# Providing Emergency Services in Internet Telephony

A SIP-based architecture for emergency calling and notification systems could increase speed, scalability, and functionality in communication services.

**Henning Schulzrinne and Knarig Arabshian**
*Columbia University*

During emergencies, service agencies rely on telecommunications to achieve quick response times and minimize loss of life and property. The existing public and government-based telephone systems provide three essential communication services: an emergency calling system, which the public uses to report problems and ask for help; emergency communications, which allow for crisis communications within agencies and between agencies and the public; and emergency alerting, which provides a channel for government agencies to notify the public when disasters occur. As we transition to a packet-switched infrastructure, we must reconsider how to provide these services. Internet-based communications offer new challenges, as old assumptions about trust, operation, and terminal location no longer hold. However, IP-based emergency services will also offer expanded services, more resilient networks, and faster response times.

About half a dozen companies are offering commercial voice-over-IP services at this time, for both commercial and residential customers. It appears likely that the transition to an all-IP phone service will be slow, as depreciation intervals for switches are measured in decades. However, third-generation wireless systems (3G) are scheduled to offer packet voice services by 2005, adding possibly hundreds of millions of such devices to the network.

As the "Emergency Communications" sidebar describes, the existing Internet architecture must be modified in two areas to support coordinated communication services. Here, however, our primary focus is on emergency calling and notification. In this article, we describe the components of the existing emergency calling and notification systems and our proposed IP-based architectures,[1] each of which uses the session initiation protocol (SIP) as the signaling framework.[2]

1089-7801/02/$17.00©2002 IEEE

## SIP: An Overview

We based our proposed emergency communications systems on SIP because it is one of the most common signaling protocols. Also, next-generation Internet telephony networks are using SIP, including those proposed by PacketCable for cable modems and the Third-Generation Partnership Project (3GPP) for next-generation wireless.

In SIP-based networks, operators identify subscribers either by E.164 telephone numbers[3] or SIP URIs, such as `sip:alice@example.com`, which are independent of the device's IP address. A caller, represented by a user agent, initiates the call by sending a SIP `INVITE` message to a local outbound proxy or a SIP server proxy in the destination domain. To convey a current network address, a user agent periodically sends a SIP `REGISTER` request to the home SIP server, with the address in the `Contact` header. Thus, a binding is created from a generic address-of-record, such as `alice@example.com,` to a current network location, such as `alice17@pc42.accounting.nyc.example.com.`

The end point identifies SIP messages by their source and destination (`From` and `To` headers), as well as by a call-and-request sequence number. Networks can carry SIP messages with the user datagram protocol (UDP), transmission control protocol (TCP), or stream control transmission protocol (SCTP). The message format is similar to HTTP messages: plain text headers followed by an opaque message body. The message body carries a session description that enumerates the call's media streams.

Developers have also extended SIP to generate event notifications[4] and instant messages.[5] Users subscribe to an event with the `SUBSCRIBE` method and receive notifications via `NOTIFY`. Although event notification is typically used for presence notification and event signaling during telephone calls, we can use it as an emergency alert mechanism.

## Emergency Calling: The Current System

Most countries have a telephone-based system that lets their citizens summon emergency help, such as an ambulance, or the police or fire departments. All such systems have four components:

- *Universal number:* Simplifies access to services by providing a single number to dial for help, such as 911 in the U.S. and Canada, or 112 in many parts of Europe.
- *Call routing*: Allows the central office to use the caller's location to determine the most appropriate emergency response center and routes the call there.
- *Caller number identification*: Allows call centers to limit prank calls, call the person back if they get disconnected, and log calls for evidence.
- *Caller location*: Speeds response and assists callers in identifying their current location. For landline phones, call centers obtain the caller's street address from the subscriber billing address; they typically identify cell-phone users' geographic location using built-in global positioning system (GPS) receivers or network-assisted solutions (based on time-of-arrival differences, for instance).[6]

The U.S. established 911 as the universal emergency number in 1968, for example. When a 911 call reaches a central office, the switch consults the selective routing database, which maps the caller's telephone number to a three- to five-digit emergency service number that includes fire, police, and emergency medical services agencies. Each emergency service number is associated with a primary and secondary public safety answering point (PSAP), which answers the call and possibly transfers it to the appropriate public safety agency. The PSAP also obtains the caller's street address using a data connection to the telephone company. Although details and terminology differ, other countries use similar approaches.

## Toward a SIP-Based Call Architecture

Once we transition to IP telephony, many existing underlying assumptions will no longer hold. The problems depend on whether we assume that the PSAPs are aware of IP telephony or are seeing IP telephony only through gateways. We refer to these as an IP-enabled PSAP and a legacy PSAP, respectively. A primary problem is that the current system assumes a central mapping from telephone number to street address, maintained by a single telephone company for each household. With IP, however, subscribers can obtain their Internet service from one company and address (such as the SIP URL, sip:alice@example.com) from another. It's indeed plausible that the same user@domain identifier could serve as both a subscriber's e-mail address and IP telephony identifier.

### Locating IP Devices

Clearly, a central problem in the transition to IP telephony is how to locate IP devices. As with e-mail addresses, SIP URLs are not associated with

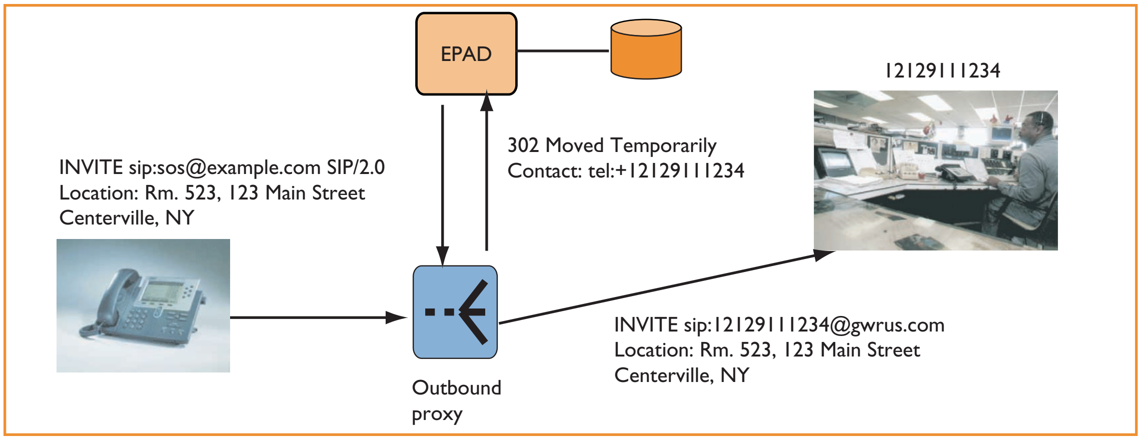
EPAD
12129111234
302 Moved Temporarily
Contact: tel:+12129111234
INVITE sip:sos@example.com SIP/2.0
Location: Rm. 523, 123 Main Street
Centerville, NY
INVITE sip:12129111234@gwrus.com
Location: Rm. 523, 123 Main Street
Centerville, NY
Outbound proxy

*Figure 1. SIP call routing for emergency calls. The outbound proxy accesses emergency services address information through an emergency provider access directory (EPAD). The address shown (Rm. 523, 123 Main, Centerville, NY) is the caller's address; the phone number (+1-211-911-1234) is for the PSAP.*

fixed locations or IP addresses. Because SIP signaling typically traverses multiple proxy servers (or network address translation devices), an IP-enabled PSAP receiving a SIP call will not necessarily have access to the caller's IP address. Thus, none of the traditional identifiers that are roughly equivalent to phone numbers can reliably identify a terminal or terminal location.

Independent of the overall architecture, we must be able to ascertain the location of indoor, wired IP devices. On some platforms, we can use the simple network management protocol (SNMP) to determine the Ethernet switch port of a particular media access control (MAC) address, which would allow us to determine each device's physical Ethernet jack and location (such as room and building). Unfortunately, this only works with managed hubs and switches, and still requires an accurate wiring database. If we only know the switch location, CAT 5 — category five of twisted-pair cabling systems — or fiber wiring can easily induce uncertainties of several floors or even miles.

Other possible approaches to device location include

- *Manual entry.* IP phones typically have one or more "owners," which the phone or a server can contact by e-mail or other mechanisms. Manual entry might provide a viable tracking mechanism if users had to enter a physical location each time they moved a phone. Although hardly ideal, this approach could be practical for phones that do not move frequently.
- *Ethernet enhancement.* We could enhance Ethernet switches by periodically sending, on each port, a broadcast packet that identifies the location. In a typical multistage, switched Ethernet, each device would receive multiple location packets, but these would provide incremental information, such as "Building 4" and "Jack Room 4F523." Such functionality would also be useful for asset management.
- *Smart jacks.* In some commercial products, such as Panduit's PanView and RiT Technologies' PatchView, jacks are active components, and we could query them for attached MAC addresses. Recently, 3Com introduced Ethernet jacks that contain an Ethernet switch.
- *Wireless-like approaches.* Although GPS doesn't work indoors, assisted GPS might. Some have suggested that we also use digital television station signals for location. Typically, however, cellular location is accurate to around 100 meters, which is insufficient for locations such as office buildings or high-rise apartments.
- *Infrared/radio frequency (IR/RF) location.* Many asset-tracking products use IR transmitters and sensors. Such approaches might work in commercial environments, although they'd add about US$50 to $100 to the cost of each device, making them impractical for widespread residential deployment.

Because each of these options has different tradeoffs in cost, reliability, and compatibility with existing systems, they're likely to be used in combination to locate IP devices.

## Legacy PSAPs

Solutions to the device-location problem also vary depending on whether PSAPs are IP-enabled or not. A legacy PSAP will identify an incoming SIP call by the terminating gateway's telephone number, but that gateway might be nowhere near the original caller. Dispatching the fire department to that location would thus be unlikely to help, and might in fact cause harm by tying up emergency service crews. Even if the caller ID reflected the correct location, the gateway connecting to the emergency number would reach a local selective routing database, which wouldn't have addresses or emergency service numbers for distant locations.

When the gateway is located in the same place as all the phones that connect to it, and the IP phones don't wander off (using virtual private networks, for example), some solutions for device location are possible. For example, Cisco has proposed assigning a unique telephone number to each Ethernet jack.[7]

There's also an intermediate approach that would work for gateways that are physically distant from the IP telephones: We could publish a directory that lists regular, non-911 numbers for each PSAP, along with the PSAP's service area. Each gateway could consult the directory based on what it knows about the IP telephone's location information. Another proposal is to assign the 911 exchange in each area code to PSAPs (such as 201-911-*XXXX*) and then attempt to determine the caller's equivalent area code (for details on this, see www.nena9-1-1.org/9-1-1%20Tutorial/9-1-1_tutorial.htm).

## A Proposed Architecture

Figure 1 shows an outline of our architecture, which supports a mixture of IP-enabled and legacy PSAPs (For a detailed specification, see www.cs.columbia.edu/sip/emergency.html). As PSAPs become IP-enabled, network administrators need only update database entries in a few places. During an emergency call, the IP phone would contact the local outbound proxy, as it does for every call. We propose, however, that it use `sip:sos@local-domain` as the universal destination for emergency calls.[8] Upon receiving this special identifier, the outbound proxy would intercept the setup request and try to determine the caller's location. An end system that can determine its own location would include the information in the request; otherwise, the outbound proxy would use the MAC-backtracking mechanism described earlier. As a last resort, the outbound proxy would assume that the device was nearby, indicating the uncertainty and relying on human interaction to determine the precise location.

Because the outbound proxy does not want to maintain a PSAP database, we propose establishing a national or regional SIP-based call router to register with the proxy as user `sos`. The SIP-based call router must be subject to appropriate authentication. In keeping with existing terminology, we'd label these routers emergency provider access directories (EPADs). Thus, any call would automatically be redirected to the appropriate EPAD. The router would then map the location information provided by the proxy to an emergency provider. The EPAD could either route the call, acting as a SIP proxy, or simply provide the SIP URI or telephone number to the receiving proxy, acting as a SIP redirect server.

For coverage redundancy, multiple EPADs could register. Normal SIP forking rules would ensure

## Emergency Communications

During emergencies, telecommunications facilities are often strained by both official and private communications: Rescue workers and law enforcement must coordinate activities, while ordinary citizens need information on the whereabouts and health of friends and relatives. To replicate existing functions, we must modify two Internet architecture layers — the IP layer and the signaling layer.[1]

### IP Layer Changes

At the IP layer, differentiated services already offer a mechanism to give better service to certain users. The main problem is authenticating access for users. Because it's impractical to add authentication information to each packet header, we need some form of boundary filtering and admission control.

For "I'm alive" notifications, it might make sense to give each device a set of tokens for elevated-priority packets, thus encouraging frugal notification options such as e-mail, instant messaging, or brief calls.

### Signaling Layer Changes

With signaling priority, there are two basic issues: accessing the existing PSTN and prioritizing resources in SIP proxy servers. Currently, military and civilian emergency networks offer multilayer preemption priority for accessing the existing PSTN. For example, the U.S. defense network defines levels ranging from "routine" to "critic-ecp" (emergency call processing). We must make similar functionality available to IP-based systems. We have proposed that a simple SIP header field indicate the desired resource access priority, addressing priority handling in proxies and gateways.[2] The same mechanism could be used for authenticated e-mail and HTTP, although its usefulness is less certain for the latter.

### References

1. F. Baker, "IEPS Requirement Statement," Internet Draft, Internet Engineering Task Force, Nov. 2001. Work in progress.
2. J. Polk and H. Schulzrinne, "SIP Communications Resource Priority Header," Internet Draft, IETF, Nov. 2001. Work in progress.

that the system contacted servers in some order, trying all servers until it located one that worked. We don't yet know how EPADs would find the proxy. While we can imagine various service-location mechanisms, simple, manual configuration might suffice because the EPADs are not likely to change frequently. The SIP-proxy owner would obtain EPAD addresses and provide EPADs with a secret to use for registration. This would prevent rogue EPADs from registering and limit EPAD registration invocation to domain owners.

Because the call itself need not traverse the proxies, any proxy could intercept the `sos` call. Thus, even if the SIP phone were misconfigured to use the owner's domain as its first-hop proxy, the system would route the call to the appropriate PSAP. The same functionality is used in telematics applications, in which cars are equipped with automatic dialers that contact an emergency call center operated by a private service provider.

In our proposed system, users could also call a PSAP directly, using a universal URL such as `sip:sos.us` that reaches a country's main PSAP. Once the head PSAP got the emergency call and determined the caller's location, it would forward the call to the PSAP nearest to the caller. This would be feasible because there would be a limited number of PSAPs in any country.

### Using other VoIP Protocols

The basic mechanisms we describe here can apply to signaling mechanisms other than SIP, including H.323[9] and Megaco (http://www.ietf.org/html.charters/megaco-charter.html). In the Megaco architecture, a media gateway controller drives one or more gateways, such as public switched telephone network (PSTN) gateways or desk phones. The gateway controllers are typically connected using SIP.

If the connection is through H.323, each zone has a gatekeeper that routes calls from local terminals. This gatekeeper intercepts emergency calls and forwards them to the appropriate location. Instead of SIP's `REGISTER`, the EPAD could register with the gatekeeper using the H.225.0 registration, admission, status (RAS) protocol.

### Benefits

Although the existing analog PSTN is likely to be around for several decades, it nonetheless makes sense to run IP-enabled PSAPs — not least because 3G wireless systems will use IP for voice communications. Transitioning to IP-enabled PSAPs adds many capabilities and permits a much richer communications environment. For example, video could help emergency operators assess emergency situations, monitor and instruct callers in first-aid efforts, and communicate with people speaking sign language. IP-based systems can easily provide text-based messaging, which is currently only available via specialized Telecommunications Device for the Deaf (TDD) equipment, and such equipment is not widely available in offices or public places. IP-based communications could also accommodate biometric data, such as from patients who are medically monitored from their homes.

In addition, current PSAPs require a highly specialized infrastructure. An IP-enabled PSAP requires only a PC and network connectivity, which makes it easy to move operations if, for example, the primary location is affected by natural disasters. Indeed, emergency operators would no longer have to be centralized. They could work from multiple locations, including their homes, assuming they had a DSL, cable modem, or other high-speed connection.

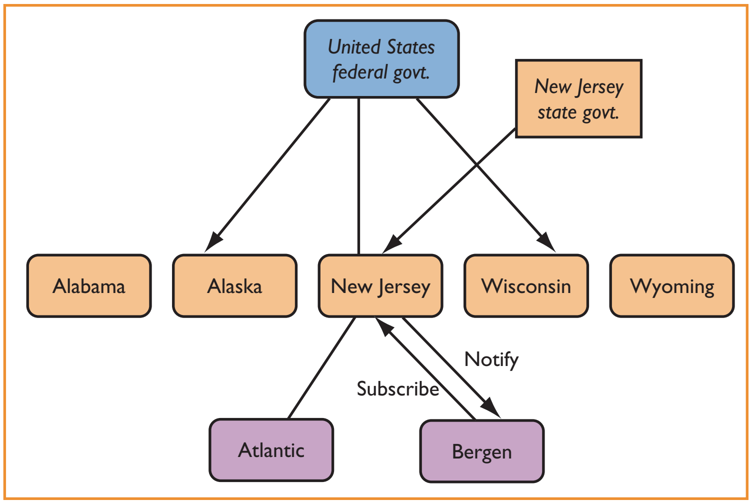

*Figure 2. SIP-based notification system's alerting hierarchy. Information flows from national government to state and local governments, and then to citizens. Users subscribe to servers at the next higher level and act as servers for lower levels.*

## Emergency Notification: The Current System

Emergency notification systems let government officials notify citizens and government agencies of emergency situations. For example, the U.S. Emergency Broadcast System (EBS) was developed in 1963 to notify citizens of emergency situations and the precautionary measures they should take.

In December 1995, the U.S. Federal Communications Commission began replacing EBS, which is limited to use by the president, with the Emergency Alert System, which state and local authorities can

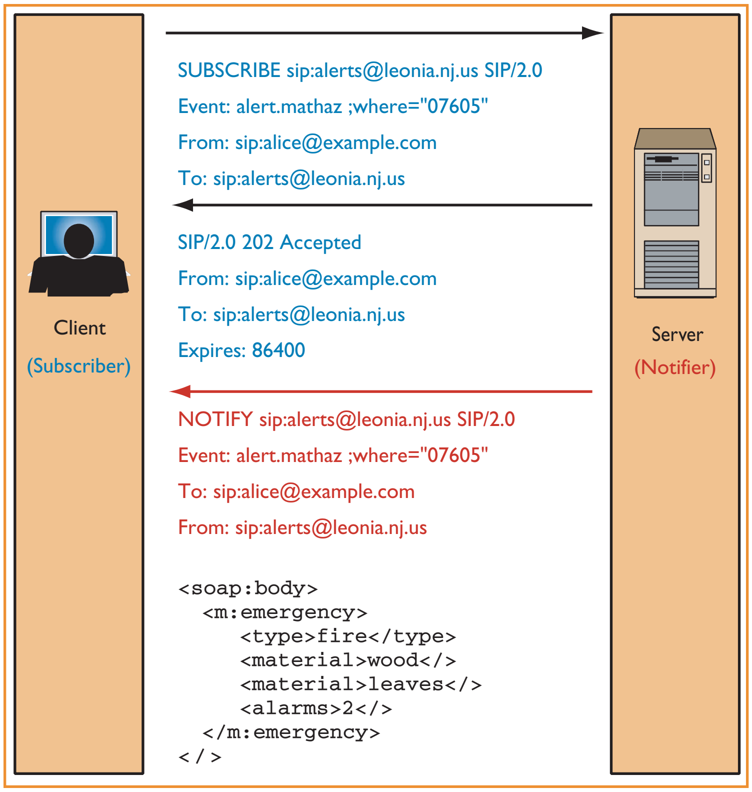

```
SUBSCRIBE sip:alerts@leonia.nj.us SIP/2.0
Event: alert.mathaz ;where="07605"
From: sip:alice@example.com
To: sip:alerts@leonia.nj.us
```

```
SIP/2.0 202 Accepted
From: sip:alice@example.com
To: sip:alerts@leonia.nj.us
Expires: 86400
```

```
NOTIFY sip:alerts@leonia.nj.us SIP/2.0
Event: alert.mathaz ;where="07605"
To: sip:alice@example.com
From: sip:alerts@leonia.nj.us

<soap:body>
  <m:emergency>
     <type>fire</type>
     <material>wood</>
     <material>leaves</>
     <alarms>2</>
  </m:emergency>
</>
```

*Figure 3. Protocol exchanges for subscriber event alerting notifications. The notification message body describes the emergency's nature using the XML-RPC schema.*

also use (for details, see www.fema.gov/pte/rep/easrep.htm). EAS distributes information across AM and FM radio stations and television stations. Each station must hear EAS alerts on at least two other stations before automatically rebroadcasting them for its local area. An emergency announcement consists of an alert tone, a frequency-shift-keying-encoded digital data stream of about eight seconds, an audio message, and an end-of-message indicator. The data stream contains information about the warning type (such as hurricane or civil disturbance), the county or part of a county it applies to, the date and time issued, and the issuing authority. The format is similar to National Weather Service weather alerts.

In addition to radio and TV EAS alerts, several emergency notification networks and products offer emergency alerts to local areas (for examples, see www.can-intl.com/ and www.inwireless.com/). Old systems used sirens, but provided minimal information content; basically, they got people to turn on their TVs or radios. Some community alert systems use loudspeakers, while others rely on telephone circuits. For example, Reverse 911 (http://www.reverse911.com) dials telephone and fax numbers from a list or within a specific geographic area. Alert systems can also be useful in private enterprises, such as chemical plants, to inform personnel when problems occur.

Unfortunately, current PSTN-based emergency notification systems are limited in scale, relatively slow, and provide only basic information. Using event-notification protocols — in this case, SIP — can both increase the scale and speed of such systems, and expand their functionality.

## Toward a SIP-Based Notification Architecture

We propose to enhance EAS and community alert systems with a SIP-based event notification system. In principle, any network-based event notification system could be used, but because end users will likely already have SIP-based notification capabilities on devices like 3G wireless handsets and PC desktops, it makes sense to use this "commodity" technology rather than invent a new one specific to emergency alerts. (Microsoft XP, for example, already includes a SIP-capable instant messaging client.)

Our basic architecture is straightforward. As Figure 2 shows, we envision a hierarchical subscription system, in which national governments disseminate information to state and local governments, and then to citizens, with information generated at any level. Users subscribe to servers at the next higher level, and might in turn become servers for the levels below. Cross subscriptions between neighboring states or countries are also necessary for notifications among peers.

As with EAS's multiple-source mechanism, lower levels of our architecture's alerting hierarchy subscribe to multiple upper levels, which increases robustness. For general alerts, each level subscribes to events generated by its children, so that events can be propagated up the hierarchy as well. Subscriber location is thus less critical and alerts are permitted across the civil hierarchy.

For example, a local police department in Alabama could generate a fugitive alert to New York authorities if they believed the person had boarded a plane bound for New York. Such a system would also prevent notification servers from having to maintain updated contact lists for law enforcement agencies and emergency response units. Instead, addressing would be determined by emergency type and geography, where each destination would determine their coverage areas and expertise.

To sign up for alerts, the client would send a `SUBSCRIBE` request to the appropriate server (see Figure 3). The request would contain the event description (Event header) and the notifications' network destination, as well as any authentication information. Upon approval, the server would add the subscriber to the appropriate event list and generate NOTIFY requests to the subscriber when an alert occurs. Subscriptions would time-out to prevent wasting network resources on users who are no longer interested in alerts or devices that are no longer capable of receiving them. The Expires header indicates the subscription's duration. To maintain a subscription, users would simply update them periodically; to end notifications, they could set the expiration time to zero.

Users interested in various events could submit multiple SUBSCRIBE requests. We might extend the subscription mechanism to include a geographic range and limit the number of notifications. Subscribers could indicate a preference for media type, such as audio or text notifications. The SUBSCRIBE request might also contain a standardized message body where subscribers could further describe their capabilities.

The system would send the emergency notification to the subscribing address, which might be a specific host (identified by a host name or IP address) or a more generic user@domain address. SIP supports request routing, wherein intermediaries (SIP proxies) can rewrite the destination address and forward the request. Thus, subscribers need not forward address changes to the source, which offers some privacy and improves system scalability.

The notification itself might contain a message body that further describes the emergency's nature in a machine-parsable way. For example, a forest fire notification to emergency personnel might detail projected fire movements, evacuation instructions, and similar information. This could then be rendered appropriately and perhaps even integrated into, say, a geographic information system. To accomplish this, we propose using the XML-RPC schema. As used for the simple object access protocol (SOAP), The XML-RPC schema offers the necessary data abstraction functionality and would permit reuse of existing emergency alert implementations.

## Finding Servers

How subscribers might locate the appropriate server is an example of the wide-area service-location problem.[10,11] Although a general solution is likely the most appropriate, lacking that, we envision various ad hoc solutions for different situations.

For citizen subscriptions, the simplest solution is to advertise the subscription address out-of-band — through Web pages, newspaper advertisements, and so on. Also, the service might have a well-known address, similar to the current 911 number. The SIP redirect server at that address would not provide notifications directly, but rather would redirect subscription requests to the appropriate server based on geographic location (say, by postal code) and the event type specified. Local agencies would then register with the central server, using the normal SIP REGISTER binding mechanism illustrated in Figure 4.

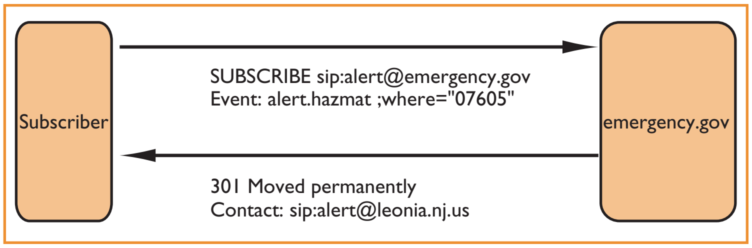
Subscriber
emergency.gov
SUBSCRIBE sip:alert@emergency.gov
Event: alert.hazmat ;where="07605"
301 Moved permanently
Contact: sip:alert@leonia.nj.us

*Figure 4. The SIP event subscription mechanism. Local agencies would use this mechanism to register with central servers, which would redirect subscription requests to local servers based on geography and event type.*

Government agencies presumably have information distribution arrangements that can provide authentication credentials and logical server addresses (such as sip:tornado@nws.noaa.gov). Because the subscription address is not subject to interruptions, such as area code changes or agencies' physical moves, it is likely to remain constant for many years.

## Authentication and Authorization

Authentication and authorization are vital for an emergency alert system, for both subscriptions and notifications. Subscriptions must be authenticated for distributing events to government officials, but authentication is also useful to prevent a single citizen from subscribing multiple times or, worse, accidentally or intentionally redirecting someone else's subscription. (You can do this by spoofing the SIP From header and inserting your own address as the Contact value.)

Different approaches will likely be required to authenticate citizens and officials. The relatively small number of emergency response agencies are likely to have mechanisms already in place for securely distributing information. To prevent creating panic with bogus alerts, the emergency alert server must prove its identity when issuing notifications.

To authenticate requests, SIP currently uses either HTTP digest or transport and network-layer security. For digest security, a server challenges the client within an error response; the client then reissues the original request, encrypting the challenge with a shared secret. This scheme protects user passwords from those listening in.

A relatively simple authentication mechanism suffices for authenticating residents. They might, for example, sign up on a Web page and receive an authentication key by e-mail. Although e-mail delivery is hardly secure, it should be good enough to prevent random users from creating nuisance subscriptions or "stealing" another user's subscription. Any secure mechanism can distribute shared secrets to authenticate government officials. Alternatively, cryptographic message syntax (CMS)-based encryption[12] can provide both authentication and confidentiality using public-key cryptography.

### Benefits

Once the authentication and notification components are implemented in a standardized manner, our SIP-based notification system would provide several benefits over existing emergency response systems.

- *Device neutrality*. The system can migrate to new devices, including IP telephones, 3G wireless handsets, and embedded devices, without being explicitly extended to handle them.
- *More information, specifically targeted*. EAS provides limited information that's hard to extend without upgrading end systems. SIP event notifications can carry detailed information tailored to different needs — ranging from alerts issued in multiple languages to those targeted to a small population during localized emergencies. We could also embed the system with an RPC-like mechanism so that the alert could trigger appropriate action in automated systems.
- *Stronger authentication*. The existing authentication mechanism relies on manual codebooks and the difficulty of spoofing an over-the-air signal. However, it would be relatively easy to drive past an EAS receiver with a small transmitter and distribute a false alarm. Our mechanism can use true cryptographic authentication, which is more amenable to automated processing and less likely to be spoofed.
- *Lower resource consumption*. A one-minute alert call consumes about 480 Kbytes (one way), while an alert notification is at most a few hundred bytes long. Our system can thus use the same amount of bandwidth to reach 1,000 times more people in the same time period. It can also leverage Web hosting and similar facilities with abundant bandwidth.
- *Integration with current systems*. Feeding EAS and the emergency digital information system into the SIP emergency alert system would be a straightforward process. Combining the systems would allow officials to reach more people, such as those on computers who are not listening to the radio or watching TV. Moreover, because we can narrowly tailor the system's reach, it can easily integrate less urgent alerts, such as traffic accidents or other police activity.
- *Out-of-area notification*. Current notification systems assume that only those in close physical proximity of the emergency event need to know about it, but that is not always the case. People might need to be informed when they are traveling, for example, if their homes are threatened in some way.

Thus, SIP-based emergency notification addresses many of the shortcomings of existing systems and offers a foundation for future automated alert routing and handling.

## Conclusion

Citizens have come to expect emergency-related services from the telephone system, and IP telephony and the associated protocol infrastructure must provide at least the existing service levels. Rather than simply replicating the existing system using packets, however, we have the opportunity to create more functional, robust, and flexible systems that can enhance existing capabilities.

In the future, we hope to build emergency calling and notification software and integrate it into our SIP user agent. We are currently implementing a prototype version of the architecture described here, integrating alerting functionality into our SIP user agent and emergency calling functionality into our SIP proxy server. Scalable authentication remains one of the challenges. If the frequency of SIP-based event notifications becomes large, we might need semantic filtering, based on message content.

We are also considering other research areas, including increasing the security and authentication process of SIP subscribe and notify messages, finding a viable way of locating IP devices, and designing a wide-scale, platform-independent event notification system.

### References

1. J. Rosenberg and H. Schulzrinne, "The IETF Internet Telephony Architecture and Protocols," *IEEE Network*, vol. 13, no. 3, May/June 1999, pp. 18-23.
2. M. Handley et al., "SIP: Session Initiation Protocol," Inter-

net Engineering Task Force RFC 2543, Mar. 1999; available at http://www.rfc-editor.org/rfc/rfc2543.txt.

3. A. Vaha-Sipila, "URLs for Telephone Calls," IETF RFC 2806, Apr. 2000; available at http://www.rfc-editor.org/rfc/rfc2806.txt.
4. A. Roach, "SIP-Specific Event Notification," Internet Draft, IETF, Nov. 2001. Work in progress.
5. J. Rosenberg et al., "SIP Extensions for Instant Messaging," Internet Draft, IETF, July 2001. Work in progress.
6. S. Tekinay, "Wireless Geolocation Systems and Services," special issue, *IEEE Comm. Magazine*, vol. 36, no. 4, Apr. 1998; available at www.comsoc.org/~ci/public/1998/apr/guested1.html.
7. "Emergency Responder Version 1.1," data sheet, Dec. 2001, Cisco Systems, San Jose, Calif.; available at www.cisco.com/univercd/cc/td/doc/pcat/erv.htm.
8. H. Schulzrinne, "Universal Emergency Address for SIP-Based Internet Telephony," Internet Draft, IETF, July 2001. Work in progress.
9. J. Toga and J. Ott, "ITU-T Standardization Activities for Interactive Multimedia Communications on Packet-Based Networks: H.323 and Related Recommendations," *Computer Networks and ISDN Systems*, vol. 31, Feb. 1999, pp. 205-223.
10. P. Castro et al., "Locating Application Data across Service Discovery Domain," *ACM/IEEE Int'l Conf. Mobile Computing and Networking (MobiCom)*, ACM Press, New York, Aug. 2001, pp. 28-42.
11. J. Rosenberg and H. Schulzrinne, "Internet Telephony Gateway Location," *Proc. Conf. Computer Comm. (IEEE Infocom)*, IEEE CS Press, Los Alamitos, Calif., 1998, pp. 488-496.
12. R. Housley, "Cryptographic Message Syntax," IETF RFC 2630, June 1999; available at www.rfc-editor.org/rfc/rfc2630.txt.

**Henning Schulzrinne** is an associate professor in the department of electrical engineering and the department of computer science at Columbia University in New York. He received a PhD in electrical engineering from the University of Massachusetts. His research interests include Internet multimedia and telephony services, signaling, network quality of service, scheduling, multicast, and performance evaluation. Schulzrinne is coauthor of the Internet standards-track protocols RTP, RTSP, and SIP.

**Knarig Arabshian** is a PhD candidate in the department of computer science at Columbia University. Her research focuses on medical applications of Internet multimedia and wide-scale Internet event notification systems.

Readers can contact the authors at {hgs,knarig}@cs.columbia.edu.

IEEE Computer Society Certified Software Development Professional

csdp

# Get CSDP Certified

Announcing IEEE Computer Society's new

## Certified Software Development Professional Program

## Doing Software Right

- Demonstrate your level of ability in relation to your peers
- Measure your professional knowledge and competence

The CSDP Program differentiates between you and others in a field that has every kind of credential, but only one that was developed by, for, and with software engineering professionals.

*"The exam is valuable to me for two reasons:*
*One, it validates my knowledge in various areas of expertise within the software field, without regard to specific knowledge of tools or commercial products...*
*Two, my participation, along with others, in the exam and in continuing education sends a message that software development is a professional pursuit requiring advanced education and/or experience, and all the other requirements the IEEE Computer Society has established. I also believe in living by the Software Engineering code of ethics endorsed by the Computer Society. All of this will help to improve the overall quality of the products and services we provide to our customers..."*
**— Karen Thurston, Base Two Solutions**

**Register Today**

Visit the CSDP web site at http://computer.org/certification
or contact certification@computer.org

IEEE Computer Society

IEEE