**Requirements for Session Initiation Protocol (SIP)-based Emergency
Calls
draft-schulzrinne-sipping-emergency-req-01**

Status of this Memo

Copyright Notice

Abstract

This document enumerates requirements for emergency calls in VoIP and
general Internet multimedia systems. We divide the requirements into
"trunk replacement" and "end-to-end". Trunking solutions only
exchange the emergency call center's circuit-switched access by an
IP-based system. The requirements for end-to-end IP-based emergency
calling address functional and security issues for determining the
correct emergency address, for identifying the appropriate emergency

call center and for identifying the caller and its location. While
we focus on systems that employ the Session Initiation Protocol
(SIP), many of the requirements may also apply to other environments,
such as those using H.248/Megaco, MGCP or H.323.

Table of Contents

[1](). Introduction

Users of telephone-like services expect to be able to call for
emergency help, such as police, the fire department or an ambulance,
regardless of where they are, what (if any) service provider they are
using and what kind of device they are using. Unfortunately, the
mechanisms for emergency calls that have evolved in the public
circuit-switched telephone network (PSTN) are not quite appropriate
for evolving IP-based voice and real-time multimedia communications.
This document outlines some of the requirements that end systems and
network elements such as SIP proxies need to satisfy in order to
provide emergency call services that offer at least the same
functionality as existing PSTN services, while hopefully making
emergency calling more robust, cheaper to implement and
multimedia-capable.

In the future, users of other real-time and near real-time services
may also expect to be able to summon emergency help. For example,
instant messaging (IM) users may want to use such services. IM is
particularly helpful for hearing-disabled users ([RFC 3351]() [[2]()]) and in
cases where bandwidth is scarce. For lack of a better term, we will
use the term "caller" or "emergency caller" to refer to the person
placing an emergency call or sending an emergency IM.

Emergency callers and ECCs expect calls to be completed reliably.
Where possible, a callback number and the current caller location
shouls be delivered to the ECC to speed up emergency response and to
limit prank calls.

The emergency calls described in this document differ from the
emergency telecommunications service (ETS) described in XXX. In ETS,
relatively small numbers of emergency workers need to maintain
communication even when parts of the infrastructure are destroyed or
disabled. Emergency calls, on the other hand, are placed by
civilians to call for emergency services such as fire, ambulance and
police services. Thus, these two services are complementary.

We distinguish two sets of requirements, one for ECC trunk
replacement use of SIP (Section XXX), where VoIP emergency callers
still use the existing PSTN, and end-to-end SIP emergency calls
(Section XXX) that terminate SIP-originated emergency calls without
transitioning the PSTN.

There is a third approach, where SIP-originated calls terminate on a
PSTN gateway in each emergency calling area. This architecture is
left for future consideration and discussed in other standardization
organizations, such as NENA, as it is strongly dependent on the
currently-deployed emergency services network architecture.

2. Terminology

In this document, the key words "MUST", "MUSTNOT", "REQUIRED",
"SHALL", "SHALLNOT", "SHOULD", "SHOULDNOT", "RECOMMENDED", "MAY", and
"OPTIONAL" are to be interpreted as described in RFC 2119 [1] and
indicate requirement levels for compliant implementations.

Since a requirements document does not directly specify an
implementable protocols, these compliance labels should be read as
indicating requirements for the protocol or architecture, rather than
an implementation.

[3](#). **Definitions**

Emergency call center (ECC): An emergency call center (ECC) receives
emergency calls within a specific geographic area and dispatches
emergency services, such as fire, police and rescue services. An
ECC may also serve as a backup for another ECC and, in backup
mode, dispatch emergency services outside of its normal service
region. In the United States and Canada, ECCs are called Public
Safety Answering Points (PSAPs).

Internet Protocol ECC (IECC): An Internet protocol emergency call
center (IECC) is an ECC that uses Internet protocols, such as SIP
for call signaling, RTP for media delivery, to receive emergency
calls.

Call taker: A call taker is an agent, typically a government
employee, at the ECC that accepts calls and may dispatch emergency
help. (Sometimes the functions of call taking and dispatching are
handled by different groups of people, but these divisions of
labor are not generally visible to the outside and thus do not
concern us here.)

Basic emergency service: Basic emergency service allows a user to
reach an ECC serving its current location, but the ECC may not be
able to determine the identity or geographic location of the
caller (except by having the call taker ask the caller).

Enhanced emergency service: Enhanced emergency services add the
ability to identify the caller identity and/or caller location to
basic emergency services. (Sometimes, only the caller location
may be known, e.g., from a public access point that is not owned
by an individual.)

Trunk replacement: In the trunk replacement architecture, the caller
uses the existing PSTN infrastructure to place an emergency call.
Only the path from the "selective router", or the equivalent
functionality outside North America, to the ECC uses IP-based
communications. The call may well be placed from a VoIP device,
but is assumed to enter the PSTN very close to the location of the
caller. The use of Internet protocols is invisible to the caller.

End-to-end emergency service: In end-to-end emergency service, the
caller and ECC both use Internet protocols end-to-end.

Selective router: A selective router or enhanced emergency call
control office. The enhanced emergency call control office is
"[t]he Central Office that provides the tandem switching of 9-1-1
calls. It controls delivery of the voice call with ANI to the
PSAP and provides Selective Routing, Speed Calling, Selective
Transfer, Fixed Transfer, and certain maintenance functions for
each PSAP. Also known as 9-1-1 Selective Routing Tandem or
Selective Router." (NENA Glossary) The term may be specific to
North America. (TBD: Find out if there are other terms.)

[4](#). **Trunk Replacement**

   In the trunk replacement architecture, an ECC replaces an analog
   (CAMA) or digital (ISDN) trunk with packet-based access, typically
   over one or more high-speed access lines such as DSL or leased lines.
   The packet-based access terminates in the "selective router" that
   normally hands off calls to the ECC. Thus, the ECC becomes an EICC,
   but no larger scale infrastructure changes are required. To amplify,
   in the trunk-replacement model, a SIP user agent calling for
   emergency assistance can NOT dial reach the ECC directly via a SIP
   session; rather, the SIP session terminates on a PSTN gateway,
   traverses the PSTN as in today's circuit-switched environment and is
   only converted to VoIP at the selective router handling the ECC.

   Motivation: Trunk replacement is motivated by cost and call setup
   considerations. It may be cheaper to use IP-based technology for the
   access link and ECC-internal communications. Also, many existing
   (US) PSAPs use analog technology (CAMA trunks), to receive emergency
   calls. These trunks, originally designed for operator positions, can
   pulse out the ten or 20-digit (for wireless) caller's number, but as
   dialed digits. Thus, they add several seconds of call setup delay.
   This can be particularly disconcerting since it affects the time
   until the call taker can pick up the call. IP-based communications,
   using, for example, SIP as a call signaling protocol, can effectively
   eliminate this extra caller identification delay. (Additional delays
   are caused by the often very low speed access to the mapping database
   that maps caller identity to geographic location.) Finally, since
   pending calls do not consume access network resources, such systems
   may be more robust in the face of overload.

   M1: Coexistence: Due to the investment required, not all ECCs will
      convert to IP-based access at the same time. Thus, emergency
      calls MUST work in a network where some ECCs use existing (analog)
      technology, some ISDN, others IP. In particular, existing back-up
      relationships between ECCs must continue to work.
   M2: Call setup delay: The call setup delay MUST NOT be no larger
      than for existing analog trunks and SHOULD be significantly
      smaller. Call setup times of two seconds or less are RECOMMENDED.
   M3: Call identification: Signaling from the PSTN switch must be able
      to convey both ten and 20-digit caller identities (ANI --
      automatic number identification) used in North America and other
      digit strings used elsewhere.
   M4: Call transfer: Call takers MUST be able to transfer active
      sessions to other call takers within the same ECC and to other
      ECCs, even those not using Internet.

M5: Simultaneous alerting: A given set of call takers must be alerted
   to any incoming emergency call.

M6: Call routing: The call may be awarded to the first call taker to
   answer or it may be routed to call agents based on policies, such
   as least-busy. Agents must be able to be assembled into multiple
   groups according to policies specified by ECC authorities. These
   groupings must be changeable by the ECC authority [4].

M7: Call queueing: It must be possible to queue calls, either in
   answered or unanswered state. Queued calls must be able to
   receive recorded announcements. ECC personnel, as directed by
   policy, should be able to modify the announcements. The call
   queue should allow automatic or manual transfer to another
   location of calls that exceed a particular expected waiting time
   [4].

M8: Call identification: The call taker MUST be able to distinguish
   the following incoming call types [4]:
   *   emergency calls dialed via a univeral emergency number;
   *   direct-dialed emergency calls;
   *   transfers from other ECCs;
   *   anonymous calls;
   *   administrative calls;
   *   call origination (wireline, wireless, telecommunication devices
       for the deaf (TDD));
   *   default-routed calls (These are calls for which selective
       routing information was unavailable, resulting in the call
       being routed to a "default" ECC based on other criteria.)

M9: Information delivery: The call setup request MUST be able to
   deliver the following information [4]:
   *   called party number (to identity ECC or type of call);
   *   calling party number, including any numbering plan digits;
   *   delivery of indication of caller ID blocking for non-emergency
       calls;
   *   location information or lookup keys;
   *   ANI on abandoned calls;
   *   indication that a terminating emergency call has been alternate
       routed from another PSAP.

M10: Agent sign-on: Agents must be able to log on and log off;
   workstations conditions should at least include "ready", "not
   ready" and "busy" [4].

M11 Conferencing: Occasionally, supervisors, translators or other
   specialists need to participate in an emergency call. Thus, it
   MUST be possible to add one or more parties, not necessarily
   located in the IECC, to any emergency call at any time.

M12: Announcements: Callers may receive automated announcements or
   other indications of call status [4].

M13: Call queues: Supervisors MUST be able to manage call queues.

M14: Call metrics: Supervisors and/or agents can measure call delays and other performance metrics [4].

M15: Monitoring and recording: In many jurisdictions, both sides of all emergency calls are automatically recorded as potential legal evidence. Thus, it MUST be possible to record and timestamp all signaling and media from all successful, queued, failed and aborted calls.

M16: Abandoned calls: ECCs need to be notified of abandoned calls, i.e., emergency calls that are dropped by the caller before being answered by a call taker.

M17: Transition to end-to-end: Protocols and architecture SHOULD be chosen so that a trunk-replacement IECC can receive emergency calls placed by IP endpoints without major system changes or hardware upgrades.

M18: Authentication of incoming calls: The IECC MUST be able to ascertain that the calls it receives are indeed originating from the selective router.

M19: Authentication of the IECC: The selective router MUST be able to be assured that the calls it places reach the desired IECC rather than an impostor.

M20: Confidentiality: Implementations MUST support confidentiality for call signaling and media streams, to protect them against unauthorized disclosure to third parties.

M21: Robustness: An IECC SHOULD be able to automatically route all incoming calls to another backup IECC, even if the access link(s) to the primary IECC are inoperative. Any such redirection MUST be authenticated.

M22: Overflow handling: An IECC SHOULD be able to automatically route calls to another IECC if the (expected) waiting time exceeds a configured threshold.

M23: Hold: The call taker MUST be able to place the a call in a status that allows him/her to handle other calls without disconnecting from the caller. A visual/audible notification should be available for the call taker to alert them that a call is on hold. The call should continue to be recorded and an optional voice message should be made available for the caller so they are aware of the status of their call [4].

M24: Forced disconnect of caller: The "forced disconnect of caller" feature allows the ECC call taker to disconnect a call when the call is in an off hook status at the calling parties end. This eliminates the possibility that emergency resources are needlessly tied up by emergency calls made and then left off hook [4].

M25: Called party hold: This feature allows a call taker to continue to stay connected to the calling party even if the calling party attempts to place their phone in an on-hook status [4].

M26: Caller ring back: Caller ring-back allows the call taker to be
      able to ring a phone back even if the destination phone is in an
      off-hook status [4].

5. End-to-End IP-Based Emergency Calls

   End-to-end emergency calls originate on an Internet device, traverse
   IP networks and terminate on an IP-capable ECC (IECC).

   As noted, emergency calls need to be identified as such Section 5.1
   and be routed to the appropriate emergency call center (see Section
   5.3). The ECC needs to determine who (Section 5.4) placed the call
   from where (Section 5.2). Emergency calls may not be subject to
   access restrictions placed on non-emergency calls. Also, some call
   features may interfere with emergency calls, particularly if triggerd
   accidentally (Section 5.5).

5.1 Emergency Address

   The emergency address is used by the emergency caller to declare a
   call to be an emergency call and to guide the call to an ECC. The
   emergency address could a be "sip", "sips" or "tel" URI, or some
   other, yet-to-be-defined URI scheme.

   A1: Universal: Each device and all network elements MUST recognize
        one or more global emergency call identifiers, regardless of the
        location of the device, the service provider used (if any) or
        other factors.
        Motivation: SIP and other call signaling protocols are not
        specific to one country or service provider and devices are likely
        to be used across national or service provider boundaries. Since
        services such as disabling mandatory authentication for emergency
        calls (S1) requires the cooperation of outbound proxies, the
        outbound proxy has to be able to recognize the emergency address
        and be assured that it will be routed as an emergency call. Thus,
        a simple declaration on a random URI that it is an emergency call
        will likely lead to fraud and possibly attacks on the network
        infrastructure. A universal address also makes it possible to
        create user interface elements that are correctly configured
        without user intervention. UA features could be made to work
        without such an identifier, but the user interface would then have
        to provide an unambiguous way to declare a particular call an
        emergency call.
   A2: Local: Since many countries have already deployed national
        emergency numbers, such as 911 in North America and 112 in large
        parts of Europe, UAs, proxies and call routers MUST recognize
        local emergency numbers. In addition, they SHOULD recognize
        emergency numbers that are found elsewhere.
        Motivation: The latter requirement is meant to help travelers
        that may not know the local emergency number and instinctively
        dial the number they are used to from home. However, it is
        unlikely that all systems could be programmed to recognize any

emergency number used anywhere as some of these numbers are used
for non-emergency purposes, in particular extensions and service
numbers.

A3: Recognizable: Emergency calls MUST be recognizable by user
agents, proxies and other network elements. To prevent fraud, an
address identified as an emergency number for call features or
authentication override MUST also cause routing to an ECC.

A5: Minimal configuration: Any local emergency numbers SHOULD be
configured automatically, without user intervention.
Motivation: A new UA "unofficially imported" into an organization
from elsewhere should have the same emergency capabilities as one
officially installed.

A6: Secure configuration: Devices SHOULD be assured of the
correctness of the local emergency numbers that are automatically
configured.
Motivation: If we assume a fixed, global emergency service
identifier that requires no configuration and only configure local
"traditional" emergency numbers, users are not likely to suddenly
dial some random number if a rogue configuration server introduces
this as an additional emergency number. The ability to override
all locally configured emergency number is of more concern.

5.2 Identifying the Caller Location

This section supplements the requirements outlined in RFC 3693 [3].
Thus, the requirements enumerated there are not repeated here. In
general, we can distinguish two modes of operation: direct and
indirect location provision. In direct location provision, the
calling end system knows its own location and can convey this
location to the ECC. In an indirect system, the caller is identified
by a permanent or temporary identifier, which the ECC then uses to
map the caller to a current location. (In the current North American
enhanced emergency calling system, the landline terminal phone number
is mapped to a location using the so-called ALI (Automatic Location
Identification) database. For wireless phones, a temporary
identifier is created and then mapped to the location information.)

(This is somewhat similar to terminal-based and network-based
location services in wireless emergency calling services. However,
even in direct location provision, the terminal may well acquire the
location information from a third party, e.g., a wireless location
beacon or a DHCP server.)

L1: Multiple location providers: For indirect locations, ECCs MUST
be able to access different location providers. The location
provider may be tied to the service provider or may be independent
of the service provider.

        Motivation: This requirement avoids that all users have to rely
        on a single location provider. This requirement is hard to avoid
        if there are no traditional national application-layer service
        providers.
   L2: Civic and geographic: Where available, both civic (street
        address) and geographic (longitude/latitude) information SHOULD be
        provided to the ECC.
        Motivation: While geographic information can usually be
        translated into civic coordinates, some coordinates, such as
        building numbers and floors, are more easily provided as civic
        coordinates since they do not require a detailed surveying
        operation. For direct location determination, it may also be
        easier for the user to check civic coordinates for correctness.
   L3: Location source identification: Sources and translations of
        location data MUST be indicated to the ECC. (Motivation: This
        allows the ECC to better judge the reliability and accuracy of the
        data and track down problems.)

[5.3](#) Identifying the Appropriate Emergency Call Center

   From the previous section, we take the requirement of a single (or
   small number of) emergency addresses which are independent of the
   caller's location. However, since for reasons of robustness,
   jurisdiction and local knowledge, ECCs only serve a limited
   geographic region, having the call reach the correct ECC is crucial.
   While an ECC may be able to transfer an errant call, any such
   transfer is likely to add tens of seconds to call setup latency and
   is prone to errors. (In the United States, there are about 6,000
   PSAPs.)

   There appear to be two basic architectures for translating an
   emergency address into the correct IECC. We refer to these as
   caller-based and mediated. In caller-based resolution, the caller's
   UA consults a directory and determines the correct IECC based on its
   location. We assume that the UA can determine its own location,
   either by knowing it locally or asking some third party for it. A UA
   could conceivably store a complete list of all ECCs across the world,
   but that would require frequent synchronization with a master
   database as ECCs merge or jurisdictional boundaries change.

   For mediated resolution, a SIP (outbound) proxy or redirect server
   performs this function. Note that the latter case includes the
   architecture where the call is effectively routed to a copy of the
   database, rather than having some non-SIP protocol query the
   database. Since servers may be used as outbound proxy servers by
   clients that are not in the same geographic area as the proxy server,
   any proxy server has to be able to translate any caller location to
   the appropriate ECC. (A traveller may, for example, accidentally or

intentionally configure its home proxy server as its outbound proxy
server, even while far away from home.)

Note that the first proxy doing the translation may not be in the
same geographic area as the UA placing the emergency call.

The problem is harder than for traditional web or email services.
There, the originator knows which entity it wants to reach,
identified by the email address or HTTP URL. However, the emergency
caller only dialed an emergency address. Depending on the location,
any of several ten thousand ECCs around the world could be valid. In
addition, the caller probably does not care which specific ECC
answers the call, but rather that it be an accredited ECC, e.g., one
run by the local government authorities. (Many ECCs are run by
private entities. For example, universities and corporations with
large campuses often have their own emergency response centers.)

I1: Correct IECC: The system MUST reach the correct IECC, that is,
     an IECC that serves the location of the caller. In particular,
     the location determination should not be fooled by the location of
     IP telephony gateways or dial-in lines into a corporate LAN (and
     dispatch emergency help to the gateway or campus, rather than the
     caller), multi-site LANs and similar arrangements.
I2: Early routing: In mediated mode, the first proxy server along a
     request path MUST attempt to route the call to the appropriate
     IECC.
     Motivation: Proxy servers close to the caller can be expected to
     have better call routing knowledge, particularly if international
     boundaries are being crossed.
I3: Choice of IECCs: The system SHOULD offer the emergency caller a
     choice as to whether he wants to reach a local private emergency
     response center, e.g., on a corporate campus, or the
     government-run emergency call center responsible for his current
     location.
     Motivation: This choice is often, but not always, provided today.
     For example, in some cases, the local campus emergency center is
     reachable by a different number or 9-911 reaches the external ECC,
     while 911 reaches campus security.
I4: Assuring IECC identity: The emergency caller SHOULD be able to
     determine conclusively that he has reached an accredited emergency
     call center.
     Motivation: This requirement is meant to address the threat that
     a rogue, possibly criminal, entity pretends to accept emergency
     calls.
     Implementations SHOULD allow callers to proceed, with appropriate
     warnings or user confirmations, if the identity of the destination
     IECC cannot be verified.

Motivation: Verification can fail for any number of reasons, such
as lack of a common certificate chain, especially when traveling,
call forwarding, or the expiration of certificates.
Accreditation, e.g., in the case of corporate or university
campuses, may not exist.

I5: Traceable resolution: Particularly for mediated resolution, the
caller SHOULD be able to definitively and securely determine who
provided the resolution answer.

I6: Assuring directory identity: The querier (UA or server) MUST be
able to assure that it is querying the intended directory.

I7: Query response integrity: The querier MUST be able to be
confident that the query or response has not been tampered with.

I8: Assuring update integrity: Any update mechanism for the
directory MUST ensure that only authorized users can change
directory information. An audit trail MUST be provided.

I9: Call setup latency: The directory lookup SHOULD add minimal
delay to the call setup. Since outbound proxies will likely be
asked to resolve the same geographic coordinates repeatedly, a
suitable time-limited caching mechanism SHOULD be supported (see
also "Ix").

I10: Multiple directories: A UA or proxy SHOULD be able to use
multiple different directories to resolve the emergency address.
We do not assume that a single directory has worldwide or even
nationwide coverage.
(Motivation: This allows competing or regional data sources.)

I11: Referral: All directories SHOULD refer out-of-area queries to an
appropriate default or region-specific directory.
Motivation: This requirement alleviates the potential for
misconfigurations to cause calls to fail, particularly for
caller-based queries.

I12: Multiple protocols: It MAY be useful if directories support
multiple query protocols, such as SIP (for proxying), IRIS, LDAP,
a SOAP-based query and others. A mandatory-to-implement protocol
MUST be specified and an over-abundance of similarly-capable
choices appears undesirable.
(Motivation: It appears likely that the resolution mechanism will
be needed by a variety of session protocols and user
applications.)

I13: Robustness: The resolution mechanism MUST allow to deploy
systems that are robust in the face of partial network and
directory server failures. Caching MAY be used to mitigate
temporary unavailability of directories or network connectivity.

I14: Incrementally deployable: An Internet-based emergency call
system MUST be able to deployed incrementally. In the initial
stages of deployment, an emergency call may not reach the optimal
ECC.

I15: Testable: A user SHOULD be able to test whether a particular
      address reaches emergency help, without actually causing emergency
      help to be dispatched or consuming ECC call taker resources. Such
      tests MUST indicate the source of any problems, including the
      validity and plausibility of civic and geospatial location
      addresses.

5.4 Identifying the Caller

Enhanced emergency call systems provide the ECC with the identity and
location of the caller. In PSTN-based systems, the identity is
represented by the number of the terminal the call is placed from.
In a SIP-based system, we have two distinct identities, namely the
address of the terminal (SIP Contact header field) and the identity
(name and/or AOR) of the person using the terminal. Depending on the
circumstances, only one of them may be available. For example, from
a public terminal (Internet payphone), only the Contact address may
be useful.

In most jurisdictions, callers do not have a choice as to whether
they want to reveal their location or identity; such disclosure is
typically mandated by law.

C1: Identity: The system SHOULD allow (but not force) to identify
     both the caller's identity and his or her terminal network
     address.
C2: Privacy override: The end system MUST be able to automatically
     detect that a call is an emergency call so that it can override
     any privacy settings that conflict with emergency calling.
     (Whether this override can be configured by the user or is
     considered a condition of service is considered a legal matter,
     not a protocol issue.)
     Motivation: Since emergency calls are often placed by children,
     by people using somebody else's end system or by people in panic,
     any configuration should be automated rather than relying on user
     interaction at the time of the call. Delaying a call until the
     user discovers that they have to answer some screen prompt or deal
     with a voice prompt in an unfamiliar language is likely to lead to
     large call setup delays or call failures. This does not preclude
     that end systems can allow, on a call-by-call basis, to configure
     special call parameters.

5.5 Call Setup and Call Features

S1: Authentication override: In many jurisdictions, emergency calls
     can be placed by any device, regardless of whether it has
     subscribed for service. Similarly, outbound proxies and other
     call filtering elements MUST be able to be configured so that they

allow unauthenticated emergency calls.

S2: Mid-call features: The end system MUST be able to recognize an
emergency call and allow configuration so that certain call
features are not triggered accidentally. For example, it may be
inappropriate to transfer the ECC or put it on hold. An end
system MAY make it more difficult to disconnect an on-going
emergency call or accept other incoming calls while in an
emergency call.
Motivation: Call transfer initiated by the emergency caller is
likely only to be a problem if a PSTN gateway or B2BUa is in the
call path. It is not clear how much effort should be expended on
preventing intentional, as opposed to accidental, disconnection,
since callers can typically find physical-layer means to terminate
the call. This feature is not generally available in the PSTN.
For example, ANSI T1.628-2001 states that "E9-1-1 Call hold is an
optional network feature provided to a PSAP which prevents a
caller from disconnecting an ESC. .... However, there is no DSS1
or SS7 support for this capability at this time."

S3: Testable: Users SHOULD be able to test the ability to place an
emergency call without actually invoking an emergency response or
tying up emergency call take resources.
Motivation: This capability is unfortunately missing from the
current PSTN.

S4: Integrity: Implementations MUST provide mechanisms that ensure
the integrity of SIP protocol component that are crucial to
providing reliable emergency call service. (This requirement
implies authentication of the caller to allow integrity protection
of the request and authentication of the ECC to allow integrity
protection of responses.)

6. Security Considerations

   Confidentiality, integrity and authentication are core requirements
   for multiple aspects of emergency calling. Threats exist at the
   infrastructure and individual call level. Security threats are
   identified throughout this document.

   An adversary could corrupt call information or ECC resolution to
   cause emergency calls to fail subtly, without the caller necessarily
   noticing. This can be done on a call-by-call basis or by corrupting
   elements that perform the resolution, including the directory
   described in Section 5.3, Internet routing tables or DNS.
   (Obviously, there are typically other ways to make emergency calls
   fail completely, an approach phone-wire cutting burglars have
   practiced for years. However, the ability to spoof an ECC requires
   physical access to the PSTN cable plant, while this may not be
   required in the IP case.)

   Here, we do not consider attacks on the emergency call infrastructure
   itself. The techniques for dealing with such attacks are likely to
   be similar as those for protecting other network infrastructure,
   although the stakes may well be higher. In particular, layered
   defenses against denial-of-service attacks, including return
   routability checks, are likely to be part of the defensive arsenal.

7. Acknowledgments

   James Polk provided helpful comments on an earlier version of this
   document.

8. References

8.1 Normative References

[1]   Bradner, S., "Key words for use in RFCs to Indicate Requirement
      Levels", BCP 14, RFC 2119, March 1997.

8.2 Informative References

[2]   Charlton, N., Gasson, M., Gybels, G., Spanner, M. and A. van
      Wijk, "User Requirements for the Session Initiation Protocol
      (SIP) in Support of Deaf, Hard of Hearing and Speech-impaired
      Individuals", RFC 3351, August 2002.

[3]   Cuellar, J., Morris, J., Mulligan, D., Peterson, J. and J. Polk,
      "Geopriv Requirements", RFC 3693, February 2004.

[4]   National Emergency Number Assocation, "NENA technical
      information document on the interface between the E9-1-1 service
      providers network and the Internet protocol (IP) PSAP", NENA
      NENA-08-501, February 2003.

Author's Address

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

   Phone: +1 212 939 7004
   EMail: hgs+sip@cs.columbia.edu
   URI: http://www.cs.columbia.edu

Intellectual Property Statement

Disclaimer of Validity

Copyright Statement

Acknowledgment