DetNet
Internet-Draft
Intended status: Informational
Expires: June 21, 2019

N. Finn
Huawei Technologies Co. Ltd
P. Thubert
Cisco

December 18, 2018

Deterministic Networking Problem Statement
draft-ietf-detnet-problem-statement-09

Abstract

This paper documents the needs in various industries to establish multi-hop paths for characterized flows with deterministic properties.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

The Deterministic Networking Use Cases [I-D.ietf-detnet-use-cases] document illustrates that beyond the classical case of industrial automation and control systems (IACS), there are in fact multiple industries with strong and yet relatively similar needs for deterministic network services with latency guarantees and ultra-low packet loss.

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) [IEEE802.1TSNTG] 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, IACS segregates the network along the broad lines of the Purdue Enterprise Reference Architecture (PERA) [ISA95], typically using deterministic local area networks for level 2 control systems, whereas public infrastructures such as Electricity Automation require deterministic properties over the Wide Area. The realization is now coming that the convergence of IT and Operational Technology (OT) networks requires Layer-3, as well as Layer-2, capabilities.

While the initial user base has focused almost entirely on Ethernet physical media and Ethernet-based bridging protocol from several Standards Development Organizations, the need for Layer-3 expressed above, must not be confined to Ethernet and Ethernet-like media. While such media must be encompassed by any useful Deterministic Networking (DetNet) Architecture, cooperation between IETF and other SDOs must not be limited to IEEE or IEEE 802. Furthermore, while the

work completed and ongoing in other SDOs, and in IEEE 802 in particular, provide an obvious starting point for a DetNet architecture, we must not assume that these other SDOs' work confines the space in which the DetNet architecture progresses.

The properties of deterministic networks will have specific requirements for the use of routed networks to support these applications and a new model must be proposed to integrate determinism in IT technology. The proposed model should enable a fully scheduled operation orchestrated by a central controller, and may support a more distributed operation with probably lesser capabilities. In any fashion, the model should not compromise the ability of a network to keep carrying the sorts of traffic that is already carried today in conjunction with new, more deterministic flows. Forward note: The DetNet Architecture [I-D.ietf-detnet-architecture] is the document produced by the DetNet WG to describe that model.

At the time of this writing, the expectation is that once the abstract model is agreed upon, the IETF will specify the signaling elements to be used to establish a path and the tagging elements to be used identify the flows that are to be forwarded along that path. The expectation is also that IETF will specify the necessary protocols, or protocol additions, based on relevant IETF technologies, to implement the selected model.

A desirable outcome of the work is the capability to establish a multi-hop path over the IP or MPLS network, for a particular flow with given timing and precise throughput requirements, and carry this particular flow along the multi-hop path with such characteristics as low latency and ultra-low jitter, reordering and/or replication and elimination of packets over non-congruent paths for a higher delivery ratio, and/or zero congestion loss, regardless of the amount of other flows in the network.

Depending on the network capabilities and on the current state, requests to establish a path by an end-node or a network management entity may be granted or rejected, an existing path may be moved or removed, and DetNet flows exceeding their contract may face packet declassification and drop.

2. On Deterministic Networking

The Internet is not the only digital network that has grown dramatically over the last 30-40 years. Video and audio entertainment, and control systems for machinery, manufacturing processes, and vehicles are also ubiquitous, and are now based almost entirely on digital technologies. Over the past 10 years, engineers

in these fields have come to realize that significant advantages in both cost and in the ability to accelerate growth can be obtained by basing all of these disparate digital technologies on packet networks.

The goals of Deterministic Networking are to enable the migration of applications with critical timing and reliability issues that currently use special-purpose fieldbus technologies (HDMI, CANbus, ProfiBus, etc... even RS-232!) to packet technologies in general, and the Internet Protocol in particular, and to support both these new applications, and existing packet network applications, over the same physical network. In other words, a Deterministic Network is backwards compatible with (capable of transporting) statistically multiplexed traffic while preserving the properties of the accepted deterministic flows.

The Deterministic Networking Use Cases [I-D.ietf-detnet-use-cases] document indicates that applications in multiple fields need some or all of a suite of features that includes:

1. Time synchronization of all host and network nodes (routers and/or bridges), accurate to something between 10 nanoseconds and 10 microseconds, depending on the application.

2. Support for Deterministic packet flows that:

    - Can be unicast or multicast;

    - Need absolute guarantees of minimum and maximum latency end-to-end across the network; sometimes a tight jitter is required as well;

    - Need a packet loss ratio beyond the classical range for a particular medium, in the range of 10^-9 to 10^-12, or better, on Ethernet, and in the order of 10^-5 in Wireless Sensor Mesh Networks;

    - Can, in total, absorb more than half of the network's available bandwidth (that is, massive over-provisioning is ruled out as a solution);

    - Cannot suffer throttling, congestion feedback, or any other network-imposed transmission delay, although the flows can be meaningfully characterized either by a fixed, repeating transmission schedule, or by a maximum bandwidth and packet size;

3. Multiple methods to schedule, shape, limit, and otherwise control the transmission of critical packets at each hop through the network data plane;

4. Robust defenses against misbehaving hosts, routers, or bridges, both in the data and control planes, with guarantees that a critical flow within its guaranteed resources cannot be affected by other flows whatever the pressures on the network - more on the specific threats against DetNet in the DetNet Security Considerations [I-D.ietf-detnet-security] document;

5. One or more methods to reserve resources in bridges and routers to carry these flows.

Time synchronization techniques need not be addressed by an IETF Working Group; there are a number of standards available for this purpose, including IEEE 1588, IEEE 802.1AS, and more.

The multicast, latency, loss ratio, and non-throttling needs are made necessary by the algorithms employed by the applications. They are not simply the transliteration of fieldbus needs to a packet-based fieldbus simulation, but reflect fundamental mathematics of the control of a physical system.

With classical forwarding latency- and loss-sensitive packets across a network, interactions among different critical flows introduce fundamental uncertainties in delivery schedules. The details of the queuing, shaping, and scheduling algorithms employed by each bridge or router to control the output sequence on a given port affect the detailed makeup of the output stream, e.g. how finely a given flow's packets are mixed among those of other flows.

This, in turn, has a strong effect on the buffer requirements, and hence the latency guarantees deliverable, by the next bridge or router along the path. For this reason, the IEEE 802.1 Time-Sensitive Networking Task Group has defined a new 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 flow or class of flows.

Robustness is a common need for networking protocols, but plays a more important part in real-time control networks, where expensive equipment, and even lives, can be lost due to misbehaving equipment.

Reserving resources before packet transmission is the one fundamental shift in the behavior of network applications that is impossible to avoid. 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 un-throttled (though bandwidth limited) flows means that bridges and routers have to dedicate buffer resources to specific flows or to classes of flows. The requirements of each reservation have to be translated into the parameters that control each host's, bridge's, and router's queuing, shaping, and scheduling functions and delivered to the hosts, bridges, and routers.

3. Problem Statement

3.1. Supported topologies

In some use cases, the end point which run the application is involved in the deterministic networking operation, for instance by controlling certain aspects of its throughput such as rate or precise time of emission. In that case, the deterministic path is end-to-end from application host to application host.

On the other end, the deterministic portion of a path may be a tunnel between an ingress and an egress router. In any case, routers and switches in between should not need to be aware whether the path is end-to-end or a tunnel.

While it is clear that DetNet does not aim at setting up deterministic paths over the global Internet, there is still a lack of clarity on the limits of a domain where a deterministic path can be set up. These limits may depend in the technology that is used to set the path up, whether it is centralized or distributed.

3.2. Flow Characterization

Deterministic forwarding can only apply on flows with well-defined characteristics such as periodicity and burstiness. Before a path can be established to serve them, the expression of those characteristics, and how the network can serve them, for instance in shaping and forwarding operations, must be specified.

3.3. Centralized Path Computation and Installation

A centralized routing model, such as provided with a Path Computation Element (PCE) (see [RFC4655]), enables global and per-flow 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,

- 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 following classical Traffic Engineering (TE) models.

To enable a centralized model, DetNet should produce a description of the high level interaction and data models to:

- report the topology and device capabilities to the central controller;

- establish a direct interface between the centralized PCE to each device under its control in order to enable a vertical signaling

- request a path setup for a new flow with particular characteristics over the service interface and control it through its life cycle;

- support for life cycle management for a path (instantiate/modify/update/delete)

- support for adaptability to cope with various events such as loss of a link, etc...

- expose the status of the path to the end devices (UNI interface)

- provide additional reliability through redundancy, in particular with packet Packet Replication, Elimination and Ordering Functions (PREOF) where the former may generate an out-of-order delivery that may need to be corrected corrected by the latter;

- indicate the flows and packet sequences in-band with the flows, this is needed for flows that require PREOF in order to isolate duplicates and reorder in the end;

3.4. Distributed Path Setup

Whether a distributed alternative without a PCE can be valuable could be studied as well. Such an alternative could for instance inherit from the Resource ReSerVation Protocol [RFC3209] (RSVP-TE) flows. But the focus of the work should be to deliver the centralized approach first.

To enable a RSVP-TE like functionality, the following steps would take place:

1. Neighbors and their capabilities are discovered and exposed to compute a path that fits the DetNet constraints, typically of latency, time precision and resource availability.

2. A constrained path is calculated with an improved version of Constrained Shortest Path First (CSPF) that is aware of DetNet.

3. The path may be installed using a control protocol such as RSVP-TE, associated with flow identification, per-hop behavior such as Packet Replication and Elimination, and blocked resources. In that case, traffic flows can be transported through an MPLS-TE tunnel, using the reserved resources for this flow at each hop.

3.5. Duplicated data format

In some cases the duplication and elimination of packets over non-congruent paths is required to achieve a sufficiently high delivery ratio to meet application needs. In these cases, a small number of packet formats and supporting protocols are required (preferably, just one) to serialize the packets of a DetNet stream at one point in the network, replicate them at one or more points in the network, and discard duplicates at one or more other points in the network, including perhaps the destination host. Using an existing solution would be preferable to inventing a new one.

4. 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.

Typical control networks today rely on complete physical isolation to prevent rogue access to network resources. DetNet enables the virtualization of those networks over a converged IT/OT infrastructure. Doing so, DetNet introduces an additional risk that flows interact and interfere with one another as they share physical resources such as Ethernet trunks and radio spectrum. The requirement is that there is no possible data leak from and into a deterministic flow, and in a more general fashion there is no possible influence whatsoever from the outside on a deterministic flow. The expectation is that physical resources are effectively associated with a given flow at a given point of time. In that model, Time Sharing of physical resources becomes transparent to the

individual flows which have no clue whether the resources are used by other flows at other times.

The overall security of a deterministic system must cover:

- the protection of the signaling protocol
- the authentication and authorization of the controlling nodes including plug-and-play participating end systems.
- the identification and shaping of the flows
- the isolation of flows from leakage and other influences from any activity sharing physical resources.

The specific threats against DetNet are further discussed in the DetNet Security Considerations [I-D.ietf-detnet-security] document.

5. IANA Considerations

This document does not require an action from IANA.

6. Acknowledgments

The authors wish to thank Lou Berger, Pat Thaler, Jouni Korhonen, Janos Farkas, Stewart Bryant, Andrew Malis, Ethan Grossman, Patrick Wetterwald, Subha Dhesikan, Matthew Miller, Erik Nordmark, George Swallow, Rodney Cummings, Ines Robles, Shwetha Bhandari, Rudy Klecka, Anca Zamfir, David Black, Thomas Watteyne, Shitanshu Shah, Kiran Makhijani, Craig Gunther, Warren Kumari, Wilfried Steiner, Marcel Kiessling, Karl Weber, Alissa Cooper, and Benjamin Kaduk for their various contributions to this work.

7. Informative References

[I-D.ietf-detnet-architecture]
Finn, N., Thubert, P., Varga, B., and J. Farkas, "Deterministic Networking Architecture", draft-ietf-detnet-architecture-09 (work in progress), October 2018.

[I-D.ietf-detnet-security]
Mizrahi, T., Grossman, E., Hacker, A., Das, S., Dowdell, J., Austad, H., Stanton, K., and N. Finn, "Deterministic Networking (DetNet) Security Considerations", draft-ietf-detnet-security-03 (work in progress), October 2018.

[I-D.ietf-detnet-use-cases]
Grossman, E., "Deterministic Networking Use Cases", draft-ietf-detnet-use-cases-19 (work in progress), October 2018.

[IEEE802.1TSNTG]
IEEE Standards Association, "IEEE 802.1 Time-Sensitive Networks Task Group", 2013, <http://www.ieee802.org/1/pages/avbridges.html>.

[ISA95] ANSI/ISA, "Enterprise-Control System Integration Part 1: Models and Terminology", 2000, <https://www.isa.org/isa95/>.

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

[RFC4655] Farrel, A., Vasseur, J., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, DOI 10.17487/RFC4655, August 2006, <https://www.rfc-editor.org/info/rfc4655>.

[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>.

Authors' Addresses

Norman Finn
Huawei Technologies Co. Ltd
3755 Avocado Blvd.
PMB 436
La Mesa, California 91941
US

Phone: +1 925 980 6430
Email: norman.finn@mail01.huawei.com

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

Phone: +33 497 232 634
Email: pthubert@cisco.com