RUCUS	H. Scholz
Internet-Draft	freenet
Intended status: Informational	June 09, 2009
Expires: December 11, 2009	

SIP endpoint security case study
draft-scholz-endpoint-security-00

Status of this Memo

Copyright Notice

Abstract

SIP endpoints are subject to unwanted communication often perceived as Spam over Internet Telephony (SPIT). This document describes caveats on various layers which can be abused to send unsolicited messages. As a result users receive a degraded experience. The issues found are based on case studies of various events seen in VoIP provider networks.

Table of Contents

1. Introduction

TOC

Broadband access providers often ship access devices (IADs) which are by default VoIP enabled. Reasons for this are manifold are out of the scope of this document. The customer may or may not know about the VoIP feature set.

The SIP user agents may suffer from implementation issues which are not necesarilly shortcomings of the SIP standards but rather poor implementations.

In order to succesfully create a SIP dialog or make a remote phone ring the attacker needs to know the destination IP and port as well as the Contact the endpoint registered itself with. If the IP or port is wrong the message would not reach the remote SIP stack. If the requested URI does not exist the UAS should return a 404 (Not Found). Other security mechanisms (packet filter) or transport limitations (NAT bindings) may prevent communication as well.

2. IP layer

TOC

In order to conduct a successful attack the attacker needs to know where to direct traffic to. In case of devices on the end customer premises the port density per device is rather low (i.e. one or two lines) compared to commercially used SIP trunks. To make up for this drawback the attacker has to target more devices and thus IP addresses.

2.1. IP Range Guessing

TOC

SIP accounts may be tied to IP access products, i.e. part of Triple Play. In these cases an attacker only needs to identify IP ranges handed out to end customers. SIP services then may also be reachable on the same public IP address. IP access providers get assigned certain network ranges which they dynamically or statically hand out to their customers. These ranges are known to the public and can be abused. An attacker would gather list of IP ranges and attack these. Success rates are higher (80% during tests) compared to random IP addresses.

2.2. Packet filtering

TOC

The SIP standard allows for direct (peer-to-peer) communication between SIP devices. In most scenarios a centralized approach with a SIP registrar is deployed, though. The SIP endpoint registers itself on a registrar server which is usually colocated in the ISPs datacenter. All signaling is inbound from that particular SIP registrar IP address or set of IP addresses.
Large deployments usually rely on such a registrar as proxy and sometimes a B2BUA. In these scenarios the SIP endpoint would only receive inbound traffic from these configured IP addresses. All other traffic can be considered unsolicited. A device used in such an environment should ignore the rogue traffic and not pass it on to the SIP stack. Filtering rogue traffic on the SIP layer (i.e. by looking at Via headers) does not work as the Via header can be spoofed. In addition filtering on the SIP layer forces the device to parse potentially malicious messages.

TOC

3. SIP layer

An attacker, just as any other SIP User-Agent, has to craft his own SIP requests and subsequent messages. He might choose to set, skip or modify header fields and values in order to confuse the victim.

3.1. Perceived Identity

TOC

The calling party ID shown to the end user (i.e. in his phone display) cannot be assumed to be a strong identity as per [RFC4474] (Peterson, J. and C. Jennings, “Enhancements for Authenticated Identity Management in the Session Initiation Protocol (SIP),” August 2006.). While not compliant with recommendations UASes do pull the Caller ID based on one or more of these headers. Trust levels or priorities do differ between devices.

*From displayname

*From URI

*Remote-Party-ID

*P-Asserted-Identity

Even when an ISP implements identity checks and verifies certain headers an atacker may still set false information inside SIP signaling. As an example an ITSP may compare the From URI against the credentials provided in the challenge authentication. But the ITSP may let the From displayname pass through to the remote end. The remote device then displays the unverified information carried in the displayname rather than the From URI.

4. Implementation Shortcomings

TOC

Even with a flawless definition of SIP signaling devices are prone to mistakes. Problems arise due incomplete or defective implementations.

4.1. Incomplete Implementations

TOC

SIP features various transport mechanisms for the caller identity. Mechanisms include the From header (as per [RFC3261] (Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R.,

Handley, M., and E. Schooler, “SIP: Session Initiation Protocol,” June 2002.)), the expired Remote-Party-ID draft and the P-Asserted-Identity suite ([RFC3323] (Peterson, J., “A Privacy Mechanism for the Session Initiation Protocol (SIP),” November 2002.), [RFC3325] (Jennings, C., Peterson, J., and M. Watson, “Private Extensions to the Session Initiation Protocol (SIP) for Asserted Identity within Trusted Networks,” November 2002.)).
A originator may choose to use an old mechnism (i.e. Remote-Party-ID) to signal a connected party. The intermediate entities are based on P-Asserted-Identityt and thus fail to check the Remote-Party-ID header. The terminating endpoint only supports Remote-Party-ID or prefers it over other headers. As a result the Remote-Party-ID can be freely set by the caller.

4.2. Defective Implementations

TOC

4.2.1. Mapping from AORs to connected phones

TOC

To assure reachability for incoming sessions SIP endpoints frequently register themselves on a SIP registrar. Registration requires a username and credientials. On the SIP network these credentials map to a reachable SIP URI often perceived as phone number. During the registration process a unique but reusable Contact URI is used to identify a certain SIP account. An incoming call from the SIP network to the endpoint carries this Contact URI.
SIP devices that connect more than just one account usually allow to map usernames (and thus Contacts) to handsets. If the Contact of an incoming request does match a known line the handset may ring. If the Contact does not match a 404 Not Found must be returned.
Multiple devices are known to not properly check the received Contact of incomfing call setups. If the Contact is not locally known the devices fall back to ring all connected handsets. This misbehaviour allows an attacker to get around the requirement to know the current Contact of the victim.

5. Case studies

TOC

This section describes real world attacks based on the issues described in the previous sections.

5.1. German 5199362832664 case

TOC

In early September 2008 German VoIP users started to complain about their phones ringing at random times while displaying the Caller ID 5199362832664, though sometimes with a modified prefix.
This case affected customers mostly based in Germany running on DSL lines. Affected devices were VoIP enabled DSL broadband routers.
Analysis of the attack by various groups lead to the same conclusions. Klaus Darilion summed up the findings in [refs.ipcom] (Darilion, K., “Analysis of a VoIP Attack,” October 2008.). The attacker crafted non-compliant SIP messages and sent these from one apparently unspoofed IP address (213.130.74.70 in Bulgaria) to millions of destinations. Destination IP addresses were blocks assigned to the large DSL and VoIP providers in Germany.
The SIP stack on DSL routers used usually runs on the external dynamic IP address and listens on port 5060/UDP. The attacker directed the messages to this port and thus reached an huge number of devices. The signaled Contact in the Request URI was not locally known but the devices fell back to ring all connected phones instead.

5.2. DSL/VoIP router 0 Byte bug

TOC

A SIP enabled DSL router tried to suppress malicous inbound SIP traffic by comparing the IP carried in the SIP Via header to the one configured in the SIP settings. This behaviour required that messages from untrusted sources had to be parsed and interpreted by the SIP stack.
The attack in this case was to send an empty UDP message to the SIP device. The device pushed the message up into the SIP stack which immediately crashed. This resulted in the registered accounts being not reachable anymore.

6. Security Considerations

TOC

This document does not introduce new risks but rather lists known problems.

7. Informative References

TOC

[RFC4474]	Peterson, J. and C. Jennings, “Enhancements for Authenticated Identity Management in the Session Initiation Protocol (SIP),” RFC 4474, August 2006 (TXT).
[RFC3261]	Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, “SIP: Session Initiation Protocol,” RFC 3261, June 2002 (TXT).
[RFC3323]	Peterson, J., “A Privacy Mechanism for the Session Initiation Protocol (SIP),” RFC 3323, November 2002 (TXT).
[RFC3325]	Jennings, C., Peterson, J., and M. Watson, “Private Extensions to the Session Initiation Protocol (SIP) for Asserted Identity within Trusted Networks,” RFC 3325, November 2002 (TXT).
[refs.ipcom]	Darilion, K., “Analysis of a VoIP Attack,” October 2008.

Author's Address

TOC

	Hendrik Scholz
	freenet Cityline GmbH
	Am Germaniahafen 1-7
	Kiel 24143
	Germany
Phone:	+49 (0) 431 9020 552
Email:	hendrik.scholz@freenet.ag
URI:	http://freenet.ag