

Internet Engineering Task Force (IETF)
Request for Comments: 9055
Category: Informational
ISSN: 2070-1721

E. Grossman, Ed.
DOLBY
T. Mizrahi
HUAWEI
A. Hacker
THOUGHT
June 2021

Deterministic Networking (DetNet) Security Considerations

Abstract

A DetNet (deterministic network) provides specific performance guarantees to its data flows, such as extremely low data loss rates and bounded latency (including bounded latency variation, i.e., "jitter"). As a result, securing a DetNet requires that in addition to the best practice security measures taken for any mission-critical network, additional security measures may be needed to secure the intended operation of these novel service properties.

This document addresses DetNet-specific security considerations from the perspectives of both the DetNet system-level designer and component designer. System considerations include a taxonomy of relevant threats and attacks, and associations of threats versus use cases and service properties. Component-level considerations include ingress filtering and packet arrival-time violation detection.

This document also addresses security considerations specific to the IP and MPLS data plane technologies, thereby complementing the Security Considerations sections of those documents.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are candidates for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9055>.

Copyright Notice

Copyright (c) 2021 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction
2. Abbreviations and Terminology
3. Security Considerations for DetNet Component Design
 - 3.1. Resource Allocation
 - 3.1.1. Inviolable Flows
 - 3.1.2. Design Trade-Off Considerations in the Use Cases Continuum
 - 3.1.3. Documenting the Security Properties of a Component
 - 3.1.4. Fail-Safe Component Behavior
 - 3.1.5. Flow Aggregation Example
 - 3.2. Explicit Routes
 - 3.3. Redundant Path Support
 - 3.4. Timing (or Other) Violation Reporting
4. DetNet Security Considerations Compared with Diffserv Security Considerations
5. Security Threats
 - 5.1. Threat Taxonomy
 - 5.2. Threat Analysis
 - 5.2.1. Delay
 - 5.2.2. DetNet Flow Modification or Spoofing
 - 5.2.3. Resource Segmentation (Inter-segment Attack) Vulnerability
 - 5.2.4. Packet Replication and Elimination
 - 5.2.4.1. Replication: Increased Attack Surface
 - 5.2.4.2. Replication-Related Header Manipulation
 - 5.2.5. Controller Plane
 - 5.2.5.1. Path Choice Manipulation
 - 5.2.5.2. Compromised Controller
 - 5.2.6. Reconnaissance
 - 5.2.7. Time-Synchronization Mechanisms
 - 5.3. Threat Summary
6. Security Threat Impacts
 - 6.1. Delay Attacks
 - 6.1.1. Data Plane Delay Attacks
 - 6.1.2. Controller Plane Delay Attacks
 - 6.2. Flow Modification and Spoofing
 - 6.2.1. Flow Modification
 - 6.2.2. Spoofing
 - 6.2.2.1. Data Plane Spoofing
 - 6.2.2.2. Controller Plane Spoofing
 - 6.3. Segmentation Attacks (Injection)
 - 6.3.1. Data Plane Segmentation
 - 6.3.2. Controller Plane Segmentation
 - 6.4. Replication and Elimination
 - 6.4.1. Increased Attack Surface
 - 6.4.2. Header Manipulation at Elimination Routers
 - 6.5. Control or Signaling Packet Modification
 - 6.6. Control or Signaling Packet Injection
 - 6.7. Reconnaissance
 - 6.8. Attacks on Time-Synchronization Mechanisms
 - 6.9. Attacks on Path Choice
7. Security Threat Mitigation
 - 7.1. Path Redundancy
 - 7.2. Integrity Protection
 - 7.3. DetNet Node Authentication
 - 7.4. Synthetic Traffic Insertion
 - 7.5. Encryption
 - 7.5.1. Encryption Considerations for DetNet
 - 7.6. Control and Signaling Message Protection
 - 7.7. Dynamic Performance Analytics
 - 7.8. Mitigation Summary
8. Association of Attacks to Use Cases
 - 8.1. Association of Attacks to Use Case Common Themes

8.1.1.	Sub-network Layer
8.1.2.	Central Administration
8.1.3.	Hot Swap
8.1.4.	Data Flow Information Models
8.1.5.	L2 and L3 Integration
8.1.6.	End-to-End Delivery
8.1.7.	Replacement for Proprietary Fieldbuses and Ethernet-Based Networks
8.1.8.	Deterministic vs. Best-Effort Traffic
8.1.9.	Deterministic Flows
8.1.10.	Unused Reserved Bandwidth
8.1.11.	Interoperability
8.1.12.	Cost Reductions
8.1.13.	Insufficiently Secure Components
8.1.14.	DetNet Network Size
8.1.15.	Multiple Hops
8.1.16.	Level of Service
8.1.17.	Bounded Latency
8.1.18.	Low Latency
8.1.19.	Bounded Jitter (Latency Variation)
8.1.20.	Symmetrical Path Delays
8.1.21.	Reliability and Availability
8.1.22.	Redundant Paths
8.1.23.	Security Measures
8.2.	Summary of Attack Types per Use Case Common Theme
9.	Security Considerations for OAM Traffic
10.	DetNet Technology-Specific Threats
10.1.	IP
10.2.	MPLS
11.	IANA Considerations
12.	Security Considerations
13.	Privacy Considerations
14.	References
14.1.	Normative References
14.2.	Informative References
Contributors	
Authors' Addresses	

1. Introduction

A deterministic IP network ("Deterministic Networking Architecture" [RFC8655]) can carry data flows for real-time applications with extremely low data loss rates and bounded latency. The bounds on latency defined by DetNet (as described in [RFC9016]) include both worst-case latency (Maximum Latency, Section 5.9.2 of [RFC9016]) and worst-case jitter (Maximum Latency Variation, Section 5.9.3 of [RFC9016]). Data flows with deterministic properties are well established for Ethernet networks (see Time-Sensitive Networking (TSN), [IEEE802.1BA]); DetNet brings these capabilities to the IP network.

Deterministic IP networks have been successfully deployed in real-time Operational Technology (OT) applications for some years; however, such networks are typically isolated from external access, and thus the security threat from external attackers is low. An example of such an isolated network is a network deployed within an aircraft, which is "air gapped" from the outside world. DetNet specifies a set of technologies that enable creation of deterministic flows on IP-based networks of a potentially wide area (on the scale of a corporate network), potentially merging OT traffic with best-effort Information Technology (IT) traffic, and placing OT network components into contact with IT network components, thereby exposing the OT traffic and components to security threats that were not present in an isolated OT network.

These DetNet (OT-type) technologies may not have previously been

deployed on a wide area IP-based network that also carries IT traffic, and thus they can present security considerations that may be new to IP-based wide area network designers; this document provides insight into such system-level security considerations. In addition, designers of DetNet components (such as routers) face new security-related challenges in providing DetNet services, for example, maintaining reliable isolation between traffic flows in an environment where IT traffic co-mingles with critical reserved-bandwidth OT traffic; this document also examines security implications internal to DetNet components.

Security is of particularly high importance in DetNet because many of the use cases that are enabled by DetNet [RFC8578] include control of physical devices (power grid devices, industrial controls, building controls, etc.) that can have high operational costs for failure and present potentially attractive targets for cyber attackers.

This situation is even more acute given that one of the goals of DetNet is to provide a "converged network", i.e., one that includes both IT traffic and OT traffic, thus exposing potentially sensitive OT devices to attack in ways that were not previously common (usually because they were under a separate control system or otherwise isolated from the IT network, for example [ARINC664P7]). Security considerations for OT networks are not a new area, and there are many OT networks today that are connected to wide area networks or the Internet; this document focuses on the issues that are specific to the DetNet technologies and use cases.

Given the above considerations, securing a DetNet starts with a scrupulously well-designed and well-managed engineered network following industry best practices for security at both the data plane and controller plane, as well as for any Operations, Administration, and Maintenance (OAM) implementation; this is the assumed starting point for the considerations discussed herein. Such assumptions also depend on the network components themselves upholding the security-related properties that are to be assumed by DetNet system-level designers; for example, the assumption that network traffic associated with a given flow can never affect traffic associated with a different flow is only true if the underlying components make it so. Such properties, which may represent new challenges to component designers, are also considered herein.

Starting with a "well-managed network", as noted above, enables us to exclude some of the more powerful adversary capabilities from the Internet Threat Model of [BCP72], such as the ability to arbitrarily drop or delay any or all traffic. Given this reduced attacker capability, we can present security considerations based on attacker capabilities that are more directly relevant to a DetNet.

In this context, we view the "conventional" (i.e., non-time-sensitive) network design and management aspects of network security as being primarily concerned with preventing denial of service, i.e., they must ensure that DetNet traffic goes where it's supposed to and that an external attacker can't inject traffic that disrupts the delivery timing assurance of the DetNet. The time-specific aspects of DetNet security presented here take up where those "conventional" design and management aspects leave off.

However, note that "conventional" methods for mitigating (among all the others) denial-of-service attacks (such as throttling) can only be effectively used in a DetNet when their use does not compromise the required time-sensitive or behavioral properties required for the OT flows on the network. For example, a "retry" protocol is typically not going to be compatible with a low-latency (worst-case maximum latency) requirement; however, if in a specific use case and implementation such a retry protocol is able to meet the timing

constraints, then it may well be used in that context. Similarly, if common security protocols such as TLS/DTLS or IPsec are to be used, it must be verified that their implementations are able to meet the timing and behavioral requirements of the time-sensitive network as implemented for the given use case. An example of "behavioral properties" might be that dropping of more than a specific number of packets in a row is not acceptable according to the service level agreement.

The exact security requirements for any given DetNet are necessarily specific to the use cases handled by that network. Thus, the reader is assumed to be familiar with the specific security requirements of their use cases, for example, those outlined in the DetNet Use Cases [RFC8578] and the Security Considerations sections of the DetNet documents applicable to the network technologies in use, for example, [RFC8939] for an IP data plane and [RFC8964] for an MPLS data plane. Readers can find a general introduction to the DetNet Architecture in [RFC8655], the DetNet Data Plane in [RFC8938], and the Flow Information Model in [RFC9016].

The DetNet technologies include ways to:

- * Assign data plane resources for DetNet flows in some or all of the intermediate nodes (routers) along the path of the flow
- * Provide explicit routes for DetNet flows that do not dynamically change with the network topology in ways that affect the quality of service received by the affected flow(s)
- * Distribute data from DetNet flow packets over time and/or space to ensure delivery of the data in each packet in spite of the loss of a path

This document includes sections considering DetNet component design as well as system design. The latter includes a taxonomy and analysis of threats, threat impacts and mitigations, and an association of attacks with use cases (based on Section 11 of [RFC8578]).

This document is based on the premise that there will be a very broad range of DetNet applications and use cases, ranging in size and scope from individual industrial machines to networks that span an entire country [RFC8578]. Thus, no single set of prescriptions (such as exactly which mitigation should be applied to which segment of a DetNet) can be applicable to all of them, and indeed any single one that we might prescribe would inevitably prove impractical for some use case, perhaps one that does not even exist at the time of this writing. Thus, we are not prescriptive here; we are stating the desired end result, with the understanding that most DetNet use cases will necessarily differ from each other, and there is no "one size fits all".

2. Abbreviations and Terminology

Information Technology (IT): The application of computers to store, study, retrieve, transmit, and manipulate data or information, often in the context of a business or other enterprise [IT-DEF].

Operational Technology (OT): The hardware and software dedicated to detecting or causing changes in physical processes through direct monitoring and/or control of physical devices such as valves, pumps, etc. [OT-DEF].

Component: A component of a DetNet system -- used here to refer to any hardware or software element of a DetNet that implements DetNet-specific functionality, for example, all or part of a

router, switch, or end system.

Device: Used here to refer to a physical entity controlled by the DetNet, for example, a motor.

Resource Segmentation: Used as a more general form for Network Segmentation (the act or practice of splitting a computer network into sub-networks, each being a network segment [NS-DEF]).

Controller Plane: In DetNet, the Controller Plane corresponds to the aggregation of the Control and Management Planes (see [RFC8655], Section 4.4.2).

3. Security Considerations for DetNet Component Design

This section provides guidance for implementers of components to be used in a DetNet.

As noted above, DetNet provides resource allocation, explicit routes, and redundant path support. Each of these has associated security implications, which are discussed in this section, in the context of component design. Detection, reporting and appropriate action in the case of packet arrival-time violations are also discussed.

3.1. Resource Allocation

3.1.1. Inviolable Flows

A DetNet system security designer relies on the premise that any resources allocated to a resource-reserved (OT-type) flow are inviolable; in other words, there is no physical possibility within a DetNet component that resources allocated to a given DetNet flow can be compromised by any type of traffic in the network. This includes malicious traffic as well as inadvertent traffic such as might be produced by a malfunctioning component, or due to interactions between components that were not sufficiently tested for interoperability. From a security standpoint, this is a critical assumption, for example, when designing against DoS attacks. In other words, with correctly designed components and security mechanisms, one can prevent malicious activities from impacting other resources.

However, achieving the goal of absolutely inviolable flows may not be technically or economically feasible for any given use case, given the broad range of possible use cases (e.g., [RFC8578]) and their associated security considerations as outlined in this document. It can be viewed as a continuum of security requirements, from isolated ultra-low latency systems that may have little security vulnerability (such as an industrial machine) to broadly distributed systems with many possible attack vectors and OT security concerns (such as a utility network). Given this continuum, the design principle employed in this document is to specify the desired end results, without being overly prescriptive in how the results are achieved, reflecting the understanding that no individual implementation is likely to be appropriate for every DetNet use case.

3.1.2. Design Trade-Off Considerations in the Use Cases Continuum

For any given DetNet use case and its associated security requirements, it is important for the DetNet system designer to understand the interaction and design trade-offs that inevitably need to be reconciled between the desired end results and the DetNet protocols, as well as the DetNet system and component design.

For any given component, as designed for any given use case (or scope of use cases), it is the responsibility of the component designer to

ensure that the premise of inviolable flows is supported to the extent that they deem necessary to support their target use cases.

For example, the component may include traffic shaping and policing at the ingress to prevent corrupted, malicious, or excessive packets from entering the network, thereby decreasing the likelihood that any traffic will interfere with any DetNet OT flow. The component may include integrity protection for some or all of the header fields such as those used for flow ID, thereby decreasing the likelihood that a packet whose flow ID has been compromised might be directed into a different flow path. The component may verify every single packet header at every forwarding location, or only at certain points. In any of these cases, the component may use dynamic performance analytics (Section 7.7) to cause action to be initiated to address the situation in an appropriate and timely manner, either at the data plane or controller plane, or both in concert. The component's software and hardware may include measures to ensure the integrity of the resource allocation/deallocation process. Other design aspects of the component may help ensure that the adverse effects of malicious traffic are more limited, for example, by protecting network control interfaces or minimizing cascade failures. The component may include features specific to a given use case, such as configuration of the response to a given sequential packet loss count.

Ultimately, due to cost and complexity factors, the security properties of a component designed for low-cost systems may be (by design) far inferior to a component with similar intended functionality, but designed for highly secure or otherwise critical applications, perhaps at substantially higher cost. Any given component is designed for some set of use cases and accordingly will have certain limitations on its security properties and vulnerabilities. It is thus the responsibility of the system designer to assure themselves that the components they use in their design are capable of satisfying their overall system security requirements.

3.1.3. Documenting the Security Properties of a Component

In order for the system designer to adequately understand the security-related behavior of a given component, the designer of any component intended for use with DetNet needs to clearly document the security properties of that component. For example, to address the case where a corrupted packet in which the flow identification information is compromised and thus may incidentally match the flow ID of another ("victim") DetNet flow, resulting in additional unauthorized traffic on the victim, the documentation might state that the component employs integrity protection on the flow identification fields.

3.1.4. Fail-Safe Component Behavior

Even when the security properties of a component are understood and well specified, if the component malfunctions, for example, due to physical circumstances unpredicted by the component designer, it may be difficult or impossible to fully prevent malfunction of the network. The degree to which a component is hardened against various types of failures is a distinguishing feature of the component and its design, and the overall system design can only be as strong as its weakest link.

However, all networks are subject to this level of uncertainty; it is not unique to DetNet. Having said that, DetNet raises the bar by changing many added latency scenarios from tolerable annoyances to unacceptable service violations. That in turn underscores the importance of system integrity, as well as correct and stable

configuration of the network and its nodes, as discussed in Section 1.

3.1.5. Flow Aggregation Example

As another example regarding resource allocation implementation, consider the implementation of Flow Aggregation for DetNet flows (as discussed in [RFC8938]). In this example, say there are N flows that are to be aggregated; thus, the bandwidth resources of the aggregate flow must be sufficient to contain the sum of the bandwidth reservation for the N flows. However, if one of those flows were to consume more than its individually allocated bandwidth, this could cause starvation of the other flows. Thus, simply providing and enforcing the calculated aggregate bandwidth may not be a complete solution; the bandwidth for each individual flow must still be guaranteed, for example, via ingress policing of each flow (i.e., before it is aggregated). Alternatively, if by some other means each flow to be aggregated can be trusted not to exceed its allocated bandwidth, the same goal can be achieved.

3.2. Explicit Routes

The DetNet-specific purpose for constraining the ability of the DetNet to reroute OT traffic is to maintain the specified service parameters (such as upper and lower latency boundaries) for a given flow. For example, if the network were to reroute a flow (or some part of a flow) based exclusively on statistical path usage metrics, or due to malicious activity, it is possible that the new path would have a latency that is outside the required latency bounds that were designed into the original TE-designed path, thereby violating the quality of service for the affected flow (or part of that flow).

However, it is acceptable for the network to reroute OT traffic in such a way as to maintain the specified latency bounds (and any other specified service properties) for any reason, for example, in response to a runtime component or path failure.

So from a DetNet security standpoint, the DetNet system designer can expect that any component designed for use in a DetNet will deliver the packets within the agreed-upon service parameters. For the component designer, this means that in order for a component to achieve that expectation, any component that is involved in controlling or implementing any change of the initially TE-configured flow routes must prevent rerouting of OT flows (whether malicious or accidental) that might adversely affect delivering the traffic within the specified service parameters.

3.3. Redundant Path Support

The DetNet provision for redundant paths (i.e., PREOF, or "Packet Replication, Elimination, and Ordering Functions"), as defined in the DetNet Architecture [RFC8655], provides the foundation for high reliability of a DetNet by virtually eliminating packet loss (i.e., to a degree that is implementation dependent) through hitless redundant packet delivery.

| Note: At the time of this writing, PREOF is not defined for the
| IP data plane.

It is the responsibility of the system designer to determine the level of reliability required by their use case and to specify redundant paths sufficient to provide the desired level of reliability (in as much as that reliability can be provided through the use of redundant paths). It is the responsibility of the component designer to ensure that the relevant PREOF operations are executed reliably and securely to avoid potentially catastrophic

situations for the operational technology relying on them.

However, note that not all PREOF operations are necessarily implemented in every network; for example, a packet reordering function may not be necessary if the packets are either not required to be in order or if the ordering is performed in some other part of the network.

Ideally, a redundant path for a flow could be specified from end to end; however, given that this is not always possible (as described in [RFC8655]), the system designer will need to consider the resulting end-to-end reliability and security resulting from any given arrangement of network segments along the path, each of which provides its individual PREOF implementation and thus its individual level of reliability and security.

At the data plane, the implementation of PREOF depends on the correct assignment and interpretation of packet sequence numbers, as well as the actions taken based on them, such as elimination (including elimination of packets with spurious sequence numbers). Thus, the integrity of these values must be maintained by the component as they are assigned by the DetNet Data Plane Service sub-layer and transported by the Forwarding sub-layer. This is no different than the integrity of the values in any header used by the DetNet (or any other) data plane and is not unique to redundant paths. The integrity protection of header values is technology dependent; for example, in Layer 2 networks, the integrity of the header fields can be protected by using MACsec [IEEE802.1AE-2018]. Similarly, from the sequence number injection perspective, it is no different from any other protocols that use sequence numbers; for particulars of integrity protection via IPsec Authentication Headers, useful insights are provided by Section 3 of [RFC4302].

3.4. Timing (or Other) Violation Reporting

A task of the DetNet system designer is to create a network such that for any incoming packet that arrives with any timing or bandwidth violation, an appropriate action can be taken in order to prevent damage to the system. The reporting step may be accomplished through dynamic performance analysis (see Section 7.7) or by any other means as implemented in one or more components. The action to be taken for any given circumstance within any given application will depend on the use case. The action may involve intervention from the controller plane, or it may be taken "immediately" by an individual component, for example, if a very fast response is required.

The definitions and selections of the actions that can be taken are properties of the components. The component designer implements these options according to their expected use cases, which may vary widely from component to component. Clearly, selecting an inappropriate response to a given condition may cause more problems than it is intending to mitigate; for example, a naive approach might be to have the component shut down the link if a packet arrives outside of its prescribed time window. However, such a simplistic action may serve the attacker better than it serves the network. Similarly, simple logging of such issues may not be adequate since a delay in response could result in material damage, for example, to mechanical devices controlled by the network. Thus, a breadth of possible and effective security-related actions and their configuration is a positive attribute for a DetNet component.

Some possible violations that warrant detection include cases where a packet arrives:

- * Outside of its prescribed time window

- * Within its time window but with a compromised timestamp that makes it appear that it is not within its window
- * Exceeding the reserved flow bandwidth

Some possible direct actions that may be taken at the data plane include traffic policing and shaping functions (e.g., those described in [RFC2475]), separating flows into per-flow rate-limited queues, and potentially applying active queue management [RFC7567]. However, if those (or any other) actions are to be taken, the system designer must ensure that the results of such actions do not compromise the continued safe operation of the system. For example, the network (i.e., the controller plane and data plane working together) must mitigate in a timely fashion any potential adverse effect on mechanical devices controlled by the network.

4. DetNet Security Considerations Compared with Diffserv Security Considerations

DetNet is designed to be compatible with Diffserv [RFC2474] as applied to IT traffic in the DetNet. DetNet also incorporates the use of the 6-bit value of the Differentiated Services Code Point (DSCP) field of the Type of Service (IPv4) and Traffic Class (IPv6) bytes for flow identification. However, the DetNet interpretation of the DSCP value for OT traffic is not equivalent to the per-hop behavior (PHB) selection behavior as defined by Diffserv.

Thus, security considerations for DetNet have some aspects in common with Diffserv, in fact overlapping 100% with respect to IP IT traffic. Security considerations for these aspects are part of the existing literature on IP network security, specifically the Security Considerations sections of [RFC2474] and [RFC2475]. However, DetNet also introduces timing and other considerations that are not present in Diffserv, so the Diffserv security considerations are a subset of the DetNet security considerations.

In the case of DetNet OT traffic, the DSCP value is interpreted differently than in Diffserv and contributes to determination of the service provided to the packet. In DetNet, there are similar consequences to Diffserv for lack of detection of, or incorrect handling of, packets with mismarked DSCP values, and many of the points made in the Diffserv Security discussions (Section 6.1 of [RFC2475], Section 7 of [RFC2474], and Section 3.3.2.1 of [RFC6274]) are also relevant to DetNet OT traffic though perhaps in modified form. For example, in DetNet, the effect of an undetected or incorrectly handled maliciously mismarked DSCP field in an OT packet is not identical to affecting the PHB of that packet, since DetNet does not use the PHB concept for OT traffic. Nonetheless, the service provided to the packet could be affected, so mitigation measures analogous to those prescribed by Diffserv would be appropriate for DetNet. For example, mismarked DSCP values should not cause failure of network nodes. The remarks in [RFC2474] regarding IPsec and Tunneling Interactions are also relevant (though this is not to say that other sections are less relevant).

In this discussion, interpretation (and any possible intentional re-marking) of the DSCP values of packets destined for DetNet OT flows is expected to occur at the ingress to the DetNet domain; once inside the domain, maintaining the integrity of the DSCP values is subject to the same handling considerations as any other field in the packet.

5. Security Threats

This section presents a taxonomy of threats and analyzes the possible threats in a DetNet-enabled network. The threats considered in this section are independent of any specific technologies used to

implement the DetNet; Section 10 considers attacks that are associated with the DetNet technologies encompassed by [RFC8938].

We distinguish controller plane threats from data plane threats. The attack surface may be the same, but the types of attacks, as well as the motivation behind them, are different. For example, a Delay attack is more relevant to the data plane than to the controller plane. There is also a difference in terms of security solutions; the way you secure the data plane is often different than the way you secure the controller plane.

5.1. Threat Taxonomy

This document employs organizational elements of the threat models of [RFC7384] and [RFC7835]. This model classifies attackers based on two criteria:

Internal vs. external:

Internal attackers either have access to a trusted segment of the network or possess the encryption or authentication keys.

External attackers, on the other hand, do not have the keys and have access only to the encrypted or authenticated traffic.

On-path vs. off-path:

On-path attackers are located in a position that allows interception, modification, or dropping of in-flight protocol packets, whereas off-path attackers can only attack by generating protocol packets.

Regarding the boundary between internal vs. external attackers as defined above, note that in this document we do not make concrete recommendations regarding which specific segments of the network are to be protected in any specific way, for example, via encryption or authentication. As a result, the boundary as defined above is not unequivocally specified here. Given that constraint, the reader can view an internal attacker as one who can operate within the perimeter defined by the DetNet Edge Nodes (as defined in the DetNet Architecture [RFC8655]), allowing that the specifics of what is encrypted or authenticated within this perimeter will vary depending on the implementation.

Care has also been taken to adhere to Section 5 of [RFC3552], both with respect to which attacks are considered out of scope for this document, and also which are considered to be the most common threats (explored further in Section 5.2). Most of the direct threats to DetNet are active attacks (i.e., attacks that modify DetNet traffic), but it is highly suggested that DetNet application developers take appropriate measures to protect the content of the DetNet flows from passive attacks (i.e., attacks that observe but do not modify DetNet traffic), for example, through the use of TLS or DTLS.

DetNet-Service, one of the service scenarios described in [DETNET-SERVICE-MODEL], is the case where a service connects DetNet islands, i.e., two or more otherwise independent DetNets are connected via a link that is not intrinsically part of either network. This implies that there could be DetNet traffic flowing over a non-DetNet link, which may provide an attacker with an advantageous opportunity to tamper with DetNet traffic. The security properties of non-DetNet links are outside of the scope of DetNet Security, but it should be noted that use of non-DetNet services to interconnect DetNets merits security analysis to ensure the integrity of the networks involved.

5.2. Threat Analysis

5.2.1. Delay

An attacker can maliciously delay DetNet data flow traffic. By delaying the traffic, the attacker can compromise the service of applications that are sensitive to high delays or to high delay variation. The delay may be constant or modulated.

5.2.2. DetNet Flow Modification or Spoofing

An attacker can modify some header fields of en route packets in a way that causes the DetNet flow identification mechanisms to misclassify the flow. Alternatively, the attacker can inject traffic that is tailored to appear as if it belongs to a legitimate DetNet flow. The potential consequence is that the DetNet flow resource allocation cannot guarantee the performance that is expected when the flow identification works correctly.

5.2.3. Resource Segmentation (Inter-segment Attack) Vulnerability

DetNet components are expected to split their resources between DetNet flows in a way that prevents traffic from one DetNet flow from affecting the performance of other DetNet flows and also prevents non-DetNet traffic from affecting DetNet flows. However, perhaps due to implementation constraints, some resources may be partially shared, and an attacker may try to exploit this property. For example, an attacker can inject traffic in order to exhaust network resources such that DetNet packets that share resources with the injected traffic may be dropped or delayed. Such injected traffic may be part of DetNet flows or non-DetNet traffic.

Another example of a Resource Segmentation attack is the case in which an attacker is able to overload the exception path queue on the router, i.e., a "slow path" typically taken by control or OAM packets that are diverted from the data plane because they require processing by a CPU. DetNet OT flows are typically configured to take the "fast path" through the data plane to minimize latency. However, if there is only one queue from the forwarding Application-Specific Integrated Circuit (ASIC) to the exception path, and for some reason the system is configured such that any DetNet packets must be handled on this exception path, then saturating the exception path could result in the delaying or dropping of DetNet packets.

5.2.4. Packet Replication and Elimination

5.2.4.1. Replication: Increased Attack Surface

Redundancy is intended to increase the robustness and survivability of DetNet flows, and replication over multiple paths can potentially mitigate an attack that is limited to a single path. However, the fact that packets are replicated over multiple paths increases the attack surface of the network, i.e., there are more points in the network that may be subject to attacks.

5.2.4.2. Replication-Related Header Manipulation

An attacker can manipulate the replication-related header fields. This capability opens the door for various types of attacks. For example:

Forward both replicas:

Malicious change of a packet SN (Sequence Number) can cause both replicas of the packet to be forwarded. Note that this attack has a similar outcome to a replay attack.

Eliminate both replicas:

SN manipulation can be used to cause both replicas to be eliminated. In this case, an attacker that has access to a single

path can cause packets from other paths to be dropped, thus compromising some of the advantage of path redundancy.

Flow hijacking:

An attacker can hijack a DetNet flow with access to a single path by systematically replacing the SNs on the given path with higher SN values. For example, an attacker can replace every SN value S with a higher value $S+C$, where C is a constant integer. Thus, the attacker creates a false illusion that the attacked path has the lowest delay, causing all packets from other paths to be eliminated in favor of the attacked path. Once the flow from the compromised path is favored by the eliminating bridge, the flow has effectively been hijacked by the attacker. It is now possible for the attacker to either replace en route packets with malicious packets, or to simply inject errors into the packets, causing the packets to be dropped at their destination.

Amplification:

An attacker who injects packets into a flow that is to be replicated will have their attack amplified through the replication process. This is no different than any attacker who injects packets that are delivered through multicast, broadcast, or other point-to-multi-point mechanisms.

5.2.5. Controller Plane

5.2.5.1. Path Choice Manipulation

5.2.5.1.1. Control or Signaling Packet Modification

An attacker can maliciously modify en route control packets in order to disrupt or manipulate the DetNet path/resource allocation.

5.2.5.1.2. Control or Signaling Packet Injection

An attacker can maliciously inject control packets in order to disrupt or manipulate the DetNet path/resource allocation.

5.2.5.1.3. Increased Attack Surface

One of the possible consequences of a Path Manipulation attack is an increased attack surface. Thus, when the attack described in the previous subsection is implemented, it may increase the potential of other attacks to be performed.

5.2.5.2. Compromised Controller

An attacker can subvert a legitimate controller (or subvert another component such that it represents itself as a legitimate controller) with the result that the network nodes incorrectly believe it is authorized to instruct them.

The presence of a compromised node or controller in a DetNet is not a threat that arises as a result of determinism or time sensitivity; the same techniques used to prevent or mitigate against compromised nodes in any network are equally applicable in the DetNet case. The act of compromising a controller may not even be within the capabilities of our defined attacker types -- in other words, it may not be achievable via packet traffic at all, whether internal or external, on path or off path. It might be accomplished, for example, by a human with physical access to the component, who could upload bogus firmware to it via a USB stick. All of this underscores the requirement for careful overall system security design in a DetNet, given that the effects of even one bad actor on the network can be potentially catastrophic.

Security concerns specific to any given controller plane technology used in DetNet will be addressed by the DetNet documents associated with that technology.

5.2.6. Reconnaissance

A passive eavesdropper can identify DetNet flows and then gather information about en route DetNet flows, e.g., the number of DetNet flows, their bandwidths, their schedules, or other temporal or statistical properties. The gathered information can later be used to invoke other attacks on some or all of the flows.

DetNet flows are typically uniquely identified by their 6-tuple, i.e., fields within the L3 or L4 header. However, in some implementations, the flow ID may also be augmented by additional per-flow attributes known to the system, e.g., above L4. For the purpose of this document, we assume any such additional fields used for flow ID are encrypted and/or integrity protected from external attackers. Note however that existing OT protocols designed for use on dedicated secure networks may not intrinsically provide such protection, in which case IPsec or transport-layer security mechanisms may be needed.

5.2.7. Time-Synchronization Mechanisms

An attacker can use any of the attacks described in [RFC7384] to attack the synchronization protocol, thus affecting the DetNet service.

5.3. Threat Summary

A summary of the attacks that were discussed in this section is presented in Table 1. For each attack, the table specifies the type of attackers that may invoke the attack. In the context of this summary, the distinction between internal and external attacks is under the assumption that a corresponding security mechanism is being used, and that the corresponding network equipment takes part in this mechanism.

Attack	Attacker Type			
	Internal		External	
	On-Path	Off-Path	On-Path	Off-Path
Delay Attack	+		+	
DetNet Flow Modification or Spoofing	+	+		
Inter-segment Attack	+	+	+	+
Replication: Increased Attack Surface	+	+	+	+
Replication-Related Header Manipulation	+			
Path Manipulation	+	+		
Path Choice: Increased Attack Surface	+	+	+	+

Control or Signaling Packet Modification	+			
Control or Signaling Packet Injection	+	+		
Reconnaissance	+		+	
Attacks on Time- Synchronization Mechanisms	+	+	+	+

Table 1: Threat Analysis Summary

6. Security Threat Impacts

When designing security for a DetNet, as with any network, it may be prohibitively expensive or technically infeasible to thoroughly protect against every possible threat. Thus, the security designer must be informed (for example, by an application domain expert such as a product manager) regarding the relative significance of the various threats and their impact if a successful attack is carried out. In this section, we present an example of a possible template for such a communication, culminating in a table (Table 2) that lists a set of threats under consideration, and some values characterizing their relative impact in the context of a given industry. The specific threats, industries, and impact values in the table are provided only as an example of this kind of assessment and its communication; they are not intended to be taken literally.

This section considers assessment of the relative impacts of the attacks described in Section 5. In this section, the impacts as described assume that the associated mitigation is not present or has failed. Mitigations are discussed in Section 7.

In computer security, the impact (or consequence) of an incident can be measured in loss of confidentiality, integrity, or availability of information. In the case of OT or time sensitive networks (though not to the exclusion of IT or non-time-sensitive networks), the impact of an exploit can also include failure or malfunction of mechanical and/or other physical systems.

DetNet raises these stakes significantly for OT applications, particularly those that may have been designed to run in an OT-only environment and thus may not have been designed for security in an IT environment with its associated components, services, and protocols.

The extent of impact of a successful vulnerability exploit varies considerably by use case and by industry; additional insight regarding the individual use cases is available from "Deterministic Networking Use Cases" [RFC8578]. Each of those use cases is represented in Table 2, including Pro Audio, Electrical Utilities, Industrial M2M (split into two areas: M2M Data Gathering and M2M Control Loop), and others.

Aspects of Impact (left column) include Criticality of Failure, Effects of Failure, Recovery, and DetNet Functional Dependence. Criticality of failure summarizes the seriousness of the impact. The impact of a resulting failure can affect many different metrics that vary greatly in scope and severity. In order to reduce the number of variables, only the following were included: Financial, Health and Safety, Effect on a Single Organization, and Effect on Multiple Organizations. Recovery outlines how long it would take for an affected use case to get back to its pre-failure state (Recovery Time

Objective, RTO) and how much of the original service would be lost in between the time of service failure and recovery to original state (Recovery Point Objective, RPO). DetNet dependence maps how much the following DetNet service objectives contribute to impact of failure: time dependency, data integrity, source node integrity, availability, and latency/jitter.

The scale of the Impact mappings is low, medium, and high. In some use cases, there may be a multitude of specific applications in which DetNet is used. For simplicity, this section attempts to average the varied impacts of different applications. This section does not address the overall risk of a certain impact that would require the likelihood of a failure happening.

In practice, any such ratings will vary from case to case; the ratings shown here are given as examples.

	PRO A	Util	Bldg	Wireless	Cell	M2M Data	M2M Ctrl
Criticality	Med	Hi	Low	Med	Med	Med	Med
Effects							
Financial	Med	Hi	Med	Med	Low	Med	Med
Health/ Safety	Med	Hi	Hi	Med	Med	Med	Med
Affects 1 org	Hi	Hi	Med	Hi	Med	Med	Med
Affects >1 org	Med	Hi	Low	Med	Med	Med	Med
Recovery							
Recov Time Obj	Med	Hi	Med	Hi	Hi	Hi	Hi
Recov Point Obj	Med	Hi	Low	Med	Low	Hi	Hi
DetNet Dependence							
Time Dependence	Hi	Hi	Low	Hi	Med	Low	Hi
Latency/ Jitter	Hi	Hi	Med	Med	Low	Low	Hi
Data Integrity	Hi	Hi	Med	Hi	Low	Hi	Hi
Src Node Integ	Hi	Hi	Med	Hi	Med	Hi	Hi
Availability	Hi	Hi	Med	Hi	Low	Hi	Hi

Table 2: Impact of Attacks by Use Case Industry

The rest of this section will cover impact of the different groups in more detail.

6.1. Delay Attacks

6.1.1. Data Plane Delay Attacks

Note that "Delay attack" also includes the possibility of a "negative delay" or early arrival of a packet, or possibly adversely changing the timestamp value.

Delayed messages in a DetNet link can result in the same behavior as dropped messages in ordinary networks, since the services attached to the DetNet flow are likely to have strict delivery time requirements.

For a single-path scenario, disruption within the single flow is a real possibility. In a multipath scenario, large delays or instabilities in one DetNet flow can also lead to increased buffer and processor resource consumption at the eliminating router.

A data plane Delay attack on a system controlling substantial moving devices, for example, in industrial automation, can cause physical damage. For example, if the network promises a bounded latency of 2 ms for a flow, yet the machine receives it with 5 ms latency, the control loop of the machine may become unstable.

6.1.2. Controller Plane Delay Attacks

In and of itself, this is not directly a threat to the DetNet service, but the effects of delaying control messages can have quite adverse effects later.

- * Delayed teardown can lead to resource leakage, which in turn can result in failure to allocate new DetNet flows, finally giving rise to a denial-of-service attack.
- * Failure to deliver, or severely delaying, controller plane messages adding an endpoint to a multicast group will prevent the new endpoint from receiving expected frames thus disrupting expected behavior.
- * Delaying messages that remove an endpoint from a group can lead to loss of privacy, as the endpoint will continue to receive messages even after it is supposedly removed.

6.2. Flow Modification and Spoofing

6.2.1. Flow Modification

If the contents of a packet header or body can be modified by the attacker, this can cause the packet to be routed incorrectly or dropped, or the payload to be corrupted or subtly modified. Thus, the potential impact of a Modification attack includes disrupting the application as well as the network equipment.

6.2.2. Spoofing

6.2.2.1. Data Plane Spoofing

Spoofing data plane messages can result in increased resource consumption on the routers throughout the network as it will increase buffer usage and processor utilization. This can lead to resource exhaustion and/or increased delay.

If the attacker manages to create valid headers, the false messages can be forwarded through the network, using part of the allocated bandwidth. This in turn can cause legitimate messages to be dropped when the resource budget has been exhausted.

Finally, the endpoint will have to deal with invalid messages being delivered to the endpoint instead of (or in addition to) a valid message.

6.2.2.2. Controller Plane Spoofing

A successful Controller Plane Spoofing attack will potentially have adverse effects. It can do virtually anything from:

- * modifying existing DetNet flows by changing the available bandwidth
- * adding or removing endpoints from a DetNet flow
- * dropping DetNet flows completely
- * falsely creating new DetNet flows (exhausting the systems resources or enabling DetNet flows that are outside the control of the network engineer)

6.3. Segmentation Attacks (Injection)

6.3.1. Data Plane Segmentation

Injection of false messages in a DetNet flow could lead to exhaustion of the available bandwidth for that flow if the routers attribute these false messages to the resource budget of that flow.

In a multipath scenario, injected messages will cause increased processor utilization in elimination routers. If enough paths are subject to malicious injection, the legitimate messages can be dropped. Likewise, it can cause an increase in buffer usage. In total, it will consume more resources in the routers than normal, giving rise to a resource-exhaustion attack on the routers.

If a DetNet flow is interrupted, the end application will be affected by what is now a non-deterministic flow. Note that there are many possible sources of flow interruptions, for example, but not limited to, such physical-layer conditions as a broken wire or a radio link that is compromised by interference.

6.3.2. Controller Plane Segmentation

In a successful Controller Plane Segmentation attack, control messages are acted on by nodes in the network, unbeknownst to the central controller or the network engineer. This has the potential to:

- * create new DetNet flows (exhausting resources)
- * drop existing DetNet flows (denial of service)
- * add end stations to a multicast group (loss of privacy)
- * remove end stations from a multicast group (reduction of service)
- * modify the DetNet flow attributes (affecting available bandwidth)

If an attacker can inject control messages without the central controller knowing, then one or more components in the network may get into a state that is not expected by the controller. At that point, if the controller initiates a command, the effect of that command may not be as expected, since the target of the command may have started from a different initial state.

6.4. Replication and Elimination

The Replication and Elimination functions are relevant only to data plane messages as controller plane messages are not subject to multipath routing.

6.4.1. Increased Attack Surface

The impact of an increased attack surface is that it increases the probability that the network can be exposed to an attacker. This can facilitate a wide range of specific attacks, and their respective impacts are discussed in other subsections of this section.

6.4.2. Header Manipulation at Elimination Routers

This attack can potentially cause DoS to the application that uses the attacked DetNet flows or to the network equipment that forwards them. Furthermore, it can allow an attacker to manipulate the network paths and the behavior of the network layer.

6.5. Control or Signaling Packet Modification

If control packets are subject to manipulation undetected, the network can be severely compromised.

6.6. Control or Signaling Packet Injection

If an attacker can inject control packets undetected, the network can be severely compromised.

6.7. Reconnaissance

Of all the attacks, this is one of the most difficult to detect and counter.

An attacker can, at their leisure, observe over time various aspects of the messaging and signaling, learning the intent and purpose of the traffic flows. Then at some later date, possibly at an important time in the operational context, they might launch an attack based on that knowledge.

The flow ID in the header of the data plane messages gives an attacker a very reliable identifier for DetNet traffic, and this traffic has a high probability of going to lucrative targets.

Applications that are ported from a private OT network to the higher visibility DetNet environment may need to be adapted to limit distinctive flow properties that could make them susceptible to reconnaissance.

6.8. Attacks on Time-Synchronization Mechanisms

DetNet relies on an underlying time-synchronization mechanism; therefore, a compromised synchronization mechanism may cause DetNet nodes to malfunction. Specifically, DetNet flows may fail to meet their latency requirements and deterministic behavior, thus causing DoS to DetNet applications.

6.9. Attacks on Path Choice

This is covered in part in Section 6.3 (Segmentation Attacks (Injection)) and, as with Replication and Elimination (see Section 6.4), this is relevant for data plane messages.

7. Security Threat Mitigation

This section describes a set of measures that can be taken to

mitigate the attacks described in Section 5. These mitigations should be viewed as a set of tools, any of which can be used individually or in concert. The DetNet component and/or system and/or application designer can apply these tools as necessary based on a system-specific threat analysis.

Some of the technology-specific security considerations and mitigation approaches are further discussed in DetNet data plane solution documents, such as [RFC8938], [RFC8939], [RFC8964], [RFC9025], and [RFC9056].

7.1. Path Redundancy

Description: Path redundancy is a DetNet flow that can be forwarded simultaneously over multiple paths. Packet Replication and Elimination [RFC8655] provide resiliency to dropped or delayed packets. This redundancy improves the robustness to failures and to on-path attacks.

| Note: At the time of this writing, PREOF is not defined for
| the IP data plane.

Related attacks: Path redundancy can be used to mitigate various on-path attacks, including attacks described in Sections 5.2.1, 5.2.2, 5.2.3, and 5.2.7. However, it is also possible that multiple paths may make it more difficult to locate the source of an on-path attacker.

A Delay Modulation attack could result in extensively exercising otherwise unused code paths to expose hidden flaws. Subtle race conditions and memory allocation bugs in error-handling paths are classic examples of this.

7.2. Integrity Protection

Description: Integrity protection in the scope of DetNet is the ability to detect if a packet header has been modified (maliciously or otherwise) and if so, take some appropriate action (as discussed in Section 7.7). The decision on where in the network to apply integrity protection is part of the DetNet system design, and the implementation of the protection method itself is a part of a DetNet component design.

The most common technique for detecting header modification is the use of a Message Authentication Code (MAC) (see Section 10 for examples). The MAC can be distributed either in line (included in the same packet) or via a side channel. Of these, the in-line method is generally preferred due to the low latency that may be required on DetNet flows and the relative complexity and computational overhead of a sideband approach.

There are different levels of security available for integrity protection, ranging from the basic ability to detect if a header has been corrupted in transit (no malicious attack) to stopping a skilled and determined attacker capable of both subtly modifying fields in the headers as well as updating an unkeyed checksum. Common for all are the 2 steps that need to be performed in both ends. The first is computing the checksum or MAC. The corresponding verification step must perform the same steps before comparing the provided with the computed value. Only then can the receiver be reasonably sure that the header is authentic.

The most basic protection mechanism consists of computing a simple checksum of the header fields and providing it to the next entity in the packets path for verification. Using a MAC combined with a secret key provides the best protection against Modification and

Replication attacks (see Sections 5.2.2 and 5.2.4). This MAC usage needs to be part of a security association that is established and managed by a security association protocol (such as IKEv2 for IPsec security associations). Integrity protection in the controller plane is discussed in Section 7.6. The secret key, regardless of the MAC used, must be protected from falling into the hands of unauthorized users. Once key management becomes a topic, it is important to understand that this is a delicate process and should not be undertaken lightly. BCP 107 [BCP107] provides best practices in this regard.

DetNet system and/or component designers need to be aware of these distinctions and enforce appropriate integrity-protection mechanisms as needed based on a threat analysis. Note that adding integrity-protection mechanisms may introduce latency; thus, many of the same considerations in Section 7.5.1 also apply here.

Packet Sequence Number Integrity Considerations: The use of PREOF in a DetNet implementation implies the use of a sequence number for each packet. There is a trust relationship between the component that adds the sequence number and the component that removes the sequence number. The sequence number may be end-to-end source to destination, or it may be added/deleted by network edge components. The adder and remover(s) have the trust relationship because they are the ones that ensure that the sequence numbers are not modifiable. Thus, sequence numbers can be protected by using authenticated encryption or by a MAC without using encryption. Between the adder and remover there may or may not be replication and elimination functions. The elimination functions must be able to see the sequence numbers. Therefore, if encryption is done between adders and removers, it must not obscure the sequence number. If the sequence removers and the eliminators are in the same physical component, it may be possible to obscure the sequence number; however, that is a layer violation and is not recommended practice.

| Note: At the time of this writing, PREOF is not defined for
| the IP data plane.

Related attacks: Integrity protection mitigates attacks related to modification and tampering, including the attacks described in Sections 5.2.2 and 5.2.4.

7.3. DetNet Node Authentication

Description: Authentication verifies the identity of DetNet nodes (including DetNet Controller Plane nodes), and this enables mitigation of Spoofing attacks. While integrity protection (Section 7.2) prevents intermediate nodes from modifying information, authentication can provide traffic origin verification, i.e., to verify that each packet in a DetNet flow is from a known source. Although node authentication and integrity protection are two different goals of a security protocol, in most cases, a common protocol (such as IPsec [RFC4301] or MACsec [IEEE802.1AE-2018]) is used for achieving both purposes.

Related attacks: DetNet node authentication is used to mitigate attacks related to spoofing, including the attacks of Sections 5.2.2 and 5.2.4.

7.4. Synthetic Traffic Insertion

Description: With some queuing methods such as [IEEE802.1Qch-2017], it is possible to introduce synthetic traffic in order to regularize the timing of packet transmission. (Synthetic traffic typically consists of randomly generated packets injected in the

network to mask observable transmission patterns in the flows, which may allow an attacker to gain insight into the content of the flows). This can subsequently reduce the value of passive monitoring from internal threats (see Section 5) as it will be much more difficult to associate discrete events with particular network packets.

Related attacks: Removing distinctive temporal properties of individual packets or flows can be used to mitigate against reconnaissance attacks (Section 5.2.6). For example, synthetic traffic can be used to maintain constant traffic rate even when no user data is transmitted, thus making it difficult to collect information about the times at which users are active and the times at which DetNet flows are added or removed.

Traffic Insertion Challenges: Once an attacker is able to monitor the frames traversing a network to such a degree that they can differentiate between best-effort traffic and traffic belonging to a specific DetNet flow, it becomes difficult to not reveal to the attacker whether a given frame is valid traffic or an inserted frame. Thus, having the DetNet components generate and remove the synthetic traffic may or may not be a viable option unless certain challenges are solved; for example, but not limited to:

- * Inserted traffic must be indistinguishable from valid stream traffic from the viewpoint of the attacker.
- * DetNet components must be able to safely identify and remove all inserted traffic (and only inserted traffic).
- * The controller plane must manage where to insert and remove synthetic traffic, but this information must not be revealed to an attacker.

An alternative design is to have the insertion and removal of synthetic traffic be performed at the application layer rather than by the DetNet itself. For example, the use of RTP padding to reduce information leakage from variable-bit-rate audio transmission via the Secure Real-time Transport Protocol (SRTP) is discussed in [RFC6562].

7.5. Encryption

Description: Reconnaissance attacks (Section 5.2.6) can be mitigated to some extent through the use of encryption, thereby preventing the attacker from accessing the packet header or contents. Specific encryption protocols will depend on the lower layers that DetNet is forwarded over. For example, IP flows may be forwarded over IPsec [RFC4301], and Ethernet flows may be secured using MACsec [IEEE802.1AE-2018].

However, despite the use of encryption, a reconnaissance attack can provide the attacker with insight into the network, even without visibility into the packet. For example, an attacker can observe which nodes are communicating with which other nodes, including when, how often, and with how much data. In addition, the timing of packets may be correlated in time with external events such as action of an external device. Such information may be used by the attacker, for example, in mapping out specific targets for a different type of attack at a different time.

DetNet nodes do not have any need to inspect the payload of any DetNet packets, making them data agnostic. This means that end-to-end encryption at the application layer is an acceptable way to protect user data.

Note that reconnaissance is a threat that is not specific to DetNet flows; therefore, reconnaissance mitigation will typically be analyzed and provided by a network operator regardless of whether DetNet flows are deployed. Thus, encryption requirements will typically not be defined in DetNet technology-specific specifications, but considerations of using DetNet in encrypted environments will be discussed in these specifications. For example, Section 5.1.2.3 of [RFC8939] discusses flow identification of DetNet flows running over IPsec.

Related attacks: As noted above, encryption can be used to mitigate reconnaissance attacks (Section 5.2.6). However, for a DetNet to provide differentiated quality of service on a flow-by-flow basis, the network must be able to identify the flows individually. This implies that in a reconnaissance attack, the attacker may also be able to track individual flows to learn more about the system.

7.5.1. Encryption Considerations for DetNet

Any compute time that is required for encryption and decryption processing ("crypto") must be included in the flow latency calculations. Thus, cryptographic algorithms used in a DetNet must have bounded worst-case execution times, and these values must be used in the latency calculations. Fortunately, encryption and decryption operations typically are designed to have constant execution times in order to avoid side channel leakage.

Some cryptographic algorithms are symmetric in encode/decode time (such as AES), and others are asymmetric (such as public key algorithms). There are advantages and disadvantages to the use of either type in a given DetNet context. The discussion in this document relates to the timing implications of crypto for DetNet; it is assumed that integrity considerations are covered elsewhere in the literature.

Asymmetrical crypto is typically not used in networks on a packet-by-packet basis due to its computational cost. For example, if only endpoint checks or checks at a small number of intermediate points are required, asymmetric crypto can be used to authenticate distribution or exchange of a secret symmetric crypto key; a successful check based on that key will provide traffic origin verification as long as the key is kept secret by the participants. TLS (v1.3 [RFC8446], in particular, Section 4.1 ("Key Exchange Messages")) and IKEv2 [RFC6071] are examples of this for endpoint checks.

However, if secret symmetric keys are used for this purpose, the key must be given to all relays, which increases the probability of a secret key being leaked. Also, if any relay is compromised or faulty, then it may inject traffic into the flow. Group key management protocols can be used to automate management of such symmetric keys; for an example in the context of IPsec, see [IPSECME-G-IKEV2].

Alternatively, asymmetric crypto can provide traffic origin verification at every intermediate node. For example, a DetNet flow can be associated with an (asymmetric) keypair, such that the private key is available to the source of the flow and the public key is distributed with the flow information, allowing verification at every node for every packet. However, this is more computationally expensive.

In either case, origin verification also requires replay detection as part of the security protocol to prevent an attacker from recording and resending traffic, e.g., as a denial-of-service attack on flow forwarding resources.

In the general case, cryptographic hygiene requires the generation of new keys during the lifetime of an encrypted flow (e.g., see Section 9 of [RFC4253]), and any such key generation (or key exchange) requires additional computing time, which must be accounted for in the latency calculations for that flow. For modern ECDH (Elliptical Curve Diffie-Hellman) key-exchange operations (such as x25519 [RFC7748]), these operations can be performed in constant (predictable) time; however, this is not universally true (for example, for legacy RSA key exchange [RFC4432]). Thus, implementers should be aware of the time properties of these algorithms and avoid algorithms that make constant-time implementation difficult or impossible.

7.6. Control and Signaling Message Protection

Description: Control and signaling messages can be protected through the use of any or all of encryption, authentication, and integrity-protection mechanisms. Compared with data flows, the timing constraints for controller and signaling messages may be less strict, and the number of such packets may be fewer. If that is the case in a given application, then it may enable the use of asymmetric cryptography for the signing of both payload and headers for such messages, as well as encrypting the payload. Given that a DetNet is managed by a central controller, the use of a shared public key approach for these processes is well proven. This is further discussed in Section 7.5.1.

Related attacks: These mechanisms can be used to mitigate various attacks on the controller plane, as described in Sections 5.2.5, 5.2.7, and 5.2.5.1.

7.7. Dynamic Performance Analytics

Description: Incorporating Dynamic Performance Analytics (DPA) implies that the DetNet design includes a performance monitoring system to validate that timing guarantees are being met and to detect timing violations or other anomalies that may be the symptom of a security attack or system malfunction. If this monitoring system detects unexpected behavior, it must then cause action to be initiated to address the situation in an appropriate and timely manner, either at the data plane or controller plane or both in concert.

The overall DPA system can thus be decomposed into the "detection" and "notification" functions. Although the time-specific DPA performance indicators and their implementation will likely be specific to a given DetNet, and as such are nascent technology at the time of this writing, DPA is commonly used in existing networks so we can make some observations on how such a system might be implemented for a DetNet given that it would need to be adapted to address the time-specific performance indicators.

Detection Mechanisms: Measurement of timing performance can be done via "passive" or "active" monitoring, as discussed below.

Examples of passive monitoring strategies include:

- * Monitoring of queue and buffer levels, e.g., via active queue management (e.g., [RFC7567]).
- * Monitoring of per-flow counters.
- * Measurement of link statistics such as traffic volume, bandwidth, and QoS.

- * Detection of dropped packets.
- * Use of commercially available Network Monitoring tools.

Examples of active monitoring include:

- * In-band timing measurements (such as packet arrival times), e.g., by timestamping and packet inspection.
- * Use of OAM. For DetNet-specific OAM considerations, see [DETNET-IP-OAM] and [DETNET-MPLS-OAM]. Note: At the time of this writing, specifics of DPA have not been developed for the DetNet OAM but could be a subject for future investigation.
 - For OAM for Ethernet specifically, see also Connectivity Fault Management (CFM [IEEE802.1Q]), which defines protocols and practices for OAM for paths through 802.1 bridges and LANs.
- * Out-of-band detection. Following the data path or parts of a data path, for example, Bidirectional Forwarding Detection (BFD, e.g., [RFC5880]).

Note that for some measurements (e.g., packet delay), it may be necessary to make and reconcile measurements from more than one physical location (e.g., a source and destination), possibly in both directions, in order to arrive at a given performance indicator value.

Notification Mechanisms: Making DPA measurement results available at the right place(s) and time(s) to effect timely response can be challenging. Two notification mechanisms that are in general use are NETCONF/YANG Notifications and the proprietary local telemetry interfaces provided with components from some vendors. The Constrained Application Protocol (CoAP) Observe Option [RFC7641] could also be relevant to such scenarios.

At the time of this writing, YANG Notifications are not addressed by the DetNet YANG documents; however, this may be a topic for future work. It is possible that some of the passive mechanisms could be covered by notifications from non-DetNet-specific YANG modules; for example, if there is OAM or other performance monitoring that can monitor delay bounds, then that could have its own associated YANG data model, which could be relevant to DetNet, for example, some "threshold" values for timing measurement notifications.

At the time of this writing, there is an IETF Working Group for network/performance monitoring (IP Performance Metrics (IPPM)). See also previous work by the completed Remote Network Monitoring Working Group (RMONMIB). See also "An Overview of the IETF Network Management Standards", [RFC6632].

Vendor-specific local telemetry may be available on some commercially available systems, whereby the system can be programmed (via a proprietary dedicated port and API) to monitor and report on specific conditions, based on both passive and active measurements.

Related attacks: Performance analytics can be used to detect various attacks, including the ones described in Section 5.2.1 (Delay attack), Section 5.2.3 (Resource Segmentation attack), and Section 5.2.7 (Time-Synchronization attack). Once detection and notification have occurred, the appropriate action can be taken to mitigate the threat.

For example, in the case of data plane Delay attacks, one possible mitigation is to timestamp the data at the source and timestamp it again at the destination, and if the resulting latency does not meet the service agreement, take appropriate action. Note that DetNet specifies packet sequence numbering; however, it does not specify use of packet timestamps, although they may be used by the underlying transport (for example, TSN [IEEE802.1BA]) to provide the service.

7.8. Mitigation Summary

The following table maps the attacks of Section 5 (Security Threats) to the impacts of Section 6 (Security Threat Impacts) and to the mitigations of the current section. Each row specifies an attack, the impact of this attack if it is successfully implemented, and possible mitigation methods.

Attack	Impact	Mitigations
Delay Attack	<ul style="list-style-type: none"> * Non-deterministic delay * Data disruption * Increased resource consumption 	<ul style="list-style-type: none"> * Path redundancy * Performance analytics
Reconnaissance	<ul style="list-style-type: none"> * Enabler for other attacks 	<ul style="list-style-type: none"> * Encryption * Synthetic traffic insertion
DetNet Flow Modification or Spoofing	<ul style="list-style-type: none"> * Increased resource consumption * Data disruption 	<ul style="list-style-type: none"> * Path redundancy * Integrity protection * DetNet Node authentication
Inter-segment Attack	<ul style="list-style-type: none"> * Increased resource consumption * Data disruption 	<ul style="list-style-type: none"> * Path redundancy * Performance analytics
Replication: Increased Attack Resource	<ul style="list-style-type: none"> * All impacts of other attacks 	<ul style="list-style-type: none"> * Integrity protection * DetNet Node authentication * Encryption
Replication-Related Header Manipulation	<ul style="list-style-type: none"> * Non-deterministic delay * Data disruption 	<ul style="list-style-type: none"> * Integrity protection * DetNet Node authentication
Path Manipulation	<ul style="list-style-type: none"> * Enabler for other attacks 	<ul style="list-style-type: none"> * Control and signaling message

		protection
Path Choice: Increased Attack Surface	* All impacts of other attacks	* Control and signaling message protection
Control or Signaling Packet Modification	* Increased resource consumption * Non-deterministic delay * Data disruption	* Control and signaling message protection
Control or Signaling Packet Injection	* Increased resource consumption * Non-deterministic delay * Data disruption	* Control and signaling message protection
Attacks on Time- Synchronization Mechanisms	* Non-deterministic delay * Increased resource consumption * Data disruption	* Path redundancy * Control and signaling message protection * Performance analytics

Table 3: Mapping Attacks to Impact and Mitigations

8. Association of Attacks to Use Cases

Different attacks can have different impact and/or mitigation depending on the use case, so we would like to make this association in our analysis. However, since there is a potentially unbounded list of use cases, we categorize the attacks with respect to the common themes of the use cases as identified in Section 11 of [RFC8578].

See also Table 2 for a mapping of the impact of attacks per use case by industry.

8.1. Association of Attacks to Use Case Common Themes

In this section, we review each theme and discuss the attacks that are applicable to that theme, as well as anything specific about the impact and mitigations for that attack with respect to that theme. Table 5, Mapping between Themes and Attacks, then provides a summary of the attacks that are applicable to each theme.

8.1.1. Sub-network Layer

DetNet is expected to run over various transmission mediums, with Ethernet being the first identified. Attacks such as Delay or Reconnaissance might be implemented differently on a different transmission medium; however, the impact on the DetNet as a whole would be essentially the same. We thus conclude that all attacks and impacts that would be applicable to DetNet over Ethernet (i.e., all

those named in this document) would also be applicable to DetNet over other transmission mediums.

With respect to mitigations, some methods are specific to the Ethernet medium, for example, time-aware scheduling using 802.1Qbv [IEEE802.1Qbv-2015] can protect against excessive use of bandwidth at the ingress -- for other mediums, other mitigations would have to be implemented to provide analogous protection.

8.1.2. Central Administration

A DetNet network can be controlled by a centralized network configuration and control system. Such a system may be in a single central location, or it may be distributed across multiple control entities that function together as a unified control system for the network.

All attacks named in this document that are relevant to controller plane packets (and the controller itself) are relevant to this theme, including Path Manipulation, Path Choice, Control Packet Modification or Injection, Reconnaissance, and Attacks on Time-Synchronization Mechanisms.

8.1.3. Hot Swap

A DetNet network is not expected to be "plug and play"; it is expected that there is some centralized network configuration and control system. However, the ability to "hot swap" components (e.g., due to malfunction) is similar enough to "plug and play" that this kind of behavior may be expected in DetNet networks, depending on the implementation.

An attack surface related to hot swap is that the DetNet network must at least consider input at runtime from components that were not part of the initial configuration of the network. Even a "perfect" (or "hitless") replacement of a component at runtime would not necessarily be ideal, since presumably one would want to distinguish it from the original for OAM purposes (e.g., to report hot swap of a failed component).

This implies that an attack such as Flow Modification, Spoofing, or Inter-segment (which could introduce packets from a "new" component, i.e., one heretofore unknown on the network) could be used to exploit the need to consider such packets (as opposed to rejecting them out of hand as one would do if one did not have to consider introduction of a new component).

To mitigate this situation, deployments should provide a method for dynamic and secure registration of new components, and (possibly manual) deregistration and re-keying of retired components. This would avoid the situation in which the network must accommodate potentially insecure packet flows from unknown components.

Similarly, if the network was designed to support runtime replacement of a clock component, then presence (or apparent presence) and thus consideration of packets from a new such component could affect the network, or the time synchronization of the network, for example, by initiating a new Best Master Clock selection process. These types of attacks should therefore be considered when designing hot-swap-type functionality (see [RFC7384]).

8.1.4. Data Flow Information Models

DetNet specifies new YANG data models [DETNET-YANG] that may present new attack surfaces. Per IETF guidelines, security considerations for any YANG data model are expected to be part of the YANG data

model specification, as described in [IETF-YANG-SEC].

8.1.5. L2 and L3 Integration

A DetNet network integrates Layer 2 (bridged) networks (e.g., AVB/TSN LAN) and Layer 3 (routed) networks (e.g., IP) via the use of well-known protocols such as IP, MPLS Pseudowire, and Ethernet. Various DetNet documents address many specific aspects of Layer 2 and Layer 3 integration within a DetNet, and these are not individually referenced here; security considerations for those aspects are covered within those documents or within the related subsections of the present document.

Please note that although there are no entries in the L2 and L3 Integration line of the Mapping between Themes and Attacks table (Table 5), this does not imply that there could be no relevant attacks related to L2-L3 integration.

8.1.6. End-to-End Delivery

Packets that are part of a resource-reserved DetNet flow are not to be dropped by the DetNet due to congestion. Packets may however be dropped for intended reasons, for example, security measures. For example, consider the case in which a packet becomes corrupted (whether incidentally or maliciously) such that the resulting flow ID incidentally matches the flow ID of another DetNet flow, potentially resulting in additional unauthorized traffic on the latter. In such a case, it may be a security requirement that the system report and/or take some defined action, perhaps when a packet drop count threshold has been reached (see also Section 7.7).

A data plane attack may force packets to be dropped, for example, as a result of a Delay attack, Replication/Elimination attack, or Flow Modification attack.

The same result might be obtained by a Controller plane attack, e.g., Path Manipulation or Signaling Packet Modification.

An attack may also cause packets that should not be delivered to be delivered, such as by forcing packets from one (e.g., replicated) path to be preferred over another path when they should not be (Replication attack), or by Flow Modification, or Path Choice or Packet Injection. A Time-Synchronization attack could cause a system that was expecting certain packets at certain times to accept unintended packets based on compromised system time or time windowing in the scheduler.

8.1.7. Replacement for Proprietary Fieldbuses and Ethernet-Based Networks

There are many proprietary "fieldbuses" used in Industrial and other industries, as well as proprietary non-interoperable deterministic Ethernet-based networks. DetNet is intended to provide an open-standards-based alternative to such buses/networks. In cases where a DetNet intersects with such fieldbuses/networks or their protocols, such as by protocol emulation or access via a gateway, new attack surfaces can be opened.

For example, an Inter-segment or Controller plane attack such as Path Manipulation, Path Choice, or Control Packet Modification/Injection could be used to exploit commands specific to such a protocol or that are interpreted differently by the different protocols or gateway.

8.1.8. Deterministic vs. Best-Effort Traffic

Most of the themes described in this document address OT (reserved)

DetNet flows -- this item is intended to address issues related to IT traffic on a DetNet.

DetNet is intended to support coexistence of time-sensitive operational (OT, deterministic) traffic and informational (IT, "best effort") traffic on the same ("unified") network.

With DetNet, this coexistence will become more common, and mitigations will need to be established. The fact that the IT traffic on a DetNet is limited to a corporate-controlled network makes this a less difficult problem compared to being exposed to the open Internet; however, this aspect of DetNet security should not be underestimated.

An Inter-segment attack can flood the network with IT-type traffic with the intent of disrupting the handling of IT traffic and/or the goal of interfering with OT traffic. Presumably, if the DetNet flow reservation and isolation of the DetNet is well designed (better-designed than the attack), then interference with OT traffic should not result from an attack that floods the network with IT traffic.

The handling of IT traffic (i.e., traffic that by definition is not guaranteed any given deterministic service properties) by the DetNet will by definition not be given the DetNet-specific protections provided to DetNet (resource-reserved) flows. The implication is that the IT traffic on the DetNet network will necessarily have its own specific set of product (component or system) requirements for protection against attacks such as DoS; presumably they will be less stringent than those for OT flows, but nonetheless, component and system designers must employ whatever mitigations will meet the specified security requirements for IT traffic for the given component or DetNet.

The network design as a whole also needs to consider possible application-level dependencies of OT-type applications on services provided by the IT part of the network; for example, does the OT application depend on IT network services such as DNS or OAM? If such dependencies exist, how are malicious packet flows handled? Such considerations are typically outside the scope of DetNet proper, but nonetheless need to be addressed in the overall DetNet network design for a given use case.

8.1.9. Deterministic Flows

Reserved bandwidth data flows (deterministic flows) must provide the allocated bandwidth and must be isolated from each other.

A Spoofing or Inter-segment attack that adds packet traffic to a bandwidth-reserved DetNet flow could cause that flow to occupy more bandwidth than it was allocated, resulting in interference with other DetNet flows.

A Flow Modification, Spoofing, Header Manipulation, or Control Packet Modification attack could cause packets from one flow to be directed to another flow, thus breaching isolation between the flows.

8.1.10. Unused Reserved Bandwidth

If bandwidth reservations are made for a DetNet flow but the associated bandwidth is not used at any point in time, that bandwidth is made available on the network for best-effort traffic. However, note that security considerations for best-effort traffic on a DetNet network is out of scope of the present document, provided that any such attacks on best-effort traffic do not affect performance for DetNet OT traffic.

8.1.11. Interoperability

The DetNet specifications as a whole are intended to enable an ecosystem in which multiple vendors can create interoperable products, thus promoting component diversity and potentially higher numbers of each component manufactured. Toward that end, the security measures and protocols discussed in this document are intended to encourage interoperability.

Given that the DetNet specifications are unambiguously written and that the implementations are accurate, the property of interoperability should not in and of itself cause security concerns; however, flaws in interoperability between components could result in security weaknesses. The network operator, as well as system and component designers, can all contribute to reducing such weaknesses through interoperability testing.

8.1.12. Cost Reductions

The DetNet network specifications are intended to enable an ecosystem in which multiple vendors can create interoperable products, thus promoting higher numbers of each component manufactured, promoting cost reduction and cost competition among vendors.

This envisioned breadth of DetNet-enabled products is in general a positive factor; however, implementation flaws in any individual component can present an attack surface. In addition, implementation differences between components from different vendors can result in attack surfaces (resulting from their interaction) that may not exist in any individual component.

Network operators can mitigate such concerns through sufficient product and interoperability testing.

8.1.13. Insufficiently Secure Components

The DetNet network specifications are intended to enable an ecosystem in which multiple vendors can create interoperable products, thus promoting component diversity and potentially higher numbers of each component manufactured. However, this raises the possibility that a vendor might repurpose for DetNet applications a hardware or software component that was originally designed for operation in an isolated OT network and thus may not have been designed to be sufficiently secure, or secure at all, against the sorts of attacks described in this document. Deployment of such a component on a DetNet network that is intended to be highly secure may present an attack surface; thus, the DetNet network operator may need to take specific actions to protect such components, for example, by implementing a secure interface (such as a firewall) to isolate the component from the threats that may be present in the greater network.

8.1.14. DetNet Network Size

DetNet networks range in size from very small, e.g., inside a single industrial machine, to very large, e.g., a Utility Grid network spanning a whole country.

The size of the network might be related to how the attack is introduced into the network. For example, if the entire network is local, there is a threat that power can be cut to the entire network. If the network is large, perhaps only a part of the network is attacked.

A Delay attack might be as relevant to a small network as to a large network, although the amount of delay might be different.

Attacks sourced from IT traffic might be more likely in large networks since more people might have access to the network, presenting a larger attack surface. Similarly, Path Manipulation, Path Choice, and Time-Synchronization attacks seem more likely relevant to large networks.

8.1.15. Multiple Hops

Large DetNet networks (e.g., a Utility Grid network) may involve many "hops" over various kinds of links, for example, radio repeaters, microwave links, fiber optic links, etc.

An attacker who has knowledge of the operation of a component or device's internal software (such as "device drivers") may be able to take advantage of this knowledge to design an attack that could exploit flaws (or even the specifics of normal operation) in the communication between the various links.

It is also possible that a large-scale DetNet topology containing various kinds of links may not be in as common use as other more homogeneous topologies. This situation may present more opportunity for attackers to exploit software and/or protocol flaws in or between these components because these components or configurations may not have been sufficiently tested for interoperability (in the way they would be as a result of broad usage). This may be of particular concern to early adopters of new DetNet components or technologies.

Of the attacks we have defined, the ones identified in Section 8.1.14 as germane to large networks are the most relevant.

8.1.16. Level of Service

A DetNet is expected to provide means to configure the network that include querying network path latency, requesting bounded latency for a given DetNet flow, requesting worst-case maximum and/or minimum latency for a given path or DetNet flow, and so on. It is an expected case that the network cannot provide a given requested service level. In such cases, the network control system should reply that the requested service level is not available (as opposed to accepting the parameter but then not delivering the desired behavior).

Controller plane attacks such as Signaling Packet Modification and Injection could be used to modify or create control traffic that could interfere with the process of a user requesting a level of service and/or the reply from the network.

Reconnaissance could be used to characterize flows and perhaps target specific flows for attack via the controller plane as noted in Section 6.7.

8.1.17. Bounded Latency

DetNet provides the expectation of guaranteed bounded latency.

Delay attacks can cause packets to miss their agreed-upon latency boundaries.

Time-Synchronization attacks can corrupt the time reference of the system, resulting in missed latency deadlines (with respect to the "correct" time reference).

8.1.18. Low Latency

Applications may require "extremely low latency"; however, depending on the application, these may mean very different latency values.

For example, "low latency" across a Utility Grid network is on a different time scale than "low latency" in a motor control loop in a small machine. The intent is that the mechanisms for specifying desired latency include wide ranges, and that architecturally there is nothing to prevent arbitrarily low latencies from being implemented in a given network.

Attacks on the controller plane (as described in the Level of Service theme; see Section 8.1.16) and Delay and Time attacks (as described in the Bounded Latency theme; see Section 8.1.17) both apply here.

8.1.19. Bounded Jitter (Latency Variation)

DetNet is expected to provide bounded jitter (packet-to-packet latency variation).

Delay attacks can cause packets to vary in their arrival times, resulting in packet-to-packet latency variation, thereby violating the jitter specification.

8.1.20. Symmetrical Path Delays

Some applications would like to specify that the transit delay time values be equal for both the transmit and return paths.

Delay attacks can cause path delays to materially differ between paths.

Time-Synchronization attacks can corrupt the time reference of the system, resulting in path delays that may be perceived to be different (with respect to the "correct" time reference) even if they are not materially different.

8.1.21. Reliability and Availability

DetNet-based systems are expected to be implemented with essentially arbitrarily high availability (for example, 99.9999% up time, or even 12 nines). The intent is that the DetNet designs should not make any assumptions about the level of reliability and availability that may be required of a given system and should define parameters for communicating these kinds of metrics within the network.

Any attack on the system, of any type, can affect its overall reliability and availability; thus, in the mapping table (Table 5), we have marked every attack. Since every DetNet depends to a greater or lesser degree on reliability and availability, this essentially means that all networks have to mitigate all attacks, which to a greater or lesser degree defeats the purpose of associating attacks with use cases. It also underscores the difficulty of designing "extremely high reliability" networks.

In practice, network designers can adopt a risk-based approach in which only those attacks are mitigated whose potential cost is higher than the cost of mitigation.

8.1.22. Redundant Paths

This document expects that each DetNet system will be implemented to some essentially arbitrary level of reliability and/or availability, depending on the use case. A strategy used by DetNet for providing extraordinarily high levels of reliability when justified is to provide redundant paths between which traffic can be seamlessly switched, all the while maintaining the required performance of that system.

Replication-related attacks are by definition applicable here.

Controller plane attacks can also interfere with the configuration of redundant paths.

8.1.23. Security Measures

If any of the security mechanisms that protect the DetNet are attacked or subverted, this can result in malfunction of the network. Thus, the security systems themselves need to be robust against attacks.

The general topic of protection of security mechanisms is not unique to DetNet; it is identical to the case of securing any security mechanism for any network. This document addresses these concerns only to the extent that they are unique to DetNet.

8.2. Summary of Attack Types per Use Case Common Theme

The List of Attacks table (Table 4) lists the attacks described in Section 5, Security Threats, assigning a number to each type of attack. That number is then used as a short form identifier for the attack in Table 5, Mapping between Themes and Attacks.

	Attack
1	Delay Attack
2	DetNet Flow Modification or Spoofing
3	Inter-segment Attack
4	Replication: Increased Attack Surface
5	Replication-Related Header Manipulation
6	Path Manipulation
7	Path Choice: Increased Attack Surface
8	Control or Signaling Packet Modification
9	Control or Signaling Packet Injection
10	Reconnaissance
11	Attacks on Time-Synchronization Mechanisms

Table 4: List of Attacks

The Mapping between Themes and Attacks table (Table 5) maps the use case themes of [RFC8578] (as also enumerated in this document) to the attacks of Table 4. Each row specifies a theme, and the attacks relevant to this theme are marked with a "+". The row items that have no threats associated with them are included in the table for completeness of the list of Use Case Common Themes and do not have DetNet-specific threats associated with them.

[illegible]

Central Administration							+	+	+	+	+	+	
Hot Swap		+	+										+
Data Flow Information Models													
L2 and L3 Integration													
End-to-End Delivery	+	+	+	+	+	+	+	+	+		+		
Proprietary Deterministic Ethernet Networks			+				+	+	+	+			
Replacement for Proprietary Fieldbuses			+										
Deterministic vs. Best-Effort Traffic	+	+	+			+	+		+				
Deterministic Flows	+	+	+			+	+		+				
Unused Reserved Bandwidth		+	+						+	+			
Interoperability													
Cost Reductions													
Insufficiently Secure Components													
DetNet Network Size	+						+	+					+
Multiple Hops	+	+					+	+					+
Level of Service									+	+	+		
Bounded Latency	+												+
Low Latency	+								+	+			+
Bounded Jitter	+												
Symmetric Path Delays	+												+
Reliability and Availability	+	+	+	+	+	+	+	+	+	+	+	+	+
Redundant Paths				+	+				+	+			
Security Measures													

Table 5: Mapping between Themes and Attacks

9. Security Considerations for OAM Traffic

This section considers DetNet-specific security considerations for packet traffic that is generated and transmitted over a DetNet as part of OAM (Operations, Administration, and Maintenance). For the purposes of this discussion, OAM traffic falls into one of two basic types:

- * OAM traffic generated by the network itself. The additional bandwidth required for such packets is added by the network administration, presumably transparent to the customer. Security considerations for such traffic are not DetNet specific (apart from such traffic being subject to the same DetNet-specific security considerations as any other DetNet data flow) and are thus not covered in this document.
- * OAM traffic generated by the customer. From a DetNet security point of view, DetNet security considerations for such traffic are exactly the same as for any other customer data flows.

From the perspective of an attack, OAM traffic is indistinguishable from DetNet traffic, and the network needs to be secure against injection, removal, or modification of traffic of any kind, including OAM traffic. A DetNet is sensitive to any form of packet injection, removal, or manipulation, and in this respect DetNet OAM traffic is no different. Techniques for securing a DetNet against these threats have been discussed elsewhere in this document.

10. DetNet Technology-Specific Threats

Section 5, Security Threats, describes threats that are independent of a DetNet implementation. This section considers threats specifically related to the IP- and MPLS-specific aspects of DetNet implementations.

The primary security considerations for the data plane specifically are to maintain the integrity of the data and the delivery of the associated DetNet service traversing the DetNet network.

The primary relevant differences between IP and MPLS implementations are in flow identification and OAM methodologies.

As noted in [RFC8655], DetNet operates at the IP layer [RFC8939] and delivers service over sub-layer technologies such as MPLS [RFC8964] and IEEE 802.1 Time-Sensitive Networking (TSN) [RFC9023]. Application flows can be protected through whatever means are provided by the layer and sub-layer technologies. For example, technology-specific encryption may be used for IP flows (IPsec [RFC4301]). For IP-over-Ethernet (Layer 2) flows using an underlying sub-net, MACsec [IEEE802.1AE-2018] may be appropriate. For some use cases, packet integrity protection without encryption may be sufficient.

However, if the DetNet nodes cannot decrypt IPsec traffic, then DetNet flow identification for encrypted IP traffic flows must be performed in a different way than it would be for unencrypted IP DetNet flows. The DetNet IP data plane identifies unencrypted flows via a 6-tuple that consists of two IP addresses, the transport protocol ID, two transport protocol port numbers, and the DSCP in the IP header. When IPsec is used, the transport header is encrypted and the next protocol ID is an IPsec protocol, usually Encapsulating Security Payload (ESP), and not a transport protocol, leaving only three components of the 6-tuple, which are the two IP addresses and the DSCP. If the IPsec sessions are established by a controller, then this controller could also transmit (in the clear) the Security Parameter Index (SPI) and thus the SPI could be used (in addition to the pair of IP addresses) for flow identification. Identification of

DetNet flows over IPsec is further discussed in Section 5.1.2.3 of [RFC8939].

Sections below discuss threats specific to IP and MPLS in more detail.

10.1. IP

IP has a long history of security considerations and architectural protection mechanisms. From a data plane perspective, DetNet does not add or modify any IP header information, so the carriage of DetNet traffic over an IP data plane does not introduce any new security issues that were not there before, apart from those already described in the data-plane-independent threats section (Section 5).

Thus, the security considerations for a DetNet based on an IP data plane are purely inherited from the rich IP security literature and code/application base, and the data-plane-independent section of this document.

Maintaining security for IP segments of a DetNet may be more challenging than for the MPLS segments of the network given that the IP segments of the network may reach the edges of the network, which are more likely to involve interaction with potentially malevolent outside actors. Conversely, MPLS is inherently more secure than IP since it is internal to routers and it is well known how to protect it from outside influence.

Another way to look at DetNet IP security is to consider it in the light of VPN security. As an industry, we have a lot of experience with VPNs running through networks with other VPNs -- it is well known how to secure the network for that. However, for a DetNet, we have the additional subtlety that any possible interaction of one packet with another can have a potentially deleterious effect on the time properties of the flows. So the network must provide sufficient isolation between flows, for example, by protecting the forwarding bandwidth and related resources so that they are available to DetNet traffic, by whatever means are appropriate for the data plane of that network, for example, through the use of queuing mechanisms.

In a VPN, bandwidth is generally guaranteed over a period of time whereas in DetNet, it is not aggregated over time. This implies that any VPN-type protection mechanism must also maintain the DetNet timing constraints.

10.2. MPLS

An MPLS network carrying DetNet traffic is expected to be a "well-managed" network. Given that this is the case, it is difficult for an attacker to pass a raw MPLS-encoded packet into a network because operators have considerable experience at excluding such packets at the network boundaries as well as excluding MPLS packets being inserted through the use of a tunnel.

MPLS security is discussed extensively in [RFC5920] ("Security Framework for MPLS and GMPLS Networks") to which the reader is referred.

[RFC6941] builds on [RFC5920] by providing additional security considerations that are applicable to the MPLS-TP extensions appropriate to the MPLS Transport Profile [RFC5921] and thus to the operation of DetNet over some types of MPLS network.

[RFC5921] introduces to MPLS new Operations, Administration, and Maintenance (OAM) capabilities; a transport-oriented path protection mechanism; and strong emphasis on static provisioning supported by

network management systems.

The operation of DetNet over an MPLS network builds on MPLS and pseudowire encapsulation. Thus, for guidance on securing the DetNet elements of DetNet over MPLS, the reader is also referred to the security considerations of [RFC4385], [RFC5586], [RFC3985], [RFC6073], and [RFC6478].

Having attended to the conventional aspects of network security, it is necessary to attend to the dynamic aspects. The closest experience that the IETF has with securing protocols that are sensitive to manipulation of delay are the two-way time transfer (TWTT) protocols, which are NTP [RFC5905] and the Precision Time Protocol [IEEE1588]. The security requirements for these are described in [RFC7384].

One particular problem that has been observed in operational tests of TWTT protocols is the ability for two closely but not completely synchronized flows to beat and cause a sudden phase hit to one of the flows. This can be mitigated by the careful use of a scheduling system in the underlying packet transport.

Some investigations into protection of MPLS systems against dynamic attacks exist, such as [MPLS-OPP-ENCRYPT]; perhaps deployment of DetNets will encourage additional such investigations.

11. IANA Considerations

This document has no IANA actions.

12. Security Considerations

The security considerations of DetNet networks are presented throughout this document.

13. Privacy Considerations

Privacy in the context of DetNet is maintained by the base technologies specific to the DetNet and user traffic. For example, TSN can use MACsec, IP can use IPsec, and applications can use IP transport protocol-provided methods, e.g., TLS and DTLS. MPLS typically uses L2/L3 VPNs combined with the previously mentioned privacy methods.

However, note that reconnaissance threats such as traffic analysis and monitoring of electrical side channels can still cause there to be privacy considerations even when traffic is encrypted.

14. References

14.1. 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>>.
- [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>>.

14.2. Informative References

- [ARINC664P7]
ARINC, "Aircraft Data Network Part 7 Avionics Full-Duplex Switched Ethernet Network", ARINC 664 P7, September 2009.
- [BCP107] Bellovin, S. and R. Housley, "Guidelines for Cryptographic Key Management", BCP 107, RFC 4107, June 2005.

<<https://www.rfc-editor.org/info/bcp107>>
- [BCP72] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", BCP 72, RFC 3552, July 2003.

<<https://www.rfc-editor.org/info/bcp72>>
- [DETNET-IP-OAM]
Mirsky, G., Chen, M., and D. Black, "Operations, Administration and Maintenance (OAM) for Deterministic Networks (DetNet) with IP Data Plane", Work in Progress, Internet-Draft, draft-ietf-detnet-ip-oam-02, 30 March 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-detnet-ip-oam-02>>.
- [DETNET-MPLS-OAM]
Mirsky, G. and M. Chen, "Operations, Administration and Maintenance (OAM) for Deterministic Networks (DetNet) with MPLS Data Plane", Work in Progress, Internet-Draft, draft-ietf-detnet-mpls-oam-03, 30 March 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-detnet-mpls-oam-03>>.
- [DETNET-SERVICE-MODEL]
Varga, B., Ed. and J. Farkas, "DetNet Service Model", Work in Progress, Internet-Draft, draft-varga-detnet-service-model-02, May 2017, <<https://datatracker.ietf.org/doc/html/draft-varga-detnet-service-model-02>>.
- [DETNET-YANG]
Geng, X., Chen, M., Ryoo, Y., Fedyk, D., Rahman, R., and Z. Li, "Deterministic Networking (DetNet) YANG Model", Work in Progress, Internet-Draft, draft-ietf-detnet-yang-12, 19 May 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-detnet-yang-12>>.
- [IEEE1588] IEEE, "IEEE 1588 Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems", IEEE Std. 1588-2008, DOI 10.1109/IEEESTD.2008.4579760, July 2008, <<https://doi.org/10.1109/IEEESTD.2008.4579760>>.
- [IEEE802.1AE-2018]
IEEE, "IEEE Standard for Local and metropolitan area networks-Media Access Control (MAC) Security", IEEE Std. 802.1AE-2018, DOI 10.1109/IEEESTD.2018.8585421, December 2018, <<https://ieeexplore.ieee.org/document/8585421>>.
- [IEEE802.1BA]

IEEE, "IEEE Standard for Local and metropolitan area networks--Audio Video Bridging (AVB) Systems", IEEE Std. 802.1BA-2011, DOI 10.1109/IEEESTD.2011.6032690, September 2011, <<https://ieeexplore.ieee.org/document/6032690>>.

[IEEE802.1Q]

IEEE, "IEEE Standard for Local and metropolitan area networks--Bridges and Bridged Networks", IEEE Std. 802.1Q-2014, DOI 10.1109/IEEESTD.2014.6991462, December 2014, <<https://ieeexplore.ieee.org/document/6991462>>.

[IEEE802.1Qbv-2015]

IEEE, "IEEE Standard for Local and metropolitan area networks -- Bridges and Bridged Networks - Amendment 25: Enhancements for Scheduled Traffic", IEEE Std. 802.1Qbv-2015, DOI 10.1109/IEEESTD.2016.8613095, March 2016, <<https://ieeexplore.ieee.org/document/8613095>>.

[IEEE802.1Qch-2017]

IEEE, "IEEE Standard for Local and metropolitan area networks--Bridges and Bridged Networks--Amendment 29: Cyclic Queuing and Forwarding", IEEE Std. 802.1Qch-2017, DOI 10.1109/IEEESTD.2017.7961303, June 2017, <<https://ieeexplore.ieee.org/document/7961303>>.

[IETF-YANG-SEC]

IETF, "YANG module security considerations", October 2018, <<https://trac.ietf.org/trac/ops/wiki/yang-security-guidelines>>.

[IPSECME-G-IKEV2]

Smyslov, V. and B. Weis, "Group Key Management using IKEv2", Work in Progress, Internet-Draft, draft-ietf-ipsecme-g-ikev2-02, 11 January 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-ipsecme-g-ikev2-02>>.

[IT-DEF]

Wikipedia, "Information technology", March 2020, <https://en.wikiquote.org/w/index.php?title=Information_technology&oldid=2749907>.

[MPLS-OPP-ENCRYPT]

Farrel, A. and S. Farrell, "Opportunistic Security in MPLS Networks", Work in Progress, Internet-Draft, draft-ietf-mpls-opportunistic-encrypt-03, 28 March 2017, <<https://datatracker.ietf.org/doc/html/draft-ietf-mpls-opportunistic-encrypt-03>>.

[NS-DEF]

Wikipedia, "Network segmentation", December 2020, <https://en.wikipedia.org/w/index.php?title=Network_segmentation&oldid=993163264>.

[OT-DEF]

Wikipedia, "Operational technology", March 2021, <https://en.wikipedia.org/w/index.php?title=Operational_technology&oldid=1011704361>.

[RFC2474]

Nichols, K., Blake, S., Baker, F., and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", RFC 2474, DOI 10.17487/RFC2474, December 1998, <<https://www.rfc-editor.org/info/rfc2474>>.

[RFC2475]

Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., and W. Weiss, "An Architecture for Differentiated Services", RFC 2475, DOI 10.17487/RFC2475, December 1998, <<https://www.rfc-editor.org/info/rfc2475>>.

- [RFC3985] Bryant, S., Ed. and P. Pate, Ed., "Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture", RFC 3985, DOI 10.17487/RFC3985, March 2005, <<https://www.rfc-editor.org/info/rfc3985>>.
- [RFC4253] Ylonen, T. and C. Lonvick, Ed., "The Secure Shell (SSH) Transport Layer Protocol", RFC 4253, DOI 10.17487/RFC4253, January 2006, <<https://www.rfc-editor.org/info/rfc4253>>.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, DOI 10.17487/RFC4301, December 2005, <<https://www.rfc-editor.org/info/rfc4301>>.
- [RFC4302] Kent, S., "IP Authentication Header", RFC 4302, DOI 10.17487/RFC4302, December 2005, <<https://www.rfc-editor.org/info/rfc4302>>.
- [RFC4385] Bryant, S., Swallow, G., Martini, L., and D. McPherson, "Pseudowire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN", RFC 4385, DOI 10.17487/RFC4385, February 2006, <<https://www.rfc-editor.org/info/rfc4385>>.
- [RFC4432] Harris, B., "RSA Key Exchange for the Secure Shell (SSH) Transport Layer Protocol", RFC 4432, DOI 10.17487/RFC4432, March 2006, <<https://www.rfc-editor.org/info/rfc4432>>.
- [RFC5586] Bocci, M., Ed., Vigoureux, M., Ed., and S. Bryant, Ed., "MPLS Generic Associated Channel", RFC 5586, DOI 10.17487/RFC5586, June 2009, <<https://www.rfc-editor.org/info/rfc5586>>.
- [RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", RFC 5880, DOI 10.17487/RFC5880, June 2010, <<https://www.rfc-editor.org/info/rfc5880>>.
- [RFC5905] Mills, D., Martin, J., Ed., Burbank, J., and W. Kasch, "Network Time Protocol Version 4: Protocol and Algorithms Specification", RFC 5905, DOI 10.17487/RFC5905, June 2010, <<https://www.rfc-editor.org/info/rfc5905>>.
- [RFC5920] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", RFC 5920, DOI 10.17487/RFC5920, July 2010, <<https://www.rfc-editor.org/info/rfc5920>>.
- [RFC5921] Bocci, M., Ed., Bryant, S., Ed., Frost, D., Ed., Levrau, L., and L. Berger, "A Framework for MPLS in Transport Networks", RFC 5921, DOI 10.17487/RFC5921, July 2010, <<https://www.rfc-editor.org/info/rfc5921>>.
- [RFC6071] Frankel, S. and S. Krishnan, "IP Security (IPsec) and Internet Key Exchange (IKE) Document Roadmap", RFC 6071, DOI 10.17487/RFC6071, February 2011, <<https://www.rfc-editor.org/info/rfc6071>>.
- [RFC6073] Martini, L., Metz, C., Nadeau, T., Bocci, M., and M. Aissaoui, "Segmented Pseudowire", RFC 6073, DOI 10.17487/RFC6073, January 2011, <<https://www.rfc-editor.org/info/rfc6073>>.
- [RFC6274] Gont, F., "Security Assessment of the Internet Protocol Version 4", RFC 6274, DOI 10.17487/RFC6274, July 2011, <<https://www.rfc-editor.org/info/rfc6274>>.
- [RFC6478] Martini, L., Swallow, G., Heron, G., and M. Bocci, "Pseudowire Status for Static Pseudowires", RFC 6478,

DOI 10.17487/RFC6478, May 2012,
<<https://www.rfc-editor.org/info/rfc6478>>.

- [RFC6562] Perkins, C. and JM. Valin, "Guidelines for the Use of Variable Bit Rate Audio with Secure RTP", RFC 6562, DOI 10.17487/RFC6562, March 2012, <<https://www.rfc-editor.org/info/rfc6562>>.
- [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>>.
- [RFC6941] Fang, L., Ed., Niven-Jenkins, B., Ed., Mansfield, S., Ed., and R. Graveman, Ed., "MPLS Transport Profile (MPLS-TP) Security Framework", RFC 6941, DOI 10.17487/RFC6941, April 2013, <<https://www.rfc-editor.org/info/rfc6941>>.
- [RFC7384] Mizrahi, T., "Security Requirements of Time Protocols in Packet Switched Networks", RFC 7384, DOI 10.17487/RFC7384, October 2014, <<https://www.rfc-editor.org/info/rfc7384>>.
- [RFC7567] Baker, F., Ed. and G. Fairhurst, Ed., "IETF Recommendations Regarding Active Queue Management", BCP 197, RFC 7567, DOI 10.17487/RFC7567, July 2015, <<https://www.rfc-editor.org/info/rfc7567>>.
- [RFC7641] Hartke, K., "Observing Resources in the Constrained Application Protocol (CoAP)", RFC 7641, DOI 10.17487/RFC7641, September 2015, <<https://www.rfc-editor.org/info/rfc7641>>.
- [RFC7748] Langley, A., Hamburg, M., and S. Turner, "Elliptic Curves for Security", RFC 7748, DOI 10.17487/RFC7748, January 2016, <<https://www.rfc-editor.org/info/rfc7748>>.
- [RFC7835] Saucez, D., Iannone, L., and O. Bonaventure, "Locator/ID Separation Protocol (LISP) Threat Analysis", RFC 7835, DOI 10.17487/RFC7835, April 2016, <<https://www.rfc-editor.org/info/rfc7835>>.
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/info/rfc8446>>.
- [RFC8578] Grossman, E., Ed., "Deterministic Networking Use Cases", RFC 8578, DOI 10.17487/RFC8578, May 2019, <<https://www.rfc-editor.org/info/rfc8578>>.
- [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>>.
- [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>>.
- [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>>.

[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, June 2021, <<https://www.rfc-editor.org/info/rfc9056>>.

Contributors

The Editor would like to recognize the contributions of the following individuals to this document.

Stewart Bryant
Futurewei Technologies

Email: sb@stewartbryant.com

David Black
Dell EMC
176 South Street
Hopkinton, Massachusetts 01748
United States of America

Henrik Austad
SINTEF Digital
Klaebuveien 153
7037 Trondheim
Norway

Email: henrik@austad.us

John Dowdell
Airbus Defence and Space
Celtic Springs
Newport, NP10 8FZ
United Kingdom

Email: john.dowdell.ietf@gmail.com

Norman Finn
3101 Rio Way
Spring Valley, California 91977
United States of America

Email: nfinn@nfinnconsulting.com

Subir Das
Applied Communication Sciences
150 Mount Airy Road
Basking Ridge, New Jersey 07920
United States of America

Email: sdas@appcomsci.com

Carsten Bormann
Universitat Bremen TZI
Postfach 330440 D-28359 Bremen
Germany

Email: cabo@tzi.org

Authors' Addresses

Ethan Grossman (editor)
Dolby Laboratories, Inc.
1275 Market Street
San Francisco, CA 94103
United States of America

Email: ethan@ieee.org
URI: <https://www.dolby.com>

Tal Mizrahi
Huawei

Email: tal.mizrahi.phd@gmail.com

Andrew J. Hacker
Thought LLC
Harrisburg, PA
United States of America

Email: andrew@thought.live