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

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

IETF 124 – DetNet WG

Carlos J. Bernardos (UC3M)
Luis M Contreras (Telefonica)
Quan Xiong (ZTE)
Alain Mourad (InterDigital)

November 2025

Aim and scope

- Initial steps exploring multi-domain in DetNet (originally looking only at wireless/RAW scenarios) in draft-bernardos-detnet-raw-multidomain
- There are scenarios requiring DetNet multi-domain
- Goal was to explore ~~RAW~~ DetNet multi-domain gaps (arch, OAM, etc) and propose solutions
  - Mostly on DetNet control plane
  - Gaps and requirements partially covered in draft-ietf-detnet-controller-plane-framework
- Feedback from the WG was to work on a PCE-based solution

I E T F

Goals and ambition

- Builds on the 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

I E T F

Scope of a DetNet Domain

- The boundaries of a DetNet domain can be defined based on several factors:
    - **Administrative Domain:** a set of network elements under the control of a single network operator or administrative entity
    - **PCE Control Domain:** a set of nodes controlled by a single PCE instance
    - **Technological Domain:** defined by the consistent use of a specific data plane technology

Defining a DetNet domain

- 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*:
  - 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

PCE-based multi-domain DetNet architectures

```
                   .------------------------------------.
                   |             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
```

PCE-based multi-domain DetNet architectures

- Problem statement
  - In a multi-domain environment, no single PCE has end-to-end visibility of the full network topology
- Hierarchical PCE (**H-PCE**) approach
  - A Parent PCE (P-PCE) has a partial, abstracted view of the child domains
  - A Child PCE (C-PCE) has full visibility of its own domain's topology and resources
- Peer-to-Peer (***stitching***) PCE approach
  - No parent-child hierarchy. PCEs from adjacent domains cooperate as peers
  - The path computation is performed sequentially from one domain to the next

Multi-Domain DetNet flow considerations

- 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
- Resource Management
  - Resources MAY be reserved in each domain for the flow
  - The control plane also needs to advertise inter-domain resource information, including bandwidth, delay, jitter with related queuing mechanisms for QoS coordination
- End-System awareness

Next steps and questions to the WG

- Is this a first good approach as indicated by the WG?
- Is there interest in pursuing this work?
  - Potential call for adoption in a future meeting?
- Please let us know if you want to work on this

Acknowledgements

- Partially funded by 6G-DATADRIVEN-04 project