

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 6, 2007

A. Charny
Cisco Systems, Inc.
J. Zhang
Cisco Systems, Inc. and Cornell
University
F. Le Faucheur
V. Liatsos
Cisco Systems, Inc.
March 5, 2007

Pre-Congestion Notification Using Single Marking for Admission and Pre-emption

[draft-charny-pcn-single-marking-01.txt](#)

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

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>.

This Internet-Draft will expire on September 6, 2007.

Copyright Notice

Copyright (C) The IETF Trust (2007).

Abstract

Pre-Congestion Notification [[I-D.briscoe-tsvwg-cl-architecture](#)] approach proposes the use of an Admission Control mechanism to limit

the amount of real-time PCN traffic to a configured level during the normal operating conditions, and the use of a Pre-emption mechanism to tear-down some of the flows to bring the PCN traffic level down to a desirable amount during unexpected events such as network failures, with the goal of maintaining the QoS assurances to the remaining flows. In [[I-D.briscoe-tsvwg-cl-architecture](#)], Admission and Pre-emption use two different markings and two different metering mechanisms in the internal nodes of the PCN region. This draft proposes a mechanism using a single marking and metering for both Admission and Pre-emption, and presents a preliminary analysis of the tradeoffs. A side-effect of this proposal that a different marking and metering Admission mechanism than that proposed in [[I-D.briscoe-tsvwg-cl-architecture](#)] may be also feasible.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

Table of Contents

1.	Introduction	4
1.1.	Changes from -00 version	4
1.2.	Background and Motivation	4
1.3.	Terminology	6
2.	The Single Marking Approach	6
2.1.	High Level description	6
2.2.	Operation at the PIN	7
2.3.	Operation at the Egress PEN	7
2.4.	Operation at the Ingress PEN	8
2.4.1.	Admission Decision	8
2.4.2.	Pre-emption Decision	9
3.	Discussion	9
3.1.	Benefits	9
3.2.	Tradeoffs, Issues and Discussion	10
3.2.1.	Restrictions on Pre-emption-to-admission Thresholds	10
3.2.2.	Assumptions on Loss	10
3.2.3.	Effect of Reaction Timescale of Admission Mechanism	10
3.2.4.	Performance Implications and Tradeoffs	11
3.2.5.	Effect on Proposed Anti-cheating Mechanisms	11
3.2.6.	Standards Implications	12
4.	Performance Evaluation Comparison	12
4.1.	Relationship to other drafts	12
4.2.	Limitations, Conclusions and Direction for Future Work	12
4.2.1.	High Level Conclusions	12
4.2.2.	Future work	13
5.	Appendix A: Simulation Details	14
5.1.	Network and Signaling Model	14
5.2.	Traffic Models	16
5.2.1.	CBR Voice (CBR)	17
5.2.2.	VBR Voice (VBR)	17
5.2.3.	"Synthetic Video": High Rate ON-OFF traffic with Video-like Mean and Peak Rates ("SVD")	17
5.2.4.	Real Video Traces (VTR)	17
5.3.	Parameter Settings	18
5.3.1.	Queue-based settings	18
5.3.2.	Token Bucket Settings	18
5.4.	Simulation Details	19
5.4.1.	Queue-based Results	19
5.4.2.	Token Bucket-based Results	23
6.	IANA Considerations	27
7.	Security Considerations	27
8.	References	27
8.1.	Normative References	27
8.2.	Informative References	27
	Authors' Addresses	28
	Intellectual Property and Copyright Statements	30

1. Introduction

1.1. Changes from -00 version

Added miscellaneous clarifications based on comments received on version -00. Added simulation results for multi-hop topologies and video trace data to the Appendix.

1.2. Background and Motivation

Pre-Congestion Notification [[I-D.briscoe-tsvwg-cl-architecture](#)] approach proposes to use an Admission Control mechanism to limit the amount of real-time PCN traffic to a configured level during the normal operating conditions, and to use a Pre-emption mechanism to tear-down some of the flows to bring the PCN traffic level down to a desirable amount during unexpected events such as network failures, with the goal of maintaining the QoS assurances to the remaining flows. In [[I-D.briscoe-tsvwg-cl-architecture](#)], Admission and Pre-emption use different two different markings and two different metering mechanisms in the internal nodes of the PCN region. Admission Control algorithms for variable-rate real-time traffic such as video have traditionally been based on the observation of the queue length, and hence re-using these techniques and ideas in the context of pre-congestion notification is highly attractive, and motivated the virtual-queue-based marking and metering approach specified in [[I-D.briscoe-tsvwg-cl-architecture](#)] for Admission. On the other hand, for Pre-emption, it is desirable to know how much traffic needs to be pre-empted, and that in turn motivates rate-based Pre-emption metering. This provides some motivation for employing different metering algorithm for Admission and for Preemption.

Furthermore, it is frequently desirable to trigger Pre-emption at a substantially higher traffic level than the level at which no new flows are to be admitted. There are multiple reasons for the requirement to enforce a different Admission Threshold and Preemption Threshold. These include, for example:

- o End-users are typically more annoyed by their established call dying than by getting a busy tone at call establishment.
- o There are often very tight (possibly legal) obligations on network operators to not drop established calls.
- o Voice Call Routing often has the ability to route/establish the call on another network (e.g., PSTN) if it is determined at call establishment that one network (e.g., packet network) can not accept the call. Therefore, not admitting a call on the packet network at initial establishment may not impact the end-user. In

contrast, it is usually not possible to reroute an established call onto another network mid-call. This means that call Preemption can not be hidden to the end-user.

- o Preemption is typically useful in failure situations where some loads get rerouted thereby increasing the load on remaining links. Because the failure may only be temporary, the operator may be ready to tolerate a small degradation during the interim failure period.
- o A congestion notification based Admission scheme has some inherent inaccuracies because of its reactive nature and thus may potentially over admit in some situations (such as burst of calls arrival). If the Preemption scheme reacted at the same Threshold as the Admission Threshold, calls may get routinely dropped after establishment because of over admission, even under steady state conditions.

These considerations argue for metering for Admission and Pre-emption at different traffic levels and hence, implicitly, for different markings and metering schemes.

Different marking schemes require different codepoints. Thus, such separate markings consume valuable real-estate in the packet header, especially scarce in the case of MPLS Pre-Congestion Notification [[I-D.davie-ecn-mpls](#)]. Furthermore, two different measurement/metering techniques involve additional complexity in the datapath of the internal routers of the PCN domain.

To this end, [[I-D.briscoe-tsvwg-cl-architecture](#)] proposes an approach, referred to as implicit pre-emption marking, that does not require separate pre-emption marking. However, it does require two separate measurement schemes: one measurement for Admission and another measurement for Preemption/Dropping. Furthermore, this approach mandates that the configured preemption rate be set to a drop rate. This approach effectively uses dropping as the way to convey information about how much traffic can "fit" under the preemption limit, instead of using a separate preemption marking. This is a significant restriction in that it results in preemption only taking effect once packets actually get dropped.

This document presents an approach that allows the use of a single PCN marking and a single metering technique at the internal devices without requiring that the dropping and pre-emption thresholds be the same. This document also investigates some of the tradeoffs associated with this approach.

1.3. Terminology

- o Pre-Congestion Notification (PCN): two algorithms that determine when a PCN-enabled router Admission Marks and Pre-emption Marks a packet, depending on the traffic level.
- o Admission Marking condition- the traffic level is such that the router Admission Marks packets. The router provides an "early warning" that the load is nearing the engineered admission control capacity, before there is any significant build-up in the queue of packets belonging to the specified real-time service class.
- o Pre-emption Marking condition- the traffic level is such that the router Pre-emption Marks packets. The router warns explicitly that pre-emption may be needed.
- o Configured-admission-rate - the reference rate used by the admission marking algorithm in a PCN-enabled router.
- o Configured-pre-emption-rate - the reference rate used by the pre-emption marking algorithm in a PCN-enabled router.
- o CLE - congestion level estimate computed by the egress node by estimating as the fraction of admission-marked packets it receives.
- o PIN - PCN internal node - an internal node in the PCN region.
- o PEN - PCN edge node - an ingress or egress edge node of the PCN region.
- o SAR - Sustainable Admission rate - the rate of unmarked traffic received by the Ingress PEN from the egress PEN
- o SPR - Sustainable Preemption rate

2. The Single Marking Approach

2.1. High Level description

The proposed approach is based on several simple ideas:

- o Replace virtual-queue-based marking for Admission Control by excess rate marking:
 - * meter traffic exceeding the Admission Threshold and mark excess traffic (e.g. using a token bucket with the rate configured to

Admission Rate Threshold)

- * at the edges, stop admitting traffic when the fraction of marked traffic for a given edge-to-edge aggregate exceeds a configured threshold (e.g. stop admitting when 3% of all traffic in the edge-to-edge aggregate received at the ingress is marked)
- o Impose a PCN-region-wide constraint on the ratio between the Admission threshold on a link and Pre-emption threshold on that link (e.g. pre-emption threshold is 20% higher than Admission threshold on all links in the PCN region)
- o The edge PCN device determines whether Pre-emption level is reached anywhere in the network **implicitly** by measuring the amount of unmarked traffic (assuming the marked traffic actually is above the threshold triggering blocking admission), i.e. the traffic that did not get admission marked. This is analogous to the notion of sustainable pre-emption rate in [[I-D.briscoe-tsvwg-cl-architecture](#)]. The rate of unmarked traffic in this case signifies the bottleneck admission threshold. The bottleneck pre-emption threshold is by design within the chosen ratio of that admission threshold. The ingress PCN edge device preempts all traffic above the bottleneck preemption threshold.

The remaining part of this section gives more detailed of a possible operation of the system.

2.2. Operation at the PIN

The PCN Internal Node (PIN) meters the aggregate PCN traffic and marks the excess rate. A number of implementations are possible to achieve that. A token bucket implementation is particularly attractive because of its relative simplicity, and even more so because a token bucket implementation is readily available in the vast majority of existing equipment. The rate of the token bucket is configured to correspond to the target Admission rate, and the depth of the token bucket can be configured by an operator based on the desired tolerance to PCN traffic burstiness.

Note that no preemption threshold is explicitly configured at the PIN, and the PIN does nothing at all to enforce it or mark traffic based on Pre-emption threshold.

2.3. Operation at the Egress PEN

The PCN Egress Node (PEN) measures the rate of both marked and unmarked traffic on a per-ingress PEN basis, and reports to the

ingress PEN two values: the rate of unmarked traffic from this ingress PEN, which we deem Sustainable Admission Rate (SAR) and the Congestion Level Estimate (CLE), which is the fraction of the marked traffic received from this ingress PEN. Note that Sustainable Admission Rate is analogous to the sustainable pre-emption rate of [\[I-D.briscoe-tsvwg-cl-architecture\]](#), except in this case it is based on the admission threshold rather than pre-emption threshold, while the CLE is exactly the same as that of [\[I-D.briscoe-tsvwg-cl-architecture\]](#). The details of the rate measurement are outside the scope of this draft.

[2.4.](#) Operation at the Ingress PEN

[2.4.1.](#) Admission Decision

Just as in [\[I-D.briscoe-tsvwg-cl-architecture\]](#), the admission decision is based on the CLE. The ingress PEN stops admission of new flows if the CLE is above a pre-defined threshold (e.g. 3%). Note that although the logic of the decision is exactly the same as in the case of [\[I-D.briscoe-tsvwg-cl-architecture\]](#), the detailed semantics of the marking is different. This is because the marking used for admission in this proposal reflects the excess rate over the admission threshold, while in [\[I-D.briscoe-tsvwg-cl-architecture\]](#), the marking is based on exceeding a virtual queue threshold. Notably, in the current proposal, if the average sustained rate of admitted traffic is 5% over the admission threshold, then 5% of the traffic is expected to be marked, whereas in the context of [\[I-D.briscoe-tsvwg-cl-architecture\]](#) a steady 5% overload should eventually result in 100% of all traffic being admission marked. A consequence of this is that for smooth traffic, the approach presented here will not mark any traffic at all until the rate of the traffic exceeds the configured admission threshold by the amount corresponding to the chosen CLE threshold. At first glance this may seem to result in a violation of the pre-congestion notification premise that attempts to stop admission before the desired traffic level is reached. However, in reality one can simply embed the CLE level into the desired configuration of the admission threshold. That is, if a certain rate X is the actual target admission threshold, then one should configure the rate of the metering device (e.g. the rate of the token bucket) to $X-y$ where y corresponds to the level of CLE that would trigger admission blocking decision.

A more important distinction is that virtual-queue based marking reacts to short-term burstiness of traffic, while the excess-rate based marking is only capable of reacting to rate violations at the timescale chosen for rate measurement. Based on our investigation, it seems that this distinction is not crucial in the context of PCN when no actual queuing is expected even if the virtual queue is full.

More discussion on this is presented later in the draft.

2.4.2. Pre-emption Decision

When the ingress observes a non-zero CLE and Sustainable Admission Rate (SAR), it first computes the Sustainable Pre-Emption rate (SPR) by simply multiplying SAR by the system-wide constant u , where u is the system-wide ratio between pre-emption and admission thresholds on all links in the PCN domain: $SPR = SAR * u$. The ingress PEN then performs exactly the same operation as is proposed in [\[I-D.briscoe-tsvwg-cl-architecture\]](#) with respect to SPR. Namely, it pre-empts the appropriate number of flows to ensure that the rate of traffic it sends to the corresponding egress PEN does not exceed SPR. Just as in the case of [\[I-D.briscoe-tsvwg-cl-architecture\]](#), an implementation may decide to slow down the pre-emption process by preempting fewer flows than is necessary to cap its traffic to SPR by employing a variety of techniques such as safety factors or hysteresis. In summary, the operation of pre-emption at the ingress PEN is identical to that of [\[I-D.briscoe-tsvwg-cl-architecture\]](#), with the only exception that the sustainable pre-emption rate is computed from the sustainable admission rate rather than derived from a separate marking. As discussed earlier, this is enabled by imposing a system-wide restriction on the pre-emption-to-admission thresholds ratio and changing the semantics of the admission marking.

3. Discussion

3.1. Benefits

The key benefits of using a single metering/marketing scheme for both Admission and Preemption presented in this document are summarized below:

- o Reduced implementation requirements on core routers due to a single metering implementation instead of two different ones.
- o Ease of use on existing hardware: given that the proposed approach is particularly amenable to a token bucket implementation, the availability of token buckets on virtually all commercially available routers makes this approach especially attractive.
- o Reduced number of codepoints which need to be conveyed in the packet header. If the PCN-bits used in the packets header to convey the congestion notification information are the ECN-bits in an IP core and the EXP-bits in an MPLS core, as is currently proposed in [\[put marking draft reference here\]](#) and [\[I-D.davie-ecn-mpls\]](#), those are very expensive real-estate. The

current proposals need 5 codepoints, which is especially important in the context of MPLS where there is only a total of 8 EXP codepoints which must also be shared with Diffserv. Eliminating one codepoint considerably helps.

- o A possibility of using a token-bucket-, excess-rate- based implementation for admission provides extra flexibility for the choice of an admission mechanism, even if two separate markings and thresholds are used.

3.2. Tradeoffs, Issues and Discussion

While the benefits of the proposed approach are attractive, there are several issues and tradeoffs that need to be carefully considered.

3.2.1. Restrictions on Pre-emption-to-admission Thresholds

An obvious restriction necessary for the single-marking approach is that the ratio of (implicit) pre-emption and admission thresholds remains the same on all links in the PCN region. While clearly a limitation, this does not appear to be particularly crippling, and does not appear to outweigh the benefits of reducing the overhead in the router implementation and savings in codepoints in the case of a single PCN domain, or in the case of multiple concatenated PCN regions. The case when this limitation becomes more inconvenient is when an operator wants to merge two previously separate PCN regions (which may have different admission-to-preemption ratios) into a single PCN region. In this case it becomes necessary to do a network-wide reconfiguration to align the settings.

3.2.2. Assumptions on Loss

Just as in the case of [[I-D.briscoe-tsvwg-cl-architecture](#)], the approach presented in this draft assumes that the Admission Threshold is configured at each link below the service rate of the traffic using PCN. This assumption is significant because the algorithm relies on the fact that if admission threshold is exceeded, enough marked traffic reaches the egress PEN to reach the configured CLE level. If this condition does not hold, then traffic may get dropped without ever triggering admission decision.

3.2.3. Effect of Reaction Timescale of Admission Mechanism

As mentioned earlier in this draft, there is a potential concern that slower reaction time of admissions mechanism presented in this draft compared to [[I-D.briscoe-tsvwg-cl-architecture](#)] may result in overshoot when the load grows rapidly, and undershoot when the load drops rapidly. While this is a valid concern theoretically, it

should be noted that at least for the traffic and parameters used in the simulation study reported here, there was no indication that this was a problem. However, the comparison has so far been done for Poisson arrivals only. Future work would look in performing a similar comparison for burstier arrivals.

3.2.4. Performance Implications and Tradeoffs

Replacement of a relatively well-studied queue-based measurement-based admission control approach by a cruder excess-rate measurement technique raises a number of algorithmic and performance concerns that need to be carefully evaluated. For example, a token-bucket excess rate measurement is expected to be substantially more sensitive to traffic burstiness and parameter setting, which may have a significant effect in the case of lower levels of traffic aggregation, especially for variable-rate traffic such as video. In addition, the appropriate timescale of rate measurement needs to be carefully evaluated, and in general it depends on the degree of expected traffic variability which is frequently unknown.

In view of that, an initial performance comparison of the token-bucket based measurement is presented in the following section. Within the constraints of this preliminary study, the performance tradeoffs observed between the queue-based technique suggested in [[I-D.briscoe-tsvwg-cl-architecture](#)] and a simpler token-bucket-based excess rate measurement do not appear to be a cause of substantial concern for cases when traffic aggregation is reasonably high at the bottleneck links as well as on a per ingress-egress pair basis. Details of the simulation study, as well as additional discussion of its implications are presented in [section 4](#).

Also, one mitigating consideration in favor of the simpler mechanism is that in a typical DiffServ environment, the real-time traffic is expected to be served at a higher priority and/or the target admission rate is expected to be substantially below the speed at which the real-time queue is actually served. If these assumptions hold, then there is some margin of safety for an admission control algorithm, making the requirements for admission control more forgiving to bounded errors - see additional discussion in [section 4](#).

3.2.5. Effect on Proposed Anti-cheating Mechanisms

Replacement of the queue-based admission control mechanism of [[I-D.briscoe-tsvwg-cl-architecture](#)] by an excess-rate based admission marking changing the semantics of the pre-congestion marking, and consequently interferes with mechanisms for cheating detection discussed in [[I-D.briscoe-tsvwg-re-ecn-border-cheat](#)]. Implications of excess-rate based marking on the anti-cheating mechanisms need to

be considered.

3.2.6. Standards Implications

The change of the meaning of admission marking for pre-congestion notification from the queue-based to excess-rate marking poses a question of coexistence of devices having different interpretation of admissions marking (and hence different metering and marking mechanisms in the core. The question of how and if the two mechanisms can co-exist in one PCN region has obvious impact on standardization efforts, and needs to be carefully considered.

4. Performance Evaluation Comparison

4.1. Relationship to other drafts

Initial simulation results of admission and pre-emption mechanisms of [[I-D.briscoe-tsvwg-cl-architecture](#)] were reported in [[I-D.briscoe-tsvwg-cl-phb](#)]. A follow-up study of these mechanisms is presented in a companion draft [draft-zhang-cl-performance-evaluation-01.txt](#). The current draft concentrates on a performance comparison of the admission control mechanism of [[I-D.briscoe-tsvwg-cl-phb](#)] and the token-bucket-based admission control described in [section 2](#) of this draft.

4.2. Limitations, Conclusions and Direction for Future Work

Due to time constraints, the study performed so far was limited to a small set of topologies, described in the Appendix. The key questions that have been investigated are the comparative sensitivity of the two schemes to parameter settings and the effect of traffic burstiness and of the degree of aggregation on a per ingress-egress pair on the performance of the admission control algorithms under study. The study is limited to the case where there is no packet loss. While this is a reasonable initial assumption for an admission control algorithm that is supposed to maintain the traffic level significantly below the service capacity of the corresponding queue, nevertheless future study is necessary to evaluate the effect of packet loss.

4.2.1. High Level Conclusions

The results of this (preliminary) study indicate that there is a potential that a reasonable complexity/performance tradeoff may be viable for the choice of admission control algorithm. In turn, this suggests that using a single codepoint and metering technique for admission and pre-emption may be a viable option.

The key high-level conclusions of the simulation study comparing the performance of queue-based and token-based admission control algorithms are summarized below:

1. At reasonable level of aggregation at the bottleneck and per ingress-egress pair traffic, both algorithms perform reasonably well for the range of traffic models considered (see [section 4.3](#) for detail).
2. Both schemes are stressed for small levels of ingress-egress pair aggregation levels of bursty traffic (e.g. a single video-like bursty SVD flow per ingress-egress pair). However, while the queue-based scheme results in tolerable performance even at low levels of per ingress-egress aggregation, the token-bucket-based scheme is substantially more sensitive to parameter setting than the queue-based scheme, and its performance for the high rate bursty SVD traffic with low levels of ingress-egress aggregation is quite poor unless parameters are chosen carefully to curb the error. It should be noted that the SVD traffic model used in this study is expected to be substantially more challenging for both admission and pre-emption mechanisms than the actual video traffic, as the latter is expected to be much smoother than the bursty on-off model with high peak-to-mean ratio we used. This expectation is confirmed by the fact that simulations with actual video traces reported in this version of the draft reveal that the performance of the video traces is much closer to that of VBR voice than of our crude SVD on-off model.
3. Even for small per ingress-egress pair aggregation, reasonable performance across a range of traffic models can be obtained for both algorithms (with a narrower range of parameter setting for the token-bucket based approach) .
4. The absolute value of round-trip time (RTT) or the RTT difference between different ingress-egress pair within the range of continental propagation delays does not appear to have a visible effect on the performance of both algorithms.

[4.2.2](#). Future work

This study is but the first step in performance evaluation of the token-bucket based admission control. Further evaluation should include a range of investigation, including the following:

- o fairness issues (how different ingress-egress pairs get access to bottleneck bandwidth)

- o interactions between admission control and preemption
- o effect of loss of marked packets
- o much more

5. [Appendix A](#): Simulation Details

5.1. Network and Signaling Model

Network topologies used in this study are shown in the Figures below. The network is modeled as either Single Link (Fig. A.1), Multi Link Network with a single bottleneck (termed "RTT", Fig. A.2), or a range of multi-bottleneck topologies shown in Fig. A.3 (termed "Parking Lot").

A --- B

Figure A.1: Simulated Single Link Network.

```

      A
      \
      B - D - F
      /
      C

```

Figure A.2: Simulated Multi Link Network.

A--B--C	A--B--C--D	A--B--C--D--E--F
D E F	E F G H	G H I J K L
(a)	(b)	(c)

Figure A.3: Simulated Multiple-bottleneck (Parking Lot)Topologies.

Figure A.1 shows a single link between an ingress and an egress node, all flows enter at node A and depart at node B. In Figure A.2, A set

of ingresses (A,B,C) connected to an interior node in the network (D). The number of ingresses varied in different simulation experiments in the range of 2-100. All links have generally different propagation delays, in the range 1ms - 100 ms. This node D in turn is connected to the egress (F). In this topology, different sets of flows between each ingress and the egress converge on the single link D-F, where pre-congestion notification algorithm is enabled. The capacities of the ingress links are not limiting, and hence no PCN is enable on those. The bottleneck link D-F is modeled with a 10ms propagation delay in all simulations. Therefore the range of round-trip delays in the experiments is from 22ms to 220ms. Our simulations concentrated primarily on capacities of 'bottleneck' links with sufficient aggregation - OC3 for voice and for SVD traffic, up to 1 Gbps. In the simulation model, a call requests arrive at the ingress and immediately sends a message to the egress. The message arrives at the egress after the propagation time plus link processing time (but no queuing delay). When the egress receives this message, it immediately responds to the ingress with the current Congestion-Level-Estimate. If the Congestion-Level-Estimate is below the specified CLE-threshold, the call is admitted, otherwise it is rejected. An admitted call sends packets according to one of the chosen traffic models for the duration of the call (see next section). Propagation delay from source to the ingress and from destination to the egress is assumed negligible and is not modeled.

Another type of network of interest is multi-bottleneck (or Parking Lot, PLT for short) topology. The simplest PLT with 2 bottlenecks is illustrated in Fig A.3(a). An example traffic matrix with this network on this topology is as follows:

- o an aggregate of "2-hop" flows entering the network at A and leaving at C (via the two links A-B-C)
- o an aggregate of "1-hop" flows entering the network at D and leaving at E (via A-B)
- o an aggregate of "1-hop" flows entering the network at E and leaving at F (via B-C)

In the 2-hop PLT shown in Fig. A.3(a) the points of congestion are links A--B and B--C. Capacity of all other links is not limiting. We also experiment with larger PLT topologies with 3 bottlenecks(see Fig A.3(b)) and 5 bottlenecks (Fig A.3 (c)). In all cases, we simulated one ingress-egress pair that carries the aggregate of "long" flows traversing all the N bottlenecks (where N is the number of bottleneck links in the PLT topology), and N ingress-egress pairs that carry flows traversing a single bottleneck link and exiting at the next "hop". In all cases, only the "horizontal" links in Fig.

A.3 were the bottlenecks, with capacities of all "vertical" links non-limiting. Propagation delays for all links in all PLT topologies are set to 1ms.

Due to time limitations, other possible traffic matrices (e.g. some of the flows traversing a subset of several bottleneck links) have not yet been considered and remain the area for future investigation. Our simulations concentrated primarily on the range of capacities of 'bottleneck' links with sufficient aggregation - above 10 Mbps for voice and 622 Mbps for SVD, up to 2.4 Gbps. But we also investigated slower 'bottleneck' links down to 512 Kbps in some experiments. In the simulation model of admission control, a call request arrives at the ingress and immediately sends a message to the egress. The message arrives at the egress after the propagation time plus link processing time (but no queuing delay). When the egress receives this message, it immediately responds to the ingress with the current Congestion Level Estimate. If the Congestion Level Estimate is below the specified CLE- threshold, the call is admitted, otherwise it is rejected. For preemption, once the ingress node of a PCN region decides to preempt a call, that call is preempted immediately and sends no more packets from that time on. The life of a call outside the domain described above is not modelled. Propagation delay from source to the ingress and from destination to the egress is assumed negligible and is not modelled.

5.2. Traffic Models

Four types of traffic were simulated (CBR voice, on-off traffic approximating voice with silence compression, and on-off traffic with higher peak and mean rates (we termed the latter synthetic video (SDV) as the chosen peak and mean rate was similar to that of an MPEG video stream, although no attempt was made to match any other parameters of this traffic to those of a video stream), and finally real video traces. The distribution of flow duration was chosen to be exponentially distributed with mean 1min, regardless of the traffic type. In most of the experiments flows arrived according to a Poisson distribution with mean arrival rate chosen to achieve a desired amount of overload over the configured-admission-limit in each experiment. Overloads in the range 1x to 5x and underload with 0.95x have been investigated. Note that the rationale for looking at the load 1and below is to see if any significant amount of "false rejects" would be seen (i.e. one would assume that all traffic should be accepted if the total demand is below the admission threshold). For on-off traffic, on and off periods were exponentially distributed with the specified mean. Traffic parameters for each type are summarized below:

5.2.1. CBR Voice (CBR)

- o Average rate 64 Kbps
- o Packet length 160 bytes
- o packet inter-arrival time 20ms

5.2.2. VBR Voice (VBR)

- o Packet length 160 bytes
- o Long-term average rate 21.76 Kbps
- o On Period mean duration 340ms; during the on period traffic is sent with the CBR voice parameters described above
- o Off Period mean duration 660ms; no traffic is sent during the off period.

5.2.3. "Synthetic Video": High Rate ON-OFF traffic with Video-like Mean and Peak Rates ("SVD")

This model is on-off traffic with video-like mean-to-peak ratio and mean rate approximating that of an MPEG-2 video stream. No attempt is made to simulate any other aspects of a video stream, and this model is merely that of on-off traffic. Although there is no claim that this model represents the performance of video traffic under the algorithms in question adequately, intuitively, this model should be more challenging for a measurement-based algorithm than the actual MPEG video, and as a result, 'good' or "reasonable" performance on this traffic model indicates that MPEG traffic should perform at least as well. We term this type of traffic SVD for "Synthetic Video".

- o Long term average rate 4 Mbps
- o On Period mean duration 340ms; during the on-period the packets are sent at 12 Mbps (1500 byte packets, packet inter-arrival: 1ms)
- o Off Period mean duration 660m

5.2.4. Real Video Traces (VTR)

We used a publicly available library of frame size traces of long MPEG-4 and H.263 encoded video obtained from <http://www.tkn.tu-berlin.de/research/trace/trace.html> (courtesy Telecommunication Networks Group of Technical University of Berlin).

Each trace is roughly 60 minutes in length, consisting of a list of records in the format of <FrameArrivalTime, FrameSize>. Among the 160 available traces, we picked the two with the highest average rate (averaged over the trace length, in this case, 60 minutes. In addition, the two also have a similar average rate). The trace file used in the simulation is the concatenation of the two. Since the duration of the flow is much smaller than the length of the trace, we need to check how does the expect rate of flow related the trace's long term average. To do so, we simulate a number of flows starting from random locations in the trace with duration chosen to be exponentially distributed with mean 1min. The results show that the expected rate of flow is roughly the same as the trace's average. Traffic characteristics are summarized below:

- o Average rate 769 Kbps
- o Each frame is sent with packet length 1500 bytes and packet inter-arrival time 1ms
- o No traffic is sent between frames.

5.3. Parameter Settings

5.3.1. Queue-based settings

All the queue-based simulations were run with the following Virtual Queue thresholds:

- o virtual-queue-rate: configured admission rate, 1/2 link speed
- o min-marking-threshold: 5ms at virtual-queue-rate
- o max-marking-threshold: 15ms at virtual-queue-rate
- o virtual-queue-upper-limit: 20ms at virtual-queue-rate

At the egress, the CLE is computed as an exponential weighted moving average (EWMA) on an interval basis, with 100ms measurement interval chosen in all simulations. We simulated the weight ranging 0.1 to 0.9. The CLE threshold is chosen to be 0.05, 0.15, 0.25, and 0.5.

5.3.2. Token Bucket Settings

The token bucket rate is set to the configured admission rate, which is half of the link speed in all experiments. Token bucket depth ranges from 64 to 512 packets. Our simulation results indicate that depth of token bucket has no significant impact on the performance of the algorithms and hence, in the rest of the section, we only present

the result with 256 packets bucket depth.

The CLE is calculated in the same way as in queue-based approach with weights from 0.1 to 0.9. The CLE thresholds are chosen to be 0.0001, 0.001, 0.01, 0.05 in this case. Note that the since meaning of the CLE is different for the Token bucket and queue-based algorithms, so there is no direct correspondence between the choice of the CLE thresholds in the two cases.

5.4. Simulation Details

To evaluate the performance of the algorithms, we recorded the actual admitted load at a granularity of 50ms, from which the mean admitted load over the duration of the simulation run can be computed. We verified that the actual admitted load at any time does not deviate much from the mean admitted load in each experiment by computing the coefficient of variation (CV is consistently 0.07 for CBR, 0.15 for VBR, 0.17 for VTR and 0.51 for SVD for all experiments). Finally, the performance of the algorithms is evaluated using a metric called over-admission-percentage, which is calculated as a percentage difference between the mean admitted load and the configured admission rate. Given reasonably small deviation of the admitted rate from the mean admitted in the experiments, this seems reasonable.

5.4.1. Queue-based Results

We found that the virtual-queue admission control algorithm works reliably with the range of parameters we simulated, for all four types of traffic. In addition, for both CBR, VBR and VTR traffic, the performance is insensitive to the parameters change under all tested topologies.

Table A.1 summarized over-admission-percentage values from 15 experiments with different [weight, CLE threshold] settings for CBR, VBR and VTR traffic. For the single bottleneck topologies (S. Link and RTT) the overload column represents the ratio of the demand on the bottleneck link to the configured admission threshold. For parking lot topologies we report the worst case result across all bottlenecks.

While in our simulations we tested the range of overload from 0.95 to 5, we present here only the results of the endpoints of this overload interval. For the intermediate values of overload the results are even closer to the expected ones than at the two boundary loads. The statistics show that for CBR, VBR and VTR traffic these over-admission-percentage values are rather similar across these traffic types, with the admitted load roughly staying within -2%+2% range of

the desired admission threshold, with quite limited variability across the experiments (see the Standard Deviation column). Note also that this is not always true for all experiments with 0.95 average load due to the variability of traffic generation.

Over Admission Perc Stats					Over	Topo	Type
Min	Mean	Max	SD	Load			
0	0	0	0	0.95	S.Link	CBR	
0.224	0.849	1.905	0.275	5			
0	0	0	0	0.95	RTT		
0.200	0.899	1.956	0.279	5			
-1.06	-0.33	-0.15	0.228	0.95	PLT		
-0.58	0.740	1.149	0.404	5			
-1.45	-0.98	-0.86	0.117	0.95	S.Link	VBR	
-0.07	1.405	1.948	0.421	5			
-1.56	-0.80	-0.69	0.16	0.95	RTT		
-0.11	1.463	2.199	0.462	5			
-3.49	-2.23	-1.80	0.606	0.95	PLT		
-1.37	0.978	1.501	0.744	5			
0	0	0	0	0.95	S.Link	VTR	
-0.53	1.004	1.539	0.453	5			
0	0	0	0	0.95	RTT		
-0.21	1.382	1.868	0.503	5			
0	0	0	0	0.95	PLT		
-0.86	0.686	1.117	0.452	5			

Table A.1 Summarized performance for CBR, VBR and VTR across different settings.

For SVD, the algorithms does show certain sensitivity to the tested parameters. Table A.2 recorded the over-admission-percentage for each combination of weights and CLE threshold.

		EWMA Weights					Over	Topo
		0.1	0.3	0.5	0.7	0.9	Load	
	0.05	-4.87	-3.05	-2.92	-2.40	-2.40		
	0.15	-3.67	-2.99	-2.40	-2.40	-2.40	0.95	
	0.25	-2.67	-2.40	-2.40	-2.40	-2.40		
								Single
C	0.05	-4.03	2.52	3.45	5.70	5.17		Link
L	0.15	-0.81	3.29	6.35	6.80	8.13	5	
E	0.25	2.15	5.83	6.81	8.62	7.95		
T	0.05	-11.77	-8.35	-5.23	-2.64	-2.35		
H	0.15	-9.71	-7.14	-2.01	-2.21	-1.13	0.95	
R	0.25	-5.54	-6.04	-3.28	-0.88	-0.27		
E								RTT
S	0.05	-5.04	-0.65	4.21	6.65	9.90		
S	0.15	-1.02	1.58	7.21	8.24	10.07	5	
H	0.25	-0.76	1.96	7.43	9.66	11.26		
E								
L	0.05	-2.51	-0.85	-0.63	0.025	-0.00		
D	0.15	-1.50	-0.63	-0.02	0.052	-0.02	0.95	
	0.25	-0.26	0.122	0.041	-0.02	0.132		
								PLT
	0.05	-3.50	0.422	1.899	3.339	3.770		
	0.15	-1.04	2.016	3.251	3.880	3.991	5	
	0.25	0.449	2.965	3.066	4.107	4.737		

Table A.2 Over-admission-percentage for SVD

It follows from these results that while performance is tolerable across the entire range of parameters, choosing the CLE and EWMA weights in the middle of the tested range appear to be more beneficial for the overall performance across the chosen range of overload, assuming the chosen values for the remaining parameters. The high level conclusion that can be drawn from Table A.2. is that (predictably) high peak-to-mean ratio SVD traffic is substantially more stressful to the queue-based admission control algorithm, but a set of parameters exists that keeps the overadmission within about -3% - +10% of the expected load even for the bursty SVD traffic.

Note that for SVD traffic these results hold even though there is no aggregation at all on a per-ingress-egress pair in the chosen RTT topology there is only a single SVD flow per ingress. In addition, there seems to be no visible influence by multiple bottleneck topology, as the results appears to be comparable to the ones in single bottleneck cases. (Note that it may seem from the data in A.2, PLT out-performs the Single Link. The reason is that the

bottleneck links in PLT happen to be twice the size of the ones in SingleLink cases. Given the same admission threshold, this implies that level of aggregation in PLT is twice as much as in the SingleLink. Hence, the better results in PLT should be viewed as the effect of the aggregation rather than the effect multi-bottleneck topology.)

5.4.2. Token Bucket-based Results

Just as for the case of the queue-based algorithm, under the under-loaded conditions for voice-like CBR, VBR and VTR traffic the sensitivity to the tested parameters remains limited for the token-bucket as well (the latter is summarized in Table A.3).

Over Admission Perc Stat				Load	Topo	Type	
Min	Mean	Max	SD				
0	0	0	0	0.95	S.Link		
0	0	0	0	0.95	RTT	CBR	
-1.17	-0.24	0.023	0.294	0.95	PLT		
-2.00	-0.78	-0.95	0.268	0.95	S.Link		
-2.83	-0.70	-1.20	0.510	0.95	RTT	VBR	
-6.19	-2.43	-1.32	1.223	0.95	PLT		
0	0	0	0	0.95	S.Link		
0	0	0	0	0.95	RTT	VTR	
0	0	0	0	0.95	PLT		

Table A.3 Summarized performance for CBR, VBR and VTR across different settings for under-loaded conditions.

However, for the overload conditions (overload greater than 1), the token bucket-based admission control algorithm shows substantially higher sensitivity to the parameter settings compared to the virtual queue based algorithm. Table A.4 shows over-admission-percentage for different settings. It is important to note here that for the token bucket-based admission control no traffic will be marked until the rate of traffic exceeds the configured admission rate by the chosen CLE. As a consequence, even with the ideal performance of the algorithms, the over-admission-percentage will not be 0, rather it is

expected to equal to CLE threshold if the algorithm performs as expected. Therefore, a more meaningful metric for the token-based results is actually the over-admission-percentage (listed below) minus the corresponding (CLE threshold * 100). For example, for CLE = 0.05, one would expect that 5% overadmission is inherently embedded in the algorithm, with the algorithm by design reacting to 5% overload (or more) only. Hence, with CLE = 0.05 a 10% over-admission in the token-bucket case should be compared to a 5% overadmission in the queue-based algorithm. When comparing the performance of token bucket (with the adjusted over-admission-percentage) to its corresponding virtual queue result, we found that token bucket performs only slightly worse for voice-like CBR VBR, and VTR traffic.

The results for SVD traffic require some additional commentary. Note from the results in Table A.4. that even for SVD traffic, in the Single Link topology the performance of the token-based solution is comparable to the performance of the queue-based scheme in table A.2, (adjusted by the CLE as discussed above). However, for the RTT topology, especially for the larger EWMA weights, the performance for SVD traffic becomes very bad, with up to 42% (adjusted by CLE) over-admission in a high overload situation (5x). We investigated two potential causes of this drastic degradation of performance by concentrating on two key differences between the Single Link and the RTT topologies: the difference in the round-trip times and the degree of aggregation in a per ingress-egress pair aggregate.

To investigate the effect of the difference in round-trip times, we also conducted a subset of the experiments described above using the RTT topology that has the same RTT across all ingress-egress pairs rather than the range of RTTs in one experiment. We found out that neither the absolute nor the relative difference in RTT between different ingress-egress pairs appear to have any visible effect on the over-load performance or the fairness of both algorithms (we do not present these results here as their are essentially identical to those in Table A.4). In view of that and noting that in the RTT topology we used for these experiments for the SVD traffic, there is only 1 highly bursty flow per ingress, we believe that the severe degradation of performance in this topology is directly attributable to the lack of traffic aggregation on the ingress-egress pair basis. We also note that even for this highly challenging scenario, it is possible to find a range of parameters that limit the overadmission case for SVD traffic to quite a reasonable range of -3% + 10% (adjusted by the CLE). Luckily, these are the same parameter settings that work quite well for the other types of traffic tested.

The overall performance on the multiple-bottleneck topology, for all types of traffic, is comparable to the ones on the SingleLink, just as shown in the queue-based results. Again, the results for the PLT

topology may appear better than those of the SingleLink in the table below. As discussed earlier, this is due to the fact that the bottleneck capacity of the PLT topology in these experiments is higher than that of the SingleLink and RTT, topologies (2.4Gbps vs 1Gbps), so the bottleneck aggregation is higher for PLT than for both single bottleneck topologies. The ingress-egress aggregation of the PLT topology corresponds to that of a SingleLink.

		EWMA Weights					Over	Topo	Type
		0.1	0.3	0.5	0.7	0.9	Load		
C	0.0001	-0.99	0.09	0.24	0.41	0.43			
L	0.001	0.02	0.37	0.43	0.46	0.45	5	S.	
E	0.01	1.37	1.32	1.32	1.31	1.31		Link	
T	0.0001	6.50	7.96	8.37	8.42	8.84			
H	0.001	7.07	8.54	8.65	8.55	8.66	5	RTT	CBR
R	0.01	7.93	9.08	8.71	8.63	9.40			
S									
H	0.0001	-1.88	0.21	0.58	0.65	0.62			
L	0.001	-0.31	0.69	0.63	0.61	0.56	5	PLT	
D	0.01	2.032	1.98	1.97	1.99	1.96			
C	0.0001	-2.95	-1.39	-0.63	0.23	0.78			
L	0.001	-1.51	-0.23	0.44	0.63	1.39	5	S.	
E	0.01	1.37	2.01	2.29	2.60	2.76		Link	
T	0.0001	-1.93	-0.23	0.99	2.09	3.38			
H	0.001	-0.75	0.89	2.07	3.12	4.27	5	RTT	VBR
R	0.01	1.91	3.42	4.35	5.36	6.38			
S									
H	0.0001	-3.78	-0.57	0.35	1.08	1.82			
L	0.001	-1.64	0.46	1.28	1.89	2.37	5	PLT	
D	0.01	1.871	2.74	3.24	3.39	3.86			
C	0.0001	-2.37	0.00	0.79	0.83	1.07			
L	0.001	-0.64	0.67	0.79	1.04	1.08	5	S.	
E	0.01	1.59	1.64	1.66	1.79	2.04		Link	
T	0.0001	-1.09	0.86	1.29	1.57	1.98			
H	0.001	0.00	1.39	1.63	1.86	2.58	5	RTT	VTR
R	0.01	1.99	2.32	2.88	3.54	4.35			
S									
H	0.0001	-1.99	0.09	0.58	0.96	1.15			
L	0.001	-0.55	0.57	1.09	1.30	1.49	5	PLT	

D	0.01	2.22	2.44	2.42	2.61	2.65				

	0.0001	-10.67	-10.58	-7.95	-6.27	-4.99				
	0.001	-8.67	-8.04	-7.61	-4.37	-2.89	0.95			
	0.01	-4.28	-2.59	-4.44	-2.13	-2.20		S.		
C	-----							Link		
L	0.0001	-16.36	-10.24	-6.50	-2.17	2.74				
E	0.001	-10.54	-5.63	-2.70	0.94	3.54	5			
	0.01	-4.11	1.26	5.38	5.75	8.82				

T	0.0001	-15.83	-10.35	-2.96	0.17	5.42				
H	0.001	-12.82	-7.62	-0.47	2.24	6.59	0.95			
R	0.01	-6.17	-0.11	2.16	5.28	10.34				
E	-----							RTT	SVD	
S	0.0001	-8.51	1.86	11.14	22.51	30.24				
S	0.001	-4.80	1.49	15.35	24.56	33.96	5			
H	0.01	0.56	8.26	25.71	35.63	42.72				
O	-----									
L	0.0001	-12.35	-8.73	-7.46	-5.97	-5.95				
D	0.001	-9.06	-7.66	-5.99	-5.19	-4.77	0.95			
	0.01	-5.13	-4.77	-4.62	-3.54	-3.52				
	-----							PLT		
	0.0001	-10.07	-5.05	-3.17	0.75	2.138				
	0.001	-7.41	-3.38	-0.53	1.89	5.071	5			
	0.01	-0.81	2.816	5.115	6.03	6.421				

Table A.4. Token bucket admission control performance.

An additional differences in token-bucket case vs the queue-based admissions in the PLT topology case revealed in our experiments is that there appears to be a consistent relationship between the position of the bottleneck link (how far downstream it is) and its over-admission-percentage. In Table A.5, we show a snapshot of the behavior with 5 bottleneck topology. Here, the over-admission-percentage displayed is an average across all 15 experiments with different [weight, CLE] setting. (We do observe the same behavior in each of the individual experiment, hence providing a summarized statistics does not invalidate the results). The data shows the further downstream the bottleneck is, the more it tends to over-admit, regardless the type of the traffic. The exact cause of this phenomenon is yet to be explained, but the effect of it seems to be insignificant in magnitude, at least in the experiments we ran.

Traffic	Bottleneck LinkId					Over	
Type	1	2	3	4	5	Load	

CBR	0.270	0.421	0.529	0.669	0.791	5	

VBR	0.152	0.512	0.715	0.901	1.169	5	

VTR	0.350	0.492	0.756	0.874	1.125	5	

Table A.5 Summarized performance for CBR, VBR and VTR across the links in the PLT topology for Token Bucket Admission

6. IANA Considerations

This document places no requests on IANA.

7. Security Considerations

TBD

8. References

8.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

8.2. Informative References

[I-D.briscoe-tsvwg-cl-architecture]
Briscoe, B., "An edge-to-edge Deployment Model for Pre-Congestion Notification: Admission Control over a DiffServ Region", [draft-briscoe-tsvwg-cl-architecture-04](#) (work in progress), October 2006.

[I-D.briscoe-tsvwg-cl-phb]
Briscoe, B., "Pre-Congestion Notification marking", [draft-briscoe-tsvwg-cl-phb-03](#) (work in progress), October 2006.

[I-D.briscoe-tsvwg-re-ecn-border-cheat]
Briscoe, B., "Emulating Border Flow Policing using Re-ECN on Bulk Data", [draft-briscoe-tsvwg-re-ecn-border-cheat-01](#)

(work in progress), June 2006.

[I-D.briscoe-tsvwg-re-ecn-tcp]

Briscoe, B., "Re-ECN: Adding Accountability for Causing Congestion to TCP/IP", [draft-briscoe-tsvwg-re-ecn-tcp-03](#) (work in progress), October 2006.

[I-D.davie-ecn-mpls]

Davie, B., "Explicit Congestion Marking in MPLS", [draft-davie-ecn-mpls-01](#) (work in progress), October 2006.

[I-D.lefaucheur-emergency-rsvp]

Faucheur, F., "RSVP Extensions for Emergency Services", [draft-lefaucheur-emergency-rsvp-02](#) (work in progress), June 2006.

Authors' Addresses

Anna Charny
Cisco Systems, Inc.
1414 Mass. Ave.
Boxborough, MA 01719
USA

Email: acharny@cisco.com

Xinyang (Joy) Zhang
Cisco Systems, Inc. and Cornell University
1414 Mass. Ave.
Boxborough, MA 01719
USA

Email: joyzhang@cisco.com

Francois Le Faucheur
Cisco Systems, Inc.
Village d'Entreprise Green Side - Batiment T3 ,
400 Avenue de Roumanille, 06410 Biot Sophia-Antipolis,
France

Email: flefauch@cisco.com

Vassilis Liatsos
Cisco Systems, Inc.
1414 Mass. Ave.
Boxborough, MA 01719
USA

Email: vliatsos@cisco.com

Full Copyright Statement

Copyright (C) The IETF Trust (2007).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).

