RC23594 (W0504-124) April 27, 2005
Computer Science

IBM Research Report

Internet Emergency Alert System

Paridhi Verma, Dinesh C. Verma
IBM Research Division
Thomas J. Watson Research Center
P.O. Box 704
Yorktown Heights, NY 10598

Research Division
Almaden - Austin - Beijing - Haifa - India - T. J. Watson - Tokyo - Zurich

INTERNET EMERGENCY ALERT SYSTEM

Paridhi Verma and Dinesh C. Verma
IBM T J Watson Research Center
19 Skyline Drive, Hawthorne, NY

ABSTRACT

The Emergency Alert System available on television and radio media provides an effective means for dissipating information to the general population during emergencies. Demographic studies in the United States have shown an increasing trend towards the use of Internet by general population, and point to the fact that increased Internet usage is associated with decreased time on television and radio. Since there is no equivalent of the TV and radio Emergency Alert System on the Internet, a significant fraction of the population may be left without prompt information in the case of an emergency. In this paper, we explore the various options for designing an Emergency Alert System for users in Internet. We address the operational and technical challenges associated with each option, present alternative designs for a Internet Emergency Alert System, and compare the relative merits of the different approaches.

INTRODUCTION

Public awareness and preparedness is a key element in the proper handling of any type of emergency. In cases of emergencies threatening public welfare at large, prompt notification and guidance to the general civilian population could significantly reduce the human and economic cost of the emergency. Existing Emergency Alert Systems (EAS) available, while originally designed for Presidential emergency messages, was converted to be used for local emergency notifications from state and local governments in 1963[1] for this purpose.

With the present day trend of increasing Internet usage, the effectiveness of existing EAS systems is significantly reduced. Demographic studies of US population [2] indicate that people using Internet spend 1.7 times as much time on the Internet than on the TV. People who use the Internet tend to watch less TV, with each hour on the Internet reducing 10 minutes of TV time on the average. In 2005, the average US citizen still spends more time on the TV than the Internet, but with the current trends in the usage of the two media, the situation may be reversed not too far in the future. Since there is no equivalent of the TV and radio EAS system on the Internet, a significant fraction of the population may be left without prompt information in the case of an emergency.

Technical differences between the Internet and conventional broadcast media have made it difficult or impossible to establish an Internet equivalent of EAS. Radio and television are characterized by a relatively limited number of centralized transmitters broadcasting a continuous stream of a relatively limited selection of entertainment and other information content to a relatively large number of user terminals in a relatively concentrated geographic area. On the other hand, the Internet is characterized by a relatively large number of geographically dispersed servers providing information content to a relatively small number of user terminals per server, with little or no attention paid to the users' geographic location. Such differences between the Internet and conventional radio or television make it difficult to implement an Internet equivalent of the Emergency Broadcast System.

This paper discusses the different issues associated with creating an analogue of EAS system on the Internet. In the next section, we discuss the characteristics required of Internet EAS; enumerate the existing mechanisms for providing emergency information on the Internet and describe their limitations in acting as an EAS equivalent. Then we consider the different alternative approaches for implementing EAS on the Internet, followed by a detailed discussion on each approach in subsequent sections. Finally, we compare the merits and demerits of each approach, and draw our conclusions.

REQUIREMENTS OF INTERNET EAS

A good EAS would have the properties of locality, automated operation, non-intrusiveness, spontaneity, ubiquity, and support for second languages. Locality implies that the emergency alert should be available to the general population in a locality or geography that is affected by the emergency, and that it should not have a perceptible effect on the population in an area that is not affected by the emergency. Automated operation implies that the EAS system should not require manual actions for a system to switch from a mode of normal operation to a mode in which alerts are being issued, when an authorized authority authorizes the alert notification. Non-intrusiveness means the alert system should not cause any disruption in the actions of the users, nor should it collect any unnecessary information of the user's activities on the Internet. Spontaneity

means that a user on the Internet should be able to see the alerts without requiring him to do any predetermined manual action. Ubiquity implies that the alert should be provided to all the users active on the Internet that are affected, and should not miss out on any user in the affected area. The support for second language implies that users more comfortable with a non-English version of alerts should have the option of switching over to their preferred language, in case they so desire.

It would be worthwhile examining the current state of the art in providing emergency information on the Internet. There are existing schemes that can provide notifications to a set of first responders in case of emergencies arising which may require the responder's attention. This has been extended to provide telephone notification to a general population without the need for explicit notification [3]. In such systems, all telephone numbers issued to people residing in a specific region are automatically registered so that an automated system can call the number when needed. However, these systems, even when not requiring explicit registration of telephone numbers fail to reach all of the users on the Internet. With an increasing use of IP telephony and cell-phones, the association of phone numbers to locality is no longer direct and clear. Furthermore, a person using a computer away from their home, e.g. in an Internet café, or on a public wireless hot-spot can not be reached by dialing their residence phone numbers. Thus, these types of existing notification systems would fail to meet the ubiquity requirements of an Internet EAS.

The other approach for exploiting the Internet for emergency notification is via the posting of breaking news and information at popular web sites. Users who are browsing those web sites that carry the emergency notification can visit designated web sites to obtain emergency information. However, this requires an explicit manual action on behalf of the operators of the web site (they need to create web site content with emergency information) as well as the user who needs to visit that web site explicitly. Internet users, who are involved in other actions at that time, e.g. writing a document, reading email, or running any of the other myriads of applications on their computers, would not be aware of any emergency unless they decide to use the browser and visit some of the designated web sites. Traveling users may have very little idea about the local web sites that may carry the emergency information. Thus, this approach would not satisfy the requirements of automated operations and spontaneity.

A good Internet EAS system would allow any computer user in the Internet to be alerted to an emergency situation that may be affecting it. Such a system can be built in one of several ways as described in the next section.

APPROACHES TO INTERNET EAS DESIGN

Radio and TV EAS were developed by means of regulations from the federal government on the operations of cable, TV and radio networks. The Internet is a global entity which does not entirely lie within the jurisdiction of any country government. However, it would be worthwhile to look at the possible regulations that can be brought into effect by a government to ensure a viable EAS system for users within its jurisdiction. We consider three types of regulation scenarios, one in which regulatory guidelines are specified for computer manufactures, one in which regulatory guidelines are specified for applications and web-servers located within the jurisdiction of the government, and one in which the regulations are specified for Internet Service Providers (ISP).

In an approach targeted at computers, the US government could mandate all computers sold within the jurisdiction of the United States to have a special application bundled with a personal computer that would have the ability to access an emergency alert site and display any alerts that are relevant to the geographic location in which it is located. Depending on the region and locale of the computer, the application on the personal computer may access sites operated by other governments. Such a system can be a standard part of the operating system. In this approach, regulatory guidelines will be needed for providers or common operating systems for computers, including Microsoft and Apple for personal computers, Microsoft and Palm for handhelds, and IBM, HP, Sun etc. for large server operating systems. Downloadable software can be made available for general users to enable emergency notification to support legacy systems. The technical details of how such a system would work are described in the next section.

In an approach targeted at ISPs, the US government could mandate that ISPs with customer access points within the US boundary support two modes of operation, a normal mode and an emergency mode. In the normal mode, network connectivity is provided as usual. However, in the emergency mode, which is determined by the appropriately authorized government entity, the ISPs would be required to redirect a subset of network traffic to an alert system which provides a user with information about any relevant emergencies. The technical details of the system operation are described subsequently.

In an approach targeted at web site operators, the US government can require the web site operators, rather than the ISPs, to operate in two modes, an emergency mode and a regular mode. In the emergency mode, the web site operators would be required to provide any user accessing the web site with information about any relevant emergency situation that may be appropriate for the user. Web sites

that are hosted within the physical boundaries of a country, and are large-scale can carry software that enables emergency notification in this manner.

Regardless of the approach used for implementing EAS, there are some unique features of Internet based communications which can be exploited to enhance the EAS infrastructure. Since the Internet provides a two-way communication medium, instead of a one-way medium as in traditional television, Internet EAS can allow the general public to provide feedback to government officials regarding any emergencies that exist. Furthermore, a much richer type of alert notification can be provided readily in all types of approach, which include specific information, detailed response measures, current status information and multilingual support. The Common Alerting Protocol [4] specifications being standardized within OASIS [5] provide a good foundation for rendering the alert information in a structured format.

In the next three sections, we discuss the technical details of each of the approaches mentioned above, followed by a discussion of the relative merits and demerits of each approach.

CLIENT SOFTWARE APPROACH

In the client software approach, the emergency notification service is provided in the following manner. Each of the client systems (PCs, laptops, handhelds, …) that are connecting to the Internet has a client software that can periodically access a well-known emergency site, and obtain any relevant emergency information about the location where the computer is found.

The US government (or the government in another country) would need to operate an Emergency Alert Site. Given the almost ubiquitous use of the HTTP protocol, and its ability to cross across firewalls, address translators, and other network devices, it would be logical to operate this site as a web-server. A well-known URL allows client software modules to access this site.

The client software running on a user's machine can determine the location of the web site from the locale information in the operating system. The client software may also provide a user with options to configure in more details, e.g. a street address, that provides a more precise location of a user's presence than can be determined by automated means. The client software can determine its IP address and if it is not a LAN-local IP address (i.e. in subnet 192.xx.xx.xx), include that as part of the request made to the web-server. The client software would connect to the well-know URL of the alert site to obtain information about any alerts that are relevant to it.

When contacted by the client software, the Emergency Alert Site would determine the IP address of the inbound connection from the client, and use IP to geographical mapping techniques to map it to a geographic locality. If the client has provided an IP address, and it matches the IP address of the IP connection, that address is used. If the IP address of the connection matches any of the well-known proxies (e.g. AOL) in the network, then the IP address provided by the client is used. Otherwise, the IP address of the connection end-point is used. IP to geographical mapping may be done by steps described in [6] or by using a software like IP2LL [7]. Although, the mapping mechanisms provided by such services are not always accurate, they can pinpoint an address to the approximate city or county in which a host is located. If the client software provides configuration information from the user, then the geographical location can be determined much more accurately.

After determining the location of the machine, the EAS assigns the user to one or more localities. Each locality is a geographical region to which the Emergency Alert Site can associate a URL to be retrieved and checked periodically by the client. The client software can then monitor the provided URL and display any alert information to the user of the machine.

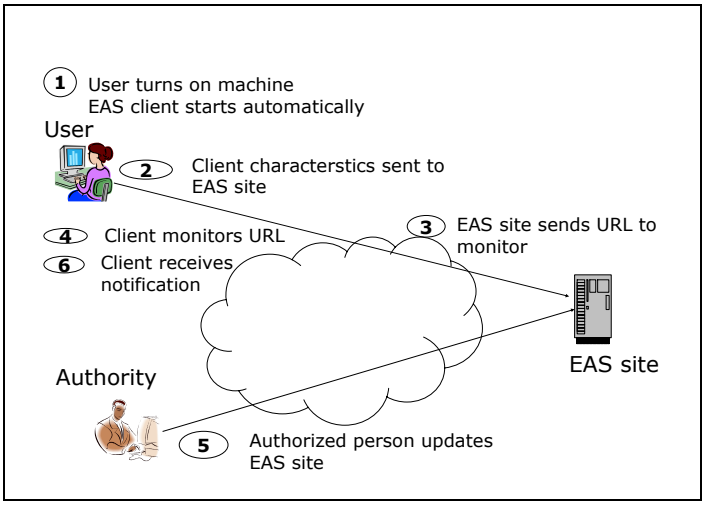

Figure 1. Client Software Approach for EAS

The operational process of the client software approach towards EAS is shown in Figure 1. The encircled numbers indicate the sequence of events. When a user turns on the machine (step 1) the EAS client starts automatically and the client characterstics are sent to EAS site (step 2). The EAS site computes the locality of the user and sends back a locality specific URL to monitor (step 3). As an example, users in New York will be sent the URL https://www.eas.com/ny_alert.xml, while users in Califor-

nia may be sent the URL https://www.eas.com/ca-alert.xml. The EAS client periodically monitors this URL (step 4). In case of an emergency an authorized personnel updates the EAS website (step 5), which results in modifying the contents of the URL being monitored, and the client automatically receives the emergency notification (step 6).

The emergency alert site operated to provide clients with the alert information needs to be operated securely and in a scalable manner. Security can be provided by restricting access only via https and authenticating the credentials of the emergency alert site to the client. The site must be provisioned so that it can handle periodic accesses by all of the active users on the Internet within a geographic domain periodically. Scalability can be obtained by means of geographical content distribution mechanisms [8]. Estimating 100 million active client machines accessing the site with 1 Kilobytes of on an average of once every 5 minutes, the site would need to support an input pipe of approximately 2.3 Gbps. By distributing the load over 5 geographically distributed sites within the US, each site would roughly need a bandwidth of less than 600 Mbps, which is easily feasible with current technology. The processing requirements can be similarly handled.

The processing and bandwidth needs of the Internet EAS site can be further reduced by leveraging peer-to-peer technologies for relaying any alert information. One way to build such a system would be to have a system analogous to the relaying mechanisms [9] used for streaming and other media. In this analogue, client systems would be used to relay alert information to other clients in the same geographical vicinity. The emergency alert site would maintain a soft registration of the clients, and provide each new client a list containing a few other clients listening on the same locality URL. When refreshing the alert information, each client randomly goes to the clients in the list, reconnecting to the alert site only if none of the specified clients are accessible. Note that public alerts relayed in this manner would need to be signed with the alert site public key for authentication.

The client based software mechanism is best at passively displaying any alert information to users on the Internet. The software can display emergency alerts regardless of the fact whether the user is actively using the computer, or is passively using the system. However, it does impose a significant scalability need on the emergency alert site that needs to be put into operation.

ISP CENTRIC APPROACH

Traditional ISPs provide connectivity to their consumers, either at home or at workplace, by means of routers that provide IP level connectivity. The ISP is responsible for assigning the IP address to any access router that is providing the connectivity to any computer at that site. As a result, by using the IP address of the access router, and the address at which an access router is located, an ISP is able to map the geographic location of its customers fairly accurately.

This knowledge of the geographic location of machines by an ISP provides a simple way to alert users of any impending geographies in the geographical vicinity. The ISP could configure each router to operate into two modes, a normal mode, and an emergency mode. In the emergency mode, the router is configured so that a subset of network traffic is directed to a special application proxy. The proxy is responsible for displaying the alert information to the user.

Either a regular proxy server or a translucent proxy server [10] may be used for the display of emergency alert information. Translucent or transparent proxy servers may be operated without requiring any modification in the normal operation of the ISP. Regular proxy servers are also used often by service providers such as AOL who can provide similar functionality. Some ISPs may have ready access to users' geographic locations, other possible proxy sites may determine geographic location from users' IP addresses as discussed in the previous section. In either case, the geographic location of a user may be cross-referenced to a list of posted local alerts. Having identified a user as being within a geographic location for which an IEAS alert has been posted, the proxy may then redirect a URL request (or a request for a target application) from such a user's browser to a public alert site displaying an urgent public announcement for the user's geographic location. Users whose geographic locations are not associated with a public alert would not have their browsers redirected but would instead access the URL (or target application) as originally requested. The operational mode is illustrated in Figure 2.

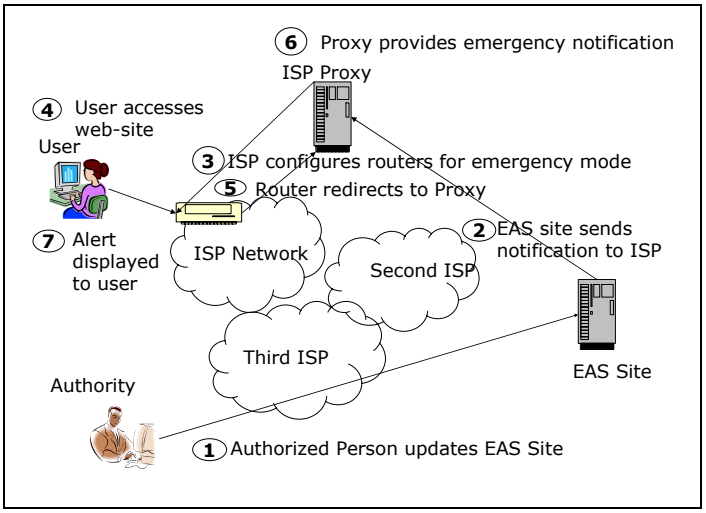

Figure 2. ISP Centric Approach for EAS

The ISP centric approach has certain advantages over the client software approach. The geographical location of users can be determined relatively easily using the identity of the access point. This is more accurate than using IP to geographical location mapping algorithms. Large ISPs may be more amenable to validate compliance from a regulatory perspective, and there is no software installation needed on a client machine with some associated issues of failures or upgrade mechanisms that are needed. ISP based notifications mechanism can be extended with IP telephony technologies to provide call-back notifications to cell-phone users bound to a cell-phone access. An approach for SIP based EAS for phone notification is described in [11]. However, the ISP based mechanism will only be activated when a user accesses a new web-page, or otherwise invokes an operation which results in the transmission of a packet on the network. A user who is simply running a simulation on his machine will not be alerted if he is not web-surfing in parallel. Translucent proxies are available only for certain types of applications, and therefore the scheme would work only for the types of applications where a translucent proxy support may be practical.

WEB SITE CENTRIC APPROACH

The third approach to Internet EAS service that we consider requires web sites that exceed a specified threshold of hits per day to support EAS notification. An EAS-enabled web site would operate in either a normal mode or an emergency mode. The web site would determine the geographical location of the user accessing their site using the IP address of the inbound connection. It would identify the geographical location on the basis of that information, and check whether an alert is in store for the locality from which the user is accessing the site. This information can then be displayed to the user in a manner that informs him of the emergency and provides a link to a site with more information.

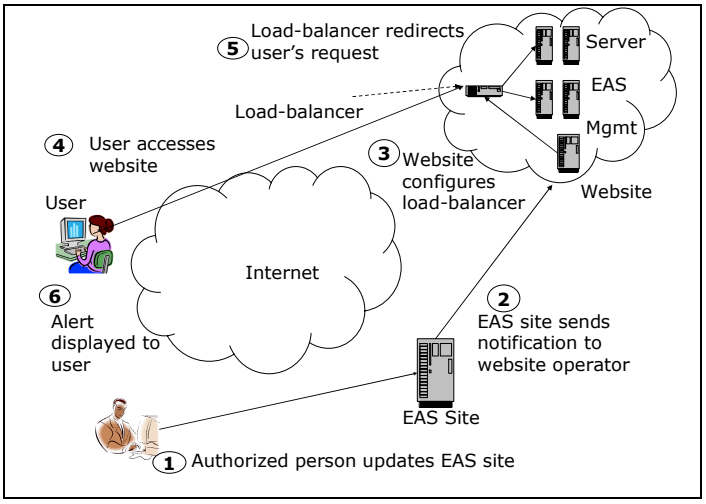

Figure 3. Web Site Centric Approach for EAS

Most large web sites operate more than one server in order to meet their scalability requirements. One option for implementing an EAS is for each large web site operator to have some of these servers running emergency alert software. The emergency alert software has two components, a management component and a proxy component. The proxy component could be implemented either as a cgi-bin script, servlet, or a Java Server Page (JSP), and is responsible for displaying the alert information to the user if needed. The management component is responsible for managing the configuration of the load balancer located in front of the servers. A web site would typically have a front-end load-balancer to distribute incoming requests among themselves. The management component may compute the IP address/subnets of any locality that may be possibly affected by the currently active set of alerts, and configure the front-end load balancer to direct those requests to servers running the emergency alert software. The operational mode for a web site centric approach is shown in Figure 3.

COMPARISON

In this section, we compare the relative merits and demerits of each of the different approaches for implementing an EAS system. We compare the different approaches on the basis of locality, automated operation, non-intrusiveness, spontaneity, regulatory ease, management overhead and security perspective.

From localization perspective, the best approach would be the ISP centric approach using the geographical location of a customer's access. While the client-side software approach where a user can configure their street address would have a higher accuracy for computers which are immobile, mobile laptops used by travelers would require reconfiguration during travel, which would not be very desirable. The web site centric approach or the client-side

software approach relying on IP to geographical address mapping would have lower accuracy due to the current state of the technology.

All of the systems discussed above show characteristics of automated operation, where the only manual step is the creation of an alert by an authorized user at an alert web site. Some ISPs and web site operators would have reservations about switching from a normal mode to an emergency mode without a human in the loop, and may prefer to make the switch manually even if automation is possible.

The ISP-centric approach and web site centric approach are non-intrusive and require little change in the normal operation of their systems. The client-side software approach requiring specific software is more intrusive, but should not interfere with the regular operation of the system.

From a spontaneity perspective, the client-side software approach is probably the best. The system would provide users alert notification within a few minutes even if the computer is not being used actively. The ISP centric and web site centric approach require the user to take an active step, e.g. access a web site before they can view the alert. These approaches can be made more spontaneous by providing a browser toolbar which can be used to perform the functions of client-side software.

From a ubiquity perspective, the client side software approach is again the best. It does not depend on a specific application being invoked by the user. The ubiquity of the ISP-centric approach depends on the number of applications for which proxy services can be deployed meaningfully, while the web site based approach is applicable only for web browsers. Although web-surfing is the most common usage of the Internet, there are other applications, e.g. peer to peer systems or IP telephony, which are also growing in significance.

From a regulatory perspective, the ISP based approach is the easiest. The physical placement of access routers can be used as a criterion to determine jurisdiction by the government. Furthermore, the number of medium-scale and large ISPs operating within the US boundaries is a few hundreds, as compared to thousands of prominent web sites. The number of operating system manufacturers and computer retailers are fairly small, and the government could mandate that emergency alert software be included in new computer systems. Since computers are sold internationally, issues related to government jurisdiction may arise. Moreover, any regulation requiring proper operation of the client PC software would not be feasible to enforce. The properties from a management perspective are analogous.

Security management would also be a significant issue in an EAS, and the system should be protected against pranksters and hackers who may subvert it to create public panic. While adequate security safeguards can be designed, the operational mode of the system exposes each system to security threats. The security threat is largest for the client-side software approach, less so for web site centric approach and the least for the ISP-centric approach. Web sites are only susceptible to being hacked into, and most sites exercise sufficient preventive measures. ISP networks are designed to be more secure and the EAS proxies on ISP networks can be made fairly secure. General public computers, however, are susceptible to several attacks arising both due to software vulnerabilities, and the inability of a layman to configure the computers properly.

As would be apparent from the discussion above, there are no clear winners among the three different approaches. Each of the three approaches for emergency alert notifications can be made to work. By examining all the different issues, the authors feel that a client-software approach may be the best possible one, followed by the ISP-centric approach and the web site centric approach.

MILITARY APPLICATIONS OF INTERNET EAS

While we have described the Internet EAS as primarily applied to a civilian emergency scenario, the same technology can be used for interesting applications in the military context. As the military moves towards achieving its goal of "network-centric operations" [12], an increasing number of soldiers will be relying on IP connected devices to perform crucial functions in their military task. An alert notification system, suitably implemented, would be valuable in alerting the soldiers to any emergency situation, probably including instructions for reacting to that emergency as well. The emergency notification system would work well in military operations located in foreign environments which are prone to attack by insurgents. Military troops in such environments may be operating in a distributed camps connected via an IP network, and could be informed quickly of an alert on their computers in an automated manner, and the notification may include directions that are customized for specific units.

A similar usage model can be envisioned for military installations that are used for peacekeeping and humanitarian missions. Any unit, when notified promptly of an emergency situation, can react effectively as long as a plan of action to react to such emergencies is in place.

REFERENCES

[1] Federal Communications Commission Emergency Alert System Fact Sheet, http://www.fcc.gov/ebeasfacts.htm.

[2] Norman H. Nie, “How do Americans use the Internet in their daily lives”, http://www.stanford.edu/group/siqss/SIQSS_Time_Study_04.pdf.
[3] Teleworks launch of First Responders, http://www.house.gov/boucher/docs/teleworks.htm, Sept. 2003.
[4] A. Botterell, “Common Alerting Protocol, Version 1.0”, OASIS Standard 200402, March 2004.
[5] Organization for the Advancement of Structured Information Standards, http://www.oasis-open.org.
[6] How do I find the geographical location of a place, given its IP address? http://www.private.org.il/IP2geo.html
[7] IP to Latitude Longitude Mapping, Software available at http://www-unix.mcs.anl.gov/~olson/IPtoLL.html.
[8] CDN Reference
[9] S Jin and A Bestavros, “Cache-and-Relay Streaming Media Delivery for Asynchronous Clients”, Proceedings of Networked Group Communication, Boston, MA 2002.
[10] P. Rodriguez, S. Sibal, and O. Spatscheck, "TPOT: Translucent Proxying of TCP", Technical report TR 00.4.1, AT&T Research Labs, 2000. http://citeseer.ist.psu.edu/rodriguez00tpot.html.
[11] H. Schulzerinne and K. Arabshian, “Providing Emergency Services in Internet Telephony”, IEEE Internet Computing, May June 2002. pp. 34-47.
[12] F. Stein, J. Garska and P. McIndoo, "Network Centric Warfare: impact on army operations", EUROCOMM 2000, Information Systems for Enhanced Public Safety and Security, Munich, Germany, 2000.