Network Working Group
Internet-Draft
Expires: April 18, 2005

H. Schulzrinne
Columbia U.
B. Rosen
Marconi

October 18, 2004

Emergency Services for Internet Telephony Systems
draft-schulzrinne-sipping-emergency-arch-02

Status of this Memo

Copyright Notice

Abstract

Summoning emergency help is a core feature of telephone networks. This document describes how the Session Initiation Protocol (SIP) can be used to provide advanced emergency services for voice-over-IP (VoIP). The architecture employs standard SIP features and requires no new protocol mechanisms. DNS is used to map civil and geospatial locations to the appropriate emergency call center.

Table of Contents

1. Requirements notation

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].

2. Terminology

(Emergency) call taker: An emergency call taker is the person that answers an emergency call, typically located in an emergency call center.

ECC (emergency control center): Facilities used by emergency organizations to accept and handle emergency calls. A PSAP (below) forwards emergency calls to the emergency control center, which dispatches polic, fire and rescue services. An ECC serves a limited geographic area. A PSAP and ECC can be combined into one facility. (ETSI SR 002 180 definition)

ESRP (emergency service routing proxy): SIP proxy that routes incoming emergency calls to the appropriate ECC.

PSAP (public safety answering point): Physical location where emergency calls are received under the responsibility of a public authority. (This terminology is used by both ETSI, in ETSI SR 002 180, and NENA.) In the United Kingdom, PSAPs are called Operator Assistance Centres, in New Zealand Communications Centres.

SIP proxy: see [RFC3261].

SIP UA (user agent): see [RFC3261].

Stationary device (user): User agent that is connected to the network at a fixed, long-term-stable geographic location. Examples include a home PC or a payphone.

Nomadic device (user): User agent that is connected to the network temporarily, for relatively short durations, but does not move significantly during the lifetime of a network connection or during the emergency call. Examples include a laptop using an 802.11 hotspot or a desk IP phone that is moved from one cubicle to another.

Mobile device (user): User agent that changes geographic location and possibly its network attachment point during an emergency call.

3. Overview

Summoning police, the fire department or an ambulance in emergencies is one of the fundamental and most-valued functions of the telephone. As telephone functionality moves from circuit-switched telephony to Internet telephony, its users rightfully expect that this core functionality works at least as well as for the older technology. However, many of the technical advantages of Internet telephony require re-thinking of the traditional emergency calling architecture. This challenge also offers an opportunity to improve

the working of emergency calling technology, while potentially lowering its cost and complexity.

It is beyond the scope of this document to enumerate and discuss all the differences between traditional (PSTN) and Internet telephony, but the core differences can be summarized as separation of signaling and media data, the emergence of application-independent carriers, and the potential mobility of all end systems, including landline systems and not just those using radio access technology.

This document focuses on how emergency call centers (PSAPs) (Section 2) can natively handle Internet telephony emergency calls, rather than describing how circuit-switched PSAPs can handle VoIP calls. However, in many cases, PSAPs making the transition from circuit-switched interfaces to packet-switched interfaces may be able to use some of the mechanisms described here, in combination with gateways that translate packet-switched calls into legacy interfaces, e.g., to continue to be able to use existing call taker equipment.

Existing emergency call systems are organized nationally; there are currently no international standards. However, Internet telephony does not respect national boundaries, and thus an international standard is required.

Furthermore, VoIP endpoints can be connected through tunneling mechanisms such as virtual private networks (VPNs). This significantly complicates emergency calling, because the location of the caller and the first element that routes emergency calls can be on different continents, with different conventions and processes for handling of emergency calls. The IETF has historically refused to create national variants of its standards. Thus, this document attempts to take into account best practices that have evolved for circuit switched PSAPs, but makes no assumptions on particular operating practices currently in use, numbering schemes or organizational structures.

This document assumes that PSAP interface is using the Session Initation Protocol (SIP). Use of a single protocol greatly simplifies the design and operation of the emergency calling infrastructure. Only peer-to-peer protocols such as H.323, ISUP and SIP are suitable for inter-domain communications, ruling out master-slave protocols such as MGCP or H.248/Megaco. The latter protocols can natually be used by the enterprise or carrier placing the call, but any such call would reach the PSAP through a media gateway controller, similar to how interdomain VoIP calls would be placed. Other signaling protocols may also use protocol translation to communicate with a SIP-enabled PSAP.

Existing emergency services rely exclusively on voice and conventional text telephony (known as TDD in the United States) media streams. However, more choices of media offer additional ways to communicate, evaluate and assist callers and call takers to handle emergency calls. For example, instant messaging and video could improve the ability to evaluate the situation and provide appropriate instruction prior to arrival of emergency crews. Thus, the architecture described here supports the creation of sessions of any media type, negotiated between the caller and PSAP using existing SIP protocol mechanisms [RFC3264].

While, traditionally, emergency services have been summoned by voice calls only, this document does not rule out the use of additional media during an emergency call, both to support callers with disabilities (e.g., through interactive text or video communications) and to provide additional information to the call taker and caller. For example, video from the caller to the PSAP may allow the call taker to better assess the emergency situation; a video session from the PSAP to the emergency caller may allow the call taker to provide instructions for first aid.

The choice of media and encodings is negotiated on a call-by-call basis using standard SIP mechanisms [RFC3264]. To ensure that at least one common means of communications, this document recommends certain minimal capabilities in Section 14 that call taker user agents and PSAP-operated proxies should possess.

This document does not prescribe the detailed network architecture for PSAPs or collection of PSAPs. For example, it does not describe where PSAPs may place firewalls or how many SIP proxies they should use.

This document does not introduce any new SIP header fields, request methods, status codes, message bodies, or events. User agents unaware of the recommendations in this draft can place emergency calls, but may not be able to provide the same user interface functionality. The document suggests behavior for proxy servers, in particular outbound proxy servers.

4. Identifying an Emergency Call

Using the PSTN, emergency help can often be summoned at a designated, widely known number, regardless of where the telephone was purchased. However, this number differs between localities, even though it is often the same for a country or region (such as many countries in the European Union). For end systems based on the Session Initiation Protocol (SIP), it is desirable to have a universal identifier, independent of location, to simplify the user experience, allow the

automated inclusion of location information and to allow the device and other entities in the call path to perform appropriate processing.

As part of the overall emergency calling architecture, we define a common user identifier, "sos", as the contact mechanism for emergency assistance. We refer to this URI as the "emergency calling URI". The calling user agent sets both the "To" header and the request-URI to the emergency URI, so that entities after the ESRP can still readily determine that this is an emergency call. Details are described in [I-D.ietf-sipping-sos]. The draft also discusses how a user agent or outbound proxy determines whether a dialed number represents an emergency number and thus should be translated into a "sos" URI.

In addition, user agents SHOULD detect emergency calls following local emergency calling conventions. There are two local conventions, namely those local to the user's SIP domain, e.g., a user's network at work, and those at the caller's current geographic location, e.g., while traveling. The former can be obtained using SIP/XCAP and DNS configuration mechanisms (Section 12).

Location information can be provided by the user agent or a proxy. If the user agent provides this information, the user agent needs to be able to determine that a call is indeed an emergency call as it is unlikely to include location information in each call.

5. Location and Its Role in an Emergency Call

5.1 Introduction

Caller location plays a central role in routing emergency calls. For practical reasons, each PSAP generally handles only calls for a certain geographic area. Other calls that reach it by accident must be manually re-routed (transferred) to the appropriate PSAP, increasing call handling delay and the chance for errors. The area covered by each PSAP differs by jurisdiction, where some countries have only a small number of PSAPs, while others devolve PSAP responsibilities down to the community level.

In most cases, PSAPs cover at least a city or town, but there are some areas where PSAP coverage areas follow old telephone rate center boundaries and may straddle more than one city.

5.2 Types of Location Information

There are four primary types of location information: civic, postal, geospatial, and cellular cell tower and sector.

Civic: Civic location information describes the location of a person or object by a floor and street address that corresponds to a building or other structure. (This is sometimes also called "civic" location information.)

Postal: Postal addresses are similar to civic addresses, but the may contain post office boxes or street addresses that do not correspond to an actual building. Also, the name of the post office sometimes does not correspond to the actual community name. Postal addresses are generally unsuitable for emergency call routing, but may be the only address available to a service provider, derived from billing records.

Geospatial: Geospatial addresses contain longitude, latitude and altitude information.

Cell tower/sector: Cell tower and sectors identify the cell tower and the antenna sector that the mobile device is currently using. (Cell/sector information could also be transmitted as an irregularly shaped polygon of geospatial coordinates reflecting the likely geospatial location of the mobile device, but since these boundaries are not sharp, transmitting the raw information is probaby preferable.) Mobile systems, possible in conjunction with the cell site location, may also transmit mobile country code (MCC) and mobile network code (MNC) of the host network. This MCC and MNC constitutes location information, in that it tells that the user (with border constraints) is in a particular country. In some cases, this may be sufficient for determining the PSAP to be used.

5.3 Sources of Location Information

Location information can be entered by the user or installer of a device ("manual configuration"), can be measured by the end system, can be conveyed to the end system or can be measured by a third party and inserted into the call signaling. We discuss these in detail below.

In some cases, an entity may have multiple sources of location information, possibly partially contradictory. This is particularly likely if the location information is determined both by the end system and a third party. This document provides no recommendation on how to reconcile conflicting location information or which one is to be used by routing elements. Conflicting location information is particularly harmful if it points to multiple distinct PSAPs. If there is no other basis for choice, the ESRP SHOULD determine the appropriate PSAP for all location objects and, if there is a conflict, route based on the most accurate one.

All location objects MUST be delivered to the PSAP. To facilitate such policy decisions, location information SHOULD contain

information about the source of data, such as GPS, manually entered or based on subscriber address information. In addition, the author of the location information SHOULD be included.

TBD: SIP system should indicate which location information has been used for routing, so that the same location information is used for all call routing decisions. Otherwise, two proxies might pick different location information from the call request, each pointing to the other one.

End systems and network elements can derive location information from a variety of sources. It is not the goal of this document to exhaustively enumerate them, but we provide a few common examples in the sections below.

5.3.1 Manually-Entered Location Information

Location information can be maintained by the end user or the installer of a network connection ("wire database"). In LANs, wire databases map Ethernet switch ports to office locations. In DSL installations, the local telephone carrier maintains a mapping of wire pairs to subscriber addresses.

Even for IEEE 802.11 wireless access points, wire data bases may provide sufficient location accuracy.

Location information added by end users is almost always inferior to measured or wire database information, as users may mistype civic location information, may not know the meaning of geospatial coordinates or may use address information that does not correspond to a recognized civic address.

Wire databases are likely to be the most promising solution for residential users where a service provider knows the customer's service address. The service provider can then perform address verification, similar to the current system in some jurisdictions.

5.3.2 End-System Measured Location Information

GPS: Global Positioning System (GPS) information is generally only available where there is a clear view of a large swath of the sky. It is accurate to tens of feet.

5.3.3 Third-party Measured Location Information

Wireless triangulation: Elements in the network infrastructure triangulate end systems based on signal strength or time of arrival. Signal strength may be reported by access points, special measurement devices or the end systems.

Location beacons: A short range wireless beacon, e.g., using BlueTooth or infrared, announces its location to mobile devices in the vicinity.

5.3.4 Conveying Location to End Systems

Unless a user agent has access to locally measured location information, it MUST use DHCP to obtain location information. DHCP can deliver civic [I-D.ietf-geopriv-dhcp-civil] or geospatial [RFC3825] information. User agents MUST support both formats. Note that a user agent can use DHCP, via the INFORM request, even if it uses other means to acquire its IP address.

In addition, link-layer mechanisms such as the Link-Layer Discovery Protocol (LLDP, IEEE 802.1ab), with proposed extensions, MAY also be used to deliver such information.

5.4 Using Location Information for Call Routing

Since all existing emergency services have limited geographic and jurisdictional coverage, all emergency calls need to be routed to the appropriate PSAP. Rather than to the geographically closest PSAP, calls need to be directed to the most jurisdictionally appropriate one, which may well be further away.

5.5 Mid-Call Location Information

Location information may not be available at call setup time. For example, if a GPS-enabled cell phone is turned on and then immediately places an emergency call, it can take an additional 20-25 seconds before the cell phone acquires a GPS fix and its location. Thus, while it is necessary and expedient to include caller location information in the call setup message, this is not sufficient in all circumstances. In some cases, the initial call setup will proceed based on, for example, cell and sector information and then add location information during the call, rather than delaying the initial call setup by an unacceptable amount of time.

In addition, the location of a mobile caller, e.g., in a vehicle or aircraft, can change significantly during the emergency call.

Location updates MAY be conveyed either in re-INVITE or UPDATE messages or the PSAP may subscribe to the location information of the caller, using SIP presence mechanisms (RFC 3265 [RFC3265]RFC 3856

[RFC3856]). Authorization for subscriptions is for future study.

5.6 Civic Address Verification

Users of SIP endpoints must be able to verify that their address is valid ahead of an actual emergency call. For example, in the United States, the Master Street Address Guide (MSAG) records all valid street addresses and is used to ensure that phone billing records correspond to valid emergency service street addresses.

There are several ways to verify this information, depending on its source. If the location information is provided by the network service provider via DHCP, SIP end systems SHOULD display this information at boot-up and at regular intervals thereafter to allow users to confirm that the information is correct.

If the DNS emergency services directory contains street-level addresses rather than just towns or counties, an end system can verify that a civic address, configured manually or via DHCP, exists.

6. Routing the Call to the PSAP

6.1 Routing the First Request

Emergency calls are routed based on one or more of the following criteria expressed in the call setup request (INVITE):

Location: Since each PSAP serves a limited geographic region and transferring existing calls delays the emergency response, calls need to be routed to the most appropriate PSAP. In this architecture, emergency call setup requests contain location information, expressed in civic or geospatial coordinates, that allows such routing. If there is no or imprecise (e.g., cell tower and sector) information at call setup time, an on-going emergency call may also be transferred to another PSAP based on location information that becomes available in mid-call.

Type of emergency service: In some jurisdictions, emergency calls for fire, police, ambulance or mountain rescue are directed to emergency-specific PSAPs. We support this mechanism by optionally labeling calls with a service identifier [I-D.ietf-sipping-sos]. Using the caller preferences [I-D.ietf-sip-callerprefs] mechanisms, ESRPs can then route labeled calls appropriately.

Media capabilities of caller: In some cases, emergency call centers for specific caller media preferences, such as typed text or video, are separate from voice systems. Also, even if media capability does not affect the selection of the PSAP, there may be call takers within the PSAP that are specifically trained, e.g., in interactive text or sign language communications. Again, we

use the callee capabilities [I-D.ietf-sip-callee-caps] mechanism to label and route such calls.

Call routing can be performed in three different ways:

1. The calling user agent can route the call to the PSAP URI it received in a DHCP or SIP configuration message (not discussed further; TBD!). This is generally only possible for stationary and nomadic devices. In that case, the DHCP server has to be able to map callers to PSAP URIs.
2. The calling user agent uses DNS to translate a geospatial or civic address into a URI identifying a PSAP or group of PSAPs. This mode can be used by stationary, nomadic and mobile devices.
3. Any SIP proxy along the call path from the mobile device to the home domain can recognize an emergency call and route it based on the location information contained in the INVITE request, using DNS or other mechanisms not defined in this document.

Each proxy receiving an emergency call request, identified as described in Section 4, attempts to route the call to the most appropriate PSAP, group of PSAPs or another, more suitable ESRP. Similarly, a user agent can also directly route emergency calls if it has location information, either obtained locally or from a redirect response provided by the outbound proxy. There are three types of routing actions: default routing, DNS-based routing and local routing. Not all routing actions can take all three dimensions (location, type of service, capabilities) into account.

ESRPs and user agents using default routing forward all emergency call requests to one designated ESRP, regardless of the location of the caller, type of service or media capabilities.

ESRPs and user agents using DNS-based routing employ the mechanism in [I-D.rosen-dns-sos] to route calls to another ESRP that is qualified to handle the emergency call.

Finally, an ESRP MAY use a local database or other query protocols to perform call routing using location, type of service or callee capabilities. The details of such a database are beyond the scope of this document.

Call routing may combine several of these methods. For example, an outbound proxy might route all emergency calls to a designated ESRP. The ESRP extracts civic location information from the request and converts the elements into a DNS query, using the "sos.arpa" domain, starting from the countrycode and adding the A1 through A6 elements of the civic location contained [I-D.ietf-geopriv-pidf-lo] in the call. It starts from the most precise location and strips location

elements if there are no entries at that level. For example, the ESRP might find that "leonia.bergen.nj.us.sos.arpa" does not exist, but that "bergen.nj.us.sos.arpa" features an entry. The ESRP identified in that entry may in turn use the location information to route the request to individual communities, without exposing this information to the public. In the extreme case, only a country-level ESRP needs to be exposed in DNS. Thus, each jurisdiction can make its own decisions as to whether it wants to use DNS or local databases to perform call routing.

If an emergency call INVITE request does not contain location information and no other location hints (such as subscriber identity) are available, the first ESRP in the call path SHOULD route it to a PSAP or group of PSAPs that is geographically local to that proxy, since no other call routing can be performed.

Jurisdictions organizing PSAPs may choose to implement multiple levels of routing based on location. For example, a state, province or county might deploy an ESRP in front of a collection of PSAPs. The information available to a VoIP carrier or enterprise ESRP may be coarse, so that any location within the state or province gets routed to that representative ESRP, with that ESRP performing the detailed routing to a specific PSAP. The routing mechanism used by the ESRP may nor may not rely on public information. Depending on choices made by the operator of the PSAP and ESRP, the PSAP may only be reachable by SIP requests routed through the ESRP.

6.2 DNS-based Mapping from Civic Coordinates to PSAP URIs

We define a hierarchy of domain names corresponding to the country name and A1 through A6 hierarchy of administrative units (e.g., state, county, and city), as subdomains below sos.arpa. For example, the domain leonia.bergen.nj.us.sos.arpa designates the town of Leonia in Bergen County in the state of New Jersey, United States. Unless the domain is the lowest one in the hierarchy, with no subdomains, it contains a PTR resource record pointing to the leaves below it. For example:

```
us.sos.arpa.         PTR al.us.sos.arpa.
us.sos.arpa.         PTR ak.us.sos.arpa.
us.sos.arpa.         PTR as.us.sos.arpa.
us.sos.arpa.         PTR az.us.sos.arpa.
...
us.sos.arpa.         PTR wi.us.sos.arpa.
us.sos.arpa.         PTR wy.us.sos.arpa.

nj.us.sos.arpa.      PTR sussex.nj.us.sos.arpa.
nj.us.sos.arpa.      PTR passaic.nj.us.sos.arpa.
nj.us.sos.arpa.      PTR bergen.nj.us.sos.arpa.
...
bergen.nj.us.sos.arpa. PTR fort_lee.bergen.nj.us.sos.arpa.
bergen.nj.us.sos.arpa. PTR leonia.bergen.nj.us.sos.arpa.
...
leonia.bergen.nj.us.sos.arpa IN NAPTR
   NAPTR 100 10 "u" "SOS" "/.*/sips:fire@leoniaboro.org/i" .
...
```

PTR records were chosen since they are designed to allow retrieval of multiple matching resource records, without doing a zone transfer.

Street names and their components (XXX in PIDF-LO) are concatenated by using a hyphen in the order .... Empty elements are omitted, including the hyphen.

6.3 Updating Location Information

Location information is needed both for routing the initial INVITE message in a call as well as possibly later during a call since location information may change or only become available later, after the call has reached a PSAP.

The caller sends UDPATE [RFC3311], either prior to completion of the initial INVITE transaction or during the call, to the destination. Care must be taken that these requests are routed to the same destination as the original call-initiating request. This is unlikely to be a problem for a re-INVITE if the Contact header field in the 200 OK indicates the PSAP address.

7. Signaling of Emergency Calls

Since emergency calls carry privacy-sensitive information, they are subject to the requirements for geospatial protocols. In particular, signaling information MUST be carried in TLS, i.e., in 'sips' mode.

Details can be found in [I-D.ietf-sipping-location-requirements].

8. Preventing Call Misdirection

We need to prevent an emergency call reaching a destination other than an PSAP. For example, a rogue UA able to intercept SIP requests might be able to impersonate an PSAP.

In the absence of a globally recognized certificate that ensures that the owner is a legitimate PSAP, we rely on a chain of trust enforced by the 'sips' URI schema. The 'sips' URI schema forces each SIP hop to route the call only to destinations supporting TLS transport. Each ESRP MUST verify that the next-hop destination chosen as described in Section 6 corresponds to the server certificate offered by that destination.

9. Including a Valid Call-Back Identifier

The call taker must be able to reach the emergency caller if the original call is disconnected. In traditional emergency calls, wireline and wireless emergency calls include a callback number for this purpose. In SIP systems, the caller SHOULD include a Contact header indicating its device URI, if available, or possibly a GRUU [I-D.ietf-sip-gruu] if calls need to be routed via a proxy.

10. Mid-Call Services and Behavior

If the called PSAP can sign the response, it can include the 'service' media feature tag in the response to indicate to the calling user agent that the call is an emergency call. The calling user agent can then modify its normal behavior to reflect the special nature of the call, e.g., to prevent accidental disconnects. A UA MUST NOT modify its behavior unless the call response is authenticated, as this could otherwise be used by malicious destinations to affect caller UA functionality.

The PSAP MAY return 403 (Forbidden) in response to a BYE request if caller hangs up before the PSAP wants to relinquish the call.

11. Requirements for SIP Proxy Servers

All ESRP SHOULD support RFC 3261 [RFC3261] with UDP, TCP, TLS transports.

User agent servers and proxy servers MUSTNOT require that the user agent client be registered or authenticated in order to place an emergency call.

For robustness, ESRPs SHOULD NOT use RFC 1918 [RFC1918] addresses, i.e., should not be behind network address translators.

12. Configuration

SIP devices do not require any additional configuration to place emergency calls. They SHOULD use the local outbound proxy, discovered via [RFC3361] or [RFC3319].

However, to acquire local dial plan numbers, the SIP configuration framework [I-D.ietf-sipping-config-framework] can be used. The format for dial plans remains to be defined. A device may retrieve dial plan information for emergency calls from two locations, namely the user's home domain and the local outbound proxy, as described in Section 3.13 of [I-D.ietf-sipping-config-framework].

Since a traveling user cannot rely on a DHCP server in the visited location to have accurate local emergency number information, we also propose a new DNS resource record, EN. Typically, this resource record will be associated with a country-level 'sos.arpa' zone, as most countries either have or are developing country-wide emergency numbers. These number strings are treated as dial strings [I-D.rosen-iptel-dialstring], not "tel" URIs. TBD: It might be possible to use NAPTR [RFC2915] records to include translations such that 112 becomes sos for de.sos.arpa. NAPTR translations are not limited to hostnames or URIs.

In the example below, the German emergency number for police is translated into an 'sos' URI. This only works if there is a designated SIP proxy that can route all emergency calls originating in Germany. There does not appear to be a way to substitute the caller's current home AOR domain, although one could conceivably adopt a convention for including this information. Note that this mechanism would also allow direct routing based on finer-grained location information, e.g., at the city level.

```
de.sos.arpa.
;;       order pre flags service        regexp     replacement
IN NAPTR 100 10 "u" "SOS" "/110/sips:sos.police@notfall.de/i" .

bonn.nrw.de.sos.arpa.
;;       order pre flags service        regexp      replacement
IN NAPTR 100 10 "u" "SOS" "/110/sips:sos.police@pol.bonn.de/i" .
```

Example NAPTR records to map dial strings to 'sos' URIs

Figure 2

In addition to the generic mechanism describe above, there may be access transport specific mechanisms for downloading this information to the user agent. For example, a 3GPP phone from release 5 onwards

can have emergency number information downloaded from visited network entities at network registration time.

13. Testing

13.1 Testing Mechanism

Since the emergency calling architecture consists of a number of pieces operated by independent entities, it is important to be able to test whether an emergency call is likely to succeed without actually occupying the human resources at a PSAP. Both signaling and media paths need to be tested since NATs and firewalls may allow the session setup request to reach the PSAP, while preventing the exchange of media.

INVITE requests to the user "sos" address and a service indicator of sos.test can be used to test if the "sos" address is valid. As in standard SIP, a 200 (OK) response indicates that the address was recognized and a 404 (Not found) that it was not. Such request cause no further action. The response MAY contain a message body describing the PSAP that was reached and may automatically. The test server SHOULD echo a limited number of RTP audio packets to test media connectivity.

User agents SHOULD perform a full call test, including media, according to Section 13.1 after a disconnect and subsequent change in IP address, as the NAT configuration may have changed.

User agents MUST NOT place a test call immediately after booting, as a widespread power outage and subsequent restoration would impose an inordinate load on the emergency call routing system.

13.2 Manual Testing

A compliant user agent implementation MUST have the capability to perform the test outlined in Section 13.1 by explicit user request.

13.3 Automatic 'sos' Resolution Testing

If a user agent does its own call routing, it MUST periodically and after every significant location change ascertain that it can still resolve its current location to a PSAP address. It does not actually have to generate a SIP request to test emergency calls.

A significant location change is defined here as a change of one degree or more in longitude or latitude or a change in the A1 or A2 level of civil locations.

The periodic test should be performed every 24 to 48 hours and MUST be randomly placed over the testing interval.

14. Requirements for SIP User Agents

14.1 Emergency call taker

To increase the likelihood that diverse user equipment can successfully communicate with the PSAP, it is recommended that call taker equipment has at least the following capabilites:

Signaling: RFC 3261 [RFC3261], with UDP, TCP and TLS (sips) support.
Media transport: RTP and RTCP according to RFC 3550 [RFC3550] and RFC 3551 [RFC3551]. SRTP according to RFC 3711. [RFC3711]
Audio codecs: G.711, GSM 06.10, DTMF support using RFC 2833 [RFC2833], with forward error correction (RFC 2733 [RFC2733]).
Interactive text: using RTP according to RFC 2793bis [I-D.ietf-avt-rfc2793bis].
Video: Support H.261, H.263 and H.264 in QCIF, CIF and 4CIF sizes.
SIP-based instant messaging: RFC 3428 [RFC3428]

14.2 Calling users

A user agent placing an emergency call SHOULD use the "sips" URI schema for all such calls, forcing these calls to use TLS as secure hop-by-hop transport. If a call cannot be established using TLS transport, the user agent SHOULD attempt a call using the "sip" URI.

If a user agent receives a redirect (3xx) response for an emergency call, it MUST include the location information contained in that response in the outgoing call. This differs from regular behavior for redirects, where the message body is not copied into the new call.

User agents MUST support blind transfer using REFER [RFC3515].

A user agent MUST check the Contact URI in redirect responses to see if it is an emergency call, as described in Section 4. If so, the behavior in the previous paragraph applies.

End systems that allow human users to initiate an emergency call with a single button press or other similar stimulus SHOULD require callers to confirm their call.

UAs SHOULD place a "Priority" header with value "emergency" in all emergency calls, but its presence cannot be relied upon to identify an emergency call.

15. Example Call Flows

TBD

16. Alternatives Considered

This is a non-normative appendix. During discussions of emergency calling, a number of suggestions are commonly made. Below, we discuss some of the reasons why these alternatives do not satisfy the requirements of emergency calling.

16.1 tel URIs

Instead of providing URIs to call routing proxies or end systems, it has been suggested that end systems be configured with a "tel" URI [I-D.ietf-iptel-rfc2806bis]. Such a "tel" URI would have to be routed to a geographically appropriate telephony gateway, as it is unlikely that every building, enterprise or residence will have its own gateway. VoIP devices can be used in networks that are completely unaware of VoIP services, with VoIP service providers that are physically far removed from the caller's network location. Thus, the use of a tel URI simply moves the problem to the outbound proxy, which has to use the caller's location to determine the appropriate telephony gateway.

In addition, emergency telephone numbers are far from universal, with some such numbers used for non-emergency purposes elsewhere. Thus, an outbound proxy would have to ascertain the location of the caller to guess whether the "tel" URI identifies an emergency call or some other number.

Thus, "tel" URIs are not likely to be appropriate or sufficient for identifying emergency calls and do not, by themselves, solve the call routing problem.

16.2 DHCP for Configuring the PSAP URI

One could add emergency calling information to network configuration protocols such as DHCP. A DHCP option could identify the appropriate PSAP URI, for example. This simple approach runs into two problems: lack of congruence of DHCP and PSAP serving areas and difficulty of DHCP server configuration.

DHCP servers may provide information to large groups of geographically dispersed users, often spanning jurisdictional boundaries. (For example, CATV plants generally do not follow community boundaries.)

The DHCP server would also have to be able to determine the appropriate URI. Unless all calls, at least within a country, are routed to a single logical proxy and that proxy maintains a national jurisdictional database, DCHP serves would have to be manually or automatically configured with regional or local PSAP information. Since the number of such DHCP servers is large and since authorities are unlikely to maintain a mailing list of DHCP server operators, it would be up to each owner of such servers to keep up with jurisdictional changes. While such changes are not frequent, they do occur, as PSAP jurisdictions are merged or as unincorporated areas are merged into neighboring municipalities.

17. Security Considerations

17.1 Caller Authentication

To prevent crank calling and to support call back, PSAPs may want to authenticate the caller. If the call is routed via an outbound proxy, the outbound proxy may be able to ascertain whether the identity provided in the call corresponds at least to the appropriate domain. However, visiting users may legitimately feature a different caller identity than the domain of the outbound proxy. Mechanisms such as the authenticated identity body [I-D.ietf-sip-authid-body] may be used to assert identities.

In keeping with established customs in circuit-switched emergency calling, authentication cannot be made a pre-requisite for routing or accepting an emergency call. However, a call taker may be more suspicious of a caller and request additional information if the call authenticity cannot be verified.

17.2 PSAP Impersonation

See Section 8.

With DNS-based call routing (Section 6), an attacker could modify the DNS entries for one or more PSAPs, re-routing calls destined for them. Thus, the use of secure DNS is RECOMMENDED.

17.3 Call Signaling Integrity

To prevent a malicious outsider from manipulating call information, SIP requests SHOULD be routed via "sips" from caller to emergency call taker.

17.4 Media Integrity and Confidentiality

Media integrity and confidentiality can be assured by the use of

SRTP.

17.5 PSAP Hiding

The issue of hiding PSAP identity has been raised in mailing list discussion. It has been argued that hiding the identity of an PSAP confers some protection against denial-of-service attacks, for example. However, it appears that this notion is based on false assumptions. Unless a B2BUA or NAT is involved, media packets will carry the IP address of the PSAP (or one of its call takers) and thus can be readily used to deduce the address of the PSAP, even if it is not advertised in DNS. (B2BUAs and NATs have known architectural, reliability and other operational disadvantages that do not recommend their use simply to hide PSAP addresses.)

Similarly, trying to protect the mapping between geographic location and PSAP is similarly difficult. Unless it is required that all location information is verified in real time, which would be close to impossible for mobile devices, end systems can simply pretend to be in different parts of the city or county and deduce which PSAP is answering the call.

18. Changes Since the Last Version

Added references to LLDP (IEEE 802.1ab) as a protocol for conveying location to end system.
Changed ECC to PSAP. ECC is also used by ETSI (ETSI SR 002 180) to designate Emergency Control Centers, which dispatch emergency assistance. ETSI uses the term PSAP, so it seemed unnecessary to create new terminology.
The description of location sources has been extended.
An non-normative section on why DHCP or tel URIs are not sufficient has been added.
Text on testing and preventing call hang-ups has been added.

19. Acknowledgements

Keith Drage provided helpful comments.

20. References

20.1 Normative References

[I-D.ietf-avt-rfc2793bis]
Hellstrom, G., "RTP Payload for Text Conversation", draft-ietf-avt-rfc2793bis-09 (work in progress), August 2004.

[I-D.ietf-geopriv-dhcp-civil]
Schulzrinne, H., "Dynamic Host Configuration Protocol (DHCPv4 and DHCPv6) Option for Civic Addresses Configuration Information", draft-ietf-geopriv-dhcp-civil-04 (work in progress), October 2004.

[I-D.ietf-geopriv-pidf-lo]
Peterson, J., "A Presence-based GEOPRIV Location Object Format", draft-ietf-geopriv-pidf-lo-03 (work in progress), September 2004.

[I-D.ietf-sip-authid-body]
Peterson, J., "SIP Authenticated Identity Body (AIB) Format", draft-ietf-sip-authid-body-03 (work in progress), May 2004.

[I-D.ietf-sip-callee-caps]
Rosenberg, J., "Indicating User Agent Capabilities in the Session Initiation Protocol (SIP)", draft-ietf-sip-callee-caps-03 (work in progress), January 2004.

[I-D.ietf-sip-callerprefs]
Rosenberg, J., Schulzrinne, H. and P. Kyzivat, "Caller Preferences for the Session Initiation Protocol (SIP)", draft-ietf-sip-callerprefs-10 (work in progress), October 2003.

[I-D.ietf-sip-gruu]
Rosenberg, J., "Obtaining and Using Globally Routable User Agent (UA) URIs (GRUU) in the Session Initiation Protocol (SIP)", draft-ietf-sip-gruu-02 (work in progress), July 2004.

[I-D.ietf-sipping-config-framework]
Petrie, D., "A Framework for Session Initiation Protocol User Agent Profile Delivery", draft-ietf-sipping-config-framework-04 (work in progress), July 2004.

[I-D.ietf-sipping-location-requirements]
Polk, J., "Requirements for Session Initiation Protocol Location Conveyance", draft-ietf-sipping-location-requirements-01 (work in progress), July 2004.

[I-D.ietf-sipping-sos]

Schulzrinne, H., "Emergency Services URI for the Session Initiation Protocol", draft-ietf-sipping-sos-00 (work in progress), February 2004.

[I-D.rosen-dns-sos]
Rosen, B., "Emergency Call Information in the Domain Name System", draft-rosen-dns-sos-01 (work in progress), July 2004.

[I-D.rosen-iptel-dialstring]
Rosen, B., "Dialstring parameter for the sip URI", draft-rosen-iptel-dialstring-00 (work in progress), June 2004.

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

[RFC2733] Rosenberg, J. and H. Schulzrinne, "An RTP Payload Format for Generic Forward Error Correction", RFC 2733, December 1999.

[RFC2793] Hellstrom, G., "RTP Payload for Text Conversation", RFC 2793, May 2000.

[RFC2833] Schulzrinne, H. and S. Petrack, "RTP Payload for DTMF Digits, Telephony Tones and Telephony Signals", RFC 2833, May 2000.

[RFC2915] Mealling, M. and R. Daniel, "The Naming Authority Pointer (NAPTR) DNS Resource Record", RFC 2915, September 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.

[RFC3264] Rosenberg, J. and H. Schulzrinne, "An Offer/Answer Model with Session Description Protocol (SDP)", RFC 3264, June 2002.

[RFC3265] Roach, A., "Session Initiation Protocol (SIP)-Specific Event Notification", RFC 3265, June 2002.

[RFC3311] Rosenberg, J., "The Session Initiation Protocol (SIP) UPDATE Method", RFC 3311, October 2002.

[RFC3319] Schulzrinne, H. and B. Volz, "Dynamic Host Configuration Protocol (DHCPv6) Options for Session Initiation Protocol (SIP) Servers", RFC 3319, July 2003.

[RFC3361] Schulzrinne, H., "Dynamic Host Configuration Protocol (DHCP-for-IPv4) Option for Session Initiation Protocol (SIP) Servers", RFC 3361, August 2002.

[RFC3428] Campbell, B., Rosenberg, J., Schulzrinne, H., Huitema, C. and D. Gurle, "Session Initiation Protocol (SIP) Extension for Instant Messaging", RFC 3428, December 2002.

[RFC3515] Sparks, R., "The Session Initiation Protocol (SIP) Refer Method", RFC 3515, April 2003.

[RFC3550] Schulzrinne, H., Casner, S., Frederick, R. and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, RFC 3550, July 2003.

[RFC3551] Schulzrinne, H. and S. Casner, "RTP Profile for Audio and Video Conferences with Minimal Control", STD 65, RFC 3551, July 2003.

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

[RFC3825] Polk, J., Schnizlein, J. and M. Linsner, "Dynamic Host Configuration Protocol Option for Coordinate-based Location Configuration Information", RFC 3825, July 2004.

[RFC3856] Rosenberg, J., "A Presence Event Package for the Session Initiation Protocol (SIP)", RFC 3856, August 2004.

20.2 Informative References

[I-D.ietf-iptel-rfc2806bis] Schulzrinne, H., "The tel URI for Telephone Numbers", draft-ietf-iptel-rfc2806bis-09 (work in progress), June 2004.

[RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G. and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.

Authors' Addresses

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

Phone: +1 212 939 7042
EMail: hgs@cs.columbia.edu
URI: http://www.cs.columbia.edu

Brian Rosen
Marconi
2000 Marconi Drive
Warrendale, PA 15086
US

EMail: brian.rosen@marconi.com

Intellectual Property Statement

Disclaimer of Validity

Copyright Statement

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.