Cisco Unified Communications Voice over Spoke-to-Spoke DMVPN Test Results and Recommendations

This document describes the interoperability of the Cisco Dynamic Multipoint VPN (DMVPN) solution with voice over IP (VoIP), which is part of the Cisco Unified Communications solution. This document provides some specific results and recommendations resulting from testing conducted by the Cisco Global Government Systems Group (GGSG) Government Systems Engineering team. The testing was specifically intended to validate the Unified Communications solution in conjunction with a specific spoke-to-spoke DMVPN solution already deployed by a service provider on behalf of a U.S. government agency; however, many of the concepts discussed are applicable to any deployment of VoIP with spoke-to-spoke DMVPN.

The following detailed design guides for the solutions tested already exist:

- Cisco Unified Communications SRND Based on Cisco Unified CallManager 4.x— http://www.cisco.com/en/US/docs/voice_ip_comm/cucm/srnd/4x/uc4_2.html
- Dynamic Multipoint VPN (DMVPN) Design Guide— http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/DMVPDG.html
- Voice and Video Enabled IPsec (V3PN) SRND— http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/V3PN_SRND/V3PN_SRND.html

This document extends the scope of the existing documents with a focus on deploying a Cisco Unified Communications solution together with a spoke-to-spoke DMVPN network. The purpose is to provide additional solution design guidance, including predicting performance of particular platforms as spoke routers and configuring the solution to ensure high voice quality.

The data in this document is the result of testing in the GGSG test lab with a sufficient scale of the DMVPN and Unified Communications solutions to mimic real-world behavior. More information on the test methodology and results is provided in References, page 69.

Americas Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

Overview

Cisco Unified Communications products are designed and tested with various types of WANs in mind. For example, Call Admission Control (CAC) takes into account the possibility of limited WAN bandwidth over which to transport media. Such features work in conjunction with QoS policies on WAN gateway routers to ensure that available media and call control signaling bandwidth is available.

The standard Unified Communications configurations apply to a deployment using the DMVPN network for media transport. This includes configuration of the Cisco Unified CallManager and Unity servers, gateways, IP phones, and LAN infrastructure.

To account for DMVPN network topology attributes and behavior, some additional requirements and best practices apply. These are discussed throughout this document.

Issues to Address with Voice Over Spoke-to-Spoke DMVPN

There are the following main concerns about the interoperability of Cisco Unified Communications with spoke-to-spoke DMVPN networks:

- Caveats exist related to voice quality because of differences in latency between the spoke-hub-spoke path and the spoke-to-spoke path when the spoke-to-spoke tunnel is initially established.
- The point-to-multipoint nature of the tunnel interfaces on the spoke routers creates the possibility that a spoke router access link (or associated low-latency queue) could be overwhelmed with traffic from multiple sources.

See References, page 69 for more information on these issues.

These issues are addressed by using the following voice solution attributes:

- Call admission control (location-based CAC)
- Quality of service (QoS)
- Automated alternate routing (AAR) for calls encountering bandwidth/CAC limits
- Survivable Remote Site Telephony (SRST) and Public Switched Telephone Network (PSTN) gateway at each remote site
- Adequate WAN performance and capacity, latency, and bandwidth

This remainder of this document addresses these issues and provides additional best practices information.

Limitations of this Document

Significant variations are possible when designing a DMVPN network and the underlying WAN. The test approach, results, measurements, and conclusions in this document are most applicable to the DMVPN network used by the target government customers with a single-cluster Cisco Unified Communications topology. Although the strategies and best practices discussed are generally applicable, the detailed characteristics of other DMVPN deployments need to be considered. Cisco did not examine other variations in topology and scale for this document.

Solution Description

This document is focused on a government network deploying voice over IP using an existing DMVPN network for media and signaling transport between the main office and branch offices. This deployment uses the following solutions in combination:

- Cisco Unified Communications (based on the Cisco Unified CallManager version 4.x)
- Voice and Video Enabled IPsec VPN (V3PN)
- DMVPN

The existing DMVPN network of the customer has a dual-hub, dual DMVPN cloud topology with a dual-tier headend architecture; and approximately 2000 spoke routers connected via the MPLS network of the WAN service provider. The same service provider operates the DMVPN network on behalf of the government customers. This includes hub complex routers at centralized data centers and spoke routers at each customer site. The second DMVPN cloud uses a different service provider network.

The DMVPN network is shared among multiple agencies in a U.S. government department. The DMVPN solution is configured to provide spoke-to-spoke tunnels between any two spoke routers. The DMVPN network has a single level of spokes subtending the hub level.

V3PN QoS recommendations were followed when deploying the DMVPN components, including Low-Latency Queuing (LLQ) and Class-Based Weighted Fair Queuing (CBWFQ) at the interfaces between the spoke routers and the service provider network. The service provider network maintains the Differentiated Services Code Point (DSCP) information in each packet as the packet traverses the WAN. Enhanced Interior Gateway Protocol (EIGRP) distributes routes to each spoke router on the DMVPN cloud, and Next Hop Routing Protocol (NHRP) resolves the DMVPN subnet (tunnel) address to the non-broadcast multiple access (NBMA) IP address of the destination spoke router. Figure 1 shows the general DMVPN topology.

Hub-and-spoke DMVPN topologies were not considered because the target customers have access only to a spoke-to-spoke DMVPN network.

Figure 1 DMVPN Topology

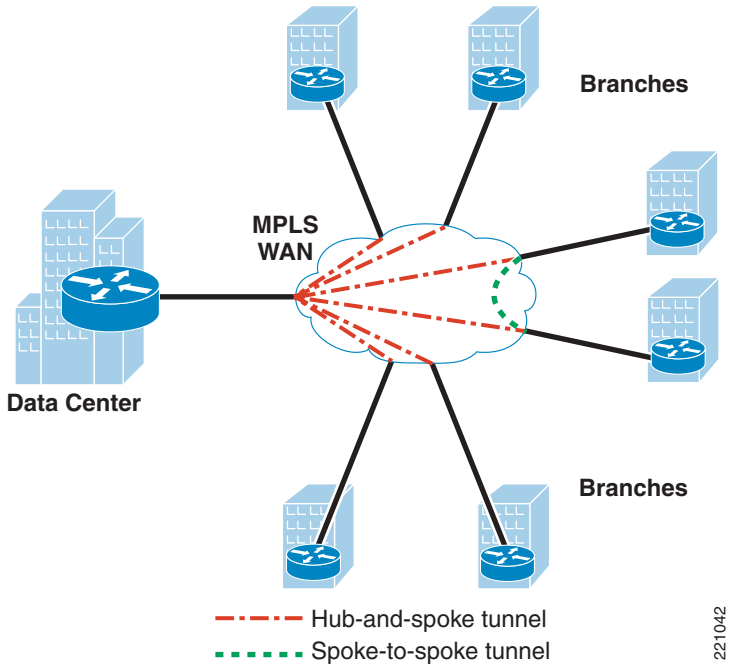

The Cisco Unified Communications solution to be deployed uses the "multisite WAN with centralized call control" topology described in the Cisco Unified Communications SRND based on Cisco Unified CallManager 4.x. The scalability needs of the target customers range from a single main site with 400 users and 3 remote sites with 200 users each (total 1000 users) to a main site with 5000 users and up to 100 remote sites with up to 250 users each (total 30,000 users). The testing did not include clustering over the WAN or multiple clusters with intercluster trunks because these topologies are not required for the deployment scale. Specifically, clustering over the WAN is an option for branch office sites requiring more lines than supported by SRST and require multiple-cluster topologies when the number of lines is more than the capacity of a single cluster or the customer has more large branch office sites than supported by the clustering-over-the-WAN topology. Figure 2 shows the multisite WAN with centralized call control topology integrated with a DMVPN WAN.

Each agency deploys a separate Cisco Unified Communications solution using PSTN connectivity between agencies.

Figure 2 Multisite WAN with Centralized Call Control Topology

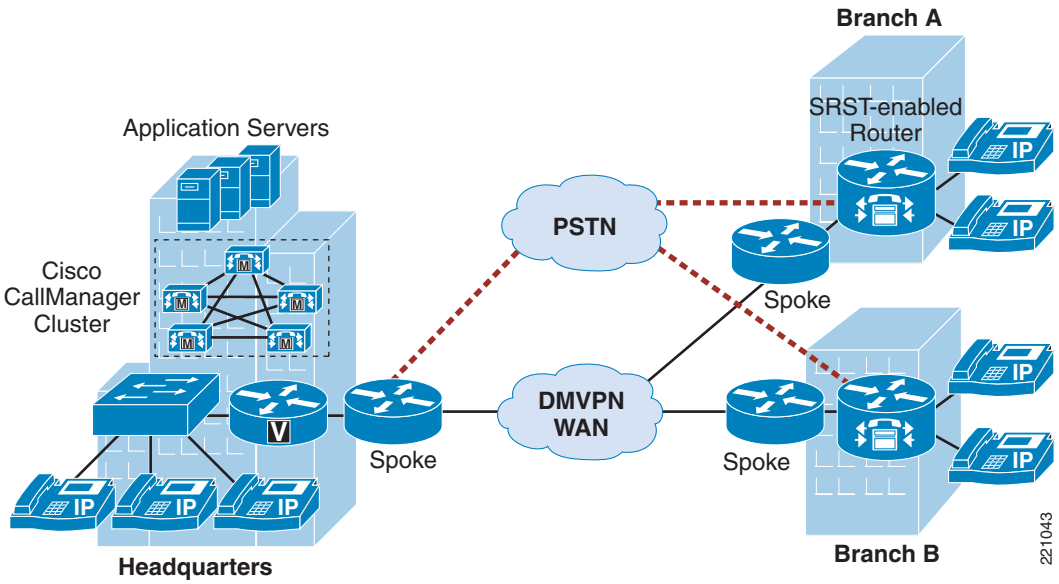

Each remote site is equipped with a PSTN gateway router with SRST capability. The new systems are replacing legacy PBX systems or upgrading standalone Cisco Unified CallManager clusters or multisite WAN with centralized call control that are currently using the PSTN for intersite traffic.

Notice that the main site connects to the DMVPN as a spoke, so communication between the main site and any remote site (or between remote sites) is via a spoke-to-spoke tunnel. The permanent tunnel to the hub from each spoke router is normally only briefly used for voice traffic while the spoke-to-spoke tunnel is first being established. The details of the DMVPN and WAN topologies are transparent to the voice solution except for their impact on latency, jitter, and packet loss.

Co-location of Voice Main Site and DMVPN Hub

This test effort did not examine the topology where the Cisco Unified CallManager cluster is located behind a DMVPN because the target customers have access only to the DMVPN via the spoke routers. This is a characteristic of this particular DMVPN network that does not apply to all DMVPN deployments. If the Cisco Unified CallManager cluster is placed behind a DMVPN hub, the voice requirements would be similar because dynamic spoke-to-spoke tunnels would still be generated between branch office sites. The main difference is the nature of the connections between the branch offices and the main site. These would be permanent static tunnels instead of dynamic tunnels kept active by call control keepalive traffic.

Benefits and Drawbacks of the Solution

Spoke-to-spoke DMVPN provides clear benefits for Cisco Unified Communications compared with a hub-and-spoke topology. Spoke-to-spoke tunnels can provide a reduction in end-to-end latency by reducing the number of WAN hops and decryption/encryption stages. In addition, DMVPN offers a simplified means of configuring the equivalent of a full mesh of point-to-point tunnels without the associated administrative and operational overhead. The use of spoke-to-spoke tunnels also reduces traffic on the hub, permitting bandwidth and processing capacity savings.

Spoke-to-spoke DMVPN networks, however, have the following caveats associated with them:

- Spoke-to-spoke tunnels are not as resilient to some forms of failure as spoke-to-hub tunnels because there is no routing protocol running through the tunnel.
- A spoke-to-spoke tunnel may take a path through the underlying WAN that is more congested than the spoke-hub-spoke path tunnel.
- Without careful planning, spoke routers may become overrun with incoming traffic from multiple remote spokes, resulting in degraded voice quality.
- The transition of Real-Time Protocol (RTP) packet routing from the spoke-hub-spoke path to the spoke-to-spoke path can create a momentary audio distortion under certain circumstances.

Influence of DMVPN on Latency, Jitter, and Packet Loss

As with any WAN-based transport for voice, the voice quality of intersite calls over a DMVPN network is impacted by latency, jitter, and packet loss. The following sections describe in more detail how the DMVPN solution influences these factors.

Latency

The following factors can contribute to latency on the DMVPN network:

- Encryption/decryption delay
- Serialization delay
- Routing/switching/queuing delay
- WAN propagation time
- Delays because of DMVPN spoke and hub congestion

A WAN connection between sites has some latency associated with normal operation. DMVPN induces an additional small amount of latency from encrypting and decrypting packets. This was measured as an average of 2 ms on the spoke-to-spoke path (a single encryption and decryption). Other sources of WAN latency are not unique to DMVPN. Latency is higher via the spoke-hub-spoke path because of additional hops, the possibility of additional propagation time, additional decryption/encryption, and the processing overhead on the hub router(s).

Jitter

The amount of jitter (delay variation between packets) is influenced by the following:

- Rerouting packets along a different path
- Congestion/queuing
- Other network conditions

One factor that introduces a transient jitter into a VoIP stream is a change in the path of the RTP stream from spoke-hub-spoke to direct spoke-to-spoke. If two spokes have an established spoke-to-spoke tunnel before initiation of the RTP stream, there is no cut-through issue. The spoke-to-spoke tunnel may have been initiated by a previous VoIP call or some data traffic between the two spokes.

Note IP phones on a spoke regularly exchange keepalive messages with the Cisco Unified CallManager servers on a different spoke. These keepalives keep the tunnel between the spokes established. Any other data traffic (for example, IP SLA probes, network management, and so on) between the two sites can also keep the tunnel established.

Spoke-to-spoke tunnels are dynamically created as the need for them arises. This process takes a relatively small amount of time (less than one second), but the sequence of events has the potential of momentarily impacting voice quality. While establishing the spoke-to-spoke tunnel, the network routes the packets via the spoke-hub-spoke path. This requires a packet to traverse the WAN twice (once from the originating spoke to the hub and another from the hub to the terminating spoke) and to be decrypted/encrypted at the hub, which results in additional latency compared with the spoke-to-spoke path.

Immediately after the spoke-to-spoke tunnel becomes available, new voice packets are transmitted along this new path. This scenario can cause out-of-sequence packets and momentary high inter-arrival jitter at the terminating spoke router as packets arrive from both paths. The jitter buffer algorithm on the terminating VoIP device attempts to account for the jitter but may sometimes be unable to conceal the discontinuity without some audible distortion. A user might describe the sounds as a glitch, gap, or splice, or the sound might not be noticed at all.

Note This effect does not apply to calls traversing an established tunnel between the two spokes.

There are jitter two scenarios caused by the path switch to the spoke-to-spoke tunnel. The usual case (see Figure 3) is when the spoke-hub-spoke path has more latency than the spoke-to-spoke path. When traffic from the spoke-to-spoke tunnel begins arriving at the destination endpoint, it overlaps with an earlier portion of the media stream that traversed the hub path. The endpoint sees the out-of-sequence packets and reorders them before playing them out; or, if the jitter buffer is too full to accommodate the new packets, discards the older packets and retains the newer packets for playout. If the latency difference is significant enough (greater than 100ms), a user may notice a portion of a syllable missing (a “splice”).

Figure 3 Jitter and Out-Of-Sequence Packets Because of Path Switch—Spoke-To-Spoke Path Faster

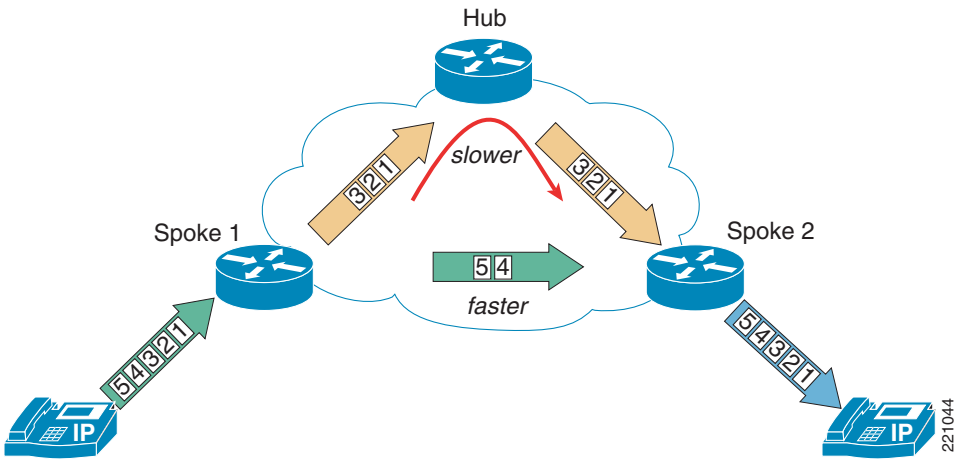

Figure 4 shows the other case. If the spoke-to-spoke path has more latency than the spoke-hub-spoke path, the terminating endpoint detects an unexpected increase in delay (jitter) between the last packet traversing the spoke-hub-spoke path and the first packet traversing the spoke-to-spoke path. If the latency difference is significant enough, a user may notice a period of silence or white noise within the speech.

Figure 4 Jitter Because of Path Switch—Hub Path Faster

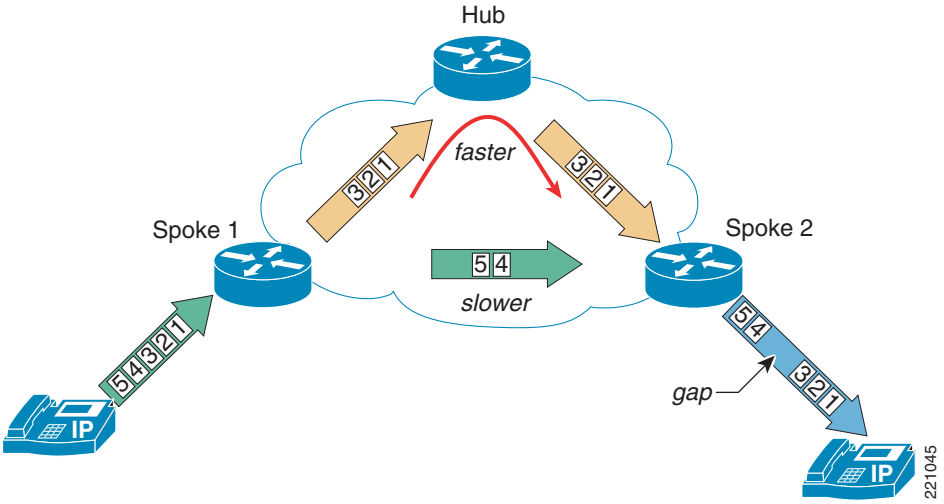

Measurements taken for this test effort show that all tested Cisco endpoints were able to avoid or conceal an audio distortion provided that the difference in latency between the spoke-hub-spoke path and the direct spoke-to-spoke is 50 ms or less. Higher latency differences up to 100 ms provided acceptable voice quality for the overlapping stream case (hub path slower) because the momentary audio distortion was almost imperceptible. When there is more latency on the spoke-to-spoke path, the tolerance is lower. Most Cisco endpoints concealed the effects of jitter up to 25 ms, but at 50 ms and higher, some distortion begins to be audible.

Although this behavior impacts voice quality, it has proven acceptable to the user community. Further details and mitigation approaches are provided in Mitigation of Path Switch Jitter Distortion, page 15.

Risk of Packet Loss

Packet loss can occur for a variety of reasons, but most often oversubscription of allocated bandwidth is the cause. This can be because of overloading physical interfaces or exceeding bandwidth allocated to QoS queues.

Because spoke-to-spoke tunnels are established dynamically, there is no QoS policy per spoke-to-spoke tunnel that queues, shapes, or polices traffic to manage bandwidth associated with that specific tunnel. Instead, per-interface policies are used. These policies maintain control over the traffic traversing the entire interface, which may include multiple spoke-to-spoke tunnels. There is the potential for multiple spokes to direct more voice traffic to a particular spoke than the interface or low-latency queue of the spoke has bandwidth to receive. Likewise, traffic sources within the local network may attempt to direct more voice traffic out through the DMVPN spoke router than there is available bandwidth. This can lead to poor voice quality.

The components of the Cisco Unified Communications solution are not aware of the specifics of the DMVPN network. CAC can deny a call over the WAN if all bandwidth, at either the originating or terminating spoke, has been allocated to other calls. CAC protects existing calls from the risk of an additional call overflowing the low-latency queue resulting in packet loss across all calls on that

interface. Although a Cisco Unified CallManager cluster can be configured to keep track of the WAN bandwidth used between regions (with one region per spoke site, in general), it is otherwise considered "topology unaware". Thus, there is a risk that changes in network topology or bandwidth availability could impact voice quality. Per-location bandwidth values must take LLQ bandwidth allocations and encryption and data link layer protocol overhead into account.

Figure 5 shows a simplified example of voice quality degradation in the absence of CAC. The low-latency queue for traffic leaving the WAN toward Spoke 4 is configured to support two calls. Each originating spoke router also has the low-latency queue bandwidth for two calls in the outgoing direction to the WAN. Assuming congestion conditions on all interfaces, each originating spoke router is easily able to route the media stream for a single call with no packet loss, but when all three streams converge at Spoke 4, one-third of the voice packets are dropped because of the policing function of the low-latency queue.

Figure 5 Degradation Of Voice Traffic Without Call Admission Control

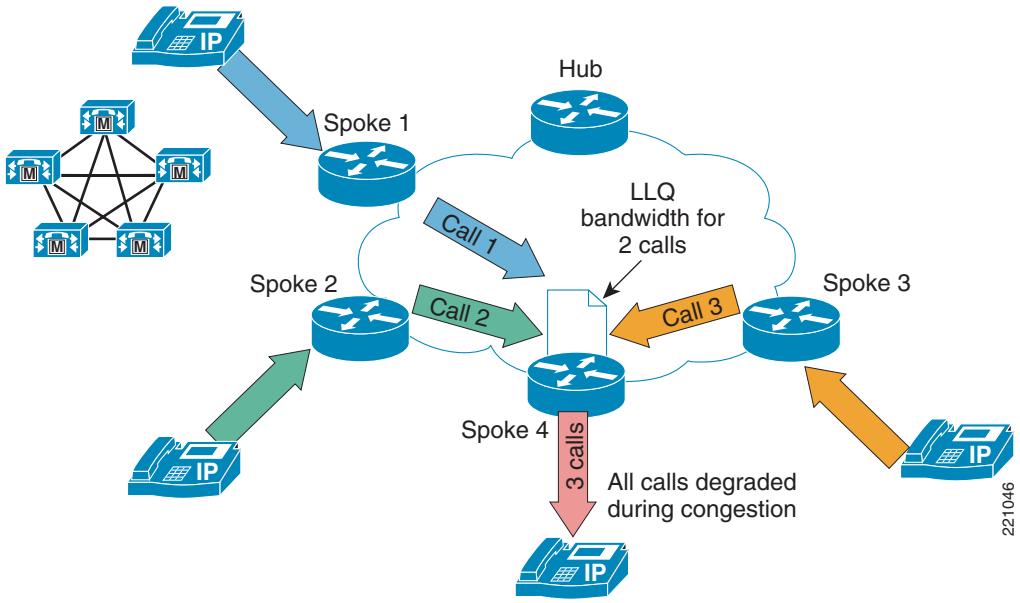

If location-based CAC is enabled (see Figure 6), a third call is denied access to the WAN so that the capacity of the low-latency queue to and from spoke 4 is not exceeded.

Figure 6 Operation Of Call Admission Control

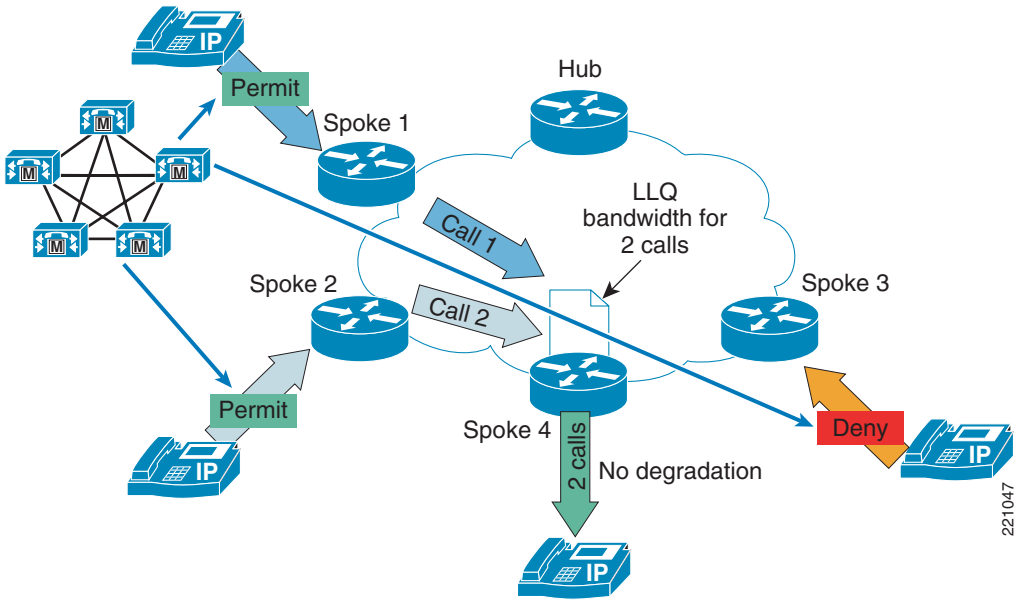

Planning and Best Practices

When you plan to deploy Cisco Unified Communications capability using a DMVPN network for signaling and media transport between sites, consider the following factors:

- Sizing of physical bandwidth from the spoke router to the WAN
- Allocation of available bandwidth for expected classes of traffic on the spoke router
- Management of traffic that uses the low-latency queue
- Assessment of WAN performance
- Selection of spoke router hardware
- Monitoring DMVPN performance to ensure voice quality

Additional WAN planning strategies are provided in Cisco Unified Communications SRND Based on Cisco Unified CallManager 4.x.

Physical Bandwidth Sizing

The amount of bandwidth between the spoke router and the WAN is based on standard traffic engineering methodology. Among the factors influencing this decision are the following:

- Average and peak traffic rates
- Protocols and applications using the bandwidth
- Expected amount of voice traffic
- Availability of alternate paths
- Traffic growth rate

- Encryption and tunneling overhead
- Layer 2 encapsulation overhead

Note The voice topology you select may lead to additional bandwidth utilization for music on hold, recorded announcements, auto attendant functionality concentrated at the main site for PSTN calls arriving at a branch site, and voice mail access.

When you calculate required bandwidth, consider that IPsec encryption and Multipoint Generic Routing Encapsulation (mGRE) tunneling for DMVPN adds significant overhead compared with the small payload size of a G.729 RTP packet. The Layer 3 data rate for a G.729 call (50 pps) is 24 kbps in each direction. Encrypting the media stream using IPsec Tunnel Mode for mGRE increases that rate to approximately 56 kbps. Added to this is the Layer 2 overhead, which varies by interface type from 4 bytes to 14 bytes per packet. See Voice and Video Enabled IPsec VPN (V3PN) SRND for more details. Table 1 lists the estimated bandwidth per call for the interfaces on the target DMVPN spoke routers for both G.729 and G.711. Use of the G.729 codec is generally recommended for bandwidth conservation on WAN links.

Note Additional overhead may be required if voice signaling and media are encrypted by the endpoints.

Table 1 Per-Call Bandwidth Including IPsec, GRE, and Layer 2 Encapsulation

	Per-call bandwidth (kbps) with IPsec, GRE, and Layer 2 Overhead	
Interface type	**G.711**	**G.729**
Gigabit Ethernet	127.2	71.2
OC3 ATM AAL5	237.6	118
2xT1 Multilink PPP	114.4	58.4
Frame Relay	112.4	56.4

In the case of the target government customers, the uplinks to the service provider WANs are already established and administered by the service provider, but deployment of additional voice traffic may require increases in bandwidth on those uplinks after considering call traffic patterns. A standard method of estimating intersite call traffic is to study current traffic patterns on the system being replaced and to add additional volume based on interviewing the customer regarding call rates, durations, peak traffic times, growth plans, and so on. See Cisco Unified Communications SRND Based on Cisco Unified CallManager 4.x and Voice and Video Enabled IPsec VPN (V3PN) SRND for more information.

Allocation of Available Bandwidth Among Traffic Classes

Each spoke site has different bandwidth requirements and differences in expected amounts of traffic for each class. In this case for the targeted government customers, the primary WAN service provider offers a range of pre-defined QoS policies that can be selected for deployment on the spoke router. Table 2 lists the bandwidth allocation by traffic class offered by the service provider. The example sites modeled for this test used profile 110, which provides 40 percent of interface bandwidth for CoS1 (DSCP 46; per-hop behavior EF) and a percentage of the remaining bandwidth for other classes of traffic (Figure 7 and Figure 8).

To make the best use of location-based call admission control (CAC), the bandwidth allocated to voice needs to be allocated symmetrically at each end of the serial link. This means that if 40 percent of the interface bandwidth is allocated to voice on the spoke router uplink to the WAN, the router at the other end needs to have the same percentage allocated for voice traffic toward the spoke router. In addition, the spoke router link to a secondary DMVPN cloud also needs to have at least the same bandwidth allocated to voice as the link to the primary DMVPN cloud.

Table 2 Example COS Packages Offered by Service Provider

CoS Package	Classes Available	Profiles Available	Profile Number	Bandwidth Allocation COS1, COS2, COS3, COS4			
Multimedia High	4	24	101	90	0	0	100
			102	80	80	10	10
			103	80	60	30	10
			104	80	40	30	30
			105	60	80	10	10
			106	60	60	30	10
			107	60	40	30	30
			108	50	0	0	100
Multimedia Low	4	16	109	40	80	10	10
			110	40	60	30	10
			111	40	40	30	30
			112	20	80	10	10
			113	20	60	30	10
			114	20	40	30	30
			115	10	80	10	10
			116	10	60	30	10
			117	10	40	30	30
Critical Data	3	7	118	0	100	0	0
			119	0	80	10	10
			120	0	60	30	10
			121	0	40	30	30
Business Data	2	3	122	0	0	100	0
			123	0	0	90	10
			124	0	0	50	50

Figure 7 Classification of Application Traffic to DSCP Marking

Application	L3 Classification			L2
	IPP	PHB	DSCP	Cos
Routing	6	CS6	48	6
Voice	5	EF	46	5
Video Conferencing	4	AF41	34	4
Streaming Video	4	CS4	32	4
Mission-Critical Data	3	AF31	26	3
Call Signaling	3	CS3	24	3
Transactional Data	2	AF21	18	2
Network Management	2	CS2	16	2
Bulk Data	1	AF11	10	1
Scavenger	1	CS1	8	1
Best Effort	0	0	0	0

221048

Figure 8 Mapping of DSCP Per-Hop Behavior to Service Provider Classes

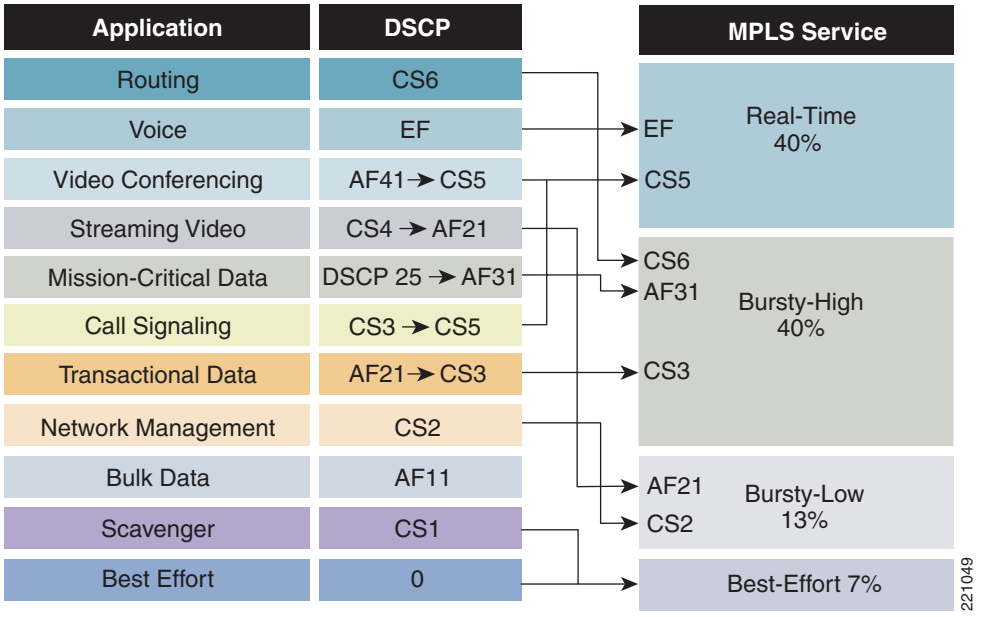

Management of Traffic Using the Low-Latency Queue

Packets marked as DSCP 46 (per-hop behavior EF) retain this marking after encryption by the DMVPN spoke router and are routed through a dedicated low-latency queue. The service provider network keeps the real-time priority across the WAN because of the marking. If adequate bandwidth is reserved for the low-latency queue, latency, jitter, and loss should be minimized for these packets. However, voice traffic can exceed reserved bandwidth and, under congestion conditions, voice quality can be impacted as packets drops increase.

Location-based CAC

For DMVPN, Cisco recommends configuring location-based CAC on the Cisco Unified CallManager. Location-based CAC tracks the bandwidth in use for media streams controlled by the Cisco Unified CallManager. When insufficient bandwidth remains to support a new media stream, calls are intercepted or rerouted. This helps prevent the low-latency queue from overflowing and causing voice quality problems because of packet loss. When a call ends, its bandwidth is made available to new calls.

Provided that all sources of DSCP 46 (per-hop behavior EF) traffic are accounted for by the value for maximum voice bandwidth for each Cisco Unified CallManager location, location-based CAC prevents voice quality issues in many situations. However, because the Cisco Unified CallManager does not have any information about the actual network topology, location-based CAC cannot account for temporary network conditions such as a partial loss of bandwidth along a certain path (for example, one link in a multilink interface), so its effectiveness has some limits. An alternate strategy using RSVP for CAC is not yet supported on DMVPN tunnel interfaces.

See the Cisco Unified Communications SRND Based on Cisco Unified CallManager 4.x (“Call Admission Control” chapter) for more information regarding the capabilities and limitations of topology-unaware CAC. See Location-Based Call Admission Control, page 19 for an example of configuring location-based CAC for DMVPN.

Call Admission Control for Internet Key Exchange

Cisco IOS-based routers provide the capability of limiting the number of Internet Security Association and Key Management (ISAKMP) security associations (SA). This can be based on the number of existing ISAKMP SAs or on the percentage of critical system resources (CPU utilization and memory buffers). IKE CAC can prevent a spoke router from being overwhelmed if it is suddenly inundated with SA requests. This feature does not restrict voice bandwidth utilization directly but it can safeguard spoke router performance from a DMVPN perspective.

See the DMVPN Design Guide and Cisco IOS Security Configuration Guide, Release 12.4—Call Admission Control for IKE for more information.

Assessment of WAN Performance

According to recommendation G.114 of the International Telecommunication Union (ITU), if end-to-end latency is kept below 150 ms, most applications experience essentially transparent interactivity. In practice, satisfactory results are attainable end-to-end latency greater than 150 ms, but the user experience approaches dissatisfaction with latency greater than 300 ms.

To meet this recommendation for IP voice over DMVPN, Cisco recommends following the guidelines in the Voice and Video Enabled VPN SRND regarding WAN service provider performance. Specifically, WAN contribution to jitter should be less than or equal to 20 ms, one-way latency across the WAN should

be less than or equal to 60 ms, and packet loss less should be less than or equal to 0.5 percent. WAN QoS, particularly low-latency queuing with guaranteed bandwidth for voice packets, is essential for voice over DMVPN deployments to meet these requirements.

Cisco recommends that you monitor the performance of your WAN to ensure that these guidelines are being met. The IP Service Level Agreements (IP SLA) command provides effective ways of measuring the performance of the WAN and the DMVPN for these recommendations. See G.729 Codec, page 28 for more information.

Because of the dynamics of establishing a spoke-to-spoke tunnel and the effect on endpoint jitter buffers, Cisco recommends that the difference in latency between the spoke-hub-spoke path and the spoke-to-spoke path be minimized. See Mitigation of Path Switch Jitter Distortion, page 15 for more information.

Mitigation of Path Switch Jitter Distortion

Even if the average one-way latency across a spoke-to-spoke DMVPN network is within the guidelines, it is possible that a brief period of higher latency will occur when the initial Real-Time Protocol (RTP) packets of a call traverse the spoke-hub-spoke path if that call triggers a new spoke-to-spoke tunnel. Test results show that momentary audio distortion because of jitter when the RTP packets switch from the spoke-hub-spoke path to the spoke-to-spoke path is avoided when the difference in latency between the paths is 50 ms or less. Latency differences between 50 ms and 100 ms produce acceptable audio quality, but if the latency difference is above 100 ms, a momentary audio distortion may be audible at the beginning of a call that triggers a new spoke-to-spoke tunnel. Latency differences higher than 100 ms are associated with increasingly noticeable momentary audio distortions at the time the spoke-to-spoke tunnel starts to be used.

In most cases, the latency difference between the spoke-hub-spoke path and the direct spoke-to-spoke path is under 100 ms. If the latency difference between paths on your network is more than 100 ms, consider one of the following approaches:

- Seek WAN performance improvements through modifications of the WAN infrastructure; for example, increasing bandwidth at bottlenecks, validating QoS configurations, reducing the number of hops, and so on. Confirm WAN service provider conformance to contractual service level agreements.
- Establish the spoke-hub-spoke path as the only path for packets to a particular remote spoke by creating a static NHRP association on the tunnel interface that directs traffic for that spoke to the hub.
- Deploy a point-to-point tunnel interface or dedicated WAN link for traffic to another spoke router.
- Establish light traffic between the involved spokes to keep the spoke-to-spoke tunnel active. The performance results in Performance Information, page 41 confirm that there is only a small overhead associated with maintaining a spoke-to-spoke tunnel carrying little traffic. A simple IP SLA operation such as UDP echo repeated with an interval shorter than the NHRP hold time of the tunnel keeps the tunnel active.

Note On a small scale, techniques used to avoid transition to the spoke-to-spoke tunnel during a voice call are feasible, but those that require destination-specific configuration on the spoke routers diminish a key advantage of DMVPN in which any-to-any tunneling is supported without destination-specific configuration commands. Therefore, you should consider one of the other approaches if an excessive momentary jitter problem is being encountered with a large number of voice-supporting remote spokes.

Selection of Spoke Router Hardware

The spoke router needs to be adequately sized for the expected level of traffic. Considering the performance information in this document and in the Voice and Video Enabled IPsec VPN SRND and the Dynamic Multipoint VPN Design Guide documents, select a platform that supports your estimates of traffic through the DMVPN spoke router.

Table 3 lists the data and voice call capacity of three spoke router platforms used on the customer DMVPN network. For each of these platforms, two physical interfaces are shown to permit easier comparison between the platform performance of each router. The number of voice calls shown is based on 50 percent of the total bit rate divided by the bit rate per G.729 call with encryption, tunneling, and Layer 2 overhead.

Table 3 Comparison of Spoke Router Hardware Performance

Platform	Interface	Number of G.729 calls (50% voice, 50% Imix) single tunnel	Packet rate @ 60% CPU (kpps) single tunnel	Bit rate @ 60% CPU (Mbps), single tunnel
7206 NPE-G1 VAM2+	ATM OC-3	181	17.1	23.2
7206 NPE-G1 VAM2+	GigE	193	18.3	24.7
2821 EPII-PLUS	2xT1 PPP Multilink	38	4.8	4.9
2821 EPII-PLUS	GigE	41	5.1	5.2
1841 BPII-PLUS	Frame relay	11	1.4	1.4
1841 BPII-PLUS	FE	18	2.3	2.3

Monitoring DMVPN Performance for Voice

There are several useful techniques for monitoring DMVPN performance. Some require access to the spoke router console or network management interface, others can be used inside the local network. The following sections discuss these techniques.

Using IP SLA for Monitoring of DMVPN Performance

Useful statistics maintained on the DMVPN spoke routers help to assess the condition of the router and the DMVPN network. For example, QoS policy map statistics, DMVPN, IPsec and ISAKMP statistics, interface counters, and CPU and memory statistics. However, in the case of a DMVPN network operated by a service provider, access to the spoke router console may not be readily available.

IP SLA operations are useful for monitoring DMVPN performance without requiring access to the DMVPN spoke router consoles. You can figure them on a Cisco router that has access to the DMVPN network. Preferably, the source and destination routers should be as close as possible to the DMVPN spoke routers so that the results can more easily be attributed to DMVPN and underlying WAN network conditions.

The following information relevant to voice performance is available through IP SLA operations:

- One-way and round-trip latency, jitter, lost and out-of-sequence packet measurements

- Estimation of audio quality given measured network conditions

In addition, you may configure the measurements to automatically repeat on a regular basis and have SNMP traps associated with pre-determined thresholds.

For more information, see the "IP SLAs--Proactive Threshold Monitoring of IP SLAs Operations" section of the Cisco IOS IP SLAs Configuration Guide.

Realtime Monitoring Tool

The Cisco Unified CallManager Realtime Monitoring Tool (RTMT) provides visual tracking of many Cisco Unified CallManager performance and configuration metrics. It is helpful to use the RTMT when configuring and monitoring location-based CAC. Bandwidth in use and calls in progress per location can be displayed as shown in Figure 9.

Figure 9 Realtime Monitoring Tool

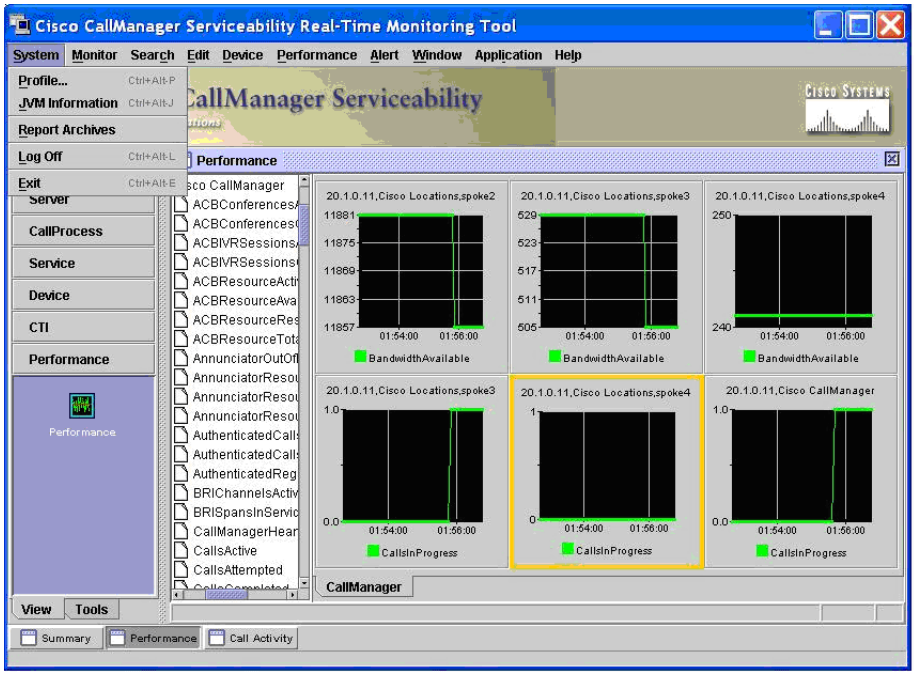

Spoke Router CLI

Table 4 lists commands that provide information relevant to voice traffic over a DMVPN network from the perspective of the spoke routers.

Table 4 Useful CLI Commands On The Spoke Router For Voice Over DMVPN

Command	Description
show dmvpn	Shows summary of DMVPN tunnels, including peers and uptime; aggregates output from related commands
show ip nhrp	Shows NHRP associations, timers, tunnel uptime

Table 4 Useful CLI Commands On The Spoke Router For Voice Over DMVPN

show crypto isakmp	Shows ISAKMP security associations
show policy-map interface	Shows QoS metrics including traffic levels for each queue, packet drops, queue size, and so on.
show interface	Shows interface status, traffic counts

Summary of Best Practices

This section summarizes the best practices for deploying Cisco Unified Communications using a DMVPN network for signaling and media transport between sites:

- Adequately size the physical bandwidth between the spoke router and the WAN
- Allocate adequate bandwidth for expected classes of traffic on the spoke router
- Classify traffic with end-to-end Diffserv QoS markings
- Dedicate sufficient bandwidth to low-latency queues along media transmission path
- Configure location-based CAC to manage traffic using the low-latency queues
- Assess and monitor DMVPN and WAN performance; minimize underlying WAN latency through selection of a Cisco Powered Network service provider
- Deploy adequately powered spoke router platforms for the anticipated traffic load
- Use available techniques to monitor DMVPN performance to ensure voice quality

Implementation and Configuration

This section provides configuration recommendations for deploying a Cisco Unified Communications voice solution with a DMVPN for intersite media and signaling transport. The recommendations are based on the assumption that an existing spoke-to-spoke DMVPN network is used. See the Dynamic Multipoint VPN Design Guide for more information on configuring DMVPN.

The configuration information in this section is based on the following assumptions:

- The customer deploys and uses Skinny (SCCP)-based IP phones and MGCP-based gateways.
- A single-cluster solution meets expected scaling requirements for these deployments.
- The DMVPN spoke routers are maintained by the service providers and are already configured with QoS policies providing low-latency queuing for IP voice packets and Class-Based Weighted Fair Queuing (CBWFQ) for other types of traffic.
- The LAN infrastructure within each spoke site follows Cisco recommendations and preserves the DSCP markings required to identify IP voice traffic.
- Adequate bandwidth is provisioned to support the expected peaks in call traffic traversing the DMVPN network.

Topology

In the multisite WAN with centralized call processing model used for this set of customers, a single Cisco Unified CallManager cluster provides call processing for all locations on the IP telephony network. The Cisco Unified CallManager cluster usually resides at the main (central or headquarters) location, along with other devices such as phones and gateways. The remote locations contain additional devices, but no Cisco Unified CallManager servers. IP WAN links connect the remote locations to the main location.

Location-Based Call Admission Control

The Cisco Unified CallManager uses "locations" to implement CAC, which enables regulation of audio quality by limiting the amount of bandwidth that is available for calls over links between the locations. For more information, see the "Call Admission Control" section in the Cisco Unified CallManager System Guide.

If CAC is not used to limit the bandwidth on WAN links, an unlimited number of calls may be active on that link concurrently. This situation can cause the audio quality of all calls to degrade as the physical interface of the local or remote DMVPN uplink becomes oversubscribed.

The main location and each remote site has a unique name configured on the Cisco Unified CallManager. The low-latency queue on the spoke router drops packets that exceed the maximum bandwidth of its queue if all bandwidth allocated to other queues is in use (that is, congestion condition exists). Cisco recommends that you specify only the bandwidth dedicated to voice packets for each location.

The Cisco Unified CallManager automatically assumes a specific bandwidth value per call. Because of the overhead associated with mGRE, IPsec, and the Layer 2 protocol being used, a conversion is required when specifying the bandwidth for the location. To limit the number of calls with a higher encapsulation overhead than the Cisco Unified CallManager assumes, the available bandwidth needs to be adjusted downward to achieve the call limit desired. Table 5 lists the per-call bandwidth assumed by the Cisco Unified CallManager in the CAC calculation.

Table 5 Bandwidth Assumed By The Cisco Unified CallManager Location-Based CAC

Codec	Bandwidth
G.711	80 kbps
G.722	80 kbps
G.723	24 kbps
G.728	16 kbps
G.729	24 kbps
GSM	29 kbps
Wideband	272 kbps

There is no provision in the Cisco Unified CallManager GUI to perform this calculation, so it must be performed outside of the GUI, and the result of the calculation must be used in place of the actual voice bandwidth available. The same concept applies to video bandwidth.

Use the following steps to determine the bandwidth for each location:

1. Determine the available (dedicated) voice bandwidth this should be the bandwidth dedicated to the low-latency queue minus the bandwidth (with overhead) any other known sources of voice traffic using that queue
2. Divide this number by the voice bandwidth per call with encryption, tunneling and Layer 2 encapsulation (refer to the "Bandwidth Provisioning for WAN Edge QoS" section in the Voice and Video Enabled VPN SRND for information on calculating this value)
3. Multiply the result by the Cisco Unified CallManager default bandwidth (Table 6):

$$Location\ bandwidth = \left(\frac{dedicated\ total\ bandwidth}{actual\ encapsulated\ bandwidth\ per\ call\ leg} \right) * CallManager\ default\ bandwidth\ per\ call\ leg$$

Table 6 and Table 7 list an example of the calculation in which the codec is G.729 and the Layer 2 protocol is Frame Relay. In this case, the low-latency queue on the spoke router is allocated 312 kbps. You need to determine this information for each physical interface type.

Table 6 Bandwidth Comparisons At Points In Network IPsec/GRE/FR

	G.729 Payload	IP + UDP + RTP + payload	Ethernet + 802.1Q + IP + UDP + RTP + Payload	F/R + IPsec + GRE + IP + UDP + RTP + Payload
Bytes	20	60	80	140
Packet rate	50 pps	50 pps	50 pps	50 pps
Bandwidth	8 kbps	24 kbps	34 kbps	56 kbps

Table 7 shows some sample adjustments for location configuration because of encapsulation overhead (Frame Relay).

Table 7 Example Adjustments for Location Configuration

	Cisco Unified CallManager Default for G.729 (from Table 5)	Actual Utilization with Encapsulation Overhead
Bandwidth per call leg	24 kbps	56 kbps
Calls per 312 kbps	312 / 24 = 13	312 / 56 = 5.57
Bandwidth to specify in location configuration for 312 kbps	312 K	5.57 * 24 kbps = 133 kbps

The value you need to input into the Cisco Unified CallManager Location Configuration screen in the Audio Bandwidth field is 133 kbps. If the actual bandwidth allocated to the low-latency queue, 312 kbps is entered without this adjustment, the Cisco Unified CallManager assumes 24 kbps per G.729 call leg, and allows 13 calls to use the 312 kbps bandwidth even though there is only capacity low-latency queue of the DMVPN uplink interface for 5 calls.

Figure 10 shows a location configuration example.

Figure 10 Location Configuration Example

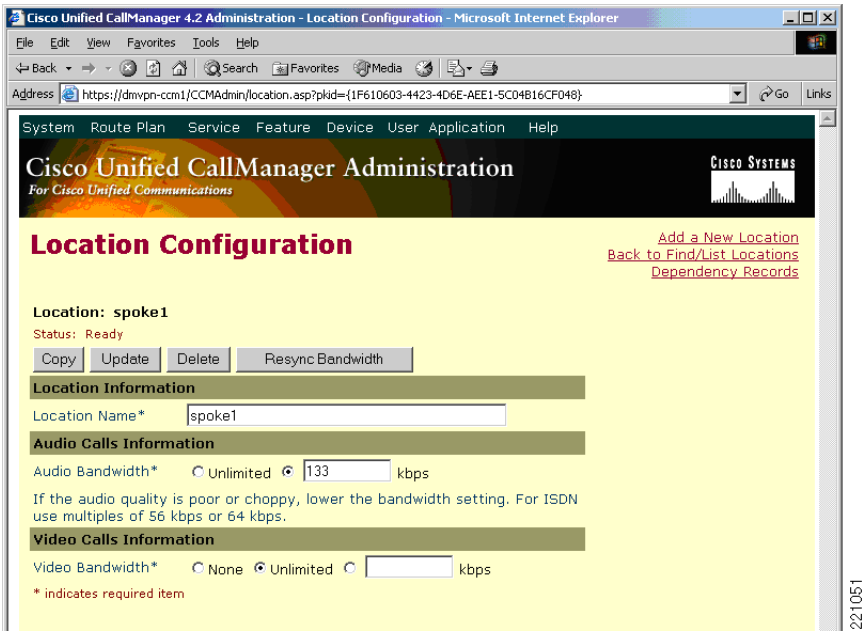

The effectiveness of the location-based CAC feature requires strict control of traffic using the low-latency queue (in this case DSCP 46; per-hop behavior EF voice traffic) through the spoke router. If necessary, the bandwidth specified for a location can be further reduced to account for other traffic using the same code point or another Cisco Unified CallManager cluster directing calls to the DMVPN via the same tunnel interface.

After the locations have been configured, you need to assign to the endpoints (gateway, IP phone, and so on) that are located behind the WAN interface.

Automated Alternate Routing

Automated Alternate Routing (AAR) is a mechanism to provide an alternate route to remote sites when insufficient bandwidth is available to route the call over the primary route. In this deployment, Cisco recommends having the WAN as the primary route for media and a PSTN connection as the alternate route. There are no DMVPN-specific configuration steps for this feature. See the “Dial Plan/AAR” section of the Cisco Unified Communications SRND Based on Cisco Unified CallManager 4.x document for configuration details.

Music On Hold

Cisco recommends that you enable the SRST-based Music on Hold (MoH) feature to minimize bandwidth consumption on the DMVPN uplinks for sites where bandwidth is constrained. This feature requires setting up the MoH server in the main/headquarters site to multicast MoH for the IP phones that are local to that location; for example, on an address such as 239.192.240.1. The SRST router is then set to multicast an audio stream on the same address so that the IP phones near the SRST router in the remote site can access the MoH audio. In this case, multicast routing of the MoH stream in the main site needs to be blocked from traversing the WAN, either through a reduced time-to-live count or an access list.

When a party at the remote site is placed on hold, the Cisco Unified CallManager instructs the IP phone at that site to listen to 239.192.240.1. It receives the music media stream from the local SRST router instead of from the main site music source across the WAN.

The following steps describe the configuration.

Step 1 Create the MoH server with the multicast option and limited hop count (see Figure 11) on the Cisco Unified CallManager.

Figure 11 Creation Of MoH Server

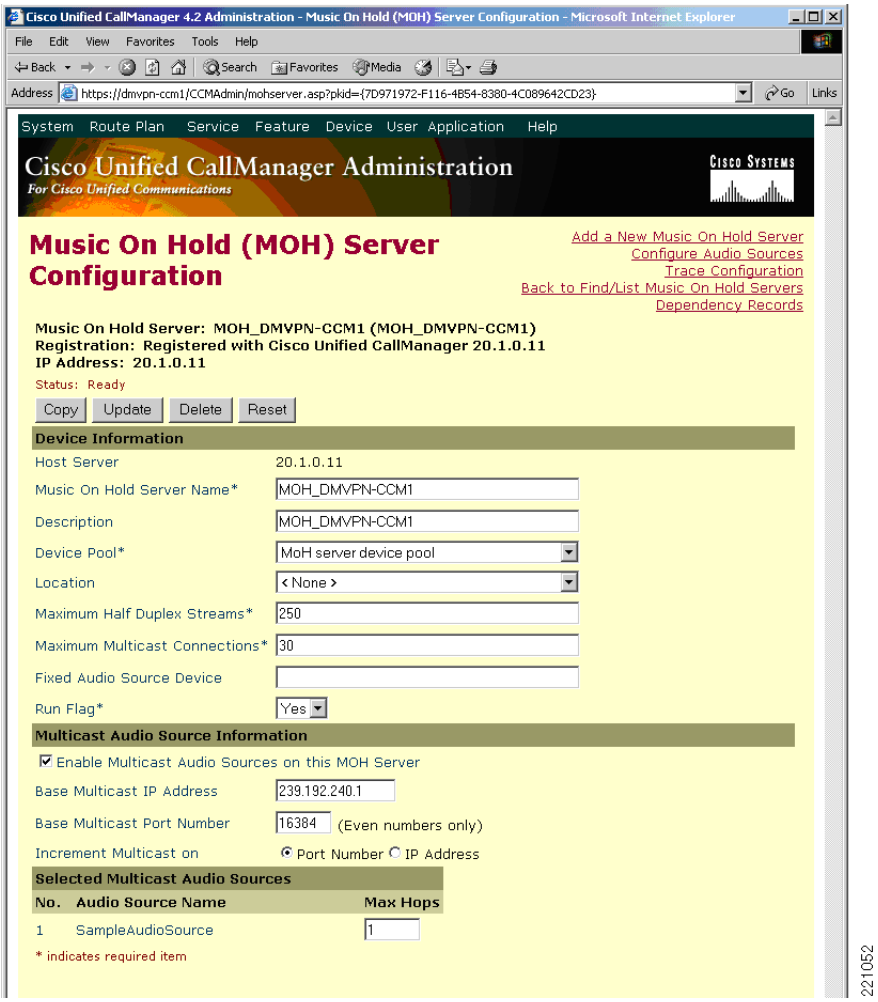

Step 2 Place the server in a media resource group (MRG). (See Figure 12.)

Figure 12 Insertion Of The MoH Server In The Media Resource Group

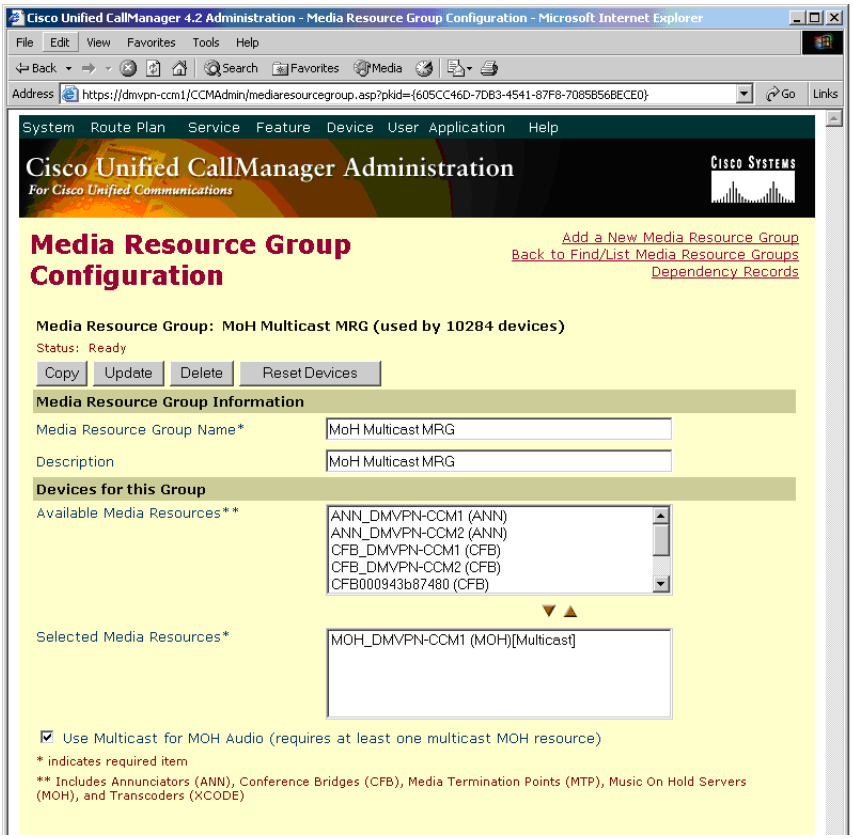

Step 3 Place the MRG in a media resource group list (See Figure 13.)

Figure 13 Insertion Of Media Resource Group Into Media Resource Group List

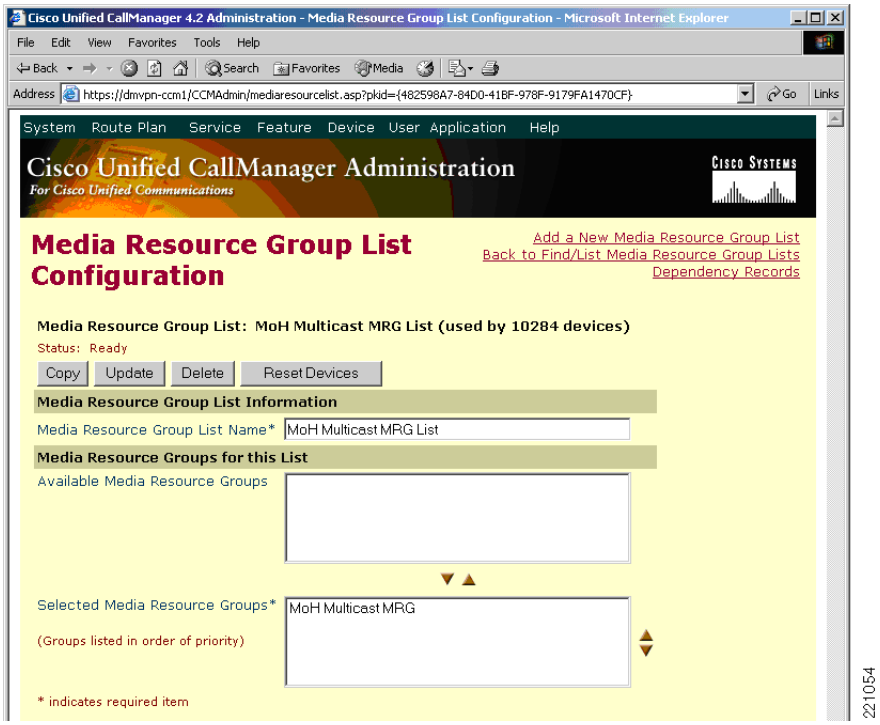

Step 4 Put the MoH server in its own region (with intersite codec selection set to G.711) via a dedicated device pool (See Figure 14.)

Figure 14 Creation Of A Dedicated Region For The MoH

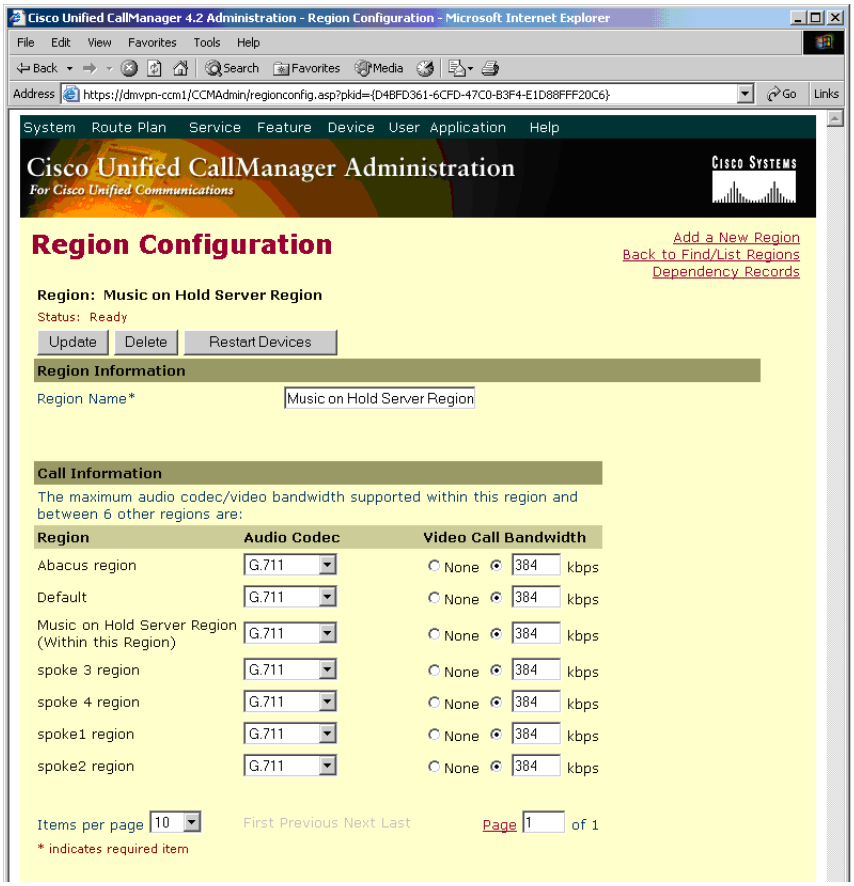

Step 5 Enable MoH multicast with the same multicast destination address as on the Cisco Unified CallManager on the SRST router.

```
ccm-manager music-on-hold
call-manager-fallback
  moh music-on-hold.au
    multicast moh 239.192.240.1 port 16384
```

If, as an alternative, you choose to have MoH media unicast across the DMVPN, you can reduce the bandwidth consumed by the music stream by using the G.729 codec and accepting a small reduction in audio fidelity compared with the default G.711 codec.

MoH servers automatically mark audio stream traffic the same as voice bearer traffic DSCP 46 (per-hop behavior EF). Therefore, as long as QoS is properly configured on the network, MoH streams receive the same classification and low-latency queuing treatment as phone-originated RTP media traffic.

For additional information and best practices for MoH, see the Cisco Unified Communications SRND based on Cisco Unified CallManager 4.x and the Integrating Cisco CallManager and Cisco SRST to Use Cisco SRST As a Multicast MOH Resource feature guide.

SRST and PSTN Gateway at Each Remote Site

Each remote site requires local PSTN access, both for PSTN-destined calls and as an alternate route for calls denied bandwidth because of location-based CAC. SRST and gateway configuration for DMVPN follows the standard procedures defined for the Multisite WAN with Centralized Call Control.

WAN Requirements

No configuration is required to verify the WAN requirements, but they are mentioned here because they may require measurement and coordination with the service provider as part of deployment. See Assessment of WAN Performance, page 14 for best practices related to WAN performance.

Differentiated Services

To minimize latency and jitter, the service provider needs to provide a priority queue for voice packets outbound from the provider edge (PE routers) of the WAN. The bandwidth allocated to the priority queue should mirror the policy on the spoke router. No additional configuration for this requirement is needed because this is already the case for the DMVPN network of the target customer.

Quality of Service

In the case of the target government customers, the QoS configuration on the DMVPN spoke routers is administered by the primary WAN service provider and already supports dedicated bandwidth for voice in the form of a low-latency queue. The following example shows the QoS configuration for the tunnel interface from a spoke router to the DMVPN. The low-latency queue configuration allocates 312 kbps for DSCP 46 (per-hop behavior EF) packets (RTP voice packets) out of an overall T1 frame relay bandwidth of 768 kbps.

```
! Standard QoS config

class-map match-any PCLASS_COS2_SAA
 match  dscp 27
class-map match-any PCLASS_COS3_SAA
 match  dscp 19
class-map match-any PCLASS_COS1_SAA
 match  dscp 47
class-map match-any PCLASS_COS4_SAA
 match  dscp 19
class-map match-any PCLASS_COS4
 match not  dscp 1
class-map match-any QCLASS_COS4
 match not  dscp 1
class-map match-any PCLASS_COS2
 match  dscp af31
class-map match-any QCLASS_COS3
 match  dscp 19
 match  dscp af21
class-map match-any PCLASS_COS3
 match  dscp af21
class-map match-any QCLASS_COS2
 match  dscp 27
 match  dscp af31
class-map match-any QCLASS_COS1
 match  dscp 47
```

```
 match  dscp ef
class-map match-any PCLASS_COS1
 match  dscp ef
class-map match-any PCLASS_NM
 match access-group 180
class-map match-any PCLASS_RP
 match  dscp cs6
class-map match-any QCLASS_NM
 match  dscp cs6


class-map match-any Prec_3
 match ip precedence 3
class-map match-any Prec_2
 match ip precedence 2
class-map match-any Prec_1
 match ip precedence 1
class-map match-any Prec_0
 match ip precedence 0
class-map match-any Prec_7
 match ip precedence 7
class-map match-any Prec_6
 match ip precedence 6
class-map match-any Prec_5
 match ip precedence 5
class-map match-any Prec_4
 match ip precedence 4
!
!
policy-map remark
 class Prec_0
  set ip dscp default
 class Prec_1
  set ip dscp af21
 class Prec_2
  set ip dscp af21
 class Prec_3
  set ip dscp af31
 class Prec_4
  set ip dscp af31
 class Prec_5
  set ip dscp ef
 class Prec_6
  set ip dscp af21
 class Prec_7
  set ip dscp af21


policy-map NM
 class PCLASS_RP
  police cir 8000 bc 8000 be 8000
    conform-action set-dscp-transmit cs6
    exceed-action set-dscp-transmit cs6
 class PCLASS_NM
  police cir 8000 bc 8000 be 8000
    conform-action set-dscp-transmit af21
    exceed-action set-dscp-transmit af21

policy-map COS4
 class PCLASS_COS4
  police cir 80000 bc 10000
    conform-action set-dscp-transmit default
    exceed-action set-dscp-transmit default
 class PCLASS_COS4_SAA
```

```
  set dscp default

policy-map COS3
 class PCLASS_COS3
  police cir 144000 bc 18000
    conform-action set-dscp-transmit af21
    exceed-action set-dscp-transmit af22
 class PCLASS_COS3_SAA
  set dscp af21

policy-map COS2
 class PCLASS_COS2
  police cir 280000 bc 35000
    conform-action set-dscp-transmit af31
    exceed-action set-dscp-transmit af32
 class PCLASS_COS2_SAA
  set dscp af31

policy-map COS1
 class PCLASS_COS1
  police cir 312000 bc 39000
    conform-action set-dscp-transmit ef
    exceed-action drop
 class PCLASS_COS1_SAA
  set dscp ef

policy-map D_768K
 class QCLASS_NM
  bandwidth remaining percent 10
  random-detect dscp-based
  random-detect dscp 18   100   200   10
  random-detect dscp 48   200   300   10
  service-policy NM
 class QCLASS_COS1
  priority 312
   compress header ip rtp
  service-policy COS1
 class QCLASS_COS2
  bandwidth remaining percent 60
  random-detect dscp-based
  random-detect dscp 26   200   300   10
  service-policy COS2
 class QCLASS_COS3
  bandwidth remaining percent 30
  random-detect dscp-based
  random-detect dscp 18   200   300   10
  service-policy COS3
```

G.729 Codec

When deploying voice in a WAN environment, Cisco recommends the lower-bandwidth G.729 codec for any voice calls that traverse WAN links because this practice provides bandwidth savings compared to G.711 at the expense of a small reduction in audio fidelity.

To configure a codec for WAN calls, configure a “region” to represent points in the network that impact codec selection. To configure G.729 for intersite calls and G.711 for intrasite calls, a unique region can be created for each remote site, and a codec option can be selected both for calls within that region and for calls to other regions. See the Cisco Unified Communications SRND Based on Cisco Unified CallManager 4.x for more information.

IP SLA Options Relevant to Voice Over DMVPN

This section provides some sample IP SLA operations that are useful in assessing DMVPN performance for voice traffic. Multiple other options are supported by IP SLA command. For more information, see the Cisco IOS IP SLAs Configuration Guide, Release 12.4T.

To obtain performance information about the DMVPN network, Cisco recommends the UDP-jitter operation. Configure the codec and ToS options to simulate G.729 voice packets with DSCP 46 (per-hop behavior EF). See the example titled "UDP Jitter with Codec and ToS" in UDP Jitter With Codec and TOS, page 30. If there is no pre-existing traffic between two sites, the most useful results are obtained if the interval between measurements is set higher than the NHRP hold time. This provides some information on the latency and jitter during DMVPN spoke-to-spoke tunnel setup. See the example in IP SLA Scheduling Compared With NHRP Holdtime, page 32.

If a router with a voice card is available to act as the IP SLA source router, the VoIP operation is also useful. This operation uses Digital Signal Processing (DSP) functions to estimate voice quality. Only the source router needs to have a voice card for this operation. An example of this command is in VoIP Operation, page 32.

In each of the following examples, the address 192.168.2.2 represents the destination of the IP SLA probes. All examples assume the following configuration on the terminating router:

```
ip sla responder
```

UDP Jitter Operation

This IP SLA option measures delay, jitter, and packet loss with UDP. In this example, 10 packets are sent at 20 ms intervals every 30 seconds.

Originating router configuration is as follows:

```
ip sla 10
udp-jitter 192.168.2.2 5000
frequency 30
ip sla schedule 10 life forever start-time now
```

Example results are as follows:

```
router#show ip sla statistics

Round Trip Time (RTT) for Index 10
Latest RTT: 17 milliseconds
Latest operation start time: 14:58:22.494 EST Mon Jan 22 2007
Latest operation return code: OK
RTT Values:
Number Of RTT: 10 RTT Min/Avg/Max: 14/17/21 milliseconds
Latency one-way time:
Number of Latency one-way Samples: 10
Source to Destination Latency one way Min/Avg/Max: 3/3/7 milliseconds
Destination to Source Latency one way Min/Avg/Max: 11/14/16 milliseconds
Jitter Time:
Number of Jitter Samples: 9
Source to Destination Jitter Min/Avg/Max: 4/4/4 milliseconds
Destination to Source Jitter Min/Avg/Max: 1/2/5 milliseconds
Packet Loss Values:
Loss Source to Destination: 0 Loss Destination to Source: 0
Out Of Sequence: 0 Tail Drop: 0 Packet Late Arrival: 0
Voice Score Values:
Calculated Planning Impairment Factor (ICPIF): 0
Mean Opinion Score (MOS): 0
```

```
Number of successes: 6
Number of failures: 0
Operation time to live: Forever

router#show ip sla statistics detail

Round Trip Time (RTT) for Index 10
Latest RTT: 17 milliseconds
Latest operation start time: 14:58:22.494 EST Mon Jan 22 2007
Latest operation return code: OK
Over thresholds occurred: FALSE
RTT Values:
Number Of RTT: 10 RTT Min/Avg/Max: 14/17/21 milliseconds
Latency one-way time:
Number of Latency one-way Samples: 10
Source to Destination Latency one way Min/Avg/Max: 3/3/7 milliseconds
Destination to Source Latency one way Min/Avg/Max: 11/14/16 milliseconds
Source to Destination Latency one way Sum/Sum2: 38/170
Destination to Source Latency one way Sum/Sum2: 140/1984
Jitter Time:
Number of Jitter Samples: 9
Source to Destination Jitter Min/Avg/Max: 4/4/4 milliseconds
Destination to Source Jitter Min/Avg/Max: 1/2/5 milliseconds
Source to destination positive jitter Min/Avg/Max: 4/4/4 milliseconds
Source to destination positive jitter Number/Sum/Sum2: 2/8/32
Source to destination negative jitter Min/Avg/Max: 4/4/4 milliseconds
Source to destination negative jitter Number/Sum/Sum2: 2/8/32
Destination to Source positive jitter Min/Avg/Max: 2/2/3 milliseconds
Destination to Source positive jitter Number/Sum/Sum2: 3/8/22
Destination to Source negative jitter Min/Avg/Max: 1/2/5 milliseconds
Destination to Source negative jitter Number/Sum/Sum2: 5/11/35
Interarrival jitterout: 0 Interarrival jitterin: 0
Packet Loss Values:
Loss Source to Destination: 0 Loss Destination to Source: 0
Out Of Sequence: 0 Tail Drop: 0 Packet Late Arrival: 0
Voice Score Values:
Calculated Planning Impairment Factor (ICPIF): 0
Mean Opinion Score (MOS): 0
Number of successes: 6
Number of failures: 0
Operation time to live: Forever
Operational state of entry: Active
Last time this entry was reset: Never
```

UDP Jitter With Codec and TOS

The UDP jitter IP SLA option can be customized with a codec. By default, 1000 packets are sent in 20 seconds with G.729 size payload and 50 ms interval. All IP SLA options can also have an optional type of service (ToS) value, shown in this example with the equivalent of DSCP 46 (per-hop behavior EF). Because the use of this option consumes low-latency queue bandwidth, the location-based CAC configuration on the Cisco Unified CallManager needs to incorporate this factor.

Originating router configuration is as follows:

```
ip sla 10
udp-jitter 192.168.2.2 5000 codec g729a
tos 184
frequency 30
ip sla schedule 10 life forever start-time now
```

Example results are as follows:

```
router#show ip sla statistics 10

Round Trip Time (RTT) for Index 10
Latest RTT: 16 milliseconds
Latest operation start time: 15:20:03.391 EST Mon Jan 22 2007
Latest operation return code: OK
RTT Values:
Number Of RTT: 1000 RTT Min/Avg/Max: 9/16/69 milliseconds
Latency one-way time:
Number of Latency one-way Samples: 1000
Source to Destination Latency one way Min/Avg/Max: 2/3/17 milliseconds
Destination to Source Latency one way Min/Avg/Max: 6/13/66 milliseconds
Jitter Time:
Number of Jitter Samples: 999
Source to Destination Jitter Min/Avg/Max: 1/2/14 milliseconds
Destination to Source Jitter Min/Avg/Max: 1/3/55 milliseconds
Packet Loss Values:
Loss Source to Destination: 0 Loss Destination to Source: 0
Out Of Sequence: 0 Tail Drop: 0 Packet Late Arrival: 0
Voice Score Values:
Calculated Planning Impairment Factor (ICPIF): 11
MOS score: 4.06
Number of successes: 3
Number of failures: 0
Operation time to live: Forever
```

```
router#show ip sla statistics 10 detail

Round Trip Time (RTT) for Index 10
Latest RTT: 16 milliseconds
Latest operation start time: 15:21:33.387 EST Mon Jan 22 2007
Latest operation return code: OK
Over thresholds occurred: FALSE
RTT Values:
Number Of RTT: 1000 RTT Min/Avg/Max: 9/16/60 milliseconds
Latency one-way time:
Number of Latency one-way Samples: 1000
Source to Destination Latency one way Min/Avg/Max: 2/3/20 milliseconds
Destination to Source Latency one way Min/Avg/Max: 6/12/57 milliseconds
Source to Destination Latency one way Sum/Sum2: 3355/13193
Destination to Source Latency one way Sum/Sum2: 12965/180127
Jitter Time:
Number of Jitter Samples: 999
Source to Destination Jitter Min/Avg/Max: 1/2/16 milliseconds
Destination to Source Jitter Min/Avg/Max: 1/3/45 milliseconds
Source to destination positive jitter Min/Avg/Max: 1/2/14 milliseconds
Source to destination positive jitter Number/Sum/Sum2: 175/357/1541
Source to destination negative jitter Min/Avg/Max: 1/2/16 milliseconds
Source to destination negative jitter Number/Sum/Sum2: 171/357/1707
Destination to Source positive jitter Min/Avg/Max: 1/3/45 milliseconds
Destination to Source positive jitter Number/Sum/Sum2: 452/1620/10212
Destination to Source negative jitter Min/Avg/Max: 1/3/43 milliseconds
Destination to Source negative jitter Number/Sum/Sum2: 427/1617/10845
Interarrival jitterout: 0 Interarrival jitterin: 0
Packet Loss Values:
Loss Source to Destination: 0 Loss Destination to Source: 0
Out Of Sequence: 0 Tail Drop: 0 Packet Late Arrival: 0
Voice Score Values:
Calculated Planning Impairment Factor (ICPIF): 11
MOS score: 4.06
Number of successes: 6
Number of failures: 0
```

```
Operation time to live: Forever
Operational state of entry: Active
Last time this entry was reset: Never
```

IP SLA Scheduling Compared With NHRP Holdtime

If you use IP SLA operations to assess spoke-hub-spoke latency, configure the IP SLA operation to have an interval longer than the NHRP hold time. In the following example, the tunnel interface has an NHRP hold time of 300 seconds and the IP SLA operation is configured to occur at intervals of 400 seconds. If there is no other traffic between the source spoke and the destination spoke than the IP SLA probes, this configuration triggers a new spoke-to-spoke tunnel each time it runs.

The initial packets of the UDP jitter test traverse the spoke-hub-spoke path. Although the statistics collected do not specifically identify the latency and jitter of the spoke-hub-spoke path compared with the spoke-to-spoke path, they include the metrics from both paths when calculating the minimum/average/maximum values.

The following example obtains the NHRP holdtime value:

```
router#show run int tunnel 0 | include interface|holdtime
interface Tunnel0
ip nhrp holdtime 300
```

Originating router configuration is as follows:

```
ip sla 10
udp-jitter 192.168.2.2 5000 codec g729a
tos 184
frequency 400
ip sla schedule 10 life forever start-time now
```

If IP SLA operations are being used to maintain a spoke-to-spoke tunnel, the frequency parameter should be assigned a value less than 2 minutes.

VoIP Operation

This IP SLA option measures latency and jitter, and estimates voice quality using DSP functions in the source router. No DSP hardware is required in the destination router for this operation.

Originating router configuration is as follows:

```
ip sla 10
voip rtp 192.168.2.2 source-ip 192.168.3.2 source-voice-port 0/0/0:23 codec g729a duration
30 advantage-factor 0
ip sla schedule 10 life forever start-time now
```

Example results are as follows:

```
router#show ip sla statistics 10

Round Trip Time (RTT) for Index 10
Type of operation: rtp
Latest operation start time: 16:29:20.278 EST Mon Jan 22 2007
Latest operation return code: OK
Latest RTT (milliseconds): 16
Source to Destination Path Measurements:
Interarrival Jitter: 0
Packets Sent: 750
Packets Lost: 0
Estimated R-factor: 82 MOS-CQ: 4.06
Destination to Source Path Measurements:
```

```
Interarrival Jitter: 6
Packets Sent: 721
Packets Lost: 0
Estimated R-factor: 82 MOS-CQ: 4.06 MOS-LQ: 0.00
Operation time to live: Forever
Operational state of entry: Active
Last time this entry was reset: Never


router#show ip sla statistics 10 detail

Round Trip Time (RTT) for Index 10
Type of operation: rtp
Latest operation start time: 16:29:20.278 EST Mon Jan 22 2007
Latest operation return code: OK
Latest RTT (milliseconds): 16
Source to Destination Path Measurements:
Interarrival Jitter: 0
Packets Sent: 750
Packets Lost: 0 Packets MIA: 0
Estimated R-factor: 82 MOS-CQ: 4.06
One way latency(avg/min/max): 2/2/13
Destination to Source Path Measurements:
Interarrival Jitter: 6
Packets Sent: 721
Packets Lost: 0 Packets OOS: 0
Packets Late: 0 Packets Early: 0
Frame Loss: 0
Estimated R-factor: 82 MOS-CQ: 4.06 MOS-LQ: 0.00
One way latency(avg/min/max): 13/7/74
Over thresholds occurred: FALSE
Operation time to live: Forever
Operational state of entry: Active
Last time this entry was reset: Never
```

UDP Echo

This IP SLA option sends two UDP packets in each direction.

Originating router configuration is as follows:

```
ip sla 10
udp-echo 192.168.2.2 5000 source-ip 192.168.3.2
ip sla schedule 10 life forever start-time now
```

Example results are as follows:

```
router#show ip sla statistics 10

Round Trip Time (RTT) for Index 10
Latest RTT: 15 milliseconds
Latest operation start time: 17:15:29.555 EST Mon Jan 22 2007
Latest operation return code: OK
Number of successes: 1
Number of failures: 0
Operation time to live: Forever

router#show ip sla statistics 10 detail

Round Trip Time (RTT) for Index 10
Latest RTT: 15 milliseconds
```

```
Latest operation start time: 17:15:29.555 EST Mon Jan 22 2007
Latest operation return code: OK
Over thresholds occurred: FALSE
Number of successes: 1
Number of failures: 0
Operation time to live: Forever
Operational state of entry: Active
Last time this entry was reset: Never
```

Test Approach

The following test concepts were used with the goal of validating the proposed solution for IP voice over DMVPN networks:

- Test range of calls and functions on a test bed representing the target customers
- Use a test bed that is scalable from the low end to the high end of expected scale
- Exercise a mix of call types, durations, features, and so on, in traffic mix
- Simulate DMVPN events impacting voice
- Monitor voice quality
- Perform stress testing, negative testing, and long duration testing
- Assess performance of a range of spoke router platforms
- No coverage in this stage for topologies not being proposed (that is, more highly scaled voice solutions)

Most of the test effort involved testing features that interact with the DMVPN WAN. Features and functions not influenced by topology or WAN performance were not tested.

Software and Hardware Environment

Figure 15 shows the test bed topology. This topology is representative of the DMVPN network deployed by the customer. It is a dual-DMVPN cloud topology with alternate service provider capability. The hub is a dual-tier headend architecture with two Cisco 7206s for one cloud and one Cisco 7206 for the other cloud with a Cisco 7609 performing IPsec with a VPN SPA.

Figure 15 Test Bed Topology

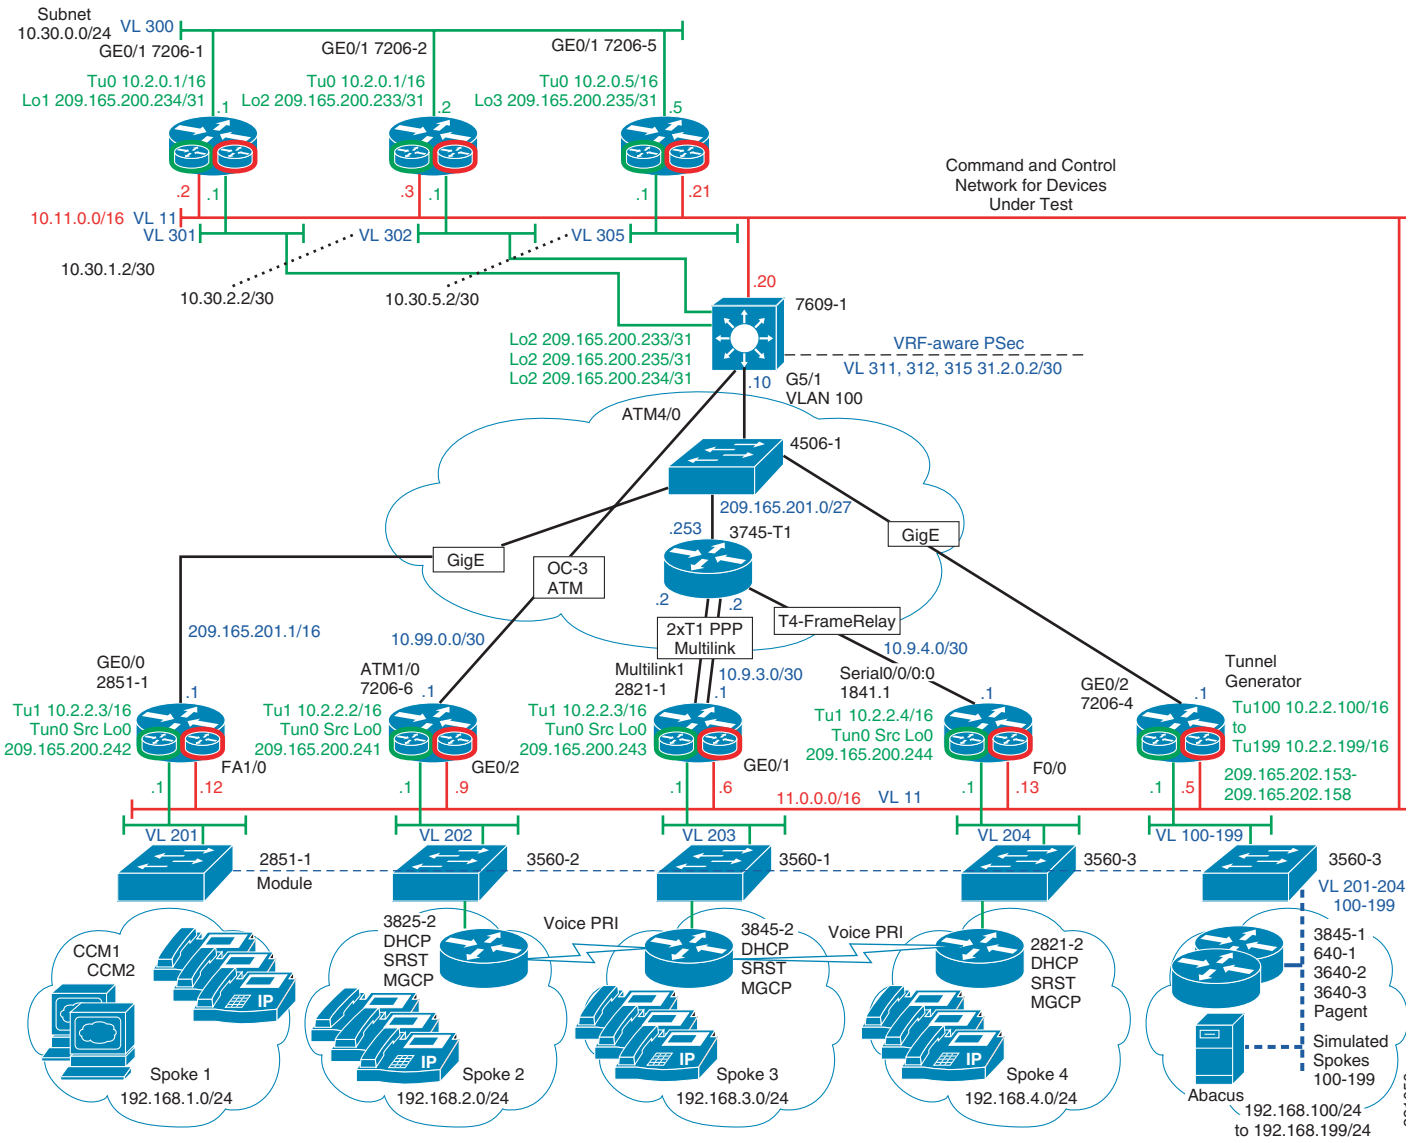

The Cisco 7206, Cisco 2811, and Cisco 1841 spoke routers are deployed on the customer DMVPN network. The MGCP gateways are the Cisco 3845, Cisco 3825, Cisco 2851, and Cisco 2821. In addition, a single Cisco 7206 was used with a multi-VRF configuration to simulate 100 separate spokes. The details of the platforms and their interfaces are listed in Table 8.

Other test environment notes area as follows:

- Routing metrics were used to designate the primary path via the primary service provider network.
- Firewalls are managed by the service provider and are transparent (no bottleneck).
- NAT functionality is not part of the test environment.

Table 8 Spoke Router Details

Spoke Description	Voice Topology Function	Platform	Interface to Service Provider WAN	Crypto Module
Spoke 1	Headquarters	Cisco 2851	OC3	Built-in
Spoke 2	Remote	Cisco 7206	4xT1 ATM IMA	VAM2+
Spoke 3	Remote	Cisco 2821	4xT1 PPP Multilink	AIM-VPN/EPII-PLUS
Spoke 4	Remote	Cisco 1841	T1 Frame Relay	AIM-VPN/BPII-PLUS
Spokes 100–199 (simulated)	Remote	Cisco 7206	Gigabit Ethernet	VAM2+

Voice Traffic Model

The voice modeled for the test effort used the following assumptions:

- IP phone busy hour call attempts (BHCA) four calls per hour; phone in use 15 percent of time (that is, 2.25 minutes per call with 12.75 minutes idle time between calls)
- Unity port BHCA 20 calls per hour, at 80 percent utilization
- MGCP gateway port BHCA 20 calls per hour, at 80 percent utilization

No more than 50 percent of calls are intersite calls, and intersite traffic is evenly distributed across all remote sites.

Because the *intrasite* traffic is not simulated, the number of simulated IP phones is 4500/2 = 2250, with 100 percent utilization and the same call duration. This reduces the required Cisco Unified CallManager server capacity of the test bed.

Given 100 percent utilization and a call duration of 2 minutes 15 seconds, the simulated phone attempts to make a call to another spoke 26.6 times per hour with no idle time between call attempts. This results in the same amount of traffic without having to simulate the full 30,000 phones on the bulk call generator.

Firmware/Software

The software and platforms included in this test effort were selected to match the equipment deployed in the customer DMVPN network. The Cisco software release for the ISR and 7206 spoke routers is 12.4(9)T2. Cisco Unified CallManager 4.2(3) was used for the test because it was the latest released version of Windows-based Cisco Unified CallManager. The Cisco ISRs and Cisco 7206 routers used Cisco IOS SW release 12.4(9)T2.

Table 9 lists the software images tested for each device in the network.

Table 9 Software Images Tested

Device	Image
Hubs	
Cisco 7206	C7200-adventerprisek9-mz.124-9.T2
Cisco 7609	S72033-ipservicesk9_wan-mz.122-18.SXF6

Table 9 Software Images Tested (continued)

Switches	
Cisco 3560	c3560-ipbase-mz.122-25.SEB4
Cisco 4506	cat4000-i9k91s-mz.122-18.EW1
Spoke Routers	
Cisco 2851	C2800nm-adventerprisek9-mz.124-9.T2
Cisco 1841	C1841-adventerprisek9-mz.124-9.T2
Cisco 2821	C2800nm-adventerprisek9-mz.124-9.T2
Gateways	
Cisco 3845	c3845-adventerprisek9-mz.124-9.T2
Cisco 3825	c3825-adventerprisek9-mz.124-9.T2
Cisco Unified CallManager	4.2(3)sr1
Endpoints	
Cisco 7902	CP7902080001SCCP051117A
Cisco 7912	CP7912080001SCCP051117A
Cisco 7960	P00308000100
Cisco 7961	SCCP41.8-0-2SR1S
Cisco 7961G-GE	SCCP41.8-0-2SR1S
Cisco 7970	SCCP70.8-0-2SR1S
Cisco 7971	SCCP70.8-0-2SR1S
Cisco 7985	CMTERM_7985.4-0-3-0
Cisco ATA 186	ATA030203SCCP051201A
Cisco IP Communicator	1.1.0.4

Test Coverage and Results

The following sections describe the test coverage and results.

Coverage

The features in the following subsections were tested.

Functional

- Individual calls to each spoke
- Calls on each type of endpoint with manual speech quality verification
- Observation of normal operation of DMVPN functions
  - Static tunnel creation
  - NHRP lookups

  - IPsec, ISAKMP SA establishment
  - Spoke-to-spoke tunnel creation
  - ISAKMP keepalives
  - NHRP cache refreshes
- Call variations
  - Codecs
    - G.729
    - G.711
  - Durations
  - T.38 fax
- Each endpoint
  - Various IP phone models
  - VT Advantage
  - MGCP gateway port (T1 PRI channel, FXS)
  - Unity voice port
- Call features
  - 3-way call over WAN
  - Call forwarding no answer to another spoke
  - Hold/resume over WAN
  - Attended/blind transfer to another spoke
- Signaling features
  - Encrypted call control TLS
  - SCCP
  - MGCP
- Media features
  - Encrypted speech SRTP
  - Conference bridge
  - Transcoding
- Network features
  - IP SLA monitoring
- DMVPN interactions with voice
  - Path switch
  - NHRP cache refresh and tunnel persistence throughout the call
  - Hub tunnel failure
  - Spoke-to-spoke tunnel failure
  - Latency/jitter increases
  - Packet loss increases
  - Multilink partial failure

  - Uplink total failure
- QoS/CAC
  - Normal operation with service provider QoS config (LLQ for voice, CBWFQ for remainder)
  - Overload of background traffic
  - Mismatch of CAC bandwidth overload of voice traffic
  - CAC interaction with call forwarding
    - Verify bandwidth returned when call forwarding no answer so that another call can use that bandwidth
    - No bandwidth allocated for call forwarding busy
  - CAC interaction with music on hold
    - No bandwidth for music on hold
- NHRP
- Static association to another spoke
  - SRST call control
    - Failover from CCM
    - Failback to CCM

Scalability and Traffic

- Solution scalability
  - Normal operation of low-end scenario
    - 1000 users, 3 SRST sites (each < 200 users)
    - Background traffic
    - Voice call type mix
    - Automatic alternate routing (AAR) to PSTN because of CAC
    - Call volume between sites as defined
  - Normal operation of high-end scenario
    - 30,000 users, 100 SRST sites (each <600 users)
    - Other items same as above
- Components
  - Number of spoke-to-spoke tunnels
  - Max voice capacity of each target spoke router platform
    - Includes typical QoS config and background data traffic
    - Identify CPU, bandwidth or other limits
  - Establish calculations for spoke router platform and bandwidth required for various anticipated call loads (snapshot measurements)
- Background data
  - IMIX

Spoke Router Performance

- Baseline
  - Spoke-to-spoke traffic mix at various rates
  - Measure voice quality, call completion, CPU
  - SP deployed interface (specific one for each platform)
  - GigE interface (to allow comparison of results)
- Voice quality
- Platform capacity
- Range of platforms
- Spoke-to-spoke tunnel scalability
  - Tunnel setup rate/time
  - CPU/memory/bandwidth/and so on, utilization
- PPS scalability (voice call capacity)
- Multilink overhead

Stress and Negative

- Stability with high traffic (call setup rate is not as important as bandwidth and PPS)
  - High call traffic
    - Automatic alternate routing (AAR) to PSTN because of CAC
  - Overload background traffic
- DMVPN events impacting voice
  - IKE admission control exceeded
  - System resource CAC exceeded
  - NHRP
    - Manual NHRP cache clearing
    - Synchronize network to achieve peak NHRP refresh traffic rate
      - Within max NHRP packet rate
      - Exceed NHRP packet rate
  - IPsec failure during calls
  - ISAKMP failure during calls
  - Spoke-hub tunnel failure during calls (route traffic to secondary DMVPN network)
- Unstable network with high traffic
  - High jitter
  - High packet loss
  - Platform reloads and physical link failures
- Tunnel creation on every call with 100 sites (allow tunnels to time out between calls)
  - Each platform

Long Duration/Endurance

- Long duration calls
- Long duration traffic run
  - Cyclic traffic volume
    - Include some long duration calls
    - Fluctuate background data traffic from low to high (higher periodic rate than the voice traffic) for example, data traffic peaks and subsides once per 5 minutes call traffic peaks and subsides once per hour

MPLS WAN Simulation Scenarios

- Repeat specific cases with NSITE DMVPN test bed with MPLS WAN

Management

- Sanity check IOS CLI commands to monitor DMVPN, ISAKMP, IPsec, QoS

Results

The use of the best practices described earlier resulted in consistent performance and very good speech quality. Performance metrics were collected and are summarized in Performance Information, page 41.

There were no blocking issues found that would invalidate interoperation of Cisco Unified Communications and DMVPN. All tests were passed except the ISAKMP and System CAC tests, which failed because of defect CSCsg36532. This issue is fixed in 12.4(11)T. All defects found during the testing are listed in Table 10.

Table 10 Defects Found During Testing

Defect	Description
CSCsh24984	CAC: Not enough bandwidth—prompt duration issues
CSCsh26690	CAC: Bandwidth not released when CFNA to busy user
CSCse70541	DMVPN debugs are turned on without enabling them
CSCsh22016	Clear IP NHRP counters also clears NHRP cache
CSCsg36532	DMVPN Phase 2: Black hole traffic when spoke-spoke tunnel fails

Performance Information

The information collected for this document is mainly associated with spoke router performance for spoke-to-spoke tunnels. See the DMVPN Design Guide for more information on DMVPN hub technologies and scalabilities. In addition, the DMVPN Design Guide provides branch office scalability results including the supported number of voice calls and throughput per ISR router with hardware encryption.

The following sections contain the results of various performance measurements taken during the testing of the DMVPN spoke routers typically deployed on the customer network.

CPU Load versus Throughput

Figure 16 and Figure 17show the throughput of the selected spoke router platforms in terms of bits per second and packets per second. The purpose of this measurement is to compare the capacities of the three platforms.

Figure 16 Mbps Throughput By Platform At 80% CPU Occupancy

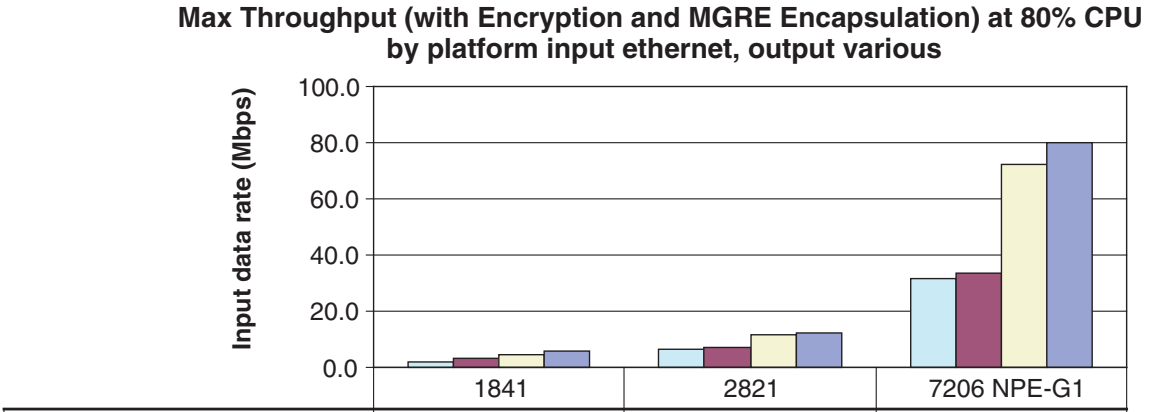

	1841	2821	7206 NPE-G1
50% Imix + 50% G.729 (WAN intf.)	1.9	6.6	31.3
50% Imix + 50% G.729 (GigE)	3.1	7.0	33.3
100% Imix (WAN intf.)	4.3	11.6	72.5
100% Imix (GigE)	5.7	12.1	79.9

221057

Figure 17 PPS Throughput By Platform At 80% CPU Occupancy

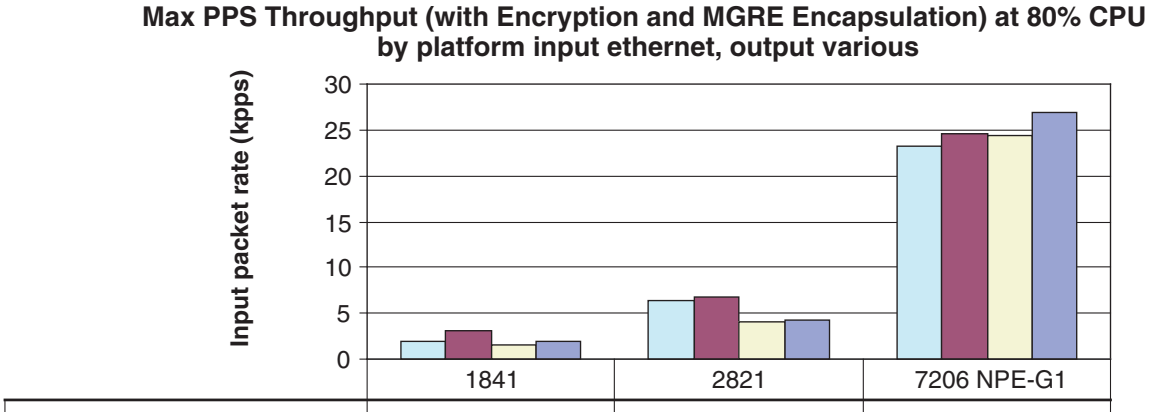

	1841	2821	7206 NPE-G1
50% Imix + 50% G.729 (WAN intf.)	2	6	23
50% Imix + 50% G.729 (GigE)	3	7	25
100% Imix (WAN intf.)	2	4	24
100% Imix (GigE)	2	4	27

221058

Tunnel Overhead for Idle Tunnels

The test results indicate that idle or very low traffic spoke-to-spoke tunnels do not consume significant CPU resources. Table 11 lists the data collected with various numbers of spoke-to-spoke tunnels active but passing minimal traffic. For 100 spoke-to-spoke tunnels, the CPU load is 1 percent or less and the process memory utilization is about 1 MB total or below 10 Kb per tunnel. The CPU load indicates that spoke-to-spoke tunnels can be maintained by IP SLA probes or other background traffic in instances where the path switchover causes unacceptable VoIP quality.

Table 11 Idle Tunnel Resource Utilization

Platform	Metric	0 tunnels	1 tunnel	10 tunnels	50 tunnels	100 tunnels
7206 NPE_G1	CPU %	0	0	0	0	0
	Process memory used	41629256	41635104	41714620	42102704	42565936
	Process memory per tunnel					9367
2821	CPU %	0	0	0	0	1
	Process memory used	23016956	22952864	23037952	23421140	23891316
	Process memory per tunnel					8744
1841	CPU %	0	0	0	0	1
	Process memory used	20172660	20178516	20179924	20579412	20976628
	Process memory per tunnel					8040

Tunnel Overhead for Active Tunnels

The purpose of measuring the overhead for active tunnels is to estimate the resource utilization of active spoke-to-spoke tunnels when selecting a platform to be a spoke router. Figure 18 compares the CPU occupancy at various packet rates over a single spoke-to-spoke tunnel versus 100 spoke-to-spoke tunnels. The result is that 100 active spoke-to-spoke tunnels increases CPU occupancy 10–15 percent compared with the same traffic over a single tunnel. Similar results were observed with the 1841 and 2821 platforms.

Figure 18 CPU Occupancy Comparison

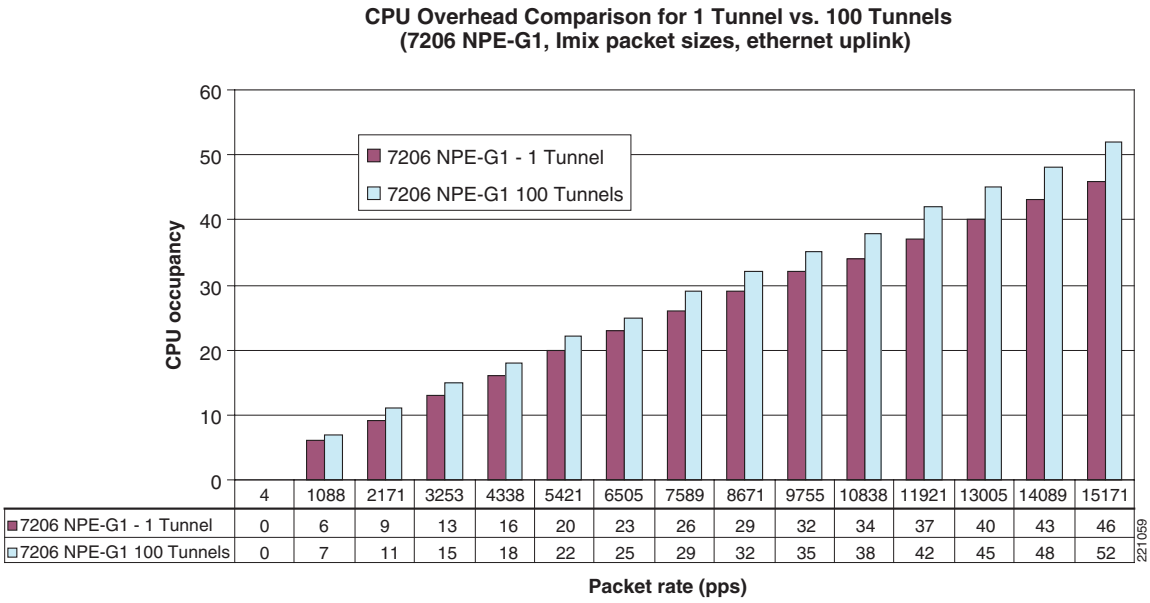

PPS versus BPS

Figure 19 shows the relationship between packet rates and CPU utilization. The CPU utilization increases with the packet rate for a fixed bit rate with variable packet sizes.

Figure 19 CPU Utilization versus Output Packet Rate

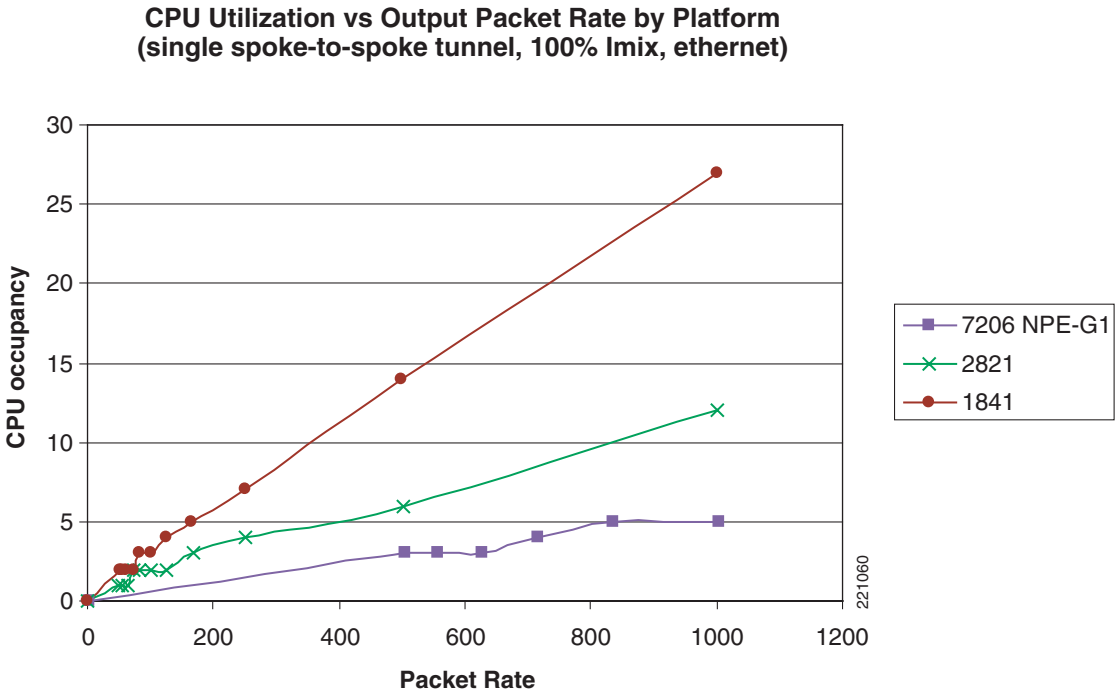

Tunnel Setup Load

Cisco measured the CPU utilization for spoke-to-spoke tunnel setup processing, which includes NHRP, ISAKMP, and IPsec operations minus the CPU load for sending the same traffic over a single established tunnel. Table 12 and Figure 20 show the peak 5-second CPU occupancy for a burst of tunnel setups at pre-determined rates. The traffic used to trigger each tunnel was 500 packets per second. The data do not factor out the 10–15 percent overhead for transmitting traffic over 100 tunnels versus 1 tunnel; however, the overhead at this packet rate is no more than 1–2 percentage points (see Figure 20).

Table 12 CPU Occupancy Because of Spoke-To-Spoke Tunnel Setup

Tunnel Setup Packet Interval	Tunnel setup rate (tunnels per second)	Peak 5 sec CPU with 500 pps traffic			Peak 5 sec CPU occupancy minus CPU occupancy because of traffic		
		7206	2821	1841	7206	2821	1841
never	0.0	9	10	20	0	0	0
2000	0.5	10	12	22	1	2	2
1000	1.0	11	14	24	2	4	4
500	2.0	12	17	28	3	7	8
333	3.0	13	19	33	4	9	13
250	4.0	14	22	36	5	12	16
125	8.0	18	33	52	9	23	32
100	10.0	20	43	62	11	33	42

Figure 20 CPU Occupancy versus Tunnel Setup Rate

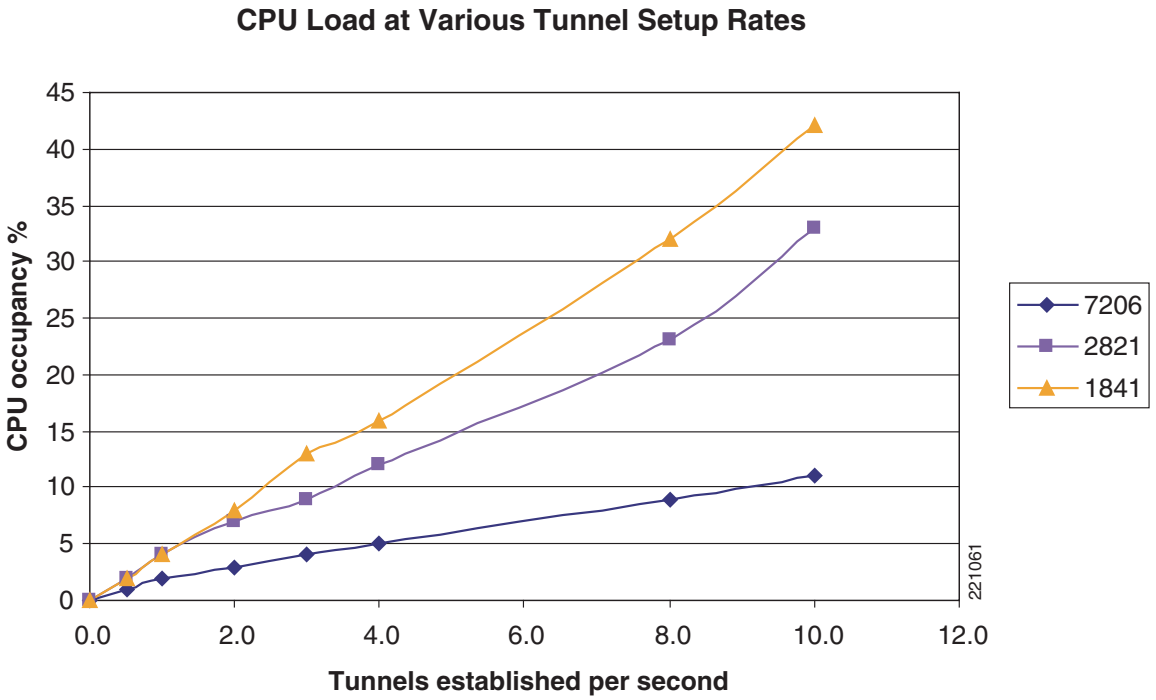

Serial Interface Overhead

The information shown in Figure 21, Figure 22, and Figure 23 helps estimate the overhead associated with the physical interfaces and data link protocols in use on the DMVPN spoke routers.

Figure 21 7206 NPE-G1 GigE versus ATM OC3

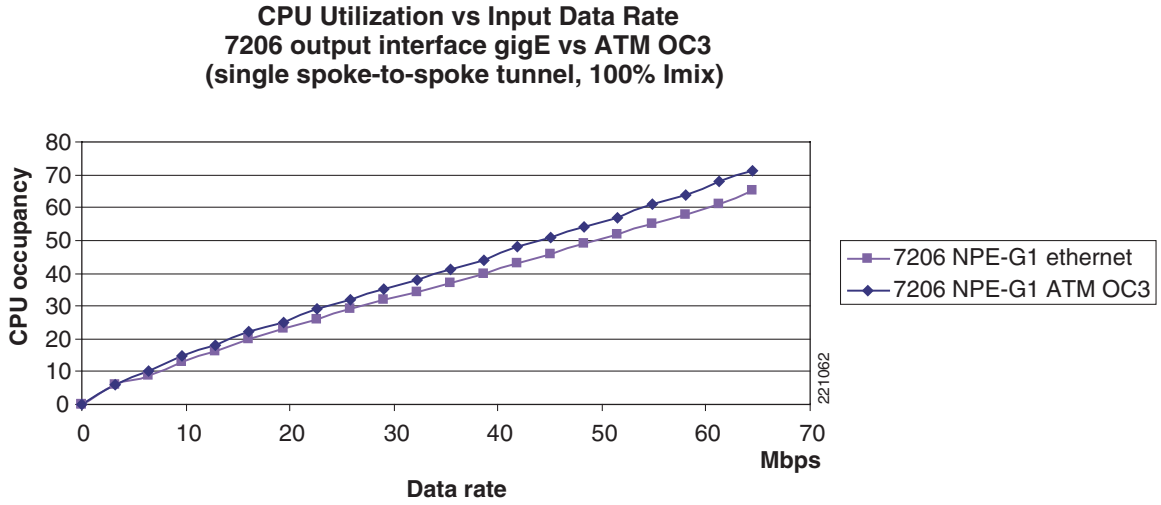

Figure 22 2821 GigE versus 2xT1 PPP Multipoint

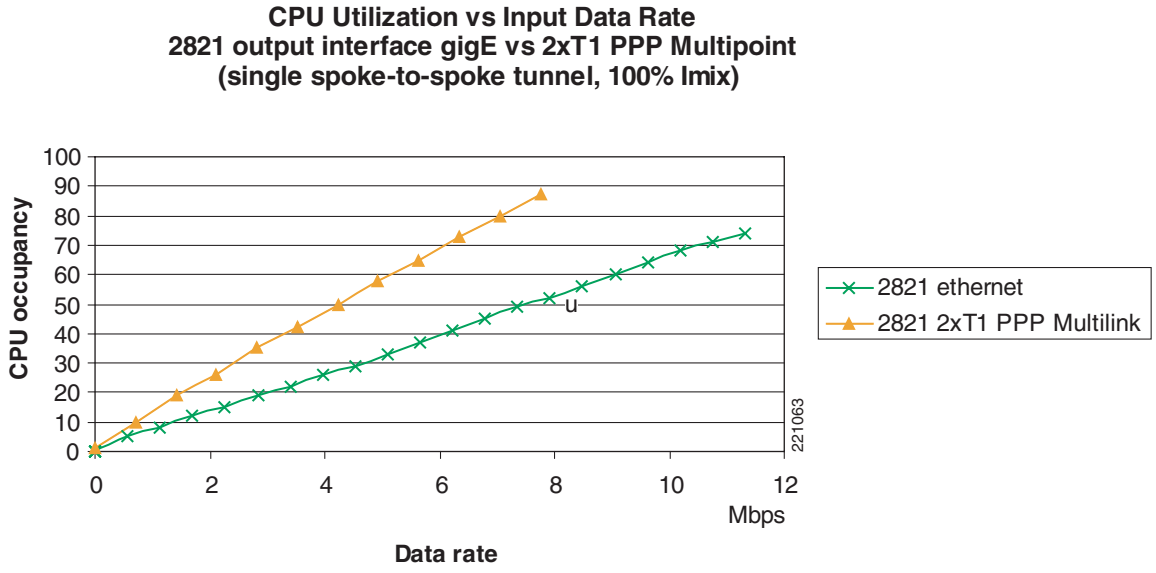

Figure 23 1841 FE versus T1 Frame Relay

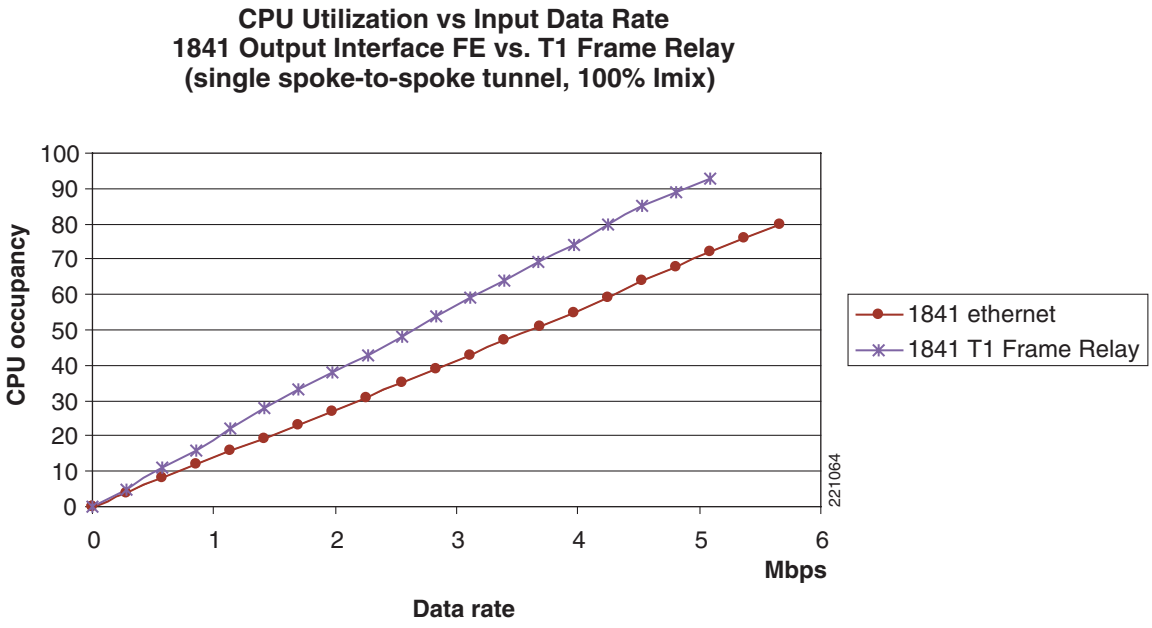

Path Switch Distortion

There is the potential that a small amount of audible distortion can occur when the spoke-to-spoke tunnel is initially established during an active IP phone call. Cisco performed tests to assess the severity of the audible distortion and the ability of Cisco IP phones to conceal the distortion.

Conditions were established similar to the customer DMVPN network, including a variety of SCCP IP phones controlled by a Cisco Unified CallManager cluster. The phones were connected to two different LANs with access to the DMVPN WAN. A network simulator controlled latency between the DMVPN network and the terminating spoke router.

A phone call was made with various latency values simulated for either the spoke-to-spoke path or the spoke-hub-spoke path. A speech sample was played into the originating phone and recorded at the terminating side. When path switch occurred to the spoke-to-spoke path, any resulting audible distortion was rated with a "momentary distortion severity" value on the scale listed in Table 13.

Table 13 Momentary Distortion Scale

Opinion	MDS Score	Customer Reaction
No distortion	0	None
Barely noticeable	1	Ignores
Noticeable	2	Comments about noise
Disruptive (lost words or multiple noticeable distortions)	3	Asks to repeat words
Severely disruptive (multiple lost words)	4	Re-originates call

The results of this test (see Figure 24) indicate that little or no audible distortion is present if the difference in latency between the spoke-hub-spoke path and the spoke-to-spoke path is 100 ms or less. Furthermore, if the difference in latency is 50 ms or less, no distortion is audible. For the customer DMVPN network, the latency difference between paths, estimated between 30 ms and 60 ms, does not impact voice quality.

Figure 24 Path Switch Distortion Results

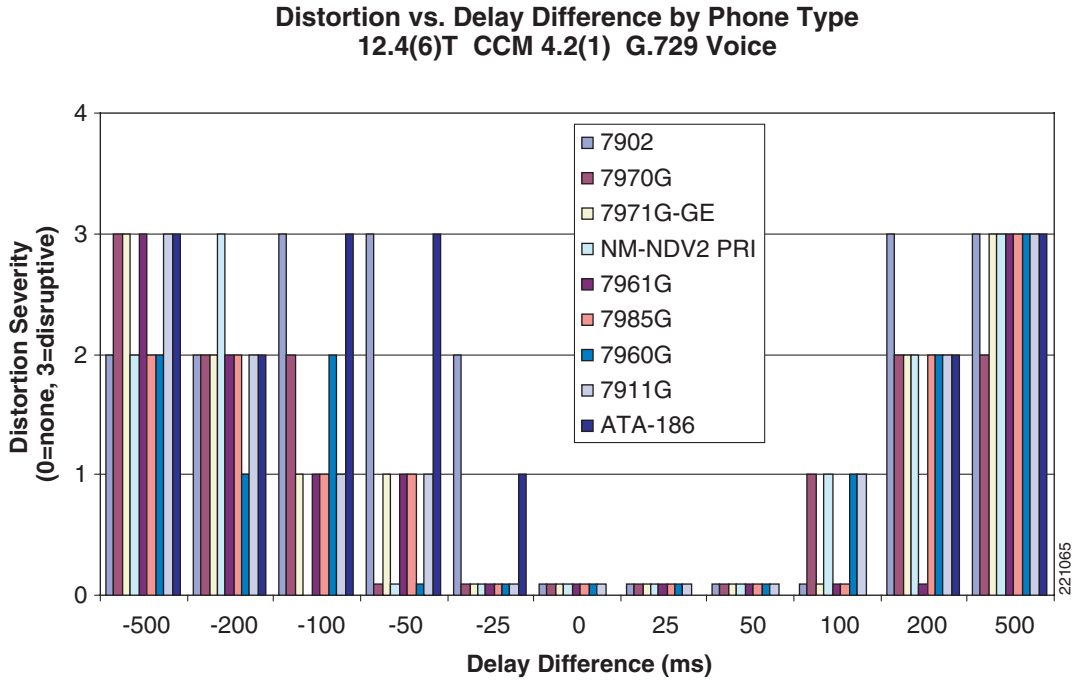

The following results were observed:

- Depending on the load on the hub and spoke routers, the path switch usually occurs within the first 500 ms after the media stream begins flowing.
- The distortion, when present, lasted much less than one second. In the most common case where the spoke-to-spoke path is faster than the spoke-hub-spoke path, the distortion can be described as a “splice” in which a very small portion of the speech (a fraction of a second) was discarded when the phone continues to play out only the packets received from the faster path.
- For calls between phones on the same cluster, the ringback tone is generated locally on the phone of the originating party, so the duration of the alerting phase of the call is not relevant and this tone does not trigger a tunnel.
- Because spoke-to-spoke tunnels are always established between the main site housing the Cisco Unified CallManager servers and any remote site, there was never any distortion because of path switching for calls between those sites.
- There is very little difference in the momentary distortion severity between calls using G.729 and G.711. There is a difference in the sound of the distortion but neither is more noticeable for low-to-moderate latency differences.

Based on the information in this section, Cisco recommends that you measure the WAN performance using IP SLA probes or another method to confirm that the difference between spoke-hub-spoke and spoke-to-spoke latencies is 100 ms or less. For more information on configuring the IP SLA function, see IP SLA Options Relevant to Voice Over DMVPN, page 29.

Typical IP Phone Statistics

Cisco IP phones maintain useful statistics related to the current call. Figure 25 shows an example of this information during a typical call between two branch offices over a spoke-to-spoke DMVPN tunnel. This information was obtained using the HTTP interface on a Cisco 7970G IP phone. The same information can be retrieved directly on the phone by pressing the help key twice or selecting call statistics from the setup menu.

Figure 25 IP Phone Statistics Example

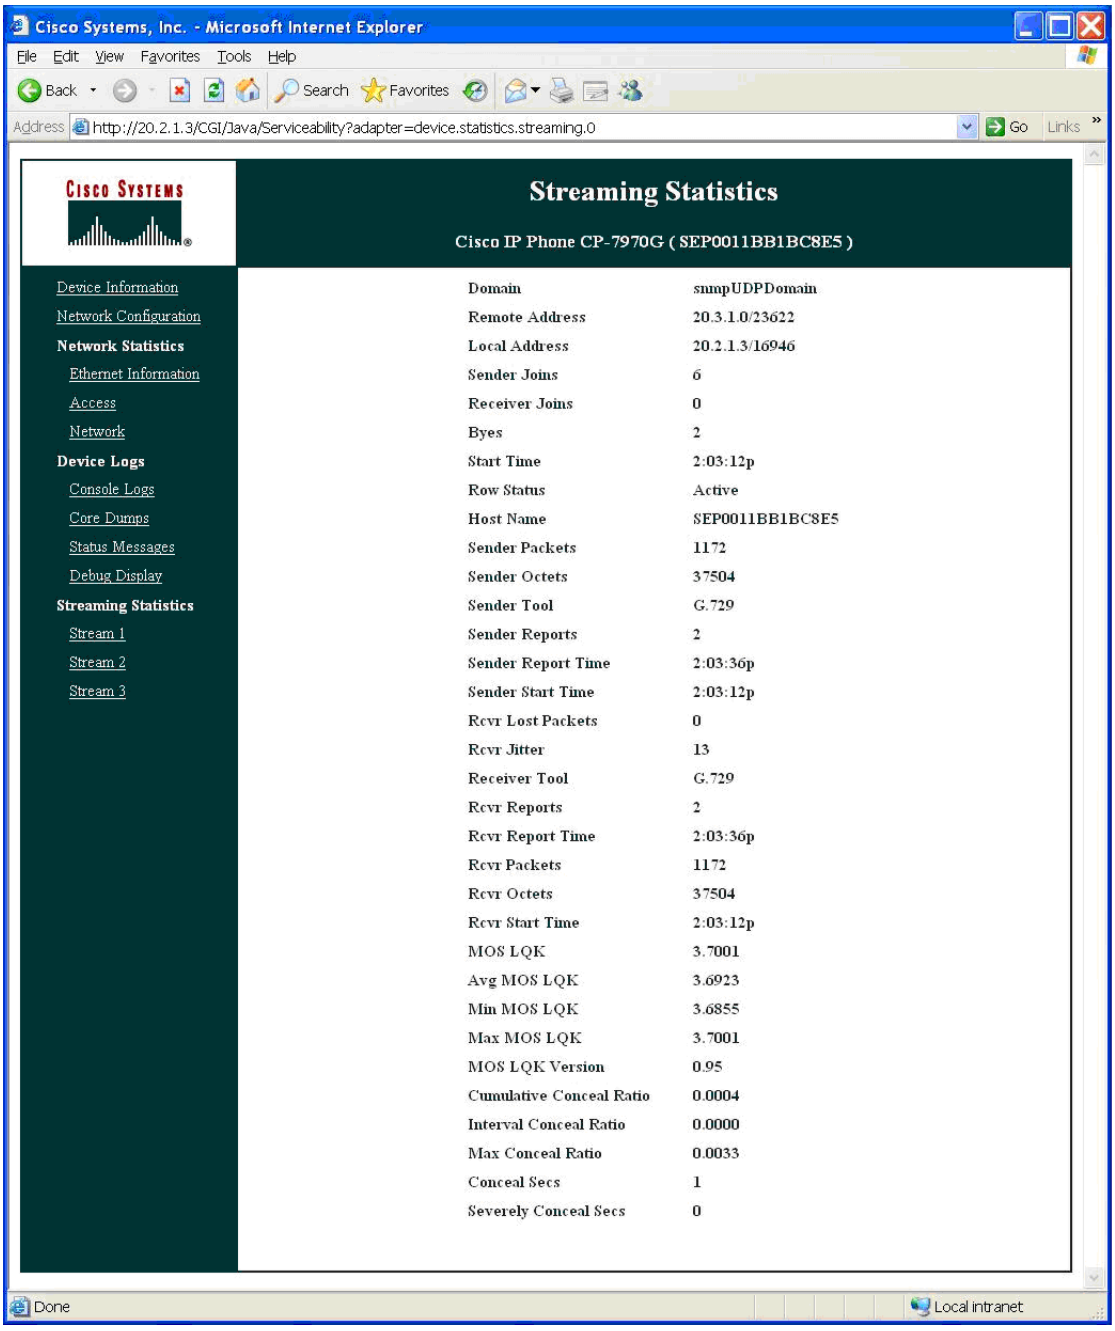

Example Configuration Files

This section contains example configuration files for the test bed routers used for the testing reported in this document.

Hub Routers Configuration

Cisco 7206 Hub Router Configuration

The following is the Cisco 7206 DMVPN hub configuration used for this testing:

```
version 12.4
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
!
hostname 7206-1
!
boot-start-marker
boot system disk2:c7200-adventerprisek9-mz.124-9.T2
boot-end-marker
!
logging buffered 1000000 debugging
no logging console
enable password lab
!
no aaa new-model
!
resource policy
!
clock timezone EST -5
clock summer-time EDT recurring
ip cef
!
!
ip vrf user
 rd 10:100
!
no ip domain lookup
ip domain name agency.com
ip multicast-routing
!
!
controller ISA 1/1
!
interface Tunnel0
 description dmvpn for user traffic
 ip vrf forwarding user
 ip address 10.2.0.1 255.255.0.0
 no ip redirects
 ip mtu 1416
 no ip next-hop-self eigrp 1
 ip nhrp authentication myGREkey
 ip nhrp map multicast dynamic
 ip nhrp map 10.2.0.2 209.165.200.233
 ip nhrp map multicast 209.165.200.233
 ip nhrp network-id 99
 ip nhrp holdtime 300
 ip nhrp nhs 10.2.0.2
```

```
 ip nhrp max-send 400 every 10
 ip nhrp registration timeout 75
 no ip split-horizon eigrp 1
 load-interval 30
 delay 1000
 tunnel source Loopback1
 tunnel mode gre multipoint
 tunnel key 123456
!
interface Tunnel1
 description dmvpn for mgmt traffic
 ip address 10.3.0.1 255.255.0.0
 no ip redirects
 ip mtu 1416
 ip nhrp authentication myGREkey
 ip nhrp map multicast dynamic
 ip nhrp map multicast 209.165.202.129
 ip nhrp map 10.3.0.2 209.165.202.129
 ip nhrp network-id 100
 ip nhrp holdtime 300
 ip nhrp nhs 10.3.0.2
 ip nhrp server-only
 ip nhrp max-send 400 every 10
 ip nhrp registration timeout 75
 ip nhrp cache non-authoritative
 no ip split-horizon eigrp 1
 load-interval 30
 delay 1000
 tunnel source Loopback21
 tunnel mode gre multipoint
 tunnel key 654321
!
interface Loopback1
 description Match with 7609-1 Lo1
 ip address 209.165.200.234 255.255.255.254
!
interface Loopback10
 ip address 209.165.200.249 255.255.255.255
!
interface Loopback21
 description Mgmt Data Match with 7609-1 Lo21
 ip address 209.165.202.130 255.255.255.254
!
interface GigabitEthernet0/1
 no ip address
 duplex full
 speed 1000
 media-type rj45
 no negotiation auto
!
interface GigabitEthernet0/1.300
 encapsulation dot1Q 300
 ip address 10.30.0.1 255.255.255.0
!
interface GigabitEthernet0/1.301
 description P2P to 7609-1
 encapsulation dot1Q 301
 ip address 10.30.1.1 255.255.255.252
!
interface GigabitEthernet0/1.3021
 description P2P to 7609-1
 encapsulation dot1Q 3021
 ip address 10.30.21.1 255.255.255.252
!
```

```
interface GigabitEthernet0/2
 no ip address
 shutdown
!
interface GigabitEthernet0/2.11
 description Labnet Management and tftp access
 encapsulation dot1Q 11
 ip address 10.11.0.2 255.255.0.0
!
interface GigabitEthernet0/2.12
 description ISP Mgmt Traffic
 encapsulation dot1Q 12
 ip address 209.165.200.226 255.255.255.248
 standby 1 ip 209.165.200.225
 standby 1 preempt delay minimum 120
!
interface GigabitEthernet0/3
 no ip address
 shutdown
!
router eigrp 1
 passive-interface default
 no passive-interface Tunnel0
 no passive-interface Tunnel1
 network 10.3.0.0 0.0.255.255
 network 209.165.200.224 0.0.0.31
 network 209.165.200.249 0.0.0.0
 no auto-summary
 !
 address-family ipv4 vrf user
 network 10.2.0.0 0.0.255.255
 no auto-summary
 autonomous-system 1
 exit-address-family
!
ip route 0.0.0.0 0.0.0.0 10.30.21.2
ip route 209.165.200.233 255.255.255.255 10.30.0.2
ip route 209.165.202.129 255.255.255.255 10.30.0.2
ip route 209.165.200.240 255.255.0.0 10.30.1.2
ip route 209.165.202.136 255.255.255.248 10.30.21.2
ip route 209.165.202.144 255.255.0.0 10.30.1.2
ip route 209.165.202.152 255.255.0.0 10.30.21.2
no ip http server
no ip http secure-server
!
!
ip sla responder
logging alarm informational
!
control-plane
!
gatekeeper
 shutdown
!
line con 0
 exec-timeout 0 0
 stopbits 1
line aux 0
 stopbits 1
line vty 0 4
 exec-timeout 0 0
 password lab
 login
!
```

```
scheduler heapcheck process memory processor io checktype lite-chunks
ntp clock-period 17180039
ntp server 10.11.0.1
!
end
```

7609 IPsec Concentrator Configuration

The following is the Cisco 7609 IPsec concentrator configuration used for this testing:

```
upgrade fpd auto
version 12.2
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
service counters max age 5
!
hostname 7609-1
!
boot system flash disk0:s72033-ipservicesk9_wan-mz.122-18.SXF6
logging buffered 1000000 debugging
no logging console
enable password lab
!
no aaa new-model
clock timezone EST -5
clock summer-time EDT recurring
ip subnet-zero
!
ip vrf vpn1
 rd 10:1
!
ip vrf vpn2
 rd 10:2
!
ip vrf vpn21
 rd 10:21
!
ip vrf vpn22
 rd 10:22
!
no ip domain-lookup
ipv6 mfib hardware-switching replication-mode ingress
vtp domain mynet
vtp mode transparent
mls ip multicast flow-stat-timer 9
no mls flow ip
no mls flow ipv6
no mls acl tcam share-global
mls cef error action freeze
!
crypto engine mode vrf
!
!
redundancy
 mode sso
 main-cpu
  auto-sync running-config
spanning-tree mode pvst
no spanning-tree optimize bpdu transmission
spanning-tree extend system-id
diagnostic cns publish cisco.cns.device.diag_results
```

```
diagnostic cns subscribe cisco.cns.device.diag_commands
port-channel per-module load-balance
!
vlan internal allocation policy ascending
vlan access-log ratelimit 2000
!
vlan 11,100,301-302,305,311-312,315,3021-3022,3025,3121-3122,3125
!
!  This QOS config is for the ATM link to spoke 2
!
class-map match-any PCLASS_COS2_SAA
  match  dscp 27
class-map match-any PCLASS_COS3_SAA
  match  dscp 19
class-map match-any PCLASS_COS1_SAA
  match  dscp 47
class-map match-any PCLASS_COS4_SAA
  match  dscp 19
class-map match-any PCLASS_COS4
  match not  dscp 1
class-map match-any QCLASS_COS4
  match not  dscp 1
class-map match-any PCLASS_COS2
  match  dscp af31
class-map match-any QCLASS_COS3
  match  dscp 19
  match  dscp af21
class-map match-any PCLASS_COS3
  match  dscp af21
class-map match-any QCLASS_COS2
  match  dscp 27
  match  dscp af31
class-map match-any QCLASS_COS1
  match  dscp 47
  match  dscp ef
class-map match-any PCLASS_COS1
  match  dscp ef
class-map match-any PCLASS_NM
  match access-group 180
class-map match-any PCLASS_RP
  match  dscp cs6
class-map match-any QCLASS_NM
  match  dscp cs6
  match access-group 180
class-map match-any Prec_3
  match ip precedence 3
class-map match-any Prec_2
  match ip precedence 2
class-map match-any Prec_1
  match ip precedence 1
class-map match-any Prec_0
  match ip precedence 0
class-map match-any Prec_7
  match ip precedence 7
class-map match-any Prec_6
  match ip precedence 6
class-map match-any Prec_5
  match ip precedence 5
class-map match-any Prec_4
  match ip precedence 4
!
!
policy-map NM
  class PCLASS_RP
```

```
   police cir 1500000 bc 150000 be 150000 conform-action set-dscp-transmit cs6
exceed-action set-dscp-transmit cs6 violate-action set-dscp-transmit cs6
  class PCLASS_NM
   police cir 150000 bc 150000 be 150000 conform-action set-dscp-transmit af21
exceed-action set-dscp-transmit af21 violate-action set-dscp-transmit af21
policy-map COS2
  class PCLASS_COS2
   police cir 45000000 bc 450000 be 450000 conform-action set-dscp-transmit af31
exceed-action set-dscp-transmit af32 violate-action set-dscp-transmit af32
policy-map COS3
  class PCLASS_COS3
   police cir 30000000 bc 300000 be 300000 conform-action set-dscp-transmit af21
exceed-action set-dscp-transmit af22 violate-action set-dscp-transmit af22
policy-map WAN_150M
  class QCLASS_NM
    bandwidth remaining percent 10
    random-detect dscp-based
    random-detect dscp 18 100 200 10
    random-detect dscp 48 200 300 10
   service-policy NM
  class QCLASS_COS1
    priority 60000
  class QCLASS_COS2
    bandwidth remaining percent 60
    random-detect dscp-based
    random-detect dscp 26 200 300 10
   service-policy COS2
  class QCLASS_COS3
    bandwidth remaining percent 30
    random-detect dscp-based
    random-detect dscp 18 200 300 10
   service-policy COS3
policy-map COS4
  class PCLASS_COS4
   police cir 15000000 bc 150000 be 150000 conform-action set-dscp-transmit default
exceed-action set-dscp-transmit default violate-action set-dscp-transmit default
!
!
crypto keyring vpn22
  local-address Loopback22
  pre-shared-key address 0.0.0.0 0.0.0.0 key 123456
crypto keyring vpn21
  local-address Loopback21
  pre-shared-key address 0.0.0.0 0.0.0.0 key 123456
crypto keyring vpn2
  local-address Loopback2
  pre-shared-key address 0.0.0.0 0.0.0.0 key 123456
crypto keyring vpn1
  local-address Loopback1
  pre-shared-key address 0.0.0.0 0.0.0.0 key 123456
!
crypto isakmp policy 1
 encr 3des
 authentication pre-share
 group 2
crypto isakmp keepalive 60
crypto isakmp profile vpnprof2
   vrf vpn2
   keyring vpn2
   match identity address 0.0.0.0
   local-address Loopback2
crypto isakmp profile vpnprof1
   vrf vpn1
   keyring vpn1
```

```
   match identity address 0.0.0.0
   local-address Loopback1
crypto isakmp profile vpnprof21
   vrf vpn21
   keyring vpn21
   match identity address 0.0.0.0
   local-address Loopback21
crypto isakmp profile vpnprof22
   vrf vpn22
   keyring vpn22
   match identity address 0.0.0.0
   local-address Loopback22
!
crypto ipsec security-association idle-time 720
!
crypto ipsec transform-set 3DES esp-3des esp-sha-hmac
 mode transport
!
crypto dynamic-map dyn-mgre2 10
 set transform-set 3DES
 set isakmp-profile vpnprof2
!
crypto dynamic-map dyn-mgre1 10
 set transform-set 3DES
 set isakmp-profile vpnprof1
!
crypto dynamic-map dyn-mgre21 10
 set transform-set 3DES
 set isakmp-profile vpnprof21
!
crypto dynamic-map dyn-mgre22 10
 set transform-set 3DES
 set isakmp-profile vpnprof22
!
!
crypto map mgre2 local-address Loopback2
crypto map mgre2 200 ipsec-isakmp dynamic dyn-mgre2
!
crypto map mgre1 local-address Loopback1
crypto map mgre1 200 ipsec-isakmp dynamic dyn-mgre1
!
crypto map mgre22 local-address Loopback22
crypto map mgre22 200 ipsec-isakmp dynamic dyn-mgre22
!
crypto map mgre21 local-address Loopback21
crypto map mgre21 200 ipsec-isakmp dynamic dyn-mgre21
!
!
!
interface Loopback1
 description Match with HUB 7206-1
 ip address 209.165.200.234 255.255.255.254
!
interface Loopback2
 description Match with HUB 7206-2
 ip address 209.165.200.233 255.255.255.254
!
interface Loopback21
 description ISP Mgmt Match with HUB 7206-1
 ip address 209.165.202.130 255.255.255.254
!
interface Loopback22
 description ISP Mgmt Match with HUB 7206-2
 ip address 209.165.202.129 255.255.255.254
```

```
!
interface GigabitEthernet1/0/1
 description VPN SPA Inside Interface
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 1,311,312,1002-1005,3121,3122
 switchport mode trunk
 mtu 9216
 no ip address
 flowcontrol receive on
 flowcontrol send off
 spanning-tree portfast trunk
!
interface GigabitEthernet1/0/2
 description VPN SPA Outside Interface
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 1,1002-1005
 switchport mode trunk
 mtu 9216
 no ip address
 flowcontrol receive on
 flowcontrol send off
 spanning-tree portfast trunk
!
interface ATM4/0/0
 no ip address
 load-interval 30
!
interface ATM4/0/0.1 point-to-point
 ip address 10.99.0.2 255.255.255.252
 no ip redirects
 crypto engine subslot 1/0
 pvc 1/40
  cbr 149760
  encapsulation aal5snap
  service-policy out WAN_150M
 !
!
interface GigabitEthernet5/1
 description WAN interface to spokes
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 100
 switchport mode trunk
 no ip address
 load-interval 30
 spanning-tree portfast trunk
!
interface GigabitEthernet5/2
 description P2P vlans to 7200 hubs
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 11,301,302,3021,3022
 switchport mode trunk
 no ip address
 load-interval 30
 media-type rj45
 spanning-tree portfast trunk
!
interface GigabitEthernet6/1
 no ip address
 load-interval 30
 shutdown
```

```
 spanning-tree portfast trunk
!
interface GigabitEthernet6/2
 no ip address
 shutdown
!
interface GigabitEthernet6/3
 no ip address
 shutdown
!
interface GigabitEthernet6/4
 no ip address
 shutdown
!
interface GigabitEthernet6/5
 no ip address
 shutdown
!
interface GigabitEthernet6/6
 no ip address
 shutdown
!
interface GigabitEthernet6/7
 no ip address
 shutdown
!
interface GigabitEthernet6/8
 no ip address
 shutdown
!
interface Vlan1
 no ip address
 shutdown
!
interface Vlan11
 ip address 10.11.0.20 255.255.0.0
!
interface Vlan100
 ip address 209.165.201.10 255.255.255.224
 crypto engine subslot 1/0
!
interface Vlan301
 ip vrf forwarding vpn1
 ip address 10.30.1.2 255.255.255.252
 load-interval 30
!
interface Vlan302
 ip vrf forwarding vpn2
 ip address 10.30.2.2 255.255.255.252
 load-interval 30
!
interface Vlan311
 description private vlan for VPN-SPA use
 ip vrf forwarding vpn1
 ip address 10.31.1.2 255.255.255.252
 load-interval 30
 crypto map mgre1
 crypto engine subslot 1/0
!
interface Vlan312
 description private vlan for VPN-SPA use
 ip vrf forwarding vpn2
 ip address 10.31.2.2 255.255.255.252
 load-interval 30
```

```
 crypto map mgre2
 crypto engine subslot 1/0
!
interface Vlan3021
 ip vrf forwarding vpn21
 ip address 10.30.21.2 255.255.255.252
 load-interval 30
!
interface Vlan3022
 ip vrf forwarding vpn22
 ip address 10.30.22.2 255.255.255.252
 load-interval 30
!
interface Vlan3121
 description private vlan for VPN-SPA use
 ip vrf forwarding vpn21
 ip address 10.31.21.2 255.255.255.252
 load-interval 30
 crypto engine subslot 1/0
!
interface Vlan3122
 description private vlan for VPN-SPA use
 ip vrf forwarding vpn22
 ip address 10.31.22.2 255.255.255.252
 load-interval 30
 crypto engine subslot 1/0
!
ip classless
ip route 0.0.0.0 0.0.0.0 209.165.201.30
ip route 209.165.200.241 255.255.255.255 10.99.0.1
ip route 209.165.202.137 255.255.255.255 10.99.0.1
ip route vrf vpn1 0.0.0.0 0.0.0.0 Vlan311
ip route vrf vpn1 209.165.200.234 255.255.255.255 30.1.0.1
ip route vrf vpn2 0.0.0.0 0.0.0.0 Vlan312
ip route vrf vpn2 209.165.200.233 255.255.255.255 10.30.2.1
ip route vrf vpn21 0.0.0.0 0.0.0.0 Vlan3121
ip route vrf vpn21 209.165.202.130 255.255.255.255 10.30.21.1
ip route vrf vpn22 0.0.0.0 0.0.0.0 Vlan3122
ip route vrf vpn22 209.165.202.129 255.255.255.255 10.30.22.1
!
no ip http server
!
logging trap debugging
logging facility local2
logging 172.18.154.98
access-list 180 permit ip host 172.19.64.1 any
!
!
!
control-plane
!
!
!
dial-peer cor custom
!
!
!
alias exec sib show ip interf brief
alias exec pol show policy-map int atm 4/0/0.1
alias exec pv show policy-map int atm 4/0/0.1 | begin COS1
!
line con 0
 exec-timeout 0 0
line vty 0 4
```

```
 exec-timeout 0 0
 password lab
 login
!
!
ntp clock-period 17180095
ntp server 10.11.0.1
no cns aaa enable
end
```

Spoke Router Configuration

The following is the Cisco 7206 spoke router configuration used for this testing:

```
version 12.4
service timestamps debug datetime msec localtime show-timezone
service timestamps log datetime msec localtime show-timezone
no service password-encryption
!
hostname 7206-6
!
boot-start-marker
boot system disk2:c7200-adventerprisek9-mz.124-9.T2
boot-end-marker
!
logging buffered 1000000 debugging
no logging console
enable password lab
!
no aaa new-model
!
resource policy
!
clock timezone EST -5
clock summer-time EDT recurring
ip cef
!
ip vrf user
 rd 10:100
!
no ip domain lookup
ip domain name agency.com
!
class-map match-any PCLASS_COS2_SAA
 match  dscp 27
class-map match-any PCLASS_COS3_SAA
 match  dscp 19
class-map match-any PCLASS_COS1_SAA
 match  dscp 47
class-map match-any PCLASS_COS4_SAA
 match  dscp 19
class-map match-any PCLASS_COS4
 match not  dscp 1
class-map match-any QCLASS_COS4
 match not  dscp 1
class-map match-any PCLASS_COS2
 match  dscp af31
class-map match-any QCLASS_COS3
 match  dscp 19
 match  dscp af21
class-map match-any PCLASS_COS3
 match  dscp af21
```

```
class-map match-any QCLASS_COS2
 match  dscp 27
 match  dscp af31
class-map match-any QCLASS_COS1
 match  dscp 47
 match  dscp ef
class-map match-any PCLASS_COS1
 match  dscp ef
class-map match-any PCLASS_NM
 match access-group 180
class-map match-any PCLASS_RP
 match  dscp cs6
class-map match-any QCLASS_NM
 match  dscp cs6
 match access-group 180
class-map match-any Prec_3
 match ip precedence 3
class-map match-any Prec_2
 match ip precedence 2
class-map match-any Prec_1
 match ip precedence 1
class-map match-any Prec_0
 match ip precedence 0
class-map match-any Prec_7
 match ip precedence 7
class-map match-any Prec_6
 match ip precedence 6
class-map match-any Prec_5
 match ip precedence 5
class-map match-any Prec_4
 match ip precedence 4
!
!
policy-map remark
 class Prec_0
  set ip dscp default
 class Prec_1
  set ip dscp af21
 class Prec_2
  set ip dscp af21
 class Prec_3
  set ip dscp af31
 class Prec_4
  set ip dscp af31
 class Prec_5
  set ip dscp ef
 class Prec_6
  set ip dscp af21
 class Prec_7
  set ip dscp af21


policy-map NM
 class PCLASS_RP
  police cir 1500000 bc 150000 be 150000
    conform-action set-dscp-transmit cs6
    exceed-action set-dscp-transmit cs6
 class PCLASS_NM
  police cir 150000 bc 150000 be 150000
    conform-action set-dscp-transmit af21
    exceed-action set-dscp-transmit af21
policy-map COS1
 class PCLASS_COS1
  police cir 60000000 bc 600000
```

```
    conform-action set-dscp-transmit ef
    exceed-action drop
 class PCLASS_COS1_SAA
  set dscp ef
policy-map COS2
 class PCLASS_COS2
  police cir 45000000 bc 450000
    conform-action set-dscp-transmit af31
    exceed-action set-dscp-transmit af32
 class PCLASS_COS2_SAA
  set dscp af31
policy-map COS3
 class PCLASS_COS3
  police cir 30000000 bc 300000
    conform-action set-dscp-transmit af21
    exceed-action set-dscp-transmit af22
 class PCLASS_COS3_SAA
  set dscp af21
policy-map WAN_150M
 class QCLASS_NM
  bandwidth remaining percent 10
  random-detect dscp-based
  random-detect dscp 18   100   200   10
  random-detect dscp 48   200   300   10
  service-policy NM
 class QCLASS_COS1
  priority 60000
  service-policy COS1
 class QCLASS_COS2
  bandwidth remaining percent 60
  random-detect dscp-based
  random-detect dscp 26   200   300   10
  service-policy COS2
 class QCLASS_COS3
  bandwidth remaining percent 30
  random-detect dscp-based
  random-detect dscp 18   200   300   10
  service-policy COS3
policy-map COS4
 class PCLASS_COS4
  police cir 15000000 bc 150000
    conform-action set-dscp-transmit default
    exceed-action set-dscp-transmit default
 class PCLASS_COS4_SAA
  set dscp default
!
!
crypto keyring test
  pre-shared-key address 0.0.0.0 0.0.0.0 key 123456
!
crypto isakmp policy 1
 encr 3des
 authentication pre-share
 group 2
crypto isakmp keepalive 60 periodic
crypto isakmp profile vpn1
   keyring test
   match identity address 0.0.0.0
crypto isakmp profile vpn2
   keyring test
   match identity address 0.0.0.0
!
!
crypto ipsec transform-set 3DES esp-3des esp-sha-hmac
```

```
 mode transport
!
crypto ipsec profile vpnprof1
 set security-association idle-time 720
 set transform-set 3DES
!
crypto ipsec profile vpnprof2
 set security-association idle-time 720
 set transform-set 3DES
!
!
!
!
!
!
interface Tunnel0
 bandwidth 1000
 ip vrf forwarding user
 ip address 10.2.2.2 255.255.0.0
 no ip redirects
 ip mtu 1416
 ip nhrp authentication myGREkey
 ip nhrp map 10.2.0.2 209.165.200.233
 ip nhrp map multicast 209.165.200.233
 ip nhrp network-id 99
 ip nhrp holdtime 300
 ip nhrp nhs 10.2.0.2
 ip nhrp cache non-authoritative
 delay 1000
 tunnel source Loopback0
 tunnel mode gre multipoint
 tunnel key 123456
 tunnel protection ipsec profile vpnprof1
!
interface Tunnel1
 description Mgmt Traffic
 bandwidth 1544
 ip address 10.3.2.2 255.255.0.0
 no ip redirects
 ip mtu 1416
 ip nhrp authentication myGREkey
 ip nhrp map 10.3.0.2 209.165.202.129
 ip nhrp map multicast 209.165.202.129
 ip nhrp network-id 100
 ip nhrp holdtime 300
 ip nhrp nhs 10.3.0.2
 ip nhrp cache non-authoritative
 load-interval 30
 delay 1000
 shutdown
 tunnel source Loopback1
 tunnel mode gre multipoint
 tunnel key 654321
 tunnel protection ipsec profile vpnprof2
!
interface Loopback0
 ip address 209.165.200.241 255.255.255.255
!
interface Loopback1
 ip address 209.165.202.137 255.255.255.255
 shutdown
!
interface Loopback10
 ip address 209.165.200.250 255.255.255.255
```

```
!
interface GigabitEthernet0/1
 no ip address
 shutdown
!
interface GigabitEthernet0/2
 no ip address
 duplex full
 speed 1000
 media-type rj45
 no negotiation auto
!
interface GigabitEthernet0/2.11
 description Labnet Management and tftp access
 encapsulation dot1Q 11
 ip address 10.11.0.23 255.255.0.0
!
interface GigabitEthernet0/2.202
 encapsulation dot1Q 202
 ip vrf forwarding user
 ip address 192.168.2.1 255.255.255.0
!
interface GigabitEthernet0/3
 no ip address
 shutdown
!
interface ATM1/0
 no ip address
 load-interval 30
 no atm ilmi-keepalive
!
interface ATM1/0.1 point-to-point
 ip address 10.99.0.1 255.255.255.252
 no ip redirects
 no snmp trap link-status
 pvc 1/40
  cbr 149760
  encapsulation aal5snap
  service-policy output WAN_150M
 !
!
router eigrp 1
 passive-interface default
 no passive-interface Tunnel0
 no passive-interface Tunnel1
 network 10.3.0.0 0.0.255.255
 no auto-summary
 !
 address-family ipv4 vrf user
 network 10.2.0.0 0.0.255.255
 network 192.168.2.0 0.0.0.255
 no auto-summary
 autonomous-system 1
 exit-address-family
!
ip route 0.0.0.0 0.0.0.0 10.99.0.2
no ip http server
no ip http secure-server
!
ip sla responder
logging alarm informational
access-list 100 deny   ip any host 239.192.240.1 log
access-list 180 permit ip host 172.19.64.1 any
!
```

```
control-plane
!
gatekeeper
 shutdown
!
line con 0
 exec-timeout 0 0
 transport preferred none
 stopbits 1
line aux 0
 stopbits 1
line vty 0 4
 exec-timeout 0 0
 password lab
 login
!
scheduler allocate 20000 1000
ntp clock-period 17179989
ntp server 10.11.0.1
!
end
```

MGCP Gateway and SRST Router Configuration

The following is the Cisco 3845 MGCP gateway and SRST configuration used for this testing:

```
version 12.4
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
!
hostname 3825-2
!
boot-start-marker
boot system flash:c3825-adventerprisek9-mz.124-9.T2
boot-end-marker
!
card type t1 0 0
logging buffered 1000000 debugging
enable password lab
!
no aaa new-model
!
resource policy
!
clock timezone EST -5
clock summer-time EDT recurring
network-clock-participate wic 0
ip cef
!
!
no ip dhcp use vrf connected
ip dhcp excluded-address 192.168.2.1 192.168.2.99
!
ip dhcp pool 0
   network 192.168.2.0 255.255.255.0
   option 150 ip 192.168.1.11
   default-router 192.168.2.1
   dns-server 192.168.1.11
   domain-name agency.com
!
!
```

```
no ip domain lookup
ip domain name agency.com
!
!
isdn switch-type primary-ni
voice-card 0
 no dspfarm
!
!
application
 global
  service alternate DEFAULT
 !
controller T1 0/0/0
 framing sf
 clock source internal
 linecode ami
 pri-group timeslots 1-24 service mgcp
!
controller T1 0/0/1
 framing sf
 clock source internal
 linecode ami
 cablelength short 133
 pri-group timeslots 1-21,24 service mgcp
!
!
!
!
!
!
interface GigabitEthernet0/0
 ip address 192.168.2.2 255.255.255.0
 duplex auto
 speed auto
 media-type rj45
 no keepalive
!
interface GigabitEthernet0/1
 ip address 10.11.0.10 255.255.0.0
 duplex auto
 speed auto
 media-type rj45
 no keepalive
!
interface Serial0/0/0:23
 no ip address
 encapsulation hdlc
 isdn switch-type primary-4ess
 isdn protocol-emulate network
 isdn incoming-voice voice
 isdn bind-l3 ccm-manager
 no cdp enable
!
interface Serial0/0/1:23
 no ip address
 encapsulation hdlc
 isdn switch-type primary-4ess
 isdn protocol-emulate network
 isdn incoming-voice voice
 isdn bind-l3 ccm-manager
 no cdp enable
!
router eigrp 1
```

```
 network 192.168.2.0 0.0.0.255
 no auto-summary
!
ip route 0.0.0.0 0.0.0.0 192.168.2.1
ip route 64.0.0.0 255.0.0.0 10.11.0.1
ip route 172.0.0.0 255.0.0.0 10.11.0.1
!
!
ip http server
no ip http secure-server
!
ip sla responder
!
!
!
!
!
!
control-plane
!
!
!
voice-port 0/0/0:23
!
voice-port 0/1/0
!
voice-port 0/1/1
!
voice-port 0/1/2
!
voice-port 0/1/3
!
voice-port 0/0/1:23
!
voice-port 1/0/0
!
voice-port 1/0/1
!
voice-port 1/0/2
!
voice-port 1/0/3
!
voice-port 1/0/4
!
voice-port 1/0/5
!
voice-port 1/0/6
!
voice-port 1/0/7
!
voice-port 1/0/8
!
voice-port 1/0/9
!
voice-port 1/0/10
!
voice-port 1/0/11
!
voice-port 1/0/12
!
voice-port 1/0/13
!
voice-port 1/0/14
!
```

```
voice-port 1/0/15
!
ccm-manager fallback-mgcp
ccm-manager redundant-host 192.168.1.12
ccm-manager mgcp
ccm-manager music-on-hold
ccm-manager config server 192.168.1.11
ccm-manager config
!
mgcp
mgcp call-agent 192.168.1.11 2427 service-type mgcp version 0.1
mgcp dtmf-relay voip codec all mode out-of-band
mgcp rtp unreachable timeout 1000 action notify
mgcp modem passthrough voip mode nse
mgcp package-capability rtp-package
no mgcp package-capability res-package
mgcp package-capability sst-package
no mgcp package-capability fxr-package
mgcp package-capability pre-package
no mgcp timer receive-rtcp
mgcp sdp simple
mgcp fax t38 gateway force
mgcp rtp payload-type g726r16 static
!
mgcp profile default
!
dial-peer voice 999011 pots
 service mgcpapp
 port 0/1/1
!
dial-peer voice 999012 pots
 service mgcpapp
 port 0/1/2
!
dial-peer voice 999013 pots
 service mgcpapp
 port 0/1/3
call-manager-fallback
 secondary-dialtone 9
 max-conferences 12 gain -6
 transfer-system full-consult
 ip source-address 192.168.2.2 port 2000
 max-ephones 336
 max-dn 500
 system message primary SRST Mode on 3825-2
 system message secondary SRST Mode on 3825-2
 moh flash:music-on-hold.au
!
!
line con 0
 exec-timeout 0 0
 stopbits 1
line aux 0
 stopbits 1
line vty 0 4
 exec-timeout 0 0
 password lab
 login
line vty 5 20
 exec-timeout 0 0
 password lab
 login
!
scheduler allocate 20000 1000
```

```
ntp clock-period 17179716
ntp server 11.0.0.1
!
end
```

References

This document references the following related documents:

- Cisco Unified Communications SRND Based on Cisco Unified CallManager 4.x— http://www.cisco.com/en/US/docs/voice_ip_comm/cucm/srnd/4x/uc4_2.html
- Cisco Unified CallManager System Guide (http://www.cisco.com/en/US/products/sw/voicesw/ps556/products_administration_guide_book09186a00806cc16b.html)
- Cisco Unified CallManager Administration Guide (http://www.cisco.com/application/pdf/en/us/guest/netsol/ns171/c649/ccmigration_09186a008074f2d8.pdf)
- Voice and Video Enabled IPsec VPN (V3PN) SRND http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/V3PN_SRND/V3PN_SRND.html
- Dynamic Multipoint VPN (DMVPN) Design Guide http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/DMVPDG.html
- Cisco IOS IP SLAs Configuration Guide (http://www.cisco.com/en/US/products/ps6441/products_configuration_guide_book09186a0080707055.html)
- Cisco IOS Security Configuration Guide, Release 12.4 - Call Admission Control for IKE (http://www.cisco.com/en/US/products/ps6350/products_configuration_guide_chapter09186a0080455adc.html)
- Cisco IOS IP SLAs Configuration Guide (http://www.cisco.com/en/US/products/ps6441/products_configuration_guide_book09186a0080707055.html)
- Integrating Cisco CallManager and Cisco SRST to Use Cisco SRST As a Multicast MOH Resource (http://www.cisco.com/en/US/products/sw/iosswrel/ps5207/products_feature_guide09186a00802d1c31.html)