Network Working Group   A. Malis
Internet-Draft   Independent
Intended status: Informational   X. Geng
Expires: January 14, 2021   M. Chen
                                                               Huawei
                                                              F. Qin
                                                         China Mobile
                                                             B. Varga
                                                             Ericsson
                                                        July 13, 2020

          Deterministic Networking (DetNet) Controller Plane Framework
               draft-malis-detnet-controller-plane-framework-04

Abstract

   This document provides a framework overview for the Deterministic
   Networking (DetNet) controller plane. It discusses concepts and
   requirements that will be basis for Detnet controller plane solution
   documents.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

   Deterministic Networking (DetNet) provides the capability to carry
   specified unicast and/or multicast data flows for real-time
   applications with extremely low data loss rates and bounded latency
   within a network domain. As discussed in the Deterministic

Networking Architecture [RFC8655], techniques used to provide this
capability include reserving data plane resources for individual (or
aggregated) DetNet flows in some or all of the intermediate nodes
along the path of the flow, providing explicit routes for DetNet
flows that do not immediately change with the network topology, and
distributing data from DetNet flow packets over time and/or space to
ensure delivery of each packet's data in spite of the loss of a path.

The DetNet data plane is defined in a set of documents that are
anchored by the DetNet Data Plane Framework
[I-D.ietf-detnet-data-plane-framework] and the associated DetNet MPLS
[I-D.ietf-detnet-mpls] and DetNet IP [I-D.ietf-detnet-ip] data plane
specifications, with additional details and subnet mappings provided
in [I-D.ietf-detnet-ip-over-mpls],
[I-D.ietf-detnet-mpls-over-udp-ip], [I-D.ietf-detnet-mpls-over-tsn],
[I-D.ietf-detnet-ip-over-tsn], and interconnection of TSN networks
[I-D.ietf-detnet-tsn-vpn-over-mpls].

While the Detnet Architecture and Data Plane Framework documents are
primarily concerned with data plane operations, they do contain some
references and requirements for functions that would be required in
order to automate DetNet service provisioning and monitoring via a
DetNet controller plane. The purpose of this document is to gather
these references and requirements into a single document and discuss
how various possible DetNet controller plane architectures could be
used to satisfy these requirements, while not providing the actual
protocol details for a DetNet controller plane solution. Such
controller plane protocol solutions will be the subject of subsequent
documents.

Note that in the DetNet overall architecture, the controller plane
includes what are more traditionally considered separate control and
management planes. Traditionally, the management plane is primarily
involved with node and network provisioning, operational OAM for
performance monitoring, and troubleshooting network behaviors and
outages, while the control plane is primarily responsible for the
instantiation and maintenance of flows, MPLS label allocation and
distribution, and active in-band or out-of-band signaling to support
these functions. In the DetNet architecture, all of this
functionality is combined into a single Controller Plane. See
Section 4.4.2 of [RFC8655] and the aggregation of Control and
Management planes in [RFC7426] for further details.

1.1. Terminology

This document uses the terminology established in the DetNet
Architecture [RFC8655], and the reader is assumed to be familiar with
that document and its terminology.

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.

2. DetNet Controller Plane Requirements

Other DetNet documents, including [RFC8655] and [I-D.ietf-detnet-data-plane-framework], contain requirements for the Controller Plane. For convenience, these requirements have been compiled here. These requirements have been organized to show those primarily related to the control plane, those primarily relate to the management plane, and those applicable to both planes.

2.1. DetNet Control Plane Requirements

The primary requirements of the DetNet Control Plane are that it must be able to:

o Support the dynamic creation, modification, and deletion of DetNet flows. This may include some or all of explicit path determination, link bandwidth reservations, restricting flows to specific links (e.g., IEEE 802.1 Time-Sensitive Networking (TSN) links), node buffer and other resource reservations, specification of required queuing disciplines along the path, ability to manage bidirectional flows, etc., as needed for a flow.

o Support DetNet flow aggregation and de-aggregation via the ability to dynamically create and delete flow aggregates (FAs), and be able to modify existing FAs by adding or deleting participating flows.

o Allow flow instantiation requests to originate in an end application (via an Application Programming Interface (API), via static provisioning, or via a dynamic control plane, such as a centralized SDN controller or distributed signaling protocols. See Section 3 for further discussion of these options.

o In the case of the DetNet MPLS data plane, manage DetNet Service Label (S-Label), Forwarding Label (F-Label), and Aggregation Label (A-Label) [I-D.ietf-detnet-mpls] allocation and distribution.

o Also in the case of the DetNet MPLS data plane, support the DetNet service sub-layer, which provides DetNet service functions such as protection and reordering through the use of packet replication, duplicate elimination, and packet ordering functions (PREOF).

   o Support queue control techniques defined in Section 4.5 of
      [RFC8655] and [I-D.finn-detnet-bounded-latency] that require time
      synchronization among network nodes.

   o Advertise static and dynamic node and link resources such as
      capabilities and adjacencies to other network nodes (for dynamic
      signaling approaches) or to network controllers (for centralized
      approaches).

   o Scale to handle the number of DetNet flows expected in a domain
      (which may require per-flow signaling or provisioning).

   o Provision flow identification information at each of the nodes
      along the path. Flow identification may differ depending on the
      location in the network and the DetNet functionality (e.g. transit
      node vs. relay node).

2.2. DetNet Management Plane Requirements

   The primary requirements of the DetNet Management Plane are that it
   must be able to:

   o Monitor the performance of DetNet flows and nodes to ensure that
      they are meeting required objectives, both proactively and on-
      demand.

   o Support DetNet flow continuity check and connectivity verification
      functions.

   o Support testing and monitoring of packet replication, duplicate
      elimination, and packet ordering functionality in the DetNet
      domain.

2.3. Requirements For Both Planes

   The following requirements apply to both the DetNet Controller and
   Management Planes:

   o Operate in a converged network domain that contains both DetNet
      and non-DetNet flows.

   o Adapt to DetNet domain topology changes such as links or nodes
      failures (fault recovery/restoration), additions and removals.

3. DetNet Control Plane Architecture

   As noted in the Introduction, the DetNet control plane is responsible
   for the instantiation and maintenance of flows, allocation and
   distribution of flow related information (e.g., MPLS label), and
   active in-band or out-of-band information distribution to support
   these functions.

   The following sections define three possible classes of DetNet
   control plane architectures: a fully distributed control plane
   utilizing dynamic signaling protocols, a fully centralized SDN-like
   control plane, and a control plane combining these two. They discuss
   the various information exchanges between entities in the network in
   each of these architectures and the advantages and disadvantages of
   each option.

   In each of the following sections, examples are used to illustrate
   possible mechanisms that could be used in each of the architectures.
   These are not meant to be exhaustive or to preclude any other
   possible mechanism that could be used in place of those used in the
   examples.

3.1. Distributed Control Plane and Signaling Protocols

   In a fully distributed configuration model, User-to-Network Interface
   (UNI) information is transmitted over a (to-be-defined) DetNet UNI
   protocol from the user side to the network side, and then UNI and
   network configuration information propagate in the network via
   distributed control plane signaling protocols. Such a DetNet UNI
   protocol are not visible in case of DetNet capable End-systems.

   Taking an RSVP-TE traffic-engineered MPLS network, where End systems
   are not part of the DetNet domain, as a theoretical example:

   1. An IGP collects topology information and DetNet capabilities of
       the network nodes

   2. The control plane of the ingress edge node receives a flow
       establishment request from the UNI and calculates one or more
       valid path(s);

   3. Using RSVP-TE [RFC3209], the ingress edge node sends a PATH
       message with an explicit route. After receiving the PATH
       message, the egress edge node sends a RESV message with the
       distributed label and resource reservation request.

   IGP in the above example would require extensions to incorporate
   DetNet capabilities. Similarly, current reservation-oriented

distributed control plane protocols, e.g. RSVP-TE, can only reserve
bandwidth along the path, while the configuration of a fine-grained
schedule, e.g., Enhancements for Scheduled Traffic
[IEEE.802.1QBV_2015], is not supported. If RSVP-TE were to be used
for serving a DetNet flow, it would require extensions in order to
support queue and scheduler reservations in addition to bandwidth
reservation.

As discussed in Section 4.9 of [RFC8655], scalability is a primary
concern for DetNet, given the large number of expected flows in a
DetNet domain. This could potentially be much larger than, for
example, the number of full-mesh MPLS traffic tunnels in a network
using MPLS traffic engineering, which would typically be N*(N-1)
tunnels, where N is the number of edge routers in the domain.

Even when flow aggregation is used, DetNet domains can be expected to
support a very large number of flows that will need particular
queuing disciplines and/or resource allocation, depending on the
requirements for each flow. This could require a large amount of
dynamic signaling, such as an RSVP-TE session to establish and
maintain each flow. Other RSVP-TE scalability concerns are further
discussed in [RFC5439].

All of the above tends to argue against a purely distributed control
plane for DetNet domains.

3.2. SDN/Fully Centralized Control Plane

In the fully SDN/centralized configuration model, flow/UNI
information is transmitted from a Centralized User Configuration or
from applications via an API or northbound interface to a Centralized
Controller, which is the sole source of routing and forwarding
information for the domain. Configurations of nodes for DetNet flows
are performed by the controller using a protocol such as NETCONF
[RFC6241]/YANG [RFC6020] or PCE-CC [RFC8283].

Taking again an MPLS network, where End systems are not part of the
DetNet domain, as a theoretical example:

1. A Centralized Controller collects topology information and DetNet
    capabilities of the network nodes via NETCONF/YANG;

2. The Controller receives a flow establishment request from a UNI
    and calculates one or more valid path(s) through the network;

3. The Controller chooses the optimal path and configures the
    devices along that path for flow transmission via PCE-CC.

Protocols in the above example may require extensions to incorporate
DetNet specific parameters.

3.3. Combined Control Plane (partly centralized, partly distributed)

In the combined model, a Controller and control plane protocols work
together to provide DetNet services, and there are a number of
possible combinations.

Using an RSVP-TE traffic-engineered MPLS network with centralized PCE
(Path Computation Engine), where End systems are not part of the
DetNet domain, as a theoretical example:

1. A Centralized Controller collects topology information and DetNet
    capabilities of the network nodes via an IGP and/or BGP-LS
    [RFC7752];

2. The Controller receives a flow establishment request from a
    Network Management System and calculates one or more valid
    path(s) through the network;

3. Based on the calculation result, the Controller distributes flow
    path information to the ingress edge node and other information
    (e.g. replication/duplicate elimination) to the relevant nodes.

4. Using RSVP-TE, the ingress edge node sends a PATH message with an
    explicit route. After receiving the PATH message, the egress
    edge node sends a RESV message with the distributed label and
    resource reservation request.

Similarly to Distributed Control Plane and SDN/Fully Centralized
Control Plane scenarios extensions of protocols are required to
incorporate DetNet specific parameters.

There are many other variations that could be included in a combined
control plane. This document cannot discuss all the possible control
plane mechanisms that could be used in combined configuration models.
Every solution has its own mechanisms and corresponding parameters
that are required for it to work.

4. DetNet Control Plane Additional Details and Issues

This section discusses some additional DetNet control plane details
and issues.

4.1. Explicit Paths

   Explicit paths are required in DetNet to provide a stable forwarding
   service and guarantee that DetNet service is not impacted when the
   network topology changes. The following features are necessary to
   have explicit paths in DetNet:

   o Path computation: DetNet explicit paths need to meet the SLA
      (Service Level Agreement) requirements and/or resource guarantees
      from the application/client, which include bandwidth, maximum end-
      to-end delay, maximum end-to-end delay variation, maximum loss
      ratio, etc. In a distributed system with IGP-TE, CSPF
      (Constrained Shortest Path First) can be used to compute a set of
      feasible paths for a DetNet service. In a system with a network
      controller, a PCE (Path Computation Engine) can compute paths
      satisfying the requirements of DetNet based on the network
      information collected from the DetNet domain.

   o Path establishment: Once the path has been computed, the options
      discussed in Section 3 can be used to establish the path. Also
      see Section 4.4.1 for some additional considerations depending on
      the details of the network infrastructure.

   o Strict or loose paths: An explicit path is strict when every
      intermediate hop is specified so that its route can't change. An
      explicit path is loose when any IGP route is allowed along the
      path. Generally, end-to-end SLA guarantees require a strict
      explicit path in DetNet. However, when the IGP route is known to
      be able to meet the SLA requirements, loose explicit paths are
      also acceptable.

4.2. Resource Reservation

   Network congestion could cause uncontrolled delay and/or packet loss.
   DetNet flows are supposed to be protected from congestion, so
   sufficient resource reservation for DetNet service is necessary.
   Resources in the network are complex and hard to quantize, and may
   include such entities as packet processing resources, packet
   buffering, port and link bandwidth, and so on. The resources a
   particular flow requires are determined by the flow's characteristics
   and SLA.

   o Resource Allocation: Port bandwidth is one of the basic attributes
      of a network device which is easy to obtain or calculate. In
      current traffic engineering implementations, network resource
      allocation is synonymous with bandwidth allocation. A DetNet flow
      is characterized with a traffic specification as defined in
      [I-D.ietf-detnet-flow-information-model], including attributes

such as Interval, Maximum Packets Per Interval, and Maximum
Payload Size. The traffic specification describes the worst case,
rather than the average case, for the traffic, to ensure that
sufficient bandwidth and buffering resources are reserved to
satisfy the traffic specification. However, in case of DetNet,
resource allocation is more than simple bandwidth reservation.
For example, allocation of buffers and required queuing
disciplines during forwarding may be required as well.
Furthermore, resources must be ensured to execute DetNet service
sub-layer functions on the node, such as protection and reordering
through the use of packet replication, duplicate elimination, and
packet ordering functions (PREOF).

o Device configuration with or without flow discrimination: The
   resource allocation can be guaranteed by device configuration.
   For example, an output port bandwidth reservation can be
   configured as a parameter of queue management and the port
   scheduling algorithm. When DetNet flows are aggregated, a group
   of DetNet flows share the allocated resource in the network
   device. When the DetNet flows are treated independently, the
   device should maintains a mapping relationship between a DetNet
   flow and its corresponding resources.

4.3. PREOF Support

DetNet path redundancy is supported via packet replication, duplicate
elimination, and packet ordering functions (PREOF). A DetNet flow is
replicated and goes through multiple networks paths to avoid packet
loss caused by device or link failures. In general, current control
plane mechanisms that can be used to establish an explicit path,
whether distributed or centralized, support point-to-point (P2P) and
point-to-multipoint (P2MP) path establishment. PREOF requires the
ability to compute and establish a set of multiple paths (e.g.,
multiple LSP segments in an MPLS network) from the point(s) of packet
replication to the point(s) of packet merging and ordering. Mapping
of DetNet (member) flows to explicit path segments has to be ensured
as well. Protocol extensions will be required to support these new
features. Terminology will also be required to refer to this
coordinated set of path segments (such as an "LSP graph" in case of
DetNet MPLS data plane).

4.4. Data Plane specific considerations

4.4.1. DetNet in an MPLS Domain

For the purposes of this document, "traditional MPLS" is defined as
MPLS without the use of segment routing (see Section 4.4.3 for a
discussion of MPLS with segment routing) or MPLS-TP [RFC5960].

In traditional MPLS domains, a dynamic control plane using
distributed signaling protocols is typically used for the
distribution of MPLS labels used for forwarding MPLS packets. The
dynamic signaling protocols most commonly used for label distribution
are LDP [RFC5036], RSVP-TE, and BGP [RFC8277] (which enables BGP/
MPLS-based Layer 3 VPNs [RFC4384] and Layer 2 VPNs [RFC7432]).

Any of these protocols could be used to distribute DetNet Service
Labels (S-Labels) and Aggregation Labels (A-Labels)
[I-D.ietf-detnet-mpls]. As discussed in
[I-D.ietf-detnet-data-plane-framework], S-Labels are similar to other
MPLS service labels, such as pseudowire, L3 VPN, and L2 VPN labels,
and could be distributed in a similar manner, such as through the use
of targeted LDP or BGP. If these were to be used for DetNet, they
would require extensions to support DetNet-specific features such as
PREOF, aggregation (A-Labels), node resource allocation, and queue
placement.

However, as discussed in Section 3.1, distributed signaling protocols
may have difficulty meeting DetNet's scalability requirements. MPLS
also allows SDN-like centralized label management and distribution as
an alternative to distributed signaling protocols, using protocols
such as PCEP and OpenFlow [OPENFLOW].

PCEP, particularly when used as a part of PCE-CC, is a possible
candidate protocol to use for centralized management of traditional
MPLS-based DetNet domains. However, PCE path calculation algorithms
would need to be extended to include the location determination for
PREOF nodes in a path, and the means to signal the necessary resource
reservation and PREOF function placement information to network
nodes. See ((?I-D.ietf-pce-pcep-extension-for-pce-controller)) for
further discussion of PCE-CC and PCEP for centralized control of an
MPLS domain.

4.4.2. DetNet in an IP Domain

For the purposes of this document, "traditional IP" is defined as IP
without the use of segment routing (see Section 4.4.3 for a
discussion of IP with segment routing). In a later revision of this
document, this section will discuss possible protocol extensions to
existing IP routing protocols such as OSPF, IS-IS, and BGP. It
should be noted that a DetNet IP data plane [I-D.ietf-detnet-ip] is
simpler than a DetNet MPLS data plane [I-D.ietf-detnet-mpls], and
doesn't support PREOF, so only one path per flow or flow aggregate is
required.

4.4.3. DetNet in a Segment Routing Domain

   Segment Routing [RFC8402] is a scalable approach to building network
   domains that provides explicit routing via source routing encoded in
   packet headers and it is combined with centralized network control to
   compute paths through the network. Forwarding paths are distributed
   with associated policy to network edge nodes for use in packet
   headers. As such, segment routing can be considered as a new data
   plane for both MPLS and IP. It reduces the amount of network
   signaling associated with distributed signaling protocols such as
   RSVP-TE, and also reduces the amount of state in core nodes compared
   with that required for traditional MPLS and IP routing, as the state
   is now in the packets rather than in the routers. This could be
   useful for DetNet, where a very large number of flows through a
   network domain are expected, which would otherwise require the
   instantiation of state for each flow traversing each node in the
   network. However, further analysis is needed on the expected gain,
   as DetNet flows may require various type of DetNet specific resources
   as well.

   In a later revision of this document, this section will discuss the
   impact of DetNet on the Segment Routing Control and Management
   planes. Note that the DetNet MPLS and IP data planes described in
   [I-D.ietf-detnet-mpls] and [I-D.ietf-detnet-ip] were constructed to
   be compatible with both types of segment routing, SR-MPLS [RFC8660]
   and SRv6 [I-D.ietf-6man-segment-routing-header]. However, as of this
   writing, traffic engineering and resource reservation for segment
   routing are currently unsolved problems.

   Editor's note: this section may be collapsed to previous sections and
   listing MPLS segment routing in the MPLS section as one of the
   possible explicit routing techniques for MPLS, and do the same for
   IP.

5. Management Plane Overview

   The Management Plane includes the ability to statically provision
   network nodes and to use OAM to monitor DetNet performance and detect
   outages or other issues at the DetNet layer.

5.1. Provisioning

   Static provisioning in a Detnet network nodes will be performed via
   the use of appropriate YANG models, including [I-D.ietf-detnet-yang]
   and [I-D.ietf-detnet-topology-yang].

5.2. DetNet Operations, Administration and Maintenance (OAM)

   The overall framework and requirements for DetNet OAM are discussed
   in [I-D.mirsky-detnet-oam]. This document currently includes
   additional OAM details that may eventually be merged into that
   document.

5.2.1. OAM for Performance Monitoring (PM)

5.2.1.1. Active PM

   Active PM is performed by injecting OAM packets into the network to
   estimate the performance of the network by measuring the performance
   of the OAM packets. Adding extra traffic can affect the delay and
   throughput performance of the network, and for this reason active PM
   is not recommended for use in operational DetNet domains. However,
   it is a useful test tool when commissioning a new network or during
   troubleshooting.

5.2.1.2. Passive PM

   Passive PM monitors the actual service traffic in a network domain in
   order to measure its performance without having a detrimental affect
   on the network. As compared to Active PM, Passive PM is much
   preferred for use in DetNet domains.

5.2.2. OAM for Connectivity and Fault/Defect Management (CFM)

   [I-D.mirsky-detnet-oam] contains requirements for connectivity and
   fault/defect detection and management in a DetNet domain.

6. Gap Analysis

   In a later revision of this document, this section will contain a gap
   analysis of existing IETF control and management plane protocols not
   already discussed elsewhere in this document for their ability (or
   inability) to satisfy the requirements in Section 2, and discuss
   possible protocol extensions to existing protocols to fill the gaps,
   if any.

7. IANA Considerations

   This document has no actions for IANA.

   Note to RFC Editor: this section may be removed on publication as an
   RFC.

8. Security Considerations

   Editor's note: This section needs more details.

   The overall security considerations of DetNet are discussed in
   [RFC8655] and [I-D.ietf-detnet-security]. For DetNet networks that
   make use of Segment Routing (whether SR-MPLS or SRv6), the security
   considerations in [RFC8402] also apply.

   DetNet networks that make use of a centralized controller plane may
   be threatened by the loss of connectivity (whether accidental or
   malicious) between the central controller and the network nodes, and/
   or the spoofing of control messages from the controller to the
   network nodes. This is important since such networks depend on
   centralized controllers to calculate flow paths and instantiate flow
   state in the network nodes. For networks that use both DetNet and
   Segment Routing with a centralized controller, this would also
   include the calculation of SID lists and their installation in edge/
   border routers.

   In both cases, such threats may be mitigated through redundant
   controllers, the use of authentication between the controller(s) and
   the network nodes, and other mechanisms for protection against DOS
   attacks. A mechanism for supporting one or more alternative central
   controllers and the ability to fail over to such an alternative
   controller will be required.

9. Acknowledgments

   Thanks to Jim Guichard, Donald Eastlake, and Stewart Bryant for their
   review comments.

10. References

10.1. Normative References

   [I-D.ietf-detnet-data-plane-framework]
             Varga, B., Farkas, J., Berger, L., Malis, A., and S.
             Bryant, "DetNet Data Plane Framework", draft-ietf-detnet-
             data-plane-framework-06 (work in progress), May 2020.

   [I-D.ietf-detnet-flow-information-model]
             Varga, B., Farkas, J., Cummings, R., Jiang, Y., and D.
             Fedyk, "DetNet Flow Information Model", draft-ietf-detnet-
             flow-information-model-10 (work in progress), May 2020.

   [I-D.ietf-detnet-ip]
             Varga, B., Farkas, J., Berger, L., Fedyk, D., and S.
             Bryant, "DetNet Data Plane: IP", draft-ietf-detnet-ip-07
             (work in progress), July 2020.

   [I-D.ietf-detnet-mpls]
             Varga, B., Farkas, J., Berger, L., Malis, A., Bryant, S.,
             and J. Korhonen, "DetNet Data Plane: MPLS", draft-ietf-
             detnet-mpls-09 (work in progress), July 2020.

   [I-D.ietf-detnet-security]
             Mizrahi, T. and E. Grossman, "Deterministic Networking
             (DetNet) Security Considerations", draft-ietf-detnet-
             security-10 (work in progress), May 2020.

   [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>.

   [RFC7426] Haleplidis, E., Ed., Pentikousis, K., Ed., Denazis, S.,
             Hadi Salim, J., Meyer, D., and O. Koufopavlou, "Software-
             Defined Networking (SDN): Layers and Architecture
             Terminology", RFC 7426, DOI 10.17487/RFC7426, January
             2015, <https://www.rfc-editor.org/info/rfc7426>.

   [RFC8174] .

   [RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L.,
             Decraene, B., Litkowski, S., and R. Shakir, "Segment
             Routing Architecture", RFC 8402, DOI 10.17487/RFC8402,
             July 2018, <https://www.rfc-editor.org/info/rfc8402>.

   [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>.

10.2. Informative References

   [I-D.finn-detnet-bounded-latency]
             Finn, N., Boudec, J., Mohammadpour, E., Zhang, J., Varga,
             B., and J. Farkas, "DetNet Bounded Latency", draft-finn-
             detnet-bounded-latency-04 (work in progress), June 2019.

   [I-D.ietf-6man-segment-routing-header]
             Filsfils, C., Dukes, D., Previdi, S., Leddy, J.,
             Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header
             (SRH)", draft-ietf-6man-segment-routing-header-26 (work in
             progress), October 2019.

   [I-D.ietf-detnet-ip-over-mpls]
             Varga, B., Berger, L., Fedyk, D., Bryant, S., and J.
             Korhonen, "DetNet Data Plane: IP over MPLS", draft-ietf-
             detnet-ip-over-mpls-06 (work in progress), May 2020.

   [I-D.ietf-detnet-ip-over-tsn]
             Varga, B., Farkas, J., Malis, A., and S. Bryant, "DetNet
             Data Plane: IP over IEEE 802.1 Time Sensitive Networking
             (TSN)", draft-ietf-detnet-ip-over-tsn-03 (work in
             progress), June 2020.

   [I-D.ietf-detnet-mpls-over-tsn]
             Varga, B., Farkas, J., Malis, A., and S. Bryant, "DetNet
             Data Plane: MPLS over IEEE 802.1 Time Sensitive Networking
             (TSN)", draft-ietf-detnet-mpls-over-tsn-03 (work in
             progress), June 2020.

   [I-D.ietf-detnet-mpls-over-udp-ip]
             Varga, B., Farkas, J., Berger, L., Malis, A., and S.
             Bryant, "DetNet Data Plane: MPLS over UDP/IP", draft-ietf-
             detnet-mpls-over-udp-ip-06 (work in progress), May 2020.

   [I-D.ietf-detnet-topology-yang]
             Geng, X., Chen, M., Li, Z., and R. Rahman, "Deterministic
             Networking (DetNet) Topology YANG Model", draft-ietf-
             detnet-topology-yang-00 (work in progress), January 2019.

   [I-D.ietf-detnet-tsn-vpn-over-mpls]
             Varga, B., Farkas, J., Malis, A., Bryant, S., and D.
             Fedyk, "DetNet Data Plane: IEEE 802.1 Time Sensitive
             Networking over MPLS", draft-ietf-detnet-tsn-vpn-over-
             mpls-03 (work in progress), June 2020.

   [I-D.ietf-detnet-yang]
             Geng, X., Chen, M., Ryoo, Y., Li, Z., Rahman, R., and D.
             Fedyk, "Deterministic Networking (DetNet) Configuration
             YANG Model", draft-ietf-detnet-yang-06 (work in progress),
             June 2020.

   [I-D.mirsky-detnet-oam]
             Mirsky, G. and M. Chen, "Operations, Administration and
             Maintenance (OAM) for Deterministic Networks (DetNet)",
             draft-mirsky-detnet-oam-03 (work in progress), May 2019.

   [IEEE.802.1QBV_2015]
             IEEE, "IEEE Standard for Local and metropolitan area
             networks -- Bridges and Bridged Networks - Amendment 25:
             Enhancements for Scheduled Traffic", IEEE 802.1Qbv-2015,
             DOI 10.1109/IEEESTD.2016.7572858, March 2016,
             <http://ieeexplore.ieee.org/servlet/
             opac?punumber=7572858>.

   [OPENFLOW]
             Open Networking Foundation, "OpenFlow Switch
             Specification, Version 1.5.1 (Protocol version 0x06)",
             ONF TS-025, March 2015, <https://www.opennetworking.org/
             wp-content/uploads/2014/10/openflow-switch-v1.5.1.pdf>.

   [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V.,
             and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP
             Tunnels", RFC 3209, DOI 10.17487/RFC3209, December 2001,
             <https://www.rfc-editor.org/info/rfc3209>.

   [RFC4384] Meyer, D., "BGP Communities for Data Collection", BCP 114,
             RFC 4384, DOI 10.17487/RFC4384, February 2006,
             <https://www.rfc-editor.org/info/rfc4384>.

   [RFC5036] Andersson, L., Ed., Minei, I., Ed., and B. Thomas, Ed.,
             "LDP Specification", RFC 5036, DOI 10.17487/RFC5036,
             October 2007, <https://www.rfc-editor.org/info/rfc5036>.

   [RFC5439] Yasukawa, S., Farrel, A., and O. Komolafe, "An Analysis of
             Scaling Issues in MPLS-TE Core Networks", RFC 5439,
             DOI 10.17487/RFC5439, February 2009,
             <https://www.rfc-editor.org/info/rfc5439>.

   [RFC5960] Frost, D., Ed., Bryant, S., Ed., and M. Bocci, Ed., "MPLS
             Transport Profile Data Plane Architecture", RFC 5960,
             DOI 10.17487/RFC5960, August 2010,
             <https://www.rfc-editor.org/info/rfc5960>.

   [RFC6020] Bjorklund, M., Ed., "YANG - A Data Modeling Language for
             the Network Configuration Protocol (NETCONF)", RFC 6020,
             DOI 10.17487/RFC6020, October 2010,
             <https://www.rfc-editor.org/info/rfc6020>.

   [RFC6241] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed.,
              and A. Bierman, Ed., "Network Configuration Protocol
              (NETCONF)", RFC 6241, DOI 10.17487/RFC6241, June 2011,
              <https://www.rfc-editor.org/info/rfc6241>.

   [RFC7432] Sajassi, A., Ed., Aggarwal, R., Bitar, N., Isaac, A.,
              Uttaro, J., Drake, J., and W. Henderickx, "BGP MPLS-Based
              Ethernet VPN", RFC 7432, DOI 10.17487/RFC7432, February
              2015, <https://www.rfc-editor.org/info/rfc7432>.

   [RFC7752] Gredler, H., Ed., Medved, J., Previdi, S., Farrel, A., and
              S. Ray, "North-Bound Distribution of Link-State and
              Traffic Engineering (TE) Information Using BGP", RFC 7752,
              DOI 10.17487/RFC7752, March 2016,
              <https://www.rfc-editor.org/info/rfc7752>.

   [RFC8277] Rosen, E., "Using BGP to Bind MPLS Labels to Address
              Prefixes", RFC 8277, DOI 10.17487/RFC8277, October 2017,
              <https://www.rfc-editor.org/info/rfc8277>.

   [RFC8283] Farrel, A., Ed., Zhao, Q., Ed., Li, Z., and C. Zhou, "An
              Architecture for Use of PCE and the PCE Communication
              Protocol (PCEP) in a Network with Central Control",
              RFC 8283, DOI 10.17487/RFC8283, December 2017,
              <https://www.rfc-editor.org/info/rfc8283>.

   [RFC8660] Bashandy, A., Ed., Filsfils, C., Ed., Previdi, S.,
              Decraene, B., Litkowski, S., and R. Shakir, "Segment
              Routing with the MPLS Data Plane", RFC 8660,
              DOI 10.17487/RFC8660, December 2019,
              <https://www.rfc-editor.org/info/rfc8660>.

Authors' Addresses

   Andrew G. Malis
   Independent

   Email: agmalis@gmail.com

   Xuesong Geng
   Huawei

   Email: gengxuesong@huawei.com

Mach (Guoyi) Chen
Huawei

Email: mach.chen@huawei.com

Fengwei Qin
China Mobile

Email: qinfengwei@chinamobile.com

Balazs Varga
Ericsson

Email: balazs.a.varga@ericsson.com