

RAW
Internet-Draft
Intended status: Standards Track
Expires: October 13, 2020

F. Theoleyre
CNRS
G. Papadopoulos
IMT Atlantique
G. Mirsky
ZTE Corp.
April 11, 2020

Operations, Administration and Maintenance (OAM) features for RAW
draft-theoleyre-raw-oam-support-02

Abstract

Some critical applications may use a wireless infrastructure. However, wireless networks exhibit a bandwidth of several orders of magnitude lower than wired networks. Besides, wireless transmissions are lossy by nature; the probability that a packet cannot be decoded correctly by the receiver may be quite high. In these conditions, guaranteeing the network infrastructure works properly is particularly challenging, since we need to address some issues specific to wireless networks. This document lists the requirements of the Operation, Administration, and Maintenance (OAM) features recommended to construct a predictable communication infrastructure on top of a collection of wireless segments. This document describes the benefits, problems, and trade-offs for using OAM in wireless networks to achieve Service Level Objectives (SLO).

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 October 13, 2020.

Internet-Draft

OAM features for RAW

April 2020

Copyright Notice

Copyright (c) 2020 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	3
1.1.	Terminology	4
1.2.	Acronyms	4
1.3.	Requirements Language	4
2.	Role of OAM in RAW	5
3.	Operation	5
3.1.	Information Collection	5
3.2.	Continuity Check	5
3.3.	Connectivity Verification	6
3.4.	Route Tracing	6
3.5.	Fault Verification/detection	6
3.6.	Fault Isolation/identification	7
4.	Administration	7
4.1.	Collection of metrics	8
4.2.	Worst-case metrics	8
4.3.	Energy efficiency constraint	8
5.	Maintenance	9
5.1.	Multipath Routing	9
5.2.	Replication / Elimination	9
5.3.	Resource Reservation	10
5.4.	Soft transition after reconfiguration	10
6.	IANA Considerations	10
7.	Security Considerations	10
8.	Acknowledgments	10
9.	Informative References	10
	Authors' Addresses	11

1. Introduction

Reliable and Available Wireless (RAW) is an effort that extends DetNet to approach end-to-end deterministic performances over a network that includes scheduled wireless segments. In wired networks, many approaches to Quality of Service (QoS) tried to implement traffic differentiation so that routers handle differently each type of packets. However, this differentiated treatment was expensive for most applications.

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.

However, wireless networks create specific challenges. First of all, radio bandwidth is significantly lower than for wired networks. In these conditions, the volume of signaling messages has to be very limited. Even worse, wireless links are lossy: a layer 2 transmission may or may not be decoded correctly by the receiver, depending on a large set of parameters. Thus, providing high reliability through only wireless segments only is particularly challenging.

Last but not least, radio links present very unstable characteristics. If the wireless networks use an unlicensed band, packet losses are not anymore temporally and spatially independent. Typically, links may exhibit a very bursty characteristic, where several consecutive packets may be dropped. Thus, providing availability and reliability on top of the wireless infrastructure requires specific layer 3 mechanisms to counteract these bursty losses.

Operations, Administration, and Maintenance (OAM) Tools are of primary importance for IP networks [[RFC7276](#)]. It defines a toolset

for fault detection and isolation, and for performance measurement.

The main purpose of this document is to detail the specific requirements of the OAM features recommended to construct a predictable communication infrastructure on top of a collection of wireless segments. This document describes the benefits, problems, and trade-offs for using OAM in wireless networks to provide availability and predictability.

In this document, the term OAM will be used according to its definition specified in [[RFC6291](#)]. We expect to implement an OAM framework in RAW networks to maintain a real-time view of the network

infrastructure, and its ability to respect the Service Level Objectives (SLO), such as delay and reliability, assigned to each data flow.

[1.1.](#) Terminology

- o OAM entity: a data flow to be controlled;
- o Maintenance End Point (MEP): OAM devices crossed when entering/exiting the network. In RAW, it corresponds mostly to the source or destination of a data flow. OAM message can be exchanges between two MEPs;
- o Maintenance Intermediate end Point (MIP): OAM devices along the flow; OAM messages can be exchanged between a MEP and a MIP;
- o Defect: a temporary change in the network (e.g. a radio link which is broken due to a mobile obstacle);
- o Fault: a definite change which may affect the network performance, e.g. a node runs out of energy.

[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

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.

Theoleyre, et al.

Expires October 13, 2020

[Page 4]

Internet-Draft

OAM features for RAW

April 2020

2. Role of OAM in RAW

RAW networks expect to make the communications reliable and predictable on top of a wireless network infrastructure. Most critical applications will define an SLO to be required for the data flows it generates. RAW considers network plane protocol elements such as OAM to improve the RAW operation at the service and the forwarding sub-layers.

To respect strict guarantees, RAW relies on an orchestrator able to monitor and maintain the network. Typically, a Software Defined Network (SDN) controller is in charge of scheduling the transmissions in the deployed network, based on the radio link characteristics, SLO of the flows, the number of packets to forward. Thus, resources have to be provisioned a priori to handle any defect. OAM represents the core of the over provisioning process, and maintains the network operational by updating the schedule dynamically.

Fault-tolerance also assumes that multiple paths have to be provisioned so that an end-to-end circuit keeps on existing whatever the conditions. The replication/elimination processes (PREOF) on a node is typically controlled by the central controller/orchestrator.

OAM is in charge of controlling that PREOF is working properly on a node and within the domain.

To be energy-efficient, reserving some dedicated out-of-band resources for OAM seems idealistic, and only in-band solutions are considered here.

RAW supports both proactive and on-demand troubleshooting.

[3.](#) Operation

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

[3.1.](#) Information Collection

Several solutions (e.g., Simple Network Management Protocol (SNMP), YANG-based data models) are already in charge of collecting the statistics. That way, we can encapsulate these statistics in specific monitoring packets, to send them to the controller.

[3.2.](#) Continuity Check

We need to verify that two endpoints are connected. In other words, there exists "one" way to deliver the packets between two endpoints A and B.

[3.3.](#) Connectivity Verification

Additionally, to the Continuity Check, we have to verify the connectivity. This verification considers additional constraints, i.e., the absence of misconnection.

In particular, the resources have to be reserved by a given flow, and no packets from other flows steal the corresponding resources. Similarly, the destination does not receive packets from different flows through its interface.

It is worth noting that the control and data packets may not follow the same path, and the connectivity verification has to be conducted in-band without impacting the data traffic. Test packets must share the fate with the monitored data traffic without introducing

congestion in normal network conditions.

[3.4.](#) Route Tracing

Ping and traceroute are two very common tools for diagnostic. They help to identify a subset of the list of routers in the route. However, to be predictable, resources are reserved per flow in RAW. Thus, we need to define route tracing tools able to track the route for a specific flow.

Wireless networks are meshed by nature: we have many redundant radio links. These meshed networks are both an asset and a drawback: while several paths exist between two endpoints, we should choose the most efficient one(s), concerning specifically the reliability, and the delay.

Thus, multipath routing can be considered to make the network fault-tolerant. Even better, we can exploit the broadcast nature of wireless networks to exploit meshed multipath routing: we may have multiple Maintenance Intermediate Endpoints for each hop in the path. In that way, each Maintenance Intermediate Endpoint has several possible next hops in the forwarding plane. Thus, all the possible paths between two maintenance endpoints should be retrieved.

[3.5.](#) Fault Verification/detection

RAW expects to operate fault-tolerant networks. Thus, we need mechanisms able to detect faults, before they impact the network performance.

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.

[3.6.](#) Fault Isolation/identification

The network has isolated and identified the cause of the fault. For instance, the quality of a specific link has decreased, requiring more retransmissions, or the level of external interference has

locally increased.

[4.](#) Administration

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

- o Packet losses: the time-window average and maximum values of the number of packet losses have to be measured. Many critical applications stop to work if a few consecutive packets are dropped;
- o Received Signal Strength Indicator (RSSI) is a very common metric in wireless to denote the link quality. The radio chipset is in charge of translating a received signal strength into a normalized quality indicator;
- o Delay: the time elapsed between a packet generation / enqueueing and its reception by the next hop;
- o Buffer occupancy: the number of packets present in the buffer, for each of the existing flows.

These metrics should be collected:

- o 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;
- o per radio channel to measure, e.g., the level of external interference, and to be able to apply counter-measures (e.g. blacklisting)
- o per device to detect misbehaving node, when it relays the packets of several flows.

[4.1.](#) Collection of metrics

We have to minimize the number of statistics / measurements to exchange:

- o energy efficiency: low-power devices have to limit the volume of monitoring information since every bit consumes energy.
- o bandwidth: wireless networks exhibit a bandwidth significantly lower than wired, best-effort networks.
- o per-packet cost: it is often more expensive to send several packets instead of combining them in a single link-layer frame.

Thus, localized and centralized mechanisms have to be combined together, and additional control packets have to be triggered only after a fault detection.

[4.2.](#) Worst-case metrics

RAW aims to enable real-time communications on top of a heterogeneous architecture. Since wireless networks are known to be lossy, RAW has to implement strategies to improve reliability on top of unreliable links. Hybrid Automatic Repeat reQuest (ARQ) has typically to enable retransmissions based on the end-to-end reliability and latency requirements.

To make correct decisions, the controller needs to know the distribution of packet losses for each flow, and each hop of the paths. In other words, the average end-to-end statistics are not enough. They must allow the controller to predict the worst-case.

[4.3.](#) Energy efficiency constraint

RAW targets also low-power wireless networks, where energy represents a key constraint. Thus, we have to take care of power and bandwidth consumption. The following techniques aim to reduce the cost of such maintenance:

piggybacking: some control information are inserted in the data packets if they do not fragment the packet (i.e., the MTU is not exceeded). Information Elements represent a standardized way to handle such information;

flags/fields: we have to set-up flags in the packets to monitor to be able to monitor the forwarding process accurately. A sequence number field may help to detect packet losses. Similarly, path inference tools such as [[ipath](#)] insert additional information in

the headers to identify the path followed by a packet a posteriori.

5. Maintenance

RAW needs to implement a self-healing and self-optimization approach. The network must continuously retrieve the state of the network, to judge 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. Multipath Routing

To be fault-tolerant, several paths can be reserved between two maintenance endpoints. They must be node-disjoint so that a path can be available at any time.

5.2. Replication / Elimination

When multiple paths are reserved between two maintenance endpoints, they may decide to replicate the packets to introduce redundancy, and thus to 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 replication/elimination process when a set of metrics passes through a threshold value.




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

[5.3.](#) 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 a schedule (changing the channel offset, allocating more timeslots, changing the path, etc.).

[5.4.](#) Soft transition after reconfiguration

Since RAW expects to support real-time flows, we have to 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.](#) IANA Considerations

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

[7.](#) Security Considerations

This section will be expanded in future versions of the draft.

[8.](#) Acknowledgments

TBD

[9.](#) Informative References

- [ipath] Gao, Y., Dong, W., Chen, C., Bu, J., Wu, W., and X. Liu, "iPath: path inference in wireless sensor networks.", 2016, <<https://doi.org/10.1109/TNET.2014.2371459>>.
- [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>>.

- [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>>.

Theoleyre, et al. Expires October 13, 2020 [Page 10]

Internet-Draft OAM features for RAW April 2020

- [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>>.
- [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>>.
- [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>>.

Authors' Addresses

Fabrice Theoleyre
CNRS
Building B
300 boulevard Sebastien Brant - CS 10413
Illkirch - Strasbourg 67400
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 Chataigneraie
Cesson-Sevigne - Rennes 35510
FRANCE

Phone: +33 299 12 70 04

Email: georgios.papadopoulos@imt-atlantique.fr

Grek Mirsky
ZTE Corp.

Email: gregimirsky@gmail.com