LSR
Internet-Draft
Intended status: Standards Track
Expires: 12 July 2026

Shaofu. Peng
ZTE

8 January 2026

IGP Extensions for Deterministic Traffic Engineering
draft-peng-lsr-deterministic-traffic-engineering-04

Abstract

This document describes IGP extensions to support Traffic Engineering (TE) of deterministic routing, by specifying new information that a router can place in the advertisement of neighbors. This information describes additional details regarding the state of the network that are useful for deterministic traffic engineering path computations.

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 12 July 2026.

Table of Contents

1. Introduction

[RFC8655] describes the architecture of a deterministic networking (DetNet) and defines the QoS goals of deterministic forwarding: 1) Minimum and maximum end-to-end latency from source to destination, timely delivery, and bounded jitter (packet delay variation); 2) A bounded packet loss ratio under various assumptions about the operational states of the nodes and links; 3) An upper bound on out-of-order packet delivery. In order to achieve these goals, DetNet use resource reservation, explicit routing, and service protection, as well as other means. A deterministic forwarding path is typically (but not necessarily) an explicit route so that it does not suffer temporary interruptions caused by the convergence of routing or bridging protocols.

The IEEE 802.1 WG has specified a set of queuing, shaping, and scheduling algorithms that are mainly used for L2 networks, such as ATS [IEEE802.1Qcr], CBS [IEEE802.1Qav], and CQF/ECQF [IEEE802.1Qch] [IEEE802.1Qdv]. There are some challenges in applying these mechanisms to IP/MPLS network, mainly due to the cost of per flow state, or limited service scale. There are also some enhanced data plane (EDP) queueing mechanisms under discussion in DetNet working group to meet large scaling requirements, such as C-SCORE [I-D.joung-detnet-stateless-fair-queuing], EDF [I-D.peng-detnet-deadline-based-forwarding], TQF

[I-D.peng-detnet-packet-timeslot-mechanism], gLBF [I-D.eckert-detnet-glbf]. According to [Net-Calculus], queueing mechanisms may be roughly classified into two categories: rate based, and delay based. For example, ATS, CBS, C-SCORE, and gLBF are rate based mechanisms, while CQF/ECQF, EDF, and TQF are delay based mechanisms. It can also be classified from other dimensions, such as work-conserving or non work-conserving. One mechanism may correspond to multiple categories simultaneously.

The latency bound provided by rate based mechanisms is generally inversely proportional to the service rate of the flow or flow aggregate, and overestimated. While the latency bound provided by delay based mechanisms is related to the time resources consumed by the flow by accurately planning the scheduling orders (e.g., slot, deadline), and is an accurate preset value.

In order to provide DetNet QoS, queueing mechanisms not only specify scheduling techniques on the data plane, but also quantify forwarding resources to provide a foundation for admission control of service flows. Generally, the successful operation of a DetNet queueing mechanism relies on admission control of the data volume and data rate allowed to be released into the network, and avoiding burst accumulation at intermediate nodes. The deterministic forwarding resources will be designed around two factors: data volume and data rate. Note that in L4S architecture [RFC9330], the endpoint of transport layer also dynamically control these two factors in response to network congestion indications to get low latency target, however, it rely on heuristics that can either undershoot or overshoot the bottleneck bandwidth, and latency bound cannot be guaranteed.

This document describes IGP extensions to advertise forwarding resources related with deterministic queueing mechanisms in the network, which may be used for the deterministic traffic engineering path computation.

2. Requirements Language

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.

3. Deterministic Forwarding Resources

One or more queueing mechanisms may be enabled on the same link, each of which may support single or multiple instances (or considered as multiple capability levels), and each instance has dedicated deterministic forwarding resources. For example, the traditional Strict Priority (SP) mechanism may support 8 traffic classes and each has the Maximum Reservable Bandwidth resource. It is known that SP faces challenges in providing DetNet QoS. For any queueing mechanism that can guarantee DetNet QoS, it is similar to support multiple instances, especially the deterministic forwarding resources involved can be summarized into two types:

* Bandwidth: refers to the share of link capacity allocated by an instance. In some contexts, bandwidth is also replaced by terms such as service rate or data rate.

* Burst: refers to the data volume that an instance is allowed to send during a busy period of scheduling.

For ATS, CBS, or gLBF, they may support multiple instances (such as 8 traffic classes), and each instance has dedicated Maximum Reservable Bandwidth (MRBan) and Maximum Reservable Burst (MRBur). DetNet flows mapped to a certain instance will consume the resources of that instance. The MRBan and MRBur corresponding to a certain instance are the dominator factors for the worst-case per-hop latency for that instance.

For C-SCORE, it may be considered to support a single instance, and have Maximum Reservable Bandwidth (MRBan) and Maximum Reservable Burst (MRBur). DetNet flows mapped to C-SCORE will consume the resources of this instance. The Maximum Reservable Burst resource provided by C-SCORE is actually determined by the physical size of the used sorted-queue, which stores all concurrent incoming bursts. However, the MRBan and MRBur of C-SCORE are only used for admitting condition check, and the worst-case per-hop latency for each flow is only determined by flow’s service rate and burst size.

For CQF/ECQF, they may support multiple instances each with specific cycle duration (e.g., 10us). Each instance has dedicated Maximum Reservable Bandwidth (MRBan) and Maximum Reservable Burst (MRBur), where MRBur = MRBan * cycle duration. Note that MRBur represents the resources of the entire instance, not the resources of a specific cycle under the instance (e.g., a instance may have cycle a, b, c). DetNet flows mapped to a certain instance will consume the resources of that instance. However, the MRBan and MRBur of CQF/ECQF are only used for admitting condition check, and the worst-case per-hop delay is determined by the cycle duration.

For EDF, it may support multiple instances each with specific delay level (e.g., 10us). Each instance has dedicated Maximum Reservable Bandwidth (MRBan) and Maximum Reservable Burst (MRBur). The MRBan and MRBur of all instances meet the shcedulability condition. DetNet flows mapped to a certain instance will consume the resources of that instance. However, the MRBan and MRBur of EDF are only used for admitting condition check, and the worst-case per-hop delay is determined by the delay level.

For TQF, it may support multiple instances each with specific orchestration period (e.g., 1ms) that containing N timeslots. Each instance has dedicated Maximum Reservable Bandwidth (MRBan) and Maximum Reservable Burst (MRBur), where MRBur = MRBan * timeslot length. Note that MRBur represents resources for individual timeslot, and in general all timeslots have the same MRBur value. DetNet flows mapped to a certain instance will consume the resources of that instance. However, the MRBan and MRBur of TQF are only used for admitting condition check, and the worst-case per-hop delay is determined by the timeslot length, as well as the mapping relationship between the incoming and outgoing timeslots.

The link should also maintain the unused resources of each capability level based on the reservation result, i.e., Unused Bandwidth and Unused Burst. Initially, unused resources are equal to the maximum available reservable resources, and only the maximum available reservable resources need to be advertised by IGP if there is no consumption of resources. The utilized resources are equal to the maximum resources minus the unused resources.

The rate based queueing mechanisms can always use MRBan and MRBur to calculate latency bound, as long as the reserved bandwidth and burst resources consumed by admitted flows do not exceed MRBan and MRBur, respectively. Alternatively, it may also use the utilized bandwidth and bursts to calculate latency bound, which however will change with the dynamic admission and release of flows.

EDF mechanism can always use MRBan and MRBur to check the schedulability condition, as long as the reserved bandwidth and burst resources consumed by admitted flows do not exceed MRBan and MRBur, respectively. Alternatively, whenever a new flow is admitted, it uses the utilized bandwidth and bursts to check the schedulability condition.

4. ISIS Advertisement of Link Scheduling Capability

A new IS-IS sub-TLV is defined: the DetNet Scheduling Capability Sub-TLV, which is advertised within TLV-22, 222, 23, 223, 141, 25. For each link, multiple DetNet Scheduling Capability Sub-TLVs can be included, depending on how many scheduling mechanisms are enabled on the link.

The following format is defined for the DetNet Scheduling Capability Sub-TLV:

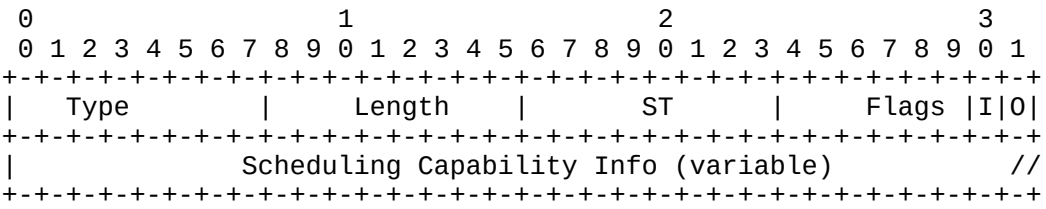

Figure 1

where:

Type: TBD.

Length: variable, depending on size of the Scheduling Capability Info field.

ST(Scheduling Type): 1 byte, represents the type of scheduling mechanism supported by the link, as below.

0: Reserved for default or unspecified scheduling mechanisms, such as SP (strict priority) that is widely used in the network. It is not recommended to explicitly advertise the detailed capability information of default scheduling mechanisms by the DetNet Scheduling Capability Sub-TLV.

1: ATS [IEEE802.1Qcr].

2: CBS [IEEE802.1Qav].

3: ATS+CBS [ATSplusCBS].

4: CQF/ECQF [IEEE802.1Qch] [IEEE802.1Qdv].

5: EDF [I-D.peng-detnet-deadline-based-forwarding].

6: TQF [I-D.peng-detnet-packet-timeslot-mechanism].

7: C-SCORE [I-D.joung-detnet-stateless-fair-queuing].

8: gLBF [I-D.eckert-detnet-glbf].

9~255: To be defined in the future.

Flags: 1 byte, currently two flags are defined as below:

I (In-time mode): indicates whether the scheduling mechanism supports in-time scheduling mode. Support if set, otherwise not support. In-time scheduling mode can be understood as sending the packet as soon as possible before its bounded latency per-hop.

O (On-time mode): indicates whether the scheduling mechanism supports on-time scheduling mode. Support if set, otherwise not support. On-time scheduling mode can be understood as sending the packet as close as possible to its bounded latency per-hop.

Scheduling Capability Info: Includes capability level information corresponding to the specific scheduling mechanism type with variable size, depending on the ST.

- If ST is one of ATS, CBS, ATS+CBS, or gLBF, the field size is 1 byte, and it contains the number (i.e., n) of traffic classes supported by the scheduling mechanism. Let the first traffic class be 0, the last traffic class be n-1, with ascending priority order from traffic class 0 to traffic class n.

- If ST is CQF/ECQF, the field size is n*2 bytes, and it contains n cycle durations, each with 2 bytes, in unit of microseconds. For example, the cycle duration may be 10 us, or 20 us, and so on. Different cycle durations represent different CQF instances.

- If ST is EDF, the field size is 6 bytes, and it contains the minimum delay level (2 bytes, in unit of microseconds), maximum delay level (2 bytes, in unit of microseconds), and delay level interval (2 bytes, in unit of microseconds) supported by the EDF scheduling mechanism. The number of supported delay levels can be deduced by n = (maximum delay level - minimum delay level) / delay level interval + 1. For example, the minimum delay level may be 10 us, the maximum delay level may be 100 us, and the delay level interval may be 10 us.

- If ST is TQF, the field size is n*8 bytes, and it contains n TQF instances, each with 8 bytes. These 8 bytes specifically include the Orchestration Period Length (4 bytes, in unit of microseconds), the amount of timeslots N (2 bytes) within the Orchestration Period, and the amount of timeslots M (2 bytes) within the Scheduling Period. The timeslot length can be deduced by Orchestration Period Length / N. For example, the Orchestration Period Length may be 1000 us, containing 100 timeslots, and the Scheduling Period may contain fewer timeslots, such as 10.

- If ST is C-SCORE, the field size is zero, and there is no need to specify capability level information. It can be considered to support a single unified instance.

For each scheduling mechanism enabled on the link, the DetNet Scheduling Capability Sub-TLV SHOULD be advertised once at most. A router receiving multiple DetNet Scheduling Capability Sub-TLVs for the same link and same scheduling mechanism, SHOULD select the first advertisement in the lowest-numbered LSP.

5. ISIS Advertisement of DetNet Maximum Reservable Bandwidth

A new IS-IS sub-TLV is defined: the DetNet Maximum Reservable Bandwidth Sub-TLV, which is advertised within TLV-22, 222, 23, 223, 141, 25. For each link, multiple DetNet Maximum Reservable Bandwidth Sub-TLVs can be included, depending on how many scheduling mechanisms are enabled on the link.

This sub-TLV contains the maximum amount of bandwidth that can be reserved in the link with the direction from this node to the neighbor, for each instance of a specific scheduling mechanism. Note that oversubscription is prohibited, so this must be less than the bandwidth of the link.

The following format is defined for the DetNet Maximum Reservable Bandwidth Sub-TLV:

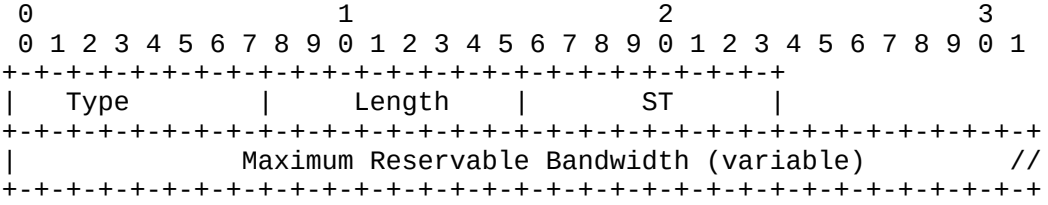

Figure 2

where:

Type: TBD.

Length: variable, depending on size of the Maximum Reservable Bandwidth field.

ST(Scheduling Type): 1 byte, represents the type of scheduling mechanism supported by the link, as defined in Section 4.

Maximum Reservable Bandwidth: includes the maximum reservable bandwidth (MRBan) corresponding to the specific scheduling mechanism type with variable size, depending on the ST.

- If ST is one of ATS, CBS, ATS+CBS, or gLBF, the field size is n*4 bytes, and it contains the MRBan per traffic class (4 bytes, in the unit of bytes per second in IEEE floating point format), from traffic class 0 to traffic class n-1.

- If ST is CQF/ECQF, the field size is n*6 bytes, and it contains n tuple <cycle duration(2B), MRBan(4B)>, where, cycle duration in the unit of microseconds, and MRBan in the unit of bytes per second in IEEE floating point format.

- If ST is EDF, the field size is n*6 bytes, and it contains n tuple <delay level(2B), MRBan(4B)>, where, delay level in the unit of microseconds, and MRBan in the unit of bytes per second in IEEE floating point format. Note that all delay levels' maximum reservable bandwidth must meet the schedulability condition.

- If ST is TQF, the field size is n*8 bytes, and it contains n tuple <OPL(4B), MRBan(4B)>, where, OPL (Orchestration Period Length) in the unit of microseconds, and MRBan in the unit of bytes per second in IEEE floating point format.

- If ST is C-SCORE, the field size is 4 bytes, and it contains the MRBan in the unit of bytes per second in IEEE floating point format.

For each scheduling mechanism enabled on the link, the DetNet Maximum Reservable Bandwidth Sub-TLV SHOULD be advertised once at most. A router receiving multiple DetNet Maximum Reservable Bandwidth Sub-TLVs for the same link and same scheduling mechanism, SHOULD select the first advertisement in the lowest-numbered LSP.

Note that oversubscription is prohibited, so that the sum of MRBan of all scheduling mechanisms must be less than the link capacity.

6. ISIS Advertisement of DetNet Unreserved Bandwidth

A new IS-IS sub-TLV is defined: the DetNet Unreserved Bandwidth Sub-TLV, which is advertised within TLV-22, 222, 23, 223, 141, 25. For each link, multiple DetNet Unreserved Bandwidth Sub-TLVs can be included, depending on how many scheduling mechanisms are enabled on the link.

This sub-TLV contains the amount of bandwidth reservable in the link with the direction from this node to the neighbor, for each instance of a specific scheduling mechanism. Initially, the unreserved bandwidth equals to the maximum reservable bandwidth.

The following format is defined for the DetNet Unreserved Bandwidth Sub-TLV:

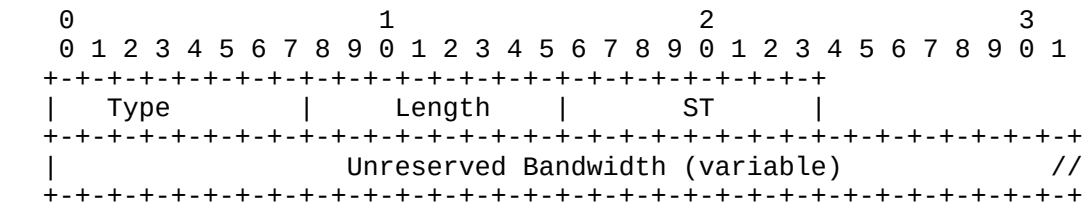

Figure 3

where:

Type: TBD.

Length: variable, depending on size of the Unreserved Bandwidth field.

ST(Scheduling Type): 1 byte, represents the type of scheduling mechanism supported by the link, as defined in Section 4.

Unreserved Bandwidth: includes the unreserved bandwidth (UBan) corresponding to the specific scheduling mechanism type with variable size, depending on the ST.

- If ST is one of ATS, CBS, ATS+CBS, or gLBF, the field size is n*4 bytes, and it contains the UBan per traffic class (4 bytes, in the unit of bytes per second in IEEE floating point format), from traffic class 0 to traffic class n-1.

- If ST is CQF/ECQF, the field size is n*6 bytes, and it contains n tuple <cycle duration(2B), UBan(4B)>, where, cycle duration in the unit of microseconds, and UBan in the unit of bytes per second in IEEE floating point format.

- If ST is EDF, the field size is n*6 bytes, and it contains n tuple <delay level(2B), UBan(4B)>, where, delay level in the unit of microseconds, and UBan in the unit of bytes per second in IEEE floating point format.

- If ST is TQF, the field size is n*8 bytes, and it contains n tuple <OPL(4B), UBan(4B)>, where, OPL (Orchestration Period Length) in the unit of microseconds, and UBan in the unit of bytes per second in IEEE floating point format.

- If ST is C-SCORE, the field size is 4 bytes, and it contains the UBan in the unit of bytes per second in IEEE floating point format.

For each scheduling mechanism enabled on the link, the DetNet Unreserved Bandwidth Sub-TLV SHOULD be advertised once at most. A router receiving multiple DetNet Unreserved Bandwidth Sub-TLVs for the same link and same scheduling mechanism, SHOULD select the first advertisement in the lowest-numbered LSP.

7. ISIS Advertisement of DetNet Maximum Reservable Burst

A new IS-IS sub-TLV is defined: the DetNet Maximum Reservable Burst Sub-TLV, which is advertised within TLV-22, 222, 23, 223, 141, 25. For each link, multiple DetNet Maximum Reservable Burst Sub-TLVs can be included, depending on how many scheduling mechanisms are enabled on the link.

The following format is defined for the DetNet Maximum Reservable Burst Sub-TLV:

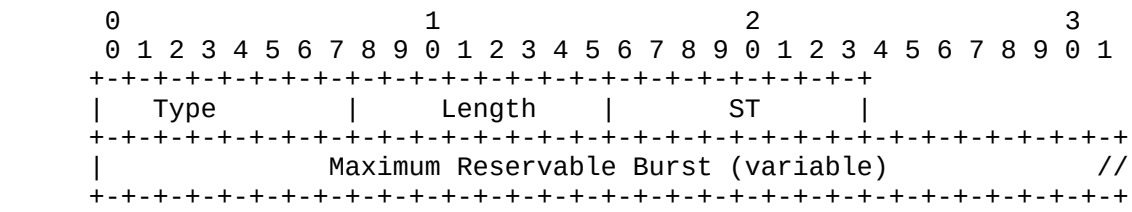

Figure 4

where:

Type: TBD.

Length: variable, depending on size of the Maximum Reservable Burst field.

ST(Scheduling Type): 1 byte, represents the type of scheduling mechanism supported by the link, as defined in Section 4.

Maximum Reservable Burst: includes the maximum reservable burst (MRBur) corresponding to the specific scheduling mechanism type with variable size, depending on the ST.

- If ST is one of ATS, CBS, ATS+CBS, or gLBF, the field size is n*4 bytes, and it contains the MRBur per traffic class (4 bytes, in the unit of bytes), from traffic class 0 to traffic class n-1.

- If ST is CQF/ECQF, the field size is n*6 bytes, and it contains n tuple <cycle duration(2B), MRBur(4B)>, where, cycle duration in the unit of microseconds, and MRBur in the unit of bytes. Note that MRBur represents the resources of the entire CQF instance, not the resources of a specific cycle (or bin) under that instance (e.g., a CQF instance may configure 3-bin mode, with bin a, b, c).

- If ST is EDF, the field size is n*6 bytes, and it contains n tuple <delay level(2B), MRBur(4B)>, where, delay level in the unit of microseconds, and MRBur in the unit of bytes. Note that all delay levels' maximum reservable burst must meet the schedulability condition equation.

- If ST is TQF, the field size is n*8 bytes, and it contains n tuple <OPL(4B), MRBur(4B)>, where, OPL (Orchestration Period Length) in the unit of microseconds, and MRBur in the unit of bytes. Note that MRBur represents resources for individual timeslot within the Orchestration Period, and in general all timeslots have the same MRBur value.

- If ST is C-SCORE, the field size is 4 bytes, and it contains the MRBur in the unit of bytes.

For each scheduling mechanism enabled on the link, the DetNet Maximum Reservable Burst Sub-TLV SHOULD be advertised once at most. A router receiving multiple DetNet Maximum Reservable Burst Sub-TLVs for the same link and same scheduling mechanism, SHOULD select the first advertisement in the lowest-numbered LSP.

8. ISIS Advertisement of DetNet Unreserved Burst

A new IS-IS sub-TLV is defined: the DetNet Unreserved Burst Sub-TLV, which is advertised within TLV-22, 222, 23, 223, 141, 25. For each link, multiple DetNet Unreserved Burst Sub-TLVs can be included, depending on how many scheduling mechanisms are enabled on the link.

The following format is defined for the DetNet Unreserved Burst Sub-TLV:

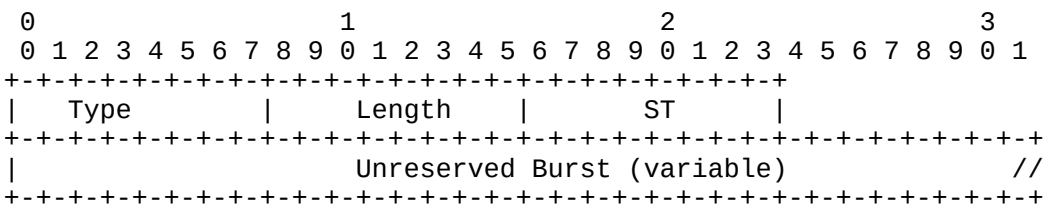

Figure 5

where:

Type: TBD.

Length: variable, depending on size of the Unreserved Burst field.

ST(Scheduling Type): 1 byte, represents the type of scheduling mechanism supported by the link, as defined in Section 4.

Unreserved Burst: includes the unreserved burst (UBur) corresponding to the specific scheduling mechanism type with variable size, depending on the ST.

- If ST is one of ATS, CBS, ATS+CBS, or gLBF, the field size is n*4 bytes, and it contains the UBur per traffic class (4 bytes, in the unit of bytes), from traffic class 0 to traffic class n-1.

- If ST is CQF/ECQF, the field size is n*6 bytes, and it contains n tuple <cycle duration(2B), UBur(4B)>, where, cycle duration in the unit of microseconds, and UBur in the unit of bytes. Note that UBur represents the resources of the entire CQF instance, not the resources of a specific cycle (or bin) under that instance (e.g., a CQF instance may configure 3-bin mode, with bin a, b, c).

- If ST is EDF, the field size is n*6 bytes, and it contains n tuple <delay level(2B), UBur(4B)>, where, delay level in the unit of microseconds, and UBur in the unit of bytes.

- If ST is TQF, the field size is 4+k*6 bytes, and it contains OPL(4B) and k (k <= N) tuple <timeslot id(2B), UBur(4B)>, where, OPL (Orchestration Period Length) in the unit of microseconds, and UBur in the unit of bytes. Note that UBur represents resources for individual timeslot within the Orchestration Period.

- If ST is C-SCORE, the field size is 4 bytes, and it contains the UBur in the unit of bytes.

For each scheduling mechanism enabled on the link, the DetNet Unreserved Burst Sub-TLV SHOULD be advertised once at most, except that TQF may advertise multiple DetNet Unreserved Burst Sub-TLVs each for a set of timeslots. A router receiving multiple DetNet Unreserved Burst Sub-TLVs for the same link and same scheduling mechanism (and same timeslot id in the case of TQF scheduling type), SHOULD select the first advertisement in the lowest-numbered LSP.

9. ISIS Advertisement of DetNet Timeslot Mapping

A new IS-IS sub-TLV is defined: the DetNet Timeslot Mapping Sub-TLV, which is advertised within TLV-22, 222, 23, 223, 141, 25. For each link, multiple DetNet Timeslot Mapping Sub-TLVs can be included, depending on how many instances are enabled on the link. Note that each instance has specific period length which contains multiple consecutive cyclic timeslots, e.g, period 100 microseconds may contain 10 consecutive cyclic timeslots (from #0 to #9) with each timeslot 10 microseconds.

[I-D.xp-ippm-detnet-stamp] defines a method to detect timeslot mapping. This mapping represents a timeslot id (e.g, x) of the output port of the local node, after link propagation delay, falls into the timeslot id (e.g, y) of the incomming port of the remote node. [ECQF] definds a method to detect cycle mapping, in this case, cycle is equivalent to timeslot.

The following format is defined for the DetNet Timeslot Mapping Sub-TLV:

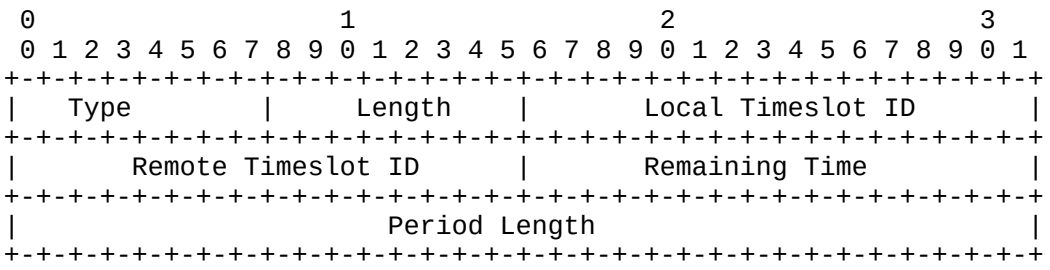

Figure 6

where:

Type: TBD.

Length: 10.

Local Timeslot ID: The timeslot id of the output port of the local node.

Remote Timeslot ID: The timeslot id of the incomming port of the remote node that mapped by Local Timeslot ID.

Remaining Time: The offset between the end of the Local Timeslot ID and the end of the Remote Timeslot ID.

Period Length: The dinstinguisher of the instance which contains multiple consecutive cyclic timeslots.

For each period instance enabled on the link, the DetNet Timeslot Mapping Sub-TLV SHOULD be advertised once at most. A router receiving multiple DetNet Timeslot Mapping Sub-TLVs for the same link and same period instance, SHOULD select the first advertisement in the lowest-numbered LSP.

10. ISIS Advertisement of DetNet Period Mapping

A new IS-IS sub-TLV is defined: the DetNet Period Mapping Sub-TLV, which is advertised within TLV-22, 222, 23, 223, 141, 25. For each link, multiple DetNet Period Mapping Sub-TLVs can be included, depending on how many instances are enabled on the link. Note that each instance has specific period length which contains multiple consecutive cyclic timeslots, e.g, period 100 microseconds may contain 10 consecutive cyclic timeslots (from #0 to #9) with each timeslot 10 microseconds.

[I-D.xp-ippm-detnet-stamp] defines a method to detect period mapping. This mapping represents the end of the priod of the output port of the local node, after link propagation delay, falls into the period of the incomming port of the remote node.

The following format is defined for the DetNet Period Mapping Sub-TLV:

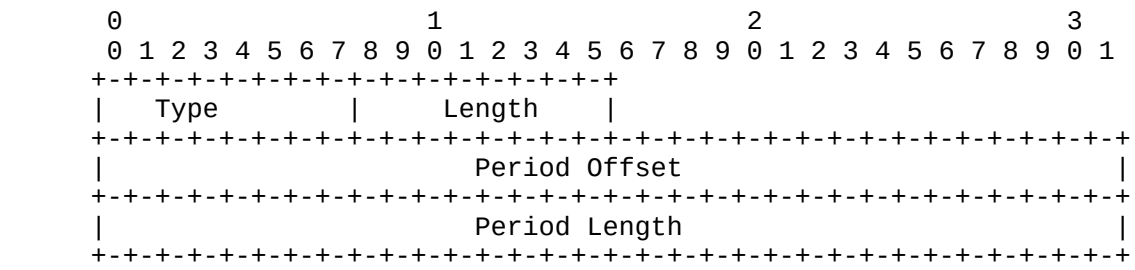

Figure 7

where:

Type: TBD.

Length: 8.

Period Offset: The offset between the end of the local period and the end of the remote period.

Period Length: The dinstinguisher of the instance which contains multiple consecutive cyclic timeslots.

For each period instance enabled on the link, the DetNet Period Mapping Sub-TLV SHOULD be advertised once at most. A router receiving multiple DetNet Period Mapping Sub-TLVs for the same link and same period instance, SHOULD select the first advertisement in the lowest-numbered LSP.

11. ISIS Advertisement of DetNet Intra-node Forwarding Delay

A new IS-IS sub-TLV is defined: the DetNet Intra-node Forwarding Delay Sub-TLV, which is advertised within TLV-242. At most one DetNet Intra-node Forwarding Delay Sub-TLV can be included. It represents the maximum forwarding delay experienced by a packet from the node’s incomming port (or the control plane of this node) to the output port, depending on the node’s implementation. It should be distinguished from the queuing delay on the output port and link propagation delay.

The following format is defined for the DetNet Intra-node Forwarding Delay Sub-TLV:

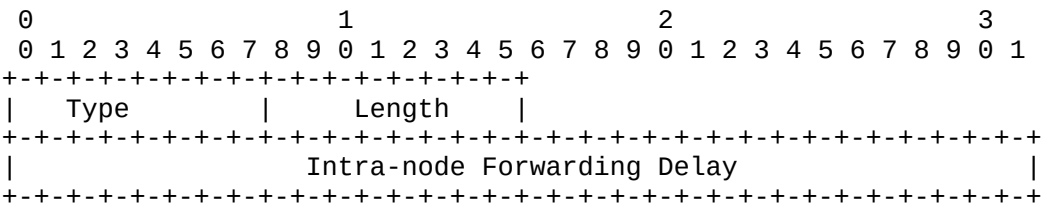

Figure 8

where:

Type: TBD.

Length: 4.

Intra-node Forwarding Delay: The maximum forwarding delay experienced by a packet from the node’s incomming port (or the control plane of this node) to the output port.

A router receiving multiple DetNet DetNet Intra-node Forwarding Delay Sub-TLVs for the same node, SHOULD select the first advertisement in the lowest-numbered LSP.

12. OSPF Advertisement of Link Deterministic Resource

Provided in next versions.

13. Announcement Suppression

To prevent oscillations and unnecessary advertisements, implementations MUST comply with the requirements found in sections 5 and 6 of [RFC8570] regarding announcement thresholds, filters, and suppression.

14. IANA Considerations

TBD

15. Security Considerations

This document introduces no new security issues. Security of routing within a domain is already addressed as part of the routing protocols themselves. This document proposes no changes to those security architectures.

The authentication methods described in [RFC5304] and [RFC5310] for IS-IS, [RFC2328] and [RFC7474] for OSPFv2 and [RFC5340] and [RFC4552] for OSPFv3 SHOULD be used to prevent attacks on the IGPs.

16. Acknowledgements

TBD.

17. References

17.1. 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>.

[RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, DOI 10.17487/RFC2328, April 1998, <https://www.rfc-editor.org/info/rfc2328>.

[RFC4552] Gupta, M. and N. Melam, "Authentication/Confidentiality for OSPFv3", RFC 4552, DOI 10.17487/RFC4552, June 2006, <https://www.rfc-editor.org/info/rfc4552>.

[RFC5304] Li, T. and R. Atkinson, "IS-IS Cryptographic Authentication", RFC 5304, DOI 10.17487/RFC5304, October 2008, <https://www.rfc-editor.org/info/rfc5304>.

[RFC5310] Bhatia, M., Manral, V., Li, T., Atkinson, R., White, R., and M. Fanto, "IS-IS Generic Cryptographic Authentication", RFC 5310, DOI 10.17487/RFC5310, February 2009, <https://www.rfc-editor.org/info/rfc5310>.

[RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, DOI 10.17487/RFC5340, July 2008, <https://www.rfc-editor.org/info/rfc5340>.

[RFC7474] Bhatia, M., Hartman, S., Zhang, D., and A. Lindem, Ed., "Security Extension for OSPFv2 When Using Manual Key Management", RFC 7474, DOI 10.17487/RFC7474, April 2015, <https://www.rfc-editor.org/info/rfc7474>.

[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>.

[RFC8570] Ginsberg, L., Ed., Previdi, S., Ed., Giacalone, S., Ward, D., Drake, J., and Q. Wu, "IS-IS Traffic Engineering (TE) Metric Extensions", RFC 8570, DOI 10.17487/RFC8570, March 2019, <https://www.rfc-editor.org/info/rfc8570>.

[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>.

[RFC9330] Briscoe, B., Ed., De Schepper, K., Bagnulo, M., and G. White, "Low Latency, Low Loss, and Scalable Throughput (L4S) Internet Service: Architecture", RFC 9330, DOI 10.17487/RFC9330, January 2023, <https://www.rfc-editor.org/info/rfc9330>.

17.2. Informative References

[ATSplusCBS] "Latency and Backlog Bounds in Time-Sensitive Networking with Credit Based Shapers and Asynchronous Traffic Shaping", 2018, <https://ieeexplore.ieee.org/document/8493026>.

[ECQF] "Enhancements to Cyclic Queuing and Forwarding", 2023, <https://1.ieee802.org/tsn/802-1qdv/>.

[I-D.eckert-detnet-glbf] Eckert, T. T., Clemm, A., Bryant, S., and S. Hommes, "Deterministic Networking (DetNet) Data Plane - guaranteed Latency Based Forwarding (gLBF) for bounded latency with low jitter and asynchronous forwarding in Deterministic Networks", Work in Progress, Internet-Draft, draft-eckert-detnet-glbf-07, 11 December 2025, <https://datatracker.ietf.org/doc/html/draft-eckert-detnet-glbf-07>.

[I-D.joung-detnet-stateless-fair-queuing]
Joung, J., Ryoo, J., Cheung, T., Li, Y., and P. Liu, "Latency Guarantee with Stateless Fair Queuing", Work in Progress, Internet-Draft, draft-joung-detnet-stateless-fair-queuing-06, 22 December 2025, <https://datatracker.ietf.org/doc/html/draft-joung-detnet-stateless-fair-queuing-06>.

[I-D.peng-detnet-deadline-based-forwarding]
Peng, S., Du, Z., Basu, K., cheng, C., Yang, D., and C. Liu, "Deadline Based Deterministic Forwarding", Work in Progress, Internet-Draft, draft-peng-detnet-deadline-based-forwarding-18, 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., Peng, G., and J. Zhao, "Timeslot Queueing and Forwarding Mechanism", Work in Progress, Internet-Draft, draft-peng-detnet-packet-timeslot-mechanism-13, 12 October 2025, <https://datatracker.ietf.org/doc/html/draft-peng-detnet-packet-timeslot-mechanism-13>.

[I-D.xp-ippm-detnet-stamp]
Min, X., Peng, S., and X. hexiaoming, "STAMP Extensions for DetNet", Work in Progress, Internet-Draft, draft-xp-ippm-detnet-stamp-03, 11 December 2025, <https://datatracker.ietf.org/doc/html/draft-xp-ippm-detnet-stamp-03>.

[IEEE802.1Qav]
"IEEE Standard for Local and metropolitan area networks -- Virtual Bridged Local Area Networks - Amendment 12: Forwarding and Queuing Enhancements for Time-Sensitive Streams", 2010, <https://ieeexplore.ieee.org/document/8684664>.

[IEEE802.1Qch]
"IEEE Standard for Local and metropolitan area networks -- Bridges and Bridged Networks - Amendment 29: Cyclic Queuing and Forwarding", 2017, <https://ieeexplore.ieee.org/document/7961303>.

[IEEE802.1Qcr]
"IEEE Standard for Local and Metropolitan Area Networks--Bridges and Bridged Networks Amendment 34:Asynchronous Traffic Shaping", 2020, <https://ieeexplore.ieee.org/document/9253013>.

[IEEE802.1Qdv]
"Draft Standard for Local and metropolitan area networks--Enhancements to Cyclic Queuing and Forwarding", 2023, <https://1.ieee802.org/tsn/802-1qdv/>.

[Net-Calculus]
"Network Calculus: A Theory of Deterministic Queuing Systems for the Internet", 2001, <https://leboudec.github.io/netcal/latex/netCalBook.pdf>.

Author’s Address

Shaofu Peng
ZTE
China
Email: peng.shaofu@zte.com.cn