

Workgroup: RAW

Published: 1 June 2021

Intended Status: Informational

Expires: 3 December 2021

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

 Cisco Systems IMT Atlantique

 L. Berger R. Buddenberg

 LabN Consulting, L.L.C.

Reliable and Available Wireless Architecture/Framework

Abstract

Reliable and Available Wireless (RAW) provides for high reliability and availability for IP connectivity over a wireless medium. The wireless medium presents significant challenges to achieve deterministic properties such as low packet error rate, bounded consecutive losses, and bounded latency. This document defines the RAW Architecture. It builds on the DetNet Architecture and discusses specific challenges and technology considerations needed to deliver DetNet service utilizing scheduled wireless segments and other media, e.g., frequency/time-sharing physical media resources with stochastic traffic.

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 3 December 2021.

Copyright Notice

Copyright (c) 2021 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents

carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

- [1. Introduction](#)
- [2. The RAW problem](#)
 - [2.1. Terminology](#)
 - [2.2. Reliability and Availability](#)
 - [2.2.1. High Availability Engineering Principles](#)
 - [2.2.2. Applying Reliability Concepts to Networking](#)
 - [2.2.3. Reliability in the Context of RAW](#)
 - [2.3. Use Cases and Requirements Served](#)
 - [2.3.1. Radio Access Protection](#)
 - [2.3.2. End-to-End Protection in a Wireless Mesh](#)
 - [2.4. Related Work at The IETF](#)
- [3. The RAW Framework](#)
 - [3.1. Scope and Prerequisites](#)
 - [3.2. Routing Time Scale vs. Forwarding Time Scale](#)
 - [3.3. Wireless Tracks](#)
 - [3.4. PAREO Functions](#)
 - [3.4.1. Packet Replication](#)
 - [3.4.2. Packet Elimination](#)
 - [3.4.3. Promiscuous Overhearing](#)
 - [3.4.4. Constructive Interference](#)
- [4. The RAW Architecture](#)
 - [4.1. The RAW Conceptual Model](#)
 - [4.2. The Path Selection Engine](#)
 - [4.3. RAW OAM](#)
 - [4.4. Flow Identification vs. Path Identification](#)
 - [4.5. Source-Routed vs. Distributed Forwarding Decision](#)
 - [4.6. Encapsulation and Decapsulation](#)
- [5. Security Considerations](#)
 - [5.1. 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 budgetted 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 IP [[IPoWIRELESS](#)]. Wireless networks operate on a shared medium where uncontrolled interference, including the self-induced multipath fading, cause random transmission losses and add new dimensions to the statistical effects that affect reachability and packet delivery.

To defeat those additional causes of transmission delay and loss, Reliable and Available Wireless (RAW) leverages scheduled transmissions with redundancy and diversity in the spatial, time, code, and frequency domains. 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.

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.

Uncontrolled interference and transmission obstacles may impede the wireless transmission, causing rapid variations of the throughput and packet delivery ratio (PDR) of the link. This uncertainty limits the volume and/or duration of traffic that can be safely transmitted on the same link while conforming to a RAW Service Level Agreement (SLA).

This increased complexity explains why the development of deterministic wireless technologies has been lagging behind the similar efforts for wired systems, both at the IEEE and the IETF.

But recent progress on scheduled radios such as TSCH and OFDMA indicates that wireless is finally catching up at the lower layers. Sitting at the layer above, RAW takes up the challenge of providing highly available and reliable end-to-end performances in a network with scheduled wireless segments.

RAW provides DetNet elements that are specialized for short range radios. From this inheritance, RAW stays 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.

The "[Deterministic Networking Architecture](#)" [RFC8655] is composed of three planes: the Application (User) Plane, the Controller Plane, and the Network Plane. The RAW Architecture extends the DetNet Network Plane, to accomodate one or multiple hops of homogeneous or heterogeneous wireless technologies, e.g. a Wi-Fi6 Mesh or parallel CBRS access links federated by a 5G backhaul.

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 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 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 called a Track. 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 covers Network Plane protocol elements such as Operations, Administration and Maintenance (OAM) to observe some or all hops along a Track as well as the end-to-end packet delivery, and in-band control to optimize the use of redundancy to achieve the required SLA with minimal use of constrained resources.

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 as a complex path with associated PAREO operations. The concept is abstract to the underlaying technology and applies to any fully or partially wireless mesh, including, e.g., a Wi-Fi 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 term OAM as defined in [RFC6291].

RAW defines the following terms:

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, constructive interference and overhearing, which can be leveraged separately or combined to increase the reliability.

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.

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

Reliability: Reliability is a measure of the probability that an item will perform its intended function for a specified interval under stated conditions. For RAW, the service that is expected is delivery within a bounded latency and a failure is when the packet is either lost or delivered too late. RAW expresses reliability in terms of Mean Time Between Failure (MTBF) and Maximum Consecutive Failures (MCF). More in [NASA].

Availability: Availability is a measure of the relative amount of time where a path operates in stated condition, in other words (uptime)/(uptime+downtime). Because a serial wireless path may not be good enough to provide the required availability, and even 2 parallel paths may not be over a longer period of time, the RAW availability implies a path that is a lot more complex than what DetNet typically envisages (a Track).

2.2. Reliability and Availability

2.2.1. High Availability Engineering Principles

The reliability criteria of a critical system pervade 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 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 parcimony.

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. This is why 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. But RAW itself operates in the Network Plane at a faster time scale. To act on the Data Plane, RAW needs live information from the Operational Plane , e.g., using [Bidirectional Forwarding Detection](#) [BFD] and its variants (bidirectional and remote BFD) to protect a link, and OAM techniques to protect a path.

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. Reliability in the Context of RAW

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 the something moves that 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 maximised 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.

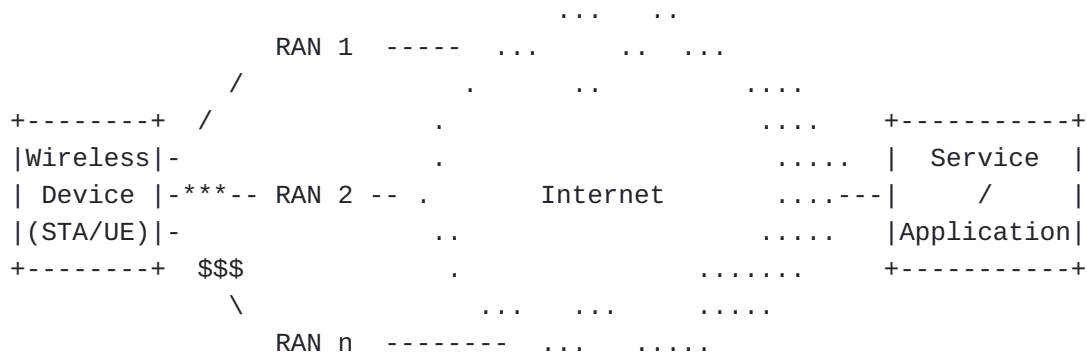
2.3. Use Cases and Requirements Served

In order to focus on real-worlds issues and assert the feasibility of the proposed capabilities, RAW focuses on selected technologies that can be scheduled at the lower layers: IEEE Std. 802.15.4 timeslotted channel hopping (TSCH), 3GPP 5G ultra-reliable low latency communications (URLLC), IEEE 802.11ax/be where 802.11be is extreme high throughput (EHT), and L-band Digital Aeronautical Communications System (LDACS). See [[RAW-TECHNOS](#)] for more.

["Deterministic Networking Use Cases"](#) [[RFC8578](#)] presents a number of wireless use cases including Wireless, such as application to Industrial Applications, Pro-Audio, and SmartGrid Automation. [[RAW-USE-CASES](#)] adds a number of use cases that demonstrate the need for RAW capabilities for new applications such as Pro-Gaming and drones. The use cases can be abstracted in two families, Loose Protection, e.g., protecting the first hop in Radio Access Protection and Strict Protection, e.g., providing End-to-End Protection in a wireless mesh.

2.3.1. Radio Access Protection

To maintain the required SLA at all times, a wireless Host may use more than one Radio Access Network (RAN) in parallel.



*** = flapping at this time \$\$\$ expensive

Figure 1: Radio Access Protection

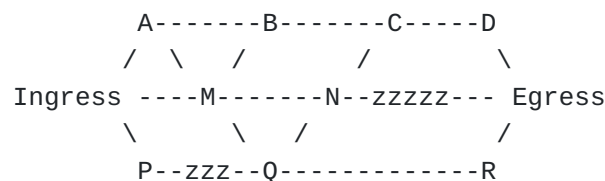
The RANs may be heterogeneous, e.g., 3GPP 5G [RAW-5G] and Wi-Fi [RAW-TECHNOS] for high-speed communication, in which case a Layer-3 abstraction becomes useful to select which of the RANs are used at a particular point of time, and the amount of traffic that is distributed over each RAN.

The idea is that the rest of the path to the destination(s) is protected separately (e.g., uses non-congruent paths, leverages DetNet / TSN, etc...) and is a lot more reliable, e.g., wired. In that case, RAW observes the reliability of the end-to-end operation through each of the RANs but only observes and controls the wireless operation the first hop.

A variation of that use case has a pair of wireless Hosts connected over a wired core / backbone network. In that case, RAW observes and controls the Ingress and Egress RANs, while neglecting the hops in the core. The resulting loose Track may be instantiated, e.g., using tunneling or loose source routing between the RANs.

2.3.2. End-to-End Protection in a Wireless Mesh

In radio technologies that support mesh networking (e.g., Wi-Fi and TSCH), a Track is a complex path with distributed PAREO capabilities. In that case, RAW operates through the multipath and makes decisions either at the Ingress or at every hop (more in [Section 3.3](#)).



zzz = flapping now

Figure 2: End-to-End Protection

The Protection may be imposed by the source based on end-to-end OAM, or performed hop-by-hop, in which case the OAM must enable the intermediate Nodes to estimate the quality of the rest of the feasible paths in the remainder of the Track to the destination.

2.4. Related Work at The IETF

RAW intersects with protocols or practices in development at the IETF as follows:

- *The Dynamic Link Exchange Protocol (DLEP) [[RFC8175](#)] from [[MANET](#)] can be leveraged at each hop to derive generic radio metrics (e.g., based on LQI, RSSI, queueing delays and ETX) on individual hops.
- *OAM work at [[detnet](#)] such as [[DetNet-IP-OAM](#)] for the case of the IP Data Plane observes the state of DetNet paths, typically MPLS and IPv6 pseudowires [[DetNet-DP](#)], in the direction of the traffic. RAW needs feedback that flows on the reverse path and gathers instantaneous values from the radio receivers at each hop to inform back the source and replicating relays so they can make optimized forwarding decisions. The work named ICAN may be related as well.
- *[[BFD](#)] detect faults in the path between an Ingress and an Egress forwarding engines, but is unaware of the complexity of a path with replication, and expects bidirectionality. BFD considers delivery as success whereas with RAW the bounded latency can be as important as the delivery itself.
- *[[SPRING](#)] and [[BIER](#)] define in-band signaling that influences the routing when decided at the head-end on the path. There's already one RAW-related draft at BIER [[BIER-PREF](#)] more may follow. RAW will need new in-band signaling when the decision is distributed, e.g., required chances of reliable delivery to destination within latency. This signaling enables relays to tune retries and replication to meet the required SLA.
- *[[CCAMP](#)] defines protocol-independent metrics and parameters (measurement attributes) for describing links and paths that are required for routing and signaling in technology-specific networks. RAW would be a source of requirements for CCAMP to define metrics that are significant to the focus radios.

3. The RAW Framework

3.1. Scope and Prerequisites

A prerequisite to the RAW operation is that an end-to-end routing function computes a complex sub-topology along which forwarding can happen between a source and one or more destinations. The concept of Track is specified in the [6TiSCH Architecture](#) [[6TiSCH-ARCHI](#)] to represent that complex sub-topology. Tracks provide a high degree of redundancy and diversity and enable the DetNet PREOF, network coding, and possibly RAW specific techniques such as PAREO, leveraging frequency diversity, time diversity, and possibly other forms of diversity as well.

How the routing operation (e.g., PCE) in the Controller Plane computes the Track is out of scope for RAW. The scope of the RAW operation is one Track, and the goal of the RAW operation is to optimize the use of the Track at the forwarding timescale to maintain the expected SLA while optimizing the usage of constrained resources such as energy and spectrum.

Another prerequisite is that an IP link can be established over the radio with some guarantees in terms of service reliability, e.g., it can be relied upon to transmit a packet within a bounded latency and provides a guaranteed BER/PDR outside rare but existing transient outage windows that can last from split seconds to minutes. The radio layer can be programmed with abstract parameters, and can return an abstract view of the state of the Link to help the Network Layer forwarding decision (think DLEP from MANET).

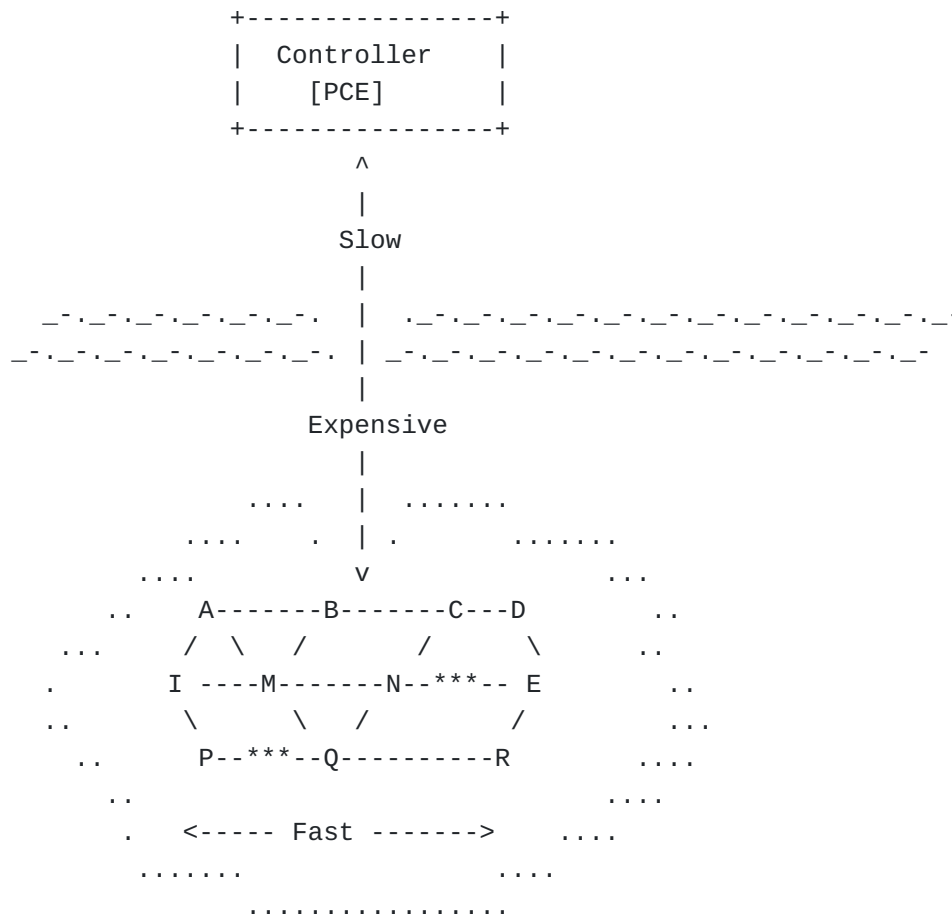
How the radio interface manages its lower layers is out of control and out of scope for RAW. In the same fashion, the non-RAW portion along a loose Track is by definition out of control and out of scope for RAW. Whether it is a single hop or a mesh is also unknown and out of scope.

3.2. Routing Time Scale vs. Forwarding Time Scale

With DetNet, the Controller Plane Function that handles the routing computation and maintenance (the PCE) can be centralized and can reside outside the network. 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 3: 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).

3.3. Wireless Tracks

The "[6TiSCH Architecture](#)" [[6TiSCH-ARCHI](#)] introduces the concept of Track. RAW extends the concept to any wireless mesh technology, including, e.g., Wi-Fi. A simple Track is composed of a direct sequence of reserved hops to ensure the transmission of a single packet from a source Node to a destination Node across a multihop path.

A Complex Track provides multiple non-equal cost multipath (NECM) forwarding solutions. The Complex Track enables to support multipath redundant forwarding by employing PRE functions [[RFC8655](#)] and the ingress and within the Track. For example, a Complex Track may branch off and rejoin over non-congruent segments.

In the context of RAW, some links or segments in the Track may be reversible, meaning that they can be used in either direction. In that case, an indication in the packet signals the direction of the reversible links or segments that the packet traverses and thus places a constraint that prevents loops from occurring. An individual packet follows a destination-oriented directed acyclic graph (DODAG) towards a destination Node inside the Complex Track.

3.4. PAREO Functions

RAW may control whether and how to use packet replication and elimination (PRE), Automatic Repeat reQuest (ARQ), Hybrid ARQ (HARQ) that includes Forward Error Correction (FEC) and coding, 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 function 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 4](#), many different paths are possible for S to reach R. 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.

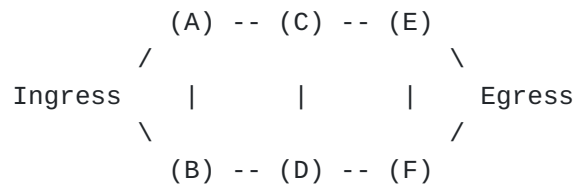


Figure 4: A Ladder Shape with Two Parallel Paths

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.

3.4.1. 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 5](#), the Ingress Node is transmitting the packet to both successors, nodes A and B, at two different times.

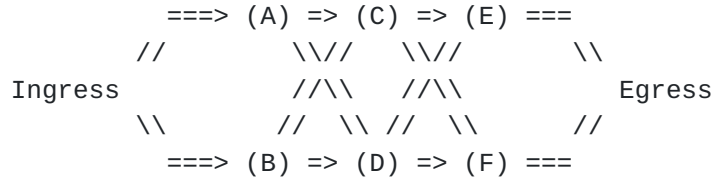


Figure 5: Packet Replication

An example schedule is shown in [Table 1](#). 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 1: Packet Replication: Sample schedule

3.4.2. 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.

3.4.3. 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 6](#), 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.

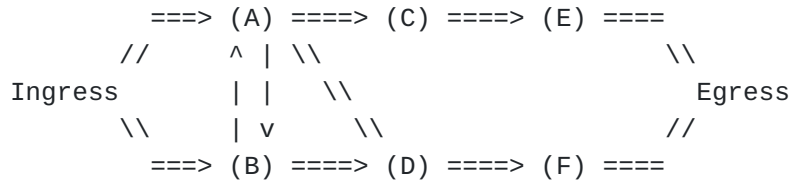


Figure 6: Unicast with Overhearing

3.4.4. 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.

4. The RAW Architecture

4.1. 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.

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 control 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.

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.

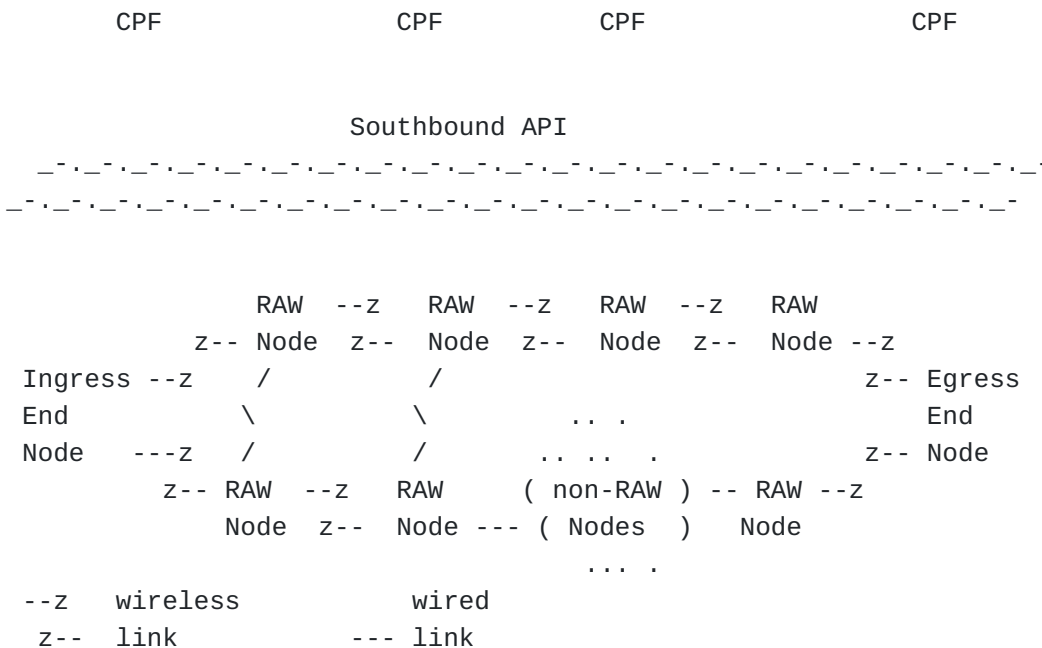


Figure 7: RAW Nodes

The Link-Layer metrics are reported to the PCE in a time-aggregated, e.g., statistical fashion. Example Link-Layer metrics include typical Link bandwidth (the medium speed depends dynamically on the PHY mode and the number of users sharing the spectrum) and average and mean squared deviation of availability and reliability figures such as Packet Delivery Ratio (PDR) over long periods of time.

Based on those metrics, the PCE installs the Track with enough redundant forwarding solutions to ensure that the Network Plane can reliably deliver the packets within a System Level Agreement (SLA) associated to the flows that it transports. The SLA defines end-to-end reliability and availability requirements, where 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 Track may comprise non-RAW segments, either interleaved inside the Track, or all the way to

the Egress End Node (e.g., a server in the Internet). 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 PAREO diversity schemes is actioned by which RAW Nodes.

Once a Track is established, per-segment and end-to-end reliability and availability statistics are periodically reported to the PCE to assure that the SLA can be met or have it recompute the Track if not.

4.2. The Path Selection Engine

RAW separates the path computation time scale at which a complex path is recomputed from the path selection time scale at which the forwarding decision is taken for one or a few packets (more in [Section 3.2](#)). RAW operates at the path selection time scale. The RAW problem is to decide, 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 counter-part 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 needs 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 needs to 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 2](#).

	PCE (Not in Scope)	PSE (In Scope)
Operation	Centralized	Source-Routed or Distributed
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 2: 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 8](#).

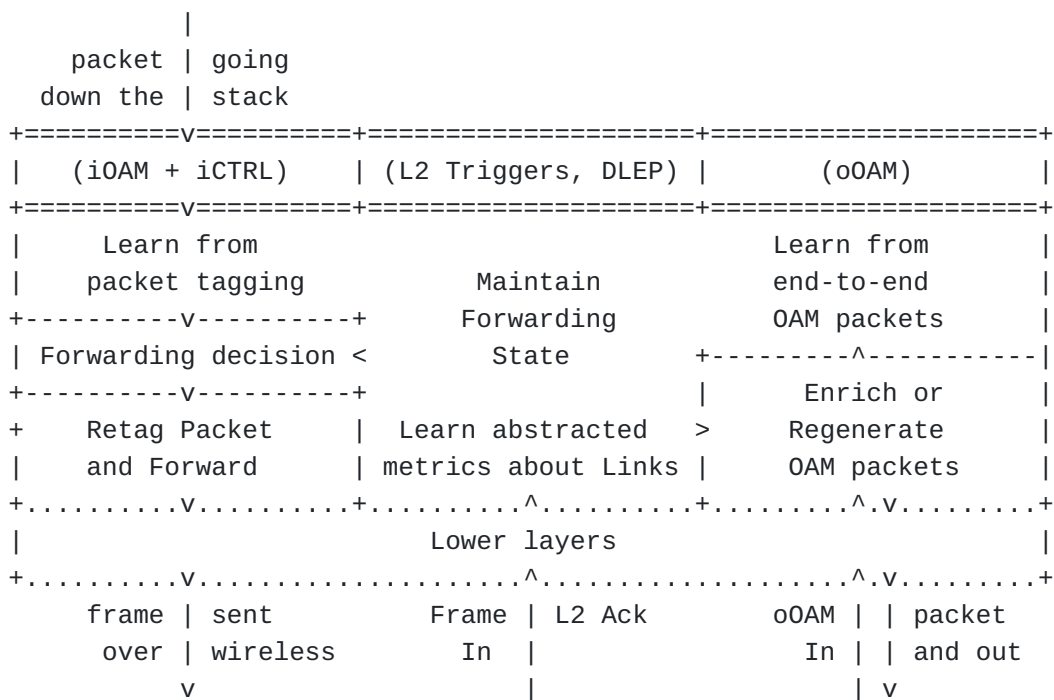


Figure 8: PSE

4.3. RAW OAM

The RAW OAM operation in the Network Plane observes either a full Track or subTracks that are being used at this time. 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.

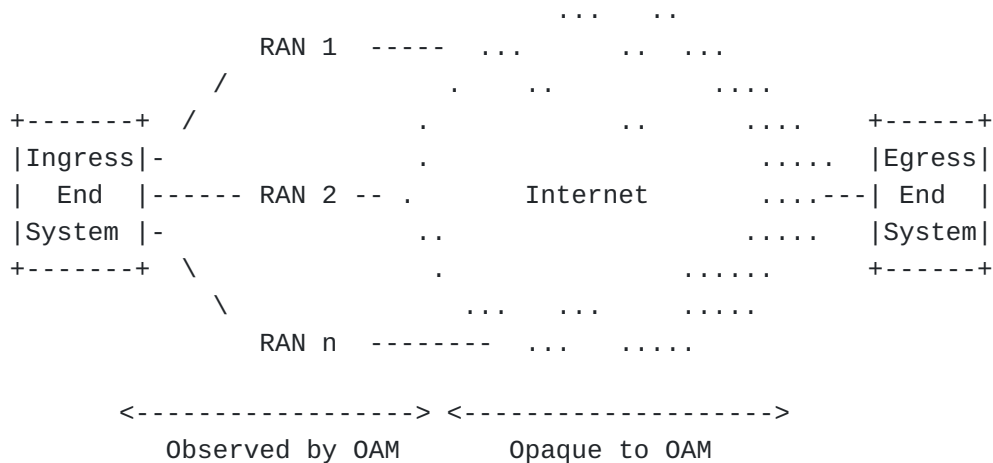


Figure 9: Observed Links in Radio Access Protection

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, the Track is Loose and in that case only the first hop is observed; the rest of the path is abstracted and considered infinitely reliable.

In the case of the Radio Access Protection, only the first hop is protected; the loss of a packet that was sent over one of the possible first hops is attributed to that first hop, even if a particular loss effectively happens farther down the path.

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 of RAN 1, RAN 2 and RAN 3, and determines whether a packet should be sent over either or a collection of those access links.

4.4. Flow Identification vs. Path Identification

Section 4.7 of the DetNet Architecture [RFC8655] ties the app-flow identification which is an application layer concept with the network path identification that depends on the networking technology by "exporting of flow identification", e.g., to a MPLS label.

With RAW, this exporting operation is injective but not bijective. e.g., a flow is fully placed within one RAW Track, but not all packets along that Track are necessarily part of the same flow. For instance, out-of-band OAM packets must circulate in the exact same fashion as the flows that they observe. It results that the flow

identification that maps to to app-flow at the network layer must be separate from the path identification that is used to forward a packet.

Section 3.4 of the DetNet data-plane framework [[DetNet-DP](#)] indicates that for a DetNet IP Data Plane, a flow is identified by an IPv6 6-tuple. With RAW, that 6-tuple is not what indicates the Track, in other words, the flow ID is not the Track ID.

For instance, the 6TiSCH Architecture [[6TiSCH-ARCHI](#)] uses a combination of the address of the Egress End System and an instance identifier in a Hop-by-hop option to indicate a Track. This way, if a packet "escapes" the Track, it will reach the Track Egress point through normal routing and be treated at the service layer through, say, elimination and reordering.

The RAW service includes forwarding over a subset of the Links that form the Track (a subTrack). Packets from the same or a different flow that are routed through the same Track will not necessarily traverse the same Links. The PSE selects a subTrack for a packet based on the links that are preferred and those that should be avoided at this time.

Each packet is forwarded within the subTrack that provides the best adequation with the SLA of the flow and the energy and bandwidth constraints of the network.

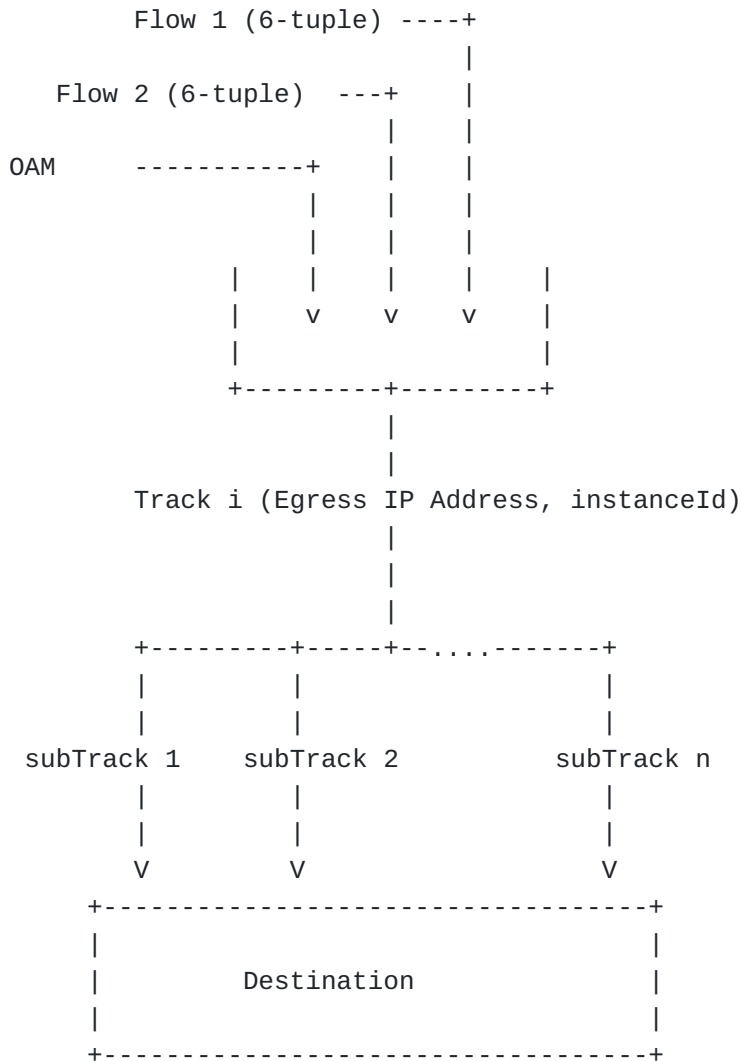


Figure 10: Flow Injection

With 6TiSCH, packets are tagged with the same (destination address, instance ID) will experience the same RAW service regardless of the IPv6 6-tuple that indicates the flow. The forwarding does not depend on whether the packets transport application flows or OAM. In the generic case, the Track or the subTrack can be signaled in the packet through other means, e.g., encoded in the suffix of the destination address as a Segment Routing Service Instruction [[SR-ARCHI](#)], or leveraging Bit Index Explicit Replication [[BIER](#)] Traffic Engineering [[BIER-TE](#)].

4.5. Source-Routed vs. Distributed Forwarding Decision

Within a large routed topology, the route-over mesh operation builds a particular complex Track with one source and one or more destinations; within the Track, packets may follow different paths and may be subject to RAW forwarding operations that include replication, elimination, retries, overhearing and reordering.

The RAW forwarding decisions include the selection of points of replication and elimination, how many retries can take place, and a limit of validity for the packet beyond which the packet should be destroyed rather than forwarded uselessly further down the Track.

The decision to apply the RAW techniques must be done quickly, and depends on a very recent and precise knowledge of the forwarding conditions within the complex Track. There is a need for an observation method to provide the RAW Data Plane with the specific knowledge of the state of the Track for the type of flow of interest (e.g., for a QoS level of interest). To observe the whole Track in quasi real time, RAW considers existing tools such as L2-triggers, DLEP, BFD and leverages in-band and out-of-band OAM to capture and report that information to the PSE.

One possible way of making the RAW forwarding decisions within a Track is to position a unique PSE at the Ingress and express its decision in-band in the packet, which requires the explicit signaling of the subTrack within the Track. In that case, the RAW forwarding operation along the Track is encoded by the source, e.g., by indicating the subTrack in the Segment Routing (SRv6) Service Instruction, or by leveraging BIER-TE such as done with [\[BIER-PREF\]](#).

The alternate way is to operate the PSE in each forwarding Node, which makes the RAW forwarding decisions for a packet on its own, based on its knowledge of the expectation (timeliness and reliability) for that packet and a recent observation of the rest of the way across the possible paths based on OAM. Information about the desired service should be placed in the packet and matched with the forwarding Node's capabilities and policies.

In either case, a per-track/subTrack state is installed in all the intermediate Nodes to recognize the packets that are following a Track and determine the forwarding operation to be applied.

4.6. Encapsulation and Decapsulation

In the generic case where the Track Ingress Node is not the source of the Packet, the Ingress Node needs to encapsulate IP-in-IP to ensure that the Destination IP Address is that of the Egress Node and that the necessary Headers (Routing Header, Segment Routing Header and/or Hop-By-Hop Header) can be added to the packet to

signal the Track or the subTrack, conforming [[IPv6](#)] that discourages the insertion of a Header on the fly.

In the specific case where the Ingress Node is the source of the packet, the encapsulation can be avoided, provided that the source adds the necessary headers and that the destination is set to the Egress Node. Forwarding to a final destination beyond the Egress Node is possible, e.g., with a Segment Routing Header that signals the rest of the way. In that case a Hop-by-Hop Header is not recommended since its validity is within the Track only.

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. 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

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

Remous-Aris Koutsiamanis: IMT Atlantique

Nicolas Montavont: IMT Atlantique

8. Acknowledgments

TBD

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>>.

[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-01, 19 February 2021, <<https://tools.ietf.org/html/draft-ietf-raw-technologies-01>>.

[RAW-USE-CASES] Papadopoulos, G. Z., Thubert, P., Theoleyre, F., and C. J. Bernardos, "RAW use cases", Work in Progress, Internet-Draft, draft-ietf-raw-use-cases-01, 21 February 2021, <<https://tools.ietf.org/html/draft-ietf-raw-use-cases-01>>.

[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>>.

[BFD] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", RFC 5880, DOI 10.17487/RFC5880, June 2010, <<https://www.rfc-editor.org/info/rfc5880>>.

[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>>.

[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>>.

[SR-ARCHI] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.

[BIER] Wijnands, IJ., Ed., Rosen, E., Ed., Dolganow, A., Przygienda, T., and S. Aldrin, "Multicast Using Bit Index Explicit Replication (BIER)", RFC 8279, DOI 10.17487/RFC8279, November 2017, <<https://www.rfc-editor.org/info/rfc8279>>.

[RFC8175]

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>>.

[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>>.

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>>.

[BIER-PREF]

Thubert, P., Eckert, T., Brodard, Z., and H. Jiang, "BIER-TE extensions for Packet Replication and Elimination Function (PREF) and OAM", Work in Progress, Internet-Draft, draft-thubert-bier-replication-elimination-03, 3 March 2018, <<https://tools.ietf.org/html/draft-thubert-bier-replication-elimination-03>>.

[DetNet-IP-OAM] Mirsky, G., Chen, M., and D. Black, "Operations, Administration and Maintenance (OAM) for Deterministic Networks (DetNet) with IP Data Plane", Work in Progress, Internet-Draft, draft-ietf-detnet-ip-oam-02, 30 March 2021, <<https://tools.ietf.org/html/draft-ietf-detnet-ip-oam-02>>.

[RAW-5G] Farkas, J., Dudda, T., Shapin, A., and S. Sandberg, "5G - Ultra-Reliable Wireless Technology with Low Latency", Work in Progress, Internet-Draft, draft-farkas-raw-5g-00, 1 April 2020, <<https://tools.ietf.org/html/draft-farkas-raw-5g-00>>.

[BIER-TE] Eckert, T., Cauchie, G., and M. Menth, "Tree Engineering for Bit Index Explicit Replication (BIER-TE)", Work in Progress, Internet-Draft, draft-ietf-bier-te-arch-09, 30 October 2020, <<https://tools.ietf.org/html/draft-ietf-bier-te-arch-09>>.

[IPoWIRELESS] Thubert, P., "IPv6 Neighbor Discovery on Wireless Networks", Work in Progress, Internet-Draft, draft-thubert-6man-ipv6-over-wireless-09, 17 May 2021,

<<https://tools.ietf.org/html/draft-thubert-6man-ipv6-over-wireless-09>>.

- [NASA] Adams, T., "RELIABILITY: Definition & Quantitative Illustration", <<https://kscddms.ksc.nasa.gov/Reliability/Documents/150814-3bWhatIsReliability.pdf>>.
- [MANET] IETF, "Mobile Ad hoc Networking", <<https://dataTracker.ietf.org/doc/charter-ietf-manet/>>.
- [detnet] IETF, "Deterministic Networking", <<https://dataTracker.ietf.org/doc/charter-ietf-detnet/>>.
- [SPRING] IETF, "Source Packet Routing in Networking", <<https://dataTracker.ietf.org/doc/charter-ietf-spring/>>.
- [BIER] IETF, "Bit Indexed Explicit Replication", <<https://dataTracker.ietf.org/doc/charter-ietf-bier/>>.
- [BFD] IETF, "Bidirectional Forwarding Detection", <<https://dataTracker.ietf.org/doc/charter-ietf-bfd/>>.
- [CCAMP] IETF, "Common Control and Measurement Plane", <<https://dataTracker.ietf.org/doc/charter-ietf-ccamp/>>.

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

Lou Berger
LabN Consulting, L.L.C.
United States of America

Email: lberger@labn.net

Rex Buddenberg

CA

United States of America

Email: buddenbergr@gmail.com