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

Z. Han
R. Pang
C. Liu
China Unicom
J. Yan
X. ZHU
ZTE Corporation
16 October 2025

Anomalous Packets Handling for DetNet
draft-han-detnet-anomalous-packets-handling-01

Abstract

In deterministic networking (DetNet), resource conflicts at flow aggregation nodes can lead to network anomalies. Existing data plane mechanisms for handling anomalous packets, such as simple discarding or treating them as Best-Effort (BE) flows, are insufficient. Consequently, the network performance can degrade to a level inferior to of traditional QoS approaches.Therefore, in order to handle the anomalous traffic, the data plane should implement an enhanced handling mechanism.

This document proposes an enhanced anomalous packet handling solution for DetNet. This solution specifies two policies: the packet squeezing policy and the packet degrading policy. These policies provide a flexible, enhanced mechanism applicable to various queuing mechanisms, ensuring the preferential scheduling and preservation of deterministic service traffic under anomalous conditions.

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 19 April 2026.

Table of Contents

1. Introduction

DetNet is capable of providing real-time application services with deterministic guarantees such as bounded latency, low jitter, and low packet loss rate, as per [RFC8655]. One of the major technologies of DetNet is resource allocation, as per [RFC8938], which reserves necessary resources for specified DetNet flows to mitigate packet loss and jitter caused by network congestion. The control plane orchestrates the paths of DetNet flows to avoid resource conflicts. The data plane then transmits DetNet flows based on this orchestration result, employing mechanisms like traffic shaping, flow admission control, and forwarding information encapsulation to

maintain the required QoS.

Each node along the end-to-end path may serve as an aggregation node. Aggregated flows that belong to the same traffic class share the reserved resources at the outgoing port. Ideally, the transmission of each flow within the same traffic class would strictly conform to the scheduling of the control plane, thereby meeting the strict deterministic requirements. However, this ideal scenario is often difficult to achieve due to the diversity of deterministic flows— such as occasional microbursts and packet size fluctuations. Allocating resources based on the maximum packet size may lead to resource waste, while basing them on the average size may cause resource conflicts. Furthermore, software and hardware limitations can introduce additional discrepancies. For instance, algorithmic flaws in the control plane may lead to resource conflicts in extreme cases, and high-priority protocol messages (e.g., ARP packets under abnormal conditions) in the data plane may preempt service packets, causing delays for lower-priority flows.

To address these network anomalies, the control plane should properly schedule resources to avoid resource conflicts at the aggregation nodes. As defined in [RFC8655], service protection solutions like PREOF (Packet Replication, Elimination, and Ordering Functions) are proposed based on multi-path transmission. Although PREOF can prevent performance reduction by reserving a large amount of redundant resources for the specified service flows, this approach may lead to poor resource utilization and potentially diminishing the value proposition of deterministic technologies. In the data plane, the existing mechanisms are relatively simple and primitive. For example, the data plane may choose to discard packets directly or buffer them until the resources allocated to its traffic class become available. Both of these approaches can result in Quality of Service (QoS) degradation that is even worse than that of Best-Effort (BE) flows.

Therefore, an enhanced, automated mechanism for handling anomalous packets in the data plane is essential for the future implementation and application of deterministic network technology.

This document proposes an enhanced anomalous packet handling policy and solution for DetNet, supporting two policies: packet squeezing and packet degrading, which can be enabled individually or in combination. The control plane and users can configure the policies' activation and associated parameters. Detailed procedures for implementing these policies across various queuing mechanisms are provided.

2. 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].

3. Terminology

The terminology is defined as [RFC8655].

4. Anomalous Forwarding Detection

The real-time detection in the data plane aims to identify anomalous forwarding behaviors. When an anomaly is detected, enhanced processing policies, such as packet squeezing and packet degrading, are applied to ensure the preferential scheduling of deterministic flows, even under abnormal conditions.

The detection process is closely associated with the queuing mechanisms employed. In general, an anomaly is detected at a node when an arriving packet's designated queue has already exceeded its allocated resource limit (e.g., buffer depth or packet count) for the current scheduling cycle or timeslot. Typically, for TQF[I-D.peng-detnet-packet-timeslot-mechanism], the target output timeslot of a packet at the current node can be determined by the upstream timeslot label and the basic timeslot mapping. For EDF[I-D.peng-detnet-deadline-based-forwarding], the target output timeslot at the current node is calculated based on the budget and delay target carried in the packet. Each output timeslot is associated with a queue. When a packet arrives, it is enqueued in the corresponding queue. For CQF, if the current scheduling timeslot is 1 and the target timeslot is 5, the packet for target output timeslot 5 will be placed into the corresponding queue preemptively. Before the packet enters the output queue, the queue depth is checked. If it does not exceed the allowable packet capacity of the queue, the packet is enqueued normally. If it exceeds the allowable capacity, it indicates an anomaly.

5. Anomalous Packets Handling Policy

The proposed solution supports two enhanced anomalous packet handling policies in the data plane:

* Squeezing Policy: Temporarily delays anomalous packets by "squeezing" them into the next timeslot while retaining their original scheduling information.

* Degrading Policy: Redirects anomalous packets to a lower-priority queue and modifies their scheduling parameters when the accumulation of anomalous packets exceeds a predefined threshold.

These policies provide flexibility in activation: they can be enabled concurrently, individually, or disabled entirely. If neither policy is enabled, the default mechanism, such as discarding the packets or treating them as a BE flow, will be utilized.

5.1. Squeezing Policy

The data plane can support the squeezing policy through the configuration of the squeezing threshold. When anomalous traffic causes the queue occupancy to exceed its allocated capacity— but remains below the squeezing threshold— the system applies the squeezing policy. Specifically, the system enqueues anomalous the packets and records the number of squeezed bits. According to the squeezing policy, packets that cannot be sent within the allocated time are squeezed into the next timeslot until the queue is emptied. The squeezing policy is compatible with various queuing mechanisms; however, the implementation details will vary depending on the specific mechanism utilized.

Assume that each timeslot permits 4000 bits, and the squeezing threshold is set to 2000 bits. Consider a service flow where the size of each packet is fixed at 1000 bits. Packets 1 to 4 are assigned to timeslot 1, while packets numbered 5 to 7 are assigned to timeslot 2. Due to the presence of aggregated traffic, assume that the current depth of queue 1 is 2000 bits.

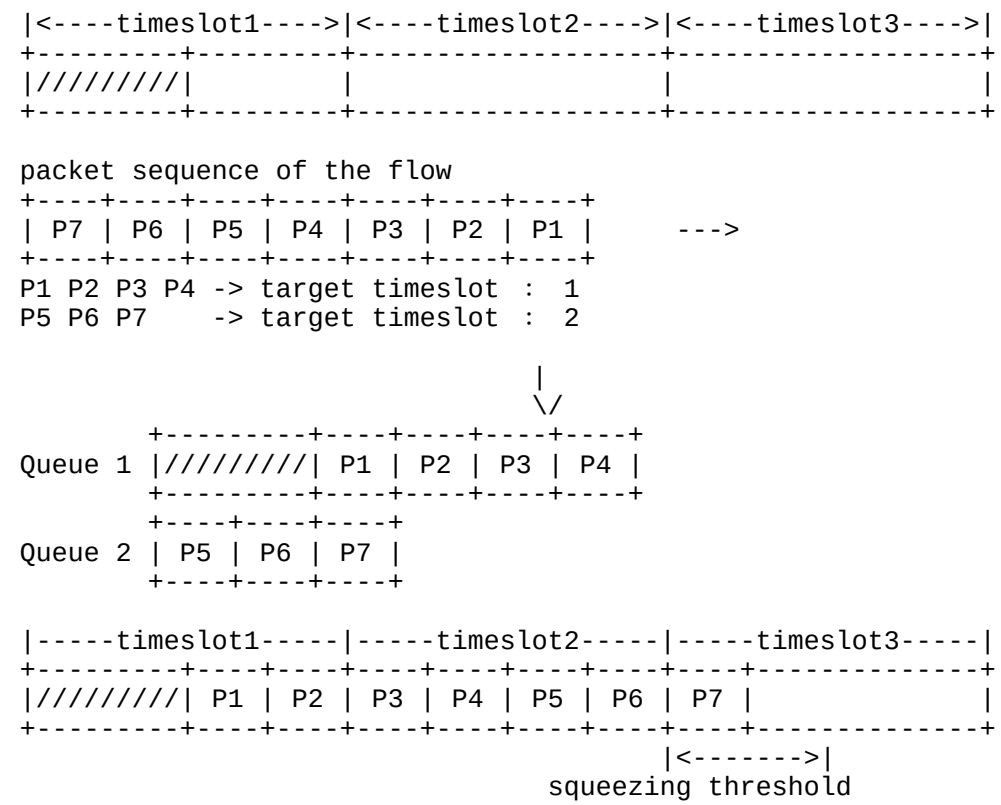

Figure 1: Squeezing policy based on timeslot-based queuing mechanism

Figure 1 illustrates the processing of service flow packets numbered 1 through 7. Packets 1 and 2 are placed into Queue 1 (associated with timeslot 1). Given the existing aggregated traffic, the addition of Packets 1 and 2 (totaling 2000 bits) causes Queue 1 to reach its allocated capacity of 4000 bits. When Packets 3 and 4 arrive, they are immediately identified as anomalous packets.

Since the squeezing policy is enabled with a threshold of 2000 bits, Packets 3 and 4 are redirected to Queue 2, while retaining their original timeslot 1 label. Based on the squeezing policy, packets 3 and 4 are now squeezed into timeslot 2 for transmission. At this point, the buffer depth of Queue 2 increases to 2000 bits. Subsequently, Packets 5, 6, and 7, which are targeted for timeslot 2, arrive and enter Queue 2. However, when Queue 2 reaches its allocated capacity of 4000 bits, Packet 7 is marked as anomalous. Packet 7 is then enqueued in Queue 2 and squeezed for transmission in timeslot 3.

At the aggregation node, continuous bursts may lead to successive squeezing, which could trigger a chain reaction. Without safeguards, packets squeezed from one timeslot into the next may accumulate indefinitely, undermining deterministic transmission guarantees. To prevent unbounded accumulation caused by consecutive squeezing, the following two safeguard mechanisms are introduced:

* Synchronization Threshold Mechanism: Defines a threshold (N) as the maximum number of consecutive timeslots permitted to be affected by squeezing. If squeezing occurs over N consecutive slots, the current queue must be resynchronized with the timeslot schedule to restore consistency and prevent unlimited delay accumulation.

* Exponential Decay Mechanism: When consecutive squeezing occurs, the allowed squeezing capacity decays exponentially. Specifically, the first affected timeslot permits a predefined squeezing capacity T; for each subsequent consecutive timeslot, the allowed squeezing capacity is reduced by 50% of the previous slot. This decay continues until the permitted capacity falls below the minimum packet size which then disallows further squeezing and triggering alternative handling (e.g., degrading).

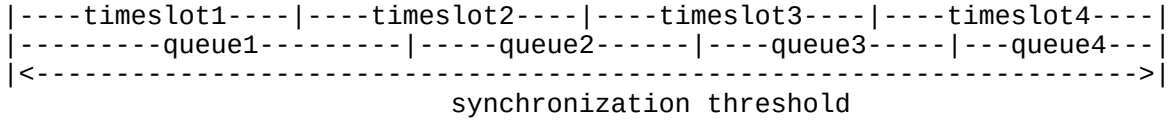

Figure 2: Illustration of synchronization threshold

5.2. Degrading Policy

The data plane supports the degrading policy and allows for the configuration of its parameters. This policy can be used either independently or in conjunction with the squeezing policy.

* When deployed with the squeezing policy, the degrading policy processes anomalous traffic that exceeds the squeezing threshold.

* When deployed independently, the degrading policy is applied directly to anomalous packets that exceed the allocated buffer capacity.

For EDF, packets are processed by adjusting their target sending time. The resultant delay time can be flexibly configured based on the congestion level at the outgoing port. For TAS/CQF and their variations, packets are redirected to a lower priority queue.

5.3. Squeezing Policy and Degrading Policy

When both squeezing and degrading policies are enabled, the node shall perform the following steps:

1. Upon packet arrival, determine whether the packet is anomalous.

2. If the squeezed resource count is below the squeezing threshold T, apply the squeezing policy to process the packet.

3. If the squeezed resource count exceeds T (or if consecutive squeezing has reached the synchronization threshold N or the exponential decay limit), immediately trigger the degrading policy by modifying the packet' s internal scheduling parameters and redirecting it to the appropriate lower-priority queue.

6. Anomalous Packets Handling Solution

6.1. Policy Selection and Configuration

The following anomaly handling policies are defined in this document:

* Degrading Policy: Process packets according to the degrading policy, which includes treating the packets as BE flow.

* Squeezing Policy: This policy provides temporary capacity expansion to avoid data loss due to unexpected traffic.

* Postponement Policy: Delays the transmission of packets until the next scheduling cycle.

* Redirection Policy: Redirect packets to a regular QoS queue.

* Discarding Policy: Discard anomalous packets.

If the squeezing or degrading policies are not enabled or are otherwise inapplicable, anomalous packets shall be processed by existing default methods, such as discarding. When the data plane supports multiple anomalous packets handling policies, the enabled policies and related parameters shall be configured by the control plane.

6.2. Anomalous Information Reporting

Once the data plane automatically handles anomalies using the squeezing policy or the degrading policy, it should promptly report these anomalies to the controller. This enables the controller to perceive detailed insights into the network anomalies and take appropriate actions, such as re-orchestration, flow entry re-configuration, resource expansion. In addition to reporting to the controller, the data plane should also transmit the anomaly information to the downstream nodes. This allows downstream nodes to adjust their forwarding behavior or restore the original parameters of the packets according to the received anomaly information. The anomaly information reported by the data plane includes, but is not limited to:

* Basic information: node ID, port ID, etc.

* Anomalous packet information: flow ID and packet sequence number, etc.

* Anomalous packet handling policy information:

  - Policy Type: Specifies the handling policy employed (e.g., squeezing, degrading, or default policies like discarding).

  - Related parameters:

    o For squeezing policy: Includes data such as the number of squeezed bits and the quantity of squeezed packets.

    o For the degrading policy: Includes data such as the priority levels before and after degrading, and the number of degraded packets.

    o For default policies: Includes information such as the number of discarded packets or treated as BE flows.

6.3. Anomalous Packets Handling Procedure

When a node in the data plane receives a DetNet packet, it first checks for anomalies. If an anomaly is detected, the node proceeds to handle the packet.

1. Identify Supported Policies. The node determines which anomalous packets handling policies are supported locally.

2. Policy-based Packet Processing.

* No Enhanced Policies Enabled: If the enhanced anomalous packets handling policies (i.e., the squeezing policy and the degrading policy) are not enabled, the anomalous packets shall be processed by the default mechanisms, such as direct discarding or treating the packets as Best-Effort (BE) flows.

* Single Policy Enabled: Process the anomalous packet using the enabled policy.

* Both Policies Enabled: If both the squeezing policy and degrading policy are enabled, the local node first checks whether the number of anomalous packets exceeds the squeezing threshold. If not, the squeezing policy is applied; otherwise, the degrading policy is applied.

3. Information Transmission

After processing the anomalous packets, the node SHOULD send the anomaly information to the controller and/or the downstream node.

7. Example

This example illustrates the anomaly detection and handling policy in the forwarding plane when the TQF is employed.

It is assumed that TQF mechanism supports three cycles (A, B, and C) at the egress ports. The timeslot size increases in powers of 2 while the number of timeslots decreases in powers of 2. Cycle A supports eight queues, and in addition, a low-priority BE queue is provided. For Cycle A, the timeslot mapping is defined as 0 -> 5; for the Cycle B, the mapping is 0 -> 3. It is assumed that each TQF timeslot in Cycle A allows a maximum capacity of 10,000 bits, Cycle B 20,000 bits, and Cycle C 40,000 bits. When the queue depth of Cycle A exceeds 10,000 bits, it indicates that an abnormal condition has occurred.

Furthermore, the control plane is configured to enable the squeezing policy on the forwarding plane with a squeezing threshold set to 15,000 bits and to enable the degrading policy, which is configured in a stepwise degrading mode.

Consider a certain service flow where each packet is 1,000 bits in size. Packets 1 to 10 use Cycle A and carry a timeslot value of 0; packets with sequence numbers 11 to 15 also use Cycle A, but carry a timeslot value of 2. When packet 1 arrives at the node, the current queue depth of timeslot 5 is 8,000 bits, and that of timeslot 7 is 0 bits.

Processing Procedure:

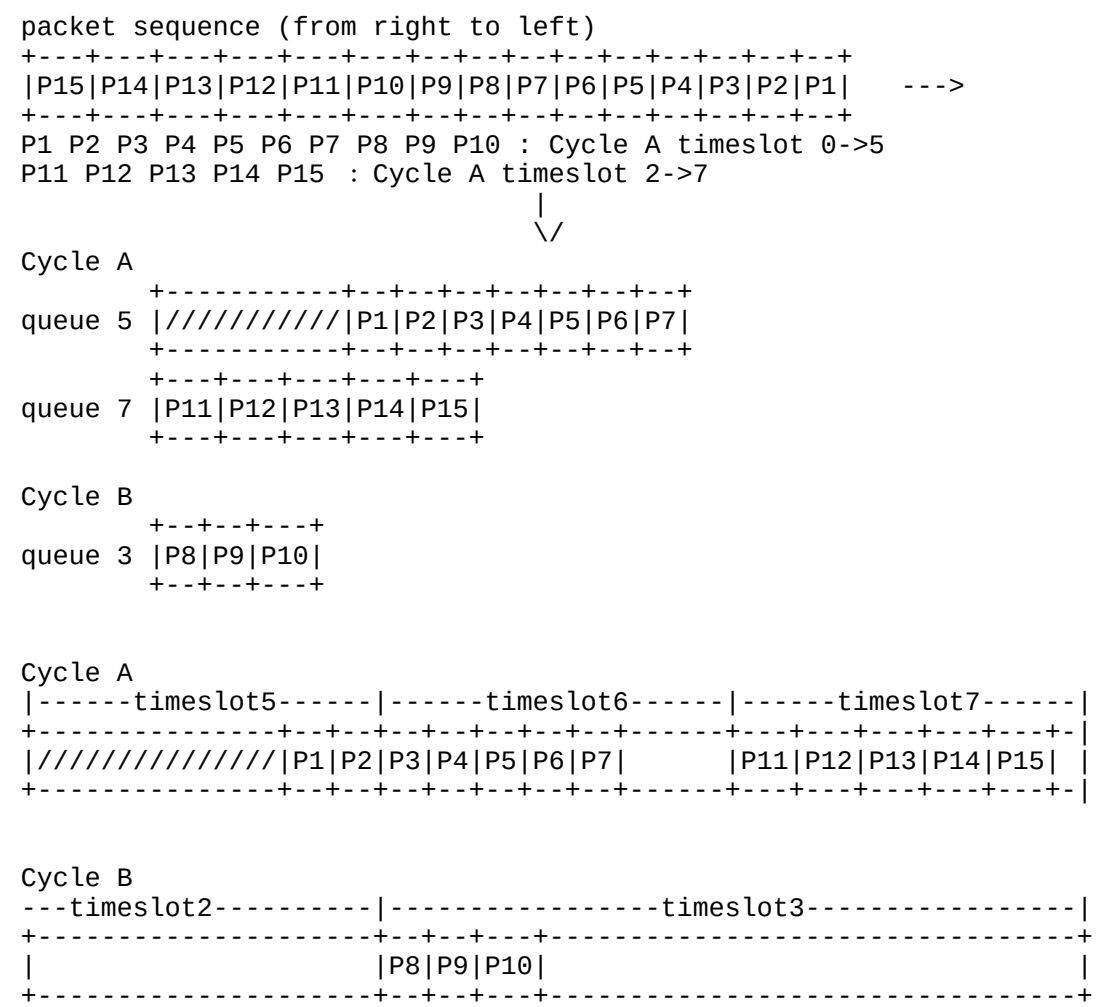

Figure 3: Example of Using the Anomalous Packets Handling Mechanism with TQF

When packets 1 and 2 are enqueued into queue 5 according to the Cycle A timeslot mapping 0 -> 5, the depth of queue 5 reaches 10,000 bits.Upon the arrival of packet 3, if it were to be enqueued using the same mapping (0 -> 5), the queue depth would exceed the 10,000-bit threshold, thereby indicating the presence of anormaly. Since the squeezing policy is enabled with a threshold of 15,000 bits, packets 3 to 7 are processed in squeezing mode and are enqueued into queue 5, retaining their original output timeslot label 5.

When packet 8 arrives, enqueuing it in queue 5 would cause the cumulative bits to exceed the 15,000-bit squeezing threshold. Consequently, the degrading policy is triggered. Packets 8 to 10 are degraded from Cycle A to Cycle B. Based on the Cycle A transmission timeslot value(0) carried in the packet, which is converted to Cycle B transmission timeslot 0, the Cycle B mapping (0 → 3) is applied. Thus, packets 810 are enqueued into Cycle B’ s Queue 3. Packets 11 to 15 mapped using timeslot 2 -> 7, are enqueued normally as the queue depth remains within the 10,000-bit capacity.

8. Security Considerations

TBA

9. IANA Considerations

TBA

10. Acknowledgements

TBA

11. References

11.1. Normative References

[I-D.peng-detnet-deadline-based-forwarding]
Peng, S., Du, Z., Basu, K., cheng, Yang, D., and C. Liu, "Deadline Based Deterministic Forwarding", 13 October 2025, <https://datatracker.ietf.org/doc/html/draft-peng-detnet-deadline-based-forwarding-18>.

[I-D.peng-detnet-packet-timeslot-mechanism]
Peng, S., Liu, P., Basu, K., Liu, A., Yang, D., and G. Peng, "Timeslot Queueing and Forwarding Mechanism", 12 October 2025, <https://datatracker.ietf.org/doc/html/draft-peng-detnet-packet-timeslot-mechanism-13>.

[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>.

[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>.

Authors' Addresses

Zhengxin Han
China Unicom
Beijing
China
Email: hanzx21@chinaunicom.cn

Ran Pang
China Unicom
Beijing
China
Email: pangran@chinaunicom.cn

Chang Liu
China Unicom
Beijing
China
Email: liuc131@chinaunicom.cn

Jinjie Yan
ZTE Corporation
China
Email: yan.jinjie@zte.com.cn

Xiangyang Zhu
ZTE Corporation
China
Email: zhu.xiangyang@zte.com.cn