

Emergency Call Services for SIP-based Internet Telephony

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of Section 10 of RFC2026.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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.”

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>

To view the list Internet-Draft Shadow Directories, see <http://www.ietf.org/shadow.html>.

Copyright Notice

Copyright (c) The Internet Society (2001). All Rights Reserved.

Abstract

If Internet Telephony is to offer a full replacement for traditional telephone services, it needs to provide emergency call services. In the United States, emergency calls are known as 911 services, based on the number dialed. This note describes some options for providing enhanced emergency service, i.e., emergency calls that allow emergency response centers to determine the address where the caller is located. This is made more difficult by the temporary nature of IP addresses, the large number of ISPs and their lack of legal responsibility for emergency services and the ability of many Internet terminals to be connected to the Internet at different locations. This note explores some of the requirements and design choices.

1 Introduction

Providing emergency communications, exemplified by the Enhanced 911 (E911) service available in the United States, encompasses a number of steps and requirements [1, 2]:

- The caller needing emergency assistance initiates a call to the appropriate Public Safety Answering Point (PSAP), typically via nationwide uniform number (911 in the United States, 000 in Australia, 112 in the European Community, with special numbers in some countries for police, ambulance or fire). This functionality, without the additional features below, is called “basic 911”, while the features below constitute “enhanced 911”.
- The call is routed by the local switch to a PSAP, based on the calling number.
- At a PSAP, emergency operators determine the nature of the emergency and contact the appropriate agency. A single PSAP is usually responsible for an area covering several independent police and fire departments. This area is called an *Emergency Service Zone* (ESZ). It is identified by an Emergency Service Number (ESN), a number consisting of three to five digits that represents a unique

combination of emergency service agencies such as law enforcement, fire and emergency medical service. The Master Street Address Guide (MSAG) is a database of street names and house number ranges within a community defining emergency service zones and the associated Emergency Service Numbers (ESNs).

Note that the appropriate PSAP may not be the one geographically closest to the caller, due to political, geographic and administrative boundaries. Also, service areas of different emergency service agencies may not be the same. For example, a whole city may be served by the same law enforcement agency, but may have a number of different emergency medical service units serving different boroughs or other divisions.

Calls to the first PSAP, the so-called primary PSAP, may be relayed to a secondary PSAP.

- Wireless phones in the United States must be able to reach 911 even if the phone is not a registered subscriber.
- Since it is possible that the caller is too confused, frightened or young to properly identify the location of the emergency, modern 911 systems [not sure about Europe] require all local telephone operators to maintain a database, called the ALI (Automatic Location Identification) database, containing current subscriber addresses, telephone numbers and supplementary emergency information for landline telephones.

For wireless service, operators must be able to identify the location of the user within 50 m (for 67% of calls).

The PSAP obtains the caller's phone number, name and address. This database lookup is called Automatic Location Identification (ALI). The caller's phone number is delivered to the PSAP regardless of any restrictions on caller ID delivery (CLID).

- Certain services, such as call-waiting or three-party calls, are restricted during (basic) 911 calls. The caller cannot hang up the call and place another call.

Internet telephony offers the opportunity not only to supply current emergency services, but also to add functionality. For example, multimedia communications can be used to provide video images or biometrics. A two-way video call allow an emergency medical technician to instruct a person at the scene of the emergency in emergency procedures and provide corrective feedback to those performing first aid.

The call setup can provide additional medical background, without having to store the information in a central database. Due to the faster call setup times and the ease of redirection, calls can be forwarded and transferred to emergency personnel en route to the caller or a nearby PSAP, in case the primary PSAP serving the caller's location is overloaded. Unlike current PSAPs, which require dedicated equipment, an Internet-based PSAP could be set up anywhere Internet access is available, providing easy relocation should the PSAP be unavailable during natural disasters. The ability to indicate language capabilities of the caller can help route the call to an operator, without the additional delay of having a general operator try to ascertain the language of the caller.

Multiple media, including audio, video and text chat, also offer improved capabilities to people with disabilities, making it more likely, for example, that a deaf person visiting somewhere can find a multimedia-capable device from which at least text (instant message) communication is possible.

Also, while 911 service is available to about 85% of the U.S. population, other countries have very different numbering schemes (e.g., 110 and 112 in Germany and other European countries). Also, many

corporate and hotel PBX systems do not support 911. Internet telephony may be able to make the service universally available.

However, many of the assumptions underlying the current 911 service do not hold for Internet telephony.

It appears likely that users of Internet telephony services, wired or wireless, expect similar behavior as in existing networks. Indeed, even if the call is terminated at the PSAP via the PSTN, difficulties can arise if the caller is using Internet telephony, as the fire department may well be summoned to the telephone company building housing the VoIP-PSTN gateway.

Simply replicating the existing phone-only emergency call system seems ill-advised. In the future, users are not likely to perceive Internet telephony as a separate service, but simply as another communications service along with email, chat and maybe even distributed games. Thus, it is desirable to allow a user to summon emergency assistance from any communication application, regardless of the underlying protocols. A general infrastructure also makes it easier for users with disabilities to summon help [3].

Emergency numbers in an Internet context require components similar to those listed above. We first describe alternatives for reaching an emergency operator in Section 2 and then discuss how the emergency operator can locate the caller in Section 3.

2 Reaching Emergency Help

The architecture described here envisions Internet PSAPs (iPSAPs) that operate in a fashion roughly similar to today's telephone-only PSAPs. These iPSAPs triage incoming communication and then contact the appropriate public safety agency. Callers reach the iPSAP in two stages, by indicating a common, location-independent address (Section 2.1), with a network entity then routing the call to the appropriate iPSAP (Section 2.2).

2.1 User-Visible Emergency Address

Emergency addresses can be defined at both the network and application layer. At the network layer, a designated scoped multicast address [4] could reach a local node with knowledge of the iPSAP's address (see below) that then forwards the request to the appropriate authority. Multicast precludes the use of TCP and application-layer forwarding obliterates the original IP identity.

As an alternative, an IPv6 anycast address can be designated for emergency communications, but a DHCP option for configuring end systems would need to be defined.

At the application layer, any protocol designed for user-to-user communication, including email, SIP, IRC and chat, should have a designated domain-specific emergency response address, similar to the current `hostmaster@domain`, `postmaster@domain`, `webmaster@domain` and similar functional addresses [5]. There is no need to restrict this to one address, however, so that a message addressed to, say, 911, sos, or 110 reach the correct destination, regardless of the host part of the URL. It is highly desirable to standardize this address Internet-wide rather than leaving this to national authorities, to avoid having to customize software applications for each national jurisdiction.

It is difficult to provision a proxy for all possible emergency numbers, as there are at least 60 different ones in use throughout the world (000 01 02 03 04 05 061 07 080 091 100 10022 101 10111 10177 102 107 110 111 112 113 117 118 119 120 122 123 131 132 133 140 144 15 150 155 158 160 166 17 171 175 18 19 190 191 192 193 196 199 911 92 93 94 990 993 994 995 997 998 999).

In addition, callers may also use the `tel` URL [6] in SIP requests with the local emergency number. However, this does not simplify the problem since the PSTN gateway still has to determine whether the

caller is within the same PSAP and thus can simply dial 911. As discussed in Section 3, PSTN gateways are likely to yield incorrect location information.

2.2 Finding the Appropriate iPSAP

Either the end system or a network entity, i.e., a SIP proxy server, can determine the appropriate iPSAP. Here, end systems include PSTN-to-VoIP gateways. As noted earlier, the appropriate iPSAP may not be the one geographically closest to the user. Except for corporate campuses with their own emergency response infrastructure, boundaries of areas served by an iPSAP (the ESZ) are very unlikely to bear any resemblance to network topology or network administrative regions.

In either case, there are a number of alternatives, including DNS, SLP or a central directory server. We discuss these options below.

For SIP-based systems, note that end system and proxy-based routing can be combined in a network, similar to how request routing works for regular SIP requests. If an end system or proxy does not know how to find the appropriate iPSAP, it routes it to another proxy that may know. If the request does not contain location information, the proxy inserts it, as described in Section 3, possibly indicating how certain it is about the accuracy of the location information.

PSTN-to-VoIP gateways need to take the location of the caller into account and thus will likely have to access existing ALI databases.

In an emergency response system, reliability is particularly important. Thus, it is desirable to minimize the number of servers and resolution steps necessary to reach the PSAP. As much as possible, the system has to work even if wide-area communications is unavailable and should minimize the delay even if the network is congested.

2.2.1 Finding the iPSAP via multicast or anycast

As discussed in Section 2.1, a multicast or anycast address can be defined which allows any system to locate the nearest iPSAP. Unfortunately, network scoping and topology are not likely to conform to the civic boundaries served by an iPSAP, but they may well be congruent within a corporate campus network.

The SIP INVITE or MESSAGE requests can be multicast. Other protocols may have to add multicast capabilities or need to use a separate multicast protocol to locate the server, which is then contacted via unicast. In that case, SLP 2.2.3 may well be appropriate.

2.2.2 Finding the iPSAP via DNS

User locations are mapped into a DNS hierarchy, with each entry referring to the appropriate iPSAP. The DNS hierarchy can be either geographically based or be based on civil designations such as postal codes or town names, or both. A DNS hierarchy based on geo location (longitude and latitude) is not likely to work well since the appropriate resolution of the entries is difficult to determine in general. For example, geographically close locations may be part of different jurisdictions or separated by a river or mountain ridge. Also, PSAPs cover areas of very different size, depending on population density. Town names or postal codes ("civil location") can be readily encoded into DNS, as in `07605.us.emergency.arpa`, `5.0.6.7.0.us.emergency.arpa` or `leonia.nj.us.emergency.arpa`. It remains to be determined whether there is sufficient overlap between PSAPs and postal zones to make this approach viable. This approach works only if the proxy or end system has accurate civil location information, as discussed below. Also, all valid postal codes or names need to be populated.

The DNS entry can point to an SRV or NAPTR record, making it easy for different protocols (SIP, H.323, chat, IRC, etc.) to reuse the same directory infrastructure and allowing services for different communications media to be handled by different iPSAPs. (For example, if a particular communications mechanism is not widely used or requires special equipment, one location may handle it for several iPSAPs.)

A DNS-based approach has the advantage that it scales well and entries can easily be delegated. However, end systems or proxies have to be configured with their civil location. Thus, either the end system or proxy has to be able to determine the correct location of the end system. This is likely to be easy for DSL, cable modem and leased-line installation, but more difficult for dial and wireless users, as discussed in Section 3.1.

A direct mapping of emergency phone numbers via the approach pursued in the enum working group [7] does not seem to be possible, since it would map every such query to the same iPSAP location. An indirect mapping to a location service, however, is necessary in any event.

2.2.3 Finding the iPSAP via SLP

Finding an iPSAP can be considered a service location problem. However, the SLP zone mechanism is not a good fit for service location, since it usually corresponds to a local network administrative zone. However, in most cases, a single SLP zone falls within a single PSAP jurisdiction, so that the mapping is likely to be valid. On the other hand, SLP does not really offer anything in that case that a simple query protocol could not do just as well. Unless entities can discover the local service directory via multicast, adding SLP would require yet another configuration mechanism.

The SLP entry would likely contain the DNS entry discussed above or, directly, the address of the iPSAP. (The latter avoids reliance on DNS, but is less flexible.)

A resolution mechanism based on SLP is not likely to be viable until almost all local networks support the service. It adds little functionality if the proxy locates the appropriate iPSAP, as it can perform the same look-up logic as the SLP server, saving an additional step.

Also, if SLP were to be deployed primarily for emergency call use, it would be less likely to be tested regularly and thus be less reliable. Reliability extensions to SLP [8] may be advisable to remove single points of failure.

Thus, the use of SLP is not likely to be viable alternative.

2.2.4 Finding the iPSAP via DHCP

As long as the address of the iPSAP is constant for all devices served by a DHCP [9] server, it is easy to configure it into DHCP. Unfortunately, it appears to be difficult for applications on standard operating systems to access DHCP configuration information; embedded devices have no such difficulty.

2.2.5 Finding the iPSAP via LDAP

An LDAP server could, given information about a caller's civil or geographical location, return the appropriate iPSAP address. Unlike the mapping from user identifier to location (see Section 3.2.3), this information is less privacy-sensitive, but updates to the data base have to be carefully controlled.

This approach requires configuring end systems or proxies with the address of the appropriate LDAP server, thus making the approach not viable for robust end system use.

3 Locating Emergency Callers

Probably the most difficult aspect of an Internet-based emergency call system is determining the civil or geographical user location. The problem has two aspects, namely how the end system or network determines the terminal's location (Section 3.1) and how the iPSAP obtains this information (Section 3.2).

3.1 Determining User Location

First, devices or the network have to determine their geo or civil location. This may well turn out to be easier for cellular devices, where all the techniques suggested for second-generation wireless systems can be applied. These include GPS or radio-based location [10]. For wireless LANs, the RADAR project [11] has shown that IEEE 802.11-based devices can be located to room accuracy within a building if radio conditions are modeled or surveyed. However, GPS generally does not work within buildings, so that even assuming desktop computers or Internet phones could be economically equipped with GPS receivers, additional methods have to be found. (GPS systems presumably report a location just before entering a building, which may well be sufficient. This is less satisfactory when entering a tunnel in the Alps or under the British Channel since the old reading may be off by a whole country. Also, GPS location or triangulation is of limited usefulness within buildings due to the large uncertainty in altitude.)

If the terminal obtains location information, it may need to update an external server, such as DNS or a specialized location service [12, 13, 14].

For devices connected to local area networks, such as switched Ethernet, it is suggested that each port periodically sends a message to the attached devices indicating the geo or civil location. The switch is configured with this information via SNMP. Since jacks are reasonably stationary in most hard-wired installations, this requires only a one-time effort, but would require a substantial upgrade effort for existing installations. However, due to the use of patch panels and hubs, there is a chance that small-scale inaccuracies creep in, but probably at the scale of room numbers rather than buildings. (In addition, this information could also be useful for asset management.)

Even for wireless LANs with base station connected to Ethernet switches, this would provide location accuracy of about 100 to 300' and generally within a single floor of a building.

Location information can also be configured directly into end systems, preferably in a manner that allows all applications access to this information. This would allow SIP, for example, to simply add a header to emergency call setup requests or mid-call INFO requests containing the geographic location. However, the likelihood that users will correctly enter or update this information is low, even if, for example, operating systems would require entry of such information. Even with modestly mobile systems, such as PCs that are moved from office to home or moved between homes, the information is likely to be out of date in a large fraction of cases.

This approach has the advantage that the end system controls the delivery of information. On the other hand, applications then have to be aware when the user contacts the emergency address.

3.2 Conveying User Location to the iPSAP

The user location can be conveyed to the iPSAP in a number of different ways, as illustrated in Fig. 1. Generally, the location information can be delivered directly by the end system into application-layer protocols, inserted by SIP proxies or be obtained by the iPSAP from either the device or a database. The first two methods depend on the application-layer protocol, while the latter can be used for any protocol employed to contact the iPSAP. We evaluate some of the alternatives below and summarize the options in Table 1.

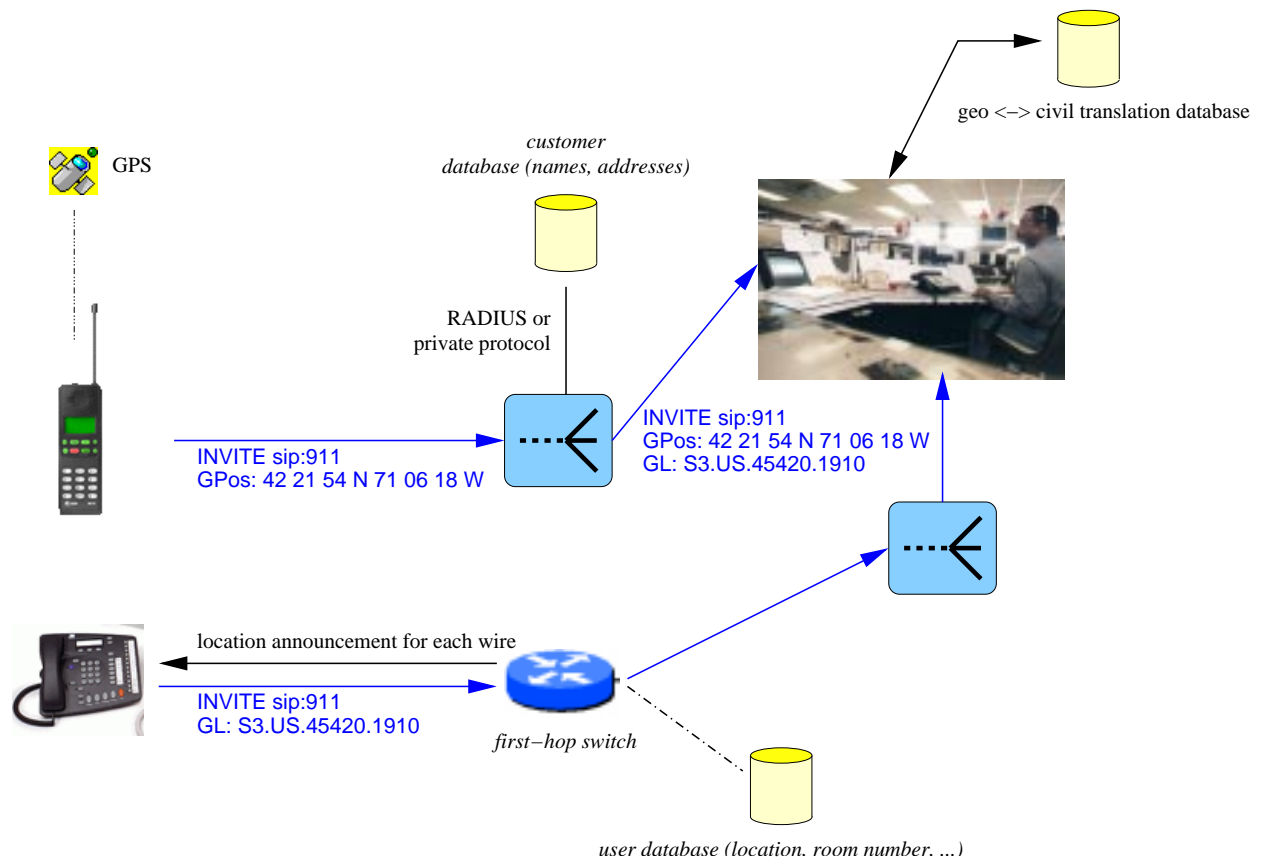


Figure 1: Architecture for obtaining and conveying location information to a PSAP

Method	Sec.	suitable for	modifications required?
DNS	3.2.1	any	no
SLoP	3.2.2	any	no
Identifier	3.2.3	any	yes
ISP	3.2.4	any	true IP address
SIP URL	3.2.5	SIP	
SIP proxy	3.2.6	SIP	

Table 1: Comparison of options for conveying system location to the iPSAP

3.2.1 DNS

Network providers or terminals can store location information in DNS, via dynamic DNS updates. DNS-based schemes have been defined that yield either a civil [15] or a geographic location, based on a caller's domain name. For example, [15] defines DNS resource records of the form below:

```
donuts A 192.188.192.1
      GL S3.US.45420.1910 "1425 Arbor Avenue, Dayton OH"
```

From the record, the country (US), 9-digit zip code (45420-1910) and the street address can be determined.

Alternatively, DNS can also contain geographic information, encoded either as strings [16], e.g.,

```
<owner> <ttd> <class> GPOS <longitude> <latitude> <altitude>
```

```
GPOS -32.6882 116.8652 10.0
```

or as binary (RFC 1876 [17]). RFC 1876 [17] also offers an indication of the precision of the location, as well as the circumference of the object. An example RFC 1876 record might contain

```
cambridge-net.kei.com. LOC 42 21 54 N 71 06 18 W -24m 30m
```

DNS mechanisms work only if the mapping between device and location is correct. If given an IP address in the SIP request, e.g., in the SDP "c=" line or the SIP Contact header, it does an reverse address lookup first to obtain a domain name. As long as both the reverse mapping and location part of the DNS entry is updated dynamically by the network provider or end system, this also works when IP addresses are assigned dynamically. It gets significantly more complicated if NATs or NAPTs are used since several different devices, in different locations, may share the same IP address. If mobile IP is used, the visited network has to do the DNS updates.

DNS information, including location information, is available to anybody, raising serious privacy concerns. DNS is fundamentally not designed to limit the distribution of information, although access control additions to DNS may be feasible.

3.2.2 SLoP (Spatial Location Protocol)

A specialized protocol for obtaining location information is being discussed within the IETF [13]. It relies on a globally unique identifier for network objects. This identifier in turn is used as a key into one of several databases containing location information.

3.2.3 End System Identifier

An end system identifier such as a MAC address or the identifier proposed by the SLoP effort [18] is mapped to a location in one or more global databases maintained by civil authorities. The database may be accessed, for example, through LDAP [19] or whois [20, 21]. This identifier would then have to be conveyed in IP or application-layer messages such as SIP INVITE or INFO. However, such an identifier raises privacy issues, even though access to the location database can be more readily secured. For non-stationary devices, the database still needs to be updated by the end system or system administrator.

3.2.4 Location Information Provided by ISP

Since network access providers generally identify their customers via PPP authentication, or by setting up a key for cable modems or via the physical line for a digital subscriber line, they have the necessary location information as part of their normal billing and maintenance records. Thus, they appear to be in the best position to supply this information. One model has the iPSAP use reverse address mapping on the IP address of the request (or some other identifier at the application layer) to the organization that was delegated this address or identifier. (This information can be obtained via whois from whois.arin.net,

whois.apnic.net or whois.ripe.net.) This organization then maintains a server that maps addresses to customer locations, using any remote query protocol. This clearly imposes additional requirements on the coupling between DHCP and customer databases. If a permanent customer identifier, scoped by AS, is conveyed by application-layer protocols making emergency calls, the mapping is static, but additional changes to DHCP and application support is needed.

The above outline of a solution glosses over a number of potential problems. For dial-in users, the calling number is needed; it is available to the modem pool unless the caller suppresses caller id and so it could be propagated along with the billing/"home" location.

NATs would have to update the organization's customer-to-IP-address database on a connection-by-connection basis. However, NATs can be ignored in many residential and small-business access networks, as each household or business is assigned one IP address, shared by a number of devices.

Depending on whether reverse tunneling is used or not, mobile IP may also obscure the "physical" IP address. For SIP, the network address of the SIP request is likely to identify the outbound proxy or some other proxy in the call path, not the IP address of the caller. Thus, the iPSAP should use the address in the SDP message instead.

3.2.5 Location Derived from SIP URL

This approach is specific to SIP-based emergency calls. The iPSAP could also query the domain indicated in the SIP URL in the request's From header. However, this is not likely to work since the SIP URL may be provided by a service provider that has no direct transport relationship with the user and, in all likelihood, does not have correct address information since it may not have a billing relationship to the user. Even assuming that this service provider has a correct address record for the customer, this approach fails for users that move. For example, a user logging in from a hotel room would still be located as being at home.

3.2.6 SIP Proxy Inserts Location Information

The outbound SIP proxy can insert the information obtained from the local authentication database into SIP requests, assuming again that all (911) SIP requests use an outbound proxy associated with a transport-based ISP that has accurate customer records.

For example:

```
GPOs: 42 21 54 N 71 06 18 W -24m 30m
GL: S3.US.45420.1910 "1425 Arbor Avenue, Dayton OH"
```

This approach has the advantage that it looks the same to the iPSAP whether information is derived via GPS or by the proxy.

3.2.7 Summary

It is likely that a combination of mechanisms will need to be deployed. For wireless end systems, either the end system or the wireless network operator can readily offer location information, either reflected back through the end system into application-layer protocol or via a system identifier and a database accessible to iPSAP authorities.

For landline devices attached to LANs, it appears easiest to enhance Ethernet switches to provide location information to their ports, as all other mechanisms are likely to be difficult to maintain as devices move

from place to place.

4 Service Restrictions

Telephony emergency services impose restrictions on call features while an emergency call is in progress. For example, a caller cannot transfer a call, switch to a call-waiting call or put the emergency operator on hold. In peer-to-peer signaling, Internet telephony end systems would have to implement these behaviors.

5 Network-Layer Priority and Preemption

It may be desirable to provide reserved resources or a higher traffic priority to emergency calls. In RSVP, this is relatively straightforward for the caller-operator direction if network operators can authenticate emergency operators. Similarly, marking packets destined for the iPSAP with a higher-priority DS value in a diffserv environment requires no additional protocol support and is a local operational issue likely to be dealt with by national regulators.

6 Summary

This note proposes investigation into how Internet communications services can be employed to summon emergency assistance. In summary, we propose a number of protocol steps.

- Define a global emergency identifier for all communications protocols, including email, SIP, instant messaging and possibly IRC. This is similar to the existing conventions for `hostmaster`, `postmaster`, etc.
- Define a mechanism that ensures that communication directed at an emergency address is delivered to the appropriate iPSAP.
- Define a mechanism that allows end systems and/or ISPs to obtain location information about the system placing an emergency call, which is then passed on to duly authorized iPSAPs. A combination of end-system provided and ISP-based mechanisms seem most likely to be scalable and work for both wired, indoor as well as wireless, outdoor end systems.
- Ensuring priority for emergency communications does not appear to require additional protocol mechanisms.

7 Security Considerations

Internet-based emergency communications shares some of the same problems that traditional “911” services have, such as the possibility of prank calls and false alarms. Thus, such a system can only be deployed successfully if end users are identified by location. Location information is highly sensitive and thus must be protected from disclosure to inappropriate parties. In the ISP-based location model, standard access controls, encryption and authentication of the iPSAP are required, but are well understood.

References

- [1] National Emergency Number Association, "9-1-1 tutorial." <http://www.nena9-1-1.org/desktop/9-1-1>
- [2] Telcordia, "E9-1-1 service description," Special Report SR-4163, Telcordia, Morristown, New Jersey, May 1997.
- [3] J. Rosenberg, H. Schulzrinne, and H. Sinnreich, "SIP enabled services to support the hearing impaired," Internet Draft, Internet Engineering Task Force, July 2000. Work in progress.
- [4] D. Meyer, "Administratively scoped IP multicast," Request for Comments 2365, Internet Engineering Task Force, July 1998.
- [5] D. Crocker, "Mailbox names for common services, roles and functions," Request for Comments 2142, Internet Engineering Task Force, May 1997.
- [6] A. Vaha-Sipila, "URLs for telephone calls," Request for Comments 2806, Internet Engineering Task Force, Apr. 2000.
- [7] P. Faltstrom, "E.164 number and DNS," Internet Draft, Internet Engineering Task Force, Aug. 2000. Work in progress.
- [8] W. Zhao, H. Schulzrinne, and E. Guttman, "mSLP - mesh-enhanced service location protocol," Internet Draft, Internet Engineering Task Force, Oct. 2000. Work in progress.
- [9] R. Droms, "Dynamic host configuration protocol," Request for Comments 2131, Internet Engineering Task Force, Mar. 1997.
- [10] J. H. Reed, K. J. Krizman, B. D. Woerner, and T. S. Rappaport, "An overview of the challenges and progress in meeting the E-911 requirement for location service," *IEEE Communications Magazine*, Vol. 36, pp. 30–37, Apr. 1998.
- [11] P. V. Bahl and V. Padmanabhan, "Radar: An in-building rf-based user location and tracking system," in *Proceedings of the Conference on Computer Communications (IEEE Infocom)*, (Tel Aviv, Israel), Mar. 2000.
- [12] J. Ruutu, H. Tang, and J. Loughney, "Problems and requirements of some IP applications based on spatial location information," Internet Draft, Internet Engineering Task Force, Feb. 2000. Work in progress.
- [13] M. Korkea-aho, "Some scenarios for an ISL architecture," Internet Draft, Internet Engineering Task Force, Mar. 2000. Work in progress.
- [14] S. Nyckelgard and J. Loughney, "ISL architectural considerations," Internet Draft, Internet Engineering Task Force, Mar. 2000. Work in progress.
- [15] A. Costanzo, "Definition of the DNS GL resource record used to encode geographic locations," Internet Draft, Internet Engineering Task Force, June 2000. Work in progress.
- [16] C. Farrell, M. Schulze, S. Pleitner, and D. Baldoni, "DNS encoding of geographical location," Request for Comments 1712, Internet Engineering Task Force, Nov. 1994.

- [17] C. Davis, P. Vixie, T. Goodwin, and I. Dickinson, "A means for expressing location information in the domain name system," Request for Comments 1876, Internet Engineering Task Force, Jan. 1996.
- [18] H. Tang, "Target naming scheme," Internet Draft, Internet Engineering Task Force, July 2000. Work in progress.
- [19] W. Yeong, T. Howes, and S. Kille, "Lightweight directory access protocol," Request for Comments 1777, Internet Engineering Task Force, Mar. 1995.
- [20] K. Harrenstien, M. K. Stahl, and E. J. Feinler, "NICNAME/WHOIS," Request for Comments 954, Internet Engineering Task Force, Oct. 1985.
- [21] S. Williamson and M. Kosters, "Referral whois protocol (rwhois)," Request for Comments 1714, Internet Engineering Task Force, Nov. 1994.

8 Acknowledgements

Matthew Cannon, Dave Devanathan, Mark Handley, Jonathan Rosenberg and Henry Sinnreich participated in an earlier discussion on this topic. Jonathan Rosenberg and John Stanaway provided helpful comments on this draft.

9 Authors' Addresses

Henning Schulzrinne
Dept. of Computer Science
Columbia University
1214 Amsterdam Avenue
New York, NY 10027
USA
electronic mail: schulzrinne@cs.columbia.edu

Full Copyright Statement

Copyright (c) The Internet Society (2001). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.