

[draft-schulzrinne-ieprep-resource-req-00.txt](#)

April 21, 2002

Expires: October 2002

Requirements for Resource Priority Mechanisms for the Session Initiation Protocol

STATUS OF THIS MEMO

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>

To view the list Internet-Draft Shadow Directories, see <http://www.ietf.org/shadow.html>.

Abstract

This document summarizes requirements for prioritizing access to PSTN and proxy resources for emergency preparedness communications using the Session Initiation Protocol (SIP).

1 Introduction

For voice-over-IP or, more generally, multimedia calls placed by authorized users during emergencies, there are four cases where certain calls may need to be treated with priority:

1. If PSTN gateway resources are scarce, authorized calls should be able to queue for resources and acquire outbound trunks with priority or, if local policy permits, preempt

an existing call.

2. If a call originates from a circuit-switched network with multiple priority levels and traverses a VoIP network using SIP signaling on its way to another circuit-switched network with multiple priority levels, it is desirable that the priority information is not lost during transit. Similarly, calls originating on a VoIP terminal and terminating in the PSTN also need to be able to invoke the appropriate priority mechanism.
3. Resources in networks may be managed based on SIP calls. (This is not always desirable or feasible due to the divergence of media and signaling paths and the limited ability of session setup messages to express traffic characteristics, but it may be the only available mechanism absent a resource reservation protocol.) In such cases, policy may grant emergency preparedness or MLPP calls prioritized access.
4. In some circumstances, it may be appropriate to prioritize access to signaling resources (e.g., CPU cycles) in SIP proxies.

Thus, SIP signaling requires a mechanism that allows the caller to indicate such priority.

In the PSTN and certain private networks, such as those run by military organizations, calls are marked in various ways to indicate priorities.

Below are some requirements for such a mechanism:

Not specific to one domain: The mechanism should not be specific to one country or particular priority mechanism. For example, there are currently at least four priority schemes in widespread use: Q.735, with five levels, the U.S. defense network and NATO with five levels, the United States GETS (Government Emergency Telecommunications Systems) scheme with implied higher priority and the British GTPS system ???.

Usage of existing namespaces: Since there are a number of existing resource priority schemes for both the PSTN and private networks, it is likely that hybrid or VoIP networks will inherit these naming schemes.

Extensibility: New naming schemes may need to be defined as

resource priority mechanism are applied in additional private networks, or if a VoIP-specific priority scheme is being defined.

Separation of indication and policy: The indication should not request a particular detailed treatment, as it is likely that this depends on the nature of the resource and local policy.

Default behavior: Network terminals configured to use priority scheme A may occasionally end up making calls in a network that does not support such a scheme or the terminal may have an older list of priorities in a namespace.

Multiple simultaneous schemes: Some user agents will need to support multiple different priority schemes, as several will remain in use in networks run by different agencies and operators. (Not all user agents will have the means of authorizing callers using different schemes, and thus may be configured at run-time to only recognize certain namespaces.)

Discovery: A terminal should be able to discover which priority name spaces are supported by a particular proxy or user agent (gateway).

Priority mechanisms can be roughly categorized by whether they use a priority queue for resource attempts or whether they preempt existing resource uses (e.g., calls). The choice between these mechanisms depends on the operational needs and characteristics of the network, e.g., on the number of active requests in the system and the fraction of prioritized calls. Generally, if the number of prioritized calls is small compared to the system capacity and the system capacity is large, it is likely that another call will naturally terminate in short order when a higher-priority call arrives. Thus, it is conceivable that the system indication can sometimes cause preemption, while otherwise it just

Some namespaces may inherently imply a preemption policy, while others may be silent on whether preemption is to be used or not, leaving this to local entity policy.

Similarly, the precise relationships between labels, e.g., what fraction of capacity is set aside for each priority level, is also a matter of local policy. This is similar to how differentiated services labels are handled.

[2](#) Relationship to Emergency Call Services

The resource priority mechanisms are used to have selected individuals place calls with elevated priority during times when the network is suffering from an emergency-induced shortage of resources. Generally, calls for emergency help placed by non-officials (e.g., "911" and "112" calls) don't need resource priority under normal circumstances. If such emergency calls are placed during emergency-induced network resource shortages, the call identifier itself is sufficient to identify the emergency nature of the call. Adding an indication of resource priority may be less appropriate, as this would require that all such calls carry this indicator. Also, it opens another attack mechanism, where non-emergency calls are marked as emergency calls. (If the entity can recognize the request URI as an emergency call, it would not need the resource priority mechanism.)

3 Security Requirements

Any resource priority mechanism can be abused to obtain resources and thus deny service to other users. While the indication itself does not have to provide separate authentication, any SIP request carrying such information has higher authentication requirements than regular requests.

4 Acknowledgements

James Polk provided helpful comments.

5 Authors' Addresses

Henning Schulzrinne
Dept. of Computer Science
Columbia University
1214 Amsterdam Avenue
New York, NY 10027
USA
electronic mail: schulzrinne@cs.columbia.edu

