Internet Engineering Task Force (IETF)
Request for Comments: 6404
Category: Informational
ISSN: 2070-1721

J. Seedorf
S. Niccolini
NEC
E. Chen
NTT
H. Scholz
VOIPFUTURE

November 2011

Session PEERing for Multimedia INTerconnect (SPEERMINT) Security Threats and Suggested Countermeasures

Abstract

The Session PEERing for Multimedia INTerconnect (SPEERMINT) working group (WG) provides a peering framework that leverages the building blocks of existing IETF-defined protocols such as SIP and ENUM for the interconnection between SIP Service Providers (SSPs). The objective of this document is to identify and enumerate SPEERMINT-specific threat vectors and to give guidance for implementers on selecting appropriate countermeasures. Security requirements for SPEERMINT that have been derived from the threats detailed in this document can be found in RFC 6271; this document provides concrete countermeasures to meet those SPEERMINT security requirements. In this document, the different security threats related to SPEERMINT are classified into threats to the Lookup Function (LUF), the Location Routing Function (LRF), the Signaling Function (SF), and the Media Function (MF) of a specific SIP Service Provider. Various instances of the threats are briefly introduced inside the classification. Finally, existing security solutions for SIP and RTP/RTCP (Real-time Transport Control Protocol) are presented to describe countermeasures currently available for such threats. Each SSP may have connections to one or more remote SSPs through peering or transit contracts. A potentially compromised remote SSP that attacks other SSPs is out of the scope of this document; this document focuses on attacks on an SSP from outside the trust domain such an SSP may have with other SSPs.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at http://www.rfc-editor.org/info/rfc6404.

Copyright Notice

Copyright (c) 2011 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (http://trustee.ietf.org/license-info) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction

With Voice over IP (VoIP), the need for security is compounded because there is the need to protect both the control plane and the data plane. In a legacy telephone system, security is a more valid assumption. Intercepting conversations requires either physical access to telephone lines or a compromise to the Public Switched Telephone Network (PSTN) nodes or the office Private Branch eXchanges (PBXs). Only particularly security-sensitive organizations bother to encrypt voice traffic over traditional telephone lines. In contrast, the risk of sending unencrypted data across the Internet is more significant (e.g., dual-tone multi-frequency (DTMF) tones corresponding to the credit card number). An additional security threat to Internet Telephony comes from the fact that the signaling devices may be addressed directly by attackers as they use the same underlying networking technology as the multimedia data; traditional telephone systems have the signaling network separated from the data network. This is an increased security threat since a hacker could attack the signaling network and its servers with increased damage potential (call hijacking, call drop, Denial-of-Service (DoS) attacks [RFC4732], etc.). Therefore, there is a need to investigate the different security threats, to extract security-related requirements, and to highlight potential solutions on how to protect against such threats.

The Session PEERing for Multimedia INTerconnect (SPEERMINT) working group provides a peering framework that leverages the building blocks of existing IETF-defined protocols such as SIP and ENUM for the interconnection between SIP servers [RFC5486]. The objective of this document is to identify and enumerate SPEERMINT-specific threat vectors and to give guidance for implementers on selecting appropriate countermeasures. Security requirements for SPEERMINT can be found in RFC 6271 "Requirements for SIP-Based Session Peering" [RFC6271]. These security requirements for SPEERMINT are derived from the threats that are detailed in this document; they have been moved from an earlier version of this document to the SPEERMINT requirements document [RFC6271]. In addition to being the base for those security requirements, this document provides to implementers advice and examples for concrete countermeasures on how to meet these security requirements for SPEERMINT with technical means. The SPEERMINT terminology outlined in [RFC5486] is used throughout this document.

In this document, the different security threats related to SPEERMINT are classified into threats to the Lookup Function (LUF), the Location Routing Function (LRF), the Signaling Function (SF), and the Media Function (MF) of a specific SIP Service Provider (SSP). Various instances of the threats are briefly introduced inside the

classification. Finally, existing security solutions for SIP and RTP/RTCP are presented to describe countermeasures currently available for such threats. Each SSP may have connections to one or more remote SSPs through peering or transit contracts. A potentially compromised remote SSP that attacks other SSPs is out of the scope of this document; this document focuses on attacks on an SSP from outside the trust domain such an SSP may have with other SSPs.

2. Security Threats Relevant to SPEERMINT

This section enumerates potential security threats relevant to SPEERMINT. A taxonomy of VoIP security threats is defined in [VOIPSATAXONOMY]. This taxonomy is comprehensive and also takes into account non-VoIP-specific threats (e.g., loss of power, etc.). Threats relevant to the boundaries of Layer 5 SIP networks are extracted from this taxonomy and mapped to the functions of the SPEERMINT architecture as defined in [RFC6406]. Moreover, additional threats for the SPEERMINT architecture are listed and detailed under the same classification of SPEERMINT functions and according to the CIA (Confidentiality, Integrity, and Availability) triad:

- Lookup Function (LUF);
- Location Routing Function (LRF);
- Signaling Function (SF);
- Media Function (MF).

2.1. Threats to the Lookup Function (LUF)

For a given request, the LUF provides a mechanism to determine the identity of the requested resource on the terminating domain. The returned identity can be used to look up Session Establishment Data (SED) using the Location Routing Function (LRF). In direct peerings, the LUF is usually hosted locally, whereas in a federation context, this function may be offered by a third party.

If the LUF is hosted locally, it is vulnerable to the same threats that affect database systems in general. If the SSP relies on a remote third party to provide the LUF functionality, confidentiality, integrity, and authenticity of the responses are at risk.

2.1.1. Threats to LUF Confidentiality

For a given request, the Lookup Function (LUF) determines the target domain to which the request should be routed. The following attacks are relevant with respect to eavesdropping on LUF messages:

- SIP URI and peering domain harvesting - an attacker can exploit this weakness if the underlying database has a weak authentication system or if SIP messages are sent unencrypted, and then use the gained knowledge to launch other kinds of attacks.

- Third-party information - a LUF providing information to multiple companies / third parties can be attacked to obtain information about third party peering configurations and possible contracts.

2.1.2. Threats to LUF Integrity

The underlying database or LUF messages could be vulnerable to input/output message modification attacks:

- Injection attack - an attacker could manipulate statements performed on the database LUF messages sent to a third party. A specific version of this attack is known as "SQL injection". An SQL injection is a code insertion into the LUF due to incorrect input validation.

2.1.3. Threats to LUF Availability

The underlying database or third party LUF service could be vulnerable to:

- Denial-of-Service attacks - For example, an attacker makes incomplete requests causing the server to create an idle state for each of them, which causes memory to be exhausted.

2.2. Threats to the Location Routing Function (LRF)

The LRF determines the location of the Signaling Function (SF) for the target domain of a given request. Optionally, it may return additional SED.

2.2.1. Threats to LRF Confidentiality

Similar to the LUF, the following attacks are related to eavesdropping on LRF messages:

- URI harvesting - the attacker harvests URIs and IP addresses of the existing User Endpoints (UEs) by issuing a multitude of location requests. Direct intrusion against vulnerable UEs or telemarketing are possible attack scenarios that would use the gained knowledge.

- SIP device enumeration - the attacker discovers the IP address of each intermediate signaling device by looking at the Via and Record-Route headers of a SIP message. Targeting the discovered devices with subsequent attacks is a possible attack scenario.

2.2.2. Threats to LRF Integrity

An attacker may modify messages, e.g., by feeding bogus information to the LRF, if the routing data is not correctly validated or sent unencrypted. Dynamic call routing discovery and establishment, as in the scope of SPEERMINT, introduce opportunities for attacks such as the following:

- Man-in-the-Middle attacks - the attacker inserts or has already inserted an unauthorized node in the signaling path modifying the SED. The result is that the attacker is then able to read, insert, and modify the multimedia communications.

- Incorrect destinations - the attacker redirects the calls to an incorrect destination with the purpose of establishing fraud communications like voice phishing or DoS attacks.

2.2.3. Threats to LRF Availability

The LRF can be the object of DoS attacks. DoS attacks to the LRF can be carried out by sending a large number of queries to the LRF or LUF, with the result of preventing an Originating SSP from looking up call routing data of any URI outside its administrative domain. As an alternative, the attacker could target the DNS to disable resolution of SIP addresses.

2.3. Threats to the Signaling Function (SF)

The Signaling Function involves a great number of sensitive information. Through the Signaling Function, User Agents (UAs) assert identities and operators authorize billable resources. Correct and trusted operation of Signaling Function is essential for service providers. This section discusses potential security threats to the Signaling Function to detail the possible attack vectors.

2.3.1. Threats to SF Confidentiality

SF traffic is vulnerable to eavesdropping, in particular, when the data is moved across multiple SSPs having different levels of security policies. Threats for the SF confidentiality are listed here:

- Call pattern analysis - the attacker tracks the call patterns of the users violating his/her privacy (e.g., revealing the social network of various users, the daily phone usage, etc.); also, rival SSPs may infer information about the customer base of other SSPs in this way;

- Password cracking - the challenge-response authentication mechanism of SIP Digest can be attacked with offline dictionary attacks. With such attacks, an attacker tries to exploit weak passwords that are used by incautious users.

- Network discovery - the attacker may learn information about the internal network structure of a peering partner that is directly or indirectly connected by looking at SIP routing information (i.e, Record-Route, Via or Contact headers).

2.3.2. Threats to SF Integrity

The integrity of the SF can be violated using SIP request spoofing, SIP reply spoofing, and SIP message tampering.

2.3.2.1. SIP Request Spoofing

Most SIP request spoofing attacks first require SIP message eavesdropping. However, some of these attacks can be also performed by estimating certain fields in SIP headers (e.g., by exploiting the fact that weak implementations may generate predictable SIP Dialog parameters) or exploiting broken implementations that do not properly verify the content of certain headers. Threats in this category are as follows:

- session teardown - an attacker can send CANCEL/BYE messages in order to tear down an existing call at the SIP layer; for such an attack, the attacker either needs to know (e.g., by eavesdropping a SIP INVITE message) the SIP Dialog of the call to be hijacked (To-tag, From-tag, Call-ID) or alternatively may rely on SIP implementations that do not properly authenticate requests based on the SIP Dialog;

- Billing fraud - the attacker can modify and replay an intercepted INVITE request in order to bill a call to a victim UE and avoid paying for the phone call;

- User ID spoofing - SSPs are responsible for asserting the legitimacy of a user ID; if an SSP fails to achieve the level of identity assertion that the federation to which it belongs expects, it may create an entry point for attackers to conduct user ID spoofing attacks;

- Unwanted requests - the attacker sends requests to interfere with regular operation, e.g., by sending a REGISTER request in order to hijack calls. The SPEERMINT architecture as defined in [RFC6406] does not require registrations between the Signaling Functions (SFs) of the connected SSPs. Hence, superfluous requests like REGISTERs should be rejected.

2.3.2.2. SIP Reply Spoofing

Threats in this category are as follows:

- Forged 199 Response - the attacker sends a forged 199 response to terminate an early dialog. The forged response will not terminate the entire session but may alter the direction of the session;

- Forged 200 Response - having seen the contents of an INVITE request, an eavesdropper can inject a 200 response, affecting the processing of the transaction of all proxies between the injection point and the originating UA and at the originating UA itself. In the extreme case, this can result in a hijacked call. In many cases, however, such an attack will leave signaling artifacts that may allow it to be detected (e.g., the element receiving the forged 200 response may also receive other SIP reply messages from the actual terminating UE);

- Forged 302 Response - having seen the contents of an INVITE request, an eavesdropper could also inject a forged "302 Moved Temporarily" reply, affecting the processing of the transaction at intermediate entities and the originating UA. This may allow the attacker to successfully redirect the call to any destination UE of his choosing;

- Forged 404 Response - having seen the contents of an INVITE request, an eavesdropper could also inject a forged "404 Not Found" reply, affecting the processing of the transaction at intermediate entities and the originating UA. Such an attack may result in disrupting the call establishment.

2.3.2.3. SIP Message Tampering

This threat involves the alteration of important field values in a SIP message or in the Session Description Protocol (SDP) body. Examples of this threat could be the dropping or modification of handshake packets in order to avoid the establishment of a secure RTP session (SRTP). The same approach could be used to degrade the quality of media session by letting a UE negotiate a poor quality codec.

2.3.3. Threats to SF Availability

- Flooding attack - a Signaling Path Border Element (SBE) is susceptible to message flooding attacks that may come from interconnected SSPs;

- Session blackholing - the attacker (assumed to be able to make Man-in-the-Middle attacks) intentionally drops essential packets, e.g., INVITEs, to prevent certain calls from being established;

- SIP Fuzzing attack - fuzzing tests and software can be used by attackers to discover and exploit vulnerabilities of a SIP entity. This attack may result in crashing a SIP entity.

2.4. Threats to the Media Function (MF)

The Media Function (MF) is responsible for the actual delivery of multimedia communication between the users and carries sensitive information. Through the media function, the UE can establish secure communications and monitor the quality of conversations. Correct and trusted operations of MF is essential for privacy and service-assurance issues. This section discusses potential security threats to the MF to detail the possible attack vectors.

2.4.1. Threats to MF Confidentiality

The MF is vulnerable to eavesdropping in which the attacker may reconstruct the voice conversation or sensitive information (e.g., PINs from DTMF tones). Some SRTP key exchange mechanisms (e.g., [RFC4568]) are vulnerable to bid-down attacks, where an attacker selectively changes key exchange protocol fields in order to enforce the establishment of a less secure or even non-secure communication.

2.4.2. Threats to MF Integrity

Both RTP and RTCP are vulnerable to integrity violation in many ways:

- Media injection - if an attacker can somehow detect an ongoing media session and eavesdrop a few RTP packets, he can start sending bogus RTP packets to one of the UEs involved using the same codec. If the bogus RTP packets have consistently greater timestamps and sequence numbers (but within the acceptable range) than the legitimate RTP packets, the recipient UE may accept the bogus RTP packets and discard the legitimate ones.

- Media session teardown - the attacker sends bogus RTCP BYE messages to a target UE signaling to tear down the media communication; please note that RTCP messages are normally not authenticated.

- Quality-of-Service (QoS) degradation - the attacker sends wrong RTCP reports advertising more packet loss or more jitter than actually experimented resulting in the usage of a poor quality codec degrading the overall quality of the call experience.

2.4.3. Threats to MF Availability

- Malformed messages - the attacker tries to cause a crash or a reboot of the Data Path Border Element (DBE)/UE by sending RTP/RTCP malformed messages;

- Messages flooding - the attacker tries to exhaust the resources of the DBE/UE by sending many RTP/RTCP messages.

3. Security Requirements

3.1. Security Requirements from SPEERMINT Requirements Document

The security requirements for SPEERMINT have been moved from an earlier version of this document to the SPEERMINT requirements [RFC6271]. The security requirements for SPEERMINT are the following, from [RFC6271]:

- Requirement #15: The protocols used to query the Lookup and Location Routing Functions SHOULD support mutual authentication.

- Requirement #16: The protocols used to query the Lookup and Location Routing Functions SHOULD provide support for data confidentiality and integrity.

- Requirement #17: The protocols used to enable session peering MUST NOT interfere with the exchanges of media security attributes in SDP. Media attribute lines that are not understood by SBEs must be ignored and passed along the signaling path untouched.

3.2. How to Fulfill the Security Requirements for SPEERMINT

Requirements #15 and #16 state that the LUF and LRF should support mutual authentication, data confidentiality, and integrity. In principle, these requirements can be fulfilled technically with Transport Layer Security (TLS) or Datagram TLS (DTLS) [RFC5246] [RFC4347] or IP layer security (IPsec) [RFC4301]. From a pure

security perspective both solutions fulfill the security requirements for SPEERMINT, just on a different layer, and both solutions are widely deployed.

However, from a more practical perspective, transport layer security (i.e., TLS or DTLS) has the advantage that the application using it is aware of whether or not security (or rather the corresponding security features) is enabled. For instance, using TLS has the consequence that the connection fails if the corresponding connection endpoint cannot authenticate properly.

While IPsec fulfills the same requirements from a security perspective, IPsec is somewhat de-coupling security from the application using it. For instance, IPsec is often provided by dedicated entities in such a way that from the application layer, it cannot be recognized whether or not IPsec or certain security features are turned on ("bump-in-the-wire").

In summary, TLS (or DTLS) has some notable advantages over IPsec for addressing the SPEERMINT security requirements. In particular, transport layer security is preferable over IPsec for SPEERMINT because with TLS (or DTLS) security is more closely coupled to the LUF or LRF. From a mere technical perspective, however, both solutions (transport layer security or IPsec) fulfill the SPEERMINT security requirements, and there may be particular cases where IPsec is a preferable solution.

4. Suggested Countermeasures

This section describes implementer-specific countermeasures against the threats described in the previous sections and for addressing the SPEERMINT security requirements described in [RFC6271]. The countermeasures listed in this section are not meant to be exhaustive; rather, the suggested countermeasures are aimed to serve as starting points and to give guidance for implementers that are trying to select appropriate countermeasures against certain threats.

The following table provides a map of the relationships between threats and countermeasures. The suggested countermeasures are discussed in detail in the subsequent subsections.

Group	Threat	Suggested Countermeasure
LUF	Unauthorized access	database security BCPs (Section 4.1), Secure Exchange of SIP messages (Section 4.5)
	SQL injection	database security BCPs (Section 4.1), Secure Exchange of SIP messages (Section 4.5)
	DoS to LUF	database security BCPs (Section 4.1), Secure Exchange of SIP messages (Section 4.5)
LRF	URI harvesting	privacy protection (Section 4.4), Secure Exchange of SIP messages (Section 4.5)
	SIP equipment enumeration	privacy protection (Section 4.4), Secure Exchange of SIP messages (Section 4.5)
	MitM attack	DNSSEC (Section 4.2), Secure Exchange of SIP messages (Section 4.5)
	Incorrect destinations	DNSSEC (Section 4.2), Secure Exchange of SIP messages (Section 4.5)
	DoS to LRF	DNS replication (Section 4.3)
SF	Call pattern analysis	Secure Exchange of SIP messages (Section 4.5), Securing Session Establishment Data (Section 4.12)
	Password cracking	Secure Exchange of SIP messages (Section 4.5)
	Network discovery	Securing Session Establishment Data (Section 4.12), Topology Hiding (Section 4.10)
	Session teardown	Secure Exchange of SIP messages (Section 4.5), ingress filtering (Section 4.6)
	Billing fraud	strong identity assertion (Section 4.7)
	User ID spoofing	strong identity assertion (Section 4.7)
	Forged 200 Response	Secure Exchange of SIP messages (Section 4.5), ingress filtering (Section 4.6)
	Forged 302 Response	Secure Exchange of SIP messages (Section 4.5), ingress filtering (Section 4.6)
	Forged 404 Response	Secure Exchange of SIP messages (Section 4.5), ingress filtering (Section 4.6)
	Flooding attack	reliable border element pooling (Section 4.8), rate limit (Section 4.9)
	Session blackholing	DNSSEC (Section 4.2)

		SIP fuzzing attack	border element hardening (Section 4.11)
	MF	Eavesdropping	Encryption and Integrity Protection of Media Stream (Section 4.13)
		Media injection	Encryption and Integrity Protection of Media Stream (Section 4.13)
		Media session teardown	Encryption and Integrity Protection of Media Stream (Section 4.13)
		QoS degradation	Encryption and Integrity Protection of Media Stream (Section 4.13)
		Malformed messages	border element hardening (Section 4.11)
		Message flooding	rate limit (Section 4.9)

4.1. Database Security BCPs

Adequate security measures must be applied to the LUF to prevent it from being a target of attacks often seen on common database systems. Common security Best Current Practices (BCPs) for database systems include the use of strong passwords to prevent unauthorized access, parameterized statements to prevent SQL injections, and server replication to prevent any database from being a single point of failure. [DBSEC] is one of many existing documents that describe BCPs in this area.

4.2. DNSSEC

If DNS is used by the LRF, it is recommended to deploy the recent version of Domain Name System Security Extensions (informally called "DNSSEC-bis") defined by [RFC4033], [RFC4034], and [RFC4035]. DNSSEC has been designed to protect DNS against well-known attacks such as DNS cache poisoning or Man-in-the-Middle (MitM) attacks on DNS queries. Essentially, DNSSEC is a set of public key cryptography extensions to DNS that provide authentication of DNS data, integrity protection for DNS entries, and authenticated denial of existence regarding non-existing DNS entries. In the context of SSP peering, DNSSEC can provide authentication and integrity regarding the location of a Signaling Function (SF) entity retrieved via DNS. Using DNSSEC can thus help to defend against MitM attacks on DNS queries invoked by the LRF, session blackholing and other attacks that lead traffic to incorrect destinations.

DNSSEC has been deployed at the root level and in several top-level domains (e.g., .com and .net). Although, at the time of this writing, DNSSEC is still not yet widely deployed on the Internet, even limited deployment can add significant integrity protection and

authentication to the LRF for Signaling Function locations received via DNS entries. Neither end users nor terminals are involved in the DNS resolution process of the LRF. Hence, if a) the sending SSP uses a DNS resolver that supports DNSSEC extensions, b) the receiving SSP stores the location of its Signaling Function cryptographically signed (using DNSSEC extensions) in the DNS, and c) the sending SSP can obtain an authentication chain (i.e., a series of linked DS and DNSKEY records) to the receiving SSP, the LRF can be secured with DNSSEC. In the context of SPEERMINT, all three of these requirements can be fulfilled even in the case of partial DNSSEC deployment. In particular, even without Internet-wide deployment of DNSSEC, it may be possible for a sending SSP to obtain a suitable trust anchor for verifying the receiving SSP's public key. For instance, a suitable trust anchor could be configured for that specific SSP's top-level domain or for the particular SSP's domain directly. If the sending and the receiving SSP use a common ENUM tree, DNSSEC use with the ENUM tree's trust anchor is "straightforward".

4.3. DNS Replication

DNS replication is a very important countermeasure to mitigate DoS attacks on the LRF. Simultaneously bringing down multiple DNS servers that support the LRF is much more challenging than attacking a sole DNS server (single point of failure).

4.4. Cross-Domain Privacy Protection

Stripping Via and Record-Route headers, replacing the Contact header, and even changing Call-IDs are the mechanisms described in [RFC3323] to protect SIP privacy. This practice allows an SSP to hide its SIP network topology, prevents intermediate signaling equipment from becoming the target of DoS attacks, as well as protects the privacy of UEs according to their preferences. This practice is effective in preventing SIP equipment enumeration that exploits LRF.

4.5. Secure Exchange of SIP Messages

SIP can be used on top of UDP or TCP as transport protocol [RFC3261]. However, look-up and SED data should be exchanged securely (see security requirements (Section 3.2)), e.g., to increase the difficulty of performing session teardown and forging responses (200, 302, 404, etc). If UDP is used to carry SIP messages, DTLS should be used to secure SIP message exchange between SSPs. If TCP is used as a transport protocol, it can be secured with TLS. Therefore, depending on the underlying transport protocol, SSPs should use either DTLS or TLS to secure SIP message delivery.

In general, encryption and integrity protection of signaling messages can be achieved on the transport layer (with TLS or DTLS) or on the network layer (with IPsec). Both solutions are technically sound, but transport layer security has some advantages. Please refer to the subsection on fulfilling the SPEERMINT security requirements (Section 3.2) for a discussion on using TLS/DTLS or IPsec for protecting the confidentiality and integrity of signaling messages. Similar to strong identity assertion, a Public Key Infrastructure (PKI) is assumed to be in place for TLS/DTLS (or IPsec) deployment so that SSPs can obtain and trust the keys necessary to decrypt messages and verify signatures sent by other SSPs.

Message-oriented protection such as [RFC3261] authentication does not fulfill the SPEERMINT requirements (e.g., mutual authentication).

4.6. Ingress Filtering / Reverse-Path Filtering

Ingress filtering, i.e., blocking all traffic coming from a host that has a source address different than the addresses that have been assigned to that host (see [RFC2827]), can effectively prevent UEs from sending packets with a spoofed source IP address. This can be achieved by reverse-path filtering, i.e., only accepting ingress traffic if responses would take the same path. This practice is effective in preventing session teardown and forged SIP replies (200, 302, 404, etc.), if the recipient correctly verifies the source IP address for the authenticity of each incoming SIP message.

4.7. Strong Identity Assertion

"Caller ID spoofing" can be achieved thanks to the weak identity assertion on the From URI of an INVITE request. In a single SSP domain, strong identity assertion can be easily achieved by authenticating each INVITE request. However, in the context of SPEERMINT, only the Originating SSP is able to verify the identity directly. In order to overcome this problem, there are currently only two major approaches: transitive trust and cryptographic signature. The transitive trust approach builds a chain of trust among different SSP domains. One example of this approach is a combined mechanism specified in [RFC3324] and [RFC3325]. Using this approach in a transit peering network scenario, the terminating SSP must establish a trust relationship with all SSP domains on the path, which can be seen as an underlying weakness. The use of cryptographic signatures is an alternative approach. "Session Initiation Protocol (SIP) Authenticated Identity Body (AIB) Format" is specified in [RFC3893]. [RFC4474] introduces two new header fields, IDENTITY and IDENTITY-INFO, that allow a SIP server in the Originating SSP to digitally sign an INVITE request after authenticating the sending UE. The terminating SSP can verify if the

INVITE request is signed by a trusted SSP domain. Although this approach does not require the terminating SSP to establish a trust relationship with all transit SSPs on the path, a PKI is assumed to be in place.

4.8. Reliable Border Element Pooling

It is advisable to implement reliable pooling on border elements. An architecture and protocols for the management of server pools supporting mission-critical applications are addressed in the RSERPOOL WG. Using such mechanisms and protocols (see [RFC5351] [RFC5352] [RFC5353] for details), a UE can effectively increase its capacity in handling flooding attacks.

4.9. Rate limit

Flooding attacks on SFs and MFs can also be mitigated by limiting the rate of incoming traffic through policing or queuing. In this way, legitimate clients can be denied the service since their traffic may be discarded. Rate limiting can also be applied on a per-source-IP basis under the assumption that the source IP of each attack packet is not spoofed dynamically. Limitations related to NAT and mobility issues apply and may result in false positives (i.e., source IP addresses blocked) when multiple legitimate clients are located behind the same NAT IP address. It may be preferable to limit the number of concurrent 'sessions', i.e., ongoing calls instead of the messaging associated with it (since sessions use more resources on backend-systems). When calculating rate limits, all entities along the session path should be taken into account. SIP entities on the receiving end of a call may be the limiting factor (e.g., the number of ISDN channels on PSTN gateways) rather than the ingress limiting device.

4.10. Topology Hiding

Topology hiding applies to both the signaling and media plane and consists of limiting the amount of topology information exposed to peering partners. Topology hiding requires back-to-back user agent (B2BUA) functionality. The most common way is the use of a Session Border Controller (SBC) as SBE. Topology hiding is explained in [RFC5853].

4.11. Border Element Hardening

To prevent attacks that exploit vulnerabilities (such as buffer overflows, format string vulnerabilities, etc.) in SPEERMINT border elements, these implementations should be security hardened. For instance, fuzz testing is a common black box testing technique used

in software engineering. Also, security vulnerability tests can be carried out preventively to assure a UE/SBE/DBE can handle unexpected data correctly without crashing. [RFC4475] and [PROTOS] are examples of torture test cases specific for SIP devices and freely available security testing tools, respectively. These type of tests needs to be carried out before product release and in addition throughout the product life cycle.

4.12. Securing Session Establishment Data

Session Establishment Data (SED) contains critical information for the routing of SIP sessions. In order to prevent attacks such as service hijacking and denial of service that exploit SED, SSPs should adopt a secure transport protocol that provides authentication, confidentiality and integrity to exchange SED among themselves. Further details can be found in [DRINKS-SPPROV].

4.13. Encryption and Integrity Protection of Media Stream

The Secure Real-time Transport Protocol (SRTP) [RFC3711] prevents eavesdropping on plain RTP by encrypting the data flow. It uses AES as the default cipher and defines two modes of operation (Segmented Integer Counter Mode and f8-mode), which is agreed upon after negotiation. It also uses HMAC-SHA1 and index keeping to enable message authentication/integrity and replay protection required to prevent media injection attacks. Secure RTCP (SRTCP) provides the same security-related features to RTCP as SRTP does for RTP. SRTCP is described in [RFC3711] as optional. In order to prevent media session teardown, it is recommended to turn this feature on. The choice of the external key management protocol is left to the deployment, a PKI is necessary to implement the security requirements of the SPEERMINT requirements document.

5. Conclusions

This document presented the different SPEERMINT security threats classified in groups related to the LUF, LRF, SF, and MF, respectively. The multiple instances of the threats were presented with a brief explanation. Finally, suggested countermeasures for SPEERMINT were outlined together with possible mitigation of the existing threats by means of them.

6. Security Considerations

This document is entirely focused on the security threats for SPEERMINT.

7. Acknowledgements

This document was originally inspired by the VOIPSA VoIP Security and Privacy Threat Taxonomy. The authors would like to thank VOIPSA for having produced a comprehensive taxonomy as the starting point of this document. Additionally, the authors would like to thank Cullen Jennings, Jon Peterson, David Schwartz, Hadriel Kaplan, Peter Koch, Daryl Malas, Jason Livingood, and Robert Sparks for useful comments to previous editions of this document on the mailing list as well as during IETF meetings.

Jan Seedorf and Saverio Niccolini are partially supported by the DEMONS project, a research project supported by the European Commission under its 7th Framework Program (contract no. 257315). The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the DEMONS project or the European Commission.

8. Informative References

[DBSEC] Gertz, M. and S. Jajodia, "Handbook of Database Security: Applications and Trends", Springer, 2008.

[DRINKS-SPPROV] Mule, J., Cartwright, K., Ali, S., and A. Mayrhofer, "Session Peering Provisioning Protocol", Work in Progress, September 2011.

[PROTOS] Wieser, C., Laakso, M., and H. Schulzrinne, "SIP Robustness Testing for Large-Scale Use", First International Workshop on Software Quality (SOQUA 2004), September 2004.

[RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, May 2000.

[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.

[RFC3323] Peterson, J., "A Privacy Mechanism for the Session Initiation Protocol (SIP)", RFC 3323, November 2002.

[RFC3324] Watson, M., "Short Term Requirements for Network Asserted Identity", RFC 3324, November 2002.

[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.

[RFC3711] Baugher, M., McGrew, D., Naslund, M., Carrara, E., and K. Norrman, "The Secure Real-time Transport Protocol (SRTP)", RFC 3711, March 2004.

[RFC3893] Peterson, J., "Session Initiation Protocol (SIP) Authenticated Identity Body (AIB) Format", RFC 3893, September 2004.

[RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, March 2005.

[RFC4034] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Resource Records for the DNS Security Extensions", RFC 4034, March 2005.

[RFC4035] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Protocol Modifications for the DNS Security Extensions", RFC 4035, March 2005.

[RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.

[RFC4347] Rescorla, E. and N. Modadugu, "Datagram Transport Layer Security", RFC 4347, April 2006.

[RFC4474] Peterson, J. and C. Jennings, "Enhancements for Authenticated Identity Management in the Session Initiation Protocol (SIP)", RFC 4474, August 2006.

[RFC4475] Sparks, R., Hawrylyshen, A., Johnston, A., Rosenberg, J., and H. Schulzrinne, "Session Initiation Protocol (SIP) Torture Test Messages", RFC 4475, May 2006.

[RFC4568] Andreasen, F., Baugher, M., and D. Wing, "Session Description Protocol (SDP) Security Descriptions for Media Streams", RFC 4568, July 2006.

[RFC4732] Handley, M., Rescorla, E., and IAB, "Internet Denial-of-Service Considerations", RFC 4732, December 2006.

[RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, August 2008.

[RFC5351] Lei, P., Ong, L., Tuexen, M., and T. Dreibholz, "An Overview of Reliable Server Pooling Protocols", RFC 5351, September 2008.

[RFC5352] Stewart, R., Xie, Q., Stillman, M., and M. Tuexen, "Aggregate Server Access Protocol (ASAP)", RFC 5352, September 2008.

[RFC5353] Xie, Q., Stewart, R., Stillman, M., Tuexen, M., and A. Silverton, "Endpoint Handlespace Redundancy Protocol (ENRP)", RFC 5353, September 2008.

[RFC5486] Malas, D. and D. Meyer, "Session Peering for Multimedia Interconnect (SPEERMINT) Terminology", RFC 5486, March 2009.

[RFC5853] Hautakorpi, J., Camarillo, G., Penfield, R., Hawrylyshen, A., and M. Bhatia, "Requirements from Session Initiation Protocol (SIP) Session Border Control (SBC) Deployments", RFC 5853, April 2010.

[RFC6271] Mule, J-F., "Requirements for SIP-Based Session Peering", RFC 6271, June 2011.

[RFC6406] Malas, D., Ed. and J. Livingood, Ed., "Session PEERing for Multimedia INTerconnect (SPEERMINT) Architecture", RFC 6406, November 2011.

[VOIPSATAXONOMY] Zar, J. and et al, "VOIPSA VoIP Security and Privacy Threat Taxonomy, Public Release 1.0", http://www.voipsa.org/Activities/taxonomy.php, October 2005.

Authors' Addresses

Jan Seedorf
NEC Laboratories Europe, NEC Europe, Ltd.
Kurfuersten-Anlage 36
Heidelberg 69115
Germany

Phone: +49 (0) 6221 4342 221
EMail: jan.seedorf@neclab.eu
URI: http://www.neclab.eu

Saverio Niccolini
NEC Laboratories Europe, NEC Europe, Ltd.
Kurfuersten-Anlage 36
Heidelberg 69115
Germany

Phone: +49 (0) 6221 4342 118
EMail: saverio.niccolini@.neclab.eu
URI: http://www.neclab.eu

Eric Chen
Information Sharing Platform Laboratories, NTT
3-9-11 Midori-cho
Musashino, Tokyo 180-8585
Japan

EMail: eric.chen@lab.ntt.co.jp
URI: http://www.ntt.co.jp/index_e.html

Hendrik Scholz
VOIPFUTURE GmbH
Wendenstrasse 4
Hamburg 20097
Germany

Phone: +49 (0) 40 688 900 163
EMail: hendrik.scholz@voipfuture.com
URI: http://voipfuture.com