Global Interoperability of National Security and Emergency Preparedness (NS/EP) Telecommunications Services

Frank J. Suraci
Technical Director
National Communications System
Arlington, VA 22201

Arye R. Ephrath
Fellow, IEEE
Mythology, Inc.
Clifton, VA 20124

John R. Wullert II
Member, IEEE
Telcordia Technologies
Piscataway, NJ 08854

Abstract - Emergency telecommunications services, which provide priority access to public telephony for those responding to emergency situations, have previously been offered in the U.S. and other countries around the world in the public switched telephone network and more recently in cellular telephone networks. This paper describes the implementation of these national security and emergency preparedness (NS/EP) services in a next-generation network (NGN) Internet-Protocol Multimedia Subsystem (IMS) environment. The implementation was tested during a global interoperability event. The tests' results, and implications of the analysis of the test data to future developments, are discussed.

I. INTRODUCTION

The Government Emergency Telecommunications Service (GETS) was established in the United States in the mid-1990s as a result of an Executive Order of the President [1]. The purpose of this program has been to provide federal, state, and local officials, as well as other authorized personnel associated with national security and emergency preparedness (NS/EP) with the means of obtaining priority in placing telephone calls via the public switched telephone network (PSTN) during times of emergency or when the telephone network is otherwise congested.

In the wake of the terrorist attacks on the U.S. of 11 September 2001, the GETS program was expanded to include a wireless priority service (WPS), whereby authorized NS/EP callers may gain priority in accessing the cellular network and placing priority calls by dialing a special prefix.

The GETS and WPS programs are managed by the National Communications System (NCS) [2], an agency of the U.S. Federal government. The NCS contracts with commercial wireline and wireless telephony carriers to provide GETS and WPS priority services on their networks. The NCS also authorizes potential GETS and WPS users and issues each user a unique personal identification number (PIN), which can then be used to place priority telephone calls in a manner similar to placing a calling-card call.

The priority of a GETS call in the PSTN is indicated by the user through the dialed access number and signaled through the network through the settings of specific parameters in Signaling System 7 (SS7) set-up messages. High probability of call completion of priority-signaled calls in a damaged or congested telephone network is achieved through a set of specialized treatments including routing calls to certain GETS-enabled public carriers, queuing calls to overcome temporary blocking situations, additional re-routing attempts to overcome link damage or route congestion, and exemption from restrictive network controls. In addition, SS7 signaling message associated with GETS calls are themselves routed with higher priority than are signaling messages associated with non-priority calls.

The number of authorized GETS users has been growing steadily; at present there are about 150,000 NCS-authorized GETS/WPS users. GETS/WPS users include officials of the federal, state and local governments, community and business leaders, and first responders such as fire-fighters, medical personnel and law-enforcement organizations.

In anticipation of the emerging migration of the public, circuit-switched telephony networks to Internet-Protocol (IP)-based packet infrastructure, the NCS embarked on an effort to develop and implement priority capabilities in the packet-network environment similar to the priority capabilities currently available in the circuit-switched PSTN. These efforts address not only voice telephony (thus extending the existing GETS features into the new infrastructure) but also take advantage of the high bandwidths and intelligent terminals associated with these next-generation networks (NGN) to define additional session-based priority telecommunications services, such as video teleconferencing over the managed public IP networks.

In recent years the NCS has been prototyping and analyzing alternate network configurations designed to achieve its objectives. These prototypes, based on existing and emerging international standards and on commercially-available networking equipment, have been used to test and demonstrate several methods of authorizing and processing priority sessions in the next-generation, IP-based network. The remainder of this paper describes the priority features and their implementation in these prototypes, the tests that the NCS conducted along with 26 commercial vendors in the context of a global, 3-continent, IP network, the results of these test and the implications of the analysis of the test data to future developments.

II. THE NETWORK

The Internet-Protocol Multimedia Subsystem (IMS), defined by the Third Generation Partnership Program (3GPP), is a functional architecture for implementing connection-oriented communication services over an IP network infrastructure. An IMS communications network is divided logically into layers: A communication service's payload (e.g., voice, video, text messages) is transported via the *transport* or *bearer layer*; end-to-end connections and sessions are set up in the *signaling* or *control layer*; and service-specific logic and processing are provided by the *application* or *service layer*. The separation of the network into layers is only a logical one; the packets and messages associated with the various layers in fact share the network's physical resources.

Ensuring that NCS's priority NS/EP sessions have a high probability of being established and maintained requires special features in each of the network's layers: The authority and identity of the originator of a priority session are verified by an application server (service layer) prior to connecting the session, the session is established with special priority features (control layer) and the payload is transported with elevated priority through the network (transport layer). The priority functions of all three layers have been prototyped, tested and demonstrated by the NCS. Some of this testing occurred during a recent Global Interoperability test event sponsored by the MultiService Forum (MSF) [3] that was designed to demonstrate deployable, operationally-ready, end-to-end IP services and networks.

The test event, known as "Global MSF Interoperability" (GMI2006) event, brought together over two dozen network equipment manufacturers and commercial carriers, and spanned four major test centers (see Figure 1). The test centers were located in the research laboratories of four global telecommunications carriers (the UK's British Telecom/Vodafone, Korea's Korea Telecom, Japan's Nippon Telephone & Telegraph and the U.S.'s Verizon), with a fifth, smaller test center provided in the U.S. by the Interoperability Laboratory of the University of New Hampshire. The network complied with the MSF R3 service architecture [4] which, in turn, is compatible with the 3GPP IMS functional architecture. Prior to the commencement of testing, the participants developed a series of mutually-binding implementation agreements. Based on established international standards, the implementation agreements specified in detail the specific functions and interfaces to be employed in GMI2006. The overall results of the GMI event are described in a white paper produced by the MultiService Forum [5].

III. TESTS AND RESULTS

NCS prototyped and demonstrated several priority services and features during GMI2006 under two overall scenarios: the first scenario ("Scenario 2") represented a network where all session participants were connected within a single network domain, while the second ("Scenario 5") addressed delivery of services to nomadic session participants roaming in network domains other than their own service provider's.

NCS demonstrated the signaling of priority status during a session's initial set-up, multiple methods of performing the on-line authentication of the session originator's identity and priority privileges, and the realization of priority treatment of properly authorized sessions via preferential admission. These were demonstrated for both telephony (priority voice) and multimedia (priority video teleconferencing) services. In
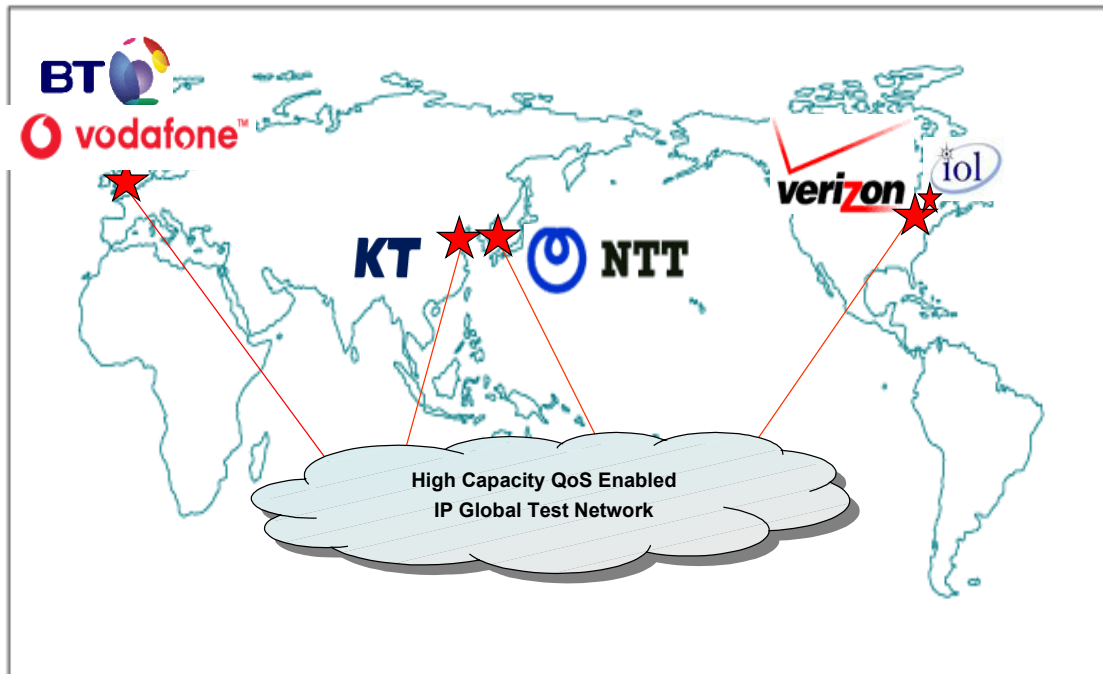

Figure 1: Topology of GMI2006 Test Centers

addition, the NCS demonstrated the interoperability of priority features between an IP-based IMS network and a circuit-switched legacy network, as well as an "origination identification restriction" (i.e., anonymity) feature. These tests are described in more detail below.

A. Resource Priority

The session initiation protocol (SIP), the signaling protocol used to establish communication sessions in IMS networks, has been extended in 2006 to support a "resource priority" header (RPH) [6]. The presence of a properly authenticated RPH in SIP messages conveys a priority status and the need for priority treatment in the establishing and maintaining associated session. In addition, a standard mapping exists [7] between the RPH in a SIP-based VoIP network and the priority indicators in SS7 call set up messages in the circuit-switched PSTN. As shown in the sample SIP message at right, the terminal indicates the need to have RPH supported with the SIP "Require" header and indicates the priority requested using the "Resource-Priority" header with a namespace ("ets") and a corresponding value ("0")

Several tests conducted by the NCS during GMI2006 were designed to demonstrate the generation and inclusion of a proper RPH in the signaling messages during the establishment of a priority NS/EP session. In addition, the behavior effected in network elements by RPH-carrying SIP messages was verified, as the presence of an RPH is expected to invoke specific actions in certain network elements, while all network elements are expected to propagate the RPH indicator downstream unchanged.

```
INVITE sip:6172802222@10.104.26.250:5060 SIP/2.0
SIP/2.0/UDP 10.80.2.10:5060; branch=z9hG4bK0dB0
From: "6172801111" <sip:6172801111@10.104.26.234>;
      tag=gK0d0144c1
To: "6172802222"
<sip:6172802222@10.104.26.235:5060>
Call-ID: 17629360_1095@10.80.2.10
CSeq: 18585 INVITE
Max-Forwards: 70
Require: resource-priority
Resource-Priority: ets.0
Contact: "6172801111"
<sip:6172801111@10.80.2.10:5060>

Content-Type: application/sdp
Content-Length:  ...
```

Of special significance is the fact that these tests were conducted with commercial global equipment vendors, on a test network comprised of commercial, standards-compliant devices. The positive results of these tests enhanced the NCS's confidence that the next-generation network will be able to support NS/EP priority communications schemes.

B. Authorization

To ensure that the use of priority capabilities is restricted to NCS-authorized users, users must be authorized prior to session establishment. In the circuit-switched PSTN the user's identity is authenticated via the touch-tone entry of a 12-digit PIN; this authorization method, which must also be supported in the IP-based NGN for reasons of backward compatibility, was demonstrated during GMI2006. Earlier NCS testing had demonstrated that the transport of digits as dual-tone, multi-frequency (DTMF) audio signals over IP-based protocols
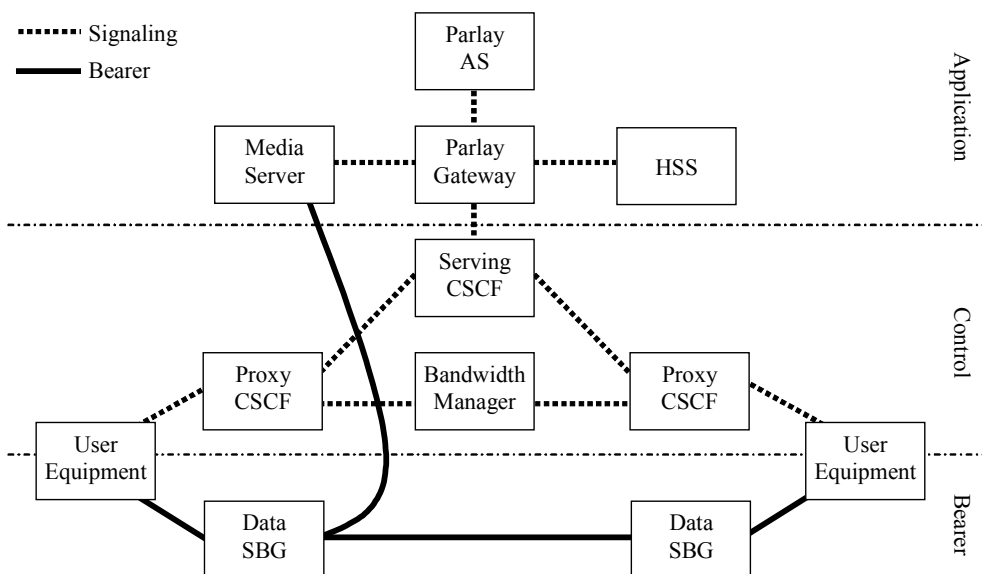

Figure 2: Partial MSF GMI Network Architecture indicating layering

poses a special challenge since dropped packets (which can be common in a congested network) can result in the corruption of the transported digits at the receiving end. Specific means to overcome this difficulty have been devised and standardized [8].

Two additional authorization methods were demonstrated. The *profile-based* method relies on the session originator's profile, stored in the home subscriber server (HSS) in the subscriber's home network, to verify the originator's identity and privileges during registration, followed by validation during call setup. In this method, the user terminal provides an encrypted digest of portions of registration message concatenated with the user ID and password. The application validates the user by calculating the same digest and comparing the results. In the *challenge-response* method, the originator's intelligent terminal is challenged to provide authorization data when attempting to place a priority call, using HTTP digest authentication [9]. The challenge includes an application-generated nonce as a security measure. The intelligent terminal then creates the security digest based on the message parameters, the nonce and the user ID and password. This digest is processed by an application server to authorize the user. This SIP messages used in this challenge-response operation are illustrated in the upper highlighted

section of Figure 3.

All three authorization methods require the use of a specialized GETS application which ran on a Parlay [10] application server and was accessed via a Parlay gateway. All three authorization schemes were demonstrated to function successfully.

The Parlay gateway implements the Parlay programming interfaces, a standardized set of APIs that provide protocol independent control of capabilities used to construct communications services. For the purposes of these experiments, the Parlay Multiparty Call Control (MPCC) API [11] was used to implement the authentication features and the Parlay User Interaction API [12] was used to prompt the user and collect digits for the in-band PIN scenarios. The SIP RPH value was mapped onto the High Probability Completion field in the MPCC API to communicate the priority information to and from the application. The generic information fields within the MPCC API were used to communicate the information used to perform the digest authentication. The application interacted with stored information (i.e., the customer profile) to validate the user information submitted either as touch-tones or SIP parameters and would then allow or reject the call as appropriate.
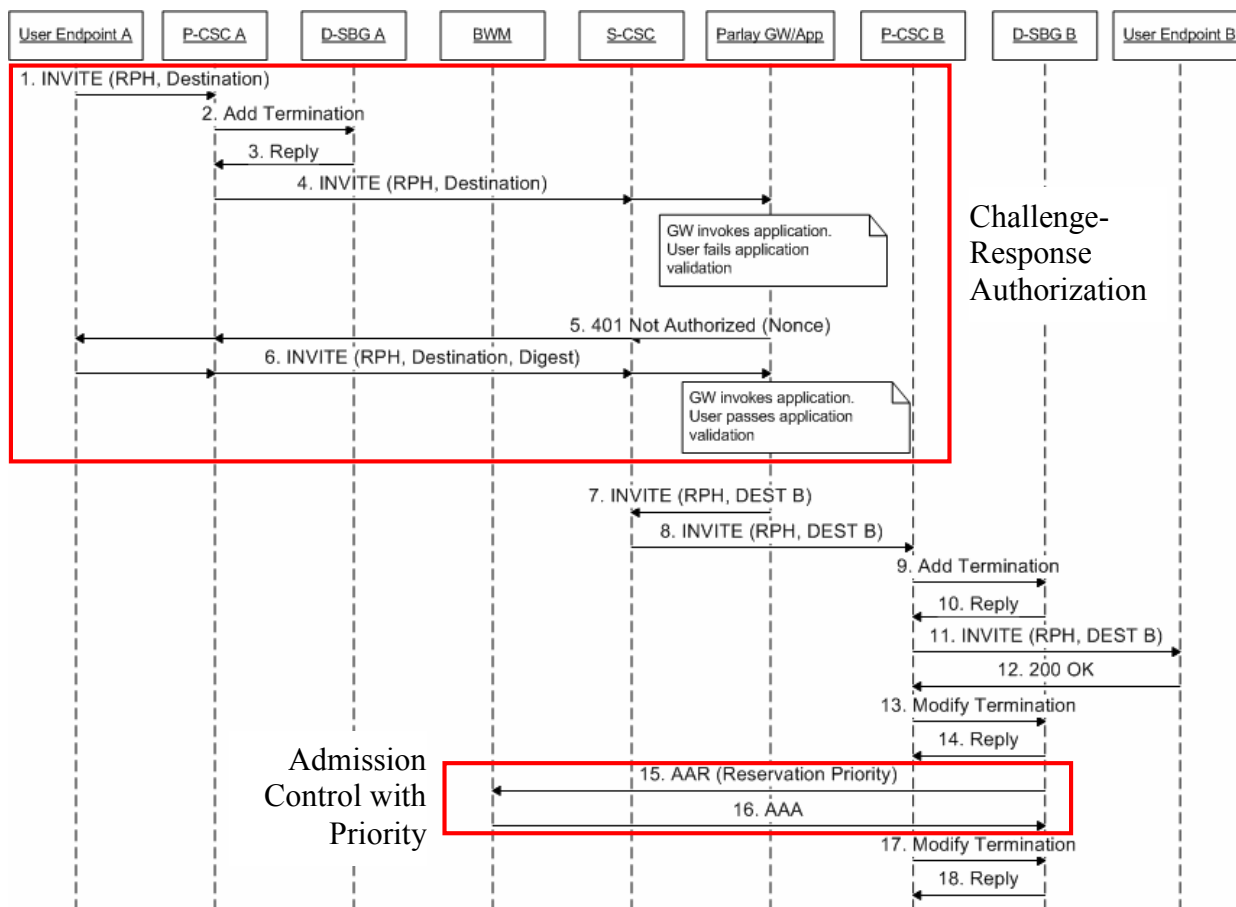

Figure 3. Partial SIP signaling flow for priority video call using challenge-response authorization

C. Connection Admission Control

At the session set up layer, priority treatment is provided by the policy decision function (PDF) at the network's border gateways or by a "bandwidth manager" network element. As the network's available bandwidth or other resources near exhaustion, new session set-up requests for non-priority sessions may be denied during an admission control process. As demonstrated at GMI2006, however, the bandwidth manager can utilize the value of the RPH parameter to discriminate among set-up requests. In particular, as illustrated in the highlighted section of Fig. 3, the RPH value was translated from the SIP message to the Resource Priority field of the DIAMETER-based bandwidth manager interface. Alternatively, the bandwidth manager can interpret the RPH value directly if the bandwidth manager is in the SIP signaling path. The bandwidth manager is configured to assign a higher priority to calls with an RPH value set. This higher priority is then taken into account during the admission control decision. Consequently, session set-up requests for NS/EP priority sessions will be granted up to and beyond the capacity limits defined for non-priority calls. During the testing, the bandwidth manager was normal calls were established up to the limit configured within the bandwidth manager. Additional non-priority calls were rejected, while priority calls were admitted.

More complex bandwidth management rules could limit priority calls to a specific fraction of the available bandwidth and/or to differentiate between priority calls with different values set in the RPH field.

D. Transport Priority

Priority treatment may be provided to data packets in the transport layer in several ways, such as special markings of each packet (which provides the packets with priority handling at the network's routers), reserving a portion of each link's bandwidth for priority traffic, or establishing separate, reserved priority routes through the network. This priority transport treatment must be provided not only to the bearer packets which carry the voice, video, or data associated with the service, but also (and perhaps more importantly) to the packets which carry the signaling messages used to control each priority session. This capability was not tested during the current effort because the emphasis of the GMI event was signaling interoperability; testing systems under load was not part of the overall test plan. However, initial prototyping and testing of transport priority had been the focus of an earlier effort and are also planned to be investigated more thoroughly in the near future.

E. PSTN Interoperability

During the expected transition from circuit-switched public telephony to the IP-based NGN, the two technologies will co-exist and interoperate. In this hybrid environment it is imperative that NS/EP priority services function seamlessly on calls between circuit-switched and IP domains.

Interoperability implies that priority sessions which originate in one domain and cross or are terminated in the other domain maintain their priority markings and priority status. Signaling and transport gateways between these two domains must perform appropriate protocol mappings, including the priority indicators, because priority is signaled via SS7 parameters in the PSTN and via the SIP RPH in an IMS domain. The proper mappings between SIP RPH and the corresponding SS7 messages were demonstrated during GMI2006 in both directions (SS7 to SIP and SIP to SS7). For example, the mapping rules from SS7 to SIP depend on the values of the Called Party Number, the Calling Party Category (used to indicate priority) and the precedence parameter (used to carry priority level). The portion of mapping decision table (from [7]) is included at the end of this paper.

F. Video Teleconferencing

For the past decade, the GETS program has provided priority in placing voice telephone calls via the PSTN. The ongoing migration to the NGN infrastructure has provided the NCS with an opportunity to develop and implement additional session-based priority telecommunications services, such as priority video teleconferencing.

The establishment of a priority video teleconferencing session requires the proper authentication of the originator's identity and privileges, in a manner similar to GETS's. One difference is that a video conferencing terminal cannot be assumed to have the capabilities to generate touch tone signaling, so the in-band transmission of PINs was not tested in this case. The priority treatment of video connections includes connection admission control and transport priority. These were verified and demonstrated by NCS at GMI2006. It should be noted that the SIP call flows for voice and video sessions look nearly identical; the sole difference is in the session description protocol, contained in the body of the SIP messages that describes the media types and codecs to be used for the session.

Beyond the current testing, the establishment of a multimedia priority session also presents additional options, such as a trade-off in a congested, bandwidth-limited network between the bandwidths allocated to the video and audio components of the priority multimedia session. While the NCS has not yet formulated a position on such additional options, various alternatives have been considered. One approach is to separate the bearer traffic associated with video from that of audio, providing different priority levels to each. An alternative is to adjust the allocation of bandwidth during session establishment by controlling the selection of audio and video codecs. While it might be assumed that the audio is more important than the video, this is not true in all situations. In-depth usability analysis is required to ensure that the resulting system fulfills the communication needs of the NS/EP user community. Future work will need to determine the proper trade-off between the video and audio channels.

G. *Session Description*

The issues described above can be made clearer through an example of a user attempting to establish a priority video session. Fig. 3 illustrates the initial portion of a SIP flow for a priority video session using challenge-response authorization. The user, interacting with priority-aware User Endpoint, enters a destination address or number and indicates that the call should be given priority. The User Endpoint, which had previously registered with the home network through a local proxy, sends an SIP invite through the Proxy CSCF (Call Session Control Function) toward the home Serving CSCF. The invite includes a Resource Priority Header. Upon receiving the SIP invite, the Serving CSCF recognizes the presence of the Resource Priority Header and directs the invite to the appropriate application server, in our case a Parlay gateway. The Parlay gateway alerts the application code which proceeds to authorize the user. In profile-based authorization, the initial SIP invite would include sufficient information to properly authorize the user. In the challenge-response authorization shown in Fig. 3, the application initially rejects the call with instructions to the User Endpoint on how to generate the appropriate information. After successful authorization, the application forwards the SIP invite toward the destination.

After the parameters describing the video session have been exchanged through the SIP signaling, the Proxy CSCFs at each end make admission control requests to the Bandwidth Manager. These requests include the priority information that has been transported via the SIP Resource Priority Header. The Bandwidth Manager incorporates the priority into the admission control decision, admitting priority calls even in situations where normal calls would be rejected. After successful admission control, the Proxy-CSCFs instruct the Data Session Border Gateways to open paths for the media, and to mark the media packets for priority treatment, either through DiffServ code points or to dedicated MPLS paths. At this point, the session is established and the users can communicate with both audio and video.

IV. DISCUSSION

The NCS has been prototyping and analyzing alternative network configurations designed to achieve its objectives, i.e., to provide both existing and new priority services in the NGN. These prototypes, which are based on existing and emerging national and international standards and on commercially-available networking equipment, have been used at GMI2006 to test and demonstrate several methods of authorizing and processing priority sessions.

The demonstrations of priority services and features during GMI2006 illustrated several critical points in the path toward implementation of these services.

Interoperability: The priority calling tests were complex, in terms of both the network elements involved and interfaces utilized. Some of the scenarios incorporated 12 different network elements from nine different vendors. In addition, many test cases were executed multiple times, incorporating equipment from different vendors. Successful interoperation was achieved with only limited reconfiguration of elements for each test. The fact that this interoperation occurred across using seven different interfaces specified by MSF implementation agreements is a strong indication that the capabilities required to implement priority services are becoming available in commercial equipment.

Priority Signaling: Despite the fact that the RFC defining the Resource Priority Header was only published in February of 2006, multiple vendors had implemented the ability to set, pass, trigger on and base decisions on the information carried within this header in time for GMI2006 in October. One shortcoming that was noted in this regard is that there are currently no audio or video clients capable of generating and populating the RPH field. For purposes of GMI2006, this was addressed by creating a custom SIP-proxy wrapper that enhanced commercial SIP clients with priority signaling capabilities. This highlights the speed and flexibility with which SIP implementations can be enhanced.

Multimedia: The IMS-based network infrastructure used during the GMI2006 event was inherently capable of handling both audio-only and audio-video calls. No special configurations or additions were required to support multimedia sessions. While there are, as discussed above, several possible ways in which multimedia sessions could be further optimized, it is clear that the emerging commercial capabilities are fully capable of supporting multimedia sessions. As these capabilities become widely available, the need for priority in establishing multimedia sessions will grow.

Advanced Terminal Capabilities: In addition to multimedia, the capabilities of intelligent terminals were leveraged during the GMI2006 testing in providing alternatives to in-band, touch-tone based signaling for user authorization. The enhanced functionality of the terminal devices facilitates an improved user interface, enabling one-click indication of the need for priority and eliminating the need to enter long PIN sequences multiple times, all while improving system reliability and overload tolerance.

In the course of prototyping and demonstrating priority services, the NCS has gained invaluable insights into the operation of the NGN. In addition, the NCS has achieved several secondary objectives as well, such as:

- raising awareness among telecommunications carrier and network equipment manufacturer communities of the need to incorporate NS/EP priority features in their networks and equipment

- participating in the detailed specification of network interfaces to support priority services under overload conditions. Such specifications will eventually influence the development of international standards

- demonstrating the feasibility of providing priority treatment to voice and multimedia sessions, of

transporting digits reliably through a congested packet network, and of incorporating the specialized GETS service logic in a third-party Parlay application server.

- expanding the potential scope of its priority services from voice telephony to IP-based multimedia, video teleconferencing, and beyond.

V. CONCLUSION

The demonstration of these priority services made several points clear and provided an opportunity for learning. A key observation that relates to the financial burden associated with deploying priority services is that such services can be added to current NGN architectures and network elements with only minor additions/modifications. The offering of priority service does not entail a large scale re-design of the underlying architecture or network elements. This is especially true when priority services are considered up front, rather than as an after-thought. Dealing with priority services during the process of specifying standards for network and functional behavior simplifies the design and implementation of the required capabilities. A related observation is that priority services designed for public service purposes can be built on features that can support commercial services as well.

These points imply that neither the financial nor the technical barriers to the deployment of priority services in the next generation network are prohibitively high.

In the course of the testing, new things were learned and existing hypotheses were confirmed. We explicitly set out to explore multiple authorization methods. These mechanisms, which leverage features that are a standard part of emerging network architectures, were not only demonstrated, but they also proved that they could be supported simultaneously within a single signaling environment. It was also demonstrated that these authorization methods, and the other basic capabilities used for priority voice, were readily extensible to video services. The process of demonstrating these capabilities also allowed us to explore some aspects of the user experience. It became clear that the addition of priority capabilities to applications must take into account not only the application's purpose but also the users' requirements.

These observations will now be fed into the process of working with the communications industry to develop an agreed-upon set of requirements for the functions and features of next generation priority services.

Table: Portion of SS7 to SIP Mapping [7]

SS7 ISUP	SIP
CdPN = Destination Number CPC not equal to NS/EP Call No PRECEDENCE PARAMETER This is a normal (non-ETS) call	R-URI/To: = Destination Number
CdPN = ETS-AN CPC = NS/EP Call No PRECEDENCE PARAMETER	R-URI/To: = ETS-AN RPH [ets.DF]
CdPN = Destination Number CPC = NS/EP Call No PRECEDENCE PARAMETER	R-URI/To: = Destination Number RPH [ets.DF]

CdPN: Called Party Number   **CPC**: Calling Party Category   **RPH:** Resource Priority Header   **ets.DF**: default value in ets namespace

REFERENCES

[1] Executive Order 12472, 3 April 1984
(http://www.fas.org/irp/offdocs/eo/eo-12472.htm)

[2] The National Communications System, http://www.ncs.gov/

[3] The MultiService Forum, http://www.msforum.org/

[4] "MSF R3 Architecture,"
http://www.msforum.org/techinfo/approved/MSF-ARCH-003.00-FINAL.pdf

[5] "GMI 2006 Validating IMS: The practical implementation of converged
networks," http://www.msforum.org/interoperability/wp2006.shtml

[6] H. Schulzrinne and J. Polk, "Communications resource priority for the
session initiation protocol," *IETF RFC-4412*, February 2006

[7] "Support of emergency telecommunications service (ETS) in IP network,"
ATIS-PP-1000010.2006

[8] H. Schulzrinne and T. Taylor, "RTP Payload for DTMF digits, telephony
tones, and telephony signals," *IETF RFC-4733*, December 2006

[9] J. Franks et al., "HTTP Authentication: Basic and digest access
authentication," *IETF RFC-2617*, June 1999

[10] The Parlay Group, http://www.parlay.org/

[11] "Open service access (OSA); Application programming interface (API);
Part 4: Call control; Sub-part 3: Multi-party call control SCF (Parlay 4),"
http://parlay.org/imwp/download.asp?ContentID=9748, (2006)

[12] "Open service access (OSA); Application programming interface (API);
Part 5: User interaction SCF (Parlay 4),"
http://parlay.org/imwp/download.asp?ContentID=9754, (2006)