

**Enterprise Session Border
Controllers – Security and More**

June 2010

A brief history of security in telephony networks

Once upon a time, an Enterprises' interface to the outside world was exclusively voice, and effective border security was provided by receptionists who forwarded important calls, took messages for other appropriate callers, and politely rejected the rest. While certainly low tech, a receptionist was (and still is) an effective mechanism for application aware dynamic call admission policy management. Or in less technical terms, they did a pretty good job of screening calls.

From a technology perspective, the voice calls were carried over a dedicated and inherently specialized network – the Public Switched Telephone Network (PSTN). In the late 1960s “Phone Phreaks” with nom-de-guerres like Captain Crunch, Evan Doorbell, and Ben Decibel developed ways to use whistles and tones to wrest signaling control from the phone system, and so became the world’s first network hackers.

Relatively benign, they did little to disrupt most subscribers’ use of the network. Their main goal was to tweak the nose of phone companies, who moved to develop technology to prevent what they feared could become major revenue leakages. Out-of-band signaling (PRI and later SS7) was the networks answer to this challenge, and while it was largely successful, it eventually entailed replacing the entire network switching infrastructure – an effort that took nearly a quarter of a century and untold billions of dollars.

And then came the internet. Emails, web surfing, file transfers, and all sorts of other signaling protocols were all sent over a decidedly unspecialized network, and unfortunately one that was conceived with little thought to security. In hindsight, the result was predictable, and malware, virus’, SPAM and other network dangers continue to plague the network in ways that make the exploits of the phone phreaks seem positively quaint.

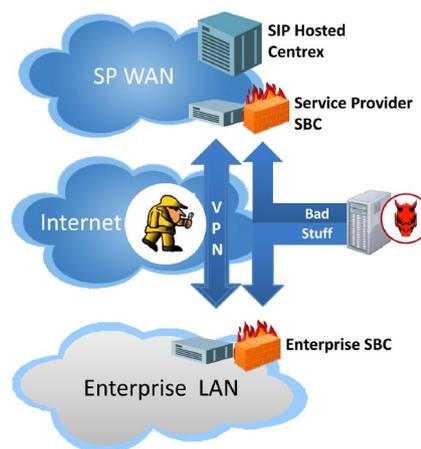
At least so far, there hasn’t been the kind of silver bullet solution for internet security that out-of-band signaling provided for the PSTN. Fortunately, thanks to firewalls with Application Level Gateways (ALGs) protecting their networks, and anti-virus software protecting their users, today’s Enterprises are reasonably well equipped to keep most network nastiness out of their corporate networks, providing a sort of security stalemate. Email, Web, and File access traffic is routed to appropriate servers when legitimate, and turned away when it’s not. Data sessions from users remote to the Enterprise’s core network are authenticated to verify the users’ identity and encrypted to protect the content of any email, IM, or files that leave the Enterprise’s secure environment.

Enter Voice over the Internet Protocol (VoIP), which does away with the (now) virtually impervious PSTN. With VoIP, voice is just another data service, and it’s a data service that commercial firewalls and ALGs don’t handle very well. Traditional firewalls are typically designed to permit data sessions only when they’ve been initiated from behind the firewall. Incoming sessions are blocked unless the SIP (Session Initiation Protocol - signaling) and RTP (Real Time Protocol – media) ports are unblocked completely, which would an unacceptable security exposure. But especially for an Enterprise, blocking incoming calls isn’t a solution either.

The Enterprise Session Border Controller

In order to resolve this issue, the firewall has to know more about the call than the Layer 3 information it has access to (eg the port it arrived on or the IP address it came from). It needs some of the information from the message itself, not just the Layer 3 header for the message. An Application Level Gateway function is required for VoIP, just like one is required for email, web, and file servers. The realization of this VoIP Application Level Gateway function is called a Session Border Controller (SBC). While typically characterized as a standalone device, it can be integrated with other existing network elements, like data routers or media gateways. However deployed, it has become an essential element in the network of any Enterprise utilizing the cost and functionality advantages that VoIP provides.

When an Enterprise connects to an Internet Service Provider (ITSP) network, the ITSP will route all calls from the ITSP network through an SBC on its way to the Enterprise. The ITSP network needs protection just like the Enterprise does, and the ITSP’s SBC provides it. But there is a crucial distinction here that is sometimes overlooked. The ITSP SBC protects the ITSP network, NOT the Enterprise. The Service Provider SBC can limit traffic out of the ITSP, but not into the Enterprise, in other words, the border role just can’t be hosted.



More importantly, even if the ITSP were willing to provision the SBC with per Enterprise screening policies (and ITSPs are typically not willing to endure that kind of churn on such a network critical element), it's not the only entity with access to the Enterprise's network connection. The connection from the Enterprise to the ITSP passes through an internet provider, supplying not just the voice, but all the data traffic for the organization as well. So even if there was no need to secure the Enterprise network from ITSP traffic, security measures are required to protect from dangers originating in the rest of the public internet.

VOIP Security

So just what exactly does the Enterprise network need to be secured against? First and foremost is that call screening function that receptionists used to supply, but scaled up to handle every call into the Enterprise. Call Admission Control (CAC) allows the specification of which calls are permitted, and which are not. This decision can be based on a variety of criteria, for example:

- Allow/deny traffic of specific types or from specific sources. The restriction might be based on IP or SIP address, network, groups of users, protocol, or other categorization
- Priority can be given to specific call types like E911
- Traffic can be limited by number of calls or packet rate to ensure that even permitted sources cannot deliver an overload delivered at wirespeed (malicious or not) into the Enterprise, and ensures users aren't "Denied Service" as a result

However, while call admission and Denial of Service attacks may be the security protections that come to mind first, they are by no means the only security related task that is required.

Topology Hiding ensures that the Enterprise's compute infrastructure is not exposed outside the LAN, reducing the opportunity that servers might be identified for attack. Also internal or personal information like names or email addresses can be eliminated or modified to ensure privacy and confidentiality.

Authentication provides assurance that traffic sources are actually who they purport to be, and not a malicious agent "spoofing" a legitimate source. Encryption can keep both signaling and media content secure not only across public internet segments, but also within the Enterprise as well. Authentication and encryption are also needed in controlling access to the SBC itself – the protector must be protected as well.

Interoperability

As important as security is, there has to be interoperability first. SIP is a complex protocol, and its specification is segmented into many different standards, typically Internet Engineering Task Force (IETF) Request for Comments (RFC) documents. Not every SIP Proxy/PBX/Softswitch has implemented exactly the same subset of specifications, or even a given specification in the same way.

For example, some systems use TCP (Transmission Control Protocol) for SIP transport, others use UDP (User Datagram Protocol), and the two are not compatible. On the media plane, G.711 might be used for clarity on the Enterprise LAN, but G.729 used on the limited bandwidth T1 connection to the Service Provider.

Either way, as intermediary devices, SBCs are well positioned to provide the signaling and media adaptations required to overcome any inter-operability issues that might exist. This makes choosing an IP-PBX or an ITSP much less restrictive – selection can be made based on merit and not compatibility.

Handling Users In Remote Locations

Enterprise operation is increasingly distributed, enabling many benefits and efficiencies. Employees work from home; call center agents can be located across the country; and Enterprise IP-PBXs can manage call traffic for remote locations as well. However, this flexibility presents some challenges for VoIP systems related to the firewalls behind which most of these remote workers are located.

The firewall behavior that is problematic is caused by the Network Address Translation (NAT) function that allows the private network behind the firewall to interwork with the public network the firewall faces. To protect the private network, the NAT will allow traffic from a particular IP source only if a session to that address has first been established from within the LAN, creating what is called a "pinhole" through the firewall's protection. This isn't a problem for outgoing calls (which are initiated from the LAN), but terminating calls are blocked. Fortunately, the first thing a phone must do is register with the PBX. This outgoing session establishes the "pinhole", needed for both outgoing and incoming traffic. Problem solved? Not quite, the pinhole expires after just a few minutes.

To keep this pinhole open, it must be "refreshed" on a periodic basis. The Enterprise SBC, through which all registrations pass, can assume this role for the PBX. The SBC needs to know about these registered users in any case to be able to associate them to User Groups upon which to base routing decisions (for example, different routing rules for users in two different branch offices).

There is also a media plane hurdle. The remote SIP user sitting behind the NAT will use its local IP address and port when exchanging media information via the SDP (Session Description Protocol), or put another way, it won't be the IP address and port on the firewall to which the media packets must be sent. However, the remote user has the public address to which to send media, and when those packets are received (from an IP address that matches the SIP exchange), the source address in the IP header yields the port on the firewall to which the media packets must be sent.

Legacy PSTN Connectivity and Wan Isolation

There are a number of reasons why an Enterprise might want to continue taking advantage of TDM technology in part of their network, including legacy TDM PBXs, local end office connectivity, and especially E911 service. Then, unless the Enterprise SBC can deal with both the PSTN and VoIP connections, a dual voice solution is required. However, integration not only offers the savings realized from fewer network elements, but integration between the two domains becomes possible. For example, if the IP link to the WAN fails, calls might be re-routed over a PSTN link preventing complete isolation and ensuring availability of critical services like E911. Finally, a universal border device allows any short term TDM investment to be re-vectored to VoIP as traffic loads shift.

In a Hosted IP Centrex configuration, loss of the WAN also cuts off call control – disabling even internal calling. A remedy for this is to route all user registrations through a local fallback device which could assume control when access to the Hosted PBX is lost. The Enterprise SBC is perfectly positioned to perform this stand alone survivability role. It is already in the signaling path for all the user registrations, and by its very nature has the capacity to perform basic call routing. In addition, it has a user registration facility for remote users behind firewalls.

Combined together, PSTN access and stand alone survivability allow the Enterprise SBC to provide critical service assurance.

Voice Quality Assurance

VoIP is subject to a host of impairments that the PSTN never faced. Packet networks introduce jitter, delay and packet loss, especially in converged networks with data traffic co-existing with voice. To deal with this, impairment mitigation strategies need to be employed, and quality measurements must be taken to ensure expected quality is being achieved.

Dynamic jitter buffers adapt to network conditions, expanding when extra delays are present to reduce packet loss, and shrinking when they're not to reduce latency. Packet Loss Concealment (PLC) minimizes the impact of lost packets when that does occur. And high quality voice processing ensures that compression codecs sound the very best they can.

Metrics such as packetloss, jitter, and Mean Opinion Score (MOS) measured per call and aggregated by route and time-of-day allow quick identification of any issues that do occur. And finally tools to discover the cause of any quality deviations are important if the issues are to be resolved quickly.

Summary

Enterprise Session Border Controllers are essential components of any corporate migration to VoIP voice services. They protect the network from harm, secure the assets of the organization, and facilitate interoperability with Service Providers and remote users. PSTN integration adds a holistic infrastructure for voice and provides for critical service assurance. Whatever the network topology, hosted centrex, call centers or SIP trunking, the E-SBC brings critical support to enterprise of bringing voice to the Enterprise.

ABOUT AUDIOCODES

AudioCodes Ltd. (NasdaqGS: AUDC) provides innovative, reliable and cost-effective Voice over IP (VoIP) technology, Voice Network Products, and Value Added Applications to Service Providers, Enterprises, OEMs, Network Equipment Providers and System Integrators worldwide. AudioCodes provides a diverse range of flexible, comprehensive media gateway, and media processing enabling technologies based on VolPerfect™ – AudioCodes' underlying, best-of-breed, core media architecture. The company is a market leader in VoIP equipment, focused on VoIP Media Gateway, Media Server, Session Border Controllers (SBC), Security Gateways and Value Added Application network products. AudioCodes has deployed tens of millions of media gateway and media server channels globally over the past ten years and is a key player in the emerging best-of-breed, IMS based, VoIP market. The Company is a VoIP technology leader focused on quality and interoperability, with a proven track record in product and network interoperability with industry leaders in the Service Provider and Enterprise space. AudioCodes Voice Network Products feature media gateway and media server platforms for packet-based applications in the converged, wireline, wireless, broadband access, cable, enhanced voice services, video, and Enterprise IP Telephony markets. AudioCodes' headquarters and R&D are located in Israel with an additional R&D facility in the U.S. Other AudioCodes' offices are located in Europe, India, the Far East, and Latin America.

International Headquarters

1 Hayarden Street,
Airport City
Lod 70151, Israel
Tel: +972-3-976-4000
Fax: +972-3-976-4040

AudioCodes Inc.

27 World's Fair Drive,
Somerset, NJ 08873
Tel:+1-732-469-0880
Fax:+1-732-496-2298

Contact us: www.audiocodes.com/info

Website: www.audiocodes.com/hdvoip

©2009 AudioCodes Ltd. All rights reserved. AudioCodes, AC, AudioCoded, Ardito, CTI2, CTI?, CTI Squared, HD VoIP, InTouch, IPmedia, Mediant, MediaPack, NetCoder, Netrake, Nuera, Open Solutions Network, OSN, Stretto, TrunkPack, VoicePacketizer, VolPerfect, VolPerfectHD, What's Inside Matters, Your Gateway To VoIP and 3GX are trademarks or registered trademarks of AudioCodes Limited. All other products or trademarks are property of their respective owners.

Ref # LTRM-80030 06/10 V.1