
Stream: Internet Engineering Task Force (IETF)
RFC: [9912](#)
Category: Informational
Published: April 2026
ISSN: 2070-1721
Author: P. Thubert, Ed.
Independent

RFC 9912

Reliable and Available Wireless (RAW) Architecture

Abstract

Reliable and Available Wireless (RAW) extends the reliability and availability of Deterministic Networking (DetNet) to networks composed of any combination of wired and wireless segments. The RAW architecture leverages and extends RFC 8655 ("Deterministic Networking Architecture") to adapt to challenges that prominently affect the wireless medium, notably intermittent transmission loss. This document defines a network control loop that optimizes the use of constrained bandwidth and energy while ensuring the expected DetNet services. The loop involves a new Point of Local Repair (PLR) function in the DetNet Service sub-layer that dynamically selects the DetNet path(s) for packets to route around local connectivity degradation.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are candidates for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9912>.

Copyright Notice

Copyright (c) 2026 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	4
2. The RAW Problem	5
3. Terminology	7
3.1. Abbreviations	8
3.2. Link and Direction	9
3.2.1. Flapping	9
3.2.2. Uplink	9
3.2.3. Downlink	9
3.2.4. Downstream	10
3.2.5. Upstream	10
3.3. Path and Recovery Graphs	10
3.3.1. Path	10
3.3.2. Recovery Graph	11
3.3.3. Forward and Crossing	13
3.3.4. Protection Path	13
3.3.5. Segment	13
3.4. Deterministic Networking	13
3.4.1. The DetNet Planes	13
3.4.2. Flow	14
3.4.3. Residence Time	14
3.4.4. L3 Deterministic Flow Identifier	14
3.4.5. Time-Sensitive Networking	14
3.4.6. Lower-Layer API	14

3.5. Reliability and Availability	14
3.5.1. Service Level Agreement	15
3.5.2. Service Level Objective	15
3.5.3. Service Level Indicator	15
3.5.4. Precision Availability Metrics	15
3.5.5. Reliability	15
3.5.6. Availability	15
4. Reliable and Available Wireless	15
4.1. High Availability Engineering Principles	15
4.1.1. Elimination of Single Points of Failure	16
4.1.2. Reliable Crossover	16
4.1.3. Prompt Notification of Failures	17
4.2. Applying Reliability Concepts to Networking	17
4.3. Wireless Effects Affecting Reliability	18
5. The RAW Conceptual Model	19
5.1. The RAW Planes	20
5.2. RAW Versus Upper and Lower Layers	22
5.3. RAW and DetNet	23
6. The RAW Control Loop	26
6.1. Routing Timescale Versus Forwarding Timescale	27
6.2. OODA Loop	29
6.3. Observe: RAW OAM	30
6.4. Orient: The RAW-Extended DetNet Operational Plane	31
6.5. Decide: The Point of Local Repair	31
6.6. Act: DetNet Path Selection and Reliability Functions	33
7. Security Considerations	33
7.1. Collocated Denial-of-Service Attacks	33
7.2. Layer 2 Encryption	34
7.3. Forced Access	34
8. IANA Considerations	34

9. References	34
9.1. Normative References	34
9.2. Informative References	35
Acknowledgments	37
Contributors	38
Author's Address	38

1. Introduction

Deterministic Networking (DetNet) aims to provide bounded latency and eliminate congestion loss, even when coexisting with best-effort traffic. It provides the ability to carry specified unicast or multicast data flows for real-time applications with extremely low packet loss rates and ensures maximum end-to-end delivery latency. A description of the general background and concepts of DetNet can be found in [[DetNet-ARCH](#)].

DetNet and the related IEEE 802.1 Time-Sensitive Networking (TSN) [[TSN](#)] initially focused on wired infrastructure, which provides a more stable communication channel than wireless networks. Wireless networks operate on a shared medium where uncontrolled interference, including self-induced multipath fading, may cause intermittent 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.

Nevertheless, deterministic capabilities are required in a number of wireless use cases as well [[RAW-USE-CASES](#)]. With scheduled radios such as Time-Slotted Channel Hopping (TSCH) and Orthogonal Frequency-Division Multiple Access (OFDMA) being developed to provide determinism over wireless links at the lower layers, providing DetNet capabilities has become possible. See [[RAW-TECHNOS](#)] for more on TSCH, OFDMA, and other technologies.

Reliable and Available Wireless (RAW) takes up the challenge of providing highly available and reliable end-to-end performances in a DetNet network that may include wireless segments. To achieve this, RAW leverages all possible transmission diversity and redundancy to ensure packet delivery, while optimizing the use of the shared spectrum to preserve bandwidth and save energy. To that effect, RAW defines protection paths that can be activated dynamically upon failures and a control loop that dynamically controls the activation and deactivation of the feasible protection paths to react quickly to intermittent losses.

The intent of RAW is to meet Service Level Objectives (SLOs) in terms of PDR, maximum contiguous losses, or latency boundaries for DetNet flows over mixes of wired and wireless networks, including wireless access and meshes (see [Section 2](#) for more on the RAW problem).

This document introduces and/or leverages terminology (see [Section 3](#)), principles (see [Section 4](#)), and concepts such as protection paths and recovery graphs to put together a conceptual model for RAW (see [Section 5](#)). Based on that model, this document elaborates on an in-network optimization control loop (see [Section 6](#)).

2. The RAW Problem

While the generic "Deterministic Networking Problem Statement" [[RFC8557](#)] applies to both wired and wireless media, the "Deterministic Networking Architecture" [[DetNet-ARCH](#)] must be extended to address less consistent transmissions, energy conservation, and shared spectrum efficiency.

Operating at Layer 3, RAW does not change the wireless technology at the lower layers. On the other hand, it 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., using diverse radio access technologies to optimize the end-to-end application experience.

RAW extends the DetNet services by providing elements that are specialized for transporting IP flows over deterministic radio technologies such as those listed in [[RAW-TECHNOS](#)]. Conceptually, RAW is agnostic to the lower layer, though the capability to control latency is assumed to ensure the DetNet services that RAW extends. How the lower layers are operated to do so (and whether a radio network is single hop or meshed, for example) are opaque to the IP layer and not part of the RAW abstraction. Nevertheless, cross-layer optimizations may take place to ensure proper link awareness (such as link quality) and packet handling (such as scheduling).

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 forwarding 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 4](#) or diverse radio access networks as illustrated in [Figure 10](#).

As opposed to wired links, the availability and performance of an individual wireless link cannot be trusted over the long term; it varies with transient service discontinuity, and any path that includes wireless hops is bound to face short periods of high loss. On the other hand, being broadcast in nature, the wireless medium provides capabilities that are atypical on modern wired links and that the RAW architecture can leverage opportunistically to improve the end-to-end reliability over a collection of paths.

Those capabilities include:

Promiscuous overhearing: Some wired and wireless technologies allow for multiple lower-layer attached nodes to receive the same packet sent by another node. This differs from a lower-layer network that is physically point-to-point, like a wire. With overhearing, more than one node in the forward direction of the packet may hear or overhear a transmission,

and the reception by one may compensate the loss by another. The concept of path can be revisited in favor of multipoint-to-multipoint progress in the forward direction and statistical chances of successful reception of any of the transmissions by any of the receivers.

L2-aware routing: As the quality and speed of a link varies over time, the concept of metric must also be revisited. Shortest-path cost loses its absolute value, and hop count turns into a bad idea as the link budget drops with the physical distance. Routing over radio requires both:

1. a new and more dynamic sense of link metrics, with new protocols such as the Dynamic Link Exchange Protocol (DLEP) and Layer 2 (L2) triggers to keep Layer 3 (L3) up to date with the link quality and availability, and
2. an approach to multipath routing, where multiple link metrics are considered since simple shortest-path cost loses its meaning with the instability of the metrics.

Redundant transmissions: Though feasible on any technology, proactive (forward) and reactive (acknowledgment-based) error correction is typical for wireless media. Bounded latency can still be obtained on a wireless link while operating those technologies, provided that link latency used in path selection allows for the extra transmission or the introduced delay is compensated along the path. In the case of coded fragments and retries, it makes sense to vary all the possible physical properties of the transmission to reduce the chances that the same effect causes the loss of both original and redundant transmissions.

Relay coordination and constructive interference: Though it can be difficult to achieve at high speed, a fine time synchronization and a precise sense of phase allows the energy from multiple coordinated senders to add up at the receiver and actually improve the signal quality, compensating for either distance or physical objects in the Fresnel zone that would reduce the link budget. From a DetNet perspective, this may translate to taking into account how transmission from one node may interfere with the transmission of another node attached to the same wireless sub-layer network.

RAW and DetNet enable application flows that require a special treatment along paths that can provide that treatment. This may be seen as a form of Path Aware networking and may be subject to impediments documented in [\[RFC9049\]](#).

The mechanism used to establish a path is not unique to, or necessarily impacted by, RAW. The mechanism is expected to be the product of the DetNet Controller Plane [\[DetNet-PLANE\]](#); it may use a Path Computation Element (PCE) [\[RFC4655\]](#) or the DetNet YANG data model [\[RFC9633\]](#), or it may be computed in a distributed fashion as per the Resource ReSerVation Protocol (RSVP) [\[RFC2205\]](#). Either way, the assumption is that it is slow relative to local forwarding operations along the path. To react fast enough to transient changes in the radio transmissions, RAW leverages DetNet Network Plane enhancements to optimize the use of the paths and match the quality of the transmissions over time.

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; thus, in the wireless case, it makes sense to provide redundant forwarding solutions along alternate paths (see [Section 3.3](#))

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. The RAW Network Plane operations happen within the scope of a recovery graph (see [Section 3.3.2](#)) that is pre-established and installed by means outside of the scope of RAW. A recovery graph may be strict or loose depending on whether each hop or just a subset of the hops is observed and controlled by RAW.

RAW distinguishes the longer timescale at which routes are computed from the shorter timescale where forwarding decisions are made (see [Section 6.1](#)). The RAW Network Plane operations happen at a timescale that sits timewise between the routing and the forwarding timescales. Within the resources delineated by a recovery graph, their goal is to dynamically select the protection path(s) that the upcoming packets of a DetNet flow shall follow. As they influence the path for the entirety of the flows or a portion of them, the RAW Network Plane operations may affect the metrics used in their rerouting decisions, which could potentially lead to oscillations; such effects must be avoided or dampened.

3. Terminology

RAW reuses terminology defined for DetNet in "Deterministic Networking Architecture" [[DetNet-ARCH](#)], e.g., "PREOF" to stand for "Packet Replication, Elimination, and Ordering Functions". RAW inherits and augments IETF recovery mechanisms such as the ones provided in DetNet [[DetNet-ARCH](#)] and in Traffic Engineering, e.g., [[RFC4090](#)].

RAW also reuses terminology defined for Operations, Administration, and Maintenance (OAM) protocols in [Section 1.1](#) of "Framework of Operations, Administration, and Maintenance (OAM) for Deterministic Networking (DetNet)" [[DetNet-OAM](#)] and in "Active and Passive Metrics and Methods (with Hybrid Types In-Between)" [[RFC7799](#)].

RAW also reuses terminology defined for MPLS in [[RFC4427](#)], such as the term "recovery" to cover both protection and restoration for a number of recovery types. That document defines a number of concepts, such as the recovery domain, that are used in RAW mechanisms and defines the new term "recovery graph". A recovery graph associates a topological graph with usage metadata that represents how the paths are built and used within the recovery graph. The recovery graph provides excess bandwidth for the intended traffic over alternate potential paths, and the use of that bandwidth is optimized dynamically.

The concept of a recovery graph is agnostic to the underlying technology and applies, but is not limited to, any full or partial wireless mesh. RAW specifies strict and loose recovery graphs 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 also reuses terminology defined for RSVP-TE in [[RFC4090](#)], such as the "Point of Local Repair (PLR)". The concept of a backup path is generalized with protection path, which is the term mostly found in recent standards and used in this document.

RAW also reuses terminology defined for 6TiSCH in [[6TiSCH-ARCH](#)] and equates the 6TiSCH concept of a Track with that of a recovery graph.

3.1. Abbreviations

RAW uses the following abbreviations.

ARQ

Automatic Repeat Request. A well-known mechanism that enables an acknowledged transmission with retries to mitigate errors and loss. ARQ may be implemented at various layers in a network. ARQ is typically implemented per hop (not end to end) at Layer 2 in wireless networks. ARQ improves delivery on lossy wireless. Additionally, ARQ retransmission may be further limited by a bounded time to meet end-to-end packet latency constraints. Additional details and considerations for ARQ are detailed in [\[RFC3366\]](#).

FEC

Forward Error Correction. Adding redundant data to protect against a partial loss without retries.

HARQ

Hybrid ARQ. A combination of FEC and ARQ.

ETX

Expected Transmission Count. A statistical metric that represents the expected total number of packet transmissions (including retransmissions) required to successfully deliver a packet along a path, used by 6TiSCH [\[RFC6551\]](#).

ISM

Industrial, Scientific, and Medical. Refers to a group of radio bands or parts of the radio spectrum (e.g., 2.4 GHz and 5 GHz) that are internationally reserved for the use of radio frequency (RF) energy intended for industrial, scientific, and medical requirements (e.g., by microwaves, depth radars, and medical diathermy machines). Cordless phones, Bluetooth and Low-Power Wireless Personal Area Network (LoWPAN) devices, near-field communication (NFC) devices, garage door openers, baby monitors, and Wi-Fi networks may all use the ISM frequencies, although these low-power transmitters are not considered to be ISM devices. In general, communications equipment operating in ISM bands must tolerate any interference generated by ISM applications, and users have no regulatory protection from ISM device operation in these bands.

PER

Packet Error Rate. The ratio of the number of packets received in error to the total number of transmitted packets. A packet is considered to be in error if even a single bit within the packet is received incorrectly.

PDR

Packet Delivery Ratio (PDR). The ratio of the number of successfully delivered data packets to the total number of packets transmitted from the sender to the receiver.

RSSI

Received Signal Strength Indication. Also known as "Energy Detection Level". A measure of the incoherent (raw) RF power in a channel. The RF power can come from any source: other transmitters using the same technology, other radio technology using the same band, or background radiation. For a single hop, RSSI may be used for LQI.

LQI

Link Quality Indicator. An indication of the quality of the data packets received by the receiver. It is typically derived from packet error statistics, with the exact method depending on the network stack being used. LQI values may be exposed to the Controller Plane for each individual hop or cumulated along segments. Outgoing LQI values can be calculated from coherent (demodulated) PER, RSSI, and incoming LQI values.

OAM

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

OODA

Observe, Orient, Decide, Act. A generic formalism to represent the operational steps in a control loop. In the context of RAW, OODA is applied to network control and convergence; see [Section 6.2](#) for more.

SNR

Signal-to-Noise Ratio. Also known as "S/N Ratio". A measure used in science and engineering that compares the level of a desired signal to the level of background noise. SNR is defined as the ratio of signal power to noise power, often expressed in decibels.

3.2. Link and Direction

This document uses the following terms relating to links and direction in the context of RAW.

3.2.1. 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 a duration of a subsecond to seconds.

3.2.2. Uplink

An uplink is the connection from end devices to 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) and a Base Station (BS) such as a 3GPP 5G gNodeB (gNB).

3.2.3. Downlink

A downlink is the reverse direction from uplink.

3.2.4. Downstream

Downstream refers to the following the direction of the flow data path along a recovery graph.

3.2.5. Upstream

Upstream refers to going against the direction of the flow data path along a recovery graph.

3.3. Path and Recovery Graphs

This document uses the following terms relating to paths and recovery graphs in the context of RAW.

3.3.1. Path

[Section 1.3.3](#) of [\[INT-ARCH\]](#) provides a definition of path:

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 [\[RFC9473\]](#) 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.

Paths are unidirectional and time dependent, i.e., there can be a variety of paths from one node to another, and the path over which packets are transmitted may change. A path definition can be strict (i.e., the exact sequence of path elements remains the same) or loose (i.e., the start and end node remain the same, but the path elements between them may vary over time).

The representation of a path and its properties may depend on the entity considering the path. On the one hand, the representation may differ due to entities having partial visibility of path elements comprising a path or their visibility changing over time.

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. 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 if D eliminates duplicates, a packet from A to B can experience two paths: A->C->E->D->B and A->C->F->D->B. Those paths are called protection paths. Protection paths may be fully non-congruent; alternatively, they may intersect at replication or elimination points.

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 this complex scenario as a DetNet path. As such, the DetNet path extends the above description of a path, but it still matches the experience of a packet that traverses the network.

With RAW, the path experienced by a packet is subject to change from one packet to the next, but all the possible experiences are all contained within a finite set. Therefore, we introduce the term "recovery graph" (see the next section) that coalesces that set and covers the overall topology where the possible DetNet paths are all contained. As such, the recovery graph coalesces all the possible paths a flow may experience, each with its own statistical probability to be used.

3.3.2. Recovery Graph

A recovery graph is a networking graph that can be followed to transport packets with equivalent treatment and is associated with usage metadata. In contrast to the definition of a path above, a recovery graph represents a potential path, not an actual one. Also, a recovery graph is not necessarily a linear sequence like a simple path and is not necessarily fully traversed (flooded) by all packets of a flow like a DetNet path. Still, and as a simplification, the casual reader may consider that a recovery graph is very much like a DetNet path, aggregating multiple paths that may overlap or fork and then rejoin, for instance, to enable a protection service by the PREOF operations.

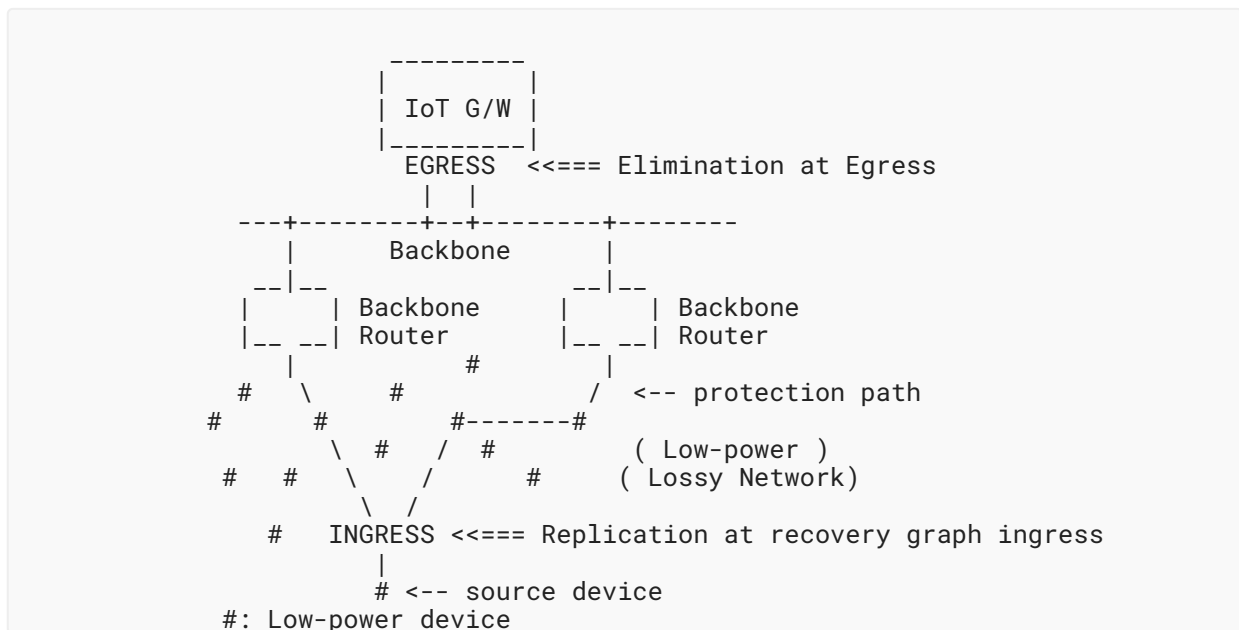


Figure 1: Example IoT Recovery Graph to an IoT Gateway with 1+1 Redundancy

Refining further, a recovery graph is defined as the coalescence of all the feasible DetNet paths that a packet with an assigned flow may be forwarded along. A packet that is assigned to the recovery graph experiences one of the feasible DetNet paths based on the current selection by the PLR at the time the packet traverses the network.

Refining even further, the feasible DetNet paths within the recovery graph may or may not be computed in advance; instead, they may be decided upon the detection of a change from a clean slate. Furthermore, the PLR decision may be distributed, which yields a large combination of possible and dependent decisions, with no node in the network capable of reporting which is the current DetNet path within the recovery graph.

In DetNet [DetNet-ARCH] terms, a recovery graph has the following properties:

- A recovery graph is a Layer 3 abstraction built upon IP links between routers. A router may form multiple IP links over a single radio interface.
- A recovery graph has one ingress and one egress node, which operate as DetNet edge nodes.
- A recovery graph 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 a recovery graph are DetNet relay nodes that operate at the DetNet Service sub-layer and provide the PREOF functions.
- The topological edges of a recovery graph are strict sequences of DetNet transit nodes that operate at the DetNet forwarding sub-layer.

Figure 2 illustrates the generic concept of a recovery graph, between an ingress node and an egress node. The recovery graph is composed of forward protection paths, forward segments, and crossing segments (see the definitions of those terms in the next sections). The recovery graph contains at least two protection paths: a main path and a backup path.

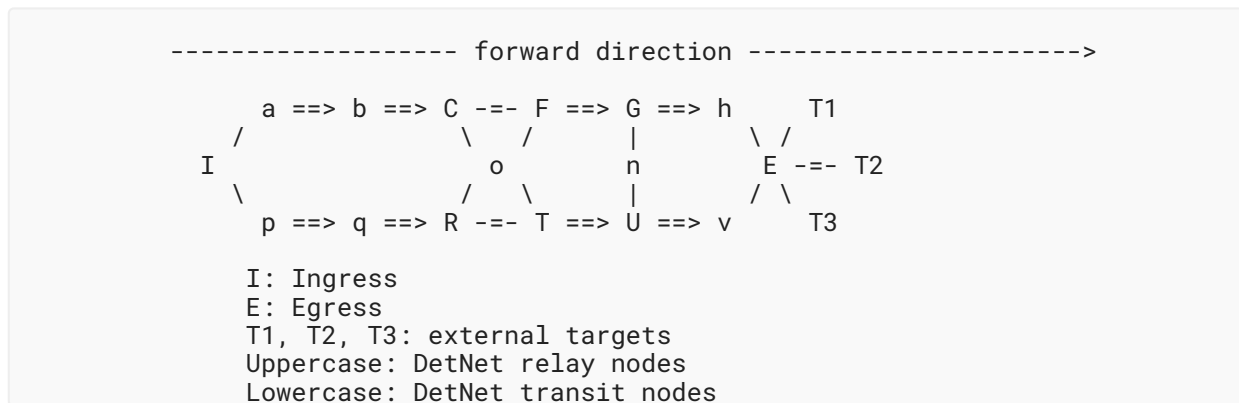


Figure 2: A Recovery Graph and Its Components

Of note:

I ==> a ==> b ==> C: A forward segment to targets F and o

C ==> o ==> T: A forward segment to target T (and/or U)

G | n | U: A crossing segment to targets G or U

I -> F -> E: A forward protection path to targets T1, T2, and T3

I, a, b, C, F, G, h, E: A path to T1, T2, and/or T3

I, p, q, R, o, F, G, h, E: A segment-crossing protection path

3.3.3. Forward and Crossing

Forward refers to progress towards the egress of the recovery graph. Forward links are directional, and packets that are forwarded along the recovery graph can only be transmitted along the link direction. Crossing links are bidirectional, meaning that they can be used in both directions, though a given packet may use the link in one direction only. A segment can be forward, in which case it is composed of forward links only, or it can be crossing, in which case it is composed of crossing links only. A protection path is always forward, meaning that it is composed of forward links and segments.

3.3.4. Protection Path

A protection path is an end-to-end forward path between the ingress and egress nodes of a recovery graph. A protection path in a recovery graph is expressed as a strict sequence of DetNet relay nodes or as a loose sequence of DetNet relay nodes that are joined by segments in the recovery graph. Background information on the concepts related to protection paths can be found in [\[RFC4427\]](#) and [\[RFC6378\]](#).

3.3.5. Segment

A segment is a strict sequence of DetNet transit nodes between two DetNet relay nodes; a segment of a recovery graph is composed topologically of two vertices of the recovery graph and one edge of the recovery graph between those vertices.

3.4. Deterministic Networking

This document reuses the terminology in [Section 2](#) of [\[RFC8557\]](#) and [Section 4.1.2](#) of [\[DetNet-ARCH\]](#) for deterministic networking and deterministic networks. This document also uses the following terms.

3.4.1. The DetNet Planes

[\[DetNet-ARCH\]](#) defines three planes: the Application (User) Plane, the Controller Plane, and the Network Plane. The DetNet Network Plane is composed of a Data Plane (packet forwarding) and an Operational Plane where OAM operations take place. In the Network Plane, the DetNet Service sub-layer focuses on flow protection (e.g., using redundancy) and can be fully operated at Layer 3, while the DetNet forwarding sub-layer establishes the paths, 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. For more information, see [Section 2](#).

3.4.2. Flow

A flow is a collection of consecutive IP packets defined by the upper layers and signaled by the same 5-tuple or 6-tuple (see [Section 5.1](#) of [\[RFC8939\]](#)). Packets of the same flow must be placed on the same recovery graph to receive an equivalent treatment from ingress to egress within the recovery graph. Multiple flows may be transported along the same recovery graph. The DetNet path that is selected for the flow may change over time under the control of the PLR.

3.4.3. Residence Time

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

3.4.4. L3 Deterministic Flow Identifier

The classic 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 Differentiated Services Code Point (DSCP) field in the packet (see [Section 3.3](#) of [\[DetNet-DP\]](#)). The IPv6 flow label is not used for that purpose.

3.4.5. Time-Sensitive Networking

Time-Sensitive Networking (TSN) denotes the IEEE 802 efforts regarding deterministic networking, originally for use on Ethernet. See [\[TSN\]](#). Wireless TSN (WTSN) denotes extensions of the TSN work on wireless media, e.g., the RAW technologies described in [\[RAW-TECHNOS\]](#).

3.4.6. Lower-Layer API

RAW includes the concept of a lower-layer API (LL API) that provides an interface between the lower-layer (e.g., wireless) technology and the DetNet layers. The LL API is technology dependent as what the lower layers expose towards the DetNet layers may vary. Furthermore, different RAW technologies are equipped with different reliability features (e.g., short-range broadcast, Multiple User - Multiple Input Multiple Output (MU-MIMO), physical layer (PHY) rate and other Modulation Coding Scheme (MCS) adaptation, coding and retransmissions methods, and constructive interference and overhearing; see [\[RAW-TECHNOS\]](#) for more details). The LL API enables interactions between the reliability functions provided by the lower layer and the reliability functions provided by DetNet. That is, the LL API makes cross-layer optimization possible for the reliability functions of different layers depending on the actual exposure provided via the LL API by the given RAW technology. The [Dynamic Link Exchange Protocol \(DLEP\)](#) [\[DLEP\]](#) is an example of a protocol that can be used to implement the LL API.

3.5. Reliability and Availability

In the context of the RAW work, reliability and availability are defined as follows, along with the following other terms.

3.5.1. Service Level Agreement

In the context of RAW, a Service Level Agreement (SLA) is a contract between a provider (the network) and a client (the application flow) that defines measurable metrics such as latency boundaries, consecutive losses, and Packet Delivery Ratio (PDR).

3.5.2. 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 packet redundancy addresses an SLO of consecutive losses in a row by augmenting the chances of delivery of a packet that follows a loss.

3.5.3. Service Level Indicator

A Service Level Indicator (SLI) measures the compliance of an SLO to the terms of the contract. For instance, it can be the statistics of individual losses or losses in a row during a certain amount of time.

3.5.4. Precision Availability Metrics

Precision Availability Metrics (PAMs) [\[RFC9544\]](#) aim to capture service levels for a flow, specifically the degree to which the flow complies with the SLOs that are in effect.

3.5.5. Reliability

Reliability is a measure of the probability that an item (e.g., system or network) will perform its intended function with no failure for a stated period of time (or for a stated number of demands or load) under stated environmental conditions. In other words, reliability is the probability that an item will be in an uptime state (i.e., fully operational or ready to perform) for a stated mission (e.g., to provide an SLA). See more in [\[NASA1\]](#).

3.5.6. Availability

Availability is the probability of an item's (e.g., a network's) mission readiness (e.g., to provide an SLA). Availability is expressed as $(\text{uptime})/(\text{uptime}+\text{downtime})$. Note that it is availability that addresses downtime (including time for maintenance, repair, and replacement activities) and not reliability. See more in [\[NASA2\]](#).

4. Reliable and Available Wireless

4.1. High Availability Engineering Principles

The reliability criteria of a critical system pervade 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 each single point 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. Both non-Internet and Internet examples are included.

4.1.1. 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 reroute to alternate routes in case of a failure. When links are cabled through the same conduit, they form a Shared Risk Link Group (SRLG) and share the same fate if the conduit is cut, making the reroute operation ineffective. The same effect can happen with virtual links that end up in the same physical transport through the intricacies of nested encapsulation. In the 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. Thus, for high availability, the use of physically link-disjoint and node-disjoint paths is required; in the wireless space, the use of the highest possible degree of diversity (time, space, code, frequency, and channel width) in the transmissions over the air is also required to combat the additional causes of transmission loss.

From an economics standpoint, executing this principle properly generally increases capital 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 efficiently.

4.1.2. Reliable Crossover

Backup equipment has limited value unless it can be reliably switched into use within the downtime parameters. IP routers execute reliable crossover continuously because the routers 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. "IP Fast Reroute Framework" [[FRR](#)] analyzes mechanisms for fast failure detection and path repair for IP Fast Reroute (FRR) 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 assured SLOs 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 FRR 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 does 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.

4.1.3. Prompt Notification of Failures

The execution of the two above principles is likely to render a system where the end user rarely sees a failure. However, a failure that occurs must still be detected in order to direct maintenance.

There are many reasons for system monitoring (Fault, Configuration, Accounting, Performance, and Security (FCAPS) is a handy mental checklist), but fault monitoring is a sufficient reason.

"Overview and Principles of Internet Traffic Engineering" [TE] discusses the importance of measurement for network protection and provides an abstract method for network survivability with the analysis of a traffic matrix as observed via a network management YANG data model, probing techniques, file transfers, 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 recovery graph based on statistical and aggregated information. RAW itself operates in the DetNet Network Plane at a faster timescale 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 DLEP, or transported over multiple hops using OAM and reverse OAM, as illustrated in [Figure 11](#).

4.2. Applying Reliability Concepts to Networking

The terms "reliability" and "availability" are defined for use in RAW in [Section 3](#), and the reader is invited to read [\[NASA1\]](#) and [\[NASA2\]](#) 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 packet but does not affect the previous packet, the next packet, or packets 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 be fully avoided, and the systems are built to resist some loss. This can be done by using redundancy with retries (as in HARQ), Packet Replication and Elimination (PRE) FEC, and Network Coding (e.g., using FEC with Static Context Header Compression (SCHC) [RFC8724] fragments). Also, in a typical control loop, linear interpolation from the previous measurements can be used.

However, the linear interpolation method cannot resist multiple consecutive losses, and a high MTBF is desired as a guarantee that the number of losses in a row is bounded. In this case, what is really desired is a Maximum Consecutive Loss (MCL). If the number of losses in a row passes the MCL, 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 as the MTBF and as a proxy to indicate an MCL, e.g., as described in Section 7.4 of "Deterministic Networking Use Cases" [RFC8578].

4.3. 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 either add up to (constructive interference) or cancel (destructive interference) the direct signal.

The affected frequencies depend on the relative position of the sender, the receiver, and all the reflecting objects in the environment. A given hop suffers from multipath fading for multiple packets in a row until 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 Fresnel zone is an elliptical region of space between and around the transmit and receive antennas in a point-to-point wireless communication. The optimal transmission happens when it is free of obstacles.

In an environment that is rich in metallic structures and mobile objects, a single radio link provides 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 in the ISM band) keeps transmitting, a continuous stream of packets are 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) leveraging diverse PHY technologies (e.g., narrowband versus spread spectrum or diverse codes). Using time diversity defeats short-term interferences; spatial diversity combats very local causes of interference 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.

5. The RAW Conceptual Model

RAW extends the conceptual model described in Section 4 of "Deterministic Networking Architecture" [DetNet-ARCH] with the PLR at the Service sub-layer, as illustrated in Figure 3. The PLR (see Section 6.5) provides additional agility against transmission loss. For example, the PLR can act based on indications from the lower layer or based on OAM.

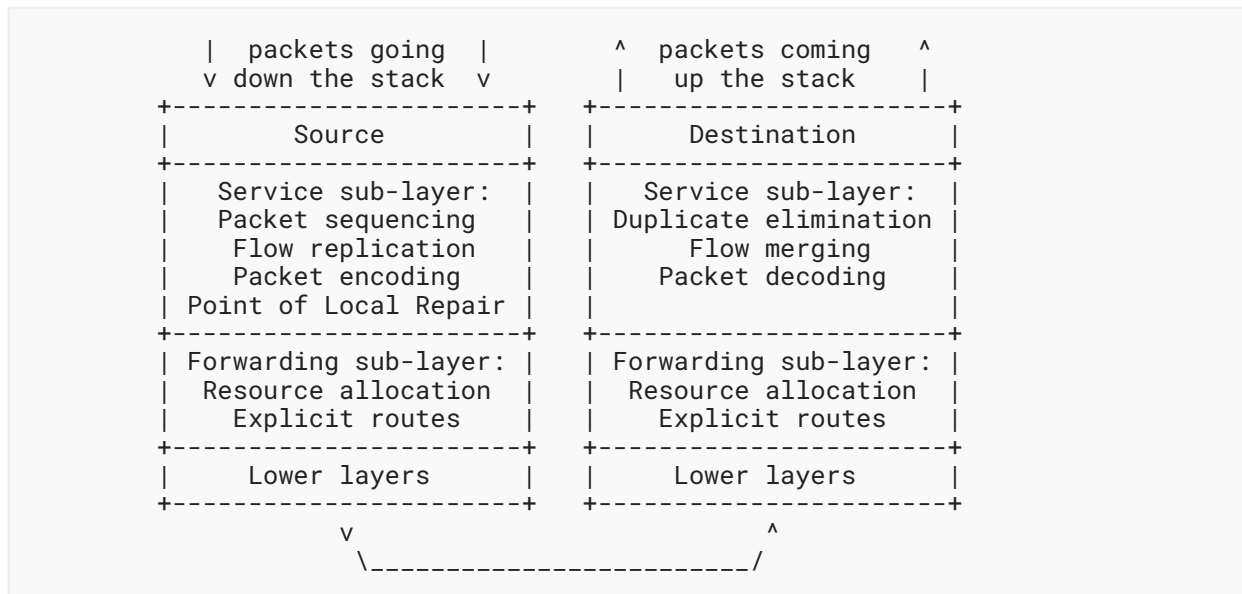


Figure 3: Extended DetNet Data Plane Protocol Stack

5.1. The RAW Planes

The RAW nodes are DetNet relay nodes that operate in the RAW Network Plane and are capable of additional diversity mechanisms and measurement functions related to the radio interface. RAW leverages an Operational Plane orientation function (that typically operates inside the ingress edge nodes) to dynamically adapt the path of the packets and optimize the resource usage.

In the case of centralized routing operations, the RAW Controller Plane Function (CPF) interacts with RAW nodes over a Southbound API. It consumes data and information from the network and generates knowledge and wisdom to help steer the traffic optimally inside a recovery graph.

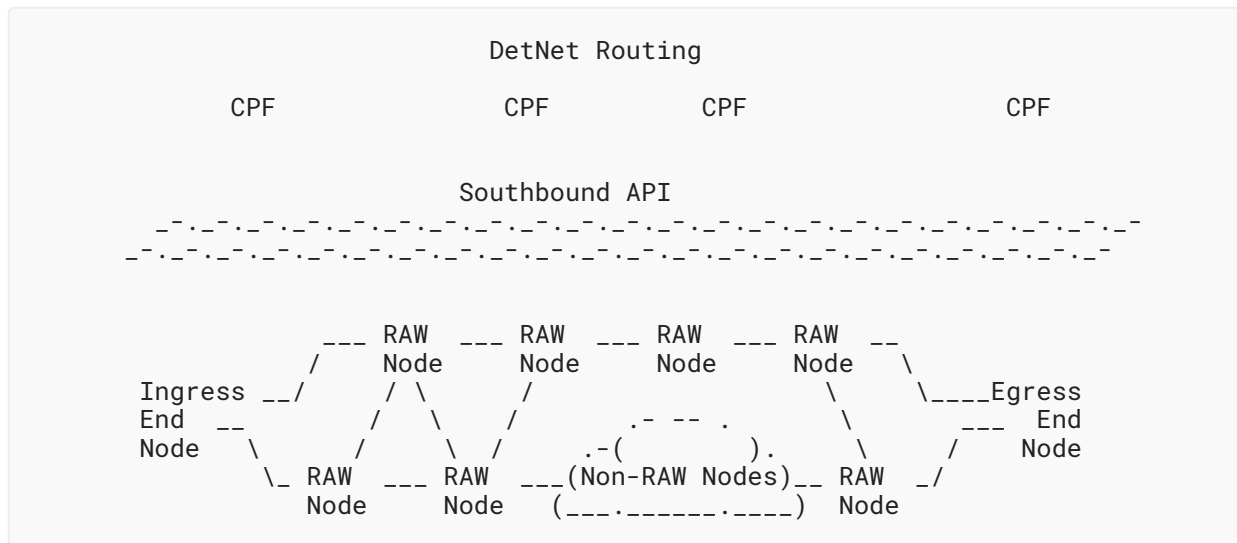


Figure 4: RAW Nodes (Centralized Routing Case)

When a new flow is defined, the routing function uses its current knowledge of the network to build a new recovery graph between an ingress End System and an egress End System for that flow; it indicates to the RAW nodes where the PREOF and/or radio diversity and reliability operations may be actioned in the Network Plane.

- The recovery graph may be strict, meaning that the DetNet forwarding sub-layer operations are enforced end to end.
- The recovery graph may be expressed loosely to enable traversing a non-RAW subnetwork as in Figure 7. In that case, RAW cannot leverage end-to-end DetNet and cannot provide latency guarantees.

The information that the orientation function reports to the routing function may be time aggregated (e.g., statistical), to match the longer-term operation of the routing function. Example information includes link-layer metrics such as link bandwidth (the medium speed depends dynamically on the mode of the PHY layer), number of flows (bandwidth that can be reserved for a flow depends on the number and size of flows sharing the spectrum), and the average and mean squared deviation of availability and reliability metrics (such as PDR) over long periods of time. It may also report an aggregated Expected Transmission Count (ETX) or a variation of it.

Based on those metrics, the routing function installs the recovery graph with enough redundant forwarding solutions to ensure that the Network Plane can reliably deliver the packets within an SLA associated with the flows that it transports. The SLA defines end-to-end reliability and availability requirements, in which reliability may be expressed as a successful delivery in order and within a bounded delay of at least one copy of a packet.

Depending on the use case and the SLA, the recovery graph may comprise non-RAW segments, either interleaved inside the recovery graph (e.g., over tunnels) or all the way to the egress End node (e.g., a server in the local wired domain). RAW observes the lower-layer links between RAW nodes (typically radio links) and the end-to-end network-layer operation to decide at all times which of the diversity schemes is actioned by which RAW nodes.

Once a recovery graph is established, per-segment and end-to-end reliability and availability statistics are periodically reported to the routing function to ensure that the SLA can be met; if not, then the recovery graph is recomputed.

5.2. RAW Versus Upper and Lower Layers

RAW builds on DetNet-provided features such as scheduling and shaping. In particular, RAW inherits the DetNet guarantees on end-to-end latency, which can be tuned to ensure that DetNet and RAW reliability mechanisms have no side effect on upper layers, e.g., on transport-layer packet recovery. RAW operations include possible rerouting, which in turn may affect the ordering of a burst of packets. RAW also inherits PREOF from DetNet, which can be used to reorder packets before delivery to the upper layers. As a result, DetNet in general and RAW in particular offer a smoother transport experience for the upper layers than the Internet at large, with ultra-low jitter and loss.

RAW improves the reliability of transmissions and the availability of the communication resources and should be seen as a dynamic optimization of the use of redundancy to maintain reliability and availability metrics within certain boundaries. For instance, ARQ (which provides one-hop reliability through acknowledgements and retries) and FEC codes (such as turbo codes which reduce the PER) are typically operated at Layer 2 and Layer 1, respectively. In both cases, redundant transmissions improve the one-hop reliability at the expense of energy and latency, which are the resources that RAW must control. In order to achieve its goals, RAW may leverage the lower-layer operations by abstracting the concept and providing hints to the lower layers on the desired outcome (e.g., in terms of reliability and timeliness), as opposed to performing the actual operations at Layer 3.

Guarantees such as bounded latency depend on the upper layers (transport or application) to provide the payload in volumes and at times that match the contract with the DetNet sub-layers and the layers below. An excess of incoming traffic at the DetNet ingress may result in dropping or queueing of packets and can entail loss, latency, or jitter; this violates the guarantees that are provided inside the DetNet Network.

When the traffic from upper layers matches the expectation of the lower layers, RAW still depends on DetNet mechanisms and the lower layers to provide the timing and physical resource guarantees that are needed to match the traffic SLA. When the availability of the physical resource varies, RAW acts on the distribution of the traffic to leverage alternates within a finite set of potential resources.

The Operational Plane elements (routing and OAM control) may gather aggregated information from lower layers (e.g., information about link quality), via measurement or communication with the lower layer. This information may be obtained from inside the device using specialized

APIs (e.g., L2 triggers) via monitoring and measurement protocols such as Bidirectional Forwarding Detection (BFD) [RFC5880] and Simple Two-way Active Measurement Protocol (STAMP) [RFC8762], respectively, or via a control protocol exchange with the lower layer (e.g., DLEP [DLEP]). It may then be processed and exported through OAM messaging or via a YANG data model and exposed to the Controller Plane.

5.3. RAW and DetNet

RAW leverages the DetNet forwarding sub-layer and requires the support of OAM in DetNet transit nodes (see Figure 3 of [DetNet-ARCH]) for the dynamic acquisition of link capacity and state to maintain a strict RAW service end to end over a DetNet Network. In turn, DetNet and thus RAW may benefit from or leverage functionality such as that provided by TSN at the lower layers.

RAW extends DetNet to improve the protection against link errors such as transient flapping that are far more common in wireless links. Nevertheless, for the most part, the RAW methods are applicable to wired links as well, e.g., when energy savings are desirable and the available path diversity exceeds 1+1 linear redundancy.

RAW adds sub-layer functions that operate in the DetNet Operational Plane, which is part of the Network Plane. The RAW orientation function may run only in the DetNet edge nodes (ingress edge node or End System), or it can also run in DetNet relay nodes when the RAW operations are distributed along the recovery graph. The RAW Service sub-layer includes the PLR, which decides the DetNet path for the future packets of a flow along the DetNet path, Maintenance End Points (MEPs) on edge nodes, and Maintenance Intermediate Points (MIPs) within. The MEPs trigger, and learn from, OAM observations and feed the PLR for its next decision.

As illustrated in Figure 5, RAW extends the DetNet Stack (see Figure 4 of [DetNet-ARCH] and Figure 3) with additional functionality at the DetNet Service sub-layer for the actuation of PREOF based on the PLR decision. DetNet operates at Layer 3, leveraging abstractions of the lower layers and APIs that control those abstractions. For instance, DetNet already leverages lower layers for time-sensitive operations such as time synchronization and traffic shapers. As the performances of the radio layers are subject to rapid changes, RAW needs more dynamic gauges and knobs. To that effect, the LL API provides an abstraction to the DetNet layer that can be used to push reliability and timing hints, like suggesting X retries (min, max) within a time window or sending unicast (one next hop) or multicast (for overhearing). In the other direction up the stack, the RAW PLR needs hints about the radio conditions such as L2 triggers (e.g., RSSI, LQI, or ETX) over all the wireless hops.

RAW uses various OAM functionalities at the different layers. For instance, the OAM function in the DetNet Service sub-layer may perform Active and/or Hybrid OAM to estimate the link and path availability, either end to end or limited to a segment. The RAW functions may be present in the Service sub-layer in DetNet edge and relay nodes.

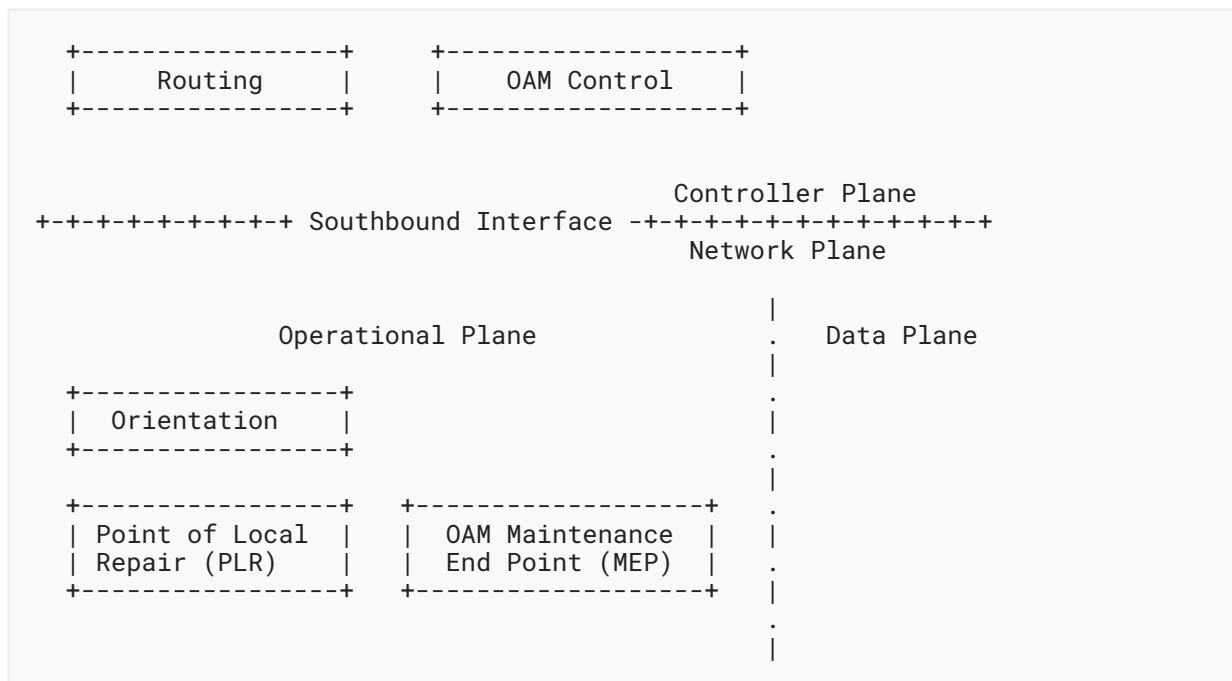


Figure 5: RAW Function Placement (Centralized Routing Case)

There are two main proposed models to deploy RAW and DetNet: strict (Figure 6) and loose (Figure 7). In the strict model, illustrated in Figure 6, RAW operates over a continuous DetNet service end to end between the ingress and the egress edge nodes or End Systems.

In the loose model, illustrated in Figure 7, RAW may traverse a section of the network that is not serviced by DetNet. RAW OAM may observe the end-to-end traffic and make the best of the available resources, but it may not expect the DetNet guarantees over all paths. For instance, the packets between two wireless entities may be relayed over a wired infrastructure, in which 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 minimal forwarding sub-layer service is provided at all DetNet nodes to ensure that the OAM information flows. DetNet relay nodes may or may not support RAW services, whereas the DetNet edge nodes are required to support RAW in any case. DetNet guarantees, such as bounded latency, are provided end to end. RAW extends the DetNet Service sub-layer to optimize the use of resources.

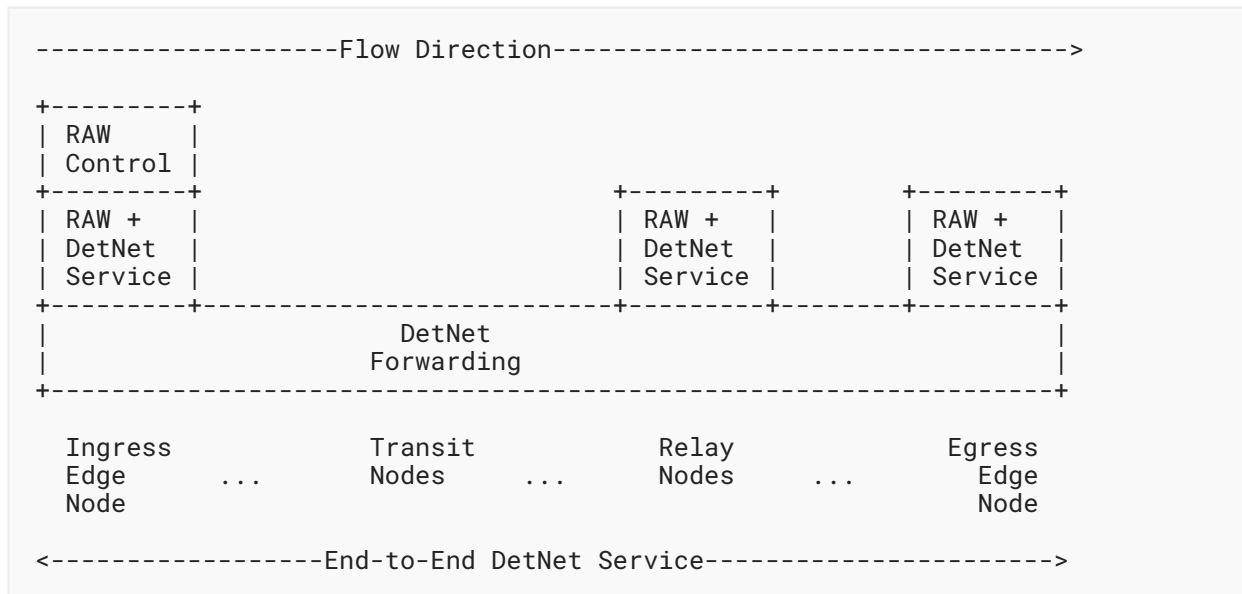


Figure 6: RAW over DetNet (Strict Model)

In the loose model (illustrated in [Figure 7](#)), 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 the 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 7: RAW over DetNet (Loose Model)

6. The RAW Control Loop

The RAW architecture is based on an abstract OODA loop that controls the operation of a recovery graph. The generic concept involves the following:

1. Operational Plane measurement protocols allow OAM to observe (like the first "O" in "OODA") some or all hops along a recovery graph as well as the end-to-end packet delivery.
2. The DetNet Controller Plane establishes primary and protection paths for use by the RAW Network Plane. The orientation function reports data and information such as link statistics to be used by the routing function to compute, install, and maintain the recovery graphs. The routing function may also generate intelligence such as a trained model for link quality prediction, which in turn can be used by the orientation function (like the second "O" in "OODA") to influence the path selection by the PLR within the RAW OODA loop.
3. A PLR operates at the DetNet Service sub-layer and hosts the decision function (like the "D" in "OODA"). The decision function determines which DetNet paths will be used for future packets that are routed within the recovery graph.
4. Service protection actions are actuated or triggered over the LL API by the PLR to increase the reliability of the end-to-end transmissions. The RAW architecture also covers in-situ signaling that is embedded within live user traffic [RFC9378] (e.g., via OAM) when the decision is acted (like the "A" in "OODA") upon by a node that is downstream in the recovery graph from the PLR.

The overall OODA loop optimizes the use of redundancy to achieve the required reliability and availability SLO(s) while minimizing the use of constrained resources such as spectrum and battery.

6.1. Routing Timescale Versus Forwarding Timescale

With DetNet, the Controller Plane Function (CPF) handles the routing computation and maintenance. With RAW, the routing operation is segregated from the RAW control loop, so it may reside in the Controller Plane, whereas the control loop itself happens in the Network Plane. To achieve RAW capabilities, the routing operation is extended to generate the information required by the orientation function in the loop. For example, the routing function may propose DetNet paths to be used as a reflex action in response to network events or provide an aggregated history that the orientation function can use to make a decision.

In a wireless mesh, the path to a routing function located in the Controller Plane can be expensive and slow, possibly going across the whole mesh and back. Reaching the Controller Plane can also be slow in regard to the speed of events that affect the forwarding operation in the Network Plane at the radio layer. Note that a distributed routing protocol may also take time and consume excessive wireless resources to reconverge to a new optimized state.

As a result, the DetNet routing function is not expected to be aware of and react to very 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.

The interaction between the network nodes and the routing function is handled by the orientation function, which reports to the routing function and sends control information in a digested form back to the RAW node to be used inside a forwarding control loop for traffic steering.

[Figure 8](#) illustrates a Network Plane recovery graph with links P-Q and N-E flapping, possibly in a transient fashion due to short-term interferences and possibly for a longer time (e.g., due to obstacles between the sender and the receiver or hardware failures). In order to maintain a received redundancy around a value of 2 (for instance), RAW may leverage a higher ARQ on these hops if the overall latency permits the extra delay or enable alternate paths between ingress I and egress E. For instance, RAW may enable protection path I ==> F ==> N ==> Q ==> M ==> R ==> E that routes around both issues and provides some degree of spatial diversity with protection path I ==> A ==> B ==> C ==> D ==> E.

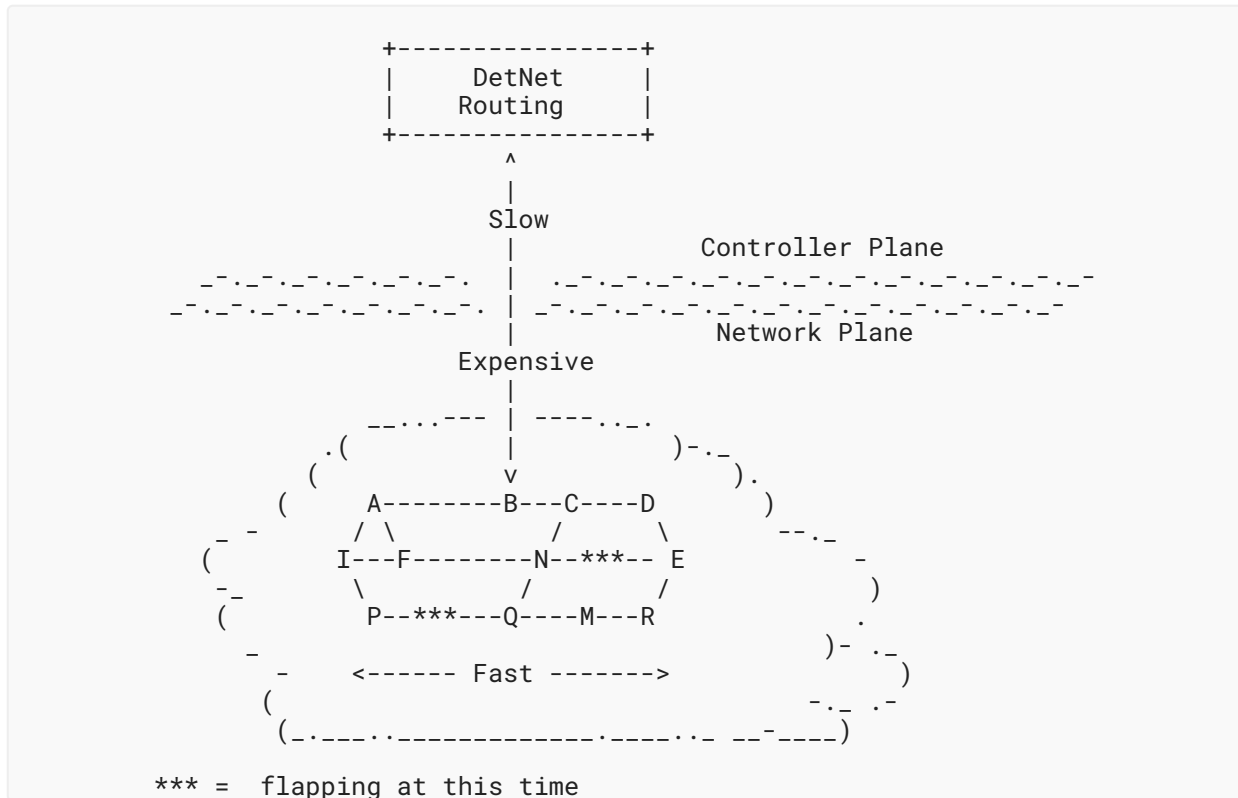


Figure 8: Timescales

In the case of wireless, the changes that affect the forwarding decision can happen frequently and often for short durations. An example of this is a mobile object that moves between a transmitter and a receiver and cancels the line-of-sight transmission for a few seconds. Another example is radar that measures the depth of a pool using the ISM band and interferes on a particular channel for a split second.

Thus, there is 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 recovery graph that enables multiple Unequal-Cost Multipath (UCMP) forwarding solutions along so-called protection paths and leaves it to the Network Plane to make the short-term decision of which of these possibilities should be used for which upcoming packets and flows.

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 Multiprotocol Label Switching - Transport Profile (MPLS-TP) or BFD over a collection of Software-Defined Wide Area Network (SD-WAN) tunnels.

RAW formalizes a forwarding timescale that may be order(s) of magnitude shorter than the Controller Plane routing timescale 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, it can ignore the cause of transient losses. On the other hand, the RAW forwarding decision is made at the scale of a burst of packets and uses information that must be pertinent at the present time for the current transmission(s).

6.2. OODA Loop

The RAW architecture applies the generic OODA model to continuously optimize the spectrum and energy used to forward packets within a recovery graph, instantiating the OODA steps as follows:

Observe: Network Plane measurements, including protocols for OAM, observe the local state of the links and some or all hops along a recovery graph as well as the end-to-end packet delivery (see more in [Section 6.3](#)). Information can also be provided by lower-layer interfaces such as DLEP.

Orient: The orientation function reports data and information such as the link statistics and leverages offline-computed wisdom and knowledge to orient the PLR for its forwarding decision (see more in [Section 6.4](#)).

Decide: A local PLR decides which DetNet path to use for future packet(s) that are routed along the recovery graph (see more in [Section 6.5](#)).

Act: PREOF Data Plane actions are controlled by the PLR over the LL API 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 is down the recovery graph from the PLR (see more in [Section 6.6](#)).

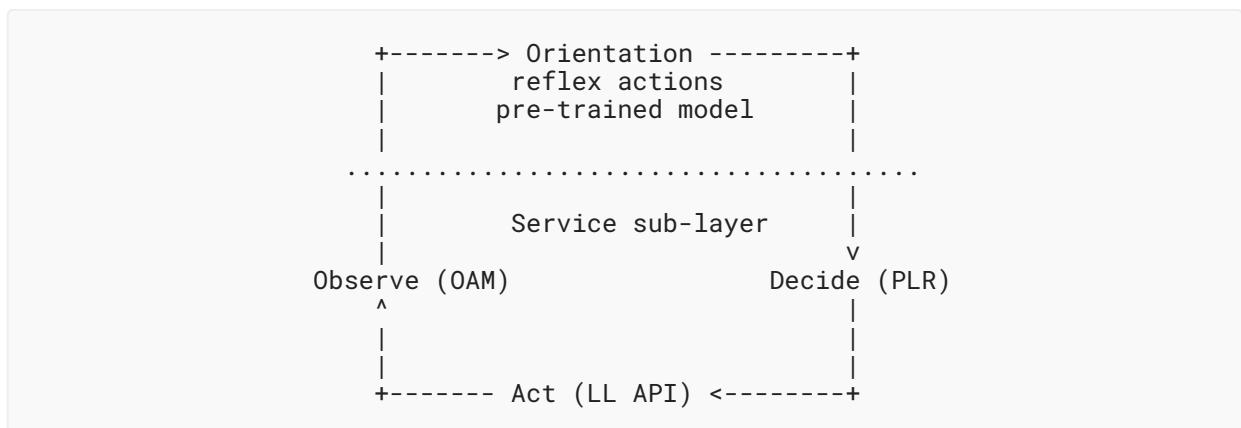


Figure 9: The RAW OODA Loop

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.

6.3. Observe: RAW OAM

The RAW in-situ OAM operation in the Network Plane may observe either a full recovery graph or the DetNet path that is being used at this time. As packets may be load balanced, replicated, eliminated, and/or fragmented for Network Coding 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 sub-layer Service Assurance may observe the individual PREOF operation of a DetNet relay node to ensure that it is conforming; this might require injecting an OAM packet at an upstream point inside the recovery graph and extracting that packet at another point downstream before it reaches the egress.

This observation feeds the RAW PLR that makes the decision on which path is used at which RAW node, for one packet or a small continuous series of packets.

In the case of end-to-end protection in a wireless mesh, the recovery graph is strict and congruent with the path so all links are observed.

Conversely, in the case of Radio Access Protection, illustrated in [Figure 10](#), the recovery graph is loose and only the first hop is observed; the rest of the path is abstracted and considered infinitely reliable. The loss of 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.

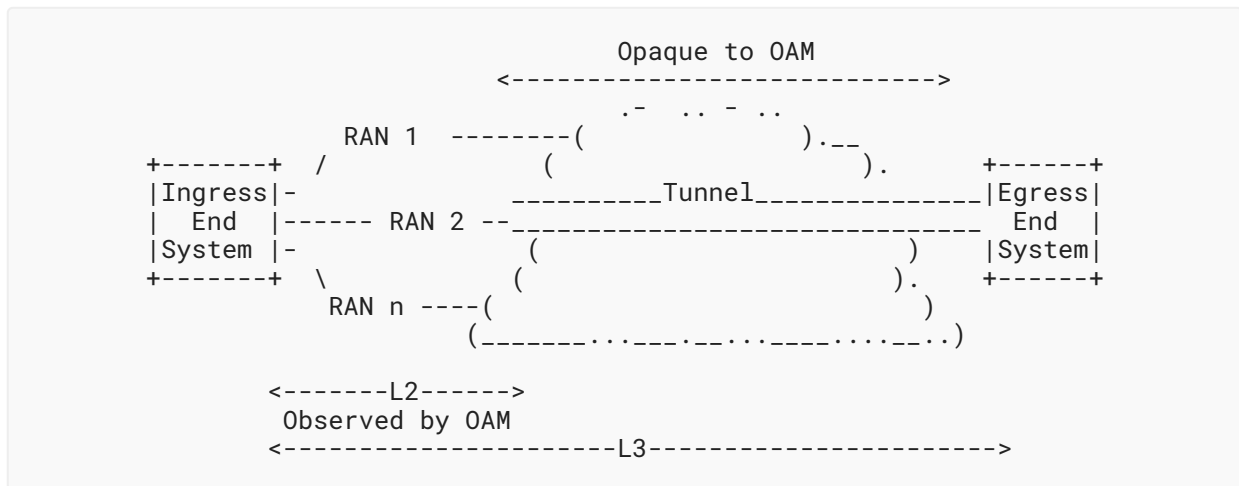


Figure 10: Observed Links in Radio Access Protection

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 captured in intermediate nodes. In the example above, the tunnel underlay is opaque and not controlled by RAW; still,

RAW OAM measures the end-to-end latency and delivery ratio for packets sent via RAN 1, RAN 2, and RAN 3, and determines whether a packet should be sent over either access link or a collection of those access links.

6.4. Orient: The RAW-Extended DetNet Operational Plane

RAW separates the long timescale at which a recovery graph is computed and installed from the short timescale at which the forwarding decision is taken for one or a few packets (see [Section 6.1](#)) that experience the same path until the network conditions evolve and another path is selected within the same recovery graph.

The recovery graph computation is out of scope, but RAW expects that the CPF that installs the recovery graph also provides related knowledge in the form of metadata about the links, segments, and possible DetNet paths. That metadata 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 metadata may include:

- A set of pre-determined DetNet paths that are prepared to match expected link-degradation profiles, so the DetNet relay nodes can take reflex rerouting actions when facing a degradation that matches one such profile; and
- Link-quality statistics history and pre-trained models (e.g., to predict the short-term variation of quality of the links in a recovery graph).

The recovery graph is installed with measurable objectives that are computed by the CPF to achieve the RAW SLA. The objectives can be expressed as any of the maximum number of packets 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.

6.5. Decide: The Point of Local Repair

The RAW OODA loop operates at the path selection timescale to provide agility versus the brute-force approach of flooding the whole recovery graph. Within the redundant solutions that are proposed by the routing function, the OODA loop controls what is used for each packet and provides a reliable and available service while minimizing the waste of constrained resources.

To that effect, RAW defines the Point of Local Repair (PLR), which performs rapid local adjustments of the forwarding tables within the path diversity that is available in that in the recovery graph. The PLR enables exploitation of the richer forwarding capabilities at a faster timescale over a portion of the recovery graph, in either a loose or a strict fashion.

The PLR operates on metrics that evolve quickly and need to be advertised at a fast rate (but only locally, within the recovery graph), and the PLR reacts to the metric updates by changing the DetNet path in use for the affected flows.

The rapid changes in the forwarding decisions are made and contained within the scope of a recovery graph, and the actions of the PLR are not signaled outside the recovery graph. This is as opposed to the routing function that must observe the whole network and optimize all the recovery graphs globally, which can only be done at a slow pace and with long-term statistical metrics, as presented in [Table 1](#).

	Controller Plane	PLR
Communication	Slow, distributed	Fast, local
Timescale (order)	Path computation + round trip, milliseconds to seconds	Lookup + protection switch, micro to milliseconds
Network Size	Large, many recovery graphs to optimize globally	Small, limited set of protection paths
Considered Metrics	Averaged, statistical, shade of grey	Instantaneous values / boolean condition

Table 1: Centralized Decision Versus PLR

The PLR sits in the DetNet forwarding sub-layer of edge and relay nodes. The PLR operates on the packet flow, learning the recovery graph and path-selection information from the packet and possibly making a local decision and retagging the packet to indicate so. On the other hand, the PLR interacts with the lower layers (through triggers and DLEP) and with its peers (through OAM) to obtain up-to-date information about its links and the quality of the overall recovery graph, respectively, as illustrated in [Figure 11](#).

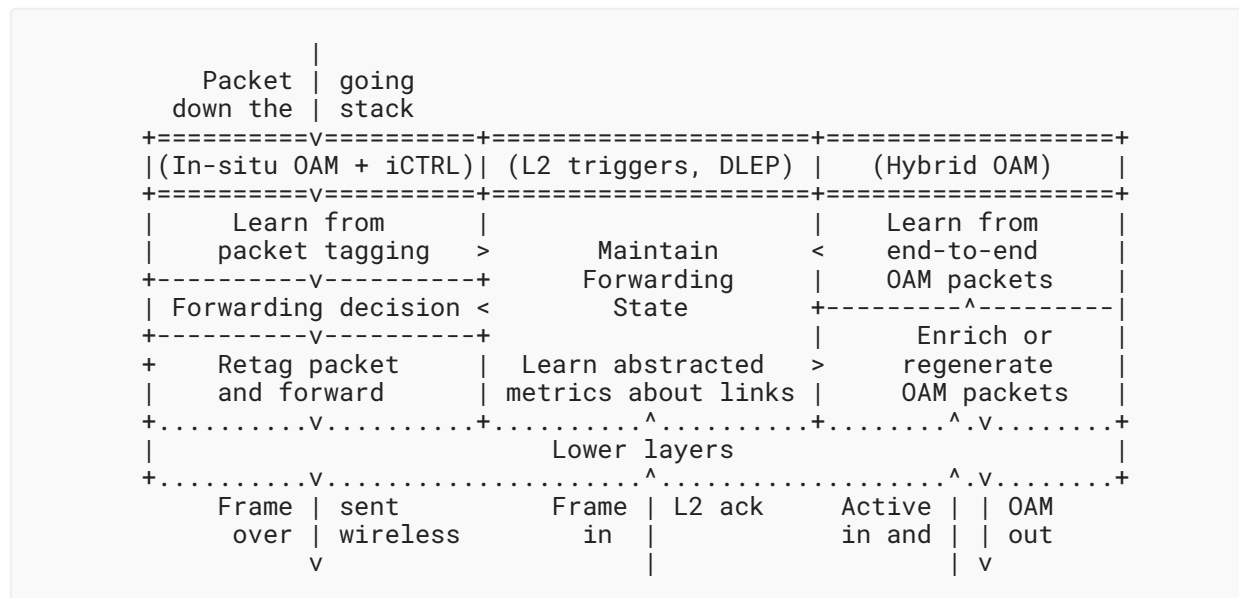


Figure 11: PLR Conceptual Interfaces

6.6. Act: DetNet Path Selection and Reliability Functions

The main action by the PLR is the swapping of the DetNet path within the recovery graph for the future packets. The candidate DetNet paths represent different energy and spectrum profiles and provide protection against different failures.

The LL API enriches the DetNet protection services (PREOF) with the possibility to interact with lower-layer, one-hop reliability functions that are more typical with wireless links than with wired ones, including ARQ, FEC, and other techniques such as overhearing and constructive interferences. Because RAW may be leveraged on wired links (e.g., to save power), it is not expected that all lower layers support all those capabilities.

RAW provides hints to the lower-layer services on the desired outcome, and the lower layer acts on those hints to provide the best approximation of that outcome, e.g., a level of reliability for one-hop transmission within a bounded budget of time and/or energy. Thus, the LL API makes possible cross-layer optimization for reliability depending on the actual abstraction provided. That is, some reliability functions are controlled from Layer 3 using an abstract interface, while they are really operated at the lower layers.

The RAW path selection can be implemented in both centralized and distributed approaches. In the centralized approach, the PLR may obtain a set of pre-computed DetNet paths matching a set of expected failures and apply the appropriate DetNet paths for the current state of the wireless links. In the distributed approach, the signaling in the packet may be more abstract than an explicit path, and the PLR decision might be revised along the selected DetNet path based on a better knowledge of the rest of the way.

The dynamic DetNet path selection in RAW avoids the waste of critical resources such as spectrum and energy while providing for the assured SLA, e.g., by rerouting and/or adding redundancy only when a loss spike is observed.

7. Security Considerations

7.1. Collocated Denial-of-Service Attacks

RAW leverages diversity (e.g., spatial and time diversity, coding diversity, and frequency diversity), possibly using heterogeneous wired and wireless networking technologies over different physical paths, to increase 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 defeats possible attacks that would impact a particular technology or a specific path.

Physical actions by a collocated attacker such as a radio interference may still lower the reliability of an end-to-end RAW transmission by blocking one segment or one possible path. However, if an alternate path with diverse frequency, location, and/or technology is available, then RAW adapts by rerouting the impacted traffic over the preferred alternates, which defeats

the attack after a limited period of lower reliability. Then again, the security benefit is a side effect of an action that is taken regardless of whether or not the source of the issue is voluntary (an attack).

7.2. Layer 2 Encryption

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

7.3. Forced Access

A RAW policy typically selects the cheapest collection of links that matches the requested SLA, e.g., use free Wi-Fi versus 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.

Similar attacks may also be used to deplete resources in lower-power nodes by forcing additional transmissions for FEC and ARQ, and attack metrics such as battery life of the nodes. By affecting the transmissions and the associated routing metrics in one area, an attacker may force the traffic and cause congestion along a remote path, thus reducing the overall throughput of the network.

8. IANA Considerations

This document has no IANA actions.

9. References

9.1. Normative References

- [DetNet-ARCH]** 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>>.
- [DetNet-OAM]** Mirsky, G., Theoleyre, F., Papadopoulos, G., Bernardos, CJ., Varga, B., and J. Farkas, "Framework of Operations, Administration, and Maintenance (OAM) for Deterministic Networking (DetNet)", RFC 9551, DOI 10.17487/RFC9551, March 2024, <<https://www.rfc-editor.org/info/rfc9551>>.
- [RAW-TECHNOS]** Thubert, P., Ed., Cavalcanti, D., Vilajosana, X., Schmitt, C., and J. Farkas, "Reliable and Available Wireless (RAW) Technologies", RFC 9913, DOI 10.17487/RFC9913, April 2026, <<https://www.rfc-editor.org/info/rfc9913>>.

- [RFC4427]** Mannie, E., Ed. and D. Papadimitriou, Ed., "Recovery (Protection and Restoration) Terminology for Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4427, DOI 10.17487/RFC4427, March 2006, <<https://www.rfc-editor.org/info/rfc4427>>.
- [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>>.
- [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>>.
- [TSN]** IEEE, "Time-Sensitive Networking (TSN)", <<https://1.ieee802.org/tsn/>>.

9.2. Informative References

- [6TiSCH-ARCH]** 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>>.
- [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>>.
- [DetNet-PLANE]** Malis, A., Geng, X., Ed., Chen, M., Varga, B., and C.J. Bernardos, "A Framework for the Deterministic Networking (DetNet) Controller Plane", RFC 9938, DOI 10.17487/RFC9938, March 2026, <<https://www.rfc-editor.org/info/rfc9938>>.
- [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>>.
- [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>>.
- [INT-ARCH]** 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>>.
- [NASA1]** Adams, T., "RELIABILITY: Definition & Quantitative Illustration", <<https://extapps.ksc.nasa.gov/Reliability/Documents/150814-3bWhatIsReliability.pdf>>.
- [NASA2]** Adams, T., "Availability", <https://extapps.ksc.nasa.gov/Reliability/Documents/160727.1_Availability_What_is_it.pdf>.

-
- [RAW-USE-CASES]** Bernardos, C.J., Ed., Papadopoulos, G., Thubert, P., and F. Theoleyre, "Reliable and Available Wireless (RAW) Use Cases", RFC 9450, DOI 10.17487/RFC9450, August 2023, <<https://www.rfc-editor.org/info/rfc9450>>.
- [RFC0791]** Postel, J., "Internet Protocol", STD 5, RFC 791, DOI 10.17487/RFC0791, September 1981, <<https://www.rfc-editor.org/info/rfc791>>.
- [RFC2205]** Braden, R., Ed., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, DOI 10.17487/RFC2205, September 1997, <<https://www.rfc-editor.org/info/rfc2205>>.
- [RFC3366]** Fairhurst, G. and L. Wood, "Advice to link designers on link Automatic Repeat reQuest (ARQ)", BCP 62, RFC 3366, DOI 10.17487/RFC3366, August 2002, <<https://www.rfc-editor.org/info/rfc3366>>.
- [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>>.
- [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>>.
- [RFC5880]** Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", RFC 5880, DOI 10.17487/RFC5880, June 2010, <<https://www.rfc-editor.org/info/rfc5880>>.
- [RFC6378]** Weingarten, Y., Ed., Bryant, S., Osborne, E., Sprecher, N., and A. Fulignoli, Ed., "MPLS Transport Profile (MPLS-TP) Linear Protection", RFC 6378, DOI 10.17487/RFC6378, October 2011, <<https://www.rfc-editor.org/info/rfc6378>>.
- [RFC6551]** Vasseur, JP., Ed., Kim, M., Ed., Pister, K., Dejean, N., and D. Barthel, "Routing Metrics Used for Path Calculation in Low-Power and Lossy Networks", RFC 6551, DOI 10.17487/RFC6551, March 2012, <<https://www.rfc-editor.org/info/rfc6551>>.
- [RFC8578]** Grossman, E., Ed., "Deterministic Networking Use Cases", RFC 8578, DOI 10.17487/RFC8578, May 2019, <<https://www.rfc-editor.org/info/rfc8578>>.
- [RFC8724]** Minaburo, A., Toutain, L., Gomez, C., Barthel, D., and JC. Zuniga, "SCHC: Generic Framework for Static Context Header Compression and Fragmentation", RFC 8724, DOI 10.17487/RFC8724, April 2020, <<https://www.rfc-editor.org/info/rfc8724>>.
- [RFC8762]** Mirsky, G., Jun, G., Nydell, H., and R. Foote, "Simple Two-Way Active Measurement Protocol", RFC 8762, DOI 10.17487/RFC8762, March 2020, <<https://www.rfc-editor.org/info/rfc8762>>.
- [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>>.
- [RFC9378] Brockners, F., Ed., Bhandari, S., Ed., Bernier, D., and T. Mizrahi, Ed., "In Situ Operations, Administration, and Maintenance (IOAM) Deployment", RFC 9378, DOI 10.17487/RFC9378, April 2023, <<https://www.rfc-editor.org/info/rfc9378>>.
- [RFC9473] Enghardt, R. and C. Krähenbühl, "A Vocabulary of Path Properties", RFC 9473, DOI 10.17487/RFC9473, September 2023, <<https://www.rfc-editor.org/info/rfc9473>>.
- [RFC9544] Mirsky, G., Halpern, J., Min, X., Clemm, A., Strassner, J., and J. François, "Precision Availability Metrics (PAMs) for Services Governed by Service Level Objectives (SLOs)", RFC 9544, DOI 10.17487/RFC9544, March 2024, <<https://www.rfc-editor.org/info/rfc9544>>.
- [RFC9633] Geng, X., Ryoo, Y., Fedyk, D., Rahman, R., and Z. Li, "Deterministic Networking (DetNet) YANG Data Model", RFC 9633, DOI 10.17487/RFC9633, October 2024, <<https://www.rfc-editor.org/info/rfc9633>>.
- [RLFA-FRR] Bryant, S., Filstis, 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>>.
- [TE] Farrel, A., Ed., "Overview and Principles of Internet Traffic Engineering", RFC 9522, DOI 10.17487/RFC9522, January 2024, <<https://www.rfc-editor.org/info/rfc9522>>.

Acknowledgments

This architecture could never have been completed without the support and recommendations from the DetNet chairs Janos Farkas and Lou Berger, and from Dave Black, the DetNet Tech Advisor. Many thanks to all of you.

The authors wish to thank Ketan Talaulikar, as well as Balazs Varga, Dave Cavalcanti, Don Fedyk, Nicolas Montavont, and Fabrice Theoleyre for their in-depth reviews during the development of this document.

The authors wish to thank Acee Lindem, Eva Schooler, Rich Salz, Wesley Eddy, Behcet Sarikaya, Brian Haberman, Gorry Fairhurst, Éric Vyncke, Erik Kline, Roman Danyliw, and Dave Thaler for their reviews and comments during the IETF Last Call and IESG review cycle.

Special thanks for Mohamed Boucadair, Giuseppe Fioccola, and Benoit Claise for their help dealing with OAM technologies.

Contributors

The editor wishes to thank the following individuals for their contributions to the text and the ideas discussed in this document:

Lou Berger

LabN Consulting, L.L.C

Email: lberger@labn.net

Xavi Vilajosana

Wireless Networks Research Lab, Universitat Oberta de Catalunya

Email: xvilajosana@gmail.com

Georgios Papadopolous

IMT Atlantique

Email: georgios.papadopoulos@imt-atlantique.fr

Remous-Aris Koutsiamanis

IMT Atlantique

Email: remous-aris.koutsiamanis@imt-atlantique.fr

Rex Buddenberg

Retired

Email: buddenbergr@gmail.com

Greg Mirsky

Ericsson

Email: gregimirsky@gmail.com

Author's Address

Pascal Thubert (EDITOR)

Independent

06330 Roquefort-les-Pins

France

Email: pascal.thubert@gmail.com