Network Working Group
Internet-Draft
Intended status: Experimental
Expires: June 2, 2021

D. Farinacci
lispers.net
S. Ouissal
E. Nordmark
Zededa

November 29, 2020

LISP Data-Plane Telemetry
draft-farinacci-lisp-telemetry-05

Abstract

This draft specs a JSON formatted RLOC-record for telemetry data which decapsulating xTRs include in RLOC-probe Map Reply messages.

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 https://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 June 2, 2021.

Table of Contents

1. Introduction

This document describes how the Locator/Identifier Separation Protocol (LISP) can obtain, measure, and distribute data-plane telemetry information. LISP is an encapsulation protocol built around the fundamental idea of separating the topological location of a network attachment point from the node's identity [I-D.ietf-lisp-rfc6830bis]. As a result LISP creates two namespaces: Endpoint Identifiers (EIDs), that are used to identify end-hosts and routable Routing Locators (RLOCs), used to identify network attachment points. LISP then defines functions for mapping between the two namespaces and for encapsulating traffic originated by devices using non-routable EIDs for transport across a network infrastructure that routes and forwards using RLOCs.

This document specifies how a decapsulating xTR returns telemetry data to an encapsulating xTR using RLOC-probe messages defined in [I-D.ietf-lisp-rfc6833bis].

Early versions of this document will define the type and format of the telemetry data and how it will be distributed. Later versions of this document will describe how telemetry measurement will be performed.

2. Definition of Terms

Encapsulating xTR is a LISP ITR, RTR, or PITR data-plane network element [I-D.ietf-lisp-rfc6830bis]. An encapsulating xTR typically sends RLOC-probe Map-Request messages to decapsulating xTRs to test for reachability of RLOC addresses. For the design scope of this specification, RLOC-probes are also sent to obtain LISP telemetry data measured by a decapsulating xTR.

Decapsulating xTR is a LISP ETR, RTR, or PETR data-plane network element [I-D.ietf-lisp-rfc6830bis]. A decapsulating xTR typically RLOC-probe replies with a Map-Reply message to an RLOC-probe Map-Request sent by an encapsulating xTR. When a decapsulating xTR does data-plane telemetry measurement, it returns measurement data in RLOC-probe Map-Reply messages to an encapsulating xTR.

Telemetry Record a telemetry record is an RLOC-record that contains telemetry data specified in this document. The telemetry data is encoded as an LCAF JSON Type specified in [RFC8060].

3. Overview

The following list of telemetry data has been identified as being useful to obtain:

- Packet Count - the number of packets received within a given time window between the encapsulating xTR and decapsulating xTR.

- Byte Count - the number bytes summed from all packets received within a given time window between the encapsulating xTR and decapsulating xTR.

- Packet Rate - the rate in packets per second an encapsulating xTR is sending encapsulated packets to a decapsulating xTR.

- Bit Rate - the bit rate per second an encapsulating xTR is sending encapsulated packets to a decapsulating xTR.

- Bandwidth - the amount of bandwidth used between encapsulating xTR and decapsulating xTR in bytes per second.

- Packet Loss - the number of packets lost within a given time window between the encapsulating xTR and decapsulating xTR.

- Packet Jitter - the amount of inter-packet time for a train of packets within a given time window between the encapsulating xTR and decapsulating xTR.

- Forward Hop-Count - the number underlay router hops from the encapsulating xTR to the decapsulating xTR.

- Forward One-Way Latency - the amount of time from the encapsulating xTR to the decapsulating xTR. Available when a universal clock and rough time synchronization is available.

- Reverse TTL - the TTL value a decapsulating xTR is using for the RLOC-probe Map-Reply. This is used to compute the return or Reverse Hop-Count or number of underlay router hops between the decapsulating xTR and encapsulating xTR.

- Reverse Timestamp - the universal clock timestamp when the decapsulating xTR sent the RLOC-probe Map-Reply message. This is used to compute the return or Reverse One-Way Latency between the decapsulating xTR to the encapsulating xTR.

4. Telemetry Record Encoding

A Telemetry Record is an RLOC-record encoded in LCAF JSON Type format [RFC8060] within the EID-record inserted in a RLOC-probe Map-Reply message. The RLOC-record is appended to the existing RLOC-records for the EID being probed.

An encapsulating xTR does not need to request telemetry data so the decapsulating xTR can provide it unilaterally by default or via configuration to enable the feature. When an encapsulating xTR receives a Telemetry Record in a RLOC-probe Map-Reply, it SHOULD NOT store it in the map-cache and not use the RLOC-record for forwarding (since there are no RLOCs in this record). The priority for this RLOC-record MUST be set to 255 and the weight MUST be set to 0.

The JSON key values imply directionality. The directionality is from encapsulating xTR to decapsulating xTR. That is, the same direction of RLOC-probe Map-Requests and encapsulated packet flow. The JSON string format is defined to be:

```
{ "type" :              "telemetry",
  "packet-count" :      "<pc>",
  "packet-loss" :       "<pl>",
  "byte-count" :        "<bc>",
  "packet-rate" :       "<pr>",
  "bit-rate" :          "<br>",
  "bandwidth" :         "<b>",
  "packet-jitter" :     "<pj>",
  "forward-latency" :   "<fl>",
  "forward-hop-count" : "<hc>",
  "reverse-ttl" :       "<ttl>",
  "reverse-timestamp" : "<ts>"
}
```

JSON data values:

JSON Value	Encoding Description
<pc>	Number of packets encoded as an integer value within a string.
<pl>	Number of lost packets encoded as an integer value within a string.
<bc>	Number of bytes encoded as an integer value within a string.
<pr>	Packet rate in packets per second encoded as an integer value within a string.
 	Bit rate in kilobits per second encoded as an integer value within a string.
<b>	Bandwidth in kilobytes encoded as an integer value within a string.
<pj>	Packet jitter in milliseconds encoded as an integer value within a string.
<fl>	Latency in milliseconds encoded as an integer value within a string.
<hc>	Hop count encoded as an integer value within a string.
<ttl>	Map-Reply IP header TTL encoded as an integer value within a string.
<ts>	Timestamp encoded in Linux UTC format as an within a string (i.e. Tue Jun 26 16:27:25 UTC 2018).

5. Security Considerations

RLOC-probe Map-Reply messages are signed to protect and authenticate the Telemetry Record according to details in [I-D.ietf-lisp-sec]. Telemetry Records can be kept confidential by encrypting RLOC-probe Map-Reply message with the asymmetric keys described in [I-D.ietf-lisp-ecdsa-auth] or the symmetric keys computed by the key exchange detailed in [RFC8061].

6. IANA Considerations

At this time there are no specific requests for IANA.

7. References

7.1. Normative References

[RFC8060] Farinacci, D., Meyer, D., and J. Snijders, "LISP Canonical Address Format (LCAF)", RFC 8060, DOI 10.17487/RFC8060, February 2017, <https://www.rfc-editor.org/info/rfc8060>.

[RFC8061] Farinacci, D. and B. Weis, "Locator/ID Separation Protocol (LISP) Data-Plane Confidentiality", RFC 8061, DOI 10.17487/RFC8061, February 2017, <https://www.rfc-editor.org/info/rfc8061>.

7.2. Informative References

[I-D.ietf-lisp-ecdsa-auth] Farinacci, D. and E. Nordmark, "LISP Control-Plane ECDSA Authentication and Authorization", draft-ietf-lisp-ecdsa-auth-04 (work in progress), September 2020.

[I-D.ietf-lisp-rfc6830bis] Farinacci, D., Fuller, V., Meyer, D., Lewis, D., and A. Cabellos-Aparicio, "The Locator/ID Separation Protocol (LISP)", draft-ietf-lisp-rfc6830bis-36 (work in progress), November 2020.

[I-D.ietf-lisp-rfc6833bis] Farinacci, D., Maino, F., Fuller, V., and A. Cabellos-Aparicio, "Locator/ID Separation Protocol (LISP) Control-Plane", draft-ietf-lisp-rfc6833bis-30 (work in progress), November 2020.

[I-D.ietf-lisp-sec] Maino, F., Ermagan, V., Cabellos-Aparicio, A., and D. Saucez, "LISP-Security (LISP-SEC)", draft-ietf-lisp-sec-21 (work in progress), July 2020.

Appendix A. Acknowledgments

The authors would like to thank the LISP WG for their review and acceptance of this draft. A special thanks to Colin Cantrell for his review, commentary and guidance.

Appendix B. Document Change Log

[RFC Editor: Please delete this section on publication as RFC.]

B.1. Changes to draft-farinacci-lisp-telemetry-05

- Posted November 2020.
- Document timer and reference update.

B.2. Changes to draft-farinacci-lisp-telemetry-04

- Posted June 2020.
- Document timer and reference update.

B.3. Changes to draft-farinacci-lisp-telemetry-03

- Posted December 2019.
- Document timer and reference update.

B.4. Changes to draft-farinacci-lisp-telemetry-02

- Posted June 2019.
- Document timer and reference update.

B.5. Changes to draft-farinacci-lisp-telemetry-01

- Posted December 2018.
- Document timer and reference update.

B.6. Changes to draft-farinacci-lisp-telemetry-00

- Initial draft posted June 2018.

Authors' Addresses

Dino Farinacci
lispers.net
San Jose, CA
USA

Email: farinacci@gmail.com

Said Ouissal
Zededa
Santa Clara, CA
USA

Email: said@zededa.com

Erik Nordmark
Zededa
Santa Clara, CA
USA

Email: erik@zededa.com