Designing MPLS Networks for VoIP

What Every User Needs to Know

Michael F. Finneran
President and Principle Consultant

dBrn Associates, Inc.
Hewlett Neck, N.Y (USA)

Telephone: (516) 569-4557
Email: mfinneran@att.net

July, 2006

The idea of carrying voice traffic on an IP network has now become an established direction in enterprise networking. However, it is important to define which part of VoIP we're describing. The generic term "VoIP" really describes two separate and distinct ideas: local and wide area. In the local area, VoIP means replacing our traditional circuit switching or time division multiplex (TDM) based PBX systems with LAN switches that use IP/Ethernet handsets in conjunction with telephony servers to provide an IP-PBX. Driven primarily by savings on cabling costs, the migration to IP-based solutions, in either a pure LAN switch or a hybrid IP-TDM configuration, has now become a foregone conclusion. The same cannot be said of wide area VoIP implementations.

Local VoIP is essentially an equipment decision, while wide area VoIP is a service decision. There appear to be a few reasons why the wide area VoIP market has been progressing more slowly. First, it is far more difficult to develop a cost justification in the wide area. Basically, we're looking at reducing long distance costs, primarily on calls placed between company facilities. With virtual private network services (i.e. the voice meaning of "virtual private network") most large customers pay only 1-1/2 to 2 cents per minute for on-net to on-net long distance. Further, when used in conjunction with an IP PBX, the delay introduced by the wide area connection makes it difficult to keep the end-to-end delay below the required 150 msec.

We are starting to see customers embarking on wide area VoIP projects. The impetus for these projects is typically the idea that if we have IP-PBXs in our major sites, the next logical step is to use IP-based services to interconnect them. Further, data networks have been migrating from frame relay to MPLS-based VPN services, and a big part of the MPLS pitch is the ability to support integrated voice/data services using MPLS's QoS capabilities.

While enterprise customers are getting the message that MPLS is the service the carriers are pushing for IP voice, there is still considerable confusion about MPLS services actually work. Further, there is even less understanding regarding the basic process for designing and implementing an MPLS voice network. The purpose of this paper is to provide a general overview about what an MPLS-based service is, how the service operates, and most importantly, the design process for implementing a voice service on MPLS.

So What is MPLS?

Mutli-Protocol Label Switching (MPLS) is probably the single most important development in TCP/IP. In a nutshell, MPLS provides a mechanism whereby IP networks can define virtual circuit services to improve security and provide a multi-level quality of service (QoS) capability. Those QoS capabilities would provide performance guarantees regarding delay, jitter, and packet loss.

Currently, MPLS is being used by Internet Service Providers (ISPs) as the basis of a new service geared toward enterprise customers. Typically referred to as an MPLS-Virtual Private Network (MPLS-VPN), the overall concept for these services is described in RFC 2547bis. We are also seeing MPLS capability being provided in enterprise routing equipment, but the primary focus today is in the area of carrier services.

These MPLS-VPNs are being marketed as an alternative to the traditional frame relay service the carriers have sold to provide wide area connections between LANs. Market researcher Vertical Systems notes that there are currently 8000 US customers using MPLS-based VPNs with about 90,000 sites connected, and that customer base is expected to double by 2010[1].

For those who are not familiar with packet switching technology, there are two basic types of packet forwarding: connectionless and connection-oriented with virtual circuits. The original IP switching technology used in the public Internet is connectionless, best effort. "Best effort" means that some of the packets may be lost due to buffer overflows or transmission errors. "Connectionless" means the service operates without virtual circuits. In traditional IP, each packet finds it own way router-to-router to its destination.

[1] Rick Malone, "Free at Last: The Move to Dedicated IP-VPN Networks", Business Communications Review, May, 2006, page 43.

The result of that connectionless operation is that packets can arrive out of order, and there is no practical means to guarantee the performance attributes for individual data flows. How could you possibly guarantee worst-case performance for delay if you don't even know the path the packets will take to their destination?

Network services like frame relay and MPLS employ virtual circuits. You can think of a virtual circuit as a software-defined pathway through the packet network; all the packets traveling between those two sites will follow that path. There are a number of number of advantages that virtual circuits can bring to packet switching:

1. As all of the packets are following the same path, they should all arrive in the same sequence as they were sent.
2. Virtual circuits provide security for traffic within the shared packet network, as hackers cannot get access to traffic on another user's virtual circuits.
3. Most importantly, as the carrier knows the path the traffic will take, they can manage the amount of capacity that has been assigned to each user, and hence guarantee the worst-case performance that traffic should experience with regard to delay, jitter, and packet loss.

The impact of those first two advantages pale in comparison to the third, particularly when we consider converging voice, video, and data services on the same network. Data traffic is typically not as time sensitive as voice or video traffic. Further, as data traffic uses TCP, when packets are lost, TCP will recover them automatically.

Time sensitive voice and video traffic are forwarded in UDP, which operates on the “send and pray” transmission philosophy; in UDP, there is no mechanism to detect or recover from lost or errored packets. TCP recovers lost packets by ordering retransmissions, but that type of process takes so long that it would be useless in a voice environment; the retransmitted packets would arrive too late to have any relevance to the conversation. In short, you’ve got only one shot at delivering a voice packet, so a network service that guarantees performance for delay, jitter, and loss has a far greater value in supporting voice and video services.

Building an MPLS-based Converged Network

The ability to support multiple traffic categories and provide separate performance parameters for each has made MPLS a natural fit for wide area VoIP services. However, there is still much confusion surrounding how the MPLS performance guarantees actually work. As enterprise users migrate to IP-based local solutions, it is inevitable that the idea of connecting them together through a wide area IP network will bubble to the surface. At that point, the user will have to do some serious research on how these MPLS-based network services handle different classes of traffic.

The basic concepts for a guaranteed packet service developed in frame relay, but we are finding that most people didn’t understand them very well in that context either. With any guaranteed packet service, there are two critical transmission rates: the access rate and the service's guaranteed capacity.

When a customer’s router is sending traffic into the network, the traffic is always sent at the access rate (i.e. if your router is connected to the network with a T-1/DS-1 rate facility, all frames are sent at 1.536 Mbps). However, the network only guarantees to deliver some portion of that traffic.

In frame relay, the guaranteed capacity is called the Committed Information Rate (CIR), and it is specified for each virtual circuit in the network. As you might expect, a virtual circuit with a higher CIR costs more. As long as the average transmission rate stays below the CIR, the carrier guarantees to deliver a very high percentage of that traffic, typically 99.99%. If the transmission rate on that virtual circuit exceeds the CIR, that additional traffic is marked discard eligible" and essentially becomes best effort. If there is capacity available in the network, that excess traffic will be delivered, but if there is a congestion condition, excess traffic can be discarded. In short, excess traffic has a higher probability of being dropped.

This same concept has been adopted in MPLS, however it has had to be modified in two important ways given the nature of the MPLS technology:

1. Where frame relay offers essentially two traffic categories, guaranteed and discard eligible, an MPLS service can offer three or more; for

convenience those categories are typically called Gold, Silver, Bronze, and Best Effort. The carrier provides different guarantees regarding delay, jitter, and packet loss for traffic sent in each category.

2. In an MPLS service, the customer does not pay for virtual circuits, they simply pay for access capacity. Every MPLS end point can communicate with every other end point (i.e. full mesh connectivity). So rather than having a CIR for each virtual circuit, in MPLS, a certain percentage of the access capacity is allocated to each traffic category; that allocation is typically referred to as the Class of Service Profile. As you might expect, CoS profiles with a higher percentage of Gold and Silver traffic are priced at a higher rate.

In an MPLS service, the customer marks each packet and so assigns it to one of the available service categories; that assignment is done by setting the Diff Serv Control Point in the IP header. Typically the highest priority is assigned to voice, while the others may be used for video and various classes of data traffic. The performance parameters are computed over the period of a month, so these should not be construed to be a hard and fast guarantee for each individual packet.

The categories and performance guarantees for AT&T and Verizon Business are summarized in the table below.

MPLS Service Performance Guarantees (US Domestic Traffic)

AT&T Enhanced VPN Service[1]				**Verizon Business Private IP Service**			
Service Class	**Performance Parameters**			**Service Class**	**Performance Parameters**		
	Jitter	**Delay** (Round Trip)	**Packet Delivery**		**Jitter**	**Delay[2]** (Round Trip)	**Packet Delivery**
CoS 1	<9 msec	<104 msec	99.9%	**Real Time/ Voice**	<5 msec	<100 msec	99.995%
CoS 2	Not Applicable	<108 msec	99.9%	**Assured Forwarding[3]**	Not Applicable	<100 msec	99.99%
CoS 3	Not Applicable	<120 msec	99.8%				
CoS 4	Not Applicable	Not Applicable	Not Applicable	**Best Effort**	<5 msec	<100 msec	99.995%

[1]- AT&T's SLA targets are defined end-to-end, and are applicable to USA Eastern region to USA Western region. They assume T1 access connections at each end point with tail circuits within 250km.

[2]- Verizon Business computes round trip delay from provider edge to provider edge, so it is not directly comparable to AT&T's delay performance

[3]- Verizon Business actually defines three sub-categories within the Assured Forwarding class, but they all provide the same delay and packet delivery parameters.

How is Excess Traffic Treated?

Most customers have not had to deal with the design of voice networks using MPLS. According to a recent study of MPLS user conducted by Forrester Research, only 20% of MPLS customers are actually using the service for voice.[2] The other 80% may be in for a big surprise when they make the move to voice.

When we take a closer look at how MPLS services actually work, we find that there are really two categories, real time and everything else. While we generically refer to the service categories as Gold, Silver, and Bronze, AT&T calls their highest category COS 1 while Verizon Business designates it Expedited Forwarding (EF). In both cases, the real time category specifies worst-case performance for delay, jitter, and packet loss. The other categories specify only delay and loss performance.

The important distinction regarding the real time category is how excess traffic is treated- it’s dropped. It’s not downgraded to a lower category or marked “discard eligible”, it’s dropped at the entry point or edge router. Traffic in the lower categories is marked "out of contract" (i.e. the equivalent of frame relay's "discard eligible"), but it is still forwarded through the network if there is capacity available.

Understanding how that excess voice traffic is treated is critical for designing a VoIP network. As we noted earlier, IP voice uses UDP transport, so there is no recovery for lost or errored packets. The impact those lost packets will have depends on the voice encoding system that is used. If we encode the voice using G.711 or 64 Kbps pulse code modulation, we can typically tolerate about 10% random packet loss before the user will note a serious degradation in voice quality. If we use the more efficient 8 Kbps G.729A voice compression, the system will only tolerate 1% to 2% packet loss.

If we configure more voice channels than the Gold service category can support, the network will begin dropping packets. As that packet dropping will be a random function, if you try to configure 15 voice channels over a service with a Gold capacity that can only support 10, you won't have 10 good trunks and 5 bad ones; you'll have 15 bad ones!

Designing a Voice Service Over MPLS

The message is that customers who are looking to carry voice traffic on their MPLS-based VPN services had better be careful about how they design and coordinate the various elements in their networks. This design process will involve coordination between the voice and data staffs.

[2] Lisa Pierce, "The Multifacteted MPLS Customer", Business Communications Review, June, 2006, page 50.

The overall design process goes like this:

1. The first step is to decide which voice calls will be carried over the MPLS network. The obvious answer is voice traffic that goes between sites that are connected to the MPLS network. As those sites will be other company locations, we're talking about voice tie lines that run between the PBXs or IP PBXs in those sites. We can potentially carry other voice traffic over those tie lines and then extend those calls through the public network to off-net locations; in tie line networks, we refer to that option as "tail-end-hop-off".
2. Next, we have to do a voice traffic study. We isolate the voice traffic that will be carried over the MPLS network, identify the busy hour, determine the amount of traffic that must be carried during that period, and compute the number of trunks that will be required to support it with an acceptable level of blocking (i.e. the P-Grade of service). That is done with the Erlang B traffic engineering formula. In the old days we sized trunk groups by poring over traffic engineering tables, but there are now Web sites like www.erlang.com/calculator/erlb/ that can compute the number of trunks if we provide the busy hour traffic and the required P-Grade of service.
3. Once we know the number of trunks that are required, it's time to shift into VoIP mode. We first determine the bit rate required for each voice trunk including all of the packet overhead. The variables in that computation are the voice encoding used (e.g. G.711, G.729A, etc.) and the size of the voice sample carried in each packet. The table below will help with that. It is important to recognize the tradeoffs involved. G.711's 64 Kbps encoding requires more capacity per channel, but it can tolerate about 10% packet loss. G.729A is more efficient, but when packet loss reaches 2%, the voice quality will degrade substantially; the voice compression also adds about 15 msec to the delay. Larger packet sizes are more efficient, but they also increase network delay. The cRTP mode cannot be used on wide area MPLS services.
4. Once we know the number of bits required per trunk and the number of trunks, we multiply them together to determine the capacity required for real time traffic. Now you can begin dealing with the subtleties.

Things get tricky when there are different number of trunks running between sites as there can be blocking and dropped traffic at the egress port. Further, if you have busy hour traffic that is substantially higher than the average traffic volume, you will have to determine if it's really cost-effective to size your trunk group for the busy hour or for the average volume. If you design for average volume, the excess traffic that occurs in the busy hour can overflow to the public network, and you pay for it on the old cost-per-minute plan. Once you start thinking about the cost of the additional MPLS-real time capacity needed to support those extra tie lines versus the cost public network services, you should start looking at the overall cost of the MPLS solution versus sticking with the public network pricing plan you currently have. You might discover that carrying voice traffic on the MPLS network doesn't really save you enough to justify the effort involved.

If you do decide that the savings justify the effort, then you have to insure that the implementation is coordinated. The PBX or IP-PBX must be configured with the correct number of tie lines on each route and the quality of service marking (Diff Serv Control Points, IEEE 802.1p LAN priority, etc.) must be coordinated end-to-end.

Packet Voice Transmission Requirements (Bits Per Second per Voice Channel)						
Codec	Voice Bit Rate	Sample Time	Voice Payload	Packets Per Second	Transmission Requirement (PPP or Frame Relay)	
					RTP	cRTP
G.711	64 Kbps	20 msec	160 bytes	50	**82.4 Kbps**	**68.0 Kbps**
G.711	64 Kbps	30 msec	240 bytes	33.3	**76.2 Kbps**	**66.6 Kbps**
G.711	64 Kbps	40 msec	320 bytes	25	**73.2 Kbps**	**66.0 Kbps**
G.729A	8 Kbps	20 msec	20 bytes	50	**26.4 Kbps**	**12.0 Kbps**
G.729A	8 Kbps	30 msec	30 bytes	33.3	**20.2 Kbps**	**10.7 Kbps**
G.729A	8 Kbps	40 msec	40 bytes	25	**17.2 Kbps**	**10.0 Kbps**

Note: RTP assumes 40-octets of RTP/UDP/IP overhead per packet
Compressed RTP (cRTP) assumes 4-octets RTP/UDP/IP overhead per packet
PPP/Frame Relay overhead adds 6-octets per packet

Conclusion-
People Cost Money Too

The interesting thing about wide area VoIP is that the strongest proponents for carrying voice traffic on MPLS services have never actually worked on tie line networks. Voice tie line networks were a big thing back in the late-1970s and early 1980s when switched voice service prices were much higher than they are today. Large business users would rent voice grade private lines to interconnect PBXs in major sites, and invest in PBX software like the Electronic Tandem Network (ETN) option on their old AT&T Dimension PBX to build a tie line network. Later we found we could reduce the cost of those dedicated circuits by using high capacity private lines and T-1 multiplexers, but the ongoing traffic engineering task remained.

What we found out was that one of the major costs involved with running a tandem network was the job of traffic engineering, and that job never ended. Voice traffic patterns change over time, and we would have to conduct traffic studies on each route periodically to insure that we still had the optimal number of trunks. If not, we would install or remove circuits and make the appropriate configuration changes in the PBX systems at each end. The bigger the network and the more dynamic the traffic patterns, the more effort that was required.

The legacy of those tie line networks was that when the carriers began offering virtual private network services (the voice meaning of “VPN”) with attractive pricing for voice calls running between company sites, customers jumped at the option. Inter-site voice traffic migrated back to the public network, and we rejoiced that we were no longer saddled with the drudgery of running the tie line network. In a voice VPN, you merely determine the number of access trunks required from each site to the carrier's network, and after that, it's the carrier’s problem.

As George Santayana wrote in <u>The Life of Reason</u>: "Those who cannot remember the past are condemned to repeat it." So the idea of voice on MPLS has led us on a circular path back to a network plan from 20-years ago; the underlying network technology is different, but a tie line is still a tie line. If voice is going to migrate to MPLS VPN services, someone is going to have to relearn the skills many of us happily forgot two decades ago.

The clear message is that a successful VoIP over MPLS installation will require both voice and data expertise, and real cooperation between the two groups.

Contact Information
dBrn Associates, Inc.
189 Curtis Road
Woodmere, NY 11598 (USA)
Tel. 516-569-4557

dBrn Associates, Inc. is an independent network consulting firm specializing in telecommunications networks and technologies.

All opinions expressed herein are based on independent research, and may not be quoted without written permission of the author.

Michael Finneran can be reached at mfinneran@att.net.