

rtgwg
Internet-Draft
Intended status: Informational
Expires: July 12, 2017

N. Kumar
C. Pignataro
D. Kumar
Cisco Systems, Inc.
G. Mirsky
Ericsson
M. Chen
Huawei Technologies
E. Nordmark
Arista Networks
S. Pallagatti
Juniper Networks
D. Mozes
Mellanox Technologies Ltd
January 8, 2017

Overlay OAM Requirements
draft-ooamdt-rtgwg-ooam-requirement-02

Abstract

This document describes a list of functional requirements for Operations Administration and Maintenance (OAM) in various Overlay and Service networks like Service Function Chaining (SFC), Bit Index Explicit Replication (BIER), Network Virtualization over Layer 3 (NV03).

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 July 12, 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.](#) Requirements notation [3](#)
- [3.](#) Terminology [3](#)
- [4.](#) Detailed Requirement List [4](#)
 - [4.1.](#) Fault Management [5](#)
 - [4.1.1.](#) Pro-active Fault Management [5](#)
 - [4.1.2.](#) On-demand Fault Management [5](#)
 - [4.2.](#) Performance Management [6](#)
 - [4.3.](#) Alarm Indication Suppression [7](#)
 - [4.4.](#) Overlay Network Resiliency [7](#)
- [5.](#) IANA Considerations [7](#)
- [6.](#) Security Considerations [7](#)
- [7.](#) Acknowledgement [7](#)
- [8.](#) References [7](#)
 - [8.1.](#) Normative References [7](#)
 - [8.2.](#) Informative References [9](#)
- Authors' Addresses [9](#)

1. Introduction

We have witnessed and participated in design of new paradigms in the networking that are aimed to address network virtualization, service function chaining, and multicast services. New paradigms require new architectural concepts, principles and components. [[RFC7365](#)] defines a framework for Data Center Network Virtualization over Layer 3 (NV03). [[RFC7665](#)] describes the architecture for creating and maintaining Service Function Chains (SFCs) in a network. [[I-D.ietf-bier-architecture](#)] defines a stateless multicast architecture for optimal multicast packet forwarding using "Bit Index Explicit Replication" (BIER). These frameworks are defined in a

flexible manner that they are transport agnostic and may be deployed on various underlay networks such as IPv4, IPv6 and MPLS.

The above mentioned new architectural concepts and principles have been combined into new network layers with distinct encapsulation headers. For example, [[I-D.ietf-sfc-nsh](#)] defines an encapsulation header as Network Service Header (NSH) to realize Service Function Path. While [[RFC7348](#)] (VxLAN) and [[RFC7637](#)] (NVGRE) are different encapsulation header proposed for NV03, [[I-D.ietf-nvo3-vxlan-gpe](#)] extends VxLAN further to be used for Service Function Chain (SFC). Similarly, [[I-D.ietf-bier-mpls-encapsulation](#)] defines the BIER encapsulation header over MPLS network and [[I-D.xu-bier-encapsulation](#)] describes the BIER encapsulation header over IP network.

Introduction of the new Overlay networks, sets forth new Operations, Administration and Maintenance (OAM) requirements that can be addressed by enhancing the existing toolset or developing new protocols. For example, [[I-D.ietf-sfc-oam-framework](#)] defines the framework for SFC OAM, [[I-D.nordmark-nvo3-transcending-traceroute](#)] proposes a way to perform traceroute in NV03 networks and [[I-D.kumarzheng-bier-ping](#)] proposes on-demand connectivity verification and fault isolation procedure (Ping and Trace) on BIER network.

The goal of this document is to identify and list the OAM requirements commonly applicable to new Overlay networks which can further be used to analyze the existing OAM tools. The identified gaps can be addressed, either through enhancing existing OAM tools and if necessary, constructing new OAM tools, that can be used as a common unified OAM toolset to support and perform various OAM functions including proactive and on-demand path monitoring and service validation on the new Overlay network.

2. Requirements notation

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 [[RFC2119](#)].

3. Terminology

ECMP: Equal Cost Multipath

SFC: Service Function Chaining

BIER: Bit Index Explicit Replication

NV03: Network Virtualization over L3

OAM: Operations, Administration and Maintenance

MPLS: Multiprotocol Label Switching

VxLAN: Virtual Extensible Local Area Network

NVGRE: Network Virtualization Using Generic Routing Encapsulation

Centralized Controller: An external standalone or virtual entity with topology awareness and with an ability to interact with network devices for OAM functionality.

Overlay nodes: Network nodes participating in the Overlay network.

Underlay Network or Underlay Layer: The network that provides connectivity between the Overlay nodes. MPLS network providing LSP connectivity between BIER nodes is an example for underlay layer.

Overlay Network or Overlay Layer: A network layer that is built on top another network layer. VxLAN-GPE over IP network is an example for Overlay layer.

4. Detailed Requirement List

This section lists the OAM requirement for different Overlay networks. The below listed requirement MUST be supported with any underlay transport network:

- REQ#1: The listed requirements MUST be supported with any type of transport layer over which the overlay network can be realized.
- REQ#2: It MUST be possible to initialize Overlay OAM between any node towards any node(s) in the overlay network.
- REQ#3: It SHOULD be possible to initialize an Overlay OAM from a centralized controller.
- REQ#4: Overlay OAM MUST support proactive and on-demand OAM monitoring and measurement methods.
- REQ#5: Overlay OAM MUST support unidirectional OAM methods for continuity check, connectivity verification and performance measurement.

- REQ#6: Overlay OAM packets SHOULD be fate sharing with data traffic, i.e. in-band with the monitored traffic, i.e. follow exactly the same overlay and transport path as data plane traffic, in forward direction, i.e. from ingress toward egress end point(s) of the OAM test.
- REQ#7: Overlay OAM MUST support bi-directional OAM methods. Such OAM methods MAY combine in-band monitoring or measurement in forward direction and out-of-band notification in the reverse direction, i.e. from egress to ingress end point of the OAM test.
- REQ#8: Overlay OAM MUST support Path Maximum Transmission Unit (MTU) Discovery from the overlay layer over any transport layer.

4.1. Fault Management

4.1.1. Pro-active Fault Management

Availability, not as performance metric, is understood as ability to reach the node, i.e. the fact that path between ingress and egress does exist. Such OAM mechanism also referred as Continuity Check.

- REQ#9: Overlay OAM MUST support pro-active monitoring of any virtual node availability in the given overlay network.
- REQ#10: Overlay OAM MUST support Remote Defect Indication (RDI) notification by egress to the ingress, i.e. source of continuity checking.
- REQ#11: Overlay OAM MUST support connectivity verification. Definition of mis-connectivity defect entry and exit criteria are outside the scope of this document.

4.1.2. On-demand Fault Management

- REQ#12: Overlay OAM MUST support fault localization of Loss of Continuity check at Overlay layer.
- REQ#13: Overlay OAM MAY support fault localization of Loss of Continuity check at transport layer.
- REQ#14: Overlay OAM MUST support tracing path in overlay network through the overlay nodes.
- REQ#15: Overlay OAM MAY support tracing path in underlay network connecting overlay border nodes.

- REQ#16: Overlay OAM MAY support verification of the mapping between its data plane state and client layer services.
- REQ#17: Overlay OAM MUST have the ability to discover and exercise equal cost multipath (ECMP) paths in its underlay network.
- REQ#18: Overlay OAM MUST be able to trigger on-demand FM with responses being directed towards initiator of such proxy request.

4.2. Performance Management

Section 3.4 and Section 3.5 of [RFC7799] defines the definition for Active and Passive mode of Performance Measurement (PM) methods. This section lists the requirements for both active and passive PM methods. Passive PM is a measurement method that should not modify the actual data packet processing behavior on underlay and overlay network. Accordingly, it should be supported by the Overlay nodes.

- REQ#19: Overlay OAM MUST support active one-way packet delay measurement.
- REQ#20: Overlay OAM MUST support passive one-way packet delay measurement.
- REQ#21: Overlay OAM MUST support active two-way packet delay measurement.
- REQ#22: Overlay OAM MUST support packet delay variation measurement.
- REQ#23: Overlay OAM MUST support active end to end packet loss measurement.
- REQ#24: Overlay OAM MUST support passive end to end packet loss measurement.
- REQ#25: Overlay OAM SHOULD support active per-segment packet delay measurement.
- REQ#26: Overlay OAM SHOULD support passive per-segment packet delay measurement.
- REQ#27: Overlay OAM SHOULD support active per-segment packet loss measurement.
- REQ#28: Overlay OAM SHOULD support passive per-segment packet loss measurement.

REQ#29: Overlay OAM MUST support delivered packet throughput measurement.

4.3. Alarm Indication Suppression

REQ#30: Overlay OAM MUST support defect notification mechanism, like Alarm Indication Signal.

REQ#31: Any virtual node in the given overlay network MAY originate a defect notification addressed to any node in that network.

4.4. Overlay Network Resiliency

REQ#32: Overlay OAM MUST support methods to enable survivability of an overlay network. These recovery methods MAY use protection switching and restoration.

5. IANA Considerations

This document does not propose any IANA consideration.

6. Security Considerations

This document list the OAM requirement for various Overlay encapsulations and may have security implications. For example, if proactive FM is required, the security implication is that a passive eavesdropper can know when the session is down. Or, proactive FM may be used either to launch DoS or to highjack session and impact state, e.g. cause protection switchover. These security implications are natural results of the requirements, and do not depend on the particular implementation. Whether existing security mechanisms of existing protocols proposed to be re-used in OAM for overlay networks are adequate or require enhancements is for further study. New OAM protocols for overlay networks must consider their security mechanism to on per-solution basis.

7. Acknowledgement

The Authors would like to thank Ron Bonico, Tal Mizrahi, Alia Atlas and Saumya Dikshit for their review and comments.

8. References

8.1. Normative References

[I-D.ietf-bier-architecture]

Wijnands, I., Rosen, E., Dolganow, A., Przygienda, T., and S. Aldrin, "Multicast using Bit Index Explicit Replication", [draft-ietf-bier-architecture-05](#) (work in progress), October 2016.

[I-D.ietf-bier-mpls-encapsulation]

Wijnands, I., Rosen, E., Dolganow, A., Tantsura, J., Aldrin, S., and I. Meilik, "Encapsulation for Bit Index Explicit Replication in MPLS and non-MPLS Networks", [draft-ietf-bier-mpls-encapsulation-06](#) (work in progress), December 2016.

[I-D.ietf-nvo3-vxlan-gpe]

Maino, F., Kreeger, L., and U. Elzur, "Generic Protocol Extension for VXLAN", [draft-ietf-nvo3-vxlan-gpe-03](#) (work in progress), October 2016.

[I-D.ietf-sfc-nsh]

Quinn, P. and U. Elzur, "Network Service Header", [draft-ietf-sfc-nsh-10](#) (work in progress), September 2016.

[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.

[I-D.kumarzheng-bier-ping]

Kumar, N., Pignataro, C., Akiya, N., Zheng, L., Chen, M., and G. Mirsky, "BIER Ping and Trace", [draft-kumarzheng-bier-ping-03](#) (work in progress), July 2016.

[I-D.nordmark-nvo3-transcending-traceroute]

Nordmark, E., Appanna, C., Lo, A., Boutros, S., and A. Dubey, "Layer-Transcending Traceroute for Overlay Networks like VXLAN", [draft-nordmark-nvo3-transcending-traceroute-03](#) (work in progress), July 2016.

[I-D.xu-bier-encapsulation]

Xu, X., somasundaram.s@alcatel-lucent.com, s., Jacquenet, C., Raszuk, R., and Z. Zhang, "A Transport-Independent Bit Index Explicit Replication (BIER) Encapsulation Header", [draft-xu-bier-encapsulation-06](#) (work in progress), September 2016.

- [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>>.
- [RFC7348] Mahalingam, M., Dutt, D., Duda, K., Agarwal, P., Kreeger, L., Sridhar, T., Bursell, M., and C. Wright, "Virtual eXtensible Local Area Network (VXLAN): A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks", RFC 7348, DOI 10.17487/RFC7348, August 2014, <<http://www.rfc-editor.org/info/rfc7348>>.
- [RFC7365] Lasserre, M., Balus, F., Morin, T., Bitar, N., and Y. Rekhter, "Framework for Data Center (DC) Network Virtualization", RFC 7365, DOI 10.17487/RFC7365, October 2014, <<http://www.rfc-editor.org/info/rfc7365>>.
- [RFC7637] Garg, P., Ed. and Y. Wang, Ed., "NVGRE: Network Virtualization Using Generic Routing Encapsulation", RFC 7637, DOI 10.17487/RFC7637, September 2015, <<http://www.rfc-editor.org/info/rfc7637>>.
- [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>>.
- [RFC7799] Morton, A., "Active and Passive Metrics and Methods (with Hybrid Types In-Between)", RFC 7799, DOI 10.17487/RFC7799, May 2016, <<http://www.rfc-editor.org/info/rfc7799>>.

8.2. Informative References

- [I-D.ietf-bier-oam-requirements]
Mirsky, G., Nordmark, E., Pignataro, C., Kumar, N., Aldrin, S., Zheng, L., Chen, M., Akiya, N., and S. Pallagatti, "Operations, Administration and Maintenance (OAM) Requirements for Bit Index Explicit Replication (BIER) Layer", draft-ietf-bier-oam-requirements-02 (work in progress), September 2016.

Authors' Addresses

Nagendra Kumar
Cisco Systems, Inc.
7200 Kit Creek Road
Research Triangle Park, NC 27709
US

Email: naikumar@cisco.com

Carlos Pignataro
Cisco Systems, Inc.
7200 Kit Creek Road
Research Triangle Park, NC 27709-4987
US

Email: cpignata@cisco.com

Deepak Kumar
Cisco Systems, Inc.
3700 Cisco Way
SJ
US

Email: dekumar@cisco.com

Greg Mirsky
Ericsson

Email: gregory.mirsky@ericsson.com

Mach Chen
Huawei Technologies

Email: mach.chen@huawei.com

Eric Nordmark
Arista Networks

Email: nordmark@acm.org

Santosh Pallagatti
Juniper Networks

Email: santosh.pallagatti@gmail.com

David Mozes
Mellanox Technologies Ltd

Email: davidm@mellanox.com