

Basic Vulnerability Issues for SIP Security

By Mark Collier
Chief Technology Officer
SecureLogix Corporation
mark.collier@securelogix.com

Introduction

The Session Initiation Protocol (SIP) is the future protocol for Voice Over IP (VoIP). SIP promises to be the universal protocol that integrates your voice and data networks and provides the foundation for new applications. SIP enables you to mix and match voice system components from different vendors and service providers. SIP is supported at some level by most voice vendors, service providers, and even vendors such as Microsoft, in their Live Communications Server (LCS) and PC messaging clients.

Most SIP development has focused on feature sets and interoperability, with less attention paid to security. A SIP system is vulnerable to general IP and VoIP attacks, as well as attacks which are unique to SIP. It is essential to understand this and take the appropriate security measures.

SIP Introduction

SIP is a session/call control protocol defined by the Internet Engineering Task Force (IETF) and documented in [RFC 3261](#). SIP is designed as an IP protocol and resembles other IP-based protocols, such as HTTP (the protocol you use for web access). SIP messages are text-based and easier to process than those used in other VoIP protocols. SIP uses the following standards to provide basic functionality:

- Session Description Protocol (SDP) ([RFC 2327](#)) – used to define parameters for the media session.
- Real-time Transport Protocol (RTP) ([RFC 3550](#)) – used to transport the media.
- RTP Control Protocol (RTCP) ([RFC 3550](#)) – used to transmit control data for the RTP stream.
- Compressors/Decompressors (CODECS) – used to encode and compress the media.
- Feature standards – SIP is highly extensible, with many supporting standards used to define feature implementation. See [RFC 3665](#) for an example.

SIP uses the following terms for its major components:

- User Agents (UAs) – an endpoint in a SIP system, typically an IP phone, softphone, media gateway, or other media processing component such as Voice Mail (VM).
- SIP Proxy – an application that enables UAs to locate and communicate with one another.
- SIP Registrar – an application with which UAs register themselves, so they can receive calls.
- SIP Redirect Server – an application that receives requests from a UA or proxy and returns a redirection response indicating where the request should be retried.

The SIP proxy, registrar, and redirect server are usually implemented on one system, typically your IP PBX. The following figure illustrates the protocol and application “stack” described above:

Proxy		User Agents		
		SDP	Codec	RTCP
SIP			RTP	
TCP	UDP			
IPv4		IPv6		

SIP Security Overview

SIP has the same IP and application-level vulnerabilities as other VoIP protocols. There are several factors, which make SIP potentially less secure:

- Maturity – the SIP standard and supporting implementations are relatively new.
- Complexity – SIP itself is moderately complex, but with all the necessary extensions, is a complicated protocol.
- Extensibility – SIP supports extensions, which are new and often fragile from a security point of view.
- Encoding – SIP uses text messages, which are easier to see with a sniffer.


SIP offers limited built-in security. [RFC 3261](#) labels any security-related item as “SHOULD,” rather than “MUST.” Some items labeled as “RECOMMENDED” are lacking in most implementations. SIP promises interoperability, allowing you to buy components from different vendors. Interoperability can be an issue from a security point of view, because all components must support the same security standard. If they don’t, you can use only a common capability (which may be weak).

Many SIP implementations still use the Universal Datagram Protocol (UDP) for transporting SIP messages. UDP is a connection-less, unreliable form of packet transfer. UDP does not use re-transmissions or sequence numbers, so it is easier for an attacker to spoof UDP packets. In contrast, the Transmission Control Protocol (TCP) is a connection-oriented, guaranteed-delivery transport. TCP is more secure than UDP, because it involves a negotiated setup and tear down, sequence numbers, and retransmissions for lost packets.

The following sections describe several inherent vulnerabilities present in most SIP systems:

Registration Hijacking

Registration hijacking occurs when an attacker impersonates a valid UA to a registrar and replaces the legitimate registration with its own address. This attack causes all incoming calls to be sent to the UA registered by the attacker. The following figure illustrates registration hijacking:



Registration is normally performed using UDP, which makes it easier to spoof requests. Authentication is often not required and if present, is often weak (just username and password). According to RFC 3261, registrars are only “RECOMMENDED” to challenge registration requests. Most registrars either do not challenge requests, or only require a simple username/password, which can be defeated with dictionary-style attacks. A dictionary-style attack is one where the attacker has one of your usernames and then steps through a list of likely passwords built based on their knowledge of your enterprise.

An external attacker can build a directory by scanning for your registerable UA addresses. He can build a list of extensions and use the SIP “OPTIONS” message to covertly build a directory of your users.


Some enterprises may use shared, weak, or “mechanically” generated passwords (such as the extension with an extra word). In these cases, an attacker who learns one of your passwords may be able to learn all of them.

Failed registrations are not always logged. Your SIP proxy will not normally detect directory scanning and registration hijacking attempts.

Registration hijacking can result in loss of calls to the legitimate UA, which may be one of your users phones or a critical resource (e.g., a media gateway, Automated Attendant (AA), Interactive Voice Response (IVR), or VM system). Also, the rogue UA can collect authentication or other key signaling information. Or the rogue UA can pose as a Voice Mail system and trick the caller into leaving a message. The rogue UA can also perform a Man-In-The-Middle (MITL) attack, where it transparently sits between the calling and called UAs, able to collect and modify both the signaling and media. Another type of MITL attack involves redirection of an inbound call to a media gateway, generating toll fraud.

Proxy Impersonation

Proxy impersonation occurs when an attacker tricks one of your SIP UAs or proxies into communicating with a rogue proxy. If an attacker successfully impersonates a proxy, he has access to all SIP messages and is in complete control of the call. The following figure illustrates proxy impersonation:




Your UAs and proxies normally communicate using UDP and do not require strong authentication to communicate with another proxy. A rogue proxy can therefore insert itself into the signaling stream through several means, including Domain Name Service (DNS) spoofing, Address Resolution Protocol (ARP) cache spoofing, and simply changing the proxy address for a SIP phone. An impersonated proxy has full control over calls and can execute the same types of attacks described for registration hijacking.

If DNS spoofing is used to redirect outgoing calls to a particular domain (e.g., “company.com”), all outbound calls to that site can be intercepted, manipulated, blocked, conferenced, or recorded.

ARP cache spoofing is an attack against a network switch that can trick a UA into communicating with a rogue proxy on the internal network. If successful, calls originating from the UA can be intercepted, manipulated, blocked, conferenced, or recorded.

Message Tampering

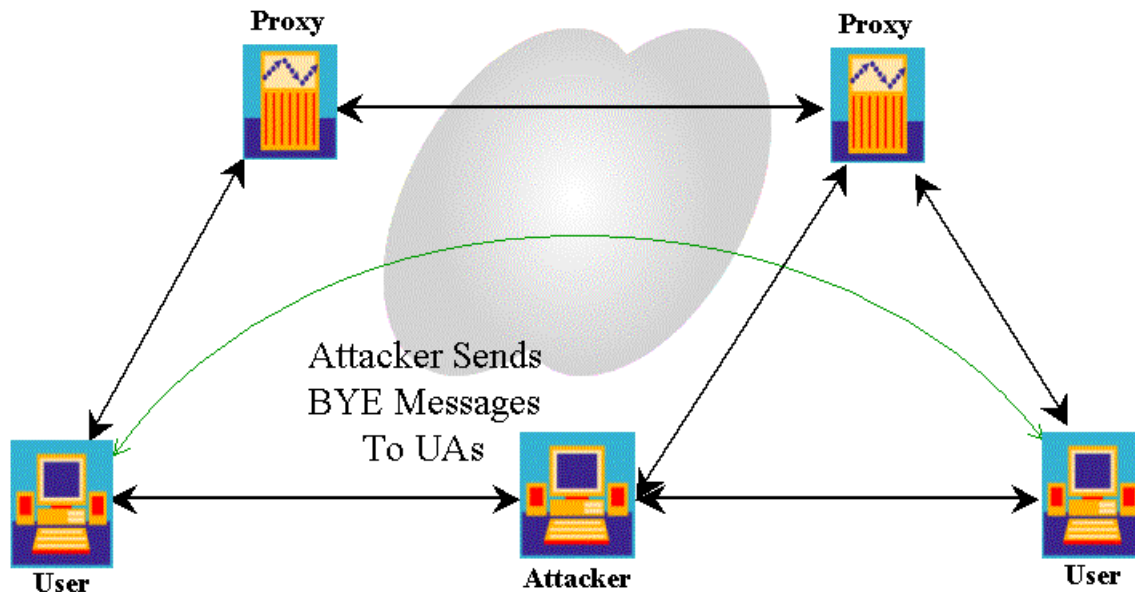
Message tampering occurs when an attacker intercepts and modifies packets exchanged between SIP components. Message tampering can occur through registration hijacking, proxy impersonation, or an attack on any component trusted to process SIP messages, such as your proxy, media gateway, or firewall. The following figure illustrates message tampering:



SIP messages have no built-in means to insure integrity. By manipulating SIP messages, an attacker can execute the same types of attacks described for registration hijacking and proxy impersonation.

Session Tear Down

Session tear down occurs when an attacker observes the signaling for a call, and then sends spoofed SIP “BYE” messages to the participating UAs. Most SIP UAs do not require strong authentication, which allows an attacker to send a properly crafted BYE messages to the two UAs, tearing down the call. The following figure illustrates session tear down:




If a UA does not check the available packet values, the attacker may not even need to observe the call signaling. If the attacker knows the address of a continually active UA (such as your media gateway, AA, IVR, trading floor phone, etc.), they can send BYE messages, causing the call(s) to be torn down. Another example of message tear down is flooding the firewall with BYE messages, possibly tearing down UDP ports opened for legitimate calls.

Denial of Service (DoS) is the primary effect of session tear down. Specific-user DoS or wholesale DoS can occur, depending upon the target. A side effect of session tear down is that the proxy may not be aware of the calls being torn down and will not have proper call records.

SIP “RE-INVITE” messages can also be used to modify media sessions. Redirecting media to broadcast addresses can cause a DoS attack. Redirecting media sessions to a media gateway can cause a DoS attack.

Denial of Service (DoS)

DoS against a SIP system can occur through any of the means described above or through additional DoS-specific attacks. Because strong authentication is rarely used, SIP processing components must trust and process SIP messages from possible attackers. The following figure illustrates some of the components vulnerable to DoS:



DoS can take the form of malformed packets, manipulating SIP states, and simple flooding, such as a “REGISTER” or “INVITE” storm (a flood of packets). Research has shown that many SIP implementations are highly vulnerable to these types of attacks. For examples, see the following link:


<http://www.ee.oulu.fi/research/ouspg/protos/testing/c07/sip/>

DoS can be especially damaging if your key voice resources are targeted (e.g., media gateways, AA, IVR, VM, and other systems). DoS can also be used to generate large numbers of toll, information (411), or emergency calls (911). Your network can also be used as a DoS launching point, from which the generated calls are directed at another enterprise.

DoS can also be directed at a firewall. SIP requires management of UDP ports for media. A DoS attack that floods the firewall with calls can prevent it from properly managing ports for legitimate calls.

Firewall/Network Address Translation (NAT) Issues

Data firewalls are commonly used to protect your trusted network from the Internet or other untrusted network. Data firewalls generally operate at the IP and UDP/TCP layers, determining what types of traffic are allowed and which systems are allowed to communicate. Data firewalls do not typically monitor the application layer – other products, such as an email content monitor or web firewall is used. Data firewalls also create issues for SIP, which uses separate IP ports for signaling and media. SIP embeds media port addresses in the SDP, which the firewall must understand in order to perform SIP-aware NAT ([RFC 2993](#)). These and other issues create the need for SIP-optimized firewalls, as shown in the following illustration:



NAT is commonly used to preserve IP addresses. For SIP to function through a firewall, the NAT must be SIP-aware, in order to modify SIP messages and to control the opening and closing of UDP ports used for media. In the future, if signaling encryption becomes commonly used, in order to function, the NAT must be able to decrypt/re-encrypt the SIP messages.

Recommendations

SIP security begins with basic IP- and VoIP-specific security. The previous article entitled “VoIP Security,” provides you with a checklist of recommendations as a starting point.

Your SIP security can be improved by using implementations that support TCP/IP for signaling, making it more difficult for an attacker to spoof SIP messages. It can also be greatly improved by using a security standard, such as the Transport Layer Security (TLS) ([RFC 2246](#)), to provide strong authentication and encryption between your SIP components. You should use these standards on every component in your SIP system, and choose vendors who support them. Once you choose a security standard such as TLS, avoid using any components that are not able to use it. Using any non-secure components—such as an inexpensive SIP phone—will allow the types of attacks described above.

The SIP community appears to be moving toward use of TLS for signaling protection and Secure RTP (SRTP) for media protection. Longer term, when SIP signaling is exchanged with proxies in an untrusted network, these same standards may be used, requiring your firewall/NAT to support TLS. If this standard (or an equivalent) is not used, then signaling will have to be accepted from untrusted/unprotected proxies. This will require use of SIP-optimized firewalls, which protect the internal SIP system from attack. Functions you should look for in a firewall include:

- Monitor inbound and outbound SIP messages for application-level attacks and discard malicious packets:
 - Monitor for directory scanning.
 - Monitor for external registration hijacking attempts.
 - Monitor for malformed SIP messages.
 - Check VIA headers and RECORD-ROUTES.
 - Block obviously malicious teardown requests.
- Support TLS and other standards-based security where possible.
- Perform SIP-aware NAT and media port management.
- Perform granular Call Admission Control (CAC). Control the number of simultaneous calls.
- Monitor for unusual calling patterns.
- Provide detailed logging of all SIP messages. Log everything for non-authenticated calls.
- Maintain QoS on all media packets. Give priority to media packets and preserve QoS markings.

Conclusions

SIP is expected to be the future VoIP protocol of choice. SIP, as with other VoIP protocols, can be difficult to secure. SIP is an evolving protocol, which does not have security built in. SIP is vulnerable to attacks common to VoIP, as well as attacks unique to SIP. Your SIP system can be best secured by following best practices for securing VoIP and using standards-based security on all system components. These same security standards should be used as SIP is exchanged with components in an untrusted network. Use SIP-optimized firewalls, which both support use of standards-based security and provide the best possible protection where system-wide standards-based security is not possible.