

DetNet
Internet-Draft
Intended status: Informational
Expires: 17 April 2026

Q. Xiong
ZTE Corporation
T. Jiang
China Mobile
J. Joung
Sangmyung University
14 October 2025

Flow Aggregation for Enhanced DetNet
draft-xiong-detnet-flow-aggregation-03

Abstract

[I-D.ietf-detnet-scaling-requirements] proposed the data plane requirements in scaling networks. This document describes the additional requirements for flow aggregation in enhanced DetNet and provides the enhancement consideration. It also discusses how the flow aggregation could be realized in a 5GS logical DetNet node participating in enhanced DetNet.

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 17 April 2026.

Copyright Notice

Copyright (c) 2025 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
1.1. Requirements Language	3
2. Terminology	3
3. Objectives & Requirements: Flow Aggregation in Enhanced DetNet	3
3.1. DetNet Services to Aggregated Flows across Domains	3
3.2. Aggregated vs. Fine-grained QoS Provisioning	4
3.3. Scale Down States Maintenance at Transit Nodes at Class-aggregated Level	5
4. Enhancement Consideration for Flow Aggregation	6
4.1. Flow Classification	6
4.2. Flow Identification	6
4.3. Flow Coordination	7
5. Realization of Flow Aggregation for 5GS DetNet	7
5.1. Realization of 5GS DetNet Service across Domains	8
5.2. 5GS QoS Provisioning: Aggregated vs. Fine-grained	8
5.3. State Maintenance at a 5GS Transit node	9
5.4. Flow Classification & Identification at 5GS node	9
6. Security Considerations	10
7. IANA Considerations	10
8. Acknowledgements	10
9. References	10
9.1. Normative References	10
9.2. Informative References	11
Authors' Addresses	12

1. Introduction

The [RFC8655] clearly states that Deterministic Networking (DetNet) operates at the IP layer and delivers service which provides extremely low data loss rates and bounded latency. The DetNet data plane must provide the aggregation of DetNet flows in order to support larger numbers of DetNet flows and improve scalability by reducing the per-hop states. The [RFC8938] introduces that the flow aggregation is the ability to aggregate individual flows along with their associated resource control into a large aggregate. It is recommended that the DetNet flow aggregation be enabled for compatible flows with the same or very similar QoS and CoS characteristics via the use of wildcards, masks, prefixes, and ranges. It also suggests the reduction of per-hop states help avoid the per DetNet-flow specific state maintenance in a transit node. It further provides arguments on how DetNet services might be realized

in term of delay bound, delay jitter and bandwidth provisioning. Furthermore, the [RFC8964] has proposed and expanded two methods of flow aggregation, one being the aggregation via LSP hierarchy and the other to aggregate DetNet flows as a new combined DetNet flow.

For enhanced DetNet, [I-D.ietf-detnet-scaling-requirements] proposed the data plane requirements and it should support the explicit aggregated flow identification. In scaling networks, it also should support aggregated flows over multi-domains and achieve different levels of co-existed applications with different SLAs requirements which requiring the fine-grained QoS provisioning through flow aggregation. Moreover, the aggregated flows still requires a large amount of control signaling to establish and maintain the states of DetNet flows in enhanced DetNet. It is required to improve the scalability and forward packets at the class-aggregate level, instead of the per-flow or flow-aggregate level.

This document describes the additional requirements of flow aggregation in enhanced DetNet and provides the enhancement consideration. It also discusses the realization of DetNet flow aggregation for 5GS.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

2. Terminology

The terminology is defined as [RFC8655].

3. Objectives & Requirements: Flow Aggregation in Enhanced DetNet

3.1. DetNet Services to Aggregated Flows across Domains

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. As per [I-D.ietf-detnet-scaling-requirements], different network implementations may be intended for different application domains, where there is no additional requirements for the coordination. As defined in [ITU-T Y.2122], the network operating parameters of a flow aggregate should be exchanged among different network domains. As shown in Figure 1, the DetNet domain B receiving an aggregated flow should identify the flow and get the service requirements such as the QoS parameters of the flow aggregate.

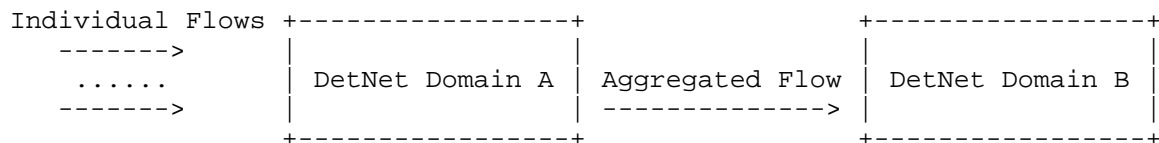


Figure 1: Aggregating DetNet Flows across Multiple Domains

3.2. Aggregated vs. Fine-grained QoS Provisioning

The draft [I-D.ietf-detnet-scaling-requirements] specifies that different levels of applications differ in the SLAs requirements such as tight jitter, strict latency, loose latency and so on. While these types of aggregated requirements might bear the coarse-grained nature, individual flows demand differentiated DetNet treatments and more granular QoS forwarding behaviors. A DetNet node or domain adopting multiple forwarding technologies needs to transmit individual flows by aggregating them into a selected treatment solution that corresponds to one of some pre-defined per-hop QoS behaviors, as shown in Figure 2. The DetNet flows with the same level of service requirements can be aggregated to receive collective treatments and forwarding behaviors. The DetNet flows can be aggregated to several pre-defined classes. For example, as per [I-D.jlg-detnet-5gs], a 5GS as a logical DetNet node requires to achieve the service requirements and service levels of the aggregated flows, along with the provisioning of fine-grained per-hop behavior (PHB) to each individual flow.

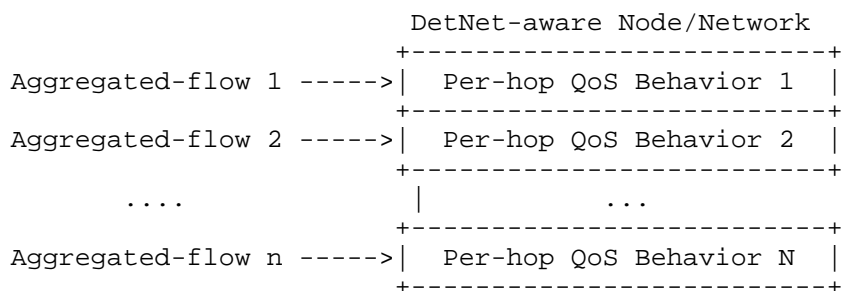


Figure 2: Aggregating DetNet flows to corresponding QoS PHBs

3.3. Scale Down States Maintenance at Transit Nodes at Class-aggregated Level

As per [I-D.ietf-detnet-dataplane-taxonomy], the treatment solutions in data plane can be categorized based on performance and functional characteristics. For example, the category of a solution can be classified based on the traffic granularity, e.g., flow aggregate vs. class level. The class level is provided to simplify the control and accommodate traffic fluctuations by combining flows requiring the same or similar levels of service requirements. The flow aggregation based on the class level could further improve the scalability. As per [I-D.ietf-detnet-scaling-requirements], there may be a large number of traffic flows in a scaling network, which makes it nearly impossible to achieve the flow-specific state identification. As shown in the Figure 3, the flow identification of aggregated-class can be used to indicate the required treatment and forwarding behaviors without the need to maintain excessive states at transit nodes.

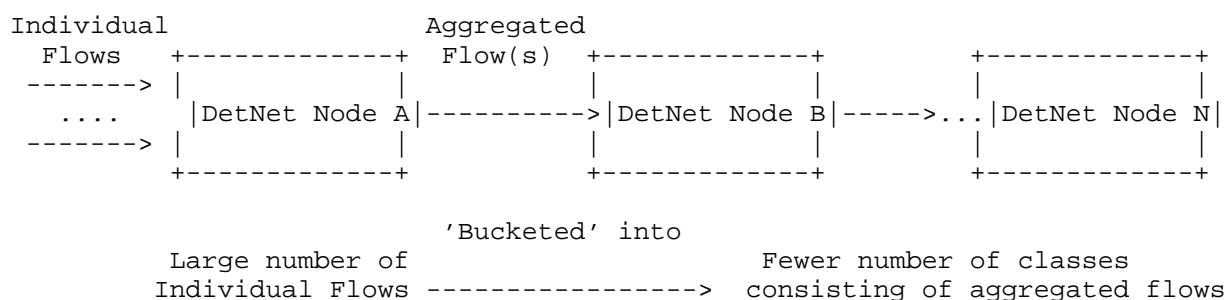


Figure 3: Aggregating DetNet Flows to Improve Scalability

When DetNet flows are aggregated based on service-class level, transit nodes provide deterministic services to a flow aggregate and go through the per-class scheduling without the burden to maintain excessive per-flow states. Still, a transit node performing aggregation should ensure all per-flow service requirements within an aggregated class are achieved. For example, the latency or jitter bounds of an aggregated class shall not exceed the corresponding metrics of any individual flow that has been bucketed into the class. The aggregation based on the class level has data plane and controller plane aspects.

4. Enhancement Consideration for Flow Aggregation

This document proposes an enhancement for DetNet flow aggregation including the functions such as flow classification, flow identification and flow coordination.

4.1. Flow Classification

The DetNet QoS can be achieved by aggregating flows based on DetNet flow-specific traffic classification and providing the traffic-forwarding treatment. The flow classification should consider the flow-specific characteristics such as traffic specification and service requirements. For example, the DetNet flows MAY be classified based on the service SLAs requirements of applications in scaling networks as per [I-D.xiong-detnet-differentiated-detnet-qos]. And the services can also be classified into tight/loose latency, large/small burst, periodic/non-periodic and large/small scale services as per [I-D.ietf-detnet-dataplane-taxonomy]. Several classes can be predefined to indicate the different levels of applications with SLAs requirements and each class demands differentiated QoS behaviors and treatment as well as different DetNet capabilities in scaling networks.

For the controller plane, the service should be provisioned on an aggregated-class level. The resources should be controlled and scheduled on a per-class basis and the routes should be established to meet the service requirements with the allocated resources. The edge nodes must be able to handle admission control for DetNet flows to an aggregated class based on the available resources. For the data plane, when DetNet flows are aggregated to a class, transit nodes provide service to the aggregate and not on a per-DetNet-flow basis. And the transit nodes supporting this type of aggregation should identify the class of the aggregated flows and ensure that individual flows receive the corresponding traffic treatment and forwarding behaviour.

4.2. Flow Identification

The flow identification is used to identify the aggregated flow and to indicate the required treatment and forwarding behaviors. As described in [I-D.ietf-detnet-scaling-requirements], it is required to be dynamic and simplified to ensure the aggregated flows have compatible DetNet flow-specific QoS characteristics. For the data plane, individual flows may be aggregated for treatment based on shared service specification on aggregated-class level which is identified by an aggregation class (A-Class). A transit node should provide the class level traffic treatment based on A-Class. The aggregation class information may be used alone or together with

other metadata to indicate the required queuing and forwarding behaviors. The encoding of the A-Class may reuse the DSCP/TC or existing field such as the TC field in A-Label as per [RFC8964]. And it also can be encapsulated with the deterministic latency information as per [I-D.xiong-detnet-data-fields-edp].

4.3. Flow Coordination

In scaling networks, flow aggregations become more prevalent, with flows frequently joining and leaving, which may potentially lead to accumulated bursts of flows across multiple hops. Such challenges can be mitigated by coordinating packets within aggregated flows such as proportional scheduling and interleaving. Proportional scheduling could allocate transmission opportunities based on flow weights, ensuring that each flow receives a fair share of network resources. Interleaving could achieve micro burst smoothing by rotating the transmission of packets across different flows through timed gates as described in [I-D.eckert-detnet-flow-interleaving].

5. Realization of Flow Aggregation for 5GS DetNet

The 3GPP in its document [TS.23.501] has defined and standardized how a 5G system (5GS) may behave as a logical DetNet node, as well as how a 5GS DetNet node may integrate into the IP-domain DetNet as described in [RFC8655]. 3GPP has realized the functionalities of the DetNet forwarding sub-layer.

As a logical DetNet transit node, a 5GS behaves as a transparent box to external DetNet entities. It can connect to either DetNet end systems or full-fledged IP DetNet domain(s) or both. The 3GPP [TS.23.501] has demonstrated a 'composite' architecture in that a 5GS could act as one or more DetNet nodes upon the integration into the IP DetNet domain. The granularity of determining a 5GS DetNet node is per UPF for each network instance, with the corresponding UPF-id identified as the 5GS DetNet node-id.

The 3GPP [TS.23.503] has implicitly specified two types of DetNet related traffic parameters. One type is the higher-level per-(logical)-node QoS requirements like the node max-latency, max-loss, etc., while the other is more granular settings with which DetNet flow attributes and specifications are mapped to the Flow Description information. The DetNet flow specifications could be based on IP-tuple information, e.g., including IP address, protocol type, ToS, TCP/DUP ports, etc. The document [I-D.jlg-detnet-5gs] has provided more details.

Please note that this draft revolves around the general discussions of the flow aggregations in enhanced DetNet across multiple domains. It emphasizes the objectives & requirements, along with insightful considerations for the possible enhancement to the matter. This indicates the generic principles that are related to the cross-domain flow aggregation as raised in the draft. While the 3GPP [TS.23.503] defines a 5GS may behave as a logical DetNet (transit) node and the 5GS does own certain advantageous features for a 'composite' DetNet instantiation, the (DetNet) flow aggregation is not an intrinsic characteristics that has been fulfilled in the 5GS. As we explain in the following subsection, the realization of flow aggregation for a 5GS DetNet 'composite' node participating in an enhanced DetNet domains requires the seamless interactions between the IETF domain (DetNet) controller and the 5GS domain counterpart.

5.1. Realization of 5GS DetNet Service across Domains

3GPP has so far standardized the forward sub-layer functionality for 5GS. It indicates a 5GS (logical) DetNet node could connect to other end systems and/or IP DetNet domains, together to form a holistic end-to-end DetNet. Thanks to the 'composite' architecture of a 5GS node, along with the interaction between an CPF:DetNet controller in IETF domain and the NF TSCTSF in 3GPP domain [TS.23.501], a 5GS node might realize much more advanced DetNet services than a traditional IP DetNet transit node.

In scenarios where the (IETF-domain) CPF:Detnet Controller could well divide the DetNet QoS service requirements that are in reality associated with an integrated DetNet domain into multiple QoS sub-requirements that together form the original end-to-end QoS equivalence, a 5GS might be considered as a standalone DetNet (sub-)domain with its own DetNet QoS (sub-)requirements that would be provisioned from the CPF:DetNet controller. The 5GS DetNet QoS (sub-)requirements serve a portion of the original requirements of the integrated DetNet domain. These together form a scaling network to realize the 5GS DetNet service across domains.

5.2. 5GS QoS Provisioning: Aggregated vs. Fine-grained

We have explained previously that the 3GPP [TS.23.503] has implicitly specified two categories of DetNet related traffic parameters. One type bears the aggregated nature for (5GS DetNet) node-level requirements, while the other addresses the more granular DetNet flow-level attributes and specifications. Evidently, with this kind of two-hierarchy architecture, a 5GS DetNet node could achieve not only the node-level aggregated QoS requirements, but also the more fine-grained flow-level QoS provisioning. This reflects the true value of applying our flow aggregation model in scaling networks to

realizing advanced DetNet services for 5GS.

Here, we want to point out that the feasibility of applying our flow aggregation scheme indeed depends on the hierarchical nature of a 5GS DetNet node. Had the same aggregation scheme been applied to DetNet nodes that do not have the similar intrinsic hierarchy, the effectiveness could be certainly impaired.

5.3. State Maintenance at a 5GS Transit node

The 5GS QoS architecture is roughly comprised of three levels, i.e., the UE, the PDU-session, and the QoS-flow levels. Technically, a 5GS DetNet node is of 'composite' nature with a large number of configuration, provisioning, operation and runtime states to maintain. At first glance, this might undermine the state-reduction objective via the flow aggregation for a 5GS DetNet transit node.

Fortunately, the 5GS DetNet work owns intrinsically a couple of aspects to handle the challenges:

First, also as we have mentioned before, a 5GS node behaves as a transparent entity participating in the DetNet domain. Thus, even having a significant number of states, this can naturally have the 5GS DetNet related states remain hidden from the external entities (and domains).

Second, the 3GPP NF TSCTSF exchanges only traffic parameters with the IETF CPF:Detnet Controller, but not the states that are maintained inside a 5GS DetNet node. The external DetNet domain does not care the inside status of a 5GS, nor can it.

5.4. Flow Classification & Identification at 5GS node

As we have explained so far, the IETF domain CPF:DetNet controller provides traffic parameters & specifications to 3GPP NF TSCTSF. Thus, the SLA requirements of applications in scaling networks could be readily pre-specified in the IETF DetNet CPF, which would then apply the flow classification mapping (to aggregated service classes) and send them over to a 5GS DetNet node to enforce. This model can also relieve the classification burden of a 5GS node in reality.

The 5GS has excellent control logics to address flow identification. For example, PDRs (Packet Detection Rules), SDF (Service Data Flow) filters (e.g., IP-filter, MAC-filter, etc.), etc., are all good tools to differentiate flows [TS.23.501]. Further, the 5GS has standardized powerful procedures for the establishment & update of PDU sessions/QoS flows, which accordingly achieves the flow dynamics (e.g., flow joining & leaving a flow-aggregate as manifested

potentially by a PDU session) [TS.23.502]. Moreover, some QoS parameters, e.g., Aggregated Bit Rate (ABR), may stand at different levels, including UE-ABR, Session-ABR, flow-ABR, etc., that would make the service differentiation & sharing on the aggregated-class (A-Class) level feasible.

6. Security Considerations

Security considerations for DetNet are covered in the DetNet Architecture [RFC8655] and DetNet security considerations [RFC9055].

7. IANA Considerations

This document makes no requests for IANA action.

8. Acknowledgements

TBA

9. References

9.1. Normative References

[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>>.

[I-D.ietf-detnet-scaling-requirements]

Liu, P., Li, Y., Eckert, T. T., Xiong, Q., Ryoo, J.,
zhushiyin, and X. Geng, "Requirements for Scaling
Deterministic Networks", Work in Progress, Internet-Draft,
draft-ietf-detnet-scaling-requirements-09, 7 September
2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-detnet-scaling-requirements-09>>.

[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>>.
- [RFC8938] Varga, B., Ed., Farkas, J., Berger, L., Malis, A., and S. Bryant, "Deterministic Networking (DetNet) Data Plane Framework", RFC 8938, DOI 10.17487/RFC8938, November 2020, <<https://www.rfc-editor.org/info/rfc8938>>.
- [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>>.
- [RFC9055] Grossman, E., Ed., Mizrahi, T., and A. Hacker, "Deterministic Networking (DetNet) Security Considerations", RFC 9055, DOI 10.17487/RFC9055, June 2021, <<https://www.rfc-editor.org/info/rfc9055>>.
- [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>>.

9.2. Informative References

- [I-D.eckert-detnet-flow-interleaving]
Eckert, T. T., "Deterministic Networking (DetNet) Data Plane - Flow interleaving for scaling detnet data planes with minimal end-to-end latency and large number of flows.", Work in Progress, Internet-Draft, draft-eckert-detnet-flow-interleaving-03, 7 July 2025, <<https://datatracker.ietf.org/doc/html/draft-eckert-detnet-flow-interleaving-03>>.
- [I-D.jlg-detnet-5gs]
Jiang, T., Liu, P., and X. Geng, "DetNet YANG Model Extension for 5GS as a Logical DetNet Node", Work in Progress, Internet-Draft, draft-jlg-detnet-5gs-01, 20 October 2023, <<https://datatracker.ietf.org/doc/html/draft-jlg-detnet-5gs-01>>.

[I-D.xiong-detnet-data-fields-edp]

Xiong, Q., Liu, A., Gandhi, R., and D. Yang, "Data Fields for DetNet Enhanced Data Plane", Work in Progress, Internet-Draft, draft-xiong-detnet-data-fields-edp-03, 1 July 2025, <<https://datatracker.ietf.org/doc/html/draft-xiong-detnet-data-fields-edp-03>>.

[I-D.xiong-detnet-differentiated-detnet-qos]

Xiong, Q., Zhao, J., Du, Z., Zeng, Q., and C. Liu, "Differentiated DetNet QoS for Deterministic Services", Work in Progress, Internet-Draft, draft-xiong-detnet-differentiated-detnet-qos-01, 27 June 2024, <<https://datatracker.ietf.org/doc/html/draft-xiong-detnet-differentiated-detnet-qos-01>>.

[TS.23.501]

"3GPP TS 23.501 (V19.0.0): System Architecture for the 5G System (5GS)", 3GPP TS 23.501, June 2024.

[TS.23.502]

"3GPP TS 23.502 (V19.0.0): Procedures for the 5G System (5GS)", 3GPP TS 23.502, June 2024.

[TS.23.503]

"3GPP TS 23.503 (V19.0.0): Policy and charging control framework for the 5G System (5GS); Stage 2", 3GPP TS 23.503, June 2024.

Authors' Addresses

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

Tianji Jiang
China Mobile
Email: tianjijiang@chinamobile.com

Jinoo Joung
Sangmyung University
Email: jjoung@smu.ac.kr