

ecrit
Internet-Draft
Intended status: Informational
Expires: March 11, 2012

B. Rosen
NeuStar
H. Schulzrinne
Columbia U.
J. Polk
Cisco Systems
A. Newton
TranTech/MediaSolv
September 8, 2011

Framework for Emergency Calling using Internet Multimedia
draft-ietf-ecrit-framework-13

Abstract

The IETF has standardized various aspects of placing emergency calls. This document describes how all of those component parts are used to support emergency calls from citizens and visitors to authorities.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on March 11, 2012.

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. Terminology	3
2. Introduction	4
3. Overview of how emergency calls are placed	7
4. Which devices and services should support emergency calls	11
5. Identifying an emergency call	12
6. Location and its role in an emergency call	13
6.1. Types of location information	15
6.2. Location determination	16
6.2.1. User-entered location information	17
6.2.2. Access network "wire database" location information	18
6.2.3. End-system measured location information	18
6.2.4. Network measured location information	19
6.3. Who adds location, endpoint or proxy	19
6.4. Location and references to location	20
6.5. End system location configuration	20
6.6. When location should be configured	22
6.7. Conveying location	23
6.8. Location updates	23
6.9. Multiple locations	23
6.10. Location validation	24
6.11. Default location	25
6.12. Location format conversion	26
7. LIS and LoST discovery	26
8. Routing the call to the PSAP	26
9. Signaling of emergency calls	28
9.1. Use of TLS	28
9.2. SIP signaling requirements for User Agents	29
9.3. SIP signaling requirements for proxy servers	29
10. Call backs	29
11. Mid-call behavior	30
12. Call termination	30
13. Disabling of features	31
14. Media	31
15. Testing	31
16. Security Considerations	32
17. IANA Considerations	32
18. Acknowledgments	32
19. Informative References	33
Authors' Addresses	36

1. Terminology

This document uses terms from [RFC3261], [RFC5222] and [RFC5012]. In addition the following terms are used:

Access network: The access network supplies IP packet service to an endpoint. Examples of access networks include digital subscriber lines (DSL), cable modems, IEEE 802.11, WiMaX, enterprise local area networks and cellular data networks.

Confidence: Confidence is an estimate indicating how sure the measuring system is that the actual location of the endpoint is within the bounds defined by the uncertainty value, expressed as a percentage. For example, a value of 90% indicates that the actual location is within the uncertainty nine times out of ten.

Dispatch Location: The dispatch location is the location used for dispatching responders to the person in need of assistance. The dispatch location must be sufficiently precise to easily locate the caller; it typically needs to be more accurate than the routing location.

Location configuration: During location configuration, an endpoint learns its physical location.

Location Configuration Protocol (LCP): A protocol used by an endpoint to learn its location.

Location conveyance: Location conveyance delivers location information to another element.

Location determination: Location determination finds where an endpoint is physically located. For example, the endpoint may contain a Global Navigation Satellite System (GNSS) receiver used to measure its own location or the location may be determined by a network administrator using a wiremap database.

Location Information Server (LIS): A Location Information Server stores location information for retrieval by an authorized entity.

Mobile device: A mobile device is a user agent that may change its physical location and possibly its network attachment point during an emergency call.

NENA (National Emergency Number Association): The National Emergency Number Association is an organization of professionals to "foster the technological advancement, availability and implementation of a universal emergency telephone number system in North America." It develops emergency calling specifications and procedures.

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

Physical location: A physical location describes where a person or device is located in physical space, described by a coordinate system. It is distinguished from the network location, described by a network address.

PSAP: Public Safety Answering Point, the call center that answers emergency calls.

Routing Location: The routing location of a device is used for routing an emergency call and may not be as precise as the Dispatch Location.

Stationary device: An stationary device is not mobile and is connected to the network at a fixed, long-term-stable physical location. Examples include home PCs or pay phones.

Uncertainty: Uncertainty is an estimate, expressed in a unit of length, indicating the diameter of a circle that contains the endpoint with the probability indicated by the confidence value.

2. Introduction

Requesting help in an emergency using a communications device such as a telephone or mobile phone is an accepted practice in many parts of the world. As communications devices increasingly utilize the Internet to interconnect and communicate, users will expect to use such devices to request help. This document describes establishment of a communications session by a user to a "Public Safety Answering Point" (PSAP), that is, a call center established by response agencies to accept emergency calls. Such citizen/visitor-to-authority calls can be distinguished from those that are created by responders (authority-to-authority) using public communications infrastructure often involving some kind of priority access as defined in Emergency Telecommunications Service (ETS) in IP Telephony [RFC4190]. They also can be distinguished from emergency warning systems that are authority-to-citizen.

Supporting emergency calling requires cooperation by a number of elements, their vendors and service providers. This document discusses how end device and applications create emergency calls, how access networks supply location for some of these devices, how service providers assist the establishment and routing, and how PSAPs receive calls from the Internet.

The emergency response community will have to upgrade their facilities to support a wider range of communications services, but cannot be expected to handle wide variations in device and service capability. New devices and services are being made available that could be used to make a request for help that are not traditional telephones, and users are increasingly expecting to use them to place emergency calls. However, many of the technical advantages of

Internet multimedia require re-thinking of the traditional emergency calling architecture. This challenge also offers an opportunity to improve the operation 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 (Public Switched Telephone Network) and IP-based telephony, but calling on the Internet is characterized by:

- o the interleaving over the same infrastructure of a wider variety of services;
- o the separation of the access provider from the application provider;
- o media other than voice (for example, video and text in several forms);
- o the potential mobility of all end systems, including endpoints nominally thought of as fixed systems and not just those using radio access technology. For example, consider a wired phone connected to a router using a mobile data network such as EV-DO as an uplink.

This document focuses on how devices using the Internet can place emergency calls and how PSAPs can handle Internet multimedia emergency calls natively, rather than describing how circuit-switched PSAPs can handle VoIP calls. 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. There are many legacy telephone networks that will persist long after most systems have been upgraded to IP origination and termination of emergency calls. Many of these legacy systems route calls based on telephone numbers. Gateways and conversions between existing systems and newer systems defined by this document will be required. Since existing systems are governed primarily by local government regulations and national standards, the gateway and conversion details will be governed by national standards and thus are out of scope for this document.

Existing emergency call systems are organized locally or nationally; there are currently few international standards. However, the Internet crosses national boundaries, and thus Internet standards are required. To further complicate matters, VoIP endpoints can be connected through tunneling mechanisms such as virtual private networks (VPNs). Tunnels can obscure the identity of the actual access network that knows the location. 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 not created 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 discusses the use of the Session Initiation Protocol (SIP) [RFC3261] by PSAPs and calling parties. While other inter-domain call signaling protocols may be used for emergency calling, SIP is ubiquitous and possesses the proper support of this use case. Only protocols such as H.323, XMPP/Jingle, ISUP and SIP are suitable for inter-domain communications, ruling out Media Gateway Controller protocols such as MGCP or H.248/Megaco. The latter protocols can 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 inter-domain VoIP calls would be placed. Other signaling protocols may also use protocol translation to communicate with a SIP-enabled PSAP. p2psip is not considered in this document.

Existing emergency services rely exclusively on voice and conventional text telephony ("TTY") media streams. However, more choices of media offer additional ways to communicate and evaluate the situation as well as to assist callers and call takers in handling emergency calls. For example, instant messaging and video could improve the ability to communicate and evaluate the situation and to 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].

This document focuses on the case in which all three steps in the emergency calling process -- location configuration, call routing, and call placement - can be and are performed by the calling endpoint, with the endpoint's Access Service Provider supporting the process by providing location information. Calls in this case may be routed via an application-layer Communications Service Provider (e.g., a Voice Service Provider), but need not be. The underlying protocols can also be used to support other models in which parts of the process are delegated to the Communications Service Provider. This document does not address in detail either these models or interoperability issues between them and the model described here.

Since this document is a framework document, it does not include normative behavior. A companion document, [I-D.ietf-ecrit-phonebcp], describes Best Current Practice for this subject and contains

normative language for devices, access and calling network elements.

Supporting emergency calling does not require any specialized SIP header fields, request methods, status codes, message bodies, or event packages, but does require that existing mechanisms be used in certain specific ways, as described below. User Agents (UAs) unaware of the recommendations in this draft may be able to place emergency calls, but functionality may be impaired. For example, if the UA does not implement the location mechanisms described, an emergency call may not be routed to the correct PSAP, and if the caller is unable to supply his exact location, dispatch of emergency responders may be delayed. Suggested behavior for both endpoints and servers is provided.

From the point of view of the PSAP, three essential elements characterize an emergency call:

- o The call is routed to the most appropriate PSAP, based primarily on the location of the caller.
- o The PSAP must be able to automatically obtain the location of the caller with sufficient accuracy to dispatch a responder to help the caller.
- o The PSAP must be able to re-establish a session to the caller if for any reason the original session is disrupted.

3. Overview of how emergency calls are placed

An emergency call can be distinguished (Section 5) from any other call by a unique Service URN [RFC5031] that is placed in the call set-up signaling when a home or visited emergency dial string is detected. Because emergency services are local to specific geographic regions, a caller obtains his location (Section 6) prior to making emergency calls. To get this location, either a form of measuring, for example, GNSS (Section 6.2.3) is deployed, or the endpoint is configured (Section 6.5) with its location from the access network's Location Information Server (LIS) using a Location Configuration Protocol (LCP). The location is conveyed (Section 6.7) in the SIP signaling with the call. The call is routed (Section 8) based on location using the LoST protocol [RFC5222], which maps a location to a set of PSAP URIs. Each URI resolves to a PSAP or an Emergency Services Routing Proxy (ESRP) that serves as an incoming proxy for a group of PSAPs. The call arrives at the PSAP with the location included in the INVITE request.

The following is a quick overview for a typical Ethernet connected telephone using SIP signaling. It illustrates one set of choices for various options presented later in this document.

- o The phone "boots" and connects to its access network.
- o The phone gets location via a Location Configuration Protocol (LCP), for example from the DHCP server in civic [RFC4776] and/or geo [RFC6225] forms, a HELD server [RFC5985] or the first level switch's LLDP server [LLDP].
- o The phone obtains the local emergency dial string(s) from the LoST [RFC5222] server for its current location. It also receives and caches the PSAP URI obtained from the LoST server.
- o Some time later, the user places an emergency call. The phone recognizes an emergency call from the dial strings and uses the "urn:service:sos" [RFC5031] URN to mark an emergency call.
- o It refreshes its location via DHCP and updates the PSAP's URI by querying the LoST mapping server with its location.
- o It puts its location in the SIP INVITE request in a Geolocation header [I-D.ietf-sip-location-conveyance] and forwards the call using its normal outbound call processing, which commonly involves an outbound proxy.
- o The proxy recognizes the call as an emergency call and routes the call using normal SIP routing mechanisms to the URI specified.
- o The call routing commonly traverses an incoming proxy server (ESRP) in the emergency services network. That proxy would route the call to the PSAP.
- o The call is established with the PSAP and mutually agreed upon media streams are created.
- o The location of the caller is displayed to the call taker.

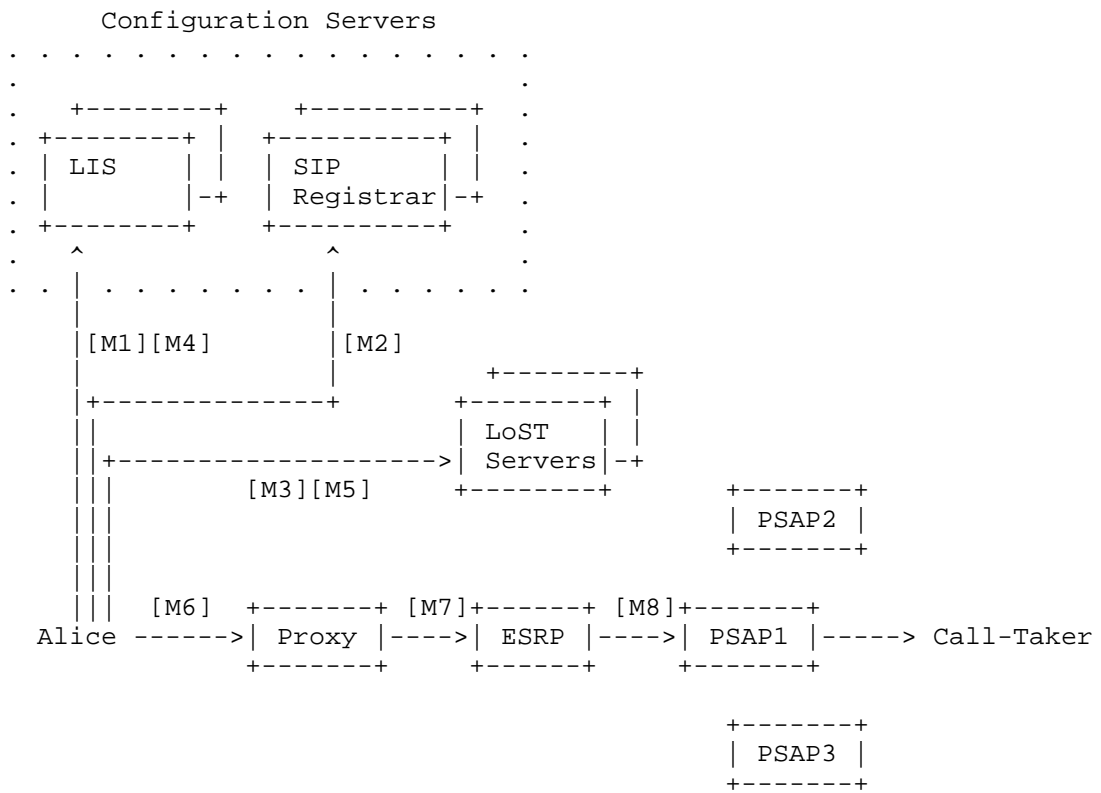


Figure 1: Emergency Call Component Topology

The typical message flow for this example using Alice as the caller:

```
[M1] Alice -> LIS: LCP Request(s) (ask for location)
      LIS -> Alice: LCP Reply(s) (replies with location)
[M2] Alice -> Registrar: SIP REGISTER
      Registrar -> Alice: SIP 200 OK (REGISTER)
[M3] Alice -> LoST Server: Initial LoST Query (contains location)
      Lost Server -> Alice: Initial LoST Response (contains
                          PSAP-URI and dial string)
```

Some time later, Alice dials or otherwise initiates an emergency call:

```
[M4] Alice -> LIS: LCP Request (update location)
      LIS -> Alice: LCP Reply (replies with location)
[M5] Alice -> LoST Server: Update LoST Query (contains location)
      Lost Server -> Alice: LoST Response (contains PSAP-URI)
[M6] Alice -> Outgoing Proxy: SIP INVITE (service URN,
                                          Location and PSAP URI)
[M7] Outgoing Proxy -> ESRP: SIP INVITE (service URN,
                                          Location and PSAP URI)
[M8] ESRP -> PSAP: SIP INVITE (service URN, Location and PSAP URI)
```

The 200 OK response is propagated back from the PSAP to Alice and the ACK response is propagated from Alice to the PSAP.

Figure 2: Message Flow

Figure 1 shows emergency call component topology and the text above shows call establishment. These include the following components:

- o Alice - who places the emergency call.
- o Configuration Servers - Servers providing Alice's UA its IP address and other configuration information, perhaps including location by-value or by-reference. Configuration servers also may include a SIP registrar for Alice's UA. Most SIP UAs will register, so it will be a common scenario for UAs that make emergency calls to be registered with such a server in the originating calling network. Registration would be required for the PSAP to be able to call back after an emergency call is completed. All the configuration messages are labeled M1 through M3, but could easily require more than 3 messages to complete.
- o LoST server - Processes the LoST request for location plus a Service URN to a PSAP-URI, either for an initial request from a UA, or an in-call routing by the proxy server in the originating network, or possibly by an ESRP.
- o ESRP - Emergency Services Routing Proxy, a SIP proxy server that is the incoming call proxy in the emergency services domain. The ESRP makes further routing decisions (e.g., based on PSAP state and the location of the caller) to choose the actual PSAP that handles the call. In some jurisdictions, this may involve another

LoST query.

- o PSAP - Emergency calls are answered at a Public Safety Answering Point, a call center.

Generally, Alice's UA either has location configured manually, has an integral location measurement mechanism, or it runs a LCP [M1] to obtain location from the access (broadband) network. Alice's UA then will most likely register [M2] with a SIP registrar. This allows her to be contacted by other SIP entities. Next, her UA will perform an initial LoST query [M3] to learn a URI for use if the LoST query fails during an emergency call, or to use to test the emergency call mechanism. The LoST response contains the dial string for emergency calls appropriate for the location provided.

At some time after her device has booted, Alice initiates an emergency call. She may do this by dialing an emergency dial string valid for her current ("local") location, or for her "home" location.

The UA recognizes the dial string. The UA attempts to refresh its location [M4], and with that location, to refresh the LoST mapping [M5], in order to get the most accurate information to use for routing the call. If the location request or the LoST request fails, or takes too long, the UA uses values it has cached.

The UA creates a SIP INVITE [M6] request that includes the location. [I-D.ietf-sip-location-conveyance] defines a SIP Geolocation header that contains either a location-by-reference URI or a [RFC3986] "cid" URL indicating where in the message body the location-by-value is.

The INVITE message is routed to the ESRP [M7], which is the first inbound proxy for the emergency services domain. This message is then routed by the ESRP towards the most appropriate PSAP for Alice's location [M8], as determined by the location and other information.

A proxy in the PSAP chooses an available call taker and extends the call to its UA.

The 200 OK response to the INVITE request traverses the path in reverse, from call taker UA to PSAP proxy to ESRP to originating network proxy to Alice's UA. The ACK request completes the call set-up and the emergency call is established, allowing the PSAP call-taker to talk to Alice about Alice's emergency.

4. Which devices and services should support emergency calls

Current PSAPs support voice calls and real-time text calls placed through PSTN facilities or systems connected to the PSTN. Future

PSAPs will however support Internet connectivity and a wider range of media types and provide higher functionality. In general, if a user could reasonably expect to be able to place a call for help with the device, then the device or service should support emergency calling. Certainly, any device or service that looks like and works like a telephone (wired or mobile) should support emergency calling, but increasingly, users have expectations that other devices and services should work.

Devices that create media sessions and exchange audio, video and/or text, and have the capability to establish sessions to a wide variety of addresses, and communicate over private IP networks or the Internet, should support emergency calls.

Traditionally, enterprise support of emergency calling is provided by the telephony service provider to the enterprise. In some more recent systems, the enterprise PBX assists emergency calling by providing more fine grained location in larger enterprises. In the future, the enterprise may provide the connection to emergency services itself, not relying on the telephony service provider.

5. Identifying an emergency call

Using the PSTN, emergency help can often be summoned by dialing a nationally designated, widely known number, regardless of where the telephone was purchased. The appropriate number is determined by the infrastructure the telephone is connected to. However, this number differs between localities, even though it is often the same for a country or region, as it is in many countries in the European Union. In some countries, there is only one uniform digit sequence that is used for all types of emergencies. In others, there are several sequences that are specific to the type of responder needed, e.g., one for police, another for fire. For end systems, on the other hand, it is desirable to have a universal identifier, independent of location, to allow the automated inclusion of location information and to allow the device and other entities in the call path to perform appropriate processing within the signaling protocol in an emergency call set-up.

Since there is no such universal identifier, as part of the overall emergency calling architecture, common emergency call URNs are defined in [RFC5031]. For a single number environment the urn is "urn:service:sos". Users are not expected to "dial" an emergency URN. Rather, appropriate emergency dial strings are translated to corresponding service URNs, carried in the Request-URI of the INVITE request. Such translation is best done by the endpoint, because, among other reasons, emergency calls convey location in the

signaling, but non-emergency calls do not normally do that. If the device recognizes the emergency call, it can include location. A signaling intermediary (proxy server) can also recognize emergency dial strings if the endpoint fails to do so.

For devices that are mobile or nomadic, an issue arises of whether the home or visited dial strings should be used. Many users would prefer that their home dialing sequences work no matter where they are. However, local laws and regulations may require that the visited dialing sequence(s) work. Therefore, the visited dial string must work. Devices may have a way to be configured or learn home dial strings.

LoST [RFC5222] provides the mechanism for obtaining the dialing sequences for a given location. LoST servers must return dial strings for emergency services. If the endpoint does not support the translation of dial strings to service URNs, the dialing sequence from the endpoint to its proxy is represented as a dial string [RFC4967] and the outgoing proxy must recognize the dial string and translate it to the equivalent service URN. To determine the local emergency dial string, the proxy needs the location of the endpoint. This may be difficult in situations where the user can roam or be nomadic. Endpoint recognition of emergency dial strings is therefore preferred. If a service provider is unable to guarantee that it can correctly determine local emergency dialstrings, wherever its subscribers may be, then it is required that the endpoint do the recognition.

Note: The emergency call practitioners consider it undesirable to have a single button emergency call user interface element. These mechanisms tend to result in a very high rate of false or accidental emergency calls. In order to minimize this issue, practitioners recommend that device should only initiate emergency calls based on entry of specific emergency call dial strings. Speed dial mechanisms may effectively create single button emergency call invocation and practitioners recommend they not be permitted.

6. Location and its role in an emergency call

Location is central to the operation of emergency services. Location is used for two purposes in emergency call handling: routing of the call and dispatch of responders. It is frequently the case that the caller reporting an emergency is unable to provide a unique, valid location themselves. For this reason, location provided by the endpoint or the access network is needed. For practical reasons, each PSAP generally handles only calls for a certain geographic area, with overload arrangements between PSAPs to handle each others'

calls. Other calls that reach it by accident must be manually re-routed (transferred) to the most 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 decentralize PSAP responsibilities to the level of counties or municipalities.

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. Irregular boundaries are common, often for historical reasons. Routing must be done based on actual PSAP service boundaries -- the closest PSAP, or the PSAP that serves the nominal city name provided in the location, may not be the correct PSAP.

Accuracy of routing location is a complex subject. Calls must be routed quickly, but accurately, and location determination is often a time/accuracy tradeoff, especially with mobile devices or self measuring mechanisms. If more accurate routing location is not available it is considered acceptable to base a routing decision on an accuracy equal to the area of one sector of a mobile cell site.

Routing to the most appropriate PSAP is always based on the location of the caller, despite the fact that some emergency calls are placed on behalf of someone else, and the location of the incident is sometimes not the location of the caller. In some cases, there are other factors that enter into the choice of the PSAP that gets the call, such as time of day, caller media requests and language preference and call load. However, location of the caller is the primary input to the routing decision.

Many mechanisms used to locate a caller have a relatively long "cold start" time. To get a location accurate enough for dispatch may take as much as 30 seconds. This is too long to wait for emergencies. Accordingly, it is common, especially in mobile systems, to use a coarse location, for example, the cell site and sector serving the call, for call routing purposes, and then to update the location when a more precise value is known prior to dispatch. In this document we use "routing location" and "dispatch location" when the distinction matters.

Accuracy of dispatch location is sometimes determined by local regulation, and is constrained by available technology. The actual requirement is more stringent than available technology can deliver: It is required that a device making an emergency call close to the "demising" or separation wall between two apartments in a high rise apartment building report location with sufficient accuracy to determine on what side of the wall it is on. This implies perhaps a

3 cm accuracy requirement. As of the date of this memo, assisted GNSS uncertainty in mobile phones with 95% confidence cannot be relied upon to be less than hundreds of meters. As technology advances, the accuracy requirements for location will need to be tightened. Wired systems using wire tracing mechanisms can provide location to a wall jack in specific room on a floor in a building, and may even specify a cubicle or even smaller resolution. As this discussion illustrates, emergency call systems demand the most stringent location accuracy available.

In Internet emergency calling, where the endpoint is located is determined using a variety of measurement or wire-tracing methods. Endpoints may be configured with their own location by the access network. In some circumstances, a proxy server may insert location into the signaling on behalf of the endpoint. The location is mapped to the URI to send the call to, and the location is conveyed to the PSAP (and other elements) in the signaling. The terms 'determination', 'configuration', 'mapping', and 'conveyance' are used for specific aspects of location handling in IETF protocols. Likewise, we employ Location Configuration Protocols, Location Mapping Protocols, and Location Conveyance Protocols for these functions.

This document provides guidance for generic network configurations with respect to location. It is recognized that unique issues may exist in some network deployments. The IETF will continue to investigate these unique situations and provide further guidance, if warranted, in future documents.

6.1. Types of location information

Location can be specified in several ways:

Civic: Civic location information describes the location of a person or object by a street address that corresponds to a building or other structure. Civic location may include more fine grained location information such as floor, room and cubicle. Civic information comes in two forms:

'Jurisdictional': refers to a civic location using actual political subdivisions, especially for the community name.

'Postal': refers to a civic location for mail delivery. The name of the post office sometimes does not correspond to the community name and a postal address may contain post office boxes or street addresses that do not correspond to an actual building. Postal addresses are generally unsuitable for emergency call dispatch because the post office conventions (for community name, for example) do not match those known by the responders. The fact that they are unique can sometimes be exploited to provide a mapping between a postal address and a

civic address suitable to dispatch a responder to. In IETF location protocols, there is an element (Postal Community Name) that can be included in a location to provide the post office name as well as the actual jurisdictional community name. There is also an element for a postal code. There is no other accommodation for postal addresses in these protocols.

Geospatial (geo): Geospatial addresses contain longitude, latitude and altitude information based on an understood datum and earth shape model (datum). While there have been many datums developed over time, most modern systems are using or moving towards the WGS84 [WGS84] datum.

Cell tower/sector: Cell tower/sector is often used for identifying the location of a mobile handset, especially for routing of emergency calls. Cell tower and sectors identify the cell tower and the antenna sector that a mobile device is currently using. Traditionally, the tower location is represented as a point chosen to be within a certain PSAP service boundary who agrees to take calls originating from that tower/sector, and routing decisions are made on that point. Cell/sector information could also be represented as an irregularly shaped polygon of geospatial coordinates reflecting the likely geospatial location of the mobile device. Whatever representation is used must route correctly in the LoST database, where "correct" is determined by local PSAP management.

In IETF protocols, both civic and geospatial forms are supported. The civic forms include both postal and jurisdictional fields. A cell tower/sector can be represented as a geo point or polygon or civic location. Other forms of location representation must be mapped into either a geo or civic for use in emergency calls.

For emergency call purposes, conversion of location information from civic to geo or vice versa prior to conveyance is not desirable. The location should be sent in the form it was determined. Conversion between geo and civic requires a database. Where PSAPs need to convert from whatever form they receive to another for responder purposes, they have a suitable database. However, if a conversion is done before the PSAP's, and the database used is not exactly the same one the PSAP uses, the double conversion has a high probability of introducing an error.

6.2. Location determination

As noted above, location information can be entered by the user or installer of a device ("manual configuration"), measured by the end system, can be delivered to the end system by some protocol or measured by a third party and inserted into the call signaling.

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. Although self measured location (e.g., GNSS) is attractive, location information provided by the access network could be much more accurate, and more reliable in some environments such as high rise buildings in dense urban areas.

The closer an entity is to the source of location, the more likely it is able to determine which location is most appropriate for a particular purpose when there are more than one location determinations for a given endpoint. In emergency calling, the PSAP is the least likely to be able to appropriately choose which location to use when multiple conflicting locations are presented to it. While all available locations can be sent towards the PSAP, the order of the locations should be the sender's best attempt to guide the recipient of which one(s) to use.

6.2.1. User-entered location information

Location information can be maintained by the end user or the installer of an endpoint in the endpoint itself, or in a database.

Location information routinely provided by end users is almost always less reliable than measured or wire database information, as users may mistype location information or may enter civic address information that does not correspond to a recognized (i.e., valid, see Section 6.10) address. Users can forget to change the data when the location of a device changes.

However, there are always a small number of cases where the automated mechanisms used by the access network to determine location fail to accurately reflect the actual location of the endpoint. For example, the user may deploy his own WAN behind an access network, effectively removing an endpoint some distance from the access network's notion of its location. To handle these exceptional cases, there must be some mechanism provided to manually provision a location for an endpoint by the user or by the access network on behalf of a user. The use of the mechanism introduces the possibility of users falsely declaring themselves to be somewhere they are not. However, this is generally not a problem in practice. Commonly, if an emergency caller insists that he is at a location different from what any automatic location determination system reports he is, responders will always be sent to the user's self-declared location.

6.2.2. Access network "wire database" location information

Location information can be maintained by the access network, relating some form of identifier for the end subscriber or device to a location database ("wire database"). In enterprise LANs, wiremap databases map Ethernet switch ports to building locations. In DSL installations, the local telephone carrier maintains a mapping of wire-pairs to subscriber addresses.

Accuracy of location historically has been to a street address level. However, this is not sufficient for larger structures. The PIDF Location Object [RFC4119] extended by [RFC5139] and [RFC5491] permits interior building/floor/room and even finer specification of location within a street address. When possible, interior location should be supported.

The threshold for when interior location is needed is approximately 650 square meters. This value is derived from USA fire brigade recommendations of spacing of alarm pull stations. However, interior space layout, construction materials and other factors should be considered.

Even for IEEE 802.11 wireless access points, wire databases may provide sufficient location resolution. The location of the access point as determined by the wiremap may be supplied as the location for each of the clients of the access point. However, this may not be true for larger-scale systems such as IEEE 802.16 (WiMAX) and IEEE 802.22 that typically have larger cells than those of IEEE 802.11. The civic location of an IEEE 802.16 base station may be of little use to emergency personnel, since the endpoint could be several kilometers away from the base station.

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 validation (see Section 6.10), similar to the current system in some jurisdictions.

6.2.3. End-system measured location information

Global Positioning System (GPS) and similar Global Navigation Satellite Systems (e.g., GLONAS and Galileo) receivers may be embedded directly in the end device. GNSS produces relatively high precision location fixes in open-sky conditions, but the technology still faces several challenges in terms of performance (time-to-fix and time-to-first-fix), as well as obtaining successful location fixes within shielded structures, or underground. It also requires all devices to be equipped with the appropriate GNSS capability.

Many mobile devices require using some kind of "assist", that may be operated by the access network (A-GPS) or by a government (WAAS). A device may be able to use multiple sources of assist data.

GNSS systems may be always enabled and thus location will always be available accurately immediately (assuming the device can "see" enough satellites). Mobile devices may not be able to sustain the power levels required to keep the measuring system active. In such circumstances, when location is needed, the device has to start up the measurement mechanism. This typically takes tens of seconds, far too long to wait to be able to route an emergency call. For this reason, devices that have end-system measured location mechanisms but need a cold start period lasting more than a couple seconds need another way to get a routing location. Typically this would be a location associated with a radio link (cell site/sector).

6.2.4. Network measured location information

The access network may locate end devices. Techniques various forms of triangulation. Elements in the network infrastructure triangulate end systems based on signal strength, angle of arrival or time of arrival. Common mechanisms deployed include:

- o Time Difference Of Arrival - TDOA
- o Uplink Time Difference Of Arrival - U-TDOA
- o Angle of Arrival - AOA
- o RF fingerprinting
- o Advanced Forward Link Trilateration - AFLT
- o Enhanced Forward Link Trilateration - EFLT

Sometimes multiple mechanisms are combined, for example A-GPS with AFLT.

6.3. Who adds location, endpoint or proxy

The IETF emergency call architecture prefers endpoints to learn their location and supply it on the call. Where devices do not support location, proxy servers may have to add location to emergency calls. Some calling networks have relationships with all access networks the device may be connected to, and that may allow the proxy to accurately determine the location of the endpoint. However, NATs and other middleboxes often make it impossible to determine a reference identifier the access network could provide to a LIS to determine the location of the device. Systems designers are discouraged from relying on proxies to add location. The technique may be useful in some limited circumstances as devices are upgraded to meet the requirements of this document, or where relationships between access networks and calling networks are feasible and can be relied upon to get accurate location.

Proxy insertion of location complicates dial string recognition. As noted in Section 6, local dial strings depend on the location of the caller. If the device does not know its own location, it cannot use the LoST service to learn the local emergency dial strings. The calling network must provide another way for the device to learn the local dial string, and update it when the user moves to a location where the dial string(s) change, or do the dial string determination itself.

6.4. Location and references to location

Location information may be expressed as the actual civic or geospatial value but can be transmitted as by value (wholly contained within the signaling message) or by reference (i.e., as a URI pointing to the value residing on a remote node waiting to be dereferenced).

When location is transmitted by value, the location information is available to entity in the call path. On the other hand, location objects can be large, and only represent a single snapshot of the device's location. Location references are small and can be used to represent a time-varying location, but the added complexity of the dereference step introduces a risk that location will not be available to parties that need it.

6.5. End system location configuration

Unless a user agent has access to provisioned or locally measured location information, it must obtain it from the access network. There are several location configuration protocols (LCPs) that can be used for this purpose including DHCP, HELD and LLDP:

DHCP can deliver civic [RFC4776] or geospatial [RFC6225] information. User agents need to support both formats. Note that a user agent can use DHCP, via the DHCP REQUEST or INFORM messages, even if it uses other means to acquire its IP address.

HELD [RFC5985] can deliver a civic or geo location object, by value or by reference, via a layer 7 protocol. The query typically uses the IP address of the requester as an identifier and returns the location value or reference associated with that identifier. HELD is typically carried in HTTP.

Link-Layer Discovery Protocol [LLDP] with Media Endpoint Device extensions [LLDP-MED] can be used to deliver location information directly from the Layer 2 network infrastructure, and also supports both civic and geo formats identical in format to DHCP methods.

Each LCP has limitations in the kinds of networks that can reasonably support it. For this reason, it is not possible to choose a single

mandatory-to-deploy LCP. For endpoints with common network connections (such as an Ethernet jack or a WiFi connection) serious incompatibilities would ensue unless every network supported every protocol, or alternatively, every device supported every protocol. For this reason, a mandatory-to-implement list of LCPs is established in [I-D.ietf-ecrit-phonebcp]. Every endpoint that could be used to place emergency calls must implement all of the protocols on the list. Every access network must deploy at least one of them. Since it is the variability of the networks that prevent a single protocol from being acceptable, it must be the endpoints that implement all of them, and to accommodate a wide range of devices, networks must deploy at least one of them.

Often, network operators and device designers believe that they have a simpler environment and some other network specific mechanism can be used to provide location. Unfortunately, it is very rare to actually be able to limit the range of devices that may be connected to a network. For example, existing mobile networks are being used to support routers and LANs behind a wireless data network WAN connection, with Ethernet connected phones connected to that. It is possible that the access network could support a protocol not on the list, and require every handset in that network to use that protocol for emergency calls. However, the Ethernet-connected phone won't be able to acquire location, and the user of the phone is unlikely to be dissuaded from placing an emergency call on that phone. The widespread availability of gateways, routers and other network-broadening devices means that indirectly connected endpoints are possible on nearly every network. Network operators and vendors are cautioned that shortcuts to meeting this requirement are seldom successful.

Location for non-mobile devices is normally expected to be acquired at network attachment time and retained by the device. It should be refreshed when the cached value expires. For example, if DHCP is the acquisition protocol, refresh of location may occur when the IP address lease is renewed. At the time of an emergency call, the location should be refreshed, with the retained location used if the location acquisition does not immediately return a value. Mobile devices may determine location at network attachment time and periodically thereafter as a backup in case location determination at the time of call does not work. Mobile device location may be refreshed when a TTL expires or the device moves beyond some boundaries (as provided by [RFC5222]). Normally, mobile devices will acquire its location at call time for use in an emergency call routing. See Section 6.8 for a further discussion on location updates for dispatch location.

There are many examples of endpoints which are user agent

applications running on a more general purpose device, such as a personal computer. On some systems, layer 2 protocols like DHCP and LLDP may not be directly accessible to applications. It is desirable for an operating system to have an API which provides the location of the device for use by any application, especially those supporting emergency calls.

6.6. When location should be configured

Devices should get routing location immediately after obtaining local network configuration information. The presence of NAT and VPN tunnels (that assign new IP addresses to communications) can obscure identifiers used by LCPs to determine location, especially for HELD. In some cases, such as residential NAT devices, the NAT is placed between the endpoint and the access network demarcation point and thus the IP address seen by the access network is the right identifier for location of the residence. However, in many enterprise environments, VPN tunnels can obscure the actual IP address. Some VPN mechanisms can be bypassed so that a query to the LCP can be designated to go through the direct IP path, using the correct IP address, and not through the tunnel. In other cases, no bypass is possible, but location can be configured before the VPN is established. Of course, LCPs that use layer 2 mechanisms (DHCP Location options and LLDP-MED) are usually immune from such problems because they do not use the IP address as the identifier for the device seeking location.

It is desirable that routing location information be periodically refreshed. A LIS supporting a million subscribers each refreshing once per day would need to support a query rate of $1,000,000 / (24 * 60 * 60) = 12$ queries per second. For networks with mobile devices, much higher refresh rates could be expected.

It is desirable for routing location information to be requested immediately before placing an emergency call. However, if there is any significant delay in getting more recent location, the call should be placed with the most recent location information the device has. In mobile handsets, routing is often accomplished with the cell site and sector of the tower serving the call, because it can take many seconds to start up the location determination mechanism and obtain an accurate location.

There is a tradeoff between the time it takes to get a routing location and the accuracy (technically, confidence and uncertainty) obtained. Routing an emergency call quickly is required. However, if location can be substantially improved by waiting a short time (e.g., for some sort of "quick fix"), it's preferable to wait. Three seconds, the current nominal time for a quick fix, is a very long

time add to post dial delay.

NENA recommends [NENAI3TRD] that IP based systems complete calls in two seconds from last dial press to ring at PSAP.

6.7. Conveying location

When an emergency call is placed, the endpoint should include location in the call signaling. This is referred to as "conveyance" to distinguish it from "configuration". In SIP, the location information is conveyed following the procedures in [I-D.ietf-sip-location-conveyance]. Since the form of the location information obtained by the acquisition protocol may not be the same as the conveyance protocol uses (PIDF-LO [RFC4119]), mapping by the endpoint from the LCP form to PIDF may be required.

6.8. Location updates

As discussed above, it may take some time for some measurement mechanisms to get a location accurate enough for dispatch, and a routing location with less accuracy may be provided to get the call established quickly. The PSAP needs the dispatch location before it sends the call to the responder. This requires an update of the location. In addition, the location of some mobile callers, e.g., in a vehicle or aircraft, can change significantly during the emergency call.

A PSAP has no way to request an update of a location provided by value. If the UAC gets new location, it must signal the PSAP using a new INVITE or an UPDATE transaction with a new Geolocation header to supply the new location.

With the wide variation in determination mechanisms, the PSAP does not know when accurate location may be available. The preferred mechanism is that the LIS notifies the PSAP when an accurate location is available rather than requiring a poll operation from the PSAP to the LIS. The SIP Presence subscription [RFC3856] provides a suitable mechanism.

When using a HELD dereference, the PSAP must specify the value "emergencyDispatch" for the ResponseTime parameter. Since typically the LIS is local relative to the PSAP, the LIS can be aware of the update requirements of the PSAP

6.9. Multiple locations

Getting multiple locations all purported to describe the location of the caller is confusing to all, and should be avoided. Handling

multiple locations at the point where a PIDF is created is discussed in [RFC5491]. Conflicting location information is particularly harmful if different routes (PSAPs) result from LoST queries for the multiple locations. When they occur anyway, the general guidance is that the entity earliest in the chain generally has more knowledge than later elements to make an intelligent decision, especially about which location will be used for routing. It is permissible to send multiple locations towards the PSAP, but the element that chooses the route must select exactly one location to use with LoST.

Guidelines for dealing with multiple locations are also given in [RFC5222]. If a UA gets multiple locations, it must choose the one to use for routing, but it may send all of the locations it has in the signaling. If a proxy is inserting location and has multiple locations, it must choose exactly one to use for routing, marking it as such (per [I-D.ietf-sip-location-conveyance]), and send it as well as any others it has.

The UA or proxy should have the ability to understand how and from whom it learned its location, and should include this information in the location objects that are sent to the PSAP. That labeling provides the call-taker with information to make decisions upon, as well as guidance for what to ask the caller and what to tell the responders.

Endpoints or proxies may be tempted to send multiple versions of the same location. For example a database may be used to "geocode" or "reverse geocode", that is, convert from civic to geo or vice versa. It is very problematic to use derived locations in emergency calls. The PSAP and the responders have very accurate databases which they use to convert, most commonly from a reported geo to a civic suitable for dispatching responders. If one database is used to convert from, say, civic to geo, and another converts from geo to civic, errors will often occur where the databases are slightly different. "Off by one" errors are serious when responders go to the wrong location. Derived locations should be marked with a "derived" method token [RFC4119]. If an entity gets a location which has a measured or other original method, and another with a derived method, it must use the original value for the emergency call.

6.10. Location validation

Validation in this context means both that there is a mapping from the address to a PSAP and that the PSAP understands how to direct responders to the location. It is recommended that location be validated prior to a device placing an actual emergency call; some jurisdictions require that this be done.

Determining the addresses that are valid can be difficult. There are, for example, many cases of two names for the same street, or two streets with the same name, but different "suffixes" (Avenue, Street, Circle) in a city. In some countries, the current system provides validation. For example, in the United States of America, the Master Street Address Guide (MSAG) records all valid street addresses and is used to ensure that the service addresses in phone billing records correspond to valid emergency service street addresses. Validation is normally only a concern for civic addresses, although there could be some determination that a given geo is within at least one PSAP service boundary; that is, a "valid" geo is one where there is a mapping in the LoST server.

LoST [RFC5222] includes a location validation function. Validation is normally performed when a location is entered into a Location Information Server. It should be confirmed periodically, because the mapping database undergoes slow change and locations which previously validated may eventually fail validation. Endpoints may wish to validate locations they receive from the access network, and will need to validate manually entered locations. Proxies that insert location may wish to validate locations they receive from a LIS. When the test functions (Section 15) are invoked, the location used should be validated.

When validation fails, the location given should not be used for an emergency call, unless no other valid location is available. Bad location is better than no location. If validation is completed when location is first loaded into a LIS, any problems can be found and fixed before devices could get the bad location. Failure of validation arises because an error is made in determining the location, although occasionally the LoST database is not up to date or has faulty information. In either case, the problem must be identified and should be corrected before using the location.

6.11. Default location

Occasionally, the access network cannot determine the actual location of the caller. In these cases, it must supply a default location. The default location should be as accurate as the network can determine. For example, in a cable network, a default location for each Cable Modem Termination System (CMTS), with a representative location for all cable modems served by that CMTS could be provided if the network is unable to resolve the subscriber to anything more precise than the CMTS. Default locations must be marked as such so that the PSAP knows that the location is not accurate.

6.12. Location format conversion

The endpoint is responsible for mapping any form of location it receives from an LCP into PIDF-LO form if the LCP did not directly return a PIDF-LO.

7. LIS and LoST discovery

Endpoints must be able to discover a LIS if the HELD protocol is used, and a LoST server. DHCP options are defined for this purpose, namely [RFC5986] and [RFC5223].

Until such DHCP records are widely available, it may be necessary for the service provider to provision a LoST server address in the device. The endpoint can also do a DNS SRV query to find a LoST server. In any environment, more than one of these mechanisms may yield a LoST server, and they may be different. The recommended priority is DHCP first, provisioned value second, and DNS SRV query in the SIP domain third.

8. Routing the call to the PSAP

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.

Type of emergency service: In some jurisdictions, emergency calls for specific emergency services such as fire, police, ambulance or mountain rescue are directed to just those emergency-specific PSAPs. This mechanism is supported by marking emergency calls with the proper service identifier [RFC5031]. Even in single number jurisdictions, not all services are dispatched by PSAPs and may need alternate URNs to route calls to the appropriate call center.

Media capabilities of caller: In some cases, emergency call centers for specific caller media preferences, such as typed text or video, are separate from PSAPs serving voice calls. ESRPs are expected to be able to provide routing based on media. 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, where routing within the PSAP based on the media

offer would be provided.

Providing a URL to route emergency calls by location and by type of service is the primary function LoST [RFC5222] provides. LoST accepts a query with location (by-value) in either civic or geo form, plus a service identifier, and returns a URI (or set of URIs) to route the call to. Normal SIP [RFC3261] routing functions are used to resolve the URI to a next hop destination.

The endpoint can complete the LoST mapping from its location at boot time, and periodically thereafter. It should attempt to obtain a "fresh" location, and from that a current mapping when it places an emergency call. If accessing either its location acquisition or mapping functions fail, it should use its cached value. The call would follow its normal outbound call processing.

Determining when the device leaves the area provided by the LoST service can tax small mobile devices. For this reason, the LoST server should return a simple (small number of points) polygon for geospatial location. This can be a simple enclosing rectangle of the PSAP service area when the reported point is not near an edge, or a smaller polygonal edge section when the reported location is near an edge. Civic location is uncommon for mobile devices, but reporting that the same mapping is good within a community name, or even a street, may be very helpful for WiFi connected devices that roam and obtain civic location from the access point they are connected to.

Networks that support devices that do not implement LoST mapping themselves may need the outbound proxy do the mapping. If the endpoint recognized the call was an emergency call, provided the correct service URN and/or included location on the call in a Geolocation header, a proxy server could easily accomplish the mapping.

However, if the endpoint did not recognize the call was an emergency call, and thus did not include location, the proxy's task is more difficult. It is often difficult for the calling network to accurately determine the endpoint's location. The endpoint may have its own location, but would not normally include it on the call signaling unless it knew it was an emergency call. There is no mechanism provided in [I-D.ietf-sip-location-conveyance] for a proxy to request the endpoint supply its location, because that would open the endpoint to an attack by any proxy on the path to get it to reveal location. The proxy can attempt to redirect a call to the service URN which, if the device recognizes the significance, would include location in the redirected call from the device. All networks elements should detect emergency calls and supply default location and/or routing if it is not already present.

The LoST server would normally be provided by the local emergency authorities, although the access network or calling network might run its own server using data provided by the emergency authorities. Some enterprises may have local responders and call centers, and could operate their own LoST server, providing URIs to in-house "PSAPs". Local regulations might limit the ability of enterprises to direct emergency calls to in-house services.

The ESRP, which is a normal SIP proxy server in the signaling path of the call, may use a variety of PSAP state information, the location of the caller, and other criteria to onward route the call to the PSAP. In order for the ESRP to route on media choice, the initial INVITE request has to supply an SDP offer.

9. Signaling of emergency calls

9.1. Use of TLS

Best Current Practice for SIP user agents [RFC4504] including handling of audio, video and real-time text [RFC4103] should be applied. As discussed above, location is carried in all emergency calls in the call signaling. Since emergency calls carry privacy-sensitive information, they are subject to the requirements for geospatial protocols [RFC3693]. In particular, signaling information should be carried in TLS, i.e., in 'sips' mode with a ciphersuite which includes strong encryption (e.g., AES). There are exceptions in [RFC3693] for emergency calls. For example, local policy may dictate that location is sent with an emergency call even if the user's policy would otherwise prohibit that. Nevertheless, protection from eavesdropping of location by encryption should be provided.

It is unacceptable to have an emergency call fail to complete because a TLS connection was not created for any reason. Thus, the call should be attempted with TLS, but if the TLS session establishment fails, the call should be automatically retried without TLS. [RFC5630] recommends that to achieve this effect the target specifies a sip URI, but use TLS on the outbound connection. An element that receives a request over a TLS connection should attempt to create a TLS connection to the next hop.

In many cases, persistent TLS connections can be maintained between elements to minimize the time needed to establish them [RFC5626]. In other circumstances, use of session resumption [RFC5077] is recommended. IPsec [RFC4301] is an acceptable alternative to TLS when used with an equivalent crypto suite.

Location may be used for routing by multiple proxy servers on the path. Confidentiality mechanisms such as S/MIME encryption of SIP signaling [RFC3261] cannot be used because they obscure location. Only hop-by-hop mechanisms such as TLS should be used. Implementing location conveyance in SIP mandates inclusion of TLS support.

9.2. SIP signaling requirements for User Agents

SIP UAs that recognize local dial strings, insert location, and perform emergency call routing will create SIP INVITE messages with the Service URN in the Request URI, the LoST-determined URI for the PSAP in a Route header, and the location in a Geolocation header. The INVITE request must also have appropriate call back identifiers (in Contact and From headers). To enable media sensitive routing, the call should include an SDP offer.

SIP caller preferences [RFC3841] can be used to signal how the PSAP should handle the call. For example, a language preference expressed in an Accept-Language header may be used as a hint to cause the PSAP to route the call to a call taker who speaks the requested language. SIP caller preferences may also be used to indicate a need to invoke a relay service for communication with people with disabilities in the call.

9.3. SIP signaling requirements for proxy servers

At least one SIP proxy server in the path of an emergency call must be able to assist UAs that are unable to provide any of the location based routing steps and recognition of dial strings. A Proxy can recognize the lack of location awareness by the lack of a Geolocation header. They can recognize the lack of dial string recognition by the presence of the local emergency call dial string in the From header without the service URN being present. They should obtain the location of the endpoint if possible, and use a default location if they can not, inserting it in a Geolocation header. They should query LoST with the location and put the resulting URI in a Route, with the appropriate service URN in the Request URI. In any event, they are also expected to provide information for the caller using SIP Identity or P-Asserted-Identity. It is often a regulatory matter whether calls normally marked as anonymous are passed as anonymous when they are emergency calls. Proxies must conform to the local regulation or practice.

10. Call backs

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 identifier for this purpose. There are two kinds of call backs. When a call is dropped, or the call taker realizes that some important information is needed that it doesn't have, it must call back the device that placed the emergency call. The PSAP, or a responder, may need to call back the caller much later, and for that purpose, it wants a normal SIP Address of Record. In SIP systems, the caller must include a Contact header field in an emergency call containing a globally routable URI, possibly a GRUU [RFC5627]. This identifier would be used to initiate call-backs immediately by the call-taker if, for example, the call is prematurely dropped. A concern arises with B2BUAs that manipulate Contact headers. Such B2BUAs should always include a Contact header that routes to the same device.

In addition, a call-back identifier as an Address of Record (AoR) must be included either as the URI in the From header field [RFC3261] verified by SIP Identity [RFC4474] or as a network asserted URI [RFC3325]. If the latter, the PSAP will need to establish a suitable spec(t) with the proxies that send it emergency calls. This identifier would be used to initiate a call-back at a later time and may reach the caller, not necessarily on the same device (and at the same location) as the original emergency call as per normal SIP rules. It is often a regulatory matter whether calls normally marked as anonymous are passed as anonymous when they are emergency calls. Proxies must conform to the local regulation or practice.

11. Mid-call behavior

Some PSAPs often include dispatchers, responders or specialists on a call. Some responder's dispatchers are not located in the primary PSAP, the call may have to be transferred to another PSAP. Most often this will be an attended transfer, or a bridged transfer. Therefore a PSAP may need to a REFER request [RFC3515] a call to a bridge for conferencing. Devices which normally involve the user in transfer operations should consider the effect of such interactions when a stressed user places an emergency call. Requiring UI manipulation during such events may not be desirable. Relay services for communication with people with disabilities may be included in the call with the bridge. The UA should be prepared to have the call transferred (usually attended, but possibly blind) per [RFC5359].

12. Call termination

It is undesirable for the caller to terminate an emergency call. PSAP terminates a call using the normal SIP call termination procedures, i.e., with a BYE request.

13. Disabling of features

Certain features that can be invoked while a normal call is active are not permitted when the call is an emergency call. Services such as call waiting, call transfer, three way call and hold should be disabled.

Certain features such as call forwarding can interfere with calls from a PSAP and should be disabled. There is no way to reliably determine a PSAP call back. A UA may be able to determine a PSAP call back by examining the domain of incoming calls after placing an emergency call and comparing that to the domain of the answering PSAP from the emergency call. Any call from the same domain and directed to the supplied Contact header or AoR after an emergency call should be accepted as a call-back from the PSAP if it occurs within a reasonable time after an emergency call was placed.

14. Media

PSAPs should always accept RTP media streams [RFC3550]. Traditionally, voice has been the only media stream accepted by PSAPs. In some countries, text, in the form of Baudot codes or similar tone encoded signaling within a voiceband is accepted ("TTY") for persons who have hearing disabilities. Using SIP signaling includes the capability to negotiate media. Normal SIP offer/answer [RFC3264] negotiations should be used to agree on the media streams to be used. PSAPs should accept real-time text [RFC4103]. All PSAPs should accept G.711 A-law (and mu-law in North America) encoded voice as described in [RFC3551]. Newer text forms are rapidly appearing, with instant messaging now very common, PSAPs should accept IM with at least "pager-mode" MESSAGE request [RFC3428] as well as Message Session Relay Protocol [RFC4975]. Video may be important to support Video Relay Service (sign language interpretation) as well as modern video phones.

It is desirable for media to be kept secure by the use of Secure RTP [RFC3711], using DTLS [RFC5764] for keying.

15. Testing

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.

[I-D.ietf-ecrit-phonebcpl] includes a description of an automated test procedure that validates routing, signaling and media path continuity. This test would be used within some random interval after boot time, and whenever the device location changes enough that a new PSAP mapping is returned by the LoST server.

The PSAP needs to be able to control frequency and duration of the test, and since the process could be abused, it may temporarily or permanently suspend its operation.

There is a concern associated with testing during a so-called "avalanche-restart" event where, for example a large power outage affects a large number of endpoints, that, when power is restored, all attempt to reboot and, possibly, test. Devices need to randomize their initiation of a boot time test to avoid the problem.

16. Security Considerations

Security considerations for emergency calling have been documented in [RFC5069] and [RFC6280].

This document suggests that security (TLS or IPsec) be used hop by hop on a SIP call to protect location information, identity, etc. It also suggests that if the attempt to create a security association fails, the call be retried without the security. It's more important to get an emergency call through than to protect the data; indeed, in many jurisdictions privacy is explicitly waived when making emergency calls. Placing a call without security may reveal user information, including location. The alternative - failing the call if security cannot be established, is considered unacceptable.

17. IANA Considerations

This document has no actions for IANA.

18. Acknowledgments

This draft was created from a draft-schulzrinne-sipping-emergency-arch-02 together with sections from draft-polk-newton-ecrit-arch-considerations-02.

Design Team members participating in this draft creation include Martin Dolly, Stu Goldman, Ted Hardie, Marc Linsner, Roger Marshall,

Shida Schubert, Tom Taylor and Hannes Tschofenig,. Further comments and input were provided by Richard Barnes, Barbara Stark and James Winterbottom.

19. Informative References

- [I-D.ietf-ecrit-phonebcip]
Rosen, B. and J. Polk, "Best Current Practice for Communications Services in support of Emergency Calling", draft-ietf-ecrit-phonebcip-20 (work in progress), September 2011.
- [I-D.ietf-sip-location-conveyance]
Polk, J. and B. Rosen, "Location Conveyance for the Session Initiation Protocol", draft-ietf-sip-location-conveyance-13 (work in progress), March 2009.
- [LLDP] IEEE, "IEEE802.1ab Station and Media Access Control", Dec 2004.
- [LLDP-MED]
TIA, "ANSI/TIA-1057 Link Layer Discovery Protocol - Media Endpoint Discovery".
- [NENAi3TRD]
NENA, "08-751 NENA i3 Technical Requirements for", 2006.
- [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.
- [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.
- [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.
- [RFC3693] Cuellar, J., Morris, J., Mulligan, D., Peterson, J., and J. Polk, "Geopriv Requirements", RFC 3693, February 2004.
- [RFC3711] Baugher, M., McGrew, D., Naslund, M., Carrara, E., and K. Norrman, "The Secure Real-time Transport Protocol (SRTP)", RFC 3711, March 2004.
- [RFC3841] Rosenberg, J., Schulzrinne, H., and P. Kyzivat, "Caller Preferences for the Session Initiation Protocol (SIP)", RFC 3841, August 2004.
- [RFC3856] Rosenberg, J., "A Presence Event Package for the Session Initiation Protocol (SIP)", RFC 3856, August 2004.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, January 2005.
- [RFC4103] Hellstrom, G. and P. Jones, "RTP Payload for Text Conversation", RFC 4103, June 2005.
- [RFC4119] Peterson, J., "A Presence-based GEOPRIV Location Object Format", RFC 4119, December 2005.
- [RFC4190] Carlberg, K., Brown, I., and C. Beard, "Framework for Supporting Emergency Telecommunications Service (ETS) in IP Telephony", RFC 4190, November 2005.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.
- [RFC4474] Peterson, J. and C. Jennings, "Enhancements for Authenticated Identity Management in the Session Initiation Protocol (SIP)", RFC 4474, August 2006.
- [RFC4504] Sinnreich, H., Lass, S., and C. Stredicke, "SIP Telephony Device Requirements and Configuration", RFC 4504, May 2006.

- [RFC4776] Schulzrinne, H., "Dynamic Host Configuration Protocol (DHCPv4 and DHCPv6) Option for Civic Addresses Configuration Information", RFC 4776, November 2006.
- [RFC4967] Rosen, B., "Dial String Parameter for the Session Initiation Protocol Uniform Resource Identifier", RFC 4967, July 2007.
- [RFC4975] Campbell, B., Mahy, R., and C. Jennings, "The Message Session Relay Protocol (MSRP)", RFC 4975, September 2007.
- [RFC5012] Schulzrinne, H. and R. Marshall, "Requirements for Emergency Context Resolution with Internet Technologies", RFC 5012, January 2008.
- [RFC5031] Schulzrinne, H., "A Uniform Resource Name (URN) for Emergency and Other Well-Known Services", RFC 5031, January 2008.
- [RFC5069] Taylor, T., Tschofenig, H., Schulzrinne, H., and M. Shanmugam, "Security Threats and Requirements for Emergency Call Marking and Mapping", RFC 5069, January 2008.
- [RFC5077] Salowey, J., Zhou, H., Eronen, P., and H. Tschofenig, "Transport Layer Security (TLS) Session Resumption without Server-Side State", RFC 5077, January 2008.
- [RFC5139] Thomson, M. and J. Winterbottom, "Revised Civic Location Format for Presence Information Data Format Location Object (PIDF-LO)", RFC 5139, February 2008.
- [RFC5222] Hardie, T., Newton, A., Schulzrinne, H., and H. Tschofenig, "LoST: A Location-to-Service Translation Protocol", RFC 5222, August 2008.
- [RFC5223] Schulzrinne, H., Polk, J., and H. Tschofenig, "Discovering Location-to-Service Translation (LoST) Servers Using the Dynamic Host Configuration Protocol (DHCP)", RFC 5223, August 2008.
- [RFC5359] Johnston, A., Sparks, R., Cunningham, C., Donovan, S., and K. Summers, "Session Initiation Protocol Service Examples", BCP 144, RFC 5359, October 2008.
- [RFC5491] Winterbottom, J., Thomson, M., and H. Tschofenig, "GEOPRIV Presence Information Data Format Location Object (PIDF-LO) Usage Clarification, Considerations, and Recommendations",

RFC 5491, March 2009.

- [RFC5626] Jennings, C., Mahy, R., and F. Audet, "Managing Client-Initiated Connections in the Session Initiation Protocol (SIP)", RFC 5626, October 2009.
- [RFC5627] Rosenberg, J., "Obtaining and Using Globally Routable User Agent URIs (GRUUs) in the Session Initiation Protocol (SIP)", RFC 5627, October 2009.
- [RFC5630] Audet, F., "The Use of the SIPS URI Scheme in the Session Initiation Protocol (SIP)", RFC 5630, October 2009.
- [RFC5764] McGrew, D. and E. Rescorla, "Datagram Transport Layer Security (DTLS) Extension to Establish Keys for the Secure Real-time Transport Protocol (SRTP)", RFC 5764, May 2010.
- [RFC5985] Barnes, M., "HTTP-Enabled Location Delivery (HELD)", RFC 5985, September 2010.
- [RFC5986] Thomson, M. and J. Winterbottom, "Discovering the Local Location Information Server (LIS)", RFC 5986, September 2010.
- [RFC6225] Polk, J., Linsner, M., Thomson, M., and B. Aboba, "Dynamic Host Configuration Protocol Options for Coordinate-Based Location Configuration Information", RFC 6225, July 2011.
- [RFC6280] Barnes, R., Lepinski, M., Cooper, A., Morris, J., Tschofenig, H., and H. Schulzrinne, "An Architecture for Location and Location Privacy in Internet Applications", BCP 160, RFC 6280, July 2011.
- [WGS84] NIMA, "NIMA Technical Report TR8350.2, Department of Defense World Geodetic System 1984, Its Definition and Relationships With Local Geodetic Systems, Third Edition", July 1997.

Authors' Addresses

Brian Rosen
NeuStar, Inc.
470 Conrad Dr
Mars, PA 16046
USA

Email: br@brianrosen.net

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

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

James Polk
Cisco Systems
3913 Treemont Circle
Colleyville, Texas 76034
USA

Phone: +1-817-271-3552
Email: jmpolk@cisco.com

Andrew Newton
TranTech/MediaSolv
4900 Seminary Road
Alexandria, VA 22311
USA

Phone: +1 703 845 0656
Email: andy@hxr.us

ECRIT
Internet-Draft
Intended status: Experimental
Expires: January 11, 2013

H. Schulzrinne
Columbia University
H. Tschofenig
Nokia Siemens Networks
July 10, 2012

Synchronizing Location-to-Service Translation (LoST) Protocol based
Service Boundaries and Mapping Elements
draft-ietf-ecrit-lost-sync-18.txt

Abstract

The Location-to-Service Translation (LoST) protocol is an XML-based protocol for mapping service identifiers and geodetic or civic location information to service URIs and service boundaries. In particular, it can be used to determine the location-appropriate Public Safety Answering Point (PSAP) for emergency services.

The <mapping> element is used to encapsulate information about service boundaries is defined in the LoST protocol specification and circumscribes the region within which all locations map to the same service Uniform Resource Identifier (URI) or set of URIs for a given service.

This document defines an XML protocol to exchange these mappings between two nodes. This mechanism is designed for the exchange of authoritative <mapping> elements between two entities. Exchanging cached <mapping> elements, i.e. non-authoritative elements, is possible but not envisioned. In any case, this document can also be used without the LoST protocol even though the format of the <mapping> element is re-used from the LoST specification.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 11, 2013.

Copyright Notice

Copyright (c) 2012 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	4
2. Terminology	5
3. A Motivating Example	6
4. Querying for Mappings with a <getMappingsRequest> / <getMappingsResponse> Exchange	12
4.1. Behavior of the LoST Sync Destination	12
4.2. Behavior of the LoST Sync Source	12
4.3. Examples	13
5. Pushing Mappings via <pushMappings> and <pushMappingsResponse>	15
5.1. Behavior of the LoST Sync Source	15
5.2. Behavior of the LoST Sync Destination	15
5.3. Example	16
6. Transport	20
7. RelaxNG	21
8. Operational Considerations	23
9. Security Considerations	24
10. IANA Considerations	25
10.1. Media Type Registration	25
10.2. LoST Sync Relax NG Schema Registration	26
10.3. LoST Synchronization Namespace Registration	27
11. Acknowledgments	28
12. References	29
12.1. Normative References	29
12.2. Informative References	29
Authors' Addresses	30

1. Introduction

Since the early days of emergency services there has been a desire to route emergency calls to Public Safety Answering Points (PSAPs) that are nearest to the location of the emergency caller. For this purpose each PSAP discloses one or multiple service boundaries so that this information can be used to select the appropriate PSAP and to route the call to it. RFC 5222 [RFC5222] defines this data structure in the following way:

A service boundary circumscribes the region within which all locations map to the same service Uniform Resource Identifier (URI) or set of URIs for a given service. A service boundary may consist of several non-contiguous geometric shapes.

RFC 5222 [RFC5222] not only defines the term but it also specifies the data structure itself: the <mapping> element.

This document re-uses this existing data structure and defines an XML-based protocol to exchange authoritative service boundaries between two entities (the LoST Sync source and the LoST Sync destination). This protocol can be used with and without the actual LoST protocol.

The rest of the document is structured as follows: Section 3 starts with an example usage of the LoST protocol. In Section 4, Section 5, Section 6, and Section 7 we describe the protocol semantics, transport considerations and the schema. Finally, we conclude with operational and security considerations in Section 8, and in Section 9.

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 RFC 2119 [RFC2119].

This document reuses terminology introduced by the mapping architecture document [RFC5582], such as 'coverage region', 'forest guide', 'mapping', 'authoritative mapping server', and 'ESRP'.

Throughout this document we use the term LoST Sync source and LoST Sync destination to denote the protocol end points of the exchange. The protocol is referred as LoST Sync within the text.

3. A Motivating Example

The LoST Sync mechanism can, for example, be used in the LoST architecture, as specified in the [RFC5582]. There, LoST servers act in different roles that cooperate to provide an ubiquitous, globally scalable and resilient mapping service. In the LoST mapping architecture, LoST servers can peer, i.e., have an on-going data exchange relationship. Peering relationships are set up manually, based on local policies. A LoST server may peer with any number of other LoST servers. Forest guides peer with other forest guides; authoritative mapping servers peer with forest guides and other authoritative servers, either in the same cluster or above or below them in the tree. Authoritative mapping servers push coverage regions "up" the tree, i.e., from child nodes to parent nodes. The child informs the parent of the geospatial or civic region that it covers for a specific service.

Consider a hypothetical deployment of LoST in two countries, for example Austria and Finland. Austria, in our example, runs three authoritative mapping servers labeled as 'East', 'West' and 'Vienna' whereby the former two cover the entire country except for Vienna, which is covered by a separate LoST server. There may be other caching LoST servers run by ISPs, universities, and VSPs but they are not relevant for this illustration. Finland, on the other hand, decided to only deploy a single LoST server that also acts as a Forest Guide. For this simplistic illustration we assume that only one service is available, namely 'urn:service:sos' since otherwise the number of stored mappings would have to be multiplied by the number of used services.

Figure 1 shows the example deployment.




Figure 1: LoST Deployment Example

The configuration of these nodes would therefore be as follows:

Forest Guide Austria: This forest guide would contain mappings for the three authoritative mapping servers (East, West and Vienna) describing what area they are responsible for. Note that each mapping would contain a service URN and these mappings point to LoST servers rather than to PSAPs or ESRPs.

LoST Server 'East': This LoST server would contain all the mappings to PSAPs covering one half of the country.

Additionally, the LoST server aggregates all the information it has and provides an abstracted view towards the Forest Guide indicating that it is responsible for a certain area (for a given service, and for a given location profile). Such a mapping could have the following structure:

```

<mapping
  xmlns="urn:ietf:params:xml:ns:lost1"
  xmlns:gml="http://www.opengis.net/gml"
  expires="2009-01-01T01:44:33Z"
  lastUpdated="2009-12-01T01:00:00Z"
  source="east-austria.lost-example.com"
  sourceId="e8b05a41d8d1415b80f2cdbb96ccf109">
  <displayName xml:lang="en">LoST Server 'East'</displayName>
  <service>urn:service:sos</service>
  <serviceBoundary profile="geodetic-2d">
    <gml:Polygon srsName="urn:ogc:def::crs:EPSG::4326">
      <gml:exterior>
        <gml:LinearRing>
          <gml:pos> ... </gml:pos>
          ..... list of coordinates for
          boundary of LoST server 'East'
          <gml:pos> ... </gml:pos>
        </gml:LinearRing>
      </gml:exterior>
    </gml:Polygon>
  </serviceBoundary>
  <uri/>
</mapping>

```

Figure 2: Forest Guide Austria Mapping XML Snippet

Note that the XML code snippet in Figure 2 serves illustrative purposes only and does not validate. As it can be seen in this example the `<uri>` element is absent and the 'source' attribute identifies the LoST server, namely "east-austria.lost-example.com".

The above-shown mapping is what is the LoST server "east-austria.lost-example.com" provides to the Austrian Forest Guide.

LoST Server 'West': This LoST server would contain all the mappings to PSAPs covering the other half of the country.

LoST Server 'Vienna': This LoST server would contain all the mappings to PSAPs in the area of Vienna.

Forest Guide Finland: In our example we assume that Finland would deploy a single ESRP for the entire country as their IP-based emergency services solution. There is only a single LoST server and it is co-located with the Forest Guide, as shown in Figure 1. The mapping data this FG would distribute via LoST sync is shown in Figure 3.

```

<mapping xmlns="urn:ietf:params:xml:ns:lost1"
  expires="2007-01-01T01:44:33Z"
  lastUpdated="2006-11-01T01:00:00Z"
  source="finland.lost-example.com"
  sourceId="7e3f40b098c711dbb6060800200c9a66">
  <displayName xml:lang="en">Finland ESRP</displayName>
  <service>urn:service:sos</service>
  <serviceBoundary profile="civic">
    <civicAddress
      xmlns="urn:ietf:params:xml:ns:pidf:geopriv10:civicAddr">
      <country>FI</country>
    </civicAddress>
  </serviceBoundary>
  <uri/>
</mapping>

```

Figure 3: Forest Guide Finland Mapping XML Snippet

An example mapping stored at the co-located LoST server is shown in Figure 4.

```

<mapping xmlns="urn:ietf:params:xml:ns:lost1"
  expires="2007-01-01T01:44:33Z"
  lastUpdated="2006-11-01T01:00:00Z"
  source="finland.lost-example.com"
  sourceId="7e3f40b098c711dbb6060800200c9a66">
  <displayName xml:lang="en">Finland ESRP</displayName>
  <service>urn:service:sos</service>
  <serviceBoundary profile="civic">
    <civicAddress
      xmlns="urn:ietf:params:xml:ns:pidf:geopriv10:civicAddr">
      <country>FI</country>
    </civicAddress>
  </serviceBoundary>
  <uri>sip:esrp@finland-example.com</uri>
  <uri>xmpp:esrp@finland-example.com</uri>
  <serviceNumber>112</serviceNumber>
</mapping>

```

Figure 4: Forest Guide Finland / Co-Located LoST Server Mapping XML Snippet

The LoST sync mechanism described in this document could be run between the two Forest Guides. Thereby, the three mappings stored in the Austria FG are sent to the FG Finland and a single mapping in the FG Finland is sent to the FG Austria. Additionally, the three Austrian LoST servers could utilize LoST sync to inform the Austrian FG about their boundaries. These three authoritative mapping servers

in Austria would be responsible to maintain their own mapping information. Since the amount of data being exchanged is small and the expected rate of change is low the nodes are configured to always exchange all their mapping information whenever a change happens.

This document defines two types of exchanges and those are best described by the exchange between two nodes as shown in Figure 5 and Figure 6. The protocol exchange always runs between a LoST Sync source and a LoST Sync destination. Node A in the examples of Figure 5 and Figure 6 has mappings that Node B is going to retrieve. Node A acts as the source for the data and Node B is the destination.

The `<getMappingsRequest>` request allows a LoST Sync source to request mappings from a LoST Sync destination.

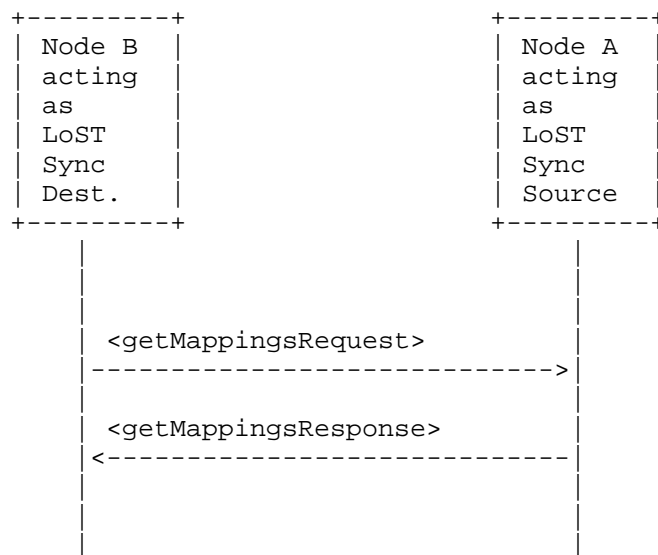


Figure 5: Querying for Mappings with a `<getMappingsRequest>` Message

Note that in the exchange illustrated in Figure 5 Node B is issuing the first request and plays the role of the HTTPS client and Node A plays the role of the HTTPS server.

The `<pushMappingsRequest>` exchange allows a LoST Sync source to push mappings to LoST Sync destination. In this example we assume that Node A has been configured maintain state about the mappings it had pushed to Node B.

No publish/subscribe mechanism is defined in this document that would allow Node B to tell Node A about what mappings it is interested in

nor a mechanism for learning to which entities mappings have to be pushed.




Figure 6: Pushing Mappings with a <pushMappingsRequest> Message

Node A issuing the first request in Figure 6 plays the role of the HTTPS client and Node B plays the role of the HTTPS server.

4. Querying for Mappings with a <getMappingsRequest> / <getMappingsResponse> Exchange

4.1. Behavior of the LoST Sync Destination

A LoST Sync destination has two ways to retrieve mapping elements from a LoST Sync source.

1. A mechanism that is suitable when no mappings are available on the LoST Sync destination is to submit an empty <getMappingsRequest> message, as shown in Figure 7. The intent by the LoST Sync destination thereby is to retrieve all mappings from the LoST Sync source. Note that the request does not propagate further to other nodes.
2. In case a LoST Sync destination node has already obtained mappings in previous exchanges then it may want to check whether these mappings have been updated in the meanwhile. The policy when to poll for updated mapping information is outside the scope of this document. The <getMappingsRequest> message with one or multiple <exists> child element(s) allows to reduce the number of returned mappings to those that have been updated and also to those that are missing.

In response to the <getMappingsRequest> message the LoST Sync destination waits for the <getMappingsResponse> message. In case of a successful response the LoST Sync destination stores the received mappings and determines which mappings to update.

4.2. Behavior of the LoST Sync Source

When a LoST Sync source receives an empty <getMappingsRequest> message then all locally available mappings MUST be returned.

When a LoST Sync source receives a <getMappingsRequest> message with one or multiple <exists> child element(s) then it MUST consult with the local mapping database to determine whether any of the mappings of the client is stale and whether there are mappings locally that the client does not yet have. The former can be determined by finding mappings corresponding to the 'source' and 'sourceID' attributes where a mapping with a more recent lastUpdated date exists.

Processing a <getMappingsRequest> message MAY lead to a successful response in the form of a <getMappingsResponse> or an <errors> message. Only the <badRequest>, <forbidden>, <internalError>, <serverTimeout> errors, defined in [RFC5222], are utilized by this specification. Neither the <redirect> nor the <warnings> messages are reused by this message.

4.3. Examples

The first example shows an empty `<getMappingsRequest>` message that would retrieve all locally stored mappings at the LoST Sync source.

```
<?xml version="1.0" encoding="UTF-8"?>
<getMappingsRequest xmlns="urn:ietf:params:xml:ns:lostsync1"/>
```

Figure 7: Example of empty `<getMappingsRequest>` message

A further example request is shown in Figure 8 and the corresponding response is depicted in Figure 9. In this example the `<getMappingsRequest>` element contains information about the mapping that is locally available to the client inside the `<mapping-fingerprint>` element (with `source="authoritative.bar.example"`, `sourceId="7e3f40b098c711dbb6060800200c9a66"`, and `lastUpdated="2006-11-01T01:00:00Z"`). The query asks for mappings that are more recent than the available one as well as any missing mapping.

```
<?xml version="1.0" encoding="UTF-8"?>
<getMappingsRequest xmlns="urn:ietf:params:xml:ns:lostsync1">
  <exists>
    <mapping-fingerprint source="authoritative.bar.example"
      sourceId="7e3f40b098c711dbb6060800200c9a66"
      lastUpdated="2006-11-01T01:00:00Z">
    </mapping-fingerprint>
  </exists>
</getMappingsRequest>
```

Figure 8: Example `<getMappingsRequest>` Message

The response to the above request is shown in Figure 9. A more recent mapping was available with the identification of `source="authoritative.bar.example"` and `sourceId="7e3f40b098c711dbb6060800200c9a66"`. Only one mapping that matched `source="authoritative.foo.example"` was found and returned.

```
<?xml version="1.0" encoding="UTF-8"?>
<sync:getMappingsResponse
  xmlns:sync="urn:ietf:params:xml:ns:lostsync1"
  xmlns="urn:ietf:params:xml:ns:lost1"
  xmlns:gml="http://www.opengis.net/gml">

  <mapping source="authoritative.bar.example"
    sourceId="7e3f40b098c711dbb6060800200c9a66"
```

```

        lastUpdated="2008-11-26T01:00:00Z"
        expires="2009-12-26T01:00:00Z">
        <displayName xml:lang="en">Leonia Police Department
        </displayName>
        <service>urn:service:sos.police</service>
        <serviceBoundary
profile="urn:ietf:params:lost:location-profile:basic-civic">
        <civicAddress
xmlns="urn:ietf:params:xml:ns:pidf:geopriv10:civicAddr">
        <country>US</country>
        <A1>NJ</A1>
        <A3>Leonia</A3>
        <PC>07605</PC>
        </civicAddress>
        </serviceBoundary>
        <uri>sip:police@leonianj2.example.org</uri>
        <serviceNumber>911</serviceNumber>
</mapping>

<mapping expires="2009-01-01T01:44:33Z"
lastUpdated="2008-11-01T01:00:00Z"
source="authoritative.foo.example"
sourceId="7e3f40b098c711dbb606011111111111">
<displayName xml:lang="en">New York City Police Department
</displayName>
<service>urn:service:sos.police</service>
<serviceBoundary profile="geodetic-2d">
  <gml:Polygon srsName="urn:ogc:def::crs:EPSG::4326">
    <gml:exterior>
      <gml:LinearRing>
        <gml:pos>37.775 -122.4194</gml:pos>
        <gml:pos>37.555 -122.4194</gml:pos>
        <gml:pos>37.555 -122.4264</gml:pos>
        <gml:pos>37.775 -122.4264</gml:pos>
        <gml:pos>37.775 -122.4194</gml:pos>
      </gml:LinearRing>
    </gml:exterior>
  </gml:Polygon>
</serviceBoundary>
<uri>sip:nypd@example.com</uri>
<uri>xmpp:nypd@example.com</uri>
<serviceNumber>911</serviceNumber>
</mapping>

</sync:getMappingsResponse>

```

Figure 9: Example <getMappingsResponse> Message

5. Pushing Mappings via <pushMappings> and <pushMappingsResponse>

5.1. Behavior of the LoST Sync Source

When a LoST Sync source obtains new information that is of interest to its peers, it may push the new mappings to its peers. Configuration settings at both peers decide whether this functionality is used and what mappings are pushed to which other peers. New mappings may arrive through various means, such as a manual addition to the local mapping database, or through the interaction with other entities. Deleting mappings may also trigger a protocol interaction.

The LoST Sync source SHOULD keep track of which LoST Sync destination it has pushed mapping elements. If it does not keep state information then it always has to push the complete data set. As discussed in Section 5.1 of [RFC5222], mapping elements are identified by the 'source', 'sourceID' and 'lastUpdated' attributes. A mapping is considered the same if these three attributes match.

A <pushMappings> request sent by a LoST Sync source MUST containing one or more <mapping> elements.

To delete a mapping, the content of the mapping is left empty, i.e. the <mapping> element only contains the 'source', 'sourceID', 'lastUpdated', and 'expires' attribute. Figure 10 shows an example request where the mapping with the source="nj.us.example", sourceId="123", lastUpdated="2008-11-01T01:00:00Z", expires="2008-11-01T01:00:00Z" is requested to be deleted. Note that the 'expires' attribute is required per schema definition but will be ignored in processing the request on the receiving side. A sync source may want to delete the mapping from its internal mapping database, but has to remember which peers it has distributed this update to unless it has other ways to ensure that databases do not get out of sync.

5.2. Behavior of the LoST Sync Destination

When a LoST Sync destination receives a <pushMappingsRequest> message then the cache with the existing mappings is inspected to determine whether the received mapping should lead to an update of an already existing mapping, should create a new mapping in the cache, or should be discarded.

If a newly received mapping has a more recent time in its 'lastUpdated' attribute, it MUST update an existing mapping that has matching 'source' and 'sourceID' attributes.

If the received mapping does not match with any existing mapping

based on the 'source' and 'sourceId' then it MUST be added to the local cache as an independent mapping.

If a <pushMappingsRequest> message with an empty <mapping> element is received then a corresponding mapping has to be determined based on the 'source', and the 'sourceID'.

If no mapping can be identified then an <errors> response MUST be returned that contains the <notDeleted> child element. The <notDeleted> element MAY contain a 'message' attribute with an error description used for debugging purposes. The <notDeleted> element MUST contain the <mapping> element(s) that caused the error.

The response to a <pushMappingsRequest> request is a <pushMappingsResponse> message. With this specification, a successful response message returns no additional elements, whereas an <errors> response is returned in the response message, if the request failed. Only the <badRequest>, <forbidden>, <internalError> or <serverTimeout> errors defined in Section 13.1 of [RFC5222], are used. The <redirect> and <warnings> messages are not used for this query/response.

If the set of nodes that are synchronizing their data does not form a tree, it is possible that the same information arrives through several other nodes. This is unavoidable, but generally only imposes a modest overhead. (It would be possible to create a spanning tree in the same fashion as IP multicast, but the complexity does not seem warranted, given the relatively low volume of data.)

5.3. Example

An example is shown in Figure 10. Imagine a LoST node that obtained two new mappings identified as follows:

o

```
source="authoritative.example"
sourceId="7e3f40b098c711dbb6060800200c9a66"
lastUpdated="2008-11-26T01:00:00Z"
```

o

```
source="authoritative.example"
sourceId="7e3f40b098c711dbb6060111111111111"
lastUpdated="2008-11-01T01:00:00Z"
```

These two mappings have to be added to the peer's mapping database.

Additionally, the following mapping has to be deleted:

- o source="nj.us.example" sourceId="123" lastUpdated="2008-11-01T01:00:00Z"

```
<?xml version="1.0" encoding="UTF-8"?>
<sync:pushMappings
  xmlns:sync="urn:ietf:params:xml:ns:lostsync1"
  xmlns="urn:ietf:params:xml:ns:lost1"
  xmlns:gml="http://www.opengis.net/gml">

  <mapping source="authoritative.example"
    sourceId="7e3f40b098c711dbb6060800200c9a66"
    lastUpdated="2008-11-26T01:00:00Z"
    expires="2009-12-26T01:00:00Z">
    <displayName xml:lang="en">Leonia Police Department
    </displayName>
    <service>urn:service:sos.police</service>
    <serviceBoundary
profile="urn:ietf:params:lost:location-profile:basic-civic">
      <civicAddress
xmlns="urn:ietf:params:xml:ns:pidf:geopriv10:civicAddr">
        <country>US</country>
        <A1>NJ</A1>
        <A3>Leonia</A3>
        <PC>07605</PC>
      </civicAddress>
    </serviceBoundary>
    <uri>sip:police@leonianj.example.org</uri>
    <serviceNumber>911</serviceNumber>
  </mapping>

  <mapping expires="2009-01-01T01:44:33Z"
    lastUpdated="2008-11-01T01:00:00Z"
    source="authoritative.example"
    sourceId="7e3f40b098c711dbb6060111111111111">
    <displayName xml:lang="en">New York City Police Department
    </displayName>
    <service>urn:service:sos.police</service>
    <serviceBoundary profile="geodetic-2d">
      <gml:Polygon srsName="urn:ogc:def::crs:EPSG::4326">
        <gml:exterior>
          <gml:LinearRing>
            <gml:pos>37.775 -122.4194</gml:pos>
            <gml:pos>37.555 -122.4194</gml:pos>
            <gml:pos>37.555 -122.4264</gml:pos>
            <gml:pos>37.775 -122.4264</gml:pos>
```

```

        <gml:pos>37.775 -122.4194</gml:pos>
      </gml:LinearRing>
    </gml:exterior>
  </gml:Polygon>
</serviceBoundary>
<uri>sip:nypd@example.com</uri>
<uri>xmpp:nypd@example.com</uri>
<serviceNumber>911</serviceNumber>
</mapping>

<mapping source="nj.us.example"
  sourceId="123"
  lastUpdated="2008-11-01T01:00:00Z"
  expires="2008-11-01T01:00:00Z"/>

</sync:pushMappings>
```

Figure 10: Example <pushMappingsRequest> Message

In response, the peer performs the necessary operation and updates its mapping database. In particular, it will check whether the other peer is authorized to perform the update and whether the elements and attributes contain values that it understands. In our example, a positive response is returned as shown in Figure 11.

```
<?xml version="1.0" encoding="UTF-8"?>
<pushMappingsResponse xmlns="urn:ietf:params:xml:ns:lostsync1" />
```

Figure 11: Example <pushMappingsResponse>

In case that a mapping could not be deleted as requested the following error response might be returned instead.


```
<?xml version="1.0" encoding="UTF-8"?>
<errors xmlns="urn:ietf:params:xml:ns:lost1"
  xmlns:sync="urn:ietf:params:xml:ns:lostsync1"
  source="nodeA.example.com">

  <sync:notDeleted
    message="Could not delete the indicated mapping."
    xml:lang="en">

    <mapping source="nj.us.example"
      sourceId="123"
      lastUpdated="2008-11-01T01:00:00Z"
      expires="2008-11-01T01:00:00Z"/>
  </sync:notDeleted>
</errors>
```

Figure 12: Example <errors> Message

6. Transport

LoST Sync needs an underlying protocol transport mechanism to carry requests and responses. This document uses HTTPS as a transport to exchange XML documents. No fallback to HTTP is provided.

When using HTTP-over-TLS [RFC2818], LoST Sync messages use the POST method. Request MUST use the Cache-Control response directive "no-cache".

All LoST Sync responses, including those indicating a LoST warning or error, are carried in 2xx responses, typically 200 (OK). 3xx, 4xx and 5xx HTTP response codes indicates that the request itself failed or was redirected; these responses do not contain any LoST Sync XML elements.

7. RelaxNG

Note: In order to avoid copying pattern definitions from the LoST Relax NG schema [RFC5222] to this document we include it as "lost.rng" (XML syntax) in the Relax NG schema below.

```
<?xml version="1.0" encoding="utf-8"?>

  <grammar ns="urn:ietf:params:xml:ns:lostsync1"
    xmlns="http://relaxng.org/ns/structure/1.0"
    xmlns:a="http://relaxng.org/ns/compatibility/annotations/1.0"
    datatypeLibrary="http://www.w3.org/2001/XMLSchema-datatypes">

    <include href="lost.rng"/>

    <start combine="choice">

      <a:documentation> Location-to-Service Translation (LoST)
        Synchronization Protocol</a:documentation>

      <choice>
        <ref name="pushMappings"/>
        <ref name="pushMappingsResponse"/>
        <ref name="getMappingsRequest"/>
        <ref name="getMappingsResponse"/>
      </choice>
    </start>

    <define name="pushMappings">
      <element name="pushMappings">
        <oneOrMore>
          <ref name="mapping"/>
        </oneOrMore>

        <ref name="extensionPoint"/>
      </element>
    </define>

    <define name="pushMappingsResponse">
      <element name="pushMappingsResponse">
        <ref name="extensionPoint"/>
      </element>
    </define>

    <define name="getMappingsRequest">
      <element name="getMappingsRequest">
```

```
        <choice>
            <ref name="exists"></ref>
            <ref name="extensionPoint"/>
        </choice>
    </element>
</define>

<define name="exists">
    <element name="exists">
        <oneOrMore>
            <element name="mapping-fingerprint">
                <attribute name="source">
                    <data type="token"/>
                </attribute>
                <attribute name="sourceId">
                    <data type="token"/>
                </attribute>
                <attribute name="lastUpdated">
                    <data type="dateTime"/>
                </attribute>
                <ref name="extensionPoint"/>
            </element>
        </oneOrMore>
    </element>
</define>

<define name="getMappingsResponse">
    <element name="getMappingsResponse">
        <oneOrMore>
            <ref name="mapping"/>
        </oneOrMore>
        <ref name="extensionPoint"/>
    </element>
</define>

<!-- error messages -->

<define name="notDeleted">
    <element name="notDeleted">
        <ref name="basicException"/>
        <oneOrMore>
            <ref name="mapping"/>
        </oneOrMore>
    </element>
</define>
</grammar>
```

8. Operational Considerations

When different LoST servers use the mechanism described in this document to synchronize their mapping data then it is important to ensure that loops are avoided. The example shown in Figure 13 with three LoST servers A, B and C (each of them acts as a sync source and a sync destination) illustrates the challenge in more detail. A and B synchronize data between each other; the same is true for A and C, and B and C, respectively.




Figure 13: Synchronization Configuration Example

Now, imagine that server A adds a new mapping. This mapping is uniquely identified by the combination of "source", "sourceid" and "last updated". Assume that A would push this new mapping to B and C. When B obtained this new mapping it would find out that it has to distribute it to its peer C. C would also want to distribute the mapping to B. If the original mapping with the "source", "sourceid" and "last updated" is not modified by either B or C then these two servers would recognize that they already possess the mapping and can ignore the update.

It is important that implementations MUST NOT modify mappings they receive. An entity acting maliciously would, however, intentionally modify mappings or inject bogus mappings. To avoid the possibility of an untrustworthy member claiming a coverage region that it is not authorized for, authoritative mapping server MUST sign mappings they distribute using an XML digital signature [W3C.REC-xmldsig-core-20020212]. A recipient MUST verify that the signing entity is indeed authorized to speak for that region. In many cases, this will require an out-of-band agreement to be in place to agree on specific entities to take on this role. Determining who can speak for a particular region is inherently difficult unless there is a small set of authorizing entities that participants in the mapping architecture can trust. Receiving systems should be particularly suspicious if an existing coverage region is replaced by a new one that contains a different value in the <uri> element. When mappings are digitally signed, they cannot be modified by intermediate LoST servers.

9. Security Considerations

This document defines a protocol for exchange of authoritative mapping information between two entities. Hence, the protocol operations described in this document require authentication of neighboring nodes.

The LoST Sync client and servers MUST implement TLS and use TLS. Which version(s) ought to be implemented will vary over time, and depend on the widespread deployment and known security vulnerabilities at the time of implementation. At the time of this writing, TLS version 1.2 [RFC5246] is the most recent version, but has very limited actual deployment, and might not be readily available in implementation toolkits. TLS version 1.0 [RFC2246] is the most widely deployed version, and will give the broadest interoperability.

Mutual authentication between the LoST Sync source and the LoST Sync destination is not necessarily required in all deployments unless an emergency service authority wants to enforce access control prior to the distribution of their mapping elements. This may, for example, be the case when certain emergency services network internal mappings are not meant for public distribution.

An additional threat is caused by compromised or misconfigured LoST servers. A denial of service could be the consequence of an injected mapping. If the mapping data contains an URL that does not exist then emergency services for the indicated area are not reachable. If all mapping data contains URLs that point to a single PSAP (rather than a large number of PSAPs) then this PSAP is likely to experience overload conditions. If the mapping data contains a URL that points to a server controlled by the adversary itself then it might impersonate PSAPs.

Section 8 discusses this security threat and mandates signed mappings. For unusual changes to the mapping database approval by a system administrator of the emergency services infrastructure (or a similar expert) may be required before any mappings are installed.

10. IANA Considerations

10.1. Media Type Registration

This specification requests the registration of a new media type according to the procedures of RFC 4288 [RFC4288] and guidelines in RFC 3023 [RFC3023].

Type name: application

Subtype name: lostsync+xml

Required parameters: none

Optional parameters: charset

Same as charset parameter of application/xml as specified in RFC 3023 [RFC3023].

Encoding considerations: Identical to those of "application/xml" as described in [RFC3023], Section 3.2.

Security considerations: This content type is designed to carry LoST Synchronization protocol payloads and the security considerations section of RFCXXXX is applicable. In addition, as this media type uses the "+xml" convention, it shares the same security considerations as described in [RFC3023], Section 10. [NOTE TO IANA/RFC-EDITOR: Please replace XXXX with the RFC number of this specification.]

Interoperability considerations: None

Published specification: RFCXXXX [NOTE TO IANA/RFC-EDITOR: Please replace XXXX with the RFC number of this specification.]

Applications which use this media type: Emergency and Location-based Systems

Additional information:

Magic number(s): None

File extension(s): .lostsyncxml

Macintosh file type code(s): 'TEXT'

Person & email address to contact for further information: Hannes
Tschofenig <Hannes.Tschofenig@gmx.net>

Intended usage: LIMITED USE

Restrictions on usage: None

Author: Hannes Tschofenig <Hannes.Tschofenig@gmx.net>

Change controller:

This specification is a work item of the IETF ECRIT working group,
with mailing list address <ecrit@ietf.org>.

Change controller:

The IESG <iesg@ietf.org>

10.2. LoST Sync Relax NG Schema Registration

Please register the schema defined in this document under the XML
schema registry at
<http://www.iana.org/assignments/xml-registry/schema.html>

URI: urn:ietf:params:xml:schema:lostsync1

Registrant Contact: IETF ECRIT Working Group, Hannes Tschofenig
(Hannes.Tschofenig@gmx.net).

Relax NG Schema: The Relax NG schema to be registered is contained in Section 7.

10.3. LoST Synchronization Namespace Registration

Please register the namespace defined in this document under the XML namespace registry at
<http://www.iana.org/assignments/xml-registry/ns.html>

URI: `urn:ietf:params:xml:ns:lostsync1`

Registrant Contact: IETF ECRIT Working Group, Hannes Tschofenig
(Hannes.Tschofenig@gmx.net).

XML:

```
BEGIN
<?xml version="1.0"?>
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML Basic 1.0//EN"
  "http://www.w3.org/TR/xhtml-basic/xhtml-basic10.dtd">
<html xmlns="http://www.w3.org/1999/xhtml">
<head>
  <meta http-equiv="content-type"
        content="text/html; charset=iso-8859-1"/>
  <title>LoST Synchronization Namespace</title>
</head>
<body>
  <h1>Namespace for LoST server synchronization</h1>
  <h2>urn:ietf:params:xml:ns:lostsync1</h2>
  <p>See <a href="[URL of published RFC]">RFCXXXX
    [NOTE TO IANA/RFC-EDITOR:
      Please replace XXXX with the RFC number of this
      specification.]</a>.</p>
</body>
</html>
END
```

11. Acknowledgments

Robins George, Cullen Jennings, Karl Heinz Wolf, Richard Barnes, Mayutan Arumaithurai, Alexander Mayrhofer, and Andrew Newton provided helpful input. Jari Urpalainen assisted with the Relax NG schema. We would also like to thank our document shepherd Roger Marshall for his help with the document.

We would like to particularly thank Andrew Newton for his timely and valuable review of the XML-related content.

We would like to thank Robert Sparks, Barry Leiba, Stephen Farrell, Brian Haberman, Pete Resnick, and Sean Turner for their AD reviews. We would also like to thank Bjoern Hoehrmann for his media type review, Julian Reschke and Martin Duerst for their applications area reviews, and Wassim Haddad for his Gen-ART review.

12. References

12.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2246] Dierks, T. and C. Allen, "The TLS Protocol Version 1.0", RFC 2246, January 1999.
- [RFC2616] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", RFC 2616, June 1999.
- [RFC2617] Franks, J., Hallam-Baker, P., Hostetler, J., Lawrence, S., Leach, P., Luotonen, A., and L. Stewart, "HTTP Authentication: Basic and Digest Access Authentication", RFC 2617, June 1999.
- [RFC2818] Rescorla, E., "HTTP Over TLS", RFC 2818, May 2000.
- [RFC3023] Murata, M., St. Laurent, S., and D. Kohn, "XML Media Types", RFC 3023, January 2001.
- [RFC4288] Freed, N. and J. Klensin, "Media Type Specifications and Registration Procedures", BCP 13, RFC 4288, December 2005.
- [RFC5222] Hardie, T., Newton, A., Schulzrinne, H., and H. Tschofenig, "LoST: A Location-to-Service Translation Protocol", RFC 5222, August 2008.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, August 2008.
- [W3C.REC-xmlsig-core-20020212] Eastlake, D., Reagle, J., Solo, D., Hirsch, F., and T. Roessler, "XML-Signature Syntax and Processing", World Wide Web Consortium Second Edition REC-xmlsig-core-20020212, June 2008.

12.2. Informative References

- [RFC5582] Schulzrinne, H., "Location-to-URL Mapping Architecture and Framework", RFC 5582, September 2009.

Authors' Addresses

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

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

Hannes Tschofenig
Nokia Siemens Networks
Linnoitustie 6
Espoo 02600
Finland

Phone: +358 (50) 4871445
Email: Hannes.Tschofenig@gmx.net
URI: <http://www.tschofenig.priv.at>

ecrit
Internet-Draft
Intended status: BCP
Expires: March 10, 2012

B. Rosen
NeuStar
J. Polk
Cisco Systems
September 7, 2011

Best Current Practice for Communications Services in support of
Emergency Calling
draft-ietf-ecrit-phonebcp-20.txt

Abstract

The IETF and other standards organization have efforts targeted at standardizing various aspects of placing emergency calls on IP networks. This memo describes best current practice on how devices, networks and services using IETF protocols should use such standards to make emergency calls.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on March 10, 2012.

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. Terminology	4
2. Introduction	4
3. Overview of how emergency calls are placed	4
4. Which devices and services should support emergency calls	5
5. Identifying an emergency call	5
6. Location and its role in an emergency call	6
6.1. Types of location information	7
6.2. Location Determination	7
6.2.1. User-entered location information	7
6.2.2. Access network "wire database" location information	7
6.2.3. End-system measured location information	8
6.2.4. Network-measured location information	8
6.3. Who adds location, endpoint or proxy	9
6.4. Location and references to location	9
6.5. End system location configuration	9
6.6. When location should be configured	10
6.7. Conveying location	11
6.8. Location updates	12
6.9. Multiple locations	12
6.10. Location validation	13
6.11. Default location	13
6.12. Other location considerations	13
7. LIS and LoST Discovery	14
8. Routing the call to the PSAP	14
9. Signaling of emergency calls	15
9.1. Use of TLS	15
9.2. SIP signaling requirements for User Agents	16
9.3. SIP signaling requirements for proxy servers	17
10. Call backs	18
11. Mid-call behavior	18
12. Call termination	18
13. Disabling of features	18
14. Media	19
15. Testing	20
16. Security Considerations	21
17. IANA Considerations	21
17.1. test service urn	21
17.2. 'test' Subregistry	21
18. Acknowledgements	22
19. References	22
19.1. Normative References	22
19.2. Informative References	25

Authors' Addresses	26
------------------------------	----

1. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

This document uses terms from [RFC3261], [RFC5012] and [I-D.ietf-ecrit-framework].

2. Introduction

This document describes how access networks, Session Initiation Protocol [RFC3261] user agents, proxy servers and Public Safety Access Points (PSAPs) support emergency calling, as outlined in [I-D.ietf-ecrit-framework], which is designed to complement the present document in section headings, numbering and content. Understanding [I-D.ietf-ecrit-framework] is necessary to understand this document. This BCP succinctly describes the requirements of end devices and applications (requirements prefaced by "ED-"), access networks (including enterprise access networks) (requirements prefaced by "AN-"), service providers (requirements prefaced by "SP-") and PSAPs to achieve globally interoperable emergency calling on the Internet.

This document also defines requirements for "Intermediate" devices which exist between end devices or applications and the access network. For example, a home router is an "Intermediate" device. Reporting location on an emergency call (see Section 6) may depend on the ability of such intermediate devices to meet the requirements prefaced by "INT-".

The access network requirements apply to those networks which may be used to place emergency calls using IETF protocols. Local regulations may impact the need to support this document's access network requirements.

Other organizations, such as the North American Emergency Number Association (NENA), define the PSAP interface. NENA's documents reference this document.

3. Overview of how emergency calls are placed

An emergency call can be distinguished (Section 5) from any other call by a unique Service URN [RFC5031], which is placed in the call set-up signaling when a home or visited emergency dial string is

detected. Because emergency services are local to specific geographic regions, a caller must obtain his location (Section 6) prior to making emergency calls. To get this location, either a form of measuring (e.g., GPS) (Section 6.2.3) device location in the endpoint is deployed, or the endpoint is configured (Section 6.5) with its location from the access network's Location Information Server (LIS). The location is conveyed (Section 6.7) in the SIP signaling with the call. The call is routed (Section 8) based on location using the Location-to-Service Translation (LoST) protocol [RFC5222], which maps a location to a set of PSAP URIs. Each URI resolves to a PSAP or an Emergency Services Routing Proxy (ESRP), which serves a group of PSAPs. The call arrives at the PSAP with the location included in the SIP INVITE request.

4. Which devices and services should support emergency calls

ED-1 A device or application that implements SIP calling SHOULD support emergency calling. Some jurisdictions have regulations governing which devices need to support emergency calling and developers are encouraged to ensure that devices they develop meet relevant regulatory requirements. Unfortunately, the natural variation in those regulations also makes it impossible to accurately describe the cases when developers do or do not have to support emergency calling.

SP-1 If a device or application expects to be able to place a call for help, the service provider that supports it MUST facilitate emergency calling. Some jurisdictions have regulations governing this.

ED-2 Devices that create media sessions and exchange real-time audio, video and/or text, have the capability to establish sessions to a wide variety of addresses, and communicate over private IP networks or the Internet, SHOULD support emergency calls. Some jurisdictions have regulations governing this.

5. Identifying an emergency call

ED-3 Endpoints SHOULD recognize dial strings of emergency calls. If the service provider always knows the location of the device (the correct dial string depends on which country you are in), the service provider may recognize them, see SP-2.

SP-2 Proxy servers SHOULD recognize emergency dial strings if for some reason the endpoint does not recognize them.

ED-4/SP-3 Emergency calls MUST be marked with a Service URN in the Request-URI of the INVITE.

ED-5/SP-4 Geographically local dial strings MUST be recognized.

ED-6/SP-5 Devices MUST be able to be configured with the home country from which the home dial string(s) can be determined.

ED-7/SP-6 Emergency dial strings SHOULD be determined from LoST [RFC5222]. Dial Strings MAY be configured directly into the device.

AN-1 LoST servers MUST return dial strings for emergency services.

ED-8 Endpoints which do not recognize emergency dial strings SHOULD send dial strings as per [RFC4967].

SP-7 If a proxy server recognizes dial strings on behalf of its clients, it MUST recognize emergency dial strings represented by [RFC4967] and SHOULD recognize the emergency dial strings represented by a tel URI [RFC3966].

ED-9 Endpoints SHOULD be able to have home dial strings provisioned.

SP-8 Service providers MAY provision home dial strings in devices.

ED-10 Devices SHOULD NOT have one button emergency calling initiation.

ED-11/SP-9 All sub-services for the 'sos' service specified in [RFC5031]. MUST be recognized.

6. Location and its role in an emergency call

Handling location for emergency calling usually involves several steps to process and multiple entities are involved. In Internet emergency calling, where the endpoint is located is "determined" using a variety of measurement or wiretracing methods. Endpoints can be "configured" with their own location by the access network. In some circumstances, a proxy server can insert location into the signaling on behalf of the endpoint. The location is "mapped" to the URI to send the call to, and the location is "conveyed" to the PSAP (and other entities) in the signaling. Likewise, we employ Location Configuration Protocols (LCPs), the Location-to-Service Mapping Protocol, and Location Conveyance Protocols for these functions. The Location-to-Service Translation protocol [RFC5222] is the Location Mapping Protocol defined by the IETF.

6.1. Types of location information

There are several forms of location. All IETF location configuration and location conveyance protocols support both civic and geospatial (geo) forms. The civic forms include both postal and jurisdictional fields. A cell tower/sector can be represented as a point (geo or civic) or polygon. Endpoints, Intermediate Devices and Service Providers receiving other forms of location representation MUST map them into either a geo or civic for use in emergency calls.

ED-12/INT-1/SP-10 Endpoints, Intermediate Devices and Service Providers MUST be prepared to handle location represented in either civic or geo form.

ED-13/INT-2/SP-11/AN-2 Entities MUST NOT convert (civic to geo or geo to civic) from the form of location the determination mechanism (see Section Section 6.2) supplied prior to receipt by the PSAP.

6.2. Location Determination

ED-14/INT-3/AN-3 Any location determination mechanism MAY be used, provided the accuracy of the location meets local requirements.

6.2.1. User-entered location information

ED-15/INT-4/AN-4 Devices, intermediate Devices and/or access networks SHOULD support a manual method to override the location the access network determines. When the override location is supplied in civic form, it MUST be possible for the resultant Presence Information Data Format - Location Object (PIDF-LO) received at the PSAP to contain any of the elements specified in [RFC4119] and [RFC5139].

6.2.2. Access network "wire database" location information

AN-5 Access networks supporting copper, fiber or other hard wired IP packet service SHOULD support location configuration. If the network does not support location configuration, it MUST require every device or intermediate device that connects to the network to support end system measured location.

AN-6/INT-5 Access networks and intermediate devices providing wire database location information SHOULD provide interior location data (building, floor, room, cubicle) where possible. It is RECOMMENDED that interior location be provided when spaces exceed approximately 650 square meters. See [I-D.ietf-ecrit-framework] Section 6.2.2 for a discussion of how this value was determined.

AN-7/INT-6 Access networks and intermediate devices (including

enterprise networks) which support intermediate range wireless connections (typically 100m or less of range) and which do not support a more accurate location determination mechanism such as triangulation, MUST support location configuration where the location of the access point is reflected as the location of the clients of that access point.

AN-8/INT-7 Where the access network provides location configuration, intermediate devices MUST either be transparent to it, or provide an interconnected client for the supported configuration mechanism and a server for a configuration protocol supported by end devices downstream of the intermediate device such that the location provided by the access network is available to clients as if the intermediate device was not in the path.

6.2.3. End-system measured location information

ED-16/INT-8 Devices MAY support end-system measured location. See [I-D.ietf-ecrit-framework] Section 6 for a discussion of accuracy of location.

ED-17/INT-9/AN-9 Devices that support endpoint measuring of location MUST have at least a coarse location capability (typically <1km accuracy) for routing of calls. The location mechanism MAY be a service provided by the access network.

6.2.4. Network-measured location information

AN-10 Access networks MAY provide network-measured location determination. Wireless access networks that do not supply network measured location MUST require every device or intermediate device connected to the network to support end-system measured location. Uncertainty and confidence may be specified by local regulation. Where not specified, uncertainty of less than 100 meters with 95% confidence is RECOMMENDED for dispatch location.

AN-11 Access networks that provide network measured location MUST have at least a coarse location (typically <1km when not location hiding) capability at all times for routing of calls.

AN-12 Access networks with range of <10 meters (e.g. personal area networks such as Bluetooth MUST provide a location to mobile devices connected to them. The location provided SHOULD be that reported by the upstream access network unless a more accurate mechanism is available.

6.3. Who adds location, endpoint or proxy

ED-18/INT-10 Endpoints SHOULD attempt to configure their own location using the Location Configuration Protocols (LCPs) listed in ED-21.

SP-12 Proxies MAY provide location on behalf of devices if:

- o The proxy has a relationship with all access networks the device could connect to, and the relationship allows it to obtain location.
- o The proxy has an identifier, such as an IP address, that can be used by the access network to determine the location of the endpoint, even in the presence of NAT and VPN tunnels that may obscure the identifier between the access network and the service provider.

ED-19/INT-11/SP-13 Where proxies provide location on behalf of endpoints, the service provider MUST ensure that either the end device is provided with the local dial strings for its current location (where the end device recognizes dial strings), or the service provider proxy MUST detect the appropriate local dial strings at the time of the call.

6.4. Location and references to location

ED-20/INT-12 Devices SHOULD be able to accept and forward location by value or by reference. An end device that receives location by reference (and does not also get the corresponding value) MUST be able to perform a dereference operation to obtain a value.

6.5. End system location configuration

Obtaining location from the access network may be preferable even if the device can measure its own location, especially indoors where most measurement mechanisms are not accurate enough. This sections requirements do not apply to devices that can accurately measure their own location.

ED-21/INT-13 Devices MUST support both the Dynamic Host Configuration Protocol (DHCP) location options [RFC4776], [RFC6225] and HTTP Enabled Location Delivery (HELD) [RFC5985]. When devices deploy a specific access network interface for which location configuration mechanisms such as Link Layer Discovery Protocol - Media Endpoint Discovery (LLDP-MED) [LLDP-MED] or 802.11v are specified, the device SHOULD support the additional respective access network specific location configuration mechanism.

AN-13/INT-14 The access network MUST support either DHCP location options or HELD. The access network SHOULD support other location

configuration technologies that are specific to the type of access network.

AN-14/INT-15 Where a router is employed between a LAN and WAN in a small (less than approximately 650 square meters) area, the router MUST be transparent to the location provided by the WAN to the LAN. This may mean the router must obtain location as a client from the WAN, and supply an LCP server to the LAN with the location it obtains. Where the area is larger, the LAN MUST have a location configuration mechanism satisfying the requirements of this document.

ED-22/INT-16 Endpoints SHOULD try all LCPs supported by the device in any order or in parallel. The first one that succeeds in supplying location MUST be used.

AN-15/INT-17 Access networks that support more than one LCP MUST reply with the same location information (within the limits of the data format for the specific LCP) for all LCPs it supports.

ED-23/INT-18/SP-14 When HELD is the LCP, the request MUST specify a value of "emergencyRouting" for the "responseTime" parameter and use the resulting location for routing. If a value for dispatch location will be sent, another request with the "responseTime" parameter set to "emergencyDispatch" must be completed, with the result sent for dispatch purposes.

ED-24 Where the operating system supporting application programs which need location for emergency calls does not allow access to Layer 2 and Layer 3 functions necessary for a client application to use DHCP location options and/or other location technologies that are specific to the type of access network, the operating system MUST provide a published API conforming to ED-12 through ED-23 and ED-25 through ED-32. It is RECOMMENDED that all operating systems provide such an API.

6.6. When location should be configured

If an endpoint is manually configured, the requirements in this section are not applicable.

ED-25/INT-19 Endpoints SHOULD obtain location immediately after obtaining local network configuration information.

ED-26/INT-20 If the device is configured to use DHCP for bootstrapping, and does not use it MUST include both options for location acquisition (civic and geodetic), the option for LIS discovery, and the option for LoST discovery as defined in [RFC4776], [RFC6225], [RFC5986] and [RFC5223].

ED-27/INT-21 If the device sends a DHCPINFORM message, it MUST include both options for location acquisition (civic and geodetic), the option for LIS discovery, and the option for LoST discovery as defined in [RFC4776], [RFC6225], [RFC5986] and [RFC5223].

ED-28/INT-22 To minimize the effects of VPNs that do not allow packets to be sent via the native hardware interface rather than via the VPN tunnel, location configuration SHOULD be attempted before such tunnels are established.

ED-29/INT-23 Software which uses LCPs SHOULD locate and use the actual hardware network interface rather than a VPN tunnel interface to direct LCP requests to the LIS in the actual access network.

AN-16 Network administrators MUST take care in assigning IP addresses such that VPN address assignments can be distinguished from local devices (by subnet choice, for example), and LISs SHOULD NOT attempt to provide location to addresses that arrive via VPN connections unless it can accurately determine the location for such addresses.

AN-17 Placement of NAT devices where an LCP uses IP address for an identifier SHOULD consider the effect of the NAT on the LCP. The address used to query the LIS MUST be able to correctly identify the record in the LIS representing the location of the querying device

ED-30/INT-24 For devices which are not expected to change location, refreshing location on the order of once per day is RECOMMENDED.

ED-31/INT-25 For devices which roam, refresh of location information SHOULD be more frequent, with the frequency related to the mobility of the device and the ability of the access network to support the refresh operation. If the device detects a link state change that might indicate having moved, for example when it changes access points, the device SHOULD refresh its location.

ED-32/INT-26/AN-18 It is RECOMMENDED that location determination not take longer than 250 ms to obtain routing location and systems SHOULD be designed such that the typical response is under 100 ms. However, as much as 3 seconds to obtain routing location MAY be tolerated if location accuracy can be substantially improved over what can be obtained in 250 ms.

6.7. Conveying location

ED-33/SP-15 Location sent between SIP entities MUST be conveyed using [I-D.ietf-sipcore-location-conveyance].

6.8. Location updates

ED-34/AN-19 Where the absolute location or the accuracy of location of the endpoint may change between the time the call is received at the PSAP and the time dispatch is completed, location update mechanisms **MUST** be implemented and used.

ED-35/AN-20 Mobile devices **MUST** be provided with a mechanism to get repeated location updates to track the motion of the device during the complete processing of the call.

ED-36/AN-21 The LIS **SHOULD** provide a location reference which permits a subscription with appropriate filtering.

ED-37/AN-22 For calls sent with location-by-reference, with a SIP or SIPS scheme, the server resolving the reference **MUST** support a SUBSCRIBE [RFC3265] to the presence event [RFC3856]. For other location-by-reference schemes that do not support subscription, the PSAP will have to repeatedly dereference the URI to determine if the device moved.

ED-38 If location was sent by value, and the endpoint gets updated location, it **MUST** send the updated location to the PSAP via a SIP re-INVITE or UPDATE request. Such updates **SHOULD** be limited to no more than one update every 10 seconds, a value selected to keep the load on a large PSAP manageable, and yet provide sufficient indication to the PSAP of motion.

6.9. Multiple locations

ED-39/SP-16 If the LIS has more than one location for an endpoint it **MUST** conform to the rules in Section 3 of [RFC5491]

ED-40 If an endpoint has more than one location available to it, it **MUST** choose one location to route the call towards the PSAP. If multiple locations are in a single Presence Information Data Format (PIDF), the procedures in [RFC5491] **MUST** be followed. If the endpoint has multiple PIDFs, and has no reasonable basis to choose from among them, a random choice is acceptable.

SP-17 If a proxy inserts location on behalf of an endpoint, and it has multiple locations available for the endpoint it **MUST** choose one location to use to route the call towards the PSAP. If multiple locations are in a single PIDF, the procedures in [RFC5491] **MUST** be followed. If the proxy has multiple PIDFs, and has no reasonable basis to choose from among them, a random choice is acceptable.

SP-18 If a proxy is attempting to insert location but the endpoint

conveyed a location to it, the proxy MUST use the endpoint's location for routing in the initial INVITE and MUST convey that location towards the PSAP. It MAY also include what it believes the location to be in a separate Geolocation header.

SP-19 All location objects received by a proxy MUST be delivered to the PSAP.

ED-41/SP-20 Location objects MUST be created with information about the method by which the location was determined, such as GPS, manually entered, or based on access network topology included in a PIDF-LO "method" element. In addition, the source of the location information MUST be included in a PIDF-LO "provided-by" element.

ED-42/SP-21 A location with a method of "derived" MUST NOT be used unless no other location is available.

6.10. Location validation

AN-23 A LIS SHOULD perform location validation of civic locations via LoST before entering a location in its database.

ED-44 Endpoints SHOULD validate civic locations when they receive them from their LCP. Validation SHOULD be performed in conjunction with the LoST route query to minimize load on the LoST server.

6.11. Default location

AN-24 When the access network cannot determine the actual location of the caller, it MUST supply a default location. The default SHOULD be chosen to be as close to the probable location of the device as the network can determine. See [I-D.ietf-ecrit-framework]

SP-22 Proxies handling emergency calls MUST insert a default location in the INVITE if the incoming INVITE does not contain a location and the proxy does not have a method for obtaining a better location.

AN-25/SP-23 Default locations MUST be marked with method=Default and the proxy MUST be identified in provided-by element of the PIDF-LO.

6.12. Other location considerations

ED-45 If the LCP does not return location in the form of a PIDF-LO [RFC4119], the endpoint MUST map the location information it receives from the configuration protocol to a PIDF-LO.

ED-46/AN-26 To prevent against spoofing of the DHCP server, entities implementing DHCP for location configuration SHOULD use [RFC3118],

although the difficulty in providing appropriate credentials is significant.

ED-47 If S/MIME [RFC5751] is used, the INVITE message MUST provide enough information unencrypted for intermediate proxies to route the call based on the location information included. This would include the Geolocation header, and any bodies containing location information. Use of S/MIME with emergency calls is NOT RECOMMENDED for this reason.

ED-48/SP-24 TLS [RFC5746] MUST be used to protect location (but see Section 9.1). All implementations MUST support TLS.

7. LIS and LoST Discovery

ED-49 Endpoints MUST support one or more mechanisms that allow them to determine their public IP address, for example, STUN [RFC5389].

ED-50 Endpoints MUST support LIS discovery as described in [RFC5986], and the LoST discovery as described in [RFC5223].

ED-51 The device MUST have a configurable default LoST server parameter.

ED-52 DHCP LoST discovery MUST be used, if available, in preference to configured LoST servers. That is, the endpoint MUST send queries to this LoST server first, using other LoST servers only if these queries fail.

AN-27 Access networks which support DHCP MUST implement the LIS and LoST discovery options in their DHCP servers and return suitable server addresses as appropriate.

8. Routing the call to the PSAP

ED-53 Endpoints who obtain their own location SHOULD perform LoST mapping to the PSAP URI.

ED-54 Mapping SHOULD be performed at boot time and whenever location changes beyond the service boundary obtained from a prior LoST mapping operation or the time-to-live value of that response has expired. The value MUST be cached for possible later use.

ED-55 The endpoint MUST attempt to update its location at the time of an emergency call. If it cannot obtain a new location quickly (see Section 6), it MUST use the cached value.

ED-56 The endpoint SHOULD attempt to update the LoST mapping at the time of an emergency call. If it cannot obtain a new mapping quickly, it MUST use the cached value. If the device cannot update the LoST mapping and does not have a cached value, it MUST signal an emergency call without a Route header containing a PSAP URI.

SP-25 Networks MUST be designed so that at least one proxy in the outbound path will recognize emergency calls with a Request URI of the service URN in the "sos" tree. An endpoint places a service URN in the Request URI to indicate that the endpoint understood the call was an emergency call. A proxy that processes such a call looks for the presence of a SIP Route header field with a URI of a PSAP. Absence of such a Route header indicates the endpoint was unable to invoke LoST and the proxy MUST perform the LoST mapping and insert a Route header field with the URI obtained.

SP-26 To deal with old user agents that predate this specification and with endpoints that do not have access to their own location data, a proxy that recognizes a call as an emergency call that is not marked as such (see Section 5) MUST also perform this mapping, with the best location it has available for the endpoint. The resulting PSAP URI would be placed in a Route header with the service URN in the Request URI.

SP-27 Proxy servers performing mapping SHOULD use location obtained from the access network for the mapping. If no location is available, a default location (see Section 6.11) MUST be supplied.

SP-28 A proxy server which attempts mapping and fails to get a mapping MUST provide a default mapping. A suitable default mapping would be the mapping obtained previously for the default location appropriate for the caller.

ED-57/SP-29 [RFC3261] and [RFC3263] procedures MUST be used to route an emergency call towards the PSAP's URI.

9. Signaling of emergency calls

9.1. Use of TLS

ED-58/SP-30 TLS is the primary mechanism used to secure the signaling for emergency calls. IPsec [RFC4301] MAY be used instead of TLS for any hop. Either TLS or IPSEC MUST be used when attempting to signal an emergency call.

ED-59/SP-31 If TLS session establishment is not available, or fails, the call MUST be retried without TLS.

ED-60/SP-32 [RFC5626] is RECOMMENDED to maintain persistent TLS connections between entities when one of the entity is an endpoint. Persistent TLS connection between proxies is RECOMMENDED using any suitable mechanism.

ED-61/AN-28 TLS SHOULD be used when attempting to retrieve location (configuration or dereferencing) with HELD. The use of [RFC5077] is RECOMMENDED to minimize the time to establish TLS sessions without keeping server-side state. IPsec MAY be used instead of TLS.

ED-62/AN-29 When TLS session establishment fails, the location retrieval MUST be retried without TLS.

9.2. SIP signaling requirements for User Agents

ED-63 The initial SIP signaling method is an INVITE request:

1. The Request URI SHOULD be the service URN in the "sos" tree. If the device does not interpret local dial strings, the Request-URI MUST be a dial string URI [RFC4967] with the dialed digits.
2. The To header field SHOULD be a service URN in the "sos" tree. If the device does not interpret local dial strings, the To: MUST be a dial string URI with the dialed digits.
3. The From header field SHOULD contain the AoR of the caller.
4. A Route header field SHOULD be present with a PSAP URI obtained from LoST (see Section 8). If the device does not interpret dial plans, or was unable to obtain a route from a LoST server, no such Route header field will be present.
5. A Contact header field MUST be globally routable, for example a GRUU [RFC5627], and be valid for several minutes following the termination of the call, provided that the UAC remains registered with the same registrar, to permit an immediate call-back to the specific device which placed the emergency call. It is acceptable if the UAC inserts a locally routable URI and a subsequent B2BUA maps that to a globally routable URI.
6. Other header fields MAY be included as per normal SIP behavior.
7. If a geolocation URI is included in the INVITE, a Supported header field MUST be included with a 'geolocation-sip' or 'geolocation-http' option tag, as appropriate. [I-D.ietf-sipcore-location-conveyance].
8. If a device understands the SIP location conveyance [I-D.ietf-sipcore-location-conveyance] extension and has its location available, it MUST include location either by-value, by-reference or both.
9. A SDP offer SHOULD be included in the INVITE. If voice is supported the offer SHOULD include the G.711 codec, see Section 14. As PSAPs may support a wide range of media types and codecs, sending an offerless INVITE may result in a lengthy return offer, but is permitted. Cautions in [RFC3261] on

offerless INVITEs should be considered before such use.

10. If the device includes location-by-value, the UA MUST support multipart message bodies, since SDP will likely be also in the INVITE.

9.3. SIP signaling requirements for proxy servers

SP-33 SIP Proxy servers processing emergency calls:

1. If the proxy interprets dial plans on behalf of user agents, the proxy MUST look for the local emergency dial string at the location of the end device and MAY look for the home dial string. If it finds it, the proxy MUST:
 - * Insert a Geolocation header field. Location-by-reference MUST be used because proxies must not insert bodies.
 - * Insert the Geolocation-Routing header with appropriate parameters .
 - * Map the location to a PSAP URI using LoST.
 - * Add a Route header with the PSAP URI.
 - * Replace the Request-URI (which was the dial string) with the service URN appropriate for the emergency dial string.
 - * Route the call using normal SIP routing mechanisms.
2. If the proxy recognizes the service URN in the Request URI, and does not find a Route header, it MUST query a LoST server immediately. If a location was provided (which should be the case), the proxy uses that location to query LoST. The proxy may have to dereference a location by reference to get a value. If a location is not present, and the proxy can query a LIS which has the location of the UA it MUST do so. If no location is present, and the proxy does not have access to a LIS which could provide location, the proxy MUST supply a default location (See Section 6.11). The location (in the signaling, obtained from a LIS, or default) MUST be used in a query to LoST with the service URN received with the call. The resulting URI MUST be placed in a Route header added to the call.
3. The proxy MAY add a Geolocation header field. Such an additional location SHOULD NOT be used for routing; the location provided by the UA should be used.
4. Either a P-Asserted-Identity [RFC3325] or an Identity header field [RFC4474], or both, SHOULD be included to identify the sender. For services which must support emergency calls from unauthenticated devices, valid identity may not be available. Proxies encountering a P-Asserted-Identity will need to pass the header to the PSAP, which is in a different domain. [RFC3325] requires a "spec(T)" to determine what happens if the "id" privacy service, or a Privacy header is present and requests privacy. In the absence of another spec(T), such proxies should pass the header unmodified if and only if the connection between the proxy and the PSAP is, as far as the proxy can determine,

protected by TLS with mutual authentication using keys reliably known by the parties, encrypted with no less strength than AES and the local regulations governing the PSAP do not otherwise specify.

5. Proxies SHOULD NOT return a 424 error. It should process the INVITE as best as it can.
6. Proxies SHOULD NOT obey a Geolocation-Routing value of "no" or a missing value if the proxy must query LoST to obtain a route. Emergency calls are always routed by location.

10. Call backs

ED-64/SP-34 Devices device SHOULD have a globally routable URI in a Contact: header field which remains valid for several minutes past the time the original call containing the URI completes unless the device registration expires and is not renewed.

SP-35 Call backs to the Contact: header URI received within 30 minutes of an emergency call must reach the device regardless of call features or services that would normally cause the call to be routed to some other entity.

SP-36 Devices MUST have a persistent AOR URI either in a P-Asserted-Identity header field or From protected by an Identity header field suitable for returning a call some time after the original call. Such a call back would not necessarily reach the device that originally placed the call.

11. Mid-call behavior

ED-65/SP-37 During the course of an emergency call, devices and proxies MUST initiate a call transfer upon receipt of REFER request within the dialog with method=INVITE and the Referred-by header field [RFC3515] in that request.

12. Call termination

ED-66 Normal [RFC3261] procedures for termination MUST be used for termination of the call.

13. Disabling of features

ED-67/SP-38 User Agents and proxies MUST disable features that will interrupt an ongoing emergency call, such as:

- o Call Waiting
- o Call Transfer
- o Three Way Call
- o Hold
- o Outbound Call Blocking

when an emergency call is established, but see ED-66 with respect to Call Waiting. Also see ED-74 in Section 14.

ED-68/SP-39 The emergency dial strings SHOULD NOT be permitted in Call Forward numbers or speed dial lists.

ED-69/SP-40 The User Agent and Proxies MUST disable call features which would interfere with the ability of call backs from the PSAP to be completed such as:

- o Do Not Disturb
- o Call Forward (all kinds)

These features SHOULD be disabled for approximately 30 minutes following termination of an emergency call.

ED-70 Call backs SHOULD be determined by retaining the domain of the PSAP which answers an outgoing emergency call and instantiating a timer which starts when the call is terminated. If a call is received from the same domain and within the timer period, sent to the Contact: or AoR used in the emergency call, it should be assumed to be a call back. The suggested timer period is 5 minutes. [RFC4916] may be used by the PSAP to inform the endpoint of the domain of the PSAP. Recognizing a call back from the domain of the PSAP will not always work, and further standardization will be required to give the endpoint the ability to recognize a call back.

14. Media

ED-71 Endpoints MUST send and receive media streams on RTP [RFC3550].

ED-72 Normal SIP offer/answer [RFC3264] negotiations MUST be used to agree on the media streams to be used.

ED-73/SP-41 G.711 A law (and mu Law if they are intended be used in North America) encoded voice as described in [RFC3551] MUST be supported. If the endpoint cannot support G.711, a transcoder MUST be used so that the offer received at the PSAP contains G.711. It is desirable to include wideband codecs such as G.722 and AMR-WB in the offer. PSAPs SHOULD support narrowband codecs common on endpoints in their area to avoid transcoding.

ED-74 Silence suppression (Voice Activity Detection methods) MUST NOT be used on emergency calls. PSAP call takers sometimes get

information on what is happening in the background to determine how to process the call.

ED-75 Endpoints supporting Instant Messaging (IM) MUST support either [RFC3428] and [RFC4975].

ED-76 Endpoints supporting real-time text MUST use [RFC4103]. The expectations for emergency service support for the real-time text medium are described in [RFC5194], Section 7.1.

ED-77 Endpoints supporting video MUST support H.264 per [RFC6184].

15. Testing

ED-78 INVITE requests to a service URN starting with "test." indicates a request for an automated test. For example, "urn:service:test.sos.fire". As in standard SIP, a 200 (OK) response indicates that the address was recognized and a 404 (Not found) that it was not. A 486 (Busy Here) MUST be returned if the test service is busy, and a 404 (Not found) MUST be returned if the PSAP does not support the test mechanism.

ED-79 In its response to the test, the PSAP MAY include a text body (text/plain) indicating the identity of the PSAP, the requested service, and the location reported with the call. For the latter, the PSAP SHOULD return location-by-value even if the original location delivered with the test was by-reference. If the location-by-reference was supplied, and the dereference requires credentials, the PSAP SHOULD use credentials supplied by the LIS for test purposes. This alerts the LIS that the dereference is not for an actual emergency call and location hiding techniques, if they are being used, may be employed for this dereference. Use of SIPS for the request would assure the response containing the location is kept private

ED-80 A PSAP accepting a test call SHOULD accept a media loopback test [I-D.ietf-mmusic-media-loopback] and SHOULD support the "rtp-pkt-loopback" and "rtp-start-loopback" options. The user agent would specify a loopback attribute of "loopback-source", the PSAP being the mirror. User Agents should expect the PSAP to loop back no more than 3 packets of each media type accepted (which limits the duration of the test), after which the PSAP would normally send BYE.

ED-81 User agents SHOULD perform a full call test, including media loopback, after a disconnect and subsequent change in IP address not due to a reboot. After an initial test, a full test SHOULD be repeated approximately every 30 days with a random interval.

ED-82 User agents MUST NOT place a test call immediately after booting. If the IP address changes after booting, the endpoint should wait a random amount of time (in perhaps a 30 minute period, sufficient for any avalanche restart to complete) and then test.

ED-83 PSAPs MAY refuse repeated requests for test from the same device in a short period of time. Any refusal is signaled with a 486 or 488 response.

16. Security Considerations

Security considerations for emergency calling have been documented in [RFC5069], and [RFC6280]. This document suggests that security (TLS or IPsec) be used hop by hop on a SIP call to protect location information, identity, etc. It also suggests that if the attempt to create a security association fails, the call be retried without the security. It's more important to get an emergency call through than to protect the data; indeed, in many jurisdictions privacy is explicitly waived when making emergency calls. Placing a call without security may reveal user information, including location. The alternative - failing the call if security cannot be established, is considered unacceptable.

17. IANA Considerations

This document registers service URNs in the Service URN Labels registry per [RFC5031] for testing.

17.1. test service urn

A new entry in the URN Service Label registry is added. The new service is "test", the reference is this document, and the description is "self test".

17.2. 'test' Subregistry

A new Subregistry is created, the "'test' Sub-Service. The registration process is Expert Review per [RFC5226]. The expert review should consider that the entries in this registry nominally track the entries in the sos sub registry, although it is not required that every entry in sos have an entry in test, and it is possible that entries in the test sub-registry not necessarily be in the sos sub registry. For example, testing of non-emergency URNs may be allowed. The Reference is this document. The initial content of the subregistry is:

Service	Reference	Description
test.sos	[this document]	test for sos
test.sos.ambulance	[this document]	test for sos.ambulance
test.sos.animal-control	[this document]	test for sos.animal-control
test.sos.fire	[this document]	test for sos.fire
test.sos.gas	[this document]	test for sos.gas
test.sos.marine	[this document]	test for sos.marine
test.sos.mountain	[this document]	test for sos.mountain
test.sos.physician	[this document]	test for sos.physician
test.sos.poison	[this document]	test for sos.poison
test.sos.police	[this document]	test for sos.police

18. Acknowledgements

Work group members participating in the creation and review of this document include Hannes Tschofenig, Ted Hardie, Marc Linsner, Roger Marshall, Stu Goldman, Shida Schubert, James Winterbottom, Barbara Stark, Richard Barnes and Peter Blatherwick.

19. References

19.1. Normative References

- [I-D.ietf-mmusic-media-loopback]
Sivachelvan, C., Venna, N., Jones, P., Stratton, N., Roychowdhury, A., and K. Hedayat, "An Extension to the Session Description Protocol (SDP) for Media Loopback", draft-ietf-mmusic-media-loopback-15 (work in progress), March 2011.
- [I-D.ietf-sipcore-location-conveyance]
Polk, J., Rosen, B., and J. Peterson, "Location Conveyance for the Session Initiation Protocol", draft-ietf-sipcore-location-conveyance-09 (work in progress), September 2011.
- [LLDP-MED]
TIA, "ANSI/TIA-1057 Link Layer Discovery Protocol - Media Endpoint Discovery".
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3118] Droms, R. and W. Arbaugh, "Authentication for DHCP Messages", RFC 3118, June 2001.

- [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.
- [RFC3263] Rosenberg, J. and H. Schulzrinne, "Session Initiation Protocol (SIP): Locating SIP Servers", RFC 3263, 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.
- [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.
- [RFC3856] Rosenberg, J., "A Presence Event Package for the Session Initiation Protocol (SIP)", RFC 3856, August 2004.
- [RFC3966] Schulzrinne, H., "The tel URI for Telephone Numbers", RFC 3966, December 2004.
- [RFC4103] Hellstrom, G. and P. Jones, "RTP Payload for Text Conversation", RFC 4103, June 2005.
- [RFC4119] Peterson, J., "A Presence-based GEOPRIV Location Object Format", RFC 4119, December 2005.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.
- [RFC4474] Peterson, J. and C. Jennings, "Enhancements for Authenticated Identity Management in the Session

Initiation Protocol (SIP)", RFC 4474, August 2006.

- [RFC4776] Schulzrinne, H., "Dynamic Host Configuration Protocol (DHCPv4 and DHCPv6) Option for Civic Addresses Configuration Information", RFC 4776, November 2006.
- [RFC4916] Elwell, J., "Connected Identity in the Session Initiation Protocol (SIP)", RFC 4916, June 2007.
- [RFC4967] Rosen, B., "Dial String Parameter for the Session Initiation Protocol Uniform Resource Identifier", RFC 4967, July 2007.
- [RFC4975] Campbell, B., Mahy, R., and C. Jennings, "The Message Session Relay Protocol (MSRP)", RFC 4975, September 2007.
- [RFC5031] Schulzrinne, H., "A Uniform Resource Name (URN) for Emergency and Other Well-Known Services", RFC 5031, January 2008.
- [RFC5139] Thomson, M. and J. Winterbottom, "Revised Civic Location Format for Presence Information Data Format Location Object (PIDF-LO)", RFC 5139, February 2008.
- [RFC5222] Hardie, T., Newton, A., Schulzrinne, H., and H. Tschofenig, "LoST: A Location-to-Service Translation Protocol", RFC 5222, August 2008.
- [RFC5223] Schulzrinne, H., Polk, J., and H. Tschofenig, "Discovering Location-to-Service Translation (LoST) Servers Using the Dynamic Host Configuration Protocol (DHCP)", RFC 5223, August 2008.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.
- [RFC5389] Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, "Session Traversal Utilities for NAT (STUN)", RFC 5389, October 2008.
- [RFC5491] Winterbottom, J., Thomson, M., and H. Tschofenig, "GEOPRIV Presence Information Data Format Location Object (PIDF-LO) Usage Clarification, Considerations, and Recommendations", RFC 5491, March 2009.
- [RFC5626] Jennings, C., Mahy, R., and F. Audet, "Managing Client-Initiated Connections in the Session Initiation Protocol

(SIP)", RFC 5626, October 2009.

- [RFC5627] Rosenberg, J., "Obtaining and Using Globally Routable User Agent URIs (GRUUs) in the Session Initiation Protocol (SIP)", RFC 5627, October 2009.
- [RFC5746] Rescorla, E., Ray, M., Dispensa, S., and N. Oskov, "Transport Layer Security (TLS) Renegotiation Indication Extension", RFC 5746, February 2010.
- [RFC5751] Ramsdell, B. and S. Turner, "Secure/Multipurpose Internet Mail Extensions (S/MIME) Version 3.2 Message Specification", RFC 5751, January 2010.
- [RFC5985] Barnes, M., "HTTP-Enabled Location Delivery (HELD)", RFC 5985, September 2010.
- [RFC5986] Thomson, M. and J. Winterbottom, "Discovering the Local Location Information Server (LIS)", RFC 5986, September 2010.
- [RFC6184] Wang, Y., Even, R., Kristensen, T., and R. Jesup, "RTP Payload Format for H.264 Video", RFC 6184, May 2011.
- [RFC6225] Polk, J., Linsner, M., Thomson, M., and B. Aboba, "Dynamic Host Configuration Protocol Options for Coordinate-Based Location Configuration Information", RFC 6225, July 2011.

19.2. Informative References

- [I-D.ietf-ecrit-framework]
Rosen, B., Schulzrinne, H., Polk, J., and A. Newton, "Framework for Emergency Calling using Internet Multimedia", draft-ietf-ecrit-framework-12 (work in progress), October 2010.
- [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.
- [RFC5012] Schulzrinne, H. and R. Marshall, "Requirements for Emergency Context Resolution with Internet Technologies", RFC 5012, January 2008.
- [RFC5069] Taylor, T., Tschafenig, H., Schulzrinne, H., and M. Shanmugam, "Security Threats and Requirements for Emergency Call Marking and Mapping", RFC 5069,

January 2008.

- [RFC5077] Salowey, J., Zhou, H., Eronen, P., and H. Tschofenig, "Transport Layer Security (TLS) Session Resumption without Server-Side State", RFC 5077, January 2008.
- [RFC5194] van Wijk, A. and G. Gybels, "Framework for Real-Time Text over IP Using the Session Initiation Protocol (SIP)", RFC 5194, June 2008.
- [RFC6280] Barnes, R., Lepinski, M., Cooper, A., Morris, J., Tschofenig, H., and H. Schulzrinne, "An Architecture for Location and Location Privacy in Internet Applications", BCP 160, RFC 6280, July 2011.

Authors' Addresses

Brian Rosen
NeuStar
470 Conrad Dr.
Mars, PA 16046
USA

Phone: +1 724 382 1051
Email: br@brianrosen.net

James Polk
Cisco Systems
3913 Treemont Circle
Colleyville, TX 76034
USA

Phone: +1-817-271-3552
Email: jmpolk@cisco.com

ECRIT Working Group
Internet-Draft
Intended status: Standards Track
Expires: May 13, 2011

M. Patel
InterDigital Communications
November 9, 2010

SOS Uniform Resource Identifier (URI) Parameter for Marking of Session
Initiation Protocol (SIP) Requests related to Emergency Services
draft-patel-ecrit-sos-parameter-11.txt

Abstract

This document defines a new Session Initiation Protocol (SIP) Uniform Resource Identifier (URI) parameter intended for marking SIP registration requests related to emergency calls and allow admission control to ensure successful initiation of emergency calls. The usage of this new URI parameter complements the usage of the Service Uniform Resource Name (URN) and is not intended to replace it.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on May 13, 2011.

Copyright Notice

Copyright (c) 2010 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. Terminology	4
3. Requirements	4
4. The "sos" URI Parameter	4
4.1. REGISTER Request	4
4.2. 2xx Response to REGISTER Request	5
4.3. Backwards compatibility issues	5
5. Formal Syntax	6
6. IANA Considerations	6
7. Security Considerations	6
8. Acknowledgements	7
9. References	7
9.1. Normative References	7
9.2. Informative References	8
Author's Address	8

1. Introduction

One way to differentiate a SIP-based emergency call from an ordinary call is by the presence of the Service URN as defined in RFC 5031 [RFC5031] (and used in the IETF emergency services architecture described in PhoneBCP[I-D.ietf-ecrit-phonebcpl]). The 3GPP IP Multimedia Subsystem (IMS) emergency services architecture, illustrated in 3GPP TS 23.167 [3GPP.23.167], specifies that the User Equipment (UE) needs to perform emergency registration prior to or during the initiation of an emergency call.

In some countries, it is a regulatory requirement that devices be able to place emergency calls in circumstances where other calls may not be permitted. When a UAC issues an emergency marked REGISTER request it indicates to the registrar that roaming and barring restrictions should not be applied for the registered address-of-record in order to successfully initiate an emergency session. Furthermore, distinguishing emergency registration from non-emergency registration allows the registrar to ensure that the contact address associated with previous registration of the address-of-record included in the emergency REGISTER request is not replaced.

Emergency registration is possible only when the UE has sufficient credentials to register with its home network and can detect that an emergency session is initiated. Unfortunately, marking of the emergency registration cannot be fulfilled by the use of the Service URN. The circumstances where such an emergency registration is beneficial are listed below:

- the UE is not registered with its home network;
- the UE is currently registered but roaming (to ensure that the emergency call is handled in the visited network, as required by some jurisdictions).

This document concentrates on a use case defined by 3GPP as described above. However, the solution proposed does not preclude other systems that require emergency registration to occur prior to placing an emergency call, to ensure that any subscription related restrictions are removed to allow successful initiation of emergency calls.

This document proposes a way to mark a REGISTER request as an emergency registration.

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 RFC 2119 [RFC2119]

3. Requirements

Req: Where emergency registration is required prior to placing an emergency call, it shall be possible to distinguish emergency registration from non-emergency registration.

4. The "sos" URI Parameter

This section provides an overview of the proposed new URI parameter to be used for marking REGISTER requests related to emergency services.

A new URI parameter "sos" is defined in this document. The "sos" parameter is appended to a URI consistent with RFC 3261 [RFC3261]. It is proposed that use of this URI parameter is restricted to the Contact header included in the REGISTER request (and the 2xx response to the REGISTER request) related to an emergency call only.

Inclusion of the "sos" URI parameter in a REGISTER request SHALL indicate that the REGISTER request pertains to emergency registration. The "sos" URI parameter MUST NOT be considered as a replacement for the Service URN for emergency calls originated by a UA.

4.1. REGISTER Request

In networks where the UA sends a REGISTER request for emergency registration prior to placing an emergency call, the "sos" URI parameter MUST be appended to the URI in the Contact header. This serves as an indication to the registrar that the request is for emergency registration thus requesting the registrar to not apply any restrictions to the user's service which might prevent emergency calls from successfully being initiated.

Example:

Contact: "Alice" <sip:alice@example.com;sos> ;q=0.7; expires=3600

In the event that more than one Contact header field is included in the REGISTER request, only the contact addresses that include the

"sos" URI parameter shall be considered as emergency registered contact addresses.

The "sos" URI parameter MUST NOT be included in non-REGISTER requests, and MUST NOT be included in REGISTER requests that do not pertain to emergency calls.

4.2. 2xx Response to REGISTER Request

If the registrar receives a REGISTER request that includes the "sos" URI parameter in the Contact header field, the registrar MUST include the "sos" URI parameter in the Contact header field in the 200 (OK) response sent by the registrar upon successful registration. The "sos" URI parameter is appended to the URI included in the Contact header.

4.3. Backwards compatibility issues

The backwards compatibility scenario considered in this document is where a legacy registrar does not support the "sos" URI parameter. In this case, if the registrar receives a REGISTER request that includes the "sos" URI parameter in the Contact header field, the registrar proceeds with registration procedures and silently ignores the URI-parameter in accordance with RFC 3261[RFC3261]. This ensures the user is registered and thus can successfully initiate an emergency call.

The drawback of proceeding with registration is if the address-of-record is for example barred or has roaming restrictions applied, then these restrictions will not be lifted and thus registration will be unsuccessful. This can limit the UA's ability to successfully place an emergency call.

If registration is successful, the 200 (OK) response from a legacy registrar includes the "sos" URI parameter in the Contact header field. Thus the UA is unaware that the registrar does not support the "sos" URI parameter. Providing the registration was successful, the UA's ability to place an emergency call is not compromised. The UA need not know that the registrar does not support the URI parameter.

The consequence of the registrar not supporting the "sos" URI parameter, in addition to the drawback pertaining to restrictions applied to the address-of-record, are as follows:

- the risk of the registrar overwriting previous registrations of the registered address-of-record, and thus disrupting any on-going non-emergency sessions associated with the UA, its address-of-record and

previously registered contact address.

- incoming calls, such as a PSAP call back (to a previously made emergency call) to the registered address-of-record might not be routed correctly to the UA that placed the emergency call, due to not suppressing any network based services such as call forwarding, or UA based services which can divert the call elsewhere, or if the address-of-record is associated to more than one contact address.

5. Formal Syntax

The following syntax specification uses the augmented Backus-Naur Form (BNF) as described in RFC 5234 [RFC5234].

The "sos" URI parameter is a "uri-parameter", as defined by RFC 3261[RFC3261].

uri-parameter =/ sos-param

sos-param = "sos"

6. IANA Considerations

This specification defines one new SIP URI parameter, as per the registry created by RFC 3969 [RFC3969]

Parameter Name: sos

Predefined Values: none

Reference: [RFCXXXX]

[NOTE TO IANA: Please replace XXXX with the RFC number of this specification.]

7. Security Considerations

As an identifier, the "sos" parameter itself does not raise any particular security issues. The semantics described by the "sos" parameter are meant to be well-known so privacy considerations do not apply to the URI parameter. The main possibility of attack involves use of the "sos" parameter to bypass the normal procedures in order to achieve fraudulent use of services or to bypass security procedures. The usage of this parameter as described in this document is purely for the purpose of the REGISTER request and hence

in presence of user authentication it is ensured that the respective user can be held accountable.

It is RECOMMENDED to log events of misuse of the "sos" URI parameter, for example by including it in a request or response not related to an emergency call.

Emergency registration can result in removing restrictions for roaming and/or barring of services. Misuse of the emergency registered AoR and contact address can be identified within the network and thus requests for unauthorized service will be rejected. Thus, no security considerations related to hijacking of services are foreseen as a result of applying a marking of emergency registrations through the use of a SIP URI parameter.

8. Acknowledgements

The author would like to thank Keith Drage, Milo Orsic, Deb Barclay, John-Luc Bakker, Andrew Allen, Hiroshi Ishikawa, Sean Schneyer, Peter Leis, Georg Mayer, Marvin Bienn, Ricky Kaura, Steve Norreys, Laura Liess, AC Mahendran, Roozbeh Atarius, Ramachandran Subramanian and Sandeep Sharma, Brian Rosen, Hannes Tschofenig, Christer Holmberg and Henning Schulzrinne for the discussions and contributions that led to this work.

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.
- [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.
- [RFC5234] Crocker, D. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, RFC 5234, January 2008.
- [RFC3969] Camarillo, G., "The Internet Assigned Number Authority (IANA) Uniform Resource Identifier (URI) Parameter Registry for the Session Initiation Protocol (SIP)", BCP 99, RFC 3969, December 2004.

9.2. Informative References

[I-D.ietf-ecrit-phonebcpl]

Rosen, B. and J. Polk, "Best Current Practice for Communications Services in support of Emergency Calling", draft-ietf-ecrit-phonebcpl-16 (work in progress), October 2010.

[RFC5031] Schulzrinne, H., "A Uniform Resource Name (URN) for Emergency and Other Well-Known Services", RFC 5031, January 2008.

[3GPP.23.167]

3GPP, "IP Multimedia Subsystem (IMS) emergency sessions", 3GPP TS 23.167 10.1.0, September 2010.

Author's Address

Milan Patel
InterDigital Communications

Email: Milan.Patel@interdigital.com

