DetNet
Internet-Draft
Intended status: Standards Track
Expires: 23 August 2025

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

19 February 2025

Anomalous Packets Handling for DetNet
draft-han-detnet-anomalous-packets-handling-00

Abstract

In deterministic networking (DetNet), there may be resource conflicts at the flow aggregation nodes, resulting in network anomalies. The existing mechanisms for handling anomalous packets in the data plane are crude, such as discarding or processing them as BE flows, so the network performance may be worse than applying traditional QoS. Therefore, in order to handle the anomalous traffic, the data plane should implement an enhanced handling mechanism.

This document proposes an anomalous packet handling solution for anomalous traffic in DetNet. This solution includes two policies: the packet squeezing policy and the packet degrading policy, which can be applied flexibly to a variety of queuing mechanisms, thereby ensuring that network traffic for deterministic services is preferentially scheduled in anomalous situations.

Status of This Memo

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]. Resource allocation reduces the packet loss and jitter caused by network congestion by allocating available resources to specified DetNet flows. The control plane orchestrates the paths of DetNet flows to avoid resource conflicts, while the data plane transmits DetNet flows based on the orchestration result from the control plane, with traffic shaping, flow admission control, and encapsulation of forwarding

information, etc., to maintain QoS.

Each node in the end-to-end path may serve as an aggregation node. Aggregated flows that belong to the same traffic class will share the reserved resources at the outgoing port. Ideally, the transmission of each flow within the same traffic class strictly conforms to the scheduling of the control plane, thus being able to meet the strict requirements of a narrowly deterministic network. However, due to the diversity of deterministic flows—such as occasional microbursts and packet size fluctuations—this ideal case is often difficult to fulfill. Allocating resources based on the maximum packet size may lead to waste, whereas 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 conflict at the aggregation nodes. As defined in [RFC8655], it proposes a service protection solution such as PREOF based on multi-path transmission. Although PREOF can prevent performance reduction by reserving a large amount of redundant resources for the specified service flows, it may cause a serious waste of resources or even a light load in the network, which further diminishes the advantage of deterministic technologies. In the data plane, the existing mechanisms are relatively simple and crude. 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 the solutions will result in even worse QoS than that of BE flows.

Therefore, the processing of anomalous packets from deterministic services should be automatically optimized in the data plane. The processing of anomalous packets is an indispensable part of the future implementation and application of the entire deterministic network technology.

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

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. Anomalous Forwarding Detection

The real-time detection in the data plane aims to identify anomalous forwarding behaviors. Upon detection, enhanced processing policies, such as packet squeezing and degrading, are applied to ensure that deterministic flows are scheduled preferentially, even under abnormal conditions.

The detection process is closely associated with the queuing mechanisms employed. Typically, for TQF[I-D.peng-detnet-packet-timeslot-mechanism], the target output timeslot of a packet at the current node can be determined based on the upstream timeslot label carried by the packet 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 preemptively placed into the corresponding queue. 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.

4. 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 packets to a lower-priority queue and modifies the scheduling parameters when the accumulation of anomalous packets exceeds a predefined threshold.

These policies provide flexibility in terms of activation; they can be enabled concurrently, selectively, or not at all. If neither policy is enabled, the default mechanism, such as discarding the packets or treating them as a BE flow will be utilized.

4.1. Squeezing Policy

The data plane can support the squeezing policy by allowing the configuration of the squeezing threshold. When anomalous traffic causes the queue occupancy to exceed its permitted capacity— but remains below the squeezing threshold— the system applies the squeezing policy. Specifically, the system will enqueue the packets and record the number of squeezed bits. According to the squeezing policy, packets that can not be sent within the allocated time should be squeezed into the next timeslot until the queue gets empty. It should be noted that the squeezing policy is compatible with various queuing mechanisms, although it may not be available in all. Regarding different queuing mechanisms, the implementation of the squeezing policy varies.

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.

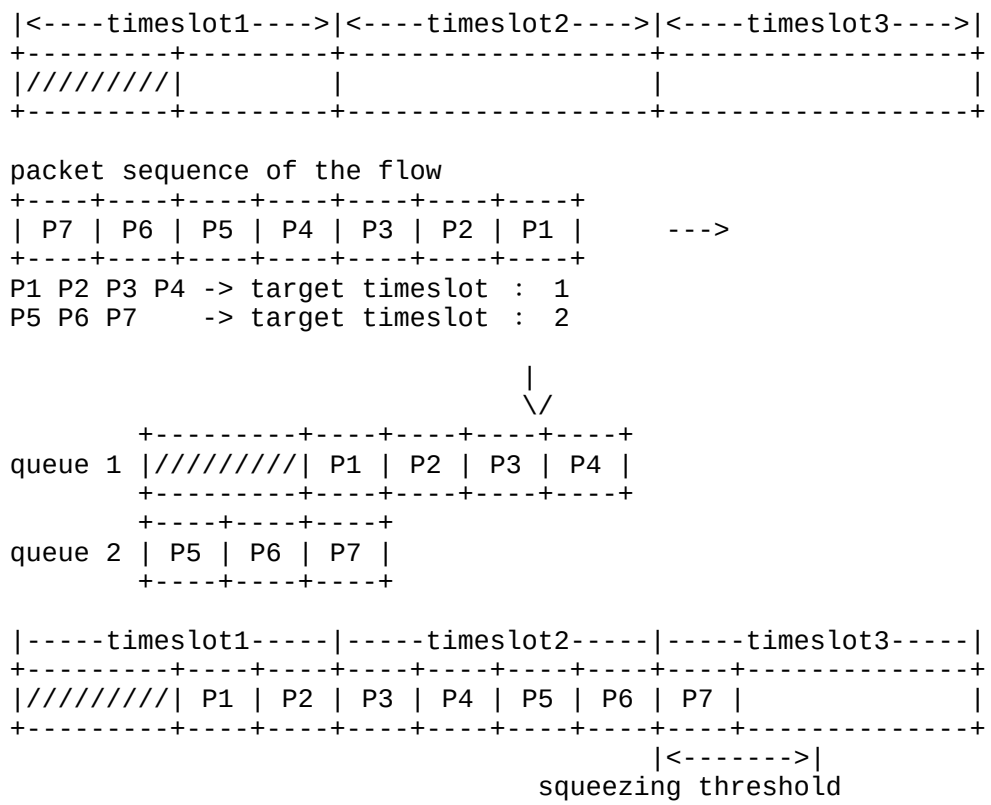

Figure 1: Squeezing policy based on timeslot-based queuing mechanism

Figure 1 illustrates the processing of packets in the service flow with serial numbers 1 through 7. Packets 1 and 2 are put into queue 1 sequentially. Therefore, queue 1 has reached the permitted carrying threshold of 4000 bits. When packets 3 and 4 arrive, they are determined to be anomalous packets.

Since the squeezing policy is enabled with a threshold of 2000 bits, packets 3 and 4 are enqueued in queue 2, while retaining their timeslot label of 1. Based on the squeezing policy, packets 3 and 4 are squeezed into timeslot 2 for transmission. At this point, the buffer depth of the queue increases to 2000 bits. Subsequently, packets 5, 6, and 7 which are targeted for timeslot 2, are allowed to enter queue 2. However, when queue 2 reaches its upper limit of 4000 bits, packet 7 is marked as an anomalous packet. It is enqueued in queue 2 and postponed for transmission in timeslot 3.

At the aggregation node, continuous bursts may lead to successive squeezing, which in turn may 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: A "synchronization threshold" is defined as the maximum number of consecutive timeslots that may be affected by squeezing. For example, if the threshold is set to N timeslots, once squeezing has occurred over N consecutive slots, the current queue must be re-synchronized with the timeslot schedule. This re-synchronization restores scheduling consistency and prevents indefinite delay accumulation.

* Exponential Decay Mechanism: In scenarios with consecutive squeezing, 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% compared to the previous slot. This decay continues until the permitted capacity falls below the minimum packet size, at which point further squeezing is disallowed, and alternative handling (e.g., degrading) is triggered.

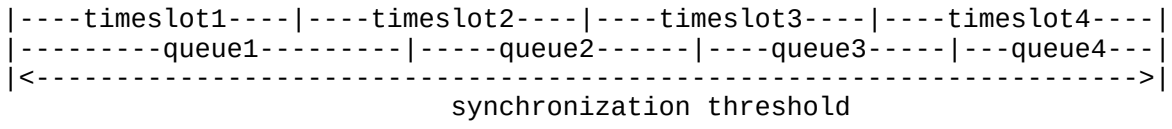

Figure 2: Illustration of synchronization threshold

4.2. Degrading Policy

The data plane supports the degrading policy and allows for the configuration of degrading parameters, and can be used either in combination or independently. When the degrading policy is used in conjunction with the squeezing policy, it processes anomalous traffic that exceeds the squeezing threshold. The degrading policy can also be deployed on its own. For anomalous packets that go beyond the allowed buffer capacity, the degrading policy can be directly applied.

For EDF, packets are delayed based on the target sending time. The delayed period can be flexibly configured due to the level of busyness at the current outgoing port. For TAS/CQF and their variations, packets are redirected to a queue with a lower priority.

4.3. 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 is anomalous packet.

2. If the abnormal accumulation is below the squeezing threshold T, process the packet by applying the squeezing policy.

3. If the abnormal accumulation exceeds T (or if consecutive squeezing has reached the synchronization threshold N or the exponential decay limit), immediately trigger the degrading policy. This involves modifying the packet's internal scheduling parameters (as detailed above) and redirecting it to the appropriate lower-priority queue.

5. Anomalous Packets Handling Solution

5.1. Policy Selection and Configuration

The following anomaly handling policies are involved in this document:

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

* Squeezing Policy: Process packets according to the squeezing policy. This policy provides temporary capacity expansion to avoid data loss due to unexpected traffic.

* Postponement Policy: Postpone packets to the next cycle.

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

* Discarding Policy: Discard anomalous packets.

If the data plane does not enable either the squeezing or degrading policy, or if neither policy is applicable, anomalous packets will be processed by the existing default methods, such as discarding. When the data plane supports multiple anomalous packets handling policies, the enabled policies and related parameters can be configured by the control plane.

5.2. Anomalous Information Reporting

Once the data plane automatically handles anomalies using either 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, etc. In addition to reporting to the controller, the data plane should also transmit the anomaly information to the downstream nodes. This allows downstream nodes 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, which could be the squeezing policy, the degrading policy, or other default policies (e.g., discarding).

  - Related parameters: For squeezing policy: Includes data such as the number of squeezed bits and the quantity of squeezed packets. For the degrading policy: Includes data such as the priority levels before and after degrading, and the number of degraded packets. For default policies: Includes information such as the number of discarded packets or treated as BE flows.

5.3. 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 initiates the procedure for anomalous packets.

1. Identify Supported Policies. The node needs to determine 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 packets will be processed by the default mechanisms which may be directly discarding, treating the packets as BE flow, processing them in a normal QoS queue, or postponing them to the next period.

   * 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, process the packet using the squeezing policy; otherwise, apply the degrading policy.

3. Information Transmission

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

6. Example

This 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. In these cycles, 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 cycle A and carry a timeslot value of 0; packets with sequence numbers 11 to 15 also use cycle 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 abnormaly. 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 the 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 so that 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.

7. Security Considerations

TBA

8. IANA Considerations

TBA

9. Acknowledgements

TBA

10. References

10.1. 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", 20 June 2024, <https://datatracker.ietf.org/doc/html/draft-peng-detnet-deadline-based-forwarding-10>.

[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", 20 June 2024, <https://datatracker.ietf.org/doc/html/draft-peng-detnet-packet-timeslot-mechanism-07>.

[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