

RAW
Internet-Draft
Intended status: Standards Track
Expires: May 6, 2020

F. Theoleyre
CNRS
G. Papadopoulos
IMT Atlantique
November 3, 2019

Operations, Administration and Maintenance (OAM) features for RAW
draft-theoleyre-raw-oam-support-01

Abstract

The wireless medium presents significant specific challenges to achieve properties similar to those of wired deterministic networks. At the same time, a number of use cases cannot be solved with wires and justify the extra effort of going wireless. This document presents some of these use-cases.

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 May 6, 2020.

Copyright Notice

Copyright (c) 2019 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	2
1.1. Terminology	3
2. Needs for OAM in RAW	3
3. Operation	4
3.1. Connectivity Verification	4
3.2. Route Tracing	4
3.3. Fault verification / detection	4
3.4. Fault isolation / identification	5
4. Administration	5
4.1. Worst-case metrics	6
4.2. Energy efficiency constraint	6
5. Maintenance	6
5.1. Multipath	7
5.2. Replication / Elimination	7
5.3. Resource Reservation	7
5.4. Soft transition after reconfiguration	7
6. Informative References	8
Authors' Addresses	8

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. The wireless and wired media are fundamentally different at the physical level. Enabling thus reliable and available wireless communications is even more challenging than it is in wired IP networks, due to the numerous causes of loss in transmission that add up to the congestion losses and the delays caused by overbooked shared resources. To provide quality of service along a multihop path that is composed of wired and wireless hops, additional methods needs to be considered to leverage the potential lossy wireless communication.

Traceability belongs to Operations, Administration, and Maintenance (OAM) which is the toolset for fault detection and isolation, and for performance measurement. More can be found on OAM Tools in [RFC7276].

The main purpose of this document is to detail the 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 Agreements (SLA), such as delay and reliability, assigned to each data flow.

1.1. Terminology

- o OAM entity: a data flow to be controlled;
- o OAM end-devices: the source or destination of a data flow;
- o defect: a temporary change in the network characteristics (e.g. link quality degradation because of temporary external interference, a mobile obstacle)
- o fault: a definite change which may affect the network performance, e.g. a node runs out of energy,

2. Needs for 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 a SLA to respect for the data flows it generates. RAW considers network plane protocol elements such as OAM to improve the RAW operation at the service and at the forwarding sub-layers.

To respect strict guarantees, RAW relies on a Path Computation Element (PCE) which will be responsible to schedule the transmissions in the deployed network. 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 path have to be provisioned so that an end-to-end circuit keeps on existing whatever the conditions. OAM is in charge of controlling the replication/elimination processes.

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. Connectivity Verification

We need to verify that two endpoints are connected with each other. Since we reserve resources along the path independently for each flow, we must be able to verify that the path exists for a given flow label.

The control and data packets may not follow the same path, and the connectivity verification has to be triggered in-band without impacting the data traffic. In particular, the control plane may work while the data plane may be broken.

The ping packets must be labeled in the same way as the data packets of the flow to monitor.

3.2. Route Tracing

Ping and traceroute are two very common tools for diagnostic. They help to identify 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.

Because the network has to be fault-tolerant, multipath can be considered, with multiple Maintenance Intermediate Endpoints for each hop in the path. Thus, all the possible paths between two maintenance endpoints should be retrieved.

3.3. 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.

We have to minimize the amount 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: 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.

3.4. 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

To take proper decisions, the network has to expose a collection of metrics, including:

- o Packet losses: the time-window average and maximum values of the number of packet losses has 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 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. Worst-case metrics

RAW aims to enable real-time communications on top of an heterogeneous architecture. Since wireless networks are known to be lossy, RAW has to implement strategies to improve the 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 take correct decisions, the controller needs to know the distribution of packet losses for each flow, and for each hop of the paths. In other words, average end-to-end statistics are not enough. They must allow the controller to predict the worst-case.

4.2. Energy efficiency constraint

RAW targets also low-power wireless networks, where energy represents a key constraint. Thus, we have to take care of the energy 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

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.

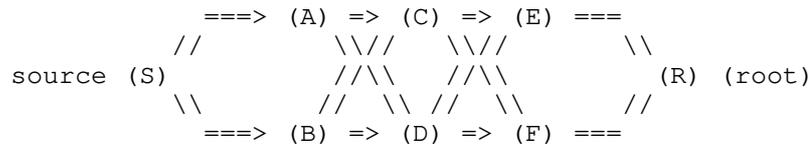


Figure 1: Packet Replication: S transmits twice the same data packet, to its DP (A) and to its AP (B).

5.3. Resource Reservation

Because the QoS criteria associated to 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 re-ordering, duplication, etc.)

6. 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>>.
- [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>>.
- [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>>.

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