

Network Working Group
Request for Comments: 5069
Category: Informational

T. Taylor, Ed.
Nortel
H. Tschofenig
Nokia Siemens Networks
H. Schulzrinne
Columbia University
M. Shanmugam
Detecon
January 2008

Security Threats and Requirements for
Emergency Call Marking and Mapping

Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Abstract

This document reviews the security threats associated with the marking of signalling messages to indicate that they are related to an emergency, and with the process of mapping locations to Universal Resource Identifiers (URIs) that point to Public Safety Answering Points (PSAPs). This mapping occurs as part of the process of routing emergency calls through the IP network.

Based on the identified threats, this document establishes a set of security requirements for the mapping protocol and for the handling of emergency-marked calls.

Table of Contents

- 1. Introduction 2
- 2. Terminology 3
- 3. Marking, Mapping, and the Emergency Call Routing Process . . . 3
 - 3.1. Call Marking 3
 - 3.2. Mapping 4
- 4. Objectives of Attackers 4
- 5. Potential Attacks 5
 - 5.1. Attacks Involving the Emergency Identifier 5
 - 5.2. Attacks Against or Using the Mapping Process 5
 - 5.2.1. Attacks Against the Emergency Response System 6
 - 5.2.2. Attacks to Prevent a Specific Individual from Receiving Aid 7
 - 5.2.3. Attacks to Gain Information about an Emergency 7
- 6. Security Requirements Relating to Emergency Marking and Mapping 8
- 7. Security Considerations 9
- 8. Acknowledgements 10
- 9. References 10
 - 9.1. Normative References 10
 - 9.2. Informative References 10

1. Introduction

Legacy telephone network users can summon help for emergency services (such as an ambulance, the fire department, and the police) using a well known number (e.g., 911 in North America, 112 in Europe). A key factor in the handling of such calls is the ability of the system to determine caller location and to route the call to the appropriate Public Safety Answering Point (PSAP) based on that location. With the introduction of IP-based telephony and multimedia services, support for emergency calling via the Internet also has to be provided. Two core components of IP-based emergency calling include an emergency service identifier and a mapping protocol. The emergency service identifier indicates that the call signaling establishes an emergency call, while the mapping protocol translates the emergency service identifier and the caller's geographic location into an appropriate PSAP URL.

Attacks against the Public Switched Telephone Network (PSTN) have taken place for decades. The Internet is seen as an even more hostile environment. Thus, it is important to understand the types of attacks that might be mounted against the infrastructure providing emergency services and to develop security mechanisms to counter those attacks. While this can be a broad topic, the present document restricts itself to attacks on the mapping of locations to PSAP URIs and attacks based on emergency marking. Verification by the PSAP

operator of the truthfulness of a reported incident and various other attacks against the PSAP infrastructure related to the usage of faked location information are outside the scope of the document.

This document is organized as follows: Section 2 describes basic terminology. Section 3 briefly describes how emergency marking and mapping fit within the process of routing emergency calls. Section 4 describes some motivations of attackers in the context of emergency calling, Section 5 describes and illustrates the attacks that might be used, and Section 6 lists the security-related requirements that must be met if these attacks are to be mitigated.

2. Terminology

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], with the qualification that unless otherwise stated, they apply to the design of the mapping protocol, not its implementation or application.

The terms "call taker", "mapping service", "emergency caller", "emergency identifier", "mapping", "mapping client", "mapping server", "mapping protocol", and "Public Safety Answering Point (PSAP)" are taken from [RFC5012].

The term "location information" is taken from RFC 3693 [RFC3693].

The term "emergency caller's device" designates the IP host closest to the emergency caller in the signalling path between the emergency caller and the PSAP. Examples include an IP phone running SIP, H.323, or a proprietary signalling protocol, a PC running a soft client or an analogue terminal adapter, or a residential gateway controlled by a softswitch.

3. Marking, Mapping, and the Emergency Call Routing Process

This memo deals with two topics relating to the routing of emergency calls to their proper destination: call marking and mapping.

3.1. Call Marking

Marking of call signalling enables entities along the signalling path to recognize that a particular signalling message is associated with an emergency call. Signalling containing the emergency identifier may be given priority treatment, special processing, and/or special routing.

3.2. Mapping

An important goal of emergency call routing is to ensure that any emergency call is routed to a PSAP. Preferably, the call is routed to the PSAP responsible for the caller's location, since misrouting consumes valuable time while the call taker locates and forwards the call to the right PSAP. As described in [RFC5012], mapping is part of the process of achieving this preferable outcome.

In brief, mapping involves a mapping client, a mapping server, and the protocol that passes between them. The protocol allows the client to pass location information to the mapping server and to receive back a URI, which can be used to direct call signalling to a PSAP.

4. Objectives of Attackers

Attackers may direct their efforts either against a portion of the emergency response system or against an individual. Attacks against the emergency response system have three possible objectives:

- o to deny system services to all users in a given area. The motivation may range from thoughtless vandalism, to wide-scale criminality, to terrorism. One interesting variant on this motivation is the case where a victim of a large emergency hopes to gain faster service by blocking others' competing calls for help.
- o to gain fraudulent use of services, by using an emergency identifier to bypass normal authentication, authorization, and accounting procedures.
- o to divert emergency calls to non-emergency sites. This is a form of a denial-of-service attack similar to the first item, but quite likely more confusing for the caller himself or herself since the caller expects to talk to a PSAP operator but instead gets connected to someone else.

Attacks against an individual fall into two classes:

- o attacks to prevent an individual from receiving aid.
- o attacks to gain information about an emergency that can be applied either against an individual involved in that emergency or to the profit of the attacker.

5. Potential Attacks

5.1. Attacks Involving the Emergency Identifier

The main possibility of attack involves use of the emergency identifier to bypass the normal procedures in order to achieve fraudulent use of services. An attack of this sort is possible only if the following conditions are true:

- a. The attacker is the emergency caller.
- b. The call routing system assumes that the emergency caller's device signals the correct PSAP URI for the caller's location.
- c. The call enters the domain of a service provider, which accepts it without applying normal procedures for authentication and authorization because the signalling carries the emergency identifier.
- d. The service provider routes the call according to the called address (e.g., SIP Request-URI), without verifying that this is the address of a PSAP (noting that a URI by itself does not indicate the nature of the entity it is pointing to).

If these conditions are satisfied, the attacker can bypass normal service provider authorization procedures for arbitrary destinations, simply by reprogramming the emergency caller's device to add the emergency identifier to non-emergency call signalling. In this case, the call signalling most likely will not include any location information, or there could be location information, but it is false.

An attacker wishing to disrupt the emergency call routing system may use a similar technique to target components of that system for a denial-of-service attack. The attacker will find this attractive to reach components that handle emergency calls only. Flooding attacks are the most likely application of the technique, but it may also be used to identify target components for other attacks by analyzing the content of responses to the original signalling messages.

5.2. Attacks Against or Using the Mapping Process

This section describes classes of attacks involving the mapping process that could be used to achieve the attacker goals described in Section 4.

5.2.1. Attacks Against the Emergency Response System

This section considers attacks intended to reduce the effectiveness of the emergency response system for all callers in a given area. If the mapping operation is disabled, then the emergency caller's device might not have the correct PSAP URI. As a consequence, the probability that emergency calls will be routed to the wrong PSAP increases. In the worst case, the emergency caller's device might not be able to obtain a PSAP URI at all. Routing to the wrong PSAP has a double consequence: emergency response to the affected calls is delayed, and PSAP call taker resources outside the immediate area of the emergency are consumed due to the extra effort required to redirect the calls. Alternatively, attacks that cause the client to receive a URI that does not lead to a PSAP have the immediate effect of causing emergency calls to fail.

Three basic attacks on the mapping process can be identified: denial of service, impersonation of the mapping server, or corruption of the mapping database. Denial of service can be achieved in several ways:

- o by a flooding attack on the mapping server;
- o by taking control of the mapping server and either preventing it from responding or causing it to send incorrect responses; or
- o by taking control of any intermediary node (for example, a router) through which the mapping queries and responses pass, and then using that control to block them. An adversary may also attempt to modify the mapping protocol signalling messages. Additionally, the adversary may be able to replay past communication exchanges to fool an emergency caller by returning incorrect results.

In an impersonation attack, the attacker induces the mapping client to direct its queries to a host under the attacker's control rather than the real mapping server, or the attacker suppresses the response from the real mapping server and sends a spoofed response.

The former type of impersonation attack itself is an issue of mapping server discovery rather than the mapping protocol directly. However, the mapping protocol may allow impersonation to be detected, thereby preventing acceptance of responses from an impersonating entity and possibly triggering a more secure discovery procedure.

Corruption of the mapping database cannot be mitigated directly by mapping protocol design. Once corruption has been detected, the mapping protocol may have a role to play in determining which records have been corrupted.

Beyond these attacks on the mapping operation itself, it is possible to use mapping to attack other entities. One possibility is that mapping clients are misled into sending mapping queries to the target of the attack instead of the mapping server. Prevention of such an attack is an operational issue rather than one of protocol design. Another possible attack is where the mapping server is tricked into sending responses to the target of the attack through spoofing of the source address in the query.

5.2.2. Attacks to Prevent a Specific Individual from Receiving Aid

If an attacker wishes to deny emergency service to a specific individual, the mass attacks described in Section 5.2.1 will obviously work provided that the target individual is within the affected population. Except for the flooding attack on the mapping server, the attacker can in theory limit these attacks to the target, but this requires extra effort that the attacker is unlikely to expend. If the attacker is using a mass attack but does not wish to have too broad an effect, it is more likely to attack for a carefully limited period of time.

If the attacker wants to be selective, however, it may make more sense to attack the mapping client rather than the mapping server. This is particularly so if the mapping client is the emergency caller's device. The choices available to the attacker are similar to those for denial of service on the server side:

- o a flooding attack on the mapping client;
- o taking control of any intermediary node (for example, a router) through which the mapping queries and responses pass, and then using that control to block or modify them.

Taking control of the mapping client is also a logical possibility, but raises no issues for the mapping protocol.

5.2.3. Attacks to Gain Information about an Emergency

This section discusses attacks used to gain information about an emergency. The attacker may be seeking the location of the caller (e.g., to effect a criminal attack). Alternatively, the attacker may be seeking information that could be used to link an individual (the caller or someone else involved in the emergency) with embarrassing information related to the emergency (e.g., "Who did the police take away just now?"). Finally, the attacker could be seeking to profit from the emergency, perhaps by offering his or her services (e.g., a news reporter, or a lawyer aggressively seeking new business).

The primary information that interceptions of mapping requests and responses will reveal are a location, a URI identifying a PSAP, the emergency service identifier, and the addresses of the mapping client and server. The location information can be directly useful to an attacker if the attacker has high assurance that the observed query is related to an emergency involving the target. The type of emergency (fire, police, or ambulance) might also be revealed by the emergency service identifier in the mapping query. The other pieces of information may provide the basis for further attacks on emergency call routing, but because of the time factor, are unlikely to be applicable to the routing of the current call. However, if the mapping client is the emergency caller's device, the attacker may gain information that allows for interference with the call after it has been set up or for interception of the media stream between the caller and the PSAP.

6. Security Requirements Relating to Emergency Marking and Mapping

This section describes the security requirements that must be fulfilled to prevent or reduce the effectiveness of the attacks described in Section 5. The requirements are presented in the same order as the attacks.

From Section 5.1:

Attack A1: fraudulent calls.

Requirement R1: For calls that meet conditions a) to c) of Section 5.1, the service provider's call routing entity MUST verify that the destination address (e.g., SIP Request-URI) presented in the call signalling is that of a PSAP.

Attack A2: Use of emergency identifier to probe in order to identify emergency call routing entities for attack by other means.

Requirement: None identified, beyond the ordinary operational requirement to defend emergency call routing entities by means such as firewalls and, where possible, authentication and authorization.

From Section 5.2.1:

Attack A3: Flooding attack on the mapping client, mapping server, or a third entity.

Requirement R2: The mapping protocol MUST NOT create new opportunities for flooding attacks, including amplification attacks.

Attack A4: Insertion of interfering messages.

Requirement R3: The protocol MUST permit the mapping client to verify that the response it receives is responding to the query it sent out.

Attack A5: Man-in-the-middle modification of messages.

Requirement R4: The mapping protocol MUST provide integrity protection of requests and responses.

Requirement R5: The mapping protocol or the system within which the protocol is implemented MUST permit the mapping client to authenticate the source of mapping responses.

Attack A6: Impersonation of the mapping server.

Requirement R6: The security considerations for any discussion of mapping server discovery MUST address measures to prevent impersonation of the mapping server.

Requirement R5 also follows from this attack.

Attack A7: Corruption of the mapping database.

Requirement R7: The security considerations for the mapping protocol MUST address measures to prevent database corruption by an attacker.

Requirement R8: The protocol SHOULD include information in the response that allows subsequent correlation of that response with internal logs that may be kept on the mapping server, to allow debugging of mis-directed calls.

From Section 5.2.2: No new requirements.

From Section 5.2.3:

Attack A8: Snooping of location and other information.

Requirement R9: The protocol and the system within which it is implemented MUST maintain confidentiality of the request and response.

7. Security Considerations

This document addresses security threats and security requirements. Therefore, security is considered throughout this document.

8. Acknowledgements

The writing of this document has been a task made difficult by the temptation to consider the security concerns of the entire personal emergency calling system, not just the specific pieces of work within the scope of the ECRIT Working Group. Hannes Tschofenig performed the initial security analysis for ECRIT, but it has been shaped since then by the comments and judgement of the ECRIT WG at large. At an earlier stage in the evolution of this document, Stephen Kent of the Security Directorate was asked to review it and provided extensive comments, which led to a complete rewriting of it. Brian Rosen, Roger Marshall, Andrew Newton, and most recently, Spencer Dawkins, Kamran Aquil, and Ron Watro have also provided detailed reviews of this document at various stages. The authors thank them.

We would like to thank Donald Eastlake for his review on behalf of the Security Area Directorate and Christian Vogt for his review as part of the General Area Review Team.

Finally, we would like to thank Jari Arkko, Jon Peterson, and Russ Housley for their IETF Last Call comments.

9. References

9.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

9.2. Informative References

[RFC3693] Cuellar, J., Morris, J., Mulligan, D., Peterson, J., and J. Polk, "Geopriv Requirements", RFC 3693, February 2004.

[RFC5012] Schulzrinne, H. and R. Marshall, Ed., "Requirements for Emergency Context Resolution with Internet Technologies", RFC 5012, January 2008.

Authors' Addresses

Tom Taylor (editor)
Nortel
1852 Lorraine Ave
Ottawa, Ontario K1H 6Z8
Canada

EEmail: tom.taylor@rogers.com

Hannes Tschofenig
Nokia Siemens Networks
Otto-Hahn-Ring 6
Munich, Bavaria 81739
Germany

EEmail: Hannes.Tschofenig@nsn.com
URI: <http://www.tschofenig.com>

Henning Schulzrinne
Columbia University
Department of Computer Science
450 Computer Science Building
New York, NY 10027
US

Phone: +1 212 939 7004
EEmail: hgs+ecrit@cs.columbia.edu
URI: <http://www.cs.columbia.edu>

Murugaraj Shanmugam
Detecon International GmbH
Oberkasseler str 2
Bonn, NRW 53227
Germany

EEmail: murugaraj.shanmugam@detecon.com

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.