Internet Engineering Task Force
Internet Draft
draft-pierce-tsvwg-pref-treat-examples-01.txt
October 20, 2004
Expires April 20, 2005
Mike Pierce
Artel
Don Choi
DISA

Examples for Provision of Preferential Treatment in Voice over IP
draft-pierce-tsvwg-pref-treat-examples-01.txt

Status of this memo

Copyright

Abstract

Assured Service refers to the set of capabilities used to ensure that mission critical communications are setup and remain connected. [Pierce] describes the requirements, one of which is to provide preferential treatment to higher priority calls. IEPS refers to a set of capabilities used to provide a higher probability of call

completion to emergency calls made by authorized personnel, usually from ordinary telephones. This also requires some form of preferential treatment. This informational memo describes some of the methods which may be applied to provide that preferential treatment.

Table of Contents

0. History

(To be removed before publication.)

This draft was originally submitted under SIPPING, then submitted under IEPREP to focus consideration and discussion in that WG in conjunction with the related discussions for IEPS. It is now submitted to TSVWG.

(SIPPING) -00 Initial version based on material removed from draft-pierce-sipping-assured-service-01.

(IEPREP) -00 Added references to IEPREP in Intro. Update references. add details about packet dropping procedure.

(IEPREP) -01 Updated references

(IEPREP) -02 Added Annexes from requirements draft.

(TSVWG) -00 Resubmitted under TSVWG. Clarified that each method by itself is not believed to be sufficient. Multiple procedures need to be used together. Expanded description of RSVP. Clarified reference to CAC.

(TSVWG -01
- Added additional description in 3.2 of how Call Admission Control fits into this framework.
- Added reference to June 2004 IEEE article by Xu.

1. Introduction

The requirements for Assured Service in support of networks requiring precedence treatment for certain calls is are described in [Pierce]. One of those requirements is Preferential treatment, which is the following:

It must be possible to provide preferential treatment to higher precedence calls in relation to lower precedence calls. Examples of preferential treatments are:

- reservation of network resources for precedence calls

- usage of higher Call Admission Control (CAC) limits for acceptance of new higher precedence calls

- preferential queuing of signaling messages based on precedence level

- preferential queuing of user data packets based on precedence level

- discarding of packets of lower precedence call

- preemption of one or more existing calls of lower precedence level

- preemption of some of the resources being used by a call of lower precedence level

- preemption of the reservation of resources being held for other traffic

Several documents describe the requirements for provision of the International Emergency Preparedness Scheme (IEPS). This service requires some types of preferential treatment for these calls, which can be viewed as a subset of the requirements for Assured Service listed above. These requirements include:

- higher probability of call completion

- lower probability of premature disconnect

- distinguish IEPS data packets from other types of VoIP Packets in order to give them "priority".

- alternate path routing

This informational memo describes some ways in which the above listed preferential treatments may be provided by utilizing current or new capabilities.

2. Background

The requirement for Precedence Level marking of a call setup attempt using SIP [RFC3261] will be met by the Resource Priority header [Resource]. The value carried in this header represents the relative precedence level of the call, and is used to control any of the following described procedures for providing Preferential Treatment.

3. Potential Preferential Treatments

The requirement to provide preferential treatment to calls may be met by applying the appropriate combination of the following procedures. Due to the complexity of the network and the protocols being used, it is not expected that any one of these procedures will be sufficient by itself.

In addition, there may be other procedures and treatments not described herein.

3.1. Reservation of Network Resources

This procedure involves pre-reserving certain network resources during periods when no higher precedence traffic is present so as to be prepared to handle a given level of high precedence traffic in the case of an emergency. While this method is already used in the circuit switched environment, it is less than desirable since it requires a tradeoff between the amount of wasted resources during non-emergency periods and the amount of emergency traffic which can be handled using those reserved facilities.

IETF defined QoS mechanisms for packet-mode operation offer some improvement to this situation by allowing the amount of reserved resources to be adjusted.

3.1.1. RSVP

3.1.1.1. Reservation of Trunk Groups

RSVP may be used to establish multiple trunk groups between switching points, with each trunk group serving a different precedence level of calls. Each trunk group would be sized based on the number of simultaneous calls of that precedence level to be supported. (In this context, a trunk group refers to a facility which can support a certain number of voice connections at a certain Quality of Service level. As noted later, the number of connections can be increased with a corresponding decease in the QoS level.)

With TE, the reserved sizes of these trunk groups could be adjusted during times of emergency.

No preemption of these trunk groups is needed. However, reducing the size of a group to near zero would prevent further calls from using it while allowing existing calls to continue.

3.1.1.2. Reservation for Individual Calls

RSVP may be used to establish paths for individual calls (packet flows) with aggregation taking place as described in RFC 3175. This also provides the ability to preempt such as flow.

3.1.2. MPLS

MPLS may be used to establish the equivalent of dedicated trunk groups between switching entities, enterprise network, etc. Each of these "trunk groups" could exist to support a specific precedence level of traffic between two points and could be setup using the procedures of CR-LDP [RFC3212] or RSVP-TE [RFC3209]. These support the signaling of the required five levels of precedence.

3.1.2.1. Constraint-based LSP Setup using LDP

CR-LDP [RFC3212] defines an extension to LDP to provide a constraint-based routing using MPLS. One of the constraints is based on the notion of a "priority" level for the new setup. It includes the signaling of a setup priority and a holding priority with the value of each being 0-7 (0 is the highest priority). When setting up an LSP as a trunk group to carry the traffic of one of the expected precedence levels defined in [Pierce], the following mapping would be used:

Assured Service Precedence Level	RFC3212 Preemption TLV	
	SetPrio	HoldPrio
Routine	4	0
Priority	3	0
Immediate	2	0
Flash	1	0
Flash Override	0	0

This mapping prevents any preemption of a trunk group for the establishment of another. Rather, it is expected that trunk groups for all precedence levels would be initially created and remain. Only their allocated size might be changed.

If actual preemption were desired, the appropriate HoldPrio values would be used.

3.1.2.2. RSVP-TE: Extensions to RSVP for LSP Tunnels

As an alternative to LDP, RSVP-TE [RFC3209] defines the use of RSVP with extensions to perform the label distribution for MPLS. It also includes the same setup and holding priorities as defined in CR-LDP [RFC3212]. When using RSVP as the label distribution protocol, the same mapping shown above for LDP would be used.

3.2. Use of Higher Call Admission Control (CAC) Limits

It is presumed that any network which might reach a congestion point (evidenced by queue overflows, packet loss, etc.) must have a means to limit the establishment of new packet flows. This is true for any system, not just those providing Assured Service. For flows used for voice calls, this function is referred to herein as "Call Admission Control (CAC)". This document does not address the methods which might be used to provide CAC. However, due to the complexity of any network and the suddenly varying traffic rates which Assured Service is specifically intended to deal with, it is further assumed that no CAC can possibly prevent all cases of congestion. At best, it is a good approximation and other techniques are still required to deal with a congestion which may still occur. It is further assumed that CAC is always based on some limits which are placed of the establishment of new packet flows for new calls, whether in terms of number of calls, or bandwidth used.

One aspect of preferential treatment may be provided by allowing higher precedence calls to be setup even when they result in exceeding the engineered traffic limit on a facility (on an MPLS LSR, for example). This operation is based on an assumption of normal traffic behavior in which calls are continuously releasing. It also presumes that the actual packet flow for the new call will not be started until some time after call setup, for example, at answer. Any exceeding of the engineered limit is expected to be short-term.

Note: "Engineered traffic limit" here is intended to mean values, either calculated or obtained through experience, of the limits on loading which can occur and still meet the desired performance, for example, packet loss rate < 0.1%. In some cases, "congestion" means going over this limit.

This procedure presumes the existence of a Call Admission Control function which is aware of the traffic loading on various links and entities, and compares these against some thresholds before allowing the establishment of a new call (packet flow).

For example, the limits for Call Admission Control for new calls could be set as depicted in the following table, where the engineered capacity of a route or facility is "x". A new call of each precedence level would be allowed only if the current load is within the limit shown:

Precedence Level	Capacity limit of
Routine	.9x
Priority	.95x
Immediate	x
Flash	1.05x
Flash-override	1.1x

Explanation of table: In this example, a new Flash call is allowed to be setup if the current traffic load for all traffic on the facility is less than 1.1x. In the example shown in this table, Routine traffic is always prevented from using the last 10% of the engineered capacity. The choice of the multipliers would be based on an analysis of the tradeoff between getting the high precedence level call through vs. sacrificing its QoS. It would depend on the voice encoding algorithms typically used and the end user expectations.

Note: As an example, the values in the above table may have been derived from a calculation that, for the codec being used, oversubscribing by 10% will lead to a certain packet loss rate which, although serious, is preferable to blocking the setup of the new Flash override call.

This procedure is based on a requirement that Flash override calls should "never" be blocked. (In a probability-based system, there is no such thing as "never".) In the circuit-switched environment this could only be guaranteed by having as many circuits as there might be Flash override calls. For IP-based service, there is no fixed number of "circuits" on any facility. The "x" referred to above is only an engineering limit based on a guarantee for the provision of a certain QoS for normal traffic, i.e., Routine and Priority. This "x" may be thought of as the number of "circuits" for normal traffic. It is preferable to allow the setup of additional higher precedence calls with reduced QoS rather than blocking their setup. For example, while a particular facility may support 100 normal calls (Routine and Priority) at the guaranteed QoS, it might support 110 calls at a reduced, yet acceptable, QoS (due to packet loss) when in an emergency situation. This could allow 10 higher precedence calls when they would otherwise be blocked.

Since the packet preferential treatment using Diff-Serv described in Section 3.5 could result in the discard or loss of the packets for the lower precedence calls, the higher precedence calls could still be provided a sufficient QoS even though they may have caused the engineered capacity of the route to be exceeded. The lower precedence calls will then experience higher packet discard rates or queuing delay times. If the discard rate or delay for these lower precedence calls is excessive, the end user will experience poor QoS and will likely disconnect, thereby freeing up the resources.

3.3. Preferential Queuing of Signaling Messages

There is no plan to apply preferential queuing to signaling messages of higher precedence calls (ahead of other signaling messages), just as this was not done in the circuit switched network. No advantage can be shown for such a procedure and it would only aggravate the problem of out-of-order messages.

3.4. Preferential Queuing of User Data Packets

It is not expected that priority queuing of user data packets (ahead of other user data packets of the same type) would provide a useful capability.

3.5. Discarding of Packets using DiffServ

Within DiffServ, Assured Forwarding [RFC2597] provides four classes and three drop precedences for each class (12 DSCP code points). One of these classes could be used for the signaling messages for session establishment and release. AF is not considered as being appropriate for audio.

Expedited Forwarding [RFC3246] defines a single class (DSCP code point) and operation, but does not include multiple drop precedences as AF does. The intention of EF is to "provide low loss, latency and jitter" and is understood to be intended for traffic such as speech, although RFC 3246 does not explicitly mention speech or voice. However, speech is less susceptible to loss than the signaling traffic and, under some traffic situations, will constitute a much larger portion of the overall load. Therefore, multiple drop precedences to alleviate overload may be more appropriate to EF than they are to AF.

The result of this use of DiffServ classes is that voice packets are always given priority over the signaling packets and all voice packets are treated the same. While this is the desired behavior in many cases, it is not desired in those cases in which a limited sized facility could become completely occupied by voice traffic (using EF). In this situation, further signaling messages (using AF), including those to setup new high precedence calls and those to release low precedence calls, would be lost or excessively delayed.

Therefore, it is necessary to reserve a small capacity for use by the AF class which serves the signaling traffic as described in Section 2.10 of EF [RFC3246].

For that portion of the capacity using EF for voice, part of the required preferential treatment for the five call precedence levels may be provided by the use of multiple drop precedence (probability) levels for packets. The procedures for these drop precedence levels would be similar to that defined currently for the three levels for each class in AF [RFC2597].

Five such levels for packet marking, using DSCPs, are needed to provide the required functionality. In the absence of "standardized" DSCP values, local values could be assigned. Based on the definitions for AF, these levels are referred to here as:

- Very low (i.e., lowest probability of being dropped)
- Low
- Medium
- High
- Very high (i.e., highest probability of being dropped)

The following possible mappings are shown to illustrate the concept of using DiffServ codepoints to assist in the provision of preferential treatment to the individual packets which make up the information transfer (both the connection setup signaling and the voice transfer) of an Assured Service call.

3.5.1. Treatment for Signaling Packets

Consideration could be given to utilization of different drop precedences for the signaling messages associated with different precedence sessions. However, using SS#7 in the PSTN as a basis, it might also be meaningful to provide different drop precedences based on the type of message rather than only based on the precedence of the call. For example, for routine traffic, those messages which cause the release of sessions could be given a lower drop precedence than those which set up new sessions in order to allow such releases to take place properly under overload conditions. High precedence calls, on the other hand could use a lower drop precedence level for session setup messages than those of routine precedence calls. The following table shows the Congestion Priority Level assignments defined for SS#7 [T1.111], including High Probability of Completion [T1.631] and MLPP [T1.619], and a suggestion of what might be used for SIP for the corresponding messages.

(Note: The highest SS#7 Congestion Priority Level, i.e., "3", is the last to be dropped during congestion.)

(Refer to RFC 3398 for mapping of ISUP to SIP messages.)

SS#7		SIP	
Message	Congestion Priority Level	Message	Drop Precedence Level
Network management	3	?	low
ANM	2	200 OK (INVITE)	medium
RLC	2	200 OK (BYE)	(note)
IAM (MLPP)	1 or 2	INVITE (AS)	low/medium
IAM (HPC)	1	INVITE (IEPS)	low
ACM	1	18x	medium
CPG	1	100 Trying/18x	medium
REL	1	BYE	low
IAM (normal)	0	INVITE (normal)	high
Others	0		

Note: For SIP, unless noted otherwise, all ACKs should have the same preferential treatment as the message they are acknowledging.

3.5.2. Treatment for Voice Packets

This example is for the case of the use of DiffServ to provide the packet forwarding preferential treatment through multiple drop precedence levels. It uses the Multi-Level Expedited Forwarding Per Hop Behavior [Silverman] which is also described in [Xu]. Each packet containing user data (voice) is marked with a unique DiffServ codepoint to indicate one of the following levels and resulting treatment:

Precedence Level	Indication in user voice packets		Drop if current queue is more than -- % full (note 1)
	Class	Drop precedence	
Routine	MLEF	Very high	80%
Priority	MLEF	High	90%
Immediate	MLEF	Medium	100%
Flash	MLEF	Low	110%
Flash Override	MLEF	Very low	120%

All voice traffic is then served by a single instance of MLEF, and served by a single (strict FIFO) queue. This results is an equal treatment in terms of delay variation (often called "jitter") for all precedence levels for those packets which are delivered, but achieves this by selective packet discard. The discard may use a simple tail dropping algorithm as shown in the above table or a form of "Random Early Detection" as described in [RFC2309] and [Xu] to drop some packets before the queue actually reaches the fill shown above. However, since the packets in this queue are not using TCP and can not be bursty or "aggressive" or of large size, there appears to be no advantage gained by the complexity of early detection and random dropping algorithms.

Note 1: The "queue full" here refers to the engineered limit, that is, the limit which needs to be applied in order to meet the requirements of the EF PHB and the desired QoS in terms of maximum delay introduced by this queue. Since this calculation of maximum queue length is based on probabilities of achieving a certain target QoS, it can be temporarily exceeded as described in Section 3.6.2. This is shown in the above table by using values greater than 100% for Flash and Flash override. It is essentially this "over-subscription" of higher precedence packets which causes packets of the lower precedence calls to be discarded. This presumes that the condition of packet drop will be temporary as calls normally release and new calls are prevented from being established.

It should be emphasized that selective packet discard based on DSCP (which is based on the call precedence level) can not by itself provide a useful service. Without effective CAC, excess offered traffic will lead to congestive collapse, and selective packet discard can not prevent this collapse.

3.6. Preemption

3.6.1. Call Preemption

If possible, actual preemption of existing calls may be provided in order to achieve the same functionality as previously available in the circuit-switched environment with MLPP, that is, use of the proper notifications sent to the users whose call is being preempted. Such preemption would have to be controlled by an entity which has knowledge of: 1) the network architecture, 2) the current loads on links, 3) which links require freed-up capacity for a higher precedence call, and 4) which packet flows need to be terminated to free-up that capacity. It would also require appropriate signaling from that entity to cause the preemption.

When interworking with circuit switched portions of the telecommunications network, preemption procedures are still required within transport facilities which are based on fixed numbers of circuits. In some cases, this preemption results in specific procedures being applied in the packet portion, such as notifications of preemption and forced disconnect of a call.

3.6.2. Preemption of Some of the Resources Being Used

The procedures described above for use of higher call acceptance limits (3.2) and selective discard of voice packets based on the precedence level of the call (3.5.2) may reduce or eliminate the need to perform preemption of existing calls within the IP domain. The statistical nature of packet transmission makes it possible to "squeeze" an additional high precedence call into an already "full" facility, as illustrated in the previous section. It should be noted that, in the extreme case, these procedures would result in a similar effect as preemption, but without the required user notification, since the resources of the lower precedence calls would be so severely degraded (via packet loss) that communication would be impossible and the users would eventually disconnect.

Because each packet flow arrives at somewhat regular intervals, it is expected that, when packet loss is occurring due to discard, the loss will not be random across all flows using the DSCP with the highest discard probability. Rather, losses will likely be bursty on each flow, with most discards being on one flow for many consecutive packets.

3.7. Preemption of the Reservation

Based on traffic engineering, the amount of resources allocated to reserved paths (e.g., MPLS or RSVP) could be adjusted. For example, when an emergency situation occurs, the need for more resources to support higher priority traffic could be recognized. The existing LSPs could be changed using the procedures of [RFC3214] to allow the size of those LSPs supporting the higher priority traffic to be increased while others are decreased.

3.8. Exemption from Network Management Controls

Network Management controls may sometimes restrict call setup, for example, during times of natural disasters a network may intentionally block calls going into that area in order to reserve facilities for calls coming from that area. One preferential treatment which may be applied to higher precedence calls is to allow them to override such Network Management controls.

4. Security Considerations

The security considerations are covered in [Pierce].

5. IANA Considerations

This document does not, by itself, specify any IANA involvement in support of provision of Preferential Treatment for Assured Service. The only referenced IANA involvement is described in [Resource].

6. References

6.1. Normative References

None

6.2. Informative References

[RFC2205] "Resource ReSerVation Protocol (RSVP)", R. Braden, et al, September 1997

[RFC2309] "Recommendations on Queue Management and Congestion Avoidance", B. Braden, April 1998.

[RFC2597] "Assured Forwarding PHB Group", J. Heinanen, et al, June 1999.

[RFC3209] "RSVP-TE: Extensions to RSVP for LSP Tunnels", D. Awduche, December 2001.

[RFC3212] "CR-LDP: Constraint-based LSP Setup using LDP", B. Jamoussi, et al, January 2002.

[RFC3214] "LSP Modification Using CR-LDP", J. Ash, et al, January 2002.

[RFC3246] "An Expedited Forwarding PHB", B. Davie, et al, March 2002.

[RFC3261] "SIP: Session Initiation Protocol", J. Rosenberg, et al, June 2002.

[T1.111] ANSI T1.111-2001, "Signalling System No. 7 (SS7) - Message Transfer Part".

[T1.619] ANSI T1.619-1992 (R1999) "ISDN - Multi-Level Precedence and Preemption (MLPP) Service Capability".

[T1.631] ANSI T1.631-1993 (R1999) "Telecommunications - Signalling System No. 7 (SS7) - High Probability of Completion (HPC) Network Capability".

[Pierce] draft-pierce-tsvwg-assured-service-req-01, "Requirements for Assured Service Capabilities in Voice over IP", October 2004

[Resource] draft-ietf-sip-resource-priority-04, "SIP Communications Resource Priority Header", Henning Schulzrinne and James Polk, August 2004.

[Silverman] draft-silverman-tsvwg-mlefphb-01, "Multi-Level Expedited Forwarding Per Hop Behavior (MLEF PHB", Steve Silverman, et al, October 2004.

[Xu] "An Investigation of Multilevel Service Provision for Voice over IP Under Catastrophic Congestion", Yang Xu, Martin Westhead, Fred Baker, June 2004 IEEE Communications Magazine.

Authors' Addresses

Michael Pierce
Artel
1893 Preston White Drive
Reston, VA 20191
Phone: +1 410.817.4795
Email: pierce1m@ncr.disa.mil

Don Choi
DISA
5600 Columbia Pike
Falls Church, VA 22041-2717
Phone: +1 703.681.2312
Email: choid@ncr.disa.mil

Disclaimer of Validity

Intellectual Property

Full Copyright Statement