

SIP, Security and Session Border Controllers



Simply Better Connected

SIP, Security and Session Border Controllers

Executive Summary

Rolling out a public SIP service brings with it several security issues. Both users and Service Providers must understand these issues, but the burden is with the Service Provider to offer a secure and reliable service to the user. This means they must show that the service does not compromise existing security and that the user's public presence is protected and managed. The Service Providers must also protect their own networks from outside attacks and service abuse.

This White Paper examines the security issues faced by users and looks at how the Service Provider can overcome these through the deployment of session border controllers in the access network and in the core.

What problems do users face?

The Firewall - a Brick Wall

When you fire up your browser and surf the net from your desk at work, have you ever wondered what is going on in the network? You take it for granted that you have the freedom to access web sites, yet at the same time you expect to be kept safe from malicious attacks on your machine. The important fact that helps to keep you safe is that you requested the information from the Web. This means that the devices keeping your network safe will allow your outgoing connection to the Web server and accept the reply returning from it. This enables the Firewall to reject incoming messages that it was not expecting.

Making a call over IP means setting up two separate connections. One connection is for signalling messages, the other for the media. Naturally, these two connections are related. The clients making and receiving the call use information in the signalling connection to learn how to make the media connection. This works well when the clients are in the same network. However, the devices such as NATs and Firewalls that separate networks are unaware that these connections are related. This means an invitation in the signalling connection to send voice to a particular address will be invisible to a Firewall. The Firewall will therefore reject the incoming voice.

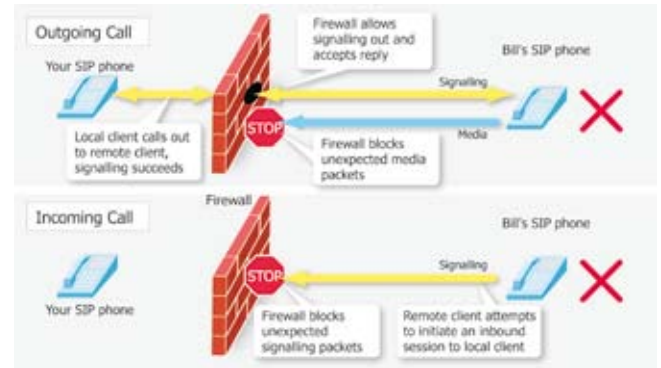


Figure 1 – The Firewall

Now think of making a phone call to Bill who is outside your network. First, your phone sends out an invitation to Bill. This goes out through the Firewall and finds Bill's IP phone. When Bill answers the call, his phone sends an acknowledgement back to your phone, which will reach you as the Firewall accepts the reply on the same port. However, when the phones send media, any Firewalls in the path will probably pass the outgoing media, but reject the incoming media. The result is the call appears to connect, but the voice path is broken.

Now think of receiving a phone call. You don't know you are going to receive the call and you don't know where it is coming from. The Firewall sees an unexpected incoming message from an unknown source so it blocks it. So, the call fails to reach you.

How do I make myself visible?

A Public Address – A Public Liability

Firstly, you need a public IP address so you can be called; this must be advertised so that you can be found. This takes care of the signalling. Secondly, you need a second IP address to exchange media. This is the address you will invite callers to send their media to.

Because your client is sitting behind a Firewall, these addresses have to be on the public side of the Firewall. The Firewall must link these addresses back to your client on the inside. This means leaving two holes in the Firewall permanently linked to your client. Now, far from being an anonymous Web user, you have advertised your presence to the world and invited them in. Unfortunately, this is like advertising your real address and leaving the front door open.

Make your Firewall Work

The Firewall is there to protect you and your network so you should make the best use of it. You can achieve this by making your IP phone call work more like your browser. That means ensuring all signalling and media connections are started outwards – even incoming calls. This may sound impossible, but that is what session border controllers do.

How can a Session Border Controller Help?

The session border controller sits within the public network and is the point to which you send your signalling. When you start your client, it registers with a server in the public network. This registration message is sent via the session border controller, which modifies the message and registers one of its own addresses with the server. Your public address is now on the session border controller. So, this is like having a PO box number; you can be reached but your real address is only known to the post office.

Now you can change your Firewall to allow the signalling to be started as an outgoing connection. You can also restrict the destination of this connection to be only the session border controller.

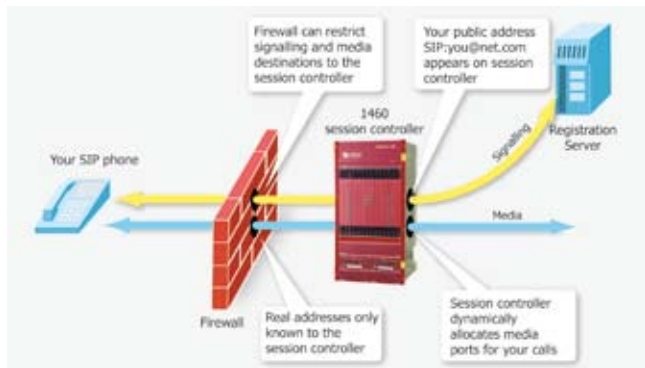


Figure 2 – Firewall Working with Session Border Controller

The session border controller can also protect the public address of the media. When you receive an invitation, the signalling travels via the session border controller. The session border controller modifies the invitation substituting one of its own addresses as the media address. This means you will send your media to the session border controller. The reply is also modified in a similar way so your caller also sends media to the session border controller. Now both clients send media to the session border controller. The session border controller learns the source addresses when the media emerges from the Firewall. This means you can change your Firewall to allow media ports to be dynamically allocated as outgoing connections. You can also restrict the destination of these connections to be only the session border controller.

Session Border Controller Benefits

A session border controller in the public network allows you to create stricter Firewall rules:

- All signalling and media connections can be dynamically opened
- All signalling and media connections are started as outbound connections
- The Firewall can restrict connections to just the session border controller
- The session border controller improves security by:
 - Hiding your real address
 - Dynamically allocating media ports
 - Policing signalling connection
 - Policing media connection

How can Service Providers Help?

If you are a Service Provider, you are probably well aware of these problems which are faced by customers who need a public SIP service. The need to connect and to create a public presence has to be weighed against the security implications. The Service Provider is ideally placed to address these issues. Deploying a carrier-class session border controller in the access network overcomes a number of issues:

- It solves the traversal problem for all NATs from customer to core
- It provides a secure connection to the user
- It works with existing Customer Premise Equipment
- It controls which customer uses which service

NATs are often used in access networks to create more IP addresses. This means that solving NAT traversal just for the customer premise is not a complete answer. The Service Provider must solve the traversal problem for multiple NATs. Placing a session border controller between the access and core networks achieves this.

With a session border controller in place, the Service Provider can offer services to any customer. The customer does not have to replace any of their equipment. The session border controller offers a secure, managed public presence for each user. The customer's Firewall can limit outgoing connections to the session border controller.

Offering a secure public presence to the customer enhances any service offering. The Service Provider is minimising the visibility and hence exposure of the customer's network. The service becomes a security enhancement rather than a security problem.

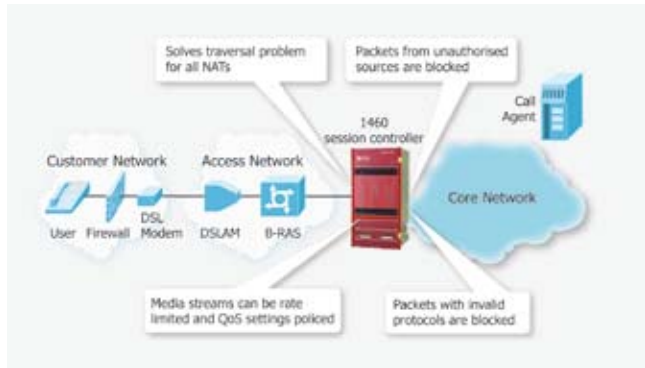


Figure 3 – Session Border Controller in the Access Network

The Newport Networks' 1460 session border controller is designed for just such duty. All signalling and media passing through the 1460 is policed. This further enhances the security of the connection:

- ♦ Packets from unauthorised sources are blocked
- ♦ Packets carrying invalid protocols can be blocked
- ♦ Media streams can be rate limited
- ♦ QoS settings can be policed

This means that the 1460 is effective in limiting the effects of port scanning. Denial of Service attacks against clients can be blocked or restricted. This protects both the access network from overload, as well as the customer. The 1460 session border controller can offer Service Providers a secure, reliable way of connecting to all customers, even those behind Firewalls, without compromising security.

Security between Networks

Securing the customer connection is not the only precaution that a Service Provider must take. Connections at peering points must also be secured. A recent Yankee Group report cited 'Network topology hiding' as one of the key drivers behind deploying session border controllers. We refer to a session border controller sited at a peering point as a Core Session Border Controller. A Core Session Border Controller performs several duties:

- ♦ It hides the real addresses of your customers from peer networks
- ♦ It hides the details of your internal network from peer networks
- ♦ It polices the connection to other Service Providers
- ♦ It can remark QoS settings between Service Providers
- ♦ It provides detailed call information

Protect your Customers

A Core Session Border Controller acts as a proxy for all users in a network. The home network's DNS ensures that all off-network calls are routed to the Core Session Border Controller. It does this by giving the address of the Core Session Border Controller as the address of any remote Call Agent. The session border controller creates new signalling and media addresses that are sent to the remote network. The called party in the remote network sees the session border controller as the source of the call. All signalling and media will be returned via the session border controller. In this way, the called party has no visibility of the user's real address.

Incoming calls are also routed via the home network's Core Session Border Controller. The remote network's DNS supplies the address of the Core Session Border Controller as the home network's Call Agent. Therefore, the Core Session Border Controller receives all calls coming into the home network. It presents its own addresses in the reply for both signalling and media.

This architecture prevents visibility of the user's real network address in the remote network. The Core Session Border Controller can prevent scanning and DOS attack at the peering point. At Newport Networks, we believe that the carrier-class 1460 session border controller is ideal for deployment in these demanding locations. Designed for high availability, it offers Service Providers a reliable method of securely interconnecting multimedia networks.

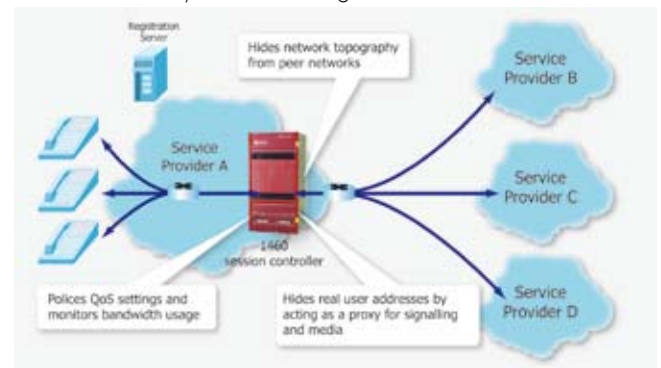


Figure 4 – Core Session Border Controller

Protect your Network

In addition to hiding the address of the user, the Core Session Border Controller hides the internal network details. The Core Session Border Controller acts as an end-point for the two legs of the SIP call: one to the home network and one to the remote network. This means that details of routing in one leg are not passed to the other. There is a clean separation between the networks. Therefore, the only information visible in the remote network is that of its own network.

Police the Border

A Core Session Border Controller connects all inter-network multimedia traffic. The Newport Networks' 1460 session border controller polices traffic flow-by-flow as it enters and leaves the network. Calls established using SIP carry an identifier of the media type. The 1460 measures the actual flow against expected flow for the requested media type. This can prevent service theft, i.e. requesting a low bandwidth connection and using high bandwidth media. If excessive data rates are seen, corrective action is taken. For example, it can dump excess traffic, it can generate an alarm or it can create punitive charging records.

The 1460 session border controller can also check and, if necessary, remark QoS bits. This can be done generically for each network, or specifically for each session. This prevents users from manipulating the quality settings of their call to get a better service than they are paying for. This also enables carriers to enforce IP-IP interconnect agreements to deliver 'end-to-end' SLAs.

Conclusion

Session border controllers enhance the security of multimedia networks both in the access network and in the core. In the access network, they hide a user's real address and provide a managed public address. This public address can be policed, minimising the opportunities for scanning and DOS attacks. Session border controllers permit access to clients behind Firewalls whilst maintaining the Firewalls effectiveness. In the core, session border controllers protect both the users and the network. They hide network topology and users' real addresses. They can also police bandwidth and QoS abuse.

The Newport Networks' 1460 session border controller is a carrier-class solution suitable for these applications. Built to provide 'five 9s' availability, it is designed for demanding deployments at both network peering points and in the access network. Service Providers who plan to roll out voice and multimedia services over IP must consider security as an integral part of the service. Deploying the right infrastructure lays the foundation stones upon which all successful future services will be built.