

Internet Engineering Task Force (IETF)
Request for Comments: 8123
Category: Informational
ISSN: 2070-1721

P. Dawes
Vodafone Group
C. Arunachalam
Cisco Systems
March 2017

Requirements for Marking SIP Messages to be Logged

Abstract

SIP networks use signaling monitoring tools to debug customer-reported problems and for regression testing if network or client software is upgraded. As networks grow and become interconnected, including connection via transit networks, it becomes impractical to predict the path that SIP signaling will take between clients and, therefore, impractical to monitor SIP signaling end-to-end.

This document describes the requirements for adding an indicator to the SIP Protocol Data Unit (PDU) or a SIP message that marks the PDU as a candidate for logging. Such a marking will typically be applied as part of network testing controlled by the network operator and not used in regular client signaling. However, such a marking can be carried end-to-end, including the SIP terminals, even if a session originates and terminates in different networks.

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 7841.

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/rfc8123>.

Copyright Notice

Copyright (c) 2017 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	3
2.	Conventions Used in This Document	3
3.	Terminology	3
3.1.	Network Boundary	3
3.2.	Trust Domain	4
3.3.	Intermediary	4
4.	Motivating Scenario	4
4.1.	Introduction	4
4.2.	Example Network Arrangement	5
4.3.	Example Debugging Procedure	6
5.	"Log Me" Marking Requirements	6
5.1.	Message Logs	6
5.2.	"Log Me" Marking	7
5.3.	Processing the "Log Me" Marker	7
6.	Security Considerations	8
6.1.	Trust Domain	8
6.2.	Security Threats	9
6.2.1.	"Log Me" Marking	9
6.2.2.	Logged Information	9
7.	IANA Considerations	9
8.	References	9
8.1.	Normative References	9
8.2.	Informative References	10
	Acknowledgments	10
	Authors' Addresses	11

1. Introduction

Service providers, enterprises, and others who operate networks that use SIP (see [RFC3261]) need the ability to debug problems reported by end users and also to run regression tests if SIP client software/hardware is upgraded. Such debugging and testing might be confined to a single service provider or network, or they may occur between the administrative domains of different network operators, including domains in different countries that are interconnected through networks belonging to one or more third parties.

A mechanism is needed to mark particular SIP sessions, i.e., those related to debugging or regression testing, as candidates for logging; this marking must be carried within the candidate SIP messages as they are routed across networks (and geographies) to enable logging at each SIP entity without having to know in advance the list of SIP entities through which the SIP signaling messages will traverse. Such marking must take into account that SIP messages might traverse different network operators, different countries, regions with different privacy requirements, and different trust domains. This document describes the requirements for such a "log me" marker for SIP signaling.

2. Conventions Used in This Document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119], except that rather than describing interoperability requirements, they are used to describe requirements to be satisfied by the "log me" marker solution.

3. Terminology

3.1. Network Boundary

A network boundary is the part of a signaling path where messages pass between entities that are under different administrative control. Figure 2 in [RFC5853] shows a network boundary between the originating gateway GW-A1 in operator A's network and the Session Border Controller (SBC) in operator B's network. A network boundary is significant in this document because manipulation of signaling at the boundary could prevent end-to-end testing or troubleshooting.

Topology hiding and protocol repair (see [RFC5853]) are two common functions that manipulate signaling at the network boundary. These functions are performed by SIP device types (see [RFC7092]) such as a Session Border Controller and Interconnection Border Control Function (IBCF).

3.2. Trust Domain

In this document, a trust domain is the set of entities that have been identified by prior agreement as the participating elements in logging, typically for the purpose of debugging or regression testing. A trust domain contains all SIP entities under configuration control of the network operator who is performing regression testing plus all SIP entities that are under configuration control of peer network operators who have agreed to participate in that regression testing. The purpose of trust domain requirements is to prevent network operators from inadvertently triggering logging in networks that are not part of any testing or troubleshooting.

3.3. Intermediary

The term "intermediary" is defined in Section 2 of [RFC7989]; it refers to any entity along the call signaling path.

4. Motivating Scenario

4.1. Introduction

Signaling for SIP session setup can cross several networks; these networks may not have common ownership and may also be in different countries. If a single operator wishes to perform regression testing or fault debugging end-to-end, the separate ownership of networks that carry the signaling and the explosion in the number of possible signaling paths through SIP entities from the originating to the terminating user make it impractical to preconfigure logging end-to-end SIP signaling of a session of interest.

4.2. Example Network Arrangement

The figure below gives an example of a signaling path through multiple networks.

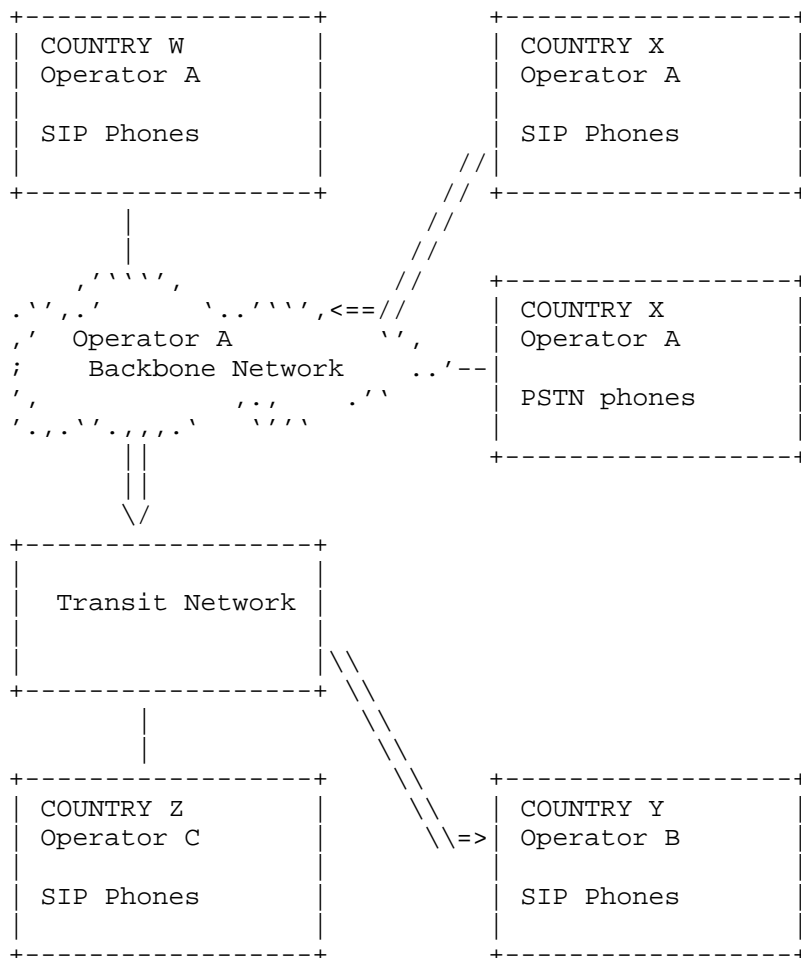


Figure 1: Example Signaling Path through Multiple Networks

4.3. Example Debugging Procedure

One possible set of steps is outlined below to illustrate the debugging procedure.

- o The user's terminal is placed in debug mode. The terminal logs its own signaling and inserts a "log me" marker into SIP requests for session setup.
- o All SIP entities that the signaling traverses, from the first proxy the terminal connects to at the edge of the network to the destination client terminal, detect that the "log me" marker is present and then log SIP requests and responses that contain the marker if configured to do so.
- o Subsequent responses and requests in the same dialog are also marked with a "log me" marker. For some scenarios, such as call transfer, related dialogs may also be marked with "log me" marker.
- o Logging stops, either because the dialog has ended or because a "stop event", typically expiry of a certain amount of time, occurred.
- o Logs are retrieved, for example, by logging on to the SIP entity or entities that contain the logs.

5. "Log Me" Marking Requirements

5.1. Message Logs

REQ1 If a SIP message is logged, then the entire SIP message (SIP headers and message body) MUST be logged using a standard logging format such as SIP Common Log Format (CLF) defined in [RFC6873].

REQ2 Header fields SHOULD be logged in the form in which they appear in the message; they SHOULD NOT be converted between long and compact forms as described in [RFC3261], Section 7.3.3.

When and how signaling logs are retrieved is out of scope of this document. Logs might be retrieved by logging on to the SIP entity that contains the logs, by sending logs to a central server that is coordinating debugging, by storing them on removable media for later manual collection, or by some other method. All log retrieval mechanisms MUST adhere to the authorization and privacy protection policies set forth by the network administrator.

5.2. "Log Me" Marking

- REQ3: It MUST be possible to mark a SIP request or response to be considered for logging by inserting a "log me" marker. This is known as "log me" marking.
- REQ4: It MUST be possible for a "log me" marker to cross network boundaries.
- REQ5: A "log me" marker MAY include an identifier that indicates the test case that caused it to be inserted, known as a "test case identifier". The test case identifier does not have any impact on session setup; it is used to collate all logged SIP requests and responses to the initial SIP request in a dialog or standalone transaction. The local Universally Unique Identifier (UUID) portion of the Session-ID described in [RFC7206] and [RFC7989] could be used as a random test case identifier.

5.3. Processing the "Log Me" Marker

- REQ6: A "log me" marker is most effective if all networks on the signaling path agree to pass it end-to-end. However, source networks should behave responsibly and not leave it to a downstream network to detect and remove a marker that it is not expecting.
- REQ7: The presence of a "log me" marker indicates that a request or response is part of debugging or regression testing.
- REQ8: It MUST be possible to insert a "log me" marker in SIP responses that correspond to SIP requests with a "log me" marker in order to ensure that the complete SIP transaction is logged. This requirement applies to endpoints, SIP/Public Switched Telephone Network (PSTN) gateways, and back-to-back user agents (B2BUAs).
- REQ9: The "log me" marker mechanism SHOULD allow a SIP intermediary to request logging SIP requests and responses on behalf of the originating endpoint. The typical use case for this requirement is for compatibility with User Agents (UAs) that have not implemented "log me" marking, i.e., when a UA has not marked a request or when responses received on a dialog of interest for logging do not contain an echoed "log me" marker. Another use case is when the session origination UA that inserted the "log me" marker is no longer participating in the

session (e.g., call transfer scenarios) and the intermediary adds a "log me" marker in related sessions to enable end-to-end signaling analysis.

- REQ10: The mechanism MUST allow stateless processing of SIP requests that contain a "log me" marker by SIP intermediaries. This requirement enables the SIP intermediaries to base the decision to log a SIP request or response solely on the presence of the "log me" marker.
- REQ11: The scope of a SIP message logging request includes all requests and responses within a given dialog. The scope can be extended to related dialogs that correspond to an end-to-end session for scenarios discussed in REQ9. The "log me" request MUST be indicated at the beginning of the dialog of interest and SHOULD continue to the dialog end without any stop and restart during the duration of the dialog.
- REQ12: The presence of a "log me" marker might cause some SIP entities to log signaling. Therefore, this marker MUST be removed at the earliest opportunity if it has been incorrectly inserted (e.g., mid-dialog or outside the configured start and stop of "log me" marking).

The definition of types of events that cause logging to stop and the configuration of SIP entities to detect such "stop events" is outside the scope of this document.

6. Security Considerations

In order to prevent any security implications of a "log me" marker, the marker itself MUST NOT contain any sensitive information, detecting its presence or absence MUST NOT reveal sensitive information, and maliciously adding a "log me" marker MUST NOT adversely affect a network. This section analyzes how to meet these requirements.

6.1. Trust Domain

Since a "log me" marker may cause a SIP entity to log the SIP header and body of a request or response, the "log me" marker MUST be removed at a trust domain boundary. If a prior agreement to log sessions exists with the next hop network, then the "log me" marker SHOULD NOT be removed.

6.2. Security Threats

6.2.1. "Log Me" Marking

The "log me" marker MUST NOT convey any sensitive information, although the "log me" marker will sometimes be inserted because a particular device is experiencing problems. The "log me" marker MUST NOT reveal any information related to any SIP user or device.

The insertion of the "log me" marker at the endpoint MUST be approved by the end user or by the network administrator. Similarly, network administrator authorization is required for a SIP intermediary to insert a "log me" marker on behalf of a UA that does not support "log me" marking.

Activating a debug mode affects the operation of a terminal; therefore, the debugging configuration MUST be supplied by an authorized party to an authorized terminal through a secure communication channel.

6.2.2. Logged Information

Logged signaling is privacy-sensitive data; therefore, signaling logs MUST NOT be readable by an unauthorized third party.

7. IANA Considerations

This document does not require any IANA actions.

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC6873] Salgueiro, G., Gurbani, V., and A. Roach, "Format for the Session Initiation Protocol (SIP) Common Log Format (CLF)", RFC 6873, DOI 10.17487/RFC6873, February 2013, <<http://www.rfc-editor.org/info/rfc6873>>.

8.2. Informative References

- [RFC3261] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, DOI 10.17487/RFC3261, June 2002, <<http://www.rfc-editor.org/info/rfc3261>>.
- [RFC5853] Hautakorpi, J., Ed., Camarillo, G., Penfield, R., Hawrylyshen, A., and M. Bhatia, "Requirements from Session Initiation Protocol (SIP) Session Border Control (SBC) Deployments", RFC 5853, DOI 10.17487/RFC5853, April 2010, <<http://www.rfc-editor.org/info/rfc5853>>.
- [RFC7092] Kaplan, H. and V. Pascual, "A Taxonomy of Session Initiation Protocol (SIP) Back-to-Back User Agents", RFC 7092, DOI 10.17487/RFC7092, December 2013, <<http://www.rfc-editor.org/info/rfc7092>>.
- [RFC7206] Jones, P., Salgueiro, G., Polk, J., Liess, L., and H. Kaplan, "Requirements for an End-to-End Session Identification in IP-Based Multimedia Communication Networks", RFC 7206, DOI 10.17487/RFC7206, May 2014, <<http://www.rfc-editor.org/info/rfc7206>>.
- [RFC7989] Jones, P., Salgueiro, G., Pearce, C., and P. Giralt, "End-to-End Session Identification in IP-Based Multimedia Communication Networks", RFC 7989, DOI 10.17487/RFC7989, October 2016, <<http://www.rfc-editor.org/info/rfc7989>>.

Acknowledgments

The authors wish to thank Jorgen Axell, Ben Campbell, Keith Drage, Vijay Gurbani, Christer Holmberg, Hadriel Kaplan, Paul Kyzivat, James Polk, Gonzalo Salgueiro, Alberto Llamas, Brett Tate, Paul Giralt, Stewart Bryant, Sean Turner, and Dan Romascanu for their constructive comments and guidance while developing this document.

Authors' Addresses

Peter Dawes
Vodafone Group
The Connection
Newbury, Berkshire RG14 2FN
United Kingdom

Email: peter.dawes@vodafone.com

Chidambaram Arunachalam
Cisco Systems
7200-12 Kit Creek Road
Research Triangle Park, NC 27709
United States of America

Email: carunach@cisco.com