TelePresence over DOCSIS

John T. Chapman, Cisco, jchapman@cisco.com
Harsh Parandekar, Cisco, harsh@cisco.com
Jeffrey Finkelstein, Cox Communications, jeff.finkelstein@cox.com

Abstract

Telepresence is a new technology category that delivers a unique, in person experience for virtual meetings. A telepresence solution integrates advanced visual, audio, and interactive technologies with broadband networking to bring people together from across the campus and around the world.

To date, telepresence has been deployed mostly by large enterprises to enable employees to conduct virtual meetings with fellow employees and customers at remote offices that are linked via a private corporate network.

As telepresence becomes more prevalent, businesses will want the ability to extend the system to locations beyond the corporate network. However, as a real-time, two-way service with high-definition video and spatial audio, telepresence places stringent requirements on the underlying network to provide a high-quality user experience.

This paper describes methods for extending telepresence to locations served by a DOCSIS access network, such as executives' homes. The challenges of delivering telepresence over DOCSIS networks are investigated, and potential solutions for addressing these challenges are proposed.

SERVICE PROVIDER DEPLOYMENT MODELS

Large enterprises were the early adopters of telepresence, which sought to increase productivity and reduce expenses by conducting virtual meetings with an in-person experience. These enterprises typically deployed telepresence over a private corporate network, or secured the appropriate service level agreements (SLAs) from service providers to support their enterprise telepresence systems.

Enterprises are now looking to extend their telepresence systems to more locations, including additional corporate offices, key customer and partner facilities, and executives' home offices. As enterprises continue to grow their telepresence systems, cable operators will increasingly have the opportunity, and the challenge, to support telepresence services on their DOCSIS access networks.

The stringent network requirements to support the real-time, bi-directional, high-definition video and spatial audio of telepresence will likely require cable operators to offer new levels of service to subscribers who want to access corporate telepresence systems via the DOCSIS network.

Although the enterprise will typically bear the cost of the network services required to provide remote access to the enterprise telepresence system, the cable

operator will likely bill the subscriber directly as part of a residential services package.

As telepresence becomes a more prevalent medium for business-to-business communications, enterprises will want to enable their systems to interoperate with their ecosystem partners. Smaller businesses will also want to deploy telepresence, but may not have the resources or know-how to manage their own telepresence system.

Cable operators and other service providers may see an opportunity to offer a managed telepresence service to these customers as part of a commercial services package. This could entail hosting a telepresence call management server (CMS) and managing access control and billing in addition to providing the network services.

As telepresence technology evolves, one can imagine consumer electronics devices incorporating telepresence capabilities for personal use. As this evolution takes place, cable operators will increasingly be able to offer the ultimate visual networking experience to their residential subscribers.

While this paper is focused primarily on enabling telepresence over DOCSIS (TPoD) as an extension of an enterprise telepresence system, the challenges and proposed solutions described herein are also applicable to the other deployment models described above.

THE TELEPRESENCE USER EXPERIENCE

Telepresence is a technology that allows people who are in physically separate locations to communicate with each other as if they were in the same room. Telepresence combines professional video, professional audio, and networking to create a real-time in-person experience.

So what does that really mean in more technical jargon? Video is the main component. Enterprise telepresence systems, as shown in Figure 1 typically use one to three 65 inch plasma monitors operating at 1080p and 30 frames per second. The monitors are set on one side of a desk in such a way that the desk at both ends of the call combines together to create one virtual desk with everyone sitting around it.

The audio uses professional

Figure 1 – Telepresence Endpoint with Three Screens

microphones and speakers. In a larger system, multiple microphones are used, and audio streams are coordinated so that audio from an individual on the left side of the room is played back on the left speaker at the far end. This is known as spatially aware audio.

Even the lighting of the room, one of the most important considerations, borrows from the professional world. Reflective or diffused lighting is recommended as direct lighting causes shadows. Light sources are located in front of the people to light up their faces, not behind them as often happens in the webcam experience.

The lighting is chosen to have the best color temperature (4100K) to make skin tones look good. The camera is a fixed focal length 1080p camera and is calibrated to the room. There are even design guidelines for the exact size of the room, and the color and composition of the walls, floor, and ceiling.

There is a method for conveying slides from a PC. This may be done through picture-in-picture or a separate monitor or projector.

And finally, there is the user interface. Classical video conferencing systems had a complicated user interface that required lots of configuration and usually didn't work. A properly designed telepresence system uses a simple user interface such as an IP phone. For example, a call can be booked in Outlook. Automated provisioning software causes the meeting to show up on the IP phone screen. The call is established by pushing a single button.

In a properly designed telepresence system, there is nothing to adjust. All adjustments have already been done. It just works.

The quality of the experience and the careful placement of monitors, cameras, and microphones create the illusion of an in-person experience.

SYSTEM OVERVIEW

Behind the Scenes

There are other components in a full telepresence system.

In the room, there is a telepresence endpoint that contains the CODEC (Coder-Decoder) and the call management software. On one side, the chassis connects to the monitor, camera, speaker, and microphone. On the other side, the chassis connects to an Ethernet network. If more than one set of monitors is needed, then this system is duplicated and then interconnected via the Ethernet network. After all, it's all IP at this point.

The telepresence endpoint may contain an auxiliary channel that sends and receives audio and video from a personal computer. This is for sharing slides and other audio-visual material.

At a remote location, there is a call manager that manages the telepresence calls and the IP phone. For telepresence systems that use SIP (Session Initiated Protocol) for signaling, the call manager is often implemented with a SIP proxy or an IP PABX.

In order to make multi-party calls, the remote location may also have a telepresence conference bridge.

To permit interoperability between vendors, the remote location may also offer gateways that perform signaling conversion. Once example of a gateway would be to convert between older H.323 systems and newer SIP based systems.

Adapting Telepresence to the Home Office

The home office is not as easily configured as the enterprise environment. The easiest solution for a home office is to use an all-in-one floor standing system and put it on the far side of the desk.

The author of the paper configured a home office telepresence system with discrete components. He put a 40" LCD TV at the end of his six foot desk and sits at the other end for calls. The camera is mounted to the top of the TV. The controller was put into a cabinet with ventilation. The phone and microphone sit on the desk.

The lights in the room were changed to get the right color, and a reflective light source was added above the TV to help the camera. The windows in the room had to get blackout shades so that they would not create back lighting.

To connect back to the enterprise environment, a good quality router with build in VPN (IPsec based) is needed. As will be seen later in the paper, the existence of this router and the VPN

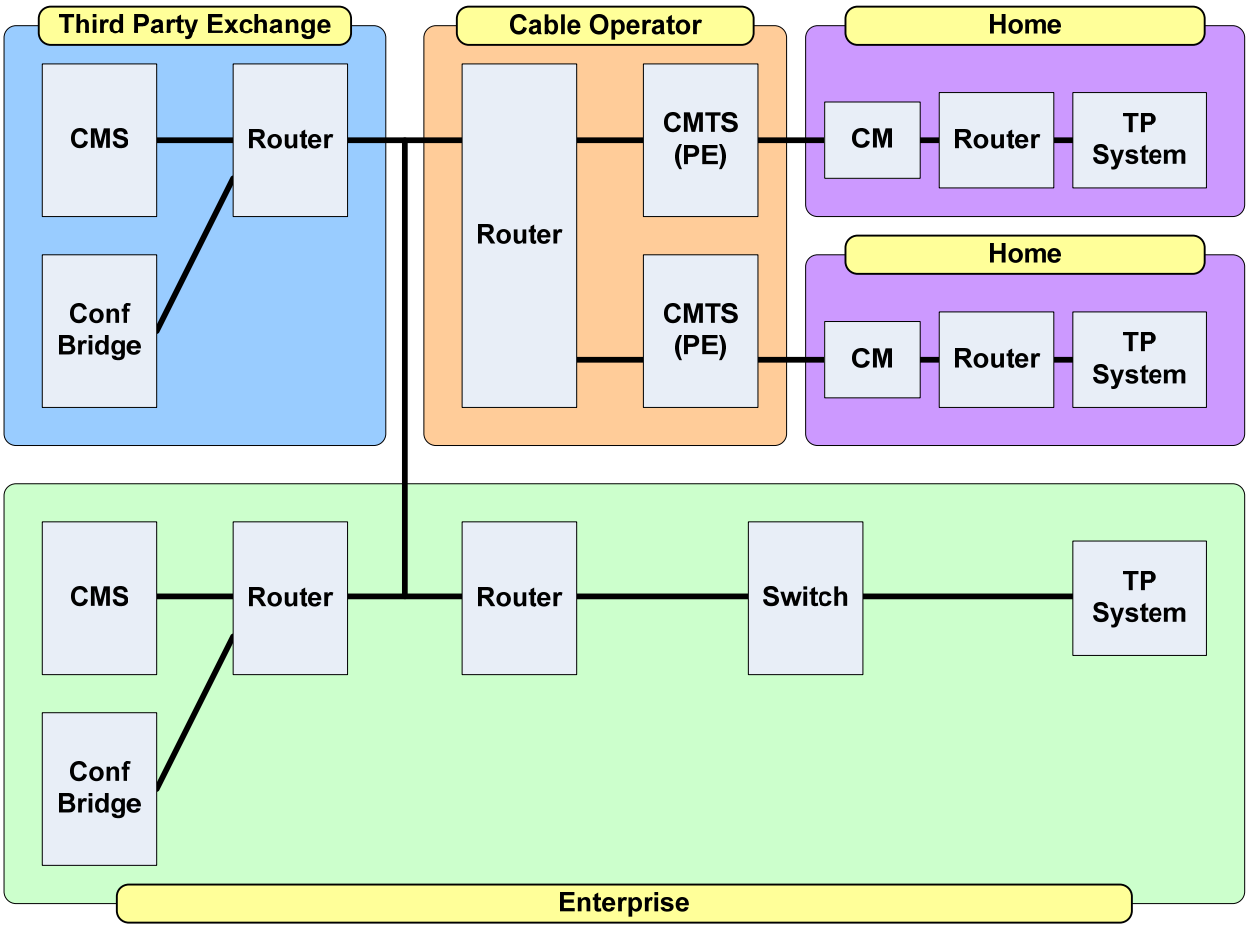

Figure 2 – TPoD End-to-End System

create one of the two major problems that have to be overcome in order to set up QoS on the DOCSIS access network.

End-to-End System

Figure 2 shows an entire TPoD system that spans both the enterprise and the service provider environment.

The enterprise contains one or more telepresence systems. It also contains its own IP PBX system (call manager system) that is used to connect IP phones and to connect the telepresence systems. There is a local telepresence conference bridge as well.

The home office as described earlier contains a telepresence system that sits behind a VPN router. This connects to the cable operator through a local cable modem (CM) and cable modem termination system (CMTS).

For TPoD systems that are deployed for intra-enterprise purposes, all calls are typically routed to the CMS managed by the enterprise. If the enterprise CMS is not compatible with the home office telepresence system, a third-party telepresence service may be used. Third-party service may also be used to enable inter-enterprise calls. The third-party service provider will require secure access to each enterprise. This is significant as it can impact call flow considerations.

HOME NETWORK CONSIDERATIONS

The home network consists of all the networked components in the home and the interconnectivity, including the telepresence endpoint and the DOCSIS CM. The connection between the telepresence system and the DOCSIS CM should be a wired path. Wireless connections are not recommended as they have a higher packet loss rate that can impact real-time video.

In a worst case scenario, there could be three routers in the home network.

1. A home router which aggregates all the traffic from the home network

2. A telepresence router that places telepresence traffic onto a VPN. A third-party service provider may remotely manage this router.

3. An enterprise telecommuter router to that is managed by the enterprise for data and VoIP connectivity. (Ideally, this is the same router as the telepresence router)

The main difference between these three routers is that a different individual or organization manages each router. In the simplest scenario, there is one router that manages both the home network and the telecommuter and telepresence networks.

One connectivity option is to have the home router and telepresence router connected separately to the CM. For this to work, the cable operator has to be able to configure the CM to classify and provide QoS treatment for telepresence traffic. This would require separate IP addresses from the cable operator for each router and NAT in the CM to be disabled.

If the telepresence router is connected through the home router, the home router should be configured to provide priority access to the CM for the telepresence service.

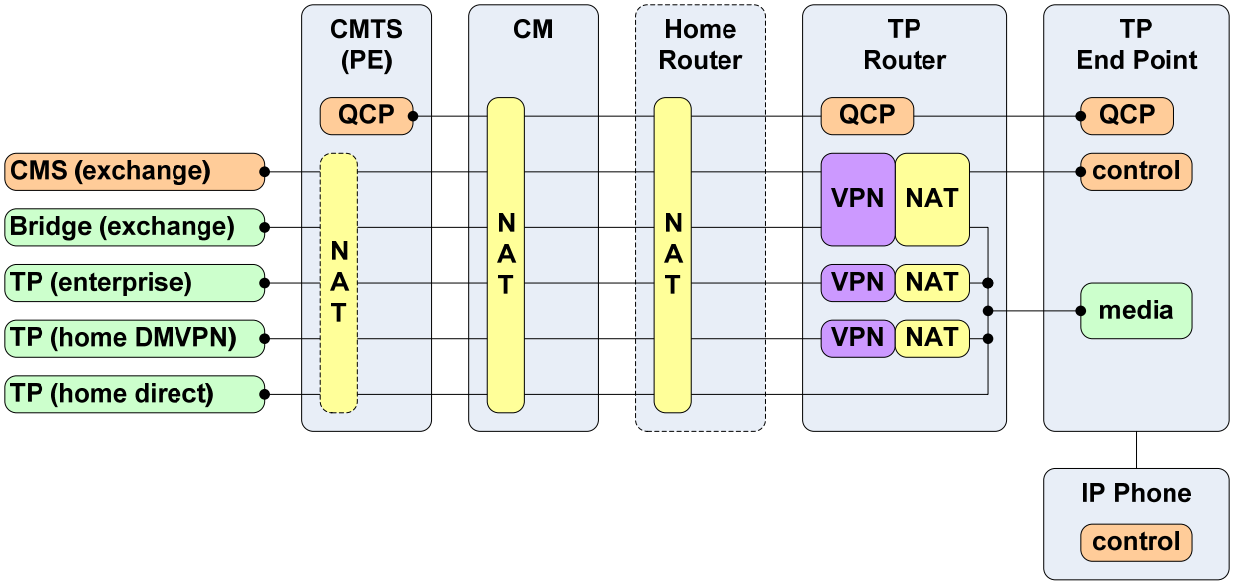

Figure 3 – VPN and NAT Scenario

For best results, the telepresence router should be directly connected to a port on the home router or the CM. This will eliminate any contentions within the home network for traffic.

ACCESS NETWORK CONSIDERATIONS

The core problem that this paper is addressing is how to deliver QoS on the DOCSIS network to support a telepresence call. To do that, the DOCSIS system needs to identify the flow of telepresence packets, separate them from the rest of the packets, and apply specific QoS algorithms.

To understand this issue better, lets look at a packet as it leaves the telepresence system and travels upstream to the CMTS.

There are two major issues with the way that telepresence is connected that prevent easy identification of telepresence packets.

- VPN
- NAT (multiple)

A VPN tunnel is built between home and enterprise routers. The VPN tunnel puts an encrypted IP packet inside another IP packet. Without the ability to read the inner encrypted packet, a SP (and the CMTS) can only use the outer IP header to route between home and the enterprise.

The home and enterprise routers terminate the VPN tunnel, remove the outer header and route only the inner packet onwards. The VPN tunnel is transparent to the telepresence system and the enterprise call manager. Since the outer IP header is only added/removed by the routers terminating the VPN tunnel, the telepresence system and the enterprise call manager only communicate using the inner IP header and have no knowledge of the outer IP address. So who is going to tell the CMTS how to build a classifier?

The situation is complicated further by NAT. Often the port on the router that is

offering a VPN service also runs a local NAT. That means that the telepresence packet had its IP addresses changed prior to being encrypted and encapsulated. This is the first level NAT.

If the VPN router is connected to the home gateway router, the home gateway router often runs a NAT. Now the outside IP address of the VPN tunnel is changed. This is the second level NAT.

Most CMs today come with built in NAT. As the packet travels through the CM, the previously NATed addresses are NATed again. This is the third level NAT.

The packet is now ready to be classified by the DOCSIS side of the CM and then be sent to the CMTS.

When the packet leaves the DOCSIS MAC domain, it may be subject to yet another NAT. It is a network NAT that may get deployed after IPv4 addresses become scarce and before IPv6 can be deployed. Although it does not impact the upstream packet, by changing the packet address yet again, it will impact the packet classifier needed on the corresponding downstream of the CMTS at the far end of the connection.

In summary, the upstream telepresence packet, once it left the safety of the telepresence endpoint is grabbed by the network, NATed, encrypted, encapsulated in a tunnel, and then NATed up to three more times. This is worse treatment than Houdini used to get.

And now the CMTS is supposed to build a classifier for this packet? How is that going to work? Before we discuss our proposed solution, lets look at the second part of the QoS equation that is the traffic characteristics of the telepresence flows.

MEDIA FLOW SPECIFICATIONS

Specifications vary between the various manufacturers of telepresence systems. The following set of tables is intended to be a reference point for a typical telepresence system. Actual performance may vary.

Video

Specification	Description
Image Size	1920 x 1080 or 1280 x 720
Frame Rate - Main Video - Aux Video	 30 fps progressive 5 fps progressive
Encoding	H.264, VBR
Encapsulation	RTP, IPsec
Video Bit Rate - 1080p - 720p	 3 – 4 Mbps average 1 – 2.5 Mbps average
Max Frame Burst	65 KB in 33 ms. (12 - 16 Mbps peak)
Aux bit rate	500 kbps avg 2 Mbps peak
Max Latency	150 ms one way including all delays.
Tolerated Jitter - packet jitter - video jitter	 10 ms 50 ms
Tolerated Packet Loss	0.05%

Table 1 – TPoD Video Specifications

In Table 1, the options for high definition video are 720p or 1080p resolution at 30 frames per second (fps). 1080i does not make sense since there is no television legacy in this system. Smaller screens may be fine with 720p. Aux video from a PC can be sent at a lower frame rate.

Ironically, 720p does not necessarily help out the cable environment. While the home monitor may be okay with 720p, the enterprise monitor is likely to be larger and do better with 1080p. Further, the executive using the TPoD system at home is usually more concerned about being presented well to other participants than about how well the others look to him.

That can lead to the odd case where a 1080p signal is sent upstream from the home camera to the corporate 60" screen, while a 720p signal is send downstream from the corporate camera to the 40" home screen. This runs contrary to the cable system bandwidth that typically has more bandwidth in the downstream than the upstream.

The common codec choice for encoding of high definition (HD) video these days is H.264. Aggressive H.264 codecs can achieve a target video bit rate of 4 Mbps or better. There is a catch. People familiar with MPEG-2 encoded video on demand (VOD) and switched digital video (SDV) systems are used to 3.75 Mbps constant bitrate (CBR) for standard definition (SD) video streams. These streams are encoded such that I-frames do not cause traffic peaks that would exceed the target 3.75 Mbps encode rate.

Maintaining a constant bit rate while encoding high-definition video in a low-latency, interactive system such as telepresence requires substantial video processing capabilities in the encoder. Thus, telepresence systems typically use variable bit rate encoding, and thus streams will see a burst when an I-frame is sent. That burst can typically be 2x to 4x the average bit rate.

The peak rate in Table 1 is calculated over a 33 ms frame interval, even though the burst less if it is spread over a longer time interval if there is room in the system latency budget. In the latest telepresence systems, these bursts can be as infrequent as every 5 minutes or more.

RTP encapsulation is used so the video streams may be multiplexed into common transports. RTP also provides better management of the data path.

The maximum one-way latency goal of 150 ms is the same latency goal that is used for voice calls. The value includes all network and equipment delays.

The jitter goal is also very similar to what is allowable for voice calls since jitter tolerance is achieved through jitter buffers that in turn add latency. Video jitter does get more complicated as there are differences between what is allowable for actual packet jitter caused by packet queuing versus video frame jitter which can be caused by serialization delays on slow links.

Audio

Specification	Description
Encoding	AAC-LD
Raw Bit Rate	64 kbps
Packet Interval	20 ms
Encapsulation	RTP
Max Rx Channels	4
Max Tx Channels	2

Table 2 – TPoD Audio Specifications

Table 2 uses an audio codec of the quality of Advanced Audio Encoding Low Delay (AAC-LD). More compressed audio codecs do not make sense since the video already takes up so much bandwidth.

A home office telepresence system may receive up to four channels of audio from the enterprise system - three channels from the three microphones and one aux channel. Assuming a single screen home system, it could send up to two channels of audio to the enterprise system – one from the microphone and one from the aux channel.

Aggregate Traffic Profile

Specification	Description
1080p	3 – 6 Mbps
720p	2 – 4 Mbps

Table 3 – TPoD Aggregate Specifications

Table 3 sums up the audio and video from previous tables. Table values are for a home office telepresence system. The presumption is up to a three-screen system at the enterprise may be communicating with a one-screen system at the home (so three audio feeds and one video feed).

Auxiliary audio and video may or may not be present. Approximately 20% overhead is added to the raw bit rates to allow for encapsulation and for signaling.

DOCSIS Configuration

The DOCSIS transmission path easily handles bursty traffic as all traffic is rate shaped with a formula that allows for traffic bursts. The common traffic parameters used are:

Downstream and Upstream

- Traffic Priority
- Max Sustained Traffic Rate (R)
- Max Traffic Burst (B)
- Min Reserved Traffic Rate
- Assumed Min Reserved Rate Packet Size.

Downstream Only:

- Downstream Peak Traffic Rate.

DOCSIS data rates are enforced by the CMTS through rate shaping of the service in the downstream, and by metering of grants in the upstream. DOCSIS uses a token bucket based rate limit for service flows. The number of bytes forwarded is limited during any time interval T by Max(T), as described by the expression:

$$Max(T) = T * (R / 8) + B$$

Where:

T = time interval under consideration

R = maximum sustained traffic rate (bits/sec) [7 - C.2.2.5.2]

B = maximum traffic burst (bytes) [7 - C.2.2.5.3]

The DOCSIS specification includes an optional parameter for the downstream called Downstream Peak Traffic Rate:

$$Peak(T) <= T * (P / 8) + 1522$$

Note that DOCSIS does not limit the instantaneous rate of a service flow. Individual packets will always travel at the native rate of the media (~40 Mbps for an Annex B downstream, and ~10 Mbps for a 3.2 MHz, 16 QAM upstream).

The minimum value for B is 1522 (one minimum size Ethernet frame) and the default is 3044 bytes although it can be as high as 20 million byes for the DS and 3

million bytes for the upstream. This is known as “Power Boost” in the cable industry.

As a general rule, the calculation is done based upon the Ethernet frame size and does not include DOCSIS framing (except for upstream concatenated frames).

To accommodate telepresence traffic as described in these tables, the value B should be greater than 65 KB and the value R should be at least 6 Mbps. The downstream peak traffic rate should be 19 Mbps (16 Mbps plus 20% packet overhead).

Different codec configurations would result in different requirements. A system in which telepresence calls are dynamically signaled and setup could optimize the values of the service flow parameters on a per call basis.

Since actual peak traffic is infrequent, admission control can be done on the average data rate numbers.

SOLUTION OVERVIEW

This section provides an overview of three proposed solutions for providing QoS for TPoD, and discusses their pros and cons. The next sections will then respectively focus on each solution and dive into the technical details.

But first, lets define exactly what information is required in any signaling messages that are trying to convey QoS information.

Problem Definition

To implement QoS with a DOCSIS Service Flow, the CMTS must have two important sets of information.

1. Packet Classification
2. Traffic Descriptor

The packet classifier is used at the CMTS in the downstream direction and the CM in the upstream direction to direct a particular flow of packets to a particular QoS queue. A DOCSIS classifier is analogous to an RSVP filterspec. The typical items of interest are:

- Destination IP address
- Source IP address (optional)
- Destination Port address
- Packet Type

The service flow encodings need to know about the size and characteristic of the media flow. The telepresence traffic has three general flows, each with multiple sub-flows. These are:

- One or more video streams
- One of more audio streams
- Signaling to one or more end points

DOCSIS has the toolset to classify each of these flows individually, and even to provide different upstream scheduling and different QoS treatment to each type of flow. But does it make sense to do so? The video flows completely dominate the other flows in a TPoD system. Video requires the most bandwidth and the least latency and jitter. The audio and signaling are along for the ride.

Also, the job of packet classification is difficult due to the existence of the VPN

and NATs. Still, the IP destination address of the IP signaling packets will be different than the IP destination address of the audio and video packets.

Each solution will have to address these issues.

Solution #1: Pre-provisioned DOCSIS

This solution tries to get the most out of the existing DOCSIS specification. The solution is focused on setting the correct parameters in the CM configuration file and on the CMTS.

The advantage of this solution is rapid time to market. It allows early TPoD systems to be deployed to see what the market interest is and how well two-way real-time video works over a DOCSIS network.

The disadvantage of this solution is the lack of visibility into when a telepresence call is set up or torn down. This complicates or eliminates the ability to do admission control. There may also be a lot of manual configuration of addresses rather than the network just figuring things out. This prevents this solution from scaling well.

Thus, this solution is targeted at field trials with limited deployment.

Solution #2: On-Path Reservation

This solution uses RSVP over UDP to communicate between the home telepresence endpoint and the CMTS. In this solution, the call manager is not involved in setting up the bandwidth reservation within the network.

This solution requires the VPN router to filter and forward the RSVP over UDP messages so that they will not be tunneled. The RSVP message has extensions that permit a bi-directional reservation.

The advantage of this approach is that by only involving the local telepresence system and the local CMTS, all DOCSIS QoS reservations can be made. This benefits the deployments in which the enterprise telepresence system is behind a firewall and unable to help out.

It also is useful for scenarios where the cable operator does not own the call manager or where the call manager cannot be easily modified for PCMM. Also, because this solution operates on-path, it can operate through NATs transparently and through a VPN gateway that performs the right signaling processing.

Because only a small number of network elements need to change, this technique can be deployed with a reasonable time to market advantage.

The main disadvantage of this system is network security. The telepresence system is an untrusted entity. Bandwidth reservations may come from that IP address that are not actual telepresence calls.

Solution #3: Off-Path Reservation

This solution looks like a classical PacketCable solution. The call manager communicates through the policy engine to the CMTS to reserve bandwidth for the call.

The advantage of this solution is that the call control is more under the control of the call management system.

The disadvantage is that a sophisticated call management system has

to be in place. Today, enterprise class call managers do not support PacketCable Multimedia (PCMM) interfaces, and service provider class call managers do not support video calls. As PCMM becomes more widely deployed, this solution provides a flexible path to providing bandwidth and QoS on demand services

SOLUTION #1
PRE-PROVISIONED DOCSIS

This solution uses provisions and hooks in the existing DOCSIS protocol in order to ensure that the telepresence traffic gets the required QoS treatment over the access network.

We also need to make sure that the introduction of telepresence units in the field does not disrupt service to other modems deployed in the field. This can be achieved in multiple ways

Option 1: Separate Upstream and Downstream for Telepresence

This option would implement a separate upstream and downstream channel on the CMTS for telepresence service. This would guarantee that service to deployed cable modems remains unchanged and the telepresence traffic is isolated from other traffic in the field. If the home telepresence subscriber is also an existing cable modem user, they would get a second modem that is dedicated for telepresence service.

While this option has severe limitation when it comes to scaling the telepresence service, it's a safe option to start field trials and to roll out the service initially. Nailing certain cable modems to certain downstream channels can be done easily using the downstream frequency parameter in the modem configuration file.

If it's too cumbersome to add the downstream frequency in the modem configuration files for existing subscribers, a simple feature can be implemented in the CMTS to designate specific upstreams and downstreams as "telepresence-only". This could be done with a specific CM Service Type TLV.

During registration, the CMTS would look for this TLV in the cable modem configuration file and, depending on whether the TLV is present or absent, it would move the cable modem to the appropriate downstream.

Since the entire upstream and downstream channel is dedicated to a telepresence setup, the QoS profile for the telepresence modem would simply dedicate the full channel bandwidth as CIR bandwidth for that modem.

Pros

1. Very simple. TPoD should always work since resources are dedicated.

2. TPoD does not interfere with the production network. Good for early field trials.

Cons

1. Inefficient use of HFC plant bandwidth and CMTS ports.

Option 2: Shared Upstream and Downstream

In this option, the telepresence cable modem shares the upstream and downstream channels with other cable

modems connecting to the same fiber node. The modem configuration file for the telepresence modem has to be carefully configured in order to minimize impact on the other cable modems.

For this option, if the telepresence user already has a cable modem for residential internet service, we could consider connecting the telepresence system to the same modem. The simplest way to set up the home network is to connect the telepresence endpoint to the home VPN router, and connect the VPN router directly to the cable modem (with NAT disabled) instead of connecting through a home gateway or NAT device. The VPN router would be statically provisioned with a routable IP address.

This address can be used to define the appropriate upstream and downstream classifier in the modem configuration file to ensure that other non-VPN CPEs in the telepresence user's home that are connected to the shared modem do not disrupt the QoS for the telepresence system. The CMTS would provide QoS for all the traffic generated by the VPN router.

The telepresence endpoint can also be configured to mark outbound (upstream) traffic with a DSCP code point. The VPN router could be configured to preserve these markings when tunneling traffic. The VPN router would do this by copying the DSCP values from the inner packet (telepresence packet) to the outer packet (VPN encapsulation).

The DSCP could then be used in the upstream classifier at the CM to identify telepresence call traffic. Note that once the VPN traffic hits the CMTS, the CMTS typically is configured to rewrite all DSCP values on outbound packets.

A hybrid option is also possible when we use multi-channel DOCSIS 3.0 modems. A subset of channels can be set aside for telepresence service and the remaining channels could be used other traffic to and from the non-telepresence CPEs in the home.

The CMTS would be configured in such a manner that the secondary upstream and downstream service flow created on the CMTS for the telepresence traffic would pick up the appropriate channel set. With this approach, a single modem would be deployed in the telepresence user's home but the telepresence traffic would still be kept isolated from other residential traffic in the cable upstream and downstream.

Once again, CIR reservations could be used to guarantee bandwidth. That would prove inefficient. A more reasonable approach is to use a priority service flow combined with a scheduling algorithm such as real time polling service (RTPS).

There is no dynamic admission control in this approach. Thus, the number of TPoD systems provisioned per channel should be managed carefully.

Pros

1. The pre-provisioned bandwidth functionality based on classifiers and service flows defined in the cable modem configuration file is compliant with standard DOCSIS behavior and hence is supported by a current CMTS.

2. Easy to provision and manage the telepresence service through the modem configuration file

3. If priority service flows are used for TPoD sessions, the bandwidth can be over-subscribed.

Cons

1. Telepresence service has strict QoS requirements. If static CIR provisioning is used to permanently reserve bandwidth for TPoD sessions, that bandwidth is not available to other CIR services such as VoIP or other TPoD calls. The bandwidth is, however, available to best effort flows when not in use by the TPoD session.

2. No admission control. Thus a newly added TPoD session could interfere with an existing TPoD session if there is not enough bandwidth on the channel.

SOLUTION #2
ON-PATH RESERVATION

Resource Reservation Protocol (RSVP) is the standard "on-path" bandwidth reservation protocol to ensure that the right QoS treatment is setup along the path of a media flow through the network.

The RSVP message contains a "filter spec" and a "session". Together, these specs allows a node in the network to identify the flow – source IP, destination IP, protocol as well as source and destination port if the traffic is UDP/TCP. The RSVP message also contains a "Flow spec" that define the traffic parameters (e.g. bandwidth) of the reservation.

There are several challenges with dynamic bandwidth reservation such as NAT traversal and IPsec based VPN tunnels. These are described in detail in the next few sections.

NAT Traversal:

Bandwidth reservation is complicated when RSVP signaling has to traverse one or more NATs. First, NAT devices may not handle raw RSVP packets properly. Hence the telepresence client will use the RSVP over UDP packet format that encapsulates the RSVP message inside a UDP packet destined to a well-known UDP ports, 1698 and 1699 for RSVP encapsulation.

The telepresence client would use the destination address of the remote telepresence client for these RSVP over UDP messages. The CMTS would have to intercept these RSVP over UDP packets on the upstream by filtering the well-know RSVP encapsulation UDP ports.

Another issue with NAT traversal is that the client requesting the bandwidth would use the "inside" NAT address and port by default, while the actual traffic that traverses across the network would have the external IP address and port.

In order to avoid this problem, the CMTS would have to use the source address of the actual packet instead of the local endpoint address specified in the filter spec of the RSVP message while creating the upstream and downstream classifiers.

There are several types of NAT algorithms. The basic NAT algorithm known as one-to-one NAT replaces a packet's source IP address and source port number with a new external IP address and external source port. There is only one internal to external mapping of the source

IP address and port for all IP destination addresses and ports.

Another NAT algorithm is known a symmetrical NAT. Symmetrical NAT changes the external mapping of the source IP address and source port for each unique instance of the destination IP address and port.

With symmetric NAT, the local port for the media flow cannot be predicted or determined easily. In order for the solution to work even with symmetric NATs, the telepresence endpoint can set a wildcard for the source port in the RSVP filter spec.

Hence, the CMTS will program the DOCSIS classifier without the local endpoint port. The upstream will match based only on the source and destination IP address, protocol type and destination port. The downstream would match based only on the source and destination IP address, protocol type and source port.

Authorization of the Telepresence Client:

As the CMTS is going to process resource reservation requests from the home, it has to be ensured that the solution is not susceptible to Denial of Service attacks from malicious cable modem users. The standard mechanism to authorize a user requesting bandwidth using RSVP is by using Pull-COPS /AAA (Authentication, Authorization and Accounting) where the CMTS would offload the authorization decision to a COPS /AAA server.

Pull-COPS /AAA authorization may not be deployed in a specific cable network. Hence another option is to add a vendor-specific TLV in the modem configuration file to indicate to the CMTS that it can accept RSVP requests from clients connected to that modem.

Teleworker VPN:

As shown in Figure 4, in most cases the telepresence endpoint would be connected to a home VPN router. In the VPN case, the CMTS will not be able to snoop packets sent from the telepresence client. This is because the VPN tunnel would extend from the home VPN router to the enterprise VPN router and the CMTS would be in the middle. By default, RSVP over UDP packets sent inside the IPsec VPN tunnel would also get encrypted along with the actual telepresence flow.

This enterprise VPN challenge can be solved by adding a new feature to the home VPN router. The VPN router will have to filter for RSVP over UDP packets and will forward the RSVP messages unencrypted outside the VPN tunnel in the upstream direction.

While this change in the VPN router would ensure that the CMTS can snoop the RSVP over UDP message, there is a further challenge in the VPN scenario. The source and destination address and port information in the RSVP message would not be useful for the CMTS, since the packets of the telepresence flow would still be encrypted within the IPsec VPN tunnel. (and therefore encapsulated with packets whose IP addresses are those of the VPN gateways).

To solve this problem the CMTS will instead have to store the source and destination IP address in the IP header of the packet carrying the RSVP over UDP message, and use those to create the classifier. When the VPN router forwards the RSVP over UDP message to the CMTS, it uses the same source and

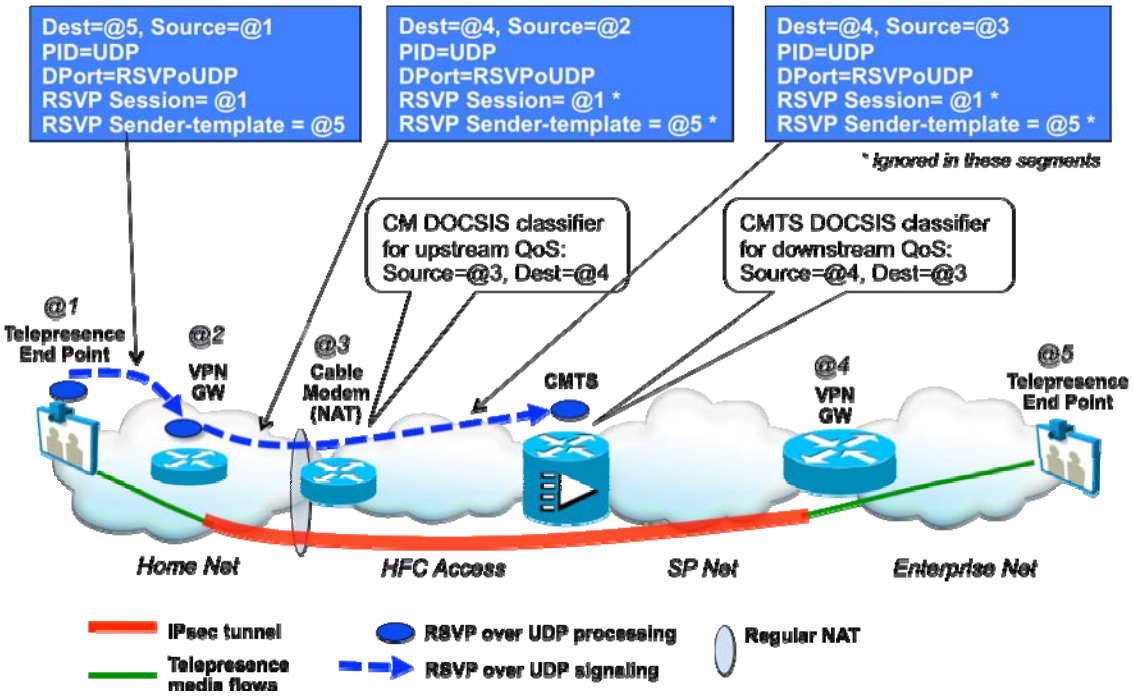

Figure 4 - RSVP over UDP Packet Flow

destination address that it will use to send the corresponding media packets within the IPsec VPN tunnel.

Hence, the CMTS can create a DOCSIS classifier that matches all the traffic within the VPN tunnel. The VPN router would have to insert a flag in the RSVP message that the CMTS would interpret as "don't use the address and port information within the RSVP message. Instead, create the DOCSIS classifier using the source and destination address of the RSVP over UDP packet".

In Figure 4, the telepresence endpoint sends an RSVP over UDP packet with a destination address of @5 (the far end telepresence endpoint) and a source address of @1. The VPN router intercepts this packet, sets the above mentioned flag in the RSVP message, and transmits the packet unencrypted using a destination address of @4 (the far end VPN router) and a source address of @2, which represent the near end of the VPN tunnel.

Figure 4 also shows an additional NAT operation that is performed by the cable modem which changes the source address of the packet to @3 and keeps the destination address unchanged as @4. The CMTS filters the RSVP over UDP packet and uses the destination address @4 and source address @3 to create the upstream and downstream classifiers to provide appropriate QoS to the bi-directional telepresence media.

Bandwidth Reservation in the Downstream Direction:

Typically RSVP messages describe unidirectional flows. If bandwidth reservation is required in both directions, it is expected that each endpoint sends RSVP requests for its own transmission in the forward direction.

Telepresence traffic is also bi-directional, but it's easier to provision the network such that the CMTS is expected to process RSVP messages only from the local telepresence endpoint. There are several issues with trying to support RSVP messages for downstream traffic from the remote telepresence endpoint

1. In case of the teleworker VPN case, traffic from the remote telepresence endpoint, including RSVP messages, would be sent encrypted inside the VPN tunnel and cannot be seen by the CMTS

2. Even for the non-VPN case, RSVP messages coming from the remote telepresence endpoint may be filtered out somewhere in the network – especially if that endpoint belongs to a different service provider.

3. The CMTS would be open to Denial of Service attacks from the Internet. It's easier to authorize the local telepresence endpoint requesting QoS than it would be to authorize a remote telepresence endpoint that the operator has no control over.

In order to avoid these issues, bidirectional RSVP can be used as defined in the PacketCable specification. This would enable the local telepresence endpoint to also request bandwidth from the CMTS for the reverse flow in the downstream direction.

Step-Wise Solution

1. As part of the call setup, the video resolution is negotiated. The local telepresence endpoint calculates the bandwidth to match this resolution and sends an RSVP over UDP message that describes the upstream and downstream bandwidth requirements. The source port is wildcarded in the filter spec to avoid issues with symmetric NATs.

2. If a VPN router is present, then the RSVP over UDP packet is intercepted by the VPN router, which changes the source and destination IP address of the packet to be the same as the source and destination of the IPsec tunnel. The VPN router also sets a flag in the RSVP message requesting the CMTS to "ignore the filter spec" within the RSVP over UDP message

3. When the CMTS intercepts the RSVP over UDP message, it ensures that the message is being received from a client behind an authorized modem. If authorization fails, the CMTS would drop the packet. If it succeeds it would continue to process the RSVP request.

4. If the "ignore the filter spec" flag is not set, the CMTS creates an upstream classifier based on source, destination IP address, protocol, and destination port. It also creates a downstream classifier based on source and destination IP address, protocol and source port.

   To avoid problems with symmetric NATs, in both classifiers the local telepresence endpoint UDP ports are not used for classification. If the "ignore the filter spec" flag is set, then the CMTS creates upstream and downstream classifiers based only on the source and destination IP address of the received RSVP over UDP packet.

5. The CMTS tries to admit the upstream and downstream service flows based on the TSPEC. If it fails, the CMTS

sends back an RSVP call admission control (CAC) reject notification. When the telepresence endpoint receives this notification, it can try to negotiate a lower video resolution and motion with the remote endpoint and send a new RSVP request, thereby restarting operation at step 2 above.

6. Once the local telepresence endpoint succeeds in reserving the required bandwidth, it completes the SIP call setup and telepresence media begins to flow between the two telepresence endpoints.

7. When the telepresence call is done, the local telepresence endpoint sends an RSVP teardown message which is again intercepted by the CMTS. When the CMTS receives this message, it deletes the upstream and downstream service flows created for the telepresence call.

8. RSVP uses soft state to manage the QoS reservation in the network. This soft state has to be periodically refreshed by the local telepresence endpoint by sending RSVP message periodically. If the CMTS does not receive these periodic RSVP messages, it will eventually timeout the bandwidth reservation and will tear down the service flows created for the RSVP request.

   This behavior ensures that bandwidth is reclaimed even in the case where the TP call is not gracefully torn down. An example where this is need would be if the local telepresence endpoint is powered down before it can end the telepresence call and send the RSVP teardown message.

Pros

1. This is an on-path bandwidth reservation solution that can be deployed independent of various flavors of off-path solutions (PacketCable, PCMM etc) that the operators may be using for existing voice services.

2. Complex NAT and VPN scenarios can be supported with this solution.

3. An on-path RSVP over UDP solution for TPoD can simultaneously exist with an off-path PCMM system for VoIP.

Cons

1. Even though the reference PCMM architecture does include support for RSVP clients that request QoS resource directly from the CMTS, this approach has not been fully defined in the standards specification. Hence initial implementations of this solution would be vendor specific or need to be further standardized.

SOLUTION #3
OFF-PATH RESERVATION

Nailing down bandwidth as recommended in solutions #1 and #2 will be a challenge for many cable operators. An alternative approach is to use the DQoS infrastructure as defined in PacketCable MultiMedia to dynamically create and tear down service flows as needed.

Rather than statically provisioning the QoS within the access network, the first initial SIP Invite from the Telepresence client to the call manager would trigger a secure sequence of events that would

allow the CMTS to dynamically provision the QoS while still meeting the goals of solutions #1 and #2. It is envisioned that there would be no need for dedicated bandwidth and that using a shared model would provide sufficient capacity.

The call manager could be located within the Service Provider domain or be provided by a third party telepresence service. The general call flow remains the same with the addition of authentication. The use of secure protocols with authenticated is a requirement as the QoS requests will be triggered by parties external to the cable operators network.

The basic flow for requesting QOS from a hosted call manager external to the cable network is as follows.

1. Local TelePresence (CPE) issues SIP Invite to call manager.

2. Call Manager sends SIP based QOS request to the Service Edge Proxy using a secure transport protocol such as SIPS, TLS or HTTPS

3. Edge Proxy validates request and sends the request to application manager.

4. Application Manager translates SIP request to PCMM (COPS) and sends to Policy Manager

5. Policy Manager validates request, determines resources needed and sends PCMM gate-set to CMTS

6. CMTS determines resource availability, creates gate and communicates with cable modem using DSx (Dynamic Service Flow) messaging

7. Cable Modem sets up service flow

8. Telepresence media flows flow bi-directionally with the proper QoS.

Multiple checkpoints are used to maintain security:

- the 3rd party call manager must have their specific certification signed by the site they are contacting,

- the edge proxy must have details on the contacting site,

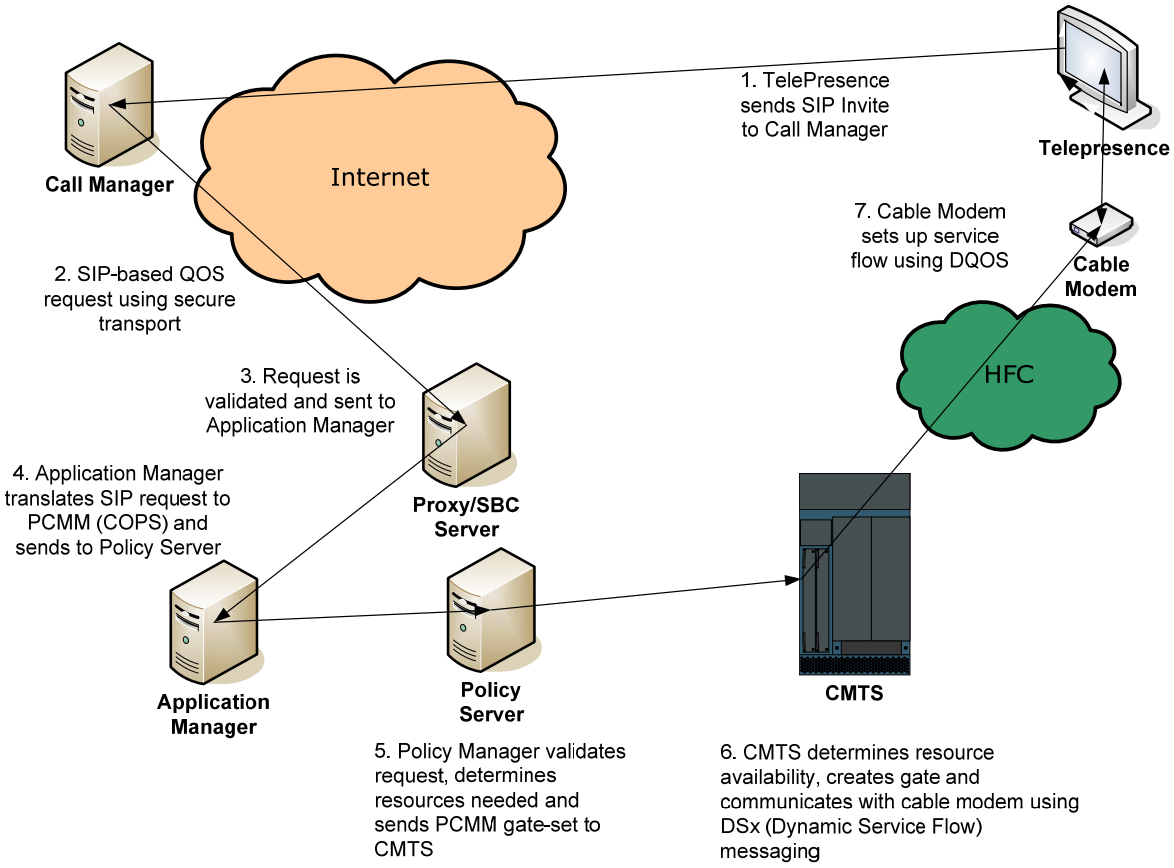

Figure 5 - TPoD PCMM System

- the application manager must have knowledge of the edge proxy and allow requests from that device,

- the policy manager must have policies that support the requests being made, and

- the CMTS must allow requests from the policy manager.

The actual flows created are equivalent to the ones created in the on-path reservation system, but rather than being static flows they are created dynamically using PCMM [4][5].

The addressing challenges from the home NAT and VPN still exist with this call setup scenario. The policy messages that arrive at the CMTS have to contain the actual IP addresses that the CMTS will see, as opposed to the IP addresses that the telepresence system uses. This is a non-trivial problem and will be the subject of additional research.

Pros

1. Uses PCMM for on demand bandwidth management. PCMM is a well-defined specification and CMTS products on the market support PCMM.

2. Does not require dedicated bandwidth

3. Will not impact the quality of voice calls

4. Secure protocols with authentication

Cons

1. Requires the operator to deploy a PCMM infrastructure.

2. Multiple points to manage increases complexity.

3. Introduces potential single points of failure.

SUMMARY

Telepresence is a new technology for conducting virtual meetings with a real-time, in-person experience. As the adoption of this new technology accelerates, users will increasingly want to communicate via telepresence in locations served by DOCSIS networks, such as executives' home offices. Cable operators will have an opportunity to offer advanced network services to support customer-owned telepresence systems, and could also offer a managed telepresence service to customers without a complete in-house telepresence system.

This paper described the challenges in supporting telepresence over DOCSIS, particularly in the areas of dynamic bandwidth reservation and quality of service. Three solutions were proposed and assessed for viability in TPoD deployments.

Pre-provisioning the DOCSIS network for telepresence services can be a suitable approach for field trials and limited deployment, but not for large-scale deployment due to the inefficient use of DOCSIS network resources.

On-path reservation enables the home office telepresence system to request DOCSIS network resources directly from the CMTS, thereby eliminating the need for integration with an external call management system. However, additional authentication mechanisms may be necessary since the TPoD end point is considered to be untrusted.

Off-path reservation enables the telepresence system to manage the reservation of end-to-end network resources, but requires further integration of DOCSIS network resource management functions with the telepresence call management system.

With further investigation and development, the on-path and off-path reservation approaches proposed can be viable solutions for large-scale TPoD systems.

ACRONYMS AND ABBREVIATIONS

AAA	Authentication, Authorization and Accounting
AAC-LD	Advanced Audio Encoding Low Delay
CAC	Call Admission Control

CIR	Committed Information Rate
CM	Cable Modem
CMS	Call Management Server
CMTS	Cable Modem Termination System
CODEC	Coder – Decoder
COPS	Common Open Policy Service
CPE	Customer Premise Equipment
DOCSIS	Data Over Cable System Interface Specification
DSCP	Differentiated Services Code Point
DSx	Dynamic Service Flow
HD	High Definition
HTTPS	Secure HTTP
IP	Internet Protocol
IPsec	IP Security
MAC	Media Access Control
NAT	Network Address Translation
PABX	Private Automated Branch Exchange
PCMM	PacketCable Multimedia
QoS	Quality of Service
RSVP	Resource Reservation Protocol
RTP	Real Time Protocol
SD	Standard Definition
SDV	Switched Digital Video
SIP	Session Initiated Protocol
SIPS	Secure SIP
SLA	Service Level Agreement
TLS	Transparent LAN Services
TLV	Type Length Variable
TPoD	Telepresence over DOCSIS
TSPEC	Transmission Specification
UDP	User Datagram Protocol
VOD	Video on Demand
VPN	Virtual Private Network

REFERENCES

1. Dierks, T. and E. Rescorla, RFC5246, The Transport Layer Security Protocol Version 1.2, http://tools.ietf.org/html/rfc5246

2. Freesoft.org, SSL V3.0 Specification, http://www.freesoft.org/CIE/Topics/ssl-draft/3-SPEC.HTM

3. Wagner, David, and Bruce Schneier, Analysis of the SSL 3.0 Protocol , The Second USENIX Workshop on Electronic Commerce Proceedings, USENIX Press, November 1996

4. Cablelabs, PacketCable™ Dynamic Quality-of-Service Specification, I03, 2007 http://www.cablelabs.com/specifications/PKT-SP-DQOS1.5-I03-070412.pdf

5. Cablelabs, PacketCable™ Network-Based Call Signaling Protocol Specification, I03, 2007, http://www.cablelabs.com/specifications/PKT-SP-NCS1.5-I03-070412.pdf

6. Cisco, Telepresence Room Requirements, 2009, http://www.cisco.com/en/US/solutions/ns669/networking_solutions_products_genericcontent0900aecd80554cb2.html

7. Cablelabs, MAC and Upper Layer Protocols Interface Specifications, DOCSIS, I09, 2009, http://www.cablelabs.com/specifications/CM-SP-MULPIv3_0-I09-090121.pdf

8. Cisco, Telepresence Network Systems 2.0 Design Guide, August 1, 2008, http://www.cisco.com/en/US/docs/solutions/Enterprise/Video/TelePresence_Network_Systems_2.0_DG.pdf