

The Addition of Explicit Congestion Notification (ECN) to IP

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>

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

Abstract

This document specifies the incorporation of ECN (Explicit Congestion Notification) to TCP and IP, including ECN's use of two bits in the IP header's DS field. We begin by describing TCP's use of packet drops as an indication of congestion. Next we explain that with the addition of active queue management (e.g., RED) to the Internet infrastructure, where routers detect congestion before the queue overflows, routers are no longer limited to packet drops as an indication of congestion. Routers can instead set the Congestion Experienced (CE) bit in the IP header of packets from ECN-capable transports. We describe when the CE bit is to be set in routers, and

describe modifications needed to TCP to make it ECN-capable. Modifications to other transport protocols (e.g., unreliable unicast or multicast, reliable multicast, other reliable unicast transport protocols) could be considered as those protocols are developed and advance through the standards process.

We also describe in this document the issues involving the use of ECN within IP tunnels, and within IPsec tunnels in particular.

One of the guiding principles for this document is that all the mechanisms specified here are incrementally deployable.

Table of Contents

1. Introduction
2. Conventions and Acronyms
3. Assumptions and General Principles
4. Active Queue Management (AQM)
5. Explicit Congestion Notification in IP
 - 5.1. ECN as an indication of persistent congestion
 - 5.2. Dropped or Corrupted Packets
6. Support from the Transport Protocol
 - 6.1. TCP
 - 6.1.1. TCP Initialization
 - 6.1.1.1. Robust TCP Initialization with an Echoed Reserve Field
 - 6.1.1.2. Robust TCP Initialization with no response to the SYN
 - 6.1.2. The TCP Sender
 - 6.1.3. The TCP Receiver
 - 6.1.4. Congestion on the ACK-path
 - 6.1.5. Retransmitted TCP packets
 - 6.1.6. TCP Window Probes.
 7. Non-compliance by the End Nodes
 8. Non-compliance in the Network
 - 8.1. Complications Introduced by Split Paths
 9. Encapsulated Packets
 - 9.1. IP packets encapsulated in IP
 - 9.1.1. The limited-functionality and full-functionality options within
 - 9.1.2. Changes to the ECN Field within an IP Tunnel.
 - 9.2. IPsec Tunnels
 - 9.2.1. Negotiation between Tunnel Endpoints
 - 9.2.1.1. ECN Tunnel Security Association Database Field
 - 9.2.1.2. ECN Tunnel Security Association Attribute
 - 9.2.1.3. Changes to IPsec Tunnel Header Processing
 - 9.2.2. Changes to the ECN Field within an IPsec Tunnel.
 - 9.2.3. Comments for IPsec Support
 - 9.3. IP packets encapsulated in non-IP packet headers.
10. Issues Raised by Monitoring and Policing Devices
11. Evaluations of ECN
12. Summary of changes required in IP and TCP
13. Conclusions
14. Acknowledgements
15. References
16. Security Considerations
17. IPv4 Header Checksum Recalculation
18. Possible Changes to the ECN Field in the Network
 - 18.1. Possible Changes to the IP Header
 - 18.1.1. Erasing the Congestion Indication
 - 18.1.2. Falsely Reporting Congestion
 - 18.1.3. Disabling ECN-Capability
 - 18.1.4. Falsely Indicating ECN-Capability
 - 18.1.5. Changes with No Functional Effect
 - 18.2. Information carried in the Transport Header

- 18.3. Split Paths
- 19. Implications of Subverting End-to-End Congestion Control
 - 19.1. Implications for the Network and for Competing Flows
 - 19.2. Implications for the Subverted Flow
 - 19.3. Non-ECN-Based Methods of Subverting End-to-end Congestion Control
- 20. The motivation for the ECT bit.
- 21. Why use two bits in the IP header?
- 22. Historical definitions for the IPv4 TOS octet

1. Introduction

TCP's congestion control and avoidance algorithms are based on the notion that the network is a black-box [[Jacobson88](#), [Jacobson90](#)]. The network's state of congestion or otherwise is determined by end-systems probing for the network state, by gradually increasing the load on the network (by increasing the window of packets that are outstanding in the network) until the network becomes congested and a packet is lost. Treating the network as a "black-box" and treating loss as an indication of congestion in the network is appropriate for pure best-effort data carried by TCP, with little or no sensitivity to delay or loss of individual packets. In addition, TCP's congestion management algorithms have techniques built-in (such as Fast Retransmit and Fast Recovery) to minimize the impact of losses, from a throughput perspective. However, these mechanisms are not intended to help applications that are in fact sensitive to the delay or loss of one or more individual packets. Interactive traffic such as telnet, web-browsing, and transfer of audio and video data can be sensitive to packet losses (especially when using an unreliable data delivery transport such as UDP) or to the increased latency of the packet caused by the need to retransmit the packet after a loss (with the reliable data delivery semantics provided by TCP).

Since TCP determines the appropriate congestion window to use by gradually increasing the window size until it experiences a dropped packet, this causes the queues at the bottleneck router to build up. With most packet drop policies at the router that are not sensitive to the load placed by each individual flow (e.g., tail-drop on queue overflow), this means that some of the packets of latency-sensitive flows may be dropped. In addition, such drop policies lead to synchronization of loss across multiple flows.

Active queue management mechanisms detect congestion before the queue overflows, and provide an indication of this congestion to the end nodes. Thus, active queue management can reduce unnecessary queueing delay for all traffic sharing that queue. The advantages of active queue management are discussed in [RFC 2309](#) [[RFC2309](#)]. Active queue management avoids some of the bad properties of dropping on queue

overflow, including the undesirable synchronization of loss across multiple flows. More importantly, active queue management means that transport protocols with mechanisms for congestion control (e.g., TCP) do not have to rely on buffer overflow as the only indication of congestion.

Active queue management mechanisms may use one of several methods for indicating congestion to end-nodes. One is to use packet drops, as is currently done. However, active queue management allows the router to separate policies of queueing or dropping packets from the policies for indicating congestion. Thus, active queue management allows routers to use the Congestion Experienced (CE) bit in a packet header as an indication of congestion, instead of relying solely on packet drops. This has the potential of reducing the impact of loss on latency-sensitive flows.

2. Conventions and Acronyms

The keywords MUST, MUST NOT, REQUIRED, SHALL, SHALL NOT, SHOULD, SHOULD NOT, RECOMMENDED, MAY, and OPTIONAL, when they appear in this document, are to be interpreted as described in [B97].

3. Assumptions and General Principles

In this section, we describe some of the important design principles and assumptions that guided the design choices in this proposal.

- * Because ECN is likely to be adopted gradually, accommodating migration is essential. Some routers may still only drop packets to indicate congestion, and some end-systems may not be ECN-capable. The most viable strategy is one that accommodates incremental deployment without having to resort to "islands" of ECN-capable and non-ECN-capable environments.

- * New mechanisms for congestion control and avoidance need to co-exist and cooperate with existing mechanisms for congestion control. In particular, new mechanisms have to co-exist with TCP's current methods of adapting to congestion and with routers' current practice of dropping packets in periods of congestion.

- * Congestion may persist over different time-scales. The time scales that we are concerned with are congestion events that may last longer than a round-trip time.

- * The number of packets in an individual flow (e.g., TCP connection or an exchange using UDP) may range from a small number of packets to quite a large number. We are interested in managing the congestion caused by flows that send enough packets so that they are still active when network feedback reaches them.

- * Asymmetric routing is likely to be a normal occurrence in the Internet. The path (sequence of links and routers) followed by data

packets may be different from the path followed by the acknowledgment packets in the reverse direction.

* Many routers process the "regular" headers in IP packets more efficiently than they process the header information in IP options. This suggests keeping congestion experienced information in the regular headers of an IP packet.

* It must be recognized that not all end-systems will cooperate in mechanisms for congestion control. However, new mechanisms shouldn't make it easier for TCP applications to disable TCP congestion control. The benefit of lying about participating in new mechanisms such as ECN-capability should be small.

4. Active Queue Management (AQM)

Random Early Detection (RED) is one mechanism for Active Queue Management (AQM) that has been proposed to detect incipient congestion [FJ93], and is currently being deployed in the Internet [RFC2309]. AQM is meant to be a general mechanism using one of several alternatives for congestion indication, but in the absence of ECN, AQM is restricted to using packet drops as a mechanism for congestion indication. AQM drops packets based on the average queue length exceeding a threshold, rather than only when the queue overflows. However, because AQM may drop packets before the queue actually overflows, AQM is not always forced by memory limitations to discard the packet.


AQM can set a Congestion Experienced (CE) bit in the packet header instead of dropping the packet, when such a bit is provided in the IP header and understood by the transport protocol. The use of the CE bit with ECN allows the receiver(s) to receive the packet, avoiding the potential for excessive delays due to retransmissions after packet losses. We use the term 'CE packet' to denote a packet that has the CE bit set.

5. Explicit Congestion Notification in IP

This document specifies that the Internet provide a congestion indication for incipient congestion (as in RED and earlier work [RJ90]) where the notification can sometimes be through marking packets rather than dropping them. This uses an ECN field in the IP header with two bits. The ECN-Capable Transport (ECT) bit is set by the data sender to indicate that the end-points of the transport protocol are ECN-capable. The CE bit is set by the router to indicate congestion to the end nodes. Routers that have a packet arriving at a full queue drop the packet, just as they do in the absence of ECN.

Bits 6 and 7 in the IPv4 TOS octet are designated as the ECN field. Bit 6 is designated as the ECT bit, and bit 7 is designated as the CE bit. The IPv4 TOS octet corresponds to the Traffic Class octet in

IPv6. The definitions for the IPv4 TOS octet [[RFC791](#)] and the IPv6 Traffic Class octet have been superseded by the DS (Differentiated Services) Field [[RFC2474](#)]. Bits 6 and 7 are listed in [[RFC2474](#)] as Currently Unused. [Section 19](#) gives a brief history of the TOS octet.



DSCP: differentiated services codepoint

ECN: Explicit Congestion Notification

Figure 1: The Differentiated Services Field in IP.

Because of the unstable history of the TOS octet, the use of the ECN field as specified in this document cannot be guaranteed to be backwards compatible with all past uses of these two bits. The potential dangers of this lack of backwards compatibility are discussed in [Section 19](#).

Upon the receipt by an ECN-Capable transport of a single CE packet, the congestion control algorithms followed at the end-systems MUST be essentially the same as the congestion control response to a *single* dropped packet. For example, for ECN-Capable TCP the source TCP is required to halve its congestion window for any window of data containing either a packet drop or an ECN indication.

One reason for requiring that the congestion-control response to the CE packet be essentially the same as the response to a dropped packet is to accommodate the incremental deployment of ECN in both end-systems and in routers. Some routers may drop ECN-Capable packets (e.g., using the same AQM policies for congestion detection) while other routers set the CE bit, for equivalent levels of congestion. Similarly, a router might drop a non-ECN-Capable packet but set the CE bit in an ECN-Capable packet, for equivalent levels of congestion. If there were different congestion control responses to a CE bit indication than to a packet drop, this could result in unfair treatment for different flows.

An additional goal is that the end-systems should react to congestion at most once per window of data (i.e., at most once per round-trip time), to avoid reacting multiple times to multiple indications of congestion within a round-trip time.

For a router, the CE bit of an ECN-Capable packet should only be set

if the router would otherwise have dropped the packet as an indication of congestion to the end nodes. When the router's buffer is not yet full and the router is prepared to drop a packet to inform end nodes of incipient congestion, the router should first check to see if the ECT bit is set in that packet's IP header. If so, then instead of dropping the packet, the router MAY instead set the CE bit in the IP header.

An environment where all end nodes were ECN-Capable could allow new criteria to be developed for setting the CE bit, and new congestion control mechanisms for end-node reaction to CE packets. However, this is a research issue, and as such is not addressed in this document.

When a CE packet (i.e., a packet that has the CE bit set) is received by a router, the CE bit is left unchanged, and the packet is transmitted as usual. When severe congestion has occurred and the router's queue is full, then the router has no choice but to drop some packet when a new packet arrives. We anticipate that such packet losses will become relatively infrequent when a majority of end-systems become ECN-Capable and participate in TCP or other compatible congestion control mechanisms. In an ECN-Capable environment that is adequately-provisioned network, packet losses should occur primarily during transients or in the presence of non-cooperating sources.

We expect that routers will set the CE bit in response to incipient congestion as indicated by the average queue size, using the RED algorithms suggested in [FJ93, [RFC2309](#)]. To the best of our knowledge, this is the only proposal currently under discussion in the IETF for routers to drop packets proactively, before the buffer overflows. However, this document does not attempt to specify a particular mechanism for active queue management, leaving that endeavor, if needed, to other areas of the IETF. While ECN is inextricably tied up with the need to have a reasonable active queue management mechanism at the router, the reverse does not hold; active queue management mechanisms have been developed and deployed independent of ECN, using packet drops as indications of congestion in the absence of ECN in the IP architecture.

5.1. ECN as an indication of persistent congestion

We emphasize that a *single* packet with the CE bit set in an IP packet causes the transport layer to respond, in terms of congestion control, as it would to a packet drop. The instantaneous queue size is likely to see considerable variations even when the router does not experience persistent congestion. As such, it is important that transient congestion at a router, reflected by the instantaneous queue size reaching a threshold much smaller than the capacity of the

queue, not trigger a reaction at the transport layer. Therefore, the CE bit should not be set by a router based on the instantaneous queue size.

For example, since the ATM and Frame Relay mechanisms for congestion indication have typically been defined without an associated notion of average queue size as the basis for determining that an intermediate node is congested, we believe that they provide a very noisy signal. The TCP-sender reaction specified in this document for ECN is NOT the appropriate reaction for such a noisy signal of congestion notification. However, if the routers that interface to the ATM network have a way of maintaining the average queue at the interface, and use it to come to a reliable determination that the ATM subnet is congested, they may use the ECN notification that is defined here.

We continue to encourage experiments in techniques at layer 2 (e.g., in ATM switches or Frame Relay switches) to take advantage of ECN. For example, using a scheme such as RED (where packet marking is based on the average queue length exceeding a threshold), layer 2 devices could provide a reasonably reliable indication of congestion. When all the layer 2 devices in a path set that layer's own Congestion Experienced bit (e.g., the EFCI bit for ATM, the FECN bit in Frame Relay) in this reliable manner, then the interface router to the layer 2 network could copy the state of that layer 2 Congestion Experienced bit into the CE bit in the IP header. We recognize that this is not the current practice, nor is it in current standards. However, encouraging experimentation in this manner may provide the information needed to enable evolution of existing layer 2 mechanisms to provide a more reliable means of congestion indication, when they use a single bit for indicating congestion.

5.2. Dropped or Corrupted Packets

For the proposed use for ECN in this document (that is, for a transport protocol such as TCP for which a dropped data packet is an indication of congestion), end nodes detect dropped data packets, and the congestion response of the end nodes to a dropped data packet is at least as strong as the congestion response to a received CE packet. To ensure the reliable delivery of the congestion indication of the CE bit, the ECT bit MUST NOT be set in a packet unless the loss of that packet in the network would be detected by the end nodes and interpreted as an indication of congestion.

Transport protocols such as TCP do not necessarily detect all packet drops, such as the drop of a "pure" ACK packet; for example, TCP does not reduce the arrival rate of subsequent ACK packets in response to an earlier dropped ACK packet. Any proposal for extending ECN-Capability to such packets would have to address issues such as the case

of an ACK packet that was marked with the CE bit but was later dropped in the network. We believe that this aspect is still the subject of research, so this document specifies that at this time, "pure" ACK packets MUST NOT indicate ECN-Capability.

Similarly, if a CE packet is dropped later in the network due to corruption (bit errors), the end nodes should still invoke congestion control, just as TCP would today in response to a dropped data packet. This issue of corrupted CE packets would have to be considered in any proposal for the network to distinguish between packets dropped due to corruption, and packets dropped due to congestion or buffer overflow. In particular, the ubiquitous deployment of ECN would not, in and of itself, be a sufficient development to allow end-nodes to interpret packet drops as indications of corruption rather than congestion.

6. Support from the Transport Protocol

ECN requires support from the transport protocol, in addition to the functionality given by the ECN field in the IP packet header. The transport protocol might require negotiation between the endpoints during setup to determine that all of the endpoints are ECN-capable, so that the sender can set the ECT bit in transmitted packets. Second, the transport protocol must be capable of reacting appropriately to the receipt of CE packets. This reaction could be in the form of the data receiver informing the data sender of the received CE packet (e.g., TCP), of the data receiver unsubscribing to a layered multicast group (e.g., RLM [MJV96]), or of some other action that ultimately reduces the arrival rate of that flow on that congested link.

This document only addresses the addition of ECN Capability to TCP, leaving issues of ECN in other transport protocols to further research. For TCP, ECN requires three new pieces of functionality: negotiation between the endpoints during connection setup to determine if they are both ECN-capable; an ECN-Echo (ECE) flag in the TCP header so that the data receiver can inform the data sender when a CE packet has been received; and a Congestion Window Reduced (CWR) flag in the TCP header so that the data sender can inform the data receiver that the congestion window has been reduced. The support required from other transport protocols is likely to be different, particularly for unreliable or reliable multicast transport protocols, and will have to be determined as other transport protocols are brought to the IETF for standardization.

6.1. TCP

The following sections describe in detail the proposed use of ECN in TCP. This proposal is described in essentially the same form in

[Floyd94]. We assume that the source TCP uses the standard congestion control algorithms of Slow-start, Fast Retransmit and Fast Recovery [RFC 2001].

This proposal specifies two new flags in the Reserved field of the TCP header. The TCP mechanism for negotiating ECN-Capability uses the ECN-Echo (ECE) flag in the TCP header. Bit 9 in the Reserved field of the TCP header is designated as the ECN-Echo flag. The location of the 6-bit Reserved field in the TCP header is shown in Figure 3 of RFC 793 [RFC793] (and is reproduced below for completeness). This specification of the ECN Field leaves the Reserved field as a 4-bit field using bits 4-7.

To enable the TCP receiver to determine when to stop setting the ECN-Echo flag, we introduce a second new flag in the TCP header, the CWR flag. The CWR flag is assigned to Bit 8 in the Reserved field of the TCP header.

| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | | |
|---|---|---|---|----------|---|---|---|---|---|----|----|----|----|----|----|---|---|
| +---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+ | | | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | U | A | P | R | S | F |
| Header Length | | | | Reserved | | | | | | | | R | C | S | S | Y | I |
| | | | | | | | | | | | | G | K | H | T | N | N |
| +---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+ | | | | | | | | | | | | | | | | | |

Figure 2: The old definition of bytes 13 and 14 of the TCP header.

| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 |
|---|---|---|---|----------|---|---|---|---|---|----|----|----|----|----|----|
| +---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+ | | | | | | | | | | | | | | | |
| | | | | | | | | C | E | U | A | P | R | S | F |
| Header Length | | | | Reserved | | | | W | C | R | C | S | S | Y | I |
| | | | | | | | | R | E | G | K | H | T | N | N |
| +---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+ | | | | | | | | | | | | | | | |

Figure 3: The new definition of bytes 13 and 14 of the TCP Header.

Thus, ECN uses the ECT and CE flags in the IP header (as shown in Figure 1) for signaling between routers and connection endpoints, and uses the ECN-Echo and CWR flags in the TCP header (as shown in Figure 3) for TCP-endpoint to TCP-endpoint signaling. For a TCP connection, a typical sequence of events in an ECN-based reaction to congestion is as follows:

- * The ECT bit is set in packets transmitted by the sender to indicate that ECN is supported by the transport entities for these packets.
- * An ECN-capable router detects impending congestion and detects

that the ECT bit is set in the packet it is about to drop.

Instead of dropping the packet, the router chooses to set the CE bit in the IP header and forwards the packet.

- * The receiver receives the packet with the CE bit set, and sets the ECN-Echo flag in its next TCP ACK sent to the sender.
- * The sender receives the TCP ACK with ECN-Echo set, and reacts to the congestion as if a packet had been dropped.
- * The sender sets the CWR flag in the TCP header of the next packet sent to the receiver to acknowledge its receipt of and reaction to the ECN-Echo flag.

The negotiation for using ECN by the TCP transport entities and the use of the ECN-Echo and CWR flags is described in more detail in the sections below.

6.1.1 TCP Initialization

In the TCP connection setup phase, the source and destination TCPs exchange information about their willingness to use ECN. Subsequent to the completion of this negotiation, the TCP sender sets the ECT bit in the IP header of data packets to indicate to the network that the transport is capable and willing to participate in ECN for this packet. This indicates to the routers that they may mark this packet with the CE bit, if they would like to use that as a method of congestion notification. If the TCP connection does not wish to use ECN notification for a particular packet, the sending TCP sets the ECT bit equal to 0 (i.e., not set), and the TCP receiver ignores the CE bit in the received packet.

For this discussion, we designate the initiating host as Host A and the responding host as Host B. We call a SYN packet with the ECE and CWR flags set an "ECN-setup SYN packet", and we call a SYN packet with the ECE and CWR flags not set a "non-ECN-setup SYN packet". Similarly, we call a SYN-ACK packet with only the ECE flag set but the CWR flag not set an "ECN-setup SYN-ACK packet", and we call a SYN-ACK packet with both the ECE and CWR flags not set a "non-ECN-setup SYN-ACK packet".

Before a TCP connection can use ECN, Host A sends an ECN-setup SYN packet, and Host B sends an ECN-setup SYN-ACK packet. For a SYN packet, the setting of both ECE and CWR in the ECN-setup SYN packet is defined as an indication that the sending TCP is ECN-Capable, rather than as an indication of congestion or of response to congestion. More precisely, an ECN-setup SYN packet indicates that the TCP implementation transmitting the SYN packet will participate in ECN as both a sender and receiver. Specifically, as a receiver, it will respond to incoming data packets that have the CE bit set in the IP header by setting ECE in outgoing TCP Acknowledgement (ACK) packets.

As a sender, it will respond to incoming packets that have ECE set by reducing the congestion window and setting CWR when appropriate. An ECN-setup SYN packet does not commit the TCP sender to setting the ECT bit in any or all of the packets it may transmit. However, the commitment to respond appropriately to incoming packets with the CE bit set remains even if the TCP sender in a later transmission, within this TCP connection, sends a SYN packet without ECE and CWR set.

When Host B sends an ECN-setup SYN-ACK packet, it sets the ECE flag but not the CWR flag. An ECN-setup SYN-ACK packet is defined as an indication that the TCP transmitting the SYN-ACK packet is ECN-Capable. As with the SYN packet, an ECN-setup SYN-ACK packet does not commit the TCP host to setting the ECT bit in transmitted packets.

The following rules apply to the sending of ECN-setup packets:

- * If a host has received an ECN-setup SYN packet, then it MAY send an ECN-setup SYN-ACK packet. Otherwise, it MUST NOT send an ECN-setup SYN-ACK packet.
- * A host MUST NOT set ECT on data packets unless it has sent at least one ECN-setup SYN or ECN-setup SYN-ACK packet, and has received at least one ECN-setup SYN or ECN-setup SYN-ACK packet, and has sent no non-ECN-setup SYN or non-ECN-setup SYN-ACK packet. If a host has received at least one non-ECN-setup SYN or non-ECN-setup SYN-ACK packet, then it SHOULD NOT set ECT on data packets.
- * If a host ever sets the ECT bit on a data packet, then that host MUST correctly set/clear the CWR TCP bit on all subsequent packets in the connection.
- * If a host has sent at least one ECN-setup SYN or ECN-setup SYN-ACK packet, and has received no non-ECN-setup SYN or non-ECN-setup SYN-ACK packet, then if that host receives TCP data packets with ECT and CE bits set in the IP header, then that host MUST process these packets as specified for an ECN-capable connection.

6.1.1.1. Robust TCP Initialization with an Echoed Reserve Field

There is the question of why we chose to have the TCP sending the SYN set two ECN-related flags in the Reserved field of the TCP header for the SYN packet, while the responding TCP sending the SYN-ACK sets only one ECN-related flag in the SYN-ACK packet. This asymmetry is necessary for the robust negotiation of ECN-capability with some deployed TCP implementations. There exists at least one faulty TCP implementation in which TCP receivers set the Reserved field of the TCP header in ACK packets (and hence the SYN-ACK) simply to reflect the Reserved field of the TCP header in the received data packet. Because the TCP SYN packet sets the ECN-Echo and CWR flags to indicate ECN-capability, while the SYN-ACK packet sets only the ECN-Echo

flag, the sending TCP correctly interprets a receiver's reflection of its own flags in the Reserved field as an indication that the receiver is not ECN-capable. The sending TCP is not misled by a faulty TCP implementation sending a SYN-ACK packet that simply reflects the Reserved field of the incoming SYN packet.

6.1.1.2. Robust TCP Initialization with no response to the SYN

ECN introduces the use of the ECN-Echo and CWR flags in the TCP header (as shown in Figure 3) for initialization. There exists some faulty equipment in the Internet that either ignores an ECN-setup SYN packet or responds with a RST, in the belief that such a packet (with these bits set) is a signature for a port-scanning tool that could be used in a denial-of-service attack. To provide robust connectivity even in the presence of such faulty equipment, a host that receives a RST in response to the transmission of an ECN-setup SYN packet MAY resend a SYN with CWR and ECE cleared. This could result in a TCP connection being established without using ECN. Similarly, a host that receives no reply to an ECN-setup SYN within the normal SYN retransmission timeout interval MAY resend the SYN and any subsequent SYN retransmissions with CWR and ECE cleared. To overcome normal packet loss that results in the original SYN being lost, the originating host may retransmit one or more ECN-setup SYN packets before giving up and retransmitting the SYN with the CWR and ECE bits cleared.

We note that in this case, the following example scenario is possible:

- (1) Host A: Sends an ECN-setup SYN.
- (2) Host B: Sends an ECN-setup SYN/ACK, packet is dropped or delayed.
- (3) Host A: Sends a non-ECN-setup SYN.
- (4) Host B: Sends a non-ECN-setup SYN/ACK.

We note that in this case, following the procedures above, neither Host A nor Host B may set the ECT bit on data packets. We further note that a host NEVER uses the reception of ECT data packets as an implicit signal that the other host is ECN-capable.

6.1.2. The TCP Sender

For a TCP connection using ECN, data packets are transmitted with the ECT bit set in the IP header (set to a "1"). If the sender receives an ECN-Echo (ECE) ACK packet (that is, an ACK packet with the ECN-Echo flag set in the TCP header), then the sender knows that congestion was encountered in the network on the path from the sender to the receiver. The indication of congestion should be treated just as

a congestion loss in non-ECN-Capable TCP. That is, the TCP source halves the congestion window "cwnd" and reduces the slow start threshold "ssthresh". The sending TCP SHOULD NOT increase the congestion window in response to the receipt of an ECN-Echo ACK packet.

TCP should not react to congestion indications more than once every window of data (or more loosely, more than once every round-trip time). That is, the TCP sender's congestion window should be reduced only once in response to a series of dropped and/or CE packets from a single window of data. In addition, the TCP source should not decrease the slow-start threshold, ssthresh, if it has been decreased within the last round trip time. However, if any retransmitted packets are dropped, then this is interpreted by the source TCP as a new instance of congestion.

After the source TCP reduces its congestion window in response to a CE packet, incoming acknowledgements that continue to arrive can "clock out" outgoing packets as allowed by the reduced congestion window. If the congestion window consists of only one MSS (maximum segment size), and the sending TCP receives an ECN-Echo ACK packet, then the sending TCP should in principle still reduce its congestion window in half. However, the value of the congestion window is bounded below by a value of one MSS. If the sending TCP were to continue to send, using a congestion window of 1 MSS, this results in the transmission of one packet per round-trip time. It is necessary to still reduce the sending rate of the TCP sender even further, on receipt of an ECN-Echo packet when the congestion window is one. We use the retransmit timer as a means of reducing the rate further in this circumstance. Therefore, the sending TCP MUST reset the retransmit timer on receiving the ECN-Echo packet when the congestion window is one. The sending TCP will then be able to send a new packet only when the retransmit timer expires.

[Floyd94] discusses TCP's response to ECN in more detail. [Floyd98] discusses the validation test in the ns simulator, which illustrates a wide range of ECN scenarios. These scenarios include the following: an ECN followed by another ECN, a Fast Retransmit, or a Retransmit Timeout; a Retransmit Timeout or a Fast Retransmit followed by an ECN; and a congestion window of one packet followed by an ECN.

TCP follows existing algorithms for sending data packets in response to incoming ACKs, multiple duplicate acknowledgements, or retransmit timeouts [RFC2581]. TCP also follows the normal procedures for increasing the congestion window when it receives ACK packets without the ECN-Echo bit set [RFC2581].

6.1.3. The TCP Receiver

When TCP receives a CE data packet at the destination end-system, the TCP data receiver sets the ECN-Echo flag in the TCP header of the subsequent ACK packet. If there is any ACK withholding implemented, as in current "delayed-ACK" TCP implementations where the TCP receiver can send an ACK for two arriving data packets, then the ECN-Echo flag in the ACK packet will be set to the OR of the CE bits of all of the data packets being acknowledged. That is, if any of the received data packets are CE packets, then the returning ACK has the ECN-Echo flag set.

To provide robustness against the possibility of a dropped ACK packet carrying an ECN-Echo flag, the TCP receiver sets the ECN-Echo flag in a series of ACK packets sent subsequently. The TCP receiver uses the CWR flag received from the TCP sender to determine when to stop setting the ECN-Echo flag.

When an ECN-Capable TCP sender reduces its congestion window for any reason (because of a retransmit timeout, a Fast Retransmit, or in response to an ECN Notification), the TCP sender sets the CWR flag in the TCP header of the first new data packet sent after the window reduction. If that data packet is dropped in the network, then the sending TCP will have to reduce the congestion window again and retransmit the dropped packet.

We ensure that the "Congestion Window Reduced" information is reliably delivered to the TCP receiver. This comes about from the fact that if the new data packet carrying the CWR flag is dropped, then the TCP sender will have to again reduce its congestion window, and send another new data packet with the CWR flag set. Thus, the CWR bit in the TCP header SHOULD NOT be set on retransmitted packets. When the TCP data sender is ready to set the CWR bit after reducing the congestion window, it SHOULD set the CWR bit only on the first new data packet that it transmits.

After a TCP receiver sends an ACK packet with the ECN-Echo bit set, that TCP receiver continues to set the ECN-Echo flag in all the ACK packets it sends (whether they acknowledge CE data packets or non-CE data packets) until it receives a CWR packet (a packet with the CWR flag set). After the receipt of the CWR packet, acknowledgements for subsequent non-CE data packets do not have the ECN-Echo flag set. If another CE packet is received by the data receiver, the receiver would once again send ACK packets with the ECN-Echo flag set. While the receipt of a CWR packet does not guarantee that the data sender received the ECN-Echo message, this does suggest that the data sender reduced its congestion window at some point *after* it sent the data packet for which the CE bit was set.

We have already specified that a TCP sender is not required to reduce its congestion window more than once per window of data. Some care is required if the TCP sender is to avoid unnecessary reductions of the congestion window when a window of data includes both dropped packets and (marked) CE packets. This is illustrated in [Floyd98].

6.1.4. Congestion on the ACK-path

For the current generation of TCP congestion control algorithms, pure acknowledgement packets (e.g., packets that do not contain any accompanying data) should be sent with the ECT bit off. Current TCP receivers have no mechanisms for reducing traffic on the ACK-path in response to congestion notification. Mechanisms for responding to congestion on the ACK-path are areas for current and future research. (One simple possibility would be for the sender to reduce its congestion window when it receives a pure ACK packet with the CE bit set). For current TCP implementations, a single dropped ACK generally has only a very small effect on the TCP's sending rate.

6.1.5. Retransmitted TCP packets

This document specifies that for ECN-capable TCP implementations, the ECT bit (ECN-Capable Transport) in the IP header MUST NOT be set on retransmitted data packets, and that the TCP data receiver SHOULD ignore the ECN field on arriving data packets that are outside of the receiver's current window. This is for greater security against denial-of-service attacks, as well as for robustness of the ECN congestion indication with packets that are dropped later in the network.

First, we note that if the TCP sender were to set the ECT bit on a retransmitted packet, then if an unnecessarily-retransmitted packet was later dropped in the network, the end nodes would never receive the indication of congestion from the router setting the CE bit. Thus, setting the ECT bit on retransmitted data packets is not consistent with the robust delivery of the congestion indication even for packets that are later dropped in the network.

In addition, an attacker capable of spoofing the IP source address of the TCP sender could send data packets with arbitrary sequence numbers, with both the ECT and CE bits set in the IP header. On receiving this spoofed data packet, the TCP data receiver would determine that the data does not lie in the current receive window, and return a duplicate acknowledgement. We define an out-of-window packet at the TCP data receiver as a data packet that lies outside the receiver's current window. On receiving an out-of-window packet, the TCP data receiver has to decide whether or not to treat the CE bit in the packet header as a valid indication of congestion, and therefore

whether to return ECN-Echo indications to the TCP data sender. If the TCP data receiver ignored the CE bit in an out-of-window packet, then the TCP data sender would not receive this possibly-legitimate indication of congestion from the network, resulting in a violation of end-to-end congestion control. On the other hand, if the TCP data receiver honors the CE indication in the out-of-window packet, and reports the indication of congestion to the TCP data sender, then the malicious node that created the spoofed, out-of-window packet has successfully "attacked" the TCP connection by forcing the data sender to unnecessarily reduce (halve) its congestion window. To prevent such a denial-of-service attack, we specify that a legitimate TCP data sender **MUST NOT** set the ECT bit on retransmitted data packets, and that the TCP data receiver **SHOULD** ignore the CE bit on out-of-window packets.

One drawback of not setting ECT on retransmitted packets denies ECN protection for retransmitted packets. However, for an ECN-capable TCP connection in a fully-ECN-capable environment with mild congestion, packets should rarely be dropped due to congestion in the first place, and so instances of retransmitted packets should rarely arise. If packets are being retransmitted, then there are already packet losses (from corruption or from congestion) that ECN has been unable to prevent.

We note that if the router sets the CE bit for an ECN-capable data packet within a TCP connection, then the TCP connection is guaranteed to receive that indication of congestion, or to receive some other indication of congestion within the same window of data, even if this packet is dropped or reordered in the network. We consider two cases, when the packet is later retransmitted, and when the packet is not later retransmitted.

In the first case, if the packet is either dropped or delayed, and at some point retransmitted by the data sender, then the retransmission is a result of a Fast Retransmit or a Retransmit Timeout for either that packet or for some prior packet in the same window of data. In this case, because the data sender already has retransmitted this packet, we know that the data sender has already responded to an indication of congestion for some packet within the same window of data as the original packet. Thus, even if the first transmission of the packet is dropped in the network, or is delayed, if it had the CE bit set, and is later ignored by the data receiver as an out-of-window packet, this is not a problem, because the sender has already responded to an indication of congestion for that window of data.

In the second case, if the packet is never retransmitted by the data sender, then this data packet is the only copy of this data received by the data receiver, and therefore arrives at the data receiver as

an in-window packet, regardless of how much the packet might be delayed or reordered. In this case, if the CE bit is set on the packet within the network, this will be treated by the data receiver as a valid indication of congestion.

6.1.6. TCP Window Probes.

When the TCP data receiver advertises a zero window, the TCP data sender sends window probes to determine if the receiver's window has increased. Window probe packets do not contain any user data except for the sequence number, which is a byte. If a window probe packet is dropped in the network, this loss is not detected by the receiver. Therefore, the TCP data sender **MUST NOT** set either the ECT or CWR bits on window probe packets.

However, because window probes use exact sequence numbers, they cannot be easily spoofed in denial-of-service attacks. Therefore, if a window probe arrives with ECT and CE set, then the receiver **SHOULD** respond to the ECN indications.

7. Non-compliance by the End Nodes

This section discusses concerns about the vulnerability of ECN to non-compliant end-nodes (i.e., end nodes that set the ECT bit in transmitted packets but do not respond to received CE packets). We argue that the addition of ECN to the IP architecture will not significantly increase the current vulnerability of the architecture to unresponsive flows.

Even for non-ECN environments, there are serious concerns about the damage that can be done by non-compliant or unresponsive flows (that is, flows that do not respond to congestion control indications by reducing their arrival rate at the congested link). For example, an end-node could "turn off congestion control" by not reducing its congestion window in response to packet drops. This is a concern for the current Internet. It has been argued that routers will have to deploy mechanisms to detect and differentially treat packets from non-compliant flows [[RFC2309](#), [FF99](#)]. It has also been suggested that techniques such as end-to-end per-flow scheduling and isolation of one flow from another, differentiated services, or end-to-end reservations could remove some of the more damaging effects of unresponsive flows.

It might seem that dropping packets in itself is an adequate deterrent for non-compliance, and that the use of ECN removes this deterrent. We would argue in response that (1) ECN-capable routers preserve packet-dropping behavior in times of high congestion; and (2) even in times of high congestion, dropping packets in itself is not

an adequate deterrent for non-compliance.

First, ECN-Capable routers will only mark packets (as opposed to dropping them) when the packet marking rate is reasonably low. During periods where the average queue size exceeds an upper threshold, and therefore the potential packet marking rate would be high, our recommendation is that routers drop packets rather than set the CE bit in packet headers.

During the periods of low or moderate packet marking rates when ECN would be deployed, there would be little deterrent effect on unresponsive flows of dropping rather than marking those packets. For example, delay-insensitive flows using reliable delivery might have an incentive to increase rather than to decrease their sending rate in the presence of dropped packets. Similarly, delay-sensitive flows using unreliable delivery might increase their use of FEC in response to an increased packet drop rate, increasing rather than decreasing their sending rate. For the same reasons, we do not believe that packet dropping itself is an effective deterrent for non-compliance even in an environment of high packet drop rates, when all flows are sharing the same packet drop rate.

Several methods have been proposed to identify and restrict non-compliant or unresponsive flows. The addition of ECN to the network environment would not in any way increase the difficulty of designing and deploying such mechanisms. If anything, the addition of ECN to the architecture would make the job of identifying unresponsive flows slightly easier. For example, in an ECN-Capable environment routers are not limited to information about packets that are dropped or have the CE bit set at that router itself; in such an environment, routers could also take note of arriving CE packets that indicate congestion encountered by that packet earlier in the path.

8. Non-compliance in the Network

This section considers the issues when a router is operating, possibly maliciously, to modify either of the bits in the ECN field. In this section we represent the ECN field in the IP header by the tuple (ECT bit, CE bit).

By tampering with the bits in the ECN field, an adversary (or a broken router) could do one or more of the following: falsely report congestion, disable ECN-Capability for an individual packet, erase the ECN congestion indication, or falsely indicate ECN-Capability. [Appendix X](#) systematically examines the various cases by which the ECN field could be modified. The important criterion considered in determining the consequences of such modifications is whether it is likely to lead to poorer behavior in any dimension (throughput,

delay, fairness or functionality) than if a router were to drop a packet.

The first two possible changes, falsely reporting congestion or disabling ECN-Capability for an individual packet, are no worse than if the router were to simply drop the packet. From a congestion control point of view, setting the CE bit in the absence of congestion by a non-compliant router would be no worse than a router dropping a packet unnecessarily. By "erasing" the ECT bit of a packet that is later dropped in the network, a router's actions could result in an unnecessary packet drop for that packet later in the network.

However, as discussed in Section X in the Appendix, a router that erases the ECN congestion indication or falsely indicates ECN-Capability could potentially do more damage to the flow than if it has simply dropped the packet. A rogue or broken router that "erased" the CE bit in arriving CE packets would prevent that indication of congestion from reaching downstream receivers. This could result in the failure of congestion control for that flow and a resulting increase in congestion in the network, ultimately resulting in subsequent packets dropped for this flow as the average queue size increased at the congested gateway.

[Appendix X](#) considers the potential repercussions of subverting end-to-end congestion control by either falsely indicating ECN-Capability, or by erasing the congestion indication in ECN (the CE-bit). We observe in the Appendix that the consequence of subverting ECN-based congestion control may lead to potential unfairness, but this is likely to be no worse than the subversion of either ECN-based or packet-based congestion control by the end nodes.

[8.1. Complications Introduced by Split Paths](#)

If a router or other network element has access to all of the packets of a flow, then that router could do no more damage to a flow by altering the ECN field than it could by simply dropping all of the packets from that flow. However, in some cases, a malicious or broken router might have access to only a subset of the packets from a flow. The question is as follows: can this router, by altering the ECN field in this subset of the packets, do more damage to that flow than if it has simply dropped that set of the packets?

This is also discussed in detail in the Appendix, which concludes as follows: It is true that the adversary that has access only to a subset of packets in an aggregate might, by subverting ECN-based congestion control, be able to deny the benefits of ECN to the other packets in the aggregate. While this is undesirable, this is not a sufficient concern to result in disabling ECN within an IP tunnel.

9. Encapsulated Packets

9.1. IP packets encapsulated in IP

The encapsulation of IP packet headers in tunnels is used in many places, including IPsec and IP in IP [RFC2003]. Currently, the ECN specification does not accommodate the constraints imposed by some of these pre-existing specifications for tunnels. This document considers issues related to interactions between ECN and IP tunnels, and specifies two alternative solutions.

Some IP tunnel modes are based on adding a new "outer" IP header that encapsulates the original, or "inner" IP header and its associated packet. In many cases, the new "outer" IP header may be added and removed at intermediate points along a connection, enabling the network to establish a tunnel without requiring endpoint participation. We denote tunnels that specify that the outer header be discarded at tunnel egress as "simple tunnels".

ECN uses the ECT and CE flags in the IP header for signaling between routers and connection endpoints. ECN interacts with IP tunnels because of the ECT and CE flags in the DS field octet in the IP header [RFC2474] (also referred to as the IPv4 TOS octet or IPv6 Traffic Class octet). [RFC2983] discusses interactions of Differentiated Services with IP tunnels of various forms. In simple IP tunnels the DS field octet is copied or mapped from the inner IP header to the outer IP header at IP tunnel ingress, and the outer header's copy of this field is discarded at IP tunnel egress. If the outer header were to be simply discarded without taking care to deal with the ECN related flags, and an ECN-capable router were to set the CE (Congestion Experienced) bit within a packet in a simple IP tunnel, this indication would be discarded at tunnel egress, losing the indication of congestion.

Thus, the use of ECN over simple IP tunnels would result in routers attempting to use the outer IP header to signal congestion to endpoints, but those congestion warnings never arriving because the outer header is discarded at the tunnel egress point. This problem was encountered with ECN and IPsec in tunnel mode, and RFC 2481 recommended that ECN not be used with the older simple IPsec tunnels in order to avoid this behavior and its consequences. When ECN becomes widely deployed, then simple tunnels likely to carry ECN-capable traffic will have to be changed.

From a security point of view, the use of ECN in the outer header of an IP tunnel might raise security concerns because an adversary could tamper with the ECN information that propagates beyond the tunnel endpoint. Based on an analysis in the Appendix of these concerns and

the resultant risks, our overall approach is to make support for ECN an option for IP tunnels, so that an IP tunnel can be specified or configured either to use ECN or not to use ECN in the outer header of the tunnel. Thus, in environments or tunneling protocols where the risks of using ECN are judged to outweigh its benefits, the tunnel can simply not use ECN in the outer header. Then the only indication of congestion experienced at routers within the tunnel would be through packet loss.

The result is that there are two viable options for the behavior of ECN-capable connections over an IP tunnel, especially IPsec tunnels:

- * A limited-functionality option in which ECN is preserved in the inner header, but disabled in the outer header. The only mechanism available for signaling congestion occurring within the tunnel in this case is dropped packets.
- * A full-functionality option that supports ECN in both the inner and outer headers, and propagates congestion warnings from nodes within the tunnel to endpoints.

Support for these options requires varying amounts of changes to IP header processing at tunnel ingress and egress. A small subset of these changes sufficient to support only the limited-functionality option would be sufficient to eliminate any incompatibility between ECN and IP tunnels.

One goal of this document is to give guidance about the tradeoffs between the limited-functionality and full-functionality options. A full discussion of the potential effects of an adversary's modifications of the CE and ECT bits is given in the Appendix.

9.1.1. The limited-functionality and full-functionality options within IP Tunnels

The limited-functionality option for ECN encapsulation in IP tunnels is for the ECT bit in the outside (encapsulating) header to be off (i.e., set to 0), regardless of the value of the ECT bit in the inside (encapsulated) header. With this option, the ECN field in the inner header is not altered upon de-capsulation. The disadvantage of this approach is that the flow does not have ECN support for that part of the path that is using IP tunneling, even if the encapsulated packet (from the original TCP sender) is ECN-Capable. That is, if the encapsulated packet arrives at a congested router that is ECN-capable, and the router can decide to drop or mark the packet as an indication of congestion to the end nodes, the router will not be permitted to set the CE bit in the packet header, but instead will have to drop the packet.

The IP full-functionality option for ECN encapsulation is to copy the

ECT bit of the inside header to the outside header on encapsulation, and to OR the CE bit from the outer header with the CE bit of the inside header on decapsulation. That is, for full ECN support the encapsulation and decapsulation processing for the DS field octet involves the following: At tunnel ingress, the full-functionality option copies the value of ECT (bit 6) in the inner header to the outer header. CE (bit 7) is set to 0 in the outer header. Upon decapsulation at the tunnel egress, the full-functionality option sets CE to 1 in the inner header if the value of ECT (bit 6) in the inner header is 1, and the value of CE (bit 7) in the outer header is 1. Otherwise, no change is made to this field of the inner header.

With the full-functionality option, a flow can take advantage of ECN in those parts of the path that might use IP tunneling. The disadvantage of the full-functionality option from a security perspective is that the IP tunnel cannot protect the flow from certain modifications to the ECN bits in the IP header within the tunnel. The potential dangers from modifications to the ECN bits in the IP header are described in detail in the Appendix.

- (1) An IP tunnel MUST modify the handling of the DS field octet at IP tunnel endpoints by implementing either the limited-functionality or the full-functionality option.
- (2) Optionally, an IP tunnel MAY enable the endpoints of an IP tunnel to negotiate the choice between the limited-functionality and the full-functionality option for ECN in the tunnel.

The minimum required to make ECN usable with IP tunnels is the limited-functionality option, which prevents ECN from being enabled in the outer header of an IPsec tunnel. Full support for ECN requires the use of the full-functionality option. If there are no optional mechanisms for the tunnel endpoints to negotiate a choice between the limited-functionality or full-functionality option, there can be a pre-existing agreement between the tunnel endpoints about whether to support the limited-functionality or the full-functionality ECN option.

In addition, it is RECOMMENDED that packets with ECT and CE both set to 1 in the outer header be dropped if they arrive at the tunnel egress point for a tunnel that uses the limited-functionality option, or for a tunnel that uses the full-functionality option but for which the ECT bit in the inner header is set to zero. This is motivated by backwards compatibility and to ensure that no unauthorized modifications of the ECN field take place, and is discussed further in the Appendix.

9.1.2. Changes to the ECN Field within an IP Tunnel.

The presence of a copy of the ECN field in the inner header of an IP tunnel mode packet provides an opportunity for detection of unauthorized modifications to the ECT bit in the outer header. Comparison of the ECT bits in the inner and outer headers falls into two categories for implementations that conform to this document:

- * If the IP tunnel uses the full-functionality option, then the values of the ECT bits in the inner and outer headers should be identical.
- * If the tunnel uses the limited-functionality option, then the ECT bit in the outer header should be 0.

Receipt of a packet not satisfying the appropriate condition could be a cause of concern.

Consider the case of an IP tunnel where the tunnel ingress point has not been updated to this document's requirements, while the tunnel egress point has been updated to support ECN. In this case, the IP tunnel is not explicitly configured to support the full-functionality ECN option. However, the tunnel ingress point is behaving identically to a tunnel ingress point that supports the full-functionality option. If packets from an ECN-capable connection use this tunnel, ECT will be set to 1 in the outer header at the tunnel ingress point. Congestion within the tunnel may then result in ECN-capable routers setting CE in the outer header. Because the tunnel has not been explicitly configured to support the full-functionality option, the tunnel egress point expects the ECT bit in the outer header to be 0. When an ECN-capable tunnel egress point receives a packet with the ECT bit in the outer header set to 1, in a tunnel that has not been configured to support the full-functionality option, that packet should be processed, according to whether CE bit was set, as follows. It is RECOMMENDED that such packets, with the ECT bit in the outer header set to 1 on a tunnel that has not been configured to support the full-functionality option, be dropped at the egress point if CE is set to 1 in the outer header but 0 in the inner header, and forwarded otherwise.

An IP tunnel cannot provide protection against erasure of congestion indications based on resetting the value of the CE bit in packets for which ECT is set in the outer header. The erasure of congestion indications may impact the network and other flows in ways that would not be possible in the absence of ECN. It is important to note that erasure of congestion indications can only be performed to congestion indications placed by nodes within the tunnel; the copy of the CE bit in the inner header preserves congestion notifications from nodes upstream of the tunnel ingress. If erasure of congestion notifications is judged to be a security risk that exceeds the congestion

management benefits of ECN, then tunnels could be specified or configured to use the limited-functionality option.

9.2. IPsec Tunnels

IPsec supports secure communication over potentially insecure network components such as intermediate routers. IPsec protocols support two operating modes, transport mode and tunnel mode, that span a wide range of security requirements and operating environments. Transport mode security protocol header(s) are inserted between the IP (IPv4 or IPv6) header and higher layer protocol headers (e.g., TCP), and hence transport mode can only be used for end-to-end security on a connection. IPsec tunnel mode is based on adding a new "outer" IP header that encapsulates the original, or "inner" IP header and its associated packet. Tunnel mode security headers are inserted between these two IP headers. In contrast to transport mode, the new "outer" IP header and tunnel mode security headers can be added and removed at intermediate points along a connection, enabling security gateways to secure vulnerable portions of a connection without requiring endpoint participation in the security protocols. An important aspect of tunnel mode security is that in the original specification, the outer header is discarded at tunnel egress, ensuring that security threats based on modifying the IP header do not propagate beyond that tunnel endpoint. Further discussion of IPsec can be found in [[RFC 2401](#)].

The IPsec protocol as originally defined in [[ESP](#), [AH](#)] required that the inner header's ECN field not be changed by IPsec decapsulation processing at a tunnel egress node; this would have ruled out the possibility of full-functionality mode for ECN. At the same time, this would ensure that an adversary's modifications to the ECN field cannot be used to launch theft- or denial-of-service attacks across an IPsec tunnel endpoint, as any such modifications will be discarded at the tunnel endpoint.

In principle, permitting the use of ECN functionality in the outer header of an IPsec tunnel raises security concerns because an adversary could tamper with the information that propagates beyond the tunnel endpoint. Based on an analysis (included in the Appendix) of these concerns and the associated risks, our overall approach has been to provide configuration support for IPsec changes to remove the conflict with ECN.

In particular, in tunnel mode the IPsec tunnel **MUST** support either the limited-functionality or the full-functionality mode outlined in Section X.

This makes permission to use ECN functionality in the outer header of an IPsec tunnel a configurable part of the corresponding IPsec

Security Association (SA), so that it can be disabled in situations where the risks are judged to outweigh the benefits. The result is that an IPsec security administrator is presented with two alternatives for the behavior of ECN-capable connections within an IPsec tunnel, the limited-functionality alternative and full-functionality alternative described earlier. All IPsec implementations MUST implement either the limited-functionality or the full-functionality alternative in order to eliminate incompatibility between ECN and IPsec tunnels, but implementers MAY choose to implement either alternative.

In addition, this document specifies how the endpoints of an IPsec tunnel could negotiate enabling ECN functionality in the outer headers of that tunnel based on security policy. The ability to negotiate ECN usage between tunnel endpoints would enable a security administrator to disable ECN in situations where she believes the risks (e.g., of lost congestion notifications) outweigh the benefits of ECN.

The IPsec protocol, as defined in [ESP, AH], does not include the IP header's ECN field in any of its cryptographic calculations (in the case of tunnel mode, the outer IP header's ECN field is not included). Hence modification of the ECN field by a network node has no effect on IPsec's end-to-end security, because it cannot cause any IPsec integrity check to fail. As a consequence, IPsec does not provide any defense against an adversary's modification of the ECN field (i.e., a man-in-the-middle attack), as the adversary's modification will also have no effect on IPsec's end-to-end security. In some environments, the ability to modify the ECN field without affecting IPsec integrity checks may constitute a covert channel; if it is necessary to eliminate such a channel or reduce its bandwidth, then the IPsec tunnel should be run in limited-functionality mode.

9.2.1. Negotiation between Tunnel Endpoints

This section describes the detailed changes to enable usage of ECN over IPsec tunnels, including the negotiation of ECN support between tunnel endpoints. This is supported by three changes to IPsec:

- * An optional Security Association Database (SAD) field indicating whether tunnel encapsulation and decapsulation processing allows or forbids ECN usage in the outer IP header.
- * An optional Security Association Attribute that enables negotiation of this SAD field between the two endpoints of an SA that supports tunnel mode.
- * Changes to tunnel mode encapsulation and decapsulation processing to allow or forbid ECN usage in the outer IP header based on the value of the SAD field. When ECN usage is allowed in the outer IP header, ECT is set in the outer header for ECN-capable

connections and congestion notifications (indicated by the CE bit) from such connections are propagated to the inner header at tunnel egress.

If negotiation of ECN usage is implemented, then the SAD field SHOULD also be implemented. On the other hand, negotiation of ECN usage is OPTIONAL in all cases, even for implementations that support the SAD field. The encapsulation and decapsulation processing changes are REQUIRED, but MAY be implemented without the other two changes by assuming that ECN usage is always forbidden. The full-functionality alternative for ECN usage over IPsec tunnels consists of the SAD field and the full version of encapsulation and decapsulation processing changes, with or without the OPTIONAL negotiation support. The limited-functionality alternative consists of a subset of the encapsulation and decapsulation changes that always forbids ECN usage.

These changes are covered further in the following three subsections.

9.2.1.1. ECN Tunnel Security Association Database Field

Full ECN functionality adds a new field to the SAD (see [\[RFC2401\]](#)):

ECN Tunnel: allowed or forbidden.

Indicates whether ECN-capable connections using this SA in tunnel mode are permitted to receive ECN congestion notifications for congestion occurring within the tunnel. The allowed value enables ECN congestion notifications. The forbidden value disables such notifications, causing all congestion to be indicated via dropped packets.

[OPTIONAL. The value of this field SHOULD be assumed to be "forbidden" in implementations that do not support it.]

If this attribute is implemented, then the SA specification in a Security Policy Database (SPD) entry MUST support a corresponding attribute, and this SPD attribute MUST be covered by the SPD administrative interface (currently described in [Section 4.4.1 of \[RFC2401\]](#)).

9.2.1.2. ECN Tunnel Security Association Attribute

A new IPsec Security Association Attribute is defined to enable the support for ECN congestion notifications based on the outer IP header to be negotiated for IPsec tunnels (see [\[RFC2407\]](#)). This attribute is OPTIONAL, although implementations that support it SHOULD also support the SAD field defined in [Section 3.1](#).

Attribute Type

| class | value | type |
|------------|-------|-------|
| ----- | | |
| ECN Tunnel | 10 | Basic |

The IPsec SA Attribute value 10 has been allocated by IANA to indicate that the ECN Tunnel SA Attribute is being negotiated; the type of this attribute is Basic (see [Section 4.5 of \[RFC2407\]](#)). The Class Values are used to conduct the negotiation. See [RFC2407, [RFC2408](#), [RFC2409](#)] for further information including encoding formats and requirements for negotiating this SA attribute.

Class Values

ECN Tunnel

Specifies whether ECN functionality is allowed to be used with Tunnel Encapsulation Mode. This affects tunnel encapsulation and decapsulation processing - see [Section 3.3](#).

| | |
|-----------|---|
| RESERVED | 0 |
| Allowed | 1 |
| Forbidden | 2 |

Values 3-61439 are reserved to IANA. Values 61440-65535 are for private use.

If unspecified, the default shall be assumed to be Forbidden.

ECN Tunnel is a new SA attribute, and hence initiators that use it can expect to encounter responders that do not understand it, and therefore reject proposals containing it. For backwards compatibility with such implementations initiators SHOULD always also include a proposal without the ECN Tunnel attribute to enable such a responder to select a transform or proposal that does not contain the ECN Tunnel attribute. [RFC 2407](#) currently requires responders to reject all proposals if any proposal contains an unknown attribute; this requirement is expected to be changed to require a responder not to select proposals or transforms containing unknown attributes.

[9.2.1.3](#). Changes to IPsec Tunnel Header Processing

Subsequent to the publication of [[RFC 2401](#)], the TOS octet of IPv4 and the Traffic Class octet of IPv6 have been superseded by the six-bit DS Field [[RFC2474](#), [RFC2780](#)] and a two-bit "currently unused" (CU) field [[RFC2780](#)], and this document supersedes the CU field by the ECN

Field.

For full ECN support, the encapsulation and decapsulation processing for the IPv4 TOS field and the IPv6 Traffic Class field are changed from that specified in [\[RFC2401\]](#) to the following:

| <-- How Outer Hdr Relates to Inner Hdr --> | | |
|--|------------------------------|------------------------------|
| | Outer Hdr at Encapsulator | Inner Hdr at Decapsulator |
| IPv4 | | |
| Header fields: | ----- | ----- |
| DS Field | copied from inner hdr (5) | no change |
| ECN Field | constructed (7) | constructed (8) |
| IPv6 | | |
| Header fields: | | |
| DS Field | copied from inner hdr (6) | no change |
| ECN Field | constructed (7) | constructed (8) |

(5)(6) If the packet will immediately enter a domain for which the DSCP value in the outer header is not appropriate, that value MUST be mapped to an appropriate value for the domain [\[RFC 2474\]](#). Also see [\[RFC 2475\]](#) for further information.

(7) If the value of the ECN Tunnel field in the SAD entry for this SA is "allowed" and the value of ECT (bit 0) is 1 in the inner header, set ECT to 1 in the outer header, else set ECT to 0 in the outer header. Set CE (bit 1) to 0 in the outer header.

(8) If the value of the ECN tunnel field in the SAD entry for this SA is "allowed" and the value of ECT (bit 0) in the inner header is 1, then set the CE bit (bit 1) in the inner header to the logical OR of the CE bit in the inner header with the CE bit in the outer header, else make no change to the ECN field.

(5) and (6) are identical to match usage in [\[RFC2401\]](#), although they are different in [\[RFC2401\]](#).

The above description applies to implementations that support the ECN Tunnel field in the SAD; such implementations MUST implement this processing of the DS field instead of the processing of the IPv4 TOS octet and IPv6 Traffic Class octet defined in [\[RFC2401\]](#). This constitutes the full-functionality alternative for ECN usage with IPsec tunnels.

An implementation that does not support the ECN Tunnel field in the SAD MUST implement processing of the DS Field by assuming that the value of the ECN Tunnel field of the SAD is "forbidden" for every SA. In this case, the processing of the ECN field reduces to:

- (7) Set the ECN field (ECT and CE bits) to zero in the outer header.
- (8) Make no change to the ECN field in the inner header.

This constitutes the limited functionality alternative for ECN usage with IPsec tunnels.

For backwards compatibility, packets with ECT and CE both set to 1 in the outer header SHOULD be dropped if they arrive on an SA that is using the limited-functionality option, or that is using the full-functionality option (i.e., and has set the ECT flag in the outer header to 1) for a packet with the ECT flag set to 0 in the inner header.

9.2.2. Changes to the ECN Field within an IPsec Tunnel.

If the ECN Field is changed inappropriately within an IPsec tunnel, and this change is detected at the tunnel egress, then the receipt of a packet not satisfying the appropriate condition for its SA is an auditable event. An implementation MAY create audit records with per-SA counts of incorrect packets over some time period rather than creating an audit record for each erroneous packet. Any such audit record SHOULD contain the headers from at least one erroneous packet, but need not contain the headers from every packet represented by the entry.

9.2.3. Comments for IPsec Support

Substantial comments were received on two areas of this document during review by the IPsec working group. This section describes these comments and explains why the proposed changes were not incorporated.

The first comment indicated that per-node configuration is easier to implement than per-SA configuration. After serious thought and despite some initial encouragement of per-node configuration, it no longer seems to be a good idea. The concern is that as IPsec is progressively deployed, many ECN-aware IPsec implementations will find themselves communicating with a mixture of ECN-aware and ECN-unaware IPsec tunnel endpoints. In such an environment with per-node configuration, the only reasonable thing to do is forbid ECN usage for all IPsec tunnels, which is not the desired outcome.

In the second area, several reviewers noted that SA negotiation is complex, and adding to it is non-trivial. One reviewer suggested using ICMP after tunnel setup as a possible alternative. The addition to SA negotiation in the document is OPTIONAL and will remain so; implementers are free to ignore it. The authors believe that the assurance it provides can be useful in a number of situations. In

practice, if this is not implemented, it can be deleted at a subsequent stage in the standards process. Extending ICMP to negotiate ECN after tunnel setup is more complex than extending SA attribute negotiation. Some tunnels do not permit traffic to be addressed to the tunnel egress endpoint, hence the ICMP packet would have to be addressed to somewhere else, scanned for by the egress endpoint, and discarded there or at its actual destination. In addition, ICMP delivery is unreliable, and hence there is a possibility of an ICMP packet being dropped, entailing the invention of yet another ack/retransmit mechanism. It seems better simply to specify an OPTIONAL extension to the existing SA negotiation mechanism.

9.3. IP packets encapsulated in non-IP packet headers.

A different set of issues are raised, relative to ECN, when IP packets are encapsulated in tunnels with non-IP packet headers. This occurs with MPLS [[MPLS](#)], GRE [[GRE](#)], L2TP [[L2TP](#)], and PPTP [[PPTP](#)]. For these protocols, there is no conflict with ECN; it is just that ECN cannot be used within the tunnel unless an ECN codepoint can be specified for the header of the encapsulating protocol. [[RFD99](#)] considered a preliminary proposal for incorporating ECN into MPLS, and proposals for incorporating ECN into GRE, L2TP, or PPTP will be considered as the need arises.

10. Issues Raised by Monitoring and Policing Devices

One possibility is that monitoring and policing devices (or more informally, "penalty boxes") will be installed in the network to monitor whether best-effort flows are appropriately responding to congestion, and to preferentially drop packets from flows determined not to be using adequate end-to-end congestion control procedures. This is discussed in more detail in the Appendix.

We recommend that any "penalty box" that detects a flow or an aggregate of flows that is not responding to end-to-end congestion control first change from marking to dropping packets from that flow, before taking any additional action to restrict the bandwidth available to that flow. Thus, initially, the router may drop packets in which the router would otherwise would have set the CE bit. This could include dropping those arriving packets for that flow that are ECN-Capable and that already have the CE bit set. In this way, any congestion indications seen by that router for that flow will be guaranteed to also be seen by the end nodes, even in the presence of malicious or broken routers elsewhere in the path. If we assume that the first action taken at any "penalty box" for an ECN-capable flow will be to drop packets instead of marking them, then there is no way that an adversary that subverts ECN-based end-to-end congestion control can cause a flow to be characterized as being non-cooperative and placed

into a more severe action within the "penalty box".

The monitoring and policing devices that are actually deployed could fall short of the 'ideal' monitoring device described above, in that the monitoring is applied not to a single flow or to a single IPsec tunnel, but to an aggregate of flows. In this case, the switch from marking to dropping would apply to all of the flows in that aggregate, denying the benefits of ECN to the other flows in the aggregate also. At the highest level of aggregation, another form of the disabling of ECN happens even in the absence of monitoring and policing devices, when ECN-Capable RED queues switch from marking to dropping packets as an indication of congestion when the average queue size has exceeded some threshold.

If there were serious operational problems with routers inappropriately erasing the CE bit in packet headers, one potential fix would be to include a one-bit ECN nonce in packet headers, and for routers to erase the nonce when they set the CE bit [[SCWA99](#)]. Routers that erased the CE bit would be unable to consistently reconstruct the original nonce, and thus repeated erasure of the CE bit would be detected by the end-nodes. (This could in fact be done without adding any extra bits for ECN in the IP header, by using the ECN codepoints (ECT=1, CE=0) and (ECT=0, CE=1) as the two values for the nonce, and by defining the codepoint (ECT=0, CE=1) to mean exactly the same as the codepoint (ECT=1, CE=0).) However, at this point the potential danger of misbehaving routers does not seem of sufficient concern to warrant this additional complication of adding an ECN nonce to protect against the erasure of the CE bit.

An ECN nonce would also address the problem of misbehaving transport receivers lying to the transport sender about whether or not the CE bit was set in a packet. However, another possibility is for the data sender to test for a misbehaving receiver directly, by occasionally sending a data packet with ECT and CE set, to see if the receiver reports receiving the CE bit. Of course, if these packets encountered congestion in the network, the TCP sender would not receive this indication of congestion, so setting the ECT and CE bits at the sender would have to be done very sparingly. In addition, the TCP sender would have to remember which packets were sent with the ECT and CE bits set, so that it doesn't react to them as if there was congestion in the network. We believe that further research is needed on possible transport-based mechanisms for verifying that the transport receiver does not lie to the transport sender about the receipt of congestion indications.

11. Evaluations of ECN

This section discusses some of the related work evaluating the use of ECN. The ECN Web Page [[ECN](#)] has pointers to other papers, as well as to implementations of ECN.

[Floyd94] considers the advantages and drawbacks of adding ECN to the TCP/IP architecture. As shown in the simulation-based comparisons, one advantage of ECN is to avoid unnecessary packet drops for short or delay-sensitive TCP connections. A second advantage of ECN is in avoiding some unnecessary retransmit timeouts in TCP. This paper discusses in detail the integration of ECN into TCP's congestion control mechanisms. The possible disadvantages of ECN discussed in the paper are that a non-compliant TCP connection could falsely advertise itself as ECN-capable, and that a TCP ACK packet carrying an ECN-Echo message could itself be dropped in the network. The first of these two issues is discussed in the appendix of this document, and the second is addressed by the addition of the CWR flag in the TCP header.

Experimental evaluations of ECN include [[RFC2884](#),[K98](#)]. The conclusions of [[K98](#)] and [[RFC2884](#)] are that ECN TCP gets moderately better throughput than non-ECN TCP; that ECN TCP flows are fair towards non-ECN TCP flows; and that ECN TCP is robust with two-way traffic (with congestion in both directions) and with multiple congested gateways. Experiments with many short web transfers show that, while most of the short connections have similar transfer times with or without ECN, a small percentage of the short connections have very long transfer times for the non-ECN experiments as compared to the ECN experiments.

12. Summary of changes required in IP and TCP

This document specified two bits in the IP header, the ECN-Capable Transport (ECT) bit and the Congestion Experienced (CE) bit, to be used for ECN. The ECT bit set to "0" indicates that the transport protocol will ignore the CE bit. This is the default value for the ECT bit. The ECT bit set to "1" indicates that the transport protocol is willing and able to participate in ECN.

The default value for the CE bit is "0". The router sets the CE bit to "1" to indicate congestion to the end nodes. The CE bit in a packet header MUST NOT be reset by a router from "1" to "0".

When viewed in terms of code points, this document has defined three code points for the ECN field, for "not ECT" (ECT=0, CE=0), "ECT but not CE" (ECT=1, CE=0), and "ECT and CE" (ECT=1, CE=1). The code point of (ECT=0, CE=1) is not defined in this document. One

possibility would be for this code point to be used, some time in the future, for some other function for non-ECN-capable packets. A second possibility would be for this code point to be used as an ECN nonce, as described earlier in the paper. A third possibility would be for the code point (ECT=0, CE=1) to be used to indicate that the packet is ECN-capable for an alternate semantics for the Congestion Experienced indication. However, at this time the code point (ECT=0, CE=1) remains undefined.

TCP requires three changes for ECN, a setup phase and two new flags in the TCP header. The ECN-Echo flag is used by the data receiver to inform the data sender of a received CE packet. The Congestion Window Reduced (CWR) flag is used by the data sender to inform the data receiver that the congestion window has been reduced.

When ECN (Explicit Congestion Notification [[RFC2481](#)]) is used, it is required that congestion indications generated within an IP tunnel not be lost at the tunnel egress. We specified a minor modification to the IP protocol's handling of the ECN field during encapsulation and de-encapsulation to allow flows that will undergo IP tunneling to use ECN.

Two options for ECN in tunnels were specified:

- 1) A limited-functionality option that does not use ECN inside the IP tunnel, by turning the ECT bit in the outer header off, and not altering the inner header at the time of decapsulation.
- 2) The full-functionality option, which copies the ECT bit of the inner header to the encapsulating header. At decapsulation, if the ECT bit is set in the inner header, the CE bit on the outer header is ORed with the CE bit of the inner header to update the CE bit of the packet.

All IP tunnels MUST implement one of the two alternative approaches described above. For IPsec tunnels, this document also defines an optional IPsec SA attribute that enables negotiation of ECN usage within IPsec tunnels and an optional field in the Security Association Database to indicate whether ECN is permitted in tunnel mode on a SA.

This document is intended to obsolete [RFC 2481](#), "A Proposal to add Explicit Congestion Notification (ECN) to IP", which defined ECN as an Experimental Protocol for the Internet Community, as well as to obsolete three subsequent internet-drafts on ECN, "IPsec Interactions with ECN", "ECN Interactions with IP Tunnels", and "TCP with ECN: The Treatment of Retransmitted Data Packets". This document is intended largely to merge the earlier documents all into a single document, for greater clarity, in preparation to becoming a Proposed Standard. The rest of this section describes the relationship between this

document and its predecessors.

[RFC 2481](#) included a brief discussion of the use of ECN with encapsulated packets, and noted that for the IPsec specifications at the time (January 1999), flows could not safely use ECN if they were to traverse IPsec tunnels. [RFC 2481](#) also described the changes that could be made to IPsec tunnel specifications to make them compatible with ECN. "IPsec Interactions with ECN" outlined these changes to IPsec tunnels in detail, and included an extensive discussion of the security implications of ECN (now included as Sections [18](#) and [19](#) of this document). The draft of "ECN Interactions with IP Tunnels" extended the discussion of IPsec tunnels to include all IP tunnels. Because older IP tunnels are not compatible with a flow's use of ECN, the deployment of ECN in the Internet will create strong pressure for older IP tunnels to be updated to an ECN-compatible version, using either the limited-functionality or the full-functionality option.

This document does not address the issue of including ECN in non-IP tunnels such as MPLS, GRE, L2TP, or PPTP. An earlier preliminary document about adding ECN support to MPLS has since expired.

This document expands on one area not addressed in [RFC 2481](#), the use of ECN with retransmitted data packets. That is, this document includes the material from "TCP with ECN: The Treatment of Retransmitted Data Packets" specifying that the ECT bit should not be set on retransmitted data packets. The motivation for this additional specification is to eliminate a possible avenue for denial-of-service attacks on an existing TCP connection. Some prior deployments of ECN-capable TCP might not conform to the (new) requirement not to set the ECT bit on retransmitted packets; we do not believe this will cause significant problems in practice.

This document also expands on the specification of the use of SYN packets for the negotiation of ECN, and specifies some optional behavior for this. In particular, the document allows a TCP host to send a non-ECN-setup SYN packet after sending a failed ECN-setup SYN packet, and precisely specifies the required behavior when both ECN-setup SYN packets and non-ECN-setup SYN packets are sent in the same connection. While some prior deployments of ECN-capable TCP might not conform to the requirements specified in this document, we do not believe that this will lead to any performance or compatibility problems for TCP connections with a combination of TCP implementations at the endpoints.

[13.](#) Conclusions

Given the current effort to implement AQM, we believe this is the right time to deploy congestion avoidance mechanisms that do not

depend on packet drops alone. With the increased deployment of applications and transports sensitive to the delay and loss of a single packet (e.g., realtime traffic, short web transfers), depending on packet loss as a normal congestion notification mechanism appears to be insufficient (or at the very least, non-optimal).

We examined the consequence of modifications of the ECN field within the network, analyzing all the opportunities for an adversary to change the ECN field. In many cases, the change to the ECN field is no worse than dropping a packet. However, we noted that some changes have the more serious consequence of subverting end-to-end congestion control. However, we point out that even then the potential damage is limited, and is similar to the threat posed by end-systems intentionally failing to cooperate with end-to-end congestion control.

14. Acknowledgements

Many people have made contributions to this work and this document, including many that we have not managed to directly acknowledge in this document. In addition, we would like to thank Kenjiro Cho for the proposal for the TCP mechanism for negotiating ECN-Capability, Kevin Fall for the proposal of the CWR bit, Steve Blake for material on IPv4 Header Checksum Recalculation, Jamal Hadi-Salim for discussions of ECN issues, and Steve Bellovin, Jim Bound, Brian Carpenter, Paul Ferguson, Stephen Kent, Greg Minshall, and Vern Paxson for discussions of security issues. We also thank the Internet End-to-End Research Group for ongoing discussions of these issues.

Email discussions with a number of people, including Alexey Kuznetsov, Jamal Hadi-Salim, and Venkat Venkatsubra, have addressed the issues raised by non-conformant equipment in the Internet that does not respond to TCP SYN packets with the ECE and CWR flags set. We thank Mark Handley, Jitendra Padhye, and others for contributions to the TCP initialization procedures.

The discussion of ECN and IP tunnel considerations draws heavily on related discussions and documents from the Differentiated Services Working Group. We thank Tabassum Bint Haque from Dhaka, Bangladesh, for feedback on IP tunnels. We thank Derrell Piper and Kero Tivinen for proposing modifications to [RFC 2407](#) that improve the usability of negotiating the ECN Tunnel SA attribute.

15. References

[AH] Kent, S. and R. Atkinson, "IP Authentication Header", [RFC 2402](#), November 1998.

[B97] Bradner, S., "Key words for use in RFCs to Indicate Requirement

Levels", [BCP 14](#), [RFC 2119](#), March 1997.

[ECN] "The ECN Web Page", URL "<http://www.aciri.org/floyd/ecn.html>".

[ESP] Kent, S. and R. Atkinson, "IP Encapsulating Security Payload", [RFC 2406](#), November 1998.

[FJ93] Floyd, S., and Jacobson, V., "Random Early Detection gateways for Congestion Avoidance", IEEE/ACM Transactions on Networking, V.1 N.4, August 1993, p. 397-413.

[Floyd94] Floyd, S., "TCP and Explicit Congestion Notification", ACM Computer Communication Review, V. 24 N. 5, October 1994, p. 10-23.

[Floyd98] Floyd, S., "The ECN Validation Test in the NS Simulator", URL "<http://www-mash.cs.berkeley.edu/ns/>", test tcl/test/test-all-ecn.

[FF99] Floyd, S., and Fall, K., "Promoting the Use of End-to-End Congestion Control in the Internet", IEEE/ACM Transactions on Networking, August 1999.

[FRED] Lin, D., and Morris, R., "Dynamics of Random Early Detection", SIGCOMM '97, September 1997.

[GRE] S. Hanks, T. Li, D. Farinacci, and P. Traina, Generic Routing Encapsulation (GRE), [RFC 1701](#), October 1994.

[Jacobson88] V. Jacobson, "Congestion Avoidance and Control", Proc. ACM SIGCOMM '88, pp. 314-329.

[Jacobson90] V. Jacobson, "Modified TCP Congestion Avoidance Algorithm", Message to end2end-interest mailing list, April 1990. URL "<ftp://ftp.ee.lbl.gov/email/vanj.90apr30.txt>".

[K98] Krishnan, H., "Analyzing Explicit Congestion Notification (ECN) benefits for TCP", Master's thesis, UCLA, 1998, URL "http://www.cs.ucla.edu/~hari/software/ecn/ecn_report.ps.gz".

[L2TP] W. Townsley, A. Valencia, A. Rubens, G. Pall, G. Zorn, and B. Palter Layer Two Tunneling Protocol "L2TP", [RFC 2661](#), August 1999.

[MJV96] S. McCanne, V. Jacobson, and M. Vetterli, "Receiver-driven Layered Multicast", SIGCOMM '96, August 1996, pp. 117-130.

[MPLS] D. Awduche, J. Malcolm, J. Agogbua, M. O'Dell, J. McManus, Requirements for Traffic Engineering Over MPLS, [RFC 2702](#), September 1999.

[PPTP] Hamzeh, K., Pall, G., Verthein, W., Taarud, J., Little, W. and G. Zorn, "Point-to-Point Tunneling Protocol (PPTP)", [RFC 2637](#), July 1999.

[RFC791] Postel, J., "Internet Protocol", STD 5, [RFC 791](#), September 1981.

[RFC793] Postel, J., "Transmission Control Protocol", STD 7, [RFC 793](#), September 1981.

[RFC1141] Mallory, T. and A. Kullberg, "Incremental Updating of the Internet Checksum", [RFC 1141](#), January 1990.

[RFC1349] Almquist, P., "Type of Service in the Internet Protocol Suite", [RFC 1349](#), July 1992.

[RFC1455] Eastlake, D., "Physical Link Security Type of Service", [RFC 1455](#), May 1993.

[RFC1701] Hanks, S., Li, T., Farinacci, D., and P. Traina, Generic Routing Encapsulation (GRE), [RFC 1701](#), October 1994.

[RFC1702] Hanks, S., Li, T., Farinacci, D., and P. Traina, Generic Routing Encapsulation over IPv4 networks, [RFC 1702](#), October 1994.

[RFC2003] Perkins, C., IP Encapsulation within IP, [RFC 2003](#), October 1996.

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

[RFC2309] Braden, B., et al., "Recommendations on Queue Management and Congestion Avoidance in the Internet", [RFC 2309](#), April 1998.

[RFC 2401] S. Kent and R. Atkinson, Security Architecture for the Internet Protocol, [RFC 2401](#), November 1998.

[RFC2407] D. Piper, The Internet IP Security Domain of Interpretation for ISAKMP, [RFC 2407](#), November 1998.

[RFC2408] D. Maughan, M. Schertler, M. Schneider, and J. Turner, Internet Security Association and Key Management Protocol (ISAKMP), [RFC 2409](#), November 1998.

[RFC2409] D. Harkins and D. Carrel, The Internet Key Exchange (IKE), [RFC 2409](#), November 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] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss, An Architecture for Differentiated Services, [RFC 2475](#), December 1998.

[RFC2481] K. Ramakrishnan and S. Floyd, A Proposal to add Explicit Congestion Notification (ECN) to IP, [RFC 2481](#), January 1999.

[RFC2581] M. Allman, V. Paxson, W. Stevens, "TCP Congestion Control", [RFC 2581](#), April 1999.

[RFC2884] Jamal Hadi Salim and Uvaiz Ahmed, "Performance Evaluation of Explicit Congestion Notification (ECN) in IP Networks", [RFC 2884](#), July 2000.

[RFC2983] D. Black, "Differentiated Services and Tunnels", [RFC2983](#), October 2000.

[RFC2780] S. Bradner and V. Paxson, IANA Allocation Guidelines For Values In the Internet Protocol and Related Headers, [RFC 2780](#), March 2000.

[RFD99] Ramakrishnan, Floyd, S., and Davie, B., A Proposal to Incorporate ECN in MPLS, work in progress, June 1999. URL "<http://www.aciri.org/floyd/papers/draft-ietf-mpls-ecn-00.txt>".

[RJ90] K. K. Ramakrishnan and Raj Jain, "A Binary Feedback Scheme for Congestion Avoidance in Computer Networks", ACM Transactions on Computer Systems, Vol.8, No.2, pp. 158-181, May 1990.

[SCWA99] Stefan Savage, Neal Cardwell, David Wetherall, and Tom Anderson, TCP Congestion Control with a Misbehaving Receiver, ACM Computer Communications Review, October 1999.

16. Security Considerations

Security considerations have been discussed in Sections [7](#) and [8](#).

17. IPv4 Header Checksum Recalculation

IPv4 header checksum recalculation is an issue with some high-end router architectures using an output-buffered switch, since most if not all of the header manipulation is performed on the input side of the switch, while the ECN decision would need to be made local to the output buffer. This is not an issue for IPv6, since there is no IPv6 header checksum. The IPv4 TOS octet is the last byte of a 16-bit

half-word.

[RFC 1141](#) [[RFC1141](#)] discusses the incremental updating of the IPv4 checksum after the TTL field is decremented. The incremental updating of the IPv4 checksum after the CE bit was set would work as follows: Let HC be the original header checksum, and let HC' be the new header checksum after the CE bit has been set. Then for header checksums calculated with one's complement subtraction, HC' would be recalculated as follows:

$$\text{HC}' = \begin{cases} \text{HC} - 1 & \text{HC} > 1 \\ 0x0000 & \text{HC} = 1 \end{cases}$$

For header checksums calculated on two's complement machines, HC' would be recalculated as follows after the CE bit was set:

$$\text{HC}' = \begin{cases} \text{HC} - 1 & \text{HC} > 0 \\ 0xFFFF & \text{HC} = 0 \end{cases}$$

[18.](#) Possible Changes to the ECN Field in the Network

This section discusses in detail possible changes to the ECN field in the network, such as falsely reporting congestion, disabling ECN-Capability for an individual packet, erasing the ECN congestion indication, or falsely indicating ECN-Capability. We represent the ECN bits in the IP header by the tuple (ECT bit, CE bit).

[18.1.](#) Possible Changes to the IP Header

[18.1.1.](#) Erasing the Congestion Indication

First, we consider the changes that a router could make that would result in effectively erasing the congestion indication after it had been set by a router upstream. The convention followed is:
(ECT, CE) of received packet -> (ECT, CE) of packet transmitted.

(1, 1) -> (1, 0): erase only the CE bit that was set.
(1, 1) -> (0, 0): erase both the ECT bit and the CE bit.
(1, 1) -> (0, 1): erase the ECT bit

The first change turns off the CE bit after it has been set by some upstream router along the path. The consequence for the upstream router is that there is a potential for congestion to build for a time, because the congestion indication does not reach the source. However, the packet would be received and acknowledged.

The potential effect of erasing the congestion indication is complex, and is discussed in depth in [Section 19](#) below. Note that the effect of erasing the congestion indication is different from dropping a packet in the network. When a data packet is dropped, the drop is detected by the TCP sender, and interpreted as an indication of congestion. Similarly, if a sufficient number of consecutive acknowledgement packets are dropped, causing the cumulative acknowledgement field not to be advanced at the sender, the sender is limited by the congestion window from sending additional packets, and ultimately the retransmit timer expires.

In contrast, a systematic erasure of the CE bit by a downstream router can have the effect of causing a queue buildup at an upstream router, including the possible loss of packets due to buffer overflow. There is a potential of unfairness in that another flow that goes through the congested router could react to the CE bit set while the flow that has the CE bit erased could see better performance. The limitations on this potential unfairness are discussed in more detail in [Section 19](#) below.

The second change is to turn off both the ECT and the CE bits, thus erasing the congestion indication and disabling ECN-Capability at the same time. The third change turns off only the ECT bit, disabling ECN-Capability.

Within an IP tunnel using the full-functionality option, the third change would not erase the congestion indication, but would only disable ECN-Capability for that packet within the rest of the tunnel. However, when performed outside of an IP tunnel, the third change would also effectively erase the congestion indication, because an ECN field of (0, 1) is undefined.

The 'erasure' of the congestion indication is only effective if the packet does not end up being marked or dropped again by a downstream router. With the first change, the packet remains ECN-Capable, and could be either marked or dropped by a downstream router as an indication of congestion. With the second and third changes, the packet is no longer ECN-capable, and can therefore be dropped but not marked by a downstream router as an indication of congestion.

[18.1.2.](#) Falsely Reporting Congestion

(1, 0) -> (1, 1)

This change is to set the CE bit when the ECT bit was already set, even though there was no congestion. This change does not affect the treatment of that packet along the rest of the path. In particular, a router does not examine the CE bit in deciding whether to drop or

mark an arriving packet.

However, this could result in the application unnecessarily invoking end-to-end congestion control, and reducing its arrival rate. By itself, this is no worse (for the application or for the network) than if the tampering router had actually dropped the packet.

18.1.3. Disabling ECN-Capability

(1, 0) -> (0, *)

This change is to turn off the ECT bit of a packet that does not have the CE bit set. ([Section 18.1.1](#) discussed the case of turning off the ECT bit of a packet that does have the CE bit set.) This means that if the packet later encounters congestion (e.g., by arriving to a RED queue with a moderate average queue size), it will be dropped instead of being marked. By itself, this is no worse (for the application) than if the tampering router had actually dropped the packet. The saving grace in this particular case is that there is no congested router upstream expecting a reaction from setting the CE bit.

18.1.4. Falsely Indicating ECN-Capability

This change would incorrectly label a packet as ECN-Capable. The packet may have been sent either by an ECN-Capable transport or a transport that is not ECN-Capable.

(0, *) -> (1, 0);
(0, *) -> (1, 1);

If the packet later encounters moderate congestion at an ECN-Capable router, the router could set the CE bit instead of dropping the packet. If the transport protocol in fact is not ECN-Capable, then the transport will never receive this indication of congestion, and will not reduce its sending rate in response. The potential consequences of falsely indicating ECN-capability are discussed further in [Section 19](#) below.

If the packet never later encounters congestion at an ECN-Capable router, then the first of these two changes would have no effect. The second change, however, would have the effect of giving false reports of congestion to a monitoring device along the path. If the transport protocol is ECN-Capable, then the second of these two changes (when, for example, (0,0) was changed to (1,1)) could also have an effect at the transport level, by combining falsely indicating ECN-Capability with falsely reporting congestion. For an ECN-capable transport, this would cause the transport to unnecessarily react to congestion. In this particular case, the router that is incorrectly changing the ECN field could have dropped the packet.

Thus for this case of an ECN-capable transport, the consequence of this change to the ECN field is no worse than dropping the packet.

18.1.5. Changes with No Functional Effect

$(0, *) \rightarrow (0, *)$

The CE bit is ignored in a packet that does not have the ECT bit set. Thus, this change would have no effect, in terms of ECN.

18.2. Information carried in the Transport Header

For TCP, an ECN-capable TCP receiver informs its TCP peer that it is ECN-capable at the TCP level, using information in the TCP header at the time the connection is setup. This document does not consider potential dangers introduced by changes in the transport header within the network. In the case of IPsec tunnels, the IPsec tunnel protects the transport header.

18.3. Split Paths

In some cases, a malicious or broken router might have access to only a subset of the packets from a flow. The question is as follows: can this router, by altering the ECN field in this subset of the packets, do more damage to that flow than if it had simply dropped that set of packets?

We will classify the packets in the flow as A packets and B packets, and assume that the adversary only has access to A packets. Assume that the adversary is subverting end-to-end congestion control along the path traveled by A packets only, by either falsely indicating ECN-Capability upstream of the point where congestion occurs, or erasing the congestion indication downstream. Consider also that there exists a monitoring device that sees both the A and B packets, and will "punish" both the A and B packets if the total flow is determined not to be properly responding to indications of congestion. Another key characteristic that we believe is likely to be true is that the monitoring device, before "punishing" the A&B flow, will first drop packets instead of setting the CE bit, and will drop arriving packets of that flow that already have the ECT and CE bits set. If the end nodes are in fact using end-to-end congestion control, they will see all of the indications of congestion seen by the monitoring device, and will begin to respond to these indications of congestion. Thus, the monitoring device is successful in providing the indications to the flow at an early stage.

It is true that the adversary that has access only to the A packets might, by subverting ECN-based congestion control, be able to deny

the benefits of ECN to the other packets in the A&B aggregate. While this is unfortunate, this is not a reason to disable ECN within an IPsec tunnel.

A variant of falsely reporting congestion occurs when there are two adversaries along a path, where the first adversary falsely reports congestion, and the second adversary 'erases' those reports. (Unlike packet drops, ECN congestion reports can be 'reversed' later in the network by a malicious or broken router.) While this would be transparent to the end node, it is possible that a monitoring device between the first and second adversaries would see the false indications of congestion. Keep in mind our recommendation in this document, that before 'punishing' a flow for not responding appropriately to congestion, the router will first switch to dropping rather than marking as an indication of congestion, for that flow. When this includes dropping arriving packets from that flow that have the CE bit set, this ensures that these indications of congestion are being seen by the end nodes. Thus, there is no additional harm that we are able to postulate as a result of multiple conflicting adversaries.

19. Implications of Subverting End-to-End Congestion Control

This section focuses on the potential repercussions of subverting end-to-end congestion control by either falsely indicating ECN-Capability, or by erasing the congestion indication in ECN (the CE-bit). Subverting end-to-end congestion control by either of these two methods can have consequences both for the application and for the network. We discuss these separately below.

The first method to subvert end-to-end congestion control, that of falsely indicating ECN-Capability, effectively subverts end-to-end congestion control only if the packet later encounters congestion that results in the setting of the CE bit. In this case, the transport protocol (which may not be ECN-capable) does not receive the indication of congestion from these downstream congested routers.

The second method to subvert end-to-end congestion control, 'erasing' the (set) CE bit in a packet, effectively subverts end-to-end congestion control only when the CE bit in the packet was set earlier by a congested router. In this case, the transport protocol does not receive the indication of congestion from the upstream congested routers.

Either of these two methods of subverting end-to-end congestion control can potentially introduce more damage to the network (and possibly to the flow itself) than if the adversary had simply dropped packets from that flow. However, as we discuss later in this section and in [Section 7](#), this potential damage is limited.

19.1. Implications for the Network and for Competing Flows

The CE bit of the ECN field is only used by routers as an indication of congestion during periods of **moderate** congestion. ECN-capable routers should drop rather than mark packets during heavy congestion even if the router's queue is not yet full. For example, for routers using active queue management based on RED, the router should drop rather than mark packets that arrive while the average queue sizes exceed the RED queue's maximum threshold.

One consequence for the network of subverting end-to-end congestion control is that flows that do not receive the congestion indications from the network might increase their sending rate until they drive the network into heavier congestion. Then, the congested router could begin to drop rather than mark arriving packets. For flows that are not isolated by some form of per-flow scheduling or other per-flow mechanisms, but are instead aggregated with other flows in a single queue in an undifferentiated fashion, this packet-dropping at the congested router would apply to all flows that share that queue. Thus, the consequences would be to increase the level of congestion in the network.

In some cases, the increase in the level of congestion will lead to a substantial buffer buildup at the congested queue that will be sufficient to drive the congested queue from the packet-marking to the packet-dropping regime. This transition could occur either because of buffer overflow, or because of the active queue management policy described above that drops packets when the average queue is above RED's maximum threshold. At this point, all flows, including the subverted flow, will begin to see packet drops instead of packet marks, and a malicious or broken router will no longer be able to 'erase' these indications of congestion in the network. If the end nodes are deploying appropriate end-to-end congestion control, then the subverted flow will reduce its arrival rate in response to congestion. When the level of congestion is sufficiently reduced, the congested queue can return from the packet-dropping regime to the packet-marking regime. The steady-state pattern could be one of the congested queue oscillating between these two regimes.

In other cases, the consequences of subverting end-to-end congestion control will not be severe enough to drive the congested link into sufficiently-heavy congestion that packets are dropped instead of being marked. In this case, the implications for competing flows in the network will be a slightly-increased rate of packet marking or dropping, and a corresponding decrease in the bandwidth available to those flows. This can be a stable state if the arrival rate of the subverted flow is sufficiently small, relative to the link bandwidth, that the average queue size at the congested router remains under

control. In particular, the subverted flow could have a limited bandwidth demand on the link at this router, while still getting more than its "fair" share of the link. This limited demand could be due to a limited demand from the data source; a limitation from the TCP advertised window; a lower-bandwidth access pipe; or other factors. Thus the subversion of ECN-based congestion control can still lead to unfairness, which we believe is appropriate to note here.

The threat to the network posed by the subversion of ECN-based congestion control in the network is essentially the same as the threat posed by an end-system that intentionally fails to cooperate with end-to-end congestion control. The deployment of mechanisms in routers to address this threat is an open research question, and is discussed further in [Section 10](#).

Let us take the example described in [Section 18.1.1](#), where the CE bit that was set in a packet is erased: $\{(1, 1) \rightarrow (1, 0)\}$. The consequence for the congested upstream router that set the CE bit is that this congestion indication does not reach the end nodes for that flow. The source (even one which is completely cooperative and not malicious) is thus allowed to continue to increase its sending rate (if it is a TCP flow, by increasing its congestion window). The flow potentially achieves better throughput than the other flows that also share the congested router, especially if there are no policing mechanisms or per-flow queueing mechanisms at that router. Consider the behavior of the other flows, especially if they are cooperative: that is, the flows that do not experience subverted end-to-end congestion control. They are likely to reduce their load (e.g., by reducing their window size) on the congested router, thus benefiting our subverted flow. This results in unfairness. As we discussed above, this unfairness could either be transient (because the congested queue is driven into the packet-marking regime), oscillatory (because the congested queue oscillates between the packet marking and the packet dropping regime), or more moderate but a persistent stable state (because the congested queue is never driven to the packet dropping regime).

The results would be similar if the subverted flow was intentionally avoiding end-to-end congestion control. One difference is that a flow that is intentionally avoiding end-to-end congestion control at the end nodes can avoid end-to-end congestion control even when the congested queue is in packet-dropping mode, by refusing to reduce its sending rate in response to packet drops in the network. Thus the problems for the network from the subversion of ECN-based congestion control are less severe than the problems caused by the intentional avoidance of end-to-end congestion control in the end nodes. It is also the case that it is considerably more difficult to control the behavior of the end nodes than it is to control the behavior of the

infrastructure itself. This is not to say that the problems for the network posed by the network's subversion of ECN-based congestion control are small; just that they are dwarfed by the problems for the network posed by the subversion of either ECN-based or other currently known packet-based congestion control mechanisms by the end nodes.

19.2. Implications for the Subverted Flow

When a source indicates that it is ECN-capable, there is an expectation that the routers in the network that are capable of participating in ECN will use the CE bit for indication of congestion. There is the potential benefit of using ECN in reducing the amount of packet loss (in addition to the reduced queueing delays because of active queue management policies). When the packet flows through a tunnel where the nodes that the tunneled packets traverse are untrusted in some way, the expectation is that IPsec will protect the flow from subversion that results in undesirable consequences.

In many cases, a subverted flow will benefit from the subversion of end-to-end congestion control for that flow in the network, by receiving more bandwidth than it would have otherwise, relative to competing non-subverted flows. If the congested queue reaches the packet-dropping stage, then the subversion of end-to-end congestion control might or might not be of overall benefit to the subverted flow, depending on that flow's relative tradeoffs between throughput, loss, and delay.

One form of subverting end-to-end congestion control is to falsely indicate ECN-capability by setting the ECT bit. This has the consequence of downstream congested routers setting the CE bit in vain. However, as we describe in the section below, if the ECT bit is changed in the IPsec tunnel, this can be detected at the egress point of the tunnel.

The second form of subverting end-to-end congestion control is to erase the congestion indication, either by erasing the CE bit directly, or by erasing the ECT bit when the CE bit is already set. In this case, it is the upstream congested routers that set the CE bit in vain.

If the ECT bit is erased within an IP tunnel, then this can be detected at the egress point of the tunnel. If the CE bit is set upstream of the IP tunnel, then any erasure of the outer header's CE bit within the tunnel will have no effect because the inner header preserves the set value of the CE bit. However, if the CE bit is set within the tunnel, and erased either within or downstream of the tunnel, this is not necessarily detected at the egress point of the

tunnel.

With this subversion of end-to-end congestion control, an end-system transport does not respond to the congestion indication. Along with the increased unfairness for the non-subverted flows described in the previous section, the congested router's queue could continue to build, resulting in packet loss at the congested router - which is a means for indicating congestion to the transport in any case. In the interim, the flow might experience higher queueing delays, possibly along with an increased bandwidth relative to other non-subverted flows. But transports do not inherently make assumptions of consistently experiencing carefully managed queueing in the path. We believe that these forms of subverting end-to-end congestion control are no worse for the subverted flow than if the adversary had simply dropped the packets of that flow itself.

[19.3.](#) Non-ECN-Based Methods of Subverting End-to-end Congestion Control

We have shown that, in many cases, a malicious or broken router that is able to change the bits in the ECN field can do no more damage than if it had simply dropped the packet in question. However, this is not true in all cases, in particular in the cases where the broken router subverted end-to-end congestion control by either falsely indicating ECN-Capability or by erasing the ECN congestion indication (in the CE-bit). While there are many ways that a router can harm a flow by dropping packets, a router cannot subvert end-to-end congestion control by dropping packets. As an example, a router cannot subvert TCP congestion control by dropping data packets, acknowledgment packets, or control packets.

Even though packet-dropping cannot be used to subvert end-to-end congestion control, there *are* non-ECN-based methods for subverting end-to-end congestion control that a broken or malicious router could use. For example, a broken router could duplicate data packets, thus effectively negating the effects of end-to-end congestion control along some portion of the path. (For a router that duplicated packets within an IPsec tunnel, the security administrator can cause the duplicate packets to be discarded by configuring anti-replay protection for the tunnel.) This duplication of packets within the network would have similar implications for the network and for the subverted flow as those described in Sections [18.1.1](#) and [18.1.4](#) above.

[20.](#) The motivation for the ECT bit.

The need for the ECT bit is motivated by the fact that ECN will be deployed incrementally in an Internet where some transport protocols and routers understand ECN and some do not. With the ECT bit, the router can drop packets from flows that are not ECN-capable, but can

instead set the CE bit in packets that *are* ECN-capable. Because the ECT bit allows an end node to have the CE bit set in a packet *instead* of having the packet dropped, an end node might have some incentive to deploy ECN.

If there was no ECT indication, then the router would have to set the CE bit for packets from both ECN-capable and non-ECN-capable flows. In this case, there would be no incentive for end-nodes to deploy ECN, and no viable path of incremental deployment from a non-ECN world to an ECN-capable world. Consider the first stages of such an incremental deployment, where a subset of the flows are ECN-capable. At the onset of congestion, when the packet dropping/marketing rate would be low, routers would only set CE bits, rather than dropping packets. However, only those flows that are ECN-capable would understand and respond to CE packets. The result is that the ECN-capable flows would back off, and the non-ECN-capable flows would be unaware of the ECN signals and would continue to open their congestion windows.

In this case, there are two possible outcomes: (1) the ECN-capable flows back off, the non-ECN-capable flows get all of the bandwidth, and congestion remains mild, or (2) the ECN-capable flows back off, the non-ECN-capable flows don't, and congestion increases until the router transitions from setting the CE bit to dropping packets. While this second outcome evens out the fairness, the ECN-capable flows would still receive little benefit from being ECN-capable, because the increased congestion would drive the router to packet-dropping behavior.

A flow that advertised itself as ECN-Capable but does not respond to CE bits is functionally equivalent to a flow that turns off congestion control, as discussed earlier in this document.

Thus, in a world when a subset of the flows are ECN-capable, but where ECN-capable flows have no mechanism for indicating that fact to the routers, there would be less effective and less fair congestion control in the Internet, resulting in a strong incentive for end nodes not to deploy ECN.

21. Why use two bits in the IP header?

Given the need for an ECT indication in the IP header, there still remains the question of whether the ECT (ECN-Capable Transport) and CE (Congestion Experienced) indications should have been overloaded on a single bit. This overloaded-one-bit alternative, explored in [Floyd94], would have involved a single bit with two values. One value, "ECT and not CE", would represent an ECN-Capable Transport, and the other value, "CE or not ECT", would represent either

Congestion Experienced or a non-ECN-Capable transport.

One difference between the one-bit and two-bit implementations concerns packets that traverse multiple congested routers. Consider a CE packet that arrives at a second congested router, and is selected by the active queue management at that router for either marking or dropping. In the one-bit implementation, the second congested router has no choice but to drop the CE packet, because it cannot distinguish between a CE packet and a non-ECT packet. In the two-bit implementation, the second congested router has the choice of either dropping the CE packet, or of leaving it alone with the CE bit set.

Another difference between the one-bit and two-bit implementations comes from the fact that with the one-bit implementation, receivers in a single flow cannot distinguish between CE and non-ECT packets. Thus, in the one-bit implementation an ECN-capable data sender would have to unambiguously indicate to the receiver or receivers whether each packet had been sent as ECN-Capable or as non-ECN-Capable. One possibility would be for the sender to indicate in the transport header whether the packet was sent as ECN-Capable. A second possibility that would involve a functional limitation for the one-bit implementation would be for the sender to unambiguously indicate that it was going to send **all** of its packets as ECN-Capable or as non-ECN-Capable. For a multicast transport protocol, this unambiguous indication would have to be apparent to receivers joining an on-going multicast session.

Another concern that was described earlier (and recommended in this document) is that transports (particularly TCP) should not mark pure ACK packets or retransmitted packets as being ECN-Capable. A pure ACK packet from a non-ECN-capable transport could be dropped, without necessarily having an impact on the transport from a congestion control perspective (because subsequent ACKs are cumulative). An ECN-capable transport reacting to the CE bit set in a pure ACK packet by reducing the window would be at a disadvantage in comparison to a non-ECN-capable transport. For this reason (and for reasons described earlier in relation to retransmitted packets), it is desirable to have the ECN-Capable bit indication on a per-packet basis.

Another advantage of the two-bit approach is that it is somewhat more robust. The most critical issue, discussed in [Section 8](#), is that the default indication should be that of a non-ECN-Capable transport. In a two-bit implementation, this requirement for the default value simply means that the ECT bit should be 'OFF' by default. In the one-bit implementation, this means that the single overloaded bit should by default be in the "CE or not ECT" position. This is less clear and straightforward, and possibly more open to incorrect implementations either in the end nodes or in the routers.

In summary, while the one-bit implementation could be a possible implementation, it has the following significant limitations relative to the two-bit implementation. First, the one-bit implementation has more limited functionality for the treatment of CE packets at a second congested router. Second, the one-bit implementation requires either that extra information be carried in the transport header of packets from ECN-Capable flows (to convey the functionality of the second bit elsewhere, namely in the transport header), or that senders in ECN-Capable flows accept the limitation that receivers must be able to determine a priori which packets are ECN-Capable and which are not ECN-Capable. Third, the one-bit implementation is possibly more open to errors from faulty implementations that choose the wrong default value for the ECN bit. We believe that the use of the extra bit in the IP header for the ECT-bit is extremely valuable to overcome these limitations.

22. Historical definitions for the IPv4 TOS octet

[RFC 791](#) [[RFC791](#)] defined the ToS (Type of Service) octet in the IP header. In [RFC 791](#), bits 6 and 7 of the ToS octet are listed as "Reserved for Future Use", and are shown set to zero. The first two fields of the ToS octet were defined as the Precedence and Type of Service (TOS) fields.

| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 |
|------------|---|---|-----|---|---|---|---|
| PRECEDENCE | | | TOS | | | 0 | 0 |

[RFC 791](#)

[RFC 1122](#) included bits 6 and 7 in the TOS field, though it did not discuss any specific use for those two bits:

| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 |
|------------|---|---|-----|---|---|---|---|
| PRECEDENCE | | | TOS | | | | |

[RFC 1122](#)

The IPv4 TOS octet was redefined in [RFC 1349](#) [[RFC1349](#)] as follows:

| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 |
|------------|---|---|-----|---|---|-----|---|
| PRECEDENCE | | | TOS | | | MBZ | |

[RFC 1349](#)

Bit 6 in the TOS field was defined in [RFC 1349](#) for "Minimize Monetary Cost". In addition to the Precedence and Type of Service (TOS) fields, the last field, MBZ (for "must be zero") was defined as

currently unused. [RFC 1349](#) stated that "The originator of a datagram sets [the MBZ] field to zero (unless participating in an Internet protocol experiment which makes use of that bit)."

[RFC 1455](#) [[RFC 1455](#)] defined an experimental standard that used all four bits in the TOS field to request a guaranteed level of link security.

[RFC 1349](#) is obsoleted by "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers" [[RFC2474](#)], in which bits 6 and 7 of the DS field are listed as Currently Unused (CU). The first six bits of the DS field are defined as the Differentiated Services CodePoint (DSCP):

| | | | | | | | | |
|---|------|---|---|---|---|---|----|--|
| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | |
| +-----+-----+-----+-----+-----+-----+-----+-----+ | | | | | | | | |
| | DSCP | | | | | | CU | |
| +-----+-----+-----+-----+-----+-----+-----+-----+ | | | | | | | | |

[RFC 2474](#)

Because of this unstable history, the definition of the ECN field in this document cannot be guaranteed to be backwards compatible with all past uses of these two bits. The damage that could be done by a non-ECN-capable router would be to "erase" the CE bit for an ECN-capable packet that arrived at the router with the CE bit set, or set the CE bit even in the absence of congestion. This has been discussed in the section on "Non-compliance in the Network".

The damage that could be done in an ECN-capable environment by a non-ECN-capable end-node transmitting packets with the ECT bit set has been discussed in the section on "Non-compliance by the End Nodes".

AUTHORS' ADDRESSES

K. K. Ramakrishnan
TeraOptic Networks, Inc.
Phone: +1 (408) 666-8650
Email: kk@teraoptic.com

Sally Floyd
Phone: +1 (510) 666-2989
ACIRI
Email: floyd@aciri.org
URL: <http://www.aciri.org/floyd/>

David L. Black
EMC Corporation
42 South St.

Hopkinton, MA 01748
Phone: +1 (508) 435-1000 x75140
Email: black_david@emc.com

This draft was created in November 2000.
It expires May 2001.