Understanding and Implementing High Definition Videoconferencing

Building a Better Telepresence Service Using Junos

by Marc Bernstein, Darpan Gogia, Vijay Kamisetty,
Taras Matselyukh, and Naveen Udoshi

JUNIPER NETWORKS
VALIDATED SOLUTION

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice. Products made or sold by Juniper Networks or components thereof might be covered by one or more of the following patents that are owned by or licensed to Juniper Networks: U.S. Patent Nos. 5,473,599, 5,905,725, 5,909,440, 6,192,051, 6,333,650, 6,359,479, 6,406,312, 6,429,706, 6,459,579, 6,493,347, 6,538,518, 6,538,899, 6,552,918, 6,567,902, 6,578,186, and 6,590,785.

Printed in the USA.

Version History: v1 Dec 2010

2 3 4 5 6 7 8 9 10

Key Contributors

Louise Apichell is a Juniper Networks Senior Technical Writing Specialist in the Solutions Marketing Group. She assisted as a content developer, chief editor and project manager in organizing, writing and editing this book. Louise specializes in writing and editing all types of technical collateral, such as white papers, application notes, implementation guides, reference architectures and solution briefs.

Mike Barker is a Juniper Networks Technical Marketing Director, Solutions Engineering and Architectures. In this role, he focuses on developing architectures and validating multi-product solutions that create business value for enterprise and Service Provider customers. Prior to this role, Mike served in various Consulting and Systems Engineering roles for Federal, Enterprise and Service Provider markets at Juniper Networks, Acorn Packet Solutions and Arbor Networks. Earlier in his career, Mike held Network Engineering positions at Cable & Wireless, Stanford Telecom and the USAF. Mr. Barker holds a Bachelors of Science Degree in Business Management from Mount Olive College and a MBA from Mount St. Mary's University.

Karen Joice is a Juniper Networks Marketing Specialist who provided the technical illustrations. Karen has been a graphic artist and marketing professional for more than 15 years, specializing in technical illustrations, Flash and Web design, with expertise in print production.

Chandra Shekhar Pandey is a Juniper Networks Director of Solutions Engineering. He is responsible for service provider, enterprise and OEM partners' solutions engineering and validation. Chandra has more than18 years of networking experience, designing ASICs, architecting systems and designing solutions to address customer's challenges in the service providers, MSO and enterprise market. He holds a bachelor's degree in Electronics Engineering from K.N.I.T, Sultanpur, India and a MBA in High Tech and Finance from Northeastern University, Boston, MA**.**

You can purchase a printed copy of this book, or download a free PDF version of this book, at: **juniper.net/books**.

Contents

Part Two

Lead Authors

Marc Bernstein is a Solutions Engineering Manager with Juniper's Solutions Engineering and Architecture team (SEA) , which integrates Juniper and non-Juniper products to solve customer challenges. He was formerly a manager within Product Management focusing on WAN technologies, Sales Engineering Specialist, and most recently, a Solutions Architect focused on triple-play service delivery across residential broadband networks. He holds an MBA from University of Michigan and a bachelor's degree in Electrical Engineering.

Darpan Gogia is a Staff Solutions Engineer in the Solutions Marketing Group at Juniper Networks. He designs and validates a wide range of solutions in both wired and wireless network deployments. His specialization is security solutions for the mobile networks, cloud/data centers, OEM alliances, enterprise and service providers markets. Darpan brings several years of relevant technical expertise and has authored a number of white papers, solution reports, application notes and design guides based on validation efforts. He holds a bachelor's degree in Entrepreneurship and Small Business Management from Delhi University, India.

Vijay Kamisetty is a Juniper Networks Staff Solutions Design Engineer in the Solutions Engineering Group. He is involved in technical solutions for IPTV-Multiplay, mobile backhaul, and application level security in the datacenter, development of managed services, adaptive clock recovery and validated mobile traffic offload services. He holds a bachelor's degree in Computer Science from JNTU HYD, India.

Taras Matselyukh is Juniper Networks Senior Solutions Architect for Business Services in Solutions Architecture and Engineering group and has more than 18 years of experience in IT and Networking. He specializes in development of innovative technical architectures and business models for delivery of new generation services to Business customers of Service Providers. He has developed reference architecture for the joint HD videoconferencing solution of Juniper Networks and Polycom and led engineering teams during creation of this solution. Taras has authored a number of white papers and solution briefs dedicated to different areas of business services based on his work and this is his first written book. He holds degrees in Electronics Engineering and Economics from Kiev National Taras Shevchenko University in Ukraine.

Naveen Udoshi is a Juniper Networks Staff Solutions Design Engineer in the Solutions Engineering Group. He is involved in designing, implementing and validating solutions that range from business and residential services to data center aggregation and cloud services. Naveen brings several years of product development experience on telecom and data networking products including Juniper's Universal Edge router. He holds a bachelor's in Electronics and Communication from Karnatak University, India.

Authors Acknowledgements

The authors would like to take this opportunity to thank the following Juniper Networks engineers:

MX Series, Junos Engineering Team

Paul Raison
Cristina Radulescu-Banu
Richard Pagliaro
Poorna Pushkala B
Bill Miller
Subrat Pani
Russ Lamoreaux
Julius Francis
Lalit Shinde

SRC Engineering Team

Lilian Moortgat
Fred Gao
Piotr Drozdzewicz
Sherine El-Medani
Greg Sidebottom

Advanced JTAC Team

Charlie (Chang-Hyun) Kim
Ameya Deshpande

Other Juniper Team Members

Gregory Dalle
Art Norton
Robert Sprouse
Andrey Petrovsky
Greg Sidebottom
Marcelino Chua
Nancy Cohen
Andrey Petrovsky
Arthur Norton
Dwight Buys
Scott Stevens
John Rolfe

Polycom Team

Jeff Adams
Clint Edwards
Piotr Drozdzewicz
Terri Gott
Philip Walenta
Randy Schultz
Megan Bouhamama

Preface

By integrating high definition (HD) video and audio technology, videoconferencing has now reached the stage where it remarkably reproduces “live” meetings around the world. Technological advances in network performance, security and accessibility have elevated HD videoconferencing onto the center stage of distributed enterprises. This is why analysts forecast explosive market growth for videoconferencing and telepresence in the next two to three years. HD videoconferencing enables enterprises to reduce the expense and lost productivity from excessive travel, while service providers who offer HD videoconferencing service will be ready to serve these expanding business needs.

This handbook describes how service providers and enterprises can cost-effectively expand their network to assure high performance and quality of HD videoconferencing services while avoiding common pitfalls along the way. It begins with a simple, high-level overview of the design principles and factors that influence the quality and reliability of the HD videoconferencing services, followed by practical use cases, which provide comprehensive network diagrams and device configurations with advanced concepts at the end of the book. This handbook also describes how service provides and large enterprises can implement assured forwarding, which ensures service quality even when there may be insufficient bandwidth. It serves as a reference tool on how to design, deploy and operate new HD videoconferencing services with assured quality, guaranteed availability and unprecedented scalability and security.

This handbook is divided into three sections:

- **Part 1** provides an overview of the HD videoconferencing market, evolution and requirements. This section is appropriate for those seeking an introduction to videoconferencing including marketing personnel and technical managers.
- **Part 2** explains concepts and design criteria starting from the early pre-design stage, provides detailed instruction and validated configurations for the network and videoconferencing devices required during the implementation stage. It also provides commands to assist in detecting network configuration issues. This section is intended for network engineers who specialize in Juniper Networks products, including network architects and network administrators engaged in design, deployment and maintenance of HD videoconferencing services. An understanding of Junos® configuration is recommended.
- The **appendices** provide primers on relevant products and technologies.

Please send us your feedback or ideas that you would like to see covered in future revisions or in other validated solutions books at: **solutions-content@juniper.net .**

Chapter 1

Evolution of Videoconferencing Services

Part One

Just as the auto industry was forced to adapt to changing demands brought about by an economic downturn and increased gasoline prices, other industries are changing their communication practices to reduce cost and speed decision-making. This has opened the door for a communications technology that is poised to revolutionize traditional communication practices, high definition (HD) videoconferencing. By providing the same sensation as if the remote party were present in the same room, HD videoconferencing holds the promise of offering productive face-to-face meetings thereby avoiding travel expenses and precious time. Closer to home, HD videoconferencing enables more effective teleworking, eliminating commute time without sacrificing the quality of human interactions.

The worldwide videoconferencing systems market is expected to grow at a compound annual growth rate (CAGR) of 17.6 percent through 2015 reaching nearly US $4.2 billion[1]. The associated managed services market is accelerating to a billion dollar business, showing a compound annual growth rate of 41.5 percent between 2008 and 2015, rising from $82.7 million to $938.3 million[2].

What's in a Name: Telepresence or HD Videoconferencing?

High definition videoconferencing refers to any videoconferencing system that uses high definition video and audio technologies. The term telepresence is also used to describe this and this usage and pre-dates Cisco's trademark spelling of their TelePresence™ products, which they have defined as applying only to high-end immersive systems. Because the concepts and configurations described in this handbook are applicable to all videoconferencing systems, it uses the term HD videoconferencing to highlight that the concepts are applicable to all types of systems.

What makes HD videoconferencing different from traditional videoconferencing? The almost life-size appearance of the remote parties accurately reproduces the effect of a face-to-face meeting. High-resolution video reveals even tiny gestures and facial expressions. Camera and microphone placement help establish eye contact and accurate sound reproduction. These technical advances combine to create the illusion of face-to-face human interaction. Numerous studies show that more than half of human communication is non-verbal. For example, one well-known analysis determined that words account for only 7 percent of all communication impact, while voice tone accounts for 38 percent and non-verbal cues account for 55 percent. Therefore, HD videoconferencing is critical to effective communications. Today's phone conference calls and traditional videoconferencing simply cannot compete with the effectiveness of face-to-face communication.

[1] Asia Pacific Video Conferencing Endpoints Markets CY2009, **frost.com/prod/servlet/report-toc.pag?ctxixpLink=FcmCtx1&searchQuery=Asia-Pacific+Video+Conferencing+Endpoints+Market&repid=P379-01-00-00-00&bdata=aHR0cDovL3d3dy5mcm9zdC5jb20vc3JjaC9jYXRhbG9nLXNlYXJjaC5kbz9xdWVyeVRleHQ9QXNpYS1QYWNpZmljK1ZpZGVvK0**

[2] Frost and Sullivan, "World Visual Collaboration Managed Services Market," November 2009

Motivators

Why is there such a strong interest in HD videoconferencing? Financial, environmental and technological factors are all drivers in HD videoconferencing emerging as a prominent player in business communications. These motivational factors include:

- **Cost Control:** Organizations are looking to rein in discretionary costs, with travel being a prime target. The cost of business travel continues to rise, with U.S. business travel averaging $349.60 per day[3].
- **Environmental:** Green initiatives are pushing businesses to reduce travel. In August 2009, the State of Virginia sponsored a Telework Day, which reduced pollutants by an estimated 82.8 tons[4]. This is also an economic issue, with the European Union discussing taxing carbon emissions on EU airspace[5].
- **Business continuity:** Iceland's Eyjafjallajokull glacier affected travel for over a month and is estimated to have cost companies billions of dollars, highlighting the need for an alternatives to business travel. Even less severe events can impact productivity, with the 2010 snowstorms in Washington DC causing an estimated productive loss of $71 million[6]. The fear of pandemics such as H1N1 influenza can also disrupt business as usual.
- **Employee satisfaction:** Increasing the productivity and effectiveness of video meetings drives teleworking, an important mechanism for attracting and retaining employees. This is most critical in key segments such as call centers, but teleworking is playing an increasing role. Eliminating commute costs also saves full-time teleworkers an average of $2000 per year[7].
- **Globalization:** Providing high quality communications simplifies working with third party organizations and remote locations.

Challenges

While the benefits are well understood, justifying the investment in the new technology can be challenging for many organizations. HD videoconferencing requires substantial bandwidth, yet at the same time introduces strict requirements for service availability.

[3] *Business Travel News*, as reported in Denver Business Journal, **denver.bizjournals.com/denver/stories/2010/04/05/daily16.html**

[4] Telework EXCHANGE, *Eliminating Gridlock,* **teleworkexchange.com/teleworkday/**

[5] As reported in **chinadaily.com.cn/bizchina/2009-10/28/content_8863537.htm**

[6] **washingtonpost.com/wp-dyn/content/article/2010/04/14/AR2010041404825.html**

[7] **www.teleworkexchange.com/teleworkdayreport/Telework_Day_2009_FINAL.pdf** (see slides 3 and 6)

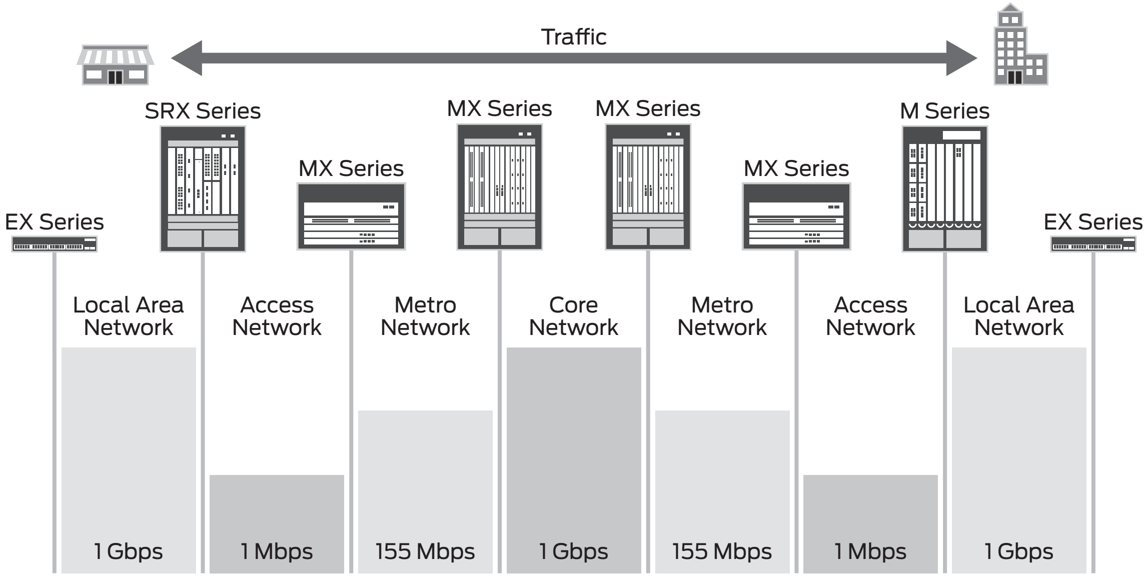

Figure 1.1 Example of bandwidth available in network domains

Broad videoconferencing adoption has the potential to overwhelm the network. IDC estimates that on average every $1 spent on videoconferencing requires roughly $3 in network upgrades[8] to eliminate the performance bottlenecks, which may limit the quality, bandwidth and number of videoconferencing connections.

Figure 1.1 illustrates the wide range of bandwidth used within a typical network. In each site, the local area network often has sufficient bandwidth to support all traffic, including videoconferencing. The network core, which is built on fiber links, is also unlikely to be the performance bottleneck. Even the metro network is often large enough to support numerous video sessions. Due to both cost considerations and lack of fiber availability, the access link that connects to each site is the most likely performance bottleneck. Specifically, since most lines are symmetric and a single click often results in a large response (a new web page, for example), it is the downstream connection to the customer site which is the most common choke point.

At the same time, videoconferencing traffic has very stringent requirements for low packet loss and latency. Because the data network was not designed to meet these requirements, network architects sometimes fall into the trap of building a separate, logical overlay network dedicated to videoconferencing. To support the required quality, the network is designed to support the peak demand. This solution is quite similar to the network silo architectures of the recent past, which is precisely the approach that most organizations shun. Figure 1.2 depicts this approach.

[8] *Worldwide Enterprise Videoconferencing and Telepresence 2010-2014 Forecast, IDC* (March 2010) (Document #221356) **idc.com/research/viewdocsynopsis.jsp?containerId=221356&sectionId=null&elementId=null&pageType=SYNOPSIS**

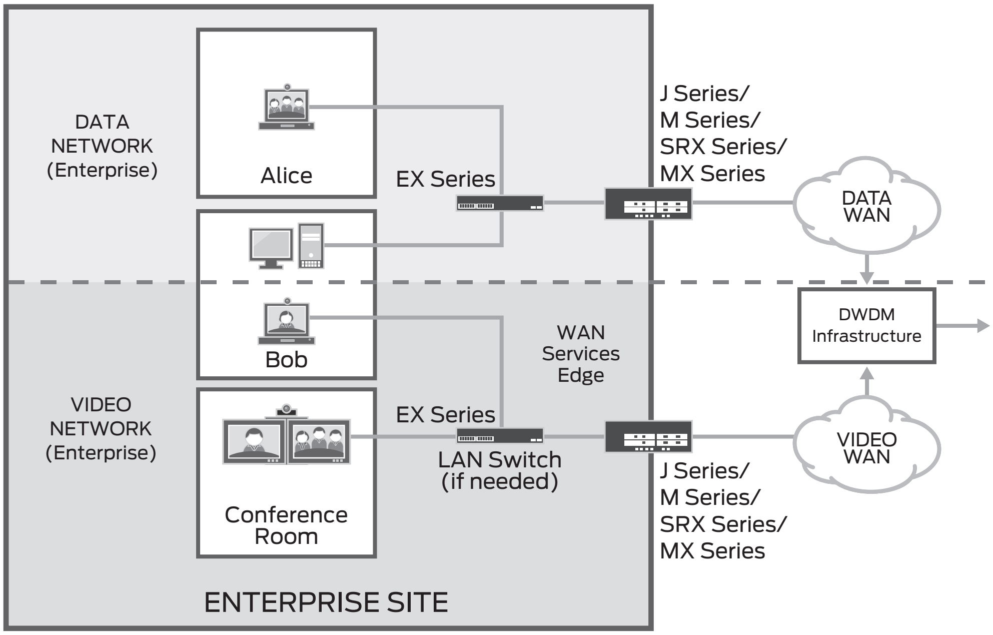

Figure 1.2 Enterprise using overlay network for videoconferencing

Justifying the overlay model is especially challenging since service usage and network utilization vary widely throughout the day and year. Daily usage typically peaks from 10-12 and 2-4 PM, with additional heavier usage during quarter-end and lowest usage during summer months. On one hand, a scalable network is required to support peak usage; on the other, videoconferencing systems and the underlying network are used relatively infrequently, making it challenging to justify the investment in an overlay network.

Figure 1.3 illustrates a simple example of this challenge. In this example, the network has been designed to support a certain amount of traffic. With the curve representing video traffic bandwidth demand, this network is sufficient at most times to meet the videoconferencing metrics. However, there are two challenges here. First, since video conferences are sporadic, the network has substantial unused bandwidth for most of the day, so the investment is not fully utilized. At the other extreme, several peak periods occur during the day, resulting in more video traffic entering the network than it can support. It takes only a few additional video conferences to impact the network from under utilization to excessive congestion, resulting in lost packets and poor quality of user experience. As video adoption increases, quality problems become more common, and users are less likely to view HD videoconferencing as a viable alternative to travel. In fact, adding “just one more” videoconferencing session could put result in more traffic being put onto the network than it can handle. Since all videoconferencing traffic is the same priority, some packets will be lost from each session, resulting in all videoconferencing sessions—the pre-existing as well as newly added video conference—suffering from packet loss and poor quality.

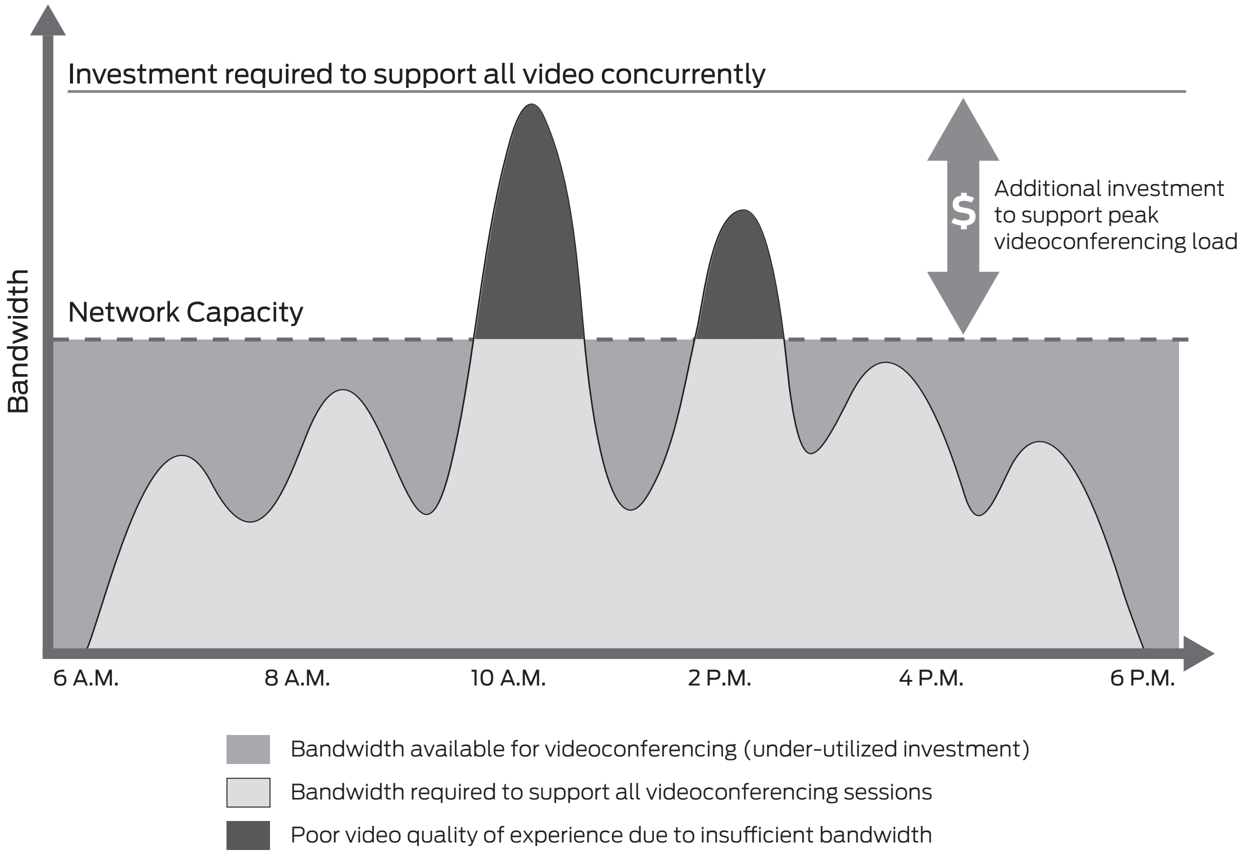

Figure 1.3 Videoconferencing investment challenge

Another challenge is service availability and reliability. If the goal is to replace costly travel with less expensive HD videoconferencing service, then the service must be highly available. This means purchasing redundant common equipment for the videoconferencing service, as well as ensuring node resiliency and redundant links. Without confidence that the network will be available, potential users will instead fall back on the traditional mechanism of heading for the airport.

Managed Videoconferencing Service

Supporting videoconferencing means adding personnel to support a range of functions from call scheduling to network troubleshooting. This increases labor costs during challenging as well as prosperous economic times, further weakening the business case. While these skills are all necessary to support the videoconferencing systems, there are not enough systems and conferences to require full-time support personnel.

One way to mitigate this issue is to outsource the management of the videoconferencing systems. In this model, the enterprise purchases the video equipment, but a third party managed videoconferencing application service provider (ASP) is responsible for the day-to-day operation of the videoconferencing system. This managed ASP can support multiple customers with a small pool of highly specialized personnel, allocating the cost of these specialists across several customers.

Employing a managed ASP reduces the people cost associated with deploying videoconferencing, as well as some related expenses such as scheduling systems. However, it does not reduce the network investment or resolve the concerns associated with service reliability.

Hosted Videoconferencing Service

In an attempt to bypass the ROI concerns, some enterprises outsource their videoconferencing capabilities to third-party hosted videoconferencing application service providers (hosted ASP) that specialize in this technology. In this hosted services model, the ASP is responsible for providing the videoconferencing service. Common equipment infrastructure, including videoconferencing bridges and servers, are housed in the service provider's premises. The hosted ASP leases circuits from the local network service provider (NSP) to interconnect customer sites back to their data center, typically getting lines at a significant discount. All videoconferencing traffic is carried on this network, which is separate from the data network of their customers that is used to interconnect PCs, clients and servers.

To the enterprise, one of the most attractive benefits of hosted video services is that it can be delivered with lower up-front costs. However, the capex investment is replaced by an on-going operational expense. While not as burdensome initially, the basic economic model—paying for bandwidth even when it is not being used—continues to make it difficult to justify paying for a videoconferencing service. Equally important, videoconferencing service must be delivered using an overlay network. This model does not allow desktop devices on the enterprise LAN, such as PC soft clients and videophones, to communicate with systems on the videoconferencing network. With this requirement hovering on the horizon, most organizations are looking to provide "desktop to boardroom" videoconferencing, allowing any video system to connect to any other system. This cannot be done when the desktop and room systems are on different networks. For example, in Figure 1.4 Alice cannot use her videoconferencing client to communicate with either Bob—who has a dedicated video system—or the shared conferencing room systems in this or any other site.

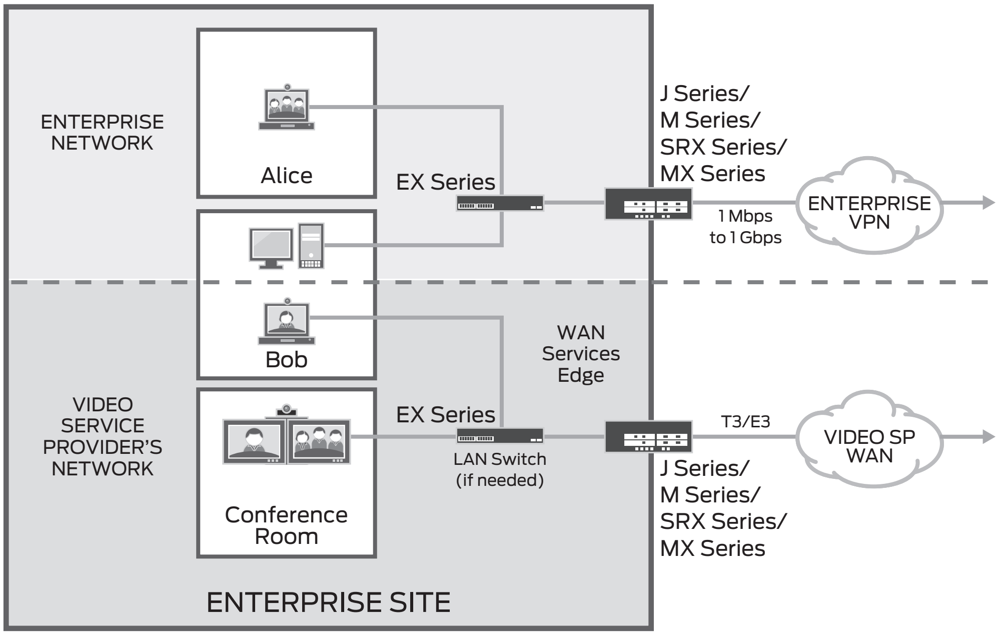

Figure 1.4 Site connectivity using video service provider

Network Service Provider Opportunity

The above challenges provide a unique opportunity for network service providers to move up the value chain and offer an integrated videoconferencing and VPN service that supports the full range of HD videoconferencing systems. For enterprises, this approach provides lower overall cost of the offering and a single contact point for customer care and service activations and upgrades. For NSPs, the benefit is in increasing their overall revenue by offering additional services. In addition, offering multiple services to a single customer reduces customer churn rates.

The unique value of NSP is the ability to support videoconferencing on a converged network, which also supports data and VoIP. Installing, upgrading and supporting one network is less costly than building a dedicated videoconferencing network. This approach also improves the competitive position of the SPs that can offer service at a lower cost. As importantly, this approach allows “desktop to boardroom” videoconferencing. Figure 1.5 illustrates that Alice is now able to establish video sessions with peers and conferencing room systems in this and other sites.

Delivering videoconferencing services across the same backbone and access link as other services also allows the NSP to bundle together the services, simplifying billing while providing a discounted price for multiple services.

The final, and perhaps most compelling, advantage provided by the NSP is its carrier-grade reliability. Since the service provider already has a high-performance, resilient core network supporting VPN services, the enterprise can feel confident that the service will be available whenever it is needed.

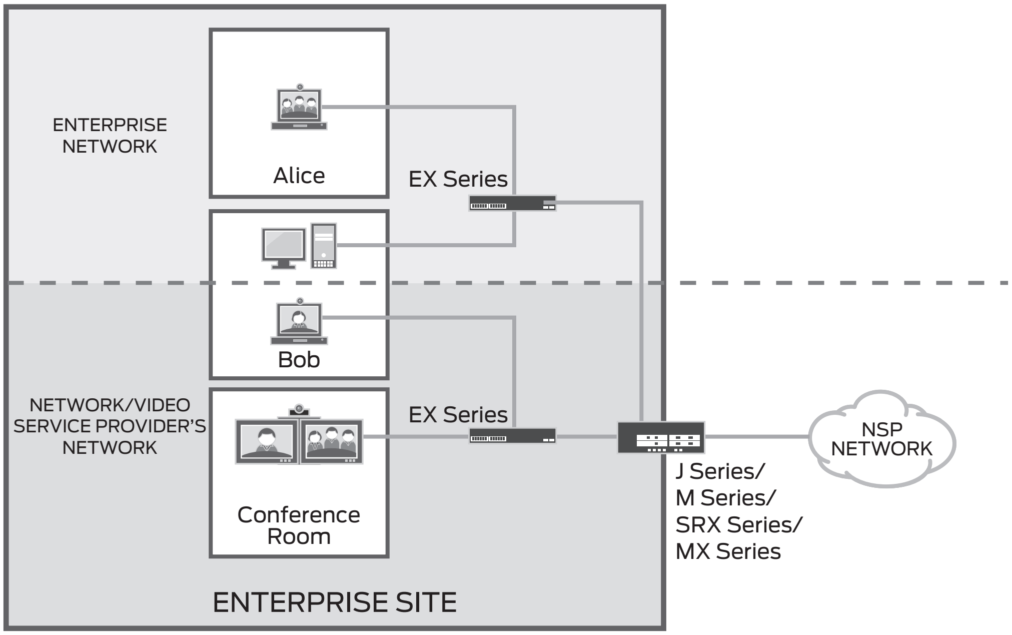

Figure 1.5 Site connectivity using converged NSP network

Ensuring high quality service delivery on a converged network requires linking the network with the conferencing application to provide assured forwarding of videoconferencing traffic. Juniper and Polycom have partnered to provide this linkage, which ensures that there is sufficient bandwidth available for each videoconferencing system. While every scenario is different, baseline comparisons show a 25 to 50 percent savings over competing solutions.

This joint capability is applicable to enterprises using their own network as well as for service providers looking to offer videoconferencing service to their customers.

Summary

Rising costs and modern world challenges are driving the need for HD videoconferencing, while higher performance audio, video and network technologies are making this a realistic alternative to traveling. There are many ways to support videoconferencing. Enterprises may deploy it themselves, videoconferencing specialists can offer the service or it can be integrated with an existing network service offering.

Offering a hosted videoconferencing service on a converged network provides many benefits such as:

- Increased revenue by moving up the food chain and by “selling bandwidth” to multiple customers
- Assured service quality even as networks become congested due to videoconferencing adoption
- Lower capex by sharing the network with other services
- Lower operational costs using a single converged network and operational infrastructure for video, VoIP and data
- Any-to-any conferencing by supporting dedicated video systems and desktop systems on a single network

Network service providers can take advantage of the emerging HD videoconferencing opportunity to leverage their network infrastructure by offering a most cost-effective and versatile videoconferencing solution. Maximizing profitability requires the ability to assure delivery of videoconferencing traffic to ensure overall service quality.

Chapter 2

Service Quality Considerations

There are several challenges to supporting large-scale videoconferencing. First and most obvious, videoconferencing requires tremendous amounts of bandwidth. Videoconferencing has the potential to become the killer application for networks, with requirements for a single session ranging from a few hundred Kbps for desktop systems to 20 Mbps for fully immersive room systems. (See *Appendix A: Videoconferencing Technology Primer on page 156.*) Second, videoconferencing demands real-time communications. High bandwidth by itself is not a critical issue. For example, it is possible to download large documents, producing bandwidth bursts, which result in packet loss and associated retransmission. Real-time does not have this luxury; packets must transfer the first time around successfully. Of course, modern IP networks support real-time VoIP, but this is possible explicitly because VoIP is such a small percentage of overall traffic. Adding videoconferencing turns this requirement on its head, demanding real-time performance for potentially a large percentage of traffic. Third, the network must scale to support both the required bandwidth and real-time responsiveness for thousands of sessions. Last but certainly not least, all of this happens as a service that must be monitored, measured and validated. An application-level service requires a different set of metrics and evaluation criteria to determine whether the service meets the end user's needs when compared to a networking service.

HD Videoconferencing Service Considerations

The end user's satisfaction with the videoconferencing capabilities is defined as service delivery quality or Quality of Experience (QoE). This high-level term does not have any technical meaning. However, the user's organization and the videoconferencing provider (which could be a third party provider or even the IT organization of the same company) will attempt to determine measurable objectives that are reasonably expected to result in acceptable service quality.

Service quality should not be confused with quality of service (often referred as QoS), which describes the capability of the network to differentiate between traffic types and treat them differently, some better than the others. However, QoS is a critical component when supporting HD videoconferencing.

Videoconferencing providers and their customers often use a Service Level Agreement (SLA) to ensure acceptable service quality. This is a contractual relationship between the end-user and the provider of the videoconferencing service—which may be a third party service provider or even the organization's own IT department. SLAs include well-defined, measurable metrics known as Key Performance Indicators (KPIs). The KPI is an exact measure that describes one aspect of service performance. Typical KPIs used to measure success for video service include:

- Percentage of video conferences that were established properly without requiring technical assistance
- Percentage of calls that remained up for the duration of the session
- Video quality, as perceived by the user, which is determined by a survey provided at the end of the conference
- The number of hours per week that videoconferencing systems were not available due to planned or unplanned upgrades

Whether using a third party service or delivered over private network, the videoconferencing service is supported by at least two different organizations: the videoconferencing applications group and the networking group. The applications group defines a service as an expected process that occurs when a packet reaches the application client or server. KPIs applicable to the videoconferencing application itself may include the:

- Number of video conferences which can be supported concurrently by the videoconferencing bridge
- Time required to process a single request for a new videoconference
- The number of call requests per second which can be processed before seeing significant delays (for example, more than 5 seconds) in call setup time or dropping the requests

In contrast, the networking group defines service quality as ensuring that all packets are delivered within a specified time. To them, timely packet arrival to the destination host equals quality. If packet loss and arrival times are within specified levels, then they consider their task accomplished. Common network metrics include:

- **Throughput:** The amount of traffic, measured in packets per second that a customer can continually put onto the network. This affects the number and type of videoconferencing systems that can be used at each location.
- **Packet loss:** The number or percentage of packets that are dropped by the network after initially being allowed into the network. Even a very low packet loss can significantly affect the quality of a videoconferencing session.
- **Latency:** The length of time for packets to traverse the network. Since videoconferencing is interactive (real time), the network must provide a very low delay.

These metrics are commonly used in SLAs for a networking service but are not frequently explicitly included in SLAs for application-level services such as videoconferencing. Instead, they become internal metrics while help ensure service quality. Poor performance in these areas results in poor quality or dropped sessions, and those metrics are covered in the SLA.

Network and Application Interactions

The networking and applications groups each seek to protect themselves from misbehavior by the other side. On the application side, the application may be written to assume that the network is infinite, meaning that it can always deliver traffic. However, too many concurrent videoconferencing sessions can overburden some portion of the network, resulting in poor quality for all videoconferencing sessions using that portion of the network. In response, the networking team protects their network from any abnormal or unexpected application behavior by policing, shaping and conditioning incoming customer traffic. When more traffic exists than is allowed—for example, when there are too many concurrent videoconferences—the network drops some traffic to protect other traffic or customers. For data traffic such as web browsing, the dropped information is retransmitted, yielding a minor delay. For videoconferencing, the result may be poor video quality and violate the customer's SLA.

Due to these potential network issues, today's videoconferencing applications are designed to accommodate some level of packet loss and delay. Most systems have jitter buffers that accept video packets and play them out to the user at a consistent rate. This mitigates the impact of network latency and jitter, although it does not compensate for consistently high latency. If packet loss is detected, videoconferencing systems also may proactively reduce the quality of the videoconference connection, for example from 1080p to 480i, which may also violate the SLA.

Ensuring service quality when there are potential performance bottlenecks therefore requires that the application and network have a common understanding of the available bandwidth.

Factors Affecting Service Quality

Successfully offering an HD videoconferencing service depends on three key criteria. First, the service must be available whenever the SLA dictates that it will be. Second, the videoconferencing information properly delivered between destinations within an acceptable time. Finally yet importantly, the service must scale to support the intended level of use.

Note the emphasis on the overall service. Success is measured by a user's overall experience, not the individual technologies which comprise the service.

Service Availability

Service availability is the ability to provide service whenever and wherever the provider has stated that it will be available. The ability to establish a videoconferencing session and maintain the videoconference connection until the user ends the call is an example of service availability. In order for videoconferencing to replace air travel as the de facto standard for meetings, there must be an assurance that the service is available when needed.

Service availability characterizes the time during which the service performs as expected by the users. It is calculated as follows:

Service availability = service uptime / total duration x 100%

Conversely, service downtime is a length of time within a year during which the service might be unavailable due to routine maintenance operations or in outages. Modern service providers promise five or six "9s" of availability. For example, five nines is a way to pronounce 99.999 percent availability (which equates to, 0.001%

unavailability). This number translates into not more than 365 days * 24 hours * 60 min * 0.00001 = 5.2 minutes of downtime per year. Six nines is much more demanding and allows approximately 30 seconds of downtime per year. Table 2.1 shows common availability terms and the associated downtime per year.

Table 2.1 Availability and Equivalent Total Downtime Per Year

Availability	Downtime per Year (365 days)
99% (two nines)	3.65 days (5256 minutes)
99.9% (three nines)	8.77 hours (525.6 minutes)
99.99% (four nines)	52.56 minutes
99.999% (five nines)	5.256 minutes
99.9999% (six nines)	32 seconds (0.5256 minutes)

Note that service availability depends upon both the network and application-level equipment. The following scenarios have the same net result, which is that the user cannot initiate a new HD videoconferencing call:

- The call server is down
- The target device is not registered with the call server
- There is a network failure which prevents traffic from reaching the call server, even though it is available
- There is a network failure which prevents videoconferencing traffic from reaching the intended remote end system

Service availability is determined by the availability of both the individual nodes (servers and network equipment) and the network that connects the nodes.

- **Node availability:** The availability of each device, which the packet traverses potentially, affects overall availability. This includes equipment over which the video traffic flows, as well as videoconferencing equipment such as call servers and videoconferencing bridges. Available methods to improve node availability include:
  - Providing redundancy for all critical components including power supplies and common equipment
  - Providing redundancy within a single chassis, such as cross-slot link aggregation
  - Making multiple nodes appear to be part of a single node, such as provided by Juniper's virtual chassis technology.
  - Using virtual addresses, which allow any system that is defined as part of that virtual address to process the information. Virtual Router Redundancy Protocol (VRRP) is one example of this.
  - Non-Stop Routing (NSR), to minimize packet loss during the switchover

  - In-Service Software Upgrades (ISSU), which provide the ability to upgrade the software without taking it out of operation.
- **Network availability:** This extends the overall system to include communications links between nodes, including both the LAN and the WAN. The ability for the forwarding software to recover from an outage is an important part of this discussion.

  Increasing network availability focuses on two key areas: detecting that an outage has occurred and determining the new path to use. For real-time traffic, all pieces must occur quickly, typically in less than 1/10 of a second. For example:

  - When using IEEE 802.3ad/802.1AX link aggregation group (LAG), loss of light signal from the remote site can quickly be detected, and traffic can quickly be directed to an alternative fiber in the same LAG bundle.
  - Ethernet's Rapid Spanning Tree Protocol (RSTP) relies on loss of keep-alive packets to detect an outage. Upon detecting the failure, RSTP first must determine the new path to follow. While the exact time to recover from a network failure is dependent on network topology and size, RTSP typically requires several seconds to recover from a network failure. When RTSP is required, segments should be kept small to minimize reconvergence time.
  - IP reconvergence is also triggered by loss of keep-alive packets of the routing protocol being used, such as OSPF or IS-IS. While the routing protocol can determine the new route more quickly than RTSP, it will still produce a noticeable delay. In addition, Bidirectional Forwarding Detection (BFD) may be used to speed failure detection time.
  - Unlike RTSP or IP, MPLS's Fast Reroute (FRR) capability determines the backup path ahead of time, resulting in faster failover. Combining MPLS with BFD yields the fastest recovery time from a network outage, and high resiliency is one key reason why MPLS is used at the core of most large networks.

It is easy to assume that the installed network is already designed with the desired level of resiliency. However, broad adoption of HD videoconferencing can quickly strain the existing infrastructure, For example, while the network was designed with two links to ensure availability, videoconferencing traffic could consume the bandwidth of both links. In this case, a single link outage would disrupt the service for at least some users.

This is a simplistic example of a complex, network design issue. The key message is that the network (most likely) was not originally designed for bandwidth requirements, or perhaps even the availability targets, required to support high definition videoconferencing. There are many tools available for building a resilient, highly available system. While resilient network design is beyond the scope of this handbook, network designers should review the existing network design to ensure it meets the demanding reliability needs of HD videoconferencing.

For additional information on configuring resiliency, see *Data Center Network Connectivity with IBM* at **juniper.net/us/en/training/jnbooks/dchandbook.html**.

Packet Delivery

Of course, an important requirement is that each piece of video information must be delivered to the intended destination in an acceptable time frame. Not meeting this criterion, results in the inability to establish new connections or in poor video/audio issues. The most common quality issue is a frozen image, although macro blocking or smearing are also possible. There are three parameters of particular interest—packet loss, latency and jitter.

Packet Loss

Packet loss is the inability to deliver some packets to the intended destination. Videoconferencing systems are quite sensitive to packet loss. Unlike data applications, because videoconferencing uses UDP, it does not retransmit packets. This is because of the real-time nature of videoconferencing—by the time the packet loss is recognized and retransmitted, a disruption has already occurred. Re-sending the packet after-the-fact will not resolve the issue and only leads to network congestion.

Many videoconferencing systems require packet loss to be under 1 percent, with some immersive systems requiring less than 0.1 percent packet loss. Drops in one packet can result in a 1-3 second degradation of a portion of the video image.

Packet Loss

Packet loss can be caused by bandwidth or equipment bottlenecks:

- **Bandwidth bottlenecks:** Network links that cannot support the amount of traffic being offered. There may be hidden bottlenecks in the network, such as the access link or metro network. It is quite possible that the LAN will attempt to send more traffic than the WAN link can handle, resulting in dropped packets.

  Even if the network seems able to handle the load, there may be momentary overloads (microbursts) since video systems send more traffic when users move more quickly and therefore more information is placed onto the network.

  When planning for videoconferencing, special attention should be paid to the existence of older network segments such as ATM, Frame Relay or T1/E1 links, as these may introduce performance bottlenecks.

- **Equipment bottlenecks:** More subtly, the networking equipment (routers, switches and security appliances) may be unable to forward packets fast enough, dropping packets even though there appears to be sufficient WAN bandwidth. The ability to move packets onto the wire is dictated by two key parameters: processing and buffer memory. Large video packets place fewer burdens on the processing component than do smaller packets since there are fewer packets per second to place onto the wire, so this is typically not the concern. However, there may be insufficient buffers to store several video packets. The network planner should ensure that all equipment can completely fill all links concurrently.

  Specifically, older equipment designed to support these lower-speed protocols may have insufficient processing and memory to support videoconferencing traffic. Even if the equipment includes gigabit interfaces, for example, it may not be able to fully fill those circuits. It is particularly important to verify that the equipment can forward packets onto high-speed interfaces at or near line rate for all ports concurrently.

Packet loss is expressed as a percentage of packets that are not successfully delivered to their destination. For example, 1.2 percent packet loss means that it is acceptable to drop 12 packets per 1000 packets in the given traffic class. Networking hosts and special service quality measurement probes are able to measure and report the total number of packets lost, variations in packets lost per time interval and the total packets transmitted. More advanced equipment can measure consecutive packets lost and calculate lost packet percentages automatically without involving the performance management systems. Since packets are sequenced, some videoconferencing endpoints can also provide information about the number of missed packets.

Most Juniper products provide line-speed forwarding on all ports. This is a key advantage for Juniper since it reduces the number of circuits required compared to competitive products that can only fill 60-70 percent of each link.

Latency

Latency (or delay) is the amount of time it takes for a packet to travel from source to destination. Videoconferencing systems react very differently to latency than traditional data applications. In data systems, latency results in a delayed delivery, but the information ultimately reaches the destination. In videoconferencing and other real-time systems, traffic that arrives "too late" to recreate the original video or audio signal is discarded, yielding the same result as packet loss. End-to-end latency must typically be less than 150 milliseconds, as longer delays affect video quality.

There are many causes of latency including:

- **Propagation delay** is the amount of time that it takes a packet to travel between the devices and is governed by the laws of physics. You may already know that the speed of light (and other electromagnetic waves that we use for telecommunications) in the vacuum of space is almost 300,000 km/sec (186,000 miles/hour). In other physical media such as air, fiber optic cables and electrical wiring, the propagation speed of electromagnetic impulses varies significantly.

  Knowing the distance between each pair of endpoints and physical media in use, the propagation delay can be calculated. For example, a link traversing the United States from east to west measures approximately 4800 km (3000 miles). The speed of light on fiber is approximately two-thirds the speed in a vacuum or 200,000 km/second. Therefore, the propagation delay for this link is 4800/200000 = 24 milliseconds.

  Satellite links yield a very high propagation, and often unacceptable, delay. Satellites sit in orbit 35,786 km/22,236 miles above the earth. Transiting a single satellite hop therefore adds 35,786 / 300,000 = 0.125 seconds of delay in each direction, or roughly ¼ second to reach its destination since the packet must go up to the satellite and back down. This is unacceptable for most real-time traffic.

- **Serialization delay** is the time it takes to put a packet onto or remove it from the wire at the line rate specified by the physical media. This value is directly linked with the physical interface type between two network nodes. For example, the serialization delay of a 1500 byte (12000 bit) packet across an E1 (2.048 Mbps) link is 12000/2048000 = 5.9 milliseconds. Serialization delay occurs both for incoming and outgoing packets; the above example therefore results in 11.8 milliseconds of both the ingress and egress links are E1. At the other extreme, a 10 Gbps link yields a serialization delay of 1.2 microseconds.

Serialization delay is not a concern in modern gigabit networks but remains a challenge when there are network segments using slower technologies, such as ATM or Frame Relay. In addition, many DSL links introduce significant serialization delay. Videoconferencing network planners must factor in all the links in the network.

- **Processing delay** is the amount of time added in each device in the system. Some amount of processing delay is added by every device in the system that a packet traverses—clients, servers, control points and network elements. Processing delay varies widely. The client must take the user's input and create one or more packets to reflect the request, while network elements must make forwarding decisions about where to send this traffic. Another delay factor is the time to access other systems, such as a subscriber database. For this discussion, we include accessing other systems, such as subscriber databases, as part of processing delay.

  Videoconferencing systems are more sensitive to processing delays than many other systems. The endpoints compress the images and convert the content into IP packets. This process inevitably introduces processing delays. Often specialized hardware is used to perform this conversion, which is one reason why full screen, high definition videoconferencing is not ubiquitously available on PCs.

  Processing delays are not directly controllable, although some system-level tuning parameters may be available to minimize these delays. Older routers and switches may also have insufficient processing (or related memory) to support scalable videoconferencing. The most common way to deal with excessive processing delays is to upgrade the sluggish equipment.

  Juniper Networks routers and switches typically provide line-speed forwarding on all ports while adding minimal delay. For example, the MX Series routers typically add less than 20 microseconds of latency.

- **Queuing delay** is the time the packet stays in the node before being passed to the system that puts the packet onto the wire. The packet must wait its turn while other packets, often those that arrived just before this one, are placed onto the wire. The amount of queuing delay depends on the level of congestion of the network interface at that moment and constantly varies. Queuing delay is most difficult to predict because it depends on actual network load at that particular time, which is uncertain. For example, being stuck behind a single jumbo (8000 bytes) packet on an E1 line adds 31 milliseconds (64000 bits / 2.048 Mbps) of delay, while the 10 Gbps link yields a delay of 6.4 microseconds.

  On a 10 Gbps backbone, queuing latency is not an issue even across a multi-hop network. However, as traffic is funneled into smaller and smaller bandwidth links, queuing delays can become the largest part of the overall latency and affect service quality. Fortunately, Juniper's Quality of Service implementation allows critical real-time traffic to receive preferential treatment by moving it to the head of the queue. Juniper provides line-rate forwarding even with QoS enabled.

Jitter

Jitter is the variation in packet arrival rates. While this inconsistent latency was a concern with early VoIP and videoconferencing systems, modern implementations include jitter buffers that accept these packets and play them out at a smooth rate. However, excessively high jitter rates can result in poor picture quality, with high-end systems requiring jitter to be under 40 milliseconds. Of course, a high jitter rate means that some packets are experiencing significant latency, and may be a precursor to packets being dropped.

Latency and jitter may be measured on the videoconferencing endpoints and can be measured by the networking equipment with assistance from traffic probes that imitate the videoconferencing traffic. However, in this case they do not take into account the last network segment between the HD videoconferencing endpoints and the first network node or the time that this traffic spent at the videoconferencing endpoint. It is also useful to monitor jitter rates as they may indicate a creeping latency issue.

Service Scalability

Service scalability is the ability of the selected hardware and network architectures to support the growth of future users. Scalability can be compared to an elastic rubber band. It will perform as normal until a certain moment. If stretched just a bit further than its designed strength, its performance will degrade rapidly and consequently rupture. As service use increases network load, performance will begin to degrade as the bandwidth capacity reaches maximum capacity and as a result, performance will rapidly degrade during intense congestion. The net result is too unacceptably long latency or too many dropped packets.

The major elements of service scalability include:

- **Bandwidth scaling:** This is the most obvious requirement for scaling of videoconferencing. If network utilization is low at a particular time, the overall perceived quality might be perfect. However, as a number of simultaneous users increase, the quality of this service can deteriorate quickly as network congestion occurs.
- **Session scaling:** The service must support the required number of concurrent sessions. This is a limit primarily of the application servers.
- **Customer scaling:** If being deployed as a service offering, the service must support many different organizations concurrently. This is determined primarily by the network, which places each customer into a separate VPN. (See *Appendix B: VPN Overview on page 165.*) Typically, MPLS-based VPNs are used by service provider networks to ensure the privacy of each customer's traffic.

Preparing for HD Videoconferencing

To ensure and successfully support HD videoconferencing, the following steps are recommended:

- Assess the network's ability to provide the required reliability, bandwidth and scalability. The key question here is whether the existing network can support videoconferencing. Of course, adding a small number of systems will most likely not significantly impact the network. However, if the goal is to get people off of airplanes, the number of required systems could have a dramatic impact on network traffic. A discussed above, this requires determining the bandwidth required for each link as well as the ability of installed equipment to support these upgrades.
- Enable QoS for important videoconferencing traffic at WAN endpoints and other potential bottlenecks. Since HD videoconferencing is sensitive to loss and delay, ensuring priority treatment for this real-time traffic is recommended for most networks. The network assessment must be performed assuming that QoS is enabled, as this function can significantly reduce the throughput of some network equipment.

  Juniper Networks Junos operating system provides an advanced QoS capability that can be deployed on routers, switches and security appliances, and are the key to reducing queuing delays. For further details concerning QoS, see *Chapter 9, Implementing Quality of Service on page 115.*
- Implement network-aware call admission control. Similar to a fast busy signal on the PSTN, this function ensures that network resources are available before allowing the connection. Without this feature, videoconferencing traffic can overwhelm the network during peak usage periods, forcing the current users (those already using the network as well any new sessions) to experience poor quality.

  This requires that the application and network work together to ensure service quality. A key requirement here is that the videoconference application must be able to understand whether the network can satisfactorily support the additional connection.
- Enable the network to dynamically add support for videoconferencing traffic by implementing dynamic policy-based prioritization at the edge of the core network.

  This is most important for service providers that are continually adding new customers and services. When a new videoconferencing session is allowed to join the network, dynamic prioritization updates the network without requiring manual reconfiguration, simplifying deployment while reducing the chance of errors.

Juniper Networks and Polycom have implemented these functions, which are collectively referred to as assured forwarding. Assured forwarding is an optional function but provides significant value by assuring service quality and reducing network investment. It is discussed in more detail in the next chapter.

Chapter 3

Introduction to Assured Forwarding

Part One

Assured forwarding ensures that once a videoconference begins sufficient bandwidth will be available throughout the call to ensure high quality service delivery. If this cannot be guaranteed, then the caller is asked to try again later. This allows the network planner to design the network to support a large percentage of requested video calls—say, 99 percent—without the large incremental investment required to support 100 percent of offered calls. The alternative is to allow all requests, potentially leading to network congestion. This in turn results in degraded service quality for all videoconferences.

Assured forwarding is not required to support high definition videoconferencing, but is recommended if there are potential bottlenecks in the network.   The benefits include:

- Reduced overall costs by offering all services on a single network
- Running the network more efficiently, since the available bandwidth can be utilized by multiple applications so there is less chance of wasted bandwidth
- Boardroom-to-desktop conferencing, since all video-enabled clients are on the same network

Figure 3.1 illustrates the process for initiating and establishing a successful call. The solid lines depict the call setup without assured forwarding, while the dotted lines represent the flows, which provide assured forwarding.

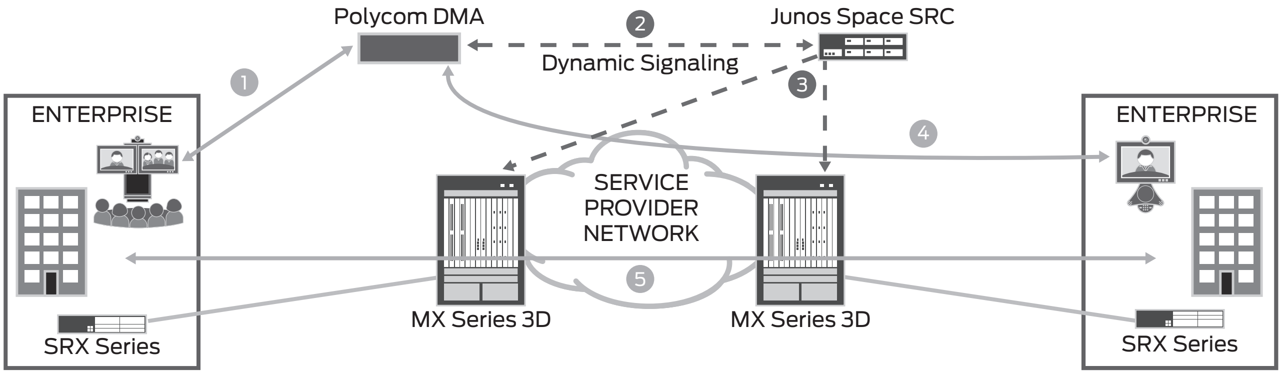

1 Initiating endpoint issues request to DMA to start session with remote system.
2 Polycom DMA forwards request to Junos Space SRC, which responds by approving or denying request (If denied, DMA notifies initiator).
3 SRC notifies MX Series (3D Universal Edge Router) to prioritize traffic.
4 DMA continues call setup process.
5 Call is successfully established.

Figure 3.1 Assured forwarding overview: CAC and dynamic QoS

The remainder of this traffic summarizes the key capabilities that comprise the assured forwarding functionality. Additional information is covered in *Chapter 10, Implementing Assured Forwarding on page 129.*

Call Admission Control (CAC)

The ability to determine whether new calls should be allowed onto the network is known as call admission control (CAC). This protects the network against exceptionally high concurrent video usage. Denying admission onto the network is the exception to the rule, not a frequent occurrence. The key assumption is that it is better to prevent a single call from occurring seldomly than to admit a call that will affect the quality of many calls.

CAC is a superset of the Application Admission Control capability provided by call control servers. This technology leaves the admission of application traffic entirely at the discretion of the application servers. System operators configure the upper level of bandwidth allowed for each session. This method can limit the number of simultaneous calls on their networks and is available today.

Despite its benefits, this method has one weakness – it is not aware of the true network state and does nothing to enforce its decisions outside its video-application realm. This in turn permits situations when some HD videoconferencing traffic may be unaccounted for and cause network congestion. For example, an application that is not using a SIP server may still call another endpoint thereby bypassing the server and consequently cause congestion.

In contrast, assured forwarding limits the number of calls or traffic flow sessions on the network. This is achieved by calculating bandwidth requirements for each application traffic flow before admitting it onto the network; verifying bandwidth availability at various potential network bottlenecks; and reserving bandwidth for all admitted traffic by maintaining the live database of current sessions. In addition, the network admission control process can dynamically deploy and enforce QoS policies on network equipment.

Assured forwarding provides the best of both worlds. The application domain is responsible for correctly identifying individual user needs of resources and subsequently signaling these requirements to the network domain. At the same time, the network domain is responsible for correctly adjusting network resources and notifying the application of resource availability.

It is worth noting that assured forwarding is not a new concept. The PSTN is not built to successfully complete every call. Instead, it is designed to allow a high percentage (well over 99.9 percent) of calls to be successfully completed—but not all calls. Ensuring 100 percent availability requires a significant incremental investment. Enterprise PBXs also provide this function.[1] Like those mechanisms, in a properly designed network calls will rarely be prevented from joining the network.

Assured forwarding's CAC function allows the organization to limit their network investment to support the desired service level.

[1] In North America, lack of network resources is signaled to the dialer by sending a fast busy signal, which consists of a repeating pattern of 0.25 seconds of tone and 0.25 seconds of silence. If the called party is using their phone, then the caller receives a standard busy signal consisting of a repeating pattern of 0.5 second tone and 0.5 seconds of silence.

Dynamic QoS

Dynamic QoS modifies the IP edge router (PE) configuration to prioritize video traffic whenever a new call is allowed onto the network. It avoids the need to manually modify each PE, reducing errors and speeding deployment time.

Dynamic QoS is especially valuable in service provider networks, where adding new customers would require manually updating the network.

Service Tiers

HD videoconferencing service has high demands should be defined as a separate traffic class which can be prioritized to minimize packet loss and jitter. The Juniper/ Polycom assured forwarding solution supports three service tiers for videoconferencing. Figure 3.2 shows one example how different service tiers can be supported.

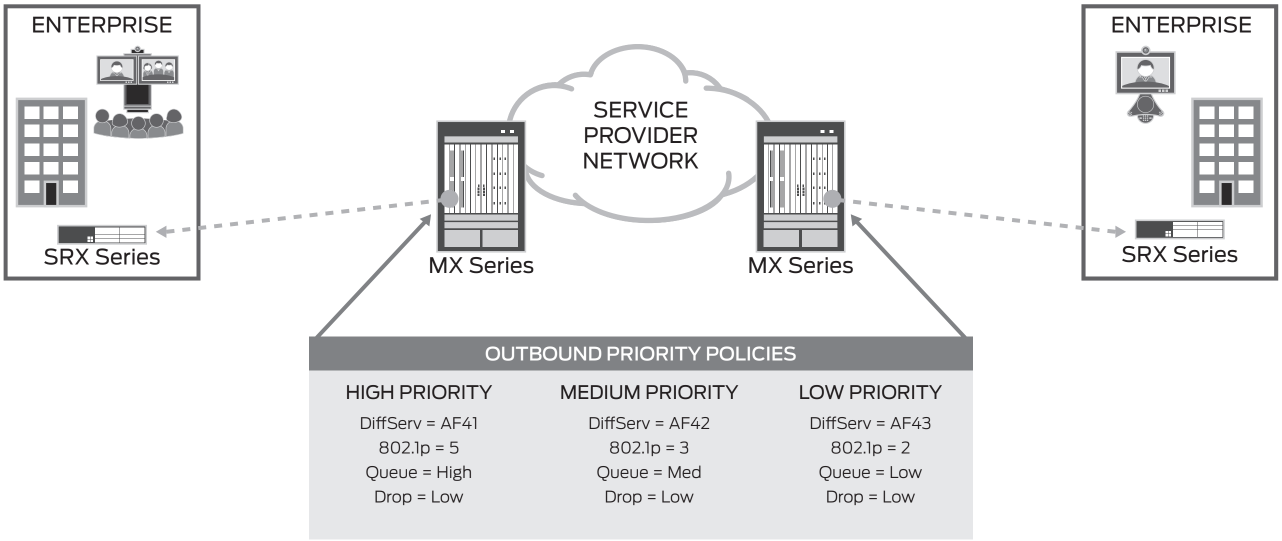

HIGH PRIORITY	MEDIUM PRIORITY	LOW PRIORITY
DiffServ = AF41	DiffServ = AF42	DiffServ = AF43
802.1p = 5	802.1p = 3	802.1p = 2
Queue = High	Queue = Med	Queue = Low
Drop = Low	Drop = Low	Drop = Low

Figure 3.2 Service tiers

The network administrator determines the implementation details of each service tier, including how much bandwidth should be reserved for video traffic on each link for each service tier. Figure 3-2 shows one possible mapping. While policies may be pushed to various points in the network, it is the downstream connection from the core to the customer site which is the most likely bottleneck and therefore the place to implement these policies. This information is configured in Juniper Networks Junos Space Session and Resource Control (SRC), allowing it to manage bandwidth on each network segment. Examples of tier assignments include:

- **Device type:** Immersive systems are given highest priority, room systems are medium priority and personal (desktop) are assigned to the lowest video priority.
- **Equipment usage:** Demo systems are the highest priority, followed by systems assigned to executives, with all other systems having the lowest video priority.

Additional Capabilities

In addition to call admission control, dynamic QoS and service tiers, the joint solution provides several unique capabilities to address this:

- **Bandwidth limiting:** The network can limit the amount of bandwidth that can be consumed by videoconferences.
- **Multiple congestion points:** The network can be designed to look at several congestion points—for example, the bandwidth available to each location—before deciding whether to allow the conference.
- **Varying bandwidth requirements:** The solution knows the real bandwidth required for each connection, instead of assuming that each call requires the same amount of bandwidth.

The solution supports SIP as well as H.323, and supports directly connected video systems as well as connections to a videoconferencing bridge.

Solution Elements

Juniper and Polycom are working together to provide assured forwarding for videoconferencing traffic. The key solution elements are:

- **Polycom Distributed Media Application (DMA):** The DMA serves as the H.323 Gatekeeper/SIP Agent control point for establishing video calls. It receives call setup requests and forwards them to Juniper Networks SRC for approval. The DMA supports any endpoints which use H.323 or SIP for call control. DMA Release 3 is required.
- **Junos Space SRC:** The SRC maintains a network topology and determines whether the request can be supported. If so, it updates its bandwidth usage tables and pushes any appropriate policies—in this case to prioritize HD videoconferencing traffic—to the network elements. If the request cannot be supported, the requester is notified by the DMA. SRC release 4 is required.
- **Juniper Networks MX Series 3D Universal Edge Routers:**The MX Series provides line-rate QoS and can accept policies from the SRC to assure service quality.

These capabilities can be added into an existing Juniper Networks infrastructure supporting HD Videoconferencing .

Why Polycom?

Juniper has partnered with Polycom to develop assured forwarding for videoconferencing. Polycom is the market leader in installed videoconferencing systems and offers the widest range of videoconferencing products, including fully immersive systems. Their products offer superb quality, including industry-leading encoding which reduces the bandwidth required. Juniper and Polycom also share a common vision of open, standards-based networking.

The first joint assured forwarding capability was demonstrated in 2008. In early 2010, Juniper and Polycom announced an alliance focused on improving the reliability, cost-effectiveness and quality of the customer experience for HD videoconferencing services. What resulted is an end-to-end solution that combines Juniper's dynamic policy management system and Junos®-powered networking and security platforms with Polycom's videoconferencing infrastructure and HD videoconferencing solutions. Customers can now intelligently use the full potential of their infrastructure to deliver assured quality for video communication. According to Ira Weinstein, senior analyst and partner, Wainhouse Research, "Juniper and Polycom are leveraging their respective strengths in carrier class network architectures and enterprise HD videoconferencing and videoconferencing to change how video traffic runs on the network among other applications".

The bottom line is that this partnership creates a network capable of dynamically managing and allocating available resources in real-time—enabling service providers to optimize existing resources to deliver cost-effective HD videoconferencing without compromising the quality of other applications.

Summary

Assured forwarding plays a valuable role in allowing customers to reduce the costs associated with supporting videoconferencing on their network. Assured forwarding provides the most value in very large networks, including service providers who offer videoconferencing services and Fortune 100 enterprises; however, it can be deployed even by medium-sized organizations looking to control costs. It may be used whenever there is a potential bottleneck in the network; most frequently, that bottleneck is the access link from the core to the enterprise site. It can be implemented initially, or can be added into the network when the existing bandwidth starts becoming depleted.

Chapter 4

Network Design Overview

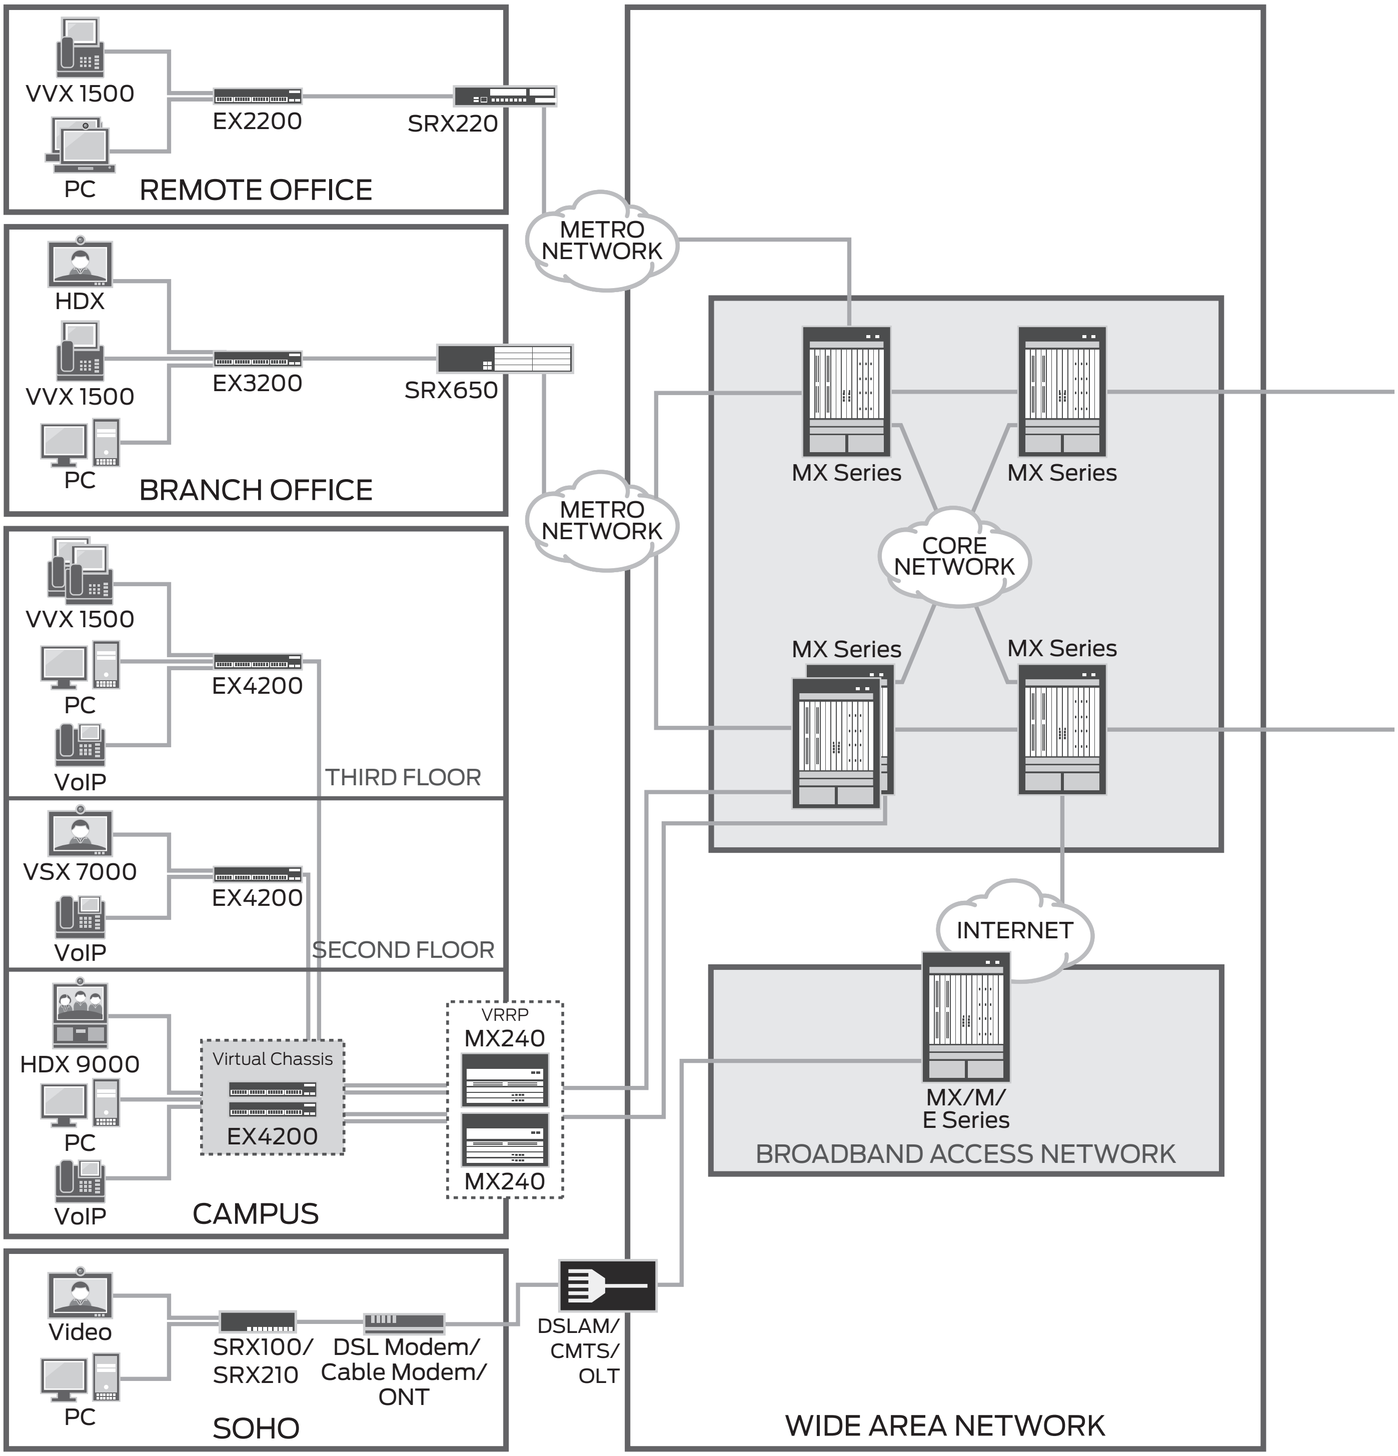

Figure 4.1 Network architecture overview

This chapter provides an overview of the network architecture described in the remainder of this handbook. Figure 4.1 provides an overview of the complete network.

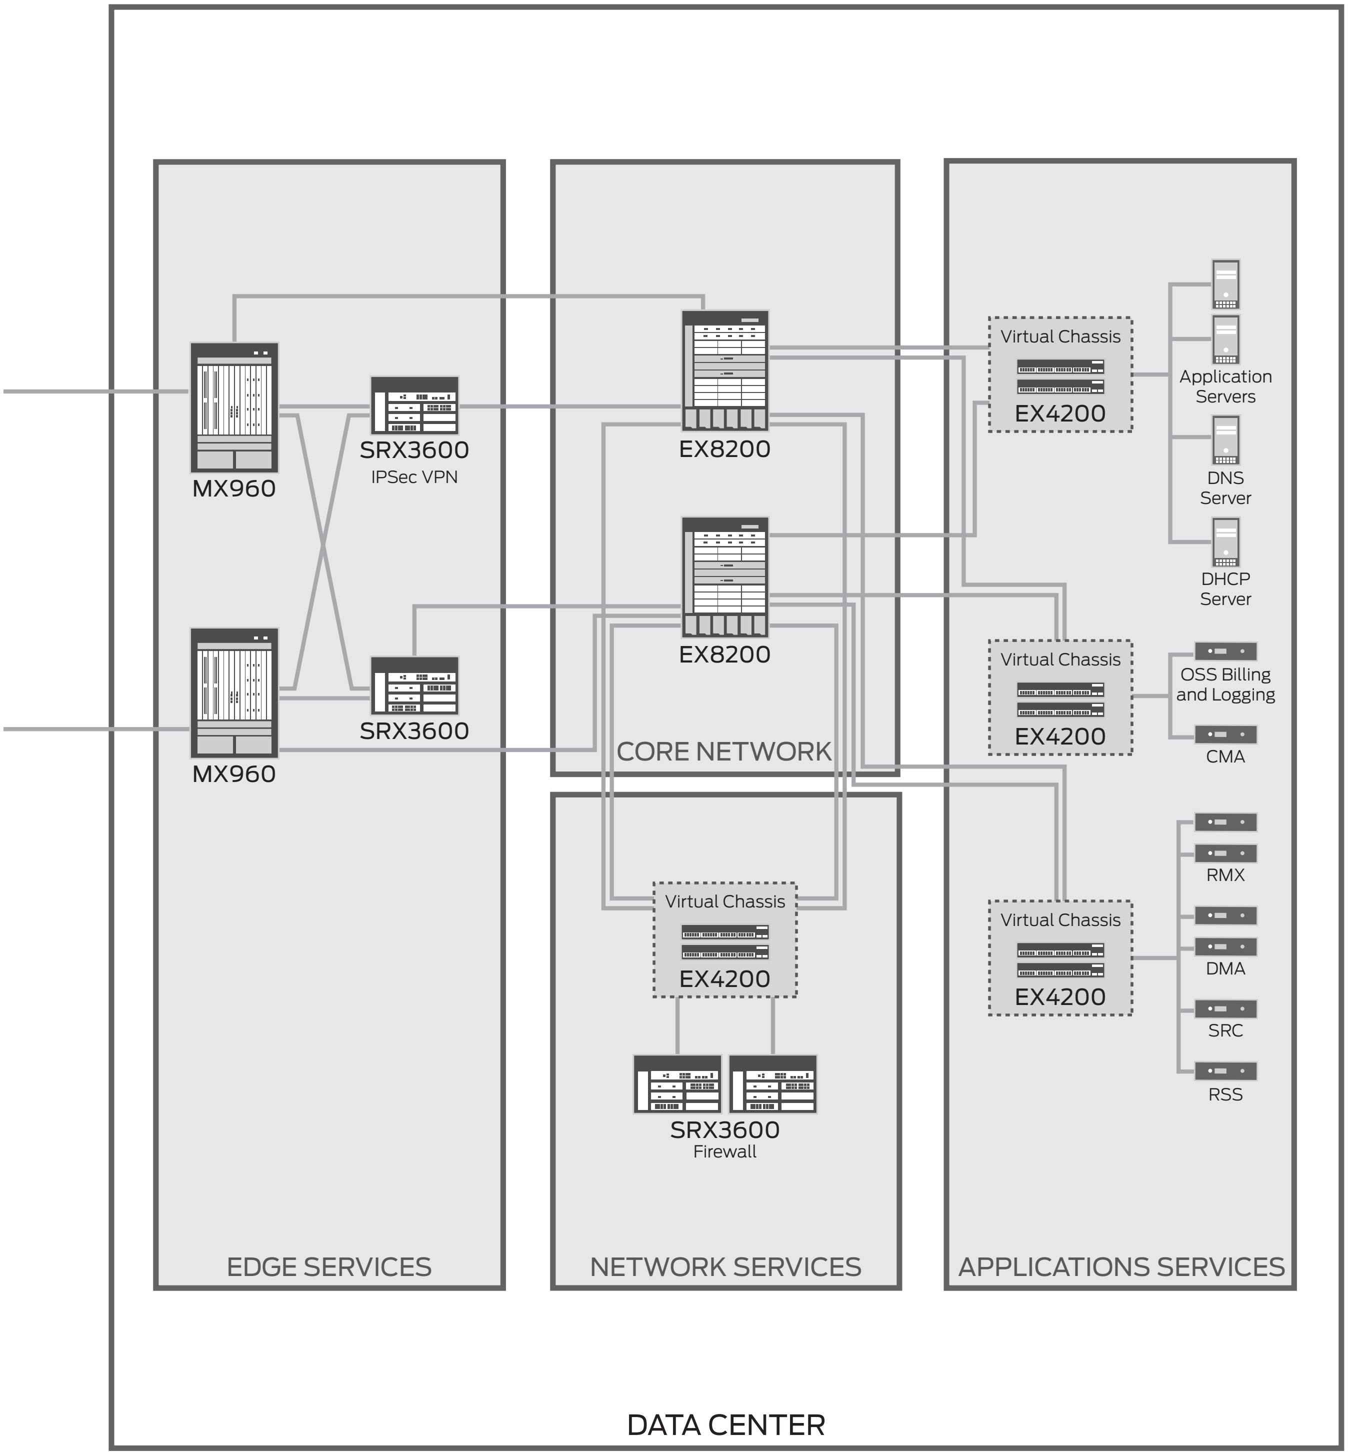

Figure 4.1 Network architecture overview (continued)

Network Domains

Core Network and IP Edge Routers

The shared backbone that interconnects all sites for all customers is the core network, which uses MPLS/BGP-based Layer 3 VPNs to ensure privacy for each customer. The key network element in the core is the Provider (P) router. These routers are responsible for moving packets across the network quickly; they are unaware of which packet belongs to which customer. Rather, the source and destination of each packet is determined by the IP Edge router, which is discussed below.

P routers are typically Juniper Networks T Series, MX Series or M Series routers. Because the P routers have no VPN awareness or functionality, they are not discussed further. If there is a potential core bottleneck, they may implement the QoS mechanism described in this handbook.

Located at the edge of the core network are MX Ethernet Service Routers. In the classical mesh VPN model, each MX is a Provider Edge (PE) device, which is defined as the device at the edge of the service provider network that interfaces with the customer. PE routers are in turn interconnected by P routers. Modern networks sometimes interconnect devices in a ring rather than a mesh, resulting in a single device serving both roles concurrently. The included network follows this latter approach. To avoid this confusion, each MX is referred to in this handbook as an IP Edge router. See *Chapter 5, Building the IP Edge Network on page 37* for further details.

Metro Network

While the customer site may be connected directly to the PE, there may also be an intervening aggregation network such as a metro Ethernet network. Often controlled by a different organization, this network must be transparent to the information flowing through it.

Two illustrative metro scenarios that have been used during this testing are VPLS using MX routers and an EX-based metro ring. In both cases, the metro network connects to both the core and the sites using standard Ethernet VLANs. Configuration details are not provided for these networks since they are not the focus of this document[1].

Data Center

Common equipment resides in a data center. The data center consists of a classic three-tier hierarchy consisting of WAN gateways (MX Series), aggregation (EX8200) and access (EX4200) switches The data center design provided in this handbook extends previously documented designs[2] by incorporating assured HD videoconferencing service.

[1] For VPLS, see **juniper.net/techpubs/en_US/junos10.2/information-products/pathway-pages/solutions/virtual-private-lan-services/index.html#configuration** .For the EX-based metro ring, see **juniper.net/us/en/local/pdf/implementation-guides/8010045-en.pdf** .

[2] See for example Implementing L2 at the Data Center Access Layer on Juniper Networks Infrastructure, **juniper.net/us/en/local/pdf/implementation-guides/8010014-en.pdf**; Implementing L3 at the Data Center Access Layer on Juniper Networks Infrastructure, **juniper.net/us/en/local/pdf/implementation-guides/8010022-en.pdf**; and Data Center Network Connectivity with IBM Servers, available at **juniper.net/us/en/training/jnbooks/dchandbook.html**

Critical sites such as data centers are directly connected to the core network. The MPLS/BGP core is extended directly to these sites. In other words, there is an IP Edge router located at the data center.

The data center setup includes equipment for call setup (Polycom DMA) and dynamically updating network policies (JuniperSRC). It also includes a videoconferencing bridge (Polycom RMX) located in the data center as well as a Polycom CMA for management. Managed service providers or enterprises may choose to put one or more RMXs at other locations. See *Chapter 6, Building the HD Videoconferencing Data Center on page 55* for further details.

Networked Sites

This network includes several types of sites. Many devices can be used as the WAN gateway router for these sites. These devices include Juniper Networks SRX Series, M Series, MX Series, EX Series or J Series as well as third party routers. For physical connectivity, there are three key major design considerations

- **Connectivity:** Whether the site is directly connected to the core MPLS network or connected using an IP/Ethernet link. In the latter case, it may be a dedicated connection to a physical backbone port, or numerous sites may be aggregated together through an aggregation or metro network. In this case, a single physical core port may connect to the various sites using different VLANs.
- **WAN links:** There can be a single WAN link from the enterprise site to the core, multiple parallel links (Ethernet LAG) to protect against a single link failure or links, which are homed to different sites to protect against a core node failure. Finally, connectivity can be through the Internet rather than across a VPN.
- **WAN gateway resiliency:** There may be a single WAN gateway router or redundant WAN gateway routers.

The network provided in this handbook offers a mixture of these options, as follows:

- **Branch office:** This type of location has a single WAN gateway, which may have a single upstream connection, or may be dual-homed. Smaller sites such as remote sales offices may connect to the core through an intermediate metro network. Many sites may be aggregated together in the metro network, with a single connection into the core. Since this segment may be provided by a different provider than the core, videoconferencing traffic must flow transparently through it. See *Chapter 8, Building the Enterprise Site on page 93* for further details.

  For an HD videoconferencing service, those sites may belong to different customers, with each customer represented as a separate VLAN. Enhanced (hierarchical) queuing capabilities must be used on the IP Edge router to implement QoS separately for each customer on the same port.
- **Campus:** Larger sites such as a campus include a redundant LAN backbone and deploy redundant WAN gateway routers that connect different backbone routers. See *Chapter 8, Building the Enterprise Site on page 93* for further details.

- **SOHO:** This site is assumed a residence or small site connected to the Internet using standard broadband technology such as DSL or cable. While the Internet does not enforce QoS, network elements at the edge can still prioritize traffic appropriately to provide the best possible experience. See *Chapter 7, Small Office/Home Office Reference Architecture on page 79* for further details.

Implementation Considerations

This remainder of this handbook describes a network supporting videoconferencing as well as data and VoIP. It includes one large campus site, one SOHO site and several branch offices. The design assumes that a network service provider is supplying a managed VPN service as well as a hosted videoconferencing service across a converged infrastructure. However, the content can also be used by an enterprise seeking to deploy videoconferencing themselves.

When building a videoconferencing-enabled network, there is a wide range of choices to consider. In addition to high availability, QoS and assured forwarding capabilities (as previously discussed in Chapters 2 and 3), this design implements the following practices and components.

VLAN Separation of Applications

Traffic from different applications is typically carried in separate VLANs. This common technique is used in the included design. In addition, each videoconferencing system and associated traffic can be assigned to one of three categories to distinguish the importance. For example, the highest priority may be assigned to immersive systems and/or systems frequently used by top management. At the other extreme, most desktop systems could be assigned to the lowest videoconferencing priority level.

This handbook uses the following VLANs within the enterprise sites:

- **Data:** VLAN 1001
- **VoIP:** VLAN 1002
- **Videoconferencing:** VLANs 100 (highest priority), 200 (medium) and 300 (lowest priority).

All signaling is carried on the same VLAN as the associated media traffic.

Encryption

Most videoconferencing endpoints can encrypt traffic to protect against snooping. This capability is enabled on our network. In addition, this encrypted traffic is passed through IPSec tunnels to ensure that doubly encrypting traffic does not add unacceptable delay. This represents a possible implementation since the video and remote access equipment may be controlled by different groups.

Address Management

Like any other device, videoconferencing endpoints can have IP addresses assigned manually or from a DHCP server. If a DHCP server is used, there are three different ways that IP addresses can be assigned. First, an IP address can be permanently reserved for a client such as an immersive system. This IP address assignment is typically based on the client's MAC address. Second, IP addresses can be assigned dynamically from an address pool used exclusively by videoconferencing systems, as is typically done for room systems. Third, IP addresses can be assigned to videoconferencing systems from the same address pool as used for PCs.

Another design question is where the DHCP server resides. Large sites typically have an on-site DHCP server to minimize WAN traffic, while smaller sites will communicate to a centralized DHCP server located in the data center. In this case, the on-site router will convert the DHCP request into a unicast request rather than a broadcast using *DHCP relay* functionality.

This network implements all of these techniques at various network sites, and all equipment performed as expected.

Security Zones

When connecting a private enterprise site to a VPN (or to the Internet), the issue of protecting each organization from the other becomes an important consideration. Adding videoconferencing equipment at the customer premise which is controlled by a videoconferencing service provider adds yet another level of security complexity. Since most attacks come from within the site, the three network sections must be protected from intentional or unintentional mischief while also having policies that enable PC-based soft clients to communicate with videoconferencing systems.

Security zone configurations are provided in the appropriate chapters.

Virtual Private Networks

A network service provider must of course implement VPNs to support multiple customers. The most common implementation is based on RFC 4364, *BGP/MPLS IP Virtual Private Networks (VPNs)*. That is the implementation used in this handbook for the core network.

This network supports two different VPNs, SmallCo and BigCo.

Maximum Transmission Unit

To ensure minimal drops and delays for videoconferencing traffic, architects must consider the entire network's Maximum Transmission Unit (MTU) and its video endpoints must be programmed accordingly. The end-to-end MTU assessment must account for Ethernet, IPSec, MPLS and VPN information.

Equipment

The following Juniper devices and specific software versions were used during validation testing.

Table 4.1 Juniper Devices and Software Versions used in Validation

Juniper Devices	Versions
MX Series 3D Universal Edge Routers	10.3R1.9 (with updates)
SRX Series Services Gateways EX Series Ethernet Switches MX Series 3D Universal Edge Routers	10.3R1.9
SRC	4.0.0R0 (with updates)
Steel-Belted Radius (SBR) Servers, Service Provider Edition	6.10

Chapter 5

Building the IP Edge Network

Part Two

This chapter describes the configuration of the IP edge network to provide high quality HD videoconferencing. It describes intra-enterprise deployment as well as multiple customer support. The IP edge device illustrated is a Juniper Networks MX960 Ethernet Services Router. Specifically, the configuration examples shown are for the router labeled as MX Series (West).Table 5.1 lists the router's physical interfaces.

Table 5.1 MX960 Router: Physical Interfaces

Remote Devices	Interface	Local IP Address	Remote IP Address
MX-North	XE-4/2/0.0	10.4.1.14/30	10.4.1.13/30
MX-South	XE-4/0/0.0	10.4.1.2/30	10.4.1.1/30
Campus-MX-A	GE-5/0/0.0	10.4.1.49/30	10.4.1.50/30

Figure 5.1 shows an example of managed WAN services and how the MX Series supports these services in the core network.

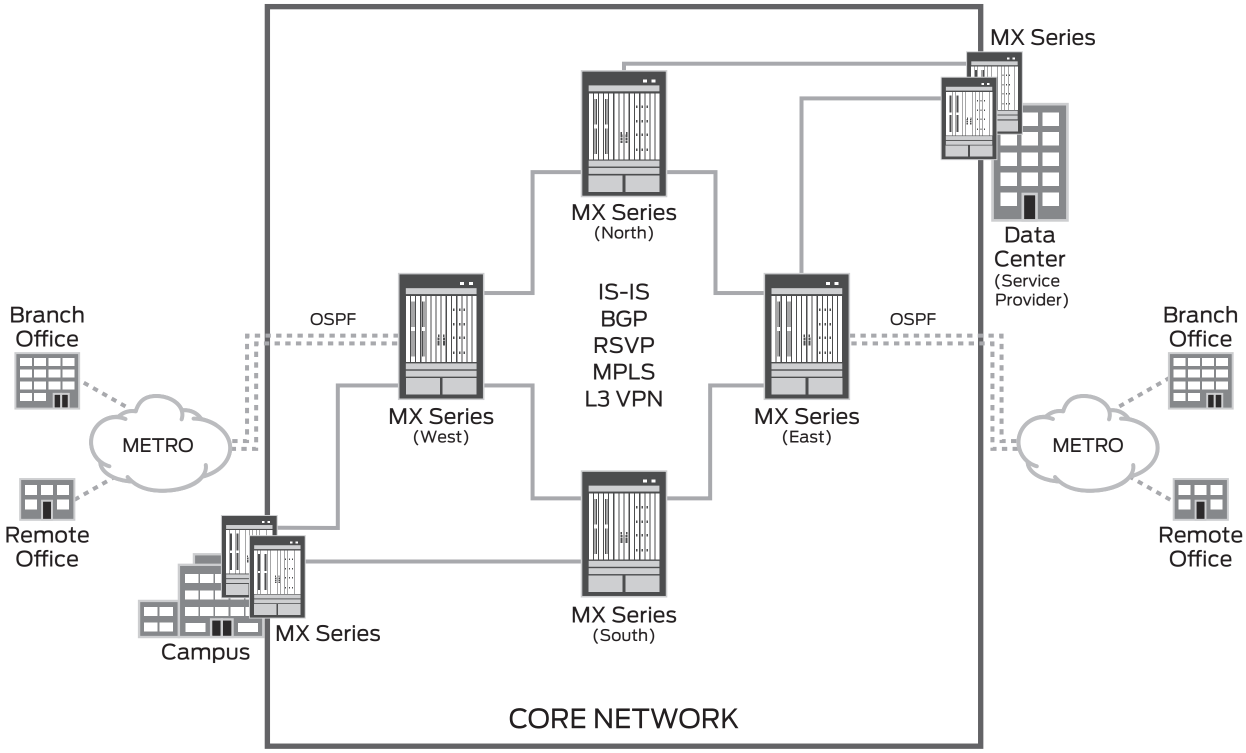

Figure 5.1 Managed WAN services in a core network

This implementation phase is built upon the existing network infrastructure, where we assume that the network administrator has preconfigured the baseline configuration that includes system level and chassis level configurations, routing engine redundancy and interface configuration for basic IP connectivity. To this, the following capabilities could be configured:

- **MPLS backbone:** The first step is to build the MPLS backbone that supports HD videoconferencing as well as other traffic. All traffic—video, VoIP and data—will use this converged backbone. If the network is supporting multiple customers or communities of interest, they too share this common MPLS infrastructure.
- **Layer 3 VPNs:** If supporting multiple customers, the next step is to define a unique VPN for each customer. A VPN separates the traffic from each customer, ensuring privacy. Figure 5.2 shows an overview of the IP edge network.
- **Quality of service (QoS):** Because of the high bandwidth and real-time requirements imposed by HD videoconferencing, it is critical to implement QoS. Even streams that do not appear to require large bandwidth often suffer from short bandwidth bursts. QoS is recommended for any network that may encounter congestion. For further details concerning QoS, see *Chapter 6, Implementing Quality of Service on page 115.*

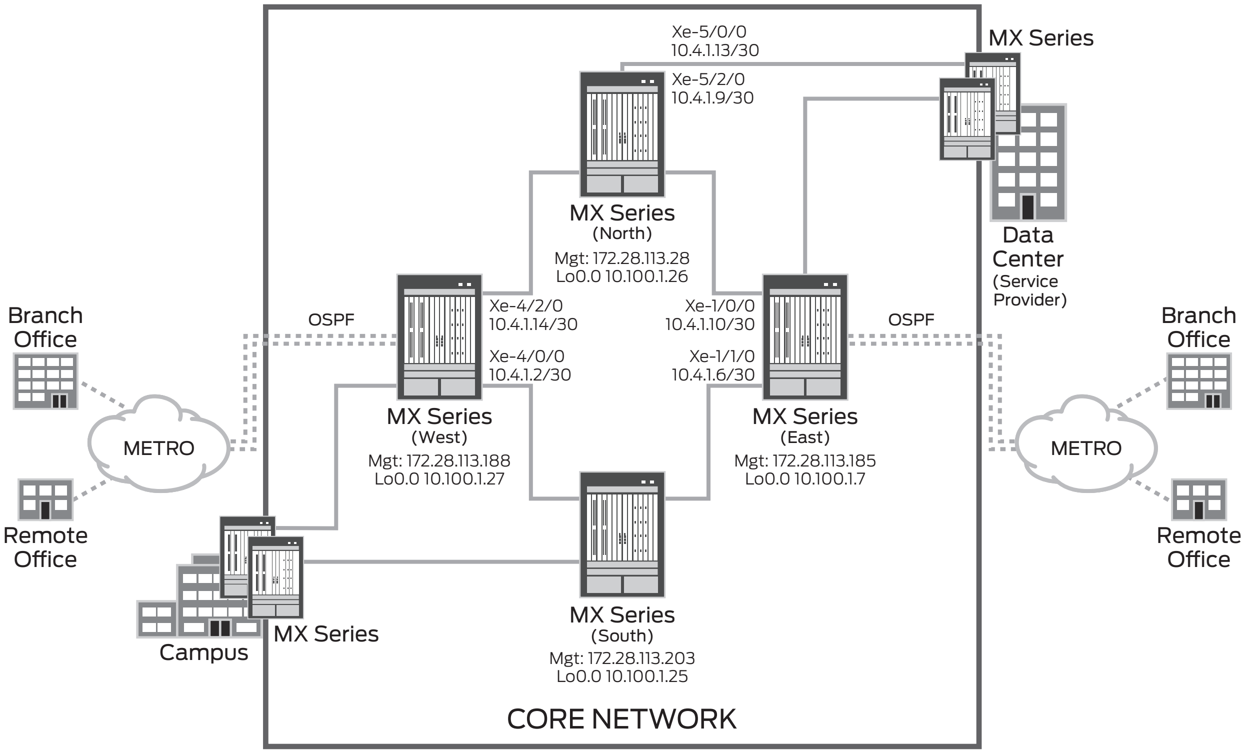

Figure 5.2 IP edge network overview

Defining the MPLS Backbone

To provision the basic HD videoconferencing services, administrators must configure the MPLS backbone tunnels through the provider's core that allows the videoconference traffic to traverse to their respective enterprise location. We will now describe the step-by-step process using configuration snippets to illustrate the required MPLS backbone configuration with displays of the show command that will help troubleshoot each step. To configure the MPLS backbone, you must configure the following:

- **Interfaces** – for backbone devices interconnectivity
- **Intermediate System-to-Intermediate System** (IS-IS) – for SP core network
- **Multiprotocol Label Switching** (MPLS) – for transporting traffic
- **Resource Reservation Protocol** (RSVP) – for MPLS traffic engineering
- **Border Gateway Patrol (BGP)** – for VPN provisioning

Configuring the Interfaces

The first step is to configure the interfaces to interconnect backbone devices. For this, we begin by assigning an arbitrary descriptor and assigning an IPv4 address to this interface. The iso keyword denotes that the Intermediate System-to-Intermediate System (IS-IS) will be used as the routing protocol between routers, and we specify that MPLS will be used on the backbone. This sequence is repeated for every interface that connects to the core network. This PE, the (MX Series (West) directly connects to two other PEs: MX Series (North) and MX Series (South).

Interface Configurations

```
## Enable reception and transmission of ISO protocol data units (PDUs) on each router
## interface in the network with the family statement, which identifies which protocol
## packets are accepted into the interfaces. Valid IS-IS packets are dropped if the
## interface is not configured with the family iso statement.
## Similarly the output shows that the interface xe-4/0/0 is configured to support MPLS.
## The family statement identifies which protocol packets are accepted into the
## interfaces. Valid MPLS packets are dropped if the interface is not configured with the
## MPLS protocol.

xe-4/0/0 {
      description “ connected to South-P/PE-203 “;
      unit 0 {
          family inet {
             address 10.4.1.2/30;
          }
## For a router to support IS-IS, you must configure an ISO Network Entity Title (NET)
## address on the router’s interfaces
          family iso;
## Configure family mpls on the logical units of the core interfaces, thereby identifying
## the interfaces that will be used for forwarding MPLS packets.
          family mpls;
      }
   }
   xe-4/2/0 {
      description “ connected to North-PE-28 “;
      unit 0 {
         family inet {
            address 10.4.1.14/30;
         }
         family iso;
         family mpls;
      }
   }
## Configure the interface connecting to the PE router Located at the campus location
   ge-5/0/0 {
      description “ connected to campus-MX-A 147 at interface ge-2/2/0”;
      unit 0 {
         family inet {
            address 10.4.1.49/30;
         }
         family iso;
         family mpls;
      }
   }
```

Verifying Connectivity to the Remote Interfaces

```
{master}
## Ping to the remote interface on  North-PE-28 to check connectivity
root@MX-West-188-re0> ping 10.4.1.13
PING 10.4.1.13 (10.4.1.13): 56 data bytes
64 bytes from 10.4.1.13: icmp_seq=0 ttl=64 time=0.608 ms
64 bytes from 10.4.1.13: icmp_seq=1 ttl=64 time=0.569 ms
64 bytes from 10.4.1.13: icmp_seq=2 ttl=64 time=0.572 ms
```

Configuring Intermediate System-to-Intermediate System (IS-IS)

Configure the provider network to run IS-IS as an Interior gateway Protocol (IGP). Each interface supporting IS-IS must be configured. In this sample, we define two interfaces using IS-IS. Both interfaces are configured to use Bidirectional Forwarding Detection (BFD) to detect line failures immediately.

IS-IS supports two routing levels. Level 1 is used for routing within a smaller area, while Level 2 interconnects different areas. As specified in the configuration, this network uses Level 2 routing between the PE and PE—that is, the backbone network.

Loopback interfaces do not actively participate in routing, so they are configured as passive participants.

```
isis {
## Configure the router with isis level 2 to route between areas and be able to use wide
## metrics, which is useful for MPLS traffic engineering
 level 2 wide-metrics-only;
   interface xe-4/0/0.0 {
## The Bidirectional Forwarding Detection (BFD) protocol is a hello mechanism that
## detects failures in a network
      bfd-liveness-detection {
## Note: Specifying an interval less than 300 milliseconds can cause undesired BFD
## flapping. In this sample we have configured it to be 100(minimum interval ) x 3
## (multiplier) = 300 ms
         minimum-interval 100;
         multiplier 3;
      }
   }
   interface xe-4/2/0.0 {
      bfd-liveness-detection {
         minimum-interval 100;
         multiplier 3;
      }
   }
   interface ge-5/0/0.0 {
   }
## configure the loopback as passive interface that should not be used for sending and
## receiving IS-IS packets
   interface lo0.0 {
      passive;
   }
}
```

Verifying ISIS Adjacencies

```
{master}[edit]
root@MX-West-188-re0# run show isis adjacency
Interface        System                        L     State      Hold (secs)     SNPA
ge-5/0/0.0       westford                      2        Up                6     0:1d:b5:a2:36:74
ge-5/0/5.600     SOLUTIONS-M7i-DC-PE           2        Up               12     0:14:f6:f2:c8:5d
xe-4/0/0.0       Metro-BROCKTON-203-RE-0       2        Up               21     0:5:85:7b:38:a5
xe-4/2/0.0       mx480-pe-28                   2        Up               23     0:1b:c0:73:5e:72
```

Configuring RSVP Signaling

BGP-based VPNs use RSVP signaling to dynamically set up label-switched paths (LSPs) through the provider's network. It is configured on PE and P routers, on the physical interfaces connecting to the core network—in other words, on interfaces connecting to other PE or P routers.

```
## The primary purpose of RSVP is to support dynamic signaling within label-switched
## paths (LSPs). When you enable both MPLS and RSVP on a router, MPLS Traffic Engineering
## becomes possible by using  traffic engineering database which is derived from the IGP
## topology of the network.
rsvp {
   interface xe-4/2/0.0;
   interface xe-4/0/0.0;
   interface ge-5/0/0.0;
   }
}
```

Verifying RSVP

```
{master}[edit]
root@MX-West-188-re0# run show rsvp neighbor
RSVP neighbor: 4 learned
Address       Idle    Up/Dn    LastChange       HelloInt    HelloTx/Rx        MsgRcvd
10.4.1.50        5    8/7            5:03              9    84619/84583       137512
10.4.1.13        0    58/57    3d 1:10:02              9    227577/227577     577814
10.4.1.25       10    1/0    3w3d 0:48:35              9    229478/229478     214713
10.4.1.1     22:16    0/0    3w3d 0:48:33              9    30250/0           352
```

Configuring MPLS Label-Switched Paths

The next step is to configure MPLS label-switched paths to every other PE that this device must communicate with (exchange information). These paths are also used to advertise which customer-facing interfaces belong to each VPN. Link protection creates a backup path to the PE, which will be used if the primary path fails.

```
## Configure MPLS to set up signaled paths by using the label-switched-path statement at
## the [edit protocols mpls] hierarchy level. Enable MPLS on all routers that will
## participate in the label switching (this should be done, on all routers that might be
## part of a label-switching path). Enable RSVP on all routers and on all router
## interfaces that take part in the LSP signaling. Configure the routers at the beginning
## of the LSP.
mpls {
   label-switched-path MX-West-to-MX-North {
      to 10.100.1.26;
      link-protection;
   }
   label-switched-path MX-West-to-MX-South {
      to 10.100.1.25;
      link-protection;
   }
   label-switched-path MX-West-to-MX-East {
      to 10.100.1.7;
      link-protection;
   }
## This is a LSP to the Extended MPLS network at the Campus location
   label-switched-path MX-West-to-MX-Campus-147 {
      to 10.102.1.3;
      link-protection;
   }
 ## This is a LSP to the Extended MPLS network at the Campus location
  label-switched-path MX-West-to-MX-Campus-114 {
       to 10.101.114.14;
       link protection;
   }
   interface xe-4/2/0.0;
   interface lo0.0;
   interface ge-5/0/0.0;
}
```

Verifying MPLS LSPs

```
{master}[edit]
root@MX-West-188-re0# run show mpls lsp
Ingress LSP: 5 sessions
To             From            State   Rt P  ActivePath       LSPname
10.100.1.7     10.100.1.27     Up      0     *                MX-West-to-MX-East
10.100.1.25    10.100.1.27     Up      0     *                MX-West-to-MX-South
10.100.1.26    10.100.1.27     Up      0     *                MX-West-to-MX-North
10.101.114.14  10.100.1.27     Up      0     *                MX-West-to-MX-Campus-114
10.102.1.3     10.100.1.27     Up      0     *                MX-West-to-MX-Campus-147
Total 5 displayed, Up 5, Down 0

Egress LSP: 6 sessions
To             From            State   Rt    Style  Labelin  Labelout  LSPname
10.100.1.27    10.100.1.52     Up      0     1      SE       3         - MX52-to-MX188
10.100.1.27    10.101.114.14   Up      0     1      FF       3         - MX-Campus-114-to-MX-West-188
10.100.1.27    10.100.1.7      Up      0     1      SE       3         - MX185-to-MX188
10.100.1.27    10.100.1.25     Up      0     1      SE       3         - MX-South-to-MX-West
```

```
10.100.1.27      10.102.1.3       Up      0  1 FF      3       - MX-Campus-147-to-MX-West-188
10.100.1.27      10.100.1.26      Up      0  1 SE      3       - MX-North-to-MX-West
Total 6 displayed, Up 6, Down 0

Transit LSP: 9 sessions
To             From            State    Rt    Style   Labelin   Labelout  LSPname
10.100.1.7     10.100.1.52     Up       0     1 SE    306640    309056    SOLUTIONS-to-WORCESTER
10.100.1.7     10.102.1.3      Up       0     1 FF    306752    309280    MX-Campus-147-to-MX-East-185
10.100.1.25    10.102.1.3      Up       0     1 FF    306736    309264    MX-Campus-147-to-MX-South-203
10.100.1.26    10.100.1.52     Up       0     1 SE    306672    3         SOLUTIONS-to-MARLBOROUGH
10.100.1.26    10.102.1.3      Up       0     1 FF    306768    3         MX-Campus-147-to-MX-North-28
10.100.1.52    10.100.1.7      Up       0     1 SE    306656    3         WORCESTER-to-SOLUTIONS
10.102.1.3     10.100.1.7      Up       0     1 FF    306720    3         MX185-to-MX-campus-147
10.102.1.3     10.100.1.25     Up       0     1 FF    306704    3         MX-South-to-MX-Campus-147
10.102.1.3     10.100.1.26     Up       0     1 SE    306688    3         MX-North-to-MX-Campus-147
Total 9 displayed, Up 9, Down 0
```

Configuring BGP

This network uses BGP/MPLS VPNs as defined in RFC 4364. Enabling BGP always requires that an autonomous system (AS) number be used to uniquely identify each domain.

```
routing-options {
   router-id 10.100.1.27;
   autonomous-system 100;
}
```

Configuring iBGP

Interior BGP (iBGP) is used to enable the routers to exchange information about routes originating and terminating in each VPN. The PE routers use this information to determine which labels to use for traffic destined to remote sites. The iBGP session runs through the loopback address, which in this setup is 10.100.1.27. Note that there is an entry for each remote PE, whether or not they are directly connected to this PE.

```
## Unlike some other routing protocols, BGP routers do not automatically discover each
## other and begin exchanging routing information. Instead, each BGP router must be
## explicitly configured with a set of BGP peers with which it exchanges routing
## information. On the PE routers, configure an IBGP session with the following
## properties:

## VPN family—To indicate that the IBGP session is used for the VPN signaling and auto
## provisioning, include the family inet-vpn statement.

## Loopback address—Include the local-address statement, specifying the local PE router's
## loopback address. The IBGP session for VPNs runs between the loopback addresses of
## participating routers. You must also configure the lo0 interface at the [edit
## interfaces] hierarchy level as shown in the interface configurations.
## Neighbor address: include the neighbor statement
bgp {
## Configure an iBGP group on the routers that belongs to the service providers network
   group hdvc-bgp-grp {
      type internal;
      local-address 10.100.1.27;
      family inet {
         any;
      }
      family inet-vpn {
         any;
      }
      local-as 100;
      neighbor 10.100.1.7;
      neighbor 10.100.1.26;
      neighbor 10.100.1.25;
      neighbor 10.101.114.14;
   }
}
```

Verifying BGP

```
root@MX-West-188-re0# run show bgp summary
Groups: 1 Peers: 4 Down peers: 0
Table          Tot Paths   Act Paths  Suppressed   History    Damp State    Pending
inet.0             0           0           0           0           0           0
inet.2             0           0           0           0           0           0
bgp.l3vpn.0       10          10           0           0           0           0
bgp.l3vpn.2        0           0           0           0           0           0
Peer                     AS     InPkt    OutPkt    OutQ  Flaps Last Up/Dwn State|#Active/Received/
Accepted/Damped...
10.100.1.7             100     10014      9952      0     2  3d 3:01:50 Establ
  inet.0: 0/0/0/0
  inet.2: 0/0/0/0
  bgp.l3vpn.0: 1/1/1/0
  bgp.l3vpn.2: 0/0/0/0
  BigCo.inet.0: 1/1/1/0
10.100.1.25            100        22        25     0     1        9:06 Establ
  inet.0: 0/0/0/0
  inet.2: 0/0/0/0
  bgp.l3vpn.0: 0/0/0/0
  bgp.l3vpn.2: 0/0/0/0
10.100.1.26            100      9972      9951      0     2  3d 3:01:29 Establ
  inet.0: 0/0/0/0
  inet.2: 0/0/0/0
  bgp.l3vpn.0: 7/7/7/0
  bgp.l3vpn.2: 0/0/0/0
  BigCo.inet.0: 7/7/7/0
10.101.114.14          100        22        23     0     3        8:17 Establ
  inet.0: 0/0/0/0
  inet.2: 0/0/0/0
  bgp.l3vpn.0: 2/2/2/0
  bgp.l3vpn.2: 0/0/0/0
  BigCo.inet.0: 2/2/2/0
```

Configuring Customers

We begin by defining the customer's VPN at each PE, which has connected sites within this VPN.

Creating a VPN

Each VPN (also known as a routing instance) must have a unique name that is used across all PEs that participate in this VPN. The same name must be used on each PE router to represent this customer's VPN. In addition, a unique route distinguisher value, consisting of the BGP Autonomous System (100) number and the PE's unique router-id is also assigned.

In addition, you must configure the following within a VPN instance:

- **`route-distinguisher:`** Each routing instance must have a unique route distinguisher associated with it. This allows BGP to distinguish between potentially identical Network Layer Reachable Information (NLRI) messages received from different VPNs.
- **`vrf-table-label:`** The vrf-table-label statement allows the examination of the encapsulated IP header at an egress PE router and administrators must enable this functionality so that they can achieve the following:
    - Forward traffic on a PE-router-to-CE-device interface, where the CE device is an L2 switch without IP capabilities, for example a metro Ethernet switch towards HD videoconferencing service VLANs.
    - Apply any IP based egress filtering actions on the HD videoconferencing service VLAN interfaces. When packets for this VPN arrive on a core-facing interface, they are treated as if the enclosed IP packet arrived on the label switched interface (LSI) and are then forwarded and filtered based on the correct table.

```
## configure the VPN in the edit routing-instance hierarchy.
BigCo {
    instance-type vrf;
    access-profile sol-188;
    interface ge-1/2/0.600;
## configure unique route distinguisher
    route-distinguisher 100:188;
## configure target community for this VPN instance
    vrf-target target:100:600;
## configure Ip Header filtering as described
    vrf-table-label;
## configure Import and Export policies as described below
    vrf-import BigCo-IMPORT;
    vrf-export BigCo-EXPORT;
    routing-options {
        router-id 10.100.1.188;
        autonomous-system 100;
        auto-export;
        }
    }
    protocols {
## OSPF runs between the PE and CE interfaces
        ospf {
            area 0.0.0.0 {
                interface ge-1/2/0.600;
            }
        }
    }
}
```

Defining Routing Policies

On each PE router, you must define policies that define how routes are imported into and exported from the router’s VPN Routing and Forwarding (VRF) table. In these policies, you must define the route target.

To configure policy for the VRF tables, perform the following configurations:

- VRF Target Community
- Route Target
- Import Policy
- Export Policy
- VRF-Table-Label

Configuring a VPN Routing and Forwarding Target Community

It is also necessary to configure a VPN Routing and Forwarding (VRF) target community. Using the **`vrf-target`** statement causes default VRF import and export policies to be generated that accept and tag routes with the specified target community.

VRF target community is specified in the format **`target:x:y`**, where x is the AS number and y is a value that the administrator assigns to this VPN.

To apply the import and export policies, include the **`vrf-export`** and **`vrf-import`** statement at the routing instance, as shown in the following code snippet.

```
routing-instances {
   BigCo {
## Applying VRF target
      vrf-target target:100:600;
      }
      protocols {
## OSPF configured between the PE to CE links
         ospf {
## applying policies to import and export the routing information for this VPN instance
            export BigCo-EXPORT;
            import BigCo-IMPORT;
            }
         }
      }
   }
}
```

Configuring a Route Target

As part of the policy configuration for the VPN routing table, you must define a route target that defines which VPN the route belongs to. When configuring VPNs for many enterprises on the same PE router, be sure to assign unique route target values to avoid the possibility of adding route and signaling information to the wrong VPN routing table.

```
policy-options {
## The route target typically consists of the AS number (100) and a unique identifier for
## the VPN (600)
   community BigCo members target:100:600;
}
```

Configuring an Import Policy

For each VPN, import policies are applied to routes learned from other PE routers to determine whether the route should be added to the PE router's **bgp.l3vpn.0** routing table. To do this, each routing instance on a PE router has a VRF import policy. If the routes match the conditions, the route is installed in the PE router's **BigCo.inet.0** VRF table.

```
policy-statement BigCo-IMPORT {
      term 1 {
         from protocol bgp ;
         from community BigCo;
         then accept;
      }
      term 2 {
## An import policy must contain a second term that rejects all other routes
         then reject;
      }
   }
```

VPN route processing differs from normal BGP route processing. In BGP, routes are accepted if they are not explicitly rejected by import policy. However, because many more VPN routes are expected, the router does not accept (and hence store) VPN routes unless the route matches at least one VRF import policy. If no VRF import policy explicitly accepts the route, it is discarded. If a VPN change occurs on a PE router, such as adding a new VRF table or changing a VRF import policy, the PE router sends a BGP route refresh message to the other PE routers (or to the route reflector if this is part of the VPN topology). This is done in order to retrieve all VPN routes so they can be re-evaluated to determine whether they should be kept or discarded.

Configuring an Export Policy

Each VPN should also have an export policy, which defines how routes that are announced by locally connected CEs are forwarded to other PEs. The export policy evaluates all routes received over the routing protocol session with the CE router. If the routes match the conditions, the specified community target (which is the route target) is added to them and they are exported to the remote PE routers. An export policy must contain a second term that rejects all other routes.

```
policy-options {
   policy-statement BigCo-EXPORT {
      term 1 {
         from protocol [ ospf | static ];
         then community add BigCo;
         then accept;
      }
      term 2 {
## An export policy must contain a second term that rejects all other routes
         then reject;
      }
```

Configuring VRF-Table-Label

The **`vrf-table-label`** statement makes it possible to map the inner label to a specific VRF routing table; such mapping allows the examination of the encapsulated IP header at an egress VPN router. We enable this functionality to achieve the following:

- Forward traffic on a PE-router-to-CE-device interface, in a shared infrastructure, where the CE device is a Layer 2 switch without IP capabilities, for example, a metro Ethernet switch. The first lookup is done on the VPN label to determine which VRF table to refer to, and the second lookup is done on the IP header to determine how to forward packets to the correct end hosts on the shared infrastructure.
- When we use the **`vrf-table-label`** statement to configure a VRF routing table, a label-switched interface (LSI) logical interface label is created and mapped to the VRF routing table. Any routes configured in a VRF routing table with the **`vrf-table-label`** statement are advertised with the LSI logical interface label allocated for the VRF routing table. When packets for this VPN arrive on a core-facing interface, they are treated as if the enclosed IP packet arrived on the LSI and are then forwarded and filtered based on the correct table.

```
[edit routing-instances]
BigCo {
   vrf-table-label;
}
```

Configuring Multiple Customers

To support multiple customers, the administrator must create a separate VPN for each customer. This is done by configuring multiple VPN instances traversing through the MPLS core network. This configuration separates routing information of different enterprises as well as the provider's routing information. To provision these services successfully requires a thorough understanding of the routing tables. Figure 5.3 illustrates a network supporting multiple customer VPNs.

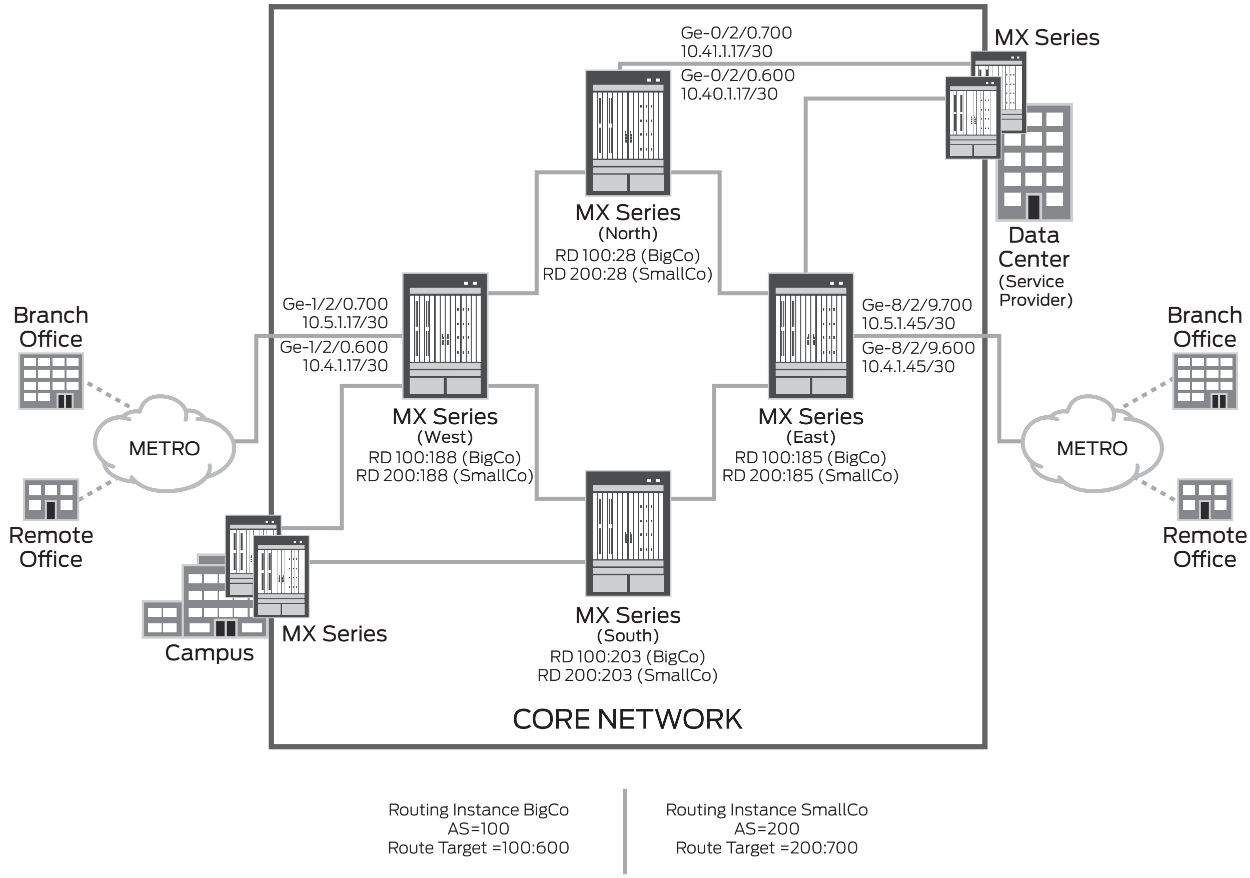

Figure 5.3 Building multiple customer VPNs

Each PE router creates a separate set of routing tables for each VPN, called VPN routing and forwarding (VRF) tables. Any customer or site that belongs to the VPN can access only the sites in the VRF tables for that VPN. Only the VRF table associated with a customer site is consulted for packets from that site.

The following tables are maintained for each customer's VPN instance:

- **`bgp.l3vpn.0:`** Stores all VPN unicast routes received from other PE routers.
- **`[Routing-instance-name].inet.0:`** Contains all explicitly configured static routes and all unicast routes received from directly connected CE routers in a routing instance, that is, in a single VPN.
- **`inet.3:`** Stores all MPLS routes learned from RSVP signaling done for VPN traffic.
- **`inet.0:`** Stores routes learned by the iBGP sessions between the PE routers.

Verifying VPN tables

```
root@MX-West-188-re0# run show route table BigCo.inet.0
BigCo.inet.0: 21 destinations, 21 routes (21 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.4.1.16/30      *[Direct/0] 3d 21:01:27
                > via ge-1/2/0.600
10.4.1.17/32      *[Local/0] 3d 21:01:27
                  Local via ge-1/2/0.600
10.8.1.44/30      *[BGP/170] 3d 21:04:27, localpref 100, from 10.100.1.7
                  AS path: I
                > to 10.4.1.13 via xe-4/2/0.0, label-switched-path MX-West-to-MX-East
10.8.1.49/32      *[Local/0] 4w1d 02:16:51
                  Reject
10.11.11.0/24     *[BGP/170] 3w0d 23:25:22, MED 2, localpref 100, from 10.100.1.26
                  AS path: I
                > to 10.4.1.13 via xe-4/2/0.0, label-switched-path MX-West-to-MX-North
10.12.12.0/24     *[BGP/170] 3w0d 23:25:22, MED 2, localpref 100, from 10.100.1.26
                  AS path: I
                > to 10.4.1.13 via xe-4/2/0.0, label-switched-path MX-West-to-MX-North
10.31.31.0/24     *[BGP/170] 3w0d 23:25:22, MED 2, localpref 100, from 10.100.1.26
                  AS path: I
                > to 10.4.1.13 via xe-4/2/0.0, label-switched-path MX-West-to-MX-North
10.40.1.16/29     *[BGP/170] 3w0d 23:25:22, localpref 100, from 10.100.1.26
                  AS path: I
                > to 10.4.1.13 via xe-4/2/0.0, label-switched-path MX-West-to-MX-North
10.100.1.5/32     *[BGP/170] 3d 21:03:42, MED 1, localpref 100, from 10.100.1.7
                  AS path: I
                > to 10.4.1.13 via xe-4/2/0.0, label-switched-path MX-West-to-MX-East
10.100.1.12/32    *[OSPF/10] 3d 21:00:42, metric 1
                > to 10.4.1.18 via ge-1/2/0.600
10.100.1.114/32   *[BGP/170] 4w1d 22:10:45, MED 0, localpref 100, from 10.102.1.3
                  AS path: I
                > to 10.4.1.50 via ge-5/0/0.0, label-switched-path MX-West-to-MX-Campus-147
10.100.1.147/32   *[BGP/170] 4w1d 22:10:45, localpref 100, from 10.102.1.3
                  AS path: I
                > to 10.4.1.50 via ge-5/0/0.0, label-switched-path MX-West-to-MX-Campus-147
10.100.1.185/32   *[BGP/170] 3w0d 23:25:21, localpref 100, from 10.100.1.7
                  AS path: I
                > to 10.4.1.13 via xe-4/2/0.0, label-switched-path MX-West-to-MX-East
10.100.1.188/32   *[Direct/0] 4w1d 22:12:39
                > via lo0.188
10.100.114.13/32  *[BGP/170] 3w0d 23:25:22, MED 1, localpref 100, from 10.100.1.26
                  AS path: I
                > to 10.4.1.13 via xe-4/2/0.0, label-switched-path MX-West-to-MX-North
192.168.41.0/24   *[OSPF/10] 3d 21:00:42, metric 2
                > to 10.4.1.18 via ge-1/2/0.600
192.168.42.0/24   *[BGP/170] 4w1d 22:10:45, localpref 100, from 10.102.1.3
                  AS path: I
                > to 10.4.1.50 via ge-5/0/0.0, label-switched-path MX-West-to-MX-Campus-147
192.168.50.0/24   *[BGP/170] 3d 21:03:42, MED 2, localpref 100, from 10.100.1.7
                  AS path: I
                > to 10.4.1.13 via xe-4/2/0.0, label-switched-path MX-West-to-MX-East
192.168.50.232/32 *[BGP/170] 3d 08:42:54, localpref 100, from 10.100.1.7
                  AS path: I
                > to 10.4.1.13 via xe-4/2/0.0, label-switched-path MX-West-to-MX-East
192.168.70.0/24   *[BGP/170] 3d 21:03:42, MED 2, localpref 100, from 10.100.1.7
                  AS path: I
                > to 10.4.1.13 via xe-4/2/0.0, label-switched-path MX-West-to-MX-East
224.0.0.5/32      *[OSPF/10] 4w1d 22:12:46, metric 1
                  MultiRecv
```

inet.3 table

```
root@MX-West-188-re0# run show route table inet.3

inet.3: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.100.1.7/32      *[RSVP/7/1] 3d 02:20:38, metric 30
                > to 10.4.1.13 via xe-4/2/0.0, label-switched-path MX-West-to-MX-East
10.100.1.25/32     *[RSVP/7/1] 01:30:54, metric 20
                > to 10.4.1.13 via xe-4/2/0.0, label-switched-path MX-West-to-MX-South
10.100.1.26/32     *[RSVP/7/1] 3d 02:20:35, metric 20
                > to 10.4.1.13 via xe-4/2/0.0, label-switched-path MX-West-to-MX-North
10.101.114.14/32   *[RSVP/7/1] 01:30:28, metric 30
                > to 10.4.1.13 via xe-4/2/0.0, label-switched-path MX-West-to-MX-Campus-114
10.102.1.3/32      *[RSVP/7/1] 01:15:52, metric 20
                > to 10.4.1.50 via ge-5/0/0.0, label-switched-path MX-West-to-MX-Campus-147
```

bgp.l3vpn.0 table

```
root@MX-West-188-re0# run show route table bgp.l3vpn.0

bgp.l3vpn.0: 14 destinations, 14 routes (14 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

100:602:10.11.11.0/24
               *[BGP/170] 3w0d 23:32:20, MED 2, localpref 100, from 10.100.1.26
                  AS path: I
                > to 10.4.1.13 via xe-4/2/0.0, label-switched-path MX-West-to-MX-North
100:602:10.12.12.0/24
               *[BGP/170] 3w0d 23:32:20, MED 2, localpref 100, from 10.100.1.26
                  AS path: I
                > to 10.4.1.13 via xe-4/2/0.0, label-switched-path MX-West-to-MX-North
100:602:10.31.31.0/24
               *[BGP/170] 3w0d 23:32:20, MED 2, localpref 100, from 10.100.1.26
                  AS path: I
                > to 10.4.1.13 via xe-4/2/0.0, label-switched-path MX-West-to-MX-North
100:602:10.40.1.16/29
               *[BGP/170] 3w0d 23:32:20, localpref 100, from 10.100.1.26
                  AS path: I
                > to 10.4.1.13 via xe-4/2/0.0, label-switched-path MX-West-to-MX-North
100:602:10.100.114.13/32
               *[BGP/170] 3w0d 23:32:20, MED 1, localpref 100, from 10.100.1.26
                  AS path: I
                > to 10.4.1.13 via xe-4/2/0.0, label-switched-path MX-West-to-MX-North
100:603:10.8.1.44/30
               *[BGP/170] 3d 21:11:23, localpref 100, from 10.100.1.7
                  AS path: I
                > to 10.4.1.13 via xe-4/2/0.0, label-switched-path MX-West-to-MX-East
100:603:10.100.1.5/32
               *[BGP/170] 3d 21:10:38, MED 1, localpref 100, from 10.100.1.7
                  AS path: I
                > to 10.4.1.13 via xe-4/2/0.0, label-switched-path MX-West-to-MX-East
100:603:10.100.1.185/32
               *[BGP/170] 3w0d 23:32:17, localpref 100, from 10.100.1.7
                  AS path: I
                > to 10.4.1.13 via xe-4/2/0.0, label-switched-path MX-West-to-MX-East
100:603:192.168.50.0/24
               *[BGP/170] 3d 21:10:38, MED 2, localpref 100, from 10.100.1.7
                  AS path: I
                > to 10.4.1.13 via xe-4/2/0.0, label-switched-path MX-West-to-MX-East
100:603:192.168.50.232/32
               *[BGP/170] 3d 08:49:50, localpref 100, from 10.100.1.7
                  AS path: I
                > to 10.4.1.13 via xe-4/2/0.0, label-switched-path MX-West-to-MX-East
100:603:192.168.70.0/24
               *[BGP/170] 3d 21:10:38, MED 2, localpref 100, from 10.100.1.7
                  AS path: I
                > to 10.4.1.13 via xe-4/2/0.0, label-switched-path MX-West-to-MX-East
100:147:10.100.1.114/32
               *[BGP/170] 4w1d 22:17:58, MED 0, localpref 100, from 10.102.1.3
                  AS path: I
                > to 10.4.1.50 via ge-5/0/0.0, label-switched-path MX-West-to-MX-Campus-147
100:147:10.100.1.147/32
               *[BGP/170] 4w1d 22:17:58, localpref 100, from 10.102.1.3
                  AS path: I
                > to 10.4.1.50 via ge-5/0/0.0, label-switched-path MX-West-to-MX-Campus-147
100:147:192.168.42.0/24
               *[BGP/170] 4w1d 22:17:58, localpref 100, from 10.102.1.3
                  AS path: I
                > to 10.4.1.50 via ge-5/0/0.0, label-switched-path MX-West-to-MX-Campus-147
```

Chapter 6

Building the HD Videoconferencing Datacenter

Part Two

The videoconferencing datacenter serves as the primary call setup hub and includes an infrastructure that helps with registering and provisioning the video endpoints, mixing multipoint calls and setting up conference bridges. The data center is also equipped with the intelligence to track and control resources that will be used by videoconferencing calls. These resources may include media conferencing equipment, network infrastructure and end-to-end bandwidth.

If the videoconferencing datacenter is part of a service provider managed solution, devices to track and report on calls – Call Detail Record (CDR) and call quality are also implemented. As with any data center supporting critical business operations, the videoconferencing datacenter has to be architected for security, performance, scale and high availability, while attempting to simplify the operational support and maintenance of the HD videoconferencing service.

This chapter serves as a reference for building a videoconferencing datacenter for hosting HD videoconferencing service when there is a mix of endpoints with different capabilities at the customer premises locations.

This chapter discusses the network architecture for a videoconferencing datacenter, demonstrates how to address key design considerations and provides implementation details.

Implementation Details

This section covers details on the network infrastructure and implementation (configuration) to deploy the HD videoconferencing videoconferencing datacenter reference framework, as depicted in Figure 6.1. The implementation details address every aspect of the design criteria listed above.

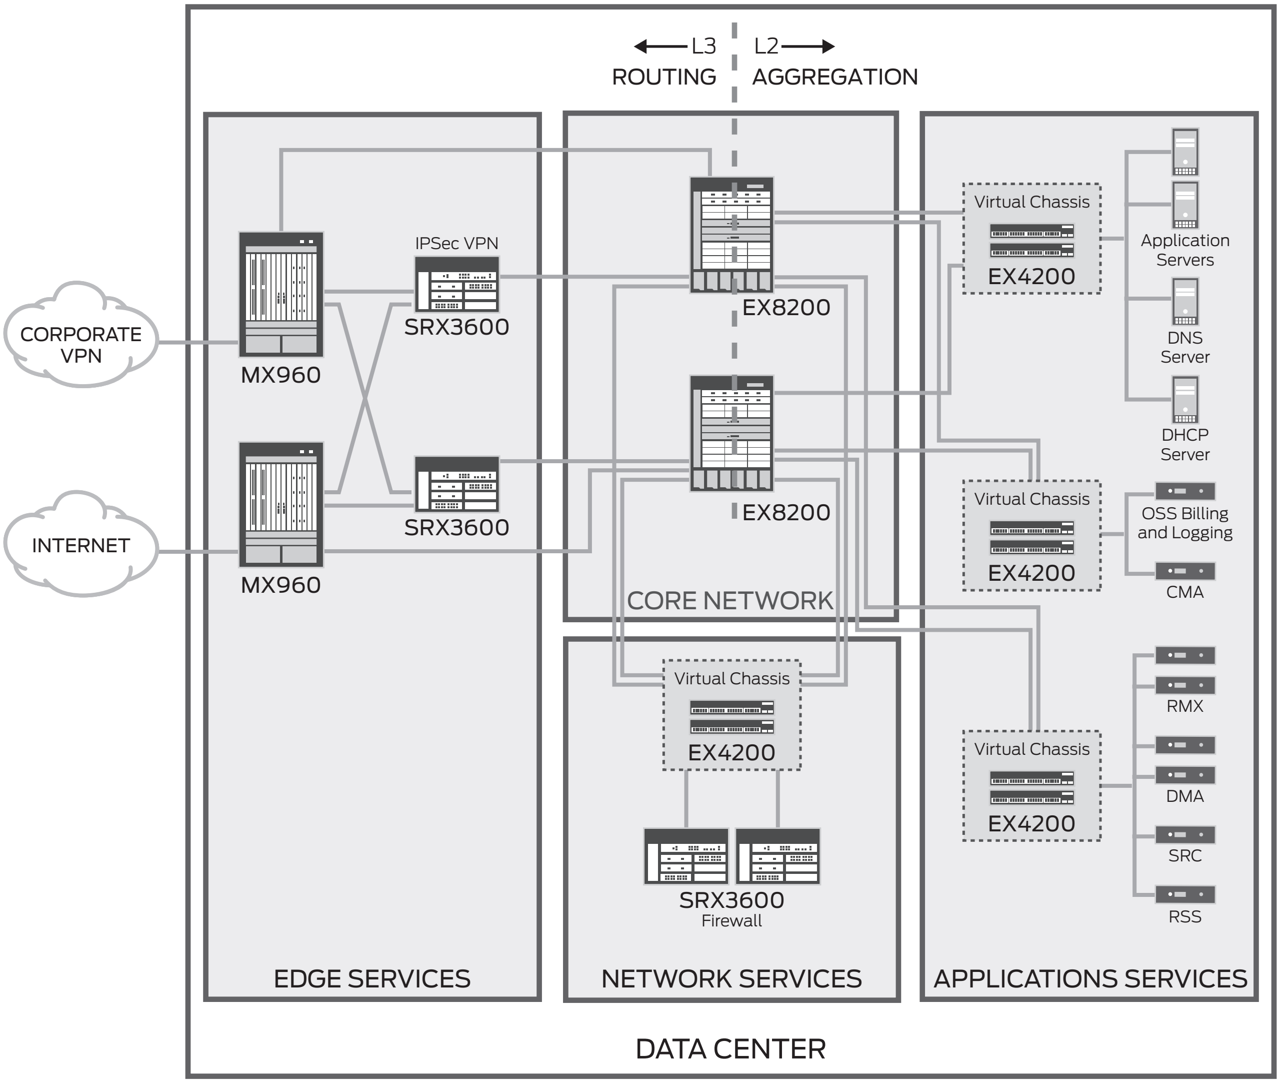

Figure 6.1 Data center overview

Data Center Infrastructure

The key equipment used in the videoconferencing datacenter as per the framework is listed below. See *Appendix C: Juniper Products section on page 171* and *Appendix D: Polycom Products section on page 176* for further details on the capabilities and specifications of the equipment used in videoconferencing datacenter. VLANs are used locally in the data center to provide L3 and L2 separation. Note that these VLANs do not extend outside of the data center.

Equipment used in this deployment include:

- **Edge services**, which provides connectivity to the wide area network, consists of two primary components. The MX480 3D Universal Edge Router provides the WAN connection, while the SRX3600 Services Gateway provides IPSec termination.
- **Data center core**, which interconnects all equipment within the data center, consists of two EX8200 Ethernet Switches.
- **Network services** which protect the data center from attack, use an SRX3600 Services Gateway acting as the firewall.
- **Application services**, which include videoconferencing infrastructure products as well as relevant networking equipment include the EX4200 Ethernet Switch, Polycom videoconferencing infrastructure and policy management products.

Videoconferencing Infrastructure

The videoconferencing infrastructure in the data center consists of the following Polycom products:

- **Distributed Media Application (DMA)** is a network-based application for managing and distributing multipoint calls across conference platforms. It also performs the role of a SIP server and H.323 gatekeeper.
- **Real-time Media Conferencing Platform (RMX)** is a multipoint conferencing platform designed for hosting of point- to- multipoint videoconferences.
- **Converged Management Application (CMA)** helps by centrally managing and deploying visual communication across the enterprise organization—from large conference rooms to individual desktops.

DMA and the RMX devices can be configured in cluster mode to improve scalability and availability. The DMA that serves as the call control engine is the gatekeeper, which is contacted by the endpoints during call setup. The DMA is pre-provisioned with information about endpoints, their capabilities and the site topology.

Appendix D: Polycom Products on page 176 includes pointers to Polycom configuration guides for the DMA 7000 and RMX 2000.

Policy Management: Session and Resource Control (SRC) Software

The Juniper Networks Session and Resource Control (SRC) software is a dynamic network resource allocation solution that enables service providers to deliver differentiated products and services. The SRC software operates on Juniper's C Series controllers. In the HD videoconferencing data center application, the SRC is network aware and helps to monitor and provision an end-to-end call while ensuring that the requested SLA is met. The SRC interacts with Polycom's DMA and performs the following functions in the HD videoconferencing solution:

- Maintains endpoint information
- Builds network infrastructure and resource information
- Provide call admission control
- Acts as SOAP gateway (interface with Polycom DMA)

The DMA communicates with Juniper's SRC, which decides whether to admit the call based on availability of network resources. Once the SRC has acknowledged the call setup request, the DMA communicates back to the requesting endpoint and the call is connected.

See *Chapter 10, Implementing Assured Forwarding on page 129* for configuring the SRC and the corresponding Junos MX routers to communicate with the SRC.

To set up Polycom's DMA to communicate with Juniper's SRC, see the *Polycom DMA – Juniper SRC Integration section on page 76* at the end of this chapter.

MPLS and L3VPN Connectivity to the Data Center

In this reference network, the data center core MX routers perform the functionality of a PE router as they participate in the extended MPLS and L3VPN network. This approach is recommended only if the data center handles traffic and routes to justify a PE router at the data center edge. This approach simplifies the network and reduces the number of hops that HD videoconferencing traffic needs to traverse from one end of the network to another. Alternatively, the edge router(s) at the data center could connect to the service provider's PE router in the core.

This section provides configuration details on the underlying protocols required to establish L3VPN. The network administrator must configure the PE router at the data center location with the following:

- Configure the MPLS Connection
  - Interface configuration
  - IGP for SP core network (for example IS-IS)
  - MPLS for transport of the traffic
  - RSVP for MPLS traffic engineering
  - BGP for VPN provisioning
- Configure the L3VPN
  - L3VPN Instance
  - L3VPN routing towards the data center core (for example OSPF)
  - VRF policies for exchange of routing information inside of VPN.

Configuring the MPLS Connections

Configuring the PE to P Router Interface

```
[edit interfaces]
xe-5/2/0 {
   description " connected to MX-West-188 on interface xe-4/2/0 ";
   mtu 4096;
   unit 0 {
      family inet {
         address 10.4.1.9/30;
      }
      ## To support ISIS, configure an ISO network entity title (NET) address on the
      ## router's interfaces
      family iso;
      ## Configure family mpls on the logical units of the core interfaces,
      ## thereby identifying the interfaces that will be used for forwarding
      ## MPLS packets.
      family mpls;
   }
}
```

Configuring IGP

```
[edit protocols]
isis {
   interface xe-5/2/0.0 {
      bfd-liveness-detection {
         minimum-interval 50;
         multiplier 3;
      }
      level 1 disable;
   }
   ## Configure the router with ISIS level 2 to route between areas and be able to
   ## use wide metrics, which are useful for MPLS Traffic engineering
   level 2 wide-metrics-only;
}
```

Configuring RSVP

```
## The primary purpose of RSVP is to support dynamic signaling within label-switched
## paths (LSPs). When you enable both MPLS and RSVP on a router, MPLS Traffic Engineering
## becomes possible by using  traffic engineering database which is derived from the IGP
## topology of the network.
[edit protocols]
rsvp {
   interface xe-5/3/0.0;
}
```

Configuring MPLS to Set up Signaled Paths using the Label-switched-path

```
[edit protocols]
mpls {
   label-switched-path MX-North-to-MX-East {
      to 10.100.1.7;
      link-protection;
   }
   label-switched-path MX-North-to-MX-West {
      to 10.100.1.27;
      link-protection;
   }
   label-switched-path MX-North-to-MX-Campus-147 {
      to 10.102.1.3;
      link-protection;
   }
   label-switched-path MX-North-to-MX-South {
      to 10.100.1.25;
      link-protection;
   }
   label-switched-path MX-North-to-MX-Campus-114 {
      to 10.101.114.14;
   }
   interface fxp0.0 {
      disable;
   }
   interface xe-5/3/0.0;
   interface xe-5/0/0.0;
   interface lo0.0;
   interface xe-5/2/0.0;
}
```

Configuring BGP

```
[edit protocols]
bgp {
   bfd-liveness-detection {
      minimum-interval 50;
      multiplier 3;
   }
   ## Configure an iBGP group on the routers that belongs to the
   ## service providers network
   group hdvc-bgp-grp {
      type internal;
      local-address 10.100.1.26;
      family inet {
         any;
      }
      family inet-vpn {
         any;
      }
      local-as 100;
      neighbor 10.100.1.25;
      neighbor 10.100.1.7;
      neighbor 10.100.1.27;
      neighbor 10.101.114.14;
      neighbor 10.102.1.3;
   }
}
```

Configuring L3VPN

The next step is to define the customer VPN at the data center PE router connected to the provider's core network. In this case, we have assigned a VPN named BigCo, and the keyword vrf that specifies that this is a L3VPN. In addition, configure the following within the VPN instance.

- **`route-distinguisher:`** Each routing instance must have a unique route distinguisher associated with it.
- **`vrf-table-label:`** The vrf-table-label statement allows the examination of the encapsulated IP header at an egress PE router.

Configuring HD videoconferencing service L3VPN instance

```
[edit routing-instances]
BigCo {
   instance-type vrf;
   interface ge-0/0/4.600;
   route-distinguisher 100:602;
   vrf-table-label;
   routing-options {
      router-id 10.100.1.26;
      autonomous-system 100;
      auto-export;
   }
}
```

Defining VRF Policies

On each PE router, you must define policies on how routes are imported into and exported from the router's VRF table. To configure policy for the VRF tables, configure the following:

- Route Target
- VRF Target Community

As part of the policy configuration for the VPN routing table, you must define a route target. The route target defines the specified VPN that the route belongs to. The route target typically consists of the AS number (100) and a unique identifier for the VPN (600).

Configuring a Route Target

```
[edit routing-instances]
BigCo {
   vrf-target target:100:600;
}
```

Configuring a VRF Target Community

It is also necessary to configure a VRF target community. For the VPN to know which routes belong to it, define a VRF target using the vrf-target statement. It causes default VRF import and export policies to be generated that accept and tag routes with the specified target community.

VRF target community is specified in the format **`target:x:y`**, where x is the AS number and y is a value that the administrator assigns to this VPN.

```
[edit policy-options]
community BigCo members target:100:600;
```

High Availability and Disaster Recovery

The following high availability features are used in the videoconferencing datacenter design:

- Virtual Router Redundancy Protocol (VRRP)
- Link Aggregation Groups (LAG)
- Bidirectional Forwarding Detection (BFD)
- Graceful Routing engine Switchover (GRES)
- Non-Stop routing (NSR) – MX Series Only
- In-service Software Upgrades (ISSU) – MX, EX Series.

The following section presents a high-level description for each of the above functionalities with sample configuration snippets. Figure 6.2 shows the devices where each technology is applicable.

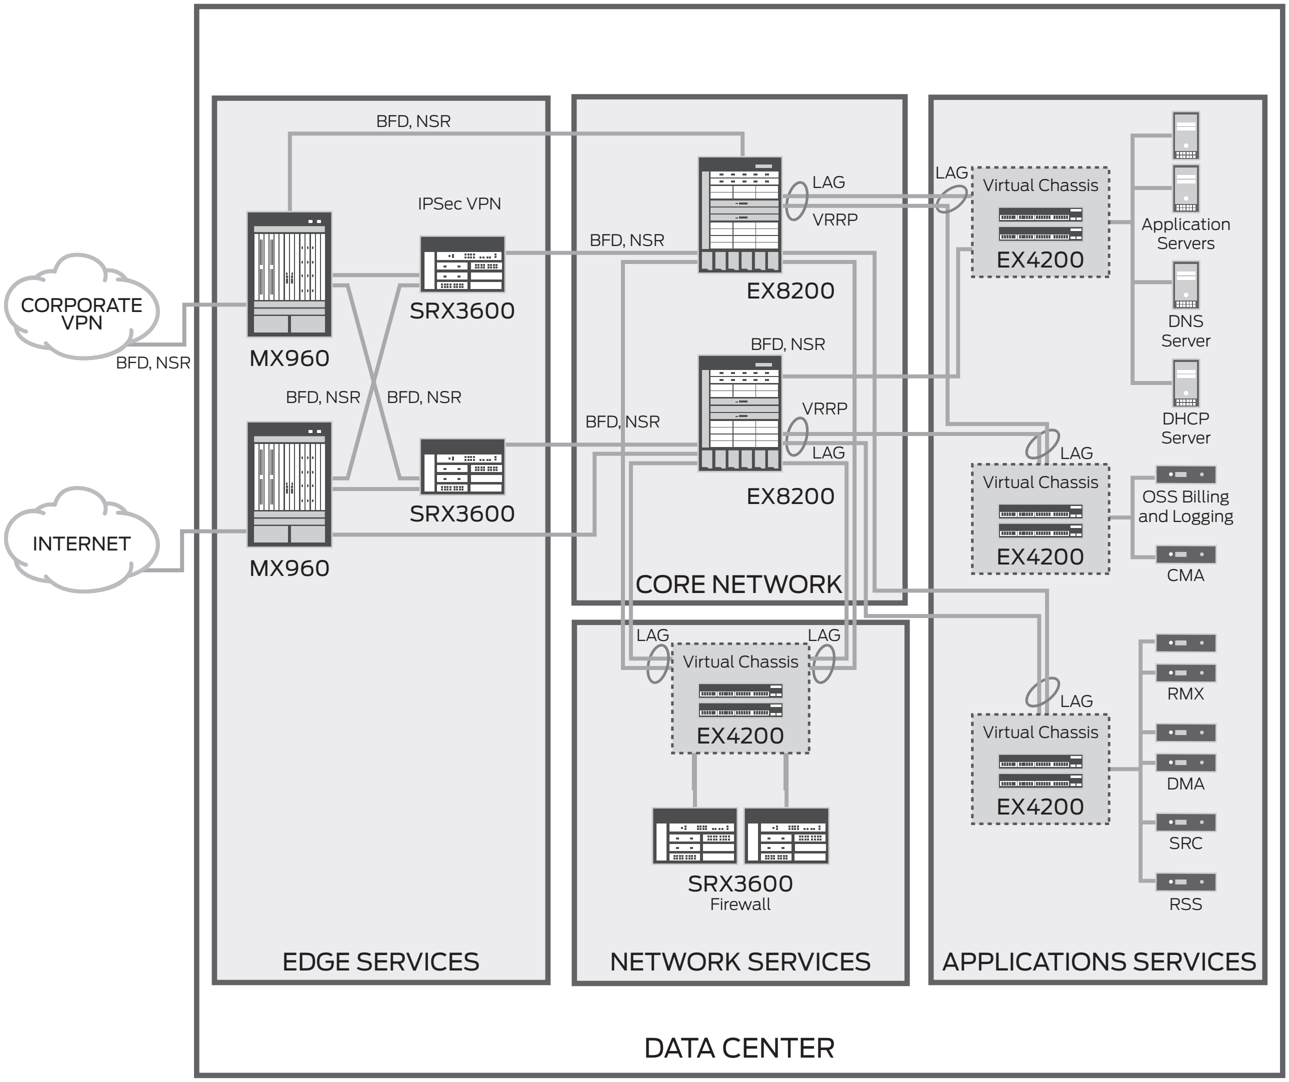

Figure 6.2 Data center resiliency

VRRP

Virtual Router Redundancy Protocol (VRRP) is a protocol, which runs on routers that are connected to the same broadcast domain. VRRP configuration assigns these routers to a group. The grouping eliminates the possibility of a single point of failure and thus provides high availability of network connectivity to the computing hosts on the broadcast domain that do not participate in IGP routing and need a default gateway for connectivity. Routers participating in VRRP share a virtual IP address and virtual MAC address. The shared virtual IP address serves as a gateway to the default route configured on the hosts.

One of the routers is elected dynamically as a default primary of the group and is active at any given time. All other participating routing devices perform a backup role. Operators can assign priorities to devices manually, forcing them to act as primary and backup devices. The primary VRRP sends out multicast advertisements to the backup devices at regular intervals (default interval is one second). When the backup devices do not receive an advertisement for a configured period, the device with the next highest priority becomes the new primary. This process occurs dynamically, thus enabling an automatic transition with minimal traffic loss. This VRRP action eliminates the dependence on achieving connectivity using a single routing platform that can result in a single point of failure. In addition, the change between the primary and backup roles occurs with minimum VRRP messaging and no intervention on the host side.

In the HD videoconferencing data center reference framework, VRRP is configured on the EX8200 interfaces interconnecting the access tier and on the interfaces to the network services layer, as shown in Figure 6.3.

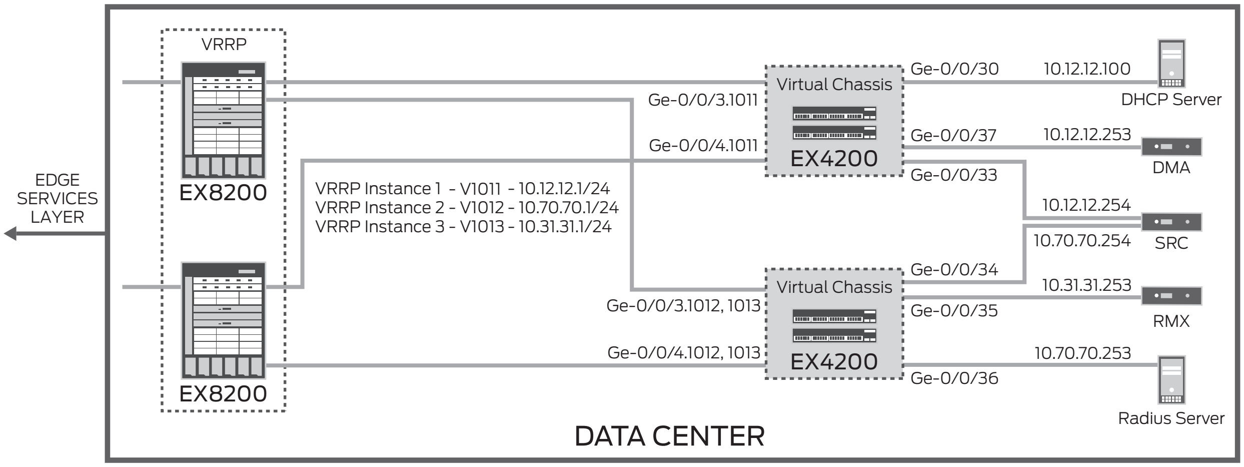

Figure 6.3 VRRP Configuration Overview

For VRRP configuration details, refer to *VRRP Configuration Hierarchy* at **juniper.net/techpubs/en_US/junos10.2/topics/reference/statement-hierarchy/vrrp-configuration-hierarchy.html**.

VRRP Configuration Snippet

The VRRP configuration snippet shows the minimum configuration required on the EX Series to enable a VRRP group.

```
## Configure the interface ge-0/0/31 on EX8200-B with an IP address of 10.12.12.31/24
## on the logical unit 1011.
## Define a VRRP group with a virtual IP of 10.12.12.1 and priority of 243.
[edit interfaces ge-0/0/31]
unit 1011 {
   family inet {
      address 10.12.12.31/24 {
         vrrp-group 1 {
            virtual-address 10.12.12.1;
            ## To enable faster failure detection and switchover configure
            ## advertisement interval to 100 ms
            fast-interval 100;
            priority 243;
            preempt {
               hold-time 0;
            }
            ## configure this  interface to accept packets destined for the virtual
            ## IP address, including the accept-data statement will make it easier
            ## to troubleshoot  as well
            accept-data;
            ## Configure the VRRP interface to track and dynamically change the
            ## priority of the VRRP group based on the state of the tracked
            ## interface, to trigger a new master router election
            track {
               ## Interface tracking to enable switchover
               interface ge-2/2/0 {
                  ## this cost is deducted from the group member priority
                  ## in case of a failure
                  priority-cost 10;
               }
            }
         }
      }
   }
}
## Configure the interface ge-0/0/11 on EX8200-A with an IP address of 10.12.12.11/24 on
## the logical unit 1011.
## Define a VRRP group with a virtual IP of 10.12.12.1 and priority of 240.
[edit interfaces ge-0/0/11]
unit 1011 {
   family inet {
      address 10.12.12.11/24 {
         vrrp-group 1 {
            virtual-address 10.12.12.1;
            priority 240;
            preempt {
               hold-time 0;
            }
            accept-data;
         }
      }
   }
}
```

Verifying VRRP operation and status of the master and backup routers

```
## verify that state of the VRRP routers
{master}[edit]
root@MX-A# run show vrrp
Interface    State      Group   VR state VR Mode   Timer    Type   Address
ge-0/0/31.1011    up           1   master   Active     A  0.011 lcl  10.12.12.31
                                                                           vip   10.12.12.1
root@MX-B # run show vrrp
Interface    State      Group   VR state VR Mode   Timer    Type   Address
ge-0/0/11.1011    up           1   backup   Active     A  0.295 lcl   10.12.12.11
                                                                           vip   10.12.12.1   mas
10.12.12.31
```

LAG

Ethernet Link Aggregation Group (LAG) is a feature that aggregates two or more physical Ethernet links into one logical link to obtain higher bandwidth and to provide link redundancy. LAG provides high link availability and capacity which results in improved performance and availability. Traffic is balanced across all links that are members of an aggregated bundle. The failure of a member link does not cause traffic disruption. Instead, because there is multiple member links, traffic continues over active links as remaining bandwidth permits.

In the data center reference framework, LAG is configured on the EX8200 interfaces connecting to the EX virtual chassis in the network and application services layer and on the EX virtual chassis aggregated interfaces.

For LAG configuration details, refer to *Understanding Aggregated Ethernet Interfaces* and LACP at **juniper.net/techpubs/en_US/junos10.0/topics/concept/interfaces-lag-overview.html**.

LAG Configuration Hierarchy

The LAG configuration snippet shows the minimum configuration required on the EX Series to prepare a group of LAG interfaces.

```
## configure the maximum number of aggregated-devices (0 - 127)
[edit chassis]
                              aggregated-devices {
                                 ethernet {
                                    device-count 2;
                                 }
                              }
                              ## include the 802.3ad reference under the desired interface
                              [edit interfaces]
                              ge-0/0/41 {
                                 ether-options {
                                    802.3ad ae0;
                                 }
                              }
                              ## create an aggregated-interface and set its properties
                              [edit interfaces ae0]
                              aggregated-ether-options {
                                 ## at least one of the links has to be up for this bundle to be marked 'up'
                                 minimum-links 1;
                                 ## all links in this AE will be GE, so link-speed is 1g
                                 link-speed 1g;
                                 ## lacp is enabled to allow for automatic addition/deletion of individual
                              links
                                 ## and for link monitoring
                                 lacp {
                                    ## LACP exchange is enabled to allow this AE to operate in active mode
                              and hence
                                    ## facilitate troubleshooting. This can be disabled as needed
                                    active;
                                    ## Set to exchange every second, so that link failures can be detected
                              and
                                    ## recovered from quickly
                                    periodic fast;
                                 }
}
```

BFD

The Bidirectional Forwarding Detection (BFD) protocol is a simple hello mechanism that detects failures in a network. Hello packets are sent at a specified, regular interval. A neighbor failure is detected when the router stops receiving a reply after a specified interval. The failure detection timers for BFD have shorter time limits than default failure detection mechanisms, providing faster detection.

In the HD videoconferencing data center reference framework, BFD is configured on the MX routers in the edge services and core network layers where it is enabled for the OSPF routing protocol.

For BFD configuration details, refer to *Configuring the BFD Protocol* in the *JUNOS® Software Routing Protocols Configuration Guide* at **juniper.net/techpubs/software/junos/junos94/swconfig-routing/configuring-the-bfd-protocol_3.html**.

BFD Configuration Snippet

```
[edit protocols ospf area 0.0.0.0]
interface ge-0/2/1.0 {
   bfd-liveness-detection {
      ## Transmit and receive intervals on the interface where the ospf neighbor
      ## is connected is set to 500 milliseconds as default is set to 10 seconds
      minimum-interval 500;
      ## Will wait for 1.5 seconds and then declare the neighbor on the interface dead
      multiplier 3;

   }
}
```

GRES

On the MX Series and EX8200 platforms, Graceful Routing Engine Switchover (GRES) redundancy can be set up if the platform has been deployed with two physical routing engines. One of the Routing Engines functions as the primary, while the other serves as a backup. When the primary Routing Engine fails, the backup Routing Engine automatically becomes the primary Routing Engine, thus increasing the availability of the device.

Any one of the following failures can trigger a switchover from the primary to the backup Routing Engine:

- **Hardware failure** – a hard disk error or a loss of power on the primary Routing Engine.
- **Software failure** – a kernel crash or a CPU lock. These failures cause a loss of keepalives from the primary to the backup Routing Engine.
- **Software process failure** –specific software processes that fail at least four times within the span of 30 seconds on the primary Routing Engine.

In the HD videoconferencing data center reference framework, GRES is configured on all the MX Series and EX8200 Series platforms.

For Graceful Routing Engine Switchover (GRES) configuration details, refer to *Configuring Routing Engine Redundancy* at **juniper.net/techpubs/en_US/junos10.2/topics/task/configuration/routing-engine-redundancy-configuring.html**.

Routing Engine Redundancy Configuration Snippet

```
## define the routing engine roles and failover mechanisms
[edit chassis]
redundancy {
   graceful-switchover;
   keepalive-time seconds;
   routing-engine 0 master;
}
## configure the automatic failover triggers
[edit chassis redundancy]{
   failover on-disk-failure;
   failover on-loss-of-keepalives;
}
## Specify the threshold time interval for loss of keepalives after which the backup
## Routing Engine takes over
[edit chassis redundancy]
keepalive-time seconds;
## Configure automatic switchover to the backup Routing Engine following a software
## process failure
[edit system processes]
routing failover other-routing-engine;
```

NSR

Nonstop Active Routing (NSR) preserves routing and interface information in a manner similar to graceful Routing Engine switchover. However, compared to graceful Routing Engine switchover, NSR goes a step further and saves the routing protocol information on the backup Routing Engine. It also preserves the protocol connection information in the kernel. Any switchover between the Routing Engines is dynamic, is transparent to the peers and occurs without any disruption to protocol peering. For these reasons, NSR is beneficial in cases where the peer routers do not support graceful Routing Engine switchover.

In the HD videoconferencing videoconferencing datacenter reference framework, NSR is configured on the MX routers connecting to the MPLS core routers.

For NSR configuration details, refer to *Configuring Nonstop Active Routing* at **juniper.net/techpubs/en_US/junos10.2/topics/task/configuration/nsr-configuring.html**.

NSR Configuration Snippet

```
## Enable graceful Routing Engine switchover under the chassis stanza.
[edit chassis redundancy]
graceful-switchover;

## Enable nonstop active routing under the routing-options stanza.
[edit routing-options]
nonstop-routing;

## Synchronize configuration changes on both Routing Engines.
[edit system]
commit synchronize;
```

```
## configure the RE to switchover to the backup when the routing protocol process (rpd)
## fails three times consecutively, in rapid intervals
[edit system processes routing failover]
routing failover other-routing-engine;
```

ISSU

In Service Software Upgrade (ISSU) facilitates software upgrades of Juniper devices in environments where there is a high concentration of users and business critical applications. Operators can use ISSU to upgrade the software from one Junos release to another without any disruption to the control plane. Any disruption to traffic during the upgrade is minimal.

ISSU runs only on platforms that support dual Routing Engines and requires that Graceful Routing Engine Switchover and NSR be enabled. Graceful Routing Engine Switchover is required because a switch from the primary to the backup Routing Engine must happen without any packet forwarding loss. The NSR with Graceful Routing Engine Switchover maintains routing protocol and control information during the switchover between the Routing Engines.

This feature significantly increases HD videoconferencing service availability to the users since systems upgrades and maintenance is one of the main contributing factors for the service downtime.

For ISSU configuration details, refer to *Unified ISSU in the JUNOS Software High Availability Configuration, Release 10.2* at **juniper.net/techpubs/en_US/junos10.2/information-products/pathway-pages/high-availability/high-availability.html**.

Verifying Conditions and Tasks Prior to ISSU Operation

Note that operational command outputs may differ based on configuration and the current state of the system.

```
Verify that the primary and backup Routing Engines are running the same software version using the show
version invoke-on all-routing-engines CLI command:
{master}
user@ MX480-DC-1> show version invoke-on all-routing-engines
re0:
---------------------------------------------------------------------------
Hostname: MX480-DC-1
Model: mx480
JUNOS Base OS boot [10.3R1.9]
JUNOS Base OS Software Suite [10.3R1.9]
JUNOS Kernel Software Suite [10.3R1.9]
JUNOS Crypto Software Suite [10.3R1.9]
JUNOS Packet Forwarding Engine Support (M/T Common) [10.3R1.9]
JUNOS Packet Forwarding Engine Support (MX Common) [10.3R1.9]
JUNOS Online Documentation [10.3R1.9]
JUNOS Voice Services Container package [10.3R1.9]
JUNOS Border Gateway Function package [10.3R1.9]
JUNOS Services AACL Container package [10.3R1.9]
JUNOS Services LL-PDF Container package [10.3R1.9]
JUNOS Services PTSP Container package [10.3R1.9]
JUNOS Services Stateful Firewall [10.3R1.9]
JUNOS Services NAT [10.3R1.9]
JUNOS Services Application Level Gateways [10.3R1.9]
JUNOS AppId Services [10.3R1.9]
JUNOS IDP Services [10.3R1.9]
JUNOS Runtime Software Suite [10.3R1.9]
JUNOS Routing Software Suite [10.3R1.9]
```

```
re1:
---------------------------------------------------------------------------
Hostname: MX480-DC-1
Model: mx480
JUNOS Base OS boot [10.2R1.8]
JUNOS Base OS Software Suite [10.2R1.8]
JUNOS Kernel Software Suite [10.2R1.8]
JUNOS Crypto Software Suite [10.2R1.8]
JUNOS Packet Forwarding Engine Support (M/T Common) [10.2R1.8]
JUNOS Packet Forwarding Engine Support (MX Common) [10.2R1.8]
JUNOS Online Documentation [10.2R1.8]
JUNOS Voice Services Container package [10.2R1.8]
JUNOS Border Gateway Function package [10.2R1.8]
JUNOS Services AACL Container package [10.2R1.8]
JUNOS Services LL-PDF Container package [10.2R1.8]
JUNOS Services PTSP Container package [10.2R1.8]
JUNOS Services Stateful Firewall [10.2R1.8]
JUNOS Services NAT [10.2R1.8]
JUNOS Services Application Level Gateways [10.2R1.8]
JUNOS AppId Services [10.2R1.8]
JUNOS IDP Services [10.2R1.8]
JUNOS Routing Software Suite [10.2R1.8]
```

Verify that graceful Routing Engine switchover and NSR are enabled using the commands:

```
MX480-DC-1> show system switchover
Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Peer state: Out of transition

MX480-DC-1> show task replication
       Stateful Replication: Enabled
       RE mode: Master

    Protocol             Synchronization Status
    OSPF                NotStarted
    BGP                 Complete
    IS-IS                NotStarted
    LDP                 Complete
```

BFD timer negotiation can be disabled explicitly during the ISSU activity using the [edit protocols bfd] hierarchy:

```
[edit protocols bfd]
no-issu-timer-negotiation;
```

Perform a software backup on each Routing Engine using the request system snapshot CLI command:

```
{master}
MX480-DC-1> request system snapshot
umount: /altroot: not currently mounted
Copying / to /altroot.. (this may take a few minutes)
umount: /altconfig: not currently mounted
Copying /config to /altconfig.. (this may take a few minutes)

The following filesystems were archived: / /config
```

Verifying a Unified ISSU. Execute the below command on the primary Routing Engine to verify the status of FPCs and their corresponding PICs after the most recent ISSU activity.

```
MX480-DC-1> show chassis in-service-upgrade
  Item          Status                  Reason
  FPC 0         Online
  FPC 1         Online
  FPC 2         Online
   PIC 0        Online
  PIC 1       Online
FPC 3         Offline              Offlined by CLI command
FPC 4         Online
  PIC 1       Online
FPC 5         Online
  PIC 0       Online
FPC 6         Online
  PIC 3       Online
FPC 7         Online
```

Virtual Chassis

A maximum of ten EX4200 chassis can be configured to form a Virtual Chassis, which improves availability and scalability. The Virtual Chassis operates as a single network entity and consists of designated primary and backup switches. Routing Engines on each of these two switches then become the master and backup Routing Engines of the Virtual Chassis, respectively. The rest of the switches of the Virtual Chassis assume the role of line cards. The master Routing Engine on the primary switch manages all the other switches that are members of the Virtual Chassis and has full control of the configuration and processes. It receives and transmits routing information, builds and maintains routing tables, and communicates with interfaces and the forwarding components of the member switches.

The backup switch acts as the backup Routing Engine of the Virtual Chassis and takes over as the master when the primary Routing Engine fails. The Virtual Chassis uses GRES and NSR to recover from control plane failures. Operators can physically connect individual chassis using either Virtual Chassis extension cables or 10G/1G Ethernet links.

In the HD videoconferencing videoconferencing datacenter reference framework, four Virtual Chassis entities are provisioned – one at the network services layer and three at the application and video services layer.

For Virtual Chassis configuration details, refer to *Configuring a Virtual Chassis with a Master and Backup in a Single Wiring Closet* at **juniper.net/techpubs/en_US/junos9.5/topics/example/virtual-chassis-basic.html**.

Virtual Chassis Configuration Snippet

```
## Define members of a virtual chassis.
[edit]
virtual-chassis {
    member 1 {
        mastership-priority 130;
    }
    member 2 {
        mastership-priority 130;
    }
}
## Set the Virtual Management Ethernet interface for OOB management of the VC.
[edit]
interfaces {
    vme {
        unit 0 {
            family inet {
                address 172.28.113.236/24;
            }
        }
    }
}
```

Security

Security can be divided into two major areas:

- Securing the infrastructure from attack, through firewalls, intrusion detection and prevention (IDP), deep packet inspection (DPI) and threat mitigation, for example protecting against viruses.
- Ensuring that the information (in this case, the videoconference) cannot be seen by unauthorized viewers or altered. A single client is protected using SSL technology, while all clients within a site are protected by IPSec. In addition, encryption can be enabled on some videoconferencing endpoints.

Firewall, IDP, DPI and Threat Mitigation

High throughput firewalls are essential to inspect and filter in-line traffic arriving at the data center from the unsecure Internet as well as from the enterprise VPNs in the case of service provider managed data centers.

In the HD videoconferencing videoconferencing datacenter framework, the Intrusion Detection and Prevention (IDP) and firewall functionality is provided by the SRX service gateway clusters in the network services layer.

For a more detailed discussion and configuration of IDP, refer to **IDP Policies** in the *JUNOS Software Security Configuration Guide* at **juniper.net/techpubs/software/junos-security/junos-security10.0/junos-security-swconfig-security/config-idp-policies-chapter.html**.

For firewall configuration details, refer to *SRX3600 Services Gateway Software Configuration Overview* at **juniper.net/techpubs/en_US/release-independent/junos/topics/concept/services-gateway-srx3600-software-configuration-overview.html**.

IDP and Firewall Configuration Snippet

```
## Download the latest signature DB from the CLI by running the following command
## (make sure the firewall is setup with a valid DNS server IP address as the download
## of signature DB requires name lookup
user@SRX-DC-1> request security idp security-package download full-update
## Verify the download status
user@SRX-DC-1> request security idp security-package download status

[edit security]
idp {
    ## Attach the pre-defined VoIP attack detection signature as the IDP policy.
    ## Note: This definition is only a snapshot. Add other application rules as per
    ## the application and server needs
    idp-policy VOIP-INSPECTION {
        rulebase-ips {
            rule 1 {
                match {
                    attacks {
                        ## Inspecting all VoIP signatures for any SRC and any DST
                        predefined-attack-groups “[Recommended]VOIP   - All”;
                    }
                }
                then {
                    action {
                        ## use pre-defined mitigation action when an attack is detected
                        recommended;
                    }
                    ## send out notifications in case an attack is detected
                    notification;
                }
            }
        }
    }
    ## set this as the active policy
    active-policy VOIP-INSPECTION;
}

##
## Firewall configuration
##
[edit policy-options]
## Define all the end-point and Data center related prefixes
prefix-list BigCo-SOHO-Site-1 {
    192.168.40.0/24;
}
prefix-list BigCo-Remote-Site-1 {
    192.168.41.0/24;
}
prefix-list BigCo-Campus-1 {
    192.168.51.0/24;
    192.168.52.0/24;
}
prefix-list dma-Managed-Srvs {
    10.12.12.0/24;
}
## Configure a firewall filter policy to allow traffic (signaling and media) to & from
## the various customer end-points to the data center video infrastructure
firewall {
    filter hdvc-through {
        ## Define separate term for each customer. This allows us for better tracking,
        ## logging and attaching distinctive actions if necessary.
        term BigCo {
            from {
                source-prefix-list {
                    BigCo-SOHO-Site-1;
                    BigCo-Remote-Site-1;
                }
                destination-prefix-list {
                    dma-Managed-Srvs;
                }
            }
            then {
                count hdvc-BigCo-through;
                forwarding-class hdvc-BigCo;
                accept;
            }
        }
        ## Define separate term for the BigCo customer – for reasons specified above
        term BigCo {
            from {
                source-prefix-list {
                    BigCo-Campus-1;
                }
                destination-prefix-list {
                    dma-Managed-Srvs;
                }
            }
            then {
                count hdvc-BigCo-through;
                forwarding-class hdvc-BigCo;
                accept;
            }
        }
    }
}
```

IPSec VPN Termination

Customers accessing the data center from SOHO and branch offices typically use IPSec VPNs to connect securely to the data center. In the HD videoconferencing videoconferencing datacenter framework, the SRX devices in the edge services layer provide IPSec VPN functionality.

For secure access configuration details, refer to the *Junos® OS Security Configuration Guide Release 10.3* at **juniper.net/techpubs/en_US/junos10.3/information-products/topic-collections/security/software-all/security/index.html**.

IPSec VPN Configuration Snippet

```
[edit]
security {
    ## The IPSec tunnel is defined in three steps.
    ## Step-1 - define a key-exchange proposal, policy and gateway
    ## Step-2 – define IPSec proposal and policy.
    ## Step-3 - Bind the IKE and IPSec together in a VPN definition and attach to an
    ## interface
    ike {
       ## The following IKE proposal should be matched up with the IPSec tunnel
       ## termination proposal at CPE in the customer’s site
        proposal video-dc-ike-proposal {
            authentication-method pre-shared-keys;
            ## Define the Diffie-Hellman group to use
            dh-group group1;
            authentication-algorithm md5;
            encryption-algorithm aes-128-cbc;
            lifetime-seconds 86400;
        }
        policy ike-policy1 {
            mode main;
            proposals video-dc-ike-proposal;
            pre-shared-key ascii-text “$9$/2i.AuBcyKxNbIENbs2GU/CtuIESre”; ## SECRET-DATA
        }
        ## Define the IPSec tunnel end-point. In this case, the CPE at the BigCo SOHO
        gateway BigCo-SOHO-Site-1 {
            ike-policy ike-policy1;
            address 192.168.40.1;
            external-interface ge-0/0/2.0;
        }
        ## Define the IPSec tunnel end-point. In this case, the CPE at the BigCo remote
        gateway BigCo-remote-Site-1 {
            ike-policy ike-policy1;
            address 192.168.41.1;
            external-interface ge-0/0/2.0;
        }
    }
    ## Define the VPN tunnel by associating the gateway and the IPSec definition
    ipsec {
        proposal video-dc-ipsec-proposal {
            protocol esp;
            authentication-algorithm hmac-md5-96;
            encryption-algorithm aes-128-cbc;
        }
        policy video-dc-ipsec-policy1 {
            proposals video-dc-ipsec-proposal;
        }
        ## Define the VPN tunnel for the SOHO site-1 at BigCo
        vpn BigCo-SOHO-Site-1-VPN {
            bind-interface st0.17;
            ike {
                gateway BigCo-SOHO-Site-1;
                ipsec-policy video-dc-ipsec-policy1;
            }
            establish-tunnels immediately;
        }
        ## Define the VPN tunnel for the remote site-1 at BigCo
        vpn BigCo-remote-Site-1-VPN {
            bind-interface st0.17;
            ike {
                gateway BigCo-remote-Site-1;
                ipsec-policy video-dc-ipsec-policy1;
            }
            establish-tunnels immediately;
        }
    }
}
```

Polycom DMA – Juniper SRC Integration

This section provides steps on how to set up Polycom's Distributed Management Application (DMA) to communicate with Juniper's SRC. This interface uses SOAP for communication.

1. Using a web browser, log in to the Polycom DMA administrator console. In this case, we are logging in to the **DMA 10.12.12.248** using administrator credentials. **https://10.12.12.248:8443/dma7000/**

2. At the **Configuration** menu, click **Director**, and then click **Service Resource Controller**, as shown in the following pane.

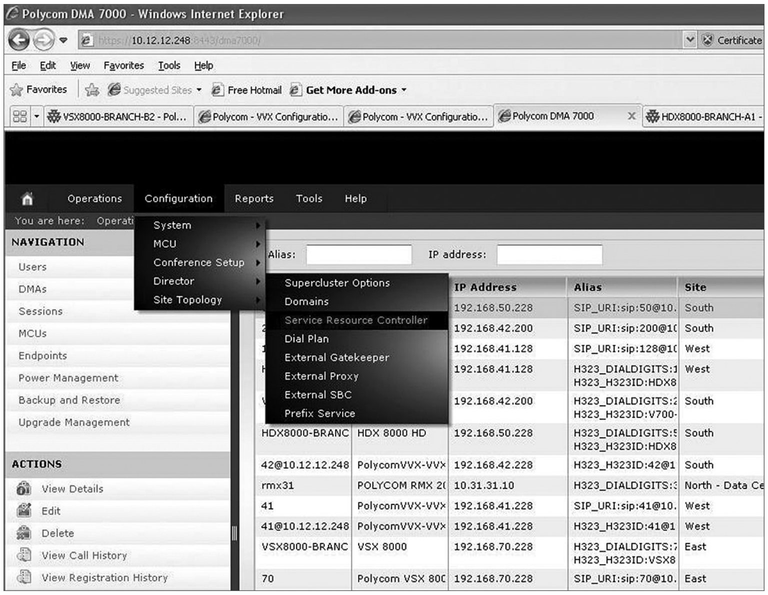

		IP Address	Alias	Site
		192.168.50.228	SIP_URI:sip:50@10.	South
2		192.168.42.200	SIP_URI:sip:200@1(	South
1		192.168.41.128	SIP_URI:sip:128@1(	West
H		192.168.41.128	H323_DIALDIGITS:1 H323_H323ID:HDX8	West
V		192.168.42.200	H323_DIALDIGITS:2 H323_H323ID:V700-	South
HDX8000-BRANC	HDX 8000 HD	192.168.50.228	H323_DIALDIGITS:5 H323_H323ID:HDX8	South
42@10.12.12.248	PolycomVVX-VVX	192.168.42.228	H323_H323ID:42@1	South
rmx31	POLYCOM RMX 2(	10.31.31.10	H323_DIALDIGITS:3	North - Data Ce
41	PolycomVVX-VVX	192.168.41.228	SIP_URI:sip:41@10.	West
41@10.12.12.248	PolycomVVX-VVX	192.168.41.228	H323_H323ID:41@1	West
VSX8000-BRANC	VSX 8000	192.168.70.228	H323_DIALDIGITS:7 H323_H323ID:VSX8	East
70	Polycom VSX 80(	192.168.70.228	SIP_URI:sip:70@10.	East

1. The following information for the SRC must be provided to ensure a proper integration.

    a. SRC IP address

    b. Port number – This will be **8088**

    c. The login, password credentials for the SRC SOAP gateway.

    d. The Subscriber URI – This is a sample query string, which helps validate that the communication worked fine. In the example below, this string is **`ip:ipAddress=192.168.41.128`**

2. Ensure that the **Enable integration to Juniper Networks** is checked.

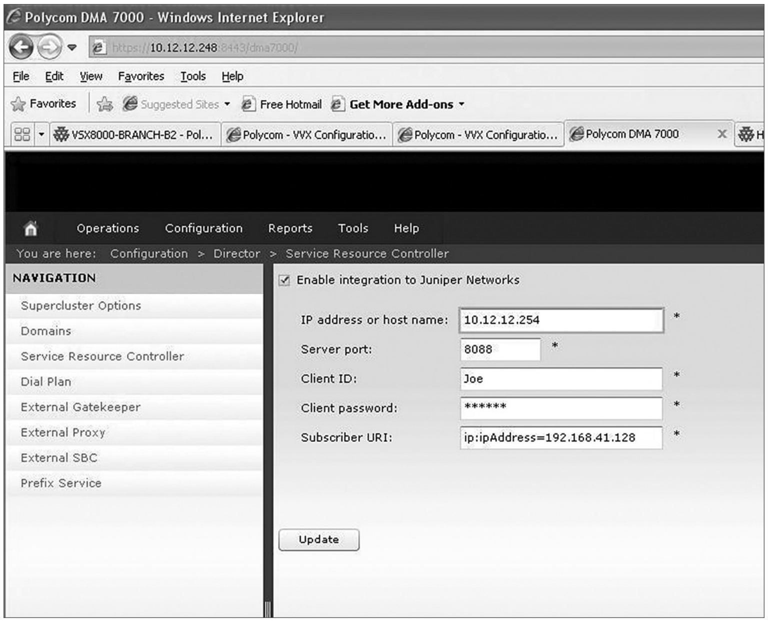
Polycom DMA 7000 - Windows Internet Explorer
https://10.12.12.248:8443/dma7000/
File Edit View Favorites Tools Help
Favorites
Suggested Sites
Free Hotmail
Get More Add-ons
VSX8000-BRANCH-B2 - Pol...
Polycom - VVX Configuratio...
Polycom - VVX Configuratio...
Polycom DMA 7000
H
Operations
Configuration
Reports
Tools
Help
You are here: Configuration > Director > Service Resource Controller
NAVIGATION
Supercluster Options
Domains
Service Resource Controller
Dial Plan
External Gatekeeper
External Proxy
External SBC
Prefix Service
Update

☑ Enable integration to Juniper Networks

IP address or host name: 10.12.12.254 *

Server port: 8088 *

Client ID: Joe *

Client password: ****** *

Subscriber URI: ip:ipAddress=192.168.41.128 *

Chapter 7

Small Office/Home Office Reference Architecture

Part Two

Small Office/Home Office (SOHO) setups have gained quite a bit of popularity in the last few years due to the advent of high-speed broadband connectivity, low cost computing equipment and innovations in secure remote access and visual communication infrastructure. In certain geographies, enterprises and large corporations are encouraging telecommuting employees to become teleworkers which allows enterprises to save on physical infrastructure costs and address environmental factors. To enhance productivity of teleworkers, reliable and cost effective videoconferencing functionality has become essential.

Figure 7.1 depicts the SOHO broadband reference architecture for a DSL and Figure 7.2 depicts a cable network.

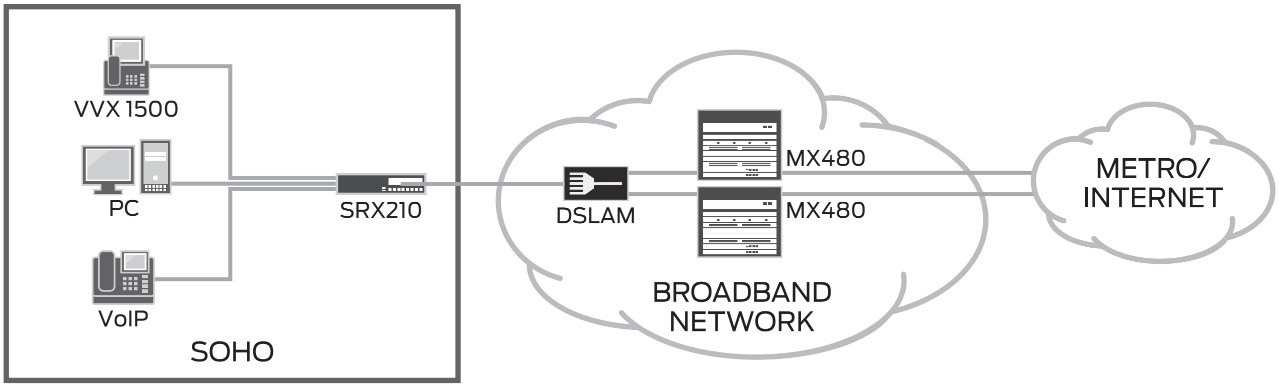

Figure 7.1 SOHO/broadband reference architecture (DSL)

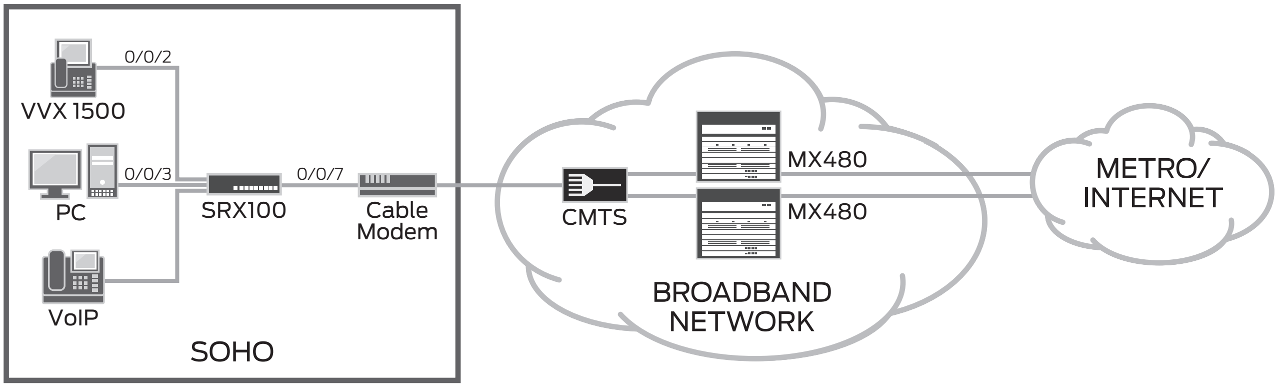

Figure 7.2 SOHO/broadband reference architecture (cable)

The modern SOHO setup typically connects a few PCs and teleconferencing equipment using the residential network (and Internet) to the high-speed broadband network. The typical network is as follows:

- A small videoconferencing system such as Polycom's VVX1500 connects to the SRX Series services gateway over a 100 Mbps Fast Ethernet link. The VVX is configured to receive its IP address automatically using DHCP from a local server running on the SRX and uses the SRX's interface IP address as its gateway. For more information on the VVX 1500, see *Appendix D: Polycom Products on page 176.*

- The SRX connects to the broadband network, using either an integrated DSL modem or an external gateway. The last mile can either be DSL based access or cable based. If DSL is used, Juniper's SRX210 CPE router with a WAN interface option (ADSL2 Annex-A mini-PIM port) can serve as the residential gateway device. This allows the SRX to serve as both a DSL modem and a security router in the SOHO premise and helps to boost security while avoiding the associated dual Network Address Translation (NAT) complexities. The ISP's residential network terminates the DSL modem on to a Digital Subscriber Line Access Multiplexer (DSLAM). This interface uses a PPPoA configuration with login credentials to negotiate and set up IP connectivity to the public Internet. For further details concerning PPPoA configuration, refer to *Configuring the ADSL Interface on SRX Series Services Gateway (CLI)*, at **juniper.net/techpubs/software/junos-security/junos-security10.2/junos-security-swconfig-interfaces-and-routing/jd0e22823.html**.

In a cable based last mile, connection from the coax to the customer premise equipment is provided via Ethernet interface through a cable modem. The modem connects to a Cable Modem Termination Device (CMTS) on the residential network. In this setup, an SRX100 Series Services Gateway can be utilized as the secure router between the customer premises LAN segment and the Internet.

NOTE: The cable based residential access model will be used as a reference example for the rest of this chapter. This assumes that the WAN link is represented with an Ethernet interface, and as such all configurations (code snippets) reflect this.

- The DSLAM or the CMTS connects to the Internet using a multi-services edge router. Optionally, the DSLAM/CMTS can be multi-homed for redundancy purposes.

On the customer premise side, the secure services gateway provides secure IPSec VPN connection to the corporate location or videoconferencing datacenter. Additionally, the SRX Series also acts as a security router protecting the SOHO equipment connected to its interfaces from the threats of Internet.

Ideally, QoS should be used across the end-to-end connection. However, the last mile may not be QoS-enabled and a congestion-free metro access network cannot be guaranteed. Consequently, end-to-end call control still presents a challenge in the SOHO environment.

SOHO deployments can be designed to use desktop end-point equipment that does not have high bandwidth needs. Polycom's VVX series is an example of this. In this case, the videoconferencing calls will be treated as best effort traffic. No special resources or admission control provisioning is required by the service provider, and therefore the HD videoconferencing service is provided with enhanced but not assured call experience.

Implementation

This section shows how to configure the SRX to support broadband access with HD videoconferencing terminals. It shows how to establish secure connectivity using IPSec, implement a local DHCP server with Network Address Translation (NAT), set up security zones, provision QoS and define interfaces.

Configuring IPSec

The SRX provides IPSec VPN connectivity for the endpoint to communicate with the videoconferencing datacenter as well as with endpoints at other sites during a point-to-point call. Figure 7.3 illustrates IPSec connectivity in a SOHO environment.

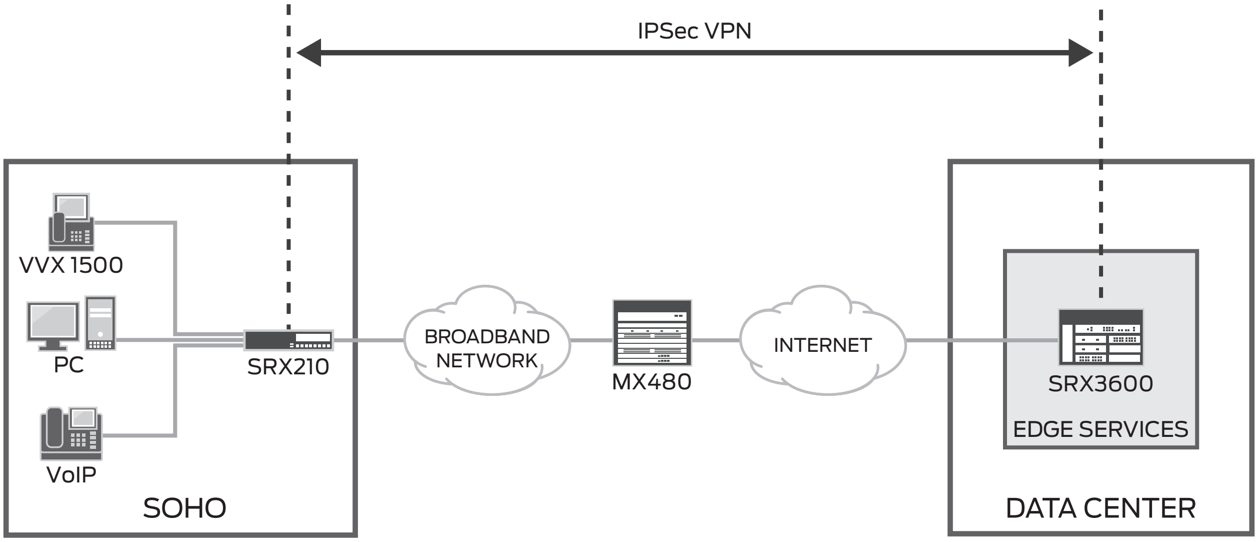

Figure 7.3 IPSec connectivity and SOHO functional overview

The following code snippet shows how to configure the SRX to implement IPSec VPN.

SRX Configuration Snippet for IPSec VPN Origination

```
[edit]
security {
   ## The IPSec tunnel is defined in three steps.
   ## Step-1 - Define a key-exchange proposal, policy and gateway.
   ## Step-2 - Define IPSec proposal and policy.
   ## Step-3 - Bind the IKE and IPSec together in a VPN definition and attach to an
   ## interface
   ike {
      ## The following IKE proposal should be match-up with the IPSec tunnel
      ## termination proposal at the SRX in the videoconferencing datacenter and at all
      ## other sites which this premise will communicate with
       proposal hdvc-ike-proposal {
          authentication-method pre-shared-keys;
          ## Define the Diffie-Hellman group to use
          dh-group group1;
          authentication-algorithm md5;
```

```
         encryption-algorithm aes-128-cbc;
         lifetime-seconds 86400;
      }
      policy ike-policy1 {
         mode main;
         proposals hdvc-ike-proposal;
         pre-shared-key ascii-text “$9$/2i.AuBcyKxNbIENbs2GU/CtuIESre”; ## SECRET-DATA
      }
      ## Define the IPSec tunnel end-point. In this case, the SRX at the data center
      gateway hdvc-dc-Site-1 {
         ike-policy ike-policy1;
         address 10.10.10.1;
         external-interface fe-0/0/7.0;
      }
   }
   ## Define the VPN tunnel by associating the gateway and the IPSec definition
   ipsec {
      proposal hdvc-ipsec-proposal {
         protocol esp;
         authentication-algorithm hmac-md5-96;
         encryption-algorithm aes-128-cbc;
      }
      policy hdvc-ipsec-policy1 {
         proposals hdvc-ipsec-proposal;
      }
      ## Define the VPN tunnel for the videoconferencing datacenter
      vpn hdvc-steer-VPN {
         ike {
            gateway hdvc-dc-Site-1;
            ipsec-policy hdvc-ipsec-policy1;
         }
         establish-tunnels immediately;
      }
   }
}
```

Configuring Security Zones

Security zones on the SRX are configured to permit all traffic to and from the physical port to which the VVX is connected. This ensures that the video endpoint can communicate (register, call signaling) with the centralized SIP proxy/ H.323 gateway which resides in the videoconferencing datacenter, as well as with other endpoints at other sites. In the SOHO scenario, we recommend configuring two security zones: Trust and Untrust. The Trust security zone is used for all customer premises devices including PCs and videoconferencing endpoints, while the Untrust zone is used on the WAN interface. This decision was made based on the assumption that there is a small number of end devices at the customer premises and additional security is not required. If it is desired to implement a more stringent security model to protect the HD videoconferencing service infrastructure from external threats, including those originating from the PC hosts at SOHO locations, see *Chapter 8, Building the Enterprise Site on page 93.*

SRX Configuration Snippet to Manage the VVX Connected Interface

```
[edit interfaces]
fe-0/0/2 {
   description “connected to VVX-1”;
   unit 0 {
      family inet {
         ## This interface will serve as the default gateway for the end-point,
         ## and hence the IP address must match the IP subnet being used to assign
         ## the end-point’s IP via DHCP
         address 192.168.40.1/24;
      }
   }
}

[edit security]
zones {
   ## All devices on the LAN side are configured in the “trust” zone and can be assigned
   ## IP addresses either statically or dynamically
   security-zone trust {
      tcp-rst;
      address-book {
         ## Configure an address book entry for the end-point’s subnet - to be used
         ## in the IPSec-VPN configuration
         address vvx-devices 192.168.40.0/24;
      }
      host-inbound-traffic {
         system-services {
            all;
         }
         protocols {
            all;
         }
      }
      interfaces {
        ## Configuring the interface connecting the VVX
        ## to receive a DHCP discover requests
         fe-0/0/2.0 {
            host-inbound-traffic {
               system-services {
                  dhcp;
               }
               protocols {
                  all;
               }
            }
         }
      }
   }
   security-zone untrust {
      screen untrust-screen;
      address-book {
         ## Configure an address book entry for the Polycom video signaling
         ## equipment subnet to be used in the IPSec-VPN configuration
         address managed-hdvc-at-dc 10.12.12.0/24;
         ## Include all other prefixes for sites which host the video endpoints
         address BigCo-remote-site-1 192.168.41.0/24;
         address BigCo-remote-site-2 192.168.44.0/24;
         ## create an address set which includes all end-points and devices that this
         ## site has to communicate with
         address-set managed-hdvc-devices {
            address managed-hdvc-at-dc;
            address BigCo-remote-site-1;
```

```
                address BigCo-remote-site-2;
            }
        }
        interfaces {
            ## Configuring the interface connecting to the WAN
            ## to receive a DHCP offers from downstream
            fe-0/0/7.0 {
                host-inbound-traffic {
                    system-services {
                        dhcp;
                        ike;
                    }
                    protocols {
                        all;
                    }
                }
            }
        }
    }
## These policies set the rules of interaction between the devices in the "trust" zones
policies {
    ## permit all types of traffic between devices in the "trust" zone
    from-zone trust to-zone trust {
        policy default-permit {
            match {
                source-address any;
                destination-address any;
                application any;
            }
            then {
                permit;
            }
        }
    }
    ## Setup policies to reach the data center as well as all other sites that this
    ## premise needs to communicate with
    from-zone trust to-zone untrust {
        ## This policy matches traffic from the end-point IP addresses destined to the
        ## video signaling devices in the data center, and directs this traffic into
        ## IPSec VPN tunnel
        policy hdvc-endpoint-steer {
            match {
                ## Refer to address-book entries that were defined earlier
                source-address vvx-devices;
                destination-address managed-hdvc-devices;
                application any;
            }
            then {
                permit {
                    tunnel {
                        ## Associate the tunnel reference which defines the tunnel
                        ## characteristics
                        ipsec-vpn hdvc-steer-VPN;
                        ## Provide a reference to the policy that will be used for
                        ## incoming traffic on this tunnel
                        pair-policy steer-to-hdvc-endpoint;
                    }
                }
            }
        }
        ## Traffic that does not match the tunnel characteristics from above will be
        ## handled as per this policy rule – which is currently set to permit everything
        policy default-permit {
            match {
                source-address any;
                destination-address any;
                application any;
            }
            then {
                permit;
            }
        }
    }
    ## Setup policies to handle the return traffic from the videoconferencing datacenter and
    ## the other sites which might have video endpoints
    from-zone untrust to-zone trust {
        ## This policy matches traffic from the video signaling devices in the data
        ## center), destined to the end-point IP addresses
        policy steer-to-hdvc-endpoint {
            match {
                ## Refer to address-book entries that were defined earlier
                source-address managed-hdvc-devices;
                destination-address vvx-devices;
                application any;
            }
            then {
                permit {
                    tunnel {
                        ## Associate a tunnel definition reference which defines the
                        ## tunnel’s characteristics
                        ipsec-vpn hdvc-steer-VPN;
                        ## Provide a reference to the policy that will be used for
                        ## outgoing traffic on this tunnel
                        pair-policy hdvc-endpoint-steer;
                    }
                }
            }
        }
        ## Traffic that does not match the tunnel characteristics from above will be
        ## handled as per this policy rule – which is currently set to permit everything
        policy default-permit {
            match {
                source-address any;
                destination-address any;
                application any;
            }
            then {
                permit;
            }
        }
    }
}
```

The SRX serves as a DHCP local server and NAT gateway for connected video end-points. It receives a public IP address from the ISP and in turn provides NAT to all traffic from the video endpoint to this IP address. The DHCP local server is configured to match the device hardware addresses of the video endpoint with a pre-defined IP address. This is important because the Polycom DMA controller in the data center is provisioned to recognize end-points using these specific IP addresses. Figure 7.4 illustrates an overview of the DHCP address assignment mechanism.

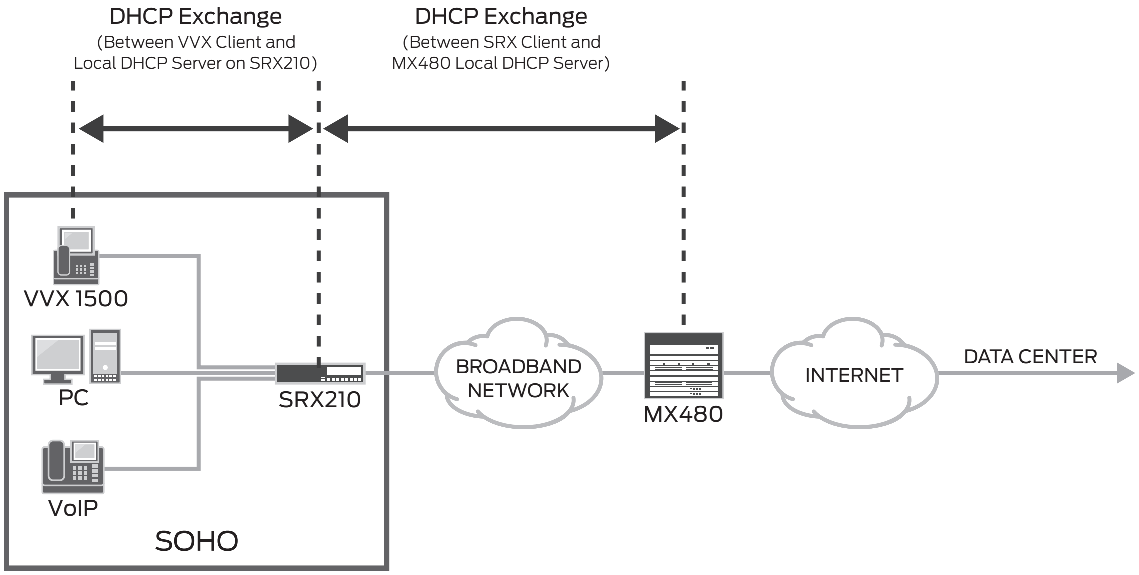

Figure 7.4 DHCP address assignment overview

The following code snippet shows how to configure the SRX to provide pre-provisioned IP addresses.

SRX Configuration Snippet to Assign IP Sddress for the VVX Device

```
## Configure the video end-point interfaces
[edit system services]
## It is important that the video end-points be pre-provisioned with IP addresses that
## the video infrastructure in the data center are aware of. One way to do this, is by
## ensuring that the end-points get assigned dynamic IP address via an hardware address
## match at the gateway
dhcp {
   ## Hardware address of the VVX device
   static-binding 00:e0:db:07:e3:df {
      fixed-address {
         192.168.40.228;
      }
      router {
         ## the router gateway is the interface's IP address
         192.168.40.1;
      }
   }
}
## Configure the WAN Interface to receive a public IP address from the
## ISP's DHCP server in the residential network
[edit interfaces]
fe-0/0/7 {
   description "connected to the WAN-Edge layer";
   unit 0 {
      family inet {
         dhcp {
            retransmission-attempt 3;
            retransmission-interval 6;
         }
      }
   }
}
```

Configuring for NAT

Configure the SRX device to NAT all traffic from the video endpoint's private IP address to the DHCP-assigned public IP address on the egress interface and vice versa. The following code snippet depicts this configuration.

SRX NAT Configuration Snippet

```
[edit security nat]
## Configure source based NAT to translate the end-point IP address to the public
## IP address which is dynamically assigned to the WAN interface
source {
   rule-set VVX-Interface-NAT {
      from zone trust;
      to zone untrust;
      rule private_net {
         match {
           ## List the range of private IP addresses which will require NAT
           ## translation
            source-address [ 192.168.40.0/24 ];
         }
         then {
            ## Specify the action to translate to the interface's IP address
            source-nat {
               interface;
```

```
                }
            }
        }
    }
}
```

Configuring IDP

The security router is also configured to protect devices in the trust zones from attacks originating from the unsecure Internet. These options help combat attacks, such as IP address sweeps, port scans, DOS attacks, ICMP, UDP, floods and many others.

SRX Configuration Snippet for Intrusion Detection

```
[edit security]
screen {
    ## The SOHO setup needs to be secured against attacks from the unsecure internet.
    ## These should typically include the following threats
    ids-option untrust-screen {
        icmp {
            ping-death;
        }
        ip {
            source-route-option;
            tear-drop;
        }
        tcp {
            syn-flood {
                alarm-threshold 1024;
                attack-threshold 200;
                source-threshold 1024;
                destination-threshold 2048;
                queue-size 2000;
                timeout 20;
            }
            ## prevent SYN attacks combined with IP spoofing
            land;
        }
    }
}

## apply the screen profile to the untrust zone
[edit security zones]
security-zone untrust {
    screen untrust-screen;
}
```

Provisioning QoS

Although it is challenging to guarantee SLA for calls which will be used over the public Internet, as a best practice, the SRX gateway is provisioned to apply static QoS on the video endpoint traffic. This is achieved by attaching a filter on the interface to which the endpoint is connected and marking ingress traffic on this interface with assured forwarding DSCP code points. Since the traffic is traversing the Internet, we can only guarantee best effort delivery and thus the traffic will be classified as Best Effort in this example. The following code snippet shows where to attach this filter. For information concerning the QoS forwarding class, which the filter definition refers to, see *Chapter 9, Implementing Quality of Service on page 115.*

SRX Configuration Snippet for Static QoS

```
[edit firewall]
## Define a filter which will enable us to associate static QoS for all traffic
## originated by the video end-point. This will be associated at the interface to which
## the end-point connects to
filter VVX-MC-Bronze-Tier {
    term All-VVX-Traffic {
        then {
            forwarding-class MC-Bronze;
            accept;
        }
    }
}
## attach the filter, QoS settings to the interface connecting the video end-points
[edit interfaces]
fe-0/0/2 {
    description “connected to VVX-1”;
    unit 0 {
        family inet {
            address 192.168.40.1/24;
            filter {
                input VVX-Bronze-Tier;
            }
        }
    }
}
```

Configuring the MTU Interface

To ensure minimal packet loss and transit delays for videoconferencing traffic, architects must consider the size of the entire network's Maximum Transmission Unit (MTU) and the video endpoints must be configured accordingly. By default, on the VVX device, the MTU is set to 1260 bytes. The end-to-end MTU assessment must account for overhead added by IPSec VPN (~52 bytes), VLAN header (4 bytes) and L3VPN (4 bytes). If the packet size exceeds that MTU of any network link, it will be fragmented in two or more fragments. This is an undesirable treatment for HD videoconferencing traffic and will result in degraded quality. In the case of IPSec tunnels, packet fragmentation is absolutely not permissible.

The default value for MTU on the Fast Ethernet interface of the SRX router is 1500 bytes. Based on the end-to-end MTU calculation, this must be changed on both the interface connecting the video endpoint and the WAN uplink if required. Typically, the MTU size in cable and Asymmetric Digital Subscriber Line (ADSL) networks is even shorter, so special care should be taken to make sure that the length of the transmitted packets does not exceed the link's MTU.

To configure the MTU, use the following command.

SRX Configuration Snippet for Setting Interface MTU

```
[edit interfaces]
## set the MTU to the appropriate value
fe-0/0/2 {
   description “connected to VVX-1”;
   unit 0 {
      family inet {
         address 192.168.40.1/24;
         filter {
            input VVX-Bronze-Tier;
         }
         ## Example value setting. Must be changed as per local setup
         mtu 1492;
      }
   }
}
```

Troubleshooting

If you are having problems getting a DHCP IP address assigned to your residential gateway device, re-initiate the DHCP request using the following operational command.

SRX Configuration Snippet for DHCP Troubleshooting

```
user@SRX-210> request system services dhcp renew

## to verify the IP address that was assigned
user@SRX-210> show system services dhcp client
user@SRX-210> show system services dhcp client statistics
```

Once you have successfully acquired an IP address at the video end-point device, ensure that you have end-to-end connectivity between the endpoint and the video signaling equipment (Polycom's DMA, RMX) in the data center.

Video endpoints provide a connectivity test functionality which is accessible using the web interface of the endpoint. If problems still exist, proceed to troubleshooting the routers.

Troubleshooting Connectivity Issues

```
## ping the video end-point from the SRX gateway
user@SRX-210> ping 192.168.40.228
## ping the DMA from the SRX gateway
user@SRX-210> ping 10.12.12.248
```

If connectivity to the video equipment in the data center has not yet been established, verify connectivity to the public internet. This can be verified by pinging the default gateway that was assigned to the SRX gateway by the ISP's DHCP server. If there are no connectivity issues, then inspect the IPSec VPN configuration on the SRX gateway and the tunnel termination configuration in the SRX cluster at the videoconferencing datacenter.

If there are any MTU related problems, ping from the endpoint using the ICMP packets with the packet length set to the size of the MTU configured on the device.

Chapter 8

Building the Enterprise Site

Part Two

This chapter provides reference network architecture for deploying videoconferencing in an enterprise site. The focus of the chapter is to cover the provisioning and configuration of Juniper's network equipment.

Enterprise sites generally fall into three categories:

- **Remote offices** (small sites), which do not have any WAN redundancy.
- **Branch offices** (medium sites), which may have multiple WAN connections.
- **Campuses** (large sites), which have redundant WAN routers.

These sites share many common characteristics. Of course, client devices are connected to a LAN switch in a wiring closet. Branch and campus sites may have a second tier of LAN switches to aggregate traffic, which ultimately passes through a security gateway and router before reaching the WAN.

The validated network includes at least one of each type of the three enterprise sites. Because all devices pertaining to the enterprise sites run on Junos, the sites are configured similarly and therefore are described only once.

The following table summarizes the sites used during validation.

Table 8.1 Enterprise Sites

Sites	Ethernet Switch	WAN Access and Security
Remote Office	EX3200	SRX240 (one WAN link)
Branch Office	EX4200 (standalone)	SRX650 (two WAN links)
Campus	EX4200 (virtual chassis)	WAN: Two MX240s (each with one WAN link) Security: SRX3600

In addition to the number of WAN links, other potential differences between the smaller and larger sites include:

- **Address assignment:** For smaller sites, clients receive their IP addresses from a centralized DHCP server. To support this, DHCP proxy should be configured on the enterprise site router. In a large site such as a campus, most devices receive their addresses from a DHCP server located on the premises rather than using a centralized DHCP server. Therefore, the DHCP proxy configuration is not required.
- The core network may optionally be extended to larger site. In other words, the PE router may reside on the customer premises, and the MPLS-based L3VPN terminates at this location. This implementation is described in *Chapter 6, Building the HD Videoconferencing Data Center on page 55*.
- **LAN resiliency:** The access layer switches may be dual-homed to redundant aggregation nodes, which are in turn connected to the WAN gateways. This can be done by extending the virtual chassis to the EX4200s on the floors or using Rapid Spanning Tree between switches, and will not be discussed further.

Branch Office/Remote Office Implementation

This section presents branch office implementation and configuration. Figure 8.1 depicts the deployed branch office reference topology, which includes the Polycom HDX 9000, VSX 7000 and VVX 1500 videoconferencing clients as well as VoIP phones and PCs. These devices are connected to Juniper's EX4200 Ethernet Switch. Connection from the WAN to the branch office uses Juniper's SRX650 Services Gateway as the customer premises (CPE) router. Redundant upstream WAN links are configured.

The remote office is configured similarly to the branch, except that different models are used, and the remote office has only one WAN connection. Aside from using different physical ports and IP addresses, the configuration is identical.

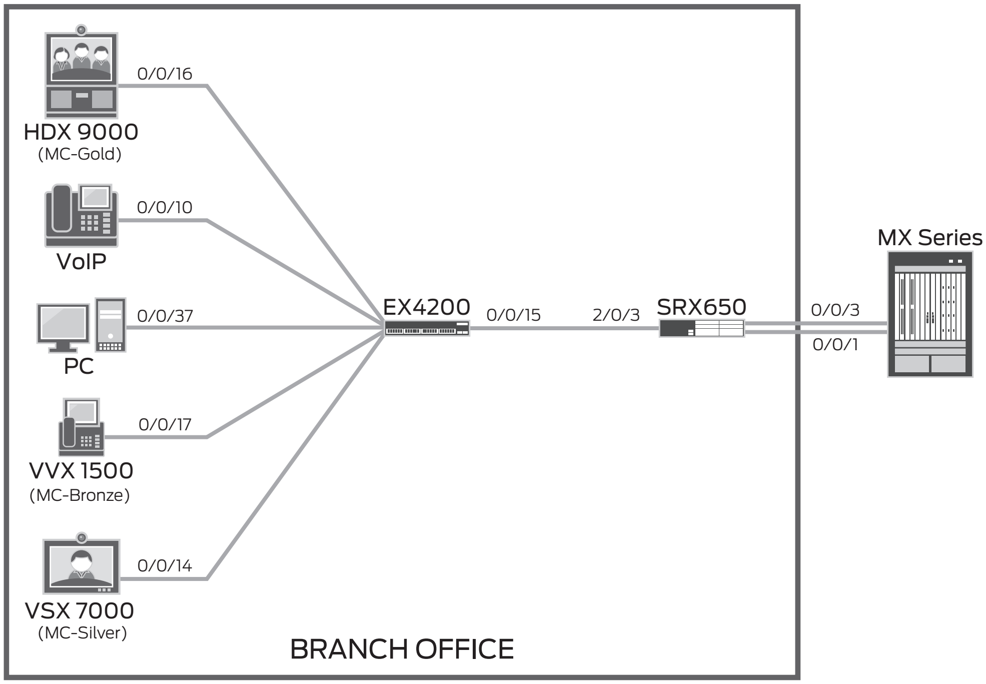

Figure 8.1 Branch office reference topology

The primary implementation and configuration guidelines discussed in this section include:

- Connecting the Videoconferencing Endpoints
- Configuring the SRX (CPE Router)
- Configuring the PE Router

Connecting the Videoconferencing Endpoints

Control traffic is carried on the same VLAN as the corresponding media traffic, so that all the video endpoints are able to communicate (register, call signaling) with the centralized SIP registrar or H.323 gatekeeper. VLAN 1001 is used for voice traffic while VLAN 1002 is used for PCs, servers, printers and data traffic.

On the WAN, the three video VLANs (100, 200, 300) are aggregated together on to a single video VLAN (600). This is a design choice to preserve VLAN addresses, which is more likely to be used if videoconferencing service is provided by a network provider supporting multiple customers. Instead, an enterprise may choose to use the same VLAN values on the WAN as are used on the LAN.

Figure 8.2 summarizes the deployed VLAN scheme.

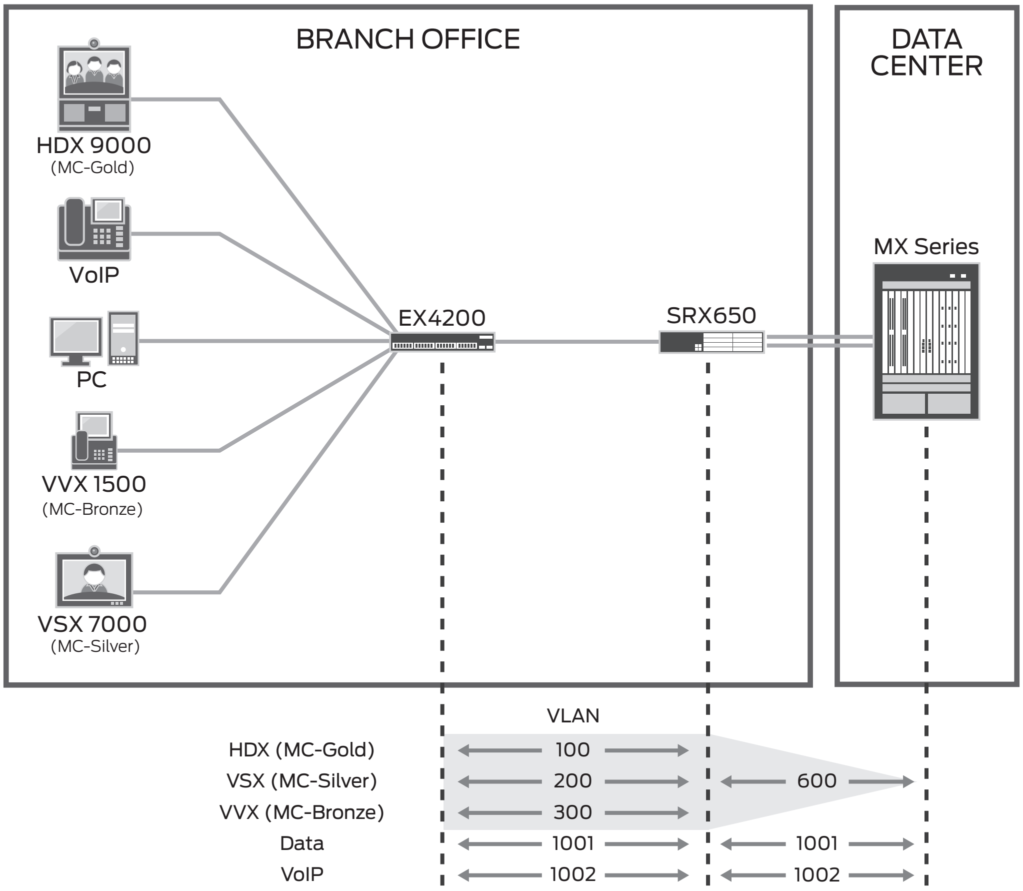

Figure 8.2 VLAN detailed scheme

Configuring the EX4200 to Connect Video Endpoints

```
## Configure the uplink trunk interface towards SRX, which carries HD videoconferencing, PC, SERVER
## etc. traffic
EX-179# run show configuration
```

```
## Configure the interface which connects to the HDX end-point
interfaces {
    ge-0/0/16 {
        description “ connected to HDX9000 video end-point “;
        unit 0 {
            family ethernet-switching;
        }
    }
}
## Configure the interface that connects the SRX
interfaces {
    ge-0/0/15 {
        unit 0 {
            family ethernet-switching {
            port-mode trunk;
            }
        }
    }
}
## Configure the interface that connects the VSX video end points
## By default auto negotiation is enabled on the interface.
EX-179# run show configuration
interfaces {
    ge-0/0/14 {
        unit 0 {
            family ethernet-switching;
        }
    }
}
## Configure the interface that connects the VVX video end points.
interfaces {
    ge-0/0/17 {
        unit 0 {
            family ethernet-switching;
        }
    }
}
## Configure the interface that connects the PC, SERVERS etc.
interfaces {
    ge-0/0/37 {
        unit 0 {
            family ethernet-switching;
        }
    }
}
## Configure the interface that connects the VOIP endpoints.
interfaces {
    ge-0/0/40 {
        unit 0 {
            family ethernet-switching;
        }
    }
}
## Configure the service vlan 300 for the interfaces that VVX video endpoints are
## connected  and SRX650 is connected.
vlans {
    MC-Bronze {
        vlan-id 300;
        interface {
            ge-0/0/15.0;
            ge-0/0/17.0;
        }
    }
}
## Configure the service vlan 200 for the interfaces that VSX video endpoints are
## connected and srx is connected.
vlans {
   MC-Silver {
## Associate the Service Vlan-Id for this group
      vlan-id 200;
      interface {
         ge-0/0/14.0;
         ge-0/0/15.0;
      }
   }
}
## Configure the service vlan 100 for the interfaces that HDX video endpoints are
## connected and srx is connected.
vlans {
   MC-Gold {
## Associate the Service Vlan-Id for this group
      vlan-id 100;
      interface {
         ge-0/0/15.0;
         ge-0/0/16.0;
      }
   }
}
## Configure the vlan 1001 for the interfaces that pc, server etc. traffic.
vlans {
   DATA {
## Associate the Service Vlan-Id for this group
      vlan-id 1001;
      interface {
         ge-0/0/15.0;
         ge-0/0/37.0;
      }
   }
}
## Configure the vlan 1002 for the interfaces that carries VOIP traffic.
vlans {
   VOIP {
## Associate the Service Vlan-Id for this group
      vlan-id 1002;
      interface {
         ge-0/0/15.0;
         ge-0/0/40.0;
      }
   }
}
```

Configuring the SRX (CPE)

VLAN 600 is configured on the WAN interface for the videoconferencing service.

Configuring the WAN Interface on SRX

```
## Interfaces connects to the PE router WEST
SRX650-BO-1# show
interfaces {
    ge-0/0/3 {
        vlan-tagging;
        description CONNECTED-TO-PE-WEST;
       ## Used to deliver HD videoconferencing service
        unit 600 {
            vlan-id 600;
            family inet {
                address 10.8.1.50/30;
            }
        }
    }
}
## Interfaces connects to the PE router SOUTH
interfaces {
    ge-0/0/1 {
        vlan-tagging;
        description CONNECTED-TO-PE-SOUTH;
        ## Used to deliver HD videoconferencing service
        unit 600 {
            vlan-id 600;
            family inet {
                address 10.9.1.50/30;
            }
        }
    }
}

interfaces {
    ge-0/0/3 {
        vlan-tagging;
        unit 1002 {
            description CONNECTED-TO-VOIP;
            vlan-id 1002;
            family inet {
                address 10.8.1.53/30;
            }
        }
    }
}
interfaces {
    ge-0/0/3 {
        vlan-tagging;
        unit 1001 {
            description CONNECTED-TO-DATA;
            vlan-id 1001;
            family inet {
                address 10.8.1.57/30;
            }
        }
    }
}
## Interfaces used to connects the HDX Video Endpoints.
SRX650-BO-1# show
```

```
interfaces {
    ge-2/0/3 {
        vlan-tagging;
        unit 100 {
            description CONNECTED-TO-HDX;
            vlan-id 100;
            family inet {
                filter {
                    ## Filter says that all the traffic comes to this
                    ## vlan will be marked with SILVER forwarding class.
                    input MC-Gold;
                }
                address 192.168.50.1/24;
            }
        }
    }
}
## Interfaces used to connects the VSX Video Endpoints.
interfaces {
    ge-2/0/3 {
        vlan-tagging;
        ## Connected to video endpoints for SILVER SERVICE
        unit 200 {
            description CONNECTED-TO-VSX;
            vlan-id 200;
            family inet {
                filter {
                    ## Filter says that all the traffic comes to this
                    ##   vlan will be marked with SILVER forwarding class.
                    input MC-Silver;
                }
                address 192.168.70.1/24;
            }
        }
    }
}
## Interfaces used to connects the VVX Video Endpoints.
interfaces {
    ge-2/0/3 {
        vlan-tagging;
        description CONNECTED-TO-VVX;
        unit 300 {
            vlan-id 300;
            family inet {
                filter {
                    ## Filter says that all the traffic comes to this
                    ## vlan will be marked with BRONZE forwarding class.
                    input MC-Bronze;
                }
                address 192.168.51.1/24;
            }
        }
    }
}
## Interfaces used to connects the VOIP Terminals.
interfaces {
    ge-2/0/3 {
        vlan-tagging;
        ## Used to connect PCs, Servers etc
        unit 1001 {
            description CONNECTED-TO-DATA;
            vlan-id 1001;
            family inet {
                address 192.168.10.1/24;
```

```
         }
      }
   }
}
## Interfaces used to connects the DATA generated Terminals like servers etc.
interfaces {
   ge-2/0/3 {
      vlan-tagging;
      ## Used to connect VoIP terminals
      unit 1002 {
         description CONNECTED-TO-VOIP;
         vlan-id 1002;
         family inet {
            address 192.168.20.1/24;
         }
      }
   }
}
## The loopback interface will be used as the router-id
SRX650-BO-1# show lo0
unit 0 {
   family inet {
      address 10.100.1.187/32;
   }
}
```

Configuring the WAN Edge Layer

The SRX650, serving as the CPE, provides the security WAN routing/gateway functionality in the branch office setup. The router connects to the service provider's PE device, which is the starting point of the L3VPN to the data center. The SRX, which is provisioned to route all traffic to the PE device, uses OSPF to selectively propagate routes into the PE domain. OSPF is configured on all interfaces that must be reachable externally.

SRX configuration snippet for dynamic routing

```
## setup OSPF on all interfaces which need to be reachable outside of the branch office
[edit protocols]
ospf {
   area 0.0.0.0 {
      interface lo0.0 {
         passive;
      }
      interface ge-0/0/1.600;
      interface ge-0/0/3.600;
      interface ge-0/0/3.1001;
      interface ge-0/0/3.1002;
      interface ge-2/0/3.100;
      interface ge-2/0/3.200;
      interface ge-2/0/3.300;
      interface ge-2/0/3.1001;
      interface ge-2/0/3.1002;
   }
}

[edit routing-options]
## set the router-id for to use
router-id 10.100.1.187;
## set the AS number as per deployment
autonomous-system 100;
```

Configuring DHCP Relay for Videoconference Endpoints

The DHCP relay on the SRX processes the DHCP discover requests from the client, and in turn communicate with the DHCP proxy server on the Junos PE router that connects this site to the core. This DHCP proxy communicates with the centralized DHCP server in the videoconferencing datacenter to receive and offer an IP address for the endpoint(s). See Figure 8-3 for an overview of DHCP implementation in the branch office.

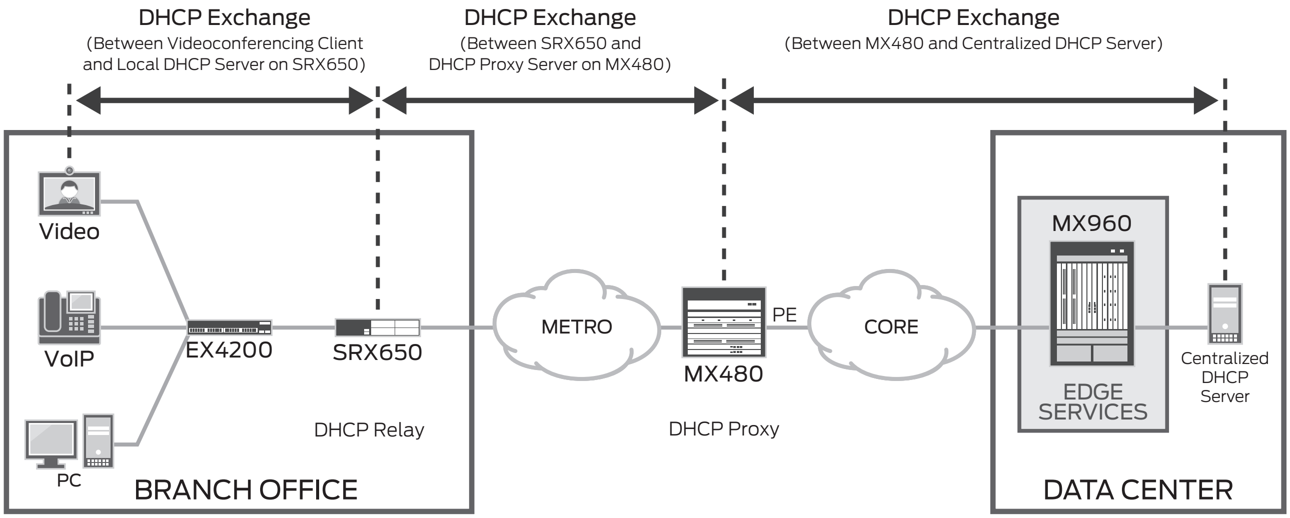

Figure 8.3 Implementing DHCP in the branch office

The HD videoconferencing video infrastructure in the data center uses the endpoint IP addresses to determine what kind of service tiers to apply for calls. To facilitate this, the endpoints must have assigned IP addresses from specific ranges based on some sort of endpoint identifier. This setup uses the MAC address as the identifier that is communicated to the data center using option-82 in the DHCP request. The following code snippet shows how to configure DHCP on the SRX.

```
## Configure the DHCP-relay functionality and specify the interfaces on which the SRX
## will accept DHCP requests
[edit forwarding-options]
helpers {
   bootp {
       description “ DHCP relay for the video end-points “;
## ensure that the requesting client’s MAC address is relayed to the DHCP server for an
## hardware address match
       dhcp-option82 {
          remote-id;
       }
## DHCP proxy server IP address on the PE router
       server 10.100.1.27;
       maximum-hop-count 16;
## will only accept DHCP requests on these interfaces
       interface {
          ge-2/0/3.100;
          ge-2/0/3.200;
          ge-2/0/3.300;
       }
   }
}
```

Configuring the Security Zone for the Endpoints

One of the key considerations in the branch office design is protecting the HD videoconferencing infrastructure from unsecure access. If voice, data and videoconferencing are controlled by different organizations, it may be necessary to configure security zones for each service, as shown in Figure. 8.4.

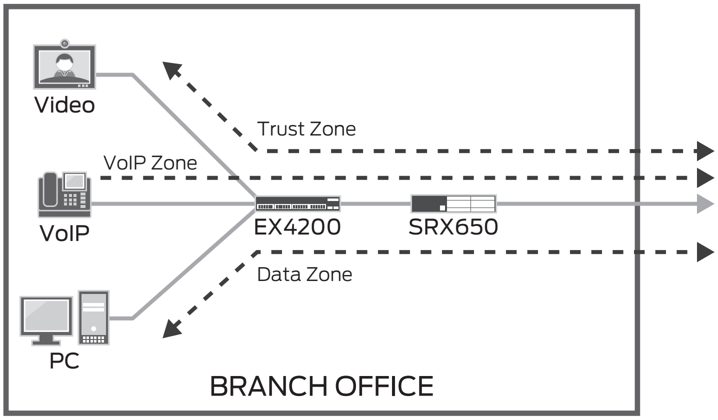

Figure 8.4 Configuring the security zone for endpoints

The design must include several security zones:

- **Trust Zone:** This zone includes the videoconference devices. This is based upon the assumption that all videoconferencing devices connecting to the trust zone are part of the control and administration domain of the service provider. Typically, only videoconferencing traffic will be sent to the trust zone.
- **Data Zone:** This zone includes the enterprise-control equipment such as PCs, servers and printers to the SRX650. This is based upon the assumption that all the devices connecting to the data zone are part of the control and administration domain of the local IT.
- **VOIP Zone:** This zone is used for all the voice traffic. This is based upon the assumption that all the devices connecting to the VoIP zone are part of the control and administration domain of the local IT.

```
SRX650-BO-1# show
security {
    zones {
## Security zone used for HD videoconferencing services only
        security-zone trust {
            host-inbound-traffic {
                system-services {
                    all;
                }
                protocols {
                    all;
                }
            }
            interfaces {
                ge-2/0/3.300 {
                    host-inbound-traffic {
                        system-services {
                            all;
```

```
                      }
                      protocols {
                         all;
                      }
                   }
                }
                ge-2/0/3.200 {
                   host-inbound-traffic {
                      system-services {
                         all;
                      }
                      protocols {
                         all;
                      }
                   }
                }
                ge-2/0/3.100 {
                   host-inbound-traffic {
                      system-services {
                         all;
                      }
                      protocols {
                         all;
                      }
                   }
                }
                lo0.0 {
                   host-inbound-traffic {
                      system-services {
                         all;
                      }
                      protocols {
                         all;
                      }
                   }
                }
                ge-0/0/3.600 {
                   host-inbound-traffic {
                      system-services {
                         all;
                      }
                      protocols {
                         all;
                      }
                   }
                }
             }
          }
## Security zone used for DATA services only
          security-zone DATA {
             interfaces {
                ge-2/0/3.1001 {
                   host-inbound-traffic {
                      system-services {
                         all;
                      }
                      protocols {
                         all;
                      }
                   }
                }
                ge-0/0/3.1001 {
                   host-inbound-traffic {
                      system-services {
```

```
                     all;
                  }
                  protocols {
                     all;
                  }
               }
            }
         }
      }
## Security zone used for VOIP services only
      security-zone VOIP {
         interfaces {
            ge-0/0/3.1002 {
               host-inbound-traffic {
                  system-services {
                     all;
                  }
                  protocols {
                     all;
                  }
               }
            }
            ge-2/0/3.1002 {
               host-inbound-traffic {
                  system-services {
                     all;
                  }
                  protocols {
                     all;
                  }
               }
            }
         }
      }
   }
   policies {
## policy allows that endpoints in the TRUST zone can communicate with each other
      from-zone trust to-zone trust {
         policy trust-to-trust {
            match {
               source-address any;
               destination-address any;
               application any;
            }
            then {
               permit;
            }
         }
      }
## policy allows that endpoints in the VOIP zone can communicate with each other.
      from-zone VOIP to-zone VOIP {
         policy VOIP-to-VOIP {
            match {
               source-address any;
               destination-address any;
               application any;
            }
            then {
               permit;
            }
         }
      }
## policy allows that endpoints in the DATA zone can communicate with each other.
```

```
        from-zone DATA to-zone DATA {
            policy DATA-to-DATA {
                match {
                    source-address any;
                    destination-address any;
                    application any;
                }
                then {
                    permit;
                }
            }
        }
    }
}
```

Configuring the PE

The DHCP proxy is configured on the PE for branch office HD videoconferencing endpoints.

```
Configuring the Interfaces on the PE Router
## Configure the interface to connect the Branch Office from PE WEST
interfaces {
    ge-1/2/2 {
        vlan-tagging;
        description "Connected to remote site-1 of smallCo";
        hierarchical-scheduler;
        unit 600 {
            vlan-id 600;
            family inet {
                address 10.8.1.49/30;
            }
        }
    }
}
## define the loopback interface for the specific customer VPN
lo0 {
    unit 188 {
        family inet {
            address 10.100.1.27/32;
        }
    }
}
```

Configuring DHCP Proxy on the PE Router for Branch Office Endpoints

```
[edit routing-instances smallCo]
forwarding-options {
    dhcp-relay {
## configure this as a proxy server
        overrides {
## allow the Proxy to receive & process DHCP unicast message as well
            allow-snooped-clients;
## set the relay in proxy mode
            proxy-mode;
        }
        server-group {
## centralized DHCP server IP address in the service providers data center
            HDVC-DC-Serv-1 {
```

```
            10.12.12.100;
         }
      }
      active-server-group HDVC-DC-Serv-1;
      group HDVC {
## interface connecting the branch office
         interface ge-1/2/2.600;
## loopback interface is used as the DHCP proxy server's IP address
         interface lo0.188;
      }
   }
}
```

Configuring the Centralized DHCP Server

Although provisioning of the DHCP server is outside the scope of this chapter, the following code snippet highlights the configuration to match the hardware address of the endpoint that is received in the option-82 field of the DHCP relay message. The DHCP server is provisioned to assign static IP addresses based on these matches so that the video infrastructure in the data center can identify these endpoints during call setup.

```
## SUBNET FOR BRONZE SERVICE ENDPOINTS
# 192.168.51.0-EndPoints
subnet 192.168.51.0 netmask 255.255.255.0 {
# --- default gateway
      option routers              192.168.51.1; ##GATEWAY IP ADDRESS FOR THE
                                           ##ENDPOINTS
      option subnet-mask           255.255.255.0;

      range dynamic-bootp 192.168.51.2 192.168.51.50; ## Poll for End points
      default-lease-time 21600;
      max-lease-time 43200;
}
## SUBNET FOR SILVER SERVICE ENDPOINTS
# Silver-EndPoints
subnet 192.168.70.0 netmask 255.255.255.0 {
# --- default gateway
      option routers              192.168.70.1; ##GATEWAY IP ADDRESS FOR THE
                                           ##ENDPOINTS
      option subnet-mask           255.255.255.0;

      range dynamic-bootp 192.168.70.2 192.168.70.50; ## Pool for End points
      default-lease-time 21600;
      max-lease-time 43200;
}
## SUBNET FOR GOLD SERVICE ENDPOINTS
# Silver-EndPoints
subnet 192.168.50.0 netmask 255.255.255.0 {
# --- default gateway
      option routers              192.168.50.1; ##GATEWAY IP ADDRESS FOR THE
                                           ##ENDPOINTS
      option subnet-mask           255.255.255.0;

      range dynamic-bootp 192.168.50.2 192.168.50.50; ## Pool for End points
      default-lease-time 21600;
      max-lease-time 43200;
}
```

Troubleshooting

After configuring the EX switch and SRX security router, ensure that there is connectivity to reach the essential elements in the network. The following steps illustrate debugging examples at various network elements in the path.

Ping the DHCP server, for example 10.12.12.100 from the SRX. If this is not reachable, perform the following steps.

1. Verify that the PE router's interface is reachable.
2. Validate if the DHCP server is reachable from the last PE router that connects to the data center's edge router.
3. Inspect the OSPF protocol configuration on the data center's edge router.
4. If all validations pass, trace the path backwards by verifying that the L3VPN (MPLS LSPs) are all established and are "up" in both directions. If not all expected LSPs are available, inspect the configuration on the PE routers.
5. If any discrepancies are found and corrected, inspect the route table on the SRX to ensure that the route to the DHCP server has been learned over OSPF.

The following code snippet is the result of performing steps 3-5.

SRX IP reachability troubleshooting

```
## Ensure that OSPF is able to advertise the route to the DHCP server
SRX-Site-1> show ospf route
Topology default Route Table:

Prefix          Path         Route          NH      Metric  NextHop         Nexthop
                Type         Type           Type            Interface       Address/LSP
10.100.1.188    Intra        Area/AS BR     IP           1  ge-0/0/2.600    10.4.1.17
10.4.1.16/30    Intra        Network        IP           1  ge-0/0/2.600
10.8.1.44/30    Ext2         Network        IP           0  ge-0/0/2.600    10.4.1.17
10.11.11.0/24   Inter        Network        IP           3  ge-0/0/2.600    10.4.1.17
10.12.12.0/24   Inter        Network        IP           3  ge-0/0/2.600    10.4.1.17
10.31.31.0/24   Inter        Network        IP           3  ge-0/0/2.600    10.4.1.17
10.40.1.16/29   Ext2         Network        IP           0  ge-0/0/2.600    10.4.1.17

## Check the status of the MPLS LSPs on the PE routers
MX-West-188-re0> show mpls lsp
Ingress LSP: 5 sessions
To             From           State   Rt      P          ActivePath  LSPname
10.100.1.7     10.100.1.27    Up       0      *                      MX-West-to-MX-East
10.100.1.25    10.100.1.27    Up       0      *                      MX-West-to-MX-South
10.100.1.26    10.100.1.27    Up       0      *                      MX-West-to-MX-North
10.101.114.14  10.100.1.27    Up       0      *                      MX-West-to-MX-Campus-114
10.102.1.3     10.100.1.27    Up       0      *                      MX-West-to-MX-Campus-147
Total 5 displayed, Up 5, Down 0
```

```
Egress LSP: 6 sessions
To             From           State    Rt   Style   Labelin   Labelout  LSPname
10.100.1.27    10.100.1.7     Up        0    1 SE         3          -  MX185-to-MX188
10.100.1.27    10.100.1.25    Up        0    1 SE         3          -  MX-South-to-MX-West
10.100.1.27    10.100.1.25    Up        0    1 SE         3          -  Bypass->10.4.1.2
10.100.1.27    10.102.1.3     Up        0    1 FF         3          -  MX-Campus-147-to-MX-West-188
10.100.1.27    10.101.114.14  Up        0    1 FF         3          -  MX-Campus-114-to-MX-West-188
10.100.1.27    10.100.1.26    Up        0    1 SE         3          -  MX-North-to-MX-West
Total 6 displayed, Up 6, Down 0

Transit LSP: 14 sessions
To             From           State    Rt   Style   Labelin   Labelout  LSPname
10.100.1.7     10.100.1.25    Up        0    1 SE    307216     299792  Bypass->10.4.1.6
10.100.1.7     10.102.1.3     Up        0    1 FF    307104     302032  MX-Campus-147-to-MX-East-185
10.100.1.7     10.100.1.26    Up        0    1 SE    307280     302128  Bypass->10.4.1.10
10.100.1.25    10.100.1.7     Up        0    1 SE    307200          3  Bypass->10.4.1.5
10.100.1.25    10.102.1.3     Up        0    1 FF    307120          3  MX-Campus-147-to-MX-South-203
10.100.1.25    10.100.1.26    Up        0    1 SE    307232          3  MX-North-to-MX-South
10.100.1.26    10.100.1.7     Up        0    1 SE    307136          3  MX185-to-MX28
10.100.1.26    10.100.1.25    Up        0    1 SE    307168          3  MX-South-to-MX-North
10.100.1.26    10.102.1.3     Up        0    1 FF    307184          3  MX-Campus-147-to-MX-North-28
10.100.1.26    10.101.114.14  Up        0    1 FF    307152          3  MX-Campus-114-to-MX-North-28
10.101.114.14  10.100.1.26    Up        0    1 FF    307264     302112  MX-North-to-MX-Campus-114
10.102.1.3     10.100.1.7     Up        0    1 FF    307072          3  MX185-to-MX-campus-147
10.102.1.3     10.100.1.25    Up        0    1 FF    307088          3  MX-South-to-MX-Campus-147
10.102.1.3     10.100.1.26    Up        0    1 SE    307248          3  MX-North-to-MX-Campus-147
Total 14 displayed, Up 14, Down 0
```

NOTE: If you experience problems getting a DHCP IP address assigned to the video endpoint, use the following operational command to verify DHCP status on the SRX.

SRX DHCP troubleshooting

```
SRX650-BO-1> show system services dhcp relay-statistics
 Received Packets:       4
Forwarded Packets            4 Dropped Packets              4    Due to missing interface in relay
database: 4    Due to missing matching routing instance: 0    Due to an error during packet read: 0
Due to an error during packet send: 0    Due to invalid server address: 0    Due to missing valid local
address: 0    Due to missing route to server/client: 0
```

Adding Redundant Routers

Figure 8.5 shows a simplified campus configuration. The campus adds more LAN access switches, but these are configured identically to those used in the branch office. In addition, an additional tier of LAN switches is added to aggregate traffic together. Depending upon the size of the site, all EX4200 switches shown can be part of a single virtual chassis.

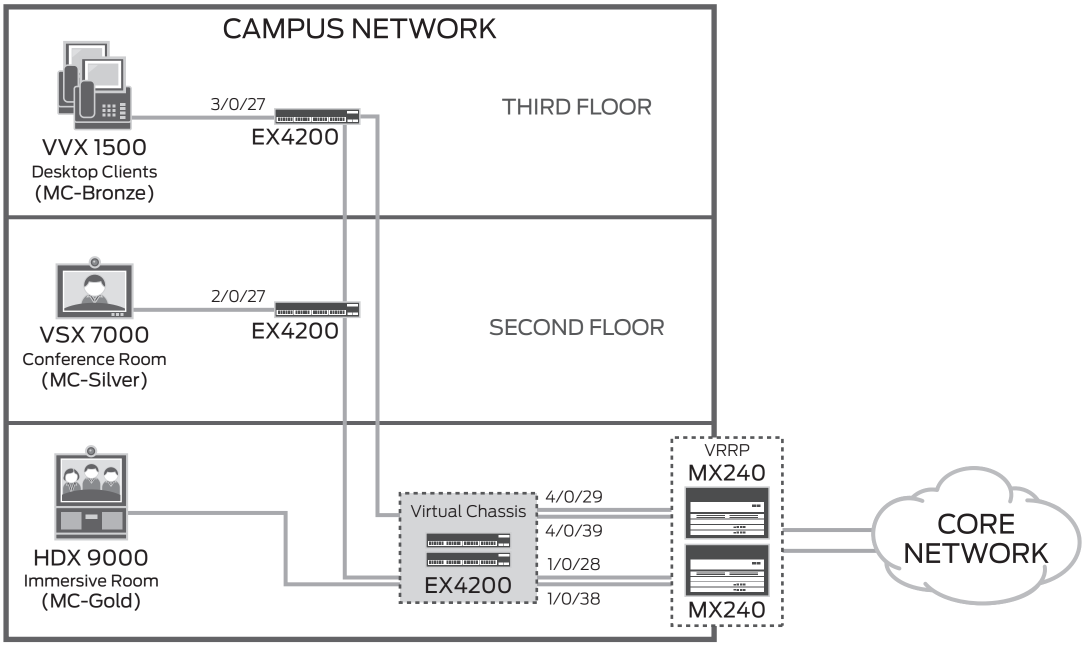

Figure 8.5 Campus implementation (simplified)

Configuring VRRP

Virtual Router Redundancy Protocol (VRRP) is configured on the LAN interfaces of the MX Series routers to provide a resilient default gateway. At any time, one of the MX Series routers is the master (active) and the other is a backup. If the master router or the tracked interface fails, the backup router becomes the new VRRP master, hence restoring a default route and providing endpoints with connectivity to the network. To minimize disruptions due to failure, the network will be configured so that gold tier and bronze tier traffic will use the top router and WAN link during normal operation, while silver tier traffic will use the bottom router and WAN link. If a node or link fails, then the remaining path will be used for all traffic. Figure 8.6 shows an overview of the VRRP setup used to provide WAN router resiliency.

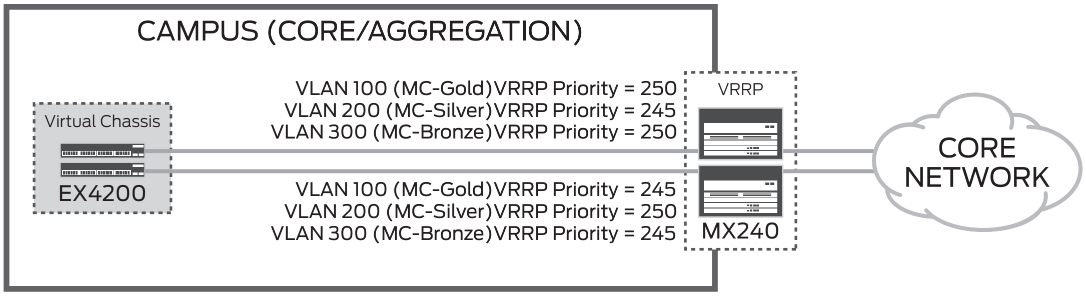
CAMPUS (CORE/AGGREGATION)
Virtual Chassis
EX4200
VLAN 100 (MC-Gold)VRRP Priority = 250
VLAN 200 (MC-Silver)VRRP Priority = 245
VLAN 300 (MC-Bronze)VRRP Priority = 250
VLAN 100 (MC-Gold)VRRP Priority = 245
VLAN 200 (MC-Silver)VRRP Priority = 250
VLAN 300 (MC-Bronze)VRRP Priority = 245
VRRP
MX240
CORE
NETWORK

Figure 8.6 Default operation using VRRP

The following MX Series code example shows how to configure redundancy using VRRP. This router is the master VRRP router for VLANs 100 (MC-Gold) and 300 (MC-Bronze), and is the backup router for VLAN 200 (MC-Silver). A separate VRRP instance is defined for each VLAN. Key parameters include:

- **Virtual IP address:** All packets are sent to this virtual address, and the VRRP-enabled routers determine which one will forward the packet. A separate virtual IP (VIP) address is used for each VLAN.
- **Priority:** To minimize the impact of a single failure, one of the routers will be the VRRP master for gold and bronze traffic, while the other router will be the master for silver traffic. By default, all traffic would traverse the same router. Configuring priorities overrides the default behavior. This configure specifies a priority of 250 for the master router and 245 for the backup router.
- **Fast-interval:** In VRRP operation, the master router sends advertisements to backup routers at (by default) one second intervals. If a backup router does not receive an advertisement for a set period, the backup router with the next highest priority takes over as master and begins forwarding packets. Using the default parameters, it takes up to 3 seconds for the restoration of the default route in the case of failure. This is too long for HD videoconferencing service. To enable faster failure detection and switchover, configure fast-interval timer to 100 ms for VRRP advertisement packets. This provides service restoration of about 300 ms, which is going to limit impact to HD videoconferencing services and prevent calls from dropping.
- **Preempt:** By default, a higher-priority backup router preempts a lower-priority master router. To explicitly enable the master router to preempt, include the preempt statement.
- **Accept-data:** Tells the interface to accept packets destined for the virtual IP address.
- **Track:** Tells VRRP which interface to track. This configuration includes two sub-parameters:
  - **Interface:** If this interface fails, then a new master router will be elected.
  - **Priority-cost:** This cost is deducted from priority when a failure occurs.

```
ge-2/0/1 {
    description “Connected to EX Virtual Chassis”;
    hierarchical-scheduler;
    mtu 4096;
    vlan-tagging;
        unit 100 {
        vlan-id 100;
        family inet {
            address 192.168.41.3/24 {
                vrrp-group 10 {
                    virtual-address 192.168.41.1;
                    ## This router will be MASTER for VLAN 100
                    priority 250;
                    ## To enable faster failure detection
                    fast-interval 100;
                    preempt;
                    accept-data;
                    track {
                        interface ge-2/2/0 {
                            priority-cost 10;
                        }
                    }
                }
            }
        }
        unit 200 {
        vlan-id 200;
        family inet {
            address 192.168.42.2/24 {
                vrrp-group 20 {
                    virtual-address 192.168.42.1;
                    ## This router will be BACKUP for VLAN 200
                    priority 245;
                    fast-interval 100;
                    preempt;
                    accept-data;
                    track {
                        interface ge-2/2/0 {
                            priority-cost 10;
                }
            }
        }
    }
        unit 300 {
        vlan-id 300;
        family inet {
            address 192.168.41.3/24 {
                vrrp-group 30 {
                    virtual-address 192.168.43.1;
                    ## This router will be MASTER for VLAN 300
                    priority 250;
                    fast-interval 100;
                    preempt;
                    accept-data;
                    track {
                        interface ge-2/2/0 {
                            priority-cost 10;
                        }
                    }
                }
            }
        }
    }
}
```

Verifying VRRP Operation

The **run show vrrp** command can be used to verify VRRP status. The LCL address, as shown in the following code snippet, represents the local interface supporting this VRRP instance. VIP is the virtual IP address that is used for this instance, and the MAS address represents the current master interface.

```
root@MX-A# run show vrrp
Interface    State      Group  VR state VR Mode  Timer   Type   Address
ge-2/0/1.100  up          10   master   Active     A  0.011 lcl    192.168.41.3
                                                   vip    192.168.41.1
ge-2/0/1.200  up          20   backup   Active     A  0.295 lcl    192.168.42.2
                                                   vip    192.168.42.1
                                                   mas    192.168.42.3
ge-2/0/1.300  up          30   master   Active     A  0.011 lcl    192.168.43.3
                                                   vip    192.168.43.1
```

Chapter 9

Implementing Quality of Service

Part Two

There is a significant difference between delivering a better HD videoconferencing service compared to delivering the best HD videoconferencing service. One key difference is in the ability to guarantee consistent Quality of Service (QoS). HD videoconferencing is a resource intensive application with low tolerance to network impairments. The most direct way of addressing the quality issues is to provide this application with all required resources when needed and exclude the possibility of resource depletion due to intensive use from other applications.

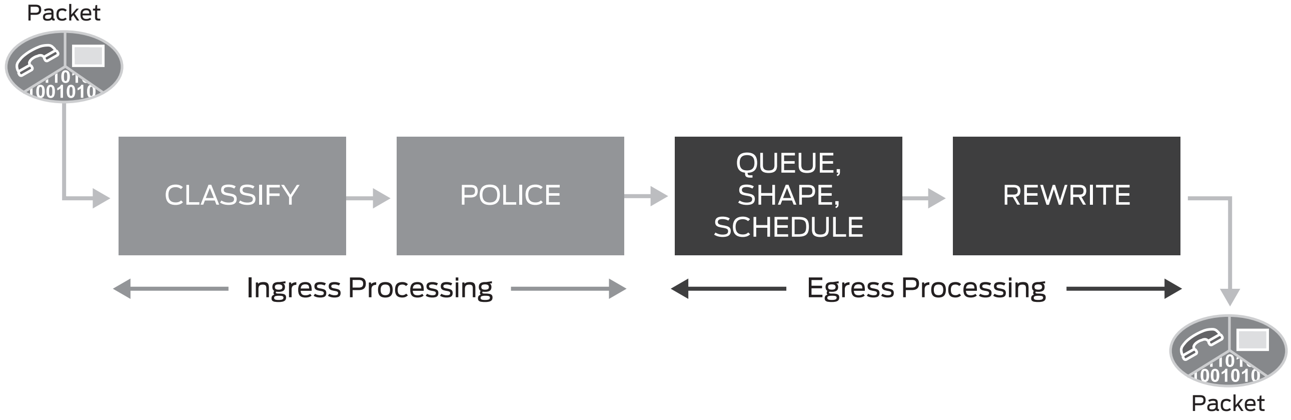

Figure 9.1 Implementing QoS

Implementing QoS is an important technique for ensuring resource availability across a converged network. Figure 9.1 illustrates the capabilities of a robust QoS implementation, which includes:

- **Ingress Classification:** In this step, the node looks at the incoming packet to determine how it should be treated. The key role of QoS classifiers is to identify the traffic class to which each packet belongs, which in turn determines how it will be treated within the node.
- **Ingress Policing:** Traffic policers enforce the agreed upon incoming data rates. The traffic that exceeds the set rate will be dropped or reclassified to a lower priority class.
- **Egress Queuing:** Each outbound packet is placed into the appropriate outbound queue based upon its traffic class. This process determines how often this type of traffic is sent onto the wire in comparison with other traffic types and how much traffic to buffer while the specified traffic waits its turn to be sent from the interface.
- **Egress Shaping:** Traffic shapers shape the outgoing data streams to make sure that these are compliant with expected data rates on the rest of the network. The traffic that exceeds the configured rate will be first queued and, if continuous congestion occurs, it will be dropped. This method provides an effective tool for protection of the HD videoconferencing traffic from the other applications that can consume too much bandwidth and starve videoconferencing flows from bandwidth.
- **Egress Scheduling:** Differentiated congestion management techniques such as random early detection (RED) help to proactively avoid heavy congestion and free the network bandwidth for HD videoconferencing application traffic.

- **Egress Rewrite:** This step makes outgoing traffic with the appropriate markings, so that the next hop device knows how to treat this packet.

The rules, which govern how each packet is treated at each node, are referred to as Per Hop Behavior (PHB), with queuing and scheduling having the greatest impact. The administrator must define an end-to-end QoS strategy, and all network nodes should implement a consistent set of QoS rules, which implement this. Each function does not necessarily need to be implemented at each device, and QoS is most critical when there are bandwidth bottlenecks that prevent all packets from being delivered to the intended destination in an acceptable time frame. The QoS configuration is generic in Junos, allowing the administrator to use a single QoS syntax across all routers and switches.

Juniper's recommended implementation is consistent with RFC 4594, which describes various service classes and offers recommendations on expected treatment. For completeness, definitions are also provided for high-priority VoIP traffic and best-efforts (BE) data traffic that share the same network.

QoS can be implemented in one of two ways. Most QoS configurations are manually provisioned, and this setup never changes. This method is commonly done at each network site's CE router. In addition, this may be done in core (P) routers if there is any possible congestion. QoS can also be manually configured at each LAN switch if desired, since their limited number of buffers may cause packets to be dropped.

The other alternative is to push QoS definitions dynamically based on network activity. For example, if one customer signs up for a videoconferencing service, it is better to automate the QoS configuration process at the PE router than manually reconfigure that node. This minimizes the possibility of misconfiguring the device. This chapter covers static QoS configuration. Dynamic QoS is covered in *Chapter 10, Implementing Assured Forwarding on page 129.* For additional information on QoS, see *Junos Class of Service Configuration, Release 10.3* at **juniper.net/techpubs/en_US/junos10.3/information-products/pathway-pages/cos/index.html**.

Planning

Since HD videoconferencing is a real-time application, it should be defined as a separate traffic class that receives high priority treatment as it traverses the network. In addition, there may be a desire to define multiple traffic classes for different tiers of HD videoconferencing usage. While the network administrator determines the number of traffic classes, this chapter illustrates an implementation that supports three traffic classes. These traffic classes are referred to collectively as multimedia-conferencing[1], with the individual classes being distinguished as gold, silver and bronze (high, medium and low priority).

There are several ways that videoconferencing traffic can be assigned to the forwarding classes. Traffic can be assigned based on the device type, for example, by assigning immersive systems to the highest (gold) class and personal systems to the lowest class. Alternatively, traffic can be prioritized based on business need, such as reserving the highest class for customer demonstration systems.

[1] This term is used in RFC 4594, *Configuration Guidelines for DiffServ Service Cla*sses

To simplify the design, we are using the first approach and assigning classes based on device type. Existing desktop videoconferencing systems are carried as their own traffic class (AF43). In many cases, soft clients are used and the video is carried as Best Efforts (BE) traffic.

Table 9.1 lists the forwarding classes used and how they are treated at each hop node.

Table 9.1 Forwarding Classes and Treatment

Forwarding Class	DiffServ[2,3]	IEEE802.1/ MPLS	Hardware Queue	Scheduler	Loss Priority
Network Control	NC	6	3	– –	Low
VoIP	EF	2	2	Strict	Low
(MC-Gold)	AF41	3			Medium-Low
MultimediaConferencing-Silver (MC-Silver)	AF42	4	1	WRR	Low
Multimediaconferencing-Bronze (MC-Bronze)	AF43	5	0	WRR	Low
Data (Best Efforts)	BE	0			High

Ingress Processing

As traffic enters the router, two functions are performed: classification and policing. Classification assigns traffic to forwarding classes, which effectively determines key characteristics of how the traffic is treated as it traverses the router. Policing limits the amount of traffic in each forwarding class.

Classification

There are three techniques for determining how to treat traffic:

- **Multi-field (MF) classifiers:** In this case, the network element looks at well-known criteria such as the source/destination IP address and application ports to determine how to treat this traffic. For example, all videoconferencing room systems may be on their own VLAN, or have IP addresses assigned from a separate address pool.
- **Behavior Aggregate (BA) classifiers:** The packet may carry a value in the DiffServ, IEEE802.1p or MPLS EXP field which can be used to determine how the packet should be treated.

[2] Assured Forwarding (AF) markings are defined in RFC 2597, Assured Forwarding PHB Group

[3] Expedited forwarding is defined in RFC 3246, An Expedited Forwarding PHB (Per-Hop Behavior)

- **Deep Packet Inspection (DPI):** In some implementations, HD videoconferencing traffic cannot be uniquely identified by examining only packet header information (IP address and or application port numbers). However, DPI specifically must be used to identify HD videoconferencing traffic. As the name implies, DPI engines look further into the packet, looking into the application data that uniquely identifies the purpose of this packet. This method has scaling limits because each packet must be individually inspected. DPI also requires separate analysis rules for each protocol. When an application upgrade occurs, the DPI equipment can stop detecting the video traffic, and such a system can suffer failures despite high implementation costs.

Juniper Networks supports all three methods. As a rule, looking deeper into the packet to determine its treatment requires additional processing. BA and MF classifiers are the most efficient, and Juniper routers support line-rate forwarding on all ports with these methods. DPI requires looking deeper into each packet, often resulting in reduced throughput.

Regardless of which type of classifier is used, the classifier:

- Assigns traffic to a forwarding-class. Forwarding classes are defined by the network administrator.
- Assigns either a low or high loss priority to each packet. High packet loss priority (PLP) means that the packet will typically be dropped upon egress. The drop mechanism is defined in the *Queuing* section in this chapter.

Multi-field Classification

Videoconferencing clients often provide the ability to assign a priority marking (typically IP DiffServ) value which will be marked on all outbound packets. However, practical QoS implementation dictates that the first hop (within a defined QoS domain) will ignore this marking. Instead, the node will look at other fields in the packet to determine how this traffic should be treated as it traverses the network. For example, the network administrator may not want every desktop video client to receive preferential treatment across the network. Therefore, MF classifiers are typically used at this first hop.

Multi-field (MF) classifiers are usually based on one of the following:

- **IP address of videoconferencing equipment:** This is common when assigning static IP addresses to equipment, as is done for high-end immersive systems.
- **IP subnets:** Videoconferencing equipment can be assigned IP addresses from a different address pool, and any traffic in the identified range can be prioritized.
- **VLANs:** Since different tiers (gold, silver, bronze) of videoconferencing equipment can be mapped to different VLANs at the wiring closet LAN switch, traffic can be prioritized based on VLAN identifiers.

The following code snippet MF classifiers, which assign traffic into forwarding classes, is based on the IP address of the videoconferencing system.

```
family inet {
## Ingress filter applied for the video end point generated traffic
    filter {
        input MC-Gold-FILTER;
    }
    address 192.168.41.1/24;
}
## Ingress filter applied for the video end point generated traffic    filter {
        input MC-Silver-FILTER;
    }
    address 192.168.70.1/24;
}
family inet {
## Ingress filter applied for the Video End Point generated traffic
    filter {
        input MC-Bronze-FILTER;
    }
    address 192.168.51.1/24;
}
## Defines that all the traffic matches will be marked with DSCP bits and forwarded to Silver Queue.
firewall {
    family inet {
        filter MC-Gold-FILTER {
            term 1 {
                then {
                    count ALL;
                    log;
                    loss-priority medium-low;
                    forwarding-class MC-Gold;
                    accept;
                    }
                }
            }
        filter MC-Silver-FILTER {
            term 1 {
                then {
                    count ALL;
                    log;
                    loss-priority low;
                    forwarding-class MC-Silver;
                    accept;
                    }
                }
            }
## Defines that all the traffic matches will be marked with DSCP bits and forwarded to Bronze Queue.
        filter MC-Bronze-FILTER {
            term 1 {
                then {
                    count ALL;
                    log;
                    loss-priority low;
                    forwarding-class MC-Bronze;
                    accept;
                    }
                }
            }
       }
}
```

Behavior Aggregate Classification

After determining the appropriate treatment using MF classifiers, the first trusted hop in the QoS domain marks the packet with the appropriate Behavior Aggregate QoS markings. Therefore, subsequent nodes can look at this field to determine how to forward traffic. This process for adding the BA markings is described in the *Rewriting QoS Markings* section of this chapter.

Differentiated Services Code Point (DSCP) (IP), 802.1p (Ethernet) and EXP (MPLS) BA classifiers are shown below, although only one is typically configured and used at each router interface. DSCP markings are typically used on IP connections from sites to the core, while MPLS markings are used within the core network. IEEE802.1p classifiers are used within sites, such as LAN segments built with EX switches.

Assuming that the incoming packet is an IP packet marked as AF41, the packet is classified as follows:

- Looking at the first classifier, traffic marked as DiffServ AF41 traffic is assigned to the forwarding-class labeled as MC-Gold.
- The classifier specifies that this packet has a low loss priority class.
- AF41 is an alias to a bit pattern in the DiffServ field of 0x100010. This is the standard bit pattern definition of AF41. Note that since we are using the generally accepted definition of AF41, the configuration step is not required but is shown for completeness in the following code-point-aliases code snippet.
- Finally, the MC-Gold traffic class is assigned to hardware queue 1, which is specified as high priority. For the MX, each queue can be identified as high or low priority.

The following code shows the classifiers used to support videoconferencing as well as VoIP and data in our scenario. At the CE, the most common implementation is to prioritize traffic based on DiffServ (DSCP) markings. At the P router, MPLS experimental (EXP) bits are used to identify traffic. At the LAN switches, IEEE 802.1p markings will be used.

```
class-of-service {
   classifiers {
      dscp DSCP-CLASSIFIER {
         forwarding-class NC {
            loss-priority low code-points nc;
         }
         forwarding-class MC-Gold {
            loss-priority low code-points ef;
            loss-priority medium-low code-points af41;
         }
         forwarding-class MC-Silver {
            loss-priority low code-points af42;
         }
         forwarding-class MC-Bronze {
            loss-priority high code-points be;
            loss-priority low code-points af43;
         }
      }
      exp MPLS-CLASSIFIER {
         forwarding-class NC {
            loss-priority low code-points nc;
         }
         forwarding-class MC-Gold {
            loss-priority low code-points ef;
            loss-priority medium-low code-points ef1;
         }
         forwarding-class MC-Silver {
            loss-priority low code-points af11;
         }
         forwarding-class MC-Bronze {
            loss-priority low code-points af12;
            loss-priority high code-points be;
         }
      }
      ieee-802.1 ENET-CLASSIFIER {
         forwarding-class NC {
            loss-priority low code-points nc;
         }
         forwarding-class MC-Gold {
            loss-priority low code-points ef;
            loss-priority medium-low code-points ef1;
         }
         forwarding-class MC-Silver {
            loss-priority low code-points af11;
         }
         forwarding-class MC-Bronze {
            loss-priority low code-points af12;
            loss-priority high code-points be;
         }
      }
   }
   code-point-aliases {
      dscp {
         be 000000;
         af41 100010;
         af42 100110;
         af43 100110;
         ef 101110;
         nc 110000;
      }
      exp {
         be 000;
         af11 100;
         af12 101;
         ef 010;
         ef1 011;
         nc 110;
      }
      ieee-802.1 {
         be 000;
         af11 100;
         af12 101;
         ef 010;
         ef1 011;
         nc 110;
      }
   }
   forwarding-classes {
      queue 0 MC-Bronze;
      queue 2 MC-Silver;
      queue 1 MC-Gold;
      queue 3 NC;
   }
```

Policing

To protect the upstream network from excessive video traffic (which might be caused by DoS attack or application misconfiguration), ingress policers on the edge of the QoS domain can be used to enforce traffic compliance to the bandwidth levels defined during the planning phase. This safeguard mechanism is only a precautionary measure. During normal operations, the policers should permit all HD videoconferencing traffic that is compliant with the specified service tier. In the case of any network anomaly, the ingress policers will drop excess HD videoconferencing traffic and protect the rest of the network from the misbehaving application.

The policing bandwidth should be set to the maximum value of expected bandwidth for MC-Gold and MC-Silver traffic tiers, as follows:

- PG=NGxBG
- PS=NSxBS

Where

- NG is the number of gold endpoints
- BG is the maximum bandwidth per gold endpoint
- NS is the number of silver video-endpoints
- BS is the maximum bandwidth per silver video-endpoint.

In this scenario, we are assigning a single HDX 8000 to the MC-Gold class. Although MPEG-4 sends a variable amount of traffic based on movement of the participants, it typically will not exceed 2.4 Mbps. Similarly a VSX is assigned to MC-Silver, and this device peaks at 1 Mbps. VVXs are assigned to MC-Bronze service, and this value should not exceed 512 Kbps.

```
policer 2.4MB {
   if-exceeding {
      bandwidth-limit 2400000;
      burst-size-limit 2800000;
   }
   then discard;
}
policer 512k {
   if-exceeding {
      bandwidth-limit 512k;
      burst-size-limit 768k;
   }
   then discard;
}
policer 1MB {
   if-exceeding {
      bandwidth-limit 1024000;
      burst-size-limit 1200k;
   }
   then discard;
}
```

Egress Processing

Egress processing consists of three parts: queuing, scheduling/shaping, and rewriting. An important piece of queuing is determining which traffic to drop when congestion occurs. Scheduling and shaping control how packets are placed onto the wire. Rewriting adds the appropriate BA header onto the packet.

Queuing

During congestion, the HD videoconferencing traffic should be queued, but only to the specific moment. All traffic that exceeds the specified buffer length should be dropped from the head of the queue by applying RED[4] drop profiles. RED drop profiles consist of two components. The fullness level represents the percentage of buffers currently in use storing packets. The associated drop probability is the likelihood that a packet will be dropped from the network.

For the higher priority (MC-Gold and MC-Silver) traffic, the RED drop profiles should be simple since other traffic will be dropped before this traffic is affected. If congestion is ever so severe that the Gold and Silver classes are affected, it is better to drop all traffic sequentially than to randomly drop packets. Therefore, at 0 percent of the queue depth, it should specify 0 percent loss probability, while at 100 percent of the queue depth, it should specify 100 percent drop probability.

Lower priority (MC-Bronze and BE-Data) traffic will be dropped much sooner. The low priority videoconferencing traffic (MC-Bronze) is not dropped until the queues are 80 percent full. In contrast, since the Best Efforts data traffic is sent through TCP and can be retransmitted, these particular packets are dropped when queue utilization reaches 5 percent.

```
drop-profiles {
    NC-DROP {
        interpolate {
            fill-level 100;
            drop-probability 0;
        }
    EF-DROP {
        interpolate {
            fill-level 100;
            drop-probability 0;
        }
    }
    MC-Gold-DROP {
        interpolate {
            fill-level 100;
            drop-probability 0;
        }
    }
    MC-Silver-DROP {
        interpolate {
            fill-level 100;
            drop-probability 0;
```

[4] For an introduction to RED, see **juniper.net/techpubs/en_US/junos10.3/information-products/pathway-pages/cos/red-drop-profiles.html**

```
          }
       }
       MC-Bronze-DROP {
          interpolate {
             fill-level [80 90 100 ];
             drop-probability [20 50 100 ];
          }
       BE-DATA-DROP {
          interpolate {
             fill-level [ 5 25 50 75 80 100 ];
             drop-probability [10 40 60 80 90 100 ];
          }
      }
 }
```

Rewriting QoS Markings

Before being put onto the network, the appropriate DiffServ (or MPLS EXP or 802.1p) markings should be added back onto the packet. For example, packets in the MC-Gold forwarding class with a medium-low loss priority are marked with the DSCP AF41 code point, thereby recreating the original marking.

```
rewrite-rules {
   dscp DSCP-RW {
      forwarding-class NC {
          loss-priority low code-point nc;
       }
      forwarding-class MC-Gold {
          loss-priority low code-point ef;
          loss-priority medium-low code-point af41;
       }
       forwarding-class MC-Silver {
          loss-priority low code-point af42;
       }
       forwarding-class MC-Bronze {
          loss-priority low code-point af43;
          loss-priority high code-point be;
       }
   }
   exp MPLS-CLASSIFIER {
       forwarding-class NC {
          loss-priority low code-point nc;
       }
       forwarding-class MC-Gold {
          loss-priority low code-points ef;
          loss-priority medium-low ef1;
       }
       forwarding-class MC-Silver {
          loss-priority low code-points af11;
       }
       forwarding-class MC-Bronze {
          loss-priority low code-points af12;
          loss-priority high code-point be;
       }
   }
   ieee-802.1 ENET-CLASSIFIER {
       forwarding-class NC {
          loss-priority low code-point nc;
       }
       forwarding-class MC-Gold {
         loss-priority low code-points ef;
         loss-priority medium-low ef1;
      }
      forwarding-class MC-Silver {
         loss-priority low code-points af11;
      }
      forwarding-class MC-Bronze {
         loss-priority low code-points af12;
         loss-priority high code-point be;
      }
   }
```

Shaping and Scheduling

The next step is to define one or more shapers, which determine how packets from the various forwarding classes are placed onto the wire. The shaper consists of two pieces: a scheduler map and the shaping rate. The scheduler map is a collection of individual schedulers, each of which specifies the rules being followed for each forwarding class. For These schedulers support both TCP and UDP protocols in this class (**`protocol any`**) regardless of what loss priority was assigned (**`loss-priority any`**).

For this configuration, there is a single scheduler map (SCHED-MAP) which defines this, and this scheduler map will be assigned to all ports in the following section. VoIP and MC-Gold traffic are assigned to HI-SCHED, which supports this high priority traffic. Similarly, data and MC-Bronze traffic are mapped to a scheduler, which supports this low priority traffic. Egress traffic is shaped (limited) to 6 Mbps.

```
   scheduler-maps {
      SCHED-MAP {
         forwarding-class MC-Bronze scheduler MC-Bronze-SCHED;
         forwarding-class MC-Silver scheduler MC-Silver-SCHED;
         forwarding-class MC-Gold scheduler MC-Gold-SCHED;
         forwarding-class NC scheduler NC;
      }
   }
   schedulers {
      NC {
         drop-profile-map loss-priority low protocol any drop-profile NC-DROP;
      }
      MC-Gold-SCHED {
         priority high;
         drop-profile-map loss-priority medium-low protocol any drop-profile MC-Gold-DROP;
         drop-profile-map loss-priority low protocol any drop-profile EF-Drop;
      }
      MC-Silver-SCHED {
         drop-profile-map loss-priority low protocol any drop-profile MC-Silver-DROP;
      }
      MC_BRONZE-SCHED {
         transmit-rate remainder;
         drop-profile-map loss-priority low protocol any drop-profile MC-Bronze-DROP;
         drop-profile-map loss-priority high protocol any drop-profile BE-DATA-DROP;
      }
         }
```

Applying QoS to Interfaces

One final task is to enable QoS on the required interfaces and VPNs. The following code sample is from the Junos router located at the enterprise's WAN edge, so DSCP classifiers and rewrite rules are used. The gigabit Ethernet link (ge-1/2/0) connects to a downstream EX4200 switch, while the 10GE link (xe-4/0/0) connects to upstream PE routers, and QoS is enabled for both upstream and downstream traffic. For egress scheduling, the same scheduler map is used on access-facing and core-facing interfaces.

```
[edit class-of-service]

traffic-control-profiles {
   CORE-MAP {
      scheduler-map SCHED-MAP;
   }
   ACCESS-MAP {
      scheduler-map SCHED-MAP;
      guaranteed-rate 50m;
   }
}
interfaces {
   ge-1/2/0 {
      unit 600 {
         output-traffic-control-profile ACCESS-MAP;
         classifiers {
            dscp DSCP-CLASSIFIER;
         }
         rewrite-rules {
            dscp DSCP-RW;
         }
      }
   }
   xe-4/1/0 {
      output-traffic-control-profile CORE-MAP;
      unit 0 {
         classifiers {
            dscp DSCP-CLASSIFIER;
         }
         rewrite-rules {
            dscp DSCP-RW;
         }
      }
   }
   xe-4/2/0 {
      output-traffic-control-profile CORE-MAP;
      unit * {
         classifiers {
            dscp DSCP-CLASSIFIER;
         }
         rewrite-rules {
            dscp DSCP-RW;
         }
      }
   }
}
```

Chapter 10

Implementing Assured Forwarding

Part Two

Juniper Networks and Polycom created a joint solution, which provides high definition videoconferencing calls with, agreed and assured service quality levels on an end-to-end converged network. This solution uses a combination of Juniper's networking and security platforms, policy management systems together with Polycom's intelligent videoconferencing infrastructure and high definition video endpoints. The solution allows service providers to offer different tiers of HD videoconferencing service with clearly defined SLA quality without the need to build a dedicated overlay network. Based on the service tiers (as listed in Table 10.1), when HD videoconferencing users make a call, the network assures sufficient bandwidth throughout the duration of the call. For specific service tiers, if bandwidth cannot be guaranteed at the time of call setup, the user is notified that the network cannot provide the assured bandwidth for this call and is prompted to try again later.

In addition to decreasing the cost of the network, the HD videoconferencing assured forwarding solution is more scalable and simplifies adding new locations to the network, thus simplifying operations as well.

This chapter provides technical details on how to implement the assured forwarding for delivery of HD videoconferencing services. Although the end-to-end solution includes the Polycom video communication infrastructure, this chapter primarily focuses on configuration of Juniper products and offers troubleshooting tips for HD videoconferencing call setups.

Assured Forwarding Overview

The key goal of the HD videoconferencing solution is to maximize the utilization of the network and the videoconferencing infrastructure at all times without degrading quality of service or excessive over provisioning of the network resources. This can be achieved by providing flexible and tiered subscription models to customers based on their service requirements and budgets. The HD videoconferencing solution recommends the three service tiers HD videoconferencing offering based on the common business rules and practices of most enterprises, specifications and capabilities of Polycom's high definition endpoints, the network resource availability and video application requirements and the cost effectiveness of the overall solution. Table 10.1 lists and defines the three types of Multimedia Conference tiers.

Table 10.1 Three Types of Multimedia Conference Tiers

Service Tier	End Users	Characteristics	Typical Polycom Endpoints
MC-Gold	Top executives and immersive room systems	· Allocation of network resources is pre-planned. No over-subscription allowed in this tier. Number of installed endpoints is strictly regulated and should not exceed the resource allocation. · Calls will never be denied. · Network will always have enough unreserved bandwidth to support this traffic	RPX, OTX, TPX (immersive systems) HDX 9000, HDX 8000
MC-Silver	Remote office employees, regular employees, everyday meeting rooms in the campus	· The number of users in this tier is flexible and combined need for bandwidth resources may be more than network available bandwidth. · Calls are admitted only if there is sufficient bandwidth to ensure service quality · Calls for which is not possible to assure bandwidth are not admitted	HDX 4000, VSX 7000
MC-Bronze	Teleworkers, desktop endpoints, SOHO locations	· Number of users in this service tier is not regulated and call admission control is not performed. · There is no reservation of the dedicated resources and it shares the bandwidth with other applications and services on the network. · This service tier may be over-subscribed experience congestion. · Quality may be impacted if too much other traffic is on the network	VVX 1500

To implement these HD videoconferencing service tiers, the key functional components required include:

- **Service tier provisioning:** When a customer subscribes or upgrades the subscription to the assured HD videoconferencing solution, the service provider provisions the customer's endpoints profiles and service tiers in the videoconferencing infrastructure and defines the network characteristics, such as topology, interfaces and bandwidth, in the policy management system. This one-time provisioning enables dynamic call setup and teardown, while preventing traffic congestion and service over-subscription.
- **Network aware call admission control:** When a video endpoint attempts to set up a call, the HD videoconferencing call-signaling infrastructure identifies the endpoint by its IP address and maps it to the subscribed service tier in the subscriber database. Based on the service tier definition for this endpoint, the system attempts to determine if the call can be permitted under the current network utilization conditions. The CAC and network Policy Management System components of the HD videoconferencing solution handle this decision.
- **Dynamic Quality of service (QoS):** If there are sufficient resources available to permit the call, the Policy Management System then communicates with the network infrastructure to reserve bandwidth for the call and signals back to the videoconferencing infrastructure to proceed with the call.

At the end of the call, the above functional components will free up the resources, thus making them available for other callers as well as for other network applications. The later sections of this chapter present further details on where these functional components are implemented.

Required Network and Video Elements

Figure 10.1 depicts the required network and video elements for enabling assured HD videoconferencing calling experience.

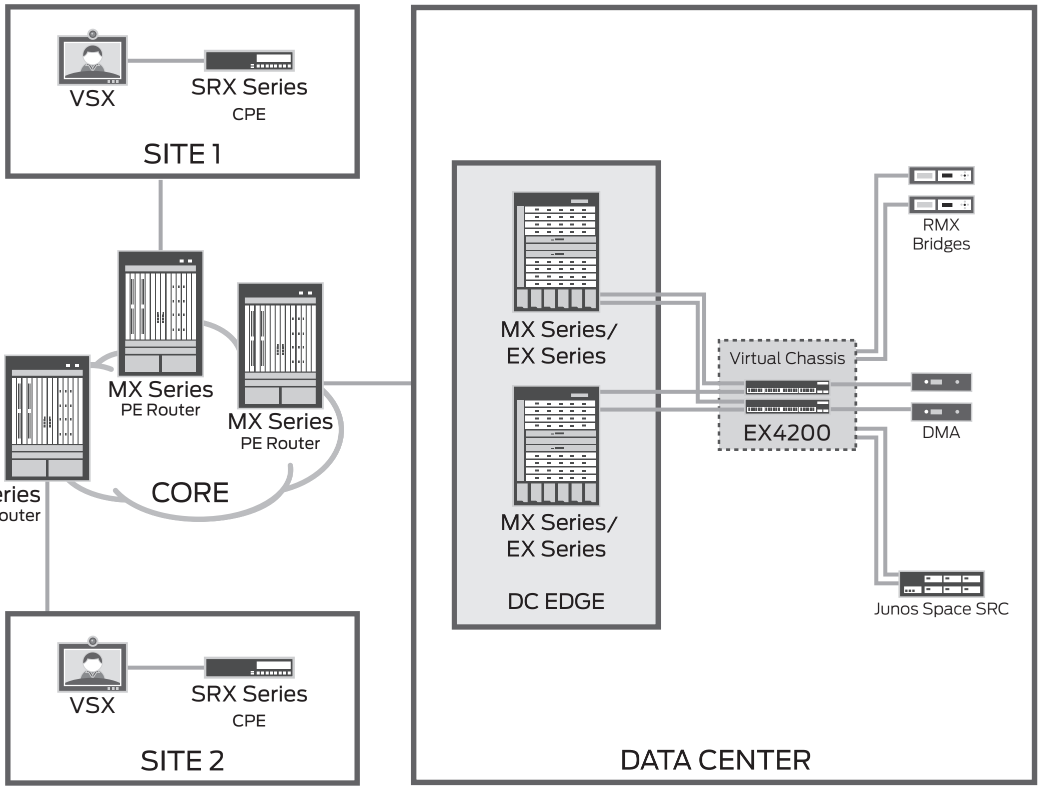

Figure 10.1 Assured HD videoconferencing– network and video infrastructure (simplified)

This section maps the functional components to the network and video elements of the HD videoconferencing solution.

- **Polycom DMA Server:** The Distributed Media Application server is Polycom's scalable call control and virtual resource pool management engine. It serves as the H.323 gatekeeper or SIP server, which the endpoints register and communicate with. The DMA is pre-provisioned with information about end-points, their capabilities and the site geographical locations. This is also the element where conference bridge calls can be setup on-demand by end customers. Currently, when offering videoconferencing on a shared infrastructure a dedicated DMA is required for each enterprise.

  The DMA communicates with Session and Resource Controller (SRC) through the SOAP interface to relay the call setup and tear down information.

- **Juniper Session and Resource Control (SRC) software:** The SRC participates in the CAC process and is required to deliver differentiated tiers of services. The SRC is aware of the network state and bandwidth resource availability and provisions an end-to-end bandwidth for the call while ensuring that the requested SLA is met. The SRC interacts with Polycom's Distributed Media Application (DMA) Server through the SOAP interface.

  The SRC communicates with the Juniper routers through the Diameter protocol to provision QoS on a per-call basis.

- **Polycom RMX:** The Real-time Media Conference Platform provides the multipoint conferencing facility to the endpoints by mixing the video and audio streams from multiple calls. When a conference call setup request is received, the DMA selects an RMX device based on the current load and communicates the call setup information to the appropriate RMX media server.

- **Juniper MX Routers:** The PE router provides the endpoints and the data center with connectivity to each other and serves as the demarcation point between SP and customer networks. This is where the dynamic QoS functionality is provisioned. These routers can be pre-provisioned with the endpoint information or can be configured to discover the endpoint(s) when they request a DHCP based IP address. When an endpoint is discovered, a temporary 'demux' logical interface is created in the router specifically for this endpoint and will be used to apply dynamic QoS policies by the SRC. The Junos routers receive the policy activation request from the SRC in the data center during call setup and in turn, apply this policy on the endpoint's 'demux' interface.

Call Setup Flow

To understand how the various network and video elements interact during a HD videoconferencing call, refer to Figure 10.2. It illustrates in details a successful point-to-multipoint conference call setup with two endpoints dialing into a conference bridge. The scenario assumes that the endpoints are MC-Silver service tier subscribers, which uses the CAC functionality. Care should be taken to allocate sufficient bandwidth on the interface to which the RMX is connected as it will be handling multiple simultaneous conference calls.

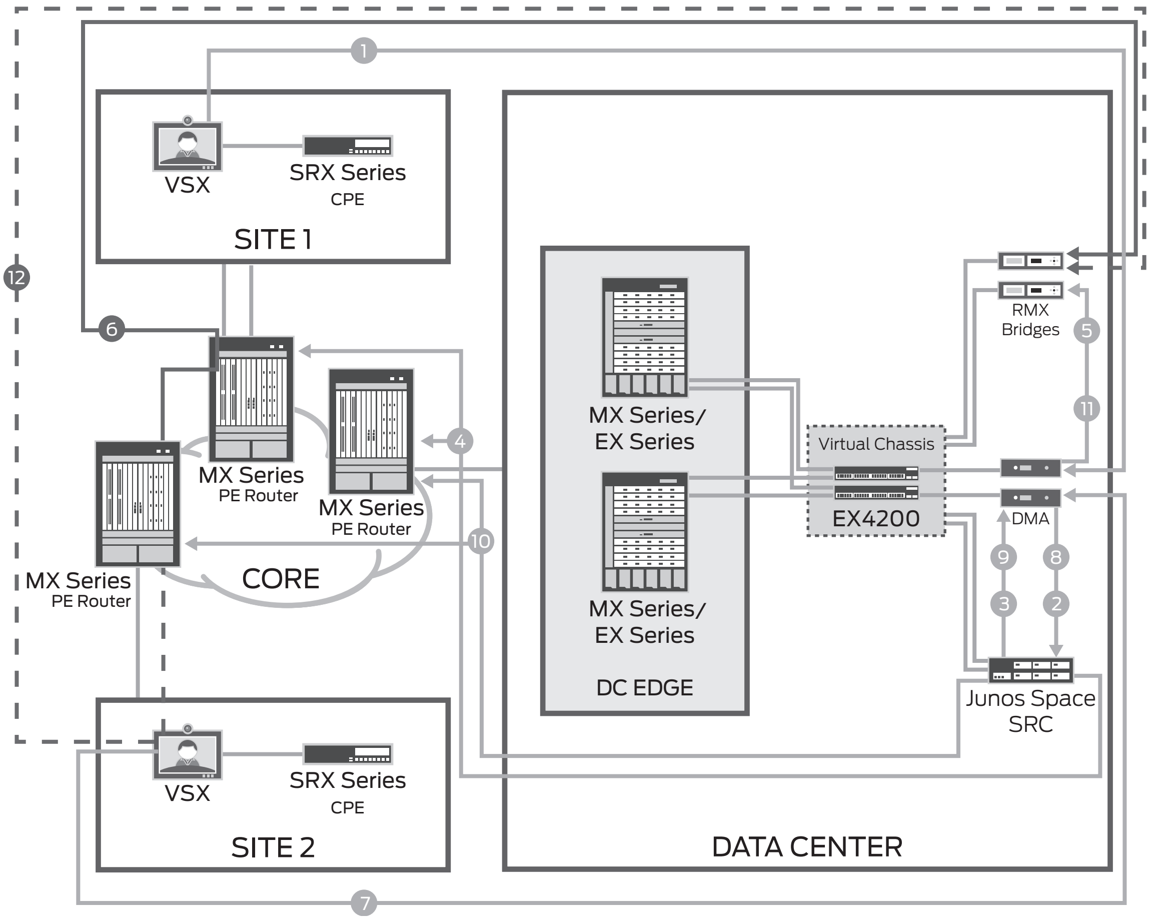

Figure 10.2 Successful point-to-multipoint call setup with CAC and Assured Forwarding

Remote Office Site 1 Dials into a Bridged Call

1. The Polycom endpoint at this premise communicates with the DMA at the data center. In the signaling portion, it sends the conference call meeting ID to the DMA.
2. The DMA identifies the endpoint, looks up the associated service tier for this endpoint and contacts the SRC (Assured Forwarding Decision Point) by sending a SOAP message. This message also includes information on the RMX that will be participating in the call. The DMA also identifies the service tier to apply to this endpoint (or meeting room). However, the call rate is negotiated as per the meeting room profile configured on the RMX bridge.
3. The SRC identifies the network routers to be used as congestion management policy enforcement points for this call, validates if bandwidth is available to allow this call and signals back to the DMA with a positive acknowledgement.
4. The SRC programs the appropriate QoS policy updates to the Juniper routers, in this case the PE router connecting the data center on one side of the call and the PE router connecting Site 1 to the core of SP network.
5. The DMA then proceeds to program the RMX that will participate in this call.

6. Site 1 is now connected to the bridge at the assured SLA.

Remote Office Site 2 Dials into a Bridged Call

7. Same as step 1 above.
8. Same as step 2 above.
9. The SRC identifies the network devices to be used for this call, validates if bandwidth is available to allow this call and signals back to the DMA with an acknowledgement.
10. The SRC programs the appropriate QoS policy updates on the ingress PE router where Site 2 is connected. A similar operation is performed on the PE router where the data center (and in turn RMX) is connected. This policy push is viewed as an incremental update at the data center PE router, and results in chaining of policies – one for each leg of the call.
11. The DMA then proceeds to update the RMX regarding the new caller joining the bridge.
12. Site 2 is now connected to the bridge at the assured SLA.

Denying a Videoconference Request

For customers who subscribe to the MC-Silver service tier, there might be times when the call request is denied due to unavailable network resources. Figure 10.3 shows a variation of the conference call setup highlighting the appropriate steps due to resource unavailability.

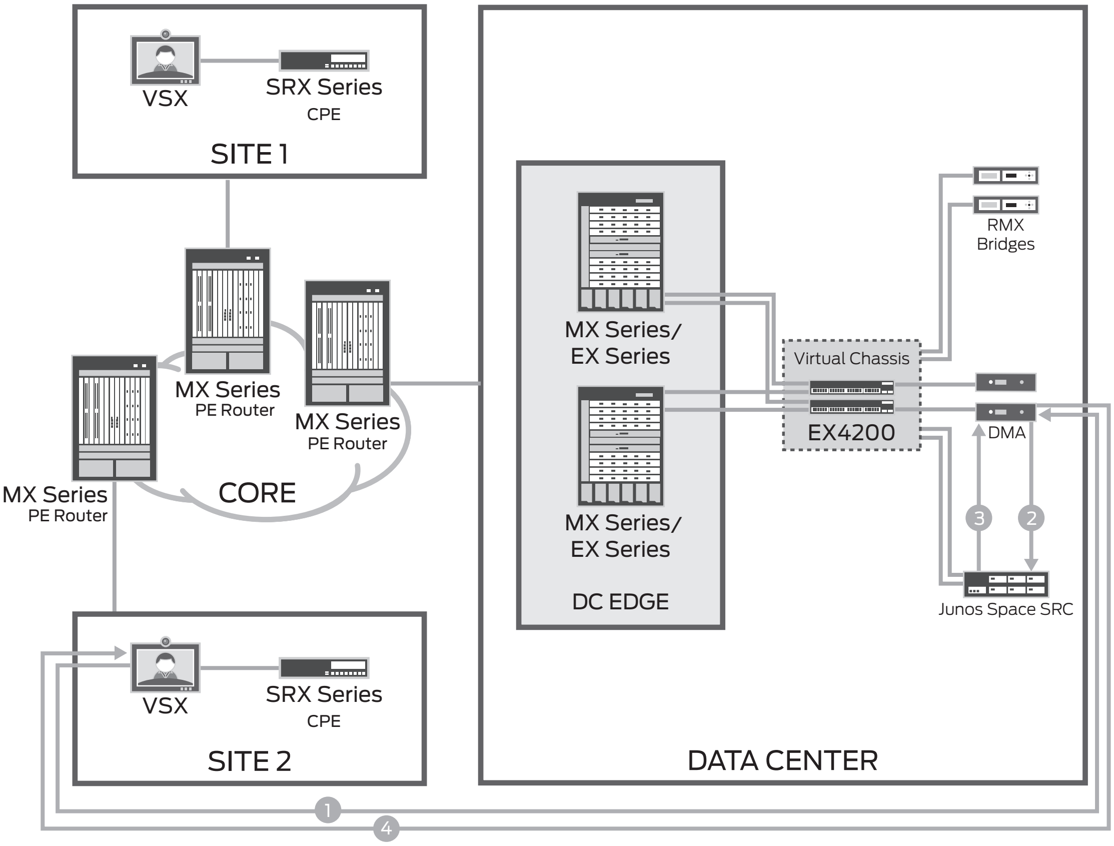

Figure 10.3 Denied call setup due to unavailable resources

NOTE: Initial steps 1-2 are the same as in the previous example.

At step 3, the SRC signals to the DMA the unavailability of the network resource through the SOAP interface.

At step 4, the DMA in turn signals back to the requesting endpoint through the SIP or H.323 signaling protocol that the call cannot be placed which results in a busy signal at the endpoint. The end user is expected to retry under these circumstances.

Implementation

This section describes the configurations required for the Juniper networking platforms to guarantee SLAs for the HD videoconferencing Assured Service levels. These SLA guarantees are established by videoconference application integration within the network policy control and management layers. This implementation enables the network to intelligently prioritize end-to-end HD videoconferencing traffic into its respective QoS, thereby enhancing user experience for those particular customers who have subscribed to a specific service tier. Figure 10.4 shows a high-level overview of where the different functional components are implemented in the HD videoconferencing network.

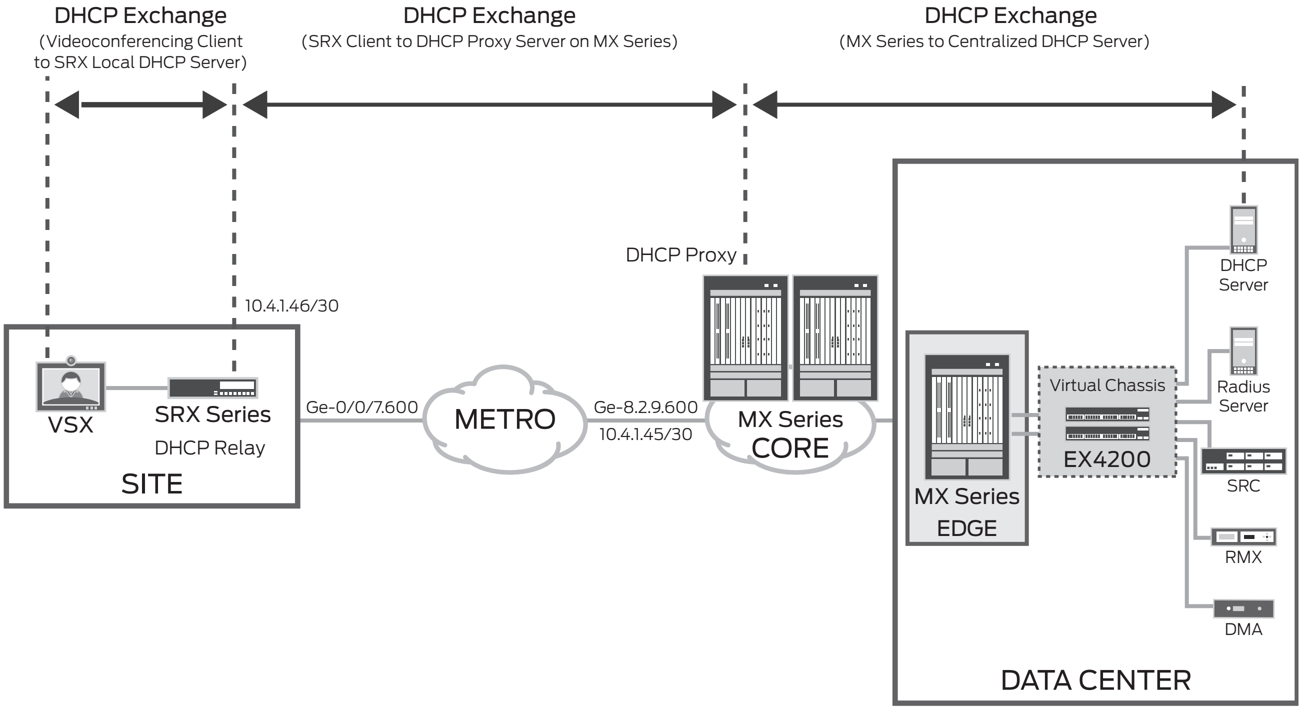

Figure 10.4 DHCP configuration for assured forwarding

In this section, we first cover the provisioning details for the PE routers, followed by the basic configurations for the CPE devices. The final part of this section focuses on the SRC configuration.

Configuring the PE Routers

The uniqueness of the HD videoconferencing solution is that it uses a converged network to deliver the assured calling experience while maximizing the usage of all network and video infrastructure. This requires that provisioning of new customers and deployments of new sites should be cost effective and scalable both at the videoconferencing datacenter and at the network elements.

The HD videoconferencing solution achieves this goal by pre-provisioning the video endpoints at the data center and letting the PE routers discover the endpoints during the DHCP address provisioning. The subscriber management process on the Junos PE router treats the video endpoint as a subscriber and creates a logical interface upon successful completion of the DHCP negotiation. When this endpoint participates in a call, QoS policies are applied to this interface resulting in prioritization of traffic, thus making it possible to meet the desired SLA.

To simplify pre-provisioning requirements of the HD videoconferencing solution, it is recommended that the endpoints be configured to request their IP address using DHCP, the CPE routers at the remote/branch/campus sites must be configured to relay the requests to the Junos PE routers that operate as DHCP proxy. These proxies communicate with the centralized DHCP server at the videoconferencing datacenter. The DHCP proxy assigns IP addresses based on MAC hardware address match rules, which in turn identifies the capabilities of the endpoint.

An alternative to discover the endpoints dynamically would be to statically provision the endpoint as a subscriber at the PE router, which would handle its traffic. Although this is not a scalable approach for videoconferencing endpoints, the number of RMX devices in the data center is well known and changes are infrequent, so static configuration is employed to discover[1] the RMX device on the PE router, which connects to the videoconferencing datacenter. The Static Subscriber Configuration (SSC) process on the Junos PE router is employed to configure subscriber interfaces for the RMX devices.

The following code snippet shows how to provision for DHCP relay at the CPE routers that connect the endpoints to the metro access network.

Configuration Snippet at the SRX for Supporting DHCP IP Address Assignment to the Endpoints

```
## Configure the DHCP-relay functionality and specify the interfaces on which
## the SRX will accept DHCP requests
[edit forwarding-options]
helpers {
    bootp {
        description “ DHCP relay for the video end-points “;
        ## ensure that the requesting client’s MAC address is relayed to the DHCP server
        ## for an hardware address match
        dhcp-option82 {
            remote-id;
        }
        ## DHCP proxy server IP address on the PE router
        server 10.100.1.185;
        ## This router will only accept DHCP requests on these interfaces
        interface {
            ge-0/0/2.600;
        }
    }
}
```

The following code snippet shows how to configure the DHCP proxy at the PE router to dynamically discover and provision the endpoint.

Configuration Snippet for Dynamically Discovering the End-points

```
[edit routing-instances smallCo]
forwarding-options {
    dhcp-relay {
        ## configure this as a proxy server
        overrides {
            proxy-mode;
        }
        server-group {
            ## centralized DHCP server IP address in the data center
            HDVC-DC-Serv-1 {
                10.12.12.100;
            }
        }
```

[1] The RMX plays an important role in conference calling as it mixes the video and voice feeds from all the participants. Traffic originating from and destined to the RMX requires the same QoS handling as the HD endpoints, and hence is statically configured as a subscriber on the Junos PE router.

```
        active-server-group HDVC-DC-Serv-1;
        group HDVC {
         ## create a logical interface based on the rules specified in this dyn-profile
            dynamic-profile campus-svlan;
            ## interface connecting the remote office
            interface ge-8/2/9.600;
            ## loopback interface is used as the DHCP proxy server's IP address
            interface lo0.185;
        }
    }
}

[edit interfaces]
ge-8/2/9 {
    description "Connected to remote site-1 of smallCo";
    hierarchical-scheduler;
    vlan-tagging;
    mtu 4096;
    unit 600 {
        ## create a logical-interface for every end-point learnt on this interface
        demux-source inet;
        vlan-id 600;
        family inet {
            address 10.8.1.45/30 {
            }
        }
    }
}
## Define the loopback interface. This will be used as the DHCP proxy server's IP address
## for the "smallCo" VR context.
lo0 {
    unit 185 {
        family inet {
            address 10.1.1.185/32;
        }
    }
}
[edit dynamic-profiles]
## This profile will dynamically discover the end-points as subscribers and creates
## logical interfaces for each of them. This is the template that will help the
## Subscriber Management process create subscriber entries for every end-point discovered
campus-svlan {
    ## All logical interfaces created will be within the context of the "smallCo"
    ## routing instance. This stanza is associating the logical interface into the
    ## desired VR context
    routing-instances {
        "$junos-routing-instance" {
            interface "$junos-interface-name";
        }
    }
    interfaces {
        demux0 {
            ## Instruct the Subscriber Management (SM) process to pick a dynamic unit
            ## number when creating the logical interface
            unit "$junos-interface-unit" {
                demux-options {
                    ## The underlying interface is used as the device to which these
                    ## logical interfaces will be attached to.
                    underlying-interface "$junos-underlying-interface";
                }
                ## The demux interface will borrow the loopback interface's IP address.
                ## The SM process will program a host specific route for the video
                ## end-point and will use the newly created logical interface as the
            ## egress to reach it.
            family inet {
               demux-source {
                  $junos-subscriber-ip-address;
               }
               unnumbered-address “$junos-loopback-interface”;
            }
         }
      }
   }
}
```

The following code snippet shows how to apply static provisioning of the RMX devices using SSC.

Configuration Snippet for Statically Provisioning of the RMX Devices on MX Router as Service Subscribers

```
[edit system services]
## static subscriber is defined and associated with the logical interface on unit 600
static-subscribers {
   group MCU-1 {
      ## Reference the access profile definition which provides the AAA process with
      ## rules to authenticate and authorize the video infrastructure at the DC
      access-profile {
         ## refer to the [edit access] configuration for a definition of this profile
         mx480-pe;
      }
      ## Reference the dynamic profile definition which provides the SM process with
      ## rules to create a subscriber entry and demux interface for the video
      ## infrastructure
      dynamic-profile {
         ## Since a static entry is being created, only a default profile is needed.
         ## This profile is internally defined.
         junos-default-profile;
      }
      ## The SSC application authenticates the video infrastructure entity with the
      ## AAA application which is needed for JSRC integration
      authentication {
         password “$9$kqfz3nCpu1zFcyKvLX”; ## SECRET-DATA
         username-include {
            domain-name mx480-pe.com;
            user-prefix MCU-1;
         }
      }
      ## The interface on which the video infrastructure’s traffic is received on this
      ## PE router
      interface ge-0/0/4.600;
   }
}

[edit interfaces]
ge-0/0/4 {
   description “ connected to data center PE router “;
   hierarchical-scheduler;
   vlan-tagging;
   mtu 4096;
   unit 600 {
      vlan-id 600;
      family inet {
         address 10.1.1.17/29;
                              }
```

```
        family iso;
        family mpls;
    }
}
```

Integrating the PE Router with the SRC

To allow the SRC to push policies to the Junos PE routers, the PE routers must be set up to communicate with the SRC. This section provides details on the configuration required to enable this communicate.

JSRC (Junos SRC) is the process in the Junos operating system that interacts with the SRC server. JSRC interacts with the Service Activation Engine (SAE) on the SRC platform using the Diameter protocol. The SAE acts as the controlling agent in the SRC environment. JSRC and the SAE jointly provide the remote control enforcement functionality (RCEF).

NOTE: Do not confuse SRC and JSRC. SRC is the standalone platform which determines whether and when to dynamically push policy updates to a Junos router. JSRC is the name of the Junos subsystem which communicates with the SRC server.

JSRC provides the following functions:

- Requests address authorization from the SAE
- Requests service activations from the SAE
- Activates and deactivate services as specified by the SAE
- Logs out subscribers as specified by the SAE
- Updates the SAE with status of new service activations and deactivations
- Synchronizes subscriber state and service information with the SAE
- Notifies the SAE when subscribers log out.

To successfully deploy and manage a reliable and scalable HD videoconferencing solution over a converged network, the network administrator must configure the following objects in the provider's network.

- JSRC functionality on the MX routers
  - JSRC Subscriber Access Profile
  - JSRC Partition
- Junos Script
  - Diameter Instance

Configuring the JSRC Subscriber Access Profile

The SRC instructs the SAE to activate and deactivate subscriber services (described by SRC policies) and log out subscribers. The SAE can control only those resources that have been provisioned through SAE. Therefore, the SAE receives information about only those subscribers for whom JSRC has requested provisioning from the SAE.

JSRC provisioning is enabled for Dynamic (discovered using DHCP) as well as static (configured using SSC) subscribers by including the provisioning-order statement. When the videoconferencing application requests AAA to activate the subscriber's session, JSRC sends an AA-Request message to request service provisioning from the SAE.

JSRC currently supports subscriber sessions on static and dynamic interfaces[2].

JSRC Subscriber Access Profile Code Snippet

```
[edit access]
## This is needed when using static subscriber authentication & authorization for JSRC
radius-server {
    10.11.11.253 secret “$9$TQnCtpBREyCAvWx7Vb”; ## SECRET-DATA
}
## Define the AAA rules and methods to use
profile mx480-pe {
    authentication-order radius;
    authorization-order jsrc;
    provisioning-order jsrc;
    radius {
        authentication-server 10.11.11.253;
    }
}
## associate the access-profile with the routing-instance – in this case “default” LS:RI
[edit]
access-profile mx480-pe;
```

Diameter Instance

The next step is to configure the diameter instance and the SRC partitions. Diameter is used as the communication protocol between the PE routers and the SRC, thus enabling the SRC to deliver granular dynamic policy enforcement on a per-service level. Before configuring the JSRC partition, configure the Diameter instance.

Diameter peers communicate over a TCP connection by exchanging Diameter messages that convey status, transmit requests and acknowledgments by means of standard Diameter Attribute Value Pairs (AVPs) and application-specific AVPs. The Diameter transport layer configuration is based on Diameter Network Elements (DNEs).

[2] JSRC Authorization for Dynamic Subscribers: JSRC authorization is enabled for dynamic (DHCP) subscribers by including the authorization-order jsrc statement. This setting causes AAA to ignore the authentication order setting in the access profile. As a result, AAA does not authenticate the DHCP subscribers. For non-DHCP subscribers, AAA ignores the authorization-order statement.

JSRC Authorization for Static Subscribers: We can associate subscribers with statically configured interfaces and provide dynamic service activation and activation for these subscribers. When the static interface comes up, the event is treated as a subscriber login. When the interface goes down, it is treated as a subscriber logout. After the subscribers are present in the session database (SDB), JSRC can report the subscribers to the SAE so that the SRC software can subsequently manage the subscribers.

Diameter Instance Code Snippet

```
[edit diameter]
## local peer characteristics
origin {
   realm solutions.juniper.net;
   host mx480-pe-28;
}
## Configure a network element by associating the peers to communicate with.
## Each DNE consists of a prioritized list of peers and a set of routes that define how
## traffic is forwarded. Each route associates a destination with a function, a function
## partition, and a metric. When an application sends a message to a routed destination,
## all routes within the Diameter instance are examined for a match. When the best route
## to the destination has been selected, the message is forwarded by means of the DNE
## that includes that route.
network-element DNE-1 {
   peer SOL-SRC {
      priority 1;
   }
   forwarding {
      route r1 {
         function jsrc partition p1;
         destination realm solutions.juniper.net host SOL-SRC;
         metric 1;
      }
   }
}
## define a peer – in this case the SRC at the data center
peer SOL-SRC {
   address 10.70.70.254;
   connect-actively;
}
```

JSRC Partition

JSRC works within a specific logical system such as routing instance context and is called a partition. Configuration for the JSRC partition consists of naming the partition and then associating a Diameter instance, the SAE host name and the SAE realm with this partition.

JSRC Partition Code Snippet

```
[edit jsrc]
partition p1 {
   diameter-instance master;
   destination-realm solutions.juniper.net;
   destination-host SOL-SRC;
}
```

Dynamic QoS on the PE Router

Once communication between the SRC and Junos PE router is enabled, a dynamic profile must be configured on the router(s) that will receive the configuration parameters based on the policy push request received from the SRC using the JSRC interface. A dynamic profile is a set of characteristics, defined in a type of a template, which is used to provide subscriber access and quality of service for HD videoconferencing application. These templates are assigned dynamically to subscriber interfaces. The dynamic-profiles hierarchy contains many configuration statements that can be normally defined statically. The dynamic profile template will substitute the values it receives in the policy push request to the variables in the fast update filter on the PE router. The policy push request references the dynamic service profile that is configured on the PE router - “HD-TP”, as shown in the following code snippet.

NOTE: The Fast Update Filter provides support for subscriber-specific filter values as opposed to classic filters, which are interface-specific. Individual filter terms can be added or removed without requiring filter recompilation after each modification.

Dynamic Profile Configuration Snippet

```
[edit dynamic-profiles]
## service profile that is used to dynamically provision the QoS
HD-TP {
   ## Variables defined here will be substituted with values received from
   ## the SRC policy push request
   variables {
      dst;
      src;
      forwarding;
      bandwidth;
   }
   ## A filter is being attached to the demux logical interface
   interfaces {
      “$junos-interface-ifd-name” {
         unit “$junos-underlying-interface-unit” {
            family inet {
               filter {
                  input SRC_Driven_filter;
               }
```

```
                }
            }
        }
    }
    ## This Fast Update Filter is constructed by substituting the variable values
    ## received in the policy push message. In this case the term identifies the
    ## end-point IP addresses and/or the RMX IP address
    firewall {
        family inet {
            fast-update-filter SRC_Driven_filter {
                ## The filter has to be tied to an interface. This is a required
                ## statement for FUF
                interface-specific;
                ## The match-order determines how the match tables for the terms are
                ## built.
                ##
                match-order [ destination-address source-address ];
                term 1 {
                    from {
                        source-address $src;
                        destination-address $dst;
                    }
                    then {
                        policer “$bandwidth”;
                        ## Count how many times this term was hit. This is helpful in
                        ## troubleshooting during initial deployment – to check if the
                        ## HD videoconferencing traffic is being processed correctly
                        count Dynamic_Policy;
                        forwarding-class “$forwarding”;
                        accept;
                    }
                }
                term 3 {
                    then {
                        ## This counter accounts for any other traffic that the
                        ## end-points are sending/receiving – an indication for
                        ## misconfiguration or network attacks
                        count Background_traffic;
                        accept;
                    }
                }
            }
        }
    }
}
```

The predefined service policies pushed down from the SRC are translated to the predefined QoS definitions on the PE routers. Please refer to the section on Egress Processing in *Chapter 9, Implementing Quality of Service on page 115* for the snippet showing how to configure QoS definitions on the PE router.

SRC Provisioning

The SRC software operates on the C Series network appliance. In the HD videoconferencing data center application, the SRC is network-aware and provisions an end-to-end QoS for the video call while ensuring that the requested SLA is met. The SRC interacts with Polycom's Distributed Media Application (DMA) Server and performs the following functions:

- Maintains information about videoconferencing endpoints
- Maintains the information about the topology of the network infrastructure and resource availability
- Provides CAC functionality to videoconferencing application
- Enforces dynamic QoS policies on the network routers
- Acts as SOAP gateway (interface with Polycom DMA).

Various components in the SRC software perform the above-listed functions. This section briefly describes each component and provides a configuration snippet to configure the component for its role in the HD videoconferencing solution. Note that these configurations are provided to educate the reader with a sense of the various moving parts in the SRC subsystem.

Network Information Collector (NIC)

The NIC collects information about the state of the network and can provide a mapping from one type of data network to another type of data network.

Pre-configured NIC Configuration Scenario.

In this scenario, the router initialization script associates each VR with the SAE that manages it.

```
[edit shared nic scenario OnePopStaticVrfIp]
nic-locators {
  vrfIp {
    resolution {
      constraints AnyString(vpn);
      expect-multiple-values;
      key-type Ip;
      resolver-name /realms/vrfIp/A1;
      value-type SaeId;
    }
  }
}
realms {
  vrfIp {
    configuration {
      custom-resolver {
        classname {
          Ip-IpPool net.juniper.smgt.gateway.nic.resolver.MultiValueCompatibleResolver;
          IpPool-Interface net.juniper.smgt.gateway.nic.resolver.MultiValueMappingResolver;
        }
      }
    }
    resolvers A1 {
      configuration {
        resolver-role RoleA;
      }
    }
    resolvers B1 {
      configuration {
        resolver-role RoleB;
        roles-list RoleA;
      }
    }
    resolvers C1 {
      configuration {
        resolver-role RoleC;
      }
    }
    resolvers D1 {
      configuration {
        resolver-role RoleD;
      }
    }
  }
}
```

Service Activation Engine (SAE)

The SAE authorizes, activates and deactivates subscriber and service sessions by interacting with systems such as Juniper Networks routers. The SAE also provides plug-ins and APIs for starting and stopping subscriber and service sessions and for integrating with other systems that authorize subscriber actions and track resource usage.

SAE Configuration Snippet

```
 [edit shared sae group POP-ID configuration]
## SAE drivers configuration
driver {
  junos-ise {
    cached-driver-expiration 600;
    keep-alive-timeout 60;
    registry-retry-interval 30;
    reply-timeout 20;
    sae-community-manager ISECommunityManager;
    sequential-message-timeout 20;
    session-store {
      maximum-queue-age 100;
    }
    thread-idle-timeout 60;
    thread-pool-size 50;
  }
}
nic-proxy-configuration ip {
  resolution {
    key-type Ip;
    resolver-name /realms/vrfIp/A1;
    value-type SaeId;
  }
  test-nic-bindings {
    key-values {
      class nicproxy;
      useNicStub false;
    }
  }
}
nic-proxy-configuration vr {
  resolution {
    key-type Vr;
    resolver-name /realms/vrfIp/A1;
    value-type SaeId;
  }
  test-nic-bindings {
    key-values {
      class nicproxy;
      useNicStub false;
    }
  }
}
```

Dynamic Service Activator (DSA)

The DSA enables external applications, such as the Polycom DMA to dynamically activate services or run scripts on an SRC's SAE through the SRC's SOAP gateway. For managing services, DSA supports a fixed set of methods and uses the SAE access interface module to access the SAE core API.

DSA Configuration Snippet

```
[edit shared dsa]
group sample {
  configuration {
    client Joe {
      permissions {
        method allocate-resource;
        method commit-resources;
        method invoke-gateway-extension;
        method invoke-script;
        method query-available-services;
        method query-client-status;
        method query-contexts;
        method query-status;
        method release-network-resource;
        method release-resource;
        method release-resources;
        method reserve-network-resource;
        method subscriber-activate-service;
        method subscriber-deactivate-service;
        method subscriber-login;
        method subscriber-logout;
        method subscriber-modify-service;
        method subscriber-read-subscription;
        method subscribers-read;
        method subscribers-read-subscriber;
        pcmm-service [ Video-Silver PCMM-Down ];
        script Echo;
      }
    }
    disable-soap-client-authentication;
    nic-proxy-configuration {
      assignedIp {
        cache {
          cache-cleanup-interval 10;
          cache-size 0;
        }
        resolution {
          key-type Ip;
          resolver-name /realms/vrfIp/A1;
          value-type SaeId;
        }
        test-nic-bindings {
          key-values {
            class nicproxy;
            useNicStub false;
          }
        }
      }
   }
```

Diameter Interface

The Diameter Interface component serves as the communication layer between the SRC and the JSRC component on the Juniper PE routers.

Diameter Configuration

Protocol used to communicate with MX routers in the service provider's network

```
## configure the local diameter settings
[edit system diameter]
active-peers;
local-address 10.11.11.254;
origin-host hdvc-src;
origin-realm solutions.juniper.net;
port 3868;
protocol tcp;
## Configure the Junos routers that the SRC will communicate via diameter peering
[edit shared network diameter]
peer mx-north {
  active-peer;
  address 10.70.70.1;
  connect-timeout 10;
  origin-host mx-north;
  port 3868;
  protocol tcp;
}
```

Admission Control Plug-In (ACP)

This plug-in authorizes and tracks subscribers' use of network resources associated with services that the SRC application manages.

ACP Configuration Snippet

```
[edit shared acp]
group config {
  configuration {
    acp-options {
      backup-database-maximum-size 100m;
      backup-directory var/backup;
      event-cache-size 1000;
      interface-tracking-filter interfaceName=*;
      mode backbone;
      network-bandwidth-exceed-message ‘2: Network BandWidth exceeded’;
      overload-method 0;
      remote-update-database-index-keys ‘interfaceName, routerName, portId’;
      reservation-timeout 10000;
      state-sync-bulk-size 100;
      subscriber-bandwidth-exceed-message ‘1: User Bandwidth exceeded’;
      tuning-factor ‘’;
    }
    nic-proxy-configuration {
      nicProxyVrToSae {
        resolution {
          key-type Vr;
          resolver-name /realms/ip/A1;
          value-type SaeId;
        }
      }
      assignedIp {
        resolution {
          key-type Ip;
          resolver-name /realms/vrfIp/A1;
          value-type SaeId;
        }
      }
    }
}
```

Troubleshooting

This section uses the simplified network reference diagram (see Figure 11.4) to highlight the commonly encountered problems and misconfigurations in deploying the solution. The assumption here is that the HD videoconferencing solution is an incremental configuration to an existing converged SP's network connecting remote, branch and campus enterprise sites.

Ensuring Physical and Network Connectivity throughout the Network

After you have incrementally installed and configured the entire network infrastructure and the data center video applications, ensure that you have connectivity between the following solutions elements:

- CPE devices in the customer premise connected to the DHCP proxy/server
- PE routers that connect to the SRC server in the data center
- PE router in the data center that connects to the RADIUS server in the videoconferencing datacenter
- SRC that communicates with the DMA server
- Videoconferencing endpoints that communicate with DMA server.

To identify connectivity issues, perform the following:

Verifying Videoconferencing Endpoint Device Connectivity

Ensure that the endpoints can acquire their IP addresses through DHCP, and that the endpoints can ping the configured DMA devices. If this works, check to ensure that devices register with the SIP server or H.323 gatekeeper on the DMA. To troubleshoot endpoint connectivity issues, refer to the individual chapters on premise connectivity and device documentation for required troubleshooting steps.

Subscriber (Endpoint) Discovery on the Junos PE Routers

The PE routers learn the video endpoints during the DHCP exchange and provision them as subscriber entries. The following operational commands provide details on how to ensure that the concerned PE router has information on the endpoint.

PE router operational commands to validate subscriber learning

```
## Ensure that the end-point has been added as a subscriber
user@PE-West> show subscribers detail
Type: dhcp
User Name: HDX8000-A1@A1.com
Logical System: default
Routing Instance: BigCo
Interface: ge-1/2/0.342342
Interface type: demux
Dynamic Profile Name: campus-svlan
State: Active
Radius Accounting ID: 84
Login Time: 2010-10-08 14:09:47 EDT
```

PE router operational command to validate subscriber statically learnt (on the data center PE)

```
## Ensure that the end-point has been added as a subscriber
user@PE-North> show subscribers detail
Type: STATIC-INTERFACE
User Name: MCU-1@mx480-pe.com
Logical System: default
Routing Instance: default
Interface: ge-0/0/4.600
Interface type: Static
Dynamic Profile Name: junos-default-profile
State: Active
Radius Accounting ID: 1281
Login Time: 2010-10-07 12:01:27 EDT
```

Monitoring Call Setup

When a call has been successfully signaled and set up, the PE router connecting the endpoint to the network will have the QoS and firewall attachments at the interface for the entire duration of the call. In the example shown below, a test counter is attached to ensure that the video traffic arriving is marked with the correct DSCP bits at the PE router, and the interface operational commands are used to verify that the PE is processing this traffic in the right queue.

Attaching a test counter at the PE router to confirm that the arriving traffic is marked correctly from the CPE

```
[edit interfaces]
ge-1/2/0 {
   description CONECTED-TO-CAMPUS-165;
   per-unit-scheduler;
   vlan-tagging;
   mtu 4096;
   unit 600 {
      vlan-id 600;
      family inet {
         filter {
            input test-1;
         }
         address 10.4.1.17/30;
      }
   }
}
[edit firewall]
filter test-1 {
   term 1 {
      from {
         dscp af41;
      }
      then {
         count af41-count;
         accept;
      }
   }
   term 2 {
      then {
         count be-count;
         accept;
      }
   }
}
{master}[edit]
chandra@MX-West-188-re0# run show firewall
Filter: __default_bpdu_filter__
Filter: test-1
Counters:
Name                                            Bytes         Packets
af41-count                                   50213292          175294
be-count                                        25784             223
```

PE router operational commands to monitor successful call setup

```
## Ensure that the interface has the QoS and filter attachments, and that the HD videoconferencing traffic
is hitting the right queues
chandra@MX-West-188-re0> show firewall
Filter: __default_bpdu_filter__
Filter: SRC_Driven_filter-ge-1/2/0.600-in
Counters:
Name                                            Bytes         Packets
Background_traffic-206                          66215             176
Dynamic_Policy-206                            6349801           18252
Policers:
Name                                          Packets
2.4MB-1-206                                         0


{master}[edit]
chandra@MX-West-188-re0# run show interfaces queue ge-1/2/0
Physical interface: ge-1/2/0, Enabled, Physical link is Up
  Interface index: 415, SNMP ifIndex: 828
  Description: CONECTED-TO-CAMPUS-165
Forwarding classes: 16 supported, 4 in use
Ingress queues: 8 supported, 4 in use
Queue: 0, Forwarding classes: MC-Bronze
  Queued:
    Packets             :                          10           0 pps
    Bytes               :                        1200           0 bps
  Transmitted:
    Packets              :                          10           0 pps
    Bytes                :                        1200           0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                           0           0 pps
     Low                 :                           0           0 pps
     Medium-low          :                           0           0 pps
     Medium-high         :                           0           0 pps
     High                :                           0           0 pps
    RED-dropped bytes    :                           0           0 bps
     Low                 :                           0           0 bps
     Medium-low          :                           0           0 bps
     Medium-high         :                           0           0 bps
     High                :                           0           0 bps
Queue: 1, Forwarding classes: MC-Gold
  Queued:
    Packets              :                       12329         378 pps
    Bytes                :                     4699486     1172792 bps
  Transmitted:
    Packets              :                       12329         378 pps
    Bytes                :                     4699486     1172792 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                           0           0 pps
     Low                 :                           0           0 pps
     Medium-low          :                           0           0 pps
     Medium-high         :                           0           0 pps
     High                :                           0           0 pps
    RED-dropped bytes    :                           0           0 bps
     Low                 :                           0           0 bps
     Medium-low          :                           0           0 bps
     Medium-high         :                           0           0 bps
     High                :                           0           0 bps
Queue: 2, Forwarding classes: MC-Silver
  Queued:
    Packets              :                           0           0 pps
    Bytes                :                           0           0 bps
  Transmitted:
    Packets              :                           0           0 pps
    Bytes                :                           0           0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets  :                           0           0 pps
     Low                 :                           0           0 pps
     Medium-low          :                           0           0 pps
     Medium-high         :                           0           0 pps
     High                :                           0           0 pps
    RED-dropped bytes    :                           0           0 bps
     Low                 :                           0           0 bps
     Medium-low          :                           0           0 bps
     Medium-high         :                           0           0 bps
     High                :                           0           0 bps
Queue: 3, Forwarding classes: NC
  Queued:
    Packets              :                           0           0 pps
    Bytes                :                           0           0 bps
  Transmitted:
    Packets              :                           0           0 pps
    Bytes                :                           0           0 bps
    Tail-dropped packets : Not Available
```

Appendices

Appendices

Appendix A: Videoconferencing Technology Primer

The quality and quantity of videoconferencing sessions depends on the capabilities of the underlying network. Not surprisingly, high definition videoconferencing requires significantly more bandwidth than traditional video. There are numerous factors which affect the amount of bandwidth required, which can vary from under 512 Kbps to upwards of 20 Mbps. Understanding the network requirements, including ways to reduce the bandwidth requirements, requires an understanding of how videoconferencing sessions are established and the key characteristics of video transmission.

This section discusses video standards and digital video technology. It introduces key videoconferencing characteristics that affect bandwidth, methods that reduce the amount of bandwidth required and the two standardized control protocols commonly used by videoconferencing systems.

Video Standards

Many attributes define digital video formats such as:

- Resolution
- Scanning
- Refresh rate
- Aspect ratio

Resolution

The video image consists of a series of dots, better known as pixels. The number of pixels transmitted determines the picture quality. Common high definition resolutions include 1920 columns by 1080 rows (1920 x 1080) and 1280 x 720. HD standards are consistent worldwide.

However, this is not the case for lower resolution standard definition (SD) video, which was designed to maintain compatibility with existing analogue video. These include 704 x 576 (for compatibility with countries that use the analog PAL or SECAM systems, including most of Europe) and 704 x 480 (for compatibility with countries that use the analog NTSC systems, including North America, Japan and South Korea).[1]

High definition videoconferencing systems typically support both HD resolutions and may support some SD resolutions as well.

Scanning Mode

Analog video fools the eye by displaying every other line (the “odd” lines) on one pass, then filling in the other (“even”) lines on the second pass. This technique, called interlaced scanning (shortened as “i”), continues to be used in some modern systems. Modern video systems display all lines consecutively. This method is called progressive (p) scanning.

[1] There are many SDTV variations based on specific application. The width may be 720 pixels instead of 704, and the NTSC height may be 486 pixels instead of 480. The values shown are most common for end-user applications including videoconferencing.

All HD videoconferencing systems support progressive scanning and some support interlaced.

Refresh Rate

Another attribute of the video stream is how many times per second the image on the screen is updated. The refresh rate for analog televisions were originally tied to the power grid, so screen updated occurred 50 (Europe) or 60 (North America) times per second. One update per second is equal to one Hertz. With interlaced scanning, two passes are required to update the screen entirely, so the refresh rate is 25 or 30 Hertz.

Many systems also support 24 Hertz, which is used for movies and other professionally generated content. This allows videoconferencing systems to deliver this content "on demand" to interested parties.

Videoconferencing systems support one or more of these rates.

Aspect Ratio

The aspect ratio describes the screen's proportions. Analog video and standard definition use a 4:3 (also called 1.33:1) aspect ratio, meaning that the width is 1.33 times that of the height. High definition video standards, including HDTV and many videoconferencing systems, have a 16:9 (which is equivalent to 1.78:1) aspect ratio.

All HD videoconferencing systems support 16:9, but many also support 4:3. See Figure A.1.

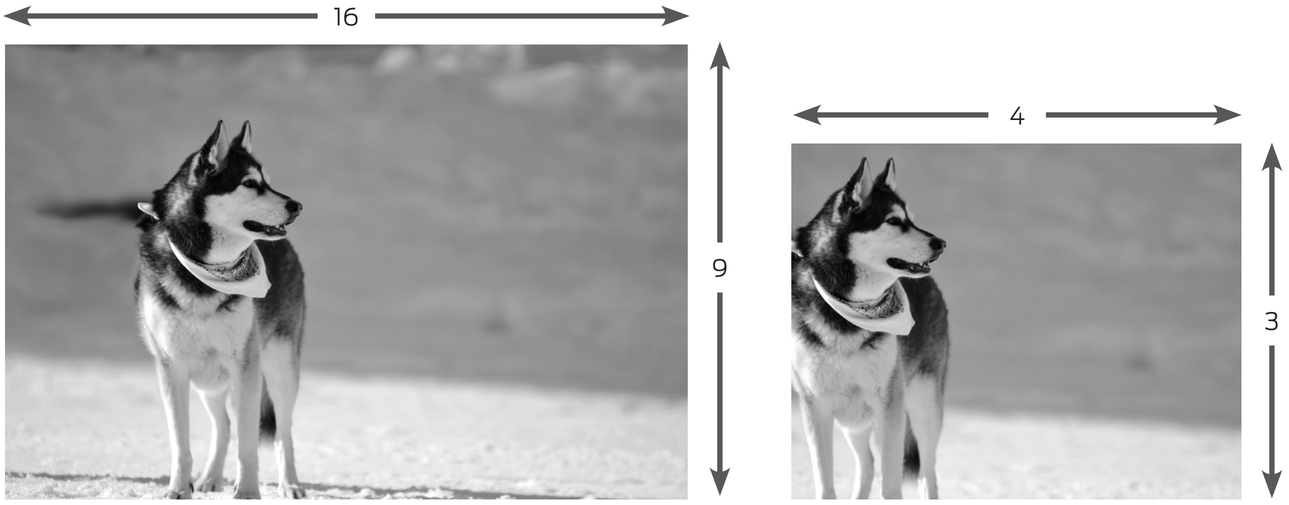

Figure A.1 Aspect ratios

Since the aspect ratio is often known, the pixel size is often written in shorthand by specifying only the number of rows, for example 1080. The number of columns can be calculated by multiplying this by the aspect ratio. The video signal is written in shorthand by specifying the resolution, scanning method and refresh rate, for example 1080p30. Another common shorthand is to omit the refresh rate since the system can support multiple refresh rates, for example 1080p. Table A.1 summarizes some common video formats used for videoconferencing as well as digital television.

Table A.1 Common Video Formats

Name	Frame Size/Scanning	Aspect Ratio	Refresh Rate (frames/second)
High Definition (HD)	1080p, 1080i, 720p	16:9	24, 25, 30, 50, 60
Enhanced Definition (ED)	480p, 576p	16:9 or 4:3	
Standard Definition (SD)	480i, 576i	16:9 or 4:3	

NOTE: These are the most common values for videoconferencing and television, but many other options exist. For example, video to cell phones may have lower resolution (such as 352 × 288) and refresh rates (15 fps). On the other hand, systems covering large areas may double or quadruple the horizontal and vertical resolution.

Adding Color

Another factor that has a major impact on bandwidth requirements is how color information is transmitted. In the simplest case, a fixed number of bits are used to determine the brightness of each of the three-color signals (red, green and blue, or RGB) which make up each pixel. Using 8 bits per color signal or 24 total bits, yields 2^{24} = 16 million different possible colors for each pixel[2].

There are numerous widely adopted techniques to minimize the amount of color information, which must be sent (or conversely, allow more colors to be represented with the same number of bits). One widely implemented technique for "consumer-level" video is to reduce how frequently each color is sampled. Fully sampling each color is written as 4:4:4 sampling. In other words, over a given period, each color is sampled four times, totaling 12 samples. To reduce bandwidth, sampling can occur less frequently, commonly using 4:2:0 sampling, totaling 6 samples. Using this technique requires $^6/_{12}$ (or $^1/_2$) the bandwidth required by full sampling.

Bandwidth Requirements

A full (uncompressed) 1080p60 video signal requires bandwidth as follows:

Bandwidth required = pixels per frame * frames per second * bits/pixel * sampling factor

= (1920 * 1080) * 60 * 24 * $^1/_2$
= 1.5 Gbps

Audio and control traffic add slightly to the bandwidth requirements, but those are insignificant compared to the video bandwidth.

This amount of information would bring most networks to their knees. Because few organizations can afford to dedicate this much bandwidth (per connection) to a service that is used sporadically, the bandwidth required must be reduced significantly in order for videoconferencing to be economically viable.

[2] Reproducing color effectively is a complex topic, which is not done justice in this simplified explanation. However, the information provided is sufficient for a discussion about bandwidth consumption.

Reducing the Bandwidth Burden

To reduce the amount of bandwidth required, video signals are compressed. The process of reducing the signal by removing unnecessary information is called encoding; at the other end the signal must be decoded to recreate the original signal. The equipment, which performs the encoding and decoding, is called a codec (which is shorthand for coder/decoder).

All codecs reduce the overall video information by transmitting only the information that has changed since the last frame. Videos depicting quickly changing images, such as a fast action basketball game or leaves blowing in the wind, require significant bandwidth. Streams conveying less motion, such as videoconferences, often consume far less bandwidth. Encoded streams are therefore naturally Variable Bit Rate (VBR), meaning that the amount of bandwidth being sent varies throughout the life of the session. Some systems will attempt to send a consistent amount of information (Constant Bit Rate, or CBR). Most systems also limit the total amount of information which can be transmitted (capped VBR). Attempting to send more encoded video information than can be handled—either exceeds the cap or is dropped by the network—results in lost packets. The viewer sees this as a freezing of some portion of the picture or as macroblocking or pixelation, as shown in Figure A.2.

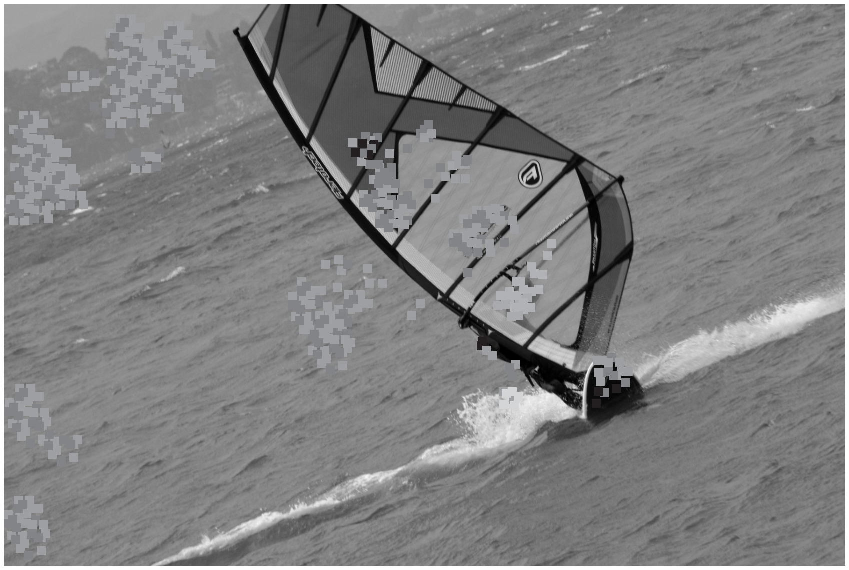

Figure A.2 Digital video with severe macroblocking

Video Codecs

There are dozens of video codecs available. Some are freeware, while others must be purchased or licensed. Most modern commercial video systems use MPEG-4 Advanced Video Codec (AVC), which is Part 10 of the MPEG-4 standard. This technique was later adopted with minor changes by the ITU-T as the H.264 standard. MPEG-4 enhances the earlier MPEG-2, which was used by earlier videoconferencing systems as well as "over the air" broadcast television. Encoded MPEG-4 streams require significantly less bandwidth than MPEG-2 streams. While the exact savings vary widely based upon content and which types of MPEG-4 techniques are implemented, the rule of thumb is that MPEG-4 reduces bandwidth by at least 50 percent.

Within MPEG-4, profiles specify combinations of techniques that are implemented, many of which affect the total bandwidth required. There are 17 profiles defined, and several levels within each profile. Fortunately, profiles are typically targeted at specific applications, and only a few are used by videoconferencing systems. Products implementing high profile require less bandwidth than those using baseline profiles.

MPEG systems transmit three types of frames. Information frames (I-frames) transmit the entire screen image and account for as much as 40 percent of the total bandwidth. Predictive (P-frames) and Bi-directional (B-frames) frames update this picture as required. A series of MPEG packets, bounded at both ends by an I-frame, is known as a Group of Pictures (GOP). Use of P- and B-frames is optional, and many real-time implementations do not use B-frames. See Figure A.3, which shows a group of P-, B- and I- frames.

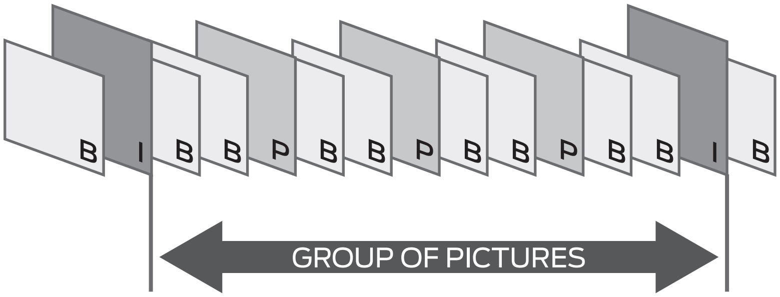

Figure A.3 Group of Pictures

Within a profile, the support level specifies the maximum picture resolution, frame rate and bit rate that can be used.

In IP networks, the encoded MPEG stream is broken down into 184-byte chucks of information. Seven chunks, plus additional control and timing information, are then packed into a single IP/Ethernet frame for transmission across the network. The additional information was originally a four-byte header added to each chunk. This was the method used by MPEG-2, and is known as an MPEG-2 transport stream (MPEG-2 TS). A more recent alternative allows this information to be carried in a single RTP header, which is added before the first 184-byte chunk. Figure A.4 shows an Ethernet/IP frame carrying MPEG-4 video, which is the format commonly used by videoconferencing systems.

Ethernet Header (6)	IP Header (20)	UDP Header (20)	RTP Header (12)	MPEG #1 (184)	MPEG #2 (184)	MPEG #3 (184)	MPEG #4 (184)	MPEG #5 (184)	MPEG #6 (184)	MPEG #7 (184)	FCS (4)
				← Payload: 7 X 184 bytes = 1288 bytes →							
← Ethernet Frame:1346 bytes →											

Figure A.4 Ethernet/IP frame carrying MPEG-4 video

Some video systems use multiple codecs to connect to multiple locations or even to send a single large image to remote locations.

Audio Codecs

Audio signals are similarly compressed. Examples of common audio codecs include:

- Advanced Audio Codec (AAC), which is often used with MPEG-4
- Siren, developed by Polycom is the latest generation. Siren 22 is also integrated into ITU G.719.
- MPEG Layer 3 (MP3), commonly used for music
- Windows Media Audio (WMA), developed by Microsoft
- AC-3, which is the audio standard for broadcast television in North America
- G.7XX (G.711, G.729), which are primarily used for speech

Encoded audio requires significantly less bandwidth that video. For example, MP3 music download services often encode music at 128 Kbps, while voice codecs can require less than 10 Kbps.

Audio codec technology is well established, and audio is a small fraction of the overall bandwidth required.

Videoconferencing Equipment

Videoconferencing equipment can be broadly broken into three categories:

- Endpoints
- Call control
- Videoconferencing bridges

NOTE: Although there is also a variety of ancillary equipment, such as recording/playback servers and Internet gateways, these topics will not be addressed because they are not required for the baseline solution.

Endpoints

Videoconferencing endpoints can be categorized into the following three types of systems:

- Fully immersive systems are the most advanced systems. These systems support HD viewing of multiple remote rooms. Participants can walk around, draw on whiteboards and otherwise interact effectively.
- Immersive (table) systems allow workers are remote locations to interact as if they are gathered around a single conference table.
- Conference room (or just “room”) systems are the next generation of the traditional videoconferencing systems, but incorporating high definition video and audio. These systems provide crystal-clear imagery while allowing the viewer to see all the remote participants in separate panes on a single display.

- Personal systems typically reside on the desktop of individual users, rather than being a shared resource. These range from PC clients to standalone desktop devices to IP phones with an integrated monitor. These typically support lower resolution, but the sheer number of devices can lead to significant bandwidth consumption.

Appendix D: Polycom Products on page 176, provides examples of each of these products.

Bridging Videoconferencing Calls

Conferences can occur in two ways: A direct connection can be established between any two endpoints on the same network, allowing only two parties to communicate. In videoconferencing lingo, this is called a point-to-point connection. However, because book is geared towards networking personnel, this term has a different meaning. Figure A.5 shows a direct (point-to-point) connection between two video endpoints.

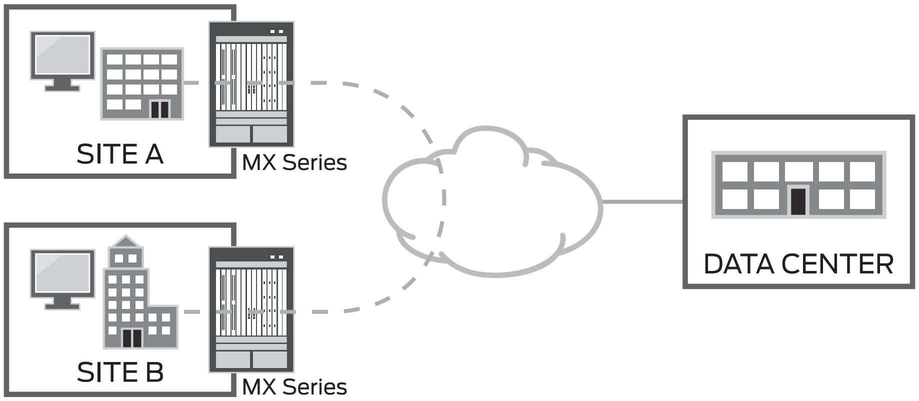

Figure A.5 Direct (point-to-point) connection

The alternative is a bridged connection. In this case, each endpoint has a connection to a videoconferencing bridge, known generically as a Multi-Conferencing Unit (MCU). In videoconferencing lingo, this is called a point-to-multipoint connection, but again this term has a different connotation for those focused on networking. Bridged video (or audio) conferences are typically established by having all users dial a conference number assigned to the meeting organizer. This number is a Virtual Meeting Room (VMR).

Figure A.6 shows three endpoints which are connected to a single MCU to create a bridged (point-to-multipoint) call. Note that each videoconference has a point-to-point connection to the MCU. Videoconferencing traffic is *not* transmitted using multicast.

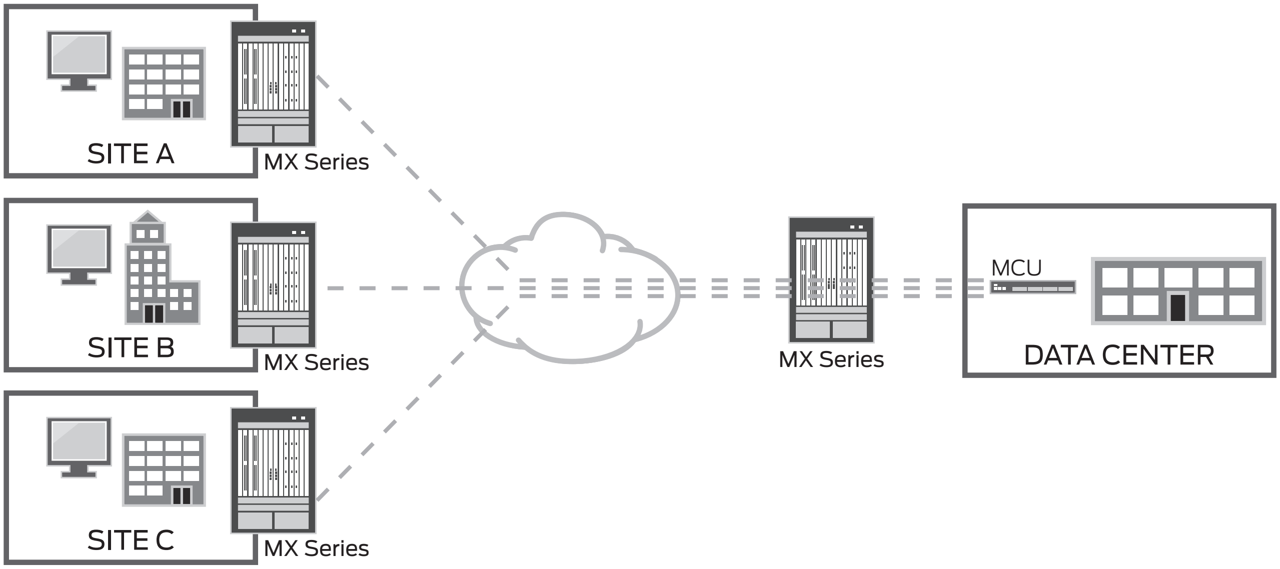

Figure A.6 Bridged (point-to-multipoint) connection

Advanced solutions allow users to connect to an MCU in their region, and the MCUs are then interconnected. This interconnection minimizes traffic across expensive WAN links. For example, two users in Hong Kong and two users in Amsterdam can connect to their local MCUs. These MCUs are then interconnected, reducing bandwidth cost by carrying each conferencing stream only once across the expensive long-haul link. This is illustrated in Figure A.7.

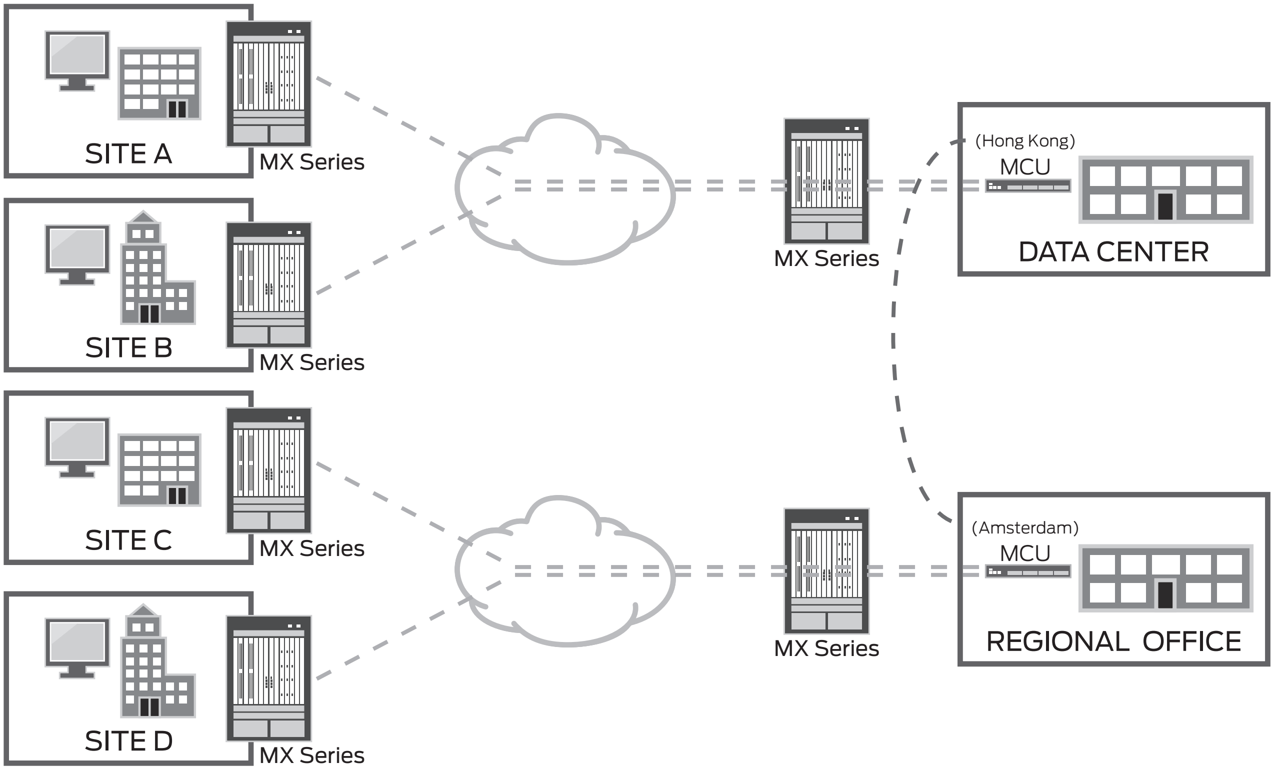

Figure A.7 Inter-domain connections

Call Control

Another important capability is the ability to establish and tear down conferencing connections. There are two common signaling standards: H.323 and SIP. These signaling protocols in turn carry information about the other protocols that are used to transport the video, audio and other capabilities.

Originally developed for VoIP call control, SIP is also used by newer videoconferencing systems. Like other IP-focused protocols, SIP information is carried as text. This lightweight protocol was designed to carry information easily about other protocols.

H.323 is the older but more common mechanism used for videoconferencing. This ITU-T standard was designed to work specifically with other ITU-T standards and has fallen out of favor due to its complexity, Information is carried in binary format, rather than as legible text. However, as the original standard signaling mechanism, which supported videoconferencing, H.323-based systems are widely deployed.

Most carriers and enterprises use SIP to support VoIP services, and prefer to eliminate H.323 from their network. However, videoconferencing systems implement additional capabilities, such as the ability to move the camera at a remote location, which are not yet supported by SIP. Therefore, many installed videoconferencing systems use H.323.

Regardless of which is being used, the key pieces of functionality remain the same. The endpoint must be registered with a centralized control server—the H.323 Gatekeeper or SIP Agent. This control point knows about all the endpoints in the network and allows endpoints to establish connections with other devices. Some of the key functions of this device include:

- Authenticating that the device is allowed to place video calls. Most commonly, the authentication process can be tied into third party equipment such as Microsoft's Active Directory. Alternatively, a list of authorized users can be maintained manually.
- For direct (point-to-point) calls, translating the "identity" of the called party—typically either an E.164 identifier (phone number) or a name such as "Sydney Conference Room"— into an IP address. Once the two parties connect, they directly negotiate the characteristics of the videoconference, such as the resolution and frame rate.
- For bridged (point-to-multipoint) calls, specify the address of the MCU, which will be used. In this case, the MCU can transcode (convert) the videoconferencing signal, allowing disparate systems to communicate.
- Limit the bandwidth and number of simultaneous sessions. An administrator typically specifies the bandwidth which may be used by videoconferences and/or number of videoconferencing sessions. It is worth noting that this does not take into consideration the multiple network hops, which support differing amounts of bandwidth, does not know about network outages and has no awareness of how much bandwidth is being used by other applications.

Appendix B: VPN Overview

This section introduces the fundamentals of VPN technologies, concepts and terminology.

VPNs are not a new concept with ATM, Frame Relay and X.25 being examples of VPN technology. VPNs allow a single network to carry traffic for multiple communities of interest. In this context, virtual means that multiple organizations share a common network infrastructure. Private means that network information is secure and cannot be viewed by others sharing that same network.

Managed network services offerings are often built on VPNs. In the managed services model, a single (shared) service provider network supports multiple enterprise customers.

There are two broad VPN use cases:

- **Tunneling:** VPN technologies partition a single physical network into multiple logical networks, ensuring the privacy of each customer's information. This is most commonly done today using MPLS-based VPNs.
- **Encryption:** VPN mechanisms can also protect user information by encrypting traffic, typically using IPSec or SSL technology.

In addition, the two capabilities can be used together, encrypting traffic before putting it onto the shared network[1].

Tunneling: MPLS-Based VPNs

An MPLS-based virtual private network (VPN) consists of two domains: the provider's network where Provider Edge (PE) and Provider (P) routers both reside and the customer's network where Customer Edge (CE) devices reside. The customer's network is commonly spread throughout multiple physical sites.

VPNs include three types of network devices. Figure B.1 shows a simplified view.

- **Provider Edge (PE) routers:** Routers in the provider's network that connect to customer edge devices located at customer sites. PE routers support VPN and label functionality. Within a single VPN, PE routers are interconnected through an MPLS label-switched path (LSP).
- **Provider (P) routers:** Routers within the core of the provider's network that are not connected to any customer sites but are used to aggregate and transport the LSP tunnel between pairs of PE routers.
- **Customer Edge (CE) devices:** Routers or switches located at the customer site that connect to the provider's network. CE devices are typically IP routers, but also can be an ATM, Frame Relay or Ethernet switch.

[1] By definition, VPN technologies allow multiple organizations to share a common packet network. Many "private networks" are built on top of shared infrastructures using time division multiplexing (TDM) technology. For example, a SONET/SDH network can carry traffic for multiple customers in different timeslots. Even though the customers share nodes and links, this implementation is considered a private (not virtual private) network because the timeslots are actually reserved to the private network and not shared. If there is no useful data transmitted at any time, the time-slot bandwidth is wasted.

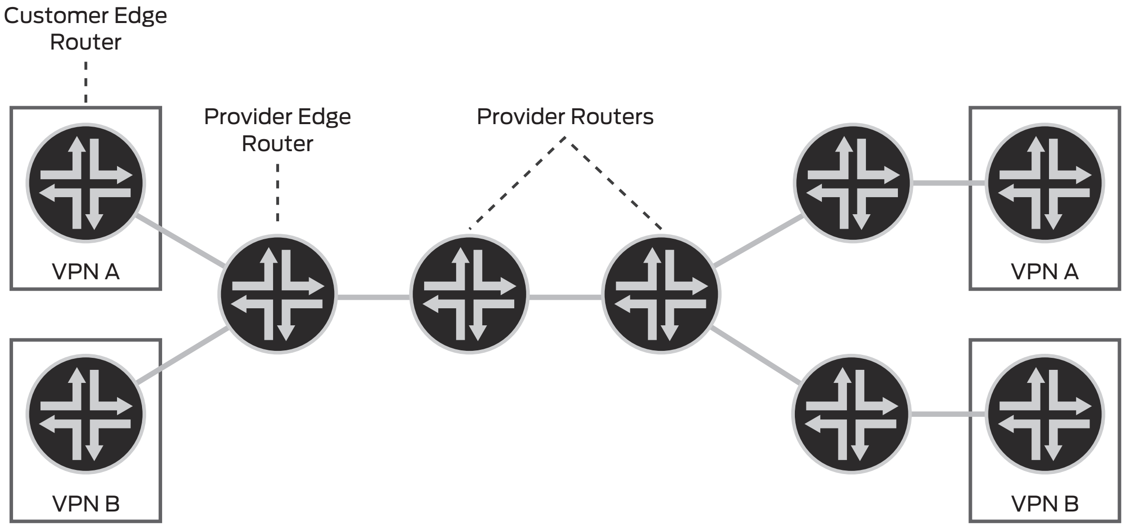

Figure B.1 Routers positioned in a VPN

Types of MPLS-based VPNs

MPLS-based VPNs fall into two general categories. L2VPNs provide a secure connection between two points. L3VPNs and VPLS—provide LAN-like "any to any" connectivity across the WAN. These latter two are of the most interest in service provider networks.

L3VPNs

In an L3VPN, forwarding is based on IP addresses. The most common L3VPN implementation is based on RFC4364. BGP/MPLS IP Virtual Private Networks (VPNs), which updates the older RFC 2547. This document defines a mechanism by which service providers can use their IP backbones to provide L3VPN services to their customers. The sites that make up a L3VPN are connected over a provider's existing backbone. A key concept within L3VPNs is the VPN Routing and forwarding (VRF) tables, which map traffic to the appropriate VPN.

VPLS

Virtual Private LAN Switch (VPLS) is a type of L2VPN that was developed for transport of Ethernet traffic between customer sites transparently. VPLS forwards traffic based on the Ethernet MAC address rather than the IP address. This means that the CE can be controlled by the enterprise and configured using their IP addressing scheme, while the PE is controlled by the service provider and can use their own IP numbering. There are two different signaling methods used for creating VPLS —BGP[2] and LDP[3]. Both techniques are supported by Juniper.

[2] See RFC4761,Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling

[3] See RFC4762,Virtual Private LAN Service (VPLS) Using Label Distribution (LDP) Signaling

Comparing L3VPNs and VPLS

At first glance, L3VPNs and VPLS are quite similar because both provide any-to-any connectivity between multiple customer locations. They use the same network topology and function similarly. A packet originating within a customer's network is sent first to a CE device. It is then sent to a PE router within the service provider's network. The packet traverses the service provider's network over an MPLS LSP. It arrives at the destination PE, which is turn hands off the packet to the destination CE. However, there are several key differences. This first and most important difference is that VPLS makes forwarding decisions based on Ethernet MAC address (Layer 2 of the OSI model) while L3VPNs make these decisions based on IP address (Layer 3).

L3VPNs are more scalable since forwarding is based on IP subnets, while VPLS must track Ethernet MAC addresses. As a result, L3VPNs are used in the core of Service Provider networks. As a service offering, selecting between these models depends on the level of service provider involvement in the customer's network operations. If the customer's goal is to use the provider network only for data transport, a Layer 2 model is better suited since the IP addressing and CE maintenance remains the customer's responsibility. This is more common for large enterprises. The Layer 3 model is appropriate if there is a requirement for the network operator to configure and maintain IP addressing for the customer, which is more typical when the customer is a medium-sized business.

Juniper Networks MX Series, M Series and T Series provide a complete suite of MPLS VPN capabilities[4,5].

MPLS VPN Components

As illustrated in Figure B.2, any VPN consists of two major components—the LSP transport tunnel and the VPN. Both originate and terminate at the PE. Note that the PE routers provide all VPN functionality; the P and CE routers have no special configuration requirements for VPNs.

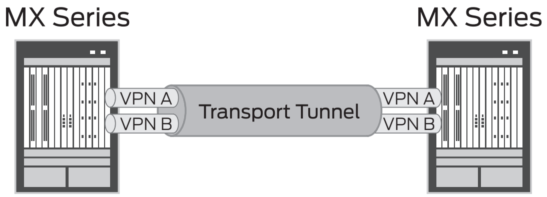

Figure B.2 VPN components

[4] See Comprehensive MPLS VPN Solutions, **juniper.net/us/en/local/pdf/solutionbriefs/3510324-en.pdf**

[5] See *Appendix C: Juniper Products on page 171* for an overview of Juniper Networks products

Transport Tunnels

A transport tunnel carries traffic between PEs. A transport tunnel between each pair of PEs can carry traffic for multiple VPNs, or there may be a separate transport tunnel for each VPN. These transport tunnels are MPLS label switched paths (LSPs). Each LSP transport tunnel is unidirectional. There are two types of LSPs—point to point and point to multipoint. Unicast VPNs use separate point-to-point transport tunnels to carry traffic in each direction, allowing for bi-directional communication. In contrast, point-to-multipoint LSPs do not have a companion carrying traffic upstream. Each transport tunnel is identified by an MPLS label carried by each packet.

VPN Tunnels

The second component is the VPN tunnel, which specifies the group (enterprise customer, for example) to which this traffic belongs. VPN traffic is carried between PEs using the transport tunnels. VPN traffic is identified by adding a second MPLS label that uniquely identifies the customer. Figure B.3 shows a high-level view of a typical VPN.

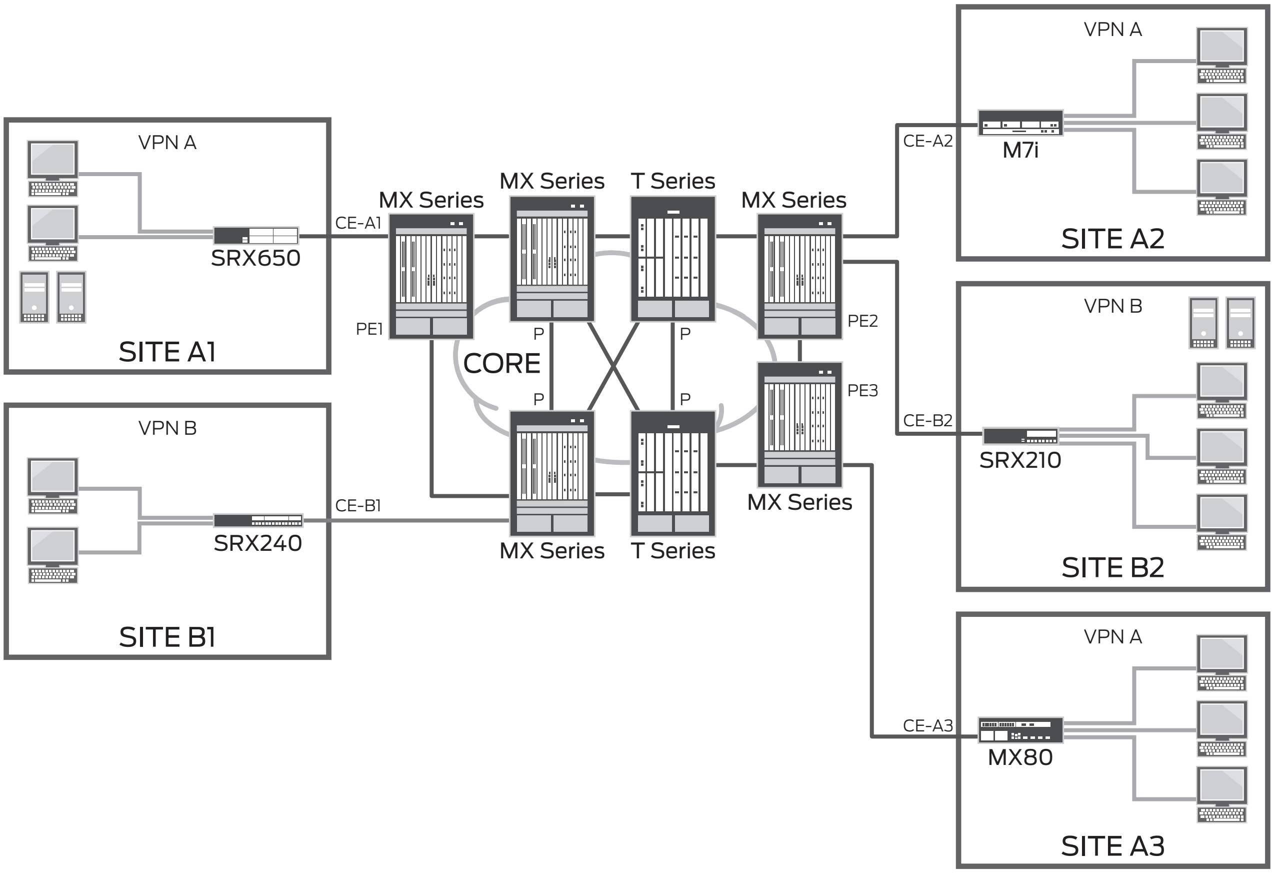

Figure B.3 VPN overview

Encrypting the Information: IPSec and SSL

It is not always possible to separate customers into separate logical networks. For example, residential customers and small businesses may share a broadband access network and subsequent connection to the Internet. Business data (as well as consumer financial transactions) must be protected from prying eyes. This is done by encrypting the information. There are two common mechanisms:

- **IP Security (IPSec)** is commonly used to protect traffic from an entire site. A router or security appliance at the customer site performs two functions[6] . First, it encapsulates the customer's IP packet into an IP packet which conforms to the carrier's numbering scheme. Hiding the business's IP address allows retaining its existing IP numbering scheme to support remote offices, while allowing the packet to be carried across the shared network. Second, the customer's entire original IP packet is encrypted. Small offices, including retail offices, typically use IPSec to protect Internet-bound traffic.

  Juniper Networks SRX Series, as well as many other Juniper products, supports IPSec.

- **Secure Socket Layer (SSL)** is typically used to protect traffic from a single device, typically a PC. Unlike IPSec, the assumption is that the customer is using a carrier-assigned IP address so it does not need to be encapsulated. Like IPSec, the packet (excluding the IP header) is encrypted. Since SSL is built into most web browsers, it is typically transparent to users. URLs that start with "https:" instead of "http" are using SSL. In addition, many corporations also provide SSL clients to connect to the network.

  SSL support is integrated into most web browsers. To encrypt all traffic on a single PC, Juniper Networks SA Series supports scalable, centralized termination of SSL VPNs[7]. Juniper Networks Odyssey Access Client with Network Connect and Junos Pulse implement SSL on the PC.

Combining MPLS Tunneling and IPSec/SSL Encryption

Traffic protected by IPSec or SSL can be sent across any IP-based network, whether or not MPLS is deployed, providing an additional layer of security.

Figure B.4 illustrates four sites: one site uses only MPLS tunneling, another uses only IPSec, a third site uses both MPLS tunneling and IPSec while the four site uses SSL to protect corporate traffic.

[6] This is the most common method of deploying IPSec, although other options exist.

[7] For additional information about Juniper SSL VPN appliances, see **juniper.net/us/en/products-services/software/security/sa-series/virtual-appliance/**. As a starting point, see *SA Series SSL VPN Appliances*, **juniper.net/us/en/local/pdf/brochures/1500023-en.pdf** and *SA Series SSL VPN Virtual Appliances*, **juniper.net/us/en/local/pdf/datasheets/1000320-en.pdf**

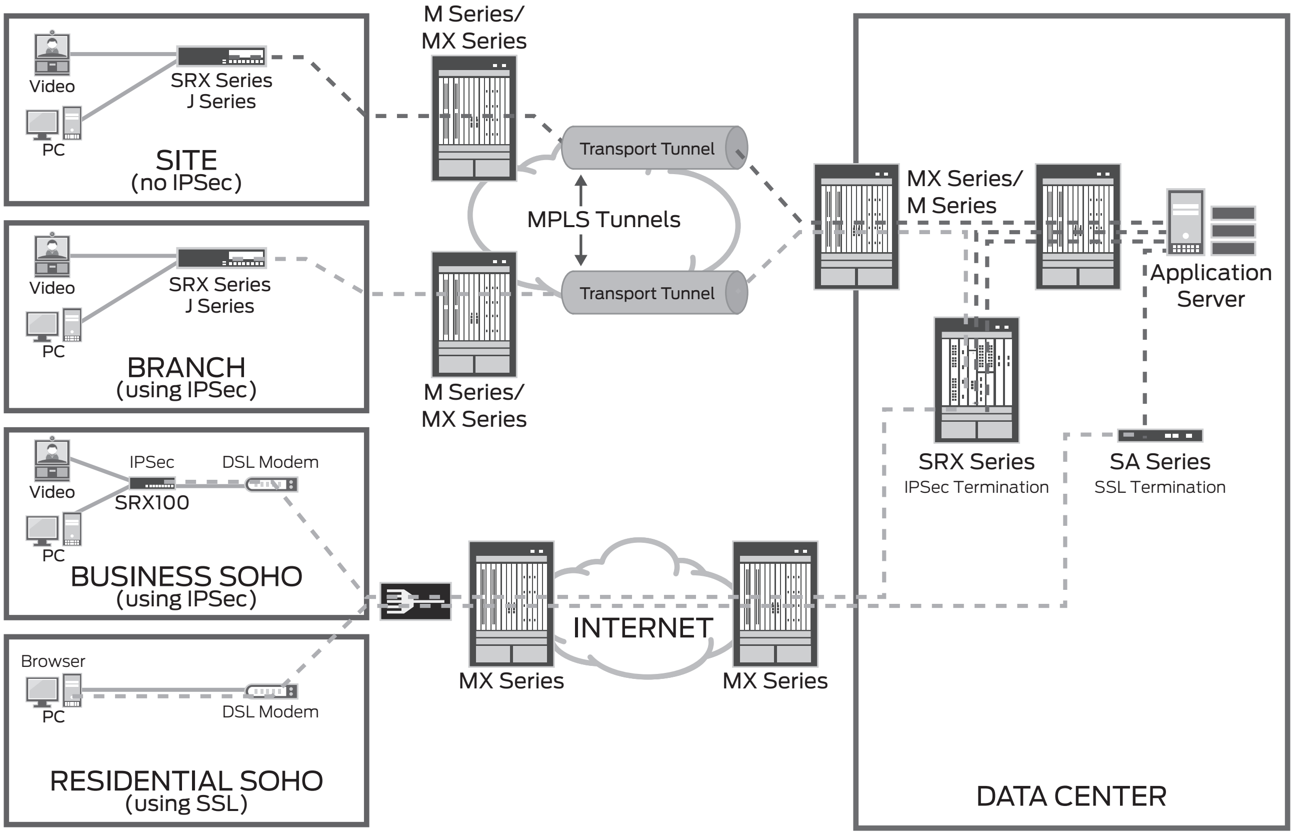

Figure B.4 Four sites showing IPSec and SSL connectivity

Appendix C: Juniper Products

This section describes the Juniper products that were used during validation.

The products discussed here are a subset of the complete line and focus on key devices tested with and/or expected to be most widely deployed. The products include:

- MX Series 3D Universal Edge Routers
- M Series Multiservice Edge Routers
- SRX Series Services Gateways
- EX Series Ethernet Switches
- Junos OS
- Session and Resource Control (SRC) software

For additional details concerning these products and any other Juniper products, visit **www.juniper.net .**

MX Series 3D Universal Edge Routers

The MX Series 3D Universal Edge Routers are a series of high-performance Ethernet routers with powerful switching features designed for enterprises and service provider networks. The MX Series provides unmatched flexibility and reliability to support advanced services and applications. It addresses a wide range of deployments, architectures, port densities and interfaces. High-performance enterprise networks typically deploy MX Series routers in high-density Ethernet LAN and data center aggregation, the data center core, and metro Ethernet aggregation and core layers.

Major features are:

- High-density routers optimized for Ethernet that function as Layer 2 switches or Layer 3 routers, or both, maximizing service flexibility and increasing investment protection
- Deliver the most advanced routing features, including network virtualization with MPLS, low latency multicast, and QoS, without compromising performance
- Provide the highest level of redundancy and resiliency to ensure that critical services and customers stay connected, allowing the enterprise to ensure customer satisfaction and lower costs
- Run Junos OS, a single network operating system that allows administrators to quickly and cost-effectively keep up with continuously changing business demands.

The MX Series provides carrier grade reliability, density, performance, capacity, and scale for enterprise networks with mission critical applications. Its high availability features ensure that the network is always up and running, including nonstop routing (NSR), fast reroute, and unified in-service software upgrade (ISSU). The MX Series also delivers significant operational efficiencies enabled by Junos OS, a collapsed architecture requiring less power, cooling, and space consumption, and open APIs for easily customized applications and services.

The MX Series support up to 2.6 Tbps. An overview of the MX Series models can be downloaded from **juniper.net/us/en/local/pdf/datasheets/1000208-en.pdf** . Additional information on the MX Series is available at **juniper.net/us/en/products-services/routing/mx-series/**.

SRX Series Services Gateways

The SRX Series Services Gateways are next-generation appliances based on a revolutionary new architecture that provides market leading scalability and service integration. Based on Juniper's Dynamic Services Architecture, the SRX Series offers service expandability as well as flexible processing and I/O scalability.

Major features include:

- Tightly integrated networking and security capabilities include firewall, intrusion prevention system (IPS), distributed denial of service (DdoS) and denial of service (DoS), routing, quality of service (QoS), Network Address Translation (NAT), and other capabilities
- Dynamic Services Architecture allows the Juniper Networks SRX Series Services Gateways to quickly enable new services and capabilities
- Carrier class reliability is based on features ranging from redundant hardware and components to Juniper Networks Junos operating system

The SRX Series include models designed for use in branch offices (SRX 100 through SRX650) as well as models designed for campuses (SRX 3000 line) and data centers (SRX 5000 line). Additional information on the SRX Series is available at **juniper.net/us/en/products-services/security/srx-series/**.

EX Series Ethernet Switches

The EX Series Ethernet Switches represent a new class of infrastructure switch for high-performance businesses. Designed to address the access, aggregation, and core layers of branch office, campus, and data center applications, EX Series switches provide the infrastructure foundation for the fast, secure, and reliable delivery of applications that support strategic business processes. EX Series switches advance the economics of networking by delivering cost saving capabilities that allow businesses to reduce capital and operational expenses. The resulting savings can fund investments in innovative initiatives that allow businesses to improve productivity, streamline operations and gain a competitive advantage.

EX Series switches are designed to address increasing demands for high availability (HA) and unified communications within high-performance enterprise networks. Working together, the EX Series switches create a standards-based network foundation that is well aligned and flexible enough to deliver all applications—everything from file services to IP telephony, messaging, presence, videoconferencing, and Web services.

EX Series switches offer sufficient scalability and performance to meet emerging requirements as well. As part of the Juniper product portfolio, the EX Series switches represent a key strategic addition that contributes to one of the industry's most complete suite of network infrastructure product offerings.

These switches all run the same Juniper Networks Junos OS, offering consistent implementation and management with time tested Juniper routers and security solutions. Junos OS adheres to a strict development process that utilizes a common source code, follows a single release train, and builds upon a modular architecture, dramatically reducing maintenance and management overhead for Junos OS-based solutions. As a result, the Junos OS-based EX Series switches ensure consistent and predictable behavior and shared feature implementation across the entire network infrastructure.

An overview of the EX Series models can be downloaded from **juniper.net/us/en/local/pdf/brochures/1500057-en.pdf** . Additional information on the EX Series is available at **juniper.net/us/en/products-services/switching/ex-series/**.

Junos Operating System

A core Juniper Networks strategy is innovation through software to integrate new value into the network and reduce complexity. The Juniper Networks Junos Platform is our open software platform that delivers on this strategy. At the foundation is the Junos operating system. The Junos OS provides a common language across our routing, switching and security devices to simplify new deployments and reduce network operation costs by up to 41%[1].

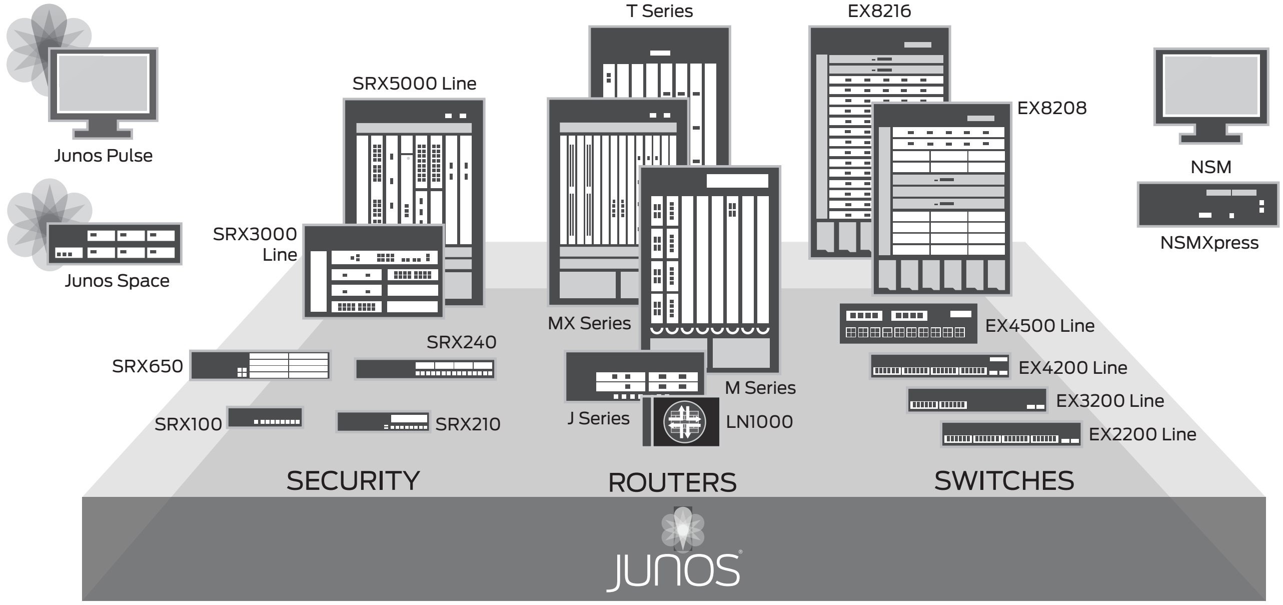

Figure C.1 Junos platform

[1] Source: A commissioned study conducted by Forrester Consulting on behalf of Juniper Networks, The Total Economic Impact™ of Juniper Networks Junos Network Operating System.

Unlike other network operating systems, Junos OS is enhanced through one software release train, and developed based on one modular architecture—the "power of one."

- One operating system across platforms reduces the time and effort to plan, deploy and operate network infrastructure.
- One release train runs a network on the same software version and provides new functionality in a steady, time-tested cadence of stable releases.
- One modular software architecture provides highly available, secure and scalable software that keeps up with changing needs.

Accelerated by the "power of one" differences, Junos OS has rapidly evolved to support a diverse set of application and service needs. Juniper platforms simultaneously scale integrated security and networking capabilities without compromising performance or reliability. Introducing Junos OS-based devices can reduce the level of complexity that would otherwise be present in a network. This reduces operational challenges and improves operational productivity.

Junos OS:

- Minimizes the impact of human factors with fail-safe mechanisms
- Speeds response and resolution of unplanned events
- Reduces configuration effort and issues
- Simplifies day-to-day operations
- Interoperates and integrates with existing systems
- Simplifies new upgrades and redeploy systems

For additional information on Junos OS, see **juniper.net/us/en/local/pdf/brochures/1500059-en.pdf** and **juniper.net/us/en/products-services/nos/junos/**.

Session and Resource Control (SRC)

The SRC connects the service layer with the network layer of service provider networks by providing a feedback loop between applications, users, and the network. Its open interfaces allow it to integrate with any network and any service offering, regardless of where the demand is generated. The SRC allows service providers to generate additional revenue on their existing network infrastructure by adding dynamically activated services.

Unlike competitive solutions that utilize AAA for policy management or solutions that deploy static policy enforcement solutions, Juniper's SRC Portfolio delivers granular dynamic policy enforcement on a per service level. This enables you to deliver revenue-generating services on top of existing sessions. The SRC readily interfaces with your existing subscriber management databases. This enables you to map available network resources to subscriber and service profiles. The result is differentiated services that are based on dynamic allocation of network resources with the ability to provide service specific accounting.

The SRC consists of the following components:

- C Series hardware appliance
- SRC Policy Engine – This is the required base SRC software
- SRC SOAP Gateway– This software provides the open, published web services interface, which allows the SRC to communicate with external equipment such as Polycom DMA.

Additional information about SRC software and hardware is available in *Chapter 10, Implementing Assured Forwarding on page 129* and in *SRC Series Session and Resource Control Modules Datasheet* at **juniper.net/us/en/local/pdf/datasheets/1000195-en.pdf**.

Appendix D: Polycom Products

This appendix introduces the Polycom® products used during validation. The products discussed here are a subset of Polycom's complete portfolio, and focus on key pieces of equipment tested with and or expected to be most widely deployed. The products discussed in this section include:

- Videoconferencing infrastructure equipment—Unified Communication Intelligent Core Solutions
  - Polycom Distributed Media Application™ (DMA™)
  - Polycom Converged Media Application™ (CMA®)
  - Polycom RMX® conference platform (MCUs)
- Videoconferencing endpoints
  - Immersive systems (Polycom RPX™, TPX®, OTX™, ATX™)
  - Room systems (Polycom HDX® video endpoints)
  - Personal systems (Polycom VVX™ 1500 business media phone, HDX 4000 personal telepresence)

Polycom provides a complete line of voice and video (UC) solutions, including a wide range of videoconferencing endpoints and related equipment such as recording servers. For additional information on Polycom products, see **polycom.com** . For technical support and configuration information, see **http://support.polycom.com/PolycomService/home/home.htm**

Polycom® UC Intelligent Core technology provides an unmatched, fully standards-based collaboration infrastructure designed to scale, perform, and optimize the way you collaborate, within your organization and with partners and customers. The key solutions within the Polycom UC Intelligent Core solution include the Polycom® RMX® platform for flexible multiparty conferencing; the Polycom DMA™ solution for call control, bridge virtualization, and load balancing; and the Polycom CMA® conferencing solution for management and reporting of video endpoints.

Polycom Distributed Media Application™ (DMA™)

The Polycom Distributed Media Application (DMA) 7000 is a network-based application that provides video endpoint / device registration, call processing (including call admission control), and the management and distribution of point-to-point and multipoint video and audio calls across globally distributed conferencing platforms and media servers. Notable DMA features include support for dual hot standby application servers, call processing of simultaneous video and audio calls, support for virtualization of up to 1,200 concurrent video / audio calls per node, integration with LDAP for centralized account management, and integrated SIP and H.323 support. The DMA can manage multipoint calls and distribute them across available conference platforms (bridges). As video network infrastructure grows, so does the need for efficient utilization of available resources. For example, conference platforms in one location might become unavailable while at the same time there is a surplus of conference resources in neighboring locations. The DMA 7000's efficient and scalable media processing capabilities help leverage investments in video infrastructure.

The DMA 7000 can be installed in a high-availability configuration in hot-standby mode. It can scale up to 1200 mixed audio and video ports and load balance them to multiple Polycom RMX or select third-party conference platforms. The DMA 7000 also serves as centralized administration and management platform and provides management of user accounts through LDAP. It also provides call logging capabilities and overall system monitoring functionalities.

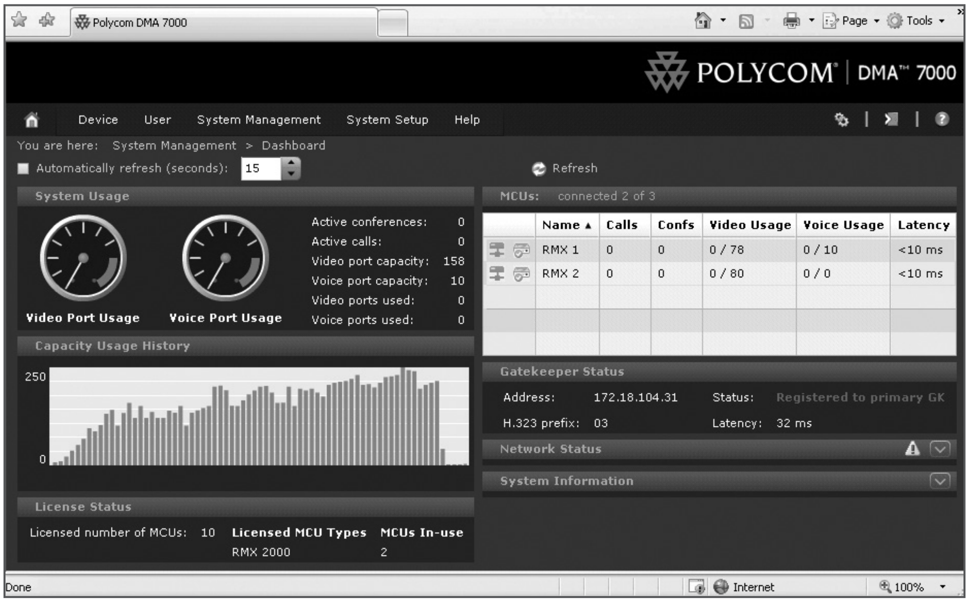

You are here: System Management > Dashboard

□ Automatically refresh (seconds): 15

Active conferences: 0
Active calls: 0
Video port capacity: 158
Voice port capacity: 10
Video ports used: 0
Voice ports used: 0

MCUs: connected 2 of 3

Name ▲	Calls	Confs	Video Usage	Voice Usage	Latency
RMX 1	0	0	0 / 78	0 / 10	<10 ms
RMX 2	0	0	0 / 80	0 / 0	<10 ms

Address: 172.18.104.31 Status: Registered to primary GK
H.323 prefix: 03 Latency: 32 ms

Licensed number of MCUs: 10

Licensed MCU Types	MCUs In-use
RMX 2000	2

Figure D.1 Polycom DMA 7000 dashboard

For information on configuring the Polycom DMA, see **polycom.com/support/network/management_scheduling/dma_7000.html** .

Polycom Converged Management Application™ (CMA®)

To centrally manage and deploy the complete suite of Polycom videoconferencing products as well as third-party systems, an administrator can deploy a Polycom Converged Management Application (CMA) 5000 or 4000. The Polycom CMA provides high-scale, powerful management and a software-based video client as simple to use as instant messaging.

The CMA 5000 is a redundant platform designed to manage larger video networks, while the CMA 4000 is a single platform with more limited scale. This system streamlines the efficiency of day-to-day tasks linked with deployment and management of video endpoints, ensures consistency of device configurations as well as software and firmware versions. It also helps to track usage of all videoconferencing resources in the organization and generates a broad set of call detail reporting which could be customized and organized per site or system. Detailed ROI calculations can also be performed as part of the reporting routine. The CMA 5000 simplifies operations of the videoconferencing infrastructure and helps troubleshoot using its world-view graphical maps of video infrastructure and network topology.

The Polycom CMA also integrates corporate directory services and manages ‘user presence’ indication for unified communication systems. The downloadable CMA desktop application offers an easy-to-use video client for enterprise PC users.

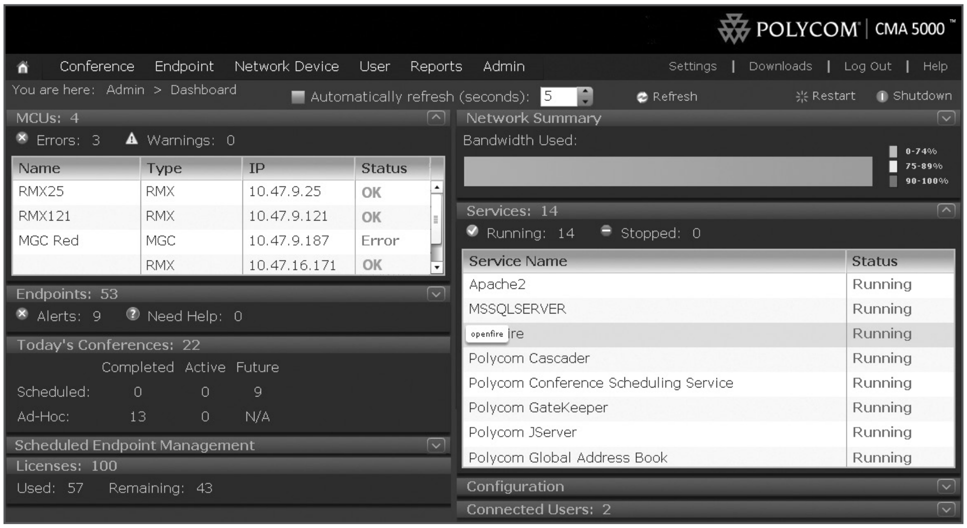
POLYCOM | CMA 5000™
Conference Endpoint Network Device User Reports Admin
Settings | Downloads | Log Out | Help
Refresh
Restart
Shutdown

You are here: Admin > Dashboard

☐ Automatically refresh (seconds): 5

MCUs: 4

Errors: 3 Warnings: 0

Name	Type	IP	Status
RMX25	RMX	10.47.9.25	OK
RMX121	RMX	10.47.9.121	OK
MGC Red	MGC	10.47.9.187	Error
	RMX	10.47.16.171	OK

Endpoints: 53

Alerts: 9 Need Help: 0

Today's Conferences: 22

	Completed	Active	Future
Scheduled:	0	0	9
Ad-Hoc:	13	0	N/A

Scheduled Endpoint Management

Licenses: 100

Used: 57 Remaining: 43

Network Summary

Bandwidth Used:

0-74% 75-89% 90-100%

Services: 14

Running: 14 Stopped: 0

Service Name	Status
Apache2	Running
MSSQLSERVER	Running
openfire re	Running
Polycom Cascader	Running
Polycom Conference Scheduling Service	Running
Polycom GateKeeper	Running
Polycom JServer	Running
Polycom Global Address Book	Running

Configuration

Connected Users: 2

Figure D.2 Polycom CMA server dashboard

Figure D.3 shows a sample report from the videoconferencing network of one large international organization highlighting daily and weekly statistics.

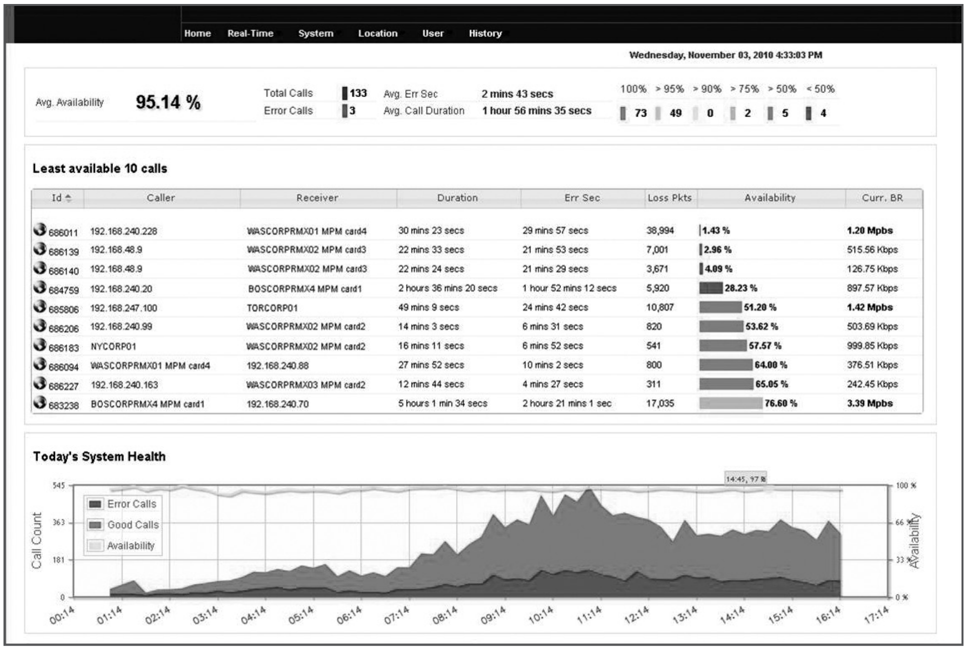

Wednesday, November 03, 2010 4:33:03 PM

Avg. Availability 95.14 %

Total Calls 133 Avg. Err Sec 2 mins 43 secs

Error Calls 3 Avg. Call Duration 1 hour 56 mins 35 secs

100%	> 95%	> 90%	> 75%	> 50%	< 50%
73	49	0	2	5	4

Least available 10 calls

Id	Caller	Receiver	Duration	Err Sec	Loss Pkts	Availability	Curr. BR
686011	192.168.240.228	WASCORPRMX01 MPM card4	30 mins 23 secs	29 mins 57 secs	38,994	1.43 %	1.20 Mpbs
686139	192.168.48.9	WASCORPRMX02 MPM card3	22 mins 33 secs	21 mins 53 secs	7,001	2.96 %	515.56 Kbps
686140	192.168.48.9	WASCORPRMX02 MPM card3	22 mins 24 secs	21 mins 29 secs	3,671	4.09 %	126.75 Kbps
684759	192.168.240.20	BOSCORPRMX4 MPM card1	2 hours 36 mins 20 secs	1 hour 52 mins 12 secs	5,920	28.23 %	897.57 Kbps
685806	192.168.247.100	TORCORP01	49 mins 9 secs	24 mins 42 secs	10,807	51.20 %	1.42 Mpbs
686206	192.168.240.99	WASCORPRMX02 MPM card2	14 mins 3 secs	6 mins 31 secs	820	53.62 %	503.69 Kbps
686183	NYCORP01	WASCORPRMX02 MPM card2	16 mins 11 secs	6 mins 52 secs	541	57.57 %	999.85 Kbps
686094	WASCORPRMX01 MPM card4	192.168.240.88	27 mins 52 secs	10 mins 2 secs	800	64.00 %	376.51 Kbps
686227	192.168.240.163	WASCORPRMX03 MPM card2	12 mins 44 secs	4 mins 27 secs	311	65.05 %	242.45 Kbps
683238	BOSCORPRMX4 MPM card1	192.168.240.70	5 hours 1 min 34 secs	2 hours 21 mins 1 sec	17,035	76.60 %	3.39 Mpbs

Today's System Health

Figure D.3 CMA daily health report

This GUI pane shows the analysis of Call Detail Records (CDR) collected by CMA for one day. The top pane shows the health status dashboard for the entire videoconferencing service of the enterprise. The average availability of the service, which is calculated automatically displays 95.14%.The top right pane shows the color bar with real time statistics pertaining to calls that met or breached SLA, which in this case was set at 95% of service availability and is represented by the green color status indicator.

The center pane “Least available 10 calls” shows 10 calls with lowest quality expressed as service availability. The availability was calculated as percentage of the seconds with errors over the total duration of the call and in absolute number of lost packets for each call respectively.

The bottom pane “Today’s System Health” which shows the videoconferencing service availability and share of Good and Error calls represented with green and red colors.

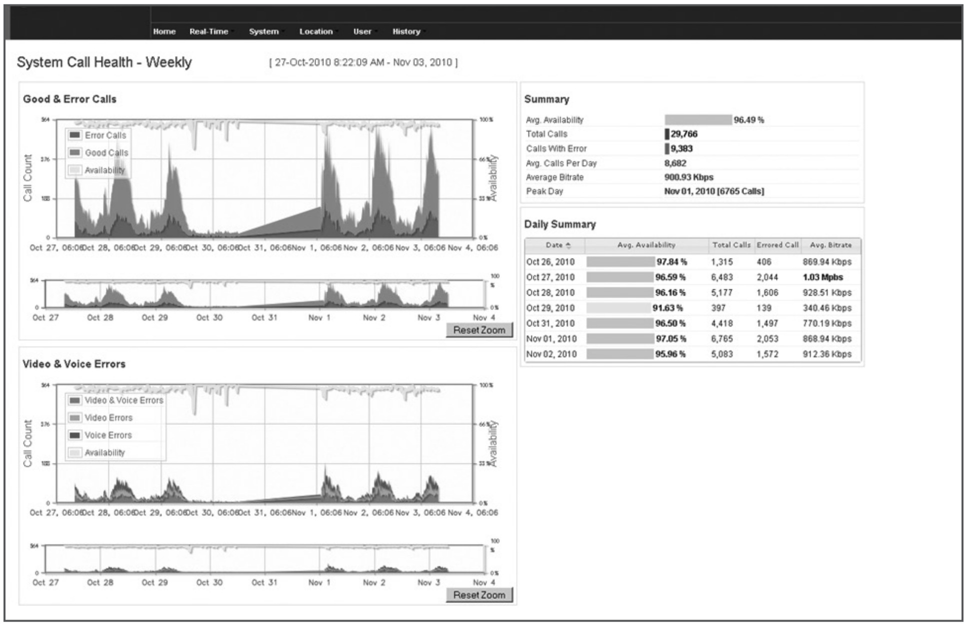

System Call Health - Weekly [27-Oct-2010 8:22:09 AM - Nov 03, 2010]

Summary

Avg. Availability	96.49 %
Total Calls	29,766
Calls With Error	9,383
Avg. Calls Per Day	8,682
Average Bitrate	900.93 Kbps
Peak Day	Nov 01, 2010 [6765 Calls]

Daily Summary

Date	Avg. Availability	Total Calls	Errored Call	Avg. Bitrate
Oct 26, 2010	97.84 %	1,315	406	869.94 Kbps
Oct 27, 2010	96.59 %	6,483	2,044	1.03 Mpbs
Oct 28, 2010	96.16 %	5,177	1,606	928.51 Kbps
Oct 29, 2010	91.63 %	397	139	340.46 Kbps
Oct 31, 2010	96.50 %	4,418	1,497	770.19 Kbps
Nov 01, 2010	97.05 %	6,765	2,053	868.94 Kbps
Nov 02, 2010	95.96 %	5,083	1,572	912.36 Kbps

Figure D.4 CMA weekly health report

Figure D.4 shows weekly call statistics. In addition to the service quality dashboards and summary statistics, the bottom left pane (Video & Voice Errors) shows the breakdown of errors overlaid with a line representing availability of the service.

These statistics represent a typical videoconferencing service that is experiencing high usage. The statistics were collected before the introduction of joint High Definition videoconferencing solution from Polycom and Juniper. We believe that this is a typical scenario that compels the IT organization to assure that different tiers of users receive differentiated and assured quality of service at any time.

The Polycom CMA server supports the following standards for call signaling, for management of videoconferencing endpoints, integration with external directories and databases: H.323, LDAP/H.350, XMPP, HTTPS/XML Provisioning and TLS – Security. It also supports Microsoft SQL Server 2005 and Microsoft Active Directory 2003.

Polycom RMX® Conference Platforms

As the number of video and voice endpoints on the enterprise network grows, so does the need for group collaboration, also known as multipoint calls. The RMX series of conference platforms is a redundant and scalable voice and videoconferencing solution that maximizes productivity by extending unified collaboration to teams throughout the enterprise and beyond. Although some video endpoints provide limited multipoint capabilities, the larger enterprises need dedicated conference bridges. The mid-range Polycom RMX 2000 platform is an example of the open architecture media-conferencing bridge which is purpose built for real-time HD collaboration. It requires a high-speed IP network and supports both H.323 and SIP signaling protocols. With its modular design, the RMX 2000 platform can support PSTN, ISDN and IP originated calls. Polycom UltimateHDTM technology on the RMX 2000 supports HD video with up to 1080p at 30fps with surround audio and content sharing. Using Polycom Video Clarity™ technology, the RMX 2000 platform can upscale and convert video streams from legacy endpoints and include them in the HD videoconference. The Polycom RMX 4000 platform expands the features of the RMX 2000 with higher scale, redundant power supplies and hot swappable control and media processing modules.

For information on configuring the Polycom RMX 2000, see **polycom.com/support/network/collaboration_conferencing_platforms/rmx_2000.html** .

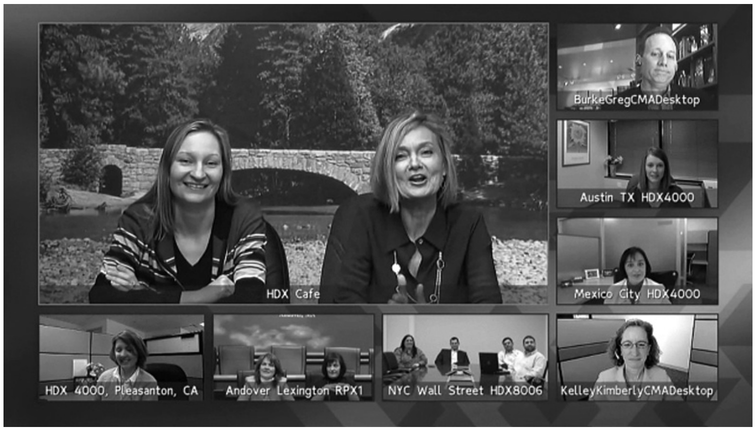

Figure D.5 Multiparty conferencing using Polycom RMX 2000

Videoconferencing Systems

Polycom provides a wide range of HD videoconferencing endpoints as well as traditional video and voice endpoints. Polycom HD videoconferencing solutions are divided in three product families: Room, Personal and Immersive systems.

Polycom Room Systems: HDX® 9000, 8000, 7000, 6000

Polycom's powerful high definition solutions for room environments expand real-time knowledge sharing and drive
faster, more informed decisions

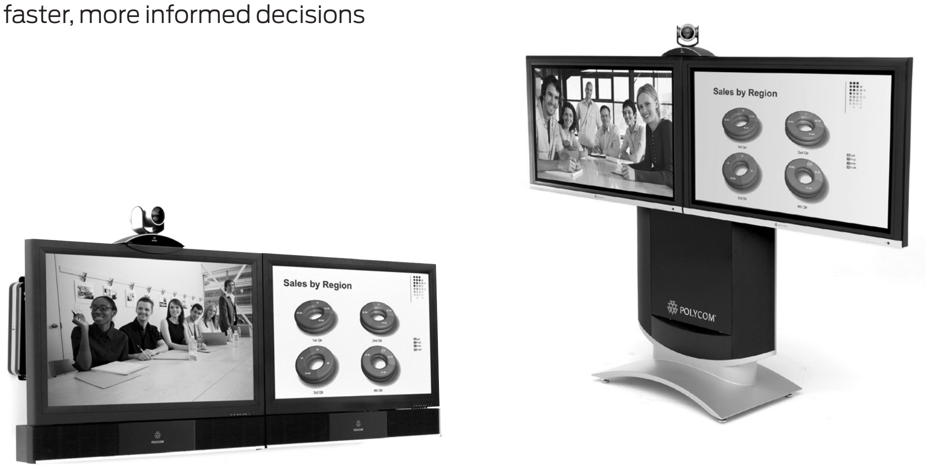

Figure D.6 Polycom Room HD videoconferencing systems

The Polycom HDX® 9000, 8000, 7000, and 6000 endpoints bring HD video collaboration to the groups and deliver exceptional performance for small to medium size environments. Support of Polycom UltimateHD™ and Polycom Siren™ 22 or Polycom StereoSurround™ audio technology allows smooth, crisp video and natural audio even during busy meetings. The HDX 8000 and 9000 systems feature video resolution at 1920x1080 (1080p) at 30 fps. The HDX 9000 is highly configurable and can connect multiple HD video sources simultaneously. Use of H.264 High Profile technology conserves bandwidth with no impact to the video quality.

For smaller HD videoconferencing deployments, the HDX 8000 system provides internal video-bridge (MCU) support for up to four participants with HD video and surround audio.

Polycom Personal Systems: HDX 4000, VVX™ 1500

Polycom personal video solutions seamlessly extend clear, high definition video to home offices, mobile users, branch sites, and beyond. The Polycom HDX 4000 series delivers the industry's premier personal HD videoconferencing experience, with high-definition video, audio, and content sharing, housed in a powerful, compact, and stylish design.

Polycom® VVX™ 1500 business media phones unify voice, video and applications capabilities into simple-to-use UC devices, providing a dynamic, real-time meeting experience in a compact format with features that make it easy to use in a home-office environment and operate over an IP network.

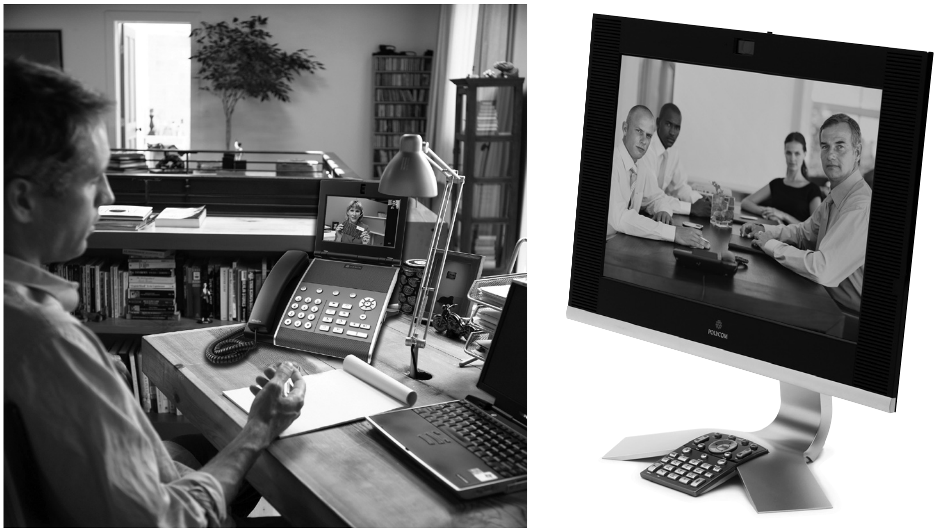

Figure D.7 Polycom VVX business media phone and HDX 4000

Polycom Immersive Telepresence Systems: RPX™, OTX™, ATX™

Polycom immersive solutions provide a natural, "across the table" experience where every meeting participant is shown in true-to-life dimensions.

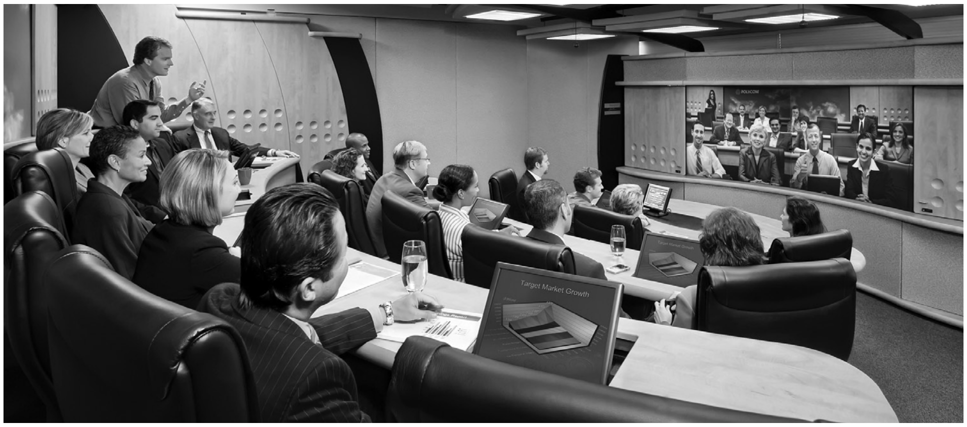

Figure D.8 Polycom RPX HD Series

The Polycom RealPresence Experience HD (RPX™ HD) system is the world's only fully immersive solution for medium to large groups, and delivers the ultimate meeting experience with a powerful combination of high definition video, audio, and content sharing featuring a cinematic, 16 foot video wall and an all-inclusive luxury environment. The system is fully integrated in the room or auditory environment and permits seating from 8 to 28 participants. The RPX system features a 48:9 aspect cinematic view screen and utilizes the Polycom UltimateHD™ and Polycom EyeConnect™ technologies.

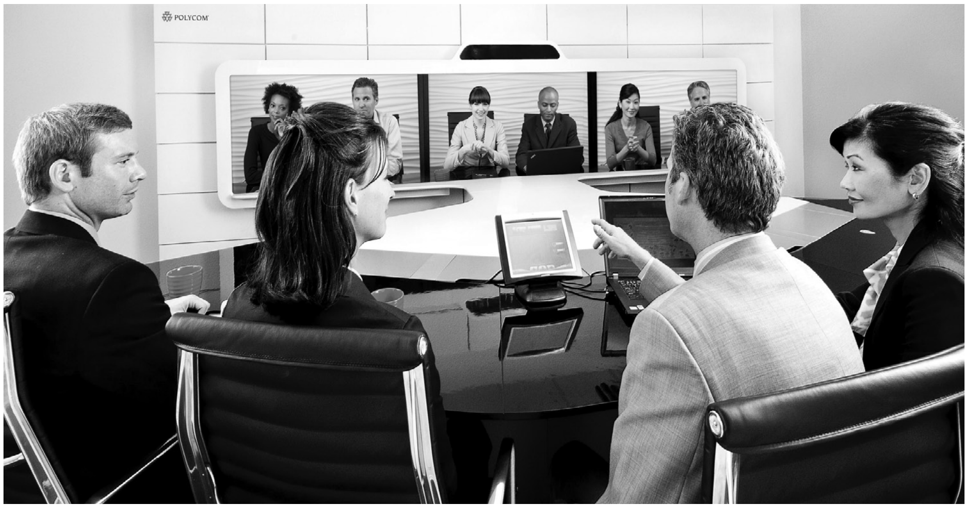

Figure D.9 Polycom OTX system

The Polycom Open Telepresence Experience™ (OTX™) High Definition 300 system combines flexible and ultra-modern room design for organizations seeking an extraordinary true-to-life meeting experience for smaller groups. It provides full HD 1080p video and content is displayed on motorized monitors that can be raised and lowered on demand, integrated hidden lighting and specially placed accessories.

The Polycom Architected Telepresence Experience™ (ATX™) system provides ultimate flexibility for deploying immersive telepresence across a wide range of room environments. The ATX system is a toolkit that includes all required components and enables the authorized Polycom partners to create any custom tailor designs for demanding customers.

Appendix E: References

Web Pages

- Juniper Networks (home page): **www.juniper.net**

Juniper Marketing Collateral: Videoconferencing

- Juniper Networks Solution for HD Videoconferencing (without assured forwarding), **juniper.net/us/en/local/pdf/solutionbriefs/3510345-en.pdf**
- Juniper-Polycom Telepresence and HD Videoconferencing Solution, **juniper.net/us/en/local/pdf/solutionbriefs/3510358-en.pdf**

Juniper Marketing Collateral: Junos

- How Operating Systems Create Network Efficiencies, **juniper.net/us/en/local/pdf/whitepapers/ lake_partners_network_efficiency.pdf**

Juniper Technical Collateral: WAN, QoS and Virtual Chassis

- Virtual Private LAN Services (web page): **juniper.net/techpubs/en_US/junos10.2/information-products/pathway-pages/solutions/virtual-private-lan-services/index.html#configuration**
- EX4200 Metro Ring with MX Series Head-End: Using Virtual Chassis Technology in the WAN, **juniper.net/us/en/local/pdf/implementation-guides/8010045-en.pdf**
- Real-Time Performance Monitoring (RPM) on Juniper Networks Devices, **juniper.net/us/en/local/pdf/app-notes/3500145-en.pdf**
- Remote Users: **juniper.net/us/en/solutions/enterprise/extended-remote/**
- Virtual Chassis Technology Best Practices, **juniper.net/us/en/local/pdf/implementation-guides/8010018-en.pdf**
- Introduction to Carrier Ethernet VPNs: Understanding the Alternatives **juniper.net/us/en/local/pdf/whitepapers/2000306-en.pdf**

Juniper Technical Collateral: Network Locations

- Cloud-Ready Data Center, **juniper.net/us/en/solutions/enterprise/data-center/**
  - Data Center Network connectivity with IBM Servers, **juniper.net/us/en/training/jnbooks/dchandbook.html**
  - Cloud-Ready Data Center Reference Architecture, **juniper.net/us/en/local/pdf/reference-architectures/8030001-en.pdf**
  - Data Center LAN Connectivity Design Guide, **juniper.net/us/en/local/pdf/design-guides/8020010-en.pdf**
  - Implementing VPLS for Data Center Interconnectivity, **juniper.net/us/en/local/pdf/implementation-guides/8010050-en.pdf**
  - Implementing L2 at the Data Center Access Layer on Juniper Networks Infrastructure, **juniper.net/us/en/local/pdf/implementation-guides/8010014-en.pdf**
  - Implementing L3 at the Data Center Access Layer on Juniper Networks Infrastructure, **juniper.net/us/en/local/pdf/implementation-guides/8010022-en.pdf**
- Campus: **juniper.net/us/en/solutions/enterprise/campus/**
  - Campus Networks Reference Architecture, **juniper.net/us/en/local/pdf/reference-architectures/8030007-en.pdf**
  - Campus LAN Design Guide, **juniper.net/us/en/local/pdf/design-guides/8020001-en.pdf**
- Branch Office: **juniper.net/us/en/solutions/enterprise/branch/**
  - Branch High Availability in the Distributed Enterprise, **juniper.net/us/en/local/pdf/implementation-guides/8010017-en.pdf**
  - Branch Office Connectivity Guide, **juniper.net/us/en/local/pdf/app-notes/3500143-en.pdf**

Juniper Technical Collateral: Security

- Junos OS Security Configuration Guide, **juniper.net/techpubs/software/junos-security/junos-security10.3/junos-security-swconfig-security/junos-security-swconfig-security.pdf**
- Securing Flows within the Cloud-Ready Data Center, **juniper.net/us/en/local/pdf/implementation-guides/8010034-en.pdf**
- Adaptive Threat Management, **juniper.net/us/en/solutions/enterprise/security-compliance/adaptive-threat-management/**
- Secure Network Access Across the Distributed Enterprise, **juniper.net/us/en/local/pdf/implementation-guides/8010031-en.pdf**
- J Series/SRX Series Multipoint VPN Configuration with Next-Hop Tunnel Binding, **http://kb.juniper.net/kb/documents/public/junos/jsrx/JSeries_SRXSeries_Multipoint_VPN_with_NHTB_12.pdf**

Junos Configuration

- *Virtual Router Redundancy Protocol (VRRP)*, **juniper.net/techpubs/en_US/junos10.2/topics/reference/statement-hierarchy/vrrp-configuration-hierarchy.html**
- *Link Aggregation (LAG)*, **juniper.net/techpubs/en_US/junos10.0/topics/concept/interfaces-lag-overview.html**
- *Bidirectional Forwarding Detection (BFD)*, **juniper.net/techpubs/software/junos/junos94/swconfig-routing/configuring-the-bfd-protocol_3.html**
- *Graceful Routing Engine Switchover (GRES)*, **juniper.net/techpubs/en_US/junos10.2/topics/task/configuration/routing-engine-redundancy-configuring.html**
- *Nonstop Active Routing (NSR)*, **juniper.net/techpubs/en_US/junos10.2/topics/task/configuration/nsr-configuring.html**
- *In-Service Software Upgrade (ISSU)*, **juniper.net/techpubs/en_US/junos10.2/information-products/pathway-pages/high-availability/high-availability.html**
- *Virtual Chassis (VC)*, **juniper.net/techpubs/en_US/junos9.5/topics/example/virtual-chassis-basic.html**
- *Intrusion Detection and Prevention (IDP)*, **juniper.net/techpubs/software/junos-security/junos-security10.0/junos-security-swconfig-security/config-idp-policies-chapter.html**
- *Point to Point over ATM (PPPoA)*, **juniper.net/techpubs/software/junos-security/junos-security10.2/junos-security-swconfig-interfaces-and-routing/jd0e22823.html**
- *Quality of Service (QoS)*, **juniper.net/techpubs/en_US/junos10.3/information-products/pathway-pages/cos/index.html**
- *Random Early Discard (RED)*, **juniper.net/techpubs/en_US/junos10.3/information-products/pathway-pages/cos/red-drop-profiles.html**

Polycom Documentation

- VSX Getting Started, **polycom.com/global/documents/support/user/products/video/vsx_series_getting_started_guide.pdf**
- VVX Admin Guide, **polycom.com/global/documents/support/setup_maintenance/products/voice/spip_ssip_vvx_Admin_Guide_SIP_3_2_2_eng.pdf**
- Polycom HDX 8000 Series, **polycom.com/products/telepresence_video/telepresence_solutions/room_telepresence/hdx8000.html**

RFCs

- RFC 2597, *Assured Forwarding PHB Group*
- RFC 3246, *An Expedited Forwarding PHB (Per-Hop Behavior)*
- RFC 4026, *Provider Provisioned VPN Terminology* (March 2005)
- RFC 4364, *BGP/MPLS IP Virtual Private Networks* (VPNs)
- RFC 4594, *Configuration Guidelines for DiffServ Service Classes*
- RFC 5880, *Bidirectional Forwarding Detection (BFD)*
- RFC 5881, B*idirectional Forwarding Detection (BFD) for IPv4 and IPv6 (Single Hop)*

Appendix F: Acronyms

A

AAA	Authentication, Authorization, and Accounting
AAC	Advanced Audio Codec
ACP	Admission Control Plug-in
ADSL	Asymmetric Digital Subscriber Line
AS	Autonomous System
ASP	Application Service Provider
AVC	Advanced Video Codec

B

BFD	Bidirectional Forwarding Detection
BGP	Border Gateway Protocol
BPDU	Bridge Protocol Data Unit
BSR	Broadband Services Router

C

CAC	Call Admission Control
CBR	Constant Bit Rate
CDR	Call Detail Record
CE	Customer Edge
CMA™	Polycom® Converged Management Application™
CMTS	Cable Modem Termination System
CoS	Class of Service

D

DDoS	Distributed Denial of Service
DHCP	Dynamic Host Configuration Protocol
DMA™	Polycom® Distributed Media Application™
DNE	Diameter Network Element
DPI	Deep Packet Inspection
DSA	Dynamic Service Activator
DSCP	Differentiated Services Code Point
DSLAM	Digital Subscriber Line Access Multiplexer

E

ESM	Ethernet Switch Module

F

FPC	Flexible PIC Concentrator
FRR	Fast Reroute
FSP	Flexible Service Processor

G

GOP	Group of Pictures
GRE	Generic Routing Encapsulation
GRES	Graceful Routing Engine Switchover

H

HBA	Host Bus Adapter
HDVC	High Definition Videoconferencing
HDX™	Polycom® High Definition Conferencing Platform
HEA	Host Ethernet Adapter
HMC	Hardware Management Console

I

IDP	Intrusion Detection and Prevention
IGP	Interior Gateway Protocol
IGMP	Internet Group Management Protocol
IPSec	Internet Protocol Security
ISDN	Integrated Services Digital Network
IS-IS	Intermediate System to Intermediate System
ISSU	In Service Software Upgrade

J

JSRC	Junos Session and Resource Control

K

KPI	Key Performance Indicator

L

LAG	Link Aggregation Group
LDP	Label Distribution Protocol
LSI	Label Switched Interface
LSP	Label Switched Path

M

MCU	Multi-Conferencing Unit
MPLS	Multiprotocol Label Switching
MSP	Managed Service Provider
MTU	Maximum Transmission Unit

N

NIC	Network Information Collector
NLRI	Network Layer Reachability Information
NSP	Network Service Provider
NSR	Non-Stop Routing
NTSC	National Television System Committee

O

OEM	Original Equipment Manufacturer
OSS	Operations Support Systems
OTX™	Polycom® Open Telepresence™ Experience

P

PAL	Phase Alternate Line
PDM	Power Distribution Module
PE	Provider Edge
PHB	Per-Hop Behavior
PIC	Physical Interface Card
PIM	Protocol Independent Multicast
PLP	Packet Loss Priority
PM	Power Module
PoE	Power over Ethernet
PVST	Per-VLAN Spanning Tree
PWE3	PseudoWire Emulation Edge to Edge

Q

QoE	Quality of Experience
QoS	Quality of Service

R

RCEF	Resource Control Enforcement Function
RED	Random Early Detection
RMX™	Polycom® Real-Time Media Conferencing Platform
RPX™	Polycom® RealPresence™ Experience
RSTP	Rapid Spanning Tree Protocol
RSVP	Resource Reservation Protocol

S

SAE	Service Activation Engine
SDB	Session Database
SECAM	Séquentiel Couleur à Mémoire, (French for "Sequential Color with Memory")
SIP	Session Initiation Protocol
SLA	Service Level Agreement
SOAP	Simple Object Access Protocol
SRC	Session Resource Control
SSC	Subscriber Service Control
SSL	Secure Socket Layer

T

TDM	Time Division Multiplexing
TWAMP	Two-Way Active Measurement Protocol

V

VBR	Variable Bit Rate
VCCP	Virtual Call Center Platform
VMR	Virtual Meeting Room
VPLS	Virtual Private LAN Service
VRF	VPN Routing and Forwarding
VRRP	Virtual Router Redundancy Protocol
VVX™	Polycom® Business Media Phone

W

WMA	Windows Media Audio

Understanding and Implementing High Definition Videoconferencing

By integrating high definition (HD) video and audio technology, videoconferencing has now reached the stage where it remarkably reproduces "live" meetings around the world. Technological advances in network performance, security and accessibility have elevated HD videoconferencing onto the center stage of distributed enterprises. This is why analysts forecast explosive market growth for HD videoconferencing in the next two to three years. HD videoconferencing enables enterprises to reduce the expense and lost productivity from excessive travel, while service providers who offer HD videoconferencing service will be ready to serve these expanding business needs.

This handbook describes how service providers and enterprises can now cost-effectively expand their network to assure high performance and quality of HD videoconferencing services while avoiding common pitfalls along the way. Beginning with a basic, high-level overview of design principles, followed by practical use cases, readers can gain a thorough understanding of Juniper Networks HD videoconferencing solution. Serving as a reference tool, the handbook illustrates how to design, deploy and operate the most demanding HD videoconferencing services using Juniper and Polycom state-of-the-art devices to assure quality, guarantee availability and provide unprecedented scalability and security.

"Organizations are looking to reap the benefits of high definition videoconferencing, but are also looking to contain their network investment. Juniper and Polycom have developed a capability which combines the best of both worlds, ensuring highest quality delivery for the most important videoconferencing traffic, in many cases without requiring additional bandwidth. This convenient handbook describes how to implement HD videoconferencing on Juniper equipment, whether you're in a home office or a traditional site, and it shows how Polycom and Juniper can deliver high-quality video by actively coordinating at the time of the conference. The content is valuable whether you're looking to implement for your own organization or you are a network service provider seeking to leverage to the existing infrastructure to offer HD videoconferencing service."

– Scott Stevens
VP Technology and Worldwide Systems Engineering
Juniper Networks

7100137-001-EN Dec 2010

JUNIPer
NETWORKS