

Network Working Group	R. Barnes	
Internet-Draft	BBN Technologies	
Intended status: Informational	B. Aboba	
Expires: April 21, 2011	Microsoft Corporation	
	J. Peterson	
	NeuStar, Inc.	
	H. Tschofenig	
	Nokia Siemens Networks	
	October 18, 2010	

[TOC](#)

Policy Considerations for Emergency Calling using Voice over IP draft-barnes-ecrit-policy-00

Abstract

The provision of emergency calling services (e.g., 911, 112) has been a critical component in the regulation of telecommunications networks. The technical architectures used by modern Voice-over-IP (VoIP) systems mean that if telecommunications regulators wish to extend emergency calling requirements to VoIP, it will likely be necessary to reconsider the ways in which such requirements are applied, both in terms of what specific mandates are imposed and which entities are subject to them. This document dicusses the fundamental technical requirements for emergency services, how these requirements can be met within the framework of VoIP, and how these solutions approaches create possibilities and limitations for regulatory involvement.

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 April 21, 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](#)
- [2. Fundamental Assumptions](#)
- [3. VoIP Architecture](#)
- [4. Obligations](#)
 - [4.1. End Hosts](#)
 - [4.2. ISP](#)
 - [4.3. VSP](#)
 - [4.4. PSAP](#)
- [5. Requirements](#)
 - [5.1. Geolocation](#)
 - [5.2. Call Routing](#)
 - [5.3. PSAP Reachability](#)
 - [5.4. Regulatory Implications](#)
- [6. Outlook](#)
- [7. Security Considerations](#)
- [8. IANA Considerations](#)
- [9. Informative References](#)
- [§ Authors' Addresses](#)

1. Introduction

[TOC](#)

For several decades, emergency calling has been a critical function of telecommunications networks. Starting in the 1960s, capabilities were created in many places around the world to allow any user of a telephone to reach emergency services simply by dialing a short string of digits (e.g., 9-1-1 in the US or 1-1-2 in Europe). Different countries have implemented these functions in their own ways, but basic emergency calling services are available throughout most of the world today. Particularly the introduction of automatic caller location,

essentially for quickly and accurately dispatching first responders, was a painful and long process that is still ongoing in different parts of the world, particularly when considering location information with better accuracy.

At the same time, an ever increasing amount of the voice traffic in the world is being carried via Voice-over-IP services. Just as with other services, the transition to VoIP entails a transition from a circuit-based model of calling to a packet-based one. Instead of a call having a dedicated channel over which signals are transmitted, in VoIP, voice signals are divided into small packets and sent through the Internet (with each packet traveling independently). From the underlying network point of view a VoIP call appears like any other application running over the Internet, rather than a core function of the network. As we will discuss in more detail below, the fact that VoIP is an application makes it significantly more flexible than existing PSTN communications facilities, but by the same token, VoIP systems cannot take advantage of certain fixed properties of the PSTN, making it much more of a challenge to implement emergency services.

So the situation is challenging: On the one hand, PSTN systems around the world are being transitioned over to VoIP systems, but on the other hand, the structure of VoIP systems makes them incompatible with the way emergency services are provided today. In their attempt to enhance voice over IP with emergency services the technical community has designed a system for enabling VoIP emergency calling. Different organizations considered different deployment approaches known as the ECRIT architecture. Even though the required technologies have already been defined the successful deployment of VoIP emergency services will require sharing of responsibilities by multiple players in the Internet ecosystem. There is thus a need to re-think emergency calling regulation to take into account the structural differences between VoIP and traditional voice services and to consider the re-align of responsibilities in this multi-stakeholder eco-system.

In this document, we review some fundamental properties of emergency calling systems and VoIP systems, and discuss how these two concepts come into conflict. We then present approaches for resolving these conflicts, and some thoughts on the role that telecommunications regulation and policy can play in fostering the deployment of VoIP emergency services.

2. Fundamental Assumptions

[TOC](#)

The basic goal of an emergency calling service is to connect the caller with an emergency response center that can help with his emergency situation. (These response centers are traditionally known as Public Safety Answering Points, or PSAPs.) A responder's ability to help is typically limited to a certain geographical region for a specific

service (such as police, fire, ambulance), either because of legal constraints (e.g., jurisdictional boundaries), because of organization structure and work split or simply because of the physical constraints on how fast a responder can travel to the scene of an emergency. These structures and responsibilities change but in a much slower pace than technology. So emergency calling is fundamentally a question of ensuring that a PSAP is reached that is responsible for the geographical area the emergency caller is currently located in order to dispatch first responders.

There are thus three fundamental functional requirements for a successful emergency call (whether over the PSTN, VoIP, or any other technology):

1. An entity routing the call must have access to information about the caller's location.
2. The same call-routing entity must know where to route the call.
3. The same call-routing entity must be able to forward a call to a PSAP.

The challenge in enabling emergency calling using a given communications system is thus to determine how each of these steps is accomplished within that system. In fixed-line PSTN networks, all three requirements were essentially met by virtue of wiring: Customer lines are connected to local switching centers, and switching centers typically cover areas that are served by a single PSAP, so any emergency call that arrives at a switching center can be routed to a dedicated line to the PSAP. (In reality, the situation is somewhat more complex, but we summarize here for simplicity.)

The advent of cellular systems forced a degree of separation among these functions, since the caller's location was not known in advance. In order to provide emergency calling, cellular networks had to deploy specific capabilities to locate their subscribers, and upgrade switches that handle emergency calls so that they can query those capabilities and route calls based on the subsequent location. Even in this case, however, the routing function can be fairly static, because a particular cellular network only covers a specific geographic area.

3. VoIP Architecture

[TOC](#)

The IP-based emergency services access architectures differentiate a few components that have different responsibilities for offering the complete end-to-end functionality. These roles are:

*Internet Service Provider (ISP)

*Voice Service Provider (VSP), or in a more generic form
Application Service Provider (ASP)

*Emergency Service Provider (that operates a PSAP) and vendors of
equipment for those.

*End Device

Note that multiple roles be provided by a single organisation.
[Editor's Note: Should we provide a short description of each of the
roles?

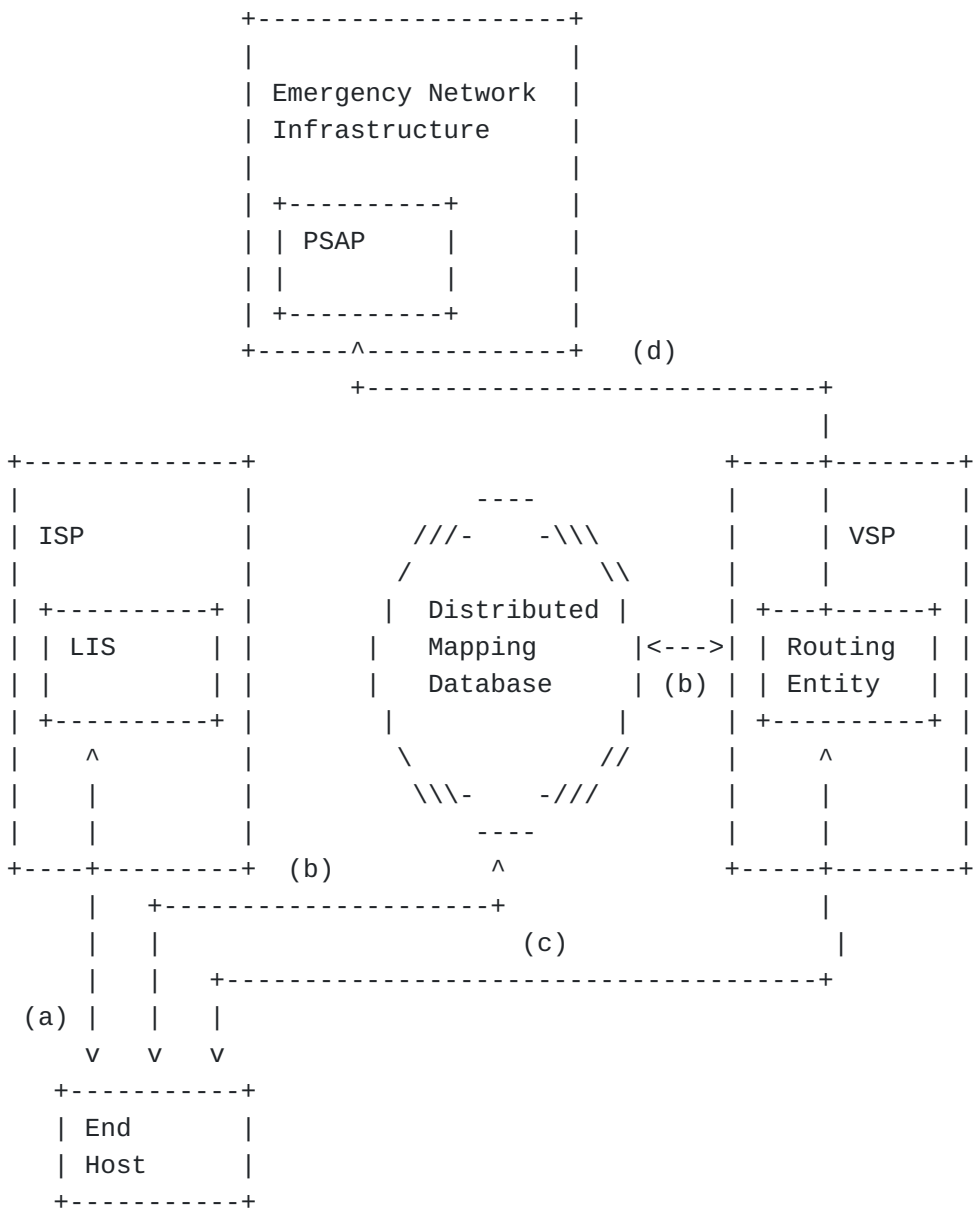


Figure 1: Stakeholders in the Emergency Services Eco-System

The references to interfaces (a), (b), (c), and (d) will be used in [Section 4 \(Obligations\)](#).

The emergency call interaction on a high level takes place as follows: the user enters an emergency services number (or potentially an emergency dial string). The end device recognises the entered sequence of digits as an emergency call attempt and determines whether location information is available locally (as part of the GPS module, for example). If no location information is available locally then various protocol extensions have been defined that allow the end host to obtain location information from a Location Information Server in the ISPs network. Then, the call setup procedure using SIP is started towards the users VSP/ASP. The VSP/ASP needs to make a route decision to determine where the call needs to be forwarded to. Often, this decision is based on a combination of the callers location information as well as other policy aspects (such as time-of-day, workload of a specific PSAP).

When considering IP-access to emergency services, one should consider the following categories and the challenges they induce:

Fixed Access: From a user point of view this scenario is characterised by a stationary usage of a computer, such as a desktop PC at someones home. Typically, these devices are often not equipped with a GPS receiver or, because of the indoor usage, these do not work very well or not at all. Information about the callers location can therefore come from manual configuration (which may be useful in cases where the location of the device indeed rarely changes or the user is careful in keeping the reported location accurate) or from the ISPs network since the attachment point is known to the operators network infrastructure.

Nomadic Access: Nomadic access is characterised in movement patterns that correspond to regular laptop usage where users switch location from time to time (e.g. go to work, use their laptop at the university or in a coffee shop, and at home). In this scenario it is not realistic to assume that users update their location manually due to the frequent change. The usage of GPS may be possible even though network operator presented location would be preferred since users will typically use their device indoors where GPS does not work well.

Mobile Access: This scenario is an enhancement of the previously presented nomadic access with the assumption that users roam while having their communication ongoing. In this scenario, devices, such as smart phones, are often equipped with GPS

receivers and make the location determination process more accurate. Manually entered location is in this scenario not possible.

Note: This is a simplified view on networking from the users point of view. As network architectures become more sophisticated the boundaries between these scenarios get more fuzzy. As an example, one may consider a traveller using a laptop with wireless LAN (as he uses at home) in a train connected. The network infrastructure in the train is connected via a cellular infrastructure to the Internet. This example blurs nomadic access and mobile access. Consider another example where a teenager uses his high-performance laptop for gaming and sometimes uses it at home as a replacement for the desktop PC and sometimes at some LAN parties to compete with other games. From software program point of view these cases are very hard to differentiate since in all cases the end device is uses WLAN technology to communicate with the network. Hence, it is left to the user to 'switch' between usage environments, which introduce problems when software developers make too restrictive assumptions about the environment where their devices will later be used or about the users awareness of the necessary configuration changes. Ideally, users should not need to configure their devices to prepare for the case of an emergency call. For the unlikely case of an emergency users should not be required to ever configure their device - zero configuration is the goal.

From user-experience point of view the overall process of establishing an emergency call begins someone dialling an emergency dial string. The exact sequence of digits depends on the infrastructure the device is connected to. While 1-1-2 became the emergency services number for Europe and 9-1-1 the emergency number for the US many countries still provide additional emergency numbers mostly for historical reasons. Furthermore, many large enterprises, university, and hotels prefix the emergency numbers with additional digits, such as 0-112.

An important part of emergency handling is in the logic of delivering emergency calls to the appropriate PSAP. With devices that may be used with different VSPs/ASPs and in cases where there is no relationship between the ISP and the VSP/ASP the automatic routing decision becomes more complex. Consider the following example where user Bob travels to from Sweden to Spain and wants to use their device to trigger an emergency real-time text interaction. Since Bob is using a real-time text provider in Sweden the messages are routed to the Swedish operator. Based on the provided location the Swedish operator notices that a PSAP in Spain has to be involved and, for example, initiates a conference bridge with Bob, a relay provider in Sweden serving as a language translator, and the PSAP operator in Spain. In order for the Swedish ASP/VSP or the Swedish PSAP to involve the appropriate Spanish PSAP their contact information needs to be available, and information about their responsibility (i.e. for which region they are responsible and which emergency service function they provide) needs to be published.

A protocol, called Location to Service Translation (LoST), has been defined to map a location together with a symbolic representation of the emergency service requested to a specific PSAP contact address (PSAP URI). Note that the returned PSAP URI does not necessarily need to be the contact information of the final PSAP but rather it will route the call closer to one.

Location information is needed by emergency services for three reasons: routing the call to the right PSAP, dispatching first responders (e.g. policemen), and determining the right emergency service dial strings. It is clear that the location has to be automatic for the first and third application, but experience has shown that automated, highly-accurate location information is vital to dispatching as well, rather than relying on the caller to report his or her location to the call taker. This increases accuracy and avoids dispatch delays when the caller is unable to provide location information due to language barriers, lack of familiarity with his or her surroundings, stress, physical or mental impairment. Location information for emergency purposes comes in two representations: geo(detic), i.e., longitude and latitude, and civic, i.e., street addresses similar to postal addresses. Particularly for indoor location, vertical information (floors) is very useful. Civic locations are most useful for fixed Internet access, including wireless hot spots, and are often preferable for specifying indoor locations, while geodetic location is frequently used for cell phones. However, with the advent of femto, and pico cells, civic location is both possible and probably preferable since accurate geodetic information can be very hard to acquire indoors. Before location can be put into a protocol for delivery and utilised it first needs to be determined. Location information can be entered by a user ("manual configuration"), measured by the end host, can be delivered to the end system by some protocol or measured by a third party. The actual process of location determination is largely outside the scope of the REACH-112 project but manual configuration, GPS usage, and the usage of location servers is relevant. Many VoIP deployments allow their users to manually enter location information for later usage with emergency services. Typically, the users enter their home address into a web-based form and this data would then be used for emergency service call routing and also delivered to PSAP operators. This mechanism is primarily suitable for users utilising fixed network deployments (such as Cable and DSL networks) rather than cellular networks where the current users location changes continuously. For nomadic users this approach already becomes very cumbersome for end users. While this mechanism clearly has limitations it is still a useful approach in absence of other techniques. For devices like laptops and in particular mobile phones the usage of GPS is a promising technique that is able to provide highly accuracy. While it has also has limitations (for example when used indoor) and may need a fair amount of time to provide the initial location fix it is a promising technique. The requirements for location accuracy differ between routing and dispatch. For call routing, city or even county-level

accuracy is often sufficient, depending on how large the PSAP service areas are, while first responders benefit greatly when they can pinpoint the caller to a particular building or, better yet, apartment or office for indoor locations, and an outdoor area of at most a few hundred meters outdoors. This avoids having to search multiple buildings, e.g., for medical emergencies.

For a reliable emergency services infrastructure utilizing the location information provided by ISPs is important, largely for dispatch purposes since the accuracy is sufficiently high to allow first responders to locate the person in need for help.

4. Obligations

[TOC](#)

In this section we discuss how the responsibilities for deployment need to be shared based on the architectural illustration from [Figure 1 \(Stakeholders in the Emergency Services Eco-System\)](#). Note that this narration focuses on the final stage of deployment and do not discuss the transition architecture, in which some implementation responsibilities can be re-arranged, with an impact on the overall functionality offered by the emergency services architecture. A few variations were introduced to handle the transition from the current system to a fully developed ECRIT architecture.

4.1. End Hosts

[TOC](#)

An end host, through its VoIP application, has three main responsibilities: it has to attempt to obtain its own location, determine the URI of the appropriate PSAP for that location, and recognize when the user places an emergency call by examining the dial string. The end host operating system may assist in determining the device location. The protocol interaction for location configuration is indicated as interface (a) in [Figure 1 \(Stakeholders in the Emergency Services Eco-System\)](#) and a number of location configuration protocols have been developed to provide this capability.

A VoIP application needs to have the ability to detect the emergency call, to obtain (potentially only locally available) location information, and to make it available to the VSP during call setup.

[TOC](#)

4.2. ISP

The ISP has to make location information available to the end point via one or more of the location configuration protocols. In order to route an emergency call correctly to a PSAP, an ISP may initially disclose the approximate location for routing to the end point and more precise location information later, when emergency personnel is dispatched by the PSAP operator. The functionality required by the IETF emergency services architecture is restricted to the disclosure of a relatively small amount of location information, as discussed in

[\[I-D.ietf-ecrit-location-hiding-req\]](#) (Schulzrinne, H., Liess, L., Tschofenig, H., Stark, B., and A. Kuett, "Location Hiding: Problem Statement and Requirements," February 2010.) and in [\[I-D.ietf-ecrit-rough-loc\]](#) (Barnes, R. and M. Lepinski, "Using Imprecise Location for Emergency Context Resolution," August 2010.).

The ISP may also operate a (caching) LoST server to improve the robustness and the reliability of the architecture. This lowers the roundtrip time for contacting a LoST server and the caches are most likely to hold the mappings of the area where the emergency caller is currently located.

In case ISPs allow Internet traffic to traverse their network the signaling and media protocols used for emergency calls function without problems. Today, there are no legal requirements to offer prioritization of emergency calls over IP-based networks. While the standardization community has developed a range of Quality of Service signaling protocols their (widespread) deployment still remains to happen.

4.3. VSP

[TOC](#)

SIP does not mandate that call setup requests need to traverse SIP proxies, i.e., SIP messages can be sent directly to the user agent. Thus, even for emergency services it is possible to use SIP without the involvement of a VSP. However, in terms of deployment, it is highly likely that a VSP will be used. If a caller uses a VSP, this VSP often forces all calls, emergency or not, to traverse an outbound proxy or session border controller (SBC) operated by the VSP. If some end devices are unable to perform a LoST lookup, VSP can provide the necessary functions as a back-up solution. If the VSP uses a signaling or media protocol that is not supported by the PSAP, it needs to translate the signaling or media flows.

VSPs can assist the PSAP by providing identity assurance for emergency calls and thus helping to prosecute prank callers. However, the link between the subscriber information and the real-world person making the call is weak. In many cases, VSPs have, at best, only the credit card

data for their customers and some of these customers may use gift cards or other anonymous means of payment.

4.4. PSAP

[TOC](#)

When emergency calling has been fully converted to Internet protocols, PSAPs must accept calls from any VSP, as shown in interface (d) of [Figure 1 \(Stakeholders in the Emergency Services Eco-System\)](#). Since calls may come from all sources, PSAPs must develop mechanisms to reduce the number of malicious calls, particularly calls containing intentionally false location information. Assuring the reliability of location information remains challenging, particularly as more and more devices are equipped Global Navigation Satellite Systems (GNSS) receivers, including GPS and Galileo, allowing them to determine their own location. However, it may be possible in some cases to the veracity of the location information provided by an end-point by comparing it against infrastructure-provided location information, e.g., LIS determined location.

5. Requirements

[TOC](#)

[[TBD: We need to add a bit of wording here in this section to your notes. I think that this section should also talk about the distribution of responsibilities among the stake holders and motivate why they want to do this. Then, we have no requirement yet for multi-media emergency calls, which would be good to have.]]

5.1. Geolocation

[TOC](#)

[[-- Location -- Basic requirement: Get location information to a call-routing entity -- Endpoint or VoIP service -- ... but probably the endpoint (security reasons) -- State of the art today -- Carrier location functions; inaccessible to end devices -- "World in a DB" functions; low-fidelity -- GPS; frequently fails -- Commercial direction is not moving quickly away from these stovepipes -- GEOPRIV model provides bridges, but not much uptake -- Also provides a location interface for PSAPs, e.g. as NENA has defined]]

[TOC](#)

5.2. Call Routing

```
[[ -- Call routing -- Need a public, open DB with a standard query
interface -- Source of the data needs to be local emergency authorities
=> Need for local authorities to coordinate to create LoST DB(s) ]]
```

5.3. PSAP Reachability

[TOC](#)

```
[[ -- PSAP reachability -- Basic requirement: Internet connectivity,
SIP reachability -- Gateways as a transition step -- Security questions;
ref to ESInet concepts ]]
```

5.4. Regulatory Implications

[TOC](#)

```
[[ -- Roles government can play -- Who can be the targets of
regulation? -- Localized (yes): PSAPs, ISPs, vertically-integrated VoIP
-- Non-localized (no): Application-only VoIP, application-only location
providers -- Enabling location: -- Several networks that provide E911
are also ISPs -- Require them to open up ALI or equivalents to
endpoint, PSAP access -- How to enable NG911 without giving away the
store: rough-loc -- Encourage ISPs in general to enable location
services for customers -- Call routing: -- Coordinate the development
of a national LoST infrastructure -- Formalize LoST as a national
standard call-routing interface -- Encourage ISPs to support LoST
discovery -- PSAP reachability: -- Support PSAP upgrades and gateways
-- Encourage VoIP vendors to integrate emergency calling into products
-- E.g., support open-source location, LoST components ]]
```

6. Outlook

[TOC](#)

In most countries, national and sometimes regional telecommunications regulators have a strong influence on how emergency services are provided; such as who pays for them and what obligations the various parties have. Regulation is, however, still at an early stage: in most countries current requirements only demand manual update of location information by the VoIP user. The ability to obtain location information automatically is, however, crucial for reliable emergency service operation, and required for nomadic and mobile devices. (Nomadic devices remain in one place during a communication session, but are moved frequently from place to place. Laptops with WiFi interfaces are currently the most common nomadic device.)

Regulators have traditionally focused on the national or, at most, the European level, and the international nature of the Internet poses new challenges. For example, mobile devices are now routinely used beyond their country of purchase and, unlike traditional cellular phones, need to support emergency calling functionality. It appears likely that different countries will deploy IP-based emergency services over different time horizons, so that a traveler may be surprised to find that she cannot call for emergency assistance outside their home country.

The separation between Internet access and application providers on the Internet is one of the most important differences to existing circuit switched telephony networks. A side effect of this separation is the increased speed of innovation at the application layer and the number of new communication mechanisms is steadily increasing. Many emergency service organizations have recognized this trend and advocated for the use of new communication mechanisms, including video, real-time text, and instant messaging, to offer improved emergency calling support for citizens. Again, this requires regulators to re-think the distribution of responsibilities, funding and liability.

Many communication systems in use today lack accountability, i.e., it is difficult or impossible to trace malicious activities back to the persons who caused it. This is not a completely new problem, as pay phones and prepaid cell phones have long offered mischief makers the opportunity to place hoax calls, but the weak user registration procedures, the lack of deployed end-to-end identity mechanisms, and the ease of providing fake location information increases the attack surface at PSAPs. Attackers also got more sophisticated over time and Botnets to generate a large volume of automated emergency calls to exhaust PSAP resources, including call takers and first responders, is not science fiction.

7. Security Considerations

[TOC](#)

[TODO]

8. IANA Considerations

[TOC](#)

This document has no actions for the IANA.

9. Informative References

[TOC](#)

[RFC2119]

	Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," BCP 14, RFC 2119, March 1997 (TXT , HTML , XML).
[I-D.ietf-ecrit-location-hiding-req]	Schulzrinne, H., Liess, L., Tschofenig, H., Stark, B., and A. Kuett, " Location Hiding: Problem Statement and Requirements ," draft-ietf-ecrit-location-hiding-req-04 (work in progress), February 2010 (TXT).
[I-D.ietf-ecrit-rough-loc]	Barnes, R. and M. Lepinski, " Using Imprecise Location for Emergency Context Resolution ," draft-ietf-ecrit-rough-loc-03 (work in progress), August 2010 (TXT).

Authors' Addresses

[TOC](#)

	Richard Barnes
	BBN Technologies
	9861 Broken Land Parkway
	Columbia, MD 21046
	USA
Phone:	+1 410 290 6169
Email:	rbarnes@bbn.com
	Bernard Aboba
	Microsoft Corporation
	One Microsoft Way
	Redmond, WA 98052
	US
Email:	bernarda@microsoft.com
	Jon Peterson
	NeuStar, Inc.
	1800 Sutter St Suite 570
	Concord, CA 94520
	US
Email:	jon.peterson@neustar.biz
	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