ECRIT
Internet-Draft
Intended status: Informational
Expires: September 7, 2010

S. Norreys
BT Group
H. Tschofenig
Nokia Siemens Networks
H. Schulzrinne
Columbia University
March 6, 2010

Requirements, Terminology and Framework for Exigent Communications
draft-norreys-ecrit-authority2individuals-requirements-04.txt

Abstract

Various agencies need to provide information to the restricted group of persons or even to the generic public before, during and after emergency situations. 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 summarizes requirements for protocols to allow alerts to be conveyed to IP-based end points.

Table of Contents

1. Introduction 3
1.1. Classical Early Warning Situations 3
1.2. Exigent Communications 3
2. Terminology 4
3. Responsible Actor Roles 5
3.1. User Actors 5
3.1.1. Author 5
3.1.2. Recipient 6
3.1.3. Return Handler 6
3.1.4. Mediator 6
3.2. Message Handling Service (MHS) Actors 7
3.2.1. Originator 8
3.2.2. Relay 9
3.2.3. Gateway 9
3.2.4. Receiver 9
3.3. Administrative Actors 10
4. Requirements 10
4.1. Communication Model Independent Requirements 11
4.2. Requirements for a Subscription Model 11
4.3. Requirements for a Push Communication Model 12
5. IANA Considerations 13
6. Security considerations 13
7. Acknowledgments 13
8. References 13
8.1. Normative References 13
8.2. Informative References 13
Authors' Addresses 13

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 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, provides sufficient context about the specific state of the lifecycle the alert message refers to.

For that purpose, the terminology utilized by the EMail architecture, see [I-D.crocker-email-arch], is applied to this context.

Three types of communication models can be envisioned:

1. Alerts may be addressed to all individuals within a certain geographic area. Today, this is often realized with the help of dedicated functionality provided by link layer technology (e.g., multicast, broadcast).

2. Alerts need to be delivered to dedicated end points via unicast messaging. Examples are displays in subway stations, or electronic bill boards. Some of these alerts may also be used to trigger automated behaviors, such as closing vents during a chemical spill or activating sirens or other warning systems in commercial buildings. Other messages may target only specific groups of individuals, such as medical personnel. These may include cases where legacy end points need to be integrated into the overall architecture and some form of protocol translation is necessary. The communication end point from an IP point of view is therefore a single gateway (or a small number of them).

3. The two models described above illustrate a push communication whereas the third model represents a subscription model where an opt-in model is used to provide further information about the type of alerts that the recipient is interested in. The information that may lead to an alert message being distributed may depend on certain factors, including certain types of events happening in a specific geographic region irrespectively of whether the entity issuing the subscription is actually located in that geographic region. 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.

This document focuses on all three types of communication models whereby a stronger emphasis is given to the subscription model since it is very powerful but less widely deployed on the Internet for exigent communication. Content-wise this document provides terminology, requirements and the architecture for IP-based protocols to enhance and complement existing authority-to-citizen warning systems.

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.

This document reuses the terminology from [I-D.crocker-email-arch]. For editorial and consistency reasons parts of the text are repeated in this document and modified as appropriate.

3. Responsible Actor Roles

The communication system used for the dissemination of alert messages builds on top of existing communication infrastructure. These distributed services consist of a variety of actors playing different roles. These actors fall into three basic categories:

- User
- Message Handling Service (MHS)
- ADministrative Management Domain (ADMD)

3.1. User Actors

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

- Authors
- Recipients
- Return Handlers
- 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.

Whenever any MHS actor sends information to back to an Author or Originator in the sequence of handling a message, that actor is a User.

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.

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.

3.1.3. Return Handler

The Return Handler is a special form of Recipient tasked with servicing notifications that the MHS generates, as it transfers or delivers the message. These notices can be about failures or completions (such as utilized by test messages) and are sent to an address that is specified by the Originator. This Return Handling address (also known as a Return address) might have no visible characteristics in common with the address of the Author or Originator.

3.1.4. Mediator

A Mediator receives, aggregates, reformulates, and redistributes alert messages among Authors and Recipients who are the principals in (potentially) protracted exchanges. When submitting a reformulated message, the Mediator is an Author, albeit an author actually serving as an agent of one or more other authors. So, a Mediator really is a full-fledged User.

The aspect of a Mediator that distinguishes it from any other MUA creating a message is that a Mediator preserves the integrity and tone of the original message, including the essential aspects of its origination information. The Mediator might also add commentary.

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

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. Exchanges that are either mediated or iterative and protracted, such as those used for communicating information about the lifetime of an alert are handled by the User actors, not by the MHS actors. As a pragmatic heuristic MHS actors 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.

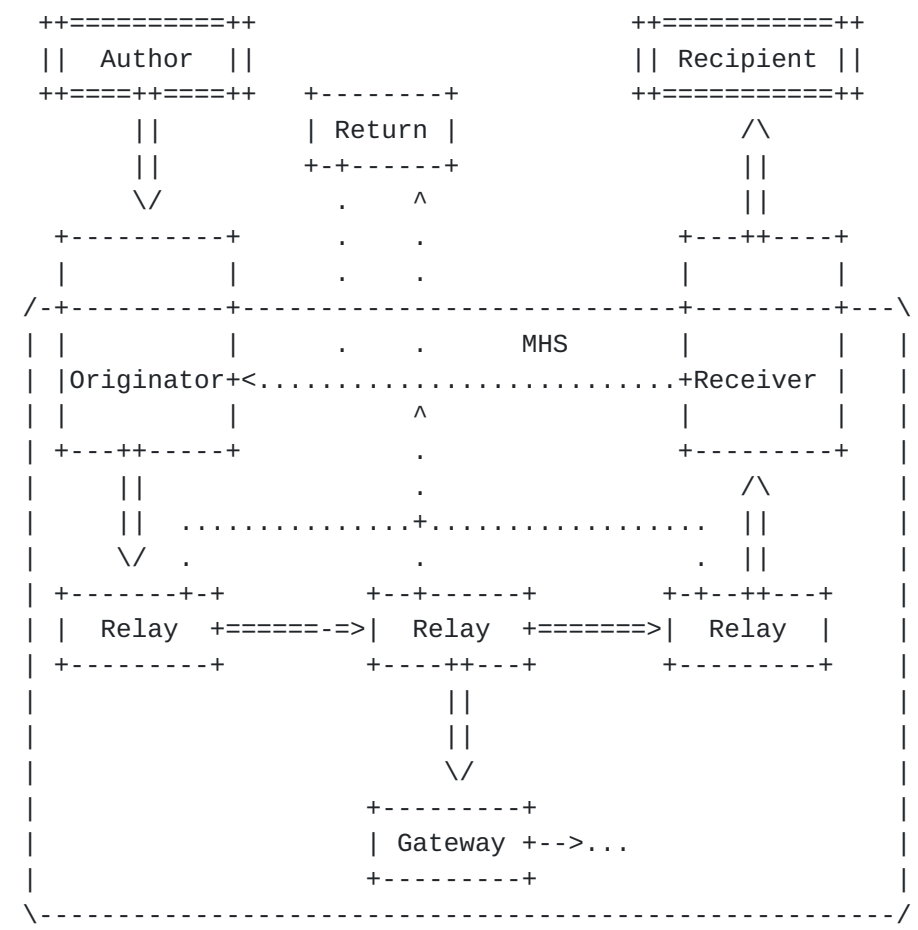

Legend: === and || lines indicate primary (possibly indirect) transfers or roles
... lines indicate supporting transfers or roles

Figure 1: Relationships Among MHS Actors

3.2.1. 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 operates with dual allegiance. It 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

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 / trace information 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 network is above any underlying packet-switching network that might be used and below any Gateways or other Mediators.

3.2.3. 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 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

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.

3.3. Administrative Actors

Administrative actors can be associated with different organizations, each with its own administrative authority. This operational independence, coupled with the need for interaction between groups, provides the motivation to distinguish among ADministrative Management Domains (ADMDs). Each ADMD can have vastly different operating policies and trust-based decision-making. One obvious example is the distinction between warning messages that are exchanged within an closed group (such as alert messages received by parents affecting the school attended by their children) and warning messages that exchanged between independent organizations (e.g., in case of large scale disasters). The rules for handling both types of traffic tend to be quite different. That difference requires defining the boundaries of each, and this requires the ADMD construct.

Operation of communication systems that are used to convey alert messages are typically carried out by different providers (or operators). Each can be an independent ADMD. The benefit of the ADMD construct is to facilitate discussion about designs, policies and operations that need to distinguish between internal issues and external ones. Most significant is that the entities communicating across ADMD boundaries typically have the added burden of enforcing organizational policies concerning external communications. At a more mundane level, routing mail between ADMDs can be an issue, such as needing to route alert messages between organizational partners over specially trusted paths.

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

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

4. Requirements

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

The requirements for the two main communication models are different

and reflected in separate sub-sections, Section 4.2 and Section 4.3 . There are, however, a few generic requirements applicable to both communication models described in Section 4.1.

4.1. Communication Model Independent Requirements

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 offer the typical communication security mechanisms. Additional security mechanisms applied to the alert message itself are outside the scope of the communication protocol and therefore outside the scope of this document.

Req-G4:

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

Req-G5:

The protocol solution MAY provide an option to return a receipt on reading message.

Req-G6:

The protocol solution MUST ensure that congestion handling is provided.

4.2. Requirements for a Subscription Model

The requirements for subscription / opt-in model require information about the type of alerts that are being asked for to be made available by the potential Recipient to the Originator or set of orginators.

Req-S1:

The protocol solution MUST allow to tailor the message to the language preferences of the receiver.

Req-S2:

The protocol solution MUST allow an indication about the geographical area the potential Recipient is interested in.

Req-S3:

The protocol solution MUST allow an indication about the type of alert the potential Recipient is interested in.

Req-S4:

The protocol solution MUST allow an indication of the media types that are understood or preferred by the potential Recipient.

The support for different media types depends to some extend on the content of the warning message but the communication protocol may be impacted as well. This functionality would, for example, be useful for those with special needs, such as hearing and vision impaired persons.

Req-S5:

The protocol solution MUST allow a potential Recipient to discover the responsible Originator or set of Originators for a certain category of warning messages.

4.3. Requirements for a Push Communication Model

The topological structure of networks is used to distribute warning notifications to all hosts that are located within a specific IP-subsetwork or multicast group.

Req-P1:

The protocol solution MUST allow network layer multicast and broadcast mechanisms to be utilized.

5. IANA Considerations

This document does not require actions by IANA.

6. Security considerations

This document outlines requirements and security requirements are a part of them.

7. Acknowledgments

This document re-uses a lot of text from [I-D.crocker-email-arch]. The authors would like to thank Dave Crocker for his work.

8. References

8.1. Normative References

[I-D.crocker-email-arch]
Crocker, D., "Internet Mail Architecture", draft-crocker-email-arch-14 (work in progress), June 2009.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

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

Steve Norreys
BT Group
1 London Road
Brentwood, Essex CM14 4QP
UK

Phone: +44 1277 32 32 20
Email: steve.norreys@bt.com

Hannes Tschofenig
Nokia Siemens Networks
Linnoitustie 6
Espoo 02600
Finland

Phone: +358 (50) 4871445
Email: Hannes.Tschofenig@gmx.net
URI: http://www.tschofenig.priv.at

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