ECRIT
Internet-Draft
Intended status: Informational
Expires: September 9, 2009

S. Norreys
BT Group
H. Tschofenig
Nokia Siemens Networks
H. Schulzrinne
Columbia University
March 8, 2009

Authority-to-Individuals Communication for Emergency Situations: Requirements, Terminology and Architecture

draft-norreys-ecrit-authority2individuals-requirements-02.txt

Status of this Memo

Copyright Notice

Abstract

Public safety agencies need to provide information to the general public before and during large-scale emergencies. 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 alert individuals within a defined geographic area.

Table of Contents

1. Introduction

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 recently been illustrated by the London underground bombing on July 7, 2006, as described in a government report [ref]. The UK authorities could only use broadcast media and could not, for example, easily announce to the "walking wounded" where to assemble.

This document summarizes key requirements for IP-based protocols to enhance and complement existing authority-to-citizen warning systems. These protocols may either directly reach the citizen or may be used to trigger more traditional alerts, such as, among many others, displays in subway stations, electronic bill boards, or SMS.

Public safety authorities need to reach, with an appropriate message, as many affected people as possible within the area impacted by the emergency, including not just residents, but also workers and travelers who may only be in the area temporarily.

In addition, people around the immediately affected area should be able to receive information and differentiated instructions, such as warnings to avoid travel or to clear roads.

Emergency alerts may be issued once for an emergency or authorities may repeat or update information during an event.

Some messages are addressed to all individuals within a certain geographic area. Other messages may target only specific individuals or groups of individuals, such as medical personnel or those particularly susceptible to an incident.

Machine-parseable 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.

At least initially, mobile and stationary devices may not have the appropriate capabilities to receive such warnings. Thus, protocols need to be designed to allow gatewaying to traditional systems, e.g., the PSTN.

We assume an event notification model, i.e., individuals subscribe to warnings that affect their current location. As a mobile device moves, the subscription may need to be updated. Thus, location information needs to be available during the subscription process.

Users may want to subscribe to warnings that do not affect their current 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.

Some technologies may also support network-level broadcast or multicast.

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 emergency alerts and early warning messages, not its implementation or application.

This document reuses the terminology introduced by [3GPP-TR-22.968] and modifies it to fit to IETF terminology:

Notification Population:

The term notification population refers to the set of warning notification consumers that receive a warning notification.

Public Warning System:

A public warning system delivers warning notifications provided by warning notification providers to IP-capable end hosts within notification areas.

Warning Notification:

Warning notifications inform recipients of the occurrence of current or pending public safety-related events and additionally may provide users with instructions on what to do and where to get help for the duration of the emergency.

Warning Notification Provider:

A warning notification provider is an agency that injects warning notifications in the communication system. Examples of such agencies are branches of governments, public transport providers or weather organizations. A private institution, such as a university, hospital or enterprise, may also provide such warnings to people located within their facilities.

Warning Notification Consumer: A warning notification consumer is the final recipient of a warning notification from an IP communication point of view. Examples are Web browsers and SIP clients on mobile phones and laptops that process warning notification messages. Once received such a warning is shown to the user (a human) via an appropriately designed user interface to

ensure that it is properly understood.

Warning Notification Aggregator:

A warning notification aggregator receives alert notifications from different providers, performs security functions (e.g., authentication, authorization) and may need to transform message into a different format before forwarding them towards warning notification consumers.

3. Architectures

The following sub-sections illustrate different architectural approaches.

3.1. Closed Warning Notification Provider and Aggregator Groups

The first architectural variant allows the distribution of warning notifications from warning notification providers to warning notification aggregators. The communication is largely in a point-to-point fashion and the number of involved players is rather small, particularly on the side of the warning notification providers. Furthermore, a new warning notification aggregator is allowed to receive warning notifications only after certain verification procedures are conducted and thereby the entire communication

infrastructure re-essembles a closed group. To ensure that the content of the warning notifications is properly understood the involved parties are very likely to agree on the detailed semantics of the warning messages prior to distributing any alert.

Figure 1 shows the involved parties graphically.

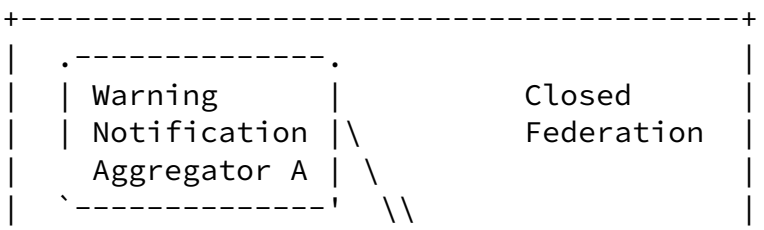

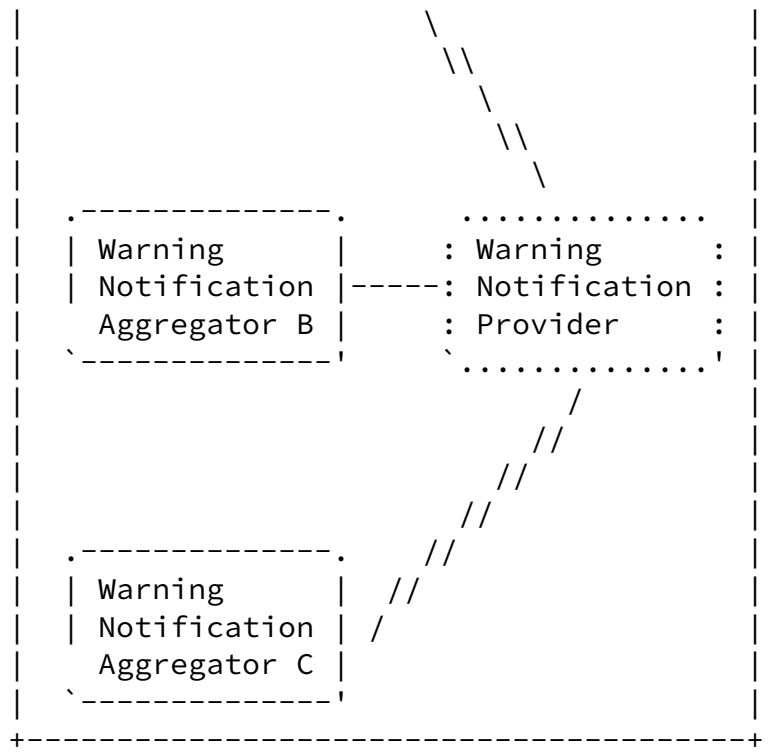

Figure 1: Closed Warning Notification Provider and Aggregator Groups

An example of this system can be seen in national early warning systems where regulatory requirements force warning notification provider and aggregators to work together.

3.2. Open Communication between Warning Notification Provider and Aggregator

This model is similar to the closed group presented in the previous section with a difference in the way how warning notification aggregators can retrieve warning notifications from warning

notification providers. When the aggregator interacts with the provider then no special client-side authentication procedure is assumed and no access restrictions are enforced. As such, the aggregator might be located anywhere on the Internet to retrieve the warning notifications. Warning notification aggregators might offer their subscribers the ability to receive warnings of a certain type (e.g., weather alerts) for a specific region (e.g., for a specific country or an entire continent). Hence, the aggregator might find

the relevant warning notification providers and interacts with them to retrieve alerts. Finally, the aggregator might use different protocol mechanisms (e.g., SMS, Web pages) towards the warning notification customers, often depending on the client capabilities. In general, the number of aggregators is very likely larger than in a closed group but there are no significant challenges with respect to scalability or congestion to expect.

Figure 2 shows the involved parties graphically.

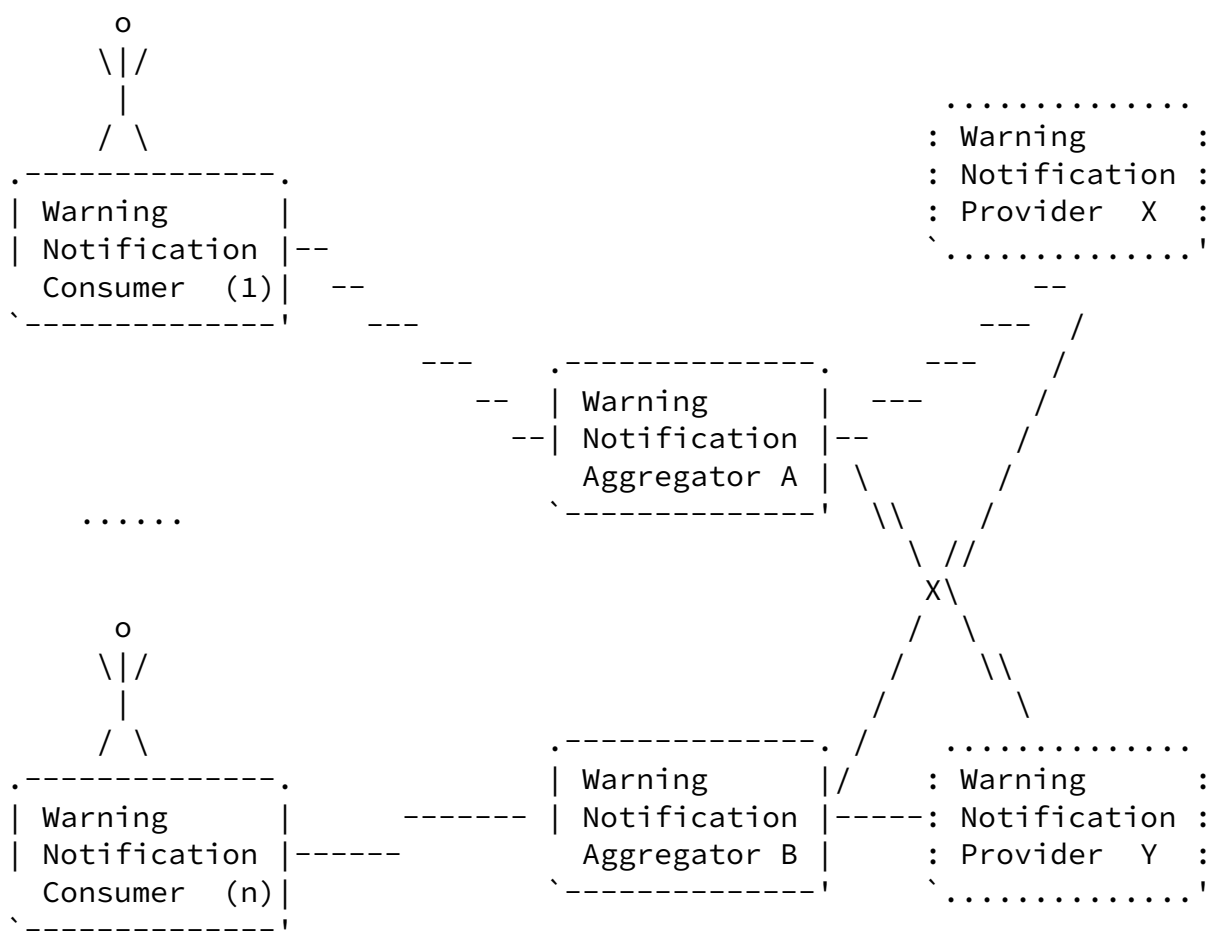

Figure 2: Open Communication between Warning Notification Provider and Aggregator

3.3. Open Communication towards Warning Notification Customers

In the following architectural variant the customers directly interact with various warning notification providers that a relevant for a specific type of alert type. In order to learn about the relevant warning notification providers it may be necessary to consult some form of discovery service; this may be done out-of-band via manual configuration or via a protocol mechanism.

Scalability and congestion control concerns arise with this scenario as this model assumes an opt-in approach from the customer. The total number of warning notification customers are provider has to serve may be considerable. There are a more security challenges since there are no pre-established trust relationships between the warning notification customers and the providers.

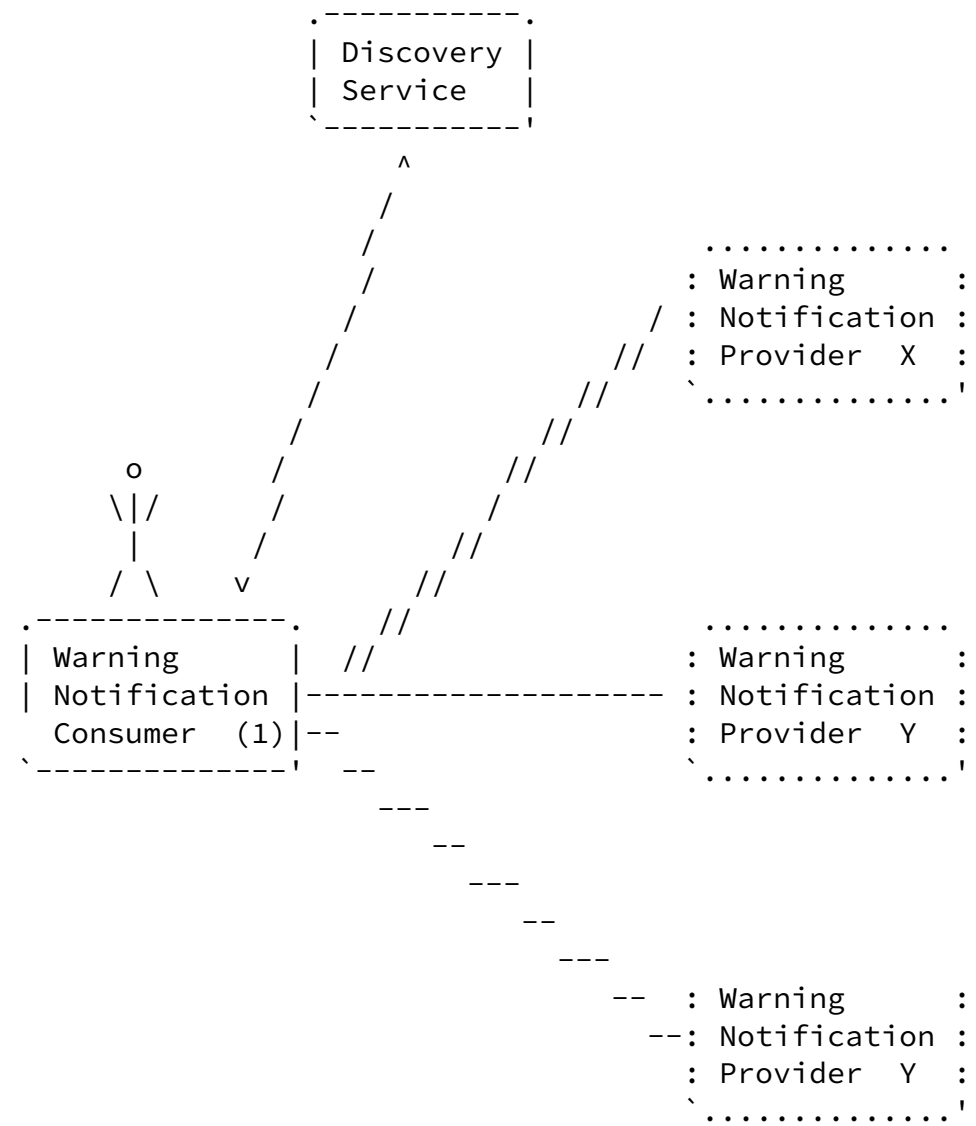

3.4. Notification Population

What warning notification customers are put into the notification population is an important architectural aspect that is largely orthongonal to the architecture presented in Section 3.1 and Section 3.2.

In an opt-in model the the warning notification customer need to provide information about what type of alerts it is interested in and, in order for the warning notification aggregator or the warning notification provider to be able to distribute warnings it is necessary for them to know the context, such as the current location, preferred language or device capabilities. This information may be provided by the customer itself or by other entities, such as the access provider.

Alternatively, the topological structure of networks is used to distribute warning notifications to all hosts that are located within a specific IP-subsetwork or hosts in a specific link layer broadcast domain. This approach typically requires co-operation from the network provider.

4. Requirements

The following requirements are related to the communication protocol and the context.

Req-1:

> The protocol mechanisms MUST allow targeting notifications to specific individuals, groups of individuals or specific geographic regions. Different regions or groups may receive different instructions for the same disaster. (For example, people very close to a chemical spill may be asked to evacuate, while those further away may be asked to close windows.)

Req-2:

> The protocol solution MUST provide real-time message delivery.

Req-3:

> The solution MUST support the delivery of warning notifications

that allow for predictable or likely events.

Req-4

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

Req-5

The protocol mechanisms MAY provide an option to return a receipt on reading message. However, the confirmation SHOULD NOT be required.

Req-6:

The protocol mechanism MUST be independent of the underlying access network technology.

Req-7:

Protocol mechanisms MUST allow to tailor the message to the language preferences of the receiver and/or deliver multiple versions in different languages within the same message, so that the recipient can choose the most appropriate one.

Req-8:

A user SHOULD be able to indicate the preferred method of communication to the public warning service, such as notification by email, different instant messaging protocols or automated voice calls.

Req-9:

The protocol conveying warning notifications SHOULD identify the warning notification provider in a secure manner.

Req-10:

The solution MUST provide a mechanism for transmitting warning notification test messages.

Req-11:

A solution MUST support delivery of notification messages (e.g., with different media types) to those with special needs, such as hearing and vision impaired.

5. IANA Considerations

This document does not require actions by IANA.

6. Security considerations

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

7. Acknowledgments

This document reuses requirements captured outside the IETF, namely ETSI (with [ETSI-TS-102-182]), and the 3GPP (with [3GPP-TR-22.968]). We would like to thank the authors of these specifications for their work. Note, however, that only a small subset of the requirements have been reflected that do not relate to specific deployments, user interface aspects, detailed regulatory requirements, management and operational considerations, and non-IP specific technologies.

We would like to thank Leopold Murhammer for his review in July 2007.

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.

8.2. Informative References

[3GPP-TR-22.968]
, ., "3GPP TR 22.968, V1.0.0 (2007-04), 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; Study for requirements for a Public Warning System (PWS) Service (Release 8)", December 2006.

[ETSI-TS-102-182]
, ., "ETSI TS 102 182, V1.2.1 (2006-12), Technical Specification, Emergency Communications (EMTEL); Requirements for communications from authorities/ organizations to individuals, groups or the general public during emergencies", December 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