TOC

ATOCA	H. Schulzrinne
Internet-Draft	Columbia University
Intended status: Informational	S. Norreys
Expires: March 28, 2011	BT Group
	B. Rosen
	NeuStar, Inc.
	H. Tschofenig
	Nokia Siemens Networks
	September 24, 2010

Requirements, Terminology and Framework for Exigent Communications draft-ietf-atoca-requirements-00.txt

Abstract

Before, during and after emergency situations various agencies need to provide information to a group of persons or to the public within a geographical area. While many aspects of such systems are specific to national or local jurisdictions, emergencies span such boundaries and notifications need to reach visitors from other jurisdictions.

This document provides terminology, requirements and an architectural description for protocols exchanging alerts between IP-based end points.

Status of this Memo

Copyright Notice

Table of Contents

1. Introduction

TOC

TOC

1.1. Classical Early Warning Situations

During large-scale emergencies, public safety authorities need to reliably communicate with citizens in the affected areas, to provide warnings, indicate whether citizens should evacuate and how, and to dispel misinformation. Accurate information can reduce the impact of such emergencies.

Traditionally, emergency alerting has used church bells, sirens, loudspeakers, radio and television to warn citizens and to provide information. However, techniques, such as sirens and bells, provide limited information content; loud speakers cover only very small areas and are often hard to understand, even for those not hearing impaired or fluent in the local language. Radio and television offer larger information volume, but are hard to target geographically and do not work well to address the “walking wounded” or other pedestrians. Both are not suitable for warnings, as many of those needing the information will not be listening or watching at any given time, particularly during work/school and sleep hours.

This problem has been illustrated by the London underground bombing on July 7, 2006, as described in a government report [July2005] (, ., “Report of the 7 July Review Committee, ISBN 1 85261 878 7,” June 2006.). The UK authorities could only use broadcast media and could not, for example, easily announce to the "walking wounded" where to assemble.

1.2. Exigent Communications

TOC

With the usage of the term 'Exigent Communications' this document aims to generalize the concept of conveying alerts to IP-based systems and at the same time to re-define the actors that participate in the messaging communication. More precisely, exigent communications is defined as:

> Communication that requirs immediate action or remedy. Information about the reason for action and details about the steps that have to be taken are provided in the alert message.
>
> An alert message (or warning message) is a cautionary advice about something imminent (especially imminent danger or other unpleasantness). In the context of exigent communication such an alert message refers to a future, ongoing or past event as the signaling exchange itself may relate to different stages of the lifecycle of the event. The alert message itself, and not the signaling protocol that convey it, provides sufficient context about the specific state of the lifecycle the alert message refers to.

There are two types of basic communication models utilized for the distribution of alert messages and relevant for this document:

> **Alert Push Communication:** With this alert communication paradigm alert messages are sent to typically many Recipients without a prior explicit communication exchange soliciting the desire to receive the alerts. Typically, the criteria for becoming a Recipient are based on current location of the Recipient itself since alerts are targeted to a specific geographical region (an area immediately relevant to the emergency event).
>
> **Alert Subscription Communication** The alert distribution in this category assumes that the Recipient has indicated interest in receiving certain type of alerts using a protocol mechanism (for example, a subscribe event). This opt-in subscription model allows Recipients to sign-up for receiving alerts independently of their current geographical location. For example, parents may want to be alerted of emergencies affecting the school attended by their children and adult children may need to know about emergencies affecting elderly parents.

Note that the Receivers of the alerts may not necessarily be the typical end devices humans carry around, such as mobile phones, Internet tablets, or laptops. Instead, alert distribution may well directly communicate with displays in subway stations, or electronic bill boards. When a Receiver obtains such an alert then it may not necessarily need to interact with a human (as the Recipient) but may instead use the alert as input to another process to trigger automated behaviors, such as closing vents during a chemical spill or activating sirens or other warning systems in commercial buildings.
This document provides terminology, requirements and an architectural description. To avoid the bias towards a specific communication model or technology this documents utilizes the EMail architecture terminology from [RFC5598] (Crocker, D., “Internet Mail Architecture,” July 2009.).

2. Terminology

TOC

The keywords "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119] (Bradner, S., “Key words for use in RFCs to Indicate Requirement Levels,” March 1997.), with the important qualification that, unless otherwise stated, these terms apply to the design of a protocol conveying warning messages, not its implementation or application.
This document reuses the terminology from [RFC5598] (Crocker, D., “Internet Mail Architecture,” July 2009.). For editorial and

consistency reasons parts of the text are repeated in this document and modified as appropriate.

TOC

3. Responsible Actor Roles

The communication system used for the dissemination of alert messages builds on top of existing communication infrastructure. At the time of writing this underlying communication infrastructure is the Session Initiation Protocol (SIP) and the Extensible Messaging and Presence Protocol (XMPP). These distributed services consist of a variety of actors playing different roles. On a high level we differentiate between the User, and the Message Handling Service (MHS) actors. We will describe them in more detail below.

TOC

3.1. User Actors

Users are the sources and sinks of alert messages. Users can be people, organizations, or processes. There are three types of Users:

- Authors

- Recipients

- Mediators

From the user perspective, all alert message transfer activities are performed by a monolithic Message Handling Service (MHS), even though the actual service can be provided by many independent organizations.

TOC

3.1.1. Author

The Author is responsible for creating the alert message, its contents, and its intended recipients, even though the exact list of recipients may be unknown to the Author at the time of writing the alert message. The MHS transfers the alert message from the Author and delivers it to the Recipients. The MHS has an Originator role that correlates with the Author role.
For most use cases the Author is a human creating a message.

TOC

3.1.2. Recipient

The Recipient is a consumer of the delivered alert message. The MHS has a Receiver role that correlates with the Recipient role.
For most use cases the Recipient is a human reading a message.

3.1.3. Mediator

TOC

A Mediator receives, aggregates, reformulates, and redistributes alert messages among
A Mediator attempts to preserve the original Author's information in the message it reformulates but is permitted to make meaningful changes to the message content or envelope. The MHS sees a new message, but users receive a message that they interpret as being from, or at least initiated by, the Author of the original message. The role of a Mediator is not limited to merely connecting other participants; the Mediator is responsible for the new message.
A Mediator's role is complex and contingent, for example, modifying and adding content or regulating which users are allowed to participate and when. The common example of this role is an aggregator that accepts alert messages from a set of Originators and distributes them to a potentially large set of Recipients. This functionality is similar to a multicast, or even a broadcast. Recipients might have also indicated their interest to receive certain type of alerts messages or they might implicitly get entitled to receive specific alerts purely by their presence in a specific geographical region. Hence, a Mediator might have additional information about the Recipients context and might therefore be able to make a decision whether the Recipient is interested in receiving a particular alert message.
A Gateway is a particularly interesting form of Mediator. It is a hybrid of User and Relay that connects to other communication systems. Its purpose is to emulate a Relay.

3.2. Message Handling Service (MHS) Actors

TOC

The Message Handling Service (MHS) performs a single end-to-end transfer of warning messages on behalf of the Author to reach the Recipient addresses. As a pragmatic heuristic MHS actors actors generate, modify or look at only transfer data, rather than the entire message.
Figure 1 (Relationships Among MHS Actors) shows the relationships among transfer participants. Although it shows the Originator as distinct from the Author and Receiver as distinct from Recipient, each pair of roles usually has the same actor. Transfers typically entail one or

more Relays. However, direct delivery from the Originator to Receiver is possible. Delivery of warning messages within a single administrative boundary usually only involve a single Relay.

```
  ++==========++                        ++===========++
  ||  Author  ||                        || Recipient ||
  ++====++====++                        ++===========++
        ||                                     /\
        ||                                     ||
        \/                                     ||
   +----------+                            +---++----+
   |          |                            |         |
 /-+----------+----------------------------+---------+---\
 | |          |     Message Handling       |         |   |
 | |Originator|       System (MHS)         |Receiver |   |
 | |          |                            |         |   |
 | +---++-----+                            +---------+   |
 |     ||                                      /\        |
 |     ||                                      ||        |
 |     \/                                      ||        |
 | +---------+         +---------+        +-+--++---+    |
 | |  Relay  +======-=>|  Relay  +=======>|  Relay  |    |
 | +---------+         +----++---+        +---------+    |
 |                          ||                           |
 |                          ||                           |
 |                          \/                           |
 |                     +---------+                       |
 |                     | Gateway +-->                    |
 |                     +---------+                       |
 \-------------------------------------------------------/

Legend: === and || lines indicate primary (possibly
            indirect) transfers or roles
```

Figure 1: Relationships Among MHS Actors

3.2.1. Originator

TOC

The Originator ensures that a warning message is valid for transfer and then submits it to a Relay. A message is valid if it conforms to both communication and warning message encapsulation standards and local operational policies. The Originator can simply review the message for conformance and reject it if it finds errors, or it can create some or all of the necessary information.

The Originator serves the Author and can be the same entity. But its role in assuring validity means that it also represents the local operator of the MHS, that is, the local ADministrative Management Domain (ADMD).

The Originator also performs any post-submission, Author-related administrative tasks associated with message transfer and delivery. Notably, these tasks pertain to sending error and delivery notices, enforcing local policies, and dealing with messages from the Author that prove to be problematic for the Internet. The Originator is accountable for the message content, even when it is not responsible for it. The Author creates the message, but the Originator handles any transmission issues with it.

3.2.2. Relay

TOC

The Relay performs MHS-level transfer-service routing and store-and-forward, by transmitting or retransmitting the message to its Recipients. The Relay may add history information (e.g., as available with SIP History Info [RFC4244] (Barnes, M., “An Extension to the Session Initiation Protocol (SIP) for Request History Information,” November 2005.)) or security related protection (e.g., as available with SIP Identity [RFC4474] (Peterson, J. and C. Jennings, “Enhancements for Authenticated Identity Management in the Session Initiation Protocol (SIP),” August 2006.)) but does not modify the envelope information or the message content semantics.

A Message Handling System (MHS) network consists of a set of Relays. This network is above any underlying packet-switching network that might be used and below any Gateways or other Mediators.

3.2.3. Gateway

TOC

A Gateway is a hybrid of User and Relay that connects heterogeneous communication infrastructures. Its purpose is to emulate a Relay and the closer it comes to this, the better. A Gateway operates as a User when it needs the ability to modify message content.

Differences between the different communication systems can be as small as minor syntax variations, but they usually encompass significant, semantic distinctions. Hence, the Relay function in a Gateway presents a significant design challenge, if the resulting performance is to be seen as nearly seamless. The challenge is to ensure user-to-user functionality between the communication services, despite differences in their syntax and semantics.

The basic test of Gateway design is whether an Author on one side of a Gateway can send a useful warning message to a Recipient on the other side, without requiring changes to any components in the Author's or Recipient's communication service other than adding the Gateway. To each of these otherwise independent services, the Gateway appears to be a native participant.

3.2.4. Receiver

TOC

The Receiver performs final delivery or sends the warning message to an alternate address. In case of warning messages it is typically responsible for ensuring that the appropriate user interface interactions are triggered.

4. Requirements

TOC

Requirements that relate to the encoding and the content of alert messages are outside the scope of this document. This document focuses on the protocols utilized to convey alert messages only.

The requirements for the two main communication models are different and reflected in separate sub-sections. For the Alert Push commnication model Section 4.2 (Requirements for a Alert Push Communication Model) the assumption is that the potential recipient's consent to provide alerts has been obtained a-priori and the message customization has externally been defined. There is no separate protocol exchange to indicate preferences. The consent may have been waived by law or has been provided when the receipient has registered for a service. As an alternative approach, the Alert Subscription communication model Section 4.1 (Requirements for a Alert Subscription Communication Model) allows the potential alert receipient to indicate preferences about the type of alerts it is interested in. This mechanism to express interest is provided as part of the protocol exchange, namely via a subscription.

> **Req-G1:**
> The protocol solution MUST allow delivery of messages simultaneously to a large audience.
>
> **Req-G2:**
> The protocol solution MUST be independent of the underlying link layer technology.

> **Req-G3:**
> The protocol solution MUST allow targeting notifications to specific individuals and to groups of individuals.
>
> **Req-R4:**
> The protocol solution MUST allow a Recipient to learn the identity of the Author of the alert message.

4.1. Requirements for a Alert Subscription Communication Model

TOC

The requirements listed below largely relate to the subscription phase when the potential recipient of alert messages indicates preferences regarding the type of alerts.

> **Req-S1:**
> The protocol solution MUST allow a potential Recipient to indicate the language used by alert messages.
>
> **Req-S2:**
> The protocol solution MUST allow a potential Recipient to express the geographical area it wants to receive alerts about.
>
> **Req-S3:**
> The protocol solution MUST allow a potential Recipient to indicate preferences about the type of alerts it wants to receive.
>
> **Req-S4:**
> The protocol solution MUST allow a potential Recipient to express preference for certain media types. The support for different media types depends on the content of the warning message but also impacts the communication protocol. This functionality is, for example, useful for hearing and vision impaired persons.

TOC

4.2. Requirements for a Alert Push Communication Model

> **Req-P1:**
> The protocol solution MUST allow delivery of alerts by utilizing he lower layer infrastructure ensuring congestion control being considered. Network layer multicast, anycast or broadcast mechanisms may be utilized. The topological network structure may be used for efficient alert distribution.

5. IANA Considerations

TOC

This document does not require actions by IANA.

6. Security considerations

TOC

With the distribution of alert messages a number of security threats need to be addressed. Because of the nature of alerts it is quite likely that end device implementations will want to provide user interface enhancements to get the attention whenever an alert arrives. This creates additional attractiveness for adversaries to exploit an alert Message Handling System. We list the most important threats below that any solution will have to deal with.

> **Originator Impersonation:**
> An attacker could then conceivably attempt to impersonate the Originator of an alert message. This threat is particularly applicable to those deployment environments where authorization decisions are based on the identity of the Originator.
>
> **Alert Message Forgery:**
> An attacker could forge or alter an alert message in order to convey custom messages to Recipients to get their immediate attention.
>
> **Replay:**
> An attacker could obtain previously distributed alert messages and to replay them at a later time in the hope that Recipients could be tricked into believing they are fresh.

Unauthorized Distribution:

> When a Receiver receives an alert message it has to determine whether the Author distributing the alert messages is genuine to avoid accepting messages that are injected by malicious entities with the potential desire to at least get the immediate attention of the Recipient.

Amplification Attack:

> An attacker may use the Message Handling System to inject a single alert message for distribution that may then be instantly turned into potentially millions of alert messages for distribution.

One important security challenge worth mentioning is related to authorization. When an alert message arrives at a Receiver, a software module at a host, then certain security checks can be performed to ensure that the message meets certain criteria. The final consumer of the alert message is, however, the Recipient, which in many cases is a human. From a security point of view the work split between the Recipient and the Receiver for making the authorization decision is important and the clarification of when to drop a message due to a failed security verfication by the Receiver. False positives may be fatal but accepting every alert message lowers the trustworthiness in the overall system.

---

7. Acknowledgments

TOC

This document re-uses a lot of text from [RFC5598] (Crocker, D., “Internet Mail Architecture,” July 2009.). The authors would like to thank Dave Crocker for his work.
The authors would like to thank Martin Thomson, Carl Reed, and Tony Rutkowski for their comments.

---

8. References

TOC

8.1. Normative References

TOC

[RFC2119] Bradner, S., “Key words for use in RFCs to Indicate Requirement Levels,” BCP 14, RFC 2119, March 1997 (TXT, HTML, XML).

[RFC5598]	Crocker, D., “Internet Mail Architecture,” RFC 5598, July 2009 (TXT, PDF).

8.2. Informative References

TOC

[July2005]	, ., “Report of the 7 July Review Committee, ISBN 1 85261 878 7,” (PDF document), http://www.london.gov.uk/assembly/reports/7july/report.pdf, June 2006.
[RFC4244]	Barnes, M., “An Extension to the Session Initiation Protocol (SIP) for Request History Information,” RFC 4244, November 2005 (TXT).
[RFC4474]	Peterson, J. and C. Jennings, “Enhancements for Authenticated Identity Management in the Session Initiation Protocol (SIP),” RFC 4474, August 2006 (TXT).

Authors' Addresses

TOC

Henning Schulzrinne
Columbia University
Department of Computer Science
450 Computer Science Building
New York, NY 10027
US
Phone: +1 212 939 7004
Email: hgs+ecrit@cs.columbia.edu
URI: http://www.cs.columbia.edu

Steve Norreys
BT Group
1 London Road
Brentwood, Essex CM14 4QP
UK
Phone: +44 1277 32 32 20
Email: steve.norreys@bt.com

Brian Rosen
NeuStar, Inc.
470 Conrad Dr
Mars, PA 16046
US
Phone:
Email: br@brianrosen.net

Hannes Tschhofenig

	Nokia Siemens Networks
	Linnoitustie 6
	Espoo 02600
	Finland
Phone:	+358 (50) 4871445
Email:	Hannes.Tschofenig@gmx.net
URI:	http://www.tschofenig.priv.at