UNIVERSITY OF ALBERTA

MINT 709 Capstone Project Report
Master of Science in Internetworking Program

Analysis of Voice Over IP deployments in Multi-site MPLS VPN Environments

Prepared by: Stefan Mocanu

Prepared for: Dr. M. H. MacGregor

September 2009

Acknowledgments

This project report is part of the requirements for the degree of Master of Science in Internetworking. The project was designed and implemented in the MINT lab at the University of Alberta during the year of 2009.

I would like to thank my project supervisor Dr. M.H. MacGregor for his guidance and support throughout the duration of the project, starting with the approval of my project definition document to live demonstration of test network in the lab.

Over the last three years, while enrolled in the M.Sc. in Internetworking Program, I had a chance to meet and learn from faculty Professors, take interesting courses and work in the lab with classmates on various technologies, as well as refine my interpersonal and leadership skills.

Finally, I would like to express my gratitude to my family for their support, encouragement and understanding during my M.Sc. studies. Our positive family attitude towards lifelong learning has now reached another milestone.

Table of Contents

Introduction

This project was undertaken as part of the requirements for the degree of Master of Science in Internetworking at University of Alberta, Department of Computer Science. The lab experiment was conducted between March - August 2009.

The project scope, Analysis of Voice Over IP deployments in Multi-site MPLS VPN Environments, included a broad range of networking concepts and technologies from MPLS, L3 MPLS VPN, hub-and-spoke VPN to OSPF and BGP while diving deep into VoIP call emulation, data transfers using IXIA IxChariot software and implementation of QoS mechanisms including traffic classification, marking, congestion avoidance and policing on Cisco devices.

The project included four major phases. **Phase I** included literature survey, network design and implementation. Network connectivity and L3 VPN connectivity in a hub-and-spoke arrangement was implemented and confirmed.

Phase II of the project implemented intra-site and inter-site Voice over IP functionality. Cisco Call Manager Express running on CE routers was used for local calls, while inter-site calls were handled via CME and a H.323 gatekeeper installed on a Cisco 3825 router located in central site. Site-facing CE router interfaces were configured as gateways which register with the gatekeeper in central site. This configuration provides better scalability and management as the gateway will simply register with the gatekeeper instead of contacting every other CE router.

Phase III of the project implemented QoS on top of existing network environment. Site level QoS was implemented using separate VLANs for VoIP and data traffic. CE routers provided inter-VLAN connectivity using virtual subinterfaces while 3750G switch ports acted as QoS domain boundary. On the provider network, QoS was implemented on PE routers in a DiffServ backbone model using QoS mechanisms like classification, marking, congestion avoidance and policing. Policies were applied in both directions on both router interfaces on all PE routers for consistency.

Phase IV of the project was dedicated to IxChariot testing, measurements, data collection and statistical analysis. Testing methodology included over 100 tests and due to project report size requirements, only a small number of tests results were included. The whole section of QoS implementation using congestion avoidance (WRED) policies for best-effort traffic was not included in the report.

The project documentation CD contains all device configuration files, test results, packet captures using WireShark and resources in electronic format. The CD root contains an index in html format that can be open in a web browser and the content can be easily accessed using hyperlinks provided.

Chapter 1
Multi Protocol Label Switching (MPLS)

1.1 MPLS Overview

Multiprotocol Label Switching has evolved from a buzzword in the networking industry to a widely deployed technology in service provider networks [13]. The success of MPLS is based on its ability to carry various types of network traffic including IP traffic, Voice over IP traffic as well as Layer 2 traffic.

MPLS is a mature and stable technology and can consolidate ATM, Frame Relay and Voice, Video and Data (so called triple play) into one unified network infrastructure which is very attractive to service providers. MPLS is capable of addressing a multitude of challenges encountered in present-day networks including speed, scalability, quality of service, traffic engineering.

MPLS also reduces router processing requirements as packets are simply forwarded based on fixed labels. MPLS provides the appropriate level of security to make IP as secure as Frame Relay in the WAN, while reducing the need for encryption on public IP networks.

MPLS combines the intelligence of IP routing with the performance of switching. Packets are forwarded based on labels which might correspond to destination IP addresses, QoS classes or source IP addresses. Labels have local significance to the routers that generated them and define label switched paths (LSPs) between endpoints.

1.2 MPLS Architecture and Operation

MPLS Header

The MPLS header is composed of a 20-bit label, three-bit experimental field, one-bit bottom of stack indicator and eight-bit TTL field (Fig. 1.2.1). MPLS assigns labels to packets moved across MPLS network and a label is assigned to a destination prefix.

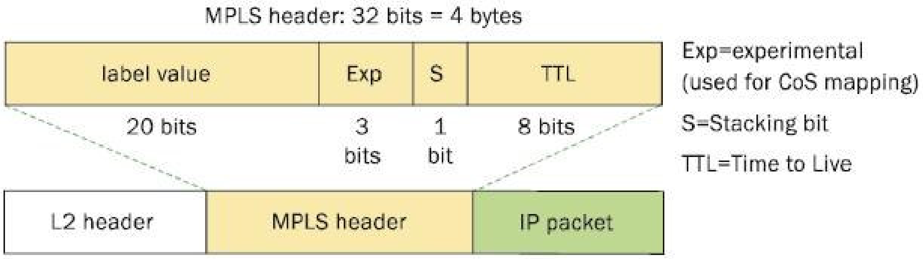

Fig. 1.2.1 MPLS header [14]

The experimental EXP field defines the CoS/QoS assigned to a Forward Equivalency Class (FEC) which is a grouping of packets that receive the same routing treatment.

A label stack (Fig. 1.2.2) is an ordered set of any number of labels and each label performs a specific function. The stacking bit or bottom of stack indicator is set to 1 to indicate the label encountered is the bottom of the label stack.

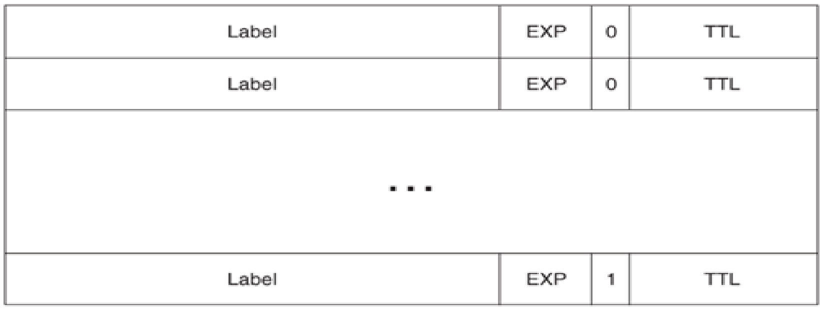

Label	EXP	0	TTL
Label	EXP	0	TTL
. . .			
Label	EXP	1	TTL

Fig. 1.2.2 Generic MPLS Label Stack [12]

Label stacks are used with MPLS applications such us MPLS VPN or MPLS Traffic Engineering (Fig. 1.2.3).

In MPLS VPN the top label (s=0) identifies the next-hop label and the bottom label (s=1) identifies the VPN label. In MPLS traffic engineering, the top label (s=0) identifies the endpoint of the TE tunnel and the bottom label (s=1) identifies the destination.

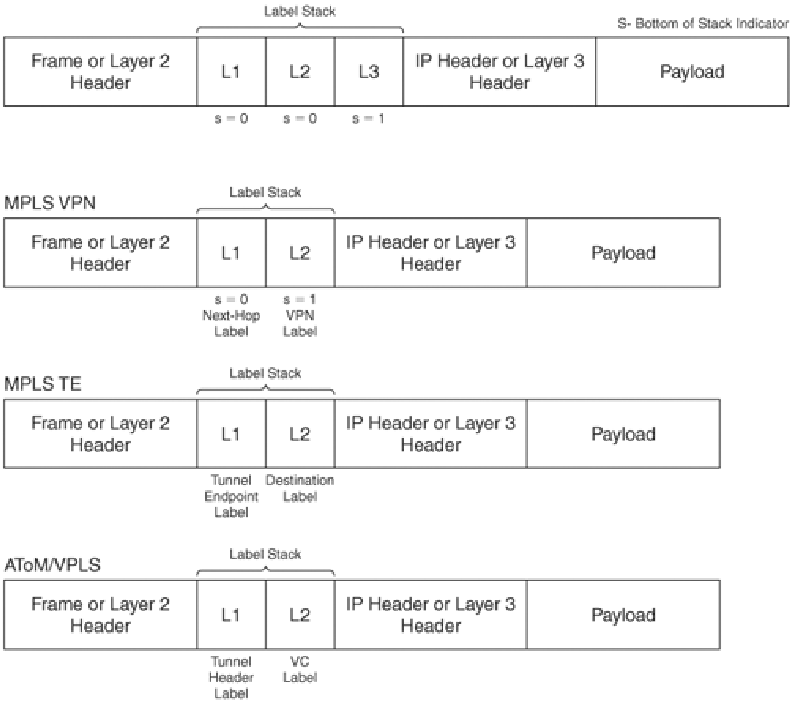

Fig. 1.2.3 MPLS Label Stack [13]

Data flow in MPLS network

The data flow in an MPLS network [14] is detailed below (Fig. 1.2.4).

1. PE routers first establish LSPs through the MPLS network to remote PE routers, before traffic is forwarded between sites.
2. Non-MPLS traffic (Frame Relay, ATM, Ethernet, etc.) is sent from a customer network, through its CE router, to the ingress PE router operating at the edge of the provider's MPLS network.
3. The PE router performs a lookup on information in the packet to associate it with a FEC, then adds the appropriate MPLS label(s) to the packet.
4. The packet proceeds along its LSP, with each intermediary P router swapping labels as specified by the information in its label information base (LIB) to direct the packet to the next hop.
5. At the egress PE, the last MPLS label is removed and the packet is forwarded by traditional routing mechanisms.
6. The packet proceeds to the destination CE and into the customer's network.

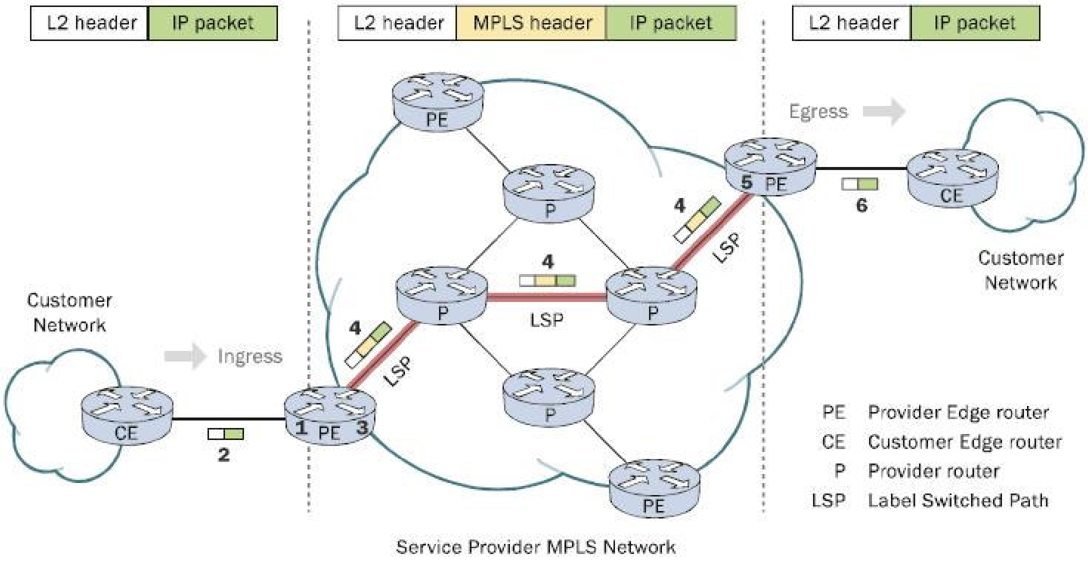

Fig. 1.2.4 MPLS Network [14]

MPLS Control and Data Plane Components

Fig. 1.2.5 presents the MPLS control and data plane components [13]. This design uses Cisco Express Forwarding (CEF) as a foundation for MPLS for improved performance and reduced overhead.

A forwarding information base (FIB) which mirrors the entire contents of the IP routing table is maintained in the data plane. A label forwarding information base (LFIB) is also maintained in the data plane and contains a local label to next-hop label mapping and the outgoing interface.

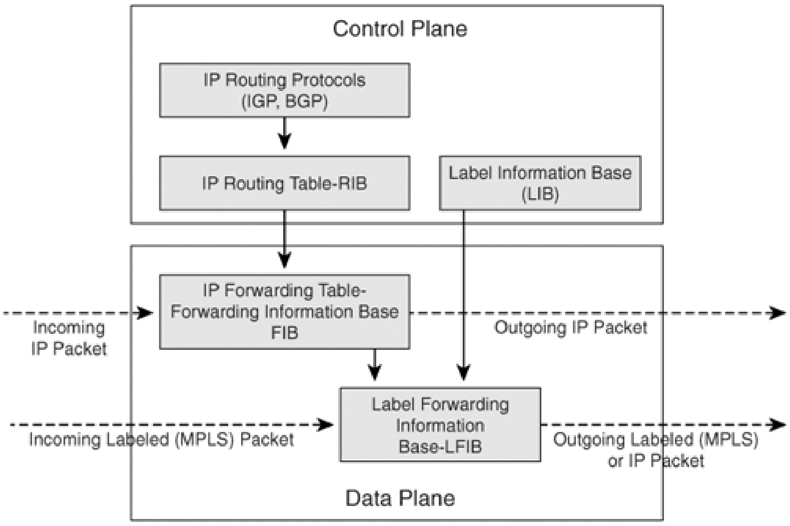
Control Plane
IP Routing Protocols (IGP, BGP)
IP Routing Table-RIB
Label Information Base (LIB)
IP Forwarding Table-Forwarding Information Base FIB
Incoming IP Packet
Outgoing IP Packet
Label Forwarding Information Base-LFIB
Incoming Labeled (MPLS) Packet
Outgoing Labeled (MPLS) or IP Packet
Data Plane

Fig. 1.2.5 MPLS Control and Data Plane Components [13]

MPLS Label Operation

A label switching router (LSR) knows how to forward a packet by looking at the top label of the received packet and corresponding entry in LFIB [12]. The router will perform one of possible three label operations: swap, push and pop (Fig. 1.2.6).

In a swap operation, the LSR will replace the top label in the label stack with another label. A push operation means that the top label is replaced with another and one or more additional labels are pushed onto the label stack. The pop operation means that the top label is removed.

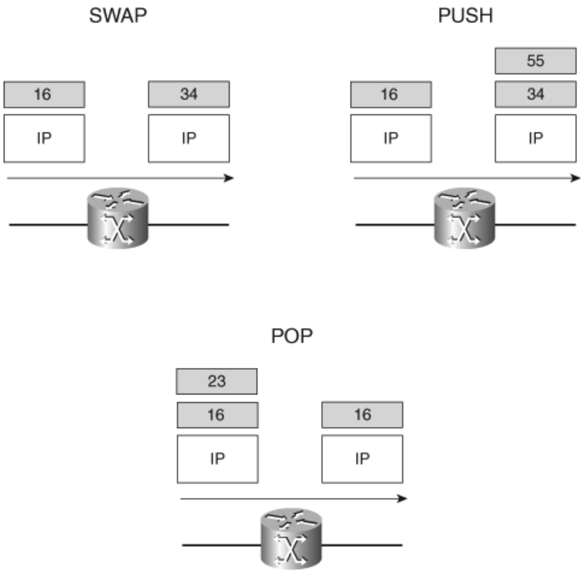

Fig. 1.2.6 MPLS Operations on Labels [12]

Tracerouting in MPLS Networks

In an IP network, a **traceroute** command on a Cisco router will send probes to the specified destination IP address. The probes are UDP packets with destination port greater than 30,000. The first probe has TTL set to 1; second probe has TTL set to 2 etc.

The TTL of the first probe will expire on the next-hop router which will reply with an ICMP "TTL exceeded" message. The process is repeated with next routers along the path which will send similar replies to the originator of the packet. When the destination device is reached (router or host) it will reply with an ICMP "port unreachable" message as it is unlikely that an application will use such a port number and generate ICMP messages.

Tracerouting in an MPLS network is similar to above except that the "TTL exceeded" message format has been extended to hold the full MPLS stack of the original packet received by the LSR replying to the ICMP request from originator. This way the label stack from each LSR is also printed in the output of the traceroute which helps in troubleshooting.

TTL behavior in MPLS networks

The following rules apply when TTL propagates from IP header to the label stack and vice versa [12]:

- When an IP packet is first labeled, TTL is copied from IP header to the TTL fields of all labels in the imposed label stack after being decremented by 1
- Label swap operation - TTL of incoming top label is copied to TTL of outgoing label after being decrementing by 1
- Label push operation - TTL field of all pushed labels is the TTL of incoming top label after being decremented by 1
- Label pop operation - TTL of incoming top label is copied to TTL of newly exposed top label after decrementing by 1. The exception is if TTL of exposed label < TTL of incoming label in which case outgoing TTL is that of exposed label after being decremented by 1
- Label stack remove operation - TTL of top label is copied to TTL of the IP header. The exception is if TTL of underlying IP packet > TTL of incoming label in which case IP TTL value remains unchanged when the packet leaves the MPLS domain.
- LSRs never change TTL values of non-exposed labels (i.e. not the top label)
- If a packet is label switched, the TTL value in the underlying IP header is never changed.

If the TTL of the top label in the label stack expires when it reaches 0, the label stack is stripped off and ICMP "TTL exceeded" message along with the label stack is returned to source IP address taken from the IP packet. One important thing to note is that the ICMP message is actually forwarded along the LSP until it reaches the end of the LSP.

In case of MPLS VPN the P core routers do not have knowledge of the underlying VPN. The only routers that can return the ICMP message are egress PE or CE routers.

Cisco IOS no mpls ip propagate−ttl command

The default behavior when an IP packet enters the MPLS domain is that TTL is copied from IP header to the TTL fields of all labels in the imposed label stack after being decremented by 1. That means a traceroute between two remote sites will reveal the topology and routers on the service provider network along the path. By using **no mpls ip propagate−ttl** command on PE routers, the IP TTL is not copied and TTL fields in labels are set to 255. The result is that the core P routers will be skipped and only ingress and egress PE routers will be visible in traceroute i.e. the MPLS domain is seen as one hop.

The above command will also make the provider topology hidden from inside which is not desired. The **no mpls ip propagate−ttl forwarded** command can be used so provider network topology is visible from inside but not from outside i.e. from the customer sites.

1.3 Cisco Express Forwarding

A packet can be forwarded through a Cisco router using process switching, interrupt switching or through an application-specific integrated circuit (ASIC) [12].

Process switching runs in router memory and it is the slowest of all methods. Interrupt switching is the opposite of process switching and runs in interrupt mode. The central CPU might be involved but the switching decision is performed within the interrupt context, not by a dedicated Cisco IOS process.

Fast switching is used as a method to speed up the router performance by using an on-demand forwarding table. The first packet for a certain destination is process switched then following packets for same destination are interrupt switched. This way a route cache is build based upon the first packet switching.

Cisco Express Forwarding (CEF) provides a prebuilt forwarding table in the hardware that is derived from the IP routing table on the router. This table is used to forward IP packets on the central processor or on the line cards when using a distributed architecture.

CEF is required on the MPLS routers as it is the only switching method that can impose labels on the IP packet.

1.4 MPLS Applications

Over the last few years MPLS has evolved into a mature technology which serves as a foundation to other applications using label-based forwarding.

MPLS VPN

Initially MPLS VPN was used as a replacement for internal Frame Relay and ATM on service provider networks. There is a growing interest for connecting provider MPLS VPN networks using inter-autonomous MPLS VPN and Carrier's Carrier implemented in large enterprise companies where it provides scalability and flexibility required to divide the network into separate smaller networks.

MPLS Traffic Engineering

Traffic Engineering is traditionally used on Frame Relay and ATM provider networks to move traffic in the most optimal way. The objective is to steer traffic in order to avoid overloaded links, load balancing using link costs has proven to be challenging.

By using label switching instead of IP routing, MPLS TE allows for source-based routing instead of IP destination-based routing. MPLS TE provides efficient movement of traffic, use of configured bandwidth of links and their attributes like delay, jitter while automatically adapts to changes in bandwidth and link attributes.

IPv6 over MPLS

Due to rapid expansion of MPLS VPN on the provider networks, there may be a need to carry some customer IPv6 traffic over the MPLS backbone.

One solution is to run dual stack IPv4 / IPv6 on the provider network routers which also require LDP support for IPv6.

Another solution is to use Any Transport over MPLS (AToM) where the MPLS payload is Layer 2 frames. In this case, the core routers do not need to run IPv6 as they only switch labeled packets.

A third solution used to carry IPv6 traffic is to use MPLS VPN. As the traffic on the existing MPLS provider network and inside VPN is IPv4, the CE routers will need to tunnel IPv6 traffic between them which adds additional overhead. Some of the tunneling methods available in Cisco IOS include [12]: IPv6 over IPv4 GRE tunnels, manual IPv6 tunnels, 6to4 tunnels, IPv4-compatible IPv6 tunnels and ISATAP tunnels.

Any Transport over MPLS

Any Transport over MPLS (AToM) was developed after the introduction of MPLS VPN and it is considered a Layer 2 VPN. AToM uses existing service provider infrastructure including leased lines, ATM, Frame Relay and only provides a point-to-point Layer 2 service. This allows AToM to carry other protocols in addition to IP protocol.

Migration to AToM is transparent for customers as they continue to use same Layer 2 encapsulation type as before and also do not need to run an IP routing protocol to the PE routers as in the MPLS VPN implementation. AToM only runs on PE routers which impose and dispose labels on the Layer 2 frames while the core routers only run LDP and switch labeled packets.

Virtual Private LAN Service

Virtual Private LAN Service (VPLS) provides a point-to-multipoint Layer 2 service across the MPLS provider backbone using pseudo-wires or virtual circuits. The customer sees the individual site LAN segments as connected together via a virtual Ethernet switch which actually is the MPLS backbone.

VPLS is used for geographically dispersed customer sites and has all characteristics of an Ethernet switch.

MPLS and Quality of Service

QoS can be implemented using IntServ or DiffServ models. IntServ uses the signaling protocol Resource Reservation Protocol (RSVP) and the hosts will tell the routers via RSVP what QoS needs are for particular traffic flows. As a result there were scalability and complexity issues which prevented IntServ from becoming a popular choice.

DiffServ does not need a signaling protocol and uses the DSCP bits in ToS field of the IP packet and MPLS EXP field to classify, mark, queue, police and shape the traffic. A typical application of QoS is triple-play networks carrying voice, video and data over IP.

Chapter 2
Layer 3 MPLS VPN networks

2.1 VPN Overview

Traditionally two major VPN models were implemented by service providers, depending on their participation in customer routing: the Overlay model and Peer-to-Peer model.

In the Overlay model, the service provider network consists of point-to-point links or virtual circuits established between customer routers. The service provider routers never see the customer routes while the customer routers act as routing peers and run routing protocols directly between them. The point-to-point links can be established at Layer 1 (TDM, E1, E3, SONET, SDH), Layer 2 (X.25, ATM, Frame Relay) or even Layer 3 (VPN services over IP using L2TP, GRE, IPSec).

In the Peer-to-Peer model, the service provider routers carry customer data but they also participate in the customer routing. The service provider routers now peer directly and form adjacencies with customer routers at Layer 3.

Security and separation of customer-specific information is achieved by implementing packet filters on the provider edge routers and IP addressing is handled by the service provider (shared peer-to-peer model). Another method is by using controlled route distribution (dedicated PE peer-to-peer model) between the core and provider edge routers via BGP with extended communities where only specific customer routes are propagated.

2.2 L3 MPLS VPN Architecture and Operation

MPLS VPN was introduced to meet the need for a scalable, efficient and cost effective peer-to-peer VPN model and it is the most popular and widespread implementation of MPLS technology [12]. It can be implemented at Layer 2 using various technologies (one of them being tunneling L2TPv3) or Layer 3 using a Peer-to-Peer model with associated routing protocols.

L3 MPLS VPN architecture includes network components and specific building blocks on PE routers which allow for VPN functionality.

L3 MPLS VPN Network Components

Fig. 2.2.1 presents a typical Layer 3 MPLS VPN composed of the following network components which perform different functions within the overall architecture framework:

- **Service Provider MPLS network** - The core MPLS/IP network administered by the service provider.
- **Provider router (P)** - MPLS/IP router deployed within the provider network with no edge service attachments.

- **Provider edge router (PE)** - Service provider edge router that provides VPN end-customer attachment and service delivery.
- **Customer network (VPN A, VPN B)** - Customer network administered by the end user attached to the Layer 3 MPLS VPN service.
- **Customer edge router (CE)** - Customer router that provides a gateway between the customer network and the Service Provider MPLS network. The CE router may be administered by the end user (and thus belong to the customer network) or may be managed by the service provider.

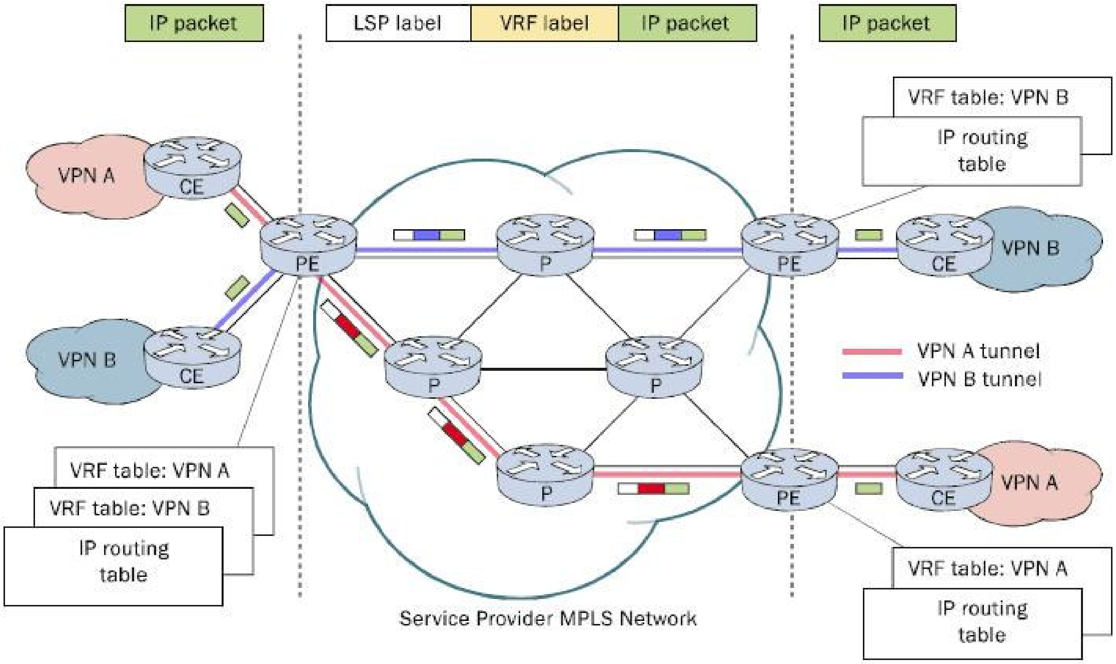

Fig. 2.2.1 Layer 3 VPN MPLS network [14]

Building blocks on PE routers [12]

- ***Virtual Routing Forwarding (VRF).*** It is a VPN routing and forwarding instance created and maintained on PE router for each attached VPN. A private VRF routing table exists for each defined VPN and is associated with a PE interface.

- ***Route Distinguisher (RD).*** RD is 64-bit field and was created to make IPv4 prefixes unique as customers might have overlapping IP addressing (Fig. 2.2.2).

  VPN prefixes are carried across provider network by Multiprotocol-BGP. An IP prefix of 172.16.2.0/24 with an RD of 1:2 (ASN:nn) will result in a VPNv4 prefix of 1:2:172.16.2.0/24. A VPN site may be connected to two PE routers resulting in two different RDs and VPNv4 routes.

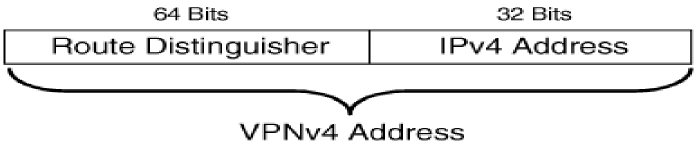

Fig. 2.2.2 VPNv4 Address Format [6]

- ***Route Target (RT).*** An RT is a BGP extended community that indicates which routes should be imported from MP-BGP into which VRFs [12]. This allows for creation of extranets where a site of company A will be able to talk to a site of company B. More than one RT might be attached to the VPNv4 route.

- ***VPNv4 Route Propagation through MP-BGP.*** A service provider network might have a large number of customers resulting in hundreds of thousands of routes. BGP is the ideal candidate as it is a proven, scalable and stable routing protocol. MP-BGP speakers are peers and exchange VPNv4 information on the MPLS VPN network (Fig. 2.2.3).

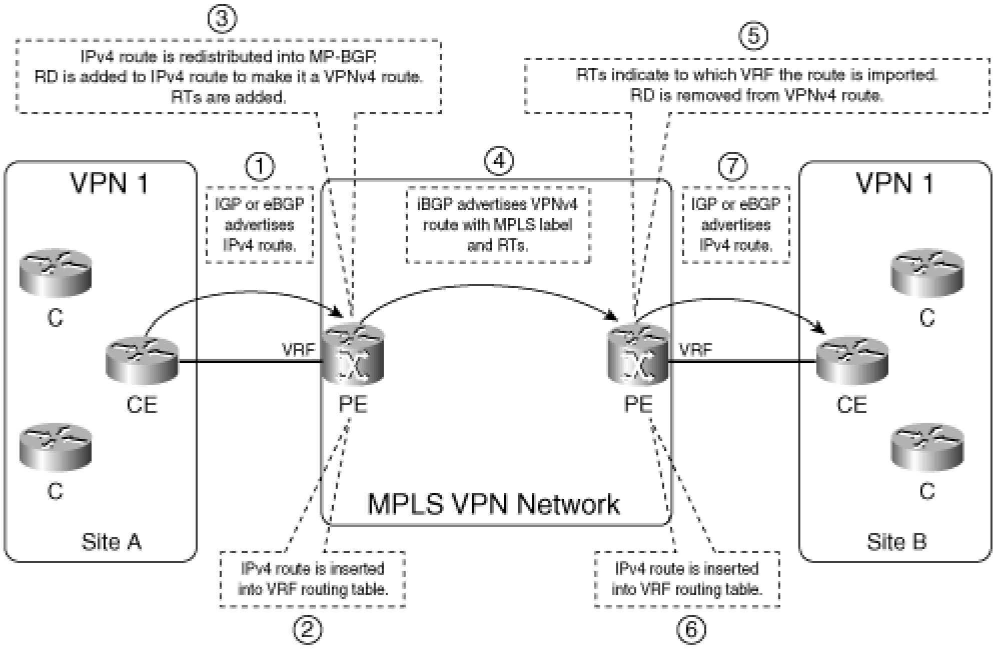
③
IPv4 route is redistributed into MP-BGP.
RD is added to IPv4 route to make it a VPNv4 route.
RTs are added.
⑤
RTs indicate to which VRF the route is imported.
RD is removed from VPNv4 route.
VPN 1
①
IGP or eBGP advertises IPv4 route.
④
iBGP advertises VPNv4 route with MPLS label and RTs.
⑦
IGP or eBGP advertises IPv4 route.
VPN 1
C
C
CE
VRF
PE
PE
VRF
CE
C
C
MPLS VPN Network
Site A
Site B
IPv4 route is inserted into VRF routing table.
IPv4 route is inserted into VRF routing table.
②
⑥

Fig. 2.2.3 Route Propagation in an MPLS VPN Network [12]

- ***Packet forwarding in MPLS VPN network.*** Label Distribution Protocol (LDP) is commonly used between P and PE routers and as a result all IP traffic is label-switched on the provider network. RSVP-TE can also be used when implementing RSVP with extensions for traffic engineering.

The label stack is composed of two labels. The top label is distributed by LDP and is called IGP label (or Next-Hop label in Fig. 1.2.3). It is used solely to forward the packet through the service provider network with no need for the P router to perform a lookup of the destination address. The P routers will only swap the top label. The bottom label is called VPN label and is advertised by MP-iBGP between PE routers. The VPN label is inserted by the ingress PE router to indicate egress PE router which VRF the packet belong to. The label stack is stripped off at the egress PE router and the packet is forwarded as an IPv4 packet onto the VRF interface toward the CE router.

As a result, provider P routers use IGP label to forward the packet to the correct egress PE router. The egress PE router uses the VPN label to forward the IP packet to the correct CE router.

The PE-CE routing protocols currently supported by Cisco IOS are static routing, RIPv2, OSPF, EIGRP, IS-IS and eBGP [12].

2.3 Hub-and-Spoke Model

In a hub-and-spoke model all inter-site traffic is routed through the customer CE router located in central (hub) site. One interesting feature of this topology is that all WAN traffic can be captured, inspected and stored for forensic analysis or to meet compliance requirements.

There is a hub router and associated spoke routers on both the provider network and customer network. In order to provide a return path for the inter-site traffic, the PE router and corresponding CE router in central site each have two separate WAN network interfaces connected between them. This can also be implemented using two separate WAN network interfaces on the PE router and two separate CE routers located in central site.

Chapter 3
QoS and VoIP in Layer 3 MPLS VPN environments

There are three types of service on IP networks: Best-Effort Service, Integrated Services (IntServ) and Differentiated Services (DiffServ).

Best-Effort Service is the default behavior of a network device with no QoS implemented.

IntServ and DiffServ are QoS models that can help service providers meet and exceed service level agreements they have with their customers for different types of traffic carried over a common provider network infrastructure.

3.1 QoS Overview

Originally, IP was specified a best-effort protocol [15]. The network delivers traffic to destination in the shortest time possible but with no guarantee of achieving it. All traffic is TCP based and in the early Internet, congestion management and Quality of Service were not important.

Over time, with the expansion of Internet, best effort service proved to be insufficient as new generations of applications emerged. Some real-time applications like voice and video have limited tolerance to variations in jitter, loss, latency and bandwidth.

IETF initially defined IntServ [RFC1633] then DiffServ [RFC2475] as two architecture models for Quality of Service. MPLS later incorporated support for the DiffServ QoS model [RFC3270].

3.2 IP QoS Models

Best-effort service

It is the standard form of connectivity without any quality of service guarantees. Devices like switches typically use first-in, first-out (FIFO) queues. The packets are simply transmitted as they arrive in a queue with no preferential treatment.

Integrated Services

IntServ provides support for real-time applications. One of the building blocks is the requirement for resource reservation and a signaling protocol (RSVP) is used to set up and refresh the session. QoS is selected at the application level and not the network.

Differentiated Services

DiffServ excluded signaling and microflow mechanisms. Instead a simple and coarse QoS approach was used which applied to both IPv4 and IPv6.

DiffServ defines classes of traffic with different service requirements. The IP ToS byte has been re-defined into a DiffServ field where the three-bit IP Precedence field (table 3.2.1) was extended to a 6-bit field called DSCP or Differentiated Services Code Point (fig. 3.2.1).

IP Precedence Value	Binary Value	Priority
0	000	Routine
1	001	Priority
2	010	Immediate
3	011	Flash
4	100	Flash Override
5	101	Critical
6	110	Internetwork Control
7	111	Network Control

Table 3.2.1 IP Precedence Values [13]

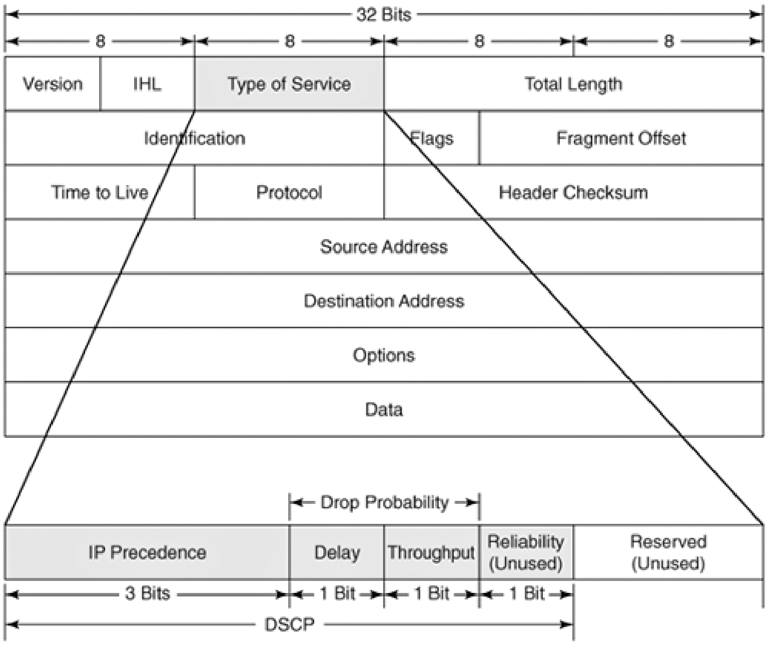

Figure 3.2.1 IP Packet Header [13]

There are three code point pools defined in the DiffServ model with 64 possible values. The first two pools (32 of 64 possible values) are reserved for experimental or local use while the third pool classes are used for QoS purpose (Table 3.2.2).

One important DiffServ class to note is the EF class (Expedited Forwarding) with a value of 101110 (101 IP Precedence) which is used for Voice over IP QoS. This is seen by end devices as a point-to-point connection which minimizes jitter, loss, latency and provides assured-bandwidth, guaranteed end-to-end service through DiffServ domains. The remaining classes are Assured Forwarding classes.

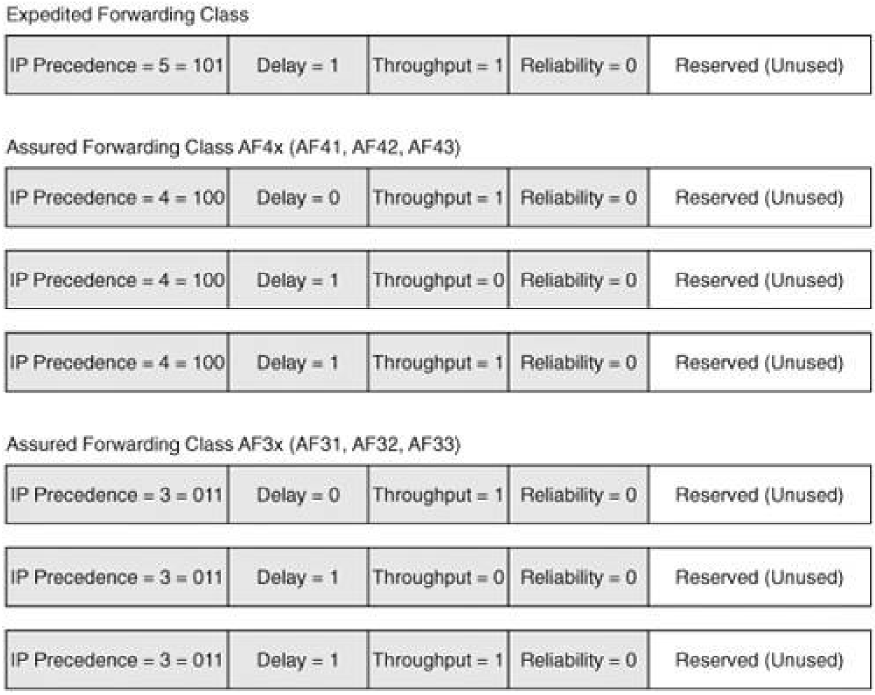

Expedited Forwarding Class

IP Precedence = 5 = 101	Delay = 1	Throughput = 1	Reliability = 0	Reserved (Unused)

Assured Forwarding Class AF4x (AF41, AF42, AF43)

IP Precedence = 4 = 100	Delay = 0	Throughput = 1	Reliability = 0	Reserved (Unused)
IP Precedence = 4 = 100	Delay = 1	Throughput = 0	Reliability = 0	Reserved (Unused)
IP Precedence = 4 = 100	Delay = 1	Throughput = 1	Reliability = 0	Reserved (Unused)

Assured Forwarding Class AF3x (AF31, AF32, AF33)

IP Precedence = 3 = 011	Delay = 0	Throughput = 1	Reliability = 0	Reserved (Unused)
IP Precedence = 3 = 011	Delay = 1	Throughput = 0	Reliability = 0	Reserved (Unused)
IP Precedence = 3 = 011	Delay = 1	Throughput = 1	Reliability = 0	Reserved (Unused)

Table 3.2.2 DSCP Classes [13]

3.3 QoS Mechanisms

QoS mechanisms at the router level (per hop behavior or PHB) include (Fig. 3.3.1):

- Traffic classification
- Traffic marking
- Congestion management (queuing)
- Congestion avoidance (WRED or weighted random early detection)
- Traffic policing (enforce bandwidth control by dropping packets)
- Traffic shaping (packet buffering/delaying according to specific traffic profile)

Traffic classification, marking and policing are typically applied on ingress while shaping is applied on egress.

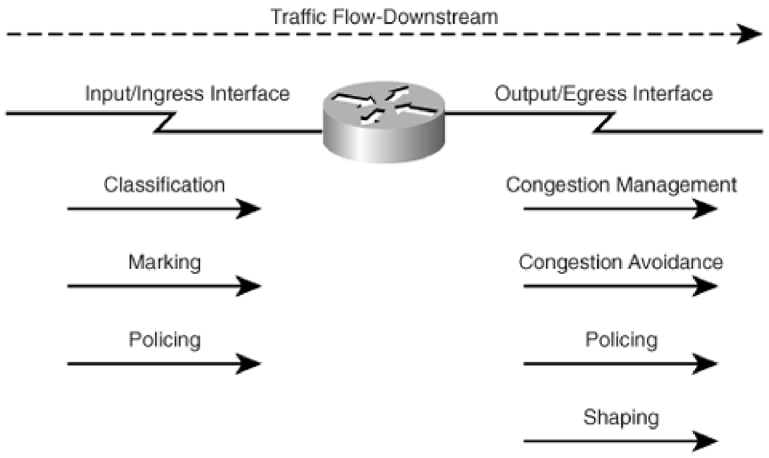

Figure 3.3.1 QoS Mechanisms [13]

3.4 MPLS QoS

MPLS support for IntServ remains undefined while MPLS supports DiffServ with minimal adjustments to MPLS and DiffServ architectures.

There are several design alternatives that can be used on a MPLS backbone [15]:

Best-effort backbone (Fig. 3.4.1)

This is the simplest approach to MPLS QoS. There is no traffic differentiation, capacity planning ensures proper utilization level. Typically the provisioning of link capacity should be more than twice the average load of traffic which will result in maximum average utilization of 50%.

Best-effort backbone with MPLS traffic engineering (Fig. 3.4.2)

Traffic Engineering provides admission control on the link which allows for traffic load management. Constraint-based routing allows traffic to avoid backbone congestion by routing traffic over underutilized links. This scenario does not offer any traffic differentiation and still relies on best-effort backbone.

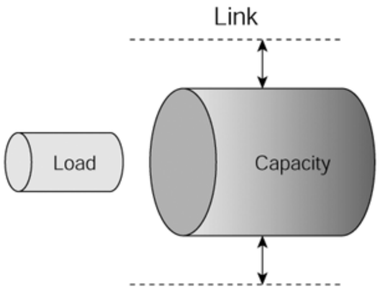

Fig. 3.4.1 [15]

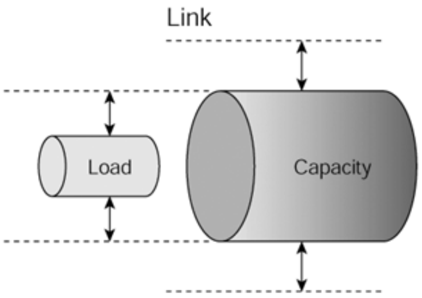

Fig. 3.4.2 [15]

DiffServ backbone (Fig. 3.4.3)

DiffServ introduces traffic differentiation. There are multiple classes and capacity planning is performed at the class level which allows for resource optimization. Deployment of DiffServ on the backbone creates multiple virtual networks with the ability to perform capacity planning independently.

MPLS field is used with typical values of five for expedited forwarding (EF) and zero for default or best-effort traffic. Values six and seven are historically used by network control traffic while remaining one through four classes are available for Assured Forwarding (AF).

DiffServ backbone with MPLS Traffic Engineering (Fig. 3.4.4)

In this scenario MPLS TE provides control over link utilization which results in better bandwidth utilization while DiffServ provides traffic differentiation.

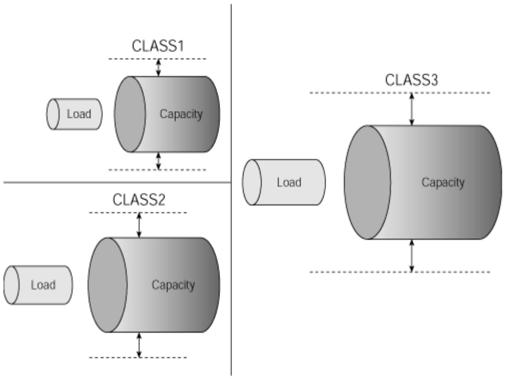

Fig. 3.4.3 [15]

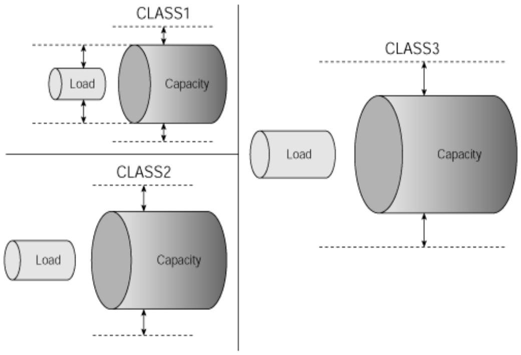

Fig. 3.4.4 [15]

DiffServ backbone with DiffServ-aware Traffic Engineering (Fig. 3.4.5)

This scenario provides the highest level of traffic differentiation and optimization. DiffServ controls class capacity while DiffServ-Traffic Engineering controls the class load.

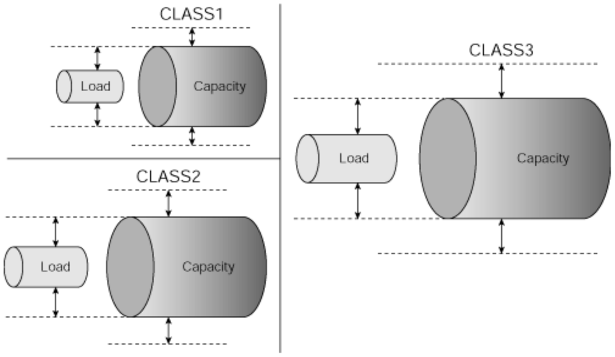

Fig. 3.4.5 [15]

DiffServ Tunneling Models over MPLS

When implementing MPLS QoS, the edge LSR router between the IP and MPLS domains performs the translation between the two QoS domains in both directions.

MPLS LSPs support for DiffServ defines three models of interaction between DiffServ markings in different layers of encapsulation [15]. For an IP packet, there is one PHB (per hop behavior) marking in the MPLS encapsulation and a PHB marking in the DiffServ field of the IP packet.

The following models are available: pipe, short-pipe and uniform. These models define the procedures that an LSR can apply to a packet (IP or MPLS) when the packet enters or exits an LSP.

The **pipe model** is used when a MPLS network connects other DiffServ domains such us when a service provider implements DiffServ QoS policies independent of the customer's QoS policies. The IP precedence of the IP packet is not copied into MPLS EXP field on ingress and remains unchanged while traversing the MPLS network.

The **short-pipe model** is a variation of the pipe model. The only difference applies at the egress router between the MPLS and IP domain. The packet's per-hop-behavior is associated with the IP Precedence/DSCP value of the underlying IP packet and not to ingress labeled EXP value.

The **uniform model** is typically applied in a managed CE scenario and acts as an extension to the DiffServ domain of the encapsulated packet. This may be required when DiffServ domains are connected via an MPLS domain and all networks need to behave as a single DiffServ domain. The IP packet IP Precedence/DSCP value is copied onto imposed label EXP value on the ingress. On egress, the topmost label EXP value is copied onto IP Precedence/DSCP of the outgoing packet.

3.5 IXIA IxChariot and VoIP

IxChariot from IXIA is a traffic generation and decision support tool used to emulate real-world application data. IxChariot provides thorough performance assessment and device testing by emulating a large number of protocols and applications using proprietary network performance endpoints. These endpoints can be host machines running various operating system platforms or run on custom built line cards installed in IXIA chassis.

The default values in an IxChariot VoIP test emulate a unidirectional voice stream. When emulating a full-duplex bi-directional voice stream like G.711, two pairs using the same codec for each voice channel are required.

IxChariot Test Process

The process that IxChariot uses to run a test is detailed in [8] and presented below and in Fig. 3.5.1. As previously mentioned, Endpoint 1 and Endpoint 2 can be two separate host computers or ports on a line card installed in the IXIA tester chassis.

- Create a test on the IxChariot Console
- Initiate execution of the test. At this point, the Console establishes communications with, and sends the test setup information to, the Endpoint 1 computer.
- The Endpoint 1 computer establishes communications with, and sends test setup information to, the Endpoint 2 computer. When Endpoint 2 has acknowledged it is ready, Endpoint 1 replies to the console. When all endpoint pairs are ready (in Figure 3.5.1 there is only one pair), the Console directs them all to start.
- The two endpoint computers execute the test.
- The Endpoint 1 computer collects the test results (timing records) and sends them to the IxChariot Console.

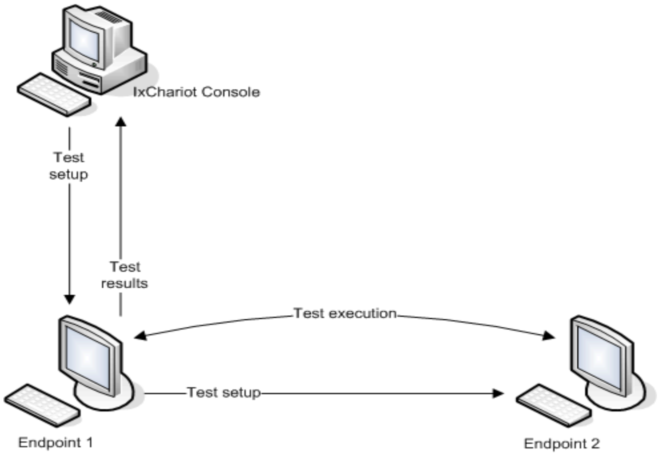

Fig. 3.5.1 - IxChariot Test Process [8]

The test definition is saved with a .tst extension and it includes all components required to either manually launch the test or schedule it to run at required intervals.

There are several types of endpoint pairs including regular (data traffic), VoIP (using one of the available codecs like G.729 and G.711u), video, multicast and IPTV.

IxChariot uses one of two methods to generate network traffic: script-generated traffic and stream-generated traffic.

IxChariot **application scripts** (Fig. 3.5.2) generate network traffic that emulates traffic patterns typical of a particular type of application and may be of two different types: streaming scripts and non-streaming scripts (for example two-way communication in a form of a database query).

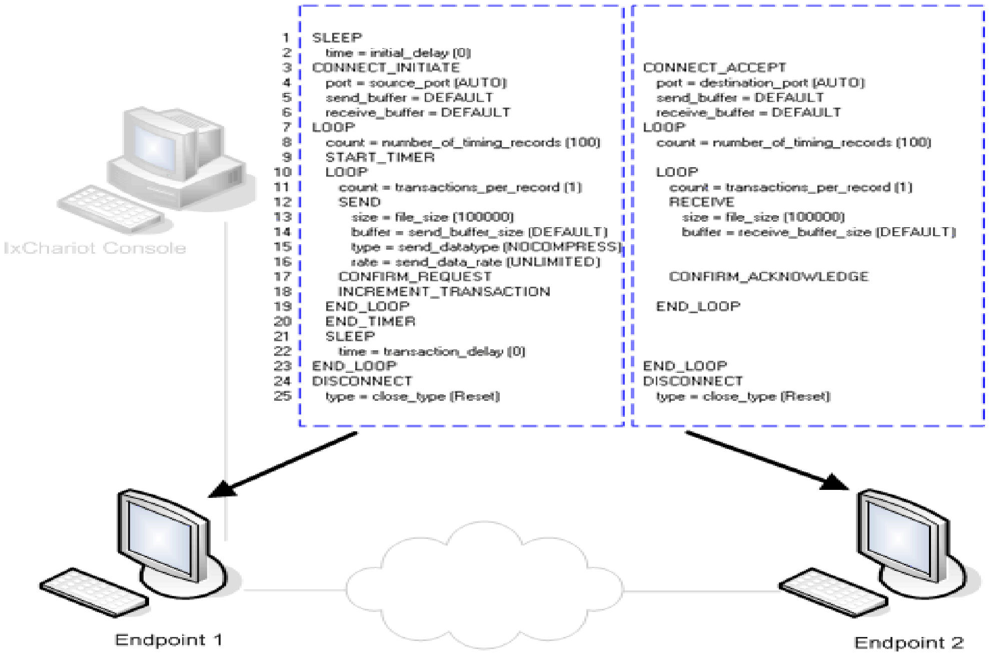

```
SLEEP
  time = initial_delay (0)
CONNECT_INITIATE
  port = source_port (AUTO)
  send_buffer = DEFAULT
  receive_buffer = DEFAULT
LOOP
  count = number_of_timing_records (100)
  START_TIMER
  LOOP
    count = transactions_per_record (1)
    SEND
      size = file_size (100000)
      buffer = send_buffer_size (DEFAULT)
      type = send_datatype (NOCOMPRESS)
      rate = send_data_rate (UNLIMITED)
    CONFIRM_REQUEST
    INCREMENT_TRANSACTION
  END_LOOP
  END_TIMER
  SLEEP
    time = transaction_delay (0)
END_LOOP
DISCONNECT
  type = close_type (Reset)
```

```
CONNECT_ACCEPT
  port = destination_port (AUTO)
  send_buffer = DEFAULT
  receive_buffer = DEFAULT
LOOP
  count = number_of_timing_records (100)

  LOOP
    count = transactions_per_record (1)
    RECEIVE
      size = file_size (100000)
      buffer = receive_buffer_size (DEFAULT)


    CONFIRM_ACKNOWLEDGE

  END_LOOP



END_LOOP
DISCONNECT
  type = close_type (Reset)
```

Fig. 3.5.2 - Sample IxChariot Script [8]

Ixia **streams** are used only when Ixia chassis ports serve as endpoints in a test. Layer 2 and layer 3 network test traffic is generated entirely in hardware, allowing tests to saturate network interfaces at up to wire speed.

IxChariot VoIP QoS Testing

IxChariot provides QoS templates that can be customized and saved for future use. The DiffServ template (Fig. 3.5.3) can be used to change settings in the ToS byte at the bit level therefore emulating various QoS classes. VoIP packets typically have a DSCP value of 101110 which represents the EF (expedited forwarding) class used in DiffServ QoS domains.

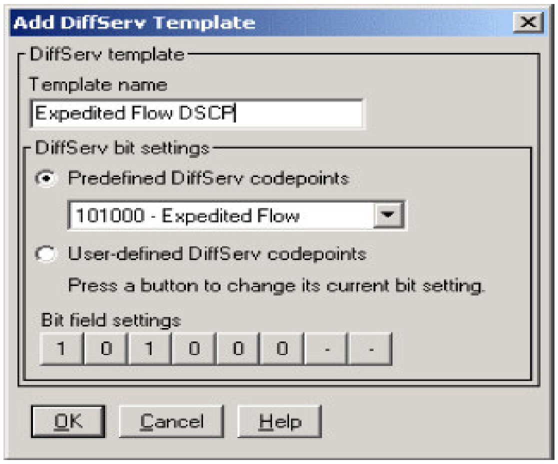

DiffServ template

Template name: Expedited Flow DSCP

DiffServ bit settings

● Predefined DiffServ codepoints: 101000 - Expedited Flow

○ User-defined DiffServ codepoints

Press a button to change its current bit setting.

Bit field settings: 1 0 1 0 0 0 - -

Fig. 3.5.3 - IxChariot DiffServ Template [8]

In the IxChariot application, only payload data is included in the calculations. As an example, 8 Kbps G729 codec (Cisco default for VoIP over WAN links) was used for experimentation in the MINT lab while throughput is typically reported by Cisco interface as up to 32Kbps (including overhead). The Cisco default VoIP codec used on high-speed networks (LAN) is G.711u with a payload bandwidth of 64 Kbps.

IXIA IxChariot application gives an indication of the relative quality of each call made during a test by treating each endpoint pair as a separate VoIP call.

IXIA uses a modified version of the ITU G.107 standard E-Model equation to calculate a MOS (mean opinion score) estimate for each endpoint pair (Table 3.5.1). A MOS score ranges from 1 for an unacceptable call to 5 for an excellent call. A typical range for Voice over IP would be from 3.5 to 4.2.

Mean Opinion Score (lower limit)	User Satisfaction
4.34	Very satisfied
4.03	Satisfied
3.60	Some users dissatisfied
3.10	Many users dissatisfied
2.58	Nearly all users dissatisfied

Table 3.5.1 - Mean Opinion Score [8]

The following factors are used to calculate the MOS estimate: one-way (network) delay, end-to-end delay, packetization delay, jitter buffer delay, additional fixed delay, data loss, jitter buffer lost datagrams.

Other VoIP call quality parameters include delay variation (jitter) as well as throughput and packet loss during data transfers. Typically jitter values in excess of 50 ms probably indicate poor VoIP call quality [8].

IxChariot automatically measures all VoIP and data transfer parameters and displays test results in graphic format on the IxChariot console. The test results can also be exported in html, csv and text format and include extensive information that may be used for further processing and analysis.

Chapter 4
Network Design and Implementation

A test network was set up in the MINT lab after the network design phase was completed (Fig. 4.1.1). The network is composed of the service provider MPLS backbone and two customer remote sites (Site 2, Site 3) are connected to their central site (Site 1) using the service provider network.

4.1 Network Design

There is a Layer 3 VPN configured in a hub-and-spoke topology on top of the provider MPLS network. All inter-site traffic is routed through customer CE-1 router. In order to provide a return path for the inter-site traffic, PE-1 and CE-1 routers have two separate network interfaces connected between them.

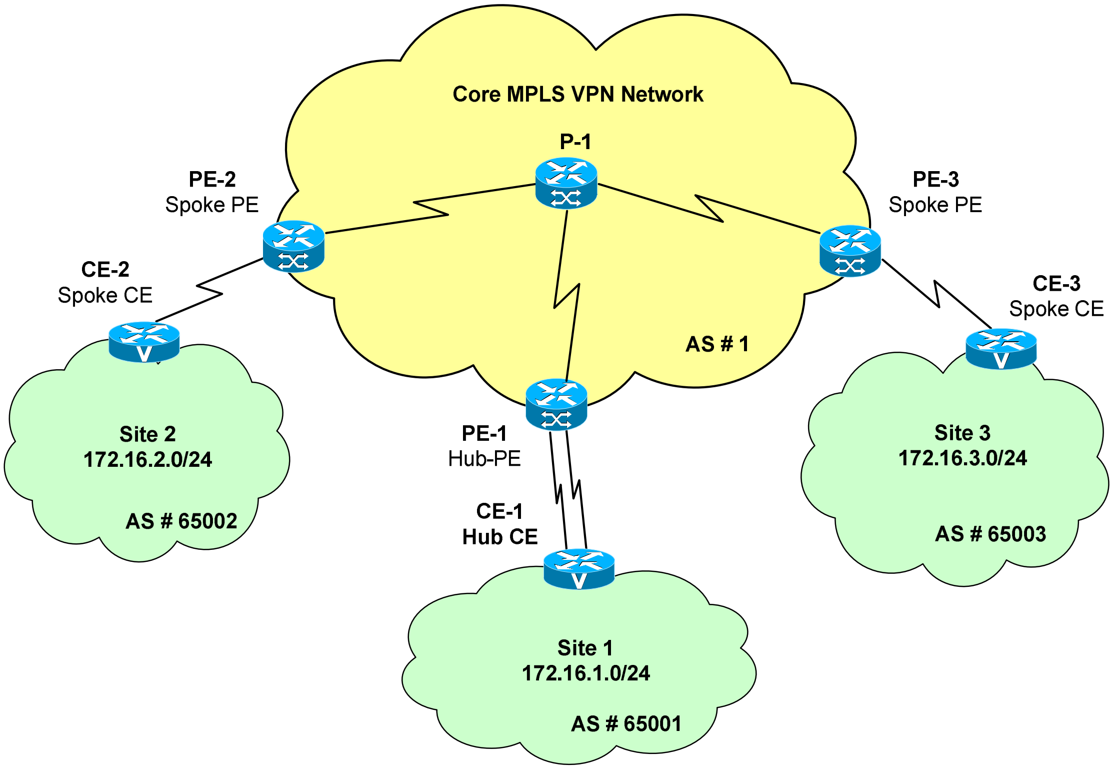

Fig. 4.1.1 Test Network

The IP addressing scheme is detailed in Appendix A - Network Diagram. On the provider network, all routers and interfaces are configured with the Cisco tag switching protocol (predecessor of MPLS) available on routers in the MINT lab. The core router is only used for packet switching. OSPF is used to provide IPv4 network reachability information and support for label switched paths (LSP) creation. iBGP is used on PE routers for VPNv4 exchange information.

On the provider PE edge routers, the IPv4 information is imported into MPLS network and VPNv4 is exported from MPLS network using virtual route forwarding (VRF) instances. The eBGP protocol is used for IPv4 routing between customer sites and provider network. In this implementation all networks use unique autonomous system numbers (AS) and PE-CE routers work as routing peers.

4.2 Network Implementation

A test network was implemented using available equipment in the MINT lab based on approved design document.

Cisco 2600 series routers were used for P-1, PE-1, PE-2 and PE-3 on the provider MPLS network and 2800 series were used for CE-1, CE-2 and CE-3. Cisco 3750G switches were used in customer sites to provide Cisco 7941G phones and IXIA tester connectivity to the network. A Cisco 3825 router in central site was used as GK-1 gatekeeper.

Two phones have been defined in Cisco Call Manager Express along with appropriate dial plans on each CE router. The internal CE router interfaces were configured as VoIP gateways which would connect to the GK-1 gatekeeper via H.323 protocol for call admission and registration.

The experiment was implemented in three consecutive phases:

- L3 MPLS VPN
- VoIP with no QoS implemented
- VoIP with DiffServ QoS implemented

The last phase, VoIP with DiffServ QoS, was implemented as two scenarios. The first scenario used VoIP traffic priority and WRED for best-effort traffic. The second scenario used VoIP traffic priority and policing to demonstrate QoS policies enforcement for a maximum of 50 available VoIP pairs available with the IxChariot license.

At the site level, QoS was implemented using a voice VLAN and a data VLAN with the CE router providing VLAN routing. On the MPLS domain, QoS was implemented as DiffServ backbone model with traffic classification, marking and policing on the PE routers. Policies were configured and applied on all three PE routers and their interfaces, in both directions, for consistency.

WireShark and session monitoring was used on each 3750G switch to capture traffic generated during various stages of the lab experiment.

MPLS Configuration

Fig. 4.2.1 below presents the MPLS bindings for P-1 core router.

```
P-1#show mpls ldp bindings
  tib entry: 10.0.0.0/30, rev 6
        local binding:  tag: imp-null
        remote binding: tsr: 10.1.1.1:0, tag: 19
        remote binding: tsr: 10.2.2.2:0, tag: imp-null
        remote binding: tsr: 10.3.3.3:0, tag: 17
  tib entry: 10.0.0.4/30, rev 8
        local binding:  tag: imp-null
        remote binding: tsr: 10.1.1.1:0, tag: imp-null
        remote binding: tsr: 10.2.2.2:0, tag: 16
        remote binding: tsr: 10.3.3.3:0, tag: 16
  tib entry: 10.0.0.8/30, rev 5
        local binding:  tag: imp-null
        remote binding: tsr: 10.1.1.1:0, tag: 16
        remote binding: tsr: 10.2.2.2:0, tag: 17
        remote binding: tsr: 10.3.3.3:0, tag: imp-null
  tib entry: 10.1.1.1/32, rev 12
        local binding:  tag: 17
        remote binding: tsr: 10.1.1.1:0, tag: imp-null
        remote binding: tsr: 10.2.2.2:0, tag: 18
        remote binding: tsr: 10.3.3.3:0, tag: 18
  tib entry: 10.2.2.2/32, rev 14
        local binding:  tag: 18
        remote binding: tsr: 10.1.1.1:0, tag: 20
        remote binding: tsr: 10.2.2.2:0, tag: imp-null
        remote binding: tsr: 10.3.3.3:0, tag: 19
  tib entry: 10.3.3.3/32, rev 10
        local binding:  tag: 16
        remote binding: tsr: 10.1.1.1:0, tag: 17
        remote binding: tsr: 10.2.2.2:0, tag: 20
        remote binding: tsr: 10.3.3.3:0, tag: imp-null
  tib entry: 10.4.4.4/32, rev 7
        local binding:  tag: imp-null
        remote binding: tsr: 10.1.1.1:0, tag: 18
        remote binding: tsr: 10.2.2.2:0, tag: 19
        remote binding: tsr: 10.3.3.3:0, tag: 20
```

Fig. 4.2.1 MPLS bindings for P-1 core router

Routing and VPN Configuration

In order to provide a returning path for the traffic between PE-1 and CE-1 router, one additional interface was installed in each router and associated link created.

A traceroute (Fig. 4.2.2) from Site 2 to Site 3 reveals that the traffic will enter provider edge router PE-2 first, then it will be routed through provider core router P-1, provider edge router PE-1 and it will reach customer CE-1 router in central site (172.16.0.6, interface Serial 0/0/0). The return path is over the second link CE-1 to PE-1 (172.16.0.1, interface FastEthernet 0/1), through core router P-1 then via PE-3 router and CE-3 router in Site 3 it will reach destination IP address of SW-3 (172.16.3.254).

```
SW-2#traceroute 172.16.3.254

Type escape sequence to abort.
Tracing the route to 172.16.3.254

  1 172.16.2.129 0 msec 0 msec 0 msec
  2 172.16.0.9 0 msec 0 msec 0 msec
  3 10.0.0.1 0 msec 0 msec 0 msec
  4 172.16.0.5 9 msec 0 msec 0 msec
  5 172.16.0.6 9 msec 0 msec 0 msec
  6 172.16.0.1 8 msec 0 msec 0 msec
  7 10.0.0.5 8 msec 0 msec 9 msec
  8 172.16.0.13 8 msec 0 msec 8 msec
  9 172.16.0.14 8 msec 0 msec 8 msec
 10 172.16.3.254 0 msec 0 msec *
```

Fig. 4.2.2 Traceroute Site 2 to Site 3 - all routers visible

By default all routers are visible on the provider network including core routers. This can be changed using **no mpls ip propagate−ttl** command in global configuration mode so only customer routers, provider ingress and egress routers are visible (Fig. 4.2.3). The provider topology is also hidden from inside which is not desired. By using **no mpls ip propagate−ttl forwarded** on PE routers, an internal traceroute will reveal internal topology which is useful in troubleshooting.

```
SW-2#traceroute 172.16.3.254

Type escape sequence to abort.
Tracing the route to 172.16.3.254

  1 172.16.2.129 0 msec 0 msec 0 msec
  2 172.16.0.9 0 msec 0 msec 0 msec
  3 172.16.0.5 0 msec 0 msec 0 msec
  4 172.16.0.6 9 msec 0 msec 0 msec
  5 172.16.0.1 25 msec 0 msec 0 msec
  6 172.16.0.13 8 msec 0 msec 8 msec
  7 172.16.0.14 9 msec 0 msec 0 msec
  8 172.16.3.254 0 msec 0 msec *
```

Fig. 4.2.3 Traceroute Site 2 to Site 3 - core router not visible

A combination of network protocols was used for the implementation of the Layer 3 MPLS VPN network: MPLS (transport protocol), OSPF (routing support for MPLS switching and forwarding), internal BGP (between PE peer routers, carry VPN information) and external BGP (between PE and CE routers, also configured as peers).

All networks involved are configured as autonomous systems including provider network AS # 1, central site AS # 65001, remote sites AS # 65002 and AS # 65003.

Below are presented BGP VPN table on router PE-1 (Fig. 4.2.4) and IP routing tables on CE-1, CE-2 and CE-3 (Fig. 4.2.5, 4.2.6 and 4.2.7 respectively).

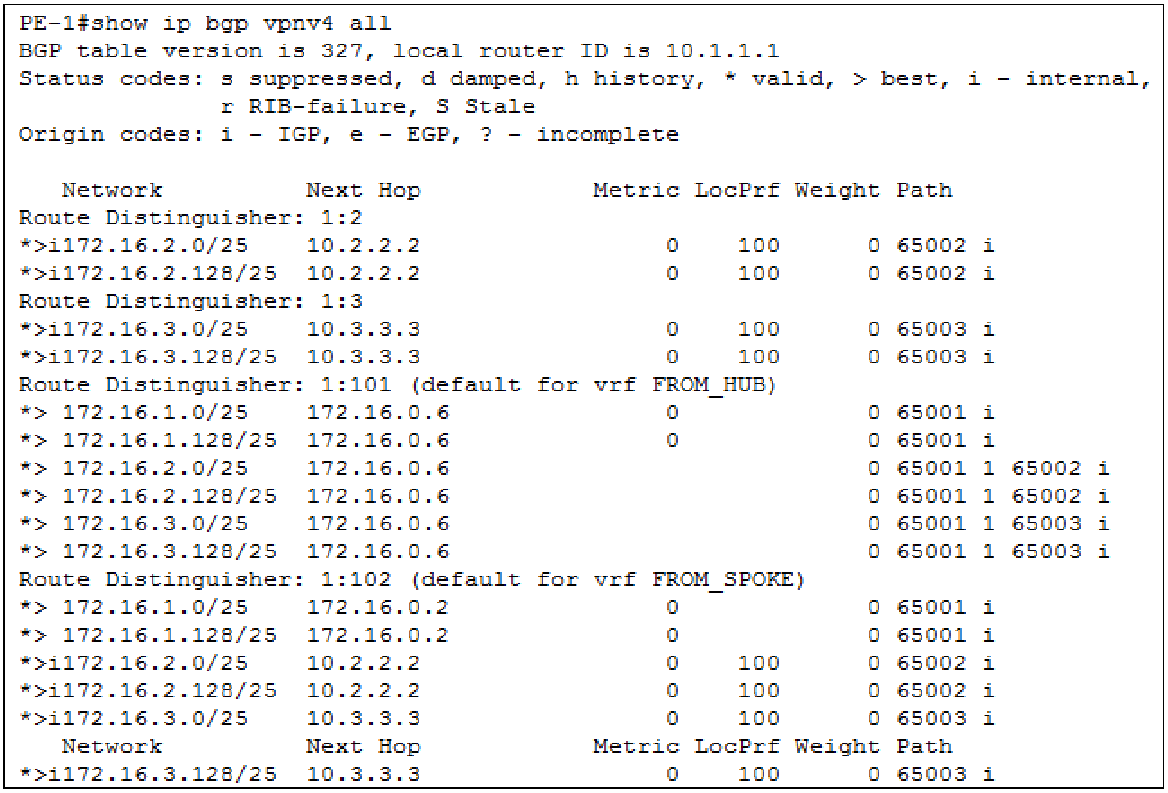

```
PE-1#show ip bgp vpnv4 all
BGP table version is 327, local router ID is 10.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:2
*>i172.16.2.0/25    10.2.2.2                 0    100      0 65002 i
*>i172.16.2.128/25  10.2.2.2                 0    100      0 65002 i
Route Distinguisher: 1:3
*>i172.16.3.0/25    10.3.3.3                 0    100      0 65003 i
*>i172.16.3.128/25  10.3.3.3                 0    100      0 65003 i
Route Distinguisher: 1:101 (default for vrf FROM_HUB)
*> 172.16.1.0/25    172.16.0.6               0             0 65001 i
*> 172.16.1.128/25  172.16.0.6               0             0 65001 i
*> 172.16.2.0/25    172.16.0.6                             0 65001 1 65002 i
*> 172.16.2.128/25  172.16.0.6                             0 65001 1 65002 i
*> 172.16.3.0/25    172.16.0.6                             0 65001 1 65003 i
*> 172.16.3.128/25  172.16.0.6                             0 65001 1 65003 i
Route Distinguisher: 1:102 (default for vrf FROM_SPOKE)
*> 172.16.1.0/25    172.16.0.2               0             0 65001 i
*> 172.16.1.128/25  172.16.0.2               0             0 65001 i
*>i172.16.2.0/25    10.2.2.2                 0    100      0 65002 i
*>i172.16.2.128/25  10.2.2.2                 0    100      0 65002 i
*>i172.16.3.0/25    10.3.3.3                 0    100      0 65003 i
   Network          Next Hop            Metric LocPrf Weight Path
*>i172.16.3.128/25  10.3.3.3                 0    100      0 65003 i
```

Fig. 4.2.4 Router PE-1 BGP VPN table

```
CE-1#show ip route
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

     172.16.0.0/16 is variably subnetted, 8 subnets, 2 masks
C       172.16.1.128/25 is directly connected, GigabitEthernet0/1.1
B       172.16.2.128/25 [20/0] via 172.16.0.1, 3d16h
B       172.16.3.128/25 [20/0] via 172.16.0.1, 01:32:38
C       172.16.0.4/30 is directly connected, Serial0/0/0
C       172.16.0.0/30 is directly connected, GigabitEthernet0/0
C       172.16.1.0/25 is directly connected, GigabitEthernet0/1.10
B       172.16.2.0/25 [20/0] via 172.16.0.1, 3d16h
B       172.16.3.0/25 [20/0] via 172.16.0.1, 01:32:39
```

Fig. 4.2.5 Router CE-1 routing table

```
CE-2#show ip route
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

     172.16.0.0/16 is variably subnetted, 7 subnets, 2 masks
B       172.16.1.128/25 [20/0] via 172.16.0.9, 3d16h
C       172.16.2.128/25 is directly connected, GigabitEthernet0/1.1
B       172.16.3.128/25 [20/0] via 172.16.0.9, 01:34:23
C       172.16.0.8/30 is directly connected, GigabitEthernet0/0
B       172.16.1.0/25 [20/0] via 172.16.0.9, 3d16h
C       172.16.2.0/25 is directly connected, GigabitEthernet0/1.10
B       172.16.3.0/25 [20/0] via 172.16.0.9, 01:34:23
```

Fig. 4.2.6 Router CE-2 routing table

```
CE-3#show ip route
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

     172.16.0.0/16 is variably subnetted, 7 subnets, 2 masks
B       172.16.1.128/25 [20/0] via 172.16.0.13, 01:37:36
B       172.16.2.128/25 [20/0] via 172.16.0.13, 01:37:36
C       172.16.3.128/25 is directly connected, GigabitEthernet0/1.1
C       172.16.0.12/30 is directly connected, GigabitEthernet0/0
B       172.16.1.0/25 [20/0] via 172.16.0.13, 01:37:36
B       172.16.2.0/25 [20/0] via 172.16.0.13, 01:37:36
C       172.16.3.0/25 is directly connected, GigabitEthernet0/1.10
```

Fig. 4.2.7 Router CE-3 routing table

VoIP Basic Configuration with no QoS implemented

Concurrent local site calls and calls over WAN links were initiated and successfully tested. Fig. 4.2.8 presents the VoIP gateways and phones registration with the GK-1 gatekeeper located in central site.

```
GK-1#show gatekeeper endpoints
GATEKEEPER ENDPOINT REGISTRATION
================================
CallSignalAddr  Port  RASSignalAddr   Port  Zone Name        Type    Flags
--------------- ----- --------------- ----- ---------        ----    -----
172.16.1.1      1720  172.16.1.1      54705 GK-1             VOIP-GW
E164-ID: 1001
E164-ID: 1002
H323-ID: CE-1
Voice Capacity Max.=  Avail.=  Current.= 2
172.16.2.1      1720  172.16.2.1      51815 GK-1             VOIP-GW
E164-ID: 2001
E164-ID: 2002
H323-ID: CE-2
Voice Capacity Max.=  Avail.=  Current.= 2
172.16.3.1      1720  172.16.3.1      51217 GK-1             VOIP-GW
E164-ID: 3001
E164-ID: 3002
H323-ID: CE-3
Voice Capacity Max.=  Avail.=  Current.= 0
Total number of active registrations = 3
```

Fig. 4.2.8 Gatekeeper registrations

QoS Policies and Implementation

QoS was implemented at the site level using two VLANs as follows: VLAN1 (default) was used for file transfers while VLAN10 (VOICE) was used for VoIP traffic only. The two VLANs were defined on CE routers using virtual sub-interfaces and 3750G switch interfaces participating in the experiment were configured with specific VoIP instructions as per Cisco documentation.

Cisco IP phones have a built-in three-port switch which will typically mark the voice datagrams with a DSCP value of 101110 (ToS 5) while the packets coming from an attached computer will be marked with a DSCP 000000 (ToS of 0).

A Cisco 3750G switch port is the boundary of the QoS domain and will automatically recognize a Cisco phone via Cisco Discovery Protocol (CDP) and trust the datagram marking. This default configuration allows for out-of-the-box QoS when using Cisco phones and 3750G switches.

As IxChariot DSCP datagram marking is not trusted by Cisco QoS mechanism by default, a **mls qos trust dscp** command was issued on the switch at interface level so DSCP marking will remain unchanged.

A DiffServ template was created in IXIA IxChariot application to mark the VoIP datagrams with DSCP value of 101110. The template can be used to change settings at the bit level therefore emulating various QoS classes.

On the provider network, only the PE routers were configured with policies for traffic classification, marking and policing. Configuration steps include:

- Create classes of traffic
    - specify what traffic is in each class
- Create policy
    - tell what to do with each class
- Apply to an interface
    - service-policy, inbound or outbound.

Two classes of traffic were created on the PE routers: MPLS-EXP-5 (voice traffic) and MPLS-EXP-0 (for file transfers and other types of traffic).

MPLS-EXP-5, with DSCP value of 101110 (EF, express forwarding) was priority reserved 1600 Kbps which corresponds to expected bandwidth required to use all 50 available IxChariot licenses i.e. 32 Kbps x 50 = 1600 Kbps (all overhead included)

MPLS-EXP-0 with DSCP value of 000000 (routine traffic) was allocated 1000 Kbps so when all 50 available licenses are used, traffic generated by file transfers will activate the appropriate policy and excess data traffic will be dropped. This will make sure VoIP calls quality is not affected.

Chapter 5
Experimental Results and Data Analysis

5.1 Test Scenarios and Methodology

The test matrix below (Table 5.1.1 and 5.1.2) was used for test scenarios including no QoS implemented and QoS with priority for VoIP calls and traffic policing for best-effort traffic. Project documentation CD contains over 100 test results for other scenarios including priority bandwidth for VoIP and WRED for best-effort traffic (data transfers).

The test results were exported in all three available formats (html, csv and text) and include extensive information on each test conducted. As an example, test C01 below refers to 1 x VoIP and 1 x Data transfer from Site 2 to Site 1 with no QoS implemented.

Group	Description	Type	SITE 2 to SITE 1	SITE 1 to SITE 2	SITE 2 to SITE 3	SITE 3 to SITE 2
A	**Data Transfer**					
	1 x Data	uni-dir	A01	A02	A03	A04
	50 x Data	uni-dir	A05	A06	A07	A08
B	**VoIP**					
	1-pair	uni-dir	B01	B02	B03	B04
	50-pair	uni-dir	B05	B06	B07	B08
C	**VoIP and 1 x Data**					
	1-pair	uni-dir	C01	C02	C03	C04
	49-pair	uni-dir	C05	C06	C07	C08
D	**VoIP and x Data**					
	5-pair, 5 x Data	uni-dir	D01	D02	D03	D04
	25-pair, 25 x Data	uni-dir	D05	D06	D07	D08

Table 5.1.1 Test Matrix, no QoS implemented

Group	Description	Type	SITE 2 to SITE 1	SITE 1 to SITE 2	SITE 2 to SITE 3	SITE 3 to SITE 2
E	**Data Transfer**					
	1 x Data	uni-dir	E01	E02	E03	E04
	50 x Data	uni-dir	E05	E06	E07	E08
F	**VoIP**					
	1-pair	uni-dir	F01	F02	F03	F04
	50-pair	uni-dir	F05	F06	F07	F08
G	**VoIP and 1 x Data**					
	1-pair	uni-dir	G01	G02	G03	G04
	49-pair	uni-dir	G05	G06	G07	G08
H	**VoIP and x DATA**					
	5-pair, 5 x Data	uni-dir	H01	H02	H03	H04
	25-pair, 25 x Data	uni-dir	H05	H06	H07	H08

Table 5.1.2 Test Matrix, QoS implemented

Testing included throughput tests (file transfers only), VoIP calls (pure VoIP network) as well as combinations of the two, starting with one pair up to maximum of 50-pair combination allowed by the IxChariot license. A “pair” refers to a set of two performance endpoints used for file transfers or unidirectional VoIP call.

5.2 Results and Data Analysis

The tests were conducted using G729 codec for VoIP calls; 100 KB file size for data transfers and all available bandwidth of 8000 Kbps when QoS was not implemented.

Call quality was severely affected even when using all available bandwidth with one data transfer and one VoIP call.

After implementing QoS with traffic classification, prioritization and policing, call quality was maintained within parameters when using only 2600 Kbps bandwidth (1600 Kbps for VoIP, 1000 Kbps for data transfer) out of 8000 Kbps available, while running concurrent 25 x VoIP calls and 25 x data transfers.

Results in graphic format from the selected test scenarios (below) involving Site 2 to Site 1 unidirectional VoIP pairs and data transfers are presented on pages 32-36. Parameters considered are MOS estimate, one-way delay, jitter (delay variation) maximum respectively. Other parameters and results are available on the documentation CD.

- B01 - 1 x VoIP only. No QoS.
- C01 - 1 x VoIP, 1 x Data. No QoS.
- D05 - 25 x VoIP, 25 x Data. No QoS.
- H05 - 25 x VoIP, 25 x Data. QoS implemented.

Test Scenario B01 (Fig. 5.2.1-3)
1 x VoIP only, Site 2 to Site 1. No QoS Implemented.

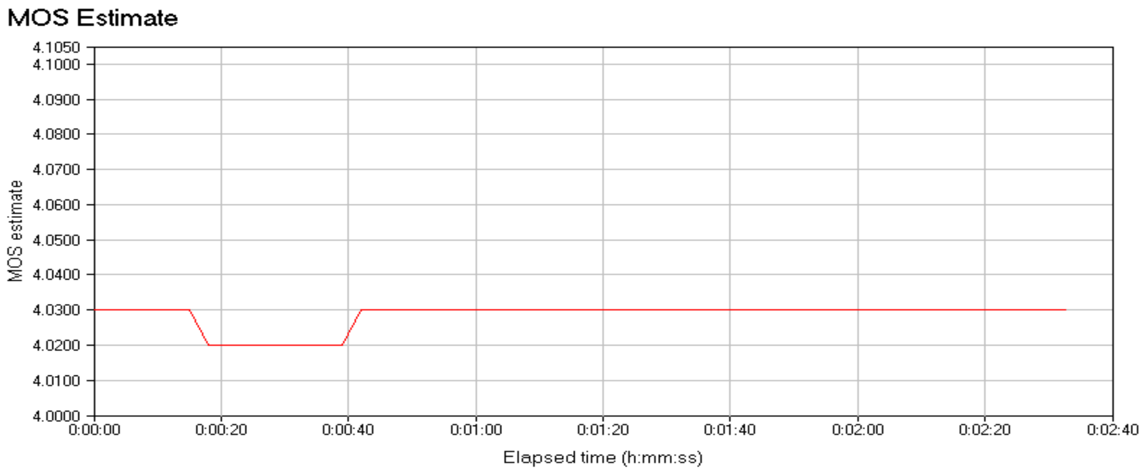

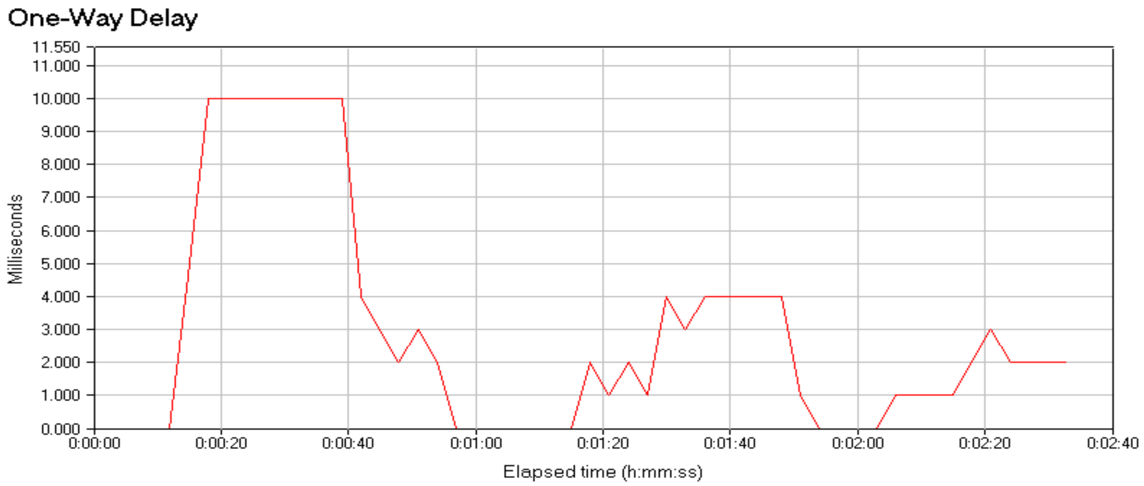

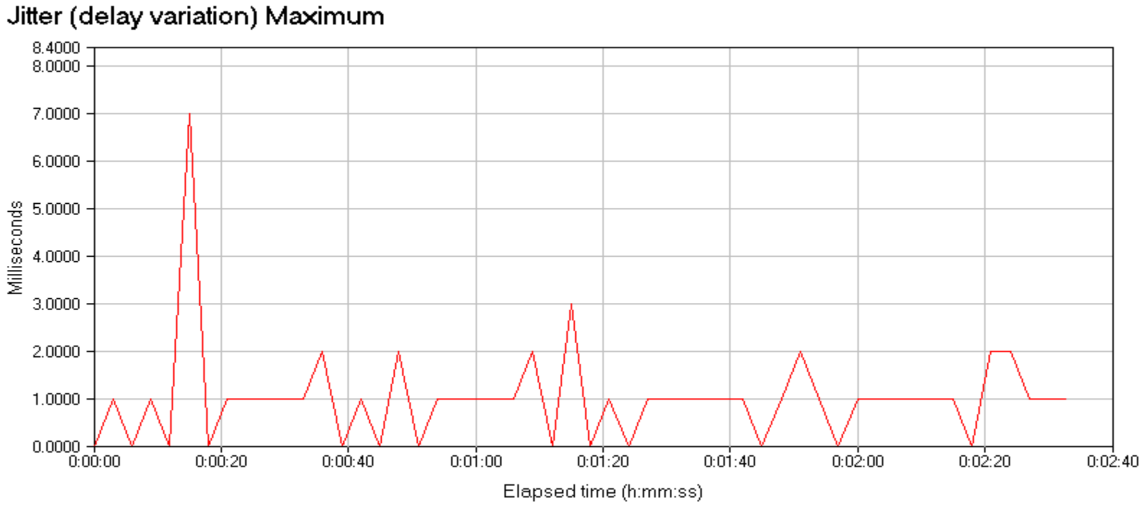

Test Scenario C01(Fig. 5.2.4-6)

1 x VoIP, 1 x Data Site 2 to Site 1. No QoS Implemented.

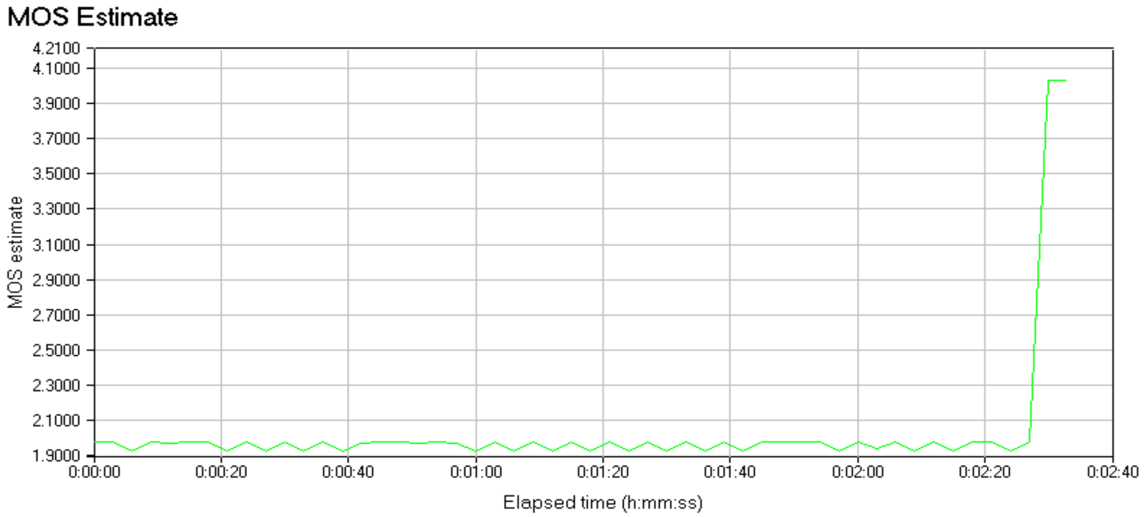

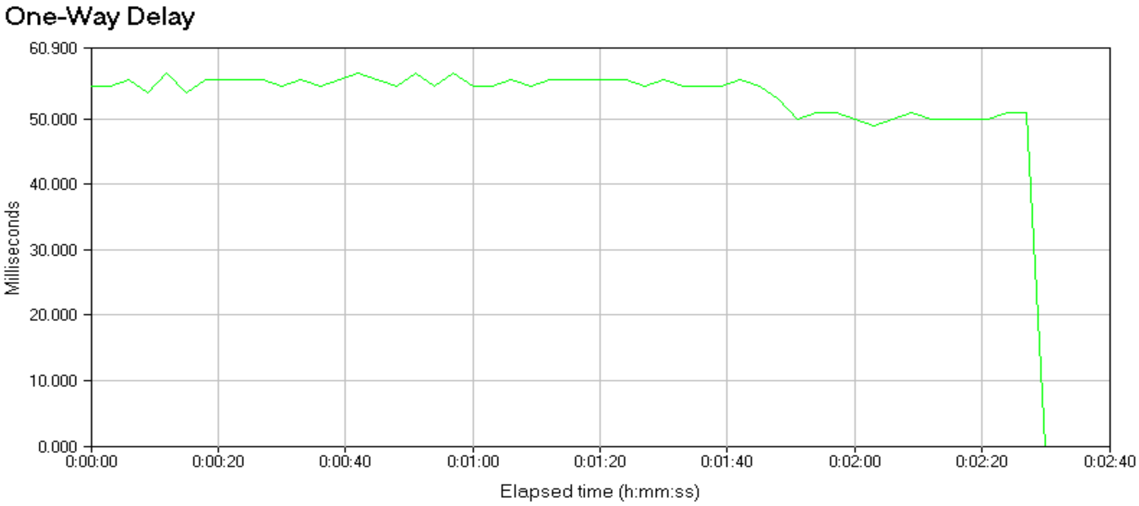

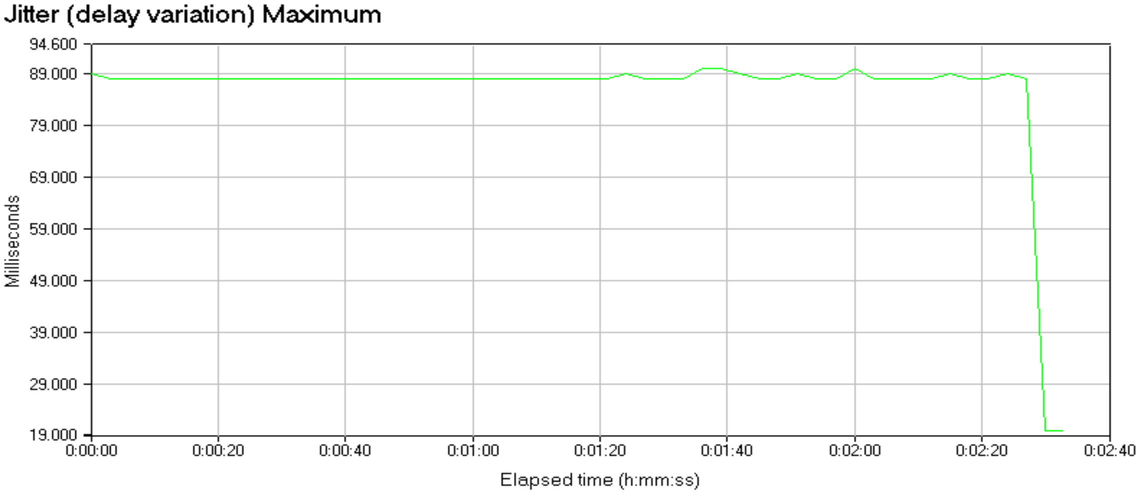

Test Scenario D05 (Fig. 5.2.7-9)
25 x VoIP, 25 x Data Site 2 to Site 1. No QoS Implemented.

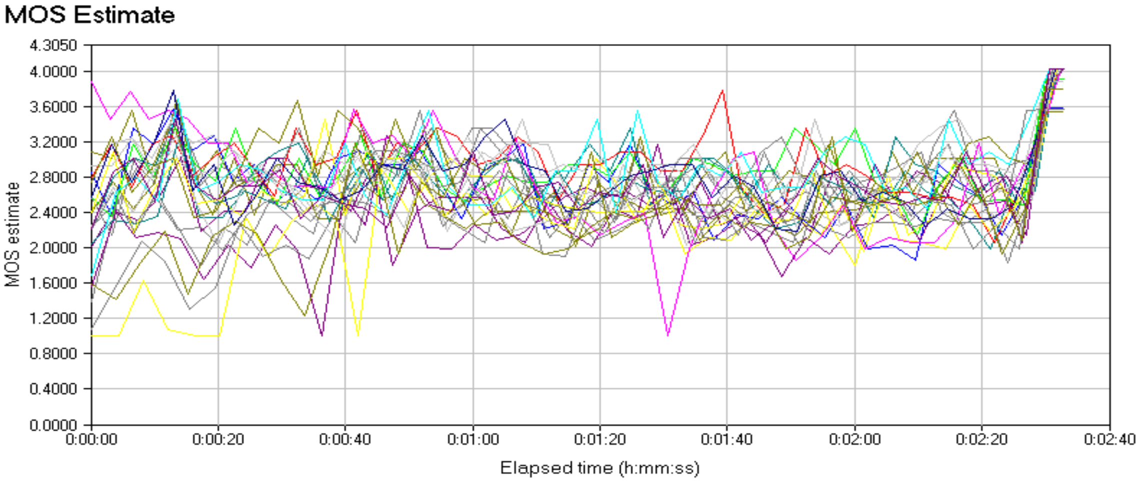

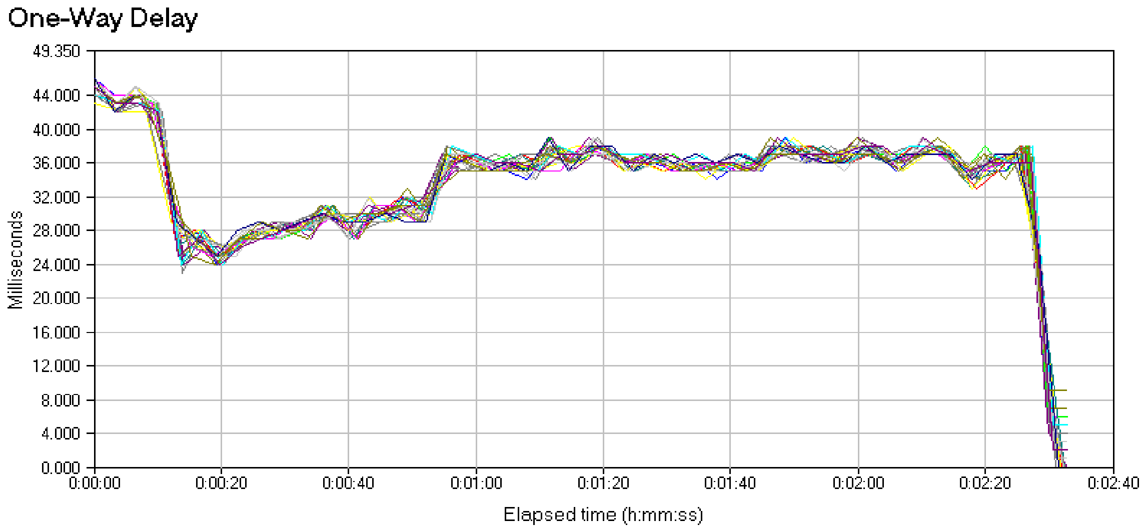

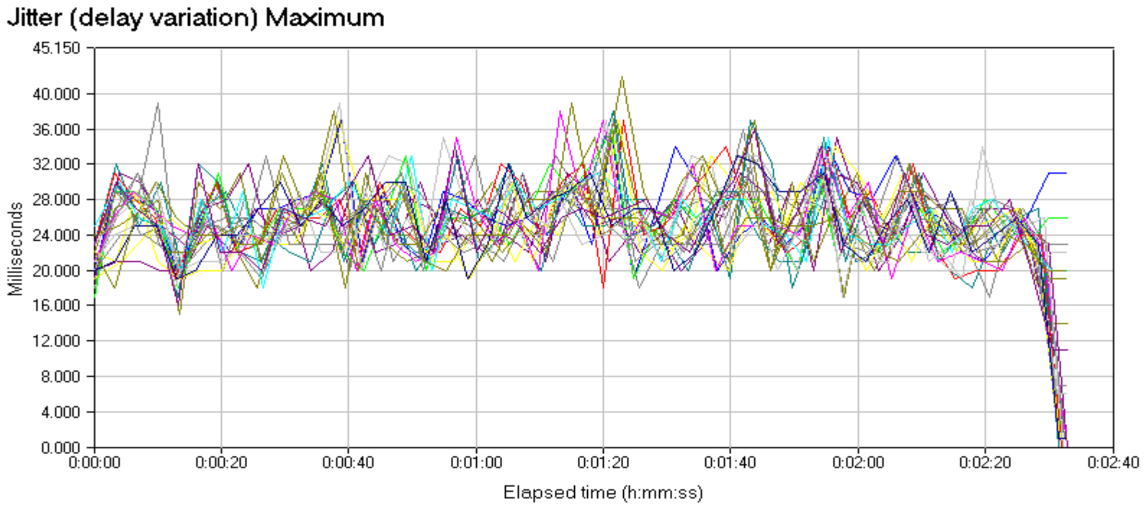

Test Scenario H05 (Fig. 5.2.10-12)
25 x VoIP, 25 x Data Site 2 to Site 1. QoS Implemented.

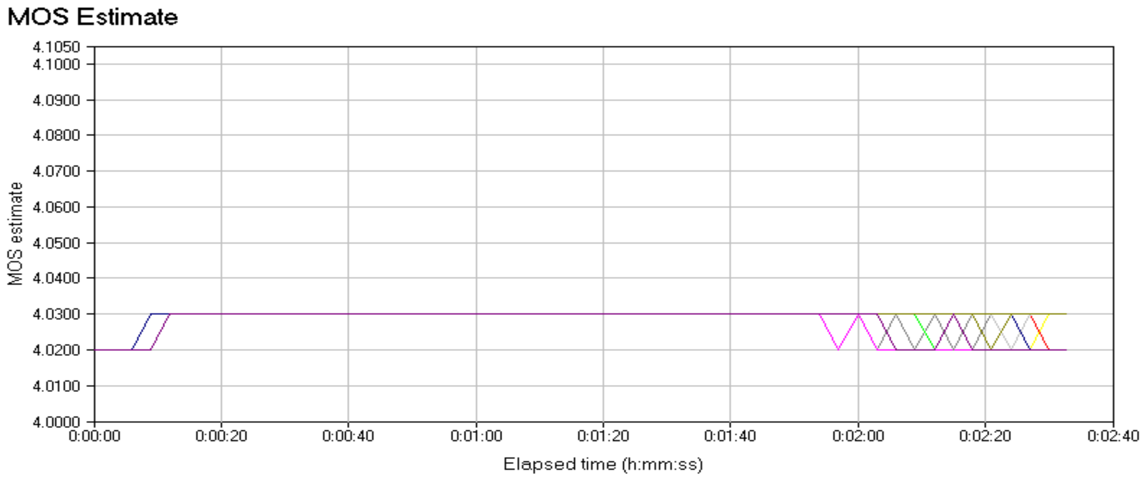

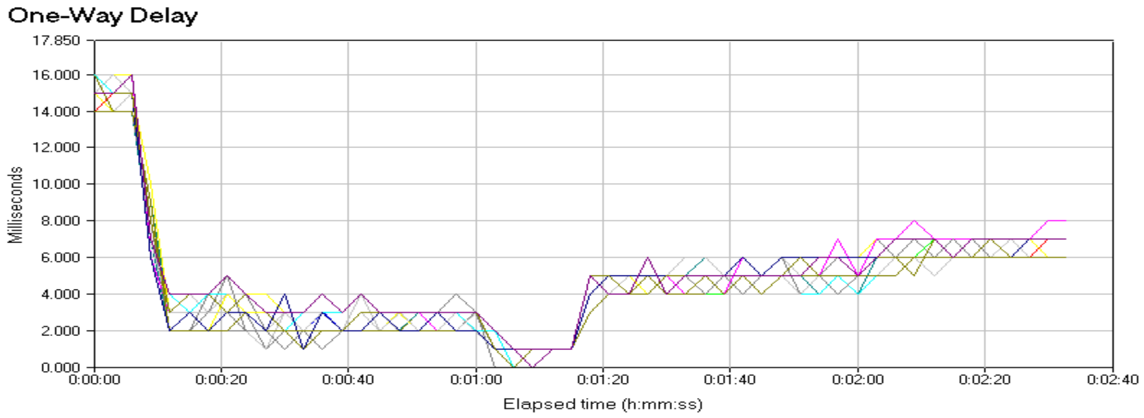

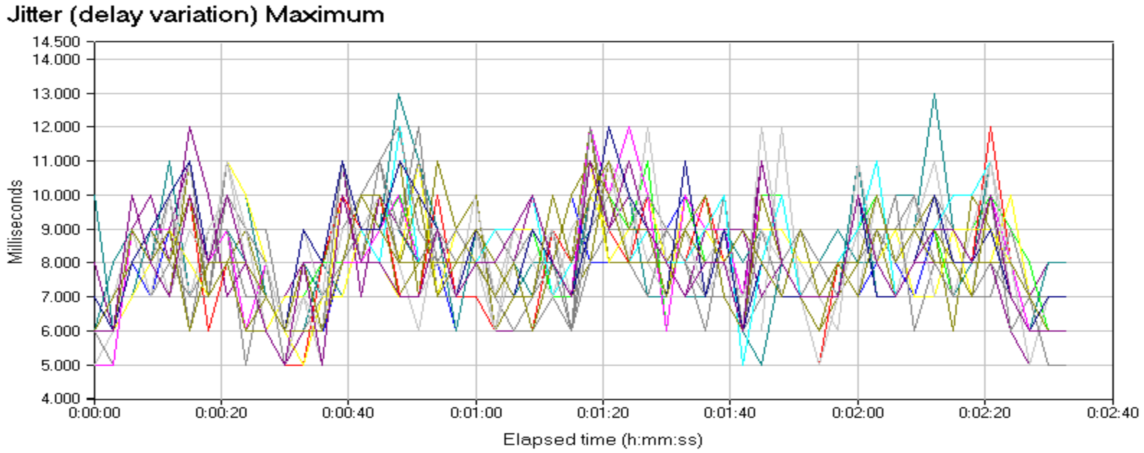

5.3 Cumulative Distribution Function

Below are the cumulative distribution function (CDF) graphs for jitter (delay variation) maximum measured in the same test scenarios presented on page 32.

Test Scenario B01 (Fig. 5.3.1)
1 x VoIP only, Site 2 to Site 1. No QoS Implemented.

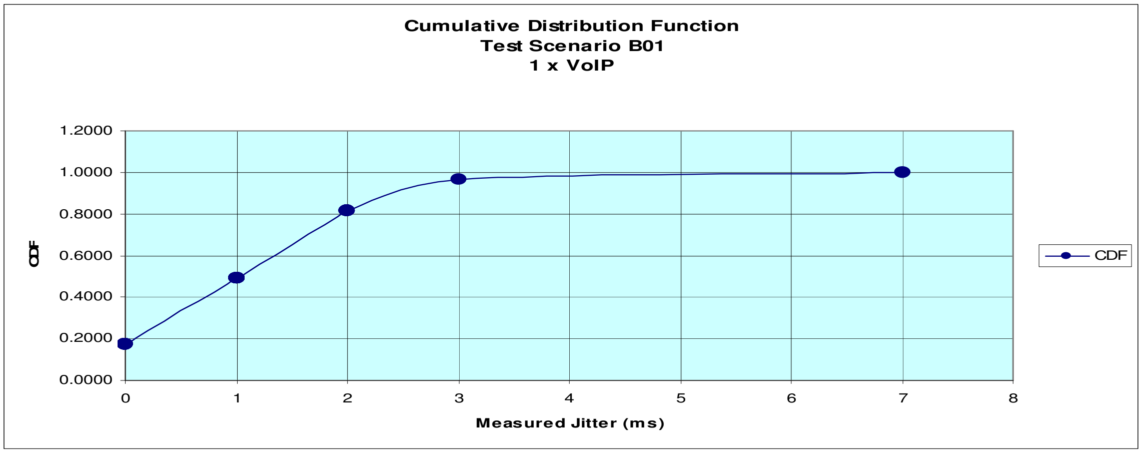

Test Scenario C01 (Fig. 5.3.2)
1 x VoIP, 1 x Data Site 2 to Site 1. No QoS Implemented.

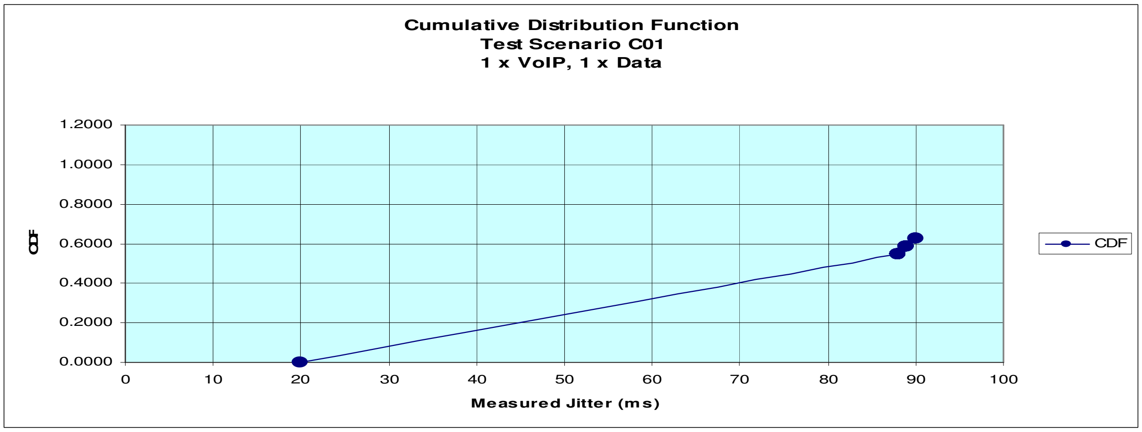

Test Scenario D05 (Fig. 5.3.3)
25 x VoIP, 25 x Data Site 2 to Site 1. No QoS Implemented.

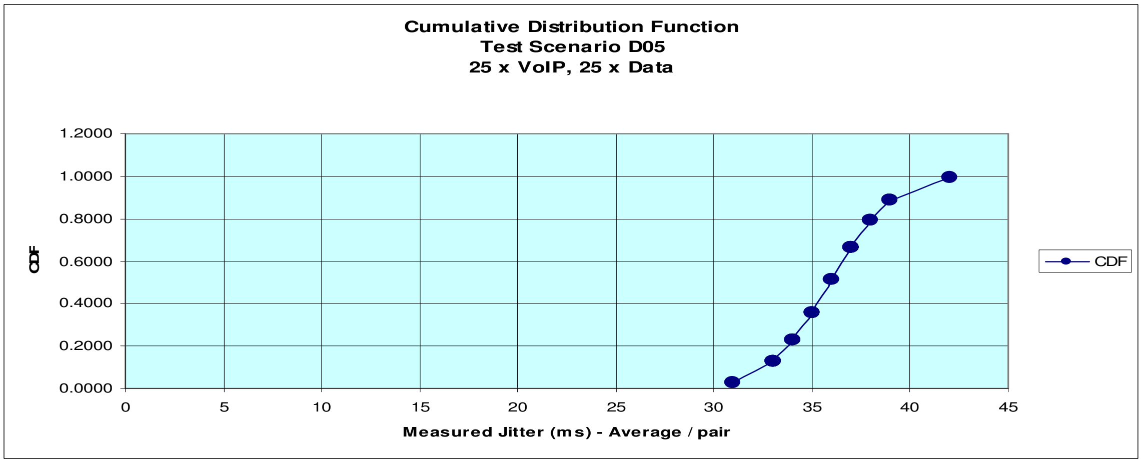

Test Scenario H05 (Fig. 5.3.4)
25 x VoIP, 25 x Data Site 2 to Site 1. QoS Implemented.

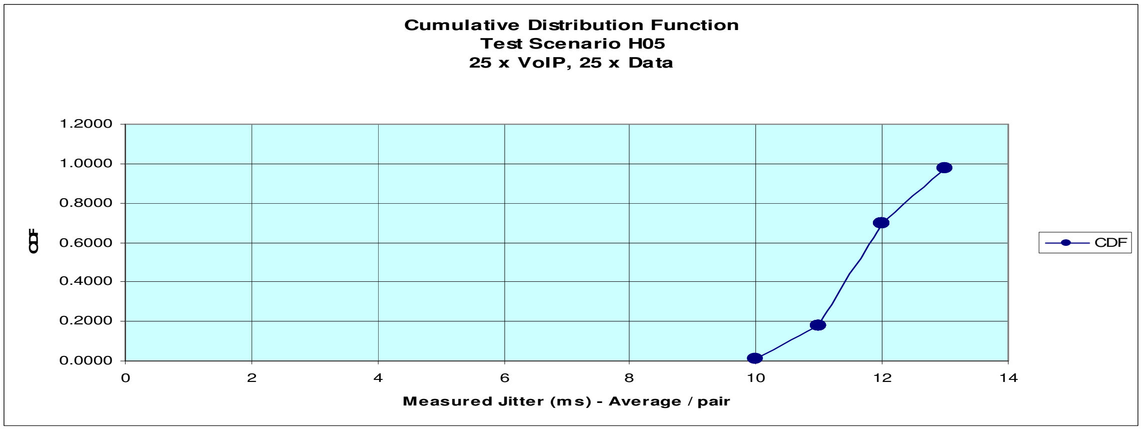

5.4 Traffic Prioritization and Policing

Fig. 5.4.1 below presents interface serial0/0 output from the PE-2 router, with relevant traffic prioritization and policing information highlighted in yellow. The test scenario is H05 where 25 x VoIP and 25 x Data transfers were used concurrently. In this test scenario, QoS was implemented and enforced which resulted in discarding data traffic exceeding 1000 Kbps.

As the default duration of a VoIP test in IxChariot is two an a half minutes (150 seconds), the default load-interval of five minutes for Cisco interface counters has been adjusted to **load-interval 30** seconds. This will still provide accurate information for the IxChariot test. As per Cisco documentation, a minimum of four load-intervals must pass before the average displayed in **show policy-map interface** will be within 2 % of the instantaneous rate of a uniform stream of traffic over that interval.

```
PE-2#show policy-map interface s0/0
 Serial0/0

  Service-policy input: FROM-P-1

    Class-map: MPLS-EXP-5 (match-all)
      0 packets, 0 bytes
      30 second offered rate 0 bps, drop rate 0 bps
      Match: mpls experimental topmost 5
      QoS Set
        qos-group 5
          Packets marked 0

    Class-map: MPLS-EXP-0 (match-all)
      12879 packets, 808378 bytes
      30 second offered rate 49000 bps, drop rate 0 bps
      Match: mpls experimental topmost 0
      QoS Set
        qos-group 0
          Packets marked 12879

    Class-map: class-default (match-any)
      70 packets, 4791 bytes
      30 second offered rate 0 bps, drop rate 0 bps
      Match: any

  Service-policy output: TO-P-1

    Class-map: MPLS-EXP-5 (match-all)
      169445 packets, 12200040 bytes
      30 second offered rate 715000 bps, drop rate 0 bps
      Match: mpls experimental topmost 5
      Queueing
        Strict Priority
        Output Queue: Conversation 264
        Bandwidth 1600 (kbps) Burst 40000 (Bytes)
        (pkts matched/bytes matched) 7430/534960
        (total drops/bytes drops) 0/0

    Class-map: MPLS-EXP-0 (match-all)
      26496 packets, 37339548 bytes
      30 second offered rate 2137000 bps, drop rate 1146000 bps
      Match: mpls experimental topmost 0
      police:
          cir 1000000 bps, bc 31250 bytes, be 31250 bytes
        conformed 12993 packets, 17111064 bytes; actions:
          transmit
        exceeded 27 packets, 31216 bytes; actions:
          drop
        violated 13612 packets, 20401812 bytes; actions:
          drop
        conformed 994000 bps, exceed 0 bps, violate 1146000 bps

    Class-map: class-default (match-any)
      91 packets, 6252 bytes
      30 second offered rate 0 bps, drop rate 0 bps
      Match: any
PE-2#show policy-map interface fa0/0
 FastEthernet0/0

  Service-policy input: FROM-CE-2

    Class-map: VOICE (match-all)
      176746 packets, 13079204 bytes
      30 second offered rate 736000 bps, drop rate 0 bps
      Match: access-group name VOICE
      QoS Set
        qos-group 5
          Packets marked 176753
        mpls experimental imposition 5
          Packets marked 176754

    Class-map: class-default (match-any)
      27717 packets, 39172609 bytes
      30 second offered rate 2151000 bps, drop rate 0 bps
      Match: any
      QoS Set
        qos-group 0
          Packets marked 27717
        mpls experimental imposition 0
          Packets marked 27717

  Service-policy output: TO-CE-2

    Class-map: QoS-GROUP-5 (match-all)
      0 packets, 0 bytes
      30 second offered rate 0 bps, drop rate 0 bps
      Match: qos-group 5
      Queueing
        Strict Priority
        Output Queue: Conversation 264
        Bandwidth 1600 (kbps) Burst 40000 (Bytes)
        (pkts matched/bytes matched) 0/0
        (total drops/bytes drops) 0/0

    Class-map: QoS-GROUP-0 (match-all)
      13503 packets, 929053 bytes
      30 second offered rate 55000 bps, drop rate 0 bps
      Match: qos-group 0
      police:
          cir 1000000 bps, bc 31250 bytes, be 31250 bytes
        conformed 13571 packets, 932981 bytes; actions:
          transmit
        exceeded 0 packets, 0 bytes; actions:
          drop
        violated 0 packets, 0 bytes; actions:
          drop
        conformed 55000 bps, exceed 0 bps, violate 0 bps

    Class-map: class-default (match-any)
      0 packets, 0 bytes
      30 second offered rate 0 bps, drop rate 0 bps
      Match: any
  PE-2#
```

Fig. 5.4.1 - PE-2 Traffic Prioritization and Policing - interface counters

Conclusion

Many service providers running MPLS VPN are looking at interconnecting their networks to other providers MPLS VPN to simplify operations, reduce cost and improve scalability. Inter-autonomous MPLS VPN and Carrier’s Carrier can be used to achieve this goal of inter-provider connectivity.

On the provider network, MPLS VPN was introduced to meet the need for a scalable, efficient and cost effective peer-to-peer VPN model and it is the most popular and widespread implementation of MPLS technology. It can be implemented at Layer 2 using various technologies (one of them being tunneling L2TPv3) or Layer 3 using a Peer-to-Peer model with associated routing protocols.

Layer 3 MPLS VPN connects customer sites as an IP only solution in a peer-to-peer arrangement. Inter-site connectivity is transparent to the customer which can leave CE configuration and management to the service provider. QoS can be implemented by service provider independently or in coordination with customer QoS policies implemented internally at the organization level. One popular application of QoS is Voice over IP which is part of the triple-play (voice, video and data) offerings from service providers.

This project report presented the work that was performed in the MINT lab on design, implement and analysis of VoIP deployments in multi-site L3 MPLS VPN environments.

Results show that when using full bandwidth of 8 Mbps for VoIP and data transfers with no QoS implemented, call quality is severely impacted starting with just one single data transfer.

After implementing QoS on the provider L3 MPLS VPN network and only using 2.6 Mbps out of 8 Mbps available bandwidth (1.6 Mbps for voice, 1.0 Mbps for data transfers), call quality was maintained within acceptable parameters for concurrent 25 VoIP calls and 25 data transfers.

The lab results and data analysis demonstrate DiffServ QoS as a viable method for providing concurrent VoIP and data services to customers while using network resources efficiently. This model can be easily extended in the lab to emulate triple play services offered by service providers running Layer 3 MPLS VPN.

References

[1] Cisco Call Manager Express (CME) 3.0 Design Guide, http://www.cisco.com, 2005

[2] Cisco Catalyst 3750 Switch Software Configuration Guide, http://www.cisco.com, Release 12.2(25)SE, 2004

[3] Cisco IOS H.323 Configuration Guide, http://www.cisco.com, 2006

[4] Cisco IOS Multiprotocol Label Switching Configuration Guide, Release 12.4, http://www.cisco.com, 2008

[5] Cisco IOS Quality of Service Solutions Command Reference, Release 12.2, http://www.cisco.com, 2006

[6] Definitive MPLS Network Designs, Jim Guichard, François Le Faucheur, Jean-Philippe Vasseur, Cisco Press, 2005

[7] Internet Routing Architectures, Second Edition, Sam Halabi, Cisco Press, 2001

[8] IxChariot User Guide Release 6.70, IXIA system documentation, 2008

[9] MINT 708 - Internet Laboratory, University of Alberta, May 2008

[10] MINT 709 Capstone Project Report - Analysis of Traffic Engineering Deployment Strategies in Core IP/MPLS Networks, Mohamed Hasan Omar, University of Alberta, April 2008

[11] MINT 715 - Advanced Routing and Network Management, University of Alberta, February 2008

[12] MPLS Fundamentals, Luc De Ghein, Cisco Press, 2007

[13] MPLS Configuration on Cisco IOS Software, Lancy Lobo, Umesh Lakshman, Cisco Press, 2006

[14] Multi-Protocol Label Switching (MPLS) Conformance and Performance Testing, http://www.ixiacom.com, 2004

[15] QoS for IP/MPLS Networks, Santiago Alvarez, Cisco Press, 2006

[16] RFC2475 - An Architecture for Differentiated Service, http://www.rfc-editor.org, 1998

[17] RFC2547 - BGP/MPLS VPNs, http://www.rfc-editor.org,

[18] RFC3270 - Multi-Protocol Label Switching (MPLS) Support of Differentiated Services, http://www.rfc-editor.org, 2002

[19] Understanding and Implementing Quality of Service in Cisco Multilayer Switched Networks, Erum Frahim, Richard Froom, Balaji Sivasubramanian, Cisco Press, 2004

[20] VoIP Testing with IxChariot, http://www.ixiacom.com, 2005

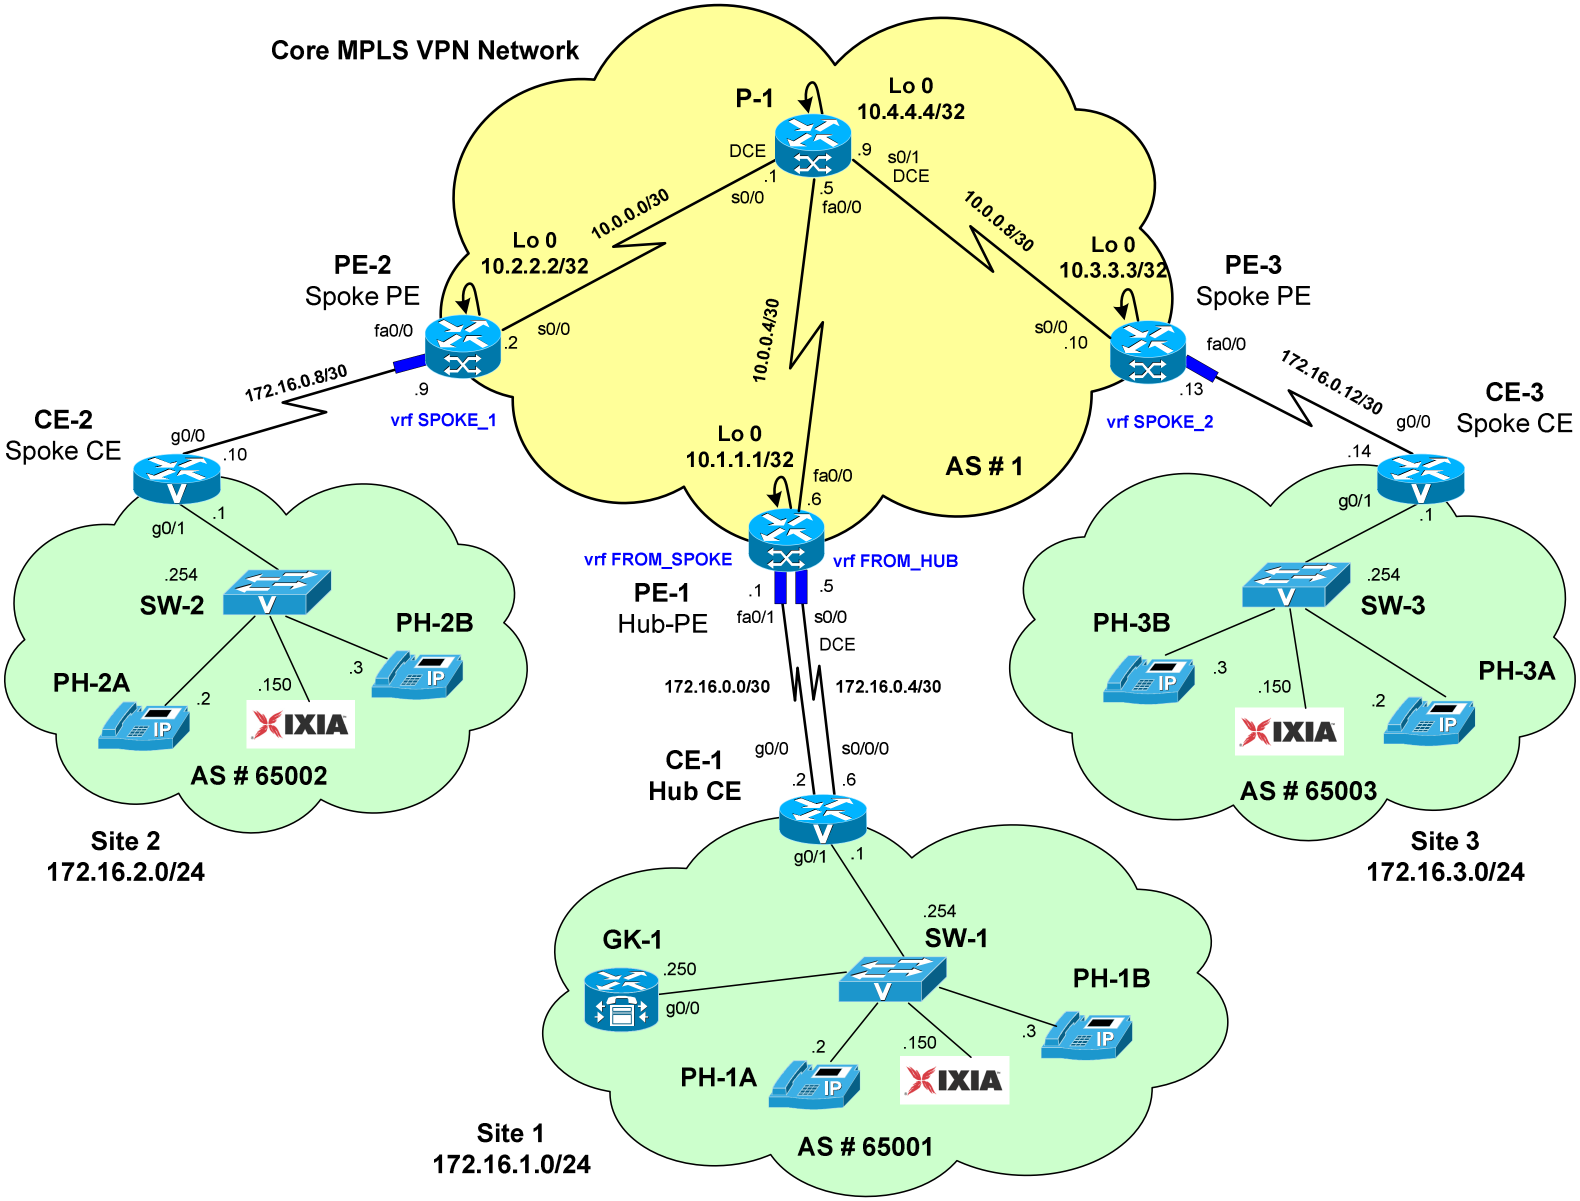
Core MPLS VPN Network
P-1
Lo 0
10.4.4.4/32
DCE
.1
s0/0
.9
s0/1
DCE
.5
fa0/0
10.0.0.0/30
10.0.0.8/30
10.0.0.4/30
Lo 0
10.2.2.2/32
PE-2
Spoke PE
fa0/0
.2
s0/0
.9
vrf SPOKE_1
Lo 0
10.3.3.3/32
PE-3
Spoke PE
s0/0
.10
fa0/0
.13
vrf SPOKE_2
172.16.0.8/30
172.16.0.12/30
CE-2
Spoke CE
g0/0
.10
CE-3
Spoke CE
g0/0
.14
Lo 0
10.1.1.1/32
fa0/0
.6
AS # 1
vrf FROM_SPOKE
vrf FROM_HUB
PE-1
Hub-PE
.1
fa0/1
.5
s0/0
DCE
172.16.0.0/30
172.16.0.4/30
CE-1
Hub CE
g0/0
.2
s0/0/0
.6
.1
g0/1
.254
SW-2
PH-2B
.3
PH-2A
.2
.150
IXIA
IP
AS # 65002
Site 2
172.16.2.0/24
g0/1
.1
.254
SW-3
PH-3B
.3
.150
PH-3A
.2
IXIA
AS # 65003
Site 3
172.16.3.0/24
g0/1
.1
.254
SW-1
GK-1
.250
g0/0
PH-1B
.3
.2
.150
PH-1A
IXIA
AS # 65001
Site 1
172.16.1.0/24

List of Figures

List of Tables

Glossary

The list of terms used in this report is taken mainly from [14] with other additions from various sources provided in References section.

Term	Description
Border Gateway Protocol (BGP)	An exterior gateway protocol defined in RFC 1267 and RFC 1268. BGP is the principal protocol used along the Internet backbone and within larger organizations.
Cisco Call Manager Express (CME)	VoIP functionality included in specific Cisco 2800 router operating system.
Class of Service (CoS)	Class of Service (CoS) is a method for managing network traffic by grouping similar types of traffic (for example, e-mail, streaming video, voice, large document file transfer) together and treating each type as a class with its own level of service priority.
Customer Edge Router (CE)	A router at the edge of a customer network, the CE interfaces to a corresponding Provider Edge (PE) router at the edge of the service provider's network.
DiffServ (Differentiated Services)	An architecture for providing different types or levels of service for network traffic.
Exterior Gateway Protocol (EGP)	A protocol that distributes routing information to the routers that connect networks.
Forward Equivalency Class (FEC)	A classification of a group of packets - all packets assigned to a FEC receive the same routing treatment. FECs can be based on IP address prefixes or service requirements for a type of packet (QoS, VPN, Traffic Engineering, etc).
Forwarding Information Base (FIB)	A table containing the information necessary to forward IP data in a router. At a minimum, the FIB contains the outbound interface identifier and next hop information for each reachable IP destination network.
Internet Gateway Protocol (IGP)	Protocol that distributes routing information to the routers within a network. The term "gateway" is historical; "router" is currently the preferred term. Example IGPs are OSPF, IS-IS and RIP.
L2 VPN	An emulation of a Layer 2 switching environment, supplied by a service provider for its customers, via a core network. In MPLS networks, L2 VPNs use LDP to signal connections for transporting Layer 2 frames over MPLS.
L3 VPN	An emulation of Layer 3 services/distribution of routes, supplied by a service provider for its customers, via a core network. L3 VPNs use BGP extensions to signal provider-provisioned VPNs per IETF Draft RFC 2547bis.
Label Distribution Protocol (LDP)	A protocol, defined in RFC 3036, used to distribute MPLS label and stream mapping information.

Term	Description
Label Edge Router (LER)	A router at the edge of an MPLS network. At the ingress side, an LER maps IP packets to LSPs, adds the appropriate MPLS header, and forwards the packet to the next hop. At the egress side, an LER strips the MPLS label(s) and forwards the packet using traditional routing mechanisms.
Label Information Base (LIB)	A table that specifies how to forward a packet in an MPLS router. This table associates each label with its corresponding FEC.
Label Switched Path (LSP)	In MPLS, a path through a network from an ingress to an egress router that has been established through the distribution of labels that define hop-by-hop forwarding treatment.
Label Switching Router (LSR)	A router, operating in the core of an MPLS network, that switches traffic based on labels.
Open Shortest Path First (OSPF)	A link-state routing protocol used by IP routers located within a single Autonomous System (AS) to determine routing paths. MPLS traffic engineering parameters can be distributed with OSPF using extensions to the protocol (OSPF-TE).
P Router (Provider Router)	A router that operates in the core of a service provider network.
Provider Edge Router (PE)	A router that operates at the edge of a service provider's network, interfacing with the corresponding Customer Edge (CE) router(s) at the edge of one or more customer networks.
Quality of Service (QoS)	A measure of performance for a transmission system that reflects its transmission quality and service availability. QoS mechanisms provide the ability to manage network traffic's bandwidth, delay, and congestion.
Upstream and Downstream	Data intended for a particular destination network always flows downstream. Updates (routing protocol or label distribution LDP/TDP pertaining to a specific prefix are always propagate upstream.
Virtual Private LAN Service (VPLS)	A class of VPN that supports the connection of multiple sites in a single bridged domain over a managed IP/MPLS network. The goal of VPLS is to provide a protocol-transparent, any-to-any, full-mesh service across a WAN.
Virtual Routing and Forwarding Table (VRF)	A VPN routing/forwarding instance. A VRF includes the routing information that defines a customer VPN site that is attached to a PE router.
Weighted Random Early Detection (WRED)	A queuing algorithm used in congestion avoidance where a single queue may have several different queue thresholds associated to IP precedence or DSCP values.