NUNTIUS

Quality of Service and Multi-Protocol Label Switching

White Paper

Carter Horney

for

Nuntius Systems, Inc.

13700 Alton Pkwy., Suite 154-266
Irvine, CA 92612
949.295.0475 voice
www.nuntius.com

INTRODUCTION

With rising costs in voice communications and the advent of the Internet, many now look to utilize Voice over Packet technology to cut telecom expenses. However, before this service becomes widely adopted, voice quality must improve. A limiting factor in quality of service (QoS) in Internet Telephony, or voice over IP (VoIP), is its inability to deliver consistent toll-quality service. With today's routers, IP packet routing is based on the connectionless paradigm in which each packet may follow an independent path through the network. The initial IP design did not include resource reservation, traffic engineering and implementation of differentiated services. To provide QoS greater than those obtained with "best-effort" mechanisms, Multiprotocol Label Switching (MPLS) was designed to include a label, such as protocol layer 2, that supports a connection-oriented switching paradigm. MPLS is designed to scale, providing traffic engineering and re-routing mechanisms within the IP domain.

Why VoIP?

Why bother with VoIP? Why not stay with circuit switch technology that has worked so well over the public switched telephone network (PSTN)? VoIP offers three primary benefits:

- It allows a single network infrastructure to carry both voice and data.

- It's more efficient than circuit switch, requiring a smaller investment in network infrastructure to carry a given amount of traffic.

- It supports private voice networks

The first of these benefits is clear. ISP's, with their extensive network investments want to enter the voice market and traditional carriers with separate voice and data networks can benefit from the huge economies of scale if they had one infrastructure for both types of traffic.

The second benefit is a consequence of the fundamental difference between circuit switched and packet-switched networks. When a call is placed in a circuit-switch network like the PSTN, a dedicated connection is "nailed up" between the calling and called parties, and it remains nailed up until call termination. Packet-switched networks work more like a postal system. The conversation is broken up into small packets that are relayed across the network between the parties. At any time, a given link can forward packets for many conversations at the same time. Thus, not only does VoIP replace two networks with one it also reduces the size and cost of the network.

The third benefit applies to commercial enterprises with offices in wide spread geographic locations. This benefit, called toll bypass allows companies to use their private Intranets to carry voice as well as data, there by saving dramatically on their phone bills, especially for long-distance calls. Today, IP virtual private networks (VPNs), which still are in the early stages of large-scale deployments, are the most approximate manifestation of the new utopian service vision.

Service providers are on a quest for the holy Grail ––– value-differentiated services that can facilitate high-margin, tiered pricing schemes and efficient use of network resources.

What are the problems with VoIP?

There are two major problems with using an IP network to carry voice traffic. The first problem is deciding how to negotiate the parameters and services for a call. The PSTN uses a protocol known as Signaling System 7 (SS7), which can be made to run over an IP network. But brings with it the assumptions of the centralized PSTN. If SS7 is used, you cannot take full advantage of the inherent flexibility of the distributed IP network.

The second problem is how to get the voice quality we expect from the telephone system over IP networks, without long and unpredictable delays. IP packets are forwarded much like letters are in the postal system. Like letters, the packets do not all follow the same route; and they all do not take the same length of time to get where they are going. They do not even necessarily arrive at their destination in the same order as sent. As a result, IP networks are subject to long (hundreds of milliseconds) and random delays. While these results don't affect data communications, they wreak havoc with delay-sensitive traffic such as voice and video traffic. Dealing with QoS is one of the most challenging aspects of VoIP.

Within the telecommunications industry, a distinction is made between the control information (signaling) flow and the media (data) flow, and in many protocols they use different paths within the network. The benefits of separating the control and media flows is that it allows a small number of intelligent, expensive signaling devices to manage a large number of simpler, cheaper media devices. For voice networks, QoS is primarily an issue for the media flow, which carries voice traffic.

TCP or UDP Protocol?

When the industry refers to TCP/IP, they generally mean the family of protocols which include TCP (Transport Control Protocol), UDP (User Datagram Protocol), IP (Internet Protocol) and some underlying WAN or LAN transport layer. Data sent using the TCP protocol is referred to as a segment; UDP data is referred to as a packet. IP is the network layer that lies under TCP and UDP. IP provides unreliable, connectionless packet delivery to a specified host address. IP packets, called datagrams, can be lost, duplicated, or delivered out of order. The advantage of IP is that an IP datagram provides a universal means for delivering data, independent of the underlying network technology.

UDP and TCP sit on top of the IP at the transport layer. Both UDP and TCP use port numbers to de-multiplex data sent to a host. A port number is specific to the application. Each UDP packet and TCP segment has a source and a destination port number. A host that waits for incoming connections is referred to as the server, and the host that initiates the connection is referred to as the client. Servers "listen" on well-known port numbers for common applications such as FTP (File Transfer Protocol), SMTP and HTTP. Clients generally choose a random source port number and connect to a server at a well-known port.

UDP provides ‘best effort’ delivery of data with an optional checksum to preserve data integrity. UDP packets have the same reliability issues as IP; packets are not guaranteed delivered in the same order at the remote host. TCP provides a reliable stream of data by using sequencing and acknowledgement numbers to recover lost data, detect out of order segments, and resolve transmission errors. Video conferencing and Voice-over IP (VoIP) are ideal examples for using UDP because throughput is most important to the real-time reception of audio and video.

TCP is connection-oriented, while UDP is connectionless. This implies that to send a UDP packet, a client addresses a packet to a remote host and sends it without any preliminary contact to determine if the host is ready for data reception. UDP has no throttling mechanism so packets can be sent a full speed - as fast as the underlying physical device can send them. On slow processors, UDP’s lack of overhead can make a large difference in the throughput when compared to TCP. The lack of end-to-end connection makes it suitable for broadcasting many-to-many type messages.

TCP provides a connection-oriented, reliable stream of data. Before sending data to a remote host, a TCP connection must be established. TCP offers a reliable data stream at the cost of more processing overhead and reduced throughput.

The Voice Packet

Either a voice or a fax data packet of ingress signals resides in a media packet, conveyed as a packet to a network layer, and reconstructs egress signals from the received packet. The packet formats are specifically designed for efficient transport over IP, frame relay, Asynchronous Transfer Mode (ATM), or proprietary networks.

Although voice is the primary packet, facsimile relay is supported along with a signaling relay. Packet voice is used to transmit high quality voice and in-band signaling over a packet network. In addition to the DSP intensive signal processing voice compression algorithms, the packet management functions of G.168 line echo cancellation (both local hybrid and a 128ms long tail) are required for packet voice applications.

The voice packet generally supports multiple compression formats. Among those commonly supported are G.711 (64kbps u-law/A-law); G.726 (16/24/32/40kbps ADPCM conversion); G.729A (8 kbps CS-ACELP) and G.723.1 (5.3 or 6.3 kbps MP-MLQ/ACELP) compression/decompression voice coding algorithms.

The fax relay typically supports modulation and demodulation of fax data for V.21 (300bps channel 2 data pump used only during negotiation), V.17 (14400bps data pump), V.29 (9600 bps data pump) and V.27ter (4800 bps data pump). This relay is covered under T.30 standards. Frame loss recovery in a data transmission session is handled using a higher-level error correction mode (ECM) defined in the annex of the T.30 standard. The purpose is to relay fax images from Group 3 equipment over the packet network. This is an essential requirement for seamless operation with standard circuit switched networks.

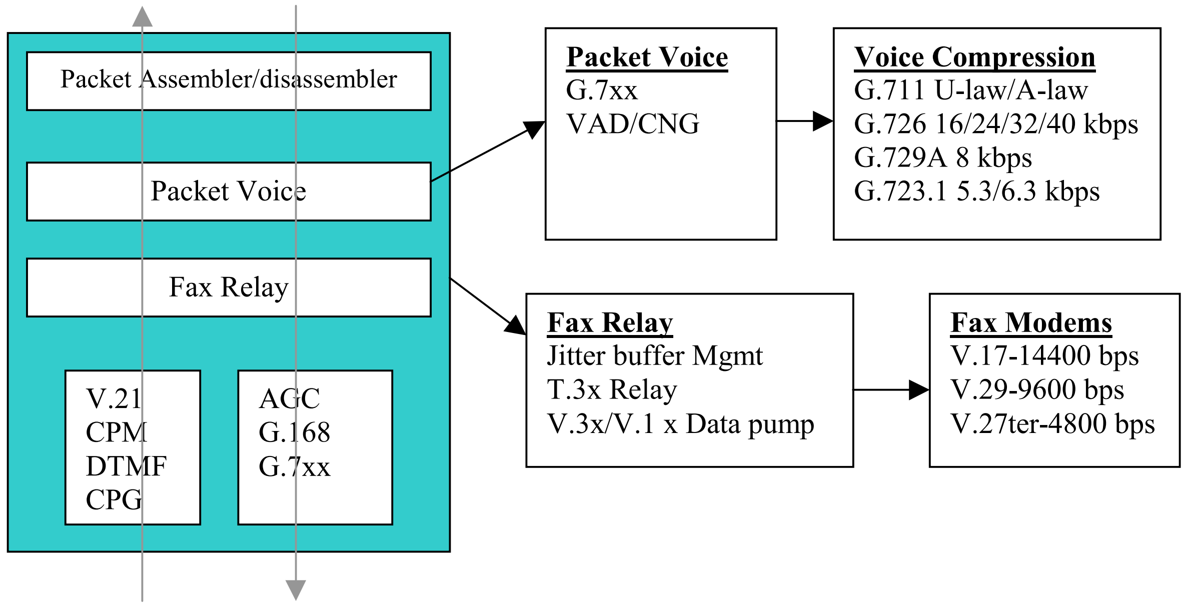

Packet voice also supports the relay of in-band (300-3600 Hz) signaling used to convey call control and call management between the local exchange and Customer Premise Equipment (CPE). The detection and generation of in-band signals is required: DTMF and call-progress tones.

Packet voice has broad applicability in Integrated Access Devices, Gateways, Cable Modems, DSL modems, LAN PBXs and Computer Telephony systems to provide convergence between traditional circuit switched networks and the emerging packet based infrastructure.

Signaling Voice over IP Networks with H.323

Two competing protocols provide signaling to set up voice calls and select standard voice codecs (5.3, 6.3, 8, 16, 48, 56 and 64 kbps) over the IP network. H.323, based on RTP, was ratified by the ITU, originally for data conferencing over corporate LANs. After H.323 was finalized, a subset of it was defined for pure Internet use where the gatekeeper was no longer required. H.323's original purpose was to facilitate interoperability between multimedia teleconference (audio, video and data) devices over LAN and WAN without a guarantee of QoS. Its intricate, multifaceted umbrella standard defines the multiple codecs, call control, and channel set up specifications to support audio, video, and data connectionless networks. H.323 software enables communications equipment manufacturers to implement Internet telephony applications in a broad range of products from small gateways to high-end carrier class switches and optional gatekeepers. The component devices within an H.323 zone are end-point Terminals, Gateways (convert digitized voice into addressed IP packets), Gatekeepers (registers users and keeps track of access policies and billing information) and the Multiple Control Units (MCU) which are used to conference multiple end points.

The protocols specified by H.323 are listed below. H.323 is independent of the packet network and the transport protocols over which it operates and does not specify them.

- Audio codecs
- Video codecs
- Q.931 (used over TCP for initial call set up)
- H.225 registration, admission and status (RAS)
- H.225.0 control signaling (for logical connection establishment)
- H.245 control signaling (capabilities negotiation and exchange)
- Real-time transfer protocol (RTP) providing end-to-end delivery services for real time data.
- Real-time control protocol (RTCP) controls monitoring of the data delivery for multicasting.

The original H.323 audio call set up consisted of: Q931 call signal, H.245 capability exchange, H.225.0 establish logical channel(s) and sending of the RTP. This connection method proved to be excessively slow. H.323 has evolved so that many of the original defects have been reduced. Two new methods of set up now exist.

- H.245 tunneling over the H.225.0 connection
- H.323v2 FastStart. (basically the SIP call start up)

RTP usage scenarios support the audio conferencing application of audio data used by each conference participant. An RTP header precedes each chunk of audio. The RTP header and data is in turn contained in a UDP packet. The RTP header specifies the type of audio encoding (such as PCM, ADPCM, or LPC) contained in each packet. Audio and (optional video) media are transmitted as separate RTP sessions and RCTP packets are transmitted for each medium using two different UDP port pairs and/or multicast addresses.

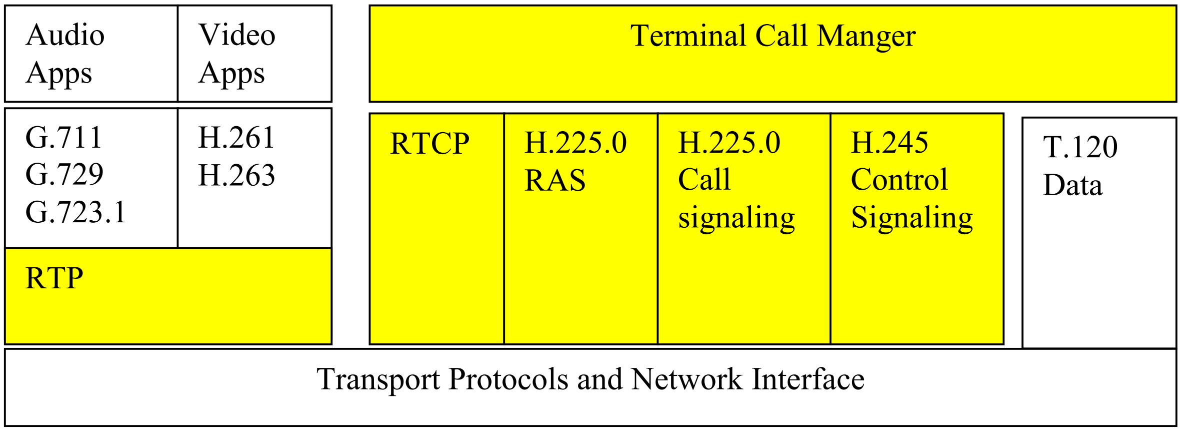

Figure H.323 Terminal Side Protocol Stack

SIP Signaling Protocol

Criticisms have been leveled at H.323 mostly for its complexity and slow call set-up rate. Although an updated version of H.323 has addressed some of these criticisms, another protocol is under development by the Internet Engineering Task Force (ITEF) - the Session Initiation Protocol (SIP). SIP, modeled after HTTP, is not as mature as H.323; its first version was completed in February 1999. The simpler to use SIP is used to locate the called party and H.323 to connect the parties. SIP primary attributes are reduced set-up time and ability to connect to a SS7 network.

The Best Current Practices for Telephony (BCP-T) extensions, allows transport of SS7 ISUP or Integrated Service Digital Network User Part, signaling information over SIP. The SIP is more scalable than H.323 and more suitable for large carrier networks. SIP is likely to become the dominant application protocol, because of the IETF backing and the protocol's relative simplicity. SIP handles the following: user location (which end system will be used), user capabilities (media available), user availability (called party availability), call set-up (call features) and call handling (transfer and termination of calls).

SIP is a text-based protocol, similar in syntax to HTTP and Real-Time Streaming Protocol (RTSP). Messages can be conveyed over UDP or TCP. SIP provides its own reliability over UDP.

MGCP and MEGACO/H.248

The architectural model that has been adopted (by both the ITEF and the ITU) is that of separate media gateways and media gateway controllers. A media gateway (MG) is a somewhat simple device that passes media flows from one network to another and is told what flows to be set up by a media gateway controller (MGC). The MGC is aware of the signaling flows and controls the MG. A protocol is required to facilitate communications between the MGC and MG. The primary standards-based option available today is the Media Gateway Control Protocol (MGCP), or its latest incarnation Megaco (also designated by ITU as H.248), a joint IETF-ITU development project.

MGCP has been deployed in some devices, but does not provide all the functionality required to provide interactions between MGC's and MG's in all types of networks. For that very reason, most existing implementations of MGCP have proprietary extensions, which are not interoperable – although organizations, such as the International SoftSwitch Consortium (ISC), are working to provide some level of consistency.

Megaco is considered to be a more complete protocol aimed at providing this function in a standard-based way. The basic standard has now been approved and the IETF and ITU are working to develop the extensions. Implementations of Megaco are already underway and the IETF and Cable Labs Packet Cable™ group, backed by the North American cable industry, has endorsed the MGCP/Megaco protocol. Megaco addresses the relationship between the MG, which converts circuit-switched voice packet-based traffic and the MGC, sometimes called a call agent or soft-switch. Megaco, considered by many to be a relatively low-level device-control protocol, instructs a MG to connect streams coming from an outside a packet or cell stream such as the RTP. Megaco is quite similar to MGCP from an architectural standpoint, but supports a broader range of networks (including ATM). Megaco is designed for intra-domain remote control of connection-aware or session-aware devices such as VoIP gateways, remote access concentrators (RACs), Digital Subscriber Line Access Multiplexors (DSLAMs), Multiprotocol Label Switching Routers, optical cross connects and PPP session aggregate boxes.

Quality of Service Deficiencies in IP Networks

In order to have guaranteed QoS in a network, all of the data packets sent in each direction, during each session, must follow the same path (in network jargon, be connection oriented) and some means for reserving resources along that path must exist. IP is not connection oriented, and IP routers don't generally have sophisticated mechanisms for committing resources at each hop; that's why ensuring a specified QoS is so difficult over an IP network. Two mechanisms have attempting to solve this problem unsuccessfully.

The Differentiated Services (DiffServ) protocol was defined to enable different levels of services to be provided across IP networks, Their protocol uses a space in the IP header to indicate different traffic types and priorities. Routers in the network are able to look at this information and prioritize traffic accordingly while DiffServ provides no guarantees. For example, congestion and queuing can increase latency, reduce available bandwidth, and thereby reduce voice quality. By itself, DiffServ is not adequate for VoIP.

The Resource Reservation Protocol (RSVP) is a signaling protocol used in IP networks to reserve resources for certain specified data flows. Although RSVP can reserve the resources, it cannot guarantee that traffic will flow along the path on which the resource was reserved: as nodes and links are added or removed in an IP network, the path along which data flows can change. RSVP attempts to recover and create an updated path reflecting the new technology, but there can be no guarantee that the QoS will be maintained, and it is possible that RSVP will fail to create an updated path.

Working towards QoS

Transport services are concerned with the transfer of user and control-plane data across networks. The services are provided by wireline or wireless access networks. A major problem at the transport layer is how to ensure reliable, predictable QoS for time-sensitive applications across an IP infrastructure originally designed for best-effort data transport. Today the problem is generally solved, by running IP-over-ATM, and using the built-in ATM QoS mechanisms. *The limitations of this approach are obvious for carriers that do not intend to implement ATM.*

The alternative has been to over-provision the IP network so that in the absence of congestion, traffic can be forwarded through the network with minimum latency, jitter and packet loss. Throwing bandwidth at the problem is certainly not a sustainable solution for a large-scale network. To deliver IP QoS, embryonic stage intelligence service layers schemes such as MPLS are in development. With MPLS, service providers can define specific packet delivery paths for traffic through the IP networks, rather than allow intermediate routers to make the packet-forwarding decisions. Conventional packet routing sends traffic along the shortest available path through the network. By moving traffic flows onto less congested paths, MPLS can better balance a networks traffic load and overall network response time and throughput.

Multi-Protocol Label Switching (MPLS) solves the QoS issue by setting up explicit paths through the network. MPLS is a technique that facilitates high-performance transport of IP traffic across Wide Area Networks. In particular, it marries connectionless IP technology to connection-oriented technologies like ATM. MPLS assigns labels to IP flows placing them in IP frames. The frames can then be transported across packet or cell-based networks and switched on the labels rather then being routed using IP address lookup. Using MPLS techniques it is possible to set up explicit routes for data flows that are constrained by path, resource availability and requested Quality of Service.

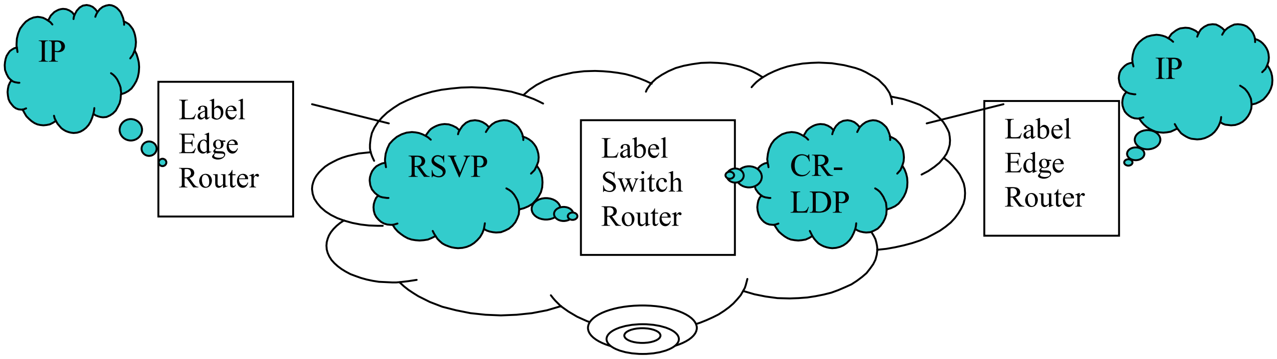

The path is defined by the sequence of IP addresses of the nodes to be traversed. All of the data that constitutes a flow is given the same label (fixed format data field inserted at the front of each packet) on entry into the MLPS network. At each node the packet is routed based on the label value and incoming interface and sent on its way with a new label value on the outgoing interface. The paths are referred to as label-switched paths (LSP). Since an LSP is a well-defined path through an IP network, it provides a means for ensuring a specified quality of service where QoS is provided by the underlying infrastructure. The multi-protocol nature of MPLS means it can be used to support IP networks over any Layer 2 infrastructure – Asynchronous Transfer mode (ATM), packet-over-SONET, Gigabit Ethernet, frame relay, etc.

A signaling protocol is used to set up the LSP. At present, two options compete for the task. One is Labels-RSVP, an extension of RSVP. A major factor is that RSVP is a proven, deployed technology. A number of companies are supplying RSVP-based MPLS networks based on RSVP. RSVP is an existing protocol (RFC2205) that has been extended by ITEF to support label distribution. RFC 2207 provides the extensions for IPSEC and RFC 2206 covers MIB management.

The Constraint-based Routing Label Distribution protocol (CR-LDP) has been co-developed from the ground up by ITEF specifically for MPLS networks. Companies are just beginning to provide MPLS devices based on the new technology. Labels-RSVP and CR-LDP both have strong advocates. However, the market is initially focused on RSVP, although it is likely that CR-LDP implementations will become available.

How an MPLS Network Works?

In the MPLS network shown in the figure, host X is sending packets to two destinations, Y (yellow packets) and Z (aqua packets). Both flows ingress into the network through label-switching router A (LSR A), which determines what label-switched path (LSP) to use for each packet and adds a label to the packet accordingly.

LSR A then forwards the packet via the appropriate interface for the chosen LSP. LSR B, an intermediate router in the network, takes each packet it receives and, based on the incoming interface and label value, decides on the outgoing interface and label value with which to forward the packet. In this example, each packet with label value 21 is forwarded to LSR D and now bears new label value 47, whereas packets with label value 17 are relabeled with value 11 and sent to LSR C.

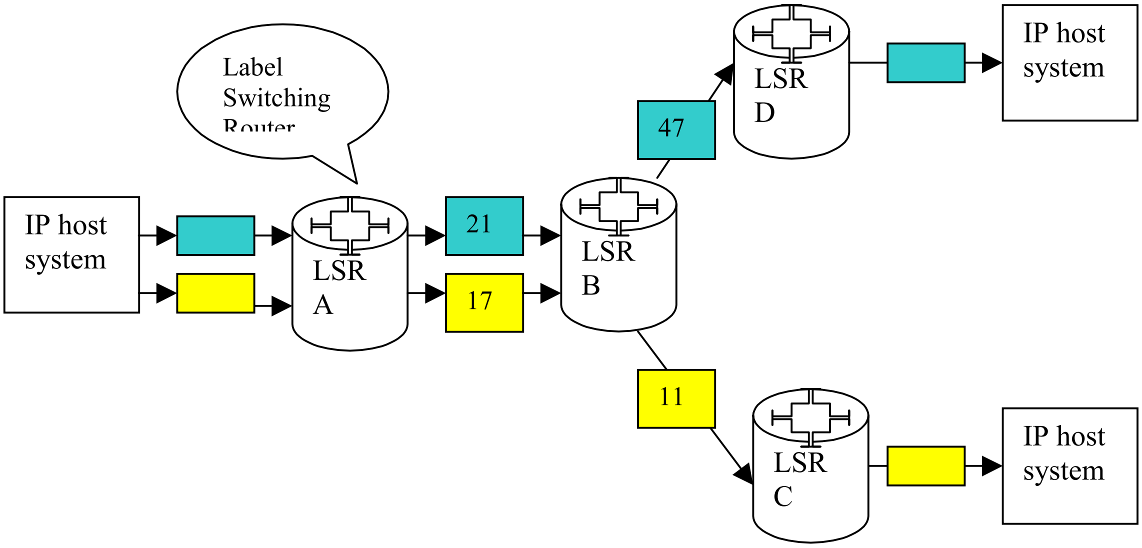

VoIP over MPLS

There is more than one way to use MPLS to implement VoIP. An individual path can be set up for each voice call, signaling the LSP at the same time as the cell is signaled. More often, system operators prefer to create a smaller number of larger-bandwidth pipes in advance, down which multiple calls can be funneled much like a T1/E1 channel works.

When sending voice over IP (or over MPLS, which is in turn running over IP), the Real-Time Protocol (RTP) is used, running over UDP. This protocol, along with the Real-Time

Control Protocol (RTCP), provides timing information in voice-packets to ensure that smooth voice reproduction can be achieved on the receiving network end. The RTP, UDP, and IP headers included in data transfer can be a significant overhead compared with the size of the voice. A voice sample typically may contain only 20 bytes of data. Whereas the UDP has an 8-byte header, an RTP header 12 bytes, an IP header 24 bytes, and an MPLS header adds 8 bytes – for a total of 52 bytes of overhead on a single voice sample.

The IP header is needed for routing a sample through the IP network when running directly over IP, but when using MPLS LSP, no IP routing is required; hence it is possible to remove the IP header and save 24 bytes. This is a central issue that has motivated the formation of the Voice-over-MPLS (VoMPLS) Discussion Group. Many issues exist, the key one is that stripping and replacing the IP header at either end of the LSP adds a performance overhead. Also, if an LSP is used for multiple voice channels, a multiplexing scheme must be chosen.

Voice-over IP Evolve towards SS7 on IP

Many equipment manufacturers and service providers are focusing on bringing the intelligence provided by SS7 signaling into the IP domain. This task is being achieved through the deployment of soft-switches for interworking within the public network.

Few telephone customers have just basic service. Most subscribe to some package of enhanced features —– call waiting, call forwarding, three-way calling for example —– that are enabled by the SS7 network in the public network. To bring these services into the IP domain and to enable delivery of new IP devices such as Internet call waiting or wireless Web browsing, sophisticated IP signaling is needed., along with tight integration of the SS& and IP network. Service creators can recreate these services in the IP network —– a daunting task that some carriers are nevertheless planning to do. Alternatively, carriers can use SS7 signaling within the IP domain (enabled by carrying SS7 over IP) to perform data base dips and provision their intelligent network services. In either case, the voice-over-IP network and the public network have to talk to each other.

Looking at the candidates for carrying SS7 over-IP, simple computer telephony protocol (SCTP) is the protocol of choice, and SCTP has several adaptation layers. Although carrying SS7 over IP is easier than translating the service, the SCTP still must interact with the gateway protocol (MGCP), H.248 or SIP). In the evolving infrastructure of the VoIP, media gateways are implemented at the network edges to translate the audio stream to and from the PSTN and to communicate signaling back to the soft-switches in the IP network. Soft-switches perform functions of traditional public network Class 4 or 5 switches in the core network.

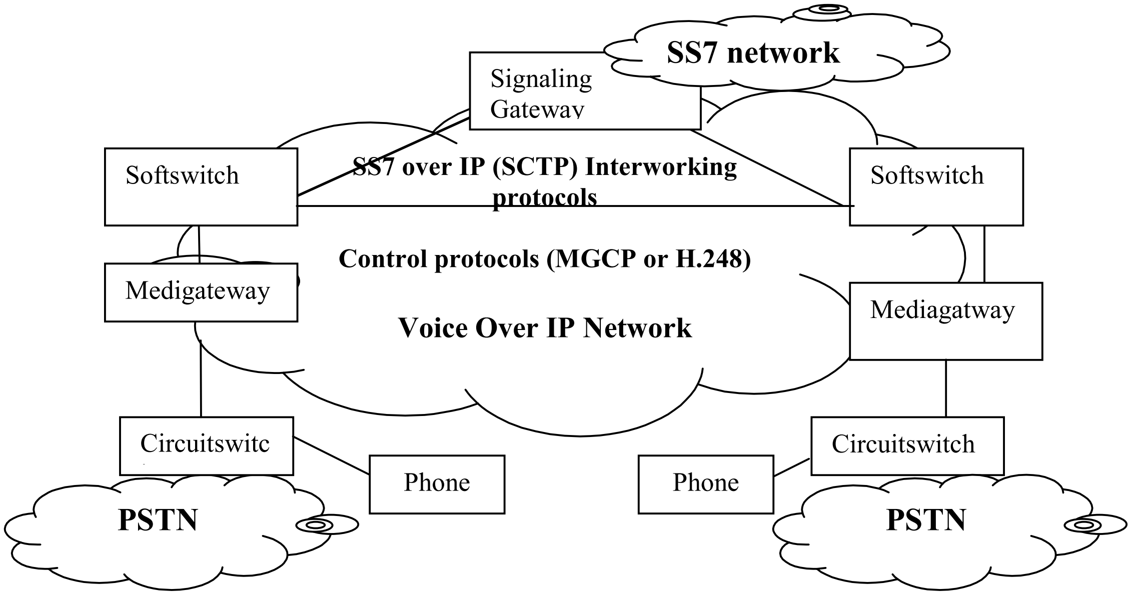

Nuntius Systems, Inc.

In-depth knowledge of, and extensive experience with IP telephony firmware and software components are required to build complete end-to-end solutions. While complex interworking with the public telephone network is handled at the edge of the network, Nuntius engineers are skilled in developing software applications in Signaling Gateways, Media Gateway Controllers, Media Gateways, Gatekeepers, Soft-switches and IP Phone Appliances. The following chart depicts the various service layers of the hybrid telephony network. It's within these layers that Nuntius engineering teams are experienced in developing software for VoIP solutions.

	Public Network	IP Network	Wireless Network
Application services	Voice, fax, modem	Digitized voice and video, fax, data, Web	Voice, fax, data
Network services	**Protocols:** SS7, ISDN, analog signaling, CAS CDRS, etc. **Devices:** Class 5 switch, SS7 STP, SSP, SCP, back office devices	**Protocols:** H.323, SIP, MGCP, H.248, MPLS **Devices:** Gateway, gatekeeper, soft-switches, back-office devices	**Protocols**: GPRS, UMTS/WCDMA **Devices:** SGSN, GGSN
Transport services	Wireline Access Analog, ISDN, DSL, Cable, leased line (TDM)		RF Access: TDMA, CDMA, UMTS/W-CDMA

When implementing a VoIP platform, one of the first issues to consider is whether to buy or build software. In the vast majority of cases, buying is the best option because the protocol stacks require very significant development efforts to build and maintain these components. When looking to acquire protocol components, it’s prudent to make sure your development partner has experience with the range of technologies you require, along with a roadmap and resources to keep your product solution ahead of the competition.

Another important consideration is whether your software vendor simply supplies standard protocol stacks or can provide a more complete solution. For example, inter-working between protocols is a significant element of the value proposition you offer your customer. A vendor who can give you the full solution for interoperation between such protocols and also provide the answer to how it will integrate with other software layers is most preferable. Outside of the standard protocol support, you need also to pay particular attention to the scalability and performance of the software.

Nuntius has a deep bench of qualified and experienced DSP professionals prepared to work with you build and support your product efforts, end-to-end.