

Workgroup: SFC WG
Internet-Draft:
draft-ietf-sfc-multi-layer-oam-15
Updates: [8300](#) (if approved)
Published: 15 October 2021
Intended Status: Standards Track
Expires: 18 April 2022
Authors: G. Mirsky W. Meng T. Ao
 Ericsson ZTE Corporation Individual contributor
 K. Leung G. Mishra
 Cisco System Verizon Inc.

Active OAM for Service Function Chaining

Abstract

A set of requirements for active Operation, Administration, and Maintenance (OAM) of Service Function Chains (SFCs) in a network is presented in this document. Based on these requirements, an encapsulation of active OAM messages in SFC and a mechanism to detect and localize defects are described.

This document updates RFC 8300. Particularly, it updates the definition of O (OAM) bit in the Network Service Header (NSH) (RFC 8300) and defines how an active OAM message is identified in the NSH.

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 18 April 2022.

Copyright Notice

Copyright (c) 2021 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 (<https://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. Terminology and Conventions](#)
 - [2.1. Requirements Language](#)
 - [2.2. Acronyms](#)
- [3. Requirements for Active OAM in SFC](#)
- [4. Active OAM Identification in the NSH](#)
- [5. Active SFC OAM Header](#)
- [6. Echo Request/Echo Reply for SFC](#)
 - [6.1. Return Codes](#)
 - [6.2. Authentication in Echo Request/Reply](#)
 - [6.3. SFC Echo Request Transmission](#)
 - [6.3.1. Source TLV](#)
 - [6.4. SFC Echo Request Reception](#)
 - [6.4.1. Errored TLVs TLV](#)
 - [6.5. SFC Echo Reply Transmission](#)
 - [6.5.1. SFC Reply Path TLV](#)
 - [6.5.2. Theory of Operation](#)
 - [6.5.3. SFC Echo Reply Reception](#)
 - [6.5.4. Tracing an SFP](#)
 - [6.6. Verification of the SFP Consistency](#)
 - [6.6.1. SFP Consistency Verification packet](#)
 - [6.6.2. SFF Information Record TLV](#)
 - [6.6.3. SF Information Sub-TLV](#)
 - [6.6.4. SF Information Sub-TLV Construction](#)
- [7. Security Considerations](#)
- [8. Acknowledgments](#)
- [9. IANA Considerations](#)
 - [9.1. SFC Active OAM Protocol](#)
 - [9.2. SFC Active OAM](#)
 - [9.2.1. Version in the Active SFC OAM Header](#)
 - [9.2.2. SFC Active OAM Message Type](#)
 - [9.2.3. SFC Active OAM Header Flags](#)
 - [9.3. SFC Echo Request/Echo Reply Parameters](#)
 - [9.3.1. SFC Echo Request/Reply Version](#)
 - [9.3.2. SFC Echo Request Flags](#)
 - [9.3.3. SFC Echo Request/Echo Reply Message Types](#)
 - [9.3.4. SFC Echo Reply Modes](#)

9.3.5.	SFC Echo Return Codes
9.4.	SFC Active OAM TLV Type
9.5.	SFF Information Record TLV Type
9.6.	SF Information Sub-TLV Type
9.7.	SF Identifier Types
10.	References
10.1.	Normative References
10.2.	Informative References
	Contributors' Addresses
	Authors' Addresses

1. Introduction

[[RFC7665](#)] defines data plane elements necessary to implement a Service Function Chaining (SFC). These include:

1. Classifiers that perform the classification of incoming packets. Such classification may result in associating a received packet to a service function chain.
2. Service Function Forwarders (SFFs) that are responsible for forwarding traffic to one or more connected Service Functions (SFs) according to the information carried in the SFC encapsulation and handling traffic coming back from the SFs and forwarding it to the next SFF.
3. SFs that are responsible for executing specific service treatment on received packets.

There are different views from different levels of the SFC. One is the service function chain, an entirely abstract view, which defines an ordered set of SFs that must be applied to packets selected based on classification rules. But service function chain doesn't specify the exact mapping between SFFs and SFs. Thus, another logical construct used in SFC is a Service Function Path (SFP). According to [[RFC7665](#)], SFP is the instantiation of the SFC in the network and provides a level of indirection between the entirely abstract SFCs 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 referred to as Rendered Service Path (RSP). The main difference between SFP and RSP is that the former is the logical construct, while the latter is the realization of the SFP via the sequence of specific SFC data plane elements.

This document defines how active Operation, Administration and Maintenance (OAM), per [[RFC7799](#)] definition of active OAM, is identified when Network Service Header (NSH) is used as the SFC encapsulation. Following the analysis of SFC OAM in [[RFC8924](#)], this document applies and, when necessary, extends requirements listed in

Section 4 of [\[RFC8924\]](#) for the use of active OAM in an SFP supporting fault management and performance monitoring. Active OAM tools, conformant to the requirements listed in [Section 3](#), improve, for example, troubleshooting efficiency and defect localization in SFP because they specifically address the architectural principles of NSH. For that purpose, SFC Echo Request and Echo Reply are specified in [Section 6](#). This mechanism enables on-demand Continuity Check, Connectivity Verification, among other operations over SFC in networks, addresses functionalities discussed in Sections 4.1, 4.2, and 4.3 of [\[RFC8924\]](#). SFC Echo Request and Echo Reply, defined in this document, can be used with encapsulations other than NSH, for example, using MPLS encapsulation, as described in [\[RFC8595\]](#). The applicability of the SFC Echo Request/Reply mechanism in SFC encapsulations other than NSH is outside the scope of this document. Also, this document updates Section 2.2 of [\[RFC8300\]](#) in part of the definition of 0 bit in the NSH.

2. Terminology and Conventions

The terminology defined in [\[RFC7665\]](#) is used extensively throughout this document, and the reader is expected to be familiar with it.

In this document, SFC OAM refers to an active OAM [\[RFC7799\]](#) in an SFC architecture.

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. Acronyms

E2E: End-to-End

FM: Fault Management

NSH: Network Service Header

OAM: Operations, Administration, and Maintenance

RSP: Rendered Service Path

SF: Service Function

SFC: Service Function Chain

SFF: Service Function Forwarder

SFP: Service Function Path

MAC: Message Authentication Code

3. Requirements for Active OAM in SFC

As discussed in [RFC8924], SFC-specific means are needed to perform the OAM task of fault management (FM) in an SFC architecture, including failure detection, defect characterization, and localization. This document defines the set of requirements for active FM OAM mechanisms to be used in an SFC architecture.




Figure 1: An Example of SFC Data Plane Architecture

The architecture example depicted in [Figure 1](#) considers a service function chain that includes three distinct service functions. In this example, the SFP traverses SFF1, SFF2, and SFF3. Each SFF is connected to two instances of the same service function. End-to-end (E2E) SFC OAM has the Classifier as the ingress and SFF3 as its egress. Segment SFC OAM is between two elements that are part of the same SFP. Following are the requirements for an FM SFC OAM, whether with the E2E or segment scope:

REQ#1: Packets of active SFC OAM SHOULD be fate sharing with the monitored SFC data in the forward direction from ingress toward egress endpoint(s) of the OAM test.

The fate sharing, in the SFC environment, is achieved when a test packet traverses the same path and receives the same treatment in the transport layer as an SFC-encapsulated packet (e.g., NSH).

REQ#2: SFC OAM MUST support monitoring of the continuity of the SFP between any of its elements.

An SFC failure might be declared when several consecutive test packets are not received within a pre-determined time. For example, in the E2E FM SFC OAM case, the egress, SFF3, in the example in [Figure 1](#), could be the entity that detects the SFP's failure by monitoring a flow of periodic test packets. The ingress may be

capable of recovering from the failure, e.g., using redundant SFC elements. Thus, it is beneficial for the egress to signal the new defect state to the ingress, which in this example is the Classifier. Hence the following requirement:

REQ#3: SFC OAM MUST support Remote Defect Indication notification by the egress to the ingress.

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

Once the SFF1 detects the defect, the objective of the SFC OAM changes from the detection of a defect to defect characterization and localization.

REQ#5: SFC OAM MUST support fault localization of the Loss of Continuity Check within an SFP.

REQ#6: SFC OAM MUST support an SFP tracing to discover the RSP.

In the example presented in [Figure 1](#), two distinct instances of the same service function share the same SFF. In this example, the SFP can be realized over several RSPs that use different instances of SF of the same type. For instance, RSP1(SFI11--SFI21--SFI31) and RSP2(SFI12--SFI22--SFI32). Available RSPs can be discovered using the trace function discussed in Section 4.3 [[RFC8924](#)] or the procedure defined in [Section 6.5.4](#).

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

The SFC OAM layer model described in [[RFC8924](#)] offers an approach for defect localization within a service function chain. As the first step, the SFP's continuity for SFFs that are part of the same SFP could be verified. After the reachability of SFFs has already been verified, SFFs that serve an SF may be used as a test packet source. In such a case, SFF can act as a proxy for another element within the service function chain.

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

4. Active OAM Identification in the NSH

The 0 bit in the NSH is defined in [[RFC8300](#)] as follows:

0 bit: Setting this bit indicates an OAM packet.

This document updates that definition as follows:

0 bit: Setting this bit indicates an OAM command and/or data in the NSH Context Header or packet payload.

Active SFC OAM is defined as a combination of OAM commands and/or data included in a message that immediately follows the NSH. To identify the active OAM message, the "Next Protocol" field MUST be set to Active SFC OAM (TBA1) ([Section 9.1](#)). The rules for interpreting the values of the 0 bit and the "Next Protocol" field are as follows:

*0 bit set and the "Next Protocol" value does not match one of identifying active or hybrid OAM protocols (per classification defined in [[RFC7799](#)]), e.g., defined in [Section 9.1](#) Active SFC OAM (TBA1).

- a Fixed-Length Context Header or Variable-Length Context Header(s) contain an OAM command or data.
- the "Next Protocol" field determines the type of payload.

*0 bit set and the "Next Protocol" value matches one of identifying active or hybrid OAM protocols:

- the payload that immediately follows the NSH MUST contain an OAM command or data.

*0 bit is clear:

- no OAM in a Fixed-Length Context Header or Variable-Length Context Header(s).
- the payload determined by the "Next Protocol" field MUST be present.

*0 bit is clear, and the "Next Protocol" field identifies active or hybrid OAM protocol MUST be identified and reported as an erroneous combination. An implementation MAY have control to enable processing of the OAM payload.

One conclusion from the above-listed rules of processing the 0 bit and the "Next Protocol" field is to avoid the combination of OAM in an NSH Context Header (Fixed-Length or Variable-Length) and the payload immediately following the NSH because there is no unambiguous way to identify such combination using the 0 bit and the Next Protocol field.

5. Active SFC OAM Header

As demonstrated in Section 4 [[RFC8924](#)] and [Section 3](#) of this document, SFC OAM is required to perform multiple tasks. Several active OAM protocols could be used to address all the requirements. When IP/UDP encapsulation of an SFC OAM control message is used, protocols can be demultiplexed using the destination UDP port number. But extra IP/UDP headers, especially in an IPv6 network, add noticeable overhead. This document defines Active OAM Header ([Figure 2](#)) to demultiplex active OAM protocols on an SFC.




Figure 2: SFC Active OAM Header

V - two-bit-long field indicates the current version of the SFC active OAM header. The current value is 0. The version number is to be incremented whenever a change is made that affects the ability of an implementation to parse or process the SFC Active OAM header correctly. For example, if syntactic or semantic changes are made to any of the fixed fields.

Msg Type - six bits long field identifies OAM protocol, e.g., Echo Request/Reply or Bidirectional Forwarding Detection.

Flags - eight bits long field carries bit flags that define optional capability and thus processing of the SFC active OAM control packet, e.g., optional timestamping. No flags are defined in this document, and therefore, the bit flags MUST be zeroed on transmission and ignored on receipt.

Length - two octets long field that is the length of the SFC active OAM control packet in octets.

6. Echo Request/Echo Reply for SFC

Echo Request/Reply is a well-known active OAM mechanism extensively used to verify a path's continuity, detect inconsistencies between a state in control and the data planes, and localize defects in the data plane. ICMP ([[RFC0792](#)] for IPv4 and [[RFC4443](#)] for IPv6 networks, respectively) and [[RFC8029](#)] are examples of broadly used active OAM protocols based on the Echo Request/Reply principle. The SFC Echo Request/Reply defined in this document addresses several

requirements listed in [Section 3](#). Specifically, it can be used to check the continuity of an SFP, trace an SFP, or localize the failure within an SFP. The SFC Echo Request/Reply control message format is presented in [Figure 3](#).

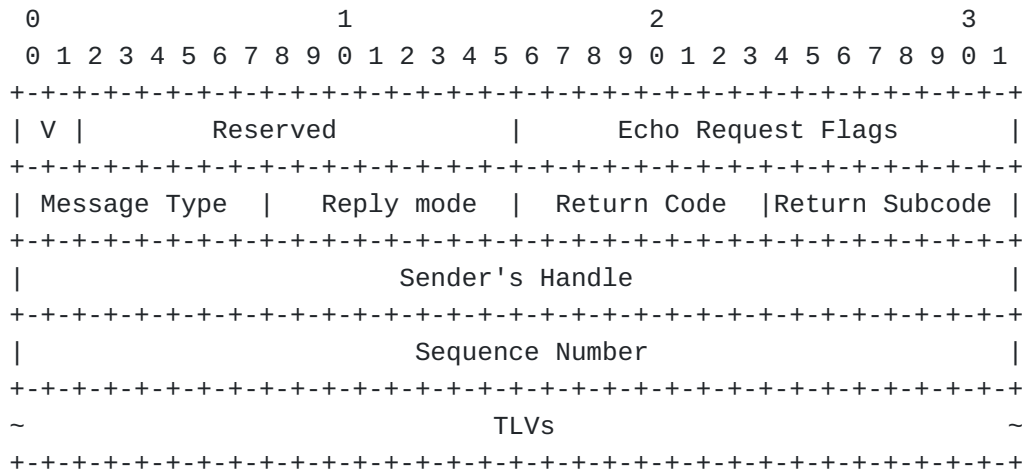


Figure 3: SFC Echo Request/Reply Format

The interpretation of the fields is as follows:

Version (V) is a two-bit field that indicates the current version of the SFC Echo Request/Reply. The current value is 0. The version number is to be incremented whenever a change is made that affects the ability of an implementation to parse or process the control packet correctly. If a packet presumed to carry an SFC Echo Request/Reply is received at an SFF, and the SFF does not understand the Version field value, the packet **MUST** be discarded, and the event **SHOULD** be logged.

Reserved - fourteen-bit field. It MUST be zeroed on transmission and ignored on receipt.

The Echo Request Flags is a two-octet bit vector field. Note that a flag defined in the Flags field of the SFC Active OAM header in [Figure 2](#) has no implication of those defined in the Echo Request Flags field of an Echo Request/Reply message.

The Message Type is a one-octet field that reflects the packet type. Value TBA3 identifies Echo Request and TBA4 - Echo Reply.

The Reply Mode is a one-octet field. It defines the type of the return path requested by the sender of the Echo Request.

Return Codes and Subcodes are one-octet fields each. These can be used to inform the sender about the result of processing its request. Initial Return Code values are provided in [Table 1](#). For

all Return Code values defined in this document, the value of the Return Subcode field MUST be set to zero.

The Sender's Handle is a four-octet field. It MUST be filled in by the sender of the Echo Request and returned unchanged by the Echo Reply sender (if a reply mandated). The sender of the Echo Request SHOULD use a pseudo-random number generator to set the value of the Sender's Handle field.

The Sequence Number is a four-octet field, and it is assigned by the sender and can be, for example, used to detect missed replies. Initial Sequence Number MUST be randomly generated and then SHOULD be monotonically increasing in the course of the test session.




Figure 4: SFC Echo Request/Reply TLV Format

TLV is a variable-length field. Multiple TLVs MAY be placed in an SFC Echo Request/Reply packet. Additional TLVs may be enclosed within a given TLV, subject to the semantics of the (outer) TLV in question. If more than one TLV is to be included, the value of the Type field of the outmost outer TLV MUST be set to "Multiple TLVs Used" (TBA12), as assigned by IANA according to [Section 9.4](#). [Figure 4](#) presents the format of an SFC Echo Request/Reply TLV, where fields are defined as follows:

Type - a one-octet-long field that characterizes the interpretation of the Value field. Type values allocated according to [Section 9.4](#).

Reserved - one-octet-long field. The value of the Type field determines its interpretation and encoding.

Length - two-octet-long field equal to the Value field's length in octets.

Value - a variable-length field. The value of the Type field determines its interpretation and encoding.

6.1. Return Codes

The value of the Return Code field is set to zero by the sender of an Echo Request. The receiver of said Echo Request can set it to one of the values listed in [Table 1](#) in the corresponding Echo Reply that it generates (in cases when the reply is requested).

Value	Description
0	No Return Code
1	Malformed Echo Request received
2	One or more of the TLVs was not understood
3	Authentication failed

Table 1: SFC Echo Return Codes

6.2. Authentication in Echo Request/Reply

Authentication can be used to protect the integrity of the information in SFC Echo Request and/or Echo Reply. In the [[I-D.ietf-sfc-nsh-integrity](#)] a variable-length Context Header has been defined to protect the integrity of the NSH and the payload. The header can also be used for the optional encryption of sensitive metadata. MAC#1 (Message Authentication Code) Context Header is more suitable for the integrity protection of active SFC OAM, particularly of the defined in this document SFC Echo Request and Echo Reply. On the other hand, using MAC#2 Context Header allows the detection of mishandling of the 0-bit by a transient SFC element.

6.3. SFC Echo Request Transmission

SFC Echo Request control packet MUST use the appropriate transport encapsulation of the monitored SFP. If the NSH is used, Echo Request MUST set 0 bit, as defined in [[RFC8300](#)]. NSH MUST be immediately followed by the SFC Active OAM Header defined in [Section 4](#). The Message Type field's value in the SFC Active OAM Header MUST be set to SFC Echo Request/Echo Reply value (TBA2) per [Section 9.2.2](#).

Value of the Reply Mode field MAY be set to:

*Do Not Reply (TBA5) if one-way monitoring is desired. If the Echo Request is used to measure synthetic packet loss, the receiver may report loss measurement results to a remote node. Note that ways of learning the identity of that node are outside the scope of this specification.

*Reply via an IPv4/IPv6 UDP Packet (TBA6) value likely will be the most used.

*Reply via Application Level Control Channel (TBA7) value if the SFP may have bi-directional paths.

*Reply via Specified Path (TBA8) value to enforce the use of the particular return path specified in the included TLV to verify bi-directional continuity and also increase the robustness of the monitoring by selecting a more stable path. [Section 6.5.1](#) provides an example of communicating an explicit path for the Echo Reply.

6.3.1. Source TLV

Responder to the SFC Echo Request encapsulates the SFC Echo Reply message in IP/UDP packet if the Reply mode is "Reply via an IPv4/IPv6 UDP Packet". Because the NSH does not identify the ingress node that generated the Echo Request, the source ID MUST be included in the message and used as the IP destination address and destination UDP port number of the SFC Echo Reply. The sender of the SFC Echo Request MUST include an SFC Source TLV ([Figure 5](#)).

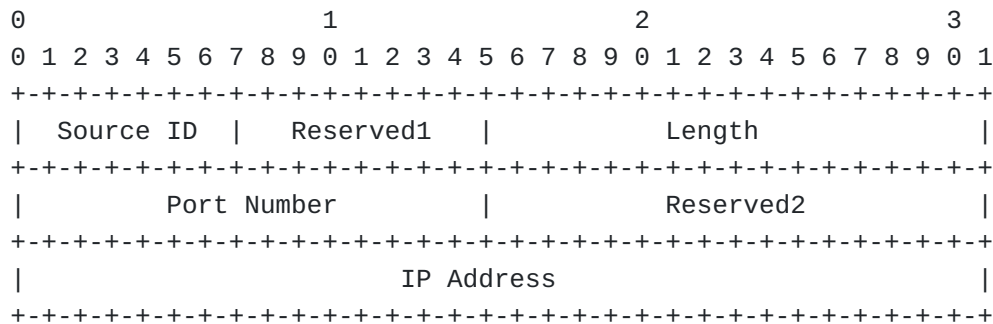


Figure 5: SFC Source TLV

where

Source ID Type is a one-octet-long field and has the value of TBA13 [Section 9.4](#).

Reserved1 - one-octet-long field.

Length is a two-octets-long field, and the value equals the length of the Value field in octets. The value of the Length field can be 8 or 20. If the value of the field is neither, the Source TLV is considered to be malformed.

Port Number is a two-octets-long field. It contains the UDP port number of the sender of the SFC OAM control message. The value of the field MUST be used as the destination UDP port number in the IP/UDP encapsulation of the SFC Echo Reply message.

Reserved2 is a two-octets-long field. The field MUST be zeroed on transmit and ignored on receipt.

IP Address field contains the IP address of the sender of the SFC OAM control message, IPv4 or IPv6. The value of the field MUST be used as the destination IP address in the IP/UDP encapsulation of the SFC Echo Reply message.

A single Source ID TLV for each address family, i.e., IPv4 and IPv6, MAY be present in an SFC Echo Request message. If the Source TLVs for both address families are present in an SFC Echo Request message, the SFF MUST NOT replicate an SFC Echo Reply but choose the destination IP address for the SFC Echo Reply based on the local policy. If more than one Source ID TLV per the address family is present, the receiver MUST use the first TLV and ignore the rest.

6.4. SFC Echo Request Reception

Punting received SFC Echo Request to the control plane is triggered by one of the following packet processing exceptions: NSH TTL expiration, NSH Service Index (SI) expiration, or the receiver is the terminal SFF for an SFP.

Firstly, if the SFC Echo Request is integrity-protected, the receiving SFF first MUST verify the authentication. Then the receiver SFF MUST validate the Source TLV, as defined in [Section 6.3.1](#). Suppose the authentication validation has failed and the Source TLV is considered properly formatted. In that case, the SFF MUST send to the system identified in the Source TLV (see [Section 6.5](#)), according to a rate-limit control mechanism, an SFC Echo Reply with the Return Code set to "Authentication failed" and the Subcode set to zero. If the Source TLV is determined malformed, the received SFC Echo Request processing is stopped, the message is dropped, and the event SHOULD be logged, according to a rate-limiting control for logging. Then, the SFF that has received an SFC Echo Request verifies the rest of the received packet's general sanity. If the packet is not well-formed, the receiver SFF SHOULD send an SFC Echo Reply with the Return Code set to "Malformed Echo Request received" and the Subcode set to zero under the control of the rate-limiting mechanism to the system identified in the Source TLV (see [Section 6.5](#)). If there are any TLVs that the SFF does not understand, the SFF MUST send an SFC Echo Reply with the Return Code set to 2 ("One or more TLVs was not understood") and set the Subcode to zero. In the latter case, the SFF MAY include an Error TLV ([Section 6.4.1](#)) that, as sub-TLVs, contains only the misunderstood TLVs. Sender's Handle and Sequence Number fields are not examined but are included in the SFC Echo Reply message. If the sanity check of the received Echo Request succeeded, then the SFF at the end of the SFP MUST set the Return Code value to 5 ("End of the SFP") and the

Subcode set to zero. If the SFF is not at the end of the SFP and the TTL value is 1, the value of the Return Code MUST be set to 4 ("TTL Exceeded") and the Subcode set to zero. In all other cases, SFF MUST set the Return Code value to 0 ("No Return Code") and the Subcode set to zero.

6.4.1. Errored TLVs TLV

If the Return Code for the Echo Reply is determined as 2 ("One or more TLVs was not understood"), the Errored TLVs TLV might be included in an Echo Reply. The use of this TLV is meant to inform the sender of an Echo Request of TLVs either not supported by an implementation or parsed and found to be in error.




Figure 6: Errored TLVs TLV

where

The Errored TLVs Type MUST be set to TBA14 [Section 9.4](#).

Reserved - one-octet-long field.

Length - two-octet-long field equal to the length of the Value field in octets.

The Value field contains the TLVs, encoded as sub-TLVs, that were not understood or failed to be parsed correctly.

6.5. SFC Echo Reply Transmission

The "Reply Mode" field directs whether and how the Echo Reply message should be sent. The Echo Request sender MAY use TLVs to request that the corresponding Echo Reply be transmitted over the specified path. [Section 6.5.1](#) provides an example of a TLV that specifies the return path of the Echo Reply. Value TBA3 is the "Do not reply" mode and suppresses the Echo Reply packet transmission.

The default value (TBA6) for the Reply mode field requests the responder to send the Echo Reply packet out-of-band as IPv4 or IPv6 UDP packet.

6.5.1. SFC Reply Path TLV

While SFC Echo Request always traverses the SFP, it is directed to using NSH, the corresponding Echo Reply usually is sent without NSH. In some cases, an operator might choose to direct the responder to send the Echo Reply with NSH over a particular SFP. This section defines a new Type-Length-Value (TLV), Reply Service Function Path TLV, for Reply via Specified Path mode of SFC Echo Reply.

The Reply Service Function Path TLV can provide an efficient mechanism to test SFCs, such as bidirectional and hybrid SFC, as defined in Section 2.2 [RFC7665]. For example, it allows an operator to test both directions of the bidirectional or hybrid SFP with a single SFC Echo Request/Echo Reply operation.

The SFC Reply Path TLV carries the information that sufficiently identifies the return SFP that the SFC Echo Reply message is expected to follow. The format of SFC Reply Path TLV is shown in [Figure 7](#).




Figure 7: SFC Reply TLV Format

where:

- *SFC Reply Path Type: is a one-octet-long, indicates the TLV that contains information about the SFC Reply path. IANA is requested to assign value (TBA23),
- *Reserved - one-octet-long field.
- *Length: is two octets long, MUST be equal to 4
- *Reply Service Function Path is used to describe the return path that an SFC Echo Reply is requested to follow.

The format of the Reply Service Function Path field displayed in [Figure 8](#)




Figure 8: Reply Service Function Path Field Format

where:

*Reply Service Function Path Identifier: SFP identifier for the path that the SFC Echo Reply message is requested to be sent over.

*Service Index: the value for the Service Index field in the NSH of the SFC Echo Reply message.

6.5.2. Theory of Operation

[RFC7110] defined mechanism to control return path for MPLS LSP Echo Reply. In SFC's case, the return path is an SFP along which the SFC Echo Reply message MUST be transmitted. Hence, the SFC Reply Path TLV included in the SFC Echo Request message MUST sufficiently identify the SFP that the sender of the Echo Request message expects the receiver to use for the corresponding SFC Echo Reply.

When sending an Echo Request, the sender MUST set the value of Reply Mode field to "Reply via Specified Path", defined in [Section 6.3](#), and if the specified path is an SFC path, the Request MUST include SFC Reply Path TLV. The SFC Reply Path TLV consists of the identifier of the reverse SFP and an appropriate Service Index.

If the NSH of the received SFC Echo Request includes the MAC Context Header, the packet's authentication MUST be verified before using any data. If the verification fails, the receiver MUST stop processing the SFC Return Path TLV and MUST send the SFC Echo Reply with the Return Codes value set to the value Authentication failed from the IANA's Return Codes sub-registry of the SFC Echo Request/Echo Reply Parameters registry.

The destination SFF of the SFP being tested or the SFF at which SFC TTL expired (as per [RFC8300]) may be sending the Echo Reply. The processing described below equally applies to both cases and is referred to as responding SFF.

If the Echo Request message with SFC Reply Path TLV, received by the responding SFF, has Reply Mode value of "Reply via Specified Path" but no SFC Reply Path TLV is present, then the responding SFF MUST send Echo Reply with Return Code set to 6 ("Reply Path TLV is

missing"). If the responding SFF cannot find the requested SFP it MUST send Echo Reply with Return Code set to 7 ("Reply SFP was not found") and include the SFC Reply Path TLV from the Echo Request message.

Suppose the SFC Echo Request receiver cannot determine whether the specified return path SFP has the route to the initiator. In that case, it SHOULD set the value of the Return Codes field to 8 ("Unverifiable Reply Path"). The receiver MAY drop the Echo Request when it cannot determine whether SFP's return path has the route to the initiator. When sending Echo Request, the sender SHOULD choose a proper source address according to the specified return path SFP to help the receiver find the viable return path.

6.5.2.1. Bi-directional SFC Case

The ability to specify the return path for an Echo Reply might be used in the case of bi-directional SFC. The egress SFF of the forward SFP might not be co-located with a classifier of the reverse SFP, and thus the egress SFF has no information about the reverse path of an SFC. Because of that, even for bi-directional SFC, a reverse SFP needs to be indicated in a Reply Path TLV in the Echo Request message.

6.5.3. SFC Echo Reply Reception

An SFF SHOULD NOT accept SFC Echo Reply unless the received message passes the following checks:

- *the received SFC Echo Reply is well-formed;
- *it has an outstanding SFC Echo Request sent from the UDP port that matches destination UDP port number of the received packet;
- *if the matching to the Echo Request found, the value of the Sender's Handle in the Echo Request sent is equal to the value of Sender's Handle in the Echo Reply received;
- *if all checks passed, the SFF checks if the Sequence Number in the Echo Request sent matches to the Sequence Number in the Echo Reply received.

6.5.4. Tracing an SFP

SFC Echo Request/Reply can be used to isolate a defect detected in the SFP and trace an RSP. As for ICMP echo request/reply [[RFC0792](#)] and MPLS echo request/reply [[RFC8029](#)], this mode is referred to as "traceroute". In the traceroute mode, the sender transmits a sequence of SFC Echo Request messages starting with the NSH TTL value set to 1 and is incremented by 1 in each next Echo Request

packet. The sender stops transmitting SFC Echo Request packets when the Return Code in the received Echo Reply equals 5 ("End of the SFP").

Suppose a specialized information element (e.g., IPv6 Flow Label [[RFC6437](#)] or Flow ID [[I-D.ietf-sfc-nsh-tlv](#)]) is used for distributing the load across Equal Cost Multi-Path or Link Aggregation Group paths. In that case, such an element MAY also be used for the SFC OAM traffic. Doing so is meant to control whether the SFC Echo Request follows the same RSP as the monitored flow.

6.6. Verification of the SFP Consistency

The consistency of an SFP can be verified by comparing the view of the SFP from the control or management plane with information collected from traversed by an SFC NSH Echo Request message. Every SFF that receives the Consistency Verification Request (CVReq) MUST perform the following actions:

- *Collect information of the traversed by the CVReq packet SFs and send it to the ingress SFF as CVRep packet over IP network;
- *Forward the CVReq to the next downstream SFF if the one exists.

As a result, the ingress SFF collects information about all traversed SFFs and SFs, information on the actual path the CVReq packet has traveled. That information is used to verify the SFC's path consistency. The mechanism for the SFP consistency verification is outside the scope of this document.

6.6.1. SFP Consistency Verification packet

For the verification of an SFP consistency, two new types of messages to the SFC Echo Request/Reply operation defined in [Section 6](#) with the following values detailed in [Table 10](#):

- *TBA16 - SFP Consistency Verification Request
- *TBA17 - SFP Consistency Verification Reply

Upon receiving the CVReq, the SFF MUST respond with the Consistency Verification Reply (CVRep). The SFF MUST include the SFs information, as described in [Section 6.6.3](#) and [Section 6.6.2](#).

The initiator of CVReq MAY require the collected information in the CVRep be sent in the integrity-protected mode using the MAC Context Header, defined in [[I-D.ietf-sfc-nsh-integrity](#)]. If the NSH of the received SFC Echo Reply includes the MAC Context Header, the authentication of the packet MUST be verified before using any data. If the verification fails, the receiver MUST stop processing the SFF

Information Record TLV and notify an operator. Specification of the notification mechanism is outside the scope of this document.

6.6.2. SFF Information Record TLV

For CVReq, the SFF MUST include the Information of SFs into the SF Information Record TLV in the CVRep message. Every SFF sends back a single CVRep message, including information on all the SFs attached to the SFF on the SFP as requested in the CVReq message.

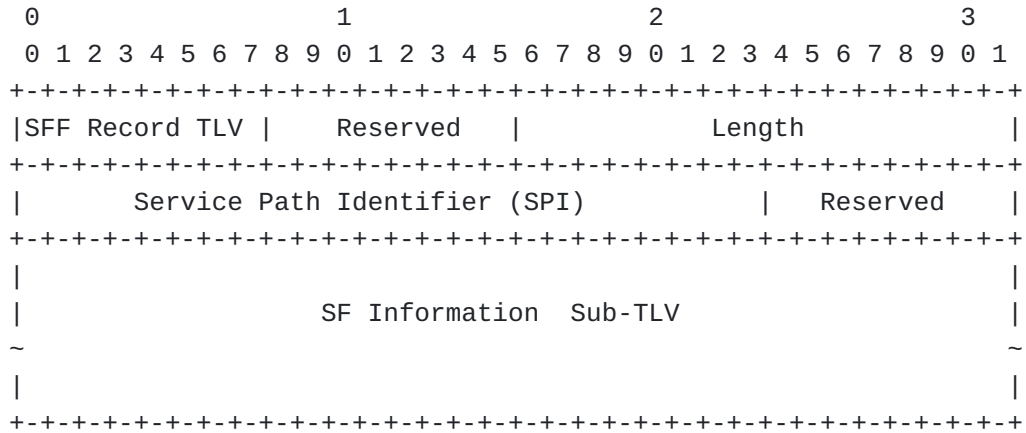


Figure 9: SFF Information Record TLV

SFF Information Record TLV is a variable-length TLV that includes the information of all SFFs mapped to the particular SFF instance for the specified SFP. [Figure 9](#) presents the format of an SFC Echo Request/Reply TLV, where fields are defined as the following:

Reserved - one-octet-long field.

Service Path Identifier (SPI): The identifier of SFP to which all the SFs in this TLV belong.

SF Information Sub-TLV: The Sub-TLV is as defined in [Figure 10](#).

6.6.3. SF Information Sub-TLV

Every SFF receiving CVReq packet MUST include the SF characteristic data into the CVRep packet. The data format of an SF sub-TLV, included in a CVRep packet, is displayed in [Figure 10](#).

After the CVReq message traverses the SFP, all the information of the SFs on the SFP is collected from the TLVs included in CVRep messages.



Figure 10: Service Function information sub-TLV

SF sub-TLV Type: Two octets long field. It indicates that the TLV is an SF TLV that contains the information of one SF.

Length: Two octets long field. The value of the field is the length of the data following the Length field counted in octets.

Service Index: Indicates the SF's position on the SFP.

SF Type: Two octets long field. It is defined in [\[RFC9015\]](#) and indicates the type of SF, e.g., Firewall, Deep Packet Inspection, WAN optimization controller, etc.

Reserved: For future use. MUST be zeroed on transmission and MUST be ignored on receipt.

SF ID Type: One octet-long field with values defined as [Section 9.7](#).

SF Identifier: An identifier of the SF. The length of the SF Identifier depends on the type of the SF ID Type. For example, if the SF Identifier is its IPv4 address, the SF Identifier should be 32 bits. SF ID Type and SF Identifier may be a list, of the SFs included in a load balance group.

6.6.4. SF Information Sub-TLV Construction

Each SFF in the SFP MUST send one and only one CVRep corresponding to the CVReq. If only one SF is attached to the SFF in such SFP, only one SF information sub-TLV is included in the CVRep. If several SFs attached to the SFF in the SFP, SF Information Sub-TLV MUST be constructed as described below in either [Section 6.6.4.1](#) and [Section 6.6.4.2](#).

6.6.4.1. Multiple SFs as hops of SFP

Multiple SFs attached to the same SFF are the hops of the SFP. The service indexes of these SFs are different. Service function types of these SFs could be different or be the same. Information about all SFs MAY be included in the CVRep message. Information about each SF MUST be listed as separate SF Information Sub-TLVs in the CVRep message.

An example of the SFP consistency verification procedure for this case is shown in [Figure 11](#). The Service Function Path(SPI=x) is SF1->SF2->SF4->SF3. The SF1, SF2, and SF3 are attached to SFF1, and SF4 is attached to SFF2. The CVReq message is sent to the SFFs in the sequence of the SFP(SFF1->SFF2->SFF1). Every SFF(SFF1, SFF2) replies with the information of SFs belonging to the SFP. The SF information Sub-TLV in [Figure 10](#) contains information for each SF (SF1, SF2, SF3, and SF4).




Figure 11: Example 1 for CVRep with multiple SFs

6.6.4.2. Multiple SFs for load balance

Multiple SFs may be attached to the same SFF to balance the load; in other words, that means that the particular traffic flow will traverse only one of these SFs. These SFs have the same Service Function Type and Service Index. For this case, the SF identifiers and SF ID Type of all these SFs will be listed in the SF Identifiers field and SF ID Type in a single SF information sub-TLV of the CVRep message. The number of these SFs can be calculated using the SF ID Type and the value of the Length field of the sub-TLV.

An example of the SFP consistency verification procedure for this case is shown in [Figure 12](#). The Service Function Path (SPI=x) is SF1a/SF1b->SF2a/SF2b. The Service Functions SF1a and SF1b are attached to SFF1, which balances the load among them. The Service Functions SF2a and SF2b are attached to SFF2, which, in turn, balances its load between them. The CVReq message is sent to the SFFs in the sequence of the SFP (i.e. SFF1->SFF2). Every SFF (SFF1, SFF2) replies with the information of SFs belonging to the SFP. The SF information Sub-TLV in [Figure 10](#) contains information for all SFs at that hop.

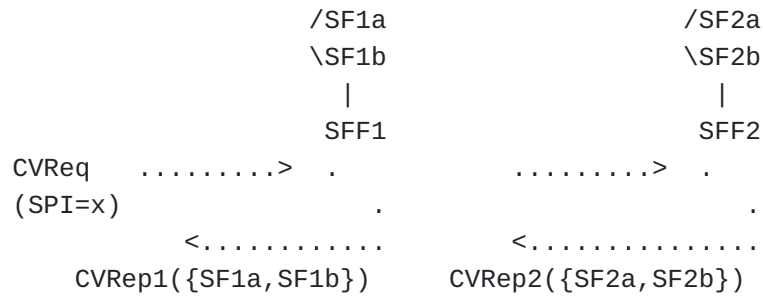


Figure 12: Example 2 for CVRep with multiple SFs

7. Security Considerations

When the integrity protection for SFC active OAM, and SFC Echo Request/Reply in particular, is required, it is RECOMMENDED to use one of the Context Headers defined in [\[I-D.ietf-sfc-nsh-integrity\]](#). MAC#1 Context Header could be more suitable for active SFC OAM because it does not require re-calculation of the MAC when the value of the NSH Base Header's TTL field is changed. The integrity protection for SFC active OAM can also be achieved using mechanisms in the underlay data plane. For example, if the underlay is an IPv6 network, IP Authentication Header [\[RFC4302\]](#) or IP Encapsulating Security Payload Header [\[RFC4303\]](#) can be used to provide integrity protection. Confidentiality for the SFC Echo Request/Reply exchanges can be achieved using the IP Encapsulating Security Payload Header [\[RFC4303\]](#). Also, the security needs for SFC Echo Request/Reply are similar to those of ICMP ping [\[RFC0792\]](#), [\[RFC4443\]](#) and MPLS LSP ping [\[RFC8029\]](#).

There are at least three approaches to attacking a node in the overlay network using the mechanisms defined in the document. One is a Denial-of-Service attack, sending an SFC Echo Request to overload an element of the SFC. The second may use spoofing, hijacking, replying, or otherwise tampering with SFC Echo Requests and/or replies to misrepresent, alter the operator's view of the state of the SFC. The third is an unauthorized source using an SFC Echo Request/Reply to obtain information about the SFC and/or its elements, e.g., SFF or SF.

It is RECOMMENDED that implementations throttle the SFC ping traffic going to the control plane to mitigate potential Denial-of-Service attacks.

Reply and spoofing attacks involving faking or replying to SFC Echo Reply messages would have to match the Sender's Handle and Sequence Number of an outstanding SFC Echo Request message, which is highly unlikely. Thus the non-matching reply would be discarded.

To protect against unauthorized sources trying to obtain information about the overlay and/or underlay, an implementation MAY check that the source of the Echo Request is indeed part of the SFP.

Also, since Service Function sub-TLV discloses information about the SFP the spoofed CVReq packet may be used to obtain network information, it is RECOMMENDED that implementations provide a means of checking the source addresses of CVReq messages, specified in SFC Source TLV [Section 6.3.1](#), against an access list before accepting the message.

8. Acknowledgments

The authors greatly appreciate the thorough review and the most helpful comments from Dan Wing, Dirk von Hugo, and Mohamed Boucadair. The authors are thankful to John Drake for his review and the reference to the work on BGP Control Plane for NSH SFC. The authors express their appreciation to Joel M. Halpern for his suggestion about the load-balancing scenario.

9. IANA Considerations

9.1. SFC Active OAM Protocol

IANA is requested to assign a new type from the SFC Next Protocol registry as follows:

Value	Description	Reference
TBA1	SFC Active OAM	This document

Table 2: SFC Active OAM Protocol

9.2. SFC Active OAM

IANA is requested to create a new SFC Active OAM registry.

9.2.1. Version in the Active SFC OAM Header

IANA is requested to create in the SFC Active OAM registry a new sub-registry called "SFC Active OAM Header Version". All code points are assigned according to the "IETF Review" procedure specified in [[RFC8126](#)]. The remaining code points to be allocated according to [Table 3](#):

Version	Description	Reference
Version 0b00	Protocol as defined by this specification	This document
Version 0b01	Unassigned	This document

Version	Description	Reference
Version 0b10	Unassigned	This document
Version 0b11	Unassigned	This document

Table 3: SFC Active OAM Header Version

9.2.2. SFC Active OAM Message Type

IANA is requested to create in the SFC Active OAM registry a new sub-registry called "SFC Active OAM Message Type". All code points in the range 1 through 32767 in this registry shall be allocated according to the "IETF Review" procedure specified in [[RFC8126](#)]. The remaining code points to be allocated according to [Table 4](#):

Value	Description	Reference
0	Reserved	
1 - 32767	Reserved	IETF Consensus
32768 - 65530	Reserved	First Come First Served
65531 - 65534	Reserved	Private Use
65535	Reserved	

Table 4: SFC Active OAM Message Type

IANA is requested to assign a new type from the SFC Active OAM Message Type sub-registry as follows:

Value	Description	Reference
TBA2	SFC Echo Request/Echo Reply	This document

Table 5: SFC Echo Request/Echo Reply Type

9.2.3. SFC Active OAM Header Flags

IANA is requested to create in the SFC Active OAM registry the new sub-registry SFC Active OAM Flags.

This sub-registry tracks the assignment of 8 flags in the Flags field of the SFC Active OAM Header. The flags are numbered from 0 (most significant bit, transmitted first) to 7.

New entries are assigned by Standards Action.

Bit Number	Description	Reference
7-0	Unassigned	This document

Table 6: SFC Active OAM Header Flags

9.3. SFC Echo Request/Echo Reply Parameters

IANA is requested to create a new SFC Echo Request/Echo Reply Parameters registry.

9.3.1. SFC Echo Request/Reply Version

IANA is requested to create in the SFC Echo Request/Echo Reply Parameters registry a new sub-registry called "SFC Echo Request/Reply Version". All code points assigned according to the "IETF Review" procedure specified in [[RFC8126](#)]. The remaining code points to be allocated according to [Table 7](#):

Version	Description	Reference
Version 0b00	Protocol as defined by this specification	This document
Version 0b01	Unassigned	This document
Version 0b10	Unassigned	This document
Version 0b11	Unassigned	This document

Table 7: SFC Echo Request/Reply Version

9.3.2. SFC Echo Request Flags

IANA is requested to create in the SFC Echo Request/Echo Reply Parameters registry the new sub-registry SFC Echo Request Flags.

This sub-registry tracks the assignment of 16 flags in the SFC Echo Request Flags field of the SFC Echo Request message. The flags are numbered from 0 (most significant bit, transmitted first) to 15.

New entries are assigned by Standards Action.

Bit Number	Description	Reference
15-0	Unassigned	This document

Table 8: SFC Echo Request Flags

9.3.3. SFC Echo Request/Echo Reply Message Types

IANA is requested to create in the SFC Echo Request/Echo Reply Parameters registry the new sub-registry Message Types. All code points in the range 1 through 175 in this registry shall be allocated according to the "IETF Review" procedure specified in [[RFC8126](#)]. Code points in the range 176 through 239 in this registry shall be allocated according to the "First Come First Served"

procedure specified in [[RFC8126](#)]. The remaining code points are allocated as specified in [Table 9](#).

Value	Description	Reference
0	Reserved	This document
1- 175	Unassigned	This document
176 - 239	Unassigned	This document
240 - 251	Experimental	This document
252 - 254	Private Use	This document
255	Reserved	This document

Table 9: SFC Echo Request/Echo Reply
Message Types

IANA is requested to assign values as listed in [Table 10](#).

Value	Description	Reference
TBA3	SFC Echo Request	This document
TBA4	SFC Echo Reply	This document
TBA16	SFP Consistency Verification Request	This document
TBA17	SFP Consistency Verification Reply	This document

Table 10: SFC Echo Request/Echo Reply Message Types Values

9.3.4. SFC Echo Reply Modes

IANA is requested to create in the SFC Echo Request/Echo Reply Parameters registry the new sub-registry Reply Mode. All code points in the range 1 through 175 in this registry shall be allocated according to the "IETF Review" procedure specified in [[RFC8126](#)]. Code points in the range 176 through 239 in this registry shall be allocated according to the "First Come First Served" procedure specified in [[RFC8126](#)]. The remaining code points are allocated according to [Table 11](#).

Value	Description	Reference
0	Reserved	This document
1- 175	Unassigned	This document
176 - 239	Unassigned	This document
240 - 251	Experimental	This document
252 - 254	Private Use	This document
255	Reserved	This document

Table 11: SFC Echo Reply Mode

All code points in the range 1 through 191 in this registry shall be allocated according to the "IETF Review" procedure specified in [[RFC8126](#)] and assign values as listed in [Table 12](#).

Value	Description	Reference
0	Reserved	
TBA5	Do Not Reply	This document
TBA6	Reply via an IPv4/IPv6 UDP Packet	This document
TBA7	Reply via Application Level Control Channel	This document
TBA8	Reply via Specified Path	This document
TBA9	Reply via an IPv4/IPv6 UDP Packet with the data integrity protection	This document
TBA10	Reply via Application Level Control Channel with the data integrity protection	This document
TBA11	Reply via Specified Path with the data integrity protection	This document

Table 12: SFC Echo Reply Mode Values

9.3.5. SFC Echo Return Codes

IANA is requested to create in the SFC Echo Request/Echo Reply Parameters registry the new sub-registry Return Codes as described in [Table 13](#).

Value	Description	Reference
0-191	Unassigned	IETF Review
192-251	Unassigned	First Come First Served
252-254	Unassigned	Private Use
255	Reserved	

Table 13: SFC Echo Return Codes

Values defined for the Return Codes sub-registry are listed in [Table 14](#).

Value	Description	Reference
0	No Return Code	This document
1	Malformed Echo Request received	This document
2	One or more of the TLVs was not understood	This document
3	Authentication failed	This document
4	TTL Exceeded	This document
5	End of the SFP	This document
6	Reply Path TLV is missing	This document
7	Reply SFP was not found	This document
8	Unverifiable Reply Path	This document

Table 14: SFC Echo Return Codes Values

9.4. SFC Active OAM TLV Type

IANA is requested to create the SFC Active OAM TLV Type registry. All code points in the range 1 through 175 in this registry shall be

allocated according to the "IETF Review" procedure specified in [RFC8126]. Code points in the range 176 through 239 in this registry shall be allocated according to the "First Come First Served" procedure specified in [RFC8126]. The remaining code points are allocated according to [Table 15](#):

Value	Description	Reference
0	Reserved	This document
1- 175	Unassigned	This document
176 - 239	Unassigned	This document
240 - 251	Experimental	This document
252 - 254	Private Use	This document
255	Reserved	This document

Table 15: SFC Active OAM TLV Type Registry

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

Value	Description	Reference
TBA12	Multiple TLVs Used	This document
TBA13	Source ID TLV	This document
TBA14	Errored TLVs	This document
TBA23	SFC Reply Path Type	This document

Table 16: SFC OAM Type Values

9.5. SFF Information Record TLV Type

IANA is requested to assign a new type value from SFC OAM TLV Type registry as follows:

Value	Description	Reference
TBA18	SFF Information Record Type	This document

Table 17: SFF-Information Record

9.6. SF Information Sub-TLV Type

IANA is requested to assign a new type value from SFC OAM TLV Type registry as follows:

Value	Description	Reference
TBA19	SF Information	This document

Table 18: SF-Information Sub-TLV Type

9.7. SF Identifier Types

IANA is requested to create in the registry SF Types the new sub-registry SF Identifier Types. All code points in the range 1 through 191 in this registry shall be allocated according to the "IETF Review" procedure as specified in [RFC8126] and assign values as follows:

Value	Description	Reference
0	Reserved	This document
TBA20	IPv4	This document
TBA21	IPv6	This document
TBA22	MAC	This document
TBA22+1-191	Unassigned	IETF Review
192-251	Unassigned	First Come First Served
252-254	Unassigned	Private Use
255	Reserved	This document

Table 19: SF Identifier Type

10. References

10.1. 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>>.
- [RFC8300] Quinn, P., Ed., Elzur, U., Ed., and C. Pignataro, Ed., "Network Service Header (NSH)", RFC 8300, DOI 10.17487/RFC8300, January 2018, <<https://www.rfc-editor.org/info/rfc8300>>.

10.2. Informative References

- [I-D.ietf-sfc-nsh-integrity] Boucadair, M., Reddy, T., and D. Wing, "Integrity Protection for the Network Service Header (NSH) and Encryption of Sensitive Context Headers", Work in Progress, Internet-Draft, draft-ietf-sfc-nsh-integrity-09, 20 September 2021, <<https://>

datatracker.ietf.org/doc/html/draft-ietf-sfc-nsh-integrity-09>.

- [I-D.ietf-sfc-nsh-tlv] Wei, Y., Elzur, U., Majee, S., Pignataro, C., and D. E. Eastlake, "Network Service Header Metadata Type 2 Variable-Length Context Headers", Work in Progress, Internet-Draft, draft-ietf-sfc-nsh-tlv-08, 1 September 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-sfc-nsh-tlv-08>>.
- [RFC0792] Postel, J., "Internet Control Message Protocol", STD 5, RFC 792, DOI 10.17487/RFC0792, September 1981, <<https://www.rfc-editor.org/info/rfc792>>.
- [RFC4302] Kent, S., "IP Authentication Header", RFC 4302, DOI 10.17487/RFC4302, December 2005, <<https://www.rfc-editor.org/info/rfc4302>>.
- [RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)", RFC 4303, DOI 10.17487/RFC4303, December 2005, <<https://www.rfc-editor.org/info/rfc4303>>.
- [RFC4443] Conta, A., Deering, S., and M. Gupta, Ed., "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", STD 89, RFC 4443, DOI 10.17487/RFC4443, March 2006, <<https://www.rfc-editor.org/info/rfc4443>>.
- [RFC6437] Amante, S., Carpenter, B., Jiang, S., and J. Rajahalme, "IPv6 Flow Label Specification", RFC 6437, DOI 10.17487/RFC6437, November 2011, <<https://www.rfc-editor.org/info/rfc6437>>.
- [RFC7110] Chen, M., Cao, W., Ning, S., Jounay, F., and S. Delord, "Return Path Specified Label Switched Path (LSP) Ping", RFC 7110, DOI 10.17487/RFC7110, January 2014, <<https://www.rfc-editor.org/info/rfc7110>>.
- [RFC7665] Halpern, J., Ed. and C. Pignataro, Ed., "Service Function Chaining (SFC) Architecture", RFC 7665, DOI 10.17487/RFC7665, October 2015, <<https://www.rfc-editor.org/info/rfc7665>>.
- [RFC7799] Morton, A., "Active and Passive Metrics and Methods (with Hybrid Types In-Between)", RFC 7799, DOI 10.17487/RFC7799, May 2016, <<https://www.rfc-editor.org/info/rfc7799>>.
- [RFC8029] Kompella, K., Swallow, G., Pignataro, C., Ed., Kumar, N., Aldrin, S., and M. Chen, "Detecting Multiprotocol Label

Switched (MPLS) Data-Plane Failures", RFC 8029, DOI 10.17487/RFC8029, March 2017, <<https://www.rfc-editor.org/info/rfc8029>>.

[RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.

[RFC8595] Farrel, A., Bryant, S., and J. Drake, "An MPLS-Based Forwarding Plane for Service Function Chaining", RFC 8595, DOI 10.17487/RFC8595, June 2019, <<https://www.rfc-editor.org/info/rfc8595>>.

[RFC8924] Aldrin, S., Pignataro, C., Ed., Kumar, N., Ed., Krishnan, R., and A. Ghanwani, "Service Function Chaining (SFC) Operations, Administration, and Maintenance (OAM) Framework", RFC 8924, DOI 10.17487/RFC8924, October 2020, <<https://www.rfc-editor.org/info/rfc8924>>.

[RFC9015] Farrel, A., Drake, J., Rosen, E., Uttaro, J., and L. Jalil, "BGP Control Plane for the Network Service Header in Service Function Chaining", RFC 9015, DOI 10.17487/RFC9015, June 2021, <<https://www.rfc-editor.org/info/rfc9015>>.

Contributors' Addresses

Cui Wang
Individual contributor

Email: lindawangjoy@gmail.com

Bhumip Khasnabish
Individual contributor

Email: vumip1@gmail.com

Zhonghua Chen
China Telecom
No.1835, South PuDong Road
Shanghai
201203
China

Phone: [+86 18918588897](tel:+86-18918588897)
Email: 18918588897@189.cn

Authors' Addresses

Greg Mirsky
Ericsson

Email: gregimirsky@gmail.com

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

Email: meng.wei2@zte.com.cn

Ting Ao
Individual contributor
No.889, BiBo Road
Shanghai
201203
China

Phone: [+86_17721209283](tel:+86_17721209283)
Email: 18555817@qq.com

Kent Leung
Cisco System
170 West Tasman Drive
San Jose, CA 95134,
United States of America

Email: kleung@cisco.com

Gyan Mishra
Verizon Inc.

Email: gyan.s.mishra@verizon.com