Telepresence and High Definition Videoconferencing on Converged IP Networks

A New Architecture for Delivering Visual Collaboration Services

December 2010

Study sponsored by:

Table of Contents

List of Figures

Introduction

From the inception of group videoconferencing until around the mid-1990s, videoconferencing calls were hosted over Integrated Services Digital Network lines (a.k.a. ISDN); a set of communication protocols or standards that digitizes data to enable traditional telephone lines to carry digital voice, video, and other traffic.

The use of ISDN data lines for videoconferencing offered several benefits including global availability, relatively low up-front and fixed monthly network costs, support for high speed video calls (via aggregation of numerous ISDN channels), and support for calls between enterprises (a.k.a. B2B videoconferencing). These benefits allowed videoconferencing to grow from a technical curiosity into a high-impact visual collaboration tool used by thousands every day.

There are also many disadvantages related to using ISDN for videoconferencing including limited reliability, relatively high usage fees, the need to manage an additional network, the need to purchase and manage expensive and complex ISDN switching equipment for large scale deployments, and the need to ISDN-enable videoconferencing systems and infrastructure devices. These disadvantages made ISDN videoconferencing difficult to scale, and the reliability issues and per-minute usage fees served to limit usage and motivate users to use low connection speeds (to cut costs) that delivered relatively poor video quality – at least by today's standards.

In the mid 90s, the first IP-capable video systems emerged allowing enterprises to host video traffic on their enterprise local area and wide area networks (LAN and WAN). This newfound ability to leverage the same data network carrying the organization's data traffic overcame virtually all of the disadvantages listed above. Note that the migration to IP is not meant to imply that ISDN videoconferencing is no longer in use today. ISDN is still heavily used by many organizations, and is especially useful for B2B / inter-enterprise video calling. However, for many enterprises, the ability to host video traffic on their IP network paved the way for both increased videoconferencing utilization AND the use of higher connection speeds to enable an enhanced user experience.

Over time, thanks to a wealth of new applications and the increased bandwidth appetite of existing applications, the demands – in terms of both capacity and performance - on the typical enterprise IP network (both LAN and WAN) have increased significantly.

Videoconferencing, in particular, places a very high demand on enterprise IP networks because it requires a combination of high capacity (bandwidth requirements tend to increase as video quality increases), low latency / delay to allow interactive conversations, and little or no packet loss or network errors.

Meeting these strict performance requirements is especially challenging for network service providers (NSPs) looking to host high quality, real-time video traffic for enterprise customers on their global networks. In many cases, the heavy burden of real-time rich media traffic is too much for the existing networks to handle. Depending upon the situation, hosting voice and video on an unworthy network can result in low quality video and audio calls (poor audio, poor video, unacceptably long latency / delay), or even dropped or interrupted sessions. In addition and perhaps worst of all, the heavy footprint of rich media traffic can impact or even block the flow of other IP traffic on the network.

This white paper provides insight into the traditional methods and a new option for supporting high bandwidth, high quality visual collaboration sessions over existing converged networks.

Traditional Methods for Hosting Video over IP Networks

Wainhouse Research (WR) has noted that interest in IP videoconferencing has increased significantly in recent years. This has resulted in several developments:

1) Expanded deployments - to make video more accessible to more people in more locations
2) The use of more bandwidth per call - to provide a high quality user experience

The bandwidth used per call has been impacted by several recent market developments including the availability of high definition (HD) video systems. Depending upon the systems in use, a single HD video call can require between 512 kbps and 4 Mbps (or more) of bandwidth. The table below demonstrates how a typical enterprise's bandwidth requirements might have changed over time.

	Yesterday	**Today**
# of Systems Installed	30	50
Typical Call Speed	384 kbps	1 Mbps
Video Resolution	CIF (320 x 240)	720p (1280 x 720)
% of Systems In Use Simultaneously	20%	30%
Total Bandwidth Utilized	2304 kbps (or 2.25 Mbps)	15360 kbps (or 15 Mbps)

The use of multi-codec / multi-display videoconferencing systems (dubbed "immersive telepresence" systems by many vendors), which require between 3 Mbps and 20 Mbps (or more) per site per session, further increases bandwidth utilization.

The options for enterprises and service providers seeking to host high quality IP video traffic fall into two general categories; 1) existing bandwidth options, and 2) bandwidth expansion options.

Existing Bandwidth Options

<u>Option 1 - Do Nothing / Make No Bandwidth or Network Changes</u>
The first option is to host the IP videoconferencing traffic on the existing data network without making any network changes or upgrades. Due to its simplicity and low (actually no) cost, this method is in use by many enterprises today.

Although not recommended, it *should* be possible to host IP video traffic over the existing LAN / WAN when:

- Videoconferencing usage is very limited
- Relatively low call speeds are used
- The existing network is sufficiently over-provisioned such that it has little or no bandwidth contention

The third item regarding over-provisioning is especially important as it means that the customers are probably paying for more network capacity or performance than they actually require.

In a world of tight budgets and cost control, it is far more likely that an enterprise's network is properly provisioned (or "right-sized"), and not heavily over-provisioned. In this situation, the addition of IP video traffic is likely to overwhelm the existing network resulting in:

- Unreliable or inconsistent performance during video calls
- Further degradation of video quality as more calls are placed on the network
- Decreased throughput or network performance for other applications

In short, for most organizations this option is a ticking time bomb.

<u>Option 2 - Implement Call Control (may be used with option #3 below)</u>
Another choice is for the organization to implement some form of call control or other mechanism to limit bandwidth utilization. This can be accomplished via the use of rule sets within videoconferencing gatekeepers, which is dubbed call admission control, by adjusting video system settings to maximize call speeds, or externally by enforcing usage policies. Typical call control / bandwidth limiting methods include:

- Limiting the number of concurrent calls within the environment
- Limiting the number of concurrent calls between specific locations
- Limiting the bandwidth used per call (which limits call quality)
- Limiting the total amount of bandwidth that can be used for video calls simultaneously

For example, an organization could protect its London facility's bandwidth by allowing only three (3) concurrent video calls of up to 768 kbps at a time or setting the maximum total bandwidth available for videoconferencing to 2 Mbps.

Although call control can be used to protect the network, it has two key disadvantages;

1) Call control limits the availability of videoconferencing to the user base, and
2) Call control is typically unaware of and therefore unable to respond to real-time changes in network availability or capacity.

Option 3 - Implement Quality of Service (QoS) on the Network (often used with option #2)
QoS involves prioritizing certain types of network traffic in an effort to provide consistent and predictable network performance (throughput, latency, packet loss, etc.). For example, the team at WR uses QoS running on a high performance, converged MPLS network to enable high-bandwidth, high-quality IP video calls between our facilities.

QoS can be implemented in a variety of ways including manually by configuring each of the routers and switches in the environment individually or centrally via the use of a policy engine that pushes out policy settings and instructions to the appropriate network devices.

QoS is a reasonable option for some enterprises and service providers seeking to host IP videoconferencing traffic on their existing data network. However, QoS is not the panacea for all situations since QoS:

- Does not eliminate the need for sufficient bandwidth to handle the video and non-video traffic
- Does not guarantee performance – even for high priority traffic. It simply prioritizes the traffic as much as possible.
- Does not protect legacy traffic from being choked off by the prioritized traffic
- May require the deployment of additional equipment or network hardware / software upgrades
- Can be complicated to deploy, configure, and manage, which increases the management burden
- Typically does not work as traffic traverses between different carrier's networks

In short, when properly implemented in a properly equipped, configured, and QoS-capable network environment, QoS can yield exceptional results. However, globally deployed QoS is out of reach for many enterprises and service providers.

Bandwidth Expansion Options

Option 1 - Add Additional Bandwidth to Existing Network
One of the most obvious options for those seeking to host IP video traffic on their data network is to implement bandwidth expansions in key locations. For example, if an organization wishes to conduct 5 Mbps of video calls between its New York and Hong Kong facilities, it could add 5 Mbps of bandwidth to each of the access links (a.k.a. local loops or last mile connections) between those locations and its service provider's core.

The advantages of this option include relative ease of deployment and limited cost as bandwidth upgrades would be performed in only certain locations. However, adding additional bandwidth in specific locations to host video traffic:

- May involve significant additional cost (depending upon location, existing local loops in place, and additional capacity required)
- May force the deployment of another large network link resulting in very high cost
- Does not guarantee performance
- Does not protect legacy traffic from being choked off by the video traffic
- Is relatively inefficient - although the new bandwidth will be available to host other traffic, the new bandwidth is only truly necessary to host video calls.
- Is likely to be consumed by bandwidth-intensive, non-real time applications (e.g. email, web browsing).

For organizations with very specific requirements (e.g. the need to conduct a limited number of video calls between a limited number of locations), the bandwidth augmentation option might make sense. However, for many organizations it is not possible to accurately predict the needs of their users.

Option 2 - Deploy an Overlay Network

This option involves the deployment of a totally separate IP data network dedicated to hosting rich media traffic. The obvious and primary benefit of this approach is that it completely isolates the IP video traffic from the existing traffic on the wide area network. Unfortunately, there are many disadvantages of this approach including:

- High cost associated with deploying an entirely new network including
    - Cost of local loops (installation, monthly fees)
    - Cost of network access / port fees (initial connections, monthly fees)
    - Cost of CPE hardware (purchase, installation, maintenance)
- Inefficiency because the new bandwidth will be used to host video traffic only
- Additional management burden associated with maintaining a second data network
- Issues related to the routing of traffic between the primary and overlay data networks
    - May require complex routing rules / tables / VLANs
    - May result in loss of security, QoS / prioritization, decreased performance, etc.

In some ways, this approach of using an isolated network is similar to the use of ISDN to host video calls in the past – without the added ISDN benefits of low cost (the majority of ISDN costs are usage-based) and native support for B2B video calls.

Many organizations have chosen to deploy overlay networks to, a) protect the video conferencing traffic from other network traffic (and vice versa), b) avoid the need to augment their existing network capacity, upgrade router / switch hardware and software, or c) implement QoS on a global basis. However, due to the high cost and management burden, most organizations view an IP overlay network dedicated to videoconferencing as a short-term solution to be used until their primary data network is capable of successfully hosting the video traffic.

Summary

As described within this section, enterprises and service providers seeking to host high quality IP video traffic on their networks have a variety of options. However, each option introduces a series of compromises such as limited reach, restrictions on call volume or quality, decreased reliability, increased cost, decreased efficiency, and additional management burden.

The fact is that in many situations, these traditional – and widely deployed – options do not address the overall requirement of successfully and reliably hosting high quality videoconferencing calls over converged networks without impacting the delivery of other data.

Solution Spotlight – Polycom / Juniper Architecture

WR believes that one of the reasons why IP video is difficult to host on existing data networks is the lack of direct and ongoing coordination between the videoconferencing environment and the network environment.

This is not to say that the video and network environments are not aware of each other and working together to some degree. For example …

1) Many video gatekeepers allow system administrators to enter information about network capacity to enable video call admission control.
2) Video devices can set traffic prioritization / QoS markings on data to tell the network how to process the video traffic.
3) Video devices can detect when packets are being lost (presumably due to network issues) and can adjust their call speed / performance accordingly.

The items above are examples of how the video and network environments work together without having any first-hand, real-time knowledge of each other. For example, the devices in item 3 above would not decrease their bit rate based on a request from the network environment. They would decrease their bit rate when they detect that their data packets are not getting through.

Ideally, the video and network environments would pass real-time status information back and forth to allow each system to pro-actively alter its behavior in response to situational changes in the other environment. Participation in this type of feedback loop would be beneficial to both environments.

Unfortunately, in the typical enterprise video calls are initiated in the video environment, and the video data traffic is simply dropped onto the network environment. The video environment does not modify its behavior to accommodate the real-time network environment. Nor does the network modify its behavior to accommodate the video traffic. As a result, running IP video over a converged network is often a high risk venture.

The sponsors of this study, Polycom and Juniper Networks, have created a joint solution that addresses many of these issues by coordinating the activities of the video environment with those of the network environment to maximize efficiency and performance. In addition, the solution provides a guaranteed SLA for the delivery of the video traffic by, a) making sure unauthorized traffic does not enter the network (a.k.a. call admission control), and b) ensuring that the approved traffic receives the appropriate level of prioritization.
The key enabler of this solution is a newfound communication path between the Polycom Distributed Media Application (DMA) product and the Juniper Junos Space Session and Resource Control (SRC) product.

Polycom Distributed Media Application (DMA)

DMA 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.

Juniper Junos Space Session and Resource Control (SRC) Platform

SRC is a network resource allocation solution system that enables network resources to be dynamically reconfigured based on user and application requirements, network status, time of day, service provider offerings, and other parameters. Within this joint solution, the SRC works with the Polycom DMA to create a real-time feedback loop between applications, users, and the network, thereby connecting the videoconferencing environment with the network layer. The SRC utilizes widely adopted standards-based open interfaces to maximize interoperability with the broadest range of network elements, applications, and management platforms. By leveraging open interfaces, the SRC can integrate with any network, application, or service. The SRC's support for dynamic policy enforcement on a per-service or per-application level allows service providers and enterprises to intelligently and optimally utilize their existing network infrastructure.

The Joint Polycom / Juniper Solution

The diagrams below illustrate the difference between how IP video traffic traverses a traditional network and how it traverses a Polycom / Juniper powered network.

Example #1 – Best Effort Traffic Only

The first diagram shows how standard (non-real-time) data traffic might traverse from one location (Enterprise Site A), through the Network Service Provider (NSPs) network core, to another location (Enterprise Site B). In this case, the traffic shown represents standard email traffic travelling from an email client user to the corporate mail server.

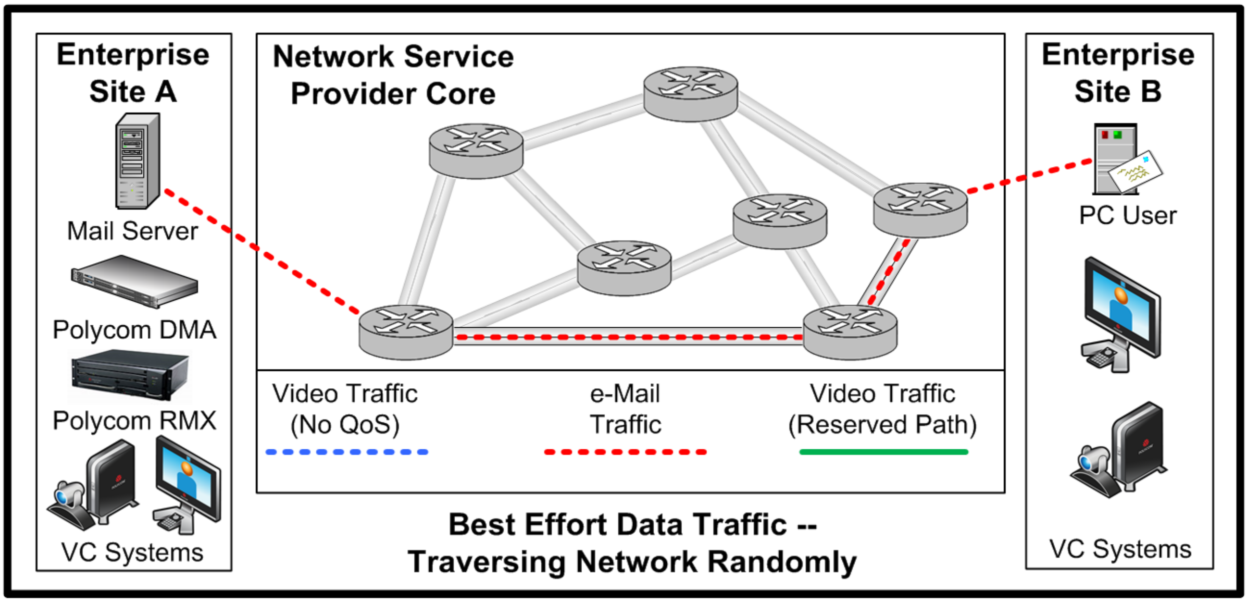

Figure 1: Best Effort Traffic Traversing a Typical WAN

<u>Example #2 – Videoconferencing Traffic Added to the Network</u>
The next diagram shows the impact of adding IP videoconferencing traffic to the network (see the dashed blue lines below) without making any network changes or implementing prioritization QoS. In this case, the combined bandwidth required for the video and email traffic exceeds the capacity of the network links. As a result, both the video and email traffic will suffer some degree of congestion-related packet loss.

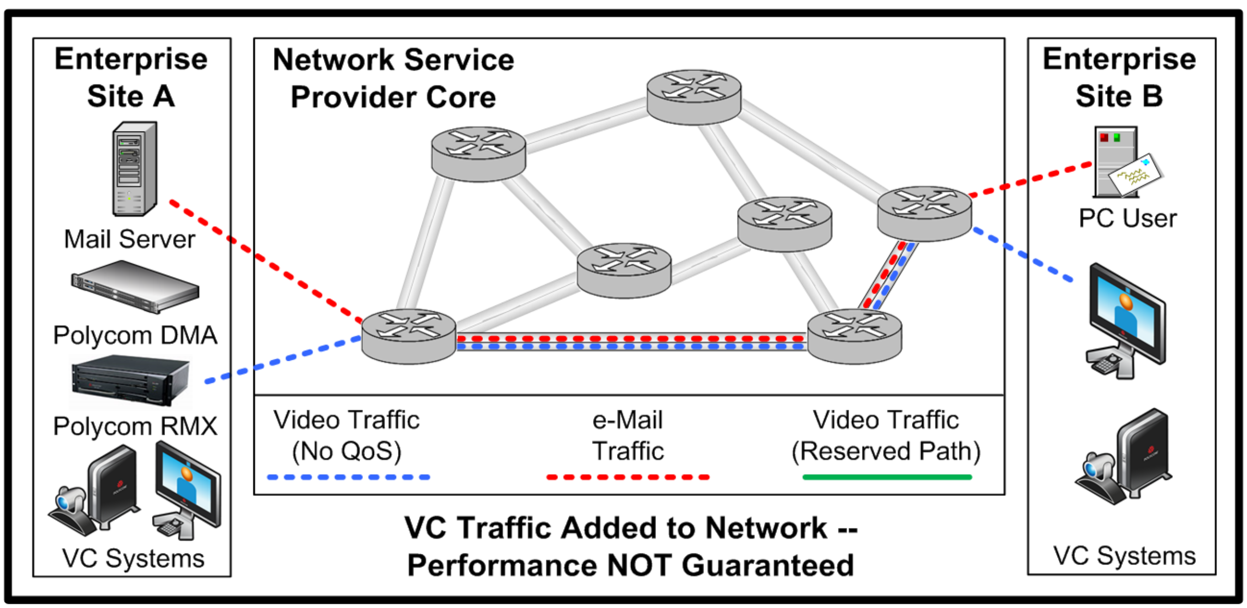

Figure 2: Videoconferencing Traffic Added to a Typical WAN

The key takeaway here is that because the video traffic is NOT prioritized in any way, the video and email traffic share common network links, and therefore the performance is not guaranteed. In effect, the video environment has simply dropped the IP traffic onto the network – without regard for whether the network is prepared or able to handle the burden of the IP video traffic.

Example #3 – Videoconferencing Traffic Traversing a Polycom / Juniper Powered Network
The secret sauce of the joint Polycom / Juniper solution is the real-time communication between the videoconferencing layer (DMA) and the network layer (SRC) within the enterprise and service provider environments which enables the creation of a reserved, guaranteed performance network path for the video traffic.

As shown in the diagram below, this architecture requires the following elements:

- Juniper Networks Edge Router(s) installed somewhere within the service provider's network
- Juniper Networks Junos Space Session and Resource Control (SRC) running within the network environment
- Polycom DMA installed within the enterprise environment (or hosted externally but dedicated to that particular enterprise client)
- Standards-based (H.323, SIP, etc.) video systems
- Polycom RMX multipoint conference platform (required for multipoint calls only).

Although not highlighted within the drawing, the Juniper SRC must be running within the environment. Note that SRC is not a device … it is a communication "signaling" layer riding on top of the Juniper environment and communicating with all Juniper routers within the network.

As shown, the Polycom / Juniper solution allows the hosting of traffic using different QoS levels within a single network link. In this case, the bottom path includes email traffic, best effort videoconferencing traffic, and "protected", high QoS videoconferencing traffic.

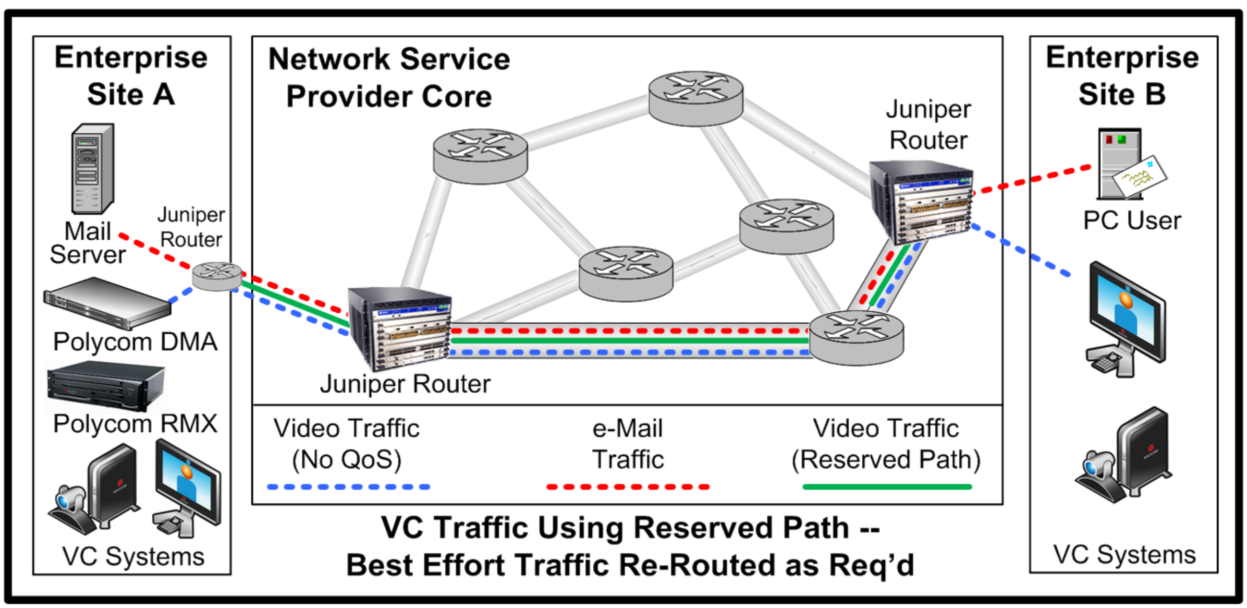

Figure 3: Videoconferencing Traffic Traversing a Polycom / Juniper Powered Environment

Additional Notes:

1) The diagram above includes a Juniper router installed on the customer premise. While not strictly necessary, the use of an onsite Juniper router extends the reach of the service guarantee beyond the network core and onto the customer's premise.

2) In the diagram above, the Video Traffic (Reserved Path) represented by the solid green line is traversing the network with assured quality (via a dynamically created QoS level) over a reserved and guaranteed network path.

The first step toward using this solution is to register all video systems and devices (video bridges, etc) to the DMA. Both H.323 gatekeeper and SIP proxy registrations are supported.

To launch a call, users simply dial the address (typically the E.164 address, H.323 name, or SIP URI) of the endpoint or video meeting room (typically a pre-defined virtual meeting room on the DMA) they are trying to reach. When the Polycom DMA detects the inbound call request, it sends a request to the Juniper SRC located in the service provider's core. For videoconferencing-savvy folks, this is essentially a call access control request. In this case, however, the request is for the required network services (bandwidth, quality of service, etc.). Upon receiving DMA's call request, the Juniper SRC assesses the network status to determine whether the call requested can be successfully hosted on the network given the current utilization of network resources by videoconferencing and other applications.

<u>Option 1 – The Call Request CAN be Accommodated</u>
If the Juniper SRC determines that the call request can be accommodated, the following occurs automatically:

1) The Juniper SRC reserves the required bandwidth across the network. Depending upon the environment, this may involve signaling the MPLS LSPs (labeled signaled paths) or the creation of a new quality of service (QoS) / class of service (CoS) level to host the IP videoconferencing network traffic.

2) Upon completion of the call, the Polycom DMA notifies the Juniper SRC that the call has been terminated, at which time the bandwidth reserved for the video traffic is relinquished and made available to host other data and/or video traffic.

Embedded within the above are two key points:

1) The above process happens automatically and without the need for manual intervention
2) Once the call is completed, the network returns to its default state

Techno-savvy folks will note that a small portion of the network path in the diagram above is NOT reserved or protected (see the dashed blue lines in the Enterprise section of the drawing). This is because the above bandwidth reservation / QoS implementation method is optimized to work between Juniper routers. The assumption is that a typical service provider environment will already include some Juniper routers, and that these routers will be relatively close (in terms of hops and transport latency) to the customer premise. Note that only those routers which connect to the enterprise sites must be Juniper routers.

<u>Option 2 – The Call Request CANNOT be Accommodated</u>
On the other hand, the Juniper SRC may determine that the DMA call request CANNOT be accommodated due to a lack of capacity, some form of network outage or issue, or the inability to reserve bandwidth / establish QoS between the required locations. In this situation, the SRC will step through a user-defined set of policies in an effort to support the requested call to the best degree possible. The options available to the SRC include:

1) Automatically adjust or augment the existing QoS / CoS policies on the network to allow the IP videoconferencing traffic to successfully traverse the network with the required level of performance. These network policy adjustments are based on the importance level of the call initiator. For example, in order to accommodate a "Gold" level user's video call, the system might decrease the bandwidth available to lower priority "Silver" level users.

   Readers should note that this represents a real-time re-allocation of the EXISTING network resources to accommodate the request of a specific application (in this case videoconferencing). Readers should also note than in any "re-allocation of network resources, some applications may be impacted negatively, i.e. yield bandwidth to another application.

2) Reject the call request because the network cannot accommodate the video traffic at the required level of performance

Note that the DMA can also be configured to allow the call to proceed at a reduced call speed if the originally requested call speed cannot be accommodated.

For example, let's assume that the Polycom / Juniper solution has received a call request for a multi-codec / multi-display (a.k.a. telepresence) system to connect to another multi-display system. In this case, the total bandwidth required would be 3 Mbps (1 Mbps per screen x 3 screens) to provide an HD1080p call experience at 30 frames per second.[1]

---

[1] The bandwidth requirement of 1 Mbps for a 1080p / 30 fps video call assumes the use of H.264 High Profile – available on Polycom HDX systems running software version 2.6 or higher. Without High Profile, a 1080p / 30 fps video call would require at least 2 Mbps.

An important feature of the Juniper SRC is that it allows users to define the policies and workflow that best addresses their specific requirements. For example, upon receiving the call request described above, the SRC could be configured to do any (or many) of the following:

1) Approve the request and reserve the required bandwidth
2) Approve the request and create a new QoS level for the video call because the current network utilization made it impossible to reserve the bandwidth
3) Adjust the QoS policies already in place (and potentially in use) on the network to make more bandwidth available for the requested call.
4) Approve the call request without providing any form of bandwidth guarantee or QoS
5) Reject the call if the network cannot provide the minimal acceptable level of performance

It is important to understand that the policies within SRC can be defined to consider the current network status and even the types of systems involved. For example, the SRC could be set to:

1) Allow up to 30% of the actual capacity of a particular network link at any moment to be used to host video traffic. This type of definition allows the policies to automatically react to network outages or errors impacting capacity.
2) Prioritize multi-codec (telepresence) calls over other video calls (depending on the defined service level in the DMA that is coordinated with network policy).

The complete set of configuration options within the Polycom / Juniper environment is beyond the scope of this document. However, readers should understand that this architecture is dynamic, network-aware, and configurable on a per-customer basis. The ability to configure the SRC with a different set of policies for each customer makes this solution well suited for service providers.

Business Case for the Network Service Provider

The Polycom / Juniper solution allows service providers to offer guaranteed performance levels for IP video and telepresence traffic over converged data networks. Key benefits for service providers include:

- Decreased cost of service delivery by helping NSPs avoid the need for CAPEX investments related to network core expansions (e.g. additional network lines, switch / router expansions, etc.)
- Ability to offer differentiated services for less money to end-user customers by avoiding the costs associated with network expansions or overlay deployments (e.g. additional local loops, additional or expanded CPE equipment, etc.)
- Increased efficiency by enabling service providers to utilize a greater portion (take advantage of higher utilization rates) of both the local loops and their network without sacrificing videoconferencing application quality and performance

- Simplified service deployment by
    - Avoiding the need to configure a second network
    - Avoiding the need to define routing, security, and QoS traversals between networks
- Expedited service deployment by leveraging existing bandwidth
- Simplified network management for both the service provider and end-user
- Simplified SLAs that are easier to define, easier to monitor, and easier to meet
- Ability to protect both the IP video traffic from other data traffic AND the data traffic from the IP video traffic
- Ability to provide guaranteed service levels and network performance for other – and even future - network-centric applications

Real World Example

A service provider has received a request from an enterprise client seeking to run high bandwidth (high definition, multi-display, etc.) video calls over the service provider's network backbone. The client's requirements include 24/7 access, support for ad-hoc video calling, guaranteed performance for certain high profile users / systems, and activation within the next four weeks.

The client's wide area network (WAN) currently consists of DS3 (45 Mbps) local loop connections between each of its locations in New York, London, and Hong Kong and the service provider's network core. At any given moment, the utilization on each DS3 ranges from 25% to 75%.

<u>Traditional Method</u>

With local loop utilization rates reaching such high levels, and a need to optimize the experience for key users, the service provider is unlikely to recommend the hosting of the IP video traffic over the existing network. Instead, the service provider would probably recommend the deployment of an overlay network dedicated to IP video traffic. This would include the deployment of an additional DS3 local loop to each location. This would result in the following costs:

One-Time Fees	Recurring (Monthly) Fees
Enterprise Costs / Fees	
Local loop installation fees	Local loop monthly fees
CPE equipment purchase fees	Network port fees
NSP installation fees (cost for programming the CPE equipment, etc.)	
Service Provider Costs / Fees	
Network equipment (e.g. core router ports, edge equipment interfaces, line cards, etc.)	Recurring support costs
Engineering / programming costs (network design, routing / security policies between networks, etc.)	

Figure 4: Typical Costs Associated with an Overlay Network

By the time the overlay network is up and running, which would probably take 8 or more weeks, the total cost to the end-user and the service provider is likely to be thousands up front and thousands per location each month on an ongoing basis. This is a hefty price to pay for circuits that are likely to be used for only a limited number of hours per month.

Polycom / Juniper Method

In a Polycom / Juniper powered environment, this client's request could be accommodated almost immediately and without the need for additional network links or router ports. In this case, the costs would be limited to the cost of a few resources on the Juniper SRC. This, of course, assumes the following:

- The service provider has already deployed SRC within its core
- The enterprise has DMA installed within its video environment

This option is not only faster to deploy, but also less expensive (the combined solution has a list price below US $75,000, which typically compares favorably to the monthly cost of adding more WAN bandwidth) and more efficient than traditional options.

Summary / Conclusion

The joint Polycom / Juniper architecture allows service providers to:

- Squeeze more revenue out of their existing network without the need for additional capital investments or network expansions.
- Provide bandwidth, quality of service, and video call control via a single offering over their existing network and infrastructure.
- Expedite service delivery by avoiding the need for additional network deployments
- Offer enhanced and differentiated network services at lower cost to their clients, thereby leading to increased customer satisfaction and retention.

Modern day service providers recognize that the days of adding customers to generate revenue to pay for CAPEX enhancements of the network are basically gone. In today's environment, the key to profitability and success is to increase margins and maximize revenue by providing greater, differentiated value without additional spending.

The joint Polycom / Juniper solution allows service providers to control costs while expanding services; a win / win scenario for providers and customers alike.

About Wainhouse Research

Wainhouse Research, LLC (WR) provides analysis and consulting on the market trends, technologies/ products, vendors, applications, and services in the collaboration and conferencing fields. Areas of coverage include hardware, software, and services related to audio, video, and web conferencing, unified communications, and enterprise social networking. The Company publishes market intelligence reports, provides customized strategic and tactical consulting and studies, and produces industry events (conferences). Additionally, the Company operates industry-focused and end user-focused Web sites and publishes a weekly sponsored bulletin for news and analysis. For more information on Wainhouse Research, visit www.wainhouse.com.

About the Author(s)

Ira M. Weinstein is a Senior Analyst and Partner at Wainhouse Research, and a 20-year veteran of the conferencing, collaboration and audio-visual industries. Prior to joining Wainhouse Research, Ira was the VP of Marketing and Business Development at IVCi, managed a technology consulting company, and ran the global conferencing department for a Fortune 50 investment bank. Ira's current focus includes IP video conferencing, network service providers, global management systems, scheduling and automation platforms, ROI and technology justification programs, and audio-visual integration. Mr. Weinstein holds a B.S. in Engineering from Lehigh University and can be reached at iweinstein@wainhouse.com.

Andrew W. Davis is a researcher, analyst, and opinion leader in the field of collaboration and conferencing. He is a co-founder of Wainhouse Research. Prior to Wainhouse Research, he held senior marketing positions with several large and small high-technology companies. Andrew has published over 250 trade journal articles and opinion columns on multimedia communications, videoconferencing, and corporate strategies as well as numerous market research reports and is the principal editor of the conferencing industry's leading newsletter, The Wainhouse Research Bulletin. A well-known industry guest speaker, Mr. Davis holds B.S. and M.S. degrees in engineering from Cornell University and a Masters of Business Administration from Harvard University and can be reached at andrewwd@wainhouse.com.

About Juniper

(Copy provided by Juniper)

Juniper Networks® delivers innovative software, silicon, and systems that transform the experience and economics of networking for global service providers, leading enterprises, and public sector organizations. Our core routers, switches, and security hardware and software run the world's largest and most demanding global networks.

From the start, Juniper saw the opportunity to help customers build and accelerate business value from their IT infrastructures. Since 1996, we have helped our customers stay ahead of the demands posed by the exponential growth in network users and endpoints, while meeting the business imperatives for high performance, reliability, and absolute security.

Thirteen years later on October 29th, 2009, the 40th Anniversary of the Internet, Juniper unveiled its vision for the next decade—a future of networking that is open, scalable, simple, secure, and automated. Marked by the company's move to the New York Stock Exchange and the unveiling of a new corporate brand identity, Juniper offers an expanded portfolio and go-to-market reach, which now includes software, silicon and systems for routing, switching, and security across enterprise and service provider markets worldwide. We also revealed Junos® as a new ingredient brand, reflecting the company's expanded software strategy, portfolio and delivery models that showcase the value of Junos software as the "brains" behind today's best networks.

Today, Juniper remains uniquely positioned to maintain industry leadership based on our core competencies in architecture, silicon design, and our open cross-network software platform that includes the popular Junos operating system, Junos Space network application platform, and Junos Pulse integrated network client. Juniper offers a broad product portfolio that spans routing, switching, security, application acceleration, identity policy and control, and management designed to provide unmatched performance, greater choice, and true flexibility while reducing overall total cost of ownership (TCO). In addition, through strong industry partnerships, Juniper is fostering a broad ecosystem of innovation across the network. Juniper has an exclusive focus on providing solutions that solve the toughest and most complex networking problems—delivering on the promise of the "New Network."

About Polycom

(Copy provided by Polycom)

Polycom, Inc. is the global leader in telepresence, video, and voice solutions and a collaborative communications visionary empowering people to connect and collaborate anytime from everywhere.

Companies choose Polycom for solutions that allow their workforces to communicate more effectively and productively over distances. Using Polycom unified communications (UC) solutions—telepresence, video, and voice solutions and services—people connect and collaborate with one another from their desktops, meeting rooms, class rooms, and a variety of mobile settings—and from anywhere in the world. In today's economy, our customers wish to cut the time, cost, and carbon emissions associated with gathering the right people in one place to solve problems. Instead of traveling, virtual teams use Polycom solutions to easily and quickly collaborate "face-to-face" wherever they are, which allows them to focus their resources, time, and energy on addressing business challenges.

Collaborating with Polycom solutions has also become a key competitive advantage for leading organizations around the globe. Our customers tell us it makes sense to use Polycom solutions and their existing business applications to communicate and share information in real time over any device and across any network. Polycom's open-standards integration with the leading unified communications (UC) platform vendors makes it possible. Quite simply, it makes good business sense for our customers to rely on the broadest offering of unified communications solutions—from Polycom—so they can improve productivity, reduce their costs, rapidly gain a return on their technology investment—and thrive.