

ATOCA
Internet-Draft
Intended status: Informational
Expires: July 19, 2011

H. Schulzrinne
Columbia University
S. Norreys
BT Group
B. Rosen
NeuStar, Inc.
H. Tschofenig
Nokia Siemens Networks
January 15, 2011

Requirements, Terminology and Framework for Exigent Communications
draft-ietf-atoca-requirements-01.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

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on July 19, 2011.

Copyright Notice

Copyright (c) 2011 IETF Trust and the persons identified as the document authors. All rights reserved.

Internet-Draft

Exigent Communications

January 2011

This document is subject to [BCP 78](http://trustee.ietf.org/license-info) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
1.1.	Classical Early Warning Situations	3
1.2.	Exigent Communications	3
2.	Terminology	5
3.	Alert Delivery Architecture	5
3.1.	Responsible Actor Roles	5
3.1.1.	User Actors	5
3.1.2.	Message Handling Service (MHS) Actors	6
4.	Requirements	9
4.1.	Requirements for Alert Subscription	9
4.2.	Requirements for Alert Message Delivery	9
5.	IANA Considerations	10
6.	Security Considerations	10
7.	Acknowledgments	12
8.	References	12
8.1.	Normative References	12
8.2.	Informative References	13
	Authors' Addresses	13

Internet-Draft

Exigent Communications

January 2011

[1.](#) Introduction

[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](#)]. 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

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 requires 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.

Communication typically occurs in two phases:

Subscription: In this step Recipients express their interest to receive certain types of alerts and happens prior to the actual delivery of the alert. This expression of interest may be in form of an explicit communication step by having the Receiver sending a subscription message potentially with an indication of the type of alerts they are interested in, the duration of the subscription and a number of other indicators. 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. The subscription step may, however, also happen outside the Internet communication infrastructure but rather by the Recipient signing a contract and thereby agreeing to receive certain alerts. Additionally, certain subscriptions may happen without the Recipient's explicit consent and without the Receiver sending a subscription. For example, a Tsunami flood alert may be delivered to Recipients in case they are located in a specific geographical area.

It is important to note that a protocol interaction initiated by the Receiver may need to take place to subscribe to certain types of alerts. In some other cases the subscription does not require such interaction from the Receiver. Orthogonal to the need to have a protocol interaction is the question of opt-in vs. opt-out. This is a pure policy decision and largely outside the scope of a technical specification.

Alert Delivery In this step the alert message is distributed to one

or multiple Receivers. The Receiver as a software module then presents the alert message to the Recipient. The alert encoding is accomplished via the Common Alerting Protocol (CAP) and such an alert message contains useful information needed for dealing with the imminent danger.

Note that alert Receivers as software modules may not necessarily only be executed on end devices humans typically 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. Note that the requirements focus on the communication

protocols for subscription and alert delivery rather than on the content of the alert message itself. With the usage of CAP these alert message content requirements are delegated to the authors and originators of alerts.

[2.](#) Terminology

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](#)], 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.

[3.](#) Alert Delivery Architecture

This section illustrates the roles useful for alert delivery.

[3.1.](#) 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.

[3.1.1.](#) User Actors

Users are the sources and sinks of alert messages. We differentiate between two types of users:

- o Authors
- o Recipients

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.

[3.1.1.1.](#) Author

The Author is a human 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.

[3.1.1.2.](#) Recipient

The Recipient is a consumer of the delivered alert message. It is a human reading the alert message.

[3.1.2.](#) Message Handling Service (MHS) Actors

The Message Handling Service (MHS) performs a single end-to-end transfer of warning messages on behalf of the Author to reach the Recipient. As a pragmatic heuristic MHS actors generate, modify or look at only transfer data, rather than the entire message.

Figure 1 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.

```

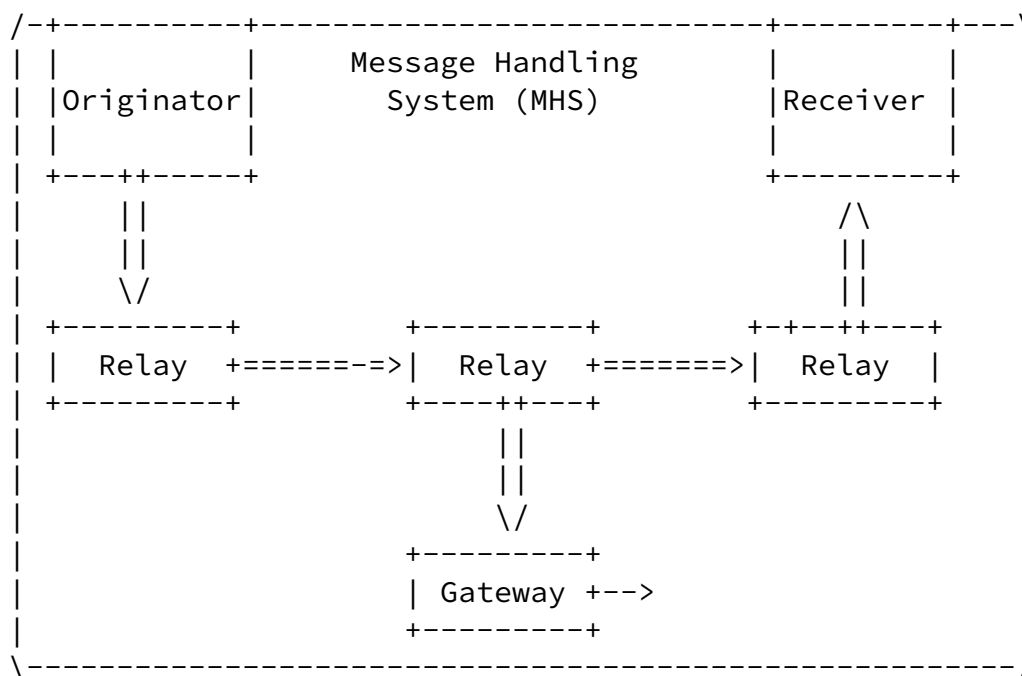
++=====++
|| Author ||
++=====++
  ||
  ||
  \ /
+-----+
|         |

```

```

++=====++
|| Recipient ||
++=====++
  /\
  ||
  ||
+-----+
|         |

```



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

Figure 1: Relationships Among MHS Actors

[3.1.2.1.](#) Originator

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 in absence of a human crafting alert messages.

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.

3.1.2.2. Relay

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](#)]) or security related protection (e.g., as available with SIP Identity [[RFC4474](#)]) 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 MHS network is above any underlying packet-switching network that might be used and below any Gateways.

3.1.2.3. Gateway

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 Receiver's communication service other than adding the Gateway. To each of these otherwise independent services, the Gateway appears to be a native participant.

3.1.2.4. Receiver

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 to interact with the Recipient.

[4.](#) Requirements

[4.1.](#) Requirements for Alert Subscription

The requirements listed below refer to the alert subscription phase.

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.

[4.2.](#) Requirements for Alert Message Delivery

The requirements listed below refer to the delivery of alerts.

Req-D1:

The protocol solution MUST allow delivery of alerts by utilizing the lower layer infrastructure ensuring congestion control being considered. Note that congestion does not only focus on over-utilization of a network caused by a large number of alerts but also in relationship with other traffic not related to exigent communication. Network layer multicast, anycast or broadcast mechanisms may be utilized. The topological network structure may be used for efficient alert distribution.

Internet-Draft

Exigent Communications

January 2011

Req-D2:

The protocol solution MUST allow delivery of messages simultaneously to a large audience.

Req-D3:

The protocol solution MUST be independent of the underlying link layer technology.

Req-D4:

The protocol solution MUST allow targeting notifications to specific individuals and to groups of individuals.

Req-D5:

The protocol solution MUST allow a Recipient to learn the identity of the Author of the alert message.

[5.](#) IANA Considerations

This document does not require actions by IANA.

[6.](#) Security Considerations

Figure 1 shows the actors for delivering an alert message assuming that a prior subscription has taken place already. The desired security properties of an MHS for conveying alerts will depend on the number of administrative domains involved. Each administrative domain can have vastly different operating policies and trust-based decision-making. One obvious example is the distinction between alert messages that are exchanged within an closed group (such as alert messages received by parents affecting the school attended by their children) and alert messages that are exchanged between independent organizations (e.g., in case of large scale disasters).

The rules for handling both types of communication architectures tend to be quite different. That difference requires defining the boundaries of each.

Operation of communication systems that are used to convey alert messages are typically carried out by different providers (or operators). Since each be in operated in an independent administrative domain it is useful to consider administrative domain boundaries in the description to facilitate discussion about designs,

policies and operations that need to distinguish between internal issues and external entities. Most significant is that the entities communicating across administrative boundaries typically have the added burden of enforcing organizational policies concerning external communications. For example, routing alerts between administrative domains can create requirements, such as needing to route alert messages between organizational partners over specially trusted paths.

The communication interactions are subject to the policies of that domain, which cover concerns such as these:

- o Reliability
- o Access control
- o Accountability
- o Content evaluation, adaptation, and modification

Many communication system make the distinction of administrative domains since they impact the requirements on security solutions. However, with the distribution of alert messages a number of additional security threats need to be addressed. Due to the nature of alerts it is quite likely that end device implementations will offer user interface enhancements to get the Recipients attention whenever an alert arrives, which is an attractive property for adversaries to exploit. Below we list the most important threats 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 is related to authorization. When an alert message arrives at the Receiver then certain security checks may need to be performed to ensure that the alert message meets certain criteria. The final consumer of the alert message is, however, the Recipient - a human. From a security point of view the work split between the Recipient and the Receiver for making the authorization decision is important, particularly when an alert message is rejected due to a failed security verification by the Receiver. False positives may be fatal but accepting every alert message lowers the trustworthiness in the overall system.

7. Acknowledgments

This document re-uses text from [[RFC5598](#)]. The authors would like to thank Dave Crocker for his work.

The authors would like to thank Martin Thomson, Carl Reed, Leopold Murhammer, and Tony Rutkowski for their comments.

At IETF#79 the following persons provided feedback leading to changes in this document: Keith Drage, Scott Bradner, Ken Carberg, Keeping Li, Martin Thomson, Igor Faynberg, Mark Wood, Peter Saint-Andre.

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC5598] Crocker, D., "Internet Mail Architecture", [RFC 5598](#), July 2009.

Schulzrinne, et al.	Expires July 19, 2011	[Page 12]
---------------------	-----------------------	-----------

Internet-Draft	Exigent Communications	January 2011
----------------	------------------------	--------------

8.2. Informative References

- [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.
- [RFC4474] Peterson, J. and C. Jennings, "Enhancements for Authenticated Identity Management in the Session Initiation Protocol (SIP)", [RFC 4474](#), August 2006.

Authors' Addresses

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

Schulzrinne, et al.

Expires July 19, 2011

[Page 13]

Internet-Draft

Exigent Communications

January 2011

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>