

9-1-1 Calls for Voice-over-IP

Ex-Parte Filing for Docket 94-102

Henning Schulzrinne
Dept. of Computer Science
Columbia University
New York, NY 10027

February 28, 2003

Contents

1	Executive Summary	2
2	Background and Motivation	3
3	What makes 9-1-1 different in VoIP?	5
3.1	VoIP Environments	5
3.2	Addressing and Naming of VoIP Terminals	8
3.3	Separation of Transport and Signaling Service	8
3.4	Lack of Locality	9
3.5	Location-Independence	9
3.6	End System Intelligence	9
3.7	Independent of Media Type	10
3.8	International Standards and Cross-National Mobility	10
4	Proposed Solutions	10
4.1	Short-Term vs. Long-Term Solutions	10
4.1.1	Short-Term (Immediate)	10
4.1.2	Mid-Term (2004?)	11
4.1.3	Long-Term	11
4.2	User location	12
4.3	Identification of Emergency Calls	12
4.4	Allowing VoIP Terminals to Reach Traditional PSAPs	13
4.5	Dual-mode PSAPs	13
4.6	Security Issues	14

4.6.1	Authentication of Callers	14
4.6.2	Authentication of PSAPs	14
5	Standardization	15
6	Conclusions and Recommendations	15

1 Executive Summary

This document enumerates some of the major opportunities and challenges for providing emergency call (9-1-1) services using IP technology. In particular, all VoIP devices are effectively mobile. The same IP telephony device works anywhere in the Internet, keeping the same external identifier such as an E.164 number or URL.

Transitioning to an IP-based E9-1-1 infrastructure not only supports emerging VoIP systems, but can also greatly improve the capabilities of the emergency call infrastructure. It can become more resilient, sets up calls faster, better supports subscribers with hearing disabilities, adds multimedia information to better direct resources, allows competition for elements of the network infrastructure, conveys more call-associated data and allow more cost-effective PSAP technology.

The document recommends:

- Facilitating the transition of PSAPs from CAMA trunks to IP-based access even before VoIP becomes widespread, as this offers cost and response time improvements;
- Establishing a universal emergency identifier for new signaling protocols such as SIP;
- An ENUM-based mechanism to allow phone numbers to be mapped to location providers, allowing multiple, competing providers of location mapping services that does not rely on the ILECs;
- Adoption of emerging IETF solutions for updating and querying location databases for E9-1-1 use;
- Making the database of PSAP geographic coverage publically accessible at reasonable cost, so that end users, carriers and enterprises can directly map caller locations to the geographically-appropriate PSAP.

The document describes short-term, mid-term and long-term solutions that allow a graceful transition from the current, technologically outdated, infrastructure to a more capable and resilient 9-1-1 infrastructure.

2 Background and Motivation

As end-to-end voice-over-IP (VoIP) services are deployed for both residential and commercial users, users of these services rightfully expect that their emergency calls are answered at least as quickly and efficiently as in today's circuit-switched network. Thus, at a minimum, VoIP services and the existing 9-1-1 infrastructure needs to be enhanced to ensure such continuity of service.

However, the transition to packet-based communication also offers new opportunities for upgrading the national 9-1-1 infrastructure. The basic 9-1-1 communications technology has not changed significantly since the 1970s, with analog CAMA trunks using in-band multi-frequency (MF) signaling and low-speed data modems dominating in current PSAPs. There are a number of likely advantages when PSAPs upgrade to packet-based technology, even before a large fraction of telephone end users employ VoIP technology. Below, we summarize some of these advantages:

Higher resilience: As described later, VoIP cleanly separates the provision of transport services from call services, i.e., the transport of IP packets from providing end-to-end voice communication. Among other advantages, this makes it possible to answer emergency calls anywhere IP packets can be routed. As an example, a PSAP call-taker could easily take calls at home, over a cable modem or DSL broadband connection, without any loss of functionality. (With current technology, only locations equipped with CAMA trunks and connected to the same selective routers get ANI delivery, for example.)

Similarly, any PSAP would be able to take over the role of any other PSAP, in case the original PSAP for a specific region became overloaded or lost connectivity due to technical failure or natural disaster. VoIP technology also greatly speeds up setting up backup PSAPs, as the only requirement for a temporary PSAP is high-speed Internet connectivity and the updating of some call routing database entries.

Because of the service-transport separation, a PSAP could readily acquire network connectivity from multiple different Internet Service Providers, so that failures in one network are less likely to disconnect the PSAP.

Since pending calls, including calls on hold, do not consume network resources, it is less likely that calls will be blocked due to trunk overload.

Faster call set up: Current CAMA technology needs to "pulse out" the ANI over an analog trunk, resulting in additional call setup times of 10 to 15 seconds. This delay not only delays emergency response and ties up resources, but may also cause the caller to abandon and redial the call.

Better support for deaf and hard-of-hearing callers: While most PSAPs support TDDs, these are not readily available in all public buildings and private residences. Also, switching from voice mode to TDD mode adds additional delay to the emergency call.

In many cases, instant (text) messaging (IM) would provide a faster and cheaper alternative. IM support is available on almost all personal computer platforms and will be available on most mobile communication devices. As discussed below, adding emergency support to instant messaging is a fairly straightforward extension. IM may also be a preferred means of communications where background noise, accented pronunciation, fear of detection by intruders or bandwidth limitations make voice calls difficult.

Multimedia support: The current 9-1-1 system is limited to voice only, with some support for TDD. With packet technology, it becomes easy to add bidirectional video communications. As third-generation wireless devices and PCs are equipped with video cameras, video can provide valuable additional information to the emergency call-taker, helping her to better assess the nature of the emergency. Also, it may become possible for the call taker to provide video instructions to the emergency caller, e.g., on how to perform first aid. While such universal availability of such capabilities may seem far in the future, even limited use may provide significant benefit.

More cost efficient PSAP operation: Current 9-1-1 technology is specialized and costly since it serves a relatively small niche application; CAMA trunks, for example, are effectively only used for 9-1-1 service. Since a CAMA trunk is limited to serving four area codes, backup PSAPs spanning a larger area become difficult. With VoIP, there is evidence that standard call handling elements require little or no modification to deal with emergency calls. Thus, a PSAP can, in effect, use standard VoIP call-center technology to handle calls, participating in the efficiency-of-scale and technology advancement benefits of off-the-shelf technology.

More call-associated data: Currently, most PSAPs use low-speed modems, from 1,200 to 9,600 b/s, to obtain location information (ALI data) from their ILEC. This limits the amount of data that can be conveyed and adds further delays. Even with wireless phones, locations have to be mapped to telephone numbers, which are then used as temporary keys into a database.

With VoIP, it becomes possible to convey call-associated data both with the call itself and to retrieve data out-of-band, on separate data lines. Thus, it becomes easy to draw on multiple databases for information such as chemical hazards, medical information on residents or floor plans.

No reliance on telephone numbers: As noted, the existing 9-1-1 technology relies on phone numbers as keys to retrieve data. This greatly complicates the task for wireless devices and will not be possible at all for devices that may not have a telephone number at all.

More competition: It is currently difficult for entities other than the ILEC to provide emergency call services. With VoIP services, the playing field between all providers is leveled. There is no need for a single integrator of MSAG and ALI data, for example.

If ALI databases are provisioned competitively, subscribers and/or service providers could choose a provider that offers the best services at the best price. (For example, di-

rectory service providers might offer subscribers the ability to enter additional medical information into the directory.)

In this note, we first describe the fundamental differences between VoIP and circuit-switched technologies that make it inappropriate to simply transfer the existing mechanisms to the new environment (Section 3). We then describe some of the detailed challenges and possible solutions to address some of these issues, pointing out where regulatory action may be helpful (Section 4).

Since most of the publically available VoIP services, including proposed third-generation wireless services, employ the Session Initiation Protocol (SIP) [1] for session setup, we focus our attention on systems using this protocol. However, most of the discussion applies to other signaling technologies as well, such as the ITU H.323 signaling system or the Megaco protocol.

3 What makes 9-1-1 different in VoIP?

3.1 VoIP Environments

The difficulty of providing 9-1-1 service depends to a large extent on the architecture of the VoIP service. For the purposes of this discussion, it may be helpful to distinguish three types of services, namely *end-to-end* (pure) VoIP, *long-distance* (hybrid) VoIP and last-mile VoIP.

In *end-to-end or pure* VoIP, all end points use packets to exchange multimedia information. The traditional telephone network is not involved at all. Such services require no application-layer service provider. In other words, the Internet Service Provider may not even be aware that the packets it carries across its network are actually part of a voice conversation. End systems in this scenario may include personal computers running multimedia conferencing applications, mobile devices using licensed and unlicensed spectrum or desktop IP phones equipped with Ethernet connectivity. Telephony services may be provided by the end users themselves, similar to how corporate or institutional users provide their own email service without relying on a classical telecommunications company. For residential users, portal services similar to the current free email services may provide call routing and naming, but without themselves forwarding any voice or video packets. Thus, in this scenario, only end users would be usefully subject to regulation.

In a *long-distance hybrid* VoIP environment, VoIP users connect primarily to PSTN terminals via gateways provided by a service provider. The gateways themselves may be located in a few places spread around the world, rather than being present in every LATA, E-911 tandem serving area or local calling area.

The pure and one-sided environment are likely to co-exist within the same provider. Calls between customers of that provider will stay purely VoIP, while calls to non-customers are handled via PSTN gateways.

Finally, in the *last-mile* scenario, a provider of cable modem or DSL service also provides packet voice service, but calls are translated back into circuit-switched calls at the cable

headend or DSLAM so that all customers for a single DSLAM or cable headend are within the same E-911 tandem serving area. (If the latter condition is not true, the last-mile scenario is transformed into the hybrid scenario.)

Also, we need to distinguish between PSAPs that are only reachable via PSTN circuits as well as PSAPs that are reachable by both VoIP and circuit-switched calls. For brevity, we refer to the former as traditional PSAPs and the latter as dual-mode PSAPs.

Providing 9-1-1 services is not possible in the combination of traditional PSAPs with unconstrained end-to-end VoIP. In other scenarios, we need to address two problems: how does the IP endpoint get connected to the correct PSAP (abbreviated as “PS” (PSAP Selection) below) and how does the PSAP obtain the correct geographic location of the caller (user location or “UL”). Table 1 summarizes which architecture and issues are connected.

PSAP selection (PS): The caller needs to reach the correct PSAP, i.e., the PSAP that can dispatch emergency help for the caller’s current location.

For the hybrid case, the gateway into the PSTN can be located anywhere. Typically, the number of gateways is orders of magnitude smaller than the number of police, fire and rescue jurisdictions. Thus, 9-1-1 calls placed from one of the gateways will end up in the PSAP serving the location of the gateway, not the location of the caller. Even if the gateway knew the geographic location of the caller, the gateway would have no way to convey this to the PSAP. This problem is very similar to call being handled by telematics emergency service providers and has similar problems in terms of determining 10-digit numbers and conveying caller ANI. (The National Emergency Number Association (NENA) is proposing that the exchange “911” be set aside in each area code for PSAPs.)

Access consolidation by multi-location businesses may become a common example of the hybrid architecture. A business with branches in several locations may handle all internal voice communications via VoIP and maintain a single gateway for the whole organization. Due to volume discounts on voice traffic, the smaller number of gateways and statistical multiplexing, maintaining a single gateway may well be significantly more cost-efficient than placing a gateway in each location. While telematics providers may be able to obtain and maintain lists of 10-digit telephone numbers for PSAPs, it is unrealistic to expect these small multi-location businesses to track and update this information. (Reaching PSAPs via a 10-digit number is, at best, a temporary kludge, since many services are not available through such numbers and such calls may not receive the same treatment as true 9-1-1 calls.)

For the last-mile architecture, the difficulty of the PSAP selection issue depends on the geographic scope of the PSTN gateway. If all subscribers for such a gateway are located within the same PSAP serving area, the gateway can just place a regular 9-1-1 call. If the DSL or cable modem provider backhauls all voice call to a small number of gateways, the PSAP selection problem becomes similar to the hybrid architecture discussed earlier.

User location (UL): In many architectures involving traditional PSAPs, it may be difficult for the PSTN gateway to first determine and then convey the caller location to the PSAP. We briefly describe the two problems separately:

A gateway or other VoIP entity may have difficulty *determining* the location since none of the information in the VoIP call is tied to geographic location. There are only two such elements that a gateway can glean from a call setup request: the layer-3 (Internet) address and the VoIP signaling address (see Section 3.2). The IP address is sufficient to identify a large Internet Service Provider or a large organization. However, many smaller Internet Service Providers obtain address blocks from their upstream providers. This delegation is not public and likely considered proprietary data by the ISP as it allows to deduce the number of customers served. Similarly, while each point of presence has a well-defined set of IP addresses, the mapping from an address to such a POP is known only to the provider. Thus, a third party other than the ISP has no reliable way of mapping an IP address to even the granularity of a metropolitan area.

Particularly for consumers, IP addresses are assigned dynamically, so that, at different times, the same address can be used by many different ISP customers.

The problem is complicated by the common use of Network Address Translation (NAT) by organizations. Using NATs, an organization may dynamically map a small set of ISP-provided IP addresses into a much larger set of local IP addresses. Thus, the same end system may well appear to have multiple different addresses at different times or different end systems may have the same address. These devices can be spread over a very large geographic area, spanning continents.

The use of virtual private networks (VPNs) further removes any geographic association from IP addresses. With VPNs, telecommuters and traveling employees “dial into” their home network and acquire an IP address from that home network. As IEEE 802.11 wireless hot spots become more popular, it appears likely that business travelers will use these to establish VPN connections and to place calls over these VPNs. (Such an arrangement may, for example, allow the traveler to use the PSTN gateway operated by their employer.) Thus, a caller whose IP address range places her in a company in Chicago may actually be physically located in a hotel in Los Angeles.

Finally, as described in Section 3.2, the signaling address has properties similar to an email address and thus conveys no geographic information.

Even if a gateway can determine the geographic location of the caller, it may not be able to convey this information to the PSAP. The best approach for traditional PSAPs is probably to treat the gateway like a PBX in a multi-line telephone system (MLTS). As in an MLTS, the gateway would maintain a map between telephone numbers (ELINs) and locations and then assign the call the appropriate telephone number. (It will likely take too long to dynamically update the MLTS location information at the time a person moves.)

The problems described make the user location problem particularly challenging for the end-to-end VoIP and the hybrid case. The user location problem is least severe for

the last-mile architecture since the network provider generally knows which device has acquired which IP address and can therefore, via its customer billing database, map this Internet address to a physical location, in a manner similar to the current circuit-switched telephone network.

	traditional	dual-mode PSAP
End-to-end	not possible	PS, UL
Hybrid	PS, UL	PS, UL
Last-mile	N/A	N/A

Table 1: PSAP Selection (PS) and User Location (UL) problems

3.2 Addressing and Naming of VoIP Terminals

VoIP terminals will likely be addressed by two kinds of identifiers, E.164 telephone numbers and Internet-only URLs. Examples of the latter include SIP URLs, such as `sip:alice@example.com` or `sip:chairman@fcc.gov`. It appears likely that many users will be able to use one or more of their email addresses as VoIP addresses, although there is no inherent technical connection between the two.

Unlike telephone numbers, these Internet URLs are plentiful, portable and cheap, just like email addresses. Users can either obtain such an address from their ISP, typically using a domain identified with that provider, from their employer or professional or affinity organization, use a freemail service or purchase a domain for themselves, their business or their family. In the latter case, they have perfect “number portability”, as they can easily switch the provider of their voice signaling service.

Similar to “800” numbers, these numbers often identify not a single individual or residence, but rather a function, for example `sip:sales@acme-widgets.com`.

The ENUM Domain Name Service allows the translation of telephone numbers into SIP URLs and other Internet-only URLs.

3.3 Separation of Transport and Signaling Service

Providing telephone service in the Internet can be cleanly separated into three independent components: delivering IP packets (“transport”), locating users (“signaling”), and auxiliary services such as conferencing, voice mail and similar services. Unlike in the PSTN, transport and signaling services have no inherent connection. In this respect, using the Internet to place voice calls is very similar to how email service works today. Just like for email service, it appears likely that these services will be provided either by the ISP, by freemail-like services or by organizations that run their own network.

These signaling services are limited to routing call setup messages. They do not need to forward or process voice packets at all and may not even be aware what kind of media (voice, text, video, etc.) is being exchanged between the caller and the callee.

Compared to email, the cost of providing telephone signaling services is likely to be lower, since the provider of such a service does not have to receive, store and forward the actual email message, but rather just relatively short signaling messages.

This fundamental principle of separation means that the notion of a voice carrier loses much of its meaning in this environment.

VoIP in a sense completes the transition from in-band signaling found in the early analog telephone system to the common channel signaling in the modern digital telephone system to a completely out-of-band signaling mechanism. Signaling messages and voice packets for a single call may only coincide at the end systems. Also, unlike the PSTN current signaling system, the signaling messages use the same high-speed network as voice data, making signaling a very cheap operation. This greatly simplifies many 9-1-1 related tasks since there is little cost in sending signaling messages to far-away places not on route to the PSAP. The common problem of “tromboning” of voice calls largely disappears. For example, an emergency call request by an American traveling in Italy could easily consult a database in the United States before being routed to the appropriate Italian emergency service agency, at an additional delay of at most a few hundred milliseconds.

This separation of transport and signaling services is the foundation for new services and offers the opportunity for increased consumer choice. Thus, it can be considered a core promise of VoIP and should be jealously guarded.

3.4 Lack of Locality

By their nature, PSAPs serve a specific geographic region defined by local jurisdictions. These service regions are almost always contained within a state. As noted, end-to-end and hybrid VoIP services have very little relationship to geography. For example, a gateway in Vermont may handle an emergency call from caller in Massachusetts, making many emergency calls *interstate calls*.

3.5 Location-Independence

With VoIP, even “stationary” devices like personal computers or IP telephones need to be treated like mobile devices. One can easily take an IP telephone and plug it into any suitable Ethernet jack within the same organization or a completely different one, without affecting the ability of the user to be reached under the same application-layer identifier.

3.6 End System Intelligence

While most services in the PSTN, from call forwarding to call blocking, are implemented in switching gear operated by carriers, all VoIP services can be implemented in equipment and software operated by either end users or their service provider. As a rough analogy,

this makes the same services available to end users as if each such user had direct access to Signaling System 7 capabilities.

3.7 Independent of Media Type

The moniker “voice-over-IP” is somewhat misleading. None of the technology used for current VoIP services is voice-specific, except naturally gateways to the PSTN. The same signaling messages that set up a voice call can just as easily set up a video session, a text chat or a multiparty game. Thus, it is likely that users will expect that media other than voice can summon emergency help.

In particular, one of the standardized mechanisms for instant messaging and presence uses the same signaling protocol (namely, SIP) as for voice calls, sharing the same signaling infrastructure, naming and request routing and appearing on the same device.

3.8 International Standards and Cross-National Mobility

While emergency calling has traditionally been considered a local or at best national issue, the use of mobile devices and increased travel has led to trans-national efforts. For example, the European Union has agreed to a single emergency number (112) valid in all member states. (Earlier nation-specific numbers continue to work.)

In addition, IP equipment and software is commonly used across national borders, as Internet standards and even mundane items like Ethernet jacks know no national distinctions.

Thus, one objective of modernizing the 9-1-1 system should be to provide a uniform mechanism across the world. With VoIP, this is not implausible, as there are no existing conflicting standards in this area. Thus, there is a short-term opportunity to achieve global interoperability.

4 Proposed Solutions

4.1 Short-Term vs. Long-Term Solutions

Given the large number of PSAPs and the slow transition to an all-IP network, it will likely take many years before all PSAPs are IP-capable. Thus, it is important to design an architecture that allows a smooth migration and interoperation of different generations of technology, without forever limiting progress to the functionality offered by the oldest technology. There are several intermediate steps.

4.1.1 Short-Term (Immediate)

In the short term, VoIP terminals can be treated as either PBX extensions (multi-line telephone system) or mobile phones. Either mode requires that VoIP terminals reach the appropriate gateway that must be within the same E9-1-1 tandem serving area. If a terminal moves,

for example from one corporate office to another, a location database needs to be updated, just as for traditional circuit-switched PBXs. This implies, however, that any VoIP carrier or corporation maintains at least one gateway in each E9-1-1 tandem serving area where it has users. (As noted, avoiding purchasing PSTN tail circuits at each branch location is one of the economic incentives for converting to VoIP.) This greatly increases the cost particularly for smaller VoIP providers unless they can contract out this service to a third party.

PSAPs should also offer 800 numbers for access by gateways outside their coverage area. This is not a substitute for a real translation and location service and offers inferior service, but appears to be the only approach that avoids the need for maintaining a gateway in each E9-1-1 tandem serving area.

No changes to PSAPs are required. Using the technology described in Section 4.2, users would be able to find the appropriate gateway. Initially, user devices may need to be configured with the network address of a suitable translation service or all calls must be routed through outbound proxies operated by the provider or enterprise. The latter appears to be common today, so it does not impose a significant new burden. The outbound proxy must then be configured to route the call to the appropriate PSTN gateway that is local to the caller. Manual configuration is possible, but undesirable, so we propose a mid-term solution below.

It is not clear whether the current Phase II wireless location system for mobile operators would scale to large numbers of enterprises. It appears that individual tests and configuration are necessary for each wireless carrier, a procedure that is unlikely to work with large-scale deployments across enterprises.

4.1.2 Mid-Term (2004?)

To allow VoIP terminals to locate the appropriate PSTN gateway, database providers that offer translation from geographic location to PSAPs should make scalable, non-proprietary IP interfaces available. The access has to be non-discriminatory and reasonably priced so that even small enterprises can access this information. Details are below in Section 4.4.

Similarly, there needs to be a way for individual users, small carriers and enterprises to update their location information. Here, ENUM DNS technology may be useful. Each number is mapped, via DNS, to a location service provider chosen by the user. The PSAP either directly contacts the DNS server to obtain this information or the database provider contracted by the PSAP performs the lookup and retrieval. The DNS ENUM entry points to a service provider that can then be queried for the current location of the user, based on an E.164 telephone number or URI. This allows full competition for providing such services. It does not raise significant privacy concerns since the ENUM entry only identifies a service provider and does not grant access to the user location itself. The PSAP or the intermediary that it authorizes has to authenticate with the lookup service.

4.1.3 Long-Term

Many of the advantages of IP-based emergency calling services can only be realized if the emergency call remains IP-based, without translation between circuit-switched and packet-

switched domains. Fortunately, once the two core directory components (user location and PSAP directory) noted above and described below are IP-enabled, the transition to IP-based PSAPs, including dual-mode ones (Section 4.5) combining circuit and packet technology, is relatively straightforward and requires little or no additional PSAP investment.

4.2 User location

We discussed earlier that transport and signaling services are independent and should remain so for sound technological and competitive reasons. However, transport service providers know the physical location of their customer circuits and connections. This knowledge may be recursive. For example, an ISP offering Internet or VPN services to a company will know the termination point of the tail circuit to that company, while the company will know or can determine the physical location of the network devices connected to their local area network.

There have recently been proposals [3, 4] how a commonly used end system configuration protocol, the Dynamic Host Configuration Protocol (DHCP), can be trivially extended to convey approximate user location to the end system. Almost all ISPs and many corporations already employ this protocol for automatic IP address assignment. With this addition, terminals can easily determine their geographic (longitude and latitude) and/or civil (street address) location, avoiding the need to configure devices manually.

It appears desirable that providers of fixed-line Internet services and any future 802.11 wireless providers should be encouraged to provide this information to devices, similar to the requirements for Phase II service for cellular phone. The cost of this additional service is likely to be minimal since it requires no new databases or communication facilities. It may be sufficient to create a voluntary “Internet 911 ready” certification program for ISPs.

For wireless providers, the transmission range of most 802.11 base stations is well below the 100 m accuracy requirement of the Phase-II wireless 911 requirements for network-based solutions. Thus, such service providers, offering Phase-II wireless 911 accuracy is technically easy and requires trivial expense, compared to traditional 2nd and 3rd generation digital wireless systems.

Unlike in the PSTN, VoIP terminal can provide user location either in-band and out-of-band. In in-band mode, the signaling request itself contains the location information, either inserted by the end system or a signaling proxy along the way. In out-of-band mode, the end system provides an identifier, which is then used by the PSAP to look up the current location information. The current PSTN operates only in out-of-band mode, where an ELIN (Emergency Location ID Number) serves as a key into a location database.

4.3 Identification of Emergency Calls

There is a large benefit if all signaling components can easily identify emergency calls. We have suggested [5] that a single universal and international identifier, “sos”, is reserved for this purpose. Among other benefits, this allows end systems to come preconfigured with an emergency call button and allows signaling servers to handle such calls appropriately.

In addition, each VoIP device should recognize the two most widely used three-digit emergency number, 911 and 112, as emergency identifiers.

Having more than one identifier imposes no appreciable burden on VoIP end systems or servers.

4.4 Allowing VoIP Terminals to Reach Traditional PSAPs

For the next few years, an increasing number of VoIP terminals will have to reach traditional, non-VoIP PSAPs. As noted earlier, the primary problems are PSAP selection and user location. This problem can be solved [5] if a number of conditions are met:

- The PSTN gateway needs access to a mapping database that maps geographic or civil coordinates to a PSAP. Since gateways are operated by a large number of non-carrier organizations, access to this database should require no more than regular Internet access and use non-proprietary, internationally standardized protocols. There can easily be multiple databases, operated by for-profit, not-for-profit and governmental agencies. The database storage volume is rather modest and the access frequency low (at most about 10 requests/second across the whole United States), so that a state-by-state database, with replication, can be maintained at low cost.
- For each traditional PSAP, the database needs to contain a ten-digit number that can reach an emergency call taker, rather than an administrator. For Internet-enabled PSAPs, a suitable contact URI should be provided.
- Internet service providers, as noted in Section 4.2, should provide end users with geographic location information.
- Vendors of VoIP end systems should allow users to configure location information, to allow use when the ISP does not provide this information.

While these proposals address the PSAP selection problem, it may be technically infeasible to convey the user location for the hybrid VoIP architecture in all cases. (The gateway may obtain the location, but this location may be outside the local PSAP service area. If the gateway places a call to a 10-digit number, the PSAP will obtain a number provided by the gateway as ANI, but cannot map this number to a location and may treat the call as an administrative call, not an emergency call. The only technically feasible solution appears to be similar to how wireless Phase II calls are handled. However, the complexity of such a solution may be prohibitive for gateways that are not owned by telecommunications carriers.)

Note that implementation of these mechanisms also allows other services, such as the NXX service codes, to work for VoIP terminals.

4.5 Dual-mode PSAPs

Even while most telephone users still use circuit-switched technology, there are numerous advantages for PSAPs to convert from CAMA trunks to VoIP connections. As noted above,

this significantly decreases the call setup latency, avoids the digit carrying limitations of CAMA trunks and provides faster data connections for ALI lookups.

The provisions needed for traditional PSAPs will also enable dual-mode PSAPs to function correctly, except that the user location issues can be resolved. The caller or service provider simply includes location information in the call setup request. Standards addressing this need are currently under development within the Internet Engineering Task Force (IETF).

In this model, the signaling server in the caller's domain recognizes the emergency call, consults a location-to-PSAP database and then simply routes the call request to that PSAP. Alternatively, it can route the call setup request to a service provider that in turns directs it to the PSAP. Since VoIP signaling is fast and efficient, the added indirection adds negligible delay.

4.6 Security Issues

Authentication of callers and PSAPs will play an increasing role for VoIP since traditional means of authentication via a trustworthy ANI do not exist.

4.6.1 Authentication of Callers

We noted that Internet identities are cheap and plentiful, so that simply requiring public key signatures on call requests offers little additional security. A crank caller can just obtain a "freemail" identity and use that to place an emergency call.

Internet service providers or financial services companies are probably in the best position to issue certificates that associate a name with a billing address. Such a certificate could then be used to sign an emergency call request.

Unfortunately, unless all callers are required to provide such proof of identity, any such authentication system does not deter crank calls.

A more serious threat are distributed denial-of-service attacks, where a set of compromised Internet hosts would conspire to reach a certain PSAP, thus interfering with its ability to function. PSAPs and their service providers may need to implement back-tracing mechanism to quickly staunch the attack.

4.6.2 Authentication of PSAPs

Since many 9-1-1 calls contain sensitive information in both the call setup messages and the conversation, the caller should be assured that she is indeed reaching a legitimate PSAP rather than an impostor. This can be ensured by transitive trust, a well-known domain name suffix for PSAPs or an attribute credential. In transitive trust, the entity that translates caller location to a particular PSAP needs to be trusted by the caller. The entity making the database request can use standard web-like security mechanisms to ensure that it is indeed connecting to the right database, that the information returned has not been tampered with and that it is forwarding the call setup request to the correct PSAP.

In the second approach, the Internet Corporation for Assigned Names and Numbers (ICANN) could designate a new top-level domain restricted to emergency call centers only. For example, all United States PSAPs could register under the “.us.sos” domain. The PSAP can then sign the response, ensuring the caller that a connection to a valid PSAP has been established.

Longer-term, so-called attribute credentials can be used to certify that a particular domain is indeed a PSAP. Both of the last two proposals are simple to implement technically, but may need some mechanism to certify PSAPs.

5 Standardization

A number of organizations are currently engaged in standardizing the technical components of a VoIP-based emergency calling system. The Internet Engineering Task Force (IETF), as the standardization body for Internet protocols, is discussing, in its SIPPING working group, proposals for identifying and routing emergency calls. The GEOPRIV working group is discussing the two DHCP-based location proposals [3, 4] mentioned earlier, as well as the format and privacy issues for conveying location information more generally [6].

The North American Emergency Number Association (NENA) has a Migration Working Group that is addressing issues and architectures for connecting PSAPs to selective routers via VoIP technology.

6 Conclusions and Recommendations

In this note, we have attempted to summarize the core differentiators between circuit-switched and packet-switched voice services and how they impact the provision of emergency calling services. We have identified a number of new features that can greatly enhance the capabilities of emergency responders, while making their communication facilities more resilient.

To address the problems of PSAP selection and user location, we have suggested a number of actions that would expedite the ability of VoIP terminals to make emergency calls and to enhance the capabilities of existing PSAPs. These recommendations include:

- Mandate the support of a single, global emergency identifier in addition to “911”, in international cooperation. We suggest the use of the internationally recognized identifier “sos”.
- Encourage the use of a single “exchange” (911) within each area code to allow 10-digit dialing of PSAPs.
- Support appropriate international and national standardization activities that allow IP-connected devices to obtain and convey their geographic location.
- Encourage Internet service providers to deliver access point location information to devices.

- Mandate open, non-discriminatory access to electronic versions of the jurisdictional (location-to-PSAP) maps since such access is necessary for the inter-state completion of emergency calls.
- Support the transition to multiple, competing user location providers, using ENUM as the neutral mediating technology.

References

- [1] J. Rosenberg, H. Schulzrinne, G. Camarillo, A. R. Johnston, J. Peterson, R. Sparks, M. Handley, and E. Schooler, "SIP: session initiation protocol," RFC 3261, Internet Engineering Task Force, June 2002.
- [2] H. Schulzrinne and K. Arabshian, "Providing emergency services in internet telephony," *IEEE Internet Computing*, vol. 6, pp. 39–47, May 2002.
- [3] H. Schulzrinne, "DHCP option for civil location," internet draft, Internet Engineering Task Force, Dec. 2002. Work in progress.
- [4] J. Polk *et al.*, "DHCP option for geographic location," internet draft, Internet Engineering Task Force, Oct. 2002. Work in progress.
- [5] H. Schulzrinne, "Emergency services for Internet telephony based on the session initiation protocol (SIP)," internet draft, Internet Engineering Task Force, Dec. 2002. Work in progress.
- [6] J. Cuellar *et al.*, "Geopriv scenarios," internet draft, Internet Engineering Task Force, Nov. 2002. Work in progress.