DetNet Working Group
Internet-Draft
Intended status: Informational
Expires: 12 July 2026

J. Joung
Sangmyung University
X. Geng
Huawei
S. Peng
ZTE Corporation
T. Eckert
Futurewei Technologies
8 January 2026

Dataplane Enhancement Taxonomy
draft-ietf-detnet-dataplane-taxonomy-05

Abstract

This draft is to facilitate the understanding of the data plane enhancement solutions, which are suggested currently or can be suggested in the future, for deterministic networking. This draft provides criteria for classifying data plane solutions. Examples of each category are listed, along with reasons where necessary. Strengths and limitations of the categories are described. Suitability of the solutions for various services of deterministic networking are also mentioned. Reference topologies for evaluation of the solutions are given as well.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

This draft is to facilitate the understanding of the data plane enhancement solutions, which are suggested currently or can be suggested in the future, for deterministic networking.

An enhancement solution can be a combination of multiple data plane functional entities, such as regulators, queues, and schedulers. A solution can also include functional entities across network nodes, e.g. traffic enforcement or regulation functions at the edge. A regulator, or equivalently a shaper, is defined as a functional entity that makes the arrival process of a flow conform to a predefined process. A packet scheduler, or simply a scheduler, is a functional entity that determines when a packet is transmitted.

The term taxonomy refers to a systematic method for classifying the sets of various solutions, how they are related, using certain criteria. A criterion is a principle or standard by which a solution can be judged or decided to be put into a certain category. A category is a subset of solutions classified into a single group using a criterion. This draft provides several criteria for classifying data plane solutions. These criteria are orthogonal to each other.

Examples of the categories are listed, along with reasons where necessary. Strengths and limitations of the categories are described.

Suitability of the solutions for various services of deterministic networking are also mentioned. The services can be classified according to the flow characteristics and the performance requirements. For example, Requirements for Reliable Wireless Industrial Services [I-D.ietf-detnet-raw-industrial-req] characterizes the services by the latency bound, the burst size, the burst transmission period, the number of nodes, etc. This document adopts this characterization rule, and classifies the services into one of tight/loose latency, large/small burst, periodic/non-periodic, and large/small scale services. For example, the display information service defined in Section 4.4. of [I-D.ietf-detnet-raw-industrial-req] is a loose latency, large burst, non-periodic, and small scale service.

The taxonomy described in this draft can be applied for the solutions of other standardization bodies, such as IEEE 802.1 TSN TG.

In this draft, the candidate solutions currently being proposed in DetNet WG are simply listed without any descriptions. The details of the solutions are intentionally omitted. Interesting readers may refer to the corresponding drafts. When necessary, the solutions from IEEE TSN TG or existing popular ones are used as examples to better understand the taxonomy and the derived categories.

The solutions raised in the DetNet WG are not entirely new concepts but rather variations of existing solutions. These deliberate approaches aim to address the scalability requirements defined in [I-D.ietf-detnet-scaling-requirements] while ensuring a degree of continuity and compatibility with the current practices. The taxonomy in this draft reflects how new solutions extend existing ones to address scalability challenges.

For instance, Cycle Specified Queuing and Forwarding (CSQF) [I-D.chen-detnet-sr-based-bounded-latency], Tagged Cyclic Queuing and Forwarding (TCQF) [I-D.eckert-detnet-tcqf], IEEE 802.1Qdv Enhanced CQF (ECQF) are enhancements built upon the foundation of Cyclic Queuing and Forwarding (CQF). Similarly, Work Conserving Stateless Core Fair Queuing (C-SCORE) [I-D.joung-detnet-stateless-fair-queuing] is an extension of Fair Queuing (FQ). Timeslot Queuing and Forwarding (TQF) [I-D.peng-detnet-packet-timeslot-mechanism] is an extension of IEEE 802.1Qbv, also known as Time Aware Shaper (TAS). Earliest Deadline First (EDF) [I-D.peng-detnet-deadline-based-forwarding] proposed to DetNet WG is a variation of the well-known solution that has the same name. Other well-known solutions that could provide bounded latency are also covered, for example Deficit Round Robin (DRR) and Asynchronous Traffic Shaping (ATS) [IEEE_802.1Qcr].

The suitable categories for achieving the goal of deterministic networking are defined. A suitable category is defined as a category in which solutions can be effective and efficient at realizing the goals of deterministic networking. The criteria defined in Section 5 are used for the decision of the suitable categories.

Reference topologies (RTs) are also listed in this document. An RT consists of a network topology and flows’ characteristics. The RTs listed in this document cover various topologies such as ring, mesh, hybrid etc. The purpose of listing the reference topologies (RTs) is to evaluate the dataplane solutions how they perform in real networks, in terms of E2E latency bound and jitter bound.

2. Terminology

2.1. Terms Used in This Document

The following terms are used in the context of deterministic networking in this document:

solution

A set of data plane functional entities, such as queue, scheduler, and regulator; which can guarantee a certain level of latency and jitter performance to a flow

regulator

A functional entity that makes the arrival process of a flow conform to a predefined process. Note that the term regulator is equivalent to shaper.

scheduler

A functional entity that determines when a packet is transmitted. Note that the term scheduler is equivalent to packet scheduler.

taxonomy

A systematic method by which a solution is put into a category

category

A set of solutions that is put together by one or more criteria

criterion

A principle or standard by which a solution can be judged or decided to be put into a certain category

level of category

The number of criteria that are used to determine the category

suitable category

A category within which solutions can be suitable for deterministic networking

time bound

The earliest or latest transmission completion time explicitly specified for a packet

2.2. Abbreviations

DetNet: Deterministic Networking

CQF: Cyclic Queuing and Forwarding

TAS: Time Aware Shaper

FQ: Fair Queuing

DRR: Deficit Round Robin

ATS: Asynchronous Traffic Shaping

RT: Reference Topology

E2E: End to End

GCL: Gate Control List

IR: Interleaved Regulator

FIFO: First In First Out

RSpec: Requested Specifications

TSpec: Traffic Specifications

IVN: In-Vehicle Network

CC: Command and Control

3. Conventions Used in This Document

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.

4. Taxonomy with Performance

Taxonomy based on the performance, such as E2E latency bounds and jitter bounds, is helpful to understand the solutions. The performance should be exhibited as a mathematical expression with the network and traffic parameters.

4.1. Per Hop Dominant Factor for Latency Bound

One possible criterion would be based on the per hop dominant factor for the latency bound. The dominant factor is defined as the largest sum term in the expression, when the network and traffic conditions are the worst. The worst condition typically means high network utilization, large packet and burst sizes, and large number of hops. Any existing solution can be put into one of three categories.

Category 1 (Max Packet Length/Service Rate): FQ and its variations like C-SCORE fall into this category, where the latency bound is primarily influenced by the ratio of a flow’s maximum packet size to its allocated service rate. This category emphasizes individual flow isolation. The consequence is that the variation of E2E latency bound for a flow is minimized with the other flows’ join and leave. Therefore, this category performs well with dynamic flows. This category also fits well to services with large bursts, since the burst sizes of flows are not the dominant factor of the latency bound.

Category 2 (Sum of Max Packet Lengths/Capacity): Solutions like DRR belong here, where the dominant factor is the sum of maximum packet lengths of all DetNet flows over the total allocated bandwidth. This category typically has less implementation complexity than Category 1 but can impact individual flow isolation. The other flows’ max packet lengths affect the latency bound, which can be altered as flows join and leave.

Category 3 (Sum of Max Burst Sizes/Capacity): CQF, TAS, their variations (including CSQF, TCQF, ECQF, TQF), and EDF fall into this category. The key influence on latency here is the total burst sizes of all DetNet flows relative to the network capacity. This category prioritizes bounded latency guarantees but may require tighter burst control mechanisms. Once the burst is controlled, for example by an extremely strict regulation, into a packet length level, then this category may be indistinguishable with Category 2. This category fits well to the services for static flows with small bursts.

As an example, assuming the capacities and maximum packet lengths are identical in all the links along the path of a flow under observation, the E2E latency bound of the flow by FQ is given as the following [STILIADIS-LRS].

$$(B-L)/r + H(Lh/Rh + L/r), \qquad (1)$$

where B, L, and r are the maximum burst size of, the maximum packet length of, and the allocated service rate to the flow, respectively; H is the number of hops; Lh and Rh are the maximum packet length and the capacity of all the links.

In this example, the term (Lh/Rh + L/r) can be seen as the per hop latency, because the max burst size, B, appears only once. The service rate of a flow, r, is likely to be much less than the link capacity, Rh, while the maximum lengths L and Lh would not differ too much. Therefore, the dominant factor here is L/r.

The dominant factor determines the level of flow isolation, as well as the level of E2E latency bound value.

5. Taxonomy with Functional Characteristics

Taxonomy based on the functional characteristics is the key to understanding the solutions. The criteria listed in this section are orthogonal to each other, if not stated explicitly.

5.1. Periodicity

If a solution has a set of consecutive time slots that is repeated periodically, and the time slots are assigned to packets based on a predefined rule, then the solution is periodic. Otherwise, the solution is non-periodic.

The set of consecutive time slots are called a period. Note that here we use the term period to avoid confusion with the term cycle used in CQF, which is equivalent to the time slot defined in this draft.

According to the above definition, IEEE 802.1Qbv TAS is a periodic solution. A finite Gate Control List (GCL) of TAS contains multiple gate control entries. Each entry represents a time slot with an assigned set of flows. A set of consecutive time slots forming a GCL is repeated periodically. Time slots can be overlapped with each other, as in ECQF.

TAS based solutions and CQF based solutions belong to periodic solutions, for example CSQF, TCQF, ECQF, TQF and so on.

Note that a periodic solution requires network synchronization. A further classification based on the level of required synchronization would also be possible. A solution could be classified as either phase synchronous, frequency synchronous, or asynchronous. However, this document does not specify this classification based on the synchronization, since only a frequency synchronization is possible in large scale networks.

Periodic solutions may fit well to periodic services, and vice versa.

5.2. Traffic Granularity

This document categorizes data plane solutions based on the granularity of their traffic control target, which refers to the size and specificity of the traffic entity they handle. Three granularity levels exist.

Flow level: Each packet is controlled based on its specific flow, which can be identified usually by the 5-tuple. Examples include FQ and its variations such as C-SCORE, which offer precise service differentiation but require potentially complex implementation.

Flow aggregate level: Flows are grouped by shared characteristics like traffic specification, service requirement, or routing path. This coarser level simplifies control but may offer less precise differentiation. Examples include interleaved regulators in ATS.

Class level: Flows are further grouped by similar service requirements, regardless of specific path or traffic details. This coarsest level simplifies control and accommodates traffic fluctuations but provides the least individual flow differentiation. Typically, time or time based information could be used for classification, such as in EDF, CQF and its variations.

For each level solution, packets within the same traffic entity receive the same treatment. For example, if a solution is flow aggregate level, then the packets within the same flow aggregate are treated identically, regardless of the flows they belong to.

There are cases in which a single solution consists of multiple functional entities that treat packets according to multiple traffic entities of different granularities. In such cases, it is defined that the functional entity with the coarsest granularity is dominant, thus the whole solution belongs to the coarsest granularity category.

For example, ATS consists of interleaved regulators (IRs) and a strict priority scheduler. An IR has a queue dedicated to a flow aggregate having the same class and the same input port. The regulation function itself is based on a flow. According to the definition above, IR is a flow aggregate level solution. On the other hand, the strict priority scheduler in ATS is class-based. Therefore, ATS as a whole is class level.

A finer granularity level solution has a benefit of a more accurate service differentiation among flows. Its limit is the larger implementation complexity. It fits to services with flows having various independent latency bound values.

Periodic solutions can further be categorized based on the traffic granularity. A time slot can be assigned per flow, per flow aggregate, or per class.

Note that TAS in 802.1Qbv is a scheduling mechanism defined in an output port with eight queues. The queues are controlled by GCL and its gate control entries. Each queue can serve a class. In an entry, queues can be either open or closed. Thus, TAS can be seen as a class level solution. However, in many cases TAS is understood as a scheduling mechanism, where the number of queues are not limited to 8. There could be a natural extension, such as TQF, which enables Qbv to allocate one queue to each flow or a flow aggregate.

Finer granularity periodic solutions have more strengths in jitter control. They also fit services with many periodic flows of independent period values.

5.3. Time Bounds

A solution can schedule a packet such that its transmission is completed within a specified interval of time. That interval can be either bounded, left-bounded, right-bounded, or unbounded. A left-bounded interval has a minimum time bound only. A right-bounded interval has a maximum time bound only. A bounded interval has both minimum and maximum time bounds. An unbounded interval does not have any finite bound.

Similarly, a solution can be categorized by this interval of allowed transmission completion time.

Bounded: A packet’s transmission completion is allowed after a minimum time bound and before a maximum time bound.

Left-bounded: A packet’s transmission completion is allowed after a minimum time bound.

Right-bounded: A packet’s transmission completion is allowed before a maximum time bound.

Unbounded: A packet’s transmission completion is allowed at any moment.

Note that unbounded solutions can also guarantee an E2E and a nodal latency bounds. They just do not have an explicitly specified nodal time bound.

FIFO, round robin schedulers, FQ and its variations like C-SCORE are examples of the unbounded solutions. TAS and CQF and their variations are bounded solutions, for example CSQF, TCQF, ECQF, and TQF. ATS and their variations with traffic shapers are left-bounded solutions. EDF can be operated either as an unbounded or a right-bounded solution, depending on whether the deadline is strictly kept or not.

Unbounded solutions have strengths in terms of average delay. They usually show smaller observed maximum latencies than the theoretical latency bound expressions suggest. They also benefit from the statistical multiplexing gain without any wasted capacity, thus more room for best effort traffic.

Bounded solutions have strengths to avoid burst accumulation and are also beneficial for jitter control. The burst size in a network of a flow can be kept similar or the same with the initial burst size. Therefore, the buffer size necessary typically is less than those in unbounded solutions.

5.4. Service Order

The packet service order within a single flow must be maintained. Data plane solutions decide the packet service order from different flows using various decision rules, categorized as follows.

Rate-based: Packets are ordered based on the allocated service rate of their flows or flow aggregates. Examples include FQ and its variations like C-SCORE, and DRR.

Arrival-based: Packets are ordered based on their arrival to a node. FIFO is an example.

Priority-based: Packets are ordered based on an assigned priority, other than the service rate or the arrival.

A rule for the service order may also be constructed with a combination of these characteristics.

According to its primary service order decision rule, a solution can be categorized into either rate-based, time-based, arrival-based, or priority-based. Any solution can also use the packet arrival time as a secondary decision rule.

The strict priority scheduler uses primarily the priority of a packet according to its class. It also uses the arrival times among packets of the same priority. In this case it is categorized as priority-based.

ATS has IRs and a strict priority scheduler. The service order among packets at an IR is arrival-based. The order among packets from different input ports are decided at the strict priority scheduler. Thus, ATS is priority-based.

Rate-based solutions have a simple admission condition check process that is dependent only on the service rates of flows. They benefit from the "pay burst only once" property, by which the maximum burst size of a flow contributes to the E2E latency bound only once, without being multiplied by the hop count. Rate-based solutions typically fit well to services with large burst and large scale services, without a need for overprovisioning, or additional burst control mechanisms.

Priority-based and arrival-based solutions benefit from the implementation simplicity. The latency and jitter differentiation among flows can be coarse, however. The services with loose latency, small burst, and non-periodic services may fit this category.

6. Suitable Categories for DetNet

This section specifies the suitable categories of solutions for deterministic networking. A category is a set of solutions that are put together by one or more criteria. A suitable category is a category in which solutions can be effective and efficient at realizing the goals of deterministic networking.

The criteria defined in Section 5 are used for the decision of the suitable categories. The seven suitable categories are defined, with the logic stated in the following.

6.1. Right-bounded category

There are four criteria in the functional taxonomy. They are the time bound, the service order, the periodicity, and the traffic granularity.

One of the most significant criteria for defining the characteristics of a solution is the time bound. There are four categories under this criterion: unbounded, left-bounded, right-bounded, and bounded.

For solutions in the right-bounded category, a packet has only a maximum time bound. In this category, the maximum bounds of packets directly decide the scheduling, or equivalently the service order of packets. No other criteria is important. No more categorization is necessary for this category. If a proper mechanism guarantees that a maximum time bound can be kept for a packet, then a solution in this category can guarantee an E2E latency bound. The category of right-bounded solutions is a suitable category.

6.2. Flow level periodic bounded category

Similarly, the bounded solutions can also guarantee E2E latency upper and lower bounds. Thus, jitter is naturally reduced. However, the existence of mechanisms to keep the time bounds for packets should be looked into.

The bounded solutions can be better described by the criterion of periodicity. The service order within an interval between time bounds is not important.

Based on the criterion of periodicity, the bounded solutions category can be further divided into the periodic bounded and the non-periodic bounded.

Periodic bounded solutions define a set of time slots, which periodically is repeated. Flows or flow aggregates can be scheduled into these slots. The slot scheduling should be executed over all the nodes in a path, and requires collaboration among nodes. Usually this process can be performed at the central control entity.

Periodic bounded solutions can be further categorized by the traffic granularity. Either flow level or class level subcategories is suitable. Note that in Section 6, the flow aggregate level category is merged into the flow level category, for simplicity. The flow level periodic bounded category is a suitable category.

6.3. Class level periodic bounded category

As stated in Section 6.2, the class level periodic bounded category is also a suitable category.

6.4. Flow level non-periodic bounded category

Non-perodic bounded solutions, as was the case with the right-bounded solutions, require a mechanism to guarantee the minimum and maximum bounds to be kept for a packet. Provided that this mechanism is present, the non-periodic bounded category is suitable. Either flow level or class level subcategories of this category is suitable. The flow level non-periodic bounded category is a suitable category.

6.5. Class level non-periodic bounded category

As stated in Section 6.4, the class level non-periodic bounded category is also a suitable category.

6.6. Flow level rate based unbounded category

Solutions without maximum time bounds cannot be described by the periodicity. They include the solutions of the unbounded and the left-bounded categories.

These two categories can be further categorized by the service order, into rate based, priority based, and arrival based. However, priority based (e.g. the strict priority scheduler) and arrival based solutions (e.g. FIFO scheduler) cannot meet tight latency bounds in large scale networks. Rate based solutions only are suitable. Therefore, the level two categories, rate based unbounded and rate based left-bounded are suitable for DetNet.

The above level two categories can be further divided with the criterion of traffic granularity. However, the class level rate based solutions are not suitable because the interference due to the burst intermix within a class is troublesome to flow with small bursts. Therefore, the flow level rate based unbounded category is a suitable category.

6.7. Flow level rate based left-bounded category

As with the same reasoning stated in Section 6.6, the flow level rate based left-bounded category is also suitable.

7. Reference Topologies

The purpose of listing the reference topologies (RTs) is to evaluate the dataplane solutions how they perform in real networks, in terms of the E2E latency bound and jitter bound. It is required to exactly calculate the E2E latency bound and jitter bound to any flow, given a dataplane solution and its parameter choices in implementation practices.

Additionally, the statistical performance evaluation results such as the average E2E latency, or the E2E latency distribution is recommended to be given. The scalability in both the data plane and the control plane are also recommended to be demonstrated. The implementation complexity of the dataplane solution, the complexity of the admission control procedure, and the slot scheduling procedure, in an environment with dynamic flows' join and leave, are the recommended performance metrics to be demonstrated.

An RT consists of a network topology and flows' characteristics. A network topology in this document specifies the abstract locations of source, destination, relay nodes and their interconnections. A flow characteristic is composed of its path, requested specifications (RSpec), and traffic specifications (TSpec). The requested specification includes the E2E latency and jitter bounds. The traffic specification includes the maximum burst size and the average rate, as if they have been shaped by a token bucket. Alternatively, a traffic can be specified by the period, the phase, and the maximum burst size. In this case, the maximum burst is transmitted at a certain fixed phase within a period of time.

By specifying the above information, other parameters such as the diameter and the maximum utilization of a network can be derived.

The RTs listed in this document cover various topologies such as ring, mesh, hybrid etc.

Some aspects of the RTs are derived from use cases, in order to reflect the current or future network deployment examples.

Based on the RTs, it is also able to check whether a dataplane solution can solve the scalability issues, e.g. those specified in [I-D.ietf-detnet-scaling-requirements]. The network diameter and the utilization level in RTs are set to examine the scalability.

The major interest of the deterministic networking is in the worst case delay, or equivalently the latency bound. Note that the reference topologies have specified the number of flows and their traffic specifications. There can be two different latency bounds with either the current scenario specified in the reference topologies, or the possible future scenario when more flows enter, thus the network becomes fully utilized.

A dataplane solution can either prepare the worst scenario from the beginning, or adjust as the flows come and go dynamically. If the solution is something flexible and tries to adjust, then the both latency bounds can be specified.

7.1. Grid

7.1.1. Network topology

A reference network topology, the grid, is shown in Figure 1. It represents a general network of partial mesh or grid topology, without considering a specific use case. A partial mesh is a common topology that can be seen in many real deployments, including datacenter networks.

```
               Src 1      Src 2      Src 3
                 |          |          |
              +------+   +------+   +------+
        Dst 1-|Node 1|<--|Node 2|-->|Node 3|-Dst 4
              +------+   +------+   +------+
                 |          ^          |
                 V          |          V
              +------+   +------+   +------+
        Dst 2-|Node 4|-->|Node 5|<--|Node 6|-Dst 5
              +------+   +------+   +------+
                 ^          |          ^
                 |          V          |
              +------+   +------+   +------+
        Dst 3-|Node 7|<--|Node 8|-->|Node 9|-Dst 6
              +------+   +------+   +------+
                 |          |          |
               Src 4      Src 5      Src 6
```

Figure 1: Grid topology

In Figure 1, arrowed links indicate the directions to follow for any traffic route. For example, from Node 2, only Node 1 and Node 3 are the next possible route.

The capacity of all the links in the topology is 1Gbps. While real deployments easily exceed 1Gbps link capacity, this RT represents a rather scaled down example in terms of the link capacity and the number of nodes.

7.1.2. Flow characteristics

In-vehicle network (IVN) is an example network which demands deterministic networking. [Buffered_Network] summarizes the flows that require deterministic networking services in IVNs as in Table 1.

Flow type	Maximum burst size	Maximum Packet length	Arrival rate	Required maximum latency
Audio	2Kbit	2Kbit	1.6Mbps	5ms
Video	360Kbit	12Kbit	11Mbps	10ms
Command and Control	2.4Kbit	2.4Kbit	480Kbps	5ms

Table 1: Flow types in in-vehicle networks

To simplify performance analysis, the flows in Table 1 are abstracted as shown in Table 2. Flows of the same type are aggregated into a single flow. For example, ten command and control (CC) flows that share the same E2E path can be considered to be a single type C flow as in Table 2. The maximum burst size and the average rate of a type C flow are about ten times those of one CC flow, in this case. Each flow type has specific destination nodes. For example, type A flows are destined only to destination 1 or 6.

Flow type	Maximum burst size	Maximum Packet length	Arrival rate	Required maximum latency	Destination in Figure 1
A	20Kbit	2Kbit	20Mbps	5ms	Dst 1, Dst 6
B	4000Kbit	10Kbit	100Mbps	10ms	Dst 3, Dst 4
C	20Kbit	2Kbit	5Mbps	5ms	Dst 2, Dst 5

Table 2: Flow type used in the reference topology

A source creates one flow to each destination for a total of 6 flows. 36 flows are created throughout the network. Table 2 describes characteristics of the three different flow types used in the RT.

7.1.3. Flow paths

The links are unidirectional as specified in Figure 1. All the flows must follow the direction of the arrows in every link. For example, a flow from source 1 to destination 5 follows the path of Src1-1-4-5-2-3-6-Dst9.

If there are more than one possible route to the destination, then the shortest path is selected. If there are more than one shortest path, then the following rules are applied.

Note that there are at most two outgoing links from a node to select. If both choices give the same distance to the destination, the node closer to the destination is selected as the next node. For example, from Src 5 to Dst 4, the selection from node 8 is to node 9, not to node 7, because node 9 is closer to Dst 4. When the above rule does not break the tie, i.e. the possible next nodes are within the same distance to the destination, then the node closer to the source is selected as the next node. For example, from Src 4 to Dst 5, the selection from node 5 is to node 8, not to node 2, because node 8 is closer to Src 4.

The above rules generate a unique route for every source and destination pair.

The reason for introducing unidirectional links is to make the network diameter large. With this configuration, the network diameter is 7 hops, which is relatively large considering a small number of nodes.

The destination of a flow decides the flow type. For example, all the flows destined to node 1 are of type A. There are 6 flows for each destination. There are 12 flows for each type. The flows with longest paths within the same flow type are of interest. Table 3 shows the path of the flows with longest paths for each flow type. For all the flow types, the number of hops in the longest paths is the same. The utilizations may differ for different links.

Flow type	Longest path
A	Src5-8-7-4-5-2-1-Dst1
A	Src2-2-3-6-5-8-9-Dst6
B	Src5-8-9-6-5-2-3-Dst4
B	Src2-2-1-4-5-8-7-Dst3
C	Src3-3-6-5-2-1-4-Dst2
C	Src6-9-6-5-8-7-4-Dst2
C	Src1-1-4-5-2-3-6-Dst5
C	Src4-7-4-5-8-9-6-Dst5

Table 3: Longest path of each flow type in the reference topology

7.1.4. Utilization

Network utilization is defined as the maximum link utilization over all the links. The RT achieves network utilization around 60%. The bottleneck links, e.g. the link 2-3, have one type A flow, five type B flows, and two type C flows. The scalability of a solution can be properly evaluated with this topology.

7.2. Hierarchical Ring-Mesh

7.2.1. Network topology

Another RT, the hierarchical ring-mesh, is illustrated over Figure 2 and Figure 3. The core network of the RT is depicted in Figure 2. A backbone node in the core network can be connected to one or two leaf network groups. A leaf network group consists of multiple leaf networks. The number of leaf networks in a group is a design parameter, but is recommended to be from 2 to 10. A leaf network of the RT is depicted in Figure 3.

This RT represents a wide area network, e.g. a state wide backbone network, having multiple subsidiary regional networks. The core network is a partial mesh with unidirectional links. Some of the nodes in the core network have the links to the leaf networks. A leaf network is a unidirectional ring with eight nodes. One of the nodes in the leaf network is linked to the core network.

```
                  Leaf NW    Leaf NW    Leaf NW
                  group 0    group 1    group 2
                     |          |          |
         Leaf NW  +------+   +------+   +------+ Leaf NW
         group 11-|Node 1|<--|Node 2|-->|Node 3|-group 3
                  +------+   +------+   +------+
                     |          ^          |
                     V          |          V
         Leaf NW  +------+   +------+   +------+ Leaf NW
         group 10-|Node 4|-->|Node 5|<--|Node 6|-group 4
                  +------+   +------+   +------+
                     ^          |          ^
                     |          V          |
         Leaf NW  +------+   +------+   +------+ Leaf NW
         group 9 -|Node 7|<--|Node 8|-->|Node 9|-group 5
                  +------+   +------+   +------+
                     |          |          |
                  Leaf NW    Leaf NW    Leaf NW
                  group 8    group 7    group 6
```

Figure 2: Topology of the core network in the hierarchical ring-mesh

The link capacity within the core network is another flexible design parameter. It can be set such that the maximum link utilization is 50%~80%. The capacity of connecting link between a leaf network and the core network is 1Gbps.

In Figure 2, arrowed links indicate the directions to follow for any traffic route. For example, from the leaf network group 0 to the group 6, the nodes 1, 4, 5, 8, 9 have to be visited.

```
                         To/From
                         Core NW
                            |
              +------+   +------+   +------+
     Src/Dest-|Node 7|-->|Node 0|-->|Node 1|-Src/Dest
              +------+   +------+   +------+
                 ^                      |
                 |                      V
              +------+              +------+
     Src/Dest-|Node 6|              |Node 2|-Src/Dest
              +------+              +------+
                 ^                      |
                 |                      V
              +------+   +------+   +------+
     Src/Dest-|Node 5|<--|Node 4|<--|Node 3|-Src/Dest
              +------+   +------+   +------+
                            |
                         Src/Dest
```

Figure 3: Topology of a leaf network in the hierarchical ring-mesh

The capacity of all the links in the leaf network is 1Gbps.

In Figure 3, arrowed links indicate the directions to follow for any traffic route. For example, for a flow to travel from the node 1 to the core network, every node in the leaf network should be visited.

7.2.2. Flow characteristics

In this RT, the flows specified in Table 1 are used again. There are three types of flows, which are the audio, video, and command & control. Let’s call these flows A, V, and C.

According to [Buffered_Network], the proportion of these three types in in-vehicle networks are such that A:V:C = 7:7:32. This RT lets seven A, seven V, thirty two C flows share the same path from source to destination. Let’s define these flows collectively as a flow set. A flow set’s total arrival rate is 103.56 Mbps.

7.2.3. Flow paths

Every flow in a leaf network travels from node i to node (i+7)mod8. Every flow travels 7 hops before leaving the leaf network. Exactly a single flow set enters into and leaves from the network by a node in the leaf network. The node 0, which is the connecting node to the core network, also has one flow set coming in from and another flow set going out to the core network.

The links in a leaf network are unidirectional as indicated by the arrows in Figure 3. All the flows follow the direction of the arrows in every link.

All the links in Figure 3 are passed by seven flow sets. This is the maximum number of flow sets in a link, over all the links. Any flow from a leaf network travels seven hops before leaving. This is the maximum number of hops in the leaf network.

In the core network, each leaf network group i sends n flow sets to the leaf network group (i+6)mod12, where n is the number of leaf networks in a leaf network group. These n flow sets are distributed to the destination leaf networks within the group evenly.

The links in the core network are unidirectional as indicated by the arrows in Figure 2. All the flows follow the direction of the arrows in every link.

At a crossroad, where a flow can choose one of two possible paths, the flow always turns left. For example, a flow from the leaf network group 1 travels to nodes 2, 3, 6, 5, 8, and then the leaf network group 7.

The link between the nodes 4 and 5 in Figure 2 are passed by (n times six) flow sets. This is the maximum number of flow sets in a link, over all the links. After joining the core network, a flow from the leaf network group 0 travels four hops before leaving. This is the maximum number of hops in the core network.

The maximum hop count in the leaf and core networks are 7 and 4, respectively. However, there are one hop each to/from the core network. The maximum hop count in this RT is 20.

7.2.4. Utilization

Network utilization is defined as the maximum link utilization over all the links. A leaf network has the network utilization of (seven times 103.56 Mbps) over 1 Gbps, which is around 72.5%. The utilization of the core network can be of any value, because the link capacity in the core network can be chosen to be any value greater than the (n times six times 103.56 Mbps), where n is the number of the leaf networks in a leaf network group. For example, if n equals to 2, then the link capacity is required to be larger than 1242.72 Mbps. If the link capacity is 1.5 Gbps, then the utilization is around 82.85%.

8. Considerations for interoperability between solutions

8.1. Background

In an E2E path of a flow, there can be multiple networks, or equivalently technology domains, in which different solutions are deployed. It is desirable to serve the flow with the promised service level without significant intervention in the middle. In this regard, it is recommended that the network consider the level of interoperability of the solution when selecting a particular solution.

A more interoperable solution is defined by the degree to which disparate technology domains can operate without specialized glue code, proprietary gateways, or manual intervention. A solution is considered more interoperable if it is closer to a "plug-and-play" without harming the promised deterministic service level.

However, the interoperability of a solution is not just a parameter of its own. In this draft, the interoperability between solutions in different categories are considered.

8.2. Interoperability between categories

The measure of the interoperability between solutions can be quantified by the following factors:

1) How difficult and complex it is to combine two different technology domains

2) How easily the requested E2E service level can be met throughout the E2E path in technology domains

These measures will be called the gateway complexity (GC) and the performance preservation level (PPL), respectively.

GC involves the complexities in both control plane and data plane. E2E admission control and cross-domain network configuration should be handled in the control plane. Metadata handling such as translation/insertion/deletion and flow reshaping can be the necessary interventions in the data plane. Note that a flow reshaping, or damping, may require additional per-flow state maintenance.

For example, an interworking gateway between a domain of flow level periodic bounded solution and another of class level non-periodic bounded solution can have low GC, since both domains can guarantee low jitter, and the arrival time of packets can be predictable. In this case, a traditional binding such as SID + SR policy can be installed on the interworking gateway, which is necessary even due to the need for E2E path concatenation.

The E2E latency or jitter bound, which is the critical performance metric in DetNet, can be just the simple sum of the bounds provided by the domains. However, in some solutions, the cooperation between domains can further improve the bounds. For example, rate-based solutions over domains can benefit from the pay-burst-only-once (PBOO) property, thus providing a better latency bound than the sum. These solutions over domains can also benefit by a unified rate-based admission control, thus having a low GC.

Table 4 is just an example illustration of interoperability measures between solutions that belong to different suitable categories. It may not reflect the actual ones. The evaluations are for further study.

	Class level non-periodic bounded	Flow level rate based left-bounded
Flow level periodic bounded	GC:Low / PPL: High	
Flow level rate based unbounded		GC:Low / PPL: High

Table 4: Interoperability Levels

8.3. Network for interoperability tests

The recommended network for interoperability test is composed of three cascaded technology domains. Each domain can use the grid topology defined in Section 7.1. If one solution is placed in the middle, and the other domains at the sides should use a different solution.

In this test network, the flows specified in Table 1 can be used. There are three types of flows, which are the audio, video, and command & control.

There are currently seven suitable categories. Any combination of these seven categories can be evaluated. The solutions in the same categories are assumed to be highly interoperable.

9. IANA Considerations

There might be matters that require IANA considerations associated with metadata. If necessary, relevant text will be added in a later version.

10. Security Considerations

This section will be described later.

11. Acknowledgements

12. Contributor

13. References

13.1. 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>.

[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>.

13.2. Informative References

[Buffered_Network] Joung, J. and J. Kwon, "Zero jitter for deterministic networks without time-synchronization", IEEE Access, vol. 9, pp. 49398-49414, doi:10.1109/ACCESS.2021.3068515, 2021.

[I-D.chen-detnet-sr-based-bounded-latency] Chen, M., Geng, X., Li, Z., Joung, J., and J. Ryoo, "Segment Routing (SR) Based Bounded Latency", Work in Progress, Internet-Draft, draft-chen-detnet-sr-based-bounded-latency-03, 7 July 2023, <https://datatracker.ietf.org/doc/html/draft-chen-detnet-sr-based-bounded-latency-03>.

[I-D.eckert-detnet-tcqf]
Eckert, T. T., Li, Y., Bryant, S., Malis, A. G., Ryoo, J., Liu, P., Li, G., and S. Ren, "Deterministic Networking (DetNet) Data Plane - Tagged Cyclic Queuing and Forwarding (TCQF) for bounded latency with low jitter in large scale DetNets", Work in Progress, Internet-Draft, draft-eckert-detnet-tcqf-12, 15 December 2025, <https://datatracker.ietf.org/doc/html/draft-eckert-detnet-tcqf-12>.

[I-D.ietf-detnet-raw-industrial-req]
Sofia, R. C., Mendes, P., Bernardos, C. J., and E. Schooler, "Requirements for Reliable Wireless Industrial Services", Work in Progress, Internet-Draft, draft-ietf-detnet-raw-industrial-req-01, 8 July 2024, <https://datatracker.ietf.org/doc/html/draft-ietf-detnet-raw-industrial-req-01>.

[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>.

[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>.

[IEEE_802.1Qcr]
IEEE, "IEEE Standard for Local and metropolitan area networks - Bridges and Bridged Networks - Amendment 34: Asynchronous Traffic Shaping", IEEE 802.1Qcr-2020, DOI 10.1109/IEEESTD.2020.9253013, 6 November 2020, <https://doi.org/10.1109/IEEESTD.2020.9253013>.

[STILIADIS-LRS]
Stiliadis, D. and A. Anujan, "Latency-rate servers: A general model for analysis of traffic scheduling algorithms", IEEE/ACM Trans. Networking, vol. 6, no. 5, pp. 611-624, 1998.

Authors’ Addresses

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

Xuesong Geng
Huawei
Email: gengxuesong@huawei.com

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

Toerless Eckert
Futurewei Technologies
Email: tte@cs.fau.de