# *BandTel's Flawless VoIP -- June 2006*

Modified and edited by J. Scott Jan 2007

## *How VOIP Got Its Start -- Early 1990's*

Back in the mid 90's a group of technology enthusiasts created a way to encode voice into IP packets, and then transport that voice data over IP networks in real time. It soon acquired a fairly large cult following. The lure was that it enabled an individual to beat the incumbent phone company on a global basis, and let one talk to his buddy for free. That's right, saving money was big for these geeks. Since its inception, saving money has always been an important motivation for those considering moving their traditional TDM/PBX environment to VOIP systems.

This VoIP cult following was very similar in spirit to the free software movement insofar as it took the large incumbent out of the picture, and enabled the average user to communicate with anyone on the planet for next to nothing.

## *Bridging Two Worlds – Impediments to VOIP implementations*

The mid 90's average user's Internet access was at a rate of 28kbs-56kbs, and typically closer to the 28kbs level. As such, trying to get a full 64kbs phone call using the industry standard G.711 Codec on the PSTN call through such a narrow modem pipeline was not going to work.

As a result, a lot of work went into building compression software, via codecs such as G.723 or G.729 that shrank the 64kbs voice stream down to 6kbps to 8kbps. The result was that you could talk, but it certainly did not sound good. The audio quality did not approach "toll quality," as is said in the PSTN world. But the technology enthusiasts did not care; it was a free call. And as they say, "free is good," especially if the caller is in the US and the called person in the UK.

Then the next phase kicked in--the techies wanted to call anyone who had a phone attached to the PSTN. Media Gateways were developed to interface the VoIP calls on the Internet to the PSTN. These gateways change the VoIP encoded call to PSTN time division multiplexed (TDM) 64kbs formats and back. But the conversion process added all sorts of new and odd sounding artifacts that made the poor quality call even worse.

So at this juncture we had heavily compressed voice running on the Internet (which at that time had its own set of problems, i.e. packet delays and packet loss) and then going though media gateways that further damaged the audio quality. Thus, VoIP got its start but did so under the curse of a bad first impression; the general public believed it to be an inferior technology.

## *21st Century VOIP Developments – Reduced Delay*

The 21st Century Internet within the continental United States has coast-to-coast delays of 100 msecs or less for IP packets. A nice site to visit to find out dynamic Internet timings is www.internettrafficreport.com, which shows round trip times from the Boston area to major Internet backbone points all over the world. The ITU-T recommendation is that one-way delays of 150 msecs or less are acceptable, and that one way delays above 400 msecs are not acceptable. Putting all this together, now in the continental US, we have acceptable delays.

The 1980's Internet ran over X.25 copper wires and dropped packets were a big problem. When VOIP packets are dropped on the Internet, they are not resent. The resulting speech is choppy.

Thanks to modern fiber transport (as opposed to copper in early 90's Internet), packet losses are typically less than 1%. The modern media gateways that bridge the internet to the PSTN can handle the issue of echo cancellation correctly.

With the advent of high speed internet in many homes, we now have a general public that on average has a bi-directional bandwidth connection to the internet of 384kbs or better. These are all the ingredients necessary for a quality VoIP call. The case for VOIP seems to be getting better, doesn't it?

If we just adopt existing standards, we still have problems. For example, a standard G.711 encoded VOIP call without echo cancellation and without silence suppression turned on takes about 155kbps of bandwidth. With proper echo cancellation and with silence suppression turned on, the same call takes only 80kbps. This is clearly doable for those homes with DSL/Cable Modem hookups whose upstream rate is 384kbps. In fact, with this much upstream rate, several properly configured VOIP calls can easily be handled and there will still be enough data space for a comfortable internet for the rest of the home users.

But even with these favorable factors in place most VoIP carriers still offer voice transport products that fail to live up to expectations. The reason? They have not dealt with the basics: end-to-end delays, echo cancellation, and quality of service at the customer premises. As a result the VoIP community is largely to blame for helping advance the myth that VoIP cannot sound as good as a toll quality PSTN call. This is simply not the case.

Let's take a quick look at three design parameters, and delve into why they are necessary to realize a toll grade VoIP call. On today's national Internet, VoIP to VoIP or VoIP to PSTN calls can sound as good as or better than a toll quality PSTN call. Here's why…

## *End to End Connections*

A long time ago, someone inside the now extinct Bell Labs came to the conclusion that end-to-end delays for a voice call would become objectionable to the human psyche if it exceeded 120 msecs. In the old days, 100-800 msec delays on coast-to-coast IP routes were common. But today in the US almost any location-to-location call can be achieved at 100 msecs or less, as the Internet has evolved within the last few years to provide genuinely acceptable delay. In fact, in many cases 30-50 msec delays are attainable, which is well within the tolerance limits of the human ear.

Even though delay in US calls is acceptable, calls outside the borders of the continental US and all bets are off. Many US based call centers have outsourced to facilities on foreign shores, such as India. These facilities then make VoIP calls back to the US, and do so with delays as high as 400-500 msecs. With this much delay, we do not have toll quality as we hear various forms of echo in these calls.

So VoIP works quite well on a national/local loop, if kept within a national IP network with low delays, such as the USA. But extended beyond its practical limit and the model starts to become questionable.

## *Echo Cancellation and VoIP (What did you say... clip, clip)*

Echo cancellation has to be the least understood design parameter associated with VoIP transport today. Let's explore the reason for this, and show how BandTel has resolved the issue.

First, why do we need echo cancellation on a voice call?

Basically, the caller using a two wire POTS telephone will insert a portion of the voice energy back to the person on the other end; i.e. phones reflect some of the audio coming to them back out to the network. This has to do with the imperfections of the phone's 2 to 4 wire converter (hybrid), and from the audio simply coupling from the speaker back to the mike. This reflected audio must be cancelled. If not, the end result is the bizarre experience of multitudes of echoes going back and forth between the two callers.

All of us have experienced echoes at one time or another, and especially in transcontinental calls, such as US to Sweden. Often this happens because the echo canceller for your call does not get turned on. So in summary, for any long distance call (20 msecs of delay or more), VoIP or PSTN, echo cancellation is mandatory. Without the cancellation, the call is an unworkable situation for the callers.

Bearing these echo cancellation issues in mind, it is possible today for users of the VoIP Session Initiation Protocol (SIP) to strike termination agreements with SIP based class-4 carriers, and have those SIP carriers terminate SIP originated telephone calls to the PSTN, or vice versa.

But a problem exists. Many SIP terminating carriers are using ISDN technology on the PSTN side of their media gateways. Consequently, they cannot handle the echo cancellation issues correctly; the caller using this type of network often hears audio clipping and long audio delays during the conversation.

The ISDN enabled media gateways deployed by these class-4 VoIP carriers have no way of controlling the echo cancellation hardware deployed inside the PSTN network. Therefore, their media gateways insert their own echo cancellation into the voice stream, but they do so indiscriminately, i.e. in the middle of the path of the call, as opposed to near the called and calling parties. The result is an unsatisfactory audio experience for the called and calling party; clips and voice delays are the norm. Correct echo cancellation functions must be applied as closely as possible to each caller's origination point; the VoIP caller's IAD (Integrated Access Device) and the PSTN caller's class-5 central office.

The only way to ensure this is to deploy the SS7-ISUP protocol, instead of ISDN, at the VoIP/PSTN interface in the VoIP media gateway. With that done one can control the echo cancellation equipment deployed in the PSTN, close to the caller. SS7-ISUP enabled media gateways simply send out the correct command, via SS7-ISUP, and the PSTN turns on its echo cancellation hardware close to the called party. The result is a telephone call with an audio quality as good as, or better than, regular PSTN to PSTN calls.

BandTel has resolved the echo cancellation issue by using SS7-ISUP gateways in key locations across it network. This gives BandTel a call quality that is superior to other termination products on the market.

### *Addressing the Last Mile*

Earlier we mentioned how the Internet bandwidth connections for the average user have greatly increased in the last several years. Now most users have access to at least 384kbs via cable or DSL connections. But bulk bandwidth alone is still no guarantee that a VoIP stream will have the Quality of Service (QoS) it needs for the call.

And before going any further, let's define QoS. QoS can mean many things, such as attempts to reserve bandwidth or VOIP calls, using queuing theory inside routers to boost VOIP packet flow onto the Internet, and even using routers to enforce fragmentation so that large packets do not dominate VOIP packets.

Even though the IP protocol used on the Internet initially provided some QoS options, these capabilities were largely ignored, especially by routers. So a general concept is that IP does not offer reserved bandwidth for any specific application, such as VoIP. Therefore a VoIP data stream gets thrown into the fray with everything else (email, http browsing, ftp file transfers, etc.. The result is that VOIP packets are treated as equal entities with email, ftp, http, and other packets. This is not a problem for the other services as they are not real time. VoIP however, needs to travel in real time.

Many of the consumer VoIP dial-tone services that surfaced over the last few years had no solution to this and simply dumped the VoIP stream onto the end user's LAN. As a result, the VoIP call sounds fine as long as nothing else is being sent over that same broadband connection. But when the VOIP user has limited bandwidth and starts using a web browser to surf the net, or a large email arrives, or some other application starts putting lots of packets onto the LAN, the VOIP call can degrade so that the call becomes choppy and quality is just not good enough to be acceptable.

There are solutions to this QoS on the horizon, with MPLS being the most likely winner long term. But ubiquitous deployment of any solution is still some time away. Fortunately, this problem is localized. On the US Internet, it is only an issue at the last mile over the end user's broadband connection. If QoS is solved on that segment of the network, the entire QoS issue is virtually solved.

BandTel has addressed the last mile, and offers a QoS solution that allows an enterprise to merge VoIP data streams with non real time bulk data (email, ftp, http, etc.), yet at the same time allows BandTel customers to hear no ill effects in the VoIP quality. This is accomplished by using the BandTel VoIP QoS certified user-agent device and QoS switch.

These two devices in tandem offer full QoS for VoIP, and guarantee a quality call, regardless of what is being sent over the customer's broadband connection.

## *BandTel's Solutions for Better VOIP*

Today's SIP based VoIP networks have grown way beyond the point of curiosity. But as SIP based VOIP networks evolve, hurdles will present themselves. With VoIP, a novel approach is now needed to overcome the issues relating to SIP throughput and redundancy—two areas that have become the latest hurdle.

## *The major problems that need to be overcome for SIP solutions are:*

1. *Bandwidth Heavy*- SIP utilizes an abnormally high rate of bandwidth for signaling. Looking at a normal SIP packet, we have 20 bytes of IP overhead, 8 bytes of UDP overhead, and 12 bytes of RTP overhead. With SIP traveling over the Local Area Ethernet, we add 14 bytes of Ethernet header overhead. Thus we have 52 bytes of overhead for each 20 bytes of voice when the codec is using 20 msec samples. Thus more than 70% of each SIP packet is overhead. That is a lot of data to process for each call!
2. *Character Clumsy*- The SIP protocol is all text, which must be parsed using relatively compute intensive text manipulation software.

As a result of these issues, today's fastest SIP servers are being challenged to match switching speeds of the purpose-built SS7-ISUP based TDM hardware.

To make a SIP call, a SIP end point, such as a VOIP phone, sends a SIP INVITE to its proxy. The general SIP solution has been to give end point SIP devices a hard coded IP address to the SIP carrier's proxy. With enough calls, even a fast SIP proxy will run out of CPU cycles, and can no longer handle any further load.

To handle this overflow, the SIP endpoint device must somehow understand how loaded its SIP proxy is. A typical VOIP phone is simply programmed to send its INVITE request to its proxy and typically has no provision to determine how loaded that proxy is.

SIP end points use the connectionless UDP protocol to send their INVITE requests to their proxy. When an end point sends its INVITE message to an overloaded SIP proxy unable to process any more messages, the INVITE message does not get processed and the SIP user cannot get calls through. Later, the proxy may have plenty of CPU cycles, so the same SIP user may get its call through with no problems. Such random phone call acceptance is acceptable for modern businesses. When the SIP user sends a SIP INVITE to an overloaded SIP proxy which does not respond, this creates an unacceptable single point of failure for the end user.

BandTel has looked at the situation, and found a way to resolve the dilemma by using a clustered architecture that uses a fixed front end so that all the SIP end points send their INVITES to this device. Behind this device is a specially setup DNS server which round-robins the requests to N different SIP proxies.

## *The Matrix and the SIP Signaling Transfer Point*

At the core of BandTel's network there are pairs of DNS servers that direct the SIP end point user agents (UA). The BandTel DNS servers resolve to (what appears to be) a BandTel proxy address for the UA. But in reality, the UA is pointed to something brand new in SIP space, which BandTel calls an STP (signaling transfer point).

This mnemonic/name was adopted from SS7 space because BandTel engineering realized it did something similar to its SS7 counterpart. Specifically, BandTel STPs (deployed in groups for redundancy) actually decide which one of N proxies in the BandTel proxy matrix will be used to process a given SIP call for a given SIP endpoint.

The STP allows BandTel to deploy a matrix of "N" SIP proxies to handle whatever load is required for any given task at hand. Thus, there is almost no limit to the CPU processing power within the BandTel signaling domain. Furthermore, there is no need for the UA's to hard code a carrier's proxy addresses into their configurations. So, the UA no longer needs to be concerned if a proxy has the capacity to handle its load.

Using this N-Plus approach also eliminates any single point of failure, and realizes redundancy to the Nth degree. BandTel's N-Plus Architecture solves some major technical hurdles just now surfacing in the SIP market place.

## *The Summation*

As mentioned prior, BandTel has deployed a matrix of proxies with each call being able to use any of the proxies in that matrix. This presents another issue; each proxy must now know registration and routing information for all the SIP endpoints on the BandTel SIP network.

To solve this problem, BandTel has created the "Synchor" and "Registrar" services. These two services run redundantly in real time, keeping the registration and routing information current and equal across all BandTel proxies. BandTel has been the industry pioneer in dealing with high capacity telephony signaling on an IP network. The ability to be able to synchronize all proxies in real time creates a single virtual switching machine with unlimited call processing potential.

Finally, BandTel ties this unique N-Plus Architecture back to the PSTN with a wide spectrum of tier one TDM national and international carriers for TDM network access. This in hand gives advanced routing and termination and origination options for all BandTel customer to or from anywhere in the world.

Stepping back and looking at the whole picture, BandTel is able to fill a critical need for high volume SIP customers by offering:

- High capacity throughput
- PSTN or better redundancy & reliability
- Multiple routes to any global destination
- Least cost route to any given destination
- Ease of connection to the PSTN

Acting as a high-capacity homologated SIP switching point for SIP endpoints, and connecting those endpoints to the world, BandTel is now the "virtual IP central office to the world" for the SIP end user.

## *In Summary*

It takes three main ingredients to get a quality voice call over IP:

-Low point to point IP network delays (< 120 msecs)

-Proper echo cancellation techniques deployed at the caller and called party locales

-QoS on that last mile

BandTel worked hard to make sure that all three of these design parameters have been addressed for its customers. BandTel is able to challenge the PSTN incumbents with VoIP quality that is second to none.

*Using these techniques, BandTel has connected 200+ million flawless VoIP calls.*