

Internet Draft
Expiration: August 1999
File: [draft-ietf-rap-signaled-priority-02.txt](#)

Shai Herzog
IPHighway

Signaled Preemption Priority Policy Element

February 13, 1999

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>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

Abstract

This document describes a preemption priority policy element for use by signaled policy based admission protocols (such as [[RSVP](#)] and [[COPS](#)]).

Preemption priority defines a relative importance (rank) within the set of flows competing to be admitted into the network. Rather than admitting flows by order of arrival (First Come First Admitted) network nodes may consider priorities to preempt some previously admitted low priority flows in order to make room for a newer, high-priority flow.

Table of Contents

| | |
|---|--------------------|
| Abstract..... | 1 |
| Table of Contents..... | 2 |
| 1 Introduction..... | 3 |
| 2 Scope and Applicability..... | 3 |
| 3 Stateless Policy..... | 4 |
| 4 Policy Element Format..... | 4 |
| 5 Priority Merging Issues..... | 6 |
| 5.1 Priority Merging Strategies..... | 7 |
| 5.1.1 Take priority of highest QoS..... | 7 |
| 5.1.2 Take highest priority..... | 7 |
| 5.1.3 Force error on heterogeneous merge..... | 8 |
| 5.2 Modifying Priority Elements..... | 8 |
| 6 Error Processing..... | 9 |
| 7 Security Considerations..... | 9 |
| 8 References..... | 10 |
| 9 Author Information..... | 10 |
| Appendix A: Example..... | 11 |
| A.1 Computing Merged Priority..... | 11 |
| A.2 Translation (Compression) of Priority Elements..... | 11 |

1 Introduction

Traditional Capacity based Admission Control (CAC) indiscriminately admits new flows until capacity is exhausted (First Come First Admitted). Policy based Admission Control (PAC) on the other hand attempts to minimize the significance of order of arrival and use policy based admission criteria instead.

One of the more popular policy criteria is the rank of importance of a flow relative to the others competing for admission into a network node. Preemption Priority takes effect only when a set of flows attempting admission through a node represents overbooking of resources such that based on CAC some would have to be rejected. Preemption priority criteria help the node select the most important flows (highest priority) for admission, while rejecting the low priority ones.

Network nodes which support preemption should consider priorities to preempt some previously admitted low-priority flows in order to make room for a newer, high-priority flow.

This document describes the format and applicability of the preemption priority represented as a policy element in [[RSVP-EXT](#)].

2 Scope and Applicability

The Framework document for policy-based admission control [[RAP](#)] describes the various components that participate in policy decision making (i.e., PDP, PEP and LDP). The emphasis of PREEMPTION_PRI elements is to be simple, stateless, and light-weight such that they could be implemented internally within a node's LDP (Local Decision Point).

Certain base assumptions are made in the usage model for PREEMPTION_PRI elements:

- They are created by PDPs

In a model where PDPs control PEPs at the periphery of the policy domain (e.g., in border routers), PDPs reduce sets of relevant policy rules into a single priority criterion. This priority as expressed in the PREEMPTION_PRI element can then be communicated to downstream PEPs of the same policy domain, which have LDPs but no controlling PDP.

- They can be processed by LDPs

PREEMPTION_PRI elements are processed by LDPs of nodes that do not have a controlling PDP. LDPs may interpret these objects, forward them as is, or perform local merging to forward an equivalent merged PREEMPTION_PRI policy element. LDPs must follow the merging strategy that was encoded by PDPs in the PREEMPTION_PRI objects. (Clearly, a PDP, being a superset of LDP, may act as an LDP as well).

- They are enforced by PEPs

PREEMPTION_PRI elements interact with a node's traffic control module (and capacity admission control) to enforce priorities, and preempt previously admitted flows when the need arises.

3 Stateless Policy

Signaled Preemption Priority is stateless (does not require past history or external information to be interpreted). Therefore, when carried in COPS messages for the outsourcing of policy decisions, these objects are included as COPS Stateless Policy Data Decision objects (see [COSP, COPS-RSVP]).

4 Policy Element Format

The format of Policy Data objects is defined in [RSVP-EXT]. A single Policy Data object may contain one or more policy elements, each representing a different (and perhaps orthogonal) policy.

The format of preemption priority policy element is as follows:

```

+-----+-----+-----+-----+
| Length (12)          | P-Type = PREEMPTION_PRI |
+-----+-----+-----+-----+
| Flags      | M. Strategy | Error Code | Reserved(0) |
+-----+-----+-----+-----+
| Preemption Priority    | Defending Priority      |
+-----+-----+-----+-----+
```


Length: 16 bits

Always 12. The overall length of the policy element, in bytes.

P-Type: 16 bits

PREEMPTION_PRI = 1

The preemption priority policy element number was assigned by IANA as defined in [[RSVP-EXT](#)].

Flags: 8 bits

Reserved (always 0).

Merge Strategy: 8 bit

- 1 Take priority of highest QoS: recommended
- 2 Take highest priority: aggressive
- 3 Force Error on heterogeneous merge

Reserved: 8 bits

Error code: 8 bits

- | | | |
|---|---------------|--|
| 0 | NO_ERROR | Value used for regular PREEMPTION_PRI elements |
| 1 | PREEMPTION | This previously admitted flow was preempted |
| 2 | HETEROGENEOUS | This element encountered heterogeneous merge |

Reserved: 8 bits

Always 0.

Preemption Priority: 16 bit (unsigned)

The priority of the new flow compared with the defending priority of previously admitted flows. Higher values represent higher Priority.

Defending Priority: 16 bits (unsigned)

Once a flow was admitted, the preemption priority becomes irrelevant. Instead, its defending priority is used to compare with the preemption priority of new flows.

For any specific flow, its preemption priority must always be less than or equal to the defending priority. A wide gap between preemption and defending priority provides added stability: moderate preemption priority makes it harder for a flow to preempt others, but once it succeeded, the higher defending priority makes

The combination of preemption priority, killer reservation and blockade state [RSVP] may increase the instability of admitted flows where a reservation may be preempted, reinstated, and preempted again periodically.

.1 Priority Merging Strategies

In merging situations LDPs may receive multiple preemption elements and must compute the priority of the merged flow according to the following rules:

- a. Preemption priority and defending priority are merged and computed separately, irrespective of each other.
- b. Participating priority elements are selected.

All priority elements are examined according to their merging strategy to decide whether they should participate in the merged result (as specified bellow).

- c. The highest priority of all participating priority elements is computed.

The remainder of this section describes the different merging strategies the can be specified in the PREEMPTION_PRI element.

1.1 Take priority of highest QoS

The PREEMPTION_PRI element would participate in the merged reservation only if it belongs to a flow that contributed to the merged QoS level (i.e., that its QoS requirement does not constitute a subset another reservation.)

A simple way to determine whether a flow contributed to the merged QoS result is to compute the merged QoS with and without it and to compare the results (although this is clearly not the most efficient method).

The reasoning for this approach is that the highest QoS flow is the one dominating the merged reservation and as such its priority should dominate it as well. This approach is the most amiable to the prevention of priority distortions such as free-riders and denial of service.

This is a recommended merging strategy.

1.2 Take highest priority

All PREEMPTION_PRI elements participate in the merged reservation.

This strategy disassociates priority and QoS level, and therefore is highly subject to free-riders and its inverse image, denial of service.

This is not a recommended method, but may be simpler to implement.

This approach effectively compresses the number of forwarded PREEMPTION_PRI elements to at most to the number of different merging strategies, regardless of the number of receivers (See the example in [Appendix A.2](#)).

6 Error Processing

A PREEMPTION_PRI error object is sent back toward the appropriate receivers when an error involving PREEMPTION_PRI elements occur.

PREEMPTION

When a previously admitted flow is preempted, a copy of the preempting flow's PREEMPTION_PRI element is sent back toward the PDP that originated the preempted PREEMPTION_PRI object. This PDP, having information on both the preempting and the preempted priorities may construct a higher priority PREEMPTION_PRI element in an effort to re-instate the preempted flow.

Heterogeneity

When a flow F1 with Heterogeneous Error merging strategy set in its PREEMPTION_PRI element encounters heterogeneity the PREEMPTION_PRI element is sent back toward receivers with the Heterogeneity error code set.

7 Security Considerations

The integrity of PREEMPTION_PRI is guaranteed, as any other policy element, by the encapsulation into a Policy Data object [[RSVP-EXT](#)].

Further security mechanisms are not warranted, especially considering that preemption priority aims to provide simple and quick guidance to routers within a trusted zone or at least a single zone (no zone boundaries are crossed).

8 References

- [RSVP-EXT] Herzog, S. "RSVP Extensions for Policy Control", Internet-Draft, [draft-ietf-rap-rsvp-ext-02.txt](#), Jan. 1999.
- [COPS-RSVP] Boyle, J., Cohen, R., Durham, D., Herzog, S., Raja, n R., Sastry, A., COPS usage for RSVP Internet-Draft, [draft-ietf-rap-cops-rsvp-02.txt](#), Jan 1999.
- [RAP] Yavatkar, R., et al., "A Framework for Policy Based Admission Control", IETF <[draft-ietf-rap-framework-02.txt](#)>, Jan., 1999.
- [COPS] Boyle, J., Cohen, R., Durham, D., Herzog, S., Raja, n R., Sastry, A., "The COPS (Common Open Policy Service) Protocol", IETF <[draft-ietf-rap-cops-05.txt](#)>, Jan. 1999.
- [RSVP] Braden, R. ed., "Resource ReSerVation Protocol (RSVP) - Functional Specification.", IETF [RFC 2205](#), Proposed Standard, Sep. 1997.

9 Author Information

Shai Herzog, IPHighway
Parker Plaza, Suite 1500
400 Kelby St.
Fort-Lee, NJ 07024
(201) 585-0800
herzog@iphighway.com

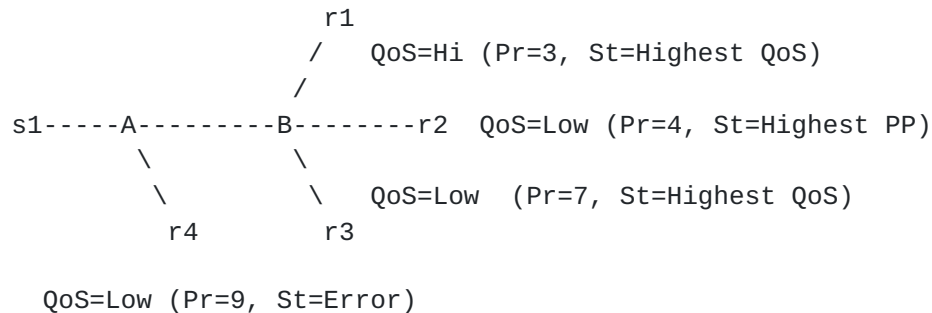
Shai Herzog

Expires August 1999

[Page 10]

Appendix A: Example

The following examples describe the computation of merged priority elements as well as the translation (compression) of PREEMPTION_PRI elements.

A.1 Computing Merged Priority

Example 1: Merging preemption priority elements

Example one describes a multicast scenario with one sender and four receivers each with each own PREEMPTION_PRI element definition.

r1, r2 and r3 merge in B. The resulting priority is 4.

Reason: The PREEMPTION_PRI of r3 doesn't participate (since r3 is not contributing to the merged QoS) and the priority is the highest of the PREEMPTION_PRI from r1 and r2.

r1, r2, r3 and r4 merge in A. The resulting priority is again 4: r4 doesn't participate because its own QoS=Low is incompatible with the other (r1) QoS=High. An error PREEMPTION_PRI should be sent back to r4 telling it that its PREEMPTION_PRI element encountered heterogeneity.

A.2 Translation (Compression) of Priority Elements

Given this set of participating PREEMPTION_PRI elements, the following compression can take place at the merging node:

From:

(Pr=3, St=Highest QoS)
 (Pr=7, St=Highest QoS)
 (Pr=4, St=Highest PP)
 (Pr=9, St=Highest PP)
 (Pr=6, St=Highest PP)

To:

(Pr=7, St=Highest QoS)

(Pr=9, St=Highest PP)

Shai Herzog

Expires August 1999

[Page 11]