NENA VoIP Technical Committee

VoIP Characteristics

Technical Information Document

NENA 08-503 - VoIP Characteristics Technical Information Document (TID)
Issue 0, June 10, 2004

Prepared by:
National Emergency Number Association (NENA) VoIP Technical Committee - VoIP Characteristics Working Group

Published by NENA
Printed in USA

NENA
TECHNICAL INFORMATION DOCUMENT

NOTICE

This Technical Information Document (TID) is published by the National Emergency Number Association (NENA) as an information source for the designers and manufacturers of systems that are used for the purpose of processing emergency calls. It is not intended to provide complete design specifications or parameters or to assure the quality of performance for systems that process emergency calls.

NENA reserves the right to revise this TID for any reason including, but not limited to, conformity with criteria or standards promulgated by various agencies, utilization of advances in the state of the technical arts or to reflect changes in the design of network interface or services described herein.

It is possible that certain advances in technology will precede these revisions. Therefore, this TID should not be the only source of information used. NENA members are advised to contact their Telecommunications Carrier representative to ensure compatibility with the 9-1-1 network.

Patents may cover the specifications, techniques or network interface/system characteristics disclosed herein. No license expressed or implied is hereby granted. This document is not to be construed as a suggestion to any manufacturer to modify or change any of its products, nor does this document represent any commitment by NENA or any affiliate thereof to purchase any product whether or not it provides the described characteristics.

This document has been prepared solely for the voluntary use of E9-1-1 Service System Providers, network interface and system vendors, participating telephone companies, etc.

By using this document, the user agrees that NENA will have no liability for any consequential, incidental, special, or punitive damages arising from use of the document.

NENA's Technical Committee has developed this document. Recommendations for change to this document may be submitted to:

National Emergency Number Association
4350 N. Fairfax Dr, Suite 750
Arlington, VA 22203
800-332-3911

Acknowledgments:

This document has been developed by the National Emergency Number Association (NENA) VoIP Technical Committee.

The following industry experts and their companies are recognized for their contributions in development of this document.

Members:	**Company**
Henning Schulzrinne	Columbia University
Jonathan Flack	Cox Communications
Deborah Stone	Intrado
Lawrence Gowin	Level 3 Communications
Brian Rosen	Marconi
Mike Aprile – WG Leader	Red Sky Technologies, Inc.
Nate Wilcox – VTC Chair	Vermont Enhanced 9-1-1
Greg Welenson	Vonage Holdings, Inc.
Larry Short	

TABLE OF CONTENTS

1 Executive Overview

1.1 Purpose and Scope of Document

The purpose of this document is to procure, create and publish a VoIP primer document to be used by individuals not familiar with VoIP technology.

1.2 Reason for Issue

This document is issued to serve as a reference for the VoIP Technical Committee as they define and create short-term and long-term standards and technical documents for VoIP, as it relates to E-911.

1.3 Reason for Reissue

NENA reserves the right to modify this document. Whenever it is reissued, the reason(s) will be provided in this paragraph.

1.4 Recommendation for Standards Development work

Not Applicable

1.5 Costs Factors

Not Applicable

1.6 Acronyms/Abbreviations

This is not a glossary! See NENA 01-002 - NENA Master Glossary of 9-1-1 Terminology located on the NENA web site for a complete listing of terms used in NENA documents.

The following Acronyms are used in this document:	
API	Application Programming Interface
ATA	Analog Terminal Adapter
CGI	Common Gateway Interface
CODEC	enCOde/DECcode
DSL	Digital Subscriber Line
DSP	Digital Signal Processing
HDTV	High-Definition Television
HTTP	Hyper Text Transfer Protocol
IAD	Integrated Access Device
IETF	Internet Engineering Task Force
IP	Internet Protocol
IPSEC	IP Security Protocol
ITU	International Telecommunication Union
LCR	Least Cost Routing
MGCP	Media Gateway Control Protocol

The following Acronyms are used in this document:	
MTA	Multimedia Terminal Adapter
PDA	Personal Digital Assistant
PPP	Point-to-Point Protocol
PSQM	Perceptual Speech Quality Measurement
QoS	Quality of Service
RAS	Remote Access Server
RTP	Real Time Protocol
S/MIME	Secure Multipurpose Internet Mail Extensions
SIP	Session Initiation Protocol
SOHO	Small Office/Home Office
TCP	Transmission Control Protocol
TLS	Transport Layer Security
UAC	User Agent Client
UAS	User Agent Service
VoIP	Voice over Internet Protocol
VPN	Virtual Private Network
XML	eXtensible Markup Language

2 Technical Description

2.1 TCP/IP Primer

TCP/IP is defined as an industry standard suite of protocols that devices use to find, access, and communicate with each other over a transmission medium. A protocol is a set of standards and rules that need to be followed. In the case of networking, a protocol is the set of standards and rules that a device's hardware and software must follow in order to be recognized and understood by other devices. The protocol suite is implemented via software most commonly known as the TCP/IP stack.

2.1.1 What is TCP/IP made of?

Under the hood, TCP/IP's architecture consists of several "layers" performing certain functions. Each layer contains protocols. There are four general layers of the TCP/IP stack:

Application layer
Transport layer
Internet layer
Physical or Network Interface layer

A full-scale description of each layer and its underlying functionality is beyond the scope of this document. However, here's a brief overview of the role that each layer plays and how the layers work together.

2.1.2 Application layer

The data that is sent starts off at the top of the TCP/IP stack in the Application layer. This layer contains network applications and services that a user interfaces with in order to use network communication. Also living in the Application layer are utilities for services like file and print services and name resolution. A good example of this is NetBIOS, an application programming interface (API) that supports a desktop operating environment. All of the utilities that work with TCP/IP live in the Application layer. Examples of some familiar protocols and their primary usage include:

dns	Domain Name Service
ftp	File transfer
http	Transporting web traffic
pop3	Retrieving email
rtp	Real-time transmission of multimedia data
smtp	Transporting email
telnet	Remote terminal access

2.1.3 Transport layer

Once the Application layer is through with the data, it is passed down the stack to the Transport layer. The two major components of the Transport layer are the Transfer Control Protocol (TCP) and the User Datagram Protocol (UDP). Entire books are available on TCP, UDP, and the Transport layer, but simply put, the Transport layer is an interface that applications use for network connectivity. The designers of TCP/IP wanted to make sure that the data that is sent gets received by the correct device, as well as the correct application running on that device and in the correct sequence. The Transport layer provides this functionality. In the Transport layer, there are mechanisms for error checking, flow control, and verification ensuring the integrity and completeness of the data that is being sent and received.

Although TCP and UDP provide similar functionality, there is one very important difference between the two. TCP is considered a connection-oriented protocol, while UDP is considered a connection*less* protocol. A connection-oriented protocol is one that establishes a connection with another machine and maintains that connection for the entire duration of data transmission. A variety of functions are built into TCP that check and recheck the data while the two devices are connected. This makes TCP a reliable, albeit slower, transmission protocol. UDP, however, does not establish a connection with the target machine at all. UDP is told by the Application layer which device it is supposed to transmit to, with no questions asked. This obviously makes UDP a much faster protocol when it comes to data transmission. But UDP has rudimentary error checking and flow control, as well as reliability issues. It is important to note that ordered reliable delivery service of TCP cannot be exploited by real-time communications because there is no time for retransmissions. For this reason, the Real-Time Transmission Protocol (RTP), which is commonly used for VoIP applications, rides over UDP.

2.1.4 Internet layer

Beneath the Transport layer is the Internet layer. Three key protocols reside in the Internet layer: Internet Protocol (IP), Address Resolution Protocol (ARP), and Internet Control Message Protocol

(ICMP). Each of these serves a specific purpose. There are also two less-used protocols, Reverse Address Resolution Protocol (RARP) and Internet Group Management Protocol (IGMP). IP addressing and address resolution occur within the Internet layer. IP addressing is a scheme that standardizes how devices are identified and differentiated from one another. This scheme allows any device running TCP/IP to communicate with other devices running TCP/IP anywhere in the world. No matter what type of machine, operating system, or network topology the devices live on, as long as both devices are using TCP/IP, they're speaking the same language. ARP's job is to resolve a logical IP address, such as *www.mywebsite.com*, into its physical equivalent address. ICMP is mostly used by routers to send information back to a source computer about a transmission that device is trying to make.

2.1.5 Physical layer

The final layer on the TCP/IP stack is the Physical layer. This layer is at the base of the stack and is the last section a packet must go through before it's sent out across the transmission medium. The Physical layer contains a collection of services and specifications that provide and manage access to the network hardware. Its responsibilities include:

- Interfacing with the device's network hardware.
- Checking for errors in incoming packets of data.
- Tagging outgoing packets with error-checking information.
- Acknowledging the receipt of a packet.
- Resending that packet if no acknowledgement is returned by the recipient.

This layer is almost totally invisible to the everyday user, which, given its complexity, is not such a bad idea.

2.2 Decentralized Implementation

When the PC revolution took hold two decades ago, we saw a movement away from the monolithic computing architecture that mainframes provided to a decentralized computing environment. Because of the Internet, companies can leverage a more distributed and mobile workforce while staying constantly in touch via cell phones, pagers, laptops, and connected organizers. Virtual Private Networks (VPNs) have made it far less expensive (compared to traditional company-owned wide area networks) and more secure to connect remote offices to headquarters. Due to the lack of networking capabilities in existing telephony systems, decentralization efforts have largely focused on data-only applications. With the emergence of IP Telephony technology, business can now decentralize their voice and data applications over a common network along with remote access.

IP Telephony allows organizations to distribute their IP call processing equipment to the locations that are most appropriate, based on the specific needs of the organization. Unlike traditional, circuit-switched technology, which is primarily a monolithic architecture, with the PBX/switch being the central and controlling device, IP Telephony services can be distributed across the entire enterprise. IP gateways, signaling servers, application servers, etc. can be located in different buildings, at remote locations or, even on the Internet or an external private network.

2.3 Physical Media Independence

In a packet network, data is routed through the network such that a set of properly configured end points can communicate. This data delivery can occur over various physical connections such as twisted pair or fiber optic cabling or even broadcast through the air. The delivery of data in packet networks is not constrained by the underlying delivery mechanism. Data between two end points can also take differing routes across the network. Routing decisions are made on each individual packet, regardless of how the previous packet was routed. This independence from a particular physical media makes IP networks particularly adaptable to different environments. It also provides for a very high level of availability by opening several different options to provide connectivity.

2.4 Network Requirements

Given that data packets can take different paths across the network to their destination, media streams like audio or video can be adversely affected. The following sections define latency, jitter, and packet loss, which are factors that affect the quality and reliability of IP Telephony.

2.4.1 Latency

Latency refers to the time, in milliseconds, that it takes for data packets to travel from the origination point to the destination point. If the traffic load on a network is well within its capabilities, the latency will likely be low. When the volume of traffic reaches levels that strain the processing power of different components, some packets may have to be queued before processing. This delay will result in increased latency. A VoIP call traversing a network with high latency figures may have a noticeable delay in the audio.

2.4.2 Jitter

Jitter is the variation in arrival times of different packets within a data stream. Since each data packet can take a different route across a network, their transmission time (or latency) can be different. If certain points of the network become congested, packets from a media stream can have widely varying latency, or high jitter. Jitter is measured in milliseconds, representing the variation in packet arrival times. Many equipment manufacturers implement what is called a jitter buffer. This buffer acts to minimize the variation in arrival times between packets, countering the effects of jitter.

2.4.3 Packet Loss

As the name implies, packet loss refers to the number of data packets which never reach their final destination. These packets may be misrouted, time out, or be excluded from a media stream, if packets are received out of order. Packet loss can be measured as a percentage or by peg count.

Packet loss manifests itself in VoIP calls as brief moments when the audio cuts out. Note that individual packet losses are unintelligible by the human ear. Only when packet loss is significant can a user recognize this problem.

2.4.4 General Guidelines

There are differing standards to what are acceptable figures for jitter, latency, and packet loss. The table below gives some general guidelines for what figures are acceptable.

General Guidelines		
	Ideal	**Maximum**
One way latency	≤ 100 milliseconds	≤ 150 milliseconds
Jitter Delay	≤ 40 milliseconds	≤ 75 milliseconds
Packet loss	≤ 1%	≤ 3%

2.4.5 Quality of Service (QoS)

Quality of service refers to a particular network's ability to transfer data when measured against the constraints below:

- Latency
- Jitter
- Packet Loss
- Availability

Latency

As discussed previously, latency refers to the amount of one-way delay in a network. If a network has sufficient capacity to transport all of the data, latency will typically be very low. Each network element can process the packets as they arrive and move them along. However, as network utilization goes up and different network components have to process more and more traffic, some packets will be queued for processing. While these packets are waiting to be processed, latency in the network rises. Packets that originally took 100ms to traverse the network may take 500ms or more, depending on the levels of congestion.

For applications that aren't time critical such as browsing the web, these delays are only a nuisance. However, time sensitive traffic like voice or video is greatly impacted by high latency figures. The transmission delay results in a pause in the audio stream, which may sound like the other person isn't responding.

Jitter

If latency were a constant value for all packets transmitted, a congested network would manifest itself only with delayed audio streams. Since IP packets can take varying paths to get to their destination, the transmission time for each packet can vary widely. Without measures to combat jitter, an audio stream would be garbled and possibly unintelligible. Jitter buffers can delay the processing of packets to even out the varying arrival times. However, the large these buffers are the more latency will be experienced by the users.

Packet Loss

Packet loss occurs mainly when congested network elements drop data packets. If the application pushing the data is using TCP, a request can be sent notifying of the loss and the packets can be resent. This is ideal for browsing the internet or file downloads. With those applications, the loss of packets can be significant. The downside of TCP is that it uses more bandwidth and isn't as agile as UDP. UDP would not request a retransmission of a lost packet. For time sensitive applications like voice or video, the occasional lost packet is less of a problem than having to keep track of every packet in a media stream. If the packet loss is significant enough, gaps in audio can be heard on VoIP calls.

Availability

Arguably the most important measure of QoS is the availability of the network. In other words, there are 31,536,000 seconds in a year. How many of those seconds was a network unable to transport data? Like a TDM network, an IP network should be highly available. Some TDM networks aim for 99.999% availability, which means that for about only 5 minutes a year calls will not be completed due to network problems. How often this goal is achieved is another question, but well designed IP networks can meet these same metrics.

2.5 Media Types

Packet loss, jitter, and latency are all considered detrimental to the performance of a network. However, these network behaviors have different impacts on different media types. Media types such as voice or audio require the timely and regular delivery of IP packets. Data on the other hand is much more forgiving. Non-time sensitive and time sensitive applications are discussed below. The effects of jitter and latency can be minimized through the use of jitter buffers and QoS mechanisms.

2.5.1 Data

Browsing content on a network does not necessarily require a connection with low jitter or packet loss. Data applications can reorder packets and handle delays with the user having little to no perception of the network impairments. Data transmissions, like one would see while using an internet browser, are not very sensitive to delay or jitter. Thus, a network that is fine for surfing the internet may not be acceptable for voice or video.

2.5.2 Voice

Voice can traverse IP networks and have a PSQM score equivalent or better to what would be seen over circuit switches. However, voice quality can be impacted with the introduction of delay, jitter, and packet loss. Long latency or delay is a nuisance on a two way voice call. When one person speaks, there is a delay while the audio is transmitted to the distant end of the connection. This delay is seen again when the other person responds. High jitter on a connection can result in the audio sounding garbled if the audio packets are processed out of order or discarded all together. Minimal packet loss is not a problem as the human ear is not precise enough perceive the periods of missing audio. However, excessive packet loss can seriously impact voice quality.

2.5.3 Video

High fidelity video can also be transmitted over IP with no degradation in quality. In fact, many of the HDTV television signals are transported over IP. However like voice, video quality will be impacted with the introduction of delay, jitter, and packet loss. Latency, if regular, isn't as much of an issue since video tends to be one way. In the case of two way video conferencing, the same awkward delays would be present that one would experience in a high latency voice connection. Jitter on a video connection can quickly result in perceivable degradation of the video image. In fact, both jitter and packet loss will negatively impact the perceived quality of a video connection.

2.6 Separation of Call Control and Media Transport

IP telephony completes the evolution from in-band signaling found in analog telephony to complete separation of signaling and media flows.

2.7 Virtual Locations and Mobility

For the most part, traditional circuit-switched telephony is based on linking the end-points (i.e. phones) to the PBX/switch using a dedicated, hard-wired connection. Through the use of patch-panels, distribution centers, wiring closets, etc., cabling is required which provides call control and media transport services to and from the phones. While some capabilities do exist that allow wired analog and digital phones to "dynamically" move from place to place without reconfiguration at the PBX/switch level, for the most part, current residential and business telephony infrastructures require that a particular phone remain connected to a particular cable unless such configuration changes are performed.

With VoIP, this is no longer a restriction. For business and residential IP Telephony service, it is not only possible to easily move IP end-points from place to place; it is one of the marketed benefits. With IP Telephony, an end point only needs to have IP connectivity with the IP Communication Server or SIP Proxy Server and the IP end-point(s) that they will be communicating with. For example, a business user with an IP phone that is connected to an IP Communication Server can simply unplug their phone from the network, walk down the hallway, or even, to another building, plug the phone back into the network, and they will be able to immediately use that phone. This example assumes that the new location where the IP phone is plugged into can provide access back to the Communication Server. For residential users, this mobility can extend to virtually anywhere that the Internet is available.

While this mobility offers many advantages, it also creates a real problem when trying to respond to a person making an emergency call from one of these phones. Because IP phones and end-points can be so mobile, the ability to utilize the current E-911 technology and infrastructure to accurately locate a caller and route the emergency call to the appropriate PSAP can be challenging or even impossible.

2.8 IP Telephony Architecture Options

One of the challenges with IP Telephony is that it must, in many cases, integrate into existing traditional telephony infrastructures that exist today. Based upon their specific goals, each organization can choose from three different approaches to begin integrating IP Telephony into their current telephony infrastructure (see Figure 1). In many cases, an organization will utilize more that one of the approaches to meet all of the IP telephony requirements within their enterprise.

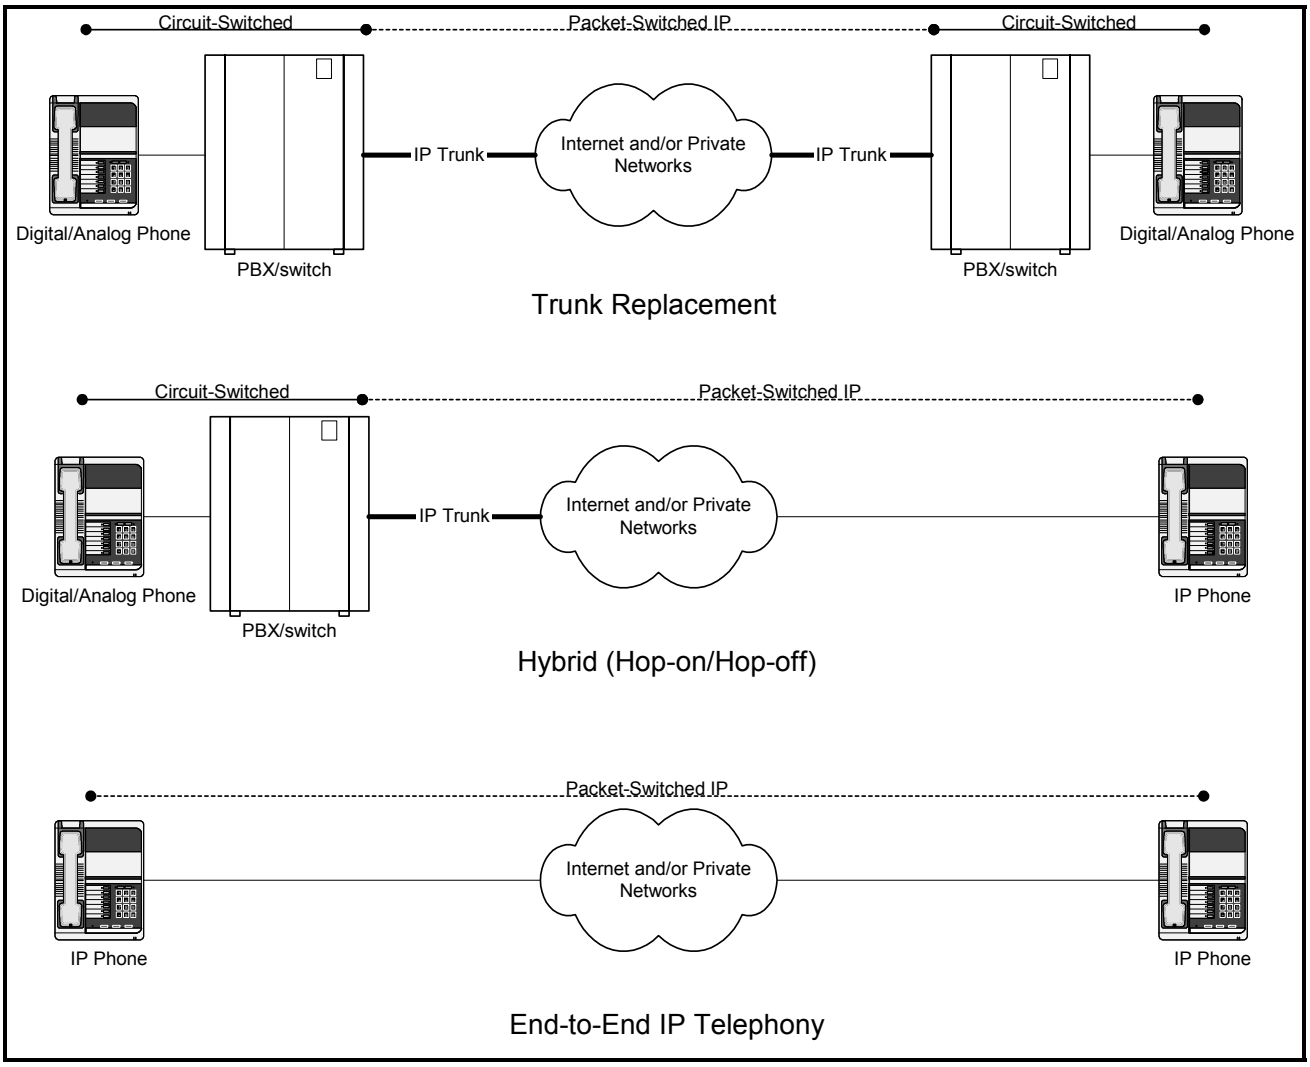

Figure 1 –IP Telephony Architecture Options

2.8.1 Trunk Replacement

In this approach, both the person originating (caller) the call and the person receiving (callee) the call continue to use circuit-switched telephone services from their end-point (i.e. phone) to the call processing equipment (i.e. PBX, switch). The call travels from the caller's endpoint to the PBX/switch, which then connects, via the public Internet, a private IP-based network, or some combination, to a PBX/switch that is close to the callee. This approach requires no changes in the phones and dialing behavior and only minimal changes at the PBX/switch. In most cases, the approach is used without the caller or callee being aware of it. Many PBX/switch vendors now offer IP trunk interfaces that simply replace a T-1 trunk with a packet-switched connection.

2.8.2 Hybrid (hop-on/hop-off)

Another approach, sometimes called hop-on or hop-off depending on the direction, directs calls from a traditional analog or digital phone to an IP-based phone or vice versa. In both cases, the IP phone

is addressed by a regular telephone number, although the phone may not necessarily be located in the geographic area typically associated with that area code. A number of companies have started to offer IP phones for residential and small-business subscribers that follow this pattern. A closely related architecture is called an *IP PBX* or *IP Communication Server*, where phones within the enterprise connect to a gateway that provides connectivity to the PSTN.

If the IP Communication Server is shared among several organizations and operated by a service provider, it is referred to as *IP Centrex*, as the economic model is somewhat similar to the Centrex service offered by traditional local exchange carriers.

2.8.3 Direct, end-to-end IP connection

The final approach dispenses with the gateways and uses direct IP-based communications end-to-end between the caller and the callee. With this approach, and IP end-point can directly communicate with another IP-endpoint, as long as they can establish an IP link over the Internet and/or one or more private networks.

2.9 Methods of Usage

Businesses have a number of options that they can choose when implementing IP Telephony. These options will depend on their specific business goals, current telephony infrastructure and level of expertise. The three primary options for business are; replace their PBXs, extend their PBXs, or utilize a 3rd party to provide IP Telephony services (IP Centrex). For residential users, the option now exists to use a broadband connection to the Internet for their telephony services.

2.9.1 Replacing the PBX

The IP Communication Server is usually implemented as software running on one or more computer servers running Windows NT/2000, Linux or another comparable operating system. One method of usage for IP Telephony is to replace the current PBX/switch equipment with IP Communication Server equipment and software. This method provides the most complete set of IP telephony features and benefits since if allows for organizations to implement all three of the architectures described in Section 2.8.

The drawbacks to this method are that it usually requires a forklift upgrade of the existing telephony equipment and all of the applications in place. Furthermore, while IP Telephony vendors have focused their attention on the features, reliability and performance of IP Communication Servers over the past few years, in some cases, they still do not meet all of the standards and expectations set by the current circuit-switches PBX/switch equipment that is available. For these reasons, IP PBXs are commonly marketed towards “green field” installations where no legacy systems are already in place.

2.9.2 Extending the PBX

For companies who have a large investment in a traditional circuit-switched equipment, adding IP telephony capabilities to their existing infrastructure can provide many of the features and functionalities that IP telephony can deliver. By implementing IP gateways alongside PBXs, organizations can employ a cost-effective and easy way to extend voice services to remote locations

or users. Because these features are integrated with their existing equipment, organizations can continue to manage them from a central location. IP gateways provide a bridge that connects a circuit switched network with a packet switched network by multiplexing and forwarding voice packets to and from remote locations over the Internet and/or private IP network. Some of the benefits of this approach are;

- Support for converged networks and applications
- Potential cost savings through toll bypass
- Flexibility to use multiple transport media (e.g., DSL, cable, frame relay, etc.) to connect to remote sites
- Preservation of existing infrastructure and investment
- Elimination of the need to purchase a separate PBX for remote locations
- An effective and economical solution for SOHO environments.

2.9.3 IP Centrex

A variation of the PBX replacement approach is to replace, or augment, existing circuit-based telephony with IP Centrex service. Like traditional Centrex, IP Centrex service decreases the initial capital investment for the enterprise and makes system maintenance the responsibility of the service provider. Unlike traditional Centrex, where each phone has its own access circuit, IP Centrex only requires that the organization have a single Internet or private IP network connection to the provider and is generally more cost-efficient. In an IP Centrex environment, the IP end-points reside within the organization while the IP Communication Server usually resides at the service provider.

2.9.4 Residential Service

Over the past few years, a growing number of residential users have begun using the Internet as one of their methods for voice communications. Initially, this was accomplished by software such as MS Net Meeting and Net2Phone. This software allows the residential user to utilize their home computer, connected to the Internet, to receive and place calls to/from other users with similar software. The users use the speakers attached to their computer or a headset to listen and a microphone connected to their computer to talk.

Soon after, a number of companies begin offering residential users a service that would allow them to place/receive calls with others who have a traditional phone connected to the PSTN. This was provided through the use of IP gateways, which can route the call from the Internet to the PSTN and vice-versa. With this approach, users still use the speakers attached to their computer or a headset to listen and a microphone connected to their computer to talk.

Recently, a number of companies have begun offering a service that works very much like the telephone that people have in their home today. A residential customer can pick up the phone, dial the number and it connects to whom they are calling. In more technical terms, this service uses the phone adapter, known as an Analog Telephone Adapter (ATA) or Multimedia Terminal Adapter (MTA) to convert the analog signal to a digital signal. The digital signal then can be sent over a high speed (i.e. cable, DSL) Internet connection. Residential customers can get a new phone number or can keep the existing number (number portability). On-net call (calls between locations connected to

the Internet can be free of charge; off-net calls are delivered to the PSTN network using IP gateway devices, with Least Cost Routing (LCR) functionality.

When someone calls a residential customer with this service, they dial a standard telephone number. Behind the scenes, this number looks very much like an e-mail address. This number instructs the call to travel over the Internet and through the network to the ATA. The phone connected to the ATA rings. From that point, the rest of the call is identical to a phone call made using an analog phone directly connected to the PSTN, via a Local Service Provider.

2.10 Hardware Components

2.10.1 IP End-Points (i.e. phones, PCs, etc.)

VoIP call typically originate on one of the three types of devices: softphones, analog phones connected to the Internet via Multimedia Terminal Adapters (MTA) and full-fledged IP Phones. Each of these devices must support some form of a VoIP signaling protocol and a protocol to transmit packetized voice. Typical choices are SIP and H.323 for signaling and RTP for voice path.

Softphones are the cheapest, but also the most limited option. A softphone is a program, which emulates IP telephone on a personal computer, without the use of Digital Signal Processors (DSPs) for voice transformation. It is an inexpensive way of permitting Internet calling, which initially could only work with other IP end-points. Service providers are beginning to support softphones by allowing users to communicate with devices connected to the PSTN through the service providers' gateways.

The benefits of softphones are their low cost, flexible user interfaces, and their attractiveness to computer users. Voice quality depends on the quality of audio system in the host computer, but is usually inferior to the other two types of end-points under the same network conditions. Recent advances make possible softphone installations on PDAs and laptops, and in combination with service provider support turn softphones into truly portable communication devices. User experience in using softphones is radically different from using analog phones in the PSTN environment.

A Multimedia Terminal Adapter (MTA) is a device that connects to the Internet and in which one or several analog telephones can be plugged in. A MTA performs all necessary voice conversions and runs one of the VoIP signaling protocols. Once the installation is complete, the user experience mimics the PSTN experience very closely, including full use of existing telephones.

Typically, a service provider will supply a customer with an MTA device and establish the customer's account, which would include assignment of a regular telephone number from the provider's pool.

Voice quality depends on the quality of the codec implementation. MTAs use DSPs and typically will support more codecs than softphones. A MTA costs more than a softphone, but in a residential service, will typically be subsidized by service providers and, at this time, has emerged as the preferred device option.

IP Phones are high-end VoIP devices, with prevalent use in the enterprise environment at this time. These phones have rich sets of features typical of office phones connected to PBXs and often replace legacy office phones, as traditional PBXs are replaced with IP PBX systems. These phones often have multiple control buttons, shortcut keys, support multiple lines and have LCD screens large enough to display short messages. Recent devices can provide excellent voice quality under favorable network conditions. These devices however are still too expensive for the residential market.

2.10.2 Access Gateways

Access Gateways enable VoIP service providers to connect to Class 5 switches as an intermediate solution before Class 5 VoIP softswitches are available. The interface protocol to Class 5 switches in North America is typically GR-303 and in the many other countries, it is V5.2. Access Gateways provide necessary support for VoIP deployment over cable and DSL infrastructures by performing necessary media conversions between IP and PSTN by converting digitized voice from an IP end-point into the PSTN.

2.10.3 Integrated Access Devices

Integrated Access Device (IAD) is the generic name for premise devices like Multimedia Terminal Adapters. IADs are meant to support multiple media types. A typical IAD will support simultaneous VoIP and data sessions. More sophisticated IADs perform resource allocation and coordination functions for better support of simultaneous sessions.

2.11 Software Components

2.11.1 Signaling Conversion

When a call originates on a VoIP terminal and terminates on a PSTN phone, signaling conversion is necessary to set up the call. VoIP signaling protocol messages need to be converted to either Q.931 messages and subsequently to SS-7 messages or to SS-7 messages directly. In carrier networks, separate Signaling Gateways usually perform the conversion function. In simpler implementations, an integrated gateway device can support both signaling and media conversions.

2.11.2 Application Server (i.e. accounting, billing, etc)

Application servers are also frequently called Feature Servers and are designed to further separate service logic from the network infrastructure. Typically, a server will contain all the software needed to perform a specific application – such as pre-paid calling card solutions or billing platforms, and will use the services of media gateways and policy servers (softswitches) to complete calls. This results in a highly distributed architectural model and allows application designers to concentrate on application logic.

2.11.3 Media Server

A media server is a programmable hardware device designed to process media events associated with the call. Carrier class media servers can handle media events for thousands of ports. Typical

events are playing customized announcements, supporting multi-user conferences, providing automatic speech recognition (ASR) and text-to-speech (TTS) capabilities. Simple tasks of this nature can be provided by media (access, trunking) gateways but media servers can scale to higher call volumes and implement additional functionality faster, with the same emphasis of separating call control and processing from supplemental functions. Typical carrier class media servers employ large volumes of DSPs to achieve the needed scalability.

2.11.4 Signaling Server Gateway

A signaling gateway performs signaling conversion from PSTN signaling, typically SS-7, to VoIP signaling (H.323, MGCP/Megaco or SIP) for calls originating on the PSTN and terminating of VoIP end-point and from VoIP signaling protocols to SS-7 for calls originating on VoIP and terminating on the PSTN. This is a demarcation point between the two signaling networks. A typical signaling gateway resides on a high availability or fault tolerant computer platform and interface with a policy server and media gateways on the VoIP side and STPs on the PSTN side.

2.11.5 Policy Server

Policy Servers are also called Call Agents and, somewhat loosely, Softswitches. This is a database server, which contains call processing scripts for various types of calls offered by the service provider. Typically, policy servers will receive signaling information from a signaling gateway on calls incoming from the PSTN and will select either a terminating gateway or pass routing instructions to an application server. Policy servers are an important element in separating signaling from media and call control. New types of services can be easily implemented by adding new scripts or, in some cases, new application servers to the database. These servers reside on high-availability/fault tolerant general purpose computer platforms and use off-the-shelf database products. They communicate with media gateways using either MGCP/Megaco or proprietary protocols and with application servers using either H.323 or SIP.

2.12 Media Encoding

When telephony started, the world was analog. The carbon microphone button in a handset varied the voltage of a current applied across the button as the user talked. The analog signal was connected by a series of wires, plugs, switches, and eventually relays to the earpiece, where a diaphragm connected to a coil of wire near a magnet reproduced the sound of the speaker for the listener.

Every since the 50's we have been digitizing voice and sending it as a series of bits. To do so, we nearly always encode the data using some form of compression, transmit a stream of bits, and then decode the stream prior to conversion back to analog.

Analog to Digital -> Encode → Transmission & Switching → Decode -> Digital to Analog

The en*co*de and *dec*code mechanisms, taken together, are often abbreviated as *codec*. Codecs have existed in telephony systems since the conversion to digital. Compression has also existing for just

as long, because the human voice has quite a range of volume (as well as pitch), but little ability to discriminate very small changes in volume. Because of these characteristics of the human ear, the Bell Labs engineers compressed the voice signal using a logarithmic coding mechanism. The specific mechanism they used is called "mu law" and uses a 3-bit characteristic and a 4-bit mantissa. A slightly different compression algorithm is used elsewhere in the world called "a law". The ITU has standardized this codec as G.711. G.711 has an 8KHz sample rate (the rate at which the analog to digital converter runs); the resulting data rate is 64Khz. Many references to G.711 claim that it is uncompressed, but in fact it does employ this logarithmic "compandor" mechanism to increase the dynamic range of the signal.

For simple speech, researchers have designed codecs that deliver the same or lower quality of speech reproduction with lower bit rates. Among the common codecs available today that have lower data rates are G.723.1, G.729 and G.726. Mobile systems use even more aggressive codecs such as AMR. These codecs compress speech to less than 3Khz, but not at the same quality as G.711. There are also codecs that deliver better than voice quality. These codecs are usually referred to as "wide band" and they have more range, precision and frequency response than G.711. Wideband codecs include G.722 and variations G.722.1 and G.722.2 also known as AMR-WB.

Video is also encoded using a different class of codec. Common video codecs include H.261 (often found in current technology videoconferencing systems), and MPEG-2. The latest video codecs are MPEG-4, a subset of which is called H.264.

So, past the codec, we have a bit stream at some data rate. In the VoIP world, we package data into blocks called packets. A packet is usually between 64 and a few thousand bytes. A packet has a header, which has the addressing and contents information that is used to get the packet from the sender to the correct recipient. The "payload" of a packet would contain the speech data. VoIP typically is deployed in a layered transmission environment, and each layer has its own header. If you were to look at bits on the wire, you would probably see:

1. The physical transport (layer 1) header. This header is dependent on the physical mechanism joining two parties. On Ethernet, for example, one would first find the Ethernet header. This header is changed from hop to hop, as it is common to find that many transports are involved in the path from the originating phone to the terminating phone.
2. The PPP header. This header is used on some transports where there are several hops between the two parties that have contractual relationships. PPP encapsulates packets that can be sent over a variety of transports, over several hops, until it gets to an entity that unwraps the encapsulation before it is routed onwards. PPP is not always present. It is often found between the residence and the Internet Service Provider, over transports such as DSL or modems. PPP can be used to route all packets from a single subscriber to a common point in the network before they are dispatched to the ultimate destination.
3. The IP header. This is the base protocol for the Internet. IP has the addresses of the source and destination of the packet. The IP header is (in the absence of NAT) created at the source, and unchanged as it arrives at the destination,
4. The UDP header. User Datagram Protocol is used for voice, unlike most web or file data, which uses Transmission Control Protocol, TCP. UDP is simpler and has much less overhead to process than TCP, although it also behaves poorer under congestion. UDP

implements "sockets" which allow a single address to handle multiple streams of data. The UDP header is always unchanged from source to destination.

5. RTP header. The Real Time Protocol provides timing, codec selection and error monitoring functions for media streams. The RTP header is always unchanged from source to destination.

A VoIP phone encodes a media stream, packetizes the stream and creates each of these headers for each of the media streams it sends out. The physical and PPP headers will be removed, and recreated, as the packet traverses the network. At the destination, the IP header will get the packet to the correct device, the UDP header will get it to the right "Call" and the RTP header will get it to the right media stream if there is more than one. The payload will then be extracted and played to the user.

2.13 Protocols

2.13.1 H.323

Originally conceived by the ITU-T to be the IP equivalent of H.320 for videoconferencing, H.323 is now widely deployed for such use. H.323 has also been used as a VoIP protocol, primarily in enterprise telephony applications. This is a peer-to-peer protocol. H.323 is entities include:

- Terminal – such as a videoconference system or a phone. Endpoints are moderately intelligent and a basic call can be completed between two endpoints without any other entity.
- Gateway – interfaces between H.323 systems and other systems such as the PSTN. A gateway typically appears to be a set of H.323 terminals to other H.323 terminals.
- Gatekeeper – a centralized registration and routing service. Terminals register with a local gatekeeper, and the gatekeeper typically routes all calls. H.323 systems are divided into "zones", each typically governed by a gatekeeper. Calls are routed by the source gatekeeper to the destination gatekeeper, and by the destination gatekeeper to the destination terminal/gateway. Once a call is established, media flows between endpoints, not through the gatekeeper. Hierarchical arrangements of gatekeepers are possible.
- Multipoint Conference Unit (MCU) – calls of more than two parties are bridged by an MCU. A terminal maintains a basic call to a port on the MCU and the MCU mixes media streams to deliver audio (and possibly video/data) to each endpoint.

H.323 consists of the main H.323 document, now at version 5, and several subsidiary protocols:

- H.225 is the call signaling protocol. There are three components, not all of which are used for all calls: "RAS", basic Call Signaling and "Annex G". The call signaling protocol is similar to ISDN Q.931
- H.245 is the "multimedia control protocol", mostly concerned with negotiating call parameters such as which codec will be used for media streams
- H.235 is the security specification for H.323
- H.350 defines a directory mechanism for finding endpoint identifiers
- H.450 defines "Supplementary Services" such Call Forward and Transfer

For further information on H.323, see http://www.packetizer.com/iptel/h323/

2.13.2 SIP

SIP is the IETF standard for peer-to-peer multimedia session control, where for this discussion a session is equivalent to a call. SIP is deployed primarily in VoIP applications, but is also the heart of the 3GPP2 wireless signaling standards, as well as Instant Messaging and Presence (SIMPLE). SIP entities include:

- User Agent – the endpoints of a call are User Agents. The originator of a call is the User Agent Client (UAC) and the terminator is the User Agent Server (UAS). A sip phone typically has both a UAS and a UAC so that it can place and receive calls. User Agents are typically very intelligent. Basic and advanced calls can be placed between two user agents without any other entities, although this is uncommon.
- Proxy Server – an intermediary in the signaling path, but not the media path. Proxy servers provide various call services to a SIP call, such as Registration, Redirection, and Routing. Signaling message typically traverse one or more proxy servers in the path from UAC to UAS. However, once a call is established, media flows directly between the UAC and the UAS.
- Media Servers – Provide media services to endpoints such as announcements, record/playback or conference bridging.

The SIP specification (RFC3261) defines the basic protocol, but there are a number of common extensions such as:

SIP-Specific Event Notification (RFC 3265)
The Session Initiation Protocol UPDATE Method (RFC 3311)
Private Extensions to the Session Initiation Protocol (SIP) for Asserted Identity within Trusted Networks (RFC 3325)
The Session Initiation Protocol (SIP) Refer Method (RFC 3515)
DHCP Option for SIP Servers (RFC 3361)

For more information on SIP, see:

http://www.cs.columbia.edu/sip/
http://www.softarmor.com/sipwg/

2.13.3 H.248/Megaco

This protocol is a joint effort of ITU and IETF. ITU calls it H.248, IETF calls it Megaco. The current release is RFC3015. An updated version is nearing completion. Version 2 can be found as an Internet Draft at http://www.ietf.org/internet-drafts/draft-ietf-megaco-h248v2-04.txt. Megaco is a Master Slave protocol. Megaco entities include:

- Media Gateway – the slave or client portion. The MG has media terminations as well as a signaling connection to the Media Gateway Controller. The MG makes connections among its terminations at the direction of the MGC.

- Media Gateway Controller – the master or server portion. The MGC controls one or more Media Gateways. MGCs often have connections to other MGCs, but they would not use

Megaco for such peer-to-peer connections. Instead, they would use a protocol such as ISUP (in the case of Bearer Independent Call Control, BICC, Q.1901), or SIP.

Megaco can be deployed in a softswitch architecture where the phones are MGs, and the softswitch is an MGC. This is the closest to existing PSTN with the phones being relatively non-intelligent, and having a master/slave relationship with the controlling MGC. In particular, a Media Gateway does not maintain call state. It does maintain media bearer state. Megaco is often deployed at a large gateway connecting to the PSTN where the cost of implementing a more complex protocol such as SIP at the scale of a gateway with 10s of thousands of ports may not be appropriate.

2.13.4 MGCP

Media Gateway Control Protocol is a direct predecessor of Megaco, but still enjoys significant deployment. Some vendors feel that Megaco deviated from the original MGCP too much, and is more complex than necessary. MGCP is defined with an Informational RFC (not Standards Track, as the SIP and Megaco standards are). The current version of MGCP is RFC3435. The MGCP entities are:

- Media Gateway – the slave or client portion. The MG has media endpoints as well as a signaling connection to the Call Agent. The MG makes connections among its endpoints at the direction of the CA.

- Call Agent – the master or server portion. The CA controls one or more Media Gateways. CA's often have connections to other CA's, but they would not use Megaco for such peer-to-peer connections. Instead, they would use a protocol such as SIP.

MGCP is often deployed between gateways, residential or trunk, and call controllers (softswitches). Like in Megaco, MGCP Media Gateways do not maintain call state, Call Agents do. Thus MGCP Media Gateways are relatively non-intelligent.

2.13.5 Others

Bearer Independent Call Control (BICC) extends SS7 signaling protocols (ISUP) to specify a VoIP or Voice over ATM bearer. The idea with BICC is to retain the entire SS7 signaling mechanism, but use it to control VoIP devices. The call signaling would traverse the existing SS7 network; the bearers would traverse IP or ATM networks. Q.1901 defines BICC as a series of minor enhancements to Q.761-764. There has yet to be significant deployment of BICC.

Skinny – Cisco Systems proprietary phone protocol. Many installations of the Cisco Call Manager use a proprietary protocol between the phones and the Call Manager known as the Skinny Client Control Protocol (SCCP).

TRIP – Telephony Routing over IP. RFC3219 describes a protocol used to identify which gateway should be used to send a call. Using a variation of the BGP route advertising protocol that is used in IP routers, TRIP is used where there are many gateways in a network, more than one of which could reach a desired phone number. TRIP would provide the information to select the "best" gateway.

2.14 Security

Security in a VoIP environment is significantly complicated by the possibility that the network used will be, all, or in part, the global Internet, and thus subject to the kinds of attacks we are all painfully aware of. Thus, most VoIP protocols include sophisticated security mechanisms primarily based on cryptographic algorithms. This document could not hope to thoroughly discuss the ins and outs of cryptographically based security mechanism, and the reader is directed to texts such as Bruce Schneier's "Applied Cryptography".

Generally, one sees mechanisms that provide:

1. Authentication - assuring the receiver of the identity of the sender.
2. Integrity - assuring the receiver that messages have not been modified by a third party, dropped, or additional messages inserted.
3. Secrecy - assuring that messages cannot be read by unauthorized entities
4. Nonrepudiation - assuring that a sender cannot falsely deny later that a message was sent.
5. Protection from Denial of Service - assuring that desired messages can be received in the face of an attack that seeks to deny the ability to do so.

Significantly, the mechanisms employed by VoIP typically do not depend on separation of facilities, or by "security by obscurity", which have proven to be problematic to maintain and deliver effective security.

It is important to recognize that while modern protocols such as Megaco and SIP have "mandatory to implement" security mechanisms written into the specification, it is still uncommon to see such implementations, and much less common to see them enabled. Nevertheless, we expect to have more of the security mechanisms fully deployed in VoIP systems in the near future.

IETF protocols (Megaco, MGCP and SIP) generally deploy one or more of the IETF developed security mechanisms such as TLS (RFC2246), IPSEC (RFCs 2401-2406), or S/MIME (RFC2632/2633). Using predefined mechanisms increases the probability of security because these mechanisms are widely deployed, and have been studied extensively. Inventing new mechanisms and algorithms is fraught with difficulty and many a "new improved" security mechanism has been found to have fatal flaws.

It is also important to recognize that security mechanisms may be deployed "end to end" or "hop by hop". In the former, the originator of a message employs the security mechanisms which are passed unmodified (and usually unverified) by any intermediary entities until the message reaches the ultimate recipient. The recipient performs the cryptographic mechanisms to assure authentication/integrity/.... In the latter, the mechanisms are deployed between adjacent entities. For example, a message send from a SIP phone, through two proxy servers to another sip phone could use TLS (which replaces SSL) security between the originating SIP phone and the first proxy, IPSEC between the two proxies, and TLS again between the second proxy and the destination SIP phone.

Some specifics:

1. Megaco (H.248, currently RFC3015) specifies IPSEC as mandatory-to-implement
2. MGCP (RFC3435) has no mandatory mechanism, but recommends IPSEC
3. SIP (RFC3261) specifies TLS as mandatory to implement hop-by-hop, and uses S/MIME for end-to-end security. SIP also uses HTTP's "Digest" authentication mechanism, which is fairly widely deployed.

Mechanisms to prevent denial of service attacks are also built into these protocols, but can be overwhelmed by rogue elements that can generate valid messages. Specific implementations can provide throttling mechanisms to mitigate this kind of problem.

2.15 Reliability

Reliability in the traditional sense of TDM-switched service is often judged by the “five nine’s” principle of 99.999% availability. With VoIP services, the operating model deployed for a certain environment can influence reliability. The standard variations of VoIP service are set forth below with considerations as to reliability for each.

2.15.1 Trunk Replacement Reliability

Reliability can be sustained at high levels in a trunk replacement model, due to the nature of IP transport characterized in section 2.10. TDM trunks are replaced with IP Trunks over a managed IP network. End users in this environment will most likely be unaware there are placing calls over an IP network. In additional to the IP network, most trunk replacement models have TDM overflow and fail over for redundancy.

2.15.2 Hybrid Architecture Reliability

Service providers may also offer VoIP on a one-way basis, offering just an origination or termination point for VoIP traffic to hop on or off the PSTN.

This model can be used in Greenfield applications where no legacy switch exists, or a Class 5 switch can be present, augmented by the softswitch or voice application server. Media Gateways may be used, but this model could be used in a native IP environment where no media conversion is necessary. Reliability hinges on redundancy of proxy servers and elimination of “single point of failure” in the originating/terminating Managed IP network, and the hand from the PSTN and the MGC.

2.15.3 End-to-End VoIP Reliability

IP networks are inherently distributed and resilient. A truly unified voice system can be distributed across multiple sites by using a peer-to-peer architecture that has no single point of failure.

IP voice switches designed specifically for voice can each incorporate a complete call processor. Each switch is a peer with a full complement of routing information safely held in local memory, and can operate as a standalone switch if its site is cut off from the IP backbone. It can make best-effort calls on its own, using a failover PSTN trunk if necessary. When switches are added or

restored to the network, they and the existing switches at all the sites discover each other and start working together.

If a switch providing PSTN access to one site were to fail, its peer switches elsewhere in the IP network would provide alternate PSTN access to the users at that site. As long as the data backbone stays up, this type of distributed voice network can't have an outage unless all the switches go down simultaneously. Reliability thus comes built in with this approach, and the "five-nines" availability requirement for voice is met. In fact, it can be increased to ten nines by installing a redundant switch with PSTN access at each site.

2.16 Programmatic Interfaces/API

In VoIP, numerous API's have been developed to facilitate the control of service creation, deployment, and management. Often application servers provide components of the operational support systems, including accounting, billing and provisioning.

2.16.1 Session Initiated Protocol – Common Gateway Interface (SIP-CGI)

Like HTTP CGI, a SIP CGI script resides in the server and passes message parameters through environment variables to a separate process. The key point here is that a new process is created on the server adding to the cost of overhead. When a request arrives at a SIP server, initiating a new transaction, the server will set a number of environment variables, and call a CGI script. The script is passed the body of the request through its standard input (*stdin*) and returns instructions back to the server through its standard output file descriptor. SIP CGI is almost identical to HTTP CGI and is particularly suitable for services that contain substantial web components. A CGI script can be written in Perl, Tcl, C, C++ or Java making it accessible to a large community of developers.

2.16.2 Call Processing Language (CPL)

A topic in VoIP service creation is how vendors can provide service providers with the ability to offer basic call services and allow the end user the comfort of customizing service. This will bring about many different combinations of call services that will require some level of standardization.

The IETF IPTL (Internet Engineering Task Force's IP Telephony Work Group) has proposed an XML-based language - CPL (Call Processing Language) to handle the vastly different service combinations as well as the different call scenarios.

Call Processing Language is a language that can be used to describe and control IP based Telephony services. CPL is by nature flexible and can be implemented on the Network Server or an Agent User Server, such as a SIP server. CPL is editable via Graphical User Interfaces and is independent of operating system or signaling protocol.

Some features of CPL are:

- Easily written and parsed
- Easily verified across all components
- Extensible as more XML tags can be added
- Easy to implement

The following diagram depicts a high level representation of a SIP architecture based on CPL. [Call Processing Language Based Service Configuration System, by Mahavir D. Karnavat and Shivaji Hogale]

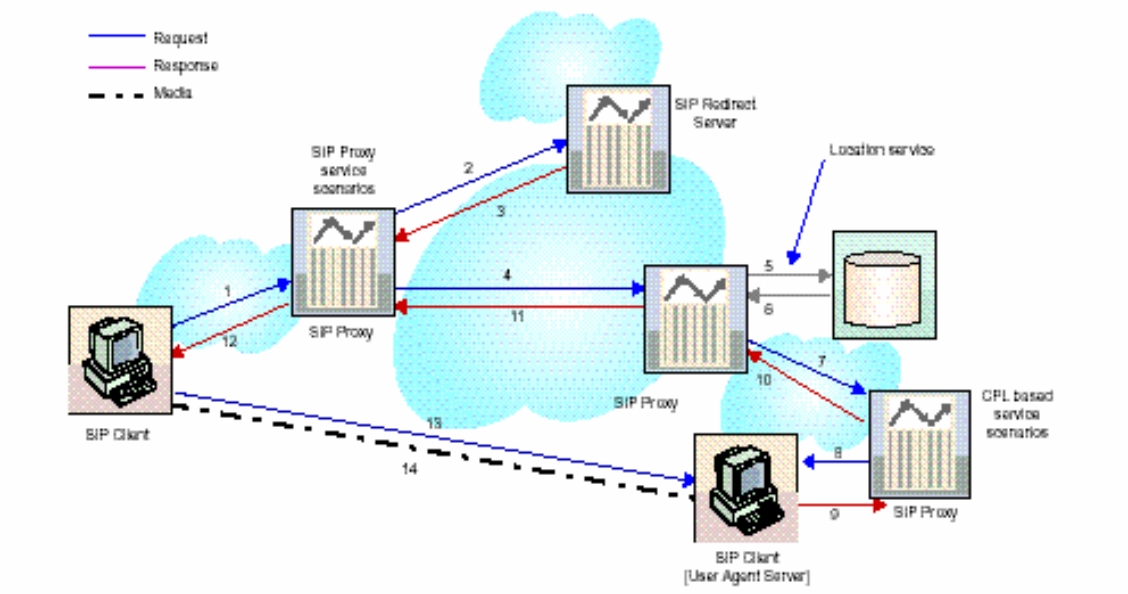

Open issues with CPL are:

- CPL is currently in the draft phase with most of the work being enhancements, but the framework is in place.
- SIP and H.323 APIs are not standardized due to the need for the ability to customize.

2.16.3 JAIN

JAIN, the Java API for Integrated Networks, is ready to apply Java's "Write once, run anywhere" philosophy to the telecommunication industry. JAIN's objective is to bring all of the proprietary telecommunication networks into one coherent network. The JAIN SIP specification is a general-purpose transaction based Java interface to the SIP protocol. It is rich both semantically and in definition to the SIP protocol. The motivation behind JAIN SIP is to develop a standard interface to the SIP protocol that can be used independently or by higher level programming entities and environments. JAIN SIP can be used in multiple ways. JAIN SIP provides a standardized interface that can be used by communications developers as a minimum to support SIP in their applications. The JAIN SIP reference implementation provides a fully functional SIP implementation that can be used by developers to talk SIP from the Java environment.

2.16.4 SIP Servlets

A Servlet is a small program that runs on a server. SIP Servlets are similar to the CGI concept but, instead of using a separate process, messages are passed to a class that runs within a JVM (Java

Virtual Machine) inside the server. SIP Servlets are very similar to HTTP Servlets; they simply enhance the interface to support SIP functions. Because they are written in Java, servlets are portable between servers and operating systems. SIP Servlets are restricted to Java, but suffer less overhead than SIP CGI. Use of a JVM for executing servlets means that the Java "sandbox" concept can be applied to protect the server from the script. Like SIP CGI, SIP Servlets closely mirror the operation of HTTP Servlets; they simply enhance the interface to support the wider array of functions a proxy can execute, as compared to an HTTP origin server.

2.17 Telephone Number Mapping

2.17.1 ENUM

While not a protocol, ENUM is important to VoIP. The mechanism in the Internet that is used to convert a "domain name" (like nena.org) to the IP address of the server that supports nena.org is called the "Domain Name System" or DNS. DNS is a large distributed database. ENUM extends the DNS to provide E.164 phone number to a URI (Uniform Resource Identifier). Protocols such as SIP and H.323 can directly use a URI to route a call. A URI can also be used to determine the carrier serving that phone number. In this way, ENUM is analogous to the LNP TCAP query, but with much more utility. Given an E.164 in ENUM, one can determine if the number is a VoIP endpoint, reachable from a VoIP endpoint via a gateway, and which carrier presently serves that number. ENUM is global. ENUM is defined in RFC2916. For more information, see http://www.enum.org/ or http://www.itu.int/osg/spu/enum/

3 References

Henning Schulzrinne, Internet Telephony, 0-8493-0052-5. © CRC Press LLC

Cisco Systems, Inc. Security in SIP-Based Networks, http://www.cisco.com/en/US/tech/tk652/tk701/technologies_white_paper09186a00800ae41c.shtml

THE SIP CENTER, http://www.sipcenter.com/files/SIPOverview.pdf

THE SIP FORUM, http://www.sipforum.org/

SUN MICROSYSTEMS, http://java.sun.com/products/jain/

JAVA WORLD, http://www.javaworld.com/javaworld/javaone00/j1-00-jain.html

Karnavat D. Mahavir, Hogale Shivaji. Call Processing Language (CPL) Based Service Configuration System, http://www.wipro.com/insights/callprocessinglanguage.htm

Susan Breidenbach, www.goshoreline.com/challenge/PDF/Reliability%20Whitepaper.pdf

Jason Pachomski, TechRepublic TCP/IP Resources, © 2001