

RAW WG
Internet-Draft
Intended status: Standards Track
Expires: 18 October 2026

CJ. Bernardos
UC3M
A. Mourad
InterDigital
16 April 2026

DetNet multidomain extensions
draft-bernardos-detnet-raw-multidomain-07

Abstract

This document describes the multi-domain DetNet problem, starting from a wireless DetNet (formerly called RAW) scope, and explores and proposes some extensions to enable DetNet 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 18 October 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 and Problem Statement	2
2. Terminology	7
3. Multi-domain extensions	8
3.1. RAW specific scenario	8
3.2. RAW-DetNet specific scenario	11
4. IANA Considerations	14
5. Security Considerations	14
6. Acknowledgments	14
7. Informative References	14
Authors' Addresses	15

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 [RFC9450] 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).

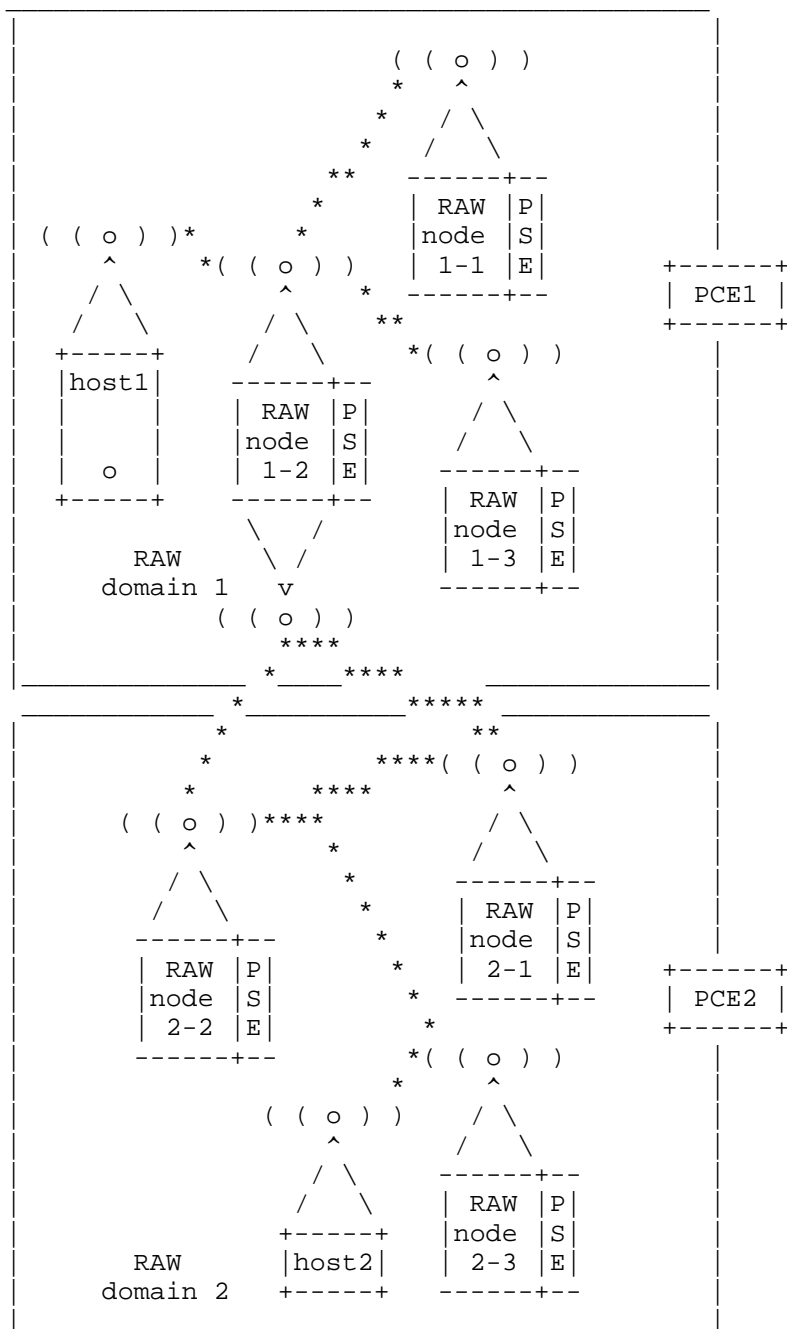


Figure 1: Exemplary scenario showing multiple RAW domains

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 domain's 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.

Note that while the figure and text above was referring to wireless (aka RAW) domains, the scope of this document includes also wireless domains, in different combinations. For example, we could consider a wireless domain connected to a wired domain, in a way that requires a host connected to one domain to have a deterministic communication with a host connected to the other domain, such as illustrated in Figure 2.

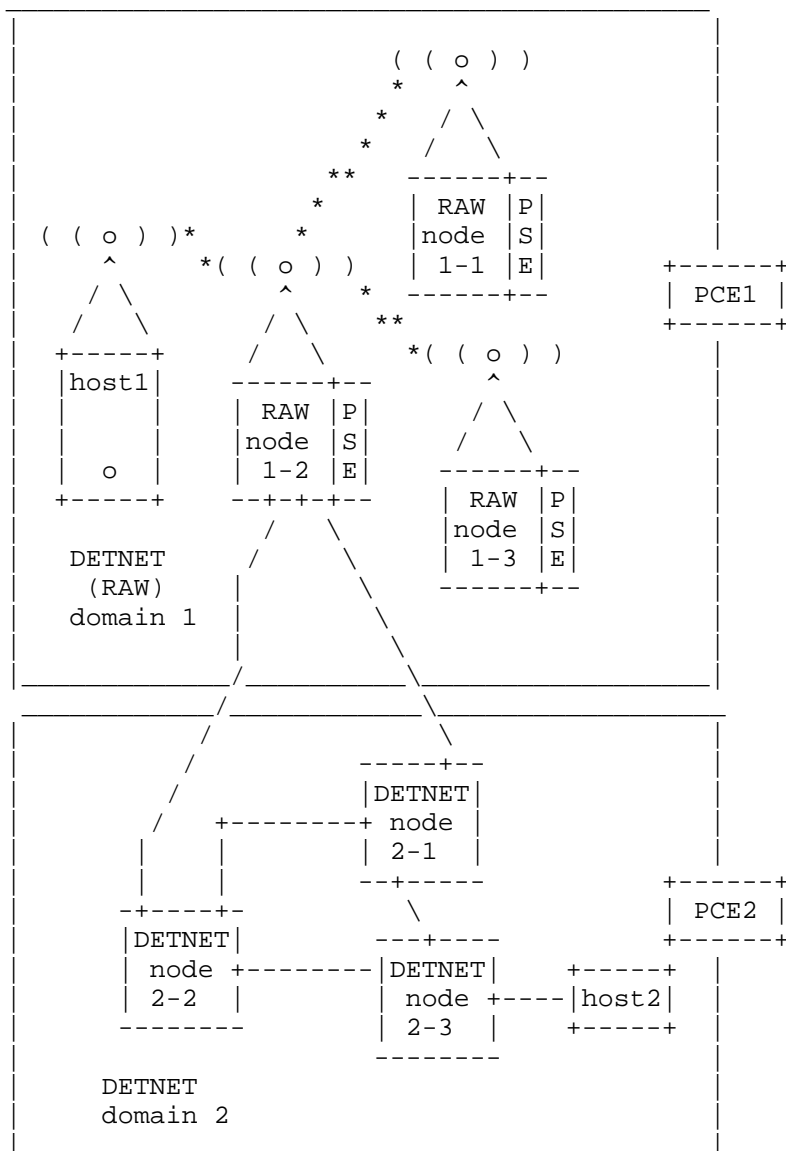


Figure 2: Exemplary scenario showing multiple wireless and wired domains

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. NOTE: this document is still using the old terminology, referring to PSE, instead of PLR. This will be updated in future revisions of the document.

3. Multi-domain extensions

3.1. RAW specific scenario

In this section we address the RAW (wireless multi-domain) scenario shown in Figure 1. Next, we specify the new mechanisms and signaling extensions to enable inter-domain RAW connectivity.

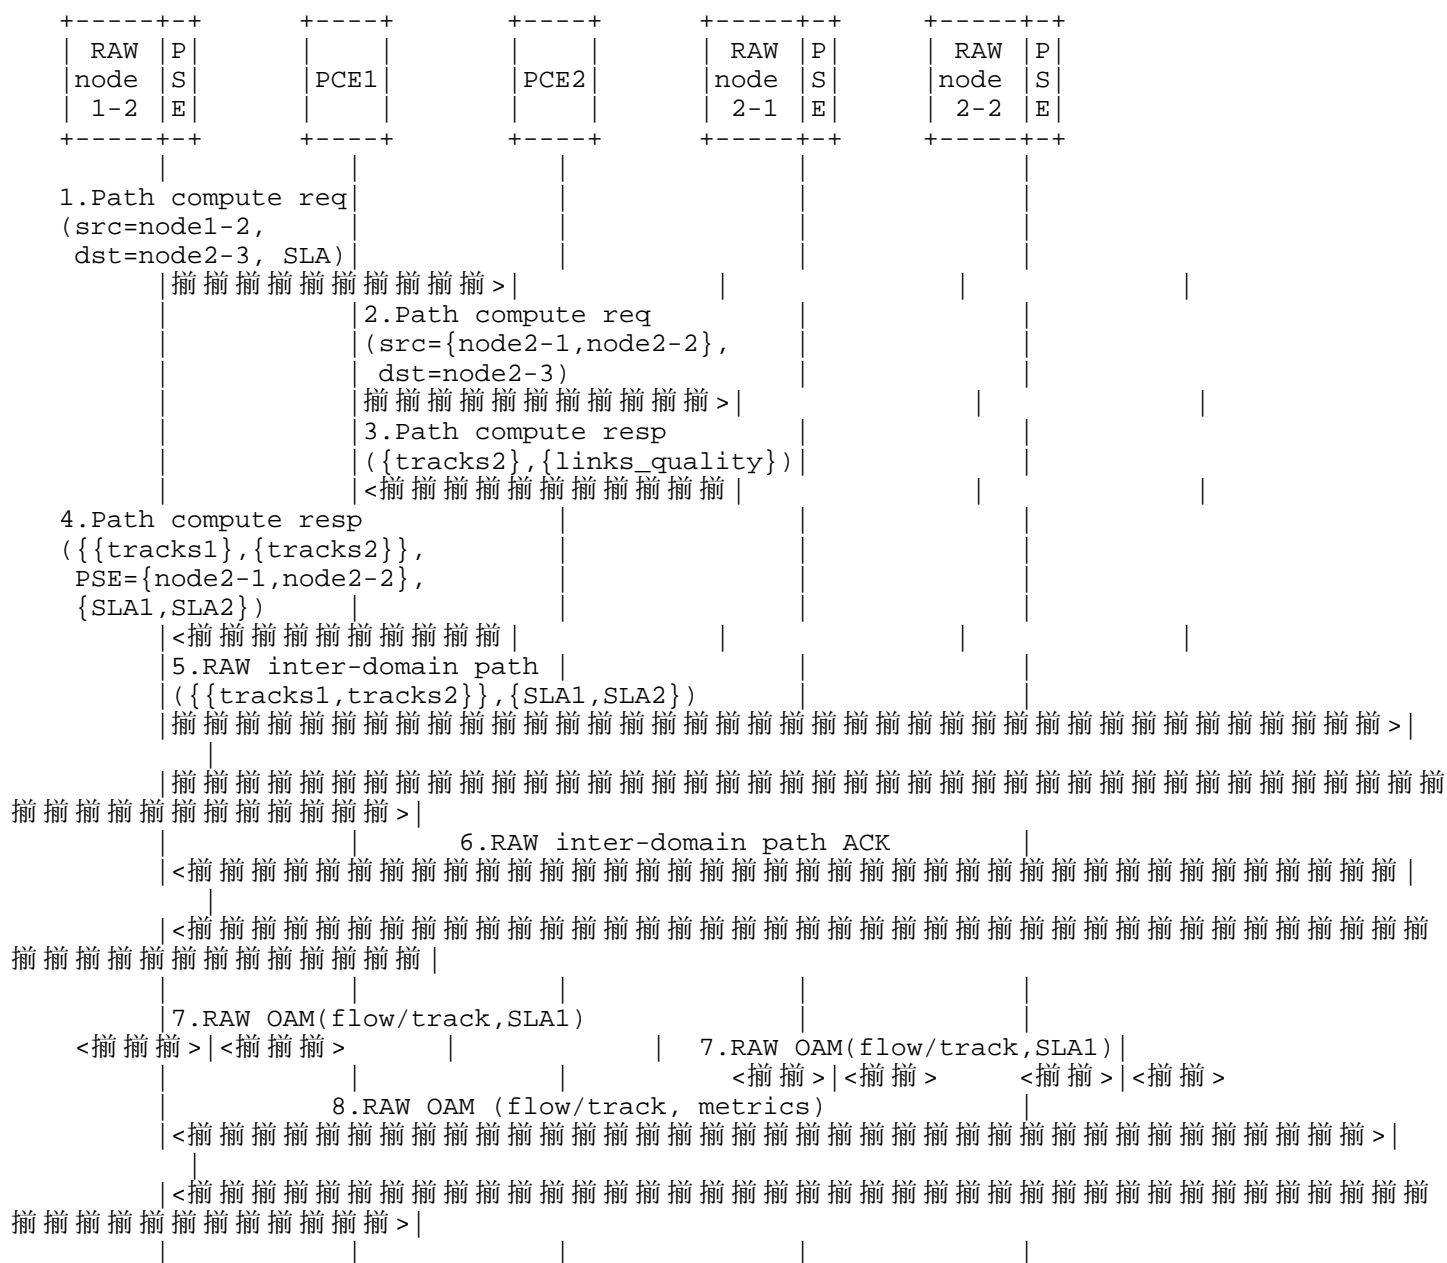


Figure 3: Multi-domain RAW signaling

Figure 3 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 hosts. 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.

3.2. RAW-DetNet specific scenario

In this section we address a more generic RAW-DetNet scenario, as shown in Figure 2. Next, we specify the mechanisms and signaling extensions to enable inter-domain DetNet connectivity, involving both wireless and wired domains.

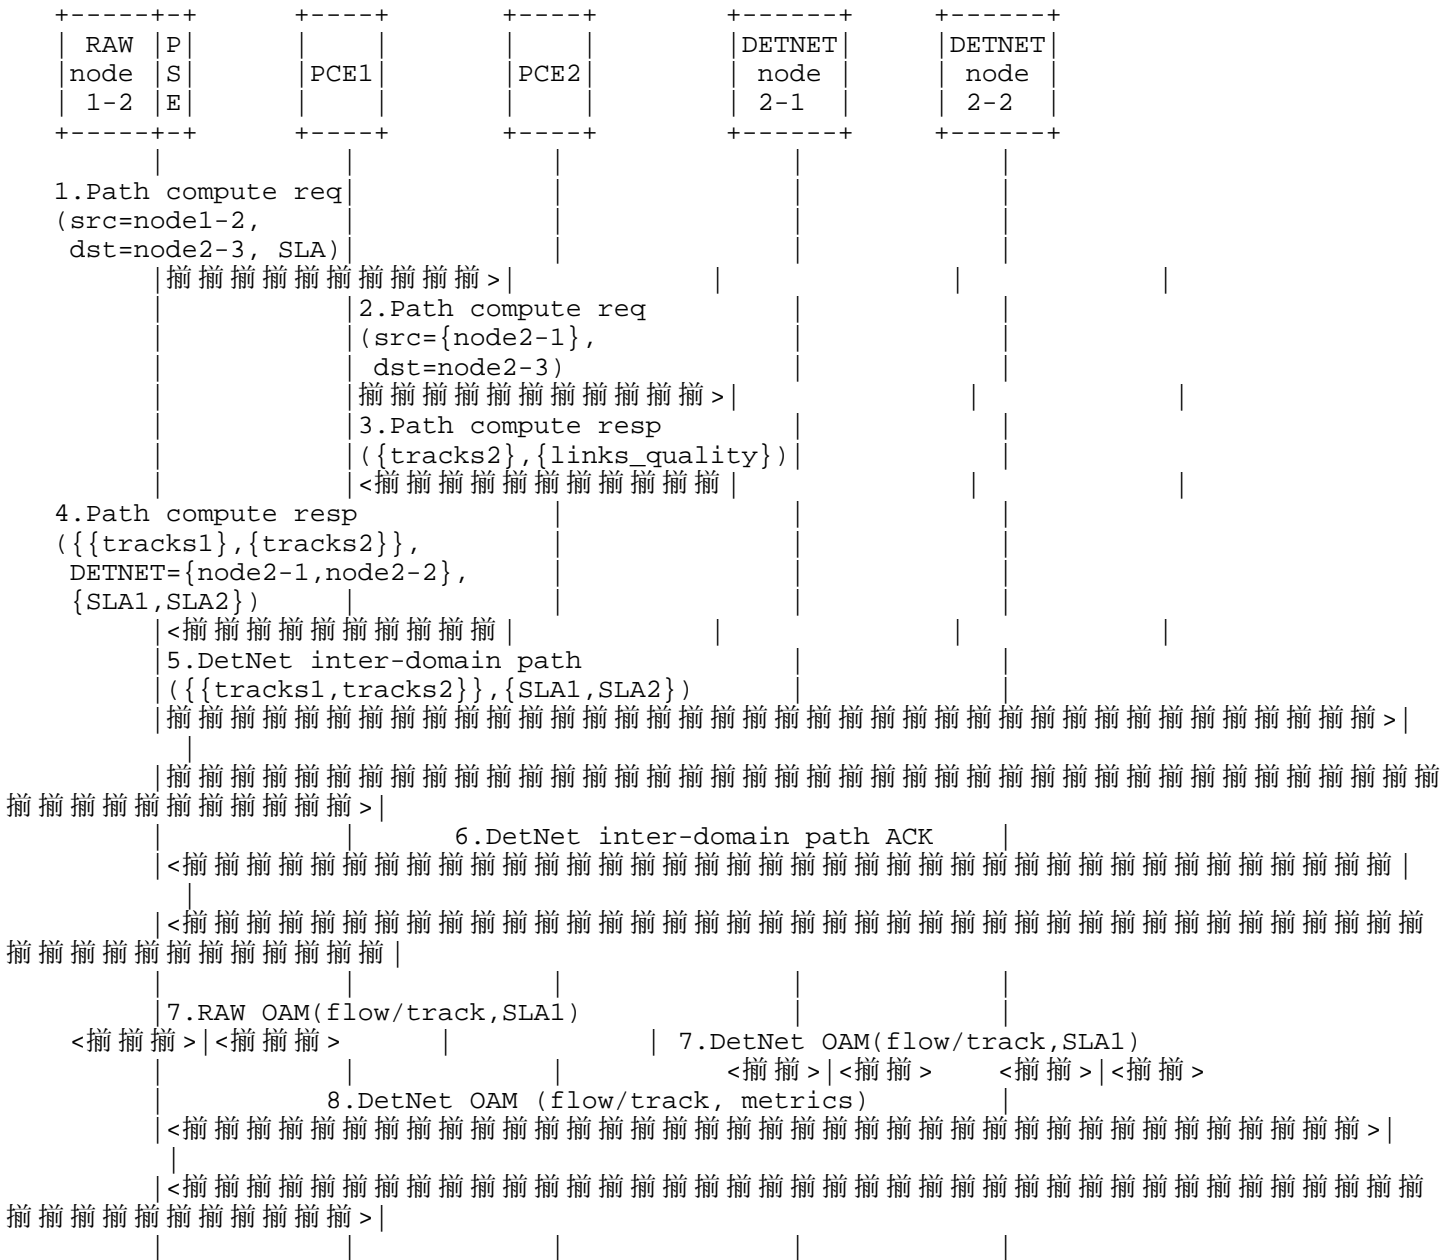


Figure 4: Multi-domain DetNet signaling

Figure 4 shows a signaling flow diagram, taking as baseline scenario the one shown in Figure 2, 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 DetNet node (node2-3) and the desired SLA in terms of reliability and availability. The source and/or destination RAW nodes might be hosts. 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 (node2-1 and node2-2). How this is done is outside of the scope of this document. PCE1 requests to PCE2 to compute the available tracks from node2-1 and node2-2 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 ingress nodes in domain2 that are connected to domain1.
5. The PSE@domain1 sends a message to each ingress node in 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 ingress DetNet nodes in the 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 and ingress DetNet nodes, at each domain, start performing OAM procedures [I-D.ietf-raw-oam-support] [RFC9551], 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 and ingress DetNet nodes 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/DetNet node would notify the domain PCE and the other PSEs/DetNet nodes, so a reaction measure can be taken (e.g., selecting different subtracks, taking different PAREO/PREOF 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), DESIRE6G (Grant 101096466), and UNICO I+D 6G-DATADRIVEN-04 project (TSI-063000-2021-132).

7. 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>>.

- [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>>.
- [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>>.
- [RFC9551] Mirsky, G., Theoleyre, F., Papadopoulos, G., Bernardos, CJ., Varga, B., and J. Farkas, "Framework of Operations, Administration, and Maintenance (OAM) for Deterministic Networking (DetNet)", RFC 9551, DOI 10.17487/RFC9551, March 2024, <<https://www.rfc-editor.org/info/rfc9551>>.

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/>