CHANNEL DYNAMICS
leverage through partnership™

Voice over IP (VoIP)

A background to the technology and its applications

Prepared by

Cam Wayland
Channel Dynamics Pty Ltd
March 2005

Executive Summary

IP Telephony, also called 'Internet telephony', is the technology that makes it possible to have a voice conversation over the Internet or a dedicated Internet Protocol (IP) network instead of dedicated voice transmission lines. This allows the elimination of circuit switching and the associated waste of bandwidth. Instead, packet switching is used, where IP packets with voice data are sent over the network only when data needs to be sent, i.e., when a caller is talking.

Its advantages over traditional telephony include:

Unlike traditional telephony VoIP innovation progresses at market rates rather than at the rates of the multilateral committee process of the International Telecommunications Union (ITU)
Lower costs per call, or even free calls, especially for long-distance calls.
Lower infrastructure costs for business & telcos: once IP infrastructure is installed, no or little additional telephony infrastructure is needed.
New advanced calling and user features included in the network
"Future proof" business phone systems since the functionality is software (protocol) based and does not require hardware replacement.
Voice over IP traffic does not necessarily have to travel over the global Internet; it may also be deployed on private IP networks for example on a LAN inside a single building.

The protocols used to carry the signal over the IP network are commonly referred to as Voice over IP or VoIP protocols.

VoIP Protocols

There are a few VoIP protocol stacks which are derived from a variety of standard bodies and vendors, namely H.323, SIP, MEGACO and MGCP.

H.323 is the ITU-T's customary, which was originally developed for multimedia conferencing on LANs, but was later extended to cover Voice over IP. The standard encompasses both point to point communications and multipoint conferences. H.323 defines four logical components: Terminals, Gateways, Gatekeepers and Multipoint Control Units (MCUs). Terminals, gateways and MCUs are known as endpoints. H322 is becoming less popular due to its complexity of implementation across multiple vendors equipment.

Session Initiation Protocol (SIP) is the IETF's regular for establishing VoIP connections. SIP is an application layer control protocol for creating, modifying and terminating sessions with one or more participants. The architecture of SIP is similar to that of HTTP (client-server protocol). Requests are generated by the client and sent to the server. The server processes the requests and then sends a response to the client. A request and the responses for that request make an operation. SIP has been the turning point in the increasing popularity of VoIP due to its simplicity and interoperability therefore driving the costs of hardware down.

Media Gateway Control Protocol (MGCP) is a Cisco and Telcordia proposed VoIP protocol that defines communication between call control elements (Call Agents or Media Gateway) and telephony gateways. MGCP is a control protocol, allowing a central coordinator to observe events in IP phones and gateways and instructs them to send media to specific addresses. In the MGCP architecture, the call control intelligence is located outside the gateways and is handled by the call control elements (the Call Agent). Also the call control elements (Call Agents) will bring into line with each other to send sound commands to the gateways under their control.

Codecs

With respect to voice over IP, a codec is an algorithm used to encode and decode the voice conversation. Since voice and sound as we hear it is analogue, it needs to be converted (or encoded) to a digital format suitable for transmission over the Internet. Once at the other end, it needs to be decoded again so the other person can hear what you are saying. There are a variety of different ways this encoding and decoding can be done - many of which utilise compression in order to reduce the required bandwidth of the conversation. A key thing to remember with VoIP, is that encoding, particularly when heavy compression is used, takes time, which adds a delay to the conversation. Thus, the holy grail is a codec which not only maintains good quality with compression, but is able to do the encoding and decoding in a minimal amount of time.

Here are some of the codecs with their bandwidth compared

- GIPS - 13.3 Kbps and up
- GSM - 13 Kbps (full rate), 20ms frame size
- ITU G.711 - 64 Kbps, sample-based Also known as alaw/ulaw
- ITU G.722 - 48/56/64 Kbps
- ITU G.723.1 - 5.3/6.3 Kbps, 30ms frame size
- ITU G.726 - 16/24/32/40 Kbps
- ITU G.728 - 16 Kbps
- ITU G.729 - 8 Kbps, 10ms frame size
- Speex - 2.15 to 44.2 Kbps
- LPC10 - 2.5 Kbps
- DoD CELP - 4.8 Kbps

The most popular of the codecs are G711, G723.1, G726 and G729a, with G729a proving to be the most popular in the voice over internet market.

VoIP for business

To set up a business VoIP system, you need several components. A central device manages the calls, the way a private branch exchange (PBX) or key system unit (KSU) does in traditional phone systems. This can be a dedicated piece of hardware such as an IP PBX, a regular PBX that has been IP-enabled, or a server running specialized software. You will also need phones and a data network. In many cases,

you may be able to use your existing digital phones and computer network, although you may need to upgrade some of your network hardware.

The most prominent benefit of an IP PBX is for businesses with multiple locations. By the advantage of VoIP, all the offices on a LAN or WAN can get the profits of having a common office phone system. The profits are - extension dialling, seamless call transfers, and other features.

In addition to making it easier to communicate, this sharing of features can enhance collaboration as employees at different locations can truly feel like they are part of the same organization. Plus, if they are on the company network, the phone calls are free - even if your offices are located thousands of miles apart. Really! It is an advantage in saving money calling between two branches of the same office.

Computer networks are designed to handle messy data: packets arrive out of order and some are even lost, but in most cases the data being sent can easily be reconstructed before it is needed. Voice conversations, though, are not as tolerant of these kinds of disturbances. Each packet of sound has to arrive in the correct order because they are being sent in real time - if packets are lost, the conversation sounds distorted, choppy, or falls off all together. This is why VoIP services that rely on the Internet to transmit calls can have uneven phone quality.

The selection of an enterprise VoIP solution is a major decision. Voice service is critical to the operation of the business, so no one wants to implement a technology that will compromise call quality or reliability in any way. On the other hand, the cost savings and value-added functionality available with VoIP makes it a compelling investment.

The VoIP phone system has a benefit for companies having multiple locations branches, telecommuters and remote sales offices. And the locations are connected with a company's Local Area Network (LAN) or Wide Area Network (WAN). In that case the companies are suitable for using a VoIP system. You can share the full features of your phone system across all your locations. In addition, even if you have one office in one place and one in different place, VoIP allows calls between them via extension dialling, making it a zero cost call. For businesses with hefty monthly long distance charges due to calls between locations, that can be a very attractive reason to upgrade.

VoIP for SOHO & Consumers

VoIP can turn a standard Internet connection into a way to place free or very low cost phone calls. The practical upshot of this is that by using some of the free VoIP software that is available to make Internet phone calls, you are bypassing the phone company (and its charges) entirely.

VoIP is a revolutionary technology that has the potential to completely rework the world's phone systems. VoIP providers like Vonage (USA) and engin (Australia) have already been around for a little while and are growing steadily. Major carriers around the world are already setting up VoIP calling plans and the FCC in the US and the

ACA in Australia are both looking seriously at the potential ramifications of VoIP service.

Above all else, VoIP is basically a clever "reinvention of the wheel." In this article explains some of the principles behind VoIP, its applications and the potential of this emerging technology, which will more than likely one day replace the traditional phone system entirely.

Calling

The interesting thing about VoIP is that there is not just one way to place a call. There are three different "flavours" of VoIP service in common use today:

ATA - The simplest and most common way is through the use of a device called an ATA (analogue telephone adaptor). The ATA allows you to connect a standard phone to your computer or your Internet connection for use with VoIP. The ATA is an analogue-to-digital converter. It takes the analogue signal from your traditional phone and converts it into digital data for transmission over the Internet. Some providers in the US are already bundling ATAs free with their service, similar to a mobile phone. You simply crack the ATA out of the box, plug the cable from your phone that would normally go in the wall socket into the ATA, connect another cable to the DSL router and you're ready to make VoIP calls. Some ATAs may ship with additional software that is loaded onto the host computer to configure it; but in any case, it is a very straight forward setup.

IP Phones - These specialized phones look just like normal phones with a handset, cradle and buttons. But instead of having the standard RJ-11 phone connectors, IP phones have an RJ-45 Ethernet connector. IP phones connect directly to your router and have all the hardware and software necessary right onboard to handle the IP call. Soon, Wi-Fi IP phones will be available, allowing subscribing callers to make VoIP calls from any Wi-Fi hot spot.

Computer-to-computer - This is certainly the easiest way to use VoIP. You don't even have to pay for long-distance calls. There are several companies offering free or very low-cost software that you can use for this type of VoIP i.e. Skype. All you need is the software, a microphone, speakers, a sound card and an Internet connection, preferably a fast one like you would get through a cable or DSL modem. Except for your normal monthly ISP fee, there is usually no charge for computer-to-computer calls, no matter the distance. One downside of this type of connection is that the voice quality is sometimes not as good as the ATA due to the specialised voice encoding hardware in the ATA. The other is that the person you are calling has to have another computer and be connected to the Internet at the time of the call.

But chances are good you are already making VoIP calls any time you place a traditional long-distance call. Phone companies use VoIP to streamline their networks. By routing thousands of phone calls through a circuit switch and into an IP gateway, they can seriously reduce the bandwidth they're using for the long haul. Once the call is received by a gateway on the other side of the call, it is decompressed, reassembled and routed to a local circuit switch.

Although it will take some time, you can be sure that eventually all of the current circuit-switched networks will be replaced with packet-switching technology (more on packet switching and circuit switching later). IP telephony just makes sense, in terms of both economics and infrastructure requirements. More and more businesses are installing VoIP systems, and the technology will continue to grow in popularity as it makes its way into our homes.

VoIP Features

The Forrester Research Group predicts that nearly 5 million U.S. households will have VoIP phone service by the end of 2006. Perhaps the biggest draws to VoIP for the home users that are making the switch are price and flexibility.

With VoIP, you can make a call from anywhere you have broadband connectivity. Since the IP phones or ATAs broadcast their info over the Internet, they can be administered by the provider anywhere there is a connection. So business travellers can take their phones or ATAs with them on trips and always have access to their home phone. Another alternative is the softphone. A softphone is client software that loads the VoIP service onto your desktop or laptop. The Vonage softphone has an interface on your screen that looks like a traditional telephone. As long as you have a headset/microphone, you can place calls from your laptop anywhere in the broadband-connected world.

Most VoIP companies are offering minute-rate plans structured like cell phone bills for as little as $30 per month. On the higher end, some offer unlimited plans for $79. With the elimination of unregulated charges and the suite of free features that are included with these plans, it can be quite a savings.
Most VoIP service providers or companies provide the features that normal phone companies charge extra for when they are added to your service plan. VoIP includes:

- Caller ID
- Call waiting
- Call transfer
- Repeat dial
- Return call
- Three-way calling

There are also advanced call-filtering options available from some carriers. These features use caller ID information to allow you make a choice about how calls from a particular number are handled. You can:

- Forward the call to a particular number
- Send the call directly to voicemail
- Give the caller a busy signal
- Play a "not-in-service" message
- Send the caller to a funny rejection hotline

With many VoIP services, you can also check voicemail via the Web or attach messages to an e-mail that is sent to your computer or handheld. Not all VoIP services offer all of the features above. Prices and services vary, so if you're interested, it's best to do a little shopping.

Now that we've looked at VoIP in a general sense, let's look more closely at the components that make the system work. In order to understand how VoIP really works and why it's an improvement over the traditional phone system, it helps to first understand how a traditional phone system works.

The Standard Phone System: Circuit Switching

Existing phone systems are driven by a very reliable but somewhat inefficient method for connecting calls called circuit switching.
Circuit switching is a very basic concept that has been used by telephone networks for more than 100 years. When a call is made between two parties, the connection is maintained for the duration of the call. Because you are connecting two points in both directions, the connection is called a circuit. This is the foundation of the Public Switched Telephone Network (PSTN).

Here's how a typical telephone call works:

1. You pick up the receiver and listen for a dial tone. This lets you know that you have a connection to the local office of your telephone carrier.

2. You dial the number of the party you wish to talk to.

3. The call is routed through the switch at your local carrier to the party you are calling.

4. A connection is made between your telephone and the other party's line using several interconnected switches along the way.

5. The phone at the other end rings, and someone answers the call.

6. The connection opens the circuit.

7. You talk for a period of time and then hang up the receiver.

When you hang up, the circuit is closed, freeing your line and all the lines in between. Let's say that you talk for 10 minutes. During this time, the circuit is continuously open between the two phones. In the early phone system, up until 1960 or so, every call had to have a dedicated wire stretching from one end of the call to the other for the duration of the call. So if you were in Sydney and you wanted to call Melbourne, the switches between Sydney and Melbourne would connect pieces of copper wire half the way across the Australia. You would use all those pieces of wire just for your call for the full 10 minutes. You paid a lot for the call, because you actually owned a 1,000 Km long copper wire for 10 minutes.

Telephone conversations over today's traditional phone network are somewhat more efficient and they cost a lot less. Your voice is digitized, and your voice along with thousands of others can be combined onto a single fibre optic cable for much of the journey (there's still a dedicated piece of copper wire going into your house, though). These calls are transmitted at a fixed rate of 64 kilobits per second (Kbps) in each direction, for a total transmission rate of 128 Kbps. Since there are 8 kilobits (Kb) in a kilobyte (KB), this translates to a transmission of 16 KB each second the circuit is open, and 960 KB every minute it's open. So in a 10-minute conversation, the total transmission is 9,600 KB, which is roughly equal to 10 megabytes. If you look at a typical phone conversation, much of this transmitted data is wasted.

While you are talking, the other party is listening, which means that only half of the connection is in use at any given time. Based on that, we can surmise that we could cut the file in half, down to about 4.7 MB, for efficiency. Plus, a significant amount of the time in most conversations is dead air -- for seconds at a time, neither party is talking. If we could remove these silent intervals, the file would be even smaller. Then, instead of sending a continuous stream of bytes (both silent and noisy), what if we sent just the packets of noisy bytes when you created them? That is the basis of a packet-switched phone network, the alternative to circuit switching.

The VoIP Phone System: Packet Switching

Data networks do not use circuit switching. Your Internet connection would be a lot slower if it maintained a constant connection to the Web page you were viewing at any given time. Instead, data networks simply send and retrieve data as you need it. And, instead of routing the data over a dedicated line, the data packets flow through a chaotic network along thousands of possible paths. This is called packet switching.

While circuit switching keeps the connection open and constant, packet switching opens a brief connection -- just long enough to send a small chunk of data, called a packet, from one system to another. It works like this:

The sending computer chops data into small packets, with an address on each one telling the network devices where to send them.

Inside of each packet is a payload. The payload is a piece of the e-mail, a music file or whatever type of file is being transmitted inside the packet.

The sending computer sends the packet to a nearby router and forgets about it. The nearby router sends the packet to another router that is closer to the recipient computer. That router sends the packet along to another, even closer router, and so on.

When the receiving computer finally gets the packets (which may have all taken completely different paths to get there), it uses instructions contained within the packets to reassemble the data into its original state.

Packet switching is very efficient. It lets the network route the packets along the least congested and cheapest lines. It also frees up the two computers communicating with each other so that they can accept information from other computers, as well.

The Advantage

VoIP technology uses the Internet's packet-switching capabilities to provide phone service. VoIP has several advantages over circuit switching. For example, packet switching allows several telephone calls to occupy the amount of space occupied by only one in a circuit-switched network. Using PSTN, that 10-minute phone call we talked about earlier consumed 10 full minutes of transmission time at a cost of 128 Kbps. With VoIP, that same call may have occupied only 3.5 minutes of transmission time at a cost of 64 Kbps, leaving another 64 Kbps free for that 3.5 minutes, plus an additional 128 Kbps for the remaining 6.5 minutes. Based on this simple estimate, another three or four calls could easily fit into the space used by a single call under the conventional system. And this example doesn't even factor in the use of data compression using an efficient codec such as G729a at 8Kbs, which further reduces the size of each call.

Let's say that you and your friend both have service through a VoIP provider. You both have your analogue phones hooked up to the service-provided ATAs. Let's take another look at that typical telephone call, but this time using VoIP over a packet-switched network:

1. You pick up the receiver, which sends a signal to the ATA.

2. The ATA receives the signal and sends a dial tone. This lets you know that you have a connection to the Internet.

3. You dial the phone number of the party you wish to talk to. The tones are converted by the ATA into digital data and temporarily stored.

4. The phone number data is sent in the form of a request to your VoIP company's call processor. The call processor checks it to ensure that it is in a valid format.

5. The call processor determines to whom to map the phone number. In mapping, the phone number is translated to an IP address (more on this later). The soft switch connects the two devices on either end of the call. On the other end, a signal is sent to your friend's ATA, telling it to ask the connected phone to ring.

6. Once your friend picks up the phone, a session is established between your computer and your friend's computer. This means that each system knows to expect packets of data from the other system. In the middle, the normal Internet infrastructure handles the call as if it were e-mail or a Web page. Each system must use the same protocol to communicate. The systems implement two channels, one for each direction, as part of the session.

7. You talk for a period of time. During the conversation, your system and your friend's system transmit packets back and forth when there is data to be sent. The ATAs at each end translate these packets as they are received and convert them to the analogue audio signal that you hear. Your ATA also keeps the circuit open between itself and your analogue phone while it forwards packets to and from the IP host at the other end.

8. You finish talking and hang up the receiver.

9. When you hang up, the circuit is closed between your phone and the ATA.

10. The ATA sends a signal to the soft switch connecting the call, terminating the session.

Probably one of the most compelling advantages of packet switching is that data networks already understand the technology. By migrating to this technology, telephone networks immediately gain the ability to communicate the way computers do.

VoIP Drawbacks

Since regular phones get all the power they need through the phone line, they continue to work if there is a power outage. In contrast, most VoIP phones need to be plugged into a power source to work. By definition, VoIP phones are also network-dependent.

To businesses where phone service is absolutely critical, this can be a concern since computer networks can occasionally be brought down by a server crash or other problem. However a good IT staff can prevent most outages and react quickly when one occurs.

It is another issue regarding pinpoint the geographic location of an IP address. In VoIP the call information is provided via an IP address. This is necessary in the case of emergency calls. Although some VoIP providers now do have call routing based from a specified address. In the US some VoIP Service Providers have implemented a system known as E911, a more advanced version of the US emergency 911 system, which provides additional location data over VoIP to emergency responders, including street address and the floor inside a building. In Australia the ACA and the government are still considering industry submissions and it is anticipated that mid to late 2005 there will be legislation covering VoIP.

Limitations

One challenge to maintaining call quality is bandwidth: high quality sound requires quite a bit of it. The technology to compress audio and to reconstruct it has been improved to the point where VoIP sound quality over a high-bandwidth connection is as good as or better than that of regular phones. But some networks that are fine for data are not up to the demands of VoIP.

Other limitation of the voice over internet protocol is the inability to use these lines for fax capabilities and a difficulty downloading electronic program that guides to personal video recorder products (such as the new Foxtel devices). Since VoIP runs through a computer, it does not function independently, meaning that VoIP will not function during a power outage, whereas regular analogue telephones do that. Many have made the argument that until VoIP is capable of offering 'all' services that PSTN can offer, that they will not be making the switch.

VoIP security issue

The legal and security issue is another drawback of the voice over internet protocol. Being the mater of information technology and cyber world what are the law exit are not sufficient to meet the legal and security issues.

It will still be at least a decade before communications companies can make the full switch over to VoIP. As with all emerging technologies, there are certain technical and regulatory hurdles that have to be overcome.

Acronyms and Definitions

Here are a few of the Internet telephony related acronyms, definitions, and terms.

ATA Analogue Telephone Adapter. These devices allow you to use standard telephones for VoIP. Think of them as a translator between what comes out of your analogue telephone and the digital voice over IP system.

CNG Comfort Noise Generation. Many VoIP devices offer silence suppression - where rather than using bandwidth to transmit silence, a message is transmitted saying there is no sound. The device at the other end then uses comfort noise generation to create background noise (since when one talks on the phone, and it is silent at the other end, it is generally not absolutely quiet, there is some background noice). Thus, CNG is an attempt to make the silence seem more natural.

DID Direct Inward Dialling (sometimes referred to as a 'virtual number'). This is a normal phone number on the PSTN (which can be in Australia, or another country) that can be called by any phone worldwide, and forwards to your VoIP phone. Many VoIP service providers offer direct inward dialling as part of their service, or as an add-on. There are also some third party companies that offer direct inward dialling, and will forward the call to your VoIP account no matter which service provider you are with.

e164 e164 offers a method of using your current telephone number as a means of being contacted in the IP world. Your telephone number gets mapped into a DNS zone, and the zone can contain your contact information (VoIP, instant messenger, email, anything). See e164.org for further details.

FXO Foreign Exchange Office. The FXO port is what the cord in the back of your telephone comes out of (on the telephone side). Your phone provides an FXO interface to the telephone network.

FXS Foreign Exchange Subscriber. Subscriber equipment (ie., telephones) plug into FXS ports. For example, the telephone port in the wall, is an FXS port.

H.323 A VoIP protocol that preceded SIP. Unlike SIP, the H.323 standard specifies the complete voice over IP protocol, and not just the signalling methods.

IAD Integrated Access Device. All-in-one VoIP adapters that also include a router and possibly an ADSL modem are known as IADs.

NAT Network Address Translator. Defined in RFC 1631, network address translation allows one public IP address to be shared between a number of devices. This is done by the device with the public address, acting as a gateway between an Internal network (running on private IP addresses - RFC 1918) and making all requests coming from the Internal network appear as though they are coming from the gateway device itself. It is commonly used in home networks.

PSTN Public Switched Telephone Network. The standard telephone network used today.

QoS Quality of Service. Exactly like it sounds. It is a way to mark some Internet traffic as being more urgent. For example, one might mark voice traffic as being "important to get there in real time, but if there is a huge delay, just drop the packets", where as downloading emails could have a QoS saying "just get the data when there is unused space on the line to my house". thus if you use both applications, whenever you use VoIP, it will be guaranteed a higher priority than emails. QoS is not implemented on an Internet-wide basis, though there is a lot of hardware that can deal with QoS parameters.

SIP Session Initiation Protocol. A protocol allowing a user to 'call' another user on their Internet phone. Think of SIP as a way to call someone, rather than 'dialling' their IP address, you call a phone number which the SIP proxy (the service provider) uses to identify them, and forward them an 'invite' message to the conversation. You can find out more about SIP from http://www.voip-info.org/wiki-SIP.

STUN Simple Traversal of UDP through NATs. STUN allows SIP devices behind NATs to discover their public IP addresses (that of the gateway), and the type of NAT in use. This is one step in creating a system whereby the VoIP device can be contacted by other Internet telephones.

VAD Voice activity detection. This feature allows VoIP clients to detect when the person is speaking versus background noise. If the VAD algorithm is not sophisticated enough, and silence suppression is enabled, then some of your voice might get cut off during a conversation.

VoIP Voice over Internet Protocol. Used to refer to all types of Internet Telephony. It encompasses just about any system that carries voice traffic over the Internet. For home users, it provides the ability to talk to other VoIP users without having to pay standard telephone usage charges (although you still have to pay for your Internet connection).

SIP and H.323 are the two main standards for VoIP, although SIP is becoming the most popular due to its generic nature, and the fact that unlike H.323, it was designed specifically for wide area Internet Telephony. There are also some custom protocols such as that used by Skype, and the talk feature on AOL Instant Messenger.