

BANANA
INTERNET-DRAFT
Intended Status: Best Current Practice

N. Leymann
C. Heidemann
Deutsche Telekom AG
J. Shen
China Telecom Co., Ltd
G. Liang
China Mobile
L. Chen
M. Zhang
Huawei
M. Cullen
Painless Security
November 20, 2017

Expires: May 24, 2018

BANdwidth Aggregation for interNet Access (BANANA)
Discard Timer Features for Bonding Tunnel Method
draft-leymann-banana-discard-timer-02

Abstract

Various tunnel based methods have been developed to realize BANdwidth Aggregation for interNet Access (BANANA). At the egress point of bonding tunnels, packets that have been delivered through different tunnels are reordered in a bonding reordering buffer. Along each tunnel, network devices may set aside buffers, which are referred as tunnel buffers, to cache packets as well. This document exposes that a time boundary, which is referred as "Discard Timer", imposed on tunnel buffers could improve the throughput and reliability of the bonding tunnels.

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>

INTERNET DRAFT

Discard Timer for Bonding

November 20, 2017

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

Copyright and License Notice

Copyright (c) 2017 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<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. Problem: Packet blocking in Tunnel Buffers	3
3. Discard Timer Feature	5
4. Security Considerations	6
5. IANA Considerations	7
6. References	7
6.1. Normative References	7
6.2. Informative References	7
Authors' Addresses	7

[1.](#) Introduction

Various bonding tunnel technologies have been proposed to realize BANDwidth Aggregation for interNet Access (BANANA). Per packet traffic distribution is used by bonding tunnels. A reordering buffer is used at the egress node to restore packet disorder. It is referred as "bonding reordering buffer" afterwards in this document.

As latency difference exists between tunnels, packets transmitted through a "faster" tunnel have to wait in the bonding reordering buffer for those packets that are being transmitted through a "slower" tunnel. Timeout limit is the maximum time that a packet can stay in the bonding reordering buffer. When this limit is reached, packets in the bonding reordering buffer MUST be forcibly delivered and un-arrived packets are regarded as lost packets. TCP endpoints would have to retransmit these lost packets.

Along each tunnel, network devices (e.g. routers) use buffers to cache data packets. These buffers are referred as "tunnel buffers" afterwards in this document. For example, an eNodeB would buffer packets transmitted through the LTE tunnel. A tunnel buffer is usually beneficial in absorbing traffic bursts on a tunnel. However, if a packet's buffering process in the tunnel buffer takes too much time so that the timeout limit of the bonding reordering buffer is reached, this packet will be regarded as a lost packet by the egress node of the bonding tunnels. Still, the tunnel buffer has to process and forward that packet, which is unnecessary. It is worthwhile to configure a Discard Timer for the tunnel buffer. Therefore, network devices will discard packets that stay in a tunnel buffer longer than the Discard Timer. This action can also help TCP to adapt to the available sending rate more quickly. The Discard Timer features are discussed in this document.

[1.1.](#) Terminology

DSL: Digital Subscriber Line

shown in Figure 2.1, the ingress point is HAAP(Hybrid Access Aggregation Point), Tunnel 1 is DSL, Tunnel 2 is LTE, the eNodeB buffer is on LTE(this is a tunnel buffer), the egress point is HG(Home Gateway) and the bonding reordering buffer is at the HG. Other possible tunnel buffers on DSL and LTE are omitted. The bonding reordering buffer has a 100 ms timeout limit.

Assuming that the LTE tunnel has a traffic burst at this moment, and the queue of the eNodeB buffer is relatively long. When packet No.2 and No.4 arrive at the eNodeB buffer, they enter the queue and wait for their turn to be processed and forwarded, as shown in Figure 2.1. Meanwhile, packet No.1 and No.3 have already reached the egress point and packet No.3 has entered the bonding reordering buffer to wait for packet No.2. If packet No.2 does not reach the egress point within 100 ms, packet No.2 will be regarded as a lost packet and packet No.3 will be forcibly delivered by the egress point. Data carried by Packet No.2 will be retransmitted by TCP afterwards. However, the eNodeB buffer will process and forward packet No.2 all the same, which is actually unnecessary because by the time packet No.2 reaches

the egress point, it will be discarded as a duplicate packet.

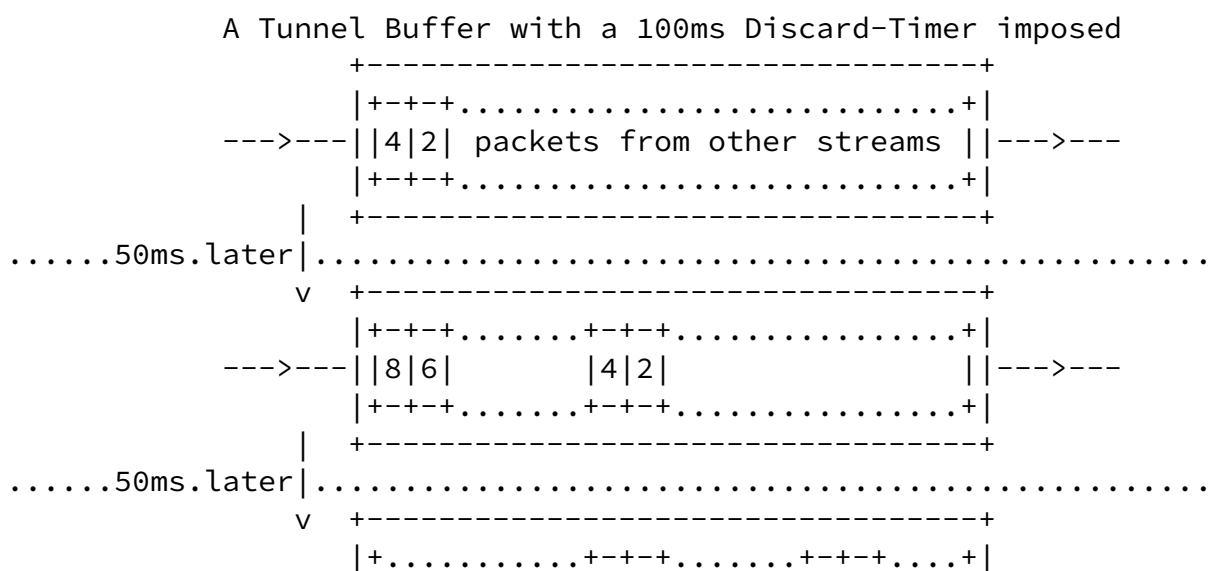
This situation, referred as "packet blocking in tunnel buffers", may continually cause violation of the timeout limit of the bonding reordering buffer, triggering a large amount of retransmission and a significant decrease of sending rate. The throughput of the bonding tunnel could be greatly reduced, even less than the throughput of a single tunnel.

RED(Random Early Detection) mechanism [[RED](#)] is one of the mechanisms commonly used in routers for congestion avoidance. RED monitors the average buffer queue size and drops incoming packets based on statistical probabilities if the queue size exceeds a specific threshold. However, if RED is implemented on eNodeB buffer in Figure 2.1, for example, packet No.2 and No.4 will not be affected for they have already been in the buffer. RED may not be suitable for solving the problem of packet blocking in tunnel buffers.

[3.](#) Discard Timer Feature

Discard Timer is proposed to control the queue length of tunnel buffers by dropping packets that have been queued for too long.

Assuming that the timeout limit of the bonding reordering buffer is 100 ms, so it does not make sense for packets to be queuing in a tunnel buffer for longer than 100 ms, unless the other tunnel coincidentally has a large latency at the moment. With a 100 ms Discard Timer imposed on the tunnel buffer, any packet that has been queued in the tunnel buffer for 100 ms MUST be forcibly discarded. As shown in Figure 3.1, after queued for 100 ms, packet No.2 and No.4 are discarded and the queue is shortened. As a result, subsequent packets, such as No.6 and No.8, are probably to be processed and forwarded to the bonding reordering buffer in time. Therefore, the implementation of the Discard Timer is helpful, especially when the tunnel buffer queue is long. The queue length would be shorten and latency difference between bonding tunnels would be reduced. Hence, the throughput and reliability of the bonding tunnels would be improved.






Figure 3.1 Discard Timer Effects on a Tunnel Buffer

It is acceptable for packets to be queuing in a tunnel buffer for longer than 100 ms if the other tunnel also has a large latency at that moment, so that the latency difference between tunnels may be less than 100 ms and the timeout limit of the bonding reordering buffer may not be violated. However, this situation indicates that both tunnels are being suffered from a traffic burst and the risk of congestion is high. With a Discard Timer implemented, some packets will be forcibly discarded at the tunnel buffer, earlier than TCP detects congestion, thus reducing the queue length as well as triggering a sending rate decrease and helping TCP to adapt to the available sending rate more quickly.

An undersized Discard Timer may cause the discarding of delayed packets which might be tolerated by the bonding reordering buffer. So, the tunnel buffer's Discard Timer SHOULD be set right to be the bonding reordering buffer's timeout limit.

4. Security Considerations

<TBD>

5. IANA Considerations

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

6. References

6.1. 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>>.
- [RFC2890] Dommety, G., "Key and Sequence Number Extensions to GRE", [RFC2890](#), DOI 10.17487/RFC2890, September 2000, <<http://www.rfc-editor.org/info/rfc2890>>.
- [RED] Floyd, S., Jacobson, V., "Random Early Detection (RED) gateways for Congestion Avoidance", IEEE/ACM Transactions on Networking, DOI 10.1109/90.251892, August 1993, <<http://www.icir.org/floyd/papers/early.twocolumn.pdf>>

6.2. Informative References

- [RFC5681] Allman, M., Paxson, V., and E. Blanton, "TCP Congestion Control", [RFC5681](#), DOI 10.17487/RFC5681, September 2009, <<http://www.rfc-editor.org/info/rfc5681>>.
- [GREbond] N. Leymann, C. Heidemann, M. Zhang, et al, "GRE Tunnel Bonding", [draft-zhang-gre-tunnel-bonding](#), 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

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

Geng Liang
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

Margaret Cullen
Painless Security

INTERNET DRAFT

Discard Timer for Bonding

November 20, 2017

14 Summer St. Suite 202
Malden, MA 02148
USA

EMail: margaret@painless-security.com

