Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: January 20, 2017

D. Dolson
K. Larose
Sandvine

July 19, 2016

OAM Mechanism for SFF-SF Connectivity Verification
draft-dolson-sfc-oam-sff-sf-cv-01

Abstract

This document describes a mechanism and packet format for allowing a Service Function Forwarder (SFF) to verify connectivity of an connected SFF or Service Function (SF).

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 20, 2017.

Table of Contents

1. Introduction

As described in Service Function Chaining (SFC) Architecture [RFC7665], Service Function Forwarders (SFFs) are responsible for forwarding traffic to connected Service Functions (SFs).

We believe there is a need for the SFFs to monitor connectivity to the SFs at the NSH layer. Rather than have the SFFs simply ping each SF's IP stack, we believe it is important to test NSH connectivity because the two protocols (IP and NSH) are often handled by different hardware or code modules.

As an example, an SFF may wish to health-check multiple connected SFs prior to load-balancing NSH traffic to the SFs, having the means to automatically remove unreachable SFs from service.

This document proposes an NSH OAM format allowing an SFF to request a neighboring SF to respond to an echo request via NSH encapsulation. This format can also be used for an SFF to request an neighboring SFF to respond to an echo request.

Note that this connectivity checking is NOT to be confused with path verification. It is in fact a feature of this mechanism that no path forwarding needs to be configured to perform the connectivity verification.

This document proposes use of the format of continuity check proposed in [I-D.ooamdt-rtgwg-demand-cc-cv] to be used within NSH frames for SFF-to-SF connectivity verification.

2. SFF-SF Connectivity Verification

An SFF may determine connectivity to an SF by means of echo request/response. However, any two NSH nodes could verify connectivity by the following mechanism.

By embedding the overlay ping packet format [I-D.ooamdt-rtgwg-demand-cc-cv] within the OAM header [I-D.ooamdt-rtgwg-ooam-header] in NSH, the packet format is that of Figure 1. The MD-type 2 is shown, since no metadata is required.

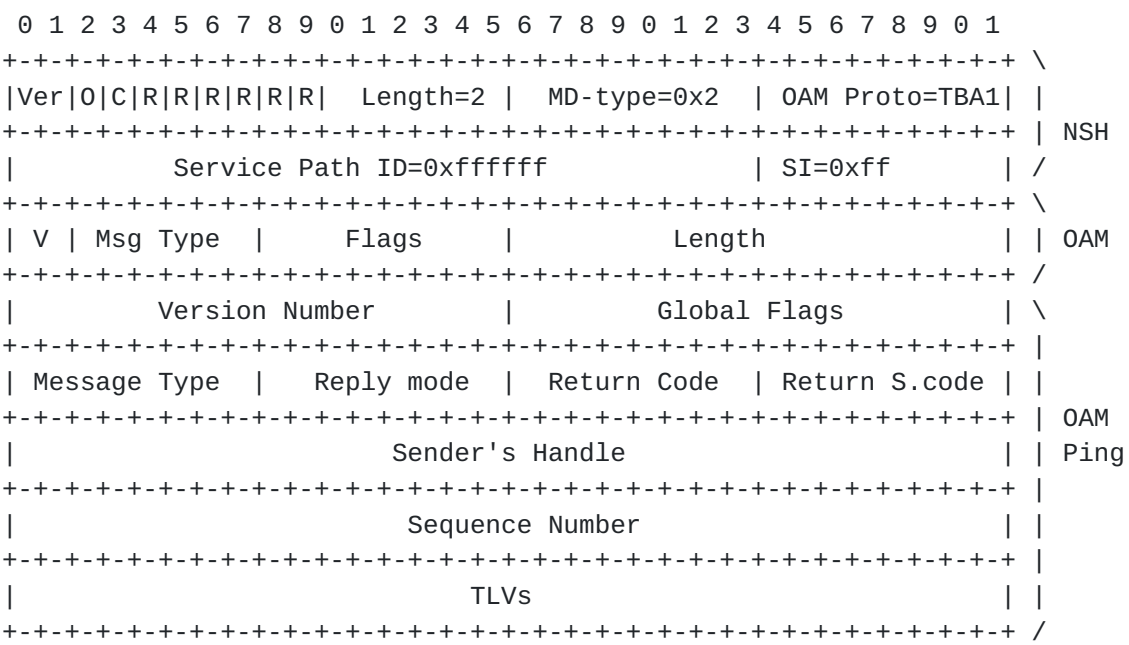

Figure 1: OAM Ping Encapsulated within NSH

The fields are:

Length: Header length in words

MD-type: metadata type of 2 is used because no metadata is required.

OAM Protocol: a value of TBA1 indicates the NSH header is followed by an OAM Header.

Service Path ID: Should be set by sender to 0xffffff. Receiver should ignore it.

Service Index: Should be set by sender to 255. Receiver should ignore it.

V: OAM header version should be set by sender to 0 (?)

Msg Type: value ? indicates OAM ping message follows the OAM header.

Flags: Value 0, indicating no extensions to the OAM header [I-D.ooamdt-rtgwg-ooam-header].

Length: Length of the OAM Ping portion of the message [I-D.ooamdt-rtgwg-ooam-header].

Version Number: refer to [I-D.ooamdt-rtgwg-demand-cc-cv]

Global Flags: refer to [I-D.ooamdt-rtgwg-demand-cc-cv]

Message Type: Initiator (SFF) is to use the code for "Overlay Echo Request" and responder (SF) is to use the code for "Overlay Echo Reply" [I-D.ooamdt-rtgwg-demand-cc-cv]

Reply Mode: code for "Reply to Sender" [I-D.ooamdt-rtgwg-demand-cc-cv] (requires extension)

Return Code: Unused at this time. MUST be ignored by initiator and responder.

Return Subcode: Unused at this time. MUST be ignored by initiator and responder.

Sender's Handle: an arbitrary handle chosen by the initiator, to be echoed by the responder. This value is generally constant across requests and may be useful for identifying the initiator in a debugging situation.

Sequence Number: a number chosen by the initiator, to be echoed by the responder. This value is generally incremeted by 1 between requests and may be useful for debugging packet loss counts.

TLVs: The initiator may place any TLVs. The responder MUST echo back the TLVs verbatim unless TLVs are specifically defined otherwise. No OAM TLVs are required for this connectivity verification. However, (a) timestamp TLVs are expected to be useful for the sender to measure round-trip time; (b) large padding TLVs are expected to be useful for verifying MTU of a connection.

3. Echo Request Responder Behavior

An NSH node receiving an echo request MUST do the following:

1. Clone the NSH packet and contents verbatim

2. Change the Message Type from "Overlay Echo Request" to "Overlay Echo Reply" [I-D.ooamdt-rtgwg-demand-cc-cv]

3. Send the NSH packet back to the initiator using the transport the echo request was received on.

Note that for this type of OAM packet, the NSH packet is NOT forwarded according to path ID and service index, rather MUST be returned to the immediate peer. The echo is expected to work even when SFF forwarding tables are empty or incomplete.

For example, an NSH packet transported directly over Ethernet MUST be returned to the MAC address from which it was received. As another example, an NSH packet received over UDP MUST be returned to the IP address and UDP port the were the source address and ports of the request.

It might not be possible to use this OAM packet if there is not an obvious way to return the packet to the sender.

4. Acknowledgements

Thanks to the Overlay OAM Design team and authors of [I-D.ooamdt-rtgwg-demand-cc-cv] for pointing us an approach in common with other overlays.

5. IANA Considerations

TODO: Need to request any codes, subcodes or TLVs?

6. Security Considerations

To reduce any attack surface, the initiator should verify that the received echo response is a response to the echo request it sent by comparing the handle and sequence number fields.

7. Normative References

[I-D.ooamdt-rtgwg-demand-cc-cv]
Mirsky, G., Nordmark, E., Pignataro, C., Kumar, N., Kumar, D., Chen, M., Yizhou, L., Mozes, D., and i. ibagdona@gmail.com, "On-demand Continuity Check (CC) and Connectivity Verification(CV) for Overlay Networks", draft-ooamdt-rtgwg-demand-cc-cv-00 (work in progress), July 2016.

[I-D.ooamdt-rtgwg-ooam-header]
Mirsky, G., Nordmark, E., Pignataro, C., Kumar, N., Kumar, D., Chen, M., Yizhou, L., Mozes, D., and i. ibagdona@gmail.com, "OAM Header for use in Overlay Networks", draft-ooamdt-rtgwg-ooam-header-00 (work in progress), July 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>.

[RFC7665] Halpern, J., Ed. and C. Pignataro, Ed., "Service Function Chaining (SFC) Architecture", RFC 7665, DOI 10.17487/RFC7665, October 2015, <http://www.rfc-editor.org/info/rfc7665>.

Authors' Addresses

David Dolson
Sandvine
408 Albert Street
Waterloo, ON N2L 3V3
Canada

Phone: +1 519 880 2400
Email: ddolson@sandvine.com

Kyle Larose
Sandvine
408 Albert Street
Waterloo, ON N2L 3V3
Canada

Phone: +1 519 880 2400
Email: klarose@sandvine.com