

Monitoring and policing the borders between your organization and the outside Internet is no easy task. Unfortunately, by its very design, Voice over IP (VoIP) using the Session Initiation Protocol (SIP) makes the task even harder. Luckily, SIP-aware devices like firewalls and intrusion detection systems/intrusion prevention systems (IDS/IPS) can dramatically improve border security and help organizations responsibly deploy Internet-wide VoIP.

A firewall is a device that stands astride a user's pathway to the Internet that denies entry and exit to all but acceptable forms of traffic. A firewall opens and closes pathways depending on the real-time needs of user applications. Very few applications communicate their intentions directly to firewalls. Instead, firewalls observe application traffic and indirectly derive the application's intention and open or close the appropriate holes to pass traffic.

For protocols that use well-known ports, such as HTTP, this is an easy task that can be handled by static rule sets. But for protocols that assign ports dynamically, particularly those based on UDP, the firewall has to inspect application-level protocol information. This adds considerable complexity to firewall software and can consume considerable firewall memory and CPU cycles.

Firewall rules typically presuppose the only services available to the "unprotected" side are those provided by a short and static list of well known servers (such as web and email servers) on well known IP addresses and ports.

SIP, however, turns every internal phone into a device that may potentially be called from outside phones. This means that SIP-aware firewalls may have to evaluate dynamic rules whenever an incoming call arrives. Or the SIP-aware firewall may delegate that responsibility to a specialized SIP device (which may be physically embedded into the firewall box) that may, in turn, impose a VoIP-specific security policy.

IDS and IPS devices can add a layer of defense by detecting attacks transiting the network. They operate on a combination of pattern analysis and protocol awareness. An IDS is a passive detection system, intended to alert administrators that an attacker is trying to exploit a vulnerability. The IPS takes this a step further by blocking malicious traffic, or even actively taking steps to stop the attack.

When used in conjunction with SIP, IDS/IPS systems monitor the SIP messages with a much greater degree of protocol awareness, detecting and potentially blocking malformed or out-of-sequence messages. They also scan the SIP message stream for usage patterns that are likely attacks, such as sequentially dialing extensions in rapid succession. If left undetected, these attacks can exploit known bugs in SIP servers and phones, or simply overload the VoIP system enough to make calls unreliable or impossible.

SIP + RTP/RTCP

Let's back up for a moment and recognize that when we talk about SIP-based VoIP, we are really talking about a bundle of separate protocols. SIP is but one protocol in this bundle. SIP is a call-setup (signaling) protocol. SIP does not itself carry voice or video media. SIP depends on the real-time protocol/real-time control protocol (RTP/RTCP) protocol (and variations for security, such as SRTP) to move the actual media content.

SIP uses UDP and TCP port 5060, and is thus easy for firewalls to intercept. However, SIP carries several types of IP address and port information that the firewall must examine. And because of the many ways that SIP can format this data, examination can be difficult.

RTP/RTCP is not anchored to any fixed port; instead, SIP dynamically assigns a port to each RTP/RTCP stream. Until the firewall parses SIP transactions, it does not know which packets contain the RTP/RTCP streams carrying a call's media. In addition, the firewall must continue to monitor the SIP activity in order to detect SIP RE-INVITES and know when the call has terminated.

Tracking SIP to discover UDP ports to be used by RTP/RTCP

A SIP-capable firewall must follow and comprehend each SIP exchange and extract from each exchange the RTP/RTCP port information that is being negotiated. The nature of SIP requires that the firewall continue to monitor the SIP exchanges through the lifetime of the call. The firewall must continue to monitor the SIP conversation in case the call is modified by subsequent SIP activity (as might occur with a call that is forwarded to voice-mail or processed by an interactive-voice-response [IVR] menu system).

Unfortunately SIP does not carry this information in a nice easy-to-digest format. Instead the media stream and port information is spread over multiple SIP packets that occur near the start of the call and may be revised in subsequent SIP packets that may occur during the life of the call.

Moreover, the way that SIP encodes data further complicates the life of a firewall. SIP is a text protocol that has blended techniques used in e-mail (SMTP) and HTTP and then added enough alternative encoding options and variations to choke a camel.

As if this were not enough, some SIP devices have adopted mechanisms, such as STUN, to detect the presence of NATs and pre-modify SIP packet contents in anticipation of NAT processing.

SIP is flexible – it can run over the reliable, sequenced, connection-oriented TCP or it can run over the unreliable, non-sequenced, connectionless UDP. Most SIP devices of today use UDP. This means that a firewall must anticipate that the SIP packets for a call might be lost, duplicated, reordered, or delayed. But even when SIP is run over TCP, the firewall must partially replicate TCP processing in order to reform the packets into a sequenced stream.

SIP is overgrown with complexity; the IETF has not yet pruned SIP to be a thing of streamlined simplicity. Firewalls, because they have to take whatever is thrown at them, must be very robust and engineered to accept and handle all of SIP's complexity.

All in all, a SIP-aware firewall must be a creature rather well endowed with intelligence, capacity, agility, and reliability.

Attacks Within the Stream

One of VoIP's greatest attractions is its ability to make and receive calls across the Internet, totally bypassing the phone company. To do this, though, an organization needs to allow inbound access from anywhere on the Internet to initiate inbound SIP calls. SIP-aware firewalls are intelligent enough to track the basics of the signaling protocol, and open only the appropriate ports. However, these firewalls lack the deep protocol knowledge to differentiate between a legitimate SIP INVITE and a malformed one intended to crash the SIP server or the handset on your desk. An IDS or IPS can detect these anomalies and potentially defend against them.

Other attacks can be much more subtle. An unscrupulous vendor may send an annoying advertisement to every voicemail box in your system. While every SIP and RTP packet in these messages may be completely standards-compliant, the pattern of their use (calling every extension in sequence) could be identified as spam over Internet telephony (SPIT) and blocked.

IDS/IPS devices provide another layer of protection on your border to catch attacks that occur within an otherwise admissible session. They work in concert with firewalls to strengthen perimeter defenses.

The effect of encryption

Several security protocols provide encryption, message integrity, and message authentication to protect data from being observed or altered as it flows from source to destination. However, this protection also creates an issue with transparency. Firewalls, IDS/IPS devices, and other systems cannot inspect data that's been encrypted. A security device that works "blind" is of little use. Thus, placement of encryption devices is a key network design concern.

Summary

SIP and RTP/RTCP form a very flexible system. Classical firewalls are only one tool among many for imposing security policy on SIP-based VoIP networks. Vendors may bundle these tools into a single physical box.

The enforcement tools for SIP and RTP/RTCP are complex and varied. But tools are just that: tools. They must enforce a policy. If experience with telephone dial plans is any guide, we can anticipate that an enterprise's VoIP security policy will be forever evolving and often very intricate.