Session Border Controllers

SBC | The Critical Component

Sangoma

WHITE PAPER

Executive Summary

As Voice-Over-IP (VoIP) has been replacing the legacy telecommunication infrastructure, the Session Border Controller (SBC) has become a crucial piece in VoIP Protection, SIP Interoperability, media transcoding and call routing. The completely new way in which VoIP interconnects with networks, versus the legacy Public Switched Telephone Network (PSTN), has exposed many pitfalls in the way communication security policies are put in place as well as how we integrate network components. Careful consideration is required in hiding network element details as well as those exposed within a typical VoIP phone call to prevent malicious attacks. Since the signaling protocol for VoIP, called SIP (Session Initiation Protocol) is not a tightly regulated protocol, the differences in how individual implement it can cause VoIP connections and phone calls to fail as the differing endpoints do not cooperate. Lastly, VoIP poses new challenges with how remote colleagues, guests and other users interact with the new technology. The integration of a Session Border Controller at the edge of VoIP networks is the perfect solution to manage all the above challenges. The SBC's ability to intercept and interpret all voice traffic in VoIP network allows it to protect the communication infrastructure from VoIP attacks and automatically deal with the SIP signaling protocol variations to provide secure, successful voice communication within a VoIP infrastructure.

Table of Contents

SBC – The Critical Component of your VoIP Infrastructure

Voice over internet protocol (VoIP) offers many operational cost, feature and flexibility advantages over the incumbent telephone system, and is rapidly replacing the public switched telephone network (PSTN).

As the compelling advantages of VoIP drive an increasing number of businesses and organizations to switch over to state-of-the-art telephony technology, the complexities and pitfalls of VoIP must be addressed.
Several issues arise when managing a VoIP system:

- Security
- Remote worker applications
- Challenges of a complicated network
- SIP interoperability / multi-vendor / bring your own device

The existing security features of a network can make the enabling of media and signal flow between communication endpoints technically challenging. Special measures are required to enable secure VoIP communications without compromising existing network security elements such as firewalls.

Corporate networks are becoming ever more complicated, due to security issues and by the variety of interoperating devices. Telecommuter applications require remote access, further complicating security issues. Meanwhile, the rise of UC (Unified Communications) means a greater variety of media need to securely use the network, including a growing use of video and other rich communication media.

New systems have new and sometimes unfamiliar or unknown vulnerabilities. Traditional communications via he switched telephone network, delivered over T1, E1, or PRI lines, may be limited and expensive, but they are remarkably reliable and fairly secure in most parts of the world.

When a business switches to VoIP and unified communications, a very high degree of security and reliability is desired and expected. But packet-switched networks, such as the public Internet, while efficient and economical, are not inherently as secure as their TDM counter parts. Additional network elements are required to achieve the expected characteristics.

While it is true that networks are getting more complicated, it is also true that the equipment to manage the network is getting more sophisticated. This is good news for organizations that care about both their ability to manage the network, and the security of their data and communications systems.

Enter the SBC

The Session Border Controller (SBC) provides a number of services that make a VoIP/UC system more secure and better able to integrate SIP-based equipment from a variety of vendors.
Let's first look at what "session border controller" means. The session border controller manages both media and signaling streams. Each session consists of the collection of signaling and media streams that connect one party to another. A session works in only one direction, so two sessions are required for a two-way phone conversation[1].

For example, a two-way telephone conversation would consist of two sessions, each containing a signaling channel and a voice channel. A video session may consist of a signaling channel, an audio channel and a video channel. The situation for conferences may be more complicated, depending on how the conference is set up. The border is the demarcation point between two networks, such as LAN and Internet, or sections of a network, such as segments in a very large and segregated corporate network.

A controller is required to manage routing, taking into account the topology of the network, traversing firewalls, and managing Quality of Service (QoS). The routing equipment controls how traffic flows through the network. If properly configured, routers will use the ToS/DSCP/DiffServ fields in the IP header to make routing decisions. It is up to the SBC to properly set these fields for each packet.

An SBC can also manage multiple SIP trunks, including load balancing, congestion avoidance and fail-over.

[1] Technically, a two-way conversation requires two sessions. However, most SBC specifications use the term session and call interchangeably

Typical SBC Use Cases

Session Border Controllers can be used in several networking scenarios. More details and more examples are provided later in this document.

ENTERPRISE

IP-PBX with SIP Trunking

The figure below show an enterprise SBC used to provide access control and security to an IP-PBX. The SBC is typically deployed in a Demilitarized Zone (DMZ) configured from existing data firewall infrastructure.

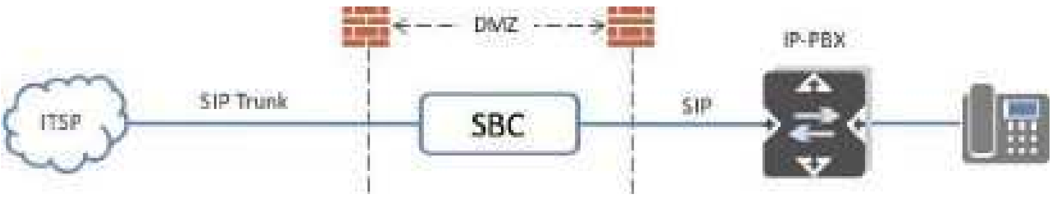

SERVICE PROVIDER

Hosted IP-PBX Services

The figure below shows a service provider using the SBC to provide access control, security and media transcoding.

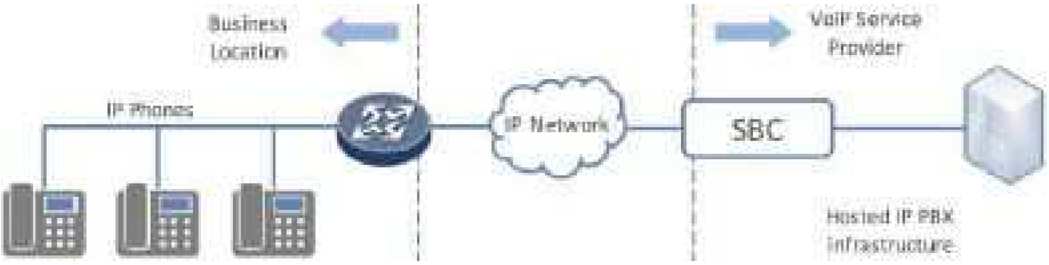

Features

The session border controller fills a number of important roles in delivering VoIP services to a local network, including:

- SIP B2BUA
- Security
- Encryption
- Policy
- Call routing
- SIP interoperability
- Media transcoding
- DTMF detection/generation
- Scalability
- High Availability

SIP B2BUA

An SBC is often referred to as a Back-to-Back-User-Agent (or B2BUA) because it is placed directly into the path of a VoIP connection, dividing it into two parts, where it mediates all the VoIP signaling between both ends of the connection. VoIP signaling is a combination of two protocols:

- The control Protocol, called Session Initiation Protocol (or SIP), which manages everything to do with how a phone call is established to how it is disconnected
- The media Protocol, called Real-Time Transport Protocol (or RTP), which handles the media for the phone call

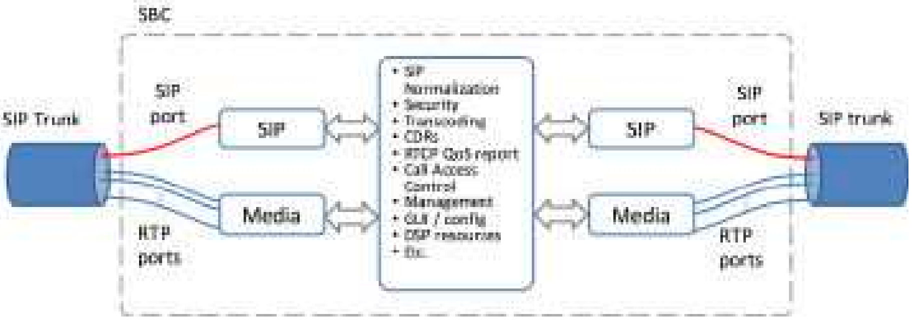

SBCs are Back to Back User Agents, in control of the VoIP flow

As a SIP B2BUA, an SBC is a very powerful device because it can 'see' all the SIP and RTP traffic coming in from both sides of the VoIP connection and call analyze it, adjust and fix it. It can greatly assist in allowing calls to penetrate networking boundaries such as between the Internet and a private LAN and can also be used to modify the encoding of the speech on each leg of the call, which is called transcoding.

Security

A key role of an SBC is to provide a robust layer of security to prevent the VoIP system from becoming a point of entry for hackers or toll fraudsters. It must do so while also smoothly routing calls through existing security systems on the network. The SBC terminates each call on one side of the network and re-originates the call on the other side. This approach allows the SBC to maintain complete control over every call, and to dynamically manage security. This hides network topology and presents a brick wall for would-be hackers intent on entering the network.

As SIP messages work their way through the network, each node adds its own 'via' field to every packet. Packets will probably pass through a PBX for example. By the time the SIP packet leaves the network, a series of 'via' fields show the route the SIP message has taken. This reveals the structure of the network behind the firewall to the outside world. Such information can help hackers break a network. The SBC removes all accumulated 'via' fields from each packet, and replaces them with a single 'via' field from the SBC, making outside SIP messages appear to have originated from the SBC. This hides the network structure from the outside world.

NAT transversal functions are used to navigate the NAT firewalls that protect typical corporate networks. The SBC opens a pinhole (single port) in the firewall for the duration of the call and performs port remapping to transfer signaling and media packets between the corporate network and internet.

As BYOD (Bring Your Own Device) becomes more popular, the security flaws that can come with BYOD grow. Smart phones may have apps installed which present a security threat. For example, the owner may download a game, which is really a Trojan horse designed to steal bandwidth from the network. By analyzing typical traffic versus rogue VoIP traffic in the SBC, the damage from such Trojans can be controlled.

Denial of service (DoS) attack is a common way for hackers to attempt to disrupt service. The SBC can detect a DoS attack by measuring traffic volume from each source, and then blocking unusual patterns. With Sangoma SBCs, the blocking is performed at the kernel level. This approach reduces the load on the SBC so that the network can operate normally during an attack.

Malformed packet attack (fuzzing) is an alternate form of DoS attack. The assailant attempts to bring down the service by sending malformed packets to the VoIP system in an attempt to break the SIP stack. The SBC detects malformed packets and blocks them. The IT manager can also configure the SBC to block any IP address that generates malformed packets.

The SBC needs to deploy other security features as well: Call access control limits the number of concurrent calls each customer can have, SIP registration security detects when too many wrong passwords are attempted within a given period of time, and known hacker user agents blocked. Toll fraud on VoIP networks is a growing problem. Toll fraud has shifted from an activity individuals undertake to make free long distance calls across the conventional phone network, to an activity carried out by more organized groups who steal minutes from unsecured corporate VoIP networks and resell them to their customers. Even the smallest PBX can be hacked and used to rack up bills of tens of thousands of dollars for unauthorized calls in a single month. Once again, an SBC can monitor rogue traffic vs normal traffic.

Encryption

Encrypted voice channels are required to prevent eavesdropping as voice packets travel public networks. It also serves the purpose of authenticating endpoints. The standard approach is to use Transport Layer Security (TLS) and Secure Real-Time Transport Protocol (SRTP) to protect signalling and voice channels respectively. Sangoma uses a hardware-based transcoding system to apply encryption. This frees the server to handle an increased call volume, allowing economical use of the SBC server for high call volumes while still providing voice encryption.

Policies

Unauthorized use of the company VoIP services can be controlled by putting appropriate policies in place. These are managed by the SBC. For example, only allowing calls between known SIP endpoints allows remote workers to access the VoIP system with a SIP phone, but prevents hackers from gaining access with an unregistered SIP phone or user agent. An integrated security library looks for patterns, such as excessive long distance calls when the office is closed.

Call Routing

Basic call routing directs each phone call so that it arrives at the intended endpoint. The flexibility of an XML-based routing file and the ability to query an internal or external database adds additional flexibility and capabilities.

Routing rules are applied to each SIP message, based on variables such as time of day; any of the variables found in the SIP header, including originating number and destination number; and SS7 parameters, if encapsulated in the SIP message.
For organizations that are using multiple SIP providers, the SBC can provide least cost routing and load balancing across trunks.

The SBC can also route around network congestion by rerouting calls when the message returned by a SIP invite indicates that the message cannot be processed.

SIP Interoperability

Although SIP is a standard protocol, it is extremely flexible, has many options, and is subject to different interpretations by different vendors. This means all SIP-compatible devices don't necessarily interoperate correctly with one another. One of the major benefits of an SBC is that it allows different devices with varied codecs and SIP protocol flavors to interconnect. This enables multi-vendor systems to operate smoothly and leaves the door open for future equipment to work seamlessly with the current infrastructure.

SIP headers can be modified by the SBC to ensure compatibility between disparate devices. In some cases, equipment such as a phone may add certain SIP headers that other equipment, such as a softswitch, cannot recognize. The SBC can remove these headers to avoid confusing incompatible hardware. SIP header modification can also be used to add extensibility, enabling custom features such as accounting, integration with a specific PBX application, call recording and more.

Media Transcoding

Transcoding is necessary to allow incompatible media types to cross the barrier between disparate devices, and to allow optimal use of available network bandwidth. For example, if bandwidth is not an issue and high quality is desired, G.722.1 is a good audio codec option. On the other hand, if bandwidth is constrained or expensive, G.729 is a better choice. In a conference scenario, there may be a mixture of codecs in the media stream from different devices, depending on endpoint equipment and individual connections. Transcoding allows a mix of codecs to work together seamlessly.

A variety of approaches are available to undertake media transcoding. Each approach has its pros and cons. The approach taken by different manufacturers varies according to their philosophy. Transcoding can be done entirely in software. This is flexible and can reduce capital costs, but it is CPU-intensive, leading to a lower call-handling capacity on the SBC.

A hardware-based transcoding approach can be used, with the transcoding hardware installed in the server. If located in the server, the footprint is reduced, but limited space in the server may restrain the number of concurrent calls the SBC can handle.
A third approach is to use externally located transcoding hardware, situated on the network at nearby or distant locations. This is more expandable than when installed in the server cabinet, but requires a bigger footprint and additional expense for enclosures, rack space, power supplies, etc.

Sangoma is the only manufacturer to offer the choice of any of these three approaches. The most appropriate configuration can be chosen depending on specific needs, and can be changed as additional capacity is required. For example, additional transcoding hardware can be purchased and added to the network as required. No matter which approach is chosen, the solution benefits from Sangoma's long history as a transcoding hardware manufacturer.

DTMF Detection / Generation

DTMF stands for Dual Tone Multi Frequency. Each key pad button on a phone is represented by a pair of tones with sinusoidal frequencies. This 1960's era in-band signaling system is still in wide use today, both for legacy phone handsets, and for interactive voice response (IVR) systems. The entrenchment of IVR systems insures that DTMF will remain an important part of corporate telephone systems for the indefinite future.

The challenge with DTMF signaling on a VoIP network is that some codecs do not reliably transmit DTMF tones due to the use of lossy bandwidth compression algorithms which are optimized for voice. While these compression algorithms can transmit voice that is intelligible to human ears, the audio processed by some codecs may remove some audio information so that DTMF tones cannot reliably be detected at the distant-end. In these cases, it is necessary to detect DTMF tones on the close-end gateway, convert the audio into data, and forward data packets representing the digits to the distant-end gateway. The distant-end gateway regenerates the DTMF tones for the end-point to "hear". RFC2833 tone relay is the standard method for handling this.

Scalability

As the number of users and voice traffic grows, whether planned or not, the network infrastructure must grow to accommodate greater capacity. Transcoding, protocol interworking and the basics of the SBC itself must scale to manage an increased workload. There are several approaches to this. One way might be to add additional SBCs to the network, and then load- balance amongst them. While such measures may be required in some situations, a more granular control over capacity is desirable so that issues with transcoding or interworking can be addressed separately, according to which subsystem is approaching capacity.

Sangoma systems decouple transcoding and interworking from other SBC systems. While it is possible to house these services in the same physical enclosure as the SBC, it is also possible to have external Sangoma interworking and transcoding appliances to manage scaling.

Transcoding can be handled in increments of 250 – 400 simultaneous calls with the addition of D150 units. These transcoders are simply plugged into the network via Ethernet connection, adding more simultaneous call capacity to the network.
The SBC function can be scaled by virtualizing the software. For organizations with a solid virtual server infrastructure, this approach is very appealing. Resources to a specific VM can be expanded and VoIP load increases or additional VM-based SBCs can be installed.

High Availability

The High Availability (HA) component provides the ability to join two SBC servers in an Active/Passive cluster with automatic migration of service in case of a hardware or network failure. The system administrator can configure both servers from the primary one and synchronize the configuration to the secondary server with a single command. In the event of a failure of one of the SBC servers, the failure is detected by the other server and the migration process begins. At that point, established calls drop and no new calls are processed. The migration period is very quick, typically lasting only 10-15 seconds.

SBC HA is a group of 2 SBC Servers that support the SBC application that can be reliably utilized with a minimum of down-time. The SBC operates by using specially designed high availability software to harness the redundant SBC Server to provide continued service when system components or network connectivity fail. Without clustering, if a SBC Server crashes, the SIP Trunking, Remote Phone or other Call Flow applications will be unavailable until the crashed SBC Server is fixed. HA resolves this situation by detecting hardware, software, and network faults, and immediately migrates the SBC application on another SBC Server without requiring administrative intervention, a process known as Failover. As part of this process, the SBC application shares and synchronizes the configuration and uses common “Floating” IP Addresses to create a singular SBC application. The Sangoma SBC makes use of Floating IPs, these are created for each of the network segments (WAN, LAN, DMZ) connected to the SBCs. These Floating IPs are common/shared on both Primary and Secondary SBCs as the common point of entry and exit for SIP Call Flow. These Floating IPs allow for the Failover to occur, when the Master Node fails, the Salve Node becomes the Master Node and take Active control of the Floating IPs.

SBC HA is often used for critical Trunking and Remote Phone applications on a Carrier network, Enterprise Network, and various VoIP business applications. SBC HA implementations add redundancy to eliminate single points of failure, including multiple network connections and Call Flow applications. There is a dedicated network connection between the two SBCs for the purpose of a “heartbeat” which is used to monitor the health and status of the other SBC Server in the cluster, as well as synchronization of the configuration within the SBC application.

Virtual SBC

As the demand for virtualized infrastructures increases, a Virtual Machine SBC is the perfect solution for Enterprises and Carriers who need VoIP security as well as transcoding benefits, while keeping their existing hardware. Depending on the vendor, the Virtual Machine SBCs typically provide the same functionality as that of a hardware-based SBCs but is 100% software. This means no additional power, space or cabling requirements to implement.

Sangoma’s Virtual Machine SBC offers the same rich functionality as that of our hardware-based SBCs, supporting up to 500 sessions. It is designed to work in leading edge virtualization platforms, including VMware, Hyper-V, KVM and Amazon Web Services. It is also compatible with most commercially available motherboards and servers so that you can install the software directly to bare metal too!.

Enterprise Use Cases

SIP Trunking

In the diagram below, the SBC provides a defined demarcation point between the internet telephone service provider, and the corporate network. It reduces interoperating issues between the ITSP (Internet Telephone Service Provider) and their clients, saving core resources from attempting to handle interoperation. It can also handle transcoding as needed. Each business also has an SBC for security and to reduce interoperating issues within the network. SBCs typically are deployed in a DMZ managed by the corporate firewall.

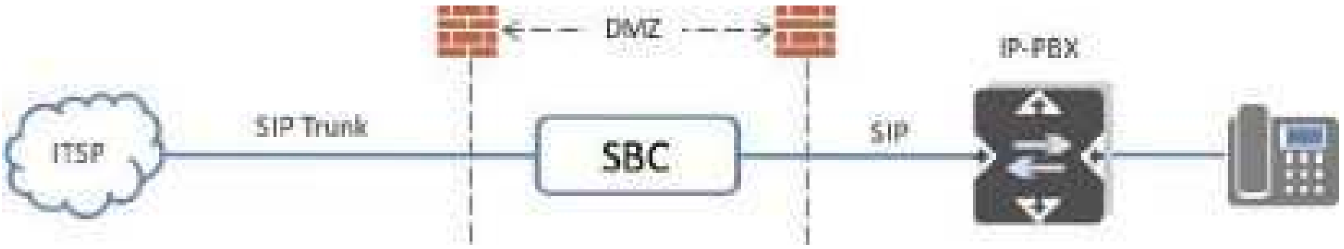

Enterprise SIP Trunking

Remote Workers

Nowadays, many businesses have workers working from home or on the road and the need to communicate as if they are in the office. With VoIP removing the geographical barriers from the old PSTN, you must take care to ensure that while allowing remote access to your IP-PBX and UC applications, you are not effectively opening the door for toll-fraud and other attacks by simply opening ports in your data firewall. The SBC here allows access to authorized endpoints on remote worker premises, while eliminating interoperation and firewall issues on corporate and remote worker networks and maintaining security. The diagram below shows a typical deployment.

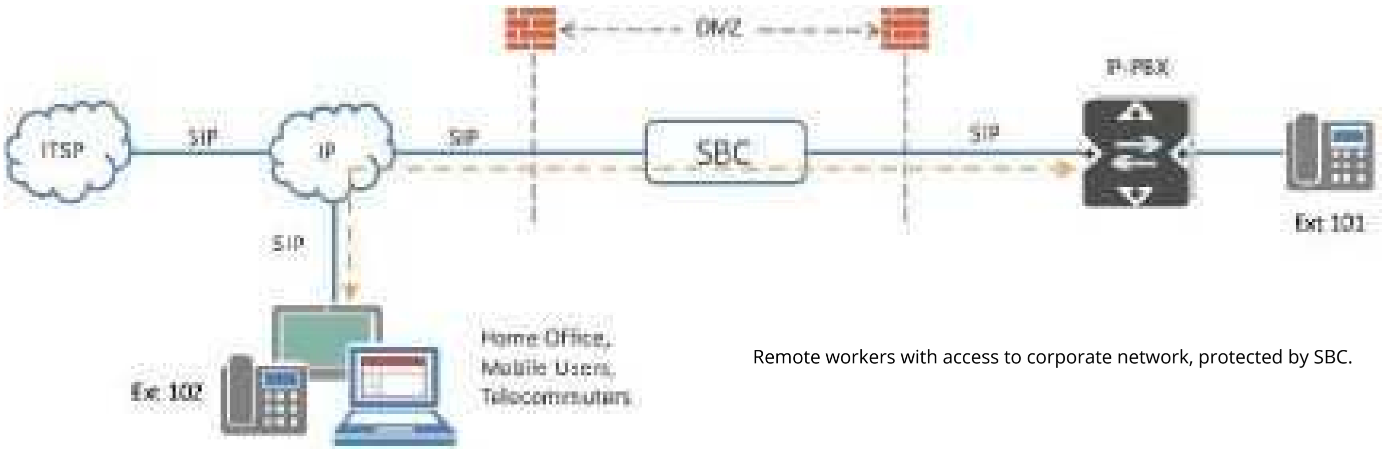

Remote workers with access to corporate network, protected by SBC.

Skype for Business Interoperability

Skype for Business (SfB) is one of the option on the market for enterprises to enable UC features. SfB however needs SBC to interconnect with the public internet, so again precautions should be taken to avoid security issues. The SBC is involved mainly for call routing and providing interoperability (such as TCP to UDP translation, but also intricate SIP call flows) between the SfB Mediation Server and the public networks. The figure below depicts a typical scenario.

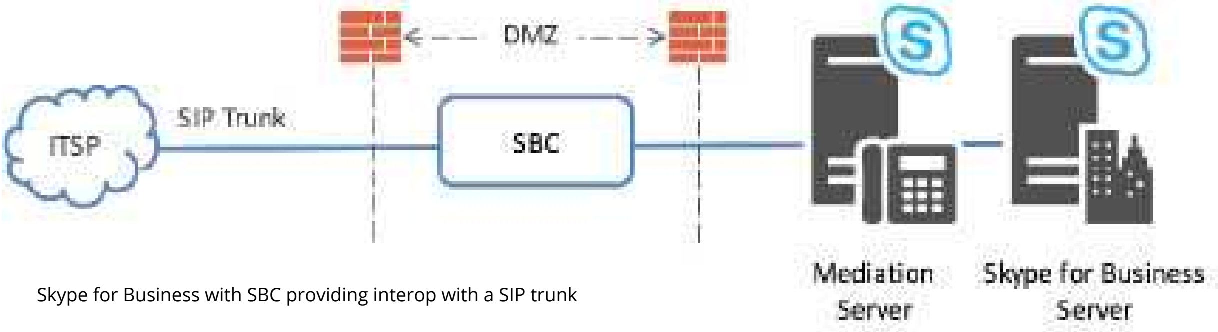

Skype for Business with SBC providing interop with a SIP trunk

Failover, Least Cost Routing, Load Balancing

For larger organizations with more sophisticated network routing (for example a large call center), several SIP trunks from different service providers can be set-up to provide failover, least cost routing or load balancing in both inbound or outbound calling scenarios. In this case, the SBC provides all the security features and also all the call routing strategies associated with these more complex calls flows. The diagram below depicts a typical scenario.

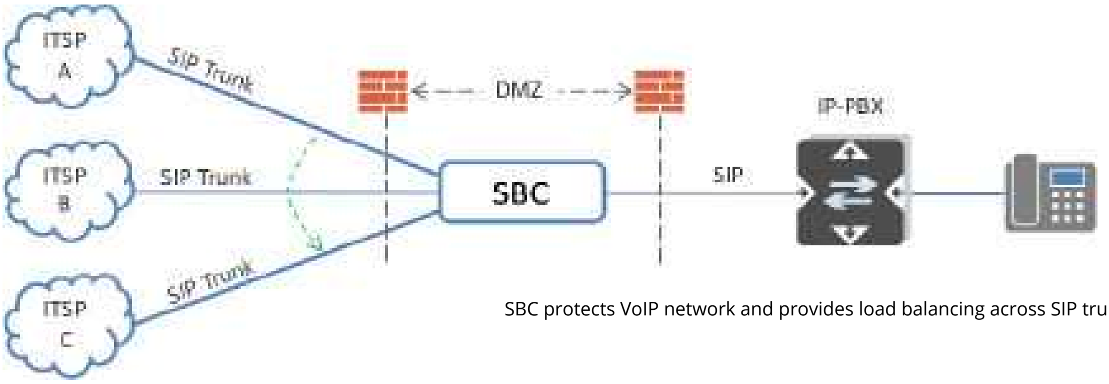

SBC protects VoIP network and provides load balancing across SIP trunks

Service Provider Use Cases

Hosted PBX

In the figure below, an access SBC protects the internet telephone service provider's network while providing access control for the IP Phones at customer sites. It also provides a defined demarcation point, and reduces interoperability issues with the core. Transcoding can be provided if required.

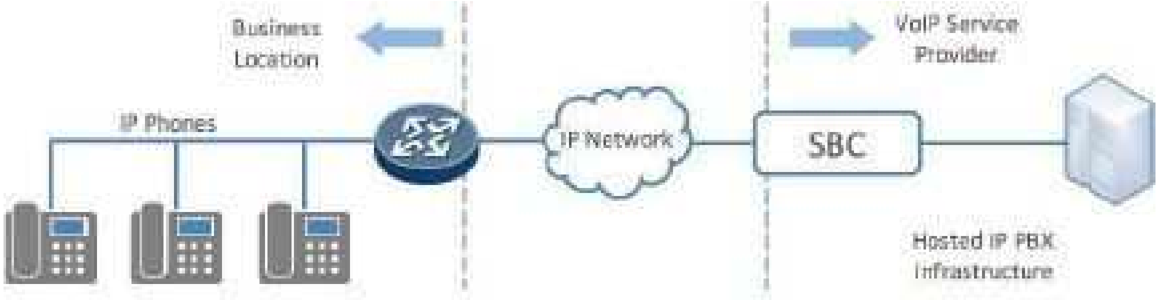

Hosted IP-PBX Access SBC

Network Peering

With the advent of VoIP, service providers are no longer using SS7 interop agreements to interconnect their networks. They are instead setting up Peering agreements, in which case large SIP trunks are set-up between their networks. The SBC acts as a peering SBC, connecting two or more IP networks (carriers, sub-networks, long distance networks, etc.) providing extensive routing capabilities, resource management and call detail records. The move to full IP communication between carriers removes TDM conversion and at same time improve voice quality.

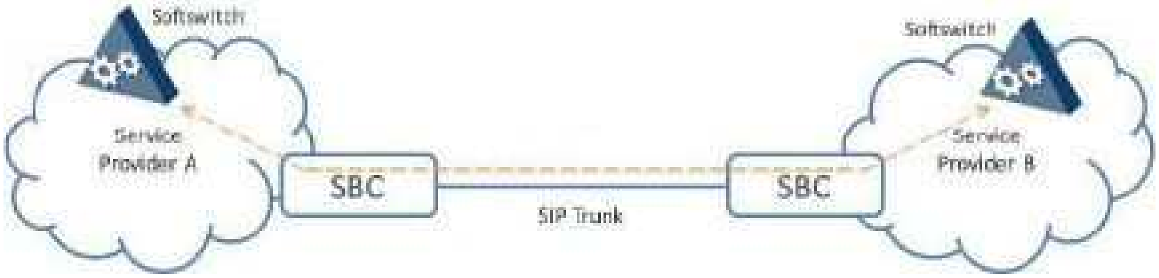

Peering Arrangement SBCs between carriers

Conclusion

When an organization makes the switch from conventional phone lines to VoIP and SIP trunks, a session border controller is highly recommended to allow VoIP calls to pass through the firewall, route VoIP traffic properly across the network, and to enhance network security. The SBC ensures that VoIP does not become a security issue in the network. The SBC also protects the network from potential security threats that BYOD (bring your own device) policies could expose.

The SBC allows disparate devices to interwork seamlessly, even when those devices use different implementations of SIP.

The SBC can perform load balancing and fail over functions between SIP trunks. This allows an organization to have multiple SIP trunk suppliers and increase the reliability of their phone service.

Meanwhile, VoIP service providers require an SBC to protect their network and correct SIP headers which may not have been properly adjusted for proper routing at the client end.

The Session Border Controller is the critical component necessary to safely and effectively utilize SIP and SIP trunks.

All businesses with significant VoIP networks, or those who rely on their VoIP network as a critical component of their business, should seriously consider installing SBCs. It is an investment that can avoid potentially costly and perhaps devastating security consequences, while increasing network efficiency and reliability; and thus, pay for itself many times over.

Sangoma

www.sangoma.com

100 Renfrew Drive, Suite 100, Markham ON L3R 9R6 Canada | +1 (256) 428-6000 or 1 (877) 344 4861 (toll free in N. America)