SPRING
Internet-Draft
Intended status: Standards Track
Expires: June 16, 2022
Y.Qiu
J.Ye
H.Li
H3C Technology Co.LTD
December 13, 2021

Data Fields Encapsulation Model of In-situ OAM in SRv6 Network
draft-qiu-spring-srv6-ioam-encap-model-01

Abstract

OAM and PM information from the SR endpoints can be piggybacked in the data packet. The OAM and PM information piggybacking in the data packets is also known as In-situ OAM (IOAM). IOAM records OAM information within the packet while the packet traverses a particular network domain. The term "in-situ" refers to the fact that the IOAM data fields are added to the data packets rather than being sent within probe packets specifically dedicated to OAM. This document defines the data fields encapsulation model of IOAM TLV in SRv6 network.

Table of Contents

1. Introduction

OAM and PM information from the SR endpoints can be piggybacked in the data packet. The OAM and PM information piggybacking in the data packets is also known as In-situ OAM (IOAM). IOAM records OAM information within the packet while the packet traverses a particular network domain. The term "in-situ" refers to the fact that the IOAM data fields are added to the data packets rather than being sent within probe packets specifically dedicated to OAM.

This document defines the data field's encapsulation model of IOAM TLV for the Segment Routing headend with H.Encaps encapsulation behavior in SRv6 network.

The IOAM data fields carried are defined in [I-D.ietf-ippm-ioam-data], and can be used for various use-cases including Performance Measurement(PM) and Proof-of-Transit (PoT).

2. Conventions

2.1. Requirement Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2.2. Abbreviations

Abbreviations used in this document:

IOAM In-situ Operations, Administration, and Maintenance

OAM Operations, Administration, and Maintenance

PM Performance Measurement

PoT Proof-of-Transit

SR Segment Routing

SRH SRv6 Header

SRv6 Segment Routing with IPv6 Data plane

3. Data Encapsulation Model of In-situ OAM

The encapsulation format of IOAM TLV and where to fill IOAM TLV in SRv6 domain are already defined in [I-D.draft-ali-spring-ioam-srv6]. It elaborates on the process of encapsulating IOAM data by individual nodes of SRV6. However, it lacks a process for how to perform IOAM measurement when encapsulating an SRV6 packet, for example, in inter-AS scenarios and in the scenario where a protected tunnel path exists.

This document defines two models for encapsulation operation of IOAM data field : Pipe Model and Uniform Model.

3.1. Pipe Model

In the Pipe Model, the SRv6 network acts like a circuit when packets carrying IOAM data traverse the network. Only the ingress and egress nodes collect IOAM information and report to analyzer with the same Flow Monitoring Identification (FlowMonID) and same type. It means that only ingress node and egress node of SRv6 network are visible to the analyzer. The analyzer can calculate the end-to-end performance of the SRV6 network.

```
          ========== SRv6 packet =======================>

              +--Transit--(d)-...-Transit--(d)---+
             /          (outer header)            \
           (d)                                    (d)
           /                                        \
>--(D)--Ingress...............(D)................Egress--(D)->
        (Push)           (inner header)           (Pop)
```

(d) represents the data field values in the outer SRH
(D) represents the data field values in the encapsulated header

This picture shows data field encapsulation method of In-situ OAM processing in the Pipe Model. The outer IOAM data fields in packet have no relationship to the inner.

The network nodes encapsulate IOAM TLV according to local configuration with a new FlowMonID and a new IOAM-Trace-Type value, and do not care about the IOAM information that is already carried in the packet.

The Pipe model is more suitable for the end-to-end measurement scenarios, since the intermediate router does not need to collect and report data.

3.2. Uniform Model

In the Uniform Model, all the nodes collect IOAM data according to the same IOAM-Trace-Type, and report IOAM data to analyzer with the same FlowMonID. So the analyzer can calculate hop-by-hop forwarding performance based on the IOAM data received from all nodes in the SRv6 network.

```
          ========== SRv6 packet ========================>

              +--Transit--(D)-...-Transit--(D)---+
             /          (outer header)          \
           (D)                                  (D)
           /                                      \
>--(D)--Ingress...............(D)................Egress--(D)->
        (Push)           (inner header)          (Pop)
```

(D) represents the data field values in the corresponding IOAM TLV

This picture shows data field encapsulation of In-situ OAM processing for a Uniform Model.

With the Uniform model, the inner and outer IOAM data-fields are synchronized, including FlowMonID IOAM-trace-Type IOAM-option-Types, etc. The contents of IOAM fields are uniform before and after tunnel encapsulation. The easy way to do it is to copy directly from the inner IOAM TLV.

Uniform model is suitable for postcard IOAM in Hop-by-Hop measurement scenario. Because cannot see how many routers are contained in another autonomous system in inter-AS scenario, Uniform mode is noneffective to passport IOAM measurement. Postcard IOAM measurement in inter-AS scenario is outside the scope of this document.

4. In-situ OAM Process Example For Uniform Model

```
       +----------------------+  +----------------------+
       |         AS1          |  |         AS2          |
+-+-+  | +-+-+  +-+-+  +-+-+ |  | +-+-+  +-+-+  +-+-+ |  +-+-+
+CE1+--+-+PE1+--+P1 +--+PE2+-+--+-+PE3+--+P2 +--+PE4+-+--+CE2+
+-+-+  | +-+-+  +-+-+  +-+-+ |  | +-+-+  +-+-+  +-+-+ |  +-+-+
       |                      |  |                      |
       +----------------------+  +----------------------+
```

Figure 1: Example Inter-AS Scenario of In-situ OAM

This figure shows an example of In-situ OAM used in across SRv6 autonomous systems. PE1, P1 and PE2 are SRv6-capable nodes in autonomous system AS1. PE3, P2, PE4 are SRv6-capable nodes in autonomous system AS2. An SRv6 instantiation of a Binding SID (BSID) of PE3 is used to cross autonomous system. When the traffic is sent from CE1 to CE2, the process is:

1) PE1 receives the packets and encapsulates SRH with a list of segments destined to BSID of PE3, which is instantiated as an ordered list of SRv6 SIDs <PE1, P1, PE2, BSID>. As part of the SRH encapsulation, AS1's ingress node PE1 adds IOAM TLV into the SRH of the packets. The IOAM TLV contains FlowMonID and IOAM-trace-Type fields. The FlowMonID is used to identify a monitored flow. IOAM-Trace-Type is a 24-bit identifier which specifies which information should be collected in this node.

2) When the packet flow arrives in P1, P1 collects the IOAM data based on the IOAM-trace-Type field in IOAM TLV of the packet, and reports the collected data to the analyzer.

3) When the packet flow arrives in PE3, PE3 also collects IOAM data based on the IOAM-trace-Type field in IOAM TLV of packet, and reports the collected data to the analyzer. After that PE3 matches Binding SID with H.encaps behavior, and pushes a outer IPv6 header with its own SRH according SRv6 policy of BSID, which contains an SID list {PE3, P2, PE4}.

4) PE3 encapsulates an outer IOAM TLV to SRH in the outer IPv6 header according local configuration and the data fields of IOAM TLV carried in packet. The outer IOAM data-fields synchronize IOAM values from the inner IOAM TLV, such as FlowMonID, IOAM-trace-Type, IOAM-option-Types and so on.

5) When the packet flow arrives in P2, the routers in AS2 collect IOAM data based on the IOAM-trace-Type in IOAM TLV of the outer SRH.

6) PE4 removes the outer IPv6 header, and recovers the inner packet. Subsequent devices continue to forward packet according to the inner IPv6 header and collect IOAM data according to the inner IOAM TLV.

Because the same FlowMonID is used throughout the forward path across multiple autonomous systems, the analyzer detects and identifies anomalies in the network based on the collected data reported by each of the devices, so as to accurately detect the delay and packet loss of each service, making the network quality service level agreement (SLA) visible in real time, and achieving rapid fault delimitation and location.

5. In-situ OAM Process Example For Pipe Model

The Pipe model is also illustrated using Figure 1. When the traffic is sent from CE1 to CE2, the process is:

1) PE1 receives the packets and encapsulates SRH with a list of segments destined to BSID of PE3, which is instantiated as an ordered list of SRv6 SIDs <PE1, P1, PE2, BSID>. As part of the SRH encapsulation, AS1's ingress node PE1 adds IOAM TLV to the SRH of the packets.

2) When the packet arrives in P1 and PE2, P1 and PE2 collect the IOAM data based on the IOAM-trace-Type field in IOAM TLV of the packet, and report the collected data to the analyzer.

3) When the packet arrives in PE3, PE3 also collects IOAM data based on the IOAM-trace-Type field in IOAM TLV of packet, and reports the collected data to the analyzer. After that PE3 matches Binding SID with H.encaps behavior, and pushes an outer IPv6 header with its own SRH according SRv6 policy of BSID, which contains a SID list {PE3, P2, PE4}.

4) If configuration requires, PE3 identifies target traffic flow that requires IOAM measurement based on the local configuration, and encapsulates the IOAM TLV in the outer SRH. Then PE3 assigns a new FlowMonID to the target flow, populates the IOAM data fields with the new IOAM-trace-Type and IOAM-option-Types.

5) When the packet flow arrives in P2, the routers in AS2 collect IOAM data based on the IOAM-trace-Type in IOAM TLV of the outer SRH.

6) PE4 removes the outer IPv6 header, and recovers the inner packet. Subsequent devices continue to forward packet according to the inner IPv6 header and collect IOAM informantion according to the inner IOAM TLV.

Because the two ASes use different FlowMonIDs for the same flow, according to the FlowMonID identified by PE1, the analyzer can only measure the forwarding performance of this flow between PE3 and PE4 in AS2. It is not possible to measure performance data between other nodes in AS2.

6. IANA Considerations

No requirements for IANA.

7. Security Considerations

TBA

8. Acknowledgements

The authors would like to thank people for their comments to this work.

9. References

9.1. 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, <https://www.rfc-editor.org/info/rfc2119>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[I-D.ietf-ippm-ioam-data] Brockners, F., Bhandari, S., Pignataro, C., Gredler, H., Leddy, J., Youell, S., Mizrahi, T., Mozes, D., Lapukhov, P., Chang, R., and Bernier, D., "Data Fields for In-situ OAM", draft-ietf-ippm-ioam-data-16 (work in progress), November 2021.

[I-D.draft-ali-spring-ioam-srv6] Ali, Z., Gandhi, R., Filsfils, C., Brockners, F., Nainar, N., Pignataro, C., Li, C., Chen, M., Dawra, G., "Segment Routing Header encapsulation for In-situ OAM Data", draft-ali-spring-ioam-srv6-04 (work in progress), July 2021.

9.2. Informative References

[I-D.ietf-6man-ipv6-alt-mark] Fioccola, G., Zhou, T., Cociglio, M., Qin, F., and R. Pang, "IPv6 Application of the Alternate Marking Method", draft-ietf-6man-ipv6-alt-mark-12 (work in progress), October 2021.

Authors' Addresses

Yuanxiang Qiu
H3C Technology Co.LTD, No.466 Changhe Rd.
Hangzhou 310008
China

Email: qiuyuanxiang@h3c.com

Jinrong Ye
H3C Technology Co.LTD

Email: jrong.y@h3c.com

Hao Li
H3C Technology Co.LTD

Email: lihao@h3c.com