255
!
interface GigabitEthernet0/0
 description Connected to RPE2
 ip address 192.168.10.21 255.255.255.252
 duplex auto
 speed 100
!
interface GigabitEthernet0/1
 description Connected to IXIA
 ip address 192.168.13.1 255.255.255.224
 duplex auto
 speed 100
!
!
router bgp 65333
 bgp router-id 3.3.3.3
 bgp log-neighbor-changes
 neighbor 192.168.10.22 remote-as 65000
 !
 address-family ipv4
 redistribute connected
 neighbor 192.168.10.22 activate
 no auto-summary
 no synchronization
 network 172.18.3.3 mask 255.255.255.255
 network 192.168.13.0 mask 255.255.255.224
 exit-address-family
!
control-plane
!
alias exec i show ip int bri
!
end
#################################################################
RCE4#sh run
Building configuration...

Current configuration : 2502 bytes
!
```

```
version 12.4
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
!
hostname RCE4
!
boot-start-marker
boot-end-marker
!
no aaa new-model
memory-size iomem 10
!
ip cef
!
no ip domain lookup
ipv6 unicast-routing
ipv6 cef
!
interface Loopback0
ipv6 address 2002:AC:13:4::4/128
 ipv6 enable
!
interface GigabitEthernet0/0
 description Connected to RPE2
duplex auto
 speed 100
 ipv6 address 2001:C0:A8:D::1/64
 ipv6 enable
!
interface GigabitEthernet0/1
 description Connected to IXIA
duplex auto
 speed 100
 ipv6 address 2003:C0:A8:E::1/64
 ipv6 enable
!
!
router bgp 65444
 bgp router-id 4.4.4.4
 no bgp default ipv4-unicast
 bgp log-neighbor-changes
 neighbor 2001:C0:A8:D::2 remote-as 65000
```

```
 !
 address-family ipv6
 neighbor 2001:C0:A8:D::2 activate
 network 2002:AC:13:4::4/128
 network 2003:C0:A8:E::1/64
 redistribute connected
 no synchronization
 exit-address-family
!
control-plane
!
alias exec i show ipv6 int bri
!
end
############################################################
R4-SW3500#sh run
Building configuration...
Current configuration:
!
version 12.0
no service pad
service timestamps debug uptime
service timestamps log uptime
!
hostname R4-SW3500
!
ip subnet-zero
!
interface FastEthernet0/1
.
.
.
interface FastEthernet0/23
 description RECEIVING MODE TO IXIA PORT-4
 port monitor FastEthernet0/11
 port monitor FastEthernet0/12
!
interface VLAN1
  no ip directed-broadcast
  no ip route-cache
!
end
############################################################
```

```
R3-SW3500#sh run
Building configuration...
Current configuration:
!
version 12.0
no service pad
service timestamps debug uptime
service timestamps log uptime
!
hostname R3-SW3500
!
ip subnet-zero
!
interface FastEthernet0/1
.
.
.
interface FastEthernet0/23
 description SENDING MODE FROM IXIA PORT-1-2-3
 port monitor FastEthernet0/11
 port monitor FastEthernet0/12
 !
interface VLAN1
  no ip directed-broadcast
  no ip route-cache
!
end
############################################################
```

Rack-3-SW3500#show port monitor

Monitor Port	Port Being Monitored
FastEthernet0/23	FastEthernet0/11
FastEthernet0/23	FastEthernet0/12

Rack-4-SW3500#show port monitor

Monitor Port	Port Being Monitored
FastEthernet0/23	FastEthernet0/11
FastEthernet0/23	FastEthernet0/12