On the State of ECN and TCP Options on the Internet*

Mirja Kühlewind[1], Sebastian Neuner[1], and Brian Trammell[2]

[1] Institute of Communication Networks and Computer Engineering (IKR)
University of Stuttgart, Germany
[2] Communication Systems Group, ETH Zürich, Switzerland

Abstract. Explicit Congestion Notification (ECN) is a TCP/IP extension that can avoid packet loss and thus improve network performance. Though standardized in 2001, it is barely used in today's Internet. This study, following on previous active measurement studies over the past decade, shows marked and continued increase in the deployment of ECN-capable servers, and usability of ECN on the majority of paths to such servers. We additionally present new measurements of ECN on IPv6, passive observation of actual ECN usage from flow data, and observations on other congestion-relevant TCP options (SACK, Timestamps and Window Scaling). We further present initial work on burst loss metrics for loss-based congestion control following from our findings.

1 Introduction

Since the initial design of TCP, there have been a number of extensions designed to improve its throughput and congestion control characteristics. Explicit Congestion Notification (ECN) is a TCP/IP extension that allows congestion signaling without packet loss. Though it has been shown to have performance benefits [1] and has been a standard since 2001 [2,3], ECN deployment lags significantly. Initial deployment problems where middleboxes cleared the ECN IP bits or even dropped packets indicating ECN-capability, as well as firewalls that would reset ECN-capable connections [4], led to mistrust of ECN.

In this work, we examine how much this situation has improved, adding another datapoint to a series of active measurements of ECN usage going back a decade. We also measured the usage of three other congestion-control-relevant TCP options: Selective Acknowledgment (SACK) [5], Timestamps (TS), and Window Scale (WS) [6]. SACK allows more precise signaling of loss, TS improves round-trip-time estimation, and WS allows a larger receiver windows.

Our measurement methodology consists of active probing of the ECN-readiness of a large set of popular web-servers (section 3.1) as well as passive measurement of ECN usage from flow data collected on a national-scale research and education network (section 3.2).

* This work is partly funded by ETICS and mPlane, FP7 research projects supported by the EU. Thanks to SWITCH for the flow data used in this study.

M. Roughan and R. Chang (Eds.) PAM 2013, LNCS 7799, pp. 135–144, 2013.

Table 1. ECN implementation status

year	OS	version
2007	Microsoft Windows	Server 2008, 7, Vista
2007	Mac OS X	10.5
2006	Cisco IOS	12.2(8)T
2001	Linux	2.4 (full support)
1999	Linux	2.3 (router support)

Table 2. History of ECN and options deployment

Reference	Date	ECN	SACK	TSOPT
Medina ea. [7]	2000	1.1%	28%	-
Medina ea. [7]	2004	2.1%	68%	30%
Langley ea. [8]	2008	1.06%	-	-
Bauer ea. [9]	2011	17.2%	-	-

Deployment of ECN and related TCP options has been periodically studied in the literature over the past decade [7,8,9]; the most relevant results for the present work are summarized in Table 2. Bauer *et al* [9] probed the same set of servers as in the present work, so these results are directly comparable. Also related are measurements on TCP extensibility, which focus on middlebox treatment of packets with TCP options. Here findings vary between 0.17% [8] and 70% [9] of hosts dropping packets with unknown options, and 4–14% of middleboxes dropping such packets [10].

We find a recent acceleration in deployment of ECN-capable servers (section 4.1) and greater ECN support on IPv6-enabled servers (section 4.2). We compare this to actual ECN usage, passively measured from flow data captured from the border of a national-scale network, and find that while ECN is more frequently deployed, it is still seldom used (section 4.3).

In section 5, we define a metric for *burst loss* taking into account the peridic probing of congestion-control algorithms, and show that different types of traffic have different burst loss characteristics. Given the continued lag of ECN usage, we advance this initial work as a way to better understand loss dynamics and its relation to application behavior. Section 6 presents our conclusions.

2 Explicit Congestion Notification (ECN): A Review

ECN allows routers using active queue management (AQM) (e.g., Random Early Detection (RED)) to mark packets in case of congestion instead of dropping them. Two bits in the IP header provide four possible marks: No-ECN (00), Congestion Experienced (CE, 11), and two codepoints for ECN-Capable Transport (ECT(0), 01; and ECT(1), 10). An ECN-capable sender sets ECT(0) or ECT(1), which can be changed to CE by a router to signal congestion.

ECN uses two additional flags in the TCP header: ECN-Echo (ECE) is set on all packets from the receiver back to the sender to signal the arrival of a CE-marked packet until the sender sets Congestion Window Reduced (CWR) to acknowledge the ECE. These flags are also used to negotiate ECN usage: a connection initiator requests ECN by setting ECE and CWR on the initial SYN, and the responder acknowledges by setting ECE on the SYN/ACK. After successfully completing the negotiation, the senders can set an ECT codepoint on all subsequent packets over the connection.