

SFC WG
Internet-Draft
Intended status: Standards Track
Expires: December 16, 2017

G. Mirsky
ZTE Corp.
W. Meng
ZTE Corporation
B. Khasnabish
ZTE TX, Inc.
C. Wang
June 14, 2017

Multi-Layer OAM for Service Function Chains in Networks
draft-wang-sfc-multi-layer-oam-09

Abstract

A multi-layer approach to the task of Operation, Administration and Maintenance (OAM) of Service Function Chains (SFCs) in networks is presented. Based on the SFC OAM requirements, a multi-layer model is introduced. A mechanism to detect and localize defects using the multi-layer model is also described.

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 16, 2017.

Copyright 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	2
2. Conventions	3
2.1. Requirements Language	3
2.2. Terminology	3
3. Multi-layer Model of SFC OAM	3
4. Requirements for SFC OAM Multi-layer Model	4
5. SFC OAM multi-layer model	5
6. Theory of Operation	6
7. Security Considerations	7
8. IANA Considerations	7
8.1. SFC TLV Type	7
8.2. SFC OAM UDP Port	8
9. References	8
9.1. Normative References	8
9.2. Informative References	9
Authors' Addresses	9

1. Introduction

[RFC7665] defines components necessary to implement Service Function Chain (SFC). These include a classifier which performs classification of incoming packets. A Service Function Forwarder (SFF) is responsible for forwarding traffic to one or more connected Service Functions (SFs) according to the information carried in the SFC encapsulation. SFF also handles traffic coming back from the SF and transports the data packets to the next SFF. And the SFF serves as termination element of the Service Function Path (SFP). SF is responsible for specific treatment of received packets.

Resulting from that SFC is constructed by a number of these components, there are different views from different levels of the SFC. One is the SFC, fully abstract entity, that defines an ordered set of SFs that must be applied to packets selected as a result of classification. But SFC doesn't define exact mapping between SFFs and SFs. Thus there exists another semi-abstract entity referred as SFP. SFP is the instantiation of the SFC in the network and provides a level of indirection between the fully abstract SFC and a fully specified ordered list of SFFs and SFs identities that the packet will visit when it traverses the SFC. The latter entity is being referred as Rendered Service Path (RSP). The main difference between

SFP and RSP is that in the former the authority to select the SFF/SF has been delegated to the network.

This document proposes the multi-layer model of SFC Operation, Administration and Maintenance (OAM) and requirements to improve the troubleshooting efficiency.

2. Conventions

2.1. Requirements 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. Terminology

Unless explicitly specified in this document, active OAM in SFC and SFC OAM are being used interchangeably.

e2e: End-to-End

FM: Fault Management

OAM: Operations, Administration, and Maintenance

RDl: Remote Defect Indication

RSP: Rendered Service Path

SF: Service Function

SFC: Service Function Chain

SFF: Service Function Forwarder

SFP: Service Function Path

3. Multi-layer Model of SFC OAM

As described in [I-D.ietf-sfc-oam-framework], multiple layers come into play to realize the SFC, including the Service layer, the underlying Network layer, as well as the Link layer, which are depicted in Figure 1:

- o The Service layer consists of classifiers and/or service functions/SFs.
- o Network and Transport layers leverage various overlay network technologies interconnecting SFs to establish SFP.
- o The Link layer is technology specific and reflects the technology used in the underlay network.




Figure 1: SFC OAM Multi-Layer model

4. Requirements for SFC OAM Multi-layer Model

To perform the OAM task of fault management (FM) in an SFC, that includes failure detection, defect characterization and localization, this document defines the multi-layer model of OAM, presented in Section 3, and set of requirements towards active OAM mechanisms to be used on an SFC.

In example presented in Figure 1 the service SFP1 may be realized through two RSPs, RSP1(SF1--SF3--SF5) and RSP2(SF2--SF4--SF6). To perform end-to-end (e2e) FM SFC OAM:

REQ#1: Packets of active OAM in SFC SHOULD be fate sharing with data traffic, i.e. in-band with the monitored traffic, i.e. follow exactly the same RSP, in forward direction, i.e. from ingress toward egress end point(s) of the OAM test.

REQ#2: SFC OAM MUST support pro-active monitoring of any element in the SFC availability.

The egress, SFF3 in example in Figure 1, is the entity that detects the failure of the SFC. It must be able to signal the new defect state to the ingress, i.e. SFF1. Hence the following requirement:

REQ#3: SFC OAM MUST support Remote Defect Indication (RDI) notification by egress to the ingress, i.e. source of continuity checking.

REQ#4: SFC OAM MUST support connectivity verification. Definition of mis-connectivity defect entry and exit criteria are outside the scope of this document.

Once the SFF1 detects the defect objective of OAM switches from failure detection to defect characterization and localization.

REQ#5: SFC OAM MUST support fault localization of Loss of Continuity check in the SFC.

REQ#6: SFC OAM MUST support tracing an SFP in order to realize the RSP.

It is practical, as presented in Figure 1, that several SFs share the same SFF. In such case SFP1 may be realized over two RSPs, RSP1(SF1--SF3--SF5) and RSP2(SF2--SF4--SF6).

REQ#7: SFC OAM MUST have the ability to discover and exercise all available RSPs in the transport network.

In process of localizing the SFC failure separating SFC OAM layers is very attractive and efficient approach. To achieve that continuity among SFFs that are part of the same SFP should be verified. Once SFFs reachability along the particular SFP has been confirmed task of defect localization may focus on SF reachability verification. Because reachability of SFFs has already been verified, SFF local to the SF may be used as source.

REQ#8: SFC OAM MUST be able to trigger on-demand FM with responses being directed towards initiator of such proxy request.

By using the multi-layer model OAM that confirms to the above listed requirements is capable to perform efficient defect localization on an SFC.

5. SFC OAM multi-layer model

Figure 2 presents a use case of applying the proposed SFC OAM multi-layer model. In this scenario operator needs to discover SFFs and SFs of the same SFC. The Layer 1 includes the SFFs that are part of the SFP. The Layer 2 - the SFs along the RSP. When trying to do SFC OAM, classifier or service nodes select and confirm which SFC OAM layering they plan to do, then encapsulate the layering information in the SFC OAM packets, and send the SFC OAM packets along the

service function paths to the destination. When receiving the SFC OAM packets, service nodes analyze the layering information and then decide whether sending these packets to next SFFs directly without being processed by SFs for Layer 1 process or sending to SFs for Layer 2 process.




Figure 2: SFC OAM multi-layering model

6. Theory of Operation

Echo Request/Reply is well-known OAM mechanism that is extensively used to detect inconsistencies between states in control plane and data plane, localize defects in the data plane. In SFC OAM Echo Request/Reply is built as extension of Overlay Echo Request/Reply functions [I-D.ooamdt-rtgwg-demand-cc-cv].

Responder to the SFC Echo Request sends the Echo Reply over IP network if the reply mode is Reply via an IPv4/IPv6 UDP Packet [I-D.ooamdt-rtgwg-demand-cc-cv]. Because SFC NSH does not identify the ingress of the Echo Request MUST include this information that to be used as IP destination address for IP/UDP encapsulation of the SFC Echo Reply. Sender of the SFC Echo Request MUST include SFC Source TLV Figure 3.




Figure 3: SFC Source TLV

where

SFC OAM Source Id Type is two octets in length and has the value of TBD1 Section 8.1.

Length is two octets long field and the valuse is equal to the length of the Value field.

Value field contains IP address of the sender of the SFC OAM control message, IPv4 or IPv6.

The UDP destination port for SFC Echo Reply TBD2 will be allocated by IANA Section 8.2.

7. Security Considerations

TBD

8. IANA Considerations

8.1. SFC TLV Type

IANA is requested to create SFC OAM TLV Type registry. All code points in the range 1 through 32759 in this registry shall be allocated according to the "IETF Review" procedure as specified in [RFC5226]. Code points in the range 32760 through 65279 in this registry shall be allocated according to the "First Come First Served" procedure as specified in [RFC5226]. Remaining code points are allocated according to the Table 1:

Value	Description	Reference
0	Reserved	This document
1- 32759	Unassigned	IETF Review
32760 - 65279	Unassigned	First Come First Served
65280 - 65519	Experimental	This document
65520 - 65534	Private Use	This document
65535	Reserved	This document

Table 1: SFC TLV Type Registry

This document defines the following new value in SFC OAM TLV Type registry:

Value	Description	Reference
TBD1	Source IP Address	This document

Table 2: SFC OAM Source IP Address Type

8.2. SFC OAM UDP Port

IANA is requested to allocate UDP port number according to

Service Name	Port Number	Transport Protocol	Description	Semantics Definition	Reference
SFC OAM	TBD2	UDP	SFC OAM	Section 6	This document

Table 3: SFC OAM Port

9. References

9.1. Normative References

[I-D.ooamdt-rtgwg-demand-cc-cv]
 Mirsky, G., Kumar, N., Kumar, D., Chen, M., Yizhou, L., and D. Dolson, "Echo Request and Echo Reply for Overlay Networks", draft-ooamdt-rtgwg-demand-cc-cv-03 (work in progress), March 2017.

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

9.2. Informative References

- [I-D.ietf-sfc-oam-framework] Aldrin, S., Krishnan, R., Akiya, N., Pignataro, C., and A. Ghanwani, "Service Function Chaining Operation, Administration and Maintenance Framework", draft-ietf-sfc-oam-framework-01 (work in progress), February 2016.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, DOI 10.17487/RFC5226, May 2008, <<http://www.rfc-editor.org/info/rfc5226>>.
- [RFC7665] Halpern, J., Ed. and C. Pignataro, Ed., "Service Function Chaining (SFC) Architecture", RFC 7665, DOI 10.17487/RFC7665, October 2015, <<http://www.rfc-editor.org/info/rfc7665>>.

Authors' Addresses

Greg Mirsky
ZTE Corp.

Email: gregimirsky@gmail.com

Wei Meng
ZTE Corporation
No.50 Software Avenue, Yuhuatai District
Nanjing
China

Email: meng.wei2@zte.com.cn, vally.meng@gmail.com

Bhumip Khasnabish
ZTE TX, Inc.
55 Madison Avenue, Suite 160
Morristown, New Jersey 07960
USA

Email: bhumip.khasnabish@ztetx.com

Cui Wang

Email: lindawangjoy@gmail.com