

DetNet
Internet-Draft
Intended status: Standards Track
Expires: September 10, 2015

N. Finn
P. Thubert
Cisco
M. Johas Teener
Broadcom
March 9, 2015

Deterministic Networking Architecture
draft-finn-detnet-architecture-01

Abstract

Deterministic Networking (DetNet) provides a capability to carry specified unicast or multicast data streams for real-time applications with extremely low data loss rates and maximum latency. Techniques used include: 1) reserving data plane resources for individual (or aggregated) DetNet streams in some or all of the relay systems (bridges or routers) along the path of the stream; 2) providing fixed paths for DetNet streams that do not rapidly change with the network topology; and 3) sequentializing, replicating, and eliminating duplicate packets at various points to ensure the availability of at least one path. The capabilities can be managed by configuration, or by manual or automatic network management.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on September 10, 2015.

Copyright Notice

Copyright (c) 2015 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 (<http://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	3
2. Terminology	4
3. Providing the DetNet Quality of Service	5
3.1. Zero Congestion Loss	7
3.2. Pinned-down paths	7
3.3. Seamless Redundancy	8
4. DetNet Architecture	9
4.1. The Application Plane	11
4.2. The Controller Plane	11
4.3. The Network Plane	12
4.4. Elements of DetNet Architecture	13
4.5. DetNet streams	14
4.5.1. Talker guarantees	14
4.5.2. Incomplete Networks	15
4.6. Data Flow Model through Systems	16
4.7. Queuing, Shaping, Scheduling, and Preemption	16
4.8. Coexistence with normal traffic	16
4.9. Fault Mitigation	16
4.10. Protocol Stack Model	17
4.11. Advertising resources, capabilities and adjacencies	17
4.12. Provisioning model	17
4.12.1. Centralized Path Computation and Installation	17
4.12.2. Distributed Path Setup	17
5. Related IETF work	18
5.1. Deterministic PHB	18
5.2. 6TiSCH	18
6. Security Considerations	19
7. IANA Considerations	19
8. Acknowledgements	19
9. Access to IEEE 802.1 documents	19
10. Informative References	20
Authors' Addresses	23

1. Introduction

Operational Technology (OT) refers to industrial networks that are typically used for monitoring systems and supporting control loops, as well as movement detection systems for use in process control (i.e., process manufacturing) and factory automation (i.e., discrete manufacturing). Due to its different goals, OT has evolved in parallel but in a manner that is radically different from IT/ICT, focusing on highly secure, reliable and deterministic networks, with limited scalability over a bounded area.

The convergence of IT and OT technologies, also called the Industrial Internet, represents a major evolution for both sides. The work has already started; in particular, the industrial automation space has been developing a number of Ethernet-based replacements for existing digital control systems, often not packet-based (fieldbus technologies).

These replacements are meant to provide similar behavior as the incumbent protocols, and their common focus is to transport a fully characterized flow over a well-controlled environment (i.e., a factory floor), with a bounded latency, extraordinarily low frame loss, and a very narrow jitter. Examples of such protocols include PROFINET, ODVA Ethernet/IP, and EtherCAT.

In parallel, the need for determinism in professional and home audio/video markets drove the formation of the Audio/Video Bridging (AVB) standards effort of IEEE 802.1. With the explosion of demand for connectivity and multimedia in transportation in general, the Ethernet AVB technology has become one of the hottest topics, in particular in the automotive connectivity. It is finding application in all elements of the vehicle from head units, to rear seat entertainment modules, to amplifiers and camera modules. While aimed at less critical applications than some industrial networks, AVB networks share the requirement for extremely low packet loss rates and ensured finite latency and jitter.

Other instances of in-vehicle deterministic networks have arisen as well for control networks in cars, trains and buses, as well as avionics, with, for instance, the mission-critical "Avionics Full-Duplex Switched Ethernet" (AFDX) that was designed as part of the ARINC 664 standards. Existing automotive control networks such as the LIN, CAN and FlexRay standards were not designed to cover these increasing demands in terms of bandwidth and scalability that we see with various kinds of Driver Assistance Systems (DAS) and new multiplexing technologies based on Ethernet are now getting traction.

The generalization of the needs for more deterministic networks have led to the IEEE 802.1 AVB Task Group becoming the Time-Sensitive Networking (TSN) Task Group (TG), with a much-expanded constituency from the industrial and vehicular markets. Along with this expansion, the networks in consideration are becoming larger and structured, requiring deterministic forwarding beyond the LAN boundaries. For instance, Industrial Automation segregates the network along the broad lines of the Purdue Enterprise Reference Architecture (PERA), using different technologies at each level, and public infrastructures such as Electricity Automation require deterministic properties over the Wide Area. The realization is now coming that the convergence of IT and OT networks requires Layer-3, as well as Layer-2, capabilities.

The present architecture is the result of a collaboration of the IETF and the IEEE and implements an abstract model that can be applicable both at Layer-2 and Layer-3, and along segments of different technologie. With this new work, a path may span, for instance, across a (limited) number of 802.1 bridges and then a (limited) number of IP routers. In that example, the IEEE 802.1 bridges may be operating at Layer-2 over Ethernet whereas the IP routers may be 6TiSCH nodes operating at Layer-2 and/or Layer-3 over the IEEE 802.15.4e MAC.

Many applications of interest to Deterministic Networking require the ability to synchronize the clocks in end systems to a sub-microsecond accuracy. Some of the queue control techniques defined in Section 4.7 also require time synchronization among relay systems. The means used to achieve time synchronization are not addressed in this document.

2. Terminology

The following special terms are used in this document in order to avoid the assumption that a given element in the architecture does or does not have Internet Protocol stack, functions as a router or a bridge, or otherwise plays a particular role at Layer-3 or higher:

bridge

A Customer Bridge as defined by IEEE 802.1Q [IEEE802.1Q-2011].

circuit

A trail of configuration from talker to listener(s) through relay systems associated with a DetNet stream, required to deliver the benefits of DetNet.

end system

Commonly called a "host" in IETF documents, and an "end station" in IEEE 802 documents. End systems of interest to this document are talkers and listeners.

listener

An end system capable of sinking a DetNet stream.

relay system

A router or a bridge.

stream

A DetNet stream is a sequence of packets from a single talker, through some number of relay systems to one or more listeners, that is limited by the talker in its maximum packet size and transmission rate, and can thus be ensured the DetNet Quality of Service (QoS) from the network.

talker

An end system capable of sourcing a DetNet stream.

3. Providing the DetNet Quality of Service

DetNet Quality of Service is expressed in terms of:

- o Minimum and maximum end-to-end latency from talker to listener;
- o Probability of loss of a packet, assuming the normal operation of the relay systems and links;
- o Probability of loss of a packet in the event of the failure of a relay system or link.

It is a distinction of DetNet that it is concerned solely with worst-case values for all of the above parameters. Average, mean, or typical values are of no interest, because they do not affect the ability of a real-time system to perform its tasks.

Three techniques are employed by DetNet to achieve these QoS parameters:

- a. Zero congestion loss (Section 3.1). Network resources such as link bandwidth, buffers, queues, shapers, and scheduled input/output slots are assigned in each relay system to the use of a specific DetNet stream or group of streams. Note that, given a finite amount of buffer space, zero congestion loss necessarily ensures a maximum end-to-end latency. Depending on the method employed, a minimum latency can also be achieved.

- b. Pinned-down paths (Section 3.2). Point-to-point paths or point-to-multipoint trees through the network from a talker to one or more listeners can be established, and DetNet streams assigned to follow a particular path or tree.
- c. Packet replication and deletion (Section 3.3). End systems and/or relay systems can sequence number, replicate, and eliminate replicated packets at multiple points in the network in order to ensure that one (or more) equipment failure events still leave at least one path intact for a DetNet stream.

These three techniques can be applied independently, giving eight possible combinations, including none (no DetNet), although some combinations are of wider utility than others. This separation keeps the protocol stack coherent and maximizes interoperability with existing and developing standards in this (IETF) and other Standards Development Organizations. Some examples of typical expected combinations:

- o Pinned-down paths (a) plus packet replication (b) are exactly the techniques employed by [HSR-PRP]. Pinned-down paths are achieved by limiting the physical topology of the network, and the sequentialization, replication, and duplicate elimination facilitated by packet tags added at the front or the end of Ethernet frames.
- o Zero congestion loss (a) alone is offered by IEEE 802.1 Audio Video bridging [IEEE802.1BA-2011]. As long as the network suffers no failures, near-zero (at best, zero) congestion loss can be achieved through the use of a reservation protocol (MSRP) and shapers in every relay system (bridge).
- o Using all three together gives maximum protection.

There are, of course, simpler methods available (and employed, today) to achieve levels of latency and packet loss that are satisfactory for many applications. However, these methods generally work best in the absence of any significant amount of non-critical traffic in the network (if, indeed, such traffic is supported at all), or work only if the critical traffic constitutes only a small portion of the network's theoretical capacity, or work only if all systems are functioning properly, or in the absence of actions by end systems that disrupt the network's operations.

There are any number of methods in use, defined, or in progress for accomplishing each of the above techniques. It is expected that this DetNet Architecture will assist various vendors, users, and/or "vertical" Standards Development Organizations (dedicated to a single

industry) to make selections among the available means of implementing DetNet networks.

3.1. Zero Congestion Loss

The primary means by which DetNet achieves its QoS assurances is to completely eliminate congestion at an output port as a cause of packet loss. Given that a DetNet stream cannot be throttled, this can be achieved only by the provision of sufficient buffer storage at each hop through the network to ensure that no packets are dropped due to a lack of buffer storage.

Ensuring adequate buffering requires, in turn, that the talker, and every relay system along the path to the listener (or nearly every relay system -- see Section 4.5.2) be careful to regulate its output to not exceed the data rate for any stream, except for brief periods when making up for interfering traffic. Any packet sent ahead of its time potentially adds to the number of buffers required by the next hop, and may thus exceed the resources allocated for a particular stream.

The low-level mechanisms described in Section 4.7 provide the necessary regulation of transmissions by an edge system or relay system to ensure zero congestion loss. Of course, the reservation of the bandwidth and buffers for a stream requires the provisioning described in Section 4.12.

3.2. Pinned-down paths

In networks controlled by typical peer-to-peer protocols such as IEEE 802.1 ISIS bridged networks or ETF OSPF routed networks, a network topology event in one part of the network can impact, at least briefly, the delivery of data in parts of the network remote from the failure or recovery event. Thus, even redundant paths through a network, if controlled by the typical peer-to-peer protocols, do not eliminate the chances of brief losses of contact. For this reason, many real-time networks rely on physical rings of two-port devices, with a relatively simple ring control protocol. This both minimizes recovery time and easily supports redundant paths. Of course, this comes at the cost of increased hop count, and thus latency, for the typical path.

In order to get the advantages of low hop count and still ensure against even brief losses of connectivity, DetNet employs pinned-down paths, where the path taken by a given DetNet stream does not change, at least immediately, and likely not at all, in response to network topology events. When combined with seamless redundancy

(Section 3.3), this results in a high likelihood of continuous connectivity.

3.3. Seamless Redundancy

After congestion loss has been eliminated, the most important causes of packet loss are random media and/or memory faults and equipment failures.

Seamless redundancy involves three capabilities:

- o Adding sequence numbers to the packets of a DetNet stream.
- o Replicating these packets and, typically, sending them along at least two different paths to the listener(s).
- o Discarding duplicated packets.

In the simplest case, this amounts to replicating each packet in a talker that has two interfaces, and conveying them through the network, along separate paths, to the similarly dual-homed listeners, that discard the extras. This ensures that one path (with zero congestion loss) remains, even if some relay system fails.

Alternatively, relay systems in the network can provide replication and elimination facilities at various points in the network, so that multiple failures can be accommodated.

This is shown in the following figure, where the two relay systems each replicate (R) the DetNet stream on input, sending the stream to both the other relay system and to the end system, and eliminated duplicates (E) on the output interface to the right-hand end system. Any one links in the network can fail, and the Detnet stream can still get through. Furthermore, two links can fail, as long as they are in different segments of the network.




Figure 1

4. DetNet Architecture

Traffic Engineering Architecture and Signaling (TEAS) [TEAS] defines traffic-engineering architectures for generic applicability across packet and non-packet networks. From TEAS perspective, Traffic Engineering (TE) refers to techniques that enable operators to control how specific traffic flows are treated within their networks.

Because of its very nature of establishing pinned-down optimized paths, Deterministic Networking can be seen as a new, specialized branch of Traffic Engineering, and inherits its architecture with a separation into planes.

The Deterministic Networking architecture is thus composed of three planes, a (User) Application Plane, a Controller Plane, and a Network Plane, which echoes that of Software-Defined Networking (SDN): Layers and Architecture Terminology [RFC7426] which is represented below:

SDN Layers and Architecture Terminology per RFC 7426




Figure 2

4.1. The Application Plane

Per [RFC7426], the Application Plane includes both applications and services. In particular, the Application Plane incorporates the User Agent, a specialized application that interacts with the end user / operator and performs requests for Deterministic Networking services via an abstract Stream Management Entity, (SME) which may or may not be collocated with (one of) the end systems.

At the Application Plane, a management interface enables the negotiation of streams between end systems. An abstraction of the stream called a Traffic Specification (TSpec) provides the representation. This abstraction is used to place a reservation over the (Northbound) Service Interface and within the Application plane. It is associated with an abstraction of location, such as IP addresses and DNS names, to identify the end systems and eventually specify intermediate relay systems.

4.2. The Controller Plane

The Controller Plane corresponds to the aggregation of the Control and Management Planes in [RFC7426], though Common Control and Measurement Plane (CCAMP) [CCAMP] makes an additional distinction between management and measurement. When the logical separation of the Control, Measurement and other Management entities is not relevant, the term Controller Plane is used for simplicity to represent them all, and the term controller refers to any device operating in that plane, whether is it a Path Computation entity or a Network Management entity (NME). The Path Computation Element (PCE) [PCE] is a core element of a controller, in charge of computing Deterministic paths to be applied in the Network Plane.

A (Northbound) Service Interface enables applications in the Application Plane to communicate with the entities in the Controller Plane.

One or more PCE(s) collaborate to implement the requests from the SME as Per-Stream Per-Hop Behaviors installed in the relay systems for each individual streams. The PCEs place each stream along a deterministic sequence of relay systems so as to respect per-stream constraints such as security and latency, and optimize the overall result for metrics such as an abstract aggregated cost. The deterministic sequence can typically be more complex than a direct sequence and include redundancy path, with one or more packet replication and elimination points.

4.3. The Network Plane

The Network Plane represents the network devices and protocols as a whole, regardless of the Layer at which the network devices operate.

The network Plane comprises the Network Interface Cards (NIC) in the end systems, which are typically IP hosts, and relay systems, which are typically IP routers and switches. Network-to-Network Interfaces such as used for Traffic Engineering path reservation in [RFC3209], as well as User-to-Network Interfaces (UNI) such as provided by the Local Management Interface (LMI) between network and end systems, are all part of the Network Plane.

A Southbound (Network) Interface enables the entities in the Controller Plane to communicate with devices in the Network Plane. This interface leverages and extends TEAS to describe the physical topology and resources in the Network Plane.

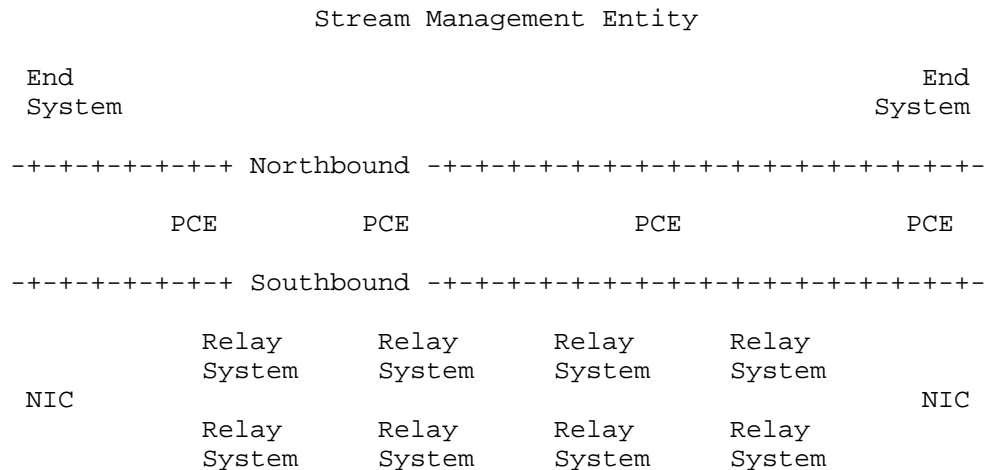


Figure 3

The relay systems (and eventually the end systems NIC) expose their capabilities and physical resources to the controller (the PCE), and update the PCE with their dynamic perception of the topology, across the Southbound Interface. In return, the PCE(s) set the per-stream paths up, providing a Stream Characterization that is more tightly coupled to the relay system Operation than a TSpec.

At the Network plane, relay systems exchange information regarding the state of the paths, between adjacent systems and eventually with the end systems, and forward packets within constraints associated to

each stream, or, when unable to do so, perform a last resort operation such as drop or declassify.

This specification focuses on the Southbound interface and the operation of the Network Plane.

4.4. Elements of DetNet Architecture

The DetNet architecture has a number of elements, discussed in the following sections:

- a. A model for the definition, identification, and operation of DetNet streams (Section 4.5), for use by relay systems to classify and process individual packets following per-stream rules.
- b. A model for the flow of data from an end system or through a relay system that can be used to predict the bounds for that system's impact on the QoS of a DetNet stream, without significantly constraining the method of implementing that system, for use by the Controllers to configure policing and shaping engines in Network Systems over the Southbound interface. The model includes:
 1. A model for queuing, transmission selection, shaping, preemption, and timing resources that can be used by an end system or relay system to control the selection of packets output on an interface. These models must have sufficiently well-defined characteristics, both individually and in the aggregate, to give predictable results for the QoS for DetNet packets (Section 4.7).
 2. A model for identifying misbehaving DetNet streams and mitigating their impact on properly functioning streams (Section 4.9).
- c. A model for the relay system to inform the controller(s) of the information it needs for adequate path computations including:
 1. Systems' individual capabilities (e.g. can do replication, can do precise time).
 2. Link capabilities and resources (e.g. bandwidth, 0 delays, hardware deterministic support to the physical layer, ...)
 3. Physical resources (total and available buffers, timers, queues, etc)

4. Network Adjacencies (neighbors)

- d. A model for the provision of a service, by end systems, or relay systems, to forward a DetNet stream over a simple or redundant path. The model includes:
 - 1. A model for an abstract relaying operation of either Routing or forwarding packets of a DetNet stream to a next-hop relay system, across Layer boundaries.
 - 2. A model of next-hop(s) information for replicating the packets of a DetNet stream, typically at or near the talker, merging and/or re-replicating those packets at other points in the network, and finally eliminating the duplicates, typically at or near the listener(s), in order to provide high availability (Section 3.3).
- e. The protocol stack model for an end system and/or a relay system should support the above elements in a manner that maximizes the applicability of existing standards and protocols to the DetNet problem, allows for the creation of new protocols where needed, thus making DetNet an add-on feature to existing networks, rather than a new way to do networking. In particular this protocol stack supports networks in which the path from talker to listener(s) includes bridges and/or routers in any order (Section 4.10).
- f. A variety of models for the provisioning of DetNet streams can be envisioned, including orchestration by a central controller or by a federation of controllers, provisioning by relay systems and end systems sharing peer-to-peer protocols, by off-line configuration, or by a combination of these methods. The provisioning models are similar to existing Layer-2 and Layer-3 models, in order to minimize the amount of innovation required in this area (Section 4.12).

4.5. DetNet streams

4.5.1. Talker guarantees

DetNet streams can be synchronous or asynchronous. The transmission of packets in synchronous DetNet streams uses time synchronization among the end and relay systems to control the flow of packets. Asynchronous DetNet streams are characterized by:

- o A maximum packet size;
- o An observation interval; and

- o A maximum number of transmissions during that observation interval.

These parameters, together with knowledge of the protocol stack used (and thus the size of the various headers added to a packet), limit the number of bit times per observation interval that the DetNet stream can occupy the physical medium.

The talker promises that these limits will not be exceeded. If the talker transmits less data than this limit allows, the unused resources such as link bandwidth can be made available by the system to non-DetNet packets. However, making those resources available to DetNet packets in other streams would serve no purpose. Those other streams have their own dedicated resources, on the assumption that all DetNet streams can use all of their resources over a long period of time.

Note that there is no provision in DetNet for throttling streams; the assumption is that a DetNet stream, to be useful, must be delivered in its entirety. That is, while any useful application is written to expect a certain number of lost packets, the real-time applications of interest to DetNet demand that the loss of data due to the network is extraordinarily infrequent.

Although DetNet strives to minimize the changes required of an application to allow it to shift from a special-purpose digital network to an Internet Protocol network, one fundamental shift in the behavior of network applications that is impossible to avoid--the reservation of resources before the application starts. In the first place, a network cannot deliver finite latency and practically zero packet loss to an arbitrarily high offered load. Secondly, achieving practically zero packet loss for unthrottled (though bandwidth limited) streams means that bridges and routers have to dedicate buffer resources to specific streams or to classes of streams. The requirements of each reservation have to be translated into the parameters that control each system's queuing, shaping, and scheduling functions and delivered to the hosts, bridges, and routers.

4.5.2. Incomplete Networks

The presence in the network of relay systems that are not fully capable of offering DetNet services complicates the ability of the relay systems and/or controller to allocate resources, as extra buffering, and thus extra latency, must be allocated at each point that is downstream from the non-DetNet relay system for some DetNet stream.

4.6. Data Flow Model through Systems

4.7. Queuing, Shaping, Scheduling, and Preemption

For this reason, the IEEE 802.1 Time-Sensitive Networking Task Group has defined a set of queuing, shaping, and scheduling algorithms that enable each bridge or router to compute the exact number of buffers to be allocated for each stream or class of streams.

4.8. Coexistence with normal traffic

A DetNet network supports the dedication of at least 75% of the network bandwidth to DetNet streams. But, no matter how much is dedicated for DetNet streams, It is a goal of DetNet to not interfere excessively with existing QoS schemes. It is also important that non-DetNet traffic not disrupt the DetNet stream, of course (see Section 4.9 and Section 6). For these reasons:

- o Bandwidth (transmission opportunities) not utilized by a DetNet stream are available to non-DetNet packets (though not to other DetNet streams).
- o DetNet streams can be shaped, in order to ensure that the highest-priority non-DetNet packet also is ensured a maximum latency.
- o When transmission opportunities for DetNet streams are scheduled in detail, then the algorithm constructing the schedule should leave sufficient opportunities for non-DetNet packets to satisfy the needs of the uses of the network.

Ideally, the net effect of the presence of DetNet streams in a network on the non-DetNet packets is primarily a reduction in the available bandwidth.

4.9. Fault Mitigation

One key to building robust real-time systems is to reduce the infinite variety of possible failures to a number that can be analyzed with reasonable confidence. DetNet aids in the process by providing filters and policers to detect DetNet packets received on the wrong interface, or at the wrong time, or in too great a volume, and to then take actions such as disabling the offending packet, shutting down the offending DetNet stream, or shutting down the offending interface.

It is also essential that filters and service remarking be employed to prevent non-DetNet packets from impinging on the resources allocated to DetNet packets.

There exist techniques, at present and/or in various stages of standardization, that can perform these fault mitigation tasks that deliver a high probability that misbehaving systemd will have zero impact on well-behaved DetNet streams, except of course, for the receiving interface(s) immediately downstream of the misbehaving device.

4.10. Protocol Stack Model

This section will be further developed. See [IEEE802.1CB], Annex C, for a description of the protocol stack. This is very much a work in progress, not a standard. See also [IEEE802.1Qcc].

4.11. Advertising resources, capabilities and adjacencies

4.12. Provisioning model

4.12.1. Centralized Path Computation and Installation

A centralized routing model, such as provided with a PCE (RFC 4655 [RFC4655]), enables global and per-stream optimizations. The model is attractive but a number of issues are left to be solved. In particular:

- o whether and how the path computation can be installed by 1) an end device or 2) a Network Management entity,
- o and how the path is set up, either by installing state at each hop with a direct interaction between the forwarding device and the PCE, or along a path by injecting a source-routed request at one end of the path.

4.12.2. Distributed Path Setup

Whether a distributed alternative without a PCE can be valuable should be studied as well. Such an alternative could for instance inherit from the Resource ReSerVation Protocol [RFC5127] (RSVP) flows.

In a Layer-2 only environment, or as part of a layered approach to a mixed environment, IEEE 802.1 also has work, either completed or in progress. [IEEE802.1Q-2011] Clause 35 describes SRP, a peer-to-peer protocol for Layer-2 roughly analogous to RSVP. Almost complete is [IEEE802.1Qca], which defines how ISIS can provide multiple disjoint paths or distribution trees. Also in progress is [IEEE802.1Qcc], which expands the capabilities of SRP.

5. Related IETF work

5.1. Deterministic PHB

[I-D.svshah-tsvwg-deterministic-forwarding] defines a Differentiated Services Per-Hop-Behavior (PHB) Group called Deterministic Forwarding (DF). The document describes the purpose and semantics of this PHB. It also describes creation and forwarding treatment of the service class. The document also describes how the code-point can be mapped into one of the aggregated Diffserv service classes [RFC5127].

5.2. 6TiSCH

Industrial process control already leverages deterministic wireless Low power and Lossy Networks (LLNs) to interconnect critical resource-constrained devices and form wireless mesh networks, with standards such as [ISA100.11a] and [WirelessHART].

These standards rely on variations of the [IEEE802154e] timeSlotted Channel Hopping (TSCH) [I-D.ietf-6tisch-tsch] Medium Access Control (MAC), and a form of centralized Path Computation Element (PCE), to deliver deterministic capabilities.

The TSCH MAC benefits include high reliability against interference, low power consumption on characterized streams, and Traffic Engineering capabilities. Typical applications are open and closed control loops, as well as supervisory control streams and management.

The 6TiSCH Working Group focuses only on the TSCH mode of the IEEE 802.15.4e standard. The WG currently defines a framework for managing the TSCH schedule. Future work will standardize deterministic operations over so-called tracks as described in [I-D.ietf-6tisch-architecture]. Tracks are an instance of a deterministic path, and the DetNet work is a prerequisite to specify track operations and serve process control applications.

[RFC5673] and [I-D.ietf-roll-rpl-industrial-applicability] section 2.1.3. and next discusses application-layer paradigms, such as Source-sink (SS) that is a Multipeer to Multipeer (MP2MP) model that is primarily used for alarms and alerts, Publish-subscribe (PS, or pub/sub) that is typically used for sensor data, as well as Peer-to-peer (P2P) and Peer-to-multipeer (P2MP) communications. Additional considerations on Duocast and its N-cast generalization are also provided for improved reliability.

6. Security Considerations

Security in the context of Deterministic Networking has an added dimension; the time of delivery of a packet can be just as important as the contents of the packet, itself. A man-in-the-middle attack, for example, can impose, and then systematically adjust, additional delays into a link, and thus disrupt or subvert a real-time application without having to crack any encryption methods employed. See [RFC7384] for an exploration of this issue in a related context.

Furthermore, in a control system where millions of dollars of equipment, or even human lives, can be lost if the DetNet QoS is not delivered, one must consider not only simple equipment failures, where the box or wire instantly becomes perfectly silent, but bizarre errors such as can be caused by software failures. Because there is essentially no limit to the kinds of failures that can occur, protecting against realistic equipment failures is indistinguishable, in most cases, from protecting against malicious behavior, whether accidental or intentional. See also Section 4.9.

Security must cover:

- o the protection of the signaling protocol
- o the authentication and authorization of the controlling systems
- o the identification and shaping of the streams

7. IANA Considerations

This document does not require an action from IANA.

8. Acknowledgements

The authors wish to thank Jouni Korhonen, Erik Nordmark, George Swallow, Rudy Klecka, Anca Zamfir, David Black, Thomas Watteyne, Shitanshu Shah, Craig Gunther, Rodney Cummings, Wilfried Steiner, Marcel Kiessling, Karl Weber, Ethan Grossman and Pat Thaler, for their various contribution with this work.

9. Access to IEEE 802.1 documents

To access password protected IEEE 802.1 drafts, see the IETF IEEE 802.1 information page at <https://www.ietf.org/proceedings/52/slides/bridge-0/tsld003.htm>.

10. Informative References

- [AVnu] <http://www.avnu.org/>, "The AVnu Alliance tests and certifies devices for interoperability, providing a simple and reliable networking solution for AV network implementation based on the Audio Video Bridging (AVB) standards.", .
- [CCAMP] IETF, "Common Control and Measurement Plane", <<https://datatracker.ietf.org/doc/charter-ietf-ccamp/>>.
- [HART] www.hartcomm.org, "Highway Addressable Remote Transducer, a group of specifications for industrial process and control devices administered by the HART Foundation", .
- [HSR-PRP] IEC, "High availability seamless redundancy (HSR) is a further development of the PRP approach, although HSR functions primarily as a protocol for creating media redundancy while PRP, as described in the previous section, creates network redundancy. PRP and HSR are both described in the IEC 62439 3 standard.", <<http://webstore.iec.ch/webstore/webstore.nsf/artnum/046615!opendocument>>.
- [I-D.finn-detnet-problem-statement]
Finn, N. and P. Thubert, "Deterministic Networking Problem Statement", draft-finn-detnet-problem-statement-01 (work in progress), October 2014.
- [I-D.ietf-6tisch-architecture]
Thubert, P., Watteyne, T., Struik, R., and M. Richardson, "An Architecture for IPv6 over the TSCH mode of IEEE 802.15.4e", draft-ietf-6tisch-architecture-06 (work in progress), March 2015.
- [I-D.ietf-6tisch-tsch]
Watteyne, T., Palattella, M., and L. Grieco, "Using IEEE802.15.4e TSCH in an IoT context: Overview, Problem Statement and Goals", draft-ietf-6tisch-tsch-05 (work in progress), January 2015.
- [I-D.ietf-roll-rpl-industrial-applicability]
Phinney, T., Thubert, P., and R. Assimiti, "RPL applicability in industrial networks", draft-ietf-roll-rpl-industrial-applicability-02 (work in progress), October 2013.

- [I-D.svshah-tsvwg-deterministic-forwarding]
Shah, S. and P. Thubert, "Deterministic Forwarding PHB",
draft-svshah-tsvwg-deterministic-forwarding-03 (work in
progress), March 2015.
- [IEEE802.1AS-2011]
IEEE, "Timing and Synchronizations (IEEE 802.1AS-2011)",
2011, <[http://standards.ieee.org/getIEEE802/
download/802.1AS-2011.pdf](http://standards.ieee.org/getIEEE802/download/802.1AS-2011.pdf)>.
- [IEEE802.1BA-2011]
IEEE, "AVB Systems (IEEE 802.1BA-2011)", 2011,
<[http://standards.ieee.org/getIEEE802/
download/802.1BA-2011.pdf](http://standards.ieee.org/getIEEE802/download/802.1BA-2011.pdf)>.
- [IEEE802.1CB]
IEEE, "Seamless Redundancy (IEEE Draft P802.1CB)", 2015,
<<http://www.ieee802.org/1/files/private/cb-drafts/>>.
- [IEEE802.1Q-2011]
IEEE, "MAC Bridges and VLANs (IEEE 802.1Q-2011)", 2011,
<[http://standards.ieee.org/getIEEE802/
download/802.1Q-2011.pdf](http://standards.ieee.org/getIEEE802/download/802.1Q-2011.pdf)>.
- [IEEE802.1Qat-2010]
IEEE, "Stream Reservation Protocol (IEEE 802.1Qat-2010)",
2010, <[http://standards.ieee.org/getIEEE802/
download/802.1Qat-2010.pdf](http://standards.ieee.org/getIEEE802/download/802.1Qat-2010.pdf)>.
- [IEEE802.1Qav]
IEEE, "Forwarding and Queuing (IEEE 802.1Qav-2009)", 2009,
<[http://standards.ieee.org/getIEEE802/
download/802.1Qav-2009.pdf](http://standards.ieee.org/getIEEE802/download/802.1Qav-2009.pdf)>.
- [IEEE802.1Qca]
IEEE, "Path Control and Reservation", 2015,
<<http://www.ieee802.org/1/files/private/ca-drafts/>>.
- [IEEE802.1Qcc]
IEEE, "Stream Reservation Protocol (SRP) Enhancements and
Performance Improvements", 2015,
<<http://www.ieee802.org/1/files/private/cc-drafts/>>.
- [IEEE802.1TSNTG]
IEEE Standards Association, "IEEE 802.1 Time-Sensitive
Networks Task Group", 2013,
<<http://www.IEEE802.org/1/pages/avbridges.html>>.

- [IEEE802154]
IEEE standard for Information Technology, "IEEE std. 802.15.4, Part. 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks", June 2011.
- [IEEE802154e]
IEEE standard for Information Technology, "IEEE std. 802.15.4e, Part. 15.4: Low-Rate Wireless Personal Area Networks (LR-WPANs) Amendment 1: MAC sublayer", April 2012.
- [ISA100.11a]
ISA/IEC, "ISA100.11a, Wireless Systems for Automation, also IEC 62734", 2011, < <http://www.isa100wci.org/en-US/Documents/PDF/3405-ISA100-WirelessSystems-Future-broch-WEB-ETSI.aspx>>.
- [ODVA]
<http://www.odva.org/>, "The organization that supports network technologies built on the Common Industrial Protocol (CIP) including EtherNet/IP.", .
- [PCE]
IETF, "Path Computation Element",
<<https://datatracker.ietf.org/doc/charter-ietf-pce/>>.
- [Profinet]
<http://us.profinet.com/technology/profinet/>, "PROFINET is a standard for industrial networking in automation.",
<<http://us.profinet.com/technology/profinet/>>.
- [RFC2205]
Braden, B., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, September 1997.
- [RFC3209]
Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, December 2001.
- [RFC4655]
Farrel, A., Vasseur, J., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, August 2006.
- [RFC5127]
Chan, K., Babiarz, J., and F. Baker, "Aggregation of Diffserv Service Classes", RFC 5127, February 2008.
- [RFC5673]
Pister, K., Thubert, P., Dwars, S., and T. Phinney, "Industrial Routing Requirements in Low-Power and Lossy Networks", RFC 5673, October 2009.

- [RFC7384] Mizrahi, T., "Security Requirements of Time Protocols in Packet Switched Networks", RFC 7384, October 2014.
- [RFC7426] Haleplidis, E., Pentikousis, K., Denazis, S., Hadi Salim, J., Meyer, D., and O. Koufopavlou, "Software-Defined Networking (SDN): Layers and Architecture Terminology", RFC 7426, January 2015.
- [TEAS] IETF, "Traffic Engineering Architecture and Signaling", <<https://datatracker.ietf.org/doc/charter-ietf-teas/>>.
- [WirelessHART] www.hartcomm.org, "Industrial Communication Networks - Wireless Communication Network and Communication Profiles - WirelessHART - IEC 62591", 2010.

Authors' Addresses

Norman Finn
Cisco Systems
170 W Tasman Dr.
San Jose, California 95134
USA

Phone: +1 408 526 4495
Email: nfinn@cisco.com

Pascal Thubert
Cisco Systems
Village d'Entreprises Green Side
400, Avenue de Roumanille
Batiment T3
Biot - Sophia Antipolis 06410
FRANCE

Phone: +33 4 97 23 26 34
Email: pthubert@cisco.com

Michael Johas Teener
Broadcom Corp.
3151 Zanker Rd.
San Jose, California 95134
USA

Phone: +1 831 824 4228
Email: MikeJT@broadcom.com