DETNET WG
Internet-Draft
Intended status: Standards Track
Expires: 19 April 2026

C.J. Bernardos, Ed.
Universidad Carlos III de Madrid
L. Contreras
Telefonica
Q. Xiong
ZTE Corporation
A. Mourad
InterDigital
16 October 2025

A PCE-based Control Plane Framework for Multi-Domain Deterministic Networking (DetNet)

draft-bernardos-detnet-multi-domain-pce-00

Abstract

Deterministic Networking (DetNet) provides the capability to carry specified unicast or multicast data flows for real-time applications with extremely low data loss rates and bounded latency over a path or network. As DetNet deployments expand, they will inevitably need to span multiple domains that may be under separate administrative or technological control. This creates a need for a control plane solution that can establish and maintain end-to-end DetNet services across these domain boundaries.

This document defines a framework for a Path Computation Element (PCE)-based control plane for multi-domain DetNet. It first establishes a working definition of a "DetNet Domain" for the purpose of path computation and control. It then describes two high-level architectural approaches for inter-domain path computation and resource reservation: a Hierarchical PCE model and a peer-to-peer PCE "stitching" model. This framework provides the foundation for more specific work on multi-domain DetNet solutions.

Status of This Memo

Table of Contents

1. Introduction 3
2. Conventions and Terminology 3
3. Defining a DetNet Domain 4
  3.1. Domain Characteristics 4
  3.2. Scope of a DetNet Domain 4
4. PCE-based Multi-Domain DetNet Architectures 5
  4.1. Exemplary Use Case 5
  4.2. Problem Statement 6
  4.3. Hierarchical PCE (H-PCE) Approach 6
  4.4. Peer-to-Peer (Stitching) PCE Approach 7
5. Multi-Domain DetNet Flow Considerations 7
  5.1. End-to-End Path Computation 7
  5.2. Resource Management 8
  5.3. End-System awareness 8
  5.4. Flow Aggregation 8
6. Security Considerations 9
7. IANA Considerations 9
8. Acknowledgments 9
9. References 9
  9.1. Normative References 9
  9.2. Informative References 9
Authors' Addresses 11

1. Introduction

The Deterministic Networking (DetNet) architecture, as defined in [RFC8655], provides a service for flows requiring bounded latency, and/or extremely low packet loss, and/or reliable service. The initial focus of DetNet has largely been on single-domain networks, where a single controller or administrative entity has full visibility and control over all network resources.

However, many use cases, such as industrial automation, professional audio/video, and smart grids, require deterministic connectivity that spans multiple networks. These networks may be operated by different providers (administrative domains), utilize different underlying link-layer technologies (technological domains), or be structured as separate control areas for scalability.

To support such scenarios, a control plane framework is needed to coordinate the establishment of end-to-end DetNet paths across these domain boundaries. The Path Computation Element (PCE) Communication Protocol (PCEP) [RFC5440] provides a standard mechanism for a PCE to compute paths and a Path Computation Client (PCC) to request them. This makes PCE a suitable candidate for building a multi-domain DetNet control plane.

This document builds on the DetNet Controller Plane Framework [I-D.ietf-detnet-controller-plane-framework] by focusing specifically on multi-domain challenges. It proposes a foundational framework by:

* Defining what constitutes a "domain" in the context of DetNet path computation.
* Describing high-level PCE-based architectures for managing multi-domain paths.
* Identifying key considerations for establishing and maintaining DetNet flows across these domains.

The goal is to establish the necessary foundational concepts before addressing specific technology implementations, such as multi-domain RAW (Reliable and Available Wireless) [I-D.bernardos-detnet-raw-multidomain].

2. Conventions and Terminology

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.

This document uses the terminology defined in [RFC8655], [RFC4655] and [RFC5440].

3. Defining a DetNet Domain

For the purpose of multi-domain DetNet control, a clear definition of a "domain" is essential. A domain represents a collection of network resources (nodes, links) that are managed and controlled as a single entity for the purpose of DetNet path computation and resource allocation.

3.1. Domain Characteristics

A DetNet domain is characterized by a set of network nodes that are subject to a single, consistent set of DetNet control and management policies. From a PCE-based control plane perspective, this typically implies that:

* A single PCE instance (or a coordinated set of redundant PCEs) has complete topological visibility within the domain.
* This PCE instance is responsible for computing paths and managing the allocation of DetNet-specific resources (e.g., buffer space, link schedules, queue reservations) for all nodes within that domain.
* There is a trusted relationship and a secure communication channel between the PCE and all the nodes it controls within the domain.

3.2. Scope of a DetNet Domain

The boundaries of a DetNet domain can be defined based on several factors, which may overlap:

Administrative Domain: A set of network elements under the control of a single network operator or administrative entity. This is the most common interpretation. Inter-domain communication occurs when a path must cross from one operator’s network to another’s.

PCE Control Domain: A domain is defined as the set of nodes controlled by a single PCE instance. This is the primary definition used within this framework. A large administrative domain might be divided into multiple smaller PCE control domains for scalability.

Technological Domain: A domain could be defined by the consistent use of a specific data plane technology (e.g., a TSN domain, an 3GPP 5G domain) or queuing mechanism (e.g. queuing solutions within the categories as per [I-D.ietf-detnet-dataplane-taxonomy]). While paths may cross technological boundaries, this document posits that this does not inherently define a control plane domain boundary. A single PCE

SHOULD be capable of managing a domain comprising multiple technologies. Similarly, the specific queuing mechanisms (e.g., [RFC9016]) supported by devices do not define a domain boundary; a single domain can contain devices supporting multiple queuing solutions, which can be used concurrently. PCE needs to select a specific queuing mechanism along the path for a DetNet flow within each domain.

4. PCE-based Multi-Domain DetNet Architectures

4.1. Exemplary Use Case

Let’s consider the scenario depicted in the figure below, where a DetNet flow is established between a source S and a destination D. The path for the flow traverses three different domains. Domain 1 is a wired domain, which could for example be a TSN-based DetNet MPLS [RFC8964] or DetNet IP [RFC8939] network. Domain 2 is a wireless (RAW) domain. Domain 3 is again a wired domain. The RAW domain provides connectivity between the two wired domains. Note that this is just an example, and other combinations of wired/wireless domains could exist (e.g., a DetNet flow traversing a wired domain providing connectivity between two RAW domains).

```
                      .-------------------------------------.
                      |             Parent PCE              |
                      '-------------------------------------'
                             ^      |      ^
                             |      |      |
                             v      v      v
        .-----------------.  .-----------------.  .-----------------.
        | Child PCE (d1) |   | Child PCE (d2) |   | Child PCE (d3) |
        '-----------------'  '-----------------'  '-----------------'
              | |                  | |                  | |
     S ---- R1 ======== R2 -------- R3 ======== R4 -------- R5 ---- D
           <-- Domain 1 -->   <---- Domain 2 ---->   <-- Domain 3 -->
              (wired)               (RAW)                (wired)

          S, D: end-systems (source and destination)
          Rx: DetNet routers/bridges
          ==: wired link
          --: wireless link
```

Figure 1: Exemplary multi-domain scenario

Each domain has its own PCE, which is responsible for path computation and resource management within the domain. These are referred to as Child PCEs (C-PCEs). Routers R2 and R3 are border routers of Domain 2 (RAW), and R3 and R4 are border routers of Domain

2 and 3, respectively. A Parent PCE (P-PCE) is responsible for the end-to-end path computation and orchestration among the different C-PCEs.

4.2. Problem Statement

In a multi-domain environment, no single PCE has end-to-end visibility of the full network topology. The challenge is to compute an end-to-end path that meets the strict latency, jitter, and loss requirements of a DetNet flow, while respecting the administrative and confidentiality boundaries of each participating domain.

Each domain’s PCE is responsible for its own internal path computation and resource allocation. The multi-domain architecture must define how these individual PCEs cooperate to create a seamless end-to-end service. Two primary models are considered: Hierarchical PCE and Peer-to-Peer PCE.

4.3. Hierarchical PCE (H-PCE) Approach

The Hierarchical PCE (H-PCE) architecture [RFC6805] defines a parent-child relationship between PCEs.

* A Parent PCE (P-PCE) has a partial, abstracted view of the child domains. It does not see the detailed topology within each child domain but knows the reachability and characteristics (e.g., available latency budget, cost) of paths to and through them.
* A Child PCE (C-PCE) has full visibility of its own domain’s topology and resources. It is responsible for all intra-domain path computations.

In a multi-domain DetNet context:

1. A request for an end-to-end DetNet path is sent to the P-PCE. This request includes the source, destination, and required QoS parameters (e.g., maximum latency).
2. The P-PCE computes a high-level, domain-sequence path. This path is a sequence of domains that the flow must traverse, along with the entry and exit boundary nodes for each domain.
3. The P-PCE then sends requests to the C-PCE of each domain in the path sequence. Each request asks for an intra-domain path segment between the specified entry and exit nodes that meets a portion of the end-to-end QoS requirements.
4. Each C-PCE computes its path segment, reserves the necessary resources locally, and reports success or failure back to the P-PCE.

5. If all C-PCEs are successful, the P-PCE confirms the end-to-end path. If any C-PCE fails, the P-PCE may attempt to find an alternate domain path.

The Parent PCE (P-PCE) would be responsible for computing the multi-domain path based on an abstracted topology of the different domains. The Child PCEs (C-PCEs) are responsible for the path computation in their own domains. A C-PCE would be aware of the specific technologies used in its domain (e.g., RAW, DetNet IP, DetNet MPLS, etc.), being able to compute a path taking into account the specific constraints of the technology. For instance, in a RAW domain, the C-PCE would be able to select the path, the schedule and the links to be used to guarantee a certain level of reliability.

4.4. Peer-to-Peer (Stitching) PCE Approach

In a peer-to-peer approach, also known as "stitching," there is no parent-child hierarchy. PCEs from adjacent domains cooperate as peers. The path computation is performed sequentially from one domain to the next. This model is described in [RFC5441] for inter-area and inter-AS TE path computation.

In a multi-domain DetNet context:

1. The PCE in the source domain (PCE-1) receives a path request for a flow destined for another domain.
2. PCE-1 computes a path from the source node to a suitable exit border node in its domain.
3. PCE-1 then sends a PCEP request to the PCE of the adjacent domain (PCE-2), specifying the entry border node (which is the exit node from domain 1) and the final destination. The request includes the remaining QoS budget.
4. PCE-2 computes a path through its domain to either the final destination (if it’s in domain 2) or to another suitable exit border node. It then "stitches" this segment to the previous one.
5. This process repeats until the PCE in the destination domain is reached. The path is confirmed backward along the chain of PCEs.

5. Multi-Domain DetNet Flow Considerations

5.1. End-to-End Path Computation

The end-to-end path is a concatenation of intra-domain path segments. The total latency and other QoS metrics are cumulative. The control plane must be able to allocate the end-to-end budget among the participating domains.

In hierarchical PCE, the P-PCE needs to collect the domain-specific information from C-PCEs and the P-PCE will divide the end-to-end budget of a DetNet flow into sub-budgets to several domains based on the capabilities (e.g. latency, jitter) within each domain.

In stitching PCE, the end-to-end budget of a DetNet will be divided from the source PCE, then to an adjacent domain, till to the destination PCE. The PCE within each domain needs to compute the latency bound as per [RFC9320] considering the bounded latency metric.

5.2. Resource Management

Resources MAY be reserved in each domain for the flow. If any domain in the path cannot provide the required resources, the end-to-end path setup fails. A mechanism for transactional, all-or-nothing resource commitment across domains is highly desirable.

The control plane also needs to advertise inter-domain resource information, including bandwidth, delay, jitter with related queuing mechanisms for QoS coordination.

5.3. End-System awareness

A critical aspect is whether the end-systems (source and destination) are DetNet-aware.

DetNet-aware End-Systems: The end-systems can signal their QoS requirements and participate in the DetNet control plane.

DetNet-unaware End-Systems: The requirements for these systems must be configured at the edge of the DetNet domain by a proxy or network management system. In a multi-domain scenario, the entry node of the first DetNet domain acts as this ingress point.

5.4. Flow Aggregation

Flow aggregation is recommended in the multi-domain scenario to achieve the end-to-end QoS guarantees for aggregated flow(s) that span across multiple domains. Multiple flows may be aggregated in a domain and disaggregated in another domain. The network parameters of an aggregated flow should be exchanged among different network domains. The path computation should consider to identify the end-to-end budget of the aggregated flow which should cover the requirements of all member flows.

6. Security Considerations

Multi-domain operations introduce significant security challenges. The communication between PCEs in different domains MUST be secured, ensuring authentication, integrity, and confidentiality. Each domain must be protected from misbehaving or compromised peer domains.

Topology and resource information exposed by a domain's PCE to an external entity (a parent PCE or a peer PCE) is a sensitive matter. The framework must allow for policy-based control over the level of abstraction and detail that is shared.

Considerations from [RFC8253] also applies.

7. IANA Considerations

This document makes no requests of IANA.

8. Acknowledgments

The work of Carlos J. Bernardos in this document has been partially supported by the Horizon Europe DESIRE6G (Grant 101096466), and UNICO I+D 6G-DATADRIVEN-04 project (TSI-063000-2021-132).

9. References

9.1. Normative References

[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>.

[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>.

9.2. Informative References

[I-D.bernardos-detnet-raw-multidomain]
Bernardos, C., et al., "Reliable and Available Wireless (RAW) based on Multi-link and Multi-domain strategies", Work in Progress, Internet-Draft, draft-bernardos-detnet-

raw-multidomain-06,
<https://datatracker.ietf.org/doc/html/draft-bernardos-detnet-raw-multidomain-06>.

[I-D.ietf-detnet-controller-plane-framework]
Saad, T., et al., "Deterministic Networking (DetNet) Controller Plane Framework", Work in Progress, Internet-Draft, draft-ietf-detnet-controller-plane-framework, <https://datatracker.ietf.org/doc/html/draft-ietf-detnet-controller-plane-framework>.

[I-D.ietf-detnet-dataplane-taxonomy]
Joung, J., Geng, X., Peng, S., and T. T. Eckert, "Dataplane Enhancement Taxonomy", Work in Progress, Internet-Draft, draft-ietf-detnet-dataplane-taxonomy-04, 7 July 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-detnet-dataplane-taxonomy-04>.

[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>.

[RFC5440] Vasseur, JP., Ed. and JL. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, DOI 10.17487/RFC5440, March 2009, <https://www.rfc-editor.org/info/rfc5440>.

[RFC5441] Vasseur, JP., Ed., "A Backward-Recursive PCE-Based Computation (BRPC) Procedure to Compute Shortest Constrained Inter-Domain Traffic Engineering Label Switched Paths", RFC 5441, DOI 10.17487/RFC5441, April 2009, <https://www.rfc-editor.org/info/rfc5441>.

[RFC6805] King, D., Ed. and A. Farrel, Ed., "The Application of the Path Computation Element Architecture to the Dark Side of the Internet", RFC 6805, DOI 10.17487/RFC6805, November 2012, <https://www.rfc-editor.org/info/rfc6805>.

[RFC8253] Lopez, D., Gonzalez de Dios, O., Wu, Q., and D. Dhody, "PCEPS: Usage of TLS to Provide a Secure Transport for the Path Computation Element Communication Protocol (PCEP)", RFC 8253, DOI 10.17487/RFC8253, October 2017, <https://www.rfc-editor.org/info/rfc8253>.

[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/rfc/rfc8939>.

[RFC8964] Varga, B., Ed., Farkas, J., Berger, L., Malis, A., Bryant, S., and J. Korhonen, "Deterministic Networking (DetNet) Data Plane: MPLS", RFC 8964, DOI 10.17487/RFC8964, January 2021, <https://www.rfc-editor.org/rfc/rfc8964>.

[RFC9016] Varga, B., Ed., Farkas, J., Berger, L., Malis, A., Bryant, S., and J. Korhonen, "Flow and Service Information Model for Deterministic Networking (DetNet)", RFC 9016, DOI 10.17487/RFC9016, May 2021, <https://www.rfc-editor.org/info/rfc9016>.

[RFC9320] Finn, N., Le Boudec, J.-Y., Mohammadpour, E., Zhang, J., and B. Varga, "Deterministic Networking (DetNet) Bounded Latency", RFC 9320, DOI 10.17487/RFC9320, November 2022, <https://www.rfc-editor.org/rfc/rfc9320>.

Authors’ Addresses

Carlos J. Bernardos (editor)
Universidad Carlos III de Madrid
Av. Universidad 30
28911 Leganes Madrid
Spain
Phone: +34 91624 6235
Email: cjbc@it.uc3m.es
URI: http://www.it.uc3m.es/cjbc

Luis M. Contreras
Telefonica
Spain
Email: luismiguel.contrerasmurillo@telefonica.com

Quan Xiong
ZTE Corporation
China
Email: xiong.quan@zte.com.cn

Alain Mourad
InterDigital Europe
Email: Alain.Mourad@InterDigital.com

URI: http://www.InterDigital.com/