Workgroup: DetNet
Internet-Draft:
draft-ietf-detnet-oam-framework-00
Published: 25 April 2021
Intended Status: Standards Track
Expires: 27 October 2021
Authors: G. Mirsky ZTE Corp.
F. Theoleyre CNRS
G.Z. Papadopoulos IMT Atlantique
CJ. Bernardos UC3M

Framework of Operations, Administration and Maintenance (OAM) for Deterministic Networking (DetNet)

Abstract

Deterministic Networking (DetNet), as defined in RFC 8655, is aimed to provide a bounded end-to-end latency on top of the network infrastructure, comprising both Layer 2 bridged and Layer 3 routed segments. This document's primary purpose is to detail the specific requirements of the Operation, Administration, and Maintenance (OAM) recommended to maintain a deterministic network. With the implementation of the OAM framework in DetNet, an operator will have a real-time view of the network infrastructure regarding the network's ability to respect the Service Level Objective, such as packet delay, delay variation, and packet loss ratio, assigned to each data flow.

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 27 October 2021.

Table of Contents

1. Introduction

Deterministic Networking (DetNet) [RFC8655] has proposed to provide a bounded end-to-end latency on top of the network infrastructure, comprising both Layer 2 bridged and Layer 3 routed segments. Their work encompasses the data plane, OAM, time synchronization, management, control, and security aspects.

Operations, Administration, and Maintenance (OAM) Tools are of primary importance for IP networks [RFC7276]. DetNet OAM should provide a toolset for fault detection, localization, and performance measurement.

This document's primary purpose is to detail the specific requirements of the OAM features recommended to maintain a deterministic/reliable network. Specifically, it investigates the requirements for a deterministic network, supporting critical flows.

In this document, the term OAM will be used according to its definition specified in [RFC6291]. DetNet expects to implement an OAM framework to maintain a real-time view of the network infrastructure, and its ability to respect the Service Level Objectives (SLO), such as packet delay, delay variation, and packet loss ratio, assigned to each data flow.

This document lists the functional requirements toward OAM for DetNet domain. The list can further be used for gap analysis of available OAM tools to identify possible enhancements of existing or whether new OAM tools are required to support proactive and on-demand path monitoring and service validation.

1.1. Terminology

The following terms are used througout this document as defined below:

* OAM entity: a data flow to be monitored for defects and/or its performance metrics measured.

* Maintenance End Point (MEP): OAM systems traversed by a data flow when entering/exiting the network. In DetNet, it corresponds with the source and destination of a data flow. OAM messages can be exchanged between two MEPs.

* Maintenance Intermediate endPoint (MIP): an OAM system along the flow; a MIP MAY respond to an OAM message generated by the MEP.

* Control and management plane: the control and management planes are used to configure and control the network (long-term). Relative to a data flow, the control and/or management plane can be out-of-band.

* Active measurement methods (as defined in [RFC7799]) modify a normal data flow by inserting novel fields, injecting specially constructed test packets [RFC2544]). It is critical for the quality of information obtained using an active method that generated test packets are in-band with the monitored data flow. In other words, a test packet is required to cross the same

network nodes and links and receive the same Quality of Service (QoS) treatment as a data packet.

* Passive measurement methods [RFC7799] infer information by observing unmodified existing flows.

* Hybrid measurement methods [RFC7799] is the combination of elements of both active and passive measurement methods.

1.2. Acronyms

OAM: Operations, Administration, and Maintenance

DetNet: Deterministic Networking

SLO: Service Level Objective

QoS: Quality of Service

SNMP: Simple Network Management Protocol

SDN: Software Defined Network

<TODO> we need here an exhaustive list, to be completed after the document has evolved.

1.3. 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. Role of OAM in DetNet

DetNet networks expect to provide communications with predictable low packet delay and packet loss. Most critical applications will define an SLO to be required for the data flows it generates.

To respect strict guarantees, DetNet can use an orchestrator able to monitor and maintain the network. Typically, a Software-Defined Network (SDN) controller places DetNet flows in the deployed network based on their the SLO. Thus, resources have to be provisioned a priori for the regular operation of the network. OAM represents the essential elements of the network operation and necessary for OAM resources that need to be accounted for to maintain the network operational.

Fault-tolerance also assumes that multiple paths could be provisioned so that an end-to-end circuit is maintained by adapting to the existing conditions. The central controller/orchestrator typically controls the Packet Replication, Elimination, and Ordering Functions (PREOF) on a node. OAM is expected to support monitoring and troubleshooting PREOF on a particular node and within the domain.

Note that PREOF can also be controlled by a set of distributed controllers, in those scenarios where DetNet solutions involve more than one single central controller.

3. Operation

OAM features will enable DetNet with robust operation both for forwarding and routing purposes.

3.1. Information Collection

Information about the state of the network can be collected using several mechanisms. Some protocols, e.g., Simple Network Management Protocol (SNMP), send queries. Others, e.g., YANG-based data models, generate notifications based on the publish-subscribe method. In either way, information about the state of the network being collected and sent to the controller.

Also, we can characterize methods of transporting OAM information relative to the path of data. For instance, OAM information may be transported out-of-band or in-band with the data flow.

3.2. Continuity Check

Continuity check is used to monitor the continuity of a path, i.e., that there exists a way to deliver the packets between two endpoints A and B.

3.3. Connectivity Verification

In addition to the Continuity Check, DetNet solutions have to verify the connectivity. This verification considers additional constraints, i.e., the absence of misconnection.

In particular, resources have to be reserved for a given flow, so they are booked for use without being impacted by other flows. Similarly, the destination does not receive packets from different flows through its interface.

It is worth noting that the test and data packets MUST follow the same path, i.e., the connectivity verification has to be conducted in-band without impacting the data traffic. Test packets MUST share

fate with the monitored data traffic without introducing congestion in normal network conditions.

3.4. Route Tracing

Ping and traceroute are two ubiquitous tools that help localize and characterize a failure in the network. They help to identify a subset of the list of routers in the route. However, to be predictable, resources are reserved per flow in DetNet. Thus, DetNet needs to define route tracing tools able to track the route for a specific flow.

DetNet with IP data plane is NOT RECOMMENDED to use multiple paths or links, i.e., Equal-Cost Multipath (ECMP) [RFC8939]. As the result, OAM in IP ECMP environment is outside the scope of this document.

3.5. Fault Verification/detection

DetNet expects to operate fault-tolerant networks. Thus, mechanisms able to detect faults before they impact the network performance are needed.

The network has to detect when a fault occurred, i.e., the network has deviated from its expected behavior. While the network must report an alarm, the cause may not be identified precisely. For instance, the end-to-end reliability has decreased significantly, or a buffer overflow occurs.

DetNet OAM mechanisms SHOULD allow a fault detection in real time. They MAY, when possible, predict faults based on current network conditions. They MAY also identify and report the cause of the actual/predicted network failure.

3.6. Fault Isolation/identification

The network has isolated and identified the cause of the fault. For instance, the replication process behaves not as expected to a specific intermediary router.

3.7. Use of Hybrid OAM in DetNet

Hybrid OAM methods are used in performance monitoring and defined in [RFC7799] as:

> Hybrid Methods are Methods of Measurement that use a combination of Active Methods and Passive Methods.

A hybrid measurement method may produce metrics as close to passive, but it still alters something in a data packet even if that is the

value of a designated field in the packet encapsulation. One example of such a hybrid measurement method is the Alternate Marking method (AMM) described in [RFC8321]. One of the advantages of the use of AMM in a DetNet domain with the IP data plane is that the marking is applied to a data flow, thus ensuring that measured metrics are directly applicable to the DetNet flow.

4. Administration

The network SHOULD expose a collection of metrics to support an operator making proper decisions, including:

* Queuing Delay: the time elapsed between a packet enqueued and its transmission to the next hop.

* Buffer occupancy: the number of packets present in the buffer, for each of the existing flows.

The following metrics SHOULD be collected:

* per virtual circuit to measure the end-to-end performance for a given flow. Each of the paths has to be isolated in multipath routing strategies.

* per path to detect misbehaving path when multiple paths are applied.

* per device to detect misbehaving node, when it relays the packets of several flows.

4.1. Collection of metrics

DetNet OAM SHOULD optimize the number of statistics / measurements to collected, frequency of collecting. Distributed and centralized mechanisms MAY be used in combination. Periodic and event-triggered collection information characterizing the state of a network MAY be used.

4.2. Worst-case metrics

DetNet aims to enable real-time communications on top of a heterogeneous multi-hop architecture. To make correct decisions, the controller needs to know the distribution of packet losses/delays for each flow, and each hop of the paths. In other words, the average end-to-end statistics are not enough. The collected information must be sufficient to allow the controller to predict the worst-case.

5. Maintenance

DetNet needs to implement a self-healing and self-optimization approach. The controller MUST be able to continuously retrieve the state of the network, to evaluate conditions and trends about the relevance of a reconfiguration, quantifying:

the cost of the sub-optimality: resources may not be used optimally (e.g., a better path exists).

the reconfiguration cost: the controller needs to trigger some reconfigurations. For this transient period, resources may be twice reserved, and control packets have to be transmitted.

Thus, reconfiguration may only be triggered if the gain is significant.

5.1. Replication / Elimination

When multiple paths are reserved between two maintenance endpoints, packet replication may be used to introduce redundancy and alleviate transmission errors and collisions. For instance, in Figure 1, the source node S is transmitting the packet to both parents, nodes A and B. Each maintenance endpoint will decide to trigger the packet replication, elimination or the ordering process when a set of metrics passes a threshold value.

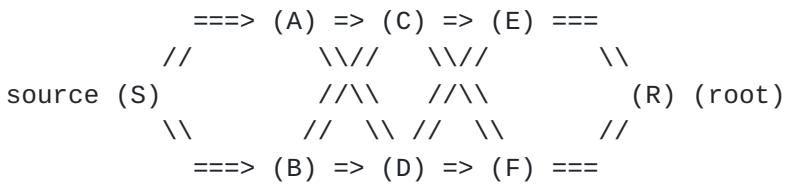

Figure 1: Packet Replication: S transmits twice the same data packet, to DP(A) and AP (B).

5.2. Resource Reservation

Because the QoS criteria associated with a path may degrade, the network has to provision additional resources along the path. We need to provide mechanisms to patch the network configuration.

5.3. Soft transition after reconfiguration

Since DetNet expects to support real-time flows, DetNet OAM MUST support soft-reconfiguration, where the novel resources are reserved before the ancient ones are released. Some mechanisms have to be proposed so that packets are forwarded through the novel track only

when the resources are ready to be used, while maintaining the global state consistent (no packet reordering, duplication, etc.)

6. Requirements

This section lists requirements for OAM in DetNet domain with MPLS data plane:

1. It MUST be possible to initiate DetNet OAM session from any DetNet node towards another DetNet node(s) within given domain.

2. It SHOULD be possible to initialize DetNet OAM session from a centralized controller.

3. DetNet OAM MUST support proactive and on-demand OAM monitoring and measurement methods.

4. DetNet OAM packets MUST be in-band, i.e., follow precisely the same path as DetNet data plane traffic.

5. DetNet OAM MUST support unidirectional OAM methods, continuity check, connectivity verification, and performance measurement.

6. DetNet OAM MUST support bi-directional OAM methods. Such OAM methods MAY combine in-band monitoring or measurement in the forward direction and out-of-bound notification in the reverse direction, i.e., from egress to ingress end point of the OAM test session.

7. DetNet OAM MUST support proactive monitoring of a DetNet node availability in the given DetNet domain.

8. DetNet OAM MUST support Path Maximum Transmission Unit discovery.

9. DetNet OAM MUST support Remote Defect Indication (RDI) notification to the DetNet node performing continuity checking.

10. DetNet OAM MUST support performance measurement methods.

11. DetNet OAM MAY support hybrid performance measurement methods.

12. DetNet OAM MUST support unidirectional performance measurement methods. Calculated performance metrics MUST include but are not limited to throughput, packet loss, delay and delay variation metrics. [RFC6374] provides excellent details on performance measurement and performance metrics.

13. DetNet OAM MUST support defect notification mechanism, like Alarm Indication Signal. Any DetNet node in the given DetNet

domain MAY originate a defect notification addressed to any subset of nodes within the domain.

14. DetNet OAM MUST support methods to enable survivability of the DetNet domain. These recovery methods MAY use protection switching and restoration.

15. DetNet OAM MUST support the discovery of Packet Replication, Elimination, and Order preservation sub-functions locations in the domain.

16. DetNet OAM MUST support testing of Packet Replication, Elimination, and Order preservation sub-functions in the domain.

17. DetNet OAM MUST support monitoring any sub-set of paths traversed through the DetNet domain by the DetNet flow.

7. IANA Considerations

This document has no actionable requirements for IANA. This section can be removed before the publication.

8. Security Considerations

This document lists the OAM requirements for a DetNet domain and does not raise any security concerns or issues in addition to ones common to networking and those specific to a DetNet discussed in [I-D.ietf-detnet-security].

9. Acknowledgments

TBD

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>.

10.2. Informative References

[I-D.ietf-detnet-security]

Grossman, E., Mizrahi, T., and A. J. Hacker, "Deterministic Networking (DetNet) Security Considerations", Work in Progress, Internet-Draft, draft-ietf-detnet-security-16, 2 March 2021, <https://tools.ietf.org/html/draft-ietf-detnet-security-16>.

[RFC2544] Bradner, S. and J. McQuaid, "Benchmarking Methodology for Network Interconnect Devices", RFC 2544, DOI 10.17487/RFC2544, March 1999, <https://www.rfc-editor.org/info/rfc2544>.

[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, <https://www.rfc-editor.org/info/rfc6291>.

[RFC6374] Frost, D. and S. Bryant, "Packet Loss and Delay Measurement for MPLS Networks", RFC 6374, DOI 10.17487/RFC6374, September 2011, <https://www.rfc-editor.org/info/rfc6374>.

[RFC7276] Mizrahi, T., Sprecher, N., Bellagamba, E., and Y. Weingarten, "An Overview of Operations, Administration, and Maintenance (OAM) Tools", RFC 7276, DOI 10.17487/RFC7276, June 2014, <https://www.rfc-editor.org/info/rfc7276>.

[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>.

[RFC8321] Fioccola, G., Ed., Capello, A., Cociglio, M., Castaldelli, L., Chen, M., Zheng, L., Mirsky, G., and T. Mizrahi, "Alternate-Marking Method for Passive and Hybrid Performance Monitoring", RFC 8321, DOI 10.17487/RFC8321, January 2018, <https://www.rfc-editor.org/info/rfc8321>.

[RFC8655] Finn, N., Thubert, P., Varga, B., and J. Farkas, "Deterministic Networking Architecture", RFC 8655, DOI 10.17487/RFC8655, October 2019, <https://www.rfc-editor.org/info/rfc8655>.

[RFC8939] Varga, B., Ed., Farkas, J., Berger, L., Fedyk, D., and S. Bryant, "Deterministic Networking (DetNet) Data Plane: IP", RFC 8939, DOI 10.17487/RFC8939, November 2020, <https://www.rfc-editor.org/info/rfc8939>.

Authors' Addresses

Greg Mirsky
ZTE Corp.

Email: gregimirsky@gmail.com, gregory.mirsky@ztetx.com

Fabrice Theoleyre
CNRS
300 boulevard Sebastien Brant - CS 10413
67400 Illkirch - Strasbourg
France

Phone: +33 368 85 45 33
Email: theoleyre@unistra.fr
URI: http://www.theoleyre.eu

Georgios Z. Papadopoulos
IMT Atlantique
Office B00 - 102A
2 Rue de la Châtaigneraie
35510 Cesson-Sévigné - Rennes
France

Phone: +33 299 12 70 04
Email: georgios.papadopoulos@imt-atlantique.fr

Carlos J. Bernardos
Universidad Carlos III de Madrid
Av. Universidad, 30
28911 Leganes, Madrid
Spain

Phone: +34 91624 6236
Email: cjbc@it.uc3m.es
URI: http://www.it.uc3m.es/cjbc/