

Network Working Group
Internet-Draft
Intended status: Informational
Expires: May 22, 2008

J. Babiarz
X-G. Liu
K. Chan
Nortel
M. Menth
University of Wuerzburg
November 19, 2007

**Three State PCN Marking
draft-babiarz-pcn-3sm-01**

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 May 22, 2008.

Copyright Notice

Copyright (C) The IETF Trust (2007).

Abstract

This document proposes a mechanism for admission control and flow termination. It is based on the concept of pre-congestion notification (PCN) using three different codepoints: "no pre-congestion", "admission-stop", and "excess-traffic" for packet marking. Therefore, the proposal is called three state marking

(3sm). The behaviour of edge nodes is presented which distinguishes from other proposals through little complexity and its ability to cope with multipath routing. Algorithms required for packet metering and marking are explained in detail.

Table of Contents

- [1. Introduction](#) [5](#)
- [1.1. Requirements Notation](#) [6](#)
- [1.2. Terminology Used in this Document](#) [6](#)
- [1.3. Adapted Terminology](#) [6](#)
- [1.4. New Terminology](#) [7](#)
- [2. The 3sm Proposal](#) [8](#)
- [2.1. Packet Marking in 3sm](#) [8](#)
- [2.2. Admission Control in 3sm](#) [9](#)
- [2.2.1. Explicit Edge-to-Edge Tunnels \(E3Tunnel\)](#) [9](#)
- [2.3. End-to-End on-Path Signalling \(End2PS\)](#) [11](#)
- [2.3.1. Operation of Standard RSVP](#) [11](#)
- 2.3.2. Modification of Standard RSVP to Perform PCN-Based Admission Control [12](#)
- [2.4. Edge-to-Edge on-Path Signalling \(Edge2PS\)](#) [12](#)
- [2.5. Flow Termination in 3sm](#) [13](#)
- [2.5.1. Marked Flow Termination \(MFT\)](#) [13](#)
- [2.5.2. Measured Rate Termination \(MRT\)](#) [14](#)
- [3. Three State PCN Marker with Marking Frequency Reduction \(MFR\) for Marked Flow Termination \(MFT\)](#) [14](#)
- [3.1. ET-Marker](#) [15](#)
- [3.1.1. Behaviour of SR-Metering and ET-Marking](#) [15](#)
- [3.1.2. Pseudo Code for the ET-Marker](#) [16](#)
- [3.1.3. Configuration of the ET-Marker](#) [16](#)
- [3.1.4. Characteristics of the Proposed ET-Marker](#) [17](#)
- [3.2. AS-Marker](#) [17](#)
- [3.2.1. Behaviour of AR-Metering and AS-Marking](#) [18](#)
- [3.2.2. Pseudo Code for AS-Marker](#) [18](#)
- [3.2.3. Configuration of the AS-Marker](#) [18](#)
- [3.2.4. Characteristics of the Proposed AS-Marker](#) [19](#)
- [3.3. Marking Codepoints](#) [19](#)
- [4. Benefits and Shortcomings of the 3sm Proposal](#) [19](#)
- [4.1. Benefits](#) [20](#)
- [4.2. Shortcomings of 3sm](#) [21](#)
- [5. Security Considerations](#) [21](#)
- [6. Changes from Previous Revision](#) [21](#)
- [7. Acknowledgements](#) [22](#)
- [8. Informative References](#) [23](#)
- [Appendix A. Overview of Token Bucket \(TB\) and Virtual Queue \(VQ\)](#) [24](#)
- [A.1. New features for marking algorithm](#) [24](#)
- [A.1.1. Virtual Queue \(VQ\)](#) [24](#)
- [A.1.2. Token Bucket \(TB\)](#) [25](#)
- [A.1.3. Tail Marking](#) [25](#)
- A.1.4. Tail Marking with Marking Frequency Reduction (MFR) . 26
- A.1.5. Tail Marking with Packet Size Independent Marking (PSIM) and Proportional Marking Frequency

Reduction (PMFR) [27](#)
[A.1.6.](#) Threshold Marking [28](#)
Authors' Addresses [29](#)
Intellectual Property and Copyright Statements [31](#)

1. Introduction

Pre-Congestion Notification (PCN) builds on the concepts of [\[RFC3168\]](#), "The addition of Explicit Congestion Notification (ECN) to IP". It is used to implement admission control and flow termination for real-time flows (such as voice, video and multimedia streaming) in DiffServ [\[RFC2474\]](#), [\[RFC2475\]](#) enabled networks. Flow admission control determines whether a new flow can be added into the network without overloading any of its links, whereas flow termination reduces the current PCN traffic load by terminating marked flows when at least one link in the network is overloaded for some reason. For a general overview, the reader is referred to [\[I-D.ietf-pcn-architecture\]](#).

This document describes the 3sm proposal which is a special implementation of the general PCN framework [\[I-D.ietf-pcn-architecture\]](#). It relies on three different packet markings: "no pre-congestion" (NP), "admission-stop" (AS), or "excess-traffic" (ET). Packets enter a network with NP (unmarked). The basic idea is as follows. Each link in a PCN domain has an admissible rate (AR). If the current PCN traffic rate of a link exceeds its AR, all PCN packets on that link are re-marked with AS. The PCN egress nodes monitor the packet markings and trigger the PCN ingress nodes to admit or reject new flow requests depending on whether they observe AS-marked packets or not. Similarly, each link has a supportable rate (SR). If the current PCN traffic rate of a link exceeds its SR, some of the PCN packets on that link are re-marked with ET. The PCN egress nodes detect ET-marked packets and pass their flow IDs to the appropriate flow termination entity. This concept is called marked flow termination (MFT).

The 3sm architecture expects the above mentioned marking behaviour from PCN interior nodes. We propose simple metering and marking algorithms for that purpose. Those are threshold marking for AS-marking and tail marking with marking frequency reduction (MFR) for ET-marking. We describe them in [Section 3](#) based on a token bucket approach. To improve the termination behaviour of MFT, we suggest packet size independent marking (PSIM) and proportional MFR (PMFR) in the [Appendix A](#) and show that the same behaviour can be specified based on a virtual queue.

The document is structured as follows. After a short section on terminology, [Section 2](#) focuses on the edge behaviour of the 3sm proposal. [Section 3](#) provides detailed algorithms for the metering and marking behaviour expected from interior nodes in 3sm. [Section 4](#) list benefits and shortcomings of the 3sm proposal. [Appendix A](#) summarizes background information about token bucket (TB) and virtual queue (VQ) metering and marking as they are the base for the proposed

metering and marking mechanisms.

1.1. Requirements Notation

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 [[RFC2119](#)].

1.2. Terminology Used in this Document

The terminology used in this document conforms to the terminology of [[I-D.ietf-pcn-architecture](#)]. However, we adapted some terminology for better readability and added some new terminology that we found useful in this document.

1.3. Adapted Terminology

[[I-D.ietf-pcn-architecture](#)] has chosen very general terms for the rate thresholds as they have different semantic meanings in different proposals for the implementation of the general PCN architecture. For better readability, we use names that are more intuitive within the 3sm proposal.

- o Admissible rate (AR) - PCN-lower-rate
- o Supportable rate (SR) - PCN-upper-rate
- o AR-overload - Condition that PCN traffic rate exceeds AR of a link
- o SR-overload - Condition that PCN traffic rate exceeds SR of a link
- o No pre-congestion (NP) - Default marking for PCN packets that have not been carried over links with any type of pre-congestion (AR-overload or SR-overload).
- o Admission-stop (AS) - PCN-lower-rate-marking
- o Excess-traffic (ET) - PCN-upper-rate-marking
- o AS-marking - Action of re-marking of NP-marked packets to AS
- o ET-marking - Action of re-marking of NP- or AS-marked packets to ET

1.4. New Terminology

We provide a brief definition of the terminology used in this document.

- o PCN - Pre-Congestion Notification meters traffic rates per service class on a link and notifies the PCN egress nodes using packet marking whether a certain rate threshold is exceeded on any link of the path through the PCN-enabled network taken by the packet. The rate thresholds may be significantly lower than the line rates such that PCN egress nodes are notified long before queues build up in the buffers and real congestion occurs. PCN is intended for the implementation of measurement-based admission control and flow termination for real-time inelastic traffic, e.g., voice. The PCN marking in the packet headers need to be standardized.
- o ECN Field - Refers to the use of the standardized two bit field in the IP header that is used for signalling Explicit Congestion Notification [[RFC3168](#)]. In the PCN framework the ECN field maybe reused to signal two levels of PCN marking.
- o Service class - By service class we mean a grouping of packets belonging to one or more applications or services that generated traffic with similar characteristics and requiring similar QoS treatment. See [[RFC4594](#)] for details.
- o Admission Control - It is the function of admitting or blocking requests of new flows or sessions for access to the network to prevent AR-overload.
- o Flow Termination - It is the function of terminating already admitted flows in the sense that they cannot continue to send PCN traffic. Only flows contributing to SR-overload are terminated to reduce SR-overload.
- o Deployment model - A method to find the PCN information for the path of a new flow that requests admission to the PCN domain and to perform the required admission decision. Deployment models take advantage of information or protocols available in the networking scenario where PCN is to be deployed.
- o Marked flow termination (MFT) - Flow termination function terminating marked flows.
- o Measured rate termination (MRT) - Flow termination function terminating a certain rate of traffic that was directly or indirectly measured by the system.

- o Marking frequency reduction (MFR) - Marked flow termination (MFT) requires that only some instead of all PCN packets exceeding the supportable rate (SR) of a link are marked. This is achieved by MFR. This can be achieved by adding a slowdown factor of "s" tokens to the fill state of the token bucket whenever a packet is ET-marked.
- o Token bucket (TB) - One base mechanism for packet metering.
- o Virtual queue (VQ) - Another, equivalent base mechanism for packet marking.

2. The 3sm Proposal

First, a high-level overview of the packet marking in 3sm is presented. Then, the edge behaviour for the admission control and flow termination function is explained.

2.1. Packet Marking in 3sm

PCN traffic can be classified by DSCP or a group of DSCPs and is forwarded by an appropriated PHB. PCN configures for each link an admissible and a supportable rate (AR, SR). PCN traffic enters the network with a "no-precongestion" (NP) mark. PCN nodes meter the PCN traffic on every link. When the PCN traffic rate on a link exceeds the corresponding AR, the PCN node re-marks all NP-marked PCN packets to "admission-stop" (AS).

Similarly, they re-mark some non-ET-marked PCN packets to "excess-traffic" (ET) when the PCN traffic rate on a link exceeds the corresponding SR. Figure 1 summarizes the relation between the AR and SR thresholds and the marking behaviour. The SR normally is at least a delta above the AR and a delta below the maximum service rate for PCN traffic for the sake of stability of the measurement-based reactive system. If the PCN traffic rate is below AR, no packets are re-marked; if it is between AR and SR, all NP-marked PCN packets are re-marked to AS; and if it is above SR, some non-ET-marked PCN packets are re-marked to ET and all other NP-marked PCN packets are re-marked to AS. Hence, the meters and markers operate in a marking-aware mode: NP-marked packets can be re-marked to AS or ET, AS-marked packets can be re-marked to ET but not to NP, and ET-marked packets cannot be re-marked at all.

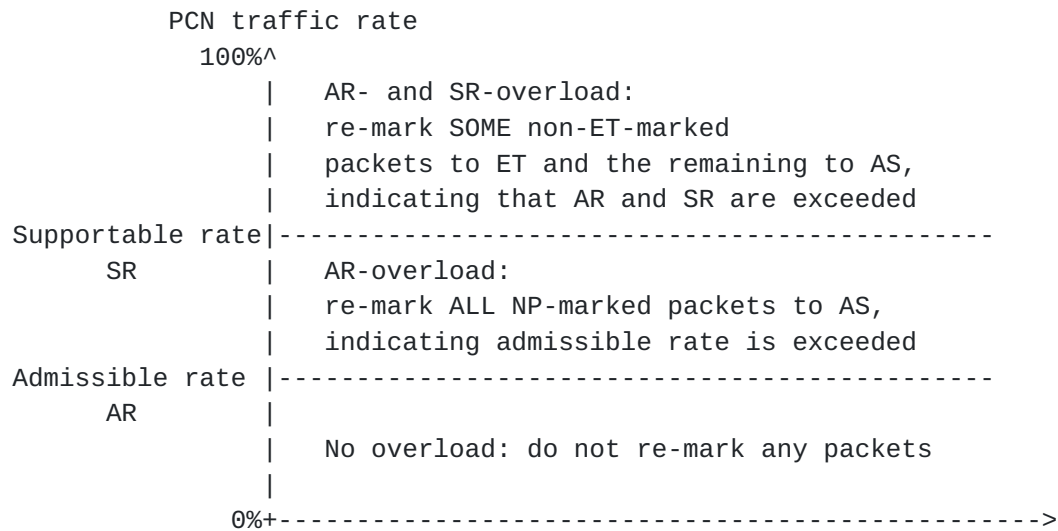


Figure 1: Packet re-marking by PCN nodes.

2.2. Admission Control in 3sm

The admission of a flow requests depends on the marking that is currently observed on the path the data packet of the future flow will take. However, it is not trivial to provide this information and may be achieved differently depending on the networking scenario. Therefore, we introduce three deployment models that use their own methods to get the PCN information for the path of a future data flow. These deployment models are adapted to special networking scenarios that are in use today and are named after the concept that helps to find the path of future data packets:

- o explicit edge-to-edge tunnels (E2Tunnel)
- o end-to-end on-path signalling (End2PS)
- o edge-to-edge on-path signalling (Edge2PS)

2.2.1. Explicit Edge-to-Edge Tunnels (E3Tunnel)

Explicit edge-to-edge tunnels provide an IP adjacency from a PCN ingress node to a PCN egress node. As a consequence, the information about the PCN egress node of a packet can be derived from the routing table. In addition, we assume that each tunnel is set up on an explicit path. This can be realized, e.g. by MPLS label switched paths (LSPs) or IP-in-IP tunnelling. We use the name E3Tunnel to abbreviate the term "explicit edge-to-edge tunnel" and as the name for the corresponding deployment model. An E3Tunnel aggregate is the ensemble of PCN traffic carries through the PCN domain through a common explicit E3Tunnel.

In the presence of E3Tunnels, the tunnel a packet is forwarded over is derived from the forwarding table of the PCN ingress node. This information is required that a new flow can request admission to the appropriate E3Tunnel. The PCN ingress and egress nodes have a context for each E3Tunnel they support. The PCN ingress node works per context in two different modes:

- o the reject mode and
- o the accept mode.

In the reject mode, the PCN ingress node rejects new admission requests while it admits them in its accept mode.

Similarly, the PCN egress node works per context in two corresponding modes:

- o the admission-stop mode and
- o the admission-continue mode.

The PCN egress node monitors the markings of the packets received from a specific E3Tunnel and depending on their marking, it switches to the admission-stop or admission-continue mode. Depending on its mode, the PCN egress node periodically sends "admission-stop" or "admission-continue" messages to the PCN ingress node belonging to the same E3Tunnel. If the PCN egress node receives an "admission-stop" message, it switches within the corresponding E3Tunnel context to its reject mode and if it receives an "admission-continue" message, it switches to its accept mode.

Different algorithms can be applied by the PCN egress node to decide when to send which control messages. We give two examples.

- o Option 1 (single-packet-based): If the PCN egress node observes a single marked packet from a specific E3Tunnel, it turns to the admission-stop mode. It returns to the admission-continue mode after some time unless it still observes marked packets occasionally. This option can be implemented without rate measurement and even without any counters requiring per-packet modification.
- o Option 2 (CLE-based): The PCN egress node tracks the number of marked and unmarked packets from a specific E3Tunnel within a measurement interval. It calculates the congestion level estimate (CLE) which is the fraction of the marked and all packets observed during this interval. If no packet was received within the last measurement interval, the PCN egress node switches to the

admission-continue mode. If the PCN egress node is in the admission-continue mode and the CLE of the last measurement interval is larger than a predefined admission-stop threshold, it switches to the admission-stop mode. If the PCN egress node is in the admission-stop mode and the CLE is smaller than a predefined admission-continue threshold, it switches to the admission-continue mode.

2.3. End-to-End on-Path Signalling (End2PS)

End-to-end on-path signalling sends PATH messages downstream to discover the path of future data packets and RESV messages upstream to trigger the admission request for the correct forward path using the gathered PATH information. The deployment model End2PS reuses the end-to-end on-path signalling protocol for probing on which the admission decision is based. We first explain the basic operation of standard RSVP and then adapt it to perform PCN-based admission control.

2.3.1. Operation of Standard RSVP

A popular protocol example is RSVP [[RFC2205](#)]. With RSVP, the data source issues a PATH message which is carried hop-by-hop over the same path future data packets will go. To that end, the PATH message uses the same source and destination address as future data packets and also all other header fields that are possible input for routing and load balancing decisions need to be the same. When a PATH message arrives at an RSVP-capable node, a PATH state is established pointing to the previous hop before the PATH message is forwarded further downstream. When the PATH message arrives at the destination, the destination triggers the end-to-end reservation for the flow by sending a RESV message upstream along the nodes that set up a PATH state. In these nodes, the RESV message is processed. In particular, resource admission control is performed for the new flow request and if it succeeds, the node forwards the RESV message to the previous hop recorded by the PATH state. This two pass signalling approach guarantees that the reservation is done on the downstream path of the future data flow. In contrast to PATH messages, RESV messages have the source address of the sending node and the destination address of the hop pointed to by the PATH state. That way, the information about the downstream next hop of the future data stream is conveyed to the previous hop and the flow-related information is stored in a RESV state. RSVP is a soft-state protocol, i.e., the PATH and RESV control messages are periodically sent to keep the PATH and RESV states alive and, thereby, the flow reservations. Admission control needs to be performed for a flow only once when no RESV state is set up, yet.

2.3.2. Modification of Standard RSVP to Perform PCN-Based Admission Control

We assume that interior nodes of a PCN domain are RSVP-disabled. That means that they just forward RSVP messages without processing them and PCN ingress and egress nodes are neighboring RSVP-capable nodes. As a consequence, PCN ingress nodes decide whether new flows can be admitted and carried through domain or not.

When the initial PATH message travels downstream, it is either marked or not, and eventually received by the PCN egress node. If no PATH state can be found for this flow at the PCN egress node, this PATH message is the first one and not a REFRESH message. If the PATH message is the first of the flow and if it is marked with "admission-stop", the RSVP engine sends back a PATHERR message to reject the flow. If the PATH message is not marked (NP), the RSVP PATH state is established at the PCN egress node and the PATH message is forwarded further downstream. REFRESH messages are just forwarded according to standard RSVP. When the PATH message arrives at the destination and a RESV message is sent. Eventually, the corresponding RESV message arrives at the PCN ingress node. When no RESV state is set up yet, this is the first RESV message and admission control must be performed. By the mere fact that the RESV message arrives, the PCN ingress node knows that the corresponding initial PATH message was not marked. Thus, it can admit any flow for which a new RESV message arrives.

A single probe is sufficient in 3sm because 3sm marks all packets when AR-overload occurs. Note that RSVP is only an example for End2PS, but End2PS works equally well with other two pass on-path signalling protocols.

2.4. Edge-to-Edge on-Path Signalling (Edge2PS)

In the absence of explicit edge-to-edge tunnels or other on-path signalling protocols, PCN information about the path future packets of a new flow will take is required for a qualified admission decision at the PCN ingress node. To that end, we propose to use an edge-to-edge on-path signalling protocol (Edge2PS).

Again, we assume that all interior nodes of the PCN domain are RSVP-disabled. When a request for the admission of a new flow arrives, e.g. signalled via SIP INVITE or other means, the PCN ingress and egress node act as RSVP proxies for the source and destination node of the data flow and set up a reservation request between PCN ingress and egress node. The source and destination address of the PATH messages are those of the actual data flow. This guarantees that they are carried on the same path as future data packet will be. The

PCN edge nodes do not forward the RSVP control messages outside the PCN domain. Instead, the PCN egress node returns a RESV message to the PCN ingress node. The RSVP messages created by PCN nodes on behalf of the flow need to be discerned somehow from regular end-to-end RSVP messages because they must not be forwarded outside the PCN domain.

With this addition, the above presented modification of standard RSVP can be used to admit flows based on a single probe message. Note that there is no stringent need to use RSVP for Edge2PS. A less complex two-pass protocol suffices.

2.5. Flow Termination in 3sm

Although flows are admitted only if the PCN traffic rate does not exceed the admissible rate (AR) on any link of their paths, it is possible that the PCN traffic rate on a link exceeds the SR, e.g., due to changed sending behaviour of admitted flows or due to route changes after a failure.

With 3sm two different options for flow termination can be supported: marked flow termination (MFT) and measured rate termination (MRT). We explain them in the following.

2.5.1. Marked Flow Termination (MFT)

Each PCN egress node monitors its received PCN traffic. If it detects an ET-marked packet, the corresponding PCN ingress node is identified and the PCN egress node sends the flow ID belonging to the ET-marked packet in a "traffic-reduction" message to the flow termination entity (e.g. the appropriate PCN ingress node) for termination. To save signalling overhead, several IDs may be signalled in a single "traffic reduction" message.

Marked flow termination can be applied with any deployment model. How the PCN ingress node of an ET-marked packet is derived depends on the deployment model:

- o E3Tunnel: the adjacency from which the packet is received defines the E3Tunnel such that the corresponding PCN ingress node is known.
- o With End2PS or Edge2PS, the PCN egress nodes can map the packets to RSVP reservations using RSVP classifiers, and the corresponding RSVP PATH state contains the address of the PCN ingress node.

Marked flow termination (MFT) requires that only some of the packets that exceed the supportable rate are ET-marked. Otherwise, MFT

terminates too many flows. With MFT, only a small fraction of the traffic is removed within a round-trip time, but the process continues if the PCN traffic rate still exceeds the supportable rate (SR). Therefore, the PCN traffic rate is gradually reduced until it drops below SR. Then, the flow termination process stops since no more packets are ET-marked.

2.5.2. Measured Rate Termination (MRT)

Measured rate termination can be applied only in the E3Tunnel deployment model. It requires that all packets exceeding SR are ET-marked. The PCN egress nodes measure the rate of ET-marked (ETR) packets per E3Tunnel and if it is larger than zero, it signals that ETR to the corresponding flow termination entity in a "traffic-reduction" message. The flow termination entity may be the corresponding PCN ingress node which then terminates a subset of the flows carried over the respective E3Tunnel such that the rate of these flows is about ETR. However, this is not an easy task since the actual rates of the flows are in general not known. If SR-overload continues in spite of the flow termination, the PCN egress node must wait some time before it sends a new "traffic-reduction" message to guarantee that the impact of the previous one is already reflected by the new measured rate of ET-marked traffic (ETR).

Measured rate termination can reduce SR-overload very quickly, but it has several drawbacks:

- o Rate measurement of ET-marked packets is complex.
- o The choice of the right subset of flows is difficult and requires that the rates of individual flows are known.

As marked flow termination (MFT) is simpler than measured rate termination (MRT), we propose MFT as preferred method for flow termination in 3sm.

3. Three State PCN Marker with Marking Frequency Reduction (MFR) for Marked Flow Termination (MFT)

The three state PCN marker (3sm) meters PCN packet streams per link and performs packet re-marking according to Figure 1. As a consequence the following three marking states are required:

- o no pre-congestion (NP),
- o admission-stop (AS), and

- o excess-traffic (ET).

In theory, the meter meters each packet and passes the packet and the metering result to the marker, and the marker marks packets according to the results of the meter. This is illustrated in Figure 2. The marking may be coded in the ECN field [[RFC3168](#)] of the packet for a specified PHB in a specific manner.




Figure 2: Block diagram of meter and marker function.

The behaviour of the two functions is often described by a single metering and marking algorithm. Therefore, we call the algorithm metering packets relative to admissible rate (AR) and re-marking them to AS simply AS-marker, and the algorithm metering packets relative to supportable rate (SR) and re-marking them to ET simply ET-marker.

In the following, we explain the behaviour of the ET- and AS-marker using token bucket (TB) based algorithms. The packet sizes counted by the meters and markers pertain to the size of the IP packet including its header bytes. Equivalent virtual queue (VQ) based algorithms are presented in [Appendix A](#). Other implementations approximating the described behaviours can be used.

3.1. ET-Marker

We describe the behaviour for SR-metering and ET-marking, present pseudo code, explain its configuration, and discuss its behaviour. We use object-oriented notation for most variables.

3.1.1. Behaviour of SR-Metering and ET-Marking

We propose an ET-marker based on a token bucket with tail marking and marking frequency reduction (see [Appendix A](#) for explanation of different options). The TB has a bucket of size TB.size which is continuously filled with tokens at rate TB.rate. When a PCN packet arrives, it is re-marked with "ET" if the fill state of the bucket (TB.fill) in tokens is smaller than its size (packet.size) in bytes and "s" additional tokens are added to the bucket; otherwise, the fill state is reduced by packet.size tokens. The slowdown parameter

"s" reduces the marking frequency of the algorithm.

3.1.2. Pseudo Code for the ET-Marker

The behaviour of the token bucket with tail marking and marking frequency reduction for SR-metering and ET-marking is expressed by the following pseudo code. It requires the time variable `TB.lastUpdate` indicating when the fill state of TB was last updated and a global variable "now" providing the current time. A PCN packet has the variables `packet.mark` showing its marking (NP, AS, ET) and `packet.size` showing its size.

Input: pcn packet

```
TB.fill = min(TB.size, TB.fill + TB.rate * (now - TB.lastUpdate));
TB.lastUpdate = now;
if (TB.fill < packet.size)
    packet.mark = ET;
    TB.fill = min(TB.size, TB.fill + s);
else
    TB.fill = TB.fill - packet.size;
endif
```

Output: void

Several enhancements of this algorithm are presented in [Appendix A.1.5](#). It marks packet independent of its size and applies MFR proportionally to the packet size. Early results show that this equalizes the termination probability for flows with different packet sizes and makes the time to remove the overload independent of packet sizes. In addition, MFR is also applied when packets are already ET-marked by previous nodes. This reduces the potential overtermination in case of multiple bottleneck links. These enhancements are simple, but still make the algorithm slightly more complex. Therefore, more performance results and discussions are needed to decide whether these enhancements should be standard marking behaviour or not [[I-D.babiarz-pcn-explicit-marking](#)], [[I-D.menth-pcn-performance](#)].

3.1.3. Configuration of the ET-Marker

The following parameters must be configured:

- o `TB.rate`: supportable rate (SR)
- o `TB.size`: supportable burst size (SBS), needs to be set appropriately

- o Slowdown parameter "s": needs to be set appropriately

3.1.4. Characteristics of the Proposed ET-Marker

The proposed algorithm can be applied with and without marking frequency reduction (MFR), i.e., $s > 0$ and $s = 0$, respectively.

a) No marking frequency reduction (noMFR, $s = 0$)

If the slowdown parameter is set to $s = 0$, MFR is switched off. As an alternative, a simplified version of the given algorithm can be used. If the PCN traffic rate on a link constantly exceeds its SR, the fill state of the TB decreases. Arriving packets for which the number of tokens in the bucket does not suffice are ET-marked. The size of the token bucket (supportable burst size (SBS)) controls how fast the marker reacts to a traffic rate above SR: if it is set to a low value, packets are already marked at a rate lower than SR in the presence of bursts, if it is set to a high value, marking starts delayed if the PCN traffic rate exceeds SR. A nice property of this option is that the rate of the ET-marked packets is exactly the rate of the traffic that exceeds SR (excess traffic rate) under the assumption that there was no packet loss. This observation can be used for "measured rate termination" (MRT): the rate of ET-marked traffic is measured per ingress-egress aggregate by the PCN egress node and signalled to the corresponding flow termination entity.

Measured rate termination (MRT) is also an option to realize flow termination in 3sm, but the preferred option for 3sm is marked flow termination (MFT). The drawback of "no MFR" arises in conjunction with MFT. Then too many packets are ET-marked and too many flows are terminated. That is the motivation to reduce the marking frequency by setting $s > 0$.

b) Marking frequency reduction (MFR, $s > 0$)

If the slowdown parameter is set to a value $s > 0$, MFR is achieved because for each marked packet up to additional "s" bytes that exceed SR can pass the ET-marker without being re-marked to ET. Thus, increasing the slowdown parameter "s" decreases the number of ET-marked packets in a short time period. In combination with MFT, a suitable "s" is required to achieve a fast termination of sufficiently many flows without terminating more flows than necessary.

3.2. AS-Marker

We explain the behaviour for AR-metering and AS-marking, present pseudo code, explain its configuration, and discuss its behaviour.

3.2.1. Behaviour of AR-Metering and AS-Marking

We propose an AS-marker based on a token bucket with threshold marking (see [Appendix A](#) for explanation of different options). The TB has a bucket of size TB.size which is continuously filled with tokens at rate TB.rate. The AR-meter and AS-marker consider only packets that are not ET-marked. When a non-ET-marked PCN packet arrives, it is re-marked to "AS" if the fill state of the bucket (TB.fill) in tokens is smaller than its size (packet.size) in bytes; otherwise, the fillstate is reduced by packet.size tokens and if the fill state is then smaller than the marking threshold (TB.threshold), the packet is also re-marked to "AS" while if the fill state is then larger than or equal to the marking threshold, the packet is not re-marked.

The AS-marker is sensitive to ET-markings in the sense that only non-ET-marked packets are considered for remarking.

3.2.2. Pseudo Code for AS-Marker

The behaviour of the token bucket with threshold marking for AR-metering and AS-marking is expressed by the following pseudo code using the same nomenclature like above.

Input: pcn packet

```
TB.fill = min(TB.size, TB.fill+TB.rate*(now-TB.lastUpdate));
TB.lastUpdate = now
if (packet.mark <> ET) //mark only non-ET-marked packets
    if (TB.fill < TB.threshold)
        packet.mark = AS;
    endif
    TB.fill = max(0, TB.fill - packet.size);
endif
```

Output: void

3.2.3. Configuration of the AS-Marker

The following parameters must be configured:

- o TB.rate: admissible rate (AR)
- o TB.size (TBS): admissible burst size (ABS) needs to be set appropriately
- o TB.threshold: needs to be set appropriately

3.2.4. Characteristics of the Proposed AS-Marker

If the AR is exceeded, the TB fill state continuously decreases, it eventually falls below its marking threshold TB.threshold, and it only increases again if the PCN traffic rate on the link falls below AR. As a consequence, all packets are AS-marked during that time and admission of further flows is stopped until the PCN traffic rate drops below AR. In particular, all probe packets are AS-marked.

[TR437] investigates the impact of the TB parameters on the marking characteristics, i.e. the percentage of marked packets depending on the PCN traffic rate relative to AR. If the PCN traffic rate is above AR, all packets must be marked with AS. To that end, TB.threshold must be set large enough. If the PCN traffic rate is below AR, no packets should be marked. To that end TB.size-TB.threshold must be set large enough. Then, we get marking with clear decisions, i.e. packets are marked if the PCN rate is above AR and they are not marked if the PCN rate is below AR. This type of marking is required for 3sm. [[I-D.menth-pcn-performance](#)] provides a summary of [TR437].

3.3. Marking Codepoints

PCN metering and marking requires classification of traffic which is subject to PCN metering and PCN marking (PCN-capable). Furthermore, PCN-aware flows that are subject to PCN-marking require at least the following codepoints:

- o "no-precongestion" (NP),
- o "admission-stop" (AS), and
- o "excess-traffic" (ET).

These signals may be encoded by re-using the two-bit ECN field or by different DS codepoints. The actual encoding is out of the scope of this document.

4. Benefits and Shortcomings of the 3sm Proposal

This section highlights some benefits of the 3sm proposal that we think are important. Some of them are based on performance results reported in [[SIM-07](#)], [[I-D.babiarz-pcn-explicit-marking](#)], [[I-D.menth-pcn-performance](#)], and [TR437].

[4.1.](#) Benefits

- o 3sm does not require the standardization of the exact algorithm in the PCN interior node but only the standardization of the behavior.
- o 3sm does not require periodic transmission of measurement results from PCN egress to PCN ingress.
- o The supportable rate of a link can be chosen independently of the admissible rate as long as it is not smaller. This is maximum freedom and does not imply any degradation in resource efficiency for resilient admission control [[PCN-Config](#)].
- o 3sm works well with multipath routing due to marked flow termination and the application of probing when needed.
 - * A flow is terminated only if it is carried over an SR-overloaded path and if one of its packets was ET-marked. Therefore, its termination reduces the SR-overload on a bottleneck link.
 - * Conversely, flows that are not carried over a bottleneck link cannot be ET-marked and, therefore, not terminated.
- o Marked flow termination is very robust
 - * The slowdown factor "s" can be reasonably chosen such that the time to remove the overload is short and that not more traffic than necessary is terminated.
 - * It works well with low and high packet loss. Packet loss just increases the time to remove the overload.
 - * It works well with a small and large number of flows. In particular, the right number of flows is terminated if there is only one flow per aggregate and an SR-overload of, e.g., 30%.
 - * It works well with constant bit rate (CBR) voice, on/off voice, and variable bit rate (VBR) traffic. It is not very sensitive to different traffic characteristics [[SIM-07](#)].
 - * It works well with multiple bottleneck links in the path of a flow.
 - * It works well with flows that have significantly different round-trip times (RTT) or flow termination delays. In particular, flows with different RTTs suffer the same flow

termination probability.

- * In case of terminating bidirectional flows, the termination of a flow entails the termination of the downstream and upstream flow. In case of overload on both the downstream and upstream path, flows are terminated on both path. This leads to increased termination aggressiveness. However, marked flow termination leads only to little overtermination as it terminates traffic only gradually.
- * Different slowdown factors "s" may be used for the configuration of marking frequency reduction on different links within a single PCN domain.
- * It leads to higher termination probabilities for flows with larger packets. This can be repaired by packet size independent marking (PSIM, in [Appendix A.1.5](#))

4.2. Shortcomings of 3sm

- o Marked flow termination leads to higher termination probabilities for flows with larger packet frequency (higher rate flows). Currently, there is no mechanism to improve this.

5. Security Considerations

The Three State PCN Marker has no known security concerns.

6. Changes from Previous Revision

Changes in version -01 compared to version -00:

- * Abstract: adapted: more general, better summary
- * shortened introduction and adapted it to new document structure
- * put "1.3 Terminology" into separate section, clarification and some minor changes
- * put "1.2 Overview of PCN" into a separate section on 3sm edge behaviour
- * added more structure and content to that section
- o deployment scenarios for admission control

- o added measured rate termination (MRT) as a non-preferred option to marked flow termination (MFT) in 3sm
- * introduced AS-meter as abbreviation for AR-metering and AS-marking function
- * introduced ET-meter as abbreviation for SR-metering and ET-marking function
- * removed comment about similarity of 3sm marking and srTCM ([RFC2697](#))
- * current 3.1 did some reformulation and added some new simulation insights, moved optimizations into [Appendix A.1.5](#)
- * current 3.2 did some reformulation and added some new simulation insights, changed the AS-marker according to
- o Joe Babiarz's comment saying all packets need to be considered for AR-metering
- o Bob Brisco's comment saying TB.fill should be set to zero if it is smaller than packet.size
- o provided insights for the setting of TB parameters to obtain desired marking behaviour
- * new section on benefits of the 3sm proposal added
- * A.1 - A.1.4 minor changes (rewording)
- * A.1.5 new section on optimization for 3sm marker (PSIM, PMFR, MFR for already marked packets)
- * A.1.6 simplified threshold marking according to Bob's comments
- * old A.6 Section on Related Work (srTCM) removed as srTCM without change cannot be reused for implementation of threshold marking
- * removed [Appendix B](#) and integrated the comments into the previous sections of the document if necessary

[7.](#) Acknowledgements

The authors would like to thank the following people for reviewing this draft or earlier versions thereof and for their suggestions to make this document more complete: Dave McDysan, Nicolas Chevrollier, Frank Lehrieder, Bob Brisco, and Ben Strulo.

8. Informative References

- [I-D.babiarz-pcn-explicit-marking]
Liu, X. and J. Babiarz, "Simulations Results for 3sM",
[draft-babiarz-pcn-explicit-marking-01](#) (work in progress),
July 2007.
- [I-D.ietf-pcn-architecture]
Eardley, P., "Pre-Congestion Notification Architecture",
[draft-ietf-pcn-architecture-01](#) (work in progress),
October 2007.
- [I-D.menth-pcn-performance]
Menth, M. and F. Lehrieder, "Performance Evaluation of
PCN-Based Algorithms", [draft-menth-pcn-performance-00](#)
(work in progress), November 2007.
- [Maglaris-88]
Maglaris et al, "Performance Models of Statistical
Multiplexing in Packet Video Communications, IEEE
Transactions on Communications 36, pp. 834-844", July
1988.
- [PCN-Config]
Menth, M. and M. Hartmann, "PCN-Based Resilient Network
Admission Control: The Impact of a Single Bit", ([http://
www3.informatik.uni-wuerzburg.de/staff/menth/Publications/
Menth07-PCN-Config.pdf](http://www3.informatik.uni-wuerzburg.de/staff/menth/Publications/Menth07-PCN-Config.pdf)"), 2007.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate
Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2205] Braden, B., Zhang, L., Berson, S., Herzog, S., and S.
Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1
Functional Specification", [RFC 2205](#), September 1997.
- [RFC2474] Nichols, K., Blake, S., Baker, F., and D. Black,
"Definition of the Differentiated Services Field (DS
Field) in the IPv4 and IPv6 Headers", [RFC 2474](#),
December 1998.
- [RFC2475] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z.,
and W. Weiss, "An Architecture for Differentiated
Services", [RFC 2475](#), December 1998.
- [RFC3168] Ramakrishnan, K., Floyd, S., and D. Black, "The Addition
of Explicit Congestion Notification (ECN) to IP",
[RFC 3168](#), September 2001.

- [RFC4594] Babiarz, J., Chan, K., and F. Baker, "Configuration Guidelines for DiffServ Service Classes", [RFC 4594](#), August 2006.
- [SIM-07] Liu, X-G. and J. Babiarz, "Simulation Results for Explicit PCN Marking and Flow Termination (http://standards.nortel.com/pcn/Simulation_EPCN.pdf)", February 2007.
- [TR437] Menth, M. and F. Lehrieder, "Comparison of Marking Algorithms for PCN-Based Admission Control, Technical Report No. 437, (<http://www-info3.informatik.uni-wuerzburg.de/TR/tr437.pdf>)", October 2007.

Appendix A. Overview of Token Bucket (TB) and Virtual Queue (VQ)

Token buckets (TB) and virtual queues (VQ) are equivalent base mechanisms for algorithms to control whether a packet is conform to its flow's indicated rate R and burst size S . Therefore, TB parameters are frequently used as traffic descriptors. TBs and VQs are dual approaches: while packets are TB-conform as long as sufficient tokens are in the bucket at their arrival times, they are VQ-conform as long as sufficient free space is available in the queue at their arrival times. Therefore, TBs and VQs can be used interchangeably and, in particular, algorithms given based on a TB description can be implemented by a VQ and vice-versa.

In the following, we explain the basic VQ and TB mechanisms (Appendix A.1.1 and [Appendix A.1.2](#)). Packets are marked depending on the state of the VQ or TB at their arrival time. There are different marking options. Only those packets that are not conform to its flow description may be marked (tail marking, [Appendix A.1.3](#)), or only some non-conforming packets may be marked (tail marking with marking frequency reduction, [Appendix A.1.4](#)), or all packets may be marked until the flow again reaches conformity (threshold marking, [Appendix A.1.6](#)).

A.1. New features for marking algorithm

A.1.1. Virtual Queue (VQ)

We use an object-oriented notation for a more intuitive readability of the algorithms. The VQ has a VQ rate (VQ.rate) and a queue which is capable to store up to VQ.size bytes. The current length of the queue is denoted by VQ.length. This length is reduced over time at rate VQ.rate. When a packet arrives, it is "accepted" by the VQ and

increments VQ.length by its size (packet.size) if there is still enough free space in the queue to accommodate it; otherwise it is "rejected". As the queue size is decreased continuously over time, the behaviour of a VQ is best described by a fluid model. However, the state of the VQ shortly after packet arrivals can be calculated based on the current time "now" and the length of the VQ at the last update time of the VQ (VQ.lastUpdate) using the following algorithm:

```
VQ.length = max(0, VQ.length - (now - VQ.lastUpdate) * VQ.rate);
VQ.lastUpdate = now;
if (VQ.length + packet.size <= VQ.size)
    VQ.length = VQ.length + packet.size;
endif
```

A.1.2. Token Bucket (TB)

The TB is basically the same mechanism, but it looks at the problem from a different angle. The TB has a rate (TB.rate) and a bucket which is capable to store up to TB.size tokens. A token is the permission to send one byte. The current fill state of the bucket is denoted by TB.fill. This fill state is increased over time at rate TB.rate. When a packet arrives, it is "accepted" and decrements TB.fill by its size (packet.size) if there are enough tokens in the bucket to send the entire packet; otherwise it is "rejected". As the fill state is increased continuously over time, the behaviour of a TB is best described by a fluid model. However, the state of the TB shortly after packet arrival can be calculated based on the current time "now" and the fill state of the TB at the last update time of the TB (TB.lastUpdate) using the following algorithm:

```
TB.fill = min(TB.size, TB.fill + (now - TB.lastUpdate) * TB.rate);
TB.lastUpdate = now;
if (TB.fill >= packet.size)
    TB.fill = TB.fill - packet.size;
endif
```

A.1.3. Tail Marking

To control whether packets of a stream with rate R and maximum burst size MBS are conform to the description R and MBS, the stream is metered either by a VQ with VQ.rate=R and VQ.size=MBS, or by a TB with TB.rate=R and TB.size=MBS. If a packet is accepted by the VQ or by the TB, it is marked in-profile. If it is rejected, it is marked out-of-profile.

The corresponding pseudo codes are for the VQ:


```
VQ.length = max(0, VQ.length - (now - VQ.lastUpdate) * VQ.rate);
VQ.lastUpdate = now;
if (VQ.length + packet.size <= VQ.size)
    VQ.length = VQ.length + packet.size;
    packet.mark = in-profile;
else
    packet.mark = out-of-profile;
endif
```

and for the TB:

```
TB.fill = min(TB.size, TB.fill + (now - TB.lastUpdate) * TB.rate);
TB.lastUpdate = now;
if (TB.fill >= packet.size)
    TB.fill = TB.fill - packet.size;
    packet.mark = in-profile;
else
    packet.mark = out-of-profile;
endif
```

A.1.4. Tail Marking with Marking Frequency Reduction (MFR)

The objective of tail marking with MFR is to mark only some of the packets that are out-of-profile. The strength of the reduction can be controlled by the slowdown parameter "s". When a packet is classified out-of-profile, the VQ length is decremented by "s" bytes and the TB fill state is incremented by "s" tokens, respectively. As a consequence, the VQ and the TB are not likely to mark consecutive packets as out-of-profile which reduces their marking frequency.

The corresponding pseudo codes are for the VQ:

```
VQ.length = max(0, VQ.length - (now - VQ.lastUpdate) * VQ.rate);
VQ.lastUpdate = now;
if (VQ.length + packet.size <= VQ.size)
    VQ.length = VQ.length + packet.size;
    packet.mark = in-profile;
else
    VQ.length = max(0, VQ.length-s); //marking frequency reduction
    packet.mark = out-of-profile;
endif
```

and for the TB:


```
TB.fill = min(TB.size, TB.fill + (now - TB.lastUpdate) * TB.rate);
TB.lastUpdate = now;
if (TB.fill >= packet.size)
    TB.fill = TB.fill - packet.size;
    packet.mark = in-profile;
else
    TB.fill = min(TB.size, TB.fill+s) //marking frequency reduction
    packet.mark = out-of-profile;
endif
```

If the slowdown parameter is set to $s=0$, the marking algorithm behaves like pure tail marking.

A.1.5. Tail Marking with Packet Size Independent Marking (PSIM) and Proportional Marking Frequency Reduction (PMFR)

For the sake of fairness, large packets should not have a larger packet marking probability; otherwise this creates incentives to send many small packets instead of a few large packets. Therefore, we propose to test a packet arrival whether an MTU can still be supported. Then, the marking probability depends only on $VQ.length$ or $TB.fill$, respectively.

For the sake of better control of the termination aggressiveness, more flows should be marked and terminated in the presence of low bit rate flows than in the presence of high bit rate flows. To that end, we propose to remove (VQ) or add (TB) the slowdown factor proportionally to the packet size.

The reason for adding "s" tokens to the bucket when a packet is marked is the anticipation of the termination of that flow. This is done to approximate the fill state with this flow already being terminated in order to avoid that too many additional flows are marked. The same condition is met when an already ET-marked packet arrives at the ET-marker. Therefore, the same slowdown factor "s" is added to the bucket of the ET-marker as if it had ET-marked the packet itself. This is done to avoid that more traffic than necessary is terminated in case that a traffic aggregate crosses several bottleneck links. However, this situation is rather unlikely and the effect is well visible only under pathologic conditions. Therefore, optimal behaviour is not crucial and the algorithm may be simplified by ignoring packets that are already ET-marked.

These are obvious enhancements of the base algorithm, but they are only recommended when solid performance results indicate that the improvements have a significant impact in practice. These studies are still ongoing.

The following pseudo codes incorporate all proposed enhancements. We give them for a VQ approach:

```
VQ.length = max(0, VQ.length - (now - VQ.lastUpdate) * VQ.rate);
VQ.lastUpdate = now;
if (packet.mark<>ET)
    if (VQ.length + MTU <= VQ.size)                // PSIM
        VQ.length = VQ.length + packet.size;
    else
        VQ.length = max(0, VQ.length-s*packet.size/MTU); // PMFR
        packet.mark = ET;
    endif
else
    VQ.length = max(0, VQ.length-s*packet.size/MTU); // PMFR
endif
```

and for the TB:

```
TB.fill = min(TB.size, TB.fill + (now - TB.lastUpdate) * TB.rate);
TB.lastUpdate = now;
If (packet.mark<>ET)
    if (TB.fill >= MTU)                            // PSIM
        TB.fill = TB.fill - packet.size;
    else
        TB.fill = min(TB.size, TB.fill+s*packet.size/MTU) // PMFR
        packet.mark = ET;
    endif
else
    TB.fill = min(TB.size, TB.fill + s*packet.size/MTU); // PMFR
endif
```

[A.1.6.](#) Threshold Marking

The objective of threshold marking is to mark all packets with, e.g., "rate-exceeded" as long as some packets are out-of-profile with respect to flow parameters R and MBS. We achieve that by setting the VQ or TB size larger than MBS. Packets are marked if the VQ length exceeds MBS or if the fill state of the TB falls below TB.size-MBS. Furthermore, the packet size is always added to the queue or removed from the bucket. Note that marking thresholds need to be configured differently for VQ and TB to obtain the same behaviour:

- o VQ.threshold = MBS;
- o TB.threshold = TB.size - MBS.

The corresponding pseudo codes are for the VQ:


```
VQ.length = max(0, VQ.length - (now - VQ.lastUpdate) * VQ.rate);
VQ.lastUpdate = now;
if (VQ.length > VQ.threshold)
    packet.mark = "rate-exceeded"; // packet out-of-profile
endif
VQ.length = min(VQ.size, VQ.length + packet.size);
```

and for the TB:

```
TB.fill = min(TB.size, TB.fill + (now - TB.lastUpdate) * TB.rate);
TB.lastUpdate = now;
TB.fill = TB.fill - packet.size;
if (TB.fill < TB.threshold)
    packet.mark = "rate-exceeded"; // packet out-of-profile
endif
TB.fill = max(0, TB.fill - packet.size);
```

Authors' Addresses

Jozef Z. Babiarz
Nortel
3500 Carling Avenue
Ottawa, Ont. K2H 8E9
Canada

Phone: +1-613-763-6098
Email: babiarz@nortel.com

Xiao-Gao Liu
Nortel
3500 Carling Avenue
Ottawa, Ont. K2H 8E9
Canada

Phone: +1-613-763-7516
Email: xgliu@nortel.com

Kwok Ho Chan
Nortel
600 Technology Park Drive
Billerica, MA 01821
USA

Phone: +1-978-288-8175
Email: khchan@nortel.com

Dr. Michael Menth
University of Wuerzburg
Institute of Computer Science
Am Hubland, D-97074 Wuerzburg, Room B206
Germany

Phone: (+49)-931/888-6644
Email: menth@informatik.uni-wuerzburg.de

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).

