

PCN
Internet-Draft
Intended status: Informational
Expires: April 29, 2010

K. Chan
Huawei Technologies
G. Karagiannis
University of Twente
T. Moncaster
BT Research
M. Menth
University of Wurzburg
P. Eardley
B. Briscoe
BT Research
October 26, 2009

**Pre-Congestion Notification Encoding Comparison
draft-ietf-pcn-encoding-comparison-01**

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [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 April 29, 2010.

Copyright Notice

Copyright (c) 2009 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents in effect on the date of publication of this document (<http://trustee.ietf.org/license-info>).

Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Abstract

A number of mechanisms have been proposed to support differential Quality of Service for packets in the Internet. DiffServ is an example of such a mechanism. However, the level of assurance that can be provided with DiffServ without substantial over-provisioning is limited. Pre-Congestion Notification (PCN) uses path congestion information across a PCN region to enable per-flow admission control to provide the required service guarantees for the admitted traffic. While admission control will protect the QoS under normal operating conditions, an additional flow termination mechanism is necessary to cope with extreme events (e.g. route changes due to link or node failure).

In order to allow the PCN mechanisms to work it is necessary for IP packets to be able to carry the pre-congestion information to the PCN egress nodes. This document collects the lessons learned as we explore the different ways in which this information can be encoded into IP packets. This document does not choose the encoding but provide information on trade offs with the encoding choices, providing guidance based on different criteria. This document provides a historical trace of the consideration on different encoding alternatives for Pre-Congestion Notification.

Table of Contents

1.	Introduction	5
2.	Encoding Environment	5
3.	Encoding Requirements	6
3.1.	Encoding States	6
3.2.	Encoding and Operating Environment	7
3.2.1.	PCN Capable (Non PCN Capable) Packet Encoding State .	8
3.2.2.	Nonce Encoding State	9
3.2.3.	Non-PCN Traffic Entering PCN Domain	9
3.2.4.	PCN Traffic Leaving PCN Domain	10
3.2.5.	PCN Encoding for Both Edge to Edge and End to End Deployment	11
3.2.6.	PCN Encoding and Alternate ECN Semantics	11
3.2.7.	PCN Encoding and Tunnels	12
3.3.	Encoding Selection Criteria	13
4.	Encoding Options	14
4.1.	Encoding Using ECN and DSCP Fields	15
4.1.1.	The Use of '01' and '10' Encoding for PCN	16
4.1.2.	The Use of '11' Encoding for PCN	16
4.1.3.	The Use of '00' Encoding for PCN	16
4.1.4.	Benefits of Using DSCP and ECN Fields	17
4.1.5.	Drawbacks of Using DSCP and ECN Fields	17
4.1.6.	Comparing DSCP and ECN Fields Encoding Options	17
4.1.7.	Concerns on Alternate Semantics for the ECN Field . . .	18
4.1.8.	Encoding Choice Considerations	20
4.2.	Encoding Using DSCP Field	21
4.2.1.	Benefits of Using DSCP Field	21
4.2.2.	Drawbacks of Using DSCP Field	22
4.2.3.	Comparing DSCP Field Encoding Options	23
5.	Encoding Recommendations	23
6.	Security Implications	24
7.	IANA Considerations	25
8.	Acknowledgements	25
Appendix A.	Encoding Using ECN Field	25
Appendix A.1.	Benefits of Using ECN Field	26
Appendix A.2.	Drawbacks of Using ECN Field	27
Appendix A.3.	Concerns on Alternate Semantics for the ECN Field .	27
Appendix A.4.	Encoding Choice Considerations	30
Appendix B.	Out-of-Band Channel as Encoding Transport	30
Appendix B.1.	Benefits of Using Out-Of-Band Channel	31
Appendix B.2.	Drawbacks of Using Out-Of-Band Channel	31
Appendix C.	Current PCN Detection, Marking and Transport Mechanisms	31
Appendix C.1.	Detection, Marking and Transport Mechanisms in CL-PHB	31
Appendix C.2.	Detection, Marking and Transport Mechanisms in Three State Marking	31

Appendix C.3.	Detection, Marking and Transport Mechanisms in Single Marking	32
Appendix C.4.	Detection, Marking and Transport Mechanisms in Load Control Marking	32
9.	Informative References	32
	Authors' Addresses	35

1. Introduction

This document provides an historical account on the examination of the different ways to encode pre-congestion notification (PCN) [[RFC5559](#)] information in IP packets for transporting the information from the PCN ingress nodes, through the PCN interior nodes, to the PCN egress nodes. Documenting the examination results to indicate the reasoning behind the approach in selection of PCN encoding in IP packets, up to the time that the encoding specifications described in [[I-D.ietf-pcn-baseline-encoding](#)] has been selected. Providing an account of reasoning and lessons learned from the different encoding choices.

The most important factor affecting the choice of encoding is the environment the encoding needs to operate in. The environment with the limited number of available bits to represent the encoding place the most stringent restrictions on the encoding choices.

This document first describes the different environments and their constraints on the encoding choices.

This document then discuss the PCN information that is required to be transported.

Lastly, this document indicates possible ways the required PCN information may be encoded and carried in the different environment with their constraints.

2. Encoding Environment

The operating environment PCN encoding may operate in are:

1. Segregated Domain using DiffServ. Within a single domain with PCN traffic segregated from all other traffic using DiffServ Per Hop Behavior (DS PHB). Using DS PHB to separate PCN traffic from all other traffic allows the use of the ECN field for PCN purpose. This is the operating environment PCN is mainly focused on. Hence the majority of this document's content is on this operating environment. Variations will also be discussed. For example, using DS PHB to separate PCN and ECN traffic, and use of ECN field to distinguish PCN/non-PCN and ECN/non-ECN traffic. Another example at one extreme is the use of variations of DSCPs to support different PCN encoding options, while not using the ECN field for PCN purposes at all.
2. Single Domain without PCN and ECN distinctions. Within a single Domain without using a PCN DS PHB. In this operating

environment, PCN and ECN traffic will be treated by the same DS PHB, and will need to handle the considerations for alternate semantics for ECN field indicated by [RFC 4774](#) [[RFC4774](#)].

3. Use of Tunnels with PCN. There have been much discussions with the use of IP and MPLS Tunnels with ECN in [[I-D.ietf-tsvwg-ecn-tunnel](#)] and [RFC 5129](#) [[RFC5129](#)]) respectively. The encodings examined by this document will take the tunneling limitations into consideration. In particular, the ECN with Tunnel [[I-D.ietf-tsvwg-ecn-tunnel](#)] proposes to make the tunneling of ECN (and hence tunneling of PCN) simpler and more efficient, unifying to use one tunneling rule for ECN (and PCN) packets for all IP-in-IP tunneling (ECN with IPsec tunneling is specified by [RFC 4301](#) [[RFC4301](#)] and ECN with MPLS tunneling is specified by [RFC 5129](#) [[RFC5129](#)]). The use of [[I-D.ietf-tsvwg-ecn-tunnel](#)] will help PCN on making the encoding choices. This approach is being taken by [[I-D.ietf-pcn-baseline-encoding](#)].
4. Multiple Domains using PCN. Currently this operating environment is out of scope for the PCN Working Group, hence this will not be discussed here. But one of the goals of the standardized PCN encoding is to use the same encoding in extended environments without modification of the standard. We believe the considerations we expressed for single domain can be extended for multiple domains when we use DS PHB to facilitate the deployment of PCN.

3. Encoding Requirements

The internal PCN encoding requirements are based on the functionality of PCN [[RFC5559](#)], and possibly how the PCN Marking Algorithms achieve the functionality. There may be external requirements depending on the environment in which PCN operates, for example co-existence with ECN as indicated by [RFC 4774](#) [[RFC4774](#)], see [Section 2](#). These are discussed secondary to the internal PCN encoding requirements because we have limited the PCN operational environment in the PCN WG's first phase charter. But we also need to take into consideration of the encoding standard should not need to be modified for PCN to work in both current charter's environment and when current charter's environment is expanded, for example, to multi-domain and end-to-end, see [Section 2](#).

3.1. Encoding States

Currently, there are a number of proposals for Pre-Congestion Detection Algorithms. The authors of the different PCN Algorithm documents have agreed to use the notion of Encoding States to

represent the information each algorithm wants to export, and hence to be carried from the interior nodes to the edge nodes for flow admission control and flow termination decisions. These Encoding States form the fundamental functional requirements for the encoding choices.

Please notice the number of "Encoding State" can be different from the number of encoding bit patterns. For example more than two "Encoding States" may be carried by two encoding bit pattern when the multiple "Encoding States" can be modulated/ multiplexed over some time domains.

For simplicity purpose, we indicate the main required encoding states for PCN capable packets:

- o Not-Marked (NM), for indication of No Pre-Congestion Indication.
- o Admission Marked (AM), for indication of Flow Admission Information.
- o Termination Marked (TM), for indication of Flow Termination Information.
- o Affected Marked (AfM), for indication of ECMP Information.
- o Not-PCN, for indication of packets that are not PCN-enabled.

A total of five main required encoding states for PCN capable packets.

There are also encoding states that may be required, depending on the environment assumptions made, these encoding states are described in the following sub sections together with their environmental considerations.

3.2. Encoding and Operating Environment

Currently the PCN Working Group Charter indicates that the operating environment being a single domain. In order to support extensibility, it is necessary to specify a consistent encoding used for the currently considered operating environments, e.g., a single PCN domain, multiple PCN domains, and PCN encoded packet reaching the IP end-point, see [Section 2](#).

In this section, we first discuss the operating environment's affect on the encoding states. We then investigate the effect of the operating environment on the encoding options.

3.2.1. PCN Capable (Non PCN Capable) Packet Encoding State

This section describes the PCN Capable packet encoding state, which is used to identify the PCN Capable from Non PCN Capable packets. This encoding allows the PCN nodes to provide the PCN treatments to only the PCN Capable packets.

The PCN Working Group considers that the PCN traffic will be identified by the DSCP codepoint it carries. But the precise meaning of this is not entirely clear. There is a question whether:

1. a DSCP is meant to only represent a scheduling behaviour and (pre-)congestion marking behaviour is an optional addition that needs to be turned on or off within each existing DSCP (as for [RFC 3168](#) [[RFC3168](#)] ECN), or
2. we redefine the meaning of the DSCP field to represent a combination of scheduling and marking behaviour.

If the first approach is used, for certain PHBs (e.g. EF [[RFC3246](#)]) PCN marking would need the congestion marking behaviour turned on by the setting of another field (e.g. the ECN field). Then there would be a need to further distinguish PCN from Not-PCN packets, both using the same DSCP. Requiring a PCN Capable/Non PCN Capable Encoding State represented by a bit pattern using bits outside of the DSCP field. Notice for this approach, we indicate Non PCN Capable bit pattern because the use of the other PCN encoding bit patterns can indicate PCN Capable. This approach may be used if we want to conserve on DSCP code points by using the same DSCP code point for both PCN Capable and Non PCN Capable packets.

In the second approach, for each scheduling behaviour needing to be combined with PCN marking, a new paired DSCP would need to be defined. Then both DSCPs would map to the same scheduling behaviour but one will and one will not receive PCN treatment. For this approach, the DSCP provides the indication of PCN Capable packets.

Hence the decision of taking approach one or approach two will indicate if Non PCN Capable Packet Encoding State will be necessary.

When the Non PCN Capable Packet Encoding State is needed, we use the encoding of Not-PCN Capable (Not-PCN) to represent this state. The use of the other required PCN Encoding States will indicate this is a PCN Capable Packet.

In this document when we discuss the encoding options of using both the DSCP field and the ECN field to represent the PCN encoding states, we assume the use of the second approach above. This allows

the separation between PCN Capable and Non PCN Capable packets be totally taken care of by the use of the DiffServ field, leaving the ECN field totally available for the other required PCN encoding states' usage.

We believe the use of the second approach is a good choice because we do not envision that PCN will be used for many different scheduling behaviours. Furthermore, the use of the second approach allow PCN to use the natural ability of DiffServ for PCN packets be handled separately, with the PCN domain viewed as a separate forwarding domain within routers that can handle multiple forwarding behaviors. This use of DiffServ also ease the adaptation of multi-domain and end-to-end PCN in the future, using inter-domain DiffServ agreements.

The proposed new DSCP for capacity-admitted traffic [[I-D.ietf-tsvwg-admitted-realtime-dscp](#)] seems like it could turn on PCN marking with EF scheduling (approach 2). In this document an early version of PCN is given as an example of schemes that might need to use the new voice-admit codepoint. The proposed new DSCP for capacity-admitted traffic [[I-D.ietf-tsvwg-admitted-realtime-dscp](#)] may be used with PCN for the combined use of DSCP and PCN scenario.

3.2.2. Nonce Encoding State

The ECN nonce [RFC 3540](#) [[RFC3540](#)] for end-to-end ECN is used to protect the sender from cheating by the receiver and/or by other down stream nodes. PCN may or may not need a mechanism like the ECN nonce. However single bit nonce schemes such as the ECN nonce require in-order, reliable data delivery to function correctly. As PCN operates at the IP layer, in-order delivery cannot be guaranteed. If PCN needs a nonce functionality, it may need to think beyond the current ECN nonce mechanism. And this is beyond the scope of this document.

Currently, the PCN work we are doing assumes the trust relationships between all the functional entities are already established. If this assumption is not true, then the trust relationships will need to be addressed, but may or may not involve the need of additional encoding states or the use of the ECN nonce mechanism.

3.2.3. Non-PCN Traffic Entering PCN Domain

One of the operating environmental concerns is the accidental handling of Non PCN packets by PCN nodes. The Non PCN packets may be:

- o Non ECN capable packets.

- o ECN capable packets.

With concerns on the impacts of such non PCN packets on:

- o the processing of PCN packets.
- o the result of PCN processing on the non PCN packets.

We first look at the impacts of PCN processing on the non PCN packet with original ECN bits:

- o '00': This indicates the original packet is non ECN capable. The best action is to drop this packet when congestion is experienced. The changing of '00' to any other bit patterns will turn such packet into an ECN capable packet for any down stream nodes.
- o '01': This indicates the original packet is ECN capable. For ECN, the only valid change is to change this packet to '11' when the offered load needs to be reduced. Changing this packet to any other bit pattern may affect down stream ECN nodes.
- o '10': This indicates the original packet is ECN capable. This packet have the same concerns as the '01' packets.
- o '11': This indicates the original packet is ECN capable. This encoding is used to indicate congestion and there is trust for the sender to reduce the sending rate when the '11' encoding is received by ECN end points.

With the assumption of using DSCP to separate PCN capable and non PCN capable packets, we have to realize the non PCN packets that are receiving PCN processing somehow are using the PCN DSCP. It may be beneficial to assume the responsibility of what packets are allowed to use the PCN DSCP inside the PCN domain rests with the PCN ingress node. Making such assumption will also allow the support of multiple PCN domains and eventually the support of end to end PCN. With the PCN encoding choice and PCN processing being friendly to non PCN packets, inside the PCN domain a second line of defense has to be supported, after the use of the correct PCN DSCP.

We defer the investigation of the impact of non PCN packets on PCN processing to the sections that describe the encoding choices.

3.2.4. PCN Traffic Leaving PCN Domain

There may be two kinds of packets leaving the PCN domain unintentionally, valid PCN packets and non PCN packets that received PCN processing. Non PCN packets that did not receive PCN treatment

are considered never entered the PCN domain.

The first line of defense is still the use of the PCN DSCP, only packets using the PCN DSCP will receive PCN treatment. Hence any PCN packets leaving the PCN domain will have the PCN DSCP. Since the PCN DSCP is unique, the only danger is for down stream domains to remark the PCN DSCP to the best effort DSCP and the PCN packets being treat as ECN packets.

3.2.5. PCN Encoding for Both Edge to Edge and End to End Deployment

It is the goal of the PCN Working Group to define a standard for PCN encoding to allow the encoding be used first in the edge to edge and then in the multi-domain and end to end deployment scenarios without the need to change the standard. This section explores this environmental consideration by indicating the requirement this consideration will place on the PCN encoding selection.

3.2.6. PCN Encoding and Alternate ECN Semantics

In order to emphasize the PCN encoding and alternate ECN semantics the following effects are discussed:

- o ECN packets leaked into the PCN domain and processed by PCN interior node.
- o PCN packets leaked into the PCN domain and processed by PCN interior node.
- o ECN packets leaked out of the PCN domain into the ECN domain and processed by ECN router and ECN end-points.
- o PCN packets leaked out of the PCN domain into the ECN domain and processed by ECN router and ECN end-points.

[RFC4774] have also required one to give consideration to what harm might be caused by the leaking of PCN traffic into a non-PCN domain. The following discussion focuses on each ECN codepoint and shows what harm, if any, would be caused when that codepoint leaks out from the PCN domain:

- o '00': The leak should be safe in all circumstances. [RFC 3168](#) compliant routers will believe such packets to be not-ECN capable and as such will drop them if the router is congested. This codepoint may be suitable for different use by PCN.
- o '01' and '10': [RFC 3168](#) compliant routers will believe these packets are from an ECN capable flow. If the routers are

congested they will mark these packets '11' (CE) instead of dropping them. If the endpoints are not ECN capable then this is not good for congestion control. The use of '01' and '10' by PCN can be a potential issue. To be completely safe, it would be best to avoid giving any PCN semantics to these codepoints.

- o '11': If the packet was already part of an ECN capable flow then receivers will believe this was an indication of congestion on the path. They will thus inform their source of this and the source will perform a congestion response. This codepoint may be suitable for different use by PCN, the degree of suitability may depend on the exact PCN encoding and the metering and marking algorithm using the encoding.

More detailed consideration of these points are provided in the sections describing the encoding options.

3.2.7. PCN Encoding and Tunnels

Tunneling techniques are used in many parts of today's Internet. It is considered tunneling could start and end at any location, e.g. outside the PCN domain, at the edges or inside the PCN domain. The use of tunneling techniques post additional restrictions in PCN encoding choices. There has been much work in considering the use of ECN field with tunnels, as indicated by [[I-D.ietf-tsvwg-ecn-tunnel](#)].

The additional restriction placed on the ECN field by tunneling rules can be summarized as follows:

1. [RFC 3168](#) [[RFC3168](#)] indicates the ECN bits of the outer IP header may be set to Not-ECT at tunnel ingress. And at tunnel egress, the ECN bits of the inner IP header are unchanged by the tunnel mechanism. This will apply packet dropping and not packet marking when congestion is encountered during tunnel.
2. [RFC 3168](#) [[RFC3168](#)] also indicates only the Not-ECT and ECT code point of the ECN bits may be copied from the inner IP header to the outer IP header at tunnel ingress and only the CE code point may be copied back to the inner IP header at tunnel egress. This limits the use of only the CE code point during tunnel.
3. [RFC 4301](#) [[RFC4301](#)] allows the complete ECN field be copied from the inner IP header to the outer header at ingress but only allow the CE code point be copied from outer IP header to the inner IP header at egress. Less restrictive than [RFC 3168](#) [[RFC3168](#)] but still limits the use of the ECN field by PCN functionalities.

4. Requires the use of '00' code point for represent Not-ECT for ECN (Not-PCN) for PCN. When transferring ECN code points between the tunnel's outside header and inside header, have strick ordering of code point transistions no change from '00' (Not-ECT); '10' (ECT(0)) to '11' (CE); '01' (ECT(1)) to '11' (CE); no change from '11' (CE). These code point transfer restrictions prohibit the use of '00' encoding for PCN specific marking purpose, and restrict the use of '11' code point for representing the most server pre-congestion marking.

These additional restrictions limit the encoding choices for PCN. [I-D.ietf-tsvwg-ecn-tunnel] proposes to make the tunneling of ECN (and hence tunneling of PCN) simpler and more efficient, unifying to use one tunneling rule for ECN (and PCN) packets for all IP-in-IP tunneling (ECN with IPsec tunneling is specified by [RFC 4301](#) [RFC4301] and ECN with MPLS tunneling is specified by [RFC 5129](#) [RFC5129]). The use of [I-D.ietf-tsvwg-ecn-tunnel] will help PCN on making the encoding choices. This approach is being taken by [I-D.ietf-pcn-baseline-encoding].

3.3. Encoding Selection Criteria

Two possible locations within the IP header have been identified as suitable for encoding PCN. These are the 2 bit ECN field whose default meaning is defined in [RFC 3168](#) [RFC3168] and the 6 bit DSCP field defined in [RFC 2474](#) [RFC2474] and [RFC 2475](#) [RFC2475]. It is already accepted that PCN traffic will be distinguished according to which DSCP codepoint it carries. The implications of this decision were discussed in [section 2.2.1](#) above. The current assumption is that PCN will need to be specified as the marking behaviour through definition of a new PCN DSCP.

There are a number of other potential issues that might affect the exact choice of encoding to be used. The key ones are:

1. The support of the required encoding states to satisfy the functional requirement of PCN. These required encoding states may need two, or three, or four encoding code points to represent.
2. Compliance with [RFC 4774](#) [RFC4774] if the ECN field is to be re-used for PCN encoding.
3. Compliance with the requirements for specifying DSCPs and DSCP per-hop-behaviour groups [RFC2474].
4. Any PCN marking has to carry the '11' codepoint in the ECN field since this is the only codepoint that is guaranteed to be copied

down into the tunneling inner header upon decapsulation. This criterion is related to the constraints that any PCN encoding needs to survive being tunnelled through either an IP in IP tunnel or an IPsec Tunnel, see [Section 3.2.7](#).

5. Co-existence of PCN and not-PCN traffic: It is important to note that the scarcity of pool 1 DSCPs coupled with the fact that PCN is envisaged as a marking behaviour that could be applied to a number of different DSCPs makes it essential that we provide a not-PCN state. Because PCN re-defines the meaning of the ECN field for such DSCPs it is important to allow an operator to still use the DSCP for traffic that isn't PCN-enabled. This is achieved by providing a not-PCN state within the encoding scheme.

Each of these are examined in further details in encoding option sections describing their usage.

With the above discussion, in addition to the criteria indicated so far, we should give higher preference to encoding options that:

- o Minimize problems if there are packet leakage by the PCN domain.
- o Is safest for wider deployment of PCN, when the current chartered environment restriction is relaxed.

4. Encoding Options

There are couple of methods to carry the encoding states. The method used affects the encoding options. Hence when we describe the different encoding options in this section, we group them based on how the encoding states are carried.

The encoding transport methods considered are:

- o using the combination of the ECN and DSCP bits of a data packet header
- o using only the ECN bits of a data packet header
- o using only the DSCP bits of a data packet header

We discuss the encoding options for each of the encoding transport methods separately in their own subsections. For shorter reading, we have moved the encoding choices the working group have agreed to not consider (Using only ECN field, Out-of-Band Channel) sections to the Appendix.

4.1. Encoding Using ECN and DSCP Fields

The use of both DSCP and ECN fields is following the second approach indicated in [section 3.2.1](#). This approach allows a clean traffic treatment separation of PCN Capable traffic and Non PCN Capable traffic. This natural use of the DSCP field, to provide treatment differentiation of packets using different DSCP encoding, is one way of providing the "PCN Capable Packet" encoding state. The using of this approach allows us to focus on encoding the four required PCN Encoding States, as indicated in [section 3.1](#), using the two ECN bits.

ECN Bits	00	10	01	11	DSCP
RFC 3168	Not-ECT	ECT(0)	ECT(1)	CE	NA
Option 1	AM	NM	NM	TM	PCN-1
Option 2	AfM	NM	NM	AM/TM	PCN-1
Option 3	NM	NA	NA	AM/TM	PCN-1
Option 4	Not-PCN	NM	EXP	AM/TM	PCN-1
Option 5	NM	NA	NA	AM	PCN-1
				TM	PCN-2
Option 6	AM	NM	TM	NA	PCN-1

Notes: NA means Not Applicable. PCN-1, PCN-2 under the DSCP column denotes specific DSCPs used to indicate PCN capable packets. AM/TM means the two encoding states are sharing the same encoding bit pattern. NM means Not-Marked to represent Not Pre-Congested. Not-PCN means that packets are not PCN enabled.

Figure 1: Encoding of PCN Information Using DSCP and ECN Fields

In Figure 1, we listed the fundamental options when both DSCP and ECN fields are used. In Option 4 the ECN codepoints '01' and '10' could both be used for NM encoding or one of them could be used NM encoding and the other for experimental encoding. There are couple of variations of the theme provided by these options. One way of comparing these options is by examining the pros and cons of the different ways the four code points provided by the two ECN bits are used. We group these discussions in the following way:

1. The '01' and '10' code points.
2. The '11' code point.
3. The '00' code point.

We discuss each of them in the following sub-sections.

4.1.1. The Use of '01' and '10' Encoding for PCN

There can be different degrees of usage of the '01' and '10' code points by PCN:

1. PCN Does NOT use the '01' and '10' code points, see Option 3 and Option 5. This will be the safest choice. But this choice will leave us with only two usable code points, unless we want to deploy more than one PCN DSCPs. Even when the PCN domain does not use these code points, the PCN domain still have to handle the receiving of '01' and '10' packets at ingress. The notion of safe comes in two flavors, first if there is any packets in the PCN domain having the '01' or '10' encoding, it is immediately known that these are packets in error, either they are leaked into the PCN domain in error or are set to '01' or '10' in error inside the PCN domain. In both cases, action can be taken. The second flavor of safe is if a legitimate PCN packet leaks out of the PCN domain, it will not have the '01' or '10' encoding and should not cause an ECN router to mistaken the PCN packets to be ECN packets.
2. PCN uses '01' and '10' code points in an ECN friendly manner, see Options 1, 2, 4, and 6. One ECN friendly manner is to have both '01' and '10' to mean "PCN Capable Packet". The determination of ECN friendliness depends on the use of code points beside '01' and '10'. Furthermore, the use of '01' and '10' codepoints allow the transport of the Not-PCN encoding, see Option 4.

4.1.2. The Use of '11' Encoding for PCN

Not using the '11' code point for PCN will be a safe choice from the ECN semantic point of view, see Option 6. However, this will reduce the possible number of encoding codepoints to three. The encoding codepoint '11' is used in Options 1, 2, 3, 4, 5.

4.1.3. The Use of '00' Encoding for PCN

The '00' codepoint are used by ECN to indicate Not ECN enabled. A safe use of '00' codepoint by PCN will be to indicate Not PCN enabled, as in Option 4. The other usage may have problem in some of the

environments.

4.1.4. Benefits of Using DSCP and ECN Fields

A major feature of using both DSCP and ECN fields is the ability to use the inherent nature of DiffServ for traffic class separation to allow PCN treatment be applied to PCN traffic, without concerns of applying PCN treatment to none PCN traffic and vise versa. This feature frees this approach for PCN encoding from some of the concerns raised by [RFC 4774](#) [[RFC4774](#)]. This feature will also keep none PCN Capable traffic out of the PCN treatment mechanisms, allowing the PCN treatment mechanisms focus on their respective PCN tasks.

This approach also leaves the ECN field available totally for PCN encoding states purposes. Removing the need to carry the Not-PCN Encoding in the ECN field.

4.1.5. Drawbacks of Using DSCP and ECN Fields

The use of both DSCP and ECN fields will require the setting aside of one (or possibly two) DSCP for use by PCN. This may add complexity to the PCN encoding standardization effort.

4.1.6. Comparing DSCP and ECN Fields Encoding Options

Here we discuss the differences between the different encoding options when both DSCP and ECN fields are used. There are many encoding options, we have provided the ones we think are favorable in Figure 1.

When DSCP is used to differentiate between PCN capable and Not-PCN capable traffic, the encoding of "Not-PCN" in the ECN field is not required. This is the motivation for Option 1 in Figure 1, where the encoding "00" for "Not-ECT" is being used for "AM" (Admission Marking) encoding state. The encodings "01" and "10" for "ECT(1)" and "ECT(0)" supports the required encoding states for "Not Pre-Congested Marking" (PCN), and reserving them for any "Nonce Marking" if necessary. With the possible additional encoding of "PCN(A)" and "PCN(T)" in place of "ECT(1)" and "ECT(0)" for indicating percentage of Admission Marked traffic and percentage of Termination Marked traffic when the algorithm benefits from such additional information.

Option 2 in Figure 1 uses the "00" encoding for "AfM". With '01' and '10' encoding the same as for Option 1, requiring the use of "11" encoding for both "AM" (Admission Mark) and "TM" (Termination Mark) states or requiring the allocation of a DSCP for encoding the "TM" state.

Option 4 is the only option that can fulfill both criteria 4 and 5, listed in [Section 3.3](#).

[4.1.7](#). Concerns on Alternate Semantics for the ECN Field

[Section 2 of RFC 4774](#) [[RFC4774](#)] raised couple of concerns for usage of alternate semantics for the ECN field. We try to address each of the concerns in this section.

1. [Section 3.1 of RFC 4774](#) [[RFC4774](#)] discusses Concern 1: "How routers know which ECN semantics to use with which packets." This use of DSCP and ECN for encoding PCN states address this by following the recommendation of [RFC 4774](#) [[RFC4774](#)] on using a diffserv codepoint to identify the packets using the alternate ECN semantics. This diffserv codepoint may possibly be a new diffserv codepoint to minimize the possible confusion between using the old per hop behavior of the codepoint and the using of the alternate ECN semantics per hop behavior of the codepoint.
2. [Section 4 of RFC 4774](#) [[RFC4774](#)] discusses Concern 2: "How does the possible presence of old routers affect the performance of the alternate ECN connections." With the notion of old routers meaning routers that performs [RFC 3168](#) ECN processing instead of PCN processing. An answer to this question is given by assuming that the environment using the alternate ECN semantics is envisioned to be within a single administrative domain, see [Section 2](#), and it has the ability to ensure that all routers along the path understand and agree to the use of the alternate ECN semantics for the traffic identified by the use of a diffserv codepoint. This uses option 2 indicated in section 4.2 of [RFC 4774](#) [[RFC4774](#)]. But incase there is a mis-configuration, the choice of encoding may make a difference:

* With encoding Option 1, the old routers will interpret:

- + '00' encoding as Not-ECT, and will drop AM marked packets. The PCN edge nodes should not admit traffic that it does not receive, hence the PCN admission functionality should be OK.
- + '01' encoding as ECT(1), which indicates ECN capable and can be remarked to '11' to indicate congestion experienced. The [RFC 3168](#) ECN CE encoding have the same functionality as the PCN TM encoding, to reduce the offered traffic load. Hence the PCN termination functionality should be OK.
- + '10' encoding as ECT(0). The discussion for '01' above applies equally to this encoding.

- + '11' encoding as CE. The old router should use this encoding to reduce the offered traffic load and should not remark this to any other ECN encoding, the same functionality the PCN TM encoding requires, hence should be OK for PCN.

The above discussion for Option 1 applies equally for PCN traffic leaked out of the PCN domain and interpreted by [RFC 3168](#) ECN nodes.

* With encoding Option 2, the old routers will interpret:

- + '00' encoding as Not-ECT, and will drop AfM marked packets. This may possibly affect the efficiency of the Affected Marking functionality.
- + '01' encoding as ECT(1), which indicates ECN capable and can be remarked to '11' to indicate congestion experienced. The [RFC 3168](#) ECN CE encoding have the same functionality as the PCN TM encoding, to reduce the offered traffic load. Depending on the PCN algorithm on how AM and TM share the same '11' encoding, this may or may not affect the functionality of PCN.
- + '10' encoding as ECT(0). The discussion for '01' above applies equally to this encoding.
- + '11' encoding as CE. The old router should use this encoding to reduce the offered traffic load and should not remark this to any other ECN encoding. Depending on the PCN algorithm on how AM and TM share the same '11' encoding, this may or may not affect the functionality of PCN.

The above discussion for Option 2 applies equally for PCN traffic leaked out of the PCN domain and interpreted by [RFC 3168](#) ECN nodes.

3. Concern 3: "How does the possible presence of old routers affect the coexistence of the alternate ECN traffic with competing traffic on the path." Within the PCN domain, the PCN (alternate ECN) traffic is separated from the other traffic using diffserv. If by mis-configuration, an old routers that does not understand PCN handles PCN traffic, the PCN traffic will get the per hop behavior as the other traffic, hence not receiving the benefits of PCN at the old router, but will not affect the coexistence of the PCN and the other traffic. If the old router uses [RFC 3168](#) ECN congestion treatment, then the discussion for Concern 2 above

applies.

4. Concern 4: "How well does the alternate ECN traffic perform."
The performance of the different proposed PCN (alternate ECN) metering and marking algorithms are currently under study with their simulation and study results described by their respective documents.

The environment using the alternate ECN semantics is envisioned to be within a single administrative domain. With the ability to ensure that all routers along the path understand and agree to the use of the alternate ECN semantics for the traffic identified by the use of a Diffserv codepoint. This uses option 2 indicated in [section 4.2 of RFC 4774](#) [RFC4774].

4.1.8. Encoding Choice Considerations

- o If three encoding states need to be separately represented, Option 1 is recommended.
- o If two encoding states need to be separately represented, for example the marking algorithm allows the AM and TM encoding states be represented using the same bit pattern, Options 2 and 3 are recommended. If the Not-PCN encoding state is required then Option 4 is recommended.
- o If [RFC4774] concerns need to be addressed by PCN encoding, then Option 1 is recommended, please see [section 3.1.4](#) for the detail discussion. Options 2 and 3 may be able to address the [RFC 4774](#) [RFC4774] concerns, but a heavier burden is placed on the metering and marking algorithms to differentiate between TM and AM meaning of the '11' encoding when a [RFC 3168](#) ECN router sets the '11' encoding.
- o If the metering and marking algorithm requires the use of Affected Marking encoding state, Option 2 is recommended. Alternatively one of the bit patterns of '01' or '10' may be used for the AfM purpose. But using '01' or '10' bit patterns for AfM may increase the interference between [RFC 3168](#) ECN and PCN encodings, please see [section 3.1.4](#) for the detail discussion.
- o If Option 1 is used and the functionality of Affected Marking encoding state is required, the metering and marking algorithms will need to provide this functionality either without the use of the Affected Marking encoding state, or using an additional DSCP to encode the Affected Marking encoding state.

- o If Option 4 is used and the functionality of Affected Marking encoding state is required, the metering and marking algorithms will need to provide this functionality using one of the NA encoding states associated with the '01' or '10' bit patterns.

4.2. Encoding Using DSCP Field

In this type of encoding and transport method the congestion and pre-congestion information is encoded into the 6 DSCP bits that are transported in the IP header of the data packets. Four possible alternatives can be distinguished, as can be seen in Figure 2, with details provided by [draft-westberg-pcn-load-control-02.txt](#) [[I-D.westberg-pcn-load-control](#)]. Option 7 needs 2 additional DSCP values, Options 8 and 9 need three additional DSCP values and Option 10 needs four additional DSCP values. Note that all additional and experimental DSCP values are representing and are associated with the same PHB. The 1st, 2nd, 3rd, and 4th DSCP values are representing DSCP values that are assigned by IANA as DSCP experimental values, see [RFC 2211](#) [[RFC2211](#)]. Furthermore, all options listed in Figure 2 are able to support the Not-PCN encoding state.

DSCP Bits	Original	Add DSCP 1	Add DSCP 2	Add DSCP 3	Add DSCP 4
Option 7	Not-PCN	UM	AM/TM	NA	NA
Option 8	Not-PCN	UM	AM/TM	AfM	NA
Option 9	Not-PCN	UM	AM	TM	NA
Option 10	Not-PCN	UM	AM	TM	AfM

Notes: Not-PCN means the packet is not PCN capable. UM for Un-Marked meaning Not Pre-Congested

Figure 2: Encoding of PCN Information Using DSCP Field

4.2.1. Benefits of Using DSCP Field

The main benefits of using the DSCP field for PCN encoding are:

- o it is not affecting the end-to-end ECN semantics and therefore the issues and concerns raised in [RFC 4774](#) [[RFC4774](#)] are not applicable for this encoding scheme.

- o it is not affected by the PCN tunneling issues discussed in [Section 3.2.7](#).
- o all 4 DSCP encoding options depicted in Figure 2 can support the PCN capable not congested/UnMarked (UM) indication, the admission control (AM) and flow termination (TM) encoding states.
- o the experimental DSCPs are lightly standardized and therefore, the rules on how to apply and use them are limited. This provides a high flexibility to network operators to apply and use them in different settings.
- o simple packet classification, since a router needs only to read the DSCP field, instead of reading both DSCP and ECN fields.
- o Option 8 and 10 support the Affected Marking (AfM) encoding, which according to [[I-D.westberg-pcn-load-control](#)], it has benefits if the PCN-domain operates ECMP routing and is not using DSCP for route selection.
- o by using an additional DSCP to encode the not congested PCN state, all PCN-ingress-nodes can be configured to encode this state into all packets that are entering the PCN domain and are PCN aware. This will solve any PCN-egress-node misconfiguration problems, which can allow a AM/TM or SM encoded packet to outgo a PCN-domain.

[4.2.2](#). Drawbacks of Using DSCP Field

The main drawbacks of using the DSCP field for PCN encoding are the following:

this type of encoding needs to use per PHB, in addition to the original DSCP and depending on the encoding option used, one, two, three, or four DSCP values, respectively. These additional DSCP values can be taken from the DSCP values that are not defined by standards action, see [RFC 2211](#) [[RFC2211](#)]. Note that all the additional DSCP values are representing and are associated with one PHB. The value of this DSCP/PHB can either follow a standards action or use a value that is applied for experimental or local use. It is important to note that the number of the DSCP values used for local or experimental use is restricted and therefore the number of different PHBs supported in the PCN domain will also be restricted.

applying the DSCP field as PCN encoding transport within an PCN aware MPLS domain, see [RFC 5129](#) [[RFC5129](#)], can be problematic due to the scarce packet header real-estate.

when the PCN-domain is operating ECMP that uses DSCP to select the routes, a risk of mis-ordering of packets within a flow might occur. The impact of this drawback depends on the following:

1. the level of deployment of ECMP algorithms that use DSCP for route selection;
2. mis-ordering of packets within a flow when there is termination marking may be acceptable;
3. the possibility of configuring the ECMP algorithms that use DSCP for route selection in the PCN-domain that the used PCN aware DSCPs are belonging to the same PHB and therefore, all these DSCP values should be converted to one preconfigured DSCP value before applying it in the ECMP routing algorithm. Note that all the additional experimental DSCPs that are used within PCN are belonging to the same PHB.

4.2.3. Comparing DSCP Field Encoding Options

Option 7 can support the basic encoding states, i.e., not PCN, not congested (UM), and the AM/TM encoding states. Option 8 can support the basic encoding states supported by Option 7, but in addition it can support the AfM state. Option 9 can support the following basic encoding states: not PCN, not congested (UM), AM and TM states. Option 10 can support the states supported by Option 9, but in addition it can support the AfM state. Furthermore, in options 7 and 8 the encoding sequence associated with Admission Control and Flow Termination is independent of each other. In options 9 and 10 a packet cannot be AM encoded if it has been earlier TM encoded. All options can support the Not-PCN encoding state.

5. Encoding Recommendations

This memo describes the selection process of the PCN encoding in IP packets, up to the time that the encoding specifications described in the [[I-D.ietf-pcn-baseline-encoding](#)] have been selected. The encoding specification given in [[I-D.ietf-pcn-baseline-encoding](#)] needs to specify how the AM and TM encoding states are represented using the same bit pattern. Furthermore, it is considered that the Not-PCN encoding state is required.

Based on the above listed [[I-D.ietf-pcn-baseline-encoding](#)] considerations and on the provided discussions in Sections [4.1](#) and [4.2](#), it can be deduced that only Option 4 fulfills the above listed [[I-D.ietf-pcn-baseline-encoding](#)] considerations and the criteria listed in [Section 3.3](#). Therefore, it is recommended that

[I-D.ietf-pcn-baseline-encoding] should use Option 4 to represent the 'Baseline Encoding and Transport of Pre-Congestion Information'.

6. Security Implications

Packets from normal precedence and higher precedence sessions [[ITU-MLPP](#)] aren't distinguishable by PCN Interior Nodes. This prevents an attacker specifically targeting, in the data plane, higher precedence packets (perhaps for DoS or for eavesdropping). However, PCN End Nodes can access this information to help decide whether to admit or terminate a flow. The separation of network information provided by the Interior Nodes and the precedence information at the PCN End Nodes allows simpler, easier and better focused security enforcement.

PCN End Nodes police packets to ensure a flow sticks within its agreed limit. This is similar to the existing IntServ behaviour. Between them the PCN End Nodes must fully encircle the PCN-Region, otherwise packets could enter the PCN-Region without being subject to admission control, which would potentially destroy the QoS of existing flows.

It is assumed that all the Interior Nodes and PCN End Nodes run PCN and trust each other (ie the PCN-enabled Internet Region is a controlled environment). For instance a non-PCN router wouldn't be able to alert that it's suffering pre-congestion, which potentially would lead to too many calls being admitted (or too few being terminated). Worse, a rogue router could perform attacks such as marking all packets so that no flows were admitted.

So security requirements are focussed at specific parts of the PCN-Region:

The PCN End Nodes become the trust points. The degree of trust required depends on the kinds of decisions it has to make and the kinds of information it needs to make them. For example when the PCN End Node needs to know the contents of the sessions for making the decisions, when the contents are highly classified, the security requirements for the PCN End Nodes involved will also need to be high.

PCN-marking by the Interior Nodes along the packet forwarding path needs to be trusted, because the PCN End Nodes rely on this information.

7. IANA Considerations

This memo includes no request to IANA.

8. Acknowledgements

We would like to acknowledge the members of the PCN working group for the discussions that generated the contents of this memo.

Appendix A. Encoding Using ECN Field

This section takes the approach 1 option indicated in [section 2.1.1](#). Which the DSCP field only indicates the packet forwarding behavior, for which both PCN Capable and Non PCN Capable traffic use/share the same DSCP. This approach requires the use of the Not PCN Capable Encoding State to be encoding using the ECN bits. Hence this section describes the encoding options that uses only the ECN field (without the DSCP field) available in the IP header of the data packets to encode the PCN states.

The use of the same DSCP for both PCN Capable and Non PCN Capable also opens the question of having PCN and [RFC 3168](#) ECN traffic using the same DSCP. Which increases the importance of satisfying the concerns indicated in [RFC 4774](#).

ECN Bits	00	01	10	11	DSCP
RFC 3168	Not-ECT	ECT(1)	ECT(0)	CE	NA
Option 11	Not-PCN	AM	PCN	TM	NA
Option 12	Not-PCN	PCN	PCN	AM/TM	NA
Option 13	Not-PCN	AfM	PCN	AM/TM	NA

Figure 3: Encoding of PCN Information Using ECN Field

In Figure 2, we listed the fundamental options when only the ECN field is used. Like in Figure 1, there are variations of the theme provided by these options. For example, when both "01" and "10" encoding are used for NPM in Option 5, they can be interpreted as PCN(A) and PCN(T) instead of just PCN. Using the PCN(A) and PCN(T) variation provides the additional information of the ratio of packets

AM marked to packets Not AM marked, and the ratio of packets TM marked to packets Not TM marked. Having these ratios being independent from one another.

For Option 11, the use of '01' for AM and '10' for PCN can be swapped and provide the same functionality. For Option 13, the use of '01' for AfM and '10' for PCN can also be swapped without change of functionality.

[Appendix A.1](#). Benefits of Using ECN Field

The using of only the ECN field for encoding PCN encoding states allow more efficient use of the DSCP field, not requiring the allocation of PCN specific DSCP values.

This approach also opens the question of possibly having both PCN and ECN traffic using the same DSCP.

When the same treatment can be provided to both ECN and PCN traffic to achieve each of ECN and PCN purpose, then not having DiffServ as separation between ECN and PCN traffic may be a benefit. Under such circumstances, having the same encoding between ECN and PCN may be desirable. But this can only be true if the requirement set forth in [RFC 4774](#) [[RFC4774](#)] for alternate ECN semantics can be satisfied.

If the same treatment can be applied to both ECN and PCN traffic, then:

- o The first issue of [RFC 4774](#) [[RFC4774](#)]: "How routers know which ECN semantics to use with which packets." may be solved because there are no difference in the treatments of ECN and PCN packets, hence they can use the same semantics.
- o The second and third issues of [RFC 4774](#) [[RFC4774](#)]: "How does the possible presence of old routers affect the performance of the alternate ECN connections." and "How does the possible presence of old routers affect the coexistence of the alternate ECN traffic with competing traffic on the path." are also solved because there are no difference in the treatment of ECN and PCN packets.
- o The forth issue of [RFC 4774](#) [[RFC4774](#)]: "How well does the alternate ECN traffic perform." are dependent on the algorithm used, and should be provided by the respective algorithm document, and not in the scope of this document.

Appendix A.2. Drawbacks of Using ECN Field

Notice this group of encoding options does not use DiffServ code points for PCN encoding. With this group of encoding options, the required states of "PCN Capable Transport"/"None PCN Capable Transport" must be encoded using the ECN field. Leaving less encoding real estate to carry the remaining required PCN encoding states. Another drawback is without the protection/separation capability provided by DiffServ, it is typically harder to satisfy the requirement set forth in [RFC 4774](#) [[RFC4774](#)] for alternate ECN semantics.

Appendix A.3. Concerns on Alternate Semantics for the ECN Field

[Section 2 of RFC 4774](#) [[RFC4774](#)] raised couple of concerns for usage of alternate semantics for the ECN field. We try to address each of the concerns in this section.

1. [Section 3.1 of RFC 4774](#) [[RFC4774](#)] discusses Concern 1: "How routers know which ECN semantics to use with which packets." When this group of PCN encodings are used without the use of DSCP, routers can not distinguished PCN encoded packets from [RFC 3168](#) ECN encoded packets. Hence there needs to be some kind of differentiation between PCN and [RFC 3168](#) ECN packets, may be using PCN for real-time traffic types (with specific DSCP) and ECN for elastic traffic (with specific DSCP). And only distinguishing PCN Capable and Non-PCN Capable packets in real-time traffic. Only distinguishing ECT and Not-ECT packets in elastic traffic. But not having PCN and ECN traffic together.
2. [Section 4 of RFC 4774](#) [[RFC4774](#)] discusses Concern 2: "How does the possible presence of old routers affect the performance of the alternate ECN connections." With the notion of old routers meaning routers that performs [RFC 3168](#) ECN processing instead of PCN processing, or drop packets instead of encoding the congestion information. The easy answer is the environment using the alternate ECN semantics is envisioned to be within a single administrative domain. With the ability to ensure that all routers along the path understand and agree to the use of the alternate ECN semantics for the traffic identified to be PCN Capable. This uses option 2 indicated in [section 4.2 of RFC 4774](#) [[RFC4774](#)]. But incase there is mis-configuration, the choice of encoding may make a difference:

* With encoding Option 11, the old routers will interpret:

- + '00' encoding as Not-ECT, and will drop Not-PCN marked packets when congestion is detected. With '00' the

encoding for Not-PCN, requiring the same functionality as Not-ECT, the presence of old routers will not affect the performance of PCN functionality.

- + '01' encoding as ECT(1), which indicates ECN capable and can be remarked to '11' to indicate congestion experienced. For Option 3, the old router can possibly remark AM to TM. This puts a burden on the metering and marking algorithms to treat TM encoded packets to indicate stop admission. This may or may not be acceptable, depending on the algorithm.
- + '10' encoding as ECT(0), which indicates ECN capable and can be remarked to '11' to indicate congestion experienced. The [RFC 3168](#) ECN CE encoding have the same functionality as the PCN TM encoding, to reduce the offered traffic load. Hence the PCN termination functionality should be OK.
- + '11' encoding as CE. The old router should use this encoding to reduce the offered traffic load and should not remark this to any other ECN encoding, the same functionality the PCN TM encoding requires, hence should be OK for PCN.

The above discussion for Option 11 applies equally for PCN traffic leaked out of the PCN domain and interpreted by [RFC 3168](#) ECN nodes.

- * With encoding Option 12, the old routers will interpret:
 - + '00' encoding as Not-ECT, and will drop Not-PCN marked packets when congestion is detected. With '00' the encoding for Not-PCN, requiring the same functionality as Not-ECT, the presence of old routers will not affect the performance of PCN functionality.
 - + '01' encoding as ECT(1), which indicates ECN capable and can be remarked to '11' to indicate congestion experienced. The [RFC 3168](#) ECN CE encoding have the same functionality as the PCN TM encoding, to reduce the offered traffic load. Depending on the PCN algorithm on how AM and TM share the same '11' encoding, this may or may not affect the functionality of PCN.
 - + '10' encoding as ECT(0). The discussion for '01' above applies equally to this encoding.

- + '11' encoding as CE. The old router should use this encoding to reduce the offered traffic load and should not remark this to any other ECN encoding. Depending on the PCN algorithm on how AM and TM share the same '11' encoding, this may or may not affect the functionality of PCN.

The above discussion for Option 12 applies equally for PCN traffic leaked out of the PCN domain and interpreted by [RFC 3168](#) ECN nodes.

* With encoding Option 13, the old routers will interpret:

- + '00' encoding as Not-ECT, and will drop Not-PCN marked packets when congestion is detected. With '00' the encoding for Not-PCN, requiring the same functionality as Not-ECT, the presence of old routers will not affect the performance of PCN functionality.
- + '01' encoding as ECT(1), which indicates ECN capable and can be remarked to '11' to indicate congestion experienced. For Option 6, the old router can possibly remark AfM to TM. This may or may not be acceptable, depending on the algorithm's Affected Marking functionality.
- + '10' encoding as ECT(1), which indicates ECN capable and can be remarked to '11' to indicate congestion experienced. The [RFC 3168](#) ECN CE encoding have the same functionality as the PCN TM encoding, to reduce the offered traffic load. Depending on the PCN algorithm on how AM and TM share the same '11' encoding, this may or may not affect the functionality of PCN.
- + '11' encoding as CE. The old router should use this encoding to reduce the offered traffic load and should not remark this to any other ECN encoding. Depending on the PCN algorithm on how AM and TM share the same '11' encoding, this may or may not affect the functionality of PCN.

The above discussion for Option 13 applies equally for PCN traffic leaked out of the PCN domain and interpreted by [RFC 3168](#) ECN nodes.

3. Concern 3: "How does the possible presence of old routers affect the coexistence of the alternate ECN traffic with competing traffic on the path." If [RFC 3168](#) ECN and PCN traffic are to be treated within a single DiffServ PHB, because with these encoding

there is no way to differentiate between the ECN packets from the PCN traffic, the metering and marking algorithm used must be totally friendly between ECN and PCN traffic, else they will affect each other in possibly non-acceptable ways. These encoding will work OK with traffic besides ECN because of the use of 'Not-PCN' encoding.

4. Concern 4: "How well does the alternate ECN traffic perform." The performance of the different proposed PCN (alternate ECN) metering and marking algorithms are currently under study with their simulation and study results described by their respective documents.

Appendix A.4. Encoding Choice Considerations

- o If three encoding states need to be separately represented, Option 11 is recommended.
- o If the marking algorithm allows the AM and TM encoding states be represented using the same bit pattern, Option 12 is recommended.
- o If the marking algorithm requires the use of Affected Marking encoding state, Option 13 is recommended. For Option 13, alternative NPM bit patterns ('01' or '10') may be used for the AfM purpose.

Appendix B. Out-of-Band Channel as Encoding Transport

In this type of encoding and transport method the congestion and precongestion information can be encoded using the IPFIX protocol [RFC 3955](#) [18], that is normally used to carry flow-based IP traffic measurements from an observation point to a collecting point. Note that this encoding scheme is denoted in this document as "IPFIX channel". An observation point is a location in a network where IP packets can be observed and measured. A collecting point can be a process or a node that receives flow records from one or more observation points. In the PCN case, each PCN-interior-node will be an IPFIX observation point and the PCN-egress-node will be the IPFIX collecting point.

The PCN-interior-node will support the metering process and the flow records. Note that in this case each flow record can be associated with the record of the congestion and pre-congestion metering information associated with each PHB. The PCN-egress-node will then support the IPFIX collecting process, which will receive flow records from one or more congested and pre-congested PCN-interior-nodes. Using this encoding method the encoding modes/states can be

aggregated and transported to the egress node by using the flow records at regular intervals or at the moment that a congestion and pre-congestion situation occurs. The used transport channel in this case is not the data path but a signaling protocol.

Appendix B.1. Benefits of Using Out-Of-Band Channel

This encoding scheme does not use the data path for encoding and transport, but it is able to transport the congestion and pre-congestion information associated with the encoding states by using a separate signaling channel. Another benefit of using this encoding scheme is that it is not affecting the end-to-end ECN semantics and therefore the issues and concerns raised in [RFC 4774](#) are not applicable for this encoding scheme.

Appendix B.2. Drawbacks of Using Out-Of-Band Channel

The "IPFIX channel" encoding mode needs a separate signaling channel for the transport of the congestion and pre-congestion information from the PCN-interior-nodes towards the PCN-egress-node. The requirement of using an additional channel increases the complexity and influences negatively the performance of the PCN-interior-nodes since each PCN-interior-node needs to support in addition to the data path a separate channel.

Appendix C. Current PCN Detection, Marking and Transport Mechanisms

This appendix indicates the different available PCN based mechanisms that can be used for congestion and pre-congestion detection and marking used at interior nodes. The requirements and characteristics of such algorithms may influence the encoding and transport of the PCN encoding states.

Appendix C.1. Detection, Marking and Transport Mechanisms in CL-PHB

Please see [draft-briscoe-tsvwg-cl-phb-03.txt](#) [[I-D.briscoe-tsvwg-cl-phb](#)] for details on the Controlled-Load PHB Algorithm.

Appendix C.2. Detection, Marking and Transport Mechanisms in Three State Marking

Please see [draft-babiarz-pcn-3sm-01.txt](#) [[I-D.babiarz-pcn-3sm](#)] for details on the Three State Marking Algorithm.

Appendix C.3. Detection, Marking and Transport Mechanisms in Single Marking

Please see [draft-charny-pcn-single-marking-03.txt](#) [[I-D.charny-pcn-single-marking](#)] for details on the Single Marking Algorithm.

Appendix C.4. Detection, Marking and Transport Mechanisms in Load Control Marking

Please see [draft-westberg-pcn-load-control-02.txt](#) [[I-D.westberg-pcn-load-control](#)] for details on the Load Control Algorithm.

9. Informative References

- [I-D.ietf-pcn-baseline-encoding]
Moncaster, T., Briscoe, B., and M. Menth, "Baseline Encoding and Transport of Pre-Congestion Information", [draft-ietf-pcn-baseline-encoding-07](#) (work in progress), September 2009.
- [I-D.ietf-tsvwg-ecn-tunnel]
Briscoe, B., "Tunnelling of Explicit Congestion Notification", [draft-ietf-tsvwg-ecn-tunnel-03](#) (work in progress), July 2009.
- [I-D.babiarz-pcn-3sm]
Babiarz, J., Liu, X., Chan, K., and M. Menth, "Three State PCN Marking", [draft-babiarz-pcn-3sm-01](#) (work in progress), November 2007.
- [I-D.charny-pcn-single-marking]
Charny, A., Zhang, X., Faucheur, F., and V. Liatsos, "Pre-Congestion Notification Using Single Marking for Admission and Termination", [draft-charny-pcn-single-marking-03](#) (work in progress), November 2007.
- [I-D.westberg-pcn-load-control]
Westberg, L., Bhargava, A., Bader, A., Karagiannis, G., and H. Mekkes, "LC-PCN: The Load Control PCN Solution", [draft-westberg-pcn-load-control-05](#) (work in progress), November 2008.
- [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.ietf-tsvwg-admitted-realtime-dscp]
Baker, F., Polk, J., and M. Dolly, "DSCP for Capacity-Admitted Traffic",
[draft-ietf-tsvwg-admitted-realtime-dscp-05](#) (work in progress), November 2008.
- [I-D.ietf-tsvwg-mllef-concerns]
Baker, F. and J. Polk, "MLEF Without Capacity Admission Does Not Satisfy MLPP Requirements",
[draft-ietf-tsvwg-mllef-concerns-00](#) (work in progress), February 2005.
- [RFC1633] Braden, B., Clark, D., and S. Shenker, "Integrated Services in the Internet Architecture: an Overview",
[RFC 1633](#), June 1994.
- [RFC2211] Wroclawski, J., "Specification of the Controlled-Load Network Element Service", [RFC 2211](#), September 1997.
- [RFC2309] Braden, B., Clark, D., Crowcroft, J., Davie, B., Deering, S., Estrin, D., Floyd, S., Jacobson, V., Minshall, G., Partridge, C., Peterson, L., Ramakrishnan, K., Shenker, S., Wroclawski, J., and L. Zhang, "Recommendations on Queue Management and Congestion Avoidance in the Internet", [RFC 2309](#), April 1998.
- [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.
- [RFC2597] Heinanen, J., Baker, F., Weiss, W., and J. Wroclawski, "Assured Forwarding PHB Group", [RFC 2597](#), June 1999.
- [RFC2702] Awduche, D., Malcolm, J., Agogbua, J., O'Dell, M., and J. McManus, "Requirements for Traffic Engineering Over MPLS", [RFC 2702](#), September 1999.
- [RFC2998] Bernet, Y., Ford, P., Yavatkar, R., Baker, F., Zhang, L., Speer, M., Braden, R., Davie, B., Wroclawski, J., and E. Felstaine, "A Framework for Integrated Services Operation over Diffserv Networks", [RFC 2998](#), November 2000.

- [RFC3168] Ramakrishnan, K., Floyd, S., and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP", [RFC 3168](#), September 2001.
- [RFC3246] Davie, B., Charny, A., Bennet, J., Benson, K., Le Boudec, J., Courtney, W., Davari, S., Firoiu, V., and D. Stiliadis, "An Expedited Forwarding PHB (Per-Hop Behavior)", [RFC 3246](#), March 2002.
- [RFC3247] Charny, A., Bennet, J., Benson, K., Boudec, J., Chiu, A., Courtney, W., Davari, S., Firoiu, V., Kalmanek, C., and K. Ramakrishnan, "Supplemental Information for the New Definition of the EF PHB (Expedited Forwarding Per-Hop Behavior)", [RFC 3247](#), March 2002.
- [RFC3540] Spring, N., Wetherall, D., and D. Ely, "Robust Explicit Congestion Notification (ECN) Signaling with Nonces", [RFC 3540](#), June 2003.
- [RFC3955] Leinen, S., "Evaluation of Candidate Protocols for IP Flow Information Export (IPFIX)", [RFC 3955](#), October 2004.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", [RFC 4301](#), December 2005.
- [RFC4594] Babiarz, J., Chan, K., and F. Baker, "Configuration Guidelines for DiffServ Service Classes", [RFC 4594](#), August 2006.
- [RFC4774] Floyd, S., "Specifying Alternate Semantics for the Explicit Congestion Notification (ECN) Field", [BCP 124](#), [RFC 4774](#), November 2006.
- [RFC5129] Davie, B., Briscoe, B., and J. Tay, "Explicit Congestion Marking in MPLS", [RFC 5129](#), January 2008.
- [RFC5559] Eardley, P., "Pre-Congestion Notification (PCN) Architecture", [RFC 5559](#), June 2009.
- [DClark] "Supporting Real-Time Applications in an Integrated Services Packet Network: Architecture and Mechanisms", Proceedings of SIGCOMM '92 at Baltimore MD, August 1992.
- [ITU-MLPP] "Multilevel Precedence and Pre-emption Service (MLPP)", ITU-T Recommendation I.255.3, 1990.
- [Reid] "Economics and Scalability of QoS Solutions", BT

Technology Journal Vol 23 No 2, April 2005.

Authors' Addresses

Kwok Ho Chan
Huawei Technologies
125 Nagog Park
Acton, MA 01720
USA

Email: khchan@huawei.com

Georgios Karagiannis
University of Twente
P.O. Box 217
7500 AE Enschede,
The Netherlands

Email: g.karagiannis@ewi.utwente.nl

Toby Moncaster
BT Research
B54/70, Sirius House Adastral Park Martlesham Heath
Ipswich, Suffolk IP5 3RE
United Kingdom

Email: toby.moncaster@bt.com

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

Email: menth@informatik.uni-wuerzburg.de

Philip Eardley
BT Research
B54/77, Sirius House Adastral Park Martlesham Heath
Ipswich, Suffolk IP5 3RE
United Kingdom

Email: philip.eardley@bt.com

Bob Briscoe
BT Research
B54/77, Sirius House Adastral Park Martlesham Heath
Ipswich, Suffolk IP5 3RE
United Kingdom

Email: bob.briscoe@bt.com

