RAW
Internet-Draft
Updates: draft-ietf-raw-oam-support-00 (if approved)
Intended status: Informational
Expires: November 25, 2021

F. Theoleyre
CNRS
G. Papadopoulos
IMT Atlantique
G. Mirsky
ZTE Corp.
CJ. Bernardos
UC3M

May 24, 2021

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

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 that 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 November 25, 2021.

Table of Contents

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 try to enable Quality of Service (QoS) by implementing traffic differentiation so that routers handle each type of packets differently. 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 broad set of parameters. Thus, providing high reliability through wireless segments is particularly challenging.

Wired networks rely on the concept of _links_. All the devices attached to a link receive any transmission. The concept of a link in wireless networks is somewhat different from what many are used to in wireline networks. A receiver may or may not receive a transmission, depending on the presence of a colliding transmission, the radio channel's quality, and the external interference. Besides, a wireless transmission is broadcast by nature: any _neighboring_ device may be able to decode it. The document includes detailed information on what the implications for the OAM features are.

Last but not least, radio links present volatile 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, isolation, and performance measurement.

The primary 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.

1.1. Terminology

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.

We re-use here the same terminology as [detnet-oam]:

- OAM entity: a data flow to be monitored for defects and/or its performance metrics measured.;
- 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 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/management/data plane: the control and management planes are used to configure and control the network (long-term). The data plane takes the individual decision. 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. Active methods may implement one of these two strategies:
  - In-band: control information follows the same path as the data packets. In other words, a failure in the data plane may prevent the control information to reach the destination (e.g., end-device or controller).

  * out-of-band: control information is sent separately from the data packets. Thus, the behavior of control vs. data packets may differ;

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

We also adopt the following terminology, which is particularly relevant for RAW segments.

- piggybacking vs. dedicated control packets: control information may be encapsulated in specific (dedicated) control packets. Alternatively, it may be piggybacked in existing data packets, when the MTU is larger than the actual packet length. Piggybacking makes specifically sense in wireless networks, as the cost (bandwidth and energy) is not linear with the packet size.

- router-over vs. mesh under: a control packet is either forwarded directly to the layer-3 next hop (mesh under) or handled hop-by-hop by each router. While the latter option consumes more resources, it allows to collect additionnal intermediary information, particularly relevant in wireless networks.

- Defect: a temporary change in the network (e.g., a radio link which is broken due to a mobile obstacle);

- Fault: a definite change which may affect the network performance, e.g., a node runs out of energy.

- End-to-end delay: the time between the packet generation and its reception by the destination.

1.2. Acronyms

OAM Operations, Administration, and Maintenance

DetNet Deterministic Networking

PSE Path Selection Engine [I-D.pthubert-raw-architecture]

QoS Quality of Service

RAW Reliable and Available Wireless

SLO Service Level Objective

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.

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 the Path Selection Engine (PSE) (as defined in [I-D.pthubert-raw-architecture] to monitor and maintain the network. As an example, a Software-Defined Network (SDN) controller may be used to schedule 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 pre-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 Packet Replication and Elimination Function ([PREF-draft]) on a node is typically controlled by the PSE. OAM mechanisms can be used to monitor that PREOF is working correctly 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.

The specific characteristics of RAW are discussed below.

2.1. Link concept and quality

In wireless networks, a _link_ does not exist physically. A device has a set of *neighbors* that correspond to all the devices that have a non null probability of receiving correctly its packets. We make a distinction between:

- point-to-point (p2p) link with one transmitter and one receiver. These links are used to transmit unicast packets.

- point-to-multipoint (p2m) link associates one transmitter and a collection of receivers. For instance, broadcast packets assume the existence of p2m links to avoid duplicating a broadcast packet to reach each possible radio neighbor.

In scheduled radio networks, p2m and p2p links are commonly not scheduled simultaneously to save energy. More precisely, only one part of the neighbors may wake-up at a given instant.

Anycast are used in p2m links to improve the reliability. A collection of receivers are scheduled to wake-up simutaneously, so that the transmission fails only if none of the receivers is able to decode the packet.

Each wireless link is associated with a link quality, often measured as the Packet Delivery Ratio (PDR), i.e., the probability that the receiver can decode the packet correctly. It is worth noting that this link quality depends on many criteria, such as the level of external interference, the presence of concurrent transmissions, or the radio channel state. This link quality is even time-variant. For p2m links, we have consequently a collection of PDR (one value per receiver). Other more sophisticated, aggregated metrics exist for these p2m links, such as [anycast-property]

2.2. Broadcast Transmissions

In modern switching networks, the unicast transmission is delivered uniquely to the destination. Wireless networks are much closer to the ancient *shared access* networks. Practically, unicast and broadcast frames are handled similarly at the physical layer. The link layer is just in charge of filtering the frames to discard irrelevant receptions (e.g., different unicast MAC address).

However, contrary to wired networks, we cannot be sure that a packet is received by *all* the devices attached to the Layer 2 segment. It depends on the radio channel state between the transmitter(s) and the receiver(s). In particular, concurrent transmissions may be possible or not, depending on the radio conditions (e.g., do the different

transmitters use a different radio channel or are they sufficiently spatially separated?)

2.3. Complex Layer 2 Forwarding

Multiple neighbors may receive a transmission. Thus, anycast Layer 2 forwarding helps to maximize the reliability by assigning multiple receivers to a single transmission. That way, the packet is lost only if *none* of the receivers decode it. Practically, it has been proven that different neighbors may exhibit very different radio conditions, and that reception independency may hold for some of them [anycast-property].

2.4. End-to-end delay

In a wireless network, additionnal transmissions opportunities are provisionned to accomodate packet losses. Thus, the end-to-end delay consists of:

- Transmission delay, which is fixed and depends mainly on the data rate, and the presence or absence of an acknowledgement.

- Residence time, corresponds to the buffering delay and depends on the schedule. To account for retransmisisons, the residence time is equal to the difference between the time of last reception from the previous hop (among all the retransmisions) and the time of emission of the last retransmission.

3. Operation

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

3.1. Information Collection

The model to exchange information should be the same as for DetNet network, for the sake of inter-operability. YANG may typically fulfill this objective.

However, RAW networks imply specific constraints (e.g., low bandwidth, packet losses, cost of medium access) that may require to minimize the volume of information to collect. Thus, we discuss in Section 4.2 different ways to collect information, i.e., transfer physically the OAM information from the emitter to the receiver.

3.2. Continuity Check

Similarly to DetNet, we need to verify that the source and the destination are connected (at least one valid path exists)

3.3. Connectivity Verification

As in DetNet, we have to verify the absence of misconnection. We focus here on the RAW specificities.

Because of radio transmissions' broadcast nature, several receivers may be active at the same time to enable anycast Layer 2 forwarding. Thus, the connectivity verification must test any combination. We also consider priority-based mechanisms for anycast forwarding, i.e., all the receivers have different probabilities of forwarding a packet. To verify a delay SLO for a given flow, we must also consider all the possible combinations, leading to a probability distribution function for end-to-end transmissions. If this verification is implemented naively, the number of combinations to test may be exponential and too costly for wireless networks with low bandwidth.

3.4. Route Tracing

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, and 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 (MIP) 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, which may quickly become untractable if we apply a naive approach.

3.5. Fault Verification/detection

Wired networks tend to present stable performances. On the contrary, wireless networks are time-variant. We must consequently make a distinction between _normal_ evolutions and malfunction.

3.6. Fault Isolation/identification

The network has isolated and identified the cause of the fault. While DetNet already expects to identify malfunctions, some problems are specific to wireless networks. We must consequently collect metrics and implement algorithms tailored for wireless networking.

For instance, the decrease in the link quality may be caused by several factors: external interference, obstacles, multipath fading, mobility. It it fundamental to be able to discriminate the different causes to make the right decision.

4. Administration

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

- 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;

- 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;

- Delay: the time elapsed between a packet generation / enqueuing and its reception by the next hop;

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

- Battery lifetime: the expected remaining battery lifetime of the device. Since many RAW devices might be battery powered, this is an important metric for an operator to take proper decisions.

- Mobility: if a device is known to be mobile, this might be considered by an operator to take proper decisions.

These metrics should be collected per device, virtual circuit, and path, as detnet already does. However, we have to face in RAW to a finer granularity:

- per radio channel to measure, e.g., the level of external interference, and to be able to apply counter-measures (e.g., blacklisting).

- per link to detect misbehaving link (assymetrical link, fluctuating quality).

- per resource block: a collision in the schedule is particularly challenging to identify in radio networks with spectrum reuse. In particular, a collision may not be systematic (depending on the radio characteristics and the traffic profile)

4.1. Worst-case metrics

RAW inherits the same requirements as DetNet: we need to know the distribution of a collection of metrics. However, wireless networks are known to be highly variable. Changes may be frequent, and may exhibit a periodical pattern. Collecting and analyzing this amount of measurements is challenging.

Wireless networks are known to be lossy, and 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.

4.2. Efficient data retrieval

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

- energy efficiency: low-power devices have to limit the volume of monitoring information since every bit consumes energy.

- bandwidth: wireless networks exhibit a bandwidth significantly lower than wired, best-effort networks.

- per-packet cost: it is often more expensive to send several packets instead of combining them in a single link-layer frame.

In conclusion, we have to take care of power and bandwidth consumption. The following techniques aim to reduce the cost of such maintenance:

on-path collection: some control information is 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.

hierarchical monitoring; localized and centralized mechanisms have to be combined together. Typically, a local mechanism should contiuously monitor a set of metrics and trigger distant OAM exchances only when a fault is detected (but possibly not identified). For instance, local temporary defects must not trigger expensive OAM transmissions.

4.3. Reporting OAM packets to the source

TODO: statistics are collected when a packet goes from the source to the destination. However, it has to be also reported by the source. Problem: resource may not be reserved bidirectionnaly. Even worse: the inverse path may not exist.

5. Maintenance

Maintenance needs to facilitate the maintenance (repairs and upgrades). In wireless networks, repairs are expected to occur much more frequently, since the link quality may be highly time-variant. Thus, maintenance represents a key feature for RAW.

5.1. Soft transition after reconfiguration

Because of the wireless medium, the link quality may fluctuate, and the network needs to reconfigure itself continuously. During this transient state, flows may begin to be gradually re-forwarded, consuming resources in different parts of the network. OAM has to make a distinction between a metric that changed because of a legal network change (e.g., flow redirection) and an unexpected event (e.g., a fault).

5.2. Predictive maintenance

RAW needs to implement self-optimization features. While the network is configured to be fault-tolerant, a reconfiguration may be required to keep on respecting long-term objectives. Obviously, the network keeps on respecting the SLO after a node's crash, but a reconfiguration is required to handle the future faults. In other words, the reconfiguration delay MUST be strictly smaller than the inter-fault time.

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.

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

[anycast-property]
Teles Hermeto, R., Gallais, A., and F. Theoleyre, "Is Link-Layer Anycast Scheduling Relevant for IEEE 802.15.4-TSCH Networks?", 2019, <https://doi.org/10.1109/LCNSymposium47956.2019.9000679>.

[detnet-oam]
Theoleyre, F., Papadopoulos, G. Z., Mirsky, G., and C. J. Bernardos, "Operations, Administration and Maintenance (OAM) features for detnet", 2020, <https://tools.ietf.org/html/draft-theoleyre-detnet-oam-support>.

[I-D.pthubert-raw-architecture]
Thubert, P., Papadopoulos, G. Z., and R. Buddenberg, "Reliable and Available Wireless Architecture/Framework", draft-pthubert-raw-architecture-05 (work in progress), November 2020.

[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>.

[PREF-draft] Thubert, P., Eckert, T., Brodard, Z., and H. Jiang, "BIER-TE extensions for Packet Replication and Elimination Function (PREF) and OAM", 2018, <https://tools.ietf.org/html/draft-thubert-bier-replication-elimination>.

[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>.

[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>.

[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>.

[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

Greg Mirsky
ZTE Corp.

Email: gregory.mirsky@ztetx.com

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

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