RAW WG
Internet-Draft
Intended status: Standards Track
Expires: 14 September 2023

CJ. Bernardos
UC3M
A. Mourad
InterDigital
13 March 2023

RAW multidomain extensions
draft-bernardos-raw-multidomain-02

Abstract

This document describes the multi-domain RAW problem and explores and proposes some extensions to enable RAW multi-domain operation.

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 14 September 2023.

Copyright Notice

Table of Contents

1. Introduction and Problem Statement

Wireless operates on a shared medium, and transmissions cannot be fully deterministic due to uncontrolled interferences, including self-induced multipath fading. RAW (Reliable and Available Wireless) is an effort to provide Deterministic Networking on across a path that include a wireless interface. 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. RAW extends the DetNet Working Group concepts to provide for high reliability and availability for an IP network utilizing scheduled wireless segments and other media, e.g., frequency/time-sharing physical media resources with stochastic traffic: IEEE Std. 802.15.4 timeslotted channel hopping (TSCH), 3GPP 5G ultra-reliable low latency communications (URLLC), IEEE 802.11ax/be, and L-band Digital Aeronautical Communications System (LDACS), etc. Similar to DetNet, RAW technologies aim at staying abstract to the radio layers underneath, addressing the Layer 3 aspects in support of applications requiring high reliability and availability.

As introduced in [I-D.ietf-raw-architecture], 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. RAW operates at the path selection time scale. The RAW problem is to decide, amongst the redundant solutions that are proposed by the Patch Computation Element (PCE), which one 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 Packet (hybrid) ARQ, Replication, Elimination and Ordering (PAREO), and scheduled transmissions at a faster time scale.

There are several use cases [I-D.ietf-raw-use-cases] where reliability and availability are key requirements for wireless heterogeneous networks. A couple of relevant examples are (i) the manufacturing sector, where a plethora of devices are interconnected and generate data that need to be reliably delivered to the control and monitoring agents; and (ii) the residential gaming, with eXtended Reality (XR).

We can refer to domains managed by a single PCE, as single-domain RAW, where nodes are typically run and managed by a single administration entity. In this scenario, the PSE can make use of tracks and paths involving only the nodes belonging to this RAW domain.

There are scenarios where hosts are connected to different RAW domains and they need to communicate to each other with certain reliability and/or availability guarantees, for example in large factories where networks might be organized in domains (per production lines or building/sites), in residential environments where there are different networks (e.g., one at home and one in the garden), or even vehicular scenarios (e.g., hosts connected to different vehicles).

```
   ___________________________________________________
  |                                                   |
  |                        ( ( o ) )                  |
  |                        *   ^                      |
  |                       *   / \                     |
  |                      *   /   \                    |
  |                    **   ------+--                 |
  |                   *     | RAW |P|                 |
  |  ( ( o ) )*      *      |node |S|                 |
  |     ^     *( ( o ) )    | 1-1 |E|       +------+
  |    / \          ^    *  ------+--       | PCE1 |
  |   /   \        / \    **                 +------+
  |  +-----+      /   \     *( ( o ) )          |
  |  |host1|     ------+--       ^               |
  |  |     |     | RAW |P|      / \              |
  |  |     |     |node |S|     /   \             |
  |  |  o  |     | 1-2 |E|    ------+--          |
  |  +-----+     ------+--    | RAW |P|          |
  |               \   /       |node |S|          |
  |      RAW       \ /        | 1-3 |E|          |
  |    domain 1     v         ------+--          |
  |              ( ( o ) )                       |
  |                 ****                         |
  |_______________ *____****     _______________|
   _____________ *___________***** _____________
  |             *               **               |
  |            *          ****( ( o ) )          |
  |           *       ****        ^               |
  |      ( ( o ) )****           / \              |
  |          ^        *        /   \             |
  |         / \        *      ------+--          |
  |        /   \        *     | RAW |P|          |
  |       ------+--      *    |node |S|          |
  |       | RAW |P|       *   | 2-1 |E|     +------+
  |       |node |S|        *  ------+--    | PCE2 |
  |       | 2-2 |E|         *              +------+
  |       ------+--        *( ( o ) )          |
  |                       *     ^               |
  |                ( ( o ) )    / \              |
  |                    ^       /   \             |
  |                   / \     ------+--          |
  |                  /   \    | RAW |P|          |
  |                +-----+    |node |S|          |
  |       RAW      |host2|    | 2-3 |E|          |
  |     domain 2   +-----+    ------+--          |
  |___________________________________________________|
```

Figure 1: Exemplary scenario showing multiple RAW domains

Figure 1 shows an example of communication involving two RAW domains. As opposed to a single-domain scenario, where a single PCE may compute all possible tracks at longer time scale, and the PSE functionality may perform subtrack selection and optimization at a shorter time scale using all information available at the domain, multidomain scenarios pose additional burdens that are not solved yet.

Each RAW domain operates independently of the other domains. While there exist inter-PCE solutions today, allowing one domains PCE to learn some inter-domain paths, this would not be sufficient, as the PSE of one domain would not have full visibility nor capability to act on the other domains (e.g., there are no multi-domain OAM solutions in place yet), limiting its capability to guarantee any given SLA. Therefore, there is a need to define inter-PSE coordination mechanisms across domains.

There exist today standardized solutions, such as the ones in the context of Path Computation Element (PCE), enabling computing multi-/inter-domain paths. As an example, the Hierarchical PCE (G-PCE) was defined in RFC 6805 [RFC6805] and is described hereafter. A parent PCE maintains a domain topology map that contains the child domains (seen as vertices in the topology) and their interconnections (links in the topology). The parent PCE has no information about the content of the child domains; that is, the parent PCE does not know about the resource availability within the child domains, nor does it know about the availability of connectivity across each domain because such knowledge would violate the confidentiality requirement and either would require flooding of full information to the parent (scaling issue) or would necessitate some form of aggregation. The parent PCE is used to compute a multi-domain path based on the domain connectivity information. A child PCE may be responsible for single or multiple domains and is used to compute the intra-domain path based on its own domain topology information.

Solutions like the above are not sufficient alone to solve the multi-domain RAW problem, as the PSEs need to have some additional information from the other involved domains to be sensitive/reactive to transient changes, in order to ensure a certain level of reliability and availability in a multi-domain wireless heterogeneous mesh network.

Within a single domain, the RAW framework architecture works, by having the PCE in charge of computing the paths (tracks) and the PSE(s) taking the short time decisions of which sub-tracks to use. Note that the PSE is assumed to be either a distributed functionality (performed by every RAW router of the path, which takes forwarding decisions based on the local and OAM information that they have), or

a centralized functionality played by the entry (ingress) router in the domain (note that if there are multiple ingress nodes, then there might be multiple PSEs), which then performs source routing.

In scenarios with multiple connected RAW domains, running uncoordinated RAW solutions in each domain is not sufficient. PSEs would need to have global end-to-end information as well as be capable of running OAM mechanisms [I-D.ietf-raw-oam-support] to monitor the quality of the selected paths.

2. Terminology

The following terms used in this document are defined by the IETF:

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.

PSE. The Path Selection Engine (PSE) 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.

3. RAW multi-domain extensions

Here we specify the new mechanisms and signaling extensions to enable inter-domain RAW connectivity.

```
   +-----+-+      +----+       +----+      +-----+-+     +-----+-+
   | RAW |P|      |    |       |    |      | RAW |P|     | RAW |P|
   |node |S|      |PCE1|       |PCE2|      |node |S|     |node |S|
   | 1-2 |E|      |    |       |    |      | 2-1 |E|     | 2-2 |E|
   +-----+-+      +----+       +----+      +-----+-+     +-----+-+
         |           |            |              |             |
   1.Path compute req|            |              |             |
   (src=node1-2,     |            |              |             |
    dst=node2-3, SLA)|            |              |             |
         |••••••••••>|            |              |             |
         |           |2.Path compute req         |             |
         |           |(src={node2-1,node2-2},    |             |
         |           | dst=node2-3)              |             |
         |           |•••••••••••>|              |             |
         |           |3.Path compute resp        |             |
         |           |({tracks2},{links_quality})|             |
         |           |<•••••••••••|              |             |
   4.Path compute resp            |              |             |
   ({{tracks1},{tracks2}},        |              |             |
    PSE={node2-1,node2-2},        |              |             |
    {SLA1,SLA2})      |            |              |             |
         |<••••••••••|            |              |             |
         |5.RAW inter-domain path |              |             |
         |({{tracks1,tracks2}},{SLA1,SLA2})      |             |
         |••••••••••••••••••••••••••••••••••••••>|             |
         |••••••••••••••••••••••••••••••••••••••••••••••••••••••>
         |           |      6.RAW inter-domain path ACK         |
         |<••••••••••••••••••••••••••••••••••••••|             |
         |<••••••••••••••••••••••••••••••••••••••••••••••••••••
         |           |            |              |             |
         |7.RAW OAM(flow/track,SLA1)             |             |
    <•••>|<•••>      |            |  7.RAW OAM(flow/track,SLA1)
         |           |            |            <••>|<••>     <••>|<••>
         |          8.RAW OAM (flow/track, metrics)            |
         |<••••••••••••••••••••••••••••••••••••••>|             |
         |<••••••••••••••••••••••••••••••••••••••••••••••••••••>
         |           |            |              |             |
```

Figure 2: Multi-domain RAW signaling

Figure 2 shows a signaling flow diagram, taking as baseline scenario the one shown in Figure 1, where host1 (connected to node1-2) wants to communicate with host2 (connected to node2-3). An ingress RAW node (node1-2) gets a request for connectivity, with a given destination RAW node (node2-3) and the desired SLA in terms of reliability and availability. The source and/or destination RAW nodes might be hostss. We next explain each of the steps illustrated in the figure:

1. The ingress node plays the role of PSE (also referred to as PSE@domain1) and requests the computation of the tracks towards the destination node2-3 with the intended SLA to the PCE of the domain (PCE1).

2. PCE1 knows that the destination is in another domain (domain2) and that the PCE of the destination domain is PCE2. PCE1 also knows the ingress nodes in domain2 that are connected to domain1. How this is done is outside of the scope of this document. These nodes (node2-1 and node2-2) play the role of PSEs@domain2. PCE1 requests to PCE2 to compute the available tracks from PSEs@domain2 to the destination, and the characteristics of the links (link_quality) forming these tracks. The detail and nature of the information provided by PCE2 regarding the links might vary depending on the deployment, and is meant to be used by PCE1 and the PSE@domain1 (node1-2) to compute how to distribute the SLA among the domains.

3. PCE2 computes the tracks and responds to PCE1, including also the characteristics of the links (link_quality). Examples of potential information elements including in the link_quality are: available bandwidth, observed reliability, delay, link variability/mobility, etc.

4. PCE1 provides to the PSE@domain1 the tracks to reach the destination, as well as the split of SLAs among domain1 and domain2 (SLA1 and SLA2). An SLA, or a Quality of Service (QoS) figure, may include aspects such as, among others: max. delay, assured BW, max. Jitter, packet loss ratio, availability ratio, etc. PCE1 also provides the PSEs@domain2.

5. The PSE@domain1 sends a message to each PSE@domain2, in order to set-up a direct communication channel to provide OAM information useful to the PSE@domain1 for computing the subtracks to use for the traffic. This message includes the SLA that each domain has to monitor and guarantee (SLA1 and SLA2).

6. Each of the PSEs@domain2 acknowledges the message. At this point, the communication channel is established and the PSE@domain1 can start taking decisions at a forwarding time scale regarding which paths (subtracks) to use.

7. All PSEs, at each domain, start performing OAM procedures [I-D.ietf-raw-oam-support], which are key to observe if traffic is meeting the desired SLAs (SLA1 and SLA2) and adapt the subtracks and tracks if needed. OAM mechanisms can be applied in-band (sharing the traffic's fate) or out-of band. Note that this per-domain distributed OAM is critical to ensure that the required SLAs (reliability and availability) are met by reacting on a short time scale at each of the involved domains.

8. PSEs share aggregated and pre-processed information among them to facilitate early detection of issues and computation of subtracks. If a violation of an SLA is detected, the respective PSE would notify the domain PCE and the other PSE, so a reaction measure can be taken (e.g., selecting different subtracks, taking different PAREO decisions, requesting the PCEs to recompute the paths and/or adjust the split of the SLAs across the domains).

Note that this example covers the direction host1-to-host2. If there is traffic in the opposite direction, the process has to be repeated in the reverse direction, as paths might not be bidirectional.

4. IANA Considerations

TBD.

5. Security Considerations

TBD.

6. Acknowledgments

The work of Carlos J. Bernardos in this document has been partially supported by the Horizon Europe PREDICT-6G (Grant 101095890) and UNICO I+D 6G-DATADRIVEN-04 project.

7. Informative References

[I-D.ietf-raw-architecture]
Thubert, P., "Reliable and Available Wireless Architecture", Work in Progress, Internet-Draft, draft-ietf-raw-architecture-11, 7 December 2022, <https://datatracker.ietf.org/doc/html/draft-ietf-raw-architecture-11>.

[I-D.ietf-raw-oam-support]
Theoleyre, F., Papadopoulos, G. Z., Mirsky, G., and C. J. Bernardos, "Operations, Administration and Maintenance (OAM) features for RAW", Work in Progress, Internet-Draft,

draft-ietf-raw-oam-support-06, 5 March 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-raw-oam-support-06>.

[I-D.ietf-raw-use-cases] Bernardos, C. J., Papadopoulos, G. Z., Thubert, P., and F. Theoleyre, "RAW Use-Cases", Work in Progress, Internet-Draft, draft-ietf-raw-use-cases-08, 22 October 2022, <https://datatracker.ietf.org/doc/html/draft-ietf-raw-use-cases-08>.

[RFC6805] King, D., Ed. and A. Farrel, Ed., "The Application of the Path Computation Element Architecture to the Determination of a Sequence of Domains in MPLS and GMPLS", RFC 6805, DOI 10.17487/RFC6805, November 2012, <https://www.rfc-editor.org/info/rfc6805>.

Authors' Addresses

Carlos J. Bernardos
Universidad Carlos III de Madrid
Av. Universidad, 30
28911 Leganes, Madrid
Spain
Phone: +34 91624 6236
Email: cjbc@it.uc3m.es
URI: http://www.it.uc3m.es/cjbc/

Alain Mourad
InterDigital Europe
Email: Alain.Mourad@InterDigital.com
URI: http://www.InterDigital.com/

RAW
Internet-Draft
Intended status: Informational
Expires: 5 September 2024

P. Thubert, Ed.
4 March 2024

Reliable and Available Wireless Architecture
draft-ietf-raw-architecture-17

Abstract

Reliable and Available Wireless (RAW) provides for high reliability and availability for IP connectivity across any combination of wired and wireless network segments. The RAW Architecture extends the DetNet Architecture and other standard IETF concepts and mechanisms to adapt to the specific challenges of the wireless medium, in particular intermittently lossy connectivity. This document defines a network control loop that optimizes the use of constrained spectrum and energy while maintaining the expected connectivity properties, typically reliability and latency. The loop involves DetNet Operational Plane functions, with a new recovery Function and a new Point of Local Repair operation, that dynamically selects the DetNet path(s) for the future packets to route around local degradations and failures.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

Deterministic Networking aims to provide bounded latency and eliminate 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, with both cost savings and complexity benefits (e.g., vs. loads of point-to-point (P2P) cables).

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

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

Wireless networks operate on a shared medium where uncontrolled interference, including the self-induced multipath fading cause random transmission losses. Fixed and mobile obstacles and reflectors may block or alter the signal, causing transient and unpredictable variations of the throughput and packet delivery ratio (PDR) of a wireless link. This adds new dimensions to the statistical effects that affect the quality and reliability of the link.

Reliable and Available Wireless (RAW) takes up the challenge of providing highly available and reliable end-to-end performances in a network with scheduled wireless segments. To achieve this, RAW can leverage multiple links and parallel transmissions, providing enough diversity and redundancy to ensure the timely packet delivery while preserving energy and optimizing the use of the shared spectrum.

Distance Vector (DV) protocols can enable more than one feasible successors along non-equal-cost multipath forwarding graphs. This provides redundancy and allows to dynamically adapt the forwarding operation to the state of the links. But this protection is limited since only a subset of the nodes along the path will have a feasible alternate successor.

RAW solves that problem by defining Protection Paths that can be fully non-congruent and can be activated dynamically upon failures. This requires additional control to take the routing decision early enough along the possible paths to route around the failure. RAW defines a end-to-end control loop that dynamically controls the activation and deactivation of the feasible Protection Paths.

In addition, RAW introduces the RAW API, which is an interface between the lower layer wireless technology and the DetNet layers. The RAW API is RAW technology [RAW-TECHNOS] dependent as it can vary what the different RAW technologies expose towards the DetNet layers. Furthermore, the different RAW technologies are equipped with different reliability features, e.g., short range broadcast, MUMIMO, PHY rate and other Modulation Coding Scheme (MCS) adaptation, (H)ARQ, constructive interference and overhearing. The RAW API enables interactions between the reliability functions provided by the wireless technology and the reliability functions provided by DetNet. That is, the RAW API makes cross-layer optimization possible for the reliability functions of different layers depending on the actual exposure provided via the RAW API by the given RAW technology.

This document presents the RAW problem and associated terminology in Section 3.2, presents a conceptual model for RAW in Section 4, and, based on that model, elaborates on an in-network optimization control loop in Section 5.2.

2. Terminology

RAW reuses terminology defined for DetNet in the "Deterministic Networking Architecture" [RFC8655], e.g., PREOF for Packet Replication, Elimination and Ordering Functions. RAW inherits and augments the IETF art of Protection as seen in DetNet and Traffic Engineering.

RAW also reuses terminology defined for MPLS in [RFC4427] such as the term recovery as covering both Protection and Restoration, a number of recovery types. That document defines a number of concepts like recovery domain that are used in the RAW works, and creates the new term recovery graph. A recovery graph associates a topological graph with usage metadata that represent how the paths within the recovery graph are built.

RAW also reuses terminology defined for RSVP-TE in [RFC4090] such as the Point of Local Repair (PLR). The concept of 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-ARCHI] and equates the 6TiSCH concept of a Track with that of a recovery graph.

In an quantic analogy, a recovery graph is to a path what an atomic orbital is to a planetary orbit, in that the electron has a probability of presence within a known shape as opposed to a deterministic trajectory.

The concept of recovery graph is agnostic to the underlaying technology and applies but is not limited to any fully or partially 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 uses the following terminology and acronyms:

2.1. Acronyms

2.1.1. ARQ

Automatic Repeat Request, a well-known mechanism, enabling an acknowledged transmission with retries to mitigate errors and loss. ARQ may be implemented at various layers in a network. ARQ is typically implemented at Layer-2, per hop and not end-to-end 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].

2.1.2. FEC

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

2.1.3. HARQ

Hybrid Automatic Repeat Request, combining FEC and ARQ.

2.1.4. MCS

Modulation and Coding Scheme. Controls the throughput of the Link to maintain reliable transmissions.

2.1.5. OAM

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

2.1.6. OODA

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

2.2. Link and Direction

2.2.1. Flapping

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

2.2.2. Uplink

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

2.2.3. Downlink

The reverse direction from uplink.

2.2.4. Downstream

Following the direction of the flow data path along a recovery graph.

2.2.5. Upstream

Against the direction of the flow data path along a recovery graph.

2.3. Path and Recovery Graphs

2.3.1. Path

Quoting section 1.1.3 of [INT-ARCHI]:

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

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

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

It follows that the general acceptance of a path is a linear sequence of links and nodes, as opposed to a multi-dimensional graph, defined by the experience of the packet that went from a node A to a node B. In the context of this document, a path is observed by following one copy or one fragment of a packet that conserves its uniqueness and integrity. For instance, if C replicates to E and F and D eliminates on the way from A to B, a packet from A to B can experience 2 paths, A->C->E->D->B and A->C->F->D->B. The terms lane is used to clarify when dealing with such path.

With DetNet and RAW, a packet may be duplicated, fragmented and network-coded, and the various byproducts may travel different paths that are not necessarily end-to-end between A and B; we refer to that complex experience 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, that experience is subject to change from a packet to the next, but all the possible experiences are all contained within a finite set. Therefore, we introduce below the term of a recovery graph 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.

2.3.2. Recovery Graph

A networking graph that can be followed to transport packets with equivalent treatment, associated with usage metadata; as opposed to the definition of a path above, a recovery graph represents not an actual but a potential, it 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, fork and rejoin, for instance to enable a protection service by the PREOF operations.

```
                  +---------+
                  | IoT G/W |
                  +---------+
                      EGR  <=== Elimination at Egress
                      | |
               /------/   \-------\    Wired backbone
               |                 |
            +--|--+            +--|--+
            |  |  | Backbone   |  |  | Backbone
            |  |  | Router     |  |  | Router
            +--|--+            +--|--+
               |                 |
            o   \     o          /  lane
          o      o      o---o---o   o      o   o  o
                  \  o /    o          o         o
           o   o   \  /       o        low power lossy network
                    \/ o           o        o
                  o  IN <=== Replication at recovery graph Ingress
                      |
                      o <- source device
```

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 the collection of all the feasible DetNet Paths that a packet which flow is assigned to the recovery graph may be forwarded along. A packet that is assigned to the recovery graph will experience 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, but 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 dependant decisions, with no node in the network capable of reporting which is the current DetNet Path within the recovery graph.

In DetNet [RFC8655] terms, a recovery graph has the following properties:

* A recovery graph is a Layer-3 abstraction built upon P2P IP links between routers. A router may form multiple P2P IP links over a single radio interface.

* A recovery graph has one Ingress and one Egress nodes, which operate as DetNet Edge nodes.

* The graph of 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 that graph are DetNet Relay nodes that operate at the DetNet Service sub-layer and provide the PREOF functions.

* The topological edges of the 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 Lanes and forward or crossing Segments, see the definition for those terms in the next sections. A Protection Path contains at least 2 Lanes as a main path and a backup path.

```
------------------- forward direction ---------------------->

      a ==> b ==> C -=- F ==> G ==> H     T1       I: Ingress
    /              \   /      |       \ /          E: Egress
  I                  o        n        E -=- T2    T1, T2, T3:
    \              /   \      |       / \            External
      p ==> q ==> R -=- T ==> U ==> v     T3         Targets

     Uppercase: DetNet Relay nodes
     Lowercase: DetNet Transit nodes

     I ==> a ==> b ==> C : an forward Segment to targets F and o
     C ==> o ==> T: an forward Segment to target T (and/or U)
     G | n | U : a crossing Segment to targets G or U
     I --> F --> E : an forward Lane 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 : lane-crossing alternate path
```

Figure 2: A Recovery Graph and its Components

2.3.3. Forward and Crossing

Forward refers to progress towards the recovery graph Egress. 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 crossing, in which case it is composed of crossing links only. A lane is always forward, meaning that is is composed of forward links and Segments.

2.3.4. Lane

An end-to-end forward lane between the Ingress and Egress Nodes of a recovery graph. A lane 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 recovery graph Segments.

2.3.5. Segment

A strict sequence of DetNet Transit nodes between 2 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.

2.4. Deterministic Networking

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

2.4.1. Flow

A collection of consecutive IP packets defined by the upper layers and signaled by the same 5 or 6-tuple, see section 5.1 of [RFC8939]. Packets of the same flow must be placed on the same 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.

2.4.2. Deterministic Flow Identifier (L2)

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

2.4.3. Deterministic Flow Identifier (L3)

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

2.4.4. TSN

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

2.5. Reliability and Availability

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

2.5.1. Service Level Agreement

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

2.5.2. Service Level Objective

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

2.5.3. Service Level Indicator

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

2.5.4. Reliability

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

2.5.5. Available

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

2.5.6. Availability

Availability is a measure of the relative amount of time where a RAW Network operates in stated condition (SLA), expressed as (uptime)/(uptime+downtime).

2.6. OAM variations

2.6.1. Active OAM

See [RFC7799]. In the context of RAW, Active OAM is used to observe a particular recovery graph, DetNet Path, or Segment of a recovery graph regardless of whether it is used for traffic at that time.

2.6.2. In-Band OAM

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

2.6.3. Out-of-Band OAM

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

2.6.4. Limited OAM

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

2.6.5. Upstream OAM

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

2.6.6. Residence Time

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

2.6.7. Lower Layer information

The RAW Operational Plane elements (PLR and OAM Supervisor) may gather aggregated information from lower layers about e.g., link quality. This information may be obtained from inside the device using specialized API (e.g., L2 triggers) or via control protocols such as BFD [RFC5880] or DLEP [DLEP]. It may then be massaged and exported through oOAM messaging, and passed to the Controller Plane using the aCPF.

2.6.8. Additional References

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

3. Reliable and Available Wireless

3.1. Reliability and Availability

3.1.1. High Availability Engineering Principles

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

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

1. elimination of 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. Examples of both non-Internet and Internet are included.

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

Intermediate network Nodes such as routers, switches and APs, wire bundles and the air medium itself can become single points of failure. For High Availability, it is thus required to use physically link- and Node-disjoint paths; in the wireless space, it is also required to use the highest possible degree of diversity (time, space, code, frequency, channel width) in the transmissions over the air to combat the additional causes of transmission loss.

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

3.1.1.2. Reliable Crossover

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

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

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

3.1.1.3. 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 an abstract 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 recovery graph based on statistical and aggregated information. RAW itself operates in the DetNet Network Plane at a faster time scale with live information on speed, state, etc... This live information can be obtained directly from the lower layer, e.g., using L2 triggers, read from a protocol such as the Dynamic Link Exchange Protocol (DLEP) [DLEP], or transported over multiple hops using OAM and reverse OAM, as illustrated in Figure 11.

3.1.2. Applying Reliability Concepts to Networking

The terms Reliability and Availability are defined for use in RAW in Section 2 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 packet, but does not affect the previous or next packet, nor 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 to some 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].

3.1.3. Wireless Effects Affecting Reliability

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

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

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

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

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

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

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

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

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

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

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

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

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

3.2. The RAW problem

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.

Operating at the Layer-3, RAW does not change the wireless technology at the lower layers. OTOH, 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 listed in [RAW-TECHNOS]. Conceptually, RAW is agnostic to the radio layer underneath though the capability to schedule transmissions is assumed. How the PHY is programmed to do so, and whether the radio is single-hop or meshed, are unknown at the IP layer and not part of the RAW abstraction. Nevertheless, cross-layer optimizations may take place to ensure proper link awareness (think, link quality) and packet handling (think, scheduling).

The "Deterministic Networking Architecture" [RFC8655] is composed of three planes: the Application (User) Plane, the Controller Plane, and the Network Plane. The DetNet Network Plane is composed a Dataplane (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 estapblishes the patsh, associates the flows to the paths, qnd ensures the availability of the necessary resources, leverages Layer-2 functionalities for timely delivery to the next DetNet system.

The RAW Architecture extends the DetNet Network Plane, to accommodate one or multiple hops of homogeneous or heterogeneous wired and wireless technologies. RAW adds reactivity to the DetNet service sub-layer to compensate the dynamics for the radio links in terms of lossiness and bandwidth. This may apply for instance to mesh networks as illustrated in Figure 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 lane 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: Because the medium is broadcast as opposed to physically point to point like a wire, 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 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 variates over time, the concept of metric must also be revisited. Shortest path loses its absolute value, and hop count turns into a bad idea as the link budget drops with the distance. Routing over radio requires both 1) a new and more dynamic sense of the link, with new protocols such as DLEP and L2-trigger to maintain L3 up to date with the link quality and availability, and 2) a new approach to multipath routing, where non-equal cost multipath becomes the norm as shortest path loses its meaning with the instability of the metrics.

ARQ, FEC and codes: Though feasible on any technology, proactive

(forward) and reactive (ARQ) error correction are typical to the wireless media. Bounded latency can still be obtained on a wireless link while operating those technologies, provided that the extra transmission happens within the budget allocated to that hop, or that the introduced delay is compensated along the path. In the case of coded fragments and retries, it makes sense to variate all the possible physical properties of the transmission to reduce the chances that the same effect causes the loss of both original and redundant transimissions.

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.

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

The establishment of a path is not in-scope for RAW. It may be, e.g., the product of a centralized Controller Plane Function (CPF) like a Path computation Element (PCE) [RFC4655], or may be computed in a distributed fashion ala Resource ReSerVation Protocol (RSVP) [RFC2205]. On the other hand, 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, and it makes sense in the wireless case to provide redundant forwarding solutions along a (see Section 2.3) and to leave it to the Network Plane to select which of those forwarding solutions are to be used for a given packet based on the current conditions.

RAW distinguishes the longer time scale at which routes are computed from the the shorter time scale where forwarding decisions are made for a limited time RAW Network Plane operations happen at a time scale that sits between the routing and the forwarding time scales, on one DetNet flow, to select a DetNet path within the resources delineated by a recovery graph (see Section 2.3.2). The recovery graph is preestablished and installed by means outside of the scope of RAW; it may be strict or loose depending on whether each or just a subset of the hops are observed and controlled by RAW.

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

1. Operational Plane measurement protocols for Operations, Administration and Maintenance (OAM) to Observe some or all hops along a recovery graph as well as the end-to-end packet delivery

2. The DetNet Controller Plane Function (CPF) is split with an optional asynchronous CPF (aCPF) that reports data and information such as link statistics to be used asynchronously by the routing CPF (rCPF) to compute, install, and maintain the recovery graphs, e.g., by generating knowledge and wisdom such as a trained model for link quality prediction, which in turn can be used by the aCPF to Orient the Path selection by the PLR within the RAW OODA loop.

3. An Operational Plane PLR that hosts the Decision function of which DetNet Paths to use for the future packets that will be routed within the recovery graph

4. Service protection actions that operate at the DetNet Service sub-layer to increase the reliability of the end-to-end transmissions. The RAW architecture also covers in-situ signaling when the decision is Acted by a node that down the recovery graph from the PLR.

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

4. The RAW Conceptual Model

RAW inherits the conceptual model described in section 4 of the DetNet Architecture [RFC8655] as illustrated in Figure 3, which also shows example reliability Functions in the different layers. RAW extends DetNet with Point of Local Repair (PLR, see Section 5.5) to provide additional agility against transmission loss. The PLR can act, e.g., based on indications from the wireless layer or based on OAM.

```
                    .
                    .
        +-----------------------------+
        |                             |
        |  DetNet Service sub-layer   | (e.g., (PREOF)
        |                             |
        +-----------------------------+
        |                             |
        | DetNet Forwarding sub-layer |
        |                             |
        |                             |
        |           _____    _____    |
        |          | PLR +--+ OAM |   |
        |           -+---    -----    |
        |            | ^              |
        |            | |              |
        |         ___v_+___          |
        +--------+ RAW API +----------+
        |         ---------           |
        |    Wireless/Radio Layer     | (e.g., (H)ARQ)
        |                             |
        +-----------------------------+
                    .
```

Figure 3: Wireless layer and DetNet sub-layers

The RAW API enables interactions between the reliability functions provided by the wireless technology and the reliability functions provided by DetNet. Thus, the RAW aPI enables cross-layer optimizations to improve reliability.

4.1. The RAW Planes

A RAW Network Plane may be strict (as illustrated in Figure 6 or loose (as illustrated in Figure 7, 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, 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.

The RAW Nodes are DetNet relays that operate in the RAW Network Plane and are capable of additional diversity mechanisms and measurement functions related to the radio interface. RAW leverages a CPF that operates inside the RAW Nodes (typically the Ingress Edge Nodes) to dynamically adapt the path of the packets and optimizes the resource usage.

An RAW-enabled rCPF 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.

```
                          DetNet Routing

        rCPF               rCPF          rCPF                 rCPF


                        Southbound API
   _-._-._-._-._-._-._-._-._-._-._-._-._-._-._-._-._-._-._-._-._-._-._-._-
 _-._-._-._-._-._-._-._-._-._-._-._-._-._-._-._-._-._-._-._-._-._-._-._-


                RAW  --/   RAW  --/   RAW  --/   RAW
            /-- Node  /--  Node  /--  Node  /--  Node --/
  Ingress --/    /          /                           /-- Egress
  End           /          /         .. .                 End
  Node   ---/   /          /       .. ..  .             /-- Node
           /-- RAW  --/   RAW     ( non-RAW ) -- RAW --/
               Node  /--  Node --- ( Nodes  )   Node
                                     ... .
  --/   wireless           wired
   /--  link           --- link
```

Figure 4: RAW Nodes

When a new flow is defined, the rCPF uses it 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 can not leverage end-to-end DetNet and cannot provide latency guarantees. The non-RAW subnetwork is neglected in the RAW computation, that

is, considered jitterless, and infinitely reliable and/or available in comparison with the links between RAW nodes, so loss and jitter that is measured end-to-end is attributed to the RAW hops (typically an access link).

A local asynchronous CPF in the RAW node reports the Link-Layer metrics to the rCPF 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), number of flows (bandwidth that can be reserved for a flow depends on the number and size of flows 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 DetNet rCPF installs the recovery graph 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 recovery graph may comprise non-RAW segments, either interleaved inside the recovery graph, 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 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 rCPF to assure that the SLA can be met or have it recompute the recovery graph if not.

4.2. RAW vs. Upper and Lower Layers

RAW improves the reliability of transmissions and the availability of the communication resources, but does not provide scheduling and shaping, so RAW itself does not provide guarantees such as latency for the application payload. Rather, it should be seen as a dynamic optimization of the use of redundancy to maintain it within certain boundaries. For instance, ARQ is operated by the lower layers and RAW will only abstract the concept and hint the lower layers on the desired outcome, as opposed to performing the retries 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. Excess of incoming traffic at the DetNet Ingress will cause either dropping, queueing, or reclassification of the packets, and entail loss, latency, or jitter, and moot 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 the lower layers to provide the timing and physical resources guarantees that are needed to match the traffic SLA. When the availability of the physical resource varies, RAW will act on the distribution of the traffic to leverage alternates within a finite set of potential resources.

4.3. RAW and DetNet

RAW leverages the DetNet Forwarding sub-layer and requires the support of in-situ OAM in DetNet Transit Nodes (see fig 3 of [RFC8655] for the dynamic acquisition of link capacity and state to maintain a strict RAW service, end-to-end, over a DetNet Network. RAW extends DetNet to improve the protection against link errors such as transient flapping that are far more common in wireless links. Nevertheless, the RAW methods are for the most part 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. The RAW Operational sub-layer typically runs only in the DetNet Ingress Edge Node or End System, though it may also run in DetNet Relay Nodes when the RAW Control sub-layer is distributed along the recovery graph. The RAW Operational sub-layer functionality includes the PLR that decides the DetNet Path for the future packets of a flows along the DetNet Path through specific signaling, and the OAM Supervisor that triggers, and learns from, OAM observations, and feeds the PLR for its next decision.

RAW extends the DetNet Stack (see fig 4 of [RFC8655]) with additional functionality at the DetNet Service sub-layer for the actuation of the PLR decision. Layer-3 in general and DetNet in particular operates on abstractions of the lower layers and through APIs to control those abstractions. For instance, DetNet already leverages lower layers for time-sensitive operations such as time synchronization and traffic shapers. Because the performances of the radio layers are subject to rapid changes, so RAW needs more dynamic gauges and knobs. To that effect, the RAW API enables interactions between the reliability functions provided by the wireless technology and the reliability functions provided by DetNet. That is, the RAW

API provides a radio abstraction to the DetNet layer. The RAW API can be used to push reliability and timing hints like suggest X retries (min, max) within a time window, or send unicast (one next hop) or multicast (for overhearing). The other way around RAW needs hints about the radio conditions like L2 triggers (RSSI, LQI, ETX) over all the wireless hops. This information is useful to both the aCPF and the PLR.

The RAW Service sub-layer also adds the OAM Propagator that (re)generates the OAM information as it is formed and propagated In-Band or Out-of-Band. The RAW Service sub-layer may be present in DetNet Edge and Relay Nodes, though the PREOF Actuator has no operation in the Egress Edge Node.

```
 +-----------------------------+ +--------------------------------+
 |                             | |                                |
.............................................................................
 |                             | |                                |
 | +-------------------------+ | | +----------------------------+ |
 | |     aCPF                | | | |         aCPF               | |
 | +-------------------------+ | | +----------------------------+ |
 | +----------+  +------------+ | | .-.-.-.-.-.-.--.  .-.-.-.-.-.-.--. |
 | | PLR      |  | OAM        | | | | Distr. PLR |  | Distr. OAM | |
 | |          |  | Supervisor | | | |            |  | Supervisor | |
 | +----------+  +------------+ | | .-.-.-.-.-.-.--.  .-.-.-.-.-.-.--. |
 |                             | |    optional         optional   |
    RAW Operational sub-layer
.............................................................................
    DetNet Service sub-layer
 |                             | |                                |
 | +----------+  +------------+ | | +------------+  +------------+ |
 | | PREOF    |  |  OAM       | | | |  PREOF     |  |  OAM       | |
 | | Actuator |  |  Observer  | | | |  Actuator  |  |  Observer  | |
 | +----------+  +------------+ | | +------------+  +------------+ |
 |                             | |                                |
    DetNet Service sub-layer
.............................................................................
    DetNet Forwarding sub-layer
 |                             | |                                |
 |               +------------+ | |                 +------------+ |
 |               |In-Situ OAM | | |                 |In-Situ OAM | |
 |               +------------+ | |                 +------------+ |
 |                             | |                                |
 +-----------------------------+ +--------------------------------+

         End System or                      Relay
       Ingress Edge Node                    Node
```

Figure 5: RAW functional posture within DetNet sub-layers

There are 2 main proposed models to deploy RAW and DetNet. In the first model (strict) 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.

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

```
--------------------Flow Direction---------------------------------->

+---------+
| RAW     |
| Control |
+---------+                          +---------+        +---------+
| RAW +   |                          | RAW +   |        | RAW +   |
| DetNet  |                          | DetNet  |        | DetNet  |
| Service |                          | Service |        | Service |
+---------+--------------------------+---------+--------+---------+
|                       DetNet                                     |
|                     Forwarding                                   |
+------------------------------------------------------------------+

  Ingress             Transit            Relay              Egress
  Edge      ...       Nodes     ...      Nodes     ...        Edge
  Node                                                          Node

<--------------------Full Guarantees------------------------------->
```

Figure 6: (Strict) RAW over DetNet

In the second model (loose), 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 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).

```
   -------------------Flow Direction--------------------------------->

   +---------+
   | RAW     |
   | Control |
   +---------+           +---------+                      +---------+
   | RAW +   |           | DetNet  |                      | RAW +   |
   | DetNet  |           |  Only   |                      | DetNet  |
   | Service |           | Service |                      | Service |
   +---------+---------------------+---+               +---+---------+
   |          DetNet                   |               |   DetNet    |
   |         Forwarding                |               | Forwarding  |
   +-----------------------------------+               +-------------+

    Ingress    Transit       Relay           Internet           Egress
    End  ...   Nodes   ...   Nodes    ...                ...       End
    System                                                         System

   <----------------------No Guarantee------------------------------->
```

Figure 7: Loose RAW

5. The RAW Control Loop

5.1. Routing Time Scale vs. Forwarding Time Scale

With DetNet, the Controller Plane Function handles the routing computation and maintenance. With RAW, the routing part of the CPF (rCPF) is segregated from the RAW Control Loop, so it may reside outside of the RAW network. To achieve RAW capabilities, the rCPF is extended to generate the information required by the local aCPF, which acts as the orientation component in the loop. The rCPF may, e.g., propose DetNet Paths to be used as a reflex action in response to network events, or by provide aggregated history that the aCPF can use to make an oriented decision.

In a wireless mesh, the path to the DetNet CPF can be expensive and slow, possibly going across the whole mesh and back. Reaching to the CPF can also be slow in regards to the speed of events that affect the forwarding operation 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 CPF is not expected to be aware of and to 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 with the (remote) RAW rCPF is handled by a (local) aCPF that builds reports to the rCPF and digests the control information back, to be used inside a forwarding control loop for traffic steering.

```
                  +----------------+
                  |     DetNet     |
                  |    Routing     |
                  |      CPF       |
                  +----------------+
                          ^
                          |
                         Slow
                          |
     _-._-._-._-._-._-.   |  .-._-._-._-._-._-._-._-._-._-._-._-
   _-._-._-._-._-._-._-.  | _-._-._-._-._-._-._-._-._-._-._-._-
                          |
                       Expensive
                          |
                    ....  |  .......
               ....    .  | .       .......
            ....          v               ...
          ..    A-------B-------C---D        ..
       ...     /  \           /      \      ..
      .       I ----M-------N--***-- E        ..
      ..       \         /         /         ...
        ..      P--***--Q-----M---R        ....
          ..                              ....
           .   <----- Fast ------->    ....
            .......                ....
                   ................
```

*** = flapping at this time

Figure 8: 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 recovery graph that enables multiple Non-Equal Cost Multi-Path (N-ECMP) forwarding solutions along so-called protection paths, and leaves it to the Network Plane to make the per-packet decision of which of these possibilities should be used.

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).

5.2. A OODA Loop

OODA (Observe, Orient, Decide, Act) is a generic formalism to represent the operational steps in a Control Loop. The RAW Architecture applies that generic 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 Operations, Administration and Maintenance (OAM), to 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, more in Section 5.3;

Orient: An asynchronous CPF that reports data and information such as the link statistics, and leverages offline-computed wisdom and knowledge to Orient the PLR for its forwarding decision, more in Section 5.4;

Decide: A local PLR that decides which DetNet Path to use for the future packet(s) that are routed along the recovery graph, more in Section 5.5;

Act: PREOF Dataplane actions are controlled from the DetNet Service

sub-layer to increase the reliability of the end-to-end transmission. The RAW architecture also covers in-situ signaling when the decision is Acted by a node that down the recovery graph from the PLR, more in Section 5.6.

```
                +-------> Orient (aCPF) -------+
                |        reflex actions        |
                |       pre-trained model      |
                |             ...              |
                |                              v
            Observe (OAM)                Decide (PLR)
                ^                              |
                |                              |
                |                              |
                +-------- Act (PREOF) <--------+
                           At DetNet
                       Service sub-layer
```

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.

5.3. Observe: The RAW OAM

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 (NC) forward error correction (FEC), the RAW In-situ operation needs to be able to signal which operation occured 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 Assurance may observe the individual PREOF operation of a 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 a small continuous series of packets.

```
                                        ...   ..
                     RAN 1  -----  ...      ..  ...
                  /              .    ..          ....
        +-------+  /             .            ..      ....    +------+
        |Ingress|-               .                     .....  |Egress|
        |  End  |------ RAN 2 -- .       Internet       ....---| End  |
        |System |-               ..                    .....   |System|
        +-------+  \              .               ......      +------+
                    \               ...   ...     .....
                     RAN n  --------  ...   .....

                <------------------> <--------------------->
                   Observed by OAM       Opaque to OAM
```

Figure 10: Observed Links in Radio Access Protection

In the case of a 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 if a packet is attributed to the first hop Radio Access Network (RAN), even if a particular loss effectively happens farther down the path. In that case, RAW enables technology diversity (e.g. Wi-Fi and 5G) which in turn improves the diversity in spectrum usage.

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

5.4. Orient: The RAW-extended DetNet Operational Plane

RAW separates the long time scale at which a recovery graph is elaborated and installed, from the short time scale at which the forwarding decision is taken for one or a few packets (see in Section 5.1) that will 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 meta data about the links, segments and possible DetNet Paths. That meta data can be a pre-digested statistical model, and may include prediction of future flaps and packet loss, as well as recommended actions when that happens.

The meta data may include:

* A set of Pre-Determined DetNet Paths that are prepared to match expected link degradation profiles, so the DDCPEs can take reflex rerouting actions when facing a degradation that matches one such profile.

* 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 rCPF to achieve the RAW SLA. The objectives can be expressed as any of maximum number of packet lost in a row, bounded latency, maximal jitter, maximum number of interleaved out of order packets, average number of copies received at the elimination point, and maximal delay between the first and the last received copy of the same packet.

5.5. Decide: The Point of Local Repair

The RAW OODA Loop operates at the path selection time scale to provide agility vs. the brute force approach of flooding the whole recovery graph. The OODA Loop controls, within the redundant solutions that are proposed by the acynchronous CPF, 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 Point of Local Repair (PLR) as a synchronous CPF that performs rapid local adjustments of the forwarding tables within the diversity that the asynchronous CPF has in store for the recovery graph. The PLR enables to exploit the richer forwarding capabilities at a faster time scale over the smaller domain that is the recovery graph, in either a loose or a strict fashion.

The PLR operates on metrics that evolve faster, but that need to be advertised at a fast rate but only locally, within the recovery graph, and reacts on the metrics 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 rCPF that must observe the whole network and optimize all the recovery graphs globally, which can only be done at a slow pace and using long-term statistical metrics, as presented in Table 1.

	rCPF	PLR (In Scope)
Operation	Typically Centralized	Source-Routed or Distributed
Communication	Slow, expensive	Fast, local
Time Scale	hours and above	seconds and below
Network Size	Large, many recovery graphs to optimize globally	Small, within one recovery graph
Considered Metrics	Averaged, Statistical, Shade of grey	Instant values / boolean condition

Table 1: CPF vs. PLR

The PLR sits in the DetNet Service sub-Layer of Edge and Relay Nodes. On the one hand, it operates on the packet flow, learning the recovery graph and path selection information from the packet, possibly making 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 iOAM and oOAM) to obtain up-to-date information about its links and the quality of the overall recovery graph, respectively, as illustrated in Figure 11.

```
                |
        packet  | going
      down the  | stack
    +==========v==========+=====================+=====================+
    |   (iOAM + iCTRL)    | (L2 Triggers, DLEP) |       (oOAM)        |
    +==========v==========+=====================+=====================+
    |     Learn from      |                     |    Learn from       |
    |    packet tagging   >       Maintain      <    end-to-end       |
    +----------v----------+      Forwarding     |    OAM packets      |
    | Forwarding decision <        State        +---------^-----------|
    +----------v----------+                     |      Enrich or      |
    +    Retag Packet     |  Learn abstracted   >     Regenerate      |
    |    and Forward      | metrics about Links |     OAM packets     |
    +.........v..........+.........^.........+........^.v........+
    |                          Lower layers                           |
    +.........v....................^....................^.v........+
         frame | sent          Frame | L2 Ack        oOAM | | packet
          over | wireless        In  |                 In | | and out
               v                     |                    | v
```

Figure 11: PLR Interfaces

5.6. 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 RAW API enriches the DetNet protection services (PREOF) with potential possiblity to interact with lower layer one-hop reliability functions that are more typical to wireless than wires, including Automatic Repeat reQuest (ARQ), Forward Error Correction (FEC), Hybrid ARQ (HARQ) that includes both, 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 RAW 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 scheduling 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 select 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 guaranteed SLA, e.g., by rerouting and/or adding redundancy only when a spike of loss is observed.

6. 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.

6.1. Layer-2 encryption

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

6.2. Forced Access

A RAW policy may typically select the cheapest collection of links that matches the requested SLA, e.g., use 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.

7. IANA Considerations

This document has no IANA actions.

8. Contributors

The editor wishes to thank the document co-authors:

Lou Berger: Lab N

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

Geogios Papadopolous: IMT Atlantique

Remous-Aris Koutsiamanis: IMT Atlantique

Rex Buddenberg: Individual contributor

Greg Mirsky: Ericsson

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

9. Acknowledgments

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

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

10. References

10.1. Normative References

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

[INT-ARCHI]
Braden, R., Ed., "Requirements for Internet Hosts - Communication Layers", STD 3, RFC 1122, DOI 10.17487/RFC1122, October 1989, <https://www.rfc-editor.org/info/rfc1122>.

[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>.

[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-08, 10 July 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-raw-technologies-08>.

[RAW-USE-CASES] Bernardos, C. J., Papadopoulos, G. Z., Thubert, P., and F. Theoleyre, "RAW Use-Cases", Work in Progress, Internet-Draft, draft-ietf-raw-use-cases-11, 17 April 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-raw-use-cases-11>.

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

[RFC6291] Andersson, L., van Helvoort, H., Bonica, R., Romascanu, D., and S. Mansfield, "Guidelines for the Use of the "OAM" Acronym in the IETF", BCP 161, RFC 6291, DOI 10.17487/RFC6291, June 2011, <https://www.rfc-editor.org/info/rfc6291>.

[RFC7799] Morton, A., "Active and Passive Metrics and Methods (with Hybrid Types In-Between)", RFC 7799, DOI 10.17487/RFC7799, May 2016, <https://www.rfc-editor.org/info/rfc7799>.

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

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

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

[RFC8655] Finn, N., Thubert, P., Varga, B., and J. Farkas, "Deterministic Networking Architecture", RFC 8655, DOI 10.17487/RFC8655, October 2019, <https://www.rfc-editor.org/info/rfc8655>.

[RFC8939] Varga, B., Ed., Farkas, J., Berger, L., Fedyk, D., and S. Bryant, "Deterministic Networking (DetNet) Data Plane: IP", RFC 8939, DOI 10.17487/RFC8939, November 2020, <https://www.rfc-editor.org/info/rfc8939>.

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

10.2. Informative References

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

[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>.

[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>.

[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>.

[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>.

[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>.

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

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

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

[DLEP] Ratliff, S., Jury, S., Satterwhite, D., Taylor, R., and B. Berry, "Dynamic Link Exchange Protocol (DLEP)", RFC 8175, DOI 10.17487/RFC8175, June 2017, <https://www.rfc-editor.org/info/rfc8175>.

[I-D.irtf-panrg-path-properties] Enghardt, R. and C. Krähenbühl, "A Vocabulary of Path Properties", Work in Progress, Internet-Draft, draft-irtf-panrg-path-properties-08, 6 March 2023, <https://datatracker.ietf.org/doc/html/draft-irtf-panrg-path-properties-08>.

[IPoWIRELESS] Thubert, P. and M. Richardson, "Architecture and Framework for IPv6 over Non-Broadcast Access", Work in Progress, Internet-Draft, draft-thubert-6man-ipv6-over-wireless-15, 8 March 2023, <https://datatracker.ietf.org/doc/html/draft-thubert-6man-ipv6-over-wireless-15>.

[DetNet-OAM] Mirsky, G., Theoleyre, F., Papadopoulos, G. Z., Bernardos, C. J., Varga, B., and J. Farkas, "Framework of Operations, Administration and Maintenance (OAM) for Deterministic Networking (DetNet)", Work in Progress, Internet-Draft, draft-ietf-detnet-oam-framework-11, 8 January 2024, <https://datatracker.ietf.org/doc/html/draft-ietf-detnet-oam-framework-11>.

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

Author's Address

Pascal Thubert (editor)
06330 Roquefort-les-Pins
France

Email: pascal.thubert@gmail.com

RAW
Internet-Draft
Intended status: Informational
Expires: 16 March 2024

P. Thubert, Ed.

L. Berger
LabN Consulting, L.L.C.

13 September 2023

Reliable and Available Wireless Framework
draft-ietf-raw-framework-02

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 following an OODA loop that involves OAM, PCE, PSE and PAREO functions. 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

Copyright Notice

Table of Contents

1. Introduction

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 deterministic layer-2 capabilities with diversity in the spatial, time, code, technology, 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.

The "RAW Architecture" [RAW-ARCHI] document presents the RAW problem and architectural concepts such as path and Tracks to provide IPv6 [IPv6] flows with Service Level Objectives (SLO) in terms of packet delivery ratio (PDR), maximum contiguous losses or latency boundaries over wireless access and meshes.

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

1. Network Plane measurement protocols for Operations, Administration and Maintenance (OAM) to Observe some or all hops along a Track as well as the end-to-end packet delivery

2. Controller plane elements to reports the links statistics to a Path computation Element (PCE) in a centralized controller that computes and installs the Tracks and provides meta data to Orient the routing decision

3. A Runtime distributed Path Selection Engine (PSE) that Decides which subTrack to use for the next packet(s) that are routed along the Track

4. Packet (hybrid) ARQ, Replication, Elimination and Ordering Dataplane actions that operate at the DetNet Service Layer to increase the reliability of the end-to-end transmission. The RAW architecture also covers in-situ signaling when the decision is Acted by a node that down the Track from the PSE.

The RAW Framework combines IETF specification to enable RAW Service Level Agreements (SLA) over a selected technologies [RAW-TECHNOS]. The framework implements the OODA loop to optimizes the use of redundancy while minimizing the use of constrained resources such as spectrum and battery.

2. Terminology

This document uses the terminology defined in the "RAW Architecture" [RAW-ARCHI].

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.

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.

```
                                   ...   ..
                 RAN 1  -----  ...     ..  ...
              /              .      ..         ....
   +--------+  /              .                   ....    +-----------+
   |Wireless|-                .                    .....  |  Service  |
   | Device |-***-- RAN 2 -- .       Internet       ....---|     /     |
   |(STA/UE)|-                ..                   .....   |Application|
   +--------+  $$$             .                .......     +-----------+
                \                ...   ...     .....
                 RAN n  --------  ...   .....

   *** = 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.

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 6).

```
                        A-------B-------C-----D
                       /  \   /       /        \
                Ingress ----M-------N--zzzzz--- Egress
                       \      \   /            /
                        P--zzz--Q-------------R

               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 enables the intermediate Nodes to estimate the quality of the rest of the feasible paths in the remainder of the Track to the destination.

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) [DLEP] 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, and obtain link characteristics such as speed in a timely manner.

* [detnet] provides an OAM framework with [DetNet-OAM] that applies within the DetNet dataplane described in [DetNet-DP], which is typically based on MPLS or IPv6 pseudowires.

* Bidirectional Forwarding Detection [BFD] and its variants (bidirectional and remote BFD) detects faults in the path between an Ingress and an Egress forwarding engines. The art of BFD expects a serial path and needs one session per path, which makes it suited to observe a Segment it is unaware of the complexity of a track , and expects bidirectionality. to protect a path. [BFD] BFD asynchronous mode considers delivery as success whereas with DetNet and RAW, the bounded latency can be as important as the delivery itself, and delivering too late is actually a failure. Note that the BFD Demand mode with unsolicited notifications may be more suitable then the Asynchronous BFD mode. The use of the Demand mode in MPLS is analyzed in [I-D.mirsky-bfd-mpls-demand] and similar considerations could apply to IP as well.

* [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.

* [IPPM] develops and maintains standard metrics that can be applied to the quality, performance, and reliability of Internet data delivery services and applications running over transport layer protocols (e.g. TCP, UDP) over IP.

5. 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.

6. 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 N-ECMP forwarding solutions. The Complex Track enables to support multi-path 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.

7. 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 network layer identification of an application layer flow (typically ther 5- or 6- tuple) 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.

```
                       Flow 1 (6-tuple) ----+
                                            |
                  Flow 2 (6-tuple)  ---+    |
                                       |    |
               OAM     -----------+    |    |
                                  |    |    |
                                  |    |    |
                             |    |    |    |    |
                             |    v    v    v    |
                             |                   |
                             +---------+---------+
                                       |
                                       |
                Track i (Ingress IP Address, Track Id)
                                       |
                                       |
                                       |
                       +---------+-----+--....-------+
                       |         |                   |
                       |         |                   |
                subTrack 1    subTrack 2          subTrack n
                       |         |                   |
                       |         |                   |
                       V         V                   V
                   +-----------------------------------+
                   |                                   |
                   |         Destination               |
                   |                                   |
                   +-----------------------------------+
```

Figure 3: Flow Injection

With 6TiSCH, packets are tagged with the same (source 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].

8. 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.

9. 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.

10. Operations Administration and Maintenance

10.1. DetNet OAM

[detnet] provides an OAM framework with [DetNet-OAM] that applies within the DetNet dataplane described in [DetNet-DP],which is typically based on MPLS or IPv6 pseudowires. How the framework applies to IPv6 is detailed in [DetNet-IP-OAM]. Within that framework, OAM messages follow the same forward path as the data packets and gather information about their individual treatment at each hop. When the destination receives an OAM message, it gets a view on the full path or at least of a segment of the path from the source of the flow.

In-situ OAM (IOAM) adds telemetry information about the experience of one packet within the packet itself [I-D.ietf-ippm-ioam-data], with the caveats that the measurement and the consecutive update of the packet interfere with the operation being observed, e.g., may increase the latency of the packet for which it is measured and into which it is stamped.

Note: IOAM and analogous on-path telemetry methods are capable of facilitating collection of useful telemetry information that characterizes the state of a system as experienced by the packet. But because of statistical character of a packet network, these methods may not be used to monitor the continuity of a path (Track) or proper connectivity of the Track (no leaking packets across Tracks).

This effect can be alleviated by measuring on the fly but reporting later, e.g., by exporting the data as a separate management packet [I-D.ietf-ippm-ioam-direct-export].
[I-D.mirsky-ippm-hybrid-two-step] proposes an hybrid two-steps method (HTS) where a trigger message starts the measurement and a follow up along the Track packet gathers the measured data.

"Error Performance Measurement" [I-D.mirsky-ippm-epm] uses Fault Management (FM) and Performance Management (PM) OAM mechanisms to determine availability/unavailability of a path according to predefined SLA.

10.2. RAW Extensions

Classical OAM typically measures information at the transmitter, e.g., residence time in the node or transmit queue size. With RAW, there is a need to combine information at the sender (number of retries) with that at the receiver (LQI, RSSI). This doubles the operating cost of an IAOM processing that would gather the experience of a single packet.

The RAW PSE may be centralized at the Track Ingress, or distributed long the Track. Either way, the PSE needs instant information about the rest of the way to the destination over the possible next-hop adjacencies along the Track in order to decide how to perform simple forwarding, load balancing, and/or replication, as well as determining how much latency credit is available for ARQ.

To provide that information timely, it makes sense that the OAM packets that gather instantaneous values from the radio senders and receivers at each hop flow on the reverse path and inform the PSE at the source and/or the PAREO relays about the state of the rest of the way. This is achieved using Reverse OAM packets that flow along the Reversed Track, West to East.

Because the quality of transmission over a wireless medium varies continuously, it is important that RAW OAM captures the state of the medium across an adjacency over multiple transmission and over a recent period of time, whether the transmitted packets belong to this flow or another. Some of the measured information relates to the medium itself. In other words, the captured information does not only relate to the experience of one packet as is the case for IOAM, but also to the medium itself. This makes an approach like HTS more suitable as it can trigger the capture of multiple measurements over a short period of time. On the other hand, the PSE needs a continuous measurement stream where a single trigger is followed by a periodic follow up capture.

In other words, the best suited OAM method to enable the PSE make accurate PAREO forwarding decisions is a periodic variation of the two-steps method flowing along the reverse Track, as an upstream OAM technique. [RAW-OAM] provides more information on the RAW OAM problem and solution approaches.

10.3. Observed Metrics

The Dynamic Link Exchange Protocol (DLEP) [DLEP] 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.

Those lower-layer metrics are aggregated along a multihop segment into abstract layer 3 information that reflect the instant reliability and latency of the observed path.

11. 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.

11.1. 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.

12. IANA Considerations

This document has no IANA actions.

13. Acknowledgments

The editor wishes to thank:

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

Remous-Aris Koutsiamanis: IMT Atlantique

Nicolas Montavont: IMT Atlantique

Georgios Z. Papadopoulos: IMT Atlantique

Rex Buddenberg: Individual contributor

Greg Mirsky: ZTE

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

14. References

14.1. 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-ARCHI]
Thubert, P., "Reliable and Available Wireless Architecture", Work in Progress, Internet-Draft, draft-ietf-raw-architecture-15, 14 August 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-raw-architecture-15>.

[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>.

[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>.

[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>.

[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>.

14.2. Informative References

[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-08, 10 July 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-raw-technologies-08>.

[RAW-USE-CASES]
Bernardos, C. J., Papadopoulos, G. Z., Thubert, P., and F. Theoleyre, "RAW Use-Cases", Work in Progress, Internet-Draft, draft-ietf-raw-use-cases-11, 17 April 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-raw-use-cases-11>.

[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. 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://datatracker.ietf.org/doc/html/draft-thubert-bier-replication-elimination-03>.

[DetNet-IP-OAM]
Mirsky, G., Chen, M., and D. L. Black, "Operations, Administration and Maintenance (OAM) for Deterministic Networks (DetNet) with IP Data Plane", Work in Progress, Internet-Draft, draft-ietf-detnet-ip-oam-09, 1 August 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-detnet-ip-oam-09>.

[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://datatracker.ietf.org/doc/html/draft-farkas-raw-5g-00>.

[BIER-TE] Eckert, T. T., Menth, M., and G. Cauchie, "Tree Engineering for Bit Index Explicit Replication (BIER-TE)", Work in Progress, Internet-Draft, draft-ietf-bier-te-arch-13, 25 April 2022, <https://datatracker.ietf.org/doc/html/draft-ietf-bier-te-arch-13>.

[RAW-OAM] Theoleyre, F., Papadopoulos, G. Z., Mirsky, G., and C. J. Bernardos, "Operations, Administration and Maintenance (OAM) features for RAW", Work in Progress, Internet-Draft, draft-ietf-raw-oam-support-06, 5 March 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-raw-oam-support-06>.

[I-D.ietf-ippm-ioam-direct-export] Song, H., Gafni, B., Brockners, F., Bhandari, S., and T. Mizrahi, "In Situ Operations, Administration, and Maintenance (IOAM) Direct Exporting", Work in Progress, Internet-Draft, draft-ietf-ippm-ioam-direct-export-11, 23 September 2022, <https://datatracker.ietf.org/doc/html/draft-ietf-ippm-ioam-direct-export-11>.

[DetNet-OAM] Mirsky, G., Theoleyre, F., Papadopoulos, G. Z., Bernardos, C. J., Varga, B., and J. Farkas, "Framework of Operations, Administration and Maintenance (OAM) for Deterministic Networking (DetNet)", Work in Progress, Internet-Draft, draft-ietf-detnet-oam-framework-08, 1 February 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-detnet-oam-framework-08>.

[I-D.mirsky-ippm-hybrid-two-step] Mirsky, G., Lingqiang, W., Zhui, G., Song, H., and P. Thubert, "Hybrid Two-Step Performance Measurement Method", Work in Progress, Internet-Draft, draft-mirsky-ippm-hybrid-two-step-15, 15 June 2023, <https://datatracker.ietf.org/doc/html/draft-mirsky-ippm-hybrid-two-step-15>.

[I-D.mirsky-ippm-epm]
Mirsky, G., Halpern, J. M., Min, X., and L. Han, "Error Performance Measurement in Packet-switched Networks", Work in Progress, Internet-Draft, draft-mirsky-ippm-epm-04, 24 October 2021, <https://datatracker.ietf.org/doc/html/draft-mirsky-ippm-epm-04>.

[I-D.mirsky-bfd-mpls-demand]
Mirsky, G., "BFD in Demand Mode over a Point-to-Point MPLS LSP", Work in Progress, Internet-Draft, draft-mirsky-bfd-mpls-demand-14, 9 May 2023, <https://datatracker.ietf.org/doc/html/draft-mirsky-bfd-mpls-demand-14>.

[I-D.ietf-ippm-ioam-data]
Brockners, F., Bhandari, S., and T. Mizrahi, "Data Fields for In Situ Operations, Administration, and Maintenance (IOAM)", Work in Progress, Internet-Draft, draft-ietf-ippm-ioam-data-17, 13 December 2021, <https://datatracker.ietf.org/doc/html/draft-ietf-ippm-ioam-data-17>.

[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/>.

[IPPM] IETF, "IP Performance Measurement", <https://dataTracker.ietf.org/doc/charter-ietf-ippm/>.

Authors' Addresses

Pascal Thubert (editor)
06330 Roquefort-les-Pins
France

Email: pascal.thubert@gmail.com

Lou Berger
LabN Consulting, L.L.C.
United States of America
Email: lberger@labn.net

RAW
Internet-Draft
Intended status: Informational
Expires: 5 June 2023

N. Mäurer, Ed.
T. Gräupl, Ed.
German Aerospace Center (DLR)
C. Schmitt, Ed.
Research Institute CODE, UniBwM

2 December 2022

L-band Digital Aeronautical Communications System (LDACS)
draft-ietf-raw-ldacs-14

Abstract

This document gives an overview of the architecture of the L-band Digital Aeronautical Communications System (LDACS), which provides a secure, scalable and spectrum efficient terrestrial data link for civil aviation. LDACS is a scheduled, reliable multi-application cellular broadband system with support for IPv6. It is part of a larger shift of flight guidance communication moving to IP-based communication. High reliability and availability of IP connectivity over LDACS, as well as security, are therefore essential. The intent of this document is to introduce LDACS to the IETF community, raise awareness on related activities inside and outside of the IETF, and to seek expertise in shaping the shift of aeronautics to IP.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

One of the main pillars of the modern Air Traffic Management (ATM) system is the existence of a communications infrastructure that enables efficient aircraft control and safe aircraft separation in all phases of flight. Current systems are technically mature but suffering from the Very High Frequency (VHF) band's increasing saturation in high-density areas and the limitations posed by analogue radio communications. Therefore, aviation strives for a sustainable modernization of the aeronautical communications infrastructure on the basis of IP.

This modernization is realized in two steps: (1) the transition of communications datalinks from analogue to digital technologies and, (2) the introduction of IPv6 based networking protocols [RFC8200] in aeronautical networks [ICAO2015].

Step (1) is realized via ATM communications transitioning from analogue VHF voice [KAMA2010] to more spectrum efficient digital data communication. For terrestrial communications the International Civil Aviation Organization (ICAO)'s Global Air Navigation Plan (GANP) foresees this transition to be realized by the development of the L-band Digital Aeronautical Communications System (LDACS). Since Central Europe has been identified as the area of the world that suffers the most from increased saturation of the VHF band, the initial roll-out of LDACS will likely start there, and continue to other increasingly saturated zones as the East and West Coast of the US and parts of Asia [ICAO2018].

Technically LDACS enables IPv6-based air-ground communication related to aviation safety and regularity of flight [ICAO2015]. Passenger communication and similar services are not supported, since only

communications related to "safety and regularity of flight" are permitted in protected aviation frequency bands. The particular challenge is that no additional frequencies can be made available for terrestrial aeronautical communication. It was thus necessary to develop co-existence mechanism/procedures to enable the interference free operation of LDACS in parallel with other aeronautical services/ systems in the protected frequency band. Since LDACS will be used for aircraft guidance, high reliability and availability for IP connectivity over LDACS are essential.

LDACS is standardized in ICAO and European Organization for Civil Aviation Equipment (EUROCAE).

This document provides information to the IETF community about the aviation industry transition of flight guidance systems from analog to digital, provides context for LDACS relative to related IETF activities [I-D.haindl-lisp-gb-atn], and seeks expertise on realizing reliable IPv6 over LDACS for step (1). This document does not intend to advance LDACS as an IETF standards-track document.

Step (2) is a strategy for the worldwide roll-out of IPv6 capable digital aeronautical inter-networking. This is called the Aeronautical Telecommunications Network (ATN)/Internet Protocol Suite (IPS) (hence, ATN/IPS). It is specified in the International Civil Aviation Organization (ICAO) document Doc 9896 [ICAO2015], the Radio Technical Commission for Aeronautics (RTCA) document DO-379 [RTCA2019], the EUROCAE document ED-262 [EURO2019], and the Aeronautical Radio Incorporated (ARINC) document P858 [ARI2021]. LDACS is subject to these regulations since it provides an "access network" - link-layer datalink - to the ATN/IPS.

ICAO has chosen IPv6 as basis for the ATN/IPS mostly for historical reasons, since a previous architecture based on ISO/OSI protocols, the ATN/OSI, failed in the marketplace.

In the context of safety-related communications, LDACS will play a major role in future ATM. ATN/IPS datalinks will provide diversified terrestrial and space-based connectivity in a multilink concept, called the Future Communications Infrastructure (FCI) [VIR2021]. From a technical point of view the FCI will realize airborne multi-homed IPv6 networks connected to a global ground network via at least two independent communication technologies. This is considered in more detail in related IETF work in progress [I-D.haindl-lisp-gb-atn] [I-D.ietf-rtgwg-atn-bgp]. As such, ICAO has actively sought out the support of IETF to define a mobility solution for step (2), which is currently the Locator/ID Separation Protocol (LISP).

In the context of the Reliable and Available Wireless (RAW) working group, developing options, such as intelligent switching between datalinks, for reliably delivering content from and to endpoints, is foreseen. As LDACS is part of such a concept, the work of RAW is immediately applicable. In general, with the aeronautical communications system transitioning to ATN/IPS, and data being transported via IPv6, closer cooperation and collaboration between the aeronautical and IETF community is desirable.

LDACS standardization within the framework of ICAO started in December 2016. The ICAO standardization group has produced the final Standards and Recommended Practices (SARPS) document as of 2022 [ICAO2022]. It defines the general characteristics of LDACS. By the end of 2023, the ICAO standardization group plans to have developed an ICAO technical manual - the ICAO equivalent to a technical standard. As such LDACS standardization is not finished yet, and therefore this document is a snapshot of current status. The physical characteristics of an LDACS installation (form, fit, and function) will be standardized by EUROCAE. Generally, the group is open to input from all sources and encourages cooperation between the aeronautical and the IETF community.

2. Acronyms

The following terms are used in the context of RAW in this document:

A/A Air/Air
A/G Air/Ground
A2G Air-to-Ground
ACARS Aircraft Communications Addressing and Reporting System
AC-R Access Router
ADS-B Automatic Dependent Surveillance - Broadcast
ADS-C Automatic Dependent Surveillance - Contract
AeroMACS Aeronautical Mobile Airport Communications System
ANSP Air Traffic Network Service Provider
AOC Aeronautical Operational Control
ARINC Aeronautical Radio, Incorporated
ARQ Automatic Repeat reQuest
AS Aircraft Station
ATC Air Traffic Control
ATM Air Traffic Management
ATN Aeronautical Telecommunication Network
ATS Air Traffic Service
BCCH Broadcast Channel
CCCH Common Control Channel
CM Context Management
CNS Communication Navigation Surveillance
COTS Commercial Off-The-Shelf

CPDLC Controller Pilot Data Link Communications
CRL Certificate Revocation List
CSP Communications Service Provider
DCCH Dedicated Control Channel
DCH Data Channel
DiffServ Differentiated Services
DLL Data Link Layer
DLS Data Link Service
DME Distance Measuring Equipment
DSB-AM Double Side-Band Amplitude Modulation
DTLS Datagram Transport Layer Security
EUROCAE European Organization for Civil Aviation Equipment
FAA Federal Aviation Administration
FCI Future Communications Infrastructure
FDD Frequency Division Duplex
FL Forward Link
GANP Global Air Navigation Plan
GBAS Ground Based Augmentation System
GNSS Global Navigation Satellite System
GS Ground-Station
G2A Ground-to-Air
HF High Frequency
ICAO International Civil Aviation Organization
IP Internet Protocol
IPS Internet Protocol Suite
kbit/s kilobit per second
LDACS L-band Digital Aeronautical Communications System
LISP Locator/ID Separation Protocol
LLC Logical Link Control
LME LDACS Management Entity
MAC Medium Access Control
MF Multi Frame
NETCONF NETCONF Network Configuration Protocol
OFDM Orthogonal Frequency-Division Multiplexing
OFDMA Orthogonal Frequency-Division Multiplexing Access
OSI Open Systems Interconnection
PHY Physical Layer
QPSK Quadrature Phase-Shift Keying
RACH Random-Access Channel
RL Reverse Link
RTCA Radio Technical Commission for Aeronautics
SARPS Standards and Recommended Practices
SDR Software Defined Radio
SESAR Single European Sky ATM Research
SF Super-Frame
SNMP Simple Network Management Protocol
SNP Sub-Network Protocol
VDLm2 VHF Data Link mode 2

VHF Very High Frequency
VI Voice Interface

3. Motivation and Use Cases

Aircraft are currently connected to Air Traffic Control (ATC) and Aeronautical Operational Control (AOC) services via voice and data communications systems through all phases of flight. ATC refers to communication for flight guidance. AOC is a generic term referring to the business communication of airlines. It refers to the mostly proprietary exchange of data between the aircraft of the airline and the airline's operation centers and service partners. The ARINC document 633 was developed and first released in 2007 [ARI2019] with the goal to standardize these messages for interoperability, e.g., messages between the airline and fueling or de-icing companies. Within the airport and terminal area, connectivity is focused on high bandwidth communications. In the En-Route domain, however, high reliability, robustness, and range are the main foci. Voice communications may use the same or different equipment as data communications systems. In the following, the main differences between voice and data communications capabilities are summarized. The assumed list of use cases for LDACS complements the list of use cases stated in [RAW-USE-CASES] and the list of reliable and available wireless technologies presented in [RAW-TECHNOS].

3.1. Voice Communications Today

Voice links are used for Air/Ground (A/G) and Air/Air (A/A) communications. The communications equipment can be installed on ground or in the aircraft, in which cases the High Frequency (HF) or VHF frequency band is used. For remote domains voice communications can also be satellite-based. All VHF and HF voice communications are operated via open broadcast channels without authentication, encryption or other protective measures. The use of well-proven communications procedures via broadcast channels, such as phraseology or read-backs, requiring well-trained personnel, help to enhance the safety of communications, but does not replace necessary cryptographical security mechanisms. The main voice communications media is still the analogue VHF Double Side-Band Amplitude Modulation (DSB-AM) communications technique, supplemented by HF single side-band amplitude modulation and satellite communications for remote and oceanic regions. DSB-AM has been in use since 1948, works reliably and safely, and uses low-cost communication equipment. These are the main reasons why VHF DSB-AM communications are still in use, and it is likely that this technology will remain in service for many more years. This however, results in current operational limitations and impediments in deploying new ATM applications, such as flight-centric

operation with point-to-point communications between pilots and air traffic control officers. [BOE2019]

3.2. Data Communications Today

Like for voice, data communications into the cockpit, are currently provided by ground-based equipment operating either on HF or VHF radio bands or by legacy satellite systems. All these communication systems use narrowband radio channels with a data throughput capacity in the order of kbit/s. While the aircraft is on the ground, some additional communications systems are available, like the Aeronautical Mobile Airport Communications System (AeroMACS) or public cellular networks, operating in the Airport (APT) domain and able to deliver broadband communications capability. [BOE2019]

For regulatory reasons, the data communications networks, used for the transmission of data relating to the safety and regularity of flight, must be strictly isolated from those providing entertainment services to passengers. This leads to a situation that the flight crews are supported by narrowband services during flight while passengers have access to inflight broadband services. The current HF and VHF data links cannot provide broadband services now or in the future, due to the lack of available spectrum. This technical shortcoming is becoming a limitation to enhanced ATM operations, such as trajectory-based operations and 4D trajectory negotiations. [BOE2019]

Satellite-based communications are currently under investigation and enhanced capabilities are under development which will be able to provide inflight broadband services and communications supporting the safety and regularity of flight. In parallel the ground-based broadband data link technology LDACS is being standardized by ICAO and has recently shown its maturity during flight tests [MAE20211] [BEL2021]. The LDACS technology is scalable, secure and spectrum efficient and provides significant advantages to the users and service providers. It is expected that both - satellite systems and LDACS - will be deployed to support the future aeronautical communication needs as envisaged by the ICAO Global Air Navigation Plan (GNAP). [BOE2019]

4. Provenance and Documents

The development of LDACS has already made substantial progress in the Single European Sky ATM Research (SESAR) framework and is currently being continued in the follow-up program SESAR2020 [RIH2018]. A key objective of these activities is to develop, implement and validate a modern aeronautical data link able to evolve with aviation needs over the long term. To this end, an LDACS specification has been produced

[GRA2020] and is continuously updated; transmitter demonstrators were developed to test the spectrum compatibility of LDACS with legacy systems operating in the L-band [SAJ2014]; and the overall system performance was analyzed by computer simulations, indicating that LDACS can fulfil the identified requirements [GRA2011].

Up to now LDACS standardization has been focused on the development of the physical layer and the data link layer. Only recently have higher layers have come into the focus of the LDACS development activities. Currently no "IPv6 over LDACS" specification is defined; however, SESAR2020 has started experimenting with IPv6-based LDACS and ICAO plans to seek guidance from IETF to develop IPv6 over LDACS. As of May 2022, LDACS defines 1536 Byte user-data packets [GRA2020] in which IPv6 traffic shall be encapsulated. Additionally, Robust Header Compression (ROHC) is considered on LDACS Sub-Network Protocol (SNP) layer (cf. Section 7.3.3.) [RFC5795].

The IPv6 architecture for the aeronautical telecommunication network is called the ATN/IPS. Link-layer technologies within the ATN/IPS encompass LDACS [GRA2020], AeroMACS [KAMA2018] and several SatCOM candidates and combined with the ATN/IPS, are called the FCI. The FCI will support quality of service, link diversity, and mobility under the umbrella of the "multilink concept". The "multilink concept" describing the idea that depending on link quality, communication can be switched seamlessly from one datalink technology to another. This work is led by ICAO Communication Panel working group WG-I.

In addition to standardization activities several industrial LDACS prototypes have been built. One set of LDACS prototypes has been evaluated in flight trials confirming the theoretical results predicting the system performance [GRA2018] [MAE20211] [BEL2021].

5. Applicability

LDACS is a multi-application cellular broadband system capable of simultaneously providing various kinds of Air Traffic Services (ATS) including ATS-B3, and AOC communications services from deployed Ground-Stations (GS). The physical layer and data link layer of LDACS are optimized for controller-pilot data link communications, but the system also supports digital air-ground voice communications.

LDACS supports communications in all airspaces (airport, terminal maneuvering area, and en-route), and on the airport surface. The physical LDACS cell coverage is effectively de-coupled from the operational coverage required for a particular service. This is new in aeronautical communications. Services requiring wide-area coverage can be installed at several adjacent LDACS cells. The

handover between the involved LDACS cells is seamless, automatic, and transparent to the user. Therefore, the LDACS communications concept enables the aeronautical communication infrastructure to support future dynamic airspace management concepts.

5.1. Advances Beyond the State-of-the-Art

LDACS will offer several capabilities, not yet provided in contemporarily deployed aeronautical communications systems. All these were already tested and confirmed in lab or flight trials with available LDACS prototype hardware [BEL2021] [MAE20211].

5.1.1. Priorities

LDACS is able to manage service priorities, an important feature not available in some of the current data link deployments. Thus, LDACS guarantees bandwidth availability, low latency, and high continuity of service for safety critical ATS applications while simultaneously accommodating less safety-critical AOC services.

5.1.2. Security

LDACS is a secure data link with built-in security mechanisms. It enables secure data communications for ATS and AOC services, including secured private communications for aircraft operators and Air traffic Network Service Providers (ANSP). This includes concepts for key and trust management, mutual authentication and key establishment protocols, key derivation measures, user and control message-in-transit protection, secure logging and availability and robustness measures [MAE20182] [MAE2021].

5.1.3. High Data Rates

The user data rate of LDACS is 315 kbit/s to 1428 kbit/s on the Forward Link (FL) for the Ground-to-Air (G2A) connection, and 294 kbit/s to 1390 kbit/s on the Reverse Link (RL) for the Air-to-Ground (A2G) connection, depending on coding and modulation. This is up to two orders of magnitude greater than current terrestrial digital aeronautical communications systems, such as the VHF Data Link mode 2 (VDLm2), provide [ICAO2019] [GRA2020].

5.2. Application

LDACS will be used by several aeronautical applications ranging from enhanced communications protocol stacks (multi-homed mobile IPv6 networks in the aircraft and potentially ad-hoc networks between aircraft) to broadcast communication applications (GNSS correction data) and integration with other service domains (using the communications signal for navigation) [MAE20211]. Also, a digital voice service offering better quality and service than current HF and VHF systems is foreseen.

5.2.1. Air/Ground Multilink

It is expected that LDACS, together with upgraded satellite-based communications systems, will be deployed within the FCI and constitute one of the main components of the multilink concept within the FCI.

Both technologies, LDACS and satellite systems, have their specific benefits and technical capabilities which complement each other. Especially, satellite systems are well-suited for large coverage areas with less dense air traffic, e.g. oceanic regions. LDACS is well-suited for dense air traffic areas, e.g., continental areas or hot-spots around airports and terminal airspace. In addition, both technologies offer comparable data link capacity and, thus, are well-suited for redundancy, mutual back-up, or load balancing.

Technically the FCI multilink concept will be realized by multi-homed mobile IPv6 networks in the aircraft. The related protocol stack is currently under development by ICAO, within SESAR, and the IETF. Currently two layers of mobility are foreseen. Local mobility within the LDACS access network is realized through PMIPv6, global mobility between "multi-link" access networks (which need not be LDACS) is implemented on top of LISP [I-D.haindl-lisp-gb-atn] [I-D.ietf-lisp-rfc6830bis] [I-D.ietf-lisp-rfc6833bis].

5.2.2. Air/Air Extension for LDACS

A potential extension of the multilink concept is its extension to the integration of ad-hoc networks between aircraft.

Direct A/A communication between aircraft in terms of ad-hoc data networks are currently considered a research topic since there is no immediate operational need for it, although several possible use cases are discussed (Automatic Dependent Surveillance - Broadcast (ADS-B), digital voice, wake vortex warnings, and trajectory negotiation) [BEL2019]. It should also be noted, that currently deployed analog VHF voice radios support direct voice communication between aircraft, making a similar use case for digital voice plausible.

LDACS A/A is currently not part of the standardization process and will not be covered within this document. However, it is planned that LDACS A/A will be rolled out after the initial deployment of LDACS A/G, then being seamlessly integrated in the existing LDACS ground-based system.

5.2.3. Flight Guidance

The FCI (and therefore LDACS) is used to provide flight guidance. This is realized using three applications:

1. Context Management (CM): The CM application manages the automatic logical connection to the ATC center currently responsible to guide the aircraft. Currently, this is done by the air crew manually changing VHF voice frequencies manually according to the progress of the flight. The CM application automatically sets up equivalent sessions.
2. Controller Pilot Data Link Communications (CPDLC): The CPDLC application provides the air crew with the ability to exchange data messages similar to text messages with the currently responsible ATC center. The CPDLC application takes over most of the communication currently performed over VHF voice and enables new services that do not lend themselves to voice communication (i.e., trajectory negotiation).
3. Automatic Dependent Surveillance - Contract (ADS-C): ADS-C reports the position of the aircraft to the currently active ATC center. Reporting is bound to "contracts", i.e., pre-defined events related to the progress of the flight (i.e., the trajectory). ADS-C and CPDLC are the primary applications used for implementing in-flight trajectory management.

CM, CPDLC, and ADS-C are available on legacy datalinks, but are not widely deployed and with limited functionality.

Further ATC applications may be ported to use the FCI or LDACS as well. A notable application is GBAS for secure, automated landings: The Global Navigation Satellite System (GNSS) based GBAS is used to improve the accuracy of GNSS to allow GNSS based instrument landings.

This is realized by sending GNSS correction data (e.g., compensating ionospheric errors in the GNSS signal) to the aircraft's GNSS receiver via a separate data link. Currently, the VDB data link is used. VDB is a narrowband single-purpose datalink without advanced security only used to transmit GBAS correction data. This makes VDB a natural candidate for replacement by LDACS [MAE20211].

5.2.4. Business Communications of Airlines

In addition to air traffic services, AOC services are transmitted over LDACS. AOC is a generic term referring to the business communication of airlines, between the airlines and service partners on the ground and their own aircraft in the air. Regulatory-wise, this is considered related to safety and regularity of flight and may therefore, be transmitted over LDACS. AOC communication is considered the main business case for LDACS communications service providers since modern aircraft generate significant amounts of data (e.g., engine maintenance data).

5.2.5. LDACS-based Navigation

Beyond communications, radio signals can always also be used for navigation. This fact is used for the LDACS navigation concept.

For future aeronautical navigation, ICAO recommends the further development of GNSS based technologies as primary means for navigation. Due to the large separation between navigational satellites and aircraft, the power of the GNSS signals received by the aircraft is, however, very low. As a result, GNSS disruptions might occasionally occur due to unintentional interference, or intentional jamming. Yet the navigation services must be available with sufficient performance for all phases of flight. Therefore, during GNSS outages, or blockages, an alternative solution is needed. This is commonly referred to as Alternative Positioning, Navigation, and Timing (APNT).

One such APNT solution is based on exploiting the built-in navigation capabilities of LDACS operation. That is, the normal operation of LDACS for ATC and AOC communications would also directly enable the aircraft to navigate and obtain a reliable timing reference from the LDACS GSs. Current cell planning for Europe shows 84 LDACS cells to be sufficient [MOST2018] to cover the continent at sufficient service level. If more than three Ground Stations (GS) are visible by the aircraft, via knowing the exact positions of these and having a good channel estimation (which LDACS does due to numerous works mapping the L-band channel characteristics [SCHN2018]) it is possible to calculate the position of the aircraft via measuring signal propagation times to each GS. In flight trials in 2019 with one

aircraft (and airborne radio inside it) and just four GS, navigation feasibility was demonstrated within the footprint of all four GS with a 95th percentile position-domain error of 171.1m [OSE2019] [BEL2021] [MAE20211]. As such, LDACS can be used independently of GNSS as a navigation alternative. Positioning errors will decrease markedly as more GSes are deployed. [OSE2019] [BEL2021] [MAE20211]

LDACS navigation has already been demonstrated in practice in two flight measurement campaigns [SHU2013] [BEL2021] [MAE20211].

6. Requirements

The requirements for LDACS are mostly defined by its application area: Communications related to safety and regularity of flight.

A particularity of the current aeronautical communication landscape is that it is heavily regulated. Aeronautical data links (for applications related to safety and regularity of flight) may only use spectrum licensed to aviation and data links endorsed by ICAO. Nation states can change this locally; however, due to the global scale of the air transportation system, adherence to these practices is to be expected.

Aeronautical data links for the ATN are therefore expected to remain in service for decades. The VDLm2 data link currently used for digital terrestrial internetworking was developed in the 1990s (the use of the Open Systems Interconnection (OSI) stack indicates that as well). VDLm2 is expected to be used at least for several decades to come. In this respect aeronautical communications (for applications related to safety and regularity of flight) is more comparable to industrial applications than to the open Internet.

Internetwork technology is already installed in current aircraft. Current ATS applications use either Aircraft Communications Addressing and Reporting System (ACARS) or the OSI stack. The objective of the development effort of LDACS, as part of the FCI, is to replace legacy OSI stack and proprietary ACARS internetwork technologies with industry standard IP technology. It is anticipated that the use of Commercial Off-The-Shelf (COTS) IP technology mostly applies to the ground network. The avionics networks on the aircraft will likely be heavily modified versions of Ethernet or proprietary.

AOC applications currently mostly use the same stack (although some applications, like the graphical weather service may use the commercial passenger network). This creates capacity problems (resulting in excessive amounts of timeouts) since the underlying terrestrial data links do not provide sufficient bandwidth (i.e., with VDLm2 currently in the order of 10 kbit/s). The use of non-

aviation specific data links is considered a security problem. Ideally the aeronautical IP internetwork, hence the ATN over which only communications related to safety and regularity of flight is handled, and the Internet should be completely separated at Layer 3.

The objective of LDACS is to provide a next generation terrestrial data link designed to support IP addressing and provide much higher bandwidth to avoid the currently experienced operational problems.

The requirement for LDACS is therefore to provide a terrestrial high-throughput data link for IP internetworking in the aircraft.

In order to fulfil the above requirement LDACS needs to be interoperable with IP (and IP-based services like Voice-over-IP) at the gateway connecting the LDACS network to other aeronautical ground networks (i.e., the ATN). On the avionics side, in the aircraft, aviation specific solutions are to be expected.

In addition to these functional requirements, LDACS and its IP stack need to fulfil the requirements defined in RTCA DO-350A/EUROCAE ED-228A [DO350A]. This document defines continuity, availability, and integrity requirements at different scopes for each air traffic management application (CPDLC, CM, and ADS-C). The scope most relevant to IP over LDACS is the Communications Service Provider (CSP) scope.

Continuity, availability, and integrity requirements are defined in [DO350A] volume 1 Table 5-14, and Table 6-13. Appendix A presents the required information.

In a similar vein, requirements to fault management are defined in the same tables.

7. Characteristics

LDACS will become one of several wireless access networks connecting aircraft to the ATN implemented by the FCI.

The current LDACS design is focused on the specification of layer one and two. However, for the purpose of this work, only layer two details are discussed here.

Achieving the stringent continuity, availability, and integrity requirements defined in [DO350A] will require the specification of layer 3 and above mechanisms (e.g., reliable crossover at the IP layer). Fault management mechanisms are similarly unspecified as of November 2022. Current regulatory documents do not fully specify the above mechanism yet. However, a short overview of the current state shall be given throughout each section here.

7.1. LDACS Access Network

An LDACS access network contains an Access Router (AC-R) and several GS, each of them providing one LDACS radio cell.

User plane interconnection to the ATN is facilitated by the AC-R peering with an A/G Router connected to the ATN.

The internal control plane of an LDACS access network interconnects the GSs. An LDACS access network is illustrated in Figure 1.

```
wireless                user
link                    plane
  AS-------------GS---------------AC-R---A/G-----ATN
    .............                 |   Router
       control   .                |
       plane     .                |
                 .                |
                GS----------------|
                 .                |
                 .                |
                GS----------------+
```

Figure 1: LDACS access network with three GSs and one AS. dashes denotes the user plane and points the control plane

7.2. Topology

LDACS is a cellular point-to-multipoint system. It assumes a star-topology in each cell where Aircraft Stations (AS) belonging to aircraft within a certain volume of space (the LDACS cell) are connected to the controlling GS. The LDACS GS is a centralized instance that controls LDACS A/G communications within its cell. The LDACS GS can simultaneously support multiple bidirectional communications to the ASs under its control. LDACS's GSs themselves

are connected to each other and the AC-R.

Prior to utilizing the system an aircraft has to register with the controlling GS to establish dedicated logical channels for user and control data. Control channels have statically allocated resources, while user channels have dynamically assigned resources according to the current demand. Logical channels exist only between the GS and the AS.

7.3. LDACS Protocol Stack

The protocol stack of LDACS is implemented in the AS and GS: It consists of the Physical Layer (PHY) with five major, functional blocks above it. Four are placed in the Data Link Layer (DLL) of the AS and GS: (1) Medium Access Control (MAC) Layer, (2) Voice Interface (VI), (3) Data Link Service (DLS), and (4) LDACS Management Entity (LME). The fifth entity resides within the sub-network layer: (5) the Sub-Network Protocol (SNP). The LDACS radio is externally connected to a voice unit, radio control unit, and via the AC-R to the ATN network.

LDACS is considered an ATN/IPS radio access technology, from the view of ICAO's regulatory framework. Hence, the interface between ATN and LDACS must be IPv6 based, as regulatory documents, such as ICAO Doc 9896 [ICAO2015] and DO-379 [RTCA2019] clearly foresee that. The translation between IPv6 layer and SNP layer is currently the subject of ongoing standardization efforts and at the time of writing not finished yet.

Figure 2 shows the protocol stack of LDACS as implemented in the AS and GS. Acronyms used here are introduced throughout the upcoming sections.

```
                      IPv6                   Network Layer
                        |
   Airborne Voice       |
   Interface (AVI)/     |               Radio Control Unit (RCU)
   Voice Unit (VU)      |
      |                 |
      |      +------------------+  +----+
      |      |        SNP       |--|    |   Sub-Network
      |      |                  |  |    |   Layer
      |      +------------------+  |    |
      |                 |            LME|
   +-----+   +------------------+  |    |
   | VI  |   |        DLS       |  |    |   LLC Layer
   +-----+   +------------------+  +----+
      |                 |             |
     DCH               DCH         DCCH/CCCH
                        |          RACH/BCCH
                        |             |
   +-------------------------------------+
   |                  MAC                |   Medium Access
   |                                     |   Layer
   +-------------------------------------+
                        |
   +-------------------------------------+
   |                  PHY                |   Physical Layer
   +-------------------------------------+
                        |
                        |
                      ((*))
                      FL/RL              radio channels
                                        separated by FDD
```

Figure 2: LDACS protocol stack in AS and GS

7.3.1. LDACS Physical Layer

The physical layer provides the means to transfer data over the radio channel. The LDACS GS supports bidirectional links to multiple aircraft under its control. The FL direction at the G2A connection and the RL direction at the A2G connection are separated by Frequency Division Duplex (FDD). FL and RL use a 500 kHz channel each. The GS transmits a continuous stream of Orthogonal Frequency-Division Multiplexing Access (OFDM) symbols on the FL. In the RL different aircraft are separated in time and frequency using Orthogonal Frequency-Division Multiple Access (OFDMA). Aircraft thus transmit discontinuously on the RL via short radio bursts sent in precisely

defined transmission opportunities allocated by the GS.

7.3.2. LDACS Data Link Layer

The data-link layer provides the necessary protocols to facilitate concurrent and reliable data transfer for multiple users. The LDACS data link layer is organized in two sub-layers: The medium access sub-layer and the Logical Link Control (LLC) sub-layer. The medium access sub-layer manages the organization of transmission opportunities in slots of time and frequency. The LLC sub-layer provides acknowledged point-to-point logical channels between the aircraft and the GS using an Automatic Repeat reQuest (ARQ) protocol. LDACS also supports unacknowledged point-to-point channels and G2A Broadcast transmission.

7.3.2.1. Medium Access Control (MAC) Services

The MAC time framing service provides the frame structure necessary to realize slot-based time-division multiplex-access on the physical link. It provides the functions for the synchronization of the MAC framing structure and the PHY Layer framing. The MAC time framing provides a dedicated time slot for each logical channel.

The MAC sub-layer offers access to the physical channel to its service users. Channel access is provided through transparent logical channels. The MAC sub-layer maps logical channels onto the appropriate slots and manages the access to these channels. Logical channels are used as interface between the MAC and LLC sub-layers.

7.3.2.2. Data Link Service (DLS) Services

The DLS provides acknowledged and unacknowledged (including broadcast and packet mode voice) bidirectional exchange of user data. If user data is transmitted using the acknowledged DLS, the sending DLS entity will wait for an acknowledgement from the receiver. If no acknowledgement is received within a specified time frame, the sender may automatically try to retransmit its data. However, after a certain number of failed retries, the sender will suspend further retransmission attempts and inform its client of the failure.

The DLS uses the logical channels provided by the MAC:

1. A GS announces its existence and access parameters in the Broadcast Channel (BCCH).
2. The Random-Access Channel (RACH) enables AS to request access to an LDACS cell.
3. In the FL the Common Control Channel (CCCH) is used by the GS to grant access to data channel resources.

4. The reverse direction is covered by the RL, where ASs need to request resources before sending. This happens via the Dedicated Control Channel (DCCH).
5. User data itself is communicated in the Data Channel (DCH) on the FL and RL.

Access to the FL and RL data channel is granted by the scheduling mechanism implemented in the LME discussed below.

7.3.2.3. Voice Interface (VI) Services

The VI provides support for virtual voice circuits. Voice circuits may either be set-up permanently by the GS (e.g., to emulate voice party line) or may be created on demand.

7.3.2.4. LDACS Management Entity (LME) Services

The mobility management service in the LME provides support for registration and de-registration (cell entry and cell exit), scanning RF channels of neighboring cells and handover between cells. In addition, it manages the addressing of aircraft within cells.

The resource management service provides link maintenance (power, frequency and time adjustments), support for adaptive coding and modulation, and resource allocation.

The resource management service accepts resource requests from/for different AS and issues resource allocations accordingly. While the scheduling algorithm is not specified and a point of possible vendor differentiation, it is subject to the following requirements:

1. Resource scheduling must provide channel access according to the priority of the request
2. Resource scheduling must support "one-time" requests.
3. Resource scheduling must support "permanent" requests that reserve a resource until the request is canceled e.g. for digital voice circuits.

7.3.3. LDACS Sub-Network Layer and Protocol Services

Lastly, the SNP layer of LDACS directly interacts with IPv6 traffic. Incoming ATN/IPS IPv6 packets are forwarded over LDACS from and to the aircraft. The final IP addressing structure in an LDACS subnet still needs to be defined; however, the current layout is considered to consist of the five network segments: Air Core Net, Air Management Net, Ground Core Net, Ground Management Net, Ground Net. Any protocols that the ATN/IPS [ICAO2015] defines as mandatory will reach the aircraft, however listing these here is out of scope. For more

information on the technicalities of the above ATN/IPS layer, please refer to [ICAO2015] [RTCA2019] [ARI2021].

The DLS provides functions required for the transfer of user plane data and control plane data over the LDACS access network. The security service provides functions for secure user data communication over the LDACS access network. Note that the SNP security service applies cryptographic measures as configured by the GS.

7.4. LDACS Mobility

LDACS supports layer 2 handovers to different LDACS cells. Handovers may be initiated by the aircraft (break-before-make) or by the GS (make-before-break). Make-before-break handovers are only supported between GSs connected to each other, usually GS operated by the same service provider.

When a handover between AS and two interconnected GS takes place, it can be triggered by AS or GS. Once that is done, new security information is exchanged between AS, GS1 and GS2, before the "old" connection is terminated between AS and GS1 and a "new" connection is set up between AS and GS2. As a last step, accumulated user-data at GS1 is forwarded to GS2 via a ground connection, before that is sent via GS2 to the AS. While some information for handover is transmitted in the LDACS DCH, the information remains in the "control-plane" part of LDACS and is exchanged between LMEs in AS, GS1 and GS2. As such, local mobility takes place entirely within the LDACS network, utilizing the PMIPv6 protocol [RFC5213]. The use of PMIPv6 is currently not mandated by standardization, and may be vendor-specific. External handovers between non-connected LDACS access networks or different aeronautical data links are handled by the FCI multi-link concept.

External handovers between non-connected LDACS access networks or different aeronautical data links are handled by the FCI multi- link concept.

7.5. LDACS Management - Interfaces and Protocols

LDACS management interfaces and protocols are currently not be mandated by standardization. The implementations currently available use SNMP for management and Radius for AAA. Link state (link up, link down) is reported using the ATN/IPS Aircraft Protocol (AIAP) mandated by ICAO WG-I for multi-link.

8. Reliability and Availability

8.1. Below Layer 1

Below Layer 2, aeronautics usually relies on hardware redundancy. To protect availability of the LDACS link, an aircraft equipped with LDACS will have access to two L-band antennae with triple redundant radio systems as required for any safety relevant aeronautical systems by ICAO.

8.2. Layer 1 and 2

LDACS has been designed with applications related to the safety and regularity of flight in mind. It has therefore been designed as a deterministic wireless data link (as far as this is possible).

Based on channel measurements of the L-band channel LDACS was designed from the PHY layer up with robustness in mind. Channel measurements of the L-band channel [SCH2016] confirmed LDACS to be well adapted to its channel.

In order to maximize the capacity per channel and to optimally use the available spectrum, LDACS was designed as an OFDM-based FDD system, supporting simultaneous transmissions in FL in the G2A connection and RL in the A2G connection. The legacy systems already deployed in the L-band limit the bandwidth of both channels to approximately 500 kHz.

The LDACS physical layer design includes propagation guard times sufficient for operation at a maximum distance of 200 nautical miles from the GS. In actual deployment, LDACS can be configured for any range up to this maximum range.

The LDACS physical layer supports adaptive coding and modulation for user data. Control data is always encoded with the most robust coding and modulation (FL: Quadrature Phase-Shift Keying (QPSK), coding rate 1/2, RL: QPSK, coding rate 1/3).

LDACS medium access layer on top of the physical layer uses a static frame structure to support deterministic timer management. As shown in Figure 3 and Figure 4, LDACS framing structure is based on Super-Frames (SF) of 240ms duration corresponding to 2000 OFDM symbols. OFDM symbol time is 120 microseconds, sampling time 1.6 microseconds and a guard time of 4.8 microseconds. The structure of a SF is depicted in Figure 3 along with its structure and timings of each part. FL and RL boundaries are aligned in time (from the GS perspective) allowing for deterministic slots for control and data channels. This initial AS time synchronization and time synchronization maintenance is based on observing the synchronization symbol pairs that repetitively occur within the FL stream, being sent

by the controlling GS [GRA2020]. As already mentioned, LDACS data transmission is split into user-data (DCH) and control (BCCH, CCCH in FL; RACH, DCCH in RL) as depicted with corresponding timings in Figure 4.

```
^
|      +--------+------------+------------+------------+------------+
|   FL | BCCH   |     MF     |     MF     |     MF     |     MF     |
|      | 6.72ms |   58.32ms  |   58.32ms  |   58.32ms  |   58.32ms  |
F      +--------+------------+------------+------------+------------+
r      <----------------- Super-Frame (SF) - 240ms ----------------->
e
q      +--------+------------+------------+------------+------------+
u   RL | RACH   |     MF     |     MF     |     MF     |     MF     |
e      | 6.72ms |   58.32ms  |   58.32ms  |   58.32ms  |   58.32ms  |
n      +--------+------------+------------+------------+------------+
c      <----------------- Super-Frame (SF) - 240ms ----------------->
y
------------------------------ Time ------------------------------->
|
```

Figure 3: SF structure for LDACS

```
^
|      +-------------+----------------+-----------------+
|   FL |     DCH     |     CCCH       |      DCH        |
|      |   25.92ms   | 2.16 - 17.28ms | 15.12 - 30.24ms |
F      +-------------+----------------+-----------------+
r      <----------- Multi-Frame (MF) - 58.32ms --------->
e
q      +--------------+---------------------------------+
u   RL |    DCCH      |                DCH              |
e      | 2.8 - 24.4ms |           33.84 - 55.44ms       |
n      +--------------+---------------------------------+
c      <----------- Multi-Frame (MF) - 58.32ms --------->
y
------------------------------ Time -------------------->
|
```

Figure 4: MF structure for LDACS

LDACS cell entry is conducted with an initial control message exchange via the RACH and the BCCH.

After cell entry, LDACS medium access is always under the control of the GS of a radio cell. Any medium access for the transmission of user data on a DCH has to be requested with a resource request message stating the requested amount of resources and class of service. The GS performs resource scheduling on the basis of these requests and grants resources with resource allocation messages. Resource request and allocation messages are exchanged over dedicated contention-free control channels (DCCH and CCCH).

The purpose of quality-of-service in LDACS medium access is to provide prioritized medium access at the bottleneck (the wireless link). Signaling of higher layer quality-of-service requests to LDACS is implemented on the basis of Differentiated Services-(DiffServ) classes CS01 (lowest priority) to CS07 (highest priority).

In addition to having full control over resource scheduling, the GS can send forced handover commands for off-loading or channel management, e.g., when the signal quality declines and a more suitable GS is in the AS's reach. With robust resource management of the capacities of the radio channel, reliability and robustness measures are therefore also anchored in the LME.

In addition to radio resource management, the LDACS control channels are also used to send keep-alive messages, when they are not otherwise used. Since the framing of the control channels is deterministic, missing keep-alive messages can thus be immediately detected. This information is made available to the multilink protocols for fault management.

The protocol used to communicate faults is not defined in the LDACS specification. It is assumed that vendors would use industry standard protocols like the Simple Network Management Protocol or the Network Configuration Protocol, where security permits.

The LDACS data link layer protocol, running on top of the medium access sub-layer, uses ARQ to provide reliable data transmission on the data channel.

It employs selective repeat ARQ with transparent fragmentation and reassembly to the resource allocation size to minimize latency and overhead without losing reliability. It ensures correct order of packet delivery without duplicates. In case of transmission errors, it identifies lost fragments with deterministic timers synced to the medium access frame structure and initiates retransmission.

8.3. Beyond Layer 2

LDACS availability can be increased by appropriately deploying LDACS infrastructure: This means proliferating the number of terrestrial ground stations. However, there are four aspects that need to be taken into consideration: (1) scarcity of aeronautical spectrum for data link communication (in the case of LDACS: tens of MHz in the L-band), (2) an increase in the amount of ground stations also increases the individual bandwidth for aircraft in the cell, as fewer aircraft have to share the spectrum, (3) to cover worldwide terrestrial ATM via LDACS is also a question of cost and the possible reuse of spectrum which makes it not always possible to decrease cell sizes and (4) the Distance Measuring Equipment (DME) is the primary user of the aeronautical L-band, which means any LDACS deployment has to take DME frequency planning into account.

While aspect (2) provides a good reason, alongside increasing redundancy, for smaller cells than the maximum range LDACS was developed for (200 Nautical Miles (NM)), the other three need to be respected when doing so. There are preliminary works on LDACS cell planning, such as [MOST2018], where the authors reach the conclusion that 84 LDACS cells in Europe would be sufficient to serve European air traffic for the next 20 years.

For redundancy reasons, the aeronautical community has decided not to rely on a single communication system or frequency band. It is envisioned to have multiple independent data link technologies in the aircraft (e.g., terrestrial and satellite communications) in addition to legacy VHF voice.

However, as of now, no reliability and availability mechanisms that could utilize the multilink architecture, have been specified on Layer 3 and above. Even if LDACS has been designed for reliability, the wireless medium presents significant challenges to achieve deterministic properties such as low packet error rate, bounded consecutive losses, and bounded latency. Support for high reliability and availability for IP connectivity over LDACS is certainly highly desirable but needs to be adapted to the specific use case.

9. Security

The goal of this Section is to inform the reader about the state of security in aeronautical communications, state security considerations applicable for all ATN/IPS traffic and to provide an overview of the LDACS link-layer security capabilities.

9.1. Security in Wireless Digital Aeronautical Communications

Aviation will require secure exchanges of data and voice messages for managing the air traffic flow safely through the airspaces all over the world. Historically Communication Navigation Surveillance (CNS) wireless communications technology emerged from the military and a threat landscape where inferior technological and financial capabilities of adversaries were assumed [STR2016]. The main communications method for ATC today is still an open analogue voice broadcast within the aeronautical VHF band. Currently, information security is mainly procedural, based by using well-trained personnel and proven communications procedures. This communication method has been in service since 1948. However, since the emergence of civil aeronautical CNS applications in the 70s, and today, the world has changed.

CCivil applications have significant lower spectrum available than military applications. This means several military defense mechanisms, such as frequency hopping or pilot symbol scrambling and, thus, a defense-in-depth approach starting at the physical layer, is infeasible for civil systems. With the rise of cheap Software Defined Radios (SDR), the previously existing financial barrier is almost gone and open source projects such as GNU radio [GNU2021] allow a new type of unsophisticated listeners and possible attackers.

Most CNS technology developed in ICAO relies on open standards, thus syntax and semantics of wireless digital aeronautical communications should be expected to be common knowledge for attackers. With increased digitization and automation of civil aviation, the human as control instance, is being taken gradually out of the loop. Autonomous transport drones or single piloted aircraft demonstrate this trend. However, without profound cybersecurity measures such as authenticity and integrity checks of messages in-transit on the wireless link or mutual entity authentication, this lack of a control instance can prove disastrous. Thus, future digital communications will need additional embedded security features to fulfill modern information security requirements like authentication and integrity. These security features require sufficient bandwidth which is beyond the capabilities of currently deployed VHF narrowband communications systems. For voice and data communications, sufficient data throughput capability is needed to support the security functions while not degrading performance. LDACS is a data link technology with sufficient bandwidth to incorporate security without losing too much user data throughput.

9.2. Security in Depth

ICAO Doc 9896 foresees transport layer security [ICAO2015] for all aeronautical data transmitted via the ATN/IPS, as described in ARINC P858 [ARI2021]. This is realized via Datagram Transport Layer Security (DTLS) 1.3 [RFC9147].

LDACS also needs to comply with in-depth security requirements, stated in ARINC 858, for the radio access technologies transporting ATN/IPS data. These requirements imply that LDACS must provide layer 2 security in addition to any higher layer mechanisms. Specifically, ARINC 858 states that [datalinks within the FCI need to provide] "a secure channel between the airborne radio systems and the peer radio access endpoints on the ground [...] to ensure authentication and integrity of air-ground message exchanges in support of an overall defense-in-depth security strategy." [ARI2021]

9.3. LDACS Security Requirements

Overall, cybersecurity for CNS technology shall protect the following business goals [MAE20181]:

1. Safety: The system must sufficiently mitigate attacks, which contribute to safety hazards.
2. Flight regularity: The system must sufficiently mitigate attacks, which contribute to delays, diversions, or cancellations of flights.
3. Protection of business interests: The system must sufficiently mitigate attacks which result in financial loss, reputation damage, disclosure of sensitive proprietary information, or disclosure of personal information.

To further analyze assets and derive threats and thus protection scenarios several threat-and risk analyses were performed for LDACS [MAE20181] , [MAE20191]. These results allowed the derivation of security scope and objectives from the requirements and the conducted threat-and risk analysis. Note, IPv6 security considerations are briefly discussed in Section 9.7 while a summary of security requirements for link-layer candidates in the ATN/IPS is given in [ARI2021], which states: "Since the communication radios connect to local airborne networks in the aircraft control domain, [...] the airborne radio systems represent the first point of entry for an external threat to the aircraft. Consequently, a secure channel between the airborne radio systems and the peer radio access endpoints on the ground is necessary to ensure authentication and integrity of air-ground message exchanges in support of an overall defense-in-depth security strategy".

9.4. LDACS Security Objectives

Security considerations for LDACS are defined by the official SARPS document by ICAO [ICAO2022]:

1. LDACS shall provide a capability to protect the availability and continuity of the system.
2. LDACS shall provide a capability including cryptographic mechanisms to protect the integrity of messages in transit.
3. LDACS shall provide a capability to ensure the authenticity of messages in transit.
4. LDACS should provide a capability for nonrepudiation of origin for messages in transit.
5. LDACS should provide a capability to protect the confidentiality of messages in transit.
6. LDACS shall provide an authentication capability.
7. LDACS shall provide a capability to authorize the permitted actions of users of the system and to deny actions that are not explicitly authorized.
8. If LDACS provides interfaces to multiple domains, LDACS shall provide capability to prevent the propagation of intrusions within LDACS domains and towards external domains.

Work in 2022 includes a change request for these SARPS aims to limit the "non-repudiation of origin of messages in transit" requirement only to the authentication and key establishment messages at the beginning of every session.

9.5. LDACS Security Functions

These objectives were used to derive several security functions for LDACS required to be integrated in the LDACS cybersecurity architecture: Identification, Authentication, Authorization, Confidentiality, System Integrity, Data Integrity, Robustness, Reliability, Availability, and Key and Trust Management. Several works investigated possible measures to implement these security functions [BIL2017], [MAE20181], [MAE20191].

9.6. LDACS Security Architecture

The requirements lead to a LDACS security model, including different entities for identification, authentication and authorization purposes, ensuring integrity, authenticity and confidentiality of data. A draft of the cybersecurity architecture of LDACS can be found in [ICAO2022] and [MAE20182] and respective updates in [MAE20191], [MAE20192], [MAE2020], [MAE2021].

9.6.1. Entities

A simplified LDACS architectural model requires the following entities: Network operators such as the Societe Internationale de Telecommunications Aeronautiques (SITA) [SIT2020] and ARINC [ARI2020] are providing access to the ground IPS network via an A/G LDACS router. This router is attached to an internal LDACS access network, which connects via further access routers to the different LDACS cell ranges, each controlled by a GS (serving one LDACS cell), with several interconnected GS spanning a local LDACS access network. Via the A/G wireless LDACS data link AS the aircraft is connected to the ground network and via the aircraft's VI and aircraft's network interface, aircraft's data can be sent via the AS back to the GS, then to the LDACS local access network, access routers, LDACS access network, A/G LDACS router and finally to the ground IPS network [ICAO2015].

9.6.2. Entity Identification

LDACS needs specific identities for the AS, the GS, and the network operator. The aircraft itself can be identified using the 24-bit ICAO identifier of an aircraft [ICAO2022], the call sign of that aircraft or the recently founded privacy ICAO address of the Federal Aviation Administration (FAA) program with the same name [FAA2020]. It is conceivable that the LDACS AS will use a combination of aircraft identification, radio component identification and even operator feature identification to create a unique AS LDACS identification tag. Similar to a 4G's eNodeB serving network identification tag, a GS could be identified using a similar field. The identification of the network operator is again similar to 4G (e.g., E-Plus, AT&T, and TELUS), in the way that the aeronautical network operators are listed (e.g., ARINC [ARI2020] and SITA [SIT2020]).

9.6.3. Entity Authentication and Key Establishment

In order to anchor trust within the system, all LDACS entities connected to the ground IPS network will be rooted in an LDACS specific chain-of-trust and PKI solution, quite similar to AeroMACS's approach [CRO2016]. These certificates, residing at the entities and incorporated in the LDACS PKI, providing proof the ownership of their respective public key, include information about the identity of the owner and the digital signature of the entity that has verified the certificate's content. First, all ground infrastructures must mutually authenticate to each other, negotiate and derive keys and, thus, secure all ground connections. How this process is handled in detail is still an ongoing discussion. However, established methods to secure user plane by IPSec [RFC4301] and IKEv2 [RFC7296] or the

application layer via TLS 1.3 [RFC8446] are conceivable. The LDACS PKI with its chain-of-trust approach, digital certificates and public entity keys lay the groundwork for this step. In a second step, the AS with the LDACS radio aboard, approaches an LDACS cell and performs a cell-attachment procedure with the corresponding GS. This procedure consists of (1) the basic cell entry [GRA2020] and (2) a Mutual Authentication and Key Establishment (MAKE) procedure [MAE2021].

Note that LDACS will foresee multiple security levels. To address the issue of the long service life of LDACS (i.e., possibly >30 years) and the security of current pre-quantum cryptography, these security levels include pre- and post-quantum cryptographic solutions. Limiting security data on the LDACS datalink as much as possible, to reserve as much space for actual user data transmission, is key in the LDACS security architecture, this is also reflected in the underlying cryptography: Pre-quantum solutions will rely on elliptic curves [NIST2013], while post-quantum solutions consider Falcon [SON2021] [MAE2021] or similar lightweight PQC signature schemes, and CRYSTALS-KYBER or SABER as key establishment options [AVA2021] [ROY2020].

9.6.4. Message-in-transit Confidentiality, Integrity and Authenticity

The key material from the previous step can then be used to protect LDACS Layer 2 communications via applying encryption and integrity protection measures on the SNP layer of the LDACS protocol stack. As LDACS transports AOC and ATS data, the integrity of that data is most important, while confidentiality only needs to be applied to AOC data to protect business interests [ICAO2022]. This possibility of providing low layered confidentiality and integrity protection ensures a secure delivery of user data over the wireless link. Furthermore, it ensures integrity protection of LDACS control data.

9.7. Considerations on LDACS Security Impact on IPv6 Operational Security

In this part, considerations on IPv6 operational security in [RFC9099] and interrelations with the LDACS security additions are compared and evaluated to identify further protection demands. As IPv6 heavily relies on the Neighbor Discovery Protocol (NDP) [RFC4861], integrity and authenticity protection on the link-layer, as provided by LDACS, already help mitigate spoofing and redirection attacks. However, to also mitigate the threat of remote DDoS attacks, neighbor solicitation rate-limiting is recommended by RFC 9099. To prevent the threat of (D)DoS attacks in general on the LDACS access network, rate-limiting needs to be performed on each network node in the LDACS access network. One approach is to filter

for the total amount of possible LDACS AS-GS traffic per cell - i.e., of up to 1.4 Mbps user-data per cell and up to the amount of GS per service provider network times 1.4 Mbps.

10. IANA Considerations

This memo includes no request to IANA.

11. Acknowledgements

Thanks to all contributors to the development of LDACS and ICAO PT-T, as well as to all in the RAW Working Group for deep discussions and feedback.

Thanks to Klaus-Peter Hauf, Bart Van Den Einden, and Pierluigi Fantappie for their comments to this draft.

Thanks to the Chair for Network Security and the research institute CODE for their comments and improvements.

Thanks to the colleagues of the Research Institute CODE at the UniBwM working in the AMIUS project funded under the Bavarian Aerospace Program by the Bavarian State Ministry of Economics, Regional Development and Energy with the GA ROB-2-3410.20-04-11-15/HAMI-2109-0015 for fruitful discussions on aeronautical communications and relevant security incentives for the target market.

Thanks to SBA Research Vienna for continuous discussions on security infrastructure issues in quick developing markets such the air space and potential economic spillovers to used technologies and protocols.

Thanks to the Aeronautical Communications group at the Institute of Communications and Navigation of the German Aerospace Center (DLR). With that, the authors would like to explicitly thank Miguel Angel Bellido-Manganell and Lukas Marcel Schalk for their thorough feedback.

12. Normative References

13. Informative References

[RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, DOI 10.17487/RFC4301, December 2005, <https://www.rfc-editor.org/info/rfc4301>.

[RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <https://www.rfc-editor.org/info/rfc4861>.

[RFC5213] Gundavelli, S., Ed., Leung, K., Devarapalli, V., Chowdhury, K., and B. Patil, "Proxy Mobile IPv6", RFC 5213, DOI 10.17487/RFC5213, August 2008, <https://www.rfc-editor.org/info/rfc5213>.

[RFC5795] Sandlund, K., Pelletier, G., and L. Jonsson, "The RObust Header Compression (ROHC) Framework", RFC 5795, DOI 10.17487/RFC5795, March 2010, <https://www.rfc-editor.org/info/rfc5795>.

[RFC7296] Kaufman, C., Hoffman, P., Nir, Y., Eronen, P., and T. Kivinen, "Internet Key Exchange Protocol Version 2 (IKEv2)", STD 79, RFC 7296, DOI 10.17487/RFC7296, October 2014, <https://www.rfc-editor.org/info/rfc7296>.

[RFC8200] 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>.

[RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <https://www.rfc-editor.org/info/rfc8446>.

[RFC9099] Vyncke, É., Chittimaneni, K., Kaeo, M., and E. Rey, "Operational Security Considerations for IPv6 Networks", RFC 9099, DOI 10.17487/RFC9099, August 2021, <https://www.rfc-editor.org/info/rfc9099>.

[RFC9147] Rescorla, E., Tschofenig, H., and N. Modadugu, "The Datagram Transport Layer Security (DTLS) Protocol Version 1.3", RFC 9147, DOI 10.17487/RFC9147, April 2022, <https://www.rfc-editor.org/info/rfc9147>.

[GRA2020] Gräupl, T., Rihacek, C., and B. Haindl, "LDACS A/G Specification", SESAR2020 PJ14-02-01 D3.3.030 , 2020, <https://www.ldacs.com/wp-content/uploads/2013/12/SESAR2020_PJ14-W2-60_D3_1_210_Initial_LDACS_AG_Specification_00_01_00-1_0_updated.pdf>.

[ARI2021] ARINC, "Internet Protocol Suite (IPS) For Aeronautical Safety Services Part 1- Airborne IP System Technical Requirements, ARINC SPECIFICATION 858 P1", June 2021, <https://standards.globalspec.com/std/14391274/858p1>.

[EURO2019] European Organization for Civil Aviation Equipment (EUROCAE), "Technical Standard of Aviation Profiles for ATN/IPS, ED-262", September 2019, <https://eshop.eurocae.net/eurocae-documents-and-reports/ed-262/>.

[ICAO2015] International Civil Aviation Organization (ICAO), "Manual on the Aeronautical Telecommunication Network (ATN) using Internet Protocol Suite (IPS) Standards and Protocols, Doc 9896", January 2015, <https://standards.globalspec.com/std/10026940/icao-9896>.

[RTCA2019] Radio Technical Commission for Aeronautics (RTCA), "Internet Protocol Suite Profiles, DO-379", September 2019, <https://www.rtca.org/products/do-379/>.

[SCH2016] Schneckenburger, N., Jost, T., Shutin, D., Walter, M., Thiasiriphet, T., Schnell, M., and U.C. Fiebig, "Measurement of the L-band Air-to-Ground Channel for Positioning Applications", IEEE Transactions on Aerospace and Electronic Systems, 52(5), pp.2281-229 , 2016.

[OSE2019] Osechas, O., Narayanan, S., Crespillo, O.G., Zampieri, G., Battista, G., Kumar, R., Schneckenburger, N., Lay, E., Belabbas, B., and M. Meurer, "Feasibility Demonstration of Terrestrial RNP with LDACS", 32nd International Technical Meeting of the Satellite Division of The Institute of Navigation, pp.1-12 , 2019.

[SCHN2018] Schneckenburger, N., "A Wide-Band Air-Ground Channel Mode", Dissertation, Technische Universitaet Ilmenau, Ilmenau, Germany , 2018.

[MOST2018] Mostafa, M., Bellido-Manganell, M.A.., and T. Gräupl, "Feasibility of Cell Planning for the L-Band Digital Aeronautical Communications System Under the Constraint of Secondary Spectrum Usage", IEEE Transactions on Vehicular Technology, vol. 67, no. 10, pp. 9721-9733 , 2018.

[MAE20191] Mäurer, N., Gräupl, T., and C. Schmitt, "Evaluation of the LDACS Cybersecurity Implementation", IEEE 38th Digital Avionics Systems Conference (DACS), pp. 1-10, San Diego, CA, USA , 2019.

[MAE20192] Mäurer, N. and C. Schmitt, "Towards Successful Realization of the LDACS Cybersecurity Architecture: An Updated Datalink Security Threat- and Risk Analysis", IEEE Integrated Communications, Navigation and Surveillance Conference (ICNS), pp. 1-13, Herndon, VA, USA , 2019.

[MAE20182] Mäurer, N. and A. Bilzhause, "A Cybersecurity Architecture for the L-band Digital Aeronautical Communications System (LDACS)", IEEE 37th Digital Avionics Systems Conference (DASC), pp. 1-10, London, UK , 2017.

[GRA2011] Gräupl, T. and M. Ehammer, "L-DACS1 Data Link Layer Evolution of ATN/IPS", 30th IEEE/AIAA Digital Avionics Systems Conference (DASC), pp. 1-28, Seattle, WA, USA , 2011.

[GRA2018] Gräupl, T., Schneckenburger, N., Jost, T., Schnell, M., Filip, A., Bellido-Manganell, M.A., Mielke, D.M., Mäurer, N., Kumar, R., Osechas, O., and G. Battista, "L-band Digital Aeronautical Communications System (LDACS) flight trials in the national German project MICONAV", Integrated Communications, Navigation, Surveillance Conference (ICNS), pp. 1-7, Herndon, VA, USA , 2018.

[ICAO2022] International Civil Aviation Organization (ICAO), "L-Band Digital Aeronautical Communication System (LDACS", International Standards and Recommended Practices Annex 10 - Aeronautical Telecommunications, Vol. III - Communication Systems, 2022 , 2022.

[SAJ2014] Haindl, B., Meser, J., Sajatovic, M., Müller, S., Arthaber, H., Faseth, T., and M. Zaisberger, "LDACS1 Conformance and Compatibility Assessment", IEEE/AIAA 33rd Digital Avionics Systems Conference (DASC), pp. 1-11, Colorado Springs, CO, USA , 2014.

[RIH2018] Rihacek, C., Haindl, B., Fantappie, P., Pierattelli, S., Gräupl, T., Schnell, M., and N. Fistas, "L-band Digital Aeronautical Communications System (LDACS) Activities in SESAR2020", Integrated Communications Navigation and Surveillance Conference (ICNS), pp. 1-8, Herndon, VA, USA , 2018.

[BEL2019] Bellido-Manganell, M. A. and M. Schnell, "Towards Modern Air-to-Air Communications: the LDACS A2A Mode", IEEE/AIAA 38th Digital Avionics Systems Conference (DASC), pp. 1-10, San Diego, CA, USA , 2019.

[CRO2016] Crowe, B., "Proposed AeroMACS PKI Specification is a Model for Global and National Aeronautical PKI Deployments", WiMAX Forum at 16th Integrated Communications, Navigation and Surveillance Conference (ICNS), pp. 1-19, New York, NY, USA , 2016.

[MAE2020] Mäurer, N., Gräupl, T., and C. Schmitt, "Comparing Different Diffie-Hellman Key Exchange Flavors for LDACS", IEEE/AIAA 39th Digital Avionics Systems Conference (DASC), pp. 1-10, San Antonio, TX, USA , 2020.

[STR2016] Strohmeier, M., Schäfer, M., Pinheiro, R., Lenders, V., and I. Martinovic, "On Perception and Reality in Wireless Air Traffic Communication Security", IEEE Transactions on Intelligent Transportation Systems, 18(6), pp. 1338-1357, New York, NY, USA , 2016.

[BIL2017] Bilzhause, A., Belgacem, B., Mostafa, M., and T. Gräupl, "Datalink Security in the L-band Digital Aeronautical Communications System (LDACS) for Air Traffic Management", IEEE Aerospace and Electronic Systems Magazine, 32(11), pp. 22-33, New York, NY, USA , 2017.

[MAE20181] Mäurer, N. and A. Bilzhause, "Paving the Way for an IT Security Architecture for LDACS: A Datalink Security Threat and Risk Analysis", IEEE Integrated Communications, Navigation, Surveillance Conference (ICNS), pp. 1-11, New York, NY, USA , 2018.

[FAA2020] FAA, "Federal Aviation Administration. ADS-B Privacy.", August 2020, <https://www.faa.gov/nextgen/equipadsb/privacy/>.

[GNU2021] GNU Radio project, "GNU radio", October 2021, <http://gnuradio.org>.

[SIT2020] SITA, "Societe Internationale de Telecommunications Aeronautiques", August 2020, <https://www.sita.aero/>.

[ARI2020] ARINC, "Aeronautical Radio Incorporated", August 2020, <https://www.aviation-ia.com/>.

[DO350A] RTCA SC-214, "Safety and Performance Standard for Baseline 2 ATS Data Communications (Baseline 2 SPR Standard)", May 2016, <https://standards.globalspec.com/std/10003192/rtca-do-350-volume-1-2>.

[ICAO2019] International Civil Aviation Organization (ICAO), "Manual on VHF Digital Link (VDL) Mode 2, Doc 9776", January 2019, <https://store.icao.int/en/manual-on-vhf-digital-link-vdl-mode-2-doc-9776>.

[KAMA2010] Kamali, B., "An Overview of VHF Civil Radio Network and the Resolution of Spectrum Depletion", Integrated Communications, Navigation, and Surveillance Conference, pp. F4-1-F4-8 , May 2010.

[KAMA2018] Kamali, B., "AeroMACS: An IEEE 802.16 Standard-based Technology for the Next Generation of Air Transportation Systems", John Wiley and Sons, DOI: 10.1002/9781119281139 , September 2018.

[NIST2013] Barker, E., "Digital Signature Standard (DSS)", National Institute of Standards and Technology (NIST), FIPS.186-4, DOI: 10.6028/NIST.FIPS.186-4 , 2013.

[SON2021] Soni, D., Basu, K., Nabeel, M., Aaraj, N., Manzano, M., and R. Karri, "FALCON", Hardware Architectures for Post-Quantum Digital Signature Schemes, pp. 31-41 , November 2021.

[AVA2021] Avanzi, R., Bos, J., Ducas, L., Kiltz, E., Lepoint, T., Lyubashevsky, C., Schanck, J.M., Schwabe, P., Seiler, G., and D. Stehle, "CRYSTALS-KYBER - Algorithm Specification and Supporting Documentation (version 3.02)", August 2021, <https://pq-crystals.org/kyber/data/kyber-specification-round3-20210804.pdf>.

[ROY2020] Roy, S.S.. and A. Basso, "High-Speed Instruction-Set Coprocessor For Lattice-Based Key Encapsulation Mechanism: Saber In Hardware", IACR Transactions on Cryptographic Hardware and Embedded Systems, 443-466. , August 2020.

[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-06, 30 November 2022, <https://www.ietf.org/archive/id/draft-ietf-raw-technologies-06.txt>.

[RAW-USE-CASES]
Bernardos, C. J., Papadopoulos, G. Z., Thubert, P., and F. Theoleyre, "RAW Use-Cases", Work in Progress, Internet-Draft, draft-ietf-raw-use-cases-08, 22 October 2022, <https://www.ietf.org/archive/id/draft-ietf-raw-use-cases-08.txt>.

[I-D.haindl-lisp-gb-atn]
Haindl, B., Lindner, M., Moreno, V., Portoles-Comeras, M., Maino, F., and B. Venkatachalapathy, "Ground-Based LISP for the Aeronautical Telecommunications Network", Work in Progress, Internet-Draft, draft-haindl-lisp-gb-atn-08, 23 September 2022, <https://www.ietf.org/archive/id/draft-haindl-lisp-gb-atn-08.txt>.

[I-D.ietf-rtgwg-atn-bgp]
Templin, F., Saccone, G., Dawra, G., Lindem, A., and V. Moreno, "A Simple BGP-based Mobile Routing System for the Aeronautical Telecommunications Network", Work in Progress, Internet-Draft, draft-ietf-rtgwg-atn-bgp-19, 6 November 2022, <https://www.ietf.org/archive/id/draft-ietf-rtgwg-atn-bgp-19.txt>.

[I-D.ietf-lisp-rfc6830bis]
Farinacci, D., Fuller, V., Meyer, D., Lewis, D., and A. Cabellos-Aparicio, "The Locator/ID Separation Protocol (LISP)", Work in Progress, Internet-Draft, draft-ietf-lisp-rfc6830bis-38, 7 May 2022, <https://www.ietf.org/archive/id/draft-ietf-lisp-rfc6830bis-38.txt>.

[I-D.ietf-lisp-rfc6833bis]
Farinacci, D., Maino, F., Fuller, V., and A. Cabellos-Aparicio, "Locator/ID Separation Protocol (LISP) Control Plane", Work in Progress, Internet-Draft, draft-ietf-lisp-rfc6833bis-31, 2 May 2022, <https://www.ietf.org/archive/id/draft-ietf-lisp-rfc6833bis-31.txt>.

[ICAO2018] International Civil Aviation Organization (ICAO), "Handbook on Radio Frequency Spectrum Requirements for Civil Aviation, Doc 9718, Volume 1, ICAO Spectrum Strategy, Policy Statements and Related Information", July 2018, <https://www.icao.int/safety/FSMP/Documents/Doc9718/Doc9718_Vol_I_2nd_ed_(2018)corr1.pdf>.

[ARI2019] ARINC, "AOC Air-Ground Data And Message Exchange Format, ARINC 633", January 2019, <https://standards.globalspec.com/std/13152055/ARINC%20633>.

[VIR2021] Virdia, A., Stea, G., and G. Dini, "SAPIENT: Enabling Real-Time Monitoring and Control in the Future Communication Infrastructure of Air Traffic Management", IEEE Transactions on Intelligent Transportation Systems, 22(8):4864-4875 , August 2021.

[SHU2013] Shutin, D., Schneckenburger, N., Walter, M., and M. Schnell, "LDACS1 Ranging Performance - An Analysis Of Flight Measurement Results", IEEE 32nd Digital Avionics Systems Conference (DASC), pp. 1-10, East Syracuse, NY, USA , October 2013.

[BEL2021] Bellido-Manganell, M.A., Gräupl, T., Heirich, O., Mäurer, N., Filip-Dhaubhadel, A., Mielke, D.M., Schalk, L.M., Becker, D., Schneckenburger, N., and M. Schnell, "LDACS Flight Trials: Demonstration and Performance Analysis of the Future Aeronautical Communications System", IEEE Transactions on Aerospace and Electronic Systems, pp. 1-19 , September 2021.

[MAE2021] Mäurer, N., Gräupl, T., Gentsch, C., Guggemos, T., Tiepelt, M., Schmitt, C., and G. Dreo Rodosek, "A Secure Cell-Attachment Procedure for LDACS", 1st Workshop on Secure and Reliable Communication and Navigation in the Aerospace Domain (SRCNAS), pp. 1-10, Vienna, Austria , September 2021.

[MAE20211] Mäurer, N., Gräupl, T., Bellido-Manganell, M.A., Mielke, D.M., Filip-Dhaubhadel, A., Heirich, O., Gerberth, D., Flux, M., Schalk, L.M., Becker, D., Schneckenburger, N., and M. Schnell, "Flight Trial Demonstration of Secure GBAS via the L-band Digital Aeronautical Communications System (LDACS)", IEEE Aerospace and Electronic Systems Magazine, 36(4), pp. 8-17 , April 2021.

[BOE2019] Boegl, T., Rautenberg, M., Haindl, R., Rihacek, C., Meser, J., Fantappie, P., Pringvanich, N., Micallef, J., Klauspeter, H., MacBride, J., Sacre, P., v.d. Eiden, B., Gräupl, T., and M. Schnell, "LDACS White Paper - A Roll-out Scenario", International Civil Aviation Organization, Communications Panel - Data Communications Infrastructure Working Group - Third Meeting, pp. 1-8, Montreal, Canada , October 2019.

Appendix A. Selected Information from DO-350A

This appendix includes the continuity, availability, and integrity requirements applicable for LDACS defined in [DO350A].

The following terms are used here:

CPDLC Controller Pilot Data Link Communication
DT Delivery Time (nominal) value for RSP
ET Expiration Time value for RCP
FH Flight Hour
MA Monitoring and Alerting criteria
OT Overdue Delivery Time value for RSP
RCP Required Communication Performance
RSP Required Surveillance Performance
TT Transaction Time (nominal) value for RCP

	RCP 130	RCP 130
Parameter	ET	TT95%
Transaction Time (sec)	130	67
Continuity	0.999	0.95
Availability	0.989	0.989
Integrity	1E-5 per FH	1E-5 per FH

Table 1: CPDLC Requirements for RCP 130

	RCP 240	RCP 240	RCP 400	RCP 400
Parameter	ET	TT95%	ET	TT95%
Transaction Time (sec)	240	210	400	350
Continuity	0.999	0.95	0.999	0.95
Availability	0.989	0.989	0.989	0.989
Integrity	1E-5 per FH	1E-5 per FH	1E-5 per FH	1E-5 per FH

Table 2: CPDLC Requirements for RCP 240/400

RCP Monitoring and Alerting Criteria in case of CPDLC:

- MA-1: The system shall be capable of detecting failures and configuration changes that would cause the communication service no longer meet the RCP specification for the intended use.
- MA-2: When the communication service can no longer meet the RCP specification for the intended function, the flight crew and/or the controller shall take appropriate action.

	RSP 160	RSP 160	RSP 180	RSP 180	RSP 400	RSP 400
Parameter	OT	DT95%	OT	DT95%	OT	DT95%
Transaction Time (sec)	160	90	180	90	400	300
Continuity	0.999	0.95	0.999	0.95	0.999	0.95
Availability	0.989	0.989	0.989	0.989	0.989	0.989
Integrity	1E-5 per FH	1E-5 per FH	1E-5 per FH	1E-5 per FH	1E-5 per FH	1E-5 per FH

Table 3: ADS-C Requirements

RCP Monitoring and Alerting Criteria:

- MA-1: The system shall be capable of detecting failures and configuration changes that would cause the ADS-C service no longer meet the RSP specification for the intended function.
- MA-2: When the ADS-C service can no longer meet the RSP specification for the intended function, the flight crew and/or the controller shall take appropriate action.

Authors' Addresses

Nils Mäurer (editor)
German Aerospace Center (DLR)
Münchner Strasse 20
82234 Wessling
Germany
Email: Nils.Maeurer@dlr.de

Thomas Gräupl (editor)
German Aerospace Center (DLR)
Münchner Strasse 20
82234 Wessling
Germany
Email: Thomas.Graeupl@dlr.de

Corinna Schmitt (editor)
Research Institute CODE, UniBwM
Werner-Heisenberg-Weg 39
85577 Neubiberg
Germany
Email: corinna.schmitt@unibw.de

RAW
Internet-Draft
Intended status: Informational
Expires: 6 September 2023

F. Theoleyre
CNRS
G.Z. Papadopoulos
IMT Atlantique
G. Mirsky
Ericsson
CJ. Bernardos
UC3M
5 March 2023

Operations, Administration and Maintenance (OAM) features for RAW
draft-ietf-raw-oam-support-06

Abstract

Some critical applications may use a wireless infrastructure. However, wireless networks exhibit a bandwidth of several orders of magnitude lower than wired networks. Besides, wireless transmissions are lossy by nature; the probability that a packet cannot be decoded correctly by the receiver may be quite high. In these conditions, providing high reliability and a low delay is challenging. This document lists the requirements of the Operation, Administration, and Maintenance (OAM) features are recommended to provide availability and reliability on top of a collection of wireless segments. This document describes the benefits, problems, and trade-offs for using OAM in wireless networks to achieve Service Level Objectives (SLO).

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

The Reliable and Available Wireless (RAW) working group aims to extends DetNet to approach end-to-end deterministic performances over a network that includes scheduled wireless segments. In wired networks, many approaches try to enable Quality of Service (QoS) by implementing traffic differentiation so that routers handle each type of packets differently.

Deterministic Networking (DetNet) [RFC8655] has proposed to provide a bounded end-to-end latency on top of the network infrastructure, comprising both Layer 2 bridged and Layer 3 routed segments. Their work encompasses the data plane, OAM, time synchronization, management, control, and security aspects.

However, wireless networks create specific challenges. First of all, radio bandwidth is significantly lower than in wired networks. In these conditions, the volume of signaling messages has to be very limited. Even worse, wireless links are lossy: a Layer 2 transmission may or may not be decoded correctly by the receiver, depending on a broad set of parameters. Thus, providing high reliability through wireless segments is particularly challenging.

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

Last but not least, radio links present volatile characteristics. If the wireless networks use an unlicensed band, packet losses are not anymore temporally and spatially independent. Typically, links may exhibit a very bursty characteristic, where several consecutive packets may be dropped because of, e.g., temporary external interference. Thus, providing availability and reliability on top of the wireless infrastructure requires specific Layer 3 mechanisms to counteract these bursty losses. Besides, Layer 3 has to be _informed_ of the physical characteristics to make the right decision, and to avoid exacerbating physical issues (e.g., overloaded link because it became unreliable, overloaded radio channels).

Operations, Administration, and Maintenance (OAM) Tools are of primary importance for IP networks [RFC7276]. They define a toolset for fault detection, isolation, and performance measurement.

The primary purpose of this document is to detail the specific requirements of the OAM features recommended to provide reliability and availability on top of a collection of wireless segments. This document describes the benefits, problems, and trade-offs for using OAM in wireless networks to provide these properties.

1.1. Terminology

In this document, the term OAM will be used according to its definition specified in [RFC6291]. We expect to implement an OAM framework in RAW networks to maintain a real-time view of the network infrastructure, and its ability to respect the Service Level Objectives (SLO), such as delay and reliability, assigned to each data flow.

We re-use here the same terminology as [I-D.ietf-detnet-oam-framework]:

* OAM entity: a data flow to be monitored for defects and/or its performance metrics measured. For such entity, we define the following terms:

  - OAM domain: a network used by the monitored flow. An OAM domain may have MEPs on its edge and MIPs within.

  - Maintenance End Point (MEP): an OAM instance that is capable of generating OAM test packets in the particular sub-layer of the OAM domain.

  - Maintenance Intermediate endPoint (MIP): an OAM instance along the flow in the particular sub-layer of the OAM domain. A MIP MAY respond to an OAM message generated by the MEP at its sub-layer of the same OAM domain.

* control/management/data plane: the control and management planes are used to configure and control the network (long-term). On a per-node basis, the data plane applies rules and policies for each packet. For example, selecting the time-frequency block or the next hop on a packet-by-packet basis. Relative to a data flow, the control and/or management plane can be out-of-band.

* Active measurement methods (as defined in [RFC7799]) modify a normal data flow by inserting novel fields, injecting specially constructed test packets [RFC2544]). It is critical for the quality of information obtained since generated test packets are in-band with the monitored data flow. In other words, a test packet is required to cross the same network nodes and links and receive the same Quality of Service (QoS) treatment as a data packet. Active methods may implement one of these two strategies:

  - In-band: control information follows the same path as the data packets. In other words, a failure in the data plane may prevent the control information from reaching the destination (e.g., end-device or controller).

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

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

We also adopt the following terminology, which is particularly relevant for RAW-specific (aka wireless) segments.

* piggybacking vs. dedicated control packets: control information may be encapsulated in specific (dedicated) control packets. Alternatively, it may be piggybacked in existing data packets, when the MTU is larger than the actual packet length. Piggybacking makes specifically sense in wireless networks, as the cost (bandwidth and energy) is sublinear with the packet size. Indeed, the cost to access the medium (e.g., early wake-up to deal with clock drifts) cannot be neglected, and is counted once, whatever the packet size.

* router-over vs. mesh under: a control packet is either forwarded directly without being processed (mesh under) or handled hop-by-hop by each router. While the latter option consumes more resources, it allows collecting additional intermediary information, particularly relevant in wireless networks. For instance, each router is a MIP and inserts its own ID in the packet’s hader, so that the destination reconstructs a posteriori the list of IDs that actually forwarded a packet.

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

* Fault: an irrevocable change which may affect the network performance, e.g., a node runs out of energy.

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

1.2. Acronyms

OAM Operations, Administration, and Maintenance

DetNet Deterministic Networking

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

QoS Quality of Service

RAW Reliable and Available Wireless

SLO Service Level Objective

SNMP Simple Network Management Protocol

SDN Software-Defined Network

1.3. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Role of OAM in RAW

RAW networks expect to provide availability and reliability over a wireless network infrastructure. Most critical applications will define an SLO required for the data flows it generates. RAW expects to exploit OAM to improve the RAW operation at the service and the forwarding sub-layers.

To respect strict guarantees, RAW relies on the Path Selection Engine (PSE) (as defined in [I-D.pthubert-raw-architecture] to monitor and maintain the L3 network. Any L2 scheduling mechanism may be used to allocate transmission opportunities, based on the radio link characteristics, the SLO of the flows, or the number of packets to forward. The PSE exploits the L2 resources reserved by the scheduler and organizes the L3 paths to introduce redundancy, fault tolerance and create backup paths. OAM represents the core of the pre-provisioning process by supervising the network. It maintains a global view of the network resources to detect defects, faults, over-provisioning, anomalies.

Fault tolerance also assumes that multiple paths must be provisioned so that an end-to-end circuit remains operational regardless of the conditions. The Packet Replication and Elimination Function ([I-D.pthubert-raw-architecture]) on a node is typically controlled by the PSE. OAM mechanisms can be used to monitor that PREOF is working correctly on a node and within the domain.

To be energy-efficient, out-of-band OAM SHOULD only be used to report aggregated statistics (e.g., counters, histograms) from the nodes using, e.g., SNMP or Netconf/Restconf using YANG-based data models. The out-of-band OAM flow MAY use a dedicated control and management channel, dedicated to this purpose.

RAW supports both proactive and on-demand troubleshooting. Proactively, it is necessary to detect anomalies, report defects, or reduce over-provisioning if it is not required. However, on-demand may also be required to identify the cause of a specific defect. Indeed, some specific faults may only be detected with a global, detailed view of the network, which is too expensive to acquire in the normal operating mode.

The specific characteristics of RAW are discussed below.

2.1. Link concept and quality

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

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

* point-to-multipoint (p2mp) link associates one transmitter and a collection of receivers. For instance, broadcast packets assume the existence of p2mp links to avoid duplicating a broadcast packet to reach each possible radio neighbor.

In scheduled radio networks, p2mp and p2p links are commonly not scheduled simultaneously to save energy and/or to reduce the number of collisions. More precisely, only a fraction of the neighbors may wake up at a given instant.

Each wireless link is associated with a link quality, often measured as the Packet Delivery Ratio (PDR), i.e., the probability that the receiver can decode the packet correctly. It is worth noting that this link quality depends on many criteria, such as the level of external interference, the presence of concurrent transmissions, or

the radio channel state. This link quality is even time-variant. For p2mp links, consequently, we have a collection of PDR (one value per receiver). Other more sophisticated, aggregated metrics exist for these p2mp links, such as [anycast-property]

2.2. Broadcast Transmissions

In modern switched networks, unicast transmissions are delivered exclusively to the destination . Wireless networks are much closer to the traditional *shared access* wired networks. Practically, unicast and broadcast frames are handled similarly at the physical layer. The link layer is just in charge of filtering the frames to discard irrelevant receptions (e.g., different unicast MAC addresses).

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

2.3. Complex Layer 2 Forwarding

Multiple neighbors may receive a transmission. Thus, anycast Layer 2 forwarding helps to maximize reliability by assigning multiple receivers to a single transmission. That way, the packet is lost only if *none* of the receivers decode it. Practically, it has been proven that different neighbors may exhibit very different radio conditions, and that reception independence may hold for some of them [anycast-property]. Anycast transmission typically exploit p2mp links.

2.4. End-to-end delay

In a wireless network, additional transmissions opportunities are provisioned to accommodate packet losses. Thus, the end-to-end delay consists of:

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

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

3. Operation

OAM features will enable reliability and availability with robust operations both for forwarding and routing purposes.

3.1. Information Collection

The model for exchanging information should be the same as for a DetNet network to ensure inter-operability. YANG may typically fulfill this objective.

However, RAW networks imply specific constraints (e.g., low bandwidth, packet losses, cost of medium access) that may require to minimize the volume of information to collect. Thus, we discuss in Section 4.1 different ways to collect information, i.e., transfer the OAM information physically from the emitter to the receiver. This corresponds to passive OAM as defined in [RFC7799].

3.2. Continuity Check

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

3.3. Connectivity Verification

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

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

3.4. Route Tracing

Wireless networks are broadcast by nature: a radio transmission can be decoded by any radio neighbor. In multihop wireless networks, several paths exist between two endpoints. In hub networks, a device may be covered by several Access Points. The network must select the most efficient path or AP, concerning specifically the reliability, and the delay.

Thus, multipath routing / multi-attachment can be viewed as making the network more fault-tolerant. Even better, we can exploit the broadcast nature of wireless networks: we may have multiple Maintenance Intermediate Points (MIP) for each of these kinds of hop. While it may be reasonable in the multi-attachment case, the complexity quickly increases with the path length. Indeed, each MIP has several possible next hops in the forwarding plane. Thus, all the possible paths between two MEPs should be retrieved, which may quickly become intractable if we apply a naive approach.

3.5. Fault detection

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

3.6. Fault identification

While DetNet already expects to identify malfunctions, some problems are specific to wireless networks. We must consequently collect metrics and implement algorithms tailored for wireless networking.

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

4. Administration

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

* Packet losses: the time-window average and maximum values of the number of packet losses have to be measured. Many critical applications stop working if a few consecutive packets are dropped.

* Received Signal Strength Indicator (RSSI) is a very common metric in wireless to denote the link quality. The radio chipset is in charge of translating a received signal strength into a normalized quality indicator.

* Delay: the time elapsed between a packet generation / enqueuing and its reception by the next hop. In wireless networks, the delay has also to take into consideration possible retransmissions.

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

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

These metrics should be collected per device, virtual circuit, and path, as DetNet already does. However, in RAW, we have to deal with them at a finer granularity:

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

* per physical radio technology / interface, if a device has multiple NICs.

* per link to detect a misbehaving link (asymmetrical link, or with a fluctuating quality).

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

RAW inherits the same requirements as DetNet: we need to know the distribution of a collection of metrics. Besides, wireless networks are known to be highly variable. Changes may be frequent, and may exhibit a periodical pattern. Collecting and analyzing this amount of measurements is challenging. OAM should find an efficient method to encode these time-series in a compact form.

4.1. Efficient measurement retrieval (Passive OAM)

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

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

* bandwidth: wireless networks exhibit a bandwidth significantly lower than wired networks.

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

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

* on-path collection: control information is inserted in the data packets if they do not fragment the packet (i.e., the MTU is not exceeded). Information Elements represent a standardized way to handle such information. IP hop by hop extension headers may help to collect metrics all along the path.

* flags/fields: we have to set-up flags in the packets to monitor to be able to monitor the forwarding process accurately. A sequence number field may help to detect packet losses. Similarly, path inference tools such as [ipath] insert additional information in the headers to identify the path followed by a packet a posteriori.

* hierarchical monitoring: localized and centralized mechanisms have to be combined together. Typically, a local mechanism should continuously monitor a set of metrics and trigger remote OAM exchanges only when a fault is detected (but possibly not identified). For instance, local temporary defects must not trigger expensive OAM transmissions. Besides, the wireless segments often represent the weakest parts of a path: the volume of control information they produce has to be fixed accordingly.

Several passive techniques can be combined. For instance, the DetNet forwarding sublayer MAY combine In-band Network Telemetry (INT) with P4, iOAM and iPath to compute and report different statistics in the track (e.g., number of link-layer retransmissions, link reliability).

4.2. Reporting OAM packets to the source (Active OAM)

The MEP will collect measurements from the OAM probes received in the monitored track. However, the aggregated statistics must then be reported to the other MEP that injected the probes. Unfortunately, the monitored track MAY be unidirectional. In this case, the statistics have to be reported out-of-band (through, e.g., a dedicated control or management channel).

It is worth noting that Active OAM and Passive OAM techniques are not mutually exclusive. In particular, Active OAM is useful when a statistic cannot be accurately acquired passively.

Besides, Active OAM may also use piggybacking techniques: the OAM packet may be piggybacked in a frame if the MTU is sufficient. Indeed, increasing the number of transmissions in radio networks may very negatively impact the performance of radio networks,

particularly for scheduled access, with fixed timeslot durations. Thus, OAM packets may be buffered until another frame has sufficient space, and has to be transmitted to the same neighbor. In conclusion, active OAM packets may be out-of-band or in-band.

5. Maintenance

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

5.1. Transient state after reconfiguration

Because of the wireless medium, the link quality may fluctuate, and the network needs to reconfigure itself continuously. During this transient state, flows may begin to be gradually re-forwarded, consuming resources in different parts of the network. OAM has to make a distinction between a metric that changed because of an usual network change (e.g., flow redirection) and an unexpected event (e.g., a fault). In a general manner, OAM mechanisms have to provide a consistent view of the OAM domain, even during the reconfiguration.

5.2. Predictions

RAW needs to implement self-optimization features. While the network is configured to be fault-tolerant, a reconfiguration may be required to keep on respecting long-term objectives. The network must continuously retrieve the state of the network, to judge about the relevance of a reconfiguration. More precisely, the OAM mechanisms have to provide enough information to predict and quantify:

* the gain of the reconfiguration: what would be the network state after the reconfiguration (e.g., reduction of the bandwidth or energy consumption)?

* the reconfiguration cost: what is the cost (energy, bandwidth) to reconfigure the forwarding and management planes?

Wireless networks exhibit non linear dependencies among links / radio channels / technologies that complexify significantly such predictions.

6. Requirements

This section lists requirements for OAM in a RAW domain:

1. Maintenance Intermediate and End Point device MUST expose a list of available metrics per flow. It MUST at least provide the end-to-end Packet Delivery Ratio, end-to-end latency, and Maximum Consecutive Failures (MCF).

2. PREOF functions MUST guarantee order preservation for a flow.

3. OAM nodes MUST provide aggregated statistics to reduce the volume of traffic for measurements. They MAY send a compressed distribution of measurements, or MIN / MAX values over a time interval.

4. Maintenance End Points SHOULD support route tracing with hybrid OAM techniques.

7. IANA Considerations

This document has no actionable requirements for IANA. This section can be removed before the publication.

8. Security Considerations

This document lists the OAM requirements for an OAM wireless domain and does not raise any security concerns or issues in addition to ones common to networking and those specific to a DetNet discussed in [RFC9055].

9. Acknowledgments

The authors express their appreciation and gratitude to the colleagues who carefully reviewed the draft and shared their comments (Xavi Vilajosana, Dominique Barthel, Pascal Thubert), and all the RAW and Detnet working group members in general.

10. Informative References

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

[I-D.ietf-detnet-oam-framework]
Mirsky, G., Theoleyre, F., Papadopoulos, G. Z., Bernardos, C. J., Varga, B., and J. Farkas, "Framework of Operations, Administration and Maintenance (OAM) for Deterministic Networking (DetNet)", Work in Progress, Internet-Draft, draft-ietf-detnet-oam-framework-08, 1 February 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-detnet-oam-framework-08>.

[I-D.pthubert-raw-architecture]
Thubert, P., Papadopoulos, G. Z., and L. Berger, "Reliable and Available Wireless Architecture/Framework", Work in Progress, Internet-Draft, draft-pthubert-raw-architecture-09, 7 July 2021, <https://datatracker.ietf.org/doc/html/draft-pthubert-raw-architecture-09>.

[ipath] Gao, Y., Dong, W., Chen, C., Bu, J., Wu, W., and X. Liu, "iPath: path inference in wireless sensor networks.", 2016, <https://doi.org/10.1109/TNET.2014.2371459>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC2544] Bradner, S. and J. McQuaid, "Benchmarking Methodology for Network Interconnect Devices", RFC 2544, DOI 10.17487/RFC2544, March 1999, <https://www.rfc-editor.org/info/rfc2544>.

[RFC6291] Andersson, L., van Helvoort, H., Bonica, R., Romascanu, D., and S. Mansfield, "Guidelines for the Use of the "OAM" Acronym in the IETF", BCP 161, RFC 6291, DOI 10.17487/RFC6291, June 2011, <https://www.rfc-editor.org/info/rfc6291>.

[RFC7276] Mizrahi, T., Sprecher, N., Bellagamba, E., and Y. Weingarten, "An Overview of Operations, Administration, and Maintenance (OAM) Tools", RFC 7276, DOI 10.17487/RFC7276, June 2014, <https://www.rfc-editor.org/info/rfc7276>.

[RFC7799] Morton, A., "Active and Passive Metrics and Methods (with Hybrid Types In-Between)", RFC 7799, DOI 10.17487/RFC7799, May 2016, <https://www.rfc-editor.org/info/rfc7799>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC8655] Finn, N., Thubert, P., Varga, B., and J. Farkas, "Deterministic Networking Architecture", RFC 8655, DOI 10.17487/RFC8655, October 2019, <https://www.rfc-editor.org/info/rfc8655>.

[RFC9055] Grossman, E., Ed., Mizrahi, T., and A. Hacker, "Deterministic Networking (DetNet) Security Considerations", RFC 9055, DOI 10.17487/RFC9055, June 2021, <https://www.rfc-editor.org/info/rfc9055>.

Authors' Addresses

Fabrice Theoleyre
Centre National de la Recherche Scientifique
Building B
300 boulevard Sebastien Brant - CS 10413
67400 Illkirch - Strasbourg
France
Phone: +33 368 85 45 33
Email: fabrice.theoleyre@cnrs.fr
URI: http://www.theoleyre.eu

Georgios Z. Papadopoulos
IMT Atlantique
Office B00 - 102A
2 Rue de la Chataigneraie
35510 Cesson-Sevigne - Rennes
France
Phone: +33 299 12 70 04
Email: georgios.papadopoulos@imt-atlantique.fr

Greg Mirsky
Ericsson
United States of America
Email: gregimirsky@gmail.com

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

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

RAW
Internet-Draft
Intended status: Informational
Expires: 19 October 2023

CJ. Bernardos, Ed.
UC3M
G.Z. Papadopoulos
IMT Atlantique
P. Thubert
Cisco
F. Theoleyre
CNRS
17 April 2023

RAW Use-Cases
draft-ietf-raw-use-cases-11

Abstract

The wireless medium presents significant specific challenges to achieve properties similar to those of wired deterministic networks. At the same time, a number of use-cases cannot be solved with wires and justify the extra effort of going wireless. This document presents wireless use-cases (such as aeronautical communications, amusement parks, industrial applications, pro audio and video, gaming, UAV and V2V control, edge robotics and emergency vehicles) demanding reliable and available behavior.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

Based on time, resource reservation, and policy enforcement by distributed shapers, deterministic networking (DetNet) provides the capability to carry specified unicast or multicast data streams for real-time applications with extremely low data loss rates and bounded latency, so as to support time-sensitive and mission-critical applications on a converged enterprise infrastructure.

Deterministic networking aims at eliminating packet loss for a committed bandwidth while ensuring a worst case end-to-end latency, regardless of the network conditions and across technologies. By leveraging lower layer (Layer 2 and below) capabilities, L3 can exploit the use of a service layer, steering over multiple technologies, and using media independent signaling to provide high reliability, precise time delivery, and rate enforcement. Deterministic networking can be seen as a set of new Quality of Service (QoS) guarantees of worst-case delivery. IP networks become more deterministic when the effects of statistical multiplexing (jitter and collision loss) are mostly eliminated. This requires a tight control of the physical resources to maintain the amount of traffic within the physical capabilities of the underlying

technology, e.g., using time-shared resources (bandwidth and buffers) per circuit, and/or by shaping and/or scheduling the packets at every hop.

Key attributes of Deterministic networking include:

* time synchronization on all the nodes,

* multi-technology path with co-channel interference minimization,

* frame preemption and guard time mechanisms to ensure a worst-case delay, and

* new traffic shapers within and at the edge to protect the network.

Wireless operates on a shared medium, and transmissions cannot be guaranteed to be fully deterministic due to uncontrolled interferences, including self-induced multipath fading. The term RAW stands for Reliable and Available Wireless, and refers to the mechanisms aimed for providing high reliability and availability for IP connectivity over a wireless medium. Making Wireless Reliable and Available is even more challenging than it is with wires, due to the numerous causes of loss in transmission that add up to the congestion losses and the delays caused by overbooked shared resources.

The wireless and wired media are fundamentally different at the physical level, and while the generic Problem Statement [RFC8557] for DetNet applies to the wired as well as the wireless medium, the methods to achieve RAW necessarily differ from those used to support Time-Sensitive Networking over wires, e.g., due to the wireless radio channel specifics.

So far, open standards for deterministic networking have prevalently been focused on wired media, with Audio/Video Bridging (AVB) and Time Sensitive Networking (TSN) at the IEEE and DetNet [RFC8655] at the IETF. But wires cannot be used in several cases, including mobile or rotating devices, rehabilitated industrial buildings, wearable or in-body sensory devices, vehicle automation and multiplayer gaming.

Purpose-built wireless technologies such as [ISA100], which incorporates IPv6, were developed and deployed to cope with the lack of open standards, but they yield a high cost in OPEX and CAPEX and are limited to very few industries, e.g., process control, concert instruments or racing.

This is now changing [I-D.ietf-raw-technologies]:

* IMT-2020 has recognized Ultra-Reliable Low-Latency Communication (URLLC) as a key functionality for the upcoming 5G.

* IEEE 802.11 has identified a set of real-applications [IEEE80211-RT-TIG] which may use the IEEE802.11 standards. They typically emphasize strict end-to-end delay requirements.

* The IETF has produced an IPv6 stack for IEEE Std. 802.15.4 TimeSlotted Channel Hopping (TSCH) and an architecture [RFC9030] that enables RAW on a shared MAC.

Experiments have already been conducted with IEEE802.1 TSN over IEEE802.11be [IEEE80211BE]. This mode enables time synchronization, and time-aware scheduling (trigger based access mode) to support TSN flows.

This document extends the "Deterministic Networking use-cases" document [RFC8578] and describes several additional use-cases which require "reliable/predictable and available" flows over wireless links and possibly complex multi-hop paths called Tracks. This is covered mainly by the "Wireless for Industrial Applications" use-case, as the "Cellular Radio" is mostly dedicated to the (wired) link part of a Radio Access Network (RAN). Whereas the "Wireless for Industrial Applications" use-case certainly covers an area of interest for RAW, it is limited to 6TiSCH, and thus its scope is narrower than the use-cases described next in this document.

2. Aeronautical Communications

Aircraft are currently connected to ATC (Air-Traffic Control) and AOC (Airline Operational Control) via voice and data communication systems through all phases of a flight. Within the airport terminal, connectivity is focused on high bandwidth communications while en-route high reliability, robustness and range are the focus.

2.1. Problem Statement

Up to 2020, civil air traffic has been growing constantly at a compound rate of 5.8% per year [ACI19] and despite the severe impact of the COVID-19 pandemic, air traffic growth is expected to resume very quickly in post-pandemic times [IAT20] [IAC20]. Thus, legacy systems in air traffic management (ATM) are likely to reach their capacity limits and the need for new aeronautical communication technologies becomes apparent. Especially problematic is the saturation of VHF band in high density areas in Europe, the US, and Asia [KEAV20] [FAA20] calling for suitable new digital approaches such as AeroMACS for airport communications, SatCOM for remote domains, and LDACS as long-range terrestrial aeronautical

communication system. Making the frequency spectrum's usage more efficient a transition from analog voice to digital data communication [PLA14] is necessary to cope with the expected growth of civil aviation and its supporting infrastructure. A promising candidate for long range terrestrial communications, already in the process of being standardized in the International Civil Aviation Organization (ICAO), is the L-band Digital Aeronautical Communication System (LDACS) [ICAO18] [I-D.ietf-raw-ldacs].

Note that the large scale of the planned low Earth orbit (LEO) constellations can provide fast end-to-end latency rates and high data-rates at a reasonable cost, but they also pose challenges such as frequent handovers, high-interference, and a diverse range of system users, which can create security issues since both safety-critical and non-safety-critical communications can take place on the same system. Some studies suggest that LEO constellations could be a complete solution for aeronautical communications, but they do not offer solutions for the critical issues mentioned earlier. Additionally, of the three communication domains defined by ICAO, only passenger entertainment services can currently be provided using these constellations. Safety-critical aeronautical communications require reliability levels above 99.999%, which is higher than that required for regular commercial data links. Therefore, addressing the issues with LEO-based SatCOM is necessary before these solutions can reliably support safety-critical data transmission [Maurer2022].

2.2. Specifics

During the creation process of new communication system, analog voice is replaced by digital data communication. This sets a paradigm shift from analog to digital wireless communications and supports the related trend towards increased autonomous data processing that the Future Communications Infrastructure (FCI) in civil aviation must provide. The FCI is depicted in Figure 1:

```
 Satellite
#         #
#          # #
#            #   #
#             #      #
#               #        #
#                #          #
#                  #            #
# Satellite-based   #              #
#   Communications   #                 #
#      SatCOM (#)     #                   #
#                      #                    Aircraft
#                       #                  %         %
#                        #               %             %
#                         #            %     Air-Air     %
#                          #         %     Communications   %
#                           #     %         LDACS A/A (%)    %
#                           #   %                            %
#                            Aircraft  % % % % % % % % % %  Aircraft
#                                  |           Air-Ground           |
#                                  |         Communications         |
#                                  |           LDACS A/G (|)        |
#      Communications in           |                                |
#    and around airports           |                                |
#         AeroMACS (-)             |                                |
#                                  |                                |
#         Aircraft-------------+   |                                |
#                              |   |                                |
#                              |   |                                |
#         Ground network       |   |         Ground network         |
SatCOM <---------------------> Airport <----------------------> LDACS
ground                          ground                          ground
transceiver                    transceiver                    transceiver
```

Figure 1: The Future Communication Infrastructure (FCI): AeroMACS for Airport/ Termina Maneuvering Area domain, LDACS A/G for Terminal Maneuvering/ En-Route domain, LDACS A/G for En-Route/ Oceanic, Remote, Polar domain, SatCOM for Oceanic, Remote, Polar domain domain communications

2.3. Challenges

This paradigm change brings a lot of new challenges:

* Efficiency: It is necessary to keep latency, time and data overhead of new aeronautical datalinks at a minimum.

* Modularity: Systems in avionics usually operate up to 30 years, thus solutions must be modular, easily adaptable and updatable.

* Interoperability: All 192 members of the international Civil Aviation Organization (ICAO) must be able to use these solutions.

* Dynamicity: the communication infrastructure needs to accommodate mobile devices (airplanes) that move extremely fast.

2.4. The Need for Wireless

In a high mobility environment such as aviation, the envisioned solutions to provide worldwide coverage of data connections with in-flight aircraft require a multi-system, multi-link, multi-hop approach. Thus air, ground and space-based datalink providing technologies will have to operate seamlessly together to cope with the increasing needs of data exchange between aircraft, air traffic controller, airport infrastructure, airlines, air network service providers (ANSPs) and so forth. Wireless technologies have to be used to tackle this enormous need for a worldwide digital aeronautical datalink infrastructure.

2.5. Requirements for RAW

Different safety levels need to be supported. All network traffic handled by the Airborne Internet Protocol Suite (IPS) System is not equal and the Quality of Service (QoS) requirements of each network traffic flow must be considered n order to avoid having to support QoS requirements at the granularity of data flows, these flows are grouped into classes that have similar requirements, following the DiffServ approach [ARINC858P1]. These classes are referred to as Classes of Service (CoS) and flows within a class are treated uniformly from a QoS perspective. Currently, there are at least eight different priority levels (CoS) that can be assigned to packets. For example, a high-priority message requiring low latency and high resiliency could be a "WAKE" warning indicating two aircraft are dangerously close to each other, while a less safety-critical message with low-medium latency requirements could be the "WXGRAPH" service providing graphical weather data.

Overhead needs to be kept at a minimum since aeronautical data links provide comparatively small data rates on the order of kbit/s.

Policy needs to be supported when selecting data links. The focus of RAW here should be on the selectors, responsible for the track a packet takes to reach its end destination. This would minimize the amount of routing information that must travel inside the network because of precomputed routing tables with the selector being responsible for choosing the most appropriate option according to policy and safety.

2.5.1. Non-latency critical considerations

Achieving low latency is a requirement for aeronautics communications, though the expected latency is not extremely low and what is important is to keep the overall latency bounded under a certain threshold. Low latency in LDACS communications [RFC9372] translates to a latency in the Forward Link (FL - Ground -> Air) of 30-90 ms and a latency in the Reverse Link (RL - Air -> Ground) of 60-120 ms. This use-case is not latency-critical from that view point. On the other hand, given the controlled environment, end-to-end mechanisms can be applied to guarantee bounded latency where needed.

3. Amusement Parks

3.1. Use-Case Description

The digitalization of Amusement Parks is expected to decrease significantly the cost for maintaining the attractions. Such deployment is a mix between multimedia (e.g., Virtual and Augmented Reality, interactive video environments) and non-multimedia applications (e.g, industrial automation for a roller-coaster, access control).

Attractions may rely on a large set of sensors and actuators, which react in real time. Typical applications comprise:

* Emergency: the safety of the operators / visitors has to be preserved and the attraction must be stopped appropriately when a failure is detected.

* Video: augmented and virtual realities are integrated in the attraction. Wearable mobile devices (e.g., glasses, virtual reality headset) need to offload one part of the processing tasks.

* Real-time interactions: visitors may interact with an attraction, like in a real-time video game. The visitors may virtually interact with their environment, triggering actions in the real world (through actuators) [KOB12].

* Geolocation: visitors are tracked with a personal wireless tag so that their user experience is improved. This requires special care to ensure that visitors' privacy is not breached, and users are anonymously tracked.

* Predictive maintenance: statistics are collected to predict the future failures, or to compute later more complex statistics about the attraction's usage, the downtime, etc.

* Marketing: to improve the customer experience, owners may collect a large amount of data to understand the behavior, and the choice of their clients.

3.2. Specifics

Amusement parks comprise a variable number of attractions, mostly outdoor, over a large geographical area. The IT infrastructure is typically multi-scale:

* Local area: the sensors and actuators controlling the attractions are co-located. Control loops trigger only local traffic, with a small end-to-end delay, typically less than 10 ms, like classical industrial systems [IEEE80211-RT-TIG].

* Wearable mobile devices are free to move in the park. They exchange traffic locally (identification, personalization, multimedia) or globally (billing, child tracking).

* Computationally intensive applications offload some tasks. Edge computing seems an efficient way to implement real-time applications with offloading. Some non-time-critical tasks may rather use the cloud (predictive maintenance, marketing).

3.3. The Need for Wireless

Removing cables helps to change easily the configuration of the attractions, or to upgrade parts of them at a lower cost. The attraction can be designed modularly, upgrade or insert novel modules later in the lifecycle of the attraction. Novelty of attractions tends to increase the attractiveness of an amusement park, encouraging previous visitors to visit regularly the park.

Some parts of the attraction are mobile, like trucks of a roller-coaster or robots. Since cables are prone to frequent failures in this situation, wireless transmissions are recommended.

Wearable devices are extensively used for a user experience personalization. They typically need to support wireless transmissions. Personal tags may help to reduce the operating costs [DISNEY15] and to increase the number of charged services provided to the audience (e.g., VIP tickets or interactivity). Some applications rely on more sophisticated wearable devices such as digital glasses or Virtual Reality (VR) headsets for an immersive experience.

3.4. Requirements for RAW

The network infrastructure must support heterogeneous traffic, with very different critical requirements. Thus, flow isolation must be provided.

The transmissions must be scheduled appropriately even in presence of mobile devices. While the [RFC9030] already proposes an architecture for synchronized, IEEE Std. 802.15.4 Time-Slotted Channel Hopping (TSCH) networks, the industry requires a multi-technology solution, able to guarantee end-to-end requirements across heterogeneous technologies, with strict SLA requirements.

Nowadays, long-range wireless transmissions are used mostly for best-effort traffic. On the contrary, [IEEE802.1TSN] is used for critical flows using Ethernet devices. However, we need an IP enabled technology to interconnect large areas, independent of the PHY and MAC layers.

It is expected that several different technologies (long vs. short range) are deployed, which have to cohabit in the same area. Thus, we need to provide layer-3 mechanisms able to exploit multiple co-interfering technologies (i.e., different radio technologies using overlapping spectrum, and therefore, potentially interfering to each other).

It is worth noting that low-priority flows (e.g., predictive maintenance, marketing) are delay tolerant: a few minutes or even hours would be acceptable. While classical unscheduled wireless networks already accomodate best-effort traffic, this would force several colocated and subefficient deployments. Unused resources could rather be used for low-priority flows. Indeed, allocated resources are consuming energy in most scheduled networks, even if no traffic is transmitted.

3.4.1. Non-latency critical considerations

While some of the applications in this use-case involve control loops (e.g., sensors and actuators) that require bounded latencies below 10 ms, that can therefore be considered latency critical, there are other applications as well that mostly demand reliability (e.g., safety related, or maintenance).

4. Wireless for Industrial Applications

4.1. Use-Case Description

A major use-case for networking in Industrial environments is the control networks where periodic control loops operate between a collection of sensors that measure a physical property such as the temperature of a fluid, a Programmable Logic Controller (PLC) that decides an action such as warm up the mix, and actuators that perform the required action, such as the injection of power in a resistor.

4.2. Specifics

4.2.1. Control Loops

Process Control designates continuous processing operations, like heating oil in a refinery or mixing drinking soda. Control loops in the Process Control industry operate at a very low rate, typically four times per second. Factory Automation, on the other hand, deals with discrete goods such as individual automobile parts, and requires faster loops, on the order of milliseconds. Motion control that monitors dynamic activities may require even faster rates on the order of and below the millisecond.

In all those cases, a packet must flow reliably between the sensor and the PLC, be processed by the PLC, and sent to the actuator within the control loop period. In some particular use-cases that inherit from analog operations, jitter might also alter the operation of the control loop. A rare packet loss is usually admissible, but typically a loss of multiple packets in a row will cause an emergency halt of the production and incur a high cost for the manufacturer.

Additional details and use-cases related to Industrial applications and their RAW requirements can be found in [I-D.ietf-raw-industrial-requirements].

4.2.2. Monitoring and diagnostics

A secondary use-case deals with monitoring and diagnostics. This data is essential to improve the performance of a production line, e.g., by optimizing real-time processing or maintenance windows using Machine Learning predictions. For the lack of wireless technologies, some specific industries such as Oil and Gas have been using serial cables, literally by the millions, to perform their process optimization over the previous decades. But few industries would afford the associated cost. One of the goals of the Industrial Internet of Things is to provide the same benefits to all industries, including SmartGrid, Transportation, Building, Commercial and Medical. This requires a cheap, available and scalable IP-based access technology.

Inside the factory, wires may already be available to operate the Control Network. But monitoring and diagnostics data are not welcome in that network for several reasons. On the one hand it is rich and asynchronous, meaning that it may influence the deterministic nature of the control operations and impact the production. On the other hand, this information must be reported to the operators over IP, which means the potential for a security breach via the interconnection of the Operational Technology (OT) network with the Internet technology (IT) network and possibly enable a rogue access.

4.3. The Need for Wireless

Wires used on a robot arm are prone to breakage after a few thousands flexions, a lot faster than a power cable that is wider in diameter, and more resilient. In general, wired networking and mobile parts are not a good match, mostly in the case of fast and recurrent activities, as well as rotation.

When refurbishing older premises that were built before the Internet age, power is usually available everywhere, but data is not. It is often impractical, time consuming and expensive to deploy an Ethernet fabric across walls and between buildings. Deploying a wire may take months and cost tens of thousands of US Dollars.

Even when wiring exists, like in the case of an existing control network, asynchronous IP packets such as diagnostics may not be welcome for operational and security reasons. For those packets, the option to create a parallel wireless network offers a credible solution that can scale with the many sensors and actuators that equip every robot, every valve and fan that are deployed on the factory floor. It may also help detect and prevent a failure that could impact the production, like the degradation (vibration) of a cooling fan on the ceiling. IEEE Std. 802.15.4 Time-Slotted Channel

Hopping (TSCH) [RFC7554] is a promising technology for that purpose, mostly if the scheduled operations enable to use the same network by asynchronous and deterministic flows in parallel.

4.4. Requirements for RAW

As stated by the "Deterministic Networking Problem Statement" [RFC8557], a deterministic network is backwards compatible with (capable of transporting) statistically multiplexed traffic while preserving the properties of the accepted deterministic flows. While the 6TiSCH Architecture [RFC9030] serves that requirement, the work at 6TiSCH was focused on best-effort IPv6 packet flows. RAW should be able to lock so-called hard cells (i.e., scheduled cells [I-D.ietf-6tisch-terminology]) for use by a centralized scheduler, and leverage time and spatial diversity over a graph of end-to-end paths called a Track that is based on those cells.

Over the course of the recent years, major Industrial Protocols (e.g., [ODVA] with EtherNet/IP [EIP] and [PROFINET]) have been migrating towards Ethernet and IP. In order to unleash the full power of the IP hourglass model, it should be possible to deploy any application over any network that has the physical capacity to transport the industrial flow, regardless of the MAC/PHY technology, wired or wireless, and across technologies. RAW mechanisms should be able to setup a Track over a wireless access segment and a wired or wireless backbone to report both sensor data and critical monitoring within a bounded latency and maintain the high reliability of the flows over time. It is also important to ensure that RAW solutions are interoperable with existing wireless solutions in place, and with legacy equipment whose capabilities can be extended using retrofitting. Maintainability, as a broader concept than reliability is also important in industrial scenarios [MAR19].

4.4.1. Non-latency critical considerations

Monitoring and diagnostics applications do not require latency critical communications, but demand reliable and scalable communications. On the other hand, process control applications involve control loops that require a bounded latency, thus are latency critical, but can be managed end-to-end, and therefore DetNet mechanisms can be applied in conjunction with RAW mechanisms.

5. Pro Audio and Video

5.1. Use-Case Description

Many devices support audio and video streaming [RFC9317] by employing 802.11 wireless LAN. Some of these applications require low latency capability. For instance, when the application provides interactive play, or when the audio plays in real time - meaning live for public addresses in train stations or in theme parks.

The professional audio and video industry ("ProAV") includes:

* Virtual Reality / Augmented Reality (VR/AR)

* Production and post-production systems such as CD and Blu-ray disk mastering.

* Public address, media and emergency systems at large venues (e.g., airports, train stations, stadiums, and theme parks).

5.2. Specifics

5.2.1. Uninterrupted Stream Playback

Considering the uninterrupted audio or video stream, a potential packet loss during the transmission of audio or video flows cannot be tackled by re-trying the transmission, as it is done with file transfer, because by the time the packet lost has been identified it is too late to proceed with packet re-transmission. Buffering might be employed to provide a certain delay which will allow for one or more re-transmissions, however such approach is not viable in application where delays are not acceptable.

5.2.2. Synchronized Stream Playback

In the context of ProAV over packet networks, latency is the time between the transmitted signal over a stream and its reception. Thus, for sound to remain synchronized to the movement in the video, the latency of both the audio and video streams must be bounded and consistent.

5.3. The Need for Wireless

The devices need the wireless communication to support video streaming via IEEE 802.11 wireless LAN for instance. Wireless communications provide huge advantages in terms of simpler deployments in many scenarios, where the use of a wired alternative would not be feasible. Similarly, in live events, mobility support makes wireless communications the only viable approach.

Deployed announcement speakers, for instance along the platforms of the train stations, need the wireless communication to forward the audio traffic in real time. Most train stations are already built, and deploying novel cables for each novel service seems expensive.

5.4. Requirements for RAW

The network infrastructure needs to support heterogeneous types of traffic (including QoS).

Content delivery with bounded (lowest possible) latency.

The deployed network topology should allow for multipath. This will enable for multiple streams to have different (and multiple) paths (tracks) through the network to support redundancy.

5.4.1. Non-latency critical considerations

For synchronized streaming, latency must be bounded, and therefore, depending on the actual requirements, this can be considered as latency critical. However, the most critical requirement of this use-case is reliability, by the network providing redundancy. Note that in many cases, wireless is only present in the access, where RAW mechanisms could be applied, but other wired segments are also involved (like the Internet), and therefore latency cannot be guaranteed.

6. Wireless Gaming

6.1. Use-Case Description

The gaming industry includes [IEEE80211RTA] real-time mobile gaming, wireless console gaming, wireless gaming controllers and cloud gaming. Note that they are not mutually exclusive (e.g., a console can connect wirelessly to the Internet to play a cloud game). For RAW, wireless console gaming is the most relevant one. We next summarize the four:

* Real-time Mobile Gaming: Different from traditional games, real time mobile gaming is very sensitive to network latency and stability. The mobile game can connect multiple players together in a single game session and exchange data messages between game server and connected players. Real-time means the feedback should present on screen as users operate in game. For good game experience, the end-to-end (E2E) latency plus game servers processing time must be the same for all players and should not be noticeable as the game is played. RAW technologies might help in keeping latencies low on the wireless segments of the communication.

* Wireless Console Gaming: while gamers may use a physical console, interactions with a remote server may be required for online games. Most of the gaming consoles today support Wi-Fi 5, but may benefit from a scheduled access with Wi-Fi 6 in the future. Previous Wi-Fi versions have an especially bad reputation among the gaming community. The main reasons are high latency, lag spikes, and jitter.

* Wireless Gaming controllers: most controllers are now wireless for a freedom of movement.Controllers may interact with consoles or directly with gaming server in the cloud. A low and stable end-to-end latency is here of predominant importance.

* Cloud Gaming: The cloud gaming requires low latency capability as the user commands in a game session need to be sent back to the cloud server, the cloud server would update game context depending on the received commands, and the cloud server would render the picture/video to be displayed at user devices and stream the picture/video content to the user devices. User devices might very likely be connected wirelessly.

6.2. Specifics

While a lot of details can be found on [IEEE80211RTA], we next summarize the main requirements in terms of latency, jitter and packet loss:

* Intra Basic Service Set (BSS) latency is less than 5 ms.

* Jitter variance is less than 2 ms.

* Packet loss is less than 0.1 percent.

6.3. The Need for Wireless

Gaming is evolving towards wireless, as players demand being able to play anywhere, and the game requires a more immersive experience including body movements. Besides, the industry is changing towards playing from mobile phones, which are inherently connected via wireless technologies. Wireless controllers are the rule in modern gaming, with increasingly sophisticated interactions (e.g., haptic feedback, augmented reality).

6.4. Requirements for RAW

* Time sensitive networking extensions: extensions, such as time-aware shaping and redundancy can be explored to address congestion and reliability problems present in wireless networks. As an example, in haptics it is very important to minimize latency failures.

* Priority tagging (Stream identification): one basic requirement to provide better QoS for time-sensitive traffic is the capability to identify and differentiate time-sensitive packets from other (like best-effort) traffic.

* Time-aware shaping: this capability (defined in IEEE 802.1Qbv) consists of gates to control the opening/closing of queues that share a common egress port within an Ethernet switch. A scheduler defines the times when each queue opens or close, therefore eliminating congestion and ensuring that frames are delivered within the expected latency bounds. Note though, that while this requirement needs to be signalled by RAW mechanisms, it would be actually served by the lower layer.

* Dual/multiple link: due to the fact that competitions and interference are common and hardly in control under wireless network, to improve the latency stability, dual/multiple link proposal is brought up to address this issue.

* Admission Control: congestion is a major cause of high/variable latency and it is well known that if the traffic load exceeds the capability of the link, QoS will be degraded. QoS degradation may be acceptable for many applications today, however emerging time-sensitive applications are highly susceptible to increased latency and jitter. To better control QoS, it is important to control access to the network resources.

6.4.1. Non-latency critical considerations

Depending on the actual scenario, and on use of Internet to interconnect different users, the communication requirements of this use-case might be considered as latency critical due to the need of bounded latency. But note that in most of these scenarios, part of the communication path is not wireless and DetNet mechanisms cannot be applied easily (e.g., when the public Internet is involved), and therefore in these cases, reliability is the critical requirement.

7. Unmanned Aerial Vehicles and Vehicle-to-Vehicle platooning and control

7.1. Use-Case Description

Unmanned Aerial Vehicles (UAVs) are becoming very popular for many different applications, including military and civil use-cases. The term drone is commonly used to refer to a UAV.

UAVs can be used to perform aerial surveillance activities, traffic monitoring (i.e., the Spanish traffic control has recently introduced a fleet of drones for quicker reactions upon traffic congestion related events [DGT2021]), support of emergency situations, and even transportation of small goods (e.g., medicine in rural areas). Note that the surveillance and monitoring application would have to comply with local regulations regarding location privacy of users. Different considerations have to be applied when surveillance is performed for traffic rules enforcement (e.g., generating fines) as compared to when traffic load is being monitored.

Many types of vehicles, including UAVs but also others, such as cars, can travel in platoons, driving together with shorter distances between vehicles to increase efficiency. Platooning imposes certain vehicle-to-vehicle considerations, most of these are applicable to both UAVs and other vehicle types.

UAVs/vehicles typically have various forms of wireless connectivity:

* Cellular: for communication with the control center, for remote maneuvering as well as monitoring of the drone;

* IEEE 802.11: for inter-drone communications (i.e., platooning) and providing connectivity to other devices (i.e., acting as Access Point).

Note that autonomous cars share many of the characteristics of the aforemention UAV case, and therefore it is of interest for RAW.

7.2. Specifics

Some of the use-cases/tasks involving UAVs require coordination among UAVs. Others involve complex compute tasks that might not be performed using the limited computing resources that a drone typically has. These two aspects require continuous connectivity with the control center and among UAVs.

Remote maneuvering of a drone might be performed over a cellular network in some cases, however, there are situations that need very low latency and deterministic behavior of the connectivity. Examples involve platooning of drones or sharing of computing resources among drones (like, a drone offload some function to a neighboring drone).

7.3. The Need for Wireless

UAVs cannot be connected through any type of wired media, so it is obvious that wireless is needed.

7.4. Requirements for RAW

The network infrastructure is composed by the UAVs themselves, requiring self-configuration capabilities.

Heterogeneous types of traffic need to be supported, from extremely critical ones requiring ultra-low latency and high resiliency, to traffic requiring low-medium latency.

When a given service is decomposed into functions -- hosted at different UAVs -- chained, each link connecting two given functions would have a well-defined set of requirements (e.g., latency, bandwidth and jitter) that must be met.

7.4.1. Non-latency critical considerations

Today’s solutions keep the processing operations that are critical local (i.e., they are not offloaded). Therefore, in this use-case, the critical requirement is reliability, and only for some platooning and inter-drone communications latency is critical.

8. Edge Robotics control

8.1. Use-Case Description

The Edge Robotics scenario consists of several robots, deployed in a given area (like a shopping mall), inter-connected via an access network to a network edge device or a data center. The robots are connected to the edge so complex computational activities are not executed locally at the robots but offloaded to the edge. This brings additional flexibility in the type of tasks that the robots do, as well as reducing the costs of robot manufacturing (due to their lower complexity), and enabling complex tasks involving coordination among robots (that can be more easily performed if robots are centrally controlled).

Simple examples of the use of multiple robots are cleaning, video surveillance (note that this have to comply with local regulations regarding user's privacy at the application level), search and rescue operations, and delivering of goods from warehouses to shops. Multiple robots are simultaneously instructed to perform individual tasks by moving the robotic intelligence from the robots to the network's edge. That enables easy synchronization, scalable solution, and on-demand option to create flexible fleet of robots.

Robots would have various forms of wireless connectivity:

* IEEE 802.11: for connection to the edge and also inter-robot communications (i.e., for coordinated actions).

* Cellular: as an additional communication link to the edge, though primarily as backup, since ultra-low latency is needed.

8.2. Specifics

Some of the use-cases/tasks involving robots might benefit from decomposition of a service in small functions that are distributed and chained among robots and the edge. These require continuous connectivity with the control center and among drones.

Robot control is an activity requiring very low latency (0.5-20 ms [Groshev2021]) between the robot and the location where the control intelligence resides (which might be the edge or another robot).

8.3. The Need for Wireless

Deploying robots in scenarios such as shopping malls for the applications mentioned cannot be done via wired connectivity.

8.4. Requirements for RAW

The network infrastructure needs to support heterogeneous types of traffic, from robot control to video streaming.

When a given service is decomposed into functions -- hosted at different robots -- chained, each link connecting two given functions would have a well-defined set of requirements (latency, bandwidth and jitter) that must be met.

8.4.1. Non-latency critical considerations

This use-case might combine multiple communication flows, with some of them being latency critical (like those related to robot control tasks). Note that there are still many communication flows (like some offloading tasks) that only demand reliability and availability.

9. Instrumented emergency medical vehicles

9.1. Use-Case Description

An instrumented ambulance would be one that one or multiple network segments to which are connected these end systems such as:

* vital signs sensors attached to the casualty in the ambulance. Relay medical data to hospital emergency room,

* radio-navigation sensor to relay position data to various destinations including dispatcher,

* voice communication for ambulance attendant (like to consult with ER doctor), and

* voice communication between driver and dispatcher.

The LAN needs to be routed through radio-WANs (a radio network in the interior of a network, i.e., it is terminated by routers) to complete the network linkage.

9.2. Specifics

What we have today is multiple communication systems to reach the vehicle via:

* A dispatching system,

* a cellphone for the attendant,

* a special purpose telemetering system for medical data,

* etc.

This redundancy of systems does not contribute to availability.

Most of the scenarios involving the use of an instrumented ambulance are composed of many different flows, each of them with slightly different requirements in terms of reliability and latency. Destinations might be either at the ambulance itself (local traffic), at a near edge cloud or at the general Internet/cloud. Special care (at application level) have to be paid to ensuring that sensitive data is not disclosed to unauthorized parties, by properly securing traffic and authenticating the communication ends.

9.3. The Need for Wireless

Local traffic between the first responders/ambulance staff and the ambulance equipment cannot be done via wired connectivity as the responders perform initial treatment outside of the ambulance. The communications from the ambulance to external services must be wireless as well.

9.4. Requirements for RAW

We can derive some pertinent requirements from this scenario:

* High availability of the inter-network is required. The exact level of availability depends on the specific deployment scenario, as not all emergency agencies share the same type of instrumented emergency vehicles.

* The inter-network needs to operate in damaged state (e.g. during an earthquake aftermath, heavy weather, wildfire, etc.). In addition to continuity of operations, rapid restore is a needed characteristic.

* The radio-WAN has characteristics similar to cellphone -- the vehicle will travel from one radio coverage area to another, thus requiring some hand-off approach.

9.4.1. Non-latency critical considerations

In this case, all applications identified do not require latency critical communication, but do need high reliability and availability.

10. Summary

This document enumerates several use-cases and applications that need RAW technologies, focusing on the requirements from reliability, availability and latency. Whereas some use-cases are latency-critical, there are also several applications that are non-latency critical, but that do pose strict reliability and availability requirements.

11. IANA Considerations

This document has no IANA actions.

12. Security Considerations

This document covers several representative applications and network scenarios that are expected to make use of RAW technologies. Each of the potential RAW use-cases will have security considerations from both the use-specific perspective and the RAW technology perspective. [RFC9055] provides a comprehensive discussion of security considerations in the context of deterministic networking, which are generally applicable also to RAW.

13. Acknowledgments

Nils Mäurer, Thomas Gräupl and Corinna Schmitt have contributed significantly to this document, providing input for the Aeronautical communication section. Rex Buddenberg has also contributed to the document, providing input to the Emergency: instrumented emergency vehicle section.

The authors would like to thank Toerless Eckert, Xavi Vilajosana Guillen, Rute Sofia, Corinna Schmitt, Victoria Pritchard, John Scudder, Joerg Ott and Stewart Bryant for their valuable comments on previous versions of this document.

The work of Carlos J. Bernardos in this document has been partially supported by the Horizon Europe PREDICT-6G (Grant 101095890) and UNICO I+D 6G-DATADRIVEN-04 project.

14. Informative References

[ACI19] Airports Council International (ACI), "Annual World Aitport Traffic Report 2019", November 2019, <https://store.aci.aero/product/annual-world-airport-traffic-report-2019/>.

[ARINC858P1]
ARINC, "INTERNET PROTOCOL SUITE (IPS) FOR AERONAUTICAL SAFETY SERVICES PART 1 AIRBORNE IPS SYSTEM TECHNICAL REQUIREMENTS", 2021, <https://www.sae.org/standards/content/arinc858p1/>.

[DGT2021] Menendez, J. M., "Drones: asi es la vigilancia", 2021, <https://revista.dgt.es/es/reportajes/2021/01ENERO/0126-Como-funciona-un-operativo-con-drones.shtml>.

[DISNEY15] Wired, "Disney's $1 Billion Bet on a Magical Wristband", March 2015, <https://www.wired.com/2015/03/disney-magicband/>.

[EIP] http://www.odva.org/, "EtherNet/IP provides users with the network tools to deploy standard Ethernet technology (IEEE 802.3 combined with the TCP/IP Suite) for industrial automation applications while enabling Internet and enterprise connectivity data anytime, anywhere.", <http://www.odva.org/Portals/0/Library/Publications_Numbered/PUB00138R3_CIP_Adv_Tech_Series_EtherNetIP.pdf>.

[FAA20] U.S. Department of Transportation Federal Aviation Administration (FAA), "Next Generation Air Transportation System", 2019, <https://www.faa.gov/nextgen/>.

[Groshev2021]
Groshev, M., Guimaraes, C., de la Oliva, A., and B. Gazda, "Dissecting the Impact of Information and Communication Technologies on Digital Twins as a Service", IEEE Access, vol. 9 , 2021, <https://doi.org/10.1109/ACCESS.2021.3098109>.

[I-D.ietf-6tisch-terminology]
Palattella, M. R., Thubert, P., Watteyne, T., and Q. Wang, "Terms Used in IPv6 over the TSCH mode of IEEE 802.15.4e", Work in Progress, Internet-Draft, draft-ietf-6tisch-terminology-10, 2 March 2018, <https://datatracker.ietf.org/doc/html/draft-ietf-6tisch-terminology-10>.

[I-D.ietf-raw-industrial-requirements]
Sofia, R. C., Kovatsch, M., and P. Mendes, "Requirements for Reliable Wireless Industrial Services", Work in Progress, Internet-Draft, draft-ietf-raw-industrial-requirements-00, 10 December 2021, <https://datatracker.ietf.org/doc/html/draft-ietf-raw-industrial-requirements-00>.

[I-D.ietf-raw-ldacs]
Mäurer, N., Gräupl, T., and C. Schmitt, "L-band Digital Aeronautical Communications System (LDACS)", Work in Progress, Internet-Draft, draft-ietf-raw-ldacs-14, 2 December 2022, <https://datatracker.ietf.org/doc/html/draft-ietf-raw-ldacs-14>.

[I-D.ietf-raw-technologies]
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-06, 30 November 2022, <https://datatracker.ietf.org/doc/html/draft-ietf-raw-technologies-06>.

[IAC20] Iacus, S.M., Natale, F., Santamaria, C., Spyratos, S., and V. Michele, "Estimating and projecting air passenger traffic during the COVID-19 coronavirus outbreak and its socio- economic impact", Safety Science 129 (2020) 104791 , 2020.

[IAT20] International Air Transport Association (IATA), "Economic Performance of the Airline Industry", November 2020, <https://www.iata.org/en/iata-repository/publications/economic-reports/airline-industry-economic-performance---november-2020---report/>.

[ICAO18] International Civil Aviation Organization (ICAO), "L-Band Digital Aeronautical Communication System (LDACS)", International Standards and Recommended Practices Annex 10 - Aeronautical Telecommunications, Vol. III - Communication Systems , 2018.

[IEEE802.1TSN]
IEEE standard for Information Technology, "IEEE 802.1AS-2011 - IEEE Standard for Local and Metropolitan Area Networks - Timing and Synchronization for Time-Sensitive Applications in Bridged Local Area Networks".

[IEEE80211-RT-TIG]
IEEE, "IEEE 802.11 Real Time Applications TIG Report", November 2018, <http://www.ieee802.org/11/Reports/rtatig_update.htm>.

[IEEE80211BE]
Cavalcanti, D. and G. Venkatesan, "802.1 TSN over 802.11 with updates from developments in 802.11be", IEEE plenary meeting , November 2020, <https://www.ieee802.org/1/files/public/docs2020/new-Cavalcanti-802-1TSN-over-802-11-1120-v02.pdf>.

[IEEE80211RTA]
IEEE standard for Information Technology, "IEEE 802.11 Real Time Applications TIG Report", November 2018.

[ISA100] ISA/ANSI, "ISA100, Wireless Systems for Automation", <https://www.isa.org/isa100/>.

[KEAV20] T. Keaveney and C. Stewart, "Single European Sky ATM Research Joint Undertaking", 2019, <https://www.sesarju.eu/>.

[KOB12] Kober, J., Glisson, M., and M. Mistry, "Playing catch and juggling with a humanoid robot.", 2012, <https://doi.org/10.1109/HUMANOIDS.2012.6651623>.

[MAR19] Martinez, B., Cano, C., and X. Vilajosana, "A Square Peg in a Round Hole: The Complex Path for Wireless in the Manufacturing Industry", 2019, <https://ieeexplore.ieee.org/document/8703476>.

[Maurer2022]
Maurer, N., Ewert, T., Graupl, T., Schmitt, C., and S. Grundner-Culemann, "Security in Digital Aeronautical Communications A Comprehensive Gap Analysis", International Journal of Critical Infrastructure Protection, vol. 38 , 2022, <https://doi.org/10.1016/j.ijcip.2022.100549>.

[ODVA] http://www.odva.org/, "The organization that supports network technologies built on the Common Industrial Protocol (CIP) including EtherNet/IP.".

[PLA14] Plass, S., Hermenier, R., Luecke, O., Gomez Depoorter, D., Tordjman, T., Chatterton, M., Amirfeiz, M., Scotti, S., Cheng, Y.J., Pillai, P., Graeupl, T., Durand, F., Murphy, K., Marriott, A., and A. Zaytsev, "Flight Trial Demonstration of Seamless Aeronautical Networking", IEEE Communications Magazine, vol. 52, no. 5 , May 2014.

[PROFINET] http://us.profinet.com/technology/profinet/, "PROFINET is a standard for industrial networking in automation.", <http://us.profinet.com/technology/profinet/>.

[RFC7554] Watteyne, T., Ed., Palattella, M., and L. Grieco, "Using IEEE 802.15.4e Time-Slotted Channel Hopping (TSCH) in the Internet of Things (IoT): Problem Statement", RFC 7554, DOI 10.17487/RFC7554, May 2015, <https://www.rfc-editor.org/info/rfc7554>.

[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>.

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

[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>.

[RFC9030] 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>.

[RFC9055] Grossman, E., Ed., Mizrahi, T., and A. Hacker, "Deterministic Networking (DetNet) Security Considerations", RFC 9055, DOI 10.17487/RFC9055, June 2021, <https://www.rfc-editor.org/info/rfc9055>.

[RFC9317] Holland, J., Begen, A., and S. Dawkins, "Operational Considerations for Streaming Media", RFC 9317, DOI 10.17487/RFC9317, October 2022, <https://www.rfc-editor.org/info/rfc9317>.

[RFC9372] Mäurer, N., Ed., Gräupl, T., Ed., and C. Schmitt, Ed., "L-Band Digital Aeronautical Communications System (LDACS)", RFC 9372, DOI 10.17487/RFC9372, March 2023, <https://www.rfc-editor.org/info/rfc9372>.

Authors' Addresses

Carlos J. Bernardos (editor)
Universidad Carlos III de Madrid
Av. Universidad, 30
28911 Leganes, Madrid
Spain
Phone: +34 91624 6236
Email: cjbc@it.uc3m.es
URI: http://www.it.uc3m.es/cjbc/

Georgios Z. Papadopoulos
IMT Atlantique
Office B00 - 114A
2 Rue de la Chataigneraie
35510 Cesson-Sevigne - Rennes
France
Phone: +33 299 12 70 04
Email: georgios.papadopoulos@imt-atlantique.fr

Pascal Thubert
Cisco Systems, Inc
Building D
45 Allee des Ormes - BP1200
06254 MOUGINS - Sophia Antipolis
France
Phone: +33 497 23 26 34
Email: pthubert@cisco.com

Fabrice Theoleyre
CNRS
ICube Lab, Pole API
300 boulevard Sebastien Brant - CS 10413
67400 Illkirch
France
Phone: +33 368 85 45 33
Email: fabrice.theoleyre@cnrs.fr
URI: https://fabrice.theoleyre.cnrs.fr/