RFC 9938
A Framework for the Deterministic Networking (DetNet) Controller Plane

Abstract

This document provides a framework overview for the Deterministic Networking (DetNet) Controller Plane. It discusses concepts and requirements for the DetNet Controller Plane, which could be the basis for a future solution specification.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

Deterministic Networking (DetNet) provides the ability to carry specified unicast or multicast data flows for real-time applications with extremely low packet loss rates and assured maximum end-to-end delivery latency. A description of the general background and concepts of DetNet can be found in [RFC8655].

The DetNet data plane is defined in the DetNet data plane framework [RFC8938] (and is further explained in the associated DetNet MPLS [RFC8964], the DetNet IP [RFC8939], and other data plane specifications [RFC9023] [RFC9024] [RFC9025] [RFC9037] [RFC9056]).

Note that in the DetNet overall architecture, the controller plane includes what are more usually considered separate control and management planes (see Section 4.4.2 of [RFC8655]). The management plane is primarily involved with fault management, configuration management, and performance management (sometimes accounting management and security management are also considered in the management plane (Section 4.2 of [RFC6632]) but they are out of the scope of this document). At the same time, 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 DetNet 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.

While the DetNet architecture and data plane documents are primarily concerned with data plane operations, they do contain some requirements and considerations for functions that would be required in order to automate DetNet service provisioning and monitoring via a DetNet Controller Plane (e.g., see Section 4 of [RFC8938]). The purpose of this document is to take these requirements and considerations into a single document and extend and discuss how various possible DetNet Controller Plane architectures could be used to satisfy these requirements, while not providing the protocol details for a DetNet Controller Plane solution. Such controller plane protocol solutions will be the subject of subsequent documents. Therefore, this document should be considered as the authoritative reference to be considered if/when protocol work on the DetNet Controller Plane starts.

2. DetNet Controller Plane Requirements

Other DetNet documents, including [RFC8655], [RFC8938], [RFC9551], and [RFC9055], among others, contain requirements for the controller plane. For convenience, these requirements have been compiled here. These requirements have been organized into three groups: 1) requirements primarily applicable to the control plane, 2) requirements primarily applicable to the management plane, and 3) requirements applicable to both planes. In addition, security requirements for the DetNet Controller Plane have been discussed in [RFC9055], and a summary of those requirements is provided in Section 2.3. For the sake of clarity, when applicable, the document in which the requirements originally appear is referenced.

2.1. DetNet Control Plane Requirements

The primary requirements for the DetNet Control Plane are as follows:

- Support the dynamic instantiation, modification, and deletion of DetNet flows. This may include some or all of explicit path determination, link bandwidth reservations, restriction of 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, the ability to manage bidirectional flows, etc., as needed for a flow [RFC8938].

- Support DetNet flow aggregation and de-aggregation via the ability to dynamically create and delete flow aggregates (FAs) and modify existing FAs by adding or deleting participating flows [RFC8938].

- 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 Software-Defined Networking (SDN) controller or distributed signaling protocols). See Section 3 for further discussion of these options.

- Manage, in the case of the DetNet MPLS data plane, DetNet Service Label (S-Label), Forwarding Label (F-Label), and Aggregation Label (A-Label) [RFC8964] allocation and distribution [RFC8938].

- Support, also in the case of the DetNet MPLS data plane, the DetNet service sub-layer, which provides DetNet service functions, such as protection and reordering through the use of Packet Replication, Elimination, and Ordering Functions (PREOF) [RFC8655].

- Support the queue control techniques defined in [RFC8655], Section 4.5 and [RFC9320] that require time synchronization among the network data plane nodes.

- Advertise static and dynamic node and link characteristics, such as capabilities and adjacencies to other network nodes (for dynamic signaling approaches) or to network controllers (for centralized approaches) [RFC8938].

- Scale to handle the number of DetNet flows expected in a domain (which may require per-flow signaling or provisioning) [RFC8938].

- 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) [RFC8938].

2.2. DetNet Management Plane Requirements

The primary requirements for the DetNet management plane are as follows:

- Monitor the performance of DetNet flows and nodes to ensure that they are meeting required objectives, both proactively and on demand [RFC9551].
- Support DetNet flow continuity check and connectivity verification functions [RFC9551].
- Support testing and monitoring of packet replication, duplicate elimination, and packet ordering functionality in the DetNet domain [RFC9551].

2.3. Requirements for Both Planes

The following requirements apply to both the DetNet control and management planes:

- Operate in a converged network domain that contains both DetNet and non-DetNet flows [RFC8655].
- Adapt to DetNet domain topology changes such as link or node failures (fault recovery/ restoration), additions, and removals [RFC8655].

In addition to the above, the DetNet Controller Plane should also satisfy security requirements derived from [RFC9055], which defines the security framework for DetNet. The following requirements are especially relevant:

- Integrity and authenticity of control/signaling packets: The controller plane should ensure that signaling and control messages cannot be modified or injected by unauthorized entities and should prevent spoofing and segmentation attacks.
- Protection against controller compromise: Mechanisms should exist to verify the legitimacy of controllers and to prevent unauthorized components from impersonating them.
- System-wide security design: The architecture must account for the possibility of compromised nodes or controllers, ensuring resilience so that the failure or subversion of a single component does not cause catastrophic impact.
- Timely delivery of control plane messages: The controller plane should ensure that control and signaling messages are delivered without undue delay to prevent disruption of DetNet services without resource leakage.

3. DetNet Control Plane Architecture

As noted in the Introduction, the DetNet Control Plane is responsible for the instantiation and maintenance of flows, the 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 types of DetNet Control Plane architectures: 1) a fully distributed control plane utilizing dynamic signaling protocols, 2) a fully centralized SDN-like control plane, and 3) a hybrid control plane containing both distributed protocols and

centralized controlling. This document describes the various information exchanges between entities in the network for each type of these architectures and the corresponding advantages and disadvantages.

The examples in the following sections illustrate possible mechanisms that could be used in each type of the architectures. They 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, the User-Network Interface (UNI) information is transmitted over a DetNet UNI protocol from the user side to the network side. Then, the UNI and network configuration information propagates in the network via distributed control plane signaling protocols. Such a DetNet UNI protocol is not necessary when the end systems are DetNet capable.

Taking an RSVP-TE [RFC3209] MPLS network as an example, where end systems are not part of the DetNet domain:

1. Network nodes collect topology information and DetNet capabilities of the network nodes through IGP.
2. The ingress edge node receives a flow establishment request from the UNI and calculates one or more valid paths.
3. The ingress node sends a PATH message with an explicit route through RSVP-TE. After receiving the PATH message, the egress edge node sends a RESV message with the distributed label and resource reservation request.

In this example, both the IGP and RSVP-TE may require extensions for DetNet.

3.2. Fully Centralized Control Plane

In the fully centralized configuration model (e.g., using an SDN controller), both flow and UNI information can be transmitted from a centralized user controller or from other applications, via an API or northbound interface, to a centralized controller. Network node configurations for DetNet flows are performed by the controller using a protocol such as NETCONF [RFC6241], YANG [RFC6020] [RFC7950], DetNet YANG [RFC9633], or a PCE-based central controller (PCE-CC) [RFC8283].

Take the following case as an 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 paths through the network.
3. The controller chooses the optimal path and configures the devices along that path for DetNet flow transmission via PCE-CC.

The protocols in the above example may require extensions for DetNet.

3.3. Hybrid Control Plane (Partly Centralized and Partly Distributed)

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

In the following case, the RSVP-TE and controller are used together:

1. A controller collects topology information and DetNet capabilities of the network nodes via an IGP and/or the Border Gateway Protocol - Link State (BGP-LS) [RFC9552].
2. A controller receives a flow establishment request through API and calculates one or more valid paths through the network.
3. Based on the calculation result, the controller distributes flow path information to the ingress edge node and configures network nodes along the path with necessary DetNet information (e.g., for replication/duplicate elimination).
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.

There are many other variations that could be included in a hybrid control plane. The requested DetNet extensions for a protocol in each possible case is for future work.

4. DetNet Control Plane for DetNet Mechanisms

This section discusses the requested control plane features for DetNet mechanisms as defined in [RFC8655], including PREOF. Different DetNet services may implement any or all of these based on the requirements.

4.1. Explicit Paths

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

- Path computation: DetNet explicit paths need to meet the Service Level Agreement (SLA) requirements of the application, which include bandwidth, maximum end-to-end delay, maximum end-to-end delay variation, maximum loss ratio, etc. In a distributed network system, an IGP with Constrained Shortest Path First (CSPF) may be used to compute a set of feasible paths for a DetNet service. In a centralized network system, the controller can compute paths satisfying the requirements of DetNet based on the network information collected from the DetNet domain.
- Path establishment: The computed path for the DetNet service has to be sent/configured/signaled to the network device so that the corresponding DetNet flow can pass through the network domain following the specified path.

4.2. Resource Reservation

DetNet flows are supposed to be protected from congestion, so sufficient resource reservation for a DetNet service could protect a service from congestion. There are multiple types of resources in the network that could be allocated to DetNet flows, e.g., packet processing resources, buffer resources, and the bandwidth of the output port. The network resource requested by a specified DetNet service is determined by the SLA requirements and network capability.

- Resource Allocation: Port bandwidth is one of the basic attributes of a network device that is easy to obtain or calculate. In current traffic engineering implementations, network resource allocation is synonymous with bandwidth allocation. A DetNet flow is characterized by a traffic specification, as defined in [RFC9016], including attributes such as Interval, MaxPacketsPerInterval, and MaxPayloadSize. 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 the 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 PREOF.
- 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 maintain a mapping relationship between a DetNet flow and its corresponding resources.

4.3. PREOF Support

DetNet path redundancy is supported via Packet Replication, Elimination, and Ordering Functions (PREOF). A DetNet flow is replicated and forwarded by 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 Label Switched Path (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 flows or 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 the case of the DetNet MPLS data plane).

4.4. Data-Plane-Specific Considerations

4.4.1. DetNet in an MPLS Domain

For the purposes of this document, "legacy 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 Transport Profile (MPLS-TP) [RFC5960].

In legacy 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 [RFC4875], and BGP [RFC8277] (which enables BGP-based MPLS Layer 3 VPNs [RFC4384], Layer 2 VPNs [RFC4664], and EVPNs [RFC7432]).

Any of these protocols could be used to distribute DetNet Service Labels (S-Labels) and Aggregation Labels (A-Labels) [RFC8964]. As discussed in [RFC8938], S-Labels are similar to other MPLS service labels, such as pseudowire and 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.

4.4.2. DetNet in an IP Domain

For the purposes of this document, "legacy IP" is defined as IP without the use of Segment Routing (see Section 4.4.3 for a discussion of IP with Segment Routing). It should be noted that a DetNet IP data plane [RFC8939] is simpler than a DetNet MPLS data plane [RFC8964] and doesn't support PREOF, so only one path per flow or flow aggregate is required. Therefore, possible protocol extensions are expected to be limited, e.g., to existing IP routing protocols.

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 policies to network edge nodes for use in packet headers. Segment Routing 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 legacy 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.

Note that the DetNet MPLS and IP data planes described in [RFC8964] and [RFC8939] were constructed to be compatible with both types of Segment Routing: Segment Routing over MPLS (SR-MPLS) [RFC8660] and Segment Routing over IPv6 (SRv6) [RFC8754] [RFC8986].

4.5. Encapsulation and Metadata Support

To effectively manage DetNet flows, the controller plane will need to have a clear understanding of the encapsulation and metadata capabilities of the underlying network nodes. This will require a control mechanism that can discover, configure, and manage these parameters for each flow.

The controller plane needs to understand and manage the encapsulation and metadata capabilities of the network nodes to provision DetNet flows effectively. This process might need a discovery phase in which the controller discovers which encapsulation types (e.g., MPLS, IP) and metadata schemes (e.g., sequencing, timestamping) that each node supports. After discovery, the controller might instruct nodes on the specific encapsulation and companion metadata to apply for a given flow. This ensures that DetNet packets are handled consistently across the network. For example, the controller might instruct a node to use an MPLS header and add a sequence number for a particular flow.

5. Management Plane Overview

The management plane includes the ability to statically provision network nodes and to use Operations, Administration, and Maintenance (OAM) to monitor DetNet performance and to detect outages or other issues at the DetNet layer.

5.1. DetNet Operations, Administration, and Maintenance (OAM)

This document covers the general considerations for OAM.

5.1.1. OAM for Performance Monitoring (PM)

5.1.1.1. Active PM

Active PM is performed by injecting OAM packets into the network to estimate the performance of the network and by then measuring the performance of those 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.

5.1.1.2. Passive PM

Passive PM, such as In Situ Operations, Administration, and Maintenance (IOAM) [RFC9197], monitors the actual service traffic in a network domain in order to measure its performance without having a detrimental effect on the network. As compared to Active PM, Passive PM is much preferred for use in DetNet domains.

5.1.2. OAM for Connectivity and Fault Management (CFM)

The detailed requirements for CFM in a DetNet IP domain and a DetNet MPLS domain are defined in [RFC9551], [RFC9634], and [RFC9546], respectively.

6. Multi-Domain Aspects

When there are multiple domains involved, multiple Controller Plane Functions (CPFs) would have to collaborate to implement the requests received from the Flow Management Entity (FME) [RFC8655] as per-flow, per-hop behaviors installed in the DetNet nodes for each individual flow. Adding multi-domain support might require some support at the CPF. For example, CPFs of different domains, e.g., PCEs, need to discover each other and then authenticate and negotiate per-hop behaviors. Furthermore, in the case of wireless domains, per-domain functions specific to Reliable and Available Wireless (RAW) [RAW-ARCH], such as Point of Local Repairs (PLRs), have to also be considered, e.g., in addition to the PCEs. Depending on the multi-domain support provided by the application plane, the controller plane might be relieved from some responsibilities (e.g., if the application plane takes care of splitting what needs to be provided by each domain).

7. IANA Considerations

This document has no IANA actions.

8. Security Considerations

This document provides a framework for the DetNet Controller Plane and does not include any protocol specifications. Any future specification that is defined to support the DetNet Controller Plane is expected to include the appropriate security considerations. For overall security considerations of DetNet, see [RFC8655] and [RFC9055].

9. References

9.1. Normative References

[RFC8655]   Finn, N., Thubert, P., Varga, B., and J. Farkas, "Deterministic Networking Architecture", RFC 8655, DOI 10.17487/RFC8655, October 2019, <https://www.rfc-editor.org/info/rfc8655>.

[RFC8938]   Varga, B., Ed., Farkas, J., Berger, L., Malis, A., and S. Bryant, "Deterministic Networking (DetNet) Data Plane Framework", RFC 8938, DOI 10.17487/RFC8938, November 2020, <https://www.rfc-editor.org/info/rfc8938>.

[RFC9016]   Varga, B., Farkas, J., Cummings, R., Jiang, Y., and D. Fedyk, "Flow and Service Information Model for Deterministic Networking (DetNet)", RFC 9016, DOI 10.17487/RFC9016, March 2021, <https://www.rfc-editor.org/info/rfc9016>.

[RFC9055]   Grossman, E., Ed., Mizrahi, T., and A. Hacker, "Deterministic Networking (DetNet) Security Considerations", RFC 9055, DOI 10.17487/RFC9055, June 2021, <https://www.rfc-editor.org/info/rfc9055>.

[RFC9551]   Mirsky, G., Theoleyre, F., Papadopoulos, G., Bernardos, CJ., Varga, B., and J. Farkas, "Framework of Operations, Administration, and Maintenance (OAM) for Deterministic Networking (DetNet)", RFC 9551, DOI 10.17487/RFC9551, March 2024, <https://www.rfc-editor.org/info/rfc9551>.

9.2. Informative References

[RAW-ARCH] Thubert, P., Ed., "Reliable and Available Wireless Architecture", Work in Progress, Internet-Draft, draft-ietf-raw-architecture-30, 25 July 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-raw-architecture-30>.

[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>.

[RFC4664]   Andersson, L., Ed. and E. Rosen, Ed., "Framework for Layer 2 Virtual Private Networks (L2VPNs)", RFC 4664, DOI 10.17487/RFC4664, September 2006, <https://www.rfc-editor.org/info/rfc4664>.

[RFC4875]   Aggarwal, R., Ed., Papadimitriou, D., Ed., and S. Yasukawa, Ed., "Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)", RFC 4875, DOI 10.17487/RFC4875, May 2007, <https://www.rfc-editor.org/info/rfc4875>.

[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>.

[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>.

[RFC6632]   Ersue, M., Ed. and B. Claise, "An Overview of the IETF Network Management Standards", RFC 6632, DOI 10.17487/RFC6632, June 2012, <https://www.rfc-editor.org/info/rfc6632>.

[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>.

[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>.

[RFC7950]   Bjorklund, M., Ed., "The YANG 1.1 Data Modeling Language", RFC 7950, DOI
            10.17487/RFC7950, August 2016, <https://www.rfc-editor.org/info/rfc7950>.

[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>.

[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>.

[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>.

[RFC8754]   Filsfils, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer,
            "IPv6 Segment Routing Header (SRH)", RFC 8754, DOI 10.17487/RFC8754, March
            2020, <https://www.rfc-editor.org/info/rfc8754>.

[RFC8939]   Varga, B., Ed., Farkas, J., Berger, L., Fedyk, D., and S. Bryant, "Deterministic
            Networking (DetNet) Data Plane: IP", RFC 8939, DOI 10.17487/RFC8939,
            November 2020, <https://www.rfc-editor.org/info/rfc8939>.

[RFC8964]   Varga, B., Ed., Farkas, J., Berger, L., Malis, A., Bryant, S., and J. Korhonen,
            "Deterministic Networking (DetNet) Data Plane: MPLS", RFC 8964, DOI 10.17487/
            RFC8964, January 2021, <https://www.rfc-editor.org/info/rfc8964>.

[RFC8986]   Filsfils, C., Ed., Camarillo, P., Ed., Leddy, J., Voyer, D., Matsushima, S., and Z. Li,
            "Segment Routing over IPv6 (SRv6) Network Programming", RFC 8986, DOI
            10.17487/RFC8986, February 2021, <https://www.rfc-editor.org/info/rfc8986>.

[RFC9023]   Varga, B., Ed., Farkas, J., Malis, A., and S. Bryant, "Deterministic Networking
            (DetNet) Data Plane: IP over IEEE 802.1 Time-Sensitive Networking (TSN)", RFC
            9023, DOI 10.17487/RFC9023, June 2021, <https://www.rfc-editor.org/info/
            rfc9023>.

[RFC9024]   Varga, B., Ed., Farkas, J., Malis, A., Bryant, S., and D. Fedyk, "Deterministic Networking (DetNet) Data Plane: IEEE 802.1 Time-Sensitive Networking over MPLS", RFC 9024, DOI 10.17487/RFC9024, June 2021, <https://www.rfc-editor.org/info/rfc9024>.

[RFC9025]   Varga, B., Ed., Farkas, J., Berger, L., Malis, A., and S. Bryant, "Deterministic Networking (DetNet) Data Plane: MPLS over UDP/IP", RFC 9025, DOI 10.17487/RFC9025, April 2021, <https://www.rfc-editor.org/info/rfc9025>.

[RFC9037]   Varga, B., Ed., Farkas, J., Malis, A., and S. Bryant, "Deterministic Networking (DetNet) Data Plane: MPLS over IEEE 802.1 Time-Sensitive Networking (TSN)", RFC 9037, DOI 10.17487/RFC9037, June 2021, <https://www.rfc-editor.org/info/rfc9037>.

[RFC9056]   Varga, B., Ed., Berger, L., Fedyk, D., Bryant, S., and J. Korhonen, "Deterministic Networking (DetNet) Data Plane: IP over MPLS", RFC 9056, DOI 10.17487/RFC9056, October 2021, <https://www.rfc-editor.org/info/rfc9056>.

[RFC9197]   Brockners, F., Ed., Bhandari, S., Ed., and T. Mizrahi, Ed., "Data Fields for In Situ Operations, Administration, and Maintenance (IOAM)", RFC 9197, DOI 10.17487/RFC9197, May 2022, <https://www.rfc-editor.org/info/rfc9197>.

[RFC9320]   Finn, N., Le Boudec, J.-Y., Mohammadpour, E., Zhang, J., and B. Varga, "Deterministic Networking (DetNet) Bounded Latency", RFC 9320, DOI 10.17487/RFC9320, November 2022, <https://www.rfc-editor.org/info/rfc9320>.

[RFC9546]   Mirsky, G., Chen, M., and B. Varga, "Operations, Administration, and Maintenance (OAM) for Deterministic Networking (DetNet) with the MPLS Data Plane", RFC 9546, DOI 10.17487/RFC9546, February 2024, <https://www.rfc-editor.org/info/rfc9546>.

[RFC9552]   Talaulikar, K., Ed., "Distribution of Link-State and Traffic Engineering Information Using BGP", RFC 9552, DOI 10.17487/RFC9552, December 2023, <https://www.rfc-editor.org/info/rfc9552>.

[RFC9633]   Geng, X., Ryoo, Y., Fedyk, D., Rahman, R., and Z. Li, "Deterministic Networking (DetNet) YANG Data Model", RFC 9633, DOI 10.17487/RFC9633, October 2024, <https://www.rfc-editor.org/info/rfc9633>.

[RFC9634]   Mirsky, G., Chen, M., and D. Black, "Operations, Administration, and Maintenance (OAM) for Deterministic Networking (DetNet) with the IP Data Plane", RFC 9634, DOI 10.17487/RFC9634, October 2024, <https://www.rfc-editor.org/info/rfc9634>.

Acknowledgments

Contributors

Fengwei Qin
China Mobile
Email: qinfengwei@chinamobile.com

Authors' Addresses

Andrew G. Malis
Independent
Email: agmalis@gmail.com

Xuesong Geng (EDITOR)
Huawei
Email: gengxuesong@huawei.com

Mach (Guoyi)Chen
Huawei
Email: mach.chen@huawei.com

Balazs Varga
Ericsson
Email: balazs.a.varga@ericsson.com

Carlos J. Bernardos
Universidad Carlos III de Madrid
Av. Universidad, 30
28911 Leganes, Madrid
Spain
Phone: +34 91624 6236
Email: cjbc@it.uc3m.es
URI: http://www.it.uc3m.es/cjbc/