Proceedings of CECNet 2022
A.J. Tallón-Ballesteros (Ed.)

doi:10.3233/FAIA220560

Measuring the State of ECN Readiness in Mail Servers on the Internet

Chun-Xiang CHEN [a,1], Tetuya SHIGEYASU [a] and Kenichi NAGAOKA [b]
[a] *Prefectural University of Hiroshima, Minami-ku, Hiroshima-shi, 734-8555, Japan*
[b] *National Institute of Technology, Ishikawa College, Ishikawa, 929-0392, Japan*

Abstract. Congestion control on the Internet has been a challenge issue so far. Since there is no one-size-fits-all scheme for TCP congestion control, many various schemes have been developed to suit different kinds of communication environments. Among them, ECN (Explicit Congestion Notification) is the one which is able to feed the congestion state in the routers along a connection path back to the TCP transmitter. To promote the deployment of ECN and provide a concrete reference parameter to the Internet industry, in this paper, we actually measure the readiness of ECN in mail servers and analyze its deployment rate on the Internet. To clarify the characteristics of ECN in the mail servers, deployment rate of ECN is compared with that of ECN in major web servers. Based on the analysis results, we show some unique features of ECN in mail servers which are quite different from that of web servers.

Keywords. TCP congestion control, explicit congestion notification, deployment rate of ECN, email server, web server

1. Introduction

Over the last three decades, congestion control of TCP (Transmission Control Protocol) on the Internet still has been a challenge issue, especially in the super high-speed and/or ultra-dense heterogeneous networks. To solve this problem, a number of congestion control algorithms with TCP have been developed [1,2]. Due to the principles of TCP/IP (Internet Protocol), the TCP sender has no way to know about the traffic state along the connection path in real time. Therefore, all the schemes, such as the early TCP Tahoe, TCP Reno, TCP CUBIC [3] and then recent BBR (Bottleneck Bandwidth and Round-trip propagation time) [4,5], judge the congestion state in the intermediate network based on the packet loss and/or the transmission delay. However, their decisions probably diverge from the reality since traffic and connections might change rapidly within a short time.

To reflect the traffic state of a connection path in the sender, ECN was proposed and has been standardized in RFC 3168. To the authors' knowledge, ECN has been the only scheme that is able to successfully feed back an imminent congestion state along the connection path and greatly improve the throughput efficiency of the whole Internet. That is one of the reasons the authors have been committing to promoting the deployment of ECN [8]-[10]. However, the previous works have targeted only at major web servers. As an indispensable supplement to the previous works, in this paper, we conduct a set of empirical measurement and investigate the readiness of ECN in major mail servers on

[1]Corresponding Author: Chun-Xiang Chen, Prefectural University of Hiroshima, Hiroshima, Japan; E-mail:chen@pu-hiroshima.ac.jp

the Internet. Moreover, by analyzing the request-response of each SMTP (Simple Mail Transfer Protocol) connection, we clarify the deployment state of ECN on the Internet. Furthermore, we demonstrate some unique features of ECN in primary mail servers by comparing with that of leading web servers.

2. Related Works and Scope of this Work

ECN uses two bits in each of the IP and TCP headers. Its implementations have been stipulated in detail in RFC 3168. To clarify the benefit and risk of using ECN, a lot of previous works have been conducted [6]-[10].

RFC 8087 [7] expounds the potential benefits of ECN-Capable connections in terms of increased throughput and reduced delay. In [6], dynamic marking threshold algorithm for ECN was proposed. But it is limited in data center networks. In [8], the characteristics of ECN with regard to the connectivity, deployment rate were analyzed. Moreover, the latent risk of an ECN-Capable connection was evaluated quantitatively. To improve the connectivity of ECN, the sender would try to establish a not-ECN-capable connection when an ECN-Capable connection can't be established by timeout period. This process is called fallback. In [9], the effectiveness of the fallback regarding ECN was evaluated. Furthermore, features related to the deployment rate and connectivity were analyzed according to the regional distribution of the domains on the Internet.

However, all the previous works only focus on the major web servers on the Internet. As we know, email traffic based on SMTP is another traditional primary medium besides web services (based on HTTP and/or HTTPS). In recent years, the spread of SNS (Social Networking Service) has led to a decrease in the amount of email communication, but email, especially email with attachments, still plays an important role in the business world. In addition, many SNS use email addresses as authentication accounts. The number of email accounts continues to grow. Therefore, to reveal the deployment state of ECN over the Internet more precisely, measurement and analysis for major mail servers are also required. This paper aims to provide a concrete deployment rate of ECN for the Internet industry which have not been revealed so far. We analyze the readiness of ECN for major mail servers via a set of empirical measurement and clarify some unique features of ECN corresponding to mail servers.

3. Measurement Method

We firstly choose the mail servers targeted to be tested from the Internet. We use the major domains published by Alexa Global as the potential domains of mail servers [11] which are the same as the domain set used in the previous studies [8,10]. The process of our measurement is summarized as follows.

Download the Global Top one million domains from Alexa [11]; Query the DNS for the MX (Mail exchange) record of each domain. Here we use Google public DNS (8.8.8.8) to retrieve MX record and its IP (IPv4 and/or IPv6) address. For example, for domain {`gmail.com`}, we get its MX records as

```
gmail.com, gmail-smtp-in.l.google.com, alt1.gmail-smtp-in.l.google.com,
alt2.gmail-smtp-in.l.google.com,...
```

Then, we resolve their IP addresses of these MX records using the tool we developed. We get the entries of mappings as

```
gmail-smtp-in.l.google.com, 142.250.157.26, 2404:6800:4008:c13::1b
```

`alt1.gmail-smtp-in.l.google.com, 142.250.141.27, 2607:f8b0:4023:c0b::1a`

Note that, MX records and their IP addresses might be varying with the DNS server and the querying time. If a domain does not have its owe MX record, it will be excluded from the domain entries. By removing the duplicated MX records of the domains, we finally obtained 220,367 and 210,440 entries of MX records (called *targeted server lists*) on 15 August 2021 and on 8 January 2022, respectively. Tools used to retrieve MX records and IP addresses and test SMTP connections, are developed in `Python`.

The measurements are performed by using a Linux host. Its specifications are as followings; CPU: Intel Xeon E5-1630v4, OS: CentOS 8.3 with kernel 4.18.0-240, Memory: 16GB, NIC: Intel 82574L Gigabit PCI interface. Measurements were conducted at the network of Prefectural University of Hiroshima. IPv6 network is connected through IPv4 tunnel provided by Hurricane Electric Internet Service. The tool used to test the targeted mail servers is developed according to the commands of SMTP. The procedures of measurements are similar to those in the previous studies [8,10]. `tcpdump` is used to capture all of the IP packets during the measurements. A simplified processes of the measurement are as follows.

1. Disable ECN by kernel setting `net.ipv4.tcp_ecn=0`; Pick up next entry from the *targeted server lists*; Initiate an SMTP connection (TCP port 25) with this server by setting flag SYN; Issue a `helo` and then quit the connection once the connection is established.
2. Enable ECN by setting `net.ipv4.tcp_ecn=1`; Initiate an SMTP connection (TCP port 25) with the targeted server by setting flags (SYN, ECN,CWR); Issue a `helo` and then quit the connection once the connection is established.
3. Repeat the above procedures 1 and 2 until all the targeted mail servers are tested.

The next task is to carefully and patiently analyze the readiness of ECN for each connection from the packets dumped by `tcpdump`.

4. Analysis and Results

In order to analyze the features of ECN in the mail servers, we need to collect these packets belonging to the same SMTP connection from all packets dumped by `tcpdump` and to extract the readiness of each connection. To do such hard work, we developed an analytical tool in `Perl`. We omitted the detailed analysis processes here due to space limitations.

This Section only presents the analysis results with regard to ECN deployment rate (shown by β_{ECN}) that is defined as the proportion of the targeted servers that are ECN-Capable to all valid servers. In order to do a comparison of the deployment rate of ECN on the web and mail servers, measurement about the web and mail server of the same TLD (Top Level Domain) are conducted at the same time.

Figures 1 and 2 show the deployment rate of ECN (β_{ECN}) by domains (IPv4) on mail servers measured on 15 August 2021 and on 8 January 2022, respectively. To compare with web servers, deployment rate of ECN on web servers is also measured and plotted in these figures. Note that the horizontal axis of Fig. 1 is sorted in descending order of β_{ECN} for mail servers in Aug. 2021. Subsequent figures are shown in the same order.

Figures 3 and 4 illustrate the variations of deployment rate of ECN on web and mail servers over a six-month period, respectively. From Figs. 1∼4, we can see that, (1) β_{ECN} on most web servers is larger than that on mail servers. This feature is more noticeable for the domains on the right area of the figure. One exception is about domain `kr`; (2) β_{ECN} on mail servers keeps increasing (Fig. 4) while that on web servers increases slightly

(Fig. 3); (3) For domain `kr`, $\beta_{\mathtt{ECN}}$ on web servers is lower than that on mail servers (Fig. 1).

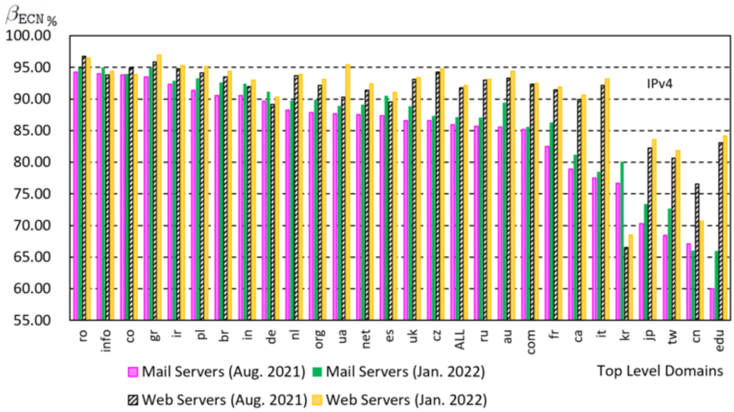

Figure 1. Deployment rate of ECN on mail servers vs. web servers (IPv4)

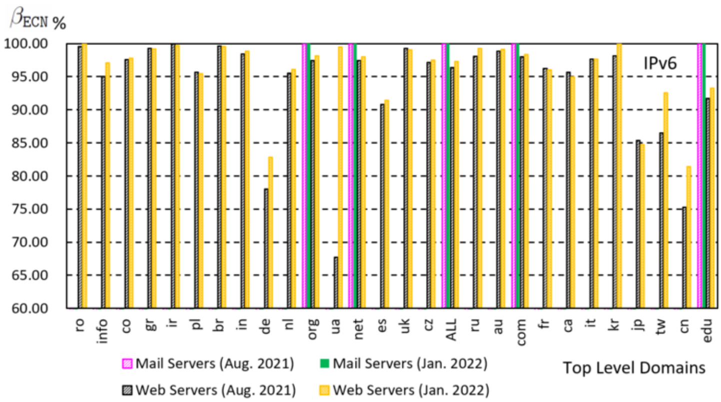

Figure 2. Deployment rate of ECN on mail servers vs. web servers (IPv6)

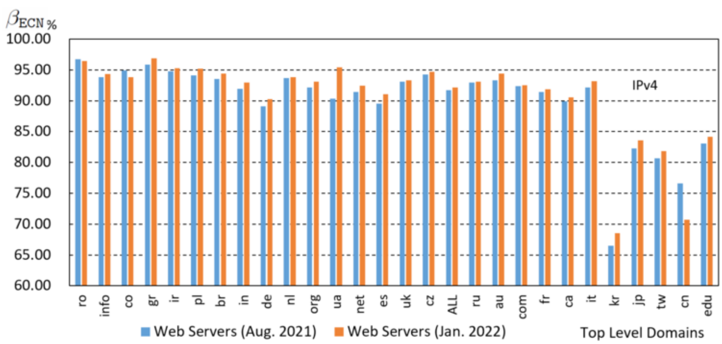

Figure 3. Deployment rate of ECN on web servers

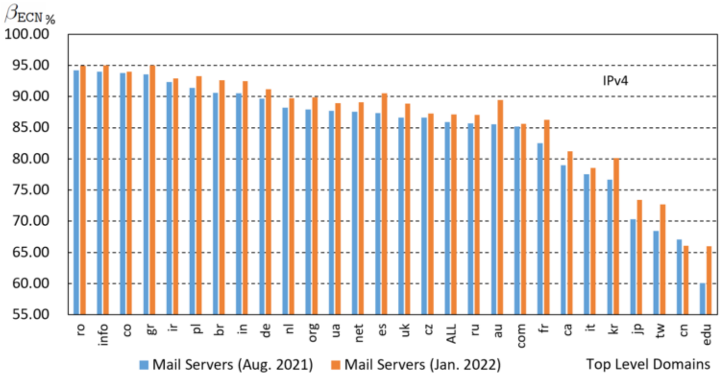

Figure 4. Deployment rate of ECN on mail servers

Deployment rates of ECN in IPv6 are illustrated in Fig. 5. It is clear that $\beta_{\mathtt{ECN}}$ on web servers keeps increasing. Especially, the range of increase on each of the domains `de`, `ua`, `tw`, and `cn` is significant. As for the deployment rate of ECN on the mail servers in IPv6, the percentage of $\beta_{\mathtt{ECN}}$ is 100% for each of domains `org`, `net`, `com` and `edu`, while it is 0% for the other domains. This means that there are very few mail servers which are running in IPv6. Furthermore, Table 1 details the number of the mail servers for those domains. Among 220,367 and 210,440 MX records of mail servers, only 16 and 14 mail servers were running in IPv6 at that time.

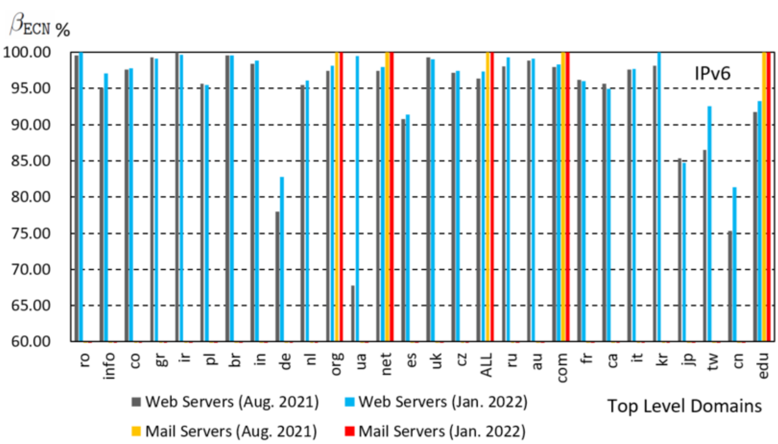

Figure 5. Deployment rate of ECN on mail and web servers (IPv6)

Table 1. Numbers of mail servers on the domains on IPv6

Time when measured	org	net	edu	com	all	Total MX records
Aug. 2021	2	1	1	12	16	220,367
Jan. 2022	1	1	1	11	14	210,440

5. Conclusion Remarks

In this paper, we measured and analyzed the state of ECN readiness in mail servers on the Internet. We also made a comparison of the deployment of ECN between mail and web servers. According to the analysis results, it is clear that (1) deployment rate of ECN in mail servers is mainly lower than that of web servers; (2) Deployment rates both on web and mail servers are increasing gradually; (3) $\beta_{\mathtt{ECN}}$ of domain `jp, cn, tw` and `kr` (in east Asia) are relatively low; (4) Unlike most web servers that work with IPv6, very few mail servers run in IPv6.

Compared to the previous studies [8,10], we also can see that deployment rates of ECN on web servers maintain stable growth over time. Therefore, it is obvious that the number of web and/or mail servers which are ECN-Capable on the Internet is increasing stably. As of Jan. 2022, deployment rates of ECN on web and mail servers at average were 92.18% and 87.16% (Fig. 2, `All`), respectively. Based on these results, we can say that ECN is almost getting ready in server side. On the other hand, considering that ECN supported servers are probably in passive mode [8], this means that the use or non use of ECN depends largely on the communication client devices. Therefore, to take the advantage of ECN, we highly recommend that the developers of OSes and/or the vendors of networking equipment should launch their products with ECN enabled at initial settings.

References

[1] R. AI-Saadi, G. Armitage, J. But, and P. Branch, "A Survey of Delay-Based and Hybrid TCP Congestion Control Algorithms," *IEEE Communi. Surverys & Tutorials*, vol.21, No.4, Feb. 2019. `DOI:10.1109/COMST.2019.2904994`

[2] H. Huang, Z. Sun, and X. Wang, "End-to-End TCP Congestion Control for Mobile Applications," *IEEE Access*, vol. 8, pp.171628-171642, Sept. 2020. `DOI:10.1109/ACCESS.2020.3024707`

[3] Ing. Marrone, Lic. Barbieri, and Mg. Robles, "TCP Performance - CUBIC, Vegas & Reno," Journal of Computer and Technology, Vol. 13, No. 01, pp.1-8, 2013.

[4] N. Cardwell, Y. Cheng, C. S. Gunn et al, "BBR: Congestion-Based Congestion Control," *ACM Queue Networks*, pp. 20–53, Sept-Oct., 2016.

[5] M. Claypool, F. Li, and J. Chung, "BBR' - An Implementation of Bottleneck Bandwidth and Round-trip Time Congestion Control for ns-3," Technical Report WPI-CS-TR-18-01, Computer Science, Worcester Polytechnic Institute, Jan. 2018.

[6] Y. Lu, X. Fan, and Lei Qian, "Dynamic ECN marking threshold algorithm for TCP congestion control in data center networks," Computer Communications, Vol. 129, pp. 197-208, 2018. `DOI:10.1016/j.comcom.2018.07.036.`

[7] G. Fairhurst and M. Welzl, "The benefits of using explicit congestion notification (ECN)," RCF 8087, March 2017. `DOI:10.17487/rfc8087`

[8] C.-X. Chen and K. Nagaoka, "Analysis of the State of ECN on the Internet," *IEICE Trans. on Inf. and syst.*, vol.E102-D, no.5, pp.910-919, May 2019. `DOI:10.1587/transinf.2018NTP0006`

[9] L. Zou, C.-X. Chen, and K. Nagaoka, "An empirical study of the effectiveness of fallback on the ECN," Proc. of *NOMS* 2020, Budapest, Hungary, pp.1-4, April 2020. `DOI:10.1109/NOMS47738.2020.9110403`

[10] L. Zou, K. Nagaoka, and C.-X. Chen, "An Evaluation of the Effectiveness of ECN with Fallback on the Internet," *IEICE Trans. on Inf. and Syst.*, vol.E104-D, no.05, pp.628-636, May 2021. `DOI:10.1587/transinf.2020NTP0002`

[11] `https://s3.amazonaws.com/alexa-static/top-1m.csv.zip`