SIP Trunking Benefits and Best Practices

White Paper

Janne Magnusson
Vice President, Product Management
Ingate® Systems

Abstract

Back in the days of wireline telephony, when all phone calls went over the PSTN, businesses would purchase "trunks" – a dedicated line or a bundle of circuits – from their service provider. Today, we have adapted the concept of "trunking" to the IP-enabled landscape resulting in lower telephony costs and rapid return on investment (ROI) plus the opportunity for enhanced communications both within the enterprise and with vendors, customers and partners.

A SIP trunk is a service offered by an ITSP to use SIP to set up communications between an enterprise PBX and the ITSP. A trunk includes multiple voice sessions – as many as the enterprise needs. While some see SIP as just voice, SIP trunking can also serve as the starting point for the entire breadth of realtime communications possible with the protocol, including Instant Messaging, presence applications, whiteboarding and application sharing.

The potential for a rapid return on investment is a key driver of SIP trunk deployments. However, maximum return on investment can be achieved when you extend VoIP outside of the corporate LAN. In terms of infrastructure purchases, SIP trunks provide an immediate cost-savings. They eliminate the need to purchase costly BRIs, PRIs or PSTN gateways.

The productivity benefits with SIP and SIP trunking are also significant. By extending the SIP capabilities of the corporate network outside the LAN, satellite offices, remote workers and even customers can use VoIP and other forms of realtime communications applications to break down barriers of geography to share ideas and increase productivity.

There are three components necessary to successfully deploy SIP trunks: a PBX with a SIP-enabled trunk side, an enterprise edge device understanding SIP and an Internet telephony or SIP trunking service provider.

Equipment based on the SIP protocol – SIP phones, IP-PBXs etc. – have been around for some time. Now that SIP trunks have gained momentum, it has become important to ensure that equipment works together. It is for this reason that standards such as SIPconnect™ have become so critical. SIPconnect was developed by the SIP Forum as a set of best practices for interfacing an enterprise PBX implementation with an ITSP that attempts to eliminate some of the unknowns and incompatibilities of mixing equipment from different vendors in a single environment.

Like any application that opens the network to the Internet, SIP trunking deployments have security considerations, but there are ways to maximize enterprise security. One of the most effective techniques is to address SIP security the same way data security is addressed - at the enterprise edge. SIP server and SIP proxy technologies offer maximum control over the flow of SIP traffic, enabling the administrator to ensure correct routing, apply verification and authentication policies and mitigate Denial-of-Service attacks.

Voice quality is not an issue with SIP trunking if proper Quality of Service (QoS) measures are applied, such as over provisioning of links, and prioritization of voice traffic. Reliability is also a moot point. In fact, SIP trunks can be more reliable than the traditional PSTN as a number of failover solutions can be implemented.

1 What is SIP trunking

Unlike in traditional telephony, where bundles of physical wires were once delivered from the service provider to a business, a SIP trunk allows a company to replace these traditional fixed Public Switched Telephony Network (PSTN) lines with PSTN connectivity via a SIP trunking service

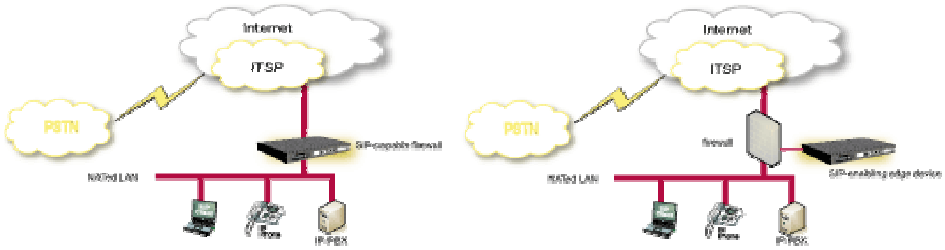

provider on the Internet.
Figure 1. Two typical SIP trunking solutions

In Figure 1, the PBX is located on the internal network. The PBX must have a SIP-enabling trunking interface. It can either be an IP-based PBX communicating to all endpoints over IP, or it may just as well be a traditional TDM PBX. The sole requirement is that an interface for SIP trunking connectivity is available.

Over the Internet, the ITSP (Internet Telephony Service Provider) provides connectivity to the PSTN for communication with mobile and fixed phones.

The PBX on the LAN connects to the ITSP via the enterprise border element. The border element may be a SIP-capable firewall or a SIP-enabling edge device, attached to an existing non-SIP-capable enterprise firewall (all these components are described in depth in section 3).

2 The benefits of SIP trunking

Many enterprises are already using VoIP, however, many are only using it for communication on the enterprise LAN. In this scenario VoIP is only being used as a one-to-one replacement for traditional wireline telephony. For all calls made to the outside of the LAN a PSTN gateway on the enterprise edge is used. These businesses realize a solid return on investment (ROI) just by lowering administrative costs and the costs associated with calls made within the company.

With SIP trunking, the potential for ROI is far greater because SIP trunking takes the idea of VoIP a step further, beyond this LAN application. The full potential for IP communications can be realized only when the communication is taken outside of the corporate LAN.

SIP trunking delivers several benefits:

•Eliminates costly BRIs (Basic Rate Interfaces) and PRIs (Primary Rate Interfaces) subscriptions
•No need to invest in PSTN gateways and additional line cards as you grow
•Edge devices offer low investment path in adding new lines as they are typically cheaper per line than the corresponding PSTN gateway
•Optimal utilization of bandwidth by delivering both data and voice in the same connection
•Maximum flexibility in dimensioning and usage of lines as you avoid having to buy capacity in chunks of 23 (T1) or 30 (E1) lines
•Flexible termination of calls to preferred providers; calls to anywhere worldwide can be made for the cost of a local one
•Redundancy with multiple service providers and links

One could argue that it is less expensive to purchase and administrate an IP-PBX than the traditional PBX, but for most companies this is not enough to motivate the investment.

Also, some telephony service providers offer free PRI subscriptions for enterprises meeting a minimum level of call minutes per month. These "deals" aren't necessarily the best investment, as SIP trunking service providers have a more efficient delivery method than traditional telephony service providers for the service and also in many cases an additional revenue stream in data traffic. These providers should be able to offer a more competitive price for the voice termination.

The cost effectiveness of a SIP trunk is such that by replacing an existing PSTN gateway/PRI installation with an edge device/SIP trunk, ROI may be achieved in a matter of months. For new installations a SIP capable edge device is most often a smaller investment than a PSTN gateway, making that investment cheaper.

2.1 Calculating the investment ROI

It is almost impossible to calculate a "standard" ROI for a SIP trunking investment, as there are far too many service providers that offer services with widely differing conditions.

This section focuses on the fundamental parameters affecting the costs and the principles of how enterprises using an IP-PBX, moving from traditional TDM PRI connections to SIP trunks, can achieve a rapid return on investment.

2.1.1 No more BRIs, PRIs or PSTN gateways

One of the immediate ways SIP trunks reduce communications costs is by eliminating the need to purchase ISDN, BRIs, PRIs or local PSTN gateways. Since the voice traffic is now routed through the Internet connection to the ITSP, no local connection to the PSTN is necessary at the enterprise location. The gateways needed to connect to the PSTN will reside in the ITSP's premises.

The devices required at the enterprise edge for SIP trunking are not only typically cheaper per line than the equivalent PSTN gateway, but they also enable the whole breadth of SIP-based realtime communication and therefore become a strategic device in the future of enterprise communication.

2.2 Bandwidth utilization

The utilization of bandwidth is often low with both telephony (TDM) and Internet lines. The telephony patterns in many organizations are distinguished by several hours a day with many calls, some with few and the rest in between. Internet data traffic, on the other hand, is for the most part erratic, with "bursts" of traffic happening throughout the day.

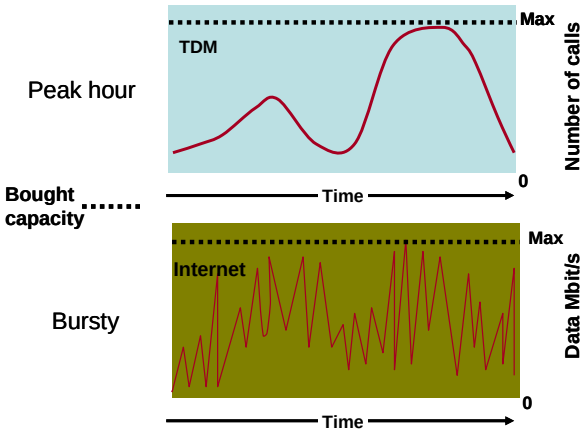

Figure 2. Typical bandwidth utilization

If we arrange the data with the time periods with the highest usage at the left and then in descending order, it becomes evident how much of the total capacity is wasted.

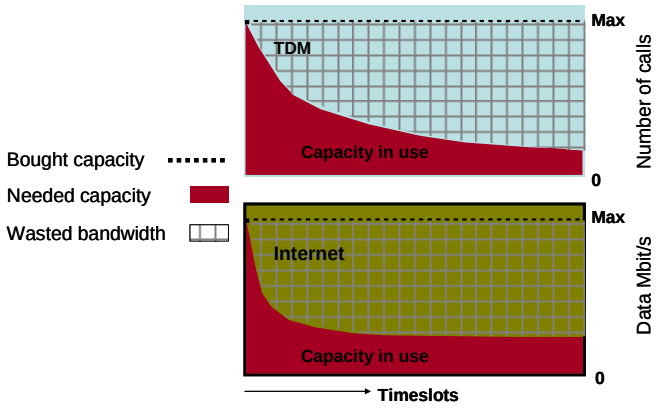

Figure 3. *Comparison of capacity in use.*

In practice, when compared to realtime communications (such as voice), data traffic is usually not as time critical. To combine the two communication forms on the same connection will give maximum use of capacity. By applying the correct Quality of Service (QoS) settings, critical voice communication can be prioritized over the data communication at all times.

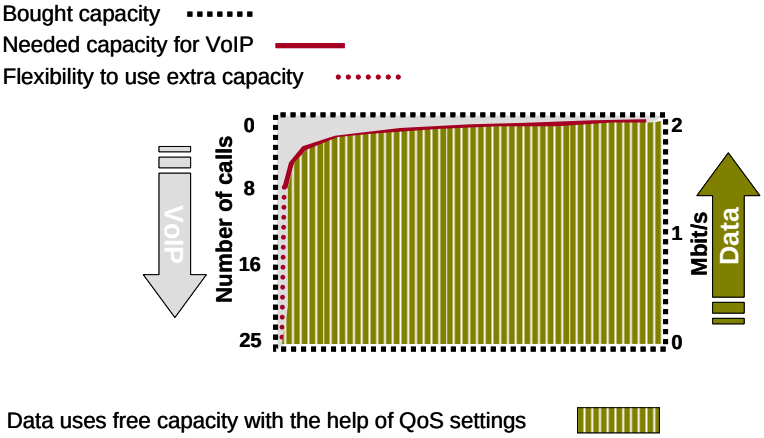

Figure 4. The capacity you need, when you need it.

With a SIP trunking solution, the capacity you need when you need it is always available. Instead of dimensioning the telephony for peak usage, it may instead be dimensioned for average usage, allowing the dynamics of QoS to make sure that voice traffic always gets the capacity it needs.

2.3 Flexibility to add new lines

Adding lines with a SIP trunk connection is quite linear.

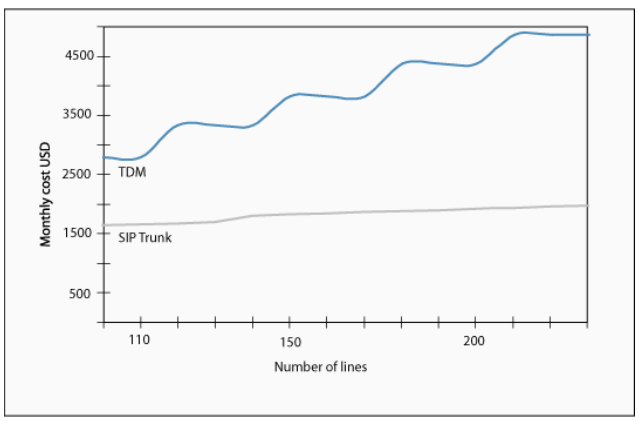

Figure 5. Costs to expand infrastructure, TDM vs. SIP trunks. All figures are based on a European operator's official price list Sept. 2006.

The prices for the communication above are based on a real case in which the enterprise has an IP-PBX installed. The costs above are from a leading European operator. This operator offers both traditional TDM PRI/BRI connections and SIP trunks. For the TDM solution the calculation is based on a traditional PSTN gateway from Cisco with PRI connections. For the SIP trunking solution an Ingate SIParator® 45 is serving as the edge device.

All investments in hardware or software are written off over 36 months.

When an enterprise using a TDM solution needs to increase its capacity, the following for each chunk of 23 (U.S.) or 30 (Europe) lines must be added:

•New PRI subscription.
•New line card for PRI in the PSTN gateway.

When the capacity of the PSTN gateway and/or PRI connection is reached, it is necessary to invest in an additional PSTN gateway and/or PRI subscription. Unfortunately this is true even if you only need one more line. Going from one E1/T1 to two always requires additional hardware and they can only be bought in steps of 23/30 lines. Even if you move from an E1/T1 to a higher level standard bundle like STM-1 the hardware will need to be replaced. The SIP edge device does not have that problem.

In a SIP trunk solution, the enterprise can increase one line at a time by:

•Purchasing additional software licenses for the edge device.
•Allocating a greater percentage of the bandwidth for voice. Only if the total bandwidth capacity is used will the Internet connection need to be upgraded.

2.4 Least Cost Routing (LCR)

The use of IP makes it possible to cost efficiently use SIP trunks from multiple service providers, depending on optimal availability and the best rates (capitalizing on time zone differences, geography etc.). In essence, the enterprise may become its own "Master Service Provider" with subscriptions to service providers in countries where they have the highest calling volumes. By routing calls to the cheapest service provider based on country codes, for example, significant savings can be achieved.

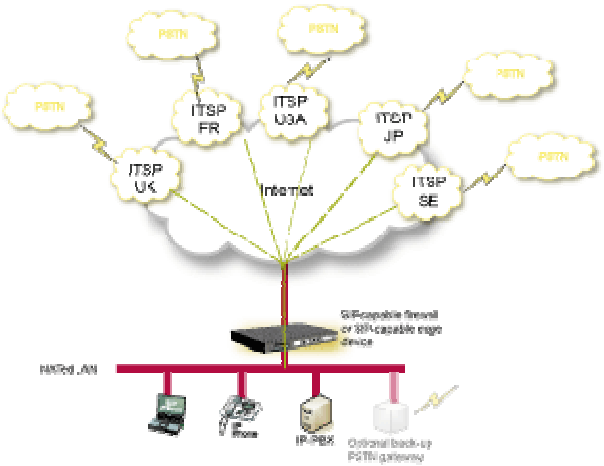

Figure 6. Utilizing multiple service providers.

These routing decisions can be made by the PBX or by the edge device. The fact that this ability can be built into the edge device means that low functionality PBXs can perform routing functionalities as well. By "outsourcing" this function to the edge device the PBX needs only to send the number as it is, and let the edge device act depending on destination etc.

Using multiple service providers, provides a higher level of security and reliability:

•Failover to secondary Internet service provider.

•Failover to secondary service provider or back-up PSTN gateway.

2.5 Making IP-to-IP calls when possible

Today, calls that could be transferred over IP all the way are connected through TDM connections instead. These situations arise when calls are routed to a PSTN gateway on the LAN. In essence the true benefits of IP communications are not only unrealized, they are defeated as quality will suffer by analog/digital transcoding several times over.

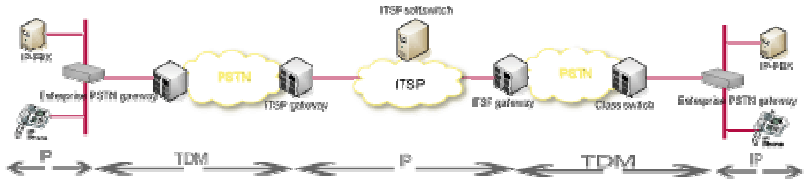

Figure 7. *Transcoding from IP to TDM.*

ENUM (Electronic Number Mapping System, also known as Telephone Number Mapping) is a standardized address translation technology adopted by the IETF (Internet Engineering Task Force) using DNS (Domain Name Service) to link a phone number to a specific SIP address. This feature is used to automatically look up phone numbers to determine if they match a known SIP address, allowing the call to be completed over the Internet (instead of transferring it to the PSTN). Since no traffic is placed on the PSTN, ENUM provides an additional means of cost savings for businesses that communicate with other enterprises also using SIP. If the number is not found in the ENUM database the edge device will route the call to the service provider for termination to the PSTN.

With the growing installed base of SIP-based IP-PBXs, the critical mass for widespread deployment of ENUM will soon be here. It will not be long before the majority of calls will be transferred directly via SIP over IP between the calling parties instead of going over the PSTN.

2.6 SIP trunking – the stepping stone to higher productivity

Even if it is far more difficult to calculate, it is perhaps with the gains achieved in improving productivity that SIP trunking delivers the fastest ROI of all. Introducing SIP-based realtime communication has a tremendous impact on how people work, collaborate and communicate now and in the future. SIP trunking is an important step in this direction as it is the feature that moves communication from the old PSTN connection to the Internet. Once that is done the field is open for adopting all of the productivity-enhancing features that SIP offers.

2.6.1 Rich communication

SIP has become the standard protocol for VoIP. However, it was originally designed to initiate all types of realtime communications over the Internet, not just voice.

These types of realtime communication include:

•Presence, to see who is currently online and available
•Instant Messaging (IM), text messaging in realtime
•File transfer
•Application sharing, collaboration on a single document
•Whiteboarding, writing and drawing on a common virtual whiteboard
•Video conferencing
•Machine-to-machine realtime communication
•The distribution of alarms

A wide palette of rich communications options enable users to exchange ideas in the best possible way for their immediate situation. For instance, for remote workers at a

An example

Productivity benefits using rich communication

Joe, a salesperson based in Los Angeles working for an international company, is pressed for time. He needs to deliver an updated sales quote to a customer by 5 pm PT. Joe needs his supervisor's approval to offer the deal, but she is based in London where it's nearly midnight.

Being the resourceful employee that he is, Joe checks his computer and sees that his supervisor is working late – she's logged in to the presence application and indicating that she is available. Joe sends her an IM, they start a voice session and collaborate on Joe's final quote together using application sharing. Joe's supervisor shows him the company's yet-to-be-announced line of products using a video session, so he can be thinking ahead when negotiating with the customer.

In this example, five different types of SIP communication are used to achieve the common objective of completing the deal. In the old days of wireline telephony this communication would not have occurred simply due to time zone restrictions. In the end, Joe would never have met his deadline, and the deal would have been lost.

WiFi-enabled hotspot, the best way to communicate with colleagues may be *via* IM, not VoIP.

2.6.2 The transparent business - road warriors and home users

One of the key benefits with rich communication applications is making businesses run more transparently. Business can be conducted from anywhere in the world – regardless of time zones, locations (remote workers, for example) – so that customers have maximum access to your staff. In addition, employees can access corporate resources from anywhere in order to save the company money, they can leverage expertise from colleagues in other offices or even other countries, or use SIP to provide customers with the best service.

This same technology for remote connectivity can be used for all clients including PC-based softphones and IP phones connected to the Internet. This is an advantage of the SIP protocol: to be able to register multiple devices with the same address i.e. phone number. A person can then, for example, use an IP phone/softphone at their home office and an IP-only phone in the corporate office, both registered to the same number. One number reaches the employee in multiple locations.

2.6.3 Dual-mode handsets supporting voice over cellular and WiFi

The demand for mobile phones equipped with both cellular and WiFi capabilities is very strong. The potential cost savings for a person who, for example, frequently travels overseas and is able to transfer expensive cellular calls into virtually-free VoIP calls when connected to the Internet is significant.

Third-party clients that decide whether to route calls over VoIP (if a strong WiFi signal exists) or over cellular networks are available. The next evolution of this capability will be to seamlessly roam between WiFi and cellular connections with no interruption in the call. This technology already exists and we will see a roll out of this very soon from different service providers.

3 SIP trunking infrastructure

This section will describe in detail the three components needed to set up a SIP trunking solution: a PBX, an edge device that can handle the traversal of SIP traffic and a SIP trunk from an ITSP.

3.1 The PBX component

This section will provide an overview of the different types of PBXs available on the market.

3.1.1 The traditional PBX

A PBX (Private Branch eXchange or Private Business eXchange) is a telephony exchange serving an enterprise or organization office. It performs the basic function of routing calls to their destination as well as provides a great number of value-added features: call transfer, music on hold, redirect when busy or no answer, etc. The traditional TDM PBX was connected to a dedicated premises network that only carried the voice traffic.

3.1.2 The line-side IP-enabled PBX

The LAN for data traffic is a much later addition to the office than the telephony network. When introduced it came in a separate parallel premises network. For many years these two network cable systems have coexisted in the office serving separate but related communications functions.

The first IP-based PBXs, or IP-PBXs, focused on making the line side of the PBX, i.e. the side connecting to the telephones, run on IP protocols. The first and very obvious gain in doing so was that the two premises networks now could be converged into one common network. By use of IP enabled telephones, these could be connected to the same physical cabling as the computers and servers, i.e. the LAN. Having made this change to a common premises infrastructure it also became possible to introduce PC-based soft clients instead of traditional telephone sets.

Some argue that voice and data traffic should not be mixed on the same LAN or at least should be run on separate virtual LANs (VLANs). The background to this

position is that voice traffic, due to its realtime nature, is sensitive to delays or lack of bandwidth in the infrastructure resulting in poor voice quality. However, this issue can be solved and should not stand in the way for realizing the benefits of converged communication as described in the previous section. The bandwidth available on most enterprise LANs, 100Mbit/s or 1Gbit/s, is more than enough for most typical enterprise applications. By using appropriate QoS techniques enterprises can easily ensure that the voice traffic gets the appropriate priority to safeguard voice quality.

3.1.3 The IP-PBX

IP telephony in the above sense, using IP-based telephones connected through the corporate LAN, has been around for quite a few years. However, whenever calls needed to flow outside the corporate LAN they had to be routed to a local PSTN gateway (or through a PSTN gateway function within the PBX) and converted to traditional TDM-based telephony. Due to the inherently proprietary nature of TDM equipment and the fact, as described in a section above, that growth in traffic inevitably leads to the need to install additional hardware, this solution is expensive.

In a world where more and more end points are running IP there is a risk of deteriorating sound quality due to repeated transcodings between IP and TDM as shown in figure 7 in section 2.5.

The next natural step, and the topic for this white paper, is to use IP for the interface to the world outside the corporate LAN. This is done by IP-enabling the trunk interface on the PBX, completing its transformation into an IP-PBX. In practice, this happens in one of two ways. For earlier TDM or IP-PBXs this can be achieved by placing an IP front-end on the trunk interface creating what is usually referred to as a hybrid IP-PBX. This PBX contains both legacy TDM and IP-enabled parts. Newer IP-PBXs, or systems that are designed from scratch, are usually built with IP technology from the ground up, without the legacy TDM part. For such systems any connection to the PSTN requires a separate PSTN gateway.

There are a number of protocols available that could be used to IP-enable the trunk interface, including MGCP, H.323 and SIP. However, SIP is now the protocol that has won the standards battle. SIP has a number of advantages over the other protocols, the most important of which is that it supports rich communication while H.323 is a voice-only protocol. The use of an IP-based trunk interface provides all the benefits described in the previous chapter and addresses the issues of sound quality and cost.

3.1.4 Benefits of IP-based PBXs over legacy systems

In the following section further advantages of IP-based PBXs, in addition to the benefits of SIP trunking, are highlighted.

3.1.5 Connect multi-vendor end points

There is a trend in the PBX market to allow equipment from different vendors to coexist within the same PBX system. This will allow the enterprise to preserve investments made in phone endpoints even if the central PBX equipment is replaced. This allows the user to select phones, media servers and switches from their preferred vendor. PBX vendors that choose to allow this believe that the customer will be more likely to swap to their system if they can keep their existing phones. Some vendors, however, continue to lock their customers in to their own end equipment by making various proprietary extensions to the system.

3.1.6 User management

One of the most obvious advantages of an IP-based PBX system is increased manageability:

•By using the existing data network the need for separate wirings for a telephony system is eliminated.
•The phone becomes a kind of computer that allows the administrator to easily make upgrades and force policies to each phone from a central management system.
•The ID and configuration of the phone will follow the phone, regardless of where it is connected to the network.
•Users may log in to the phone when they arrive at a new desk; user profiles and information will automatically be loaded into the phone allowing greater flexibility.

Developing a Working Solution End-to-End
by John Casselman, Shoretel Inc

Providing a working SIP trunking solution from end-to-end can be a difficult task. Industry-wide vendors are partnering with other technology manufacturers and service providers to address issues such as security, NAT traversal *etc.* This partnering is a tremendous benefit to end-users, who reap the rewards of behind-the-scenes interoperability testing. With SIP trunking, these partnerships translate into easier deployments with bundled solutions – a plug-and-play package that works end-to-end.

This is why we have forged relationships with industry leaders (such as Ingate Systems) to provide complete solutions that are not only fully interoperable, but address security issues as well -- a critical factor for many of our customers.

We've teamed up with Ingate for a solution to connect the ShoreTel system to Internet Telephony Service Providers (ITSP). A box needs to sit between the LAN and WAN (Ingate) to perform the following functions; otherwise complications can occur between the ShoreTel system and the ITSP:

- SIP NAT traversal
- SIP REFER to reINVITE messaging conversion plus other B2BUA tweaks to help with interoperability between ITSP and ShoreTel
- Dial Plan manipulation may be required
- Digest Authentication / Registering
- Routing - When one has multiple sites without the Ingate box, RTP can potentially go out another GW which may or may not support SIP. With the Ingate box, SIP messaging, as well as RTP, always goes out the one Ingate GW.
- Full SIP firewall, plus the option of enabling the box to handle data firewall security
- Optional module allows for "Far End" NAT traversal – provides ability for SIP messaging to go through a non-SIP aware firewall.

ShoreTel and Ingate are working hard to perform joint testing as well as a full suite of tests for each ITSP the solution connects to. Once the solution involves ShoreTel, Ingate and the ITSP is completed and an application note is created which documents how each piece is configured along with contacts for each company. It also provides a summary of the 100 plus tests cases that are performed which documents the pass, fail and N/A of each test case. The document also provides which versions of code were used and any known issues at the current time. ShoreTel will also generate other supporting documentation. ShoreTel and Ingate's goal is to provide the customer peace of mind that the solution has been tested and will be straightforward to deploy.

3.1.7Integration with other IP-based applications

The SIP IP-PBX serves as the primary registrar of SIP users and utilizes this information for routing purposes. But the fact that the PBX is now IP-based also means that it can be integrated with other communications applications running on servers on the LAN. One of the best examples of this is converged communications soft clients that can integrate voice capabilities with applications such as presence, instant messaging, file transfer, white boarding, etc. Through such integration the PBX becomes part of a greater converged communication system that enables the enterprise to benefit from productivity enhancing communications applications.

3.2 The enterprise edge component

The enterprise edge component can either be a firewall with complete support for SIP or an edge device connected to the firewall, handling the traversal of the SIP traffic.

3.2.1 Firewall/NAT traversal

When moving to VoIP, the telephones, as the PCs, are connected to the Internet. It is imperative to safeguard the system from attacks and other unwanted access. This is especially critical if the PCs and the phones are always connected to the Internet, for example via a broadband connection or a fixed line. A firewall protects the PC by rejecting attacks and illegal data packets, allowing only approved traffic. On a local area network where several PCs or other equipment are connected, it is common to have private IP addresses on the LAN and a single common public IP address to the Internet. This functionality is called NAT (Network Address Translation) and is usually integrated into the firewall.

Firewalls and NAT routers are designed for data traffic that is initiated from the inside of the private network. Because malicious attacks on the network frequently originate from outside of the private network, firewalls and NAT routers protect the enterprise by blocking this kind of traffic. The problem, however, is that SIP traffic is "misunderstood" by traditional enterprise firewalls and NAT routers as being unwanted traffic.

The biggest hurdle for IT managers looking to SIP-enable their network is preparing the system to handle the traversal of SIP traffic across the firewall. The majority of current firewalls and NAT-routers are still not designed to handle full person-to-person communication, which will not reach users on the LANs unless the enterprise firewall has specific SIP support. SIP traversal of firewalls and NATs is becoming a commodity in the sense that most vendors advertise support for the protocol. However, the basic SIP support offered by most of these vendors does not have the richness of features to fulfill the needs of a complex enterprise environment. It is critical that IT managers evaluate their current firewall solution to ensure there is proper SIP support when new firewalls and NAT routers are installed.

One problem is that the media streams (e.g. voice) are transferred over dynamically assigned UDP ports that are generally closed. The firewall must be able to dynamically open and close ports based on the transferred SIP signaling. Another problem is that the SIP clients inside the firewall cannot be reached by IP addresses since these most often are private and local to the LAN. Communication simply does not take place, unless there is specific SIP support in the firewall.

Several methods and equipment have been suggested to resolve the issue of reaching users on the LAN. One such method solves the problem where it occurs – within the firewall itself. Firewalls that have a SIP server, with SIP proxy, SIP registrar and possible B2BUA (Back to Back User Agent), that dynamically control the firewall have been available for many years. This solution provides optimal flexibility as SIP signaling can be rewritten and processed in a very flexible way ensuring correct routing and interoperability with other systems built to RFC 3261 and related standards.

Several firewall vendors develop models with SIP ALG (Application Layer Gateway). ALGs usually work at a lower level than a proxy, adjusting the data packets "on the fly." Cisco is developing firewalls with ALGs that also handle incoming calls to multiple users, while other more simple implementations may only support a single SIP user on the LAN. One limitation of the ALG

architecture is that it cannot handle secure SIP signaling via TLS (Transport Layer Security). This architecture also lacks the ability to rewrite SIP signaling in several ITSP scenarios.

3.2.2 Mediation between PBX and service provider equipment

Most basic call scenarios in a SIP trunking solution, using equipment from different vendors, work well. However, when more advanced features such as call transfer are used, problems occur when the standard is not strictly adhered to by all vendors. In addition, SIP is a flexible standard that leaves some room for adjustments. This means that, at times, two clients can have difficulties talking to each other even though none of them directly violate the standard.

To make the situation even more complex, some ITSPs and PBX vendors only implement parts of the standard. Or, they add vendor-specific extensions to the standard.

While performing traversal and security these SIP-capable edge devices can also mediate between the PBX and service provider, offering an important function. They can process the SIP signaling and media in a way that is understood and expected by both the ITSP and the PBX.

3.2.3 Reliability – survival features

Thanks to its architecture with a full SIP proxy and registrar, an edge device can perform basic call routing functions. These functions can be used in order to increase the reliability and overall uptime of the total VoIP communications system.

There are edge devices that have a built-in watchdog function that detects if the contact with the central server is lost. The central server in this case could either be the carrier equipment at the SIP trunking service provider or it could be an IP-PBX located at headquarters serving several branch offices. (Such an intra-company connection is sometimes also referred to as a SIP trunk, even though this is not the scope of this white paper). The detection process will work whether contact is lost due to the fact that the central server is down or because the connection between has failed.

In the case of failure the edge device will take over the task of basic call routing, and depending on where the failure took place, enable the reinstitution of service partly or fully. For example, if the central equipment fails the edge device can route calls to alternative PSTN connectivity providers or a local PSTN gateway. If the problem occurs because the last mile Internet connection goes down the edge device can at least make sure that local, intra-office, communication can still flow.

3.2.4 Security from the edge device

SIP-enabling edge devices can also add a layer of security to enterprise communications, specifically in securing SIP media. Most security administrators will have serious concerns connecting a PBX system directly to the public Internet without any SIP-aware firewall in front of it. Like any server on the LAN, it needs to be protected by a firewall. A PBX is not built to withstand or recover from denial-of-service attacks and, in most cases, does not have filtering capabilities available to reduce traffic (requiring processing power to only the appropriate traffic). The enterprise edge device can secure the SIP media as well as data traffic.

The edge device can also protect the network from eavesdropping. Solutions for encryption of media and signaling using IETF proposed standards are recommended. These solutions include TLS (similar to SSL used for https) for signaling and SRTP (Secure Real Time Protocol) for media. Both are recommended in the SIPconnect initiative.

3.2.5 Enabling remote workers

To be able to extend the PBX features to remote workers in various locations, it is necessary to address the NAT traversal issue with SIP at the remote client end as well. While many businesses are either replacing their existing firewalls with SIP-capable firewalls, or deploying SIP-enabling edge devices to solve this problem internally, the NATs at remote sites (wireless hotspots, hotels etc.) are usually not SIP-capable. The result is that rich communication is not possible at remote locations.

Who is listening to my calls?

by Per Cederqvist, Chief Architect, Ingate Systems

IP telephony may be a great way to save money, but isn't it easier to eavesdrop on IP telephony than regular telephony? That is a question that is often raised when SIP is discussed.

The fact is that, properly used, the SIP standards suite can make it a lot harder to eavesdrop on your conversations by using SRTP to encrypt the media, making it virtually impossible to listen in on the conversation. Compare this to an old-style PSTN conversation, where it is easy to attach a wiretap to your phone line.

I worked for the Swedish phone company for a short while back in 1990, repairing phone lines. Even then, we had a device that fit easily in your hand that could pick up phone conversations when touched to one wire of a call. We used it to avoid cutting the wires while somebody was talking. We never used it to listen in on an interesting call. We never misused the device. Not one of us.

In today's deregulated world, where phone companies regularly outsource maintenance, it can be hard to recognize legitimate phone company workers. There is a telephony connection box a hundred meters from my office. I often see somebody there, doing something. I have never checked whether or not they are authorized to do whatever they do. Presumably they connect phone lines to new companies in the area around me -- but they could be installing wiretaps. When I was with the Swedish phone company, I only had to show an ID in two places: when installing a phone in a goldsmith's shop, and when fixing a problem in an office of the phone company itself. Everybody else trusted me just because I was dressed as a telephony company employee.

Tapping the old-style PSTN connections requires you to be physically close to the circuit you want to tap. Tapping a VoIP call can, in theory, be done from anywhere, as long as you manage to take control of a core router. However, doing so is no easy task. And if all the security mechanisms built into SIP are employed, you would still not be able to listen in on the conversation. SIP's protections against eavesdropping are based on several standardized building blocks: AES (Advanced Encryption Standard) performs the encryption, and SHA-1 (Secure Hash Algorithm) makes it tamperproof.

SRTP (Secure Realtime Transport Protocol) specifies how the generic algorithms AES and SHA-1 are used to protect RTP streams. The Security Descriptions document specifies how you exchange the keys needed by SRTP. TLS (Transport Layer Security) is used to protect the key exchange. AES, SHA-1, SRTP, sdescriptions, TLS. Can you trust something that uses so many cryptic abbreviations? The good news is that the core components are used in many cryptographic applications. That means that they are continuously scrutinized by security experts. Any flaw becomes widely known in the technical fora. For instance, it is well-known that SHA-1 isn't quite as good as previously thought, and as a consequence work is underway on providing a replacement.

SRTP and sdescriptions are built so that they are extensible. Once better replacements for AES and SHA-1 are available, they can be easily adopted with minimal effort.

The fact that SIP security is based on open standards also means that vendors are testing their implementations against one another. If a vendor doesn't get the algorithms right, you will hear noise or silence when using that product with a compliant product. This means that any problems that influence what is sent over the wire will be fixed.

The fact that all these standards are created in an open process also helps ensure that they really are secure. Time and time again we have seen cryptos created by a single company or using a closed process being broken: the DVD encryption and the WEP encryption just to mention two famous examples. Public scrutiny is needed to find all the flaws so that they can be fixed.

There are connectivity solutions available – software solutions deployed in the enterprise edge device – that provide the necessary functionality to allow remote workers to connect to the central PBX. These include different ways to traverse the remote NAT without any special requirements on the client or server outside the scope of the SIP standards. This far-end NAT traversal works well at wireless hot spots etc., but does not work with enterprises with strict security policies as the ports need to be open from the inside for this to work.

3.2.6 Branch office interconnect

When the PBX is IP-based, a whole host of new possibilities open up since communication between the PBX and other devices (including phones) are using a protocol (SIP) that works just as well over the Internet as on the corporate LAN. This means it is now possible to connect with other offices within the same organization or with partners and customers via IP without the need to traverse the PSTN network and without the need for dedicated circuits. This actually enables an entire, multi-site enterprise to use one centrally located IP-PBX instead of installing separate PBXs at each site.

When doing branch office interconnect of SIP-based systems, the same problems of traversing the corporate firewalls and NATs as with SIP trunking itself will occur. A SIP-capable enterprise edge device will solve this problem as for SIP trunking. Some people even refer to such an inter-office connection within an enterprise as a SIP trunk.

3.3 The service provider component

A traditional voice telephony service provider typically offers one or more T1/E1 trunks to the enterprise for fulfilling its needs for voice communication outside its own premises. The service provider is then connected to what is sometimes referred to as "the world's biggest machine:" the worldwide PSTN or Public Switched Telephony Network. Connectivity between the networks of the different service providers that constitute this "machine" is achieved by bilateral interconnect agreements between the various service providers. There are also wholesale service providers that aggregate the traffic from several local service providers and make the interconnect agreements for all of them collectively.

The SIP trunk offering is just another way of connecting the enterprise subscriber to the network. The interconnect and wholesale aspects remain the same. In a SIP trunk the traditional T1/E1 interface ("trunk") is replaced by a SIP-based connection that runs over the Internet connection to the enterprise. Nowadays, most enterprises already have such a connection to be used for their data traffic. As a SIP trunk is software- and IP based, it is much easier to manage remotely and therefore cheaper for the service provider to maintain than the traditional connections. It also typically does not require the service provider to deliver and take responsibility for any additional customer premises-based equipment. That, too, adds to the simplicity and cost effectiveness of SIP trunks as a means of delivering PSTN connectivity.

3.3.1 Different types of SIP trunking service providers

Long gone are the days when there was only one carrier available to offer telephony services. These "old" incumbents are, however, still there and they do offer SIP trunking services. These service providers typically have their own facilities all the way down to the subscriber which means that they have greater control over the quality of the service delivered. However, as discussed in the quality-of-service section that is by no means the only way to ensure that voice quality is maintained in a VoIP network.

Among the newer entrants to the voice market offering SIP trunking and other VoIP services are both facilities-based and facilities-less providers. Generally there are only a few major companies that have their own network infrastructure while others are reselling traffic that will travel on another party's (be it a "new" IP wholesaler or an incumbent) network. The number of such VoIP resellers is increasing rapidly because in the IP environment delivering such a service is relatively simple, at least in comparison to the old TDM world. With this arrangement the customer gets the best of all worlds: the facilities-based operator can focus on operating a high volume large network in the most efficient way while the reseller can focus on customer support, billing simplicity and other customer-related features of the service.

The move to IP also enables service providers to created bundled offers. There are

A Service Provider's Perspective on SIP Trunking

By Cary Tengler, Director Level3 Partner Program

The carrier community has historically had a relatively simple charter in the delivery of telephony services – design, build, operate and maintain high-performance, reliable and scalable networks, with a focus on Layers 1-4 in the OSI model.

Standards development, product innovation, and feature enhancements have typically been driven by hardware, software, and other vendors that are responsible for enterprise and end-user solutions, and with a focus on OSI Layers 5-7.

When a new technology, service or standard – SIP and SIP trunking for example – comes along and promises to increase network traffic, that's good for the carrier business. Assuming the traffic increase is expected to be significant and sustainable, carriers will be encouraged to support that product, service or standard with infrastructure investments.

SIP trunks appear to be one of these "significant" new products that offer a variety of technical and economic benefits which are described further in this document. Of equal importance, however, is the broad industry support for the underlying SIP standard and for a variety of advanced, SIP-based communications solutions. The combination of technical and economic benefits and strong vendor support has created a "perfect storm" of market influences that will promote awareness and education and ultimately drive widespread adoption of SIP trunks.

As is the case with many emerging technologies, however, the adoption decision is complicated by a number of factors.

The evolution and development of the SIP standard and SIP trunking solutions is effectively a community effort, relying on a large number of hardware and software companies and numerous service providers. Some of these vendors have a proprietary interest in maintaining compatibility with existing standards as well as protecting and migrating their installed base and many have chosen to develop extensions to SIP that are unique to their own product and not supported by the broader community.

Furthermore, the highly complex and rapidly evolving nature of telecom services dictates that no one company can deliver all of the necessary infrastructure, hardware, software, support and services required for advanced communications solutions. Non-standard "standards", proprietary extensions, competing vendor agendas, and an overall lack of awareness and education add up to a real dilemma for service providers and end-users alike.

The graphic below, from XChange Magazine, shows the considerable complexity and number of vendors involved in delivering VoIP solutions.

Level 3 Communications, with one of the world's most advanced, IP-optimized networks, was an early adopter of SIP and has been a strong supporter of VoIP and SIP-based communications solutions, with a keen interest in promoting SIP trunks to both its wholesale and enterprise customers.

→

several cases where an Internet Service Provider (ISP) adds a telephony service to its offering. Such an Internet Telephony Service Provider (ITSP) can create attractive bundles of data and voice capacity making use of the bandwidth utilization benefits.

3.3.2 PSTN connectivity

A SIP trunking service provider aggregates the traffic from many enterprise customers. The traffic passed to the PSTN is of much larger volume than the traffic from any individual enterprise. This means that the SIP trunking service provider can acquire the call minutes from the PSTN service providers at a lower rate than the individual enterprise. The network charge for the IP part of the call is typically not traffic-dependent so there are significant gains to be made here.

3.3.3 Local breakout

The use of the IP networks for part of the route of the call means that a service provider with several points of presence around the world, or that has agreements with other service providers to exchange traffic with, can allow the call to stay on the IP network for as long as possible. The call is transferred to the PSTN at the point of presence closest to the destination of the call. This process, sometimes referred to as "local breakout," allows the service provider to make maximum use of local PSTN call rates rather than paying international or long distance charges. This contributes to making SIP trunking a very cost-effective solution for the enterprise as well as for the SIP trunking service provider.

4 Interoperability

Open standards are key to the success of voice over IP adoption. Back in the mid-1990s both email and Web browsing became ubiquitous practically overnight, driving the majority of people in the industrialized world to connect to the Internet. The exponential growth was contributed to by many different organizations, companies, universities and individuals. With the large amount of participants with different interests and goals this growth would not have been possible without SMTP and HTTP, both open standards developed by the IETF.

As mentioned before, the open standard for VoIP is the IETF standard, SIP. We should expect that SIP will give rise to another period of exponential growth in Internet usage. SIP-based realtime communication over the Internet will be the source of the third wave of Internet growth.

4.1 SIP Standards

SIP stands for Session Initiation Protocol; the name describes well what it does. It is used for setting up sessions between endpoints. Endpoints are often end-user devices or servers. SIP differs from the signaling protocol of the PSTN domain in that it allows for locating much more intelligence in the endpoints rather then in centralized network elements.

SIP is specified in a growing number of IETF RFCs. In order to aid the reader to navigate through the various RFCs a "hitchhiker's guide" to SIP has been created by the IETF. That guide, and an extensive list of references, can be found at the end of this white paper.

Different groups with varied interests have taken part in adapting these standards. Some are PSTN operators who (in some cases) try to redesign the PSTN world on top of SIP. Mobile operators, 3GPP IMS as well as companies focused on data communication or IT push for support of features like IM, presence, file sharing, video etc.

As the SIP standard is comprised of a large number of specifications, most vendors do not implement all of them. SIPconnect is an example of how a specific subset of these specifications can be used for defining a limited feature set (in this case, SIP trunking).

4.2 SIP trunking by means of SIPconnect

SIPconnect was developed by the SIP Forum as a set of best practices for interfacing an enterprise PBX implementation with an ITSP that attempts to eliminate some of the unknowns and incompatibilities. The SIPconnect specification defines how a PBX located at an enterprise or organization can connect to a VoIP service provider. The primary service to be delivered by means of SIPconnect is audio-based PSTN

As a networking company, however, Level 3 does not provide all of the required customer premise equipment (CPE) and software and therefore needed to create a program to enable the development, delivery, and support for SIP trunks and a variety of other voice, data, and converged services.

Through its TAP (Technology Alliance Program) and Master Reseller programs, Level 3 has established an ecosystem of partners whose solutions have been interop tested and deployed in a wide variety of production environments. The technology partners and resellers in the two programs work with Level 3 and one another, in a collaborative fashion, to develop and deliver a wide variety of creative SIP trunking solutions for SMB through F500 enterprises, supporting many of the leading IP-PBXs currently on the market.

We believe that encouraging and investing in partner ecosystems has the potential to create a virtuous economic cycle as it relates to SIP trunking and other emerging VoIP and unified communications solutions. Multi-partner collaboration leads to the creation of new services and features, broader availability, and better delivery and support. This leads to higher adoption rates and higher customer satisfaction, thus leading to improved financial performance for all of the partners in the ecosystem, ultimately driving more investment in the creation of new services.

And when the industry produces reliable, cost-effective solutions to real customer problems, we all win.

For more information on Level 3 and the Technology Alliance and Master Reseller programs, please visit www.level3.com/partners

call origination and termination (voice).

SIPconnect refers to a number of existing IETF RFC specifications. Thus, SIPconnect provides a minimum set of requirements that are needed to be implemented at the SIP trunking service providers' end as well as at the enterprise in order to ensure interoperability. SIPconnect covers requirements in the following areas:

•DNS.
•Signaling security.
•Firewall and NAT traversal.
•Authentication and accounting.
•PSTN and SIP addressing.
•Quality of service (QoS).
•Handling of media.

Compliance with SIPconnect is vital to the overall success of SIP trunking deployments as it directly addresses, or eliminates, issues of interoperability. Compliance also future proofs the network; as new technologies based on SIP trunking are introduced, enterprises industry-wide will have the ability to leverage whatever the next big thing will be.

4.3 Interoperability

Even though the SIP standard is written with interoperability in mind, integrating SIP equipment from different vendors always takes time because, all too frequently, there are minor inconsistencies with regard to how the different vendors interpret the SIP specifications.

With regard to SIP trunking, different operators will utilize equipment from many gateway vendors who have varied requirements when it comes to the authentication of the SIP trunk user. If a company is looking to use SIP trunks from more than one vendor, e.g. in order to implement least-cost routing, they would normally have to deal with the complexities of interoperating with several SIP trunks that each behave in different ways.

As mentioned above, enterprise edge devices can mitigate these issues by addressing the complexities of interoperability. These details and the different ways to handle authorization for the SIP trunks are handled by the device. From the inside, the edge device will appear as one SIP trunk, even though it will then distribute traffic to several SIP trunks from different vendors on the outside. As the customer device that is located closest to the operator, an edge device is well placed to handle this type of operation.

Another interoperability problem common with SIP trunking is when one endpoint is located behind a SIP-unaware NAT box (home user, hotel, etc.). When the edge device is the first point of contact for such an endpoint, remote connectivity technology can enable such users to participate in both outbound and inbound calls even though they are behind a SIP-unaware NAT.

Call transfer represents another interoperability problem. Some operators do not support this feature, and some SIP user agents do not support it either. Additionally, a user who has a phone that can support call transfer cannot detect if the phone in the other end does so as well. If a call transfer attempt is made and fails, the call is often dropped.

Edge devices can detect when a call is being made to or from an endpoint that does not support call transfer. If someone still attempts to transfer a call to or from that endpoint, the device can perform the transfer itself, in lieu of the endpoint that is not able to. The call will be transferred, and the edge device makes sure that the media is sent to another destination. By using B2BUA in the device the party that does not support call transfer will still think that they called the intended person.

SIP Connect- An industry collaboration success story
By Chris Gatch, CTO, Cbeyond

Introduction

SIPconnect, the SIP Forum's Technical Recommendation for native IP connectivity between an IP PBX and VoIP Service Provider, is growing in awareness and impact among service providers, IP PBX providers, and business users. Successful deployments are increasing in number, and the family of service providers and PBX supporting SIPconnect has expanded substantially in the past 12 months.

With its widespread industry adoption, a quick view of the evolution of SIPconnect provides valuable insight into the necessity that led to its development, as well as to the challenges that still remain.

Background

In 2003 and early 2004, while most VoIP service providers touted the virtues of IP Centrex and hosted PBX services, Cbeyond asked a different question: "How can we build a VoIP-based communications service that complements the deployment of IP PBXs?" Cbeyond was not deaf to the buzz surrounding the hosted model. The company was fascinated by the growth in the IP PBX market, and market data that projected over half of deployed PBXs would be IP-enabled by 2008. Cbeyond's analysis led to an internal R&D effort to produce a new version of its BeyondVoice® services bundle of local and long distance voice, mobile and broadband Internet services designed specifically for IP PBXs. The company knew that a key foundation of the new version of BeyondVoice was native IP connectivity between Cbeyond's 100% VoIP network and the growing number of VoIP-based PBXs being deployed.

As Session Initiation Protocol (SIP) began to dominate VoIP protocols, it was the logical technology choice to achieve Cbeyond's development goal. The company began a prototyping effort with some of its close partners, including Cisco, Broadsoft, and Talkswitch, a small and nimble IP PBX company based in Canada. All of its partners believed they supported SIP 'out of the box,' but testing soon showed that "SIP" meant different things to different companies. Cbeyond's first attempts at SIP Trunking revealed a myriad of interoperability challenges.

As Cbeyond continued to work through issues, it quickly became apparent that companies supporting SIP all implemented different combinations of the underlying Internet Engineering Task Force (IETF) Request for Comments (RFCs) and Internet Drafts that composed the overall SIP standard. Additionally, SIP's flexibility supported many ways to perform a single task, and each company had their own thoughts regarding which way was best. Finally, beyond the SIP standards, there were other areas of interoperability challenges such as codec selection, fax, modem and QoS. Cbeyond and its team of partners kept careful notes as they worked through these challenges, and a list of 'best practices' began to emerge. In late 2004 and early 2005 the efforts of the team starting expanding to include additional business partners, and a consensus was reached to publish the team's work.

→

5 Security considerations for SIP trunking

5.1 Threats

Connecting a device to the Internet exposes the entire network to many types of threats. One example is a brute force attack where the intruder tries to log into a service using a user/password database trying a huge number of username and password combinations until the intruder finally succeeds in finding the right one. Once access has been granted the intruder may be able to launch other types of attacks based on known vulnerabilities to the service in question and in this way get access to other services or data.

Another example of a threat would be Denial of Service (DoS) attack where the attacker uses many different hosts or "zombies" to send a large number of packets to make the host drown or crash due to the vast amount of traffic.

The above are two examples of traditional data communication attacks. These and many others can easily be transformed into attacks on VoIP equipment. The VoIP Security Alliance or VOIP-SA has categorized possible attacks and threats on a VoIP system and made this information publicly available. This document is a resource for understanding what threats needs to be taken into account when it comes to securing VoIP in SIP trunking scenarios.

5.2 Importance of a stable platform

Firewall vendors have developed significant expertise in securing data communication. They know how to design stable systems that are locked down to only admit services that have been configured to pass. Firewalls inspect and log traffic and, if intelligent enough, they can even block suspected attacks including traffic from known malevolents.

Firewalls alone cannot prevent DoS attacks, but they can be built to withstand attacks, making them harder to occur. Firewalls can also lay the foundation for a swift recovery. More importantly, they can be built to protect the enterprise LAN from being reached by the DoS attack.

5.3 SIP signaling

Firewalls with a SIP server and full SIP proxy play a critical role in maintaining enterprise security, and securing SIP trunks. They can rewrite SIP signaling and process in a very flexible way, ensuring correct routing and interoperability with other systems built to RFC 3261 and related standards.

One important part of the SIP proxy is the SIP parser. The SIP parser verifies that the SIP message is valid and that it may be forwarded to the local LAN. Malformed SIP messages are discarded. The SIP parser must be robust enough to withstand any types of malformed SIP messages without crashing. Also, to mitigate DoS attacks, the parser should be able to process a very large number of packets.

The SIP proxy should include support for the optional loop detection mechanism defined in the SIP specification. This mechanism discerns whether a SIP message is looping (sending the SIP message to itself) and, if so, aborts this behavior. This detection mechanism also protects against DoS attacks where a SIP message is constructed to create loops and thus keep the SIP proxy too busy to engage in useful processing.

In order to protect resources, e.g. a PSTN gateway, authentication of SIP users should be supported. The standard means of authentication of SIP users is via the Digest protocol. SIP users' credentials should be stored in a centralized database e.g. on a RADIUS server. This is more secure and likely easier to maintain.

SIP signaling consists of messages in ASCII text (plain text), and are therefore easy to read and manipulate. It is strongly recommended to encrypt and authenticate SIP signaling. This is normally achieved by supporting TLS or MTLS. MTLS is the most secure method as both server and client mutually authenticate each other using CA-signed certificates or certificate chains.

This was logical step since the value of interoperability guidelines or standards lies almost exclusively in how widely they're adopted. Placing a bookmark on this lengthy period of private work, the SIPconnect Interface was announced in a press released during VoiceCON in February of 2005 with Cbeyond, Avaya, Broadsoft, Cisco, Mitel and Talkswitch participating. The resulting interest in the standard was strong, and recommendations soon followed to move the SIPconnect effort into a more mainstream industry organization where additional companies and technologists could debate and refine the technical aspects of the specification.

SIPconnect soon found a new home in the SIP Forum where a technical working group was formed to leverage the broad input of Forum members to refine and improve SIPconnect. By November of 2006 the work was complete, and the improved SIP Forum version of SIPconnect was announced at the VON Conference on September 11, 2006. This version enjoyed broad input from the industry as well as many SIP luminaries from the IETF who authored the SIP standards that form the basis of SIPconnect.

The Road Ahead

The impact of SIPconnect has been substantial. Today most new IP PBXs support a majority of SIPconnect requirements and many service providers offer SIP Trunking services. By any objective measure interoperability has improved. Today the SIP Forum is looking toward the future of SIPconnect. Starting in 2007 the SIP Forum will offer a certification mark licensing program designed to raise awareness of SIPconnect among enterprises by giving service providers and IP-PBX companies a "SIPconnect Compliant" logo to use in their marketing materials. The SIP Forum is also looking at new ways to expand and refine SIPconnect and plans to publish new technical documents dealing with more complex calls flows and provisioning this year.

If you are interested in getting involved, visit the SIP Forum website at www.sipforum.org. Membership is free and open to all interested parties.

In order to provide greater and more flexible protection mechanisms, filters are useful features. A typical filter would include the following:

•SIP methods can be allowed or prohibited per network.
•Authentication can be enabled or disabled per network and SIP method.
•SIP messages can be filtered on content type.
•Incoming callers can be restricted to a white list; this list can be individually enabled/ disabled per user.
•A filter based on from/to header may be used to allow or disallow processing.

5.4 Controlling media

SIP proxy technology is an excellent way to add a level of control to the flow of SIP media. This control offers tremendous advantages with regard to security.

The main purpose of SIP is to set up a media session between clients. Media is handled by other protocols (often RTP). For media to traverse the enterprise edge, the SIP proxy must dynamically open the media ports for media to flow during the duration of the call. As soon as the call is completed the media ports are closed. This behavior is much more secure than solutions with non–SIP-aware firewalls/border elements where a media port range constantly needs to be open. In general the SIP proxy approach is more secure than the IETF specified STUN/Turn/ICE methods, which requires that ports are left open from the inside of the firewall to allow media port negotiation to succeed.

In addition to the dynamic opening and closing of media ports, the edge device should only accept incoming media from the endpoint that receives media from the edge device. This protects against hackers trying to inject media from other endpoints or devices.

To protect media from being overheard by unauthorized persons, media encryption comes into play. The industry seems to have chosen SRTP using sdescriptions for key exchange as the de facto standard for media encryption. Using SRTP to encrypt media traversing the Internet effectively stops eavesdropping. The integrity of the call is much stronger than ever possible on PSTN.

6 Quality and reliability issues

One of the main concerns about VoIP and SIP trunking is with regard to Quality of Service and reliability. Will voice quality be good enough? Will the telephony service be available when I need it? The answer to both questions is definitively yes. In fact, many people who use traditional PBXs are using VoIP without knowing it, as many service providers use IP in their backbone networks.

Clearly, IP is not the issue. How the network is managed and planned is what makes the difference.

6.1 QoS – Different service provider approaches

The bottleneck on the Internet is often the last mile connection to the enterprise premises. There are two methods used by service providers to deliver adequate Quality of Service. In theory only the service provider controlling the link the entire way will be able to guarantee an adequate level of Quality of Service. However, in practice, the service provider relying on the over-provisioning of links will also be able to offer excellent quality.

6.1.1 Service provider controlling the connection all the way

In this case the service provider owns the connection and can control the equipment all the way from the enterprise to their SIP trunking PSTN termination point. This makes it possible to prioritize the voice traffic over data and also to give different Service Level Agreements (SLA) for different customers.

6.1.2 Over-provisioning of links

Here, the SIP trunking service provider facilitates the connection all the way to the subscriber. Any Internet connection is possible as long as there is enough bandwidth. Good voice quality is achieved by over-provisioning of the link so that the last mile never becomes a problem.

6.2 Prioritization of voice traffic

To maximize the utilization of a given capacity, both data and voice should be delivered in the same connection. However, this makes prioritization of the voice traffic necessary.

Prioritization, which can take place in the firewall or edge device, can be based on:

•Services (protocol and port).
•Packet size.
•SIP traffic.
•IP-address and segments.

This prioritization should be possible for both outbound and inbound traffic. It should also be dynamic so that bandwidth dedicated for voice can automatically be used for data when it is available.

The setting of Type of Service (TOS) and/or DiffServ bits on packet level will make it possible for routers on the Internet to make prioritizations. There is no guarantee, however, that all equipment on the Internet are using these settings for prioritization. In this case it will of course help if the service provider controls the communication all the way out to the customer premises.

6.3 Call admission control

Call admission control, also implemented in the edge device, makes sure that it is not possible to initiate more calls than what should fit into the link. The administrator defines the amount of bandwidth that is dedicated for voice and the bandwidth per call based on the codec used for voice. The edge device then keeps track of all calls; and when the dedicated bandwidth is used no more calls can be made or received. The response from the edge device in this case will be "service unavailable." It is important to reserve call slot(s) for emergency calls.

6.4 Poor voice quality can be a client problem, or based on the internal LAN

Poor voice quality is often a client problem. It is commonly known that the general performance of a PC degrades over time due to badly managed software installations and fragmented hard disks. These issues affect voice quality.

In addition, many PCs (especially laptops) may not have a sound card optimized for voice. It is highly advised to invest in a high-performance headset with a built-in sound card if the PC is meant to be used as the primary phone.

Another often overlooked factor is the QoS on the internal LAN. If the LAN is the bottleneck the voice quality will be poor no matter how good the quality of the Internet connection may be.

6.5 MPLS

Many operators offer MPLS as a means of delivering QoS in a VoIP service. The MPLS network is a service provider-managed VPN. However, it is as easy to achieve good Quality of Service in an open standards-based SIP trunking connection as with MPLS. One of the most important factors is whether the service provider controls the links all the way from the enterprise to the PSTN termination or not, not which protocol is used.

Also, SIP trunks are sometimes delivered over an MPLS connection for voice only. This means there is no support for global SIP connectivity over the Internet and the solution can never be more than just a one-to-one replacement of the traditional TDM lines.

6.6 Reliability of SIP trunks

Another argument commonly heard is that a SIP trunking connection is not as reliable as the traditional TDM. It is true that Internet connections are more dependent on electrical power, and TDM lines may have a slightly better average uptime in many parts of the world. However, many enterprise telephony systems also rely on electrical power, so a policy with uninterruptible power supply (UPS) that corresponds to the desired uptime is a must. Furthermore a TDM line, when down,

is truly dead. With SIP trunks alternative backup solutions are available.

The migration to SIP trunks will not happen overnight, so the enterprise might optionally choose to keep some traditional TDM/PSTN gateway capacity as a backup system.

With the right choice of redundancy features and service provider, SIP trunking may even offer higher reliability than many TDM-based networks.

6.7 SIP Trunking may be more reliable

Due to the inflexibility in the TDM in terms of number of lines, it is tempting to have a common PRI pool of lines at the headquarters also serving the branch offices with PSTN connectivity.

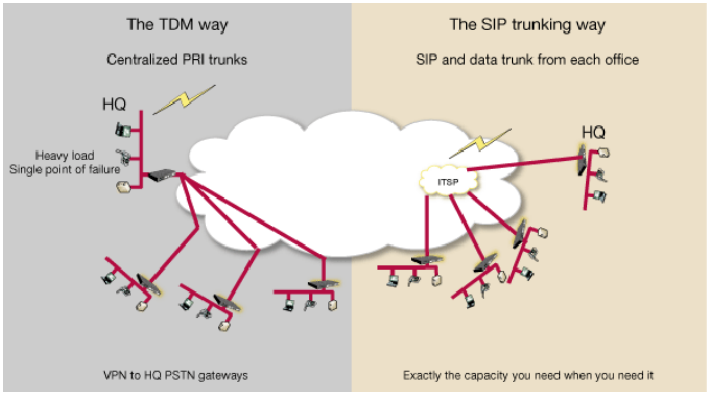

Figure 8. Comparison of TDM solution and a SIP trunking solution.

Many IP-PBX installations look like the left side of the figure above.

This will provide a single point of failure combined with an unnecessary high load at the headquarters. The SIP trunking scenario on the right offers higher reliability (here, with the different sites independently connected to the SIP trunking provider).

In many cases a SIP trunking connection may be more reliable than the traditional TDM in itself. It offers more backup alternatives including the ones described in the following sections.

6.7.1 Failover to secondary SIP trunking provider

With SIP trunking it is possible to utilize multiple service providers for PSTN termination. The edge device handling the SIP trunking connection should be able to automatically failover to a secondary (or tertiary, and on and on) SIP trunking provider if the connection to the primary service provider fails.

In addition, to make the switch triggered by a failed call, the device should be able to monitor the primary service provider by periodically sending SIP option messages and make the switch if the service provider fails to answer.

6.7.2 Failover to secondary Internet service provider

The edge device should also be able to failover to a secondary Internet service provider if the primary goes down. It is important to be able to automatically switch back to the primary once it is operational again. This will make it possible to have a cheaper backup Internet service provider.

It should be noted that many service providers share the last mile, so there is really no point in having multiple service providers if they all use the same equipment.
Ideally, the different connections should be divided e.g. the primary Internet connection delivered with an optical fiber and the secondary as an xDSL line.

6.7.3 Failover to secondary edge equipment

Finally, it should be possible to install the edge device in failover pair so it can failover to a secondary unit if the primary experiences a hardware error. The type of failover capability depends on the individual needs of the enterprise and can be divided into three levels:

1. Plain hardware failover where both registrations and on-going calls are lost.
2. Failover with registrations maintained.
3. Failover with both registrations and ongoing calls maintained.

To make it meaningful to have a failover unit in a SIP trunking scenario, the unit should at least have failover with registrations maintained, since with level one (above) it may take time for the phones to realize that they need to re-register and thus it will take time to become operational again. It might be acceptable losing the ongoing call in the case of a hardware failure, but it must be possible to redial again immediately when the failover unit is activated.

7 Summary

In the end, it all boils down to this: can we trust SIP trunking? The answer is yes, indeed we can trust SIP trunking and its applications as long as we employ the right measures to secure media, ensure interoperability/future proof the network with standards-based equipment, and are smart about the way SIP trunks are deployed. By including a SIP-capable edge device as part of the deployment, security, QoS and interoperability issues can be reduced significantly. This translates into excellent voice quality, an easier deployment and seamless interoperability – i.e., an overall better experience.

We see SIP trunks as paving the way to an all IP, all SIP world where businesses can work without geographical constraints, employees can contribute equally regardless of location, and everyone is reachable anywhere and anytime as long as there's access to an Internet connection. This is the vision the IETF had when they first introduced the SIP protocol, the idea of true global connectivity. SIP trunking extends the notion of seamless connectivity within a business to customers, remote employees, anyone working outside the corporate network. This is the next evolution of telecommunications – we look forward to sharing it with you.

About the author

Janne Magnusson, VP Operations, Ingate Systems
Formerly a project manager at Cendio Systems, Janne Magnusson led the team responsible for developing the product that was the father to Ingate's series of firewalls. When Cendio teamed with Intertex Data to form Ingate, Janne became Ingate's Director of Development, responsible for merging Cendio's firewall product with Intertex's SIP-capability to creating a suite of products that would enable small to midsized businesses to utilize SIP-based IP communications such as presence, instant messaging, video/audio conferencing and IP telephony. Prior to joining Cendio Systems, Janne accrued a wide range of communications and network experience in both the business and military sectors, including one of the first implementations of the SAAB 39 Gripen fighter aircraft's Command and Control System. Janne holds a degree in electrical engineering from the University of Kalmar in Sweden.

About Ingate® Systems

Ingate® Systems develops firewall technology and products that enable SIP-based live communication for the enterprise while maintaining control and security at the network edge. Ingate has a long history of developing next-generation firewall technology that solves the NAT/firewall traversal issue with SIP communications. In addition to an extensive line of Ingate Firewalls®, the company also produces the award-winning Ingate SIParator®, a device that connects to an existing network firewall to seamlessly enable SIP communications. Ingate products currently protect the networks of retail companies, financial institutions, industrial firms, government agencies and small-to-large enterprises throughout Europe, Asia and North America. Ingate Systems AB is headquartered in Sweden with offices in Stockholm and Linköping. Its wholly-owned subsidiary, Ingate Systems Inc., is located in Hollis, New Hampshire, with a U.S. technology center in Frisco, Texas. For more information on Ingate Systems, visit www.ingate.com.

About the authors – guest editorials

Chris Gatch, Chief Technology Officer, Cbeyond
Chris Gatch is the Chief Technology Officer and a founder of Cbeyond, Inc., an IP-based managed services provider focused on small businesses that started in 1999 and is now publicly traded on the NASDAQ under the ticker CBEY. Chris is a contributor to the industry effort to standardize SIP Trunking and serves as an editor of the SIPconnect technical specification published by the SIP Forum. He has served on the Service Provider Board of the International Packet Communications Consortium (IPCC), and presently serves on the Board of the SIP Forum.

Cbeyond
Cbeyond, Inc. (NASDAQ: CBEY) is a leading IPbased managed services provider that delivers integrated packages of local and long-distance voice along with mobile and broadband Internet services to more than 27,000 small businesses in Atlanta, Chicago, Dallas, Denver, Houston, Los Angeles and San Diego. Cbeyond offers more than 20 productivity-enhancing applications including BlackBerry(R), voicemail, email, Web hosting, fax-to-email, data backup, file-sharing, and VPN. Cbeyond manages these services over a private, 100-percent Voice over Internet Protocol (VoIP) facilities-based network. For more information on Cbeyond, visit www.cbeyond.net.

Per Cederqvist, Chief Architect, Ingate Systems
Having been on the firewall development team since 1996, Per Cederqvist is one of the people who has worked with the Ingate Firewall product line the longest. At the University of Linköping Per studied computer technology and, while at University, he co-founded the company Cendio Systems (then known as Signum Support). During his time at school Per successfully participated in several international programming contests, as well. Per has also been working as a programming consultant and written the CVS manual.

Ingate Systems
Ingate® Systems develops firewall technology and products that enable SIP-based live communication for the enterprise while maintaining control and security at the network edge. Ingate has a long history of developing next-generation firewall technology that solves the NAT/firewall traversal issue with SIP communications. In addition to an extensive line of Ingate Firewalls®, the company also produces the award-winning Ingate SIParator®, a device that connects to an existing network firewall to seamlessly enable SIP communications. Ingate products currently protect the networks of retail companies, financial institutions, industrial firms, government agencies and small-to-large enterprises throughout Europe, Asia and North America. Ingate Systems AB is headquartered in Sweden with offices in Stockholm and Linköping. Its wholly-owned subsidiary, Ingate Systems Inc., is located in Hollis, New Hampshire, with a U.S. technology center in Frisco, Texas. For more information on Ingate Systems, visit www.ingate.com.

Cary Tengler, Director Level3 Partner Program, Level3
Cary Tengler joined Level 3 in 2004 and currently serves as Director of Partner Programs, where he is responsible for the company's Master Reseller and Technology Alliance Programs and global alliances.

Prior to joining Level 3, Cary gained extensive experience managing strategic channels and alliance partnerships in the IT, media and telecommunications markets for Apple Computer, CBS, Requisite Technology, and Verint.

Level3
Level 3 Communications, Inc. (Nasdaq: LVLT), an international communications company, operates one of the largest Internet backbones in the world. Through its customers, Level 3 is the primary provider of Internet connectivity for millions of broadband subscribers. The company provides a comprehensive suite of services over its broadband fiber optic network including Internet Protocol (IP) services, broadband transport and infrastructure services, colocation services, voice services and voice over IP

services. These services provide building blocks that enable Level 3's customers to meet their growing demands for advanced communications solutions. The company's Web address is www.Level3.com.

John Casselman, Shoretel

John Casselman joined ShoreTel in 2003 in Product Marketing. In this capacity he works with all products from ShoreTel along with business vendor partners building solutions.Casselman brings more than 15 years of high tech experience to ShoreTel, serving in different technology roles. Before joining ShoreTel, he served as a Technical Marketing Engineer for Extreme Networks and 3Com. Prior to those two companies working for Quantum Corporation in the networking team.

Shoretel
ShoreTel is a leading provider of IP telephony solutions worldwide. ShoreTel voice systems provide enterprise customers with a number of key benefits, including ease of use, manageability and lower total cost of ownership than alternative solutions. ShoreTel's distributed software architecture and switched-based hardware platform extend enterprise-class voice services to every office and outpost, keeping employees fully connected wherever they go.

Founded in 1998, ShoreTel has achieved broad industry recognition for its technology and high customer satisfaction. For the last three years, IT executives surveyed by Nemertes Research, an independent research firm, have rated ShoreTel highest in customer satisfaction among leading enterprise telecommunications systems providers. A select, worldwide group of channel partners provide service and support. ShoreTel is headquartered in Sunnyvale, California and has regional offices in the United Kingdom, Sydney, Australia and Munich, Germany. For more information, visit http://www.shoretel.com/ or call 1-877-80SHORE.

References

IETF SIP Specifications

There are a great many IETF RFCs and drafts that together define the SIP standard. Rather than listing them all here we have chosen to point to an excellent summary, provided by the IETF, called, "A Hitchhikers Guide to SIP."

This summary can be found here:
http://www.ietf.org/internet-drafts/draft-ietf-sip-hitchhikers-guide-02.txt

The SIP Connect Spec

SIPconnect related specifications:

Hitchhikers guide to SIP, an overview of important SIP RFC's and drafts.
http://www.ietf.org/internet-drafts/draft-ietf-sip-hitchhikers-guide-01.txt

sf-draft-twg-IP_PBX_SP_Interop-sibley-sipconnect

SIP Connect related specifications:

Standard ID	Description
E.164	ITU-T Recommendation E.164 *Defines the international public telecommunication numbering plan used in the PSTN and some other data networks. It also defines the format of telephone numbers. E.164 numbers can have a maximum of 15 digits and are usually written with a + prefix*
RFC 2246	The TLS Protocol Version 1.0 *This document specifies Version 1.0 of the Transport Layer Security (TLS) protocol. The TLS protocol provides communications privacy over the Internet. The protocol allows client/server applications to communicate in a way that is designed to prevent eavesdropping, tampering, or message forgery.*
RFC 2833	RTP Payload for DTMF Digits, Telephony Tones and Telephony

	Signals *This memo describes how to carry dual-tone multifrequency (DTMF) signaling, other tone signals and telephony events in RTP packets.*
RFC 2782	A DNS RR for specifying the location of services (DNS SRV) *This document describes a DNS RR which specifies the location of the server(s) for a specific protocol and domain.*
RFC 3261	SIP: Session Initiation Protocol
RFC 3262	Reliability of Provisional Responses in Session Initiation Protocol (SIP) *This document specifies an extension to the Session Initiation Protocol (SIP) providing reliable provisional response messages.* *This extension uses the option tag 100rel and defines the Provisional Response ACKnowledgement (PRACK) method.*
RFC 3263	Session Initiation Protocol (SIP): Locating SIP Servers *The Session Initiation Protocol (SIP) uses DNS procedures to allow a client to resolve a SIP Uniform Resource Identifier (URI) into the IP address, port, and transport protocol of the next hop to contact. It also uses DNS to allow a server to send a response to a backup client if the primary client has failed. This document describes those DNSprocedures in detail.*
RFC 3264	An Offer/Answer Model with Session Description Protocol (SDP) *This document defines a mechanism by which two entities can make use of the Session Description Protocol (SDP) to arrive at a common view of a multimedia session between them. In the model, one participant offers the other a description of the desired session from their perspective, and the other participant answers with the desired session from their perspective. This offer/answer model is most useful in unicast sessions where information from both participants is needed for the complete view of the session. The offer/answer model is used by protocols like the Session Initiation Protocol (SIP).*
RFC 3311	The Session Initiation Protocol (SIP) UPDATE Method *UPDATE allows a client to update parameters of a session (such as the set of media streams and their codecs) but has no impact on the state of a dialog. In that sense, it is like a re-INVITE, but unlike re-INVITE, it can be sent before the initial INVITE has been completed. This makes it very useful for updating session parameters within early dialogs.*
RFC 3323	A Privacy Mechanism for the Session Initiation Protocol (SIP) *This document defines new mechanisms for the Session Initiation Protocol (SIP) in support of privacy. Specifically, guidelines are provided for the creation of messages that do not divulge personal identity information. A new "privacy service" logical role for intermediaries is defined to answer some privacy requirements that user agents cannot satisfy themselves. Finally, means are presented by which a user can request particular functions from a privacy service.*
RFC 3324	Short Term Requirements for Network Asserted Identity *A Network Asserted Identity is an identity initially derived by a Session Initiation Protocol (SIP) network intermediary as a result of an authentication process. This document describes short term requirements for the exchange of Network Asserted Identities within networks of securely interconnected trusted nodes and to User Agents securely connected to such networks.*
RFC 3325	Private Extensions to the Session Initiation Protocol (SIP) for Asserted Identity within Trusted Networks *This document describes private extensions to the Session Initiation Protocol (SIP) that enable a network of trusted SIP servers to assert the identity of authenticated users, and the application of existing privacy mechanisms to the identity problem. The use of these extensions is only applicable inside an administrative domain with previously agreed-upon policies for generation, transport and usage of such information. This document does NOT offer a general privacy or identity model suitable for use between different trust domains, or use in the Internet at large.*
RFC 3489	STUN - Simple Traversal of User Datagram Protocol (UDP) Through Network Address Translators (NATs) *Simple Traversal of User Datagram Protocol (UDP) Through Network Address Translators (NATs) (STUN) is a lightweight protocol that allows applications to discover the presence and types of NATs and firewalls between them and the public Internet. It also provides the ability for applications to determine the public Internet Protocol (IP) addresses allocated to them by the NAT. STUN works with many existing NATs, and does not require any special behavior from them. As a result, it allows a wide variety of applications to work through existing NAT infrastructure.*

RFC 3581	An Extension to the Session Initiation Protocol (SIP) for Symmetric Response Routing *The Session Initiation Protocol (SIP) operates over UDP and TCP, among others. When used with UDP, responses to requests are returned to the source address the request came from, and to the port written into the topmost Via header field value of the request. This behavior is not desirable in many cases, most notably, when the client is behind a Network Address Translator (NAT). This extension defines a new parameter for the Via header field, called "rport", that allows a client to request that the server send the response back to the source IP address and port from which the request originated.*
RFC 3725	Best Current Practices for Third Party Call Control (3pcc) in the Session Initiation Protocol (SIP) *Third party call control refers to the ability of one entity to create a call in which communication is actually between other parties. Third party call control is possible using the mechanisms specified within the Session Initiation Protocol (SIP). However, there are several possible approaches, each with different benefits and drawbacks. This document discusses best current practices for the usage of SIP for third party call control.*
RFC 4028	Session Timers in the Session Initiation Protocol (SIP) *This document defines an extension to the Session Initiation Protocol (SIP). This extension allows for a periodic refresh of SIP sessions through a re-INVITE or UPDATE request. The refresh allows both user agents and proxies to determine whether the SIP session is still active. The extension defines two new header fields: Session-Expires, which conveys the lifetime of the session, and Min-SE, which conveys the minimum allowed value for the session timer.*

VoIPSA Security taxonomy
http://www.voipsa.org/Activities/taxonomy.php