

DetNet WG
Internet-Draft
Intended status: Standards Track
Expires: 18 November 2026

CJ. Bernardos
UC3M
A. Mourad
InterDigital
17 May 2026

Extensions to enable wireless reliability and availability in multi-
access edge deployments
draft-bernardos-detnet-raw-mec-05

Abstract

There are several scenarios involving multi-hop heterogeneous wireless networks requiring reliable and available features combined with multi-access edge computing, such as Industry 4.0. This document describes solutions integrating IETF RAW and ETSI MEC, fostering discussion about extensions at both IETF and ETSI MEC to better support these scenarios.

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 18 November 2026.

Copyright Notice

Copyright (c) 2026 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
2. Terminology	4
3. Problem Statement	5
4. RAW and MEC integration	7
4.1. MEC app request for RAW	8
4.2. RAW OAM triggering MEC app migration	11
4.3. MEC OAM for RAW updates	13
5. IANA Considerations	15
6. Security Considerations	15
7. Acknowledgments	15
8. Informative References	15
Authors' Addresses	16

1. Introduction

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.

Multi-access Edge Computing (MEC) -- formerly known as Mobile Edge Computing -- capabilities deployed in the edge of the mobile network can facilitate the efficient and dynamic provision of services to mobile users. The ETSI ISG MEC working group, operative from end of 2014, intends to specify an open environment for integrating MEC capabilities with service providers' networks, including also applications from 3rd parties. These distributed computing capabilities will make available IT infrastructure as in a cloud environment for the deployment of functions in mobile access networks.

One relevant exemplary scenario showing the need for an integration of RAW and MEC is introduced next. One of the main (and distinct) use cases of 5G is Ultra Reliable and Low Latency Communications (URLLC). Among the many so-called "verticals" that require URLLC features, the Industry 4.0 (also referred to as Wireless for Industrial Applications) is probably the one with more short-term potential. As identified in [RFC9450], this scenario also calls for RAW solutions, as cables are not that suitable for the robots and mobile vehicles typically used in factories. This is also a very natural scenario for MEC deployments, as bounded, and very low latencies are needed between the robots and physical actuators and the control logic managing them. Figure 1 depicts an exemplary scenario of a factory where terminals (robots and mobile automated guided vehicles) are wirelessly connected. Control applications running in the edge (implemented as MEC applications) require bounded low latencies and a guaranteed availability, despite the mobility of the terminals and the changing wireless conditions. A heterogeneous wireless mesh network is used to provide connectivity inside the factory.

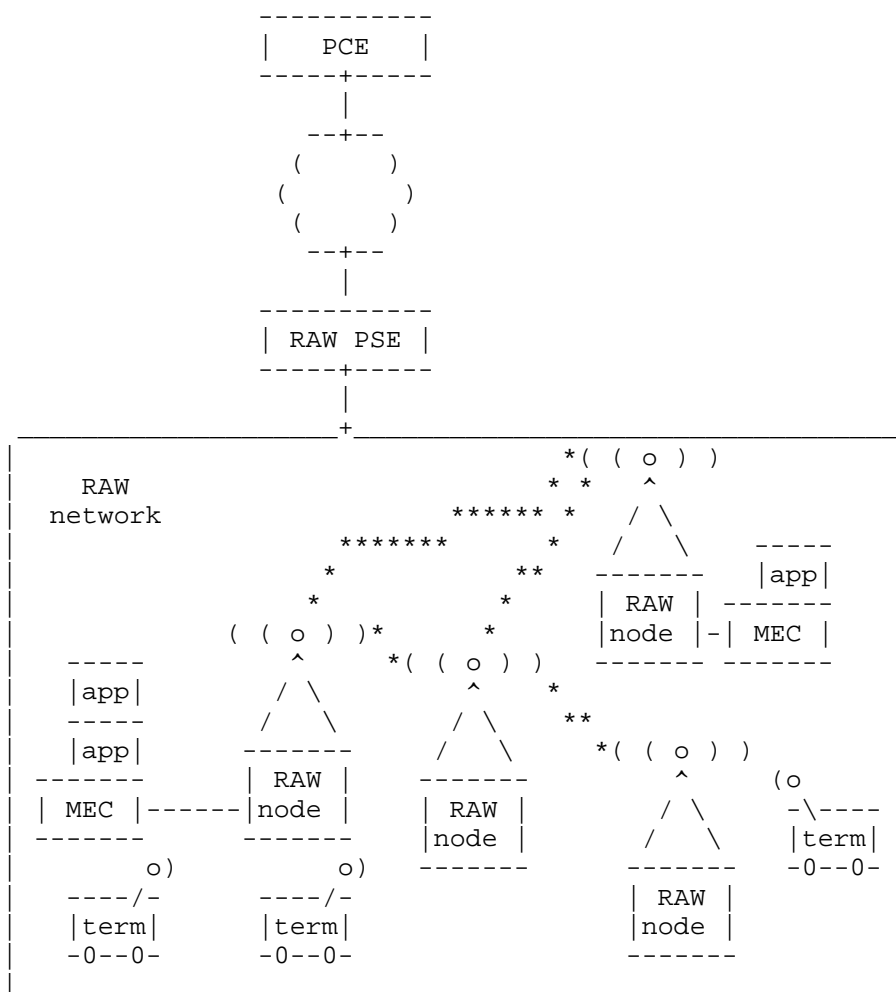


Figure 1: Exemplary scenario depicting MEC and RAW in an industrial environments

This scenario includes a wireless domain, involving multiple MEC platforms to ensure low latency to applications, by being able to host MEC applications in several locations, and dynamically migrate the apps as the terminals move around and the serving MEC platform might no longer be capable of meeting the latency requirements.

2. Terminology

The following terms used in this document are defined by the ETSI MEC ISG, and the IETF:

MEC host. The mobile edge host is an entity that contains a mobile edge platform and a virtualization infrastructure which provides compute, storage, and network resources, for the purpose of running mobile edge applications.

MEC platform. The mobile edge platform is the collection of essential functionalities required to run mobile edge applications on a particular virtualization infrastructure and enable them to provide and consume mobile edge services.

MEPM. MEC Platform Manager.

MEC applications. Mobile edge applications are instantiated on the virtualization infrastructure of the mobile edge host based on configuration requests validated by the mobile edge management.

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. Problem Statement

With current standards, the MEC platform(s) would have to interact with a Path Computation Element (PCE) for data plane requests and updates. This tremendously limits the capabilities to guarantee real-time forwarding decisions, as it will make it challenging and not possible to manage forwarding decisions in real or near-real time.

RAW solutions and approaches being explored today consider the role of the PSE, which computes at a short time scale which of the available paths (called tracks) -- computed by a PCE -- should be used per flow/packet and also which PAREO functions can be used, in order to provide the flow with the required availability and reliability features. The PSE interacts with the PCE and with the RAW nodes so they can setup the required per-flow state, to recognize the flow and determine the forwarding policy to be applied. These RAW forwarding decisions can be distributed among the necessary nodes

using in-band signaling (e.g., extending Segment Routing, BIER-TE or DETNET tagging) or can be taken autonomously by each forwarding node locally (based on its knowledge of the status of the network, gained via OAM RAW-specific mechanisms).

Figure 1 shows an exemplary scenario, depicting an industrial environment where different nodes are wirelessly connected to provide connectivity to several stationary and mobile terminals (i.e., robots). Industry environments are a good example of applications where reliability and availability are critical. Ensuring this in wireless heterogeneous and multi-hop networks requires using multiple paths, using PAREO functions and even using dual or multiple connectivity. Terminal mobility makes it even more challenging to guarantee certain reliability and availability levels, due to the dynamic and fast changes that this needs at the data plane level. The short-time scale forwarding decisions that are required to ensure reliability and availability with terminal mobility cannot be managed if MEC platforms can only interact with the data plane through the PCE.

The PCE is responsible for routing computation and maintenance in a network and it is typically a centralized entity that can even reside outside the network. It is meant to compute and establish redundant paths, but not to be sensitive/reactive to transient changes, and therefore is not capable of ensuring a certain level of reliability and availability in a wireless heterogeneous mesh network, especially if some of the nodes (e.g., the end terminals) might be mobile. With currently standardized solutions, a MEC platform could only interact with the RAW network through the PCE, most possibly through the Mp2 reference point defined by ETSI MEC. This reference point is defined between the MEC platform and the data plane of the virtualization infrastructure to instruct the data plane on how to route traffic among applications, networks, services, etc. This reference point is not further specified by ETSI MEC, but it would be the one that could be used by current solutions to allow for MEC to request the data plane on the RAW network a certain behavior (in terms of availability and reliability) for MEC app traffic flows. With existing solutions, the PCE would be the entry point for configuring and managing the RAW network, probably through the Mp2 reference point. Note that the PCE might reside outside the RAW network, the path between the RAW network and the PCE might be expensive and slow (e.g., it might need to traverse the whole RAW network) and reaching to the PCE can also be slow in regards to the speed of events that affect the forwarding operation at the radio layer.

Additionally, the MEC architecture as currently defined by ETSI does not have any component designed to deal with the specifics of an heterogeneous wireless multi-hop networks (such a RAW one), and

therefore, it is very limited in terms of what a MEC app (through the MEC platform) can request to the data plane of an heterogeneous wireless multi-hop network. Besides, this lack of RAW-aware component at the ETSI MEC architecture prevents any enhancement at either the MEC side (e.g., MEC app migrations triggered by RAW status updates) or the RAW side (e.g., PAREO function updates triggered by MEC app/terminal mobility).

Because of all these aforementioned reasons, it is needed to define a new RAW-enabled component at the ETSI MEC architecture, aimed at enabling a more direct interaction between the MEC platform and the RAW network, allowing the MEC platform to notify events and/or request actions to the RAW network quick enough. This involves some challenges, as the RAW PSE has to understand the needs from the running MEC applications, so it can properly configure the RAW nodes so the data plane provides the required reliability and availability.

4. RAW and MEC integration

This document defines a new entity inside the MEC platform: the RAW ctrl. This entity is responsible for computing what to instruct the RAW PSE, based on the requirements of the MEC apps, as well as to take decisions at the MEC side (e.g., migration of apps) based on information about the RAW network status.

As a result of the introduction of the RAW ctrl and the actions it is responsible of, new interactions and interface semantics are added. These interactions and semantics can be terminated at either the PCE or the RAW PSE, depending on the requirements of the MEC apps. For near real-time coordination and control between MEC and RAW mechanisms, the interactions are between the RAW ctrl and the RAW PSE. We mostly refer to this deployment model from now on, as it is the one that allow for near real-time updates on the forwarding plane, but note that an alternative deployment model in which the RAW ctrl interacts with the PCE is also possible, though only supporting non real-time interactions.

The MEC-RAW new interface semantics/extensions depicted in Figure 2 allows the MEC platform to issue requests to the RAW network, through the RAW PSE, so it can behave as required by MEC apps.

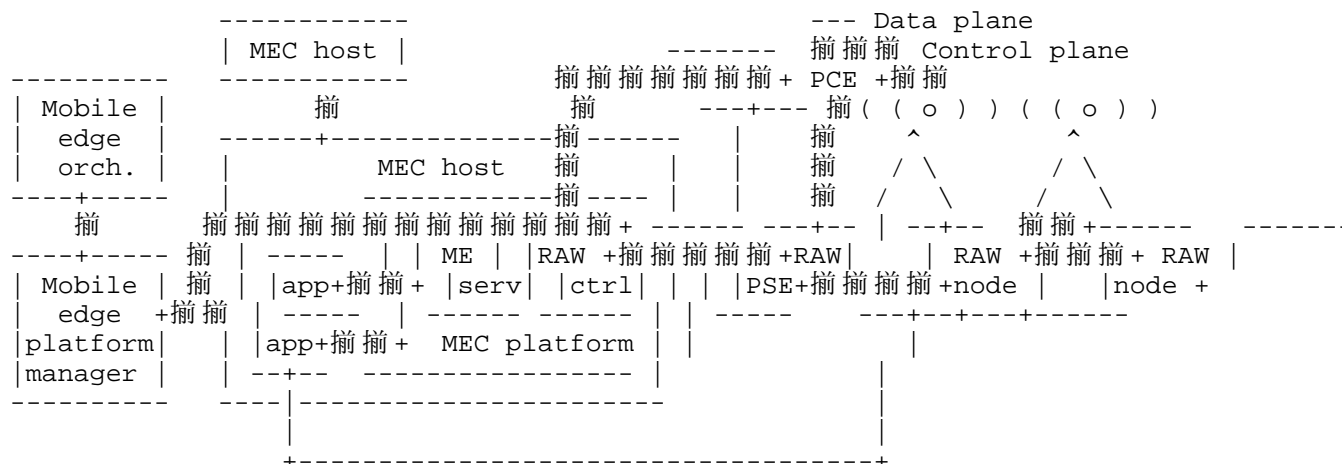


Figure 2: Block diagram

The new semantics of the interface between the MEC platform and the RAW PSE do not only serve to convey the requests, but also to synchronize the status and topology of the RAW (relevant portion of the) network, enabling to perform real-time or near-real time forwarding decisions. In the sequel, we show different exemplary signaling diagrams for the most relevant procedures.

4.1. MEC app request for RAW

Here we specify the interface extensions and signaling procedures needed to enable a MEC app describe and request its needs in terms of availability and reliability. As it will be further developed in other subsections, the wireless network conditions could also have an impact back on the MEC platform (e.g., by triggering the migration of the MEC app).

Figure 3 shows an exemplary signaling flow diagram, in which a certain MEC app request a given behavior for the treatment of the packets the app generates. We consider an industrial wireless application scenario, as the one used in previous sections, as an example for the sake of describing the interface and specified interactions.

The MEC platform is enhanced with a RAW ctrl entity, as introduced in the previous section. The RAW ctrl is a RAW-aware component within the MEC architecture that enables the required interactions with the RAW network, through the RAW PSE. This allows MEC apps to: (i) adapt to RAW conditions (e.g., if the requested reliability and availability is not possible), and (ii) dynamically modify the requested flow forwarding to the RAW network, based on the MEC app and mobility conditions.

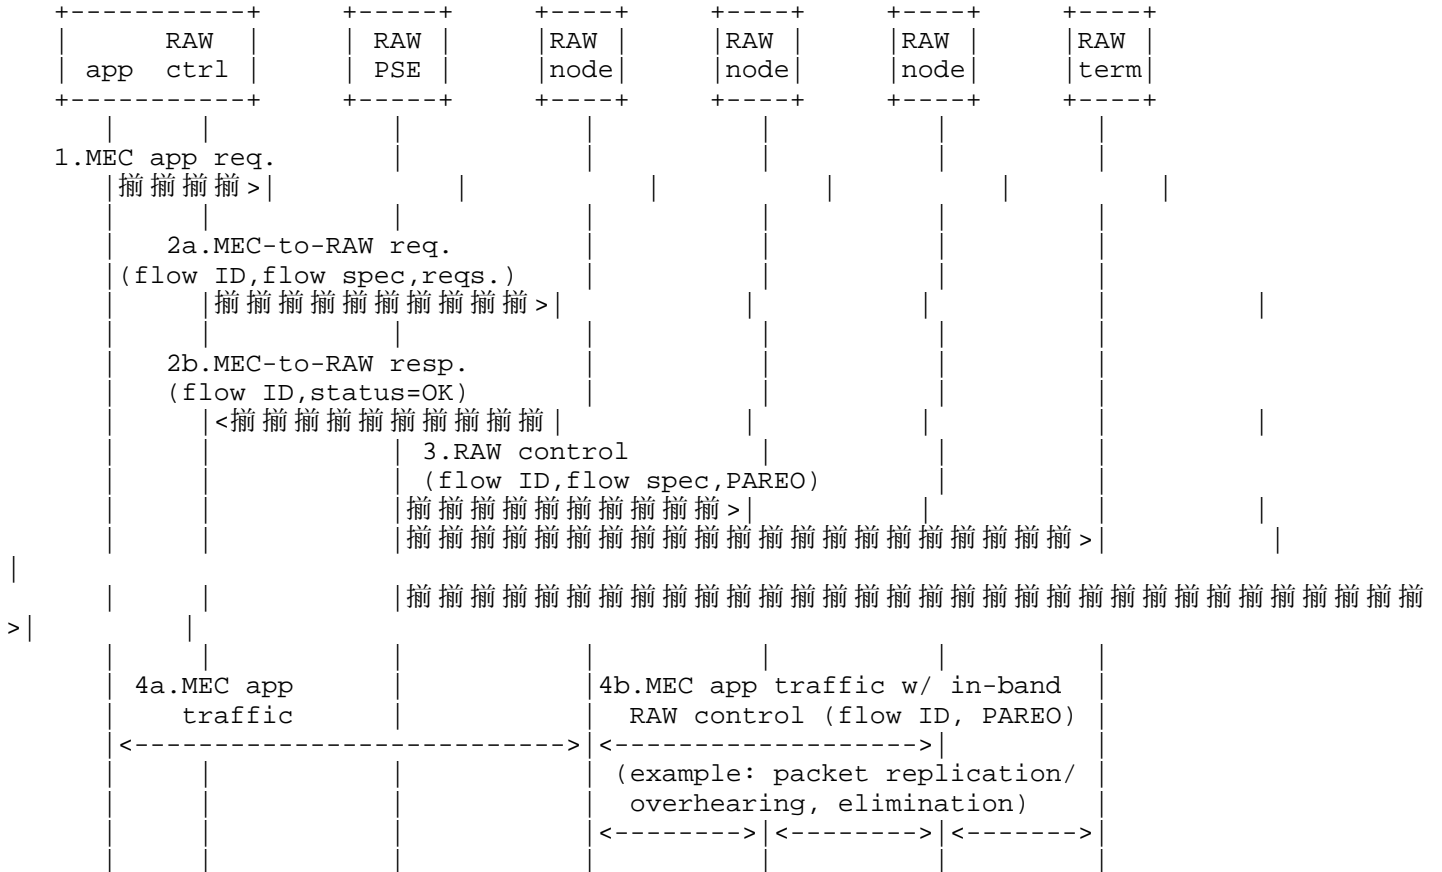


Figure 3: MEC app request for RAW

We next explain each of the steps illustrated in the figure:

1. A MEC app which is going to be consumed by a given terminal (or set of terminals, though in this example we show just one consumer), specifies to the MEC platform the characteristics of the traffic is going to generate and its associated requirements.

2. The MEC platform -- namely the RAW ctrl component, which is RAW-aware and knows the characteristics of the deployment -- analyzes the characteristics of the MEC app traffic and the provided requirements, and generates a new request -- over a new interface between the MEC platform and the RAW PSE -- that includes, among others, the following parameters:
 - * An ID for the given flow, which can be used for future near real-time update/configuration operations on the same flow.
 - * The flow specification, describing the characteristics of the packets, allowing an efficient identification of flow(s) based on well-known fields in IPv4, IPv6, and transport layer headers like TCP and UDP. An example of how to do this is defined in the Traffic Selectors for Flow Bindings [RFC6088].
 - * The requirements of the flow in terms of reliability and availability.
3. The RAW PSE processes the request, and based on its knowledge of the network (topology, node capabilities and ongoing flows) computes the best way of transmitting the packets over the RAW network (using the available paths/tracks, previously computed by a PCE). Note that at this point it might be possible that the RAW PSE realizes that it is not possible to provide the requested reliability and availability characteristics, and would report that back to the MEC platform (which might issue a new request with less requirements). The RAW PSE sends control packets to each of the RAW nodes involved, instructing how to deal with the packets belonging to the MEC app flow. This includes:
 - * assigning an ID to the flow;
 - * instructing the entry point in the RAW network to tag packets with that ID;
 - * specific PAREO functions to be used by each of the RAW nodes. This might be signaled to each of the RAW nodes, or just to some of them (e.g., only the entry point) who can then use in-band signaling to specify it.
4. The MEC app generates traffic (step 4a in the figure) which arrives at the RAW entry point in the network, which following the instructions of the RAW PSE, encapsulates and tags the packet, and might also include in-band signaling (e.g., using Segment Routing). Some PAREO functions are applied to the MEC app traffic packets (step 4b in the figure) to achieve the required levels of reliability and availability. In the figure,

as an example, packets are replicated (this could be done by means of wireless overhearing) at one point of the network, and then later duplicated packets eliminated.

4.2. RAW OAM triggering MEC app migration

Here we specify the mechanisms for MEC to benefit from RAW OAM information, for example, to trigger the migration of a MEC application to a different MEC platform, to ensure that the requirements of the MEC app continue to be met.

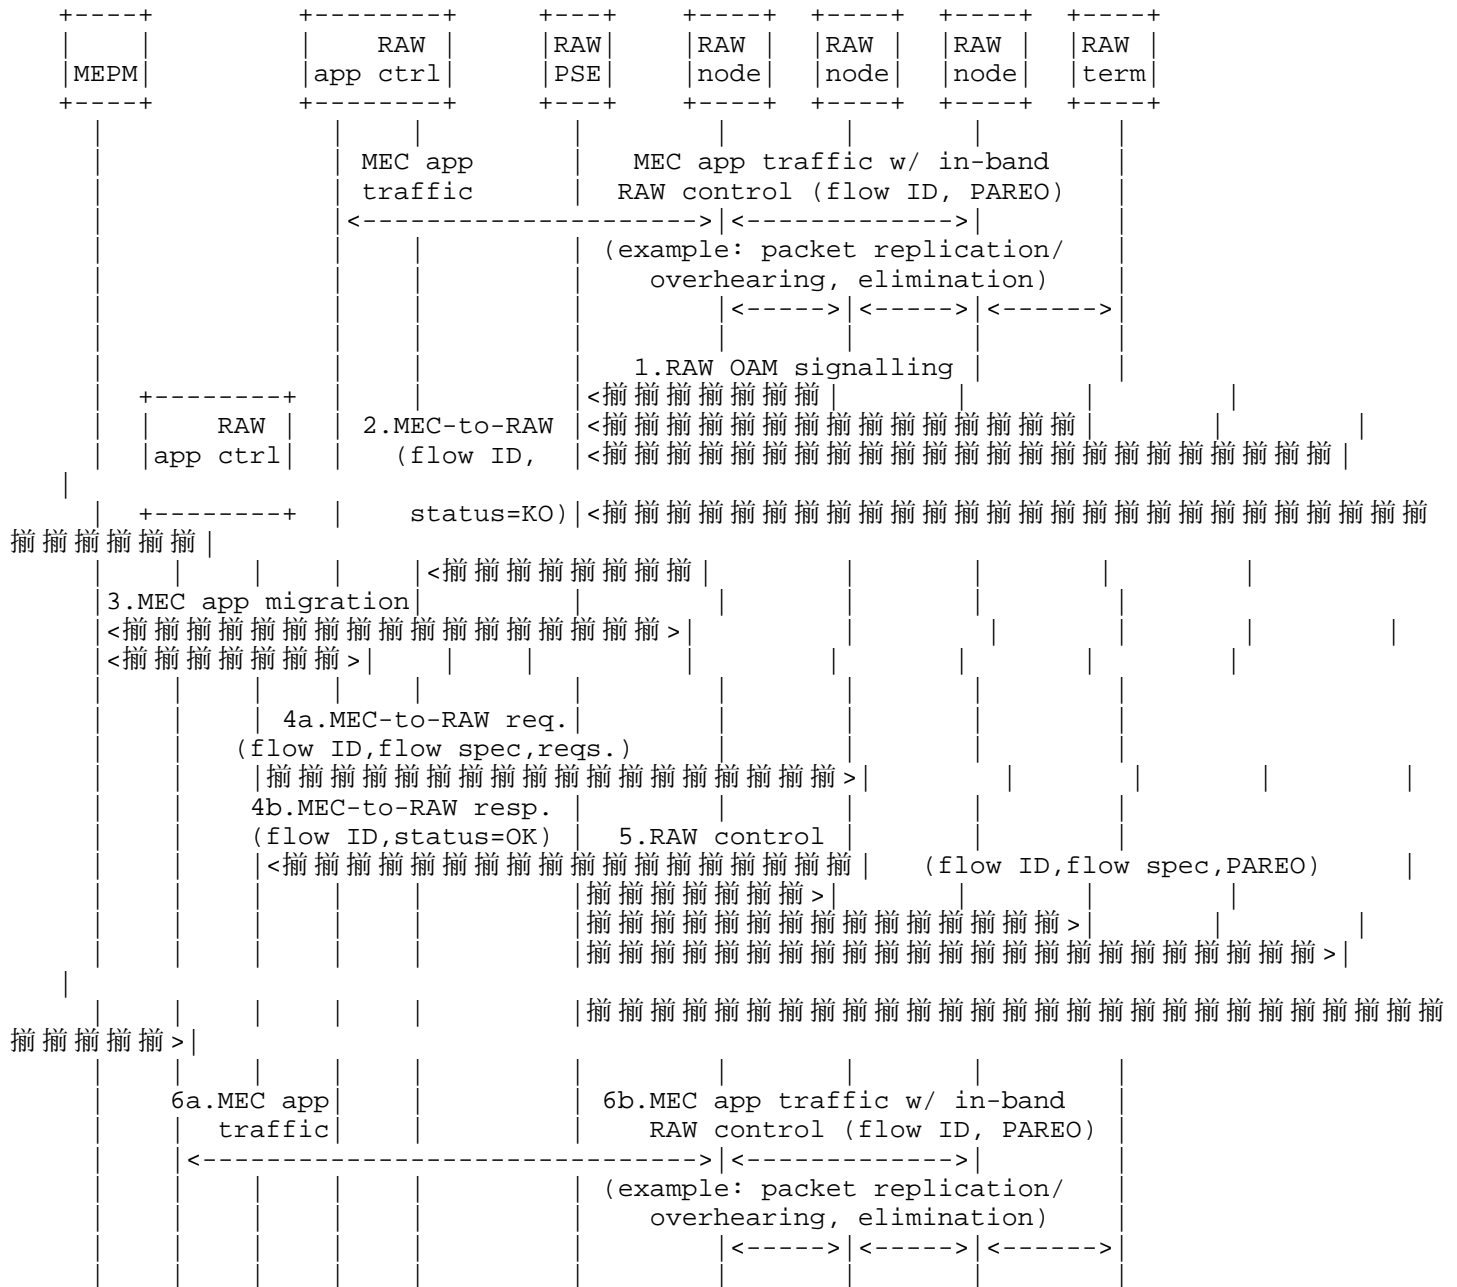


Figure 4: RAW OAM triggering MEC app migration

Figure 4 shows an exemplary signaling flow diagram, in which changes in the RAW network, detected thanks to RAW OAM, trigger the migration of a MEC app. We assume there is already a MEC app deployed and traffic is already flowing (i.e., all the procedures explained in the previous section took already place). We next explain each of the steps illustrated in the figure:

1. RAW OAM signaling is periodically and reactively exchanged between the RAW nodes and the RAW PSE [I-D.ietf-raw-oam-support].
2. If the conditions of the network get worse (e.g., because of changes in the radio propagation of a critical link), making it impossible to guarantee the required levels of reliability and availability, this generates a message from the RAW PSE to the MEC platform, indicating that a given MEC app flow can no longer be served.
3. The currently serving MEC platform triggers a MEC app migration to a different MEC platform. This involves the MEC platform manager. Note that the MEC platform might provide suggestions regarding where to migrate the MEC app, based on its knowledge of the RAW network status, acquired by the RAW ctrl through interactions with the PSE.
4. The same steps 2-3-4 specified in the procedure described in Section 4.1 take place. In this step, the MEC platform generates a new request to the RAW PSE.
5. The RAW PSE processes the request, and based on its knowledge of the network computes the best way of transmit the packets over the RAW network. The RAW PSE sends control packets to each of the RAW nodes involved.
6. The MEC app generates traffic, arriving at the RAW entry point in the network, which following the instructions of the RAW PSE, encapsulates and tags the packet.

4.3. MEC OAM for RAW updates

There are scenarios and situations where -- due to the mobility of the terminals or the nodes hosting the MEC platform hosting a given MEC app -- it might be needed to take actions on the RAW network: e.g., to update the paths, apply different PAREO functions, migrate the MEC app (thus involving a change in the RAW forwarding). This triggers the need for mechanisms enabling the RAW PSE to get and use MEC OAM information to update traffic forwarding at the RAW network as needed to adapt to varying conditions, e.g., due to node mobility.

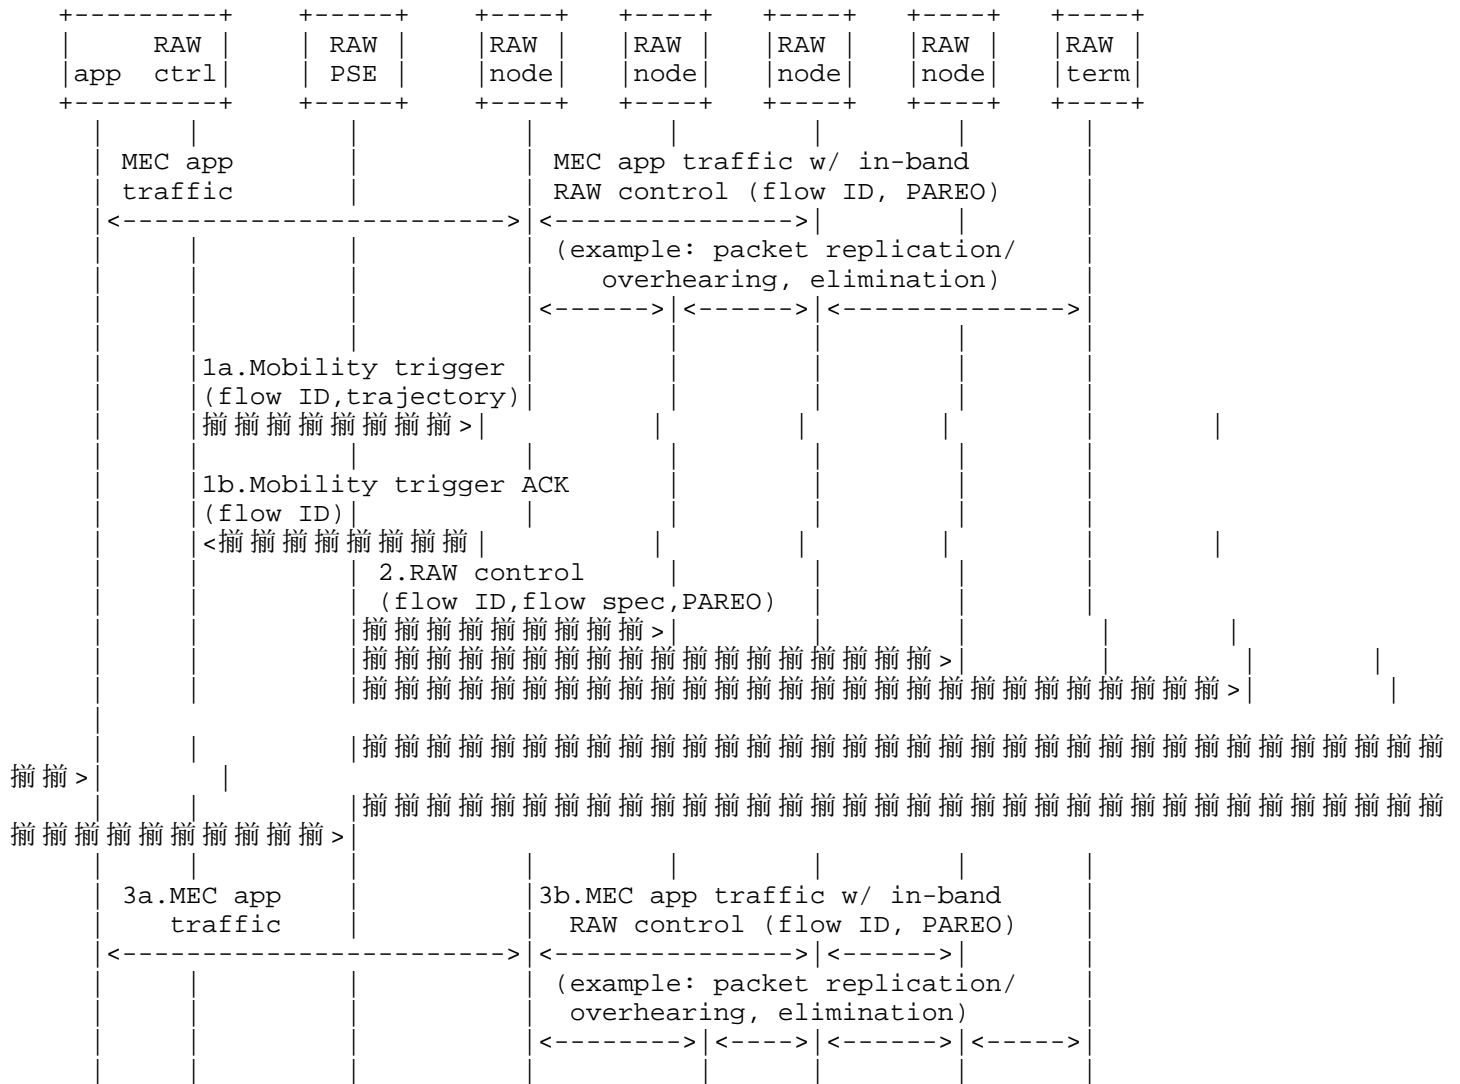


Figure 5: MEC OAM for RAW updates

Figure 5 shows an exemplary signaling flow diagram, in which the mobility of the a node (in this case the terminal, but it could have been the MEC platform hosting the MEC app) triggers the need for updating RAW forwarding mechanisms. As in the previous section, we assume there is already a MEC app deployed and traffic is already flowing (i.e., all the procedures explained in Section 4.1 took already place). We next explain each of the steps illustrated in the figure:

1. The MEC platform notifies that the terminal consuming the MEC app is moving (note that it other events can be notified, such as the mobility of the MEC platform itself), including the expected trajectory (if can be known or predicted in advance, as it will be the case in most cases in several scenarios, such as industrial use cases).
2. The RAW PSE uses this information to re-compute how to best provided the reliability and availability needed by the MEC app traffic flow, and updates the RAW nodes about the PAREO functions that they have to apply.
3. After this, traffic from the MEC app benefits from updated policies, computed according to the new conditions, and ensuring that the requirements of the MEC app continue to be fulfilled.

5. IANA Considerations

TBD.

6. Security Considerations

TBD.

7. Acknowledgments

The work of Carlos J. Bernardos in this document has been partially supported by the UNICO I+D 6G-DATADRIVEN-04 project (TSI-063000-2021-132).

8. Informative References

[I-D.ietf-raw-architecture]

Thubert, P., "Reliable and Available Wireless Architecture", Work in Progress, Internet-Draft, draft-ietf-raw-architecture-30, 25 July 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-raw-architecture-30>>.

[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-07, 19 January 2026, <<https://datatracker.ietf.org/doc/html/draft-ietf-raw-oam-support-07>>.

- [RFC6088] Tsirtsis, G., Giarreta, G., Soliman, H., and N. Montavont, "Traffic Selectors for Flow Bindings", RFC 6088, DOI 10.17487/RFC6088, January 2011, <<https://www.rfc-editor.org/info/rfc6088>>.
- [RFC9450] Bernardos, CJ., Ed., Papadopoulos, G., Thubert, P., and F. Theoleyre, "Reliable and Available Wireless (RAW) Use Cases", RFC 9450, DOI 10.17487/RFC9450, August 2023, <<https://www.rfc-editor.org/info/rfc9450>>.

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/>