SESSION BORDER CONTROLLERS

ENABLING THE VOIP REVOLUTION

February 2005

Jon Hardwick
Data Connection Ltd
jon.hardwick@dataconnection.com

Data Connection Limited
100 Church Street
Enfield, EN2 6BQ, UK
Tel: +44 20 8366 1177
http://www.dataconnection.com/

TABLE OF CONTENTS

1. INTRODUCTION AND OVERVIEW

Session Border Controllers (SBCs) have become an important element of modern Voice over IP (VoIP) networks, as service providers look to protect the integrity of their networks and business models while offering diverse services to their customers.

Most people would agree that an SBC is a kind of firewall for Voice over IP traffic. However, as soon as you start to look beyond this initial consensus, there is considerable disagreement as to what an SBC actually is, and what function it should offer! This is partly because SBC vendors are pushing out to cover a wide variety of niches in order to compete for market share, and partly due to the genuine range of scenarios where service providers are looking for solutions.

This white paper starts by looking at the deployment scenarios where SBCs have a role today. It then goes on to examine the breadth of technology and functionality that different companies claim should be delivered by an SBC, assessing which functions are most important for each type of deployment. Finally, it looks in more detail at the key elements common to most SBC products.

1.1 Document Roadmap

This white paper is structured as follows.

- Chapter 2, **Introducing SBCs**, takes a high-level look at the role of SBCs within a network, and explains where they are most likely to be deployed.
- Chapter 3, **SBC function overview**, gives more detail of the types of function offered by SBCs.
- Chapters 4 to 6 then discuss the key features of SBCs in depth.
    - Chapter 4, **DMZ processing**, explains SBCs' firewall and Network Address Translation (NAT) function.
    - Chapter 5, **Firewall and NAT traversal**, describes how SBCs can be used to direct VoIP packets behind firewalls in adjacent networks.
    - Chapter 6, **Call Admission Control**, lists the types of call processing performed by SBCs.
- Chapter 7, **Conclusion**, provides a summary of the analysis.
- Chapter 8, **About Data Connection**, contains information about the author of this paper, Data Connection, and Data Connection's range of portable software (including DC-SBC).
- Chapter 9, **Glossary**, contains a glossary of some of the important terms used in this paper.
- Chapter 10, **References**, provides details of references made in this paper.

2. INTRODUCING SBCS

An SBC is a VoIP session-aware device that controls call admission to a network at the border of that network. Optionally (depending on the device), it can also perform a host of call-control functions to ease the load on the call agents within the network.

This chapter provides a general introduction to SBCs.

- Section 2.1, **Internal structure of an SBC**, describes the two distinct components that make up an SBC.
- Section 2.2, **The Demilitarized Zone**, describes the part of the network in which SBCs are mostly (though not always) used.
- Section 2.3, **Applicable network scenarios for SBCs,** describes the five main network scenarios where SBCs are deployed.

2.1 Internal structure of an SBC

An SBC device breaks down into two logically distinct pieces.

- The Signaling SBC function (SBC-SIG) controls access of VoIP signaling messages to the core of the network, and manipulates the contents of these messages. It does this by acting as a Back-to-Back User Agent (B2BUA).
- The Media SBC function (SBC-MEDIA) controls access of media packets to the network, provides differentiated services and QoS for different media streams, and prevents service theft. It does this by acting as an RTP proxy.

Some SBC devices offer both functions in a single box (referred to hereafter as single-box SBCs). Others take a distributed approach, and separate SBC-SIG and SBC-MEDIA onto separate machines (referred to hereafter as dual-box SBCs), using call control protocols such as H.248 and COPS-PR to link the two. See section 3.10, **Comparison of single-box and dual-box model**, for a discussion of the relative advantages and disadvantages of these different approaches.

2.2 The Demilitarized Zone

The Demilitarized Zone (DMZ) is the conceptual term for a small subnetwork (or individual device) that sits between a trusted private network, such as a corporate private LAN, and an untrusted public network, such as the public Internet[1]. Typically, the DMZ contains devices directly accessible to Internet traffic, such as web servers, FTP servers, or SBCs. The purpose of the DMZ is to prevent hostile or unwanted traffic from entering (or, in some cases, leaving) the private network.

The rest of this section describes how single-box and dual-box SBC deployments function within the DMZ. For more detailed information on the DMZ, see chapter 4, **DMZ processing**.

2.2.1 Single-box SBC deployments

Single-box SBCs are deployed in the DMZs of VoIP-enabled service provider (SP) networks. A schematic diagram of a DMZ containing a single-box SBC is shown below.

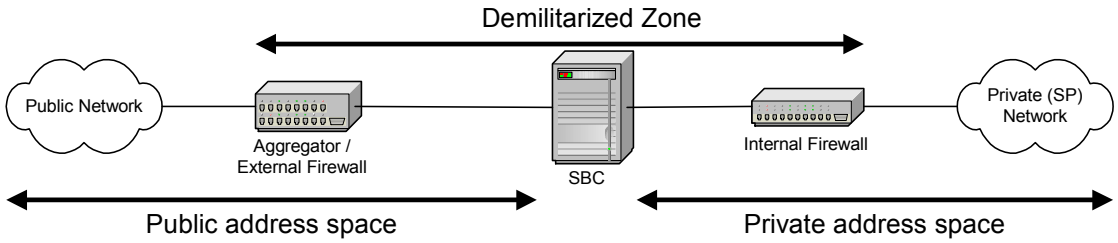

Figure 1: The Demilitarized Zone

The following devices are shown in the diagram above.

- The **external firewall or aggregator** prevents unwanted traffic arriving from the public network (or networks) from entering the DMZ, based on its configured set of packet-inspection rules.
- The **SBC** performs an Application Layer Gateway (ALG) role, translating the addresses and ports in signaling data (found within the IP packet payloads) between the internal and external addressing schemes.
- The **internal firewall** prevents unauthorized traffic from leaving the private network and entering the DMZ. In some DMZs, this device is omitted, and all traffic is permitted to leave the network.

[1] The term "Demilitarized Zone" comes from military use, meaning a buffer area between two enemies.

Although these devices are shown separately in the previous diagram, many SBCs incorporate firewall functionality, in which case a single SBC device replaces the external and internal firewalls, and the DMZ looks like this. (Note that in this case the external firewall reverts to being a simple edge router.)

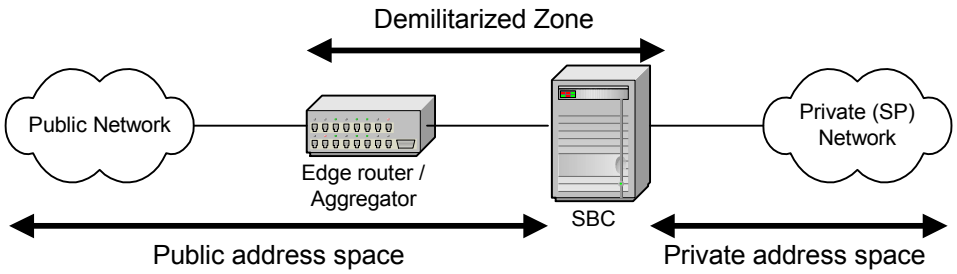

Figure 2: Reduced Demilitarized Zone

In both of these scenarios, the SBC is the only application-aware device in the DMZ. As such, when applications within the private network require IP traffic to traverse the DMZ, the SBC takes responsibility for ensuring that the other equipment within the DMZ allows that traffic through. It does this as follows.

- If the firewalls are separate devices, then either
    - the firewalls are statically configured by the network operator to permit all traffic addressed to the SBC (on both the public network side and the private network side), *or*
    - the SBC dynamically configures the firewalls across the network.
- If the firewalls are incorporated within the SBC, then the SBC programmatically alters their packet inspection tables to permit the appropriate application traffic.

See chapter 4, **DMZ processing**, for more information on how the SBC interacts with firewalls.

2.2.2 Dual-box SBC deployments

Dual-box SBCs (described in section 2.1, **Internal structure of an SBC**) may also be deployed entirely within the DMZ. However, the distributed nature of the dual-box device allows an alternative deployment, where the SBC-MEDIA devices are deployed in the DMZ and the SBC-SIG devices are deployed within the private network. The advantage of this deployment is that it allows a single SBC-SIG box to manage multiple SBC-MEDIA boxes. For more information on the advantages and disadvantages of dual-box SBCs, see section 3.10, **Comparison of single-box and dual-box model**.

In this scenario, all VoIP signaling traffic received from the public network is allowed through the DMZ and routed to SBC-SIG.

A schematic of a dual-box SBC deployment is shown below.

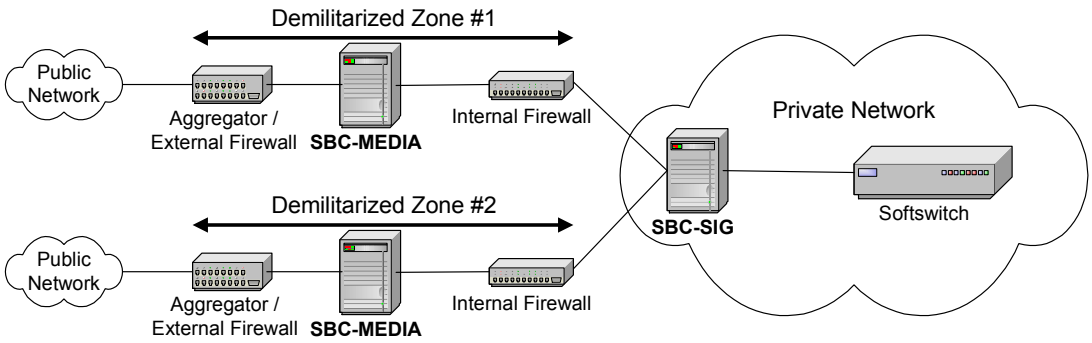

Figure 3: Dual-box SBC components in the DMZ

2.3 Applicable network scenarios for SBCs

SBCs can be deployed in five main network scenarios. In the first three of these, the SBC is part of the DMZ, while in the last two scenarios, it is in the core of a network.

- On the border between a provider and their customer (this can be thought of as providing a User Network Interface or UNI)
- On the border between two providers with a reciprocal agreement with respect to VoIP traffic (a Network-to-Network Interface or NNI) – this is where the majority of the SBC market is today
- Within a provider offering VPN services to its customers, to bridge calls across the customers' VPN sites
- In the core of a network as a means to overcome internal topology issues
- As a centralized codec transcoder

These scenarios are discussed in more detail in the following sections.

SBCs are typically deployed at the edge of service provider networks, but they can also be used at the edge of enterprise networks. To illustrate this, some of the enterprise networks shown below have SBCs deployed at their edge, and others have a more basic firewall without any SBC function. The latter type of customer network presents particular challenges to the provider's SBC, as special arrangements must be made to allow VoIP signaling and media to traverse the customer's basic firewall; see chapter 5, **Firewall and NAT traversal**, for more details on this.

In each of the following sections, a reduced DMZ (in which the firewall is contained within a single-box SBC) is shown to simplify the diagrams, although a separate firewall and/or a dual-box SBC could equivalently be shown in each scenario.

2.3.1 UNI scenario

SBCs are often used in the DMZ between a service provider and its customers. The SBC acts as an ALG replacement for the SP, and also provides Call Admission Control (CAC).

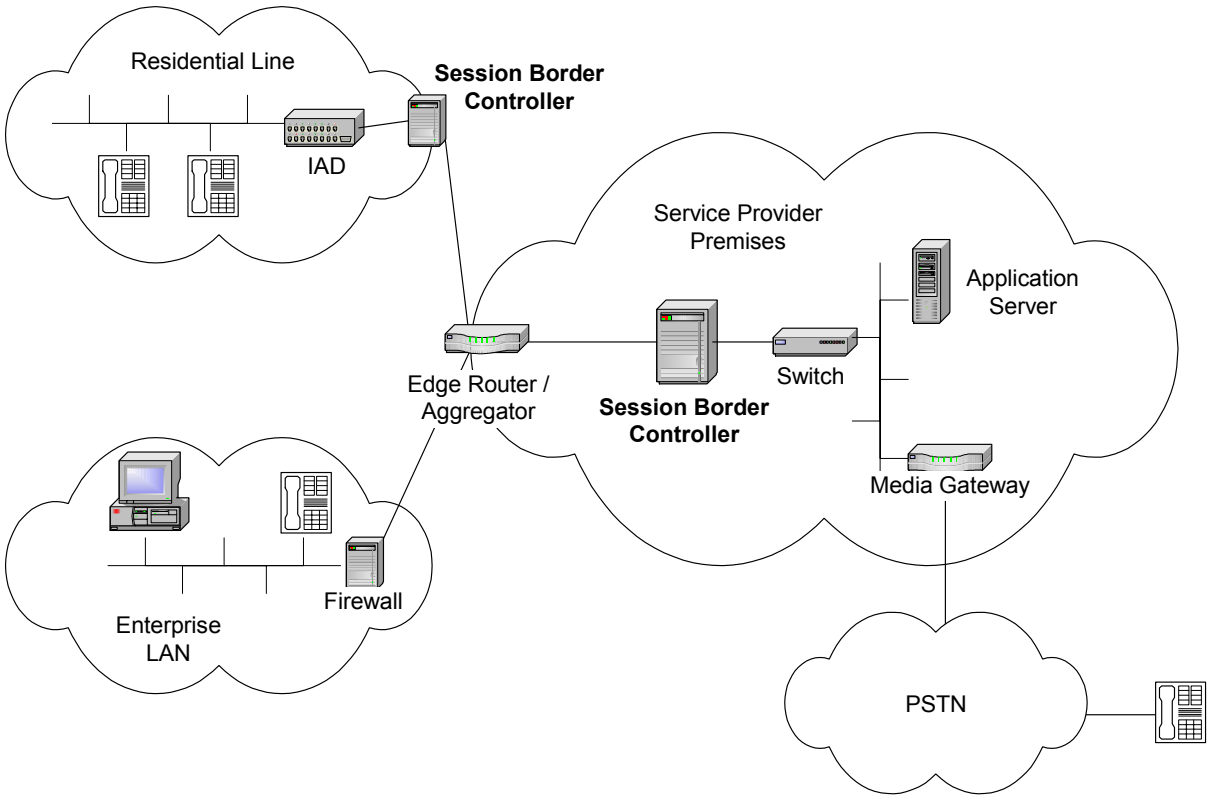

Figure 4: SBC on the provider/subscriber boundary

In this diagram, the service provider is providing VoIP services to its customers. VoIP calls to non-directly-connected networks or servers are internally routed to the PSTN using a media gateway.

For long-distance VoIP calls, the service provider may instead route calls to a second service provider. When doing this, it will want to keep the details hidden from its customer, to prevent the customer from approaching the second service provider directly for a cheaper price. To this end, the SBC removes any routing headers that would reveal the underlying SP from the signaling messages that it forwards to the customer networks.

The SBC performs Call Admission Control (CAC), to prevent the SP's servers from being overloaded by calls, and to ensure that nobody exceeds their Service Level Agreement (SLA). More information on CAC is given in chapter 6, **Call Admission Control**.

2.3.2 NNI scenario

SBCs can be used at the border between two service providers, where one or both have agreed to carry the other's traffic.

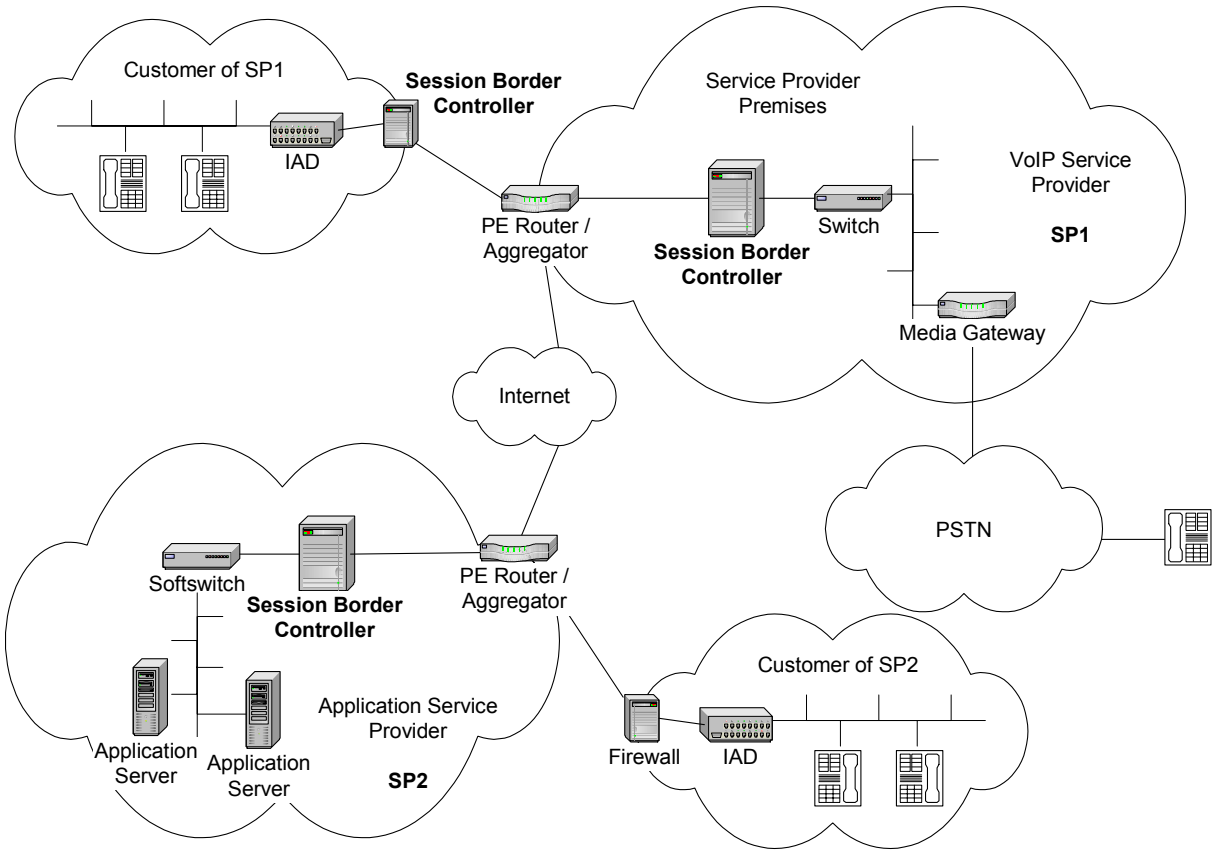

Figure 5: SBCs on a provider/provider boundary

In this example, SP1 provides VoIP calling services to its customer. SP2 provides a VoIP application to its customer, but has no connection of its own to the PSTN. These SPs have a reciprocal agreement, whereby SP2 agrees to provide the application service for SP1's customer, whilst SP1 agrees to offer a VoIP calling service to SP2's customer.

Each SP has an SBC guarding its border. The SBCs can do any or all of the following.

- Normal ALG, topology-hiding, and CAC functions.
- Policing of the reciprocal agreements between the SPs. For example, SP2 may restrict the number of calls that SP1 can make to its application servers per day.
- Call routing. The SBCs route the calls coming in to the network, farming calls out to the other SP where necessary, to prevent the servers and softswitches in the SP network core from having to deal with calls which are not intended for that SP.
- Gathering Call Detail Reports (CDRs), which are used to bill customers, and allow SPs to calculate who made how much use of the other's services during a given period.

2.3.3 VPN scenario

An interesting application of SBCs is to allow VoIP calls between customers in different IP Virtual Private Networks (VPNs) provided by the same SP. Please see the Data Connection white paper *VPN Technologies - A Comparison* for more information on VPNs.

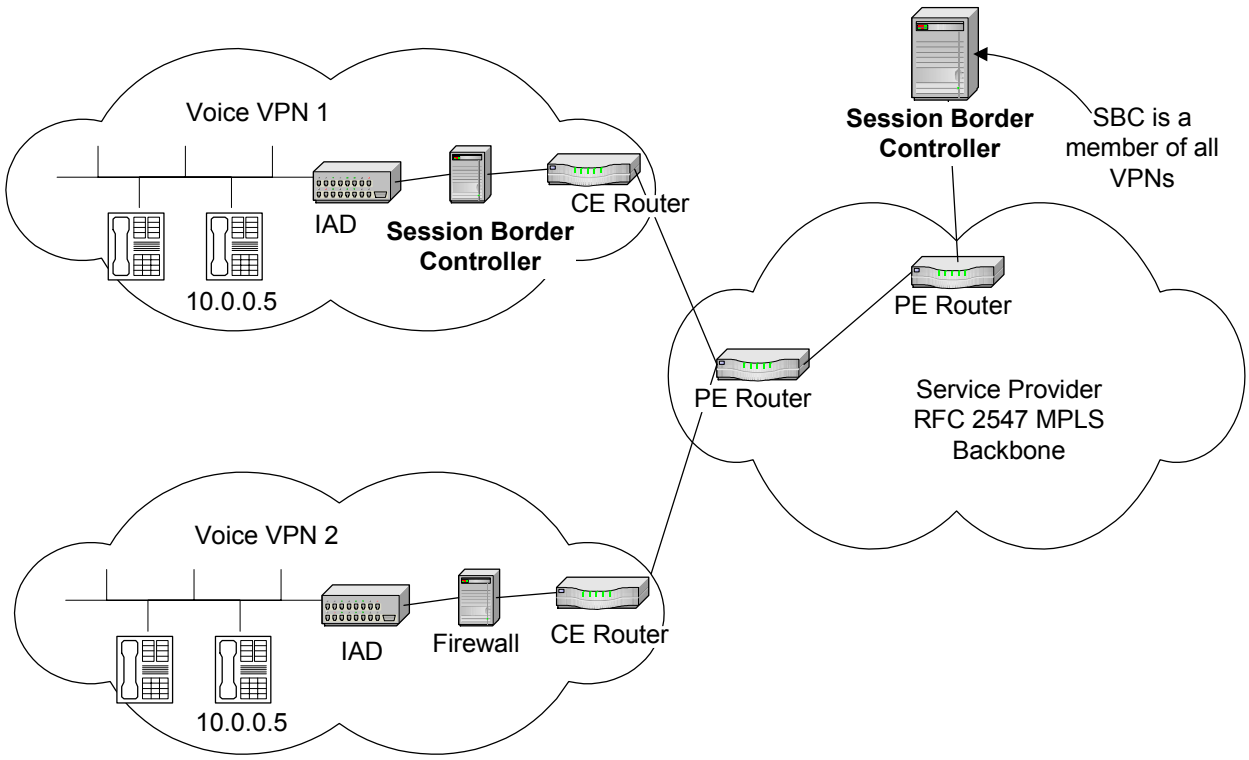

Figure 6: SBC bridging two VPN sites

In this scenario, the SP provides VPN services to two customers. Each customer has a number of VPN sites connected via MPLS tunnels in the SP's backbone network (only one site per customer is shown here; site 1 is using an SBC to protect itself, as described in section 2.3.1, **UNI scenario**). Endpoints at different sites within a given customer VPN can call each other directly, since they appear to be on the same LAN. However, it is not possible for an endpoint in customer network 1 to make a direct call to an endpoint in customer network 2, for the following reasons.

- Since the endpoints are in different VPNs, traffic cannot be routed out of one and into another using standard IP routing.
- Even if it could, the VPNs' address spaces may overlap, so the caller's and callee's addresses may clash (either with each other, or a third device).

These problems are solved by an SBC. The SBC in the provider network is configured to be a member of all VPNs (and hence is routable to in each VPN). It uses VLAN tags on the link between itself and the nearest PE router to emulate a different connection to the PE router for each VPN. It therefore appears to the PE router to be multiple logical devices.

The phones in each of the VPN sites configure the SBC to be their outbound proxy for external calls. The SBC acts as a B2BUA, and propagates the calls from one VPN into another. It also channels media between the VPNs. When it does this, it sorts out any ambiguities in the address space by rewriting IP and SIP (or H.323, or MGCP, or Megaco) headers in the signaling and media packets.

It is possible that the SBC and its nearest PE router could be combined in a single device, to eliminate a link from the network diagram above.

Although the above diagram shows a single SBC, it would be possible to deploy multiple SBCs, each of which managed a particular subset of VPNs. The SBCs would communicate amongst themselves to bridge VPN signaling and media across the provider backbone. This would distribute the SBC processing amongst the PE routers, and thereby improve scalability.

2.3.4 Solving internal topology issues

As well as being deployed at the network border, SBCs can be used within a single network to control the use of network resources by VoIP traffic. For example, the following diagram shows a "dumbbell" network topology, where two high-bandwidth sites are connected by a low-bandwidth backbone link.

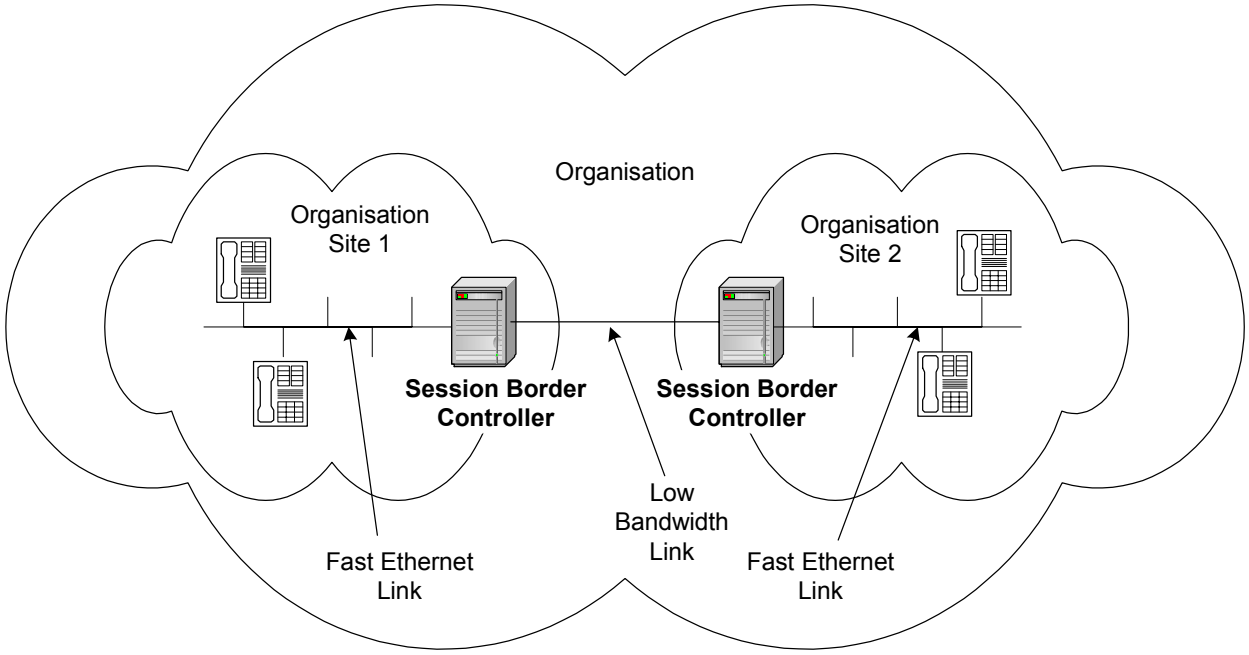

Figure 7: SBCs controlling network resource use by VoIP traffic

In this network, the SBCs perform CAC to prevent the low bandwidth link from becoming oversubscribed with voice traffic.

2.3.5 Centralized codec transcoding

Since many SBCs perform codec transcoding at the media bridge, they can be used by an organization as a centralized codec-transcoding server to overcome any issues of interoperability between equipment with different capabilities. The diagram below shows one such organization, which has IP phones using both the G.711 and G.729 codecs. In this network, the SBC may be used to bridge all intra-organizational calls that require codec transcoding.

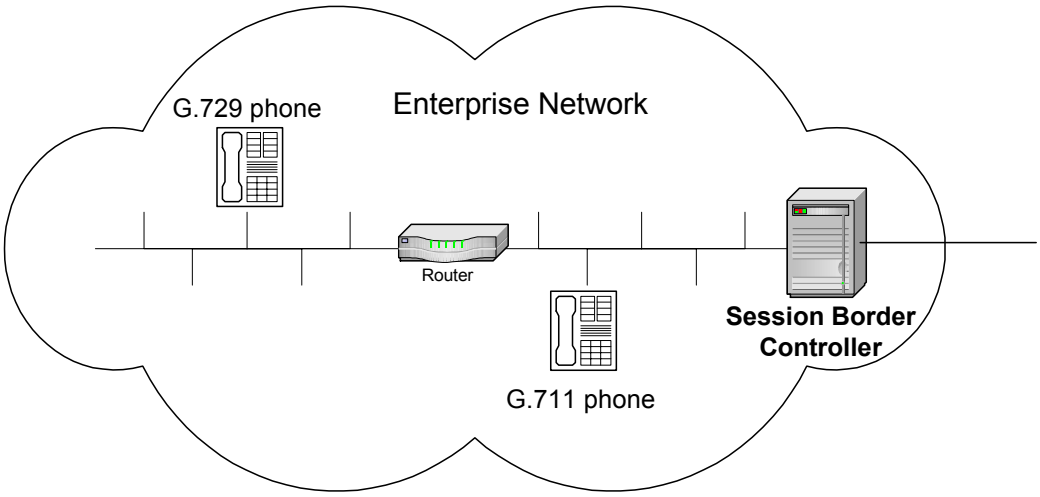

Figure 8: SBC as a centralized codec transcoding server

It is unlikely that a network operator would purchase an SBC purely to do codec transcoding, because cheaper devices can be purchased to perform this task. However, in the scenario above, the SBC is also fulfilling a border control function, and the operator of the network is leveraging the SBC's codec transcoding ability to avoid having to purchase a separate transcoder device.

3. SBC FUNCTION OVERVIEW

This chapter describes the breadth of function either provided by SBCs today or targeted for deployment in the near future. As noted earlier, not all SBCs on the market perform all of these functions. Indeed, there are nearly as many different function sets as there are SBC products!

- Sections 3.1-3.9 describe the various functions that current SBCs support.
- Section 3.10, **Comparison of single-box and dual-box model**, discusses the advantages and disadvantages of single-box versus dual-box SBCs.
- Section 3.11, **Configuration models**, describes how SBCs are configured.
- Finally, Section 3.12, **Feature matrix**, provides a feature matrix showing which functions are required for various SBC deployment scenarios.

3.1 DMZ processing

The SBC performs several functions to secure the service provider's network boundary.

- It acts as a Network Address Translator (NAT) for the SP.
- It may either act as a firewall, or cooperate with existing firewall devices in the DMZ. It may open pinholes in the firewall to allow VoIP signaling and media to pass through (or alternatively, the firewall may be statically configured to allow this).
- It performs a topology-hiding function to prevent customers or other service providers from learning details about how the SP's network is configured, or how calls being placed through the SP are routed. It does this by rewriting VoIP signaling messages that traverse it (for example, by deleting SIP Via headers).
- It eliminates bad VoIP signaling and media protocol at the network boundary.

This functionality is offered by all SBCs. See chapter 4, **DMZ processing**, for more details.

3.2 Firewall and NAT traversal

The SBC enables VoIP signaling and media to be received from and directed to a device behind a firewall and NAT (Network Address Translator) at the border of an adjacent network, without requiring the device or firewall to be upgraded. In brief, the SBC achieves this by rewriting the IP addresses and ports in the call signaling headers and the SDP blocks attached to these messages.

This functionality is offered by all SBCs. See chapter 5, **Firewall and NAT traversal**, for more details.

3.3 Call Admission Control (CAC) and DoS protection

The SBC-SIG device (which can be part of a single- or dual-box deployment, as described in section 2.1, **Internal structure of an SBC**) controls which calls may be signaled through the network, and gracefully rejects calls when necessary. This serves to protect the service provider, and specifically the softswitches in the SP's network, from the following.

- Various types of Denial-of-Service (DoS) attack that may be perpetrated on the network
- Massive spikes in the rate of incoming call setup requests, which may result from cataclysmic events, TV and radio phone-in competitions, and so on
- General congestion of the network

The SBC-SIG device achieves this by rate-limiting the calls that are set up through it, per subscriber and per group of subscribers, and also by rate-limiting or blacklisting calls to particular numbers.

CAC also allows SBC-SIG to guarantee and police SLAs, and to ensure that subscribers keep their call setup rates within limits that the backbone can handle.

- It tracks the bandwidth being consumed per subscriber in the access network. It rejects new calls from that subscriber if such calls would exceed the bandwidth limit set in their SLA.
- It monitors the total number of calls per subscriber, again to prevent the subscriber exceeding the limits set in their SLA.

SBC-SIG can also be configured to whitelist certain numbers (e.g. emergency services numbers), so as always to permit these calls through the network. It may pre-emptively discard active non-emergency calls to ensure that enough network resources can be dedicated to an emergency call being set up.

CAC is basic functionality offered by all SBCs (although, to date, the quality and flexibility of CAC offered by SBC devices is fairly crude, and leave lots of room for improvement by future devices). CAC function is primarily what differentiates SBCs from Application-Layer Gateways (ALGs).

See chapter 6, **Call Admission Control**, for more details on CAC.

3.4 Quality of Service (QoS)

The SBC reserves service provider network resources to handle calls being set up. There are a number of ways in which the SBC can achieve this.

- It can set DiffServ code points in the IP headers of media packets that are forwarded onto the network.

- It can set up MPLS LSPs for varying QoS levels, and use these LSPs to aggregate calls and carry them across the network. See the Multiservice Switching Forum (MSF) document MSF2003.105.00 for a network architecture capable of performing this function[2].

The SBC also reserves "signaling bandwidth" for emergency services calls (i.e. makes sure that the softswitch behind it always has capacity to handle a given number of emergency services calls), and prioritizes these calls appropriately (i.e. allows them through Call Admission Control, even when other calls are being rejected).

Some SBCs provide this function, but not all.

3.5 Media bridging

The SBC always routes the media for calls that it handles through an SBC-MEDIA device (which can be part of a single- or dual-box deployment, as described in section 2.1, **Internal structure of an SBC**). This allows the SBC to monitor bandwidth consumption, needed to police SLAs and prevent bandwidth theft. It also allows it to provide a lawful intercept function for both signaling and media, which all service providers are required to provide (as defined by CALEA for the US).

SBC-SIG devices rewrite the SDP in the signals that they forward, to ensure that the media of a call is routed through the appropriate SBC-MEDIA device. The SBC-MEDIA device is then responsible for bridging the media streams in both legs of the call.

Media bridging involves (at a minimum) rewriting IP headers in the media stream, and may also require transcoding of the media to enable interworking between devices or over links using different standards. This processor-intensive function occurs directly on the golden path (critical area for high performance) of the data plane, so must be done very efficiently. Therefore, SBC-MEDIA devices normally use dedicated network processors for media bridging, and programmable hardware-based packet filters that perform at least the IP-header-rewriting portion of this processing.

Nearly all SBCs provide this function, with some rare exceptions.

---

[2] There are some RFCs in this area, for example RFC 3313. See the References section for more details.

3.5.1 Voice over IP media bridging

VoIP media is transported by RTP, and so the SBC-MEDIA device must support RTP, and specifically, proxying of RTP media packets. This proxying can be done in software ("slow-path") by terminating the media stream at the IP layer, and then re-initiating it. Higher-performance devices will perform RTP proxy in network processor microcode ("fast-path") by rewriting each RTP packet's IP headers and forwarding it on without the packet leaving the network processor.

RTCP, the RTP control protocol, must be examined by the SBC (either in software or on the Network Processor) in order to allow it to monitor bandwidth-usage and quality characteristics of the VoIP media stream. RTP and RTCP are defined in RFC 3550.

Existing SBC-MEDIA devices support a wide variety of codecs. For example, G.711 and G.729 (as specified by the ITU-T) are two codecs commonly supported. Some also support transcoding for codec interworking.

3.5.2 Fax over IP media bridging

Most SBC-MEDIA devices now support two types of Fax over IP transmission, as defined in the following specifications.

- The ITU-T specification, T.38
- Cisco's proprietary Fax Relay protocol

Lack of support for Fax over IP was a barrier for uptake of many SBCs during their initial roll-out. It is somewhat different to supporting Voice over IP, because Fax over IP media is not carried by RTP.

In T.38, fax media is transported using either

- UDP-TL, which is a lightweight transport protocol for fax media that runs over UDP, *or*
- media transported directly over TCP, using TPKT headers to provide framing.

During call setup, devices decide which of these transport protocols to use, and also which addresses and ports to use for sending and receiving the data. SBCs that support T.38 must support one or both of these transport protocols.

T.38 media can be signaled by H.323, SIP, or H.248. In the latter two cases there are some relevant extensions to the SDP protocol, defined in T.38. These extensions allow applications to negotiate various fax-specific parameters, as well as the choice of TCP or UDP, and the choice of IP address and port for the media transport. As explained in section 4.3, **How VoIP media packets traverse the DMZ**, the SBC must be careful to indicate its external (NAT-translated) IP address and port in the SDP, as opposed to its internal address and port.

More information on Cisco Fax Relay should be obtained directly from Cisco.

3.5.3 Modem over IP media bridging

Similarly, most SBCs also support Modem over IP, as defined in the ITU-T specification, V.150.1. Again, this differs from voice traffic in that the modem media is not carried by RTP. Instead, it is carried by a tailor-made transport protocol (defined in V.150.1 specifically for the purpose) called SPRT (Simple Packet Relay Transport).

In addition, V.150.1 requires that the SBC support

- transport of DTMF, telephony tones and telephony signals in-band within RTP media (RFC 2833) in parallel with the modem media, *and*
- a tailor-made in-band signaling protocol defined within V.150.1 called State Signaling Event Protocol (SSEP), which operates over RTP in a very similar way to that specified in RFC 2833.

There have been discussions within the ITU-T and the TIA (Telecommunications Industry Association) on the topic of allowing interworking between V.150.1 and T.38 media streams (see for example TR-30.5/03-02-006). For this reason, it is desirable for an SBC supporting codec conversion also to support procedures for V.150.1-T.38 interworking.

3.6 Fault Tolerance

Most SBCs are fault-tolerant.

- In the majority of deployments, there will be redundant SBC devices for each access network. Fault-tolerant SBCs all provide a 1:1 redundancy mechanism, and most carriers want their SBCs to support a 1:N redundancy mechanism.
- Some SBCs dynamically replicate state information to each other, for example the addresses of pinholes on the exterior of customer's firewalls. This is to continue to allow VoIP traffic to traverse the customer's firewall prior to the firewall pinhole being refreshed (see chapter 5, **Firewall and NAT traversal**, for more details), thereby offering an uninterrupted service to the customer.
- The SBCs also need to implement some mechanism for deciding which device is the primary, i.e. which device owns the public IP address to which customers address their call signaling. This can either be via some private interface, or via a standard protocol such as VRRP (Virtual Router Redundancy Protocol).
- Some SBCs support Hot Software Upgrade (HSU) and Downgrade (HSD).

When the SBC fails over

- existing calls must be maintained, and the failover must be transparent to the user
- calls in progress may be dropped.

Carriers require that their SBCs provide an availability of 99.999% (five nines). Any carrier-class SBC device must meet this availability requirement.

3.7 Policy-based call routing

Some SBCs on the market provide a policy interface to allow calls that do not need to be processed by the SP's network elements to be routed intelligently to an exit from the network (e.g. to one of several candidate carriers for a long-distance call, depending on which carrier is cheaper for a particular call at a particular time of day).

This functionality is equivalent to that offered by a softswitch. Effectively, SBCs that do this are cut-down softswitches that live in the DMZ of the network and route certain calls away from the network as early as possible, so that the softswitches in the network core do not have to cope with the load.

3.7.1 Crankback and re-initiation of call setups

One particularly important aspect of policy-based call routing in the NNI scenario (described in section 2.3.2, **NNI scenario**) is crankback of call setup, which is a key feature for SBCs targeting the NNI. This feature involves correct handling of failed call attempts by a carrier's carrier ("supercarrier"), where the supercarrier network failed to handle the call setup.

Instead of propagating a call release back to the carrier, the supercarrier's SBC routes the call to another supercarrier network, with which it has a reciprocal agreement. This continues until the supercarrier finds another carrier that can handle the call (or until it runs out of alternatives, in which case the only choice is to release the call).

The supercarrier that reroutes the call may not make much of a profit on that call, but that is still better than releasing the call altogether, because a released call may result in the supercarrier's customer not routing any more calls to them for a period (or for ever after!), causing a loss of future business.

3.8 Signaling protocol interworking

SBCs all support SIP. They also commonly support H.323. On the boundary between service providers, it is possible for an SBC to provide interworking functions (such as between H.323 and SIP, or between variants of H.323) if the service providers use different signaling protocols.

SBCs also commonly support MGCP and/or Megaco, although typically no interworking is done involving these two protocols.

3.9 Call billing

The SBC-SIG device may track the progress of each call for the purposes of billing (although again, most softswitches will perform an equivalent function).

- It may produce Call Detail Reports (CDRs), documenting the details of each call, which are used for billing and for capacity planning (some SBC vendors offer software tools to post-process the CDRs produced by their machines, for accounting and network load visualization).

- It may run a session timer, and disconnect calls that have not terminated correctly.

- It may instruct the firewall to prevent any RTP traffic not associated with a valid, billed session from entering the network. In particular, it shuts down the flow of RTP traffic on a session as soon as the end of the session is signaled, to prevent service theft (which occurs when the callers continue to send media even after signaling the end of the call).

The SBC-MEDIA device monitors the bandwidth consumed by each call in the access link, to guard against bandwidth theft (which occurs when the callers silently upgrade their media from, say, voice-only to full video, to attempt to make a video call for the price of a voice call).

The default mode of operation of the SBC-MEDIA is not to forward RTP traffic, unless explicitly directed to do so by SBC-SIG, and then only to forward traffic of the type specified (e.g. within certain per-session bandwidth limits). This design principle prevents unauthorized media from gaining access to the carrier's network.

3.10 Comparison of single-box and dual-box model

As was noted in section 2.1, **Internal structure of an SBC**, both single-box model and dual-box model are valid architectures for SBCs, with different advantages and disadvantages for each. This section compares the two approaches, and discusses how they affect the function offered by the SBC.

3.10.1 Single-box SBCs

Most SBC vendors today offer a single-box solution. The majority of these devices are simple two-port devices, with one port dedicated to the external network, and the other dedicated to the internal network. Each port is used for both signaling and media. However, using the same port for signaling and media creates a high risk of signaling packets being dropped. Therefore, a more sophisticated SBC will offer multiple ports, one or more for signaling, and multiple separate ports for media. Ideally, the signaling and media ports will be controlled by separate processors.

The advantage that single-box SBCs hold over dual-box SBCs are that they are less complex, easier to make, and easier to configure and deploy. There does not need to be a communication protocol operating between the two devices, so a private programmatic API will suffice, which does not require any interoperability testing.

3.10.2 Dual-box SBCs

Some SBC vendors are now offering a dual-box SBC solution.

The dual-box model allows for greater scalability. A single SBC-SIG device can handle call control for several different access points, and control an SBC-MEDIA device at each access point. SBC-MEDIA devices are lower-cost than SBC-SIG devices, so this provides for a more affordable solution where high scaling is important.

However, this approach also diminishes the effectiveness with which the SBC-SIG can defend against Denial of Service (DoS) attacks and other spikes in network activity, because it exposes a single SBC-SIG to signaling traffic from multiple network access points. In order for an SBC to satisfy the basic requirement of protecting the core of the network from DoS attacks, each SBC-SIG device must only be exposed to a fraction of the total signaling load of the network, and so each should manage only a small number of SBC-MEDIA devices.

An advantage of this model is that it permits separate development of interoperable devices. Therefore, a company may specialize in one type of device or the other. There are standards being developed to allow interoperability between dual-box devices, in particular by the Multiservice Switching Forum (MSF). The MSF proposes that H.248 be used as the control protocol, although other candidates are in the offing, such as COPS-PR.

Another advantage is that, where several SBC-MEDIA devices serve a single network, the SBC-SIG device can load-balance new calls across them. This improves the scalability of the service that can be offered to individual customers.

Finally, the dual-box SBC architecture offers the opportunity of subsuming the SBC-SIG function into the SP's softswitches. This offers obvious advantages in terms of reducing the number of boxes needed in the network, reducing complexity and OAM overhead. However, it intensifies the disadvantages described for the dual-box architecture. The softswitch itself is now both responsible for the workload involved in managing the SBC-MEDIA devices (potentially tying up more expensive resources than necessary) and exposed directly to potential Denial of Service or overload incidents.

3.11 Configuration models

Existing SBC devices are configured using a variety of methods.

- Standard CLI (command line interface).
- MIB management via SNMPv3 (SNMPv2 and v1 may be desirable but are usually considered optional).
- Web-based configuration GUI using XML and SOAP.
- COPS (common open policy service).
- Configuration via a standard CORBA interface.

Of these, the most common is CLI via telnet, followed by SNMP.

3.12 Feature matrix

The matrix overleaf shows which of the features described in this chapter are typically supported by an SBC in each of the five deployment scenarios listed in section 2.3, **Applicable network scenarios for SBCs**. The level of feature support is classified as follows.

- "Yes" means that the feature must be supported by an SBC deployed in the corresponding scenario, or else the SBC will not be viable.
- "Maybe" means that the feature may be desirable in that scenario for some types of customer or network, but that an SBC device can still be viable there without this feature.
- "No" means that there is no requirement for a feature to be supported in that scenario.

A superscripted number in a cell refers to a footnote at the end of the table, where more information is provided.

Table 1: SBC feature matrix

Feature Description	Section	UNI	NNI	VPN	Topology	Transcoder
On-board firewall / firewall control	3.1	Yes	Yes	Yes	No	No
Bad protocol detection	3.1	Yes	Yes	Yes	Maybe[1]	Maybe[1]
Firewall/NAT traversal (aliases)	3.2	Yes	Maybe[2]	Yes	No	No
Firewall/NAT pinhole keepalive	3.2	Yes	No [3]	Yes	No	No
Call Admission Control	3.3	Yes	Yes	Yes	Yes	Maybe
DiffServ QoS	3.4	Yes	Yes	Yes	Yes	No
MPLS-based QoS	3.4	Maybe[4]	Maybe[4]	Maybe[4]	Maybe[4]	No
Voice over IP media bridge	3.5.1	Yes	Yes	Yes	Yes	Yes
Fax over IP media bridge	3.5.2	Yes	Yes	Yes	Yes	Yes
Modem over IP media bridge	3.5.3	Yes	Yes	Yes	Yes	Yes
Media transcoding	3.5.1	Maybe[5]	Maybe[5]	Maybe[5]	Maybe[5]	Yes
CALEA wiretapping	3.5	Maybe[6]	Maybe[6]	Maybe[6]	No	No
Redundant backup SBC	3.6	Yes	Yes	Yes	Yes	Yes
High availability (five nines)	3.6	Yes	Yes	Yes	Yes	Yes
HSU and HSD	3.6	Maybe	Maybe	Maybe	Maybe	Maybe
Call crankback and re-initiation	3.7.1	Maybe	Yes	Maybe	No	No
General policy-based call routing	3.7	Maybe	Maybe	Maybe	Maybe	No
SIP : H.323 interworking	3.8	Maybe	Maybe	Maybe	No[7]	No[7]
CDR generation for billing	3.9	Maybe	Maybe	Maybe	No[7]	No[7]
Dual-box SBC	3.10.2	Maybe[8]	Maybe[8]	Maybe[8]	Maybe[8]	Maybe[8]

(1) Bad protocol ideally will not be found in the core of the network, because all network devices are trusted. However, it may be desirable to check protocol to guard against malfunctioning devices.

(2) At the NNI, the SBC is presumably communicating with another SBC, which will support the procedures in section 3.1, **DMZ processing**. If the peer device supports these procedures, then the procedures outlined in section 3.2, **Firewall and NAT traversal**, are unnecessary. However, it may still be desirable to support the features described in section 3.2, because in back-level networks, the peering device may not be an SBC.

(3) At the NNI, pinholes for signaling are statically configured at start of day. REGISTER and RSIP messages are not exchanged.

(4) MPLS QoS is a feature that will be important in next-generation networks but is not currently being widely used. SBCs that do not support it ought to have it on their roadmaps.

(5) Not all SBCs transcode but it is a desirable feature.

(6) CALEA interception could be offered by SBCs, and would be desirable in networks where there is not already a softswitch providing equivalent function.

(7) Only required at network edge.

(8) Depends on scalability requirements.

4. DMZ PROCESSING

The DMZ is the "demilitarized zone" between two networks, as described in section 2.2, **The Demilitarized Zone**. This chapter provides more detail on the role played by SBCs in the DMZ.

- Section 4.1, **Devices in the DMZ**, shows single-box and dual-box SBCs in the DMZ and introduces firewalls and network address translation.
- Sections 4.2 and 4.3 explain how the SBC cooperates with the firewalls to ensure that VoIP signaling and media traverses the DMZ without compromising the security of the trusted network.
- Section 4.4, **Other DMZ processing**, describes topology hiding and bad protocol detection.

4.1 Devices in the DMZ

All SBCs fall into one of the following two categories.

- Those that do not perform firewall processing in the DMZ, but instead rely on an external and internal firewall.
- Those that do perform firewall processing (i.e. that have a firewall on-board).

If an SBC does not perform firewall processing, then the DMZ looks like this (see also Figure 1).

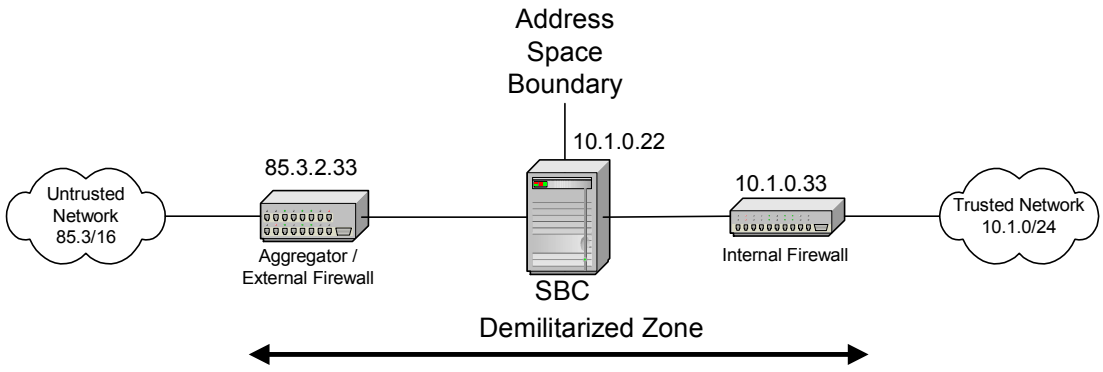

Figure 9: SBC and firewalls are distinct

If the SBC is decomposed into the dual-box model, then the SBC-MEDIA device resides in the position indicated above, and the SBC-SIG device resides in the network core. In this case, SBC-SIG controls the firewall to allow VoIP signaling and media to pass though (see section 5.2, **SBC pinhole solution**, for details).

If the SBC does perform firewall processing, then the same DMZ looks like this (see also Figure 2).

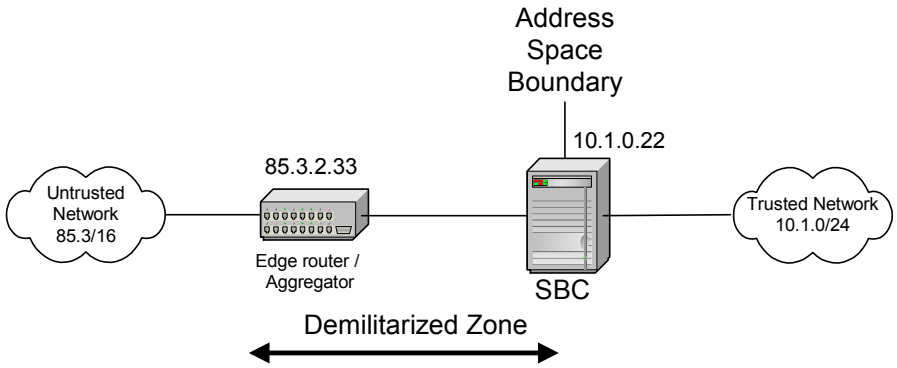

Figure 10: SBC does firewall processing

If the SBC is decomposed into the dual-box model, then SBC-MEDIA implements a firewall that screens out unwanted signaling and media packets.

All SBCs typically incorporate NAT (Network Address Translator) function. The remainder of this section describes the role of firewall devices in the DMZ, and the role of the NAT component of the SBC.

4.1.1 The firewall

Firewalls prevent unwanted traffic from entering, or leaving, a network by performing basic packet filtering. Note that firewalls filter packets purely by examining packet headers, and do not parse or understand the payload of the packets. Therefore, they do not filter out all types of unwanted traffic. For example, firewalls do not perform Call Admission Control – SBCs do that. However, firewalls are valuable because they efficiently filter out large categories of unwanted traffic, leaving application-aware devices such as SBCs with much less work to do.

The external firewall in Figure 9 filters packets from the external network, but allows all packets from the internal network to pass through unfiltered. The internal firewall filters packets from the internal network, but allows all packets from the external network to pass through unfiltered (since they have already passed the external firewall).

Firewalls by default do not accept packets from the network, but are configured with rules that allow them to select and accept certain packets. Therefore, packets are admitted to (or from) the network based on *explicit configuration*, and not on default configuration.

Firewalls are configured either

- by the network operator, using a human interface, *or*
- by trusted software, using an API.

There are no standards-defined APIs for configuring firewalls; however, the IETF's MIDCOM working group is evaluating suitable protocols for this task. SNMP, RSIP, Megaco, Diameter, and COPS are all being considered. In addition, the MSF have made some steps towards defining their own protocol for firewall control (MSF2003.113.00 – Draft IA for RTP Proxy / FW Control Protocol).

4.1.2 The NAT

SBCs typically incorporate NAT function. NATs separate a network into distinct address spaces. In Figure 9, the NAT component of the SBC separates the internal network address space 10.1.0/24 from the external network address space 85.3/16. A few addresses from the 85.3/16 domain are used to represent all machines within the 10.1.0/24 domain, as described below.

The NAT maintains a table of mappings from {external address, port} to {internal address, port} and vice versa[3]. The table is a dual-index table, so a particular mapping can be looked up given either the internal or external addressing information. The NAT uses this table to rewrite the headers of the IP packets that it forwards.

- On receiving an IP packet from the external network, the NAT looks in its table for the destination address and port of the packet (which will be an address from the external address space). If a mapping is found, then the destination address header in the IP packet is changed to contain the corresponding internal address and port from the table, and the packet is forwarded towards the internal network. If no mapping is found, the packet is discarded.
- On receiving an IP packet from the internal network, the NAT looks in its table for the source address and port of the packet (which will be an address from the internal address space). If a mapping is found, then the source address header in the IP packet is changed to contain the corresponding external address and port from the table, and the packet is forwarded towards the external network. If no mapping is found, then a new mapping is created: the NAT dynamically allocates a new external address and port from the external address space for the packet (and all future packets from this source address and port tuple).

[3] These mappings are also known as "pinholes."

Mappings in the table are created in one of two ways.

- By packets traversing the NAT from the internal network towards the external network, as described in the second bullet point above.

- By configuration, either from the network operator via a human interface, or programmatically from trusted software via an API.

4.2 How VoIP signaling packets traverse the DMZ

The NAT component of the SBC and the firewalls in the DMZ is configured (at start of day) as follows.

- The NAT is configured with a mapping that converts between the SBC's internal address (10.1.0.22 in Figure 9) and the port it uses for signaling, and some address and port taken from the external network's address space. This external address and port is used to identify the SBC in the public network. Packets sent from the external network and destined for the SBC are sent to this address and port.

- The external firewall is configured to permit IP packets whose destination address header contains the address and port that identify the SBC in the external network.

- The internal firewall is configured to permit IP packets whose destination address header contains the internal address of the SBC, and the port that it uses for signaling.

Note that this configuration could either come from human input, or from the SBC by programmatic API if the SBC and firewalls are collocated, or from the SBC by network protocol if the devices are separate.

This configuration allows all signaling packets addressed to the SBC to traverse the DMZ devices and reach the SBC, whether the packets originate from the internal or external network. In addition, it allows the SBC to send signaling messages towards either the internal or external networks.

This scheme relies on the fact that the external address and port that is used to identify the SBC on the public network for signaling is well-known to VoIP devices on the public network. Typically, this is achieved by using DNS records to associate this address with the SBC's hostname in the public network.

This scheme also relies on the SBC knowing its external IP address and port for signaling, because it must use these in the VoIP signaling headers that it sends in requests to the external network (as these fields are usually used to route the signaling response). This can be configured on the SBC.

4.3 How VoIP media packets traverse the DMZ

The situation with media packets is a little more complex than with signaling, because the media packets in a given call originate from, and are sent to, addresses and ports that are dynamically allocated by the RTP protocol when the call is established.

The SBC acts as a signaling Back-to-Back User Agent (B2BUA) and a media bridge, and so it terminates the media of a call on both the internal and external network sides. The ports that it uses to send and receive media on each side are allocated dynamically when the call is established.

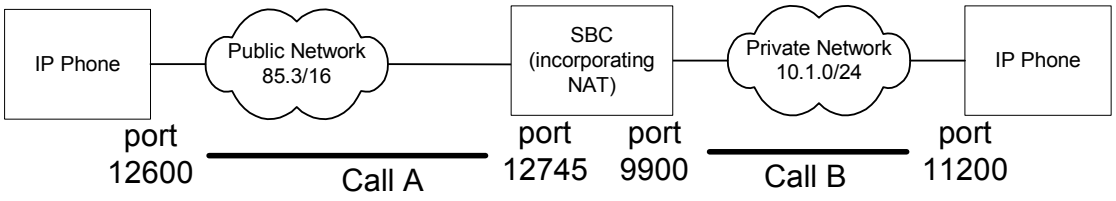

Figure 11: VoIP media packets traversing the DMZ

This causes a couple of problems.

- The internal firewall needs to be configured to permit IP traffic sent to port 9900 on the SBC. Since 9900 is a dynamically allocated number, this must be done automatically during call set-up (i.e. without human intervention).

- The external firewall needs to be configured to permit IP traffic sent to the SBC's external address and port 12745. Again, this must happen automatically during call set-up.

If the firewall and SBC are on the same device, then these problems are easily overcome by implementing a programmatic interface that allows the SBC to dynamically configure the firewall software.

If the firewalls are separate from the SBC, then either

- the SBC dynamically configures the firewall over the network (although, as noted above, there are no standards for this at present), *or*

- the firewalls must be configured to permit all traffic sent to any port on the SBC (or at least, any port in the range used by the RTP protocol).

4.4 Other DMZ processing

SBCs also perform other DMZ-related processing, as described in the following sections.

4.4.1 Topology hiding

VoIP signaling messages convey information that can allow the recipient to determine both the internal topology of a network, and the route taken by a call across that network (and possibly out the other side). For example, the Via headers in SIP signaling messages carry this kind of information.

It is often undesirable to expose this information to users outside a network. For example, if you are a service provider who uses a second service provider to act as a carrier for your calls, you do not want to expose the identity of the carrier SP to your customers in case they approach the carrier SP directly for a better price.

To solve this problem, SBCs can remove sensitive information by rewriting the VoIP headers in the signaling messages that they send across the network boundary. SBCs achieve this by acting as B2BUAs. They terminate the VoIP signaling that they receive from within the private network, and signal a new call towards the public network. Since this is a new call, it does not require any of the routing information from the previous call (for example, none of the SIP Via headers are carried over into the public network call leg).

4.4.2 Bad protocol detection

The SBC processes all signaling and media that enter or leave the network. It can therefore screen the network from bad protocol within signaling or media packets, discarding or sending negative responses to badly-formed packets. This has two advantages.

- It reduces the load on the VoIP servers within the network, which can be significant if someone is attempting to mount a DoS attack on the network by sending poorly-formed packets.

- It reduces the likelihood of the badly-formed messages causing a crash on a key piece of VoIP infrastructure within the network.

The amount of checking that an SBC does on signaling messages should be configurable. For example, it could be configured to check only those fields that it itself needs to process the message, or it could check all fields in the message, or anywhere in between.

5. FIREWALL AND NAT TRAVERSAL

SBCs also enable VoIP signaling and media to be received from and directed to a device behind a firewall/NAT at the border of an *adjacent* network, *without* needing the device or firewall to be upgraded. (See chapter 4, **DMZ processing**, for discussion of firewalls and NATs.)

- Section 5.1, **The VoIP firewall/NAT traversal problem**, explains why this is not possible without an SBC (or other device offering equivalent functionality).

- Section 5.2, **SBC pinhole solution**, describes the methods that SBCs use to solve this problem.

5.1 The VoIP firewall/NAT traversal problem

One of the early barriers to rolling out VoIP was that VoIP signaling protocols and media cannot traverse NATs and firewalls without significant help. The diagram below is a simplification of a VoIP service provider network and its customer.

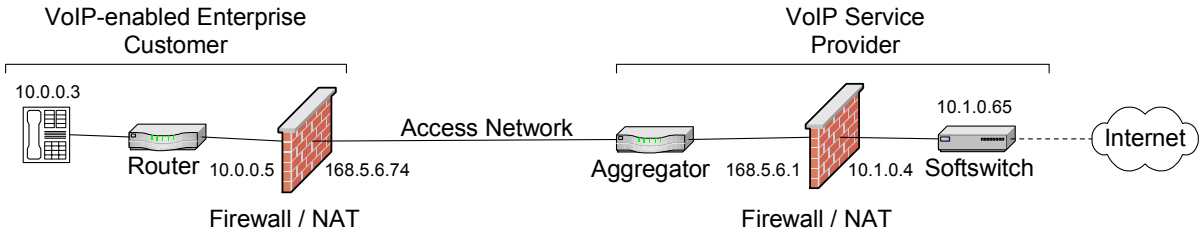

Figure 12: VoIP firewall/NAT traversal

The customer network has its own private address space (10.0.0/24), and is protected by a firewall/NAT device. The access link between the customer and provider is subnet 168.5.6/24. The service provider network itself has a private address space 10.1.0/24, and is again protected by a firewall/NAT.

There are two problems with this network arrangement with respect to VoIP. We illustrate these problems using SIP; identical problems exist for H.323 and MGCP/Megaco.

- The customer's firewall/NAT blocks inbound call signaling. The softswitch in the service provider cannot send a SIP INVITE request to the IP phone, because the IP phone is not addressable from the service provider. Even if the phone were addressable from the service provider (i.e. the customer's firewall were not also a NAT), the customer's firewall would still block the INVITE message.

Note that this problem does not usually apply to outbound call signaling, because the softswitch will typically have a well-known external IP address, which is statically mapped by the provider's NAT to the internal IP address (10.1.0.65 in the above diagram). The provider's firewall will be configured to accept unsolicited packets received on the external address on certain ports designated for VoIP signaling traffic (e.g. port 5060 for SIP).

- The customer's firewall/NAT blocks inbound call media. When the customer phone makes a call, it sends a SIP INVITE with an SDP body containing its IP address. When the call is established, the callee sends its media to that IP address. This causes problems because the IP address in the SDP is not routable from the Internet and, even if it were, the customer's firewall would still block traffic sent to it.

5.2 SBC pinhole solution

The key to solving this problem is the fact that the customer's NAT has to open pinholes in order to allow the IP phone to send signaling packets and media packets to the public network, and the customer's firewall has to allow these packets through. Inbound signaling and media from the public network can therefore be made to traverse the customer's firewall and NAT by directing them at the pinhole's address and port on the public network side of the customer's NAT.

The pinholes for signaling and media have different lifetimes.

- The signaling pinhole, once created, is reused for all call signaling.
- The media pinhole is created anew for each media stream, because the source and destination ports of the media stream are dynamically allocated per call.

The signaling pinhole is ideally created when the IP phone first comes online, and then kept open until the phone goes offline again. Media pinholes are created when the IP phone first sends a media packet on each established media session.

To solve the VoIP firewall/NAT traversal problem, the SBC replaces the provider's firewall/NAT, as shown below.

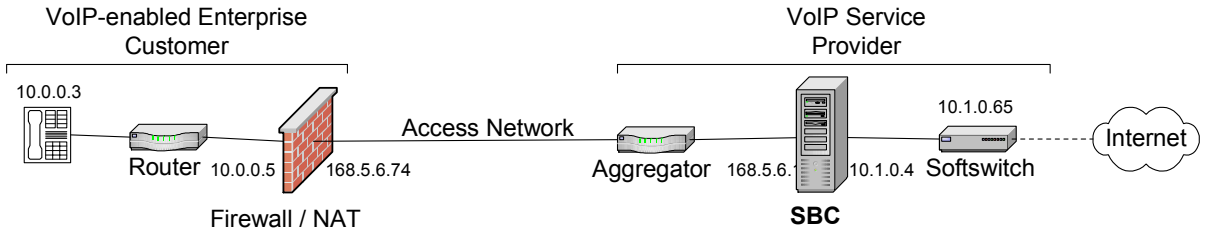

Figure 13: SBC solution to firewall/NAT traversal problem

5.2.1 The signaling pinhole

In the scenario illustrated in Figure 13, the IP phone is configured with the SBC's public IP address as its outbound SIP proxy. The IP phone thus sends a SIP REGISTER message to the SBC when the phone comes online. The SBC typically forwards the REGISTER to the softswitch (assuming it permits the REGISTER to gain access to the network). It also internally stores a mapping from the device name (e.g. ipphone@enterprise.com) to the "pinhole-traversal details." This is dependent on the underlying transport being used to carry the signaling.

- For signaling over UDP, the SBC maps the device name to the pinhole's public address and port, as learned from the IP header of the REGISTER datagram.

- For signaling over TCP (including TLS), the SBC maps the device name to the TCP connection that has been established between itself and the IP phone.

Thereafter, the SBC uses the mapping to direct inbound signaling messages either over the established TCP connection, or towards the public address and port of the pinhole for UDP.

In the following network diagram, the IP phone is replaced by a POTS phone and an IAD.

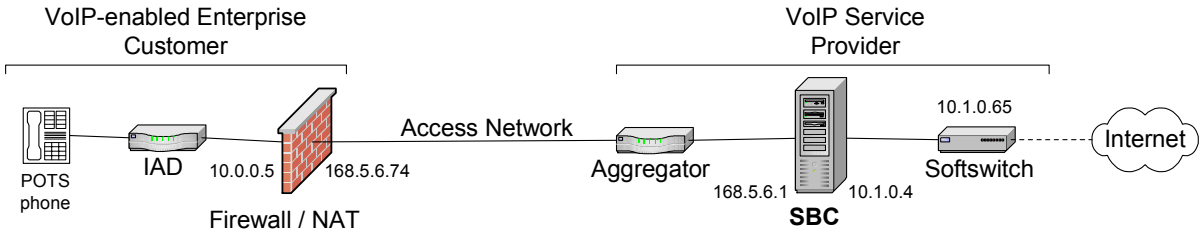

Figure 14: SBC solution with POTS phone and IAD

In this case, a SIP REGISTER is not sent. Instead, the IAD communicates with the SBC using MGCP (or Megaco). It sends an RSIP message to the SBC when it boots up, which opens the signaling pinhole in exactly the same way as the SIP REGISTER.

5.2.2 Keeping the signaling pinhole open

If a firewall sees no traffic on a pinhole for a period of time (typically a few minutes), it will time it out and close it, to minimize the security exposure. However, for the purposes of the SBC, once a signaling pinhole is opened, it must be kept open. How this is done depends on the signaling protocol in use.

- If the signaling protocol is SIP, then the SBC may periodically send OPTIONS requests to the SIP endpoint. Each time it does, the customer's firewall restarts the timer on its pinhole, which keeps it open. Note that the SIP endpoint does not even need to respond positively to the OPTIONS request.

- If the signaling protocol is MGCP, then the SBC periodically sends AUEP messages to the MGCP endpoint, which has the same effect on the firewall pinhole as the OPTIONS request in the SIP case. A similar mechanism applies for Megaco.

Further, if the signaling pinhole closes for any reason (for example, if the customer's firewall restarts), then it must be reopened without requiring the signaling endpoint to be restarted. This requires some cooperation on the part of the signaling endpoint, because the pinhole can only be opened from within the customer's network.

- If SIP is in use, then the IP phone periodically sends a REGISTER message to the SBC. The SBC rate-limits these REGISTER messages, and only forwards some of them on to the call agent. If the IP phone is configured to do this, then there is no need for the SBC to periodically send OPTIONS to it. Most (if not all) SIP endpoints can be configured to periodically re-register themselves. To ensure that REGISTERs are sent frequently enough, the SBE may modify the expiration time of each REGISTER to a low value on the response. A frequency of one REGISTER every 30 seconds is typically required.

- If MGCP is in use, then this is not so easily achieved. It is not possible for an MGCP to send repeated RSIP requests to maintain the binding, as this will cause the call agent to drop any calls that are currently in progress for that endpoint.

An alternative to the MGCP solution described above is to make use of an MGCP keepalive package. The X-NET/ping event is from a proprietary package supported on the Cisco ATA (Analogue Telephone Adaptor). The ATA sends a NTFY command with observed event "O: X-NET/ping" to the SBC if no signaling messages are exchanged with it for a configured period. Alternatively, the NAT/ka event is from the NAT Package defined in Internet-Draft draft-aoun-mgcp-nat-package (currently expired draft). Such packages are supported by several devices.

Again, if the MGCP endpoint supports one of these packages, there is no need for the SBC to send periodic AUEP messages to it.

5.2.3 The media pinhole

The media does not travel through the signaling pinhole; another pinhole needs to be opened for it. Pinholes must be opened from within the customer's network. Therefore, the media pinhole opens when the customer sends the first RTP packet in their media stream.

When the SBC receives this media packet, it correlates it with an existing call. If it fails to do so, it drops the packet (as this may be an attempt to steal service from the provider). If it succeeds, then it creates a mapping between the call and the external pinhole address and port in the firewall through which the RTP packet was received (which it learns from the RTP packet's IP header). It then redirects any inbound RTP packets belonging to this call to the external pinhole address and port.

The media pinhole stays open for as long as media flows on the call, and then closes at the end of the call. No keepalive mechanism is necessary, provided that media continues to flow.

One factor that can complicate this scenario is the fact that many calls contain some element of one-way media flows.

For example, calls to the PSTN typically result in early media flowing from the media gateway back to the caller, to play a ringing tone on the caller's handset. This is known as backwards in-band alerting. This media flow is one-way. If the caller is behind a firewall, then (since no media flows outwards from the caller during the early media phase of the call) the media pinhole will not open, and the early media will not reach the caller.

To circumvent this problem, SBCs rewrite the SDP describing the early media to make it appear to the caller that the call is two-way (by changing a=sendonly to a=sendrecv in the SDP). An SBC terminates the media flow from the caller, and maintains a one-way media flow with the media gateway.

A similar situation is where the caller makes a call to a one-way media server, for example the speaking clock. In this case, no media will flow outwards from the caller during the call itself and so the media pinhole will not open, unless the SBC rewrites the SDP as explained in the paragraph above.

5.2.4 Example SIP flow

This message flow illustrates the signaling pinhole being opened in the customer firewall, when SIP is the signaling protocol in use. The IP addresses used in this flow (and subsequent flows in this section) match the network diagram in Figure 13.

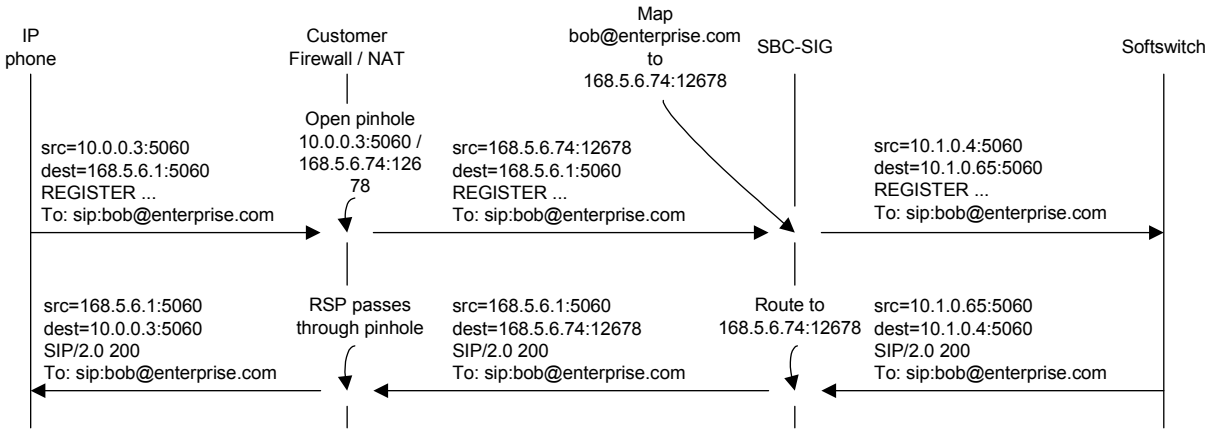

Figure 15: Opening signaling pinhole in customer firewall

- The customer phone sends a SIP REGISTER. The customer's firewall opens a pinhole for it and rewrites the IP header as it forwards the datagram.
- The SBC also acts as a firewall and NAT. The firewall component of the SBC is configured to allow through unsolicited packets sent to the SBC's public address and arriving on port 5060.
- The SBC is a SIP Back-to-Back User Agent (B2BUA). On receiving the REGISTER, it (i) remembers that the response must be sent to the firewall's external address (rather than the address found in the Via header of the REGISTER), and (ii) sends a new REGISTER on to the softswitch, taking care to rewrite the headers to ensure that it remains on the media path. Note that in step (ii) the REGISTER will acquire new Via and Contact headers, to ensure that the SBC remains on the path of subsequent signaling requests sent to this endpoint.
- When the SBC receives a response to its REGISTER, it sends a response to the original REGISTER, directing it to the correct address and port on the firewall. Providing the pinhole in the firewall is still open, this will be allowed through the firewall to the customer phone. The IP header in the response is rewritten as it traverses the customer's firewall.

Once the pinhole is opened, inbound calls to the IP phone can be directed through the pinhole, as shown below (the ACK is not shown, but traverses the firewalls identically to the INVITE).

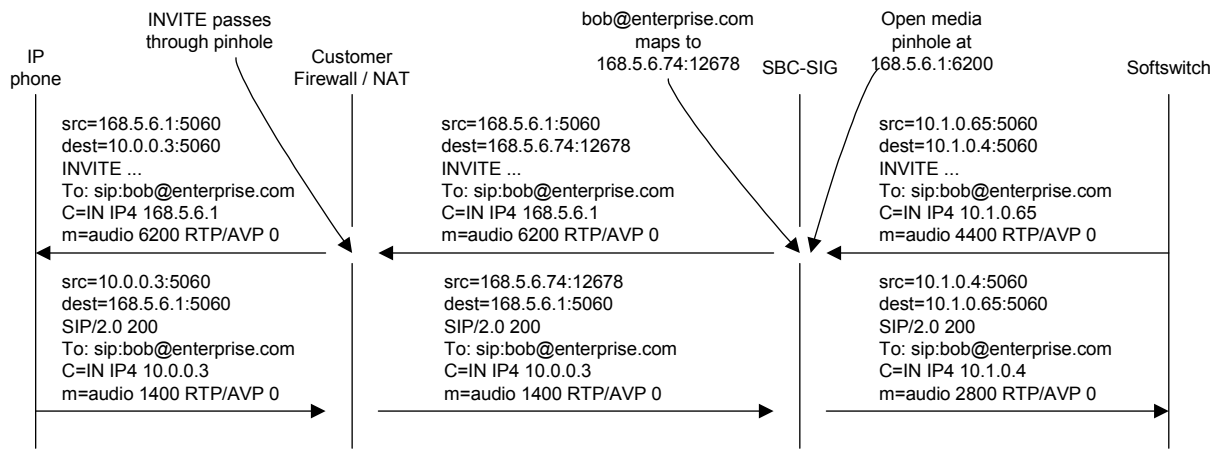

Figure 16: Inbound calls to IP phone directed through pinhole

- The softswitch detects an incoming call for bob@enterprise.com, and generates a SIP INVITE, which it sends to Bob's registered Contact address (i.e. the SBC).
- The SBC recognizes that bob@enterprise.com maps to 168.5.6.74:12678 (learned earlier from the INVITE exchange), and sends an INVITE to that address and port.
- The SBC's RTP stack allocates port 6200 on the SBC's external IP address for it to receive media from the SIP endpoint on. The SBC opens a pinhole in its own firewall (which is usually resident on the SBC device itself) to allow the media for the call to be received on this address and port.
- The INVITE traverses the customer's firewall. The response to the INVITE travels over the reverse path through the network. The ACK (not shown) follows the same path as the INVITE.

Once the inbound call has been signaled (as shown previously), media begins to flow.

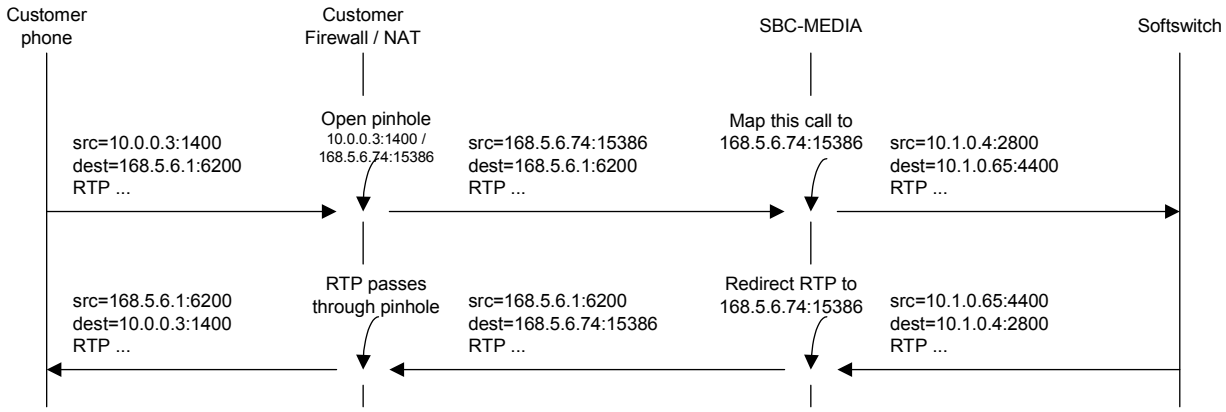

Figure 17: Media flowing through pinhole

- When the customer's IP phone sends its first RTP packet, the customer firewall opens a pinhole to allow that packet out, and rewrites its IP headers as it passes through. This RTP packet also passes through the hole in the provider's firewall that the SBC opened when the call was signaled. The RTP packet received by the SBC has a source address and port (dynamically allocated by the firewall on opening the customer pinhole) that differs from the source address and port in the SDP. When sending RTP packets back to the customer, the SBC must send them to this address and port, rather than using the address and port from the SDP.

- When RTP packets flow back from the softswitch, the SBC directs them at the external IP address and port of the media pinhole in the customer's firewall, as learned in the previous step.

6. CALL ADMISSION CONTROL

This chapter describes the various types of Call Admission Control (CAC) that an SBC may perform.

- Section 6.16.1, **DoS and DDoS attack prevention**, describes various types of Denial of Service (DoS) attack, and the methods SBCs use to prevent them.
- Section 6.2, **Reacting to network congestion**, describes how SBCs respond to excessive incoming requests.
- Section 6.3, **Policing SLAs**, describes how SBCs ensure that customers do not exceed the limits of their Service Level Agreement (SLA).
- Section 6.4, **Preventing theft of service and bandwidth**, describes the methods SBCs use to prevent customers stealing service and bandwidth.
- Section 6.5, **Emergency services calls**, describes special handling for these calls.

6.1 DoS and DDoS attack prevention

VoIP networks need to be protected against the following types of Denial of Service attack.

- A continuous stream of signaling messages coming from, or destined to, random (and invalid) endpoints
- A continuous stream of signaling messages with valid endpoint names, but incorrect or absent authentication headers
- A continuous stream of badly-formed signaling messages
- A massive, continuous stream of media packets being directed through an established call, in an attempt to consume all the bandwidth on the link behind the aggregating router

An attacker may try to launch a DoS attack by sending a stream of packets in one of the above categories from a single IP address. Alternatively, the attacker may launch a Distributed Denial-of-Service (DDoS) attack by spoofing the source IP address of each packet, making the packets appear to come from many different addresses.

SBC-SIG can identify badly-formed messages. If it is configured with the subscribers' identities, it can also identify invalid endpoint names or authentication headers.

As well as DoS attacks, there are other situations in which a large volume of signaling traffic may be directed towards the network. For example, numbers for phone-ins on TV and radio shows may receive massive spikes of calls during a show. While this does not constitute a deliberate DoS attack, it must still be deflected by the SBC to prevent the softswitches in the network core from being overloaded.

There are a number of ways in which an SBC can prevent DoS attacks and other signaling spikes from adversely affecting the network. The SBC can protect the network using any or all of the following methods.

- As per section 4.4.2, **Bad protocol detection**, the SBC parses incoming messages, rejects incorrectly formatted ones (if possible, and if configured to do so), and shuts down signaling traffic from the source IP address of the offending packets if several are received consecutively.

- The SBC inspects the responses sent from the softswitch. If it sees a (configurable) number of consecutive "404 Not Found" responses, or responses indicating that multiple requests to the same address were unauthorized, it shuts down all subsequent traffic coming from the source IP address of the corresponding requests.

- The administrator can define groups of subscribers for the SBC to monitor (where the group may contain one, many, or all subscribers). The SBC monitors the rate at which calls are made from or to these groups, and rejects calls that are signaled faster than the maximum rate, as measured over some period of time. The SBC also enforces a minimum interval that must pass in between individual call attempts per group, and rejects calls that are signaled too soon after the previous call. If the problem persists, then all traffic to or from these groups is shut down. This method can also be applied to consecutive registrations received from a group.

- The SBC also monitors the total number of calls that are made from or to a group, and caps the total number of concurrent calls permitted to the group. Calls are rejected if the call limit would be exceeded by accepting a new call, and all traffic from or to the group is shut down if the problem persists.

- The SBC monitors the bandwidth consumed by media from an individual call. If the rate of arrival of media packets for a given call exceeds a certain limit, then all media for that call is discarded and the call is terminated.

- The external firewall (or the SBC, if the external firewall is configured to permit all RTP packets sent to the SBC) does not permit media packets to enter the network that do not belong to an existing call (and that have not been spoofed to appear to belong to an existing call).

When traffic from or to a given address is "shut down," it means that signaling messages from that IP address are discarded. However, since the attacker may have spoofed the source IP address to match that of a legitimate caller, the shutdown applied to that address cannot be absolute and eternal.

- The shut-down runs for a configurable time period and then ceases, provided that rogue messages are not still being received (alternatively, it may be cancelled at any time by explicit management action).

- Emergency services calls from the address that has been shut down are still usually permitted.

The above techniques are not completely effective against DDoS attacks, where the source IP address of the attacker's packets changes at random (although the rules that filter streams of packets sent *to* a given IP address or endpoint/DN will defend against certain classes of DDoS attack).

The only way for the SBC to defend reliably against DDoS attacks is to place a global rate-limit on the number of registrations, new calls, and other types of signaling message that it will permit either from or to any subscriber. If this rate is exceeded, then the SBC can be configured either to respond negatively to, or to shut down all traffic received from outside the network for a period. If the SBC is also an aggregating edge router, then it can monitor the rate of signaling on each individual access link, and shut down all traffic received just from a given access link, which carries the advantage of continuing to provide a service over the other links.

When dealing with a DDoS attack, a much greater class of traffic must be shut down (either "all traffic" or "all traffic from a given access link"). It is tempting to think that this would be more efficiently accomplished by the external firewall than by the SBC. However, this cannot easily be made to work, because the same caveats about shutting down calls apply to the DDoS case. For example, most (if not all) firewalls would be unable to identify reliably and permit signaling messages pertaining to emergency services calls.

6.1.1 Summary of limiting options

This section summarizes the limiting options described in the previous section, without reference to the specific DoS attacks or other events that are being defended against.

- **Rate limits of VoIP signaling traffic**. The network operator can configure the SBC with rate limits to control the rate at which all VoIP signaling messages is allowed to enter or leave the network. Separate limits can be applied on the following basis (with separate limits configurable for messages sent and messages received).
    - A limit per session
    - A limit per port
    - A limit per subscriber
    - A limit per group of subscribers
    - A limit per network (i.e. rate-limit the messages originating from or being sent to an address within a configured prefix block)
    - A limit per customer VPN
    - A global limit for all VoIP signaling messages

- **Rate limit per message type**. The network operator can configure the SBC with separate rate limits for each type of VoIP signaling message (for example, one rate limit for SIP INVITEs, another for SIP REGISTERs, etc.). Separate limits can be applied on the following basis (with separate limits configurable for messages sent and messages received).
    - A limit per message type per session
    - A limit per message type per port
    - A limit per message type per subscriber
    - A limit per message type per group of subscribers
    - A limit per message type per network
    - A limit per message type per customer VPN
    - A global limit per message type
- **Session limits**. The network operator can configure the SBC with maximum concurrent session limits, to prevent total network resource use from exceeding the network's capacity, and to prevent subscribers from exceeding their SLAs. The following limits can be applied per port, per subscriber, per subscriber group, per network, per VPN and globally.
    - A maximum number of concurrent sessions
    - A maximum number of concurrent sessions, per session codec type
    - A maximum for the bandwidth consumed by all concurrent sessions (calculated by adding together the nominal bandwidth requirements of the codecs used for each session)

6.2 Reacting to network congestion

The SBC monitors the total rate of all incoming requests, from and to any endpoint. A maximum permissible rate is configured on the SBC, and it rejects additional requests that would exceed the rate limit (while giving priority to in-call requests and emergency services calls, as before).

The SBC also monitors responses sent by the softswitch. If the softswitch responds negatively to any request with an error code that indicates that it is congested, then the SBC throttles back the rate at which future incoming requests are permitted for a time period.

6.3 Policing SLAs

Service Level Agreements (SLAs) are contracts between service providers and their customers that guarantee the customer a certain level of service, and also restrict the maximum service that can be offered.

SBCs can police the SLA by ensuring that each customer does not exceed the limits set out in their SLA. The service provider may place limits on the total number of calls that can be made concurrently by each customer (either a global limit, or a per-call-type limit). There may also be a limit on the total bandwidth consumed by call media. The processing on the SBC may therefore involve

- counting the number of concurrent calls made by each customer
- counting the number of calls of each type (voice, video, etc.) being made by each customer
- counting the total bandwidth consumed by all calls from a customer.

If a new call is signaled from the customer to the SBC, and accepting it would cause the SLA to be violated, then the SBC rejects the call. In addition, the SBC may scan the SDP of each new call before accepting it, and remove any codecs that either the SLA or the provider does not allow the customer to use.

6.4 Preventing theft of service and bandwidth

Customers can steal service and bandwidth from service providers by

- signaling the end of the call, but keeping the media stream open
- signaling a voice call, but then upgrading the call media to video without signaling a change to the call.

The SBC can prevent this from happening. SBC-MEDIA by default does not forward RTP, and will only do so during an authorized call, which prevents the media stream from being kept open after call termination is signaled. In addition, if a firewall control protocol is in use, SBC-SIG closes the media pinhole in the SP's firewall when each call is terminated, thus preventing unauthorized media from entering the network after call termination.

SBC-MEDIA monitors the rate and type of media packets in a call, and communicates this information to SBC-SIG for billing purposes. It sends a message to SBC-SIG requesting that the call be terminated (and stops forwarding media) if the media parameters are changed without authorization from SBC-SIG.

6.5 Emergency services calls

As mentioned previously, the SBC places emergency services calls on a permanent “whitelist” of calls that will always be permitted, except possibly from addresses or interfaces that have been shut down because of suspected DoS attacks.

If placing an emergency services call would cause an SLA to be violated, or the total number of calls or volume of media being forwarded to exceed some predefined limit, then the SBC terminates one or more non-emergency calls to allow the emergency call to take precedence.

7. CONCLUSION

Session Border Controllers play different roles and offer different functionality in a variety of scenarios. SBCs may be used in UNI or NNI scenarios, to connect VPNs, to solve internal topology issues or for centralized codec transcoding. All SBCs offer DMZ processing, firewall/NAT traversal, and some degree of Call Admission Control. Other functions that SBCs may offer include media bridging, policy-based call routing, signaling protocol interworking, and call billing.

In each case, the issues that the SBC tries to resolve are caused by boundaries of trust, administration, and policy. Indeed, the requirements for the SBC come from the clash between these boundaries and the peer-to-peer model envisioned by VoIP protocol designers. In this way, the development of the SBC looks poised to play an essential role in full-scale commercial deployment of VoIP and integration into the existing PSTN.

8. ABOUT DATA CONNECTION

Data Connection Limited (DCL) is the leading independent developer and supplier of Unicast and Multicast IP Routing, MPLS, SIP, MGCP/Megaco, ATM, and SNA portable products, as well as Web Conferencing, Voicemail, Unified Messaging, and Directory applications. The company's MetaSwitch division supplies the industry's leading Class 5 Softswitch, with both integrated (Compact Softswitch) and distributed architecture options including call agent, feature server, media gateway and signaling gateway functionality. The MetaSwitch Class 5 Softswitch has been widely deployed among incumbent and competitive local exchange carriers, as well as operators of broadband wireless, cable and fiber networks.

Net profits have exceeded 20% of revenue every year since the company was founded in 1981. Customers include BT, Cisco, IBM, Lucent, Microsoft and Verizon. With over 300 employees, Data Connection is headquartered in London, UK with US offices in Alameda, CA, Boxborough, MA, Dallas, TX, and Reston, VA. For more information please see www.dataconnection.com.

Data Connection's SBC solution

Data Connection's Session Border Controller solution, DC-SBC, supports the functionality required for deployment into all of the existing scenarios, and is designed to be simple to extend as new requirements emerge.

Our products share a common architecture to facilitate their integration. This architecture allows components to be combined in a variety of ways to create routers with specific functionality, and allows seamless integration with Data Connection's suite of routing and switching products.

All of the Data Connection protocol implementations are built with scalability, distribution across multiple processors and fault tolerance architected in from the beginning. We have developed extremely consistent development processes that result in on-time delivery of highly robust and efficient software. This is backed up by an exceptionally responsive and expert support service, staffed by engineers with direct experience of developing the products.

About the author

Jon Hardwick is the senior architect for Data Connection's SBC solution. He plays a leading role in product architecture and standards-based development in Data Connection's Networking Protocols Group. He has six years' experience in the field of communications protocols, having worked on IP Multicast, BGP, MPLS, ATM, SIP, Megaco, MGCP, SNA, and APPN.

9. GLOSSARY

This section provides a brief explanation of some of the terminology used in this white paper.

1:1 redundancy	Mechanism to provide redundancy by ensuring that for each piece of hardware there is a backup that can take over non-disruptively.
1:n redundancy	Mechanism to provide redundancy by ensuring that for each n identical pieces of hardware, there is a single backup that can take over non-disruptively in the case of a single failure.
ALG	Application Layer Gateway. A bridge for traffic between two networks. It has knowledge of, and operates at the level of, the application generating the traffic.
B2BUA	Back-to-Back User Agent. This is a piece of software which links together the signaling flows for two legs of a call, providing a bridge between them with local termination for each leg.
CAC	Connection/Call Admission Control. This is the set of actions taken by a network during the call set-up phase in order to determine whether a connection request should be accepted or rejected.
CALEA	Communications Assistance for Law Enforcement Act. Passed in 1994, CALEA requires telecommunications carriers in the United States to modify their equipment, facilities, and services to ensure that they are able to comply with authorized electronic surveillance.
CDR	Call Detail Record/Report. The billing record for a phone call.
CE	See PE.
Codec	Compressor/decompressor. A codec is any technology for compressing and decompressing data, typically voice or video.
COPS-PR	Common Open Policy Service. This is an IETF standard, supplying network switches and hubs with policy rules to help maintain Quality of Service.
CORBA	Common Object Request Broker Architecture. CORBA is an architecture and specification for creating, distributing, and managing distributed program objects in a network.
DiffServ	Differentiated Services. A mechanism for marking IP traffic with different priorities.

DoS	Denial of Service. A malicious attempt to overload a piece of hardware in some way.
DMZ	Demilitarized Zone. This is a small subnetwork that sits between a trusted private network, such as a corporate LAN, and an untrusted public network, such as the public Internet.
Firewall	A system designed to protect a computer network from unauthorized access, especially via the Internet.
H.248	H.248 (or Megaco) is a VoIP signaling protocol, usually used between a dumb (or slave) device and a clever (or master) controller. It is similar in functionality (if not syntax) to MGCP. It is defined in ITU standard H.248 and RFC 3525.
H.323	A protocol used for signaling for VoIP.
HSD	Hot Software Downgrade.
HSU	Hot Software Upgrade.
IAD	Integrated Access Device. An IAD is a one-box DSL voice and data solution equipment typically installed at the customer's site.
LSP	Label Switched Path. The name for a single traffic flow in MPLS.
Megaco	See H.248.
MGCP	Media Gateway Control Protocol. This is a VoIP signaling protocol, usually used between a dumb (or slave) device and a clever (or master) controller. It is similar in functionality (if not syntax) to H.248/Megaco. It is defined in RFC 2705.
MPLS	MultiProtocol Label Switching. Protocols used for network traffic flow shaping and management.
NAT	Network Address Translator. This is a program or piece of hardware that converts an IP address from a private address to a public address in real time. It allows multiple users to share a single public IP address.
NNI	Network to Network Interface. The border between two carriers.
OAM	Operation, Administration and Maintenance.
PE	Provider Edge. This is a piece of equipment situated at the edge of a service provider’s network, typically contrasted with Customer Edge (CE) equipment.

POTS	Plain Old Telephone Service. This is the standard telephone service that most homes use. It is also referred to as the PSTN (qv).
PSTN	Public Switched Telephone Network. The world's collection of interconnected voice-oriented public telephone networks.
RSIP	Realm-Specific Internet Protocol. An IP address translation technique that is an alternative to NAT. RSIP lets an enterprise safeguard many private Internet addresses behind a single public Internet address.
RTCP	Real Time Control Protocol. A protocol to carry information on the performance of RTP traffic.
RTP	Real Time Protocol. This is the dominant protocol for carrying VoIP media data. It is defined in RFC 3550.
SDP	Session Description Protocol. A syntax for describing key features of media streams, including codecs, IP addresses/ports, bit rates, and other information. It is defined in RFC 2327.
SIP	Session Initiation Protocol. A protocol used for signaling for VoIP.
SLA	Service Level Agreement. The contract between a service provider and their customer which specifies the level of service that will be provided.
SNMP	Simple Network Management Protocol. An Internet standard that defines methods for remotely managing active network components such as hubs, routers, and bridges.
SOAP	Simple Object Access Protocol. A way for a Web server to call a procedure on another, physically separate Web server, and get back a machine-readable result in standard XML format.
SP	Service Provider.
TCP	Transmission Control Protocol. The connection-oriented, transport level protocol used in the TCP/IP suite of communications protocols.
TLS	Transport Layer Security. A protocol that provides data integrity and privacy on a communications link over the Internet. It allows client-server applications to communicate and is designed to prevent eavesdropping, message forgery and interference.
Transcoder	Technology for converting between different codecs (qv).

UDP	User Datagram Protocol. This is a transport layer protocol in the TCP/IP protocol suite used in the Internet. UDP is used at the two ends of a data transfer. It does not establish a connection or provide reliable data transfer like TCP.
UNI	User to Network Interface. The border between a service provider and their customer.
VoIP	Voice over IP.
VPN	Virtual Private Network.

10. REFERENCES

The following documents provide more information on the topics covered by this white paper. All RFCs and current Internet-Drafts may be downloaded from the IETF web site at http://www.ietf.org/.

Note that all Internet-Drafts are work in progress and may be subject to change or may be withdrawn without notice.

10.1 Media

MSF2003.113.00	Draft IA for RTP proxy / FW Control Protocol
RFC 2833	RTP Payload for DTMF Digits, Telephony Tones and Telephony Signals
RFC 3550	RTP: A Transport Protocol for Real-Time Applications
RFC 3685	Real Time Control Protocol (RTCP) attribute in Session Description Protocol (SDP)
T.38	Procedures for real-time fax communication over IP networks
V.150.1	Procedures for real-time modem communication over IP networks
TR-30.5/03-02-006	Proposed modification to T.38 to support interworking with V.150.1 (V.MoIP)
G.711	Pulse code modulation (PCM) of voice frequencies
G.723/G.726	40, 32, 24, 16 kbit/s adaptive differential pulse code modulation (ADPCM)
G. 729	Coding of speech at 8 kbit/s using conjugate-structure algebraic-code-excited linear-prediction (CS-ACELP)
RFC 3313	Private Session Initiation Protocol (SIP) Extensions for Media Authorization

10.2 Signaling

H.248	Media Gateway Control (Megaco)
H.323	Packet-based multimedia communications systems
H.450	Supplementary Services for H.323
MSF2003.105.00	Quality of Service for next generation VoIP networks framework
RFC 2705	Media Gateway Control Protocol (MGCP) Version 1.0
RFC 3525	Gateway Control Protocol Version 1.0
RFC 3261	SIP: Session Initiation Protocol
RFC 3263	Locating SIP Servers
draft-ietf-sip-session-timer	SIP Session Timer
RFC 3966	The tel URI for Telephone Numbers
RFC 3924	Cisco Architecture for Lawful Intercept in IP Networks
RFC 2327	Session Description Protocol