

DETNET WG
Internet-Draft
Intended status: Informational
Expires: 9 November 2026

C.J. Bernardos, Ed.
Universidad Carlos III de Madrid
L. Contreras
Telefonica
Q. Xiong
ZTE Corporation
A. Mourad
InterDigital
8 May 2026

A Control Plane Framework for Multi-Domain Deterministic Networking
(DetNet)
draft-ietf-detnet-multi-domain-framework-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 generic framework for a multi-domain DetNet control plane. 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 model and a peer-to-peer "stitching" model. While a Path Computation Element (PCE)-based realization is used as an illustrative example, the framework is designed to be applicable to any controller-plane technology that satisfies the stated functional requirements. This framework provides the foundation for more specific work on multi-domain DetNet solutions.

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 9 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	3
2. Conventions and Terminology	4
3. Defining a DetNet Domain	4
3.1. Domain Characteristics	4
3.2. Scope of a DetNet Domain	5
4. Multi-Domain DetNet Architectures	5
4.1. Exemplary Use Case	5
4.2. Functional Requirements	6
4.3. Problem Statement	7
4.4. Hierarchical Coordination Model	7
4.5. Peer-to-Peer (Stitching) Model	8
5. Multi-Domain DetNet Flow Considerations	9
5.1. End-to-End Path Computation	9
5.2. Resource Management	10
5.3. End-System Awareness	10
5.4. Flow Aggregation	10
6. Security Considerations	10
7. IANA Considerations	11
8. Acknowledgments	11
9. References	11
9.1. Normative References	11
9.2. Informative References	11
Authors' Addresses	13

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 domain boundaries. The DetNet Controller Plane Framework [I-D.ietf-detnet-controller-plane-framework] provides the basis for controller-plane operations; this document extends that work specifically to multi-domain scenarios.

The framework defined here is intentionally technology-agnostic with respect to the controller-plane implementation. Centralized path computation elements, distributed controllers, Software-Defined Networking (SDN) controllers, and other approaches can all be mapped onto the architectural models described herein. The Path Computation Element (PCE) and its associated protocols [RFC4655] [RFC5440] are used as a concrete example throughout this document to illustrate how the framework can be realized, but they do not represent the only valid instantiation.

This document proposes a foundational framework by:

- * Defining what constitutes a "domain" in the context of multi-domain DetNet control.
- * Describing high-level architectural models for managing multi-domain paths.
- * Identifying key considerations for establishing and maintaining DetNet flows across domain boundaries.

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]. The following additional terms are used:

Domain Controller: A logical entity responsible for path computation and resource management within a single DetNet domain. A Domain Controller has full visibility into the topology and resources of its own domain. A PCE is one example of a technology that can fulfil the Domain Controller role.

Hierarchical Controller: A logical entity that has an abstracted, partial view of multiple domains and is responsible for coordinating inter-domain path computation by interacting with the Domain Controllers of each constituent domain. A Parent PCE (P-PCE) in the H-PCE architecture [RFC6805] is one example of a Hierarchical Controller.

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 control plane perspective, this typically implies that:

- * A single Domain Controller instance (or a coordinated set of redundant controllers) has complete topological visibility within the domain.
- * This Domain Controller 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 Domain Controller 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.

Controller Domain: A domain is defined as the set of nodes controlled by a single Domain Controller instance. This is the primary definition used within this framework. A large administrative domain might be divided into multiple smaller controller 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, a 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 Domain Controller 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. The Domain Controller needs to select a specific queuing mechanism along the path for a DetNet flow within each domain.

4. Multi-Domain DetNet Architectures

4.1. Exemplary Use Case

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).

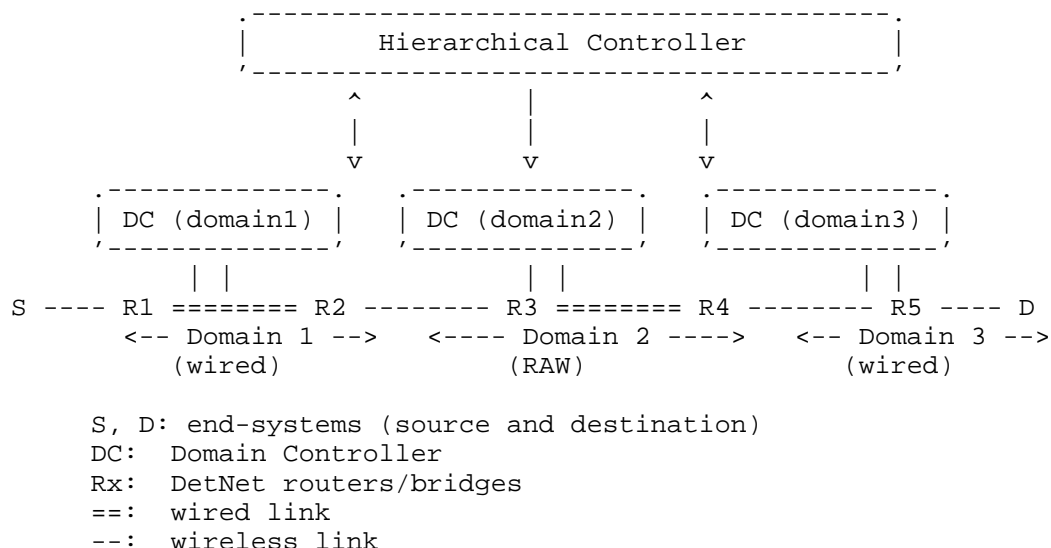


Figure 1: Exemplary multi-domain scenario

Each domain has its own Domain Controller (DC), which is responsible for path computation and resource management within that domain. Routers R2 and R3 are border routers of Domain 2 (RAW), and R3 and R4 are border routers of Domains 2 and 3, respectively. A Hierarchical Controller is responsible for end-to-end path computation and orchestration across the different Domain Controllers.

In a PCE-based realization of this framework, each Domain Controller maps to a Child PCE (C-PCE) and the Hierarchical Controller maps to a Parent PCE (P-PCE) as defined in [RFC6805]. Other controller-plane technologies (e.g., YANG/NETCONF-based SDN controllers, BGP/PLS-based systems) can equally fill these roles, provided they satisfy the functional requirements described in Section 4.2.

4.2. Functional Requirements

Regardless of the specific technology used to realize the control plane, any multi-domain DetNet framework MUST satisfy the following high-level functional requirements:

- * Each domain's controller must be able to compute an intra-domain path segment that meets a specified portion of the end-to-end DetNet QoS budget (latency, jitter, loss).

- * A mechanism must exist for domain controllers to advertise an abstracted representation of their domain's capabilities (reachability, available latency budget, bandwidth) to an inter-domain coordination entity, without exposing confidential internal topology details.
- * An inter-domain coordination mechanism must be able to compose an end-to-end path from individual per-domain path segments, allocating the end-to-end QoS budget across domains.
- * The framework must support both a centralized (hierarchical) coordination model and a distributed (peer-to-peer) coordination model.
- * Resource reservation, whether performed transactionally or sequentially, must be coordinated across all participating domains.

4.3. Problem Statement

In a multi-domain environment, no single controller 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 controller is responsible for its own internal path computation and resource allocation. The multi-domain architecture must define how these individual controllers cooperate to create a seamless end-to-end service. Two primary coordination models are considered: Hierarchical and Peer-to-Peer.

4.4. Hierarchical Coordination Model

In the Hierarchical model, a Hierarchical Controller has a partial, abstracted view of the child domains. It does not see the detailed topology within each domain but knows the reachability and characteristics (e.g., available latency budget, cost) of paths to and through them. Each Domain Controller has full visibility of its own domain's topology and resources, and is responsible for all intra-domain path computations.

The H-PCE architecture [RFC6805] is one well-defined realization of this model, where the Hierarchical Controller is a Parent PCE and each Domain Controller is a Child PCE. Other SDN-based hierarchical controller architectures follow the same conceptual structure.

In a multi-domain DetNet context, path establishment under this model proceeds as follows:

1. A request for an end-to-end DetNet path is sent to the Hierarchical Controller. This request includes the source, destination, and required QoS parameters (e.g., maximum end-to-end latency).
2. The Hierarchical Controller computes a high-level, domain-sequence path: a sequence of domains the flow must traverse, together with the entry and exit boundary nodes for each domain.
3. The Hierarchical Controller then sends requests to the Domain Controller 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 Domain Controller computes its path segment, reserves the necessary resources locally, and reports success or failure back to the Hierarchical Controller.
5. If all Domain Controllers succeed, the Hierarchical Controller confirms the end-to-end path. If any Domain Controller fails, the Hierarchical Controller may attempt to find an alternate domain-sequence path.

Each Domain Controller is aware of the specific technologies used in its domain (e.g., RAW, DetNet IP, DetNet MPLS), and can compute a path taking into account the specific constraints of those technologies. For instance, in a RAW domain, the Domain Controller can select the path, the schedule, and the links to be used to guarantee a certain level of reliability.

4.5. Peer-to-Peer (Stitching) Model

In the peer-to-peer model, also known as "stitching," there is no parent-child hierarchy. Domain Controllers from adjacent domains cooperate as peers. Path computation is performed sequentially from one domain to the next. The BRPC procedure defined in [RFC5441] for inter-area and inter-AS TE path computation is one example of how this model can be realized in a PCE context.

In a multi-domain DetNet context, path establishment under this model proceeds as follows:

1. The Domain Controller in the source domain (DC-1) receives a path request for a flow destined for another domain.

2. DC-1 computes a path from the source node to a suitable exit border node in its domain.
3. DC-1 then sends a request to the Domain Controller of the adjacent domain (DC-2), specifying the entry border node and the final destination, along with the remaining QoS budget.
4. DC-2 computes a path through its domain to either the final destination (if it is within domain 2) or to another suitable exit border node. It then "stitches" this segment to the previous one.
5. This process repeats until the Domain Controller in the destination domain is reached. The path may be confirmed backward along the chain of Domain Controllers.

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.

The end-to-end path computation should select one of the end-to-end bounded latency metrics for all domains as defined in [I-D.xiong-pce-detnet-bounded-latency]. The calculation of the bounded latency will differ depending on the queuing solutions and categories as per [I-D.ietf-detnet-dataplane-taxonomy]. The end-to-end bounded delay will be the sum of the per-domain delays, and the end-to-end variation will be either the sum of the per-domain variations or the maximum variation among all domains, depending on the queuing mechanism used.

In the Hierarchical model, the Hierarchical Controller collects domain-specific capability information from each Domain Controller and divides the end-to-end QoS budget of a DetNet flow into per-domain sub-budgets, based on the capabilities (e.g., latency, jitter) available within each domain.

In the peer-to-peer stitching model, the end-to-end budget is divided starting from the source domain's controller and passed along to each successive adjacent domain, until reaching the destination domain's controller. The controller within each domain computes 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, and 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 end-to-end QoS guarantees for aggregated flows that span across multiple domains. Multiple flows may be aggregated in one domain and disaggregated in another. The network parameters of an aggregated flow should be exchanged among different domain controllers. Path computation should consider the end-to-end budget of the aggregated flow, which must cover the requirements of all its member flows.

6. Security Considerations

Multi-domain operations introduce significant security challenges. The communication between Domain Controllers 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 Controller to an external entity (a Hierarchical Controller or a peer Domain Controller) is sensitive. The framework must allow for policy-based control over the level of abstraction and detail that is shared, so that internal topology information is not unnecessarily disclosed across administrative boundaries.

Where a PCE-based realization is used, the security considerations of [RFC8253] also apply.

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 MultiX (Grant Agreement No. 101192521) and DISCO6G-CM.

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. J. and A. Mourad, "DetNet multidomain extensions", Work in Progress, Internet-Draft, draft-bernardos-detnet-raw-multidomain-07, 16 April 2026, <<https://datatracker.ietf.org/doc/html/draft-bernardos-detnet-raw-multidomain-07>>.

[I-D.ietf-detnet-controller-plane-framework]

Malis, A. G., Geng, X., Chen, M., Varga, B., and C. J. Bernardos, "A Framework for Deterministic Networking (DetNet) Controller Plane", Work in Progress, Internet-Draft, draft-ietf-detnet-controller-plane-framework-15, 24 September 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-detnet-controller-plane-framework-15>>.

[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-05, 8 January 2026, <<https://datatracker.ietf.org/doc/html/draft-ietf-detnet-dataplane-taxonomy-05>>.

[I-D.xiong-pce-detnet-bounded-latency]

Xiong, Q., Liu, P., and R. Gandhi, "PCEP Extension for Bounded Latency", Work in Progress, Internet-Draft, draft-xiong-pce-detnet-bounded-latency-07, 29 January 2026, <<https://datatracker.ietf.org/doc/html/draft-xiong-pce-detnet-bounded-latency-07>>.

[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., Zhang, R., Bitar, N., and JL. Le Roux, "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 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>>.

- [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/info/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/info/rfc8964>>.
- [RFC9016] Varga, B., Farkas, J., Cummings, R., Jiang, Y., and D. Fedyk, "Flow and Service Information Model for Deterministic Networking (DetNet)", RFC 9016, DOI 10.17487/RFC9016, March 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/info/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/>