Routing Working Group
Internet-Draft
Intended status: Standards Track
Expires: December 12, 2015

A. Mishra
M. Jethanandani
A. Saxena
Ciena Corporation
S. Pallagatti
Juniper Networks
M. Chen
Huawei
P. Fan
China Mobile

June 10, 2015

BFD Stability
draft-ashesh-bfd-stability-03.txt

Abstract

This document describes extensions to the Bidirectional Forwarding Detection (BFD) protocol to measure BFD stability. Specifically, it describes a mechanism for detection of BFD frame loss.

Requirements Language

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 <xref target="RFC2119">RFC 2119</xref>.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

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."

This Internet-Draft will expire on December 12, 2015.

Table of Contents

1. Introduction

The Bidirectional Forwarding Detection (BFD) protocol operates by transmitting and receiving control frames, generally at high frequency, over the datapath being monitored. In order to prevent significant data loss due to a datapath failure, the tolerance for lost or delayed frames (the Detection Time as described in RFC 5880) is set to the smallest feasible value.

This document proposes a mechanism to detect delayed or lost frames in a BFD session in addition to the datapath fault detection mechanisms of BFD. Such a mechanism presents significant value with the ability to measure the stability of BFD sessions and provides data to the operators.

This document does not propose BFD extension to measure data traffic loss or delay on a link or tunnel and the scope is limited to BFD frames.

2. Use cases

Legacy BFD can't detect any BFD frame delay or loss if delay or loss does not last for dead interval. Frequent delay or loss of BFD frames could potentially lead to flap.

It may be possible that network has healthy link or tunnel but only BFD frames are getting dropped or delayed. This potentially leads to network convergence or use of suboptimal path when fast reroute is enabled such as:

Routing protocols with LFA enabled, BFD is used to monitor the link.

Aggregate Ethernet with BFD to monitor each member link.

Primary and protected tunnels with BFD to monitor tunnels.

This proposal will help BFD session to give more information to operator about the health of BFD session that could be used to avoid BFD session flap with faulty BFD path on a healthy link or tunnel.

In a faulty link or tunnel scenario operator can use BFD health information to dynamically run delay and loss measurement OAM protocol (CFM or LM-DM) to further isolate the issue.

3. BFD Null-Authentication TLV

The functionality proposed for BFD stability measurement is achieved by appending the Null-Authentication TLV to the BFD control frame.

The Null-Authentication TLV (called 0-Auth in this document) extends the existing BFD Authentication TLV structure by adding a new Auth-Type of <IANA Assigned>. This TLV carries the Sequence Number for frame loss measurement and optional sender timestmap.

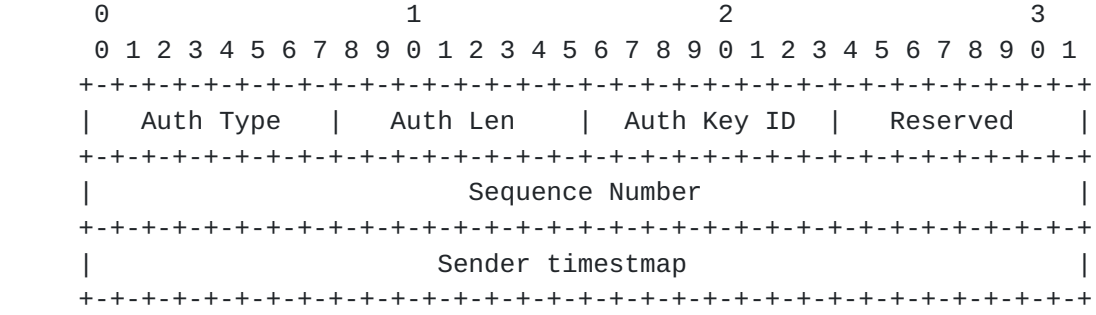

where:

Auth Type: The Authentication Type, which in this case is <IANA assigned> (Null Authentication).

Auth Len: The length of the Authentication Section, in bytes. Length depends on the Auth Key ID.

Auth Key ID: The Authentication Key ID is used to control optional feature. Vales are:

```
              0 -- BFD loss measurement
              1 -- BFD loss and delay measurement
          2-255 -- Reserved for future use
```

when value is set to 0 then last 4 bytes of this TLV MUST not be present in the packet, Auth Len MUST be set to 8 bytes. When set to 1 Auth Len MUST be set to 12.

Sequence Number: This indicates the sequence number for this packet and MUST be present in every 0-Auth TLV. This value is incremented by 1 for every frame transmitted while the session state is UP. A value of 0 indicates a request by sender to reset the sequence number correlation logic at the receiver. The first frame transmitted by the sender MAY set this field to 0.

Sender timestamp: MUST be set to time when packet is about to leave the sender system. Sender system MAY time stamp this as close to wire when packet is about to leave system. Details of how sender system timestamps is out of the scope of this document.

4. Theory of Operations

This mechanism allows operator to measure the loss and delay of BFD CC frames.

4.1. Loss Measurement

This measurement counts the number of BFD control frames missed at the receiver due to a transient change in the network such as congestion. Frame-loss is detected by comparing the Sequence Number field in the 0-Auth TLV in successive BFD CC frames. The Sequence Number in each successive control frame generated on a BFD session by the transmitter is incremented by one.

The first BFD 0-Auth TLV processed by the receiver that has a non-zero sequence number is used for bootstrapping the logic. Each successive frame after this is expected to have a Sequence Number that is one greater than the Sequence Number in the previous frame.

BFD being aggressive protocol, sequence number may wrap to 0 within few hundred days. Sender MUST ensure that when sequence number is wrapped, it starts with value 1. Receiver MUST accept this BFD packet and adjust his next anticipated sequence number.

4.2. Delay Measurement

Delay measurement can be done in two ways.

Using sender timestamp in 0-Auth TLV:

If AuthKey ID in 0-Auth TLV is set to 1 then sender timestamp MUST be set. Delay measurement is the difference between the sender timestamp on any two consecutive BFD CC frames that carry the 0-Auth TLV with AuthKey ID set to 1 for a session. This is a key metric to determine transient changes in stability of BFD transmission engine or to determine the systems capability of handling the existing load. A significant deviation from the negotiated transmission interval on the local node indicates potential instabilities in the BFD transmission engine. Based on the timestamp measurements, the operator MAY take action to configure the system to maintain normal operation of the node.

Similar delay measurements on the receiver can be made using timestamps in the meta data when packet is received. In conjunction with sender delay measurements, these can indicate delays caused by data-path. While a constant delay may not be indicator of instability, large transient delays can decrease the BFD session stability significantly.

Using centralized controller:

When AuthKey ID in 0-Auth TLV is set to 0 then sender timestmap will not be present in the packet. Peers MAY still choose to do delay measurement by sending their packet sent timestamps to central control unit. Central control unit MAY gather all timestamp information and can do delay calculation for a BFD session. Details of how BFD component sends timestamps to central unit is outside the scope of this document.

5. IANA Requirements

IANA is requested to assign new Auth-Type for the Null-Authentication TLV for BFD Stability Measurement. The following number is suggested.

Value Meaning

6 Null-Authentication TLV

6. Security Consideration

Other than concerns raised in [RFC5880] there are no new concerns with this proposal.

7. Contributors

Manav Bhatia
manav@ionosnetworks.com
Ionos Networks
Bangalore, India

8. Acknowledgements

Authors would like to thank Nobo Akiya, Jeffery Haas, Peng Fan, Dileep Singh, Basil Saji, Sagar Soni and Mallik Mudigonda who also contributed to this document.

9. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", RFC 5880, June 2010.

Authors' Addresses

Ashesh Mishra
Ciena Corporation
3939 North 1st Street
San Jose, CA 95134
USA

Email: mishra.ashesh@gmail.com
URI: www.ciena.com

Mahesh Jethanandani
Ciena Corporation
3939 North 1st Street
San Jose, CA 95134
USA

Email: mjethanandani@gmail.com
URI: www.ciena.com

Ankur Saxena
Ciena Corporation
3939 North 1st Street
San Jose, CA 95134
USA

Email: ankurpsaxena@gmail.com

Santosh Pallagatti
Juniper Networks
Juniper Networks, Exora Business Park
Bangalore, Karnataka 560103
India

Phone: +
Email: santoshpk@juniper.net

Mach Chen
Huawei

Email: mach.chen@huawei.com

Peng Fan
China Mobile
32 Xuanwumen West Street
Beijing, Beijing
China

Email: fanp08@gmail.com