Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: January 4, 2018

S. Aldrin
Google
C. Pignataro, Ed.
N. Kumar, Ed.
Cisco
N. Akiya
Big Switch Networks
R. Krishnan
A. Ghanwani
Dell

July 3, 2017

Service Function Chaining Operation, Administration and Maintenance Framework draft-ietf-sfc-oam-framework-02

Abstract

This document provides reference framework for Operations, Administration and Maintenance (OAM) for Service Function Chaining (SFC).

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 RFC 2119 [RFC2119].

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 January 4, 2018.

Table of Contents

1. Introduction

Service Function Chaining (SFC) enables the creation of composite services that consist of an ordered set of Service Functions (SF) that are to be applied to packets and/or frames selected as a result of classification. Service Function Chaining is a concept that provides for more than just the application of an ordered set of SFs to selected traffic; rather, it describes a method for deploying SFs in a way that enables dynamic ordering and topological independence of those SFs as well as the exchange of metadata between participating entities. The foundations of SFC are described in the following documents:

- SFC Problem Statement [RFC7498]

- SFC Architecture [RFC7665]

The reader is assumed to familiar with the material in these documents.

This document provides reference framework for Operations, Administration and Maintenance (OAM, [RFC6291]) of SFC. Specifically, this document provides:

- In Section 2, an SFC layering model;

- In Section 3, aspects monitored by SFC OAM;

- In Section 4, functional requirements for SFC OAM;

- In Section 5, a gap analysis for SFC OAM.

1.1. Document Scope

The focus of this document is to provide an architectural framework for SFC OAM, particularly focused on the aspect of the Operations component within OAM. Actual solutions and mechanisms are outside the scope of this document.

2. SFC Layering Model

Multiple layers come into play for implementing the SFC. These include the service layer at SFC layer and the underlying Network, Transport, Link, etc., layers.

- The service layer, refer to as the "Service Layer" in Figure 1, consists of classifiers and service functions, and uses the overlay network reach from a classifier to service functions and service functions to service functions.

- The overlay network layer, refer to as the "Network" in Figure 1, extends in between various service functions and is mostly transparent to the service functions. It leverages various overlay network technologies interconnecting service functions and allows establishing of service function paths.

- The underlay network layer, refer to as the "Transport" in Figure 1, is dictated by the networking technology of the PSN. It may be either based on MPLS LSPs or IP.

- The link layer, refer to as the "Link" in Figure 1, is dependent upon the physical technology used. Ethernet is a popular choice for this layer, but other alternatives are deployed (e.g. POS, DWDM etc...).

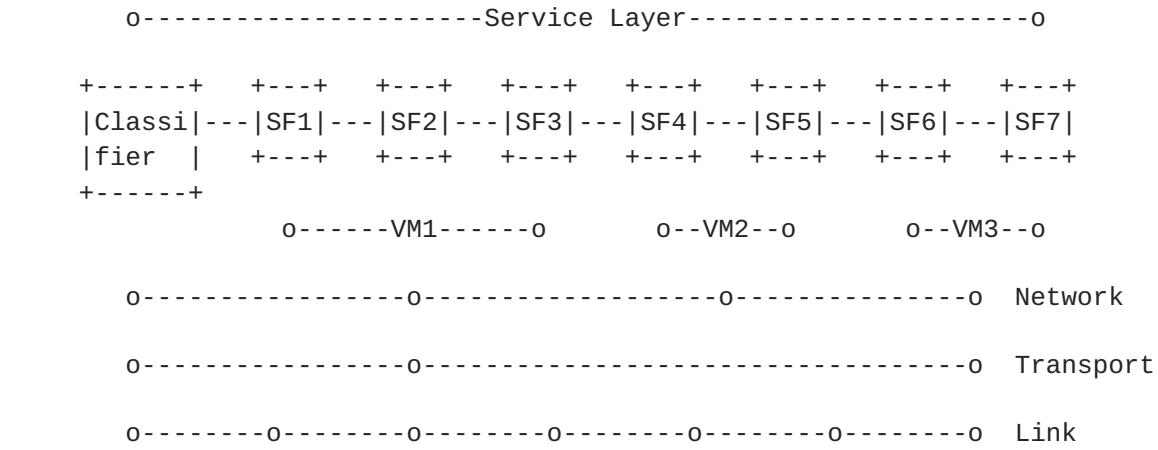

Figure 1: SFC Layering Example

3. SFC OAM Components

The SFC operates at the service layer. For the purpose of defining the OAM framework, the service layer is broken up into three distinct components.

1. Service function component: A function providing a specific service. OAM solutions for this component are to test the service functions from any SFC aware network devices (i.e. classifiers, controllers, other service nodes).

2. Service function chain component: An ordered set of service functions. OAM solution for this component are to test the service function chains and the service function paths.

3. Classifier component: A policy that describes the mapping from flows to service function chains. OAM solutions for this component are to test the validity of the classifiers.

Below figure illustrates an example where OAM for the three defined components are used within the SFC environment.

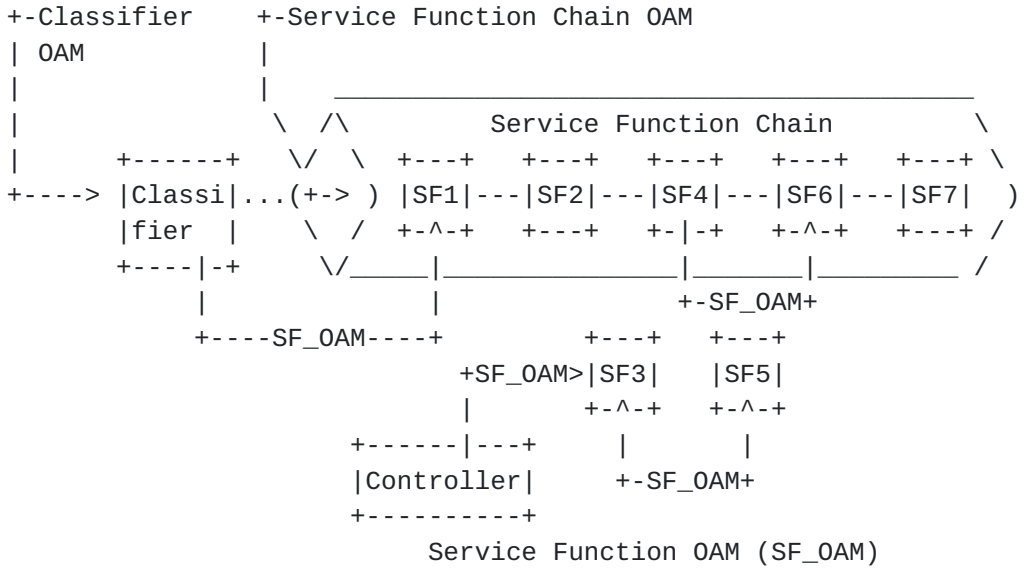

Figure 2: SFC OAM for Three Components

It is expected that multiple SFC OAM solutions will be defined, many targeting one specific component of the service layer. However, it is critical that SFC OAM solutions together provide the coverage of all three SFC OAM components: the service function component, the service function chain component and the classifier component.

3.1. Service Function Component

3.1.1. Service Function Availability

One SFC OAM requirement for the service function component is to allow an SFC aware network device to check the availability to a specific service function, located on the same or different network

devices. Service function availability is an aspect which raises an interesting question. How does one determine that a service function is available? On one end of the spectrum, one might argue that a service function is sufficiently available if the service node (physical or virtual) hosting the service function is available and is functional. On the other end of the spectrum, one might argue that the service function availability can only be concluded if the packet, after passing through the service function, was examined and verified that the packet got expected service applied.

The former approach will likely not provide sufficient confidence to the actual service function availability, i.e. a service node and a service function are two different entities. The latter approach is capable of providing an extensive verification, but comes with a cost. Some service functions make direct modifications to packets, while other service functions do not make any modifications to packets. Additionally, purpose of some service functions is to, conditionally, drop packets intentionally. In such case, packets will not be coming out from the service function. The fact is that there are many flavors of service functions available, and many more flavors of service functions will likely be introduced in future. Even a given service function may introduce a new functionality within a service function (ex: a new signature in a firewall). The cost of this approach is that verifier functions will need to be continuously modified to "keep up" with new services coming out: lack of extendibility.

This framework document provides a RECOMMENDED architectural model where generalized approach is taken to verify that a service function is sufficiently available. TBD - details will be provided in a later revision.

3.1.2. Service Function Performance Measurement

Second SFC OAM requirement for the service function component is to allow an SFC aware network device to check the loss and delay of a specific service function, located on the same or different network devices. TBD - details will be provided in a later revision.

3.2. Service Function Chain Component

3.2.1. Service Function Chain Availability

Verifying an SFC is a complicated process as the SFC could be comprised of varying SF's. Thus, SFC requires the OAM layer to perform validation and verification of SF's within an SFC Path, as well as connectivity and fault isolation.

In order to perform service connectivity verification of an SFC, the OAM could be initiated from any SFC aware network devices for end-to-end paths or partial path terminating on a specific SF within the SFC. This OAM function is to ensure the SF's chained together has connectivity as it is intended to when SFC was established. Necessary return code should be defined to be sent back in the response to OAM packet, in order to qualify the verification.

When ECMP exists at the service layer on a given SFC, there must be an ability to discover and traverse all available paths.

TBD - further details will be provided in a later revision.

3.2.2. Service Function Chain Performance Measurement

The ingress of the service function chain or an SFC aware network device must have an ability to perform loss and delay measurements over the service function chain as a unit (i.e. end-to-end) or to a specific service function through the SFC.

3.3. Classifier Component

A classifier defines a flow and maps incoming traffic to a specific SFC, and it is vital that the classifier is correctly defined and functioning. The SFC OAM must be able to test the definition of flows and the mapping functionality to expected SFCs.

4. SFC OAM Functions

Section 3 described SFC OAM operations required on each SFC component. This section explores the same from the OAM functionality point of view, which many will be applicable to multiple SFC components.

Various SFC OAM requirements provides the need for various OAM functions at different layers. Many of the OAM functions at different layers are already defined and in existence. In order to support SFC and SF's, these functions have to be enhanced to operate a single SF to multiple SF's in an SFC and also multiple SFC's.

4.1. Connectivity Functions

Connectivity is mainly an on-demand function to verify that the connectivity exists between network elements and the availability exists to service functions. Ping is a common tool used to perform this function. OAM messages SHOULD be encapsulated with necessary SFC header and with OAM markings when testing the service function chain component. OAM messages MAY be encapsulated with necessary SFC

header and with OAM markings when testing the service function component. Some of the OAM functions performed by connectivity functions are as follows:

- Verify the MTU size from a source to the destination SF or through the SFC. This requires the ability for OAM packet to take variable length packet size.

- Verify the packet re-ordering and corruption.

- Verify the policy of an SFC or SF using OAM packet.

- Verification and validating forwarding paths.

- Proactively test alternate or protected paths to ensure reliability of network configurations.

4.2. Continuity Functions

Continuity is a model where OAM messages are sent periodically to validate or verify the reachability to a given SF or through a given SFC. This allows monitor network device to quickly detect failures like link failures, network failures, service function outages or service function chain outages. BFD is one such function which helps in detecting failures quickly. OAM functions supported by continuity check are as follows:

- Ability to provision continuity check to a given SF or through a given SFC.

- Notifying the failure upon failure detection for other OAM functions to take appropriate action.

4.3. Trace Functions

Tracing is an important OAM function that allows the operation to trigger an action (ex: response generation) from every transit device on the tested layer. This function is typically useful to gather information from every transit devices or to isolate the failure point towards an SF or through an SFC. Some of the OAM functions supported by trace functions are:

- Ability to trigger action from every transit device on the tested layer towards an SF or through an SFC, using TTL or other means.

- Ability to trigger every transit device to generate response with OAM code(s) on the tested layer towards an SF or through an SFC, using TTL or other means.

- Ability to discover and traverse ECMP paths within an SFC.

- Ability to skip un-supported SF's while tracing SF's in an SFC.

4.4. Performance Measurement Function

Performance management functions involve measuring of packet loss, delay, delay variance, etc. These measurements could be measured pro-actively and on-demand.

SFC OAM framework should provide the ability to perform packet loss for an SFC. In an SFC, there are various SF's chained together. Measuring packet loss is very important function. Using on-demand function, the packet loss could be measured using statistical means. Using OAM packets, the approximation of packet loss for a given SFC could be measured.

Delay within an SFC could be measured from the time it takes for a packet to traverse the SFC from ingress SF to egress SF. As the SFC's are generally unidirectional in nature, measurement of one-way delay is important. In order to measure one-way delay, the clocks have to be synchronized using NTP, GPS, etc.

Delay variance could also be measured by sending OAM packets and measuring the jitter between the packets passing through the SFC.

Some of the OAM functions supported by the performance measurement functions are:

- Ability to measure the packet processing delay of a service function or a service function path along an SFC.

- Ability to measure the packet loss of a service function or a service function path along an SFC.

5. Gap Analysis

This Section identifies various OAM functions available at different levels. It will also identify various gaps, if not all, existing within the existing toolset, to perform OAM function on an SFC.

5.1. Existing OAM Functions

There are various OAM tool sets available to perform OAM function and network layer, protocol layers and link layers. These OAM functions could validate some of the underlay and overlay networks. Tools like ping and trace are in existence to perform connectivity check and tracing intermediate hops in a network. These tools support

different network types like IP, MPLS, TRILL etc. There is also an effort to extend the tool set to provide connectivity and continuity checks within overlay networks. BFD is another tool which helps in detection of data forwarding failures.

Layer	Connectivity	Continuity	Trace	Performance
Underlay N/w	Ping	E-OAM, BFD	Trace	IPPM, MPLS
Overlay N/w	Ping	BFD, NVo3	Trace	IPPM
SF	None	+ None	+ None	+ None
SFC	None	+ None	+ None	+ None

Figure 3: OAM Tool GAP Analysis

Layer	Configuration	Orchestration	Topology	Notification
Underlay N/w	CLI, Netconf	CLI, Netconf	SNMP	SNMP, Syslog
Overlay N/w	CLI, Netconf	CLI, Netconf	SNMP	SNMP, Syslog
SF	CLI	+ CLI	+ None	+ None
SFC	CLI	+ CLI	+ None	+ None

Figure 4: OAM Tool GAP Analysis (contd.)

5.2. Missing OAM Functions

As shown in Figure 3, OAM functions for SFC are not standardized yet. Hence, there are no standard based tools available to verify SF and SFC.

5.3. Required OAM Functions

Primary OAM functions exist for network, transport, link and other layers. Tools like ping, trace, BFD, etc., exist in order to perform these OAM functions. Configuration, orchestration and manageability of SF and SFC could be performed using CLI, Netconf etc.

As seen in Figure 3 and 4, for configuration, manageability and orchestration, providing data and information models for SFC is very much needed. With virtualized SF and SFC, manageability of these functions has to be done programmatically.

6. SFC OAM Model

This section describes the operational aspects of SFC OAM at Service layer to perform the SFC OAM function defined in Section 4 and analyze the applicability of various existing OAM toolsets in the Service layer.

6.1. SFC OAM packet Marker

SFC OAM function described in Section 4 performed at service layer or overlay network layer must mark the packet as OAM packet that can be used by the relevant nodes to differentiate the OAM packet from data packets. The base header defined in Section 3.2 of [I-D.ietf-sfc-nsh] assigns a bit to indicate OAM packets. When NSH encapsulation is used at the service layer, the O bit must be set to differentiate the OAM packet. Any other overlay encapsulations used in future must have a way to mark the packet as OAM packet.

6.2. OAM packet processing and forwarding semantic

Upon receiving OAM packet, SF may choose to discard the packet if it does not support OAM functionality or if the local policy prevent it from processing OAM packet. When SF supports OAM functionality, it is desired to process the packet and respond back accordingly that helps with end-to-end verification. To avoid hitting any performance impact, SF can rate limit the number of OAM packets processed.

Service Function Forwarder (SFF) may choose not to forward the OAM packet to SF if the SF does not support OAM function or if the policy does not allow to forward OAM packet to SF. SFF may choose to skip the SF, modify the header and forward to next SFC node in the chain. How SFF detects if the connected SF supports or allowed to process OAM packet is outside the scope of this document. It could be a configuration paramater instructed by the controller or can be a dynamic negotiation between SF and SFF.

If the SFF receiving the OAM packet is the last SFF in the chain, it must send a relevant response to the initiator of the OAM packet. Depending on the type of OAM solution and tool set used, the response could be a simple response (ICMP reply or BFD reply packet) or could include additional data from the received OAM packet (like stats data consolidated along the path). The proposed solution should detail it further.

The classifier will normally be the node that initiates the OAM packet in order to validate the local classification policy or to validate the SFC or SFP. When the classifier initiates OAM packet, it must set the OAM marker in the overlay encapsulation.

6.3. OAM Function Types

As described in Section 4, there are different OAM functions that may require different OAM solution or tool sets. While the presence of OAM marker in overlay header (For ex: O bit in NSH header) indicates it as OAM packet, it is not sufficient to indicate what OAM function the packet is intended for. We can use the Next Protocol field in NSH header to indicate what OAM function is it intended to or what toolset is used.

6.4. OAM toolset applicability

As described in Section 5.1, there are different tool sets available to perform OAM functions at different layers. This section describes the applicability of some of the available tool sets in service layer.

6.4.1. ICMP Applicability

[RFC0792] and [RFC4443] describes the use of ICMP in IPv4 and IPv6 network respectively. It explains how ICMP messages can be used to test the network reachability between different end points and perform basic network diagnostics.

ICMP could be leveraged for basic OAM functions like SF availability or SFC availability. Initiator can generate ICMP echo message and control the overlay encapsulation header to get the response from relevant node. For example, a classifier initiating OAM can generate ICMP echo message can set the TTL field in NSH header to 255 to get the response from last SFF and thereby test the SFC availability. Alternately, Initiator can set the TTL to other value to get the response from specific SF and there by test the SF availability. Alternately, Initiator could send OAM packets with sequentially incrementing the TTL in NSH header to trace the Service Function Path.

It could be observed that ICMP at its current stage may not be able to perform all SFC OAM functions, but as explained above, it can be used to test the basic OAM functions.

6.4.2. Seamless BFD Applicability

[RFC5880] defines Bidirectional Forwarding Detection (BFD) mechanism for fast failure detection. [RFC5881] and [RFC5884] defines the applicability of BFD in IPv4, IPv6 and MPLS networks. [RFC7880] defines Seamless BFD (S-BFD), a simplified mechanism of using BFD. [RFC7881] explains its applicability in IPv4, IPv6 and MPLS network.

S-BFD could be leveraged to perform SF or SFC availability. Classifier or Initiator could generate BFD control packet and set the "Your Discriminator" value as last SFF in the control packet. Upon receiving the control packet, last SFF will reply back with relevant DIAG code. We could also use the TTL field in NSH header to perform the SF availability. For example, Initiator can set the "Your Discriminator" value to the SF that is intended to be tested and set the TTL field in NSH header in a way that it will be expired on the relevant SF. How the initiator gets the Discriminator value of the SF is outside the scope of this document.

6.4.3. In-Situ OAM

[I-D.brockners-proof-of-transit] defines the mechanism to perform proof of transit to securely verify if a packet traversed the relevant path or chain. While the mechanism is defined inband (i.e, it will be included in data packets), it can be used to perform various SFC OAM functions as well.

In-Situ OAM could be used with O bit set and perform SF availability, SFC availability of performance measurement.

6.4.4. SFC Traceroute

[I-D.penno-sfc-trace] defines a protocol that checks for path liveliness and trace the service hops in any SFP. Section 3 of [I-D.penno-sfc-trace] defines the SFC trace packet format while section 4 and 5 of [I-D.penno-sfc-trace] defines the behavior of SF and SFF respectively.

Initiator can control the SIL in SFC trace packet to perform SF and SFC availability test.

6.5. Security Considerations

SFC and SF OAM must provide mechanisms for:

- Preventing usage of OAM channel for DDOS attacks.
- OAM packets meant for a given SFC should not get leaked beyond that SFC.
- Prevent OAM packets to leak the information of an SFC beyond its administrative domain.

6.6. IANA Considerations

No action is required by IANA for this document.

6.7. Acknowledgements

TBD

7. References

7.1. Normative References

[I-D.brockners-proof-of-transit]
Brockners, F., Bhandari, S., Dara, S., Pignataro, C., Leddy, J., Youell, S., Mozes, D., and T. Mizrahi, "Proof of Transit", draft-brockners-proof-of-transit-03 (work in progress), March 2017.

[I-D.ietf-sfc-nsh]
Quinn, P. and U. Elzur, "Network Service Header", draft-ietf-sfc-nsh-13 (work in progress), June 2017.

[I-D.penno-sfc-trace]
Penno, R., Quinn, P., Pignataro, C., and D. Zhou, "Services Function Chaining Traceroute", draft-penno-sfc-trace-03 (work in progress), September 2015.

[RFC0792] Postel, J., "Internet Control Message Protocol", STD 5, RFC 792, DOI 10.17487/RFC0792, September 1981, <http://www.rfc-editor.org/info/rfc792>.

[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>.

[RFC4443] Conta, A., Deering, S., and M. Gupta, Ed., "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", RFC 4443, DOI 10.17487/RFC4443, March 2006, <http://www.rfc-editor.org/info/rfc4443>.

[RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", RFC 5880, DOI 10.17487/RFC5880, June 2010, <http://www.rfc-editor.org/info/rfc5880>.

[RFC5881] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD) for IPv4 and IPv6 (Single Hop)", RFC 5881, DOI 10.17487/RFC5881, June 2010, <http://www.rfc-editor.org/info/rfc5881>.

[RFC5884] Aggarwal, R., Kompella, K., Nadeau, T., and G. Swallow, "Bidirectional Forwarding Detection (BFD) for MPLS Label Switched Paths (LSPs)", RFC 5884, DOI 10.17487/RFC5884, June 2010, <http://www.rfc-editor.org/info/rfc5884>.

[RFC7498] Quinn, P., Ed. and T. Nadeau, Ed., "Problem Statement for Service Function Chaining", RFC 7498, DOI 10.17487/RFC7498, April 2015, <http://www.rfc-editor.org/info/rfc7498>.

[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>.

[RFC7880] Pignataro, C., Ward, D., Akiya, N., Bhatia, M., and S. Pallagatti, "Seamless Bidirectional Forwarding Detection (S-BFD)", RFC 7880, DOI 10.17487/RFC7880, July 2016, <http://www.rfc-editor.org/info/rfc7880>.

[RFC7881] Pignataro, C., Ward, D., and N. Akiya, "Seamless Bidirectional Forwarding Detection (S-BFD) for IPv4, IPv6, and MPLS", RFC 7881, DOI 10.17487/RFC7881, July 2016, <http://www.rfc-editor.org/info/rfc7881>.

[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, <http://www.rfc-editor.org/info/rfc8029>.

7.2. Informative References

[RFC6291] Andersson, L., van Helvoort, H., Bonica, R., Romascanu, D., and S. Mansfield, "Guidelines for the Use of the "OAM" Acronym in the IETF", BCP 161, RFC 6291, DOI 10.17487/RFC6291, June 2011, <http://www.rfc-editor.org/info/rfc6291>.

Authors' Addresses

Sam K. Aldrin
Google

Email: aldrin.ietf@gmail.com

Carlos Pignataro (editor)
Cisco Systems, Inc.

Email: cpignata@cisco.com

Nagendra Kumar (editor)
Cisco Systems, Inc.

Email: naikumar@cisco.com

Nobo Akiya
Big Switch Networks

Email: nobo.akiya.dev@gmail.com

Ram Krishnan
Dell

Email: ramkri123@gmail.com

Anoop Ghanwani
Dell

Email: anoop@alumni.duke.edu