

Workgroup: RAW

Published: 4 July 2022

Intended Status: Informational

Expires: 5 January 2023

Authors: P. Thubert, Ed. G.Z. Papadopoulos

 Cisco Systems IMT Atlantique

Reliable and Available Wireless Architecture

Abstract

Reliable and Available Wireless (RAW) provides for high reliability and availability for IP connectivity across any combination of wired and wireless network segments. The RAW Architecture extends the DetNet Architecture and other standard IETF concepts and mechanisms to adapt to the specific challenges of the wireless medium. This document defines an architecture element for the RAW data plane, in the form of an OODA loop, that optimizes the use of constrained spectrum and energy while maintaining the expected connectivity properties. The loop involves OAM, PCE, and PREOF extensions, and a new component called the Path Selection Engine (PSE).

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 5 January 2023.

Copyright Notice

Copyright (c) 2022 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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

- [1. Introduction](#)
- [2. The RAW problem](#)
 - [2.1. Terminology](#)
 - [2.1.1. Acronyms](#)
 - [2.1.2. Link and Direction](#)
 - [2.1.3. Path and Tracks](#)
 - [2.1.4. Deterministic Networking](#)
 - [2.1.5. Reliability and Availability](#)
 - [2.1.6. OAM variations](#)
 - [2.2. Reliability and Availability](#)
 - [2.2.1. High Availability Engineering Principles](#)
 - [2.2.2. Applying Reliability Concepts to Networking](#)
 - [2.2.3. Wireless Effects Affecting Reliability](#)
 - [2.3. Routing Time Scale vs. Forwarding Time Scale](#)
- [3. The RAW Conceptual Model](#)
 - [3.1. The RAW Planes](#)
 - [3.2. RAW vs. DetNet](#)
- [4. The OODA Loop](#)
 - [4.1. Observe: The RAW OAM](#)
 - [4.2. Orient: The Path Computation Engine](#)
 - [4.3. Decide: The Path Selection Engine](#)
 - [4.4. Act: The PAREO Functions](#)
 - [4.4.1. Packet Replication](#)
 - [4.4.2. Packet Elimination](#)
 - [4.4.3. Promiscuous Overhearing](#)
 - [4.4.4. Constructive Interference](#)
- [5. Security Considerations](#)
 - [5.1. Layer-2 encryption](#)
 - [5.2. Forced Access](#)
- [6. IANA Considerations](#)
- [7. Contributors](#)
- [8. Acknowledgments](#)
- [9. References](#)
 - [9.1. Normative References](#)
 - [9.2. Informative References](#)
- [Authors' Addresses](#)

1. Introduction

Deterministic Networking is an attempt to emulate the properties of a serial link over a switched fabric, by providing a bounded latency and eliminating congestion loss, even when co-existing with best-effort traffic. It is getting traction in various industries

including professional A/V, manufacturing, online gaming, and smartgrid automation, enabling cost and performance optimizations (e.g., vs. loads of P2P cables).

Bringing determinism in a packet network means eliminating the statistical effects of multiplexing that result in probabilistic jitter and loss. This can be approached with a tight control of the physical resources to maintain the amount of traffic within a budgeted volume of data per unit of time that fits the physical capabilities of the underlying network, and the use of time-shared resources (bandwidth and buffers) per circuit, and/or by shaping and/or scheduling the packets at every hop.

This innovation was initially introduced on wired networks, with IEEE 802.1 Time Sensitive networking (TSN) - for Ethernet LANs - and IETF DetNet. But the wired and the wireless media are fundamentally different at the physical level and in the possible abstractions that can be built for IPv6 [[IPv6](#)], more in [[IPoWIRELESS](#)]. Nevertheless, deterministic capabilities are required in a number of wireless use cases as well [[RAW-USE-CASES](#)]. With new scheduled radios such as TSCH and OFDMA [[RAW-TECHNOS](#)] being developed to provide determinism over wireless links at the lower layers, providing DetNet capabilities is now becoming possible.

Wireless networks operate on a shared medium where uncontrolled interference, including the self-induced multipath fading cause random transmission losses. Fixed and mobile obstacles and reflectors may block or alter the signal, causing transient and unpredictable variations of the throughput and packet delivery ratio (PDR) of a wireless link. This adds new dimensions to the statistical effects that affect the quality and reliability of the link. Multiple links and transmissions must be used, and the challenge is to provide enough diversity and redundancy to ensure the timely packet delivery while preserving energy and optimizing the use of the shared spectrum.

Reliable and Available Wireless (RAW) takes up the challenge of providing highly available and reliable end-to-end performances in a network with scheduled wireless segments. To defeat those additional causes of transmission delay and loss in wireless transmission, RAW requires and leverages deterministic layer-2 capabilities. Operating at the layer-3, RAW can further increase diversity in the spatial, time, code, and frequency domains by enabling multiple link-layer wired and wireless technologies in parallel or sequentially, for a higher resilience and a wider applicability. RAW can also provide homogeneous services to critical applications beyond the boundaries of a single subnetwork, e.g., controlling the use of diverse radio access technologies to optimize the end-to-end application experience.

While the generic ["Deterministic Networking Problem Statement"](#) [RFC8557] applies to both the wired and the wireless media, the methods to achieve RAW must extend those used to support time-sensitive networking over wires, as a RAW solution has to address less consistent transmissions, energy conservation and shared spectrum efficiency.

RAW provides DetNet elements that are specialized for transporting IP flows over deterministic radios technologies such as listed in [RAW-TECHNOS]. Conceptually, RAW is agnostic to the radio layer underneath though the capability to schedule transmissions is assumed. How the PHY is programmed to do so, and whether the radio is single-hop or meshed, are unknown at the IP layer and not part of the RAW abstraction. Nevertheless, cross-layer optimizations may take place to ensure proper link awareness (think, link quality) and packet handling (think, scheduling).

The ["Deterministic Networking Architecture"](#) [RFC8655] is composed of three planes: the Application (User) Plane, the Controller Plane, and the Network Plane. The DetNet Network Plane is composed of a DetNet service sub-layer that focuses on flow protection (e.g., using redundancy) and can be fully operated at layer-3, and a DetNet forwarding sub-layer that associates the flows to the paths, ensures the availability of the necessary resources, and leverages layer-2 functionalities for timely delivery to the next DetNet system.

The RAW Architecture extends the DetNet Network Plane, to accommodate one or multiple hops of homogeneous or heterogeneous wired and wireless technologies. RAW adds reactivity to the DetNet service sub-layer to compensate the dynamics for the radio links in terms of lossiness and bandwidth. This may apply for instance to mesh networks as illustrated in [Figure 3](#), or diverse radio access networks as illustrated in [Figure 8](#).

RAW and DetNet route application flows that require a special treatment along the paths that will provide that treatment. This may be seen as a form of Path Aware Networking and may be subject to impediments documented in [RFC9049].

The establishment of a path is not in-scope for RAW. It may be the product of a centralized Controller Plane as described for DetNet. As opposed to wired networks, the action of installing a path over a set of wireless links may be very slow relative to the speed at which the radio conditions vary, and it makes sense in the wireless case to provide redundant forwarding solutions along a complex path (see [Section 2.1.3](#)) and to leave it to the Network Plane to select which of those forwarding solutions are to be used for a given packet based on the current conditions.

RAW distinguishes the longer time scale at which routes are computed from the shorter forwarding time scale where per-packet decisions are made. RAW operates within the Network Plane at the forwarding time scale on one DetNet flow over a complex path delineated by a Track (see [Section 2.1.3.2](#)). The Track is preestablished and installed by means outside of the scope of RAW; it may be strict or loose depending on whether each or just a subset of the hops are observed and controlled by RAW.

The RAW Architecture is based on an abstract OODA Loop (Observe, Orient, Decide, Act). The generic concept involves:

1. Network Plane measurement protocols for Operations, Administration and Maintenance (OAM) to Observe some or all hops along a Track as well as the end-to-end packet delivery
2. Optional Controller plane elements to reports the links statistics to a Path computation Element (PCE) in a centralized controller that computes and installs the Tracks and provides meta data to Orient the routing decision
3. A Runtime distributed Path Selection Engine (PSE) that Decides which subTrack to use for the next packet(s) that are routed along the Track
4. Packet (hybrid) ARQ, Replication, Elimination and Ordering Dataplane actions that operate at the DetNet Service Layer to increase the reliability of the end-to-end transmissions. The RAW architecture also covers in-situ signaling when the decision is Acted by a node that down the Track from the PSE.

The overall OODA Loop optimizes the use of redundancy to achieve the required reliability and availability Service Level Agreement (SLA) while minimizing the use of constrained resources such as spectrum and battery.

This document presents the RAW problem and associated terminology in [Section 2](#), and elaborates in [Section 4](#) on the OODA loop based on the RAW conceptual model presented in [Section 3](#).

2. The RAW problem

2.1. Terminology

RAW reuses terminology defined for DetNet in the "[Deterministic Networking Architecture](#)" [[RFC8655](#)], e.g., PREOF for Packet Replication, Elimination and Ordering Functions.

RAW also reuses terminology defined for 6TiSCH in [[6TiSCH-ARCHI](#)] such as the term Track. A Track associates a complex path with PAREO

and shaping operations. The concept is agnostic to the underlaying technology and applies but is not limited to any fully or partially wireless mesh. RAW specifies strict and loose Tracks depending on whether the path is fully controlled by RAW or traverses an opaque network where RAW cannot observe and control the individual hops.

RAW uses the following terminology and acronyms:

2.1.1. Acronyms

2.1.1.1. ARQ

Automatic Repeat Request, enabling an acknowledged transmission and retries. ARQ is a typical model at Layer-2 on a wireless medium. ARQ is typically implemented hop-by-hop and not end-to-end in wireless networks. Else, it introduces excessive indetermination in latency, but a limited number of retries within a bounded time may be used within end-to-end constraints.

2.1.1.2. OAM

OAM stands for Operations, Administration, and Maintenance, and covers the processes, activities, tools, and standards involved with operating, administering, managing and maintaining any system. This document uses the terms Operations, Administration, and Maintenance, in conformance with the ['Guidelines for the Use of the "OAM" Acronym in the IETF'](#) [[RFC6291](#)] and the system observed by the RAW OAM is the Track.

2.1.1.3. OODA

Observe, Orient, Decide, Act. The OODA Loop is a conceptual cyclic model developed by USAF Colonel John Boyd, and that is applicable in multiple domains where agility can provide benefits against brute force.

2.1.1.4. PAREO

Packet (hybrid) ARQ, Replication, Elimination and Ordering. PAREO is a superset Of DetNet's PREOF that includes radio-specific techniques such as short range broadcast, MUMIMO, PHY rate and other Modulation Coding Scheme (MCS) adaptation, constructive interference and overhearing, which can be leveraged separately or combined to increase the reliability.

2.1.2. Link and Direction

2.1.2.1. Flapping

In the context of RAW, a link flaps when the reliability of the wireless connectivity drops abruptly for a short period of time, typically of a subsecond to seconds duration.

2.1.2.2. Uplink

Connection from end-devices to a data communication equipment. In the context of wireless, uplink refers to the connection between a station (STA) and a controller (AP) or a User Equipment (UE) to a Base Station (BS) such as a 3GPP 5G gNodeB (gNb).

2.1.2.3. Downlink

The reverse direction from uplink.

2.1.2.4. Downstream

Following the direction of the flow data path along a Track.

2.1.2.5. Upstream

Against the direction of the flow data path along a Track.

2.1.3. Path and Tracks

2.1.3.1. Path

Quoting section 1.1.3 of [[INT-ARCHI](#)]:

At a given moment, all the IP datagrams from a particular source host to a particular destination host will typically traverse the same sequence of gateways. We use the term "path" for this sequence. Note that a path is uni-directional; it is not unusual to have different paths in the two directions between a given host pair.

Section 2 of [[I-D.irtf-panrg-path-properties](#)] points to a longer, more modern definition of path, which begins as follows:

A sequence of adjacent path elements over which a packet can be transmitted, starting and ending with a node. A path is unidirectional. Paths are time-dependent, i.e., the sequence of path elements over which packets are sent from one node to another may change. A path is defined between two nodes.

It follows that the general acceptance of a path is a linear sequence of links and nodes, as opposed to a multi-dimensional

graph, defined by the experience of the packet that went from a node A to a node B.

With DetNet and RAW, a packet may be duplicated, fragmented and network-coded, and the various byproducts may travel different paths that are not necessarily end-to-end between A and B; we refer to that experience as a complex path. The complex path does not fit the traditional description of a path, and is subject to change from a packet to the next. Therefore we introduce below the term of a Track as the overall topology where the possible complex paths are all contained.

In the context of this document, a path is observed by following one copy or one fragment of a packet that conserves its uniqueness and integrity. For instance, if C replicates to E and F and D eliminates on the way from A to B, a packet from A to B experiences 2 paths, A->C->E->D->B and A->C->F->D->B.

2.1.3.2. Track

A networking graph that can be followed to transport packets with equivalent treatment; as opposed to the definition of a path above, a Track represents not an experience but a potential, is not necessarily a linear sequence, and is not necessarily fully traversed (flooded) by all packets of a flow. It may contain multiple paths that may overlap, fork and rejoin, for instance to enable the RAW PAREO operations.




Figure 1: Example IoT Track to an IoT gateway with 1+1 redundancy

In DetNet [\[RFC8655\]](#) terms, a Track has the following properties:

- *A Track is a layer-3 abstraction built upon P2P IP links between routers. A router may form multiple P2P IP links over a single radio interface.
- *A Track has one Ingress and one Egress nodes, which operate as DetNet Edge nodes.
- *A Track is reversible, meaning that packets can be routed against the flow of data packets, e.g., to carry OAM measurements or control messages back to the Ingress.
- *The vertices of the Track are DetNet Relay nodes that operate at the DetNet Service sub-layer and provide the PAREO functions.
- *The topological edges of the graph are serial sequences of DetNet Transit nodes that operate at the DetNet Forwarding sub-layer.

2.1.3.3. SubTrack

A Track within a Track. The RAW PSE selects a subTrack on a per-packet or a per-collection of packets basis to provide the desired reliability for the transported flows.

2.1.3.4. Segment

A serial path formed by a topological edge of a Track. East-West Segments are oriented from Ingress (East) to Egress (West). North/South Segments can be bidirectional; to avoid loops, measures must be taken to ensure that a given packet flows either Northwards or Southwards along a bidirectional Segment, but never bounces back.

2.1.4. Deterministic Networking

This document reuses the terminology in section 2 of [[RFC8557](#)] and section 4.1.2 of [[RFC8655](#)] for deterministic networking and deterministic networks.

2.1.4.1. Flow

A collection of consecutive IP packets defined by the upper layers and signaled by the same 5 or 6-tuple, see section 5.1 of [[RFC8939](#)]. Packets of the same flow must be placed on the same Track to receive an equivalent treatment from Ingress to Egress within the Track. Multiple flows may be transported along the same Track. The subTrack that is selected for the flow may change over time under the control of the PSE.

2.1.4.2. Deterministic Flow Identifier (L2)

A tuple identified by a stream_handle, and provided by a bridge, in accordance with IEEE 802.1CB. The tuple comprises at least source MAC, destination MAC, VLAN ID, and L2 priority. Continuous streams are characterized by bandwidth and max packet size; scheduled streams are characterized by a repeating pattern of timed transmissions.

2.1.4.3. Deterministic Flow Identifier (L3)

See section 3.3 of [[DetNet-DP](#)]. The classical IP 5-tuple that identifies a flow comprises the source IP, destination IP, source port, destination port, and the upper layer protocol (ULP). DetNet uses a 6-tuple where the extra field is the DSCP field in the packet. The IPv6 flow label is not used for that purpose.

2.1.4.4. TSN

TSN stands for Time Sensitive Networking and denotes the efforts at IEEE 802 for deterministic networking, originally for use on Ethernet. Wireless TSN (WTSN) denotes extensions of the TSN work on wireless media such as the selected RAW technologies [[RAW-TECHNOS](#)].

2.1.5. Reliability and Availability

In the context of the RAW work, Reliability and Availability are defined as follows:

2.1.5.1. Service Level Agreement

In the context of RAW, an SLA (service level agreement) is a contract between a provider, the network, and a client, the application flow, about measurable metrics such as latency boundaries, consecutive losses, and packet delivery ratio (PDR).

2.1.5.2. Service Level Objective

A service level objective (SLO) is one term in the SLA, for which specific network setting and operations are implemented. For instance, a dynamic tuning of the packet redundancy will address an SLO of consecutive losses in a row by augmenting the chances of delivery of a packet that follows a loss.

2.1.5.3. Service Level Indicator

A service level indicator (SLI) measures the compliance of an SLO to the terms of the contract. It can be for instance the statistics of individual losses and losses in a row as time series.).

2.1.5.4. Reliability

Reliability is a measure of the probability that an item will perform its intended function for a specified interval under stated conditions (SLA). RAW expresses reliability in terms of Mean Time Between Failure (MTBF) and Maximum Consecutive Failures (MCF). More in [[NASA](#)].).

2.1.5.5. Available

That is exempt of unscheduled outage or derivation from the terms of the SLA. A basic expectation for a RAW network is that the flow is maintained in the face of any single breakage or flapping.

2.1.5.6. Availability

Availability is a measure of the relative amount of time where a RAW Network operates in stated condition (SLA), expressed as $(\text{uptime}) / (\text{uptime} + \text{downtime})$. Because a serial wireless path may not be good enough to provide the required reliability, and even 2 parallel paths may not be over a longer period of time, the RAW availability implies a journey that is a lot more complex than following a serial path.

2.1.6. OAM variations

2.1.6.1. Active OAM

See [[RFC7799](#)]. In the context of RAW, Active OAM is used to observe a particular Track, subTrack, or Segment of a Track regardless of whether it is used for traffic at that time.

2.1.6.2. In-Band OAM

An active OAM packet is considered in-band for the monitored Track when it traverses the same set of links and interfaces and if the OAM packet receives the same QoS and PAREO treatment as the packets of the data flows that are injected in the Track.

2.1.6.3. Out-of-Band OAM

Out-of-band OAM is an active OAM whose path is not topologically congruent to the Track, or its test packets receive a QoS and/or PAREO treatment that is different from that of the packets of the data flows that are injected in the Track, or both.

2.1.6.4. Limited OAM

An active OAM packet is a Limited OAM packet when it observes the RAW operation over a node, a segment, or a subTrack of the Track, though not from Ingress to Egress. It is injected in the datapath and extracted from the datapath around the particular function or subnetwork (e.g., around a relay providing a service layer replication point) that is being tested.

2.1.6.5. Upstream OAM

An upstream OAM packet is an Out-of-Band OAM packet that traverses the Track from egress to ingress on the reverse direction, to capture and report OAM measurements upstream. The collection may capture all information along the whole Track, or it may only learn select data across all, or only a particular subTrack, or Segment of a Track.

2.1.6.6. Residence Time

A residence time (RT) is defined as the time period between the reception of a packet starts and the transmission of the packet begins. In the context of RAW, RT is useful for a transit node, not ingress or egress.

2.1.6.7. Additional References

[[DetNet-OAM](#)] provides additional terminology related to OAM in the context of DetNet and by extension of RAW, whereas [[RFC7799](#)] defines the Active, Passive, and Hybrid OAM methods.

2.2. Reliability and Availability

2.2.1. High Availability Engineering Principles

The reliability criteria of a critical system pervades through its elements, and if the system comprises a data network then the data network is also subject to the inherited reliability and availability criteria. It is only natural to consider the art of high availability engineering and apply it to wireless communications in the context of RAW.

There are three principles [pillars] of high availability engineering:

1. elimination of single points of failure
2. reliable crossover
3. prompt detection of failures as they occur.

These principles are common to all high availability systems, not just ones with Internet technology at the center. Examples of both non-Internet and Internet are included.

2.2.1.1. Elimination of Single Points of Failure

Physical and logical components in a system happen to fail, either as the effect of wear and tear, when used beyond acceptable limits, or due to a software bug. It is necessary to decouple component failure from system failure to avoid the latter. This allows failed components to be restored while the rest of the system continues to function.

IP Routers leverage routing protocols to compute alternate routes in case of a failure. There is a rather open-ended issue over alternate routes -- for example, when links are cabled through the same conduit, they form a shared risk link group (SRLG), and will share the same fate if the bundle is cut. The same effect can happen with virtual links that end up in a same physical transport through the games of encapsulation. In a same fashion, an interferer or an obstacle may affect multiple wireless transmissions at the same time, even between different sets of peers.

Intermediate network Nodes such as routers, switches and APs, wire bundles and the air medium itself can become single points of failure. For High Availability, it is thus required to use

physically link- and Node-disjoint paths; in the wireless space, it is also required to use the highest possible degree of diversity (time, space, code, frequency, channel width) in the transmissions over the air to combat the additional causes of transmission loss.

From an economics standpoint, executing this principle properly generally increases capitalization expense because of the redundant equipment. In a constrained network where the waste of energy and bandwidth should be minimized, an excessive use of redundant links must be avoided; for RAW this means that the extra bandwidth must be used wisely and with parsimony.

2.2.1.2. Reliable Crossover

Having a backup equipment has a limited value unless it can be reliably switched into use within the down-time parameters. IP Routers execute reliable crossover continuously because the routers will use any alternate routes that are available [[RFC0791](#)]. This is due to the stateless nature of IP datagrams and the dissociation of the datagrams from the forwarding routes they take. The "[IP Fast Reroute Framework](#)" [[FRR](#)] analyzes mechanisms for fast failure detection and path repair for IP Fast-Reroute, and discusses the case of multiple failures and SRLG. Examples of FRR techniques include Remote Loop-Free Alternate [[RLFA-FRR](#)] and backup label-switched path (LSP) tunnels for the local repair of LSP tunnels using RSVP-TE [[RFC4090](#)].

Deterministic flows, on the contrary, are attached to specific paths where dedicated resources are reserved for each flow. Therefore each DetNet path must inherently provide sufficient redundancy to provide the guaranteed SLA at all times. The DetNet PREOF typically leverages 1+1 redundancy whereby a packet is sent twice, over non-congruent paths. This avoids the gap during the fast reroute operation, but doubles the traffic in the network.

In the case of RAW, the expectation is that multiple transient faults may happen in overlapping time windows, in which case the 1+1 redundancy with delayed reestablishment of the second path will not provide the required guarantees. The Data Plane must be configured with a sufficient degree of redundancy to select an alternate redundant path immediately upon a fault, without the need for a slow intervention from the controller plane.

2.2.1.3. Prompt Notification of Failures

The execution of the two above principles is likely to render a system where the user will rarely see a failure. But someone needs to in order to direct maintenance.

There are many reasons for system monitoring (FCAPS for fault, configuration, accounting, performance, security is a handy mental checklist) but fault monitoring is sufficient reason.

["An Architecture for Describing Simple Network Management Protocol \(SNMP\) Management Frameworks"](#) [STD 62] describes how to use SNMP to observe and correct long-term faults.

["Overview and Principles of Internet Traffic Engineering"](#) [TE] discusses the importance of measurement for network protection, and provides abstract an method for network survivability with the analysis of a traffic matrix as observed by SNMP, probing techniques, FTP, IGP link state advertisements, and more.

Those measurements are needed in the context of RAW to inform the controller and make the long term reactive decision to rebuild a complex path based on statistical and aggregated information. RAW itself operates in the Network Plane at a faster time scale with live information on speed, state, etc... This live information can be obtained directly from the lower layer, e.g., using L2 triggers, read from a protocol such as the [Dynamic Link Exchange Protocol \(DLEP\)](#) [DLEP], or transported over multiple hops using OAM and reverse OAM, as illustrated in [Figure 9](#).

2.2.2. Applying Reliability Concepts to Networking

The terms Reliability and Availability are defined for use in RAW in [Section 2.1](#) and the reader is invited to read [[NASA](#)] for more details on the general definition of Reliability. Practically speaking a number of nines is often used to indicate the reliability of a data link, e.g., 5 nines indicate a Packet Delivery Ratio (PDR) of 99.999%.

This number is typical in a wired environment where the loss is due to a random event such as a solar particle that affects the transmission of a particular frame, but does not affect the previous or next frame, nor frames transmitted on other links. Note that the QoS requirements in RAW may include a bounded latency, and a packet that arrives too late is a fault and not considered as delivered.

For a periodic networking pattern such as an automation control loop, this number is proportional to the Mean Time Between Failures (MTBF). When a single fault can have dramatic consequences, the MTBF expresses the chances that the unwanted fault event occurs. In data networks, this is rarely the case. Packet loss cannot never be fully avoided and the systems are built to resist to one loss, e.g., using redundancy with Retries (HARQ) or Packet Replication and Elimination (PRE), or, in a typical control loop, by linear interpolation from the previous measurements.

But the linear interpolation method cannot resist multiple consecutive losses, and a high MTBF is desired as a guarantee that this will not happen, IOW that the number of losses-in-a-row can be bounded. In that case, what is really desired is a Maximum Consecutive Failures (MCF). If the number of losses in a row passes the MCF, the control loop has to abort and the system, e.g., the production line, may need to enter an emergency stop condition.

Engineers that build automated processes may use the network reliability expressed in nines or as an MTBF as a proxy to indicate an MCF, e.g., as described in section 7.4 of the ["Deterministic Networking Use Cases"](#) [RFC8578].

2.2.3. Wireless Effects Affecting Reliability

In contrast with wired networks, errors in transmission are the predominant source of packet loss in wireless networks.

The root cause for the loss may be of multiple origins, calling for the use of different forms of diversity:

Multipath Fading A destructive interference by a reflection of the original signal.

A radio signal may be received directly (line-of-sight) and/or as a reflection on a physical structure (echo). The reflections take a longer path and are delayed by the extra distance divided by the speed of light in the medium. Depending on the frequency, the echo lands with a different phase which may add up to (constructive interference) or cancel the direct signal (destructive interference).

The affected frequencies depend on the relative position of the sender, the receiver, and all the reflecting objects in the environment. A given hop will suffer from multipath fading for multiple packets in a row till a physical movement changes the reflection patterns.

Co-channel Interference Energy in the spectrum used for the transmission confuses the receiver.

The wireless medium itself is a Shared Risk Link Group (SRLG) for nearby users of the same spectrum, as an interference may affect multiple co-channel transmissions between different peers within the interference domain of the interferer, possibly even when they use different technologies.

Obstacle in Fresnel Zone The optimal transmission happens when the Fresnel Zone between the sender and the receiver is free of obstacles.

As long as a physical object (e.g., a metallic trolley between peers) that affects the transmission is not removed, the quality of the link is affected.

In an environment that is rich of metallic structures and mobile objects, a single radio link will provide a fuzzy service, meaning that it cannot be trusted to transport the traffic reliably over a long period of time.

Transmission losses are typically not independent, and their nature and duration are unpredictable; as long as a physical object (e.g., a metallic trolley between peers) that affects the transmission is not removed, or as long as the interferer (e.g., a radar) keeps transmitting, a continuous stream of packets will be affected.


The key technique to combat those unpredictable losses is diversity. Different forms of diversity are necessary to combat different causes of loss and the use of diversity must be maximized to optimize the PDR.

A single packet may be sent at different times (time diversity) over diverse paths (spatial diversity) that rely on diverse radio channels (frequency diversity) and diverse PHY technologies, e.g., narrowband vs. spread spectrum, or diverse codes. Using time diversity will defeat short-term interferences; spatial diversity combats very local causes such as multipath fading; narrowband and spread spectrum are relatively innocuous to one another and can be used for diversity in the presence of the other.

2.3. Routing Time Scale vs. Forwarding Time Scale

With DetNet, the Controller Plane Function (CPF) handles the routing computation and maintenance. With RAW, the CPF also performs the PSE orientation, proposing SubTracks to use in response to network events. The CPF can be centralized in a PCE, and can reside outside the network. This is how the remainder of this document depicts it, though the CPF could be implemented otherwise without affecting the architecture. In a wireless mesh, the path to the PCE can be expensive and slow, possibly going across the whole mesh and back. Reaching to the PCE can also be slow in regards to the speed of events that affect the forwarding operation at the radio layer.

Due to that cost and latency, the Controller Plane is not expected to be sensitive/reactive to transient changes. The abstraction of a link at the routing level is expected to use statistical metrics that aggregate the behavior of a link over long periods of time, and represent its properties as shades of gray as opposed to numerical values such as a link quality indicator, or a boolean value for either up or down.



*** = flapping at this time

Figure 2: Time Scales

In the case of wireless, the changes that affect the forwarding decision can happen frequently and often for short durations, e.g., a mobile object moves between a transmitter and a receiver, and will cancel the line of sight transmission for a few seconds, or a radar measures the depth of a pool and interferes on a particular channel for a split second.

There is thus a desire to separate the long term computation of the route and the short term forwarding decision. In that model, the routing operation computes a complex Track that enables multiple Non-Equal Cost Multi-Path (N-ECMP) forwarding solutions, and leaves it to the Data Plane to make the per-packet decision of which of these possibilities should be used.

In the wired world, and more specifically in the context of Traffic Engineering (TE), an alternate path can be used upon the detection of a failure in the main path, e.g., using OAM in MPLS-TP or BFD

over a collection of SD-WAN tunnels. RAW formalizes a forwarding time scale that is an order(s) of magnitude shorter than the controller plane routing time scale, and separates the protocols and metrics that are used at both scales. Routing can operate on long term statistics such as delivery ratio over minutes to hours, but as a first approximation can ignore flapping. On the other hand, the RAW forwarding decision is made at the scale of the packet rate, and uses information that must be pertinent at the present time for the current transmission(s).

3. The RAW Conceptual Model

RAW inherits the conceptual model described in section 4 of the DetNet Architecture [[RFC8655](#)]. RAW extends the DetNet service layer to provide additional agility against transmission loss.

3.1. The RAW Planes

A RAW Network Plane may be strict or loose, depending on whether RAW observes and takes actions on all hops or not. For instance, the packets between two wireless entities may be relayed over a wired infrastructure such as a Wi-Fi extended service set (ESS) or a 5G Core; in that case, RAW observes and controls the transmission over the wireless first and last hops, as well as end-to-end metrics such as latency, jitter, and delivery ratio. This operation is loose since the structure and properties of the wired infrastructure are ignored, and may be either controlled by other means such as DetNet/TSN, or neglected in the face of the wireless hops.

A Controller Plane Function (CPF) called the Path Computation Element (PCE) [[RFC4655](#)] interacts with RAW Nodes over a Southbound API. The RAW Nodes are DetNet relays that are capable of additional diversity mechanisms and measurement functions related to the radio interface, in particular the PAREO diversity mechanisms. RAW leverages a CPF that operates inside the RAW Nodes (typically the Ingress Edge Nodes) to dynamically adapt the path of the packets and optimizes the resource usage.

The PCE defines a complex Track between an Ingress End System and an Egress End System, and indicates to the RAW Nodes where the PAREO operations may be actioned in the Network Plane. The Track may be expressed loosely to enable traversing a non-RAW subnetwork. In that case, the expectation is that the non-RAW subnetwork can be neglected in the RAW computation, that is, considered infinitely fast, reliable and/or available in comparison with the links between RAW nodes.

assure that the SLA can be met or have it recompute the Track if not.

3.2. RAW vs. DetNet

RAW leverages the DetNet Forwarding sub-layer and requires the support of in-situ OAM in DetNet Transit Nodes (see fig 3 of [\[RFC8655\]](#) for the dynamic acquisition of link capacity and state to maintain a strict RAW service, end-to-end, over a DetNet Network.

RAW extends the DetNet Stack (see fig 4 of [\[RFC8655\]](#)) with additional functionality at the DetNet Service sub-layer for the PSE operation. PREOF is extended with the PAREO capabilities (see [Section 4.4](#)) and the RAW PAREO Actuator manages dynamically the PAREO operations. The RAW Service sub-layer also adds the OAM Propagator that (re)generates the OAM information as it is formed and propagated In-Band or Out-of-Band. The RAW Service sub-layer may be present in DetNet Edge and Relay Nodes, though the PAREO Actuator has no operation in the Egress Edge Node.

RAW also adds a Control sub-layer that operates in the DetNet Controller Plane. The RAW Control sub-layer typically runs only in the DetNet Ingress Edge Node or End System, though it may also run in DetNet Relay Nodes when the RAW Control sub-layer is distributed along the Track. The RAW Control sub-layer functionality includes the PSE that decides the subTrack for the next packets of a flows and controls the PAREO Actuators along the subTrack through specific signaling, and the OAM Supervisor that triggers, and learns from, OAM observations, and feeds the PSE for its next decision.

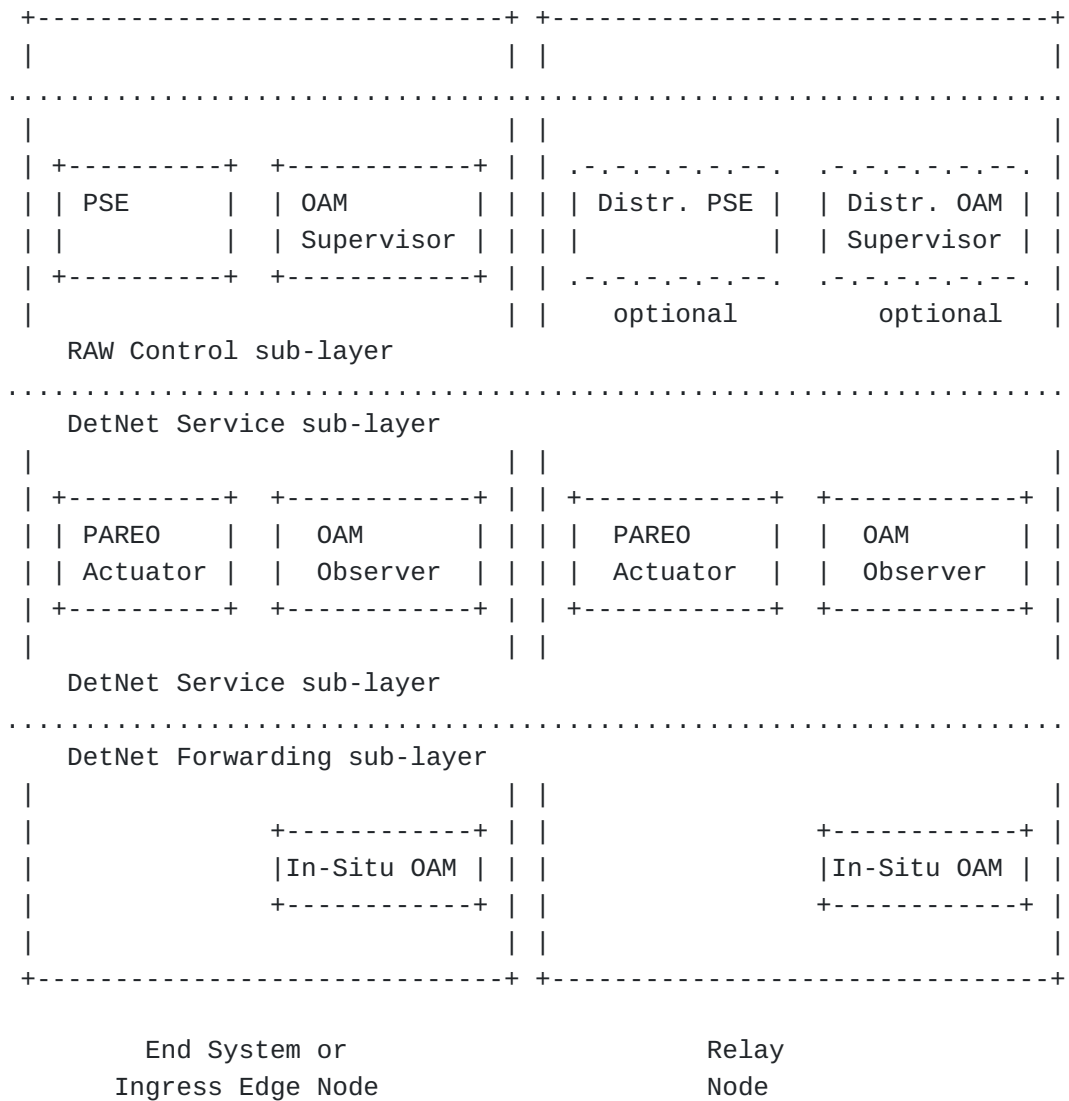


Figure 4: RAW DetNet Services

There are 2 main proposed models to deploy RAW and DetNet. In the first model (strict) illustrated in [Figure 5](#), RAW operates over a continuous DetNet Service end-to-end between the Ingress and the Egress Edge Nodes or End Systems.

A minimal Forwarding sub-layer service is provided at all DetNet Nodes to ensure that the OAM information flows. Relay Nodes may or may not support RAW services, and the Edge nodes do support RAW. DetNet guarantees such as latency are provided end-to-end, and RAW supports the DetNet Service to optimize the use of resources.

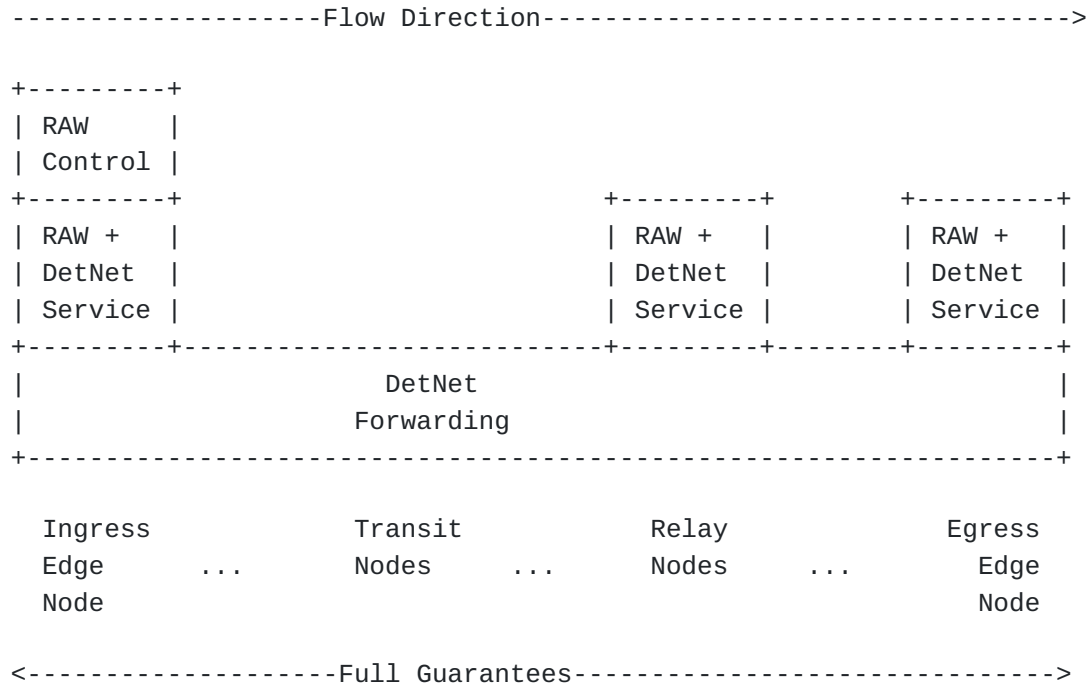


Figure 5: (Strict) RAW over DetNet

In the second model (loose), illustrated in [Figure 6](#), RAW operates over a partial DetNet Service where typically only the Ingress and the Egress End Systems support RAW. The DetNet Domain may extend beyond the Ingress node, or there may be a DetNet domain starting at an Ingress Edge Node at the first hop after the End System.

In the loose model, RAW cannot observe the hops in network, and the path beyond the first hop is opaque; RAW can still observe the end-to-end behavior and use Layer-3 measurements to decide whether to replicate a packet and select the first hop interface(s).




Figure 6: Loose RAW

4. The OODA Loop

The RAW Architecture is structured as an OODA Loop (Observe, Orient, Decide, Act). It involves:

1. Network Plane measurement protocols for Operations, Administration and Maintenance (OAM) to Observe some or all hops along a Track as well as the end-to-end packet delivery, more in [Section 4.1](#);
2. Controller plane elements to report the links statistics to a Path computation Element (PCE) in a centralized controller that computes and installs the Tracks and provides meta data to Orient the routing decision, more in [Section 4.2](#);
3. A Runtime distributed Path Selection Engine (PSE) thar Decides which subTrack to use for the next packet(s) that are routed along the Track, more in [Section 4.3](#);
4. Packet (hybrid) ARQ, Replication, Elimination and Ordering Dataplane actions that operate at the DetNet Service Layer to increase the reliability of the end-to-end transmission. The RAW architecture also covers in-situ signaling when the decision is Acted by a node that down the Track from the PSE, more in [Section 4.4](#).

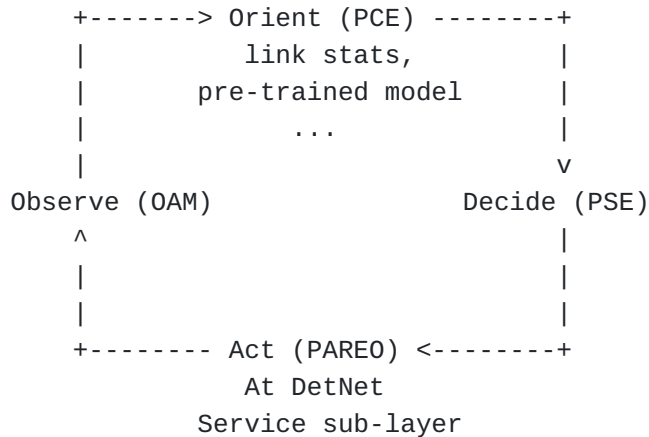


Figure 7: The RAW OODA Loop

The overall OODA Loop optimizes the use of redundancy to achieve the required reliability and availability Service Level Agreement (SLA) while minimizing the use of constrained resources such as spectrum and battery.

4.1. Observe: The RAW OAM

RAW In-situ OAM operation in the Network Plane may observe either a full Track or subTracks that are being used at this time. As packets may be load balanced, replicated, eliminated, and / or fragmented for Network Coding (NC) forward error correction (FEC), the RAW In-situ operation needs to be able to signal which operation occurred to an individual packet.

Active RAW OAM may be needed to observe the unused segments and evaluate the desirability of a rerouting decision.

Finally, the RAW Service Layer Assurance may observe the individual PAREO operation of a relay node to ensure that it is conforming; this might require injecting an OAM packet at an upstream point inside the Track and extracting that packet at another point downstream before it reaches the egress.

This observation feeds the RAW PSE that makes the decision on which PAREO function is actioned at which RAW Node, for one a small continuous series of packets.




Figure 8: Observed Links in Radio Access Protection

In the case of a End-to-End Protection in a Wireless Mesh, the Track is strict and congruent with the path so all links are observed.

Conversely, in the case of Radio Access Protection illustrated in [Figure 8](#), the Track is Loose and only the first hop is observed; the rest of the path is abstracted and considered infinitely reliable. The loss if a packet is attributed to the first hop Radio Access Network (RAN), even if a particular loss effectively happens farther down the path. In that case, RAW enables technology diversity (e.g. Wi-Fi and 5G) which in turn improves the diversity in spectrum usage.

The Links that are not observed by OAM are opaque to it, meaning that the OAM information is carried across and possibly echoed as data, but there is no information capture in intermediate nodes. In the example above, the Internet is opaque and not controlled by RAW; still the RAW OAM measures the end-to-end latency and delivery ratio for packets sent via each if RAN 1, RAN 2 and RAN 3, and determines whether a packet should be sent over either or a collection of those access links.

4.2. Orient: The Path Computation Engine

RAW separates the long time scale at which a Track is elaborated and installed, from the short time scale at which the forwarding decision is taken for one or a few packets (see in [Section 2.3](#)) that will experience the same path until the network conditions evolve and another path is selected within the same Track.

The Track computation is out of scope, but RAW expects that the Controller plane protocol that installs the Track also provides related knowledge in the form of meta data about the links, segments

and possible subTracks. That meta data can be a pre-digested statistical model, and may include prediction of future flaps and packet loss, as well as recommended actions when that happens.

The meta data may include:

- *Pre-Determined subTracks to match predictable error profiles
- *Pre-Trained models
- *Link Quality Statistics and their projected evolution

The Track is installed with measurable objectives that are computed by the PCE to achieve the RAW SLA. The objectives can be expressed as any of maximum number of packet lost in a row, bounded latency, maximal jitter, maximum number of interleaved out of order packets, average number of copies received at the elimination point, and maximal delay between the first and the last received copy of the same packet.

4.3. Decide: The Path Selection Engine

The RAW OODA Loop operates at the path selection time scale to provide agility vs. the brute force approach of flooding the whole Track. The OODA Loop controls, within the redundant solutions that are proposed by the PCE, which will be used for each packet to provide a Reliable and Available service while minimizing the waste of constrained resources.

To that effect, RAW defines the Path Selection Engine (PSE) that is the counterpart of the PCE to perform rapid local adjustments of the forwarding tables within the diversity that the PCE has selected for the Track. The PSE enables to exploit the richer forwarding capabilities with PAREO and scheduled transmissions at a faster time scale over the smaller domain that is the Track, in either a loose or a strict fashion.

Compared to the PCE, the PSE operates on metrics that evolve faster, but that need to be advertised at a fast rate but only locally, within the Track. The forwarding decision may also change rapidly, but with a scope that is also contained within the Track, with no visibility to the other Tracks and flows in the network. This is as opposed to the PCE that must observe the whole network and optimize all the Tracks globally, which can only be done at a slow pace and using long-term statistical metrics, as presented in [Table 1](#).

	PCE (Not in Scope)	PSE (In Scope)
Operation	Centralized	Source-Routed or Distributed

	PCE (Not in Scope)	PSE (In Scope)
Communication	Slow, expensive	Fast, local
Time Scale	hours and above	seconds and below
Network Size	Large, many Tracks to optimize globally	Small, within one Track
Considered Metrics	Averaged, Statistical, Shade of grey	Instant values / boolean condition

Table 1: PCE vs. PSE

The PSE sits in the DetNet Service sub-Layer of Edge and Relay Nodes. On the one hand, it operates on the packet flow, learning the Track and path selection information from the packet, possibly making local decision and retagging the packet to indicate so. On the other hand, the PSE interacts with the lower layers and with its peers to obtain up-to-date information about its radio links and the quality of the overall Track, respectively, as illustrated in [Figure 9](#).

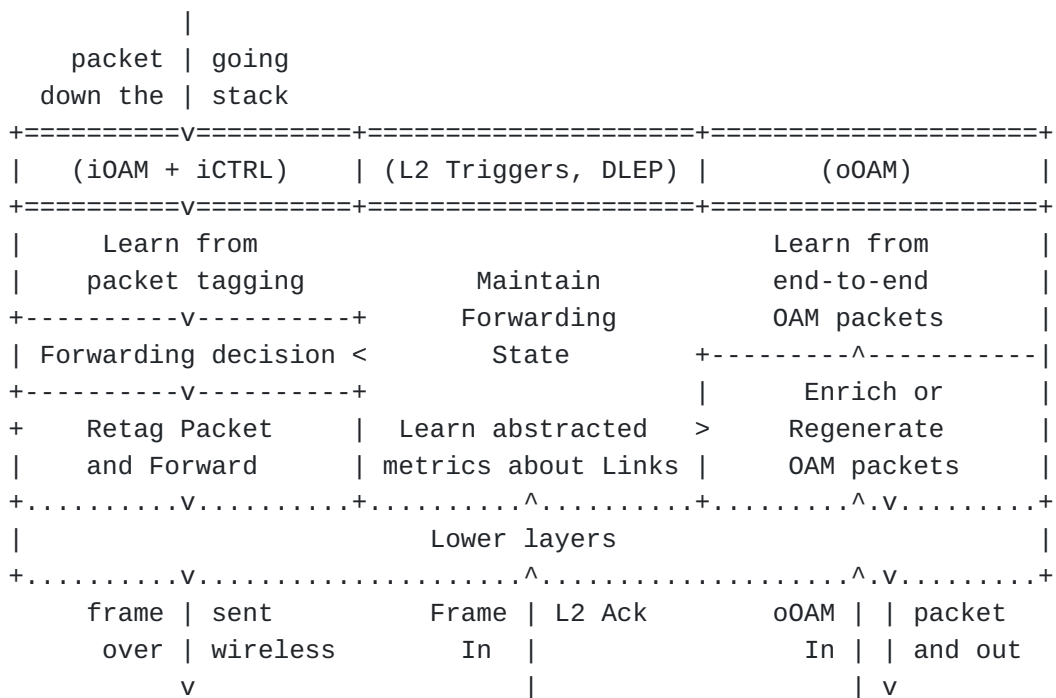


Figure 9: PSE

4.4. Act: The PAREO Functions

RAW may control whether and how to use packet replication and elimination (PRE), fragmentation, and network coding, and how the lower layers performs Automatic Repeat reQuest (ARQ), Hybrid ARQ (HARQ) that includes Forward Error Correction (FEC), and other wireless-specific techniques such as overhearing and constructive

interferences, in order to increase the reliability and availability of the end-to-end transmission.

Collectively, those functions are called PAREO for Packet (hybrid) ARQ, Replication, Elimination and Ordering. By tuning dynamically the use of PAREO functions, RAW avoids the waste of critical resources such as spectrum and energy while providing that the guaranteed SLA, e.g., by adding redundancy only when a spike of loss is observed.

In a nutshell, PAREO establishes several paths in a network to provide redundancy and parallel transmissions to bound the end-to-end delay to traverse the network. Optionally, promiscuous listening between paths is possible, such that the Nodes on one path may overhear transmissions along the other path. Considering the scenario shown in [Figure 10](#), many different paths are possible to traverse the network from ingress to egress. A simple way to benefit from this topology could be to use the two independent paths via Nodes A, C, E and via B, D, F. But more complex paths are possible by interleaving transmissions from the lower level of the path to the upper level.




Figure 10: A Ladder Shape with Two Parallel Paths

PAREO may also take advantage of the shared properties of the wireless medium to compensate for the potential loss that is incurred with radio transmissions.

For instance, when the source sends to Node A, Node B may listen promiscuously and get a second chance to receive the frame without an additional transmission. Note that B would not have to listen if it already received that particular frame at an earlier timeslot in a dedicated transmission towards B.

The PAREO model can be implemented in both centralized and distributed scheduling approaches. In the centralized approach, a Path Computation Element (PCE) scheduler calculates a Track and schedules the communication. In the distributed approach, the Track is computed within the network, and signaled in the packets, e.g., using BIER-TE, Segment Routing, or a Source Routing Header.

4.4.1. Packet Replication

By employing a Packet Replication procedure, a Node forwards a copy of each data packet to more than one successor. To do so, each Node (i.e., Ingress and intermediate Node) sends the data packet multiple times as separate unicast transmissions. For instance, in [Figure 11](#), the Ingress Node is transmitting the packet to both successors, nodes A and B, at two different times.




Figure 11: Packet Replication

An example schedule is shown in [Table 2](#). This way, the transmission leverages with the time and spatial forms of diversity.

Channel	0	1	2	3	4	5	6
0	S->A	S->B	B->C	B->D	C->F	E->R	F->R
1		A->C	A->D	C->E	D->E	D->F	

Table 2: Packet Replication: Sample schedule

4.4.2. Packet Elimination

The replication operation increases the traffic load in the network, due to packet duplications. This may occur at several stages inside the Track, and to avoid an explosion of the number of copies, a Packet Elimination procedure must be applied as well. To this aim, once a Node receives the first copy of a data packet, it discards the subsequent copies.

The logical functions of Replication and Elimination may be collocated in an intermediate Node, the Node first eliminating the redundant copies and then sending the packet exactly once to each of the selected successors.

4.4.3. Promiscuous Overhearing

Considering that the wireless medium is broadcast by nature, any neighbor of a transmitter may overhear a transmission. By employing the Promiscuous Overhearing operation, the next hops have additional opportunities to capture the data packets. In [Figure 12](#), when Node A is transmitting to its DP (Node C), the AP (Node D) and its sibling

(Node B) may decode this data packet as well. As a result, by employing correlated paths, a Node may have multiple opportunities to receive a given data packet.




Figure 12: Unicast with Overhearing

Variations on the same idea such as link-layer anycast and multicast may also be used to reach more than one next-hop with a single frame.

4.4.4. Constructive Interference

Constructive Interference can be seen as the reverse of Promiscuous Overhearing, and refers to the case where two senders transmit the exact same signal in a fashion that the emitted symbols add up at the receiver and permit a reception that would not be possible with a single sender at the same PHY mode and the same power level.

Constructive Interference was proposed on 5G, Wi-Fi7 and even tested on IEEE Std 802.14.5. The hard piece is to synchronize the senders to the point that the signals are emitted at slightly different time to offset the difference of propagation delay that corresponds to the difference of distance of the transmitters to the receiver at the speed of light to the point that the symbols are superposed long enough to be recognizable.

5. Security Considerations

RAW uses all forms of diversity including radio technology and physical path to increase the reliability and availability in the face of unpredictable conditions. While this is not done specifically to defeat an attacker, the amount of diversity used in RAW makes an attack harder to achieve.

5.1. Layer-2 encryption

Radio networks typically encrypt at the MAC layer to protect the transmission. If the encryption is per pair of peers, then certain RAW operations like promiscuous overhearing become impossible.

5.2. Forced Access

RAW will typically select the cheapest collection of links that matches the requested SLA, for instance, leverage free WI-Fi vs. paid 3GPP access. By defeating the cheap connectivity (e.g., PHY-layer interference) the attacker can force an End System to use the paid access and increase the cost of the transmission for the user.

6. IANA Considerations

This document has no IANA actions.

7. Contributors

The editor wishes to thank:

Xavi Vilajosana: Wireless Networks Research Lab, Universitat Oberta de Catalunya

Remous-Aris Koutsiamanis: IMT Atlantique

Nicolas Montavont: IMT Atlantique

Rex Buddenberg: Individual contributor

Greg Mirsky: ZTE

for their contributions to the text and ideas exposed in this document.

8. Acknowledgments

The authors wish to thank Dave Cavalcanti and Fabrice Theoleyre for their in-depth reviews during the development of this document.

9. References

9.1. Normative References

[6TiSCH-ARCHI] Thubert, P., Ed., "An Architecture for IPv6 over the Time-Slotted Channel Hopping Mode of IEEE 802.15.4 (6TiSCH)", RFC 9030, DOI 10.17487/RFC9030, May 2021, <<https://www.rfc-editor.org/info/rfc9030>>.

[INT-ARCHI] Braden, R., Ed., "Requirements for Internet Hosts - Communication Layers", STD 3, RFC 1122, DOI 10.17487/

RFC1122, October 1989, <<https://www.rfc-editor.org/info/rfc1122>>.

[RAW-TECHNOS] Thubert, P., Cavalcanti, D., Vilajosana, X., Schmitt, C., and J. Farkas, "Reliable and Available Wireless Technologies", Work in Progress, Internet-Draft, draft-ietf-raw-technologies-05, 2 February 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-raw-technologies-05>>.

[RAW-USE-CASES] Bernardos, C. J., Papadopoulos, G. Z., Thubert, P., and F. Theoleyre, "RAW use-cases", Work in Progress, Internet-Draft, draft-ietf-raw-use-cases-05, 23 February 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-raw-use-cases-05>>.

[RFC4655] Farrel, A., Vasseur, J.-P., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, DOI 10.17487/RFC4655, August 2006, <<https://www.rfc-editor.org/info/rfc4655>>.

[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>>.

[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>>.

[RFC8578] Grossman, E., Ed., "Deterministic Networking Use Cases", RFC 8578, DOI 10.17487/RFC8578, May 2019, <<https://www.rfc-editor.org/info/rfc8578>>.

[IPv6] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <<https://www.rfc-editor.org/info/rfc8200>>.

[RFC8557] Finn, N. and P. Thubert, "Deterministic Networking Problem Statement", RFC 8557, DOI 10.17487/RFC8557, May 2019, <<https://www.rfc-editor.org/info/rfc8557>>.

[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>>.

[RFC9049]

Dawkins, S., Ed., "Path Aware Networking: Obstacles to Deployment (A Bestiary of Roads Not Taken)", RFC 9049, DOI 10.17487/RFC9049, June 2021, <<https://www.rfc-editor.org/info/rfc9049>>.

9.2. Informative References

[RFC0791]

Postel, J., "Internet Protocol", STD 5, RFC 791, DOI 10.17487/RFC0791, September 1981, <<https://www.rfc-editor.org/info/rfc791>>.

[TE]

Awduche, D., Chiu, A., Elwalid, A., Widjaja, I., and X. Xiao, "Overview and Principles of Internet Traffic Engineering", RFC 3272, DOI 10.17487/RFC3272, May 2002, <<https://www.rfc-editor.org/info/rfc3272>>.

[STD 62]

Harrington, D., Presuhn, R., and B. Wijnen, "An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks", STD 62, RFC 3411, DOI 10.17487/RFC3411, December 2002, <<https://www.rfc-editor.org/info/rfc3411>>.

[RFC4090]

Pan, P., Ed., Swallow, G., Ed., and A. Atlas, Ed., "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", RFC 4090, DOI 10.17487/RFC4090, May 2005, <<https://www.rfc-editor.org/info/rfc4090>>.

[FRR]

Shand, M. and S. Bryant, "IP Fast Reroute Framework", RFC 5714, DOI 10.17487/RFC5714, January 2010, <<https://www.rfc-editor.org/info/rfc5714>>.

[RLFA-FRR]

Bryant, S., Filsfils, C., Previdi, S., Shand, M., and N. So, "Remote Loop-Free Alternate (LFA) Fast Reroute (FRR)", RFC 7490, DOI 10.17487/RFC7490, April 2015, <<https://www.rfc-editor.org/info/rfc7490>>.

[DetNet-DP]

Varga, B., Ed., Farkas, J., Berger, L., Malis, A., and S. Bryant, "Deterministic Networking (DetNet) Data Plane Framework", RFC 8938, DOI 10.17487/RFC8938, November 2020, <<https://www.rfc-editor.org/info/rfc8938>>.

[DLEP]

Ratliff, S., Jury, S., Satterwhite, D., Taylor, R., and B. Berry, "Dynamic Link Exchange Protocol (DLEP)", RFC

8175, DOI 10.17487/RFC8175, June 2017, <<https://www.rfc-editor.org/info/rfc8175>>.

[I-D.irtf-panrg-path-properties] Enghardt, T. and C. Krähenbühl, "A Vocabulary of Path Properties", Work in Progress, Internet-Draft, draft-irtf-panrg-path-properties-05, 7 March 2022, <<https://datatracker.ietf.org/doc/html/draft-irtf-panrg-path-properties-05>>.

[IPoWIRELESS] Thubert, P., "IPv6 Neighbor Discovery on Wireless Networks", Work in Progress, Internet-Draft, draft-thubert-6man-ipv6-over-wireless-11, 15 December 2021, <<https://datatracker.ietf.org/doc/html/draft-thubert-6man-ipv6-over-wireless-11>>.

[DetNet-OAM] Mirsky, G., Theoleyre, F., Papadopoulos, G. Z., Bernardos, C. J., Varga, B., and J. Farkas, "Framework of Operations, Administration and Maintenance (OAM) for Deterministic Networking (DetNet)", Work in Progress, Internet-Draft, draft-ietf-detnet-oam-framework-06, 13 June 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-detnet-oam-framework-06>>.

[NASA] Adams, T., "RELIABILITY: Definition & Quantitative Illustration", <<https://kscddms.ksc.nasa.gov/Reliability/Documents/150814-3bWhatIsReliability.pdf>>.

Authors' Addresses

Pascal Thubert (editor)
Cisco Systems, Inc
Building D
45 Allée des Ormes - BP1200
06254 MOUGINS - Sophia Antipolis
France

Phone: [+33 497 23 26 34](tel:+33497232634)
Email: pthubert@cisco.com

Georgios Z. Papadopoulos
IMT Atlantique
Office B00 - 114A
2 Rue de la Chataigneraie
35510 Cesson-Sevigne - Rennes
France

Phone: [+33 299 12 70 04](tel:+33299127004)
Email: georgios.papadopoulos@imt-atlantique.fr