

BANANA
INTERNET-DRAFT
Intended Category: Standards Track

N. Leymann
C. Heidemann
Deutsche Telekom AG
J. Shen
China Telecom Co., Ltd
L. Geng
China Mobile
L. Chen
M. Zhang
X. Geng
Huawei
March 19, 2018

Expires: September 20, 2018

BANdwidth Aggregation for interNet Access (BANANA)
ECN Operations for Bonding Tunnels
draft-leymann-banana-ecn-01

Abstract

This document specifies a Bonding Tunnel ECN Mechanism that uses Explicit Congestion Notification (ECN) in bonding tunnels to notify congestion of a tunnel so that the load-balancing strategy of the tunnel ingress can be adjusted accordingly. Attributes for the control protocol of BANANA are defined to support this mechanism.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

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

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

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

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

INTERNET DRAFT

ECN Operations for Bonding Tunnels

March 19, 2018

Copyright and License Notice

Copyright (c) 2018 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](http://trustee.ietf.org/license-info) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
1.1.	Terminology	3
2.	ECN Features of IP-in-IP Bonding Tunnels	3
2.1.	ECN Features	3
2.2.	ECN Features of IP-in-IP Tunnels	4
2.3.	Bonding Tunnel ECN Mechanism	5
3.	ECN Capability in Bonding Tunnels	6
4.	Congestion Notification in Bonding Tunnels	6
5.	Security Considerations	7
6.	IANA Considerations	7
7.	References	7
7.1.	Normative References	7
7.2.	Informative References	7
	Authors' Addresses	8

[1.](#) Introduction

Conventionally, ECN allows end-to-end notification of network congestion without dropping packets [[RFC3168](#)], and the sender reduces its transmission rate when it receives the congestion indication. ECN may be used between two ECN-enabled endpoints when the underlying network infrastructure also supports it. [[RFC6040](#)] redefines how the ECN field of the IP header should be constructed on entry to and exit from any IP-in-IP tunnel.

This document, however, focuses on load-balancing adjustment between bonding tunnels rather than end-to-end transmission rate adjustment. When establishing the bonding tunnels, the local BANANA box and the remote BANANA box negotiate whether the Bonding Tunnel ECN Mechanism is supported. When this is successfully negotiated, an ECN-aware router may set a mark on the ECN field of the outer IP header of any packets in the tunnel. As soon as the bonding tunnel egress (one of the BANANA boxes) receives the packet with that mark, it will send an Congestion Notification to the bonding tunnel ingress (the other BANANA box) to inform congestion so that the ingress can change the load-balancing strategy accordingly.

ECN Capability and Congestion Notification are two attributes for the control protocol of BANANA defined to support the Bonding Tunnel ECN Mechanism.

[1.1.](#) Terminology

AQM: Active Queue Management

CE: Congestion Experienced

ECN: Explicit Congestion Notification

ECT: ECN-Capable Transport

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

[2.](#) ECN Features of IP-in-IP Bonding Tunnels

[2.1.](#) ECN Features

The ECN field in the IP header has two bits, making four ECN codepoints, '00' to '11', as shown in Figure 1. The not-ECT (ECN-Capable Transport) codepoint '00' indicates a packet that is not using ECN. The ECT codepoints '10' and '01' are set by the data

sender to indicate that the end-points of the transport protocol are ECN-capable. Senders are free to use either the ECT(0) or ECT(1) and routers treat ECT(0) and ECT(1) as equivalent. AQM allows routers to use the CE (Congestion Experienced) codepoint '11' in a packet header as an indication of congestion, instead of relying solely on packet drops. [[RFC 3168](#)]

+-----+-----+		
ECN FIELD		
+-----+-----+		
0	0	Not-ECT
0	1	ECT(1)
1	0	ECT(0)
1	1	CE

Figure 1 The ECN Field in IP

[2.2.](#) ECN Features of IP-in-IP Tunnels

While the outer header of an IP packet can encapsulate one or more IP headers for IP-in-IP tunneling, routers using ECN to signify congestion only mark the immediately visible outer IP header. When the tunnel decapsulator later removes this outer header, it follows rules to propagate congestion markings by combining the ECN fields of the inner and outer IP header into one outgoing IP header. [[RFC 6040](#)]

Figure 2 shows an example about how ECN works in the IP-in-IP tunnel

scenario.

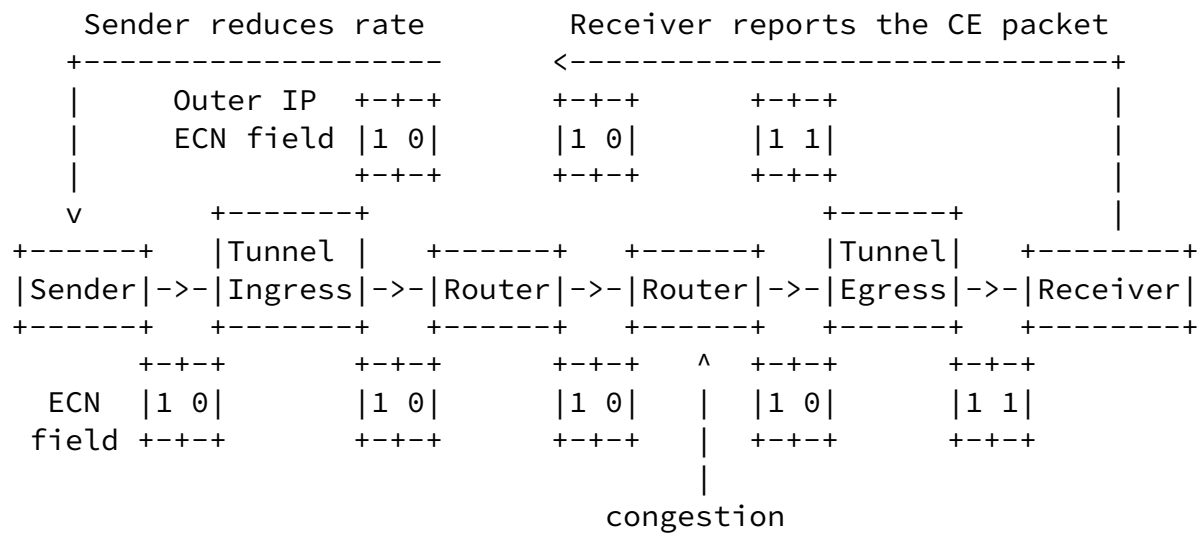



Figure 2 An IP-in-IP Tunnel ECN example

2.3. Bonding Tunnel ECN Mechanism

In the IP-in-IP bonding tunnel scenario, the tunnel ingress has an additional load balancing function compared to the single tunnel scenario. Thus, ECN can be used to notify congestion within the bonding tunnels. As Figure 3 shows, the tunnel egress receives a packet with the CE codepoint from Tunnel 2. Then, the tunnel egress reports this situation to the tunnel ingress by sending an Congestion Notification through Tunnel 2 so that the tunnel ingress can change its load balancing strategy, e.g., temporarily reducing the load-balance proportion for Tunnel 2. As the tunnel ingress may receive more than one Congestion Notification during a certain time period, the load-balance strategy of the next time period can be made based on the number of received Congestion Notifications.



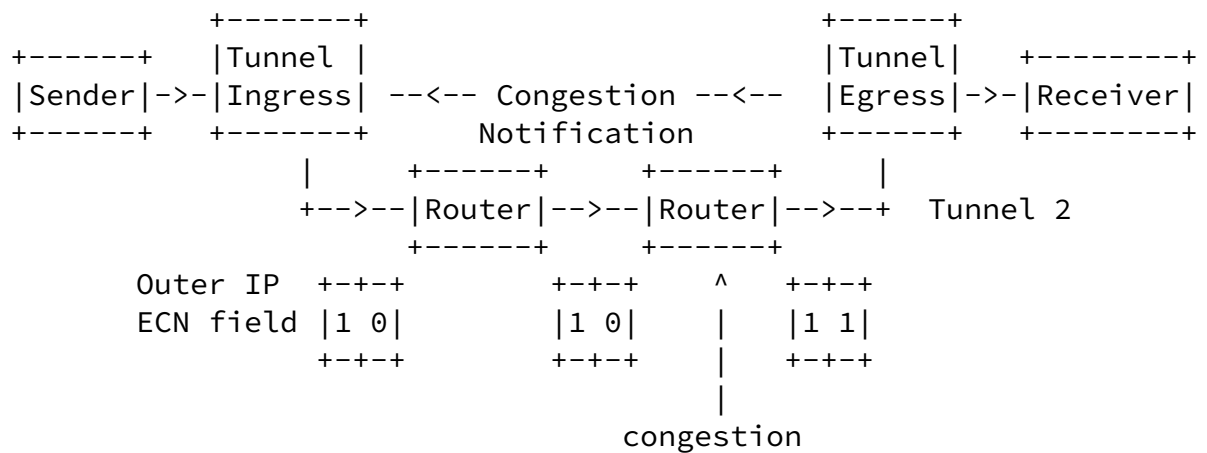


Figure 3 An IP-in-IP Bonding Tunnel ECN example

At the tunnel ingress, the ECN field of the incoming packets will be copied to the inner IP headers. The outer IP headers will be set to the ECT or not-ECT codepoint, according to whether the bonding tunnel supports the Bonding Tunnel ECN Mechanism or not.

At the tunnel egress, if the outer IP headers from Tunnel 1 and Tunnel 2 are both CE and the inner IP headers are ECT, the ECN field of the outgoing packet will be set to CE. Otherwise the ECN field of the outgoing packet will be copied from the inner IP headers.

3. ECN Capability in Bonding Tunnels

The local BANANA box (could be either the tunnel ingress or the tunnel egress) uses the ECN Capability to notify the remote BANANA box (could be either the tunnel egress or the tunnel ingress) that the local BANANA box supports the Bonding Tunnel ECN Mechanism. The first GRE Tunnel Setup Request message [[RFC8157](#)] MAY include the ECN Capability attribute.

```

+---+---+---+---+---+
|Attribute Type | (1 byte)
+---+---+---+---+---+
| Attribute Length | (2 bytes)
+---+---+---+---+---+

```

Attribute Type
ECN Capability, set to 36.

Attribute Length
Set to 0

If the remote BANANA box receives the GRE Tunnel Setup Request message with the ECN Capability attribute included, the remote BANANA box could use the ECN Capability to inform the local BANANA box that the remote BANANA box supports the Bonding Tunnel ECN Mechanism as well. The first GRE Tunnel Setup Accept message MAY include the ECN Capability attribute.

The remote BANANA box activates the Bonding Tunnel ECN Mechanism when it sends out the ECN Capability attribute. The local BANANA box activates the Bonding Tunnel ECN Mechanism when it receives the ECN Capability attribute from the remote BANANA box.

[4.](#) Congestion Notification in Bonding Tunnels

The tunnel egress (could be either the local BANANA box or the remote BANANA box) uses the Congestion Notification to notify congestion on Tunnel 1 or Tunnel 2 to the tunnel ingress. GRE Tunnel Notify messages sent over both Tunnel 1 and Tunnel 2 MAY include the Congestion Notification attribute.

```
+-----+
|Attribute Type |                (1 byte)
+-----+
| Attribute Length |            (2 bytes)
+-----+
```

Attribute Type
Congestion Notification, set to 37.

Attribute Length
Set to 0.

[5.](#) Security Considerations

<TBD>

6. IANA Considerations

No IANA action is required in this document. RFC Editor: please remove this section before publication.

7. References

7.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC3168] Ramakrishnan, K., "The Addition of Explicit Congestion Notification (ECN) to IP", [RFC3168](#), DOI 10.17487/RFC3168, September 2001, <<http://www.rfc-editor.org/info/rfc3168>>.
- [RFC6040] Briscoe, B., "Tunnelling of Explicit Congestion Notification", [RFC6040](#), DOI 10.17487/RFC6040, November 2010, <<http://www.rfc-editor.org/info/rfc3168>>.
- [RFC8157] Leymann, N., "Huawei's GRE Tunnel Bonding Protocol", [RFC8157](#), DOI 10.17487/RFC8157, May 2017, <<http://www.rfc-editor.org/info/rfc8157>>.

7.2. Informative References

- [RFC2784] Farinacci, D., "Generic Routing Encapsulation (GRE)", [RFC2784](#), DOI 10.17487/RFC2784, March 2000, <<http://www.rfc-editor.org/info/rfc2784>>.
- [RFC2890] Dommety, G., "Key and Sequence Number Extensions to GRE", [RFC2890](#), DOI 10.17487/RFC2890, September 2000,

[TSVWG-ECN]

Briscoe, B., "Layered Encapsulation of Congestion Notification", [draft-briscoe-tsvwg-ecn-tunnel-01](#), <<https://tools.ietf.org/pdf/draft-briscoe-tsvwg-ecn-tunnel-01.pdf>>.

[TSVWG-TCF]

Wei, X., "Tunnel Congestion Feedback", [draft-ietf-tsvwg-tunnel-congestion-feedback-05](#), <<https://www.ietf.org/id/draft-ietf-tsvwg-tunnel-congestion-feedback-05.txt>>.

[BANANA-signaling]

Leymann, N., Heidemann, C., et al, "BANdwidth Aggregation for interNet Access (BANANA) The Control Protocol of Bonding Tunnels", [draft-leymann-banana-signaling](#), work in progress.

[BANANA-attributes]

Leymann, N., Heidemann, C., et al, "BANdwidth Aggregation for interNet Access (BANANA) Attributes for the Control Protocol of Bonding Tunnels", [draft-leymann-banana-signaling-attributes](#), work in progress.

Authors' Addresses

Nicolai Leymann
Deutsche Telekom AG
Winterfeldtstrasse 21-27
Berlin 10781
Germany

Phone: +49-170-2275345
EMail: n.leymann@telekom.de

Cornelius Heidemann
Deutsche Telekom AG
Heinrich-Hertz-Strasse 3-7
Darmstadt 64295
Germany

Phone: +4961515812721
EMail: heidemannc@telekom.de

Jun Shen

China Telecom Co., Ltd
109 West Zhongshan Ave, Tianhe District
Guangzhou 510630
P.R. China

EMail: shenjun@gsta.com

Liang Geng
China Mobile
32 Xuanwumen West Street,
Xicheng District, Beijing, 100053,
P.R. China

EMail: gengliang@chinamobile.com

Lihao Chen
Huawei Technologies
No.156 Beiqing Rd. Haidian District,
Beijing 100095
P.R. China

EMail: lihao.chen@huawei.com

Mingui Zhang
Huawei Technologies
No.156 Beiqing Rd. Haidian District,
Beijing 100095
P.R. China

EMail: zhangmingui@huawei.com

Xuesong Geng
Huawei Technologies
No.156 Beiqing Rd. Haidian District,
Beijing 100095
P.R. China

EMail: gengxuesong@huawei.com

